

Humans Encode Qualitative Topological Relations

Andrew Lovett, Steven Franconeri

Visual Cognition Lab, Northwestern University
andrew.lovett@northwestern.edu

Abstract

The human visual system relies heavily on categorical representations, similar to the qualitative representations in computing. Here we examine categorical topological relations between objects. When asked to detect changes between object arrangements, participants were better at detecting those changes that crossed hypothesized category boundaries, such as 'overlapping', or 'touching', compared to equally-sized changes that did not. These effects were magnified at increased memory load, presumably because categorical relations forms a more efficient code. This finding, predicted by previous computational modeling work, suggests that categorical relations are critical for remembering and comparing complex images.

Introduction

Visual comparison plays a central role in our mental lives. We learn by comparing what we see to what we've seen before, whether it is an abstract scientific diagram or a complex physical environment. As researchers, it is imperative that we better understand the representations and processes that guide visual comparison – how and when do we identify commonalities and differences?

One key factor is the type of information being compared. Visual information can be divided into two types: *metric* and *categorical*. Metric information describes continuous values, such as an object's location or orientation in space. Categorical, or qualitative, information divides the continuous world into discrete categories, indicating that an object is **right of** another, or that an object's orientation is **vertical**. During comparison, it is easier to distinguish two stimuli if there is a categorical difference between them. For example, it is easier to distinguish **blue** from **green** than to distinguish two shades of green (Bornstein & Korda, 1984). This effect has been labeled *categorical perception*.

Categorical perception of object properties like color has been heavily studied. However, categorical perception of relations *between* objects has seen less attention. We address this omission, directly demonstrating that categorical relations support difference detection. The relations tested are drawn from previous computational modeling work, which

has suggested a range of powerful qualitative relations for discretizing the continuous world (Randall, Cui, & Cohn, 1992; Lovett & Forbus, 2011b).

The work also introduces a new paradigm that varies stimulus complexity, to examine its interaction with categorical perception. We find that as stimuli become more complex, participants increasingly rely on categorical relations. This finding suggests categorical relations may play a pivotal role in supporting difficult visual comparisons.

We begin with background on the categorical/metric distinction, including evidence that these two information types are represented independently in the brain. We then describe computational modeling work which makes concrete predictions about the categorical relations used in visual comparison. We present two behavioral experiments which test the model predictions. We close by considering directions for future work.

Background

Metric representations, also known as *coordinate* or *quantitative*, describe objects in a mental coordinate system analogous to the real world. Categorical representations, also known as *qualitative*, carve up this space into regions and apply labels to them. Categorical representations are redundant, containing the same information found in metric representations. However, research suggests people rely on categories to help them remember and compare images.

For example, when participants are asked to remember a dot's location in a circle and later generate it, their memories are biased by the quadrant in which the dot was located (Huttenlocher, Hedges, and Duncan, 1991). If the dot was in the upper left quadrant, then participants tend to generate a location skewed towards the center of the upper left quadrant. This suggests that participants integrate their metric and categorical representations of the dot's location, using the two together to help them remember.

Categorical perception of color suggests a similar integration (Bornstein & Korda, 1984). Color patches are designed such that the quantitative difference in their hues is always




Figure 1. Stimuli from Kim & Biederman (2012).

the same. In some cases they lay on either side of a category boundary (**blue** and **green**), while in other cases they are in the same category. Participants are better able to distinguish patches with different color categories, indicating that they use the categories to aid them in making the comparison.

Categorical perception of object features has received significant attention (e.g., color: Roberson, Pak, & Hanley, 2008; Regier & Kay, 2009; size: Kosslyn et al., 1977; location: Maki, 1982; angles: Rosielle & Cooper, 2001). However, to our knowledge only one study (Kim & Biederman, 2012) has demonstrated categorical perception of between-object relations. Here, participants performed a match-to-sample task, where they had to distinguish an identical image from a distractor image (Figure 1: Given the top image, decide which of the lower images is the same). It was easier to distinguish the distractor image if it contained a different categorical relation, e.g., one object was now resting on the other. However, the images were always tilted to different orientations. This may have biased participants towards using categorical information which is less orientation-dependent.

