

Illuminac: Simultaneous Naming and Configuration for Workspace Lighting Control

*Ana Ramirez Chang
John F. Canny*

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2008-119

<http://www.eecs.berkeley.edu/Pubs/TechRpts/2008/EECS-2008-119.html>

September 17, 2008



Copyright 2008, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Acknowledgement

We would like to thank the study participants from all three studies and Lora Oehlberg for her sketch of the person speaking into a microphone. The first author was supported by a National Science Foundation Graduate Student Fellowship during this work.

Illuminac: Simultaneous Naming and Configuration for Workspace Lighting Control

Ana Ramírez Chang, John Canny
Berkeley Institute of Design and Computer Science Division
University of California, Berkeley
{anar, jfc}@cs.berkeley.edu

ABSTRACT

We explore natural and calm interfaces for configuring ubiquitous computing environments. A natural interface should enable the user to *name* a desired *configuration* and have the system enact that configuration. Users should be able to use familiar names for configurations without learning, which implies the mapping from names to configurations is many-to-one. Instead of users learning the environment’s command language, the system simultaneously learns common configurations and infers the keywords that are most salient to them. We call this the SNAC problem (Simultaneous Naming and Configuration). As a case study, we design a speech interface for workspace lighting control on a large array of individually-controllable lights. We present an approach to the SNAC problem and demonstrate its applicability through an evaluation of our system, Illuminac.

Author Keywords

Natural Speech Interfaces, Non-negative Matrix Factorization, Environment Control

ACM Classification Keywords

H.5.2 User Interfaces: {Voice I/O, Natural language}

INTRODUCTION

The number of electronic devices in our environment is ever increasing. While this brings greater flexibility and control, configuring each individual device becomes ever more tedious. For example, in the home, to prepare the environmental state for cooking, one might want to turn the radio on and turn it to the news, turn the volume of the speakers up since the kitchen will be noisy, and make the kitchen lights bright. Then, to prepare the environmental state for watching television, one might dim the lights in the living room, turn on the television, turn the volume of the speakers to medium, and close the blinds. Controlling all of these devices—the lights, the radio, the television, the speaker volume, the television, and the blinds—to achieve a desired environmental state is quite tedious. It is widespread best practice to use activity-

specific “configurations” of many devices rather than setting each device individually. The user can then invoke the configuration with a single action— i.e. a key press or a spoken command.

As with any interface, an interface for controlling the environmental state in the home should match the user’s mental model. That is, the user should only need to specify an intuitive *name* for the environmental state rather than the *configuration* of each individual device needed to achieve the desired state. In the home example, the user should be able to indicate “set cooking mode” or “apply TV mode” rather than specifying the configuration of the radio, speakers and lights to achieve the cooking or television watching environmental state.

In addition to matching the user’s mental model, we want the interface to be “calm.” That is, the interface in a ubicomp environment, like environment control, should be almost invisible except during direct (focal) interaction (as advocated by Weiser [24]). In this context, a speech-based interface seems like a good option. With distributed microphone technology, the physical interface all but disappears, but jumps fluidly to the foreground when the system responds to spoken input. Furthermore, speech is often considered the most natural form of human expression and has the potential to address certain accessibility concerns.

In the home example, the user should be able to say “I’d like the cooking mode please”, or “Apply the mode for preparing dinner,” or similar variation, and have the system give a similar response. They should also be able to say “cooking mode” or “TV mode” and get a different response. These terms are widely shared by people, and their repeated use during training allows a system to learn them as well.

Note that this problem is more challenging than simply memorizing command strings and the appropriate device settings. In the latter case, the system will be extremely brittle, and will respond only when exact training strings are provided. By *simultaneously* learning commands and device settings, the system becomes both more robust and better able to generalize. For instance, there will be many training strings for watching television that include the word “TV” and many other filler words, but which all specify a similar light and window shade pattern. Since the system looks for common patterns in names and configurations, the word “TV” will be strongly present in a pattern that includes presentation

light settings. Thus it is able to infer that “TV” is a salient keyword for watching television vs. “please” or “now” that may also occur in command strings. Similarly, “movie” and “TV” will typically occur with similar patterns of devices, and the system will be able to infer that they are aliases in this context. We called this the SNAC problem (Simultaneous Naming And Configuration).

We believe this theme occurs in many ubicomp environments. For example, in the workspace the SNAC problem arises in mapping the “semantic gap” from a configuration command like “presentation” to control the lights, projector, and sound devices. In a workspace with an array of individually controllable lights instead of a bank of lighting all controlled by one light switch, mapping from a configuration command like “Turn Joe’s desk lights on please” to turn on the lights over Joe’s desk presents the same challenge.

In this paper, we describe the design, deployment and testing of a natural speech interface in an open-plan workspace. The system runs live and controls 79 devices (individually controllable lights) from 25 users who have both their own and shared environment names and configurations. We make the following contributions:

- We identify and describe an approach to the Simultaneous Naming and Configuration problem (SNAC) — a problem which arises in natural environment control. Our approach results in a calm interface which learns intuitive *names* for commonly used *configurations* of devices in a home or work environment instead of requiring the user to learn hard coded, static names for hard coded configurations of devices, or worse, hard coded names for individual devices. (SNAC Approach section)
- We identify an appropriate learning algorithm for the simultaneous naming and configuration problem: non-negative matrix factorization (NMF) (SNAC Approach Learning Model section).
- We show the applicability of our approach to the SNAC problem on natural speech interfaces for workspace lighting control by implementing and deploying a SNAC-based system, Illuminac. With Illuminac, we see that one or two training points is often sufficient to produce environmental states with mostly correct configurations, thereby providing good accuracy with little training (Illuminac section).

