

# A language of thought for the mental representation of geometric shapes

Mathias Sablé-Meyer<sup>\*1,2</sup>

Kevin Ellis<sup>3</sup>

Josh Tenenbaum<sup>4</sup>

Stanislas Dehaene<sup>1,2</sup>

## Abstract

Why do geometric shapes such as lines, circles, zig-zags or spirals appear in all human cultures, but are never produced by other animals? Here, we formalize and test the hypothesis that all humans possess a compositional language of thought that can produce line drawings as recursive combinations of a minimal set of geometric primitives. We present a programming language, similar to Logo, that combines discrete numbers and continuous integration in higher-level structures based on repetition, concatenation and embedding, and show that the simplest programs in this language generate the fundamental geometric shapes observed in human cultures. On the perceptual side, we propose that shape perception in humans involves searching for the shortest program that correctly draws the image (program induction). A consequence of this framework is that the mental difficulty of remembering a shape should depend on its minimum description length (MDL) in the proposed language. In two experiments, we show that encoding and processing of geometric shapes is well predicted by MDL. Furthermore, our hypotheses predict additive laws for the psychological complexity of repeated, concatenated or embedded shapes, which are experimentally validated.

## Keywords

Geometry; program induction; shape perception; Language of Thought; Complexity in Cognition

## Highlights

- Some geometrical shapes are universally attested in humans, but not other animals
- We formalize the hypothesis of a compositional language of thought for shapes
- We argue for modeling human shape perception as program induction
- We provide empirical evidence of our proposal and outline the space of alternatives

## Introduction

We could never know the geometric triangle through the one we see traced on paper, if our mind had not had the idea of it elsewhere.

## René Descartes

The cognitive origins of geometric knowledge remains heavily debated. While several animal species possess sophisticated neural circuits for spatial navigation (including head direction, place, grid and border cells), only humans seem capable of conceiving formal, symbolic geometric structures. The formalization of geometry is traditionally dated to Euclid's *Elements*, itself rooted in Egyptian and Babylonian precursors. Yet various lines of evidence suggest that an intuitive sense of geometry is much more ancient, and that many, possibly all human cultures share a drive towards creating geometric designs (Van der Waerden, 2012). Throughout the world, at geographically distant and presumably unrelated sites, humans have produced parallel lines, circles, squares, zig-zags or spirals, in activities as diverse as drawing, pottery, body paintings, rock art, land art (e.g., Nazca lines), stone-cutting (e.g., bifaces), or large-scale constructions (e.g., Stonehenge). Many Neolithic sites contain square, circular or rectangular buildings as well as large circles of stones (cromlechs) whose axes are often systematically oriented relative to geographical or astronomical landmarks (Pimenta & Tirapicos, 2015). Before the advent of artificial flight, the shape of these large structures could not be directly apprehended: from the ground, they would be perceived as a distorted quadrilateral or ellipse, at best. The fact that squares and circles appeared at many different scales suggests

---

Mathias Sablé-Meyer  <https://orcid.org/0000-0003-0844-0775>

Stanislas Dehaene  <https://orcid.org/0000-0002-7418-8275>

\*Corresponding author [mathias.sable-meyer@ens-cachan.fr](mailto:mathias.sable-meyer@ens-cachan.fr)

<sup>1</sup> Unicog, CEA, INSERM, Université Paris-Saclay, NeuroSpin center, 91191 Gif/Yvette, France

<sup>2</sup> Collège de France, Université Paris-Sciences-Lettres (PSL), 75005 Paris, France

<sup>3</sup> Cornell University, United States

<sup>4</sup> Massachusetts Institute of Technology, United States

that their human designers possessed an abstract mental concept of geometric shape that guided their architectural, artistic or practical creations.

### Human and animal sensitivity to geometric patterns: a brief review

Evidence for abstract concepts of geometry, including rectilinearity, parallelism, perpendicularity and symmetries, is widespread throughout prehistory. About 70,000 years ago, *homo Sapiens* at Blombos cave carved a piece of ochre with three interlocking sets of parallel lines forming equilateral triangles, diamonds and hexagons (Henshilwood et al., 2002). Much earlier, approximately 540,000 years ago, *homo Erectus* in Java carved a zig-zag pattern on a shell (Joordens et al., 2015). Such a zig-zag may look simple, but it approximately respects geometric constraints of equal lengths, equal angles and parallelism, and is undoubtedly attributed to the *homo* genus. Even earlier, since ~1.8 million years, ancient humans have been carving spheroids (sphere-like stones) and bifaces — stones possessing two orthogonal planes of symmetry (Le Tensorer, 2006). The vast number of bifaces, their near-perfect symmetry (which is not required for them to operate as efficient tools), and the archeological evidence that many were never used as tools, suggest that an aesthetic drive for symmetry was already present in ancient humans.

Contemporary cognitive anthropology corroborates those findings. Cognitive tests performed in relatively isolated human groups such as the Mundurucu from the Amazon, the Himba from Namibia, or aborigine groups from Northern Australia, show that in the absence of formal education in mathematics, adults and even children already possess strong intuitions of numerical and geometric concepts (Amalric et al., 2017; Butterworth et al., 2008; Dehaene et al., 2006; Izard et al., 2011b; Pica et al., 2004; Sablé-Meyer et al., 2021). (Amalric et al., 2017; Butterworth et al., 2008; Dehaene et al., 2006; Izard et al., 2011b; Pica et al., 2004; Sablé-Meyer et al., 2021) Indeed, those uneducated adults and children share a large repertoire of abstract geometric concepts (Dehaene et al., 2006) and use them to capture the regularities in spatial sequences (Amalric et al., 2017) and quadrilateral shapes such as squares or parallelograms (Sablé-Meyer et al., 2021). They even possess sophisticated intuitions of how parallel lines behave under planar and spherical geometry, such as the unicity of a parallel line passing through a given point on the plane (Izard et al., 2011a).


Another piece of evidence arises from developmental data. Preschoolers and even infants have been shown to possess sophisticated intuitions of space (Hermer &

Spelke, 1994; Landau et al., 1981; Newcombe et al., 2005), spatial sequences (Amalric et al., 2017), and mirror symmetry (Bornstein et al., 1978). Indeed, even preschoolers' drawings already show a tendency to represent abstract properties of objects rather than the object itself. Although they look primitive, drawings of a house as a triangle on top of a square, or a person as a stick figure with a round head, suggest a remarkable capacity for abstracting away from the actual shape and attending to its principal axes, at the expense of realism. Numerous tests leverage this geometric competence to measure a kid's development by counting the number of correct or incorrect abstract properties, for instance when asked to draw a person (Goodenough, 1926; Harris, 1963; Long et al., 2019; Prewett et al., 1988; Reynolds & Hickman, 2004). There is some evidence, however limited, that this ability may be specifically human: when given pencils or a tablet computer, other non-human primates do not draw any abstract shapes or recognizable figures, but mostly generate shapeless scribbles (Saito et al., 2014; Tanaka et al., 2003).

We recently compared the perception of quadrilateral geometric shapes in humans and in baboons, using the very same task (Sablé-Meyer et al., 2021). We used the intruder test (Dehaene et al., 2006), which involves viewing an array of pictures and clicking on the one that looks distinctly different from the others, and is well within the grasp of human adults, children and baboons. All humans, regardless of age, culture and education, exhibited a striking effect of shape regularity: squares and rectangles were processed better than other, more irregular quadrilaterals, and there was a continuous ordering of complexity, from squares and rectangles to parallelograms, trapezoids, and fully irregular shapes. Strikingly, this geometric regularity effect was absent in baboons. Baboon behavior was quite consistent across individuals and could be captured by neural network models of the ventral visual pathway for object recognition. Modeling the human perception of quadrilaterals, however, required an additional assumption, namely the existence of discrete symbolic concepts of parallelism, right angle, equal length, or equal angle. We therefore argued that two strategies can be used to perform the outlier task: a visual one, available to all primates, and an abstract, symbolic one that may be unique to humans (Sablé-Meyer et al., 2021).

### Summary of our approach and hypotheses

In the present paper, we formalize and put to an empirical test the hypothesis that geometry is one of the manifestations of the specifically human ability to represent and manipulate recursively embedded languages



**Figure 1**

**Geometric shapes in human cultural history.** *A*, examples of small- and large-scale geometric drawings and constructions (From left to right and top to bottom: an engraved slab from Blombos caves dating about 70.000 years ago (Henshilwood et al., 2002); zigzag pattern engraved on a shell in Java approximately 540.000 years ago (Joordens et al., 2015); Boscawen-Ūn’s Bronze Age elliptical cromlech in Cornwall; spiral stone engraving on Signal Hill in Saguaro National Park, Arizona, dated 550 to 1550 years ago; geometrical shapes below the painting of a Megaloceros in Lascaux, France, typically dated to be 17,000 years old). *B*, basic shapes found in many human cultures and which exemplify the main components of the proposed language for geometric shapes: line (continuous integration), circle and spiral (integration with continuous variation in parameters such as direction or speed), zigzag (concatenation), square (discrete repetition), and square of circles (embedding of a subprogram).

(Dehaene et al., 2015; Fitch, 2014; Fodor, 1975; Frankland & Greene, 2020; Hauser et al., 2002; Piantadosi, 2011).

Fodor (1975) famously introduced the language of thought hypothesis, according to which an inner combinatorial language underlies high-level cognition in humans and allows the creation of a vast space of mental representations by recursive recombination of preexisting ones. Hauser, Chomsky and Fitch (Hauser et al., 2002) hypothesized that recursion might be the single uniquely human ingredient that explains the emergence

of the human language faculty. Fitch (2014) and Dehaene et al. (2015) later argued that recursion is not limited to linguistic communication, and that various “languages of thought”, all based on a basic capacity for recursive syntax and compositional semantics, could underlie many other uniquely human abilities such as music, mathematics or theory of mind. Here, we apply this idea to the domain of geometric shape perception.

Our proposal builds upon the seminal work of Leeuwenberg and colleagues (Boselie & Leeuwenberg, 1986; Leeuwenberg, 1971), who proposed a formal coding language for 2- and 3-dimensional shapes, and showed that it could account for data on human shape perception. Furthermore, Leeuwenberg’s formal code output sequences of numbers which, when interpreted as patterns, could also predict auditory pattern perception in humans. Later, Leyton (1984, 2003) argued that the shapes that humans generate arise from a set of primitives (points, lines, planes) together with the repeated mental application of a series of group transformations that duplicate, stretch, rotate, or skew them. While these elegant proposals have had a considerable influence in the design of graphics software, it is fair to say that they remained partially disconnected from the experimental psychophysical or neurophysiological literature on shape perception (for exceptions, see Brincat & Connor, 2004, 2006; Hung et al., 2012).

In our previous work, we introduced a much more restricted, yet more precise, language of thought for geometric sequences. Our work focused on capturing the psychological complexity of all the sequences of 8 locations that can be generated by drawing without repetition from the vertices of an octagon (Amalric et al., 2017). The basic building blocks of our proposed language were the arithmetic primitive of discrete number, the geometric primitives of symmetry and rotation, and a single recursive operation of repetition (possibly with variations). These operations could be embedded, thus allowing for repetitions of repetitions in a nested manner. For instance, the repeated application of a symmetry operation, each time with an increment in the starting point, could generate a zig-zag pattern. A square could be generated by a 4-fold repetition of moving by 2 vertices around the octagon. As a more complex example, a sequence of two squares could be generated by two nested “for loops”, i.e. a 2-fold repetition (while changing the starting point) of the 4-fold repetition that draws a square. These examples illustrate the generative power of the concept of **repetition with variation**, which is indeed formally indistinguishable from symmetry, a key concept on which mathematics is founded.

Amalric et al. (2017) measured empirically the difficulty

that preschoolers and adults (including Mundurucu adults) had in predicting or memorizing spatial sequences of locations on an octagon. Across 11 geometric sequences, psychological complexity was determined by the complexity of their internal representation in the proposed language. Working memory was not determined by sequence length (which, indeed, was fixed at 8 items), but by the capacity to compress the sequence into a compact internal representation using the proposed language. The central concept here, as already proposed by many others (Chater & Vitányi, 2003; Feldman, 2000, 2003; Li & Vitányi, 1997; Mathy & Feldman, 2012; Romano et al., 2013), is that psychological complexity in humans depends on **minimum description length (MDL)**, i.e. the length of the shortest mental representation which can encode the sequence, rather than the literal length of the sequence. In the case of our language for geometric sequences, MDL was also shown to tightly predict brain activity in both functional magnetic resonance imaging (fMRI; Wang et al., 2019) and magneto-encephalography (Al Roumi et al., 2021). The very same language was also successfully extended to account for the perception of simple auditory sequences made of two discrete sounds (Planton et al., 2021).

