Using Quantitative Information to Improve Analogical Matching Between Sketches

Maria D. Chang, Kenneth D. Forbus
Qualitative Reasoning Group, Northwestern University
2133 Sheridan Road, Evanston, IL 60208
maria.chang@u.northwestern.edu, forbus@northwestern.edu

Abstract

Qualitative representations are suitable for sketch understanding systems because they highlight important relationships while leaving out details that are not essential for conceptual understanding. These representations can be used to perform spatial analogies between sketches, which determine qualitative similarities and differences. However, there are cases where including quantitative information is necessary for accurately representing a sketch. We describe a method for using quantitative information to constrain qualitative spatial analogies. The utility of this method is demonstrated in the context of a sketch-based educational software system. Importantly, using quantitative information to improve analogical matches is not domain-specific. It can be used in any situation where qualitative and quantitative spatial information must be combined to accurately interpret a sketch. This approach has the potential to improve sketch understanding in educational software applications for highly spatial domains.

Introduction

Sketching is an excellent tool for communicating spatial ideas. When we externalize spatial concepts into a sketch or diagram, spatial inferences become easier and working memory demands decrease (Larkin and Simon 1987). Sketching is pervasive in design settings and in classrooms. For highly spatial domains, like science, technology, engineering and mathematics (STEM fields), sketching is useful for teaching spatial concepts and for assessing a student's knowledge (Ainsworth, Prain and Tytler 2011, Jee et al. 2009, Kindfield 1992).

One of the benefits of sketching is its flexibility. Sketches can be rough, inexact, and not drawn to scale. For this reason, qualitative representations are well-suited for sketch understanding because they break continuous quantities into discrete units that can be reasoned about more easily, eliminating irrelevant quantitative details. Consider, for example, a sketch of the solar system. Qualitative relations that describe containment are sufficient to determine if the order of the planets is correct. Mercury and its orbit must contain the Sun, Venus and its orbit must contain Mercury and so on. Quantitative information (e.g. the location of the ink in a 2D coordinate plane or the raw distance between the planets) is not necessary for understanding this particular sketch.

Qualitative representations can also be used to create spatial analogies, which are used to detect similarities and differences between spatial representations. Spatial analogies that are based on qualitative representations are stable because they highlight important relationships while leaving out details that are not needed for a meaningful, human-like comparison. This provides a powerful tool for applications. For example, in *sketch worksheets* (Yin et al. 2010), a student's sketch is compared with an instructor's sketch, and the differences between them, which are found via analogical comparison, are used to generate feedback.

However, there are cases where purely qualitative representations are not enough. When annotating a photograph, for example, the annotation of a feature must actually be at the location of that feature in the photograph, and be of roughly the correct size and shape. As Yin et al. (2010) outlines, this can be done by prescribing an optional tolerance associated with entities in the instructor's sketch. We refer to these quantitative criteria as *quantitative ink constraints*. In Figure 1a (left), for example, the instructor has specified a tolerance around the glyph for the right ventricle. When the student sketch in Figure 1a (right) is compared with this sketch, the student's drawing of the right ventricle is within the tolerance region, and hence the quantitative ink constraint is satisfied. As Figure 1b

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). Toronto, Canada.

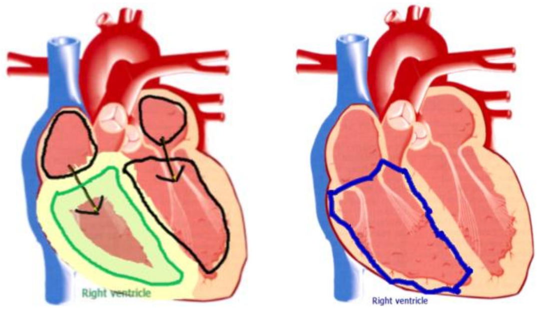

Figure 1A: An example of a quantitative ink constraint in a heart anatomy exercise. The sketch on the left has the chambers labeled along with blood flow arrows. The buffer region around the right ventricle is the tolerance region for the quantitative ink constraint. The sketch of the right shows an acceptable drawing of the right ventricle.

