Fractal Representations and Core Geometry

Keith McGregor KEITH.MCGREGGOR@GATECH.EDU
Ashok Goel ASHOK.GOEL@CC.GATECH.EDU

College of Computing, Georgia Institute of Technology, Atlanta, GA 30308 USA

Abstract

Dehaene (2006) has posited that humans possess a set of core geometric abilities regardless of culture, language or education. In this paper, we develop a computational model (called CoreGeo) of core geometric abilities based on a more general theory of fractal analogical reasoning. Fractal analogical reasoning enables the calculation of confidence in an answer and the automatic adjustment of level of resolution if the answer is found to be ambiguous. We present results from running CoreGeo on Dehaene's test of core geometry problems. We also compare our results with Dehaene's results from different cultures and show that CoreGeo performs about as well as humans at core geometry tasks.

1. Mathematical Reasoning and the World around Us

What is the origin of human mathematical abilities? Does geometry constitute a core set of abilities present in humans, regardless of culture, language or education? Lakoff & Núñez (2000) raised the first of these questions in their analysis of human mathematical abilities from a cognitive science perspective. Their arguments, principally that the embodied mind of humans creates mathematics and therefore it is subject to analysis using cognitive science methodologies, suggest that there may be innate principles in the mind which afford mathematical and geometric reasoning capabilities. Lakoff & Núñez discuss number discrimination in human babies, and in particular look at subitizing (the ability to determine the number of objects presented from a glance), drawing on the work of others (among them, Mandler & Shebo 1982; Trick & Pylyshyn, 1993, 1994), to note that subitizing is not a pattern recognition process. They point to Dehaene & Cohen's (1994, 1996) work with patients who have suffered injury which prevents them from attending to things in their environment in a serial fashion (and therefore cannot count them), but nonetheless perform limited subitizing. Not unexpectedly, statements such as these, and others made in their book have attracted much criticism to Lakoff & Núñez. For example, Schiralli & Sinclair (2003) appear to take issue with the use of metaphor rather than pattern recognition as the basis of the Lakoff & Núñez argument, citing repeated examples of the derivation of mathematical principles precisely due to the discovery of patterns. Insofar as we are aware, the debate of whether the embodied mind gives rise to mathematics, or whether mathematical principles are otherwise transcendent continues without resolution.

Dehaene and colleagues (2006) posed the second question above: do humans possess a set of core geometric abilities? They designed and conducted a study of spontaneous geometrical knowledge of the Mundurukú, an Amazonian indigene group. The study looked at two nonverbal tests designed to probe conceptual primitives of geometry. It is the first of these tests, inspired by a test administered in an earlier study by Franco & Sperry (1977) of the hemispheric localization of geometric processing in patients with a surgical disconnection between the left and right hemispheres of their brain, that is of interest here. This particular test was designed to probe the Mundurukú's intuitive comprehension of the basic concepts of geometry, including points, lines, parallelism, and the like (Dehaene et al. 2006). For each of these concepts, Dehaene et al. designed an array of six images, five of which incorporated some desired concept, while the sixth image did not. In essence, each of these problems was a test of perceiving visual oddity.

Dehaene et al. (2006) report that the Mundurukú fared very well with core concepts of topology, Euclidean geometry, and basic geometrical figures, but they experienced more difficulty in detecting symmetries and metric properties. The Mundurukú faired poorly on two domains, each of which Dehaene et al. (2006) point out involve the mental transformation of one shape into another, followed by a second-order judgment about the nature of that transformation. Dehaene suggests that perhaps geometric transformations are more inherently difficult mathematical concepts, or that the detection of such transformations may be more difficult in static images. Dehaene also tested for comparison a group of American children and adults, and found that both the American group and the Mundurukú group shared some of the same difficulties on the task, although American adults performed at a higher level overall. This led Dehaene et al. to conclude that there exists some shared (possibly innate) competence for basic geometrical concepts, regardless of culture, language or education (Dehaene et al. 2006).

Dehaene et al. however do not provide an information-processing account of the core geometric abilities: what kinds of knowledge, representations, and reasoning might enable the core geometric abilities? In this paper, we develop a computational model called CoreGeo which is based on a more general theory of fractal analogical reasoning (McGreggor, Kunda & Goel 2011; McGreggor & Goel 2012). Fractal analogical reasoning enables the calculation of confidence in an answer, as well as the automatic adjustment of level of resolution if the answer is found to be ambiguous. We present results from running CoreGeo on Dehaene's test of 45 core geometry problems. We also compare our results with Dehaene's results from different cultures and show that CoreGeo performs about as well as humans at core geometry tasks.

2. Reasoning based upon the Fractal Representation

In our research, we explore the extent to which the fractal representation affords analogical reasoning (McGreggor et al. 2011; McGreggor & Goel 2011, 2012); we briefly present the fractal representation in an appendix below, as it has appeared in previous papers. The representations of the visual input given to our system represent the whole of the scene, and do not segment or otherwise seek to discriminate between objects in the scene: there is no notion of either qualitative or quantitive in the fractal representation, nor is there any distinction between a shape and its edge. Furthermore, the general strategy we developed provides a singular method of using similarity scores (chiefly determined via recall of objects indexed via features from memory) as a means both for declaring analogically derived answers as well as providing evidence for those cases in which the scene should be represented at a different level of abstraction. There are no additional special purpose mechanisms to augment the reasoning.

