Grounding a Computational Model of Place

Matthew Hockenberry

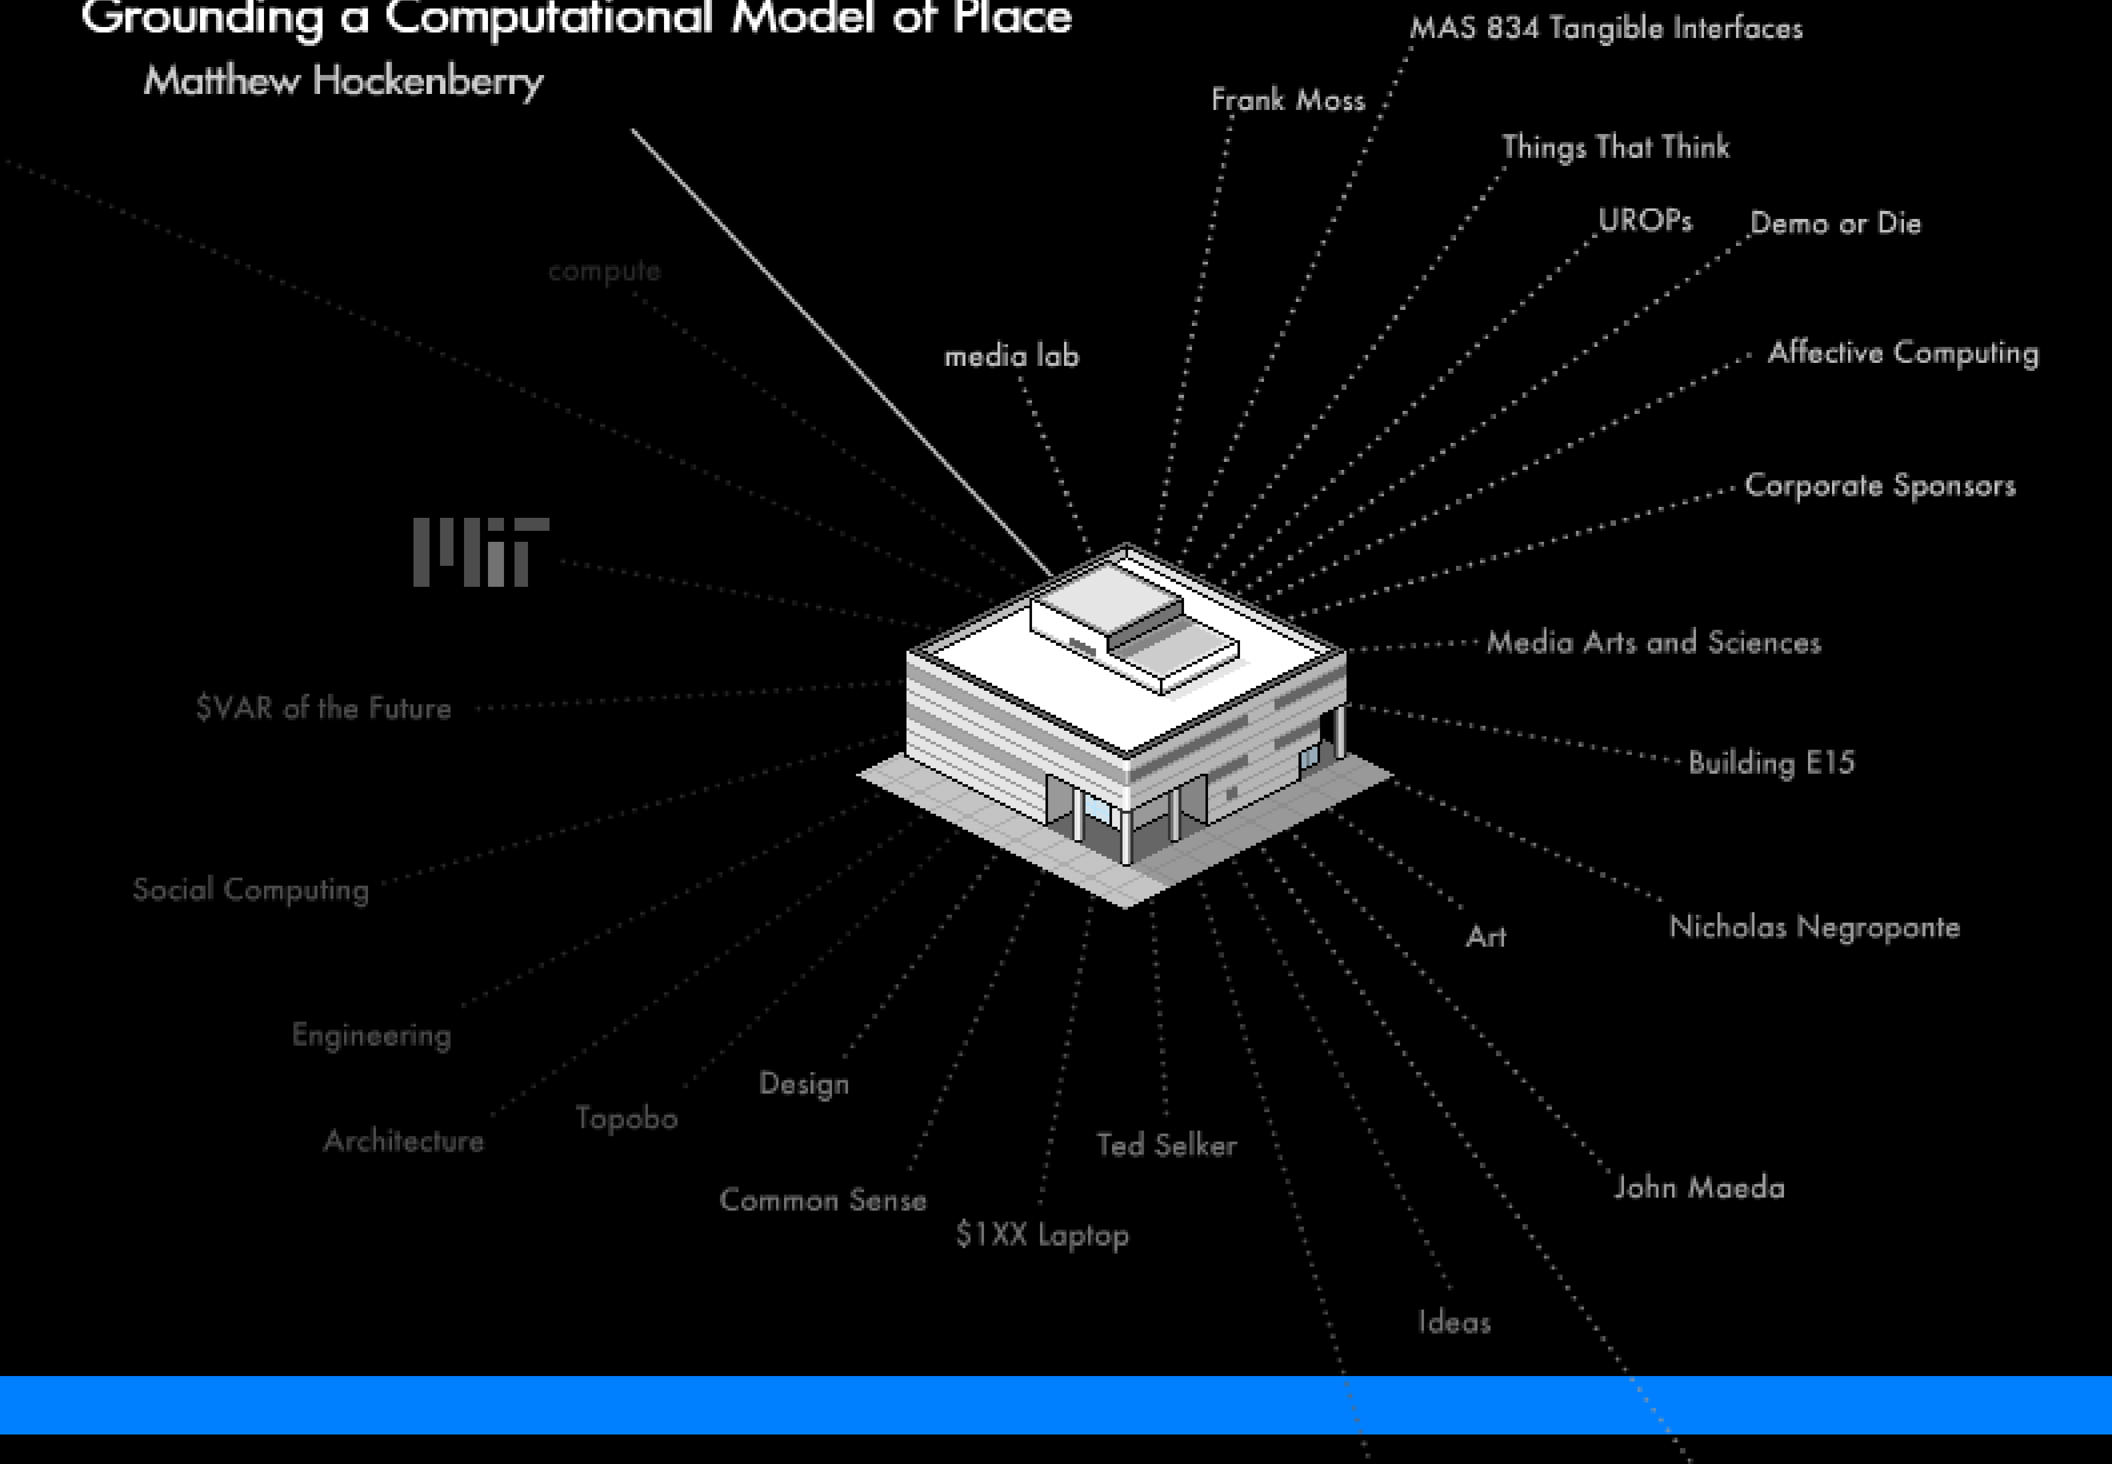

Grounding for a Computational Model of Place

Abstract
Places are spatial locations that have been given meaning by human experience. The sense of a place is its support for experiences and the emotional responses associated with them. This sense provides direction and focus for our daily lives.

Physical maps and their electronic decedents deconstruct places into discrete data and require user interpretation to reconstruct the original sense of place. Is it possible to create maps that preserve this sense of place and successfully communicate it to the user?

This thesis presents a model, and an application upon that model, that captures sense of place for translation, rather than requiring the user to recreate it from disparate data. By grounding a human place-sense for machine interpretation, new presentations of space can be presented that more accurately mirror human cognitive conceptions. By using measures of semantic distance, a user can observe the proximity of place not only in distance but also by context or association. Applications built upon this model can then construct representations that show places that are similar in feeling or reasonable destinations given the user's current location.

To accomplish this, the model attempts to understand place in the context that a human might by using commonsense reasoning to analyze textual descriptions of place, and implicit statements of support for the role of these places in natural activity. It produces a semantic description of a place in terms of human action and emotion. Representations built upon these descriptions can offer powerful changes in the cognitive processing of space.

Thesis Committee

Advisor
Ted Selker
Associate Professor of Media Arts and Sciences
Massachusetts Institute of Technology

Thesis Reader
John Maeda
Allen Professor of Media Arts and Sciences
Massachusetts Institute of Technology

Thesis Reader
Henry Lieberman
Research Scientist in Media Arts and Sciences
Massachusetts Institute of Technology

Table of Contents

Acknowledgements

This is not a singular work. The contribution of many colleagues, mentors, and scholars in the field can be felt in this work and their efforts should be recognized.

The author would like the many undergraduate students who have contributed to this work. Of particular note are Rob Gens and Jeff Hoff. Without these students it would have been unable to accomplish as much as we have. Other undergraduate students include Kirupa Chinnathambi, Brandon Yoshimoto, Zachary Ozer, Praveen Pamidimukkala, and Lauren McCarthy.

The author would also like to thank the following: Leonardo Bonanni for his constant advice and insight. Ernesto Arroyo for his perspective. Henry Lieberman for the conversations that helped to mentally frame and organize this work. Ted Selker for his mentorship.

Additional recognition should go to the MIT Media Laboratory, the sponsors of the lab, and the MIT iCampus Initiative for funding and supporting this work.

1. Introduction

Place is a complex notion. Very few of us human beings perceive the world as satellite photographs. We tend not to describe locations in terms of their latitude and longitude points. Instead we spend most of the time talking about things like “the lab,” “our house” or “that chic new restaurant on Newbury.” What is truly miraculous is that the people we are speaking to usually understand, often very quickly, exactly what spatial region we are talking about and receive a rich sense of what that area is like. The reason for this is simple: We are not sharing geographic spaces, we are sharing things we call places.

Place qualifies most of human behavior. “Where do we go?” “Go there?” “I don’t like that place, there’s no room” or “that place is boring.” How do we make these determinations and how do they affect our daily action? There is no simple answer, and over time attempts to quantify sense of place have met with difficulty. The question of *how can we craft artificial systems that understand place and make use of this understanding* remains unanswered. This thesis presents a computational model of place within this tradition, while also presenting an exploration in design and application that led to the construction of this model. At the heart of this work (Chapter Seven) is a system that provides crisp mechanisms for the identification and interpretation of palatial knowledge, directed towards its practical usage in spatial applications.

"The Argument"

1. There is a longstanding philosophical tradition of place-making through active perception as the primary means of spatial awareness and cognition.

2. This tradition has evolved into a rich psychological and neurophysiologic understanding of the role of place-making and sense of place in spatial cognition.

3. The tradition of constructive spatial representation has divorced itself from place-based conceptualizations in order to avoid the inherent subjectivity of place-construction.

4. This process has resulted in a modern tradition of deconstructed representations that fail to match common human representations of the spatial world, requiring significant effort in reconstruction.

5. Efforts to construct new spatial representations that successfully match human cognitive perceptions rest on the ability to relate personal subjective place sense with an artificial machine understanding of place.

6. Human experiential accounts of place exist and are available for machine collection, aggregation and semantic interpretation through commonsense understanding.

7. The aggregate of these accounts represents a general translation—not a deconstruction of placial thought—that is general and practical in purpose and able to be subjectively interpreted by humans.

8. The conception of a 'naïve geography'—a human common-sense conceptualization of space for machines can only be realized by considering these common-sense accounts of place sense.

9. A model of place, grounded in the concept of affordance and the interpretation of generalized human experiential accounts of place, can be used to construct commonsense representations of space by machines for average human consumption.

Notes on Structure

Following the introduction the thesis proceeds to a short section in Chapter Two on the practical motivation of understanding place and sense of place for human purposes. A short study into the philosophical tradition of place, and its grounding in neurophysiology and cognitive psychology, leads into a presentation of traditional representations in Chapters Three and Four.

Motivated by the failures and limitations of these representations, the thesis then explores a design iteration in Chapter Five of alternative representations that are closely modeled on human cognitive foundations. The success of these representations, as they relate to information search, is then presented. The limitations in success suggest the need for a richer understanding of sense of place as it relates to representation. While there are some interesting technologies and alternative approaches presented in Chapter Six, most have fundamental limitations.

From this, Chapter Seven presents the PlaceMap system and CampusMap as an application that exemplifies the model of light, rich place-building applications. PlaceSense, a general system design to organize and semantically interpret this information, is offered as a mechanism for interpreting the accounts formed through this application.

This body of work leads up to a presentation in chapter eight of Naïve Geography, a common-sense presentation of geography and spatial issues. Chapter Nine finally offers a computational model of place is offered to provide a concrete underpinning to this presentation and Chapter Ten addresses the implications of this in some concluding remarks.

2. Motivating a Human Place Sense

The task of developing a concrete computational theory of place requires some initial motivation. While the introduction claims to assert that sense of place qualifies most of human behavior, the obvious concern arises that 'humans are already good at getting on with their own sense of place, what utility is found in an artificial and possibly overly complex model?" This is a question that serves as a backdrop for most of the work in this thesis. It is true, indeed, that humans do have an innate sense of place. It is also true that often that sense of place is hard to quantify or categorize. However one must ultimately ask, why would we want to?

The suggestion that this thesis will offer is that, yes—humans do have a remarkably complex, efficient, and powerful understanding of the sense of a place. The most pragmatic motivation that can be offered is purely social. While a particular human may have a rich sense of a place, it is often difficult to translate that existing sense to another human. This difficulty increases when we look to larger groups of humans, and very rapidly becomes problematic when we begin to include our mechanical counterparts. Computers have been demonstrably ill equipped at understanding the things that come easily to humans, although they have had some success in other fields. As others have demonstrated, despite the fact that computer-kind has in its possession a rich supply of geospatial information, this information about space does not necessarily translate into having a rich knowledge of *place*.

This is awkward because these computer systems (particularly geographic information systems) are increasingly being used to communicate to human beings information necessary for decision making in spatial temporal context. Simple questions, such as "Where might one find a cup of coffee?" are certainly possible. It is the more complex kind of question, one that might be asked of a fellow human, that becomes difficult. The question, "Where can I have a relaxing cup of coffee?" is significantly more difficult. While the computer may be able to understand what is meant by relaxing, it would be hard pressed to identify the particular components of geospatial information that lead a place to be relaxing. To do so would also requiring significantly more detailed information at a much lower level of granularity than that which is available in most geospatial databases.

Consider purely human social sharing of placial information. To ask another human where the best cup of coffee might be found is a complex demand. The other party first has to understand the basic nature of the request (I want coffee, I want the best coffee). They then must determine what the qualifier 'best' means and they also need to have some understanding of the inquiring party. What is best for a busy person with no time to sit may not be the best for a relaxing smoker who keeps odd hours. Some degree of mental translation occurs, which results in a suggestion, based on the answering party's own knowledge and personal bias, adjusted to favor the asking party. Can machines learn what this means?

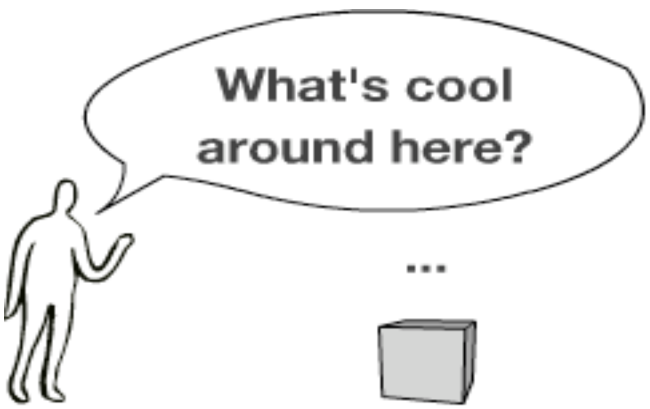

3. A Philosophical Tradition of Place

The concept of place and what place means in human thought and understanding has been a concern throughout the history of the intellectual discussion. This concept is not new, and many underlying themes continue to re-emerge in these discussions. Answers, and an appreciation for the role of these themes, however, remain lacking. The first question that naturally arises is, why is place a different concept than space?

Early philosophical thought introduced the concept of place as being distinct and unique from traditional notions of geographic space. In his work, Heidegger puzzles over the nature of a dwelling and notes that not all buildings are dwellings and not all dwellings are buildings (Heidegger 1971). This distinction is common to casual thought and is evident in common expressions such as "home is where your heart is" and a "house is not a home."

But while this distinction is a noted and important part of the human experience, a concrete definition of the concept has been notoriously difficult to formalize. This philosophical tradition begins with early explorations into semantics leading to understanding in situated cultural and social roles, as well as communication and work in theories of artificial intelligence. Ultimately these traditions reveal a complex dialogue of constructed, situated sense of place that is difficult to express in social constructions resulting from it.

The aspects of one's culture that are anchored in the body or daily practices of individuals form the notion of habitus, a notion introduced by Marcel Mauss and further developed by Norbert Elias (Mauss 1934; Elias 1978). The French philosopher, Pierre Bourdieu later appropriated and expanded this concept to include a broader (and arguably more rigid) notion of Habitus (Bourdieu 1977). Here one finds habitus described as “a sense of one's place... a sense of the other's place.” This describes both our perceptions of space and place and the impact of these perceptions on human action and socialization. This not only results from shaping environments, but simply from the experience and interaction within a place. This implies that a web of complex processes inseparably links the physical, the social and the mental that can explain processes of place making in relation to practices of the built environment.

This philosophical notion is descended from a longstanding tradition originating in the thought of Aristotle and of the medieval Scholastics, that was retrieved and reworked after the 1960s by sociologist Pierre Bourdieu to forge a dispositional theory of action (Wacquant 1992).

There are prominent and recurring themes within this dialogue that are valuable to consider in exploring issues of sense of place. The first is simply that places are human constructions, subjectively developed through a human actor's association with a particular spatial region. Places, in this tradition, are not environmental or geographical features so much as they are mental ones. Places are created, not found. They are not made by contractors and bulldozers, but rather through a process of social and personal interaction within a space.

A second, possibly disturbing result of this is the inherent subjective nature of place. This is a subjectivism that is constrained by social and cultural limitations (and less divergent than what might otherwise be believed), but a subjective perspective nonetheless.

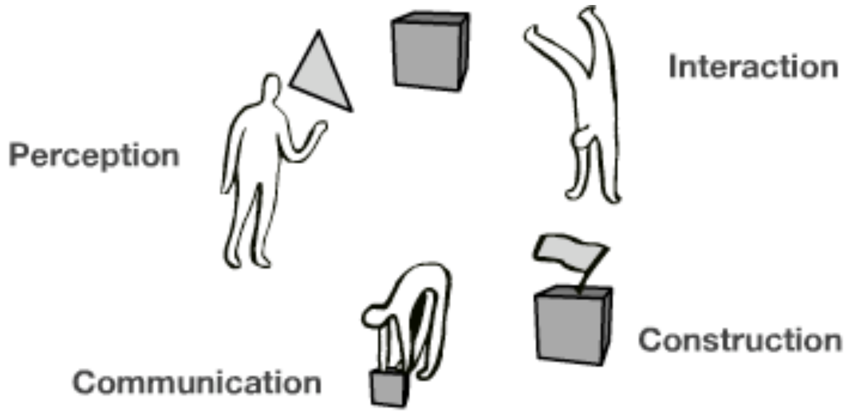

These two themes result in a constant undercurrent found in the dialogue on place since late classical thought. Place is subjective construction of human minds, created from interaction and perception within physical space and constrained by cultural and social perspectives. What is perhaps most fascinating is the eventual divergence of place from social constructions of space because of its subjectivity, despite the fact that place making remains the primary mechanism through which humans interpret, categorize, and communicate about the spatial world.

3.1 Perception of Place

Human perception has an extraordinary role in our organization and identification of places. Humans are born to be natural mapmakers and geographers, and children almost always experience a period of intense spatial representation and sharing (Siegel 1975). One question that arises is, exactly what are we attempting to understand—places or spaces? Compared to modern concepts of geography, the actual process through which the human mind seems to identify and organize space seems prone to error and exaggeration.

Cognitive maps are the personal, individual constructions people use to organize thought. They are rarely as precise as modern maps, and in fact they seem to serve an altogether different purpose. Errors identified within these cognitive maps (as compared to actual measures and maps of spatial regions) are almost always metrical, and only in very rare circumstances topological (Lynch 1960). This suggests that cognitive maps are primary tools for place making for human understanding and not necessarily tools for wayfinding and precise spatial communication. Indeed, the topological structure and gestalt used for spatial reasoning are very different from modern representations concerned with precision and shared perspectives (Stevens 1978).