In the present work, we move beyond discrete sequences made of points and straight lines, and tackle the mental representation of static geometric shapes such as a square, a circle or a spiral. As noted above, the square is easily captured by a language with discrete integers and repetition (“for loops”). However, continuously varying shapes such as circles and spirals raise interesting issues that arguably require more than integers. In computer languages such as Logo, such drawings are implemented using a discrete repetition instruction with a very small increment, thus drawing a quasi-continuous curve which is in fact made of straight lines. However, we find implausible the idea that humans intuitively think of such an infinitesimal and inherently discrete representation when thinking of a circle. Furthermore, computationally, the unbounded nature of such infinitesimal loops would allow short programs to generate visually complex shapes. Instead, we argue that the crucial notion of “repetition with variation” introduced in our previous work can be helpful again, but now in a continuous version. We propose that, whenever a mental primitive is available, for instance for drawing a straight line, mental control structures in humans are available to either keep its parameters constant, or to continuously vary them over time. Thus, the new version of our proposed languages includes, not only discrete repetition (“for loops”), but also continuous repetition (i.e. integration). As a result, the language can conceive of a curve with a fixed amount of turning at any moment – a circle –, or a curve

where the amount of turning increases continuously – a spiral –, etc. Our proposal implies that both discrete repetition and continuous path integration are primitive concepts in the human language of geometry. Indeed, a key hypothesis of the present work is that the human mind can encode discrete as well as continuous changes and integrate them within a single language of thought. This part of our proposal is deeply related to the near-universal existence of a system of aspect in human languages, thus betraying the existence of continuous and discrete concepts of time and repetition (compare for instance the imperfect, e.g. “the curve was turning”, with the perfect, e.g. “the curve turned”) (Comrie, 1976).

In summary, we propose that the human mental representation of geometric shape involves a language of thought that can produce virtually all the geometric line drawings observed in human cultures as combinations of a minimal set of geometric primitives. Our core hypothesis is that perceiving a shape, in humans, consists in finding the shortest program that suffices to reproduce it. Our proposal thus connects shape perception to the problem of program induction, i.e. the identification of a program that produces a certain output. In line with much previous work (e.g. Chater & Vitányi, 2003; Feldman, 2000, 2003; Li & Vitányi, 1997; Mathy & Feldman, 2012; Romano et al., 2013), we hypothesize that the perceived complexity of a shape is determined by its minimum description length (MDL) in the proposed language.

Below, we describe the proposed language in detail, list its predictions, and test them in two experiments. First, we show that our language predicts which shapes are judged simple. Second, we show that any such language has to satisfy a set of additive relationships for repeated, concatenated or embedded shapes, and that those universal laws can be experimentally validated.

### A generative language for geometric shapes

Here, we make our proposal concrete by introducing a specific language, somewhat similar to Logo’s turtle language, for generating a variety of geometric line drawings. The language we propose is based on two postulates. First, we assume that all humans possess a set of primitive operators or “mental routines” (Ullman, 1984) that serve as building blocks for more complex programs. We included elementary primitives that have been proven to be present in human children or adults in the absence of formal education; several of them are likely to be inherited from primate evolution. Our primitives comprise the concepts of

- Small exact integers (Feigenson et al., 2004) which

can be minimally generated by the successor function (Izard et al., 2008),

- Fractions, i.e. ratios of those integers (Jacob & Nieder, 2009; Siegler et al., 2011)
- Straight line (see e.g. Izard et al., 2011b),
- Heading direction (Muller et al., 1996) and how it changes when we turn,
- Path integration (Dehaene et al., 2006; Gallistel, 1990; Leeuwenberg, 1971; McNaughton et al., 2006; O’Keefe & Nadel, 1978),
- Right angle turn (Dehaene et al., 2006; Dillon et al., 2019; Izard et al., 2011b)

Extensions of this list, for instance to large approximate numbers, would be straightforward and are considered in the discussion, but as we shall see, those primitives appear to suffice to account for a broad variety of geometric shapes that humans universally consider simple.

Our second postulate is that, in humans only, a compositional language of thought allows these primitive operators to be combined into larger programs. We suppose that three composition instructions are available: concatenation; repetition; and call to a subprogram.

### Program instructions

The full language, described in Figure 2, contains the following instructions. First, as in the “logo” language (Abelson et al., 1974), drawing instructions dictate the movements of a pen that can move and trace curves on a plane. Those instructions are *Turn*, which changes the current heading of the pen; *Move*, which changes the position of the pen by a certain amount in the current direction without tracing; and *Trace*, which traces a curve by integrating over a set of parameters (duration, speed, acceleration, and turning speed).

Second, the three control structures are *Concatenate* (also denoted by “;”) which executes one program and then another; *Repeat*, which repeats a program a certain number of times (twice by default); and *Subprogram* which saves the current state, executes a given program, and resumes the previous state for the rest of the execution.

Third, since these instructions require either discrete or continuous arguments, the language contains a number system, with integers (*Int*) and numerical (*Num*) types. For computational simplicity, in order to avoid a huge combinatorial explosion that would prevent the enumeration of all minimal programs, we did not include a full algebra, in spite of recent evidence that humans may

possess one (Grace et al., 2020). Instead, the *Int*’s are built using Peano arithmetic starting from 1 (the language has a *one* primitive and a *successor* primitive, and the *Num* are either signed integers (positive or negative), or signed fractions of two integers. This is enough to generate rational numbers, but prevents nesting of fractions.

While the numbers generated by our language are unitless, they are interpreted differently depending on the functions in which they are evaluated. For the *Turn* function, an argument of 1 is interpreted as “one right angle” (i.e. the unit for angle is “right angle”). Similarly, *Move* and *Trace* instructions use implicit units of length and speed, such that the default values (1) on duration and turning speed imply turning by a full circle. These hypotheses, while plausible, are not crucial, since changing them would only minimally change the predicted shape complexities (e.g. if the default turn was by 180°, a right-angle turn would still be available at a minimal cost, as half of it).

### Calculation of minimum description length

Our language does not guarantee that each shape can be generated by a single program. Quite the opposite: whenever a shape can be generated at all, an infinity of programs are available to generate it. Our third key postulate is therefore that humans search for the shortest program that draws a given shape. We refer to the complexity of the shortest program for a given shape as its Minimum Description Length (MDL), and the corresponding program(s) as the Minimal Program(s). Notice that this is compatible with a Bayesian framework, or probabilistic program induction (Lake et al., 2015): since the number of programs of MDL  $n$  increases exponentially with  $n$ , the log likelihood of a given program will be proportional to its MDL.

The complexity of a program is defined as the number of nodes in its syntax tree, or equivalently the number of primitives in the program, with two exceptions. First, whenever a signed number is required, for instance when turning by a certain angle, an additional node is needed (indicating + or -). This node was not counted, thus preventing the cost of signed values from being systematically higher than that of unsigned values (e.g. one, vs. +one) Second, concatenations did not increase MDL. This feature arose as a by-product of our implementation, which used continuations (any program can take a program as an argument, and if it does then it executes it once it is done, which is equivalent to concatenation). We checked that the results did not change dramatically when adding a cost for concatenations, as the results should not hinge on this implementation detail.

<b>Program</b> :=			
<b>Program</b> ; <b>Program</b>		Concatenate : run one program and then another	control
<b>Repeat</b> ([ <b>Int</b> =2]) { <b>Program</b> }		Repeat a program a certain number of times	
<b>Subprogram</b> { <b>Program</b> }		Execute a program, then restore the original state	
<b>Trace</b> ([ <b>t</b> = <b>Int</b> =1], [ <b>speed</b> = <b>Num</b> =+1], [ <b>acceleration</b> = <b>Num</b> =+0], [ <b>turningSpeed</b> = <b>Num</b> =+0])		Trace a curve by moving according to the parameters	drawing
<b>Move</b> ([ <b>t</b> = <b>Num</b> =+1])		Move a certain distance without tracing anything	
<b>Turn</b> ( <b>angle</b> = <b>Num</b> )		Rotate the current heading	
<b>Int</b> :=			arithmetic
<b>one</b>		Number 1	
<b>Next</b> ( <b>Int</b> )		Successor function	
<b>Num</b> :=			
<b>+Int</b>   <b>-Int</b>		Return a signed number	
<b>+Int/Int</b>   <b>-Int/Int</b>		Return the signed fraction of two integers	

Figure 2

*Proposed language of thought for the mental representation of geometric shapes.* The figure lists all primitive operators and their parameters. As indicated in the right column, *control* primitives act on programs, *drawing* primitives move the pen on the plane in various ways, and *arithmetic* instructions generate integers and fractions that are passed as parameters. Green, instructions; pink, types; blue, named parameters; gray, default values for optional arguments (denoted by brackets).

### Examples

The minimal program to draw a square is `Repeat { Repeat { Trace ; Turn(angle=+one) } }` (where the infix operator “;” denotes concatenation). This program works because, in the absence of any argument, *Repeat* defaults to 2 repetitions. Since a turn of one means a right-angle turn, this program concatenates four segments, each ending with a right-angle turn. The cost of this program is 5 (repeat + repeat + trace + turn + one).

As a slightly more complex example, the following program draws a triskelion (🌀), a classic Celtic figure:

```
Repeat(next(next(one))) {
  Subprogram {
    Trace(acceleration=
one/next(next(one)),
turningSpeed=one)
  };
  Turn(angle=next(next(next(one)))/next(next(one)))
}
```

This program draws three identical inward spirals using the *Trace* instruction. Note that the number 3 is coded as `next(next(one))` (again, this assumption is adopted for simplicity; adding primitives for numbers 2 and 3

would only minimally change the predicted cost). The *Subprogram* instruction ensures that, after drawing a single spiral, the position is reset to the origin. The *Turn* instruction, which takes 4/3 as its argument, ensures that the three spirals are oriented at 120 degrees from each other.

### Simulation results

We first examined the shortest programs in the proposed language, and whether they always generate shapes that are simple and frequently attested in human cultural history. To this end, we wrote a program that systematically enumerates all possible programs in order of increasing MDL, and draws the corresponding shapes (we eliminated, automatically as well as manually, the shapes that could be generated by a simpler program). Figure 3 shows a subsample of the resulting shapes, sorted by MDL. The simplest, lowest complexity shapes are extremely simple: they consist of a line segment (MDL=1), then a circle (MDL=2) and a spiral (MDL=3). The low-complexity shapes with MDL = 4 or 5 are also excellent candidates for cultural universals: repeating circles, dashed lines, spirals with various numbers of loops, and other simple mathematical shapes such as the square, the half-circle, or two tangent circles. At this stage, the concatenation instruction also gener-

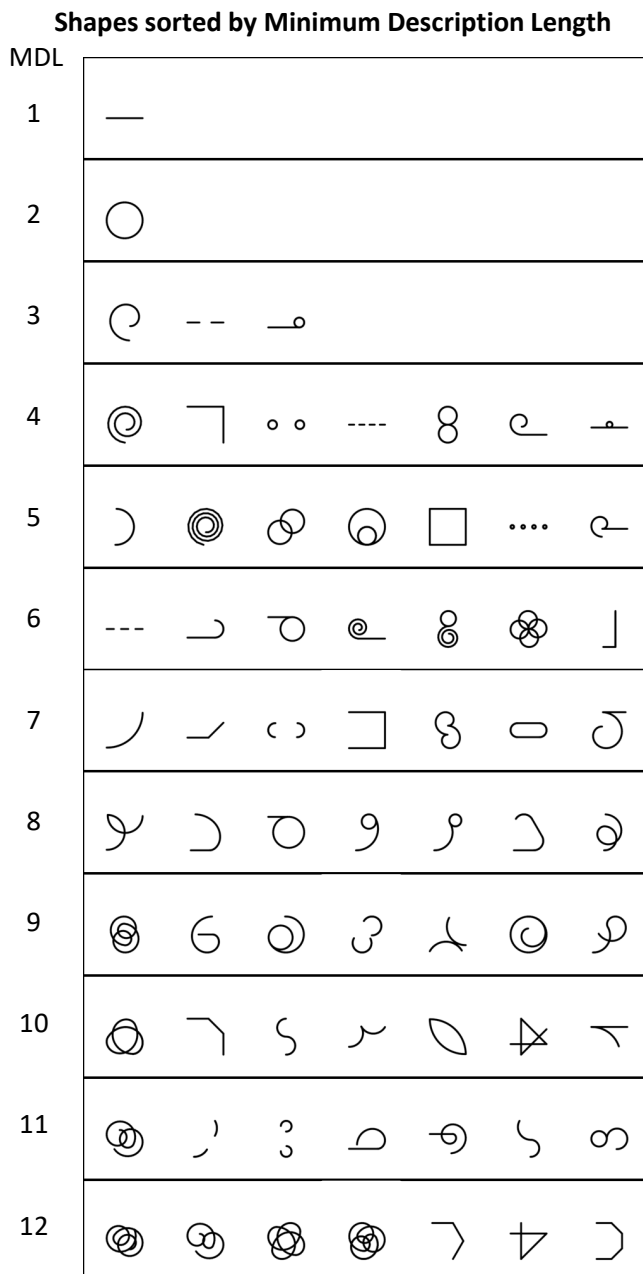


ates less intuitive, but still culturally attested shapes such as a “sigma” (segment + circle) or a “crosier” (segment + spiral). As MDL increases, the huge combinatorial explosion of programs results in an enormous variety of shapes, only some of which are culturally observed. Nevertheless, the shapes with low MDL remain introspectively simple (this intuition is tested formally in experiment 1 further below). This observation is in stark contrast with most other such languages where the combinatorial explosion creates short programs with complex, unintuitive graphic outputs (e.g. logo).

