

Sketch Worksheets in STEM Classrooms: Two Deployments

Kenneth D. Forbus,¹ Bridget Garnier,³ Basil Tikoff,²
Wayne Marko,³ Madeline Usher,¹ Matthew McLure¹

¹Computer Science, Northwestern University, 2133 Sheridan Road, Evanston, IL, 60208, USA

²Department of Geoscience, University of Wisconsin-Madison, 1215 W. Dayton Street, Madison, Wisconsin, 53706

³Earth & Planetary Sciences, Northwestern University, 2145 Sheridan Road, Evanston, IL, 60208, USA
forbus@northwestern.edu

Abstract

Sketching can be a valuable tool for science education, but it is currently underutilized. *Sketch worksheets* were developed to help change this, by using AI technology to give students immediate feedback and to give instructors assistance in grading. Sketch worksheets use visual representations automatically computed by CogSketch, which are combined with conceptual information from the OpenCyc ontology. Feedback is provided to students by comparing an instructor's sketch to a student's sketch, using the Structure-Mapping Engine. This paper describes our experiences in deploying sketch worksheets in two types of classes: Geoscience and AI. Sketch worksheets for introductory geoscience classes were developed by geoscientists at University of Wisconsin-Madison, authored using CogSketch and used in classes at both Wisconsin and Northwestern University. Sketch worksheets were also developed and deployed for a knowledge representation and reasoning course at Northwestern. Our experience indicates that sketch worksheets can provide helpful on-the-spot feedback to students, and significantly improve grading efficiency, to the point where sketching assignments can be more practical to use broadly in STEM education.

Introduction

Sketching is a natural way for people to think through spatial ideas, and to communicate about these ideas with others. This makes it attractive for STEM (Science, Technology, Engineering, and Mathematics) education, and there is indeed evidence that it can both help students learn and can be used to help assess student knowledge (e.g. Ainsworth et al. 2011; Jee et al. 2014). In geoscience, for example, sketching is heavily used by practitioners. Paradoxically, a survey of geoscience instructors indicates that, while most of them believe that sketching is useful for students as well, they use sketching assignments less than they desire (Garnier et al.

2017). One reason is that grading sketches is very time-consuming, compared to short-answer question forms. Providing support for semi-automatic grading of sketches can potentially increase their use in STEM education. The same mechanisms, if employed correctly, can also provide on-the-spot feedback to students, anytime and anywhere, similar to the functionality provided by cognitive tutors (Koedinger et al. 1997) for non-spatial subjects.

Sketch worksheets (Yin et al. 2010) are our approach to providing these capabilities for education. As described below, Sketch Worksheets use ideas from cognitive science and AI technology to visually analyze student sketches, and provide advice based on comparisons with instructor sketches via analogy. Importantly, domain experts and instructors author sketch worksheets, as opposed to AI experts or software developers needing to be involved. They have now been developed and deployed in two very different kinds of classes, including the use of sketch worksheets developed by one university in classes at another. This paper summarizes these experiences in deployment. We start by outlining the underlying technology, to set the stage. Next we describe their deployment in geoscience, including the process of developing them for an introductory geoscience class, and lessons learned from using them in classes. Then we describe how they have been used in a Knowledge Representation & Reasoning course, a topic far from the disciplines that motivated their original development. We close with conclusions and future work.

Sketch Worksheets: The basics

Sketch worksheets are built on CogSketch, an open-domain sketch understanding system (Forbus et al. 2011). This section provides just enough information about CogSketch to understand the rest of this paper¹.

Most AI efforts on sketch understanding focus on recognition, and this has led to successful deployed educational software systems for structural mechanics (Lee et al. 2007; Valentine et al. 2012), aspects of analog electronics (De Silva et al. 2007), and chemistry (Cooper et al. 2009). Our open-domain approach is fundamentally different. Recognition makes sense in domains like handwriting recognition or circuit diagrams, where the mapping from abstract entities to shapes is one-to-one, and the particular spatial properties of the layout of the visual symbols constituting a sketch are irrelevant. By contrast, for most STEM domains, the mapping between abstract entities and shapes is many-to-many, and the specific spatial properties of what is drawn are often crucial to understanding. For example, three concentric circles might depict the layers of the earth, planetary orbits, or the cross-section of a heat exchanger. The context of a specific exercise might enable discarding most of these interpretations, but even within an interpretation, an educational system should not presume that the student gets it right: The student might have the order of the layers mixed up, for instance. Hence it is important that students label what they draw, in order for the instructor (and the software) to accurately assess their knowledge. Thus CogSketch requires students to draw visual objects, called *glyphs*, and label them with their intended meaning. The labels are drawn from an underlying knowledge base², with textual renderings chosen by the instructor so that they can customize it.

