2006 CCRTS
THE STATE OF THE ART AND THE STATE OF THE PRACTICE

Knowledge Glyphs: Visualization Theory Development to Support C2 Practice

Topics:
Cognitive Domain Issues, C2 Architecture, C2 Concepts and Organizations

Name of Author(s)

Randall Whitaker, Ph.D
Gina Thomas-Meyers, Ph.D

Point of Contact

Gina Thomas-Meyers, Ph.D

Name of Organization

Air Force Research Laboratory (AFRL/HECS)

Complete Address

AFRL/HECS
2698 G Street
Wright Patterson AFB, OH 45433-7604

Telephone / Fax

(937) 255-9474 / (937) 255-6555

E-mail Address

Gina.Thomas-Meyers@wpafb.af.mil

Abstract: Recent and ongoing research into the subject of knowledge glyphs illustrates how the states of the art in theory and practice interact and opens up new requirements for innovation with respect to each. This paper reports research that was motivated by recognition of the need to augment current "best practices" in C2 symbology in order to accommodate the complexities of modern warfare. The notion of knowledge glyphs derives from an understanding of shortcomings of current visualizations and a desire to create more sophisticated visualization elements that allow interaction between information regarding the entities portrayed and the contexts in which decisions must be made regarding those entities.

This paper lays out the most recently developed definitions of knowledge glyphs, which were derived from the notion that a model for knowledge glyphs would have to be internally coherent, referentially comprehensive, and explicitly linked to the visual interface elements that were the foci of our interest. Our examination of this model suggests a subsequent state of the art in visualization technique linked to an analogous new state of the art in this technique's usage. In this way we have 'circulated' from practice to theory and back again.

Introduction

The history of our most recent research reflects this year's CCRTS theme by illustrating how the states of the art in both 'theory' and 'practice' not only interact but also open up new requirements for innovation with respect to each. The research reported in this paper concerns advanced information visualization as a means for improving C2 decision making. This research was motivated by recognition of a need to augment current 'best practices' in C2 symbology to better accommodate the complexities of modern warfare and the information domain within which commanders must operate. In other words, this effort can be seen as theorization in response to practices, which themselves embody prior theorization. This effort illustrates the reciprocity between theory and practice at any given point as well as the manner in which these factors interact to guide progress over time.

This illustration involves inaugural Air Force Research Laboratory (AFRL/HECS) research into the subject of *knowledge glyphs*. Dr. Thomas-Meyers chairs the AFRL/HE Visualization Group that has undertaken this research, and Dr. Whitaker has served as a principal researcher. The notion of a knowledge glyph derives from an understanding of current visualization tools' shortcomings and a desire to create a more sophisticated class of visualization elements. The main thrust of our 2005 work focused on establishing a framework for defining basic constructs and interrelating them with conventional interface or visualization concepts. This would set the stage for demonstration applications and experimental studies to be pursued in 2006. In the following sections we shall review the issues motivating this work, the results of the 2005 effort, and our emerging vision of how this new class of visualization artifacts would be employed.

We would like to acknowledge that this work has not been done in isolation. The insightfulness of Dr. Michael Young of AFRL/HECS initiated the formation of the AFRL/HE Visualization Group. Key members of the group have helped to develop notional visualizations that have been instrumental to the success of this undertaking. In particular, Dr. Dan Repperger of AFRL/HECP, Ms. Denise Aleva of AFRL/HECV, and Mr. Steve Fullenkamp of General Dynamics have added much to this project through both the conduct of research on current C2 symbology (funded by the Air Force Office of Scientific Research (AFOSR)) and the development of novel and innovative visualization concepts. In addition, they have contributed to the development of knowledge glyphs through extensive participation in the Visualization Group. Other key participants in this group include Dr. Paul Havig, Mr. George Reis, and Ms. Francine Schaffner of AFRL/HECV and Mr. Gil Kuperman of AFRL/HECS. In addition to the funding by AFOSR, we would like to acknowledge the support of an Air Mobility Command program, Work-Centered Interface Distribution Environment (WIDE), which is managed by Mr. Jeff Wampler of AFRL/HECS. (Complementary work completed through the WIDE program was reported at the 2005 ICCRTS conference.) Knowledge glyph development also has the current support of the Commanders' Predictive Environment (CPE) program, which is co-managed by Dr. Janet Miller of AFRL/HECS and Jerry Dussault of AFRL/IFS. Current and on-going movement of knowledge glyphs from theory to practice is being facilitated by the Human Systems Integration Information Analysis Center (HSIAC).