OVERVIEW OF WORKSPACE LIGHTING CONTROL

Before we present the approach to the SNAC problem, we first introduce the workspace lighting control domain for which Illuminac is designed.

Many large open-plan workspaces have extensive banks of lights controlled by just one light switch. Therefore, the lighting control is not flexible enough to respond to occupancy or daylighting. Many lights are turned on for just a few occupants, and lights next to a window cannot be turned off without turning off the lights away from the window. More granular control over the lighting in large workspaces

could enable a reduction in energy consumption by allowing unnecessary lights to be turned off. In addition to turning off unnecessary lights, Moore *et. al.* show when users are given control over their lights, they often set the lights to a luminance level lower than the level recommended by the Chartered Institution of Building Services Engineers (CIBSE) with an average lamp output 55% of the maximum [15].

Proprietary granular lighting control systems for both retrofit and new construction have been available for some years [23]. Most recently, a low-cost industry standard has emerged: DALI: Digital Addressable Lighting Interface [7]— is a bus system for individual light control which has been adopted by most lighting manufacturers, and deployed in major installations (e.g. Heathrow terminal 5). It is an increasingly popular option in energy-conscious building design. However, the increase in flexibility implies an increase in control complexity. The current state-of-the-art in complex lighting control is panels of wall buttons with (often cryptic) scene names, or touch screens with complex menu hierarchies. Neither of these approaches address the issue of lighting in reconfigurable workspaces, where many more configurations are possible. There is clearly a need for more flexible, intuitive and scalable lighting control.

Intelligent systems for controlling workspace lights that adapt to daylight from windows and occupancy levels are being developed using flexible lighting control. At the same time, there is still a need for user interfaces that allow users to directly control flexible workspace lights and override such intelligent systems when appropriate. Escuyer *et. al.* found users didn’t mind the occasional error in an automatic lighting control system as long as they had an easy manual way to correct the lighting scene [8]. Love notes that automatic controls were commonly disabled by users [13]. Whether it’s in conjunction with an intelligent system, or on it’s own, a manual control for flexible lighting is necessary.

We propose a speech-based interface for lighting control in shared workspaces, as they are natural and allow for a calm interface. To do so, users should be able to customize the system to work with names of lighting scenes that are natural to them. For example, based on our experience with Illuminac, we found that one user says “turn on my lights” to turn on the two lights over her desk, while another user says “all on” to turn on the four lights around her desk.

As alluded to earlier, to support customized speech commands for personalized lighting configuration, we must address the SNAC problem. Specifically, we need to discover the lighting scenes commonly used in the space and the names for those lighting scenes. A lighting scene is a set of lights and an intensity setting for each of the lights in the set. A name for a lighting scene is a set of words. For example, a user might have two intensity settings for the three lights over his desk—one for desk work and one for computer work. He might say “Please turn Jason’s reading lights on” or “can I get Jason’s computer lights on?”. The desk work lighting scene would be the three lights over Jason’s desk set to 100% intensity and the computer lighting scene would

be the three lights over his desk set to 50% intensity. The name of the desk work lighting scene might be {Jason's, reading} and the name of the computer lighting scene might be {Jason's, computer}.

In our approach to the SNAC problem we use a learning-based system that is trained on two kinds of data from the users simultaneously: the commands and the configurations (lighting scenes). While the details of the design and evaluation of Illuminac are discussed in the subsequent sections, in the remainder of this section, we sketch how the user interacts with our system.

To add a command and lighting scene to Illuminac, users train the system by first recording their command. Then, the user demonstrates the desired lighting configuration and identifies herself. The novel aspect of our system is that rather than simply storing this mapping from command to configuration, we combine the recorded speech command and lighting configuration into a common representation to provide as input for a standard machine learning algorithm. Intuitively, the system uses the learning algorithm to identify structure across the space of command-configuration pairs, not just the space of commands. Once the user has trained the system on a few examples, the user can say her command into any of the microphones in the room, and the system changes the lighting scene by applying the trained model to the user's command. Because the model is trained on commands and configurations specific to the workspace, we expect to be able to perform reasonable lighting actions with less command training. For example, when a visitor who has never provided training input to the system comes into the lab, she can try her command and potentially get reasonable behavior because regular users may have already trained the system on similar commands. Of course, if the resulting behavior is undesired, she can manually change the lighting scene, thereby giving the system another training data point suited to her.

APPROACH TO THE SNAC PROBLEM

As we described above, the challenge in supporting customized commands for configuration tasks is not only discovering how users give commands. The challenge is in simultaneously discovering (1) the commands natural to the users, (2) the configurations naturally used in the domain, and (3) the mapping between the commands and configurations.

Traditional Approach to Designing Speech Interfaces

Discovering the names naturally used by users is a common problem in the design of speech interfaces. Speech interface designers commonly use the wizard-of-Oz data collection technique [6] to gather formative data including sample utterances, which helps inform the design of the speech interface toward supporting more natural speech input. The wizard-of-Oz data collection technique allows more realistic sample usage data to be collected before a prototype of the system is ready. A human (called the wizard) acts as the speech interface by responding to the user's commands. The wizard often communicates with the user through a speech

synthesizer to make the user think she is using a working system and invoke more realistic responses.

