

# *Discovering The Community Of Our Future*

## **Decision Support and Digital Library tools supporting Informed Community Decision- Making in the Georgia Basin Region**

**Rob Harrap, Queen's University  
Murray Journeay & Sonia Talwar, Geological Survey of Canada  
Boyan Brodaric, Penn State University**

---

People want informa-  
tion *now*, but what they  
often get is information  
that has no context,  
and so leads to flawed  
decisions.

### *Introduction*

---

We live in an age of instant information gratification, where increasingly the public views access to information as a key to their ability to understand the ever changing world. But as has been pointed out repeatedly, there is no guarantee that the quality of information is sufficient in theory or practice to guide decision making, and that the data as presented can be understood without context established by formal education. Most critically, there is a complete naivete on information consumers' part, brought on by lack of formal training in part and the desire for immediate answers more crucially, that one 'factoid' found using an ad hoc search of a dubious global opinion database (aka the Internet) is itself sufficient to make an *informed* decision, all other information being irrelevant.

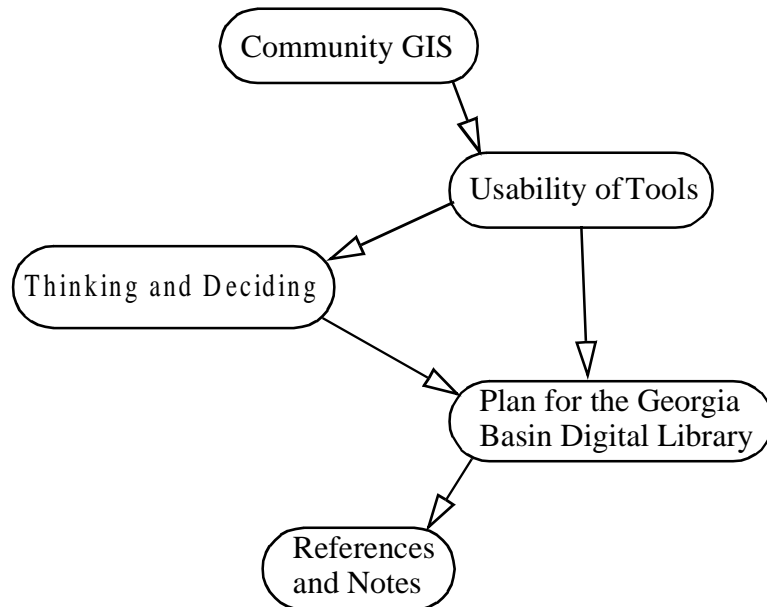
The world, even when isolated in a laboratory setting, can rarely be characterized by a single factor of importance, and even if one factor dominates a context, the bounds of that context are not obvious. Quick and painless information access leads to naive and ultimately painful, shortsited decisions. The education community is running scared before the onslaught of web-based free information and training materials, but these as currently in existence are in practice mostly disinformation. We can only hope that the public will soon realize this and demand information access tools of higher quality, better context, and more balanced presentation of ideas.

With regard to decision making in an urban environment, the different contexts that knowledge relevant for planning falls into are staggering, but there are common elements that provide a stage for a more global meta-context. These are, at a modest scale, maps and familiar (metaphorical) tools for information access such as newspapers and television documentaries. At a less modest scale, these could include virtual reality presentations of the urban setting, video game like modeling of changes in this environ-

ment, and a semi-intelligent agent-based help system that tracks access to information, spots common problems, and suggests topics for further investigation. But a question arises: exactly how do we come up with the underlying knowledge for such a system? Even the simplest map has deep contextual meaning, and whereas the meaning of expertly constructed maps is well constrained (though not necessarily well communicated), the overarching context is poorly constrained and cannot in practice be constrained by any small group of domain experts. The overall context is *what the community wants*, which is an instance of the general subject *what people want*, and attempts to characterize this anthropological framework for urban design cannot arise from theory, they must arise from past practice. The historical context guides us, as we back up into the future.

The purpose of this document is to present some ideas about how a specific community planning and decision making toolset might be built for the Georgia Basin area of Southwestern British Columbia. This toolset is novel only in that it proposes that several different approaches to information access are not separate but in fact can be readily integrated, with significant benefits arising first from their integration, and second from the fact that a digital integration of these tools (some of which are traditionally paper-based) will allow entirely new means of understanding the context of factors in urban decision making. We take the viewpoint here that the grounding for the use of such a system should be a high school education at the most.

The material hereafter is organized into modules that can be read in any order. An overview map (which is active as hyperlinks in the online version of this document) provides some navigation. If in doubt, read the document in the order presented.



## Community GIS

---

*Jennie picks up her assignment for Social Studies and reads the first question: What are the problems with building large 'Box Store' shopping centres, both to the environment and to transportation flow in cities? She thinks about this for a moment, recalling a recent trip to Chapters. Turning to her computer, she fires up a web browser and hits the hotkey icon for Search. She runs a quick search on the Internet using the key words Box Store Traffic. She spends the next hour tracking down enough information to fill a page. There is a lot of information, but she has a hard time seeing how it is all related. Most of sites, unfortunately, are trying to sell her something or are details about one particular aspect of the problem. At least she got something. Her friend Ann, who doesn't have a computer at home, must really be struggling...*

Where does the information come from?  
What assumptions are we making about who will access information and how they will do so?

Any tools we build must reach multiple communities if they are to bring communities together in an informed decision making process.

What exactly is the community, anyway? We need some sense of this if we are going to address their needs. The assumption, somehow, is that the community is everyone who is interested in their surroundings, but in fact this isn't accurate since it misses several unspoken assumptions about information and access to information, about knowledge and tools, and about participation in process.

First, consider the assumption of access to information. The widespread assumption that has arisen over the last five years is that everyone who is interested in learning and access to information:

1. Has a computer
2. Has access to the World Wide Web
3. Will spend hours using these to learn something about the background to issues of community planning.

We propose herein a web-based solution to access to information, one with substantial improvements to usability and so on. Nevertheless, we will miss a significant portion of the community, disproportionately the elderly and the socioeconomically disenfranchised, with such a system. The system we develop might then include some capacity to have a subset of the information and tools available in printed form, or at the very least amenable to presentation in information kiosks at museums, libraries, city hall, and so on.

Second, consider the assumption about knowledge and tools. The assumption that follows from the 'has-computer, will browse' direction of development also comes with an assumption of minimum education, access to hardware and software, and feeling of empowerment. As indicated, we can partially address the access issue through printed versions or public access sites. The feeling of empowerment is more difficult, but the very act of involving the community in decision making contributes to this, and many parallel public relations efforts can work in conjunction with a digital planning tool to partially address these concerns. The education level issue is more difficult. And of course this is related to the other concerns as well (education and ownership of comput-

ers and feelings of empowerment are all related...). The education issue leads us to a general principle or goal.