For these reasons, exploring the effects of the fractal representation and reasoning on the Dehaene problem set seemed prudent. In addition, we wished to explore the extent to which strictly visual representations could afford that agent rudimentary geometric reasoning capacity.

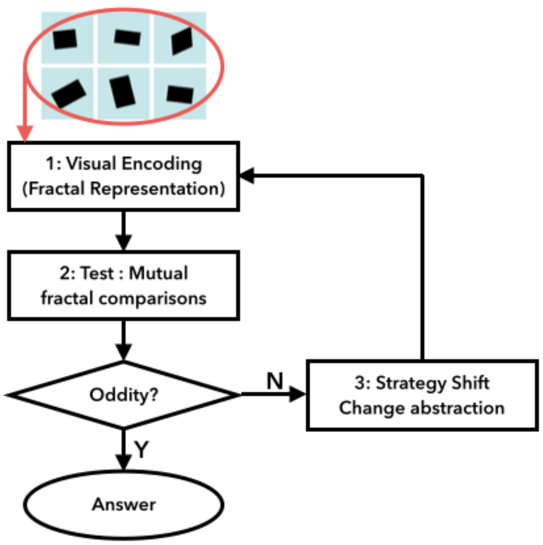

Figure 1. The CoreGeo computational model.

2. 1 The CoreGeo algorithm

Previously, we developed a general strategy of fractal reasoning for solving oddity problems (McGreggor & Goel 2012). We now present an algorithm which is derived from that strategy as a means for direct application to the Dehaene set of core geometry problems. This algorithm, which we call CoreGeo, consists of three phases: a preparatory phase in which the given problem is segmented into its constituent images and the images encoded and represented as mutual fractals; a reasoning phase which determines the similarity between those mutual fractals to see if any of them may prove odd; and a re-representation phase in which the representational strategy is shifted automatically to a different level of abstraction should no single answer stand out. A schematic of the model is shown in Figure 1 above.

In practice and as shown in our work on the Odd One Out (McGreggor et al. 2011), the re-representation phase is managed in close concert with the representation phase as a nested loop. In this way, the two phases of reasoning and re-representation occur as a unified execution phase. Upon the conclusion of the iterative reasoning, the most visually odd item may be indicated.

The CoreGeo algorithm also incorporates the similarity distribution technique used in our work on the Odd One Out problem domain as a means for calculating individual object similarity. The CoreGeo algorithm attempts to isolate the novel object as being the statistical outlier. The threshold value E corresponds to a deviation from the mean equivalent to some desired confidence level.

To determine from a group of geometrically related objects M which of the objects does not share the sought-for relationship. The specifics of the relationship are not known at the outset of the problem.

PREPARATORY

Let M := { O_1, O_2, ... O_n } represent a group of objects.

Let A := { a_1, a_2, ... a_l } represent an ordered range of abstraction, from most coarse (at a_1) to finest (at a_l).

Let G := { g_1, g_2, ... g_m } represent an ordered range of complexity groupings, from 2-combinations (at g_1) to (n-1)-combinations (at g_m). Thus |G| = |M|-2.

Let E be a real number which represents the number of standard deviations beyond which a value's answer may be judged as "confident." E then is the threshold value which groups objects into the TYPICAL or ODD sets according to their similarity value.

REASONING & RE-REPRESENTATION

For each complexity g □G:
  For each abstraction a □A:

- Form a set of relationships R from the objects in M according to g and a
- Derive the set of similarity values S := { S_1, S_2, S_3, S_4, ... S_n } from the set of relationships R, using the similarity distribution algorithm
- Set $\mu \leftarrow$ mean (S)
- Set $\sigma_\mu \leftarrow$ stdev (S)/$\sqrt{n}$
- Set D ← { D_1, D_2, D_3, D_4, ... D_n } where $D_i = (S_i-\mu)/\sigma_\mu$
- Set TYPICAL □ □
- Set ODD □ □
- Distribute the objects O_i □M by this rule:
  if D_i > E , then ODD □ ODD ∗ { O_i }
  else TYPICAL □ TYPICAL ∗ { O_i }
- If |ODD| = 1, return the object O_i □ODD
- otherwise there exists ambiguity, and re-representation must occur.

If no answer has been returned, then no answer may be given unambiguously.

Algorithm 1. The CoreGeo Algorithm.

2. 2 An example

Let us illustrate the CoreGeo algorithm by working through one of the Dehaene problems in detail. The chosen problem is Dehaene #35, which seeks to determine an understanding of symmetry transformation about a mixed axis. The problem can be seen in figure 2.

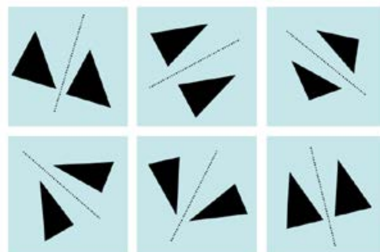

Figure 2. Dehaene #35, Transformation with mixed axial, symmetry.

