Blind spots in the mind's eye: Mental imagery often lacks detail and coherence

Andreas Arslan (Arslan_Andreas@phd.ceu.edu)
Department of Cognitive Science, Central European University (CEU), Vienna, Austria

Jonathan F. Kominsky (kominskyj@ceu.edu)
Department of Cognitive Science, Central European University (CEU), Vienna, Austria

Abstract

The workings and products of the imagination are often described in visual terms: We speak of 'mental images' and the 'mind's eye.' To what extent is this metaphorical? Should imagination be conceived of as a process of quasi-perceptual simulation or is it more sparse and abstract? Building on recent findings that suggest mental imagery tends to lack detail, we investigate how complete and coherent people's imagined scenes are. In Experiment 1, we presented a riddle-like vignette to participants and found that they on average only imagined 54% of the simple features we asked them about. Moreover, successfully finding a solution was unrelated to the number of features imagined. In Experiments 2a – 2c we found that participants often did not notice spatial contradictions in text descriptions (2a), even when scaffolded with a map (2b), and that spotting these contradictions was unrelated to performance on a mental rotation task (2c).

Keywords: imagination; mental imagery; event simulation

Introduction

Many of our everyday experiences largely take place in our imagination. We are especially reliant on it whenever we have to make sense of events we do not perceive directly – be it while reading a novel, skimming the news, or listening to a story someone is telling us. The language used to describe the workings and products of imagination often likens it to vision: We talk of mental images, picturing situations in our head, and seeing people or objects in our mind's eye. Are these merely convenient metaphors? Or is to imagine, at least occasionally, to mentally simulate scenes in a format, and with a vividness, that parallels direct experience?

Philosophers have long wondered about the nature of mental imagery. Descartes used the example of a 'Chiliagon,' a thousand-sided polygon, to illustrate the difference between mentally conceiving of an object and literally envisioning it (Descartes, 2013). Much later, Ryle roundly rejected the notion of mental images that hang suspended in some internal 'gallery' (Ryle, 2009) and suggested the only thing that can be said with certainty is that people *behave as though* they were seeing a kind of image. By contrast, Block argued that the occasional imperfection or crudeness of mental images is not evidence of them being non-pictorial, calling the claim that it is the 'photographic fallacy' (Block, 1983).

In a famous early psychological investigation of the topic, Galton related that many of his scientific colleagues claimed 'mental imagery was unknown to them,' while members of the general public said they frequently experienced distinct and colorful mental scenes (Galton, 1880). Perhaps surprisingly, one of the strongest contemporary claims about the simulative capacity of imagination is being made by researchers working on memory: In the last two decades, the view that the *episodic memory system* is also responsible for imagination-related processes has become increasingly popular. Some of the empirical findings (Hassabis et al., 2007) and theoretical perspectives (De Brigard, 2014) related to this idea concern imagining one's own future or hypothetical actions and plans in a relatively narrow sense. Others, however, have posited that the episodic memory system could be viewed as a more general event simulation system (Addis, 2020) and that episodic memories are simply those event simulations that have been metacognitively tagged as memories (Mahr et al., 2023). A corollary of such views is that imagination is generatively powerful and creates scenes that are potentially pseudo-experiential in character. This conflicts somewhat with recent findings of Bigelow et al. (2023), who instructed participants to imagine fairly short sentences as vividly as they could: In their experiments, features of scenes that would be difficult to overlook in real-life situations (for example the build, age, or clothes of a person one is directly and *intently* observing) often were not imagined by participants.

Intrigued by this disparity between theoretical perspectives and empirical results, we set out to build on the work of Bigelow et al. (2023) and study how comprehensively and coherently people imagine scenes. Our approach was simple: We asked participants to imagine short narratives we presented to them. Although we evidently could not measure what exactly they imagined, it allowed us to get a sense of what they left out. In Experiment 1, we investigated participants' ability to imagine an ambiguous scene coherently and whether doing so successfully related to its reported completeness. In Experiments 2a – 2c we tested participants' capacity to detect spatial contradictions in the texts they were asked to imagine.