**The digital library toolset must accommodate learning at all scales: learning principles of analysis, learning about the data that exists, and learning about the significance and context of the data and principles.**

Building tools that are available in multiple forms and venues, that are self-documenting and built to educate at a wide range of background competencies, and that promote community feelings of involvement represents the context of what follows.

---

### *Usability of Tools*

---

*Jennie has pretty much got it together for the 'Box Store' question. The next question is worse: How is crime and building density related? She tries another search, but most of the results are kind of wierd. Some are about Jamaica. But there is one about a thing called a GIS, which is a tool for understanding how things are arranged, as far as she can tell. She makes a note to herself to ask her teacher about this. She thinks that the computer lab at school has some GIS stuff, but it is probably way too complicated for her.*

Geographic Information Systems (GIS) provides a tool building framework for a community decision support system. But these are often poorly designed.

Usability of tools must improve if they are to reach more members of the community.

Clarity of data and operations must be improved if the community is to understand what the tool is telling them.

The setting for decision making in a community is explicitly spatial. While the information needed might include texts, video, audio, and so on, the common element of these is a geographic region, and the issue of spatial distribution and context of elements of the information cannot be ignored. While information about some topics might be a natural fit for storage in a conventional database or as Web pages, spatial information needs to be managed in a setting that preserves this spatial context. Geographic Information Systems technology (GIS) includes a range of tools and methods that not only preserve spatial context, but include modeling operations that make use of spatial configuration, such as proximity and adjacency, while also offering supporting more traditional presentations. GIS tools have been around for about 20 years, with a large increase of availability and application in the mid 1990's as software was moved from workstation settings to personal computers and even handheld devices.

Current academic research in GIS focuses on issues that are central to the development of new modeling techniques, improved workflow, and more rigorous data models. These have very little impact on the consumers of the decision support process: managers and the public, who either through direct (usually non-expert) use or exposure to the products of GIS need a different set of problems to be addressed: usability of tools, and clarity of meaning. The issue of usability is currently a hot topic with application vendors who see it as a means to extending the audience for their software; clarity of meaning on the other hand has received little attention to date in the GIS community.

Usability refers to the ease of use of tools, the ease of learning new techniques, and the basis of these on consistency in user interface design, workflow design, and conformance to community standards of these. Many interface 'features,' for example, are

demonstrably confusing and yet standard in current generation windowing tools. Usability looks to balance the use of these standard interface and workflow elements with the need for tools to be both elegant in design and use. Too often, this design process is instead redesign: broken products are submitted for corrective surgery in a post-coding attempt at repair.

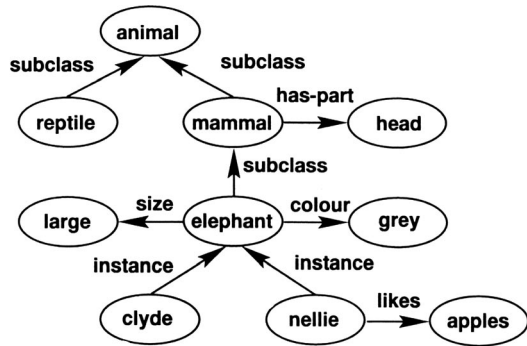
GIS tools are by no means usable. In fact, they are prime offenders, mixing interface metaphors from pre-windowing interface days with ill-conceived interface elements and attempts at flashy interface hacks that are intended to somehow improve usability. And tools that are redesigned from the interface perspective - Arc/Info 8 being a prime example - are such a dramatic improvement that the cost of total recoding from scratch to fix legacy issues is justified. The sad fact that large parts of rewritten code is simply in place to maintain compatibility with old data models, macro language code, and hybrid tools keeps rewrites from being a total success, but nonetheless the efforts are well justified. Web-based GIS tools are, in general, more usable than their precursors because they are explicitly designed for non-expert use and have for the most part post-dated both widespread community standards in GUI design and to some extent standards in Web usability. Unfortunately they often suffer from overly 'cool' interfaces.

Clarity of meaning has another name: semantics. Many GIS tools, even those with improved or rewritten and now-usable interfaces, are unclear precisely because the semantics of the system and especially of the data are imprecise or downright inconsistent. Semantics refers to meaning: it is the counterpart to syntax, which is adherence to grammatical rules. A GIS dataset that 'works' and adheres to the requirements of the tool it is used with could be said to be syntactically correct. But if there was no clear sense of what the data was or how it related to other data (and the 'real world') then to a novice user of that data, there is no clarity and no sense of how things connect and so can be used meaningfully. The problem only gets worse for the semantics of the tools themselves - every tool has underlying assumptions of proper use, and these interact with data in a complex way that largely determines whether a process will produce garbage or useful results.

This issue has been addressed on two fronts: metadata and community data models. Adherence to community data models (such as the North American Data Model for geologic data, or the recent group of data models developed by groups in collaboration with ESRI) provides some chance at providing system level metadata about domain specific tools and their proper place. Metadata provides a place for rigorous documentation about data and proper use, albeit in a rather dry way.

As with many other aspects of computer applications that have come into and out of vogue over the last two decades (GUI's and Object Oriented Programming being prime examples) the issue of semantics has been addressed in detail in the black sheep discipline of computer science, artificial intelligence research. And AI researchers have made considerable, though often contradictory, inroads into semantics and formal ontologies (catalogues of what can be, and how things that can be relate). Some of these advances need to be brought out of the research lab and into widespread application, and GIS, and community GIS in particular, is a prime area to do exactly this.

One of the mainstays of knowledge representation, the AI subdiscipline that looks at how we can and should demonstrate relations in knowledge, is the semantic network, a



**Figure 1:** A simple semantic network with concepts and instances.

Note that this corresponds to two levels of discussion, or a language and metalanguage. One level refers to actual things or concepts, and one level refers to their relations. In practice semantic nets, or ontologies, operate on multiple scales of discussion. This imposes a certain brittleness on their automated application. Nevertheless they are a very useful method of semi-formal knowledge representation.

simple graph drawing that shows concepts linked by (often labeled) lines. A properly labeled semantic network can efficiently show the semantic relations in concepts, and in data as well, and has the additional advantage of being transparent - a simple semantic network takes no training to use, is its own interface metaphor, and is trivial to implement in a GUI.

The key issue in community GIS is the relation between concepts (whether they be data concepts or modeling concepts) and how these interact. Without some knowledge of tools and the domain details of a dataset (the geological details behind a map, both in data terms and in geologic theory terms) it is impossible to use that data and those tools in any meaningful way. The response is to include voluminous help files, or pointers to literature, that in general is organized for expert, not community, use.

Underlying any domain, though, lies the semantics of the domain, and these can be represented by semantics networks, which as alluded to above, are conceptually simple and easily implemented in a GUI. Behind the semantics - the network of concepts - lie the details of these concepts. If some level of granularity is chosen for both the concepts in the network and the details behind them, it is quite simple to build a combined 'textbook' and semantic network browser that not only provides details on data and concepts, but shows how things are connected, which is totally lacking in textbook forms alone.

