CogSketch

Ken Forbus, Andrew Lovett, Kate Lockwood, Jon Wetzel, Camillia Matuk, Ben Jee and Jeffrey Usher

Qualitative Reasoning Group, Northwestern University

2133 Sheridan Rd

Evanston, IL 60208

{forbus, andrew-lovett, kate, jw, c-matuk, b-jee, usher}@northwestern.edu

Abstract

CogSketch is an open-domain sketch understanding system. CogSketch takes a unique approach to sketch understanding that focuses less on low-level recognition and more on high-level reasoning with sketches. In addition to demonstrating the basic system, we will showcase applications to cognitive simulation and education.

Overview of CogSketch

CogSketch[1] is an open-domain sketch understanding system built on the nuSketch architecture (Forbus & Usher, 2001). The idea behind nuSketch is that sketching is often a multimodal interaction where users use a combination of drawing and language to convey ideas. Many sketching systems rely on bottom-up recognition of the objects being sketched (e.g. Avarado & Davis, 2001). Such systems have been shown to be quite useful. However they often place restrictions on drawing conventions and are limited to domains for which recognition libraries have been constructed. We take a different approach, based on the observation that recognition is not essential for sketch-based reasoning. Humans are often not great artists in real time and rely on language to provide clues to the conceptual content of their sketches.

We provide a complimentary approach to recognition based sketching systems by allowing users to conceptually label objects they have sketched, telling the system what each object represents. There are two benefits to this approach: 1) by removing recognition, we are free to focus on deeper reasoning about objects in a sketch and 2) we do not place any domain or depiction constraints on users.

The basic unit of a CogSketch sketch is a *glyph*. Glyphs contain *ink* and *content.* The ink is the set of points drawn by a user and the content is the knowledge about what the glyph represents. The content consists of one or more conceptual labels chosen by the user. Conceptual labels are links to concepts in an underlying knowledge base. In CogSketch the knowledge base is currently a subset of the OpenCyc[2] KB containing over 1.8 million facts, over 58,000 concepts, and over 17,000 predicates. For a more detailed description of the nuSketch architecture underlying CogSketch see (Forbus & Usher, 2001; Forbus et al, 2004)

CogSketch automatically computes a number of qualitative spatial relationships between glyphs in a sketch. These include topological relations, relative position of glyphs, and relative size. CogSketch uses the topological relations to identify groups of glyphs that are connected or contained within a single outer glyph. For more on the spatial relationships computed in CogSketch see (Forbus & Usher, 2003).

While we rely on users to draw their ink as glyphs, CogSketch can also decompose glyphs into their component edges and build up structural representations of a glyph's shape, to allow deeper spatial reasoning. For example, the shapes of glyphs can be compared, and CogSketch can identify transformations between shapes, such as rotations and reflections.

Analogy and Similarity

A central feature of CogSketch is our use of analogical processing based on Gentner's structure mapping theory (Gentner, 1983). We use the Structure Mapping Engine (SME) (Falkenhainer, Forbus and Gentner, 1989) to allow users to compare objects in a sketch and detect similarities and differences. Analogies in CogSketch are based on both the visual and the conceptual material in a sketch. There is psychological evidence that structural alignment occurs in visual processing (Markman & Gentner, 1996; Gentner & Markman, 1997), and SME captures many aspects of this processing accurately. This provides a powerful mechanism for CogSketch to determine when things in a sketch will look alike to users.

Application Domains

Cognitive Simulation

CogSketch provides a convenient platform for simulating results from cognitive psychology experiments that use visual stimuli. If original stimuli were created in PowerPoint, they can be directly copied and pasted into CogSketch. For other formats, CogSketch can be used to resketch the stimuli. CogSketch and its predecessor sKEA

---

[1] CogSketch is available online at
  http://spatiallearning.org/projects/cogsketch_index.html

[2] OpenCyc is available online at http://www.opencyc.org/

have been used to simulate results from spatial language use (Lockwood, et al, 2006) and visual analogies (Tomai, et al, 2005). Current work is examining diagram interpretation and depiction conventions, simulation of spatial tasks that evaluate intelligence and spatial reasoning abilities in humans (Lovett, et al, 2007; Lovett, et al, 2008), and expanding work in spatial language and visual analogy.

