

MAGI: Analogy-based Encoding Using Regularity and Symmetry

Ronald W. Ferguson
Qualitative Reasoning Group
Institute for the Learning Sciences
Northwestern University
Evanston, Illinois 60202
ferguson@ils.nwu.edu

Abstract

Analogy has always been considered a mechanism for interrelating distinct parts of the world, but it is perhaps just as important to consider how analogy might be used to break the world into comprehensible parts. The MAGI program uses the Structure-Mapping Engine (SME) to flexibly and reliably match a description against itself. The resulting mapping pulls out the two maximally consistent parts of the given description. MAGI then divides out the parts of the mapping and categorizes the mapping as symmetrical or regular. These parts may then be used as the basis for new comparisons. We theorize that MAGI models how people use symmetry and regularity to facilitate the encoding task. We demonstrate this with three sets of examples. First, we show how MAGI can augment traditional axis detection and reference frame adjustment in geometric figures. Next, we demonstrate how MAGI detects visual and functional symmetry in logic circuits, where symmetry of form aids encoding symmetry of function. Finally, to emphasize that regularity and symmetry detection is not simply visual, we show how MAGI models some aspects of expectation generation in story understanding. In general, MAGI shows symmetry and regularity to be not only pretty, but also cognitively valuable.

Introduction: Why regularity and symmetry aren't (just) pretty


Regularity and symmetry are phenomena strangely divided between disciplines. Researchers in computer vision (Witkin & Tenenbaum, 1983) and perceptual psychology (Palmer, 1985; Rock, 1983) have long recognized regularity and especially symmetry as important, but have understood it strictly as a perceptual effect. Although researchers in analogy might easily agree that symmetry and regularity must involve some form of self-similarity, this community has produced little work in the area, perhaps due to an emphasis on problem solving and learning, rather than encoding.

This paper is an attempt to bridge this gap by recasting symmetry and regularity as analogical processes that operate on structured but undivided representations in the world. There are two central theoretical claims in the MAGI model. The first is that regularity and symmetry are like analogy--they work by mapping a maximal common set of structurally interconnected relations, but within a


single description instead of between separate base and target descriptions. The second is illustrated by Figure 1. Regularity and symmetry are not strictly perceptual, but may be found in any task involving the encoding of relational knowledge structures. For example, regularity and symmetry may be found in imperfect figures (a), in diagrams (b), or in story narratives (c). To support these two claims, we have constructed MAGI, a system that uses SME to detect regularity and symmetry, and can handle all these cases.



A. Goblet figure



B. SR Latch Diagram



C. Representation of The Gift of the Magi
by O. Henry (simplified example)

Figure 1: Three types of symmetry handled by MAGI

Self-similarity Mapping Stage

Map a description to itself using the I-SME algorithm (Forbus, Ferguson & Gentner, to appear) based on the literal similarity algorithm in SME (Falkenhainer, Forbus, & Gentner, 1989). In addition to the constraints provided by I-SME, block matches between an entity or expression and itself, but allow such an item to match to itself when it is aligned as the argument of two non-identical expressions.

Mapping Analysis Stage

Create mirrored mapping pairs. For each mapping M_i , create a pair of mirrored mappings. A mirrored mapping pair (M_i, N_i) is a pair of mappings such that each mapping kernel (submapping) in M_i between two expressions has a twin mapping kernel in N_i that maps those expressions in the opposite direction. In many cases for regular mappings, SME will produce the mirrored mappings as two separate, equally-scored mappings (see text). Otherwise, M_i must be split to produce the mirrored mapping pair.

To split a single mapping. Divide mapping into pairs of mirrored mapping kernels, starting with the largest kernel. When kernels overlap with previously paired kernels, align the left-right pairing such that the left side of the new pair has maximum overlap of match hypotheses with left sides of those already paired.

Find symmetry set within mirrored mapping pairs. For each mirrored mapping pair (M_i, N_i) , use M_i . Collect *symmetry core* SR_i by collecting all kernels in M_i that contain at least one pair of symmetrically related entities. The *symmetry set* S_i is then defined recursively as all root mappings in M_i that are either in SR_i or overlap S_i .

