CogMaster's M2 Internship Memoir
Visual Sequence Primitives in Humans

Mathias Sablé Meyer under the supervision of Stanislas Dehaene

University year 2016-2017

Abstract

Objects such as points, infinitely thin lines or spirals emerge across history and cultures.

Our hypothesis is that understanding geometry goes beyond visual perception. Humans are able to understand deep structures through mere visual stimulation and without being prompted to do so, much like they do with language. Subsequently, by understanding not only the surface but also the structure of visual stimulation, humans gain the ability to extrapolate, generalise, compress or classify efficiently information.

We designed a new paradigm where subject are presented with structured sequences of dots anywhere on a plan and asked to detect rule violations if any. The nature of the structures allows for wide generalisations in terms of visual stimulation — rotations, dilations, translations — that are orthogonal to the task.

Our first experiment shows that there is a fondamental difference between order violation and spatial violation within the sequence. The second experiment highlights subjects' accurate ability to implicitely use abstract rules in discriminating sequences. Our third experiment ranks abstract rules by complexity and our fourth experiment confirms and improves the findings of the third experiment.

This research was par of my internship with Stanislas Dehaene in NeuroSpin for my CogMaster's Master, 2nd year.

Contents

CogMaster's requirement

Originality of this work

This work is both a generalisation of previous work on a larger geometrical scale and a comprehensive study of the possible primitives available for human geometry. Our questions are not focused on an exhaustive study of visual the areas: they target concepts available at a higher level and can be studied mostly behaviourally, although brain imaging could be latter used to gain more hindsight.

This was a studied topic in the lab when I arrived: I pursued and extended the reach of the work that has already been done. To this purpose I designed several experiments targeting specific hypothesis, trying to answer new questions related specifically to the nature and variety of the available rules in the domain of geometry. While most of the analysis for the main corpus of my work are fairly standard, the targeted questions are new, as are the experiments themselves.

I also started another — more theoretical — project that will take more time and effort to accomplish: it aims at program induction for a language of though in the domain of shapes, and both the ideas and the methods are new.

Contribution

- The original questions were raised and specified by Marie Amalric, Timo Van Kerkoerle, Stanislas Dehaene and myself.
- The bibliographic research as well as the initial methodology choices were conducted by Timo Van Kerkoerle and Stanislas Dehaene.
- The first experiment was designed and programmed by Timo Van Kerkoerle and run by Marie Amalric on French Children, by Pierre Pica on Mundurukus and by myself on French Adults.
- The following experiments were designed by me with the help of both Timo Van Kerkoerle and Stanislas Dehaene, programmed, run and analysed by myself. Subject were recruited by me through the RISC platform.
- The writing of this present document, including all the figures and the data, is from me, latter corrected by Stanislas Dehaene.
- In appendix you will find two other projects I worked on during my time in the lab. Here's a breakdown of the contributions:
  - The Language Of Shape project originates from exchanges between Marie Amalric, Stanislas Dehaene and myself. The latter design was a common work of Stanislas Dehaene and myself. The implementation, testing and publishing is my work with advice from Stanislas Dehaene.
  - My contribution to the primate work in the lab was discussed with Timo Van Kerkoerle and Béchir Jarraya for the supervision and directions and Edouard Chazel and Julien De Sloovere for more practical questions. Now Maxime Roustan is pursuing this work in the lab.

Understanding geometrical sequences

Humans are able to develop abstract notion that go beyond what is physically possible to perceive. For example the notion of point as defined by Euclid in the *Elements*, "$\Sigma\eta\mu\varepsilon\iota o\nu$ $\varepsilon\sigma\tau\iota\nu$, $o\upsilon$ $\mu\varepsilon\rho o\varsigma$ $o\upsilon\theta\varepsilon\nu$", which translates to "A point is that which has no part", seems to defy accounts of geometry based on perception[1]. Likewise, were found non figurative paintings alongside figurative ones in prehistoric drawings, trigonometry emerged independently in several places at different periods, almost perfect spheres were carved as early as 2 Million years ago, body paintings share abstract similarities across aboriginal cultures, and so on. Yet these perfect abstractions evade visual perception by their very nature.

Furthermore, humans' ability to integrate, compress and process information in various modalities is striking. Upon seeing complex visual patterns one is able, when it makes sense, to extract some underlying structure that goes beyond pure visual or perceptual learning, allowing one to integrate easily certain complex patterns.

This is then reflected in language, as expression such as "a spiral of squares", "a line of dots" or "a grid of squares where one was swapped for a triangle" intuitively lead a subject to recognise — or draw — a corresponding, *complex*, geometrical figure.

We know that, through language, humans are able to manipulate abstract rules applied to object of symbolic nature [1]. Likewise, it has been shown (see [2]) that predictions arise at various level for auditory sequences for either local or global deviants. We therefore want to know how much is predicted by a subject when stimuli come in the form of visual sequences as well what kind of predictions a subject is able to make base on various geometrical regularities.

Michael Leyton [4] [5] argues that geometry can be perceived through nested structures of control that maximise the re-usability of elements of a given level of embedding at higher levels. This is very similar to Fodor's "Language of Thought" [6] where complex representations arise from small sets of atomic symbols, the *lexicon*, together with combinatorial rules. See Figure 1 for an example of such possible compression with simple geometrical shapes organised in to form a bigger organised complex shape.

These abstract rules would precede and encompass more subtle distinctions in the core cognitive systems for representing shapes such as the one studied by Spelke *et. al* in [7], [8]. Preceding the small-scale/large-scale distinction within core knowledge, the ability to generate complex structures based on combinatorial rules would explain one's ability to structure the world.

A recent article from the lab [9] studied this phenomenon with the hypothesis of a "language of thought" for geometry. Such a language would allow subjects to compress sequences into programs with embedded structures and symbols. Within the restrictions of the design it was shown that primitives of symmetries and rotations were spontaneously detected by various population, as well as hierarchical embedding.

In particular, eye tracking data showed that even before the first repetition of a

[1] Definition 1 of the *Elements*. The same goes for the line: "A line is breadthless length.", definition 2, and it follows that "The ends of a line are points.", definition 3.

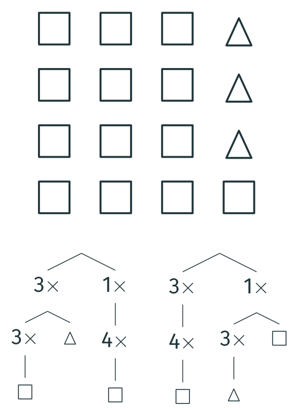

Figure 1: Two compatible compressions for a given complex shape, horizontaly (left) and verticaly (right)

given sequence, subject were anticipating the sequence as it unfolded according to hypothesis compatible with the sequence so far.

The article used a fixed set of position on the vertices of an octogone and restricted all possible sequences positions on this subset of space. Our paradigm is more general in the sense that everything is defined through structure rather than position, allowing stimuli to be rotated, dilated, translated, without changing the hypotheses and the results.

Our goal is to study symmetry, parallelism, right angles and matching lengths as rules for the sequences. For any sequences with underlying regularities and we want to understand how much is picked up by subjects. This will subsequently lead to a better understanding of the nature of abstract geometry as it allows answering questions such as “is parallelism a primitive of a language of thought for geometry” and to look for the emergence of *perfect* concepts such as a point or an infinitely thin line that our senses cannot directly account for.

The goal of the first experiment is to show that the perception of sequences has spontaneous orthogonal components with temporal order on one side and spatial position on the other.

Our second experiments was designed to accurately map subject’s spatial precision in discriminating sequences rule violations.

The next two experiments’ purpose was to disentangle the relative influences of the various rules in the subjects’ ability to distinguish rules from outliers through the comparison of structurally different sequences.

Through this sequence of experiments we refined our questions and dug deeper into the understanding of some specific geometrical properties such as parallelism and right angles. The main topic of this study is to specifically explore this, while my other, more theoretical project during this internship targeted the program induction side of this hypothesis: see appendix for more details.

Definitions and notations

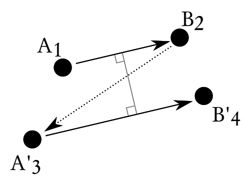

Figure 2: The general structure of sequences

In this study we will explore the perception of geometrical sequences, here defined by a sequence of visual stimuli grouped in time together to form a given geometrical shape.

All experiments used sequences of four dots and simple shapes such as segments and various regular quadrilaterals. The order of a sequence is not defined by the underlying geometry, thus a single quadrilateral can correspond to up to 24 different sequences. To avoid any confusion we will use the same notation everywhere: letters denote spatial position and subscript number represent sequential position.

Figure 2 is an example of a *trapezoid* sequence where it makes sense to group some pair of points which is reflected in the notation: the "A"s thus represents the beginning of a visual cue for a line while the "B"s denote the end, with emphasis on parallelism.

