




ARTICLE

Teleofunction in the Service of Computational Individuation

Nir Fresco¹, Marc Artiga² and Marty J. Wolf³

¹Departments of Cognitive & Brain Sciences and Philosophy, Ben-Gurion University of the Negev, Beer Sheva, Israel, ²Department of Philosophy, University of Valencia, Valencia, Spain and ³Department of Mathematics and Computer Science, Bemidji State University, Bemidji, MN, USA

Corresponding author: Nir Fresco; Email: nfresco@bgu.ac.il

(Received 12 September 2023; revised 12 April 2024; accepted 27 July 2024; first published online 18 October 2024)

Abstract

One type of computational indeterminacy arises from partitioning a system's physical state space into state types that correspond to the abstract state types underlying the computation concerned. The mechanistic individuating strategy posits that computation can be uniquely identified through either narrow physical properties exclusively or wide, proximal properties. The semantic strategy posits that computation should be uniquely identified through semantic properties. We develop, and defend, an alternative functional individuating strategy that appeals—when needed—to wide, distal functions. We claim that there is no actual computation outside of a functional context. Desiderata for the underlying notion of teleofunction are discussed.

1. Introduction

Computational indeterminacy has been discussed by various authors, even if by a different name. They include, in chronological order, Block (1990), Shagrir (2001, 2020), Bishop (2009), Sprevak (2010), Piccinini (2015), Coelho Mollo (2018), Dewhurst (2018), Fresco, Copeland, and Wolf (2021), and Curtis-Trudel (2022). The underlying idea is that a single physical computational system may legitimately be described as computing several distinct mathematical functions. There is no apparent fact of the matter that determines the computational identity of the system in the face of such simultaneous computation. Thus an important question arises: what determines the computational identity of a physical system that simultaneously computes multiple mathematical functions?

Computational indeterminacy is commonly demonstrated by appealing to logic gates. Consider an electronic gate that operates systematically on voltage ranges, such as 1–4 V and 5–9 V, and produces the higher voltage as output only when both inputs are the higher voltage. Such a gate can be used to compute conjunction when the lower range is mapped onto logical 0 and the higher onto logical 1 (see table 1).

Table 1. Truth table of an AND-gate

First input	Second input	Output
0	0	0
0	1	0
1	0	0
1	1	1

Table 2. Truth table of an OR-gate

First input	Second input	Output
0	0	0
0	1	1
1	0	1
1	1	1

It can likewise be used to compute disjunction when the mapping is reversed (see table 2). Broadly speaking, two different strategies have emerged to address such indeterminacy: the *mechanistic approach* (e.g., Coelho Mollo 2018; Dewhurst 2018; Piccinini 2015; Tucker 2018) and the *semantic approach* (e.g., Bishop 2009; Harbecke and Shagrir 2019; Shagrir 2022; Sprevak 2010).

We introduce a new individuating strategy that addresses shortcomings of the mechanistic approach but does not require a commitment to representational properties, as does the semantic approach. We draw distinctions between narrow and wide individuation and couple those with the relative distality of candidate individuating functions. When a system is individuated without consideration of features external to the system, the individuation is narrow; otherwise, the individuation is wide. In either case, by our account, the individuating function of the system is the most proximal one.

The article proceeds as follows. In section 2, we spell out variations of computational indeterminacy and some details of existing individuating strategies. We argue that Piccinini’s mechanistic strategy cannot adequately address the indeterminacy found in Shagrir’s tri-stable system. To develop this argument, we introduce two conceptual dimensions along which computation may be individuated. In section 3, we introduce our wide teleofunctional strategy of computational individuation, which appeals to wide, distal functions in the face of computational indeterminacy. This appeal need not collapse to semantic content. To defend these claims and exemplify how indeterminacy of biological computation can be addressed, we introduce a hypothetical plankton case. In section 4, we address four objections to the teleofunctional individuating strategy and, in doing so, illuminate its nuances. In section 5, we specify some constraints on the notion of teleofunction that can be used to individuate computation using our approach. Section 6 concludes the article.

2. The indeterminacy of computation and how to address it

This section provides the backdrop for extending current teleomechanistic approaches to computational individuation. In section 2.1, we briefly discuss two types of computational indeterminacy and identify the type that is the focus of our analysis. In section 2.2, we introduce two conceptual dimensions along which our proposed strategy individuates computation. Finally, in section 2.3, we evaluate two versions of current individuation strategies: the *mechanistic* and the *semantic*.

2.1. Two types of computational indeterminacy

We distinguish two types of computational indeterminacy as identified in Papayannopoulos, Fresco, and Shagrir (2022). The first type, *abstract* indeterminacy, may arise when specifying the abstract (functional) organization of the system according to its systematic behavior. It concerns how the system's physical state space is partitioned into physical state types that correspond to the abstract state types underlying the computation concerned. The indeterminacy question here is which physical states should be grouped together as a type to correspond with a particular abstract state type. The second type of indeterminacy, *interpretative* indeterminacy, may arise when interpreting the values of the computational states. It is typically exemplified by the two logic gates introduced previously (see tables 1 and 2), though it extends beyond Boolean logic. These two types of indeterminacy are related, but distinct.

Piccinini (2020) has recently claimed that interpretative indeterminacy poses no challenge to the mechanistic individuation of computation. "Relabeling state types is a purely observer-dependent maneuver. It does not alter how the system works or what it's capable of" (152). The real problem of computational indeterminacy is *supposedly* finding the correct grouping of physical microstates (Piccinini 2020, 153; Tucker 2018, 14–15). There is disagreement about this claim in the literature (cf. Fresco, Copeland, and Wolf 2021). Nevertheless, Piccinini identifies abstract indeterminacy as the one that challenges individuation strategies, and Shagrir's tri-stable system (which is an instance of abstract indeterminacy) is developed to provide a compelling reason for preferring semantic individuation strategies over mechanistic ones. Therefore, for the purpose of developing a nonrepresentational functional individuation strategy, we focus simply on abstract indeterminacy.

Let us explain abstract indeterminacy by means of an example originally developed by Shagrir (2001) and later elaborated by Harbecke and Shagrir (2019). Consider an electronic circuit that operates on voltage levels (just like standard logic gates) in the range between 1 and 6 V. The circuit is connected to a robotic arm, and based on the output produced by the circuit, the arm moves. This circuit can be described as computing (at least) two mathematical functions based on how the input voltages (i.e., physical states) are grouped. The tabular specification (see table 3) describes the circuit as operating on *three* voltage ranges (a tri-stable system): 1–2 V (low voltage), 3–4 V (medium voltage), and 5–6 V (high voltage).