Experiment 1

In Exp. 1, we investigated how people imagine scenes based on short, written narratives. Specifically, we were interested in how receiving ambiguous and apparently conflicting information about a scenario would influence what they imagine. We focused on two aspects in particular: completeness and coherence. A scene is 'complete' if it is as detailed and vivid as an equivalent external scene that is perceived through one's senses. 'Coherent,' for our present

In D. Barner, N.R. Bramley, A. Ruggeri and C.M. Walker (Eds.), *Proceedings of the 47th Annual Conference of the Cognitive Science Society*

purposes, means 'free of contradictions.' If someone reports that they are imagining a prisoner escaping through a window in a room that is explicitly described as windowless, we would label this scene 'incoherent.' The central question of Exp. 1 was: Is there a systematic relationship between completeness and coherence?

As the completeness of a mental scene cannot be measured directly, we had to rely on participants' reports. To this end, we used a modified version of the method employed by Bigelow et al. (2023) and quantified *how many* in a list of predetermined features were part of a participant's mental image. If imagined scenes are just as rich in detail as what we perceive of the external world through our senses, this method arguably would be too crude to offer interesting insights. Many people would consistently report having imagined every feature listed and we would gain no information about individual variability. Conversely, if imagined scenes are often highly incomplete, people should report having not imagined a substantial proportion of several seemingly basic features of a scene, as Bigelow et al. indeed found – a phenomenon they called non-commitment.

In principle, simple descriptive sentences are suitable for assessing the completeness of imagined scenes. The same is not true for coherence, as it is improbable that many people would spontaneously introduce contradictions into their mental renditions of sentences such as 'A person opens a door and walks into a room.' To probe both completeness and coherence, we had to use text stimuli that were at least in some respects ambiguous. Riddles seemed a promising candidate, as finding their solution often requires coming up with ways to reconcile seemingly contradictory pieces of information.

In this first experiment, we presented a single short riddle-like text to participants. It was a lightly adapted version of a vignette used in Bar-Hillel et al.'s 2018 paper on the psychology of riddles:

> Please imagine the following scene, as vividly as you can: A big brown cow is lying down in the middle of a country road. The street lights are not on, the moon is not out, and the skies are heavily clouded. A truck is driving towards the cow at full speed, its headlights off. The driver sees the cow from afar easily, and avoids hitting it, without even having to brake hard.

Because of how the text is worded, many are initially led to assume that this event takes place at night. Figuring out the solution (and thereby envisioning a coherent scene) hinges on how a single feature is imagined: *time of day*.

Rather than challenging participants to solve the riddle, we instructed them to imagine the text as vividly as they could. They were then asked how many in a list of nine features they had imagined. Finally, we inquired whether they had been able to explain to themselves how the truck avoided the cow while they were imagining the text. This design allowed us to analyze whether and how completeness relates to coherence. Based on the aforementioned findings of Bigelow et al. (2023), we expected that people's mental scenes would be rather incomplete – only a minority would report having imagined all nine features we asked about. More importantly, we hypothesized that the completeness of an imagined scene would be entirely unrelated to its coherence (preregistered at: https://osf.io/r6fyg/?view_only=ca05a8b9f01440f6806aaec799ea6355). In other words, we did not expect that the lack or abundance of detail in a participant's mental 'simulation' would correlate with their finding the correct solution of the riddle.

Method

Both Exp. 1 and Exp. 2a – 2c were approved by CEU's Psychological Research Ethics Board.

Participants 41 adult participants took part in Exp. 1. They were recruited via Prolific (https://www.prolific.com) and received £0.90 for their participation. In this and all subsequent experiments, participants were required to be native English speakers and to have an approval rating of 98% or higher on Prolific. Three participants were excluded from analysis because they indicated already having been familiar with the riddle that was presented in the experiment, resulting in N=38.

Materials and procedure The experiment was a survey hosted on the Qualtrics platform (www.qualtrics.com). After they had given their consent, participants were informed that the experiment was about imagination. Then the short text was presented to them. They were instructed to imagine the scene described therein as vividly as they could and were only able advance to the next survey page after at least 60 seconds had elapsed.

Once they did, participants were asked whether they had imagined certain features of the scene. We used a modified version of the approach described in Bigelow et al. (2023), presenting a terse description of some potentially relevant aspect of the scene, for example, *The type of truck*, and offering two response options: 'No, it was not part of my mental image.' or 'Yes, it was part of my mental image.' If someone selected 'Yes,' a new text box appeared, accompanied by a prompt like 'Describe in your own words what type of truck you imagined.'. We asked people about nine such features (see Fig. 1). Each question was presented on a separate page of the survey and the order of presentation was randomized.