The actual representation of these cognitive maps is quite a bit looser than one might imagine, but altogether reasonable in later reflection. When Lynch asked individuals to draw their own maps of the city in his *Image of the City*, he noticed that while each map appeared to be very different with regard to metrics of distance, direction and proportion, the general classification of relationships remained constant (Lynch 1960). Other researchers have noted the preservation of category, relative position and structure in both large-scale spatial reasoning and small scale, positing that for many simpler relations take the place of more precise representations. (Cohn 1997).

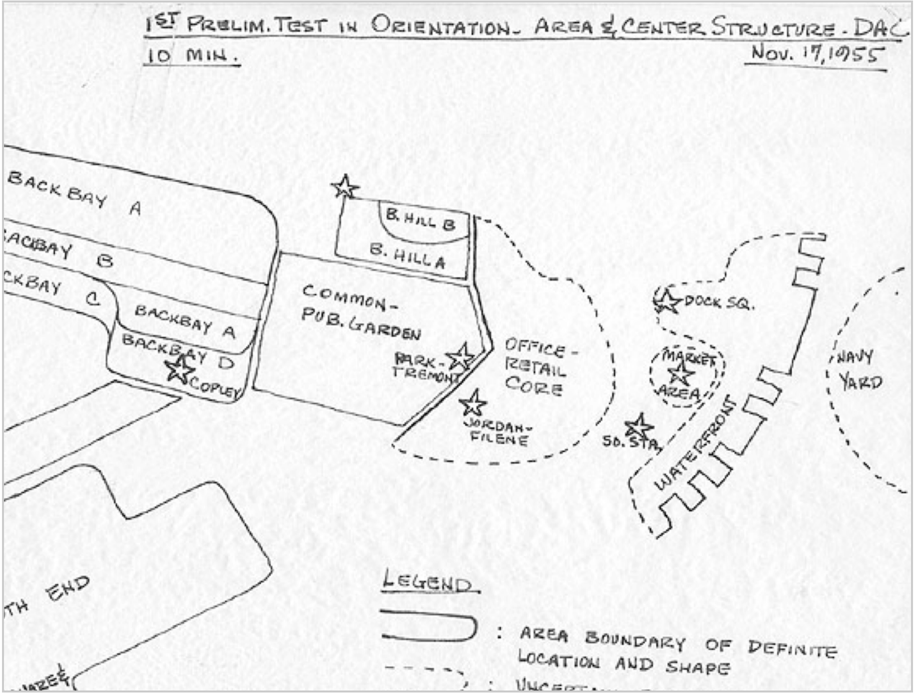

***Figure 1**. Sketch map of Boston, 1955 from Kevin Lynch's Image of the City. Although drawn to precision, differences in cognitive perception remain evident.*

One possibility, which is supported by this thesis, is that this results from a model of perception that is fundamentally tied to human action and interaction. Actions, interactions, and situated goals form the mechanism through which places are constructed, encoded and linked. This forms the native state through which humans attempt to construct representations, share, and interpret them. Not only does this seem to be a reasonable interpretation based on personal experience and empirical studies, there is ample physiological evidence that supports this conclusion.

3.2 Action in Place

In his Neurophysiology of Human Spatial Cognition, Mike Kahana notes that there is little evidence for purely allocentric (map-like) representations (Kahana 2004). Indeed there seems to be a direct link between the existing perception of place, the action or goal occurring, and the perception information being experienced (Kahana 2004). This further is elaborated upon in additional work that suggests the importance of order and topological distinctiveness of the goal or action. While many behaviors are complex, so long as they can be constructed as subgoals or related actions they still can be represented as a single encoded goal.

While this follows a tradition of embodied action in language and thought, the relationship between action and place construction is directly evident in internal representations. This is found in wayfinding, particularly with regard to the construction of landmarks, but also in information seeking. Decision making is rooted in location. In order to make effective decisions in spatial contexts, familiar as well as not, information organization based on action or perceived action becomes critical. While these judgments ("That looks like a nice place") can be relatively superficial, they are based on complex organizations of actively embodied places and their associated topologies.

While there is a history of associating action with place making, it is likely that the meaning, kind, and justification for this remains unclear. There is strong evidence, from a variety of sources, that people conceptualize geographic spaces differently from manipulable, table-top spaces (Kuipers 1978; Zubin 1989; Mark 1992a; Montello 1993; Pederson 1993; Mark 1995) and that the kind of active construction that occurs in those situations changes fundamentally when considering larger geographic spaces.

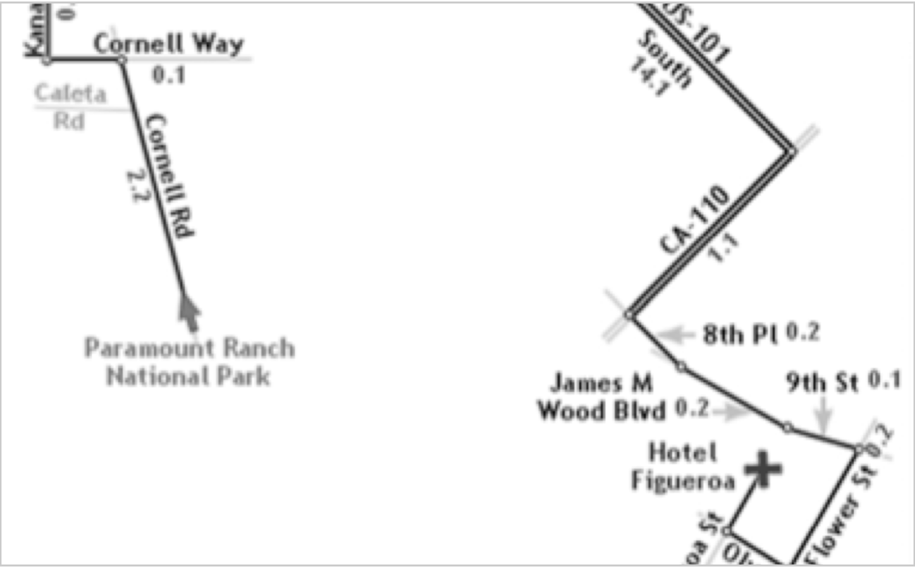

Figure 2. *Sketch map of LineDrive. LineDrive is a software system and associated algorithms that organize directive information by topological changes in action. Metrical distance is reduced in the representation to a supplemental property of the active descriptions (Agrawala 2001).*

The requirement of action suggests a concept of place outside of spatiality. Relph (1976) describes place as a unique instance of a pattern, composed of physical features and appearances, observable activities and functions—ritual routines. Places overlap and interpenetrate and this brings with it the notion of insideness vs. outsideness. “A gypsy camp is a place regardless of surroundings and locative coordinates” (Qtd. in Jordon 1998). In this tradition, Curry's (1996) theory of place asserts that places do not have natural boundaries that "existed long before man" but the place is a "location that has been given shape and form by people." Carl Sauer stated that there are no natural places. Places are human inventions.

3.3 The Personal Place

Since our world is spatial and three-dimensional, notions of space pervade our everyday experience (Tuan, 1977). A strong model of the individual is necessary to supplemental these conclusions. Particularly, this thesis focuses on behavior-based approaches rather than categorical models.

Regardless of the approach, the intuition is not obvious from a scientific standpoint. The emphasis on this unique personal perspective and even modeling it creates difficulties with larger social constructs and miscommunication is prevalent. The difficulty in relying on individual perspective in order to come up with a scientific concept of place is that it is necessary to accommodate the relatively objective view of the theoretical scientist (decentered) with the subjective view of the individual (centered) who directly experiences the place.

Entrikin (1991) suggests that "understanding place in a manner that captures its sense of totality and contextually is to occupy a position that is between the objective pole of scientific theorizing and the subjective pole of empathic understanding." Experiences of place involve perception, cognition, and affection. Again, this must be integrated both in location and meaning in the context of personal action.

3.4 The Social Self

The social presentation of space and place is perhaps the most developed of the traditions.

In the physical world, a place is simply a space that is invested with understandings of behavioral appropriateness, cultural expectations, and so forth. We are located in "space," but we act in "place." Furthermore, "places" are spaces that are valued. Harrison and Dourish continue Heidegger's thoughts: "The distinction is rather like that between a 'house' and a 'home'; a house might keep out the wind and the rain, but a home is where we live." (Harrison 1996)

Places provide a context for everyday action and a means for identification with the surrounding environment. They help inform our own sense of personal identity (Entrikin 1991) and make us identifiable to others. Behavior is linked to place. Judgments of what is appropriate are based on the place of an act (Therborn 1980; Cresswell 1996). Meanings given to places are a fundamental component of social interaction (Goffman 1959)

Shields has investigated the role of a social theory of spatiality focused on the role of space in cultural formation (Shields, 1992). Place is both broader and more specific than space. Conversely, the same location—with few changes in its spatial organization or layout—may function as a different place at a different time. "An office might act, at different times, as a place for contemplation, meetings, intimate conversation and sleep" (Harrison, 1996). This suggests that a place may be more specific than a space. "A space is always what it is, but a place is how it's used" (Harrison, 1996).

This meaning can change based on our social or cultural role. Carnegie Hall becomes a worksite for an electrician, a performance site for a dancer, a place of entertainment for the audience, a revenue source for the owner, a destination for a taxi driver or even a bad memory for one who was jilted one night after leaving.

Analysts of social action have been concerned with notions of place, and with the settings that convey cultural meaning and frame behavior. Goffman uses a theatrical metaphor, where "frontstage" and "backstage" distinguish different modes of behavior and action in interpersonal interaction. He points explicitly to "regions" as one of the elements that contribute to the framing of these different styles of action. However behavior can be framed as much by the presence of other individuals as by the location itself. In other words, the "place" is more than simply a point in space.

Giddens (1984) adopts the term "locales" to capture a similar sense of behavioral framing. Again, these are more than simply spaces; he observes, "it is usually possible to designate locales in terms of their physical properties . . . but it is an error to suppose that locales can be described in those terms alone." For Giddens, again, the critical feature of these settings is the way in which "features of settings are used, in a routine manner, to constitute the meaningful content of interaction." This strengthens the role of human action in how it is framed not only by spaces, but by the pattern of understandings, associations and expectations with which they are infused (Harrison, 1996).

The impact of this cultural and social framing may have deeply ingrained consequences. Knez (2005) showed a significant link proceeding from residential time to place attachment to place identity. This latter result indicates that prolonging one's stay at a place intensifies one's emotional bond to that place, which in turn means that a place becomes more a part of one's conceptual and extended selves (Neisser 1998). In this context a part of the impact of place on social action and of social action on place become deeply rooted in one's social consciousness and awareness.

4. Place Making and Social Constructions

Artificial social constructions, such as maps, globes, paintings and poems, have been used to share spatial and placial information for the history of humankind. While early traditions were more closely modeled on direct human perception and interpretation, that perspective has waned in recent decades, during which the focus has shifted to more objective 'pure' representations of space.

4.1 Maps and Mapmaking

Maps have always been a useful construct to help humans understand and make use of spatial information. Maps help us find where we are going and what to expect when we get there. Maps have served a variety of purposes throughout history. They have been tools to aid in navigation, works of art, symbols of power, and even methods of political control. Still, the fundamental purpose of a map has remained the organization and sharing of spatial information in a clear concrete representation. In some cases this has extended the use of the map to influence an individual's natural perspective and introduce bias (Vertesi 2005). Nevertheless the directive and informative qualities of maps, situated in a clear and presentable framework, influence decision making in the spatial world and spatial thought.

The rise of powerful tools of precision and computing in the last half century has brought new capabilities to a fairly traditional art. This has meant a new power to construct accurate and precise maps that can harness large repositories of spatial information. Even recently, new web mapping applications have offered the capability for the dynamic spatial information presentation to almost anyone.

4.2 The Geographic Information Systems Approach

Geographic Information Systems (GIS) are tools and technologies used to view and analyze information from within a geographic perspective. The primary focus of these applications is to link information to location and enable the visualization of large sets of spatial data. Typically, a GIS application presents images that have been captured by sensors, terrestrial cameras, and so on. It then supports the manipulation of these images by zooming, panning and layering additional sources of information (Lanter 1991). More sophisticated applications represent this as vector information to be rendered at run time. This allows the addition or removal of certain parts of the geographic content independently (showing and hiding roads, buildings, parks and so on).

The typical interaction in GIS applications is the query. This is where a user specifies a set of geographic information to serve as a base structure and then layers supplemental geographic information on top (Lanter 1991). For example, we might look at only the rivers in a geographic region and then layer information such as presence and type of trees and soil structure in order to predict riverbank erosion. This kind of approach is very powerful, especially when developed with modern design techniques.

There are a large number of benefits to the GIS approach. It focuses on displaying accurate information, which is of absolute necessity in certain kinds of applications (Miles 1999). The layer metaphor scales well and supports the view and manipulation of large amounts of information that may or may not be obviously related. In this respect, the GIS approach is very flexible. Many aspects of the world can be captured in GIS; Spaces full of discrete spatial objects, measures of the attributes and relations between these objects, or even continuous measurement of several different properties or themes within a concrete spatial region (Egenhofer 1995).

There are, however, fundamental limitations to the GIS approach, and many difficulties in implementing it successfully. The most serious of these still remains the serious distance between the system preconceptions and the user's understanding of the goals in interacting with geographic information (Aime 1999). This can often result in usability problems that are tied to failures in interpretation and gaps between user task conception and GIS query implementations (Prado 2000).

These problems are well addressed by Traynor and Williams (2005) in their survey of several GIS systems, while attempting to understand how the usability of these systems affected users. They chose a selection of common tasks, such as opening a map and analyzing multiple layers of spatial information. They concluded that the GIS applications had three distinct problems when used by non-specialists. They often rely on technical terminology, they require a strong mental model of the software architecture to be effective, and there is no strong attachment between the final compound representations of spatial information and how that information was generated (Traynor 1995).

This last point is the most troubling, as it means that while users of GIS applications are capable of creating rich displays of spatial information, they lack a solid understanding of how the information is being displayed and (consequently) how to use and manipulate it. In order to rectify this, Traynor and Williams propose a more task-centered design for GIS applications that helps ground the information more concretely to the user's needs and the fixed spatial perspective (the underlying map), and less to some arbitrary interface structured for data professionals (Traynor 1995).

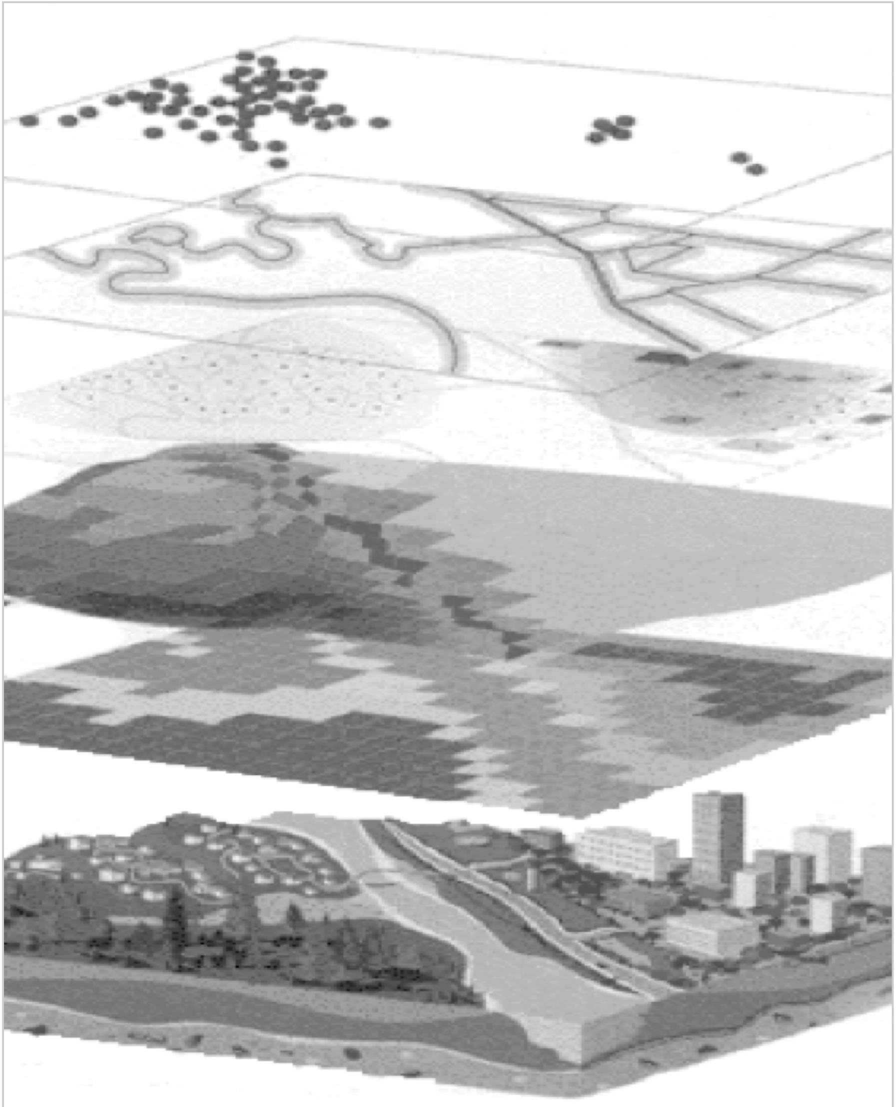

Figure 3. *Layer metaphor visualization of a typical GIS.*

One might compare this solution to the approach now taken by databases when designing information display for non-experts. In this

case, the GIS application is designed for individuals who possess expert skill at dealing with the manipulation and organizing such data and 'thin client' applications that need to be constructed to present this information to casual end users. This speaks to the fundamental limitation of GIS applications, as they are concerned about the precise output of large sets of data. While this makes them well suited to data professionals, they limit the end user population only to these professionals. The use of layers to categorize disparate sets of information speaks to the inability to establish deep meaningful relationships between this information and an inability to tie it to the geographic display in more than a very limited fashion (Traynor 1995). This is the kind of approach that makes casual users ask, "What is it about GIS software that makes it so hard to use, so hard to get the information out?" (Schuurman 2000).

4.3 The Rise of the Web Map

The Map 2.0 approach is the name that has been given to the relatively recent availability of web-based mapping applications that offer increasingly powerful APIs (Application Programming Interfaces) that enable outside developers to build their own maps (the so-called mashup). These maps often showcase widely disparate displays of spatial information in a powerful web-based geographic display. Google Maps and Virtual Earth represent good examples of these mapping applications. These kinds of maps are very similar to the approach employed by traditional GIS applications with a few key differences.

These applications dismiss the need for sorting through widely disparate information within a single application and instead offer a map based on the particular spatial information needs of the user. If you need to see a map with all the cabs in New York City, go to this address; if you are interested in a map with apartment listings from Craig's list, go to this address. In a sense, each layer in a GIS application becomes a new instance of a Map 2.0 application. These maps also incorporate the idea that spatial information sources can be inherently dynamic. GIS applications rely on decidedly more static reserves of information—large databases collected for specific purposes. (Foresman 1997)

A Map 2.0 application is perfectly content with scouring new sources of information from the web at run time. While earlier web-based maps were more clearly directive (with some limited informative capabilities) these maps embrace the idea of a map-based information display in an unprecedented way; anyone can display any kind of spatial information they desire. This is a powerful approach and within months after the first Map 2.0 applications launched, hundreds of different maps displaying all kinds of dynamic spatial information have become available (Google Maps Mania 2006).

In some respects, however, these maps are a step back. They forego the complex layer-based approach of GIS applications in favor of tailored unique displays. This necessarily limits their scalability. Programmers using these technologies must incorporate disparate spatial information on their own, with only the capability of displaying that information on these applications. In short, these maps offer a powerful front end for the display of spatial information, but not a mechanism for building relationships between that spatial information. They fail to support the kind of complex relationship between geographic information and supplemental spatial information that a developer might desire. One can add "spatial information pins" to a map, but cannot change how the underlying image is displayed based on differing spatial information.

4.4 Issues and Concerns in the Tradition

The question of how humans organize spatial information has already been explored to some degree. From this understanding, it appears that these maps are not how humans actually organize spatial information. These maps are a useful construct for us, because it represents a fixed view of spatial information that everyone can share (MacEachren 2004), despite the fact that individual mental maps of space tend not to look like this (Hayward 1995). While certain particular salient relationships may be preserved across multiple users, it is more likely that different kinds of people will create altogether different maps. Even in casual consideration, obvious differences arise. A person with a car, for example, may have a very different understanding of the city than a person who travels by foot and subway (Vertesi 2005). The second person will have a much more fragmented set of spatial relationships that are centered in proximity around subway stations.

There is also the idea that we capture relationships between spatial objects as they relate to their relevance in our own lives. One person may be very familiar with important landmarks between certain office buildings because he or she travels frequently between them. Another person may have no strong mapping of those landmarks because even though he or she travels that same space, there is no particular tie to any of the buildings there.

While GIS applications represent the traditional use of computers and mapping, the limitations of this and the new Map 2.0 approach, coupled with knowledge about the organization and use of spatial information in psychology and decision making, suggest new approaches that have yet to be properly explored.

4.5 Outside Perspectives in Urban Planning & Architecture

While this perspective comes with limitations, there are other approaches that address the problem from different angles. Organization of city spaces and building and planning a city (instead of mapping) suggests new insights into spatial representation.

Place, as we have described it here, is a central concern for architects and urban designers. For example, Whyte (1998) provides detailed descriptions of the life of the street in a modern city. His comprehensive descriptions of the use of the street-side plazas highlight the issues between places which “work” and those which do not; whether or not people want to be there. The approach to place for Whyte (1998), and for many in the field, becomes the practical concern of place construction. Designing a place or a collection of places is focused on human need and experience within spatiality.

Similarly, while Christopher Alexander’s “patterns” seem to describe principles of physical design, the focus is not on the structure of buildings and cities, but more on the living within them. He comments, “Those of us who are concerned with buildings tend to forget too easily that all the life and soul of a place, all of our experiences there, depend not simply on the physical environment, but on the pattern of events which we experience there” (Alexander 1977). This highlights a strong similarity with the notion of activity in place and its effect on mental representation from the philosophical traditions.

Architects and urban designers are concerned not simply with designing three-dimensional structures; they are concerned not with ‘spaces’, but with places and the human activity that occurs within them. (Harrison 1996)

A common conception expressed by Harrison and Dourish is the idea that place derives from a tension between connectedness and distinction. Connectedness is the degree to which a place fits in with its surroundings, strengthening the pattern of the surrounding environment (color, material or form are obvious, but also is the need and relationship between human action). Successful places respond to those patterns, even if they do not maintain the patterns completely. It is when these relationships are broken down that we say that something is 'out of place' (Harrison 1996).

A measure of this placeness also noted by Harrison and Dourish is "the degree to which a place reinforces—or even defines—the pattern of its context". These relationships are not static or flat, for a place must also be distinct from its context. The tension can be addressed by defining the distinctiveness of a place in terms of the surrounding context, and vice versa. This model of place, in the tension between connectedness and distinction as they relate to human experience, is indeed a valuable way to think about and design places in computational space as well as physical space (Harrison 1996).