Abstract indeterminacy manifests in this imagined circuit when the voltage range is bipartitioned in different ways. For example, one possible grouping is 1–4 V (low voltage) and 5–6 V (high voltage). Another possible grouping is 1–2 V (low voltage) and 3–6 V (high voltage). The system is thus amenable to different ways of grouping

Table 3. An electronic circuit—connected to a robotic arm—that yields either two (movement or no movement) or three types of possible arm movement: no movement (between 0 and 45 degrees), medium movement (between 45 and 90 degrees), and high movement (more than 90 degrees)

Input channel 1 (V)	Input channel 2 (V)	Output channel (V)	Arm movement (ternary) (deg)
1–2	1–2	1–2	None (e.g., 0–45)
1–2	3–4	3–4	Medium (e.g., 45–90)
1–2	5–6	3–4	Medium (e.g., 45–90)
3–4	1–2	3–4	Medium (e.g., 45–90)
3–4	3–4	3–4	Medium (e.g., 45–90)
3–4	5–6	3–4	Medium (e.g., 45–90)
5–6	1–2	3–4	Medium (e.g., 45–90)
5–6	3–4	3–4	Medium (e.g., 45–90)
5–6	5–6	5–6	High (e.g., >90)

Note. The possible “binary movement” outputs are described in tables 4 and 5.

Table 4. A compact description of the circuit described in table 3 in which the medium voltage range (3–4 V) and the high voltage range (5–6 V) are grouped together as the high range (3–6 V)

Input channel 1 (V)	Input channel 2 (V)	Output channel	Arm movement
1–2	1–2	1–2	None
3–6	1–2	3–6	Movement
1–2	3–6	3–6	Movement
3–6	3–6	3–6	Movement

Note. Redundant rows are omitted. The threshold is thus set at 3 V. If the low voltage range is treated as logical 0 and the newly partitioned high voltage range is treated as logical 1, then the corresponding function is disjunction.

physical states together without there being any principled reason for preferring one way over the other. Thus it can be simultaneously used, without changing its underlying physical structure or operation, as two different binary circuits—one computing disjunction (see table 4) and another computing conjunction (see table 5). But this is *not* an instance of interpretative indeterminacy, because in *both* cases, the low voltage range is treated as logical 0 and the high voltage range is treated as logical 1.¹ It is the different groupings of the physical states that result in the system computing conjunction or disjunction. Therefore the same underlying circuit

¹ Nevertheless, we note that once the specific grouping is *fixed* (e.g., the tri-stable circuit being used to compute conjunction), interpretive indeterminacy looms (e.g., while, on one interpretation of the physical values processed by the circuit, it performs conjunction, on another, the very same circuit performs disjunction).

Table 5. A compact description of the circuit described in table 3 in which the low voltage range (1–2 V) and the medium voltage range (3–4 V) are grouped together as the low range (1–4 V)

Input channel 1 (V)	Input channel 2 (V)	Output channel	Arm movement
1–4	1–4	1–4	None
5–6	1–4	1–4	None
1–4	5–6	1–4	None
5–6	5–6	5–6	Movement

Note. The threshold is thus set at 4 V. If the newly partitioned low voltage range is treated as logical 0 and the high voltage range is treated as logical 1, then the corresponding function is conjunction.

computes the binary functions (disjunction and conjunction) as well as the ternary function listed in table 3.

Proponents of the mechanistic and semantic views of computation individuation offer different answers to the question of what the computational identity of this circuit is. Before evaluating their different answers and the implications for computational individuation, we turn to distinguishing between two important conceptual dimensions of individuation: (1) narrow and wide features and (2) proximal and distal features.

2.2. Conceptual machinery: Narrow versus wide individuation and proximal versus distal functions

Different individuation strategies fix a system's computational identity along two conceptual dimensions: (1) the boundaries of the physical system concerned (call it *S*) and (2) the relative distality of inputs and outputs. Our individuation strategy relies on considering both with respect to *S*'s inputs and outputs. The first distinction concerns the question of whether computational individuation supervenes only on properties, states, processes, and structure that are internal to *S*. If so, then the individuation is narrow (Fresco 2021, 14004). Narrow individuation does not refer, either implicitly or explicitly, to the states or structure of the environment.² If *S*'s environment, including any larger system encompassing *S* as a subcomponent, can affect the computational identity of *S*, then the individuation is wide. Simply, narrow and wide individuation concern the system's boundaries.

The second distinction concerns the relative distality of inputs and/or outputs that play a role in the computational process. This dimension is a relative matter and applies whether the individuation is narrow or wide. In teleomechanistic views of computation (e.g., Coelho Mollo 2018; Fresco 2021; Piccinini 2015), it concerns the distality of teleofunctions. A function is only more or less distal relative to another function and may not even exceed the boundaries of the system concerned. The heart, for example, has many functions. Most proximally, it rhythmically expands and contracts; more distally, it pumps blood; yet more distally, it distributes oxygen and nutrients to cells in the body, and so on. Therefore one can draw the proximal/distal

² Coelho Mollo (2018), Dewhurst (2018), and Tucker (2018) all similarly argue that the environment cannot (or does not) affect how the system concerned is computationally individuated.




Figure 1. (right) The amygdala can be described as receiving inputs at increasing levels of distality, both narrowly, from the second visual cortex (V2), or retinal photoreceptors, and widely, from a large visual stimulus in the environment, or some bear at a short distance. (left) Likewise, it can be described as contributing to outputs produced at increasing levels of distality, both narrowly, such as the emission of serotonin, or leg muscle contraction velocity, and widely, such as running, or escaping from a nearby bear.

distinction between narrowly individuated functions, between widely individuated functions, or between a combination of the two.

The proximal/distal distinction can be drawn between *effects* (some of which might qualify as functions) and between *causes*. To make these conceptual dimensions more palpable—also with respect to the relevant inputs and outputs—let us consider a rather straightforward example: the amygdala, the primary neural processing center for emotions (see figure 1). On the input side, activation in the secondary visual cortex is a proximal cause of activity in the amygdala, whereas retinal photoreceptor activation is a more distal trigger. Both causes are within the organism’s boundaries and thus can be part of a narrow individuating strategy. But the inputs can also be widely individuated. The amygdala may be triggered by an undetermined large visual stimulus in the environment or by a nearby bear.³ The former is more proximal than the latter, yet both are outside the organism’s boundaries. Therefore what triggers the amygdala may be individuated either narrowly or widely, and more proximally or less so.

Let us consider next the individuation of outputs. The amygdala may contribute to narrowly individuated activities, such as the emission of serotonin (also associated with the regulation of energy expenditure) or leg muscle contraction velocity—the latter being more distal than the former. However, it may also contribute to widely individuated activities, such as running on a treadmill in the gym or escaping from a nearby bear. Escape behavior will likely involve running, but not every run is an escape. On the output side, too, then, the (selected) effects of the amygdala may be individuated either narrowly or widely and either more distally or less so.

To reiterate, computational individuation can be narrow or wide and, simultaneously, more distal or less so. It is narrow if the computational inputs and outputs do not go beyond the computational system itself, that is, beyond the system’s sensors (e.g., photons hitting photoreceptors in the retina) and effectors

³ The semanticist may complain that the sight of a bear invokes semantic content. That is fine for present purposes, but, importantly, the plankton case study in section 3.2 does not.

(e.g., muscle fiber movements). Otherwise, it is wide (i.e., the inputs or outputs need not be restricted to the system's boundaries—as in the case of V2 neuronal activity or retinal receptor activity triggering the amygdala). Second, within a narrow or wide distinction, there can be a supplementary appeal to different proximal and distal functions.