Such a system simultaneously provides information and context, reminding users what relationships exist within and between items of interest.

The question of granularity is a question of audience. How much description is needed? If we can assume an undergraduate education in something, then we can use quite sophisticated concepts compared to say, a public school education. The problem is actually open ended. Even a high level semantic network for experts, and high level detail descriptions, expands at some level to being a description of all of knowledge, since at some level all things are related. For example, looking at the details of a mineral might lead to chemistry, to physics, to instrumentation, to funding of technology development, to politics, and onward. Bounding the system is a fuzzy concept, because extra-domain concepts are needed more by novices, and yet require much finer granularity for novices.

A system that, having addressed the granularity problem, provides a combined semantic representation for knowledge and details, also needs a GIS display of some sort. So we have a tri-element display: GIS view, Concept View, and Concept Details. All of these ideally should interact at three levels: the concept of domain level, the concept of data element level, and the concept of the tool itself level. So a geological browser organized in this fashion would have semantics networks and details for relevant geology concepts, represented geologic data, and geologic GIS methods.

Of these, only the represented data information is unique to a project, as geology as a subject, while evolving, has consistent semantics (inconsistencies in geologic theory are equally inconsistent regardless of field area, so the domain of semantics is consistent, even if the applied semantics are generally problematic) and the tools are perhaps even less of a problem (depending on attention to usability). And given that to a large degree the issues in the semantics of represented geologic data are addressed in a data model such as the NADM, we can in principle subsume the semantics of a larger community data model into a tripartite GIS browser and so further diminish the uniqueness of a particular data sets' expressed semantics and details.

To summarize, an ideal community GIS might be a system that provides three windows into a dataset relevant to that community: a map viewer showing the area in cartographic form, a Web-like hypertext viewer showing supporting information such as reports and pictures, and a visual network browser showing the interrelations between map specific concepts, map elements, and extant scientific theory. This system promotes a deep understanding of context: spatial, inter-subject, and intra-subject.

If we intend to build tools that reach a broad audience and build long term understanding of issues, we need to understand how people learn and make decisions.

Procedural instructions are a poor substitute for intuitive understanding and a theoretic framework for building mental models. How do we turn procedural instructions into deep, theory-constrained understanding?

---

## Thinking and Deciding

*Jennie is still stuck on the crime and buildings question. The Internet sources are way too complicated for her to really understand. Then she remembers something about this from class. No, it was from her brother's video game, SimCity. But she has a basic idea. She digs up the CD for SimCity and after a few hours of play, has a basic understanding of how things fit together, at least in the game. Why can't learning be more like this at school? Then she realizes she has a problem - her teacher doesn't want to hear about SimCity. She has to somehow extract some fancy sounding words from her gut feelings about the game. Yeah, right.*

What exactly is learning? What exactly is knowledge? How do humans make decisions? If we are interested in helping communities use GIS data in particular, or knowledge in general, then it is useful to take a quick look at exactly what knowledge is. This is a contentious topic, but even a crude attempt at characterizing what knowledge is and decisions are may shed some light on aspects of the problem of a community decision support tool design.

As a first step, consider three kinds of knowledge. One is knowledge gained through experience in the natural environment of a skill - on the job knowledge. One is theoretic knowledge of that environment - what some would call academic understanding, or

abstract knowledge. One is procedural knowledge for carrying out tasks in that environment - checklists for novices (and experts?) to use when faced with a problem. At first glance one might rank these in terms of effectiveness as procedural, then experience, then abstract knowledge, especially if one was an academic trained to think in abstract terms. In fact (as public opinion about academics usually suggests) on the job knowledge is far more useful in dealing with situations that are unbounded or urgent than either procedural or abstract knowledge. Snap decisions made by experienced people using intuition are far more likely to be accurate than prescriptive or theoretic knowledge. We need to understand why this is, as it strongly affects how we empower the community with a decision support system.

There is no doubt whatsoever that theoretic knowledge is far superior to on-the-job knowledge in dealing with new situations, and in explaining deep causes for phenomena. If this weren't the case, no one would bother with the painful abstraction of general principles from specific cases. And procedures derived from theory and experience are very useful in cases where the person faced with a decision has neither on-the-job knowledge nor theoretic grounding. These constitute the prescriptive or procedural knowledge that fares poorly, except in comparison to no knowledge at all.

In community decision making this leads us to a general statement and a general problem: it would be nice to have decision makers gain on-the-job like intuition, but given the difficulty of this, a mix of procedures and some theoretic background are a good second best choice. And the second best choice is not only likely, but inevitable, because no one is an expert, or on-the-job experienced, in all aspects of the problems facing communities. If we wish to empower communities to buy-in to decision making processes, some method of giving a combination of theory and procedures, and some measure of on-the-job simulation, are a powerful combination. In other words, we need to deliver prescriptive 'recipes', access to theory, and access to data, and some form of transparent representation mechanism like semantic nets must be present to keep the context grounded. But these must be geared in such a way as to promote building on-the-job like deep understanding.

The internalized understanding that is one of our goals is not completely unconscious and unknowable. People use a complex weave of knowledge structures that blend symbolic ways of knowing with preconscious, often perception-dominated ways of knowing. The resulting recipes for problem solving, for tasks as simple as opening a door or as complex as judging intention in a critical conversation, have some elements in common with procedural knowledge, and certainly build on a theoretic knowledge in some cases, but are nevertheless unconscious once internalized. There is a great deal of research into these structures, called image schemas, which show up in everyday language (such as I am now using) in the form of metaphorical constructions. Cognition, schemas, and computer models of cognition are discussed in detail in an appendix.

One of the lessons from both the theoretic study of schemas and cognition, and from the more practical reality of the metaphorical grounding of everyday language, is that both understanding and communication are built on a fabric of learned constructions, and that these constructions are story-like in their nature and simplicity. At the simplest level, they describe operations of everyday life. At a higher level, they relate to patterns of behaviour and the recognition of junctures in situations. The crucial point here is that communication of the understanding of the complex, real world is handled by stories

much more easily than by abstractions. This point is reflected in history. Humans had stories long before anyone bothered with logic. One definition of a community would be those who share life stories.

If image schema can be communicated in some fashion, and if they underlie human understanding in general, then they form a powerful counterpart to theoretic understanding. Moreover, the recognition of these schema, even if we can only partially reveal them at work, provides a theoretic metastructure for understanding deep learning in general. Research on this topic is ongoing (see Lakoff and Johnson 1999, for example). For our purposes we need only recognize that these structures are present, and then find some way of building representations of knowledge related to our goal of decision making. Unfortunately, high level schema such as we most desire to reveal are difficult to recognize at best, and in any one individual any attempt at drawing them out becomes a study in personal bias. The schema we most desire to possess are those that are not personal bias, but instead human bias, or at least the bias of a culture or larger community.

