MILLEE: Mobile and Immersive Learning for Literacy in Emerging Economies

Matthew Boon Tian Kam

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2008-155
http://www.eecs.berkeley.edu/Pubs/TechRpts/2008/EECS-2008-155.html

December 13, 2008

Acknowledgement

(See actual dissertation.)

MILLEE: Mobile and Immersive Learning for Literacy in Emerging Economies

by

Matthew Boon Tian Kam

B.A. (University of California, Berkeley) 2001
B.S. (University of California, Berkeley) 2001

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Computer Science

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:
Professor John Canny, Chair
Professor Eric Brewer
Professor Glynda Hull

Fall 2008

The dissertation of Matthew Boon Tian Kam is approved:

Chair Date

Date

Date

University of California, Berkeley

Fall 2008

MILLEE: Mobile and Immersive Learning for Literacy in Emerging Economies

by

Matthew Boon Tian Kam

Abstract

MILLEE: Mobile and Immersive Learning for Literacy in Emerging Economies

by

Matthew Boon Tian Kam

Doctor of Philosophy in Computer Science

University of California, Berkeley

Professor John Canny, Chair

Literacy levels in most developing countries remain shockingly low and formal education is making little progress. MILLEE improves literacy through language learning games on cellphones – the "Personal Computers of the developing world" – which are a perfect vehicle for new kinds of out-of-school language learning. Games bring children into rich, immersive environments where they can acquire and use language naturally, while encouraging them to transfer their language skills outside the game.

The MILLEE research project focuses on developing scalable, localizable design principles and tools for language learning. The challenges are (i) to integrate sound learning principles, (ii) to provide concrete design patterns that integrate entertainment and learning, and (iii) to account for cultural and learning differences in children in developing regions. In this thesis, we will describe a framework called PACE that addresses these challenges and eight rounds of fieldwork that had contributed to its development. We will also describe a tool to expedite audio-only learning (Pimsleur Generator), a very important niche for developing regions. We discuss our most recent work which patterns learning games after local children's traditional village games and the benefits this approach offers. Finally, we describe the complex adoption ecology in

developing regions, and how MILLEE preserves learning principles while supporting rich localization and customization at multiple stages in the adoption hierarchy.

____

Professor John Canny
Dissertation Committee Chair

"My son will read and open the books, and my son will write and know writing. And my son will make numbers, and these things will make us free because he will know – he will know and through him we will know." – the non-literate Kino, in John Steinbeck's *The Pearl*

Table of Contents

Acknowledgments

First and foremost, I wish to acknowledge a debt of gratitude to my grandparents and parents. During the summer after I had completed my basic degrees at the University of California, Berkeley, I was working in a challenging position in my home country that drew my attention to the inadequacies of my formal training. I understood that one way to realize my full potential is to complete my schooling via a PhD degree, so that I would be better equipped to succeed in future and more demanding responsibilities. My parents did not have the opportunity to attend college – much less understand what a PhD program is about. But they nevertheless pledged their utmost support when I indicated my interest in pursuing a PhD. My father recalled that two of his classmates who finished their PhDs had gone on to become "movers and shakers" in my country, and perhaps have similar hopes for me. But my country's government opposed my intentions to enroll for a PhD and wanted me to return permanently as soon as possible. My appeals to the highest level of government were not successful. That experience provided me with invaluable insights into the workings of bureaucrats and politicians. It was also a most stressful ordeal, and I am grateful to my parents and grandparents for standing by me throughout this period.

I would like to thank my thesis committee in Computer Science (Eric Brewer and John Canny) and Education (Glynda Hull) for a supportive environment that enabled me to pursue doctoral research across traditional disciplinary boundaries. As a driving force behind the "Technology and Infrastructure for Emerging Regions" research group, Eric evangelized our mission – among the many other things that he did – and helped us gain greater acceptance in mainstream CS. Eric truly provided us with the "air cover" that we doctoral students needed while we were conducting our research on the ground. Growing up as a teenager, I was inspired by the economists who contributed to my country's economic miracle and dreamt of becoming an outstanding development economist after their footsteps. My encounter with John changed my life's direction. It was John who introduced me to "human-centered computing" and showed how CS can integrate the social and behavioral sciences to target relevant human needs, when I did undergraduate research with him. The experience made me turn down other graduate programs to stay on at Berkeley, and to contemplate CS as a serious career – and not only a hobby. As my thesis advisor afterward, John was initially worried about the difficulty of field research in a developing country setting, but gave his support anyway. I have always admired his courage to set the trend for others to follow. I hope that our results have shown that it is indeed possible to make progress in "technology for developing regions" research even when ground conditions are difficult, and offer some lessons to future graduate students who wish to embark on similar work.

This thesis would not have happened without Glynda. It was her introducing me to Urvashi Sahni in April 2004 that led to the genesis of the MILLEE project, as narrated in Chapter 2. Glynda has continually encouraged me to acquire a background in language and literacy studies, and I eventually took enough classes in Education for a minor in my doctoral program. This background was instrumental in enabling me to frame this thesis as CS in the service of humanity, as opposed to technology for technology's sake. I thank my classmates and instructors (Anne Cunningham, Sarah Freedman, Claire Kramsch and Laura Sterponi, in addition to Glynda) at the Graduate School of Education at Berkeley who welcomed me into their midst and provided a supportive environment for an outsider to learn their craft. Our class discussions helped me to grasp the material in an entirely new field (my basic degrees were in Economics and CS), and my class term papers were subsequently revised to become Chapters 4 and 5.

An international project would not be possible without our collaborators in India. In the last 4½ years, we have cooperated with the Mysore Literacy Trust (Babu Mathew, Soundar Rajan and M.L. Ramanarasimha), Sesame Workshop India (Sashwati Banerjee), and Suraksha (Rahul Chatterjee, Pratim Kumar, Shalini Mathur and Urvashi Sahni). The MLT and Suraksha hosted the field studies described in Chapters 2, 3, 7, 8, 9, 10 and 11. I thank them for their support and helping us to establish trust with the local communities

whom they seek to empower. It has been enriching to interact with and observe how the children in rural and urban slums communities made use of our prototypes. I learned as much from watching them as I did from my graduate coursework and a part-time diploma course on Teaching English as a Foreign Language. Most of all, my interactions with the persons on the street during the course of my fieldwork in developing regions gave me ample opportunities to further my education in Economics in the real world.

I could not have pulled off the MILLEE project without my team members. I have worked with 30+ undergraduate research assistants at both UC Berkeley and in India, and some of them were a tremendous help. Ruth Alexander, Asya Grigorieva, Maksim Lirov, Priyanka Reddy and Monish Subherwal were part of the group who experimented with the pedagogical design patterns described in Chapter 5, and endured my "thinking aloud" as I tried to articulate the emerging PACE framework from Chapter 6 to them.

Aishvarya Agarwal, Varun Devanathan, Anjali Koppal, Anuj Kumar, Siddhartha Lal, Akhil Mathur, Anand Raghavan, Vijay Rudraraju and Anuj Tewari accompanied me on my fieldwork, and provided the manpower to help with data collection when I could not do everything by myself within the short time on the ground. In addition, Aishvarya and Varun digitized the images and voice recordings for our prototypes, and in doing so, helped to establish the structure for our subsequent content development process. Anuj Kumar, Siddhartha, Akhil, Vijay and Anuj Tewari programmed the games that we field-tested and described in Chapters 7, 8, 9 and 10. Some of the games were based on older code by Dimas Guardado, Christopher Hom, Aaron McKee and Anand Raghavan. Anjali also lent her linguistic background in Hindi, Kannada and English to help me understand the Pimsleur audio units, which allowed me to design and program Pimsleur Generator. It has been a pleasure to have worked with the above undergraduate research assistants, and they have in turn helped me to stretch and grow as a mentor.

Our local personnel who were based in India are: Mehnaaz Abidi, Aman Anand, Siddharth Bhagwani, Jatin Chaudhary, Shirley Jain, Neelima Purwar, Gautam Singh and

Kavish Sinha. Shirley was a teacher of English as a Second Language in India. It was Shirley who developed the curriculum for our games. Her practical experience as an English teacher dovetailed with my theoretical knowledge of second language acquisition and reading science, and facilitated a productive cross-disciplinary collaboration. Aman served as our content developer and digitized the curriculum that Shirley developed. The remaining local personnel listed above were responsible for the operational management of our pilot deployments and implementing the data collection procedures that I had in mind.

Graduate students Jane Chiu, David Nguyen and Divya Ramachandran went with me on some of those field studies described in Chapters 2, 3 and 8. I have always been amazed by their ability to interact effectively with children, and their presence helped to facilitate the smooth running of our user studies. My lab-mate and "technical whiz" (even among CS PhDs!) Jingtao Wang helped us troubleshoot the esoteric technical challenges that we encountered.

A special round of thanks goes to Lauren Bailey, who has been the indefatigable energy behind the administrative processes at UC Berkeley for the MILLEE project. If not for her persistence, we would not have overcome the bureaucratic hurdles for signing a Memorandum of Understanding between UC Berkeley and Suraksha, which is required for the pilot deployment and summative evaluation described in Chapter 10.

Sangeeta Anand, Susan Dray, Jerome Feldman, Alastair Iles, Thomas Kalil and Annie Yeh have my gratitude for acting as our advocates. In particular, Alastair, my friend and mentor, contributed his time and wisdom unstintingly with invaluable advice throughout the project. Tom Kalil has been generous in introducing us to his network of contacts. For example, it was Tom who introduced me to Linda Roberts, a board member of Sesame Workshop whom he got to know during his days as a Clinton advisor, and her enthusiasm for the MILLEE project, which led to our collaboration and the inclusion of some popular "Sesame Street" characters in the MILLEE games. Tom also connected us with Verizon, who eventually funded us to draw on our 4½ years of lessons in India and explore similar applications with Spanish-speaking cellphone subscribers in the US.

The MILLEE project has received major funding from the MacArthur Foundation (Digital Media and Learning award), Microsoft (Digital Inclusion award), US National Science Foundation (Grant No. 0326582), Qualcomm (Wireless Reach award) as well as Verizon. Other sources of financial support are the "Big Ideas @ Berkeley" competition on serious games, Edioma and Intel undergraduate research program. I thank Seeqpod for travel sponsorship and Sony Creative Software for their sponsorship-in-kind. There were several moments when we needed funding to expand the MILLEE project, and many of the above funding and solicitations came when we needed them most. On a related note, I thank Leora Eisen, Eleisha McNeill and the filming crew for their enthusiasm in filming MILLEE on-site at an Indian village for a television documentary on novel cellphone applications in the developing world, which was aired by the Canadian Broadcasting Corporation in Canada on April 3 and June 5, 2008. I hope that the documentary will raise the MILLEE project's visibility and help us secure resources for future expansions, especially in this difficult financial downturn.

Finally, I thank the parents in the urban slums and villages of India who gave their consent for their children to participate in our numerous field studies, in the hopes that their children – together with millions of other poor children throughout the world – can realize their fullest potential.

1 Introduction

Low literacy is one of the grand challenges in the developing world. Despite huge improvements in recent decades, literacy in many poor regions remains shockingly low. But even more challenging is the tension between regional languages and global "power languages," such that economic opportunities are often closed to those literate only in a regional language. For instance, India is a country with 22 regional and 2 national languages, that is, Hindi and English (Figure 1.1). But English, together with computer skills, are the two most sought-after skills in surveys of poor parents (Pal, Lakshmanan and Toyama 2007). The English language is widely perceived to be a socioeconomic enabler:

> [Indians] secretly believe, if not openly say, that competence in English makes a considerable difference in their career prospects…. Politicians and bureaucrats denounce the elitism of [English-medium] schools but surreptitiously send their children to them. (Gupta 1995: 76)

English is not only the language of all the professions and higher education, but is also crucial for middle-level service jobs such as retail, clerical, teaching and law enforcement that are the most common rungs along the social ladder above menial labor. The value of English is widely recognized by most ordinary Indians (Shukla 1996), and it is in fact the poorest citizens who are lobbying most strongly to expand English teaching. As a result, English is taught in almost every school in India: as a second language in public schools, and as a first language and the medium of instruction in most private schools.

Figure 1.1: A map of India that shows how states are divided along linguistic lines, such that English plays a crucial role as the official language of business and government.

English is thus the language of power in India, such that mastery of the language can almost be associated with membership in the middle and upper classes in India (Faust and Nagar 2001; Kishwar 2005):

> English has the constitutional status of an 'associate official language' in a highly multilingual national context and is the dominant medium of higher-level administration, higher education, the learned professions, large-scale industry and commerce, and a considerable part of literary and artistic activity. Indians who use English are estimated to constitute only about 5 per cent of the nation's population, but this group forms a very large proportion of those who are in leadership roles and are concentrated in the largest cities in the country, where English functions as a lingua franca (Prabhu 1987: 5).

The economists Munshi and Rosenzweig observe that in the Indian city of Mumbai, "schooling in [the regional language] Marathi channels the child into working class jobs, while more expensive English education significantly increases the likelihood of obtaining a coveted white-collar job" (2003: 2). They estimate that English speakers in Mumbai experience returns on investment in schooling that are between 24% and 27%. On the other hand, non-English speakers with similar characteristics experience returns that are about 10%.

For marginalized communities in other developing regions elsewhere, the power language could be another language such as Spanish, Mandarin or French that is spoken by the dominant social classes in the population. This power language is not likely to be spoken by the underserved communities as their native languages. As an example, French is spoken as the language of business in postcolonial Morocco where Arabic and other native languages are used. The economists Angrist and Lavy (1997) found that switching the medium of instruction from French to Arabic in public schools appeared to be associated with a decline in French writing skills, and hence a reduction of almost 50% in the returns to schooling. The focus of the MILLEE (Mobile and Immersive Learning for Literacy in Emerging Economies) project is therefore on second language acquisition, which entails an overlapping focus on developing early literacy in the target language. But we believe many of our lessons will apply to other languages and developing regions, in spite of the project being presently centered around English as a Second Language (ESL) in India.

Unfortunately, the public school systems in developing regions such as India have poor outcomes. In India, the Azim Premji Foundation and Pratham are two of the non-governmental organizations (NGOs) at the forefront of spearheading initiatives to make educational opportunities more accessible to underprivileged children living in rural areas and urban slums. Two dismal statistics from the "Annual Status of Education Report 2007" national survey administered by Pratham (2007) stand out: 1) more than 26% of rural children could not read the English alphabet after one year of instruction, and 2) teacher absenteeism rate exceeded 25%. These statistics are highly consistent with our experience on the ground. As we will describe in Chapter 3, at one of the rural schools where we conducted early fieldwork, the principal had specially selected his best students for us to interact with, so as to impress us, who were his foreign visitors. However, we soon realized that many of them struggled to read the letters in the English alphabet despite having taken three years of classes on English as a subject. At another rural government school we visited, there were three teachers responsible for teaching the students there. We soon observed that on any given day, one of the teachers would be taking his turn to go on “unofficial leave.”

In our field studies in India, we were consistently unable to converse in English with the teachers who were responsible for teaching English in poor schools without the help of our interpreters. More troubling, it appears that it is difficult to make progress on the problem purely through the formal schooling system. According to a literature review commissioned by the Azim Premji Foundation (2004b), public schooling is out of reach for more than 43% of school-going age children in rural areas who cannot attend school regularly due to their need to work for the family in the agricultural fields or households. As such, India has a two-tier education system that reproduces existing social structures, such that very good private schools serve the upper and middle classes and teach entirely in English, while low-quality government schools are where the overwhelming majority of children struggle to learn:

> English has come to play the role of a divider rather than a unifier in the larger Indian society…. It has had a stratifying effect because those who are not sufficiently literate in English are essentially unable to participate in a number of domains such as higher education or global business. English proficiency and access to English may also be a strong factor in the polarization between urban and rural communities. (Sheorey 2006: 18)

On the other hand, we hypothesize that e-learning applications on cellphones can create a more engaging learning experience when they take the form of e-learning games. Furthermore, we hypothesize these applications can make language and literacy learning more effective when their instructional designs draw on the latest research by education researchers in second language acquisition and the science of reading. Most importantly, cellphones are increasingly adopted in developing regions to the extent that the cellphone has become the fastest growing technology platform in the developing world (Vodafone 2005). A growing fraction of these phones feature programmable support and multimedia capabilities for gaming, which make these devices a promising vehicle for out-of-school learning that can complement the formal schooling system. We therefore hypothesize that we can dramatically expand the reach of language and literacy learning in the developing world by using portable mobile devices such as the cellphone as the target platform, so as to potentially enable children with work commitment to access educational resources anytime and anywhere, at places and times that are more convenient than school alone.

Using e-learning videogames to target the language and literacy learning needs of underserved communities is not far-fetched. Gee (2004) argues that games can make the learning process more engaging while incorporating good educational principles at the same time. Furthermore, the Azim Premji Foundation and Pratham have implemented e-learning games in their educational initiatives for children in rural areas and urban slums in India. Most importantly, a large-scale longitudinal randomized experiment with more than 10,000 urban slums children in India over more than two years, carried out between Pratham and economists from the MIT Poverty Action Lab (Banerjee et al. 2005), has shown significant learning benefits from games that target mathematics when children played these games in twice-weekly sessions. We believe that similar learning outcomes can be replicated in the case of e-learning games that target language and literacy.

1.1 Project Summary (2004 to 2009)

At the time of writing, MILLEE is an ongoing project that is in the middle of its fifth year. However, this thesis will only cover the first four years (i.e. mid-2004 to mid-2008) and report the steps currently taken in Year 5 under future work (Chapter 13). To date, we have made eight trips to India, with the time spent on the ground totaling 7½ months. Figure 1.2 provides a graphical depiction of this timeline in terms of the field studies that represent significant milestones in the project, and expected expansion path for late 2008 and beyond. We should also be clear that even though we carried out needs assessment and exploratory studies in the beginning of the project lifecycle, we continued to acquire a more nuanced understanding of the local context after every trip to India. In other words, the overall process should not be misinterpreted as entirely linear. Similarly, although the graphical timeline gives the impression that the overall process was smooth, the reality is that every step in the project is more accurately characterized as "two steps forward, one step back."

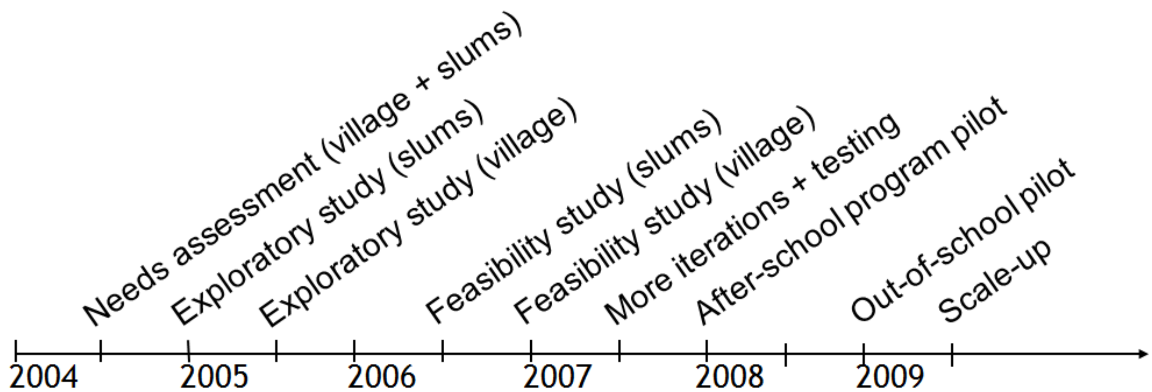

Figure 1.2: The timeline of the MILLEE project and the field studies in India from 2004 to 2008 that constitute significant milestones. The activities indicated for late 2008 and after are tentative, and represent our projected expansion path.

The reader will also observe that we addressed the challenges of conducting user studies in this difficult cross-cultural design environment by running every field study with students in the slums before we attempt a similar study at a more comprehensive level with rural children. This systematic and cautious approach enabled us to iron out initial kinks in our research designs and prototypes with children in the urban slums, in preparation for subsequent studies with rural children, who were considerably harder to conduct user studies with given their lower levels of familiarity with technology and the logistical obstacles in traveling to the villages.

We summarize our human-centered design process and the field studies that we have performed so far:

July 2004 (2 weeks in Lucknow, Uttar Pradesh) – qualitative interviews at five rural schools near Lucknow and an afternoon school program for girls living in the slums in Lucknow. The goal of conducting these visits was to gain a first-hand understanding of ground level conditions in poor schools in India, instead of starting the MILLEE project with preconceived ideas about what we would do. We also did qualitative interviews with teachers, low-income students and their parents to understand their attitudes toward computer-aided learning, based on an initiative that a local non-profit partner started in 2000, and their everyday obstacles in obtaining access to high quality education. Respondents consistently cited English as one of the three hardest subjects for teachers in poor schools to teach, and this difficulty is compounded by the limited ability of rural teachers to make themselves understood in English when interacting with us. Most important, we gained a deep appreciation of how infrastructural constraints in terms of limited building space and irregular electricity make it challenging for computer-aided learning on desktop computers to scale.

November 2004 (2 weeks in Lucknow) – exploratory studies on multimedia storytelling at the above afternoon program for urban slums girls. Our goal in carrying out exploratory studies was to examine the feasibility of some ideas for computer-aided learning that arose from our needs assessment, in order to determine if it is indeed realistic to proceed further with these ideas. We made use of off-the-shelf software that embodied these ideas, as opposed to implementing our own prototypes for the exploratory studies. In this way, we reduced our emotional attachment to our ideas and made it easier to discard non-practical ideas. One of the most salient lessons was that it was very difficult for participants to find time to attend our sessions. We did not appreciate it completely at that time, but these incidents underscored the potential for mobile learning.

August 2005 (3 weeks in Lucknow) – like the exploratory studies in November 2004, we carried out similar exploratory studies, but this time, at a government rural school. We experimented with digital cameras, authoring applications and e-learning games that targeted English as a Second Language, all on the Tablet PC platform. We found that a game which targeted vocabulary demonstrated positive outcomes in terms of short-term retention. We also learned that the participants in our study, who were the brightest students handpicked by their principal to impress us (who were his visitors from abroad), struggled with the English alphabet despite having already taken three years of ESL classes. This experience made us realize that it is essential to return to the basics in English teaching, and that a highly structured syllabus – as opposed to a more open-ended constructionist approach – is more feasible, at least in the short term.

August 2006 (2½ weeks in Lucknow) – building on the above early success with a vocabulary building game for ESL, we conducted our first feasibility study of ESL learning games that we developed for the cellphone platform. This study was carried out at the above afternoon school for urban slums students. We designed these games based on the best practices used in successful language learning applications on the commercial market, and saw significant learning gains on short-term vocabulary retention. More important, we learned that the ESL baseline for this group differed dramatically from the rural students we worked with in August 2005. This would be our first – but not last – time that we observe highly uneven ESL baselines across low-income learners in India, such that there is little of the lock-step correlation between age and grade level that is the norm in the industrialized world. This experience led us to propose the PACE framework (Chapter 6) for structuring ESL learning applications into separate modules for curriculum and learning activities, so that existing content can be inexpensively adapted for local conditions, and hence scale for a different community by reusing as many existing modules as possible.

January 2007 (2 weeks in Mysore, Karnataka) – a second feasibility study with cellphone games. Unlike the feasibility study in August 2006, this study took place at a village school in South India – which enabled us to observe learning dynamics in a different state in India – and focused on gameplay enjoyment as opposed to learning. We had designed these games based on common patterns found in successful Western digital games, and observed that these features do not usually match the understanding and expectations that rural children have about games.

May 2007 to August 2007 (8 weeks in Mysore, followed by 5 weeks in Lucknow) – building on previous fieldwork that have studied ESL learning (August 2006) and gameplay enjoyment (January 2007) separately, the games piloted in this extended fieldwork integrated earlier lessons on ESL learning and gameplay enjoyment into the same designs. We field tested the new designs with three communities of rural children, i.e. one in Mysore and two in Lucknow. This study was substantially longer than earlier studies since we were performing frequent iterative prototyping in the field in response to field observations. We also studied the traditional village games that rural children play, in order to understand how their games differ from Western videogames.

December 2007 (4 weeks in Lucknow) – drawing on the formative assessments that we performed in the summer of 2007, we proceeded with a summative evaluation throughout the spring of 2008 that spanned 38 sessions over 5 months. This assessment took the form of a pilot deployment in an after-school program at a village school which is an hour's drive from Lucknow. It aimed to examine the learning impacts of cellphone-based ESL learning on a more extensive scale compared to previous summative studies. We spent our time in Lucknow during December 2007 working with the local community and pilot personnel to begin the pilot.

June 2008 (2 weeks in Lucknow) – this trip to India was meant to wind down the pilot that spanned the spring of 2008. Specifically, we worked with pilot personnel to resolve ambiguities in the collected data and to assemble the data into a form that is amenable for analysis. We also made use of this period to advance the MILLEE project into the post-doctoral thesis stage, by conducting initial ethnographic studies that aimed to inform how we could expand our work into out-of-school settings other than an after-school program.

1.2 Structure of Thesis

The rest of this thesis is structured in the following way: Chapter 2 narrates the genesis of the MILLEE project. It provides an account of the needs assessment that we conducted with teachers, parents and their children in villages and urban slums in India, the ground conditions that make desktop computing less than conducive, and ideas for computer-aided learning based on our ground observations. Chapter 3 describes the key results from our exploratory studies with children from the urban slums and rural areas in India. In these studies, we experimented with various ideas for computer-aided learning which seemed plausible from our needs assessment. These ideas included constructionist learning by having children author digital stories as producers of digital learning content, as well as children playing e-learning games that target English as a Second Language as consumers of digital content. In the process, however, our direct observations of the children's dismal English baseline made us realize that we needed to defer lofty goals such as constructionism, and that our immediate priority is to return to the basics – such as alphabet and rudimentary vocabulary. We also present initial results that suggest that e-learning games for ESL learning is a promising direction towards which to steer the project – and the rest of this thesis.

Chapter 4 elaborates on our theoretical framework, namely, task-based language teaching, for second language instruction. We describe how the task provides a structure as an object of design, and how tasks can be composed with other tasks and activities to form larger instructional sequences. Following that, we review the empirical evidence for task-based language teaching, its successes with children in India so far, and some of its criticisms. We propose that the criticisms can be reconciled with the successes observed by reframing task-based language teaching within a sociocultural perspective of language development. In Chapter 5, we show how we made use of task-based language teaching as our analytical lens for reviewing more than 35 state-of-the-art commercial applications for language learning. We summarize over 50 design patterns that we distilled from this review for representing the best practices employed in the commercial products.

The challenge, however, is that "one-size-fits-all" games cannot accommodate the complexity of the adoption ecology and the requirement for local adaptation. In Chapter 6, we make the case for micro-localization and describe the PACE framework, such that educational applications structured according to this framework can have their learning activities and curriculum reused as much as possible in the interests of scalability. Next, although we describe the design patterns for language pedagogy in Chapter 5, gameplay design is another important design dimension in the MILLEE project. We describe our experiment with game design patterns in Chapter 7, where we found that such patterns can both help and hinder game design. It seems that patterns require adequate knowledge of the cultural context to be employed effectively. However, it remains difficult to articulate what these cultural factors are in a principled manner.

In Chapter 8, we embark on the difficult task of integrating the dual concerns of instructional and videogame design. We show how the pedagogical design patterns from Chapter 5 were applied within the broader context of the PACE framework to the design of ESL learning games. In addition to drawing on principles from language acquisition, these designs also incorporated heuristics for enjoyable gameplay. We describe our field trials with children from both the urban slums and rural areas. In particular, whereas the initial designs were fairly successful with urban slums children, the same designs needed numerous iterations before we ironed out usability issues with rural children. It appears that rural children may face more usability problems than their urban slum counterparts because the former had lesser exposure to high technology and videogames. Although we were not able to identify how to use game design patterns effectively, some of the later game designs were nevertheless based on videogames that were found to be popular with rural children.

However, some of the game mechanics in these popular digital games continue to elude rural children, whose understanding and expectations about games differ from our Western notion of games. In Chapter 9, we take a step back and study the characteristics of 28 traditional village games that children play in rural regions. We analyze how their traditional village games vary from contemporary Western videogames, and show that designing digital games that account for these distinctions appear to yield promising results. Since language learning is a long-term process, however, the formative studies in Chapter 8 do not adequately investigate the learning possibilities that cellphone-based ESL learning aims to address. Chapter 10 is our first step in this direction. It describes a summative evaluation in which we piloted games on cellphones that target ESL learning with 27 children in an after-school program – which we ran over the course of an entire semester at a village school in North India. We found that in an underdeveloped region where rural children do not have alternative access to quality ESL instruction, as a group, participants exhibited significant post-test gains. But the gains were unevenly distributed, with the academically advanced students accruing more of the benefits. Future research will have to examine how e-learning games in this underserved context can be designed to better scaffold the less advanced learners.

Audio-only language learning holds the promise that poor children in developing regions who spend many hours doing repetitive manual work may be able to learn from audio material that is played in the background. Chapter 11 describes an application for audio-only learning called Pimsleur Generator. Preliminary field trials show that it was less popular with rural children compared to e-learning games, which suggest that more research into audio-only learning is required. Chapter 12 situates the MILLEE project in the related literature. Chapter 13 presents our conclusions and directions for future research. Our bibliography is given in Chapter 14.

2 Needs Assessment

We uphold a human-centered view of computing, in which we believe that it is critical to understand the users and their needs before we begin to design and implement technology solutions for and with them. The idea of performing a needs assessment, as opposed to imposing a technology solution from outside without sensitivity to local needs and conditions, has been accepted within the international development community as a step that is crucial to the success and adoption of an initiative. On the other hand, the predominant approach in the computer science community is a *technology push* mindset, in which technology solutions are devised and evaluated based on technical novelty. This approach is not without its benefits, but is less likely to succeed with interactive systems that call for a high degree of human-computer interaction (HCI).

It is the HCI sub-community in computer science that is advocating a *social pull* approach to computer systems design, in which social considerations about the context of use influence design decisions. In other words, technology designers ought to investigate and understand users before designing computing systems that target their needs. A needs assessment becomes even more crucial in the MILLEE (Mobile and Immersive Learning for Literacy in Emerging Economies) project, since we aim to empower users from a very different cultural background with significantly lower levels of literacy and exposure to technology compared to users in industrialized societies. It is this focus on international development that makes MILLEE radically different from most previous work in HCI, which have historically focused on computing applications in the developed world.

This Chapter describes the genesis of the MILLEE project. It originated from the author's first visit to India, which constituted his first exposure to the conditions in rural schools and the urban slums. The rest of this Chapter summarizes the key considerations for technology design based on observations at the ground level, and ideas for computer-aided learning that result from these observations. These ideas may have the potential to address relevant local learning needs in directions yet to be adequately explored, and we follow up on some of them in exploratory studies which we describe in the next Chapter.

2.1 Background

Among the states of India, Uttar Pradesh (UP) has the highest number of residents below the poverty line in terms of income (Srivastava 2002). Dr. Urvashi Sahni, who is the founder of one of our non-government organization partners, and an alumnus of the University of California, Berkeley (PhD in Education, 1994), has worked on educational issues in rural UP for about two decades. She is widely recognized for her efforts to reform education in India and improve educational opportunities for girls. Her prior work includes initiating and managing a school reform project involving 62 schools, 16,000 students and 258 teachers in rural UP, as well as launching an in-service program for the United Nations Children's Fund (UNICEF) with 30,000 kindergarten and first-grade teachers in 28 UP districts. Her other accomplishments include: initiating and managing a school reform project involving 62 schools, 16,000 students, and 258 teaching staff in rural areas of UP, launching an in-service program for United Nations Children's Fund (UNICEF) with 30,000 kindergarten and first-grade teachers in 28 districts in UP, working as the director of an action research project to bolster education for girls and serving on the UP state government's Girl Child Mission to promote girls' education, founding and running a highly regarded school in Lucknow (UP) and its affiliated afternoon school program for girls from the neighboring slums.

In 2001, Dr. Urvashi Sahni pioneered the use of computers in rural schools in UP. In this pilot, which was called Digdarshan, a computer was introduced into a rural school, after which custom courseware was developed for this computer and distributed via CD-ROM. The pilot was gradually extended to six other rural schools – one private and six government-run – as well as the above urban afternoon school program for girls in the urban slums. This model has since been replicated by the state government in 1,400 rural schools. The experiences from the Digdarshan pilot provided us with initial lessons on computer-aided learning with low-income children in India, and served as a starting point for informing the MILLEE project when we began to collaborate with this NGO in 2004. In particular, we are especially keen to explore innovative designs of systems and their applications that can potentially revolutionize educational delivery in the interest of poor children around the world.

Specifically, the Digdarshan pilot commenced on June 30, 2000, when a Pentium III desktop computer with 128 MB RAM and a 20 GB hard-drive was placed in one of the above rural schools and set up with its screen facing outdoors, according to “The Hole in the Wall” configuration (Mitra et al. 2003). This computer came with the Windows 98 operating system and a 52X CD-ROM drive installed. It was preloaded with a variety of common English and Hindi applications. But being situated in a rural region, it lacked Internet connectivity. With the aid of three solar panels, the computer could be used for up to five hours per day. Students were given three days of training, with each session lasting up to 30 minutes. The training covered basic computer concepts and terminology, basic applications such as Microsoft Paint, and how to install software applications from CD-ROMs. Computer literacy was introduced into the school curriculum as a required course, such that some class time was set aside for students to use the computer. These classes were placed on the same level of importance as other required subjects. Students worked in groups to learn how to use the computer. Some “computer geniuses” emerged and taught other students to use the machine. This outcome was consistent with outcomes in the original “The Hole in the Wall” experiment.

Within five weeks into the Digdarshan pilot, between 25 and 30 children were able to use applications such as Microsoft Paint. By April 2002, about 239 students were using the computer at various pilot schools, with most of these users belonging to grades 5 to 8. All of them were able to turn it on and off, and use the mouse. 125 students knew how to use Microsoft Paint and Word, and navigate file directories. But only 10 students could use PowerPoint. These statistics were consistent with what we learned when we visited the same rural schools. We were told that about 90% of the students can use the computer, especially Microsoft Paint and Word. PowerPoint appeared to be significantly harder to master, in that only a few students knew how to use it, and by the time they learned to use it, it was almost time for them to graduate.

Teachers told us that they were not the primary users of the computers. But they supervised their students' use of the computers, in order to ensure that they do not get electrocuted or hit the wrong button. Teachers accepted the computer in the classroom because they believed that it benefited their students, and were confident that the machine did not undermine their authority in the classroom. Both teachers and students came to see themselves as owners of the technology. Respondents told us that it is a source of pride because no other schools in the vicinity have a computer.

But teachers and students began to realize the limits of the standard applications that came with the Windows operating system. Both the teachers and students asked for educational software in Hindi for a variety of subjects. In response, Dr. Urvashi Sahni put together a team of two government school primary teachers, two recent graduates from the teacher training institute, one 9th grade student from the same school where the pilot commenced, and two Flash developers. This team made science lesson units in Hindi for grades 5 to 8 based on the official curriculum. The decision to focus on science was due to the shortage of teachers in village schools who were originally trained in science. Through the courseware, the computer became integrated into the curriculum, such that teachers used the courseware as a lecturing tool to present course material to the entire class, who would sit facing the computer monitor.

According to one student, the courseware has additional content that was not found in existing textbooks, and she claimed to have learned it. For the teachers, the courseware helped to fill in the gaps in their subject matter knowledge. This was because it was not always possible for teachers in village schools to be fully qualified to teach every subject, especially when there could be as little as two teachers in a school of about 250 students and teachers were responsible for teaching every subject. There were some limitations with using the computer in this manner to teach an entire class, however. The same student told us that she had difficulty seeing the screen from the back of the class. She preferred her textbooks to the screen since textbooks are individually accessible.

On top of being used as a lecturing tool, the computer was also used to facilitate small-team learning. Under a teacher's supervision, students in small groups helped one another to understand the material in the courseware. Furthermore, older students who were accustomed to teaching their juniors (due to the shortage of teachers) found the courseware to be a useful teaching aid. The courseware included an assessment module, and students enjoyed taking the electronic quizzes to the extent that they kept re-taking them until they obtained a perfect score. In the process, students helped one another to understand why their initial answers were incorrect.

Among the few students who attained significant computer literacy – to the level of Microsoft PowerPoint – tghree of them authored their own rudimentary courseware in the form of multimedia stories. Such stories can be used to communicate one's life story, but can also focus on academic content, such that digital stories are a narrative description of an academic topic using a combination of still images, text and voice recordings. We were told that at least one student disliked studying, but was so motivated to create a digital story using PowerPoint to explain soil erosion that he persisted at the task for three days to complete it. But digital stories seemed to be difficult to create for reasons that we will describe below. Despite the difficulty of composing digital content, however, such creations appeared to be enthusiastically received by both content producers and content consumers. One student told us that she felt a sense of pride in being able to show off what she created using Microsoft Paint to her teachers and peers. Similarly, digital stories were very well received by students because they were excited to see that their classmates were capable of authoring these digital artifacts.

Methods

Lucknow is the state capital of Uttar Pradesh. When we visited Lucknow for the first time in the summer of 2004, one of the NGO's staff members who was instrumental to the Digdarshan pilot acted as our local cultural guide and interpreter. We conducted semi-structured interviews with 8 teachers and 6 students. Our informants belonged to the NGO's afternoon school for urban slums girls and five of the above village schools near Lucknow that participated in the Digdarshan pilot. In addition, we interviewed 2 families from the slums whose daughters were in the afternoon school program. All interviews were conducted in Hindi with our local guide serving as the interpreter. If he was not available, a teacher in the afternoon school, who was effectively bilingual in Hindi and English, acted as our interpreter. (Teachers in the afternoon program were significantly better equipped to discharge their teaching duties than the regular teacher in a rural public school, due to the NGO's commitment to delivering high quality education.)

On top of the interviews, we videotaped a session in which the shared computer was used in a village school, and a session in which the shared computer was used in the after-school program. The video data was subsequently transcribed and then translated from Hindi to English so that it could be analyzed.

The interviews focused on informant attitudes towards education and technology, including computer-aided learning and how they perceived the shared computer in the above pilot program. In the case of students, we also asked about the challenges that they face in improving their educational prospects, both with and without using the computer. In our interviews and observations of ground conditions, we sought to understand both the everyday constraints and enabling conditions that defined the "space" of possibilities for computer-aided learning in a resource-poor context.

2.2 Ground Conditions in Village Schools

We were told that in villages, there may be only one rural public school within 15 kilometers. Whereas children in remote rural areas have to walk to school, we learned that students in less rural (i.e. peri-urban) areas have family members who take them to school on bicycle and other means of transport. The schools that we visited have as many as 250 students and as little as 2 teachers. Low teacher-student ratio was therefore a perennial problem that applied to several schools. More troubling, teachers were not necessarily equipped to teach the subjects that they had to teach. For instance, although teachers in rural schools were responsible for teaching English, we were only able to communicate with them with the help of the interpreter. The language that is spoken predominantly in Uttar Pradesh is Hindi. We were informed that Hindi is spoken by 40% of the country, whereas English is spoken by only 15% of the population. The teachers whom we interviewed consistently cited English, mathematics and science as the three most difficult subjects to teach.

2.2.1 School Life

The teachers whom we interviewed estimated the monthly household income of a rural family to be approximately 3,000 Indian rupees (approximately US$65). We also learned that rural families usually have as many as seven members in the household, to the extent that each member lives on less than $2 per day.

One of the rural schools that we visited offered classes from grades 1 to 8, while the other four schools offered classes from grades 1 to 5 only. Classes run on weekdays and Saturdays. School starts at 7am with prayers, and classes run from 8am to 12 noon. Each class session lasts 30 minutes. There is a mid-day break from 10 to 10:30am, during which students play. Public schools in Uttar Pradesh did not provide free mid-day meals, unlike some other states in India that have implemented a similar program and witnessed increased attendance.

We observed that each student carried a bag containing items such as stationery, textbooks and notebooks. The availability of these stationery and notebooks suggest that students may be more receptive to a pen-based tablet device that allows the user to write using a stylus, as opposed to a computer that requires the keyboard for user input. As we will highlight later below, we observed students struggling with usability problems on the keyboard. This observation suggests that computer-aided learning may be more feasible if the keyboard is not indispensable for user input. Textbooks were provided free by the state government, who was aware that poor students cannot afford to pay for them. A textbook cost about 16 rupees, which was the price of newspapers for about a week. But most respondents appeared to share the consensus that textbooks are of poor educational quality.

There is an annual syllabus mandated by the state's education board. Classes in village schools are taught using Hindi as the medium of instruction and subjects include: English, Hindi, Sanskrit, science (physics, chemistry, biology), mathematics (arithmetic, algebra, geometry), history, social studies, home science, agriculture, general knowledge, arts and craft, bookcraft, computer literacy, and games (both indoor and outdoor) such as khoko, kabaddi, and cricket. Some of the subjects are optional. Annual exams take place in April and results are released in May. In addition to annual exams, there are monthly and half-yearly exams, both of which count towards the final grade.

2.2.2 Low School Attendance

Two of the schools we visited had an attendance chart written on a blackboard (Figure 2.1). The attendance rate at the schools was about 50% during the summer when we visited. A teacher at one of the schools told us that attendance was low during the summer because it was the mango-picking season, and students had to help their parents with the harvesting.

Figure 2.1: This photograph shows a blackboard (left) with the school's attendance chart written on it, and a cabinet (right) where the shared computer in the pilot program could be kept safely under lock and key. The computer is also connected to an Uninterrupted Power Supply system seen at the bottom of the cabinet.

When asked to name the most difficult challenge that they face in their work, teachers from at least two schools cited the lack of parental support for their children's schooling. According to teacher respondents, most students were eager to attend classes, but their parents were not always willing to send them to school. This was because the parents themselves were usually illiterate and felt that their children could be earning money. One teacher estimated that 60% of all parents held this attitude. Teachers had to engage with these parents and educate them on the importance of formal schooling. Interestingly, after the shared computer was introduced in the Digdarshan pilot, parents became more willing to allow their children to attend classes at the pilot schools more regularly because they believed it was important for their children to be computer literate. This outcome is supported by surveys of rural parents reported in Pal, Lakshmanan and Toyama (2007), which found that rural parents associated computer-aided learning with social advancement.

At the policy level, the government of Uttar Pradesh rewarded each child who had an 80% school attendance rate with rice. As the last resort, teachers made use of this incentive to persuade parents to send their children to school. We were told that the rice given was comparable to what a child would earn in a month. In fact, a teacher at a private rural school told us that a few students had dropped out in order to transfer to a government school, since only government schools offer the rice incentive.

There was also a gender dimension to school attendance. Some parents prevented their daughters from attending school for fear of sexual assaults, since there was little social and legal recourse for victims of sexual crimes in Indian society. More important, some teachers told us that parents were reluctant to send their daughters to a school that did not have a female teacher. Fortunately, the mindset regarding education for girls has been changing. Some parents in the rural areas are expressing an interest in having their daughters study so that they can achieve something in their lives. A possible explanation offered by respondents was the recent "Education for All" (i.e. *Sarva Shiksha Abhiyan*) government program, which required teachers to go to every home and get every child under age 14 registered for school. In addition to this policy, there was another scheme in which a girl who passed grade 12 will receive 20,000 rupees. This sum was considerable and made up slightly more than a rural family's half-yearly income.

2.2.3 Limited Building Space

Building space for classrooms in rural schools is limited (Figure 2.2), to the extent that students from multiple grades are often squeezed into the same room and lessons may even take place in the open area outside the classrooms. Moreover, in every school that we visited, the computer is perceived as a precious resource that must be kept locked in a wooden cabinet (Figure 2.1). In some schools, the computer is kept in a room dedicated to the computer. We were told that the room is too small to be used as a regular classroom (for reasons unknown to us), such that it would have otherwise gone disused if it is not used for computer sessions. In other village schools where there are no “unused” classrooms, the computer is kept in a regular classroom. It appears that this constraint presents an inconvenience in that the computer, being a scarce resource, cannot be used by other students when a class is taking place in the room where the computer is kept.

Figure 2.2: At one of the rural schools where we performed fieldwork, the cluster of students in the upper-left, upper-right and lower-left hand corner of the photograph correspond to students in grades 3, 4 and 5 respectively. Students belonging to all three grades are squeezed into the same classroom for their school lessons.

The more crucial point is that even if there were financial means for schools to obtain more desktop computers, including recycled computers, the shortage of built space in schools – in the form of classrooms that can be securely padlocked – usually prevents them from housing more than one computer. Security is a pressing concern, such that the three solar panels that were used to power the computer at one of the above pilot schools were stolen after they had been in place for three years. The implication is that even though existing initiatives by several NGOs favor the desktop computer as the platform for educational delivery, infrastructural constraints such as limited secure building space (and electricity, which we discuss in the next subsection) render the desktop infeasible in most village schools unless additional classrooms can be built.

On the other hand, there is plenty of open space outside the classroom. As such, if there is no particular reason to stick with desktop computers, it is worth considering mobile devices such as cellphones, low-cost laptops and Tablet PCs. Students can, for example, share these devices in small groups as they sit and work on the grass patch next to the school. There are of course possible caveats. For instance, mobile computers need strong contrast displays to be viewable outdoors under the sun's glare. It is also important to investigate the extent to which rural children can use mobile technology effectively or take responsible care of the mobile devices – which are not necessarily a familiar part of their everyday lives – especially during inclement weather. We address the latter point in the next chapter.

2.2.4 Lack of Regular Electricity

Electricity is the most serious impediment to deploying information technologies in village schools (Figures 2.3 to 2.5). 4 out of the 5 schools that we visited suffered from intermittent power supply, and on average, received a few hours of electricity each day – sometimes only at night. In fact, one of the schools has not had electricity for five days. On the other hand, we note that another school has fairly reliable access to electricity.

Figure 2.3: The school shown in this photograph is not connected to the electricity grid. Whenever the students need to use the computer, they connect the school to a neighbor's home, by taking a long coil of wire and connecting one end to the school. The students in this photo are using masking tape to join the coil of wire to the white wires, which are a permanent fixture of the school.

The temperature in Uttar Pradesh can be as high as 118 degrees Fahrenheit (48 degrees Centigrade). We learned from the students that they woke up as early as 4am to study before the day became hot. As such, one way to tell if a village school has reliable electricity supply is to look upwards and see if there are ceiling fans installed. As a few teachers among our respondents have told us, if they had more regular electricity, they would install ceiling fans to help their students cope with the heat.

Figure 2.4: There was considerable division of labor in wiring up the school in Figure 2.3 for electricity. Even as the students in Figure 2.3 were busy taping the wires together, other students were laying out the cable all the way to the neighbor's home, and placing bricks on the cable at regular intervals to keep it in place.

We learned from students in the pilot schools that on average, each of them got to use the computer in a group for 10 to 15 minutes each day, and only when electricity was available. Students did not have more time to use the computer since it has to be shared during school hours. Uninterrupted Power Supply (UPS) was the most common solution for intermittent power supply, but was inadequate because it only provided 20 minutes of power in the event of a power outage. This duration was enough only for students to save their work and shut down the computer properly.

Figure 2.5: This photograph shows what happened at the other end of the cable in Figure 2.4. Some students had succeeded in climbing over a wall – which is much taller than them – into their neighbor's home to tap on the latter's electricity supply. One respondent told us that the students stole the neighbor's electricity, while another person told us that the students had asked for permission. In any case, the entire process of connecting the school to the electricity grid took 20 minutes.

2.2.5 Usability Challenges

One of the responsibilities of our local guide was to visit the Digdarshan pilot schools periodically and check that the computers were working. Moreover, if there were electricity, he would also train students to use the computer. Finally, he was responsible for showing teachers how to use the science courseware. The latter sessions typically lasted between a day and a week.

Our guide told us that he introduced the students to computers by teaching them to use Microsoft Paint, after which they progress to Microsoft Word, and then Microsoft PowerPoint. He observed that students were more excited to use PowerPoint than Word because it is graphical, even though most students did not eventually succeed in learning PowerPoint due to its difficulty. Specifically, despite the motivation of children in the Digdarshan pilot program to create multimedia stories, only 3 out of 239 students succeeded. These students needed intensive guidance – both technical and subject matter – from our guide throughout the authoring process that spanned three days. It appeared that PowerPoint was too difficult for the average child to learn and use because he or she lacked access to the computer to practice using it. For example, we found out that 125 out of the 239 children had learned to use Word to some extent after the pilot had been in operation for a year, whereas there were only 10 out of 239 students who learned to use PowerPoint. Moreover, we understand that it was almost time for the same students to graduate from the school when they finally learned to use PowerPoint.

In addition to usability challenges with software, there were also usability issues associated with computing hardware. In particular, we observed that students were not proficient at typing, which was only to be expected given the limited time they have with the computer and the duration required to acquire typing skills. In fact, our guide believed that it was more urgent to show them how to use the space, backspace and Enter keys to move the cursor, as opposed to teaching them to use the cursor keys in the beginning. Based on his experience, with up to 15 minutes of computer usage per day, he estimated that an average student would take 6 months to become proficient with basic keyboard controls.

