

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, Chaki Ng
MIT Media Lab

Abstract

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 physically interact 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 developed using dimensionality reduction and common sense knowledge which automatically classifies and organizes the information.

The GI is currently in daily use at the lab. We discuss the AI models development, the integration of AI into an HCI interface, and the use of the GI during the labs peak visitor periods. We show that the GI is used repeatedly by lab visitors and provides a window into the workings of the organization.

There is an evident trend toward adding interactivity and computing to the physical spaces in which we work, live and travel. In part this may be due to the ease with which places can now be networked, and in part by the decreasing cost of installed equipment. This is evident in signage, instrumentation of transportation systems, digital cameras, and portable applications that connect to these devices. Many of them expose information about the environment to individuals, either directly or by means of a mobile device.

We are building the Glass Infrastructure (GI) – a new type of visitor information kiosk for spaces rich in interesting artifacts, such as stores, museums, and research laboratories. Our approach preserves 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 displayed at each

kiosk relates to the location of the kiosk and the user or users standing in front of the kiosk.

The main contribution of this paper 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 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 navigating 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 inhabited in 2010. Full of glass and open space that let one observe physically neighboring research groups from the perspective of visiting one particular group, the building drove us to want to help visitors also see research in the conceptual neighborhood of what they were already exploring. We present the key aspects of this design below, and then discuss more about their motivation and execution later.

1. **Object permanence:** Infants discover early on that fleeting images on their retinas are actually caused by objects that persist over time, and these objects soon become key organizers of their world view (Piaget 1977). In a similar way, technological artifacts in our lab space – even more than the people who make them – become the key organizers of visitors' understandings of our lab. Therefore, the main elements in the GI user experience are persistent representation of these artifacts. Every view is a view onto "artifact space" – there are no maps, or org charts unless they are used to organize and show artifacts. Furthermore, the representation 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 on one of them to see thematically-related projects, ones that are already in view simply move over into their new arrangement.
2. **Interests drive perceptual salience:** Our interests effect



Figure 1: A student shares their work with a Media Lab sponsor using the Glass Infrastructure.

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 introduced 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.

3. **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 common. They may then exchange charms to collaboratively explore their common interests.

Besides welcoming, informing, and entertaining visitors, we believe a place-based information system also help organizations with a common but difficult problem—the siloing of information related to the organization's various sub-components. Frequently, people are unaware of what co-workers are doing down the hall or across the building.¹ Our system through its interaction design and data model is designed to expose and encourage the exploration of the links

¹Since deploying the GI at the MIT Media Lab, we have prototyped a GI install at a large financial software company to help company employees connect the companies' own ideas, projects, and people. Additionally, at the request of our users, we have install at GI screen at the headquarters of a large bank to enable them to better navigate the changing Media Lab projects.

between people, what they are doing, and the underlying ideas.

Such a system needs to be flexible, responding quickly to changes in the lab's research focuses, adding new content and projects quickly, and helping the user to not become overwhelmed by the large amounts of information and demos available upon 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 (if not the first) 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.

Modeling the Media Lab

To create the Glass Infrastructure (GI), we needed a way to model the relationships that make up the Media Lab. In the lab there are relationships between 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 have chosen to combine a "base layer" of common sense information with domain-specific Media Lab information using a technique called blending (Havasi et al. 2009). Blending extends a machine learning technique across multiple different domains so that the contributions of the domains to the resulting model are equal (Arnold and Lieberman 2010). Common sense background knowledge allows one to find semantic connections in textual information by bridging the gaps between concepts (such as "bakery" and "coffee shop" or "bride" and "cake") which would typically require mining very large data sources to discover (Lieberman et al. 2004). The Media Lab has only 360 active projects which traverse the boundaries of technical domains, providing a challenge for simple bag-of-words or keyword-based techniques. We use blending to create a model which balances background knowledge with learned Media Lab-specific information.

We use blending over a dimensionality-reducing machine learning method called Spectral Association (Havasi, Speer, and Holmgren 2010) to create a multi-dimensional model representing the semantic relatedness of projects in the lab, or a "semantic space". Closeness in this space signifies similarity of people, projects, focus areas, and ideas (Speer, Havasi, and Lieberman 2008).

We can help users navigate the lab by recommending nearby projects to visit based on projects they've seen and liked, tell visitors which projects are similar across the entire lab, and show them how a project fits into a series of lab-wide initiatives. This living model of the lab aids visitors in retaining more of the ideas they encounter here.