To access these deep schema, or patterns of knowledge and decision making, we need to find a method or set of methods that identifies them through an anthropological approach, seeking and revealing universals. We also need to understand what level of complexity exists in these schema in terms of decision making: are hundreds of factors relevant when a group of schema interoperate?

The anthropological study of schema in community decision making can be traced to the development of Pattern Language theory in the 1960's and 1970's. Pattern Languages are moreover a powerful medium for communication, blending stories and relations between schema into a complex but compelling narrative. The study of the level of complexity of decision making itself has been studied recently as part of the ABC Project. These two ways of knowing *about knowing* are discussed below.

## **PATTERN LANGUAGES**

The Pattern Language concept can be traced back to the 1960's. Christopher Alexander, a graduate student working on industrial design and architecture, began examining a formal basis for design that he hoped would eventually automate quality design. The project failed, largely because quality of design is a complex issue much as urban planning and decision making are. Alexander remained hopeful, however, that some formal aspect of language would prove useful in the design process. As a professor of architecture at Berkeley, he began working on a design language, but from an anthropological point of view. By studying quality of life considerations as controlled by architecture, and carrying out this study across many cultures, the Berkeley group painstakingly derived a series of design elements, or patterns, that are culturally invariant and associated with quality of life. The resulting Pattern Language, published in 1977, became a cult hit, unpopular with architects but very popular with members of the larger community. Several books on the theory behind the Pattern Language, and on the process of building the language, followed.

The general structure of a Pattern in the language is:

1. Referents: What concept leads you to need this pattern.
2. Title: A very short title that captures the sentiment. A ranking system indicates how common the pattern was across cultures and is appended to the title.

3. Illustration: A photograph or artistic presentation intended to evoke the deep structure of the problem at hand.
4. Problem: What real-world problem does the pattern address.
5. Discussion: A story-telling-like discussion of the problem as encountered, solutions as tested, and research results consulted.
6. Solution: A law-like statement of how to address the problem. Combined with the Title, the Solution can be read as a short form for the whole pattern. In this way the entire pattern language can be read in a short time.
7. Follow-ups: Concepts that naturally follow from this problem, and in most cases, operate at a more specific or local scale. The early patterns in the language are therefor global in context, the last, very local.

The entire Pattern Language contains approximately 250 patterns, and covers architecture from the scale of a city down to the scale of a piece of furniture. The early patterns in the language are directly applicable to urban planning: in fact, many of the patterns address urban planning at the scale of community building.

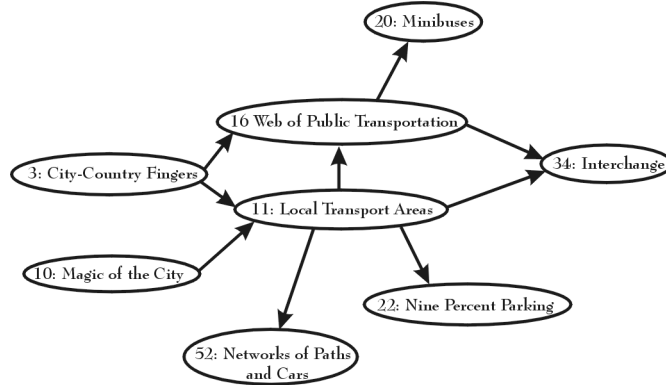
Although Pattern Languages are best applied to situations where design is a factor, even apparently ahistorical concepts from science can be portrayed using this general form. This process profoundly changes the mode of communicating (and so teaching) science because it emphasizes storytelling and historical factors instead of just solutions as abstract rules. In other words, textbooks now portray science in the form of the Title and Solution of a Pattern Language, with perhaps a small measure of history. Pattern Languages are completely about historical context, regional context, and crucially, relating concepts to other concepts in a web of knowledge and process.

Even though the Pattern Language concept predates even the definition of the hypertext concept (although not by any means the first proposal of a web-based method of teaching science) the similarities between the presentation in the book form and well-designed web sites are very strong. Pattern Languages are meant to be Web sites.

Since the book was published, the Pattern Language concept has become wildly successful in an unanticipated domain: software design. The issues in design of object-oriented software at the large scale turn out to be profoundly anthropological, rather than purely syntactical or formal. Pattern Languages of Program Design are now widely discussed, and in recent years Pattern Languages for Human-Computer Design have also appeared.

There is a strong relation between semantic networks and Pattern Languages, but important differences as well. Pattern Languages are anthropological constructions, profoundly human, and addressing the cognitive fit between humans and the environment. Semantic networks attempt to capture the ontological relations within a domain with only passing reference to cognitive fit, and in the view of many artificial intelligence researchers, the less fit the better. A semantic network based on a Pattern Language approach might not serve as a basis for automated reasoning, but it combines the best elements of human understanding (stories and ecological rules) with communication clarity (a visual knowledge representation scheme). Recall from the discussion of learning: procedural knowledge works, theoretic understanding works better, and an internalized understanding works best. Virtual experience, as provided by visual displays and

stories, provides the strongest mechanism for turning procedural knowledge into internalized understanding. Theoretic knowledge suffers from a lack of clear context, but as portrayed in a web of knowledge (pattern based or not) is contextualized.



**Figure 2:** Fragment of Alexander’s Pattern Language for urban planning, focussed on the pattern for Local Transportation. Many of the links are omitted for clarity.

This particular pattern has the statement of problem:

*Cars give people wonderful freedom and increase their opportunities. But they also destroy the environment, to an extent so drastic that they kill all social life.*

**SIMPLE HEURISTICS THAT MAKE US SMART (THE ABC PROJECT)**

While the process of building a Pattern Language - examining procedural intelligence (i.e., design) via anthropology - is a proven method for capturing and communicating relationships and properties of the world, the method does not place any strong constraints on how many factors are at work in any decision, or what the process of decision making resembles. Pattern Languages are the building blocks of decision making, but only indirectly speak to the complexity of the decisions themselves. And Pattern Languages do not give us a clear sense of discriminatory decisions are made, instead focussing on constructive decisions.

The schema approach to understanding cognition obviously places some constraints on decision making, since decision making is therein understood to involve the interaction of relatively simple, experiential rule sets. What is the granularity of these rule sets? This can be addressed both as a question of theory - how do rule sets of various sizes perform - and as a question of actuality - can we say definitively how people *actually* make decisions, what they factor in, and what they ignore.

The ABC project addresses this problem by studying, through simulation and cognitive models, what type of heuristics could be operating in humans making decisions. The research proposes very simple operational rules for decision making, and suggests that these rules are learned by experience: they are ecological rules. In other words, the ABC project suggests that at least some schema, and specifically those used in everyday problem solving, are very simple heuristic rules. The research addresses how these simple rule sets compare both in success and mode of failure with human performance. Simple heuristics, involving such trivial cognitive decisions as “what worked last time will work this time” and “predict performance using a set of yes/no decisions based on rapidly evaluable factors,” turn out to strongly mimic human performance in everyday decision making, and to completely avoid the serious limitations of abstract processes like global optimization.