The computers in the pilot program were not used entirely for academic purposes, however. For instance, Microsoft Paint and computer games were encouraged as a means of fostering children's interest in computers at an early age. But it remained important to take the perceived distinction between entertainment and education that parents have into consideration, because this will have a major implication on the adoption of educational technology. One teacher narrated an incident when she screened a video-clip in class on the computer and never wanted to do it again. That was because her students told their parents that they had watched a movie in school, and parents subsequently complained to teachers that their children should not be attending school to watch movies. This incident underscored the importance of helping parents appreciate the educational value in novel (from a rural perspective) computing applications.

2.3 Ground Conditions in Urban Slums School

In addition to examining the ground conditions at five schools in the villages, we also visited an afternoon school that targets girls living in low-income households in the urban slums. This afternoon program takes place in the premises of a three-story, urban private school in Lucknow (Figure 2.6). This private school caters to students from urban middle-income families and their classes are held in the morning. Classes range from pre-nursery through pre-college. The facilities on the premises include a swimming pool, a netball court, a basketball court, science laboratories, a music room and two computer laboratories. As such, regular classes run in the morning for middle-class children (using English as the medium of instruction), and the premises are available in the afternoon to conduct lessons for girls from the slums who would otherwise not have a chance to attain formal schooling (using Hindi as the medium of instruction). The afternoon school is cross-subsidized by the revenue from the regular morning classes.

Figure 2.6: Classes in the afternoon school program that target girls from the nearby urban slums.

2.3.1 Life in School

The afternoon school program offers lessons up to grade 8 and classes take place on weekday afternoons from 1pm to 4:30pm. This schedule accommodates the students, who have to perform housework in the morning and evenings. A teacher in this program felt that the students in this program lagged in their studies by a year compared to their middle-class counterparts. Students live in the nearby slums, which is a 10-minute walk away.

One student told us that the program offers a better environment in which to study. Prior to enrolling there, she was in a public school where she neither received a basic education (such as algebra) nor was she taught practical skills. A teacher in the afternoon program attributed the problem to teachers in the government school, who she believed were either ill-equipped or not motivated to teach. She narrated anecdotes about government school teachers who chat or quarrel all day among themselves.

Another student recounted how she has improved in arithmetic after participating in the program, to the extent that whenever she has to go to a provision store to purchase goods, she is capable of calculating the change correctly. Most mathematical operations that she has to perform involve additions and working with estimates. Even though the numbers that she has to work with in everyday transactions are up to four digits, she is confident of working with up numbers that have up to five digits. She attributed her improvement to how mathematics was taught in the program using the blackboard, in comparison to her prior school, where explanations were communicated verbally only.

The students have acquired some mastery of English, although we were told that their knowledge leaned towards written English as opposed to spoken English. Students could speak basic phrases such as “Welcome,” “Can I go to the bathroom?” and “Sorry,” and our NGO partner who administers this program is encouraging the teachers to focus more on conversational English.

Learning also took place in non-academic terms. One student indicated that she learned to be better disciplined and to maintain a higher level of hygiene. Another student said that education dispels superstitions, in that her family no longer perceives diseases as a curse from God. From now onwards, whenever they fall seriously ill, they will head for treatment at a clinic. Most importantly, the afternoon school appears to have achieved some attitudinal changes. One mother observed her daughter has become more confident, to the extent that she believes herself to be capable of accomplishing anything that she sets her mind on.

The shared computer was used in the same way in the afternoon school program and the rural public schools. In particular, teachers appreciated that the courseware saved them from having to spend time writing on the blackboard. In the computer session that we observed with students from grades 7 and 8 (Figure 2.7), the teacher in this photo was using the courseware to present and explain concepts in microbiology (e.g. dehydration) with the aid of a monitor placed in front of the class. The electronic content included still images and video-clips. The teacher paused the courseware occasionally to elaborate on some concepts, while gesturing at the screen. In the process, she and Dr. Sahni (who was also present during the session we observed) directed questions at students and engaged in discussions with them. We saw that this interactive form of teaching engaged the students' rapt attention. Students were so eager to answer questions that they continually interrupted one another.

After the teacher had covered a module in the courseware, every student needed to take turn to stand in front of the class and paraphrase the same material to classmates. The session concluded with students coming forward, either individually or in pairs, to attempt the matching puzzle included in the courseware. We were told that there were puzzles other than matching games, such as crossword puzzles, fill-in-the-blanks quizzes and multiple-choice questions. The teacher observed that the urban slums students found the puzzles to be fairly intuitive and learned the rules after one round.

Figure 2.7: A class session in the urban slums program where a teacher (left) uses the shared computer to present science concepts to students. In this photograph, the teacher has finished explaining the concepts for the day, and it is every student's turn to come up to the front of the class and explain the same concepts to her peers.

After covering the module for the above session, the teacher tested the students on their retention with a set of multiple-choice questions. She summarized her impressions of teaching using the shared computer:

> Generally, I repeat the same topic two or three times before the class understands and remembers it. However, using the computer, within just one class, everyone seems to have grasped the topic very well indeed…. Computers make it interesting and exciting for the teacher to teach their students the required lessons. Normally, they are least interested in what is going on or being taught in class. Today was a big change.

One student added that she experienced better retention of the material when concepts were presented using computer graphics and animation, as compared to the textbook. In fact, when she was interviewed a week after the above session, she could recall that she was taught five types of micro-organisms (i.e. fungi, algae, protozoa, bacteria and virus) during the class. Before that, she was unaware that small organisms cause diseases.

After the above session, students sat in pairs at each computer. The goal of this post-session was to provide students with follow-up opportunities to engage further with the same material covered earlier. In this particular case, we observed the teacher asking her students to Google for images of bacteria, so that seeing the visual forms would help them better relate to the material. (Whereas there was no Internet connectivity in the above villages where the computers were piloted, the premises where the afternoon school program was held had a 128 kbps shared connection.) We observed usability problems with Google, however, and two staff members were at hand to render assistance to the students. Moreover, some students needed the teacher's help to spell the search term "bacteria" in English since search terms had to be given in English. The teacher eventually devised a solution by providing students with a list of English words related to "bacteria." Another frustration was the lack of educational content in the local language (Hindi) on the Internet. For example, we saw our NGO partner's futile attempts to search online for video-clips related to bacteria.

Despite these obstacles, one student indicated that she enjoyed using the computer because of the Internet, especially Google and email. In addition, we were informed that students have learned to use Microsoft Paint, Word and PowerPoint to some degree, and that they will be learning to make digital stories soon.

2.3.2 Life at Home

On one evening after classes at the afternoon school program were over, we accompanied two students back to their homes in the urban slums to better understand their family conditions. Both homes were infested with flies, although one was much less infested. Our hosts in both families kept waving a large fan to keep the flies away from us. The typical dwelling space for an entire household is a room which is the size of 3 to 4 meters by 6 to 7 meters (Figure 2.8). The average household comprised the parents and 4 to 5 children. This space was used for multiple purposes: living, cooking and eating.

(However, one of the two families spent the night at a clinic, where the father performed shift duty as a paramedic. But the space that they shared at the clinic with the families of other staff members was even smaller than their slum quarters.) At least one of the two families occupied a dwelling space with a bathroom, but it lacked running water and they had to obtain water from the tap outside the home. There were no electric bulbs installed because of the lack of electricity. But one of the homes owned a battery-operated black-and-white television set and used a kerosene lamp at night for lighting.

Figure 2.8: There is very little furniture in a typical slums dwelling except those for cooking and dining, such as cookstoves, crockery and a dining table with a makeshift tablecloth. There are some pots and pans next to the cookstoves. Indoor pollution from cooking is a problem. We could hardly breathe when we were there. In one small corner is a worship corner with a statue. The wall is plastered with newspaper and magazine clippings that form a wallpaper.

We learned that the students are required to perform chores in the morning and evening. For instance, in one family, the daughters were required to cook dinner because their mother had to work as a foodstuff vendor in the evening. Other household chores included washing the dishes, laundry and house cleaning. In another family, a daughter was required to help her parents sell vegetables in the afternoon, until a teacher from the afternoon school program intervened and insisted that she be given a chance to attend formal lessons. Due to their housework duties, the students are free from 1 to 4:30pm in the afternoons, which is why the afternoon program takes place during this time. For the same reason, students do not necessarily have adequate time at home for homework.

To provide a deeper glimpse into the gender bias that permeates Indian society, we draw attention to the observation that while the daughters in one of the two families attended the afternoon school program, the same family had mustered the resources to enroll their only son in a private school. The mother told us that this school charged 170 rupees per month for school fees and 1,000 rupees per month for stationery. As a basis for comparison, the latter was equivalent to one-third of the family's income. In addition, the parents bought many computer books for their son even though he had no access to computer training classes and had highly limited access to the two personal computers in his school.

Both families lived on daily wages. The average monthly wage for a household was about 3,500 rupees, and monthly expenditure on food was at least 1,500 rupees. In the second family, the father worked as a paramedic while the mother was a cleaner. She expressed that she has no reluctance in sending her children to school, and since she is illiterate (but her husband had studied up to grade 10), she wanted them to receive an education so that they could live a better life. She told us that she and her husband had in fact uprooted themselves from their village so as to provide their children with better educational opportunities in the city. They had also transferred their son out of a school that they deemed unwholesome for children. Furthermore, there was a time when her husband stopped talking to one of his daughters as he had used to, after she disappointed him by performing poorly in grade 7 in her previous school. That was a privately-run school in which teachers used to beat the students. But he was very happy to learn that his daughter was enjoying her classes at the afternoon school program. Unfortunately, due to the parents' work schedule, they were mostly not at home to help their children with their schoolwork.

2.4 Ideas for Computer-Aided Learning

During our interviews with informants at both the village schools and urban slums school, we tried to probe the propositional and procedural knowledge that teachers think students should acquire. It seemed practical skills were highly valued, such that teachers perceived education to yield benefits that include the ability to make calculations needed for everyday activities (e.g. shopping), literacy tasks in everyday activities (e.g. reading signs and directions when traveling), and help illiterate parents make decisions. In short, the motivation for learning seemed to be more extrinsic than intrinsic. One student in the afternoon school program even told us that she liked to use the shared computer because acquiring computing literacy will enhance her employment prospects.

One possibility is to design videogames as an extension of the interactive quizzes in the courseware deployed in the Digdarshan pilot program. These e-learning games will aim to make the learning experience more interesting than existing interactive exercises such as multiple-choice questions, and will build on prior observations of rural children's enthusiasm to replay the exercises until they have attained the highest possible score. The challenge with e-learning games is that rural parents need to perceive these applications as educational and not entertainment, as the earlier incident regarding parental complaint about movie screening on a Digdarshan pilot computer illustrates. One way to encourage greater acceptance for e-learning games is to design them such that their content focuses on an educational topic that parents believe to have practical value.

In some regions such as Indian villages, rich oral traditions have made storytelling a dominant mode of communication. Moreover, since rural children have been observed to be motivated and persistent in creating multimedia digital stories that explain academic concepts, we believe that many more children can be given authoring tools to create such content. But given the usability obstacles in existing tools such as Microsoft PowerPoint that we have observed above, such tools must be significantly more usable, compared to existing tools such as PowerPoint or Adobe Premiere, for children aged 8 to 10 with little computing exposure.

We therefore envision a suite of authoring tools that enable children to work collaboratively in small teams to author digital stories and simple e-learning games. The primary objective is to foster active learning, in which learners discover and address gaps in their understanding when explaining concepts to others. Constructionist learning, in which children-producers design digital artifacts such as games that aim to teach their peers, have been shown to bring positive learning outcomes to the producers themselves (Harel 1991; Kafai 1994). Similarly, cooperative learning in small teams, which has been demonstrated to yield learning benefits (Johnson and Johnson 1989), is another objective. Other expected learning outcomes include teamwork skills, independent thinking (vs. rote learning), as well as writing and communication skills.

In the longer-term, student authored content can plausibly supplement existing textbooks, which respondents perceived to have mediocre educational quality. The same content can also address the shortage of local language courseware and can be used as teaching aids by older students to coach their juniors.

However, the limited time that students have to use shared computers suggests that user interfaces for educational applications in the context of rural education must be designed to be usable, learnable and efficient. That is, users who lack practice or enough time to use a computer should nevertheless make few errors, understand and remember how to use the computer in the absence of regular practice, and be able to perform their tasks without taking excessive steps or time. These requirements, especially the last one, represent broad research opportunities in human-machine interaction methods (especially for user input) and user interfaces. In particular, our observations about children’s lack of proficiency with the keyboard (and Microsoft PowerPoint) suggest that a pen-based user interface such as an inexpensive Tablet PC which allows the user to write (vs. type) on it may be more usable.

At the same time, the shortage of building space in rural Indian schools that can be secured against theft implies that it is difficult to house more computers in existing school premises even if the financial means are available to obtain more machines, unless more classrooms could be built concurrently. Moreover, the shortage of stable electricity inside rural classrooms, alongside the abundance of open fields inside and around rural schools, lead us to believe that with the appropriate educational models and software applications for outdoor spaces, battery-powered mobile devices can overcome space-related and other infrastructural constraints. Most importantly, given the limited time (i.e. 10 to 15 minutes) that rural students have in school to use the shared computers, it is not clear if acquiring more machines will lead to greater computer usage in schools. On the other hand, it is possible that students may have more time at home and other out-of-school settings to use mobile devices if children could borrow the equipment home. This offers more learning opportunities because learners will then have more time to interact with the educational artifacts.

We expect that shared computers in Internet kiosks and schools will remain the most economical means in the conceivable future for providing computing access for children whose families cannot afford these equipment. But we speculate that as scarce resources, the computers could be used more efficiently if augmented by paper-based practices, which are relatively more prevalent. Here is an illustrative scenario: children draw, on paper, the individual slides that comprise a digital story, either at home or in school. A low-cost digital camera is used to capture these sketches once they are finished, after which children color, manipulate and organize these sketches electronically into a slideshow. By preparing as much of the multimedia story as possible offline, the shared computers can be made available for more users. In a *déjà vu* sense, this shared usage of computers in developing regions parallels the batched use of mainframes when they first appeared. In terms of learning, having students perform more tasks offline and restricting the amount of time that they have on the computer will help them avoid the "perfectionist urge” to keep tweaking their work when composing it online. When students are less able to spend time perfecting the aesthetics of the end product, they can focus better on their original learning tasks.

We investigate the feasibility of some of the above ideas in subsequent fieldwork, which we describe in the next Chapter.

3 Exploratory Studies

Our needs assessment described in Chapter 2 had given us some ideas on how we could develop more innovative computing software applications to empower low-income learner communities in both rural areas and the urban slums in India. These ideas include constructionist learning by having students author digital stories and e-learning games as producers of digital e-learning content, and students playing educational games that target English as a Second Language as consumers. However, given that high technology is not an everyday feature in the lives of our target users and their family members, we feel it is only imperative that we tread cautiously when introducing technology into the fabric of their social lives. Specifically, it is necessary to explore the feasibility of the above ideas through exploratory studies, so as to investigate their acceptability in the targeted cultural context, before we proceed to design and implement them as a viable research agenda.

This Chapter describes the lessons from the exploratory studies that we conducted to investigate the feasibility of the above ideas. These studies involved children from the urban slums and a village in North India, in which we introduced technological artifacts that embodied the above ideas. Our objective was to examine the short-term uses of these artifacts *in situ*, in which these artifacts constitute a part of the broader cultural repertoire. In other words, we view culture as a phenomenon that is mutually co-constituted and co-constructed by various actors using the technological tools at their disposal, such that the introduction of computing technology facilitates potentially profound changes in culture.

On the whole, we found constructionist learning to be challenging to implement – at least in the short-term – in resource-poor areas owing to the lack of adults and other people who could scaffold the children's learning processes. Constructionist learning also represents a dramatic departure from existing rote-based teaching methods and appears to require nothing less than major systemic educational reforms for it to succeed. On the other hand, we obtained promising results with interactive educational games on desktop computers. Most importantly, there were early signs that mobile learning could be an effective means for complementing the current formal schooling system.

3.1 Digital Storytelling

We experimented with multimedia storytelling using KidPix, which is a software published by Broderbund that can be thought of as “Microsoft PowerPoint for children.” KidPix bears similarities with the digital story authoring application that we envision, and our objective was to assess how well prospective end-users will accept such an authoring tool if it contains much less usability challenges for children compared to PowerPoint. A related objective was to identify limitations in KidPix that can be addressed in a proposed redesign. Our observations in Chapter 2 about rural children’s lack of proficiency with the keyboard (and PowerPoint) suggests that a pen-based user interface such as the Tablet PC which allows the user to write (vs. type) on it may be more usable. As such, for our exploratory study with digital storytelling, we made use of Tablet PCs as the platform for running KidPix. The Tablet PCs either came with built-in microphones or were attached to low-cost microphones. The study also included two inexpensive digital cameras. These cameras were simple and have only two buttons, namely, an on/off switch and a button for taking photographs.

Our exploratory study on multimedia storytelling was conducted at the afternoon school for urban slums girls, and the reader can refer to Chapter 2 for more details on this program. In total, 15 teenage girls from this program participated in the study, for 3 hours per child on average. They were taught to use KidPix and the Tablet PCs, after which they were asked to share the Tablets in small groups to create simple digital stories using KidPix. Similarly, we showed participants how to use the cameras to take photographs of their everyday lives and to import the photographs into their digital stories, so as to help their peers, who were the audience, better relate to the content in the digital stories.

In practice, the children learned how to use the cameras in less than 10 minutes. More importantly, we found that the children were responsible with the cameras. We had loaned the cameras to them to take photographs of their surroundings, and they returned the equipment promptly and in their original condition after the KidPix sessions. Most of all, within three hours, we succeeded in training participants to write on the Tablet PC screens using the stylus, make audio-clips from their voice recordings, create individual slides, import photographs from the digital cameras into slides, and arrange the slides into a slideshow. Participants reacted positively to the Tablet PCs due to their compact form factor, in that the slate Tablet PC that we used in the study (i.e. the Compaq TC-1000 model) was more conducive to the smaller size of the female palm: “I like this because it can be just held in the hand.” Participants added that they preferred the Tablets to desktop computers. Furthermore, the portability of the Tablet PCs made the participants intrigued with the possibility of taking these devices outside the premises of the afternoon school to use.

Whereas we observed students experiencing usability problems with PowerPoint during the needs assessment (Chapter 2), in our exploratory study, students were creating attractive slides using KidPix on the Tablet PCs. But participants did not use KidPix to organize individual slides into a digital story. Neither did participants annotate the slides with text using the stylus. We were not expecting participants to create multimedia stories spontaneously since we have yet to train them to come up with a narrative script for the digital story and to compose individual slides into a slideshow based on the script. But we were surprised to observe that participants focused on beautifying individual slides,which we later realized to be less cognitively demanding compared to the more challenging task of organizing slides into a coherent sequence. We also learned that the children did not enjoy writing and prefer to draw. These observations suggest that digital storytelling on a large scale in developing regions is not feasible unless there are sufficient well-educated adults and other mentors in the community who can scaffold children throughout the authoring and learning process.

Worse, there appears to be limited time under the present curriculum and lifestyle for computer usage. For example, interviews with parents from the urban slums indicated that their daughters were expected to perform time-consuming household chores as a "rite of passage" that prepares them for married life. Moreover, participants failed to show up for two KidPix sessions because they were participating in festivals. As a consequence, we were forced to cancel both sessions. Our exploratory studies marked our first attempt to schedule time with children in developing regions to observe their use of computers, and these difficulties reinforce the observation that it can be challenging for them to set aside structured blocks of time to use computers.

3.2 Game Authoring

Our initial experiences in working with urban slums children to create their own digital stories helped us to appreciate the scaffolding that we needed to provide children if they were to learn as designers. We followed up on these early experiences by running a related exploratory study on game authoring with rural children. Our goal was to guide them to design e-learning games, so that in the process, we could identify the primitives and basic abstractions that game authoring tools for this user group will need to support.

This exploratory study took place at the first village school which participated in the Digdarshan pilot program. This school was located in a remote rural area and we had to travel 75 kilometers (1½ hours) each way to reach it. It has about 250 Hindi-speaking students (grades 1 to 8) whose classes were housed in over five classrooms in two single-story buildings. Regular classes took place only in the morning. The computing facility at this school has expanded from one computer during the Digdarshan pilot to a computing center with three computers when we did this study. Computer lessons were scheduled to be held on a weekly basis for the students, but these classes were often disrupted by the frequent electricity outage. As a result, computer classes were either cancelled or reverted from practical, hands-on training sessions to the lecture format. The study took place in the afternoon so as not to disrupt classes in the morning, and spanned a total of 10 days.

Each afternoon sessions lasted between 2½ and 3 hours.

We expected that it was possible to recruit up to 12 rural students from the same grade, and hence indicated our preference to have 12 participants from grade 5 such that they reflected a balance in terms of sex, academic performance and computer literacy. However, these criteria were difficult to achieve in practice because the principal wanted to impress his foreign visitors (i.e. us) and we learned after the study that he intervened in the selection procedure. As such, our participants comprised 10 girls and 2 boys from grades 4 to 8. Two of them did not know their age, while the remaining participants were between 10 and 16 years old. They were chosen by their principal because they excelled among their peers where academic performance and computer literacy were concerned. We learned that their computing experience was limited to Windows Paint, some features of Microsoft Word and some videogames, such that only one of them could use the basic features in PowerPoint. As a token of appreciation, each participant received US$20 of stationery.

To give each participant enough opportunities to contribute to the design process, we separated the 12 of them into 4 groups of 3 each. Each group was handed a Tablet PC and low-tech prototyping material. The 10-day period was divided into a warm-up phase (3 days), a low-tech prototyping phase (3 days), a hi-tech prototyping phase (3 days) and a final day for interviewing participants on their prototyping experiences.

3.2.1 Warm-Up Phase

This phase was intended to prepare participants to design and prototype the e-learning games. Firstly, we needed to get those participants who lacked practice or confidence with computers "up to speed," because basic familiarity with the graphical user-interface as a user or tester is a first step to becoming a design informant or partner. Secondly, we wanted to engage participants' creativity through an appropriate warm-up exercise. To meet both goals, we guided them to create photo collages by using Windows

Paint to compose multiple photographs taken with inexpensive digital cameras. As a side benefit, this exercise allowed us to understand their computing background better.

Thirdly, we gave participants the opportunity to play at least one e-learning game, so as to provide them with a concrete example of an educational game that will be helpful when they have to brainstorm ideas for e-learning games that target English as a Second Language. We chose Word Munchers after pretesting a selection of ESL learning games with a 10-year-old Taiwanese boy who was learning ESL. Word Munchers develops the learner's vocabulary through inductive learning, and will be described in greater detail in the next section.

Fourthly, we wanted to better understand the everyday lives of the participants but lacked time for more extensive fieldwork. Inspired by how digital cameras were used as cultural probes (Gaver, Dunne and Pacenti 1999), we asked participants to take turn to borrow the cameras home overnight and over the weekends. In particular, we asked them to take photographs of everyday activities which they think would be interesting to us, so that they could later describe selected photographs to us. There was a day with a heavy downpour. Yet, the rural children with the cameras were responsible in taking good care of the devices and kept them from getting wet. This incident lends further support for the idea of mobile learning in this context.

Finally, we took measures to help rural parents appreciate the educational value of the study, so that they would continue to support their children's participation throughout the duration of the study. One of the steps that we took was to wrap up every session with a 5-minute debriefing, during which we helped participants to review what they learned in that session. In this way, participants were prepared to "report" what they had learned about English and computer literacy to their parents when they reached home, should their parents inquire.

3.2.2 Low-Tech Prototyping Phase

We began this phase by asking participants to think, individually, of fun ways to teach English as a warm-up for designing the same games in groups. Participants seemed to find it stressful to brainstorm ideas, however. We thought that the difficulty was due to their limited command of English, and hence helped every group to mock-up a game on paper that aimed to teach us six Hindi words within the context of a shopping scenario, as an intermediate step. We then helped participants to adapt their Hindi learning games to target ESL. We observed that participants were highly frustrated throughout this phase; it took tremendous effort before we designed two word matching games, a kinesthetic game similar to but more elaborate than "hopscotch" and a skit-based game.

It seemed that low-tech prototyping was frustrating because participants found it difficult to come up with initial ideas and to iterate on their initial designs. To illustrate the tensions in these moments, we show a typical low-tech prototyping scene between a participant (P), a researcher (R) and facilitator (F):

> P: I can just tell you the meanings whenever you want.
>
> R: But I want to learn the meanings from the game itself. I like this game, but how can I learn the words from the game itself?
>
> (10 minutes pass; participants take frequent breaks and look unenthusiastic.)
>
> P: I can't think of anything.
>
> F: Nothing? Try and change the game that you have.
>
> P: I don't understand.
>
> F: So what don't you understand? You can't learn meanings from this game, can you?
>
> P: No, I don't know what I can change.

Even though participants should have found low-tech prototyping to be easier by drawing inspirations from Word Munchers, they continued to respond with blank stares and silence after we asked them to recall the features in this game. As such, we spent an additional day exposing them to more ESL learning videogames. We chose the collection of ESL learning games in Clifford Reading Pack since it included word matching games that resembled that word matching games that participants came up with. By introducing participants to Clifford, we hoped to show participants how their designs might look and behave after being implemented as software.

3.2.3 Hi-Tech Prototyping Phase

We switched from low-tech to hi-tech prototyping due to the above unsuccessful experience with low-tech prototyping. We speculated that participants may find hi-tech prototyping to be more enjoyable since they asked, throughout the low-tech prototyping phase, when they could resume using the Tablet PCs. For the hi-tech prototyping phase, we selected Stagecast Creator – an end-user programming for children and a run-time environment for interactive simulations – as our prototyping tool for two reasons. First, it has an extensive library of themed backgrounds and characters that allowed participants to leverage their computer skills from the earlier collage building exercise to create hi-tech prototypes by dragging virtual characters around on themed backgrounds. Second, Scaife et al. (1997) found that arranging laminated cut-out characters on a background scene was found to be an effective low-tech prototyping technique with children in terms of helping them to generate ideas, and our application of Stagecast Creator was based on this model.

We worked with participants to design and prototype matching games for ESL learning using Stagecast Creator. Although Creator was among the most usable end-user programming tools that we were aware of, it was still relatively complex. As such, we not only used example simulations to show participants what they could build using Creator, we also used Creator directly on their behalf on most occasions. Similarly, due to time constraints, we only implemented those ideas that were easy to implement using Creator. Nonetheless, although the hi-tech prototypes were mostly left unfinished, they embodied far richer ideas than the low-tech prototypes. More importantly, participants exhibited a dramatically higher level of enthusiasm for hi-tech prototyping, in comparison with low-tech prototyping. A characteristic hi-tech prototyping scene looked like this:

> R: So now you know how to do animations. How can you change your game to be as fun as Clifford?
>
> P1: Using animations for example? We can make it move left and right…. we can make things go left and right like in Clifford.
>
> R: Right now, you're matching labels to pictures. How can you change that to be more fun… What did Clifford do every time you got something right?
>
> P2: (nods)…When we see a door, we can make the door open. The lion in the picture can roar. The man can fall at some point.
>
> P3: If we write the man's name correctly in the box, maybe he can climb up the ladder. Can we do that?
>
> R: Sure, yes you could do that. What else would you want to do?
>
> P2: When the elephant comes on the screen it should trumpet.

Participants also expressed pride and ownership over their hi-tech prototypes:

> F: How does [this hi-tech prototype] compare to Clifford?
>
> P: I like this game more.
>
> F: What is in this game that you like that's not in Clifford?
>
> P: The numerical point system and the messages constantly spoken by the central character.

We observed that participants used primitives such as animations, audio playback and textual labels in Stagecast Creator as building blocks in their designs. Additional comparative studies between low-tech and hi-tech prototyping, as well as more research into the latter as a candidate methodology for participatory design with rural children, are needed since the latter departs from several previous participatory design attempts with children. At a first glance, our findings ran counter to most of the literature (Druin 1999; Scaife et al. 1997), which showed and/or assumed that low-tech prototyping was effective with children. It was likely that rural students had little exposure to software and could not imagine how software worked via low-tech representations. We also concede that our hi-tech prototyping attempt might have been just as unsuccessful as low-tech prototyping if we had started the former on a "blank slate", as opposed to showing participants how they could import characters from Stagecast Creator's libraries and drag these characters around on a themed background. It was likely that our attempt at high-tech prototyping resulted in relatively more creative designs, because it was carried out in accordance to the laminated cut-out approach proposed in Scaife et al. (1997), which was found to be effective in stimulating children to generate ideas for low-tech prototypes. In other words, our relative success with hi-tech prototyping could be attributed to us employing low-tech prototyping best practices with a hi-tech prototyping tool.

3.3 E-Learning Game

Word Munchers (Figure 3.1) is a vocabulary building game based on inductive learning. The game board is initialized such that each grid contains an object. A word or phrase (e.g. "red things") is then shown at the top of the screen. Some of the objects in the grids match the given word or phrase while others do not. The learner is required to move the Muncher character around the board to gobble all objects that meet the given description, in order to win the game. The learner starts with three lives, and loses one life each time he gobbles an object that fails to match the given description. Based on the player's successes and failures, he is expected to induce the meaning of the given word or phrase.

Figure 3.1. A screenshot of the Word Munchers Deluxe game for vocabulary building, in which the player's goal is to move the Muncher character around the game board to consume all objects that match the word or phrase displayed.

Pre- and post-tests administered for Word Munchers Deluxe showed an average post-test gain of 1.3 out of a total of 6 points on a t-test ($p < 0.01$, $\sigma = 1.4$). These results are by no means conclusive, but indicate that ESL learning games for developing regions are feasible and merit further study. More importantly, even though rural participants had attended ESL classes for three years on average and were the "star" students who were specially hand-picked by their principal for us to interact with, we observed that many of them struggled to read some letters in the English alphabet or decode words phonetically with accuracy. We had originally wanted to work with the oldest primary school children so that we could explore a greater range of learning possibilities with technology, such as tools that facilitate student authoring of digital artifacts. But our direct observation of the participants' dismal English baseline, coupled with the above obstacles behind student authoring – possibly attributed to the fact that their schooling revolved so much around rote learning (UNESCO 2005) that they have very little exposure to design activities or classroom exercises involving creativity – made us realize that our top priority is to return to the basics, such as the English alphabet and rudimentary vocabulary.

But for any solution to be scalable, it must overcome barriers in the form of poor infrastructure and limited learner time to access educational resources, including school. We believe that mobile learning on devices such as cellphones – which are increasingly being adopted in developing regions – hold the key for dramatically expanding the reach of educational delivery, such that poor children with work commitments can access ESL learning in out-of-school settings, in places and times that are far more convenient than school. We have also witnessed that poor children who do not own mobile equipment such as digital cameras in their homes can nonetheless be trusted to care for the devices.

Now that we have given an account of the exploratory studies that shaped the direction of the MILLEE project, the rest of this thesis describes our design approach and findings.

4 Pedagogical Framework

In this chapter, we present an overview of the Task-Based Language Teaching (TBLT) framework that we draw on to inform our instructional design. We choose to situate our technology and instructional design processes within the TBLT curriculum development framework for two reasons. First, the task is an appealing object of design and analysis for both technology designers and language teachers. As a construct, the task is an artifact for design with an inherent degree of structure, which technology designers and language educators can follow. The task in fact fits well with existing work practices of designers and educators: just as the language instructor can plan her teaching tasks prior to her classroom lessons, the instructional designer can devise tasks for computer-assisted learning systems that are eventually deployed with learners.

Second, and more importantly, while it remains to be shown that task-based instruction is superior to other language teaching approaches, it is nonetheless relatively well understood to the extent that there is a substantial body of literature which describes its applications and limitations. In particular, for the purposes of our project, TBLT has been implemented in India with some success.

In the rest of this chapter, we first explain what a pedagogic task is, and describe the characteristics and components that constitute a task. Second, we outline how pre-task and post-task activities may accompany tasks, as well as how tasks can be composed with other tasks to form larger instructional sequences. Third, we present an overview of the empirical basis for TBLT and address some of the criticisms commonly leveled against it. Finally, we explain how we envision the place of technology and instructional design in a TBLT framework. Due to limited space, it is not possible to provide a thorough treatment of the broader discipline of second language acquisition. Instead, we refer the reader to introductory texts such as Lightbown and Spada (2006).

4.1 Task Characteristics and Components

In TBLT, the learner engages with a series of pedagogic tasks that take the form of goal-directed activities. Tasks may involve problem-solving activities, communication strategies (e.g. the notion of strategic competence proposed in Canale and Swain 1980), puzzles, pictures, dialogues, decision-making, role playing and matching. Prabhu (1987) outlines problem-solving activities in which the learner has to communicate information to another learner who does not already know it (i.e. an information-gap activity), draw inferences from information that is encoded in language (i.e. a reasoning-gap activity), or provide a personal response (i.e. an opinion-gap activity). It can therefore be observed that TBLT belongs to the family of Communicative Language Teaching approaches. In CLT, language is viewed as a means for communication that is centered around meaning, as opposed to a set of phonological, grammatical or lexical elements.

Several definitions of the task have been proposed in the literature. For instance, Prabhu defines a task as "an activity which required learners to arrive at an outcome from given information through some process of thought, and which allowed teachers to control and regulate that process" (1987: 24). Despite the differences between various definitions, however, Ellis (2003), Nunan (2004) and Prabhu (ibid) appear to converge on the task as an activity that requires learners to employ their linguistic, cognitive and other resources so as to process the target language. More important, Skehan (1998a) observes that the diverging definitions nevertheless agree on five points about how the task differs from any other learning activity, in that a task:

1. is focused on meaning and not the manipulation of form,
2. culminates in creative (vs. reproductive) usage of the target language,
3. bears a relationship to a real-world task or discourse,
4. emphasizes the completion of a (communicative) goal, and
5. is evaluated on the basis of its outcome.

Nunan (2004) distinguishes between two types of pedagogic tasks, namely, the rehearsal task and the activation task. A rehearsal task resembles a real-world task and is intended to provide the learner with opportunities to practice a simplified version of the real-world task. On the other end of the continuum is the activation task, which is meant to provide the learner with opportunities to develop her linguistic ability, from using the language reproductively to using it creatively.

Nunan (ibid) synthesizes earlier work on TBLT and propose that tasks are made up of the following components:

Goals – which reflect the learning objectives behind the task. Nunan believes that goals should be broader than exact performative outcomes but be more specific than the transactional, interpersonal and aesthetic macrofunctions proposed by Halliday (1985).

Input – which represents an optimal combination of authentic (e.g. newspapers) and pedagogic material (e.g. the Ladybird *Key Words Reading Scheme* and Scott Foresman series of pedagogic readers, and Voice of America radio broadcasts expressed using Special English). Krashen's controversial "comprehensible input" hypothesis states that language acquisition takes place when the learner is exposed to input slightly beyond his current stage of language development (1987). Long (1985) adds that it is the act of negotiating meaning with other speakers that renders the input comprehensible. Unlike Krashen, Swain (1985) argues that it is the attempt to produce comprehensible output, and not input processing, which taps syntactic processing behind language development. As such, both Long and Swain hold that output production – in addition to input processing – has a role to play in language acquisition.

Procedures – which indicate the steps that learners are required to take with the input. According to Skehan (1998a), procedures can focus on accuracy, fluency or complexity with the language. Schmidt (1990) adds that linguistic forms in the input can be acquired only if it is "noticed." Procedures which draw the learner's attention to such forms can help in this respect.

Roles – which learners and teachers are expected to perform toward the learning goals, as well as the relationships between them.

Settings – which include the classroom or out-of-school learning environment, whether the task is meant for the entire class, small groups, pairs or individual learners, and the extent to which the task is self-paced vs. teacher-directed.

Tasks can vary in their difficulty depending on the task variables. Nunan (2004) indicates that a task becomes harder as the input for the learner to process becomes more authentic. Similarly, a task that involves more authentic procedures, i.e. procedures that require the learner to employ the language in ways that are closer to real-world tasks as opposed to reproducing the language given in the input, is more difficult (ibid). Where task type is concerned, Prabhu (1987) believes an information-gap activity is less difficult than a reasoning-gap activity. Finally, Skehan (1998a) advocates that task difficulty be adjusted (through their characteristics and components, as well as the pre- and post-task activities) so as to balance the learner's development between complexity in the language involved, fluency in employing the language, and accuracy in terms of the correctness of the language used.

4.2 Instructional Sequences around Tasks

Now that we have described what a task is, it is time to consider how individual activities and tasks may be assembled into a more complete curriculum or lesson plan. Nunan (2004) offers the following instructional sequences around tasks, in which the first 5 steps constitute pre-task activities while step 6 is the actual task:

- Step 1: Schema-building exercises that introduce some of the key vocabulary, linguistic forms and context for the task.

- Step 2: Communicative exercises that provide controlled practice in using the vocabulary and linguistic units that make up the task.
- Step 3: Listening to how the above vocabulary and forms are used by actual speakers in authentic settings.
- Step 4: Language exercises that cover the rules behind the linguistic units in the task.
- Step 5: Freer practice (compared to the controlled practice in step 2) through activities such as the information gap, reasoning gap or opinion gap activity, possibly with learners interacting with each other in pairs.
- Step 6: The pedagogic task proper, to be carried out in possibly larger groups.

Nunan notes that although steps 1 to 5 may resemble a syllabus that emphasizes form over meaning, the learners have "seen, heard and spoken the target language within a communicative context" (ibid: 32). In other words, even though the pre-task activities may focus more on form than meaning, they are nonetheless essential for developing the learner's linguistic foundation to engage in communicative tasks. By focusing the learner on the phonological, lexical and grammatical elements in the target language, or having him practice with limited communicative activities, pre-task activities help to prepare him for the eventual task. Likewise, Skehan (1998a) argues that pre-task activities can be used to focus the learner on form through attentional manipulation, i.e. Schmidt's notion of "noticing" specific linguistic forms (1990), which is crucial for their uptake, and hence the restructuring of the learner's interlanguage (McLaughlin 1990). Conversely, post-task activities can engage the learner in reflecting on what was covered earlier and in pointing out areas for improvement, so as to create further opportunities for restructuring.

Nunan (ibid) and Prabhu (1987) outline some principles that are employed when designing instructional sequences around tasks. Activities should build on the material in earlier activities, provide an appropriate level of scaffold for the language targeted in the current activity, transition from receptive (i.e. listening and reading) uses to reproductive uses (i.e. speaking and writing), provide repeated exposure and opportunities to use the targeted linguistic items, integrate the linguistic items so as to elucidate the connection between semantics, syntax and pragmatics, transition the learner from reproducing the target language to using it creatively, as well as opportunities to reflect on what has been covered and how well it has been learned. It is also common to adopt a communication-oriented syllabus (Canale and Swain 1980) that is functionally organized, even as new grammar rules are progressively introduced and reinforced. In fact, the same rules can be repeated across multiple tasks, so as to promote further development of automaticity and restructuring (McLaughlin 1990). It is therefore observed that the above principles reflect the Piagetian concrete-abstract progression that has become an accepted notion in the learning sciences.

4.3 Empirical Basis and Criticisms of Task-Based Instruction

Various aspects of task-based language teaching are based on previous research in psycholinguistics (c.f. Nunan 2004). On the other hand, TBLT has been examined to a significantly lesser extent in sociolinguistics research (c.f. Lantolf 2000). Similarly, there remains limited empirical evidence in the literature on the implementation of tasks and their results in actual classrooms. Van den Branden, Van Gorp, and Verhelst 2007 is an edited volume that was recently put together to address this gap. It comprises a collection of papers that describe the actual tasks, their implementation in classroom settings from Europe to China with learners who include primary school and kindergarten children, and the learning processes observed.

This edited volume provides empirical support from classroom studies for task-based language teaching. These studies show that the benefits which learners accrue from a task depend on the intensity with which the input language is processed, the volume and quality of the language produced by learners and facilitators, and opportunities to reflect on the task. The results lend further credence to the psycholinguistic theories underlying task-based instruction. They suggest that interventions that aim to help learners negotiate meaning, so as to make input more comprehensible, output more comprehensible, fluent or accurate, and/or language processed more intensively, ensure stronger outcomes.

In the Indian context, Prabhu (1987) describes the well-known Bangalore/Madras Communicational Teaching Project, in which task-based approaches for teaching English as a Second Language in India were experimented with children aged between 8 and 13 years (grades 3 to 8) over a five-year period. In an independent evaluation, Beretta and Davies (1985) found that children who were taught using TBLT performed significantly better than their counterparts in the control group on tests of transfer that evaluated them on contextualized grammar, dictation, and comprehension (both listening and reading). Whereas TBLT focuses on meaning, students in the control group were taught using the traditional form-focused Structural-Oral-Situation method, which involved substitution drills with linguistic structures on a syllabus that is structurally and lexically graded.

This chapter will not be complete without addressing the criticisms against task-based instruction. The most common criticism is perhaps that there is limited empirical support for the psycholinguistic theories behind task-based teaching. For instance, Swan (2005) observes that although Schmidt and Frota (1986) propose the hypothesis that the "noticing" of linguistic forms in the input is necessary for restructuring the learner's interlanguage, this hypothesis is only supported by Schmidt's diary studies of how he had learned Portuguese over a six-month period. Similarly, Ellis (2000) notes that there is little empirical evidence to support Swain's "comprehensible output" hypothesis (1985) and Long's "interaction hypothesis" (1985).

Next, Swan (2005) argues that "acquisition" in task-based learning is usually construed narrowly to mean the acquisition of grammar, to the exclusion of other types of linguistic items. As such, task-based instruction inadvertently upholds a view of linguistic competence that privileges grammatical competence over competencies such as discourse competence, sociolinguistics competence and strategic competence pointed out by Canale and Swain (1980). This bias toward grammatical competence has dire implications when learners exhibit a tendency to complete tasks as soon as possible by using fewer linguistic forms in engaging with one another over tasks, since this tendency leads to impoverished language (Seedhouse 1999). That is, learners need to communicate with one another to accomplish a task, and their focus is on accomplishing the task rather than the language that they use in it.

How should we interpret the above criticisms, especially that of limited empirical evidence, given the above classroom studies that have demonstrated benefits from task-based teaching? In our opinion, more research is needed to investigate and account for the exact psychological processes behind second language acquisition and how positive learning outcomes are attained. Our view resonates with Ellis (2000), who concludes that psycholinguistic approaches have yet to show how a task's design corresponds directly to the acquisition of communicative competency. Another significant observation is that the learner may view the task differently from how the task designer views it (Nunan 2004). Most importantly, the above criticisms suggest that the psycholinguistic perspective may not be the only way to explain the above classroom results – much less offer an adequate account for the divergence between learning processes and the above outcomes observed. Furthermore, the criticisms expose the limitations in a purely psycholinguistic approach, which neither considers learner agency nor how the learner views the task. In the next section, we introduce an alternative view of language learning, i.e. the sociocultural view, and contrast it against the psycholinguistic tradition.

4.4 Tasks in the Psycholinguistic vs. Sociocultural Traditions

Task-Based Language Teaching is a methodology that is surprisingly compatible with both psycholinguistic and sociocultural theories of language learning. However, Kinginger has argued that the psycholinguistic and sociocultural traditions do not share a "complementary relationship" (2001: 418). In fact, she states that the two traditions hold very different and distinct assumptions about language, learning and mind. At the most fundamental level, the psycholinguistic view asserts that cognition is a phenomenon which exists in the individual, whereas the sociocultural perspective views the mind as a social construct that arises as the result of semiotic mediation. Despite the incompatibility between both traditions at the philosophical level, scholars in language and literacy have attempted to reconcile both traditions in their respective frameworks. This section will briefly examine their differences and how they pertain to task-based instruction.

According to the psycholinguistic tradition, the task is a device that comprises input for the learner to process according to a set of procedures, such that the learner acquires the target language by processing this input, possibly by having to deploy her interlanguage. Work on task-based instruction in the psycholinguistic tradition have tried to identify properties about tasks that influence how they are processed, as a basis for sequencing tasks according to their cognitive demands on the learner. On the other hand, in the sociocultural perspective, the task is the basic unit that forms the basis for social interaction among learners and their facilitators, i.e. the zone of proximal development (Vygotsky 1978, 1986). According to this tradition, language development thus comes about not through input processing, but is an emergent phenomenon that results from the internalization of verbally-mediated communication.

Ellis (2000) notes that psycholinguistics has the tendency to view the task as having a predictive – and sometimes even deterministic – outcome based on the task's features, whereas sociocultural theories of learning imply that the learners' interpretation of and orientation toward the task determines its outcome. Donato (2000) adds that the student's orientation to the task is more crucial than the task's outcome. Furthermore, learner orientation may change in the middle of a task (Appel and Lantolf 1994), or when the means of mediation changes (Thorne 1999). As such, Ellis distinguishes between the task as a "workplan," which is the lesson activity planned by the task designer prior to the class, and the task as a "process," which refers to how the planned task is actually enacted by its participants in the course of the class (2000). Contextual factors and learner agency are therefore salient since the task-as-process may unfold differently from how the task designer had originally conceived of the task-as-workplan.

The sociocultural view thus represents a qualitative shift from focusing on the intrinsic properties of the task to examining the social processes and interactions that learners engage in around the task, given the underlying social factors. However, Ellis (ibid) maintains that task characteristics do influence task performance – even if not in a deterministic fashion – suggesting that a psycholinguistic notion of task-based learning can be incorporated into a broader sociocultural framework of task-based learning when the task is considered as an artifact that mediates social interaction. In short, Ellis (ibid) seems to argue that task-based language learning calls for a balance between the planning vs. improvisational ends of the continuum in teaching (Van Lier 1996).

4.5 The Place of Technology in Task-Based Instruction

The above discussion has underscored that one way to resolve the differences between the psycholinguistic and sociocultural views of task-based language teaching is to subsume the former within the latter, such that it is the general genetic law of cultural development (Vygotsky 1978, 1986) and not input processing that accounts for language acquisition. That is, previous research in task-based instruction in the psycholinguistics tradition can potentially be reinterpreted from the sociocultural standpoint. This position implies that tasks can be designed to foster communication between learners and other interlocutors, such that it is from this social interaction and scaffolding that language development emerges.

More importantly, digital activities such as word-picture matching activities and conversations with non-player characters in game settings can be designed as pre-task activities that prepare learners to eventually engage in the communicative tasks proper. Similarly, technological artifacts can display customized information that learners have to draw on in performing reasoning-gap or information-gap activities in the tasks proper.

Finally, digital activities can pose questions for learners to reflect on or provide high-level feedback (which would have been disruptive if given during the task itself) as part of their post-task activities.

In other words, we envision that technology can play supporting roles in task-based teaching. For instance, technology can take the form of digital activities during the pre- and post-task stages that scaffold the learner in building foundational linguistic skills and reviewing the target linguistic units respectively. Technology can also take the form of artifacts that constitute one of the means for semiotic mediation in the actual tasks themselves. We believe that these roles for technology in language learning are not only applicable to task-based language teaching, but are sufficiently broad and inclusive to be compatible with other pedagogical frameworks.

5 Pedagogical Design Patterns

In Chapter 4, we have described the Task-Based Language Teaching framework which guides our instructional design. The challenge, however, is that theories are often too abstract for technology designers to grasp easily, much less translate into software designs. Most technology designers also lack a background in language acquisition and teaching English as a Second Language. At a cursory level, the design process calls for multidisciplinary collaboration between designers and educators, in addition to specialists from other domains such as game design. But in our experience, the typical language teacher may not know how to draw on her teaching experiences and background, so as to imagine concrete designs for computer-assisted language learning applications.

On the other hand, there are existing commercial language learning products that include games and other software. Some of these products incorporate language learning research and methods into their designs, and are successful on the commercial market. There is therefore some indirect evidence for their pedagogical effectiveness. We do not claim that they are perfect. But we propose that their designs can be a starting point that we revise and iterate on, based on educational theory, classroom experiences and/or field testing. With this principle, we take the state-of-the-art in existing commercial language learning products as the initial basis for our instructional design and avoid reinventing the wheel.

How can we capture existing knowledge on designing effective language learning software? In some domains, design patterns (Alexander 1977) are becoming increasingly popular as compact representations of problems and the solutions that have been found to address these design problems adequately. The primary benefit of a design pattern is to encourage the reuse of existing solutions to problems that are frequently encountered. As such, we can represent and improve on a significant fraction of today's best practices for designing effective language learning software by extracting a set of design patterns from reviewing commercial products, after which we revise these patterns through learning assessments conducted with ESL learners. More important, we believe that it is possible to extract patterns from commercial software in a principled manner by using a language instruction framework, such as task-based language teaching that we covered in Chapter 4, as our analytical framework.

In the rest of this chapter, we explain what design patterns are in greater detail and provide an example. Second, we outline the process that we took when extracting design patterns from commercial language learning packages. Third, we summarize the patterns that we found. Throughout the chapter, we will relate our methodology for distilling the design patterns, and the patterns themselves, to the literature on language acquisition and the psychology of reading whenever appropriate.

5.1 What are Design Patterns?

A design pattern provides insights into a frequent design problem by describing the design problem that is being addressed, the essence of the solution to the problem, the rationale for the solution, how to apply the solution, some of the tradeoffs in applying the solution, and related design patterns (Alexander 1977). Because this description is somewhat abstract, we will next describe how three commercial software applications aim to promote word recognition in terms of developing the learner's sight vocabulary, before showing the pattern that is common across all three instances.