The visual language of CogSketch provides three kinds of glyphs. *Entity glyphs* depict specific objects, concrete or abstract, e.g. an orbit. *Relation glyphs* depict binary relationships between the entities depicted by glyphs, e.g. *owns* in a knowledge graph. *Annotation glyphs* provide a way of specifying non-visual properties in the sketch, e.g. the temperature of an object. These glyphs are compositional, i.e. relation glyphs can apply to relations, to describe dependency structures, and annotation glyphs can apply to relation glyphs, e.g. to indicate rate of flow between carbon reservoirs.

In CogSketch, a sketch consists of *subsketches*, each providing related information about the subject of the sketch. In Sketch Worksheets, there is one or more *solution subsketch*, where the instructor depicts a correct solution to the problems posed by that worksheet³. Another subsketch

is used by students in doing their work, and unless the worksheet is unlocked, all solution subsketches are hidden from the student. Instructors author worksheets by selecting what concepts need to be used (by browsing the KB) and drawing their solution sketch. As always, CogSketch automatically performs a visual analysis of what they draw incrementally. If instructors desire more types of visual relationships to be computed, they can be specified via a menu. Instructors can choose particular facts to be marked as important facts for tutoring. Tutoring facts can have an associated piece of feedback to be provided when that fact isn't true in a student's sketch, and a point value that is used in grading. A second kind of tutoring fact are *quantitative ink constraints* where the instructor specifies an error tolerance for a glyph, and CogSketch will use a numerical comparison between the instructor's ink and the student's ink, providing feedback based on how they mismatch. These are used when the student is assigned to mark up a photograph or diagram.

Feedback to students is generated on demand, to avoid interrupting them while thinking. When feedback is requested, any optional visual computations needed to derive tutoring facts are performed if needed, and the instructor and student sketches are compared via the Structure-Mapping Engine (SME), a cognitive model of analogical matching (Forbus et al. 2016). Mismatches are detected by analyzing SME's candidate inferences, and any associated advice is retrieved and presented to the student, hyperlinked with the glyphs involved to help them make sense of it.

Instructors also provide text posing the problem(s) to the student, and optionally provide a background image as part of the problem (e.g. a photograph or diagram to be annotated). They can also provide multiple-choice questions to be asked of students before and after they complete a worksheet, as additional assessments and opportunities for reflection. CogSketch also includes a gradebook, which does batch processing of a directory of sketches, e.g. as downloaded from a course management system when assignments are turned in. It uses a web interface to enable instructors to browse each student's work, including its full history, and the automatically assigned points, based on the instructor-provided rubrics.


Sketch Worksheets in Geoscience

Geoscience is one of the most spatially intense STEM disciplines, hence it is a natural discipline for sketching. Here we discuss how sketch worksheets were developed for an introductory course at University of Wisconsin-Madison,

¹ For more details, please see (Forbus et al. 2017).

² The contents of the KB are derived from OpenCyc.


³ Multiple solution sketches, and sketches depicting common misconceptions, are also supported and have been useful in classroom experiments, but were not used in the deployed sketch worksheets discussed here.



and deployed in an introductory course at Northwestern University.

Developing Sketch Worksheets for Geoscience

Geoscience sketch worksheets were developed through a collaboration between geoscientists, cognitive scientists, and CogSketch developers. Collaboration was necessary to ensure that each worksheet: a) used results of cognitive-science based research to support the learning of spatially complex topics through sketching; b) properly utilized the spatial, feedback, and grading capabilities of CogSketch, which results in an activity that is beyond a typical paper sketching worksheet; and c) required primary authorship by a domain expert to ensure that the worksheet was rigorously correct and relevant to geoscience instruction.



Twenty-six sketch worksheets were developed for introductory geoscience courses (typically, physical geology). In developing sketch worksheet content, geoscientists identified specific concepts commonly taught in introductory geoscience courses that are difficult for students to grasp. In addition, cognitive scientists helped identify four spatial skills that are necessary for success in the geosciences: disembedding, reasoning about dynamic processes (Figure 1), penetrative thinking (Figure 2), and scaling (see Garnier et al., 2017). Geoscience concepts and spatial skills are important to develop at the introductory level to aid student understanding and possible success in future STEM courses.

