

Research Statement: Andrew Lovett

I model the representations and processes that underlie spatial reasoning. This domain includes 2D visual problem-solving (Figure 1A), mentally transforming spatial representations (Figure 1B), and categorizing complex visual stimuli (Figure 1C). Such skills are important in academic and professional development—while problem-solving is commonly used to evaluate intelligence, the other skills are critical in many advanced disciplines, including geoscience, meteorology, chemistry, surgery, and dentistry.




Figure 1. Spatial reasoning tasks. A) Solve for the missing image in the 3x3 matrix. B) Which answer best depicts the cross-section of the 3D image? C) Can you recognize and characterize the fault in this slice of earth? B & C are geoscience examples from the Spatial Intelligence and Learning Center.

While other researchers have studied particular spatial tasks, I am cataloguing the representations and processes used across spatial reasoning, i.e., the primitives of spatial thought. My objectives are 1) to develop precise computational models of each primitive, and 2) to show how these primitives can work together to perform spatial tasks. This approach is useful for at least two reasons. Firstly, we can identify the base factors that make one person better or worse at spatial tasks; this could revolutionize education in many disciplines. Secondly, we can build computational agents that act as instructors or collaborators within spatial disciplines. An agent that sees an image the way a person does can better understand a person's mistakes and provide appropriate feedback.

My methodology is heavily interdisciplinary. I build computational models based on psychological theories of perception, visual comparison, and spatial problem-solving. The models focus on the interplay between symbolic and metric spatial representations. Symbolic representations are automatically computed from two-dimensional sketches and used to perform abstract comparisons between sketches. When necessary, the models return to a sketch's two-dimensional coordinate space to gather new information or compute spatial transformations. The use of sketches as input allows the models to focus on high-level perception and problem-solving, bypassing low-level visual problems like edge detection.

My models are evaluated by comparing their problem-solving performance to human performance on the same stimuli. In addition, the models generate novel predictions about human spatial reasoning. These predictions can be tested in new psychological studies, which can lead to revisions of the models and new predictions. At any point in this cycle, aspects of the models can be spun off as computational agents for educational software.

This approach relies on collaborators in several other areas. I work with psychologists to test model predictions and develop new theories. I also work with educators and software designers to develop applications for intelligent spatial agents. In the following section, I describe my current collaborations in these areas. In the future, I hope to develop new collaborations and expand the influence of this work.

Current Research

My current focus is on visual comparison and spatial problem-solving, primarily in two dimensions. For my thesis, I created Spatial Routines, a framework for modeling spatial problem-solving tasks. The premise of Spatial Routines, which is based on Ullman's [11] visual routines, is that we begin with a set of spatial operations, cognitive operations that any person can do. These are our spatial primitives. The operations can be combined in different ways to create a near-infinite range of routines.

The key operations are of three types: perception, visual comparison, and visual inference. *Perception* generates a spatial representation from a sketch. Representations are hierarchical: they can include information at different levels of abstraction. For example, a face might be represented as a single, holistic entity, or as a set of parts (eyes, nose, mouth) with spatial relations between them. Beyond that, one might focus on a single part (a nose) and represent its shape as the set of edges along its contour. I believe a key step in problem-solving is determining the level of abstraction that provides the necessary information.

In addition, spatial representations are hybrid: there are qualitative and quantitative components. The qualitative information is more abstract and thus more manageable: rather than representing an object's exact location, size, and orientation, it represents the object's size and location relative to other objects. Whenever possible, Spatial Routines uses qualitative representations to perform a task.

Visual comparison compares spatial representations to identify commonalities and differences. My model is based on the psychological claim that people compare images via structure-mapping [5][8]. Structure-mapping is a domain-general process of comparing two cases by aligning their common relational structure. I use the Structure-Mapping Engine [3], a well-established computational model.

Visual inference transforms a shape or image representation to produce a new one, such as the answer to a problem. For example, a routine could solve Figure 2A by applying the A/B differences to C to produce D', a representation of the solution image. Visual inference is my first pass at modeling *spatial visualization*, an important spatial skill (see Figure 1B).

I have built routines for solving three types of spatial problems: geometric analogy (Figure 2A), Raven's Progressive Matrices (Figure 2B), and the oddity task (Figure 2C). Because they are built in the Spatial Routines framework, I can determine precisely what is the same or different in the three models. They use identical spatial representations and similar comparison processes.