2. 2 1 Segmentation phase

First, the algorithm must segment the problem into its constituent subimages, which we shall label O_1 through O_6. In this example, the problem is given as a 720 x 540 .PNG image in the RGB color space. The subimages are arrayed in a 3x2 grid within the problem image. At this resolution, we found that each subimage fits within a 210x210 pixel image. Note that the matrix arrangement of the subimages is immaterial to the problem at hand: the one which does not belong would be determined as not belonging without regard to its specific position.

Even though the objects or shapes in the subimages appear to be regular geometric shapes, the algorithm and the representation do not interpret them in any manner other than as mutual fractals. In addition, some noise and image artifacts are inevitably present, even though they may not be evident in the illustrations here.

2. 2 2 Representation Phase

Given these six subimages, the strategy now groups subimages as 2-combinations, pairs, such that each subimage is paired once with the other five subimages, to form 10 distinct 2-combinations. The strategy then calculates the mutual fractal representation R_{ij} for each pair of objects O_i and O_j as described above.

The level of abstraction used initially is identical to the largest possible pixel dimension, in this case 210x210 pixels. For this example, the algorithm shall determine the finer levels of abstraction by uniform subdivision of the subimages into block sizes of 105x105, 52x52, 26x26, 13x13, and 6x6 pixels.

2. 2 3 Reasoning Phase

To determine the subimage which does not possess the same geometric relationship as the others, the algorithm must determine the dissimilarity of each subimage. The CoreGeo algorithm calculates the dissimilarity via the distributed similarity technique. At the coarsest level of abstraction, these are the distributed similarity values for the subimages in the example problem.

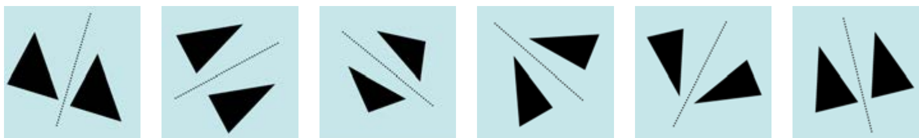

0.2273	0.2416	0.2303	0.2484	0.2419	0.2381

Table 1. Similarity distribution for the initial 210x210 level of abstraction.

2. 3 Concluding the example

Once the distribution of similarity is accomplished, the algorithm must examine the resulting values to see if any subimage has a substantially lower score than the others. The distinction between the lowest and the next-lowest score can be quite close (in this case, 0.2273 vs. 0.2303, a delta of only 0.003). If we calculate the mean of the similarity values and standard deviation of each of these values from that mean, a different picture emerges. In this example, the similarity mean is 0.2379 and the standard deviation is 0.0032. Therefore, we get the following table of deviations and subsequent confidences:

3.302	1.151	2.379	3.257	1.227	0.045
-99.9%	75.0%	-98.3%	99.9%	78.0%	3.6%

Table 2. Deviations and confidences for the initial abstraction level.

As can be seen, both the first and third answers are values well below the standard deviation and therefore at a strong confidence level. The negative values here are to indicate the least similar outliers; the strong positive confidence in the fourth answer, in contrast, indicates a strong prototype of the group. Thus, one is left with an ambiguous answer at this level of abstraction, and the algorithm must re-represent each of the 2-combinations at a finer level and try again.

2. 4 Refinement

By re-representing the 2-combinations at increasing levels of abstraction, it is possible to determine an unambiguous answer. Table 3 illustrates the results of that successive refinement of abstraction.

	μ	σ_μ						
210x210	0.2379	0.003	-99.9%	75.0%	-98.3%	99.99%	78.0%	3.6%
105x105	0.189	$5.5x10^{-5}$	41.1%	96.3%	-39.6%	-27.3%	98.9%	-99.9%
52x52	0.271	0.005	-95.2%	-74.3%	99.9%	-65.4%	-93.9%	92.6%
26x26	0.337	0.002	-98.9%	-99.7%	99.9%	42.6%	3.7%	82.6%
13x13	0.399	0.001	-99.9%	-91.9%	99.9%	90.1%	69.6%	-60.2%
6x6	0.494	0.001	-98.3%	-99.8%	99.9%	80.0%	53.4%	2.8%

Table 3. Mean, Standard Deviation, and Confidence for various levels of abstraction

As can be seen in the table, there are three levels of abstraction for which a singular answer stands out as significantly odd, while for the other levels there exists ambiguity. If the algorithm strictly followed the philosophy of proceeding from coarsest to finest abstraction until a value stands out, then the CoreGeo algorithm would select answer 6 at the abstraction denoted by 105x105 partitioning. This answer, unfortunately, is incorrect: the proper answer is answer 1. Inspection of the results shows that for all levels of abstraction except for the 105x105 level, answer 1 is among the chosen values which exceed a 95% confidence. Why would this not be true at the 105x105 level?

2. 5 The advent of significance

A closer inspection unveils the mystery. At the 105x105 level, the standard deviation from the mean for all values is remarkably low (in this case, $5.5x10^{-5}$). That deviation is two orders of magnitude smaller than all other abstraction levels. Thus, while the signal at that level of abstraction is unambiguously in favor of answer 6, the signal itself is too weak to merit consideration. In contrast, the unambiguous signal for answer 1 at the 52x52 level of abstraction is three orders of magnitude stronger.

