STUDIA SEMIOTYCZNE (SEMIOTIC STUDIES), t. XXXV, nr 1 (2021), s. 73–94
ISSN 0137-6608, e-ISSN 2544-073X
DOI: 10.26333/sts.xxxv1.05
Article

PIOTR KOZAK [*]

THE ANALOG-DIGITAL DISTINCTION FAILS TO EXPLAIN THE PERCEPTION-THOUGHT DISTINCTION: AN ALTERNATIVE ACCOUNT OF THE FORMAT OF MENTAL REPRESENTATION[1]

SUMMARY: The format of mental representation is the way information is organized in the mind. The discussion surrounding the format of representation addresses the problem of what representational primitives are and the rules of information processing.

In philosophy, the discussion is dominated by the distinction between analog and digital representational systems. It is thought that this distinction can bring us closer to an understanding of the nature of perceptual and discursive representations.

I argue that the analog-digital distinction cannot meet that expectation. The analog-digital distinction is neither sufficient nor necessary to explain the distinction between perceptual and discursive representations (and perception and thinking, respectively). I propose an alternative interpretation of the concept of representational format which provides us a better understanding of the difference between iconic and discursive representations. I explain the differences between formats of representations in terms of differences in information processing. I demonstrate, how this alternative interpretation of the concept of the representational format can explain the constraints put on the contents of representational systems.

KEYWORDS: mental representation format, analog, digital, perception, thought, iconic representations, discoursive representations.

[*] University of Bialystok, Institute of Philosophy. E-mail: piotr.kozak1@gmail.com. ORCID: 0000-0001-9734-4640.

[1] This work was supported by the research grant “What is Thinking with Images?”, SONATA 10, granted by the National Science Centre, Poland, based on decision No. 2015/19/D/HS1/02426.

A mental representation format is the way information is organized in the mind.[2] A discussion of mental representation formats addresses the question of how the information in our mind is stored and processed. It concerns the structure of representations interpreted as a set of representational primitives and combinatorial principles. Thus, to describe a representational format, one has to describe the structure of representation, that is, one has to explain what the primitive elements are and the possible operations that can be carried out with them.

There are two philosophical traditions of thinking about the format of representation. First, following Goodman (1976), we distinguish between analog and digital systems of representation. Second, Goodman's distinction between two formats of representational systems has been adopted in the philosophy of mind as the basis for the distinction between perceptual and discursive representations. Following Dretske (1981), it is argued (Peacocke, 1989) that perception consists of iconic representations and is analog in format. In contrast, beliefs and thoughts are discursive representations and are encoded in digital format.[3] In this paper, I will use the terms "iconic" and "discursive representation" as referring to representations that describe perceptual and discursive mental phenomena, respectively.

Let me give two examples of the discussion surrounding representational format. First, the early stage of the so-called imagery debate (from the 70s to the beginning of the 90s) was mostly devoted to issues surrounding the way our minds encode mental images. On the one side, pictorialists (e.g., Kosslyn, 1980) have held that the format of mental images is perceptual-like and analog. On the other side, descriptionalists (e.g., Pylyshyn, 1973, 1981) have argued that mental images are formed out of structured descriptions that are digital in format. Most descriptionalists have argued that mental images are epiphenomena of some internal discursive, language-like processes.

Second, the debate on representational format underpins philosophical discussions on the nature of perceptual representations. On the one hand, some philosophers, most notably Sellars (1997) and McDowell (1996), have held that perceptual representations are propositional. On the other, there are philosophers, such as Crane (2009), Dretske (1969), and Travis (2013), who have argued that perceptual representations are distinct in kind from discursive representations. For those who deny that perceptual representations are propositional, the question arises what kind of representations they can be. The most common answer is that perception has an iconic structure. It consists of iconic elements and relations between these elements. The "atoms" of perception are iconic representations that are most often described as having an analog nature and as being deprived of ca-

[2] The concept of the format of mental representation is different from the concept of the vehicle of representation. Information can be stored in the same format in different types of vehicles. Using a computer metaphor, the same .jpg file format can be stored on both magnetic vehicles, as in the case of a floppy disk, and optical vehicles, as in the case of a CD.

[3] That does not mean that there are no intermediate representations, such as maps and pictographs, see Casati and Giardino (2013).

nonical decomposition. Interactions between these elements are based on causal relations between iconic representations. In contrast, interactions between discursive representations are based on logical transitions (e.g., Matthen, 2005).

The difference between analog and digital formats of representation is intuitive but conceptually blurred. According to Goodman (1976), to be an analog representational system means to be both a syntactically and semantically dense representational system. An example of a dense representational system is an old-fashioned clock that represents time continuously, unlike a digital clock that represent time discreetly. Moreover, an analog representational system is relatively replete. A representational system is relatively replete if, in comparison with other systems, many of its members' features are relevant to determining what they represent. The system of old-fashioned analog clocks is not replete, since only the position of the clock's hands matter. In comparison, in the case of images, such features as colour, shape, and size are relevant. However, for reasons that will not be covered here (e.g., Kulvicki, 2006), it is doubtful whether Goodman succeeded in adequately explaining analog and digital formats of representation.

In the last 50 years, the distinction has been variously interpreted and explicated (e.g., Fodor, Pylyshyn, 1981; Haugeland, 1998; Lewis, 1971; McGinn, 1989). Across those approaches, digital representations are generally understood to be discrete entities. Numerals provide a good example. "0" and "1" are discrete because they indicate distinct and separable entities. For every representational token, it is clear which type it instantiates. In contrast, analog representations do not admit definite type-identity. For example, the colour value of a given colour patch is measured on a continuous rather than a discrete scale (Dretske, 1981). There is always room for the question of whether a given colour patch is more blue-like or dark blue-like. In contrast, whereas there is no room for the question of what number is represented by "0".

Iconic representations are believed to be analog structures. For instance, there is no way to determine a discrete point where a blue colour patch ends and a dark blue one begins. The structure of discursive representations is believed to be based on digital operations. For instance, if one believes that the king is dead, one thinks about the king, and not *more or less* about the king, ascribing the property of being dead and not being *more or less* dead to the king. In contrast, a mental image of the dead king more or less resembles the king being dead.

However, this understanding of the analog-digital distinction is far from being clear (e.g., Lorenzo, Rubiera, 2019; Maley, 2011). A colour patch can be represented in analog format but can be represented digitally as well, namely, as a set of colour values in the RGB colour model. Musical notes C and C# are discrete when playing piano, but there is a continuum of notes between C and C# when playing a violin. A film frame is discrete, but the events depicted with the help of film frames are indiscrete.