We want to be able to build this space in a dynamic and unsupervised manner. For this purpose, we have chosen to use the Open Mind Common Sense (OMCS) common sense

platform(Havasi et al. 2009) and it's associated dimensionality reduction based inference algorithms(Kenneth Arnold and Robert Speer and Catherine Havasi and Jason Alonso and Jayant Krishnamurthy 2008) to provide a consistent basis of background knowledge. We will explain the creation and use of the GIs model following a discussion of user interactions.

The Glass Infrastructure

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 users 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 an 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.

User Experience

Visiting the Media Lab can be a confusing and overwhelming process; it can be difficult to remember and manage all of the interactions and projects not to mention decide which projects to visit or understand the connections between projects. “Visiting the Media Lab is like drinking from the firehose for those of us visiting,” a sponsor commented during our user survey. The GI is designed to help solve this problem.

When a visitor arrives, they receive a badge with a RFID identifier. When a user approaches one of the many touch screens around the building, the system recognizes them and prompts them to log in. They are presented with a view of the lab relative specific to the geographic location of the screen where the user begins their engagement; a list of research groups that occupy the surrounding space are presented as buttons on the screen.

Let's say the user is deciding which project to view next. The user is 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 and other related projects both in this group and others throughout the lab taking into account her location, interests, and current project. The machine learning system clusters projects according to their semantic connections with several predefined topic area which represent the major interest areas of lab visitors. The user 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 are stored in a personal profile she can access after the visit. Projects that are her favorites are demarcated with heart shaped badges.



Figure 2: Two users logged into a screen. Image shows research group detail view with favorited projects indicated by heart shaped badges, in molecule configuration clustered by OCMS

When a second user approaches the screen and logs in, their 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 they exchange 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, users are able to log into a “portfolio” via 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.

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 core visual mechanism for handling the changes in screen configuration is animated transitions that crystallize selected information and keep the users gaze on known entities as the remainder of the screen re-arranges. For example, the default state of a screen is a list of research groups. Clicking on 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 re-arranges. The user is 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 the user has 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 users collection to the inline presentation. Items that are favorites are demarcated with heart shaped badges.

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 user 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 users gaze on the selected element. A back button allows the user 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 dependant interface.

Related Systems and Research

Contributions related to our work have been made in both the private sector and academia. 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 users and screens physical location.

The United States Library of Congress has begun to deploy a series of information kiosks that have appeared in several of their 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 portal². 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 which 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, which are immediately reflected in the users digital account. The kiosks utilize touch screens, which allow users to identify themselves using email addresses³. 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 uni-directional. 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.

Common Sense Backend

The model behind the Glass Infrastructure must be derived from material which the lab already uses and maintains — the adoption of a new system, such as tagging, would require motivating the lab to maintain and create the tags and to update the tags for existing projects. Additionally, hand-tagging would make it difficult to change focus areas or maintain project entries for students who have graduated. 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.

For each project in the PLDB, we can locate the project's authors, the research group in which the project resides, and a short abstract description of the project. The information in PLDB is updated by researchers continually, with two primary pushes before the lab's bi-annual sponsor week.

The information written in the PLDB is primarily in unstructured natural language and thus is not readily processed by a computer. 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

²See <http://www.cwhonors.org/CaseStudy/viewCaseStudy2009.asp?NominationID=175&Username=LiBry>

³See <http://www.greenerpackage.com/recycling/video%20pepsico%20reverse-vending%20%80%98dream%20%80%99%20encourage%20recycling>

toolkit.

Building a Semantic Space

In order to build a multi-dimensional representation of the concepts with Media Lab projects using the mostly textual information in the PLDB, Luminoso must process the text and blend this text with the information in ConceptNet. We need to find and isolate patterns in all facets of our data set: those in the domain specific PLDB, those in common sense (like those discovered in AnalogySpace), and those that are only apparent in a combination between common sense and PLDB information. We can use these patterns to find correlations between objects in our space, and from those similarities we can find meaning in our data which will help the Glass Infrastructure's users find projects which fit their interests. Luminoso and OMCS have been used successfully to combine information in this way before, for examples see (Havasi et al. 2009),(Havasi, Speer, and Pustejovsky 2009),(Mueller 2006).

Since OMCS's concepts were created from and are stored in natural language, it's fairly easy to make connections between the PLDB text and the information in the common sense database. These techniques build on the power of LSA (Deerwester et al. 1990) which performs dimensionality reduction on the bag of words models described earlier, both by modifying the algorithm used to perform the decomposition and by including a basis of common sense knowledge. Common sense can recognize when two words are semantically close to one another even when this closeness is not apparent from the document concurrence. Thus common sense would be able to tell that "user interface" and "computer" are related and it would be able to distinguish between different different topic areas that exist independently of the input data, such as "action verbs", "household items", "computer terminology", and "things people don't want". This allows more of the intuitive or rough meaning of the project descriptions to be captured in the vector space.