With this conceptual machinery in hand, let us now return to evaluating two predominant individuating strategies that have been suggested in the literature: the *mechanistic* and the *semantic*.

2.3. Does wide, proximal individuation of computation suffice?

Both the mechanistic and the semantic views aim to clearly distinguish computational systems from noncomputational systems and to provide a means to develop a taxonomy of computational—and, ultimately, cognitive—systems based on their (mechanistic or semantic) properties. Within the mechanistic view, we set aside individualist mechanistic views of computational individuation (e.g., Coelho Mollo 2018; Dewhurst 2018; Tucker 2018), because they are, arguably, either *too narrow*⁴ or *vulnerable to other problems* (Fresco 2021) in rejecting the claim that computational indeterminacy is a matter of logical function. We thus focus here on Piccinini's wide, proximal individuating strategy. How does the wide, proximal strategy individuate computational states and systems?

On Piccinini's view, while the identity of computation that the system performs may be affected by the environment, it is only by the *immediate* environment that it may be affected. According to Piccinini (2015, 139–40), to understand the nature of wide individuation, one should draw an epistemic distinction between functionally irrelevant and relevant properties of the computational system. Moreover, this epistemic distinction requires knowing (1) *which* of the properties are relevant to the system's computational inputs and outputs and (2) *how* they pertain to the computational explanandum. The relevant functional properties, Piccinini argues, are not very wide: these properties concern the system's interaction with its immediate context through the system's input/output transducers. "Identifying broad [= wide] functions requires looking at the relation between a mechanism and its context, [not] solely at the intrinsic properties of a system" (155). A wide computational individuation of the relevant properties of nervous systems, however, "does not even reach into the organisms' environment; it only reaches sensory receptors and muscle fibers" (140).

Nevertheless, such individuation is not always sufficient for determining the computational identity of a system. The tri-stable circuit (see table 3) demonstrates one example. Consistent with Piccinini's suggestion that wide computational individuation need not extend beyond the system's sensors and effectors, the circuit's outputs are functionally connected to a robotic arm moving up and down (Harbecke and Shagrir 2019). Arm movement may be individuated in two ways, depending on the system's

⁴ Tucker (2018), like Dewhurst (2018), rejects any appeal to teleological functions in individuating computation, and narrowly individuates the *actual* behavior of computational systems. However, he argues that widely individuated norms, which are *external* to the definition of computation, may fix the function of computational systems. One problem with this view is that it admits of some form of pancomputationalism.

“thresholds.” If the system has two thresholds (i.e., it implements a ternary function), then its movements are finer grained. Below 45 degrees, arm motions qualify as “no movement.” Arm motions between 45 and 90 degrees qualify as “medium movement.” And above 90 degrees, arm motions qualify as “high movement.” However, if the system has only one threshold (i.e., it implements a binary function), then it gives rise to one of two types of movement. Table 5 describes the circuit as yielding arm movement only when both inputs are within the high voltage range (5–6 V). Table 4, however, describes the circuit as *not* activating arm movement only when both inputs are within the low voltage range (1–2 V). Table 4 describes an implementation of disjunction, and table 5 describes an implementation of conjunction. Why can the wide, proximal individuating strategy not distinguish between these two computational identities?

The short answer is that the “sensory” inputs and “effector” outputs are insufficient for distinguishing between the ternary versus binary versions of the circuit and between disjunction and conjunction. In both cases, the physical outputs are the same: the voltage outputs are the same and the corresponding arm movements are identical. Simply by observing the voltage inputs that the circuit receives, and the robotic arm movements that follow, a scientist cannot distinguish between the circuit’s computing conjunction or disjunction. Piccinini argues that computational systems are usually employed to perform their “maximal task,” that is, the task that exploits all the system’s functionally distinct inputs and outputs. If that is so, then the circuit’s computational identity is the ternary function.

Is this conclusion justified? A complete answer also depends on the relation between computational individuation and scientific explanation. Piccinini’s claim should be assessed relative to whether his view is committed to computational explanation being completely mechanistic (i.e., an explanation that specifies that the mechanism is thereby complete). If that is so, then his view of computational explanation may fit into one of two categories. For one category, computational explanation should also specify the causal mechanism that connects the relevant contextual factors and the computational system. For the other category, although computational explanation is mechanistic, at least some computational explanations also have nonmechanistic explanatory aspects (for a further discussion of these two possible views, see Harbecke and Shagrir 2019, sec. 5). It thus remains unclear how the robotic arm movements should be explained in the broader context.

And although the ternary computation is indeed a live option, there are reasons to be skeptical of computational systems always performing their maximal tasks. Human-engineered systems (e.g., the circuit plus the robotic arm) that do not perform their maximal tasks may certainly be “unnecessarily costly and cumbersome to build” (Piccinini 2015, 41–42). But, first, computational systems are not always engineered to accomplish their maximal tasks: the computer’s arithmetic logic unit (a subcomponent of the CPU) is typically designed with more inputs than are necessary to select from among its various functionalities. Hence, in many contexts, the arithmetic logic unit does not perform its maximal task. Second, evolutionary selection does not guarantee an optimal design of organisms (e.g., a predator may not be quick enough to catch its prey every time it is hungry). Therefore we do not claim that scientists should attribute *either* conjunction or disjunction to the aforementioned circuit. Some external factors may, however, be needed to computationally individuate systems

that exhibit abstract indeterminacy (e.g., does the circuit perform ternary or binary computation?). Our hypothetical plankton case study (in section 3.2) makes this claim even clearer.

Computational semanticists (e.g., Bishop 2009; Sprevak 2010; Harbecke and Shagrir 2019; Shagrir 2022)—who endorse wide, distal individuating strategies—argue that semantic properties are needed to individuate computational systems. On their view, even if some computational systems may be individuated by mechanistic properties, some must be individuated by semantic properties. Shagrir (2022), for example, distinguishes between *implementation* and *computation*. He holds that implementation is nonsemantic but claims that there is more to computation than implementation. According to Shagrir, computation is individuated semantically (210–11). To determine the computational identity of the aforementioned circuit, and like systems, in the context of theories of computation and cognition, Shagrir argues that one has to appeal to the contents of the system’s relevant physical states. Thus, for example, the computational identity of the binary version of the circuit (i.e., whether it computes conjunction or disjunction) depends on the (distal) content carried by the output voltage ranges (e.g., “no movement” vs. “movement”). Must we accept a full-blown representational individuating strategy? The short answer is no; the long one consists of two parts: sections 3.3 and 4.3.

3. A nonrepresentational functional individuating strategy

In this section, we formulate the functional strategy of computational individuation that appeals to wide, distal functions in the face of computational indeterminacy (section 3.3). Before proposing this strategy, we respond to a possible objection that proper functions (of traits) are always the most proximal ones (section 3.1); we then develop a hypothetical case study that shows how computational identity can be determined by wide, distal functions (section 3.2).