An alternative way to interpret the analog-digital distinction can be put in terms of constraints that the representational system puts on representational content. It can be illustrated with notational systems in mathematics. Although

the same mathematical magnitudes can be recorded in different notational systems, such as Roman or Arabic numeral systems, these systems are not computationally equivalent. For example, it is more efficient to carry out a calculation with large numbers in Arabic than the Roman numeral system. Analogously, it may be more efficient to represent the values of a linear function with the help of a graph than with the help of a numeral matrix.

Thus, the difference between representational systems can be interpreted in terms of being capable of expressing different kinds of content. That means that different representational systems put constraints on the content a representational system can carry and the range of possible transitions between different contents. So, for example, an analog representation can represent the value of a magnitude but not an integer (Beck, 2015), iconic representations cannot be used to represent a negation (Crane, 2009), etc.

Yet these two interpretations are linked, for if one wants to explain why some representational systems put constraints on representational content, then one has to describe the features of the representational structure. For instance, one can explain why the Arabic numeral system is preferred over the Roman numeral system by pointing out the fact that the Roman numeral system does not have the concept of zero. Analogously, iconic and discursive representational systems put constraints on their representational contents. A theory of representational format should explain where these constraints come from.

To put it more generally, the question is what should the concept of the format of mental representation explain and how it can do that. I claim that the problem with the analog-digital distinction is not that it is not clear. Even if it were clear, it would still be doubtful whether it could explain what it should explain, namely, the difference between thoughts and perceptions.

In this paper, I propose an alternative interpretation of the representational format. I claim that it provides us a better understanding of what the difference between iconic and discursive representations is. In the next section, I show what any theory of representational format should be able to explain. I demonstrate how to interpret the difference between iconic and discursive representations. I claim that discursive representations can meet the requirements of the so-called Generality Constraint, while iconic representations cannot. Next, I explain how one can understand the difference between iconic and discursive formats of representation. I put it in terms of differences of information processing in cognitive systems. Last but not least, I demonstrate how the alternative interpretation of the concept of a representation format can explain the constraints put on the contents of representational systems.

1. Generality Constraint

One of the distinctive features of discursive representations is that they are systematically structured. Entertaining a thought of one kind entails a capacity to entertain a thought of another kind. For instance, entertaining the thought that

John is happy and that Mary is sad is systematically connected with the cognitive ability to entertain the thought that Mary is happy and that John is sad. Having the thought that John is happy entails a capacity to think that someone is happy. In Evans' words (1982, p. 104), "if a subject can be credited with the thought that *a* is *F*, then he must have the conceptual resources for entertaining the thought that *a* is *G*, for every property of being *G* of which he has a conception."

The same rule applies to inferences (e.g., Fodor, Pylyshyn, 1988). If I think that it is dark and cold and raining, I can infer that it is cold and raining; for from *P* & *Q* I can infer that *P* (or *Q*). By the same token, I must be able to infer from it is cold and raining that it is raining. If I am unable to do so, I do not know what inference is.

Thus, discursive representations are systematically co-related (e.g., Heck, 2000; Peacocke, 1992), which means that they are systematic in nature. Evans (1982) calls this requirement the Generality Constraint.

To meet this requirement, discursive representations have to consist of recombinable constituents that can build more complex structures. It means that discursive representations are compositional in nature. Compositionality of discursive representations means, first, that the meaning of complex structures is determined by the meaning of their constituents. The constituents of discursive representations are parts of the representations that are canonically distinguishable, for not every partition of the representation makes sense. The idea is that canonically decomposed parts are syntactically and semantically meaningful units. The thought that John loves Mary can be decomposed into John loves and Mary, but not into John…Mary (e.g., Fodor, 2008).

Second, the meaning of complex discursive representations must come from the meaning of their canonically distinguishable parts together with the rules of composition, for not all combinations are allowed. The recombination of the parts must be meaningful. John loves Mary can be recombined into Mary loves John, but not into John Mary loves.[4] These rules are recursive. If I have a thought that John loves his mother, I must be capable of having the thought that John loves his mother's mother, etc. Putting it together, discursive representations have a recursive syntax that combines canonically distinguishable parts according to combinatorial rules (e.g., Pagin, Westerståhl, 2010).

Language seems to be systematic and compositional. It has syntax and distinguishable syntactic and semantic parts. Thus, one may infer that discursive representations are language-like representations (e.g., Devitt, 2006). In contrast, iconic representations lack systematicity and compositionality, and therefore they

[4] According to Evans (1982), systematicity is constrained by semantic conditions of appropriateness. For instance, thinking that JOHN FELL INTO THE LAKE need not entail a capacity to think that THE LAKE FELL INTO JOHN. However, even if a well-formed string of thoughts is a semantical absurdity, it does not mean that it cannot express thoughts. For one thing, we can entertain absurd thoughts. For another, an absurd but well-formed string of thoughts can be the basis of inferences in logic. See, e.g., Camp (2004).

do not have the metaphysical properties we are looking to ascribe to discursive representations.

Iconic representations are neither systematic nor compositional, for they lack syntactic structure.[5] According to Fodor, they lack canonical decomposition. According to Frege, they lack logical form. Therefore, iconic representations do not meet the Generality Constraint.

Fodor's argument (2007; 2008) from lack of canonical decomposition takes the form of the so-called Picture Principle. According to the Picture Principle, iconic representations can be distinguished topologically: although pictures have interpretable parts, they lack canonical decomposition. It means, loosely, that we can cut up a picture however we like, and each picture-part will represent a relevant part of the represented object. Thus, every part of the representation represents some part of the scene represented by the whole representation (e.g., Green, Quilty-Dunn, 2017; Quilty-Dunn, 2016; 2020; Sober, 1976). In contrast, discursive representations have a canonical decomposition, which means that they cannot be cut into pieces however we like. Discursive representations have constituent parts. For instance, the content of the proposition snow is white can be decomposed into the parts snow and is white, but not into snow…white, which means that the expression "snow…white" does not possess independent semantical value. Thus, although iconic representations can be decomposed, they cannot be canonically decomposed. However, if they can be composed and decomposed however one wants, then they lack syntactical structure.

Frege's argument from lack of logical form is based on the observation that icons are unable to express logical relations. For logical relations to hold, the elements of the relation have to possess a logical form, i.e., syntactically fixed structure, such as a set of logical constants and variables, with determined transformational rules that preserve the logical values of its components. If A implies B, then basing on transformational rules it is possible to transform the truth of the first into the truth of the second. Propositional logic shows how the truth of complex propositions depends on the truth of simple ones. Truth-functions operate on propositions that can be negated, disjoined, and conjoined; they can imply one another or be equivalent. One of the main reasons for talking about propositions at all is that they explain how things can stand in these logical relations.