Participants were then asked: '**While reading the story and imagining it**, did you come up with an explanation for why the driver had no problem spotting the cow?' Finally, participants were asked whether they had already encountered this kind of riddle previously (the three who answered in the affirmative were excluded from analysis).

Results

Participants on average reported imagining 4.87 out of 9 features (*SD* = 1.70), or 54% of the features they were asked about. *Time of day* was the feature most frequently imagined and *driver's clothes* the least imagined (Fig. 1). No participant reported imagining all 9 features, and no one imagined fewer than 2.

27 out of 38 participants responded with 'No' when asked whether they had an explanation for what happened in the vignette. Out of the 11 'Yes' responses, we classified 8 as correct. All technically correct responses pointed out that it was still day – there was no solution among them that solved the riddle in unforeseen and 'unintended' ways[1].

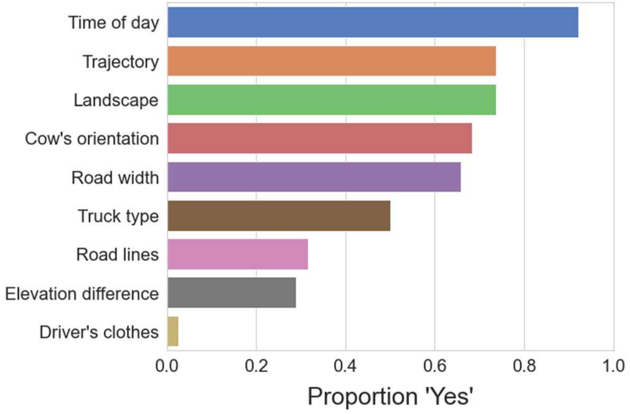

Figure 1: Results of Experiment 1. Histogram of target features showing the proportion of participants that reported imagining them (including 'explicitly' imagining their absence).

To test our first preregistered hypothesis that if and how participants imagined *time of day* would predict whether they could solve the riddle, we first pooled 'No'-responses and incorrect responses together. We then analyzed participants' responses to the specific question in the survey that asked whether they had imagined the time of day. 35 participants reported that they had. If they explicitly stated that they imagined it was day, we coded their input as 'imagined daylight.' As preregistered, we also coded more ambiguous terms describing transitional periods such as 'twilight' or 'dusk' as 'imagined daylight.' Those that clearly stated it was nighttime or reported to not have imagined the feature at all were coded as 'did not imagine daylight.'

A $\chi 2$-test of independence demonstrated a significant difference, $\chi 2(N=38,1) = 8.57, p < 0.001, \varphi_c = 0.47$, between those who did imagine daylight and those who did not with regard to their ability to solve the riddle (Table 1). As we hypothesized, a much higher proportion of those who imagined daylight (7 out of 16) gave correct responses, compared to those who did not imagine daylight (1 out of 22). The 6 participants who reported imagining that it was dusk or evening all failed to provide the correct response. Although we appeared to be too conversative in our classification, the results show that how and whether *time of day* was imagined is predictive of people's ability to solve the riddle.

Table 1: Results of Experiment 1. Correct and incorrect responses and 'time of day' feature.

Time of day	**Correct**	**No or incorrect solution**
Imagined daylight (including 'dusk' or 'twilight')	7	9
Did not imagine daylight (night or no answer)	1	21

Our second preregistered hypothesis stated that the number of features imagined is unrelated to whether someone is able to find the riddle's solution. To test it, we calculated the point-biserial correlation coefficient between the number of features participants imagined ($M = 4.87$, $SD = 1.70$) and the binary value of whether they had provided a correct solution (7 out of 38 had), finding no significant relation, $r_{pb}(36) = 0.02$, $p = 0.91$.

Experiment 2a

In Exp. 1, we investigated the completeness and coherence of imagined scenes by asking participants to vividly picture an *apparently* contradictory state of affairs in their mind's eye. We found that successfully imagining the scene in a way that resolved its ambiguities was entirely unrelated to the level of detail with which they imagined it. Experiments 2a – 2c turned this approach on its head: We again challenged participants to imagine short vignettes as vividly they could, but now the texts we presented did *objectively* contradict themselves. This changed task design would allow us to study whether people imagine how the various elements of their imagined scenes relate to each other.

