ORIGINAL RESEARCH

Long-arm functional individuation of computation

Nir Fresco[1,2,3]

Abstract

A single physical process may often be described equally well as computing several different mathematical functions—none of which is explanatorily privileged. How, then, should the computational identity of a physical system be determined? Some computational mechanists hold that computation is individuated only by either narrow physical or functional properties. Even if some individuative role is attributed to environmental factors, it is rather limited. The computational semanticist holds that computation is individuated, at least in part, by semantic properties. She claims that the mechanistic account lacks the resources to individuate the computations performed by some systems, thereby leaving interesting cases of computational indeterminacy unaddressed. This article examines some of these views, and claims that more cases of computational indeterminacy can be addressed, if the system-environment interaction plays a greater role in individuating computations. A new, long-arm functional strategy for individuating computation is advanced.

Keywords Computation · Individuation · Teleological function · Mathematical function · Indeterminacy · Cognitive science · Semantics · Mechanism · Computational explanation

✉ Nir Fresco
    nfresco@bgu.ac.il

[1] Present Address: Department of Brain and Cognitive Sciences, Ben Gurion University of the Negev, 8410501 Beer Sheva, Israel

[2] Department of Philosophy, Ben Gurion University of the Negev, Beer Sheva, Israel

[3] Zlotowski Centre for Neuroscience, Ben Gurion University of the Negev, Beer Sheva, Israel

1 Introduction

Computational explanations are typical in the cognitive sciences. The identification of *the* mathematical function being computed by a physical system—be that a brain circuit, or a single neurone—may be complicated by the fact that some such functions have *other* "isomorphic copies"—a term that will shortly be clarified using a simple example. Thus, any computational description that is based on identifying such a mathematical function as part of explaining the explanandum may be based on an indeterminate computation.

If that is true, then computational explananda in the cognitive sciences are likewise susceptible to exhibiting this phenomenon. Neurocognitive explanations regularly confront the question: 'What does a specific neural structure do and how does it do it?'. Answers that hypothesise that the structure concerned computes some specific function, which has an isomorphic copy, should arbitrate between the possible competing functions. At least prima facie, neither mathematical function is epistemically privileged as a description of the structure's physical behaviour. We turn next to describe a simple physical system whose behaviour is multiply specifiable using two Boolean functions, yet this phenomenon is not limited to *Boolean* functions (see Fresco et al., 2021).

Consider a simple, electrical Boolean gate *G* with two input-channels and a single output-channel. *G*'s physical behaviour is described in Table 1.

If the voltage range 1–3 V represents False and 4–6 V represents True, Table 1 turns out to be the standard truth-table for Boolean conjunction. *G*, thus, computes conjunction (or is an AND-gate). However, if the voltage range 1–3 V represents True and 4–6 V represents False, *G* computes an isomorphic copy of conjunction, namely: inclusive disjunction (or is an OR-gate). (Conjunction and disjunction are considered *dual* functions in Boolean logic.) The gate's computing either conjunction or disjunction illustrates the indeterminacy of computation. A similar moral applies to an electric gate that computes either the NAND or the NOR function and a gate that computes either the XOR or the XNOR function (and to other dual Boolean functions).

There seems to be agreement that it is necessary to identify relevant properties that determine what computation a given physical system, *P*, actually performs. In philosophy of computation, whether a given account of computation is able to settle the indeterminacy of computation when it arises has been deemed a litmus test for the adequacy of that account. However, what those properties *are* is an open question. Bishop (2009), Sprevak (2010), and Shagrir (2020), for example, appeal to *semantic* properties that render it determinate as to what computation is performed by *P*. Some proponents of the mechanistic view of computation appeal

Table 1 Electrical gate *G*'s inputs and corresponding outputs specified in volts	Input-channel 1	Input-channel 2	Output-channel
	1–3 V	1–3 V	1–3 V
	1–3 V	4–6 V	1–3 V
	1–3 V	1–3 V	1–3 V
	4–6 V	4–6 V	4–6 V

to *narrow physical* or *functional* properties for computational individuation. Dewhurst, for example, argues that the relevant properties that render it determinate as to what *P* computes are purely *physical* (2018). Coelho Mollo argues that the computational identity of *P* is determined by *functional* properties (2017).