While their approaches are interesting, most architects and urban planners are soundly focused on incorporating in the city technology and understanding, not on extracting it from the city. In his work City of Bits, and later publications, Mitchell (1996) imagines spaces that have become augmented by technology and information flow. This perspective, although interesting in its application, often neglects the obvious reversal—how can one get the rich perspective and already complex picture of place back out into the digital world?

4.6 Translation Vs. Deconstructions

These approaches lead one to wonder why these points aren't concretely represented in the map representations of space. The answer is complicated. While a sense of place may be somewhat subjective, factual elements of a spatial location are not. This is geospatial data like elevation, vegetation, population density, or street congestion.

Geographic information systems approaches, including the newer web maps, look at a place and reduce it to base spatial data. At its most detailed, this information is a deconstruction of the original place, neither as complete nor as flexible as the representation a human uses, but free from subjective impressions (MacEachren 2004).

The pragmatic difficulty is that this information must be re-represented to the user. Looking at a map, a user sees only a collection of data plotted on a coordinate space. Some basic relationships may be developed visually, such as land usage and elevation, forming the base of a perspective contour map, but interpretation and synthesis is left to the user. This produces layers of discrete information that can often be reduced to the observation of digital pins on a static image. Professionals, looking to discover meaning and draw conclusions on there own, are capable of doing this, but a typical end user trying to get a sense of the relaxing parts of a city has some work to do.

In attempting to develop new spatial representations, one focus will be on representations that translate existing human perspective rather than deconstructing it.

4.7 A Working Definition of Place

While the last two chapters have focused extensively on what place is and what place means, many of the theories both conflict and compliment one another. This is reasonable; given the significant advances in spatial theory and practice over time, the concept of place has waxed and waned in importance. While the importance of place and place construction has never been ignored, the difficulties in adequately capturing the meaning and practical importance of place have led to varying formal definitions of place. For the purposes of this thesis and the work conducted in the next several chapters, the working definition of place originates with discussions from Tuan (1977) and Lakoff & Johnson (1980). Here:

> *"Places are spatial locations that have been given meaning by human experience. The sense of a place is its support for experiences and the emotional responses associated with them."*

This defines place actively, as the result of human perception and subsequent activity. The personal and social sense of place is constructed by human encounters with spatial regions and their subsequent fitness for particular human activities.

5. Representations for Space & Place

Having suggested that places form the cornerstone of human spatial representations, and noticing the lack place-based representations in social constructions such as maps, one may begin to wonder what a place-based social representation would look like.

5.1 Motivating a Design Exploration

This question formed the motivation for the initial design exploration in place-based representations—representations that are more closely modeled on the cognitive maps that humans use to organize spatial information and placial knowledge. Consider, as an example, Figure 4. While intended to be a humorous cartoon, the representation is not such an unreasonable one. It captures, quite rightly, the fact that an individual has a significantly more detailed representation of the area with which they are familiar. It seems to have reasonable topological organization. Even with regard to wayfinding, it is not entirely useless. Although clearly distorted in distance, the basic directive qualities (across the Hudson, across the Pacific, far away from Ninth Avenue) are reasonable.

Why is this not a reasonable representation?

The early design exploration was also encouraged by the limitations and oversights of existing mapping trends in answer to this question. This initially led to the consideration of what maps should do, and what the problems are with current mapping solutions. Maps should be directive; they should tell one how to get somewhere. They must also be informative; they should tell one what is at a location or what that location is. Finally, maps should also be enjoyable and easy to use; they must meet certain aesthetic and usability requirements.

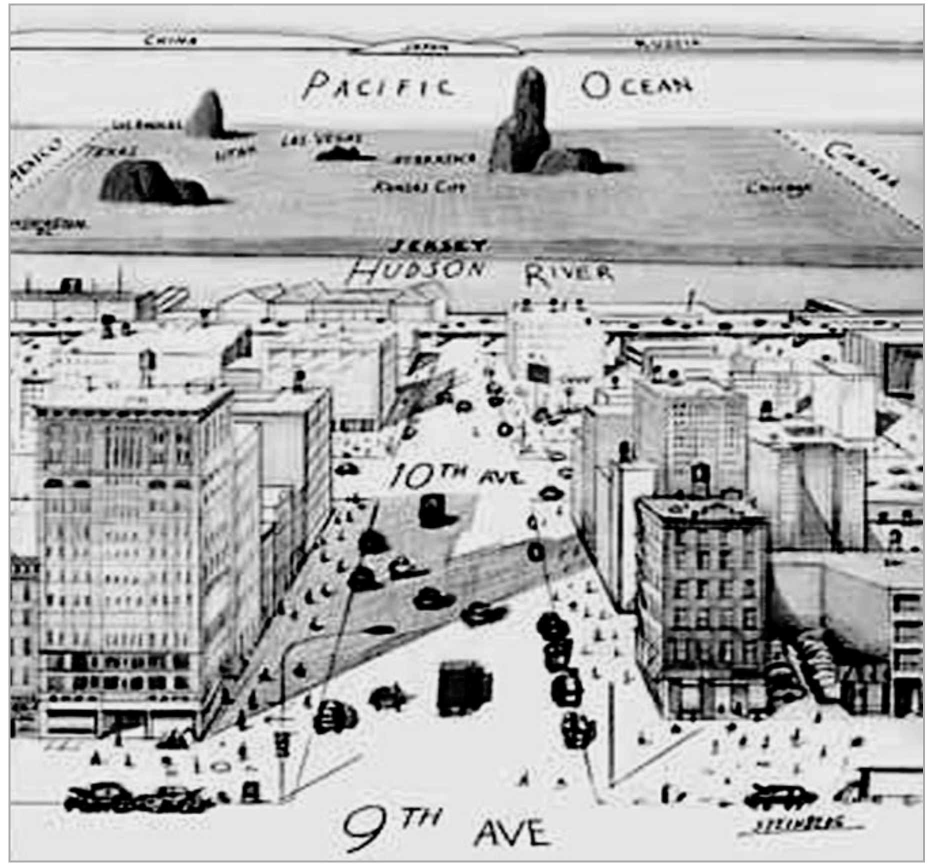

***Figure 4**. Saul Steinberg's "View of the World from Ninth Avenue" showcases a humorous, but realistic, interpretation of the mental model people adopt in visualizing large geographic regions.*

After reviewing modern maps, certain limitations and problems were discovered that failed to live up to these requirements. In general the focus was almost exclusively on the directive component of map making (and indeed, a very limited automobile-centric component of that). They were (arguably) not designed towards ease of use and aesthetic guidelines. They also had difficulties encoding informative qualities in anything but a superficially shallow manner. The closer the representation of this information mirrors a user's existing conceptual representation, which arguably must be similar to those utilized in cognitive map making and place making, the more successful and efficient these structures will be for user action and behavior (Goguen 1999).

This seemed to stem from a locative tradition of map-making within GIS. (Bleecker 2005). Here the focus is on displaying accurate information. This is a reasonable goal, but it is not necessarily the only goal of mapmaking. Overwhelmingly, the metaphor for displaying large amounts of information fell to a layer-based metaphor. This, coupled with a primary concern about geographic accuracy, led to the criticism of the systems to meet the needs of the average user (Traynor 1995; Schuurman 2000). While the locative approach does focus on using relevant information and concerns itself with information usage by the user, there were few widely employed mechanisms for determining relevance or for incorporating this information visually.

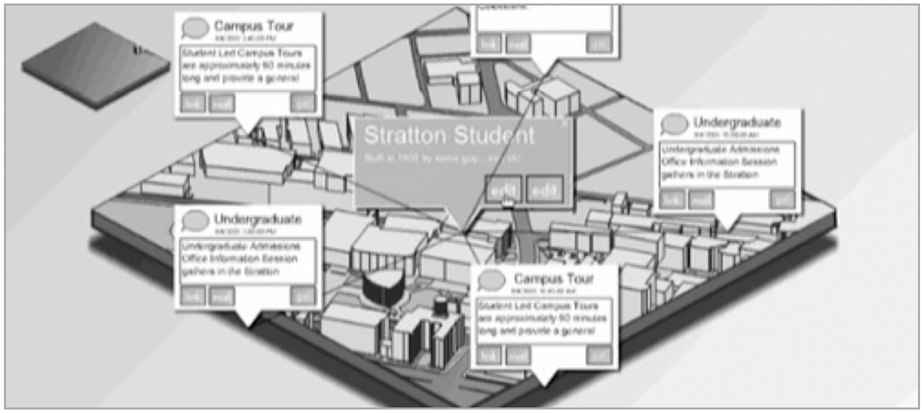

Figure 5. *Early PlaceMap visualizations of MIT's Campus highlighting an expanded view of ongoing events.*

5.2 Considering New Representation Methodologies

The limitations of these existing approaches led to the exploration, from within this work (and from others), into new approaches to mapping that would better serve the necessary aims. This includes:

Object-Oriented Mapping: This approach focuses on an alternative to the layer metaphor common within GIS. Here the aim is to incorporate holistically information into the representation. This treats spatial objects as service providers or option enablers and encodes information within the space, rather than tacking it on and grouping it within layers (Egenhofer 1992; Kidner 1994; Camara 1996; Lurie 2002; Worboys 2004). Discrete spatial objects are primary considerations. These include geographic features such as buildings, parks, and streets. Secondary features form a part of these discrete objects. This includes the events, people, activities and services within the spaces. The goal is not to associate this information with only a point in space, but to encode this information (and its impact) directly within the spatial object. This allows us to visualize the data in different ways and observe the interactions between the data and with the user.

User-Centered Mapping: User-centered mapping eschews the traditional 'data-oriented' focus of GIS and is focused on user-centric goals and methods. Here the user is at the center of the interaction. Implementations vary, but there is a trend of focusing on changes to the map layout or presentation based on the user model. This is not a particularly new idea, and the use of visualizations that capture this concept has been widespread (although not concrete) (Jordon 1988; Lanter 1991; Virrantaus 2001; Hockenberry 2006). This relies on limiting the social considerations of maps to some degree and instead focusing on constructing a spatial representation from a particular user's perspective. While this sacrifices universality, the user is able to getter a better sense of place local to his or her particular needs.

Active Context Mapping: Incorporates features of the above approaches in an attempt to visualize implicit spatial information directly (Chen 2000; Cheverst 2000a; Cheverst 2000b). In some sense this attempts to create artificial happenstance. Information is directly encoded and an interest structure can be built on top of the spatial representation to describe use interest and semantic relationships between spaces. Another goal is to create a more organic, living representation of space where new objects and interactions are created based on the intersection of information and interest. Given a building that provides a food area, for example, if you and a friend (with high interest) are both child nodes of the building around lunchtime a new 'Eat lunch with so-and-so' event is automatically created.

5.3 Design Goals

In addition to exploring these approaches, the principle goal is simply to consider representations like Steinberg's and the message they impart—'the user is a mediating force for spatial information.' Essentially what this is describing is an appreciation for the fact that one user's map need not (and often should not) look the same as another's. Each user has a different spatial orientation and places different demands on the information in his or her space. A person looking for amusement parks to visit has a very different map from someone looking for late-night pizza places. The first map may show very little of the surrounding urban area because it supports a very different kind of spatial expectation. This map may expect to encompass a larger spatial perspective because it expects the user to allow more time for the trip. The map of nearby pizza places, however, would be interested very specifically in spatial proximity to the user and would know which places were open and which delivered. These examples illustrate the need take several key things into consideration when constructing these new representations.

Build a strong model of the user: A strong model of the user is highly desirable. This model may be derived in a variety of ways, but it needs to be able to capture adequately user intent within the limitations of the application's focus. Ideally, this model can be adjusted over repeated use, although this is not always necessary or appropriate.

Establish relationships between the spatial information: A system that understands relationships between spatial information can more fully support novel decisions in the presentation of that information. An abstract data structure such as a graph could be ideal for modeling not only the pure spatial relationships between objects, but would also allow the system to project interest or task appropriateness onto the model (Shneiderman 1997).

Incorporate the user model with the spatial information: In the design exploration this is represented this by weighting the data structure (Dudek 1993) that describes spatial relationships by an outside measure of need based on the user's profile or task. This measure could be an interest function based on social and temporal demands, or some outside measure of task suitability. In a sense, it describes the effect of the spatial context (Zipf 2002).

Use a holistic representation: While the GIS approach treats spatial information as a static image with information to be layered upon, it can be more reasonable to consider the base geography as a collection of semi-discrete objects. Not only does this accurately mirror how people really consider such things, but also we will eventually show this to be a more convenient way to manipulate spatial relationships.

This initial design exploration resulted in three design phases. The first was an initial proof of concept designed simply to explore the integration of dynamic spatial information on a web-based map. The second iteration was a more complex and sophisticated iteration of this. This design predated the current generation of web maps (Google Maps, Yahoo! Maps, etc.) and the open APIs that make this relatively easy and available given current technology. The interactive flat map was an effective system comparable (in both timing and capabilities) to the newly emerging web-mapping solutions. The final iteration was a more complex system designed to explore this information within multiple representations. This includes both traditional spatial perspectives as well as graph-like models, linear representations, and altered or augmented spatial representations where additional information is allowed to warp and distort the spatial viewpoint.

The idea of using multiple representations of space is a natural one. People rarely rely on a fixed model, and tend to switch (alternate) between models often. This can be partly explained by changes in objective and circumstance, but also by the differences between perceptual and cognitive space. (Couclecis and Gale 1986)

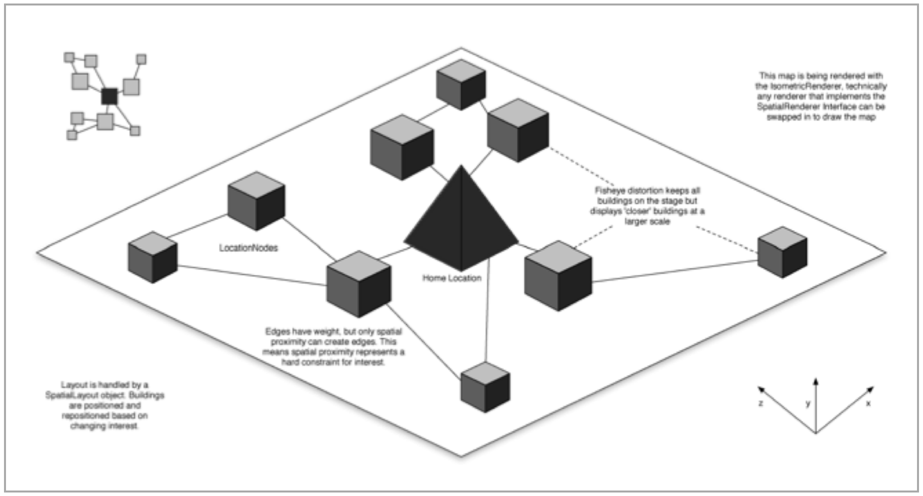

Figure 6*. Early concept sketch of semantic graph visualization and representation for PlaceMap design explorations.*

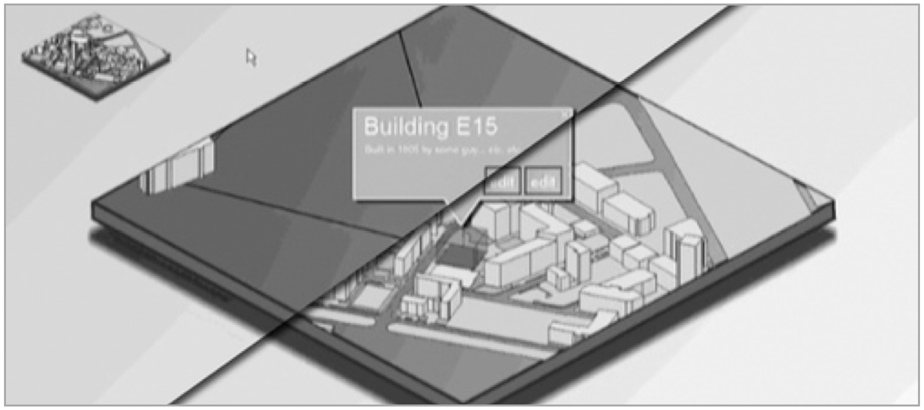

Figure 7*. Crosscut of early PlaceMap visualizations showing differences between traditional spatial view and the same area as a weighted semantic graph.*

5.4 Architecture for Representation Exploration

The exploration was built both to gather and to present spatial information. Some information is gathered on the server and stored locally on the server's MySQL database. Additional information was piped directly from the MIT Data Warehouse's Oracle database, which provides access to all available MIT information about a particular building, service, or user. Some of this information was stored, while other information (particularly sensitive user data) was kept only during the user's working session and then immediately expunged. This information was made available as a web service layer constructed in ASP.Net to provide a gateway to remote applications. While the principle application was the constructed design exploration, the possibility of outside and future applications with access to this information was an important consideration.

A lightweight client environment was written in Actionscript for the Flash platform (initially targeting Flash Player 7). All transactions occurred over with SOAP (Simple Object Access Protocol) over HTTPS (Hypertext Transfer Protocol with an additional authentication layer) because of the sensitive nature of some of the user information. The Flash platform was chosen because it offers near seamless web integration, with powerful processing capabilities. This was necessary because some of the visualizations can be complex and computationally demanding. At the same time it is a relatively ubiquitous mechanism for presenting rich multimedia content relatively seamlessly compared to other solutions available at the time.

The architecture was designed to incorporate the ideas present in the various approaches to mapping. Instead of mapping to arbitrary points in space, spatial regions were constructed as discrete spatial objects that housed events. Events included explicit activities, possibilities, and opportunities for user interaction. People were incorporated directly into the system based on a simple IP (Internet Protocol Address) lookup that, given the construction of MIT's Network, determined the person's current (or in the case of wireless access, closest) building. The spatial regions were designed to be flexible so that they could be organized into loose hierarchies and generally included objects such a buildings, parks, and squares—and even more specific determinations such as rooms or similar areas.

The system is divided into a variety of classes that govern different parts of the interaction. Data connectivity from the Flash Client is managed by a DataBroker class that is responsible for connecting to server-provided web services and maintaining changes in that information over time. This model worked particularly well, and it was one that would be followed in later implementations.

The other classes that make up the principle architecture are intended to be flexible and modular. This is necessary in order to support a variety of different viewpoints and representations. This requires having classes that can be flexibly swapped in order to provide these

Figure 8. *Early PlaceMap visualization showing weighted semantic content and connections within discrete spatial regions. Three buildings are linked and contain spatial information or services that have been sized by an interest calculation based on an existing user profile. The base geography is represented in terms of objects. Objects are buildings, roads, parks, parking lots, and so on. A spatial object is anything that could house spatial information or other objects. This allows different levels of granularity in the representation. A city may be an appropriate spatial object at some level, and a neighborhood at another. These are not purely placial objects at this point, but they are similar to places (Rumbaugh, 1991).*

viewpoint changes, while at the same time preserving an underlying consistency of structure and interaction.

For this exploration, this distinction was divided into an underlying data organization structure, a particular renderer for specific elements, and a layout manager to organize those elements. For the most part, the view generally utilized a SpatialGraph class to store general data. This includes information about spatial objects, their spatial data information such as latitude and longitude, height and shape. This information was organized in a graph-like way, simply because it was

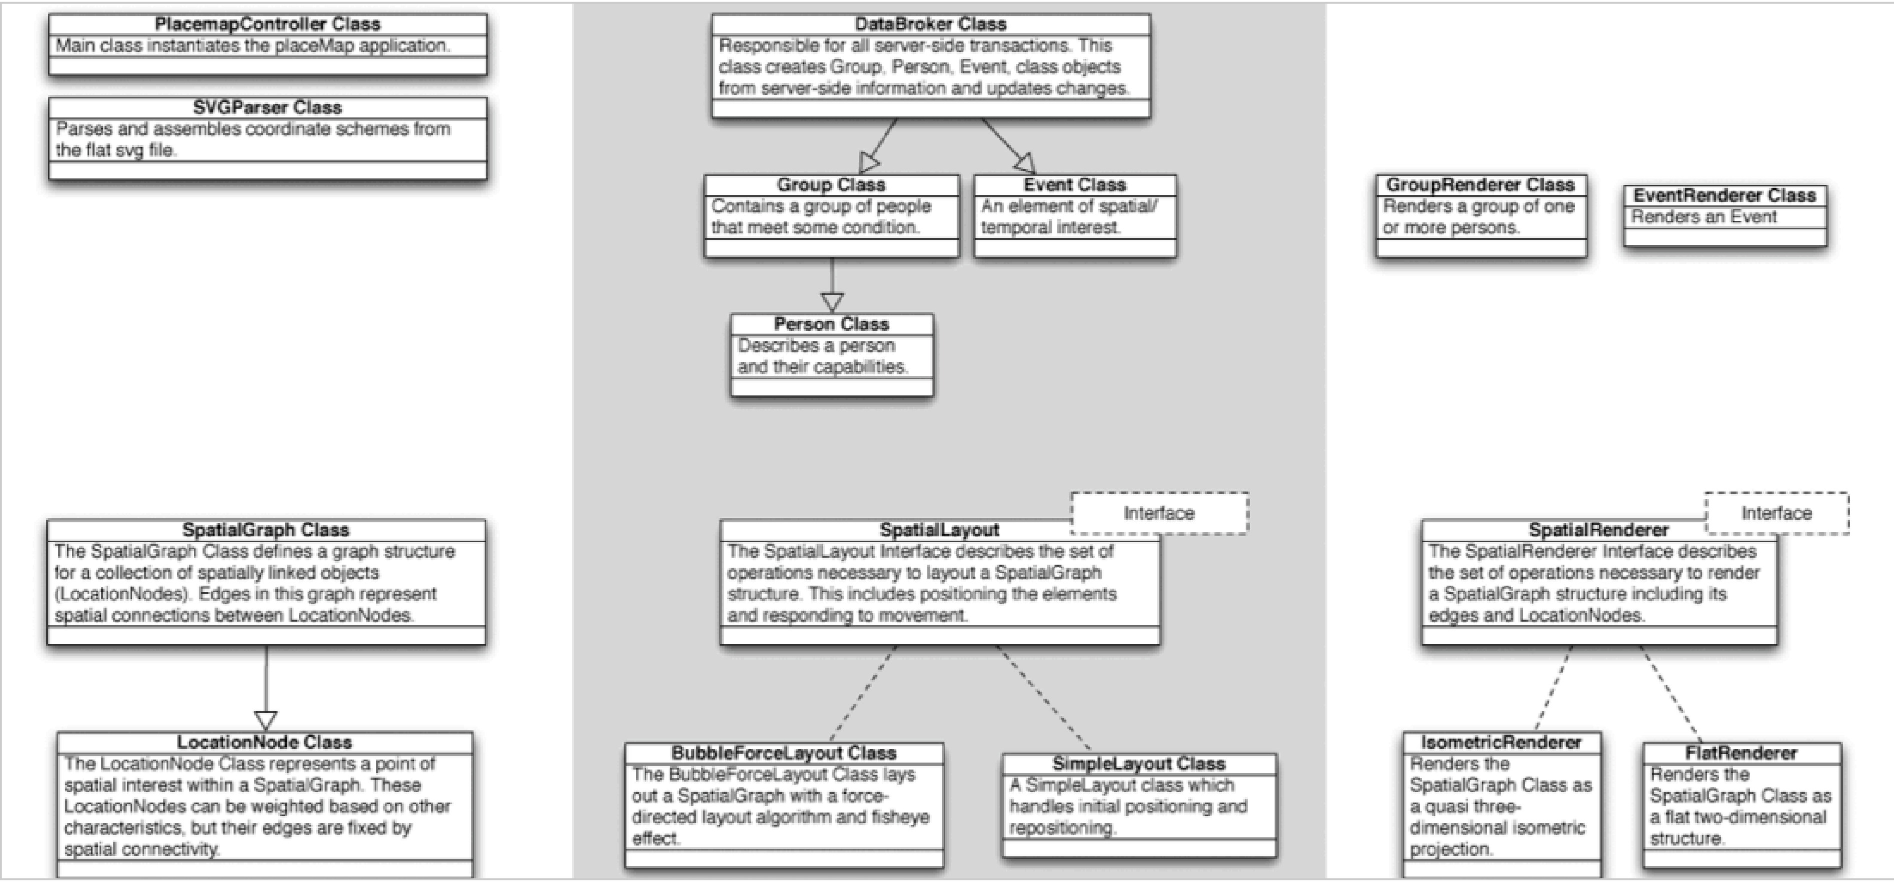

Figure 9. *Class Diagram of PlaceMap Exploration Framework, this primary framework impacts greatly on the later applications.*

computationally beneficial in the various views that were more graph like and also because it wasn't a significant detriment when the view was not graph like.