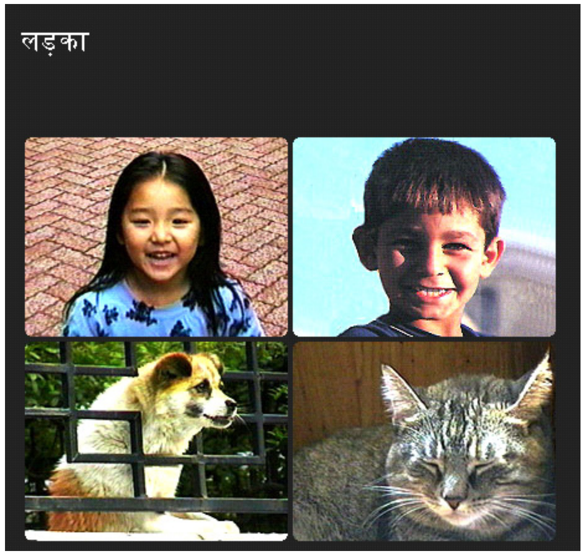

Figure 5.1. An activity in Rosetta Stone (Hindi Level 1: Personal Edition) for vocabulary building and testing.

A vocabulary building activity in the Hindi-learning version of the Rosetta Stone application could work like this: first, it displays the Hindi word for “girl” in the top left hand corner of the screen and highlights the picture of the girl (see Figure 5.1). Next, it displays the Hindi word for “boy” in the top left hand corner of the screen and highlights the picture of the boy. After that, it displays the Hindi word for “dog” even as it highlights the dog’s picture, before displaying the Hindi word for “cat” and highlighting the cat’s picture. Now that the activity has completed the “training phase,” it proceeds to the “testing phase,” during which it shows the Hindi word for either girl, boy, dog or cat on the top left hand corner of the screen and requires the learner to select the picture that corresponds to the word being shown.

Figure 5.2. An activity in Word Munchers Deluxe that takes the inductive learning approach for vocabulary building and testing.

Although we personally found the vocabulary building exercise in Rosetta Stone to be effective in helping us learn to read common Hindi words, it is arguably to boring in that the activity resembles a typical classroom multiple-choice quiz. The Word Munchers Deluxe game (Figure 5.2) attempts to address the same vocabulary building goal but in a presumably more interesting way. It does not have a "training phase" where the learner is explicitly taught the meanings of new words through word-picture associations. Instead, Word Munchers uses an "inductive learning and testing phase," in which the learner is shown a word or phrase (e.g. "red things") at the top of the screen. The game board is also initialized such that each grid contains an object. Some objects match the given word or phrase (e.g. "red things") while others do not fit this label. The learner is required to move his Muncher character around the board to gobble all objects that meet the given description, so as to win the game. The learner starts with three lives, and loses one life each time he gobbles an incorrect object that does not match the given description.

Figure 5.3. Vocabulary testing activity in Who is Oscar Lake? (Spanish learning version).

The *Who is Oscar Lake?* game goes one step further than Word Munchers Deluxe in the extent to which it embraces the whole language instruction approach, as opposed to teaching vocabulary outside the context of use. Figure 5.3 shows a scene in which the learner is directed through Spanish text uttered by a non-player character to buy a train ticket. The English text displayed in Figure 5.3 is an optional scaffold to help the novice Spanish reader, whose is assumed to speak English as his native or near-native language. In this example, the learner is tested on his vocabulary and understanding of the Spanish word for "ticket." But unlike Rosetta Stone and Word Munchers, *Who is Oscar Lake?* is more challenging – and arguably more fun – in that the learner is not simply presented a set of options on the screen to choose from. Instead, he has to move about and explore the game world (Figure 5.4) to locate the ticket seller, which is the option that reflects the correct answer.

Figure 5.4. The ticket seller in the Who is Oscar Lake? game world, which is the "correct answer" for the vocabulary testing "question" posed in Figure 5.3.

Despite superficial differences across the above activities for vocabulary building and/or testing in Rosetta Stone, Word Munchers and *Who is Oscar Lake?*, these activities share a common underlying pattern at the abstract level. To provide the reader with an example of how a design pattern looks, we explicate the following pattern (Figure 5.5) from the above activities. In this example design pattern, steps 1 and 2 correspond to the optional training phase whereas steps 3 and 4 correspond to the testing phase:

Pattern name: Written Word->Semantics Association

Problem: vocabulary building, word recognition

Solution:

Suppose X=4:

1. Displays a word and also displays its meaning pictorially. As an optional step, the meaning of the word can also be conveyed orally and/or textually in the learner's native language
2. Repeat step 1 for X-1 more times
3. Displays one of the X words that was previously displayed during steps 1-2
4. Presents the learner with at least X pictures to choose from, and provide learner with feedback on whether or not his choice was correct or incorrect
5. Repeat steps 3 and 4 for X-1 more times

To reduce level of difficulty:

- read aloud the word in step 3 to the learner, so as to help him learn to decode it

To increase level of difficulty:

- steps 1 and 2 can be omitted
- X can take on a higher value
- the sequence of the X pictures presented in step 4 can be randomized each time step 4 is repeated
- limit amount of time learner is given in step 4
- replace the word in step 1 with a phrase or sentence

Figure 5.5. An example design pattern called "Written Word → Semantics Association" that aims to help the learner associate the written form of a word with its meaning, usually through a word-picture matching activity.

5.2 Sample

We reviewed more than 35 commercial language learning applications in total. We selected this sample based on the following factors, which we adopted as our proxy indicators for educational quality: a large professional customer base, highly-educated users (in the case of adult learners) or parents (in the case of children's software), as well as glowing reviews and/or ratings on home schooling, e-commerce, etc. websites.

Applications were also selected such that the overall sample reflected a balance between listening, reading, speaking and writing skills. However, we should note that the distinction between oral language and written language is not well defined. For instance, there is a reciprocal relationship between spoken and written language, such that word identification processes depend on individual's phonological processing subsystems (Ball and Blachman 1988; Bradley and Bryant 1978; Bradley and Bryant 1983, Fox and Routh 1976; Lundberg, Olofsson and Wall 1980; Tunmer, Herriman and Nesdale 1988; Wagner and Torgesen 1987). Conversely, knowledge about the written language helps to promote metalinguistic awareness (Olson 1994), including phonological awareness (Morais et al. 1979; Shankweiler and Lundquist 1992). As such, some of the patterns that we eventually distilled pertain to language learning tasks across more than one modality.

Our sample included English as a Second Language learning software packages that are developed specifically for non-English-native low-income students from the rural areas and urban slums in developing regions (e.g. the series of software developed by the Azim Premji Foundation in India for use in over 15,000 affiliated rural Indian primary schools), best-sellers in the foreign language learning market (e.g. *Rosetta Stone*, Simon & Schuster's *Pimsleur*, Topic Entertainment's *Instant Immersion* and Auralog's *Tell Me More* series), as well as early literacy games (e.g. the Learning Company's *Reader Rabbit* and Scholastic's *Clifford: The Big Red Dog* series).

5.3 Methodology

Our goal was not only to extract individual patterns from commercial language learning packages, but also to organize these patterns into a comprehensive collection, i.e. a pattern language (Alexander 1977), so that they would be more useful and accessible to technology designers. To our knowledge, the closest work is Clegg, Ogange and Rodseth 2003, who reviewed a sample of digital materials that target English language learning for both native and non-native English speakers in Western Europe and North America, and evaluated the suitability of this sample for ESL learners in Africa. In particular, one of the steps that they carried out was to evaluate the extent to which each software application supports common language learning tasks (cf. task-based language teaching, described in more detail in Chapter 4) for developing skills in listening, reading, writing and speaking. In Table 5.1, we reproduce a subset of their results to better acquaint the reader with the scope of their review.

Task type	Language focus	Used in digital materials sampled
Pre-reading tasks		
Discuss the topic briefly with the learners	All aspects	
Ask the learners to predict what the text will be about	All aspects	
Introduce some of the key vocabulary in the text	Vocabulary	
While- and post-reading tasks		
Read text	All aspects	****
Read text and look at visuals / follow a picture sequence	All aspects	****
Read text and sequence pictures	Vocabulary; discourse	
Match a picture (sequence) with a caption (sequence)	Vocabulary; discourse	
Sequence a set of sentences / captions	Vocabulary; discourse	*
Sequence a set of paragraphs / texts	Discourse	
Link sentences with connections	Discourse	
Read a text and link reference items with antecedents	Discourse	
Read and make/complete notes	Discourse	
Punctuate a text	Discourse	
Sort cards	Vocabulary	***
Classify words in a chart	Vocabulary	**
Match word and picture	Vocabulary	***
Word search	Vocabulary	*
Crossword	Vocabulary	*
Match word and definition	Vocabulary	*
Fill gaps in a text	All aspects	**
Read a text and fill in a chart	All aspects	**
Read text and label a visual (diagram, map, etc.)	All aspects	*

*rarely **sometimes ***often **** very often

Table 5.1. The extent to which activities in a sample of English language learning software support various types of common learning tasks for developing reading skills (reproduced from Clegg, Ogange and Rodseth 2003).

We extended Clegg, Ogange and Rodseth (2003) by undertaking a similar review. But instead of determining the *extent* to which each task is supported by the applications in our sample, we went one step further to identify *how* every type of task is supported by each application, and then noted the actual tasks themselves. More specifically, each time we encountered – for the first time – a task which is supported by an application in our sample, we identified the steps that the application takes to engage with the user (from an instructional design and scaffolding standpoint), and then codified these steps into a new design pattern. Following that, when we encountered an application which supports a task that we have already seen, we expanded the existing pattern for that task by generalizing the pattern description, such that the pattern subsequently included the steps employed by this application.

In refining the pattern description, we aimed to account for how the pattern can be applicable to linguistic units at multiple levels (e.g. words vs. phrases vs. sentences), and the parameters that influence how the pattern can be instantiated differently (e.g. for varying levels of difficulty). Another consideration was to keep the number of patterns in our collection as low as possible, so that the pattern language does not become unwieldy. However, since each pattern pertains to a specific task, if we encountered a series of steps that resemble those in an existing pattern but targeted a different task, we created a new pattern instead of incorporating those steps into the existing pattern. Doing so enabled us to account for differences in learning goals between tasks. On a related note, even though we could have used a different set of learning tasks as our categories, we chose to adopt the same taxonomy from Clegg, Ogange and Rodseth (2003) because these task types are often encountered in the literature on teaching ESL. Most importantly, adhering to the same taxonomy made it possible for us to compare the tasks that we found against Clegg et. al's as a "consistency check." In total, we reviewed every software application over 2 to 3 rounds before we finalized the patterns that we distilled from it.

5.4 Summary of Patterns Found

In total, we identified more than 50 design patterns, including over 30 patterns for promoting reading skills. In this section, we summarize the patterns that we had found. For readability, we use **bold print** when referring to a pattern by its name. Similarly, although we use the imperative voice to provide the reader with recommendations on how these patterns can be employed, it is precisely our observations of the patterns being used in these specific ways in our sample of software applications – consistent and possibly even motivated by the literature on language acquisition and the psychology of reading – that prompted us to abstract the steps in these concrete instances into their more generalized pattern forms.

5.4.1 Spoken Language Acquisition

According to Skehan's (1998b) cognitive framework for language acquisition, the learner needs to acquire a rich set of formulaic expressions for fluency, as well as develop complexity and accuracy in terms of rule-based grammatical competence. As such, the **Initiation-Reply Sequence** introduces a sequence of statements in a dialogue which are usually formulaic expressions. This sequence, for example, may focus on a sequence of question-and-answer illocutionary acts, such that the learner may be given more than one candidate response to choose from. Similarly, in introducing the learner to new language, the software should explain the grammatical rules involved and how they can be used in communicative contexts. Learners can also reduce formulaic expressions analytically into smaller chunks over time, such that they develop rule-based grammatical knowledge in the process.

However, given the focus on meaning over form in task-based language teaching, the **Oral Word → Semantics Association** and **Semantics → Oral Word Association** patterns can be employed to help the learner associate spoken words and their meanings respectively with each other. The **Graduated Interval Recall** pattern helps to promote long-term retention of the words, phrases, etc. that are targeted, by introducing these linguistic items to the learner repeatedly in increasing (i.e. graduated) time intervals, as informed by the literature (Nation 2001) on how long-term memory retention works.

Extensive language input is a necessary condition for language acquisition to take place (Krashen 1987). The learner can be supplied with language input via oral narration of a text, using the **Text Story (Listen Only)** pattern. In case this narration is difficult for the learner to comprehend, the learner can be provided with extralinguistic context in the form of pictures and videos that accompany the narration, i.e. the **Picture Story (Listen Only)** and **Video Story (Listen Only)** patterns respectively. After which, the learner can be tested on his understanding of the narration by being presented a set of oral questions to answer, such as in the **Story + Oral Multiple-Choice Question** pattern.

Most importantly, language input is a necessary but not a sufficient condition for language acquisition. The work of second language acquisition researchers such as Swain (1985) has shown that production is an essential element in language development. At the same time, the **Initiation-Reply Sequence** pattern is sometimes criticized for rote-based teaching of language in a decontextualized approach. The **Anticipation** pattern addresses this shortcoming by first presenting the learner with a communicative context, before helping him to associate the context with the illocutionary sequences that are appropriate for the context. The **Pronunciation Exercise** facilitates communication by providing the learner with adequate drills to improve his pronunciation of difficult words and phrases. Finally, to help the learner prepare to use the target language in communicative contexts, the **Storytelling (From Pictures)** pattern provides him with controlled communicative practice by having him tell a story from a sequence of pictures.

5.4.2 Letter-Sound Knowledge

In languages such as English that have (phonetic) alphabetic writing systems, the basis for literacy stems from the alphabetic principle, in that skilled readers have a good command of grapheme-phoneme mappings, i.e. letter-sound knowledge. The **Animated Grapheme** design pattern introduces the learner to a grapheme (e.g. letters, numerals and punctuation marks) by playing a video clip that shows how to write it. This pattern can also recite the grapheme aloud via an audio clip in order to teach the learner how to read it aloud.

But **Animated Grapheme** presents letter-sound correspondences to the learner without testing him to check if he actually knows how to read the graphemes aloud. The **Grapheme → Phoneme Association** design pattern improves on **Animated Grapheme** by displaying a letter, after which the computer reads aloud a set of phonemes for the learner to choose from. Similarly, **Phoneme → Grapheme Association** pattern aims to help the learner form similar associations in the converse direction, i.e. when given a phoneme, the learner is tested to determine if he knows the corresponding letter. More important, this pattern is inspired by the sublexical route in dual-route models of spelling (Romani, Olson and Di Betta 2005), and thus goes beyond reading skills to target spelling skills as well. In any case, the **Grapheme → Phoneme Association** and **Phoneme → Grapheme Association** patterns establish the basis for the learner to grasp the alphabetic principle.

Lastly, writing systems such as the English alphabet contain graphemes that are "equivalent." For instance, the learner must be able to associate capital letters with their lower-case equivalents. The **Grapheme → Grapheme Association** pattern targets this learning need by showing (and testing) the learner how to match such letters.

5.4.3 Foundation Literacy

In the foundation literacy phase, the learner develops the ability to recognize very simple words by sight. He also learns to read monosyllables (i.e. graphemes involving 2 or 3 letters), bisyllables (i.e. graphemes with 3 or 4 letters) and simple non-words through grapheme-phoneme decoding. Despite architectural differences between various dual route models of word identification (Frost 2005), they share the common thesis that word meanings are accessed from the learner's lexicon in two ways: either directly from print (i.e. the lexical – also known as the visual – route) or by computing a letter string's phonological representation and using this information to access meaning (i.e. the sub-lexical or phonological route). These theories, however, generally assume that the visual route is more efficient and hence dominates skilled reading processes (Frost 2005), which accentuates the importance of the sight vocabulary. This point is especially salient for languages such as English whose writing system is characterized as a deep orthography (vs. a shallow orthography where there are one-to-one grapheme-phoneme relationships), in which a fair number of irregular words appear frequently in print and can be learned by sight.

Given the perceived importance of the sight vocabulary, the **Printed Phrasebook** and **Electronic Phrasebook** patterns aim to develop the learner's sight vocabulary by displaying written words alongside their definitions, such as by using pictures or showing definitions in the learner's native language (if he is literate). The **Printed Phrasebook** is advantageous in that it entails making printouts of vocabulary lists from the software's vocabulary bank, which imposes a lower requirement for computing infrastructure. On the other hand, the **Electronic Phrasebook** requires that the user has a computer to look up word definitions. Following which, the **Written Word → Oral Word Association** pattern can be used to show (and quiz) the learner how the written forms of these words – which can be irregular in terms of their orthographic mappings – can be read aloud, such that the learner is taught to associate the written form of a word with its spoken form.

Both the **Printed Phrasebook** and **Electronic Phrase** patterns could be criticized as being heavily biased toward the rote learning of vocabulary. The **Written Word → Semantics Association** pattern is a complementary pattern which we believe is a step closer to promoting active learning. In this pattern, the computer shows a written word to the learner alongside its meaning (e.g. pictorially). This step is repeated a few more times to introduce the learner to a few more written words, after which the learner is quizzed on the meanings for one or more of these words. The reverse pattern **Semantics → Written Word Association** aims to develop sight vocabulary in the converse direction.

Despite the importance of sight reading, however, sublexical processing is argued to be the default strategy employed by skilled readers in transparent orthographies and by novice readers (including second language learners) in a deep orthography like English (Frost 2005). As such, we cannot ignore the phonological route entirely in favor of the lexical route. The **Grapheme → Phoneme Association** and **Phoneme → Grapheme Association** patterns, which we describe above, can be used to help the learner acquire simple decoding skills for graphemes that are monosyllables, bisyllables and/or simple non-words (as opposed to how we have previously shown that the same patterns could be used for single letters).

5.4.4 Orthographic Literacy

In the orthographic literacy phase, the learner builds on his existing structures for letter-sound correspondences and word identification, so as to develop his understanding of how spelling is related to syllabic- and subsyllabic-level (e.g. onsets, peaks and codas) linguistic units. The words targeted in this phase are mostly at the primary school level, whereas non-words would include orthographic complexities, e.g. multi-letter graphemes, inconsistencies and irregularities.

The **Printed Phrasebook**, **Electronic Phrasebook**, **Written Word → Semantics Association** and **Semantics → Written Word Association** patterns that we cover above can be used to help extend the learner's sight vocabulary. **Text-Text Matching** is a new pattern which is another means of testing the learner's sight vocabulary. For example, the learner could be asked to match written words with their definitions, words with their synonyms, words with their antonyms, etc.

More fundamentally, phonological awareness is important to word identification (Ball and Blachman 1988; Bradley and Bryant 1978; Bradley and Bryant 1983, Fox and Routh 1976; Lundberg, Olofsson and Wall 1980; Tunmer, Herriman and Nesdale 1988; Wagner and Torgesen 1987). In this respect, the **Phonics Rule Matching** pattern can be used to introduce the learner to a phonetic rule that works at the level of the syllable, subsyllable or phoneme, as well as test his understanding of the rule after that (e.g. his ability to apply this rule in a new context). The **Subtitles** pattern shows the learner how written words can be segmented into syllables and the grapheme-phoneme mapping for each syllable. Likewise, the **Segmentation** pattern tests his understanding on how to segment written words at their syllable boundaries. Once the learner has attained a higher level of decoding skills, the above **Written Word → Oral Word Association** pattern can be used to show (and quiz) the learner how (regular) whole words can be read aloud. More important, through the **Oral Word → Semantics Association** pattern, the **Written Word → Oral Word Association** pattern can be used to help the learner enhance his lexicon of spoken words by assisting him to read the words aloud.

We conclude with a discussion on patterns for spelling, since the learner's grasp of the alphabetic principle is sometimes demonstrated more in spelling than reading. At the most basic level, the **Oral Word → Written Word Association** design pattern aims to familiarize the learner with the spellings of a certain set of words, such that a word is read aloud, and the learner is given a set of written words to choose from. As such, he only needs to recognize the written form of the word that was read aloud, but does not need to spell it. As a recognition task, the difficulty level can be raised with the **Similar Sounding Words** pattern, such that the learner is given several homophones (i.e. words that sound close enough) to choose from. In this way, the learner is tested on his ability to decode words via the lexical (vs. the sublexical) route in dual-route models of reading (Coltheart 2005).

The harder goal of scaffolding (and testing) the learner to spell correctly is left to the **Dictation Exercise** pattern, which involves the computer dictating words aloud to the learner and providing the learner with appropriate feedback. The difficulty level of the spelling task can be reduced by supplying some of the letters in the word and requiring the learner to spell the remaining letters. Similar to the **Similar Sounding Words** pattern, the **Find the Misspelled Word** pattern involves giving the learner a set of words to choose from, such that the given words have visual spellings that are very close to one another but only one set of spelling corresponds to the correct answer. This pattern is compatible with observations that learners have more difficulty spelling irregular words (vs. regular words and non-words), and as informed by dual-route models of spelling, can be used to test if learners have impoverished lexical representations (Romani, Olson and Di Betta 2005).

5.4.5 Morphographic Literacy

Morphographic literacy refers to the learner developing his representations of complex multisyllabic words and morphemes. It is significant because knowledge about word forms (i.e. morphological awareness) is crucial for reading comprehension. Despite differences at the superficial level, the **Text-Text Matching**, **Fill in the Blanks Exercise** and **Complete the Word** patterns can all be used to develop (and quiz) the learner's knowledge of morphemes. For example, the learner can be given a root word and options in the form of morphemes for completing the root word (e.g. "ed" or "ing,"), such that the completed word meets the required grammatical form (e.g. past tense, present continuous tense, etc.). The **Fill in the Blanks Exercise** is potentially more challenging than the other two patterns in this respect since the learner might need to infer the required grammatical form from the text around the blank(s), as opposed to being told the required grammatical form outright. Finally, the **Segmentation** pattern can be used to test the learner's understanding on how to segment written words at the boundaries between the root word and morpheme.

5.4.6 Reading Comprehension

In Kintsch's influential model of reading comprehension (Kintsch and Rawson 2005), reading comprehension refers to the learner's ability to reconstruct the situational model that is expressed by decontextualized statements of text. The **Text Story (Read**

Only) pattern is the most basic pattern for text comprehension, in which the learner is given a passage of text to read without any aid from the computer. This pattern can be used in more sophisticated ways, however. For example, it can be used to expose the learner to different genres of text, such that he enhances his sensitivity to story structure, which is a factor behind reading comprehension (Perfetti, Landi and Oakhill 2005). The **Gloss** pattern is frequently used in conjunction with the **Text Story (Read Only)** pattern to help the learner build his vocabulary when he encounters an unknown word. In these situations, the learner has the option of selecting the word to look up its meaning as used in the given context. The **Story + Open Question** and **Story + Written Multiple-Choice Question** patterns build on the **Text Story (Read Only)** pattern by posing questions to the learner to test him on his understanding of the text. The **Text Story (Read Only)** pattern may be too difficult for the novice reader, however. In this case, the text can be accompanied by audio-video, video (but without audio) or pictures via the **Video Story (Read + Listen)**, **Video Story (Read Only)** or **Picture Story (Read Only)** patterns respectively to provide the learner with extratextual context to help him construct the situational model.

Syntactic processing is a potent facilitator of reading comprehension (Coltheart 2005; Perfetti, Landi and Oakhill 2005). Hence, while some of the above patterns address morphemic awareness (i.e. grammar at the sublexical level), it is also crucial to develop the learner's awareness of grammar and ability to process text at the lexical and super-lexical levels. One pattern that aims to do the latter is **Word Classification**: for example, the computer presents the learner with a set of written words, such that it provides him with feedback as he classifies these words according to various grammatical categories (e.g. proper noun, adjective, verb, interjection, etc.). Another pattern that targets grammar is **Word Order Exercise**, which has the learner arranging a set of words and punctuation marks into syntactically well-formed sentences. Lastly, although the **Dictation Exercise** pattern can be used to test the learner on his knowledge of word spelling, this pattern can also be used to test him on his understanding of letter capitalizations and punctuations.

We did not encounter patterns that target other skills for reading comprehension, such as textbase formation (Kintsch and Rawson 2005), inference making (Perfetti, Landi and Oakhill 2005) and comprehension monitoring (ibid). One exception is the above **Fill in the Blanks Exercise** pattern, which can be used to promote *basic* comprehension monitoring in addition to developing awareness about grammar. For instance, the learner can be asked to choose from a few options for a blank, such that even though any option results in a syntactically well-formed sentence when used to fill in the blank, only one option constitutes a semantically meaningful sentence or a clause which is semantically consistent with the rest of the text. To conclude, the patterns identified for reading skills from our sample appear to be much more relevant to word identification than reading comprehension.

6 PACE Framework

In co-designing e-learning applications with community partners to address their local language learning needs, we need to be mindful not to incur content development costs that exceed the limited budgets of a typical community development project. When we commenced the MILLEE project, we were initially under the impression that the unit of design is the activity and its user-interface. We soon realized that the requirements of a language learning application necessitated the curriculum as a separate unit of design. In other words, content development is the real and more difficult challenge. The priority is to integrate the user-interface and curriculum design processes in a more comprehensive content development framework. Most important, given the cost of content development, such a framework needs to promote the reuse of existing knowledge and material as far as possible.

The task-based language teaching framework (Chapter 4) and pedagogical design patterns (Chapter 5) serve as a starting point for such a content development framework. In this chapter, we present the PACE (Pattern-Activity-Curriculum-Exercise) framework that integrates best practices in language instruction with considerations about curriculum design. The essence of the PACE framework is that an e-learning software is viewed as being structured into multiple layers, such that there are distinct layers that pertain to best practices for learning and engagement (i.e. patterns), interaction design (i.e. activity), the syllabus to be learned (i.e. curriculum) and the basic unit that realize an activity with its curriculum (i.e. exercise). Modular design along the lines of this structure allows the units of an educational software application to be reused as much as possible, such that only those units that need to be adapted for local needs are customized. The emphasis on modular design and reuse is inspired by "learning objects" (Wiley 2002), which are digital resources that can be reused and recombined in other educational contexts so as to be repurposed in supporting new learning needs. Furthermore, this structure facilitates multidisciplinary collaboration such that team members only need to focus on those modules that are aligned with their areas of expertise and responsibilities.

The rest of this chapter proceeds as follow. First, we explain how the complexity of the adoption ecology calls for a high degree of micro-localization. Second, we describe the components in the PACE framework that correspond to the various modules in an educational software application that are amenable to localization. Third, we explain how structuring an e-learning software into these components promote reuse. We conclude by summarizing our formative experiences with the pedagogical design patterns – in which we assembled a large team to employ the patterns to design English learning games for low-income children in India. This experience led us to have a better appreciation for the other components that are needed to complement the design patterns. It served to inform us about the rest of the components in the PACE framework that we have to put together in order to arrive at a complete language learning game.

6.1 Micro-localization

Figures 6.1 and 6.2 depict the policy making and implementation structures in India's formal education system respectively. The reader will observe that both diagrams highlight the complexity behind the adoption ecology in terms of the numerous entities in various layers in the hierarchy, and how they interact with one another to discharge their responsibilities.

National: Ministry of Human Resource Development (Minister)

Department of Education

Department of School Education and literacy

Department of Higher Education

Elementary Education

Secondary Education

Adult Education & Literacy

{Headed by Secretary, Joint Secretary and Add. Secretary (IAS Officer)}

Sarva Shiksha Abhiyan (SSA)

State: Educational Secretary (IAS, Indian Administrative Service Officer)
SSA (State SSA Body headed by IAS called State Project Director)

District: District Education Officer (State Civil Servant)
SSA (District Project Coordinator)

Block: Block Education Officer (BEO)

Figure 6.1: The entities in various levels of India's government that are responsible for education policy making.

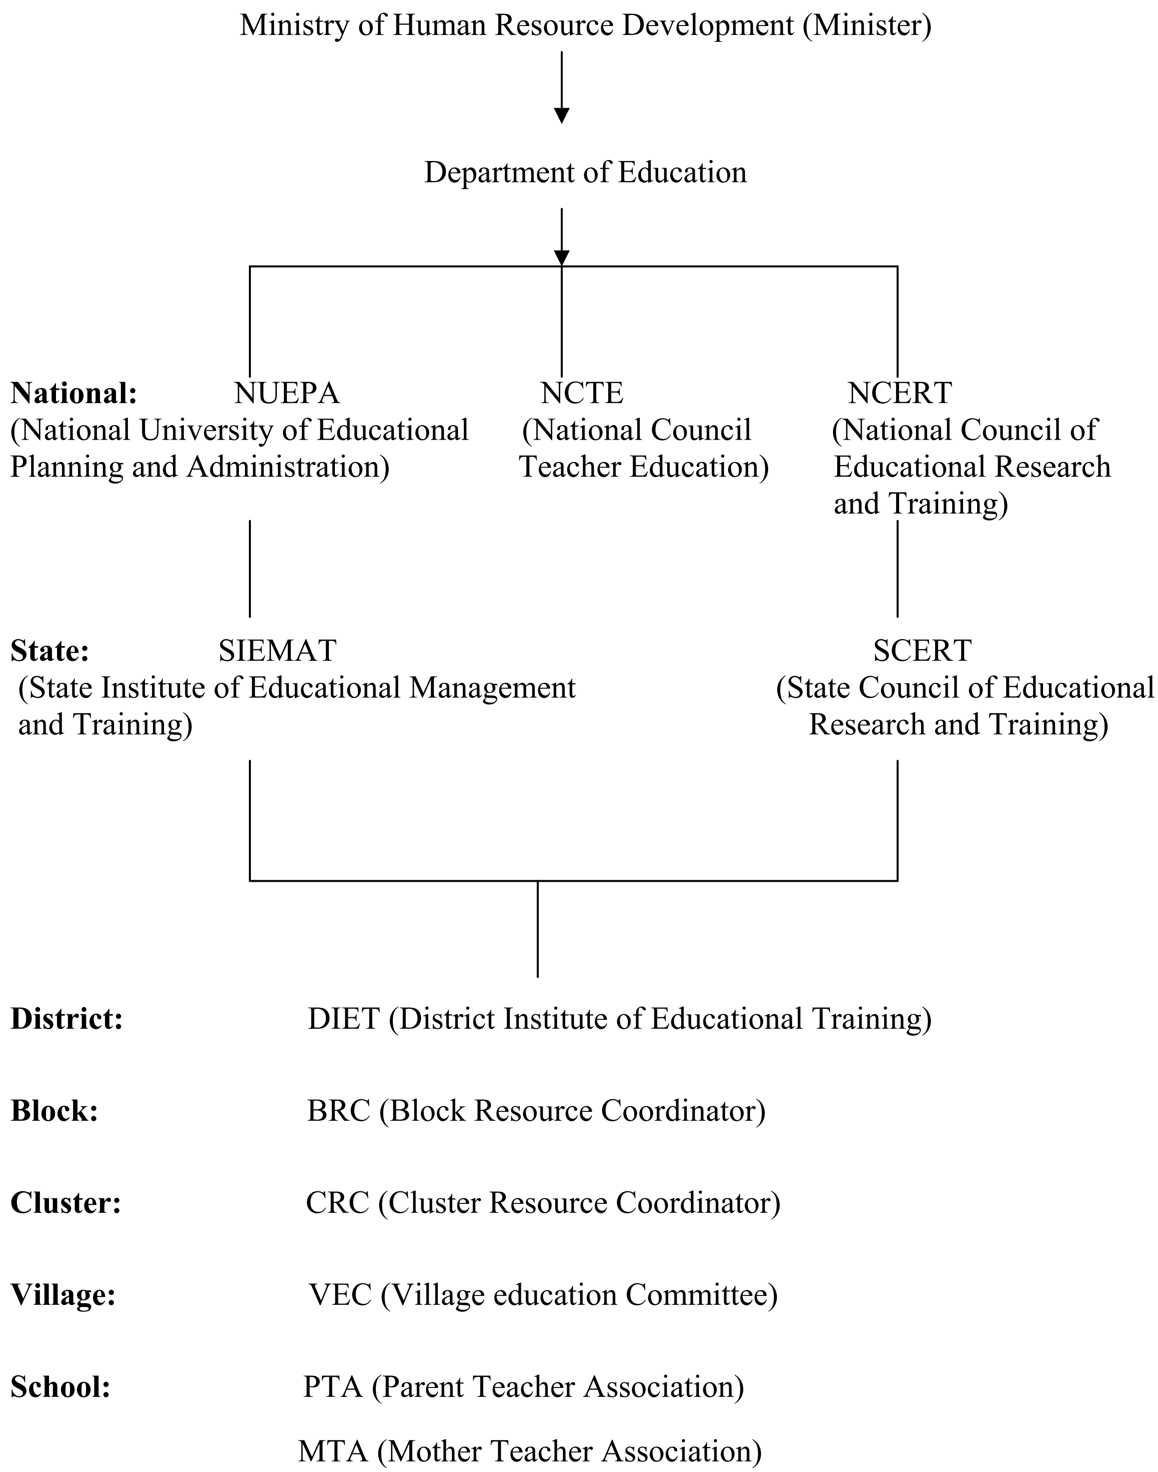

Figure 6.2: The entities in various levels of India's government that are responsible for education policy implementation.

To add to the complexity in Figures 6.1 and 6.2, we note that the exact details in the hierarchical structure can vary significantly across states. As such, although English is taught in schools, when seen from the top-down direction, every state implements its own curriculum for English teaching. To compound the heterogeneity, each state not only has its own education board, but also enacts distinct educational policies. For instance, the grade level in which English is introduced into the curriculum differs across states, which means that English material of comparable difficulty is taught to students in different grade levels and ages depending on the state.

From the bottom-up direction, the state curriculum for English teaching may not be adhered to for various practical reasons. For example, in Chapter 1, we have narrated the challenges associated with teacher absenteeism and training. Similarly, in Chapter 10, we recount how one teacher felt unmotivated to teach well. Hence, students in a location may receive higher quality lessons than children elsewhere due to differences in teacher effort and English proficiency. Consequently, non-government organizations working on the ground to improve English learning among rural children are sometimes forced to customize the state curriculum, which is too difficult for rural children given their level of preparation. In our experience, NGO volunteers also adapt the state curriculum when they feel that it is too urbanized for rural children to relate to.

The localization of the English curriculum therefore takes place at multiple levels in the adoption hierarchy. Given the heterogeneity in the learners' English baseline, it is not scalable to develop "one-size-fits-all" e-learning games when they need to be adapted for specific local needs and conditions. We call this requirement *micro-localization*. This notion of micro-localization varies from existing practices of localization in the computer science community, whereby the latter meant adapting a user-interface for a different culture (Marcus and Gould 2000; Smith et al. 2004; Yeo 2001), or more conventionally, localizing a user-interface for speakers of a different language. In the unique situation of language learning, we not only need to customize the interface, but also have to adapt the curriculum, until both components are appropriate for the specific community. This level of local adaptation takes place at a micro level that is much more granular than the user-interface or its language. We next describe the components in the PACE framework and how they facilitate micro-localization along the multiple levels of pedagogy, curriculum and interaction design.

6.2 Components of the PACE Framework

In this section, we describe the Pattern-Activity-Curriculum-Exercise components that make up the PACE framework, from the most abstract (i.e. the Pattern) to the most concrete (i.e. the Exercise). We have also learned from our experience that this sequence reflects, to a large extent, the same steps in the process that a design team takes when instantiating abstract design patterns into concrete learning exercises.

6.2.1 Pattern

A design pattern (Alexander 1977) is a "template" description of a solution to a problem that is previously encountered and solved. In the language learning domain, a pattern allows us to represent, in a skeletal form, the steps that current language learning software take to implement a learning task, which the learner engages with to develop her language skills. Patterns also capture contextual information (e.g. domain applicability) and some rationale. Patterns hence scaffold designers who lack specialized backgrounds in language teaching to implement language learning tasks in the form of software. They may also capture tacit knowledge about the domain or local context after having evolved through several iterative design cycles. In this way, there is no need to reinvent the wheel. The reader may refer to Chapter 5 for further details on the pedagogical design patterns that make up the PACE framework.

6.2.2 Activity

Design patterns are necessary but insufficient, because our objective is not only to promote the reuse of patterns but also to reuse those learning activities that instantiate these patterns. By learning activity, we mean an interaction design that constitutes a user-interface and experience for the language learner. This extent of reuse is critical for developing regions because activities and their user-interfaces may require several rounds of iterations before they are finally appropriate for the users' cultural backgrounds and low levels of familiarity with computing usage.

Figure 6.3. An activity and its user-interface that is an instantiation of the "Written Word → Semantics Association" design pattern from Figure 5.5. Starting from the top left, in clockwise order: (a) the player is taught the meaning of "stop" with a corresponding picture, (b) tested if she remembers what "stop" means, (c) selects the wrong picture for "stop", and (d) finally selects the correct picture.

Figure 6.3 shows an example of a learning activity that implements the "Written Word → Semantics Association" pedagogical design pattern given in Figure 5.5 in the form of a game. This game begins with a receptive phase where the player is presented with words and pictures that represent the meanings of those words (top left in Figure 6.3). In the activation phase, a word appears while pictures scroll by (top right). The player presses the "Enter" button when the current picture matches the word, and receives feedback for being correct (bottom left) or wrong (bottom right).

An advantage of separating patterns from activities is that content developers with expertise in language teaching can focus on the design patterns, while developers with programming backgrounds focus on implementing the learning activities.

6.2.3 Curriculum and Exercise

Learning activities correspond to particular curricula. For example, the curriculum for the "Written Word → Semantics Association" learning activity shown in Figure 6.3 comprises a list of words, images that represent their meanings graphically and sounds for their pronunciations. That is, a curriculum comprises both the language syllabus that the learner should be taught, and the digitized form of this syllabus. Finally, an exercise is a realization of an activity that is associated with a curriculum. In other words, whereas a learning activity is the interaction design and its user-interface that is independent of the actual curriculum, an exercise consists of a learning activity matched with a digitized curriculum. It is hence the exercise – and not the learning activity – that is the software which the learner eventually interacts with and uses. Chapter 8 will provide examples of how we design and iterate on educational applications that are structured according to the PACE framework. These examples will make the components in this framework, and the distinction between them, clearer.

Separating the learning activity from curriculum promotes reuse and scalability in at least four ways. First, curricula such as high-frequency word lists in commercial-grade products (e.g. Ladybird's *Key Word Reading Scheme* and Voice of America's Special English radio broadcasts) are relevant to most learners of the language – independent of their cultural backgrounds – and hence do not need to be reinvented. Such a curriculum can be easily reused for a new learner population by matching it with a different activity whose user-interface is more culturally appropriate for the given learner population.

Second, one of the common problems encountered in some developing regions is that learners in the same grade level or language baseline may belong to a relatively wide age bracket due to inequity in their opportunities to gain access to high quality education. For example, a language curriculum could be relevant to both male teenagers and middle-aged mothers who had to discontinue formal schooling during their teenage years due to marriage. As such, a curriculum for the learners belonging to the same grades or baseline can be relevant to the same learners despite their age differences. This curriculum can be reused by associating it with multiple learning activities, such that some activities and their user-interfaces appeal to older learners while others aim to engage younger learners, for instance.

Third, learnability and usability are serious issues in underdeveloped regions due to limited exposure to computing technologies. Hence, as soon as an activity and its user-interface are intuitive to the target learners, presumably as a result of extensive user-testing, it makes sense to reuse the same learning activity as much as possible. Such an activity can be associated with multiple curricula (e.g. each curriculum is the vocabulary list for a different topic) so that the same user-interface is reused for a new curriculum.

Fourth, a localized curriculum like a vocabulary list that is culturally meaningful for a given learner community and hence meets its learning needs can be reused by associating it with multiple activities in order to address different aspects of their learning needs. For instance, one activity can attend to spelling skills, one activity can concentrate on listening comprehension skills, another activity can focus on pronunciation, etc. for the same curriculum. In this way, a comprehensive package that targets multiple language competencies for the same syllabus can be created inexpensively from a core curriculum.

In summary, content development is a prohibitively expensive process that incurs recurring cost. The shortage of educational content for both digital and paper media that is relevant to learner communities in underdeveloped regions remains a huge problem. In view of the above advantages of scalability, reuse and learnability, the PACE framework and process promises to streamline the costs of repurposing existing language learning resources for new audiences.

6.3 Formative Design Experience

Armed with our observations from the exploratory studies described in Chapter 3, in the spring of 2006, we set out to design our first batch of English learning games that would run on cellphones. Since this was our first attempt at designing for a domain that was not well understood, we did not expect these early designs to be successful. Instead, we treated the process as a formative experience and prelude to future designs.

In the spring of 2006, our group comprised two graduate student researchers and fourteen undergraduates supervised by a faculty advisor. Our team members came from backgrounds that included HCI, development studies and education. In addition to the above exploratory field studies in North India, the author had conducted fieldwork in East Africa while the other graduate student researcher had done fieldwork in South India. We pooled our collective experiences and shared them with the rest of the team in these forms:

- ~80 pages of trip reports, which included interviews with teachers in rural India and other local stakeholders, as well as observations of how children from the slums and rural areas used existing software applications (general-purpose applications such as Microsoft Office as well as educational courseware) described in Chapter 3,
- ~60 pages of secondary literature on the present state of elementary education in rural India, i.e. the report commissioned by the Azim Premji Foundation (2004b),
- Personas of the above Indian children from the urban slums and rural areas,

- Scenarios of the children's everyday lives, and
- >400 MB of photographs from the above fieldwork to help team members better visualize ground conditions.

One of our team members in that semester had grown up in India and came to the United States two years earlier on a scholarship for his college education. He provided us with invaluable perspectives on the local culture. In particular, he helped to identify a list of functional literacies that were both relevant to the above children and entailed written or conversational English (e.g. traveling or visiting the pharmacy). This functional view of literacy, in which language is relevant to specific local contexts and social practices (Scribner and Cole 1981) – as opposed to viewing a language as uniform and universal for all of its speakers – is widely acknowledged to be critical for language acceptance in developing regions. We subsequently targeted the selected functional areas as learning objectives in our designs.

Our aim was to complete the design documents for a suite of ESL learning games by the end of spring 2006. On top of the above resources, we drew on best practices from two sources. First, we adopted best practices from the computer games industry, such as sample design documents, case studies and interviews with experienced game designers (Hallford and Hallford 2001). Second, we incorporated best practices from commercial language learning packages in the form of the pedagogical design patterns summarized in Chapter 5. By the end of the semester, we had designed over 30 e-learning games for English as a Second Language. These games included crossword puzzles, hangman, fill-in-the-blank games, multiple choice games, matching games and word scrambles.

However, the above processes and resources were clearly inadequate. The designs suffered from localization-related problems. For instance, it was difficult for many team members, who were unfamiliar with the functions of English in Indian society, to expand the given functional literacies into a suitable syllabus. Similarly, the ESL textbooks that we have from India were not helpful because they were not structured according to the shortlisted literacies. Finally, the game settings and user-interfaces designed by team members contained several Western biases.

More troubling, there were also issues with the educational quality of the games. Before the patterns were introduced to the team members to try out as design tools, many team members came up with designs that were interesting but lacked pedagogical value. After the patterns were provided as a design tool, members began to design games that were arguably more educational. However, the patterns also imposed constraints on their creativity, and we noticed that on the whole, the games were less exciting to play than before the patterns were introduced. Some of the later games resembled multiple-choice quizzes taken straight out of paper-and-pencil tests. A plausible explanation is that some members adhered rigidly to the formulaic structure in which the patterns were presented, without considering that patterns are only skeletal structures that serve as starting points for further creative input.

We acknowledge that our process of giving team members design patterns based on learning software biased our team members to favor pedagogy over other qualities, including fun and engagement. Furthermore, because most of the team members lacked specialized knowledge about language acquisition, they could not easily depart from the patterns without the risk of losing their pedagogical value. In future, we believe it is important to also prime members with design patterns for successful games, independent of pedagogy. We will take up the topic of game design patterns in the next Chapter.

On reflection, the above experience indicated that in addition to the patterns, we needed to take considerations about the learning activity (which included games and their settings) and curriculum into account in order to have a comprehensive design process. Such a process will maintain a sufficient balance in designing e-learning games in terms of interaction design, game design and curriculum. It was this experience that in turn inspired the components that comprise the PACE framework.

7 Game Design Patterns

A design pattern is a "template" description of a solution to a recurring problem that has been solved. In this way, there is no need to reinvent the wheel, and new designs can leverage on earlier designs that were successful. A pattern also captures contextual information such as domain applicability and its rationale. A pattern may also capture tacit knowledge on the domain after having evolved through iterative design cycles. In the case of the MILLEE project, a critical challenge is: how can we design mobile games that are engaging and fun for rural children in India to play? The challenge of designing enjoyable games becomes even more pronounced when we have not found any studies on how children in rural settings interact with mobile games, which we could use to inform our work.

In Chapter 5, we described a review of successful commercial language learning products and the pedagogical design patterns that we distilled from these products. In this Chapter, we employ the design pattern approach in parallel, but from the standpoint of gameplay design, in which we take advantage of game design patterns that are usually found in successful games as basic building blocks for composing new game designs (Kreimeier 2002). Our hypothesis is that games that are *consciously* designed using game design patterns as design tools are more fun and engaging to play, in comparison with games whose design process did not involve the patterns. This qualification is critical because patterns are arguably pervasive to the extent that novice designers with extensive gaming experience can *unthinkingly* apply patterns encountered in prior gameplay experience to their work. Our hypothesis assumes, however, that there are still benefits to taking patterns explicitly into account during the design process, e.g. in doing so, the designer can deliberate on the rationale and tradeoffs for each pattern.

This Chapter is structured as follows. First, we describe our fieldwork in India, where we conducted an experiment with 24 children at a rural school to evaluate the efficacy of game design patterns as design tools. Among the 8 mobile games that we deployed, we had designed and implemented 3 of them with the aid of patterns, while the remaining 5 games were obtained off-the-shelf and did not consciously factor the patterns into their designs. Second, we present the data collection procedures for eliciting rural children's opinions on how they had found each game to be fun and engaging. These procedures evolved out of pre-testing with the same children, and are invaluable for other game studies researchers wishing to work with similar rural demographics. Third, we report our qualitative observations on how participants played these 8 games, in order to draw lessons on designing future games for this unusual user group.

7.1 Hypothesis

Although our broad hypothesis is that games which were consciously designed using game design patterns are more fun and engaging to play, in comparison with games whose design process did not involve game design patterns, this statement needs to be framed in a more nuanced manner, as we soon learned in the process of selecting games for our comparison group. More specifically, the mobile games that we piloted came under three categories, i.e. those which we:

- consciously designed using the game design patterns,
- obtained off-the-shelf that were designed by amateur game developers, and
- obtained off-the-shelf that were designed by professional game developers,

such that we hypothesize that games which we (who are not specialists in game experience design) designed with the aid of game design patterns would be more engaging than those which were designed by amateurs, e.g. hobbyists who develop games without a profit motive. On the other hand, we expect that games designed using patterns would not be as enjoyable as those designed by professional designers. In other words, while patterns can scaffold non-specialists in designing enjoyable games, patterns do not substitute for specialized knowledge and experience in game design. The implication is that patterns are likely to be more useful to novice game developers. This does not mean, however, that patterns have limited usefulness in general since they can help novice designers transition to expert status.

7.2 Games

We believe that casual games are especially appropriate for developing regions since their low time overhead is a good fit with children's work commitment. The sample of casual games that we evaluated at a village school included three games (Crocodile Rescue, Dancer and Train Tracks) that we had designed with the help of game design patterns, after which we implemented in Flash Lite. The sample also comprised three amateur games (Floored, Beginner Land and Turtle Boat) and professional games (Toy Factory, Jump Bot and Critter Crossing). The off-the-shelf games developed by amateur and professional game developers were selected such that they were comparable to the games that we had designed, in terms of play complexity, age appropriateness, cognitive demand and the rural children's ability to relate to the games culturally. We also ensured that the games have comparable animation effects such that participants could not distinguish commercial games from non-commercial ones simply from their user-interfaces. The off-the-shelf games were developed for the phone platform in Adobe's Flash Lite or native code. Figure 7.1 shows screenshots of the games that we deployed. Due to lack of time in the field, we could not deploy Turtle Boat.