3.1. Are proper functions always the most proximal ones?

If narrow individuation does not suffice for individuating a system’s computational identity, one should go wide and possibly distal. A narrow individuation of the outputs of the tri-stable circuit that has the robotic arm attached is insufficient for determining whether the circuit is computing conjunction or disjunction. The proximal inputs (the different voltage levels that the circuit receives) are also insufficient for identifying the computation performed by the system. At this point, however, it is important to address a possible objection. It has been argued that, at least in the biological context, the proper function of a trait is its *most proximal* one (Garson 2019, chap. 7). Hearts beat, thereby enabling them to move blood around the body, which in turn enables the delivery of nutrients and oxygen throughout the body, thereby contributing to the organism’s survival. If no single function of the heart is privileged, a function indeterminacy problem arises. If we are to address computational indeterminacy by appealing to teleofunctions, it surely does not help if we are faced with another indeterminacy problem. As we note in section 3.3, when proximal functions are insufficient for computational individuation, distal functions are invoked. We thus cannot endorse the claim that the proper function of a trait is its most proximal one.

Although a full discussion of Garson's arguments exceeds the scope of this article, we make three relevant observations. First, as Garson (2019, 117) admits, restricting functions to the most proximal ones goes against standard practice in biology and is, in a sense, a revisionary proposal. It is common practice in all areas of biology and cognitive science to attribute functions at different levels of distality, so, in a sense, we are simply relying here on a default assumption in these disciplines.

Second, recognizing multiple functions is compatible with the claim that some of them are more central or explanatory than others. One could coherently hold that some functions are singled out as more important (and are even described as *the* function of a trait) without rejecting that other effects also qualify as genuine functions. Neander (1995, 120) and Papineau (1998), for example, hold that there are different possible functional descriptions of a given trait, though some of them have a certain priority over others. This view could be made compatible with our approach in different ways. One way is not to require that the function with highest priority is always the most proximal effect. Another is to hold that computational individuation need not be determined by the most important function but by the one that resolves the computational indeterminacy in the first place. Hence functional pluralism can include relevant differences (even hierarchies) between functions.

Finally, as we explain in section 5, one of the central properties of function attributions is their "explanatory depth": functions purport to give causal explanations for the existence of specific traits. If only the most proximal effects can qualify as functions, then functions are much less explanatory than usually thought. A more complete explanation of the existence of the pancreas, for example, mentions not only the secretion of hormones like insulin but also the fact that it helps regulate blood glucose, which in turn contributes to avoiding hyperglycemia. On a standard approach, all of these are functions of the pancreas (even though some of them require the activation of many other parts of the organism, of course), and all of them contribute to explaining the existence of the pancreas. On Garson's approach, only the first effect is functional, and hence function provides a much thinner explanation. All in all, we submit that an appeal to functions at multiple levels is fully justified.

3.2. The hypothetical plankton case study

We now introduce a hypothetical plankton-based system that is an analogue to the tri-stable circuit discussed earlier. Rather than responding to voltages as inputs, the plankton responds differentially to varying degrees of light intensity: no light, dim light, and bright light (see table 6). The relevant visuomotor system in the plankton may be likewise described as a tri-stable system with two thresholds: dim and bright light. Different combinations of inputs lead to three distinct outputs: drifting movement, sideways movement, and upward movement.⁵ Given that not all plankton

⁵ How are movements individuated? What determines that a particular change counts as an instance of sideways movement rather than of drift? A similar question arises for the robotic arm movements: why does a particular grouping of states, in the ternary case, qualify as "no movement," for example? (Perhaps this could be some noise threshold in a given range of, say, 0–5 degrees). This is an important metaphysical question that we cannot address here. In the plankton case, we are assuming that there is some "natural" way of grouping changes into movements that somehow carves nature at its joints. Whether the groupings of inputs and outputs are necessarily organism-dependent remains an open

Table 6. The plankton’s visuomotor system, which accepts three possible inputs (no light, dim light, and bright light) and yields three possible motor outputs (drifting movement, sideways movement, and upward movement)

First input (photoreceptor)	Second input (photoreceptor)	Motor output
No light	No light	Drifting movement
No light	Dim light	Sideways movement
No light	Bright light	Sideways movement
Dim light	No light	Sideways movement
Dim light	Dim light	Sideways movement
Dim light	Bright light	Sideways movement
Bright light	No light	Sideways movement
Bright light	Dim light	Sideways movement
Bright light	Bright light	Upward movement

Note. For the sake of brevity, the input and output values are already categorized (e.g., “dim light” and “upward movement”). However, a more accurate, lower-level physical description can also be given. For example, “no light” may correspond to less than 50 micromol quanta per meters squared per second (micromol quanta $m^{-2} s^{-1}$), “dim light” to 300–800 micromol quanta $m^{-2} s^{-1}$, and “bright light” to greater than 1,000 micromol quanta $m^{-2} s^{-1}$.

propel themselves in the manner posited by our construction, we assume, for simplicity, that the plankton described hereinafter is like the *arrow worm* (i.e., *Chaetognatha*). The output of the plankton’s visuomotor system is not of the same *type* as the input (i.e., it is not light).⁶ For ease of exposition, we also omit the details of the intermediate electrochemical outputs that are transduced to the relevant motor command (i.e., in a manner akin to the robotic arm from earlier). Our argument is not affected by this omission.

The main thrust of the argument concerns how the visuomotor’s abstract indeterminacy arises due to the possible ways that the varying degrees of light are grouped together. The first option is grouping together “no light” and “dim light” as “insufficient light” (typed as logical 0), whereas “bright light” is deemed “sufficient light” (typed as logical 1). If “drifting” and “sideways movement” are likewise grouped together and typed as logical 0 (and “upward movement” is typed as logical 1), the resulting description of this behavior, after omitting duplicate rows, corresponds to conjunction’s specification (see table 1). The second option is typing “no light” as logical 0 and grouping “dim light” and “bright light” as “sufficient light” (typed as logical 1). Furthermore, if “drifting movement” is typed as logical 0 and “sideways” and “upward

question. While plankton movements are more plausibly organism-dependent, it is unclear whether their grouping is functional (or even *wide* functional) or structural, for example. The grouping of inputs, however, may arguably even be organism-independent (e.g., the fact that fewer than x photons hit a given surface area at time t is independent of any organism—blind or not—that happens to be present in that area).

⁶ Likewise, the output of Shagrir’s tri-stable circuit is also not identical to its input: the former is movement, and the latter is voltages. In Harbecke and Shagrir’s (2019) tri-stable example of an aircraft marshaller—who conveys signals to pilots by moving either one arm or both arms simultaneously—the inputs are beeps, whereas the outputs are hand movements.

movement” are grouped together and typed as logical 1, the resulting description of this behavior (after omitting duplicate rows) corresponds to disjunction (see table 2). The two “competing” descriptions of the very same “physical” behavior are computationally distinct and thus indeterminate.⁷

3.3. Appealing to wide, distal functions, if needed

In a nutshell, our proposal is that if a narrow individuating strategy does not suffice for fixing the computational identity of a specific system, then the strategy needs to go wider and farther. The individuation of many computational systems (e.g., those computing a single mathematical function that is not susceptible to indeterminacy, such as the majority function) may not require any factors beyond the internal states, properties, processes, or structure of the system itself. However, as in the cases of both the tri-stable circuit and our hypothetical plankton, whether the system has the proximal function to compute one function or another (e.g., conjunction or disjunction), in *a given context*, may depend on its wide, distal function. Let us unpack this claim.