The primes underline the fact that while the corresponding points are conceptually similar and should be members of a similar group, they will appear at different spatial positions.

Once a sequence is defined, there are two possible types of outliers, respectively spatial where a given dot is not at the right geometrical position, or sequential where all dots are at the right location but the order is disturbed. Having both at the same time was never used in our experiments.

In all experiments we measured subjects' ability to accurately detected whenever a sequence, that was already explicitly defined, had an outlier, be it spatial or sequential. This relates strongly to predictive coding as soon as we make the natural hypothesis that once a rule is given, either explicitly or through many repetitions, a participant naturally goes back and forth through periods of acquisition of visual information and periods of prediction of the futur possible outcomes.

Procedure common to all experiments

Several experiments were designed and conducted to address this question and we will present in this document the general progression toward a better understanding. Therefore, all the experiments will be thoroughly described as well as their associated results, and then a more general overview will be given with what was cumulatively learnt.

The experiments shared the common following structure, described in figure 3. All of them had participants facing a computer screen connected to a green-or-red input device. After a description of the task, they were shown different blocks associated to different sequences, with a pause between blocks. Within a block, the structure was the following:

1. The subject was shown a few examples of sequences following the rule
2. Then the test started:
    1. A full sequence was shown on the screen
    2. The subject had to answer
        - green if he believed the sequence to be compatible with the rule
        - red if he believed the sequence to violate the rule
    3. The subject received feedback about the correctness of his answer — not about the nature of the shown sequence
    4. After a random delay, next test started.

At the end of the last block the subject was invited to verbally describe what he could about the task, the nature of the sequences he was displayed, and more specifically about potential strategies. The discussion, if need be, was oriented toward potential cues such as parallelism.

Importantly, the subject is never given explicit shapes such as "the sequence will match the edges of a rectangle", and apart from experiment 3 the subject never sees several points at the same time on the screen.

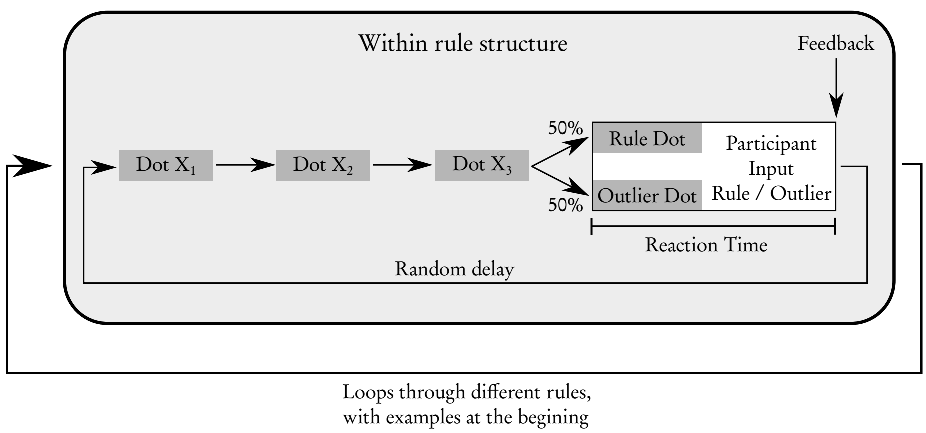

Figure 3: General sctructure of the experiments

Experiments and results

Experiment 1 — Sequential vs. Spatial, several populations

The goal of this initial experiment was to explore how well various populations could detect parallels and parallelism in a visual sequences, and how one population

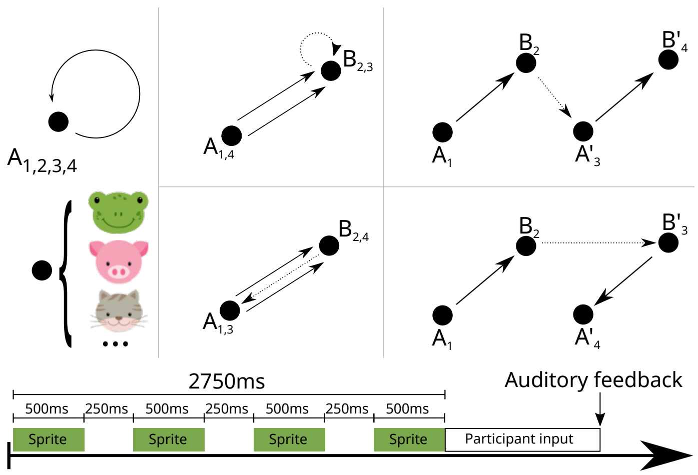

Figure 4: Sequence stimuli used for experiment one. Each cell is a block and its associated sequence structure. For the children to attend the task, dots were replaced — for *all* populations — by animal-like faces sprites. The *point* (top-left) is a control and a training. The four other sequences form a 2×2 design with "Repeat" on the top row, "Mirror" on the bottom row, "Segment" on the left column and "Rectangle" on the right column.

compared to another with regard to different kind of violations — respectively geometrical and sequential as described before.

It was short and relatively easy for a French adult by design, as the relative difference between population and between outliers was the relevant information.

Method

To investigate this, a discrimination task was designed where subject were shown a sequence as rule a few times, then asked whether other sequences were violating the structure of the example. Three populations were tested: French adults, French children and Mundurukus.

In order for the children — entering "CP" thus around 5-6 years old — to understand the task, it was described as a dance that one animal, visually young, had to learn from its ancestors, visually old but of the same specie. The children had to explicitly say whether the child's replication of the dance was correct or not with regard to the examples given by the ancestor: he had to supervise the learning. This explanation was then given for all the populations.

Five cases were studied as shown in figure 4, and we will refer to them as *point*, *segment repeat and mirror* and *rectangle repeat and mirror*.

We used two kinds of outliers for any given rules: spatial, where the visual position of an element was disturbed, or sequential where the order of appearance of the elements was disturbed. Spatial outliers always occurred on the last dot.

The *point* sequence has no sequential outliers as the four positions are visually identical, the spatial outlier was defined as the fourth point being on a different position as the three others on the screen.

The two *segment* were each other's outliers as exchanging points 3 and 4 goes from one to the other, and so were the two *rectangle*.

For these two cases the spatial outlier was defined as the last vector $\overrightarrow{A'_3B'_4}$ or $\overrightarrow{B'_3A'_4}$ being rotated by 30°.

The overall task lasted for about fifteen minutes including the explanations, therefore the French adults were also recruited for a follow-up task that will be described in the next part.

Detail of the participating population:

- Adults: N = 27, one subject was removed from the analysis because of a strong strabismus, age from 19 to 52 (mean = 25.3, std = 6.25). They were recruited through the RISC mailing list and wore no glasses — lenses were allowed. Some of them had this experiment first, then experiment 2, and other had this experiment followed with experiment 3. On average, they had 3.67 years of education post bachelor (std = 1.72) with 14 women and 13 men.

- Children: N = 13, entering CP thus aged 5-6. Note that these are partial results as the task will be run again soon to have a wider population to study.

- Mundurukus: N = 37, age from 11 to 65 (mean = 38, std = 18.6). This population is very inhomogeneous in terms of education background, age and geographical origin but the data was not subdivided any further. The experiment was run by Pierre Pica.

All three populations were tested on the very same computer with the same program, although the feedback sounds for the Mundurukus were recordings of Munduruku's words for "*good*" and "*not so good*" while sounds with pitch going up (resp. down) were used with the French population with the good (resp. bad) meaning.

We were expecting the adults to perform well with this task, and we were interested in the detection rates of the different shapes — more specifically how it compared to the detection rate of the sequential outliers.

In adults, we expected the sequential outliers to be trivially ruled out, while the geometrical one would be slightly harder — still relatively easy for this angle but they can be made difficult at will (see Experiment 2).

In children, we expected a conflict between the attention on spatial position and sequential order leading to the collapse of one in favour of the other, while for Mundurukus the results were harder to predict: the perceptual and cognitive abilities on such a task were expected to match French adults but the lack of training with exercises of this nature could result in the same kind of conflict as in the children.

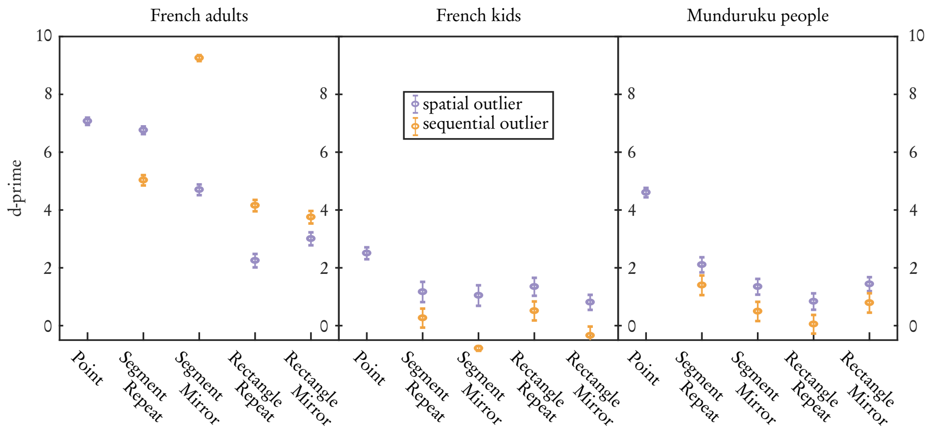

Figure 5: Average d-prime for all populations and conditions, $\pm$ standard-error.

