The Two Visual Systems Hypothesis and Contrastive Underdetermination

Grünbaum, Thor

Published in:
Synthese

DOI:
10.1007/s11229-018-01984-y

Publication date:
2018

Document version
Peer reviewed version

Final draft. Accepted 2018 in *Synthese*.
Please quote the published version: https://doi.org/10.1007/s11229-018-01984-y

The Two Visual Systems Hypothesis and Contrastive Underdetermination

By Thor Grünbaum

Abstract

This paper concerns local yet systematic problems of contrastive underdetermination of model choice in cognitive neuroscience debates about the so-called two visual systems hypothesis. The underdetermination problem is systematically generated by the way certain assumptions about the representationalist nature of computation are translated into experimental practice. The problem is that behavioural data underdetermine the choice between competing representational models. In this paper, I diagnose how these assumptions generate underdetermination problems in the choice between competing functional models of perception– action. Using the tools of philosophy of science, I describe the type of underdetermination and sketch a possible cure.

1. Introduction

The notion of contrastive underdetermination is central to a number of different discussions and disciplines. In philosophy of science, the possibility of global forms of underdetermination of theories by data plays a pivotal role in arguments for and against scientific anti-realism. In experimental sciences, say, in experimental psychology, the possibility of local and transient forms of underdetermination has an important motivating function. The goal of designing new experiments will often be to overcome some state of underdetermination by available data of the choice between two or more competing models. If it turns out the choice between two or more experimental hypotheses or models is more permanently underdetermined, experimentalists tend to "despair" and discard the choice as ill-defined or experimentally unconstrained.[1]

This paper concerns a particular problem of contrastive underdetermination of model choice in cognitive neuroscience. I will identify and describe a basic problem of model underdetermination in research on visuomotor control. The problem is local but systematic. It is caused by a set of assumptions about the representationalist nature of computation in the brain. This problem of underdetermination is confined to domains of research where such assumptions are made. As will

[1] This seems to have been the situation in the 70-80s debates about the choice between serial and parallel processing models of attention and executive control (see Townsend, 1990). The same kind of "despair" was manifest in the mental imagery debate in the 90s before the boom of brain imaging techniques (Ganis & Schendan, 2011). This also seems to have been the situation in the debates between theory-theory and simulation-theory about mind-reading in social cognitive neuroscience (Apperly, 2008).

become apparent, the underdetermination problem is generated by the way certain assumptions about the representationalist nature of computation are translated into experimental practice. Looking more closely at this translation allows us to see a way to avoid some of the underdetermination problems. We can avoid adopting an anti-realist stance with respect to computationalist models as well as escape the experimentalist's "despair".

I will proceed as follows. First, I frame the discussion of underdetermination problems in cognitive neuroscience by introducing underdetermination arguments from philosophy of science. The aim is to pinpoint relevant forms of underdetermination and some of their possible roles in arguments for anti-realism. In section 3, I introduce our target debate in cognitive neuroscience, the debate concerning visual systems for manual action and visual cognition. In section 4, I argue that representationalist assumptions about the nature of computational mechanisms in the brain affect experimental practise in our target debate. Section 5 turns to the local problem of contrastive underdetermination. I argue that the choice between computationalist models in our target debate is systematically underdetermined by behavioural data. In section 6, returning to the general types of underdetermination, I discuss the nature and scope of our local underdetermination problem. In section 7, I suggest a possible cure by introducing a new kind of experimental paradigm (Christiansen et al., 2014). Using this strategy, my aim is not only to say something about this particular debate in cognitive neuroscience. The more ambitious hope is that by diagnosing the features generating the underdetermination in this case, we can learn a more general lesson about underdetermination of models by data in cognitive neuroscience. [2]

[2] The present paper is related to but is importantly different from the argument in Grünbaum (2017). As in my (2017) paper, the present paper points to problems of empirically distinguishing between two different functional models of visuomotor processing. Grünbaum (2017) located the cause of the problems in a set of metaphysical assumptions about the individuation of computational mechanisms. The present paper focuses on a different set of problematic methodological assumptions. The present paper locates parts of the cause of the underdetermination problem in the way in which certain computationalist assumptions are translated into experimental practice. Furthermore, where the 2017-paper placed the discussion in the context of philosophical arguments about consciousness, the present paper places the discussion in the broader context of model-choice and underdetermination by data in philosophy of science. Seen in this context, the importance of the contribution in the present paper is threefold: 1) I provide a new analysis of common methodological assumptions in cognitive psychology; 2) I use insights from philosophy of science to describe forms of underdeterminations and their consequences for model choice in a domain of cognitive neuroscience; and 3) I use the empirical literature surrounding the TVSH to describe a form of underdetermination often overlooked by philosophers of science (local but systematic underdetermination).

Final draft. Accepted 2018 in *Synthese*.
Please quote the published version: https://doi.org/10.1007/s11229-018-01984-y

2. Forms of underdetermination

Underdetermination of scientific theories by empirical data provides a challenge to scientific realism: If empirical data underdetermine which theory of a set of mutually conflicting scientific theories is the correct one, then we are not justified in believing any of the theories' claims about unobservable entities. For underdetermination to be a threat to scientific realism, it has to be global (Earman, 1993; Hoefer & Rosenberg, 1994). It would have to be the case that for any actual or possible scientific theory, data underdetermine whether to accept the theory as true (or at least approximately true). Local and practical forms of underdetermination are familiar to scientists (Kitcher, 2001, 30). Practical underdetermination is a challenge to be surpassed in the progress of experimental science. Local and practical forms of underdetermination pose no threat to scientific realism (Norton, 2008).[3]

In the literature on underdetermination arguments, it is customary to distinguish between two forms of underdetermination of theories by observational scientific data: 1) Holistic underdetermination and 2) contrastive underdetermination.[4]

Holistic underdetermination arguments are well-known from Duhem (1962) and Quine (1951). According to their view, scientific confirmation is always holistic in nature. The core claims of a theory can only generate predictions when conjoined with a set of auxiliary hypotheses. Consequently, if the predictions are disconfirmed, it is impossible to know whether the error is located in the core theory or the auxiliary hypotheses. If that is true, then the choice between accepting and rejecting any scientific theory will always be underdetermined by data. In the following, I will set aside holistic underdetermination for a number of reasons. First, proponents of this argument assume that the complete network of core theory and auxiliary hypotheses is being tested in any single test. However, auxiliary hypotheses have usually been tested independently (Sober, 1999). Second, often scientists could not rationally reject auxiliary hypotheses without undermining the scientific enterprise itself (Laudan, 1990). Third, the case of interest in the present paper is best categorized as a contrastive form of underdetermination.

Contrastive underdetermination is the situation in which two rival theories that have equal prior probability assign equal probability to the same set of scientific data. In a situation like this, the scientific data are evidentially irrelevant to the choice between the theories because the data will boost (or not boost) the probability of each model to the same degree. In such a situation, the rival theories can be said to be equally supported by the set of scientific data and the choice between such rival

[3] Though practical underdetermination might be unproblematic for the scientific realist, it is by no means a trivial or theoretically uninteresting phenomenon. For some discussion, see Biddle, 2013.
[4] The terminology is from Stanford, 2017

Final draft. Accepted 2018 in *Synthese*.
Please quote the published version: https://doi.org/10.1007/s11229-018-01984-y

theories is underdetermined by the data. Consequently, we are not justified in accepting the existence of unobservable entities posited by any of the rival theories.

It has turned out to be hard to come up with convincing arguments for global forms of contrastive underdetermination. Probably, the most promising version of a contrastive underdetermination argument against scientific realism focuses on actual cases from the history of science.[5] Generally, the argument works by convincing us that important cases demonstrate contrastive underdetermination and then inductively showing us that this is the general condition of science.[6] Consequently, we would never be justified in accepting the existence of the unobservable entities posited by any particular scientific theory.

A number of important problems has been raised for this kind of contrastive underdetermination argument (Dellsén, 2017; Godfrey-Smith, 2008). Here I will focus on two. First, when looking closer at the actual cases of underdetermined rival theories, they are usually not cases of empirical equivalence but rather cases of being equivalent with respect to one particular type of data. Consequently, the theories are not equivalent with respect to all their possible observational consequences but only with respect to one particular type of data (Worrall, 2011). For instance, rival psychological theories of mental imagery might be "data equivalent" relative to behavioural data without being equivalent with respect to all their possible observational predictions (Anderson, 1978). Second, some cases of contrastive underdetermination in the history of science turn out to be cases where the rival theories are so close in their theoretical structure that we cannot "preclude the possibility that they are variant formulations of the same theory" (Norton, 2008, 33). Thus, the rival theories posit the existence of the same unobservable entities.

