TOPICS
TOPICS IN COGNITIVE SCIENCE

The Cognitive Science of Sketch Worksheets

Kenneth D. Forbus,[a] Maria Chang,[b] Matthew McLure,[a] Madeline Usher[a]

[a]*Electrical Engineering and Computer Science, Northwestern University*
[b]*IBM T.J. Watson Research Center, New York*

Abstract

Computational modeling of sketch understanding is interesting both scientifically and for creating systems that interact with people more naturally. Scientifically, understanding sketches requires modeling aspects of visual processing, spatial representations, and conceptual knowledge in an integrated way. Software that can understand sketches is starting to be used in classrooms, and it could have a potentially revolutionary impact as the models and technologies become more advanced. This paper looks at one such effort, Sketch Worksheets, which have been used in multiple classroom experiments already, with students ranging from elementary school to college. Sketch Worksheets are a software equivalent of pencil and paper worksheets commonly found in classrooms, but they provide on-the-spot feedback based on what students draw. They are built on the CogSketch platform, which provides qualitative visual and spatial representations and analogical processing based on computational models of human cognition. This paper explores three issues. First, we examine how research from cognitive science and artificial intelligence, combined with the constraints of creating new kinds of educational software, led to the representations and processing in CogSketch. Second, we examine how these capabilities have been used in Sketch Worksheets, drawing upon experiments with fifth-grade students in biology and college students in engineering design and in geoscience. Finally, we examine some open issues in sketch understanding that need to be addressed to better model high-level aspects of vision, and for sketch understanding systems to reach their full potential for supporting education.

Keywords: Sketch understanding; Education; Analogy; Qualitative reasoning; Knowledge representation; High-level vision

Correspondence should be sent to Kenneth D. Forbus, EECS Department, Northwestern University Ford 3-313, 2133 Sheridan Road, Evanston, IL 60208, USA. E-mail: forbus@northwestern.edu

1. Introduction

Sketch understanding is an important scientific problem because it involves aspects of visual processing, spatial representations, and conceptual knowledge in a tightly integrated way. To illustrate, consider a simple drawing of three concentric circles (Fig. 1). In a science education context, they might stand for the layers of the Earth, a nucleus and two orbiting electrons, or the cross-section of a heat exchanger. The relationships between shapes and concepts in sketches is many to many, which is why people use labels (often verbal, sometimes written) when sketching with each other. These conceptual labels help participants interpret visual relationships in terms of spatial relationships, for example, in the layers of the Earth and the heat exchanger, the space inside the circles is filled with material, while for the orbiting electrons, it is empty (Lockwood, Lovett, Forbus, Dehghani, & Usher, 2008). These subtleties are what make sketch understanding such an important problem for understanding the interaction of vision and cognition. Better cognitive models of the representations and processes involved in sketch understanding can, in turn, lead to new kinds of educational software, combining the power of sketching to communicate spatial ideas (Ainsworth, Prain, & Tytler, 2011; Jee et al., 2014) with the on-demand feedback of intelligent tutoring systems (VanLehn, 2011).

This paper describes research on *Sketch Worksheets* (Yin, Forbus, Usher, Sageman, & Jee, 2010a), a new kind of sketch-based educational software system, from the perspective of the cognitive models that underlie them. These cognitive models are embodied in *CogSketch* (Forbus, Usher, Lovett, Lockwood, & Wetzel, 2011). The idea behind Sketch Worksheets is simple. An expert, that is, a teacher or curriculum designer, uses CogSketch to draw a solution to an exercise, such as drawing the layers of the Earth.[1] They segment their ink into entities as they draw, and label them. For example, the student might draw three circles to start with, as in Fig. 1, and label the innermost circle "Inner Core" and the outermost circle "Crust." These labels, along with labels for the other parts and distractors (here, "Lava" and "Rock"), come from a list provided by the worksheet author, grounded in an underlying knowledge base. CogSketch's visual system
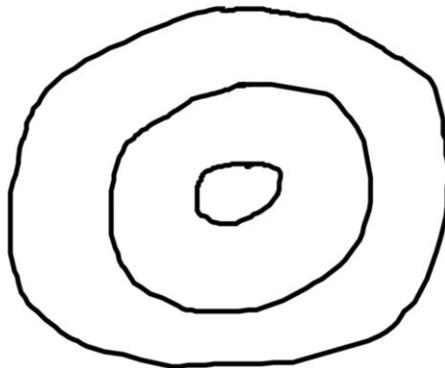

Fig. 1. These three shapes could mean many things, depending on context.

automatically constructs a variety of relationships between the visual entities (and, in some cases, decomposes them into finer-grained representations, as discussed below). The worksheet author marks which of these relationships are important and what advice to provide if the relationship does not hold. Here, for instance, the inner core must be geometrically inside the outer core, so the worksheet author enters natural language advice that will be shown to a student if his or her sketch does not meet this requirement. When students are doing a Sketch Worksheet, they are given instructions on what they should draw. In this worksheet, for example, the instructions are "Draw a sketch showing the layers of the Earth. Label the radius of the innermost layer." Students then start drawing their answer, asking for feedback at any time. When they ask for feedback, CogSketch compares what they have drawn to the solution sketch (exactly how this is done is detailed below), and if any of the important relationships do not hold, the corresponding advice is provided. For example, if they had labeled the three circles (starting from the inside) as Inner Core, Outer Core, and Crust, the Worksheet tutor would point out that they should consider adding the Mantle and the radius of the inner core. The student can then modify their sketch until they are satisfied. The author can also provide grading rubrics associated with the advice, and a gradebook provides tools for exploring the time course of a student's sketching as well as grading them.

Underneath this simple idea lies some sophisticated cognitive modeling. CogSketch must provide a set of visual relationships that is natural and rich enough to enable authors to select what they think is important. Its computation of these relationships must be robust enough that it can find them in student sketches, even if those sketches vary substantially in unimportant ways from what the author drew. Sketch Worksheets have been used by over 500 students to date in a variety of formative evaluations, across different age ranges and different domains (Garnier et al., this issue; Miller, Cromley, & Newcombe, 2014; Yin et al., 2010a). This paper focuses on how this work sheds light on the representations and reasoning that people are using in a sketch understanding. Instructors expect students to be able to see important spatial properties (or learn to do so), and hence understanding what representations and processes suffice to do this provides constraints that inform us about how human to human sketching, and cognition more broadly, works.

We begin by summarizing the cognitive models that power Sketch Worksheets, starting with the representations, outlining both the hierarchical visual representations CogSketch uses and the conceptual representations that tie visual entities and relationships to the world, thereby providing spatial meaning. We also describe how analogical comparison is used to find differences between a student's sketch and an expert's sketch that form the basis for generating advice.

2. Representations in sketch worksheets

CogSketch provides feedback by comparing a student's sketch to the instructor's sketch. But what properties matter? This section outlines the visual and conceptual representations that CogSketch uses.