The similarity calculations arise from the comparison of features present in the fractal representations of the relationships being examined. As noted earlier, in the chapter on visual similarity, the number of features available gives rise to the presence or absence of ambiguity, either through a sparsity of features or an increase in the homogeneity of features. In this example, we find evidence of both of these, yet the data itself has yielded now a clue for detecting homogeneity: the coefficient of variation (CV), a normalized measure of the dispersion of the values.

We now can offer an amendment to the CoreGeo algorithm. The selection of a threshold confidence value itself is not enough: the signal must be unambiguous and strong before the algorithm declares a solution. With this addition, we capture a powerful notion of analogy making: the analogy must be significant enough to warrant notice.

	CV σ / μ						
210x210	0.0126	-99.9%	75.0%	-98.3%	99.99%	78.0%	3.6%
105x105	0.0003	41.1%	96.3%	-39.6%	-27.3%	98.9%	-99.9%
52x52	0.0185	-95.2%	-74.3%	99.9%	-65.4%	-93.9%	92.6%
26x26	0.0059	-98.9%	-99.7%	99.9%	42.6%	3.7%	82.6%
13x13	0.0025	-99.9%	-91.9%	99.9%	90.1%	69.6%	-60.2%
6x6	0.0020	-98.3%	-99.8%	99.9%	80.0%	53.4%	2.8%

Table 4. CV and Confidence for various levels of abstraction

3. Results of CoreGeo on the Dehaene set

The CoreGeo algorithm was run against the 45 problems provided to me by Dehaene via personal correspondence, and are the same problem set examined in Lovett et al. (2008). The problems are grouped into categories of mathematical or geometric reasoning.

3. 1 Preparation of the material

The Dehaene problem set was given as individual slides contained within a PowerPoint document. Each slide was exported into a single image in the .PNG format. Each problem image was 720 x 540 pixels, in the RGB color space. Each problem consists of six subimages, each of which upon inspection was found to fit well within a 210 x 210 boundary.

3. 2 Levels of abstraction considered and calculations performed

The levels of abstraction used ranged from a coarse partition of 210x210 pixels, down to a fine partitioning of 6x6 pixels, giving 6 levels of abstraction, using a strategy of halving the pixel dimension at each successively finer level of abstraction. The CoreGeo algorithm is capable of examining all combinations of the six subimages; however, only 2-combination relationships were examined in this experiment. This restriction was made only to serve the interests of experimental time, and illustrate the use of the algorithm on the problem set. It is important to note that the goal of the experiment was not to improve upon any prior computational model results.

At each level of abstraction for each problem, the algorithm determined the similarity value as distributed amongst the six candidate images. For these calculations, the algorithm used the Tversky formula (Tversky 1977) and set alpha and beta to 1.0, thus conforming the model of Gregson and Sjöberg (Gregson, 1976; Sjöberg, 1971), as it is unclear which of the difference relationships to favor. As prescribed by the CoreGeo algorithm, calculations continued until the confidence in an answer exceeded a given threshold of 95%, or until all levels of abstraction were calculated. For this experiment, that threshold value was able to be varied. After the initial pass of calculations, any problems found to be either incorrect or ambiguous were reprocessed, using a fine stepping down of partitioning, from 210x210 pixels down to 10x10 pixels, in increments of 10 each time, for a total of 21 levels of abstraction. The results reported below reflect the outcome of these passes.

The Java code and entire Dehaene set of problems are available via inquiry at our lab's research site, to facilitate replication of these results and future studies.

3. 3 Performance

The overall results are that the CoreGeo algorithm detected the correct answer at a 95% or higher level of confidence on all of the 45 problems. Of these problems where the correct answer was detected, 15 were ambiguously so.

3. 4 Discussion of the specific results

As can be plainly seen in the table above, there are certain problems and categories for which the CoreGeo algorithm successfully identifies a correct answer unambiguously, and others for which the algorithm is consistently ambiguous.

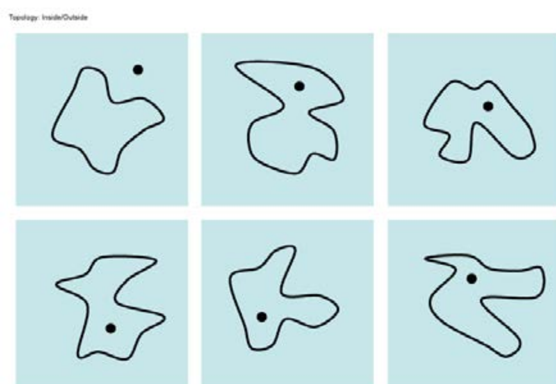

Figure 2. Dehaene problem #4, topological inside/outside

3. 4 1 Topological reasoning

The CoreGeo algorithm performs well on the problems involving topological reasoning except for one type: inside vs outside. Figure 2 illustrates the Dehaene problem #4.

As in all the examples, the features over which the CoreGeo algorithm reasons are derived from the fractal representation. Between the features themselves, there is no connection, and therefore no ability to directly associate the location of the dot in the figures above as either within or without the closed line.

3. 4 2 Geometrical reasoning

The CoreGeo algorithm performs well on problems involving geometric shapes and geometric reasoning, with two notable exceptions: reasoning about trapezoids, and reasoning about parallel lines. We note that both of these problems involve the notion of parallelism.