Education

In many educational settings, sketching is used to test students on physical and spatial concepts or in the brainstorming phases of design work. We believe sketch-based educational software could have significant benefits for education at all levels. Electronic worksheets are a simple example. CogSketch includes a simple worksheet facility, where students can create sketches for assignments and get feedback on their work. The feedback is generated by comparing the student's sketch to a solution sketch, provided by a teacher or curriculum designer. The solution sketch is annotated with the relationships that are important for the student to reproduce. SME is used to compare the student's sketch to the solution sketch, and the differences it finds are used along with the annotations to generate feedback. Curricula and problems are loaded via external configuration files, making it possible for anyone to create their own CogSketch assignments.

In collaboration with instructors of the Engineering Design & Communications (EDC) course at Northwestern, we are exploring how CogSketch might be used to help their students. In EDC, small teams of students tackle design problems for real customers, often patients at the Rehabilitation Institute, who need specialized artifacts to help them handle everyday chores and gain independence. One of the biggest problems students have in the course is learning to use sketches to communicate with their teammates, their customers, and potential "funders", i.e., the instructors. We are using CogSketch to create a "crash-test dummy" which students will use to practice communicating their designs. Using specialized interfaces and the usual CogSketch facilities, students will express the intended behavior and functions of the designs. CogSketch will use qualitative mechanics (Kim, 1993) to reason through the sketch's possible behaviors and use SME to compare the behaviors it finds with the student's intentions. The differences will be examined to generate questions and advice for the student.

Data Collection

In many domains (e.g. spatial cognition, diagram understanding, visual analogy) sketching provides a natural medium for examining human cognition. We are working closely with colleagues in psychology and learning sciences to extend CogSketch to support collecting data from human subjects. For example, timing data for ink is automatically collected by CogSketch, and can be exported in a delimited file for analyses by other tools.

References

Alvarado, C and Davis, R. (2001). Resolving ambiguities to create a natural sketch based interface. *Proceedings of IJCAI-2001*.

Falkenhainer, B., Forbus, K., and Gentner, D. (1989). The Structure-Mapping Engine: Algorithm and examples. *Artificial Intelligence*, 41, pp1-63.

Forbus, K., Ferguson, R. and Usher, J. (2001). Towards a computational model of sketching. *IUI'01*,January 14-17, Santa Fe, New Mexico.

Forbus, K., Tomai, E., and Usher, J. (2003). Qualitative spatial reasoning for visual grouping in sketches. *Proceedings of the 17th International Workshop on Qualitative Reasoning*, Brasilia, Brazil, August.

Forbus, K., Lockwood, K., Klenk, M., Tomai, E., and Usher, J. (2004). Open-domain sketch understanding: The nuSketch approach. To appear in *AAAI Fall Symposium on Making Pen-based Interaction Intelligent and Natural*, October, Washington, DC.

Gentner, D. 1983. Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155-170

Gentner, D., & Markman, A. B. (1997). Structure mapping in analogy and similarity. *American Psychologist*, *52*, 42-56.

Kim, H. (1993). Qualitative reasoning about fluids and mechanics. Ph.D. dissertation and ILS Technical Report, Northwestern University.

Lockwood, K., Forbus, K., Halstead, D. & Usher, J. (2006). Automatic Categorization of Spatial Prepositions. *Proceedings of the 28th Annual Conference of the Cognitive Science Society*. Vancouver, Canada.

Lovett, A., Forbus, K., & Usher, J. (2007). Analogy with qualitative spatial representations can simulate solving Raven's Progressive Matrices. *Proceedings of the 29th Annual Conference of the Cognitive Science Society*. Nashville, TN.

Lovett, A., Lockwood, K., & Forbus, K. (2008). A computational model of the visual oddity task. *Proceedings of the 30th Annual Conference of the Cognitive Science Society*. Washington, D.C.

Markman, A. B., & Gentner, D. (1996). Commonalities and differences in similarity comparisons. *Memory & Cognition, 24*(2), 235-249.

Tomai, E., Lovett, A., Forbus, K., & Usher, J. (2005). A Structure Mapping Model for Solving Geometric Analogy Problems. *Proceedings of the 27th Annual Conference of the Cognitive Science Society*. Stressa, Italy.