2.1. Visual representations in sketch worksheets

There is evidence that relational qualitative representations are important in human visual processes, including categorization (Biederman, 1987; Marr & Nishihara, 1978), comparison and similarity judgments (Lovett, Gentner, Forbus, & Sagi, 2009a; Markman & Gentner, 1996), symmetry and regularity detection (Ferguson, 1994), and visual problem solving (Lovett & Forbus, 2011). Furthermore, these visual representations seem to be organized hierarchically (Palmer, 1977). As described below, CogSketch is capable of generating a hierarchy of relational representations, inspired by models of these visual phenomena in humans, to describe sketches. These levels are the object level, edge level, and group level (described in more detail below). Sketch Worksheets use two of the levels in this hierarchy (object level and edge level) to compare an instructor's sketch to a student's sketch. The rest of this section discusses (1) how symbolic entities are derived at these levels—that is, perceptual organization, and (2) how the set of qualitative attributes and relationships are computed over these entities to encode structured representations at each level—that is, perceptual encoding.

2.1.1. Perceptual organization

CogSketch enables users to draw digital ink with pen strokes, where each stroke produces a *polyline*, that is, a series of time-stamped points considered as a unit. These polylines can be manually split and grouped into conceptually meaningful entities called *glyphs*. The user interaction for defining glyphs is very flexible: Either they can start drawing and use the "Finish Glyph" button to indicate when they have completed a glyph, or they can use ink editing tools provided in CogSketch to merge or split their ink if they change their minds. In Fig. 1, for example, a student would have drawn three glyphs, one per circle, because they are intended to represent different entities. CogSketch's manual segmentation and labeling interface provides a stand-in for what happens in human to human sketching, where we talk and gesture while we sketch, identifying what we are drawing is intended to mean. CogSketch does not guess what a student intends a glyph to depict; it waits until the student tells it. This means that worksheet students (and authors) do a little more work, but it provides far more flexibility than recognition-based systems have.

Glyphs constitute the primitive elements in CogSketch's *object level* of representation. This is the default level of representation, and the majority of Sketch Worksheets use only this level of representation. CogSketch also can compute *group-level* representations, based on Gestalt principles (Lovett & Forbus, 2010), but to date these representations have not been used in Sketch Worksheets so we do not consider them further here. However, some Worksheets also use *edge-level* representations, which decompose an object into a network of edges and junctions. Edges are polylines that denote perceptually atomic ink segments. By atomic, we mean that they can be seen as having a uniform curvature, relative to some scale, and lack sharp inflection points (compare Fig. 2 left to middle). Inflection points and intersections between ink give rise to *junctions* that provide the end-points of edges. Visual properties of edges are computed as part of the
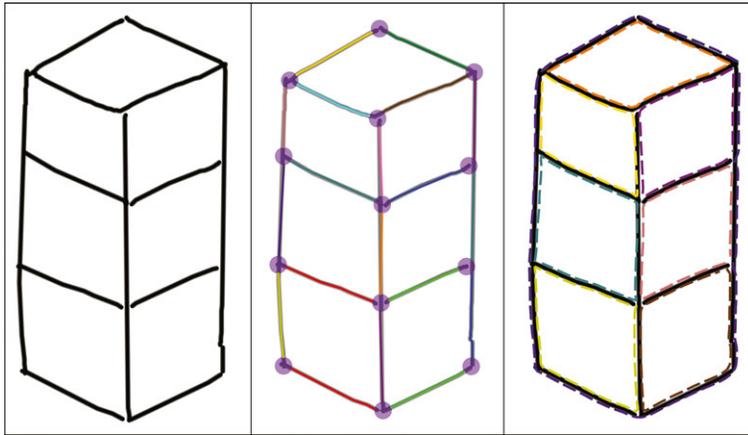
Fig. 2. Object, edges, and edge cycles for a simple glyph.

construction of edge-level representations. Junctions are the connections between the end-points of edges.

Edges form the most primitive level of representation, decomposing objects so that new visual entities can be constructed by grouping them according to various constraints. For example, edge-level representations enable CogSketch to construct *edge-cycles* (McLure, Friedman, Lovett, & Forbus, 2011; Fig. 2, right). An edge-cycle is a closed region in the network of edges, that is, the faces in its planar embedding. Such regions have been proposed in human vision research as an early organizational scheme (Palmer & Rock, 1994).

2.1.2. Perceptual encoding

By default, CogSketch only constructs object-level representations. Basic visual attributes of glyphs are computed automatically, such as whether they are open or closed, their major axis, and their size relative to the rest of the glyphs in the sketch. The key kinds of qualitative spatial relationships computed between them concern relative position, relative length, and topological relationships. Relative position includes rightOf and above, which are only computed between pairs of glyphs that are adjacent, as determined via a glyph-level Voronoi diagram.[2] Relative length includes shorterThan, longerThan, and sameLength. These three predicates use the major axis of a glyph to qualitatively compare its length to other glyphs, which is useful for comparing lines and arrows (e.g., in worksheets involving graphs and diagrams). For topology, we use RCC8 (Cohn, Bennett, Gooday, & Gotts, 1997), which classifies pairs of glyphs as disconnected, edge-connected, partially overlapping, equivalent, inside, or inside and touching. Combined with quantitative ink constraints (described below), this small set of representations is surprisingly expressive.

However, sometimes internal properties of visual objects are crucial in the feedback that is needed for a worksheet, for example, that the edges of an object drawn from a

viewpoint are all straight. When more detailed visual representations are needed in a worksheet, edge-level representations are computed for specific glyphs. Once edges and edge cycles are detected, CogSketch constructs structured relational representations characterizing them, like it does at the object level. Table 1 summarizes these representations. They include attributes that describe the relative length, curvature, and orientation of every edge. Relative length is stratified into four qualitative values. Curvature is described more coarsely, that is, straight versus curved. Orientation attributes label an edge with one of four qualitative directions, which vary depending on whether the edge is straight or curved. If a straight edge is aligned with the X or Y axis, it is labeled as *horizontal* or *vertical*, respectively. It also receives the more general *AxisAligned* attribute. If the edge is oblique, it is labeled with whether it points upward or downward, moving from left to

Table 1
Vocabulary for edge-level representations

Edge Attributes	Edge Relations
Length:	**Adjacency:**
`Tiny` `Short` `Medium` `Long`	`connected` `intersect` `intersects`
Curvature:	**Relative orientation:**
`Straight` `Curved`	`parallel` `perpendicular` `collinear`
Arc length:	**Positional:**
`MinorArc` `Semicircle` `MajorArc` `Ellipse`	`rightOf` `above`
Orientation:	**Cycle angles:**
(for straight edges) `AxisAligned` `Vertical` `Horizontal` `Oblique-Upward` `Oblique-Downward` (for curved edges) `LeftBumped` `RightBumped` `UpBumped` `DownBumped`	`convexCorner` `concaveCorner`
	Adjacent corner relations:
	`cycleAdjacentAngles` `acuteToObtuse` `obtuseToAcute`
	Corner attributes:
	`perpendicularCorner`

right. Every curved edge is labeled with the direction in which its "bump" points—that is, the normal of the curve at its apex, pointing in the convex direction—up, down, left, or right. Curved edges additionally receive attributes to reflect arc length. Curves that roughly complete 180-degree turn from beginning to end are labeled as semicircular. Curves that complete less of a turn are minor arcs, those that complete more are major arcs, and those that circle back on themselves such that their ends connect are labeled as elliptical.