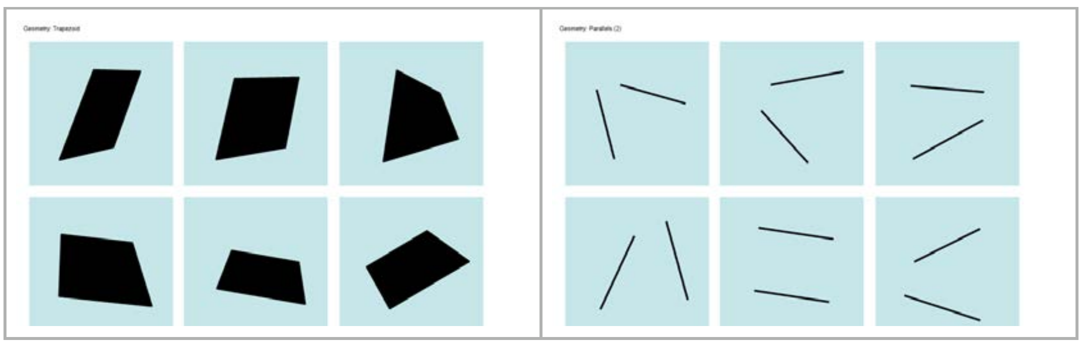

Figure 3. Dehaene problems #26 and #40

In the case of Dehaene problem #26, while the shape lacking non-parallel edges is apparent, at the fractal feature level, the comparisons would be with respect to finding similarity between the angles themselves. Each of the other shapes contains at least one oblique angle and two acute angles, and so the ambiguity for this problem would seem to indicate that the numerosity of the angle kinds is not readily inferred from the fractal representation. In Dehaene problem #40, as in Dehaene problem #4 above, the satisfactory answer would imply that comparisons be made between the whole line shapes, rather than their constituent parts (that is, the fractal representation of the images would note that line segments may be formed via the collage of other line segments). In both problems, the oddity lies in the presence or absence of parallel edges.

problem /category		Core Geo	Lovett et al.	American	Mundurukú
	Total noted correct	45	39	43	40
	Correct but ambiguous	15	0	12	10
1	Training Color	yes	yes	yes	yes
2	Training Orientation	yes	yes	yes	yes
3	Topology Holes	yes	yes	yes	yes
4	Topology Inside/Outside	ambiguous	yes	yes	yes
5	Topology Closure	ambiguous	yes	yes	yes
6	Topology Connexity	yes	yes	yes	yes
7	Topology Belongs To	ambiguous	yes	yes	yes
8	Geometry curved lines	yes	yes	yes	yes
9	Geometry Convexity	yes	yes	yes	yes
10	Geometry straight lines	yes	yes	yes	yes
11	Geometry aligned	ambiguous	yes	yes	yes
12	Geometry quadrilateral	ambiguous	yes	yes	yes
13	Geometry right angle triangle	ambiguous	yes	ambiguous	yes
14	Geometry right angle cross	yes	yes	yes	yes
15	Geometry right angle abut	yes	yes	yes	yes
16	Geometry distance	yes	yes	yes	yes
17	Geometry circles	yes	yes	yes	yes
18	Geometry center of circle	yes	yes	yes	yes
19	Geometry midpoint	ambiguous	yes	ambiguous	ambiguous
20	Geometry Equilateral triangles	yes	yes	yes	yes
21	Geometry Proportion 1:3	ambiguous		ambiguous	yes
22	Geometry Diagonals	yes			ambiguous
23	Geometry Square	yes	yes	yes	yes
24	Geometry Rectangle	ambiguous	yes	yes	yes
25	Geometry Parallelogram	ambiguous	yes	ambiguous	yes
26	Geometry Trapezoid	ambiguous	yes	yes	yes
27	Transformation vertical axial	yes	yes	ambiguous	
28	Geometry vertical axial symmetry	yes	yes	ambiguous	yes
29	Geometry horizontal axial symmetry	yes	yes	ambiguous	ambiguous
30	Geometry random axial symmetry	yes	yes	yes	yes
31	Transformation translation	yes	yes	yes	ambiguous
32	Transformation point symmetry	ambiguous	yes	yes	ambiguous
33	Transformation horizontal axial	ambiguous	yes	ambiguous	ambiguous
34	Transformation rotation	yes		ambiguous	
35	Transformation mixed axial symmetry	yes	yes	ambiguous	ambiguous
36	Transformation homothety	yes	yes	yes	ambiguous
37	Geometry Parallels	yes	yes	yes	yes
38	Geometry Chirality 1	yes		ambiguous	
39	Geometry Proportions	ambiguous			
40	Geometry Parallels 2	ambiguous	yes	yes	yes
41	Geometry Chirality 2	yes	yes	yes	yes
42	Geometry Chirality 3	yes	yes	yes	yes
43	Series Arithmetic	yes	yes	yes	ambiguous
44	Geometry Chirality 4	yes		ambiguous	
45	Series Geometric	yes	yes	yes	ambiguous