Figure 2: Summary of MAGI Algorithm

The MAGI system and Algorithm

How does MAGI work? MAGI uses the Structure Mapping Engine (SME) to match a structured description to itself, taking the description as both the base and target. MAGI then uses the resulting mapping to find two maximally similar subareas of the original description, and classifies the mapping as symmetrical or regular.

Figure 1(c) is illustrative of the process. This figure shows the results of a MAGI run on a simplified example representation of its namesake, O. Henry's short story "The Gift of the Magi." (1992) In this portion of the story, a husband sells his watch to buy an expensive comb for his wife's hair, while the wife sells her hair to purchase a chain for her husband's watch. This single highly-interconnected plotline, when presented to MAGI, produces a mapping that divides the story into two symmetrical subplots, one concerning the husband's actions and the other the wife's. MAGI then interprets the mapping as symmetrical, rather than regular, by finding a core of symmetrical matches between the two **MOTIVE** expressions, which are augmented by regular matches between the **HELD-BY** and **EXPENSIVE** expressions.¹

This section explains the algorithm MAGI uses, beginning with the new version of the Structure Mapping Engine (SME) at the core of MAGI, then moving to the

¹A note on the distinction between regularity and symmetry. Regularity in descriptions comes from the self-similarity of one part to another part. Symmetry is a subtype of regularity that maps one part to another part, but also contains a "core" of relations that map some entities symmetrically, i.e. when A is mapped to B, B is also mapped to A. "Bob likes ice cream and Susan like cake" is regular. "Bob knows Susan and Susan knows Bob" is symmetric. Likewise, when the core symmetry is linked to other regular expressions, the result is a symmetric as a whole. "Bob knows Susan likes cake, and Susan knows Bob likes ice cream" is, as a whole, symmetric, even though the mapping between cake and ice cream is regular.

extensions MAGI makes to SME, and finally ending with the routines MAGI uses to analyze the mapping and classify the mapping as either regular or symmetric. Figure 2 contains a summary.

The Self-Similarity Mapping

The first stage of MAGI simple uses I-SME (Forbus, Ferguson & Gentner, to appear) to create a mapping between a description and itself, using the literal similarity rules described in (Falkenhainer, Forbus & Gentner, 1989) which are modified to discourage matches between an item and itself. This is because the literal similarity algorithm, if left unmodified, would produce a verbatim mapping, with every description item (i.e. every expression and entity) mapped to itself. To avoid this rather uninformative result, I-SME could simply block all self-matches, but this also has unfortunate consequences, since self-matches are sometimes needed. For example, suppose that I wanted to find regularity in a story that began with Bob missing his train and walking to the Art Institute in Chicago, and Susan missing a movie and also visiting the Institute. The intuitive regularity in the story maps Bob to Susan, and their mutual misfortunes to each other, but also maps the ^{Art}Institute to itself. The natural interpretation of the regularity in the story would be missed if we blocked all self-matches. Because of cases like this, MAGI takes a flexible approach by allowing self-matches only when the self-match occurs due to the alignment of two non-identical expressions. So (to return to our story) given the expressions (**TRANS Bob Institute**) and (**TRANS Susan Institute**), Bob could not be self-mapped, nor could Susan, but **Institute** could because the two non-identical **TRANS** expressions align it as the second argument.² MAGI's matching rules are a

²Commutative expressions, such as (**EQUAL-LENGTH line1 line2**), or (**MARRIED Jim Della**), are a special case. Such an expression can be matched to itself, but only if the arguments are reversed.

compatible extension of I-SME's rules, as they do not modify I-SME's mapping of separate base and target descriptions.