Although a general set of coordinates was used to describe position and shape (provided in an xml listing offered by the XMLCoordManager), it was sometimes necessary to adjust these coordinates to provide different perspectives (such as the difference between a traditional top-down perspective and an isometric). A coordinate renderer serves as an intermediary between the raw coordinates and other renderers in order to perform these transformations.

Other objects, such as people and events, are drawn by their own renderers. This allows changes to particular types of objects. While generally the rendering of these objects depends somewhat on the overall perspective and view, it is possible to arbitrarily mix and match.

Given a particular perspective coordinate space, the other primary adjustment to the view occurs in the LayoutManager. The layout manager determines how sub elements are positioned within the perspective. This can be a traditional layout based on spatial perspective (essentially the raw coordinate information) or it can be an arbitrary graph-like layout based on models of user interest, semantic similarity, and so on. Practical perspectives tended to fall in between; a traditional spatial perspective that has been modified or distorted to reflect other measures of user viewpoint.

5.5 Implementation Details & Discussion

Organizing these spatial objects as some sort of graph structure was a useful representation. Each node in the graph represents a spatial object and all of the spatial information contained therein. An object also allows arbitrary links to other graphs, and this allows expansion for an adjustment to granularity. The edges represent some degree of spatial connectivity and the method of connection. Each edge also has the capability to be weighted based on user task or interest. A building with great importance to a user's task would have the edge (or series of edges) connecting it to the user's current location weighted heavily. The spatial information within a node is used to help calculate this weight. Spatial objects are modeled as both containers and service providers (Rodden 2003). A building may contain a variety of events, people, or other pieces of spatial information. At the same time, it may offer certain services that may be of use in certain kinds of tasks. A pizza place allows you to eat and sit, but not to shower. These services also form the basis of interesting spatial relationships between objects. Buildings that will sell food, but offer no seating, may have a strong linking with nearby buildings that offer seating, but no food.

One approach was the use of distortion to mediate visually spatial information. When it comes to visual display, particularly computer-augmented visual display, there is some flexibility in altering presentation when supporting different kinds of users and different kinds of use. One might think this is not appropriate in the world of mapping. A subway map is a perfect example of a map limited by user needs. The user can travel only between discrete points. There is no need here to consider the geographic position of the destinations—the only concern is the relationships between the spatial information (Vertesi 2005). Many tourist maps also disregard geographic accuracy. They show popular destinations with great distortion and only major routes to connect them. These kinds of distortions are effective because of constraints. In subway maps we assume a very limited spatial mobility, while in tourist maps we assume a small finite model of interest.

These real-world examples are why some have begun to consider the use of distortion in maps, and have provided further motivation for this aspect of the exploration. The fact is that distortion-oriented presentation techniques work well when dealing with large amounts of visual information. Instead of suffering from information overload the user is allowed to focus on what visual information is relevant (Churcher 1995). Because spatial information is often a large data set, previous researchers have been interested in using distortion-oriented techniques to present GIS information. Much of this early work was inspired by Furnas' paper on distortion-oriented presentation techniques (Furnas 1986). Rather than faithfully reproduce a map, one could use a distortion-oriented system to emphasize the connection between parts of the map.

Churcher points out the value of the degree of interest function to determine the amount of distortion. This concept relies on some sort of base measure of interest for a node based on contributions depending on its distance from the focus (Churcher 1995). When presenting these exploratory representations, these techniques are used to draw attention to spatial objects that are more likely to be interesting or appropriate to the user's task. Interesting buildings will be larger or closer and uninteresting buildings will be smaller or farther away. Using the underlying graph structure, increased complexity can be integrated successfully. Certain buildings may become more interesting because spatially adjacent buildings are. For example, a nearby sporting event may increase the relational interest of nearby restaurants. Adding temporal considerations to this model results in spaces that only become more interesting near the beginning or end of the events. These techniques were explored in a variety of ways and strong efforts were made to follow goals and suggestions based on previous work and research (Leung 1994; Janecek 2002):

Meaningful and transparent distortion: Because these representations allow distortion that is controlled by the system, its meaning should be made clear to the user. If the user cannot look inside the spatial objects easily, he or she may have no idea why the objects are being distorted.

Anchor the distortion to the real world: While distortion is a useful technique for highlighting semantic relationships, the spatial objects exist in a fixed geography. It is useful to some degree to anchor the positions of spatial objects to their traditional geographic positions, which means that while other factors can pull and distort spatial objects, the system attempts to retain some semblance of the traditional map structure at least in part. This is particularly necessary if users are already familiar with the traditional geography (Bouquet 1999).

Support for multiple foci and clustering: Users may have more than one area of interest in a map. This can be preserved by using multiple foci. The ability to cluster spatial objects should also be taken into consideration. This allows the system to further preserve strong spatial relationships such as physical neighborhoods (Sarker 1994; Janecek 2003).

5.6 Evaluation of the Design Visualizations

The success of this design visualization work and the resulting representations built on this methodology were primarily evaluated by their effect on performance times and behaviors with regard to information search. Egenhofer and Mark (1995) point out that information search in these situations can act as a sort of “litmus test,” helping to understanding the cognitive framework through which the representation is perceived. The goal of the evaluation was to observe changes in the effectiveness and behavior of information search in a spatial world.

It was also clear that the constructed system focuses almost exclusively on information search—the design goals and work are not focused on directive information. To that end a clear understanding of the kinds of information search that were worth consideration were built into the evaluation. The results were significant observations in the effects of these spatial representations on spatial information search. The model of information search represented in the system targets several notions of search behavior, much of which is addressed generally by Morville (2002) and Maurer (2006).

Known Item Recollection:

Known Item Recollection is the most simplistic information-seeking behavior. In a known-item task, the user knows what he or she is looking for, how to describe it, and generally has a good impression of where to look for the item. In many circumstances the user will be happy be happy with the first item he or she finds that matches some of the search criteria.

Search is the traditional solution to this kind of task. When a user is able to articulate what he or she needs, he or she is able to type it into a search box. So long as the results indicate the terms in context and the result description is clear, the user is likely to recognize an appropriate answer to the search. In the real world this may be more difficult, as most humans lack the comprehensive index that search engines have, but the concept is the same. If you ask someone for the item and they have the appropriate knowledge, they will return possible results, and if they do know the answer this will quickly become obvious.

Indices are another good solution, as are quick links, if there is sufficient contextual information to suggest what may be a reasonable answer at a given time. This can be particularly useful with a strong user model or well-developed task model. Basic navigation can also support this feature, but only if the navigation matches the user’s preconceived organization of where he or she would go to get this information. The primary goal in matching this kind of information is efficiency.

Exploratory Search:

In an exploratory information search task, the user generally has some idea of what he or she needs to know or where he or she wants to go. However, the user is usually unable to correctly articulate this, either not knowing how to articulate it, or not having the right words to use. Many kinds of spatial information search followed this behavior. Here the user also probably doesn't know where to look for the information. Despite this, the user will generally recognize when he or she has found what he or she is looking for, but may be unaware of the scope or extent of the required information.

In these tasks, search is generally not an acceptable solution. It can be useful, but because the user is unable to successfully articulate exactly what is being looked for, he or she may have difficulty getting the right answers. Some initial searching can, however, be useful for investigating the domain and learning about how to articulate what one is looking for.

A more appropriate avenue for search is simple exploration. Navigation, when well designed, can allow some exploratory investigation that prompts discovery and learning. When this is successful the user can learn the necessary information to articulate his or her main goal. Providing related information during this exploratory procedure can also be useful. This is particularly helpful when the user has chosen one option (incorrectly) that offered clues similar or related to the ultimate goal path for information search. The distorted representations seemed to have clear impact on more successful exploratory searches.

Undirected Search:
Here the user doesn't know he or she needs to know. The key concept behind this mode is that people often don't know exactly what they need to know. They may think they need one thing but need another; or, they may be looking at an area without a specific goal in mind.

Re-finding:
This mode is relatively straightforward—people looking for things they have already seen. They may remember exactly where it is, remember where it was, or have some intuitions about where it was. Here a successful representation needs to be very transparent, or have a very clear understanding of the user.

Design solutions can be active (where the user takes explicit action to remember an item) or passive (where the user takes no action but items are remembered). Active solutions work well but require a conscious effort from the user, who needs to know that he or she will want to return to an item in the future.

A good passive solution allows users to see items they have seen before, order them by frequency of use, easily get to the content, and the information within it persists over time. In the system this is represented by adjustments to presentation based on past use and the presentation of available user history.

5.6.1 Evaluation Results
Generally, the evaluation results suggest some large improvements in certain areas of spatial information search, while there are some limitations in others. The underlying cognitive implications of the results show that when presented with a map distorted by semantic relationships, when compared to traditional maps, information search behavior changes, and in certain key areas improves significantly with regard to search time.

One hundred (100) users were presented with either a map of a fixed geography and one incorporating the interest, or a task-based distortion map. The study was a within subjects comparison with a randomized ordering of condition. The users were asked to make several claims and complete certain tasks in each condition. These included:

- Spatial tasks that involved finding a specific place (such as finding the student center)
- Meeting a specific goal (such as locating reference material on biology)
- Solving an unspecific or undirected goal (such as finding something fun to do).
- A general report of user satisfaction.
- A report on the perceived social awareness of the display.

For the report of social awareness, a score of 100 represents a sense of complete social awareness and a score of 0 represents a lack of any social awareness, where social awareness is described as "an awareness about the social situation of other people, i.e. what they are doing, whether they are engaged in a conversation and can be disturbed, and of who is around and what is up." (Prasolova-Forland 2002; Tollmar 1996)

The user's reported social awareness was much greater in the "user-centered" map that was distorted by task, interest, or other semantic considerations. It was also noted that while it was more difficult for users to find a specific place, it was much easier for them to complete goals—often in unique ways (instead of simply looking for the library they might look for professors of biology, for example). Users were also to perform undirected behavior (finding something fun to do) more quickly and found that, in general, they enjoyed the task more in the distorted representation. The average reported social awareness was 63.4 for the tradition map. These are users reporting that they felt "aware of their surrounding social climate." The reported average was 84.8 with a map distorted by user interest and incorporating people and events in the spatial locations. Reported satisfaction was similar, with an average 58.1 (out of 100) satisfied with the traditional map and 79.3 satisfied with the distorted map.

Times for the task are particularly enlightening (D is the distorted representation and S is the static one):

- ➔ Finding a specific location: (D - 17.45s, S - 5.77s).
- ➔ Meeting a specific goal: (D - 8.76s, S - 24.33s).
- ➔ Meeting an undirected goal: (D - 9.98s, S - 36.87s).

This suggests that information seeking is altered and augmented by representations that present more contextually grounded representations that are adjusted to user needs and behavior.

5.7 Challenges, Solutions and Limitations

There were a number of challenges encountered in this initial design exploration. Some of these could be addressed, but others highlight remaining limitations in the approach.

With regard to pure computational power, it was difficult to get the kind of performance necessary to perform arbitrary adjustments to the layout and rendering of the large number of objects, even just with the amount of spatial objects in the MIT Campus.

For example, the system had initially attempted to use a force-directed graph layout that created a natural adjustment when switching views. This was desirable from a user and development perspective, as the distortion could occur and be visually mapped to the change from the original perspective to the new one simply by changing the forces of repulsion and attraction for particular nodes. This was implemented successfully, but even with only geographic elements from the MIT campus it was clear that the amount of computational power to scale this to a reasonable number of nodes was lacking. Eventually the system implemented a visualization that used a basic geometric layout with tweening to supplement the animation.

Data loading also proved to be difficult. The information necessary to capture MIT's campus represents about 10,000 different coordinate points and pathing information. This implementation occurred before the release of new web maps, which prerender bitmap sections at various zoom levels to deal with this problem. This option was available, but given the kind of transformations, distortions, and alterations that were attempted (beyond mere zoom adjustment), this was not a reasonable solution. This resulted in the implementation of some basic path smoothing, face trimming, and other visual adjustment algorithms to decrease initial load time. This included folding edges and detecting (and removing) occluded edges (with help from Markosian 1997). The newly available feature from Flash Player 8 that allowed caching vector objects as bitmaps to perform basic manipulations (movement, scaling, etc.) without rerending the objects was also used. The data itself was loaded with highly refined XPath (XML Path Language) queries and staged to eliminate initial performance bottlenecks.

In allowing the user to access rich sets of data about themselves there was the risk of possible exposure of sensitive data to other users because of illegal access, data connection errors, and simple user error. In order to eliminate the risk associated with this, MIT personal certificate authentication was employed. This was challenging simply because Flash seemed ill equipped to handle this kind of authentication and security scheme. The working solution was to utilize a Perl stub to manage authentication and retain the security settings in the Flash application.

These specific challenges were eventually overcome, but there were a number of challenges that met with limited or unsatisfactory solutions. Basic terrain was presented (below spatial object level) in the traditional spatial representation. Any distortion, however, was too computationally difficult to apply to the terrain. Terrain deformation is important, because it provides additional context and grounding, something that was lacking somewhat in additional viewpoints. While the local distortion that occurred (animated) within the user's perspective was relatively easy to follow, the greater implications of these distortions were not as clear in the user's mind. This was partly because the visual adjustments were not rich enough to provide this capability, but also because the system lacked a sufficiently holistic representation of each node to provide a rational justification for the distortion. In this design exploration there was a reliance on set, structured pieces of information that were ranked and coded by the database or set by user behavior. While this could eventually produce representations that adjusted meaningfully to users' understanding of place, it became clear that this required significant time and user investment. This was investment that was not motivated during the early uses of the system.

There is also the concern that in some cases the user can lose orientation because the semantic representation employed by the distortion conflicts with the spatial representation with which the user is familiar. This is the danger in having a system that overloads the spatial dimension with additional information; it is possible to distract and dilute analysis if the new representation fails to match the user's internal one. While flat spatial representations may not be the best representation, they are heavily ingrained by years of use and exposure and it is possible that it is easier in some circumstances for the user to adjust to this perspective rather than a semantic conceptual perspective that differs from his or her own. Adding more user control over the input to this representation and ensuring its generalization become key concerns.

At the heart of this is the fact that the system simply lacked a sufficient semantic understanding of spaces to provide holistic representation of elements and a sufficiently clear adjustment to contextual representations. The system could organize spaces, but it did not understand places. It is one thing to say to a user that a graph-like representation is organized semantically and weighted to their interests, but quite another to actually be able to understand the semantics of place for each location and the relationship between how a user truly feels. It is this driving limitation that prompted the construction of a general-purpose research tool and structure for gathering and interpreting place and human accounts of place in order to support semantic understanding in these representations.

.

6. Related Work and Alternative Approaches

This problem of how to capture a semantic sense of place has been touched on and addressed in other work. Some of these works only indirectly touch on this need, while others offer approaches that limit the necessity of sense of place in representation. In some regard this has already been touched on in the previous chapters, particularly the discussion on GIS and related technologies. These topics are reviewed here with the focus placed on their answer to this problem.

6.1 Building Blocks

There are a number of approaches and technologies that, in and of themselves, don't offer complete solutions to the problem of capturing place. This does not imply, however, that they are without consideration. Many of these approaches offer capabilities that can be built on in order to design systems that are able to capture and interpret place. These 'building blocks' form some of the necessary ingredients for answering the problem, without offering the complete recipe themselves.

Location Awareness:

Location awareness usually refers to approaches that understand where a user is, either through network monitoring, special hardware such as GPS, or combinations of these approaches with user input. Generally speaking, this can relate to notions of ubiquitous computing or augmented reality, but the scale that we are concerned about falls away from these approaches. The precision of these techniques is rapidly increasing, with GPS systems already capable of identifying coordinates to accuracies of less than a few meters, given optimal conditions. Technologies such as wireless triangulation and wireless positioning are rapidly becoming able to approach these levels of precision without the need for external sensors, instead relying on common hardware such as wireless access cards and other networking features and technologies.

Exemplar: Skyhook Wireless offers a service called Loki (Skyhook 2006) that exists as a plugin for the Firefox web browser. This relies on access to wireless access information. Comparing signal strengths and system conditions with observed database trends of user behavior can be very precise, and offer more focused precision than GPS under suboptimal conditions. This technology is available to any device with a supported wireless card and access to the Firefox browser. This includes desktops, laptops, and similar personal devices. Unique features are available, such as location-aware content streams and hooks into web maps to provide location-aware directions and information search.

Technology like location awareness doesn't offer a solution to understanding place in and of itself. It does, however, provide some excellent tools for helping to link information to particular spatial regions. The presented work assumes that the logical extension of the work done in location awareness is that, eventually (and possibly very soon) it will be able to identify precisely almost any location. However, the necessary granularity for most tasks comes down to place—not to a number of meters. Still, this can (and for the system presented in the next chapter, does) allow the linking of descriptive pieces of information with spatial regions and further the identification of specific places. Of particular note are solutions that rely on wireless technology, because they can very seamlessly integrate user position with relevant places and events. A similar approach forms one of the key components of the PlaceMap system and its utilization in CampusMap allows us to generate implicit accounts of place based on this information.

Web Mapping APIs:
As mentioned in the chapter on social constructions, web mapping APIs (or Map 2.0) are the direct decedents of GIS style approaches to spatial representation. There are a number of key differences, however, which separate them from GPS to some degree and make them attractive as possible building blocks for applications that can gather and interpret placial information. The main areas of interest are the lightness of the web maps when compared to traditional GIS and the ease with which varied and diverse information sources can be incorporated and realized, the result of which is the so-called mashup.

Exemplar: Google Maps are perhaps the best known of the web mapping APIs and offer a very diverse set of features. Google Maps can be deployed on any web site (given a Google-approved API key) and can incorporate information from any source. This includes event listing sites, personal ads, online databases and essentially any piece of locative information that can be computationally identified. Users are also able to interact with these pieces of information with capabilities provided by developers. Additional functionality, such as seamless navigation, spatial interaction, and drawing capabilities are also provided (Google 2006).

Web maps such as those offered by Google provide a rich foundation for the display of spatial information. The primary focus of this is the presentation of such information. These web maps don't provide the capabilities for aggregating outside data or interpreting it. While visual trends in spatial data may become apparent, this does not actually translate into the system being aware of those trends or preserving this information for subsequent use. These capabilities, if they are desired, must be recreated by each developer. This results in a system that is useful with regard to presentation, but does little of the legwork necessary to deeply understand and make use of spatial information or to draw inferences about place related to that information. As previously stated, a system functionally similar to this was developed early in the course of the work and offered similarly rich presentation, while also lacking in deep understanding of what data was being presented. The utility of these systems is in their successful ability to serve as a visual representation that meaning can be built upon.

6.2 Alternative Solutions

There are some real solutions to attempting to capture semantic sense of place in and for use with spatial applications. Some of these solutions are interesting, but for a variety of reasons fail to offer as general an approach as one might like. This is not to say that these approaches are without merit, and there are certainly circumstances where their use may be preferable. In general, however, they don't offer the same practical application that one might desire and which is incorporated into the PlaceMap approach.

The Geosemantic Web:
The Geosemantic Web is an attempt to incorporate geographic and spatial information in a semantically meaningful markup for the web. This is related to the general conceptions of the semantic web. Specifically, meaningful semantic geodata and metadata are structured into web documents with the intent that they are human readable, but also with direction for them to be machine readable.

Exemplar: The Open Guide network (Open Guide 2006) is a geosemantically structured set of city guides. Unlike similar listings provided by corporate sites, these encode rich semantic markup in the form of RDF or XML. The Open Guides represent a project within this approach that serves a practical purpose (city information), is of a significant size (covering over ten major cities—mostly in the United Kingdom—by contribution of altruistic individuals), and is well-structured practical semantic markup with direct human representation and machine instruction.

One could imagine a world where all of the information related to spaces and places were carefully associated with correct geosemantic meaning. It would, however, be a much more perfect world than today. While the goals of projects like Open Guides are admirable, and useful (the PlaceMap system is able to consider them as sources of data) they are not, and likely will not be, able to capture all of the important places. Relying on well-structured markup can be efficient from a parsing perspective, but it is difficult from the content creation perspective. When this support is available, systems like Open Guides can offer incredible comprehension and transformation capabilities for a large set of spatial data. When they cannot, they encounter missing pieces of space, seemingly devoid of understanding and limited by the insights and interests of the user base. This is one of the reasons why PlaceMap offers a more general natural language parser that doesn't rely on such explicitly structured documents. They are useful, but they represent a small portion of the available resources.

Smart Geographic Information Systems:
The GIS approach is covered extensively in Chapter 4, but it is worth reviewing here. GIS is focused on concrete data collection, with an emphasis on objective spatial data. This usually involves methods of data acquisition involving human agents with specialized devices, but these are giving way to mobile data acquisition and satellite photo analysis.

Exemplar: Environmental Systems Research Institute, Inc., commonly known as ESRI (2006), has emerged as the premier GIS solution in the commercial sector. Their solutions, such as ArcGIS, offer support for numerous kinds of data sources, manipulation capabilities, and advanced queries. This allows expert users to make significant research efforts into geographical problems. Recently trends in web mapping have resulted in sharing capabilities that offer interactions similar to those found in Google Maps. This allows a full cycle of data collection, interpretation, analysis, and sharing. There have also been recent trends towards ‘smarter’ GIS systems that offer models of behavior that have preserved some existing human interpretations of geography. ArcGIS has begun to embrace these, but support remains very limited.

Good GIS systems such as ArcGIS are good for a particular kind of user, the expert user. They offer the support for spatial data tied to coordinates and relationships between these elements. However, no matter how much data a traditional GIS system collects, there is little emphasis on its importance in specific human tasks and no emphasis on the perceptions of the average user. In general, the system does not attempt to understand the data itself. Rather its job is to provide the set of tools necessary for an expert human agent to make these observations themselves by adjusting the data and conducting analysis of the results. This is often necessary for these expert tasks, but sacrifices the general needs of an average user. Instead of being concerned about this, the approach outlined in the next chapter focuses on only the data that average humans actually encode through communication. While this might not be as complete as what can be gathered, it can be more meaningful in general situations.