Following Worrall (2011), we can conclude that a successful version of the contrastive underdetermination argument must at least "establish not only (i) for any accepted scientific theory there is always another that is 'equally empirically successful' but also (ii) that there is some reason why the realist could not reasonably regard the alternative as 'approximately true' just like the

[5] See, Stanford (2001, 2006) and Tulodziecki (2013). Stanford convincingly argues that the algorithmic version (Kukla, 1996) of contrastive underdetermination is a form of global skepticism and is too remote from the practice of actual science. I therefore set aside this kind of contrastive underdetermination. In recent discussions of underdetermination of psychological models by fMRI data, Loosemore and Harley (2010) present a close cousin of Kukla's argument. They invent a toy model and argues that it and standard models are equivalent relative fMRI data. Bechtel and Richarson (2010) rightly argues that the toy model used by Loosemore and Harley is a purely imaginary construct and has nothing to do with real science. Hence, the underdetermination of the choice between the toy model and the standard psychological models has no consequence for the standard models and the use of fMRI data.

[6] Important cases include theory choices concerning formalism, ontology and methodology in quantum mechanics (Belousek, 2005), the absolute state of rest in Newtonian physics (van Frassen, 1980, 46-47), and models in cosmology (Butterfield, 2012).

Final draft. Accepted 2018 in *Synthese*.
Please quote the published version: https://doi.org/10.1007/s11229-018-01984-y

accepted theory" (2011, 165). "Equally empirically successful" means that the theories are not just data equivalent relative to a specific kind of data but are equivalent with respect all possible observational predictions.

Only very few cases (if any at all) will be able to satisfy Worrall's condition concerning empirical equivalence. This makes the condition less interesting when trying to understand the structure of actual scientific disputes about models. In what follows, I will use not Worrall's strict notion of empirical equivalence but a more relaxed notion of *comprehensive data equivalence*. Two rival theories are comprehensively data equivalent if all currently available forms of data equally support the theories and we have no empirical evidence supporting the claim that decisive data of a new kind is forthcoming in a foreseeable future. This allows that the theories could differ with respect to some possible observational consequences.

I will assume that comprehensive data equivalence of two rival theories can justify adopting some form of *epistemic* anti-realism with respect to at least some unobservable entities posited by the two theories. If the rival theories are comprehensively data equivalent and there is no reason to think that both theories could be "approximately true", then the right attitude towards at least some of the unobservable entities would be agnosticism. If the contrasted theories are comprehensively data equivalent and could not both be approximately true, we have no basis on which to judge whether some of the unobservable entities posited by the theories exist.

In what follows, I will not be concerned with global forms of realism and underdeterminism. I will be concerned with local forms of realism relative to particular kinds of entities, viz. the computational mechanism involved in visuomotor control and visual cognition. The aim will be to show that there are local yet systematic (hence "quasi-permanent") forms of contrastive underdetermination that threaten realism with respect to computational mechanism in a domain of cognitive neuroscience. This form of underdetermination occupies a logical space in between the philosophers' strong underdetermination and the scientitists' practical underdetermination. It is a strong form of underdetermination because it flows from core assumptions made by scientists. Given certain core theoretical assumptions about the nature of computation and methodological assumptions inherent in experimental practice, the rival theories in the domain will always be comprehensively data equivalent. Having established this fact, the question then becomes whether the empirical deadlock can be broken by rejecting some of the core assumptions.

Final draft. Accepted 2018 in *Synthese*.
Please quote the published version: https://doi.org/10.1007/s11229-018-01984-y

3. TVSH, the 2-systems model, and the 1-system model

The *Two Visual Systems Hypothesis (TVSH)* is a hypothesis concerning the visual systems driving visuomotor control and visual cognition (Goodale & Milner, 1992; Milner & Goodale, 2006). Debates about this hypothesis are particularly instructive for our present concerns for a number of reasons. First, it is a clear case where competing functional models differ with respect to types and numbers of computational machinery. Second, as in cognitive psychology generally, researchers involved in this debate use results from behavioural experiments to make inferences about underlying representational formats and types of computational mechanisms. Third, the TVSH is a provocative yet appealing hypothesis, which has had a tremendous influence in both cognitive neuroscience and philosophy of mind.[7]

TVSH consists of three independent but tightly connected models:

I. Anatomical model: Neural projections from visual area V1 are divided into two separate neural streams. Even though researchers accept cross-talk between the brain systems, "the distinction between a dorsal stream projecting to the posterior parietal lobule and a ventral stream projecting to the inferotemporal cortex has been repeatedly confirmed" (Milner & Goodale, 2006, 39-40).[8]

II. Functional model: The two separate anatomical projections from V1 are the neural substrates of two functionally independent visual systems: The dorsal stream is a system for visual control of manual action and the ventral system is a system for visual control of cognition (judgement, planning, and memory). The two systems are independent of each other in the sense that each system embodies computational mechanisms that can be doubly dissociated from each other and are operating by different kinds of processing principles.

[7] Concerning TVSH's influence on cognitive neuroscience, see substantial review chapters in prominent handbooks (Goodale, 2014; Goodale & Ganel, 2015; Westwood, 2009). Concerning the influence on philosophy, see debates about the nature of perception and its content (Briscoe, 2008; Ferretti, 2018; Matthen, 2005; Wu, 2014), the nature and function of consciousness (Brogaard, 2011; Clark, 2001), the nature of action and control (Clark, 2007; Kozuch, 2015; Mole, 2008; Wallhagen, 2007; Wu, 2013), the nature and status of folk-psychology (Bermudez, 2006; Grünbaum, 2012), and the interface problem (Butterfill & Sinigaglia, 2014; Mylopoulos & Pacherie, 2016; Shepherd, 2017)

[8] The anatomical model is endorsed by proponents of TVSH (Ganel & Goodale, 2017). At present, it is hard to tell whether they are right to do so. Even if the strict anatomical division into separate processing streams has been challenged recently (see, for instance, Galletti & Fattori, 2018; Rossetti, Pisella, & McIntosh, 2017), the anatomical model is supported by data from anatomical, electrophysiological, and psychophysical studies mainly with monkeys (for reviews, see Culham & Valyear, 2006; Kravitz et al., 2011).

III. NCC model: Only the ventral stream is the neural correlate of visual consciousness (NCC). Visual processing in the dorsal stream system is non-conscious.

This paper is concerned with disputes over the choice of the functional model (model II). According to the functional model of TVSH, "the set of object descriptions that permit identification and recognition may be computed independently of the set of descriptions that allow an observer to shape the hand appropriately to pick up an object" (Goodale & Milner, 1992, 20). By "object descriptions" the authors mean representations of features such as size, location, and orientation. This model claims that there are two independent mechanisms, one for computing the object descriptions driving grasping and another for computing the object descriptions driving visual cognition. Call this functional model *the 2-systems model*.

The canonical way to understand the functional 2-systems model of TVSH is with respect to the location and job description of the mechanisms computing the object descriptions. The two mechanisms are located downstream from the anatomical split. Both computational mechanisms receive the same information about the same object features, but "each stream uses this visual information in different ways" (Goodale & Haffenden, 1998, 162). The ventral mechanism delivers representations that can play a role in "the identification of objects and enable us to classify objects and events, attach meaning and significance to them, and establish their causal relations" (Goodale & Haffenden, 1998, 162). In what follows, I will for convenience lump identification, categorisation, judgement, etc. under the general label *cognition*. The job of the ventral stream is to compute visual representations for use in cognitive tasks. By contrast, the dorsal mechanism "deal[s] with moment-to-moment information about the location and disposition of objects in egocentric coordinates and thereby mediate the visual control of skilled actions, such as manual prehension, directed at those objects" (Goodale & Haffenden, 1998, 162). In what follows, I will mainly focus on grasping. The job of the dorsal stream is to compute visual representations for manual action.

This functional model builds on the assumption that cognitive tasks and grasping tasks are two radically different forms of task. When identifying and categorising an object, I rely on long-term representations of the object. The object stays the same irrespective of its retinal size, and its location and orientation in space with respect to the subject. By contrast, when grasping an object, I rely on dynamic representations that are sensitive to changes in orientation and location, and the object features must be represented in such a way as to reflect the nature of the "effector system". One way to spell out these differences is to say that visual cognition requires that the object features are encoded in an *allocentric* or scene-based representational format, whereas grasping requires that the features are encoded in an *egocentric* or effector-relative representational format (see Foley, Whitwell, & Goodale, 2015). Consequently, despite receiving similar retinal information about the environment, the dorsal and ventral mechanisms are independent of each other because they have different jobs to do: the dorsal mechanism has the job of encoding object features in an egocentric format and the ventral mechanism the job of encoding them in an allocentric format. To take the example of size, the ventral system encodes information about the object's size in a relational metric representing the feature relative to sizes of other objects, whereas the dorsal system is supposed to encode the information using a metric that represent its size relative only to the effector (Goodale & Ganel, 2015).