A simple, narrow analysis of the plankton may not be explanatory of the plankton’s behavior under the relevant environmental conditions. Its proximal behavior appears to be uniform when just looking at the system’s inputs and corresponding motor output.⁸ Recall that on the wide, proximal individuating strategy, the wideness of the relevant properties required for computational individuation of the plankton’s behavior would only reach the plankton’s sensory receptors, as well as its body and tail fins—responsible for propulsion and steering. However, such a strategy cannot explain the difference between the two competing mathematical functions that are at play here: conjunction and disjunction.

Our proposal is that the wider, distal function is required to fix the computational identity of the plankton’s visuomotor system in each context. Such an individuating strategy need not appeal to the representational contents of the plankton’s physical states—as semanticists argue with respect to the tri-stable circuit (Harbecke and Shagrir 2019). Weighing in possible ecological considerations may lead to discovering the different adaptive behaviors that are underpinned by the visuomotor system computing conjunction or disjunction. One such plausible consideration in the “conjunction scenario” may be searching for food closer to the water’s surface. Only when the plankton receives in the two photoreceptors “bright light” as input does it

⁷ Although conjunction and disjunction are dual functions, the indeterminacy in our example is of the abstract, and not the interpretative, type. They give rise to interpretative indeterminacy, just like when the mappings described in tables 4 and 5 are reversed. A similar construction could result in a competition between *inclusive* and *exclusive* disjunction, which are *not* dual functions (see, e.g., Papayannopoulos *et al.* 2022).

⁸ Although the present analysis focuses on simple, Boolean computation, it can be extended to algorithmic computation. Selection processes may likewise settle on a complex algorithm from sensory inputs to appropriate motor output. Computational teleofunctions—of the sort described herein in the plankton subsystems—likewise apply to “mediating states and processes in light of the computational contributions they make to the performance of the complex algorithm it is the teleofunction of the whole system to compute” (Coelho Mollo 2021, 6891). A teleo-based analysis of computation need not face a teleosemantic-like problem of content ascription that is devoid of explanatory purchase below some level of complexity (Cao 2012).

move upward toward the water's surface in search of food. Otherwise, it either drifts idly by (when no light is received as input) or moves sideways (when only dim light is received as input) as part of a random exploration process. On the other hand, a plausible consideration in the "disjunction scenario" is dispersing when conspecifics are abundant and nearby. If any light is received, the plankton does not simply drift idly by but rather moves either sideways or upward to distance itself from other conspecifics competing for the same resources. Thus the distal functions of the visuomotor system may fix its relevant computational identity in the presence of competing computational descriptions.

The present proposal differs from the wide, proximal individuating strategy à la Piccinini. The system's inputs may extend beyond the boundaries of its sensors, and its outputs may likewise reach out beyond its effectors (recall the amygdala example above). The proximal function of the plankton's visuomotor system is to respond in specific ways (i.e., movement patterns) to varying degrees of light in the environment. Its distal function, however, varies in different contexts (e.g., seeking food or avoiding competition). Therefore, pace Piccinini's view, the visuomotor system's outputs should be examined beyond the immediate interface between the effectors and the nearby environment: the system and the environment have been codesigned.⁹ There is also more explanatory pressure to attribute to the visuomotor system one of the binary computations depending on the competing behavioral explananda (i.e., food seeking or competition avoidance), rather than the ternary computation corresponding to some maximal task. The upshot is that a system's computational identity may be fixed by wider functional facts: the system's interaction with the wider, distal environment.

Before concluding this section, we respond to a potential worry about our appeal to wide, distal functions for computational individuation being ad hoc. What kind of principle unifies both narrow and wide strategies for individuation of computation?¹⁰ Our hypothetical example shows that by considering different wide functions, the plankton computes conjunction in one distal context and disjunction in another. Were they selected simultaneously? If not, it is hard to see (1) why both functions would be biologically justified and (2) how the proposed individuating strategy can be generalized.

The challenge is to show how the relevant ecological or evolutionary considerations are seamlessly integrated into our proposed individuating strategy. Different theoretical options are compatible with our main claim that in some cases computational individuation depends on wide functions. The proposal we have been developing thus far is that when abstract indeterminacy cannot be settled merely by appealing to narrow functions, computational individuation may, at least partly, depend on wide, distal functions. On the one hand, selected mechanisms have, in general, a hierarchy of functions rather than just one (section 3.1), and on the other hand, some computational systems are not subject to indeterminacy. Their computations can be individuated merely by inspecting the system itself. These

⁹ We note that our individuating strategy is about identifying the computation that is actually performed by the system and not about which evolutionary explanation fits each one of the possible computational descriptions.

¹⁰ We thank Coelho Mollo for raising this objection.

two considerations combined suggest that narrow, proximal functions should be considered first. (In section 4.1, we also discuss the possibility that a system may actually compute two proximal nondual functions simultaneously.) Only when such an individuating strategy has been exhausted should wide functions be considered. For similar reasons, wide proximal functions ought to be considered before more distal ones. The upshot is that narrow individuation is privileged over wide individuation, and likewise for proximal functions over more distal functions.¹¹

Different wide functions could bear important relations between them. Suppose that the seek-food behavior was selected first (i.e., being the beneficial effect of exploiting the physical structure for computing conjunction). In that context, conjunction is the computational identity of the visuomotor system—despite it also physically supporting disjunction. So, although the seek-food behavior was originally selected, the avoid-competition behavior may, after a while, be coselected: no further physical change to the visuomotor system is required. (In section 4.1, we introduce a slightly more intricate version of the plankton case to further elaborate this last point.) Because phylogenetic selection takes time, it is also possible that the two behaviors coevolve. In sum, the same system can “evolve” to have two or more wide, distal functions, depending on the specific context.

4. Responding to objections and elaborating the functional individuating strategy

In this section, we address four objections to the functional individuating strategy. Showing how our individuating strategy handles these problems not only buttresses the strategy but also helps clarify important nuances about it. In section 4.1, we respond to an objection according to which appealing to downstream consumers within a given system enables narrow individuation in the face of indeterminacy. In section 4.2, we respond to an objection according to which our proposed strategy entails that computation is always contextual. In section 4.3, we conclude our response to the claim that the functional individuating strategy is a semantic one. And finally, in section 4.4, we respond to the Swamp-consumer objection.