Iconic representations cannot express logical relations, for they lack logical form. There are no truth-preserving transformation rules for imagistic representation. There is no pictorial negation (Crane, 2009; Sainsbury, 2005), conjunction, or disjunction (Heck, 2007); images cannot express implications or quantifications (Frege, 1984), etc.

[5] That does not mean that they lack construction rules. They are obviously rule-governed. That is why if one understands how to interpret one Venn-diagram, then one understands how to understand another. However, a representational system can have construction rules without syntactic structure. I can construct a triangle according to the rules of construction, but that does not mean that triangles have a syntax.

Frege's argument implies that a clear line between iconic and discursive representations can be drawn. Discursive representations are interpretable in logical terms, icons are not. Discursive representations are "inferentially promiscuous" (Stich, 1978), which means that they can figure as premises in logical transitions.

Moreover, lack of logical form renders iconic representations a-rational; they are neither rational nor irrational—the concept of rationality simply does not apply to them. Relations between iconic representations are not logical; these relations are usually understood as a causal chain of associations. An image of a mother can evoke a memory image of a family home, but the link between these two images is not a matter of a logical consequence. We can speak of the temporal or causal sequence of images but rationality is not based on temporal or causal links. It is a matter of following logical rules and reasons. Therefore, iconic representations cannot be rational, and if someone like Frege thinks that the core of thinking is rational thinking, then icons cannot be constituents of thoughts.

Two remarks are required here. First, it may seem that Frege's argument can be easily refuted by pointing out straightforward counterexamples. For instance, if I want to negate that John has red hair, I can depict him as blond. If I depict a green and a red apple, I express an alternative of a green and red apple. If one places two pictures next to each other, much like in a comic book, then one can say that their content is conjoined or implies one another (Westerhoff, 2005).

These examples are, however, misleading. The role of logical form is determining the truth-conditions of its elements. In the case of iconic representations, truth-conditions cannot be determined. Having a picture of John with blond is either a negation of having red hair or black hair, etc., for the content of not-redheaded is not simply being blond but an infinite alternative of a form: being blond or having black hair, or having green hair, etc. No image can represent infinitely many properties.

By the same token, conjunction, disjunction, and implication do not simply represent a sequence of elements; they set up a logical link between them. In the case of two pictures, there is no way to determine the nature of this link—whether it is a temporal sequence, causal link, spatial transformation, or if it is simply a set of two unrelated pictures. In all these cases, the pictorial form is the same.

Let me illustrate these problems with the mental model theory of reasoning (e.g., Barrouillet et al., 2000; Byrne, 2005; Byrne and Johnson-Laird, 1989; Johnson-Laird, 1983). A mental model is a schematic representation of a possible state of affairs. It represents the elements of a set as well as possible spatial and causal relations between the relevant elements of the set. Manipulation of the spatial and causal properties of a model allows one to reason about the properties of the set's element. As an example, let us try to solve the following syllogism:

(1) Some artists are beekeepers.
All beekeepers are chemists.
What follows?

To do this task, we can form a mental model of an artist who is at the same time a beekeeper and a chemist. The task is easy: some artists are chemists.

On the surface, it may be tempting to interpret mental models as iconic, predominantly visual, representations (e.g., Johnson-Laird, 1983). The idea is that we visualise the elements of the syllogism that can help us to solve it. However, visual models are not logical structures (e.g., Hintikka, 1987; Johnson-Laird, 1998; Knauff and Johnson-Laird, 2002). They lack the generality and precision that is required by logical operations. For instance, notice that the same mental model can be used in the cases of the following distinct syllogisms:

(2) Some artists are beekeepers.
All beekeepers are chemists.
Ergo: All artists are chemists.

(3) Some artists are beekeepers.
All beekeepers are chemists.
Ergo: Some artists are chemists.

There is also nothing that would separate this reasoning from fallacious reasoning, such as:

(4) Some artists are beekeepers.
All beekeepers are chemists.
Ergo: All artists are beekeepers.

The problem is that it is impossible to depict the difference between the claim that $\exists x\mathrm{P}(x)$ and $\forall x\mathrm{P}(x)$. Therefore, iconic representations lack logical form.

Second, one might object that iconic representations can exhibit systematicity. For one thing, map-like representations seem to be systematic (e.g., Camp, 2007; Braddon-Mitchell, Jackson, 1996). A part of a map that represents that London is west of Berlin also represents that Berlin is east of London. For another, as Matthen (2005) notes criticizing Evans' Generality Constraint, if one can imagine a blue circle and a red square, then one can imagine a red circle and a blue square. In other words, if a representational system is capable of representing multiple features together, then it can represent different configurations of these features.

These objections, however, miss the mark, for the systematicity of discursive representations comes paired with compositionality. For discursive representations to be systematic, we have to be able to distinguish between the meanings of the constituents and the meanings of the complex structures they form. For instance, the thought that John loves Mary is built out of the concepts John, Mary, and love, which can be distinguished as separate semantical units. In the case of iconic representations, such separation cannot be carried out, for they lack canonical decomposition.

Does the analog-digital distinction help us to understand the iconic-discursive distinction? It seems that it does not, as the iconic-discursive distinction does not overlap with the analog-digital one.

For one thing, as von Neumann (1958) demonstrates, discursive operations, such as computations, can be functions of digital and analog processes. From the h a r d w a r e point of view, discursive representations can be encoded either in digital or analog format. Thus, the distinction between analog and digital format is irrelevant for determining whether we are dealing with an iconic or discursive representational system. For another, this distinction falsely implies that there are no digital iconic representations. There obviously are. Therefore, being analog or digital is neither necessary nor sufficient for being an iconic or discursive representation, which sometimes leads to the conclusion that the analog-digital distinction is only notational (e.g., Johnson, 2015; Szabo, 2012).

However, the difference between iconic and discursive representational systems can be described in terms of different ways iconic and discursive information is processed in the mind. In the next section, I present an alternative understanding of the concept of representational format.

2. Mental Representation Format as a Way of Processing Information

There are at least two marks that distinguish the structure of iconic and discursive representations. Concerning the first, mechanisms of information processing in iconic representations are domain-specific. The domain-specificity of iconic representations can be understood as a joint alternative of two theses. First, it means that the mechanism of information processing varies depending on the nature of the vehicle of representation. Second, mechanisms of information processing depend on the modality of representation. In contrast, mechanisms of information processing of discursive representations are domain-general, which means here that they do not depend on the features of the vehicle of representation nor the modality of representation.

Concerning the second mark, iconic and discursive representations employ different predicative functions. The structure of iconic content is organised non-hierarchically and is based on holistic data. In contrast, the structure of discursive content is organised hierarchically and is based on discrete chunks of information.