The main claims defended in this article are that (1) computational individuation is (often) wide, and (2) the inputs and/or outputs of the computation may not be physically realised in the computing system itself. These claims yield a new, long-arm functional individuative strategy. The underlying idea is that some interesting cases of indeterminacy in biological systems cannot always be settled by appealing to activities in the relevant sensory receptors and motor neurones—as the short-arm functional individuative strategy suggests. Real things in the world may count as the inputs and outputs of the computation performed (see also Table 5 below). As a backdrop for developing this strategy, the main contenders in the debate on computational individuation are systematically examined.

The article is organised as follows. Section 2 discusses two influential mechanistic strategies that individuate computation very narrowly. In Sect. 3, we examine the semantic individuative strategy, according to which semantic constraints are needed, at least in some interesting cases of computational indeterminacy. Section 4—the central part of this article—advances a long-arm functional individuative strategy as an intermediate between the short-arm mechanistic view and the semantic view of computational individuation. The long-arm strategy is defended against the objection that it collapses into the semantic strategy (Subsection 4.4). Section 5 responds to another important challenge to the long-arm functional individuative strategy. Section 6 concludes the article.

2 Mechanistic individuation of computation

Although some mechanists agree that the system-environment interaction can play a role in computational individuation (e.g., Coelho Mollo, 2019; Miłkowski, 2017; Piccinini, 2015), they diverge on *what* this role is. In this section, we examine two different mechanistic positions. But, first, what do we mean by a mechanism? In a nutshell, a "mechanism for a phenomenon consists of entities and activities organised in such a way that they are responsible for the phenomenon" (Illari & Williamson, 2012, p. 120). The mechanistic explanatory strategy, then, is to decompose the explanandum into its spatiotemporal constituent parts, and to discover how their causal interactions and structural relations are responsible for producing (or maintaining) the explanandum.

When it comes to computational phenomena, proponents of the mechanistic view diverge on the type of properties that are relevant for individuation and how these properties contribute to avoiding potential indeterminacies. Table 1 above may be said to provide a systematic specification of the system's physical behaviour: in response to voltages in specific ranges, the system produces voltages in a specific range. Once the voltage ranges are mapped onto logical True (or 1)

Table 2 Hydraulic gate *H*'s inputs and corresponding outputs specified in water pressure measured in Litre per Second (LpS)	Input-channel 1	Input-channel 2	Output-channel
	0.1–0.5 LpS	0.1–0.5 LpS	0.1–0.5 LpS
	0.1–0.5 LpS	1–1.5 LpS	0.1–0.5 LpS
	1–1.5 LpS	0.1–0.5 LpS	0.1–0.5 LpS
	1–1.5 LpS	1–1.5 LpS	1–1.5 LpS

and False (or 0), the corresponding table provides a logical specification of the system's behaviour (describing either conjunction or disjunction). Here, computational indeterminacy arises as described above. Let us, next, examine two different mechanistic individuative strategies for avoiding such indeterminacy.

2.1 Dewhurst's individuative strategy: losing computational equivalence.

Dewhurst proposes to individuate computation on the basis of physical properties alone (2018). As a computational mechanist, he can specify the computational identity of a given system in virtue of three key ingredients (Fresco & Miłkowski, 2019). The first ingredient is digits: what are the unique digits processed by the system and how many are there? (The answer is '2' in relation to gate *G* above: [1–3 V] is one digit, and [4–6 V] is the second, distinct digit.) The second ingredient is the processing unit(s) that operates on these digits: how many processing units are in use by the given system? (The answer is '1' in relation to *G*. But a more complex computing system that comprises many Boolean circuits may have many distinct processing units.) The third, and important, ingredient is the input–output relations in which the digits partake in the encompassing system. (Table 1 specifies *G*'s input–output relations.) On Dewhurst's view, the computation performed by *G* can be fully spelled out in terms of these three ingredients: computational individuation does not require any logical or semantic content.

