

# The Glass Infrastructure — Using Common Sense to Create a Dynamic, Place-Based Social-Information System

*Catherine Havasi, Richard Borovoy, Boris Kizelshteyn,  
Polychronis Ypodimatopoulos, Jon Ferguson, Henry Holtzman,  
Andrew Lippman, Dan Schultz, Matthew Blackshaw, Greg Elliott*

■ Most organizations have a wealth of knowledge about themselves available online, but little for a visitor to interact with on-site. At the MIT Media Lab, we have designed and deployed a novel intelligent signage system, the Glass Infrastructure (GI), that enables small groups of users to interact physically through a touch-screen display with this data and to discover the latent connections between people, projects, and ideas. The displays are built on an adaptive, unsupervised model of the organization and its relationships developed using dimensionality reduction and commonsense knowledge that automatically classifies and organizes the information.

The GI is currently in daily use at the lab. We discuss the AI model's development, the integration of AI into a human-computer interaction (HCI) interface, and the use of the GI during the lab's peak visitor periods. We show that the GI is used repeatedly by lab visitors and provides a window into the workings of the organization.

Each semester, many people visit the MIT Media Lab: visiting faculty and students, press and nonprofit leaders, and representatives of the companies that sponsor our research. These visitors find a building full of more than 300 active research projects and more than 20 research groups. It's difficult to figure out which research projects one should visit, and how the projects and groups are interrelated and relate to the lab's dominant research themes. This problem is especially notable during "member's week," when hundreds of visitors come to the lab for a week of talks and demos. Even if one visits during member's week with a plan and agenda, it can be difficult to remember and manage all of the interactions and projects one has had during the week. "Visiting the Media Lab is like drinking from a firehose for those of us visiting," a sponsor commented during our user survey.

We have built the Glass Infrastructure (GI)—a new type of visitor information kiosk for spaces rich in interesting artifacts, such as stores, museums, and research laboratories. We take the novel perspective that an information kiosk should preserve key aspects of exploring collections of artifacts in physical space, such as object permanence and social engagement, while also giving users access to alternative spatial organizations of those

artifacts. In the particular context of our research lab—where physical space is structured around organizational groupings of researchers—we wanted to let visitors simultaneously explore research in a thematically organized space, to help them discover work related to their interests across research groups. Kiosks are context dependent, so the information on display at each kiosk relates to the location of the kiosk and the user or users standing in front of the kiosk.

The Glass Infrastructure allows visitors to catalogue, share, and organize the projects they've visited as they travel through the lab. It organizes these projects without input from the students who upload the material and keeps the categorization up to date as the lab's areas of focus change. It helps users share projects they have visited and recommends projects based on previous interests.

The main contribution of this article is an AI approach to automatically computing the structure of a thematically organized map of artifacts, using only short descriptions of the artifacts that are likely already available. This map is then used to discover the evolving relationships among the artifacts, their creators, the research ideas, and the visitors. This approach eliminates a key obstacle to providing visitors with a current and coherent thematic interface to artifacts: the difficult and often contentious human process of organizing the artifacts thematically. This is particularly true in a context such as our research lab, where new projects may be added every day, and the ideas that connect the projects are themselves emerging, fading, and shifting.

Another important contribution is a collection of design decisions that harmonize the navigation of a space of artifacts with one's fingers on a large touch-screen kiosk with navigating with one's feet in a physical space. In part, we were motivated by the new building designed for our research lab that we moved into in 2010. The new Media Lab building was designed around the principles of visibility and open space. Much of the building is made of glass that lets one observe physically neighboring research groups while visiting one particular group. The building drove us to want to help visitors to also see research in the conceptual neighborhood of what they were already exploring by being able to see so much of the building's activity from any one location in the building.

## System Description

The Glass Infrastructure is currently deployed on 30 screens throughout the MIT Media Lab, with each location strategically chosen to be within proximity of one or more lab spaces. The research groups that occupy a particular lab space are displayed as the default view on the corresponding

screen. This serves as both the entry point for the user's experience of the system and as a means for the research groups to link their project collateral to the physical space they inhabit. From the proximal view, the user is able to shift focus to either a global view of the lab, drill down into greater detail of each research group, or direct the experience based on concepts that interest them. Users equipped with a radio-frequency identification (RFID) tag are able to "favorite" projects to see their conceptual overlap with other projects, see other users who expressed interest in the same projects, and review them after they leave the lab.

## Physical Components

The Glass Infrastructure has several major physical components: kiosks, RFID tags, existing databases, and a server. *Kiosks* are the touch-screen displays deployed throughout the lab that run the Glass Infrastructure. A UI application called the Project Browser is the central application of the Glass Infrastructure and is running on the majority of the screens. It displays the information and connections between the projects in the lab. Kiosks have an RFID reader that captures the identity of visitors and lab members who pass by and interact with the screens.

Each lab member and visitor is issued an *RFID tag* in their name badge with a unique identifier the system uses to understand who is interacting with it.

Another major component is the collection of *existing databases* that provide information about projects and groups and the people associated with those groups. Another system provides information about people (lab members and visitors) who interact with the system through their assigned RFID tags.

The Glass Infrastructure is powered by a *server* that runs the back-end processes and that organizes the information from the existing database and feeds it to the Project Browsers on the kiosks. It also collects and processes the data from the RFID tag readers and records information about the usage of the system. The physical system and information flow is illustrated in figure 1.

## User Experience

When a visitor arrives, she receives a badge with an RFID identifier. When she approaches one of the many touch-screens around the building, the system recognizes her and prompts her to log in. She is presented with a view of the lab relative to the geographic location of the screen where she has begun her engagement; a list of research groups that occupy the surrounding space are presented as buttons on the screen.