There are several types of edge-level relations. Adjacency relations are the first to be computed because they constrain which edges are considered for other relations. Edges can be adjacent in three different ways: Their endpoints can meet at a junction (*connect*), they can cross at an X-junction (*intersect*), or the endpoint of one can lie along the other at a T-junction (*intersects*). Two straight edges in the same cycle can also be parallel or perpendicular. Positional relationships (e.g., *above*) are computed between all adjacent edge pairs. Convex–concave angle relationships are computed between all connected edge pairs in the cycle (note that the cycle itself is necessary to determine concavity).

With cycle angles represented as relations, consecutive corners in the cycle are related, using the binary higher order relation *cycleAdjacentAngles*, as well as the more specific *acuteToObtuseAngles* and *obtuseToAcuteAngles*. Right angles, considered particularly salient, are represented with a predicate that takes a cycle angle relationship as its single argument.

Edge-cycle representations (Table 2) include attributes to distinguish atomic cycles (cycles that contain no other cycles) from perimeter cycles (cycles that envelop the islands of connected ink). Each edge-cycle is assigned an attribute to reflect its relative area, another to reflect its relative complexity in terms of number of edges, and a third to label the orientation of its major axis as either horizontal, vertical, or oblique. Some edge-cycle attributes are aggregations of attributes of their parts; for example, an edge-cycle is straight when all of its edges are straight, or convex when all of its corners are convex. Two edge cycles are adjacent if they share at least one junction or one edge. Additional relationships are computed for adjacent pairs, for example, positional relationships.

2.1.3. Quantitative representation

The relational vocabulary described so far has been used in a number of cognitive simulations and performance systems, so there is ample evidence for their general utility in modeling human visual problem solving (e.g., Lovett & Forbus, in press; Lovett, Forbus, & Usher, 2010; Lovett, Tomai, Forbus, & Usher, 2009b; Lovett et al., 2009a). However, the task requirements of sketch worksheets led us to add a new kind of visual representation: *quantitative ink constraints* (Fig. 3). Some sketching exercises involve annotating a background image or diagram, such as the circulatory system worksheet at the upper right of Fig. 4. A correct answer for this worksheet requires not just drawing the chambers of the heart, but drawing them at the appropriate places on the diagram. To express such constraints, the worksheet author draws the glyphs in the places that they should be and marks them as quantitative ink constraints. They

Table 2
Vocabulary for edge-cycle-level representations

Attributes	Relations
Type:	**Adjacency:**
• Atomic	• shareJunction
• Perimeter	• shareEdge
Area:	• edgeSubset
• Tiny	**Relative orientation:**
• Small	• parallel
• Medium	• perpendicular
• Large	**Positional:**
Edge complexity:	• rightOf
• High	• above
• Medium	• enclosesVertically
• Low	• enclosesHorizontally
• VeryLow	**Special junction relations:**
Major axis:	• joinedAtJct-Y
• Vertical	• dominatesAtJct-T
• Horizontal	• joinedAtJct-T
• Oblique	• dominatesAtJct-Arrow
All edges:	• joinedAtJct-Arrow
• Straight	
• Curved	
• Ellipse	
• AxisAligned	
All corners:	
• Convex	
• Perpendicular	

also provide a tolerance that specifies how close a student's glyph must be to be satisfactory. When computing relationships between a quantitative ink constraint and the student's glyph for it (identified via analogical mapping, described below), CogSketch computes a qualitative characterization of problems with the match which can be used for providing more guidance. That is, CogSketch can provide either a general suggestion (e.g., "Your drawing of the left atrium is in the wrong location") or it can use qualitative representations for relative position to give the student more specific advice (e.g., "Your drawing of the left atrium is too far to the left.") If the location of individual edges is important, as is often the case with orthographic projection worksheets (bottom right, Fig. 4), then quantitative ink constraints can be defined at the edge level as well as the object level.
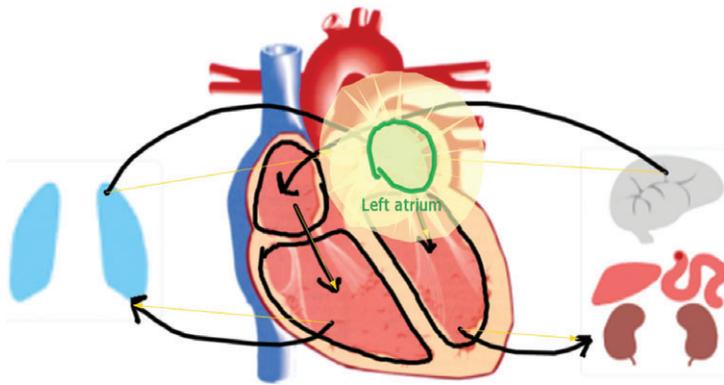
Fig. 3. The visual display for an object-level quantitative ink constraint on the left atrium. The worksheet author defines a numerical tolerance, which CogSketch displays over their glyph (in yellow above).
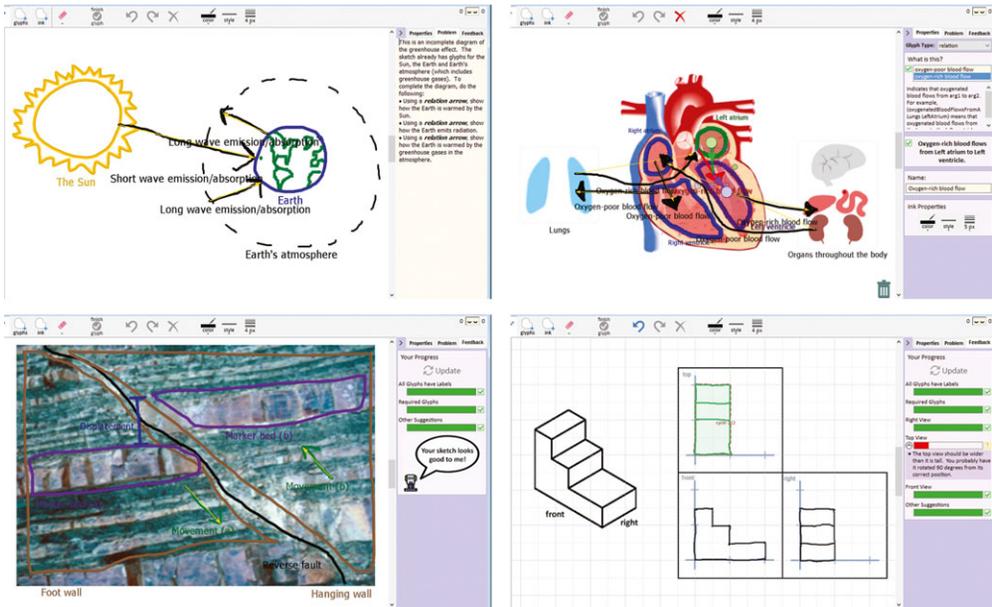
Fig. 4. Four sketch worksheets used as examples in this paper. Clockwise from top left: greenhouse effect worksheet, circulatory system worksheet, orthographic projection worksheet, and geological fault identification worksheet (containing a diagonal fault line and two rectangular polygons indicating marker beds).