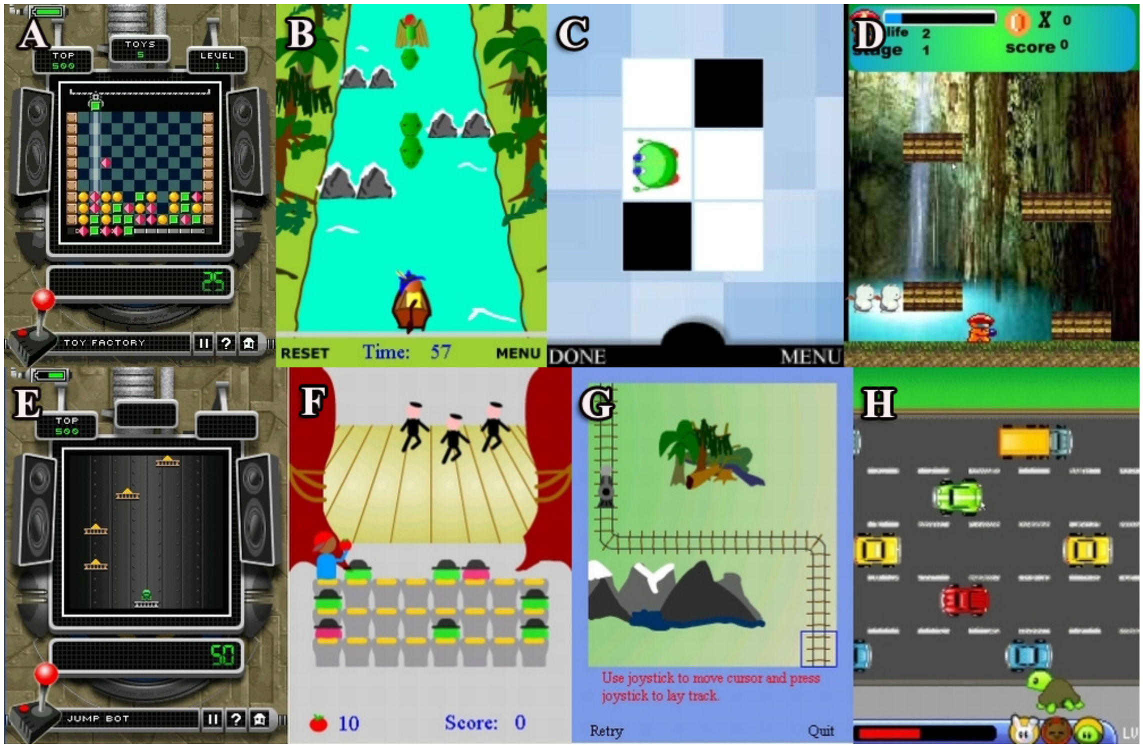

Figure 7.1: Screenshots of the 8 mobile games that were deployed in the study, in the order that they were introduced to the participants, namely, (A) Toy Factory, (B) Crocodile Rescue, (C) Floored, (D) Beginner Land, (E) Jump Bot, (F) Dancer, (G) Train Tracks and (H) Critter Crossing.

7.2.1 Games Implemented Based on Design Patterns

Since our focus is on casual games, we reviewed a series of casual games that appeared on bestseller charts in order to identify those patterns found in them. In total, we identified slightly more than 30 patterns, which we classified into 4 categories: core mechanics, story elements, goal states and reward mechanisms. Next, we randomly picked 1-3 patterns from each category and used them as “germs” to brainstorm our designs. We repeated this process to culminate in three game designs (Crocodile Rescue, Dancer and Train Tracks), with Table 10.1 indicating the actual patterns that we applied:

Crocodile Rescue – the player rescues the drowning boy by moving his boat around a two-dimensional map to the boy. The level is completed once the player moves onto the same location as the boy within the time limit. As a reward for completing the level, the player is challenged with more obstacles that take the form of crocodiles, which he must bait out of his path using chunks of meat, in order to clear a path to the boy.

Dancer – the player moves among the audience on a 2D map in order to throw tomatoes at dancers on the stage. The player seeks to maximize his scores within the time limit by hitting as many dancers as possible at least once, and completes the level once every dancer is hit. As such, the player assumes the coveted identity of a trouble-maker that is not necessarily possible in real life.

Train Tracks – the player meets his need for self-expression by laying railway tracks to extend the existing track in any direction he desires, so long as there is no obstacle. The goal is to get the moving train from the top left-hand corner of the map to the bottom right-hand corner before the train derails at the end of an unfinished railway track.

The game settings for Crocodile Rescue, Dancer and Train Tracks were chosen such that rural children could relate to them readily. As a consistency check with the literature, we also verified that the patterns we had used could also be found in Björk and Holopainen (2005) in some form.

Table 7.1
Patterns Applied to Crocodile Rescue, Dancer and Train Tracks

	Crocodile Rescue	Dancer	Train Tracks
Core Mechanics	2D movable object	2D moveable object, Projectile	Polyline construction
Story Elements	Rescue	Dodging, Dancing	Racing, Building
Goal States	Trigger condition	Trigger condition, Maximize	Trigger condition
Reward Mechanisms	Deliberate obstacle	Coveted player identity, Create time pressure	Self expression, Create time pressure

7.2.2 Comparison Games

Among the games that we piloted, those that were designed by amateur game designers were:

Floored – the goal is to flip the colors of each tile on the board until all the tiles share the same color. Each time the player moves from one tile to another, the color of the previous tile changes. The rules may become more difficult at higher levels, e.g. the player is not allowed to backtrack to the tile that he was last on.

Beginner Land – the player moves around the game world (which spans multiple screens) to collect enough coins before he is allowed to fight the boss. It is in the player's best interests to shoot the white ghosts who fly around since they will hurt him when they draw nearer.

The three games in our sample which were designed by professional game developers were:

Toy Factory – in this Tetris-like game, the player needs to drop toys on an existing heap of toys with the same color before the heap overflows. This game differs from Tetris in that there is a toy dropper at the top of the screen that oscillates from left to right, and that it is not possible to rotate the toys released by the toy dropper.

Jump Bot – the goal is to scale the highest possible height by jumping onto the next highest platform when it hovers to float above the player. The player will plummet to his death if he misses the platform when he jumps.

Critter Crossing – the goal is to help the various creatures at the bottom of the screen cross the road, river, etc. without getting hit by a vehicle within the time limit.

7.3 Experiment

The experiment took place over 10 days in January 2007 at a rural school near Mysore, India. The experiment was assisted by 4 bilingual local adults whom we hired as interpreters.

7.3.1 Participants

We obtained consent from parents, the village head, school teachers and government officials for all grade 1-2 students (i.e. ages 6-7, on average) to be excused from classes in the mornings, so that the students could participate in our study and resume classes after lunch. In total, there were 9 1st-graders (5 male and 4 female) and 15 2nd-graders (6 male and 9 female) enrolled at this school. About 6 participants did not appear to be regular school-goers and showed up only after hearing about our study involving mobile games. Even though none of them reported having used cellphones previously, all of them understood what cellphones are. None had prior experience with electronic games either. They could write simple words in their native language Kannada. English was not taught in school until 5th grade, but a few students have learned the English alphabet from their parents or siblings. None of the participants spoke English beyond simple greetings.

7.3.2 Instruments

We adapted our data collection instruments from Read's Fun Toolkit (2002), which are techniques for standardized questionnaire interviews with children aged 5-10 to elicit their views on the extent to which their play experiences with particular games were fun and engaging. Although there is other work in the game studies literature to develop methodologies for "measuring fun," the Fun Toolkit is the only methodology that we know of which is specifically for and has evolved through iterative testing with children.

Figure 7.2: The Smileyometer, reproduced from Read, MacFarlane and Casey (2002).

Two instruments in the Toolkit that are especially relevant are the Smileyometer (Figure 7.2) and the Again-Again table (Figure 7.3). The rationale behind the former is that asking a child to rate the level of fun that he had with a game on a 5-point scale may be more appropriate for adult respondents. The Smileyometer was designed with some children, such that child respondents are shown five faces corresponding to a 5-point Likert scale and asked to pick a face that best matches their experience with a game, with the help of this visual tool. The Again-Again table is used to determine the extent to which a child finds a game engaging such that he would like to play it again.

However, since Read's instruments evolved from iterative testing with children in urban developed country settings, we are concerned about the ecological validity of our results if we do not adapt these instruments for a rural setting. In particular, we know from earlier cross-cultural psychology studies (Cole 1996) that children in other cultures may fail to provide valid responses to questions that are not culturally meaningful or which they had never been exposed to.

7.3.3 Pre-Test

The above concerns prompted us to pretest our instruments with our participants prior to introducing them to the 8 mobile games illustrated in Figure 7.1. The pre-test occupied the initial 3 days of the 10-day field study. Since our participants had no prior experience with video games, we introduced 4 mobile games (i.e. Tangram, Litterbug, Captain Gravity and Over the Fire; see Figure 7.4 for screenshots) during the pre-test before conducting the questionnaire interviews based on the same prefatory games, which were chosen so that we could introduce the children progressively to basic controls found in most games, e.g. the joystick button, the shooting action, etc.

Although we could have pre-tested with a different group of rural children, we chose to refine our instruments with our participants so as to familiarize them with the questions and train them to give meaningful responses. We worked with our interpreters to "tease out" what the children actually meant, as opposed to taking their answers literally. In the process, the interpreters rephrased our questions repeatedly until they conveyed our intentions without cultural biases or translation errors.

Would you like to do it Again?

	Yes	Maybe	No
Visit U Boat	✓		
Puppet show		✓	

Figure 7.3: Part of an Again-Again table, reproduced from Read, MacFarlane and Casey (2002).

We found that it was necessary to scrap the Smileyometer because the participants kept selecting smileys for "really good" and "brilliant". We learned that they never chose "awful" or "not very good" because they found smileys to be aesthetically more appealing than frowns. We eventually asked them to give a score from 1 (lowest) to 5 (highest) to rate their enjoyment with each game. We learned that they initially rated the games on a reversed scale because it was a customary school practice to rank students in class, such that "1" denotes the top student, etc. However, we retained our original scale when participants understood it after our explanations. We also streamlined our interview process by asking the "Would you like to play this game again?" question from the Again-Again table immediately after the above rating question, and removed the "Maybe" option after finding that participants struggled with this ambiguity.

Figure 7.4: Screenshots of the 4 games that were used to pre-test the data collection procedures, namely, (A) Tangram, (B) Litterbug, (C) Captain Gravity and (D) Over the Fire.

To reduce recall failure and interference from exposure to newer games, our preference was to ask the above questions for each game immediately after participants had finished it. This would also make it easier to seek clarifications on the reasons that participants had for assigning the ratings that they gave. However, due to their limited experience with electronic games, we found that they understood each game better – and how to play it with greater successes – after they had learned more games. For instance, participants did not appear to understand Over the Fire and disliked it until they had gained exposure to more games, whereupon they developed more positive impressions about the game. As such, the compromise that we adopted was to ask the above questions for each game after participants had been exposed to 1-2 subsequent games.

7.3.4 Procedure

On each day, we introduced an average of two new games and conducted a revision of 1-2 earlier games. The 8 games were introduced according to a sequence that interleaved those designed by us with those games from amateurs and professionals. In the ideal within-subjects experiment, we would have introduced games in a different sequence for each participant to minimize ordering effects. But we lacked manpower for this level of classroom management.

Figure 7.5: A boy looks on as he waits for his turn.

Three researchers conducted the interviews using questions refined during the above pre-test phase, as well as observed how participants played the 8 selected games. To minimize observer bias, we rotated ourselves among the respondents. On the last day, after most participants had become familiar with the games, we selected 3 of them and asked to video-tape them as they played every game. They were selected based on the understanding of the games that they exhibited in the above interviews. They were also selected to ensure a variation in age (one 1st-grader and two 2nd-graders) and sex (2 boys and 1 girl) among the participants whom we videotaped. But their prior performance in the games was not a criterion for selection because we wanted to observe any usability problems that they may encounter.

Given our concerns that the above self-reported data may not be reliable, we performed one more round of interviews on the last day of the study. More specifically, we showed printouts of screenshots from each game and asked every child to pick the top two games that he or she would like to play again. This data provided us with an additional source of data to triangulate our analysis with the above numerical ratings and video observations.

The cellphones used in the experiment were the i-mate SP5 Windows Mobile 5.0 smartphones. Since we only had 9 phones, students took turns to play (Figure 7.5).

7.4 Results and Analysis

How useful are game design patterns as a design tool? Table 2 shows the average rating assigned by participants to every game on a 5-point scale. We can make three striking observations. First, none of the "top three" games (Floored, followed by Critter Crossing and Jump Bot) were designed using patterns. Second, and more interesting, one of the games in the "top three" (Floored) was in fact designed by amateurs. Third, among the three least popular games, one was designed using patterns (Crocodile Rescue), one was designed by amateurs (Beginner Land) and another was designed by professionals (Toy Factory). Clearly, designing engaging mobile games for rural Indian children is more complex than labels such as "amateur" and "professional," or even the direct application of tools such as patterns.

Table 7.2
Average Ratings Assigned by Participants to Games (From Most Popular to Least Popular)

Game Title	Mean Rating (Out of 5)
Floored	4.5
Critter Crossing	4.5
Jump Bot	4.1
Dancer	3.9
Train Tracks	3.9
Beginner Land	3.7
Crocodile Rescue	3.4
Toy Factory	3.3

7.4.1 Three Most Popular Games

Among the three most popular games in the case of Floored, during the interviews, 5 participants specifically volunteered that they liked it because it was an easy game. We observed that participants succeeded at initial levels through a divide-and-conquer strategy, where they made “local” moves in specific regions of the game board until the tiles in the region had switched to a uniform color before proceeding to do likewise with other regions. That is, it was possible to work on different parts of the board as independent regions without incurring significant cognitive overhead for strategic reasoning or thinking ahead. 13 participants added that they enjoyed the game because they liked the colors of the tiles.

However, while aesthetics play a role, we also noted that many other games in our sample were equally – if not more – colorful, yet few participants liked them because of their colors alone. Based on feedback from respondents that they enjoyed changing the tile colors, we believe that it was this ability to change the aesthetics of the game elements in addition to the colors themselves that contributed to the excellent ratings. We also note that colors and household drawings have a dominant role in Indian cultural festivals.

For Critter Crossing, 4 participants enjoyed it because they found it easy to help the critters cross the road, although 2 participants disliked the game because they found the road crossing goal to be difficult to achieve, presumably due to heavy road traffic. 4 more children cited liking the critters and their appearances as their reason for liking the game.

Jump Bot received positive ratings because 9 participants enjoyed performing the jump action. On the other hand, Jump Bot received negatively ratings from 4 participants who found it difficult. In particular, the game was not well designed for player error in that missing the platform during a jump results in the player plummeting to destruction and having to restart. It seemed that Jump Bot would have been more popular if the consequences of a failed leap had been less dire.

7.4.2 Three Least Popular Games

Turning our attention to the three games with the lowest average ratings, 5 participants indicated that they did not like Beginner Land for its difficulty because it was hard to shoot the ghosts before they draw closer to the player. A possible reason was that participants had limited familiarity with arcade games that required hand-eye coordination. It was also possible that the joystick button on the cellphone was difficult to use, due to hardware usability problems. We also believed the software user-interface could be improved in terms of usability and aesthetics. In terms of its positive aspects, 5 participants enjoyed the shooting action.

As to Toy Factory's unpopularity, 3 participants indicated that it was difficult to play, plausibly because it is an arcade game like Beginner Land. It was also possible that Toy Factory received the lowest average rating among all the 8 games because it was the first to be introduced, when the participants were still learning how to play mobile games.

In the case of Crocodile Rescue, the biggest problem was that 4 children did not like the crocodile sprites in the game. In particular, two girls found the crocodiles frightening, especially when they open their jaws in our animations. We understand that crocodiles are often villains in Indian mythology, and soon learned that crocodiles are perceived as dangerous by villager dwellers. In contrast, 3 participants enjoyed the Crocodile Rescue game because they fared well in it while another 3 identified with the hero whose goal was to rescue the drowning boy from the crocodiles.

7.4.3 Other Results

Among the remaining two games designed using patterns, which received moderate mean ratings, Dancer was popular because 4 children enjoying the sight of the dancers moving about on stage while 3 players loved to throw tomatoes at the dancers.

Players appeared to identify with the coveted identity of the mischief maker whose goal in the game was to ruin the dancers' stage performance. The game was not as well-liked on the whole, however, because 4 respondents found it difficult to hit the dancers accurately.

Similarly, for the Train Tracks game, 8 participants enjoyed watching the train move on the railway track. But 4 children found the game difficult for reasons such as the obstacles on the game map.

Lastly, we observed children handing over the cellphones to the adult interpreters to restart the game or advance them to the next level. In some cases, usability was the issue since some games such as Floored required the player to advance to the next level by passing through three "next level" screens. But in other games that required pressing a button at only one screen to advance to the next level, we also observed similar behavior with children seeking adult assistance. More interestingly, in some of these situations, we observed an asymmetry in that some children restarted the game in the event of a "Game Over" but solicited help when they did well in the games and needed to advance to the next level. We hypothesize that these solicitations for help were in fact attempts to gain the approval of adults by showing them one's accomplishment in a game.

7.5 Discussion and Lessons Learned

We found that it is too simplistic to make generalizations about patterns, which are akin to decontextualized formulas. We realized that there are contextual factors that we should have taken into account when we employed patterns to design mobile games for our target users in rural India. On one hand, patterns can be building blocks for successful games. As an example, the **coveted player identity** of the hero and trouble-maker in Crocodile Rescue and Dancer respectively appealed to the children. Similarly, a reason for Critter Crossing's popularity was the fact that its design incorporated **characters** such as rabbits that appealed to children.

On the other hand, failure to apply patterns in a manner that is contextually and culturally appropriately is likely to bring about poor gameplay experiences. For instance, introducing **deliberate obstacles** in Crocodile Rescue as crocodiles conflicts with the typical village sentiment (and sensibility) about these dangerous creatures. Equally important is the observation that **creating time pressure** may be a common approach to making games more challenging and engaging, as recommended by Malone (1980), but negative comments about the time limits from our participants suggest that electronic games for rural children, who have less exposure to video games as their counterparts in developed countries, should not be designed to be overly difficult.

In fact, this point about difficulty relates to Lazzaro's (2004) distinction between Hard Fun and Easy Fun. By the former, she means an engaging play experience that challenges and rewards the player for each tangible progress, similar to Malone's (1980) challenge heuristic, whereas Easy Fun refers to the player's enjoyment of a game when his attention is focused on enjoyment of the experience as opposed to the winning conditions or final goals. An analogy might be the enjoyment of life as a journey instead of the destination(s).

From our study with rural children in India, it seemed that Easy Fun plays a significantly more dominant role in this cultural context than Hard Fun in unlocking an engaging play experience. We saw that participants reacted favorably to Floored, Crocodile Rescue and Critter Crossing because they could immerse themselves in the flow of the gameplay experience, while they found Beginner Land, Toy Factory, Jump Bot, Train Tracks and Dancer to be less than pleasant because various factors, including challenges coming from time limits, impeded the optimal flow that is required for an engaged play experience. We therefore recommend that the designer adopt Easy Fun as a primary principle in designing electronic games for this user group.

But we also caution that Easy Fun alone is inadequate. As we have observed in Crocodile Rescue, participants liked Easy Fun, but our mis-application of the **deliberate**

obstacles pattern resulted in this game receiving the second lowest average rating among those games in our sample. The broader lesson is that there will be interaction effects between patterns, which highlights the importance of combining patterns holistically such that the whole is greater than the sum of its parts. A pattern which is a poor fit with the other patterns is sufficient to impair the play experience drastically.

To complement design patterns, the above results indicate that aesthetics continue to matter to a significant extent. Our experience with Floored suggests that games be designed with rules that allow the player to change the aesthetics of game elements, especially in the case of a culture where vibrant colors have a major place in festivities and other cultural events. The aesthetics of game elements that move on their own in the game world will also be welcome by players who enjoy being spectators.

Finally, when designing mobile games for rural children in India, and possibly other rural contexts, it is vital to take the existing power structures into account. Teachers and other adults are undoubted authority figures, especially in a community where utmost respect is accorded to seniority. Games can be designed such that they facilitate the player in gaining adult approval and strengthen community ties. For instance, screens between levels can feature the player's performance and/or improvement over previous sessions prominently. A pause feature can also be added to allow players to show off their proudest moments to others. However, unlike the pause feature in the typical game, which occludes the active game screen, a pause feature that is more visually transparent is more suitable in this context.

In conclusion, Alexandrian patterns have been gaining popularity in design communities because they facilitate the reuse of existing knowledge about successful solutions to common problems. When we began the reported comparative study to examine the extent to which patterns can leverage prior lessons and promote the design of successful games, we observed that patterns can both help and hinder good designs. In the process, we gained a better appreciation of how patterns are analogous to "knowledge constructs” such as formulas in that they are decontextualized artifacts that are to be guided and framed by broader cultural and contextual knowledge in their use. We hope that the contextual factors which we have identified for the effective applications of patterns in a rural Indian context will encourage other designers and researchers to pursue work on the “situational” dimension of design patterns, which have hitherto been viewed mostly as abstract representations of solutions that are applicable to other contexts to a large extent.

8 Design Iterations

Iterative design is undoubtedly integral to successful technology design. Iterative design is even more important when designing educational software for and with users in underdeveloped regions, for several reasons. Firstly, it is difficult to obtain an accurate understanding of the user's educational baseline since it could deviate sharply from the official syllabus or accounts of local informants due to huge variations between locations in factors such as access to quality schooling. Secondly, the user is likely to have very limited computing experience, which is a strong impetus for conducting additional rounds of iterations until the system is sufficiently usable. Thirdly, and most importantly, local stakeholders and designers may not share the same cultural background, which imply the need for continuous co-learning and iterating until the design is consistent with the local culture and social norms.

In this chapter, we show how the pedagogical design patterns from Chapter 5 and the PACE framework introduced in Chapter 6 can be applied to the design of e-learning games. We also describe the receptive-practice-activation cycle that forms the conceptual model for their designs. This model is consistent with the instructional sequence for task-based language teaching that we describe in Chapter 4. We report our early experiences from field-testing these games among urban slums children. While these games are fairly intuitive to children from the slums, rural children found them to be considerably more difficult, possibly because of their relatively lower levels of exposure to technology. We narrate the user acceptance issues that rural children exhibited and how we iterated on the game designs in response to our experiences in the field. Some of the redesigns took our observations with game design patterns from Chapter 7 into account. On the whole, it appears that maintaining a distinction between learning and fun to some extent is crucial to effective design.

8.1 Application of the PACE Framework

Building on the formative experiences with the pedagogical design patterns in the spring of 2006 (see Chapter 6), we embarked on a second round of design in the summer of 2006. In this round of design, we adopted the PACE process as well as structured our games according to the PACE framework. Most of the games that we designed in this round were created from scratch, although we consolidated some ideas from the spring's designs. This section explains how we applied the PACE framework to the design and implementation of mobile e-learning games that we subsequently field tested with urban slums children in North India.

8.1.1 Pattern

In Chapter 5, we had described over 50 design patterns for language teaching that we had earlier distilled from more than 35 commercial language learning applications. The applications were selected such that the overall sample reflects a balance between listening, reading, speaking and writing skills. We then distributed handouts describing the patterns to our team members. We also shortlisted those patterns that are appropriate for learners in the initial stages of the second language learning trajectory. In total, we identified 11 such patterns that are balanced between spoken and written English. These patterns focus on phonetic decoding (i.e. how to "sound out" letters and syllables, and segment words into their syllables), pronunciation, listening comprehension and sight reading. Patterns that target the last two aspects involve word-picture matching activities, such as the pattern shown in Figure 5.5. Other patterns involve matching visual symbols of letters or syllables with their sounds. One more pattern is Pimsleur's graduated interval recall principle, which is based on theories of how human memory works and is popular in the Pimsleur series of audio-only language learning recordings.

8.1.2 Activity

Since two of the patterns, i.e. pronunciation and Pimsleur's graduated interval recall principle, could be combined with other shortlisted patterns, there were a total of 9 learning activities out of the 11 patterns that we could be implement as independent learning activities. We prioritized them for development on .NET Compact Framework 2.0 for the i-Mate SP5 smartphone. We target this expensive phone because the ease of prototyping on .NET CF facilitates numerous rounds of iterations, which is indispensable when we expected to make substantial modifications in response to user observations and feedback from local stakeholders. The same activities can be subsequently ported to more inexpensive phones once user acceptance issues have been resolved.

Two over-arching considerations guided how we designed the learning activities. First, we expected that target learners could be first-time users of cellphones. Hence, we designed the games around a conceptual model that we reused in every game. Having a conceptual model common to every game is expected to promote the learnability of their user-interfaces. From the instructional design standpoint, we designed every activity such that it consists of receptive-activation cycles in which the linguistic items (for example, vocabulary words) are introduced to the learner are selected based on Pimsleur's recall principle. More specifically, every cycle comprises:

a receptive phase – which targets one area of competency in English (for example, sight vocabulary or the alphabet) , followed by

an activation phase – which tests the player on the targeted language competency. This phase tests the player on some or all of the linguistic items that were "taught" in the preceding receptive phase. However, the player would be tested on the items in a sequence – usually random – different from the sequence in which they were presented in the receptive phase.

To avoid overwhelming the player with too much new material at any time, the conceptual model was designed around several short receptive-activation cycles. In this way, the player is presented with a small amount of material in each cycle. She is then tested on it before she receives feedback on her answers. Only then is a new cycle started to cover more material – both new and repeated (from previous cycles) – until she obtains correct responses on the entire curriculum.

Let us illustrate the receptive-activation cycle by way of the "catch the parrot" game (Figure 8.1), which is intended to help learners improve their grasp of letter-sound correspondences, i.e. relate letters with their sounds. More specifically, the parrot game attempts to teach three letters in every receptive-activation cycle, such that each receptive phase focuses on the three selected letters for that phase. The game combines visuals with sound: when teaching a letter, it shows the symbol for the letter within the speech bubble of the human sprite when the human moves from one parrot to the next, such that every parrot repeats the targeted letter aloud as soon as the human moves beside it. The game is done with the current letter once the human has moved off the screen, and the remaining letters are covered until the game has completed all three letters for the current cycle. The player then transitions to the activation phase, in which she is tested on the same three letters.

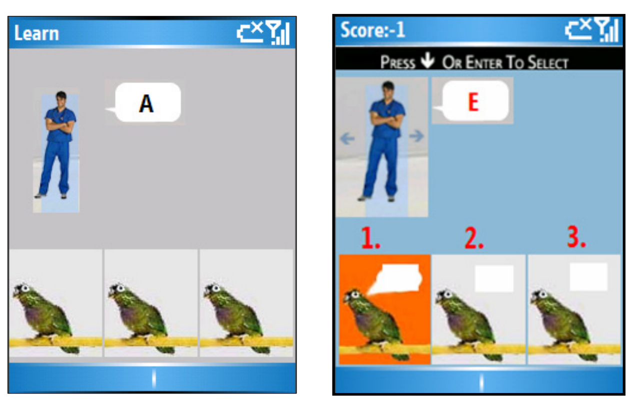

Figure 8.1. The receptive phase (left) and activation phase (right) in the initial parrot game. The learner was being taught "A" among other letters in the receptive phase and was tested on the letter "E" in the activation phase shown.

In the corresponding activation phase (Figure 8.1), for the letter that the player is currently tested on, the game shows its symbol in the speech bubble of the human sprite. Unlike the receptive phase, however, every parrot corresponds to a different letter, and only one parrot stands for the letter that the player is being tested on. The player can move the human sprite from one parrot to another, such that the parrot that is currently beside the human says aloud the letter that it represents. The player is required to choose the parrot that correctly corresponds to the letter that she is tested on.

Our conceptual model was designed so as to revolve around multiple receptive-activation cycles. It was informed by the above review of over 35 commercial language learning applications. To what extent is the receptive-activation format that we gleaned from bestselling commercial software applications for industrialized markets applicable to low-income learners in underdeveloped regions? What adaptations are required for this format to be more relevant for the latter users? To address these questions, we will turn to our experiences in the field shortly.

Our second over-arching consideration was a conscious attention not to follow the patterns rigidly in a formulaic manner this time. We also realized that most of the designs from the spring of 2006 did not seem interesting to play. As such, we applied Malone's heuristics (1980) on how games could be designed to be engaging and fun. For instance, we introduced difficulty levels in which time limits were imposed to make the games more challenging. We also adopted Malone's recommendation to situate our games in fantasy settings, so that game settings do not resemble paper-and-pencil tests too closely. In consultation with Indian members in our team, we came up with game settings from everyday Indian experiences. Figure 8.2 shows two activities that are based on culturally-relevant settings. Indian mythology is replete with heroes who excel in archery, hence the bow-and-arrow setting of the first game. We chose the train setting for the second game because the train is a mode of transport that is familiar to all income levels in India.

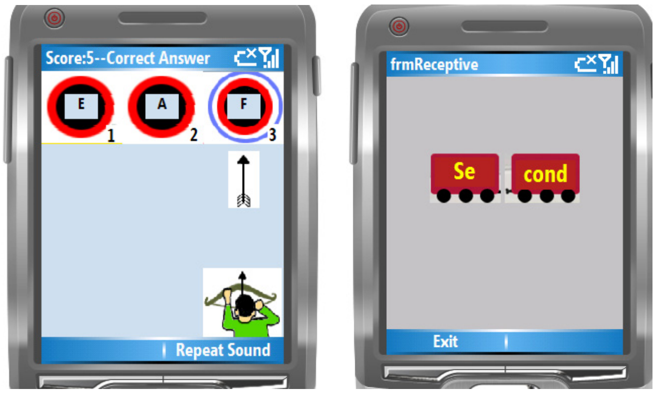

Figure 8.2. The bow-and-arrow setting for the learning activity that instantiates the "Phoneme → Grapheme Association" pattern (left), and the train setting for the activity that instantiates the "Syllable Segmentation" pattern (right). In the former activity, the game plays the sound of a letter (i.e. phoneme) and the player shoots the archery board that shows its corresponding visual symbol (i.e. grapheme). In the latter, train carriages represent multi-syllable words that the player breaks up at syllable boundaries into smaller carriages.

When designing the games, we also draw on lessons from our exploratory studies (see Chapter 3), in which we were told that rural students persisted in replaying electronic quizzes until they obtained the satisfaction of achieving the maximum score. As such, we added score-keeping (penalties) to reward the player for selecting the correct (wrong) answers. In total, we implemented 6 out of the 9 learning activities in time for fieldwork in the late summer of 2006. The activities that we implemented included those activities shown in Figures 6.3, 8.1 and 8.2.

8.1.3 Curriculum and Exercise

Mindful of the difficulties that we faced in our design attempts in the spring of 2006 to identify words that are culturally appropriate for the shortlisted functional areas, we turned to a Hindi phrasebook that targeted the needs of English-speaking tourists who visit north India (where Hindi is the lingua franca). Since the functional literacies covered in this phrasebook overlapped with the functional literacies that we have shortlisted to a large extent, we referred to it for the concepts (expressed in Hindi) that are meaningful to our target learners. After which, we thought of the English words that expressed the same ideas. These English words constituted our initial syllabus.

We then divided this syllabus into a total of 21 curricula for the selected learning activities. The curricula targeted the alphabet and functional literacies such as numbers, dates and time, shopping, traveling, nature and social situations. For the English alphabet, we maintained separate curricula for the lower-case and upper-case letters. We also kept the curriculum for each functional literacy separate. This separation enabled us to benefit from the advantages of learnability, reuse and scalability that the PACE framework promotes. The curricula were implemented in XML, with the audio and images stored in separate binary files. In particular, the images were first obtained from commercial clipart libraries and online repositories, after which they were edited for local appropriateness. For example, the uniforms of the police officer and bus conductor shown in Figure 6.3 were edited such that their colors matched those of Indian institutions. The curricula also included Hindi voiceovers that provided explanations in the native language for the target English words.

8.2 Experiences with Urban Slums Children

We conducted a fourth round of fieldwork in Uttar Pradesh in the late summer of 2006 to evaluate the games that we had designed and implemented earlier that summer. Our study took place in an afternoon school that was founded and directed by one of our non-government organization partners. This program targets girls who live in the nearby slums, who would otherwise not have an opportunity to receive formal schooling. Classes are as such free-of-charge. Each class session lasts 3½ hours and is held every afternoon since students have to perform household chores in the mornings. Students are recruited into the afternoon school program through a combination of two ways. Firstly, teachers make home visits to convince some parents about the importance of formal education for their daughters. Secondly, other parents hear about the program through word-of-mouth and enroll their daughters subsequently.

This field study lasted an average of 2 hours every morning over a 2-week period. We hired two local adults as interpreters. Children were videotaped playing the mobile games. We obtained a total of 4.5 hours of video recordings, which we later transcribed and translated into English. We also visited 5 students at their homes to better understand their backgrounds, which helped us to better contextualize our observations. This study raised questions that were interesting from a human-computer interaction point of view: how quickly can children in an underdeveloped region who have never used cellphones learn to use them? How well can the same children use cellphones for language learning? Will they find the game-like experience engaging? These questions were important since this study was our first attempt to deploy e-learning games on cellphones with children in a developing country context, and hence constituted a feasibility study.

8.2.1 User Study Sessions with Kindergarten and First Grade Students

On arriving in India, we demonstrated the games to our NGO partner. After some consultation, we concluded that the games for the English alphabet were suitable for the 14 kindergarten and 1st grade students who were enrolled in her afternoon school. They were aged between 4 and 6 years old.

This was a surprise to us. We had originally designed our activities and their user-interfaces for 10-year old children, since we expected the ESL baseline of children in the slums to be close to that of the rural children whom we interacted with in summer 2005 (see Chapter 3). On hindsight, we realized that even though the average student from the slums will not attain an ESL competency comparable to her middle-class counterpart in India, the average participant in the afternoon school program is nonetheless significantly more advanced than her rural counterparts who study in government-run schools. This is attributed to our NGO partner's commitment to education and administration of a high-quality program. As a result, we had to change the user-interfaces to suit users who were considerably younger than what we expected. In particular, we altered the feedback for correct and incorrect answers from text displays to smileys and frowns (see Figure 6.3).

We spent one morning to train participants on how to use the cellphones and play the English learning games. In general, the games were introduced in the following way: at the start of each session, we briefed every child on the learning goals for that day. We next divided the learners into smaller groups with about 7 to 8 children in each group, after which we showed them how to play the game that we had scheduled for that day. This demonstration lasted 10 minutes, after which a cellphone was handed to each child to play the given game. We observed some usability problems that we fixed during the afternoon of the same day. For example, we removed the time limit for the "easy" level of difficulty because children appeared to have trouble with the joystick button on the smartphone, which resulted in them taking longer than expected to choose their answers. We also introduced more animations to make the receptive phase of one activity more intuitive.

With the training over, participants played the same games for the next two days. Since we did not have enough smartphones, participants needed to share the equipment. Hence, as they took turns to use the smartphones, they reached across each other for the phones. They gestured and shouted their answers collectively. They also always repeated each letter aloud after the game, as they might do with a teacher in a classroom setting. Even though we made sure that each participant got her turn, the children continued to demand that their neighbors hand over the cellphones back to them. Moreover, none of them noticed our camera when we videotaped them because they were engrossed in the games. From the above observations, we concluded that participants found the games to be highly engaging.

One notable usability problem arose on the first day after the training day. On the bow-and-archery game (Figure 8.2, left image), children sometimes missed the letter that was read aloud. Not knowing which of the visual letters they were supposed to select to match the audio prompt, learners had no choice but to choose any letter so as to go on to the next question (and audio prompt) in the activation phase. We fixed this problem on the same day by adding a "repeat sound" menu option. Gameplay became less frustrating the next day, which made it easier for learners to focus on learning the alphabet.

In comparison to usability, it was less clear if users learned the alphabet. A pre-test indicated that 3 of the 14 students did not know any of the letters since they enrolled in the afternoon school only a week ago, and it was their first schooling experience. Other students knew two-thirds of the letters, on average. After both days of gameplay, one participant completed all 6 exercises that we developed; most children completed only 2 or 3 of the 3 exercises that collectively covered the English alphabet. The latter provided some indirect evidence that some of the children had learned all the letters, because it was necessary to answer every question correctly in order to complete the 3 exercises which targeted the alphabet. A post-test would be more conclusive. But we decided to cease our study with this group of children after two days. This was because their teachers were present throughout our study, to the extent that we noticed one of the teachers providing her students with occasional hints when they were playing the games. As such, even if we could find time to conduct a post-test, there is nevertheless potential for bias arising from teacher intervention. It seemed that teachers may not have understood the objective of the research study and were under the impression that we were evaluating how well they had taught the children, as opposed to us evaluating the e-learning games. As a result, one or more teachers may have introduced and be teaching similar material in class concurrently with our study, which made it difficult to determine if learning gains were attributed to the teachers or cellphone games.

8.2.2 User Study Sessions with Sixth Grade Students

In consultation with our NGO partner, the games that targeted vocabulary were evaluated with her 11 6th grade students. They were aged 11-15 and had attended school regularly for 3 to 6 years. Although they had learned English throughout this period, we still needed interpreters to communicate with them beyond simple greetings. Our NGO partner also opined that some words in our original curricula were too simple. Hence, we first conducted a pre-test to determine those curricula whose words participants fared least well on. In other words, our activities for this second study would focus on teaching the words in these functionally-organized curricula. We spent two days after the pre-test to add more difficult words, as well as create their image and audio files, to the selected curricula at our partner's request.

In all, we piloted 12 exercises over 5 days, after which we conducted a post-test. Every pair of exercises introduced about 10 words, such that one exercise in each pair involved the user matching pictures with given words while the other exercise involved matching words with given pictures (i.e. the converse). In other words, the games introduced at this stage were based on the "Semantics → Written Word Association" and "Written Word → Semantics Association" patterns from Chapter 5. A notable challenge was that it was not easy to conceive of graphics that intuitively conveyed what their corresponding words meant. We had to consult a native in Uttar Pradesh to understand the cultural conventions that the illustrations needed to embody, and spent days prior to the study iterating on the images.

The most common usability issue, which cropped up more than 4 times as often as other problems, was the joystick button reported above. Other than that, participants were generally making good progress once they understood the conceptual model behind the games. But one of the participants left the study mid-way because she was struggling with usability problems and felt discouraged when she saw that her peers were making more progress. The problem that we most expected to see, but did not arise, was the login screen. We had implemented this screen at the start of each activity so that we could track learners' scores (for data analysis) and thought that we will be entering their names for them. We were amazed that each participant learned to perform text input after observing us do it within the first 1 to 2 days, especially when this was their first experience with cellphones.

During the first two days, there were several instances when we asked participants if they wanted a break. However, they would refuse, stating that they wanted to continue with the games and finish playing them. On the third day onwards, participants began to finish the 12 exercises. As more participants finished the games, those participants who were lagging behind showed reluctance to continue playing. It seemed the games were appealing until the atmosphere became competitive. In contrast to our experiences with the teachers of the previous group of students, the teacher of the 6th graders was more comfortable with us since we had met during previous field studies in Uttar Pradesh. This made a more reliable post-test possible since she did not find it necessary to accompany her students throughout our study. The learners exhibited post-test gains of 4.3 out of 12 points on a t-test ($p < 0.001$, $\sigma = 3.7$) when we showed them words from the curricula and asked them to read those words as well as explain their meanings.

8.3 Initial Experiences with Rural Children

The above "catch the parrot" game was among those games which we evaluated during the fourth field study in August 2006 with kindergarten and 1st-grade girls living in the urban slums. Our experience in piloting the parrot game with rural children in India would turn out to be significantly more challenging. It was deployed in the fifth study in January 2007 with 24 1st- and 2nd-grade children at a village school in Southern India, in collaboration with a different NGO partner. Like participants from the urban slums in the August study, none of the participants in the January study reported any experience with cellphones but understood what cellphones were.

We consulted our NGO partner, who believed that participants were beginning to learn English and should learn phonetics as part of their early literacy curriculum so that they are equipped to decode words phonetically (i.e. read aloud by sounding out). We had not designed this game originally to teach phonics, but the NGO had seen this game in an earlier presentation and advised how it could be adapted to cover phonemes (i.e. the basic units of sound) and their respective graphemes (i.e. symbols for individual or clusters of letters). Moreover, our NGO partner provided a syllabus of the grapheme-phoneme pairs that we should cover. Finally, the NGO recommended that the parrots in the activation phase (Figure 8.1, right image) be modified until they became visually distinct from one another. This change was meant to avoid causing potential confusion since each parrot represented a different phoneme.

8.3.1 Experiences with Parrot Game (Second Iteration)

We piloted the parrot game as soon as we have iterated on it based on the above feedback from the NGO. In contrast to our relatively smooth experience with urban slums children in August 2006, guiding rural children of a similar age bracket to play and learn with it proved to be a more formidable challenge. While the former could use the parrot game on their own to learn how to pronounce the letters in the English alphabet, the latter required our interpreters to sit beside them to coach and provide feedback on the pronunciations of the various phonemes. We do not think that the reason stemmed from usability or learnability issues since we had been familiarizing participants with mobile games and cellphones for more than a week. In fact, we observed that they were coping reasonably with the cellphone games in the January 2007 study by the time we introduced the parrot game.

Instead, we believe that rural participants needed adult facilitators to assist them in learning with the games because the former had a lower English baseline compared to their slums counterparts, who had enough exposure to the sound system of the English language to benefit from the instruction provided by the parrot game. Some background is necessary here: while slums students in the afternoon school receive English lessons in kindergarten and more advanced classes and are also exposed to the use of English by urban dwellers, the rural school introduces English only in the 5th grade and hardly any villager speaks English. As such, the parrot game appeared to build on the knowledge of the English phonemes that urban slums children had, by reinforcing their understanding and retention as they repeated aloud after the game.

In contrast, rural learners appeared to lack this background. They depended on our interpreters to perform the role of facilitators, who repeated the phonemes aloud from the game and helped participants to distinguish between phonemes that sounded similar. The rural children also articulated the phonemes less often and with more mispronunciations than slums learners. In fact, it seemed that they realized their errors only when facilitators corrected them. It came as no surprise that the given design for the receptive phase was inadequate for preparing rural participants to be tested in the subsequent activation phase, and we observed that most of them were not able to select the correct parrot.

8.3.2 Third Iteration

The above observations that we made on the first day prompted us to make two major modifications in the next iteration, which came in time for further field testing on the second day. First, we introduced a "practice phase" (Figure 8.3) after every receptive phase and before its respective activation phase, so as to provide the learner with plenty of exposure and practice with linguistic items in the curriculum for as long as she wants. Unlike the receptive phase, in which items are introduced to her at a pre-programmed pace, in the practice phase, she has the freedom to switch freely between the items for the current cycle by moving the human sprite to any parrot to listen to its corresponding phoneme. In this way, she gains the flexibility to focus on those phoneme-grapheme mappings she is weaker on in a *self-paced* manner, so that she could better prepare for the ensuing activation phase. In the particular case of the parrot game, this flexibility allowed her to listen repeatedly to more than one parrot until she could easily tell their phonemes apart.

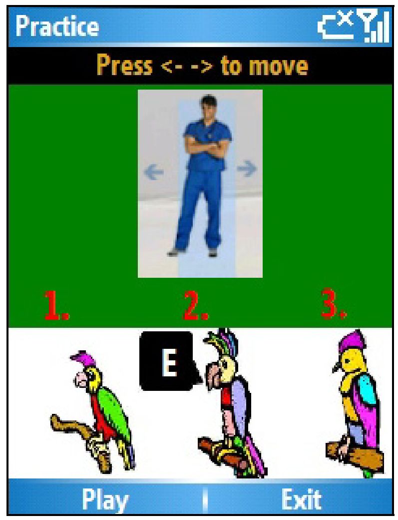

Figure 8.3. The practice phase in the redesigned parrot game. The human was next to the parrot that was echoing the /ɛ/ phoneme (pronounced as "air") for the "E" grapheme.

Second, we introduced user-interface controls for the facilitator in the receptive, practice and activation phases via the left and right soft keys on the cellphone keypad (Figures 8.3 and 8.4). The soft keys were labeled "Exit," "Start" and "Play" and provided the facilitator with more control over the flow of the parrot game. For instance, the Exit functionality was added to make it easier for the facilitator to quit and restart the game should she learn that the learner lacked the prerequisite knowledge for the current phase and needed to return to an earlier phase to relearn the basics. Similarly, the Start button was added to the receptive phase at the request of the facilitators to enable them to start this phase explicitly. They commented that students were not always ready to concentrate on the pre-programmed instructional sequence in the receptive phase and were often caught off-guard when a receptive phase began, possibly because rural learners lacked familiarity with interactive software to realize that the end of an activation phase is always associated with – and followed by – the start of the receptive phase for a new receptive-practice-activation cycle.

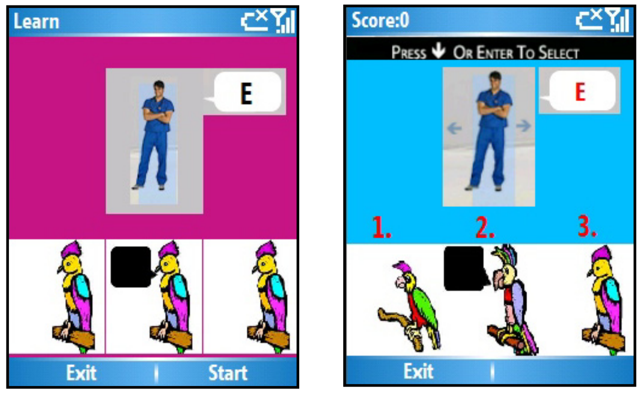

Figure 8.4. The user-interface controls for the adult facilitator in the receptive phase (left) and activation phase (right) of the redesigned parrot game.

We also made other modifications in response to feedback from the facilitators. For example, since learners needed time to repeat the phonemes aloud after the game, we slowed down the audio playback of the phonemes in the receptive phase by increasing the pause intervals between phonemes. The additional delay gave learners, especially those who were less confident or prepared, more time to struggle with the grapheme-phoneme associations and the pronunciations. It also enabled facilitators to provide more thorough explanations, corrections and feedback. Likewise, since the cellphones were held closer to the children than their facilitators in many cases, we increased the font sizes to enhance visibility for the latter.

After the parrot game was deployed on the second day with the above revisions, the facilitators told us that the changes made it easier for them to guide the rural children in learning phonetics. The facilitators also reported that the children were learning what they covered. But based on the results of a recall test that we conducted immediately after the deployment, we found that 13 of the 25 students who had played the game did not remember any of the 7 phonemes that were covered. Two additional students did not feel confident enough to respond to the test questions. Among the 10 remaining students, the average number of phonemes that each of them recalled correctly was 1.5 out of 7. It was plausible that learners needed more than a day to develop a stronger grasp of the material before learning effects would show up on formal tests. However, the more likely reason for the poor test performance according to the facilitators was that learners had perceived the parrot game as an application to be played for pleasure, vs. learning. Facilitators drew this conclusion after they observed that participants were focused on having an enjoyable time with the game without putting in the time and effort to remember the material. This observation reminded us about the sociocultural tradition of task-based language teaching that emphasized learner orientation towards the learning task, as opposed to task design *per se* (see Chapter 4). It made us better appreciate the need to remind participants about the educational objectives of the games in future briefings, so that they would be more likely to approach the games as an educational (vs. entertainment) goal-directed activity.

The refined conceptual model around the receptive-practice-activation cycle was not without its drawbacks. The most significant problem was that the game with three phases was comparable to an application that has three different modes, and modes have been widely acknowledged as a common cause of usability problems. Worse, because we wanted to situate all three phases in the "catching the parrot" setting, we were constrained to maintain as much of the same "look and feel" for all three phases. As a result, the rural learners found it difficult to see which phase they were in, possibly due to their limited computing experience. We tried to address this issue by having three distinct background colors for every phase, and with affordances such as arrow signs for the human sprite in phases where the user is allowed to move it, but it remained difficult for the learners to associate the current screen with the phase that they are in. It would take one more field study to find a satisfactory solution to the above problems.

8.4 Further Experiences with Rural Children

Our next field study took place throughout eight weeks in South India, at the same rural school where we worked in January 2007. We designed and implemented a broader range of applications so that we could target a more comprehensive curriculum beyond the alphabet and phonetics that the parrot game targeted. After which, we iterated on the games based on our experiences with them in the field. Among the applications we tested in this period, we focus on Floored and Frogger in the remainder of this Chapter because their designs build on the above lessons from the parrot game. Our experiences with both Frogger and Floored are expected to yield further insights for designing e-learning games with lower-income learners in India that balance the goals of entertainment and learning.

Whereas the parrot game targeted letter-sound correspondences, the Frogger and Floored games aimed to develop the learner's vocabulary in spoken English. Their syllabi were adapted from the English curriculum recommended by the state government. In the summer of 2007, we came with more cellphones than in January so that we could include all the students from the school in our research study. The larger sample size enabled us to avoid evaluating subsequent iterations of the games with the same students each time, i.e. learning effects. In all, 47 students from grades 2 to 5 took part in the study. 1st grade students had to be excluded since they were in the process of registering for school at the beginning of the academic year.

8.4.1 Initial Game Designs

Frogger and Floored were not designed by their original creators to teach English as a Second Language. We adapted the game designs and incorporated ESL curricula into them after we learned that they were the two most popular games among the eight games that we previously tested with the same children (Chapter 7). Our adaptations revolved around the receptive-practice-activation cycle that had evolved from the above field trials with the parrot game. The advantage of this cycle is that we can introduce new words to learners and give them repeated exposure to the same material at their own pace, before testing their newly-acquired knowledge by requiring them to apply it to accomplish goals in a game setting. In contrast, educational software that target markets in the developed world do not usually include features that resemble the receptive and practice phases. The omission was not expected to be a serious problem for learners in industrialized countries who come from middle-income families and typically have well-educated parents and teachers.