Since most geoscience concepts involve the use of at least one spatial thinking skill, it was important to incorporate cognitive science-based research into sketching activities to support simultaneous learning of geoscience concepts and spatial skills development. Once a geoscience concept was chosen, sketch activities were created based on activities that would support development of the spatial skill, often taken from techniques or activities in educational research of geoscience, or other STEM disciplines.

Sketch activities were also tailored to take advantage of CogSketch's interactive and spatial capabilities. Each worksheet incorporated tasks that allow students to move and rotate objects, draw or annotate on top of photos and diagrams, free sketch, and draw arrows to show motion or relationships. Worksheets included a background image and/or objects to manipulate, directions to complete various spatial tasks, and multiple choice questions to answer when the worksheet is complete. All worksheets are currently accessible on the SERC website (<http://serc.carleton.edu>).

Creating the solution sketch and feedback was an important part of worksheet development. The goal of solution sketches and feedback was to lead students to the correct answer through a trial and error process. Preliminary testing

of each worksheet with student volunteers showed that poor solutions/feedback lead to student frustration if correct answers were identified as incorrect, if feedback was not helpful, and if students could not fix an incorrect sketch. These errors may initially be seen as errors with the SME or tutor, where in fact, it is a reflection of errors in the authoring process. For example, if a student correctly traces a fault but part of the line goes outside of the ink tolerance set by the worksheet author, the student will receive feedback for a correct sketch. To correct this situation, a worksheet author would increase the ink tolerance to include acceptable human error when drawing a line. Worksheet authors must thoroughly test and continually update solution sketches as a wider range of errors are observed. This process can be time-consuming but results in improved solution/feedback advice for each worksheet.

All worksheets were authored in CogSketch by a geoscientist, using the authoring environment. The first few worksheets each took 1-1.5 weeks to fully develop, from idea to completed worksheet. Development time greatly decreased with each worksheet, to the point where it took ~3 days per worksheet. About three-quarters of development time was spent creating the material (i.e., worksheet idea, images, text) and the remainder was spent authoring the worksheet in CogSketch, as well as testing the worksheet and correcting problems.

Initial testing of 16 developed sketch worksheets was conducted in an introductory geoscience course at the University of Wisconsin-Madison, spring 2014. By analyzing completed student worksheets, we learned that: a) the worksheet tutor allowed students to use different strategies to complete worksheets; b) the tutor identified and helped correct common mistakes; and c) almost no instructor time was needed to grade and provide feedback on sketch worksheets. Quantitatively, it took on average five hours per week to grade 196 paper worksheets (~1.5 minutes/worksheet; which includes hand-writing feedback messages similar to those that students automatically receive from the CogSketch tutor) versus only 7.5 minutes per week to grade 66 sketch worksheets (~0.11 minutes/worksheet), an order of magnitude difference. Therefore, we saw that sketching opportunities could increase in courses and save time for instructors, while still providing helpful, effective feedback for students.

Through the development process, we also learned that the CogSketch program and sketch worksheets are adaptable and able to change with continual use. Instructors also have the capabilities to make changes to sketch worksheets to better serve their courses. This well positions the program and sketch worksheets for continued growth and usage as new tasks and worksheets develop in courses.

Deploying Sketch Worksheets in Geoscience

Sketch worksheets were used alongside paper exercises in physical geology laboratory sections (Northwestern's Earth

201) during three 10-week quarters. Course enrollments ranged from 20 to 41 students. Students completed worksheets on topographic maps, geologic time, geologic structures, earthquakes, floods and flood recurrence, and glacial movement. This section briefly describes the pedagogical approach to using worksheets in the lab, summarizes qualitative observations from the classroom, and presents preliminary quantitative estimates of grading times.

Worksheets were selected from the assignments developed by Garnier et al. (2017; see the previous section), based on alignment with class learning objectives and estimated time needed to incorporate the exercises into the laboratory session. Nearly 500 worksheets were completed, submitted, and graded. One or two CogSketch Worksheets were typically completed by students in class on tablet PCs after they finished their paper exercises and were submitted as image files that were uploaded to the campus learning management system. Students used the graphical feedback meter that is part of the program to evaluate their progress and to determine if their sketches met the minimum requirements to receive credit for their work. Sketches were graded using a pass/fail scheme. For a passing grade, the following criteria needed to be met: all components of the worksheet were attempted and at least 70% of the sketch was correct. Initial grading of the sketches was done in the course learning management system.