If people convert a text into a perception-like representation that encodes the spatiotemporal relations between various features, they should be very sensitive to any contradictions present in that source text. After all, contradictions would prevent, or at least severely interfere with, the construction of a coherent, pseudo-perceptual simulation. Even if people are capable of embedding contradictory information into their mental scenes, one would still expect it to be quite salient – perhaps similar to seeing a striking illusion or an impossible image à la some of M.C. Escher's artworks. If, however, people do not tend to form an 'integrated' mental representation of scenes – despite, presumably, trying hard to do so – they might not even become aware of such contradictions.

[1] An example of a response we labelled as correct: 'The scene took place during daylight.' Incorrect responses either focused on irrelevant elements of the scenario or contradicted what was stated in the text: 'I just imagined him being a bit higher in the cab [so] he was able to see the road ahead.' All data are available at the OSF repository for this experiment: https://osf.io/prkn5/?view_only=156a1413f1b844f69c0ff3a58a737ff4

Naturally, there are many ways a text can contradict itself. In Exps. 2a – 2c, we introduced contradictions into short vignettes by giving conflicting descriptions of the *spatial arrangement* of multiple objects. We focused on space, because it seemed to be the dimension of primary importance with regard to the imagination of concrete, physical *scenes*. Furthermore, because we were interested in representations of relations between two or more things or locations, simply giving inconsistent descriptions of, say, the color of an object would not have been sufficient, as this kind of contradiction could be noticed by remembering one item of interest in isolation.

Exp. 2a was a first exploratory study to determine how sensitive people are to conflicting spatial information in a short text. We designed two vignettes that each contained objectively contradictory descriptions. Seeing that previous research and Exp. 1 suggested that people's mental scenes are frequently rather sparsely furnished, we expected that a non-negligible fraction of participants would not notice the error.

Method

Participants For Exp. 2a, we recruited 30 participants on Prolific. Each received £0.75 for their participation.

Materials and procedure We created two new vignettes for this experiment. Each was a short narrative that featured a human character and an animal. Both vignettes contained descriptions of, or pertaining to, spatial arrangements that contradicted earlier parts of the text. This being the first test of this 'contradiction detection' task design, we sought to create vignettes that varied with regard to the subtlety with which spatial descriptions were conveyed.

Vignette 1 was about a woman who sat on a couch, got up to look at her cat, closed a window, and then sat down on the couch again. At one point, a description of the relative positions of couch and window contradicted what was said previously. Stylistically, the text bluntly stated where the various elements of the room were located in relation to the woman. Vignette 2 was about a man waiting for someone else, looking out of the windows of his house as he paced back and forth. It also abounded with descriptions of the man's environment but employed less blatant language. The contradiction in Vignette 2 was that the text suggested that the sun was concurrently visible in two opposite directions (Table 2). Vignette 1 was 191 words long, Vignette 2 211 (see OSF repository for Exps. 2a – 2c's full vignettes: https://osf.io/379zx/?view_only=4787509bbd4c4d2aab8700ee691dbfc7).

As in Exp. 1, participants first had to read the vignette and spend at least 60 seconds imagining it as vividly as they could before being able to advance to the next page of the survey. They were asked a short distractor question that was the same in all versions of Exp. 2 ('Did you imagine the color of the cat's fur?'). This was followed by three questions aimed at assessing whether they had noticed the contradiction in the text. The first one (Q1) was '*Did you notice anything strange about the story?*', the second (Q2) '*There was an aspect of the story that did not make sense at all - an objective error. Knowing this, can you identify the error?*', and the final one (Q3) revealed the error and asked '*If you didn't notice it, does this seem apparent in hindsight?*'. The first two questions participants could answer with either yes or no; if they chose yes, they had to fill in a text box, describing what exactly was strange or erroneous.

Table 2: Excerpts from Vignette 2 illustrating the contradiction. Sentences in square brackets provide contextual information – they were not part of the vignette.

	Initial statement	Contradictory statement
Vignette 2	[The man looks out of a window:] He has to squint to avoid being blinded by the sun that shines right into his eyes.	[The man looks at his cat on the windowsill on the *opposite side* of the room:] Her figure casts a long, strange shadow onto the floorboards.

Results and discussion

Half of the 30 participants who took part in Exp. 2a were presented Vignette 1, the other half Vignette 2. Their responses to questions Q1 and Q2 were analyzed, and the solutions they had written in the text box were coded as either correct or incorrect. Because Q1 was phrased intentionally vaguely and participants may very well have considered multiple aspects of the vignettes 'strange,' there was no objectively correct answer to this first question. The central quantity of interest was therefore how many participants pointed out the contradiction in their response to *either* Q1 or Q2.