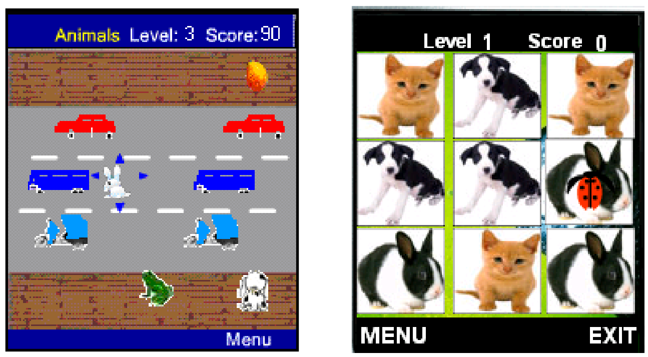

Figure 8.5. The activation phase in the initial iteration of the adapted Frogger (left) and Floored (right). The player was helping the selected animal (i.e. the rabbit) to cross the road in Frogger, and flipping tiles to show the cat in Floored.

For clarity of exposition, we will first describe the activation phases for both the Frogger and Floored games (Figure 8.5) before we describe their preceding receptive phases (Figure 8.6). In the original Frogger game, the goal was to help a group of animals cross the road one at a time without getting run over by a vehicle. In our adaptation of the Frogger game, we wanted to target the English vocabulary for everyday animals. When an activation phase begins, the game says the word for one of the animals whose sprites are shown at the bottom of the screen. The objective of the player is to select and help the correct animal cross the road.

In the original Floored, the goal was to flip the floor tiles until every tile is facing up with the same side showing. There are two or more states associated with each tile, such that when the player moves the ladybird sprite to a new tile, the destination tile flips to show its next state. We adapted Floored such that whenever an activation phase begins, the game plays a word aloud. The player has to flip the tiles such that all of them show the picture for that word. That is, we adapted Floored for vocabulary building, especially for nouns and verbs that could be conveyed pictorially. We reused Frogger's curriculum in Floored so that the latter also targeted vocabulary related to animals.

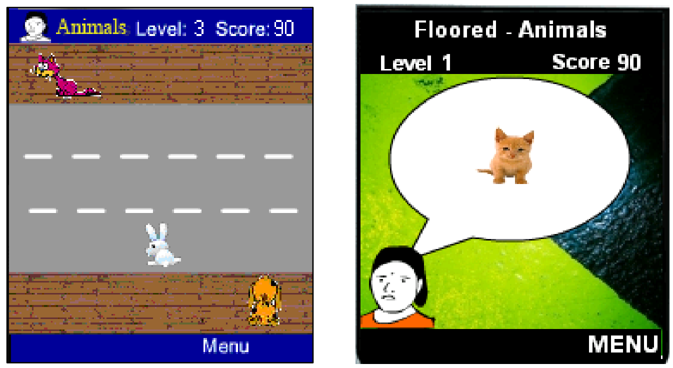

Figure 8.6. The receptive phase in the initial iteration of the adapted Frogger (left) and Floored (right). In Frogger, the player had been taught "cat" and was being taught "rabbit." In Floored, she was being taught "cat."

In Frogger, in any receptive phase, the spoken English words for three animals are taught. The animals start at the bottom of the screen, and when each animal crosses the road, the game says the word corresponding to it. The audio playback takes place twice – once at the normal pace and then again at a slower pace to help the learner follow its pronunciation more accurately. As we have discovered in January 2007, pronunciation is particularly difficult for rural learners and we expected the slower audio replay to benefit them. In addition, to make it more intuitive to the player that the audio playback refers to the animal that is being taught, a caricature of the player's school-teacher is shown in the top left-hand corner. The caricature enlarges and then reverts to its normal size whenever a word is played back, so as to create the impression that the teacher character is teaching the spoken word for the animal. We will touch more on the caricature below.

Once all three animals have crossed the road, the receptive phase transitions to the practice phase, where a bicycle moves onto the screen. The bicycle acts as a "pointer" which the player can move to the left or right such that it is always next to one of the three animals who are now at the of the screen. Each time the bicycle moves beside an animal, the game shows a larger image for that animal and plays the audio for its word. In this way, the player gets as much exposure to the spoken English words for the animals that she needs. Unlike the activation phase, however, there is no vehicle traffic in the receptive and practice phases to avoid distracting her from vocabulary learning.

On the whole, our adaptation of Floored differed from Frogger in three significant ways. First, Frogger's receptive phase was situated in the "road crossing" game setting in the hope that immersing vocabulary learning within a game setting would make learning more pleasurable. In contrast, we had situated Floored's receptive phase in the classroom setting, such that the caricature of the teacher at the bottom left-hand corner of the screen would say the spoken English word for the picture in her speech bubble. We were not satisfied with this design at that time, but could not ideate a less mundane receptive phase that leverage Floored's tile-based game setting. The differences in the receptive phases in both games turned out to have important consequences, as we will discuss shortly.

To add, coming back to the caricature of the teacher, the English words in Frogger and Floored were recorded in the voice of the teacher in the caricature. This was because in an earlier fieldwork with the same participants, we found that they did not always heed the audio prompts in another application. Those prompts were recorded in the voice of a team member whom the children did not view to be an authority figure, to the extent that they teased, bullied and even pinched her. Hence, we wanted to investigate if the voice recordings of a more confident stranger, or that of an authority figure (e.g. the children's teacher) accompanied by a visual representation (e.g. her caricature), could better engage the children's attention. This idea was contributed by our NGO partner, who believed that participants would take the audio recordings more seriously when they recognized the recordings to be in their teacher's voice.

Second, mindful of the difficulties that players in January had with distinguishing between the three phases in the receptive-practice-activation cycle, we integrated the practice phase with the activation phase, i.e. we did not design a separate practice phase. Instead, in the activation phase, each movement of the ladybird is accompanied by the audio playback of the word corresponding to the new picture displayed on the destination tile. The player therefore receives additional exposure to the words for the tiles each time she moves the ladybird. The audio playback in the activation (and practice) phase took place at the normal pace while playback in the receptive phase took place at a slower pace. We would have liked to cut down on the number of separate phases in Frogger, but we needed to work within the constraints of the game rules (since we wanted to build on a game that has been shown to appeal to our target users) and could not think of how the settings for any two phases could be integrated naturally.

Third, we needed to make the transition from an activation phase to a receptive phase more evident so that learners would not be caught unprepared when an activation phase ended and they needed to focus again on the instructional sequence in the ensuing receptive phase. Our solution in Floored was a simple "You Win!" splash screen that required the user to hit a button to continue to the receptive phase. It was informed by the January study, when users showed off the level completion screens in their mobile games to adults like us, as a means of gaining our approval and recognition for their gameplay achievement (see Chapter 7). We implemented this feature for the summer to investigate if the same behavior would repeat in an e-learning game. In contrast, for our adaptation of Frogger, we wanted to experiment with a fancier transition and had the animal walk off the screen after it reached the destination lane.

Even though the above contrasts between our adaptations of Frogger and Floored were sometimes unintentional and arose from serendipity, they seemed to have profound impacts on how participants interacted differently with both games. Their differences and similarities constituted the basis for some interesting design lessons, which we will next discuss.

8.4.2 Key Questions

The parrot game and the ensuing analysis above had raised some questions: Can rural children tell the three phases in the receptive-practice-activation cycle apart? Were the first two phases effective in equipping users with the knowledge to succeed in the activation phase? Were the caricature and voice of the teacher sufficiently “authoritative” in drawing learner attention to the material taught? Were games designed based on this three-phase model engaging and interesting to play? Were players learning through the games? Although we had some preliminary results with the parrot game, we sought to address these issues more conclusively by testing Frogger and Floored in the field. Most important, although more research is needed to arrive at generalizable lessons, we argue that those observations that repeated across more than one game spoke to a greater degree of generalizability for designers.

8.4.3 Lessons From First Iteration

We evaluated Frogger and Floored with 29 students on 3 days and 16 students on 3 days respectively. A session for any child lasted an hour on average and took place under natural conditions at the rural school. On average, a player took about 15 minutes to familiarize herself with a game, and the process usually involved playing its activation phase more than three times. The rest of her session was spent on gaining more practice with the game, and players began to win the game after they were first introduced to it for 30 minutes. For the players who completed all the levels in the game before their sessions ended, we asked that they replay the entire game so that they could potentially become more familiar with the targeted English material.

We were naïve in thinking that situating the receptive and practice phases for Frogger within its "road crossing" setting would help to make the learning experience more pleasurable. Players seemed to grow bored with the receptive and practice phases after playing these phases more than thrice, and were eager to skip ahead to the activation phase to play the game proper. In other words, even though we did not expect formal schooling to be interesting to rural children, situating our learning episodes within a game setting did not help to make the curriculum or learning experience any more appealing.

Worse, by situating both the receptive and practice phases in the "road crossing" setting, we made the learning goals behind both phases less obvious to rural children, who appeared to associate education more closely with a school setting as opposed to a game setting. In particular, during the practice phase, participants focused on moving the bicycle "pointer" in the leftward direction, more so than paying attention to the audio playback of the English words or the enlarged image of the animal referred to by the pointer. We offer two plausible reasons. First, players found it more intuitive to move the pointer to the left, versus right, because the bicycle in the icon was facing left, and back-pedaling in the real world was possibly inconceivable to them. (However, it strained our imagination to devise a pointer interaction technique that was both a good fit with this setting as well as afforded (i.e. provided cues for) the left and right movement actions.) Second, the game setting together with the opportunity to ride a bicycle in this fantasy world was a possible distraction to learning.

More troubling, in the receptive phase, players did not appear to understand that the animations of the animals crossing the road were meant to teach them vocabulary, and that they should paying attention to the audio recordings. Alternatively, learners may knew about the vocabulary building objective but chose instead to focus their attention on watching the animated creatures cross the road, as opposed to concentrating on the audio playback.

There were also navigational difficulties with the three phases in Frogger. For instance, we observed that players did not always know that they had to move the bicycle "pointer" off the screen in the practice phase to proceed to the activation phase. Similarly, learners were sometimes confused between the words tested in the activation phase, and would benefit from having an option to return to the practice phase from the activation phase for additional instruction. Having consistent user-interface controls that provide the user more control over her transition between phases is expected to enhance the usability and learnability of the user-interface, and in turn, English learning. We implemented such controls in the next iteration.

In comparison to Frogger, rural children appeared to understand the receptive and activation phases, as well as their distinction, in the Floored game significantly better. Even though it was possible that the users made sense of the two phases in Floored due to the game having fewer phases (than Frogger), we argue that it were the dramatically different settings – beyond cosmetic differences in background colors or variations on the same fantasy theme – which helped to accentuate the sharp distinctions between phases. Most importantly, it was intuitive to players that they needed to pay attention to the voice recordings and visuals in Floored's receptive phase because it was situated in a classroom setting, which players instinctively associated with learning and schoolwork. In fact, we observed that learners paid more attention to the receptive phase in Floored than the same phase in Frogger. For example, when playing Floored, children eventually began saying aloud the words for the animals as soon as their pictures appeared in the speech bubble, before the game's audio playback of the same words, even though we found that the same learners did not know any of those words during their pre-test.

On a related note, for both Frogger and Floored, players did not recognize that the voice for the audio recordings belonged to their teacher, when we asked whose voice the recordings were from. In fact, none of them recognized that the caricature was that of their teacher either.

Figure 8.7. Rural children displaying their "You Win!" screens. Success in the English language learning games, and mastery of the technology, was a considerable incentive and source of pride for these children. These kinds of displays were repeated over and over.

Next, each time the "You Win" splash screen appeared in Floored, players were visibly excited and would clamor to show their cellphone screens to us and their peers (Figure 8.7). On the other hand, the animal walking off the screen in Frogger appeared to be less intuitive in communicating to the players that they had won, since they would look perplexed when the game transitioned to the receptive phase. Most importantly, it was precisely the design of the "You Win!" screen, which required the player to press a button to continue to the next phase, that made it possible for the children to bring their cellphones around the classroom to show off their exploits. While we had observed similar behavior in a few children in January 2007 with games that focused entirely on entertainment without any educational goals (see Chapter 7), the fact that this behavior repeated in the summer of 2007 in the specific context of *e-learning* games played by a larger sample of children showed that the simple splash screen design could potentially leverage social relationships as a powerful motivator of learning.

8.4.4 Frogger and Floored Redesign

We redesigned Frogger and Floored based on the above lessons. For Frogger, we changed its receptive phase by situating it within the classroom setting, such that its new receptive phase resembled the receptive phase in Floored (Figure 8.6, right image) with a teacher supplying the English word for the image in her speech bubble.

Similarly, we modified Frogger's practice phase and situated it in the classroom setting too. However, to keep its practice phase distinguishable from its receptive phase, we designed the practice phase to have a significantly different background (Figure 8.8). Specifically, the practice phase contained a blackboard which displayed the pictures for the three vocabulary words belonging to the given receptive-practice-activation cycle. There was an icon of the participants' teacher pointing to one of the three pictures with a stick, such that the user could move this icon to the left and right to point to a different picture by pressing the left and right arrow buttons. Whenever the player moves the teacher to point at a new picture, the game would say its corresponding English word. With the exception of the picture that is currently selected, the other two pictures would be dimmed out.

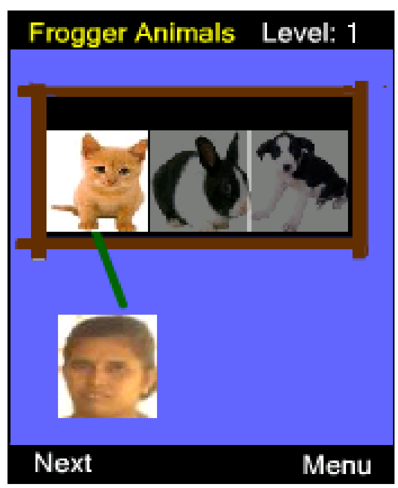

Figure 8.8. The practice phase in the second iteration of Frogger.

Since participants could not recognize their teacher from her caricature in earlier designs, we replaced her caricature with her photograph in the receptive phases of the Frogger and Floored games. For the same reason, we made use of her photograph for the above pointer icon in Frogger's practice phase, as opposed to using her caricature. Lastly, we modified the activation phases in both Frogger and Floored, such that when the game says the word that the player is tested on, the teacher's photograph would appear at the same moment to create the impression that she was quizzing the player.

Like the third iteration of the parrot game, we implemented user-interface controls in Frogger and Floored for flow control through the left and right soft keys. Among the navigational shortcuts that we could think of, we expected that players will appreciate having the freedom to transition to the next phase most. Hence, we made this shortcut highly accessible by associating it with the left soft key. Conversely, shortcuts such as exiting the game or reverting to the previous phase would frustrate the player if she were to press their buttons by mistake and had to repeat the game. As such, we made them less accessible by requiring the player to reach them via a menu that would appear when the right soft key is pressed.

8.4.5 Experience with Second Iteration

We tested the second iteration of Frogger and Floored with the same 47 students between grades 2 and 5. With the new receptive and practice phases that were situated in a classroom setting, it seemed that our players understood that they were supposed to focus on learning the words covered in these phases. Once the phases became intuitive, players figured out for themselves what they needed to do in both phases and attempted to learn the curriculum. In comparison, when participants were interacting with the prior iteration of the receptive and practice phases, it did not appear to occur to them that the game was attempting to convey the English words for the animals on the screen through the audio playback.

More important, our concerns that participants may not be able to tell the phases apart once there were three phases turned out to be unfounded. We had earlier believed that incorporating three phases in Frogger would not confuse the users if the phases were distinctively different, and the observations in the above paragraph supported this earlier design decision to retain three separate phases in Frogger. We also note that it was easier for us to think of a design for Floored (vs. Frogger) in which the practice and activation phases were combined seamlessly. The design implication is that when it is not easy to come up with a fantasy setting that integrates any two consecutive phases naturally and in a meaningful manner, keeping the overall design simple by having the receptive-practice-activation cycle span three entirely separate phases is not such a bad idea despite the lack of so-called creativity.

On a related note, participants told us that they were now able to recognize their teacher in the receptive and practice phases from her photograph. Since the teacher is arguably the main source of formal knowledge and instruction in this cultural context, where school is usually "teacher-centered" as opposed to centered on the learner's needs and prior knowledge, we argue that her photo went hand in hand with the background images in both phases to reinforce the classroom atmosphere that we wanted to establish in the gaming experience. Similarly, her photo may have helped to establish expectations on the part of participants that the teacher character was there in the game to teach them. Users appeared to pay the same amount of attention to the audio playback before and after the caricature was replaced with the photo. But they were visibly excited to see their teacher's face in the game. Some of them even showed us the cellphones and pointed out their teacher's photo to us.

On the flip side, many participants mispronounced the audio recordings when they repeated words aloud after the teacher's voice. For example, children pronounced "rabbit," "dog" and "frog" as "dabbit," "odd" and "frock" respectively. The most likely cause was that the teacher at the rural school could not pronounce some English words well, as we have observed. There were various design options that we could take. But more importantly, the interrelated issues of authority and the language model surfaced in our design considerations. We have incorporated the teacher into both the receptive and practice phases so as to lend a stronger air of authority to the designs for both phases and to encourage players to focus on the curriculum more seriously. More importantly, since rural participants had very little access to English language learning resources in their everyday environments, it was imperative that our games serve as a good language model for learners. We could train the above teacher to improve on her pronunciation, but it would be some time before we could obtain high-quality audio recordings. In our view, a more feasible approach would be to use the voice of a more qualified English speaker, while retaining the photo of the teacher at the rural school. This option would work if users could not identify that the voice in the audio recordings no longer matched that of the teacher in the photo. We needed to do more field-testing to investigate the feasibility of this idea.

The facilitator controls were used to adapt both games to the English baseline of each player. In particular, the left soft key for skipping to the next phase was used increasingly as the learners became more familiar with the targeted vocabulary words through repeated gameplay, and simply wanted to advance to the activation phases to test their vocabulary knowledge. In contrast, the shortcut for returning to the previous phase was seldom used, unless the facilitator or player advanced to the next phase by mistake. On the whole, the facilitators did not comment on these controls, which implied that these controls met their navigational shortcut needs and that they were not cumbersome to use.

In terms of quantitative measures, participants scored an average of 1.96 out of 5 on the pre-test and 3.85 out of 5 on the post-test. The effect size was 1.33 and post-test gains were significant ($p < 0.001$, $\sigma = 1.42$) on a one-tailed t-test.

8.4.6 Final Designs

We decided to apply the above lessons and ideas to Frogger and Floored, but with a new curriculum for each game, to investigate the extent to which we could replicate the above results. In the third iteration of Frogger, we modified the activation phase in order to target vocabulary words for everyday vehicles. Specifically, instead of helping the correct animal cross the road, the goal in the game was to help the animal get onto the vehicle designated by the teacher character in the game. Two other major changes that we made were to implement the same facilitator user-interface controls and the "You Win!" splash screen that were found in Floored.

For Floored, we changed its curriculum from vocabulary of everyday animals to lexical verbs such as climb, jump and run. To improve the quality of the pronunciations for the verbs, we did the audio recordings using the voice of a team member who spoke excellent English with an Indian accent. We did not make other substantial modifications to either Frogger and Floored since we did not observe any further outstanding design problems to be addressed.

8.4.7 Experience with Third Iteration

Since there were no more students at the rural school who had not yet played the earlier two iterations of Frogger, we tested the third iteration with all 47 students between grades 2 to 5. Similarly, we evaluated the latest iteration of the Floored game with 16 students from grade 3.

Unfortunately, we learned during the pre-test for Frogger that participants could name the English words for the vehicles in their everyday lives. We did not anticipate this issue when we designed the third iteration of Frogger, but on further reflection, we did not find it surprising that these English words had been adopted into their native language as loan words. In terms of design implications, we noticed that children made use of the facilitator controls to skip directly to the activation phase. In fact, they used these controls in this iteration of Frogger more frequently than in earlier versions of Frogger or Floored. This observation was highly consistent with the earlier observation that players skipped ahead to the next phase more frequently as they became more familiar with the targeted vocabulary words through continued play, thus suggesting a strong positive correlation between familiarity with the curriculum and frequency of using the navigation shortcuts.

In earlier iterations of Floored, we observed that learners mispronounced many of the words for animals. In fact, at least two participants were not able to follow the audio playback at all. In comparison, learners were able to repeat the words aloud with a more accurate pronunciation after we switched the voiceover from the teacher's voice to the voice of the above team member. We learned from the pre-test that the students did not know many of the verbs covered in Floored. Yet, we observed that learners had difficulty pronouncing only one verb ("cook"), and that was presumably because they had confused this word with a similar word ("cock") from the previous day's English learning activity. This result supported our design assumption that voiceovers from a more fluent English speaker would make it easier for rural learners to acquire the spoken English targeted in the game.

Quantitatively, the 16 learners from class 3 who played Floored scored an average of 3.50 out of 8 on the pre-test and 7.06 out of 8 on the post-test. The effect size was 2.73 and post-test gains were statistically significant ($p < 0.001$, $\sigma = 1.30$) on a one-tailed t-test. More importantly, none of the children pointed out that the voiceovers in the latest iteration of Floored did not belong to their teacher. It seemed that they did not notice that the voice in the recordings no longer matched that of their teacher's. Finally, we observed that participants reacted enthusiastically to the "You Win!" splash screen in Frogger as they did when playing Floored.

8.5 Summary

The receptive-activation format is a conceptual model found in several successful commercial language learning applications in developed country markets. In this Chapter, we have described a design case study in which three language learning games based on this conceptual model have collectively underwent nine rounds of iterations as a result of our field-testing with children from the urban slums and rural areas of India. We learned that the learning goals in our e-learning games were more obvious to rural learners when we maintained a distinction between pleasure and education in game settings, such that the learning phases were situated within a setting that is associated with learning, such as in the form of the teacher and blackboard. These cues contained the connotations of an educational setting and appeared to help users understand that they needed to focus their attention on the curriculum, as opposed to playing the game for the sake of satisfaction. The fantasy setting is a heuristic which is often recommended to enhance the enjoyment value of games (Fullerton 2008; Malone 1980), and our observations suggested that game settings which ensured a good balance between pleasure and learning are necessary in the case of games that target education.

However, even though an authority figure such as the teacher can be incorporated into game designs to scaffold the learning process, more research is needed to understand the cues that are required in the case of rural children who have less familiarity with the school setting. We also found that the games needed to be an accurate language model in contexts such as Indian villages where learners do not have other forms of access to good English speakers. More importantly, the language model can be kept separate from the teacher and other elements in the classroom setting that constitute the cue for the learning phases of the games, without the distinction being noticeable to players.

Finally, we discussed how the receptive-activation model should be expanded to include a practice phase that gave rural learners repeated exposure to the curriculum in a self-paced manner. The reader will also observe that these phases are related to the pre-task activities in task-based language teaching (Chapter 4), such that the receptive phase aims to introduce the vocabulary and other linguistic items for the task proper while the practice phase provides controlled practice on these linguistic items. We also note that the winner splash screen leverages the social aspects of motivation, which is highly crucial for keeping children engaged with the e-learning games. Finally, although we presented our key findings in the context of a project on language acquisition, we believe that our lessons are relevant to other educational domains in the developing world.

9 Traditional Village Games

In Chapter 7, we describe how 8 mobile games were tested with rural children in India. A few of the games were designed by amateurs, while others were designed by professional game designers. Unfortunately, independent of who designed those games, when these games were taken as a whole, we reported how users did not necessarily find them intuitive, exciting or free from playability problems. It appeared that rural children have relatively little exposure to these videogames, whose designs were influenced by Western cultural traditions that their designers unconsciously incorporated into the game design processes. Existing videogames thus arguably fail to match the understandings or expectations that rural children have about games.

On the other hand, rural children are more familiar with the traditional games that they play in their villages. It is timely to study the latter games so that we can address these questions: what can we learn about and from traditional games so that we can better design digital games that rural children can relate to more readily? What are some of the game mechanics in traditional games that we can incorporate into game designs for this user group? How do village games differ from contemporary videogames? In addition to iterative design and field-testing – which is undoubtedly important and which we have documented in Chapter 8 – what are some emerging principles that we can identify from traditional games to inform initial game designs for rural children?

In the remainder of this section, we first summarize how rural children in India interacted with 4 Western-style mobile games, and the playability problems which they encountered. Second, we report on 28 traditional games that rural children play in India, based on contextual interviews with 3 rural communities in North and South India. Third, we analyze the elements in this sample of games and identify the game mechanics found in traditional village games. We then compare these elements against recurring patterns in existing videogames to show what makes village games different from contemporary

Western videogames. Finally, we describe a videogame that we designed based on these lessons, and described preliminary results which show how rural children found it to be more intuitive and engaging.

9.1 Earlier Experiences

In this section, we describe the observations from piloting 4 Western mobile games with three communities of rural children in both North and South India. Our goal in this section is not to be comprehensive, but to establish a context for the rest of this Chapter by sharing a few concrete observations.

9.1.1 Games

The games (Figure 9.1) that we had earlier designed with the goal of investigating gameplay design for rural children are:

Crocodile Rescue – the player rescues the drowning boy by moving his boat to the boy within the time limit. The player is challenged with obstacles in the form of crocodiles, which he baits out of his path using chunks of meat, in order to clear a path to the boy.

Dancer – the player seeks to maximize his score within the time limit by throwing tomatoes to hit as many dancers as possible. The green- (red-) colored members in the audience give (take) a tomato to (from) the player when he moves to them. Members in the audience switch between green and red as the game proceeds.

Train Tracks – where the player extends the railroad from the top-left hand corner of the game world to the bottom-right hand corner, before the train derails at the end of an unfinished railway track.

Frogger – the player is quizzed on various English words for animals. His goal is to select the correct animal and help it to cross the road without getting run over.

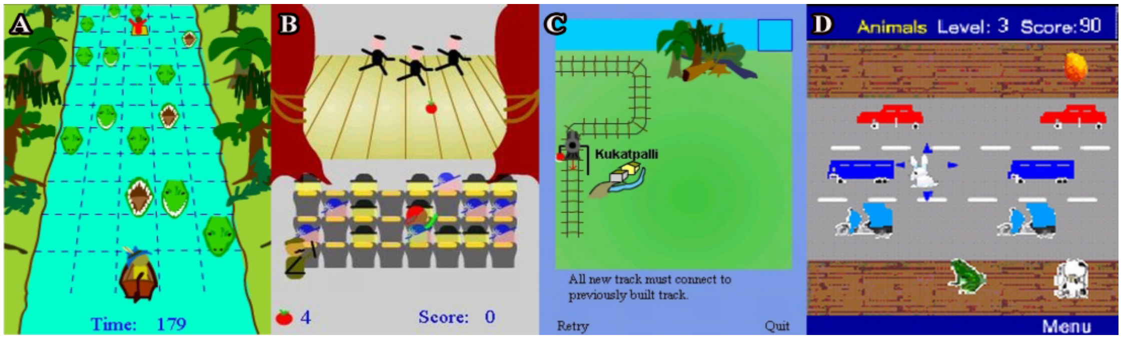

Figure 9.1: The games that we designed for rural children in India, namely, (A) Crocodile Rescue, (B) Dancer, (C) Train Tracks and (D) Frogger.

The game settings for Crocodile Rescue, Dancer and Train Tracks were chosen such that rural children could relate to them readily. We did not have any principles for designing games for rural children at that time, and instead borrowed features from successful casual games in Western societies. More details on our design process are given in Chapters 7 and 8.

The version of Frogger that we initially tested with the first community does not target English. But after observing that it tied with another game as the most popular game, we adapted it for ESL learning and tested it with all three rural communities. A plausible explanation for its popularity is that its goal (i.e. to cross the road) is situated in a familiar everyday setting – it is challenging to cross roads in India due to the lack of traffic lights and pedestrian crossings. In contrast, most games that we tested in the first community had goals that did not resonate with players, e.g. having to collect enough coins before being allowed to fight the boss.

9.1.2 Three Participant Communities

We tested the above games in the following communities:

MYSORE – a rural public school in Mysore (Karnataka, South India) in a remote village that is seldom visited by outsiders. We tested Crocodile Rescue, Train Tracks and Dancer with 24 students in grades 1 to 2 (ages 6 to 7) and Frogger with all 47 students in grades 2 to 5 (ages approximately 7 to 13). 20 of the 47 (43%)

children, especially older ones, have played the "snake" game on cellphones before. This game involves moving a snake around to eat various food while avoiding touching itself or the screen's boundaries.

LUCKNOWPRIVATE – a private rural school in Lucknow (Uttar Pradesh, North India). Parents can afford the fees since they are landowners (vs. daily-wage farm laborers) who grow cash crops. This school has a computing center backed up by a power generator. We tested each game with 11 children in grades 2 to 3 (ages 8 to 10). Due to higher parental income, 2 of the children own low-cost, portable gaming devices. Similarly, unlike participants from the other 2 communities who may have only played "snake" on the cellphone, 4 students in this group have played between 1 and 3 other mobile games on cellphones.

LUCKNOWPUBLIC – children attending public schools near LUCKNOWPRIVATE because their parents cannot afford to pay private fees. They participated in our research study after hearing about the opportunity to learn English in our project from neighbors who attend LUCKNOWPRIVATE. We tested each game with 10 children in grades 2 to 7 (ages 7 to 12). 3 of the 10 children have only played the "snake" game on cellphones previously.

As the wealthiest rural community among the three groups, LUCKNOWPRIVATE provides a point of contrast which helps us to contextualize our observations better. In particular, even though it is not representative of average rural children, it allows us to observe how rural children's interaction with videogames varies with prior videogaming experience.

9.1.3 Key Observations

For each game, we introduced it by demonstrating it to the children in small groups. After which, each child receives a cellphone to play the game for about 1½ hours. We looked for playability problems and assigned severity ratings to the critical incidents.

After the sessions, we asked participants what they understood and perceived about the game. The most salient observations are:

Goals – almost every player immediately understood the goal of Frogger. On the other hand, more than 50% of the participants in all three groups did not appear to find the process of exploring the Train Tracks game world to lay the railroad to be evident. Similarly, they seemed to be confused with the “bait” subgoal in Crocodile Rescue and did not attempt to achieve it in many cases.

Player actions – despite grasping the goal of Frogger, 33 out of 47 children in MYSORE attempted to find a straight route by waiting until there is no traffic on the next lane. They did not move sideways (i.e. left and right) along the current lane to slip into the gaps between vehicles on the next lane. Moreover, a child commented that his tortoise cannot cross the road because it is a slow-moving animal. In comparison, participants in LUCKNOWPRIVATE readily understood the need to move sideways although people usually do not cross roads in this way. It appears that the latter users could take advantage of this action, which is usually found in videogames, due to their prior exposure to electronic games. Children who lacked this exposure seem to relate videogames to their daily experiences instead.

Difficulty level – just like the impassable terrains in Train Tracks, children playing Crocodile Rescue appear to find it frustrating to navigate a path to the destination square, especially when the number of obstacles (e.g. crocodiles) increases in the more difficult levels. More than 50% of the children in all 3 groups indicated that they disliked the crocodiles that obstructed their movements. Part of the problem was that participants in all 3 groups did not understand the “bait” action by the time their sessions ended. Some also seemed to struggle to determine if there was a clear path to the boy. Our observations for Dancer were consistent: participants appeared to become more stressed when there were more dancers to hit in the more difficult levels.

Resource management – in Dancer, although users threw tomatoes enthusiastically at the dancers, only children in LUCKNOWPRIVATE could explain that they can gain more tomatoes by going to the green-colored members in the audience. In the other two communities, all but 3 users did not appear to know what to do once they run out of tomatoes, or that red-colored members in the audience would take their tomatoes, despite repeated explanations.

Score-keeping – most of the children in LUCKNOWPUBLIC and MYSORE did not pay attention to their scores during gameplay. For instance, when asked about why they liked Dancer, none of the MYSORE participants cited the score-keeping mechanism as an appealing factor, despite being excited each time they complete a game level. In contrast, children at LUCKNOWPRIVATE were very competitive and compared their scores with their peers after finishing each game. A plausible reason is that test scores matter a lot in this school, unlike most government rural schools. We observed similar behaviors with Frogger.

9.2 Traditional Games

Clearly, designing a game to be culturally appropriate goes beyond situating its goals in an everyday setting. Frogger is instructive in showing that there are other components in a game such as player actions which must also be culturally meaningful. More generally, the above games did not entirely match rural children's understanding of games. In order to design videogames that our target users can better relate to, it is timely to understand the traditional games that they play in their villages. We expect that digital games which incorporate familiar game mechanics from everyday village games will be more culturally meaningful to rural children.

9.2.1 Data Collection

We conducted contextual interviews with 87 children in the above 3 rural communities over 12 days, in which we asked participants to recall the everyday games that they love and to play these games for us to videotape. Participants did not appear to be self-conscious about us observing them since we had spent enough time interacting with them previously. Our contextual interviews were adapted for children (Druin 1999), in that the brighter children helped us to understand the games that their peers – who were not always able to explain the rules of their games coherently – played in front of us. We had the chance to observe both outdoor and indoor games because the interviews took place during the monsoon.

The games appear to be a marker of social identity, in that children in the Lucknow studies were initially reluctant to tell us about their village games and shifted the interviews toward games that are more urbanized. It was only after we told them about some traditional games in South India and pointed out what we liked about them, that our participants became less conscious about class differences and felt more encouraged to show us more of those games that they play with other children in their social classes.

We learn that some games have variations for different ages or genders, but we do not count these variations as separate games. The variations are significant, however, for e.g., in enabling us to understand how a game's difficulty could be adjusted to suit older players. Children from wealthier rural families also know about games that require relatively more expensive equipment (e.g. Chinese Checkers). However, we did not follow up on these games in our interviews because the majority of our respondents do not play or know them.

9.2.2 Games Overview

We observed a total of 23 outdoor and 5 indoor games, in addition to 4 variations on 3 outdoor games. In Figure 9.2, we show 4 games that are more exotic to a Western audience. 17 outdoor games belong to the family of “tag” games, in which there is generally at least one player designated “it” who has to “tag” players in the opposing team by touching them, either with a hand or an object. In particular, 2 and 3 “tag” games belong to the “cops and robbers” and “hide and seek” sub-families respectively. The 6 outdoor games that do not belong to the “tag” family include tug-of-war, kite flying, marbles, hopscotch and the spinning top.

Figure 9.2: Some of the traditional village games that children in India play, namely, (A) Gucchi Garam, (B) Siya Satkana, (C) Giti Phod and (D) Halla Guli Mane. In (A), a player stands in the safety of his circle and deflects a ball before it hits him. In (B), a player reaches to pick up the stick while avoiding being touched by the “it” player. In (C), a team attempts to rebuild the rocks into a heap while avoiding a ball thrown by the opposing team. In (D), players aim to move all the tamarind seeds into one crevice on the tray.

The indoor games can be generally classified as “tabletop” games, although 3 of them (Halla Guli Mane, Chor Sipahi and Trump Cards) can be also played while sitting on the floor, i.e. only 2 of them (Pen Pick and Pen Fight) require a table. In Halla Guli Mane (Figure 9.2), there is a tray with 14 crevices. Players take turn to empty one of the crevices that belong to them by distributing all the tamarind seeds in it among other crevices. The game ends when all the seeds on the tray are in one player’s crevices. In Chor Sipahi, each of the 4 players randomly chooses a folded slip of paper, after which one of them tries to guess who has the two slips with “chor” (i.e. thief) and “sipahi” (i.e. soldier) written on them. In Trump Cards, a pack of cards showing the attributes of World Wrestling Federation wrestlers is distributed evenly among everyone. Players take turns to compare an attribute on their topmost cards and the winner takes all the topmost cards. The game ends when one person acquires every card.

In Pen Pick, 6 to 8 pens are randomly scattered in a heap on a table, and each player has to take a pen without displacing other pens. In the Pen Fight game, there are 2 players, each of whom has a pen on a table. The goal is to flick one's pen as a "striker" to dislodge the opponent's pen from the table, while ensuring that one's striker does not overshoot and fall off the table.

9.2.3 Consistency Checks

We do not describe the above games in greater detail due to space constraints. The games in the "tag" family (including Kabbadi and Kho-Kho), as well as hopscotch, have their descriptions on Wikipedia. Similarly, websites such as India Visit Information describe ancient Indian games. We refer the reader to these sources. As a consistency check, however, we notice that the traditional games we observed differ from their online descriptions, which appear to be urban variations on the same games. It seems that the variations played in villages involve simpler rules and less equipment.

As another consistency check, we do not observe significant differences between the traditional games in North India vis-à-vis South India. However, Southern traditional games typically involve players performing a skit before the game proper begins, so as to situate the gameplay in a story, whereas none of the Northern games do. We do not notice other major regional differences. More importantly, 14 of the 28 games can be classified into 6 sub-families (e.g. "marbles," "cops and robbers"), such that each sub-family is represented by at least one game from each region. In other words, more than 6 of the traditional games have variations that exist across geographic regions.

9.3 Elements of Traditional Games

To understand what traditional games are made up of, we examined each of the above 28 games and identified the elements that belonged to the following categories:

Players – what are the various forms of interaction among player characters? How many players are there, and is this number variable or fixed? What roles do players take, and are these roles uniform or different across players?

Resources – what are the resources and objects in the game that players can leverage to attain their goals?

Goals – what does each player have to accomplish in the game? How does the game progress from one round to the next? What factor(s) determine that the current round is the last? How is the overall winner determined after the last round?

Actions – what are the actions that players in various roles may perform, possibly with the resources in the game, to attain his goals and sub-goals?

Rules – what are the constraints on a player's actions and their outcomes that render it challenging for him to attain his goals? How do these constraints relate to the resources and objects in the game?

The above categories for the formal elements comprising a game are provided in standard texts on games such as Fullerton (2008) or Björk and Holopainen (2005). Since our objective is to understand village games in order to inform the design of videogames for rural children, as opposed to studying traditional games for their own sake, when we analyzed traditional games for their elements, we abstracted away those details that do not transfer readily to digital media (e.g. hopping on only one leg in hopscotch).

9.3.1 Players

The village games involved players interacting with one another in these ways:

Pair-Wise Competition – two players compete with each other, i.e. each player makes up one of the two opposing sides. If the game allows the number of players to increase beyond two, the pattern of interaction becomes that of multilateral competition or team competition, which we discuss next.

Multilateral Competition – three or more players compete with one another, i.e. the number of opposing sides is equivalent to the total number of players.

Team Competition – players form two or more groups that compete with one another. In most traditional games that are team-based, teams have the characteristic of being small (two or three players in each team) or medium in size (up to six players per team). An exception is "tug-of-war" in which larger teams can be supported. Another characteristic of most team-based traditional games is that children try to balance the number of members across groups, out of fairness. An exception, however, is a variation of "cops and robbers" in which two players team up as cops while remaining players assume the role of robbers. Some traditional games involve only two teams, while others may allow more than two. In team-based traditional games, every team constitutes an opposing side, i.e. there is no cooperation between teams.

Unilateral Competition – one player competes with everyone else (who comprise two or more players), i.e. there are two opposing sides, out of which an opposing side has disproportionately more members than the other side.

The above details are relevant to videogame design because the number of players and their attributes (e.g. scores and positions) are part of the overall game state, as well as influence other aspects of the game design such as the rules. Most of the games involve a variable number of players, with only Chor Sipahi, Pen Fight and the Southern variation of "marbles" in the minority that involve a fixed number of players. A minimum of 3 to 6 players in total is required in most team games to ensure that there are enough members in every team. Conversely, most of the traditional games – especially those that involve unilateral competition – do not impose an upper-bound on the number of players, which is more often determined by practical resource constraints such as the size of the playing field or amount of physical objects available as game equipment.

In terms of the states that a player can be in, the most common is "active." A player can also be "at rest" and waiting to become an active player, e.g. an active player swaps his state with an inactive player, who is catching his breath, by touching the latter. A player can be caught and "imprisoned," i.e. frozen on his current spot and forbidden from moving to another location until he is rescued by his teammates. Finally, players can be "out" of the game and waiting to resume gameplay in the next round or session. In contrast to players in other states, an "active" player may possibly have a significantly greater range of actions that he is allowed to perform.

In addition to player states, players may also be in sub-states that are orthogonal to the above states. For example, a pursuing player may be "hindered" by his opponents who are grabbing him, so that he cannot move as fast as he normally would. It is also possible for a player to be "hidden" from sight when playing a "tag" game which belongs to the "hide-and-seek" sub-family. Finally, a player may be "invulnerable" to capture. Player sub-states can be transitive (e.g. a player who is touching an invulnerable player becomes invulnerable himself).

A player may assume one of the following roles: play a *primary* role such as steer the direction of a flying kite, play a *supporting* role such as holding the rest of the thread for the kite, play an *offensive* role, or take a *defensive* role. In unilateral games, the player who is competing with everyone else always takes on a role which is different from other players. In traditional games that involve other types of player interactions, players either have roles that are uniform – that is, every player is striving toward the same goal – or different from other players.

These "either or" distinctions become more complex in the case of team games, in which players may have similar or different roles from their team mates (such as a cop deciding between the primary role of pursuing other robbers or playing a supporting role of standing guard over captured robbers). Similarly, players in a team may have roles that are either similar or different from players in the opposing teams (in some traditional games, all teams are either on the offense or defense only; in other games, there is at least an offensive team and at least another defensive team). Finally, in a team, some members can play a primary role (such as pursuing their opponents) while other members play a supporting role (such as taking a rest, so that they can exchange roles with the pursuer(s) when fully rested). On the other hand, there are no players in a supporting role in games that involve pair-wise, multilateral or unilateral competition.

Similarly, this distinction breaks down in those traditional games where players can adopt both an offensive and defensive role. For instance, a player who is attempting to cut the threads that tether other flying kites to their owners has to ensure that no other player cuts his kite loose. Similarly, a player who is pursuing his opponents may have to be careful not to be hindered from returning to his territory. Moreover, roles do not have to be static across consecutive rounds. For instance, players or teams may take turns to assume each role. Alternatively, players or teams swap roles as a result of a victory or penalty, such as when someone or a team defeats another player or team. Finally, in the case of primary vis-à-vis supporting roles, players in the same team have the flexibility to swap roles as the situation demands.

9.3.2 Resources

Broadly, the resources in traditional games that advance and thwart players in their actions to achieve their goals fall into three broad categories: (i) attributes of player characters, (ii) objects and their attributes, as well as (iii) territories that the playing field is divided into. A player character in traditional games has very few attributes, which can be his score, x and y coordinates, current direction that he is facing, and the objects in the game that he possesses.

The objects that children use in playing their traditional games are the everyday objects that are readily available in villages. They include: the ball, stone, rocks, sticks, marbles, pens, paper slips, cards, handkerchief, tray with holes, a top with a string to spin it, and a kite with an attached string. Again, like player characters, objects in traditional games have very little attributes. Two properties of game objects that are almost universal to all traditional games are its owner, as well as its x, y and z coordinates. Other attributes relate to the configuration of the rocks (to what extent are they stacked up in a heap with smaller rocks on top of larger ones?, in the Giti Phod "tag" game), configuration of the seeds in the Halla Guli Mane tray (who owns each hole in the tray?, how many seeds are there in each hole?), velocity and/or acceleration (is the stop spinning or has it stopped?), and the condition of the kite's string (has it been cut by an opposing team?). In short, the attributes for most objects relate to their physical properties. However, objects can also have non-physical properties. For example, cards and strips of paper have information written or printed on them (e.g. Trump Cards show wrestler attributes such as height).

Lastly, players in traditional village games can be assigned resources in the form of territories, which could be an area on the playing field that adopts a shape such as a circle. But a territory does not have to apply literally to an area on the field. For instance, in Halla Guli Mane, players claim ownership over various holes in the tray. Similarly, a territory that players can flee to for refuge does not need to be on level ground, and can instead be a tree that they climb up to be safe from capture.

9.3.3 Goals

Some traditional games revolve around a single round, while other village games comprise a series of consecutive rounds. Like videogames, the player's objectives in each round in a traditional game are organized according to a hierarchy of goals and sub-goals. In the traditional games that we have observed, the goals for a player or team tasked with an offensive role are:

Eliminate – defeat all opponents from the game (e.g. by keeping one's top spinning for the longest duration, cutting the thread that keep every opponent's kites within his reach while in the air, hitting every opponent with an object or projectile, or capturing all opponents).

Overcome – defeat one opposing player (e.g. by hitting an opponent successfully with a ball, capturing her, hindering her from returning to her territory within a time limit, dropping a handkerchief behind her without her awareness, or guessing who holds a given slip of paper). This goal differs from the above elimination goal in that the player on the offense only needs to defeat one – and not every – player. In some games, a player on the offense aims to overcome someone before an opponent completes his goal. Similarly, in some traditional games, a player (or team) who successfully overcomes an opponent gets to swap roles with him (or the opposing team) in the next round.

Furthermore, the following goals are applicable to players in both offensive and defensive roles. In other words, we have not encountered goals that belong exclusively to players on the defense:

Survival – the player can aim to be the last player to be defeated, or conversely, not to be the first player to be defeated (e.g. by dodging all oncoming projectiles so that one is not the first player to be hit). This goal can involve territories (e.g. the player can reach the refuge of a specific location or tree). In team games, the team can also attempt to prolong the length of the current round (e.g. evading capture for as long as possible) so as to gain a higher score than the opposing side.

Acquisition – this goal is applicable to objects, territories, and points. For instance, the player could be required to acquire one or more objects (e.g. by scoring a hit on all marbles using a "striker" marble; picking up a stick which is lying on the ground, possibly while evading capture at the same time; picking up a pen from a heap of pens such that no other pens are moved; or guessing the Wrestler's attribute on one's card that has the highest value among the values for the same attribute among other players' cards). Similarly, the player can gain possession of one or more territories (e.g. entering an opponent's turf when it is left unguarded). Lastly, the child could aim to be the first player to acquire a certain score.

Object Manipulation – keep an object in the game in a certain state for the longest time as compared to other players (e.g. keep a top spinning), or arrange objects into a given configuration before a team mate is defeated (e.g. team members have to arrange seven stones into a heap such that each stone is atop a larger stone, before one of them is hit by a ball thrown by the opposing team).

Once a player achieves his goals for the current round through a successful action, one or more of the following happens:

Point Award – the successful player is awarded one or more points. It is also possible for his or her team mates to receive one or more points at this time.

Round Resumes – the current round resumes with the current assignment of player roles. In games with this characteristic, the current round ends only when the player or team on the offense fails to maintain his (its) winning streak.

Round Ends – the current round ends. In most traditional games like those in the "tag" game, out of fairness, the successful player (and/or team) swaps roles with the defeated player (and/or team) for the next round, especially if the successful player is taking an offensive role. If there are more than one defeated players (i.e. in a unilateral game), the successful player exchanges roles with the first opponent he defeats. Furthermore, as an alternative to exchanging roles, roles for the next round can be assigned randomly. In the case of card games, the offensive role for the next round goes to the player who holds the card with the highest value in the current round.

Those traditional games that comprise a single round end as soon as the goal state is satisfied through the direct actions of players in offensive and/or defensive roles. For those traditional games that span multiple rounds, end conditions are more complex, and may be one of the following:

Predetermined Rounds – a predetermined number of rounds have elapsed. This number of rounds could be an arbitrary, or in the case of those traditional games that promote fairness, it could equal the total number of players (teams), i.e. roles are rotated among the players (teams) such that every player (team) gets the chance to perform every role.

Goal Satisfaction – this condition differs from predetermined rounds, in that the game ends with the current round only when the goal has been satisfied (e.g. when opponents are eliminated or overcome, or when a certain number of points or other resources are accumulated). Otherwise, players or teams exchange roles with their opponents and proceed to the next round.

Equal Chance – the game ends when each player has a chance to assume every role.

Mutual Consent – some traditional games have no explicit end conditions, and conclude upon the mutual consent of every player.

Traditional Indian games that end through mutual consent do not have the notion of an overall winner, since players mutually decide to stop when they are tired of playing. (However, traditional games that end with mutual consent span multiple rounds. Hence, despite the lack of an overall winner by the time the rounds have ended, there is still the notion of successful and unsuccessful players for each round, since this metric of success is used to determine when a round is over.) In traditional games with end conditions other than mutual consent, the winner is determined using one of the following criteria:

Highest Score – the winner is the player or team with the highest score.

Shortest Time – the winner is the player or team who took the shortest time to attain the goal state (e.g. cut all flying kites from the threads that hold them, or eliminate all opponents).

First Mover – the winner is the first player or team to attain the goal state (e.g. eliminate all opponents, or attain a certain number of points).

Longest Time – the winner is the player who maximizes the duration in which the game is kept in a certain state (e.g. keep the top spinning for the longest time).

Most Possessions – the winner is the player or team who possesses the highest quantity of a certain object or territory.

9.3.4 Actions

The direct actions that a player in an offensive role can take to defeat a player or team are one or more of the following:

Pursue – pursues an opponent by moving in his direction. This action may involve the player venturing into the opponent's territory.

Capture – capture an opponent by grabbing and/or touching him, after which the latter's state switches from "active" to "out of the game" or "imprisoned."

Seek Out – move around and/or look for opponents who are in hiding.

Utter – utter a word or phrase – which can be an opponent's name – correctly and at the appropriate time. This utterance can be a part of another direct action (e.g. an utterance that has to be correct for the accompanying action to succeed), or a substitute for another action (e.g. an utterance that is said in order to capture a hidden opponent who has just been spotted, as a substitute for having to move to where he is and capturing him physically).

Guard – guard an imprisoned opponent by standing nearby, so that his team mates would not dare to rescue him for fear of being captured.

Pick Object – successfully pick up one or more objects, such as a ball or stick from the ground, a pen from the table top, or a tamarind seed from a hole in a tray.

