

Towards a Computational Model of Sketching

Kenneth D. Forbus, Ronald W. Ferguson, Jeffery M. Usher

Qualitative Reasoning Group, Northwestern University

1890 Maple Avenue, Evanston, IL, 60201, USA

Voice: 847-491-7699; Fax 847-491-5258

Email: forbus@nwu.edu; ferguson@nwu.edu; usher@ils.nwu.edu

ABSTRACT

Sketching is a powerful means of interpersonal communication. While many useful multimodal systems have been created, current systems are far from achieving human-like participation in sketching. A computational model of sketching would help characterize these differences and help us better understand how to overcome them. This paper is a first step towards such a model. We start with an example of a sketching system (*nuSketch COA Creator*) designed to aid military planners, to provide context and a source of examples. We then describe four dimensions of sketching, *visual understanding*, *conceptual understanding*, *language understanding*, and *drawing*, that can be used to characterize the competence of existing systems and identify open problems. The issues involved will be illustrated by examples from our experience with *nuSketch*. Three research challenges are posed, to serve as milestones towards a computational model of sketching that can explain and replicate human abilities in this area.

Categories and Subject Descriptors

I.2.1[Artificial Intelligence] Applications and Expert Systems.

General Terms

Theory, Design

Keywords

Sketching; multimodal interfaces; intelligent front-ends to knowledge-based systems

1. INTRODUCTION

Person-to-person communication often involves diagrams, charts, and drawings on white boards and other shared surfaces. People point, mark, highlight, underscore, and use other gestures to help disambiguate what they are saying. Being able to use multiple modalities, i.e., speech and gesture, to communicate ideas is especially crucial for spatial information [1,5,7,16,27]. The ability to understand spatial representations, and to use them appropriately in dialogue, is a critical skill that we need to embed in software, in order to create systems that better understand the users they are interacting with.

We focus here on *sketching*, by which we mean a communication activity involving a combination of interactive drawing plus

linguistic interaction. The drawing carries the spatial aspects of what is to be communicated. The linguistic interaction provides a complementary conceptual channel that guides the interpretation of what is drawn and provides information that is not easily depicted spatially. Most people are not artists, and even artists cannot produce, in real time, drawings of complex objects and relationships that are recognizable solely visually without breaking the flow of conversation. The verbal description that occurs during drawing, punctuated by written labels, compensates for inaccuracies in drawing. Follow-up questions may be needed to disambiguate what aspects of a drawing are intended versus accidental.

There is now a substantial body of research on multimodal interfaces [24]. Sketching is clearly a form of multimodal interaction, but not all multimodal interactions are sketching. Many multimodal interfaces focus on placement of predefined entities, e.g., selecting a location, often via pointing (cf. [1,7]). Such selection operations require a minimal shared understanding on the part of the participants, and hence has provided a natural starting point for multimodal interface research. Work that comes closer to sketching (cf. [9,16,29]) incorporates more domain semantics, to increase the level of shared understanding. This progression suggests that to achieve the kind of flexible interaction that sketching provides in human-to-human communication, multimodal research will rely heavily upon, and even drive, AI research. This paper examines sketching in that light, to provide a framework for understanding the phenomena and suggesting new research directions.

The rest of this paper describes our progress towards a computational model of sketching. We start with an example, our *nuSketch* multimodal interface architecture, showing how it has been used to create a system for creating and reasoning about military courses of action sketches (*nuSketch COA Creator*). We then step back and describe a framework for sketching, motivated by a combination of constraints from computation and from cognitive science research. We end by identifying three challenges for research on sketching, as a potential way to benchmark progress in the area.

2. nuSketch: A MULTIMODAL ARCHITECTURE FOR SKETCHING

nuSketch is designed as a general-purpose multimodal architecture to support sketching. The best way to illustrate *nuSketch*'s abilities is through an example application. Military planners use a *Course of Action sketch* (COA sketch) when planning an operation. COA sketches express the gist of a plan, before many details, such as timing, have been worked out. Traditionally such sketches are created using acetate overlays on maps, or on paper starting with hand-drawn abstractions of critical

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IUI'01, January 14-17, 2001, Santa Fe, New Mexico.

Copyright 2001 ACM 1-58113-325-1/01/0001...\$0.00.