Research Motivation: Shortcomings in Best Practices

The most concise description of our knowledge glyphs effort is theory formulation specifically directed toward mitigating perceived shortcomings in current best practices. The state of the art in C2 visualization capabilities continues to advance. However, each advance either induces or suggests opportunities for further improvement. In this case, providing consistency in C2 visualization elements set the stage for perceived opportunities for improvement. MIL STD 2525 ('B'), issued in 1999 and revised in mid-2005 (DOD, 2005), represents current best practice in battlespace symbology. It specifies a structured set of elements which - singly or in integrated combination - provide a basis for illustrating (e.g.):

- Units
- Equipment / vehicles
- Installations
- A limited set of operations involving the illustrated elements
- Some limited indication of spatial movement for mobile elements
- Meteorological and oceanographic (METOC) factors
- Signals Intelligence (SIGINT) factors
- Military Operations Other Than War Symbology (MOOTW) factors

MIL STD 2525's symbology set was designed for application in the context of geospatial representations - i.e., geographical maps. It is therefore reasonable to characterize MIL

STD 2525 as representing a 21st Century refinement on a Napoleonic era commander's battle map. So long as one assumes a 21st Century commander is adequately supported by a 19th Century information artifact, MIL STD 2525 constitutes an advance. Such an assumption is naturally open to question, and it is in the context of questioning this assumption that the prospects for constructively augmenting the standard appear.

To illustrate and examine both the limitations of a 19th Century representation and the directions in which constructive improvements lie, we must consider not just the map but also what role the map plays. The classical battle map is a representation of a context (in this case - a geospatial context) within which a competitive 'game' is being or is to be played out. In this sense, the battle map can be construed as the commander's chessboard. For the sake of illustration, let us proceed with the metaphor of a chessboard being employed in a game between remotely-located adversaries, neither of whom has direct contact with the other, and each of whom must use his / her own chessboard as the sole tool for maintaining situation awareness (SA).

First, let us consider the ways in which a chessboard remains as viable a battlespace metaphor in the 21st Century as it was in the 19th. The chessboard depicts the configuration of a set of elements (pieces) as well as the location of each discrete element. Absent any other information (e.g., historical data on the game's moves to date) the chessboard is a static picture of where the pieces are and nothing more. If one projects upon the chessboard the allowable next moves for each piece, one obtains a static picture of the entire set of next states in the game. If one writes upon each piece data about its basic identity and general capabilities, the set of elements is augmented so as to be self-descriptive upon examination. Figuratively speaking, this depiction is equivalent to the maximum information one can presume to obtain from a MIL STD 2525 - enabled map.

Now let us consider the ways in which the chessboard (and, by extension, a static battle map) falls short in supporting adversarial decision making of the sort performed in modern command and control. A chessboard may have been a reasonable metaphor for a battlefield in a time of highly-stylized warfare conducted with limited weaponry. It is not, however, a completely effective metaphor for modern warfare. The parallelism and complexity of battlespace events are completely different from the linearly stepwise moves of a chess game. Multiple 'pieces' may be moving in unexpected and novel ways - either individually or in concert. Battlespace elements (e.g., vehicles) can shift capabilities from one 'move' to the next. One must always second guess the represented state of the battlespace, because the adversary is capable of masking his / her assets. One must always bear in mind that one or another 'chess piece' may not be there or may not be what one thinks it is. Moreover, one must also second guess the state of one's own local chessboard; there is always the possibility of errors in assembling the current representation (e.g., from misreading the opponent's move as written on a postcard).

All these factors are important in comprehending the adversary's intentions and plans. None of these factors are capable of adequate portrayal on our metaphorical chessboard, meaning none of them are capable of adequate portrayal on the equivalent battle map -

even with MIL STD 2525 enhancements. This implies that additional symbolic refinement is required to accommodate:

- Uncertainty with regard to the current state of the battlespace
- Uncertainty with regard to the current state of the battlespace representation itself
- Uncertainty with regard to the probable next state of the battlespace
- Richer representation of element dynamics
- Richer representation of element characteristics.

This set of new capabilities became the targets for the next step in our exploratory effort.

Proto-Theoretical Artifacts: Examples of Next-Generation Practices

The next step would be to develop illustrative examples of tactics (i.e., new 'practices') providing these sorts of capabilities. This need gave the Visualization Group a concrete basis for (a) creating and evaluating candidate innovations as well as (b) identifying key factors that must be addressed in theory formulation.