The categorical perception paradigm is valuable because it provides direct evidence that humans are encoding and comparing a particular set of visual categories. Thus, the lack of research on relational categorical perception is concerning. It is also surprising, given the evidence that our brains may be hardwired to encode categorical relations.

Distinct neural pathways for relational categories?

Kosslyn and colleagues have argued that the right brain hemisphere specializes in metric spatial relations, while the left hemisphere specializes in categorical relations. This was first demonstrated with behavioral studies, showing that participants judge coordinate relations (e.g., “Are the objects X inches apart?”) faster when stimuli are presented to the left visual field, connected to the right hemisphere. On the other hand, participants judge categorical relations (e.g., “Is one object *above* another?”) faster when stimuli are presented to the right visual field (Kosslyn et al., 1989). A larger set of follow-up studies have largely supported these findings, though the results have not always been consistent (see Jager & Postma, 2003 for a review).

More recently fMRI studies have found that spatial judgments generate activity in SPL (superior parietal lobule), with greater activation in the left hemisphere for categorical and in the right hemisphere for metric (Trojano et al., 2002;

van der Ham et al., 2009). In addition, some categorical judgments trigger activation in the left IPL (inferior parietal lobule) (Trojano et al., 2002; Amorphanth, Widick, & Chatterjee, 2010). The SPL and IPL are part of the dorsal stream, associated with encoding spatial location.

Biederman and colleagues measured brain activity when participants compared images to determine if they were the same or different (Kim & Biederman, 2012; Hayworth, Les-croart, & Biederman, 2011). Some differences changed a categorical relation, while others were purely metric. They found that categorical relation changes triggered activity in the lateral occipital complex (LOC), part of the ventral stream associated with object recognition.

These two sets of findings suggest categorical relations may be encoded in disparate parts of the visual pipeline, used for locating objects in space and identifying the objects. If so, these relations may play a critical role in visual perception. In the following section, we describe computational modeling work addressing the question of *how* categorical relations are used.

Evidence from qualitative spatial reasoning

There is also evidence for relational categories from the field of qualitative spatial reasoning. In this field, researchers often develop spatial calculi, i.e., sets of categorical relations for describing space in a computational system. For example, the RCC8 relations (Randall, Cui, & Cohn, 1992) describe topological relationships between pairs of two-dimensional objects. While most research in this area is not concerned with human cognition, Klippel et al. (2012) have demonstrated a sensitivity to RCC8 relations in humans.

In the Klippel studies, participants are given a series of animations and asked to group them together however they’d like. The animations, which contain simple two-dimensional objects, all start the same, but they end on different frames. In particular, they vary in the topological relation that exists between the objects in the final frame. The experiments find that animations that end on a common RCC8 relation tend to be grouped together, suggesting, for example, that participants make a categorical distinction between when two objects are touching at the edges, and when they are overlapping.

While the evidence from the grouping studies is important, it is somewhat limited by the complexity of the task—given an instruction to group animations together, it is difficult to say what criteria people may use. Thus, the evidence may be strengthened by using a simple comparison task as in the categorical perception studies.

Computational Modeling Work

We previously used computational models to explore categorical relations and their role in visual problem-solving

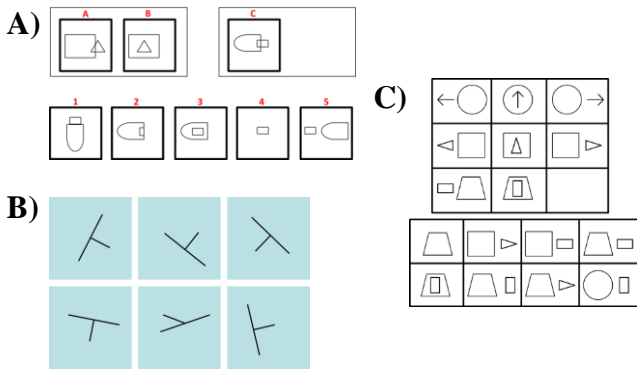


Figure 2. Visual problem-solving tasks. A) Geometric Analogy. B) Oddity Task. C) Matrix Problem.