The mechanisms of information processing are domain-specific if the operations defined on the elements of the structure depend on the area of application. The relevant mechanisms are domain-general if the operations do not depend on the area they are applied to. For instance, rules of addition are domain-general; regardless of what one is adding, the rules are the same. In contrast, heuristic rules are domain-specific, for there is no general heuristic that could help us solve every type of cognitive task. Depending on what one is trying to solve, one uses different heuristics.

The mechanisms of information processing in iconic representations are domain-specific. This claim consists of a joint alternative of two theses. For one

thing, the mechanism of information processing depends on the features of the vehicle; for another, it depends on the modality of representation. Let us dub them "vehicle-specificity" and the "modality-specificity", respectively. I distinguish between vehicle-specificity and modality-specificity mostly because I do not want to settle whether discursive representations are amodal here (Prinz, 2002). In other words, one can hold that discursive representations, such as concepts, are modally-specific; thus, modality-specificity is not necessarily a valid criterion for distinguishing between iconic and discursive representations. However, acknowledging vehicle-specificity as a criterion of iconic format does not imply that one has to acknowledge modality-specificity as a relevant criterion (although the implication is the other way round).

Vehicle-specificity means that the features of the vehicle of representation determine the way we process the information. First, access to the information varies depending on whether the information is displayed on an external or internal vehicle of representation. External vehicles of representations are central for using pictures, maps, diagrams, or gestures. Internal vehicles of representations are central for mental imagery and perception. Second, access to the information varies depending on the way information is displayed on the vehicle of representation.

The distinction between internal and external vehicles of representation corresponds to different mechanisms for how iconic and discursive information is processed in the mind. Let us illustrate it with a mental imagery example. In comparison to external images, it is widely believed that the content of mental images is not subject to interpretation but is displayed as already interpreted (e.g., Chambers, Reisberg, 1985; Reisberg, 1996; Reisberg, Heuer, 2005; Sartre, 1962; Slezak, 1995). In contrast, we can always reinterpret the content of picture perception. It means that the meaning of mental images is fixed, while the meaning of external images may change accordingly to the way we perceive the image.

Two clarifications are needed. The inability to reinterpret mental images is subject to scientific dispute. Contrary to the classic positions in cognitive psychology represented, for instance, by Chambers and Reisberg (1985), it seems that we can reinterpret the content of mental images. However, the reinterpretation of mental images is more cognitively loaded and less efficient. For example, we can reinterpret so-called bistable figures (such as the duck-rabbit picture) displayed on an external and an internal representation but with significant differences in the effectiveness of solving the task (Mast, Kosslyn, 2002; Kamermans et al., 2019). That suggests, first, that the difference between mental imagery and picture perception is vaguer than we previously thought. It is not a difference in kind, rather a difference in degree. This fact is easier to understand if we assume that mental imagery and pictorial perception share the same format but that the format is vehicle-specific. Second, it indicates that the descriptivist positions (e.g., Pylyshyn, 1973) that assume only the discursive format of representation in the imagery debate are wrong, for if mental images were encoded in a discursive format, then it would be difficult to explain why mental images are subject to reinterpretation.

Moreover, it is important to keep in mind that even proponents of the pictorial theory of mental imagery, such as Kosslyn, have never claimed that the nature of mental images is the same as picture perception. At most, they speak of the quasi-perceptual nature of imagery. In other words, even if it is not clear what the difference between perception and mental imagery is, no one claims that there is no such difference. For instance, as was pointed out by Hinton (1979),[6] if we form a mental image of nine letters put randomly into a 3×3 grid, it is difficult for us to read the imagined string of letters. However, it is a trivial task if we write down the same letters on paper. Thus, even if information encoded in internal representation is processed similarly to the way information encoded in external representation is processed, they are not of the same processes (Ittelson, 1996). It does not mean that the format is different. It means that the format is vehicle-specific.

Furthermore, the different ways in which a piece of information is displayed on the vehicle of representation affects the accessibility of the piece of information. For instance, a line graph (a) and a bar graph (b), as shown in Figure 1, can represent the same information but in different ways. Graph (a) makes it easier to understand the relationship between the amount of exercise and weight loss. The data is connected by an increasing function, whereas graph (b) makes it easier to understand the relationship between exercise and the number of calories burned because the bars comparing the data of calories and the amount of exercise are closer to one another.

Broadly speaking, we can assume that the features of the vehicle can affect the mechanisms of information processing in iconic representations. The features of the vehicle do not merely provide input for certain internal processes. The interpretation of the same content displayed on different types of vehicles involves separate mental processes. In other words, the type of vehicle of representation determines what information we have access to and what mental processes are involved in processing that information. Using Larkin and Simon's (1987) formulation, iconic representations are computationally inequivalent due to the type of vehicle of representation.

[6] Kosslyn (1994) argues that Hinton's results indicate the mere fact that mental images are displayed in the mind's eye too short, and, therefore, the results show only how memory limitations affect imagery. However, the time that is not sufficient to solve Hinton's tasks is the same time that suffices to solve mental rotation and mental zooming tasks in Kosslyn's classic research on mental imagery. Therefore, the time for which information is available does not seem to be a relevant factor.