To begin to build the space, some simple pre-NLP processing is applied to the project descriptions. Common words such as "and" and "the" are removed, verbs are stripped of inflections, contractions are unpacked, and the space is normalized so that projects with longer descriptions are not given more weight in the final model.

For its dimensionality reduction, Luminoso uses a technique call Spectral Association (Havasi, Speer, and Holmgren 2010) which finds relational connections (such as "puppy" and "cute" or "sadness" and "crying") that would not readily appear in the standard AnalogySpace inference algorithm. Spectral association uses a symmetric matrix that expresses which terms are related, rather than AnalogySpace's term-by-feature matrix. This technique works on large sets of related words at one time, which gives Luminoso the ability to generalize from extremely sparse input data, even as sparse as a single word (Cambria et al. 2010).

In the PLDB data, terms are strongly associated when they appear in the same sentence, and more weakly associated when they appear in the same project description. ConceptNet is blended into this matrix as an association matrix of concepts by concepts. A model that overlaps the two matrices

is then found using the standard "rough blending" method (Arnold and Lieberman 2010)(Havasi et al. 2009) to create our semantic space

Processing the Text

There are a number of steps involved in creating a vector space from the information contained in the PLDB:

- Query the PLDB to discover the researchers, research group, and project description for each active project.
- 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, so that words which are common across many documents are given less weight than those which appear in only a subset.
- Drop terms that appear fewer than a threshold number of times, to simplify the model. We use a threshold of three.
- Scan over the documents again. 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 together the matrix above with the ConceptNet association matrix.
- Find the spectral decomposition of the blended matrix; this is your semantic space (PLDB-space).
- Create a projection of the project matrix into the semantic space, by multiplying them along the axis of terms.

Finding similarity

The result of these calculations is a vector space where similar projects are represented by nearby vectors (Speer, Havasi, and Lieberman 2008). 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 charms, or favorited projects, which have been

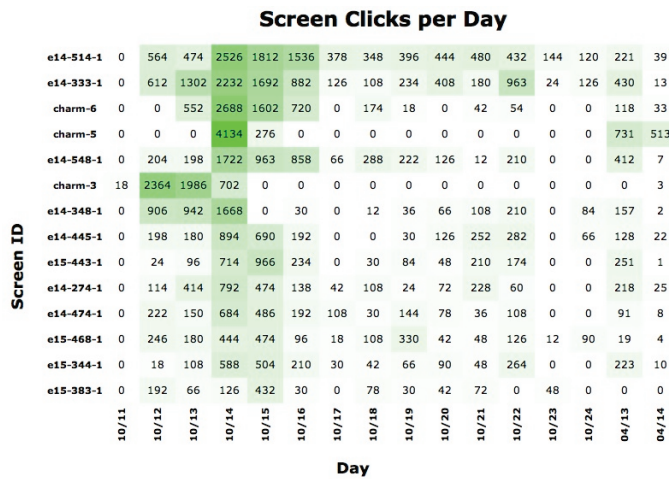


Figure 4: Heat-map showing GI usage at different screens over a period of two weeks during our second sponsor week deployment.

acquired by the user or users currently using the Charms system. By adding the vectors for each of the user’s favorited projects together to form a new vector called an ad-hoc category vector (Havasi, Speer, and Pustejovsky 2009). We let that vector represent the users 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 charms, placing a greater weight on charms that both users have.

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 pre-defined 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 is a quick and unsupervised (from the AI’s point of view) 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 lab-wide 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 cat-

egory’s name, but since many of the Lab’s category names are idiosyncratic, we proved several words which would help define, and in some cases target, the model’s definition of the particular lab direction. For example, the category “Re-ability” focuses on helping those with disabilities and contains words such as “prosthesis” and “autism”, focuses of the lab’s research in that area. In Figure 3, we can see the area defined by the canonical documents “Cognition” and “Music”.

Before the launch of the Glass Infrastructure, researchers and students were invited to check the categorization of their project. Most of the Media Labs researchers showed up to check the categorization and there were no complaints of incorrect categorization.

Usage at the Media Lab

	Fall 2010	Spring 2011
Click events	4883	3720
Unique user logins	215	201
Logins at multiple locations	60	99
Co-logins	95	111
Favoriting events	332	568
Favoriters	81	92
Unique projects favorited	141	221
Average time spent on-screen	25sec	25sec

Table 1: “Logins at multiple locations” reflects ubiquity of system use by one person. “Co-logins” 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.