Suppose she is deciding which project to view

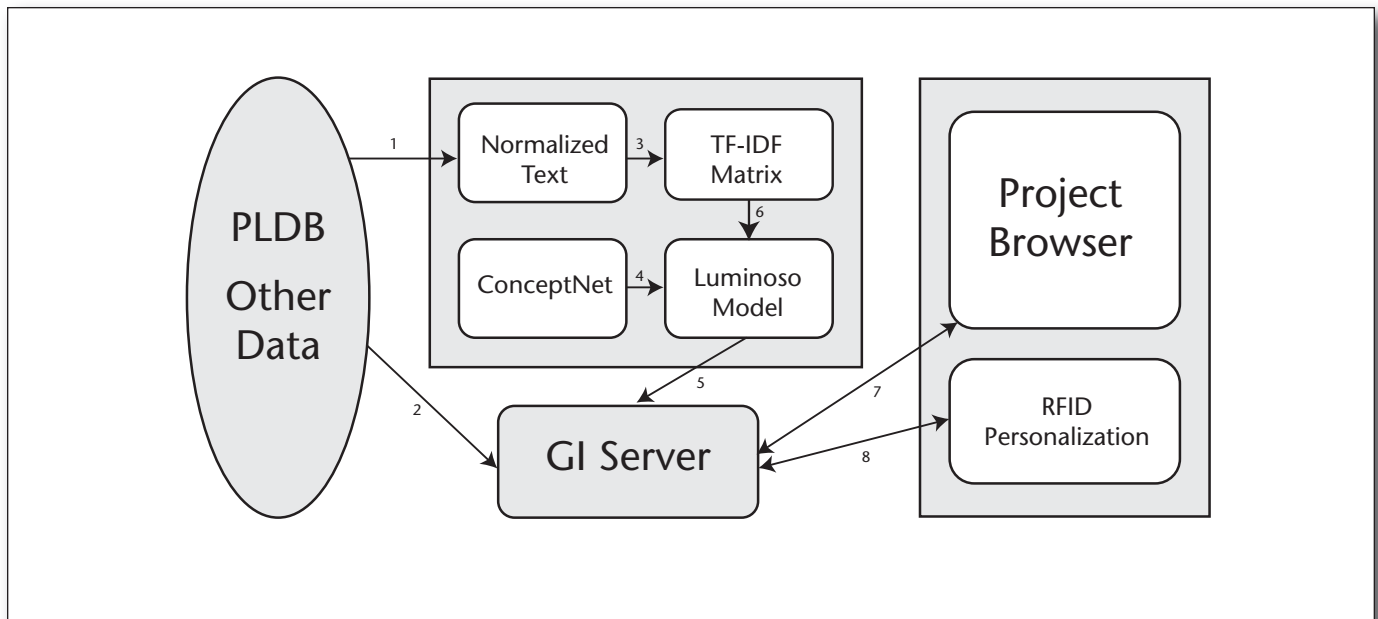


Figure 1. A System Diagram of the Information Flow in the Glass Infrastructure.

Edges are labeled for reference in the text.

next. She would be able to zoom into a research group to view the people and the projects that constitute it. She picks a project she is interested in and selects it. She is then able to see a video demo of the project in action, alongside other related projects in this group and throughout the lab, taking into account her location, interests, and currently selected project. The machine-learning system clusters projects according to their semantic connections with several predefined topic areas that represent the major interest areas of lab visitors. She becomes interested in a topic area and chooses to browse other local projects in that topic area. The topic area places the projects she has seen before in context, helping her plan what to see next.

She can then “favorite” projects she is interested in, and these favorites, or “charms” (as they are referred to in the GI), are stored in a personal profile she can access after the visit. Favorite projects are demarcated with heart-shaped badges.

When a second user approaches the screen and logs in, that user’s favorites are also displayed, and items that the two users have in common are highlighted by having their heart-shaped badges stacked. Users are able to exchange favorites with each other. One of the core innovations of this user experience is that it is social; we have successfully encouraged the sharing of a single screen, a space historically reserved for solitary interaction. Figure 2 shows a screen detailing the Viral Communications research group, with two users logged in.

During public events, we display a large leaderboard, or community mirror (Borovoy et al. 1998), showing a series of interesting statistics about the activity of the Glass Infrastructure including which screens have the most activity, which projects are the most popular, who has favorited the most projects, and a ticker that displays the social activity of the system. When a user favorites a project, this is reflected on the leaderboard counts. When a user exchanges a favorite with another user, the activity feed is updated with a statement such as “Pol Y. shared CitiCar with Boris K. at the BT Laboratory.”

After leaving the Media Lab, the users are able to log in to a “portfolio” through the Internet, where they see a log of the projects they favorited and the people they shared screens with. The users may connect with the researchers whose projects interested them and the people they met while interacting with the system. This living model of the lab aids visitors in retaining more of the ideas they encounter here.

## Technical Approach

To create the Glass Infrastructure we needed a way to model the relationships that make up the Media Lab. In the lab there are relationships among projects, researchers, research groups, the lab’s main directions, and ideas or concepts that inform a project. Each of these connections is important in building an accurate picture of the lab’s interactions. We want to be able to build this space in a