2.2. Conceptual representations in sketches

When people sketch, they typically say what the sketched entities are intended to depict. Recognition can catalyze this process; for example, stick figures of humans in standard poses are easy for people to recognize. But for most entities in most domains, there are not standard visual symbols. Hence, CogSketch provides an interface that lets

users label their glyphs, drawing on concepts from the OpenCyc ontology. The OpenCyc ontology is quite large, over 58,000 concepts, providing a broad set of material to work with. This richness can be daunting, so it is not exposed to students. For Sketch Worksheets, the worksheet author selects a small set of concepts that are relevant to likely student sketches (including distractors and likely misconceptions), so that students just pick from a small list.

For education, having students explicitly label their sketches with their intended concepts is important for a second reason. An expert seeing an unlabeled student sketch like Fig. 1 might think it is correct because he or she is interpreting the glyphs differently from the student. In drawing the layers of the Earth, for example, the sketch of a student who swapped the mantle and crust would look the same as the sketch of a student who got it right: Both sketches consist of a set of concentric circles. Conceptual labels make the intended meaning of the students' ink clear.

In addition to visual relationships, conceptual relationships are often crucial in communicating one's understanding via sketching. A common convention in concept maps is to indicate binary relationships via labeled arrows. CogSketch supports this convention. For example, the arrows in Fig. 3 indicate blood flow between parts of the heart and the rest of the circulatory system. The possible labels are chosen from a set provided by the worksheet author, using binary relations from the OpenCyc ontology.[3] When a student tells CogSketch that a glyph indicates a relationship, CogSketch assumes that the glyph is an arrow and attempts to ascertain which part is the head and which part is the tail. It then uses these hypotheses to find the closest glyph that satisfies the constraints on the relationship's arguments (if its label is already known) and shows the student a sentence corresponding to its conjecture about his or her intent. If CogSketch is correct, the student doesn't need to do anything. Otherwise, the student can easily correct it by editing the head and tail with drag and snap icons, explicitly choosing the relationship arguments from a drop-down menu, or editing the glyph to make the head and tail clearer.

The third type of conceptual representations used in CogSketch are *annotations*. Annotations are graphical indicators of relevant properties of a glyph. For example, annotations are used to provide quantitative information about spatial or physical properties—the width of a chasm, the rate of a flow—that would not be convenient to indicate spatially. Drawing things to scale requires considerable drafting and artistic skill, and it is often impossible for the kinds of entities that one wants to consider together in a sketch (e.g., a person and a planet). There are several kinds of annotations. *Quantity annotations* provide a quantitative value. In a worksheet about the Carbon Cycle, for example, flows between different reservoirs in the system are represented by relationship glyphs. Each flow arrow is annotated in turn to indicate how many petatons of carbon the student believes is involved in that flow annually. *Force annotations* have an associated direction and are used to indicate applied forces or global forces like gravity. (Supplying a quantitative value is optional with force annotations.) *Direction annotations* indicate directions of linear or circular motion, again by CogSketch interpreting what is drawn as an arrow and inferring the appropriate direction.

2.3. Relationships used in sketch worksheets so far

When a worksheet is created, the worksheet author draws an ideal solution to the worksheet and identifies correctness criteria for the sketch. The correctness criteria can include any of the representations computed by CogSketch. With the exception of quantitative ink constraints, the visual and conceptual relationships described in sections 2.1 and 2.2 are presented to the author as natural language facts that can be inspected and marked as important. If a fact is marked as important, it is included in the correctness criteria for the sketch. Quantitative ink constraints are defined separately because, unlike visual and conceptual relationships between existing glyphs, there is no cognitive model for automatically proposing them, so they must be defined explicitly by the author for each applicable glyph. Table 3 shows what representations have been used as correctness criteria to date.

There have been 49 different worksheets used in classroom settings so far. These worksheets span three general domains: biology (4), engineering (19), and geoscience (26). The four biology worksheets involve drawing and labeling parts and flows of the human circulatory system and use a combination of conceptual relationships (e.g., blood flows) and object-level quantitative ink constraints for correctness criteria. The engineering worksheets involve exercises in free body diagrams and drawing orthographic projections. Although the engineering worksheets are within the same general domain, they use different types of representations. The free body diagrams use force arrows to convey

Table 3
How often specific relationships were used in correctness criteria for classroom worksheets. For quantitative constraints, how they are used is summarized

Representation	Example Meaning(s)	Count
Conceptual relationship or attribute		129
Biology	`bloodFlowsFromAndTo` `biologicTransfer-Oxygen`	22
Physics	`forceAssumed`	37
Geoscience	`emitsLongWaveRadiationTo` `photosynthesis-CarbonTransfer` `transMotion*`	70
Qualitative spatial relationship or attribute		112
Positional	`above` `rightOf`	33
Shape- or edge-level description	`cyclesShareJunction-Y` `longerThan` `parallelElements` `perpendicularCorner`	79
Quantitative relation or attribute		197[a]
Object-level ink constraint	Location of a heart chamber or fault line	182
Edge-level ink constraint	Location of an edges in orthographic projection	280
Quantity annotation	Associated quantity	15

Note. [a]Each edge-level ink constraint is itself part of an object-level constraint, so the total number of quantitative representations used is the sum of object-level ink constraints and quantity annotations.

assumptions about forces. The force arrows are represented qualitatively because the students are not required to enter numerical values for them. Instead, the arrows they draw need to be oriented in the correct qualitative direction, and they need to have the appropriate relative length. In contrast, worksheets that involve orthographic projection use edge-level representations to capture student drawings with greater detail. This allows the worksheet author to identify correctness criteria at the level of individual edges, either qualitatively (e.g., that two edges need to be perpendicular to each other) or quantitatively (e.g., that an edge needs to be in a specific location). In turn, the worksheet tutor can provide students with more detailed, targeted feedback. Thus far, the widest variety of worksheets has been developed for the geosciences, which use multiple object-level quantitative ink constraints and conceptual relationships to determine whether or not a student's sketch is correct. Many of the geoscience worksheets involve annotating a background image to identify relevant geologic features and any conceptual relationships that might hold between them. Quantitative ink constraints are useful for such worksheets, as are conceptual relationships that can be used to convey movement and physical transfer (e.g., carbon transfer). In many cases, both types of representations are used in a single worksheet.