Table 6. Comparison of the CoreGeo Algorithm to other studies

3. 5 Comparison against another computational model

In Lovett et al. (2008) and Lovett & Forbus (2011), they describe a computational model for Dehaene's test viewed as a visual oddity task. Their model employs four tools: CogSketch to construct the qualitative representations (Forbus et al. 2008); SME to model comparison and similarity (Falkenhainer et al. 1986); MAGI to model symmetry detection (Ferguson 1994); and SEQL to model generalization (Kuehne et al. 2000). In their approach, a chosen Dehaene problem is first segmented into the individual images, and those images are then represented via CogSketch. Next, SEQL is used to create a generalization of those representations. The individual images are compared against the generalization and scored via SME. If one image is noticeably less similar to the generalization, it is deemed the one that doesn't belong. Lovett et al. note that the actual processing consists of a series of these "generalize and compare" trials, selecting subsets of the individual images from which to create generalizations. They appear to limit these subsets to be either the top three or bottom three images. Further, since their model uses representations of either the shape or the edges of the shape but not both, the system decides which version to use based upon an examination of the first image in the problem: if the image contains multiple shapes or a shape with a single edge (e.g. a circle), then the shape qualitative representation is used; otherwise, the edge qualitative representation is used. They do note that their system will abandon the edge representation if SEQL cannot generate a sufficient generalization (for example, if the images subject to generalization contain varying numbers of edges). The system looks for a sufficiently distinct candidate (they suggest a confidence of 95% as sufficient). If one is not found, then the system attempts additional trials, switching between shape and edge representation or varying the similarity scoring.

Lovett et al. (2008) and Lovett and Forbus (2011) report that their system correctly solves 39 of the 45 problems. On the six problems that the Lovett et al. model misses, the CoreGeo algorithm answers correctly on 4 of them and ambiguously correct on the remaining two. A comparison with respect to ambiguity between the models is not possible, as Lovett et al. do not report the confidence of their answers or provide a means for calculating the confidence from the individual scores per answer choice.

3. 6 Comparison against human performance

Table 6, above, contains a direct comparison of answers on the test between our CoreGeo algorithm, the Lovett et al. model, and human results, at a confidence level of 95%. The human performance data was derived by the authors from data given in a chart in Lovett et al. (2008).

As we were interested in comparing not just the answers between models on a correct versus incorrect stance, but rather on a confidence basis, we imposed a system of interpretation upon the human data. To that end, we interpreted the human data as correct if the accuracy value (as given in the chart in Lovett et al. 2008) fell above 0.6 (about 1 standard deviation of confidence), ambiguous if between 0.6 and 0.2 (approximately random), and incorrect if below 0.2. In this manner, we established that the American data provided 43 total correct answers with 12 ambiguously correct, while the Mundurukú data provided 40 total correct answers with 10 ambiguously correct.

As can be seen in Table 6, in the cases where the human results (either American or Mundurukú) are interpreted as ambiguous, our CoreGeo algorithm also generally but not consistently responds ambiguously (e.g. problems #13, #19, #21, #25, #32, and #33 are mutually ambiguous, but the rest are not).

As noted in the specific discussion of results above, the CoreGeo algorithm is less confident than human subjects on aspects of topology. However, it is more confident than human subjects on those tests involving symmetry, transformations, and chirality. We believe that this is attributable to the specific nature of the fractal representation, in that the features derived from those representations and upon which analogical assessments made are essentially spatial transformations. We direct the reader to the first appendix for a brief description of the fractal representation.

Conclusions

The origin of cognitive abilities, such as geometric abilities, is a complex and open issue. Dehaene et al.'s test suggests humans may have a core set of geometric abilities irrespective of culture, language or education. In this paper, we exploit the analogical nature of the problems on Dehaene et al.'s test to develop a new computational model of geometric analogies which, like all theories of analogy more generally, strongly relies on the notion of similarity, and, in particular, on visual similarity.

Our computational model of geometric analogies uses fractal representations. This has the advantage that the reasoner can work at multiple levels of spatial resolution, zooming in and out as needed, because the geometric figures considered at those resolution levels are similar to one another. Thus, if at a given level of resolution, an answer is ambiguous because the confidence in the answer is about the same as the confidence in another answer, then our reasoner can shift strategies, changing the level of abstraction to find a spatial resolution level at which an answer may emerge unambiguously.

We view perception as playing an important role in cognition. Further, we view the fractal representation as a knowledge representation. In earlier work, we have developed a computational theory of fractal analogical reasoning (FAR). We have also used FAR for a variety of tasks such as geometric analogies on intelligence tests, visual oddity tests, as well as perception in simulated reactive agents.

In this paper, we showed how the CoreGeo algorithm, based on FAR's ABR* algorithm, can address problems on the Dehaene et al.'s test. Both humans - American as well as Mundurukú - and CoreGeo compute many answers unambiguously and others with some ambiguity. Our algorithm is parsimonious in that it doesn't need additional mechanisms, representations, or components.

More importantly to us, our CoreGeo computational model performs at about the level of humans, both educated Americans and the Mundurukú tribe. Of course, it is much too early to suggest that computational model is also a cognitive model of core geometric abilities. At this stage of development, our work simply shows that CoreGeo can address the Dehaene et al.'s test at about the level of human performance, that it enables calculation of confidence values in answer, and that it allows automatic adjustment of the level of resolution.

The characteristics of the CoreGeo model results in some predictions about human performance. As Table 6 indicates, humans as well as CoreGeo have difficulty finding answers to some problems on the Dehaene's test because of the ambiguity of the answers to the problems. Based on our work with CoreGeo, we predict that humans have difficulty choosing answers on specific problems on the Dehaene's test (the ones on which CoreGeo finds ambiguous answers), and that when they do have such difficulties, they change the level of resolution/abstraction at which they analyze the given problem. We hope to design experiments that can test these predictions.