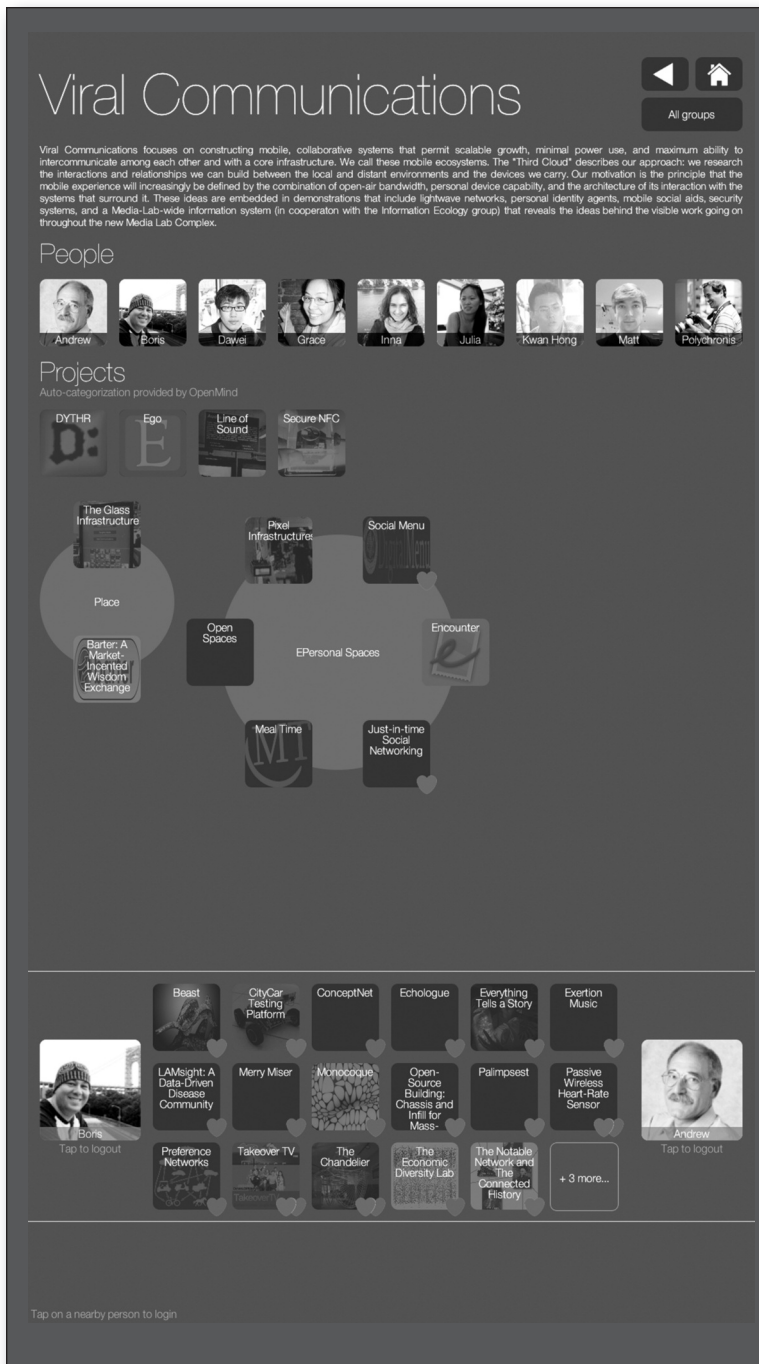


Figure 2. Two Users Logged into a Screen.

The image shows the research group detail view with favorited projects indicated by heart-shaped badges, in the “molecule” configuration clustered by Open Mind Common Sense (OMCS).

dynamic and unsupervised manner. For this purpose, we used the Media Lab’s abstracts as an existing base of text to create a semantic space to model the connections that make up the Media Lab. We represent this text in a way that is easy for a computer to understand in an automatic fashion

by focusing on the relationships between the words and concepts in each abstract. Then we add some world knowledge, in the form of common-sense statements, to help in the text understanding. The result combines this knowledge to form a multidimensional space where concepts, people, groups, and projects are all represented as vectors. From that space we retrieve information relevant to lab visitors—dynamically creating their presence in the vector space by creating a vector from the projects they have chosen as favorites. We then use the vector space to determine the relevance of objects in the space to each other—determining which projects are similar, which projects would be good fits for a lab visitor, and which projects fit which lab themes. Additionally, we have designed a user interface that makes this system easy and social to interact with.

The following subsections discuss our approach to interface design, our methods for extracting semantic information from the text base, and for assessing similarity of user interests with that knowledge.

### User Interface Design

The interface design of the GI adroitly navigates some challenging information-presentation scenarios. We are able to manage a dynamically shifting display that rearranges a large amount of information in response to where the screen is located and who is interacting with it, in a clear and consistent manner. The three key aspects of this design include object permanence, interests drive perceptual salience, and social engagement.

**Object Permanence:** Technological artifacts in a museum or our lab space—even more than the people who make and curate them—become the key organizers of visitors’ understandings of our lab. Therefore, the main elements in the GI user experience are persistent representations of these artifacts. Every view is a view onto “artifact space”—there are no maps or organization charts unless they are used to organize and show artifacts. Furthermore, the representations of these artifacts are designed to persist across the user experience whenever possible. When a user is viewing the projects in a particular group and clicks one of them to see thematically related projects, ones that are already in view simply move over into their new arrangement.

**Interests Drive Perceptual Salience:** Our interests affect what we see (Egeth and Yantis 1997). When we explore artifacts in physical space, our interests make relevant properties of those artifacts seem more salient, and this makes navigating the space more manageable. In the flat world of the GI, where small images are substituted for the artifacts themselves, there is much less information to guide user salience. Therefore, we explicitly intro-





*Figure 3. A Student Shares Her Work with a Media Lab Sponsor Using the Glass Infrastructure.*

duced a means to call attention to artifacts related to users' expressed interests, called "charms." Users can charm artifacts they're interested in, and these charms persist across time (users wear RFID badges that identify them to each screen, and allow their charms to be saved and reloaded). In this way, charming an artifact is similar to bookmarking or favoriting. The key distinction is our ubiquitous use of charms throughout the interface to call attention to people, groups, ideas, and projects related to users' interests (somewhat similar to Google's +1 buttons).

**Social Engagement:** We know that visitors often explore our laboratory in groups, and that they learn a lot from each other. Therefore, we designed the GI to also support social exploration. When two people are in front of a screen, both their charm sets display in a way that makes it immediately apparent what interests they have in com-

mon. They may then exchange charms to collaboratively explore their common interests (figure 3).

Besides welcoming, informing, and entertaining visitors, we believe a place-based information system also helps organizations with a common but difficult problem: the siloing of information related to the organization's various subcomponents. Frequently, people are unaware of what co-workers are doing down the hall or across the building.<sup>1</sup> Our system through its interaction design and data model is designed to expose and encourage the exploration of the links between people, what they are doing, and the underlying ideas.

Such a system needs to be flexible, responding to changes in the lab's research focuses and the daily addition of new content and projects, and it should help the user to not be overwhelmed by the large amounts of information and demos available during their visit to the lab. Underlying the Glass

Infrastructure is a text-understanding system that drives this interaction. Rather than rely solely on the graphic design and position of the stations, we combine this with an analytical process that dynamically tailors the information to the ongoing activities in the lab, as recognized from text about those activities, and ultimately personalizes the perspective presented to the people using the system at any moment. We believe this is a novel application of such techniques to open, public displays. In the following sections, we describe the manner by which the data is processed, presented, used, and altered.

The core visual mechanism for handling the changes in screen configuration is animated transitions that crystallize selected information and keep the user's gaze on known entities as the remainder of the screen rearranges. For example, the default state of a screen is a list of research groups. Clicking the name of the research group initiates a transition whereby the text of the button grows and moves to the top left of the screen to become the header of the page. At all times the text remains visible even as the rest of the screen completely rearranges. Users are always aware of what they clicked, and after following the text to the top of the screen they can now take in the newly arranged display. If users have favorited projects that are part of the currently displayed research group or are conceptually related to the projects they are browsing, the icon of the project will transition from the user's collection to the inline presentation. Favorited projects are demarcated with a heart-shaped badge.

