Demo: How CogSketch uses Qualitative Representations and Analogy for Visual Reasoning and Learning

Kenneth D. Forbus
Qualitative Reasoning Group, Northwestern University
forbus@northwestern.edu

Abstract

Qualitative representations and analogical reasoning have been used to help explain and simulate aspects of human visual reasoning for decades. CogSketch combines many of these research results to provide a new level of modeling for human visual reasoning and for building systems that use visual reasoning for tasks. This demonstration will show how CogSketch has been used to model visual reasoning tasks, to help generate advice for students learning, and as a new data-efficient, incremental, and inspectable method for learning to recognize objects.

1 Introduction

CogSketch is an *open-domain* sketch understanding system. That is, most sketch understanding systems focus on recognition, limiting the kinds of sketches that they can understand. CogSketch, by contrast, computes representations of digital ink, including both representations commonly found in computer vision (e.g., edges and junctions) and in qualitative spatial reasoning (e.g., RCC8). The goal is for it to see sketches as people do, sufficiently well enough to provide human-like performance in tasks. This demonstration will show three aspects of CogSketch. First, it will show how CogSketch can perform simple visual reasoning in a way that fits human performance. Second, it will show how CogSketch can use analogy to provide helpful advice to students learning via sketching activities. Third, it will show how CogSketch is being used to develop new algorithms for learning to recognize objects, using qualitative representations with analogical generalization to provide data-efficient, incremental, and inspectable learning. This abstract summarizes each part of the demonstration in turn.

2 Simulating Human Visual Reasoning

CogSketch currently provides the best simulations of three visual reasoning tasks: Geometric analogies, an Oddity task, and Raven's Progressive Matrices (Forbus & Lovett, 2021). Figure 1 illustrates. These simulations exhibit human levels of performance, correctly predict what specific problems are easy or hard for people, and in some cases predict properties of human reaction times. The demonstration will focus on geometric analogies.

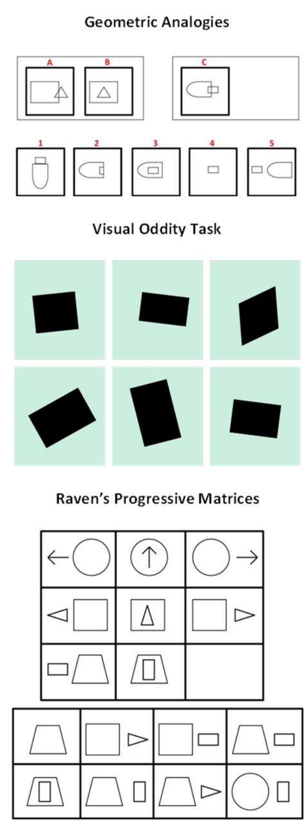

Figure 1: Some Human Visual Reasoning Problems CogSketch Simulates

3 Helping Students Learn

CogSketch was developed to be a platform for new forms of sketch-based educational software as well as an effective human-like visual reasoner. The most successful type of these systems are *sketch worksheets* (Forbus et al. 2017), which are sketch-based activities that are authored by domain experts and instructors. The idea is that the instructor has in mind a problem that can be solved by drawing (or modifying) a sketch. They sketch their intended solution using CogSketch, which then computes a variety of representations based on the visual properties of the instructor's ink and their conceptual labeling of it. The instructor chooses some of these facts to mark as important, either providing grading rubrics if they are not correct, or advice to give if they are not correct, or ideally both. When a student tackles a sketch worksheet, they attempt to solve the problem posed by drawing their own sketch or modifying a sketch provided as a starting point. Figure 2 illustrates a simple example. At this point dozens of sketch worksheets have been written by domain educators (Garnier et al 2017), many of them publicly available[1]. This demonstration will show how a simple sketch worksheet operates.