Artificial Space:
Ironically, perhaps some of the most interesting work in understanding place comes from research into artificial space. In the realm of computer-supported cooperative work and complex data visualizations, spatial metaphors have been useful for the communication and presentation of large amounts of data. To that end, significant effort has gone towards understanding the role of place construction with an eye towards practical investment of placial knowledge.

Exemplar: The work done by Dourish and Harrison (Harrison 1996) is significant, as is the work in the Data Mountain project (Robertson 1998). Here spatial memory is utilized for organizing documents and there is clear observable place construction in resultant user behavior. These insights offer predictive power for the developers of such systems. Place construction is a key component in the virtual world, as well as the physical, and designing with this understanding creates systems that are able to support larger amounts of data, increased efficiency, and support of communication.

It is interesting to read the literature from artificial spaces, computer-supported cooperative work, and data visualization and information. The considerations about place, and how place construction influences behavior, are evident in most of these works. However, much like work in urban planning, the focus is on who this will be designed for, not how to identify this and make use of it within the system. There is not an active role in the system for place identification and subsequent utilization of this information. These would be systems that actively capture placial determinations with the goal of reincorporating them into the system. This is something PlaceMap supports, and is a natural, although unexplored, extension of virtual spaces.

Spatiality for Robots and Rats:

Significant work has been done with regard to spatial understanding in systems less vocal (and presumably less intelligent) than humans. From robots seeking to navigate unfamiliar environments with limited sensors, to rats moving through mazes, the history of these efforts is rich. The focus here is usually on small-scale space and (almost exclusively) on navigation. There is a strong focus in studying information search that is relatively simplistic (such as pure retrieval for rats in a maze) or where it can be clearly encoded (for robots). Information search in these situations becomes focused primarily on identification of spatial position and simple recollection.

Exemplar: Projects such as those proposed by Werner (1997) include navigating wheelchairs and robot office navigation. These devices employ interesting algorithms for the identification of features (corners, obstacles, etc.) and serve as useful aids in navigation and identification of basic spatial features that form the core of visualizing small-scale spaces such as rooms or even buildings.

This kind of spatial work is interesting, and deserves consideration simply because of the significant amount of time and effort that has been invested in it. However, as chapter eight points out, the differences between small-scale space and larger geographic space are poorly understood and may be more profound than originally offered. This suggests that the work into spatial navigation for robots and rats and the supplemental identification of features, recollection, and aggregation offer few insights into place identification. These systems are limited to the identification of the most basic features, not large-scale conceptions of place. The reality is that models for small-scale spatial navigation and even manipulation are different enough from large-scale space that, while there are lessons to learn from this approach, it fails to suggest a real methodology for place identification and understanding.

Common Sense Collection:
While not intuitively obvious, common-sense knowledge systems provide insight into the kind of approach that will be followed in the PlaceMap system. These systems attempt to capture common-sense facts about the world, similarly to how one might capture common-sense understandings of place.

Exemplar: Open Mind Common Sense (Singh 2002) is a system that depends on web-based entry of structured common-sense statements. These can be statements like “it is cloudy when it rains.” While these statements are not always true, they often are (or are often perceived casually by humans to be).

While systems like Open Mind offer an interesting approach, they rely on altruistic data entry. They also tend to be less specific than the accounts of place systems like PlaceMap are attempting to identify (they are usually more general, with specific persons or places rarely identified). Some systems tend to be significantly more structured as well, relying on data input from knowledge engineering rather than casual use. While the PlaceMap system follows Open Mind’s tradition to some degree, the primary focus is not on a special ‘place knowledge data entry’ but on a more flexible approach that can be embedded in general spatial applications. Here the focus becomes implicit inference, and not data entry and collection.

6.3 Moving to the PlaceMap Approach

Ultimately these approaches are all useful, and parts of them are contained within the PlaceMap system. The particular approaches, however, all generally fail to capture the expressed goals of the PlaceMap system: The ability to create a crisp set of interactions that supports the construction of placial knowledge and identification of place within general purpose spatial applications.

7. PlaceMap: System for Place Identification

7.1 Overview of the PlaceMap Architecture

The PlaceMap system was designed as an effective but flexible method for acquiring meaningful semantic spatial knowledge and to interpret that knowledge in meaningful ways. This system is primarily divided into two distinct components, The PlaceMap interaction architecture and component system and the PlaceSense semantic parser.

The PlaceMap interaction architecture is designed to be a flexible collection of components suited to the acquisition, annotation, and collection of meaningful supplemental spatial knowledge. For the most part this is the kind of human-mediated information previously described, records of human conversations or semantic adjustments to such conversations (such as tagging, annotating, and specification). The primary installed set of these components is found in the MIT CampusMap system, which was constructed to perform these activities as they relate to the distribution of meaningful spatial knowledge at MIT's campus.

The PlaceSense semantic parser is a backend system designed to organize and structure the gathered spatial knowledge in a way that can be used for later information visualization or incorporation back into the PlaceMap system. It is useful to think of the structure create by PlaceSense as our modeled 'sense of place' with regard to the source and type of the inputted spatial knowledge. To facilitate initial spatial information construction, the parser has a limited spidering capability for gathering web-based spatial knowledge.

The design of both of these components is directed toward research interests rather than practical ones (although there is some significant overlap)

7.2 PlaceMap Architecture in MIT CampusMap.

This section describes the specific instantiation of the PlaceMap interaction architecture for place-capture in a specific application. That application is MIT CampusMap, ostensibly intended to be a lightweight social information-sharing and event system for MIT.

7.2.1 Goal and Intentions

The PlaceMap interaction architecture is designed to be instantiated in a particular working application, MIT CampusMap. CampusMap is intended, from a user perspective, to be an interactive window into the events, individuals and locations within MIT. It is also intended to foster the social goals of successful communication and sharing about these features. From a research perspective, CampusMap is intended to gather communicative experiential accounts of place, ranging from low-level communication situated in spatial locations to structured annotation of place and place-based features.

The system is intended to be modular, with a well-designed user interface to promote rapid use and allow for significant expansion in the future. It is also necessary that it allow effortless data collection of user behavior for future experimentation and research. Ideally the design is general enough that CampusMap can be used by other universities and similar organizations (such as corporate campuses).

The overarching design methodology is such that the system appears tailored to specific user demands and capabilities, while at the same time supporting the investigation of user accounts of place. For the most part the goals are complimentary, and when they are not user goals are favored, so long as they don't interfere with the research aims, and the possibility of philanthropic behavior is encouraged. The design is user centric and user focused, giving the user control over data and its presentation through filtering, customization, and annotation.

In short, the several key goals of constructing this system are:

➔ From a **research perspective**, to present a system that is capable of using crisp interaction techniques and technology to capture semantic meanings of place dynamically in order to provide support for richer contextual decision making and representation.

➔ From a **developer perspective**, provide a toolkit, set of design methodologies, patterns, and interaction techniques that support the capture of semantic meaning of place in spatial applications that offers practical benefit to end-user behavior.

➔ From a **user perspective**, to present a compelling social application that fosters communication about places as well as the events, people, and features that define those places in practical human usage.

These goals ensure a system that computationally embeds notions of place transparently in spatial applications and offers practical and tangible benefit to existing and future applications.

7.2.2 Design Overview

The principle user interaction layer is built on top of the Yahoo! Map API. The API is similar to those developed by other providers such as Microsoft and Google. Technically, the interaction of the API is very similar to the system developed early in the design exploration for PlaceMap. The reason the commercial system is used (and not our own) is simply because we are able to offload a large amount of data storage and bandwidth, and it works almost identically. The Yahoo! Map API is particularly advantageous because it (like the earlier work) is based on the Flash platform, allowing the incorporation of rich animation, interactive media, and effects, with minimal programming.

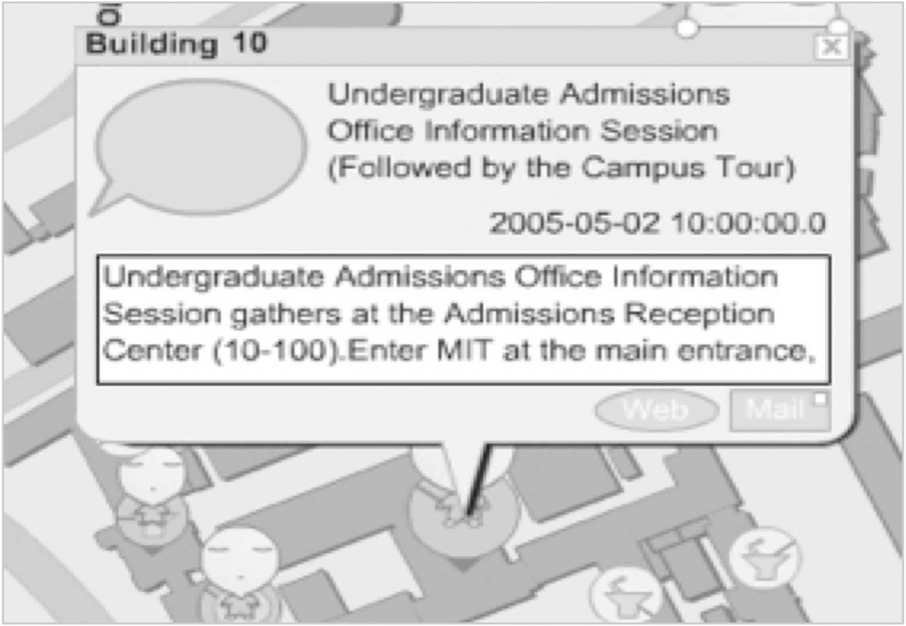

***Figure 10**. This early iteration was functionally equivalent to the new generation of web maps prior to their release.*

The provided API allows the presentation of traditional map views, satellite views, and an annotated hybrid of the satellite view. It also allows the research platform to be deployed as a cross-platform solution and attract as many users within the target demographic as possible (the target demographic being the MIT community).

Functionally, the API allows the population of the base dataset with any arbitrary object we can construct within the Flash platform, similar to the methodology utilized by GIS solutions (but lacking the necessity of a layer-based metaphor for interaction).

Using the provided functionality of the API also facilitates the ability for the user to choose a satellite, standard map, or hybrid view of that location, further customizing the user's experience. These factors all aided in the decision to use the Yahoo! Maps API, which implements the Flash platform; however, the core reasons are guided by the goals of the project. First, the Flash-enabled API allows cross-platform accessibility without any additional effort. Next, using Yahoo!'s map allows the application literally to post objects at any location in the world without limit. This ability ensures that this part of the system is portable to any other organization that wishes to use it. Lastly, anyone can download the software needed to run our application easily and free of charge, including as many potential users as possible.

On top of this base API, we incorporate our own base layer of MIT building information drawn from our own GIS repositories. Data is transferred to the representation through the web service layer constructed in ASP.Net and we user a local Jabber server to support the notion of 'presence' and allow real-time user communication.

The client is a complex web application developed in a mixture of Actionscript, JavaScript, and ASP.NET. The principle visualization component built on top of Yahoo! Maps is programmed as a Flash Component in Actionscript. This itself is embedded in a traditional HTML framework with a supplemental visualization frame at the top of the page. Messages are passed back and forth between this frame and the map component over JavaScript. The top visualization component is generated by an ASP.Net module. Because of its design we refer to this upper visualization component as "The Billboard."

The basic user interaction works in the following way. The user opens the application and is confronted with a user registration or log in. New users go through an email process that confirms their Athena identity for MIT before they are allowed to log in. There is also guest access.

After a user has logged in he or she is presented with a perspective that showcases the map component with a situated billboard at the top of the application. Initially the map is focused on the user's current position at MIT and the billboard presents general information about the social community and status. From this point the user then has several options, all of which are reasonable interactions with the application.

The user may elect to engage in social dialogue. The system supports a friend system with the notion of presence and instant messaging. They are also able to add friends, remove friends, and see events that their friends are interested in or attending, or even where their friends are currently located on campus.

The second principle use is for the user to browse the events that are occurring on campus. The user can see where the events are occurring as they are placed spatially on the map; he or she may also elect to view additional information about the events in the billboard. This includes expanded information about the location at which the event is occurring, the nature of the event, including information about accessibility, price, and so on, as well as social information indicating which users have expressed interest or committed to the event.

There are several interactions that allow the user to participate explicitly in the social landscape. They may flag events, a light unstructured action that can be attributed to interest, willingness to attend, or simple notification to others. They may also tag buildings, events, and users with additional resources. This is a behavior similar to linking where a user may associate one of his or her friends, a location, or an event with a blog entry, web page, or similar net resource. This is also relatively unstructured, but in general, it takes the form of associations where the subject is explicitly mentioned.

Another interaction is the user profile. Each user is able to capture his or her photograph via webcam (or upload photos if desired) and edit information about him or herself. This allows a rich expression of the user's personality and gives us additional user contextual information. Some profile information that is available in standard MIT databases (first and last name, Athena name, campus address, etc.) is fixed and cannot be edited. Additionally, support for social site profiles is also a capability. This assures us that any additional information about a given user can be linked to a specific person.

7.2.3 Design Patterns

There are a number of design patterns and elements that have been incorporated into CampusMap in order to explore certain areas of interest, fulfill its primary role as a social system, or for the purposes of increasing the amount of captured placial knowledge. These are elements that serve as archetypical components, either from a functional or design perspective that can be incorporated into other applications to support the acquisition of spatial knowledge and basic user goals.

Creating a new application is as easy as the following:

```
var app = new SpatialApplication(MITDataBroker,
YMapView, SimpleInterface, ...);
```

This sets up the basic spatial application class. The main classes, the databroker, the view, and the interface (controller) can all be swapped out for custom classes or (as written above) use the MIT defaults.

7.2.3.1 Perspective

A subtle design feature in CampusMap is the use of a slight perspective, while the more richly dimensional perspective employed in the design exploration is not used fully. This is partly the result of limitations from the Yahoo! Map data and partly a decision to mimic more closely the presentation of current web maps. This perspective, coupled with the billboard feature and the objects placed on the base layer, create the sense that the user is looking "into" rather than on top of the map.

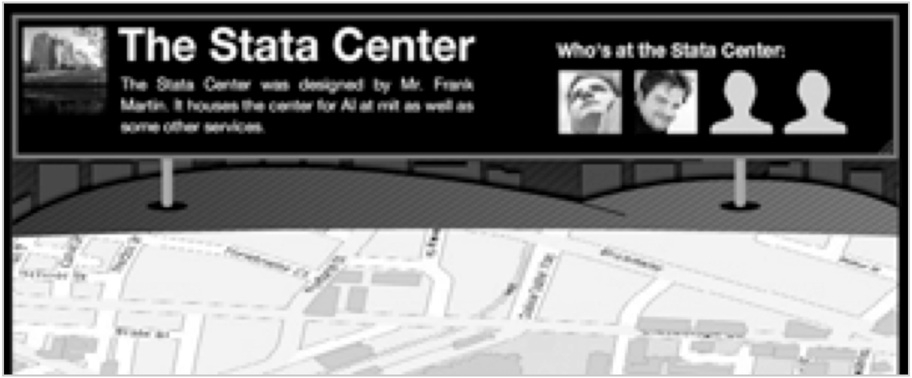

***Figure 11**. Overview of perspective usage in CampusMap.*

Not only does this promote an immersive quality in the representation, it more adequately mirrors the 2 1/2 D representation used in computational vision, which, in fact, is quite similar to how the mind seems to view large spaces. Human cognition, rather than representing relationships as three dimensional (as we might assume), or even from a two-dimensional perspective, more closely follows as a horizontal representation with an additional characteristic of position (Marr 1982). This accounts for the larger errors of measurement related to the third-dimensional qualities (evident in estimations of steepness and depth) and provides justification for a more cognitively native perspective.

7.2.3.2 Map Component

Most information is presented to the user spatially, positioned at the places where it actually occurs. There are various goals to this end. One goal is simply based on investigations that show that spatial proximity (and available time) is the most limiting factor in decision making. It is also intended to promote spatial investigation, exploration, and thought in a natural way. Rather than simply interacting with information outside of a spatial context, users are forced to consider spatial reality in their decision making and interaction. This adds value to user accounts and information presentation.

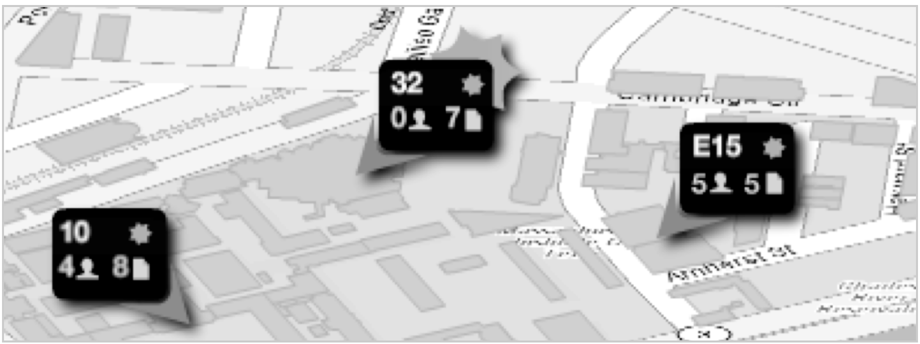

***Figure 12**. Basic CampusMap Markers showing number of events and people at a place.*

The most complex component, the View component, is usually implemented in CampusMap as an extension of Yahoo!'s Map Component. At application creation the following lines

```
var ymap = new YMapView();
ymap.onMapInit(data);   //callback function
ymap.initialize();
```

construct the map. Markers are constructed by calling the mark events function, which in turn passes off the data to specific markers:

```
ymap.addCampusOverlay();
ymap.markEvents(eventByLocation, friendByLocation,
locationList);

if(markerExists(loc)) {loc.addToMarker(data);}
else {ymap.addMarker(loc, data);}
```

7.2.3.3 Marker

On the map, information is presented to the user as a placial marker. These markers appear at distinct spatial locations, but for the most part, they are constructed and organized by place. If the Media Lab is a referenced place, it appears vaguely at the media lab's location, even if the event is nearby or outside (as long as it is referenced as occurring at the media lab). This means that although markers are positioned spatially they are organized placially.

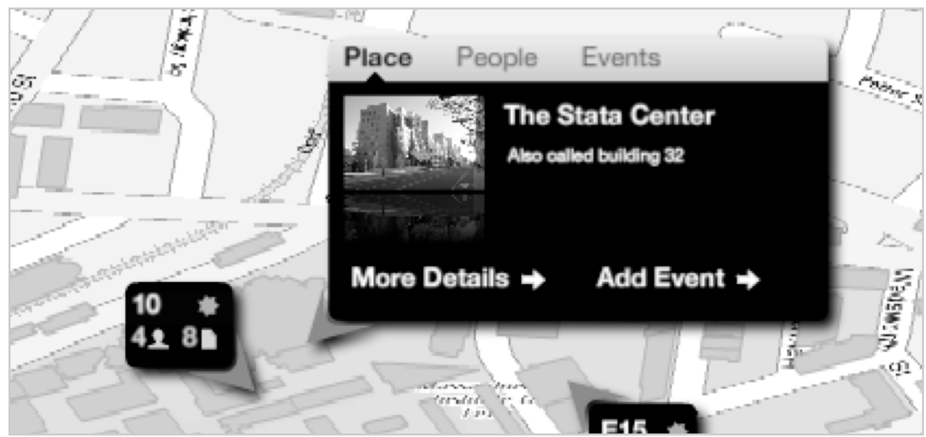

***Figure 13**. Expanded CampusMap marker showing detailed information about a place with possible user interaction opportunities.*

Each marker has a number of tabs. One of these tabs is a description about the place. It shows a photograph and allows interactions (adding an event, getting more information). For the most part the interaction occurs in the other principle tabs, events and people. Events provide an interactive listing of events occurring during the day at that location. Users are able to flag events directly from the marker. The people tab gives an interactive listing of individuals who are associated with the user as friends, colleagues and so on. Here a user is able to message their friends through an interactive messaging system built on Jabber. They are also able to get more information about their friends.

Marker construction, initially prompted by the successful retrieval of data from the DataBroker and passed to the View component, is simple:

```
m.initialize(eventList, personList, loc);
```

This information can later by updated (usually by the DataBroker):

```
m.addEvent(event);
```

Additionally various helper functions can directly make use of this data:

```
var numSportsEvents = m.contains("sport");
var lastChatMsg = m.lastMessage(friendName);
```

7.2.3.4 Billboard Display

The billboard is a dynamic web page that resides in a frame on top of the rich map component. The principle purpose of the billboard, unlike the map component, is expanded information presentation, rather than interaction. When looking at a building here the user is able to see which users are currently available at the building. When viewing a friend he or she is able to see which events that person has flagged. Similarly, viewing an event shows a listing of which friends have flagged it. When the user first logs in, the billboard presents a summary of the social activity that is occurring for that given day, showing, for example, the most flagged events and the most active users.

Figure 14*. Billboard display showing expanded information about a place, as well as a listing of who is there. From here a user can see what events a person has tagged today or in the past, what events are currently active at the place, and so on. The system also records this information.*

Essentially the billboard component is designed to feel like an integrated but distinct component of the application. From a design perspective, the decision was made to move away from more traditional (and clearly nonintegrated) side tabs to display this information. The billboard is under the control of the user and any information displayed inside is the result of user interaction and behavior. The size of the billboard can also be altered based on user demands.

In addition to these components there are a small number of features that have been added purely for user experience. This includes a number of inspectors (such as a friend inspector and an event inspector) that provide a traditional (nonspatial) view of the relevant information. It also includes a data-changing capability to browse and view future or past events.

To load the appropriate billboard data, the client, upon successful user action, makes a call over javascript to update the billboard to the necessary target (person, place, event, etc.):

```
addToBillboard(target);
```

On the server side application, an ASP.NET page loads via

```
initPage();
```

dynamically populating the necessary view information. Actions in the billboard simply trigger xml data updates over javascript, or additional pages that are rendered by the `initPage` method.