Simultaneously Discovering Naming and Configuration

For workspace lighting control, we not only need to discover the names users use to refer to lighting scenes, but also the configurations of lights that achieve the named lighting scene. Thus, in addition to collecting sample utterances from users, we also collect sample configurations for those utterances.

Since the configurations are not known *a priori*, a wizard is not able to accurately respond to a user's commands. In fact, if a wizard is used, the users will alter their commands so that the wizard will understand the configurations to which they are referring. We discovered this effect during an early phase of data collection in which users sent messages to a wizard asking her to change the lighting scene. Thus, in our approach, users train the system directly by demonstrating their configurations in addition to recording their commands.

In order to make the system both more robust to small variances in the commands and better able to generalize, we use a supervised learning algorithm. The learning algorithm is trained on data provided by the users that includes both the users' personalized commands and their customized configurations.

This approach also allows an interface to support more natural interaction through customized commands and configurations in a new workspace or with a new set of users without having to have a designer perform the customization.

LEARNING MODEL

We select a learning algorithm that allows factors to overlap. This property is important in our problem because both configurations and commands can overlap. In the home example, a cooking environmental state and a television watching environmental state could both set the speaker volume up. For workspace lighting, the space is shared and users who sit next to each other often have overlapping sets of lights in their lighting scenes. The names of the configurations can also overlap. In the workspace lighting control example, many users refer to their lights as "my lights." We disambiguate such similar names with the speaker's identification.

Non-negative matrix factorization (NMF), a standard learning algorithm, finds non-orthogonal or possibly overlapping, factors in the training data. Each data point (a command-configuration pair) can be explained by an additive combination of the factors. NMF is more appropriate than a clustering learning algorithm such as *k*-means that would assign a light to one cluster, not allowing it to be part of multiple clusters.

Since we need to learn the commands, configurations, and the mapping, we train our model on command-configuration pairs. That is, we train the learning model on two kinds of data simultaneously: the commands and the light configurations. To record a new training point, users provide both the




Figure 1. Picture and floor plan of the workspace for which we have implemented a natural speech interface for the lighting control. Each of the 79 individually-controllable lights as well as the eight microphones are shown.

command and the configuration, specifically they

1. record their command using a microphone and
2. demonstrate the lighting scene named by their command by manually changing the lighting scene (using a web-based graphical user interface).

By collecting both kinds of data simultaneously, we are collecting data about the commands that are natural to users, the configurations that naturally occur in the target domain, and the ground truth mapping between the two.

ILLUMINAC

We designed and deployed a natural speech interface for workspace lighting control, Illuminac, to show the applicability of our approach to the SNAC problem. In this section, we describe the workspace for which Illuminac was designed, outline the iterative design steps used in the design process, explain the design of the system, discuss the evaluation of the system, and present the results from the evaluation.

Workspace Details

Illuminac was designed for and is deployed in a 2,300 square foot open-plan shared workspace with about twenty-five regular occupants (nineteen of whom have permanent desks in the workspace; the rest have permanent desks in the adjacent room). The workspace has six graded-awareness cubicles (i.e., cubicles with walls of varying heights from full height to desk height) that occupy half the room. The other half is a multi-use space for meetings, presentations, ad-hoc team meetings or individual work.

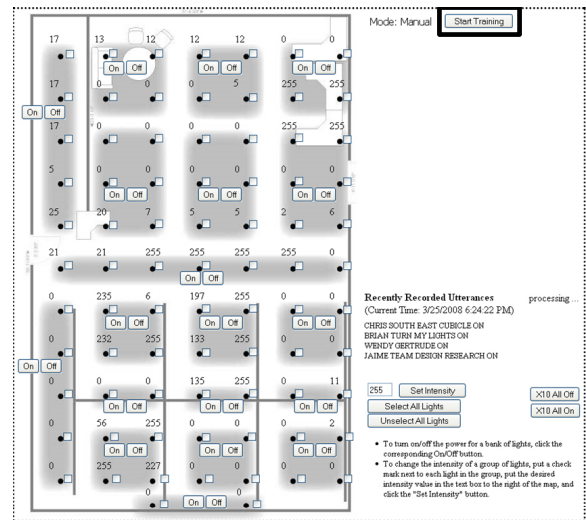


Figure 2. Web-based graphical user interface used to manually configure the array of lights in the workspace.

The multi-use space has a presentation screen, a “soft space” with a couch and chairs, and a set of four computers for visitors of the lab to use. There is also a tool shop in one corner of the room. Figure 1 shows a picture and the floor plan of the room. The room has 79 individually-controllable compact fluorescent lights mounted overhead. The intensity of each of the lights can be controlled over the network via a web interface. All occupants of the lab have access to the web interface (Figure 2). They can access it from their personal computers or from one of the public machines. The public machine next to the entrance (labeled “Main Public Machine” in Figure 1) always has the web interface open.

There are eight desk microphones throughout the room to allow easier access to the lighting control system. There is one microphone in each cubicle, one at the desk in the corner, and one at the main public machine where the graphical interface to the lights is always open in a browser window. Each microphone has a clearly labeled on/off switch so that residents may control what is and is not recorded.

Iterative Design Steps

In the design of our natural speech lighting control system, we followed an iterative design process beginning with a text-based wizard-of-Oz study, followed by a training data collection study, and finally a deployment of the live system.