3. Processing in sketch worksheets

This section describes the visual and analogical processing underlying Sketch Worksheets. Importantly, Sketch Worksheets do not do domain-specific reasoning. We have, in other experiments, modeled such reasoning (e.g., geologic interpretation [Yin, Chang, & Forbus, 2010b], conceptual physics problems [Chang, Wetzel, & Forbus, 2014]), and qualitative reasoning about mechanics within CogSketch has also been used in a system aimed at helping engineering students learn to explain their sketches (Wetzel & Forbus, 2010). By avoiding domain-specific reasoning, Sketch Worksheets can be used across a wide range of domains. However, this does mean that their visual processing and comparison capabilities must be robust and human-like enough to enable CogSketch to see students' sketches like an instructor might. This section explains how that works.

3.1. Visual processing in CogSketch

The goal of CogSketch's visual processing routines is to produce structured, relational representations that correspond reasonably well to what people are using, so that the visual distinctions that experts find relevant are detected reliably. Given the differences between today's computers and human visual processing capabilities, we have not attempted to model the details of human visual processing. Instead, we have aimed for what we hope is reasonable input/output fidelity.

To avoid computing potentially irrelevant information, only a subset of the spatial relationships described in section 2.1 are computed by default. CogSketch automatically computes topological relations between all glyphs in a sketch. Positional relations between

glyphs are automatically computed between adjacent glyphs only, as noted in section 2.1.2. The rest of the visual relations can be computed on demand, depending on whether or not they are needed to understand the contents of the worksheet.

As described in section 2.3, the worksheet author identifies properties that are important for drawing the sketch correctly. These properties may include spatial relationships that are not computed by default. To gain access to those relationships, the worksheet author can ask CogSketch to elaborate its representations of the sketch. For example, in orthographic projection worksheets (bottom right, Fig. 4), important properties of the solution need to be described using edge-level representations. The author can request that more detailed representations be computed so that some of them may be marked as important. Crucially, when the student's sketch is evaluated with respect to the solution, the elaboration requests made by the author are used on the student sketch to ensure that the information needed to evaluate the student's sketch is available.

3.2. Comparing sketches to generate advice

When a student requests feedback, his or her sketch is compared to the solution sketch via analogy. This requires a comparison mechanism that can support the wide range of student sketches possible. The student may request feedback from the software at any time, so his or her sketch may be nearly complete (with one or two differences from the solution) or it may be wildly different (with incorrectly labeled items or missing glyphs). It is therefore important that the comparison process focuses on relevant visual and conceptual information, even when there are many differences.

We use the Structure Mapping Engine (SME, Falkenhainer, Forbus, & Gentner, 1989; Forbus, Ferguson, Lovett, & Gentner, 2016) to find important differences between the solution and student sketch and to offer advice based on those differences. SME takes as input the solution sketch and the student sketch and creates a mapping between the two. Each mapping consists of (1) *correspondences*, which indicate how things in the solution correspond to things in the student sketch; (2) *candidate inferences*, which indicate important differences between the solution and the student sketch; and (3) a *structural evaluation score*, which is used to select the best mapping if there is more than one. The correspondences and candidate inferences are used to determine what advice should be given to the student.

The correspondences of the mapping are used to generate advice based on violated quantitative ink constraints. If a glyph in the solution sketch has a quantitative ink constraint associated with it, and it has a corresponding glyph in the student sketch, then the student's glyph is evaluated with respect to the solution glyph's quantitative ink constraint.

The candidate inferences of the mapping indicate visually and conceptually salient differences. If one of the differences involves a fact that was marked as important by the worksheet author, then it triggers the natural language advice string associated with it. Keying advice to differences in correctness criteria means that only relevant visual aspects of a sketch matter. This provides considerable flexibility and robustness. Consider

again the Greenhouse Effect worksheet (Fig. 4, top left), where the purpose is to see whether the student understands the paths and types of thermal radiation in the Sun heating the Earth. The relative positions of the Sun and the Earth don't matter. The Earth could be drawn as seen from a distance, for example, a sphere, or from a closer view. Since the relationships here are all conceptual relationships, and there aren't spatial relationships among the key facts for the worksheet, it could even be drawn as a concept map! As long as the entities and relationships are correctly labeled, it means that the student understands it (Fig. 5). This general approach works well in most cases. However, in expanding to wider range of tasks, we have found that the matching process is even more robust if we constrain analogical mappings with conceptual and quantitative information, and use recursive, hierarchical matching for different levels of representation. We discuss these in turn.

3.2.1. Match constraints

Once pre-comparison reasoning is complete, the analogical mapping between the solution and student sketch is constructed. In addition to the constraints dictated by structure-mapping theory, *match constraints* can be defined to further guide comparison. Comparisons in sketch worksheets are within-domain analogies: Things of the same type should match to each other. For this reason, we use *within-partition constraints,* which are automatically created based on conceptual labels entered into the sketch. Each label is considered a separate partition, and cross-partition matches are not allowed. For example, in the fault identification worksheet (Fig. 1 bottom left), glyphs representing a fault can only match with other glyphs representing a fault. Where there are multiple glyphs of the same type (e.g., two marker beds), then the spatial representations derived by CogSketch guide the matching process. In most cases, the qualitative visual relationships described in section 2.1.2 are sufficient for constructing an accurate mapping. However, in worksheets where the student is drawing on a background image or diagram, we have found that quantitative ink constraints can provide evidence for putting two particular entities into correspondence. In other words, given more than one possible mapping, we prefer matches where quantitative ink constraints are satisfied. This is done by looking at the competing correspondences constructed by SME, and evaluating them with respect to quantitative matching, when available. Close quantitative matches are translated into additional match constraints requiring those entities to be placed into correspondence in subsequent comparisons. This allows the matching process to, in effect, take quantitative spatial information into account when it is relevant to the task. We have found that this technique significantly improves the accuracy of matches on a corpus of worksheets where students draw on a background image (Chang & Forbus, 2012).