4.1. *Narrow individuation (sometimes) suffices for computational individuation*

The first objection is that in plausible biological scenarios, narrow individuation should suffice for determining the identity of a computationally indeterminate system by appealing to downstream consumers of that system. To unpack this claim, we need to introduce some further complexity to the hypothetical plankton example as described in figure 2. The computational identity of the visuomotor system need not be fixed widely but rather by downstream consumers within the plankton that simply exploit one function or another.¹² If a consumer is designed with an activation threshold above “drifting movement,” such that “sideways” and “upward movement” are grouped together, then it exploits the computation of disjunction. If it is designed with an activation threshold above “sideways movement,” then it exploits the computation of

¹¹ However, one may accept that wide functions are always privileged. Fully assessing these, or other, alternatives exceeds the scope of this article.

¹² We thank Copeland for raising this objection.

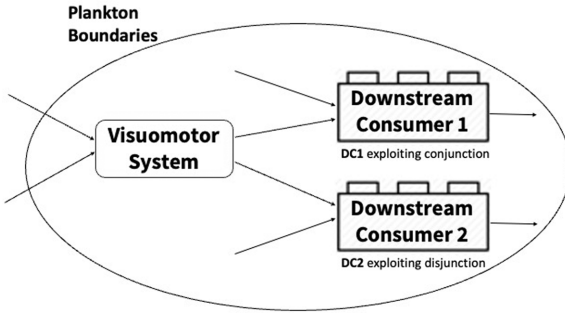


Figure 2. In the present version of the hypothetical plankton, the visuomotor system is connected to two downstream consumers: DC1 and DC2. DC1 exploits the conjunction computation, whereas DC2 exploits the disjunction computation that the visuomotor system performs.

conjunction. Once a specific computation is exploited in this interaction between the visuomotor system and the downstream consumer, the computational identity is fixed, thereby making redundant any appeal to the plankton's behavior in the environment.

First, we agree that in cases in which narrow individuation is sufficient for fixing the computational identity of a given system, we should look no further. As already noted in section 3.3, the computational identity of some systems may be fixed by the internal states, properties, processes, or structure of the system itself. Thus, if the computational identity of the downstream consumer (e.g., DC1 or DC2 in figure 2) is fixed, then we agree with the preceding description of the interaction between the visuomotor system and the downstream consumer as fixing the computational identity of the former. However, if the downstream consumer itself also exhibits computational indeterminacy (e.g., by computing either exclusive disjunction or logical equivalence—which are dual Boolean functions), then it is unclear whether we can avoid appealing to wide, distal functions to fix the computational identity of the systems concerned (i.e., both the producer and the consumer).

Second, on our proposed strategy, the interaction between a computational system S and its downstream consumer, if such exists, partly determines S' computational identity. In the present case, the visuomotor system may support two computations simultaneously due to its physical properties with the appropriate degrees of freedom on which its state transitions depend (Coelho Mollo 2018, 3482). It *has the function* to compute conjunction (or, likewise, disjunction) if it has the right evolutionary history (depending on the downstream consumer with which it has worked). When DC1 exploits the computation of conjunction, the visuomotor system *actually* computes conjunction. When DC2 exploits the computation of disjunction, the visuomotor system actually computes disjunction. This is so because in the right context, the interaction between the visuomotor system and the downstream consumer enables the successful performance of the visuomotor system's (teleological) function to compute one of these two mathematical functions. If in *another* context (different from the one described earlier), both downstream consumers exploit the computation of two nondual functions simultaneously (e.g., conjunction and exclusive disjunction), then the visuomotor system works as two logic gates (e.g., an AND-gate and a XOR-gate) concurrently; the indeterminacy of computation dissolves. Computation is, thus, contextual.

4.2. Computation is contextual

Our response to the second objection, which also draws on the construction depicted in figure 2, is that, yes, computation is inherently contextual. We can identify a significant amount of support for embracing this objection. Suppose that the plankton contains only one of the two downstream consumers, say, DC1, which exploits the computation of conjunction. Let us further assume that half of the plankton population somehow relocates to a different continent, where it continues to thrive, explore, and reproduce. Suppose that in the migrated plankton, DC1 still contributes to the plankton's fitness but does so without regard to the computational output of the visuomotor system. In this case, our strategy entails that the visuomotor system computes in the original environment but not in the new environment. That is so even though nothing has changed in the plankton's physical structure.¹³

Given that the plankton's visuomotor system—with the right evolutionary history—has the function to compute conjunction, it actually computes conjunction in its original environment when DC1 exploits that computation. In the new environment, although it still has that teleological function, it does not actually compute conjunction, because that mathematical function is not exploited. This is a consequence of our proposed strategy being grounded in an etiological notion of teleological function. Similarly, over sufficient evolutionary time, the visuomotor system may even lose its (teleological) function to compute conjunction, if it is not exploited. The visuomotor system—in this type of plankton—has the function of computing conjunction by virtue of belonging to the type *visuomotor*, which has been selected for computing conjunction.

To generalize the conclusions drawn from our plankton examples, we introduce the *contextual computation thesis*: there is no actual computation, in itself, outside of a functional context.¹⁴ Computation is an effect of a system that has been selected for during short- or long-term evolutionary change to process medium-independent vehicles according to a rule F (corresponding to a mathematical function f) that is sensitive only to the vehicles' degrees of freedom. The teleological function of the visuomotor system is to compute conjunction (and possibly disjunction in some versions of our example) due to the evolutionary history that shaped the system to maximize the adaptive benefits to the plankton in exploiting that computation in its environment. In the original construction discussed in section 3.2, the visuomotor system computes conjunction and disjunction because of how the plankton's behavior in the environment depends directly on it. In the modified version discussed in section 4.1, it computes conjunction, disjunction, or both, depending on how the two downstream consumers exploit its computations. And in the present section, it computes conjunction in one environment but does not compute at all in the other

¹³ We thank Shenker and Hemmo for raising this objection.

¹⁴ Our contextual view of computation individuation aligns with the noncausal contextual view of computational explanation (Harbecke and Shagrir 2019), agreeing that satisfactory computational explanations should typically consider contextual features. Nevertheless, arguably, not all relevant contextual factors determine computational identity. Sometimes, narrow, proximal functions suffice, and when wide, distal functions are needed, many contextual factors (e.g., the temperature at the water's surface for our plankton) will not impact the system's computational identity. Further discussion exceeds the scope of this article.

environment. These are the implications of our proposed individuating strategy being functional and etiological (more on that in section 5).

4.3. The functional individuating strategy is (not) representational

In this section, we conclude our response to the question raised in section 2 concerning whether our individuating strategy is a full-blown representational one. The present objection is that our wide, distal individuating strategy implicitly appeals to representations. Our account sometimes individuates computational states by appealing to wide functions, and according to the most popular naturalistic theory of representation, teleosemantics, representations are grounded on wide functions. As a result, by appealing to wide functions, we might inadvertently commit to a semantic theory of computational individuation. Thus, although one of our aims was to avoid a full-blown semantic individuation of computational states, our appeal to wide functions implies such a view.