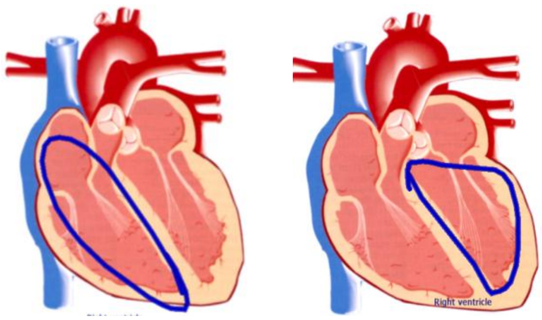

Figure 1B: Two drawings for the right ventricle that violate the quantitative ink constraint.

illustrates, a student may draw the right ventricle in the wrong location or without sufficient overlap with the instructor's drawing, In this case, the quantitative ink constraint is violated and the system would provide feedback to the student.

This method relies on the accuracy of the analogical match, which in turn relies on the qualitative structure found for a student sketch being close to that of the teacher's sketch. Because students are still learning the domain, this is often not the case, and mismatches can lead to students getting inaccurate advice from the tutor.

This paper describes a general technique for using quantitative information to repair mismatches in analogies between sketches. We begin with an overview of CogSketch, our open-domain sketch understanding system. We then describe how we use quantitative information to improve matches. Using a corpus of student sketches from a classroom experiment, we show that it yields significant improvements in matching accuracy. We conclude with related and future work.

CogSketch

CogSketch is an open-domain sketch understanding system that incorporates models of human spatial and analogical reasoning to understand sketches in human-like ways (Forbus et al. 2011). The basic building blocks of a sketch are called *glyphs*. Glyphs can be used to represent *entities*, *relationships* and *annotations*. The chambers of the heart in Figures 1a and 1b are examples of entities. Relationships conceptually connect entities. For example, relationship arrows are used in Figure 1a (left) to indicate flows between chambers of the heart. Attributes can be represented with annotations, e.g. the rate of a flow.

What a glyph represents is specified explicitly by the user, using interfaces that provide student-friendly access to a subset of concepts in CogSketch's knowledge base[1]. No recognition is required. This is what we mean by open-domain sketch understanding, and it is an important design decision: today's statistical recognizers require training, typically on a per-user basis, and work best when there is a small vocabulary of pre-determined symbols. The conceptual labeling interface that CogSketch provides enables users to pick concepts easily. As Figure 1 illustrates, many STEM domains do not use a fixed library of visual symbols: the geometry of the parts in a sketch often matters.

Qualitative Representations

CogSketch uses visual processing techniques to construct qualitative spatial representations of the ink. Topological relations are automatically computed (Cohn et al. 1997). Positional relations (e.g. above, rightOf) are automatically computed between adjacent glyphs and can be computed on demand between non-adjacent glyphs. These relationships capture the essence of the spatial properties of the sketch, without relying on quantitative measures that humans would typically ignore (Huttenlocher et al 1991).

To compare sketches, CogSketch uses the Structure-Mapping Engine (SME) (Falkenhainer et al 1989), which is based on Gentner's Structure-Mapping theory of analogy (Gentner 1983). Structure-mapping takes as input two descriptions (a base and a target), which are structured, relational representations. It produces one or more *mappings*, consisting of *correspondences* that describe what entities and statements are aligned within that mapping, a (possibly empty) set of *candidate inferences* that describe differences between the inputs, and a *structural evaluation score* (SES) that provides a numerical estimate of match quality.

Another optional input to SME consists of *match constraints*. Two kinds of match constraints are used here:

[1] CogSketch uses contents derived from Cycorp's OpenCyc knowledge base, plus extensions for qualitative and analogical reasoning.

partition constraints on concepts indicate that only instances of that concept can match with each other, e.g. ventricles can only match with ventricles. *Required correspondences* indicate that any mapping must include a correspondence involving the given pair of items (both entities or both statements). Here required correspondences are generated using our new algorithm, described below. SME uses a "middle-out" matching process, i.e. an early stage finds a large set of candidate correspondences, proposing them based on local matches between statements and their arguments. This forest of possible matches is winnowed down via psychologically motivated constraints[2], and combined via a greedy merge algorithm into one or more structurally consistent global mappings (Forbus & Oblinger, 1990). The approximate nature of this process is the main source of mismatches, but such approximations are essential for tractability.