It is important to notice that both egocentrically encoded and allocentrically encoded object features are visual representations. The egocentric representation is not itself a motor instruction. The thought is that because it has an egocentric format, it is easily and accurately transformed into a motor instruction executable by the motor system. Similarly, in cognitive tasks, allocentric representations have to be transformed into various types of behaviours (e.g. manual size estimation, line drawing, choosing figures, or verbal reports). Again, the idea would be that these cognitive tasks are more easily and accurately subserved by visual representations in an allocentric format.

The 2-systems model of TVSH is not the only plausible functional perception-action model. Various researchers have proposed instead a *1-system model* according to which object descriptions are computed by common mechanisms located before an anatomical split into two separate streams. On this 1-system model, there are common mechanisms for computing representations of an object's features, say, its size. These common representations are subsequently sent along both processing streams, where final transformations into behaviour take place. Proponents of the 1-system model deny that visual cognitive tasks and manual action are radically different forms of task requiring radically different forms of visual processing. The most common version of the 1-system model describes the brain as encoding common visual representations for grasping and judging in dynamic and task-dependent ways.[9]

It is important to notice that both models agree on a number of important points. First, arguably, they are both consistent with the same anatomical model. Assuming this consistency, the two models compare in the following way. The 1-system model claims that the mechanism computing

[9] A 1-system model of perception-action has been defended by Franz et al., 2000, Franz and Gegenfurtner, 2008, Christiansen et al., 2014, De la Mala, Smeets, & Brenner, 2018, and Medendorp, De Brouwer, and Smeets, 2018.

visual representations of object features (such as size, shape and orientation) is located before the anatomical split and sends a common representation of object features along dorsal and ventral streams. By contrast, the 2-systems model claims that two independent mechanisms located after the anatomical split compute differently encoded representations of object features for action (dorsal mechanism) and cognition (ventral mechanism).

Second, irrespective of whether we compare the single mechanism of the 1-system model with the mechanisms of the 2-systems model or the two mechanisms of 2-systems model with each other, the input to the mechanisms is always the same. That is, the input is always the ‘raw’ visual information of V1 before extensive visual processing.

Third, both models agree that important and often complex transformational operations occur when the visual representations are transformed into motor instructions or various forms of cognition. Thus, on both models, we need to distinguish between *core computational processes* belonging to the computational mechanisms of the functional model and *transformational processes* translating into actual behaviour.

4. Methodological principles: Inferring representational formats from data

The choice between 2-systems and 1-system models is determined by deciding whether we have evidence for positing one common or two independent computational mechanisms. Let us assume that neural computational mechanisms are individuated by their representational formats.[10] That is, the computational mechanism for encoding the visual representations driving grasping are distinguished from the computational mechanism encoding the visual representations driving perceptual judgements by their respective formats of encoding. The computational mechanisms take the same input, namely, the “raw” visual information from V1. Had the computational mechanisms for manual action and visual cognition produced the same output, that is, had they encoded the input information in the same format, they would have been the same kind of computational mechanism performing the same job. If that were the case, there would have been no reason for postulating two independent mechanisms. The epistemic situation therefore seems to be this: we need to infer the nature and number of computational mechanisms from facts about the formats of the encoding of visual information. The question is now: how do we infer the facts about the involved representational formats?

[10] This claim is controversial. See Piccinini, 2015. It is generally assumed in the target debate under consideration. See Grünbaum, 2017.

Final draft. Accepted 2018 in *Synthese*.
Please quote the published version: https://doi.org/10.1007/s11229-018-01984-y

The standard answer in experimental psychology is that we can infer the underlying format of the information driving the performance from systematic response profiles. Take the example of mental imagery and the debate between the "pictorial" and the "propositional" accounts. One influential form of reasoning was that if the computational mechanism subserving mental imagery computes over "pictorial" representations, then we should expect mental analogues of rotation and scanning. That is, we should expect the response time of the performance of certain tasks to be a linear function of the degrees of angular rotation of inner mental "pictures" or the distances between locations on inner mental "maps". By contrast, if the computational mechanisms compute over "propositional" representations (using some language of thought), then we should not expect a linear relation between, on the one hand, degrees of angular rotation or distances between locations and, on the other hand, response time (for a discussion of this experimental logic, see Kosslyn, Ball, & Reiser, 1978; Pylyshyn, 1981).

There are two important points to make about such cases. First, confirmation of a model by experimental data is usually contrastive, and second, data only serve as evidence for a particular kind of representational model (say, a "pictorial" model) given background assumptions about cognitive processes. Let me discuss these two points in turn.

Confirmation is contrastive if data that serve as evidence for a given hypothesis H_1 are by the same token serving as evidence against some competing hypothesis H_2. Hypotheses or models are not confirmed or tested in isolation. If we were to assess in complete isolation from alternative hypotheses whether some data confirm H_1, we would need to know whether the data boost the probability of H_1. In Bayesian terms, in order to make this assessment, we would need to know the value of the prior probability of H_1, the probability of data conditional on the truth of H_1 (the likelihood), and the unconditional probability of the data. The acquisition of these types of knowledge is confronted with a set of well-known problems (see, e.g., Sober, 2008, Ch. 1).

It is important to notice that, even if we cannot determine the numerical values of the prior probability of H_1 and the unconditional probability of the data, we might be able to make qualitative contrastive estimates of likelihoods of the competing hypotheses.[11] We might not be able to assess the numerical probabilities of the "pictorial" and the "propositional" models in isolation, but we might still be able to assess in qualitative terms whether some experimental outcome would be more probable according to one model than it would be according to the other. Given a certain set of background assumptions about the analogies between internal "pictures" and external pictures, the

[11] For discussion, see Salmon, 1990; Sober, 1999, 2008, Ch. 1.

Final draft. Accepted 2018 in *Synthese*.
Please quote the published version: https://doi.org/10.1007/s11229-018-01984-y

types of operation that can be performed on pictures, and the types of operation that can be performed on propositions, certain experimental outcomes could be regarded as more probable conditional on the "pictorial" model than on the "propositional" model (Kosslyn, Ball, & Reiser, 1978; Kosslyn, 1994).

Turning to our target debate about the number of computational mechanisms in the correct perception-action model, we see the application of the above points. Even if we are not able to assign a numerical value to the probability of TVSH in isolation, we might still be able to assess in a more qualitative and contrastive way the probability of some experimental outcome given TVSH (the 2-systems model) and given the 1-system model. This kind of contrastive model testing can only take place on the background of a set of shared background assumptions. It is probably not easy to spell out exhaustively exactly what these shared background assumptions are. Here I will focus on a set of beliefs concerning the inference from experimental data to representational formats.

A general and common assumption is that types of representational format can be inferred from systematic differences in performance in experimental tasks. In the Ebbinghaus illusion, central circles look larger or smaller than they really are depending on the size of the surrounding circles. If you have two circles of identical size but one is surrounded by small satellites and the other by big satellites, then the first circle will normally look larger than the second. In a seminal study by Aglioti and colleagues (1995), the central circles were physical discs that participants could pick up. By comparing participants' size judgements with their maximal grip aperture when grasping one of the central discs, Aglioti et al. found a striking difference between judgement and grasping, where the effect of the size-contrast illusion on grasping was "significantly smaller […] and more variable" (1995, 683) than the effect on perceptual judgements. The authors used the data as evidence for the claim that "accurate [grip] calibrations required for skilled actions may depend on visual computations that are different from those driving perceptual judgements" (1995, 684). More specifically, this result "has been interpreted as evidence for the idea that vision-for-action makes use of real-world metrics while vision-for-perception uses relative or scene-based metrics" (Goodale & Ganel, 2015, 674).

Both proponents and opponents of TVSH usually share the belief that the validity of this inference from systematic performance differences between grasping and judgement tasks to different representational formats involved in grasping and judging depends on a crucial premise. According to this premise, the two tasks (grasping and judging) have to be similar in all relevant respects. If the two tasks are too dissimilar, then too many confounding factors (such as working memory load, attentional load, general difficulties, and subjective strategies) might explain the systematic performance differences. One line of response to the original Ebbinghaus illusion studies has consequently been to insist that the grasping task and perceptual judgement task are too dissimilar to warrant the inference from performance differences to separate representational formats (Franz & Gegenfurtner, 2008; Franz, Hesse, & Kollath, 2009; Schenk, 2012). According to this line of reasoning, if you minimize the difference between the grasping and judgement tasks, then the systematic performance difference tends to disappear. In this way, it has recently been argued that if the tasks are sufficiently similar, both grasping and judging show similar illusion effects (Kopiske et al., 2016). Proponents of TVSH deny this last claim. They accept the assumption that the inference from performance difference to a difference in representational formats relies on task similarity, but they deny that the grasping tasks and perceptual judgement tasks are too dissimilar (see Goodale & Ganel, 2015, 675-676; Haffenden & Goodale, 1998; Whitwell & Goodale, 2016).