To support free exploration of all the projects in the laboratory from macro, proximal, or interest-based entry points we faced a daunting challenge with respect to navigation. We needed the users to always be aware of why they were seeing the current screen configuration and how they could return to something they had seen before. We solved this problem by employing a "molecule" metaphor for arranging the items onscreen. Projects within a research group are arranged around the circumference of a circle representing the concept by which they are all linked. Touching a project focuses the screen around that project; the content on the screen will relate to that project. The concept to which the project belongs and the other projects related to it by that concept also remain on screen at all times, while those no longer relevant fade out and new associations fade in. Transitions always maintain elements that don't change onscreen and maintain the user's gaze on the selected element. A back button allows users to retrace their steps, while selecting another project continues the users' exploration.

The GI makes heavy use of hardware-accelerated CSS3 animations to create an application-like

interface that runs within a WebKit browser. This platform allows us to leverage the portability of web standards while simultaneously developing a pageless, smooth animation dependent interface.

## Building a Semantic Space from Text

We can help users navigate the lab by recommending nearby projects to visit based on projects they've seen and favorited, tell visitors which projects are similar across the entire lab, and show them how a project fits into a series of labwide initiatives. This requires building a semantic understanding of the relations between projects and user interests.

## Inputs: Common Sense and Project Information

The model behind the Glass Infrastructure must be derived from material that the lab already uses and maintains. As a result, we have chosen the Media Lab's Project List Database (PLDB) as the source of our data. The PLDB's central feature is an abstract-length description of each project written by the students, faculty, and staff involved in the project. It also notes which researchers and groups are affiliated with each project.

The information in the PLDB is primarily unstructured natural language text, which presents several issues for automatically extracting the relations among projects. In particular, researchers have a strong propensity to refer to an object or action by many different terms, and to describe the same or similar problems using different terminology or expressions (for example, "cell phone," "cell," "phone," "mobile phone," "smart phone," "smartphone," and "mobile device" are all terms used in Media Lab abstracts to refer to the same thing). As texts, the project descriptions also rely on implicit background knowledge to communicate meaning, which is available to the reader but not necessarily to the machine. We will employ a commonsense reasoning component to "fill in the gaps."

A more traditional way of quickly processing unstructured text would be the "bag of words" approach, where each project is distilled to a list of how frequently each word appears in the description, and projects are compared to each other in that form. However, the PLDB data set is small, only 325 active projects with only a paragraph abstract each, and a bag of words approach would fail to find many fundamental connections.

We might also have employed tagging and classification to process text, but this would require motivating the lab to maintain and create the tags and to update the tags for existing projects. Additionally, handtagging would make it difficult to change focus areas or maintain project entries for students who have graduated.

## Building a Semantic Space

We process information in the PLDB using the Open Mind Common Sense (OMCS) platform (Havasi et al. 2009), which includes a representation of commonsense background knowledge called ConceptNet, and an analysis tool for creating and exploring semantic spaces from textual information, called Luminoso (Speer et al. 2010).

The creation of a semantic space begins by preprocessing the text in the project descriptions. Common words such as *and* and *the* are removed, verbs are stripped of inflections, and contractions are unpacked. We find the terms and phrases that appear in each document, and place the results in a matrix whose rows are Media Lab projects and whose columns are term occurrence counts. We normalize these counts across projects, so that words that are common across many documents are given less weight than those that appear in only a subset, and eliminate terms of sufficiently low count. Finally, we use this representation of projects to extract the implicit relations among the terms they contain. This data takes the form of a symmetric matrix indexed by terms, whose entries are occurrence counts adjusted by proximity information, that is, the counts are increased when terms appear in the same or adjacent sentences.

To add background knowledge to this domain-specific data, we use Luminoso to construct a similar representation of the relations between terms in ConceptNet. For its semantic model, Luminoso uses a technique called Spectral Association (Havasi, Speer, and Holmgren 2010) to find weighted connections between related words (such as “puppy” and “cute” or “sadness” and “crying”), even when these connections require multiple steps of association. Spectral Association constructs a symmetric matrix that expresses which terms are related, and expands this matrix through power iteration to find connections that require multiple steps of association.

We employ a technique called blending (Havasi et al. 2009, Arnold and Lieberman 2010) to merge the domain-specific term associations derived from the PLDB data with the more general appreciation of the relations among concepts obtained from ConceptNet. Blending performs joint reasoning over multiple sources of data simultaneously by taking advantage of the overlap between them. It inputs two matrices with some overlap in their row and/or column labels, and outputs a single matrix that can be processed by more standard techniques. In effect, blending allows associations / analogies to propagate over different forms of information.

We employ a form of dimensionality reduction called Singular Value Decomposition (SVD) to transform this blended matrix into a semantic model we call AnalogySpace (Speer, Havasi, and

Lieberman 2008). The SVD expresses concepts (such as “car” or “happiness”) and features (“made of metal”) in terms of a core set of axes, or principal components, that are automatically selected by the algorithm to represent the most variance in the data. This process simultaneously generalizes over related words and smoothes over missing information, while reducing the memory footprint of the model. The resulting space of axes provides a basis for computing semantic similarity of projects and user interests through operations on vectors that represent concepts or features in the compressed space.

Overall, our approach to building a semantic space works on large sets of related words at one time, which gives the system the ability to generalize from extremely sparse input data; it can construct a meaningful representation of a topic even when that topic is identified by a single word (Cambria et al. 2010).

## Finding Similarity

If we want to find projects similar to a given project, we simply find the project vector in our PLDB-space and find projects whose vectors are closest in the semantic space. Here, closeness can be measured by the angle between the two vectors—projects with a small angle between them are similar.

We can personalize our recommendations by taking into account the projects that have been favorited by the user or users currently using the kiosk. By adding the vectors for each of the user's favorited projects together we form a new vector called an ad hoc category vector (Havasi, Speer, and Pustejovsky 2009). We let that vector represent the user's preferences in the space and then use the similarity metric described above to find projects similar to that category vector. This extends to recommendations for multiple users as well: we can simply combine their favorites, placing a greater weight on items that more than one user have favorited.

## Focus Areas

A central part of the Glass Infrastructure is its automatic categorization of projects into several categories that serve as the lab's focus areas or directions. These categories, which cut across multiple lab groups, were predefined by the Media Lab's leadership. As the lab's focus areas change from time to time (the categories have changed twice during our deployment), and new projects arise that need to be categorized, it's important that calculating these categories be a quick and unsupervised process.