3.1.2. Hierarchical matching

Some worksheets, such as engineering drawing worksheets (e.g., Fig. 4 lower right) involve using more intricate edge-level representations. Carrying out matches at multiple levels at once can lead to mismatches, and even when matches are perfectly accurate,
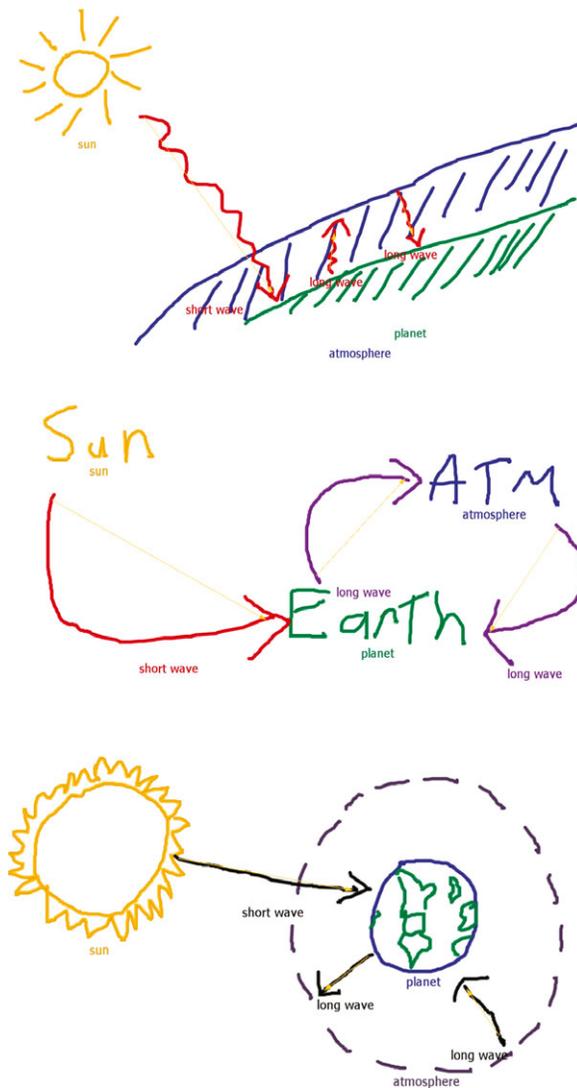
Fig. 5. Greenhouses. These three sketches all are valid solutions to the Greenhouse effect worksheet. Cog-Sketch views them as equivalent because they all satisfy the properties the author specified as important.

deluge students with too much advice. To increase accuracy in mapping, we use matches between more abstract representations to guide matches involving more detailed representations. For example, the engineering drawing worksheets (Fig. 4, bottom right) provide three regions (top, front, left) for students to draw their different views in, indicated via special glyphs, called *ink-segmenting glyphs*, in the author's solution. This causes CogSketch to restrict comparisons to glyphs found in the same ink segmentation glyphs. At the level of glyph matching, if tutoring advice is generated, then that advice is given to

the student and no further attention is paid to that pair. Otherwise, the worksheet tutor drills down to the edge level of representation. It first looks for differences in edge cycle properties, for example, if it should be straight but isn't, that are worth addressing. When the edge cycle level representations provide no guidance, comparisons between particular edges for that cycle are carried out, to give more detailed feedback when the author has provided important distinctions at that level.

3.1.3. Misconception sketches

Sometimes there are canonical misconceptions in a domain. For example, in identifying the four chambers of the heart (Fig. 4, top right), a students who have right and left confused in their sketch would benefit more from feedback that indicates the reversal, rather than simply being told that they are incorrect. To enable authors to communicate such misconceptions and provide advice based on them, sketch worksheets also support *misconception sketches* in addition to the solution sketch. When a worksheet has misconception sketches, the tutor first checks if the student's sketch was a strong match for a known misconception. If so, it provides advice specific to that misconception that was provided by the author. This provides the opportunity to give much more informative feedback.

4. What knowledge and skills can be tested or learned via sketch worksheets?

Our research on the utility of sketch worksheets is ongoing, but we already have some insights as to what types of activities have been successful, based on our experience in both observing students in laboratory experiments and in working with instructors and students in classroom experiments. Sketch worksheets have shown promise in helping students learn spatial layouts, conceptual relationships that can be conveyed visually, and scaling when qualitative differences in quantity can be readily observed and aligned visually. In contrast, sketch worksheets do not appear to be the right platform for exercises where quantities or relationships are represented abstractly and not grounded visually.

The worksheets that have been developed on the human circulatory system (Fig. 4, top right) provide an example of spatial layouts and conceptual relationships working together in a single sketch. Over a series of four worksheets, students are asked to draw and label the four chambers of the heart, use arrows to indicate how blood flows, and use arrows to indicate when oxygen is added to and removed from blood. An in-class experiment with these worksheets showed that students made significant gains on two out of three learning assessments on heart structure and function after using these worksheets (Miller et al., 2014). They improved their ability to correctly identify chambers of the heart and accurately identify how blood flows through the circulatory system. They did not make significant gains on an assessment on oxygen flow. Interestingly, the oxygen flow assessment was the only one of the three that did not have a visual component. We suspect that one factor in the success of these worksheets was the detection of common mistakes via

misconception sketches, which enabled the worksheet tutor to provide more targeted feedback.

In geoscience, we have found through laboratory experiments that sketching exercises using CogSketch can help assess student knowledge (Jee et al., 2014). In working with instructors and students in classroom studies (e.g., Garnier et al., this issue; Yin et al., 2010a) we observed that worksheets can be effective for helping students with scaling (i.e., understanding quantities across different scales and orders of magnitude). However, it is important that visual representations of quantity are visually alignable and that multiple levels of feedback are defined for location-specific glyphs. For example, in teaching students about the geologic time scale, one approach might be to have geologic events plotted on a single time scale, with the ability to zoom in and out to detect small differences in geologic time (e.g., 100,000 years) and large differences (e.g., 1 billion years). However, zooming in to one particular order of magnitude makes quantities from other orders of magnitude invisible. Instead of this approach, geoscience worksheets about scale have multiple scales aligned visually (i.e., a time scale with 100,000-year units aligned with a time scale with 100 million year units), so students can more easily see the relative length of geologic time periods without losing sight of very large or very small differences. To support the wide range of possible student answers, it is important that advice is flexible enough to let the student know what is wrong and how it might be fixed. For these worksheets, positional relations between student glyphs and quantitative ink constraints were used to provide hints about how the location of the student glyph should change (e.g., "Your mark for the Cenozoic era is too early."). Note that the worksheet tutor does not understand the interpretation of visual properties in terms of temporal relationships in the domain being taught—that mapping is the responsibility of the worksheet author, to keep the tutor as general as possible.

We have also observed that worksheets where quantities are expressed completely independent of ink are not as effective. For example, in a worksheet on the Earth's carbon cycle, students used arrows to indicate carbon transfers and used quantity annotations to denote the amount of carbon transferred. Each carbon mass annotation could be drawn in any way. Some students used dots, others wrote numeric symbols, and others used more arrows. However, the ink did not spatially correspond to the magnitude of the annotation, and students found this confusing or not useful. In contrast, in worksheets on free body diagrams, the lengths of force arrows are proportional to the magnitudes of the forces they represent (even when no numerical value is assigned to the force arrow). This seems to be a more effective approach since relative lengths (and thus relative magnitudes) are easy to compute visually.