Text-Based Wizard-of-Oz Study

We began our design of Illuminac with a low-fidelity wizard-of-Oz study to better understand the lighting control domain including the kinds of lighting scenes used in the space and the types of commands used to refer to the lighting scenes. By low-fidelity, we mean natural language text input instead of speech input. Participants were asked to send the wizard (a researcher in the lab) an instant message whenever they wanted to change the lighting scene. The wizard would change the lighting scene using the web interface based on the participant’s message. If the wizard was not available

to change the lighting scene, the participants were asked to type the message they would have sent to the wizard into the web interface and then change the lighting scene themselves with the web interface. This way data could be collected even if the wizard was not at his desk. The GUI interface was similar to the one in Figure 2, but it also included a text box for entering the messages.

We collected three weeks of data including 230 command-configuration pairs from ten participants who were regular occupants of the space.

This study confirmed our hypothesis that the lighting configurations are sometimes overlapping and are not all disjoint sets of lights. After the formal study concluded, some participants expressed the desire to continue being able to ask the wizard to change the lighting scene, but instead of sending instant messages, they wanted to ask the wizard with a spoken command. This provided anecdotal evidence that speech would be a good fit for lighting control in the space.

Formative Training Data Collection

After the low-fidelity wizard-of-Oz study, we collected two weeks of high fidelity training data to design and tune the learning algorithm, which estimates lighting scenes given a spoken command. The study included sixteen participants who were regular occupants of the space. Each participant was asked to record a command and demonstrate the desired system response as if they were training the system to understand their personalized commands. They were asked to complete the following three steps each time they wanted to change the lighting scene:

1. say their command to change the lighting scene;
2. type their name into the text box on the web page; and
3. change the lighting scene with the web interface.

We did not tell them when or how to change the lighting scene in the lab. The participants could use any of the eight microphones throughout the space to record their command. Then they used the web interface to demonstrate their desired change in the lighting scene. We followed the data collection with individual interviews and asked the participants to reflect on their lighting control preferences and their experience with the study.

We collected 120 command-configuration pairs. For each pair, we collected the user's name, the audio clip of the command and the intensity values before and after the command. Each audio clip was transcribed manually as well as with an automatic speech recognizer. The complete vocabulary included about 350 unigrams and bigrams (i.e., single words and two-word phrases).

This high-fidelity formative training data once again demonstrated both the configurations and commands are overlapping, which strongly suggests the selection of NMF as the basis for the learning algorithm. This data directly informed the design of the learning algorithm, which in turn enabled us to build and deploy a live system.

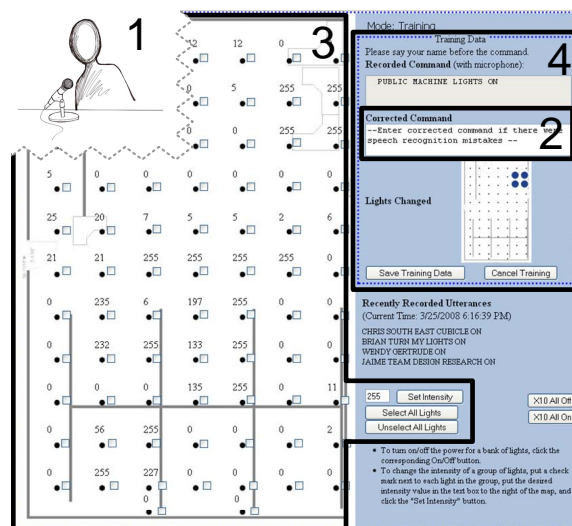


Figure 3. The four steps users follow in the training mode to train the system on their personalized command and lighting scene.

System Design

Illuminac has two modes: a training mode where users train the system on their personalized lighting configurations and a running mode where users can use their personalized commands to change the lighting scene. The system is in running mode by default meaning it is ready to process a command and change the lighting scene. To switch to the training mode, the user uses the “Start Training” button on the web-based graphical user interface (highlighted in Figure 2).

Training Mode

In the training mode, users complete the following four steps to add a training point (shown in Figure 3):

1. record her command using one of the microphones around the room (prefaced by her name to simulate speaker identification);
2. correct any recognition errors in the automatic speech recognition transcript using the web-based graphical user interface;
3. demonstrate her desired lighting scene using the web-based graphical user interface to manually set the lighting configuration; and
4. verify the command and configuration are correct before saving the training point.

After the user adds a new training point, the system retrains the model on the old training data plus the new data point. It also adds the command to the speech recognition language model to increase the recognition accuracy for the command the next time it is used. The new model is trained and ready to be used in the running mode in just a few seconds. Even though we are not using an online training algorithm, we can retrain the model fast enough when we get a new training point that users can think of the system as an online training system.

Running Mode

In the running mode, a user can use one of her speech commands to change the lighting scene. To do so, the user says her command into any one of the microphones around the room. The system transcribes the command, applies the trained model to the transcript, and changes the lighting scene according to the estimated lighting scene.

If the estimated lighting scene is incorrect, the user can say “undo,” “cancel,” or “wrong” to undo the last lighting scene change.

While the system is transcribing a command, the text “processing ...” is displayed in the GUI and an LED next to each microphone lights up (typically takes about two seconds). This lets the user know the system “heard” them and is processing the command.

Learning Model

The system uses a learning algorithm, namely least-squares non-negative matrix factorization (NMF) to estimate the new lighting scene given a command.