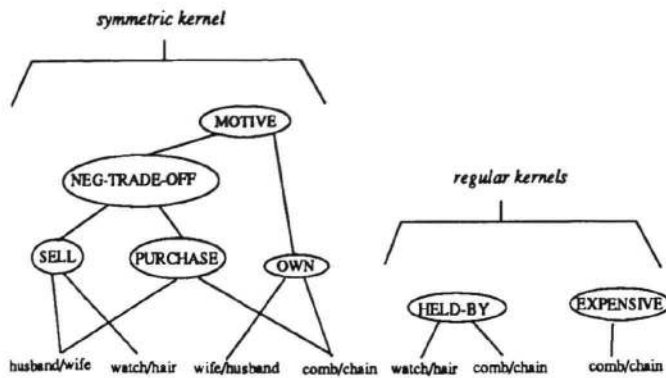


Figure 3 Mapping kernels from Figure 1(c).

The Mapping Analysis Stage

After the mappings are created, they must be analyzed. This involves several steps. First, each mapping must be split into a pair into "mirrored" mappings³. The mappings produced by the self-similarity stage are "mirrored" with one half from the base to the target, and another from the target to the base. For example, for the Magi story, (SELL wife hair) is mapped to (SELL husband watch), but (SELL husband watch) is also mapped to (SELL wife hair). Pairing them removes this redundancy.

After a mapping is paired into two mirrored mappings, the properties of either of the mirrored mappings determines the type of that mapping. Symmetric mappings will contain at least one kernel that will symmetrically match at least one pair of entities. This set of kernels is called the *symmetry core* of the mapping. All kernels that are not part of the symmetry core are considered regular. The symmetry core propagates symmetry to all kernels that overlap it. Thus, the *symmetry set* of the mapping consists of all kernels that are either in the symmetry core, or overlap a kernel within the symmetry set. If all the kernels of the mapping are in the symmetry set, the mapping itself is symmetrical. If some kernels are in the symmetry set, but not all, the mapping is partially symmetrical. Finally, if there are no kernels in the symmetry core, and thus none in the symmetry set, the mapping is regular.

Returning to our initial example from O. Henry, Figure 3 shows the kernels from one half of the first mirrored mapping pair. Figure 3(a) is the symmetry core, since it maps *wife* to *husband* and *husband* to *wife*. The other two kernels are regular, but are included in the final

³ Actually, it is not always necessary to split a mapping to get the two mirrored mappings. When a figure is regular, rather than symmetric, a resulting mapping may be inconsistent with its mirror, and thus will that mapping and its mirror emerge from the self-similarity mapping stage as separate mappings with identical scores. MAGI checks for this condition before executing the algorithm outlined above.

symmetry set because they share entity mappings with the core. Because all kernels in this mapping are included in the symmetry set, this mapping is symmetric.

Examples using MAGI

In this section we will describe three classes of examples where MAGI does regularity and symmetry detection. These three cases span the range from purely perceptual to purely conceptual representations. The first example shows how MAGI can detect axes of symmetry in simple geometric figures, and how these axes are helpful in building representations. The second example, where MAGI is used to construct representations of logic diagrams, takes the results of the perceptual examples to show how symmetry is useful in real world tasks. More specifically, it shows that symmetry of form may facilitate encoding of symmetrical function in logic circuit diagrams. Finally, after doing perceptual and perceptual/conceptual examples, we briefly demonstrate MAGI in a conceptual domain by showing how MAGI may model the detection of regularity and symmetry in the narrative structure of fables.

Using MAGI with GeoRep to model spatial reasoning.

Before beginning the examples, we should briefly note how the geometric representations are created. The described spatial representations were created using a system called GeoRep. GeoRep is a diagram representation tool kit currently in development that takes FIG diagram files (from the public domain drawing program Xfig, by Brian V. Smith) and turns them into predicate calculus representations used by a logic-based truth maintenance system (LTMS) and its associated rule engine (Forbus and deKleer, 1993). The LTMS then can infer further representations that can be fed into MAGI. The mappings created by MAGI may also be fed in turn into the LTMS for further inferencing.

GeoRep begins with a diagram consisting of groups of lines, circles, and arcs. It constructs a representation based on the types of connections that may occur among geometric elements (corners, mid-connections, and intersections of lines, and connections, tangents, and abutments for lines and shapes) and on interval relationships (as in Allen, 1983) between proximate parallel lines and other proximate objects within a reference frame. It also has a rudimentary representation of polylines based on an analogy to members of a set. GeoRep can also run domain-specific rules that use the visual representation to build a representation based on what the diagram itself represents. Figure 4 shows the visual and a functional representations generated for a single NAND gate.