Rubric	Paper-based		Tutor
	Pass/Fail (n=7)	Tallied (n=7)	Tallied (n=6)
Time (min.)/ sketch	.15(.025)	.40(.090)	.42(.094)

Note: n=number of batches; 2σ values shown in parentheses.

Table 1. Estimates of average grading times.

Using tablet PCs with sketch worksheets had impacts on both social dynamics and efficiency. In social dynamics, both benefits and drawbacks were observed. The drawbacks include decreased content-related questions, a shift toward technical support of tablets or worksheets, and a decrease in student-student interaction. However, an observed benefit was increased group focus when students switched to using tablets. Within the context of this deployment, group focus is defined as the apparent collective increase in student engagement with course content. Observations that support this interpretation included fewer and quitter student conversations and individual student attention directed toward tablets. Similar observations have been made in other lab environments where sketch worksheets and tablets were not used. However, the transition from paper to digital exercises appeared to be more abrupt

in this deployment and suggests that there may be a causal relationship between student behavior and the transition to a different content delivery medium. Additional research is needed to explore such phenomenon.

Sketch worksheets on tablet PCs influenced teaching efficiency in the areas of lab preparation, in-class use, and grading. In terms of lab preparation, tablet PCs needed to be maintained and access provided to assignments via the course management system. Occasional bugs when the software tutor did not recognize changes made by students were observed. These were easily fixed by saving and then restarting the worksheet, and selecting “update” or by quickly redrawing the sketch in a new worksheet. The small screen size (~10.5 inches) sometimes made it more challenging to read feedback messages and to find errors in sketches, especially in exercises that involved scaling (e.g., requiring the user to zoom in and out) and the positioning of many small glyphs. Benefits occurred in the areas of in-class use and grading. There were lower barriers to starting worksheets. Few if any students had questions about how to get started on a worksheet. It was also easier to read labels on sketches and there was a reduction in paper assignments to track and hand back in class.

None of the drawbacks presented above were significant enough to deter further use of CogSketch Worksheets in future lab sections and plans for new deployments are being made. Tasks such as tablet maintenance will need to be absorbed into the teaching process and the issue of diagnosing problems with sketches may be solved in several ways: using larger tablets or laptops, using a classroom management system to observe student work, or by designing worksheets to optimize the length of feedback and minimize the use of small glyphs.

The gradebook software was not originally used in the Northwestern classes in part because paper worksheets were also still being used during lab sessions, which had to be graded by hand anyway, and the pass-fail rubric on images was fairly efficient to grade. We used the Northwestern data to perform an additional evaluation of the efficacy of the software gradebook, using 20 batches of sketches composed of 10 to 19 exercises, sampling across the exercise types used in class. Batches of exercises were regraded using the tutor and its built-in quantitative rubric for which points were deducted and tallied, and in a paper format using the same rubric. Grading times varied both by exercise and grading method. Average results from these rubrics and methods are presented in Table 1. Single sketches were graded in less than a minute and batch times ranged from two to ten minutes. On average, the pass/fail method was the quickest per sketch. Paper-based and machine grading using tallied points produced similar grading times per sketch.

These times are substantially less than those found in the UW-Madison deployment, for two reasons. First, all students at Northwestern were using sketch worksheets, so they already benefited from the feedback they provide, thereby reducing the number of mistakes in what was turned in and simplifying the manual grading process. Second, the grading rubric at Northwestern, as noted above, was pass-fail, whereas for the paper worksheets at UW-Madison, not only did grades have to be assigned but the instructor also had to provide feedback, since that was their only source of help. This is further evidence that sketch worksheets can help with grading efficiency. Another potential advantage is that, unlike people grading paper worksheets, the software never suffers from fatigue.

Sketch Worksheets for KR&R

Arguably, Computer Science is one of the least spatial disciplines in STEM. Spatial models are often used in introductory courses, e.g. contour models for depicting variable scope, box-and-pointer notation for describing data structures, and process diagrams for describing the flow of computation. However, textual media dominates the everyday work of computer scientists, at least in terms of what they share with each other most often. Nevertheless, we believe that spatial models provide an important part of the tacit knowledge that computer scientists use. In knowledge representation, the rise of *knowledge graphs* provides a new opportunity for using spatial learning to support students. CogSketch’s visual language can be used to express knowledge graphs, by using entity glyphs to denote concepts and relation glyphs to express the links between them. Moreover, unlike existing concept map tools, arbitrary ink can be used to depict nodes, hence providing additional scaffolding and mnemonics. However, the same issues of grading efficiency arise in Knowledge Representation classes, even more so with the current flood of computer science enrollments. Consequently, we developed and deployed sketch worksheets in EECS 371: Knowledge Representation and Reasoning, taught by Forbus at Northwestern University. This section summarizes what worksheets were developed and our experiences in deploying them. We describe each in turn.