NMF has been shown to work well for pure text clustering, as well as image segmentation [11]. In our problem, we have both text (the commands), as well as an image (the grid of lights).

As the name suggests, NMF imposes a non-negative constraint on the data. The commands are naturally represented with non-negative values in a term-frequency vector. The lighting scene is also naturally represented with non-negative values for the intensity of each light.

The NMF algorithm is fast enough to allow us to retrain the model between uses. The algorithm runs in 0.89 seconds on two weeks of training data.

Given the similarity of our data with text and images, the non-orthogonal factors in our data, and the naturally non-negative values in our data representation, we selected NMF to learn which lights should change given a new command and the current lighting scene.

Data Representation

Commands

We represent the commands (tagged with the user’s name) with a term-frequency vector. A term is a word (unigram) or pair of consecutive words (bigram) that appears in any of the commands, and each term is represented with an entry in the vector. Thus the length of the vector is the number of unique unigrams or bigrams in all of the commands. The value of each entry in the vector is the number of times the unigram or bigram appears in the command. We include bigrams to be able to capture phrases, such as “soft space,” “kitchen area,” or “public machines.”

Lighting Scene

We represent the lighting scene as a vector of intensity values: one for each light. The intensity values that did not

change when the user demonstrated the new lighting scene are set to zero to allow the algorithm to learn to which lights a command refers, as well as to what intensity to set the lights.

Users do not train the system on “off” commands, instead they are derived from the “on” commands. This is because we want to make sure the lights are set to intensity zero when they are turned off, not just a low intensity. Based on data from the formative data collection studies, the words “off” or “out” are added to the “on” command to derive the “off” command. Note, “on” commands do not need to have the word “on” in them, but to turn the lights off, the command must have the word “off” or “out” in it (which was always true in our sample data).

Application of NMF

Our trained model is a set of orthogonal, possibly overlapping factors. To get the factors that describe our training data, we use NMF to factorize our data into two matrices, one of which describes the factors in our data.

NMF is an algorithm that finds a positive factorization of the given matrix [12, 11]:

$$\mathbf{X} \approx \mathbf{U}\mathbf{V}^T$$

where \mathbf{X} represents the training data, \mathbf{V}^T represents the factors in our data, and \mathbf{U} tells the strength of each factor in each of the data points.

Each row in \mathbf{X} is a training point (i.e., a command-configuration pair). Thus each row vector is comprised of the command term-frequency vector concatenated with the lighting configuration vector. The matrix \mathbf{X} has dimensions $m \times n$ where m is the number of training points in the model and n is the length of the term-frequency vectors (q) plus the number of lights (p). The matrix \mathbf{V}^T has dimensions $k \times n$ where k is the number of factors. We explain shortly how we chose k in the system. Each row in \mathbf{V}^T represents one of the factors. The matrix \mathbf{U} has dimensions $m \times k$. The i^{th} row in \mathbf{U} gives the coefficients to the additive sum of the factors to describe the i^{th} training point.

Selecting the Number of Factors (k)

A good factorization should explain the data with the smallest number of factors. We used the formative training data to calculate the accuracy rate for a large range of k values (between 1 and 200) and plotted the accuracy as a function of k . We chose the value of k where increasing k had negligible improvement in accuracy (i.e., the elbow of the curve), which we found to be 32.

Lighting Configuration Estimation

To estimate a lighting configuration given a new command, we use NMF again to find the factors from our training data that are present in the new command. Figure 4 depicts how to apply the trained model to a new command. As described above, our trained model is a matrix that represents the factors in the training data \mathbf{V}^T . Each factor has a “command” component and a “configuration” component. The command



Train Model:Factorize training data (X) into matrices U and V using NMF**Calculate factor coefficients (u'_{run}):**Given a new command (x'_{cmd})Use NMF, holding V_{cmd}^T and x_{cmd} constant**Estimate new lighting scene (x'_{config}):**Multiply factor coefficients (u'_{run}) by factors (V_{config}^T)

Figure 4. Given a new command, we use the factors matrix from our training data and NMF to estimate the new lighting configuration. In each step, matrices that are given are in black, and matrices that we calculate are in gray.

components are in the left half of V^T (columns 1 to q in V^T), and the configuration components are in the right half of V^T (columns $q + 1$ to n in V^T). We will call these matrices V_{cmd}^T and V_{config}^T respectively.

$$V^T = [V_{cmd}^T, V_{config}^T]$$

The matrix we need to factorize contains the term-frequency vector for the new command, which is only one row. We run NMF on the command term-frequency vector

$$x'_{cmd} \approx u'_{run} \times V_{cmd}^T$$

(x'_{cmd} has dimensions $1 \times q$, u'_{run} , has dimensions $1 \times k$), holding V_{cmd}^T constant to get u'_{run} , which tells us which factors are present in the new command. We can then multiply u'_{run} by the configurations component of the factors to get the lighting configuration vector:

$$x'_{config} = u'_{run} \times V_{config}^T$$

The final step in estimating the new lighting configuration is to estimate which intensity values to change and to what value to change them. We use two NMF models to do so. To estimate which intensity values to change, we train a model on a data matrix where the values in the configuration vectors are boolean values that indicate which lights are part of a lighting configuration. To estimate to what intensity value to set a light, we use the original data matrix described above. When we apply these two models to a new command, we get two parts of the lighting configuration estimate: $x'_{configbool}$ and x'_{config} . The $x'_{configbool}$ vector has values between 0 and 1. We perform cross validation on the formative training data to select a threshold value to convert the values to

boolean values. We change the intensity value for each light that has a one in $x'_{configbool}$ to the estimated intensity value for that light in x'_{config} . In other words, we only change the intensities of the lights involved in the lighting scene referred to in the new command.