Summing up, experimental outcomes serve as evidence for a particular model to the extent that the data weigh against competing models. This kind of contrastive testing is only possible given a set of assumptions shared among the contrasted models. In the debates for and against TVSH, proponents and opponents broadly share assumptions concerning the validity and warrant of the inference from systematic performance differences in grasping and judgement tasks to different forms of visual encodings. These assumptions of validity and warrant are problematic, or so I will argue in the next section.

5. Underdetermination of models by behavioural data

In this section, I turn the spotlight onto another set of shared assumptions, but this time assumptions that are seldom openly recognized. A cognitive function executed by some computational mechanism is defined by, on the one hand, types of representational input and output, and, on the other hand, the rule-governed operation by which input representations are transformed into output representations. Consequently, when we determine the cognitive function performed by some computational mechanism, we determine two factors, namely, the representation and the operation. This has a number of interesting consequences for our ability to use experimental data as evidence for particular representational formats.

Formats are often thought to determine the character of the operations. Following this line of thought, some formats would seem to make certain types of operations more probable. For instance, propositional formats are thought to require certain types of digital, syntactic processing. If we already knew the representational format, we could make an informed, good guess about the type of operation (or if we knew the type of computational operation, we could point to a probable type of format). The challenge is that researchers know neither format nor operation in advance. A computational model is in this context a model that uses claims about representations and operations to predict behaviour in various kinds of task.[12] It is well known in the psychological literature that for such models, it is always possible for one type of model to mimic the behavioural output of another type of model (see Anderson, 1978, on how propositional models can mimic pictorial models, and Townsend, 1990, on how parallel models can mimic serial models).

If a cognitive computational mechanism is defined by the cognitive function it is computing, then it is defined by two factors, namely representation and rule-governed operation. We cannot infer what kind of representation is driving behaviour in a particular task independently of the operation (i.e., the rule-governed way in which input is mapped onto output). We are thus confronted with a problem with two unknown factors. Systematic performance difference in two different types of task (say, grasping vs. manual size estimation) could be explained either by a difference in the representational formats, a difference in the rule-governed operations, or both. The systematic performance difference will underdetermine which of these factors explain the behaviour. That is, systematic differences between performance in grasping tasks and in size estimation tasks can be taken as evidence in favour of an "egocentric vs. allocentric"-model over a "common representation"-model, only if the two competing models agree on common processing constraints. Otherwise, it will always be possible to explain the systematic performance difference by a systematic difference in operations transforming the visual representations into motor commands and cognitive responses. The behavioural data underdetermine the choice between a model according to which separate representational formats produce different behavioural tendencies and a model according to which a common representational format is processed in two different ways giving rise to different behavioural tendencies.[13]

One possible objection to this underdetermination argument is that the distinction between representation and operation is an idle theoretical distinction that finds no resonance in the experimental literature. That is, cognitive neuroscientists would count a difference in processing as a

[12] By "computational model" I mean a functional model of a supposed computational system. I do not use the term in the sense of a mathematical "representation of the natural system" (Palminteri, Wyart, & Koechlin, 2017). Only some computational models in the latter sense are models of supposed computational systems.

[13] A general version of this argument was proposed and discussed in detail by Anderson (1978).

Please quote the published version: https://doi.org/10.1007/s11229-018-01984-y

difference in representation. In particular, the objection would go, proponents and opponents of TVSH actually do not distinguish between representation and process.

This is a mistake. Proponents and opponents of TVSH distinguish between representation and process. Before taking a closer look at a recent discussion illustrating this fact, recall the distinction between *core computational processes* and *transformational processes*. The 2-systems model (TVSH) and the 1-system model are equally committed to the idea that complex transformational operations occur when the visual representations are transformed into motor commands or various forms of cognitive response. That is, both models accept that once the visual object-descriptions (say, representations of size, location, and shape) have been computed from the 'raw' visual information, distinct and separate processes take place transforming the visual representations into a manual motor response and a cognitive motor response. Independently of how the visual object-descriptions are computed, both models accept that the transformational process involved in grasping is separate from the transformational process leading to a cognitive response (like a manual size estimation or a verbal judgement).

These issues are nicely illustrated by a recent exchange between proponents and opponents of TVSH. In a series of visuomotor and psychophysical experiments, Ganel, Chajut, and Algom (2008) compared perceptual (i.e. cognitive) response with grasping. Ganel and colleagues presented objects of different sizes in two tasks. In the perceptual task, participants had to match the size of the object by the aperture of their index and thumb. In the visuomotor task, participants simply had to grasp the objects. In the latter task, the maximum grip aperture (MGA) was measured. The MGA is the maximum opening of the hand in-flight before closing in on the object, and it is standardly believed to reflect the representation of size computed from visual information before beginning the movement (Jeannerod, 1981). In both conditions, a measure of the finger aperture was the dependent variable. In both conditions, this measure was taken as a measure of the visual representation driving the responses.

By varying the sizes of the objects in the two tasks, Ganel et al. could compare the pattern of variable errors. In the perceptual task, it was found that the variable errors measured by standard deviation of the distributions of the finger aperture grew with object size. Perceptual responses thus follow Weber's law according to which the smallest quantity of change in stimulus intensity that causes a noticeable change in sensation (the so-called "just noticeable difference", JND) is a constant ratio of the stimulus intensity. This means that the JND increases for larger stimulus magnitudes. The larger the stimulus magnitude, the larger the change in stimulus intensity needed to produce a reliably noticeable change in sensations. As a consequence, the larger the stimulus magnitude, the larger the standard deviation of the stimulus discrimination will be. According to the authors, "this variance gauges the "area of uncertainty" for which the observer is unable to tell the difference between the size of the comparison and the target object" (Ganel, Chajut, & Algom, 2008, R600). By comparing the statistical mean measures of finger aperture in the two tasks, Ganel and colleagues found that the perceptual responses follow Weber's law (the larger the object, the larger standard deviation of the distribution of finger aperture), whereas grasping responses were reported to violate Weber's law. When grasping objects of varying sizes, the standard deviation of the MGA remained more or less constant. That is, in the visuomotor task, "JND was unaffected by the […] variation of size of the referent objects" (Ganel, Chajut, & Algom, 2008, R599).

According to Ganel et al., "these findings document a fundamental difference in the way that object size is computed for action and for perception and suggest that the visual coding for action is based on absolute metrics even at a very basic level of processing" (R599). They conclude that what they "have discovered is that one important feature of human action is coded based on absolute object size and hence is inconsistent with Weber's law" (R601). Ganel's results have frequently been cited (see, for instance, Goodale and Ganel, 2015) as important evidence for the claim that "the visual codings of spatial properties in the two functional modules obey different principles, resulting in context-sensitive, object-relative [i.e. allocentric] representations in perception but in context-blind [so-called "absolute metric"], effector-relative [i.e. egocentric] representations in action" (Bruno et al., 2016, 328). That is, the different patterns of variable errors in the two tasks are taken to be evidence for two different kinds of formats of visual encodings, on the basis of which it is inferred that separate and independent functional modules are responsible for computing them.

These distinct patterns of variable errors in grasping and size estimation tasks can be explained in other ways as well. Recently, it has been suggested that it might be the case that various kinds of sensory noise might mask the fact that visual representations driving grasping are governed by Weber's law (Utz et al., 2015; Schenk, Utz, & Hess, 2017). If that were the case, vision-for-action and vision-for-perception as measured in the grasping and size estimation tasks would not differ from each other. According to Utz and colleagues, grasping involves not only processing of visual information but also of various kinds of proprioceptive and tactile information that could have an effect on the distribution of errors involved in grasping. Grasping might violate Weber's law not because visual representations are encoded in an absolute metric that is not susceptible to Weber's law but because the system is using proprioceptive and tactile information to correct for the increased uncertainty involved with larger stimulus magnitudes.