Analysis & Results

The point was a control where the sequential outliers were nonexistent and the position change was just a similar/different experiment on a single position on screen. Unsurprisingly the results are good for all populations with d-primes over 2. It was always the first to be presented.

For the other conditions the results are striking: the results of the children and the Mundurukus follow a similar trend where the spatial outlier is harder than the temporal one *in every condition* with a difference of the order of magnitude of one d-prime.

Conversely, in adults, the sequential outliers are easier to detect that the spatial one.

Here is a breakdown of the relevant statistical points.

- French Adults
  - The performance within types of sequences are not statistically different but they are between groups, $p < 0.005$ for all pairs of conditions (point, segment, rectangle), tested using a Wilcoxon signed-rank test.
  - The Wilcoxon signed-rank test over subject for the different outliers are all significant, $p < 0.005$ for all pairs (rule, angle and sequence)

Thus, for French adults all conditions can be distinguished just through the results hence they are sensible to all cues and outliers independently.

- French Children
  - The Wilcoxon signed-rank test over subject within segment and within rectangle are not significant, but one group against the other is ($p = 0.0281$)

    - The Wilcoxon signed-rank test p-values over subject for the different outliers are the following:
        * Significant for angle against rule ($p = 0.0046$) and angle against sequence ($p = 0.0073$)
        * Not significant for sequence vs rule ($p = 0.74$)

This can be interpreted in the following way: they can detect geometrical violation but are not accurate when it comes to sequence order. This also explains why a sequence and its mirror are not significantly different in the success rate, but they can still discriminate one sequence against a completely different one as expected.

- Mundurukus
    - The Wilcoxon signed-rank test over subject for performance between rectangles is not significant ($p = 0.08$)
    - The Wilcoxon signed-rank test over subject for performance between segments is significant ($p = 0.0051$), as is the Wilcoxon signed-rank test between rectangles and segments ($p = 0.00024$)
    - The Wilcoxon signed-rank test over subject for the different outliers are all significant with $p < 0.005$ for all pairs (rule, angle and sequence)

What we learn from this is that albeit the performances of the Mundurukus look closer those of French children than to French adults, they have a better understanding of the task and especially of the sequential violation.

Discussion

Stepping back, these analyses tell us that all populations can do the discrimination task on the spatial outliers. Furthermore, it appears to be difficult for both the children and the Mundurukus to discriminate the sequential outliers from the rules in this context. This may be due to several factors: our main hypothesis is that, as the task was described as a dance unfolding on a screen, the accent was put on the spatial location rather than the sequential one — which is reinforced by the fact that during the first block, the *point*, **no** sequential outlier could occur — prompting the subject to attend one over the other, despite the negative feedback.

Because in this experiment the underlying symbolic rule was parallelism this conclusion is highly relevant to our main question: all the studied populations could detect a violation in their expectations about an abstract, non explicit rule, that involved being able to *translate* along a vector.

The absence of conclusion on the sequential outliers could lead to several interpretations: it may be that cognitive attendance during the task was focused too strongly on the spatial domain by the nature of the task hence errors in the sequential domain. The constant feedback seems to exclude a misinterpretation of the task — while it happened that French adults made an error on the sequential outlier during the training, it was usually made only once as the negative feedback was automatically interpreted correctly and integrated for the rest of the experiment.

This task was simple and informative but did not allow us to explore the spatial precision of humans' predictions about the last point of the sequence. This is the goal of our next experiment.

Experiment 2 — Detailed psychophysics of the rectangle

Our next hypothesis to test was that the predictive coding was approximate and therefore subject to confusion: we could make the same task as hard as we wanted by changing the distance between the correct position and the position of appearance of the outlier.

Our other hypothesis was that the predicted position would be well-defined in certain directions, for example if moving along these directions broke an important property such as parallelism, but less accurate in others where no strong cues existed.

Method

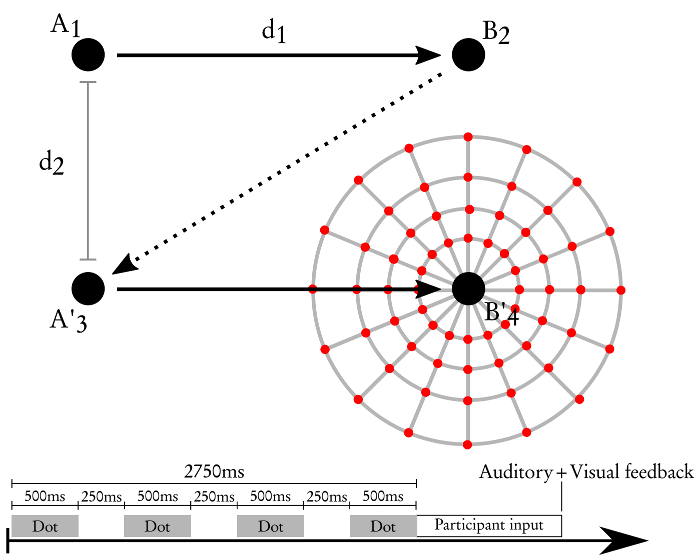

Figure 6: Experiment 2 — Method

To explore this we focused on a small subset of the previous experiment where only the rectangular sequence was explored. There were no sequential outliers — i.e. the order remained constant —, and the last point was half the time correct and half the time on a nearby position.

16×4 possible outlier positions were defined along four concentric circles and 16 radial axis, with 10 repetitions per possible outlier, leading to a total of 1280 data points per subject. As many data points were collected per subject we limited the study to a small population of 10 French adults, recruited through the RISC mailing list for two hours. All of them were tested on experiment 1 first, and then were described experiment 2 as a longer and more difficult run of a specific case of experiment 1.

The general orientation of the rectangle was completely random. The position on the screen was randomly distributed in the following way: a grid of nine positions equally distributed and centered on the screen was defined, and the first dot of each sequence was displayed on one of these positions[2].

The distances $d1$ was randomly drawn between 150 and 200 pixels. We wanted the *aspect* of the underlying rectangle to stay constant thus $d2$ was always bigger than $d1$: it was randomly drawn between $d1$ and 200 pixels.

The concentric rings were centered on the last position and their radius were respectively of 15%, 30%, 45% and 60% of d_1[3]

Between two sequences there was 1500ms and the reaction time was recorded from the apparition of the fourth dot as participants could answer from this time point on.

After each sequence the subject could answer either that the last point was correct according to the explicit rule or not, and received auditory feedback that would be either up or down going pitch, explained during the training, and a visual feedback with one or two points on screen: the correct one and, in the case of an outlier, the outlier one.

We expected the d-prime to increase as we went further away from the last correct point, meaning that it was indeed accurately predicted.

We also hypothesise that the global orientation, the position and to some extent the size has no effect on the accuracy : if our more general hypothesis is right then this task's strategy should be more conceptual than visual. Although some visual cues remain in the world outside of the task — horizontal and vertical line of the screen for example — it should not significantly impact the success rate.

A second level effect hypothesis was that the pattern of error would not be rotation invariant: strong cues were given mainly along two axes, respectively $\overrightarrow{A_1 B_2}$ and $\overrightarrow{A_1 A_1'}$ allowing the subject to fully specify the last position, but outliers breaking both those cues should be easier to detect than outliers along one of them.

An alternative hypothesis would be that outliers visually closer to the previously shown point, or to the barycenter of the three previous points, would be more easily recognised as such and discriminated. This would mean that to some extent low level perceptual cues are easier to detect than abstract properties in this context.

Analysis & Results

The first result is provided by the error map shown in Figure 7: the further away from the correct position, the higher the d'. For pair of adjacent circles we can already see that the accuracy is statistically different: indeed the three student tests across subjects for pairs of adjacent circles leads to difference with p-values in the

[2]For following experiments the designed was improved as it was the *last*, and not the *first*, point of the sequence that was displayed on one of these possible positions

[3]This is already a hypothesis that needs to be tested separately, and it would have been better to have it a function of both d_1 and d_2.