We found that 7 out of 15 participants detected the contradiction in *Vignette 1*, 5 of them in response to Q1 (i.e., before being told there was an objective error in the text). 3 participants indicated they saw the error in hindsight. The inconsistency in *Vignette 2* was detected by only 2 of the 15 participants. Both gave the correct solution in response to Q2, not Q1. 9 replied to Q3 that they could see the problem in hindsight.

Experiment 2b

Exp. 2a showed that people are not especially sensitive to contradictory descriptions of spatial arrangements in short texts. The 'error' in Vignette 2 seemed to be particularly difficult to detect and was only spotted by a small fraction of participants. In Exp. 2b, we aimed to investigate what could account for this difficulty other than the possibility that humans' ability to construct integrated models of scenes described in writing is rather limited.

First, we considered that the findings in Exp. 2a could be explained by forgetting or misremembrance. By the time Q2 informs participants that there was an error in the text, they may not remember the scene well enough to retrospect and 'search' for an inconsistency. We therefore added a memory test comprising 8 questions to the survey. If participants have poor memory for the description, then it is little surprise that they would fail to notice the contradiction. On the other hand, if participants remember the scene well, it also ensures they paid attention during the imagination phase.

Second, we wondered whether scaffolding and thus constraining participants' imagination with an external visual aid might help them detect the contradiction. To test this, we designed a second experimental condition in which participants were shown a map of the location described in the text (Fig. 2). They then had to indicate in which compass direction they thought the sun was visible at the moment. Vignette 2 – the only stimulus used in Exp. 2b – gave contradictory information with regard to the sun's current position (Table 2). Accordingly, the only correct response to this question would be to select the option that stated that the information provided by the text did not allow a reader to determine the sun's position.

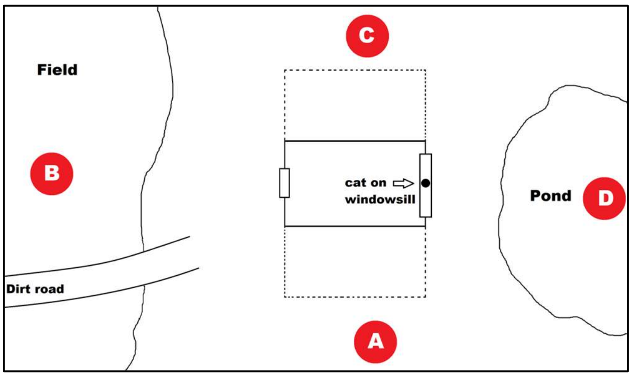

Figure 2: The crude map shown to participants. Vignette 2 suggested the sun was simultaneously present at B and D.

Method

Participants We recruited 40 participants via Prolific. They received £0.75 for their participation.

Materials and procedure The experimental design was a modified version of the survey used in Exp. 2a. Participants were randomly assigned to a *text* or a *map* condition. Those in either group went through an imagination phase that was identical to that in the previous experiment. Participants in the text group then were asked questions Q1- Q3, which were the same as in 2a. In the map condition, however, they were instead presented with a simple map that gave a top-down view of the locations/landmarks mentioned in the text (Fig. 2). A number of spots were labeled to orient them, and four red disks, marked with the letters A, B, C, and D, were placed close to the edges of the map – one in each cardinal direction. Participants were instructed to select the disk they believed corresponded best to the location of the sun. They were given *six* response options they could click on. In addition to the four labeled disks, they could select '*Somewhere else*' or '*The text does not provide the information necessary to answer this question.*'. The vignette was contradictory because it provided equal evidence for answer options *B* and *D,* making the 'insufficient information option' the only correct one. Regardless of which option participants picked, they had to 'explain their reasoning' in an attached text box.

In the subsequent memory test, eight sentences were presented to participants in random order. Each was a statement about the text that was either true or false (three of them were true, five false). The response options for each of these items were the same: either 'True' or 'False / not mentioned'.

Results and discussion

Out of the 40 participants, 23 had been assigned to the *text* and 17 to the *map* condition. Only 2 participants in the text condition successfully detected the contradiction. 0 of the 17 participants in the map condition selected the correct response. The majority of participants, 13, selected *B*; 3 the entirely wrong option *A*; and only one person picked *D*. Participants performed well on the memory task (which comprised 8 questions), the average number of errors being $M = 1.07$ ($SD = 1.52$). 20 participants (including the two who successfully detected the contradiction) made 0 errors.