If these simple heuristics are accurate, then a community decision making system must provide procedural equivalents to these heuristics, context for their application, and opportunities for internalizing the heuristics as learned schema for continued use. What

the ABC research shows us is that these procedural heuristics need not be complex. They should be as simple as possible, and involve yes/no decisions rather than factor evaluation.

**PATTERN LANGUAGE  
THEORY AS A COMMUNITY  
APPLICATION OF  
OPERATIONAL HEURISTICS**

Obviously there is a lot of overlap between heuristics as manifest in Pattern Languages, simple heuristics discovered through research, and the idea of image schema operating in our everyday mental lives. Given that we want to communicate understanding about processes active in an environment in order to empower a community to make decisions about the future of that environment, we can (based on theory and analogues to the research discussed) make a few points about how we might approach building the knowledge component of the decision support system. These are:

1. People, even experts, make everyday decisions based on very few factors evaluated in a simple manner. We need to isolate what these pivotal factors are for various questions, and present these along with contextual background.
2. Procedural knowledge, even if it perfectly summarizes the methodology of experts, will be used much less efficiently than internalized equivalents of the same knowledge. Stories and examples are much more effective at internalizing even simple procedures than abstract presentations.
3. Some decisions are not based on expert knowledge, but instead on community norms and deep and often opaque anthropological norms. Our process for establishing both the content of and relations between elements of the decision support system must somehow capture these intangibles.
4. The Pattern Language method accomplishes this capture process, surrounds procedural knowledge with both examples (often as stories) and a framework of contextual relations with other elements of the web of knowledge.
5. The Pattern Language result - a web of knowledge - is well suited to presentation using a combined semantic network browser and text/content window.
6. This Pattern Language or more likely, series of Pattern Languages needs to be integrated with the spatial domain which is so vital in understanding community decision making. We need to combine a Pattern Language for learning *about the factors* with a GIS for *showing the distribution of relevant factors*.

---

*The Georgia Basin Digital Library*

---

*The next question on Jennie's assignment is really strange. It has a web site address and the question is simply: What do you think? About what? She's not sure. She goes to the website and it turns out to be a picture of the Vancouver area, and a separate window with a poem. She reads the poem. Actually, it's not a poem, more of a dreamlike quote from some book. She clicks on the map and the poem fades out, replaced by a floating spiderweb of words and lines. The words are in short sentences: Cultural Diversity, Social Equity, that kind of thing. The area she clicked on has a small label now: Mix of Housing Types. When she moves the mouse over it a menu appears, with options like Pattern, Model, and*

In the Georgia Basin Digital Library project in particular, what are we trying to communicate?

What first-generation tools can we bring together that incorporates the idea of a decision support system that promotes deep learning?

What are the key enhancements that need to be made to build a second-generation tool?

*Links. The spiderweb is also focussed on Mix of Housing Types. She clicks on another spot in the web: Connected Wildlands and the background of the map changes. A section of the shoreline is now brighter, and a new group of labels on the map fade in. She settles down for a good look at this site.*

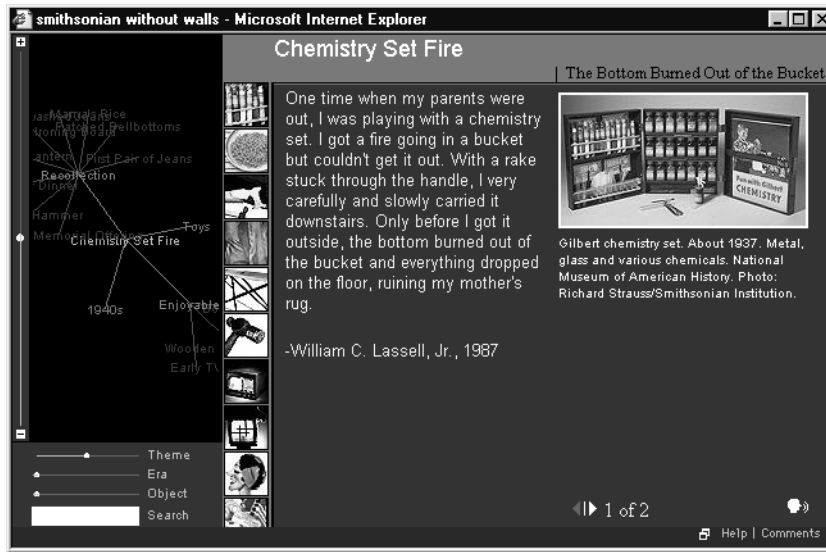
The application of these principles are being studied in a particular frame of reference: the development of community decision support tools for investigating sustainable development in the Georgia Basin region, in and around Vancouver, B.C. The tools should integrate a high level expression of issues in urban planning - perhaps simply the elements of the Alexanderian Pattern Language on Urban Planning - and be flexible enough to expand into other domains as the need arises. For example, in the Georgia Basin region the system could include a well constrained, geological domain such as the application of geologic data in evaluating landslide risk, a relevant community problem that is poorly understood by the public. The choice of exact domain is not crucial: the Pattern Language approach is general, as shown by its origin in architecture and urban planning, and current use in computer programming.

The system itself - GIS tools, browsers, and associated dynamic tools - itself represents a domain that will be somewhat opaque to new users. The help system for the tools we propose should themselves be expressed in the framework of the system. The system must be self-documenting both at the level of interface use and the level of informed use of the advanced features available.

The audience for the system as a whole is public-school educated but interested members of the public. The chosen venues include an exhibit in a science museum and a digital library website, and perhaps stand-alone versions of the system for use in an off-line setting. The first-order consumers of this system will be those with access to a modern personal computer capable of supporting a Java-enabled Web Browser, allowing the tools to include dynamic behavior programmed in, or linked to, the Java programming language.

The chosen toolset at this point reflects a pragmatic reliance on proven solutions rather than newly developed custom tools. We will extend the existing GBDL, which presently includes a GIS browser that shows dynamic views of two-dimensional maps. The detail browser level will be handled through a web database front end tool, though in principle this would ideally be done using XML and RDF methods. This browser will present all aspects of the Patterns relevant to the problems at hand, including those that document the system itself. The connectivity of the Pattern Language will be handled by ThinkMap, a state of the art GUI toolkit written in Java that provides on-the-fly network generation from underlying database tables and has been used to build systems that provide simultaneous semantics and details for museum installations already.

The key innovation that we bring using these tools is in fact two innovations, one pragmatic and one experimental. The integration of a Web GIS tool with THINKMAP based browsing of a Pattern Language meant to deepen user's appreciation of the issues in urban models is a pragmatic goal that should provide community GIS on the web, and/or in the museum, with the ability to see the relations between concepts, learn about the concepts themselves, and see how these concepts are played out in the actual setting of



**Figure 3:** THINKMAP application at the Smithsonian Museum. Note the semantic network at left and the detail browser for the selected item in the network. The icons in the centre are keyed to specific landmarks in the semantic net.