We represented the categories as canonical documents in Luminoso. In the graphical version of Luminoso, a canonical document serves as a “signpost” in the visual representation of the semantic



space and as a focus area during Luminoso's statistical analysis of the space (Speer et al. 2010). In the Glass Infrastructure, a canonical document serves as a focal point for a cluster; the resulting clustering helps the user understand the big ideas behind a project and how that project relates, through that big idea, to labwide research focuses and other projects.

Canonical documents behave like the ad hoc categories (Havasi, Speer, and Pustejovsky 2009) described above, but for a set of words that is defined by Luminoso's users—in this case the architects of the Glass Infrastructure. Luminoso can often discover categories simply given the category's name, but since many of the lab's category names are idiosyncratic, we provided several words that would help define, and in some cases target, the model's definition of the particular lab direction. For example, the category "reability" focuses on helping those with disabilities and contains words such as "prosthesis" and "autism," which are focuses of the lab's research in that area.

Before the launch of the Glass Infrastructure, researchers and students were invited to check the categorization of their projects. Most of the Media Lab's researchers showed up to check the categorization and all researchers agreed with their categorization.

### Summary: Text Analytics from Project Descriptions

Here, again, are the steps involved in creating and reasoning with a vector space from the information contained in the PLDB and ConceptNet. We'll tie these steps back to the labeled edges in figure 1.

Query the PLDB to discover the researchers, research group(s), and project description for each active project (labeled Edge 1). Use ConceptNet to find the terms and phrases that appear in each document such as "cell phone" or "user interface."

Remove stop words and inflections, but preserve a reverse mapping for the purpose of outputting the phrases. For example, "Opera of the Future" should appear in that form, not as the system's internal phrase of "opus future."

Account for negative relations: words that appear between a negation token (such as not) and punctuation are counted negatively. The system can handle double negatives, should they occur.

Create special words inside each document to identify the research group and researchers involved in the project, such as Lifelong Kindergarden or minsky. Place the counts in a matrix whose rows are Media Lab projects and whose columns are terms. This is the "project matrix."

Normalize the matrix with tf-idf,<sup>2</sup> so that words that are common across many documents are given less weight than those that appear in only a subset (Edge 3).

Drop terms that appear fewer than a threshold number of times, to simplify the model. We use a threshold of three.

Create a new symmetric matrix containing the implicit associations between terms in the documents, A. Whenever terms appear in the same sentence, increment their entry in A by 1. Whenever terms appear in adjacent sentences, increment their entry in A by 1/2.

Use blending to combine the symmetric matrix with the ConceptNet association matrix. This and the two following steps are described earlier (Edges 4 and 6).

Find the spectral decomposition of the blended matrix; this is the semantic space (PLDB-space). Create a projection of the project matrix into the semantic space, by multiplying them along the axis of terms (Edge 5).

Discover the relationships between people, projects, and focus areas (Edge 7) that are displayed alongside pictorial representations of these elements from the PLDB (Edge 2).

Personalize the GI content for users as they acquire charms and interact with screens (Edge 8).

## Usage at the Media Lab

The system was first deployed in the Media Lab building in Spring 2010, in time for the biannual sponsor event that the Media Lab hosts every fall and spring. Each of these events is attended by 500–1000 people, including representatives from nearly a hundred companies and Media Lab students, faculty, and staff. The GI has now been in place across four sponsor meetings, and the bulk of the activity clusters around those events.

Figure 4 shows a more detailed look of GI activity across time and space during—and immediately following—the Fall 2010 Sponsor Event. Each number in this heat-map reflects the number of times a user clicked a UI object on a particular screen on a particular date. Higher numbers correspond to a darker square; for example, the chart reflects peak activity on October 14 when screen charm-5 was clicked 4134 times. This diagram shows that on October 14 and 15, 2010, there was substantial use of most of the GI screens. This use tapered off on the days that followed, as the Media Lab returned to a more normal activity pattern.

To use a uniform performance metric across all three events, we introduce "engagement"  $E$ , defined as:

$$\%VAF = 100 \times \left( 1 - \frac{\sigma_e^2}{\sigma_r^2} \right)$$

where  $E$  is the engagement metric for variable  $V$  during event  $e$ ,  $|V|$  is the observed value for variable  $V$ , and  $N_e$  is the number of visitors participat-