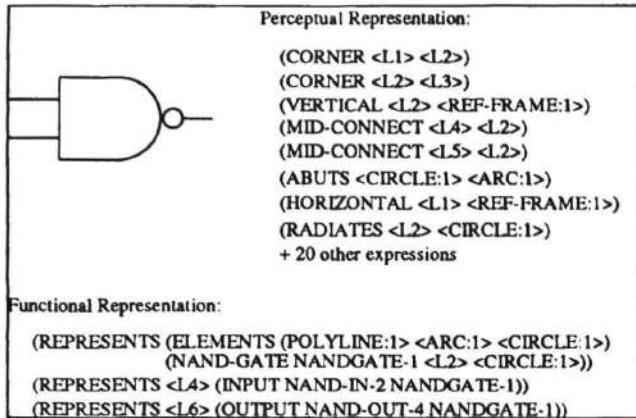


Figure 4 Sample representation for NAND gate figure

Finally, GeoRep has an axis-detection mechanism that uses the self-similarity mapping produced by the MAGI. The mechanism is similar to the standard Hough transform used in computer vision (Duda and Hart, 1983), but for pairs of lines rather than for single lines. Each mapped pair of lines votes for a particular pair of axes and frame of reference. When the vote is tallied, the winning reference frame is adopted, along with any compatible axis with a vote greater than half that of the winning frame of reference.

Example Set 1: Geometric Figures




Figure 5: Symmetry need not be exact

Geometric figures are traditionally used to illustrate aspects of symmetry and regularity, but the geometric examples given here demonstrate how MAGI is different from traditional approaches. Symmetry in computer vision has conventionally meant the detection of an axis using a transform acting on points in the boundary of a figure (e.g., Blum, 1973; Brady, 1984). Neither of these methods handle figures that contain symmetry within their boundaries (such as in Figures

1(a) and 1(b)). Each of the following figures (Figures 5-7) were represented using GeoRep and then processed by MAGI. Each set of figures shows a different aspect of how thinking of regularity as analogical mapping lends it a power not traditionally associated with regularity or symmetry.

The sketchy goblet drawn in Figure 5 indicates the most obvious advantage of using propositional representations--symmetry need not be exact. The relational structure is important, and contains the same flexibility as normal SME matches. Secondly, because SME produces multiple mappings, MAGI can produce multiple interpretations, as Figure 6 illustrates. In mapping the top figure (A), it is

possible to see the squares as regular or symmetrical (either they are two squares in a row or as two squares facing each other). MAGI produces both interpretations, but favors the regular interpretation. In the bottom figure (B), the facing trapezoids are interpreted as symmetrical, with no regular interpretation.




Figure 6: MAGI can produce both regular and symmetric interpretations, sometimes for the very same descriptions

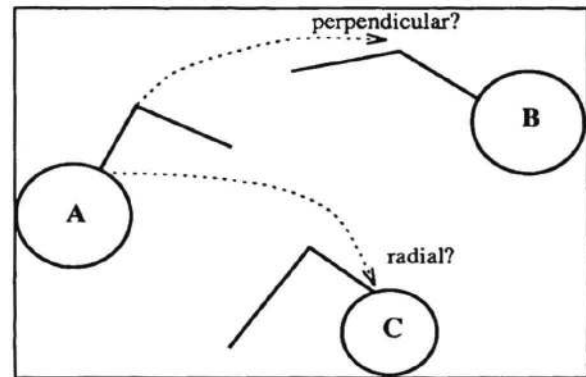


Figure 7: Candidate inferences from MAGI mappings

Finally, mappings produced by SME not only list correspondences but also generate candidate inferences when the mapping intersects base structure (Falkenhainer, Forbus & Gentner, 1989). The same effect occurs in regular or symmetric mappings in MAGI. Figure 7 contains a triad of three circles with attached polylines. Object A contains a polyline with a perpendicular corner, and that line is radially connected to the circle. When given this figure, MAGI produces two mappings--one from