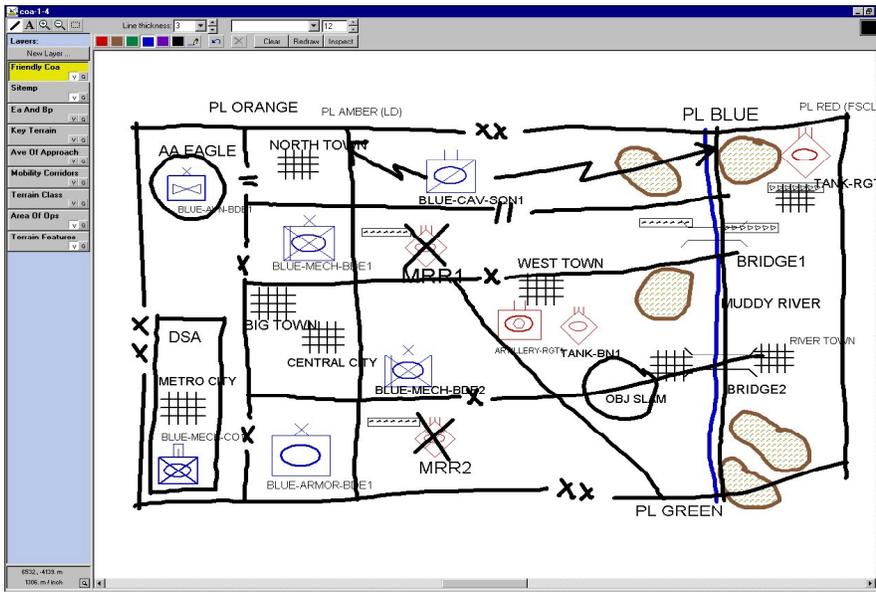


Figure 1: A Course of Action Sketch

terrain features. A well-worked out vocabulary of visual symbols is used to represent terrain features, military units, and tasks assigned to units.

Figure 1 illustrates a course of action drawn using the nuSketch COA Creator. A *layer* metaphor organizes the interface. Like acetate layers, each nuSketch layer corresponds to some category of domain information, such as terrain analysis, enemy disposition, disposition of your units, and so forth. Switching between layers is accomplished by clicking on the tabs to the left. Multiple layers can be displayed at once, or hidden, or grayed out as convenient. Maps as backgrounds for sketching are provided by bitmap layers, which are registered against geospatial coordinate systems.

While multiple layers can be visible at any time, only one layer at a time is *active*. The choice of active layer determines how user inputs are interpreted. For example, the Terrain Features layer enables planners to draw key properties of an area, such as rivers, mountains, cities, roads, and so on.¹ The Terrain Class layer enables users to mark regions that restrict vehicular movement (*trafficability*), due to factors such as slope, soil type, or vegetation. Items are drawn on layers via spoken commands (e.g., “Add severely restricted terrain”) accompanied by gestures, whose interpretations depend on the command. For adding regions, the curve drawn is taken to be the boundary of the region, so it is closed and filled with the appropriate texture to indicate that the command was understood. For adding standard symbols, e.g. towns or military units, the user’s gesture indicates a bounding box, and the appropriate glyph is retrieved from the KB and displayed there, scaled appropriately. Military tasks are more complex, requiring several gestures to specify which unit is doing the task, what the task is being done to, and often properties such as where it is being done and what path should be taken.

¹ Many planners prefer to work from this description, which abstracts out the significant features of an area, instead of the more cluttered view that a map background provides.

When a new glyph is created, a set of assertions constituting the system’s conceptual understanding of that visual element and what it represents in domain terms is also created. This conceptual understanding facilitates reasoning to support the user. For instance, geographic queries are made by dragging and dropping sketch elements onto a simple parameterized dialog (Figure 2). These queries are answered by using qualitative and visual reasoning to interpret the spatial entities and relationships in the sketch [12]. Similarly, users can request critiques based on analogies with prior plans, with the application of the advice to their plan illustrated by the system highlighting the appropriate visual elements of the sketch (Figure 3). The analogies are based on the visual and conceptual descriptions constructed during sketching, which are fed to a general-purpose analogy matcher, the Structure-Mapping Engine (SME) [10,14].

The cases are also created via sketching, with additional information about the relative wisdom of particular decisions stored in the case as a basis for providing advice.