The authors tested this hypothesis by manipulating these additional sources of information. Even though Utz et al. (2015) found no positive evidence supporting the hypothesis that also vision-for-grasping is governed by Weber's law, they found that data supported a kind of inverse Weber's law for grasping. The larger the object to be grasped, the smaller was the standard deviation of MGA (Utz et al., 2015). Their proposed explanation is that MGA is governed by biomechanical constraints. Simply put, when objects reach a certain size, then the fingers cannot open much wider. When an object has a width of 2 cm, then the correlated mean MGA might be approximately 7 cm (Castiello, 2005). The standard deviation of the finger aperture will be standardly distributed above and below the 7 cm. However, as the object sizes get larger, the aperture of the fingers reaches the point where a person cannot open her grip any more due to biomechanical constraints. This means that standard deviation shrinks. This simple fact explains the fact that for a certain set of object sizes, it looks as if vision-for-grasping is not governed by Weber's law. This leaves it open that for small objects, grasping would still show an increase in the variance of errors for increasing object sizes. This possibility has recently been supported by experiments by Bruno et al. (2016).

This discussion back and forth concerning whether visual representations driving grasping are governed by Weber's law clearly demonstrates the difficulties involved in inferring representational formats from systematic performance differences. The systematic performance differences (i.e. the distinct patterns of distributions of measured finger apertures) are taken as evidence for different kinds of representational encodings by Ganel and colleagues (Ganel, Chajut, & Algom, 2008; Goodale & Ganel, 2015) but as evidence for separate transformational processes by Utz et al. (2015) and Bruno et al. (2016). According to the latter authors, visual core computations could be producing a common visual representation used in grasping and size estimation. This fact is, however, masked by the transformational processes involved in grasping – transformational processes involving various kinds of sensory information and biomechanical constraints not present in manual size estimation. Difference and sameness of behavioural performance in perception and visuomotor tasks can thus be explained either by different representational formats or same representational formats but different transformational processes.

Final draft. Accepted 2018 in *Synthese*.
Please quote the published version: https://doi.org/10.1007/s11229-018-01984-y

6. Local yet systematic

We can sum up the discussion so far in the following way. The choice between the 2-systems model and the 1-system model of perception-action is *locally* yet *systematically* underdetermined by behavioural data. The choice is *locally* underdetermined because the underdetermination is generated by a set of local assumptions about the nature of computational mechanisms in the brain and experimental methodological principles. The problem of underdetermination follows from the individuation of computational mechanisms by formats of encoding and a set of methodological principles about how to infer encoded formats from behavioural differences in contrasted task situations. The particular underdetermination problem I have raised in this paper is therefore relative to areas of research where these assumptions are endorsed. We cannot be sure that the problem generalises to other discussions and explanatory models of cognitive neuroscience. There is no way to tell without analysing the actual explanatory and evidential practises of the scientific debates.

The choice is *systematically* underdetermined by data because the underdetermination is generated by core assumptions. By "systematic", I mean that the underdetermination problem is automatically produced by the combination of the view of individuation and certain methodological principles guiding the inference from behavioural data to representational formats. This problem will only go away if researchers are willing to make some hard, theoretical choices. In section 7, I suggest a way to avoid the underdetermination problem by rejecting the methodological assumptions.

The underdetermination of the choice between 2-systems and 1-system perception-action models is systematic in the sense that if you adopt the view of individuation and endorse the methodological assumptions about inferring representational formats, then you are confronted with a serious problem of underdetermination. For any given pair of contrasted psychological models, behavioural data will always underdetermine whether the behavioural results should be explained by representational formats or by transformational processes. If there are no objective reasons for ascribing a higher prior probability to either of the models (and there are no reasons for saying one of the models is simpler or more plausible than the other), then since the two models assign the same probability to the behavioural data (i.e. they have equal likelihoods), the behavioural data will be evidentially irrelevant to the choice between the models. This undermines the inference from behavioural data to representational formats. Consequently, given the view of individuation of computational mechanisms and standard methodological principles, behavioural data underdetermine the choice between a 2-systems model (TVSH) and a 1-system model.

Final draft. Accepted 2018 in *Synthese*.
Please quote the published version: https://doi.org/10.1007/s11229-018-01984-y

The preceding sections have established that our two rival models are data equivalent relative to data from behavioural experiments. This section addresses two related questions: (1) Are the two models comprehensively data equivalent? (2) If they are, what kind of local anti-realism with respect to computational mechanism would this justify? Recall that contrastive underdetermination can support anti-realism with respect to the entities posited by the rival theories only if the rival theories are comprehensively data equivalent and the rival theories could not both be approximately true.

The 2-systems and 1-system models are not empirically equivalent. The models differ with respect to some of their possible observational consequences. The two competing models are committed to different descriptions of the brain (the brain implements either two or one mechanisms for computing object-descriptions located either after or before the anatomical split), so given a sufficiently fine-grained functional imaging technique it might be possible to detect these differences.[14] Nevertheless, there are reasons for thinking that these observational consequences will not be cashed out in actual data in any near future. That is, there are reasons for thinking that the models are comprehensively data equivalent.

Due to space limitations, this part of the argument will have to be more conjectural than substantial. Three forms of data could constrain our choice of computational model: anatomical, physiological, and behavioural. I have already shown that, given the metaphysical and methodological assumptions, the model choice will always be underdetermined by behavioural data. What about anatomical data? Even if some researchers argue that the rejection of the anatomical model renders the functional model of TVSH implausible (see de Haan, Jackson, & Schenk, 2018), there are reasons for accepting that the anatomical model is in certain ways independent of the functional model. Perfect knowledge of the anatomy is in principle possible without knowing what the various parts of the brain are doing. Thus, even when we take into account all of the data in favour of the anatomical model of two separate neural streams, "it still remains to be established how the functional 'roles' of the two streams should best be characterized" (Milner & Goodale, 2006, 40).[15] On the assumption that anatomical structure by itself tells us very little about its psychological function and that both models accept the same anatomical model, we can set aside also anatomical data.

---

[14] This makes the underdetermination transient in Laudan and Leplin's (1991) sense.

[15] Evidence for the mutual independence of the anatomical and functional models also comes from the fact that there are some studies that seem to support a version of the functional model of TVSH while rejecting the anatomical model (e.g., Freud, Plaut, & Behrmann, 2016) and some studies reject the functional model while accepting the anatomical one (e.g. Christiansen et al., 2014).

Final draft. Accepted 2018 in *Synthese*.
Please quote the published version: https://doi.org/10.1007/s11229-018-01984-y

What about physiological data? Are the models equivalent relative to neurophysiological data from neuroimaging studies? Here I will focus exclusively on data from fMRI studies. Let us follow Kriegeskorte and Bandettini (2007) in distinguishing fMRI studies with respect to whether they employ a univariate *activation-based* or a multivariate *information-based* analysis of neuroimaging data. The activation-based analysis method works by averaging the activity across all voxels within some region of interest. The information-based analysis method determines degrees of statistical correlation between experimental conditions and patterns of activation in the individual voxels of the region of interest.

Traditionally, the activity-based analysis has been combined with some kind of subtraction paradigm. This allows researchers to determine the areas with most activation relative to particular experimental conditions (Amaro & Barker, 2006). This method might settle some model issues in cognitive neuroscience, if two psychological models make conflicting predictions about the coarse-grained localisation of activation patterns (see Colombo, 2014; Henson, 2005). This seems, however, not to be the case with the present model dispute about 2-systems and 1-system perception-action models. Both models are perfectly consistent with the there being differences in areas of activation between perception and action. After all, the 1-system model accepts that cognition and action embody different types of task. The claim is that a common type of visual representation is used by the brain to solve the different types of processing problem. This is consistent with there being many other differences between contrasted task conditions (for instance, differences in the involvement of attention, memory, modalities of sensory processing, and motor processing).

Maybe some of these limitations with the traditional use of activity-based analysis methods can be surpassed by using the information-based analysis methods. Kriegeskorte and Bandettini (2007, 658) claim that this analysis method can be used to determine the "representational content" of particular brain regions. The claim seems to be that since the voxel patterns correlated with particular experimental conditions carry information about the condition in question (from the voxel pattern alone one might be able to tell the experimental condition, i.e. what the participant saw or did), it is the brain's own representation of the stimulus or the behaviour that is conveyed by the fMRI data analysis. If that is the case, then maybe different voxel patterns in regions of interest might allow us to determine the brain's visual representation in cognitive and action tasks. Could this perhaps constrain the choice between our two rival perception-action models?

At present, it is hard to see that the information-based method of analysis would make a difference to our particular underdetermination problem. Take a recent study by Freud et al. (2018).

They contrasted four conditions: grasping a real 3D block vs "grasping" the 2D shape on a screen and reaching to a real 3D block vs reaching to a 2D shape. Using imaging data from the left anterior intraparietal sulcus in the dorsal stream, the authors were able to train a pattern classifier reliably (i.e. above-chance) to tell whether a participant was grasping or reaching and whether she was seeing a real block or a 2D image. From a region in the ventral stream, they were only able to train the pattern classifier reliably (above-chance) to tell the type of stimulus (real block vs 2D image).