Portrayal of Uncertainty

Because static representations portray a fixed set of elements in a fixed state, they are ill-suited for portraying uncertainty in elements or in their interrelationships. Unfortunately, such uncertainties are certain to be encountered in military operations. The multiple ways in which uncertainty affects C2 decision criteria translate into multiple ways in which the representation of uncertainty must be addressed. One form of uncertainty pertains to the current state or configuration of the elements being portrayed. A knight either is or is not occupying a particular space on a chessboard physically configured in the present. Physical location connotes certainty on the chessboard, because 'a knight is a knight is a knight', and the identity of a knight is clear and unambiguous. In a real world battlespace the 'pieces' may not be what they appear or are construed to be.

A second form of uncertainty pertains to the prospective or possible state(s) of the elements being portrayed. Modern C2 is future-oriented, in the sense that one is always contemplating tomorrow's battlespace state and working with respect to it. A knight may or may not occupy another given space on its next or a subsequent move. One can populate a chessboard with additional knights to illustrate the possible next locations for the current knight. However, there is no straightforward way to depict the relative probability of the current knight's next move. Similarly, a static display (a la MIL STD 2525) can do no more than give the commander a static picture of one possible future state. It would be more flexible and less misleading to give the commander a picture indicating the range of possible future states for one or more elements.

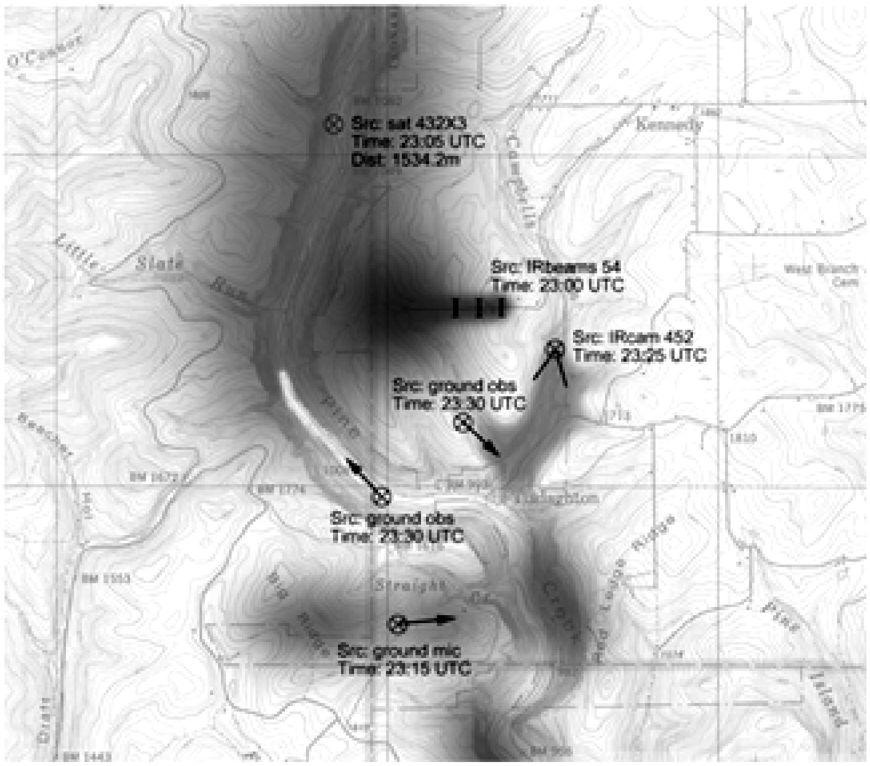

Figure 1: Display of Probabilistic Data

Figure 1 illustrates a geospatial map upon which a variety of symbology is overlaid. Among this symbology are icons denoting discrete entities. The shaded 'blobs,' though, are representations of a different type. They represent the probabilistic envelope within which one or another entity might be expected to move or have moved during a given period of time. The entity’s prospective movements are plotted, not the entity per se. Similar such representations might be employed to depict (e.g.) estimated ranges of effective fire, distribution of mass assets, and the like.

A third form of uncertainty relates to the correctness of the portrayal itself. The reason for using the metaphor of a chess game between distant players is because (a) this is the form of chess in which each player is dependent upon his / her own local representation and (b) it is therefore the variant most figuratively akin to the situation in which modern commanders operate. In such a scenario, each player is responsible for the correctness of his / her representation of the state of the game. There is always the possibility of error or misrepresentation, and this is unavoidable. However, to the extent the representation manager is aware of qualifications pertaining to the elements being portrayed (or, more to the point - the data underlying the portrayal of those elements) it would be useful to cue the user on these qualifications. Such qualifications might be documented and hence portrayed for factors such as: age of the data; deficiencies (gaps, omissions) in the data; manually-flagged indicators of uncertainty or suspicion; and automatically-generated cues alerting the user to discrepancies with other data. The notion of such additions leads to the focal topic of the next section.