The nuSketch COA Creator has been developed using a joint applications development model, using a combination of Army



Figure 2: nuSketch provides an interface to intelligent support tools, such as geographic reasoning. Results can be spoken or shown via highlighting on the sketch, as in this query about the shortest trafficable path that enemy tank battalion 1 can take to reach River Town.

interns on the development team and frequent formative feedback from other military personnel. It has been distributed to a small group of alpha testers², with very encouraging results. As a consequence, experiments with the system are planned for

² Our alpha testers include active US military personnel from several commands, retired military personnel associated with DARPA’s Command Post of the Future program (“graybeards”), and researchers from several institutions.

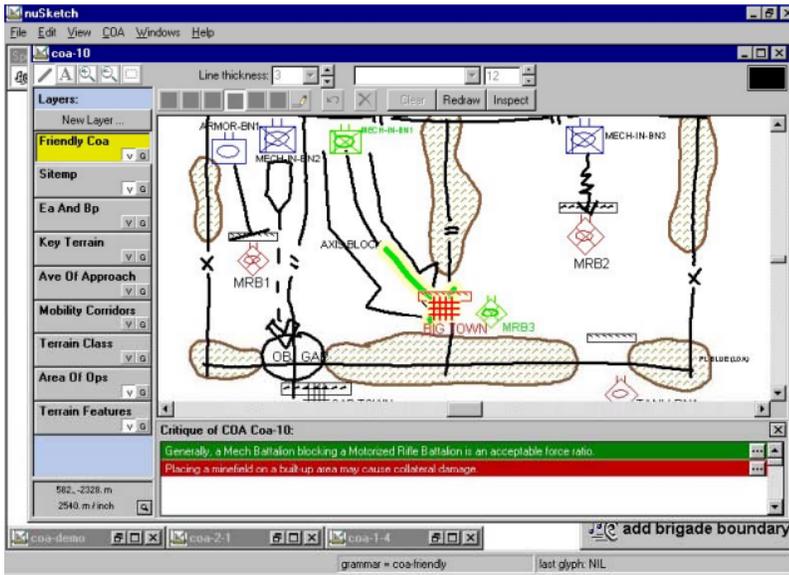


Figure 3: Analogies with prior COAs can be used to provide advice, with results displayed by highlighting elements of the planner's sketch. Here, the advice in green indicates a reasonable choice (i.e., an acceptable force ratio for the task) and advice in red indicates a problem with the COA (i.e., putting a minefield near a town)

November 2000 at the US Army's Battle Command Battle Laboratory in Ft. Leavenworth, Kansas. The COA Creator, in turn, is being used as a component in a larger application, a *Tactical Decision Coach*, a case-based instructional systems for tactics training that was made available to alpha testers in Fall of 2000.

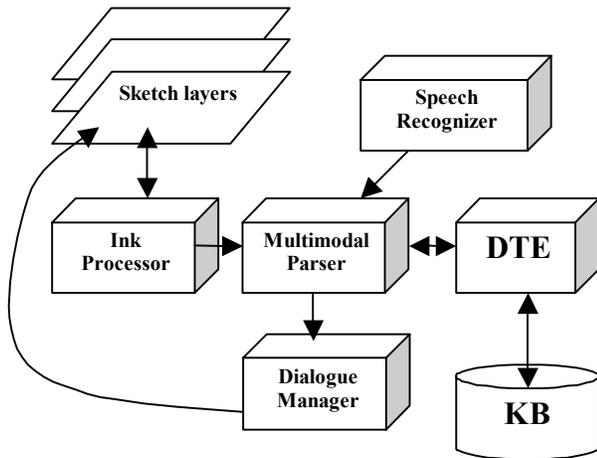


Figure 4: nuSketch architecture

3. The nuSketch Architecture

Figure 4 shows the nuSketch architecture. The Ink Processor accepts pen input, does simple signal processing, and passes time-stamped data to the Multimodal Parser. The other input to the multimodal parser is from a commercial speech recognizer, which produces time-stamped text strings. The Multimodal Parser uses

grammars that include both linguistic and gesture information, to produce propositions that are interpreted by the Dialogue Manager. The Dialogue Manager and the KB contents are the only application-specific components of nuSketch. The Dialogue Manager is responsible for interpreting propositions and supplying grammars to the speech recognizer and Multimodal Parser based on context (as determined by its own state and the active layer). Central to nuSketch is the use of a knowledge-based reasoner (DTE³), which provides integrated access to a number of reasoning services, including analogical reasoning and geographic reasoning. The Dialogue Manager uses DTE for its reasoning, and all domain-specific knowledge is stored in the KB. For example, the glyphs corresponding to the visual symbols in a domain are stored as part of the knowledge base, so that how something is depicted can be reasoned about (e.g., if a glyph is not available for a specific unit type, a glyph corresponding to a more general type of unit is used).