### Program induction using DreamCoder

A crucial aspect of our proposal is that humans encode a shape mentally by inferring a simple program that could generate it. Thus, the perception of a simple shape is an act of “program induction”. Yet it is implausible that humans scan through thousands of programs before recognizing a square. Otherwise, the time required to recognize a shape would grow exponentially with the length of its shortest program. Thus, it is important to show that such an inference is, at least approximately, computationally feasible in our specific case. While program induction remains a difficult challenge for computer science, we leveraged a state-of-the-art program induction technique, the DreamCoder algorithm (Ellis et al., 2021). This algorithm is given programming problems via examples of the desired behavior, and searches for the simplest program that performs the task. Here, a task reduces to a shape, and DreamCoder has to find the shortest program that generates it. DreamCoder internally represents the language as a probabilistic grammar and enumerates programs according to their likelihood using the probabilistic weights of the grammar.

Two features of DreamCoder speed up the search. First, the weights are task-dependent and are suggested by a neural network for a given task. For instance, DreamCoder may learn that shapes with straight lines call the Trace instruction without any TurningSpeed argument. The neural network can be trained without any environmental input or supervision, simply by sampling a random program, generating the corresponding shape (called “dreaming”) in a top-down manner, then using this internally generated shape-program data pair in supervised learning in order to adjust the bottom-up weights from the shape to its program representation (see also Lake et al., 2017). Second, DreamCoder builds new *abstractions* for pieces of programs that are often used for a given set of tasks: for example, if the shapes contain many right angles, it may create a new abstraction Turn(angle=1), thereby increasing the likelihood of programs that use it. In a Bayesian sense, this cor-



**Figure 3**

*Sample shapes generated by the enumeration of all programs in the proposed language. Programs were listed by increasing MDL. Identical or perceptually indistinguishable shapes that could be generated by a simpler program were eliminated. Starting at MDL=4, only a limited sample of 7 shapes is shown, as the number of shapes increases exponentially with MDL.*

responds to updating the priors over the space of programs. This abstraction mechanism is useful to capture

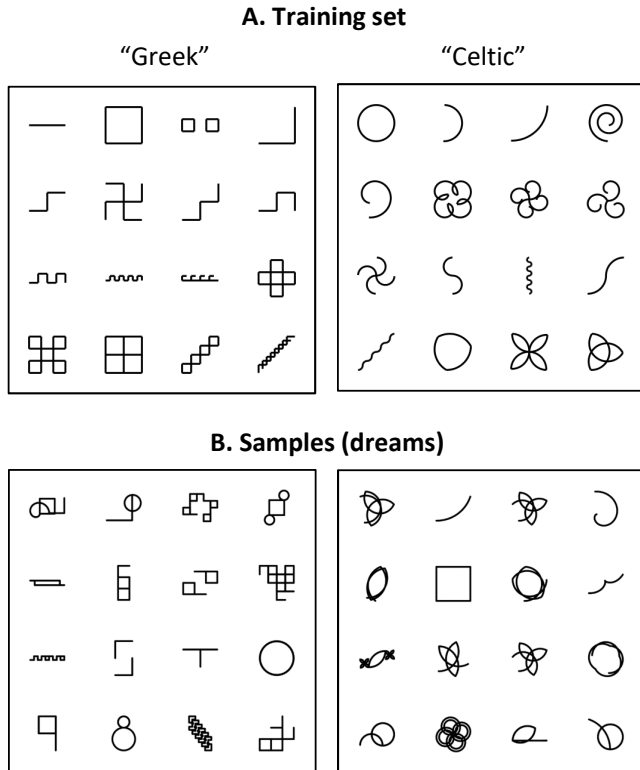


Figure 4

**Testing the DreamCoder algorithm for program induction.** *A, shapes used in a training phase. We verified that, in response to all 32 shapes shown, the algorithm was able to identify a short, presumably minimal program that could generate it. The algorithm was trained either with the square shapes at left ("Greek" style) or with the circular shapes at right ("Celtic" style). B, examples of additional shapes spontaneously generated by sampling from the grammar learned during the training set, and therefore biased towards a certain geometric style.*

regularities in a corpus of shapes: subprograms used to draw the simplest ones can be reused to draw more complex ones. Interestingly, these two mechanisms interact. As the grammar becomes biased towards using certain program schemas, the neural network also becomes biased towards recognizing them; for example, the neural network might increase the probability of the "turn" primitive when it sees angles, or that of "repeat" when it notices repeating patterns.

We found that, together, those two mechanisms made program induction feasible for our language, at least for relatively simple shapes. In Figure 4, we present two separate corpora, one with mostly rectilinear shapes (referred to as "Greek"), and one with mostly curvilinear

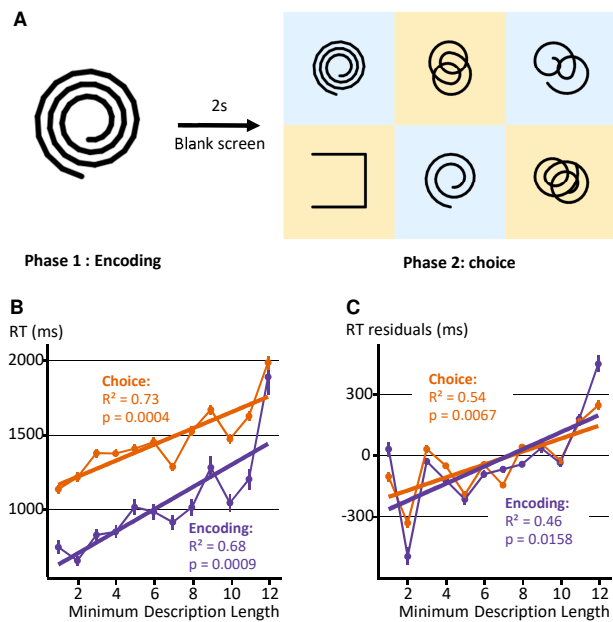
shapes (referred to as "Celtic"). After learning, DreamCoder could solve all of these tasks. Interestingly, the abstractions it created were different. This could be visualized by sampling "dreams" from the resulting grammars. All simple shapes (circles, lines, squares) appear in both cases, but the grammars exhibit a bias towards shapes that resemble the training set (bottom row in figure 4).

The DreamCoder approach opens up a number of perspectives on how human cognition could efficiently address the problem of program induction. First, it naturally accounts for cultural drifts: while shapes such as circles and squares are universally shared, cultures are also characterized by the frequent use of specific patterns (e.g. linear Greek friezes versus curvilinear Celtic spirals). This arises even though the geometric primitives are universal, because each culture adopts, initially by chance, some preferred combination of primitives, which are then internalized as frequent subprograms or program fragments and progressively cement a specific style of geometric patterns. Second, DreamCoder may explain how simple geometric shapes may be efficiently recognized and used by young children in the absence of much or any training (poverty of the stimulus argument). This is because the top-down system (from programs to shapes) can be used to train the bottom-up system (from shapes to programs) via the use of "dreams", i.e. internally generated training data. In an improved version of DreamCoder, the bottom-up neural network could be repurposed to directly retrieve the most plausible program for a given shape, thus providing a possible mechanism by which humans can quickly identify simple shapes. Finally, another promising aspect of this proposal, which remains to be fully explored, is the possibility of creating reusable abstractions or program templates. While the square, for instance, is not a primitive of the original language, a square-drawing program schema may become abstracted over time, thus allowing the participant to easily understand concepts such as "a square of circles" or "a square twice larger than the previous one", etc. At present, however, for simplicity, such named subprograms are not part of the current language, but solely of the DreamCoder program-induction software.

### Experiment 1: predicting geometric complexity

Does the proposed language have any psychological reality? In experiment 1, we tested the simplest prediction of our proposal: the perceived complexity of a shape should be determined by its minimal program length. If humans represent shapes as mental programs, then for tasks involving shape perception and manipulation, the MDL of the shape should predict the difficulty of the





**Figure 5**

**Procedure and results for experiment 1. A, task structure.** On each trial, participants pressed a key for as long as they needed to memorize a shape (encoding phase), then, after a 2-second delay, had to select the corresponding shape among a 2x3 grid (choice phase). Note that there was a size change between phase 1 and phase 2. **B, correlation between MDL and behavior.** Average encoding time (purple) and choice time (orange) are plotted as a function of MDL. Error bars represent one standard error across participants. **C same**

task. For example, as MDL increases, it should take longer to encode the shape in working memory, and to compare it with other shapes.

This prediction should hold only provided that other simple perceptual strategies do not suffice to perform those tasks. In previous work, we found that the perception of quadrilaterals could be based on two systems: a list of symbolic rules akin to those arising from the current language (e.g. right angle, equal sides), and only available to humans; and a classical invariant shape recognition system, well captured by a convolutional neural network model of the ventral visual pathway, and available to both human and non-human primates (Sablé-Meyer et al., 2021). Thus, to properly test the existence of the first system, it is important to cancel out the potential contributions of the second.

Here, we asked participants to memorize a sample geometric shape and perform a delayed match-to-sample

task where, after a 2-second delay, that shape had to be selected from an array of 6 possible choices, some of which were perceptually quite similar (Figure 5A). We measured the choice time, but also the encoding time by letting participants view the sample shape for as long as they wished, holding down the space bar until they were ready to decide; as soon as they released the space bar, the sample shape disappeared, then after a fixed delay, the choices appeared. We hypothesized that both encoding time and choice time would be predicted by MDL.

## Methods

**Participants.** Participants were 125 adults tested online (53 females, 69 males, 3 preferred not to answer; age range 20 to 78, mean and median 44 years old), recruited and tested online via messages on social media. The task was approved by our university’s committee for ethical research (CER-Paris-Saclay-2019-063) and participants gave informed consent.

**Procedure and Stimuli.** On each trial, we showed participants a sample shape for a variable duration, then an empty screen for 2 seconds, and finally a 2x3 choice screen with 6 shapes. Participants were asked to click on the shape that was shown originally. Participants controlled how long they looked at the reference shape: when they were ready to start a trial, they clicked at the center of the screen (thus centering their mouse), then pressed the space bar on their keyboard. The sample shape was shown for as long as the space bar was depressed. Upon releasing the space bar, the shape disappeared and they were left with a blank screen. Instructions insisted on the self-paced nature of the task: “Keep the spacebar depressed for as long as you deem necessary to remember the shape well”. We call the duration of the press the encoding time, as it is an indirect measure of the time that participants needed to encode the target shape.

The choice screen comprised six different shapes that were displayed on an isoluminant blue/yellow checkerboard (Figure 5A). At that point, participants could click on a shape, which ended the trial. We measured both accuracy and response time, which we refer to as choice time.

The experiment comprised 6 initial training trials, then 68 trials with 68 unique testing shapes, each appearing only once as a sample. The 68 testing shapes are shown in figure 3, while the training shapes were 6 additional shapes with MDL = 5 in our language. During training, the distractors were always the same six shapes and participants had each of them as a target once. Piloting indicated that the choice of distractors

was crucial for the performance to vary from one shape to the other: if the distractors were too dissimilar from the sample shape, participants learned to press and release the spacebar as fast as they could, knowing that a purely perceptual strategy sufficed for virtually perfect accuracy. To mitigate this strategy, we selected distractors closely matched to each shape. We computed, for each shape, the four closest ones in figure 3 as defined by two metrics: (i) the average gray level (average pixel value of an image), and (ii) the difference in the vector codes of the shapes within the last layer of a convolutional neural network of object recognition, CORnet (Kubilius et al., 2019). CORnet is a convolutional neural network that figures amongst the top-scoring models of the ventral stream according to BrainScore (Schrimpf et al., 2018, 2020). Then, on each trial, we presented on the choice screen, at random locations: (1) the correct target shape; (2) two of the four closest shapes according to CORnet; (3) two of the four closest shapes according to average gray level; (4) a last shape uniformly sampled among the remaining test shapes. The selection algorithm ensured that all 6 choice shapes differed. The choice of shapes and their placement were fully randomized for each participant, independently within training and within testing.

## Results

Overall error rate was very low (1.82%), so we concentrated our analysis on response times. We removed all participants who failed on 5 or more trials, as well as participants whose overall average encoding time or choice time was higher than the group mean plus three standard deviations (9 participants removed in total; 116 remaining). We also removed, for each shape, trials where the encoding time exceeded the average encoding time of that shape plus three standard deviations (and similarly for choice time). This procedure removed 3.8% of the total number of trials.