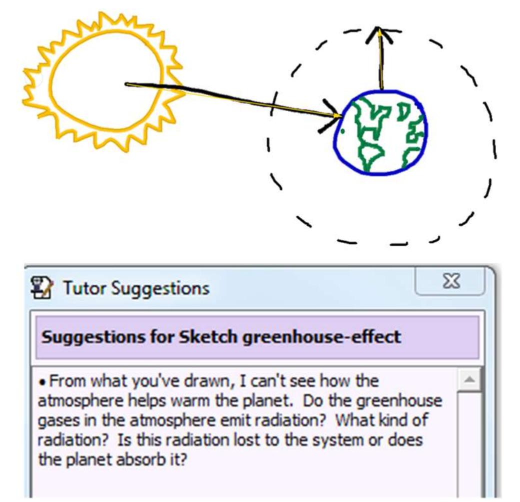

Figure 2: A Sketch Worksheet

4 Visual Recognition by QR and Analogy

Today's deep learning systems have revolutionized computer vision. Nevertheless, they require massive amounts of data and learn models that are not inspectable. Consequently, we are exploring the hypothesis that qualitative visual relationships plus analogical learning provides both better data-efficiency and interpretability for visual learning. We call our approach the *Hybrid Primal Sketch*, in honor of David Marr's (2010) original Primal Sketch. Figure 3 illustrates. The "Hybrid" in the name comes from the inclusion of deep learning modules for detecting object boundaries and bounding boxes, along with edge-finders from the computer vision community. These algorithms provide digital ink for CogSketch to process. When verbal labels are provided by a module, that information can be passed on using natural language resources built into CogSketch's knowledge base.

This idea has been explored in a number of experiments, including learning MNIST and learning visual relation detection. For the demonstration, we will use the Coloring Book

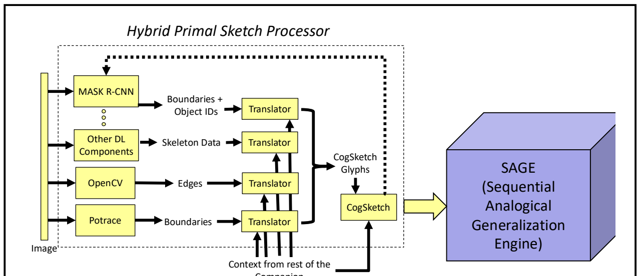

Figure 3: The Hybrid Primal Sketch

[1] https://serc.carleton.edu/serc/search.html?search_text=CogSketch&endpoint=%2Fserc%2Fsearch.html

Objects dataset (Chen et al. 2019), which contains only 10 examples per category total – a serious test of data-efficiency! Figure 4 illustrates some examples from this dataset, to illustrate how highly varied they are. Our latest visual learning algorithm is at 34% accuracy, whereas LeNet-5 is at chance.

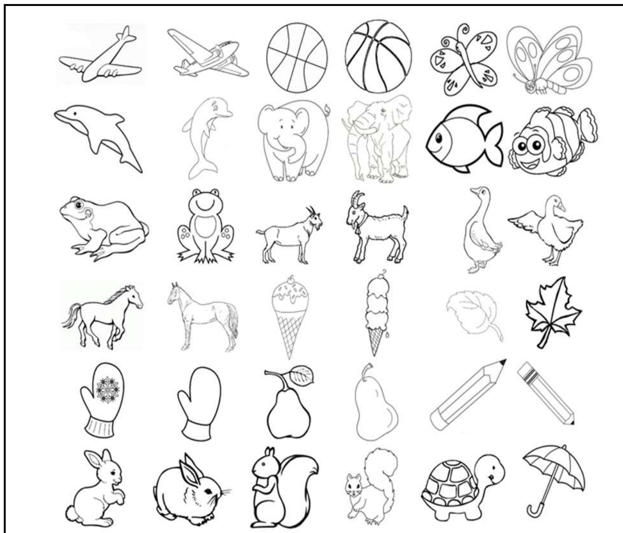
Figure 4: Some Examples from the Coloring Book Objects dataset

Acknowledgments

This research is sponsored by the US Office of Naval Research under grant number N00014-20-1-2447 and the US National Science Foundation.

References

Chen, K., Rabkina, I., McLure, M., & Forbus, K. (2019) Human-like Sketch Object Recognition via Analogical Learning. AAAI 2019.

Forbus, K. D., Chang, M., McLure, M. & Usher, M. (2017). The Cognitive Science of Sketch Worksheets. Topics in Cognitive Science. DOI: 10.1111/tops.12262

Forbus, K. & Lovett, A. (2021). Same/Different in Visual Reasoning. Current Opinion in Behavioral Sciences, 37:63-68, https://doi.org/10.1016/j.cobeha.2020.09.008

Garnier, B., Chang, M., Ormand, C., Matlen, B., Tikoff, B., & Shipley, T. (2017). Promoting Sketching In Introductory Geoscience Courses: CogSketch Geoscience Worksheets. Topics in Cognitive Science, 9:4(943-969). doi: 10.1111/tops.12291.

Marr, D. (2010). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. MIT Press, Cambridge MA.