The motivation for this individuative strategy becomes apparent once we realise that the type of indeterminacy described above does not even manifest itself. The indeterminacy arises *only* once the system's behaviour is *logically* specified. *G*—as it is described by Table 1—indeterminately computes conjunction *or* disjunction depending on how the variables (voltage ranges) are mapped onto True and False. But *if* Table 1 (plus 'digits' and 'processing units') provides all the theoretical posits necessary for individuating *G*'s computation, then at least prima facie—as Occam's razor dictates, this individuative strategy is appealing: computation is individuated by non-semantic transformations of digits, and indeterminacy is thereby avoided.

But the economic efficiency of this strategy comes at a cost, namely giving up computational equivalence and multiple realisability. Computational equivalence is a central idea in computer science: two physical systems may compute the same function even if the physical magnitudes on which they operate are different. One system may traffic in voltages, the other in *different* voltages or even fluid pressure, and they may still both compute logical conjunction. Thus, the hydraulic gate *H*,

described in Table 2, may be computationally equivalent to *G*. If both the low voltage range and, likewise, the low water pressure are mapped onto False, *G* and *H* compute conjunction (and *mutatis mutandis* they may both compute disjunction). Since Tables 1 and 2 describe *distinct* physical behaviours of the respective systems, Dewhurst's strategy individuates them as *different* computations.

The problem cuts even deeper. For "the physical structure of two computing mechanisms is always going to be distinct, and it is unclear whether we can draw any non-arbitrary boundary between the structures that are relevant or irrelevant to computational individuation" (Dewhurst, 2018, p. 110). Any two conventional AND-gates in one's smartphone—made of the same materials, based on the same blueprint, by the same manufacturer—turn out to be *computationally distinct* (since any minute difference in their voltage range is enough). The idea of computational equivalence is, thus, lost in Dewhurst's strategy.

The closely related idea of multiple realisability is likewise threatened by this individuative strategy.[1] Cognitive explananda that are multiply realisable bestow an explanatory edge to *causes* over physical *constituents*. For the physical constituents in one realisation that constitute a given cognitive phenomenon (e.g., visual object identification) will not necessarily be the constitutive elements in another realisation of the same phenomenon. "There is little reason to believe that cognitive and neural entities and activities must be similarly organized. In complex systems, what looks stable and robust at one scale may not be so at another scale" (Stinson, 2016, p. 1603). Insofar as *distinct* neural structures can give rise to the *same* cognitive function by computing a *specific* mathematical function, the mathematical computation may be more stable as a cause than the particular constituents. The present computational individuative strategy is incompatible with the common and compelling explanation of multiple realisability of cognitive functions in terms of computational functions.

In sum, for those who think that computational equivalence and multiple realisation should be preserved as important principles in the computational sciences, including cognitive science, "the physical level is [simply] too fine-grained" (Coelho Mollo, 2017, p. 3493) and so the present individuative strategy fails to deliver the goods.

2.2 Coelho Mollo's individuative strategy: identifying equivalence classes.

Realising that this is too great a price to pay, Coelho Mollo supplements Dewhurst's strategy with a *teleofunctional* analysis of computing systems "in which the only structural considerations at play are having appropriate degrees of freedom" (Coelho Mollo, 2017, p. 3494). By classifying computational phenomena as a proper subset of teleofunctional phenomena, Coelho Mollo's individuative strategy gains an important explanatory advantage; it draws a boundary between computing and

---

[1] Some computational mechanists deny that multiple realisability is an essential feature of physical computation (e.g., Miłkowski, 2016).

non-computing systems. Thus, planetary motions, hurricanes, and tides are excluded as non-computational phenomena.