To test for the predicted effect of MDL on behavior, we performed simple linear regressions on the encoding times and choice times as a function of the MDL. Both measures were significantly correlated with the MDL of the target shape (Figure 5B; encoding time:  $R^2=.68$ ,  $p<.001$ ; choice time:  $R^2=.73$ ,  $p<.001$ ). We also performed between-subjects ANOVAs on the linear effect of MDL (numerical factor) across participants. The main effect of MDL was again highly significant (encoding time:  $F(1,115)=148.5$ ,  $p<.001$ ; choice time:  $F(1,115)=480.7$ ,  $p<.001$ ).

Since MDL showed a small but significant partial correlation with gray level ( $R^2=.07$ ,  $p=.027$ ), we replicated this analysis by first removing the main effect of

gray level on response times, then performing a simple linear regression on the residuals. The effect of MDL was again significant (Figure 5.C; encoding time:  $R^2=.46$ ,  $p=.016$ ; choice time:  $R^2=.54$ ,  $p=.006$ ; between-subjects ANOVAs, encoding time:  $F(1,115)=118.2$ ,  $p<.001$ ; choice time:  $F(1,115)=248.1$ ,  $p<.001$ ).

To compare the effects of MDL and gray level, we also performed a multiple-regression analysis with both variables (normalized) as predictors, across the 68 test shapes. For encoding time, both predictors were significant (both  $p<.0001$ ;  $R^2=.59$ ; betas = 157.7 (Standard Error [SE]=37.4) for MDL and 272.0 (SE=37.4) for gray level). For choice times, both predictors were also significant (MDL:  $p=.0003$ ; gray  $p<.0001$ ;  $R^2=.53$ ; betas = 108.7 (SE=28.8) for MDL and 182.75 (SE=28.8) for gray level).

We also controlled for additional visual features (see Figure S1). For each shape, we manually counted the number of extremities (variable  $N_{ex}$ ), intersections where two lines meet (singularities, variable  $N_{i2}$ ), intersections where three lines meet ( $N_{i3}$ ), intersections where four or more lines meet ( $N_{i4+}$ ), disconnected parts ( $N_{dis}$ ), and finally the presence of a closed shape or not (closure). To confirm that the effect of MDL did not result from a spurious correlation with these features, we first ran a model comparison between a mixed-effect model with all features, the gray level, and the average spatial frequency, with participants as random intercepts, and the same model plus the MDL. The second model was significantly better than the first one for both dependent measures (Likelihood Ratio Test, encoding time  $\chi^2 = 106.4$ ,  $p < .001$ ; choice time  $\chi^2 = 169.8$ ,  $p < .001$ ). Figure S1 shows the breakdown of each shape’s property, as well as a comparison of the predictors associated to each normalized term in the full model. As a final control, we computed the residuals from the first model for each participant, then averaged those residuals over participants and MDL levels, and examined the correlation of those residuals with MDL. This linear model was not significant for the encoding time ( $r^2=.04$ ,  $p=.51$ ), but it was for choice time ( $r^2=.58$ ,  $p=.003$ ). Inspection of the residuals show that the “segment” shape, which is the only shape with MDL=1, had a very high residual for the encoding time. After removing this outlier item, both encoding time and choice time residuals were now significantly predicted by (encoding time:  $r^2=.5$ ,  $p=.015$ ; choice time:  $r^2=.63$ ,  $p=.003$ ).

## Discussion of experiment 1

As predicted, in a delayed match-to-sample task with geometric shapes, participants were influenced by the shape’s Minimum Description Length (MDL) in our pro-

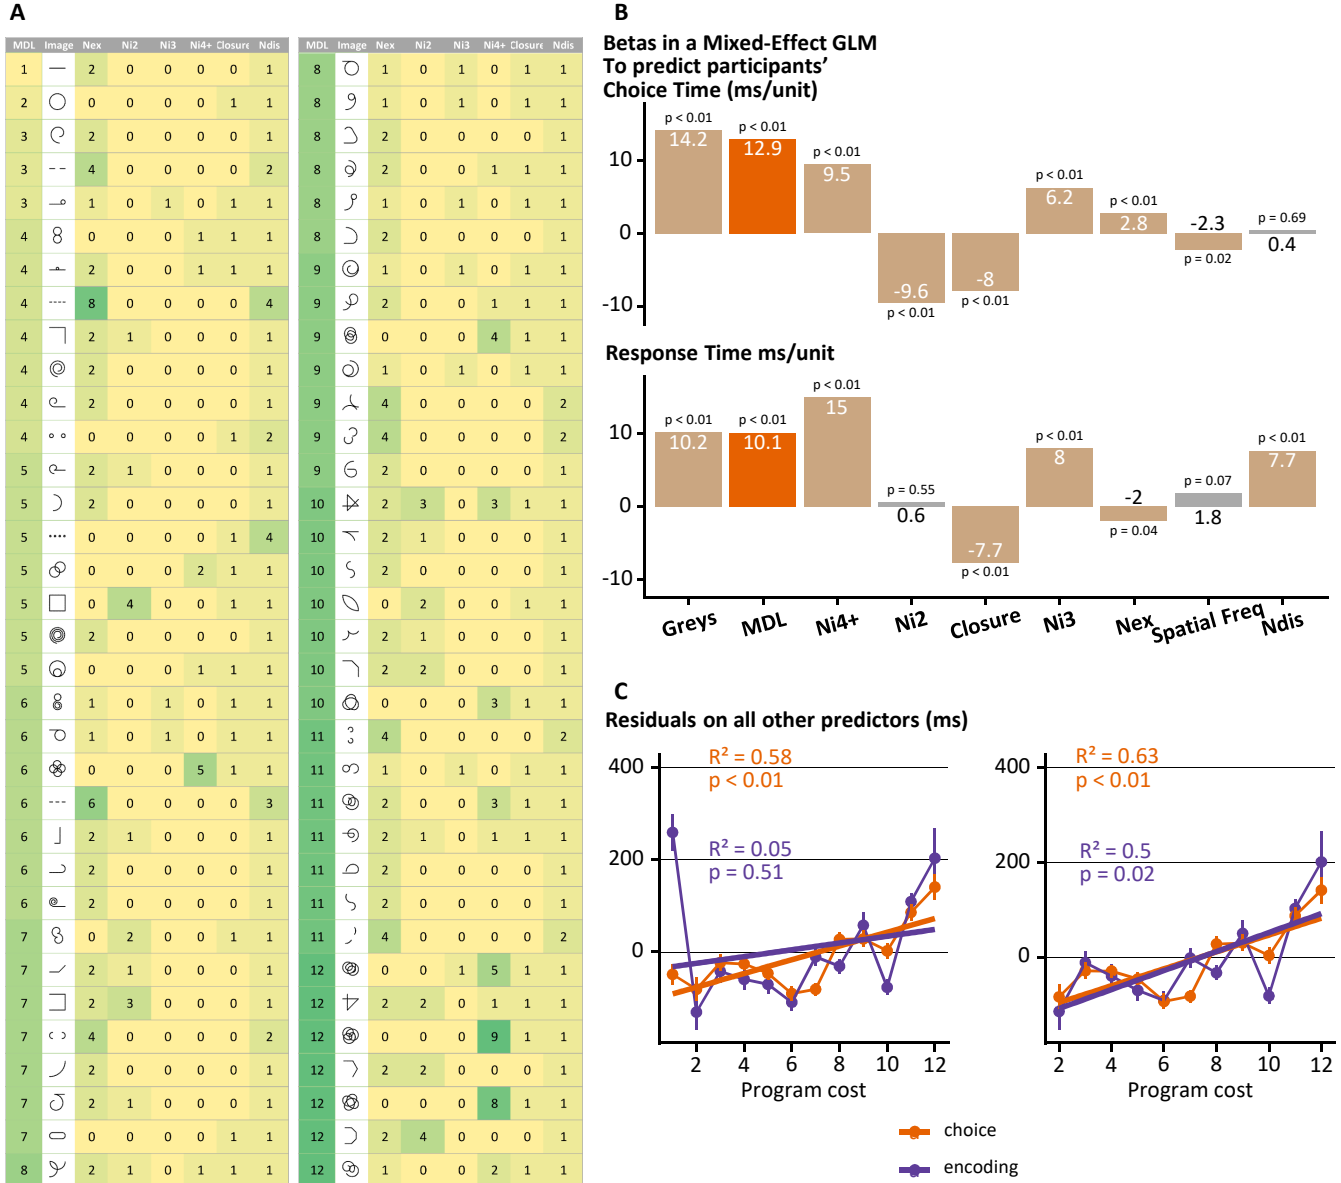


Figure S1

*Detailed analysis of Exp. 1 A, breakdown of the properties. For each shape, describes the manually annotated or computed visual features. B, Predictive power of the visual features.\*\* We ran a mixed-effect GLM with all the features, normalized for comparison, as well as the cost, and display the fitted coefficient and its significance, for both choice time and response time. Greyed out predictors were not significant in the regression, and bright orange indicates our regressor of interest, the MDL. C. The simple linear model MDL and either choice or response time, on the residuals of a mixed-effect model with all visual features, after averaging per MDL value. On the right, we display the same regression after removing the condition MDL=1.*

posed language. This result held at individual levels, and applied to both encoding time and choice time, with similar effect sizes. As shown in figure 3 and in the results, the effect of MDL was not trivially correlated with other image properties such as number of parts, intersec-

tions, closure, etc. While behavior was also influenced by the perceptual property of average gray level, the effect of MDL remained even when controlling for this low-level effect and all other variables. This result supports the prior suggestion that there are two strategies

available to perform the task (Sablé-Meyer et al., 2021): one that encodes shapes at a visual level, based on perceptual properties such as gray level, intersections, etc; and one that represents them at an abstract, symbolic level, which is well captured by our proposed language.

While this result is promising, there are many reasons to believe that our proposal is not complete, and that it would be possible to find shapes for which the fit would be poor. As we further examine in the general discussion, there are shapes that the language cannot easily describe (e.g. ovals), or properties it cannot easily encode (e.g. while an equilateral triangle has a low MDL, there is no way to describe “any triangle”). Thus, while experiment 1 tested the specifics of our language, experiment 2 was designed to test regularities that any such reasonable language should verify.

### Experiment 2: Fundamental laws of repetition, concatenation and embedding

To sidestep the constraints that come with choosing a specific language, we designed an experiment that tests, independently, the three most fundamental aspects of our proposition, namely the existence of operations of *repetition*, *concatenation* and (*nested*) *embedding* (that is, recursive call to a subprogram). Those operations, available in any modern programming language, have the highest impact on the compressibility of shapes and help decorrelate the length of a program and the amount of “ink” of a shape. They allow a program to be extremely short and yet the shape to be complex, as long as it is highly regular. For instance, from the programs for squares and circles, a single additional instruction suffices to generate a circle of squares, or a square of circles – and thus, this addition should just have an additive effect on MDL.

These observations lead to the following quantitative prediction: in any language for geometric shapes which includes primitives of repetition, concatenation and embedding, the cost of complex shapes should be the sum of the lengths of (1) the program(s) that are being repeated, concatenated or embedded, plus (2) a fixed cost for the instruction itself and, if necessary, its operands. Consider for instance a figure formed by two shapes placed side by side: we predict its complexity to be the sum of the complexity of each shape plus some constant for the concatenation instruction. Likewise, the cost for a repetition of a shape should only depend on the cost of that shape, plus a constant to express the repetition, and a cost for the parameter “number of repetitions”. Finally, the complexity of a figure consisting of one shape embedded in another (e.g. a square of circles) should be the cost of each shape plus a constant for the “embed”

instruction.

In summary, the following relations should hold (at least for  $x \neq y$ , see below):

$$[1] \text{Cplx}(\text{repeat}(x, n)) = \beta_0 + \beta_1 * \text{Cplx}(x) + \beta_2 * \text{Cplx}(n)$$

$$[2] \text{Cplx}(\text{concat}(x, y)) = \beta_0 + \beta_1 * \text{Cplx}(x) + \beta_2 * \text{Cplx}(y)$$

$$[3] \text{Cplx}(\text{embed}(x, y)) = \beta_0 + \beta_1 * \text{Cplx}(x) + \beta_2 * \text{Cplx}(y)$$

Note that the multipliers  $\beta_1$  and  $\beta_2$  should be close to 1 if the length of the program is the only factor that comes into play, but might exceed 1 if additional factors intervene (e.g. interference between the two shapes in working memory).