Several aspects of nuSketch are inspired by Quickset [7], a multimodal interface system for setting up military simulations. Like Quickset, we use off-the-shelf speech recognition and time-stamp ink and speech signals to facilitate integrating information across modalities. Quickset incorporates ink-recognition schemes that nuSketch does not (as a matter of principle; see below). Because Quickset was designed as an interface for legacy computer systems, it lacks an integrated reasoning system. As the discussion below will make clear, this significantly limits QuickSet's potential as a model of sketching. For example, it does not reason about depiction as nuSketch can.

4. Dimensions of sketching

The power of sketching in human communication arises from the high bandwidth it provides [27]. There is high *perceptual bandwidth* because the shared drawing is interpreted by the participants' powerful visual apparatus. There is high *conceptual bandwidth* because the combination of visual and linguistic channels facilitates the interaction needed to create a shared conceptual model.

Sketching covers a wide span of activities that occur under a variety of settings. A computational model of sketching must identify what knowledge and skills the participants need for such activities. We characterize these competencies along four dimensions: *visual understanding*, *language understanding*, *conceptual understanding*, and *drawing skills*. Variations along these dimensions determine how many different types of interactions something having those skills can participate in. We describe each in turn.

Visual understanding. This dimension characterizes how deeply the spatial properties of the ink are understood. The simplest level of understanding is recognizing gestures. Gestures indicate locations or sizes, often including an action to be taken with regard to something at that location (e.g., selecting or

³ DTE (*Domain Theory Environment*) is a reasoning system that combines a prolog-style query-driven inference system with a logic-based TMS to enable heterogeneous inference systems to interoperate.

deleting) [2,7,30]. We do not consider a system with only this level of visual understanding to be capable of sketching, since it does not understand the spatial relationships between visual elements.

The next level of visual understanding is the use of a visual symbology, i.e. a collection of glyphs, representing conceptual elements of the domain whose spatial properties can also convey conceptual meaning. Schematic diagrams in various technical fields and formal visual languages such as the military task language illustrated above are two examples. Is a CAD system a participant in sketching? We argue no, for two reasons. First, it is not taking an active role as a participant. In multimodal interactions, even during data entry the system is engaged in recognizing the kinds of entities and actions the user intends. Second, the time and conceptual overhead needed to deal with menus prevents the maintenance of a conversation-like flow [4,7,27]. By contrast, multimodal systems that use recognition procedures to “parse” ink automatically (cf. Quickset), or use speech plus gesture to create entities (cf. nuSketch) keep the interaction more like dealing with another person, someone capable of looking at what you are drawing and hearing what you are saying, and responding appropriately. The amount of visual

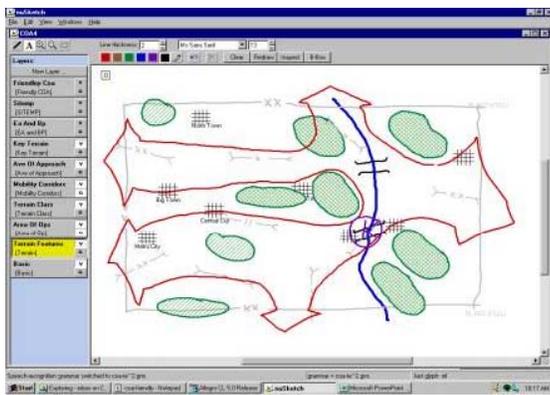


Figure 5: Visual symbols can be complex

reasoning needed in a sketching interaction depends in part on the desires of the participants. For example, to our surprise, the military personnel we have interacted with typically did not want automatic glyph recognition for hand-drawn unit symbols. To them, it was easier to just say what they wanted and use gesture for placement and sizing. However, we also discovered that, even in what seems to be a tightly restricted domain, human use of visual symbols requires richer visual processing than we had anticipated.

Most multimodal systems rely on a combination of speech and black-box recognition algorithms (e.g., hidden Markov models or neural nets) operating on digital ink to identify a user’s intent (cf. [7,16]). While certainly useful in some applications, we claim that they are detours from paths that will lead to human-like sketching capabilities. The reason is that people are very flexible in their use of visual symbols, and they expect the same flexibility from their partners. Three examples from the military planning domain illustrate this point. First, Figure 5 illustrates how complex visual symbols can be.