In response, we make two observations. First, our proposal does not imply the semantic view: we certainly individuate computational mechanisms by appealing to functions, but it is unlikely that functions alone ground representations (even according to most teleosemantic theories). Hearts, wings, and legs have functions, but they do not represent. According to most teleosemantic theories of content, representation also requires some *other condition* (e.g., sender–receiver systems, constancy mechanisms, or robust tracking). Sterelny (1995), for example, claims that a distinctive feature of representational states is the tracking of the same *distal* stimulus using *diverse* cues. Similarly, Schulte (2015) holds that, in addition to the system's function of tracking some environmental feature, constancy mechanisms are required. Millikan (1984), Shea (2018), and others appeal to sender–receiver systems, among other conditions. Functional individuation of computation is not thereby semantic. Thus our proposal requires *less* than a full-blown semantic theory to address the indeterminacy problem.

Second, although our account is not committed to the semantic view, it is compatible with it. Accordingly, in principle, one could hold the semantic view and accept our arguments. Nonetheless, one of the central motivations for endorsing a semantic view of computation depends on its capacity to address the indeterminacy problem outlined previously. If our account can satisfactorily solve this problem by relying only on functions, then our approach undermines one of the key motivations for the semantic view.

4.4. The Swamp-consumer

The final objection to be addressed is the Swamp-consumer.¹⁵ Our functional strategy assumes that the plankton has undergone a similar process of natural selection as its conspecifics. Thus the visuomotor system might have the teleological function to compute disjunction, conjunction, or both, depending on the downstream consumers that have evolved to exploit its computations. Suppose that, in line with the construction in section 4.1, the plankton has a visuomotor system and two downstream consumers that exploit the computations of conjunction and disjunction.

¹⁵ We thank Shagrir for raising this objection.

Suppose further that in another plankton in which only DC1 has evolved, a Swamp-consumer, which is a physical duplicate of DC2, springs into existence. Unlike DC1, the Swamp-consumer has not undergone a selection process. The visuomotor system, to which the Swamp-consumer is attached, supports both conjunction and disjunction—due to its physical properties with the appropriate degrees of freedom. However, it has the teleological function to compute only conjunction, and it works as an AND-gate when DC1 exploits this computation. The problem is that when comparing the plankton with the evolved DC1 and DC2 and the plankton with the Swamp-consumer, their *computational capacities differ*, although they are *physically identical* by stipulation.

We accept this implication. In the plankton with the evolved DC1 and DC2, the visuomotor system has the teleological functions to compute conjunction and disjunction. In the plankton with the evolved DC1 and the Swamp-consumer, the visuomotor system has the teleological function to compute only conjunction, and it works as an AND-gate (when properly exploited by DC1), even if the Swamp-consumer exploits the computation of disjunction. However, there is room for this to gradually change as the plankton interacts with its environment over ontogenetic time. The visuomotor system may acquire the additional teleological function to compute disjunction, if the Swamp-consumer systematically exploits this computation such that it contributes to the plankton's fitness. Such acquisition of a new teleological function is supported by different teleological theories of function (e.g., Artiga 2021; Garson 2019; Shea 2018).

5. Constraints on a suitable notion of teleofunction

Our individuating strategy relies on some notion of teleofunction, but which theory of function best suits our purposes? We identify some specific features (or desiderata) that any satisfactory theory of function should possess if it is to play the envisaged individuating role.

Three constraints on any theory of function—including nonetiological notions of function—are widely endorsed (Garson 2019, chap. 1): (1) the function/accident distinction, (2) normativity, and (3) explanatory depth. Let us briefly specify them.

1. **Function/accident distinction.** *Functions should be distinguished from accidental beneficial effects.* Functions are specific kinds of effects, but not every effect counts as a function. They should be distinguished from accidentally beneficial effects. Our proposal requires that computations (e.g., conjunction and disjunction) be functions of the relevant systems (e.g., the plankton's visuomotor system), rather than mere effects. Any adequate theory of function should draw this distinction appropriately.
2. **Normativity.** *It should be possible for an item to fail to perform its function.* Functions are effects that an item is *supposed* to have; failing to perform this effect qualifies as either a malfunction or a dysfunction. Accounting for this aspect is especially important because a primary reason for employing the notion of function is to accommodate the possibility of *miscomputation*. If the plankton propels itself upward to the water's surface when "no light" is received in both its photoreceptors, the visuomotor system miscomputes.

3. **Explanatory depth.** *Functions play an important explanatory role: they purport to give causal explanations for the existence of specific traits.* Although the exact nature of this explanatory role is contentious, a standard way of understanding it is that functions partly explain the existence of the trait. The fact that hearts pump blood partly explains why many living beings have this organ, whereas making thumping noises fails to play this explanatory role. Many effects fail to count as proper functions partly because they typically fail to exhibit this explanatory depth. In our basic plankton example, some visuomotor systems computed conjunction and disjunction, and some did not. Those that did contributed to the plankton's fitness by better sourcing food and avoiding competition with conspecifics. Owing to the mechanisms of inheritance, computational visuomotor systems were maintained in the population.

To these three standard features we add another important one deriving from our preceding discussion: the proximal/distal distinction.

4. **Proximal/distal distinction.** *The theory of function should be compatible with the existence of proximal and distal functions.* Our individuating strategy entails that an adequate solution to computational indeterminacy must distinguish between proximal and distal functions. For instance, in the plankton case in section 3.2, the proximal function is to compute conjunction and disjunction, whereas the distal functions are moving to the water's surface to seek food and moving away from conspecifics. On the one hand, computing conjunction is a function that the visuomotor system is supposed to perform, and failing to have this effect counts as a miscomputation. On the other hand, this function is determined by the distal effects it has in a wider environment (beyond the plankton's body). However, not every effect determines computational identity: only effects that the computational system is *supposed* to have count. Hence an adequate approach to function should accommodate a distinction between proximal and distal functions.

Many strategies for addressing computational indeterminacy might be committed to some sort of distinction between proximal and distal functions. If computation is a function of specific systems, and addressing computational indeterminacy requires appealing to some other factor, this additional aspect often turns out to depend on functions. The semantic view, for example, would be committed to the same distinction if its proponents were to endorse a teleosemantic theory of representation. If computational systems have the function to compute (as the mechanistic view claims), computational identity depends on representational content, and representational content is defined in functional terms, then one needs to distinguish between proximal (i.e., computational) and distal (i.e., representational) functions. Therefore other theories of computation may have to make a similar distinction between proximal and distal functions.

These desiderata combined constrain the list of theories of function suitable for use with our individuating strategy. Garson (2019, 93), for one, develops a promising theory, according to which "function of a trait is an activity that led to its differential reproduction, or its differential retention in a population." Unfortunately, while

others have accepted that proximal functions have some sort of privileged status compared with distal ones (e.g., Neander 1991, 2017; Papineau 1993), on his theory, there are no distal functions, only proximal ones. Hence Garson's theory does not meet the proximal/distal distinction desideratum, so it is incompatible with our framework. However, we assert, without further argument,¹⁶ that some version of selected effects theory would do the trick.