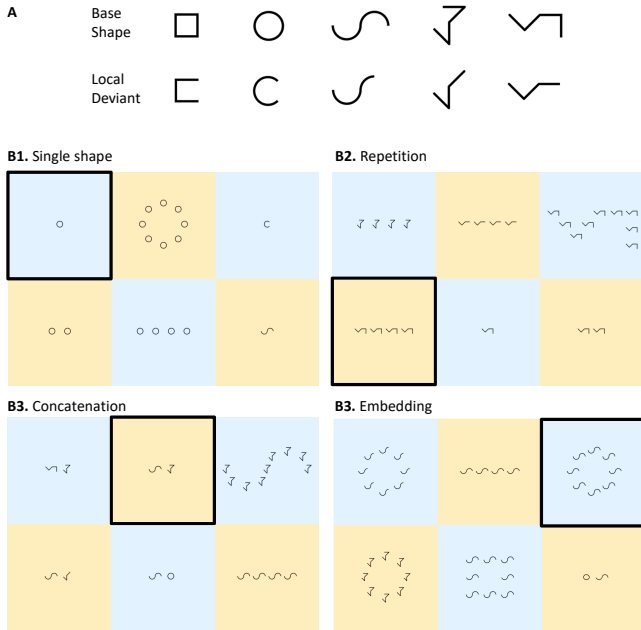
In experiment 2, we therefore replicated the delayed match-to-sample with new stimuli. We selected five base shapes spanning a broad range of predicted complexities, and used them to build new stimuli through repetition, concatenation and embedding, with the goal of testing whether their complexity could be predicted from the complexity of their base shapes. Thus, we designed a total of 60 images that served as samples in the task : (1) five base shapes, shown in figure 6; (2) five repetitions of those base shapes, generated by showing side by side four copies of the base shape; (3) twenty-five concatenations corresponding to all 5x5 pairs of base shapes placed side by side; and (4) twenty-five embeddings of a base shape inside another, generated by drawing the outline of one shape using 8 or 9 copies of the other. Example stimuli are shown in figure 6.

In a separate group of participants, we also measured the subjective complexity of the same shapes. We presented participants with those 60 shapes in random order and asked them to evaluate each shape’s complexity on a scale from 0-100.

### Method

**Participants.** Participants were recruited via Twitter. One hundred and seventy adults participated in the main delayed match-to-sample task (71 male and 99 females), with a breakdown of 16 participants in the 18-25 age group, 77 in the 25-40, 68 in the 50-60 and 9 in the 60+. Participants were not compensated for participating in this study. An additional 27 adults participated in the subjective ratings (15 females and 12 males). Both tasks were approved by CER-Paris-Saclay-2019-063, and participants gave informed consent.

**Procedure and Stimuli for the delayed match-to-sample task.** The procedure for this experiment was identical to that of our first experiment, and only the



**Figure 6**

**Stimuli used in experiment 2.** **A** the 5 base shapes which were presented alone and in combinations, and the corresponding deviant shapes. **B**, Examples of screens presented during the choice phase in the four conditions of the experiment (randomly intermixed): single shape, repetition of a base shape, concatenation of two base shapes, and embedding of a base shape into another. In each case, the sample image which was presented during the encoding phase is highlighted, and the other cells of the 2x3 grid illustrate the diversity of distractors. A distinct group of subjects also gave subjective ratings of complexity for each stimulus.

stimuli changed. The stimuli were generated from five base shapes that were piloted to vary in encoding and choice time (square, circle, S, sigma, and square root, of respective MDL 5, 2, 10, 15, 13; see Figure 6). All shapes had similar amounts of ink, or gray level: the square, sigma and square root all had four segments of identical length, arranged differently; the circle matched the square in length; and the S shape comprised two semi-circles and was therefore matched with the circle. Those base shapes were then used to generate the 60 target stimuli for the four experimental conditions: *single shape* (5 stimuli), *repetition* (a string of 4 identical shapes; 5 stimuli), *concatenation* (2 shapes side by side;  $5 \times 5 = 25$  stimuli), and *embedding* (an inner shape was presented at the usual size, but in 8 or 9 copies that formed the outlined of a second, outer shape;  $5 \times 5 = 25$  stimuli).

Each of those 60 stimuli served as samples for the delayed match-to-sample task. During the choice period, the sample stimulus was intermixed with 5 distractors, which were generated in order to prevent short-cut strategies and maximize the need for a full identification of the target shape. For each base shape, we designed a local deviant by removing one fourth of the shape (see Figure 6). For the *single shape* and the *repetition* conditions, the distractors were (1) a distractor from the same condition, but using the local deviant instead of the target, (2) a distractor from the same condition, but using a different shape, (3) three distractors drawn from each of the other three conditions, and sharing at least one shape with the target. For the *concatenation* condition, the distractors were (1) a *concatenation* distractor where one of the two shapes was replaced by its local deviant; (2) two *concatenation* distractors where either the left or right shapes were replaced by a different base shape, (3) one distractor from the *embedding* condition, using the same two shapes as the left and right shapes, assigned randomly to embedded and embedding, (4) one distractor from the *repetition* condition, using randomly either the left or the right shape. Similarly, for the *embedding* condition, the distractors were (1) a distractor from the *embedding* condition, with the same outer shape but the inner shape replaced by its local deviant; (2) two *embedding* distractors where either the outer or the inner shape was replaced by a different base shape, (3) one distractor from the *concatenation* condition, but using the same two shapes as the inner and outer shapes, (4) one distractor from the *repetition* condition, using either the inner or the outer shape. Our logic was that this set of distractors covered a broad range of programs in the proposed language, each with a small local change or “bug” – thus forcing subjects to search for the shape whose description exactly matched the sample. Figure 6B shows, for each condition, an example of a target and its five distractors.

In addition to those 60 trials, 10 initial training trials allowed participants to get used to the task flow and to the difficulty level of the memory task. Training trials were generated similarly, but using three different base shapes and the same exact procedure. The experiment proceeded seamlessly from training to testing trials, without any notice.

**Procedure and stimuli for subjective complexity ratings.** Upon clicking on a link from the twitter message, participants landed on an experiment designed with jsPsych (De Leeuw, 2015). The experiment started with a consent form as well as a small demographic questionnaire for age group and sex. Then they were presented with instructions for the task: using sliders from

0-100, they had to give a complexity rating to each of the 60 sample shapes. To familiarize them with the type of shapes, the instructions included 24 shapes that did not appear afterward. We highlighted in the instructions that they should focus on trying to be consistent across shapes. Participants were then presented with all 60 shapes in a shuffled order. Participants could freely look at the shapes in any order and change their rating until they were satisfied. The task took a median time of 6:34 minutes to answer (1<sup>st</sup> quartile, 4:48; 3<sup>rd</sup> quartile, 9:36.).

## Results

None of the subjective rating data were rejected. For the delayed match-to-sample task, the error rate was low (3.39%), and we removed data following the same strategies as in experiment 1 (7 participants and 3.16% trials rejected).

Figure 7 shows the encoding time, choice time, and subjective complexity for each of the 60 stimuli, as a function of the base shape(s) used to generate them. First, we verified that, in the single-shape condition, all three dependent measures varied across the 5 base shapes we selected (ANOVAs where participants with missing data were removed, all  $p < .01$ ). For encoding and choice times, the increase was roughly monotonic: square and circle were roughly on par, and then the response times increased for the S, sigma and square-root shapes, in this order (note that this is close to, but not strictly identical to the MDL ordering, which was circle, square, S shape, square root and sigma). The subjective ratings followed a noisier profile, but still ranking the last two shapes as more complex than the first 2.

We then examined whether, when those shapes entered into more complex stimuli, the complexity still varied in a similar way, predictable from equations [1-3]. In the *repetition* condition with 4 identical shapes, a significant, monotonic profile of response was observed for all three dependent measures (ANOVAs where participants with missing data were removed, encoding time  $F(4,576)=53.88$ ; response time  $F(4,576)=36.87$ ;  $F(4,104)=6.63$ ; all  $p < .001$ ). Similarly, we analyzed just the stimuli of the *concatenation* condition where two identical shapes were placed side by side, thus corresponding to a repetition of 2. Again, for all dependent measures, a significant, nearly monotonic increase was seen across the 5 base shapes (participants with missing data removed; encoding time,  $F(4,576)=32.20$ ; choice time  $F(4,576)=41.01$ ; ; subjective  $F(4,104)=15.29$ ; all  $p < .001$ ). Importantly, the curves for single shape, 2 repetitions and 4 repetitions were nearly parallel to each other, as predicted by our equations for MDL. To test

this idea, we entered all 3 conditions into a mixed-effect linear model with two main factors and their interaction, i.e. the repeated shape (5 levels), and the number of repetitions (numerical factor spanning 1, 2 or 4 repetitions). As shown in Table 1, we found main effects for each measure, in agreement with equation [1]. There was no significant interaction with the number of repetitions for both subjective rating and choice time, in agreement with equation [1], but encoding time significantly interacted with the number of repetitions.

For the concatenation and embedding conditions, as shown in Figure 7, encoding time, choice time, and subjective complexity also increased with each of the two shapes involved (left/right or inner/outer). Equations [2] and [3] predicted that these effects should be similar to the single shape condition and should not interact. Table 1 shows the results of mixed effect linear models, but now with the shapes as two 5-level factors. Starting with the Concatenation condition, for all 3 dependent variables, we found significant effects of both the shape on the left and the one on the right, as predicted, but the interaction terms were also significant. However, there was a very simple explanation: the 5x5 design matrix for concatenation included 5 stimuli in which the left and right shapes were identical. In this case, our theory predicts that these stimuli should be more compressible, and therefore easier to perceive than other concatenation stimuli. This is exactly what we found: as seen in the middle panels of figure 7 (where, for simplicity, those data points are connected by a gray line), the repeated stimuli stood out as faster and subjectively less complex than the corresponding concatenation stimuli. When we removed those conditions from the mixed-effect model (Table 1), the F values for the main effects increased while those for the interaction decreased massively, although a small but significant interaction remained for both encoding and choice time.

Finally, we ran the same analysis on the *embedding* condition, using both embedded and embedding shape as the 2 factors, plus their interactions. Again, both main effects were significant on all 3 dependent measures, as predicted by equation [3]. However, as in the concatenation condition, the interaction term had a significant effect on encoding time and choice time, which our additive equations did not predict. Inspired by our analysis of concatenation, we plot separately the “self-embedding” trials, in which the same shape was used at the inner and outer levels (e.g. a square of squares, a circle of circles, etc). In figure 7 (right panels), we can see that those data points again yielded lower values than the others (i.e. lower subjective complexity, faster encoding and choice times). When we excluded those



Repetition		Shape		Number		Interaction	
	Subjective Rating	$F_{4,369}=2.95$	$p=.020$	$F_{1,369}=11.56$	$p<.001$	$F_{4,369}=0.88$	n.s.
	Encoding Time	$F_{4,2215.25}=4.66$	$p<.001$	$F_{1,2215.10}=185.21$	$p<.001$	$F_{4,2215.17}=14.86$	$p<.001$
	Choice Time	$F_{4,2215.41}=28.31$	$p<.001$	$F_{1,2215.23}=35.75$	$p<.001$	$F_{4,2215.31}=0.76$	n.s.
Concatenation		Left shape		Right shape		Interaction	
	Subjective Rating	$F_{4,624}=22.16$	$p<.001$	$F_{4,624}=16.71$	$p<.001$	$F_{16,624}=3.83$	$p<.001$
<i>Including self-concatenated shapes</i>	Encoding Time	$F_{4,3796.63}=52.88$	$p<.001$	$F_{4,3796.54}=45.02$	$p<.001$	$F_{16,3796.56}=27.67$	$p<.001$
	Choice Time	$F_{4,3797.13}=76.72$	$p<.001$	$F_{4,3796.97}=58.94$	$p<.001$	$F_{16,3797.01}=25.77$	$p<.001$
<i>Excluding self-concatenated shapes</i>	Subjective Rating	$F_{4,494}=25.85$	$p<.001$	$F_{4,494}=20.80$	$p<.001$	$F_{11,494}=0.45$	n.s.
	Encoding Time	$F_{4,3004.67}=57.78$	$p<.001$	$F_{4,3004.59}=50.76$	$p<.001$	$F_{11,3004.66}=1.85$	$p=.041$
	Choice Time	$F_{4,3005.32}=78.42$	$p<.001$	$F_{4,3005.18}=63.97$	$p<.001$	$F_{11,3005.31}=2.16$	$p=.014$
Embedding		Outer shape		Inner shape		Interaction	
	Subjective Rating	$F_{4,624}=66.58$	$p<.001$	$F_{4,624}=28.54$	$p<.001$	$F_{16,624}=0.90$	n.s.
<i>Including self-embedded shapes</i>	Encoding Time	$F_{4,3786.68}=119.23$	$p<.001$	$F_{4,3786.67}=61.62$	$p<.001$	$F_{16,3786.70}=3.64$	$p<.001$
	Choice Time	$F_{4,3787.10}=115.12$	$p<.001$	$F_{4,3787.08}=118.95$	$p<.001$	$F_{16,3787.12}=6.74$	$p<.001$
<i>Excluding self-embedded shapes</i>	Subjective Rating	$F_{4,494}=55.57$	$p<.001$	$F_{4,494}=24.90$	$p<.001$	$F_{11,494}=0.42$	n.s.
	Encoding Time	$F_{4,2996.89}=103.49$	$p<.001$	$F_{4,2996.79}=55.88$	$p<.001$	$F_{11,2996.82}=1.20$	n.s.
	Choice Time	$F_{4,2997.45}=111.20$	$p<.001$	$F_{4,2997.29}=110.49$	$p<.001$	$F_{11,2997.33}=3.74$	$p<.001$

**Table 1**

The table shows the statistics of mixed effect models applied to our three main conditions (Repetition, Concatenation and Embedding) and to our three dependent variables (Subjective Rating, Encoding time and Choice time). In each case our model had two main effects and their interaction, plus a single random effect of the participant. P-values were computed using the Kenward-Roger approximation for degrees of freedom. For Repetition, the main effects were the repeated shape (5 levels), and the number of repetitions (numerical factor in 1, 2 or 4, respectively for the shape alone, for the concatenation of two identical shapes, and for the repetition condition). For Concatenation, the main factors were the shape on the left and the shape on the right. For embedding, the main effects were the outer/embedding shape, and the inner/embedded shape. For Concatenation and Embedding, we ran the model a second time after removing the stimuli in which the same base shape was used twice (e.g. two squares, or a square of squares).