Current theories of analogy approach regard the process of analogy-making as if it were a singular process based upon three core ideas: that analogy depends upon the capture of not just features but the relationships between objects or concepts (Holyoak & Thagard, 1996); that propositional representations (and not imagistic representations) are crucial for the concise expression of those transformations and relationships; and that the fit of an analogy depends upon the alignment between the structure of representations, not necessarily the content of knowledge represented (Gentner, 1983). In contrast, our research group has a long history of building an alternative account of analogy (Goel, 1997). In our view, analogy is composed of multiple processes, some of which are based on just features (e.g. Kunda et al., 2013) and others capturing relationships (e.g. Yaner & Goel, 2008), some more focused on propositional representations (Davies et al., 2009) and others on imagistic representations (e.g. Kunda et al., 2013), and some intent on exploiting the organization of knowledge into abstraction hierarchies (e.g. Goel & Bhatta, 2004; Davies et al., 2009) while others examine the content of knowledge at specific abstraction levels (e.g. McGreggor et al. 2012). This paper shows that the fractal representation as an imagistic representation and the ABR* algorithm as a featural analogical process offer an example of a content-based theory of analogy.

Acknowledgements

This work has benefited from many discussions with Maithilee Kunda and our colleagues in the Design and Intelligence Laboratory of the School of Interactive Computing at the Georgia Institute of Technology. We are grateful to the US National Science Foundation for its support of our research through NSF (RI) Grant #1116541, entitled "Addressing visual analogy problems on the Raven's intelligence test."

Appendix 1 Fractal Representation

We have previously published the details of the fractal representation (e.g. McGreggor et al. 2011; McGreggor & Goel, 2012). We include here a brief review of the representation, however.

The mathematical derivation of fractal image representation expressly depends upon the notion of real world images, i.e. images that are two dimensional and continuous (Barnsley & Hurd, 1992). A key observation is that all naturally occurring images appear to have similar, repeating patterns. Another observation is that no matter how closely one examines the real world, one may find instances of similar structures and repeating patterns. These observations suggest describing images in terms that capture the observed similarity and repetition alone, without regard to shape or traditional graphical elements.

Computationally, determining fractal representation of an image requires the use of the fractal encoding algorithm. For a target image D and a source image S, the algorithm seeks to discover a set of transformations T.

First, systematically partition D into a set of smaller images, such that D = {d_1, d_2, … }.

For each image d_i:

- Examine the entire source image S for an equivalent image fragment s_i such that an affine transformation of s_i will likely result in d_i.
- Collect all such transforms into a set of candidates C.
- Select from the set C that transform which most minimally achieves its work, according to some predetermined metric.
- Let T_i be the representation of the chosen transformation associated with d_i.

The set T = {T_1, T_2, T_3, … } is the fractal encoding of the image D.

Algorithm 2. Fractal Encoding of D from S

The fractal encoding of the transformation from a source image S into a destination image D is tightly coupled with the partitioning scheme P of the destination image D, which may be achieved through a variety of methods. In our present implementation, we merely choose to subdivide D in a regular, gridded fashion. Note that the cardinality of the resulting set T is determined solely by the partitioning P. In this fashion, the fractal representation is derived.

The analogical relationship between source and destination images may be seen as mutual; that is, the source is to the destination as the destination is to the source. However, the fractal representation is decidedly one-way (e.g. from the source to the destination). To capture the bidirectional, mutual nature of the analogy between source and destination, we introduced the notion of a mutual fractal representation. Let us label the representation of the fractal transformation from image A to image B as T_{AB}. Correspondingly, we would label the inverse representation as T_{BA}. We shall define the mutual analogical relationship between A and B by the symbol M_{AB}, given by this equation: $M_{AB} = T_{AB} * T_{BA}$

Appendix 2 Similarity Distribution

In McGreggor and Goel (2012), we developed the notion of the distribution of similarity. In that work, we determined a score for a grouping of images G, but needed a way to assign individual scores for the participating images.

To determine the image scoring, we distribute the similarity score of the grouping equally among the participating images. For each of the images, a score is generated which is proportional to its participation in the similarity of the grouping's similarity vectors. If an image is one of the two images in a pairing, as an example, then the image's similarity score receives one half of the pairing's calculated similarity score.

Given a set of M images $I = \{ I_1, \ldots I_m \}$ and a set of N representations $R = \{ R_1, \ldots R_n \}$, where each R_i is the mutual fractal representation between two (or more) of those images.

Let Q be a vector of cardinality M, initialized to 0: $Q \leftarrow 0, |Q| = M$

For each representation R_i □R:
- **Let S be an vector of cardinality N, initialized to 0: $S \leftarrow 0, |S| = N$**
- **For each representation R_k □R:**
    - If $i = k$, then $S_k = 1 \because R_i$ is identical to itself
    - If $i \neq k$, then calculate S_k using the Tversky formula:
      $S_k \leftarrow f(R_i \cap R_k) / [f(R_i \cap R_k) + \alpha f(R_i\text{-}R_k) + \beta f(R_k\text{-}R_i)]$
- **Let V be a scalar value, set to the normalized magnitude of S: $V \leftarrow \|S\| / |S|$**
- **For each subimage I_j which is represented by R_i:**
    - $Q_j \leftarrow Q_j + V$

The vector Q may then be said to contain the distributed similarity of each image to one another.