The analogue we propose is that the contextual information for using the modeling and information browsing tools in the Georgia Basin Digital Library would be categorized as a pattern language - like web.

the Georgia Basin. It should illustrate concepts in a procedural manner but follow this with story and example based support meant to build both theoretical background and intuitive, image-schema deep understanding of the issues.

The key experimental aspect is the issue of granularity, and the meta-issue of the granularity of the Pattern Languages and associated semantic networks themselves. The usability of this system, no matter how well conceived, needs to be tested by real people in a real community, and we expect to learn a great deal by test deployment of this system as well as initial deployment of some components in a museum setting.

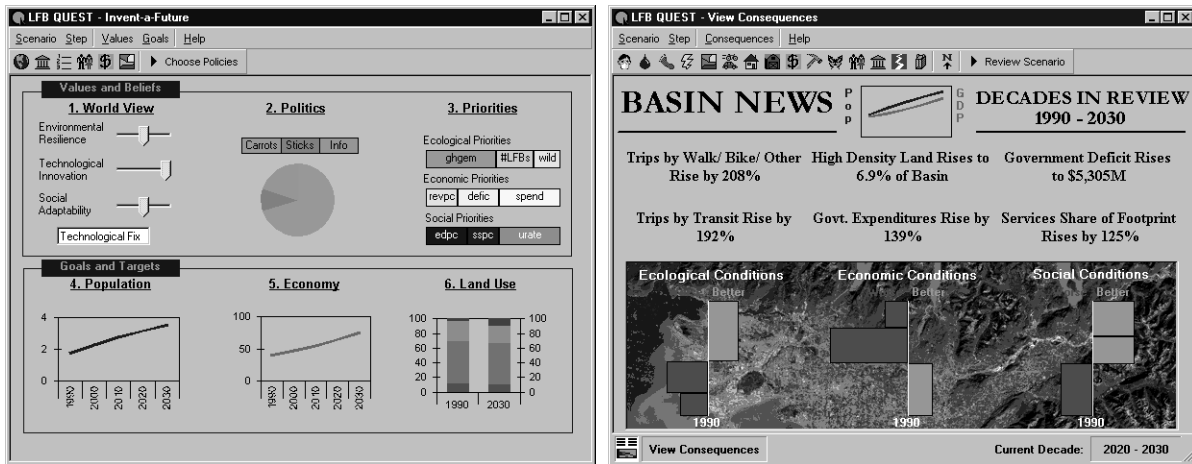
## THE QUEST TOOL AND FUTURE WORK

Although outside of the scope of heuristics and Pattern Languages as detailed above, there is another approach to community GIS that we hope to integrate with the system as described. This consists of dynamic modeling tools that make forecast and backcast models of a community given input on key factors. In other words, these models operationalize the kinds of knowledge we hope to incorporate into the Pattern Language for learning purposes, to build a predictive modeling environment accessible to users.

Obviously this has great implications for learning and decision making, since there is nothing more immersive and immediate than a simulation of the changes in a realistic environment, except for observing that environment itself. Since the modeling can compress time, or even reverse it, then modeling tools that build on and alongside the learning tools we have described so far have the potential to accelerate learning and help move understanding into the internalized schema realm that we have described.

As part of ongoing work on sustainable development at the University of British Columbia, researchers have already constructed a tool of this type. The tool, called LFBQUEST, presents a series of planning priority decisions to users (generally members of the community) and then presents a simple GIS display of how these planning priorities may be reflected in changes to the Georgia Basin environment. The current system is very strong on modeling, but provides essentially no semantics nor ability eas-

ily to gain background knowledge for any aspect of the modeling process. Note that this



**Figure 4:** LFBQUEST tool with priorities panel and run results. Note that the tool does show changes in conditions within one model.

document specifically refers to LFBQUEST, and some points made here may not apply to versions of the QUEST tool now in development.

There is a natural fit between the triadic, Pattern Language based GIS browser we are developing for GBDL and the modeling tools in QUEST. QUEST represents a possible component, or perhaps replacement, for the GIS component of the triadic browser. In this case both the conceptual semantics and data semantics embodied in LFBQUEST would be integrated into the Pattern Language map and detail browsers, so that these could provide details on the modeling process, providing a deeper engagement for users. In this mode, the tool might reflect the semantics of planning priorities as those decisions were being made - showing a semantic network or Pattern Language fragment for concepts in urban planning during decisions concerning downtown cores, and a fragment for shoreline ecosystems during decision making about shoreline land use and green space priorities. The level of engagement, and the potential for real learning, would be heightened dramatically by this enhancement to the LFBQUEST environment. And the utility of the redesigned GBDL would be dramatically enhanced as well through the incorporation of dynamic behavior linked to the data and decision support tools we have described.

This linkage would require API-level access to LFBQUEST to bridge it to THINKMAP and a browser window, which leads to a deeper problem.

There is, unfortunately, one central and critical flaw in the design of LFBQUEST that needs to be addressed if the maximum benefit is to be delivered. LFBQUEST as it stands is a linear tool. The planning process occurs first. The tool then calculates forward models. The user can then back up the process and change parameters, rerun the model, and so on. This violates one of the fundamental concepts of usability and thus learning: people have very poor short term memory, and short term memory for abstract spatial relations is among the weakest of human abilities. Having done a model in LFBQUEST, the user can back up and do another one. However, many of the key

insights from modeling come from successive models being compared, and humans are just not good at remembering GIS-style output between successive runs, remembering what the parameters were from these trials, and learning from the result. LFBQUEST needs to be non-linear, with an ability to graphically compare annotated models. And these annotations need to be linked to the Pattern browsers so that the deep causes for differences between successive runs can be examined, and real learning can occur.

The ability to not only store results from modeling trials in QUEST but also compare them would have the added benefit that LFBQUEST would then provide a strong framework for not only explanatory learning but also case-based learning. Cases, presented in a Pattern browser and with associated heuristics and semantics clearly illustrated, could trigger exploratory runs in LFBQUEST by activating stored sets of parameters, and the user could then continue the exploration process at will. Cases could be compared using newly added trial comparison tools, and at any point the concepts and data in question could be browsed.

As emphasized in the discussion of learning, storytelling has a unique and profound ability to communicate deep issues. A case-based system can be generalized to a storytelling system in which scenario models are presented as human-centred stories but linked to supporting models and background information via the LFBQUEST process and the Pattern representation of the web of knowledge. Given that the story would be embodied in a dynamic environment, non-linear stories would be simple to implement.

At a more abstract and speculative level, the process of deploying such a tool into an environment for use itself represents an opportunity for the research community to learn about decision support tool usability, knowledge representation, and new modes of information presentation on the Web. If the process and progress of use by all users is tracked, along with some measures of their background and interests, then the pattern of their use itself represents an anthropological artifact akin to the process by which the original Pattern Language was developed - the examination of universals - and also a direct reflection of the issue of granularity immanent in use patterns. These benefits will only arise if the system is deployed, rather than staying as a proof-of-concept prototype.