Figure 6: Two arrows

In terrain analysis, broad red arrows indicate avenues of approach. This multi-headed “arrow”, drawn by a military officer, accurately conveys where units might move. However,

it is hard to see how any statistical recognizer could be trained up in advance to recognize such a complex figure. The second example is Figure 6, which shows two arrows that are very similar, except for the style in which they are drawn. Most visual symbologies assign different meanings to dashed versus solid arrows, so it would not be enough to simply recognize both of these as arrows. The richness of visual properties that can arise even with very stylized visual symbologies is illustrated by the third example, Figure 7, which shows a pair of attacks that has been automatically identified as symmetric by high-level visual reasoning [11, 12]. Understanding such emergent configurations

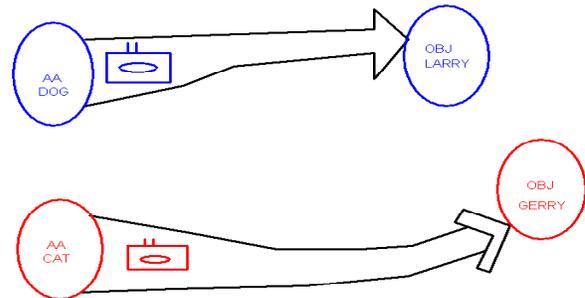


Figure 7. Higher-level visual constructs such as symmetry often convey important information

requires computing spatial relationships across multiple visual elements as they are drawn, and understanding the conceptual import of these relationships.

These examples suggest that powerful visual skills are one of the keys to human-like sketching. We suspect that work like Saund’s [28] on perceptual organization will play a major role in bringing sketching systems closer to human capabilities.

Conceptual understanding: As a communicative act, sketching requires common ground [3]; the depth of representation of what is sketched is probably the single strongest factor determining how flexible communication can be. To be sure, there must be enough visual and language understanding, and these can be traded off against each other somewhat, but it is the degree of shared conceptual model that ultimately limits what can be communicated, no matter what modalities are available. As might be expected, this is the weakest area for current systems.

The simplest level of conceptual understanding for sketching is the ability to handle a fixed collection of types of entities and relationships (cf. [6,7,16,30]). It is also the level most commonly used, since it suffices to issue commands to other software systems, the primary purpose of most existing multimodal interfaces. Type information is often used to reduce ambiguity, e.g., if a gesture indicating the argument to a MOVE command might be referring to a tank or a fence, the latter is ruled out.

Moving beyond identifying an intended command and its arguments requires broader and deeper common ground. Domain-specific systems (e.g., Quickset, particular nuSketch applications) obviously need knowledge about their domain. But there are areas of knowledge that cut across multiple domains of discourse that seem to be necessary to achieve flexible communication via sketching:

- *Qualitative representations of space.* Being able to reason about regions, paths, and relative locations is important in every spatial domain [13,8].

- *Qualitative representations of shape.* The ability to abstract away minor differences in order to describe important properties facilitates recognition [9].

We claim that qualitative representations are crucial for several reasons. First, they are well suited for handling the sorts of approximate spatial descriptions provided by hand-drawn figures, layouts, and maps. Second, the level of description they provide is close to the descriptions of continuous properties common in human discourse [15,29]. The nuSketch COA Creator, for instance, relies on qualitative representations to understand geographic questions and as part of the encoding of a situation that facilitates retrieval for generating critiques via analogy.

Other types of general knowledge are needed for flexible sketching as well:

- *Graphical conventions.* Many conventions used in drawings are deliberately unrealistic, e.g., cutaways to show the internal structure of a complex object. Using sequences of snapshots to depict dynamics requires interpreting spatial repetition as temporal progression. Understanding these conventions is necessary for many types of sketches [21].
- *Standard visual symbols.* Part of our shared visual language consists of simplified drawings that convey complex concepts easily. Stick-figure drawings and many other types of cartoons (cf. [25]) are examples.

Graphical conventions and visual symbols require combining visual/spatial knowledge with conceptual knowledge, and thus we suspect are a crucial area for improvement to create better sketching systems.