Drop Object – drop an object, such as a handkerchief behind an opponent without her knowing, or a tamarind seed into a hole in a tray.

Shoot Object – there are three variations on this action. First, the player can throw a projectile such as a ball at an opponent with the aim of hitting him. Second, the player flicks a projectile, such as a marble (or pen), at another marble (or pen), in the hope of hitting the marble (and dislodging the pen from the table). The projectile can also be thrown at other objects (e.g. a heap of stones in Giti Phod) with the purpose being to mess up their arrangement. Third, the player can throw a projectile like a stone into a territory (e.g. a region in a hopscotch pattern on the ground) with the aim of landing the projectile in the territory.

Guess – successfully guess which strip of paper one or more players are holding, or the attribute on one's Wrestler card that has a higher value than the same attribute value shown on other players' cards.

Cut Kite's Thread – fly a kite such that it successfully cuts the thread of another kite in the air.

Pull – pulls a member of the opposing team into one's territory (i.e. tug-of-war).

Occupy – successful occupy a territory, which could have belonged to an opponent who left it unguarded, or which currently does not have an owner (e.g. hopscotch). The player occupies the territory by moving directly into it, or by performing an elaborate ritual of hopping through every territory.

Evade Hindrance – the pursuing player moves away from his opponents so that he is not hindered from moving. The hinder action differs from the capture action in that the former is intended to achieve a sub-goal, but never a goal. That is, players on the defense do not hinder a pursuing player to attain the goal state directly, but as part of an indirect strategy to invalidate the latter's captures. A "hindered" player also differs from an "imprisoned" player, in that the former is still allowed to move, but with more hindrance (e.g. at a slower speed and/or having to break free of those opponents who hinder him).

Rest – rests to catch a breath, so that the player has more strength to chase his opponents later. In case the game supports both primary and supporting roles (e.g. one player is responsible for pursuing members in the opposing team while all his remaining team mates are at rest), the player who wishes to take a rest may, depending on the game rules, swap roles with a team mate anytime or after the current round.

Conversely, a player in a defensive role may perform one or more of these direct actions:

Evade Capture – evade capture by a pursuing opponent by moving away from him. This action can also involve objects (e.g. to steer one's kite so as to avoid its thread being cut by other kites).

Protect – protect oneself from capture. Depending on the specific game rules, a player is invulnerable to capture when he is touching a tree, touching a team mate who is already vulnerable, or holding hands with any team mate.

Hinder – hinder a pursuing opponent by grabbing him, for example, so as to prevent him from returning to his territory within the time limit.

Preempt – preempt an opponent who is in an offensive role. This action can be carried out in two ways. First, in games where the opponent is searching for players in hiding, the player can preempt him by touching him or hitting him on the back before he finds the player. Second, in games with one or more opponents who are pursuing players in the opposing team(s), the latter can preempt their opponents by successfully picking up an object like a stick that is lying on the ground, before the latter is captured. The preempt action is found in traditional games with the "round restart" rule (see the "Rules" sub-subsection below).

Rescue – rescue an opponent who is in the "imprisoned" state, and release him back into the game as a player in the "active" state, usually by touching him.

Dodge Projectile – dodges an incoming projectile (e.g. a ball) by moving out of its path.

Deflect – deflect an oncoming projectile (e.g. a ball) with a bat.

Sweep – sweep aside a fallen projectile (e.g. a ball lying on the ground). To introduce an element of uncertainty into the outcome, this action may require the player to move close enough to the projectile to sweep it farther away beyond the reach of an opponent (who can pick it up and throw it) When moving toward the projectile, the player is at risk of being captured.

Manipulate Object – manipulate one or more objects (e.g. to spin a top, arrange some stones into a heap).

Seek Refuge – seek refuge from a pursuing opponent by climbing onto a tree or moving into an area (e.g. the player's turf, or a neutral territory that does not belong to the pursuer) that is out of bounds to the pursuer.

Stake Ownership – stake ownership over one's turf by standing in it, so that an opponent cannot claim it by moving into this territory.

In terms of the score-keeping mechanism in traditional games, players may gain one point or a possession for performing a difficult action successfully (e.g. capturing an opponent or hitting an opponent with a projectile). When a player's score exceeds a threshold, he may gain the privilege to perform an action at that moment. This action is a one-time action (i.e. this action cannot be deferred) which he cannot otherwise carry out under normal circumstances in the game.

9.3.5 Rules

Games are sometimes argued to be fun when the attainment of the goals in them is uncertain (Fullerton 2008, Malone 1980). In the traditional games that we observed, the following rules make it harder for players in an offensive role to succeed in performing actions to achieve their goals:

Capture Restrictions – there is a limit on the number of opponents that a player may capture in each round or each time he ventures into the opponents' territory.

This rule has the effect of requiring the player to take multiple rounds or trips into enemy territory to capture every opponent. Likewise, the player could be prohibited from capturing an invulnerable opponent so long as the latter is protecting himself from capture.

Restricted Movement – this rule manifests itself in a few different ways. First, when the player is holding a projectile (e.g. a ball), he is not allowed to move another step. This restriction makes it more difficult for him to hit someone with the projectile when he is farther from his target. Second, this rule can be applied to situations where the player is pursuing an opponent. For example, he can be required to follow a certain path and/or prohibited from moving to his prior location. This restriction makes it harder for him to catch his opponent, who can predict the latter's next location with greater ease. This restriction can also be relaxed to result in less difficult gameplay by allowing the player to choose from more than one paths. Alternatively, the player could be required to start his pursuit from a specific location (which is presumably less accessible), or only after a countdown, so that his opponents gain a headstart in fleeing from him. Another instance of this rule is that the player could be required to touch a certain line in the opponent's turf before he is allowed to leave this territory. Third, the player can be forbidden from moving outside certain boundaries (or region), or prohibited from moving into certain territories or onto trees which only evading players can climb for refuge. This restriction makes it impossible to capture opponents once they reach the safety of certain regions. Fourth, a player who has to pick up an object may have to satisfy the constraint that he does not inadvertently move another object (e.g. when picking up a pen out of a heap in the Pen Pick game).

Stamina – the player finds it harder to perform certain actions such as pursue opponents once he becomes exhausted.

Specific Target – there are two ways in which this rule is reified. First, when throwing a projectile at an opponent, the thrower may be required to hit a specific body part, as opposed to anywhere on the opponent's body, for the hit to count as a successful one. Second, when throwing a projectile, the player does not get to choose the target, which is instead selected by the opponent. The latter is expected to deliberately choose a more difficult target (e.g. in Pen Pick, when the player is required to take a pen from a heap without shifting the other pens in the heap, it is the opponents who specify which pen to pick up).

Time Limit – the player is given a time limit to perform an action. For example, the player is given a time limit to return to his turf after pursuing his opponents in their territory. During this moment, the latter can hinder him from returning to his turf on time.

Penalty – there are four ways in which this rule could be applied. First, in a game where winners are determined based on scores, when a player (team) fails to achieve his (their) goals for the current round, one or more points can be awarded to the opponent (opposing team). This rule can involve physical boundaries. For instance, in Pen Fight, if the player flicks his pen such that it not only fails to dislodge the opponent's pen, but also falls off the table, his opponent gains a point. Similarly, a player can lose a point for each opponent who reaches the boundaries of a refuge. Second, in a game in which the goal is to acquire objects, the player can lose all of his possessions to the game (e.g. in a game of marbles, a player tries to acquire other marbles by flicking another marble to hit them successfully. If the "striker" marble and/or target marble ends up on a certain region on the ground, the player loses all the marbles that he has acquired so far). Third, in game where the player needs to make an utterance at an appropriate time, he is eliminated from the current round if he forgets to say this utterance or says it incorrectly. Fourth, in a multilateral game, at the start of each round, the player could be required to throw a projectile (e.g. a stone) and make sure it hits the target (e.g. an area on the ground). If he misses, he loses his turn and the game proceeds to the next round, in which the offensive role goes to the next player.

Undo – this rule manifests itself in two ways. First, there can be a direct action which an opponent can perform to undo the consequences of a player's action (e.g. a robber can rescue and thereby release an imprisoned robber whom the player had captured). This rule can be accompanied by an expiry time, after which the player's action can no longer be undone (e.g. after the player has dropped a handkerchief behind an opponent, the latter could catch him to invalidate his action, as long as he has not returned to his original position). Second, the player must remember to utter a keyword – and correctly – when performing an action. Otherwise, the action is undone.

Restart Round – an opponent can cause all the achievements that the player has strove to accomplish so far in the current round to be a wasted effort, by causing the current round to be restarted as a result of a direct action from the former (e.g. the opponent preempts the player who is searching for and catching children in hiding by touching or hitting him before he sees the opponent).

Conversely, players in a defensive role experience the following rules that make it more challenging for them to attain their goals through their actions:

Restricted Movement – the player is prohibited from moving out of certain boundaries (or region), or is required to run along a certain line or in a certain direction (e.g. clockwise in a circle), for example, when dodging oncoming projectiles or evading capture. This restriction gives him less physical space to maneuver.

Turf Obligations – the player is obligated to protect his territory by standing in it, so that an opponent cannot step into it to claim it. In the latter case, the player loses the current round and has to swap roles with the opponent. This turf obligation puts the player at a risk should he attempt an action with the potential payoff of disadvantaging an opponent.

Limited Sight – the player finds it difficult to see what an opponent is doing owing to limited sight, for instance, when the latter is performing actions (e.g. dropping a handkerchief) behind the former.

9.4 Differences from Western Videogames

The above section has covered the game elements present in traditional Indian games. However, in order to better understand how traditional games might be unique, it is equally important to examine how the game mechanics in Western-style videogames are *not* present in village games, or vice versa. We use the adjective "Western" broadly to include videogames that target affluent, urban consumers in non-Western regions, such as Japan and the urban middle-classes in India.

We perform our comparison between traditional Indian and Western games by contrasting the above elements in the former against Björk and Holopainen (2005), which is a huge compendium of 296 design patterns commonly observed in videogames. In other words, we assume that this sample of 296 patterns is comprehensive to the extent that it is representative of the features in contemporary games. A pattern is a "template" description (Alexander 1977) of a solution to a recurring problem that has been solved. In this way, there is no need to reinvent the wheel, and new game designs can leverage on earlier games that have been successful. Patterns also capture contextual information such as domain applicability and its rationale. A pattern may also capture tacit knowledge on the domain after having evolved through iterative design cycles.

The patterns in Björk and Holopainen (2005) cover numerous aspects of game design such as goals, rules, actions and heuristics for structuring games to be enjoyable. Although traditional games appear to overlap most with casual games or mobile games, we nevertheless consider patterns from other genres, under the rationale that confining our comparative analysis to a few genres is expected to prematurely narrow the space of future design possibilities that we seek to inform. For every pattern given in Björk and Holopainen (2005), we check if it is instantiated in a traditional game. For those patterns that are absent, we try to offer an explanation for why they are missing. For those patterns that are present, we examine if they are instantiated in a way that is qualitatively different in Western videogames.

In total, 74 out of the 296 patterns are completely absent from or instantiated differently in traditional Indian games. On closer examination, this proportion (25%) is lower than what we originally anticipated because both categories of games incorporate several of the same design patterns on structuring game designs for enjoyment. That is, in spite of noticeable differences between both categories of games, it appears that a fair proportion of the principles and heuristics for enjoyable play are fairly universal across some cultures. To avoid overburdening the reader, this section will focus on only the most salient differences.

9.4.1 Players

Even though both traditional games and videogames can involve multiple players, traditional games are never massively multiplayer, unlike some Western videogames, due to the (lower) population density in rural areas. Conversely, while Western games – both digital and non-digital – can be played with other players or in solitary, every traditional Indian game that we have observed is inherently multiplayer. By implication, there are other differences in the structure of interactions between players and the game, as well as between players themselves. Western games can be played alone, while traditional games require the presence of playmates. Moreover, some Western games, including selected (non-digital) boardgames, are designed such that humans play against the game system, either in isolation or in cooperation with one another. In contrast, Indian village games must be social experience in which humans only play against other humans.

More deeply, social engagement and interaction appear to be significantly simpler in traditional games. Cooperation is required among team members in traditional games, but teams and players do not form alliances (whether covert or overt) with the opposing side, even when there can be more than two opposing sides in some traditional games. In addition, we do not observe complex inter-team interactions such as trading, negotiating or reciprocating obligations in Indian village games. More broadly, team membership is determined before the game starts and remains stable for the duration of the game. On this note, friendly-fire or betrayal is not possible since traditional games do not define actions that make it possible for players to inflict negative consequences on team mates.

9.4.2 Resources

Since objects in traditional Indian games are everyday objects which are readily available in villages, traditional games do not involve objects such as money, armor and magic items that are found in Western games.

With the exception of Trump Cards, none of the traditional games involve skills or statistics that enable a character to perform an action better than other characters. Furthermore, even though Trump Cards involve statistics, the statistics are static and do not change permanently over the game. As such, character or team development has no central role in traditional games, unlike videogames (e.g. role-playing ones) in which it is desirable to raise a skill level or gain the ability to perform an action that a character could not do previously. In other words, in traditional games, state variables for player characters do not track skills because skill acquisition is non-applicable. (By "skill" acquisition, we mean the acquisition of skills on the part of the player characters, not the players themselves. Players in both traditional and non-traditional games improve their gameplay skill with practice.) On a related note, whereas a player character's "number of lives" is a common attribute in many classic videogames, traditional games do not have this attribute.

Similarly, since player characters in traditional games do not have attributes that differentiate them significantly from one another, there are no boss monsters for players to defeat. More interestingly, although there are scores in both traditional and videogames alike, the score-keeping mechanism differs significantly. While players in videogames typically gain many points for each sub-goal that is successfully achieved and accumulate total scores that are stratospherically high (i.e. with several digits), only one traditional game has a similar reward structure. In the other 6 village games that implement score-keeping, a player's score increases only by one at each increment, and he has to expend a tremendous amount of effort to earn this point (e.g. having to run and catch one – or every – opponent). Hence, the total score rarely exceeds a digit. A plausible explanation is that larger numbers are difficult for rural children to count and manage, given their limited numeracy skills.

Lastly, in village games, no resources are hidden or secret, i.e. the information pertaining to each object's ownership is known to everyone. Moreover, although some actions are carried out with certain objects (e.g. deflecting an oncoming ball with a stick in Gucchi Garam), no object in the game world is sought after because it imbues its owner with a special ability (i.e. a "power-up"). That is, skill-enhancing objects do not exist. In fact, objects in traditional games are mostly meant to be acquired or moved around only.

9.4.3 Goals

The goal structure in traditional games is much less complex. More specifically, some contemporary Western games may contain goals that are:

Ephemeral – existing for a limited period of time only.

Continuous – player(s) having to maintain the game in a certain goal state for a certain period of time for the goal to be counted as attained.

Excluding – a goal becoming meaningless once another goal has been achieved.

Incompatible – the game states for two or more goals are mutually exclusive, such that achieving one sub-goal or goal precludes the attainment of other goals.

Selectable – player(s) can choose to achieve one or more goals from a set of options.

Unknown – goals that are not known to the players.

On the other hand, none of the traditional Indian village games we observed have the above goals. More important, in traditional games, once players have been assigned their roles, their goals also become known.

At the more concrete level of specific goals, since the playing field in traditional games is visible to all, exploring an unknown territory is not a meaningful goal, even though this goal is found in some videogames. Similarly, having to clear a path to a destination by removing obstacles (e.g. in Sokoban) is never a primary goal in traditional games. Interestingly, while there are non-digital Western games in which the player wins by inducing the opponent(s) to break a rule, we have not observed this goal in our sample of traditional Indian games.

More important, while there are goals in traditional games, many of the village games do not have end conditions, i.e. play continues after the goal is achieved, and the game only ends through mutual consent when everyone is tired or bored. In contrast, with the exception of a few videogames such as massively multi-player ones, the majority of videogames end after certain conditions are attained. In other words, it seems that many traditional games are played to experience the play process, without having to experience the euphoria of final victory after all the rounds are over.

The relationship between goals and sub-goals in traditional games is particularly interesting. While videogames can be designed to offer more difficult levels of challenges through more – and possibly faster moving -- enemy characters, in kinesthetic games, the number of players and their speed have physical constraints and cannot be increased arbitrarily to make the game any more difficult. Instead, in variations of one "tag" game, we have observed that the version played by older children is more difficult because it introduces an additional sub-goal, i.e. having to touch a marked spot in the opponent's turf before one is allowed to return to one's turf and win that round. This observation suggests that introducing appropriate sub-goals into a videogame for rural children is a natural approach (to this user group) for raising the game's difficulty level.

9.4.4 Actions

The actions that players perform in traditional Indian games are governed by rules that are simpler than Western games. For example, the probability of success for actions in village games is determined based on rules that are known to every player. In contrast, in videogames, the probability of an action succeeding is determined behind the scenes by the computer, based on rules that can be non-transparent to players. Moreover, some of these rules for determining success can be more sophisticated in Western games – both digital and non-digital alike – when they have to take into account those attributes that reflect the characteristics of the player character (and object) performing (and used to perform) the action. On the other hand, player characters and objects in traditional Indian games do not have such attributes. Finally, unlike real-time videogames where the game can advance according to a preprogrammed logic without players having to do anything, the game state in traditional games advance only when the players perform actions. In fact, traditional games do not have system-generated events (e.g. earthquakes) which are independent of the players' actions.

9.4.5 Rules

Since village games hardly involve elaborate score-keeping mechanisms, player characters do not have "hit" points. As such, while it is possible to eliminate a player entirely, it is not possible to "damage" him partially. In other words, an offensive action has an atomic outcome, i.e. it either eliminates entirely or fails to do so. Subsequently, since players cannot be wounded, they cannot be healed, unlike Western videogames like role-playing games. Moreover, there are no extended actions in traditional games, i.e., assuming that an action succeeds, it takes effect as soon as it is performed. Most of all, even though traditional games share several similarities with arcade games, there are no "combination actions" in the former. What we mean is that when multiple actions are performed in traditional games, the cumulative effect is the additive effect of each separate action, whereas with "combination actions," there are non-linear outcomes that are more than sum of the individual effects of the individual actions.

Traditional games have elaborate rituals associated with physical movements and the playing space, which are usually not observed in videogames. One of them concerns that of the "turf," which is a region that is often associated with an individual owner or team. A turf can confer an advantage to its owner when he is standing inside (e.g. the state of invulnerability, as in the Gucchi Garam game), by being "off limits" to members of the opposing team, or imposing a time limit on how long members of the opposing team can spend in it. Conversely, a turf may be disadvantageous in that it imposes restrictions, e.g. requiring that the player only runs along a circle in either a clockwise or anti-clockwise direction only, or requiring that he runs only in the area on one side of a line. More interestingly, these restrictions are not completely fixed by the game rules, in that the player can choose which direction or side he to run on for the current round. Finally, a player may be prohibited from performing certain actions outside his turf.

On the same note, while players may move around freely in village games that entail movements in the outdoors, these games do not have playing fields that contain inaccessible areas or obstacles. This is a contrast to the videogames that contain such challenges in the game world, which make it necessary for the player to navigate a path around the obstacles. It is also notable that while videogames may include traps in the game world, traditional games do not involve traps.

Next, the objects in traditional games are physical equipment that are available in fixed and highly limited quantities. As such, while they can be eliminated from the game, they are also non-renewable. That is, new resources cannot be created in a game session.

For instance, unlike those videogames where resource management is a central objective, in traditional games, it is not possible to create new units in the middle of gameplay (e.g. creating new military units in digital war games like WarCraft). By implication, resources in traditional games cannot be invested to reap more resources in the future.

9.5 Game Redesign and Preliminary Results

Based on the insights that we have gleaned from the above analysis, we designed a game called Tree-Tree (Figure 9.3), which was mirrored after the "tag" family of traditional games. Tree-Tree aimed to help the player acquire an English vocabulary of everyday fruits in India. Every tree in the game stood for a unique fruit, and a computer-controlled character would call out one of these fruits. The player had to move to and touch the corresponding tree, so that he became invulnerable from capture by a computer-controlled opponent. While moving, he needed to evade capture from the computer-controlled opponent.

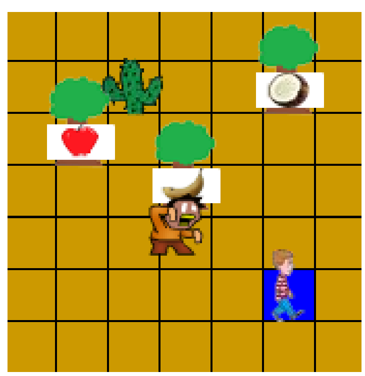

Figure 9.3: The Tree-Tree game, mirrored after village games in the "tag" family.

We evaluated Tree-Tree with 21 children aged 8-14 from LUCKNOWPUBLIC. Each child had up to 3½ hours with the game. 12 children were videotaped, after which their video recordings were transcribed. Whereas we sometimes spent days to explain Train Tracks and Crocodile Rescue to rural children in prior studies, almost every participant grasped the goal and rules of Tree-Tree after we demonstrated the game to him or her. In particular, some participants were visibly excited when we showed them Tree-Tree, since they could relate easily to the game. Furthermore, 4 players found Tree-Tree to be easy due to their familiarity with similar traditional games, and asked if it could be made more difficult.

Similarly, players seemed to concentrate hard when moving to the correct tree, while doing their utmost to avoid the opponent at the same time. At least 3 participants went a step further and enjoyed enticing the opponent to draw near, before they ran away in another direction. Children often employed this strategy in their village games, and this strategy transferred to their approach in playing Tree-Tree. More importantly, this degree of engagement was a noticeable contrast from their frustration or boredom with previous games. In fact, at least 2 players were so engrossed in the game that they failed to notice what was happening around them in their physical environment.

Most important, at least 2 players repeated the spellings and pronunciations of the fruits aloud after the game, without any prompting from us. In future research, it would be imperative to explore how we can design educational games – modeled after selected aspects of traditional games – which scaffold rural children to *transfer* the subject matter from the digital game back to their everyday village games. For example, players could be encouraged to play a non-digital variation of Tree-Tree among themselves, in which they spell out the English words for fruit trees in the playground. In this way, the digital game could potentially transform everyday, non-digital play into educational experiences. The possibilities are exciting.

9.6 Directions for Future Experimentation

This section has provided a tool for designing videogames that target children in rural India, and potentially other rural regions. By providing a detailed description of the elements in traditional village games, we have provided the reader with a “palette” of game elements that game designers can draw from to put together new game designs for children in rural developing regions. As such, this tool is *generative* in that it facilitates new designs. A thorough validation of this tool in any meaningful sense will take years. We encourage other researchers and practitioners who are working on videogames that aim to improve lives for poor children in the developing world to experiment further with this tool. As a first step, our preliminary results with Tree-Tree suggest that designing videogames with the same game mechanics as those found in traditional games, while leaving out those mechanics that are absent, ensures the most successful videogame designs that rural children can relate to more readily.

More broadly, what are some of the emerging principles that can be explored in future work? We will attempt to outline some of our thoughts here. Through our earlier attempts to design videogames intuitive to rural children, we observed that situating our game designs in everyday settings that are familiar to rural children is necessary but insufficient. As such, we were prompted to learn more about 28 traditional village games through our contextual interviews, analyze the elements in these games and study how they differed from contemporary Western videogames.

Our analysis has enabled us to interpret playability issues with Frogger, Dancer, Train Tracks and Crocodile Rescue. We now believe that our prior experiences with Western-style games led to associated cultural expectations creeping into our processes and design outcomes for these 4 games. Dancer is reminiscent of “tag” games that require the player to eliminate opposing players by hitting them with a ball. The goal in Frogger is similarly intuitive. On the other hand, having to explore an unknown world in Track Tracks to lay railroads is a non-intuitive goal. Users also appear to understand the actions in Dancer (e.g. throwing tomatoes, and moving among the audience) more readily than the “bait the crocodile” action in Crocodile Rescue. Similarly, there were players who struggled to make sense of the sub-goals related to resource management in Dancer, in which they gain more tomatoes by approaching green-colored members of the audience.

These observations are highly consistent with the above analysis on the game mechanics that are found – and *not* found – in traditional games.

A critical ingredient for a game to be enjoyable is that it has to be challenging without being too difficult. Even though we designed Crocodile Rescue such that there were more stationery crocodiles in higher levels, players did not find it easy to navigate the “maze.” Similarly, navigating around impassable terrains in Train Tracks appeared to be an unfamiliar task. Players also did not like the fact that there were more dancers on stage to hit in the more difficult levels in Dancer. In other words, increasing a game’s difficulty by increasing the number of obstacles or enemies is inconsistent with how the difficulty in traditional games is scaled up. If we had to design these games again, we would design the advanced levels to be harder by featuring more sub-goals in them, or by drawing on other rules and penalties that we described above.

On a concluding note, although we take the position that videogames patterned after traditional games are more intuitive to rural children, we similarly recognize that there is room to innovate on digital games and experiment with videogames that differ from traditional games. Only time will tell how soon rural children will learn how to play videogames that do not match their expectations about games, and more important, which game features are easier for them to grasp.

10 After-School Program

The challenge with evaluating any language learning project is that language acquisition is a long-term process on the learner's part. Worse, with a novel technology solution that has yet to be institutionalized, there were tremendous logistical obstacles in running a pilot deployment over a non-trivial duration. After three years into this project, which started with needs assessments and feasibility studies, followed by subsequent rounds of field testing interleaved with numerous iterations on our technology designs, we have established the necessary relationships with local partners to undertake such an evaluation. This Chapter describes the results from a semester-long pilot study – the longest so far in this project – that took place during the project's fourth year. The study involved 27 rural children who participated in an after-school program which we implemented in their village.

The pilot study was carried out in collaboration with a non-government organization (NGO) in North India under the terms of a Memorandum of Understanding. The study took the form of an after-school program, which we held during the afternoons at the premises of a private village school affiliated with this NGO. However, our goal was to investigate learning impacts that English as a Second Language learning games on cellphones have on lower-income rural children. As such, students who were already enrolled in this school were ineligible to participate in the study. Instead, we invited those parents who could not afford the fees for this private school – and hence had no choice but to send their children to the less expensive schools in the same vicinity – to give consent for their children to participate.

In the context of the after-school program, on average, we ran three sessions per week. Each session lasted two hours in the afternoon. Children from neighboring villages attended the after-school sessions after finishing their regular classes in the morning. In the after-school program's sessions, we loaned cellphones preloaded with ESL learning games to participants. The after-school program took place from late December 2007 to early April 2008, and spanned sessions on 38 days in total.

10.1 Data Collection

As our preparation for this pilot study, we made two trips to India, i.e. once in the summer of 2007 to familiarize ourselves with the pilot location and end-user community, and a second time in December 2007 to kick-off the actual pilot. 4 local staff members were hired to run the after-school sessions on an everyday basis. It was critical to have local personnel because we needed adequate manpower to run the program, local children were more comfortable in the presence of local staff who understood their culture and spoke their language, and it was not practical for foreign researchers to be stationed on-site in the long term. 3 staff members were engineering undergraduates in their last semester, while the last had graduated a few years ago. We spent two weeks training them to run the after-school sessions and perform data collection. We continued to coordinate with them regularly via conference calls and emails after we left India.

We interviewed participants on their demographics such as their ages and the grades they were currently enrolled in in school. During the interviews, we also asked other questions, such as the number of cellphones that their households owned, what they currently and/or had previously used cellphones for, their television watching habits and frequency, as well as their parents' occupations. The questions on media and technology exposure were included because these variables were expected to impact participant ability to learn using cellphone games.

To ensure that each participant has the basic numeracy and ESL literacy to benefit from cellphone-based learning, participants were required to pass a qualifying test, that is, obtain at least 50% of the total score. The test required them to complete one-word blanks using English words about themselves (e.g. name, age, school, grade). They were also asked to fill in the missing letters in the alphabetic sequence, write numbers in the

Arabic notation, match words with their pictures, spell the words for everyday objects, and describe a picture of a market scene with short sentences. The qualifying test was designed such that an average child in India with no learning disabilities who has finished 1st grade in a reputable urban school would obtain a perfect score on it.

By using the qualifying test as a screener, we ensured participants were numerate. This was important because we had previously found it difficult to teach children to use the cellphone keypad's to play e-learning games when they could not read the numbers from 0 to 9 in Arabic notation. Similarly, by ensuring that participants were familiar with the English alphabet, we could target a more advanced syllabus that went beyond the alphabet. We made this decision since Horowitz et al. (2006) had already investigated the efficacy of cellphone-based learning for the English letters with preschool children in the USA.

Moreover, since success in acquiring a second language is correlated with literacy in one's native language, we administered a test which evaluated participants on their ability to read in Hindi. Every child was given a short passage that described a diet for promoting dental health. Each child was then asked to read the passage aloud so that we could observe his or her fluency and accuracy. These sessions were videotaped. After that, every participant was asked to write answers to written questions that tested his or her comprehension of the passage, in Hindi. We designed this test such that an average child who has completed 3rd grade in a reputable urban school in India would obtain a perfect score on it.

Our primary method of assessment was to administer pre- and post-tests which evaluated participants on their ability to spell the common nouns that the curriculum for the pilot study targeted. Although the curriculum targeted other competencies such as listening comprehension and recognition of written words, our assessment emphasized spelling, which as a recall task was cognitively more difficult than recognition tasks.

We maintained attendance records for the participants for every session. We also videotaped each session so as to have contextual data that could potentially account for their test performances. The video recordings captured the classroom proceedings and individual participants' interactions with the games. The latter recordings captured participants' levels of engagement with the games as shown in their facial and body expressions. The pilot staff member who was responsible for videotaping the sessions tried to ensure that every participant was videotaped playing at least one level in the curriculum per session. The recordings were later transcribed and translated from Hindi to English.

Finally, for every session, we asked pilot personnel to write a report summarizing what happened in that session, as well as how well each participant interacted with the games. The latter not only covered usability and learning obstacles, but also included pilot staff's observations on the attitude and persistence that each child demonstrated towards learning.

10.2 Participants

Owing to the strong relationships that our NGO partner had built with the local community over more than a decade, we were able to generate a high level of support among parents in this rural community. Parents associated the private school where we held the after-school program at with high quality education, and in fact, most of them could not afford to enroll their children in the regular morning classes there. The school caretaker was an influential figure within the community. He was supportive about our work and spread the news about the program personally. One of the pilot staff members was also working at this school as a teacher and knew the community well.

In short, the after-school program started with a high level of support among the parents, due to our association with this school and its staff. Interested parents were invited to attend a briefing session, where we explained the goals of the program and obtained their consent. In the case of parents who were interested but were unable to attend the session due to work, pilot staff visited them at their homes to obtain their consent. At the briefing, many mothers wanted to enroll in the study so that they could resume their schooling, which was disrupted when they had to marry at a young age. We were gratified by their enthusiasm but emphasized that the pilot study targeted children only.

In total, we obtained consent for 47 children to participate in the study. However, we needed to turn 16 of them away; 15 children did not pass the qualifying test while the 16th was attending private tuition for English, which represented a confounding variable. Of the 31 children whom we started the pilot with, 6 of them left the program mid-way. Reasons for attrition include time conflicts with private tuition lessons (2 children) and disinterest in attending the pilot sessions (the other 2 children). From post-deployment interviews, we understand that the latter was attributed to caste tensions between the 2 children, who were members of the lower castes, and some upper-caste children in the program.

10.2.1 Demographics

The 27 children who participated in the study until it ended were aged 7 to 14 (mean = 11½ years) and belonged to grades 2 to 9 (mean = 6th grade). There were 11 boys and 16 girls. 5 children came from the upper castes while others belonged to the lower castes. The gender and caste breakdown in this sample seemed to mirror the demographics in the community. Every participant attended between 8 and 29 sessions (mean = 20) in the after-school program, broken down according to the following three functions:

- Cellphone training: 0 to 5 sessions (mean = 4) where we taught participants how to use the cellphones, perform alphanumeric input and play mobile games,

- ESL learning: 4 to 17 sessions (mean = 10) in which participants played ESL learning games on the cellphones, and
- Assessment: 4 to 7 sessions (mean = 6) for administrative tasks and data collection, e.g. demographics interviews and various tests.

In India, traditionally, only the upper castes owned land. As such, the upper castes earn their livelihood on the land or run small businesses, while lower castes graze their goats, work as daily-wage laborers or perform menial jobs in the homes of the upper castes. Land-owning and non-land-owning families told us that they earn up to one 1 lakh (US$2,500) and 50,000 rupees (US$1,250) respectively per year.

10.2.2 Hindi and English Baseline

26 of the 27 participants were enrolled in the same school, where Hindi is the medium of instruction. The 27th participant was a school dropout. Assuming regular school attendance, the typical participant would have taken classes on Hindi and English for 5½ and 3½ years respectively prior to the study.

We devised a grading rubric to evaluate each participant on the Hindi literacy test and qualifying test. On the former test, participants scored 7.9 out of 18 on average (σ = 4.5, n = 19). 2 participants turned in blank answer sheets. We observed the following problems in the answers submitted by the remaining test-takers:

- Wrong answers due to poor comprehension of the questions (5% of the test-takers) or passage (53%), or responses that simply repeated the questions (32%)
- Spelling errors (16%)
- Grammatical errors, i.e. using the incorrect form of the verb for the subject's gender (21%), or the incorrect form of the noun for the subject's singularity vs. plurality (5%)
- Inability to phrase responses in complete sentences (5%)

Table 10.1
Breakdown of Participant Performance on Qualifying Test

Test Section	Poor	Fair	Good
About myself	13% left blanks empty or filled them in Hindi	32% filled in blanks with at least 1 misspelling	55% filled in blanks with correct spellings for most questions
Alphabet	5% filled in less than 7 blanks in alphabetic sequence correctly	13% filled in ~11 blanks in alphabetic sequence correctly	82% filled in at least 12 out of 13 blanks in alphabetic sequence correctly
Word recognition	14% matched up to 2 out of 6 words with correct pictures	14% matched 4 out of 6 words with correct pictures	72% matched at least 5 out of 6 words with correct pictures
Spelling	41% spelt up to 2 out of 6 words correctly	18% spelt about 3 out of 6 words correctly	41% spelt at least 4 out of 6 words correctly
Picture description	18% made no attempt to answer this section	50% wrote answers as individual words, not full sentences	32% wrote intelligible answers in sentences

On the qualifying test, on average, participants scored 44.0 out of 50 ($\sigma = 5.5$, $n = 22$). Our grading rubric indicated how participants should be classified as "poor," "fair" and "good" on every section of the test. Table 10.1 gives the breakdown of how test-takers were distributed across categories for selected sections, and descriptions of the categories. In summary, the average participant had a good knowledge of the alphabet and a fair vocabulary of written words that she could read. On the other hand, she was weak in recalling and spelling everyday nouns, and even weaker in constructing complete sentences with these words. Despite the wide range in the ages of the participants, it appeared from their performance on the above tests that the variation in their English proficiency was much narrower. Specifically, we estimated the average participant was comparable to an urban child in India who had taken between 1 and 2 years of English classes.

Notably, only 10 children (45%) could spell their names correctly in English on the qualifying test papers. We had a chance later to interview the teacher who taught them English in their school. She revealed that her pedagogical approach revolved around having students copy sentences from English textbooks into their notebooks. She felt that it was not worth putting in more effort to teach English since she believed she was underpaid.

10.2.3 Technology Baseline

Among the 27 participants, 25 of them came from families who owned at least one cellphone; 5 participants belonged to families that owned 2 phones each while 2 participants came from families which owned 3 phones each. The cellphone was usually used by the eldest male member in every family, and in fact, 3 boys aged 13-14 possessed their own cellphones. 8 participants came from families that owned cellphones with a color screen, as opposed to monochrome display. Two of the above cellphones – both of which belonged to 2 of the above 3 boys – contained built-in cameras.

Nonetheless, cellphone ownership and access were separate issues. Among the 25 children whose families owned at least one phone, 6 of them – 5 girls aged 7-11 and a boy aged 12 – were prohibited by their parents from using the phone, either entirely or most of the time. In general, in poorer families, it seemed that children were allowed to receive (free) incoming calls, but not play mobile games lest they drop the devices. On the other hand, in wealthier families, children were allowed to play on the phones. As such, although most participants were familiar with cellphones, it appeared that rural parents were more willing to entrust these relatively costly devices to their sons (vs. daughters). In total, 15 of the 27 participants (56%) reported having played cellphone games before.

10.3 Curriculum and Game Design

One of the major challenges with carrying out a pilot study over a non-trivial timeframe was that we needed to develop sufficient digital content that could last throughout its entire duration. We ensured that our syllabus was aligned with local ESL learning needs in India by recruiting a local ESL teacher as our curriculum developer. She had a decade's experience as an ESL teacher at a prestigious urban school, located in the same geographic region as the after-school program.

10.3.1 Curriculum Design

Given the above attendance rate, the ESL curriculum for the pilot was designed to be comparable to the amount of material that a qualified teacher could reasonably cover in 18 hours with rural children in a regular classroom. We chose to situate our syllabus within the classroom theme, which participants could readily relate to. Concretely, the syllabus included:

- Common nouns that are found in the typical classroom, e.g. chair, table, door.
- Verbs that can be performed with the above nouns, e.g. sit, write, open, close.
- Sentence structures for constructing sentences out of the above nouns and verbs, e.g. "This is a __."
- Sentence structures for phrasing question-and-answer sequences with the above nouns and verbs, e.g. "What is this?", "Where is the __?"

In designing the curriculum, we took participants' performance on the qualifying test into account. The curriculum was also informed by our attempts to converse with participants informally, during which we learned that they did not comprehend simple questions about themselves, could not recall the English words for most objects around them in the classroom, and made grammatical errors (e.g. those pertaining to pronouns). The curriculum hence targeted the above syllabus in terms of listening comprehension, word recognition (of the written word), sentence construction and spelling.

10.3.2 Game Design

We designed a set of ESL learning games for the cellphone platform that targeted the above curriculum, and piloted them in the after-school program. Our designs drew on three resources, namely:

1. recurring patterns from state-of-the-art commercial software applications for language learning, which represented the best practices that we reused to avoid reinventing the wheel (Chapter 5),
2. elements in traditional village games, which more closely matched the expectations and understandings that rural children have about games, compared to contemporary videogames that were largely Westernized (Chapter 9), and
3. lessons from several previous rounds of field-testing and iterations with rural children elsewhere in India (Chapter 8).

In this subsection, we walk the reader through a subset of the screen designs.

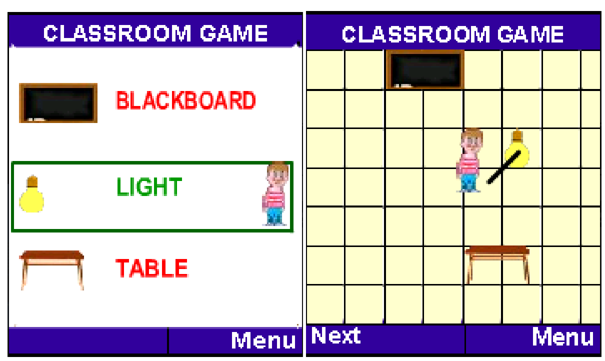

Figures 10.1 and 10.2: Figure 10.1 introduces the English vocabulary for common nouns in the classroom. Word-picture association is a technique employed by many successful commercial language learning software. As the boy moves to each object, the software highlights the corresponding word in a different (green) color and says the word aloud. Figure 10.2 situates these objects in a classroom scene and builds on the earlier screen by demonstrating how to use the nouns in complete sentences. As the boy moves to each object, the software says the "This is a __" phase aloud for the corresponding object.

In earlier field studies, we observed that rural children did not readily associate a game with learning. It seemed that they viewed a game as an activity to be played purely for pleasure, and did not pay attention to the educational content embedded within game activities. In contrast, when educational content such as English words and phrases were introduced in non-interactive screens which were separate from interactive game screens, rural children appeared to grasp more intuitively that the software was trying to teach them those English words and phrases. Learners subsequently paid more attention to the latter. Figures 10.1 to 10.4 showed selected screenshots in which we introduced words and phrases – both written and spoken – to the learners.

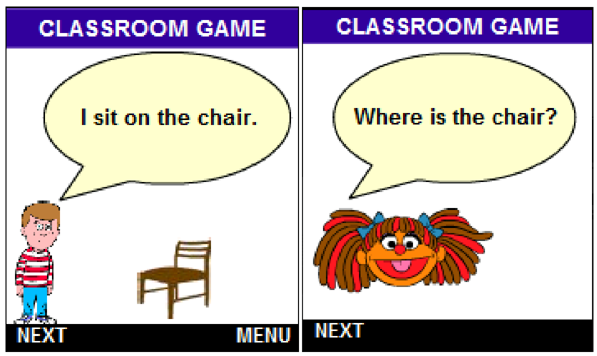

Figures 10.3 and 10.4: Both figures introduce additional phrases that the nouns and verbs in the syllabus can be used in. Figure 10.3 teaches a phrase that associates the verb "sit" with the noun "chair." Figure 10.4 shows how to ask questions using the "Where" keyword. Abstract phrases and function words such as "where" are difficult to convey graphically. Hence, when they are taught for the first time, the software explains their meanings orally in Hindi.

The games tested players on their comprehension and recall of the words and phrases. For example, the game shown in Figure 10.5 says the word aloud for one of the objects displayed on screen. The player needs to identify the correct object and push it onto the area that is blinking blue. At the same time, he needs to avoid the balls thrown by the computer-controlled opponent. This game was an adaptation of Giti Phod, which was one of the traditional games that children play in Indian villages. In Giti Phod, players in a team have to arrange some objects (e.g. rocks) into a given configuration (e.g. a heap), while avoiding being hit by a ball thrown by members in the opposing team. In our experience so far, we have observed that rural children found it more intuitive to understand videogame rules when the designs of these videogames drew on the rules found in the traditional village games that they play everyday.

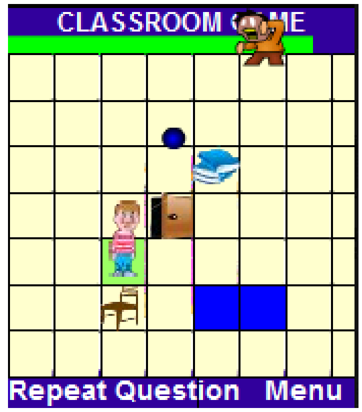

Figure 10.5: A word-picture matching game which is an adaptation of one of the traditional village games that children in rural Indian play everyday.

Given that television has become a pervasive media among all economic classes in India, it only made sense to draw on popular culture in India to make our designs more appealing to children there. One of these sources is Sesame Street, which is a successful television program for young children that has local co-productions around the world – in both industrialized and developing countries. Its producers in India have found some of its localized characters to be popular with children in India, and we incorporated the same characters into our game designs for teaching (Figure 10.4), quizzing (Figure 10.6) and congratulating (Figure 10.7) the user.

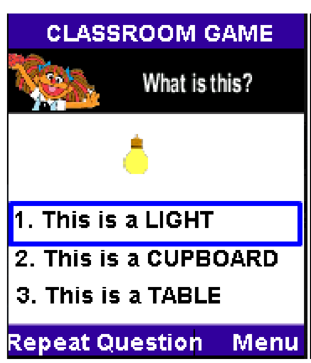

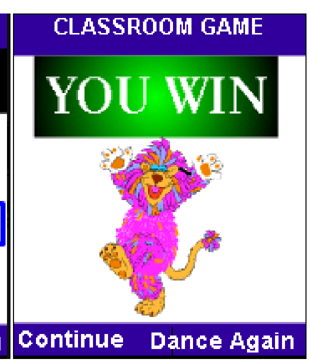

Figures 10.6 and 10.7: Localized characters from the Indian production of Sesame Street tested the player on his ability to engage in question-and-answer style dialogues, and performed a victory dance for the player upon successful completion of each level in the game.

The activity that targets spelling skills is illustrated in Figure 10.8. The player is given an image (e.g. blackboard) and is required to spell the word corresponding to it.

Some of the letters in the word are displayed, while blanks are shown for the remaining letters. The player moves between blanks with the arrow keys. Once he has filled in all the blanks and submitted his answer, the correct and wrong letters are displayed in green and red respectively. This feedback constitutes the first level of hints that we have designed to help the learner arrive at the correct spelling. If the player spells the word correctly, he proceeds to the next game. Otherwise, all blanks are cleared after a short pause and he is required to spell the same word again.

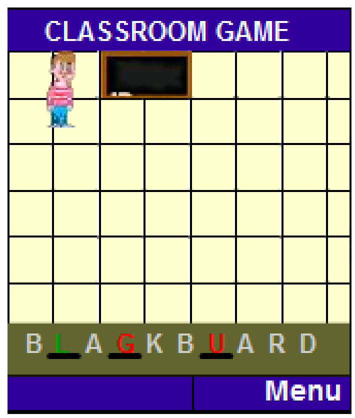

Figure 10.8: The first level of hints in the spelling activity. After the player has tried to spell the word by filling in the blanks with letters, the correct and incorrect letters are shown in green and red respectively.

If the player is unsuccessful in spelling a word correctly after two attempts, the second level of hints (Figure 10.9) appears to provide him with additional learning support. Based on the blank that the cursor is currently located at, the game displays a set of possible letters for him to narrow down the choice of candidate letters.

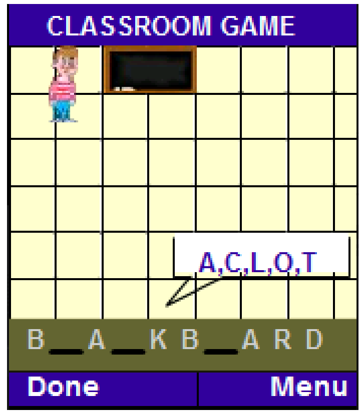

Figure 10.9: The second level of hints in the spelling activity. For every blank, a set of possible letters are displayed to provide the player with some assistance, if he was unable to spell the word correctly after two attempts.

The curriculum is broken up into a total of 6 levels in the games. On every screen, the player can access a menu through a shortcut button. Among various options, this menu permits him to move to an earlier level in the curriculum to repeat the material, as well as to move to higher levels in the curriculum. The software was designed so as not to require airtime, which was expensive for most rural families. We implemented the above games on Adobe's Flash Lite and Qualcomm's BREW (Binary Runtime Environment for Wireless) platforms. We piloted the games on the Motorola Razr V3m cellphone model, which has a fairly large screen for children to view comfortably.

10.4 Pilot Sessions

In those sessions where participants were taught how to use the cellphone, they were shown how to move their sprites with the arrow buttons. They were also taught how to perform alphanumeric text entry, since most of them did not already know how to do this. Sprite movement and text input were the skills essential for playing the above games that we designed. Pilot staff were therefore asked to write some simple, short sentences on the blackboard, and ensure that each participant demonstrated his ability to enter those sentences via alphanumeric text input.

Some other sessions focused on administrative tasks, such as the above tests and demographics interviews. We learned that a few participants had difficulty reading a small subset of the English alphabet despite having passed the qualifying test. We spent two sessions coaching them on those less-frequently encountered letters, so that they would be better prepared for the syllabus targeted in the pilot. Next, at least 8 children had seen the localized Sesame Street characters on television, but did not know their names. To help participants better relate to the characters, so that our games would appeal to them even more, we introduced the characters at the start of the semester. We also screened three localized episodes on separate occasions for participants to watch. These episodes were chosen such that they were educational but did not target English learning. Each episode lasted ½ hour, and we observed that participants enjoyed the humorous acts performed by the characters.

Figure 10.10: In the after-school sessions, each participant was loaned a cellphone preloaded with English language learning games. Participants were taught how to start the games, and were asked to focus on learning English when playing the games on their own.

The remaining – and majority of the – sessions focused on ESL learning. A two-hour session was typically structured as follow: after an exchange of greetings, pilot staff took attendance and briefed participants on the learning objectives for that day. If new games were deployed that day, pilot staff explained and demonstrated how to play them to the participants in small groups. Each participant was then handed a cellphone to play the games on her own (Figure 10.10), and were told to focus on learning the English syllabus that the games covered. Participants who were absent from previous sessions received help from pilot personnel in learning how to play those games that they were unfamiliar with. Pilot staff members were limited to providing technical support; and were explicitly instructed not to teach English or communicate with the participants in English. There was a short break of 10 to 15 minutes in the middle of each session. At the end of each session, pilot staff took back the phones so that they could recharge the batteries overnight and download new games onto the devices. Participants received no compensation other than a small packet of biscuits after every session.

10.5 Quantitative Results

On the pre- and post-tests, test-takers were awarded one point for each common noun in the syllabus that was spelt correctly.

10.5.1 Post-Test Gains

The mean pre-test score was 5.2 out of 18 ($\sigma = 3.3$, $n = 27$) while the mean post-test score was 8.4 out of 18 ($\sigma = 5.5$, $n = 24$). On average, participants exhibited statistically significant post-test gains on a one-tailed t-test with an effect size of 0.71 ($p = 0.007$). We present the frequency histograms for both pre- and post-test scores in Figure 10.11. They illustrate that the score distribution had shifted toward the higher end of the spectrum after the deployment.