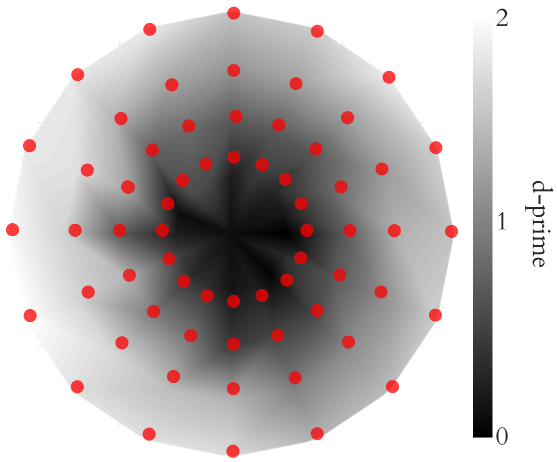

Figure 7: Experiment 2 — Error Map, d-prime computed on the various points and then interpolated for the whole 2-D surface. It is null for the central point, that is to say that there is no possible accurate detection of outliers too close to the actual position — by definition— then gradually increases as we go further away from this.

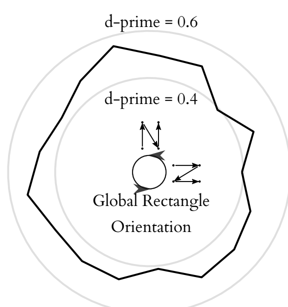

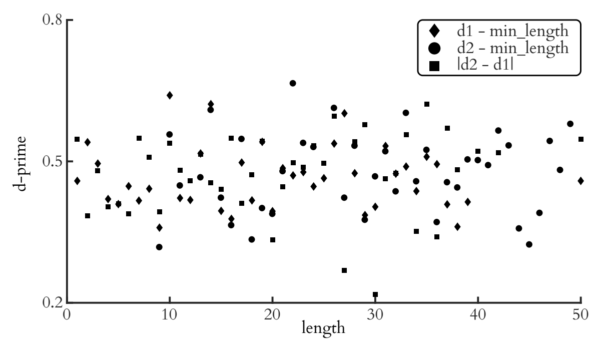

Figure 8: Left pannel: influence of the orientation of the sequence on the d-prime value, mean across subjects, angle from 0 to 360 averaged with a step size of 20°. Right pannel: influence of the global size of the sequence on the d-prime, averaged across subjects, step-size of 1 pixel.

order of magnitude of e-4 — respectively, from the inner to the outer comparisons: 5.3e-4, 5.8e-4 and 6.2e-4.

Another way to test our first hypothesis is to look at a one way ANOVA across concentric circles over subjects, see ANOVA's Table 1.

Table 1: One way, between subjects ANOVA to compare the effect of distance to correct location on subjects' d-primes. There was a significant effect of the distance to the correct location for the four conditions [$F(3, 32) = 12.1, p < 0.005$].

Source	SS	df	MS	F	Prob>F
Columns	10.7	3	3.57	12.1	1.9e-05
Error	9.48	32	0.30		
Total	20.2	35			

Therefore, we know that subject accurately predict the position of the last dot when the sequence unfolds and the difficulty of the task directly relates to the proximity of the outlier.

We can also check whether the global orientation, as well as the various lengths, has a significant effect on the success rate of participants. We would hope that his is not the case as our hypothesis is that the task relies on structural properties rather than visual ones.

For the lengths the correlations are not statistically significant, i.e. the influence of the relative sizes to the ability to solve the task is irrelevant. The same goes for the overall orientation.

To test for this effect the orientation and the length were discretized and we tried to fit an appropriate model: the correlation were poor and the p-values against a constant model were high.

The model for length was linear with while the model for the orientation was periodic of period $\frac{\pi}{4}$ to account for visual preference for vertical and horizontal lines [10].

Table 2: Linear model against constant hypothesis for the lengths, sinusoidal model against constant hypothesis for the orientation (model of the form $y \sim x{*}sin(x){+}b{*}cos(x){+}c$. This is enough thanks to the classic linear equality $a{\cdot}cos(x){+}b{\cdot}sin(x) = c \cdot sin(x + \varphi)$ for well-chosen c and φ only function of a and b).

	Step size	F-value	p-value
Orientation	20°	0.323	0.729
$d1$ length	5 pixels	1.46	0.261
$d2$ length	5 pixels	0.122	0.736

The same goes for the different position on screen that appears to be orthogonal to the performance as we hypothesized, see ANOVA's Table 3.

Table 3: One way, between subject ANOVA to compare the effect of screen position of the stimuli on subjects' d-primes. There was no significant effect [$F(8, 72) = 0.224, p = 0.985$].

Source	SS	df	MS	F	Prob>F
Columns	0.11	8	0.015	0.224	0.985
Error	4.79	72	0.066		
Total	4.91	80			

The higher level result is to look for more specific pattern of error to see whether this gives us information about how subject *compute* their prediction. This will be done by comparing the 16 possible radial positions, and it appears that this appears not to be significant either, see ANOVA's Table 4

Table 4: One way, between subjects ANOVA to compare the effect of outlier position on subjects' d-primes. There was no significant effect [$F(15, 128) = 0.777, p = 0.701$].

Source	SS	df	MS	F	Prob>F
Columns	3.03	15	0.20	0.777	0.701
Error	33.32	128	0.26		
Total	36.35	143			

This second result will be replicated in a latter experiments and the results will concur with this conclusion.

Discussion

Subjects are accurate in their predictions and can address this task, and there is an error gradient around the correct position. Due to the nature of this task and the rigid structure of rectangular sequences, this is not enough to disentangle various rules as all were present in the rectangle, but subjects can accurately predict positions for a given structure through rotation, dilation and translation. This results is extremely important for the following experiments.

Further explorations are possible with this data, for example by comparing outliers along main axis to other outliers, or by comparing outliers close to the barycenter of the sequence to others. We chose to run more specific experiments targeting exactly these questions rather than to dissect the data into too many sub-cases.

Experiment 3 — pining down the relevant rules

From previous experiments we know that subject are able to use geometrical cues to solve a sequence discrimination task. What we want to do now is to explore systematically which cues are available, and how extensively they can be used.

Method

Our goal for the third experiment was to have a $2\times2\times2$ design where the relevant variables were geometrical cues, respectively parallelism, right angles and symmetry. One case is absent as a sequence with a single right angle would be such that the right angle would give no usable information to the subject about the remaining point.

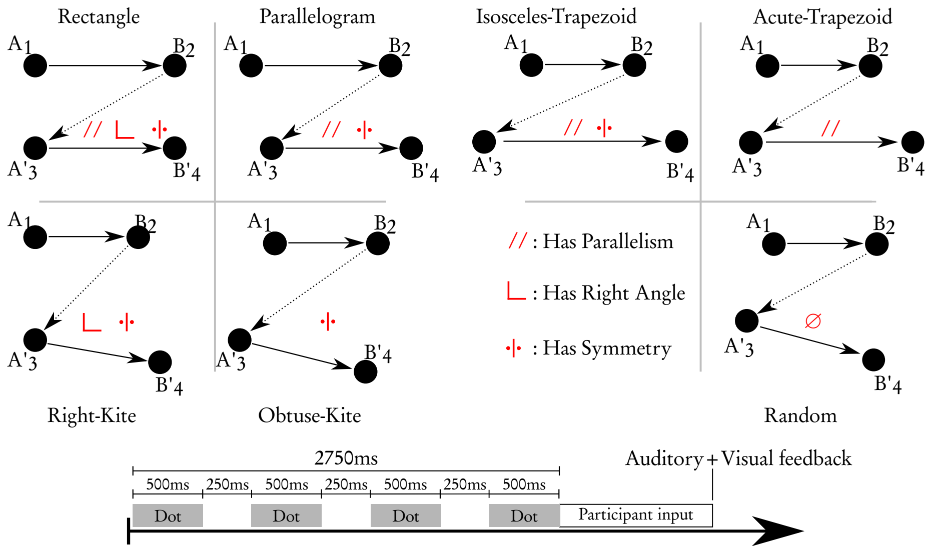

Figure 9: The set of shapes that we chose for experiment 3: the various relevant cues are specified with the red symbols.

We decided upon the following set of shapes, represented in figure 9, where names in bold indicate fully determined sequences[4]:

- **Rectangle**: has all the relevant cues: four right angles, two pairs of parallel lines, two axial symmetry and a point symmetry. This is also our reference shape as we already have experiment 2 with extensive study of this pattern.
- **Parallelogram** misses the right angles but still has a point symmetry and two pairs of parallel lines.
- Right-Kite has two pairs of segments of equal lengths that are connected, respectively in A_1 and B_4' as well as a right angle in A_1. It has a single axial symmetry along A_1B_4'
- Kite is similar to Right-Kite but lacks the right angle.
- **Isosceles Trapezoid** is mainly defined by its two parallel segments A_1B_2 and $A_3'B_4'$, but it also has an axial symmetry along the bisector of A_1B_2 making A_1A_3' and B_2B_4' of equal length. It is fully determined though harder to build.
- Acute Trapezoid lacks the symmetry but is otherwise similar to the Isosceles Trapezoid

[4]that is to say sequences where the information about the three first dots is enough to mathematically predict the fourth one without additional information about length, angle or position.