Language Understanding: Language provides several services during sketching. It can ease the load on vision by labeling entities and specifying what type of thing is being drawn. It is used for stating what spatial relationships are essential versus accidental, and describing entities and relationships not depicted in the drawing. Speech is the most common linguistic modality used during sketching because it enables visual attention to remain on the diagram, although short handwritten labels are often used as well. Existing multimodal systems tend to use off-the-shelf speech recognition systems, limiting them to finite-state or definite clause grammars (cf. [5, 6, 24]). Given the differences in complexity between spoken and written text, such grammars, albeit with multimodal extensions, are likely to remain sufficient [1]. Consequently, the most important dimension for characterizing language understanding in sketching systems concerns dialogue management [18, 21]. Most systems have been command-oriented, with some support for system-initiated clarification questions. We know of no sketching systems that use full mixed-initiative dialogs. We suspect two reasons for this. First, when multimodal interfaces are grafted onto legacy software, the existing output presentation systems are often used. Second, the relatively shallow conceptual understanding used in most systems does not enable them to do much on their own, so they are less likely to be able to interject anything.

Mixed-initiative dialogues will become crucial as the complexity of sketching increases. For example, we have discovered that the standard technique of using timeouts as the means for identifying when a user has stopped drawing a symbol (cf. [7,16]) fails completely when the spatial shape or configuration of a symbol must carry complex information. The avenues of approach arrows in Figure 5 illustrate this phenomenon, as do most complex paths drawn on a map. Consider what happens when the user is thinking

hard about a path, or tracing a path through complex terrain. Such drawings must be done carefully, since incorrect paths can have serious consequences. Interpreting sketched paths requires care, since it requires detecting and removing unintentional overlaps with terrain features that make the path as literally drawn impossible. The same system timeouts that keep a dialogue moving along when simple stereotyped glyphs are being drawn are far too short in these interactions, leading to intense frustration. It is essential to understand when the user is engaged in a complex drawing, requiring sophisticated interpretation, versus when the user is drawing something that can safely be interpreted as a simple glyph. As the scope of applications tackled becomes broader, e.g., using sketching in knowledge acquisition, the need for richer dialogue models becomes even stronger. For example, as the topics of sketches become more open-ended, the need to engage in clarification dialogues will be more frequent (e.g., “Is the specific position of the nucleus inside the cell significant?”).

Drawing capabilities: Sketching is a two-way street; ideally visual and linguistic expression should be modalities available to all participants (cf. [25]). The state of the art in natural language generation and text to speech is constantly improving, and such improvements will of course benefit sketching systems. Visual expression by sketching programs provides some new challenges. The simplest forms of visual expression are highlighting and performing operations on human-drawn elements (e.g., moving, rotating, or resizing). Some systems complement their human partners by neatening their diagrams (cf. [20]), while in other applications maintaining the informal, original ink is vital [21]. The ability to modify a user’s sketch as part of asking a clarification question, and generating new sketches to start a dialog, are beyond the present state of the art. Significant progress has been made on expressing the visual skills needed for graphical production tasks such as layout (cf. [24]), the key barriers for sketching are the lack of understanding of both the domain and visual representations, as outlined above.

5. CHALLENGES FOR COMPUTATIONAL MODELS OF SKETCHING

The discussion of the dimensions of sketching should make it clear that, while currently software can be built that participates in sketching in a limited way, the state of the art is far from creating systems that have the depth and flexibility of a human partner. In the spirit of encouraging progress, we suggest three challenges as useful benchmarks to measure progress in the area.

Integrated compositional semantics: The preponderance of systems that use “black-box” recognizers is more a function of them being easily available and of the limited range of tasks tackled to date than their suitability for use in sketching. An example provides the best illustration. The symbols used on military maps are highly standardized, with thick books providing visual symbols for almost every conceivable occasion. Nevertheless, during military exercises unique visual symbols are sometimes generated to cover special needs. Figure 8 shows an icon, drawn on a post-it, that appeared in several places on an intelligence map in a recent US Army exercise.



Figure 8: A novel, but easily understandable, visual symbol

This symbol represents a downed US pilot at its location. Although this symbol cannot be found in any military manual, it is

quite easy to interpret. Even non-military people tend to get it after one or two leading questions (what is the thing on the right? Okay, it's a crashed airplane. Who might that be?). There are degrees of ambiguity in the interpretation: Some people interpret the dashed lines coming out of the pilot's head as sweat rather than tears, and some think the person is a passenger. However, no one who is told what the symbol means has trouble identifying the airplane and the pilot and the pilot's unhappy/stressed state as a consequence of a hypothesized crash.