Do these results tell us that different forms of visual computations are being implemented in the dorsal and the ventral streams? The answer is surely no. From the fact that the classifier can tell the difference between grasping real blocks and "grasping" 2D shapes, we cannot infer that different visual representations are used in these two conditions.[16] Several things differ across these conditions, and it might be impossible to rule out that the pattern classifier is sensitive to these differences (see Popov, Ostarek & Tenison, 2018; Ritchie, Kaplan & Klein, 2018). As in the case of data analysed with the traditional activation-based methods, evidence from information-based fMRI-studies seems to support only the more general claim that different task situations (reaching vs. grasping, moving towards real objects vs. 2D images) are associated with differences in localised activation patterns (Freud et al., 2018).[17] Such widespread differences in activation patterns are equally consistent with a 2-systems model and a 1-system model.[18]

Let me sum up these considerations. The 2-systems (TVSH) and the 1-system perception-action models are comprehensively data equivalent (i.e., they are equivalent relative to anatomical, physiological, and behavioural data). Furthermore, the two models cannot both be approximately true with respect to their core claims (the number of computational mechanisms). Consequently, given the metaphysical and methodological assumptions, there is no way in which we can determine whether there is one or two computational mechanisms for computing object descriptions from "raw" visual information. This is a local form of underdetermination relative to research domains where the assumptions are endorsed. Therefore, if the underdetermination supports a form of anti-realism, it

[16] Proponents of TVSH often assume that pantomime grasping and grasping 2D shapes are driven by the ventral system. See Freud & Ganel (2015).

[17] According to Popov, Ostarek and Tenison (2018), just like univariate activity-based methods, multivariate information-based methods can tell us something about *where* information is represented or processed but very little about *how* the information is represented or processed.

[18] It is important to notice that I have not presented a general argument against the use of neuroimaging data to constrain and qualify psychological processing models. In many cases, anatomical and physiological data should play a role in developing psychological models (Bechtel & Richardson, 2010; McGeer, 2007), and, given rival models, all kinds of predicted observational consequences are potentially relevant (Henson, 2005). The underdetermination problems I have presented are relative to the choice between perception-action models in the debates about TVSH.

could at most be a local form of anti-realism: anti-realism with respect to the computational mechanisms in question.

Even if the underdetermination problem confronting the choice between 2-systems and 1-system models of perception-action is systematic and involves comprehensive data equivalence, it provides us with no reasons for accepting *ontological anti-realism* with respect to computational mechanisms. Ontological anti-realism about computational mechanisms is the view that there are no computational mechanisms in the world independently of our descriptive practises (Putnam, 1988; Searle, 1990). The fact that anatomical, physiological, and behavioural data underdetermine the choice between models distinguished by their different numbers of computational mechanisms gives us no reason for denying that there are real computational mechanisms. Since the problem is derived from both a set of metaphysical assumptions (the individuation of computational mechanisms by formats of their input and/or output) and epistemic assumptions (experimental-methodological principles about how to infer formats from experimental data), it is an option to block the problem by rejecting only the experimental-methodological principles and leaving the view of individuation intact.[19]

One might think that a local, yet systematic and comprehensive form of underdetermination should motivate a local form of *epistemic anti-realism* with respect to computational models. If dominant and unavoidable background assumptions make it the case that that anatomical, physiological, and behavioural data underdetermine the choice between models distinguished by their different numbers of computational mechanisms, then we cannot determine the truth or falsity of any of these models. Given that the models are comprehensively data equivalent, we ought to be instrumentalists with respect to their central constructs, i.e. the computational mechanisms.[20] This would seem the right response only if these are really unavoidable background assumptions. Let us accept for the sake of argument that the view of individuation is unassailable. What about the experimental-methodological principles?

7. A new kind of experimental paradigm

Assuming that there are computational mechanisms in the brain, we can ask: What properties could these mechanisms have that would allow us to empirically determine their number? We have seen that representational format is not a helpful answer. What other properties might computational

[19] This might not alleviate ontological anti-realism motivated by the concerns of Hacking (1982). According to Hacking, only the possibility of manipulation and use of an entity gives us reason to postulate that it exists.
[20] This seems to be the position of Dennett (1987). It might also be the position of Shagrir (2010).

mechanisms in the brain have? Plausibly, they might have metabolic properties. Given that we do not have the techniques and methodologies at present to measure these spatiotemporally fine-grained differences, we are forced to look for properties that might manifest themselves behaviourally. One possible option is to focus on informational properties.

A computational mechanism has the job of transforming input information into output information. Without making any assumption about the involved representational formats, we might focus on the reliability of this information transfer. Take the following analogy. Imagine a version of the game of "Telephone". In this version, we are interested in comparing two different set-ups. In the first set-up, a Whisperer whispers a message ("We're going to advance") to a Writer, who writes down what she takes to be the message ("We're going to a dance") and passes it on to two separate and independent Readers, who each reads the written message out loud ("We're going to a dance"). Assume that the written message is always printed in a clear and legible manner. In the second set-up, a Whisperer whispers a message ("We're going to advance") to two separate and independent Writers, who each writes down what she takes to be the message. One writes "There's going to be an avalanche" and the other "We're going to a dance". Each of the Writers sends her message to a Reader. The two separate and independent Readers read out the messages. We can describe each of the functional roles (Whisperer, Writer, and Reader) as a computational mechanism. The Whisperer, Writer, and Reader each has the job of transforming input information in a rule-governed way into output information. The two set-ups are two different "architectures" that differ from each other only with respect to the number of Writer-mechanisms.

These two different "Telephone" architectures have different informational properties. A whisper is a message with a degraded informational quality in the sense that there is a high probability of misinterpretation. Let us assume that in our version of the game a Writer only rarely writes down the correct message. That is, assume that the Writer is mainly guessing. We now start to see interesting differences between the two architectures. In the 1-Writer architecture, on any given trial, the Writer writes down what she guesses to be the message ("We're going to a dance") and sends the same message to two Readers. The two Readers give voice to the same message. In the 2-Writer architecture, on a given trial, there is a high probability that the two writers will write down different messages ("There's going to be an avalanche" and "We're going to a dance") that are each sent to their respective Readers. The result is that two different messages will be read aloud. This means that if we do not know what kind of "Telephone" architecture we are dealing with, we can make our choice on the basis of the probability of sameness of response (on a trial-by-trail basis) by the Readers.

On a given trial, the probability of same response by Readers given a 1-Writer architecture is much higher than the probability given a 2-Writer architecture. That is, on a trial-by-trial basis, given a 1-Writer architecture, same response by the two readers should be above chance, whereas given a 2-Writer architecture it should be at chance level.

In a recent study (Christiansen et al., 2014), we implemented this logic in a new experimental paradigm to determine the choice between 2-systems and 1-system models of perception-action. Participants were instructed (in the same trial) to both grasp and judge (length and orientation) rectangles on a computer monitor. As in my version of "Telephone", the goal of the study was to determine whether object-descriptions used to drive the two responses (grasping and judgement) are computed by one common or two independent mechanisms. As in my story, when presented with a degraded stimulus, it should become possible to distinguish between 2-systems and 1-system models. Interestingly, the results demonstrated that, on a trial-by-trial basis, grasping and judging responses to length and orientation of a rectangle are in agreement significantly above chance level. On trials where participants successfully grasp the object, they also make correct judgements, whereas on trials where they do not succeed in grasping the object, their judgements are wrong. Furthermore, we showed that the accuracy of grasping and judging responses, on a trial-by-trial basis, was very similar for both length and orientation. These findings are best explained by assuming that grasping and judging responses are driven by feature information processed by a common mechanism.

If this logic and the results of the study are correct, then counting mechanisms for computing object-descriptions need not depend on the determination of representational formats. Counting the number of mechanisms for computing the representations of features can depend only on measuring the probability of sameness-of-response on trial-by-trial basis of the grasping and judging performance. Consequently, on its own, determining the number of computational mechanism in virtue of their informational properties would not help us determining the representational formats of the visual representations (egocentric or allocentric). This is consistent with a view according to which computational mechanisms are what they are in virtue of their representational formats. It just so happens that besides their representational formats, they also have other properties that would allow us to say something about their nature and number.

8. Concluding remarks

Let me end by summing up the argument. Progress in some areas of cognitive neuroscience is obstructed by local yet systematic problems of contrastive underdetermination: behavioural data underdetermine the choice between competing functional models. This situation is the case in debates between proponents and opponents of TVSH. The underdetermination problem is systematically generated by the joint assumption of a metaphysics of computational mechanisms and a set of methodological principles. The metaphysics of computation says that computational mechanisms are individuated by the formats of visual representations. This motivates a set of methodological principles that would allow cognitive scientists to infer representational formats from systematic performance profiles in behavioural experiments. The problem is that behavioural data underdetermine the choice between competing representational models when these are determined only by their representational formats. At present, anatomical and physiological data cannot help with this underdetermination problem.