- Random was designed to have no geometrical property that could help a subject, thus having to rely solely on perceptual learning.

Outliers could appear on six possible positions, all equidistant from the *correct* position, respectively along three main axes: $\overrightarrow{A_1B'_4}$, $\overrightarrow{A'_3B'_4}$ and $\overrightarrow{B_2B'_4}$, in order to have a second level analysis of the pattern of error.

This time the shapes were fixed in the sense that the ratio of lengths were kept constant through blocks — they could still be dilated.

In order to have them visually similar we designed them in such a way that the smallest circle comprising all the points of a given sequence was of same diameter for all cases.

Both the orientation and the position on the screen was random. The block were displayed in a random order with one block per sequence and with a 50% chance of outlier. Participants received auditory feedback after each trial and the timing was similar to experiment 2: 500ms with dots on screen and 250ms between dots.

N = 18 subject were recruited from the RISC mailing list with a mean age of 14.6 (std = 3.0), 11 women and 7 men, with on average 3.8 years of education post bachelor (std = 1.9).

Participants were trained on the task through experiment 1 beforehand and were not told about the specificities of the different sequences, but were described the task including the random orientation and position. They received auditory feedback in the form of a good/error sound, as well as a visual feedback shown the four *correct* dots and if need be the outlier one[5].

Our hypothesis here is that performances will differ from one sequence to the other with the ordering on performances resembling the order on geometrical constraints as subject use these to solve the task, the alternative being that the performances are identical, meaning that solely visual rote learning was at play.

Additionally, we hypothesise that a finer analysis should exhibit patterns in the errors at the level of shapes, where certain shapes sharing a property will have a given pattern of outlier error, the alternative being that all shapes share a common pattern that rely only on visual perception — proximity for example.

Analysis & Results

The first observation from Figure 10 is that people can still solve the task in the absence of any cue — see results for the *random* sequence — which means that a certain amount of the predictive coding relies solely on perceptual learning.

At the same time d-primes value differ between shapes — see ANOVA Table 5 for the associated test — meaning that for a given sequence, the underlying geometrical structure has a strong influence on the subject's ability to address this task.

[5]This may be a design error and was removed from experiment 4. Indeed, instead of always seing only the sequences, in this version they could see the whole shape at once, and while it indeed helped them to solve the task it could change the way they address the task

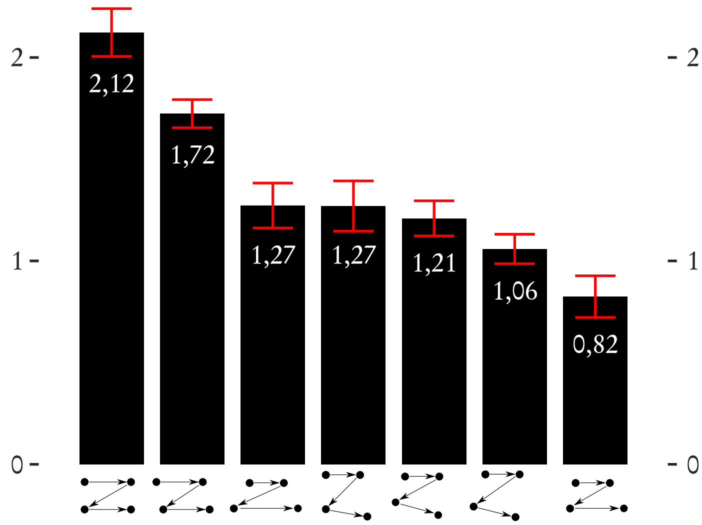

Figure 10: Mean d-primes across subjects for the different sequences, $\pm$ standard error after removal of the mean.

Table 5: One way, between subjects ANOVA to compare the effect of the sequence on subjects' d-primes. There is a significant effect $[F(6, 112) = 6.26, p < 0.005]$.

Source	SS	df	MS	F	Prob>F
Columns	19.2	6	3.20	6.26	1.1e-05
Error	57.2	112	0.51		
Total	76.5	118			

Discussion

An issue with this experiment, that hints at a control for stimuli not rigorous enough, is that while *random* is in all respect the hardest sequence to structure, it is not true here that it was the most difficult for subjects: two other sequences proven to be harder. Likewise, two other sequences were of similar difficulty. While this could mean that the structure is not as significant as simpler visual properties, we realised through this result that the way the lengths were adjusted was not visually constraining enough.

Indeed, a very flat kite that would share its longest diagonal length with a square would be much easier to recognise than the latter as B_2 and B'_4 would be almost identical — while both would be similar according to our similarity constraint as defined earlier. This is the main issue we wanted to address with our next experiment.

Because of this issue we moved directly to the next experiment, and as the results on the outliers were not significant on the improved experiment we postpone studying the effect of the shape on the pattern of error in this experiment.

Experiment 4 — improving on experiment 3

We concluded from the previous experiments that some rules could be used to generalise sequence recognition and prediction, but only several cues were used at the same time, making it difficult to disentangle them. While expriment 3 tried to answer this question, some design flaws made it hard to interpret as such.

We pursued our systematic work from experiment 3 with more shapes that had a better design and were matched more accurately.

Experiment 3 and 4 are very similar in terms goal and design, but experiment 4 had a better, more explicit design.

Method

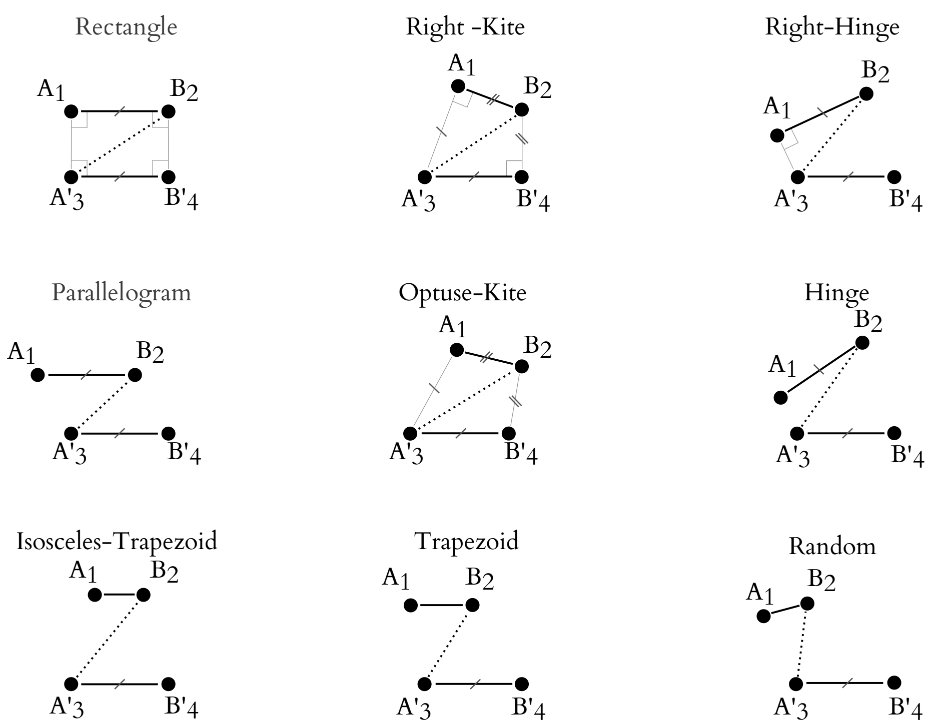

Figure 11: The set of shapes we defined

We defined a set of nine relevant shapes that had various relevant properties and asked the participants to do a similar task for the different shapes. Figure 11 shows the exact representation of the sequences as they were shown and the table below breaks down the different targeted properties:

Here's a detailed list of the shapes:

- The rectangle is both our baseline thanks to Experiment 2 and the richest shape of our list: it has all relevant properties, is highly regular and should

Table 6: Parallelism is true whenever two sides at least are parallel. Relevant lengths gives info about how many visual cues were given, in terms of length, to specify the final position. Right angle counts the number of right angles, and symmetry (n, m) counts the number of axial symmetry n, and the number of point symmetry m. The elements are ranked by d-prime on the experiment.