This seemingly simple interpretation problem requires an enormous amount of knowledge: Of airplanes, pilots, and their relationships, of the visual appearance of airplanes and how modifications of that appearance might be interpreted (i.e., a crashed airplane), and of conventions for depicting people and their states. The ability to combine visual and conceptual understanding in a compositional way to decode everyday sketches is, we believe, a major challenge for computational models of sketching.

User-extensible visual symbologies: As less restrictive domains are attempted, confining users to a predefined vocabulary of visual symbols will be infeasible. Asking a user to provide dozens to hundreds of samples to train a statistical recognition system in order to add a new glyph, for instance, is quite unnatural. Some systems add a mode that enables users to describe the specific pattern of strokes they will use to draw a new symbol (cf. [16]), but this method does not scale to complex glyphs (because there are too many ways to draw them), and does not provide a method for tying the new glyph to representations in a knowledge base. When a new glyph is introduced in human-to-human sketching, the introducer may have to linguistically mark its occurrence for a while when first used, but over time the other participants learn to recognize it. The ability to engage in a mixed-initiative dialogue seems essential for specifying the conceptual meaning of new symbols. Being able to interactively specify the domain semantics of a new glyph, and have the software start picking up how to recognize it through normal interactions, will be an important benchmark since it will enable the bootstrapping of sketching systems.

Visual analogies: Being able to compare sketches is an important aspect of comparing what the sketches are about (e.g., comparing engineering designs or comparing COAs). Shared history provides an important form of common ground, so the ability to recognize when aspects of the current sketch have been seen before will enable software participants to take on more of a community memory role. Currently there are domain-specific systems that do sketch-based retrieval [9,17], but these only operate in narrow domains. Some progress has been made on using similarity in visual encoding, particularly to detect symmetry and regularity in line drawings [11], but using these and other *analogical encoding* techniques in visual understanding is currently an area of active research.

Sketching systems that can carry out such visual analogies and retrievals in a broad range of domains will be another important benchmark in modeling sketching.

6. DISCUSSION

Sketching is a powerful human-to-human means of communication, and a powerful metaphor for human-computer interaction. We have argued that this power remains mostly untapped, i.e. that the state of the art is still far from creating software that participates in sketching with the same fluency as

humans. We argued that two key areas of improvement are depth of conceptual understanding and visual processing. The three challenges we outlined provide, we believe, benchmarks that would mark significant advances towards more human-like sketching systems. Even leaving aside its importance as an interface modality, research on sketching provides an arena for investigating the intersection of conceptual knowledge, visual understanding, and language, making it a valuable area for investigation in order to understand human cognition. We hope that this paper encourages more research in this area.

7. ACKNOWLEDGMENTS

This research was supported by DARPA under the High Performance Knowledge Bases, Command Post of the Future, and Rapid Knowledge Formation programs. Major Robert Rasch (US Army) and Captain Bill Turmel (US Army) played essential roles in the design and development of the nuSketch COA Creator. We thank Brigadier General Keith Holcomb (USMC, retired) and Colonel Michael Heredia (US Army) for many excellent suggestions. We also thank Herb Clark, Brian Dennis, Pat Hayes, Jim Mahoney, Stanley Peters, Eric Saund, and Bill Scherlis for insightful discussions.

8. REFERENCES

1. Allen, J.F. et al, The TRAINS Project: A Case Study in Defining a Conversational Planning Agent, *Journal of Experimental and Theoretical AI*, 1995.
2. Bolt, R. A. (1980) Put-That-There: Voice and gesture at the graphics interface. *Computer Graphics*. **14**(3), 262-270.
3. Clark, H. H. 1996. *Using language*. Cambridge University Press.
4. Clark, H.H. 1999. Speaking in Time. Proceedings of the ESCA workshop on dialogue and prosody, September 1-3, Veldhoven, the Netherlands.
5. Cohen, P. (1992) The role of natural language in a multimodal interface. *UIST92*, pp 143-149.
6. Cohen, P., Dalrymple, M., Moran, D., Pereira, F., Sullivan, J., Gargan, R., Schlossberg, J. and Tyler, S. (1989) Synergistic use of direct manipulation and natural language. *Proceedings of CHI-89*, pp 227-233.
7. Cohen, P. R., Johnston, M., McGee, D., Oviatt, S., Pittman, J., Smith, I., Chen, L., and Clow, J. (1997). QuickSet: Multimodal interaction for distributed applications, Proceedings of the Fifth Annual International Multimodal Conference (Multimedia '97), (Seattle, WA, November 1997), ACM Press, pp 31-40.
8. Cohn, A. (1996) Calculi for Qualitative Spatial Reasoning. In *Artificial Intelligence and Symbolic Mathematical Computation*, LNCS 1138, eds: J Calmet, J A Campbell, J Pfalzgraf, Springer Verlag, 124-143, 1996.
9. Egenhofer, M. (1997) Query Processing in Spatial-Query-by-Sketch in *Journal of Visual Languages and Computing* 8(4), 403-424 pp.
10. Falkenhainer, B., Forbus, K., Gentner, D. (1989) The Structure-Mapping Engine: Algorithm and examples. *Artificial Intelligence*, 41, pp 1-63.
11. Ferguson, R.W. and Forbus, K.D. 2000. GeoRep: A Flexible Tool for Spatial Representation of Line Drawings. *Proceedings of AAAI-2000*. Austin, Texas.