The data collected by sketch worksheets also have potential for assessing student knowledge and strategies. Using a corpus of student sketches, we have used analogical reasoning to cluster sketches with similar structural characteristics to discover what common mistakes were being made (Chang & Forbus, 2014). In addition to providing an aggregate view of how students were completing a particular worksheet, the analysis provided insights into how feedback could be improved. In the biology worksheet activity

described earlier, Miller et al. (2014) performed a cluster analysis on students' interactive behaviors with the software, finding three distinct groups of students depending on how much and how often they requested feedback. Interestingly, the learning gains across these three groups were not consistent, indicating that learner strategy played a significant role in student learning. Students who scored relatively high on the pre-tests did not request feedback as often and did not make large learning gains compared to students who scored lower on the pre-test. The students who scored lower on the pre-tests sought feedback more often and caught up to the higher-achieving students by the end of the experiment.

5. Related Work

The analogy-based approach to automated feedback described above is closely related to some approaches in sketch recognition, particularly those that begin by decomposing ink into geometric primitives and describing their arrangement. LADDER (Hammond & Davis, 2003) organizes ink into points, lines, curves, and arcs, similar to the edge-level entities in CogSketch, whereas the region adjacency graphs (RAGs) of Lladós, Martí, and Villanueva (2001) resemble the edge-cycle level of organization. These approaches are similar in their representations but not in processing. LADDER recognized a visual symbol by testing its encoded representations against a set of strict, universally quantified geometric constraints. RAGs are compared via an inexact graph-matching algorithm based on edit distance, which is more similar to our analogical approach. McLure et al. (2011) perform sketched object recognition using some of the same representations and analogical processing models used here.

Most applications of sketch understanding in education have focused on automatically recognizing objects in a student's drawing (e.g., De Silva et al., 2007; Valentine et al., 2012). This can provide fluid interaction in domains where the situations are described via a small number of entity types and the mapping between shapes and entity types is one to one, such as electronic circuits or truss diagrams. Unfortunately, this approach requires that each system be crafted to handle a particular class of subproblems within a specific domain. It does not work when the mapping between shapes and entities is many to many (e.g., Fig. 1) and the shapes of entities are spatialized instead of conventionalized, for example, geoscience or biology. By contrast, the open-domain approach that CogSketch uses trades some naturalness of interaction for the ability to use the same system across multiple domains, including domains for which conventionalized visual symbols do not exist.

6. Discussion

This paper has focused on Sketch Worksheets: what they are, how they work, and what we have learned so far from laboratory and classroom experiments focused on education.[4] While the experiments to date are far from exhaustive in terms of domains even

within STEM education, these studies provide evidence that the particular visual representations CogSketch is computing are useful for understanding sketches in several STEM education domains. They also indicate that tutoring via analogical comparison over structured, relational representations of visual and spatial properties can provide useful feedback.

As is the case with the design of any instructional materials, worksheets are most effective when any knowledge about common answer patterns is known in advance. This is not necessary, but it is helpful in determining how to best detect and target specific mistakes. As we have seen in the many worksheets that use quantitative ink constraints, it is far better to be able to remark on the relative position of a student's drawing relative to the correct answer than to simply notify the student that he or she is incorrect. Similarly, as we saw with circulatory system worksheets, we were able to design worksheets that target specific types of errors by leveraging existing findings on common misconceptions about the circulatory system. Exploiting the hierarchical nature of perceptual organization also plays a role in generating useful feedback as it can dictate how to choose what suggestions should be shown to a student and what suggestions should be suppressed. We have found this approach to be useful in worksheets that use multiple levels of representation.

6.1. Future work

There is still a lot to be learned about the representations that have been useful so far and how they transfer to other domains. So far, our evaluations on learning have been in geoscience and biology. We plan to explore additional domains as well, for example, collecting data on free body diagram worksheets in classroom settings. Similarly, we are planning to expand our use of edge-level representations beyond the current set of orthographic projection exercises, for example, providing feedback on geoscience block diagrams. While we expect that new domains will present additional challenges, we are encouraged that the number of new representations needed has, in our experience, been shrinking over time, indicating a convergence for many kinds of 2D representations useful in STEM education.

We see two open issues that will be important to be addressed in the future. The first is better handling of spatially distributed entities, such as flow lines, force fields, and air masses. These are represented visually by multiple elements. Entering glyphs to depict such entities is easy with the current CogSketch interface: There is no requirement that glyphs consist of connected ink. However, decomposing them in visually meaningful ways that correspond to conceptually important distinctions will require more visual processing than we currently handle. Two examples are the changes in flow orientation as air moves around an airfoil (Kim, 1993) and the gradient of pressure implied by distances between air masses. Ideally, CogSketch would provide a summary of these visual properties to the worksheet author and be able to recognize them in student sketches. The second open issue is incorporating more 3D representation and reasoning capabilities in CogSketch, to better interpret sketches about surfaces and complex shapes, which often occur in engineering, biology, and geoscience. Since sketches are 2D, much can currently

be done by crafting advice based on 2D visual properties, but potentially more 3D representations could enable students to choose a wider variety of viewpoints when sketching. We have done preliminary experiments in this area, using analogical matching over line-labeling representations from computer vision (Lovett, Dehghani, & Forbus, 2008), but much remains to be done.

The domain generality of Sketch Worksheets is an important strength, but in one way it is also a weakness. The relationship between visual differences and what that means in terms of domain understanding remains implicit in the mind of the worksheet designer. Since there is no explicit model of expert performance in the domain, knowledge tracing of the kind performed by cognitive tutors (Koedinger, Anderson, Hadley, & Mark, 1997) is not possible. It might be possible to add an additional layer that makes this expert model explicit, particularly if such reasoning were done off-line and used to sequence worksheets given to students, just as student model updates in cognitive tutors are used to pose new problems to students. Another important dimension to explore is integrating CogSketch's representation and processing capabilities into more open-ended tutoring systems, like a tutoring system that teaches via analogy or uses a combination of sketching and conversational interaction to engage in Socratic tutoring.

Acknowledgments

Notes

1. We thank Kim Kastens for suggesting this worksheet.
2. A Voronoi diagram partitions a plane into regions based on distance to points in a plane. This can easily be converted into a graph of nodes (glyphs) and links (adjacency relationships between glyphs whose regions touch).
3. Currently it has 11,243 binary relations. Like concepts, the author chooses a small subset for students to use.
4. We have ignored the broader theoretical concerns motivating CogSketch, its usage in simulating human visual problem solving, and other educational applications—see Forbus et al. (2011) for an early summary of those aspects of the project.

References

Ainsworth, S., Prain, V., & Tytler, R. (2011). Drawing to learn in science. *Science*, *333*(6046), 1096–1097.
Biederman, L. (1987). Recognition-by-parts: A theory of human image understanding. *Psychological Review*, *94*(2), 115.

Chang, M. D., & Forbus, K. D. (2012). Using quantitative information to improve analogical matching between sketches. In *Proceedings of the 24th Annual Conference on Innovative Applications of Artificial Intelligence (IAAI)* (pp. 2269–2274). Toronto, Canada: Association for the Advancement of Artificial Intelligence.

Chang, M. D., & Forbus, K. D. (2014). Using analogy to cluster hand-drawn sketches for sketch-based educational software. *AI Magazine*, *35*(1), 76–84.