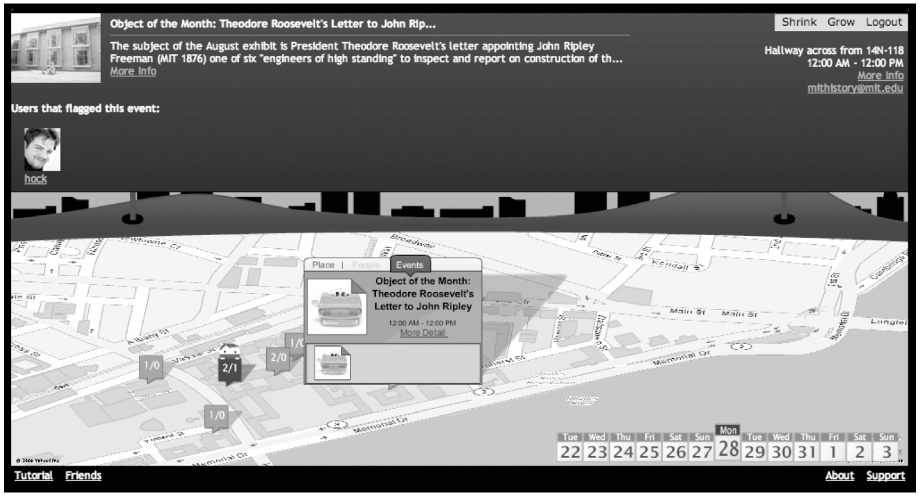

Figure 15 (above). Live version of CampusMap available now. ***Figure 16 (below)*** *a simple illustration of PlaceSense generating semantic concepts for place information.*

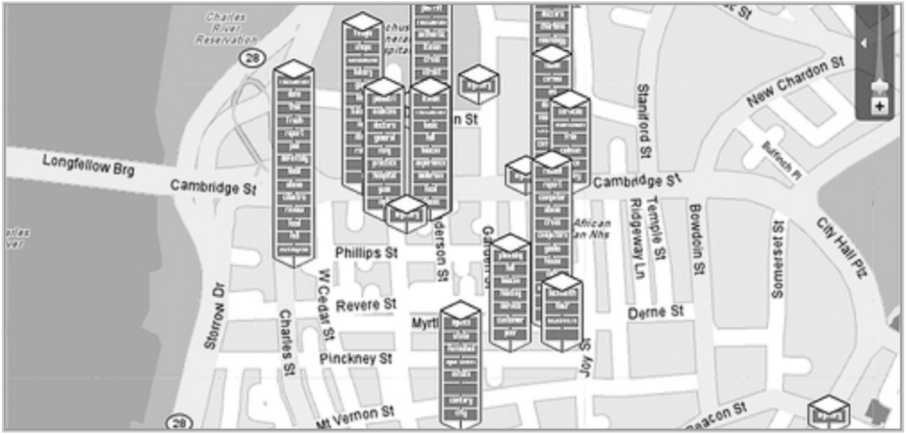

Sample End User Use Case

Manny, an MIT sophomore, logs on to the CampusMap system as his class is finishing. He has a short 45 minute break until his next class and he's looking to find some free food.

After logging on, the system the map loads and centers itself around Manny's current location, the Stata Center. He begins browsing the events in the markers near his location. As he is doing so, he notices a lecture happening later at the Stata Center. He flags the event because he is interested in it and wouldn't mind going. Manny hopes that one of his friends will see the flag and be interested in joining him.

Manny notices that one of the markers near his location is throbbing. Clicking on it, he sees that his friend Guido has messaged him to tell him there is a lot of free food available in Building 10. Clicking on Guido's icon, Manny sees all of the events Guido has flagged, including the lecture in the Stata Center later. He tells Guido that he was thinking about going as well, but now he has to get something to eat.

Clicking on the marker for Building 10, Manny sees that there is a festival going on. He also notices that another one of his friends is signed on in Building 10; hopefully they can catch up.

What's happening here:

- ➔ The user was able to flag events for social sharing with his friends.
- ➔ The user was offered a perspective of the events and people around his location, to view information about them, and to interact.
- ➔ The user was able to engage in a spatially motivated chat.
- ➔ The user was able to see what events and places his friends are, or might be, interested in attending throughout the day.
- ➔ The user was able to browse in a system that was aware of his current location, and adjust its information display accordingly.

7.2.4 Backend Components

Not all of the hard work is done just for the user experience. Most of the system work goes toward ensuring the proper association of information and place-capturing capabilities.

7.2.4.1 Login/Registration

There is a basic user registration and authentication method. While previous applications have used more complex security (MIT Certificates), this application has moved away from this approach. Instead, user authentication is done primarily by providing one's Athena ID to the system. An email is then sent to the associated account, and user access is enabled. This ensures the validity of the user, associates them with their MIT identity, and provides basic authentication.

As stated, the current login implementation is very simple and all of the hard work is handled on the server and only the result is passed:

```
login.checkUserCallback = function(result) {
if(result) {app.startUser(user, pass)}; }

login.checkUser(user, pass);
```

7.2.4.2 Data Brokers

Functionally, all of the transaction information from the backend to the client is handled by a set of brokers. Placial information, for example, is drawn at run time to allow the creation of user-defined places or group social places that may not be a fixed part of the community perspective. Event information and permission listings for users are also passed to the client through this mechanism. The brokers are responsible for keeping the client up to date. When an event is added by the user the EventBroker updates the backend event listing and new event information can then be transferred to other users.

There are four principle data brokers—FriendBroker, LocationBroker, EventBroker, and MitDataBroker. The **MitDataBroker** is a container class to hold the data brokers, while the other three perform all tasks related to retrieving information from the centralized server, and storing it internally to the application for efficient use.

FriendBroker: The FriendBroker class is responsible for handling all of the user information, including friends, current users, and potential friends. The FriendBroker allows for such functionalities as getting a list of friends, adding/removing friends, and updating user information. One of the most important functions of the FriendBroker is the ability to determine both the location of the user and the locations of his or her friends. The correct updating of all person-marker information is dependant on this.

LocationBroker: The LocationBroker handles all information regarding locations stored in the database memory. This is heavily dependent on a Location class that stores all information about a given Location. The LocationBroker is highly needed for the methods to retrieve all buildings on the map, and reverse lookup of buildings based on building number. An example of this would be placing a friend on the map at the proper location, after just receiving their data.

EventBroker: Possibly the most important of the data brokers, the EventBroker communicates with the database to retrieve and process the events from the database. EventBroker handles all events as it retrieves, sorts by location, and displays the events on the map. This forms the basic functionality of the CampusMap system.

Each broker is responsible for gathering a semantically significant series of data, keeping it updated and organizing it for client use. They seem complex, but in usage they function in a very practical manner:

```
broker = new EventBroker();
broker.initialize();   //calls retrieve function
```

Each broker has a handful of utility functions (such as `broker.refresh()` or `broker.expunge()`) that allow the data to be updated, cleared, and otherwise manipulated. When the appropriate method is called via initialize or refresh, the broker will query the server and collect the appropriate data:

```
broker.retrieveCurrentEvents();
```

This function makes a query to the server, where the server side method getPublicEvents is able to query the database, and return a SOAP WSDL formatted piece of data for the broker to deal with in its callback function:

```
broker.retrieveEventsCallback = function(result) {}
```

At this point the data is essentially organized for efficient local usage. It is available via some public data member arrays and various helper functions.

For the most part, information handled by the brokering system is relatively static. This is not to say that it cannot change frequently, as users are prone to update event information and edit friend permissions with some frequency. Rather, it is not as necessary from the user perspective to keep a real-time transaction about this information. A newly added event does not need to be available to all users instantly after the submission button is clicked (although it will be, shortly thereafter). Real-time information is managed by the Jabber server. Much of this information is handled by the PlaceSense parser and gatherer before and after its utilization in the CampusMap.

7.2.4.3 Presence and Communication

We run a local Jabber server that each user is logged into when they access the system. This handles presence information, user-to-user communication via instant messaging, and other functions that are necessarily "real time." It is important from a user's perspective to know when a friend is available and that messages sent via instant message actually do arrive instantly. It also offers a useful future mechanism for similar real-time capabilities.

The Jabber authentication is separate from the login, but functions similarly. Technically, the jabber username/password is different from the system username/password, but in practice they are the same:

```
chathandler.initialize(jabberuser, jabberpass);
```

From this point the basic interaction is to send a message with:

```
sendMessage(person, message);
```

Where the message is either a simple text message, or a command (request authorization, etc.).

7.2.5 Data Collection

The system that has been described appears to be an interesting system for building a spatially centered social application. That is one of its goals, but not the overriding one. The system is primarily a research platform. It is useful that it is a feature-rich and usable social application, simply to allow us to gather as much information as possible about the user's placial conceptions.

The system gathers rich sets of data about usage. It considers which events users have looked at, flagged, requested more information about, and so on. It also stores user conversations for parsing, including between which users these conversations occurred and where they were located. As much effort as possible is made to provide contextual information for data points. A user doesn't just look at an event, that user looks at an event given their behavior history, the current event and person climate, the time, their location, any ongoing conversations with other users, and so on. This contextual information is intended to allow better analysis of the impact of user actions as "experiential accounts" or statements about a place.

PlaceSense is the backend parser and interpreter for considering user data and supplied information. The general premise is that we are looking for human experiential accounts of place. These can be subtle implict statements, such as textual chat involving a place "that's a cool place," which must be interpreted based on the user conception of coolness. The information available about the user, event, or place that already exist form a contextual model for interpretation. While we may rely on a general impression of what coolness is, we also introduce user actions and behavior in attempting to offer suggestions about cool.

These experiential accounts form the basis for our understanding and reasoning about place. PlaceSense takes these accounts and performs after they have been semantically analyzed and interpreted with common-sense knowledge. Once stored, it forms the working corpus of our claims about a particular sense of place.

This is very different from most traditional GIS systems. While these systems rely on objective, tabular data or various finely developed xml specifications of spatial data, the primary consideration in this framework is human textual account. Essentially any piece of text that can be associated with a physical location or even an abstract (but verifiable) place is relevant input for the system.

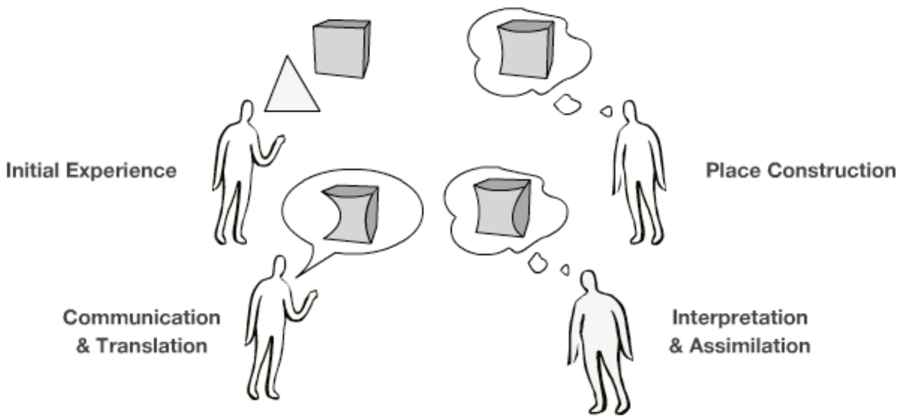

While PlaceSense is designed to work with information from CampusMap, it is capable of doing some general parsing and spidering for information. This process begins with a trusted data source. In CampusMap input this represents information that has been associated by a user with a particular place. This may represent a file, a web page, a blog, and so on. This information is linked with a particular abstract place, which may (or may not, but which usually is) be associated with a particular spatial description. This could be a latitude-longitude point, a street address, a postal code, and so on. Generally, this information is accurate enough to estimate its position in geographic space. As will be described, precise accuracy isn't completely necessary because we can use semantic proximity for some location correction. For a web page some structure can be retained, but when necessary formatting can be stripped to leave the base textual description.

Once this information has been gathered or inputted into PlaceSense, the system needs to process it for meaning. The primary goal of this processing is to understand what the account means in a general sense, the way an average user might. This does not claim that this particular account is average (it may, in fact, be a very divergent account) but at this point the system treats it as a general average account of the place.

The general behavior here is that we feed some of this information into a processor built on top of ConceptNet. ConceptNet is a large-scale commonsense knowledge base with a natural language processing toolkit that performs textual reasoning tasks over the input. It is useful to imagine this as a large graph structure of various concepts (nodes) with associated links representing different kinds of relations over those concepts. Given a particular input we attempt to 1) generalize it into the highest level but still meaningful concept and 2) determine what relations are necessary for determining the actions and perceptions associated with the particular account.

System Response to Use Case

Consider how the system responds to the actions that occurred in the presented end user use case scenario:

- ➔ The system records Manny as being at the Stata Center at the given time, it associates his profile with the Stata Center.
- ➔ When Manny flags the lecture event, it associates his profile with that event. This is added to the existing associations with the location (the Stata Center) and any other users or events that have been associated with that location, event, or any of the users.
- ➔ It records the details of the chat between Manny and Guido, this account is associated with identified places (Building 10) and the users themselves.
- ➔ Manny's friend who is already in Building 10 has his profile associated with that building, as well as the free food event if he has flagged it.

How these accounts impact understanding will be discussed in the following sections.

7.2.6 Concept expansion

This context expansion makes use of ConceptNet's contextual neighborhoods. The contextual neighborhood around a particular concept is formed by spreading activation of the dataset and finding relevant related nodes from the source concept node. This makes relatedness of a particular node a function of its link distance from the source.

In order to determine the more general a concept is, we simply follow this linking as we find nodes with an increasing number of links. Concepts like "dine" can expand upward until we find a more general concept like "eat."

Of course it is possible that we find related concepts that are not actually expansions of meaning, but rather new kinds of meetings in general. "Dine" may just as easily expand to "socialize." We employ the concept of realm filtering to distinguish between expanded concepts and related concepts by relying on estimates of importance and weighting each semantic relationship type for the domain-specific tasks. This allows focusing on temporal, spatial and in particular action-only neighborhoods.

Still, to ensure expansion as a distinct focus from relation, we simply monitor the kinds of linking relationships we follow. To expand "eat breakfast" we may follow only generalization, but we may also be interested in following the relationships that lead us to the physical requirements (kitchen table) and their general location (a house or dining hall)

Occasionally, unique words that generate no matches in OMCS and conform to a simple rule schema are accepted as meaningful descriptors. The goal of this is to include important concepts of a place (such as Italiano, which may not have appeared in OMCS) while excluding meaningless words.

Input is further reduced by comparing it to a stop list of common spatial terms (street, avenue, etc.) that add no additional semantic meaning. At this point the output of this process is a set of tags representing meaning. The tags are tuples containing a word and an associated rank. Tags ranked below a certain threshold are omitted. This allows us to capture the high-level concepts (food, restaurant), accurate lower-level concepts (pizza, Italian), and ignore less relevant lower-level concepts (sauce, cheese).

This description of the content already serves as a reasonable basis for consideration of a place. However, we are particularly interested in active and perceptual features that may need to be teased out. For example, "food" is a useful descriptor of a place, but it is not explicitly descriptive. It would be more useful to understand explicitly what kind of food and why one might care if the food is contained at the place. Food at a restaurant and food at a grocery store carry very different meanings.

Possible Developer Use Case

Manny is developing a spatial application that presents a representation of open real estate data on a map representation. This is not an entirely uncommon application, as several others exist. Manny wants to set his apart so that it more accurately allows buyers to find not only the price and physical requirements of a property, but a better understanding of the actual semantic climate of the area. This is particularly important in the real estate market, as it is a major purchase and must be ideally suited to the buyer's needs.

Employing some of the capabilities of the PlaceMap system, Manny is able to achieve this goal. Instead of basing this application on one of the available web mapping APIs, Manny decides to base it on the more robust PlaceMap system. There is sufficient motivation to do this simply in the ease with which the PlaceMap system handles many common developer tasks (organizing markers containing real estate information, supporting user login and association of past data histories and richer visualization capabilities). There is also significant benefit in the kinds of displays that the backend PlaceSense parser can help generate.

Manny constructs a basic application that supports, but does not require, user login. This allows the history of viewed properties and their association with user profile (including requirements, price range, etc.). This also allows comparison between what the system identifies as similar users. Manny also utilizes a side display that functions similarly to the billboard, but without the integrated look and perspective.

Deciding to employ the PlaceSense system in the background, Manny focuses on accounts of place that are related to common real estate questions. These are sometimes components that are hard to quantify: How 'nice' a neighborhood is, or how exciting it is to live there. Even general concerns such as the identification of a region as 'suburban' or 'rural' can be made to some degree. With this information Manny then associates an overlay based on color. Based on user search or profile information, the 'niceness' of the region on the map is indicated in color. Manny also places sliders on the interface that allow the manipulation of aspects of this, readjusting the weights for the associated accounts of place.

7.3 Experiential Accounts of Place

What we actually want these experiential accounts to look like is a valid question. What we end up capturing with this loose tagging is not bad, and in general returns reasonable results.

Example: Consider the above 'problem,' where there is the concern that we can't distinguish between groceries and restaurant food. If we perform a search option for "food," it is true that we will receive both tag matches. However if we search for "eat out," "dinner" or "American restaurant," we won't. Similarly, "buy some groceries" returns grocery store profiles and not restaurants. We can perform concept expansion on search terms as well, and functionally tagging represents decent representations of space.

Still, it seems like there should be a better way of capturing this as a semantic representation. Let's look at a protostructure for these experiential accounts of place.

Experiential Pseudo-Account of Place

Place: Student Center -> Stratton Student Center -> Building W20
Textual Account: "I like the student center, it is a pretty good place to get a meal."

Actor: User 16 -> Athena Account information
Object of Action -> Meal -> Food
Action:- Eat -> To Eat
Underlying Motivation: Hunger
Related Actions: Drinking, Socialization, Dancing -> Socialization
Possible actions and support: Eating, socialization
Valuation: Good -> Not Great, Greater than Average, Not Bad

This actually measures up quite well with emergent affordance-based models of place. In these models a means-end hierarchy is often the presented model. This model captures purpose, abstract function, generalized function, physical function, and physical form. Jordan et al. consider a reduced model of this for clarity, the why-what-how model. Our representation of the student center here would more closely resemble the example they present:

Simplified Means-End Hierarchy

Why: To Lunch To Snack
What: Socialize, News, Food Consumption
How: Talk, Eat, Read Newspaper, Observe Others.

Experiential Account of Place

Place: Student Center -> Stratton Student Center -> Building W20 **Textual Account**: "I like the student center, it is a pretty good place to get a meal."
Account Confidence: 9.6 (High Confidence)

Supporting Actor: User 16 -> Athena Account information ->
Link to all associated user accounts.

Object of Supported Action: Meal -> Food (Importance: 2.5)
Supported Action: Eat ... -> To Eat (Importance: 2.3)
Underlying Motivation: Hunger ...
Related Support: Drinking, Socialization, Dancing ... -> Socialization (Importance 7.8)
Resultant Actions with Suggested Support: Eating, Socialization

Valuation: Good -> Not Great, Above Average, Not Bad (Support: 6.6)

As we will see, it is not particularly necessary for us to utilize a means-end hierarchy and the protoaccount is a reasonable description of these experiential accounts. By utilizing the linking relationships to necessary relations within ConceptNet, we can find the nodes necessary to provide this information. We also associate weights with elements of the account, and with the account in general as it compares to other accounts. These weights can be adjusted based on a particular query, but in general they represent our confidence in particular parts of this account or within the account itself. It is also reasonable to archive the information that was used to generate the account so that at any time in the future a new account could be generated from altered parameters or the original narrative, meta-information, and user information being recalled.

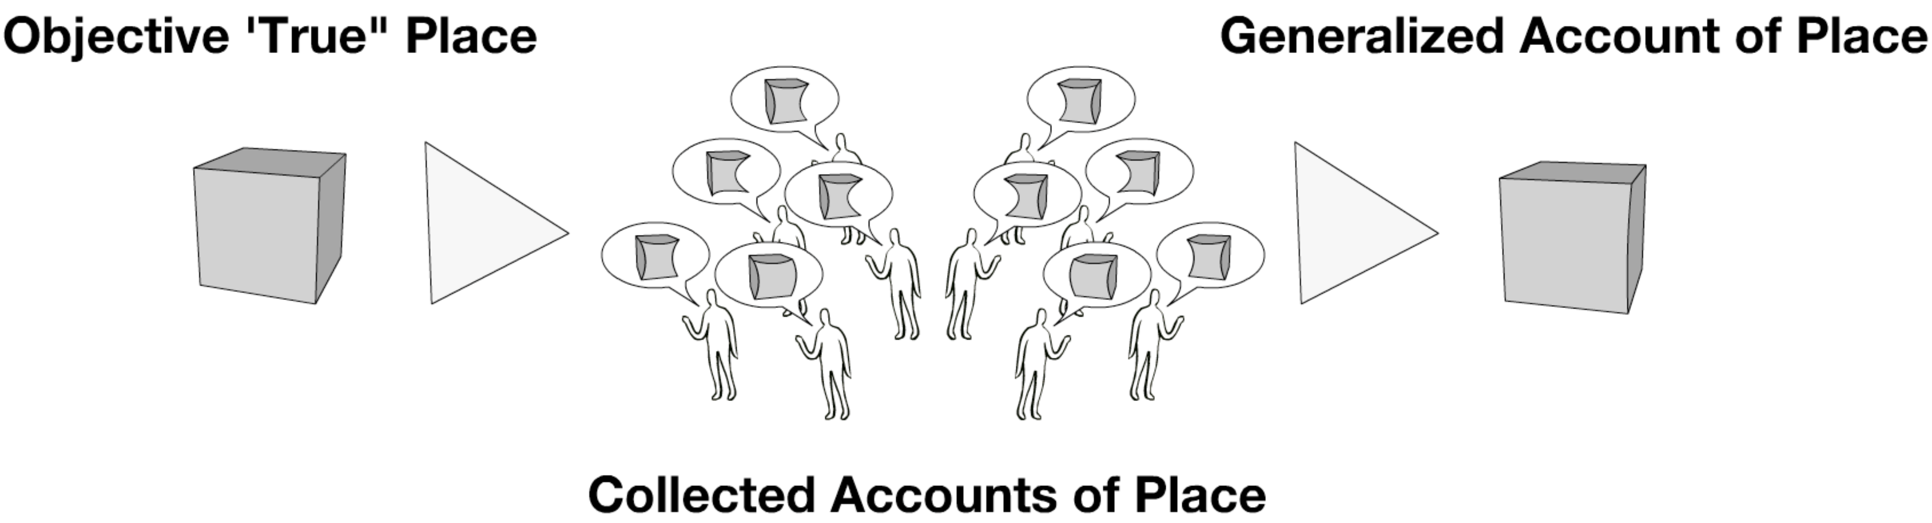

Here the system is highly confident that this is a reasonable account of this place based on success in parsing, its evaluation of the user, and other factors. The aggregate determination is that food is of relatively low importance and that the valuation of its support for food and eating is above average. Socialization, a related activity it also supports, is significantly more important to the makeup of the place. These weights (out of 10 points) represent importance to the general makeup of the place, the place's support for the given account activities, and our general confidence about the account. While a particular account wouldn't be reweighted automatically, the general aggregate account could be if a particular set of actions were favored by the given user's profile, weighting the aggregate account information based on a particular context.

The mechanism used for determining accounts of people or events are slightly more general. They are not structured as much and are more generally represented. The goal of the system is to produce accounts about place, not about users and events. Models for implicit actions (flagging an event at a place, being at a place, etc.) are interpreted as statements of support with low confidence. For example, a user who likes dancing and is constantly at a place suggests it may support dancing, but this is assigned a low confidence and could be bad data.

A dancing event held at a place, however, represents high confidence (with significantly more investment). These filter down so that a user flagging an event represents with high confidence a statement about the user and the event, with some confidence a statement about the event and the place, and with low confidence a statement about the user and the place. This attempts to capture the increasing distance of support for action as the distance from perceived support for that action. One aim here is to move commonsensical accounts away from data entry and towards implicit inference of support.

7.4 A Complex Accounting

An interesting component of the research platform is the utilization and reliance on an emerging model for community social systems that is developing within so-called Web 2.0 style applications on the Internet. This model is based on the different roles individuals migrate into when they become part of an interactive online community. This model, which can be observed in social media sites such as Flickr, Upcoming.org, Del.icio.us, increases in importance as questions of scalability arise. This is also a model that can be observed in older social services, such as message boards and even early bulletin systems.