Automatic Speech Recognition

We are using Carnegie Mellon’s Sphinx 3.5 speech recognizer [19] to transcribe the commands. Sphinx 3.5 is a state-of-the-art fully-continuous acoustic model recognizer, which is open for experimental use.

We trained an acoustic model on the ICSI Meeting Corpus [9], which features all the imperfections of natural speech: pauses, um’s and ah’s, truncated words, grammar errors, sentence and phrase restarts, etc. It also features a variety of nationalities and accents. This acoustic model is a better match for our user base than the publicly available acoustic models (HUB4, WSJ, RM1), which feature artificially clean speech being read from transcripts.

The language model is trained on the commands from the training data. The recognizer runs fast enough to transcribe the commands almost in real-time (it runs in 1.2x real time on a 3 GHz dual processor machine with 1.5 GB of RAM). As described above, we use a status LED to let users know the system “heard” them while the recognizer is processing the command.

Evaluation*Study Description*

We deployed Illuminac with ten of the twenty five regular occupants of the lab for one week. The system ran with live speech recognition on the audio from the microphones around the room, live training mode where the model was re-trained after each training point was collected, and live running mode which applied a lighting scene when a user spoke a command into one of the microphones. We started with no training data and asked the participants to train the system on their commands again¹. Participants were instructed to use the system whenever they wanted to change the lighting scene. The first time they used the system for a particular command they were asked to record a training data point. Subsequent times they were asked to test their commands, but recording more training points if the system did not respond as expected. When the participants tested a new command they recorded the accuracy of the results on a paper log next to the microphone. They recorded the accuracy of the system response by circling one of the following options:

4	3	2	1	0
Correct	Partially Correct	Some Correct, Some Wrong	Nothing Happened	Wrong

At the end of the study, the participants were asked to complete an anonymous web questionnaire about their experiences with the system.

¹Many of the participants in this study also participated in previous studies and had already provided training points.

Training and Test Data Per Command Type

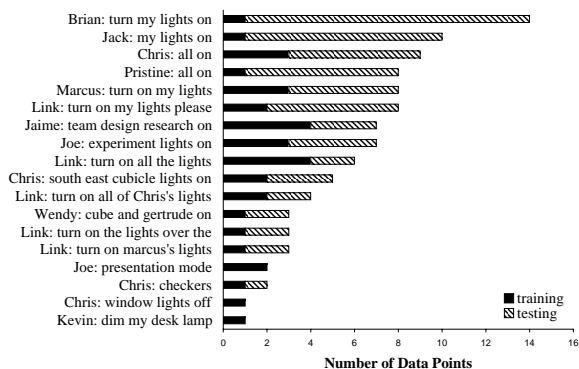


Figure 5. The number of training and test data points collected for each group of similar commands. The command group labels are representative commands recorded by the participants.

Data Collected

We collected 43 training points and 81 test points from the ten participants. For each training point, we collected the state of the lighting scene before and after the training session, as well as the command (transcribed using the automatic speech recognizer and corrected by the participant). For each test point, we collected the state of the lighting scene before and after the system changed the lighting scene, the command that the participant provided using one of the microphones, and the participant’s evaluation of the system’s response. Figure 5 shows the number of data points collected for each group of similar commands.

Study Results and Discussion

To analyze the results we manually assigned each training and testing point to a group of similar commands. For example all of the commands Jack used to refer to the lights over his desk were put into one group. Figure 5 lists a representative command for each group of similar commands.

The average testing score plateaued between “correct” and “partially correct” when commands were tested with 1, 2, or 3 training points (see Figure 6). We believe these results could be improved by changing the training GUI to alleviate a common confusion about which lights were being saved as a lighting scene. The interface only recorded the intensity values that changed during the training mode, but users thought it was saving the intensity values for each light selected. As a result some of the training data was not correct, which we believe impacted the accuracy results negatively.

Although the average score is closer to “partially correct” than “correct”, when asked in the post questionnaire, “After the study is over, would you like to continue using the system?” eight out of ten participants responded *Yes*, and two responded *Maybe*—the options were *Yes*, *Maybe* and *No*. One of the participants who responded *Maybe* said the microphone was too far away, and he was too lazy to get to a microphone, which is a limitation in our experimental setup that could be overcome in a commercial deployment. Right now each cubicle with three people shares one microphone,

Average Lighting Scene Change Score vs. Training Length




Figure 6. The average test point score as reported by the participants versus the training length for the specific type of command when the test point was recorded

which could be remedied by giving each user a microphone at their desk, making the microphone convenient to access. The other participant who responded *Maybe* tends to sit in the public area most of the time where the furniture moves around quite a bit, and he does not often use the same set of lights. In such an open space, a location-based speech approach would likely work better where users could say “lights on here.” Such a location based approach could be implemented with distributed microphone array technology overhead, though such technology would not be desirable in the cubicle area for privacy reasons. With overhead microphones, users cannot control what is being recorded, but with desk microphones, users have the power to turn the microphone on their desk off.

When asked, “How many training data points would you be willing to provide to be able to use speech to control the lights” participants responded with an average of 3.9 (min 2, max 5). On average participants recorded 1.9 training points per command group during the study. Since the average number of training points participants would be willing to provide is higher than the average number they recorded during a formal study, we believe the performance of the system outside a formal study would be at least similar to the performance during the study.