(Lovett et al., 2009; Lovett & Forbus, 2011; Lovett & Forbus, in prep). Figure 2 presents the types of problems explored. Each problem requires comparing images and thinking about their commonalities and differences.

The models build on structure-mapping theory (Gentner, 1983; Falkenhainer, Forbus, & Gentner, 1989), a domain-general theory of comparison and analogy in humans. Structure-mapping compares two relational representations, aligning the common relational structure to identify corresponding elements. It provides a fast, efficient way of to compare complex stimuli and identify commonalities and differences. For it to work on images, the images must be represented as a set of objects, categorical features for each object, and categorical relations between the objects (Sagi, Gentner, & Lovett, 2012).

Considering the tasks in Figure 2, our hypothesis was that people would be guided by categorical features and relations, and thus when they compared the images, their attention would be drawn to categorical commonalities and differences. To model the tasks, it was necessary to develop a categorical vocabulary, a set of features and relations that would be automatically encoded by the models during problem-solving. This vocabulary was chosen based on the constraints of the problem-solving tasks, as well as psychological theories and previous modeling work.

Figure 3 presents an example, three categorical relations for describing topology. The relations are: 1) **Intersect**: The edges of two objects intersect. 2) **Overlap**: There is a common region found within two objects. 3) **Contains**: One object contains another. Note that these three relations make similar distinctions to the RCC8 relations, and indeed they were inspired by those relations. However, unlike RCC8, these relations are not mutually exclusive; for example, if one object **contains** another, they may or may not have edges that **intersect**. For a complete description of the categorical vocabulary used in the models, see (Lovett & Forbus, 2011b).

Overall, the models proved effective at matching human performance on the tasks. In addition, the models generated

several predictions about categorical relations and how humans use them. Two of these predictions were:

Prediction 1: The models make concrete predictions about which categorical relations people will encode. For example, they predict the **Intersect/Overlap/Contains** relations described above.

Prediction 2: The models predict that as a task becomes more difficult, or as visual stimuli become more complex, people will increasingly rely on categorical relations. There are two reasons for believing this: a) Categorical relations can be compared quickly and efficiently via structure-mapping. Thus, they can guide individuals in finding the commonalities and differences in complex stimuli. b) Categorical relations are encoded in a simple, abstract manner, using symbols such as **intersect**. Thus, they may be remembered more easily than precise metric values (Rosielle & Cooper, 2001). As memory demands increase, categorical relations should be increasingly useful for recording an image's critical features.

Below, we describe two experiments designed to test these predictions, focusing on the three topological relations.

Experiments

We tested relational categorical perception using pairs of circles. The stimuli (Figure 4) were designed such that the smaller circle's position changes the same amount between each adjacent pair. However, in some cases this results in a categorical change to the relations between the circles. For example, when the change causes the two circles to touch, an **intersect** relation now exists between them.

Experiments 1 and 2 tested participants' ability to notice a change between adjacent circle pairs in Figure 4. There are 15 total pairs: the 8 pairs in Figure 4, and 7 more pairs created by moving the smaller circle to the left of the large circle. There are 14 intervals between these adjacent pairs. Of these, 8 intervals introduce a categorical change (e.g., a change from **contains** to **intersect + contains**), and 6 intervals are purely metric (e.g., the two adjacent pairs that are both **intersect + overlap**).

A key question was whether the advantage for detecting categorical changes would increase with stimulus complexity. For these experiments, we operationalized complexity as set size. Participants had to detect differences in one, two, or three circle pairs (Figure 5B).




Figure 3. Topological relations.




Figure 4. Experimental stimuli covering a range of topological relations.




Figure 5. A) Typical trial. B) Displays with one, two, or three circle pairs.

General Methods

These experiments used a sequential same-different paradigm. Figure 5A depicts a typical trial with two pairs. Participants were cued to attend to two quadrants. After 500 ms, circle pairs appeared in each quadrant and were visible for 2500 ms. There was a 1000 ms delay (with a mask for the first 250 ms), and then the two circle pairs reappeared. Participants pressed one key if the pairs were the same, and another key (with their other hand) if either pair was different. When participants were incorrect, the word “Wrong” appeared on the screen for 1000 ms. During this time, if a pair actually had changed, the display flipped between the original and changed pair repeatedly to highlight the difference the participant had failed to notice.