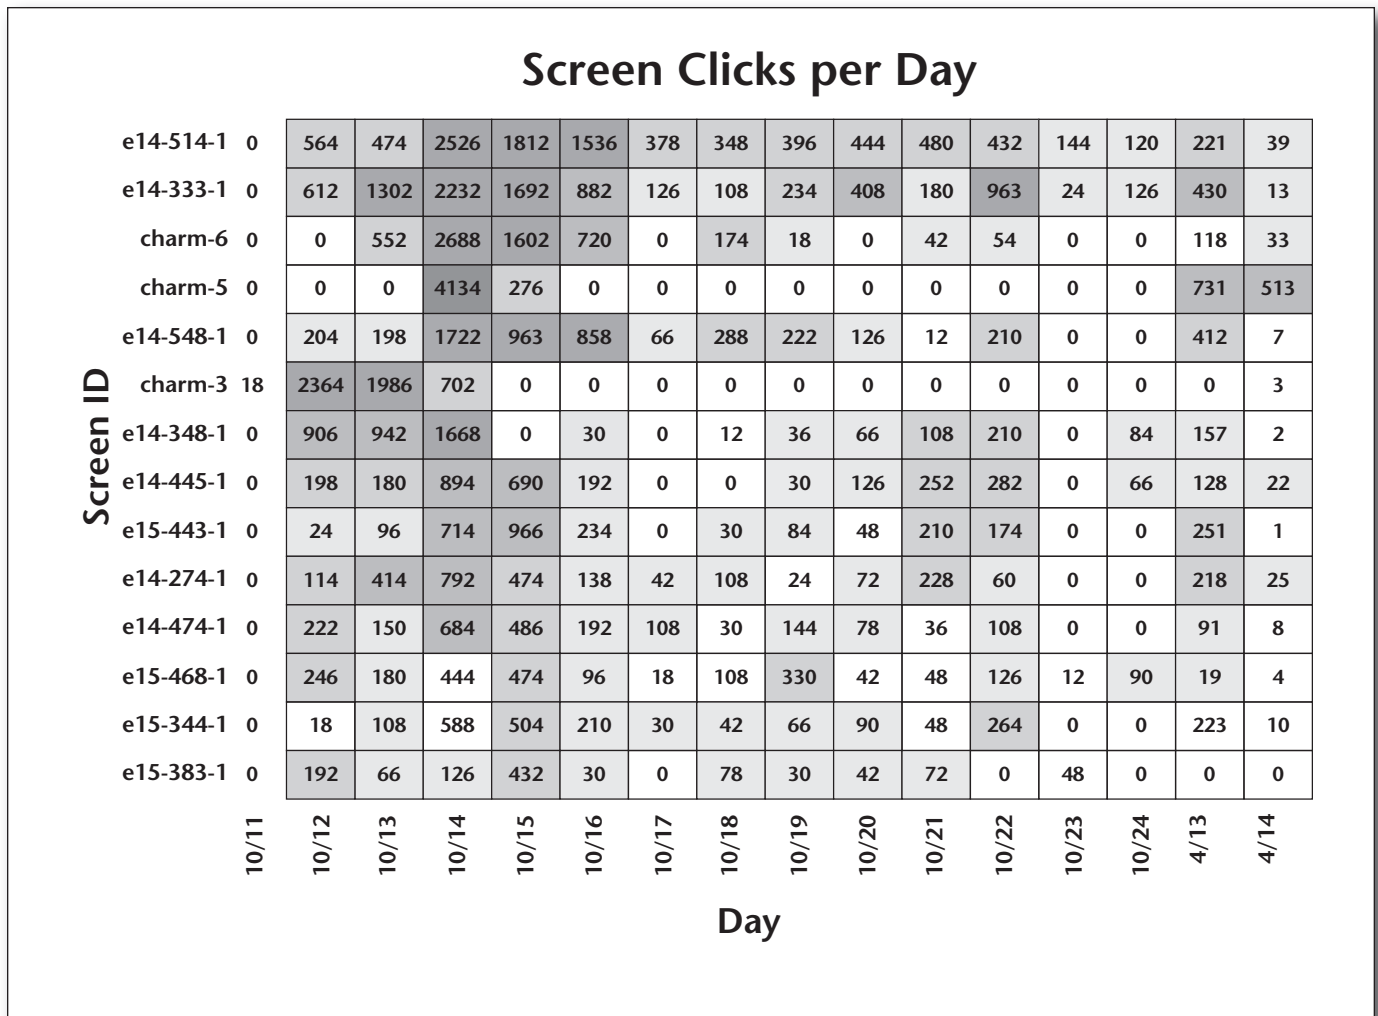


Figure 4. Heat-Map.

This heat-map shows GI usage at different screens over a period of two weeks during our second member's week deployment in fall 2010.

ing in event *e*. Essentially, this metric normalizes the observed value of each variable shown in table 1 by dividing by the number of visitors at each respective event.

Figure 5 shows that per capita engagement between fall 2010 and fall 2011 has almost doubled. While the absolute number of navigation events remained constant in the last event, user engagement was richer in terms of location ubiquity (more logins at multiple locations by the same user), more in-depth exploration of the displayed information (more users and favorited projects), and more social engagement (more cologins). We postulate this increase was due to repeat users returning to member's week and using the GI more often and at greater depth. Additionally, a new format to member's week removed many of the overview presentations, requiring visitors to orient themselves more than in previous weeks.

Based on the data and by our own observations, the GI added substantial value during all three member's week periods. There were often crowds of people gathered around the displays, and we witnessed many conversations between sponsors about favorite projects while collaboratively exploring the GI and exchanging charms with each other. GI use and value fell off steeply in the times between member's weeks, when we found it more challenging to maintain the system, to get visitors set up with RFID badges, and to orient them to the system's functionality.

During our third deployment, we performed a small user study of visitors who used the screens and trip reports. Seventeen users took our study; for some of these users it was their first time at the lab, while others had been to multiple (some up to 10) member's weeks. There was no significant difference in response from those who were visiting