Richer Representations for Entities Portrayed on the Battlespace Display

The second innovation theme was a desire to enrich C2 visualizations – particularly with respect to the visual elements used to portray discrete entities. In a conventional windowing environment, such entities are likely to be graphically portrayed as icons - i.e., graphical units generated in 1-to-1 correspondence with the 'real world' entities they denote. Icon features may be coded to allow the user to identify what type of entity the icon designates, as done via shape and other characteristics in MIL STD 2525. However, denoting the entity is often the extent of what an icon can do. On a geospatial map, a polyhedron conventionally associated with a tank does no more than indicate that "a tank is located at this set of spatial coordinates." This is informative about the tank being situated in the represented geospace, but it tells the user little about the tank itself. Is it active? Is it moving? Is it known to have been destroyed or disabled already?

These and other relevant questions cannot be answered by a simple icon - unless, of course, the icon itself is 'loaded up' with additional features or characteristics indicative of such additional information. Unfortunately, the arity of the set of all relevant data items about a given entity can easily exceed the arity of the possible set of features which can be coded on an icon. As a result, it is commonly the case that additional information about an entity has to be obtained from a display element other than the icon itself. For example, a user may 'drill down' into the information on a given entity by manipulating its icon to call up an additional palette or window within which this information is given.

The AFRL Visualization Group generated two illustrative prototypes of such enriched icon plus drilldown capabilities and demonstrated them to the group in spring 2005. One such illustrative prototype was a multi-level representation system for entities such as a tank on a battlefield. At a top level, this system provided an iconic element representing the tank itself, augmented by visual cues and add-on elements used to indicate the real world state of the tank. This iconic representation was augmented by an element which serves to organize and characterize the range of additional data available on the entity. This graphical element - termed a *polytope* (Figure 2) - offered multiple sub-elements (in this case, 'rays' emanating from a central core figure). Each of these sub-elements was associated with a feature or factor pertaining to the entity, and each was enabled to use as a hot link to call up additional data (as available) relating to that feature or factor.

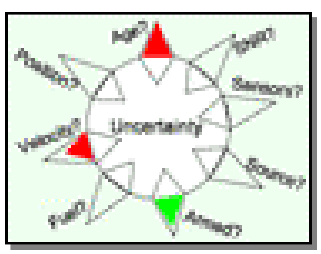

Figure 3: Polytope

A different approach to extended data access was taken with the 'Cube' concept (Figure 3). This concept used a 3-dimensional cubic figure as an enriched representation for a given battlespace entity - in this case a military unit. By rotating the cube, the user would be presented with up to 6 different 'facets,' each one of which contained a data representation. The Cube was therefore configured to offer maximal total display area in as small a space as possible.

Figure 3: The 'Cube' Representational Concept

Even though the Cube affords additional data on the entity being portrayed, its on-screen size would have to be huge to permit it to display all the available data that might be needed. As a result, the Cube was configured to permit a face or facet to serve as a hot link to a pop-up window providing more extensive data. An illustration of this two-stage drilldown protocol is given in Figure 4 below. In the illustration, the 'ammunition' aspect of the military unit is graphically addressed on one face of the Cube. This face provides a simple graphical cue on how many types of ammunition are stocked by the unit and the relative quantities of each estimated to be on hand. This presentation can be manipulated to invoke a more detailed textual display providing specific numerical estimates on the ammunition inventory.

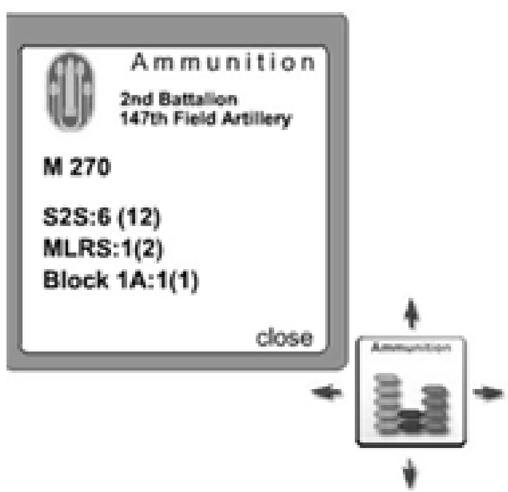

Figure 4: Data Drilldown in the Cube Concept

These illustrative prototypes' features, combined with the earlier themes, provided guidance for creating a list of characteristics attributable to the prospective innovation that was our objective.