Half the trials were *different* trials where one of the pairs changed, while the other half were *same* trials. Half the different trials were categorical differences, while the other half were purely metric differences. The difference appeared equally often in the four quadrants of the screen. Other factors were randomized: each trial used either the pairs in Figure 4 or the corresponding pairs with the smaller circle on the left. The non-critical pairs (the ones which appeared on the screen but did not change), were randomly chosen from this set, subject to the following constraint: no display should contain two instances of the same pair.

The large and small circles were assigned different colors, either red/green or blue/yellow. This was to decrease the chance that they would be perceptual grouped as a single object. The colors of the circles, as well as the keys for responding “same” and “different,” were counter-balanced across participants.

Experiment 1

Experiment 1 contained two trial types: trials with a single pair, and trials with two pairs (Figure 5B). This provided initial evidence for an interaction between categorical perception and stimulus complexity. The trial types were intermixed over a total of 256 trials. Participants received one break, after the first 128 trials.

15 Northwestern University students took part in this study for class credit.

Results

Figure 6 shows the accuracies across conditions. We analyzed accuracy on different trials via a two-way repeated measures ANOVA with number of pairs (one vs. two) and difference type (categorical vs. metric). There was a main effect for number of pairs, $F(1,11) = 31.6$, $p < .001$, partial $\eta^2 = .758$, and a main effect for difference type, $F(1,11) = 43.8$, $p < .001$, partial $\eta^2 = .693$. These indicate that a) it was easier to spot a difference with only one pair, and b) it was easier to spot a categorical difference. This finding demonstrates categorical perception of the topological relations.

Finally, there was a significant interaction between number of pairs and difference type, $F(1,11) = 17.7$, $p = .001$, partial $\eta^2 = .559$. This interaction matches the prediction and is apparent in Figure 6: as the number of pairs increased, there was a small cost for categorical differences, but a much greater cost for metric differences.

Experiment 2

Experiment 2 contained two trial types: trials with two pairs, and trials with three pairs. This allowed us to test how our predictions scale up with increased complexity (note that it would have been difficult to run all three set sizes in a single experiment due to time constraints). Due to the higher difficulty level, participants received three breaks. The experiment was otherwise the same as Experiment 1.

15 participants, aged 18-35, took part in this study for \$10.

Results

Figure 6 shows the accuracies across conditions. We analyzed accuracy on different trials via a two-way repeated measures ANOVA with number of pairs (two vs. three) and difference type (categorical vs. metric). There was a main effect for number of pairs, $F(1,11) = 32.7$, $p < .001$, partial $\eta^2 = .700$, and a main effect for difference type, $F(1,11) = 35.2$, $p < .001$, partial $\eta^2 = .715$. Again, the task was easier when there were fewer pairs, or when the difference was categorical.

There was a significant interaction between number of pairs and difference type, $F(1,11) = 5.4$, $p = .036$, partial $\eta^2 = .278$. Note, however, that the effect size was considerably smaller than in Experiment 1 (.559 vs. .278). In this experiment, though the cost for increased stimulus complexity was greater for metric differences, there was also a high cost for categorical differences.

Discussion

Overall, the results demonstrate categorical perception of topological relations. Participants identified differences far more accurately when those differences changed the categorical relations between the circles. Furthermore, as predicted by the computational models, the reliance on categorical relations increased when the stimulus complexity (the number of pairs) increased.

Interestingly, the interaction was more pronounced going from one to two pairs than going from two to three. This may indicate a qualitative difference in participants' strategies for one vs. multiple pairs. Faced with only a single pair, participants may have encoded both categorical and metric information in detail, allowing them to detect both difference types with high accuracy (though post-hoc analysis

found higher accuracy for categorical differences even here, $t = 3.3$, $p = .005$). Faced with two or three pairs, participants may have focused on the more easily remembered categorical information, resulting in a large cost to detection of metric differences.