Sketch Worksheets

The qualitative and quantitative representations described above are used heavily in *sketch worksheets*, a sketch-based educational software system built within CogSketch (Yin et al. 2010). Sketch worksheets are inspired by traditional paper and pencil worksheets, which are common tools for teaching and learning in domains that require spatial skills (i.e. the STEM fields). Unlike traditional worksheets, sketch worksheets use spatial and conceptual reasoning to provide on-demand feedback so the student can iteratively revise his or her sketch until either the system has no more suggestions or the student is satisfied with their sketch. Sketch worksheets are not tied to any particular domain. The main knowledge representation requirement is that the problem solution can be represented with a sketch.

Each sketch worksheet contains a solution sketch, which is hidden from the student. Authoring a worksheet includes drawing that solution sketch and providing conceptual labels for all of the glyphs in it. The relationships automatically computed by CogSketch can be flagged as important, i.e. they must hold for any student sketch to be correct. Quantitative ink constraints are also specified when relevant for glyphs in the solution. Advice to be given if a constraint is violated is provided via text strings associated with that constraint.

An example worksheet solution from an undergraduate structural geology class is shown in Figure 2A. The task for this worksheet is to identify the main fault line, the hanging wall and foot wall, the direction of movement along the fault line and the four prominent marker beds (indicated by numbers 1-4). Quantitative ink constraints are defined for the marker beds and the main fault line because their location relative to the background image is important.

When the worksheet is distributed to students, they sketch their candidate solution. At any time, they can request feedback to get advice from the system. Sketch worksheets have been used in several in-class experiments at Northwestern, plus an experiment at Carleton College. These experiments are providing the data that we need to refine the system in order to better support student learning.

On-Demand Feedback

Feedback is generated by comparing the student's sketch to the pre-defined solution sketch, using an analogical mapping computed by SME. The base and target consist of the qualitative spatial representations that CogSketch computes, along with the attributes specified for each glyph via conceptual labeling (e.g. that a glyph represents a fault or a marker bed) and any conceptual relations provided by relationship and annotation glyphs.

One of the challenges of using analogy in this task is that the sketches being compared are often very different. This arises both due to lack of knowledge on the students' part, but also because they can ask for feedback at any time, even during the early stages of sketching their solution. To improve mapping accuracy, the tutor includes partition constraints for each concept in the sketch. This exploits the fact that the matches are all within-domain, i.e. it makes no sense to have a fault correspond with a marker bed. If there is only one instance of a concept per sketch, partition constraints suffice to eliminate mismatches. However, as Figure 2 illustrates, this is often not the case when there are multiple instances of a concept.

The analogy is used to find differences between the teacher's sketch and the student's sketch. If the student's sketch is missing an important attribute or relationship, it will show up as a candidate inference. All candidate inferences are scanned to see if there is advice associated with the base statement that generated them, and if so, the advice is added to the pool of feedback provided to the student. Quantitative constraints are handled by finding the corresponding entity in the student's sketch, and seeing if that entity's ink lies entirely within the tolerance region for the corresponding solution glyph. In other words, the qualitative structure tells us what quantitative tests to do.

Sometimes the qualitative and conceptual relationships and attributes are not enough to create an accurate mapping. Consider the sketch in Figure 2B, which only has the main fault line and 4 marker beds. The qualitative spatial representations capture the relative location of the marker beds, leading SME to map the upper left marker bed to the upper left marker bed in the solution, the upper

[2] The psychological constraints used by SME are based on evidence that people prefer analogies with structurally consistent systems of relations (where matching relations have matching arguments) and with greater systems of nested relations, i.e. deep relational structure.

right marker bed to the upper right marker bed in the solution, and so on. Since the locations of these glyphs matter, they each have a quantitative ink constraint, and the tutor advises the student that all four are incorrect. This is bad advice on the tutor's part: human instructors recognize that two of the glyphs are in the correct positions, while the other two are not. This is an example of the kind of mismatch our technique addresses.