In conclusion, providing participants with visual materials that scaffold their imagination did not – at least in the manner implemented here – help them detect an inconsistency in the way spatial relations are described in the vignette. Participants' performance on the memory task suggests that an inability to recall details of the vignette is not the issue: Most in the subset of participants who responded correctly to all 8 questions were unable to detect or point out the contradiction (18 out of 20).

Experiment 2c

The findings of Exps. 2a and 2b imply that imagining the layout of a space described in a text, and the spatial relations between objects located in that space, is challenging. While in both preceding experiments only a small minority of participants was able to detect the contradiction in Vignette 2, responses to Vignette 1 in Exp. 2a showed quite a bit more individual variability. In this final experiment, we wanted to investigate what might account for that variability.

Specifically, we wanted to test whether people's spatial visualization ability is related to how comprehensively they imagine the spatial relations in a scene and, consequently, whether they notice the contradiction in the description. To this end, we reused the design of Exp. 2a but followed it up with a mental rotation task. This would us allow to analyze whether performance on that task predicts the detection of the contradiction. We furthermore had participants complete a Stroop task (for an overview, see MacLeod, 1991), to rule out that a potential positive correlation between mental rotation performance and giving the correct solution might not merely reflect higher task engagement. If people's spatial visualization ability influences how they imagine the vignette, performance on the mental rotation task should be a distinct predictor of success.

Method

Participants We recruited 43 participants via Prolific. They received £2.11 for their participation. 13 were excluded based on their task performance (see next section for criteria). We thus arrived at our preregistered sample size of N=30 (https://osf.io/yq6f5/?view_only=54dcbb102f9a4d69920a21cb24071780).

Materials and procedure The experiment began with an imagination phase identical to that in Exp. 2a. This was followed by a memory task that had the same structure as the one described in Exp. 2b, the only difference being that the eight questions were now about Vignette 1. Participants then completed a mental rotation task and a Stroop task (their order was counterbalanced across participants).

The materials we used in the mental rotation task were 24 images of angular 3D shapes. They were a subset of the stimuli created, validated, and made freely available by Ganis and Kievit (2015). Each of the images showed two shapes side by side. In half of them, the shapes were identical – if oriented the same way, they would perfectly overlap. The rest of the shape pairs were different, though very similar (pseudo mirror images). There were four possible angular disparities between the two shapes: 0 (identical orientation), 50, 100, or 150 degrees. These offsets were brought about by clockwise rotations around the vertical axis. The subset of images we picked comprised 24 completely different shapes. In it, all angular disparities were represented with equal frequency. In the online survey, participants saw these images in randomized order and had to report whether they were identical or not by pressing 's' (same) or 'd' (different) on their keyboard. Participants had to respond in less than 10 seconds for their key press to be counted as correct.

In the Stroop task participants had to react to 96 items that were presented to them in randomized order. In our implementation, participants saw color words (like 'red') that half of the time were of a color that did not match the word's meaning (e.g., 'red' being green). They had to press a key that matched the *color*, not the meaning, within 5 seconds or less.

To not confound the influence of mental rotation ability with potential memory effects, we excluded participants who made more than two mistakes in the memory task. We also excluded those who responded incorrectly to more than 10 of the 24 items in the mental rotation task.

Results and discussion

Of the 43 participants who completed the online experiment, we excluded 5 who made three or more errors in the memory task, 7 because they gave fewer than 14 out of 24 possible correct responses in the mental rotation task, and one participant because they made 48 errors in the Stroop task. This left us with us with 30 participants.

After determining that 9 participants (30%) had reported the contradiction in response to either Q1 or Q2, we tested whether performance on the mental rotation task correlated with success. First, we looked at a potential association with accuracy (the number of correct responses divided by the total number of items), $M = 0.81$, $SD = 0.13$. A point-biserial correlation analysis, $r_{pb}(28)= 0.07$, showed no significant relation, $p = 0.72$. The same was true for the average response time (only correct responses), $M = 3.36$ seconds, $SD = 1.16$, $r_{pb}(28)= 0.03$, $p = 0.86$. Findings were similar in the Stroop task: There was no correlation between accuracy, $M = 0.98$, $SD = 0.02$, and the detection of the contradiction, $r_{pb}(28)= 0.008$, $p = 0.97$. Likewise, response time (only correct responses), $M = 1.08$ seconds, $SD = 0.26$, did not significantly correlate with reporting the contradiction, $r_{pb}(28)= 0.06$, $p = 0.75$.