Mean accuracy for detecting metric differences with three pairs was only .567. This might appear to be near-chance performance. However, accuracy on same trials with three pairs was considerably higher, .826. This indicates a response bias in the participants. We hypothesize that participants generally responded "same" if they were unable to detect a difference, only responding "different" if the difference was salient. Thus, although participants only detected about half the differences in the hardest condition, they did not appear to be guessing.

There was one additional finding of interest. After the experiment, participants were asked if they used any particular strategy to perform the task. One third (5/15) of the participants in Experiment 2 reported that they named the different stimuli. For example, the rightmost two pairs in Figure 4 might be named "near" and "far" based on the distance between the circles. This strategy appears to be an attempt to *categorize* the metric information, applying a label to it so that it can be remembered more easily. No participants reported using this strategy in Experiment 1, perhaps because the task was easier.

Conclusion

These experiments help to bridge the gap between qualitative representations in computational models and categorical representations in humans. They provide evidence that humans encode topological relations categorically, and use them to remember and compare images. When images be-

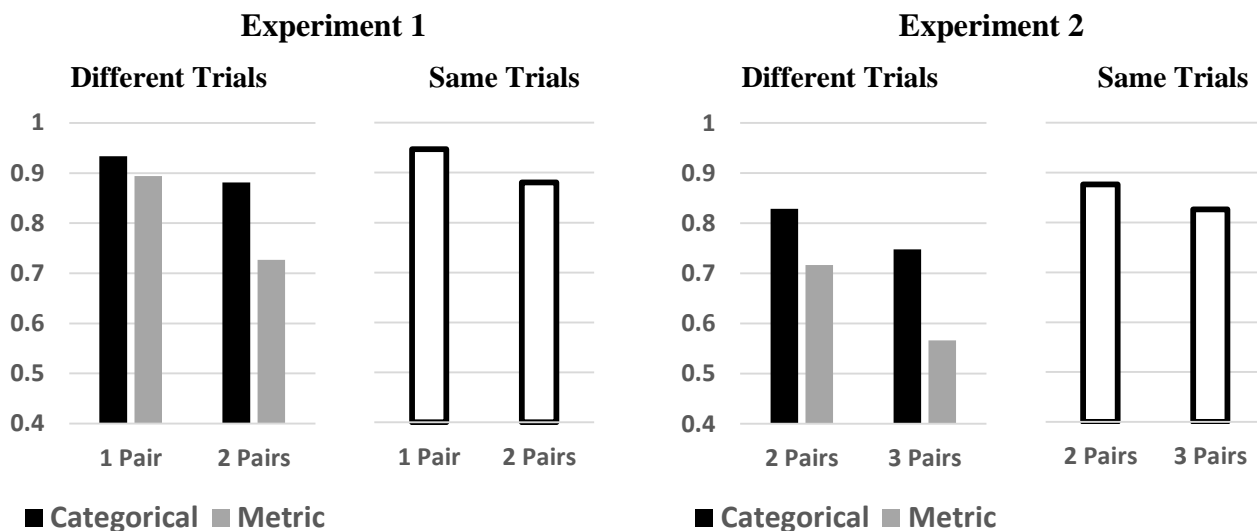


Figure 6. Mean accuracies in the two experiments.

come more complex, as with two or more shape pairs, people rely more on the easily remembered categories, failing to notice changes in metric properties. The present paper focuses on the categorical/metric distinction predicted by computational models. Further analysis is required to tease apart each individual relation's contribution, e.g., the ease of remembering an **overlap** relation versus an **intersect** relation.

When remembering complex visual stimuli, individuals may rely on verbal memory to encode key features. In the future we plan to run this experiment with a verbal interference task, where participants must repeat certain words in their mind throughout each trial. Verbal interference disrupts verbal coding of visual features. We foresee two possible outcomes: A) The categorical advantage might increase because participants can no longer assign names to metric properties (as some did in Experiment 2). B) The categorical advantage might *decrease* because relational categories, associated with symbolic processing in the left hemisphere (Kosslyn et al., 1989), might require the same cognitive resources as are used in verbal rehearsal. The second outcome would be particularly illuminating, as it would tell us more about how categorical relations are encoded in the brain.