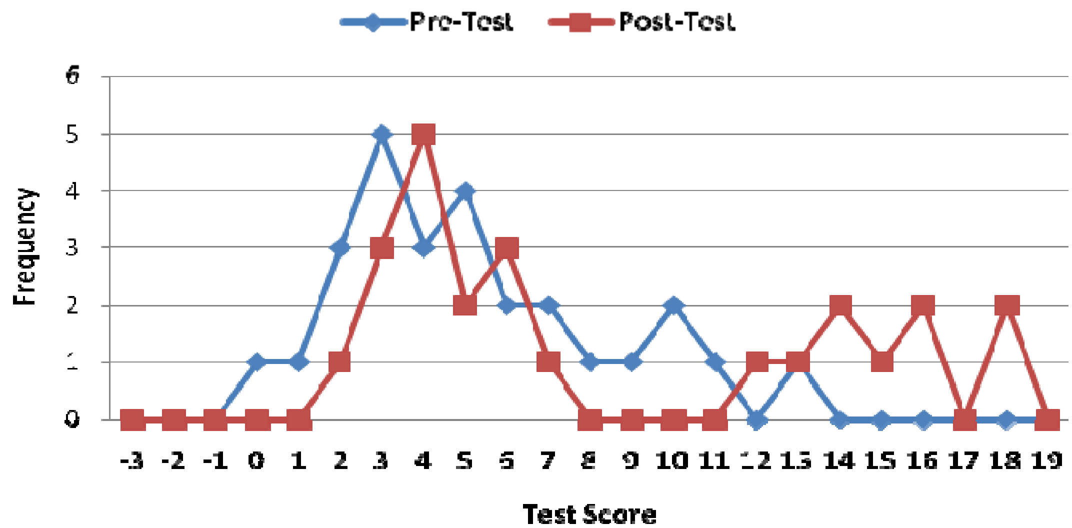

Figure 10.11: Frequency histogram of participant scores on the pre- and post-tests.

The average post-test gains was 3.4 out of 18 (σ = 3.3, n = 24). The gains exhibited a fairly large variation, and ranged from -2 (two participants exhibited negative gains) to 9 out of 18. We present the frequency histogram for post-test gains in Figure 10.12.

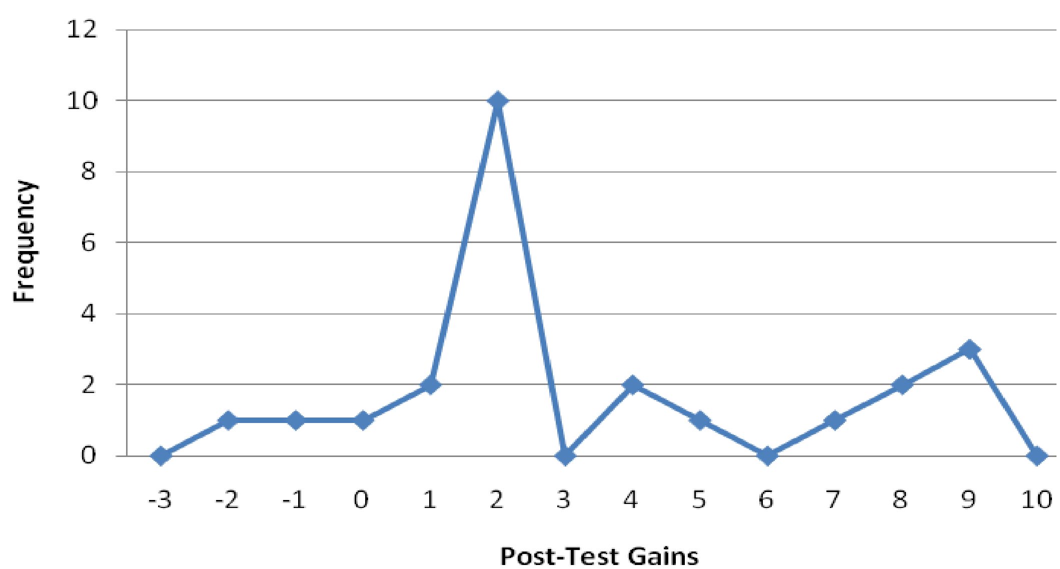

Fig. 10.12: Frequency histogram of participant post-test gains.

10.5.2 High-Gains vs. Low-Gains Learners

We sought to understand how participants' post-test gains were correlated with their demographics and performance on other tests. We also categorized participants into two groups, namely, high-gains learners and low-gains learners, based on their post-test gains. A child whose post-test gains exceeded the mean of 3.4 was categorized as a high-gains learner, else he was classified as a low-gains learner. In total, 9 participants were classified as high-gains learners while 15 children were categorized as low-gains learners. 3 of the 27 participants could not be classified since they were absent on the day when the post-test was administered.

On a normalized scale, when the 27 participants were taken as one group, the average pre-test score was 29% whereas the average post-test score was 47%. The latter score did not seem high in absolute terms, i.e. on average, a participant could not spell over half of the common nouns targeted in the syllabus by the end of the intervention. However, once the participants had been classified, on a normalized scale, high-gains learners scored 41% (80%) on the pre-test (post-test) whereas low-gains learners scored 19% (27%) on the pre-test (post-test), on average. In other words, high-gains learners not only showed larger post-test gains but also appeared to "start" with a higher mean pre-test score. More importantly, the post-test gains for both high-gains ($p < 0.001$) and low-gains learners ($p = 0.076$) were statistically significant, with effect sizes of 0.54 and 2.24 respectively. That is, both groups of participants exhibited learning gains as a whole. (But the average post-test gains for the low-gains learners were only marginally significant, owing to the 2 participants who obtained lower scores on the post-test compared to their pre-test.)

Table 10.2
High-Gains vs. Low-Gains Learners in terms of Demographics

		Age (Years)	Grade Enrolled in School	Days Spent Learning ESL Games	Days Spent Playing ESL Games
Low-gains learners (n=15)	Mean	10.7	5th	3.9	10.5
	σ	1.8	1.9	1.4	4.2
	Min.	7	2nd	0	5
	Max.	14	9th	5	16
High-gains learners (n=9)	Mean	12.8	8th	3.8	10.2
	σ	1.2	0.9	0.8	5.2
	Min.	11	7th	2	4
	Max.	14	9th	5	17
Is difference between means significant?		Yes ($p = 0.002$)	Yes ($p < 0.001$)	No ($p = 0.4$)	No ($p = 0.4$)
Correlation with post-test gains (r)		0.45	0.61	0.11	0.10

In Tables 10.2 and 10.3, we examined how high-gains learners may differ from low-gains learners in terms of demographic information. On the whole, high-gains learners did not appear to differ significantly from low-gains learners in terms of the number of days that they had spent on learning how to play the cellphone-based games ($p = 0.4$) or actually playing the games to learn English ($p = 0.4$). Instead, high-gains learners belonged to higher ages ($p = 0.002$) and were enrolled in more advanced grades in school ($p < 0.001$). In fact, post-test gains exhibited high correlation with grades enrolled in school ($r = 0.61$) and medium correlation with age ($r = 0.45$).

Table 10.3
High-Gains vs. Low-Gains Learners in terms of Demographics

	Sex	Caste	Media Exposure	Attitude*
Low-gains learners (n=15)	67% (33%) were females (males)	77% (23%) belonged to lower (upper) castes	73% have played games on cellphones prior to pilot; 60% (40%) watched less (equal to or more) than 1 hour of TV per day	24%, 38% & 38% were described as below average, average and above average learners respectively
High-gains learners (n=9)	44% (56%) were females (males)	88% (12%) belonged to lower (upper) castes	56% have played games on cellphones prior to pilot; 40% (60%) watched less (equal to or more) than 1 hour of TV per day	29% and 71% were described as below average and above average learners respectively

*The Attitude column is based on the observations that pilot personnel have on the seriousness and aptitude that participants exhibited as learners throughout the pilot. These qualitative comments were subsequently coded into the "below average", "average" and "above average" learner categories.

The proportions in Table 10.3 were presented for the sake of completeness. We were unable to perform any statistical tests on these proportions due to the small sample size, which for example did not satisfy the standard binomial requirement. We therefore caution the reader against drawing any firm conclusions from these statistics. However, when examining individual learners to identify surprising cases, we took the demographic variables in Table 10.3 into consideration. The qualitative analysis is deferred to the next section.

Table 10.4
High-Gains vs. Low-Gains Learners in terms of Test Scores

		Qualifying Test (out of 50)	Qualifying Test, Spelling Section (out of 6 words)*	Hindi Literacy Test (out of 18)	Pre-Test (out of 18)	Post-Test (out of 18)
Low-gains learners (n=15)	Mean	42.9	1.2	6.3	3.5	4.8
	σ	2.9	1.1	4.2	2.3	2.4
	Min.	37	0	0	0	2
	Max.	46.5	4	14	10	12
High-gains learners (n=9)	Mean	47.1	3.4	12.0	7.4	14.4
	σ	1.8	1.6	1.7	3.1	3.6
	Min.	43.5	2	10.5	2	6
	Max.	49	6	14	13	18
Is difference significant?		Yes ($p < 0.001$)	Yes ($p = 0.001$)	Yes ($p < 0.001$)	Yes ($p = 0.003$)	Yes ($p < 0.001$)
Correlation with post-test gains(r)		0.57	0.70	0.45	0.46	0.86

*In this column, we present the number of words that participants spelt correctly on the spelling section of the qualifying test, out of a total of 6 words.

Table 10.4 presents a comparison of the high-gains and low-gains learners in terms of their test scores. The former group outperformed the latter on the Hindi test ($p < 0.001$). Next, we analyzed the qualifying test results at two levels, namely, the score on the entire test as well as the score on the spelling section. We found that the high-gains learners outperformed low-gains learners on the entire test ($p < 0.001$) as well as the spelling section ($p = 0.001$). High-gains learners also obtained higher scores on the pre-test ($p = 0.003$) and post-test ($p < 0.001$), vis-à-vis the low-gains learners. In fact, participants' post-test gains exhibited a high degree of correlation with their qualifying test scores, for both the entire test ($r = 0.57$) and spelling section ($r2 = 0.70$). On the other hand, post-test gains had a lower degree of correlation with Hindi literacy levels (r = 0.45) and pre-test scores (r = 0.46).

10.6 Qualitative Results

The above quantitative results suggested that current levels of spelling proficiency and grades enrolled in school were the strongest predictors of success in learning how to spell new words through the cellphone-based games which we designed. Higher levels of Hindi literacy and academic preparation were also associated with higher post-test gains.

On the other hand, the number of sessions that participants had with the cellphone games – both for learning how to play the ESL learning games and learning ESL through the games – were not associated with post-test achievements. Among the 24 children whom we have post-test gains data on, 5 of them were classified as high-gains learners despite having played the games on only 4 to 7 days (mean = 6 days). Conversely, 7 participants were classified as low-gains learners in spite of having played the e-learning games for 13 to 16 days (mean = 14.7 days). More important, pilot personnel described 6 of the 7 low-gains learners as "hardworking" or "serious" about learning ESL. Similarly, we were curious about how the two students who exhibited negative post-test gains, as well as the school dropout, interacted with the cellphone-based games.

The above quantitative trends raise the following questions: How did some of the high-gains learners played the games such that they benefited despite lower attendance? In the case of some low-gains learners, why did they improve little on the post-test despite spending numerous days with the games and being perceived as diligent? To address such questions at the interaction design level, we reviewed our video recordings and daily reports. We hope to identify possible improvements to the technology designs and/or the after-school setting.

10.6.1 Interaction Patterns with the Technology

At first glance, it seemed that participants needed to attend the after-school program for more days. Our video recordings showed that only 3 of the participants reached the last level in the curriculum by the last session in the program. This was a surprise. Given that the curriculum was designed for 18 hours of instruction, we expected an average attendance rate of 10 gameplay sessions to constitute enough time with the games. On examining the video recordings more closely, we saw that at least 8 participants were using the game menu to skip ahead to other levels when they were unable to spell the words in the current level correctly after a few attempts. (We should note that the menu was not necessarily a negative feature. Among those 8 participants, at least 2 of them used the menu to skip those words that they already knew how to spell.)

We needed to understand why learners gave up on retrying the spelling activity for difficult words despite the hints in the spelling activity. On the whole, we observed 4 different levels of behavior associated with the spelling activity in the videos:

1. When students encountered a word that they could spell, they pressed the keypad buttons quickly and with ease to fill in the blanks for the missing letters.
2. When students saw a word that they did not know how to spell, some of them learned to spell it correctly with the help of the first level of hints.
3. Some of those students who failed to learn how to spell a word with the first level of hint eventually learned how to spell it correctly with the help of the second level of hints.
4. Other students never succeeded in learning how to spell certain words despite both levels of hints.

In general, we observed that high-gains learners succeeded in learning how to spell words after having seen their written forms displayed on earlier screens (i.e. such learners were able to spell those words correctly – without requiring any hints – on their first attempt in the spelling activity), or with only the first level of hints. It seemed that they did not require much scaffolding support from the software. In fact, from the video recordings of 9 high-gains learners, we saw that 5 (56%) and 1 (11%) of them depended on the first and second levels of hints respectively. In contrast, 12 (80%) and 8 (53%) out of the 15 low-gains learners who were videotaped relied on the first and second levels of hints respectively. It seemed that the low-gains learners, as compared to the high-gains learners, were less able to rectify their errors in filling in the blanks for the missing letters through the first level of hints, and required the second level of hints to attain the correct spellings.

Worse, the inability on the part of the low-gains learners to spell correctly with help from only the first level of hints made some of them visibly unhappy or bored when the second level of hints appeared. The reason for this distress was unclear. The learner could be frustrated that he was spending too much time to learn how to spell the word. Alternatively, on seeing the second level of hints show up, he could be demoralized that he had just been relegated to the ranks of the most inferior learners and needed the second level as a "crutch" in order to succeed.

Moreover, some low-gains learners struggled despite both levels of hints. In the videos, two of them turned to their neighbors and asked for the correct letters, and/or to chat. In some cases, participants were embarrassed to ask their neighbors for help again after so soon, and hence used the menu to skip to other levels in the games.

More important, we observed that participants – especially among the low-gains learners – may be able to spell the words in the spelling activity, but were not able to spell the same words on the post-test. We offer two plausible explanations. Firstly, some children may have learned to spell the words by their last session in the program, but had forgotten their spellings between their last session and the post-test. (The post-test was conducted a week after the last gameplay session.) Secondly, some children had never learned to spell the words in their entirety through the spelling activity, which involved filling in a few blanks and did not require the learner to spell the entire word eventually.

Nowhere in the video recordings did we observe any participant struggling with usability problems.

10.7 Discussion

Our reactions to the above results of the learning assessment were mixed. In an underdeveloped region where rural children did not have access to high quality English instruction in their regular school or elsewhere, we were excited to see the participants – both high-gains and low-gains learners – in the after-school program exhibit statistically significant post-test gains that could be reasonably attributed to our cellphone-based English learning games. On the other hand, the educational benefits were uneven among participants. This could be a cause for concern.

To begin with, high-gains learners outperformed low-gains learners on the pre-test, qualifying test and Hindi literacy test. In fact, participants' post-test gains appeared to be highly correlated with their existing levels of spelling proficiency (as measured by their performance on the spelling section of the qualifying test) and the grades that they are currently enrolled in. This observation suggests that those rural children who have a stronger academic foundation are the same children who are most well positioned to take advantage of the benefits that cellphone-based learning confers.

Our results are consistent with the outcomes of a study with rural and urban low-income children in India described in He et al. (2008). This study showed that weaker students benefited more from a teacher-directed pedagogical intervention, while stronger students benefited more from a self-paced, machine-based approach to English learning. These results should not, however, be interpreted to mean that we rule out technology-augmented learning entirely in the context of low-income children. Horowitz et al (2006) reported a study on videos for learning the letters in the English alphabet streamed over cellphones. In this study, a greater proportion of lower-income parents, vis-à-vis their higher-income counterparts, perceived the videos to have improved their children's knowledge of the alphabet.

In the face of the above overwhelming odds, what can we do to promote more equitable educational opportunities in the developing world? One possible – and perhaps cautiously optimistic – interpretation of the above results is that future research needs to be directed at understanding how instructional software can provide more scaffolding support for those rural children who have less academic preparation. As an example, the spelling activity needs to be redesigned such that the learner is guided such that he has to spell the entire word eventually. With this redesign, however, gameplay becomes prolonged and can potentially increase player frustration, as we have witnessed above. One remedy is to have e-learning games track learner performance, so that the software can be adaptive in skipping stages that are similar to those that the player has previously performed well in. Another implication for instructional design, which calls for additional investigation, is scaffolds such as hints that are less conspicuous, so that their appearing on screen does not diminish the learner's sense of self-esteem or achievement.

Unfortunately, adaptive educational applications require the application state to be stored and retrieved on the same mobile device. From the logistics standpoint, this requirement is more difficult to implement in developing regions since it is harder to ensure that the same learner uses the same phone – which stores his performance from prior session(s) – across multiple sessions. For instance, in an after-school program where attendance fluctuates from session to session, it would be prudent to keep a shared pool of cellphones, such that children who show up for the day's session can draw from. In these circumstances in which it is not possible to reserve a cellphone for each child, a wireless networked mechanism for synchronizing application state across all cellphones may be necessary. This, and the other issues that we have raised above, require further investigation for cellphone-based literacy learning to be more effective in targeting less academically advanced rural children.

As we think more widely beyond the cellphone to consider it as a component in the broader learning environment, such as the after-school program model that is readily replicable, we encourage the reader to adopt and experiment with the lessons from this paper. Our results suggest that the cellphone – which remains a relatively scarce resource in the developing world – is most effectively utilized with children in more advanced grades. This limitation may be a necessarily evil until we gain a deeper understanding of how to design instructional scaffolds for less well-prepared rural children.

Next, children's tendency to seek help from their neighbors can be channeled productively if the latter are taught to offer help appropriately (e.g. instead of only telling their neighbors the correct spelling, help them to associate and remember the correct spelling). Such peer coaching strategies are especially crucial since cooperative group learning is unfamiliar to many rural children, whose schools (if they attend one) are more likely to implement rote learning. Alternatively, such an after-school program can hire facilitators to provide academically less prepared learners with similar coaching. Most of all, in regions such as rural India where caste politics continue to be pervasive, caste differences ought to be kept in mind when organizing after-school programs.

11 Pimsleur Generator

Audio-only language learning occupies an important niche within the space of technological solutions for developing regions. While most MILLEE games use graphics, we have a priority to explore audio-only applications. The latter are "eyes-free/hands-free" and can be used by a child even when he or she is working. This direction is especially compelling in the case of children who spend long hours in repetitive, manual work, since they can potentially use work time more efficiently for language learning. An ideal solution would, for instance, empower learners in developing regions to develop their conversational skills in a second language in a self-paced approach in out-of-school settings.

One promising candidate is the Pimsleur series of cassette tapes and audio CD-ROMs – arguably the most effective product on the commercial market for developing *conversational skills* in a second language that is supported by previous research (Nation 2001). But two obstacles hinder the applicability of existing Pimsleur units to developing regions. Firstly, their costs are prohibitive: 15 hours of pre-recorded audio cost more than US$200, partly due to the labor by language specialists to develop audio modules that implement the Pimsleur principles for language learning. Secondly, and more important, existing offerings on the commercial market emphasize formal language registers and do not focus enough on the informal language employed in everyday colloquial speech. In fact, existing offerings focus on functional contexts (e.g. asking for directions in urban areas such as New York City) that are irrelevant to the above learner communities, who usually come from less developed settings and are likely to be viewed as unprofitable markets by for-profit content developers.

In this Chapter, we describe an application known as Pimsleur Generator and the computational mechanisms that make it possible. Pimsleur Generator allows third-party developers in developing regions (e.g. language teachers, curriculum developers funded by international development agencies or governments, etc.) to automate the process of generating Pimsleur-like modules that meet local conversational learning needs. Pimsleur Generator is especially appealing for our purposes because the resulting audio output can be distributed on audio CD-ROMs and cassette tapes, which is a relatively accessible technology in underdeveloped regions even as audio cassette tapes are no longer used widely in developed countries. We had originally developed Pimsleur Generator to target English as a Second Language, but we believe its computational abstractions are largely applicable to other target languages.

The rest of this Chapter proceeds in the following sequence. First, we discuss the language learning principles employed in commercial Pimsleur products and present a transcription of its audio-only interface to show how these principles are instantiated in practice. Next, we present an overview of Pimsleur Generator, including the theory and methodology that we used to arrive at it. We give the representation that we used for input data to the Generator, and algorithms that enable the original Pimsleur principles to be supported in computational terms. We conclude with results from an early field evaluation at a village school in India.

11.1 Pimsleur Audio Modules

The Pimsleur series of audio cassettes and audio compact discs are named after the late linguist Paul Pimsleur, who proposed four principles in the 1960's for learning a second language. Pimsleur is arguably the most popular product for people who wish to learn how to speak a second language, according to anecdotal accounts, e.g. customer satisfaction is very high, as seen in the glowing reviews on e-commerce sites such as amazon.com where Pimsleur modules are sold. Most importantly, Nation (2001) surveys the literature on vocabulary learning, which cites several studies of graduated recall in support of Pimsleur's memory schedule and argues that graduate recall should be used.

11.1.1 Language Learning Principles

The principles (Farber 2002) employed in every Pimsleur audio unit are:

Core vocabulary – for most everyday usage of a language, especially in spoken (vs. written) contexts, only a small vocabulary of "core" words is used with high frequency. The implication is that each Pimsleur unit does not have to focus on building an extensive vocabulary and should instead reinforce the learner's retention and mastery of this core vocabulary (related to the graduated interval recall principle below).

Organic learning – the learner is constructing her interlanguage (Selinker 1972), i.e. her cognitive structure of the target language, and is not likely to have a complete understanding of a linguistic element after her first introduction to it. This implies that it is more effective to teach various aspects of the target language (e.g. pronunciation, vocabulary and grammar) in tandem to help her form connections between various aspects, as opposed to teaching vocabulary, etc. in isolation.

Anticipation – while most audio packages involve the learner repeating aloud verbatim after the narrator, the anticipation principle uses an active learning approach in which the learner is prompted with a challenge, after which she is given a short time to think of an answer, before the narrator provides feedback by reading aloud the correct response.

Graduated interval recall – the learner's retention of new words and phrases is reinforced by quizzing her on them repeatedly and in increasing time intervals until they are transferred into long-term memory. She is challenged to recall a word in the target language every few seconds, then minutes, then hours, etc. This principle is used in conjunction with the anticipation principle.

The above principles reinforce one another and appear to be applicable to any pairwise combinations of native and target languages. We next illustrate how the above principles are instantiated in a Pimsleur unit's content and its audio user interface.

11.1.2 Example of a Pimsleur Unit

Each Pimsleur audio unit has a short length (30-40 minutes) consistent with the duration of the human attention span. In every unit, the learner is guided through a conversation in the target language between two speakers, with the narrator providing instructions and explanations in the learner's native language in the background. Figure 11.1 reproduces a portion of a Pimsleur unit which aims to help English speakers learn basic conversational Hindi.

00:00	Man:	Shamakijae.
	Man:	Kya aarp Angrasi shamasthehae?
	Woman:	Geenahee Shriman.
	… (rest of conversation exchange)	
00:20	Narrator:	In the next few minutes, you will not only learn how to understand this conversation, but will also learn to take part in this conversation yourself.
00:28	Narrator:	Imagine that an American man meets a Hindi-speaking woman.
00:36	Narrator:	The American man wants to say, ‘Excuse me.’
	Man:	Shamakijae. (pause)
00:40	Narrator:	The Hindi speaker is going to repeat this, part-by-part, starting from the end. You are to repeat each part after him, trying to make your pronunciation sound exactly like his. Make sure that you repeat aloud after him.
00:52	Hindi speaker:	ae (pause) jae (pause) jae (pause) ki (pause) ki-jae (pause) ki-jae (pause) ki-jae (pause) sha-ma (pause) ma (pause) sha (pause) sha (pause) sha-ma (pause) sha-ma (pause) sha-ma-ki-jae (pause) sha-ma-ki-jae
01:45	Narrator:	How do you say, ‘Excuse me’ in Hindi? (pause)
01:52	Man:	Shamakijae. (pause)
01:56	Man:	Shamakijae. (pause)
02:00	Narrator:	Now he wants to ask her if she understands English.
02:04	Narrator:	First, the word ‘English.’ Listen and repeat.
02:07	Man:	Angrasi. (then proceeds to cover pronunciation for every syllable in ‘Angrasi’, like the 00:52 to 01:45 segment for ‘Shamakijae.’)
	… (quizzes the learner on how to say ‘English’ in Hindi, like the 01:45 to 02:00 segment for ‘Excuse me.’)	
02:53	Narrator:	Say, ‘Excuse me.’ (pause)
02:58	Man:	Shamakijae.
03:03	Narrator:	You should repeat after the speaker gives the correct response, trying to make your pronunciation sound like his. (pause)
03:10	Man:	Shamakijae.
03:15	Narrator:	Say again, ‘English.’ (pause)
03:19	Man:	Angrasi.
03:24	Narrator:	Now he wants to ask, ‘Do you understand?’
03:27	Narrator:	‘Understand’ is ‘shamasthehae.’
	… (covers pronunciation for every syllable in ‘shamasthehae’, like 00:52 to 01:45 segment)	
04:23	Narrator:	Here’s how to say, ‘You understand.’ Listen and repeat.
04:25	Man:	Aarp shamasthehae. (pause)
	…	
04:55	Narrator:	Do you remember how to say, ‘English’? (pause)
05:00	Man:	Angrasi.
05:05	Narrator:	You want to say, ‘You understand English.’ In Hindi, you say literally, ‘You English understand.’ Try it. (pause)
05:18	Man:	Aarp Angrasi shamasthehae. (pause)
05:25	Man:	Aarp Angrasi shamasthehae. (pause)
05:33	Narrator:	Here is the word that makes a question in Hindi. Listen and repeat. Kya. (pause)
	…	
05:46	Narrator:	Any statement can be made into a yes/no question just by putting this in front of it.
	…	
06:04	Narrator:	And now, try to ask, ‘Do you understand?’
06:12	Man:	Kya aarp shamasthehae?

Figure 11.1: Extracts from a transcription of a Pimsleur audio recording for English speakers to learn Hindi.

A typical module begins by presenting a conversation in the target language (00:00 to 00:20 in Figure 11.1), before presenting the context for this conversation (00:20 to 00:36). For each sentence in the conversation exchange, the learner is first taught its meaning before she learns how to express it in the target language. The organic learning principle is employed by teaching vocabulary, pronunciation and grammar on an “as needed" basis. For e.g., she is told that "Shamakijae" is Hindi for "Excuse me" (00:36), after which she is taught to pronounce each syllable in it (00:52).

Similarly, the learner is not taught explicit grammar rules as in a grammar classroom, but rather, grammatical structures whenever she has to understand them in order to construct sentences. For example, she learns that "Kya" can be placed at the start of a sentence to turn it into a Yes/No question (05:33 to 06:12). Likewise, she is taught that Hindi uses the Subject-Object-Verb structure as opposed to English, which uses the SVO form, and is communicated the "You English understand" pattern to help her remember it (05:05 to 5:33). In fact, when she is taught to express "Do you understand English?" in Hindi (02:00), she is led through its subphrases in a step-by-step sequence, i.e. "English" (02:04 to 02:53), followed by "Do you understand?" (03:24), etc. For the latter, she is first taught "understand" (03:27 to 04:23), followed by "You understand" (04:23) and finally "Do you understand?" proper (06:04 and 06:12). In this manner, a Pimsleur module acquaints the learner with the linguistic structures for various grammatical forms in an order that resembles traversing a syntax tree.

The anticipation principle makes use of two different types of challenges in order to quiz the learner on her retention and understanding, as opposed to simply "feeding" her with information in a rote learning fashion. The learner can be prompted to recall a phrase in Hindi (01:45 to 02:00) as well as to construct a sentence using a formulaic form that has been previously introduced to her (06:04 to 06:12). The anticipation principle reflects everyday speaking situations in which the learner needs to recall what is appropriate to say given a situation, as posed in the form of a challenge. The anticipation principle is usually used together with the graduated interval recall principle, so as to quiz the learner on her memory of a word or phrase that she was previously taught. For example, she is quizzed if she remembers how to say "Excuse me" (02:53 to 03:15) and "English" (03:15 to 03:24, and again from 04:55 to 05:05) in Hindi.

More important, the use of the anticipation and graduated interval recall principles makes Pimsleur radically different from most audio materials on the commercial market for language learning products, which are flat “single-rep” (double-rep) audio recordings where each phrase is given in the native language before it is repeated once (twice) in the target language. In contrast, although a Pimsleur unit cannot vary its feedback according to the learner’s response, the anticipation prompts (which are scheduled according to the graduated interval recall principle) allow traditional audio playback-only technologies (which lack support for speech input) to deliver a more *interactive* experience, in which the user engages actively with the audio prompts and derives satisfaction from giving the correct answer to a challenge, (as a tangible sign of learning progress), compared to rote repetition in flat-rep recordings. We argue that the potential for low-cost, audio playback-only user-interfaces to support a user experience with some level of interactivity and active learning, especially in developing world applications, is underappreciated and deserve further exploration.

Lastly, each module is purely audio and has no accompanying visuals due to Pimsleur’s belief that language is primarily speech (vs. writing). This tension is a longstanding debate in literacy theory and is outside the scope of this paper. More important, this characteristic is both a strength and weakness. On one hand, the learner does not need to invest time and effort into learning how to read in a foreign language if all that she wants is to learn how to speak it. On the other hand, visual learners would find it easier to learn the spoken material if there were text and visuals for them to refer to. We believe that the visuals can be distributed on accompanying paper handouts, so that we can continue to focus on low-cost audio-playback technologies.

11.2 Pimsleur Generator

Although commercial Pimsleur modules based on the above principles are highly popular with customers, existing units suffer from the above drawbacks of price and nonrelevance to rural and peri-urban communities. Pimsleur Generator is an application that takes a conversational exchange between two or more parties (Figure 11.2) that lasts a few minutes in duration as its input. It then generates an audio sequence as its output which supports the above principles. The duration of the output audio is substantially longer than that of the input audio mostly because of the graduated interval recall principle, although some of the length increase is due to the anticipation prompts and native language explanations. The output can target the learner's specific needs when the local curriculum developer provides a culturally relevant input.

Raju to Sheela:	When do you go to bed?
Sheela to Raju:	I go to bed at eight in the night.
Raju to Sheela:	When do you wake up?
Sheela to Raju:	I wake up at six in the morning.
Raju to Sheela:	When do you eat breakfast?
Sheela to Raju:	I eat breakfast at eight thirty in the morning.

Figure 11.2: An example of a conversation exchange between a boy (Raju) and a girl (Sheela).

More specifically, the input to Pimsleur Generator contains both the textual and acoustic form of the conversational exchange. Due to the complexity in generating speech with the appropriate intonation, speech generation is currently outside the scope of our work and Pimsleur Generator will rely on the user-supplied acoustic data to "patch together" the corresponding audio output recording. This output can be distributed on storage media such as audio cassette tapes, which is a low-cost technology that is more widespread in underdeveloped regions and requires minimal electricity.

We provide more details on the input representation for Pimsleur Generator and its algorithms in the following two sections. For the remainder of this section, we provide an overview of how we devised the algorithms for Pimsleur Generator and its main data abstraction, i.e. the pattern. (The "pattern" that we refer to in this Chapter refer to those patterns found in formulaic speech, and not design patterns discussed in earlier Chapters in this thesis. We discuss formulaic speech in more detail below.)

11.2.1 Methodology

We began by reviewing 4 units for Hindi speakers to learn ESL and 2 units for English speakers to learn Hindi. In the process, we transcribed 3 of the former units and one of the latter, such as the transcript shown in Figure 11.1, in order to perform a fine-grain analysis of the units. We had varied the native and target languages for the units in our sample so as to observe possible differences depending on the target language. In particular, we chose the Hindi versions of Pimsleur since we have Hindi speakers in our team who could help to decipher their content.

During the transcription process, we carefully took notes about details such as the length of pauses (especially when the learner was given time to respond to the narrator, e.g. a challenge posed under the anticipation principle), the exact words in the audio records, the speaker who uttered those words (i.e. was the speaker the narrator, the male character or the female character?), and whether or not the utterances were made in the direct or indirect speech. It was critical to note these details because Pimsleur Generator needed to reproduce similar utterances.

After that, we reviewed our transcripts to understand how the Pimsleur principles were employed in practice, before devising algorithms that reproduce the implementation of these principles. In some segments of the transcriptions, we could not always identify an apparent or consistent logic that we could express in computational terms. In some of these cases, we took the liberty to impose a uniform logic in line with our knowledge of second language acquisition. In other cases, our algorithms deviated from the transcripts when our observations conflicted with our understanding of learning science principles, e.g. we inserted pauses when the learner was challenged to deploy her interlanguage but was given insufficient time to give a well-thought response.

11.2.2 Formulaic Speech and Patterns

Since a Pimsleur module "tracks" the grammatical structure of each sentence in any conversational exchange, such that the learner is challenged to construct new sentences using grammatical forms that she was taught earlier and receives explanations on those differences in the grammar between the native and target languages, we need an appropriate abstraction to represent grammatical knowledge. We were originally (mis)led by our background in computer science to consider the Chomskian generative grammar and its derivatives (Allan 1995; Jurafsky and Martin 2000), because the grammar rewrite rule is an abstraction which is used widely for representing grammatical knowledge.

However, we abandoned the Chomskian grammar after struggling for two months to express our test input and algorithms using this abstraction. We found that sentences in *oral* conversation are not always "grammatically well-formed" according to the standards of written language, which makes it hard to express grammar rules without an advanced background in linguistics (the classic "ignorance" problem in knowledge representation). Similarly, for any conversational exchange, we need to declare a critical mass of rewrite rules to make its processing possible (the classic "laziness" problem in knowledge representation). There is also the problem of having multiple syntax trees for a sentence, which calls for heuristics to resolve the ambiguity. We expect the average curriculum developer to face similar challenges and decided to adopt a simpler representation, i.e. the pattern.

A pattern is a sentence with possibly one or more moveable parts (e.g. "When do you __?" where the blank can be filled in with a verb phrase contingent on the actual context of interaction). In computational terms, blank(s) in a pattern are variables, such that a pattern that has been matched with a sentence becomes a *pattern instance* that holds variable-value mappings for each variable in the pattern declaration. Patterns are sufficiently expressive for linguistic forms that include Wh-questions, negations, Yes/No questions, etc. Patterns and pattern matching algorithms (Norvig 1992) are not new in computer science. But in developing Pimsleur Generator, it is necessary to extend the usual pattern abstraction in ways that we describe in the next section.

At a more conceptual level, in second language acquisition, patterns are known as “formulaic language.” Whereas the Chomskian grammar offers a top-down view of language in terms of well-defined rules, patterns constitute a bottom-up view of language (Weinert 1995). Patterns are advantageous over rules in that the former allow formulas (e.g. common phrases found in phrasebooks for tourists) and politeness conventions (e.g. honorifics) to be expressed and taught to language learners, without having to explicitly teach them grammar rules. The latter approach makes language lessons resemble formal classes and unappealing. Patterns also allow us to express collocations (i.e. words that co-occur extremely frequently), which some language teachers believe should be taught. In short, there is some theoretical basis that teaching formulas is a reasonably fast approach to equipping language learners with a core repertoire of vocabulary and linguistic forms for common communication needs.

11.3 Input Representation

The declarative language for Pimsleur Generator’s input is based on XML, which we selected for its flexibility and pervasive parser support. Our goal is to provide the average curriculum developer with an intuitive and yet expressive representation for expressing a conversation exchange as input to Pimsleur Generator. (We concede that an XML-based declarative language is inaccessible to the average language curriculum developer, but believe that an appropriate front-end authoring application will alleviate this obstacle so long as there is an intuitive mental model behind the input data representation.) The rationale for certain aspects of our input representation will not be entirely clear until we reach the “Algorithms” section below; we ask the reader to be patient in the meantime.

The input comprises three sections, i.e. sentences that make up a conversation exchange, patterns for sentences in the conversation exchange, and syllable-level pronunciations of words and phrases found in the conversation exchange.

11.3.1 Conversational Script

Sentences in the conversational script are either *specific* to or *independent* of a conversation exchange. For example, in the case of the exchange given in Fig. 1, a specific sentence is "I wake up at six in the morning" while independent sentences based on it might be "I wake up at seven in the morning", "I wake up in the morning", etc. Unlike specific sentences, independent sentences are not directly part of the conversation exchange between the various parties carrying out the dialogue, but provide the curriculum developer with a valuable abstraction for declaring sentences outside the exchange that are variations on sentences in the exchange. Independent sentences are intended to give the learner more exposure to how specific sentences can have their linguistic structures varied (by substituting words, e.g. pronouns with other pronouns) to communicate other meanings. This distinction between specific and independent sentences provides Pimsleur Generator with the necessary metadata to filter out independent sentences when walking the learner through the conversation exchange, e.g. during a debrief at the end of a module.

Specific sentences can also be marked as a *conversation sequence*, i.e. a sequence of questions and responses. This abstraction is meant to support the graduated interval recall principle and will be justified in the next section.

11.3.2 Patterns

Every sentence in the conversation exchange must have a matching pattern. If a sentence has more than one matching pattern, it is matched according to the first pattern. Unlike typical patterns, our pattern declaration includes the speaker and hearer of its corresponding utterance. For e.g., a pattern may be: "Raju to Sheela: When did you wake up, Madam?" such that the pattern matches an utterance if the addressee is female. Encoding the speaker and hearer in every pattern allows us to account for the fact that formulas can not only express syntax, but also pragmatics such as honorifics (e.g. "Madam?") and aspects relevant to the addressor-addressee relationship in terms of formality, gender, seniority, etc. In this way, we argue that the pattern abstraction holds an advantage over abstractions derived from the Chomskian grammar, which is mostly for syntax and less expressive for pragmatics that language learners also need to learn about.

However, we would still like to have a means to provide a phrase structure for the pattern without resorting to explicit grammar rules. Our approach is to "overlay" a syntax tree over the pattern, such that leaf nodes correspond to variable and non-variable elements that could be the subject, verb, object, etc. while non-leaf nodes correspond to higher-level structures such as phrases and sentences. In this way, the language instructor can organize the pattern's elements according to linguistic form(s) without having to declare the underlying grammar rules. This phrase structure provides Pimsleur Generator with a sequence for teaching the elements of a pattern according to phrases.

More importantly, this n-ary tree allows simpler linguistic structures to be composed into more complex structures. Each node can reference audio recordings provided by the curriculum developer in the learner's native language that explain how the node's linguistic structure differs between the native and target languages. These explanations can be devised based on experience with the difficulties that prior learners encounter with these grammatical structures. The developer can also specify the order in which each non-leaf node's children are composed into the corresponding phrase in the native language, in case this sequence differs in the target language. The above metadata is useful, for example, when the learner's native language adopts the Subject-Verb-Object convention whereas she is learning a language that adopts the SOV order, and could benefit from knowing how the target language adopts a different linguistic form.

Each leaf node also contains a reference to the speech audio file (e.g. in the .WAV format) that corresponds to how the respective word is to be pronounced in the target language. However, because speech synthesis is outside our current scope, to avoid having to synthesize individual audio files into a single stream when composing many phrases into a higher order phrase or sentence, every non-leaf node also references the speech audio file that corresponds to how the phrase represented by the respective subtree is pronounced.

Moreover, every leaf (non-leaf) node also references audio files that communicate the meaning of the word (phrase) represented by the corresponding leaf (sub-tree) using the learner's native language, in the direct speech form. These audio recordings are used by Pimsleur Generator to explain the meaning of new words in the target language (e.g. "The English word for 'Yaavaagaa' is 'When'."), and to prompt the learner to recall the phrase's equivalent in the target language (e.g. "How would you say this in English: 'Yaavaagaa'?"). A related observation is that with the use of translations to explain new words in the target language using the learner's native language, it is not always possible to explain the meaning of a new phrase in terms of the meanings of its constituent words, since the meaning of the whole can vary from the meanings of its individual parts. The implication is that it is sometimes necessary to teach a new phrase as an "atomic" unit. In computational terms, a leaf node can thus represent a phrase – and not only a word – that is taught as a unit that cannot be decomposed further.

11.3.3 Pronunciation of Syllables and Variable Phrases

In addition to specifying an audio file for the pronunciation of each syllable, the curriculum developer can also indicate a difficulty index for each syllable, so that the resulting audio module spends more time on those syllables that the learner may find harder to pronounce, e.g. based on factors such as the phonological system of her native language. For simplicity, we only support three difficulty levels at this time: easy, medium (default) and difficult. The curriculum developer can also provide audio recordings that explain how difficult syllables can be pronounced.

For phrases that can be substituted for variables in patterns, their pronunciation audio files are referenced outside of the above pattern declaration, since the variable-value bindings for pattern instances may not be known at declaration time.

11.4 Algorithms

To render this subsection easy for the reader to follow, we first present the basic algorithm of the Pimsleur Generator, before we elaborate on the more complex ways in which the organic learning, graduated interval recall and anticipation principles are supported in algorithmic terms. We will not cover the core vocabulary principle since it is a function of the conversational exchange in the input. The algorithms presented here have undergone simplifications following several rounds of re-thinking and revisions.

11.4.1 Basic Loop

At its core, Pimsleur Generator revolves around the read-evaluation-output loop in Figure 11.3. We will not elaborate on our pattern matching algorithms as they are straightforward adaptations of the algorithms in Norvig (1992). Similarly, we omit book-keeping details, which are basic.

```
for each sentence in conversational script {
   find matching pattern for this sentence;
   get pattern instance’s variable-value bindings;
   get root of this pattern instance’s tree;
   if (learner was taught linguistic structures &
      meaning of all phrases in instance) {
      anticipation (root);
      insert root into priority queue;
      graduated_interval_recall ();
   }
   else {
      teach (root);
      graduated_interval_recall ();
   }
}
```

Figure 11.3: The algorithm for the read-evaluation-output loop that constitutes the core of the Pimsleur Generator.

For each sentence in the conversational script, the sentence is taught if the learner has not previously been taught the meanings of all the phrases in it as well as its grammatical structure.

Otherwise, the learner already has the requisite vocabulary and knowledge of the linguistic form, in which case she is challenged to apply them to construct the sentence, under the anticipation principle. This approach works best when the curriculum developer wants to expose the learner to variations on sentences in the conversational exchange (i.e. *specific sentences*), by providing *independent sentences* as variations. Since the learner may not necessarily construct the sentence correctly or recollect it accurately afterwards, the independent sentence is inserted into a priority queue, i.e. our primary data structure for the graduated interval recall principle, so that she will be quizzed on it again later and given more opportunities to arrive at the correct reply.

11.4.2 Organic Learning Principle

The core algorithm for the organic learning principle is given in Figure 11.4. We distinguish between two cases, i.e. when the item to be taught is “atomic” vs. a more complex item such as a phrase or sentence.

```
teach (node) {
   if (node is a leaf) {
      if (learner has not been taught meaning of
         node’s phrase) {
         say node’s meaning in native language;
         say node’s phrase in target language,
            and asks learner to repeat aloud;
         pause;
         teach_pronunciation (node);
         anticipation (node);
         insert node into priority queue;
      }
   }
   else {
      if (learner has not been taught meaning of
         phrase’s meaning or linguistic structure
         pertaining to node) {
         if (first child of node is a leaf)
            say node’s meaning in native language;
         for each child of node {
            teach (child);
            graduated_interval_recall ();
         }
         if (learner has not encountered linguistic
            structure pertaining to node & node has
            explanation for this structure)
            say node’s explanation in native language;
         anticipation (node);
         insert node into priority queue;
      }
   }
}
```

Figure 11.4: The algorithm for teaching according to the organic learning principle.

In the first case, Pimsleur Generator builds the learner’s vocabulary by explaining the meaning of a new word (or phrase) in her native language before conveying the word itself in the target language, e.g.: “The Kannada word for ‘When’ is ‘Yaavaagaa’” (assuming that a Kannada speaker is attempting to learn English). Next, it shows how this word is to be pronounced. It then challenges the learner (i.e. the anticipation principle) if she remembers it, e.g.: “How do you say in English, ‘Yaavaagaa?’” There is a pause for her to think and to say her response aloud, after which she hears the correct response, e.g. “When.” This word is then noted for subsequent graduated interval recall.

In the case of a more complex item such as a sentence or phrase, Pimsleur Generator first conveys its meaning and explains how its grammatical structure differs from that of a similar phrase in the learner's native language. (The observant reader will notice that the phrase's meaning is conveyed only if the first child is a leaf node. The reason is that it is possible for the algorithm to keep traversing the left child of each node for a while before it actually reaches a leaf. As such, it would continue to convey the meaning of each clause along the recursion tree without actually giving the words in the target language necessary for articulating the respective semantics. For simplicity, we decided to skip communicating this meaning in this situation since it will be conveyed anyway during the anticipation principle step after returning from the recursive calls for the child nodes.) It then tries to teach (recursively) the meaning of every word or phrase that makes up this item. She is then challenged to put the item's constituent words and sub-phrases together based on the appropriate grammatical structure. This item is also noted for subsequent graduated interval recall.

```
teach_pronunciation (leaf) {
   for each syllable in leaf {
      if (syllable is easy)
         count = 1;
      else
         count = 2;
      repeat for count times {
         say syllable;
         pause;
      }
   }
   say leaf's phrase;
   pause;
   for each syllable in leaf that is difficult {
      say explanation for syllable;
      pause;
      say syllable;
      pause;
   }
}
```

Figure 11.5: The algorithm for teaching pronunciation.

The learner is also taught to pronounce every new word based on the algorithm in Figure 11.5. Basically, every syllable is repeated once (twice) if it is marked as easy (medium or difficult). The entire word is then repeated as a whole. For those syllables that are difficult, the learner receives more explanation on how to pronounce them effectively.

In conclusion, vocabulary, pronunciation and grammar are integrated in our algorithm for Pimsleur's organic learning principle. Furthermore, words and phrases are skipped if they appear in earlier sentences in the conversational script.

11.4.3 Graduated Interval Recall Principle

A good time to quiz the learner on a word or phrase which she was previously taught is when she has just been taught a new word (see Figures 11.3 and 11.4), and immediately before she is taught the next word or phrase in the conversation script.

The algorithm is given in Figure 11.6. The primary data structure is a priority queue which stores nodes pertaining to words, phrases and sentences that are ordered according to a field for timestamp, i.e. the time for the learner to be quizzed on the respective item. As such, when an item is inserted into the queue (as indicated in the above pseudocode), what happens is that multiple instances are actually inserted, such that the timestamp for each instance is computed based on a memory schedule. We currently adopt Pimsleur's 1967 schedule, in which items are quizzed at the 5-second, 25-second, 2-minute and 10-minute graduated intervals.

```
graduated_interval_recall () {
   while (there is node at head of priority queue)
      if (node's timestamp <= current time) {
         anticipation (node);
         if (node is first sentence in a
            conversation sequence)
            for each subsequent sentence in sequence
               if (learner was taught linguistic
                  structures & meaning of all phrases
                  in sentence)
                  anticipation (sentence's root);
               else
                  break;
         remove node from priority queue;
      }
}
```

Figure 11.6. The algorithm for the graduated interval recall principle.

Three issues influenced the design of our algorithm. First, *backlog* in the queue, i.e. there is more than one item to quiz the learner on whose timestamps have expired (that is, precede the current time). We could spread the backlog over the entire module but decided to clear the backlog immediately, since it was overdue for the learner to be challenged on these items. Second, the issue of *hierarchical promotion*, i.e. once the learner has been taught a phrase, subsequent recall intervals should focus on the phrase and not its constituent words and phrases. As such, when a node is inserted into the queue, existing nodes in the queue are removed and substituted with the former node if existing nodes in the queue are the former's descendents.

Third, our algorithm not only quizzes the learner on isolated sentences, but also quizzes her on the sentences that follow in the same *conversation sequence* if they had already been taught. In the example of the conversation exchange in Fig. 2, after the learner is quizzed on how to say "When do you go to bed?" in the target language, she is also quizzed on how to respond to this question in the target language: "I go to bed at eight in the night." The aim is to reinforce her ability to *respond appropriately* when faced with questions, instead of merely focusing on sentences in a conversational exchange in isolation without drawing her attention to the connection between related sentences.

11.4.4 Anticipation Principle

There are two cases for the anticipation principle in Figure 11.7, depending on whether the item to quiz the learner on is "atomic" or a higher-level structure such as a sentence or phrase. In the former case, the learner is challenged to recall how to say a word or phrase in the target language, and is given a short pause to think before she receives feedback on the correct answer. The latter case is similar to the former, except that the learner is conveyed a sentence or phrase in her native language, after which she is challenged to form its corresponding sentence or phrase in the target language.

```
anticipation (node) {
   if (node is a leaf) {
      asks if learner remembers node’s phrase,
         using either direct or indirect speech;
      pause;
      say node’s phrase in target language;
      pause;
      if (this is first or second time that learner
         is quizzed on node) {
         if (number of medium or difficult syllables
            in node>threshold_pronunciation_accuracy)
            teach_pronunciation (node);
         else
            say node’s phrase in target language;
         pause;
      }
   }
   else {
      asks learner to say node’s phrase,
         using direct speech;
      pause;
      say node’s phrase in target language;
      pause;
      if (this is first or second time that learner
         is quizzed on node) {
         say node’s phrase in target language;
         pause;
      }
   }
}
```

Figure 11.7: The algorithm for the anticipation principle.