The underdetermination problem gives us no reasons to be ontological or epistemic anti-realist about computational mechanisms in the brain. Ontological anti-realism is avoided because the problem is partly an epistemic problem. Epistemic anti-realism about computational mechanisms is avoided because the problem is in part generated by methodological principles that can be sidestepped. These principles dictate that we infer facts about representational formats from behavioural data. Even if we accept the received view of computational mechanisms in the brain, representational formats might not be the only kinds of property that would allow us to empirically distinguish and count computational mechanisms.

If the arguments of this paper are on the right track, we have to accept that, given common assumptions in the field, the rival functional perception-action models are comprehensively data equivalent. Consequently, data from the type of paradigm exemplified by the studies by Aglioti et al. (1995) and Ganel et al. (2008) turn out to be evidentially irrelevant to the choice for or against TVSH. Fortunately, this is not the only possible type of experimental paradigm. We should move away from a focus on representational formats to a focus on the informational properties of computational mechanisms. Data from experimental paradigms that enable us to analyse the probability of making same response on a trial-by-trial basis would seem to be evidentially relevant to the model choice in question. The data so far weigh in favour of the 1-system model and against TVSH.

Acknowledgements: The author is indebted to Adrian Alsmith, Somogy Varga, and a number of reviewers for the journal. I am in particular grateful to reviewer 1 for insisting on larger revisions and providing me with advice and literature. The ideas have been presented at workshops and seminars

in Copenhagen, St. Andrews, Oxford, Aarhus, and I am grateful for the feedback I received on those occasions.

References

Aglioti, S., DeSouza, J. F., & Goodale, M. A. (1995). Size-contrast illusions deceive the eye but not the hand. *Current biology*, *5*(6), 679-685.

Anderson, J. R. (1978). Arguments concerning representations for mental imagery. *Psychological Review*, *85*(4), 249.

Apperly, I. A. (2008). Beyond simulation–theory and theory–theory: why social cognitive neuroscience should use its own concepts to study “Theory of Mind”. *Cognition*, *107*(1), 266-283.

Bechtel, W. & Richarson, R.C. (2010). Neuroimaging as a Tool for Functionally Decomposing Cognitive Processes. In S. J. Hanson & M. Bunzl (eds.), *Foundational Issues in Human Brain Mapping*, pp. 241-262. Cambridge, MA: The MIT Press.

Bermúdez, J. L. (2006). Arguing for eliminativism. In B. Keeley (Ed.), *Paul Churchland: Contemporary Philosophy in Focus* (pp. 32-66). Cambridge: Cambridge University Press.

Belousek, D. W. (2005). Underdetermination, realism, and theory appraisal: An epistemological reflection on quantum mechanics. *Foundations of Physics*, *35*(4), 669-695.

Biddle, J. (2013). State of the field: Transient underdetermination and values in science. *Studies in History and Philosophy of Science Part A*, *44*(1), 124-133.

Briscoe, R. (2008). Another look at the two visual systems hypothesis: The argument from illusion studies. *Journal of Consciousness Studies*, *15*(8), 35-62.

Brogaard, B. (2011). Conscious vision for action versus unconscious vision for action? *Cognitive Science*, 35, 1076–1104.

Bruno, N., Uccelli, S., Viviani, E., & de'Sperati, C. (2016). Both vision-for-perception and vision-for-action follow Weber's law at small object sizes, but violate it at larger sizes. *Neuropsychologia*, *91*, 327-334.

Butterfield, J. (2012). Underdetermination in Cosmology: An Invitation. *Proceedings of the Aristotelian Society*, *Supplementary Volumes*, 1-18.

Butterfill, S. A. and Sinigaglia, C. (2014). Intention and Motor Representation in Purposive Action: Intention and Motor Representation in Purposive Action. *Philosophy and Phenomenological Research*, 88(1), 119–145.

Castiello, U. (2005). The neuroscience of grasping. *Nature Reviews Neuroscience*, *6*(9), 726-736.

Christiansen, J. H., Christensen, J., Grünbaum, T., & Kyllingsbæk, S. (2014). A common representation of spatial features drives action and perception: grasping and judging object features within trials. *PloS one*, *9*(5).

Clark, A. (2001). Visual experience and motor action: Are the bonds too tight? *Philosophical Review*, *110*(4), 495-519.

Clark, A. (2007). What reaching teaches: Consciousness, control, and the inner zombie. *The British Journal for the Philosophy of Science*, *58*(3), 563-594.

Colombo, M. (2014). For a Few Neurons More: Tractability and Neurally Informed Economic Modelling. *The British Journal for the Philosophy of Science*, *66*(4), 713-736.

Culham, J. C., & Valyear, K. F. (2006). Human parietal cortex in action. *Current opinion in neurobiology*, *16*(2), 205-212.

de Haan, E. H. F., Jackson, S. R., & Schenk, T. (2018). Where are we now with 'What' and 'How'?. *Cortex*, *98*, 1-7.

Dennett, D. (1987). Instrumentalism reconsidered. In his *The intentional stance*. Cambridge, MA: MIT Press.

de la Malla, C., Smeets, J. B., & Brenner, E. (2018). Errors in interception can be predicted from errors in perception. *Cortex*, *98*, 49-59.

Dellsén, F. (2017). Realism and the absence of rivals. *Synthese*, *194*(7), 2427-2446.

Duhem, P. (1962). *The aim and structure of physical theory*. New York: Atheneum.

Earman, J. (1993). Underdetermination, realism, and reason. *Midwest studies in philosophy*, *18*(1), 19-38.

Ferretti, G. (2018). Visual feeling of presence. *Pacific Philosophical Quarterly*, *99*, 112-136.

Foley, R. T., Whitwell, R. L., & Goodale, M. A. (2015). The two-visual-systems hypothesis and the perspectival features of visual experience. *Consciousness and Cognition*, *35*, 225-233.

Franz, V., & Gegenfurtner, K. (2008). Grasping visual illusions: Consistent data and no dissociation. *Cognitive Neuropsychology*, 25, 920–950.

Franz, V. H., Gegenfurtner, K. R., Bülthoff, H. H., & Fahle, M. (2000). Grasping visual illusions: No evidence for a dissociation between perception and action. *Psychological Science*, 11, 20–25.

Franz, V. H., Hesse, C., & Kollath, S. (2009). Visual illusions, delayed grasping, and memory: no shift from dorsal to ventral control. *Neuropsychologia*, *47*(6), 1518-1531.

Freud, E., & Ganel, T. (2015). Visual control of action directed toward two-dimensional objects relies on holistic processing of object shape. *Psychonomic Bulletin & Review*, *22*(5), 1377-1382.

Freud, E., Macdonald, S. N., Chen, J., Quinlan, D. J., Goodale, M. A., & Culham, J. C. (2018). Getting a grip on reality: Grasping movements directed to real objects and images rely on dissociable neural representations. *Cortex*, *98*, 34-48.

Galletti, C., & Fattori, P. (2018). The dorsal visual stream revisited: Stable circuits or dynamic pathways?. *Cortex*, *98*, 203-217.

Ganel, T., Chajut, E., & Algom, D. (2008). Visual coding for action violates fundamental psychophysical principles. *Current Biology*, *18*(14), R599-R601.

Ganel, T., & Goodale, M. A. (2017). Still holding after all these years: An action-perception dissociation in patient DF. *Neuropsychologia*.

Ganis, G., & Schendan, H. E. (2011). Visual imagery. *Wiley Interdisciplinary Reviews: Cognitive Science*, *2*(3), 239-252.

Godfrey-Smith, P. (2008). Recurrent transient underdetermination and the glass half full. *Philosophical Studies*, *137*(1), 141-148.

Goodale, M. A. (2014). Visual Control of Action. In K. N. Ochsner and S. M. Kosslyn (eds.), *The Oxford Handbook of Cognitive Neuroscience, Vol. 1*, pp. 273-295. Oxford: Oxford University Press.

Goodale, M. A., & Ganel, T. (2015). Different modes of visual organization for perception and action. In J. Wagemans (Ed.), *Oxford handbook of perceptual organization* (pp. 672-690). Oxford: Oxford University Press.

Goodale, M. A., & Haffenden, A. (1998). Frames of reference for perception and action in the human visual system. *Neuroscience & Biobehavioral Reviews*, *22*(2), 161-172.

Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends in neurosciences*, *15*(1), 20-25.