12. Ferguson, R.W., Rasch, R.A., Turmel, W., & Forbus, K.D. (2000) Qualitative Spatial Interpretation of Course-of-Action Diagrams. *Proceedings of the 14th International Workshop on Qualitative Reasoning*. Morelia, Mexico. June, 2000.
13. Forbus, K. 1995. Qualitative Spatial Reasoning: Framework and Frontiers. In Glasgow, J., Narayanan, N., and Chandrasekaran, B. *Diagrammatic Reasoning: Cognitive and Computational Perspectives*. MIT Press, pp. 183-202.
14. Forbus, K., Ferguson, R. and Gentner, D. (1994) Incremental structure-mapping. *Proceedings of the Cognitive Science Society*, August.
15. Forbus, K., Nielsen, P. and Faltings, B. "Qualitative Spatial Reasoning: The CLOCK Project", *Artificial Intelligence*, **51** (1-3), October, 1991.
16. Gross, M. (1996) The Electronic Cocktail Napkin - computer support for working with diagrams. *Design Studies*. 17(1), 53-70.
17. Gross, M. and Do, E. (1995) Drawing Analogies - Supporting Creative Architectural Design with Visual References. in *3d International Conference on Computational Models of Creative Design*, M-L Maher and J. Gero (eds), Sydney: University of Sydney, 37-58.
18. Grosz, B.J., and Sidner, C.L., "Attention, Intentions, and the Structure of Discourse", *Computational Linguistics*, 12:3 (1986).
19. Henderson, K. 1999. *On Line and On Paper: Visual representations, visual culture, and computer graphics in design engineering*. MIT Press.
20. Julia, L. and Faure, C. 1995. Pattern Recognition and Beautification for a pen-based interface. *Proceedings of ICDAR'95*: Montreal (Canada), pp 58-63.
21. Landay, J. and Myers, B. 1996. Sketching storyboards to illustrate interface behaviors. CHI'96 Conference Companion: Human Factors in Computing Systems, Vancouver, Canada.
22. Lakin, F. 1986. Spatial parsing for visual languages. In Chang, S., Ichikawa, T. and Liguemendes, P. (Eds.) *Visual Languages*. Plenum Press, NY.
23. Luperfoy, S. (editor) *Automated Spoken Dialogue Systems*. MIT Press (in preparation).
24. Maybury, M. and Whalster, W. 1998. *Readings in Intelligent User Interfaces*. Morgan Kaufmann.
25. McCloud, Scott. 1993. *Understanding Comics: The Invisible Art*. Kitchen Sink Press, NY.
26. Moran, D.B., Cheyer, A.J., Julia, L.E., Martin, D.L., and Park, S. (1997) Multimodal user interfaces in the Open Agent Architecture. *Proceedings of IUI97*. pp 61-68.
27. Oviatt, S. L. 1999. Ten myths of multimodal interaction, *Communications of the ACM*, Vol. 42, No. 11, November, 1999, pp. 74-81.
28. Saund, E., and Moran, T. (1995) Perceptual Organization in an Interactive Sketch Editing Application. *ICCV'95*
29. Stahovich, T. F., Davis, R., and Shrobe, H., "Generating Multiple New Designs from a Sketch," in *Proceedings Thirteenth National Conference on Artificial Intelligence*, AAAI-96, pp. 1022-29, 1996.
30. Waibel, A., Suhm, B., Vo, M. and Yang, J. 1996. Multimodal interfaces for multimedia information agents. *Proc. of ICASSP 97*.