How does Coelho Mollo's individuative strategy work? In essence, it draws on the principle of 'equivalence classes', which is a technical notion in logic. In this strategy, however, "[e]quivalence classes are defined by input values that lead to uniform behaviour of the whole device—the differences in value to which the device is sensitive and which are uniformly transformed into new values" (ibid). To see how this definition works, consider another gate, *G**, which is very similar to *G* (described by Table 1), but whose voltage ranges are (2–4 V) and (5–7 V) instead. *G** and *G* are computationally equivalent. Why? Because both *G* and *G** respond to two distinct equivalence classes of acceptable physical inputs (voltages) and yield the same equivalence classes of physical outputs (voltages) in response. Each such equivalence class is a *digit*. In *G*, the first equivalence class (or digit) is (1–3 V) and in *G** it is (2–4 V); *G*'s second equivalence class is (4–6 V), and *G**'s is (5–7 V). Thus, to some extent computational equivalence is preserved on Coelho Mollo's individuative strategy.

Another explanatory advantage of this strategy is that computational equivalence also holds between systems of different physical makeups. The hydraulic gate *H* above (described by Table 2) is functionally equivalent to *G* and *G**. Because what matters to computational individuation is the overall functional profile that defines these three gates. *H* shares the same functional profile of *G* and *G**, since it is sensitive to, and responds uniformly and in the same way to, the same number of equivalence classes. Whilst the *physical* descriptions of *G*, *G**, and *H* are clearly *distinct*, their *functional* descriptions are *identical*. Whether the equivalence classes are based on voltages or water pressure is irrelevant for computational individuation.

Nevertheless, this individuative strategy raises two main worries. The first one concerns the notion of an equivalence class.[2] In logic, an equivalence relation—over a given set A—is one that is reflexive, symmetric, and transitive.[3] That is, it satisfies specific logical properties. An equivalence relation divides A into equivalence classes based on these properties. The equivalence relation in Coelho Mollo's strategy, though, is not characterised as rigorously. 1.32 V and 2.31 V, for example, are classified as belonging to the same equivalence class with respect to *G*. Because *G*'s behaviour is sensitive to, and responds in the same way to both values. The equivalence relation is one that partitions the set [(1–3 V)∪(4–6 V)] into the two corresponding ranges. In that sense, it is a trivial matter that each is an equivalence class: they are so defined (on the basis of *G*'s systematic physical behaviour to be sure!). But which properties should 1.32 V and 2.31 V (and *every other* possible voltage value) satisfy to belong to the *same* equivalence class? One might further object that the partitioning method is *ex post facto*: first, identify the *logical function(s)* that the physical gate computes, and, then, fix the equivalence relation on that basis. How

---

[2] I am indebted to Marty Wolf and Philippos Papayannopoulos for a useful discussion about this worry.

[3] Suppose that '∼' is a binary equivalence relation on A. Reflexivity means that for all a∈A, a∼a. Symmetry means that for all a, b∈A, if a∼b, then b∼a. And transitivity means that for all a, b, c∈A, if a∼b and b∼c, then a∼c.

else can we determine which functional differences "are not relevant to the regimented input–output transformations of equivalence classes of physical states across the system" (Coelho Mollo, 2017, p. 3496)?

Coelho Mollo's reply would be that the equivalence relation is deteremined by the behaviour to be explained (2017, p. 3490). First, the system should be functionally decomposed in light of the behaviour to explained (e.g., producing a specific output when only two inputs are above a certain threshold). Second, in the case of *G*, for example, the functional component is identified as taking two inputs and behaving differentially when they are received (i.e., systematically producing a *specific* output for these *specific* ranges of inputs). This observed differential behaviour partitions the system into two equivalence classes, if they have two different functional roles. This discovery method, however, partitions the system's behaviour into equivalence classes only *ex post facto*: producing a specific output when only two inputs are above a certain threshold, for example, is simply another way of describing conjunction.

A second worry is that this strategy seems to divorce computational *individuation* from computational *explanation* (Shagrir, 2020, p. 4098). A cognitive capacity that is explained by means of specifying the mathematical or logical function that a mechanism computes typically qualifies as a *computational* explanation. (Coelho Mollo would describe it as a *mathematical model* explanation, though.) A discovery that the locust's visual neurone computes *multiplication* in order to trigger an escape response to a looming object at just the right time (Gabbiani et al., 2002) *is* explanatory of that capacity (under the relevant theoretical and empirical constraints). Coelho Mollo, however, claims that "logical individuation is at least one step above computational individuation" (2017, p. 3495). It follows, then, that computational individuation is separate from computational explanation, even when the latter is couched in purely formal terms.