Criteria for Theory Development: Key Characteristics of the New Concept

The knowledge glyph precedents and examples discussed earlier are not identical among themselves. Still, there are some points upon which most or all of them evidence features which appear to connote a new perspective on visualization artifacts. Review and analysis of the above-cited evidence led to the identification of a set of features or themes judged to be intrinsic characteristics of the knowledge glyph concept. In this section these perceived characteristics will be presented and discussed. In the next section a knowledge glyph definition will be generated that accommodates these characteristics.

A Knowledge Glyph is a Form of 'Glyph'

Whatever a knowledge glyph may be, it is a particular form of 'glyph.' This assumption is implied by the two-word format of the term itself, but that trivial implication is not the entire basis for this claim. As will be more fully illustrated in the subsequent sections, a knowledge glyph is to be distinguished from a conventional icon, both in terms of its role (within an interface display) and in terms of its utility (to the user). To distinguish this new concept from an icon it will be necessary to develop a definition for a 'glyph.' To distinguish the particular type of glyph meeting the utility requirements derived from the factors discussed above, it will be necessary to develop a definition for what makes a 'glyph' a 'knowledge glyph.'

A Knowledge Glyph is a 'Super-Icon'

All the examples presented earlier treat a knowledge glyph as a discrete interface element connoting some specific entity. In this sense, a knowledge glyph is or can be construed as analogous to a conventional icon. What, then, distinguishes a knowledge glyph from an icon? Each of the examples illustrates how a knowledge glyph does not simply denote the presence of an entity within the coordinate or referential space of the dataspace. Each of the examples goes above and beyond simple entity denotation to provide additional or deeper information about the entity it denotes.

A battlespace map icon for an entity (e.g., a battalion or a tank) need do no more than denote 'battalion' (or 'tank') via its format and the geospatial location of that entity via its position in the 2-dimensional map representation. The Cube goes beyond this simple denotative functionality by providing a modest set of varied battalion data on its face. The earlier specimen similarly goes beyond simple denotation by providing graphical cues indicating tank status or characteristics. The probabilistic force movement 'blob' does not simply indicate where the subject unit is at one point in time but where it may be over an extended timeframe and how likely those relative positions are estimated to be.

Even though a knowledge glyph can serve as an icon, it is demonstrably more. As a result, the first clue to a knowledge glyph definition is the fact that it is a 'super-icon' whose presentational utility is not limited to denoting the simple fact of an entity's existence and relative 'position' in and of itself.

A Knowledge Glyph is a 'Micro-Interface'

The illustrative examples of the last section included capabilities for drilldown to additional data. These capabilities were user-invoked through proactive interface manipulations (e.g., clicking on a graphical element). These capabilities imply a knowledge glyph is not just a static display element but rather a means for accessing additional data (and, conceivably, additional functionalities relative to those data). As such, a knowledge glyph is capable of serving as a 'micro-interface' allowing the user to invoke informative material via the entity representation itself.

Furthermore, those examples illustrate how a knowledge glyph serves as more than a 'one-shot' micro-interface - i.e., as a device allowing discretionary access to more than a single pro forma data presentation. With the Cube, multiple drilldown paths are offered by the Cube's six faces. Similarly, multiple paths are afforded by the polytope invocable from the top-level representation. This flexibility and multi-path capability distinguishes this notion of a 'micro-interface' from the 'pop-up' or 'tool tip' windows commonly encountered in modern desktop applications.

The implication is that, by definition, a knowledge glyph is capable of dynamic manipulation and that this manipulation is minimally capable of invoking additional data on the entity denoted by the glyph. The reason for treating data invocation as the one essential type of dynamic response intrinsic to the knowledge glyph is derived from the notion that knowledge glyphs facilitate examination and exploration of the multiple coordinate or representational spaces their denoted entities intersect. This point will be discussed further in the following subsection.

A Knowledge Glyph Interrelates Referential Contexts

One key feature evident in the AFRL examples cited above is the manner in which a knowledge glyph affords the user access to information about the entity denoted by the glyph within the referential context of a display (e.g., a tank depicted on a map) even though the information provided is not properly contained within that initial referential context. All the examples allow a user to access data having nothing to do with the geospatial context within which the subject matter is initially presented. The probabilistic 'blob' example addresses geospatial location, but breaks out of the map context by portraying data across a period of time and with respect to a probabilistic qualification extrinsic to the delineation of the map upon which it is projected.

Command and control decision making cannot proceed solely on the basis of knowing what entities are where within the battlespace. The implications of those entities' deployments need to be evaluated. Some such implications relate to the entities themselves, as addressed with the drilldown data. Other implications relate to the entities as they are or as they may conceivably be in the wider context of reference, as addressed by the probabilistic 'blob' example.