We hope the present paradigm can be used more broadly beyond topological relations. We see this work as establishing a template. By applying this template across a range of objects and relations, researchers can arrive at a better understanding of which categorical relations people encode and how those relations are used to support memory and comparison.

Acknowledgements

This research was supported by an NIH training grant in Human Cognition at Northwestern University (T32 NS047987).

References

Amorapanth, P. X., Widick, P., & Chatterjee, A. (2010). The neural basis for spatial relations. *Journal of Cognitive Neuroscience*, 22(8), 1739-1753.

Borstein, M. H., & Korda, N. O. (1984). Discrimination and matching within and between hues measured by reaction times: Some implications for categorical perception and levels of information processing. *Psychological Research*, 46, 207-222.

Chen, L. (2005). The topological approach to perceptual organization. *Visual Cognition*, 12(4), 553-637.

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

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

Hayworth, K. J., Lescroart, M. D., & Biederman, I. (2011). Neural encoding of relative position. *Journal of Experimental Psychology: Human Perception and Performance*, 37(4), 1032-1050.

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

Jager, G., & Postma, A. (2003). On the hemispheric specialization for categorical and coordinate spatial relations: A review of current evidence. *Neuropsychologia*, 41, 504-515.

Lovett, A., & Forbus, K. (in prep). Modeling comparison and abstraction during visual problem-solving.

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

Lovett, A. & Forbus, K. (2011b). Organizing and representing space for visual problem-solving. *Proceedings of the 25th International Workshop on Qualitative Reasoning*. Barcelona, Spain.

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

Kim, J. G., & Biederman, I. (2012). Greater sensitivity to nonaccidental than metric changes in the relations between simple shapes in the lateral occipital cortex. *NeuroImage*, 63, 1818 - 1826.

Klippel, A., Yang, J., Wallgrün, J. O., Dylla, F., & Li, R. (2012). Assessing similarities of qualitative spatiotemporal relations. In C. Stachniss, K. Schill, & D. H. Uttal (Eds.), *Spatial Cognition 2012* (pp. 242-261). Berlin: Springer.

Kosslyn, S. M., Koenig, O., Barrett, A., Cave, C. B., Tang, J., & Gabrieli, J. D. E. (1989). Evidence for two types of spatial representations: hemispheric specialization for categorical and coordinate relations. *Journal of Experimental Psychology: Human Perception and Performance*, 15, 723-735.

Kosslyn, S. M., Murphy, G. L., Bemesderfer, M. E., & Feinstein, K. J. (1977). Category and continuum in mental comparisons. *Journal of Experimental Psychology: General*, 106(4), 341-37.

Maki, R. H. (1982). Why do categorization effects occur in comparative judgment tasks? *Memory & Cognition*, 10(3), 252-264.

Randell, D. A., Cui, Z., & Cohn, A. G. (1992). A spatial logic based on regions and connection. *Proceedings of the 3rd International Conference on Knowledge Representation and Reasoning*.

Regier, T., & Kay, P. (2009). Language, thought, and color: Whorf was half right. *Trends in Cognitive Science*, 13(10), 439-446.

Roberson, D., Pak, H., & Hanley, J. R. (2008). Categorical perception of colour in the left and right visual field is verbally mediated: Evidence from Korean. *Cognition*, 107, 752-762.

Rosielle, L. J., & Cooper, E. E. (2001). Categorical perception of relative orientation in visual object recognition. *Memory & Cognition*, 29(1), 68-82.

Sagi, E., Gentner, D., & Lovett, A. (2012). What difference reveals about similarity. *Cognitive Science*, 36(6), 1019-1050.

Trojano, L., Grossi, D., Linden, D. E., Formisano, E., Goebel, R., Cirillo, S., & Di Salle, F. (2002). Coordinate and categorical judgments in spatial imagery. An fMRI study. *Neuropsychologia*, 40(10), 1666-1674.

van der Ham, I. J., Raemaekers, M., Van Wezel, R. J., Oleksiak, A., & Postma, A. (2009). Categorical and coordinate spatial relations in working memory: An fMRI study. *Brain Research*, 1297, 70-79.