Figure 1

Differences in Information Processing in Graphs

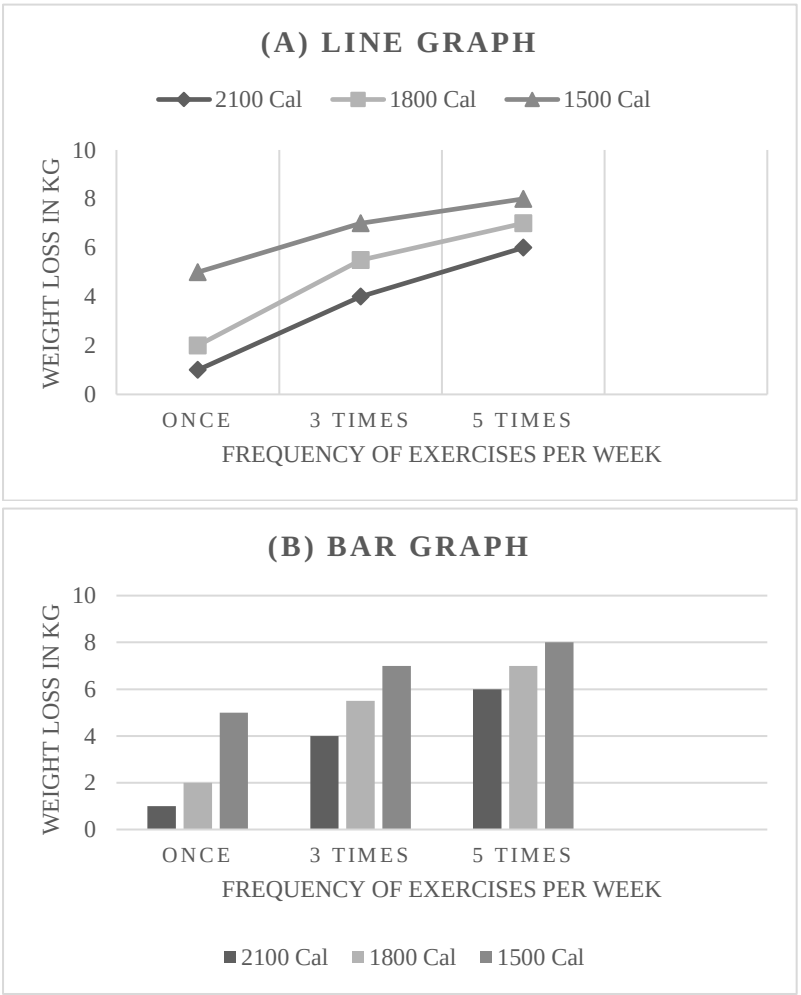

Note. Line graphs (a) and bar charts (b) convey the same information but in a different way, affecting the accessibility of the information.

In contrast, in discursive representations, access to the content of the belief that snow is white is the same as access to the content of the sentence "snow is white". The mechanism for processing the information that snow is white represented by the thought snow is white and represented by the sentence "snow is white" is the same because the propositions contained in the belief and expressed

in the sentence are the same. In other words, the same predicative function predicating white of snow is expressed in both the internal and external representation.

Moreover, there is shared meaning in the case of different external representations. The expression "snow is white" expresses the same information as its equivalent in German "Schnee ist Weiss". The same content can be expressed by two syntactically different representations. The cognitive content of the sentence "snow is white" is the same as "the colour of snow is white".

We face a similar situation in reasoning. For instance, the function 2+2=4 is carried out in the same way regardless of whether it is done in one's head or on paper. The reasoning is the same because the cognitive content is the same. This means that the information is processed in the same way because the operations on the sets are the same. The representations are computationally equivalent. Thus, content is not affected by differences in the features of vehicles of discursive representations.

One remark is required. From the given description, it does not follow that there is no difference in cognitive access between the mental content inside the mind and the content expressed by an external (e.g., linguistic) representation of a thought. A thesis of this type would be false for obvious reasons. For example, calculations performed on a piece of paper may be cognitively less loaded than those performed in thought. However, this does not mean that the structure of information is different. To use a computer metaphor, the .jpg format can be supported by more or less computationally efficient computers. The type of computer, however, does not affect the format of the file.

The second way of understanding the domain-specificity of the iconic format of representation refers to the concept of modality-specificity. The mechanisms of information processing depend on the modality of representation. For instance, a visual and gustatory representation of wine are two different systems of representation—the information is processed differently. In contrast, discursive representations are amodal—discursive representations of the colour and taste of wine are different in content but the information corresponding to the the colour and taste is processed in the same way.

Processing iconic information involves different mechanisms that are responsible for processing information of different sensory types (interoceptive, visual, tactile, auditory, etc.), which is associated with activation of different neurobiological systems. This point can be illustrated with individual differences in visual imagery and spatial cognition tasks (e.g., Hegarty, Waller, 2005; Kozhevnikov, Blazhenkova, Becker, 2010). In short, there are large individual differences in tasks where people are asked to imagine a sunny day and tasks which measure spatial abilities, such as when they are asked to imagine the spatial transformation of mental images. These dimensions are uncorrelated (e.g., Kosslyn et al, 1984) and can negatively affect each other. For example, visualising the content of a problem can lower the effectiveness of reasoning in spatial and abstract problem-solving tasks, which is known as the visual impedance effect (e.g., Knauff, Johnson-Laird, 2002).

The details of different mechanisms of information processing here are of less importance. I only wish to emphasize the fact that iconic representations are modality-specific. In contrast, discursive representations are amodal. The proposition snow is white remains the same when it is expressed in a spoken or written form, in English or in German.

The other detail that distinguishes iconic and discursive representations concerns the problem of predication. It is believed that discursive representations, such as beliefs, have a predicative nature. When I have a belief that snow is white, I attribute the property white to the object snow, where the terms "snow" and "white" work as arguments of a predicative function expressed in the proposition. Thanks to the predicative nature of the proposition, we can distinguish the proposition from a list of terms. The proposition snow is white differs from the list of terms "snow", "white", and "is" because the proposition has a predicative structure that carries the denotations of the terms into a truth-value (e.g., Rescorla, 2009), while the list of terms does not. It means that the proposition snow is white can say something about the world, while the list of terms cannot.

It may seem that iconic representations can be predicative too (e.g., Blumson, 2012; Matthen, 2005). If I form an image of a red triangle, I attribute the property of redness to the triangle. And the image of a red triangle can be distinguished from the conjunction made of an image of a triangle and an image of a red patch. However, to assess whether or not we are dealing with equivocation here, we must have a clear understanding of what predication is.

First, in the case of discursive representations, predication is based on a compositional and combinatorial mechanism. The compositional and combinatorial character of a discursive representational system means that if I possess the propositions snow is white and a triangle is red, I can form the structurally similar propositions snow is red and a triangle is white. Moreover, if I know that snow is white is true, I can infer that the proposition snow has a colour is also true. It means that an output of one operation of predication can be an input of another higher-order operation. These operations are hierarchically organised (e.g., Camp, 2018). From the proposition snow is white, I can infer that snow has a colour, but from the proposition snow has a colour I cannot infer that snow is white.

Second, the output and input information comes in discrete chunks. It means that the proposition snow is white attributes the property whiteness and no other property to snow. It says nothing about the hue of the colour or the shape of snow, for there is one-to-one correspondence between a vehicle and a content. Every chunk of information needs a separate vehicle to be expressed. Thus, discursive representations are hierarchically organised and are based on a structure made of discrete chunks of information.

In contrast, iconic representation is neither hierarchically organised, nor is it based on discrete chunks of information. Hierarchical organisation of representational structure means, for instance, that the proposition snow is white implies snow has a colour but not another way round. In the case of iconic representation, both pieces of information are processed simultaneously. An image of white

snow represents both that snow is white and that snow has a colour. That is why it does not matter from which point we start analysing an image of white snow—whether from thinking of snow as having a colour or of snow being white—both starting points lead to the same result.

The non-hierarchical structure of iconic representations can be illustrated with the help of maps. On maps, all of the pieces of information about the locations of objects are displayed simultaneously. The information that London is west of Berlin is simultaneously displayed on the map with the information that London is west of Warsaw. In the case of discursive representation, the information that London is west of Berlin does not contain the information that London is west of Warsaw. It must be inferred from the conjunction of the propositions Berlin is west of Warsaw and London is west of Belin. In contrast, a map displays information about all of the possible spatial relations simultaneously.