Exp. 4	Parall.	Relevant lengths	R.-Angle	Symmetry
Rectangle	yes, 2	yes, 2	yes, 4	yes, (2,1)
Parallelogram	yes, 2	yes, 2	no	yes, (0,1)
Trapezoid	yes	no	no	no
Obtuse-Kite	no	yes, 2	no	yes, (1,0)
Hinge	no	yes, 1	yes, 1	no
Isosceles-Trapezoid	yes	yes, 1	no	yes, (1,0)
Rusted-Hinge	no	yes, 1	no	no
Right-Kite	no	yes, 2	yes, 2	yes, (1,0)
Random	no	no	no	no

yield the highest results

- The Parallelogram just remove a single constraint from the rectangle, while breaking some symmetries, it will allow us to study the atomic importance of right angle in this context
- The Trapezoid is the purest design using parallel vectors: it has no other relevant property and can only be solved using this. The isosceles version adds some visual symmetry as well as a usable cue to solve the task.
- The Kite is geometrically very close to the rectangle: one way to construct it is to start from a rectangle and swap two non-identical segments. Because of this property, all results about lengths are preserved, as are right angle if need be, but the parallelism is absent
- The hinges is the length equivalent of the trapezoids: they have a single pair of equal length to solve the task, and one additionally has a right angle, but no other property can be used.

This list was chosen as it has a lot of what we will refer to as *critical pairs*, that is to say pairs of shapes that only differ by one atomic property, which allows us to study the effect of this property in isolation but in various contexts. Of course many more shapes could have been added but this is a concise compromise we can test subjects on. Here are a few examples of critical pairs:

- Parallelism:
  - Trapezoid and Random
  - Isosceles Trapezoid and Optuse Kite
  - Parallelogram and Optuse Kite
  - Rectangle and Right Kite
- Right angle:

- Rectangle and Parallelogram
  - Hinge and Right Hinge
  - Optuse Kite and Right Hinge

- Symmetry
  - Obtuse Kite and Rusted Hinge
  - Hinge and Right Kite

- Pairs of equal length
  - Rusted Hinge and Random

Because these constraints left at least two degree of freedom on all shapes we could match two visual properties to prevent external factor to come into play and to correct our mistake of experiment 3. Therefore, the last length $|A'_3B'_4|$ is constant across shapes in order to minimise the last segment different between the different conditions, and the mean of the four sides and the two diagonales is constant across shapes to keep them as close as possible visually.

Outliers were on four possible positions, equidistant from the correct position and either along $A'_3B'_4$ or equidistant from A'_3, see Figure 12 for more details. This was chosen to investigate second level cues in the sequence prediction: a hypothesis is that when parallelism is involved, detecting violation in parallelism is easier than detection violation in the distance.

To avoid biases in the discrimination task, the presented stimulus was half the time the correct sequence and half the time an outlier one with a random outlier.

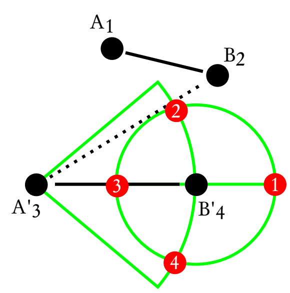

Figure 12: The four possible positions for the outliers: all are equidistant from B'_4, two of them are additionally along $A'_3B'_4$ and two of them are additionally equidistant from A'_3

Sequences' size could vary but the ratio of length was kept constant across trials. Both the orientation and the position on the screen was random. The block were displayed in a random order with one block per shape and 150 trials per block with a 50% chance of outlier. Participants received auditory feedback after each trial and the timing was slightly quicker than the previous experiment : 400ms per point and 200ms between points.

The population for this experiment was comprised 20 participants out of which three were pilots whose results were used to tune the difficulty of the experiment.

Subject were recruited through the RISC mailing list, mean age was 23.1 (std=2.55), 9 women and 11 men, with a mean of 3.44 years of post bachelor education (std=1.5)

Our first hypothesis is that the success rate is a monotonous function of the available cues: the more geometrical cues, the better the results.

As in experiment 2 we can also look for differences between the different outliers, within a sequence. For each sequence the prediction is different but it can be summed up in the following was: outliers that break cues are easier to detect than others.

Analysis & Results

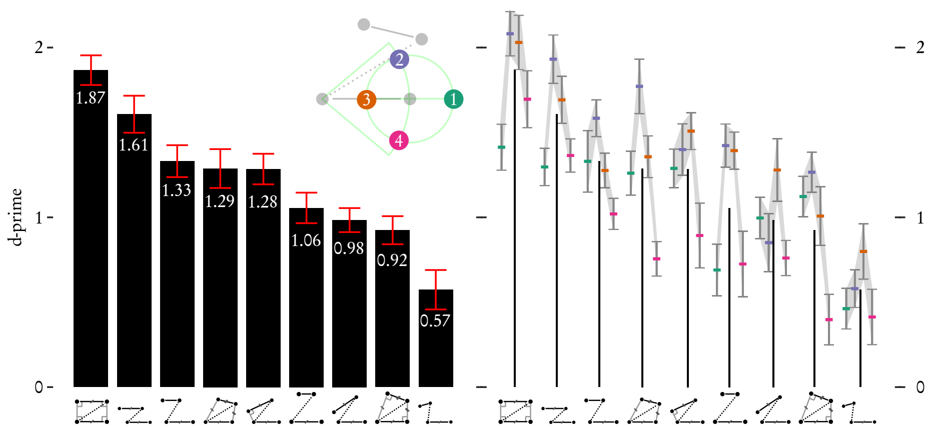

Figure 13: Left pannel: Mean d-prime across subjects ± standard error after removal of the mean, sorted by d-prime value. Right pannel: detail within each shape of the mean d-prime across subjects for each possible outlier.

For this experiment two levels of analysis are relevant, *between sequences* and *within sequences, between outliers*. The first one is our primary goal.

The most important analysis is to know how much sequences were statistically different in the participants' results. A one way ANOVA confirms they are indeed different, see ANOVA's Table 7.

Table 7: One way, between subjects ANOVA to compare the effect of the sequence on subjects' d-primes. There is a significant effect [\$F(8,133) = 3.85, p < 0.005].

Source	SS	df	MS	F	Prob>F
Groups	15.6	8	1.96	3.85	0.00041
Error	67.6	133	0.50		
Total	83.3	141			

Conversely, the difference between the various possible outliers is not statistically significant — this was expected as the predictions differ between sequences hence looking at a significant difference across all sequences shouldn't raise anything particular, see ANOVA's Table 8

Table 8: One way, between subjects ANOVA to compare the effect of the nature of the outliers on subjects' d-primes. All sequences were averaged here. There is no significant effect [$F(3,60) = 2.43, p = 0.0738].

Source	SS	df	MS	F	Prob>F
Columns	2.9	3	0.97	2.43	0.0738
Error	24.1	60	0.40		
Total	27.1	63			

Another intuition emerges from figure 13 (right panel): it would appear that the two outliers closer to the geometrical center of the sequence were easier to detect than the two other, irrelevant of the shape. This can be tested but is not statistically significant: a two sample student test across subjects for the mean of two outliers close to the shape against the two others yields a p value of 0.073.

A more extensive study of the influence of the outliers on the success rates for each sequence did not yield consistent results about our individual hypotheses for each sequence.

Discussion

The main take home message from this experiment is that more geometrical cues leads to a better violation detection. People are able to cumulatively use cues in order to solve the task, and when you remove a given cue they fall back to a simpler prediction, down to the level where pure perceptual learning allows them to still solve the task but with degraded performances.

The next step is to build a model taking into account the whole table above with all the available and to see how much each cue explains the variance in the data. This is ongoing work.

General results and overlook

This sequence of experiments told us that abstract notion were used as tools in order to solve a visual task, as opposed to raw perceptual learning.

More importantly, the notions available to a subject are fondamental in the mathematical sense, even though sensory experiences cannot account for a perception of parallelism or right angle as abstract is defined here.

Let's take a step back and review the general progress we made through these experiments.

Table 9: While all the experiment shared a common paradigm, the designs differed and it is important to keep in mind how so. This table is a summary of the most relevant differences. Experiments with a star were solely run as pilot so far and are described in the next section.

	Exp 1	Exp 2	Exp 3	Exp 4	Exp 5*	Exp 6*
Constant ratio	yes	no	yes	yes	no	no
Few possible orientations	yes	no	no	no	no	no
Had sequential outliers	yes	no	no	no	no	no
Had different shapes	yes	no	yes	yes	no	no
Required fixation	no	no	no	no	yes	yes

A summary of the experiments and their parameters

The conclusions from these experiments

These experiments give us certainty about humans ability to extract and use abstract rules when it comes to geometry, as well as hints as to where to look for in the search for both the nature of these rules and their localisation in the brain.

Most importantly, we confirmed our hypothesis that while pure perception learning is important in a visual sequence perception task, higher level elements are spontaneously picked up and used to understand the profound structure of stimuli of this nature.