The model is described by Horowitz (2006) as the Content Production Pyramid. The idea is that, given a large enough user population, certain user groups will naturally fall into several distinct (but possibly overlapping) roles. Given a user population of one hundred (100) users, the breakdown might be as follows:

Creators: Out of the hundred users, one user actually might be the creator of new (novel) content. This can include posting an event, creating a discussion group, adding a new capability to the system, and so on. They are performing tasks involved in creating and authoring new material without outside input, but motivated by their own desires and drive.

Commenter: Ten users might actively participate in the activity. This includes notifying others that they attend an event, writing comments about an event, and associating information with existing places.

Consumers: The rest of the user population takes less of an active role in the content production and editing, but they benefit (as do all users) from the actions of others. These include people who log in and access the system, and might have a few conversations, but for the most part simply absorb and use the content privately.

This model highlights that a social system doesn't necessarily need a large percentage (and certainly not a majority percentage) to be highly valued and generate value. It also suggests some interesting concerns and statements about the user group. One of these is that even with the most limited and relaxed barrier to entry, not everyone is going to be the most active participant in the community. The consequence of this is that we should try to harness their energy with more implicit

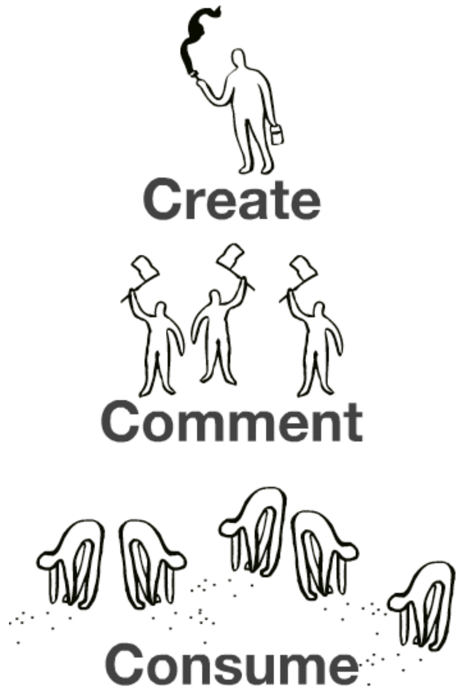

creation, where simple consumption and use encourages useful data. The area most apt to use is to try to grow the number of commenters, but even casual user data can suggest a lot of implicit value.

Specific to the research architecture, there has been a substantial attempt to create a low barrier to entry for casual use and commentary. This includes associating (linking) information to a person, place or event. This also includes the messaging system and event flagging capabilities. However, even without active commentary, significant sources of information can be generated from casual use. Given that we know a person's MIT account information, even without a deliberate profile created we have access to a lot of information

This information can be used and associated with particular events, places, or people. At some level we take a person's profile information to be another account of these things. When a particular person looks at an event, for example, we associate that profile information with the event (and ultimately with the place the event occurs at) as a statement about the event. If this is a casual glance, this may not be a statement that carries much weight. Measuring the length of observance, noticing repeated observances within a fixed period of time, or noticing resultant interactions such as conversations about the event or flagging the event, the system can be more confident in its association of this particular profile with this event.

This results in a complex series of inferences. We assume that (given certain conditions) this event then supports (or at least interests) this person, perhaps this kind of people. The event is another higher-level statement about the place.

These user populations, those that create events, are more similar to the top one percent discussed in the creation model above. The investment of having an event is large. We take the act of having an event at a particular place to be a statement of support for that place's capability to support that kind of event. For example, a dance is probably not held at a lecture hall. It is possible that one or two dances are held there, but in the aggregate consideration it will be clear that the lecture hall is not the kind of place to support such an event.

Thinking about these various actions as statements of support is an almost Darwinian account. Places that are well suited to certain kinds of activities will (in the long run) support those activities more often than other places. Similarly the kind of person who is at a place or event is a statement for that place's support for that kind of person. While there will be some outliers and random occurrences, there will also (given enough data) be trends. By tying this information to explicit statements about place nature and support, we begin to arrive at a fairly good understanding of what a place means in an active way.

This concept is similar to other interactions in social landscapes, such as in the interestingness algorithm implemented at Flickr. Instead of relying on explicit determinations (rate this photo) the decision was made that this metric utilized was prone to exaggeration and self interest (in a negative way), and failed to address a significant population of the user base. Instead this measure is based on natural activity and traversal through the site. Although its explicit implementation is unknown, it is based on a number of factors such as who views a picture, and how many times it is tagged and commented on. Without explicit voting and without disruption, the user population is nudged toward the middle group.

“Without anyone explicitly voting, and without disrupting the natural activity on the site, Flickr surfaces fantastic content in a way that constantly delights and astounds. In this case lurkers are gently and transparently nudged toward remixers, adding value to others’ content” (Horowitz 2006).

7.5 Evaluating PlaceMap

PlaceMap is offered as a private system with limited small-scale use, although this will eventually expand to support a larger user base. The placial knowledge collected by PlaceMap can be directly transferred to the kinds of representations evaluated in the design exploration. These representations have demonstrated significant cognitive gain in information search, and presumably suggest semantic constructions more closely aligned with human ones. As the amount of usage for PlaceMap increases, the reliability of these semantic accounts increases and brings the resultant representations closer to human conceptions. This is significant because even the more general representations used in the design explorations showed significant improvements to the efficiency of knowledge search.

User response to the PlaceMap systems is generally positive. Since it is a more simplified version of the representations explored in the design exploration, it offers the same kind of results. This is supported by user response to the system. It grounds the user at the center of the social world, understands action, and embodies action in its representation. It supports a human conception of place. Additionally, many users report that the functionality of the system is sufficient for their continued use, and that they find the features to be natural without being intrusive. Here are several (distinct) user reports of their experience in the system.

User Response to PlaceMap

➔ "I found the sharing of events to be nice. It makes sense to be able to look directly at a building, see what is going on, and just click to tell people you are going to be somewhere or even message them if you want."

➔ "Adding information to a place is something everyone can do. I've done it a few times, it is not hard—you just copy and paste a link. Sometimes I do it just so other people can know where to go."

➔ "I still use Google Maps, but whenever I really need to know about what's going on and not get directions I go here. It just makes a lot of sense to me to be able to see everything that's going on and where everyone is. I don't mind sharing where I am when I'm here [on MIT campus]."

1
Object of the Month: Theodore Roosevelt's Letter to John Rip...
The subject of the August exhibit is President Theodore Roosevelt's letter appointing John Ripley Freeman (MIT 1876) one of six "engineers of high standing" to inspect and report on construction of th...
More Info
2
Shrink Grow Logout
Hallway across from 14N-118
12:00 AM - 12:00 PM
4
mithistory@mit.edu
Users that flagged this event:
3
hock
5
6
7
Events
Object of the Month: Theodore Roosevelt's Letter to John Ripley
12:00 AM - 12:00 PM
More Detail
8
9
10
11
Wed Thu Fri Sat Sun Mon Tue Wed Thu Fri Sat Sun
22 23 24 25 26 27 28 29 30 31 1 2 3
Tutorial Friends
About Support
12
13

Figure 17. Live Version of MIT CampusMap Application (A Sample PlaceMap Application)

1) **Billboard area**, showing information about the place and possible event and user details.
2) **Event title and description**, drawn from MIT events page.
3) **Event sharing**, showing users who have flagged events. Clicking on the user picture shows the user profile and a listing of all the events they have tagged and their current location.
4) **Meta event information**. Options appear here for events, places, and people that allow manipulation such as adding events, or associating web pages with people and places with image data drawn from Yahoo! Maps.
5) **Perspective horizon**, producing an inward perspective rather than a top-down one.
6) **Map display**, the main map display showcases places, events, and people spatially.
7) **Open marker**, showing a flagged event. Interactions here display information in the billboard. Multiple events are scrolled in the bottom shelf.
8) **MIT campus overlay**, showing MIT data provided for buildings and other spatial features overlaid on the map display.
9) **Unopened marker**, showing the number of active events and the number of active users.
10) **Active marker**, showing current user location and also showing available people in buddy icon situated on top of the marker.
11) **Time slider widget**, showing the current data in red with other dates in the future and past selectable to show previous and upcoming events.
12) **Inspectors**, links which launch additional information generally related to non-spatial views of the relevant data.
13) **Meta links**, showing general links for additional information and support for the application

8. Towards a Naïve Geography

A primary goal of this work is to address the clear discrepancy between the traditional map representation models and actual human conceptions. This limitation is evident even in determinations about the regions or boundaries of place. In traditional map systems, class set-theoretic notion of categories and discrete views of space are employed (termed by Burrough as a "double-crisp" model) (Burrough, 1996). Fuzzy categories and spatial continuity cannot be captured in these models, and an average user's conception of space is ignored.

There is significant interest in attempting to mediate the differences between user conceptions and actual models in the GIS community (Burrough 1996; Burrough and Frank 1996). This includes models that incorporate more relaxed or fuzzy areas, or areas of "indeterminate boundaries" similar to those found in place construction; people are thinking about places, not spaces. Towards this end, qualitative spatial reasoning and modeling of this knowledge have become significant subjects of research (e.g. Hernandez 1994).

Clearly this discussion is very complex. By asking for insight into place, one must consider the wider impact this has on models of geographic space and their usage. The model that will serve well to encapsulate this viewpoint is the so-called ‘Naïve Geography.’

8.1 Life in a Naïve Geography

Naïve Geography is the idea of developing formal models of commonsense geographic worlds. Formal models of commonsense knowledge have been examined by philosophers (Smith, 1994), and common-sense physics, or naïve physics, has been an important topic in artificial intelligence (Hayes, 1978). Egenhofer and Mark suggest that formal models of commonsense geography are a necessary prerequisite to the development of geographic information systems that are truly intuitive.

This approach hopes to:

- ➔ Identify basic elements of common-sense conceptualizations of geographic space, entities, and processes, and develop an integrating framework.
- ➔ Investigate users' reactions to intuitive geographic inferences, and compare the inferences with the results obtained with current technology.

According to Smith (1994), Horton (1982) theories can be divided into Primary Theory and Secondary Theory. Primary Theory of the world is self-evident and unquestioned. It describes those aspects of the world for which scientific theories and common sense agree. Naïve Geography continues the call outlined by Hayes in 'The Naïve Physics Manifesto' for a focus on aspects of the secondary theory. These are commonsensical notions that may be incorrect or incomplete from a scientific perspective, but which represent useful working human constructs.

Egenhofer and Mark (1995) introduced the term of Naïve Geography to refer to what might otherwise have been called the Naïve Physics of Geographic Space. Modifying Hardt's (1992) definition of Naïve Physics, they defined Naïve Geography as follows (Egenhofer 1995):

Naïve Geography is the body of knowledge that people have of the surrounding geographic world.

They went on to state that Naïve Geography captures and reflects the way humans think and reason about geographic space and time. (Egenhofer 1995)

Tobler and Egenhofer reiterate some anecdotal (though well supported) elements of Naïve Geography as presented by Egenhofer and Mark (1995):

- Naïve Geographic Space is Two-Dimensional
- The Earth is Flat
- Maps are More Real Than Experience
- Geographic Features are Ontologically Different from Enlarged Table-Top Objects
- Geographic Space and Time are Tightly Coupled
- Geographic Information is Frequently Incomplete
- People use Multiple Conceptualizations of Geographic Space
- Geographic Space has Multiple Levels of Detail
- Topology Matters, Metric Refines
- People have Biases Toward North-South and East-West Directions
- Distances are Asymmetric
- Distance Inferences are Local, Not Global
- Distances Don't Add Up Easily

The authors highlight that a key topic in understanding 'Naïve Geography' "may be a search for the principles, schemata, and heuristics that allow people to find things in novel environments."

Here information seeking again becomes a key litmus test of common-sense understanding of place. The way in which individuals make common generalizations in order to meet these spatial information-seeking goals reveals significant insight into the mental processes and constructions we hold about spatial representation. This is well represented in this excerpt of "Finding Things in 'First World' Economic Systems."

> Travelers often are faced with the need to find goods and services, such as telephones, Automated Teller Machines (ATMs), beer, stamps, or toiletries. To what extent are members of the public aware of general principles of retail and service location, as well as the systematic distribution of retail and service functions across inter- or intra-urban hierarchies? Concepts of thresholds and ranges for goods and services, and the 'order' of places and goods, systematically documented and formalized by Central Place Theory, might be part of geographic common sense, and might be used by people when they have to find such goods or services in unfamiliar places. Such schemata, however, might break down in other places, due to cultural differences or local regulations. Buying a bottle of wine involves very different search strategies in different parts of North America, mainly because of differences in regulations for the sale of alcoholic beverages. Also, goods or services may have different associations in different cultures and jurisdictions. For example, in most countries, stamps are sold in certain types of outlets other than post offices; however these associations may differ from country to country. In Spain, stamps are sold at stores whose primary function is the sale of cigarettes—without local knowledge or guidebooks, it may take the traveler quite a while to discover that in order to find stamps outside of post office business hours, one must look for brown-and-yellow signs saying "tabac." Such principles could become heuristics in vehicle information systems that could operate in areas without local information about shops or services, but having only road networks and basic census data.

Still, there remain breaks with understanding of what a common-sense view of space covers and how it might be generated.

Geographic space is large space, space outside the traditional table-top conception of smaller spatial relationships. This concept alone has been difficult to define. Kuipers and Levitt offer that geographic space is space that cannot be observed from a single viewpoint (Kuipers 1978; Kuipers and Levitt 1988). Fundamentally this implies a view of geographic space that is larger than what a single person can actually perceive at a given time. Pederson (1993) argues that a better definition might be space that does not contain objects that humans generally think of as manipulable objects.

This definition is interesting to consider, because while it appears to be superficially valid (most humans are not skilled at moving buildings and mountains, nor do they think of these activities as valid day-to-day occurrences), it is somewhat misleading. The work done with the CampusMap system defines place mostly by the ability or lack of ability to enable people to perform certain kinds of activities. While these are not direct activities, this definition is more active than that suggested by this definition. Geographic space is indeed space that is large, can be navigated, and often requires multiple viewpoints.

The suggestion made by some that these multiple viewpoints are pieced together mentally like a puzzle is also misleading. One conclusion, however, is that the act of investigation for determining the nature of geographic space is difficult. In a smaller table-top-like space the observer is able to investigate objects actively and gather additional information about them to see touch- and measure-relevant parts. This is significantly more difficult in large-scale geographic spaces where perception is based on human accounts, observation, and to a much more limited degree one's personal experience.

This leads human common-sense geographic experience to rely, to a high degree, on our reductive bias. This is an affinity to construct overly simplistic understandings and categories. Egenhofer and Mark (1995) acknowledge that people's common knowledge "may be contrary to objective observations in the physical world." This includes, for example, the general disregard for the curvature of the earth and acting perception (though not actual belief) of a flat two-dimensional world.

The reasoning that a geographic system may contain errors and inconsistencies is disturbing. While many of these errors can be disregarded as outliers (mistaking one building for another in conversation, mislabeling, etc.), some of these errors may actually be widely held, but incorrect, beliefs. In general purpose systems these errors, even if they are intended to be eventually correct, can result in conceptual models that are difficult to correct and overcome.

However, for most purposes the system is intended to be used by average users for average goals. Any indication of reductive bias that is held by the majority of experiential accounts is not necessarily a 'fault' with the system, so much as it is an actual statement about user belief. The statement that "there is nothing to eat in that building," even if it is found to be false, can actually reflect a poor general knowledge of available eateries or a statement about its quality. Incorporating user behavior with a 'real' or 'true' model in this circumstance actually suggests that the building is poorly suited to eating, and is a valuable conclusion despite the (apparently) incorrect factual statement. Rather than ignore these inconsistencies, we can look to them as powerful statements about the failure of an intended place to offer certain user behavior successfully.

8.2 Experiential Account Methodology

The idea of using textual accounts, as CampusMap does, to determine a rich semantic model is not new. First employed by Silverman in semiotics, the use of topology of viewpoint space acquired by semantic mining with an eye towards inferring a perspective location in these spaces by psychoanalytic reading has been employed heavily by Liu (Silverman 1983; Liu 2006)

Liu models viewpoint as an individual's psychological locations within latent semantic spaces that represent cultural taste, aesthetics, and opinions. This is represented as a computational theory of point-of-view, building closely on existing semiotic/cultural theories of viewpoint, aesthetics, culture, and taste (Bourdieu 1984; Liu 2006).

This approach has also been employed more closely to the domain of geography, albeit with a focus on more concrete goals. Egenhofer demonstrates an approach for formalizing user actions and constructing task-oriented ontologies:

Verbs and nouns are extracted from a document that depicts user actions during GIS tasks. The conceptual structure of the user actions is formalized through a combination of Formal Concept Analysis and Entailment theory. The subconcept-superconcept relations between user actions are then refined. The approach is intended to strengthen the consideration of user tasks in geographic information applications (Egenhofer 2004).

As presented, however, it is a fairly limited model focused on tasks that require direct use of GIS information from expert actors (Egenhofer 2004).

A similar approach is presented by Kuhn. Here a method is proposed to derive ontologies of geographical domains from natural language texts that describe human activities. Through its textual grounding, the method addresses the issue of from where to take the contents of ontologies. Through its focus on actions afforded by domain objects, it establishes a criterion for selecting the contents. The actions are organized into a hierarchical theory of human activities in the domain (Kuhn 2001).

It is reasonable to suggest that textual accounts, in and of themselves, are simply not enough. This is a reasonable suggestion and other experiential accounts of place are reasonable inputs for CampusMap. If commonsense semantics of place are grounded in human construction alone, how can we represent this?

9. A Computational Model of Place

The goal of gathering and analyzing this placial knowledge has led to the development of a working computational model of place. This model is intended to be computational, in the sense that it could be computed mathematically. Realistically, given the complex set of information involved in the generation of placial knowledge, this is not practical. Rather, this model intends to build upon the results of the research and existing account of place to offer a working model of place. While a particular place may never be fully modeled, it can certainly be modeled in general and specific areas to necessary degrees of precision.

The presented theory attempts to address the goals of implementing a naïve, common-sense geography directly as the result of human action, expectation of this action, and a generalized semi-objective representation of its perception.

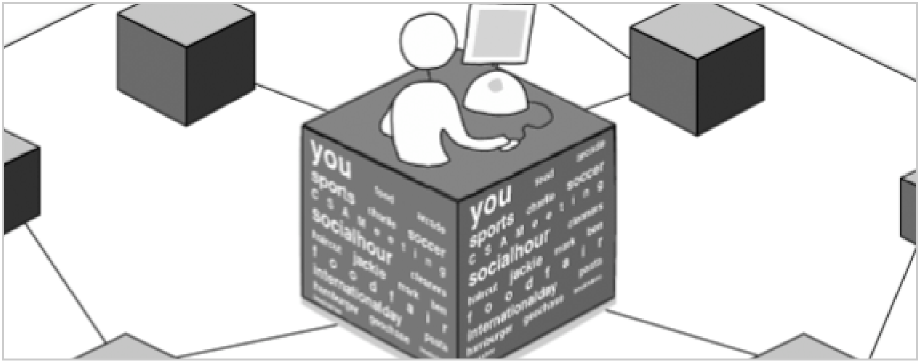

9.1 Affordance Planes

This computational model of place is described as an affordance plane. The concept of affordance can be described as "what objects or things offer people to do with them." This concept was introduced by Gibson (1979), who described the process of perception as the extraction of invariants from the stimulus flux and called these invariants affordances. Affordances create activities for humans to do. This idea was influenced by Koffka's (1935) work on Gestalt psychology, where he states, "Each thing says what it is." One may consider that a doorknob, for example, is particularly suited to the concept of grasping (for which it has a high affordance). It may be less clear how much affordance is offered for turning and opening. The idea of affordance is complex. One might say *a priori* that an out-turned shape sized to the human hand affords grasping quite regularly. It does not afford turning *a priori*, but after years of experience or initial contact this affordance may be very obvious. Opening (in particular, determining the direction the door will move) is afforded as well, but less clearly and often with some necessary experimentation.

In reality, the notion of affordance can be described very concretely. The concept of agent-environment mutuality (Gibson 1979, Zaff 1995) suggests that various aspects of agents (actors) and their environments need to be understood in terms of the relationships between them. According to Zaff, "They [affordances] are measurable aspects of the environment that can only be measured in terms of the individual." By understanding the action-relevant properties of the environment in terms of values intrinsic to the agent, the affordance can be determined. (Qtd. in Jordon 1998).

For example, Warren (1995) shows that the "climbability" affordance of stairs is more effectively specified as a ratio of riser height/leg length. This was demonstrated empirically where subjects of different heights perceived stairs as 'climbable' depending on their own leg length. Other such low-level affordances have been studied extensively. Additional dynamic contextual information is also shown to be a factor. For example, the act of walking produces movement that impacts one's ability to pass through a door, and accordingly to perceive this affordance.

The chief argument against this theory is that it neglects the process of cognition. Lakoff (1987) and Norman (1988) recast affordances as the result of mental interpretation of thing, based on past knowledge and experience to offer an experiential view of space (Lakoff 1988, Kuhn 1996). This framework represents both the physical environment and the contextual or situational interpretation of this environment by the actors.

Here we consider places as affording certain activities. This is a larger scale than described in most work about affordance, but this is not the first work to look to affordance as a methodology to model place (Jordon 1998). This is an attractive methodology because it is rooted in action and active thought. Interactions with places are based on the meaning people assign to them, meaning rooted in past, present and future activity within the place. As Jordan writes, "Modeling places with affordances integrates cognitive and engineering aspects, therefore leading to a knowledge-representation that comes closer to the user." Specifically, it represents the integration of location and the meaning of that location with the context of human action. The general model developed by Jordan describes a place as defined according to the user, for a given task. The difficulty with the model developed by Jordan et al. comes in the actual description of affordance for a particular place. They apply Rasmussen's (1986) means-ends abstraction hierarchy to represent the environment, along with an object aggregation model. This is done to provide a 2D mechanism to determine a set of possible purposes or functions or some configuration of GIS data. Functional user requirements would return the configuration suited to a user's place needs. This model, particularly the object-aggregation model, requires knowing the constituent parts of the given place to make successful determinations. While this is possible in specific circumstances, it is difficult to scale, requires access to relatively low-level granular information about a place, and makes large breadth determinations difficult. They offer an interesting model but there is little suggestion to actually construct such a model automatically through mechanical or computational means. The experimental PlaceMap system addresses this by providing a mechanism for the semantic aggregation of accounts of place.

In the model of sense of place presented here, the aim is to avoid the labor inherent in approaches that require descriptors to be constructed, and instead focus on what can be observed and taken from an existing system. We take each user account of place, given the user profile and context, to be an experiential claim about the true nature of the place. While it is unlikely that a particular account will model the whole set of affordances that may meet any given task, given user, or given experience, the aggregate sum of these will begin to approximate it. It may be useful to compare this to general statistical sampling, where the population remains unknown and we imagine that, given a large enough sampling of a random group, we can make a claim about the population in general. This is a useful if somewhat misleading comparison. As the sense of place is rooted in human experience, it is difficult to say that there is some objective 'true' set of affordances. It is more appropriate to say that there are intended or general affordances that are difficult to measure directly. These are those affordances more like the original description posed by Gibson. A capability for action may exist in general, even if it is not utilized in practice.