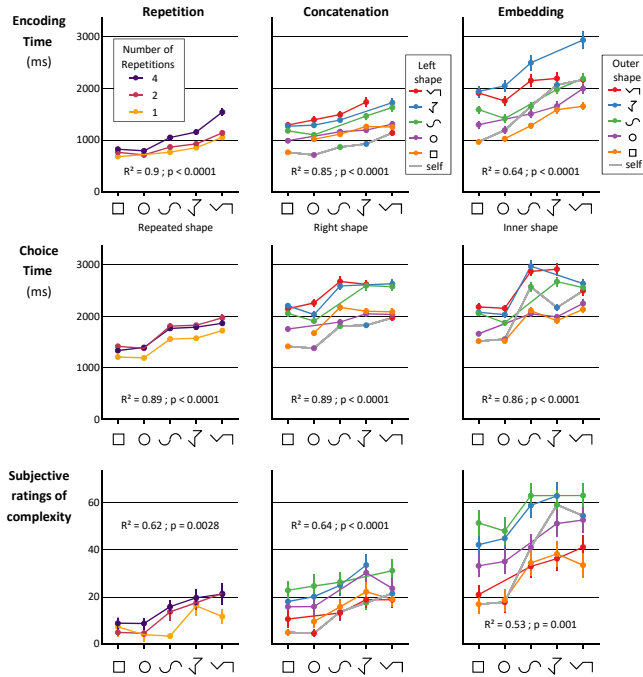


Figure 7

**Results of experiment 2.** The three behavioral measures (encoding time, choice time, and subjective complexity) are shown for the 60 different stimuli, grouped into repetition (left), concatenation (middle), and embedding (right). The conditions are sorted as a function of base shape complexity. Note that the single-shape condition is shown in the repetition panel (number of repetitions = 1), and that the data from the concatenation of two identical shapes appears twice (as number of repetitions = 2 in the left panels, and as dots connected by a gray line in the middle panels). Similarly the data for “self-embedded” shapes (e.g. a square of squares) appear with a gray line in the right panels. In each panel, we show the coefficient of determination and statistical significance of the associated model (equations [1-3] in the main text). Error bars indicate standard errors across subjects.

self-embedding trials from the mixed-effect model, the interaction terms again collapsed while the main effects did not (the interaction did remain significant for choice time). Although those observations were not predicted, they can easily be accommodated: it appears that the mental representation of a “square of squares” involves a saving, because both the embedded and the embedding shapes are identical, and thus presumably the same mental program is used twice.

We further tested the predictions of equations [1-3] using General Linear Models (GLMs). Those equations imply

that we should be able to accurately reconstruct the complexity of composite shapes from that of the five base shapes. To test this, we fitted generalized linear models on each of our dependent variables after averaging data across participants for each item (Figure 7; all fitted values are provided in Table 2). First, we modeled the repetition conditions (base shape, two shapes, and four shapes) by predicting, for each dependent measure, its value for a given trial as a linear function of its value in the single shape condition and the number of repetitions. This model was significant for each of our three dependent variables, and all predictors were significantly different from 0 (all  $p < .05$ ; subjective rating,  $R^2 = .62$ ; encoding time,  $R^2 = .90$ ; choice time,  $R^2 = .89$ ). The fitted coefficients for the shape term were all close to one (non-significantly so for subjective complexity and choice time, and only slightly larger than 1 for encoding time), suggesting that base shape complexity was directly reflected in the complexity of the repeated shape.

Similarly, we modeled the concatenation condition by predicting the complexity of a trial with a linear combination of the complexity of the left shape, the complexity of the right shape, and a dummy variable (termed IsSelf) for whether the left and right shapes were identical (as these trials are instances of repeat and should have a lower perceptual complexity). This model was significant for each of our three dependent variables, and all predictors were significantly different from 0 (see Table 2; all  $p < .05$ ; subjective,  $R^2 = .64$ ; encoding,  $R^2 = .85$ ; choice,  $R^2 = .89$ ). The coefficients for both left and right shapes were not significantly different from 1 at the  $p < .05$  level, indicating that both contributed equally and directly, as predicted from equation 2. The IsSelf predictor was always significantly negative, indicating a saving when both shapes are identical.

Finally, we applied the same model to our embedding condition. Again, a good fit was found, with significant effects of both shapes for all dependent measures (see Table 2; all  $p < .05$ ; subjective,  $R^2 = .53$ ; encoding  $R^2 = .64$ ; choice  $R^2 = .86$ ). In this between-items analysis with only 25 average data points, unlike more powerful mixed-model approaches, the impact of self-embedding did not reach significance (subjective rating  $p = .33$ ; encoding time  $p = .14$ ; response time  $p = .06$ ). This time, the slopes tended to be higher than 1, although this reached significance only in the case of encoding time for the outer shape. We also observed a trend towards a larger influence of the outer shape compared to the inner shape, consistent with the outer shape’s greater visual impact, but again this effect was only significant in a single dependent measure (subjective ratings). Those minor trends notwithstanding, the main finding is that

the data supported equation [3]: even when we made the visual pattern much more complex by embedding one shape inside another, thus creating for instance a circle of squares, the final complexity was still only a linear function of the complexity of the individual shapes.

### Discussion of experiment 2.

Participants' behavior matched the additive properties predicted by equations 1 through 3: the complexity of complex patterns could be decomposed into the sum of the complexity of constituents. This property was observed in three different metrics that we collected as proxies for the complexity of the mental representation of shapes: subjective complexity ratings, encoding times and choice times in a delayed match-to-sample task.

More specifically, and following the equations, the following three results were verified: [1] the complexity of the mental representation of  $n$  identical shapes is the complexity of the mental representation of the shape, together with a cost that increases with  $n$ ; [2] the complexity of the mental representation of two different shapes side by side is the sum of their respective complexities; and [3] the complexity of the mental representation of a shape drawn in outline using several small copies of a different shape (which we called embedding) is predicted by the sum of their respective complexities. The latter finding is the most interesting, as intuition alone might have predicted a product operation – after all, the overall pattern comprises as many copies of the inner shape as needed to draw the outer shape. However, the prediction from the “language of thought” perspective is clear enough: drawing a square of circles is not much more complex than drawing a square itself – it merely requires stopping the square program at regular intervals to call a subprogram for drawing a circle, and in first approximation, such embedding only has a linear effect on total complexity. Of course, this is only true in first approximation – our “square of circles”, for instance, comprised additional circles not only at the vertices of the square, but also in the middle of its sides, thus requiring a slightly more complex square-drawing program. Such subtleties, which require further investigation, may explain why the slope measuring the impact of the outer shape on the complexity of the overall pattern tended to be larger than 1 for embedded shapes (see Table 2).

Two other salient effects emerged, one which could be predicted from equations [1-3] and another which could not. First, when concatenating two identical shapes, the resulting shape can be either described as a repetition or as a concatenation – but our language of thought predicts that programs involving repetition are shorter,

and therefore that the complexity of a pair of shapes should be lower than predicted by concatenation alone. Our data supports this prediction: concatenations of two identical shapes have a lower complexity than predicted by the sum of the complexities of the two shapes, each with their respective coefficients, indicating that identical shapes induce a saving. Second, unexpectedly, the same phenomenon occurred in the embedding condition when a shape was outlined using a smaller version of itself (“self-embed” trials, e.g. a square of squares). Once again, the data points to the lower complexity of those trials, compared to those using two different shapes. Such a saving is not captured by equation [3]. It suggests that in the mental representation of e.g. “a square of squares”, the two squares may be represented by a single mental program or at least by some degree of sharing of working memory resources. This interesting finding supports the idea that human mental representations allow for recursive calls or higher-order functions (functions over functions) (Dehaene et al., 2015; Fitch, 2014; Hauser et al., 2002).

Overall, Experiment 2 goes beyond experiment 1 in showing that, over and above the specific complexity predicted by the particular geometrical language we proposed, there are several properties of additivity that must be satisfied by any such language, and that these properties are true of the human working memory for geometric shapes.

### General Discussion

Previous research has emphasized that all humans inherit, from evolution, core knowledge of space and number that they share with many other animal species (Dehaene et al., 2006; Feigenson et al., 2004). Here, we propose that, in humans, those core systems can also be recombined using a language of thought in order to form complex mental programs. As a result, humans are able to form complex, compositional thoughts such as “three parallel lines”, “repeat a pattern four times”, or “arrange some circles in the shape of a square”.

In the present work, we argued that such combinatorial mental representations underlie human perception and working memory for geometric shapes, and we put this hypothesis to several tests. First, we proposed a concrete language inspired by observations on prehistoric and ethnographic geometric patterns (featuring abstract patterns, right angles, parallel lines, circles and spirals) as well as elements from the core-knowledge literature (number sense). We tested this language in Experiment 1 and showed that it could predict the behavior of participants in a shape memory task, above and beyond lower-level perception mechanisms. In Experiment 2,

<b>Repeat</b>		<b>Shape</b>	<b>Number of repetitions</b>	
	Subjective	$0.88 \pm 0.23^{**}$	$1.98 \pm 0.89^*$	
	Encoding	$1.37 \pm 0.15^{***}$	$87.69 \pm 16.45^{***}$	
	Choice Time	$1.05 \pm 0.11^{***}$	$47.31 \pm 19.1^*$	
<b>Concatenate</b>		<b>Left shape</b>	<b>Right shape</b>	<b>IsSelf</b>
	Subjective	$0.72 \pm 0.2^{**}$	$0.67 \pm 0.2^{**}$	$-8.33 \pm 2.37^{**}$
	Encoding	$0.92 \pm 0.17^{***}$	$0.89 \pm 0.17^{***}$	$-433.15 \pm 55.29^{***}$
	Choice Time	$0.96 \pm 0.12^{***}$	$0.81 \pm 0.12^{***}$	$-521.96 \pm 63.28^{***}$
<b>Embedding</b>		<b>Outer shape</b>	<b>Inner shape</b>	<b>IsSelf</b>
	Subjective	$1.84 \pm 0.43^{***}$	$0.95 \pm 0.43^*$	$-5.26 \pm 5.27$
	Encoding	$1.94 \pm 0.45^{***}$	$1.83 \pm 0.45^{***}$	$-222.82 \pm 148.07$
	Choice Time	$1.29 \pm 0.16^{***}$	$1.23 \pm 0.16^{***}$	$-169.22 \pm 85.43$

**Table 2**

The table shows, for each shape type (repeat, concatenate, and embedding), each dependent variable (subjective rating, encoding time or choice time), and for each variable in the model (in columns), the value of the regression slope, associated standard error, and whether the effect is significantly different from 0 (\*,  $p < .05$ ; \*\*,  $p < .01$ ; \*\*\*,  $p < .001$ ).

we further showed that theoretically motivated additive equations for the complexity of composite shapes characterize humans’ subjective ratings and objective behavior in a delayed match-to-sample task, thereby constraining the properties that any proposition for a language of thought for geometric shape must satisfy.

The proposed theory of shape perception requires an explanation for how one can efficiently go from a visual percept to a mental program, a problem referred to as program induction. The idea that humans infer mental programs was previously shown to successfully account for human concept learning (Lake et al., 2015), including hand-drawn sketches (Ellis et al., 2018). However, program induction is also a computational challenge in computer science. Enumerating all possible programs until a match is discovered is not a plausible strategy, as the search time would scale exponentially with the size of the program. We show how recent work from the program-induction literature helps tackle this problem by using DreamCoder (Ellis et al., 2021) to find the best representation for several shapes. DreamCoder uses a bottom-up neural network to speed up the search for the relevant program. In DreamCoder, the network is trained to map each visual shape onto biases that accelerate the search for the relevant program. Future versions could incorporate direct mappings from shapes to programs, thus leading to the immediate recognition of shapes close enough to the training set. What is remarkable about this idea is that the system does not

need any external training data (although successfully solved problems will be used when training, much like replays, in addition to “dreams”): it can generate its own supervised learning dataset by sampling programs, executing them to produce the corresponding shapes, and then training a neural network to perform the backward inference from shape to program. The notion of “inner training” is an interesting metaphor for how humans may explore, in a purely mental manner, the domain of geometry and discover interesting properties on their own, without external inputs. This is possible if we assume, as René Descartes did in the introductory citation, that our mind already has the ideas, at least in the form of a large space of potential mental programs.