These kinds of linkages and enhancements would, unfortunately, require fairly major rewriting of the QUEST tools. Given the benefits of trial comparison and case-based learning, the changes may be justifiable even if a browser as described is not added. These same changes would move LFBQUEST a long way towards being a collaborative tool (since the foundation for trial comparisons by an individual, and examining causal factors, as described is equally a foundation for the same process *simultaneously* by multiple users). Finally, the addition of a rigorous API to LFBQUEST, necessary for integration at any serious level, would furthermore open up LFBQUEST for collaborative development and expansion into related markets such as community decision making related to environmental remediation and political decision making in developing countries, since the needed additional tools could be added via an API.

## Conclusion

---

*Jennie is older now, a mother with a busy job working with the Vancouver Symphony. Her son Ryan is now struggling with the same social science questions that she once struggled with. He is sitting on the back porch with his Book, which appears to be simultaneously showing the latest episode of Ryan and Fred's Dubious Adventure (Ryan's media project), and a simulation of urban change. The simulation is two dimensional to her, but she knows he is seeing it in three dimensions, and probably getting audio cues as well. The Book appears to be showing urban effects of the decline of box stores in response to the New Urbanism law of 2009. She picks up her shopping bag and heads out to the local grocer and market garden, a couple of blocks down the footpath that leads by the front door. As the door shuts, she hears the porch audio turn up as the Book senses her departure. She suspects the Book is now being instructed to put all of its resources into rendering Ryan's Adventure in 3d at full volume.*

Although many of the ideas presented here come from an academic and somewhat theoretic background, a very pragmatic goal is proposed both as a useful tool for residents of the Georgia Basin and as a mechanism for studying dynamic decision making tools in general. This document represents a snapshot of a rapidly evolving project, and we welcome feedback of all kinds to move the project towards implementation. For updates and further information, please contact Rob Harrap, harrap@geol.queensu.ca, 613 533 6000 x74951. This version complete as of January 16, 2001.

## References

---

### BOOK SOURCES:

Alexander, C., Ishikawa, M., Silverstein, M. 1977: A Pattern Language: Towns, Building, Construction. Oxford University Press, 1171pp.

The classic work on Pattern Languages, and also a very interesting narrative about urban planning.

Campbell, Jeremy 1989 The Improbable Machine: What the Upheavals in Artificial Intelligence Research Reveal About How the Mind Really Works. Simon and Schuster.

A popular account of the cognitive debate, including a skeletal overview of the power of schema theory.

Gigerenzer, G., Todd, P.M., et al 2000 Simple Heuristics that Make Us Smart. Oxford University Press, 416pp.

A collection of papers detailing the work on the level of complexity of decision making, done by the ABC research group.

Klein, G. 1999 Sources of Power: How People Make Decisions. MIT Press, 352 pp.

---

## References

A classic, field-based study of human decision making performance, ranking procedural vrs. theoretic vrs. internalized knowledge.

Lakoff, G., Johnson, M. 1999 *Philosophy in the Flesh: The Embodied Mind and Its Challenge to Western Thought*. Basic Books, 550 pp.

An exhaustive treatment of the problem of cognition in Western philosophy, and a detailed overview of image schema and related theory.

Lakoff, G., Johnson, M. 1983 *Metaphors We Live By*. University of Chicago Press, 242pp.

A classic treatment of how everyday language reveals the body-based metaphorical nature of thought and communication.

Stephenson, Neal 1995 *The Diamond Age: A Young Ladies Illustrated Primer*. Bantam Spectra.

A science fiction novel that examines the questions of what education could be (with a high-tech dynamic book) and what communities are (with a discussion of neotribalism). The discussion of the book and how it guides learning parallels many of the points made in this paper.

Wierzbicka, A. 1996 *Semantics - primes and universals*. Oxford University Press, 500 pp.

A technical overview of the existence of universal semantic primitives from a linguistic perspective. There are many similarities between the 'primes' here and the image-schema of Johnson and Lakoff.

## WEB SOURCES:

[www.thinkmap.com](http://www.thinkmap.com) - information on the THINKMAP toolset.

[www.envisiontools.com](http://www.envisiontools.com) - information on the QUEST tool.

[www.esri.com](http://www.esri.com) - information on Arc IMS3 toolset.

<http://web2.si.edu/revealingthings/> - Smithsonian THINKMAP exhibit.

<http://www-abc.mpib-berlin.mpg.de/users/ptodd/SimpleHeuristics.BBS/>

- Simple Heuristics Summary Paper

<http://www.designmatrix.com/pl/ecopl/index.html> - Ecosystem Pattern Language Example

<http://www.patternlanguage.com> - Alexander's Pattern Language site.

## *Decisions and Learning: Background Details*

---

As summarized in the main body of the report, we need to understand three types of knowing: procedural understanding, theoretical (abstract) understanding, and internalized, intuitive understanding. To expand on and justify the ranking of utility of these levels of knowing, we need to look in detail at the subject as seen by philosophers, neuroscientists, artificial intelligence researchers, and linguists (collectively, these disciplines form the larger domain of *cognitive science*). Recall the ranking:

1. Procedural knowledge is better than no knowledge, but very slow to use and easily misapplied.
2. Theoretic knowledge is considerably better than procedural knowledge - in fact, procedural knowledge is often theoretic knowledge in prescriptive form - but slow and often difficult to align with the complex 'real world.'
3. Internalized, intuitive knowledge is better in the complex 'real world' than either procedural or theoretic knowledge, especially under time constraints. It is often misapplied, difficult to place bounds of relevance on, and impossible to communicate directly. It is also responsible for many maligned flaws of reasoning, such as stereotypes, prejudices, and other biases.

Learning involves building theoretic and internalized knowledge out of received procedural knowledge, received communication, and experience. In absence of education, internalized knowledge dominates, and educators often attempt to suppress internalized knowledge in favor of theory. The western intellectual tradition is one of suppressing feeling in favor of logic. Logic is, unfortunately, poorly suited to decision making in the complex, ever-changing world.

### **ABOUT MODELS OF LEARNING**

The community of philosophers, and the community of researchers into artificial intelligence, have each addressed the issue of what knowledge and skills are. Recently, the computer has become a metaphor for mind, and the appropriateness of different versions of this metaphor outline a raging debate in academia.

There are two approaches to systematizing thought in machines, and each is believed to be important to some extent in the human mind. They are profoundly different in implication. And they correspond directly to our modes of decision making.