The general sets of affordances are these base capabilities for action supported by a place, removed of situational perspective and demands. A particular account will always fall within this model. We call this the general affordance plane (see figure) and represent it as a multi-edged plane where each edge represents a particular affordance and objective measure of affordability compared to some quantifiable maximum. A particular experiential affordance plane contains the set of all user-observed affordances, and the associated degree. The aggregate of ongoing experiential accounts creates a perceived affordance plane that, again, must fall within the confines of the general plane. Points represent discrete activities, such as eating and drinking, and are ideally organized closest with other points similar in meaning. The particular construction of an affordance plane depends on the intended use of the model, and semantic overlap makes it difficult for this particular visualization to represent effectively all possible activities. In practice, this is not a problem.

While only a few accounts are insufficient to make a general claim, when considering a few hundred accounts the perceived affordance plane quickly stabilizes into a statistically significant model, where the impact of additional input quickly approaches asymptotic limits.

Affordance planes represent an interesting visualization; however we can also simply output a textual description of the relevant information. A partial listing (of some interesting affordances) for the MIT Media Lab appears right in tag cloud format listed action / and affordance rank—shown relatively. This account can be broken up to aggregate experiential accounts, or even individual ones and their sources.

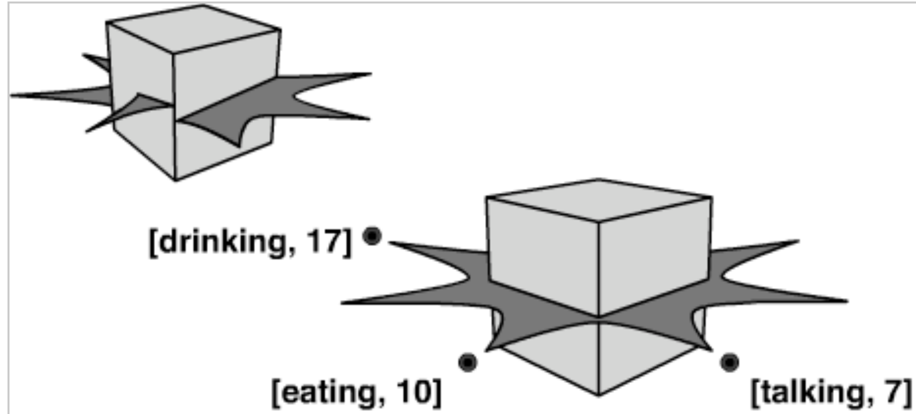

Figure 18. *Simple model of affordance planes in placial objects. Here each affordance specifies an afforded action and degree of affordance.*

The Media Laboratory, MIT
20 Ames St. Cambridge, MA 02139

Demo (6.7) / Publish (3.2) / Educate (5.6) / Communicate (7.3) / Get Press (8.9) / Express Oneself (5.9) / Research (9.1) / Make Agents (1.3) / Do Web Stuff (3.2) / Design (3.0) / Perish (0.8) / Transcend (7.8) / Die (0.2) / Make Robots (1.2) / Make AI (3.7) / Grow (4.1) / Adapt (7.9)

9.2 Cultural Exclusion

As Harrison and Dourish point out, "a conference hall and a theatre share many similar spatial features (such as lighting and orientation); and yet we rarely sing or dance when presenting conference papers, and to do so would be regarded as at least slightly odd (or would need to be explained)" (Harrison 1996). This behavior is not necessarily out of space but it is clearly out of place; not only does this prescribe behavior, it may also modify our perception of the place (we may interpret singing and dancing as part of the presentation, for example).

"It is a sense of place, not space, which makes it appropriate to dance at a concert, but not at a Cambridge college high table; to be naked in the bedroom, but not in the street; and to sit at our windows peering out, rather than at other people's windows peering in. Place, not space, frames appropriate behavior. (Harrison, 1996).

This sentiment suggests a distinction between what may be technically afforded by a particular space, and what affordances are perceived as aspects of a place. In this model we identify this discrepancy as "cultural exclusion"; the segment between the intended affordance of a space and the perceived affordance of a place.

The degree of cultural exclusion directly translates into the idea of 'fake' artificial or placeless places. While some cultural exclusion is to be expected, when the cultural exclusion expands the utility of the place decreases.

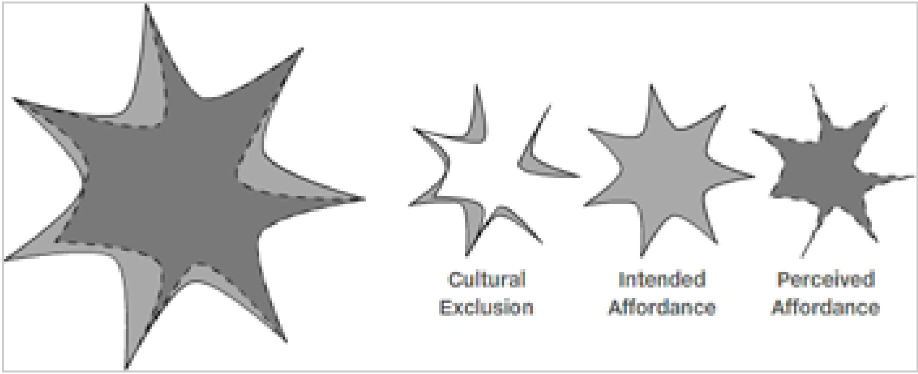

Figure 19. *Diagram showing cultural exclusion as the difference of intended and perceived affordance planes.*

Possible Applications

One can easily imagine future systems built upon this model and some of the techniques, employing a variety of visualizations.

Some possible applications of this model could include:

- ➔ Applications that utilize place knowledge and contextual awareness to transform between map representations and linear suggestions. Knowing you are at a restaurant, and employing common-sense systems, allows the suggestions of possible future branches—perhaps coffee, perhaps drinks, perhaps dancing.
- ➔ The limitations of relying on pure data entry are removed. It is possible to show the 'bad' sections of the city without relying on police databases, or to show the boring sections of the city and the exciting ones.
- ➔ Search and suggestive capabilities increase, allowing events to be created based on interactions between who a person is, what they like to do, and what capabilities are afforded by a place. Imagine event networks that create "go hiking" events without someone needing to construct such an event.

10. Concluding Remarks

This thesis has framed two years of exploration into the representation, visualization, and interpretation of space and place, and how they relate to the average human being. Early work in the visualization of spatial representations that are more closely modeled on human understanding established the utility of these representations and the resultant exploration into machine-readable but human interpretable place-sense provides a mechanism for placial representation through computation, through which these representations can be constructed.

Space and place are central to human experience. This centricity demands that computer systems not only understand space, but place as well. Place is created by use, and this use can be derived by observation of textual description and implicit support in applications. This understanding, this model, can construct visualizations that present human conceptions of place.

The model presented is a simple representation intended for practical usage. The presented CampusMap application provides light application directions that could enable any spatial application to gather spatial knowledge. The early representation exploration provides direction for this knowledge. Having created a system that understands place allows for the successful construction of these representations.

The presented model inherits the philosophical tradition of a grounded, active, and very human conception of place construction. To that end, we embrace a model of place grounded in the active human conception of affordance written in the language of the machine.

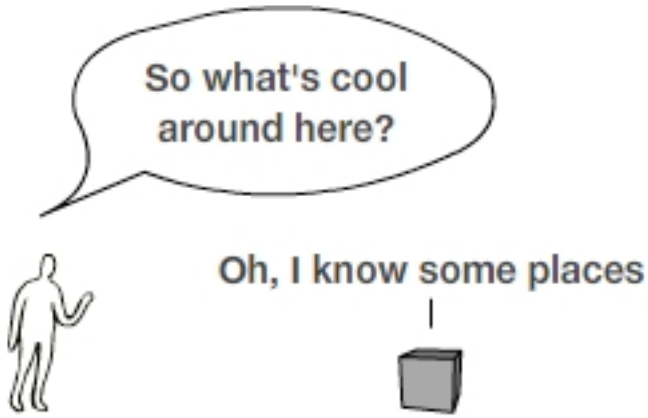

11. References

1. Aggarwal, C.C., Yu, P.S. *Data Mining Techniques for Personalization*. IEEE Data Engineering Bulletin. (2000).
2. Agrawala, M., Stolte, C. *Rendering Effective Route Maps: Improving Usability Through Generalization*. International Conference on Computer Graphics and Interactive Techniques. (2001).
3. Aime, A., Bonfati, F., Monari, P.D. *Making GIS closer to end users of urban environment data*. ACM GIS'99. (1999).
4. Alexander, C., Ishikawa, S., Silverstein, M. *A Pattern Language*. New York: Oxford University Press. (1977).
5. Bleecker, J. *A Design Approach for the Geospatial Web*. Where 2.0. (2005).
6. Bouquet, P., Warglien, M. *Mental models and local models semantics: the problem of information integration*. European Conference on Cognitive Science. (1999).
7. Bourdieu, P. *Outline of a Theory of Practice*. (1977)
8. Bugayevskiy, L.M., Snyder, J.P. *Map Projections*. London: Taylor & Francis. (1995).
9. Burrough, P. A. and Frank, A. U. (Eds.*) Geographic Objects with Indeterminate Boundaries*. London: Taylor and Francis. (1996).
10. Burrough, P.A.; van Gaans, P. F. M.; Hootsmans, R. *Continuous classification in soil survey: spatial correlation, confusion and boundaries*. Geoderma. (1996).
11. Camara, G., Souza, R.C.M., Freitas, U.M. Garrido, J. *Spring: integrating remote sensing and GIS by object-oriented data modeling*. Computers & Graphics. (1996).
12. Chen, G., Kotz, *A Survey of Context-Aware Mobile Computing Research*. Dartmouth Computer Science Technical Report. (2000).
13. Cheverst, K., Davies, N., Mitchell, K., Friday, A. *Experiences of developing and deploying a context-aware tourist guide: the GUIDE project*. In Proc. International Conference on Mobile Computing and Networking. (2000a).
14. Cheverst, K., Davies, N., Mitchell, K., Smith, P. *Providing Tailored (Context-Aware) Information to City Visitors*. In Proc. Adaptive Hypermedia and Adaptive Web-Based Systems. (2000b).
15. Churcher, N. *Applications of Distortion-Oriented Presentation Techniques in GIS*. In Proc. New Zealand Conference on Geographical Information Systems and Spatial Information System Research. (1995).
16. Cohn, A.G. *Qualitative spatial representation and reasoning techniques*. Advances in Artificial Intelligence. (1997).
17. Couclelis, H., Golledge, R. G., Gale, N., and Tobler, W. *Exploring the Anchor Point Hypothesis of Spatial Cognition*. Journal of Experimental Psychology. (1987).
18. Curry, M.R. *The Work in the World - Geographic Practice and the Written Word*. (1996).
19. Cresswell, T. *In Place / Out of Place*. (1996).
20. Dudek, G., Jenkin, M.R.M., Milios, E.E., Wilkes, D. *Map Validation and Self-location in a Graph-like World*. International Joint Conferences on Artificial Intelligence. (1993).
21. Egenhofer, M., Burns, H. *Visual Map Algebra: a direct-manipulation user interface for GIS*. In Proc. Working Conference on Visual Database Systems. (1995).
22. Egenhofer, M., Mark, D. *Naïve Geography*. In Frank, A.U. and Kuhn, W., (Eds.) Spatial Information Theory: A Theoretical Basis for GIS. Lecture Notes in Computer Sciences (1995).
23. Egenhofer, M., Frank, A.U. *Object-Oriented Modeling for GIS*. URISA Journal. (1992).
24. Egenhofer, M., Freksa, C., Miller, H.J. (Eds.) *Formalizing User Actions for Ontologies*. (2004).
25. Elias, N. *The Civilizing Process*. The History of Manners. (1978).
26. Entrikin, J. *The Betweenness of Place - Towards a Geography of Modernity*. (1991).

27. ESRI. *The GIS Software Leader*. http://www.esri.com (2006).
28. Foresman, T. *The History of GIS (Geographic Information Systems)*. Prentice Hall Series in Geographic Information Science. (1997).
29. Furnas, G.W., Morristown, N.J. *Generalized fisheye views*. In Proc. SIGCHI conference on Human Factors in Computing Systems. (1986).
30. Giddens, A. *The Constitution of Society*. (1984).
31. Goguen, J. *An Introduction to Algebraic Semiotics, with Applications to User Interface Design*. In Nehaniv, C. (Ed.) Computation for Metaphors, Analogy and Agents. Lecture Notes in Artificial Intelligence. (1999).
32. Goffman, E. *The Presentations of Self in Everyday Life*. (1959).
33. Google. *Google Maps*. http://maps.google.com (2006).
34. Google Maps Mania. http://googlemapsmania.blogspot.com (2005).
35. Hackos, J.A.T., Redish, J.C. *User and Task Analysis for Interface Design*. IEEE Transactions on Professional Communications. (1999).
36. Hardt, S. *Naïve Physics*. In Shapiro, S. (Ed) Encyclopedia of Artificial Intelligence. (1992).
37. Harley, J.B., Laxton, P., Andrews, H.J. *The New Nature of Maps: Essays in the History of Cartography*. (2002).
38. Harrison, S., Dourish, P. *Re-Place-ing Space: The Roles of Place and Space in Collaborative Systems*. In Proc. ACM Conference on Computer Supported Cooperative Work. (1996).
39. Hayes, P. *Naïve Physics Manifesto*. In Michie, D. (Ed.) Expert Systems in the Microelectronic Age. (1978).
40. Hayward, W.G., Tarr, M.J., Wallace, W.P., Stewart H.L. *Spatial language and spatial representation*. Cognition 55. (1995).
41. Heidegger, M. *Poetry, Language, Thought*. Translated by Hofstadter, A. (1971).
42. Hernandez, D. *Qualitative Representation of Spatial Knowledge*. Lecture Notes in Computer Science. (1994).
43. Hockenberry, M., Hoff, J., Selker, T. *A Mindset for User-Centered Spatial Applications*. Intl. Conference on Spatial Cognition. (2006).
44. Hockenberry, M., Selker, T. *A Sense of Spatial Semantics*. Conference on Computer Human Interaction (2006).
45. Hockenberry, M., Gens, R. *PlaceMap: Building Community through Active Context Mapping*. Siggraph. (2005)
46. Horowitz, B. *Creators, Synthesizers, and Consumers*. Elatable.com / Yahoo Inc. (2006).
47. Janecek, P., Pu, P. *A Framework for Designing Fisheye Views to Support Multiple Semantic Contexts*. Conference on Advanced Visual Interfaces. (2002).
48. Janecek, P., Pu, P. *Visual interfaces for opportunistic information seeking*. In Proc. of the International Conference on Human Computer Interaction. (2003).
49. Jankowski, P., Andrienko, N., Andrienko, G. *Map-Centered Exploratory Approach to Multiple Criteria Spatial Decision Making*. International Journal Geographical Information Science. (2001).
50. Jordan, T., Raubal, M., Gartrell, B., Egenhofer, M. *An Affordance-Based Model of Place in GIS*. In Proc. International Symposium on Spatial Data Handling. (1998).
51. Kahana, M. *The Neurophysiology of Human Spatial Navigation*. Psychology Science Convention. (2004).
52. Kidner, D., Jones C. *A Deductive Object-Oriented GIS for Handling Multiple Representations*. In Proc. International Symposium on Spatial Data Handling. (1994).
53. Knez, I. *Attachment and identity as related to a place and its perceived climate*. Journal of environmental psychology. (2005).

54. Kramer, J., Noronha, S., Vergo, J. *A user-centered design approach to personalization*. Communications of the ACM. (2000).
55. Koffka, K. *Principles of Gestalt Psychology*. (1935).
56. Kuhn, W. *Defining semantics for spatial data transfers*. Proceedings of the International Symposium on Spatial Data Handling. (1994).
57. Kuhn, W. *Ontologies in support of activities in geographical space*. International Journal of Geographical Information Science. (2001).
58. Kuipers, B. *Modeling Spatial Knowledge*. Cognitive Science. (1978).
59. Kuipers, B., Levitt, T. *Navigation and Mapping in Large Scale Space*. AI Magazine. (1988).
60. Lakoff, G., Johnson, M. *Metaphors We Live By*. (1980).
61. Lanter, D.P., Essenger, R. *User-Centered Graphical User Interface Design for GIS*. Technical Report for the National Center for Geographic Information and Analysis. (1991).
62. Leung, Y.K., Apperley, M.D. *A review and taxonomy of distortion-oriented presentation techniques*. ACM Transactions on Computer-Human Interaction. (1994).
63. Liu, H., Singh, P. *ConceptNet: a practical commonsense reasoning toolkit*. BT Technology Journal. (2004).
64. Liu, H. *Computing Point-of-View: Modeling and Simulating Judgments of Taste*. (2006).
65. Lieberman, H. *Powers of Ten Thousand: Navigating In Large Information Spaces*. Conference on User Interface Software Technology. (1994).
66. Lurie, G., Sydelko, P., Taxon, T. *An Object-Oriented GIS Toolkit for Web-Based and Dynamic Decision Analysis Applications*. Journal of Geographic Information and Decision Analysis. (2002).
67. Lynch, K. *The Image of the City*. (1960).
68. MacEachren, A.M. *How Maps Work: Representation, Visualization, and Design*. (2004).
69. Mark, D. *Spatial Metaphors for Human-Computer Interaction*. International symposium on spatial data handling. (1992).
70. Mark, D. *Counter-Intuitive Geographic "Facts:" Clues for Spatial Reasoning at Geographic Scales*. In Frank, A. Campari, I., Fiormentini, U. (Eds.) *Theories and methods of spatio-temporal reasoning in geographic space*. Lecture notes in Computer Science. (1992b).
71. Mark, D., Comas, D., Egenhofer, M., Freundshuh, S. Gould, M., Nunes, J. *Evaluating and Refining Computational Models of Spatial Relations Through Cross-Linguistic Human-Subjects Testing*. COSIT. Lecture Notes in Computer Science. (1995).
72. Markosian, L., Kowalski, M.A., Goldstein, D., Trychin, S.J. *Real-time nonphotorealistic rendering*. In Proc. Computer graphics and interactive techniques. (1997).
73. Maurer, D. *Four Modes of Seeking Information and How to Design for Them*. boxesandarrows.com (2006).
74. Mauss, M. *Les Techniques du corps*. Journal de Psychologie. Reprinted in Mauss, Sociologie et anthropologie. (1934).
75. Miles, S.B., Ho, C.L. *Applications and Issues of GIS as Tool for Civil Engineering Modeling*. Journal of Computing in Civil Engineering. (1999).
76. Montello, D. *Scale and Multiple Psychologies of Space*. Spatial Information Theory: A Theoretical Basis for GIS. Lecture Notes in Computer Science. (1993).
77. Morville, P. Rosenfeld, L. *Information Architecture for the World Wide Web: Designing Large-Scale Web Sites*. (2002).
78. Neisser, U. *Five Kinds of Self-Knowledge*. *Philosophical Psychology*. (1998).
79. Norman, D.A., Draper, S.W. *User Centered System Design; New Perspectives on Human-Computer Interaction*. Lawrence Erlbaum Associates. (1986).

80. Open Guides. *OpenGuides: the guides made by you*. http://openguides.org (2006).
81. Pederson, E. *Geographic and Manipulable Space in Two Tamil Linguistic Systems*. Spatial Information Theory: A Theoretical Basis for GIS. Lecture Notes in Computer Science. (1993).
82. Peng, C. *E-Service as a New Paradigm for Interactive Multidimensional City Modeling*. E-Technology, e-Commerce and e-Service. (2004).
83. Prado, A.B., Baranauskas, M.C.C., Medeiros, C.M.B. *Cartography and Geographic Information Systems as Semiotic Systems: A Comparative Analysis*. ACM Symposium on GIS. (2000).
84. Prasolova-Forland, E. *Supporting Social Awareness in Education in Collaborative Virtual Environments*. (2002).
85. Relph, E. *Place and Placelessness*. (1976).
86. Rivest, S., Bedard, Y., Marchand, P. *Towards better support for spatial decision making: Defining the characteristics of Spatial On-Line*. Geomatica. (2001).
87. Robertson, G., Czerwinski, M., Larson, K., Robbins, D.C., Thiel, D., van Dantzich, M. Data Mountain: using spatial memory for document management. In Proc. Symposium on User Interface Software and Technology. (1998).
88. Rodden, T., Benford, S. *The evolution of buildings and implications for the design of ubiquitous domestic environments*. In Proc. CHI Conference on Human Factors in Computing Systems. (2003).
89. Rumbaugh, J., Blaha, M., Premerlani, W., Eddy, F. *Object-oriented modeling and design*. (1991).
90. Sarker, M., Brown, M.H. *Graphical fisheye views*. Communications of the ACM (1994).
91. Schuurman, N. *Trouble in the heartland: GIS and its critics in the 1990s*. Progress in Human Geography. (2000).
92. Selker, T. Style and Function of Graphic Tools. Graphics. (1999).
93. Shneiderman, B. *Designing the User Interface: Strategies for Effective Human-Computer Interaction*. (1997).
94. Siegel, A.W., White, S.H. *The development of spatial representations of large-scale environments*. Advances in child development and behavior. (1975).
95. Silverman, K. *The Subject of Semiotics*. (1983).
96. Singh, P. *The Open Mind Common Sense Project*. KurzweilAI.net. (2002).
97. Skyhook. *Skyhook Wireless*. http://www.skyhookwireless.com (2006).
98. Smith, B. *The formal ontology of space: An essay in mereotopology*. In Hahn, L. (Ed.) The Philosophy of Roderick Chisholm. (1994).
99. Stevens, A., Coupe, P. *Distortions in judged spatial relations*. Cognitive Psychology. (1978).
100. Therborn, G. *The Ideology of Power and the Power of Ideology*. (1980).
101. Tollmar, K., Sandor, O., Schömer, A. *Supporting social awareness work design and experience*. (1995).
102. Traynor, C., Williams, M.G. *Why are geographic information systems hard to use?* In Conference Companion of Human Factors in Computing Systems Conference. (1995).
103. Tuan, Y. *Space and Place*. The Perspective of Experience. (1977).
104. Vertesi, J. *Mind The Gap: The 'Tube Map' as London's User Interface*. hciresearch.org. (2005).
105. Virrantaus, K. Markkula, J., Garmash, A., Terziyan, V. *Developing GIS-supported location-based services*. Web Information Systems Engineering. (2001).
106. Waquant, Loic. *An Invitation to Reflexive Sociology* (with Pierre Bourdieu). (1992).

107. Werner, S., Krieg-Bruckner, B., Mallot, H.A., Schweizer, K., Freksa, C. *Spatial Cognition: The Role of Landmark, Route and Survey*. Informatik. (1997).
108. Whyte, W. *City: Rediscovering the Center*. (1988).
109. Worboys, M., Duckham, M. *GIS: A Computing Perspective*. (2004).
110. Zipf, A. *User-Adaptive Maps for Location-Based Services (LBS) for Tourism*. In Proc. ENTER. (2002).
111. Zubin, D. *Natural Language Understanding and Reference Frames*. In Mark, D., Frank, A., Egenhofer, M., Freundschuh, M., McGranaghan, M., White, R.M. (Eds.) Languages of Spatial Relations: Initiative Two Specialist Meeting Report. National Center for Geographic Information Analysis. (1989).