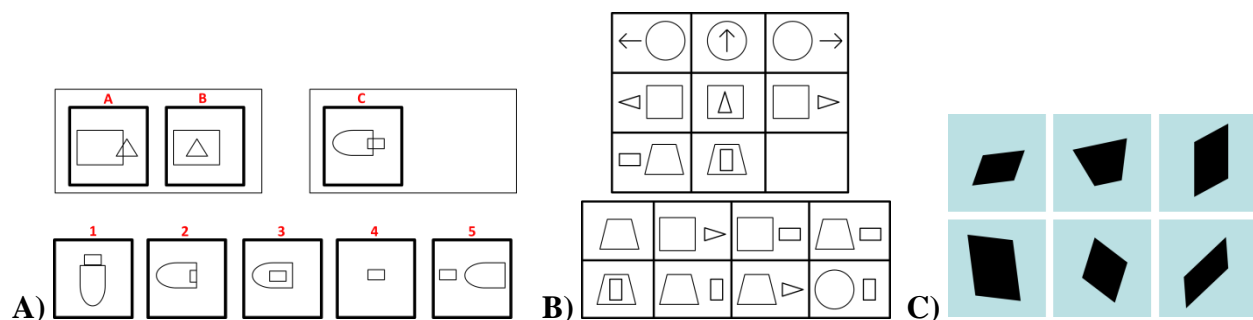


Figure 2. A) Geometric analogy: "A is to B as C is to..." [2]. B) Matrix problem: "Pick the answer that completes the 3x3 matrix." C) Oddity task: "Pick the image that doesn't belong." [1].

I have compared each task model against human performance. The models, which automatically compute representations and solve the problems, perform at least as well as adult humans. Additionally, problems

that are hard for the models are also hard for people. Finally, I have used the models to generate novel predictions about human performance.

For example, [1] originally gave the oddity task to two groups: North Americans, and the Mundurukú, a South American indigenous population. By modifying the oddity task routine and comparing it to each group, I determined a possible difference in the groups' spatial representations. The North Americans are more holistic, representing shapes or shape groups well, while the Mundurukú are more analytic, focusing on the parts within each shape [6]. This hypothesis suggests important cultural differences that could be examined in new psychological studies.

Collaborations

I have collaborated extensively with psychologists looking at visual comparison and the role of structure-mapping [7][10]. This collaboration involves both modeling comparison tasks and designing new psychological studies to test predictions of the model.

I have also collaborated with computer science researchers on CogSketch [4], a sketch understanding system. CogSketch is a platform for both cognitive science research and educational software. Our objective is to take the representations and processes from the research and build it into the software, creating intelligent educational agents.

One basic educational application in CogSketch is worksheets. Each worksheet requires students to sketch out a solution for a problem description—given a slice of earth, geoscience students might sketch and label the fault and marker beds. The student sketch is compared to the teacher's solution sketch, and CogSketch gives tutoring advice based on any differences. This technology relies on its spatial representations—if it represents an image the same way a person would, then it can identify the same differences that a teacher might identify. To facilitate this, we are transferring my models' representations into CogSketch worksheets.

Future Research Directions

Modeling Spatial Skills

Understanding spatial skills and how they can be taught is becoming a priority for educational research. The National Research Council [9] released a report stating, "Spatial thinking must be recognized as a fundamental and necessary part of the process of K-12 education." In 2006, the National Science Foundation funded the Spatial Intelligence and Learning Center, a multi-university, multidisciplinary center that studies spatial skills and spatial thinking. My dissertation research was developed within this center.

In my continuing research, I want to model new spatial skills by expanding my catalog of spatial primitives. One important skill type is spatial visualization: transforming spatial representations. Spatial Routines currently supports limited visualization, transforming a two-dimensional shape or image representation. However, advanced disciplines like geoscience and surgery require transforming between two and three dimensions (e.g., Figure 1B). I want to develop new spatial primitives to support this difficult task. This includes three-dimensional representations and processes for moving between dimensions, as well as processes for transforming within three-dimensional space.