Discussion

In two experiments, we asked participants to use their imagination to translate short texts into vivid mental scenes. In Exp. 1, participants' reports about what they imagined suggested these scenes tend to lack detail, corroborating the findings of Bigelow et al. (2023). The average participant reported imagining only 54% of the features we asked about (most of which were arguably central features of the described event). Exp. 1 not only tested participants' ability to 'paint a picture' in their imagination but also implicitly challenged them to resolve an apparent contradiction in the riddle-like text. We found that successfully imagining the scene coherently, i.e., 'solving the riddle,' was unrelated to the number of features they envisioned. Moreover, our results show that the difficulty of this particular riddle does not arise from people neglecting to imagine a decisive feature (time of day) but imagining it *incorrectly*. Apparently the process of mentally constructing a scene is quite opaque to the average imaginer in this case: Adding a feature to one's mental picture does not imply one is aware of it, much less able to systematically vary it ('Wait, I am imagining it is night – what if the sun were shining?').

Exp. 2 demonstrated that people are not especially sensitive to contradictions in spatial descriptions. Even when these descriptions were the overt focus of a narrative, fewer than half of the participants in Exps. 2a and 2c noticed the contradictions. The detection rate for the subtler Vignette 2 was much lower still. If people converted the texts they were instructed to imagine as 'vividly as they can' into something akin to 3D simulations of the spaces described therein, one might expect contradictions to be highly salient. The text, after all, is what guides the imagination process. Here it of course has to be noted that we only used a limited number of vignettes in these preliminary investigations which restricts the generalizability of these findings.

The data so far suggest that the scenes people imagine when reading a text are not just a few missing details away from being simulacra of reality. Instead of an integrated scene where all relations between parts are specified or at least inferable, mental images may frequently be composed of multiple, somewhat independent, feature patches that together create the illusory impression of an organized whole.

References

Addis, D. R. (2020). Mental time travel? A neurocognitive model of event simulation. Review of *Philosophy and Psychology, 11(2),* 233-259. https://doi.org/10.1007/s13164-020-00470-0.

Bar-Hillel, M., Noah, T., & Frederick, S. (2018). Learning psychology from riddles: The case of stumpers. *Judgment and Decision Making, 13(1),* 112-122. https://doi.org/10.1017/S193029750000886X

Bigelow, E. J., McCoy, J. P., & Ullman, T. D. (2023). Non-commitment in mental imagery. *Cognition*, *238*, 105498. https://doi.org/10.1016/j.cognition.2023.105498

Block, N. (1983). The photographic fallacy in the debate about mental imagery. *Nous,* 651-661.

De Brigard, F. (2014). Is memory for remembering? Recollection as a form of episodic hypothetical thinking. *Synthese, 191,* 155-185. https://doi.org/10.1007/s11229-013-0247-7

Descartes, R. (2013). *Meditations on first philosophy*. Broadview Press.

Galton, F. (1880). Statistics of mental imagery. *Mind*, *5*(19), 301-318.

Ganis, G., & Kievit, R. A. (2015). A New Set of Three-Dimensional Shapes for Investigating Mental Rotation Processes: Validation Data and Stimulus Set. *Journal of Open Psychology Data*, *3*(1), e3. https://doi.org/10.5334/jopd.ai

Hassabis, D., Kumaran, D., Vann, S. D., & Maguire, E. A. (2007). Patients with hippocampal amnesia cannot imagine new experiences. Proceedings of the National Academy of Sciences, 104(5), 1726-1731. https://doi.org/10.1016/j.tics.2007.05.001

MacLeod, C. M. (1991). Half a century of research on the Stroop effect: an integrative review. *Psychological bulletin*, *109*(2), 163.

Mahr, J. B., van Bergen, P., Sutton, J., Schacter, D. L., & Heyes, C. (2023). Mnemicity: A Cognitive Gadget? *Perspectives on Psychological Science*, *18*(5), 1160-1177. https://doi.org/10.1177/17456916221141352

Ryle, G., & Tanney, J. (2009). *The concept of mind*. Routledge.