Algorithm 3. Similarity distribution

References

Barnsley, Michael, & Hurd, Lyman. (1992). *Fractal Image Compression.* Boston, MA: A.K. Peters.

Dehaene, Stanislas, & Cohen, Laurent. (1994). Dissociable mechanisms of subitizing and counting: neuropsychological evidence from simultanagnosic patients. *Journal of Experimental Psychology: Human Perception and Performance*, 20(5), 958.

Dehaene, Stanislas, & Cohen, Laurent. (1996). Towards an anatomical and functional model of number processing. *Mathematical cognition*, 1, 83-120.

Dehaene, Stanislas, Izard, Véronique, Pica, Pierre, & Spelke, Elizabeth. (2006). Core knowledge of geometry in an Amazonian indigene group. *Science*, 311(5759), 381-384.

Davies, Jim, Goel, Ashok K, & Nersessian, Nancy J. (2009). A computational model of visual analogies in design. Cognitive Systems Research, 10(3), 204-215.

Falkenhainer, Brian, Forbus, Kenneth D, & Gentner, Dedre. (1986). *The structure-mapping engine*: Department of Computer Science, University of Illinois at Urbana-Champaign.

Ferguson, Ronald W. (1994). *MAGI: Analogy-based encoding using regularity and symmetry.* Paper presented at the Proceedings of the sixteenth annual conference of the Cognitive Science Society.

Franco, Laura, & Sperry, RW. (1977). Hemisphere lateralization for cognitive processing of geometry. *Neuropsychologia*, 15(1), 107-114.

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

Goel, Ashok K. (1997). Design, analogy, and creativity. IEEE Intelligent Systems, Special Issue on Qualitative Reasoning, 12(3), 62-70.

Goel, Ashok K., & Bhatta, Sambasiva R. (2004). Use of design patterns in analogy-based design. Advanced Engineering Informatics, 18(2), 85-94.

Gregson, RAM. (1976). A comparative evaluation of seven similarity models. *British Journal of Mathematical and Statistical Psychology*, 29(2), 139-156.

Holyoak, Keith J, & Thagard, Paul. (1996). Mental leaps: Analogy in creative thought: MIT press.

Kuehne, Sven, Forbus, Kenneth, Gentner, Dedre, & Quinn, Bryan. (2000). *SEQL: Category learning as progressive abstraction using structure mapping.* Paper presented at the Proceedings of the 22nd Annual Meeting of the Cognitive Science Society.

Kunda, Maithilee, McGreggor, Keith, & Goel, Ashok K. (2013). A computational model for solving problems from the Raven's Progressive Matrices intelligence test using iconic visual representations. Cognitive Systems Research, 22-23, 47-66.

Lakoff, George, & Núñez, Rafael E. (2000). *Where mathematics comes from: How the embodied mind brings mathematics into being*: Basic books.

Lovett, Andrew, Lockwood, Kate, & Forbus, Kenneth. (2008). Modeling cross-cultural performance on the visual oddity task. *Spatial Cognition VI. Learning, Reasoning, and Talking about Space* (pp. 378-393): Springer.

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

Mandler, George, & Shebo, Billie J. (1982). Subitizing: an analysis of its component processes. *Journal of Experimental Psychology: General*, 111(1), 1.

McGreggor, Keith, & Goel, Ashok. (2011). *Finding the odd one out: a fractal analogical approach.* Paper presented at the Proceedings of the 8th ACM conference on Creativity and cognition.

McGreggor, Keith, & Goel, Ashok. (2012). Fractal analogies for general intelligence. *Artificial General Intelligence* (pp. 177-188): Springer.

McGreggor, Keith, Kunda, Maithilee, & Goel, Ashok. (2011). *Fractal Analogies: Preliminary Results from the Raven's Test of Intelligence.* Paper presented at the Proc. International Conference on Computational Creativity, Mexico City, Mexico, April.

Sinclair, N, & Schiralli, M. (2003). A constructive response toWhere mathematics comes from'. *Educational Studies in Mathematics*, 52(1), 79-91.

Sjöberg, Lennart. (1971). Three models for the analysis of subjective ratios. *Scandinavian Journal of Psychology*, 12(1), 217-240.

Trick, Lana M, & Pylyshyn, Zenon W. (1993). What enumeration studies can show us about spatial attention: evidence for limited capacity preattentive processing. *Journal of Experimental Psychology: Human Perception and Performance*, 19(2), 331.

Trick, Lana M, & Pylyshyn, Zenon W. (1994). Why are small and large numbers enumerated differently? A limited-capacity preattentive stage in vision. *Psychological review*, 101(1), 80.

Tversky, Amos. (1977). Features of similarity. *Psychological review*, 84(4), 327-352.

Yaner, P. W., & Goel, A. K. (2008). Analogical recognition of shape and structure in design drawings. International Journal of AI in Engineering Design, Analysis and Manufacturing, Special Issue on Multimodal Design, 22(2), 117-128.