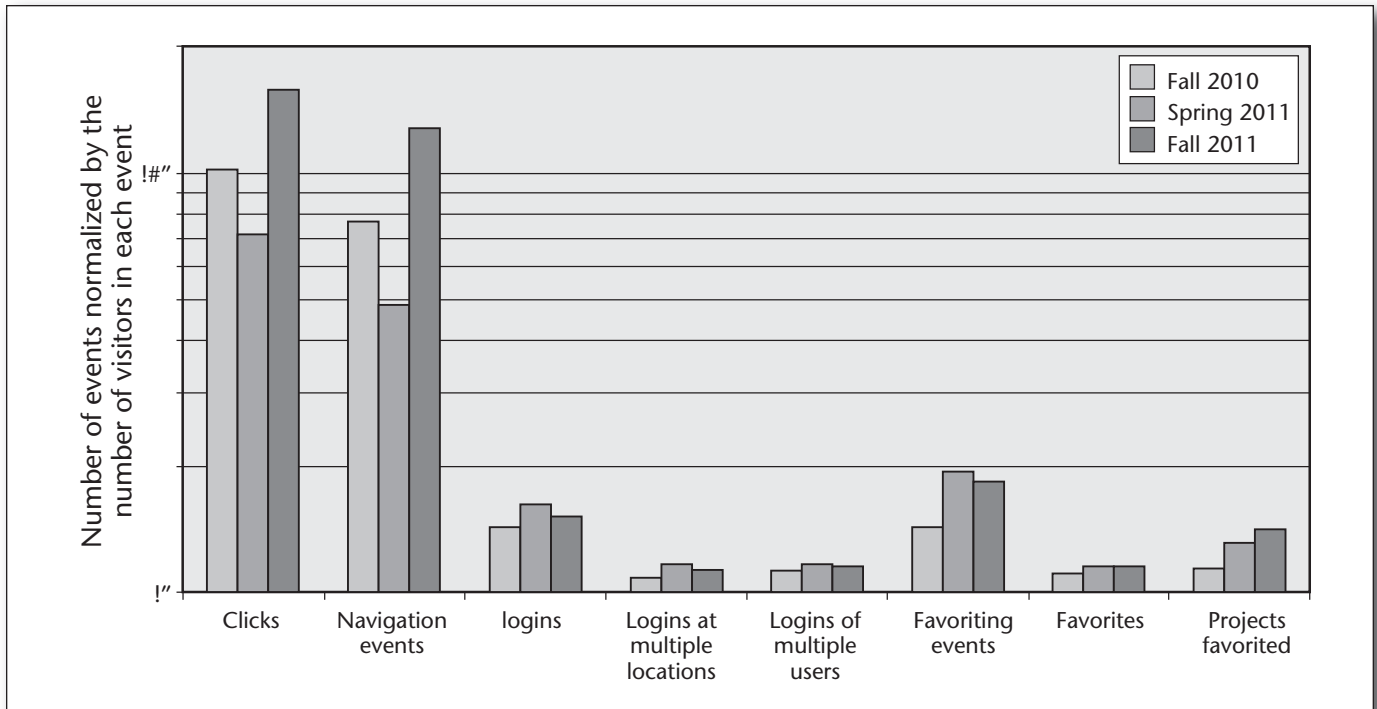


Figure 5. Usage Statistics.

This figure depicts usage statistics for the events of fall 2010, spring 2011, and fall 2011, in which there were 882, 590, and 474 visitors respectively. The variables observed in each column are normalized by the number of visitors in each event. As a result, the graph reveals the degree of per capita engagement  $E_e$  across each variable. While the absolute number of events observed during fall 2010 and fall 2011 are approximately the same, the number of visitors in the latter were almost half and therefore users in the latter event exhibited higher engagement.

	Fall 2010	Spring 2011	Fall 2011
Click events	8167 (9.26)	3694 (6.16)	7041 (14.85)
Navigation events	5915 (6.71)	2317 (3.86)	5604 (11.82)
Unique user logins	383 (0.43)	374 (0.62)	247 (0.52)
Logins at multiple locations	75 (0.09)	99 (0.16)	61 (0.13)
Cologins	114 (0.13)	100 (0.16)	77 (0.16)
Favoriting events	389 (0.44)	568 (0.95)	398 (0.84)
Favoriters	93 (0.11)	91 (0.15)	79 (0.17)
Unique projects favorited	124 (0.14)	188 (0.31)	199 (0.42)
Average time spent on-screen	25 seconds	25 seconds	30 seconds
Number of visitors	882	590	474

Table 1. "Logins at Multiple Locations" Reflects the Ubiquity of System Use by One Person.

"Cologins" is the number of unique pairs that visitors formed when logging in at the same time on the same screen. "Favoriters" are the unique visitors who have favorited at least one project. Values in parentheses reflect the respective normalized value by the number of visitors in each event.

for the first time and the more experienced visitors.

In table 2, we can see the various GI goals and users' responses to those goals. Using the AI system to discover projects was the most highly rated goal, with an average rating of 4.0, followed closely by using the system to remember and organize visited projects. In table 3 we can see that 16 out of 17 users got value out of the Glass Infrastructure.

## Related Systems and Research

Our work was influenced by and is related to trends in both the private sector and academe. These come in two general forms: utility-based kiosks that are designed to accomplish or incentivize specific tasks, and information-based kiosks that are designed to provide contextual access to digital content. GI fills both of these roles by promoting exploration and social interaction while presenting custom displays based on the user's and screen's physical location.

The United States Library of Congress has begun to deploy a series of information kiosks that have appeared in several of its exhibits. The kiosks provide interfaces where patrons can identify their

favorite artifacts and associate them with a personal account. More information about these artifacts can then be accessed from home through the myloc.gov web portal.<sup>3</sup> The project shares several GI concepts, such as facilitated exploration through digital “mementos,” but there is no social component to this system.

Pepsi is producing a recycling kiosk system that allows consumers to receive incentives based on the amount they personally recycle. Their system allows users to log in and provide physical inputs at the kiosk that are immediately reflected in the user’s digital account. The kiosks utilize touchscreens, which allow users to identify themselves using email addresses.<sup>4</sup> Kiosk location is not a factor in this system, which speaks to the fact that its primary focus is not based on information display.

A system utilizing RFID linked to student schedules was deployed to explore the concept of personalized maps on a college campus. The system was placed in two locations within the same building. It could detect nearby students with RFID badges, look up their schedules, and direct them to their next destination (Luse 2010). Unlike GI, the information transfer at this kiosk is unidirectional. The data is customized based on location and user, but there is no way for the user to update that data from the kiosk itself.

## The Future of the Glass Infrastructure

The semantic space itself will need to be updated during the life of the project as lab research focuses shift and new projects are created. Since the information is read in an unsupervised manner from the internal database, the model is automatically updated periodically. This allows for domain knowledge to be added and updated in an unsupervised manner.

In addition to continuously improving the reliability and ease of use of the Glass Infrastructure, we are focused on making the system more valuable for everyday use. As reflected in the data, the GI added more value during member’s weeks than other, more normal patterns of lab activity in part because the RFID functionality is only active during member’s week. We are also adding a variety of new functions to address this, including making it easier for lab researchers to use the GI to show off their work to sponsors; making it easier to register visitors in the system; and making the GI more useful for locating and communicating with lab members.

Since this research began, we’ve been working to expand these techniques to areas such as ad hoc social networking and locationally aware user interfaces. The Glass Infrastructure has become a research platform here at the Media Lab and will

GI Goal	1	2	3	4	5
Discover new projects	1	0	2	9	5
Find out more about interesting projects	1	0	4	11	1
Connect with other sponsors	2	5	8	2	0
Connect with Media Lab researchers	1	0	3	13	0
Organize my visit	1	2	3	8	3
Remember interesting projects	1	2	1	7	6

*Table 2. Users Opinions of the GI’s Value Doing Different Tasks.*

A score of 1 indicates strong disagreement while a score of 5 indicates strong agreement.

	1	2	3	4	5
I got value out of the GI	1	0	0	11	5
The screens were helpful	1	0	0	12	4
The trip report was helpful	1	1	3	9	3

*Table 3. Users Evaluate the GI.*

A score of 1 indicates strong disagreement while a score of 5 indicates strong agreement.

continue to grow and expand as a way to bring place and artificial intelligence into a variety of research projects.