As we have seen in our 7 years of work-centered support system (WCSS) work with Air Mobility Command, modern military C2 personnel must employ multiple legacy systems, each one of which contains critical data on one or another aspect of a particular item of interest. Each of these legacy systems organizes and presents its data in a manner consistent with the perspective or vantage appropriate to its specialization. One can say a given entity (e.g., an aircraft) has 'projections' into all the databases within which data about it are stored. Alternatively, one could say that a single aircraft entity is fully described (within the bounds of the available IT and data) at the intersection of the dataspaces realized by those systems. The sample knowledge glyphs cited above illustrate how a visualization element on one system's interface can serve as a representation of this juncture among dataspaces or referential contexts. The examples that provide 'micro-interface' functionality illustrate how this element can allow the user to examine the entity from another system's contextual vantage without having to exit the current system and go interact with the other.

Summary of Pre-Theoretical Foundations

By midsummer 2005, the following points had come to circumscribe the notion of knowledge glyphs:

- Whatever a 'glyph' may be, it needs to be defined in terms of being a visualization element associated with a given thing or 'entity.'
- This assumption means the definition must account for the referential context underlying the visualization at hand as well as any additional contexts in which the entity is to be portrayed.
- A 'glyph' needs to be defined as something more than an 'icon.'
- We need to explain how the 'micro-interface' functional attribution fits into the otherwise 'structural' definitional framework.
- We need to account for implicit distinctions among 'data,' 'information,' and 'knowledge' (for our purposes).

These conclusions set the basis for the subsequent work. By the late summer of 2005, we had proceeded on the basis of these conclusions to derive the following points as the basis for constructing a model defining a knowledge glyph:

- A knowledge glyph is a specialized or qualified instance of a glyph.
- A glyph is a visualization element distinct from an 'icon.'
- A glyph is an enhanced version of an icon
- A glyph can trivially serve as an icon.
- As such, all glyphs are 'iconic,' but icons are not 'glyphic.'
- A knowledge glyph is distinguished from a glyph by virtue of affording users additional leverage on the semantics of the entity being depicted.
- The semantic support afforded by a knowledge glyph lies in its ability to interrelate the denoted entity's projection or occurrence within multiple referential contexts.

In late September 2005, a presentation of the first-generation working definition was briefed to the Visualization Group. There was no substantial criticism or request for modification raised by the group at that time, and this working definition was therefore adopted for exploratory application. In the following subsections the first-generation knowledge glyph definition will be presented.

Regarding Icons

Over the last 20 years, we have become accustomed to the term 'icon' denoting a discrete graphical element presented on an interface display. An icon is not just any graphical unit, though some loose usage of the term might lead one to suspect that. Instead, an icon is an icon by virtue of the fact that it denotes an entity of interest in the context of the visual presentation. This characterization of an icon as a form of abstract representation is actually relatively modern, having originated with the semiotic framework of the American philosopher Charles Sanders Peirce. The relevance with respect to Peirce's definition lies in the idea that an icon signifies or designates a single thing in and of itself, 'as is.' An icon does not qualify or describe the entity it signifies; it only denotes it.

Adhering to Peirce's precise usage of the term means that an icon carries little or nothing in the way of 'semantics.' The implication of the icon lays not so much in the icon itself as in the referential context in which it is placed. For example, a small dot on a sheet of paper does not automatically denote a place where people live. That same dot on a 2-dimensional page of a road atlas denotes a settlement at a particular position in the geographical coordinate system being represented. The dot tells nothing about the settlement itself. The dot at a particular position in the geo-referenced grid indicates that the settlement is 'there.' For the purposes of this definitional exercise, then, the term 'icon' will be strictly circumscribed to mean no more than what Peirce originally intended. As such, an 'icon' is a display element that serves as no more than a datum within the presentation context. Phrased another way:

> *An icon is a presentational element which designates an object of reference and specifies its 'locus' in the presentation context.*

Regarding 'Glyphs'

There are few, if any, clear precedents upon which to base a characterization of a 'glyph' as a distinct class of interface elements. The term 'glyph' itself (derived from the Greek for 'to carve' or 'carving') is used in reference to diverse objects ranging from carved symbols to typographic marks. As a root component, 'glyph' contributes no more than the notion of 'carved in stone' to compound terms such as 'petroglyph' or 'hieroglyph.' The term 'glyph' has been sporadically applied to a variety of IT-related innovations. However, few of these have to do with interface or visualization issues, and arguably none of these few align with the intent or thrust of the AFRL knowledge glyph research. Examples of such 'glyph' allusions include:

- 'Glyphs' providing programmers a toolkit of simple modular elements for building an interface's graphic structure (Calder and Linton, 2003).
- 'Glyphs' as components of a typographical set (Microsoft Typography Standards).
- 'DataGlyphs' as printed digital coding augmenting readable text on document hardcopy (Xerox PARC, 2004).
- Pseudo-three-dimensional objects employed for rendering mathematical and other abstract data (e.g., Kindlmann, 2004; Chuah and Eick, 1998).