The non-hierarchical organisation of information processing is often confused with the holistic nature of the components of iconic structure. These two concepts, however, have to be separated, since we can have non-hierarchically organised processes based on non-holistic components. For instance, parallel computing is non-hierarchically organised and is based on discrete chunks of information. However, there is no clear account of what "holistic representation" means. There are at least three interpretations of this term.

First, the concept of the holistic nature of iconic representation is often interpreted as indicating the fact that iconic representations are informationally rich (Dretske, 1981; Kitcher, Varzi, 2000) and fine-grained (Tye, 2005). It means, first, that iconic content conveys so much information that it cannot plausibly be expressed with a finite set of propositions and, second, that iconic content is detailed and determined. In contrast, the content of discursive representations is general and abstract. For instance, seeing white snow is having an experience of a determined shade of white—thinking that snow is white does not determine the shade of the colour.

Although the concepts of information richness and being fine-grained seem intuitive, they are far from clear. First, there are discursive representations that are rich in content, such as the symbol π, and iconic representations that are informatively primitive, such as an image of a dot. Moreover, even if the richness of detail we are dealing with cannot plausibly be expressed by a finite set of propositions, it does not mean that it is impossible. A potentially infinite set of complex propositions can express any amount of information. Second, discursive content can be more fine-grained than iconic content. Pictures of aqua and cyan objects are often not detailed enough to see the difference between them; the propositions x is aqua and y is cyan are. Therefore, informative richness and fineness of grain do not determine whether we are dealing with iconic or discursive representations.

Second, we can interpret the concept of holistic representation as indicating the fact that pieces of information are entangled in a representation. For instance, Camp (2018) understands the holistic nature of information processing in maps

as a matter of structural linkage of the pieces of information. It means that changing the informational content of one of the map's elements changes the informational content of every other element. For example, moving the position of London on a map changes its distance to every other point on the map. In contrast, changing the informational content of *p* does not have to change anything about the content of *q*.

However, the problem is the scope of this thesis. Camp's argument certainly applies to information concerning spatial relations. However, it only says that spatial properties are relational, which is trivial. It does not apply to non-spatial information. For instance, changing the size of a circle representing London's population does not change the information about the size of a circle representing the population of Berlin.

Third, the holistic nature of representation can be interpreted (as it is interpreted here) as indicating the relation between the content and the vehicle of representation. To my knowledge, the first person to draw attention to this idea was Kazimierz Twardowski (1965), who ascribed the feature of concreteness to imaginings. He understood concreteness as the combination of multiple properties in a single representation. Similarly, the concept of holism is sometimes understood (e.g., Green, Quilty-Dunn, 2017; Kulvicki, 2020) as the thesis that multiple pieces of information expressing the content of a representation are assigned to the same vehicle of representation. It means that there is no separate vehicle for every chunk of information corresponding to different representational properties. In other words, there is no one-to-one correspondence between parts of the information and parts of the vehicle. For instance, the part of an icon that represents the colour of a triangle is the same part that represents its shape and location. Likewise, in the case of maps. The part of a topographic map that represents the location of London on the map represents its height above sea level too, along with a host of other things.

Thus, although iconic and discursive representations are described as expressing predicate functions, the way in which they process information is different. They express distinct predicative functions.

To sum up, iconic representations can be distinguished from discursive representations based on differences in the structure of information processing. The structure of iconic representations (or their format) is domain-specific; it processes information in a non-hierarchical manner and is based on holistic components. The structure of discursive representations is domain-general; it processes information hierarchically and is based on discrete chunks of information.

3. Canonical Decomposition and Lack of Logical Form

In the first section, I claimed that the distinction between different formats of representation should be able to provide at least a partial explanation of why iconic representations lack canonical decomposition and logical form, and why discursive representations are canonically decomposable and inferentially promiscuous. In

this section, I explain how the description of the representational structure I presented in the second section can help us to better address these questions.

First, Fodor's argument regarding the lack of canonical decomposition holds that iconic content cannot be decomposed into canonically distinguished parts. However, in this minimal form, the argument is clearly false, for the claim that iconic representations lack syntactic structure works as both a premise and the conclusion of the argument.

To fully present the structure of Fodor's argument, it is necessary to supplement it with a metaphysical premise concerning the format of iconic representation. According to the metaphysical premise, iconic representations have no structure linking the represented properties to the vehicle of representation, and therefore they cannot be canonically decomposed.

Let me illustrate the metaphysical premise with an image of a red square. The image of a red square can represent the content of the concept red or square, as well as the content of the proposition some squares are red. Yet, if we have a mental image of a red square, there is no way to determine whether we are thinking of the concept square or the content of the proposition some squares are red. Discursive representations can distinguish between concepts and structures composed of concepts, such as propositions. Iconic representations are unable to do so.

Likewise, in the case of maps. A map representing that Warsaw is east of Berlin represents Warsaw, Berlin and the fact that Warsaw is east of Berlin. We can distinguish between the representations Warsaw and Berlin and the representation Warsaw is east of Berlin only if we have the concept east. The concept east allows us to isolate the spatial relation property from all other represented properties. However, this means that we need a representation that can isolate a particular bit of information and assign it to the relevant representation vehicle. The discursive representation east does exactly that.

Iconic representations are unable to do so because there is no one-to-one correspondence between the information and its vehicle. As I claimed, the structure of iconic representations is based on holistic components. Multiple pieces of information are displayed on the same vehicle of representation. An image of a red square represents both redness and squareness, as well as the fact that some squares are red. In other words, iconic representations lack constituent structures and cannot be canonically decomposed, for there is no way to assign a single bit of information to a corresponding distinct vehicle of information.

Moreover, the domain-specificity of iconic representations renders them unable to specify what the canonical decomposition of iconic representations could be. Compare spatial and non-spatial iconic representations. Although we can cut a map into spatial pieces, not all iconic representations have spatial parts. Gustatory representations do not. For instance, the taste of a meal can be described as savory or sweet, but it is not dividable into spatial pieces. Therefore, there is no general way to decompose iconic representations.

In contrast, discursive representations can be decomposed canonically since their structure is domain-general. The proposition snow is white can be divided into canonically isolated parts regardless of whether it is expressed in English or in German.

Second, the systematicity of discursive representations means that they are inferentially promiscuous (Peacocke, 1992; Stich, 1978). For instance, beliefs and thoughts can figure as premises in inferences. Discursive representations have a logical form and can be modelled according to the rules of logic. Yet, according to Frege's argument, iconic representations lack logical form. They cannot be inferred or negated. They are not inferentially promiscuous. Why is that so?