## Acknowledgments

A deployed project such as this could not be possible without the help of so many people, not the least of which are the hundreds of sponsors and Media Lab students who have used the Glass Infrastructure. Additionally, this project would not have been possible without the help and support of all of Necsys, who maintain the Media Lab’s computing infrastructure. We’d like to thank Deborah Widener and Felice Gardner who help to organize member’s week. Ken Arnold of the Common Sense Computing Initiative provided additional support on transforming Luminoso into a deployed AI application.

## Notes

1. Since deploying the GI at the MIT Media Lab, we have prototyped a GI installation at a large financial software company to help company employees connect the company’s own ideas, projects, and people. Additionally, at the request of our users, we have installed GI screens at the headquarters of a large bank that wishes to further engage the Media Lab in collaboration to enable them better to navigate the changing Media Lab projects.
2. Term frequency-inverse document frequency (tf-idf) is a measure of how interesting a word is in a document set by taking into account not only how related a word is to

a document but also how often the word occurs in the document set. This allows words that differentiate topics to come into the forefront and tends to deemphasize words that appear frequently, such as stop words.

3. See [www.cwhonors.org/CaseStudy/viewCaseStudy2009.asp?NominationID=175&Username=LiBry](http://www.cwhonors.org/CaseStudy/viewCaseStudy2009.asp?NominationID=175&Username=LiBry).

4. [www.greenerpackage.com/recycling/video\\_pepsico\\_unveils\\_reverse-vending%E2%80%99dream\\_machine%E2%80%99encourage\\_recycling](http://www.greenerpackage.com/recycling/video_pepsico_unveils_reverse-vending%E2%80%99dream_machine%E2%80%99encourage_recycling).

## References

Arnold, K., and Lieberman, H. 2010. Scruffy Cross-Domain Inference. In *Commonsense Knowledge: Papers from the 2010 AAAI Fall Symposium*. Technical Report FS-10-02. Menlo Park, CA: AAAI Press.

Borovoy, R.; Martin, F.; Vemuri, S.; Resnick, M.; Silverman, B.; and Hancock, C. 1998. Meme Tags and Community Mirrors: Moving from Conferences to Collaboration. In *Proceedings of the 1998 ACM Conference on Computer Supported Cooperative Work*, 159–168. New York: Association for Computing Machinery.

Cambria, E.; Hussain, A.; Havasi, C.; and Eckl, C. 2010. SenticSpace: Visualizing Opinions and Sentiments in a Multidimensional Vector Space. In *Knowledge-Based and Intelligent Information and Engineering Systems*, ed. Rossitza Setchi, 385–393. Berlin: Springer.

Egeth, H., and Yantis, S. 1997. Visual Attention: Control, Representation, and Time Course. *Annual Review of Psychology* 48(3): 269–297.

Havasi, C.; Speer, R.; and Holmgren, J. 2010. Automated Color Selection Using Semantic Knowledge. In *Commonsense Knowledge: Papers from the 2010 AAAI Fall Symposium*. Technical Report FS-10-02. Menlo Park, CA: AAAI Press.

Havasi, C.; Speer, R.; and Pustejovsky, J. 2009. Automatically Suggesting Semantic Structure for a Generative Lexicon Ontology. Paper presented at the Fifth International Conference on Generative Approaches to the Lexicon, 17–19 September, Pisa, Italy.

Havasi, C.; Speer, R.; Pustejovsky, J.; and Lieberman, H. 2009. Digital Intuition: Applying Common Sense Using Dimensionality Reduction. *IEEE Intelligent Systems* 24(4): 24–35.

Luse, A. 2010. Improving Direction-Giving through Utilization of an RFID-Enabled Kiosk. *International Journal of Advanced Pervasive and Ubiquitous Computing* 2(2): 45–56.

Speer, R.; Havasi, C.; Treadway, N.; and Lieberman, H. 2010. Finding Your Way in a Multidimensional Semantic Space with Luminoso. In *Proceedings of the 15th International Conference on Intelligent User Interfaces*. New York: Association for Computing Machinery.

Speer, R.; Havasi, C.; and Lieberman, H. 2008. AnalogySpace: Reducing the Dimensionality of Common Sense Knowledge. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence*. Menlo Park, CA: AAAI Press.

**Catherine Havasi** is a computational linguist working to help computers understand people. She leads the Digital Intuition group at the MIT Media Lab and organizes the Common Sense Computing Initiative (OMCS), an international open-source AI project.

**Richard Borovoy** is an entrepreneur and research scientist with a passion for face-to-face community building, social networking, and mobile technology. He was previously a research scientist at the MIT Media Lab and now works at Google.

**Boris Kizelshteyn** is an alumni of the MIT Media Lab currently working at Nokia.

**Polychronis Ypodimatopoulos** is a PhD candidate at the MIT Media Lab. He received his BSc in computer science from Aristotle University of Thessaloniki (1999), MSc in computer science from King's College London (2000), and MSc in media arts and sciences from MIT (2008). He has worked in the research departments of IBM, Ferrari, and Intracom.

**Jon Ferguson** is a senior research systems programmer at the MIT Media Lab. Ferguson has worked at the Media Lab for 15 years.

**Henry Holtzman** is the MIT Media Lab's chief knowledge officer, codirector of the Digital Life consortium, and director of the Information Ecology research group. In addition, Holtzman codirects the lab's CE 2.0 initiative, a collaboration with more than 40 Media Lab sponsor companies to formulate the principles for a new generation of consumer electronics that are highly connected, seamlessly interoperable, situation aware, and radically simple.

**Andrew Lippman** served as the founding associate director of the MIT Media Laboratory and is currently a senior research scientist at MIT. He directs a research consortium titled Digital Life that addresses bits, people, and community in a wired world. In addition, he is a principal investigator of the MIT Communications Futures Program. He holds 11 patents in television, digital image processing, and interface technologies. His current research interests are in the design of scalable wireless systems for personal expression and generalized mobile systems that cooperate with the physical environment.

**Dan Schultz** is a graduate student at the MIT Media Labs Information Ecology group as well as the MIT Center for Civic Media. Before coming to the lab Schultz was trained to think in terms of systems at Carnegie Mellon University and was awarded a Knight News Challenge grant in 2007 to write about "Connecting People, Content, and Community" on the PBS Idea Lab.

**Matthew Blackshaw** is interested in making spaces more dynamic, social, adaptive, and engaging. Prior to having all kinds of fun at the Media Lab he worked at Microsoft where he helped build the next generation of user interfaces for Windows products.

**Greg Elliott** is a PhD candidate at the MIT Media Lab.