To further accelerate its search, the DreamCoder algorithm looks for subprograms that are reused across several shapes. This mechanism is useful to go from simple shapes to more and more complex ones, as each new success offers the possibility of discovering new abstractions. This mechanism changes the topography of the search space by bringing certain shapes (those that leverage the discovered abstractions) closer to the effectively searchable threshold for program induction: while all shapes remain accessible, some of them become much simpler as they can be expressed more succinctly using the discovered abstractions. This mechanism could explain cultural drifts in the style of geometrical patterns, whereby a given human culture focuses on certain shapes and their variants, thus producing, over time, a variety

of similar-looking patterns.

The core of our proposal builds upon seminal work by Leyton (Leyton, 1984, 2003), who proposed a generative theory of shapes. Leyton’s theory stipulates that all shapes are constructed in a bottom-up fashion by a sequence of operations, called “control groups”, starting from a single point. For example, the mental representation of a cube would be the extrusion across the z-axis of a square, which itself would come from turning a segment 4 times around a central point. The segment itself would be built by translating the starting point along an axis. The unpublished experiments briefly mentioned by Leyton (1984), however, seem to have probed those hypotheses only indirectly, for example by asking participants to perform intuitive judgements about the stability of certain shapes. While less general and confined to two dimensions, our proposal is supported by direct empirical tests of the mental complexity of shapes.

Earlier work by Leeuwenberg (Leeuwenberg, 1971) also introduced a language for shapes, particularly focused on the idea of nested repetitions with variations, and the concept of continuous integration as a complement of discrete repetition capable of tracing curves. However, Leeuwenberg’s language only satisfied one direction of the intended correspondence between mental and linguistic complexity: the language was such that low-complexity mental representation corresponded to short programs, but the converse was not true as some short programs generated shapes that were not easily parsed by humans. Nevertheless, the behavioral results he reported are in line with ours: he found the length of the shortest program is a good predictor of subjective ratings of complexity, as well as objective performance in shape copying and other similar tasks (Boselie & Leeuwenberg, 1986; Leeuwenberg, 1969, 1971; see also Restle, 1970, 1973).

More recently, several articles highlighted the importance of the notion of repetition with variation in the human perception of geometric and auditory sequences (Amalric et al., 2017; Piantadosi et al., 2016; Planton et al., 2021; Roumi et al., 2021; Simon, 1972). The language first proposed by Amalric et al. (2017) included two distinct notions of repetitions with variations, either changing the starting point with each repetition, or changing a parameter with each repetition. Our geometric language generalizes this idea to the case of continuous curve tracing. Our tracing primitive can be considered a particular case of infinitesimal repetition with variation: at each time step, parameters are updated for the computation of the next time step. Importantly, we did not allow for arbitrary infinitesimal repetition like what is present in the Logo language (Abelson et al.,


1974), where at each time step arbitrary computations can occur. While such computations were required in the original logo in order to draw curves by infinitesimally changing the heading at each time step, it also opened the possibility of very small programs having extremely complex outputs. By limiting ourselves to simple, linear variations of speed or turning angle over time, our language fits with the universal presence of a limited set of shapes (mostly lines, circles and spirals) in human geometric patterns.

Another, more speculative aspect of our proposal is that a recursively compositional language of thought for geometry is unique to humans. As also proposed by others, only humans would possess a recursive compositional capacity (Dehaene et al., 2015; Fitch, 2014; Hauser et al., 2002; Penn et al., 2008). Our recent work (Sablé-Meyer et al., 2021) shows that, for quadrilaterals, human behavior differs strikingly from baboons and is characterized by a symbolic geometrical regularity effect: in humans only, regular quadrilaterals (which can be compressed in the present language) are easier to perceive than less regular ones. Importantly, in this work, baboon behavior was not random, but could be captured by models of the ventral visual pathway for object recognition. Furthermore, in preschoolers and in human adults without formal education, both the object recognition model and the symbolic geometry model were jointly needed to account for behavior. Thus, two strategies seem to be available for geometric shape perception: a purely visual strategy, available to both baboons and humans, and a symbolic geometry strategy, putatively available only to humans. It is important to take this dual-route model into consideration when testing the present idea: if the stimuli can be too easily discriminated by object recognition alone, then humans may not make the effort to encode them symbolically, in the proposed language of thought, and MDL may cease to determine performance (as we observed in pilot data).

We do not yet have non-human primate data testing the current proposal, but a preliminary data set corresponding to experiment 1 was recently collected in baboons, in collaboration with Joel Fagot, and the results conform to our hypothesis: baboon behavior was driven by visual similarity alone, without any influence of MDL in our language of thought, contrary to humans.

### Limits of our language

Simple arguments show that the proposed language can generate most of the shapes that humans find simple and that are frequently attested in human cultures as well as in the history of geometry. Any regular polygon, for instance, such as an equilateral triangle, can be

generated by a program similar to the square, but using fractions of a right-angle turn; the pseudocode would read: Repeat  $p$  times { Trace ; Turn( $\text{angle}=4/p$ ) }. Stars with various numbers of branches can be similarly generated. Less regular polygons, such as a rectangle approximating the proportions of the famous golden section, can be generated using fractions (e.g.  $5/3$  or  $8/5$ ). Symmetrical patterns and friezes, such as (svastika) or  arise naturally from recursive combinations of repetitions and alternations. Finally, using concatenation or embedding, these patterns can be combined to generate, for instance, a pentagram (star inside a circle), a circle of circles, etc.

Nevertheless, some simple figures remain difficult to generate with a short program. A trapezoid, for instance, comprises two parallel sides interrupted by two arbitrary segments. Drawing a trapezoid in our language requires a turn by an arbitrary angle  $\alpha$ , followed by a second turn by  $2\text{-}\alpha$  to restore parallelism. However, our language does not have variables that could store the value  $\alpha$ , and hence does not find this shape simpler than an arbitrary quadrilateral with turns  $\alpha$  and  $\beta$ . In general, our approach is unable to encode a partial regularity inside an otherwise arbitrary figure. The addition of local variables could address this limitation, but exploration of this idea suggested that straightforwardly adding variables has a high cost: while they allow one to express otherwise hard to describe simple shapes, they also make very complex shapes easy to describe, an undesirable feature that is hard to keep in check.

There are also shapes for which our language proposes implausible programs. For instance, the minimal program that draws a  $+$  shape repeats four segments starting from the center, instead of drawing the shape using two intersecting segments. As for continuous shapes, our program cannot account for some simple shapes such as the ellipse or parabola. These shapes might be better represented as visual transformations of the outputs of another program (e.g. compressing a circle to get an ellipse). Such a dual-mode system, combining the generative and transformative capacities of mental imagery, has been proposed by others (e.g. Kosslyn, 1980; Leyton, 2003; Shepard & Cooper, 1982). The addition of a buffer in which mental operations such as rotations or shearing could be applied would be an important addition to the present work.

### Future directions

We see two promising ways forward to go beyond the present work. One, drawing inspiration from Kosslyn, Shepard and Leyton's work, would allow for further mental manipulation of the outputs of the present language,

thus performing operations such as deformations, rotations, or even extrusions as a post-processing step. This would open up to the modeling of 3-D shape, for instance using rotation around a fixed axis. Another possibility would be to integrate ideas from a different programming paradigm. Currently, our programs are imperative in nature: they describe, in every detail, the sequence of operations needed to draw a shape. However, not all programming languages work that way. Logic programming languages such as Prolog describe the logic of the computation and its constraints rather than the details of its execution flow. Such constraint-based programming may be closer to how humans think, particularly in the mathematical domain. The canonical definition of a circle, for instance, involves a constraint (equidistance from a center point) rather than a generative program (turn by a fixed curvature, as in the present language). Integrating both declarative and imperative elements may provide a better account of specific shapes, such as trapezoids or generic triangles, which only possess some properties (e.g. two parallel sides, or three sides) while leaving the other details unspecified. Some geometrical shapes could thus be defined in terms of properties that they satisfy, others in more detailed imperative instructions, and both could be reused in control structures such as embed or repeat. Meanwhile, we surmise that the present proposal merely brings us one step closer to understanding how mathematical objects are mentally encoded.

### Acknowledgements

### Funding

This research was supported by a European Research Council (ERC) grant "NeuroSyntax" to S.D., as well as funding from INSERM, CEA, Collège de France and Fondation Bettencourt-Schueller. M.S.M. was supported by a doctoral grant from Ecole Normale Supérieure. J. T. was supported by an Air Force Office of Scientific Research grant (FA9550-19-1-0269), an MIT-IBM Watson AI Lab grant (W1771646), and an NSF grant (CCF-1231216). K.E. was supported by an NSF GRFP Award.

### Data Availability

The data from the experiment, the scripts to generate the figures and statistical analyses performed, an interpreter for our language, and a modified version of DreamCoder to handle it are provided in an Open Science Framework repository: [https://osf.io/zcb64/?view\\_only=7abbac57edf34b5ab98e976da1a79b6b](https://osf.io/zcb64/?view_only=7abbac57edf34b5ab98e976da1a79b6b)

The latest version of DreamCoder is available online on GitHub <https://github.com/ellisk42/ec>



The reader can test our language for geometric shapes by using the following online interpreter: [https://private.unicog.org/msm/LoT\\_Geometric\\_Shapes.html](https://private.unicog.org/msm/LoT_Geometric_Shapes.html)

### References

- Abelson, H., Goodman, N., & Rudolph, L. (1974). *Logo manual*.
- Amalric, M., Wang, L., Pica, P., Figueira, S., Sigman, M., & Dehaene, S. (2017). The language of geometry: Fast comprehension of geometrical primitives and rules in human adults and preschoolers. *PLoS Computational Biology*, *13*(1), e1005273.
- Bornstein, M. H., Gross, C. G., & Wolf, J. Z. (1978). Perceptual similarity of mirror images in infancy. *Cognition*, *6*(2), 89–116.
- Boselie, F., & Leeuwenberg, E. (1986). A test of the minimum principle requires a perceptual coding system. *Perception*, *15*(3), 331–354.
- Brincat, S. L., & Connor, C. E. (2004). Underlying principles of visual shape selectivity in posterior inferotemporal cortex. *Nat Neurosci*, *7*(8), 880–886.
- Brincat, S. L., & Connor, C. E. (2006). Dynamic Shape Synthesis in Posterior Inferotemporal Cortex. *Neuron*, *49*(1), 17–24. <https://doi.org/10.1016/j.neuron.2005.11.026>
- Butterworth, B., Reeve, R., Reynolds, F., & Lloyd, D. (2008). Numerical thought with and without words: Evidence from indigenous Australian children. *Proc Natl Acad Sci U S A*, *105*(35), 13179–13184.
- Chater, N., & Vitányi, P. (2003). Simplicity: A unifying principle in cognitive science? *Trends in Cognitive Sciences*, *7*(1), 19–22.
- Comrie, B. (1976). *Aspect: An Introduction to the Study of Verbal Aspect and Related Problems*. Cambridge University Press.
- De Leeuw, J. R. (2015). jsPsych: A JavaScript library for creating behavioral experiments in a Web browser. *Behavior Research Methods*, *47*(1), 1–12.
- Dehaene, S., Izard, V., Pica, P., & Spelke, E. (2006). Core knowledge of geometry in an Amazonian indigene group. *Science*, *311*, 381–384.
- Dehaene, S., Meyniel, F., Wacogne, C., Wang, L., & Pallier, C. (2015). The Neural Representation of Sequences: From Transition Probabilities to Algebraic Patterns and Linguistic Trees. *Neuron*, *88*(1), 2–19. <https://doi.org/10.1016/j.neuron.2015.09.019>
- Dillon, M. R., Duyck, M., Dehaene, S., Amalric, M., & Izard, V. (2019). Geometric categories in cognition. *Journal of Experimental Psychology: Human Perception and Performance*. <https://doi.org/10.1037/xhp0000663>
- Ellis, K., Ritchie, D., Solar-Lezama, A., & Tenenbaum, J. B. (2018). Learning to Infer Graphics Programs from Hand-Drawn Images. *ArXiv:1707.09627 [Cs]*. <http://arxiv.org/abs/1707.09627>
- Ellis, K., Wong, C., Nye, M., Sablé-Meyer, M., Morales, L., Hewitt, L., Cary, L., Solar-Lezama, A., & Tenenbaum, J. B. (2021). DreamCoder: Bootstrapping inductive program synthesis with wake-sleep library learning. *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*, 835–850. <https://doi.org/10.1145/3453483.3454080>
- Feigenson, L., Dehaene, S., & Spelke, E. (2004). Core systems of number. *Trends Cogn. Sci.*, *8*(7), 307–314.
- Feldman, J. (2000). Minimization of Boolean complexity in human concept learning. *Nature*, *407*(6804), 630.
- Feldman, J. (2003). The simplicity principle in human concept learning. *Current Directions in Psychological Science (Wiley-Blackwell)*, *12*(6), 227–232. <https://doi.org/10.1046/j.0963-7214.2003.01267.x>
- Fitch, W. T. (2014). Toward a computational framework for cognitive biology: Unifying approaches from cognitive neuroscience and comparative cognition. *Physics of Life Reviews*, *11*(3), 329–364. <https://doi.org/10.1016/j.plrev.2014.04.005>
- Fodor, J. A. (1975). *The language of thought* (Vol. 5). Harvard University Press.
- Frankland, S. M., & Greene, J. D. (2020). Concepts and Compositionality: In Search of the Brain’s Language of Thought. *Annual Review of Psychology*, *71*(1), 273–303. <https://doi.org/10.1146/annurev-psych-122216-011829>
- Gallistel, C. R. (1990). *The organization of learning*. MIT Press.
- Goodenough, F. L. (1926). *Measurement of intelligence by drawings*.
- Grace, R. C., Carvell, G. E., Morton, N. J., Grice, M., Wilson, A. J., & Kemp, S. (2020). On the origins of computationally complex behavior. *Journal of Experimental Psychology: Animal Learning and Cognition*, *46*(1), 1–15. <https://doi.org/10.1037/xan0000227>
- Harris, D. B. (1963). *Children’s Drawings as Measures of Intellectual Maturity: Revision of Goodenough Draw-a-man Test*.