Another skill type is recognizing complex visual categories (e.g., Figure 1C). This sometimes requires learning new representational primitives—for example, in geoscience it might be useful to encode a relation saying that one object lies within the convex hull of another object. To support this, I want to develop a visual routines framework to accompany the Spatial Routines framework. While Spatial

Routines reason over spatial representations, visual routines generate particular features in a spatial representation. For example, a visual routine might codify a particular strategy for determining that one object lies “between” two other objects. While Spatial Routines are grounded in spatial operations like perception and visual comparison, visual routines are grounded in low-level visual operations like curve-tracing and region-filling [11].

Learning Spatial Skills

Both Spatial Routines and visual routines codify established strategies. They do not describe how someone learns a new strategy. To address this important problem, I want to develop a third, metacognitive level. This level would support running a routine on a task, evaluating its results, and modifying the routine to improve performance. Importantly, this requires that spatial routines and visual routines each be coded in a declarative manner, so that the metacognitive level can reason about the operations and how they fit together.

Closing the Psychological Loop

My Spatial Routines for problem-solving have produced several novel predictions about human spatial cognition. Thus far, there have been few opportunities to test these hypotheses. In the future, I plan to expand my collaborations with researchers in the Spatial Intelligence and Learning Center, while developing new collaborations with researchers in such areas as vision, comparison, and problem-solving.

My objective is to close the psychological loop: build models, which produce hypotheses, which are tested on human participants, leading to refinements of the models. While I have some experience designing and running psychological experiments, I hope to develop new relationships with psychologists whose own research can benefit from computational models. As this work expands, cognitive neuroscience may also play an important role.

New Collaborations in Educational Software

In recent years, there has been increasing interest in educational software that thinks like people. For example, the Pittsburgh Science of Learning Center is developing *cognitive tutors*. Because these tutors model human thought processes, they can better determine how those thought processes have led to a misconception during learning. At the University of Southern California’s Institute for Creative Technologies, researchers develop human-like agents to support a more realistic experience in military training simulators.

In developing CogSketch, we have made important progress on spatial tutoring systems. However, there is still more work to be done. I would like to build cognitive tutors that provide step-by-step feedback on difficult spatial tasks, such as transforming between two and three dimensions. Such tutors could help college freshmen build a base of spatial skills to assist in various advanced disciplines.

This work requires collaborators for both developing educational software and evaluating software in the classroom. Fortunately, the National Science Foundation has shown a ready willingness to fund these types of collaborations. Both the Spatial Intelligence and Learning Center and the Pittsburgh Science of Learning Center are heavily invested in combining cognitive models with educational software. I plan to continue supporting CogSketch while pursuing new collaborations in the science of learning.

References

- [1] Dehaene, S., Izard, V., Pica, P., & Spelke, E. (2006). Core knowledge of geometry in an Amazonian indigene group. *Science*, 311, 381-384.

- [2] Evans, T. (1968). A program for the solution of geometric-analogy intelligence test questions. In M. Minsky (Ed.), *Semantic Information Processing*. Cambridge, MA: MIT Press.
- [3] Falkenhainer, B., Forbus, K. D., Gentner, D. (1989). The Structure-mapping engine: Algorithm and examples. *Artificial Intelligence*, 41, 1-63.
- [4] Forbus, K., Usher, J., Lovett, A., Lockwood, K., & Wetzel, J. (2011). CogSketch: Sketch understanding for cognitive science research and for education. *Topics in Cognitive Science*, 3(4), 648-666.
- [5] Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155-17
- [6] Lovett, A., & Forbus, K. (2011). Cultural commonalities and differences in spatial problem-solving: A computational analysis. *Cognition*, 121(2), 281-287.
- [7] Lovett, A., Gentner, D., Forbus, K., & Sagi, E. (2009). Using analogical mapping to simulate time-course phenomena in perceptual similarity. *Cognitive Systems Research* 10(3): Special Issue on Analogies - Integrating Cognitive Abilities, 216-228.
- [8] Markman, A. B., & Gentner, D. (1996). Commonalities and differences in similarity comparisons. *Memory & Cognition*, 24(2), 235-249.
- [9] National Research Council. (2006). *Learning to Think Spatially Spatially*. Washington, D.C.: National Academies Press.
- [10] Sagi, E., Gentner, D., & Lovett, A. (in press). What difference reveals about similarity. *Cognitive Science*.
- [11] Ullman, S. (1984). Visual routines. *Cognition*, 18, 97-159.