Quantitative Constraints on Analogy

To improve the accuracy of analogical mappings between sketches, we developed a strategy for using quantitative constraints to improve mappings. It works as follows:

1. Run SME to compare the teacher's (base) and student's (target) sketches.
2. For each base glyph G_b that has a quantitative ink constraint Q,
    a. If the corresponding target glyph G_t satisfies Q, do nothing.
    b. Otherwise, for each competing correspondence, test its target glyph G_a to see if it satisfies Q. If so, extend the set of match constraints to require that G_b correspond to G_a.
    c. If no quantitative match is found, then Q is violated.
3. Repeat until the set of match constraints stops growing.

Step 2b is efficient because SME automatically computes all of the potential competing matches in its initial phase of operation. This step also assumes that no two glyphs in the sketch have exactly the same ink, which is reasonable given the nature of sketches. CogSketch uses a truth-maintenance system, so that when glyphs are moved or edited, their spatial properties are automatically recomputed. The required correspondence in Step 2b is justified via assertions about spatial properties of the ink, hence they will automatically be retracted if the student improves their sketch.

Given the pair of sketches in Figure 2, this algorithm first runs SME to compare the base (teacher's sketch, Figure 2A) to the target (student's solution, Figure 2B). Recall that SME uses qualitative spatial relations to put entities into correspondence. For instance, the following two statements are true in the base and target, respectively:

```
(above B1 B3)
(above T1 T3)
```

These facts (and others) can be used as support for putting B1 into correspondence with T1 and B3 into correspondence with T3. Using qualitative relationships like these, SME arrives at the following set of entity correspondences for marker beds (other entities omitted for brevity):

Base Item	Target Item
B1	T1
B2	T2
B3	T3
B4	T4

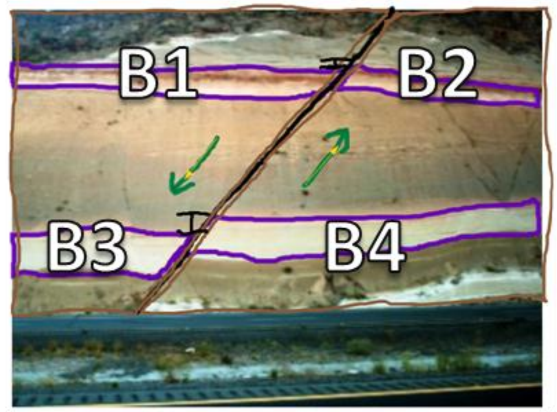

Figure 2A: An example teacher's solution from a worksheet for an undergraduate structural geology class. The marker bed outlines are indicated by numbers 1-4.

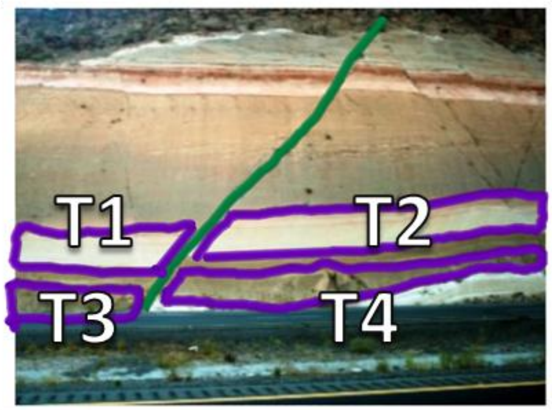

Figure 2B: An example candidate solution with only marker beds and a fault line. Without using quantitative constraints to improve the analogy, the system would determine that all four marker beds were incorrect.