- Hauser, M. D., Chomsky, N., & Fitch, W. T. (2002). The faculty of language: What is it, who has it, and how did it evolve? *Science*, *298*(5598), 1569–1579. <https://doi.org/10.1126/science.298.5598.1569>
- Henshilwood, C. S., d’Errico, F., Yates, R., Jacobs, Z., Tribolo, C., Duller, G. A. T., Mercier, N., Sealy, J. C., Valladas, H., Watts, I., & Wintle, A. G. (2002). Emergence of Modern Human Behavior: Middle Stone Age Engravings from South Africa. *Science*, *295*(5558), 1278–1280. <https://doi.org/10.1126/science.1067575>
- Hermer, L., & Spelke, E. S. (1994). A geometric process for spatial reorientation in young children. *Nature*, *370*(6484), 57–59.
- Hung, C.-C., Carlson, E. T., & Connor, C. E. (2012). Medial Axis Shape Coding in Macaque Inferotemporal Cortex. *Neuron*, *74*(6), 1099–1113. <https://doi.org/10.1016/j.neuron.2012.04.029>
- Izard, V., Pica, P., Spelke, E. S., & Dehaene, S. (2008). Exact Equality and Successor Function: Two Key Concepts on the Path towards Understanding Exact Numbers. *Philosophical Psychology*, *21*(4), 491–505. <https://doi.org/10.1080/09515080802285354>
- Izard, V., Pica, P., Spelke, E. S., & Dehaene, S. (2011a). Flexible intuitions of Euclidean geometry in an Amazonian indigene group. *Proceedings of the National Academy of Sciences*, *108*(24), 9782–9787.
- Izard, V., Pica, P., Spelke, E. S., & Dehaene, S. (2011b). Flexible intuitions of Euclidean geometry in an Amazonian indigene group. *Proceedings of the National Academy of Sciences of the United States of America*, *108*(24), 9782–9787. <https://doi.org/10.1073/pnas.1016686108>
- Jacob, S. N., & Nieder, A. (2009). Notation-independent representation of fractions in the human parietal cortex. *J Neurosci*, *29*(14), 4652–4657.
- Joordens, J. C. A., d’Errico, F., Wesselingh, F. P., Munro, S., de Vos, J., Wallinga, J., Ankjærgaard, C., Reimann, T., Wijbrans, J. R., Kuiper, K. F., Mûcher, H. J., Coqueugnot, H., Prié, V., Joosten, I., van Os, B., Schulp, A. S., Panuel, M., van der Haas, V., Lustenhouwer, W., ... Roebroeks, W. (2015). Homo erectus at Trinil on Java used shells for tool production and engraving. *Nature*, *518*(7538), 228–231. <https://doi.org/10.1038/nature13962>
- Kosslyn, S. M. (1980). *Image and Mind*. Harvard University Press.
- Kubilius, J., Schrimpf, M., Kar, K., Rajalingham, R., Hong, H., Majaj, N., Issa, E., Bashivan, P., Prescott-Roy, J., & Schmidt, K. (2019). Brain-Like Object Recognition with High-Performing Shallow Recurrent ANNs. *Advances in Neural Information Processing Systems*, *32*, 12805–12816.
- Lake, B. M., Salakhutdinov, R., & Tenenbaum, J. B. (2015). Human-level concept learning through probabilistic program induction. *Science*, *350*(6266), 1332–1338. <https://doi.org/10.1126/science.aab3050>
- Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and Brain Sciences*, *40*, e253. <https://doi.org/10.1017/S0140525X16001837>
- Landau, B., Gleitman, H., & Spelke, E. (1981). Spatial knowledge and geometric representation in a child blind from birth. *Science*, *213*(4513), 1275–1278.
- Le Tensorer, J.-M. (2006). Les cultures acheuléennes et la question de l’émergence de la pensée symbolique chez Homo erectus à partir des données relatives à la forme symétrique et harmonique des bifaces. *Comptes Rendus Palevol*, *5*(1–2), 127–135.
- Leeuwenberg, E. L. (1969). Quantitative specification of information in sequential patterns. *Psychological Review*, *76*(2), 216.
- Leeuwenberg, E. L. (1971). A perceptual coding language for visual and auditory patterns. *The American Journal of Psychology*, 307–349.
- Leyton, M. (1984). Perceptual organization as nested control. *Biological Cybernetics*, *51*(3), 141–153.
- Leyton, M. (2003). *A generative theory of shape* (Vol. 2145). Springer.
- Li, M., & Vitányi, P. (1997). *An introduction to Kolmogorov complexity and its applications*. Springer Heidelberg. <http://cs.ioc.ee/yik/lib/7/Li1.html>
- Long, B., Fan, J., Chai, Z., & Frank, M. C. (2019). *Developmental changes in the ability to draw distinctive features of object categories* [Preprint]. PsyArXiv. <https://doi.org/10.31234/osf.io/8rzku>
- Mathy, F., & Feldman, J. (2012). What’s magic about magic numbers? Chunking and data compression in short-term memory. *Cognition*, *122*(3), 346–362. <https://doi.org/10.1016/j.cognition.2011.11.003>
- McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I., & Moser, M. B. (2006). Path integration and the neural basis of the “cognitive map.” *Nat Rev Neurosci*, *7*(8), 663–678. <https://doi.org/10.1038/nrn1932>
- Muller, R. U., Ranck, J. B., & Taube, J. S. (1996). Head direction cells: Properties and functional significance.

*Current Opinion in Neurobiology*, 6(2), 196–206.

Newcombe, N. S., Sluzenski, J., & Huttenlocher, J. (2005). Preexisting knowledge versus on-line learning: What do young infants really know about spatial location? *Psychol Sci*, 16(3), 222–227.

O’Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Clarendon Press.

Penn, D. C., Holyoak, K. J., & Povinelli, D. J. (2008). Darwin’s mistake: Explaining the discontinuity between human and nonhuman minds. *Behav Brain Sci*, 31(2), 109–130; discussion 130–178. <https://doi.org/10.1017/S0140525X08003543>

Piantadosi, S. T. (2011). *Learning and the language of thought* [Thesis, Massachusetts Institute of Technology]. <https://dspace.mit.edu/handle/1721.1/68423>

Piantadosi, S. T., Tenenbaum, J. B., & Goodman, N. D. (2016). The logical primitives of thought: Empirical foundations for compositional cognitive models. *Psychological Review*, 123(4), 392–424. <https://doi.org/10.1037/a0039980>

Pica, P., Lemer, C., Izard, V., & Dehaene, S. (2004). Exact and approximate arithmetic in an Amazonian indigenous group. *Science*, 306(5695), 499–503.

Pimenta, F., & Tirapicos, L. (2015). Megalithic Cromlechs of Iberia. *Handbook of Archaeoastronomy and Ethnoastronomy*, 1153–1162.

Planton, S., Kerkoerle, T. van, Abbih, L., Maheu, M., Meyniel, F., Sigman, M., Wang, L., Figueira, S., Romano, S., & Dehaene, S. (2021). A theory of memory for binary sequences: Evidence for a mental compression algorithm in humans. *PLOS Computational Biology*, 17(1), e1008598. <https://doi.org/10.1371/journal.pcbi.1008598>

Prewett, P. N., Bardos, A. N., & Naglieri, J. A. (1988). Use of the Matrix Analogies Test-Short Form and the Draw a Person: A quantitative scoring system with learning-disabled and normal students. *Journal of Psychoeducational Assessment*, 6(4), 347–353.

Restle, F. (1970). Theory of serial pattern learning: Structural trees. *Psychological Review*, 77(6), 481–495. <https://doi.org/10.1037/h0029964>

Restle, F. (1973). Serial pattern learning: Higher order transitions. *Journal of Experimental Psychology*, 99(1), 61–69. <https://doi.org/10.1037/h0034751>

Reynolds, C. R., & Hickman, J. A. (2004). *Draw-a-person Intellectual Ability Test for Children, Adolescents and Adults*. Par, Incorporated.

Romano, S., Sigman, M., & Figueira, S. (2013).  $LT^2C^2$ : A language of thought with Turing-computable Kolmogorov complexity. *Papers in Physics*, 5, 050001.

Roumi, F. A., Marti, S., Wang, L., Amalric, M., & Dehaene, S. (2021). Mental compression of spatial sequences in human working memory using numerical and geometrical primitives. *Neuron*, 0(0). <https://doi.org/10.1016/j.neuron.2021.06.009>

Sablé-Meyer, M., Fagot, J., Caparos, S., Kerkoerle, T. van, Amalric, M., & Dehaene, S. (2021). Sensitivity to geometric shape regularity in humans and baboons: A putative signature of human singularity. *Proceedings of the National Academy of Sciences*, 118(16). <https://doi.org/10.1073/pnas.2023123118>

Saito, A., Hayashi, M., Takeshita, H., & Matsuzawa, T. (2014). The Origin of Representational Drawing: A Comparison of Human Children and Chimpanzees. *Child Development*, n/a-n/a. <https://doi.org/10.1111/cdev.12319>

Schrumpf, M., Kubilius, J., Hong, H., Majaj, N. J., Rajalingham, R., Issa, E. B., Kar, K., Bashivan, P., Prescott-Roy, J., Geiger, F., Schmidt, K., Yamins, D. L. K., & DiCarlo, J. J. (2018). *Brain-Score: Which Artificial Neural Network for Object Recognition is most Brain-Like?* [Preprint]. Neuroscience. <https://doi.org/10.1101/407007>

Schrumpf, M., Kubilius, J., Lee, M. J., Murty, N. A. R., Ajemian, R., & DiCarlo, J. J. (2020). Integrative Benchmarking to Advance Neurally Mechanistic Models of Human Intelligence. *Neuron*, 108(3), 413–423. <https://doi.org/10.1016/j.neuron.2020.07.040>

Shepard, R. N., & Cooper, L. A. (1982). *Mental images and their transformations*. MIT Press.

Siegler, R. S., Thompson, C. A., & Schneider, M. (2011). An integrated theory of whole number and fractions development. *Cognitive Psychology*, 62(4), 273–296. <https://doi.org/10.1016/j.cogpsych.2011.03.001>

Simon, H. A. (1972). Complexity and the representation of patterned sequences of symbols. *Psychological Review*, 79(5), 369.

Tanaka, M., Tomonaga, M., & Matsuzawa, T. (2003). Finger drawing by infant chimpanzees (Pan troglodytes). *Animal Cognition*, 6(4), 245–251. <https://doi.org/10.1007/s10071-003-0198-3>

Ullman, S. (1984). Visual routines. *Cognition*, 18, 97–159.

Van der Waerden, B. L. (2012). *Geometry and algebra in ancient civilizations*. Springer Science & Business

Media.

Wang, L., Amalric, M., Fang, W., Jiang, X., Pallier, C.,  
Figueira, S., Sigman, M., & Dehaene, S. (2019). Rep-

resentation of spatial sequences using nested rules in  
human prefrontal cortex. *NeuroImage*, *186*, 245–255.  
<https://doi.org/10.1016/j.neuroimage.2018.10.061>