A to B, and another from A to C--with both mappings




Figure 8: Symmetrical and non-symmetrical flip-flop diagrams

aligning the corners, polylines and circles in the obvious way. In the mapping between figures A and B, because the corners match, MAGI produces a candidate inference that B's corner is also perpendicular. In the second mapping between A and C, for similar reasons, the radial connection of A is expected at C as well. A geometric reasoning system could use these candidate inferences to guide further exploration of the figure by attempting to confirm the candidate inferences. If the inference turn out to be true, the inferences have cut down on the search space; if they are not true, they may become alignable differences (Markman and Gentner, 1993) that are salient.

Even at the geometric level, then, understanding symmetry and regularity in terms of analogical mapping offers a wealth of new possibilities that are simply not addressed by conventional forms of symmetry detection.

Example Set 2: Flip-flop diagrams: Symmetrical form suggests symmetrical function

Geometric figures are helpful for demonstrating the various capabilities of MAGI, but say little about its cognitive usefulness, or about encoding larger more complex figures. To understand encoding in a more realistic domain, we have turned to logic circuits. We are currently using diagrams of various latches and flip-flops (Figure 8). Flip flops and latches are among the first sequential logic devices taught to electrical engineering and computer science students, and the functioning of these

devices is symmetrical (the two input/output pairs can be switched without changing the functioning of the device). Diagrams for flip-flops and latches tend to be symmetrical in form. It seems intuitive that the symmetry of form should simplify the recognition of the symmetry of function.

We tested this hypothesis on three different flip-flop circuits--an SR latch, a clocked SR flip-flop, and a JK flip-flop. For each circuit, two functionally equivalent diagrams were given to GeoRep (shown in Figure 8), where one was symmetrical and the other asymmetrical. The symmetrical diagram was taken from a standard digital design textbook (Johnson & Karim, 1987), and was then modified to produce the asymmetrical version. GeoRep was then given two sets of rules--the perceptual representation rules described above, and a set of function recognition rules specific to logic diagrams (e.g. describing the NAND gate, the device inputs and outputs, and the feedback paths). Links between form and function were made explicit by expressions of the form (REPRESENTS X Y), where X was some geometric object in the diagram and Y was the logical path or connection represented (see Figure 4).

MAGI was then run twice on each figure. First, MAGI ran on the perceptual representation only. Then, after the functional rules were run, the new functional information was added to the description, and MAGI incrementally extended the mapping using I-SME. In each case, the result was the same as in Figure 8. The initial run on the perceptual representation produced a clear symmetry mapping on the symmetrical diagram, while the asymmetrical diagram would manage only a partial symmetry. Usually the NAND gates themselves and the device inputs and outputs would map correctly, but the other polylines would map in erratic ways (see Figure 8 for the mappings). Then the functional representation was added, and the mapping incrementally extended. For the symmetrical diagram, the symmetry of function in the flip-flop cleanly integrates into the symmetry of form for the symmetrical diagram. In the nonsymmetrical diagram, the functional representation could be added, but the links between form and function were not consistent with the initial perceptual mapping, and were thus lost. Also, the mapping for the function is clearly lead astray, mapping logical paths in the figure that have nothing to do with each other (see dashed lines in Figure 8(B)). If the nonsymmetrical representation was remapped from scratch, however, with both the functional and perceptual parts of the description, the functional relationships would carry more weight, and the proper functional mapping produced.

Example Set 3: Using regularity and symmetry in understanding social narratives

Symmetry makes revenge sweeter, and returned favors more satisfying. Reciprocity of this type plays a key role in many social situations, and thus are central to understanding stories about social situations (Lehnert, 1981). To test whether the perception of regularity and symmetry set up expectations in the reader, we gave

subjects sets of short fables. Half the fables were symmetrical; the others were non-symmetrical variants with the same ending and characters. The question of interest is whether subjects will generate inferences based on the symmetric structure of the stories. For each fable, the subjects were asked to choose one inference from a list of four. One of these was designed to be the symmetric inference; the others were designed to be plausible inferences. The results for 24 subjects showed sensitivity to symmetric structure: subjects choose the symmetric inference 60 percent of the time for the symmetric stories, but only 38 percent of the time for non-symmetric stories. We are currently working on replicating these effects with MAGI.