Of these examples, it is the last which comes closest to matching the visualization subject matter at issue here. Even in that case, there is no close correspondence, because 'glyphs' are taken to be visualization data forms whose shape and / or size parameters are exploited to enrich their capacity for portraying complex information. This lack of definition means the field was wide open for deploying a sense of 'glyph' suited to our research purposes. By autumn 2005, we had begun defining a glyph in contrast to an interface icon. Like an icon, a glyph denoted a particular entity of interest. Unlike a basic icon (as characterized above), a glyph afforded the user information about that entity above and beyond its presence and its coordinate position in the referential context the visualization represents. This assumption led to our first-generation working definition of a glyph:

> *A glyph is an icon which additionally affords the user access to information about the denoted entity on its face. This information may be data concerning the denoted entity or data about available redirection to another presentational device where such additional information may be obtained.*

Because the provision of additional information may take either or both of two forms (on the face of the glyph or outboard of the glyph itself) it will be a good idea to introduce labeling for these two types of information presentations. Information presented on the face of the glyph itself will be termed *intraglyphic*, and information invoked via the glyph (but presented in some fashion outside the glyph) will be termed *extraglyphic.* Intraglyphic presentations are realized on or within the glyph artifact or adjacent to it within the glyph's referential / display context. Extraglyphic presentations are realized on or within a separate referential / display context, which may be implemented as an overlay or other coherent addendum projected onto or otherwise correlated with the glyph's referential / display context.

It is important to note that, under this definition, an icon capable of being manipulated to invoke outboard data displays qualifies as a glyph if and only if the icon includes some visible cue or mechanism for this extended data access. Lack of a visible cue or mechanism (button, hot spot, etc.) renders the invocation action a feature of the operating system or a subsuming software application - not a feature associated with the form of the visual element (icon) itself.

This visible cueing qualification is not as arbitrary as it may first seem. The point is to define a class of interface artifacts which are clearly distinct from simple icons. This distinction should be evident on the face of the interface artifact itself - even when viewed in isolation from the rest of the on-screen elements. This distinction also relates to the fact that one of the characteristics cited above for knowledge glyphs is their capacity for serving as 'micro-interfaces.' Hidden or secret features are poor interface design practice, and the ascription of fault to such features pertains to 'micro interfaces' as well as to conventional ones. As a result, the sort of pop-up 'tool tip' mechanism embedded in many Windows applications does not qualify an icon as a glyph under our definition.

Regarding 'Knowledge Glyphs'

The final step in this definitional exercise was to specify what makes a glyph a 'knowledge glyph.' For the purposes of the inaugural definition, the unique aspect of 'knowledge glyph' was framed with regard to the mode in which additional information on the denoted entity is afforded the user by virtue of its being a glyph. The provision of no more than descriptive textual data qualifies an icon to be a glyph. To qualify as a *knowledge* glyph, the visualization element has to be configured so as to provide a vantage on the denoted entity within a distinct referential context. As a result:

> *A knowledge glyph is a glyph affording its user the ability to access extraglyphic information in such a form that the extraglyphic presentation is anchored with respect to the same entity (or other discrete object of reference) denoted iconically by the originating glyph.*

Under this definition, then, a knowledge glyph implements an intersection among two or more distinct referential contexts in such a manner as to treat the denoted entity as the point of juncture among them all as well as the focal point of reference in the extraglyphic ones. This completes the three-level definitional framework illustrated in Figure 5.