6. Conclusion

When a system has two (or more) possible computational identities that are equally epistemically legitimate (because they are simultaneously implemented), some factor is needed to fix the system's computational identity. We have argued that computational identity sometimes depends only on the system's narrow properties—thereby siding with one version of the mechanistic strategy of computational individuation—and sometimes it *also* depends on wider properties. These wider properties, however, are *wide, distal functions*, rather than *semantic* properties—as the semantic strategy of computational individuation asserts. Our claim requires a theory of etiological function that (1) adequately distinguishes between proximal and distal functions and (2) gives distal functions a pivotal place in the theory. Under our proposed strategy, first, computation is contextual (in a functional sense), and second, the problem of computational indeterminacy dissolves: a physical system may compute two mathematical functions simultaneously

Acknowledgements. Our gratitude goes to all those who contributed to this article both through lively discussions and helpful comments on earlier versions. They include Dimitri Coelho Mollo, Jack Copeland, Lotem Elber-Dorozko, Jens Harbecke, Meir Hemmo, Marcin Miłkowski, Jonathan Najenson, Manolo Martinez, Gualtiero Piccinini, Michael Rescorla, Richard Samuels, Oron Shagrir, and Orly Shenker. Thanks are also due to the audience at the 2023 Dubrovnik 'Representation and Computation in the Study of Language and Mind' conference and in the 2023 ISF 'The Indeterminacy of Computation' workshop. We are grateful to several anonymous reviewers for *Philosophy of Science* (and thank anyone else whom we may have inadvertently omitted). This research was supported by ISF grant 386/20 to the first author. Part of this research was undertaken in the context of the MeReX research project funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – project number: 520194884. Research conducive to this article was also made possible by the projects “The Representational Penumbra” (PID2021-127046NA-I00) funded by the Ministerio de Ciencia, Innovación y Universidades and the projects “Deceptive Representations” (CISEJI/2023/51) and “Autonomy as Address” (PROMETEO-CIPROM/2023/55) funded by the Conselleria d'Educació, Universitats i Ocupació, Generalitat Valenciana.

References

- Artiga, Marc. 2021. “Biological Functions and Natural Selection: A Reappraisal.” *European Journal for Philosophy of Science* 11 (2):54. <https://doi.org/10.1007/s13194-021-00357-6>.
- Bishop, John Mark. 2009. “A Cognitive Computation Fallacy? Cognition, Computations and Panpsychism.” *Cognitive Computation* 1 (3):221–33. <https://doi.org/10.1007/s12559-009-9019-6>.
- Block, Ned. 1990. “Can the Mind Change the World?” In *Meaning and Method: Essays in Honor of Hilary Putnam*, edited by G. Boolos, 137–70. Cambridge: Cambridge University Press.
- Cao, Rosa. 2012. “A Teleosemantic Approach to Information in the Brain.” *Biology and Philosophy* 27 (1):49–71. <https://doi.org/10.1007/s10539-011-9292-0>.

¹⁶ This remains a task for future work.

- Coelho Mollo, Dimitri. 2018. "Functional Individuation, Mechanistic Implementation: The Proper Way of Seeing the Mechanistic View of Concrete Computation." *Synthese* 195:3477–97. <https://doi.org/10.1007/s11229-017-1380-5>.
- Coelho Mollo, Dimitri. 2021. "Why Go for a Computation-Based Approach to Cognitive Representation." *Synthese* 199:6875–95. <https://doi.org/10.1007/s11229-021-03097-5>.
- Curtis-Trudel, André. 2022. "The Determinacy of Computation." *Synthese* 200:43. <https://doi.org/10.1007/s11229-022-03568-3>.
- Dewhurst, Joe. 2018. "Individuation without Representation." *British Journal for the Philosophy of Science* 69 (1):103–16. <https://doi.org/10.1093/bjps/axw018>.
- Fresco, Nir. 2021. "Long-Arm Functional Individuation of Computation." *Synthese* 199:13993–4016. <https://doi.org/10.1007/s11229-021-03407-x>.
- Fresco, Nir, B. Jack Copeland, and Marty J. Wolf. 2021. "The Indeterminacy of Computation." *Synthese* 199:12753–75. <https://doi.org/10.1007/s11229-021-03352-9>.
- Garson, Justin. 2019. *What Biological Functions Are and Why They Matter*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/9781108560764>.
- Harbecke, Jens, and Oron Shagrir. 2019. "The Role of the Environment in Computational Explanations." *European Journal for Philosophy of Science* 9 (3):37. <https://doi.org/10.1007/s13194-019-0263-7>.
- Millikan, R. (1984). *Language, Thought and Other Biological Categories*. Cambridge: The MIT Press.
- Neander, Karen. 1991. "Functions as Selected Effects: The Conceptual Analyst's Defense." *Philosophy of Science* 58 (2):168–84. <https://doi.org/10.1086/289610>.
- Neander, Karen. 1995. "Misrepresenting and Malfunctioning." *Philosophical Studies* 79 (2):109–41. <https://doi.org/10.1007/BF00989706>.
- Neander, Karen. 2017. *A Mark of the Mental: In Defense of Informational Teleosemantics*. Cambridge, MA: MIT Press. <https://doi.org/10.7551/mitpress/9780262036146.001.0001>.
- Papayannopoulos, Philippos, Nir Fresco, and Oron Shagrir. 2022. "On Two Different Kinds of Computational Indeterminacy." *The Monist* 105 (2):229–46. <https://doi.org/10.1093/monist/onab033>.
- Papineau, David. 1993. *Philosophical Naturalism*. Malden, MA: Blackwell.
- Papineau, David. 1998. "Teleosemantics and Indeterminacy." *Australasian Journal of Philosophy* 76 (1):1–14. <https://doi.org/10.1080/00048409812348151>.
- Piccinini, Gualtiero. 2015. *Physical Computation: A Mechanistic Account*. Oxford: Oxford University Press.
- Piccinini, Gualtiero. 2020. *Neurocognitive Mechanisms: Explaining Biological Cognition*. Oxford: Oxford University Press.
- Schulte, P. 2015. "Perceptual Representations: A Teleosemantic Answer to the Breadth-of-Application Problem." *Biology & Philosophy*, 30(1), 119–136.
- Shagrir, Oron. 2001. "Content, Computation and Externalism." *Mind* 110 (438):369–400. <https://doi.org/10.1093/mind/110.438.369>.
- Shagrir, Oron. 2020. "In Defense of the Semantic View of Computation." *Synthese* 197:4083–108. <https://doi.org/10.1007/s11229-018-01921-z>.
- Shagrir, Oron. 2022. *The Nature of Physical Computation*. Oxford: Oxford University Press. <https://doi.org/10.1093/oso/9780197552384.001.0001>.
- Shea, Nicholas. 2018. *Representation in Cognitive Science*. Oxford: Oxford University Press.
- Sprevak, Mark. 2010. "Computation, Individuation, and the Received View on Representation." *Studies in History and Philosophy of Science, Part A* 41 (3):260–70. <https://doi.org/10.1016/j.shpsa.2010.07.008>.
- Sterelny, Kim. 1995. "Basic Minds." *Philosophical Perspectives* 9 (A1):251–70. <https://doi.org/10.2307/2214221>.
- Tucker, Chris. 2018. "How to Explain Miscomputation." *Philosophers' Imprint* 18 (24):1–17.