The first approach is to treat problems of knowledge as problems of logic. In fact, the development of modern logic was specifically intended to be a method for systematizing, or axiomatizing thought. Knowledge in this symbolic world is rules and operations on rules and data about the world. The tradition of symbolic representation began in earnest with Frege in the 19th century and reached its peak with the wild optimism for the potential of expert systems, computer based rule manipulation systems, in the 1980's. Expert systems, and symbol systems in general, are compact, self-explanatory, easily communicated, and very brittle in application. Even the slightest ambiguity in what terms mean, what data is, and how the world works leads to complete failures of these systems. On the other hand, many of the techniques developed for representing information in these systems are highly efficient as human tools, because they root out ambiguity and show relationships clearly. Semantic networks, discussed earlier, are an example of a logic-based communications and reasoning mechanism.

The second approach is to use a model of thought based on the structure of the human brain. In the brain, a dense and interconnected network of neurons process information by taking many inputs and filtering these efficiently into a few outputs. The inputs may be senses or other neurons, and clearly must include all aspects of language and meaning, including symbolic meaning. The details of how this is done for any specific problem is unclear: the wiring represents a black box approach to problem solving. What is clear is that with some mechanism for rearranging the wiring (actually, strengthening or weakening connections) the network of neurons will over time 'train' to be a very effective custom processor. Computer analogues to these human neural networks, called artificial neural networks, can be built and indeed are extremely powerful tools for discrimination and categorizing noisy data. This is usually called the connectionist approach to reasoning with computers. Artificial neural networks are messy, opaque, but very robust tools for dealing with a real-world level of complexity. They suffer from a serious flaw with regards to our quest for a toolkit: they offer no means of communicating knowledge or methodology.

It is clear that the human mind uses both neural network approaches and somehow hosts a symbolic system on top of the neural network. Language is a symbolic system, as is math, and humans clearly use these for some level of decision making. Theoretic knowledge, even if to some extent internalized, is largely a symbolic structure for understanding the world. Procedures are clearly symbolic. Expert intuition on the other hand clearly involves a heavy connectionist component. Communication is clearly symbolic. Deep learning is clearly a mix of the symbolic (to get the information in and manipulate it formally) and connectionist approaches (intuition takes over and dramatically outperforms symbolic thought in real world cases).

## **SO WHAT IS DECISION MAKING?**

Clearly things have been oversimplified, and in the real world humans use a combination of low-level connectionist processing and high-level symbolic processing. What sort of hybrid structures exist in between neurons and abstract thought? The evidence from this comes from the errors people make in real world situations, which reveal the processing going on behind the scenes.

First of all, remember that intelligence as it evolved is not about science, or complex problem solving in a city environment. It is about survival in a complex, ever-changing natural setting full of dangers and with scarce resources. This is also the setting for the development (likely co-evolution with intelligence) of language. Intelligence is about getting up a tree when a predator approaches, and building simple tools to manipulate the environment. It is a testimony to the power of the foundations of intelligence that we have done so much more than this, but it is simply ridiculous to assume that science as we now know it, or urban decision making or the like, was the point of intelligence.

What we do know about the middle layer of knowledge is that it is for the most part based on simple prescriptive relationships. These are apparently directly coded in neural structures, but include recognition of symbolic or abstract concepts. Many simple language structures are built around these relationships, and one of the key purposes of language is to tell stories, wherein these prescriptive relationships are transferred between individuals. Neural structure becomes symbolic structure and then again, neural structure. These prescriptive relationships are usually called schemas. They are prescriptive in the sense of offering a simple shortcut for decision making based on past experience,

but not in the sense of being a conscious mental check list as was the case with the procedural checklists mentioned earlier. Schemas are cultural universals found in stories, found in behavior, and guiding decision making in real world settings. Without schemas, which simply point out what variables are important in decision making, we would be overwhelmed by irrelevant information when making crucial decisions. Faced with a predator, we need not worry about the details of its coloring. Faced with a model for recognizing predators at a distance, we abstract only the relevant features (teeth, body plan) and not the irrelevant ones (exact color) in order to build higher level survival oriented structures.

Knowledge, schema-based, is profoundly operational. Schemas are prescriptions for dealing with everyday problem solving. An expert is simply someone with a large collection (hundreds, thousands,?) of relevant schemas. Intuition is the subconscious activity of schemas.

Symbolic thought at a higher level was very likely originally a schema transfer system, a system for passing real-world data, and to some extent (and increasingly) a system for more abstract knowledge. Humans learn first by operating in the environment, as young children, and then by stories, and finally by building more complex abstract systems such as math, logic, and other subjects at a conscious level. As discussed earlier, operational and story based schema knowledge is far more accurate in complex world decision making than high level procedural knowledge and/or theory. The human neural architecture is vastly superior to expert-system like formal reasoning when it comes to selecting relevant variables, making quick decisions, and the like.

On the other hand, schema knowledge is over-applied routinely. Stereotypes, prejudices, and biases are all ways of seeing what may be locally-correct schema knowledge outside of relevant domains. The fact that many of these are passed on through storytelling in a cultural setting highlights the role of social interaction in learning. We cannot unlearn knowledge schema. You cannot look at a letter and not see the letter, instead the lines. Similarly, a trained scientist has schema for understanding their domain and cannot undo these. Differences of scientific theory are often differences of world view; different schema give incommensurable results. Science changes by major paradigm shifts and these are routinely generational. It is in fact through the over application of schema that we can most easily study their activity.

We can't all be experts, we can't totally rid ourselves of learned (schema-) biases, so what exactly is the point of all of this?

First, theoretic knowledge in and of itself is inferior to practical, on-the-job style knowledge. It is both slower and less accurate. It is, on the other hand, visible and can be generalized. Perhaps most importantly, it can be easily communicated.

Second, people learn prescriptive decision making through stories, and examples, far better than through consideration of abstract principles. The more immediate, or 'experience-like' the examples and stories are made, the more clearly the prescriptive knowledge will be translated into schemas. We must maximize opportunities to translate formal/procedural knowledge into virtual experience. We must not underestimate the power of stories and examples: there are transcriptions of visceral experience.

Third, there is a universal, anthropological, and profoundly operational set of mechanisms at work in learning. If we can understand even a fraction of the schema and operations on schemas at work we can make great inroads into understanding how to make our examples and prescriptive knowledge, and even theoretic knowledge, more effective. Since communication is really about sharing schemas and real-world data, these mechanisms should also maximize our ability to communicate effectively.

This last point deserves more attention. How can we 'get at' the mechanisms and schemas that are at work? Examining our own subconscious is clearly not an option. Studying the thought of a single individual is also dangerous. Each of us has strange idiosyncrasies as well as downright wrong views of large parts of the world. One approach is to view the problem as anthropology - how do communities of humans, worldwide, treat specific problems. Invariants across communities should better reflect underlying mechanism than details of any one communities or individuals approach to decision making. Another approach is to postulate simple operational rules - heuristics- and to evaluate their performance in a realistic setting.

And this anthropological view is what leads us back to Pattern Languages, simple heuristics, and the blending of stories and theory advocated in the main flow of the report.