Next, the algorithm checks the quantitative ink constraints on each base glyph with respect to its corresponding target glyph (step 2). Using the third correspondence as an example, we see that T3 does not satisfy the quantitative ink constraint for B3 because it is lower than it should be. The algorithm checks competing correspondences T1, T2 and T4 for potential matches. It finds that T1 satisfies the quantitative ink constraint for B3 and thus asserts a required correspondence between B3 and T1. The required correspondences that result from this

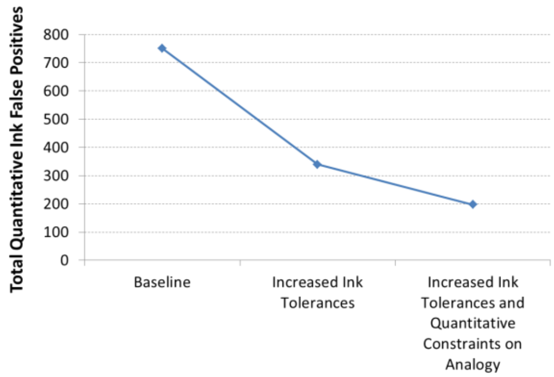

Figure 3: Quantitative ink false positives. Errors were greatly reduced by using increased quantitative ink comparison tolerances and by using quantitative information to constrain analogical mappings.

step are added to the set of match constraints for this analogy.

Steps 1 and 2 are repeated until no new match constraints are found. After using these quantitative ink constraints to restrict the analogical mapping, the system arrives at the following set of marker bed correspondences:

Base Item	Target Item
B1	T3
B2	T4
B3	T1
B4	T2

In the new mapping, only two target items violate the quantitative ink constraints: T3 and T4. Indeed, these are the only two incorrect marker beds in the student's sketch. As a result, the tutoring system provides feedback to the student, indicating that these two marker beds are drawn incorrectly.

Evaluation

To evaluate the utility of our technique, we tested it on a corpus of sketch worksheets on fault identification (e.g. Figure 2A,2B). All sketches were drawn by undergraduate students at Northwestern University as part of a structural geology homework assignment.

A total of 120 sketches were submitted by students. Over the 120 sketches, students requested feedback a total of 834 times. Each sketch comes with a history, which saves what action students did when. The history data was used to reconstruct each sketch as it was when a feedback request was initiated by the student. This provided us with 834 sketches that represent the scenarios where students requested feedback. Each of these sketches was visually inspected to determine the suggestions that should have been given by the tutor. These gold standard suggestions were then compared to the feedback that the student actually received. There were several types of mismatches, the most common being a large number of false positives for quantitative ink constraint violations (Figure 3). The original number of false positives was 751 out of 3,360 possible, or a 22% error rate. Further visual inspection revealed that a slight increase in quantitative tolerance could yield a substantial improvement, dropping the number of false positives to 340 (10%). Doing so increased the number of false negatives slightly, i.e., by 58, but the total number of errors dropped from 756 to 398.

By using the algorithm above, the number of false positives drops from 340 to 197 (5.8%). This is statistically significant: the average number of quantitative ink false positives per feedback request decreased from 0.41 to 0.24 ($t(833) = 9.76$, $p < 0.001$).

Related Work

Several sketch understanding systems have been developed but they rely on ink recognition to understand the contents of the sketch (Lee et al. 2007, de Silva et al. 2007). Ink recognition can make sense when the domain is tightly limited to a small number of visual symbols, and users are either experts who are willing to invest in training the system (and themselves) on that vocabulary, or they are trying to learn how to draw those symbols, e.g. Kanji or Mandarin phonetic symbols (Taele and Hammond 2009, Taele and Hammond 2010).

Most computational models of analogical processing today focus on connectionist modeling (e.g. Hummel & Holyoak 2003; Larkey & Love 2003), and have capacity limits which make them incapable of matching sketches of the complexity needed for STEM education problems. Our algorithm can be viewed as a variation of Falkenhainer's (1987) *map-analyze cycle*, where a partial mapping is analyzed to provide constraints to improve the mapping. Falkenhainer's work concerned modeling the learning of qualitative domain theories via analogy, and did not handle spatial representations nor quantitative constraints, nor was it ever applied to an application such as education.

Discussion and Future Work

These results demonstrate that using quantitative information to constrain qualitative analogical mappings can improve the interpretation of sketches. The evaluation we used is specific to Sketch Worksheets but the approach is not. This approach may be used in any situation where a combination of qualitative and quantitative information is necessary for understanding a sketch.