Inferential promiscuity requires propositional structure. It involves a kind of relation between the vehicle of representation and the content. To infer from a is F and if a is F, then a is G that a is G, one has to be able to assign distinct contents to the vehicles of representation expressed by logical variables. Moreover, the chunks of information have to be hierarchically organised. From the thought that a is G and that if a is F, then a is G I cannot infer that a is F. From the thought that snow is white I can infer that it has a colour, but from snow has a colour I cannot infer that it is white. Discursive representations are hierarchically organised.

In contrast, images are organised non-hierarchically. The information is processed simultaneously. When I see a picture of white snow, I see that it has a colour; when I see a colourful picture, I can see the specific colour of the picture.

Moreover, inferential promiscuity requires a representational structure that can abstract from the nature of the vehicle of representation. For instance, the reasoning if A then B and A then B is correct regardless of whether the reasoning is conducted in the mind or on paper. Discursive representations meet this requirement since they are domain-general. In contrast, iconic representations are domain-specific. The nature of the vehicle of iconic representation affects the way information is processed. For instance, tasting wine and imagining tasting wine are informationally two different representations.

To sum up, iconic representations lack syntactic structure and do not meet the requirements of the Generality Constraint. They are neither systematic nor canonically decomposable. These facts are easier to understand if we hold that iconic representations are domain-specific, that they process information in a non-hierarchical fashion, and that their structure is based on holistic components. In contrast, discursive representations are domain-general, they process information hierarchically, and their structure is based on discrete elements. Thus, discursive representations are systematic and canonically decomposable.

The argument presented here does not show that the distinction between analog and digital representational systems is useless. However, it demonstrates that this distinction is insufficient for distinguishing perception and thought, for it does not provide any explanation of the source of the differences between iconic and discursive representations. In contrast, thinking of iconic and discursive representations in terms of the way they structure information helps us to better understand why they differ. According to the view presented here, the different

functional properties of iconic and discursive representations follow from different informational structures.

REFERENCES

Barrouillet, P., Grosset, N., Lecas, J.-F. (2000). Conditional Reasoning by Mental Models: Chronometric and Developmental Evidence. *Cognition*, *75*(3), 237–266.

Beck, J. (2015). Analogue Magnitude Representations: A Philosophical Introduction. *The British Journal for the Philosophy of Science*, *66*(4), 829–855.

Blumson, B. (2012). Mental Maps. *Philosophy and Phenomenological Research*, *85*(2), 413–434.

Braddon-Mitchell, D., Jackson, F. (1996). *Philosophy of Mind and Cognition*. Oxford: Blackwell.

Byrne, R. M. J. (2005). *The Rational Imagination: How People Create Counterfactual Alternatives to Reality*. Cambridge, MA: MIT Press.

Byrne, R. M. J., Johnson-Laird, P. N. (1989). Spatial Reasoning. *Journal of Memory and Language*, *28*, 564–575.

Camp, E. (2004). The Generality Constraint and Categorial Restrictions. *Philosophical Quarterly*, *54*(215), 209–231.

Camp, E. (2007). Thinking with Maps. In J. Hawthorne (Ed.), *Philosophical Perspectives 21: Philosophy of Mind* (pp. 145–182). Oxford: Wiley-Blackwell.

Camp, E. (2018). Why Maps Are Not Propositional. In A. Grzankowski, M. Montague (Eds.), *Non-Propositional Intentionality* (pp. 19–45). Oxford: Oxford University Press.

Casati, R., Giardino, V. (2013). Public Representation and Indeterminacies of Perspectival Content. In Z. Kondor (Ed.), *Enacting Images* (pp. 111–126). Köln: Herbert von Halem Verlag.

Chambers, D., Reisberg, D. (1985). Can Mental Images Be Ambiguous? *Journal of Experimental Psychology: Human Perception and Performance*, *11*(3), 317–328.

Crane, T. (2009). Is Perception a Propositional Attitude? *The Philosophical Quarterly*, *59*, 452–469.

Devitt, M. (2006). *Ignorance of Language*. Oxford: Oxford University Press.

Dretske, F. (1969). *Seeing and Knowing*. Chicago: University of Chicago Press.

Dretske, F. (1981). *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press.

Evans, G. (1982). *The Varieties of Reference*. Oxford: Oxford University Press.

Fodor, J. (2007). The Revenge of the Given. In B. P. McLaughlin, J. D. Cohen (Eds.), *Contemporary Debates in Philosophy of Mind* (pp. 105–116). Oxford: Basil Blackwell.

Fodor, J. (2008). *LOT 2: The Language of Thought Revisited*. Oxford: Oxford University Press.

Fodor, J., Pylyshyn, Z. (1981). How Direct Is Visual Perception? Some Reflections on Gibson's 'Ecological Approach'. *Cognition*, *9*, 207–246.
Fodor, J., Pylyshyn, Z. (1988). Connectionism and Cognitive Architecture: A Critical Analysis. *Cognition*, *28*(1–2), 3–71.
Frege, G. (1984). Thoughts. In B. McGuinness (Ed.), *Collected Papers on Mathematics, Logic, and Philosophy* (pp. 351–372). Oxford: Basil Blackwell.
Goodman, N. (1976). *Languages of Art* (2nd Ed.). Indianapolis: Hackett.
Green, E. J., Quilty-Dunn, J. (2017). What Is an Object File? *The British Journal for the Philosophy of Science*. doi:10.1093/bjps/axx055
Haugeland, J. (1998). *Having Thought: Essays in the Metaphysics of Mind*. Cambridge: Harvard University Press.
Heck, R. G. (2000). Nonconceptual Content and the 'Space of Reasons'. *The Philosophical Review*, *109*, 483–523.
Heck, R. G. (2007). Are There Different Kinds of Content? In B. P. McLaughlin, J. Cohen (Eds.), *Contemporary Debates in Philosophy of Mind* (pp. 117–138). Oxford: Blackwell.
Hegarty, M., Waller, D. A. (2005). Individual Differences in Spatial Abilities. In P. Shah, A. Miyake (Eds.), *The Cambridge Handbook of Visuospatial Thinking* (pp. 121–169). New York: Cambridge University Press.
Hintikka, J. (1987). Mental Models, Semantical Games, and Varieties of Intelligence. In L. Vaina (Ed.), *Matters of Intelligence: Conceptual Structures in Cognitive Neuroscience* (pp. 197–215). Dordrecht: D. Reidel.
Hinton, G. (1979). Some Demonstrations of the Effects of Structural Descriptions in Mental Imagery. *Cognitive Science*, *3*(3), 231–250.
Ittelson, W. H. (1996). Visual Perception of Markings. *Psychonomic Bulletin & Review*, *3*(2), 171–187.
Johnson, K. (2015). Maps, Languages, and Manguages: Rival Cognitive Architectures? *Philosophical Psychology*, *28*(6), 815–836.
Johnson-Laird, P. N. (1983). *Mental Models: Towards a Cognitive Science of Language, Inference and Consciousness*. Cambridge: Harvard University Press.
Johnson-Laird, P. N. (1998). Imagery, Visualization, and Thinking. In J. Hochberg (Ed.), *Perception and Cognition at the Century's End* (pp. 441–467). San Diego: Academic Press.
Kamermans, K. L., Pouw, W., Mast, F. W., Paas, F. (2019). Reinterpretation in Visual Imagery Is Possible Without Visual Cues: A Validation of Previous Research. *Psychological Research: An International Journal of Perception, Attention, Memory and Action*, *83*(6), 1237–1250.
Kitcher, P., Varzi, A. (2000). Some Pictures are Worth 2[aleph]0 Sentences. *Philosophy*, *75*(3), 377–381.
Kosslyn, S. M. (1980). *Image and Mind*. Cambridge, MA: Harvard University Press.
Kosslyn, S. M. (1994). *Image and Brain: The Resolution of the Imagery Debate*. Cambridge, MA: MIT Press.