Problems with the MAGI model, and areas to explore

Currently, there are several limitations to the MAGI model. It does not handle regularity or symmetry that is not binary in some sense. For example, it could not handle a diagram of a brick wall, nor separate out the plot line of 'The Three Little Pigs.' Also, there is a distinct tendency for MAGI to produce rotational mappings on visual figures (as noted, for example, in Figure 6). Such rotational mappings are technically correct and somewhat novel, but people seem to generate them much less than MAGI does, which makes them suspect.

Conclusion

Any cognitive model of symmetry and regularity must meet three criteria. First, it must show why regularity and symmetry are useful. Second, it should show a clear distinction between regularity and symmetry, since this distinction is clear to people, and thus must explain the role of an axis in symmetrical representations. Finally, there should be an account of how asymmetry is handled, and why it is sometimes stark and other times less noticeable. We believe that the MAGI model satisfies these criteria.

The maturity of the current models of analogical mapping provide the means for a congruent model of analogical encoding using symmetry and regularity. Symmetry and regularity should be understood, not as perceptual, but as analogical processes of encoding that often operate on perceptual representations. MAGI's success in a variety of domains provides clear evidence for the strength of this view. Symmetry and regularity, from a cognitive perspective, are not gems, but pick-axes. Though they are aesthetically pleasing, underneath the pretty covers are incredibly useful tools that help us pull apart and understand the world.

Acknowledgments

Special thanks to my advisors, Ken Forbus and Dedre Gentner. Thanks also to Laura Allender, Alex Aminoff, Larry Birnbaum, Paul Cooper, Ric Feifer, Will Fitzgerald, Art Markman, and David Uttal for helpful feedback on this topic. The fable pilot study was conducted in collaboration with Dedre Gentner and Alex Aminoff. This research was

supported by the Cognitive Science Division of the Office of Naval Research and by a Cognitive Science Interdisciplinary Research Fellowship from Northwestern University to the author.

References

- Allen, J.F. (1983). Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26:1, 832-843.
- Blum, H. (1973). Biological shape and visual science (part 1). *Journal of Theoretical Biology*, 38,205-245.
- Brady, M. (1983). Criteria for representations of shape. In J. Beck, B. Hope, & A. Rosenfeld (Eds.), *Human and Machine Vision*. New York: Academic Press.
- Duda, R.O., & Hart, P.E. (1987). Use of the Hough transformation to detect lines and curves in pictures. *Communications of the ACM*, 15:1, 11-15.
- Falkenhainer, B., Forbus, K.D., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence*, 41,1-63.
- Forbus, K.D., & de Kleer, J. (1993). *Building Problem Solvers*. Cambridge: MIT Press.
- Forbus, K.D., Ferguson, R.W., & Gentner, D. (to appear). Incremental structure-mapping. *Proceedings of the Cognitive Science Society*, August, 1994..
- Henry, O. (1992). *The Gift of the Magi*. Downloadable extext. Project Gutenberg.
- Johnson, E. L., & Karim, M.A. (1987) *Digital Design*. Boston: PWS Publishers.
- Lehnert, W.G. (1981). Plot units and narrative summarization. *Cognitive Science*, 4, 293-331.
- Markman, A., & Gentner, D. (1993). Splitting the differences: A structural alignment view of similarity. *Journal of Memory and Language*.
- Palmer, S. (1985). The role of symmetry in shape perception. *Acta Psychologica*, 59, 67-90.
- Rock, I. (1983). *The Logic of Perception*. Cambridge, Mass: The MIT Press.
- Witkin, A. P., & Tenenbaum, J.M. (1983) On the role of structure in vision. In J. Beck, B. Hope, & A. Rosenfeld (Eds.), *Human and Machine Vision*. New York: Academic Press.