Additionally, we know that these cues can be used cumulatively, which hints strongly in the direction of a notion of compression through embedding and a "language of thought" in the visual domain for sequences. Indeed, it appears that it requires fewer efforts to assimilate a sequence whose underlying structure can be efficiently compressed through structure, as opposed to rote learning of visual information such as length or angles.

While we hoped we could find explicit error patterns correlating these conclusions within given sequences, no clear effect was exhibited as of yet and further study need to be run both on the currently available data and on new experiments.

Discussion & Conclusion

Next stepping stones

This research exhibits a notion of abstract rule in the visual understanding of sequences, leading to a vast amount of finer grained research for each possible rules, as well as a brain imaging version of them: for each relevant abstract rule, how much relies only only visual areas, and how much is prefrontal for example.

Toward pupil and MEG

Another experiment was designed to target specifically parallelism, breaking free of several constraints of the previous experiments. Its design does not rely on shapes but on the very notion of two segments being parallel, as opposed to having a fixed angle.

It is possible to have either an explicit recognition version with 50% of the time parallelism and 50% of the time a fixed angle, but we opted for a passive version to look for rule violation: a given block was 70% of a given rule — either parallel, 30° angle or 60° angle — and the 30% other were equally distributed onto one of the two other possibilities.

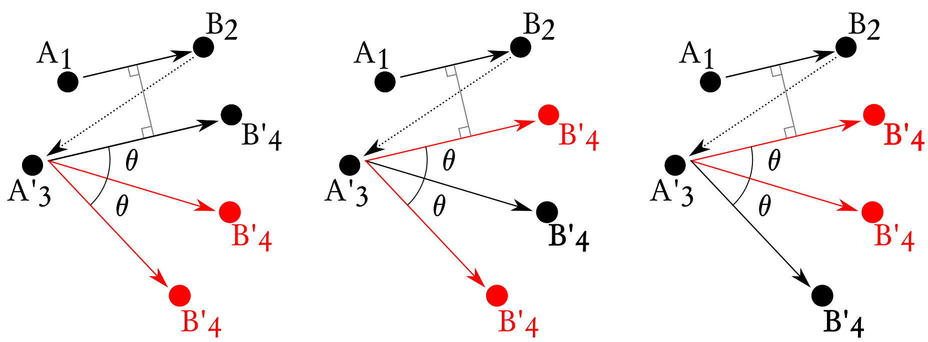

Figure 14: The design next experiment: θ is fixed and adjusted to match a target difficulty, and there are three blocks with each time a single rule and two possible outliers.

Behaviourally, by asking the participant to explicitly discriminate the rule, we can estimate how important parallelism is by comparing it to another fixed rule, but in this case we can also look for mismatch negativities in MEG and to see if the effect is, as we expect:

- Stronger for $2 \cdot \theta$ than for θ when the rule is parallelism
- Stronger for θ when the rule is parallelism than for $2 \cdot \theta$ when the rule is θ
- Stronger for parallel than for $2 \cdot \theta$ when the rule is parallel

This relies on the idea that parallelism is somehow special and is picked up early on and accurately.

This has been successfully piloted using pupillometry as a second hand indicator of surprise, the analysis of the results is ongoing work. If proven to be significant, it will be used as an MEG experiment.

What is not addressed in this study nor directly with the MEG is the location in the brain of the different effects. One interesting future direction would be to design a passive, display only experiment in fMRI based on the current design for pupillometry and MEG, and then to subtract the conditions we want to compare : θ difference between parallel and θ angle versus θ difference between θ angle and

$2 \cdot \theta$ angle for example, to see how much of parallelism we can find in the prefrontal cortex — or more generally outside the visual area.

Explicit predictions

One of the next experiment we are ready to run is very similar to experiment 2 but with a strong *explicit* prediction. Instead of having people give a single bit of information on each trial through a yes-or-no mechanism, we would have them use a touch screen computer to give their best guess about the last point of the sequence.

After rescaling and realigning of the conditions this would give us a direct heat map of the predictions, rather than a deduced one as in experiment 2. Then we could purposefully dilate sequences along various axes and see how the heat map dilates accordingly — if it does.

The main reason this was pushed back was to have data we could compare between primates and humans but things will soon be ready in the lab to run this kind of paradigms on macaques — see next section.

Toward primates

One of the key point of this design was to be able to compare the results with primates: as it only required two possible input it is relatively easy to adapt it for macaques and even to run not only behaviour but also brain imaging experiments. The training has been ongoing for a few month and the first meaningful experiments are approaching.

Our main hypothesis is that non-human primates would be as good as humans, if not better, when it only boils down to perception learning. Conversely, the perception of sequence mirroring, or high level abstract geometrical cues, would not be picked up and their performances would collapse when compared to humans in those cases.

The lab is also preparing a new paradigm to test non-human primates with touch screen computers, which would give us more freedom in the design for example by asking subject to explicitly predict the sequence rather than simply discriminating outliers. My contribution to this work is to be found in appendix.

A more comprehensive language

In the long run, the main achievement will be to propose a language of thought for geometry with both the lexicon and the structures, at which point what we are studying here will be used as a basis and justification for the primitives.

My work in the lab toward this language is to be found in appendix: it is still the babbling of what a comprehensive model will look like but it's also a huge step forward as it offers an explicit and testable model we can work with and improve through predictions — and failures to predict.

Final word — Thanks

I greatly enjoyed my time in the lab, not only thanks to the colossal and constant intellectual stimulation from everyone around but also for how welcoming it was. I take this — albeit small — opportunity to thank people that participated in this, may you continue to welcome newcomers way you did!

References

[1] R. B. Lees and N. Chomsky, "Syntactic structures," *Language*, vol. 33, no. 3 Part 1, pp. 375–408, 1957.

[2] C. Wacongne, E. Labyt, V. van Wassenhove, T. Bekinschtein, L. Naccache, and S. Dehaene, "Evidence for a hierarchy of predictions and prediction errors in human cortex," *Proceedings of the National Academy of Sciences*, vol. 108, no. 51, pp. 20754–20759, 2011.

[3] T. A. Bekinschtein, S. Dehaene, B. Rohaut, F. Tadel, L. Cohen, and L. Naccache, "Neural signature of the conscious processing of auditory regularities," *Proceedings of the National Academy of Sciences*, vol. 106, no. 5, pp. 1672–1677, 2009.

[4] M. Leyton, *A generative theory of shape*, vol. 2145. Springer, 2003.

[5] M. Leyton, "Perceptual organization as nested control," *Biological Cybernetics*, vol. 51, no. 3, pp. 141–153, 1984.

[6] J. A. Fodor, *The language of thought*, vol. 5. Harvard University Press, 1975.

[7] E. Spelke, S. A. Lee, and V. Izard, "Beyond core knowledge: Natural geometry," *Cognitive science*, vol. 34, no. 5, pp. 863–884, 2010.

[8] S. A. Lee and E. S. Spelke, "Two systems of spatial representation underlying navigation," *Experimental Brain Research*, vol. 206, no. 2, pp. 179–188, 2010.

[9] M. Amalric, L. Wang, P. Pica, S. Figueira, M. Sigman, and S. Dehaene, "The language of geometry: Fast comprehension of geometrical primitives and rules in human adults and preschoolers," *PLoS computational biology*, vol. 13, no. 1, p. e1005273, 2017.

[10] S. Appelle, "Perception and discrimination as a function of stimulus orientation: The' oblique effect' in man and animals." *Psychological bulletin*, vol. 78, no. 4, p. 266, 1972.

Appendix

A proposal for a Language of Shape

Rationale

Let's imagine someone drawing a non-figurative shape on a piece of paper, holding a pencil on a piece of paper. We argue that behind the representation of what this person is drawing lies a language with primitives and combinations that is *executed* by this person.

A language is defined by its primitive and its syntax. To identify those components, we started from the literature on prehistoric drawings, mathematics, children's drawing and traditional shapes across cultures. On these bases we adopted a reference set of 6 shapes as the target for the minimal expressiveness of our language, with the idea that our language should be powerful enough to capture those shape with a low complexity program.

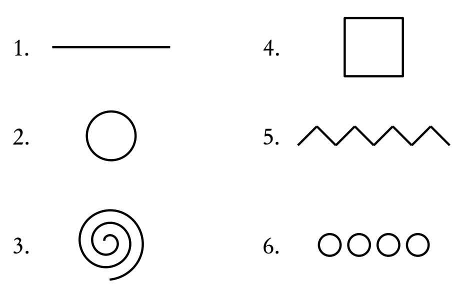

Figure 15: The reference set of shapes

Once we had generated a simple language with as few primitives as possible we checked its ability to generalise to other shapes by exploring the range of programs and the associated shapes.