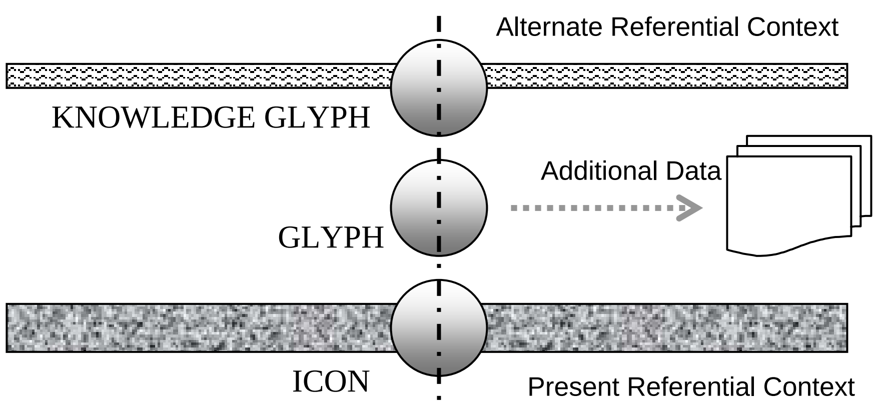

Figure 5: Overview of the Knowledge Glyph Framework

As illustrated in Figure 5, a visual element operates as an icon if it does no more than occupy a meaningful position in the present referential context (e.g., a point on a geospatial grid such as a map). This visual element escalates to the status of a glyph if it affords the viewer additional data either intra- or extraglyphically. The visual element qualifies as a knowledge glyph if and only if it portrays the referenced entity's meaningful position in an alternate referential context (e.g., a timeline rather than a geospatial grid).

The construction of a coherent model interrelating these three concepts constitutes the outcome of the abstract theorization work. This model can then be evaluated in terms of its ability to constructively describe, organize, and explain distinctions among visualization elements. Such model evaluation is the point at which our knowledge glyph research effort stands at the time of this writing.

Closing the Loop: Projecting Subsequent 'Best Practice'

The model developed in the theoretical work illustrates not only the form of the conceptual framework to be applied in assessing interface elements but also the nature of a use protocol amenable to exploratory implementation. Recall that the original intent was to provide a means for augmenting existing best practice to accommodate richer presentation of subject matter. The model in Figure 5 was crafted to 'telescope' in the sense of there being a coherent progression from icon through glyph through knowledge glyph along which any given visual element may 'escalate' in accordance with its capacity for affording the user additional data on the designated entity. This progression holds in the opposite direction, because a knowledge glyph can be 'de-escalated' to serve as a glyph, which in turn can be de-escalated into an ordinary icon.

This model opens up the possibility of crafting interfaces on which icons can be manipulated to escalate all the way to knowledge glyphs portraying the designated entity within an entirely distinct referential / presentational context *without the focal entity's representation being displaced from its on-screen location and hence without requiring the user to shift his / her gaze from one point on the display.* Phrased another way, our model provides a basis for describing interface design strategies that permit the user to focus on an object of reference and 'rotate' pertinent referential contexts into and out of view.

This ability in turn raises the prospect of interfaces permitting extended examination of specific entities while minimizing the number of distinct windows cluttering up the display. This prospect constitutes a next-generation 'best practice' in advanced information visualization which closes a cycle extending back to the examination of the previous best practice afforded by existing windowing environments and graphic protocols such as MIL STD 2525.

Conclusion

During the past year we have undertaken research on a concept labeled 'knowledge glyphs'. There are any number of ways we could have approached the subject matter, delineated research objectives, and evaluated success. In this case, we grounded what could have been extremely abstract theorization with respect to the state of the art in C2 visualization practice. We then constrained the course of theoretical research with respect to specific issues or themes discerned with the best practice or scope of best practices that extant state of the art entailed. We then demanded of ourselves that a model for these knowledge glyphs would have to be internally coherent, referentially comprehensive, and explicitly linked to the visual interface elements that were the foci of our interest. As of the time of this writing, this analytical and theoretical work has produced a model meeting these requirements. Our examination of this model has already suggested a subsequent state of the art in visualization technique linked to an analogous new state of the art in this technique's usage. In this way, we have 'circulated' from practice to theory and back again.

REFERENCES:

Calder, P. R., and M. A. Linton (2003). Glyphs: Flyweight Objects for User Interfaces. Stanford University Center for Integrated Systems report available WWW: http://www.fresco.org/docs/papers/glyphs.pdf

Chuah, M. C., and S. G. Eick (1998). Information rich glyphs for software management data. *IEEE Computer Graphics and Applications, 18:4* (July-August), 24-29.

Department of Defense (2005). *MIL-STD-2525B: Department of Defense Interface Standard: Common Warfighting Symbology*. Washington DC, Revised Change 1 edition, July.

Kindlmann, G. (2004). Superquadric Tensor Glyphs. In Deussen, O., Hansen, C., Keim, D., and D. Saupe (Editors), *Joint EUROGRAPHICS - IEEE TCVG Symposium on Visualization (2004).* Available WWW: http://www.cs.utah.edu/~gk/papers/vissym04/

Xerox PARC (2004). *Technical Overview of DataGlyphs®*. Report on research in progress at: http://www.parc.com/research/projects/dataglyphs/techoverview.html