Using Sketch Worksheets for KR

Sketch worksheets were used in four out of the five homework assignments in the course. (The fifth assignment was an update on students' progress on their term projects, which absorbed their energies for the second half of the 10 week course.) The first assignment asked students to express relationships between concepts (Dog, Cat, Animal, Plant, Carnivore, Organism) via Venn diagrams, using containment and disjointness to represent what (in Cyc) would be *genls* and *disjointWith* relationships, thereby connecting these new ideas with prior learned models (see Figure 3). All but two students received a perfect score, which is not surprising given that this was essentially a warm-up exercise. Even so, feedback was needed, since students often forgot about the existence of carnivorous plants. Students used the feedback 4 times on average during this assignment.




Figure 3: Student misconception and corrected version

The second assignment gave them practice in representing everyday situations. It consisted of two sketch worksheets. The first provided a drawing in the background, of a person standing on a floor, with a ceiling overhead that had a light connected to it. Students were expected to add entities representing what was depicted, choosing the most appropriate concept to represent each one of them, and to draw a specific subset of the relationships that held between them. The set of concepts and relationships was chosen from the KB to have both the correct concept and a reasonable set of tempting distractors. Similarly, a second worksheet asked them to draw landmarks on Northwestern's North Campus "the way you might draw them to explain their layout to a visiting friend" and then add relation arrows to indicate the spatial relationships between adjacent objects (five in all). The only wrinkle with this worksheet was that new blank relationship had to be added and renamed to look like the

actual Cyc relationships that should be used. Otherwise, CogSketch automatically inferred the correct geospatial relationships and declared the sketch to be finished before the student drew anything. 57 out of 58 students received perfect scores, accessing feedback 6 times and 9 times on average, for the two worksheets respectively.

The third assignment required them to fill out a worksheet on a mythical soap opera (The Eternal Turmoil), whose contents had been informally specified via a student-driven discussion in class. Soap operas, as noted by (Brachman & Levesque 2004), provide marvelous scope for practice with representing events, relationships, and causality. This story included an event ("Leo is murdered in an abandoned gym





Figure 4: Encoding a murder mystery. Note the glyphs inside a region, indicating a microtheory corresponding to one of the participant's desires, rather than the world itself.

with a candlestick"), which is depicted by an entity for the event itself, entities for the roles in the event, and role relations connecting them. A subtlety, supported by Cyc's use of microtheories - which provide Cyc's notion of local context - is handling desires. Here, as shown in Figure 4, "Kathy wants to murder Leo because Leo killed her twin sister." This desired murder must be distinct from the actual murder that occurred, which is done in the formalism via a separate microtheory, linked to the person who wants it via the *desires-Microtheory* relationship. Microtheories can be depicted via entity glyphs, with glyphs inside them depicting the facts specific to that microtheory. 78% of the students received perfect scores, but even so, 22% turned in worksheets with one or more problems, typically concerning a missing event or relationship. The additional complexity of this assignment can be seen from the use of the feedback system 40 times on average by each student during its completion.

The final sketch worksheet had students encode a different murder mystery from a homework assignment in the textbook⁴, to give them practice in producing DNF clauses from natural language specifications (Figure 5). Each clause was depicted by an entity glyph which was connected to its terms by relationships arrows (*posDisj* indicates that a proposition is a positive disjunct in the clause, and *negDisj* indicates that a proposition is a negative disjunct). This worksheet went beyond the limit of the built-in analogical matching support in sketch worksheets, with the ambiguity in the multiple *posDisj* relations often leading to errors in mapping (and hence feedback and grading) in pilot testing⁵, so we turned feedback off (warning students of this, since by now they expected it) and grading the worksheets by eyeballing them using the gradebook. Students did indeed do worse, with only 38% of them achieving perfect scores without on-the-spot feedback.