Using both qualitative and quantitative information increases the flexibility of sketch understanding and allows the system to harness the benefits of both types of representations. Qualitative representations are needed to describe sketches at a level of detail that makes analogical mappings stable and robust. However, for cases where qualitative representations are not enough, quantitative representations provide just enough extra information needed to get the mapping right. This essentially allows the system to fine tune the analogical mapping to come up with the optimal interpretation. This approach is inspired by the way people incrementally interpret a sketch. In educational settings, instructors will often give students the "benefit of the doubt" by reinterpreting the sketch based on multiple sources of information.

The current approach only uses one type of quantitative information to constrain the analogical mapping. Other types of quantitative information include those specified by annotations entered by the user, and lengths of segments when a shape is decomposed into edges. Understanding how this information could be used to improve analogical mappings could be helpful as well.

Our future goals include making sketch worksheets widely available across STEM domains, by enabling domain experts and educators to create their own worksheets. This requires extremely robust matching, which also needs to be human-like in order to support applying instructor-provided grading rubrics. Consequently, we plan to evaluate this method in sketch worksheets for other domains. This may help reveal other potential strategies for finding optimal spatial analogies, which might be used for improving sketch understanding more generally.

Acknowledgements

Thanks to Brad Sageman and Andrew Jacobson and their students, for providing us with expertise, advice, and data. This work was supported by the Spatial Intelligence and Learning Center (SILC), an NSF Science of Learning Center (Award Number SBE-1041707).

References

Ainsworth, S., Prain, V. and Tytler, R., 2011. Science education. Drawing to learn in science. *Science*, 333(6046): 1096-1097.

Cohn, A. G., Bennett, B., Gooday, J. and Gotts, N. M., 1997. Qualitative spatial representation and reasoning with the region connection calculus. *GeoInformatica*, 1(3): 275-316.

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.

Falkenhainer, B., 1987. An examination of the third stage in the analogy process: Verification-based analogical learning. In Proceedings of the International Joint Conference of Artificial Intelligence (IJCAI). Los Altos, CA.

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

Forbus, K. and Oblinger, D. 1990. Making SME greedy and pragmatic, *Proceedings of the Cognitive Science Society*, July, 1990.

Forbus, K. D., 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: 648-666.

Gentner, D., 1983. Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7(2): 155-170.

Hummel, J. E., & Holyoak, K. J. (2003). A symbolic-connectionist theory of relational inference and generalization. *Psychological Review*, 110, 220-264.

Huttenlocher J., Hedges L. V., Duncan S. (1991) Categories and particulars: Prototype effects in estimating location. *Psychological Review*, 98(3), 352-376.

Jee, B., Gentner, D., Forbus, K. D., Sageman, B. and Uttal, D. H., 2009. Drawing on experience: Use of sketching to evaluate knowledge of spatial scientific concepts. *Proceedings of the 31st Annual Conference of the Cognitive Science Society*.

Kindfield, A. C. H., 1992. Expert diagrammatic reasoning in biology. In Proceedings of the AAAI Spring Symposium. Palo Alto, CA.

Larkey, L. & Love, B. (2003). CAB: Connectionist Analogy Builder. *Cognitive Science* 27,781-794.

Larkin, J. H. and Simon, H. A., 1987. Why a diagram is (sometimes) worth 10000 words. *Cognitive Science*, 11(1): 65-99.

Lee, W., De Silva, R., Peterson, E. J., Calfee, R. C. and Stahovich, T. F., 2007. Newton's pen – a pen-based tutoring system for statics. In Proceedings of the Eurographics Workshop on Sketch-Based Interfaces and Modeling. Riverside, CA.

Taele, P. and Hammond, T., 2009. Hashigo: A next-generation sketch interaction system for japanese kanji. In Proceedings of the 21st Annual Conference on Innovation Applications of Artificial Intelligence (IAAI). Pasadena, CA.

Taele, P. and Hammond, T., 2010. Lamps: A sketch recognition-based teaching tool for mandarin phonetic symbols i. *Journal of Visual Languages and Computing*, 21(2): 109-120.

Yin, P., Forbus, K. D., Usher, J., Sageman, B. and Jee, B., 2010. Sketch worksheets: A sketch-based educational software system. In Proceedings of the 22nd Annual Conference on Innovative Applications of Artificial Intelligence (IAAI). Portland, OR.