RELATED WORK

We situate our work in the literature related to designing natural speech interfaces, existing smart home and home IT projects, workspace lighting control, multimodal interfaces, and previous applications of factorization algorithms such as NMF.

Natural Speech Interface Design

In speech interfaces where all of the objects that can be referred to are known *a priori*, the designer can use the wizard-of-Oz technique [6] to support names for those objects that are natural to the users. For example, in the RoomLine [3] system, a speech interface that allows users to reserve rooms over the phone, all of the rooms, the sizes of the room, and the equipment available in the rooms were known at the time the speech interface was designed. Such speech inter-

faces can be grammar-based like many successful commercial systems (e.g., BeVocal [18], Tellme [17]), they can be statistical language model-based [4], or they can use a combination of both [21]. In our system, we do not know the objects (configurations) *a priori*, but we still want to be able to support names that are natural to the users.

Smart Homes and Home IT

Controlling lighting and other devices in the home is not new, nor is using speech to do so, but we believe using speech to control configurations of devices is novel. Quesada et al. address the interface challenge in the home machine environment [20] with a speech interface for lighting control in the home, but each light is controlled with individual commands, and the speech interface is designed specifically for the particular set of lights. Mozer et. al. [16] have studied predictive light automation, which as mentioned above, could be combined with a speech interface for situations when the predictive light automation does not do a good job with the prediction. Juster and Roy use situated speech and gestures to control a robotic chandelier, Elvis [10]. Their work focuses on how to move the chandelier arms to achieve a desired lighting scene given input from photo sensors. They focus less on the speech interface to the chandelier, and it is designed to support a static set of lighting configurations. We focus on the speech interface and supporting customized commands and configurations for each new set of devices without having to have a designer perform the customization.

Workspace Lighting Control

Commercial as well as research intelligent workspace lighting control systems exist, but they do not address the interface for manual control when users wish to override the system. Much of the work focuses on figuring out to what intensity to set the lights based on a variety of constraints, such as the amount of daylight present, the user's preferences and the building manager's requirements [25, 22]. Whether a manual interface is used in conjunction with an intelligent system (when the users want to override the system) [8] or a manual interface is used instead of an intelligent system (some users disable intelligent systems [13]), the manual interface should be natural and easy to use. If it is not easy to use, users will be less likely to use it, either resulting in a waste of energy (the default intensity is too high), or a potential decrease in productivity (the default intensity is too low). In Illuminac, our focus is on designing a natural, easy to use manual interface for the lighting control. Illuminac can be used on its own, or in conjunction with an intelligent lighting control system.

Multimodal Interfaces

It is well established that speech and gesture work well together [14], and many of our users mentioned they would like to be able to point at the lights for which they would like to turn on. Wilson's work on the XWand [26] demonstrates such a system where the user can control different devices by pointing at the device and saying a predetermined utterance. The position of the wand is determined with the use of at least two calibrated cameras and a blinking LED at the

end of the wand. The speech recognition system uses a simple command and control style grammar. Wilson acknowledges "while speech clearly has enough expressive power to make the wanted unnecessary, relying on speech alone can be difficult in practice." We aim to address this difficulty. Wilson also acknowledges "the acceptance of the XWand or a related device is limited by the limitations imposed by the installation and calibration of the cameras." The authors address this limitation by trying a wand with audio feedback to aid in pointing tasks without cameras available to track the position of the wand. They find it is possible without the cameras, but the pointing takes more thought on the part of the user and requires the targets be more widely spaced, which is not the case in our workspace. Our proposed approach could be combined with the XWand to develop a multimodal gesture and speech interface for configuration tasks where the configurations are not able to be predetermined and can also be used in cases where the calibration and installation of cameras is not possible.

Applications of Matrix Factorization

In recent years, alternative factorization methods such as least-squares NMF (Non-negative Matrix Factorization) have found favor over singular value decomposition (SVD) in cases where patterns are not orthogonal. This choice is especially true when the patterns appear "non-negatively," which is indeed the case for both light intensities and for word frequencies in user commands. Thus, NMF seems like a natural candidate for SNAC analysis. Furthermore, NMF has been shown to be superior to SVD for pure text clustering [11] or for image segmentation [11], which is similar to our light grouping problem. Barnard et al. match words and pictures [1] using an "aspect" probabilistic model, which is another type of factor model. Other candidate factor models include Latent Dirichlet Analysis (LDA) [2] and GaP [5]. These methods add prior probabilities to the factors and use likelihood measures (rather than least squares) to fit the original data. We did not use these methods because (i) when there is enough training data, the factor priors have little or no effect and (ii) least-squares NMF has been shown to produce better (more independent) factors compared to KL-divergence (likelihood) fitting methods [11].

CONCLUSION AND FUTURE WORK

This paper introduced an approach to the challenge of simultaneously learning names and the configurations named in natural speech interfaces for environment and device control. We demonstrated this challenge and introduced a solution in the domain of workspace lighting control. Our system allows users to have customized "light switches" for their lighting scene configurations through the use of natural language speech commands. The users train the system not only on their personalized commands but also on the configurations to which their commands refer.

Results from the live deployment of our natural speech interface for workspace lighting control are promising, and we look forward to seeing how well our approach works in other workspaces, especially ones foreign to the authors. In addition to other workspaces, we would like to evaluate our

approach with other devices in the workspace in addition to lights such as projectors, and speakers. In other ubicomp domains, such as home environment control and mobile device configuration.