Conclusions and Future Work

Our experience with deploying sketch worksheets indicates that the technology has reached the point of achieving most of its goals. Specifically, they can be used by students in more than one discipline (geoscience, AI), and they can be authored by domain experts and instructors (as indicated by the geoscience experience). Grading efficiency is enhanced, as is the ability for an instructor to gain a deeper understanding by browsing through the history of a student's work on a sketch, something which is simply unavailable with pencil and paper sketches (barring video analysis and drawing with

multiple color pens, two laboratory practices that are completely impractical for classroom-scale use and impossible for homework assignments). The one remaining goal to be demonstrated is showing that using sketch worksheets actually improves student learning, compared to both non-sketching exercises and sketching on pencil and paper. There is already evidence that sketching can provide gains over verbal self-explanation in understanding texts (Scheiter et al. 2017), and the sketching experience for students is sufficiently fluent that we would expect it to hold for sketch worksheets as well. But such experiments remain to be done, ideally with randomized controlled trials across balanced classrooms. Removing the bottleneck of grading burden should facilitate those experiments being done in the future. However, we note that the geoscience worksheets focus on implementing research-based techniques that have previously been shown to improve learning. In addition, the fact that CogSketch and worksheets can be used in an actual course and greatly reduce grading time for instructors is a major accomplishment that not all educational tools can claim. Finally, many of the lessons concerning feedback in cognitive tutors may be applicable to sketch worksheets, but again, this is a subject for future experimentation. CogSketch's visual analysis capabilities provides the prospect of using sketches as a medium for educational data mining, and using analogical generalization to help instructors identify common patterns of misconceptions (Chang and Forbus, 2014). We hope that these deployments are just the next step of helping spread sketching more broadly through STEM education.

Acknowledgements

This work was supported by the NSF-funded Spatial Intelligence and Learning Center, Award # SBE-1041707.

References

- Ainsworth, S., Prain, V., & Tytler, R. (2011). Drawing to learn in science. *Science*, 333(6046), 1096-1097.
- Brachman, R. & Levesque, H. (2004). *Knowledge Representation and Reasoning*. Morgan-Kaufmann.
- Chang, M. and Forbus, K. (2012) Using Quantitative Information to Improve Analogical Matching Between Sketches. *Proceedings of IAAI 2012*, Toronto, Canada.
- Chang, M. D. and Forbus, K. D. (2014). Using analogy to cluster hand-drawn sketches for sketch-based educational software. *AI Magazine*, 35(1), 76-84.
- Cooper, M., Grovea, N., Pargas, R., Bryfczynski, S., and Gatlin, T. (2009). OrganicPad: An interactive freehand drawing application

⁴ Problem 3, Chapter 4, in (Brachman & Levesque, 2004)

⁵ Since the material was abstract, the quantitative grounding techniques introduced previously (Chang & Forbus, 2012) were inapplicable.

for drawing Lewis structures and the development of skills in organic chemistry. *Chemistry Education Research and Practice*, 10(4), 296-301.

De Silva, R., Bischel, T. D., Lee, W., Peterson, E. J., Calfee, R. C. and Stahovich, T., 2007. Kirchhoff's pen: A pen-based circuit analysis tutor. In *Proceedings of the 4th Eurographics workshop on Sketch-based interfaces and modeling*.

Forbus, K., Usher, J., Lovett, A., Lockwood, K., and Wetzel, J. (2011). CogSketch: Sketch understanding for Cognitive Science Research and for Education. *Topics in Cognitive Science*, 3(4), pp 648-666.

Forbus, K., Ferguson, R., Lovett, A., & Gentner, D. (2016) Extending SME to Handle Large-Scale Cognitive Modeling. *Cognitive Science*. DOI: 10.1111/cogs.12377, pp 1-50.

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

Garnier, B., Chang, M., Ormand, C., Matlen, B., Tikoff, B., & Shipley, T. (2017) Testing the efficacy of CogSketch geoscience worksheets as a spatial learning and sketching tool in introductory Geoscience courses. *Topics in Cognitive Science*

Jee, B., Gentner, D., Uttal, D., Sageman, B., Forbus, K., Manduca, C., Ormand, C., Shipley, T., and 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.

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.

Lee, W., de Silva, R., Peterson, E., Calfee, R., and Stahovich, T. (2007) Newton's Pen – A Pen-based Tutoring System for Statics. *EUROGRAPHICS Workshop on Sketch-based interfaces and modeling*.

Scheiter, K., Schleinschok, K, & Ainsworth, S. (2017) Why Sketching May Aid Learning from Science Texts: Contrasting Sketching with Written Explanations. *Topics in Cognitive Science*, DOI: 10.1111/tops.12261.

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)*.

Yin, P., Forbus, K., Usher, J., Sageman, B. and Jee, B. (2010). Sketch Worksheets: A Sketch-based Educational Software System. *Proceedings of the 22nd Annual Conference on Innovative Applications of Artificial Intelligence*.