More importantly, our algorithm design for the latter case is consistent with psycholinguistic theories of second language acquisition. For example, Swain (1985) argues that having learners engage in output production provides them with opportunities to develop their interlanguage. Similarly, learners employ strategies such as transferring knowledge of grammatical structures from their native language and generalizing on their growing linguistic knowledge in the target language (Selinker 1972). Our algorithms build on these strategies by giving the learner opportunities to construct new sentences based on linguistic structures that she was taught earlier. In fact, when the algorithm repeats the pronunciation (former case) and syntactical structure (latter case) for an item if it is the first or second time that the learner is quizzed on it, we actually give the learner the chance to develop her phonological or syntactic accuracy. This algorithm is based on Skehan's (1998a; 1998b) cognitive framework for task-based language teaching, which argues for balancing the development of the learner's fluency with accuracy, such that the first round of feedback in the anticipation algorithm aims to develop her fluency, while the repetition draws her attention to the same item to help her notice the salient features essential for accuracy.

11.5 Initial Experience

We need to understand how expressive our abstractions are for representing real-world conversational syllabi and how users in underdeveloped regions react to audio units created using Pimsleur Generator.

11.5.1 Expressiveness

We evaluate the expressiveness of our data abstractions by attempting to represent a set of input conversational scripts using them. We make use of conversation scripts developed by our NGO partners in India, who have been teaching ESL in India for over three decades. These scripts are localized to meet the learning needs of Kannada-speaking children in grades 1-8 who live in villages in the state of Karnataka.

Our review of the above conversation scripts indicates that our augmented pattern abstraction is sufficiently expressive for the various linguistic forms found in the given syllabi. For a more rigorous check, we selected scripts from the two extreme endpoints (i.e. grades 1 and 8) and digitized them for testing with Pimsleur Generator. In total, we digitized 2 units for grade 1 and one unit for grade 8. While the scripts for grade 1 were simple, the script for grade 8 was far more complex. The latter contained grammatical structures for English that differed from Kannada and atomic phrases whose constituent words cannot be explained in English.

11.5.2 Language Learning

We deployed the above 2 units that we digitized for 1st grade students at a village school. A total of 24 children (11 male and 13 female) participated in the study. To keep logistics manageable, we deployed our audio modules on the cellphones instead of using audio tape players, since we were already using the former in the broader study. Students listened to the audio playback using hands-free kits for the cellphones. On average, each student listened to each of the 2 modules for 1-2 times. The reader who is familiar with Pimsleur may object to an evaluation with child learners since the original Pimsleur is targeted at adults. However, an earlier study with children in the urban slums in India involving commercial Pimsleur units showed promising results. We found that these children reacted well to the Pimsleur principles, but that commercial units target situations in the developed world that are not necessarily culturally meaningful to our target users. In turn, this earlier study motivated our work on Pimsleur Generator.

Participants seemed to enjoy the audio playback-only user-interface. Each module was about 16 minutes long and was able to hold the children’s attention for this duration. In fact, users were spontaneous in repeating the Kannada and English portions aloud after the audio units. Their facial and nonverbal expressions suggested that they were having fun with the experience. The participants also seemed to readily understand the content in the audio modules. For example, when a boy repeated "I (am a) girl" aloud in English after the module, he covered his face with his hand in a shocked expression; he seemed to comprehend that he was not a girl despite what he was repeating after the module. The lack of support for speech synthesis did not appear to result in any serious usability problem.

However, we observed that participants experienced three usability problems with the audio playback user-interface. First, children struggled to pronounce the English words in the modules. In fact, we observed that one girl was writing the phonetics for the English words on the blackboard in the Kannada alphabet. Second, several participants could not understand some of the Kannada recordings due to accent differences between urban and rural Kannada. The audio was recorded by the second author who speaks Kannada as her native language and who lives in a city district less than an hour's drive from the village school. Yet, we learned that accent differences can be somewhat distinct within such a small geographic region. Third, most users remained silent and did not respond to the challenge prompts.

11.6 Discussion and Future Work

We plan to extend our XML-based input representation to support phonetic transliterations (using the alphabet of the native language), so that Pimsleur Generator can generate paper handouts that help learners follow the pronunciations of phrases in the corresponding audio units. We did not provide this feature initially because of concerns about low native language literacy, but the above experience suggests that native literacy rates can be adequately high in some villages to make transliterations feasible.

To address the problem with accent differences, we plan to record input audio clips in future with rural children whose ESL competency is slightly higher than the target learners. In fact, this approach takes us closer to our vision of giving communities the means to create their own learning tools.

We found the lack of response to the anticipation prompts to be somewhat surprising, since our earlier study involving commercial Pimsleur units with 10-year-old urban slums children showed a significantly higher level of engagement with the challenge prompts. The youth of our participants in the above study could be a major factor. There could also be other responsible social factors. For e.g. participants were not afraid of the second author and continually teased her; they may not be predisposed to following instructions that were recorded in her voice. We believe that recording the audio instructions using a more confident or authoritative-sounding voice, such as that of the participants' school teachers, may have been more effective. We could have also provided participants with more directions about how to respond to the prompts when we introduced Pimsleur Generator to them.

In future fieldwork, we plan to evaluate the audio material produced by Pimsleur Generator with a wider age range that includes adult learners. We also plan to conduct learning assessments after resolving interface-related issues.

12 Related Work

Most work on computer-aided learning – much less technology-assisted language learning – in the developing world does not explore the convenience that the cellphone's mobility offers. Banerjee et al. (2005) report on a large-scale randomized evaluation with mathematics learning games on desktop computers, which was carried out over 2 years with urban slums children in India. Mitra et al. (2003) describe a study in India with slums children over a 5-month period, in which a "hole-in-the-wall" public computer was installed with speech-to-text software for English learning. Dias et al. (2005) describe a project in Ghana around computer-based tutoring software for improving reading skills in English as a Second Language. Kothari's karaoke-like approach of using "same language subtitling" in Bollywood movies targets native language – not second language – literacy using television sets (2004).

With the cellphone's increasing ubiquity in Africa, Brown (2003) argues that it is timely to envision a future where the cellphone plays a pivotal role in education in Africa. The only education-related projects we know of in the developing world that leverage the cellphone are Islam et al. (2006) in Bangladesh, and Librero et al. (2007) in Mongolia and the Philippines. Both projects rely heavily on the Short Messaging Service (SMS) infrastructure to work and target university students, unlike our approach, in which the cellphone is the runtime device for game software applications.

To date, Horowitz et al. (2006) is the only study we know of which examines the applicability of the cellphone for promoting literacy. However, this study took place in an industrialized country – the United States – even though participants included households below the poverty line. In the study, Sesame Street videos that target the English alphabet were streamed over-the-air to preschool children over cellphones throughout an 8-week period. Our work is thus a contribution to the literature as the first learning assessment on cellphone-based language learning software in developing regions.

He et al. (2008) describe a 2-year randomized evaluation of a LeapPad-like device that supports custom software modules for English language learning. This interactive system involves a paper book that is attached to a stylus and which supports audio output. It lacks a visual display, unlike a cellphone, but overlaps with our goal of making literacy learning more accessible in the developing world using portable devices. We expect some convergence between our respective projects in the near future.

Among the learning technologies designed to meet the challenges in developing countries, one of the most novel solutions is the multiple-mice computer (Pawar et al. 2007). It was intended for collocated learning by a group of children gathered around a computer, with a mouse input device for each child. This work has since been extended to distance learning by Moraveji et al. (2008). The evaluations in both papers were short-term, and underscored the tremendous difficulty in conducting a learning assessment for a novel educational technology over a substantial timeframe when it is not yet integrated into the everyday operations of a formal entity.

PACE is inspired by previous work on learning objects (Wiley 2002) and reusable software kits (Roschelle and Jackiw 1999). But there is very little demonstration of their benefits in real-world case studies. On the other hand, this thesis gives concrete examples of how their ideas on reuse and modular architecture have been applied and found to be indispensable for rapid prototyping and localization in a developing region.

Adapting software for a local context is often described as a two-step process of *internationalization* and *localization* in which (i) culturally independent components are isolated from culturally dependent components, after which (ii) the latter is adapted for the local context. Prior work on localization (e.g. Marcus and Gould 2000; Smith et al. 2004; Yeo 2001) focuses mostly on localization at the level of the user-interface, whereas the PACE framework covers the localization of both the user-interface and educational content. For our purpose, a design pattern is an abstract representation that captures a best practice currently employed in state-of-the-art language learning software. The pattern therefore enables a best practice to be isolated and represented as a culturally independent component that can be adapted for various local contexts.

Design patterns have been employed in the domains of urban planning (Alexander 1977), software engineering (Gamma et al. 1995), games design (Björk and Holopainen 2005), interaction design (Borchers 2001), website design (Van Duyne, Landay and Hong 2002), computer science education (Pedagogical Patterns), science education (Linn and Eylon 2006) and other social applications of computing (Liberating Voices!). We have neither encountered design patterns in language learning instructional design nor earlier localization work.

Our review of commercial language learning products – both software and audio recordings – is preceded by earlier work. Clegg, Ogange and Rodseth 2003 reviewed a sample of digital materials developed for English language learning in Africa, the United Kingdom and North America. However, their goal was to determine the extent to which the sample is appropriate for African learning needs. Other studies such as Wood 2001 examined the extent to which commercial language learning applications are designed to reflect well-accepted educational principles and theories. On the other hand, in taking the design pattern approach, we have not only reviewed the extent to which the applications in our sample are consistent with prevalent language learning theories. More importantly, we recognize that there is room for improvement in existing applications – both in terms of their theoretical underpinnings and fluidity in the social settings where they are used – and that the pattern is an appropriate abstraction for representing existing best practices. This attempt to capture existing knowledge constitutes the crucial first step for codifying this knowledge explicitly, so that it can be improved and reused.

Where the field of videogame studies is concerned, Björk and Holopainen (2005) is the most comprehensive compendium of game design patterns. These design patterns can be used in the creative design of new games according to a structured brainstorming process, as well as to analyze existing games. Our work builds on Björk and Holopainen

(2005) in two ways. First, Björk and Holopainen (2005) and Kreimeier (2002) make the case for game design patterns but do not provide case studies where game design patterns have been applied in actual real-world settings, whereas we present a comparative study that examines the extent to which such patterns are useful in practice. Second, we distill lessons on the contextual and cultural factors that must be considered when applying the patterns to designing games for rural children, especially those in India. After all, design patterns are mostly abstract – and hence decontextualized – representations of common solutions to frequent problems.

Lazzaro (2004) identifies four "keys," namely, Hard Fun, Easy Fun, Altered States and The People Factor, which the game designer can use in tandem for unlocking an engaging play experience. Gee (2005) argues that games are fun because they give the player opportunities to grow by learning and solving problems. The player can learn and solve problems in the game when he can assume an identity to act in it. At the most basic level, the player can attribute meaning to his play experience when story elements such as actions, states, events and characters are merged with the game's abstract rule system. Malone (1980) proposes that games become fun when they are designed to challenge the player, are situated in a fantasy setting and arouse the player's curiosity. Norman (2003) discusses how a designer can design for pleasure by combining design at the visceral, behavioral and reflective levels. In our view, all of the above work complement design patterns in that they are less concrete than individual design patterns and thus provide the game designer with higher level principles that guide him in applying the patterns.

On the other hand, Salen and Zimmerman (2003) believe that game design is a second-order design problem such that it is not possible to design the play experience directly. Instead, all that the designer can design is the game's core mechanic, and that the play experience emerges from this structure. Notwithstanding their position, we argue that game design patterns comprise the basic building blocks which the designer can use to compose the core mechanics for an engaging game, particularly when these elements are already present in previously successful games.

The first work to argue that games can be used for education was the classic by Abt (1970). Its content is dated as of today, and in any case, it does not offer practical advice for designers. Instead, there is more recent work in the emerging "serious games for education" literature such as Bergeron (2006) as well as Michael and Chen (2005) that provides a high-level survey of e-learning games but is not directly related to design. Other work is mostly anecdotal (e.g. Prensky 2000) or examines e-learning games from perspectives of non-designers such as epistemology (Shaffer 2006) or the socio-historical analysis of literacy development (Gee, Hawisher and Self 2007), or is limited to specific game genres such as simulations (Aldrich 2003). In our opinion, the most relevant recent work is (Gee 2003), which offers a set of learning science principles that could be applied to game design. Its limitation, however, is that it does not connect these principles with the potential for videogames to motivate the learner. And in any case, none of the above references focus on e-learning games that target low-income users in developing regions.

We have not come across previous studies on how rural children in India – or another developing region – engage with mobile games in terms of interaction design and gameplay. We believe that this user group deserves attention from videogame designers due to the out-of-school learning opportunities that we have raised. Taking a longer-term view, a recent news article reports on the phenomenal sales of mobile games in India, and adds that although mobile games in India are currently targeted at urban consumers in India, "the real market … lies in rural areas" (The Times of India 2006).

Sutton-Smith's work on children games (Sutton-Smith 1976; Herron and Sutton-Smith 1971) is among the most established in the field, but it does not focus on the games that children play in villages in developing regions. Instead, work such as Sutton-Smith (1976) and Herron and Sutton-Smith (1971) focuses on traditional games that children in industrialized countries play. With the exception of Narayan (1995), most descriptions of traditional games either do not cover South Asia (e.g. Bell 1979; Topics Online Magazine Vol. 11) or cover urbanized variations of the same games (c.f. India Visit Information; Wikipedia). Our work hence contributes to existing documented accounts of traditional village games. Narayan (1995) details the traditional games played by tribal children in India, but does not attempt any analyses to inform digital game design.

We are the not the first researchers to perform a comparison between electronic and traditional games using design patterns in the compendium by Björk and Holopainen (2005). Loh and Seah identified the differences between WarCraft III and the Risk board game, to explore how successful traditional games can be adapted into digital games (2006). Our work has a significantly broader scope than Loh and Seah (2006), in that the traditional games in our comparison goes beyond one game in the tabletop family of traditional games, and in fact cover 28 village games which belong to other families.

Björk and Holopainen (2005) is not the only work on design patterns for games. For example, there is the pattern library by Folmer. But we do not use Folmer in our comparison since it emphasizes game design patterns for accessibility and usability at the user-interface level, whereas our focus is at the level of gameplay mechanics and rules. Similarly, there is other work on game design patterns (e.g. Lankoski and Björk 2007a; Lankoski and Björk 2007b; Mor et al. 2007; Peitz and Björk 2007), but our comparison does not use these patterns because they target specialized game topics such as believable non-player characters, social networks, mathematical games and pervasive games, which are less relevant to poor children in the developing world. Björk and Holopainen (2005) is currently the most comprehensive encyclopedia on *typical* patterns found in Western video and traditional games. Performing a comparison using the patterns in Björk and Holopainen (2005) is thus more practical for the purpose of drawing design ideas that are more readily applicable to affordable gaming technologies in the developing world.

Work in the emerging field of videogame studies (e.g. Fullerton 2008; Salen and Zimmerman 2003) provides a thorough treatment on game mechanics, including those in non-electronic games. But it does not systematically study how mechanics in videogames and traditional games differ. Finally, designing videogames based on traditional village games is not the same as designing sports videogames (e.g. Rollings and Adams 2003), in that the former have mechanics that reflect a higher degree of informality.

Finally, among projects whose goals are to promote children's education in the developing world using technology, Pawar et al. (2007) aims to design e-learning games on multiple-mice desktop computers while the Playpower Foundation (Lomas, Douglass and Rehn 2008) focuses on building an open-source development community around a $12 TV-computer. The Azim Premji Foundation (2004a) and Pratham (2008) target similar demographics in their computer-aided learning initiatives. The design tools and results from this thesis will inform the above projects as well as new initiatives by other actors who aim to improve education in the developing world through e-learning games.

13 Conclusion and Future Work

This is the customary Chapter in a doctoral thesis in which the author is expected to end with concluding remarks. At the time of writing, the MILLEE project is in its fifth year. But this is hardly the opportune moment for a conclusion in the usual sense of the word. Instead, we feel that our progress to date has only been a "warm-up," and that the *real* project is only truly starting now. That is, a huge distance remains between where we currently are at and getting MILLEE games into the hands of millions of underprivileged children worldwide, before we can even conceive of making a dent on global poverty.

In summary, the challenges in the MILLEE research project are to (i) draw on sound learning principles, (ii) provide design patterns which allow the twin concerns of entertainment and learning to be integrated, and (iii) account for cultural differences in children in developing regions. As such, the end product of our work is not the games themselves but a suite of tools and methods for localizing and scaling the games for other underserved learning contexts. In the past 4½ years, we have developed and refined some foundational games on the cellphone platform that target English as a Second Language, and perfected a framework (called PACE) which streamlines the development of new content. We do not believe it is practical to develop "one-size-fits-all" games for general or even national use. Rather, the games need to be tailored to local customs and practices, and in particular, we adapt rural children's games with which they are already familiar.

In this doctoral thesis, we have taken the reader on a journey that has crossed the boundaries of many traditional disciplines. It is a journey that reflects our development at both the personal and intellectual levels. The reader will recall that our journey began in the villages and urban slums in India, where we were conducting needs assessment and exploratory studies to narrow down on a thesis topic which is relevant to local needs. We narrate the process that we took in search of a compelling thesis topic to reinforce the message, as well as to encourage future researchers in the field of "information and communication technologies and international development" (ICTD), that initial ideas usually turn out to be non-culturally appropriate, and that it requires perseverance and commitment to arrive at a thesis topic that is consistent with local conditions. At the more practical level, these formative experiences taught us about the ground conditions in rural schools and how to address the challenges in performing field research in a developing country setting.

As a thesis that first and foremost focuses on literacy as opposed to technology for technology's sake, it is crucial to ground our work in the psychology of reading processes and second language acquisition. The reader ventured with us into some psycholinguistic and sociocultural theories of language learning, in which we showed how they converge on the task-based language teaching framework. We argued that Christopher Alexander's design pattern is an appropriate abstraction for representing best practices in design knowledge. We demonstrated how design patterns could be applied in the domains of language pedagogy and game design, and gave a walkthrough of our process for distilling language pedagogy patterns from successful commercial language learning applications. Most important, we expanded on the design pattern to come up with the PACE (Pattern-Activity-Curriculum-Exercise) framework, which allows the representation of design knowledge from one context for adaptation to other local contexts.

The reader was with us when we crossed the disciplinary threshold from language and literacy studies back to human-computer interaction. We designed and prototyped a total of 10 mobile games for ESL learning, which we evaluated in formative studies with children from the urban slums and villages in India. The reader shared our frustrations as we learned that our designs, which were relatively more usable by children in the slums, were less intuitive to rural children who had much less experience with high technology. The reader accompanied us on each step of the way as we iterated on our designs to make them more meaningful to rural children. In the process, we observed that our experiences with Western games unconsciously crept into our game designs, thus rendering them less consistent with the expectations and understandings that rural children have about games.

We sought to understand more about the traditional village games that rural children play, and what made village games qualitatively different from contemporary Western games. The reader came with us as we crossed disciplinary boundaries once more, this time, into videogame studies.

Armed with all this knowledge, we made our first foray into running a semester-long summative evaluation. This step is particularly crucial for obtaining more conclusive evidence on the extent to which cellphone-based ESL learning can make a difference in the education of disadvantaged children in rural areas who have no alternative access to quality ESL instruction. Our journey ends with Pimsleur Generator, which is the work on audio-only learning in this thesis that perhaps best fits the mold of "traditional" computer science.

Most important, in this journey, we learned to work with local partners at various levels of the adoption hierarchy. The adage "think global, act local" has never been truer. In the remainder of this Chapter, we share some of our lessons on partnerships and doing research in human-computer interaction in a developing country context. We conclude by discussing some of our future research steps.

13.1 Partnerships

Building strong relationships with reliable local partners is critical to the success of research projects in human-computer interaction that target social empowerment in the developing world. In our experience, partnerships can occur at various levels, including government, corporations, non-government organizations, local undergraduates and end-user communities. However, we only share our experiences for the last three categories of entities in the rest of this section since we have only begun to engage with the first two categories of actors at this stage in the MILLEE project.

In any research project on educational technologies, summative assessments are crucial for demonstrating the extent to which the proposed solutions result in learning outcomes. But in our experience, summative studies turned out to be more complex than expected to carry out due to the cross-border nature of the financial transactions involved, which entail having to comply with the laws in both the sponsor and host countries. The length and scale of any summative assessment meant that enough funds have to be at hand in the host country to cover staffing and other costs. On one hand, it is not possible to pay local staff salaries *directly* from a university that is based in the United States such as the University of California, Berkeley because we are not set up to employ project personnel who are non-US taxpayers. This constraint indicated the need for an institution with a local presence in India who could undertake the funds disbursements, in addition to its regular responsibilities as an NGO collaborator.

However, in our attempts to transfer funds to a partner organization in India, we learned that there are laws that regulate money transfers from overseas. Specifically, the Foreign Contribution (Regulation) Act was established in 1976 to ensure that funds from foreign sources are not used in ways that undermine the interests of India as a sovereign nation. For our purposes, when the MILLEE project reached the stage in the fourth year when we were ready to commence a summative study, we had to collaborate on it with a local partner organization that was approved by the Indian government to receive foreign funds under FCRA. The procedure of signing a Memorandum of Understanding between UC Berkeley and our NGO partner to legitimize this collaboration in the eyes of the FCRA took a total of 15 months. This lengthy duration accounted for various false starts in the administrative process, which was only expected when a research collaboration of this nature that involved a substantial pilot deployment with actual learners – especially in a developing country – was fairly unprecedented in a computer science department. The lesson is that although it is conventional in the ICTD community to think about the size of a NGO in terms of its financial resources, manpower pool and scale of operations, another metric for an NGO's size is the forms of regulatory approvals that it has received from the federal and state governments.

Since the MILLEE project started in mid-2004, we have had highly motivated UC Berkeley undergraduates accompany us to India for field studies, so that there are enough researchers present to manage the user study sessions and collect field data. But there are also tremendous benefits to involving local undergraduates from India in our fieldwork. Familiarity with the local languages, cultures, and institutions is an enormous advantage when undertaking field research. Furthermore, the participation of local students creates cross-cultural learning opportunities for the non-Indian members in the team, and it is this experience which is important in an increasingly globalized world.

But the benefit that is perhaps least appreciated is the rapport that we can build with community partners. Although it is difficult to validate for sure from our interactions with existing and prospective NGO partners, we nevertheless believe that local partners such as NGOs may understandably perceive a project to be driven by "outsiders" and hence hesitate to lend their complete support. On the other hand, an active role by local undergraduates helps to reinforce our message that we are committed to working with the locals and giving back to the local community by facilitating cross-institutional learning opportunities.

What about involving other locals, such as graduate students or professional field researchers in a cross-cultural team instead? There is no doubting the value that more-experienced personnel can bring to a cross-cultural design project. However, in previous field research that we have conducted in other developing regions, we have observed that lowly-literate users in marginalized communities have more often than not perceived educated visitors like us as "outside experts" or authority figures, and have felt more comfortable in learning about the technology that we were attempting to investigate – not from us – but from their peers whom they view as their equals. Hence, local youths and undergraduates can often find it easier to overcome these barriers.

To date, we have worked with nine undergraduates from reputable universities in India. Their familiarity with the local languages and cultural norms means that they can be invaluable in providing local support. In our view, however, their participation should not be restricted to the roles of interpreters and cultural guides. Rather, by investing the time and energy to mentor them on those skills related to conducting user studies, we can develop their capabilities to undertake greater responsibilities as research assistants.

From our experience, when recruiting local undergraduates, the first criterion to consider is their level of commitment to community service or rural development. More important, some universities in India encourage or make community service a mandatory requirement for graduation. As such, it becomes imperative to learn what a candidate's responsibilities and contributions in earlier community service projects were, in addition to simply looking for evidence of prior volunteer experience. For instance, several local undergrads who work with us have described how they lived for weeks under austere conditions in villages when they were administering baseline surveys for NGOs.

Next, while traits such as technical competence and academic achievement are no doubt important, the ability to engage with end-users, stakeholders, and NGO partners on a mature level is the most essential quality. For example, the ability to enlist the support of influential figures in the community is instrumental to encouraging the participation of community members in our user studies. While our NGO partners had to introduce us to these leaders at the beginning, language barriers had prevented us from developing these relationships to the fullest possible extent. In the end, it was the local undergraduates and personnel in our research team who helped by liaising tactfully and professionally with community stakeholders in their native languages.

We did not only have to manage relationships with adults in the community, but we also had to establish rapport with the children themselves. One issue arose when the children at one of the rural schools wanted to take us to visit the temple in their village. In fact, they became upset with us and lost their enthusiasm for the user study until we obliged them. On our way to the temple, the children insisted that we skirt around a part of the village which was inhabited by some of their classmates who are "dalits" (i.e. the "untouchables" caste). The local undergraduates in our team were sensitive to the cultural implications behind this action; they decided that on our return journey, the team should walk through the area inhabited by the dalits. This gesture cheered up the dalit children who were disappointed that their peers wanted us to avoid their residences and helped to promote a more inclusive atmosphere in our user study.

We do not claim, however, that local team members never encounter barriers in interacting with end-user communities. However, local team members are arguably more informed and sensitive towards the local culture to work around these obstacles. For example, in Chapter 9, we described our contextual studies of the traditional games that rural children play in the villages. Our goal behind these studies was to design e-learning games patterned after these games. It was the local undergraduates who first noticed that the children were only showing us the urban games in India. It seemed that games were a marker of social identity, to the extent that our young informants feared that we, as urban dwellers, might look down on their village games. The local undergraduates eventually coaxed the children to reenact how they played their everyday games by enthusiastically describing those village games that we knew about.

The MILLEE project has come to be spread across three geographic locations, namely: Berkeley (California, USA) where systems design and data analysis takes place, Dhirubhai Ambani Institute for Information and Communication Technology (Gujarat, India) where most of the computer programming activities are carried out, and Lucknow (Uttar Pradesh, India) which is the base for the curriculum development and deployment testbed described in Chapter 10. Development of the game applications and monitoring of the pilot activities continues to be coordinated across two time zones. In this current division of labor, UC Berkeley undergraduates contribute to design brainstorming while local Indian undergraduates at DA-IICT work on application development. To a large degree, this division of responsibilities seems attributed to differences in the institutions for higher education between USA and India. We observed that Berkeley undergraduates are better prepared to bounce design ideas off, possibly due to the liberal arts education system in an American university environment and availability of an undergraduate class on human-computer interaction. On the other hand, undergraduates in India have a head-start in their engineering curriculum, which stresses technical skills right from the start of their first year in college.

From the broader perspective, human-computer interaction is only beginning to take root – if at all – in the mainstream undergraduate curriculum in India and many other so-called developing countries. Most undergraduate students would hence graduate and enter the workplace or graduate school without having been exposed to HCI. As such, in relation to the wider implications for HCI education in the developing world, we believe that involving local undergraduates from developing regions in international development projects is one way to give them meaningful exposure to HCI that they are not likely to have otherwise. Their participation also provides them with an outlet for their creativity and ambition to achieve.

We encourage other researchers, educators, and professionals in our community to explore similar symbiotic arrangements. Our efforts may not be on the same scale as formal institutionalized approaches to HCI education in the developing world, but they may nevertheless contribute to local capacity building in HCI and facilitate the eventual adoption of HCI in mainstream curricula. More important, such efforts may set the stage for nurturing the next generation of HCI practitioners who can contribute to the growth of our community with a more diverse international outlook and representation. In doing so, we lay the foundation, both for HCI and computing as a field that is making increasingly pervasive social impacts, to become a more positive force in economic development.

13.2 Future Work

With all that we have learned about managing a three-location project, conducting field studies with underserved communities, second language acquisition, game studies, technology design and learning assessment in the context of an international development project in the last 4½ years, we are now more confident of scaling greater challenges. It is time to take bolder (and wiser) steps as we seek to bring the MILLEE project forward as well as start new projects under the wider umbrella of "literacy technologies."

First, our observations of how the less academically advanced children struggled with the games in our summative assessment (Chapter 10) implied that they needed better scaffolding support in the instructional design. We plan to investigate how we can design adaptive learning games that track learner performance in out-of-school settings, in which scaffolding supports such as hints are designed with more sensitivity so as not to diminish the learner's sense of self-esteem or achievement. At a more fundamental level, we have to understand the cognitive implications of schooling (Scribner and Cole 1981), and how to take these implications into account when designing instructional applications for out-of-school children who have been considerably less socialized to school-based practices.

Second, we need to benchmark learning outcomes from MILLEE games against existing standardized English tests in India, in order to obtain additional support from parents, teachers, policy makers and NGOs. In line with this goal, we plan to scale-up our summative evaluation from one village site to a significantly larger sample involving 800 children in 20 village sites. A pilot deployment on this scale is critical for attracting the attention of governments and major non-government organizations. If this study is able to replicate the positive learning benefits that we have found in earlier smaller-scale studies, we can make a compelling case to governments and major NGOs to adopt our work. Strong learning results will also help to convince cellphone manufacturers and wireless carriers that a customized, ultra-low cost cellphone model can be developed specifically for rural education. Similarly, positive results will enable us to make a stronger case to third-party non-profit Indian content developers such as the Azim Premji Foundation and Sesame Workshop India on the feasibility of targeting the cellphone platform in their computer-aided learning initiatives. We believe that having third-party developers adopt our tools and methods, so as to adapt them for their local needs, is one way for our work to scale.

Third, as we prepare to conduct a large-scale randomized evaluation with 800 rural students in a classroom setting, we remain cognizant that close to half of all school-going age children in rural India cannot attend school regularly due to their need to work for the family in agricultural fields or do housework. We expect this challenge to apply to several other developing regions. It is therefore imperative to identify plausible scenarios in the everyday lives of rural children that are potentially conducive for mobile learning, as well as the cultural factors which impact learning in these scenarios. We have made a start with ethnographic studies in this direction, and initial results suggest that technology designs for out-of-school learning in rural India must take the social fault lines of gender and caste into account. Only then can we understand what it means to design applications that are culturally appropriate for everyday rural life.

Finally, most of the literature on our current pedagogical framework – task-based language teaching – is based on the classroom setting. Future work will need to extend it for out-of-school learning. In this thesis, we have designed our e-learning games as pre-task activities and have not adequately addressed what the tasks themselves ought to be. We believe that the traditional games themselves can be modified to become the eventual tasks, in which players interact with one another using the target language. As Prabhu observes about task-based instruction: “The teacher and the learners are both bound by the rules of the task and the source of authority is, in a limited but real sense, the task not the teacher” (1987: 51). In other words, tasks and games share some parallels in that the actions of the participants are restricted by the rules of the tasks and games respectively. The expected benefit of adapting village games into tasks in that they are a spontaneous context which rural children already enjoy interacting in. Most importantly, traditional games are rich systems of meaning that can help children relate to the target language, such that they can potentially transfer the language that they acquire from the pre-task e-learning games back into their everyday, non-digital games.

Work on some of the above future directions is already underway. The author is perhaps most gratified that as he exits his role as the primary graduate student leading the MILLEE project and transitions to spearheading MILLEE in a different capacity, a new generation of graduate students have appeared on the scene to help take MILLEE to the next level.

14 References

Abt, C.C. (1970). *Serious Games*. Viking Adult.

Aldrich, C. (2003). *Simulations and the Future of Learning: An Innovation (and Perhaps Revolutionary) Approach to E-Learning*. Pleiffer.

Alexander, C. (1977). *A Pattern Language: Towns, Buildings, Constructions*. Oxford University Press.

Appel, G., and Lantolf, J. (1994). Speaking as Mediation: A Study of L1 and L2 Text Recall Tasks. In *The Modern Language Journal*, Vol. 78, pp. 437-452.

Angrist, J., and Lavy, V. (1997). The Effect of a Change in Language of Instruction on the Returns to Schooling in Morocco. In *Journal of Labor Economics*, Vol. 15, No. 1, pp. S48-S76.

Azim Premji Foundation. (2004a). *Impact of Computer Aided Learning on Learning Achievements: A Study in Karnataka and Andhra Pradesh*.

Azim Premji Foundation. (2004b). *The Social Context of Elementary Education in Rural India*.

Banerjee, A., Cole, S., Duflo, E., and Lindon, L. (2005). Remedying Education: Evidence from Two Randomized Experiments in India. *NBER Working Paper No. 11904*.

Ball, E.W., and Blachman, B.A. (1988). Phoneme Segmentation Training: Effect on Reading Readiness. In *Annals of Dyslexia*, Vol. 38, pp. 208-225.

Bell, R.C. (1979). *Board and Table Games from Many Civilizations*. Dover Publications.

Beretta, A., and Davies, A. (1985). Evaluation of the Bangalore Project. In *ELT Journal*, Vol. 39, No. 2, pp. 121-127, April 1985.

Bergeron, B. (2006). Developing Seriosu Games. Charles River Media.

Björk, S., and Holopainen, J. (2005). *Patterns in Game Design*. Charles River Media.

Borchers, J. (2001). *A Pattern Approach to Interaction Design*. Wiley.

Bradley, L., and Bryant, P.E. (1978). Differences in Auditory Organization as a Possible Cause of Reading Backwardness. In *Nature*, Vol. 271, pp. 746-747.

Bradley, L., and Bryant, P.E. (1983). Categorizing Sounds and Learning to Read – A Causal Connection. In *Nature*, Vol. 301, pp. 419-421.

Brown, T.H. (2003). The Role of M-Learning in the Future of E-Learning in Africa? In *Proceedings of the 21st ICDE World Conference*.

Canale, M., and Swain, M. (1980). Theoretical Bases of Communicative Approaches to Second Language Teaching and Testing. In *Applied Linguistics*, Vol. 1, No. 1, pp. 1-47.

Clegg, J., Ogange, B., and Rodseth, V. (2003). Evaluating Digital Learning Material for English Language Development in African Primary Classrooms. IMFUNDO KnowledgeBank Paper, April 2003.

Cole, M. (1996). *Cultural Psychology: A Once and Future Discipline*. Harvard University Press.

Coltheart, M. (2005). Chapter 1: Modeling Reading: The Dual-Route Approach. Snowling, M.J., and Hulme, C. (Eds.), *The Science of Reading: A Handbook*, Blackwell Publishing Ltd.

Dias, M.B., Mills-Tetty, G.A., and Mertz, Joseph. (2005). The TechBridgeWorld Initiative: Broadening Perspectives in Computing Technology Education and Research. In *Proceedings of the International Symposium on Women and ICT: Creating Global Transformation*.

Donato, R. (2000). Sociocultural Contributions to Understanding the Foreign and Second Language Classroom. In Lantolf, J. (Ed.), *Sociocultural Theory and Second Language Learning*, Oxford University Press.

Druin, A. (1999). Cooperative Inquiry: Developing New Technologies for Children with Children. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 592-599.

Ellis, R. (2000). Task-Based Research and Language Pedagogy. In *Language Teaching Research*, Vol. 4, No. 3, pp. 193-220.

----------. (2003). *Task-Based Language Teaching and Learning*. Oxford University Press.

Faust, D., and Nagar, R. (2001). Politics of Development in Postcolonial India: English-Medium Education and Social Fracturing. In *Economic and Political Weekly*, India, July 28, 2001.

Folmer, E. Interaction Design Pattern Library for Games. http://www.helpyouplay.com/

Fox, B., and Routh, D.K. (1976). Phonetic Analysis and Synthesis as Word Attack Skills. In *Journal of Educational Psychology*, Vol. 68, pp. 70-74.

Frost, R. (2005). Chapter 15: Orthographic Systems and Skilled Word Recognition Processes in Reading. Snowling, M.J., and Hulme, C. (Eds.), *The Science of Reading: A Handbook*, Blackwell Publishing Ltd.

Fullerton, T. (2008). *Game Design Workshop: A Playcentric Approach to Creating Innovative Games, 2nd Edition*. Elsevier Inc..

Gamma, E., Helm, R., Johnson, R., and Vlissides, J. (1995). *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley Professional.

Gaver, B., Dunne, T., and Pacenti, E. (1999). Design: Cultural Probes. In *ACM interactions*, Vol. 6, No. 2.

Gee, J.P. (2003). *What Video Games Have to Teach Us About Learning and Literacy*. Palgrave Macmillan.

Gee, J.P. (2005). *Why Video Games Are Good For Your Soul: Pleasure and Learning*. Common Ground Publishing.

Gee, J.P., Hawisher, G.E., and Self, C.L. (2007). *Gaming Lives in the Twenty-First Century: Literate Connections*. Palgrave Macmillan.

Gupta, R.K. (1995). English in a Postcolonial Situation: The Example of India. In *Profession*, pp. 73-78.

Hallford, N., and Hallford, J. (2001). *Swords & Circuitry: A Designer's Guide to Computer Role-Playing Games*. Prima Publishing.

Halliday, Michael A.K. (1985). *An Introduction to Functional Grammar*. Arnold.

Harel, I. (1991). *Children Designers: Interdisciplinary Constructions for Learning and Knowing Mathematics in a Computer-Rich School*. Ablex Publishing.

He, F., Linden, L.L., and MacLeod, M. (2008). *How to Teach English in India: Testing the Relative Productivity of Instruction Methods within the Pratham English Language Education Program*. Working Paper.

Herron, R.E., and Sutton-Smith, B. (1971). Child's Play. John Wiley & Sons.

Horowitz, J.E., Sosenko, L.D., Hoffman, J.L.S., Ziobrowski, J., Tafoya, A., Haagenson, A., and Hahn, S. (2006). *Evaluation of the PBS Ready to Learn Cell Phone Study: Learning Letters with Elmo*. Report Prepared by WestEd.

India Visit Information. Sports in Ancient India. http://www.indiavisitinformation.com/india-sports/sports-ancient-india.shtml

Islam, Y.M., Rahman, Z., Razzaq, S.S., Sayed, M.A., and Zaman, S. (2006). Effect of Feedback During Lecture Style Delivery Both in a Face-to-Face Classroom and During a Distance Education TELEVISION Session in a Developing Country Like Bangladesh Without the Use of Internet. In *Proceedings of 6th IEEE International Conference on Advanced Learning Technologies*.

Johnson, D., and Johnson, R. (1989). *Cooperation and Competition: Theory and Research*. Interaction Book Company.

Kafai, Y. (1994). *Minds in Play: Computer Game Design as a Context for Children's Learning*. Lawrence Erlbaum.

Kinginger, C. (2001). i+1 ≠ ZPD. In *Foreign Language Annals*, Vol. 34, No. 5, pp. 417-425.

Kishwar, M.P. (2005). Deprivations's Real Language, *The Indian Express*, 2005. http://www.indianexpress.com/printerFriendly/12662.html#.

Kothari, B., Pandey, A., and Chudgar, A.R. (2004). Reading Out of the "Idiot Box": Same-Language Subtitling on Television in India. In *Journal of Information Technologies and International Development*, Vol. 2, No. 1, Fall 2004, pp. 23-44.

Krashen, Stephen D. (1987). *Principles and Practice in Second Language Acquisition*. Prentice Hall International.

Kreimeier, B. (2002). *The Case for Game Design Patterns*. Gamasutra, March 13, 2002. http://www.gamasutra.com/features/20020313/kreimeier_01.htm.

Lankoski, P., and Björk, S. (2007a). Gameplay Design Patterns for Believable Non-Player Characters. In *Proceedings of 3rd Digital Games Research Association International Conference*.

Lankoski, P., and Björk, S. (2007b). Gameplay Design Patterns for Social Networks and Conflicts. In *Proceedings of Computer Game Design and Technology Workshop*.

Lantolf, J.P. (2000). *Sociocultural Theory and Second Language Learning*. Oxford University Press.

Lazzaro, N. (2004). Why We Play Games: Four Keys to More Emotion Without Story. White Paper, XEO Design, March 8, 2004. http://www.xeodesign.com/whyweplaygames.html.

Liberating Voices! Pattern Language Project. http://trout.cpsr.org/program/sphere/patterns/

Librero, F., Ramos, A.J., Ranga, A.I., Triñona, J., and Lambert, D. (2007). Uses of the Cell Phone for Education in the Philippines and Mongolia. In *Distance Education*, Vol. 28, No. 2, pp. 231-244.

Lightbown, P., and Spada, N. (2006). *How Languages are Learned, 3rd Edition*. Oxford University Press.

Linn, M.C., and Eylon, B. (2006). Science Education: Integrating Views of Learning and Instruction. Alexander, P.A., and Winne, P.H. (Eds.), *Handbook of Educational Psychology, 2nd Edition*, Lawrence Erlbaum Associates.

Loh, S., and Seah, H.S. (2006). Comparing computer and traditional games using game design patterns. In Proceedings of International Conference on Game Research and Development.

Lomas, D., Douglass, J., and Rehn, D. (2008). Playpower: Radically Affordable Computer-Aided Learning with $12 TV-Computers. In *Proceedings of Meaningful Play 2008 Conference*.

Long, M. (1985). Input and Second Language Acquisition Theory. In Gass, S., and Madden, C. (Eds.), *Input in Second Language Acquisition*, Newbury House.

Malone, T.W. (1980). What Makes Things Fun to Learn? Heuristics for Designing Instructional Computer Games. In *Proceedings of the 3rd ACM SIGSMALL Symposium and the First SIGPC Symposium on Small Systems*, pp. 162-169.

Marcus, A., and Gould, E. (2000). Cultural Dimensions and Global Web User-Interface Design: What? So What? Now What? In *Proceedings of the 6th Conference on Human Factors and the Web*.

McLaughlin, B. (1990). Restructuring. In *Applied Linguistics*, Vol. 11, No. 2, pp. 113-128.

Michael, D., and Chen, S. (2005). *Serious Games: Games That Educate, Train, and Inform*. Course Technology PTR.

Mitra, S., Tooley, J., Inamdar, P., and Dixon, P. (2003). Improving English Pronunciation: An Automated Instructional Approach. In *Journal of Information Technologies and International Development*, Vol. 1, No. 1, pp. 75-84.

Mor, Y., Winters, N., Pratt, D., and Björk, S. (2007). Tools for Developing Design Patterns for Mathematical Computer Games. In *Proceedings of Game in Action Conference*.

Moraveji, N., Kim, T., Ge, J., Pawar, U.S., Inkpen, K. and Mulcahy, K. (2008). Mischief: Supporting Remote Teaching in Developing Regions. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*.

Munshi, K., and Rosenzweig, M. (2006). Traditional Institutions Meet the Modern World: Caste, Gender, and Schooling Choice in a Globalizing Economy. In *American Economic Review*, Vol. 96, No. 4, pp. 1225-1252.

Narayan, S. (1995). *Play and Games in Tribal India*. Commonwealth Publishers, New Delhi, India.

Nation, I.S.P. (Ed.). (2001). *Learning Vocabulary in Another Language*. Cambridge University Press.

Norman, D. (2003). *Emotional Design: Why We Love (or Hate) Everyday Things*. Basic Books.

Nunan, D. (2004). *Task-Based Language Teaching*. Cambridge University Press.

Pal, J., Lakshmanan, M., and Toyama, K. (2007). "My Child Will Be Respected": Parental Perspectives on Computers in Rural India. In *Proceedings of 2nd IEEE/ACM International Conference on Information and Communication Technologies and Development*.

Pawar, U.S., Pal, J., Gupta, R., and Toyama, K. (2007). Multiple Mice for Retention Tasks in Disadvantaged Schools. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*.

Pedagogical Patterns. http://www.pedagogicalpatterns.org

Peitz, J., and Björk, S. (2007). Understanding Pervasive Games Through Gameplay Design Patterns. In *Proceedings of 3rd Digital Games Research Association International Conference*.

Pimsleur, P. (1967). A Memory Schedule. In *Modern Language Journal*, Vol. 51, pp. 73-75.

Prabhu, N.S. (1987). *Second Language Pedagogy*. Oxford University Press, Oxford, UK.

Pratham. (2008). Annual Status of Education Report (Rural) 2007 – Provisional. Pratham Resource Center, Mumbai, India.

Prensky, M. (2000). *Digital Game-Based Learning*. McGraw-Hill Companies.

Read, J., MacFarlane, S., and Casey, C. (2002). Endurability, Engagement and Expectations: Measuring Children's Fun. In *Interaction Design and Children*, Shaker Publishing.

Rollings, A., and Adams, E. (2003). *Andrew Rollings and Ernest Adams on Game Design*. New Riders.

Roschelle, J., and Jackiw, N. (1999). Technology Design as Educational Research: Interweaving Imagination, Inquiry, and Impact. In Kelly, A. and Lesh, R. (Eds.), *Research Design in Mathematics & Science Education*, Kleuwer.

Salen, K., and Zimmerman, E. (2003). *Rules of Play: Game Design Fundamentals*. The MIT Press.

Scaife, M., Rogers, Y., Aldrich, F., and Davies, M. (1997). Designing For or Designing With? Informant Design for Interactive Learning Environments. In *Proceedings of ACM Conference on Human Factors in Computing Systems*.

Schmidt, R., and Frota, S. (1986). Developing Basic Conversational Ability in a Second Language: A Case Study of an Adult Learner of Portuguese. In Day, R.R. (Ed.), *Talking to Learn: Conversation in Second Language Acquisition*, Newbury House.

Schmidt, R. (1990). The Role of Consciousness in Second Language Learning. In *Applied Linguistics*, Vol. 11, pp. 129-158.

Scribner, S., and Cole, M. (1981). *The Psychology of Literacy*. Harvard University Press.

Seedhouse, P. (1999). Task-Based Interaction. In *ELT Journal*, Vol. 53, No. 3, pp. 149-156.

Shaffer, D.W. (2006). *How Computer Games Help Children Learn*. Palgrave Macmillan.

Sheorey, R. (2006). *Learning and Teaching English in India*. Sage Publications.

Shukla, S. (1996). From Pre-Colonial to Post-Colonial: Educational Transitions in Southern Asia. In *Economic Political Weekly*, Vol. 31, No. 22, pp. 1344-49.

Skehan, P. (1998a). Task-Based Instruction. In *Annual Review of Applied Linguistics*, Vol. 18, pp. 268-86.

Skehan, P. (1998b). *A Cognitive Approach to Language Learning*. Oxford University Press.

Smith, A., Dunckley, L., French, T., Monocha, S., and Chang, Y. (2004). A Process Model for Developing Usable Cross-Cultural Websites. In *Interacting with Computers*, Vol. 16, pp. 63-91.

Srivastava, R.S. (2002). Anti-Poverty Programmes in Uttar Pradesh: An Evaluation. Report of the Planning Commission Research Project, Institute of Human Development, New Delhi, India.

Sutton-Smith, B. (1976). *The Psychology of Play*. Arno Press.

Swain, M. (1985). Communicative Competence: Some Roles for Comprehensible Output in its Development. In Gass, S., and Madden, C. (Eds.), *Input in Second Language Acquisition*, Newbury House.

Swan, M. (2005). Legislation by Hypothesis: The Case of Task-Based Instruction. In *Applied Linguistics*, Vol. 26, No. 3, pp. 376-401.

The Times of India (2006). *Bollywood Films Fuels Mobile Gaming Industry*. May 8, 2006.

Thorne, S. (1999). *An Activity Theoretical Analysis of Foreign Language Electronic Discourse*. Ph.D. dissertation. University of California, Berkeley, USA.

Traditional Children's Games from Around the World. Topics Online Magazine, Issue 11. http://www.topics-mag.com/edition11/games-section.htm

UNESCO. (2005). *Working Guidelines on Teacher Development for Pedagogy – Technology Integration*, UNESCO Bangkok, Working Draft.

Van den Branden, K., Van Gorp, K., and Verhelst, M. (2007). *Tasks in Action: Task-Based Language Education from a Classroom-Based Perspective.* Cambridge Scholars Publishing.

Van Duyne, D.K., Landay, J., and Hong, J. (2002). *The Design of Sites: Principles, Processes and Patterns for Crafting a Customer-Centered Web Experience*. Addison-Wesley.

Van Lier, L. (1996). *Interaction in the Language Curriculum: Awareness, Autonomy and Authenticity*. Addison Wesley.

Vodafone (2005). Moving the Debate Forward. *Vodafone Policy Paper Series*, No. 2, March 2005.

Vygotsky, L.S. (1978). *Mind in Society*. Harvard University Press.

-------- (1986). *Thought and Language*. The MIT Press.

Weinert, R. (1995). The Role of Formulaic Language in Second Language Acquisition: A Review. In *Applied Linguistics*, Vol. 16, No. 2, pp. 180-205.

Wikipedia. Hopscotch. http://en.wikipedia.org/wiki/Hopscotch

Wikipedia. Kabbadi. http://en.wikipedia.org/wiki/Kabaddi

Wikipedia. Kho-Kho. http://en.wikipedia.org/wiki/Kho_kho

Wikipedia. Tag (game). http://en.wikipedia.org/wiki/Tag_(game

Wiley, D.A. (2002). Connecting Learning Objects to Instructional Design Theory: A Definition, A Metaphor, and a Taxonomy. In Wiley, D.A. (Ed.), *The Instructional Use of Learning Objects*, Association for Instructional Technology and Association for Educational Communications and Technology.

Wood, J. (2001). Can Software Support Children's Vocabulary Development? In *Language Learning & Technology*, Vol. 5, No. 1, pp. 166-201.

Yeo, A.W. (2001). Global-Software Development Lifecycle: An Exploratory Study. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*.