Kosslyn, S. M., Brunn, J., Cave, K. R., Wallach, R. W. (1984). Individual Differences in Mental Imagery Ability: A Computational Analysis. *Cognition*, *18*, 195–243.

Kozhevnikov, M., Blazhenkova, O., Becker, M. (2010). Trade-off in Object Versus Spatial Visualization Abilities: Restriction in the Development of Visual-Processing Resources. *Psychonomic Bulletin & Review*, *17*(1), 29–35.

Knauff, M., Johnson-Laird, P. N. (2002). Visual Imagery Can Impede Reasoning. *Memory and Cognition*, *30*, 363–371.

Kulvicki, J. (2006). *On Images: Their Structure and Content*. Oxford: Clarendon Press.

Kulvicki, J. (2020). *Modelling the Meanings of Pictures: Depiction and the Philosophy of Language*. Oxford: Oxford University Press.

Larkin, J. H., Simon, H. A. (1987). Why a Diagram Is (Sometimes) Worth Ten Thousand Words. *Cognitive Science*, *11*, 65–99.

Lewis, D. (1971). Analog and Digital. *Noûs*, *5*, 321–328.

Lorenzo, G., Rubiera, E. (2019). On Iconic-Discursive Representations: Do They Bring Us Closer to a Humean Representational Mind? *Biosemiotics*, *12*, 423–439.

Maley, C. J. (2011). Analog and Digital, Continuous and Discrete. *Philosophical Studies*, *155*, 117–131.

Mast, F. W., Kosslyn, S. M. (2002). Visual Mental Images Can Be Ambiguous: Insights From Individual Differences in Spatial Transformation Abilities. *Cognition*, *86*, 57–70.

Matthen, M. (2005). *Seeing, Doing, and Knowing: A Philosophical Theory of Sense Perception*. Oxford: Clarendon Press.

McDowell, J. (1996). *Mind and World*. Cambridge, MA: Harvard University Press.

McGinn, C. (1989). *Mental Content*. Oxford: Blackwell Publishers.

Pagin, P., Westerståhl, D. (2010). Compositionality I: Definitions and Variants. *Philosophy Compass*, *5*, 250–264.

Peacocke, C. (1989). Perceptual Content. In J. Almog, J. Perry, H. Wettstein (Eds.), *Themes from Kaplan* (pp. 297–329). New York: Oxford University Press.

Peacocke, C. (1992). *A Study of Concepts*. Cambridge, MA: MIT Press.

Prinz, J. J. (2002). *Furnishing the Mind. Concepts and their Perceptual Basis*. Cambridge, MA: MIT Press.

Pylyshyn, Z. W. (1973). What the Mind's Eye Tells the Mind's Brain: A Critique of Mental Imagery. *Psychological Bulletin*, *80*, 1–25.

Pylyshyn, Z. W. (1981). The Imagery Debate: Analogue Media Versus Tacit Knowledge. *Psychological Review*, *88*(1), 16–45.

Quilty-Dunn, J. (2016). Iconicity and the Format of Perception. *Journal of Consciousness Studies*, *23*(3–4), 255–263.

Quilty-Dunn, J. (2020). Perceptual Pluralism. *Noûs*, *54*(4), 807–838.

Reisberg, D. (1996). The Nonambiguity of Mental Images. In C. Cornoldi, R. H. Logie, M. A. Brandimonte, G. Kaufmann, D. Reisberg (Eds.), *Stretching the Imagination: Representation and Transformation in Mental Imagery* (pp. 119–172). New York: Oxford University Press.

Reisberg, D., Heuer, F. (2005). Visuospatial Images. In P. Shah, A. Miyake (Eds.), *The Cambridge Handbook of Visuospatial Thinking* (pp. 35–80). New York: Cambridge University Press.

Rescorla, M. (2009). Predication and Cartographic Representation. *Synthese*, *169*(1), 175–200.

Sainsbury, R. M. (2005). *Reference without Referents*. Oxford: OUP.

Sartre, J.-P. (1962). *Imagination: A Psychological Critique*. Ann Arbor: University of Michigan Press.

Sellars, W. (1997). *Empiricism and the Philosophy of Mind*. Harvard: Harvard University Press.

Slezak, P. (1995). The 'Philosophical' Case Against Visual Imagery. In P. Slezak, T. Caelli, R. Clark (Eds.), *Perspectives on Cognitive Science: Theories, Experiments and Foundations* (pp. 237–271). Norwood: Ablex Publishing.

Sober, E. (1976). Mental Representations. *Synthese*, *33*, 101–148.

Stich, S. P. (1978). Beliefs and Subdoxastic States. *Philosophy of Science*, *45*(4), 499–518.

Szabo, Z. (2012). The Case for Compositionality. In M. Werning, W. Hinzen, E. Machery (Eds.), *The Oxford Handbook of Compositionality* (pp. 64–80). Oxford: Oxford University Press.

Travis, C. (2013). *Perception—Essays After Frege*. Oxford: Oxford University Press.

Twardowski, K. (1965). Wyobrażenia i pojęcia. In K. Twardowski, *Wybrane Pisma Filozoficzne* (pp. 114–197). Warsaw: PWN.

Tye, M. (2005). Non-Conceptual Content, Richness, and Fineness of Grain. In T. Gendler, J. Hawthorne (Eds.), *Perceptual Experience* (pp. 504–526). Oxford: Oxford University Press.

Von Neumann, J. (1958). *The Computer and the Brain*. New Haven: Yale University Press.

Westerhoff, J. (2005). Logical Relations between Pictures. *Journal of Philosophy*, *102*(12), 603–623.