Chang, M. D., Wetzel, J. W., & Forbus, K. D. (2014). Spatial reasoning in comparative analyses of physics diagrams. In C. Freksa, B. Nebel, M. Hegarty & T. Barkowsky (Eds.), *Spatial Cognition* 2014, LNAI 8684 (pp. 268–282). Bremen, Germany: Springer.

Cohn, A. G., Bennett, B., Gooday, J., & Gotts, N. M. (1997). Qualitative spatial representation and reasoning with the region connection calculus. *Geoinformatica*, *1*, 275–316.

De Silva, R., Bischel, T. D., Lee, W., Peterson, E. J., Calfee, R. C., & Stahovich, T. (2007). Kirchhoff's pen: A pen-based circuit analysis tutor. In M. van de Panne & E. Saund (Eds.), *Proceedings of the 4th Eurographics Workshop on Sketch-based Interfaces and Modeling* (pp. 75–82). Riverside, CA: Association for Computing Machinery, Inc.

Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial intelligence*, *41*, 1–63.

Ferguson, R. W. (1994). MAGI: Analogy-based encoding using regularity and symmetry. In A. Ram & K. Eiselt (Eds.), *Proceedings of the 16th Annual Conference of the Cognitive Science Society* (pp. 283–288). Atlanta, GA: Lawrence Earlbaum Associates.

Forbus, K., Ferguson, R., Lovett, A., & Gentner, D. (2016). Extending SME to handle large-scale cognitive modeling. *Cognitive Science*, *40*, 1–50. doi:10.1111/cogs.12377.

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.

Garnier, B., Chang, M., Ormand, C., Matlen, B., Tikoff, B., & Shipley, T. (this issue). Testing the efficacy of CogSketch geoscience worksheets as a spatial learning and sketching tool in introductory Geoscience courses. *Topics in Cognitive Science*, this Topic.

Hammond, T., & Davis, R. (2003). LADDER: A language to describe drawing, display, and editing in sketch recognition. *Proceedings of the IJCAI 2003* (pp. 461–467).

Jee, B., Gentner, D., Uttal, D., Sageman, B., Forbus, K., Manduca, C., Ormand, C., Shipley, T., & Tikoff, B. (2014). Drawing on experience: How domain knowledge is reflected in sketches of scientific structures and processes. *Research in Science Education*, *44*(6), 859–883.

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

Koedinger, K. R., Anderson, J. R., Hadley, W. H., & Mark, M. A. (1997). Intelligent tutoring goes to school in the big city. *International Journal of Artificial Intelligence in Education*, *8*, 30–43.

Lladós, J., Martí, E., & Villanueva, J. (2001). Symbol recognition by error-tolerant subgraph matching between region adjacency graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *23*(10), 1137–1143.

Lockwood, K., Lovett, A., Forbus, K., Dehghani, M., & Usher, J. (2008). A theory of depiction for sketches of physical systems. *Proceedings of QR 2008*. Boulder, CO.

Lovett, A., Dehghani, M., & Forbus, K. (2008). Building and comparing qualitative descriptions of three-dimensional design sketches. *Proceedings of the 22nd International Qualitative Reasoning Workshop*. Boulder, CO.

Lovett, A., & Forbus, K. (2010). Shape is like space: Modeling shape representation as a set of qualitative spatial relations. *Proceedings of the AAAI Spring Symposium on Cognitive Shape Processing* (pp. 21–27). Palo Alto, CA: Association for the Advancement of Artificial Intelligence.

Lovett, A., & Forbus, K. (2011). Cultural commonalities and differences in spatial problem-solving: A computational analysis. *Cognition*, *121*(2), 281–287.

Lovett, A., & Forbus, K. (2017). Modeling visual problem-solving as analogical reasoning. *Psychological Review*, *124*(1), 60–90.

Lovett, A., Forbus, K., & Usher, J. (2010). A structure-mapping model of Raven's Progressive Matrices. In S. Ohlsson & R. Catrambone (Eds.), *Proceedings of CogSci-10* (pp. 2761–2766). Portland, OR: Cognitive Science Society.

Lovett, A., Gentner, D., Forbus, K., & Sagi, E. (2009a). Using analogical mapping to simulate time-course phenomena in perceptual similarity. *Cognitive Systems*, *10*(3), 216–228 [Special Issue on Analogies – Integrating Cognitive Abilities].

Lovett, A., Kandaswamy, S., McLure, M., & Forbus, K. (2012). Evaluating qualitative models of shape representation. *Proceedings of the 26th International Workshop on Qualitative Reasoning*. Los Angeles, CA.

Lovett, A., Tomai, E., Forbus, K., & Usher, J. (2009b). Solving geometric analogy problems through two-stage analogical mapping. *Cognitive Science*, *33*(7), 1192–1231.

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

Marr, D., & Nishihara, H. K. (1978). Representation and recognition of the spatial organization of three-dimensional shapes. *Proceedings of the Royal Society of London*, *200*, 269–294.

McLure, M. D., Friedman, S. E., Lovett, A., & Forbus, K. D. (2011). Edge-cycles: A qualitative sketch representation to support recognition. *Proceedings of the 25th International Workshop on Qualitative Reasoning*. Barcelona, Spain.

Miller, B., Cromley, J., & Newcombe, N. (2014). Using CogSketch to support student science learning through sketching with automatic feedback. *Proceedings of AERA*.

Palmer, S. E. (1977). Hierarchical structure in perceptual representation. *Cognitive Psychology*, *9*, 441–474.

Palmer, S. E., & Rock, I. (1994). Rethinking perceptual organization: The role of uniform connectedness. *Psychonomic Bulletin & Review*, *1*(1), 29–55.

Valentine, S., Vides, F., Lucchese, G., Turner, D., Kim, H., Li, W., Linsey, J., & Hammond, T. (2012). Mechanix: A sketch-based tutoring system for statics courses. *Proceedings of the Twenty-Fourth Innovative Applications of Artificial Intelligence Conference (IAAI)* (pp. 2253–2260). Toronto, Canada: Association for the Advancement of Artificial Intelligence.

VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*, *46*(4), 197–221.

Wetzel, J., & Forbus, K. (2010). Design buddy: Providing feedback for sketched multi-modal causal explanations. *Proceedings of the 24th International Workshop on Qualitative Reasoning*. Portland, Oregon.

Yin, P., Chang, M. D., & Forbus, K. D. (2010b). Sketch-based spatial reasoning in geologic interpretation. *Proceedings of the 24th International Workshop on Qualitative Reasoning*. Portland, Oregon.

Yin, P., Forbus, K., Usher, J., Sageman, B., & Jee, B. (2010a). Sketch worksheets: A sketch-based educational software system. In *Proceedings of the 22nd Annual Conference on Innovative Applications of Artificial Intelligence* (pp. 1871–1876). Atlanta, GA: Association for the Advancement of Artificial Intelligence.