ACKNOWLEDGMENTS

We would like to thank the study participants from all three studies and Lora Oehlberg for her sketch in Figure 3. The first author was supported by a National Science Foundation Graduate Student Fellowship during this work.

REFERENCES

1. K. Barnard, P. Duygulu, N. de Freitas, D. Forsyth, D. Blei, and M. I. Jordan. Matching words and pictures. *Journal of Machine Learning Research*, 3:1107–1135, 2003.
2. D. Blei, A. Ng, and M. Jordan. Latent Dirichlet Allocation. In T. G. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems 14*. MIT Press, 2002.
3. D. Bohus. Roomline: a spoken dialog system that provides assistance for conference room reservation and scheduling, as of 19 Sept 2007. <http://www.ravenclaw-olympus.org/roomline.html>.
4. P. F. Brown, J. Cocke, S. A. Della Pietra, V. J. Della Pietra, F. Jelinek, J. D. Lafferty, R. L. Mercer, and P. S. Roossin. A statistical approach to machine translation. *Comput. Linguist.*, 16(2):79–85, June 1990.
5. J. Canny. GAP: a factor model for discrete data. In *SIGIR*, pages 122–129. ACM Press, 2004.
6. N. Dahlbäck and A. Jönsson. Wizard of oz studies – why and how. In *IUI*, pages 193–200, 1993.
7. DALI. Digital addressable lighting interface (dali), as of 24 June 2008. http://en.wikipedia.org/wiki/Digital-Addressable_Lighting_Interface.
8. S. Escuyer and M. Fontoynt. Lighting controls: a field study of office workers’ reactions. *Lighting Research and Technology*, 33(2):77–94, June 2001.
9. A. Janin, D. Baron, J. Edwards, D. Ellis, D. Gelbart, N. Morgan, B. Peskin, T. Pfau, E. Shriberg, E. Shriberg, A. Stolcke, A10, C. Wooters, and A11. The icisi meeting corpus. In *ICASSP*, volume 1, pages I–364–I–367 vol.1, 2003.
10. J. Juster and D. Roy. Elvis: situated speech and gesture understanding for a robotic chandelier. In R. Sharma, T. Darrell, M. P. Harper, G. Lazzari, and M. Turk, editors, *ICMI*, pages 90–96. ACM, 2004.
11. D. D. Lee and S. H. Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, October 1999.
12. D. D. Lee and S. H. Seung. Algorithms for non-negative matrix factorization. In *NIPS*, volume 13, pages 556–562, 2000.
13. J. Love. Field performance of daylighting systems with photoelectric controls. In *Proceedings of Right Light Three, 3rd European Conference on Energy Efficient Lighting.*, pages 75–82, 1995.
14. D. McNeill. *Hand and mind: What gestures reveal about thought*. University of Chicago Press., 1992.
15. T. Moore, D. J. Carter, and A. I. Slater. A field study of occupant controlled lighting in offices. *Lighting Research and Technology*, 34(3):191–202, September 2002.
16. M. C. Mozer. Lessons from an adaptive house. In D. Cook and R. Das, editors, *Smart environments: Technologies, protocols, and applications*, pages 273–294. Wiley & Sons, 2005.
17. T. Networks. Fundamentally improving how people and business use the phone, as of 19 Sept 2007. <http://www.tellme.com>.
18. Nuance. Formerly bevoval, a leading provider of hosted application systems for customer self-service, as of 19 Sept 2007. <http://www.bevoval.com>.
19. P. Placeway, S. Chen, M. Eskenazi, U. Jain, V. Parikh, B. Raj, M. Ravishankar, R. Rosenfeld, K. Seymore, M. Siegler, R. Stern, and Thayer. The 1996 hub-4 sphinx-3 system. In *In DARPA Speech Recognition Workshop*, Chantilly, VA, February 1997.
20. J. F. Quesada, F. Garcia, E. Sena, J. A. Bernal, and G. Amores. Dialogue management in a home machine environment: Linguistic components over an agent architecture. In *Spanish Society for Natural Language Processing*, volume 27, pages 89–98, September 2001.
21. M. Rayner, P. Bouillon, B. A. Hockey, N. Chatzichrisafis, and M. Starlander. Comparing rule-based and statistical approaches to speech understanding in a limited domain speech translation system. In *TMI*, 2004.
22. V. Singhvi, A. Krause, C. Guestrin, J. H. Garrett, and S. H. Matthews. Intelligent light control using sensor networks. In *SenSys ’05: Proceedings of the 3rd international conference on Embedded networked sensor systems*, pages 218–229, New York, NY, USA, 2005. ACM.
23. A. Technologies, as of March 2008. <http://www.aduratech.com>.
24. M. Weiser and J. S. Brown. The coming age of calm technology. In *Beyond calculation: the next fifty years*, pages 75–85. Copernicus, New York, NY, USA, 1997.
25. Y.-J. Wen, J. Granderson, and A. M. Agogino. Towards embedded wireless-networked intelligent daylighting systems for commercial buildings. In *In SUTC, 2006*, pages 326–331, Washington, DC, USA, 2006. IEEE Computer Society.
26. A. Wilson and S. Shafer. Xwand: Ui for intelligent spaces. In *CHI*, pages 545–552, New York, NY, USA, 2003. ACM Press.