1. The segment is our most simple target shape, it requires the ability to draw something or as we will see later, to integrate a set of parameters for a fixed, arbitrary duration.
2. The circle should be as simple as the segment, it is somehow the simplest — in terms of invariances and high level properties — closed line, and in our language it's represented with another set of parameters to integrate that defines the curvature of the drawn line.
3. The spiral is a circle that would accelerate along the integration, thus missing the starting point: experiments should be run here to confirm that indeed this is cognitively plausible — are people acceleration along the spires of a spiral?
4. The square is the first shape that requires explicit discrete repetition, thus mixing the infinitesimal repetition of the integration and the explicit repetition for the four sides. It also requires the concatenation, as opposed to the embedding, of two instruction: drawing a segment and turning.

5. The zigzag needs to deal with a repetition, but also with turning in either a direction or another, thus introducing simple arithmetic and variable manipulation to do that concisely — i.e. without explicitly giving all the cases.
6. The line of dots was added to be able to use one shape — the line — as a cue for another — the circle — and thus embedding program seamlessly.

In all the above, we postulate the existence of fixed units specific to each operation as well as a minimal arithmetic capable of generating all integers at least. This arithmetic can be later extended at will to include cognitively relevant primitives — 5 as a primitive of measurement for exemple.

Here is the proposed code for those six reference shapes:

1. Segment

```
Integrate
```

2. Circle

```
Integrate(angularSpeed=unit)
```

3. Spiral

```
Integrate(accel=unit, angularSpeed=unit)
```

4. Square

```
Repeat(Double(Double(unit))) {
  Integrate ;
  Turn
}
```

5. Zigzag

```
alpha = unit ;
Repeat(indefinite) {
  Integrate ;
  Turn(angle=alpha) ;
  alpha = Opposite(alpha)
}
```

6. Line of dots

```
Repeat(Double(Double(unit))) {
  Integrate(t=Half(unit),pen=off) ;
  Embed {
    Integrate(angularSpeed=unit)
  }
}
```

Syntax

We propose the following syntax for such a language of shapes:

```
Var     ::= | unit
            | Double(Var) | Half(Var)
            | Next(Var) | Prev(Var)
            | Oppos(Var)
            | Divide(Var,Var)

Program ::= | Program ; Program
            | Turn([angle=Var])
            | Embed { Program }
            | Repeat([Var]) { Program }
            | Integrate([d=Var], [pen={on,off}],
                        [speed=Var],
                        [accel=Var],
                        [angularSpeed=Var],
                        [angularAccel=Var])
```

Terms within `[...]` have default values and can be omitted while still being valid with regard to the syntax.

Semantics

The three most important operators are the `Repeat` and the `Integrate` instruction. Here is a detailed breakdown of what every instruction does:

- `;` is just the concatenation of programs: it executes the left-hand side in the current environment and the right-hand side in the environment returned by the left-hand side.
- `Repeat(Var){ Program }` evaluates Var in the current context to a number `n` deterministically we'll explain later, and then executes the following program:

```
Program ;
Program ;
Program ;
...
... [n times]
...
Program ;
```

Note that while the result is syntactically equivalent to an explicitly written version — and because variables can always be manually evaluated as there is no side effect with regard to any form of outside world input — one could malignly decide not to use `Repeat` ever. Of course in many cases it simplifies a program a lot, and when it comes down to costs it will in most cases significantly shrunk the cost of a given shape.

- `Integrate([...])` is the only instruction here that actually draws anything. The idea behind it is to move the pen, on or off the paper, during a given amount of time, with a given set of parameters. The integration *does* refer to an infinitesimal repetition in the formal semantics, although for obvious reasons the operational semantics is an approximation. It takes a few arguments, let's see the detail:

  - `t` is the duration of the integration.
  - `pen` is whether the pen should be touching the paper or not
  - `speed` indicates the speed at which the hand is moving
  - `accel` indicates how the speed should change across time
  - `angularSpeed` indicates at each time point how the direction of the drawing should change — it is therefore already the derivative of an internal value that represents the facing direction, the same way speed represents the derivative of the x/y position.
  - `angularAccel` indicates how much the previous value should change over time

  All these variables are in arbitrary units that were chosen so that the default values lead to simple programs as described in the beginning of the document.

- `Turn(angle=[...])` is an instantaneous change in the facing direction. This allows non-differentiable angles, although it could be seen as a less expensive syntactic sugar for an `Integrate` at null speed.

- `Embed {[...]}` allows one to insert a program within another one seamlessly and without breaking the surrounding environment. More precisely it executes the given program in the current environment but when it returns all modifications are erased. It is at the core of the compositional aspect of the language.

Future — ongoing — work

The cost of a program is straightforwardly defined as the number of instructions in the program.

This leads to a few important questions:

- What do *simple* programs look like? This is easily answered by generating them and having a look at the result. Experiments will hopefully show that they match what subjects would agree to describe as *simple*
- For a given shape, what is the simplest program, and is it the one inferred by humans in all cases? In most cases? This requires to locally solve the backward problem of inferring a program from an output, which is ongoing work through deep learning and program inferences.

- Can we map program complexity to an objective mesure of required attention through brain imaging? This is postponed as it requires the two previous questions to have solid answers in order not to target the wrong language.

An experiment is currently being designed to explore intuitive complexity of given shapes and its correlate to minimal description length in this language, that will compare several populations.

Sandbox

An implementation of the language is available online to experiment with at the following address: www.dptinfo.ens-cachan.fr/ ~ msableme/LoG/.

Moving the experiments to Primates

Initially, a key notion behind these experiments was that we could test them behaviourally on primates. There is a currently existing setup but several researchers from the lab, including Stanislas Dehaene, Timo Van Kerkoerle, Béchir Jarraya and myself wanted to move to a setup using touchscreens and to train them to explicit their predictions when it comes to visual sequences.

I was honoured to be part of this decision process and got to participate in the design of the new setup on several levels, from the practical questions to the task pattern design. I helped in building the automatic reward process, and implemented both electronically and on MATLAB© the system to couple any experiment to a reward dispenser.

Although this was not central in my research internship and is not directly relevant to the results shown in the main body of this report, I will display here some information about these technicalities for they are part of integrating a lab's decision process.

The goal

We needed an automatic reward system to train the subjects, that would give a reward when positive feedback was desired. The interaction took place through a touchscreen and it was important to abstract the reward mechanism from a given experiment, i.e. to create a high level abstracted API that could be plugged into any experiment.

The tools

In order to do so, an Arduino board was coupled to a food dispenser and is controlled through any program on a computer that handles the experiment as well as the touchscreen.

These components were not designed to work together and therefore some engineering was required for the low level handling of the information, so that an end user could simply run a command of the form `give_reward(reward_amount)` and not have to know about the low level process.

Were used:

- An ENV-203IR Pellet Dispenser from Med Associates Inc.
- An Arduino Uno board
- A few hours of designing, welding and coding

The result

A single small controller was produced, that would be plugged to the computer and the dispenser and control the reward flow. I also produced the documentation for the controller, as well as the programing part. The system is currently operational and should be used soon.

From an operational point of view, a single function `give_reward` is implemented, as well as a constructor and a destructor for the associated object. I will not detail here the electronic part of the circuit, nor the programing detail, but a documentation is written and available.

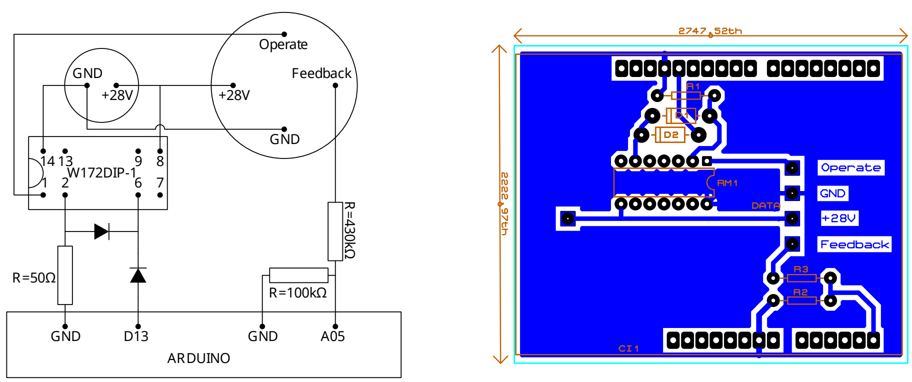

Figure 16: The circuit and its PCB

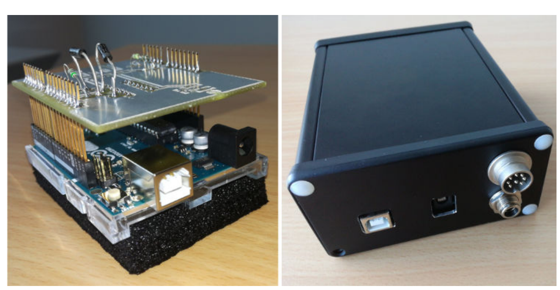

Figure 17: The resulting controller