The system was first deployed in the Media Lab building in May 2010, in time for the bi-annual sponsor event that the Media Lab hosts every Fall and Spring. Each of these events is attended by roughly one thousand people, including representatives from nearly a hundred companies and Media Lab students, faculty, and staff. The GI has now been in place across three 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 (really touched) a UI object on a particular screen on a particular date. Higher numbers correspond to brighter green coloring; for example, the chart reflects peak activity on October 14 when screen charm-5 was clicked 4134 times. This diagram shows that on October 14th and 15th 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.

These numbers indicate a relative increase in user engagement during the second event. Increased engagement is also evident from the usage statistics which show that while the absolute number of navigation event decreased in the last event, user engagement was richer in terms of location ubiquity (more logins at multiple locations by the same user), more

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 five 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 five indicates strong agreement.

in-depth exploration of the displayed information (more users and favorited projects) and more social engagement (more co-logins).

Based on the data and by our own observations, the GI added substantial value during all three Sponsor 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 Sponsor Weeks, when we found it more challenging to maintain the system, to get visitors set up with RFID badges, and to introduce them the systems 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 (sometimes up to 10) sponsor weeks. There was no significant difference in response from those who were visiting 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 positive 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 sixteen out of seventeen users got value out of the Glass Infrastructure.

The Future of the Glass Infrastructure

The AI model 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.

In the next version of the system, we plan to incorporate adhoc social networking for sponsors. The AI will not only recommend projects the user should visit, but people the user should talk to and other visitors which the visitor might want

to network with.

References

- Arnold, K., and Lieberman, H. 2010. Scruffy cross-domain inference. In *AAAI Fall Symposium on Common Sense Knowledge*, 2–7.
- 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, CSCW '98*, 159–168. New York, NY, USA: ACM.
- Cambria, E.; Hussain, A.; Havasi, C.; and Eckl, C. 2010. SenticSpace: Visualizing opinions and sentiments in a multi-dimensional vector space. In *Knowledge-Based and Intelligent Information and Engineering Systems*, 385–393.
- Deerwester, S. C.; Dumais, S. T.; Landauer, T. K.; Furnas, G. W.; and Harshman, R. A. 1990. Indexing by latent semantic analysis. *Journal of the American Society of Information Science* 41(6):391–407.
- Egeth, H., and Yantis, S. 1997. Visual attention: control, representation, and time course. *Annu. Rev. Psych.* 48(3):269–297.
- Grice, P. 1975. Logic and conversation. In *Syntax and Semantics 3: Speech Acts*. Academic Press. 41–58.
- 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.
- Havasi, C.; Speer, R.; and Alonso, J. 2007. ConceptNet 3: a flexible, multilingual semantic network for common sense knowledge. In *Recent Advances in Natural Language Processing*, 27–29.
- Havasi, C.; Speer, R.; and Holmgren, J. 2010. Automated color selection using semantic knowledge. In *Proceedings of the AAAI 2010 Symposium on Common Sense Knowledge*.
- Havasi, C.; Speer, R.; and Pustejovsky, J. 2009. Automatically suggesting semantic structure for a generative lexicon ontology. In *Proceedings of the Generative Lexicon Conference*.
- Kenneth Arnold and Robert Speer and Catherine Havasi and Jason Alonso and Jayant Krishnamurthy. 2008. Divisi: Commonsense Reasoning over Semantic Networks. Project web site.
- Lieberman, H.; Liu, H.; Singh, P.; and Barry, B. 2004. Beating common sense into interactive applications. *AI Magazine* 25(4):63–76.
- Luse, A. 2010. Improving direction-giving through utilization of an rfid-enabled kiosk. *International Journal of Advanced Pervasive and Ubiquitous Computing (IJAPUC)* 2(2):45–56.
- Mueller, E. T. 2006. *Commonsense Reasoning*. Morgan Kaufmann.
- Piaget, J. 1977. *The Essential Piaget*. Academic Press, New York.
- Speer, R.; Havasi, C.; Treadway, N.; and Lieberman, H. 2010. Finding your way in a multi-dimensional semantic space with luminoso. In *Proceedings of the 15th International Conference on Intelligent User Interfaces*.
- Speer, R.; Havasi, C.; and Lieberman, H. 2008. AnalogySpace: Reducing the dimensionality of common sense knowledge. *Proceedings of AAAI 2008*.