Grünbaum, T. (2012). Commonsense psychology, dual visual streams, and the individuation of action. *Philosophical Psychology*, *25*(1), 25-47.

Grünbaum, T. (2017). The Perception-Action Model: Counting Computational Mechanisms. *Mind & Language*, 32(4), 416-445.

Hacking, I. (1982). Experimentation and Scientific Realism. *Philosophical Topics*, *13*(1), 71-87.

Haffenden, A. M., & Goodale, M. A. (1998). The effect of pictorial illusion on prehension and perception. *Journal of Cognitive Neuroscience*, *10*(1), 122-136.

Hoefer, C., & Rosenberg, A. (1994). Empirical equivalence, underdetermination, and systems of the world. *Philosophy of science*, *61*(4), 592-607.

Jeannerod, M. (1981). Intersegmental coordination during reaching at natural visual objects. In Long, J., Baddeley, A. (Eds.), *Attention and Performance IX* (pp.153–168). Hillsdale, NJ: Erlbaum.

Kitcher, P. (2001). *Science, truth, and democracy*. Oxford: Oxford University Press.

Kopiske, K. K., Bruno, N., Hesse, C., Schenk, T., & Franz, V. H. (2016). The functional subdivision of the visual brain: Is there a real illusion effect on action? A multi-lab replication study. *Cortex*, *79*, 130-152.

Kosslyn, S. M. (1994). *Image and brain*. Cambridge, MA: The MIT Press.

Kosslyn, S. M., Ball, T. M., & Reiser, B. J. (1978). Visual images preserve metric spatial information: evidence from studies of image scanning. *Journal of experimental psychology: Human perception and performance*, *4*(1), 47.

Kozuch, B. (2015). Dislocation, not dissociation: The neuroanatomical argument against visual experience driving motor action. *Mind & Language*, 30(5), 572–602.

Kravitz, D. J., Saleem, K. S., Baker, C. I., & Mishkin, M. (2011). A new neural framework for visuospatial processing. *Nature Reviews Neuroscience*, *12*(4), 217-230.

Kriegeskorte, N., & Bandettini, P. (2007). Analyzing for information, not activation, to exploit high-resolution fMRI. *Neuroimage*, *38*(4), 649-662.

Kukla, A. (1996). Does every theory have empirically equivalent rivals?. *Erkenntnis*, *44*(2), 137-166.

Laudan, L. (1990). Demystifying underdetermination. In C. W. Savage (Ed.), *Scientific theories, Minnesota studies in the philosophy of science* (Vol. 14) (pp. 267–297). Minneapolis: University of Minnesota Press.

Laudan, L., & Leplin, J. (1991). Empirical equivalence and underdetermination. *The Journal of Philosophy*, *88*(9), 449-472.

Loosemore, R. & Harley, T. (2010). Brains and Minds: On the Usefulness of Localization Data to Cognitive Psychology. In S. J. Hanson & M. Bunzl (eds.), *Foundational Issues in Human Brain Mapping*, pp. 217-240. Cambridge, MA: The MIT Press.

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

McGeer, V. (2007). Why neuroscience matters to cognitive neuropsychology. *Synthese*, *159*(3), 347.

Medendorp, W. P., de Brouwer, A. J., & Smeets, J. B. J. (2018). Dynamic representations of visual space for perception and action. *Cortex*, *98*, 194-202.

Milner, A. D., & Goodale, M. A. (2006). The visual brain in action (2nd ed.). Oxford: Oxford University Press.

Mole, C. (2009). Illusions, Demonstratives, and the Zombie Action Hypothesis. *Mind*, *118*(472), 995-1011.

Mylopoulos, M., & Pacherie, E. (2017). Intentions and motor representations: The interface challenge. *Review of Philosophy and Psychology*, *8*(2), 317-336.

Norton, J. (2008). Must evidence underdetermine theory? In M. Carrier, D. Howard, & J. Kourany (Eds.), *The challenge of the social and the pressure of practice: Science and values revisited* (pp. 17–44). Pittsburgh, PA: University of Pittsburgh Press.

Quine, W. V. O. (1951). Two dogmas of empiricism. *The Philosophical Review*, 60, 20–43.

Palminteri, S., Wyart, V., & Koechlin, E. (2017). The Importance of Falsification in Computational Cognitive Modeling. *Trends in Cognitive Sciences*, 21(6), 425-433.

Piccinini, G. (2015). *Physical Computation: A Mechanistic Account*. Oxford: Oxford University Press.

Popov, V., Ostarek, M., & Tenison, C. (2018). Practices and Pitfalls in Inferring Neural Representations. *NeuroImage*, *174*, 340-351.

Putnam, H. (1988). *Representation and Reality*. Cambridge, MA: MIT Press.

Pylyshyn, Z. W. (1981). The imagery debate: Analogue media versus tacit knowledge. *Psychological Review*, *88*(1), 16.

Ritchie, J.B., Kaplan, D.M. & Klein, C (2018). Decoding the Brain: Neural Representation and the Limits of Multivariate Pattern Analysis in Cognitive Neuroscience, *The British Journal for the Philosophy of Science*, axx023, https://doi.org/10.1093/bjps/axx023

Rossetti, Y., Pisella, L., & McIntosh, R. D. (2017). Rise and fall of the two visual systems theory. *Annals of physical and rehabilitation medicine*, 60, 130-140.

Salmon, W. (1990). Rationality and objectivity in science or Tom Kuhn meets Tom Bayes. *Minnesota Studies in the Philosophy of Science*, 14, 175-204.

Schenk, T. (2012). No dissociation between perception and action in patient DF when haptic feedback is withdrawn. *The Journal of neuroscience*, *32*(6), 2013-2017.

Schenk, T., Utz, K. S., & Hesse, C. (2017). Violations of Weber's law tell us more about methodological challenges in sensorimotor research than about the neural correlates of visual behaviour. *Vision research*, *140*, 140-143.

Searle, J. R. (1990). Is the brain a digital computer?. In *Proceedings and addresses of the American Philosophical Association* (pp. 21-37). American Philosophical Association.

Shagrir, O. (2010). Brains as analog-model computers. *Studies In History and Philosophy of Science Part A*, *41*(3), 271-279.

Shepherd, J. (2017). Skilled action and the double life of intention. *Philosophy and Phenomenological Research*. https://doi.org/10.1111/phpr.12433.

Sober, E. (1999). Testability. In *Proceedings and Addresses of the American Philosophical Association* (Vol. 73, No. 2, pp. 47-76). American Philosophical Association.

Sober, E. (2008). *Evidence and evolution: The logic behind the science*. Cambridge: Cambridge University Press.

Stanford, P. K. (2001). Refusing the devil's bargain: What kind of underdetermination should we take seriously?. *Philosophy of Science*, *68*(S3), S1-S12.

Stanford, P. K. (2006). *Exceeding our grasp*. Oxford: Oxford University Press.

Stanford, K. (2017). Underdetermination of Scientific Theory. *The Stanford Encyclopedia of Philosophy*, E. N. Zalta (ed.), URL=<https://plato.stanford.edu/archives/win2017/entries/scientific-underdetermination/>.

Townsend, J. T. (1990). Serial vs. parallel processing: Sometimes they look like Tweedledum and Tweedledee but they can (and should) be distinguished. *Psychological Science*, *1*(1), 46-54.

Tulodziecki, D. (2013). Underdetermination, methodological practices, and realism. *Synthese*, *190*(17), 3731-3750.

Utz, K. S., Hesse, C., Aschenneller, N., & Schenk, T. (2015). Biomechanical factors may explain why grasping violates Weber's law. *Vision research*, *111*, 22-30.

Van Frassen, B. (1980). *The scientific image*. Oxford: Clarendon Press.

Wallhagen, M. (2007). Consciousness and action: Does cognitive science support (mild) epiphenomenalism? *The British Journal for the Philosophy of Science*, *58*(3), 539-561.

Westwood, D. A. (2009). The visual control of object manipulation. In E. Morsella, J. A. Bargh, & P. M. Gollwitzer (Eds.), *The Oxford handbook of human action* (pp. 88-103). New York: Oxford University Press.

Whitwell, R. L., & Goodale, M. A. (2016). Real and illusory issues in the illusion debate (Why two things are sometimes better than one): Commentary on Kopiske et al.(2016). *Cortex*, *88*, 205-209.

Worrall, J. (2011). Underdetermination, realism and empirical equivalence. *Synthese*, *180*(2), 157-172.

Wu, W. (2013). The Case for Zombie Agency. *Mind*, *122*(485), 217-230.

Wu, W. (2014). Against division: Consciousness, information and the visual streams. *Mind & Language*, *29*(4), 383-406.