Why is that problematic? Coelho Mollo need *not* deny, of course, that describing the locust's visual neurone as computing multiplication is *explanatory*. Rather, he would claim that computational individuation is what allows us to make sense of the mathematical model explanation in terms of the neurone's multiplication. By showing that some component performs specific functional state transformations, e.g., on the basis of specific dynamic coupling between the neuronal input region—soma and dendrites—and the spike-generating output region, we can explain how the neurone computes multiplication. Coelho Mollo, thus, has to reject the common claim that explanatory practices in the cognitive sciences are roughly aligned with Marr's tri-level explanatory hierarchy (e.g., Anderson, 2015; Blokpoel, 2018; Hardcastle & Hardcastle, 2015)—as others have recently suggested (e.g., Bickle, 2015; Love, 2015).

The computational indeterminacy described above arises at Marr's top-level, which specifies the problem solved by the system in terms of input–output relations. Identifying the relevant input–output relations is an important step in figuring out *why* the system does what it does and *how*. Input–output (I/O)

equivalence, however, does not entail functional equivalence. Any two I/O equivalent algorithms may go through different sequences of states intermediate between these inputs and outputs. Hence, it is "pertinent to [further] inquire as to which state(s) the system occupies in the process of producing its output(s) for some given input(s)" (Buller, 1993, p. 158). And, indeed, scientists have to figure out which algorithm is likely used to compute that function and propose a plausible biophysical model that supports their hypothesis (Jones & Gabbiani, 2012). How does Coelho Mollo's individuative strategy fit these computational practices?[4]

In sum, the present individuative strategy fares better than Dewhurst in preserving a narrow version of computational equivalence by foregoing some implementational details. It remains unclear, however, how a *functional* equivalence relation may be further regimented in an analogous manner to its *logical* counterpart. Moreover, adopting this strategy, arguably, comes at the cost of being at odds with some explanatory practices in computational cognitive science.

3 Semantic individuation of computation

We now turn to the semantic individuative strategy, according to which representation is necessary for computation. A strong version of this individuative strategy may require that *only* semantic properties figure in the computational individuation of a state (or process). However, both Sprevak (2010), and Shagrir (2001, 2020) advance a weaker, and thus more plausible, individuative strategy, according to which a computational state (or process) is partially individuated by semantic properties and partially by non-semantic properties. Thus, Shagrir, for example, accepts that the relation of implementing an automaton by a physical system need not be individuated semantically, but claims that computational individuation proper does require semantic individuation (2020, p. 4088). Accordingly, he argues that G (and H for that matter) simultaneously implement both AND and OR automata. To determine the computational identity of such systems, on the present view, semantic constraints are required.

Relatedly, Sprevak argues that I/O equivalence is a necessary, but not a sufficient, condition for computational identity (2010, p. 269). The "respective inputs and outputs […of our G and H gates] are different, […] so different as to not have any physical or functional properties in common" (2010, p. 268). Sprevak asks what their I/O equivalence may consist in. Coelho Mollo's response is 'similarity in equivalence classes and their respective degrees of freedom'. The former, however, claims that the I/O equivalence of G and H consists in their respective inputs and outputs *representing* the same thing. Nevertheless, the *type* of representational content is

---

[4] A possible answer is that this individuative strategy somehow specifies the system's algorithmic level. The "different functional profiles [of two computing systems] would [result in a difference] in their capacity to carry out logical and mathematical functions" (Coelho Mollo, 2017, n. 20) (3495, fn 20). Evaluating this answer, though, exceeds the scope of this article.

Table 3 Electrical gate *G*'s with arm movement that is triggered only when both inputs are within the high voltage range	Input-channel 1	Input-channel 2	Output-channel	Arm Movement
	1–3 V	1–3 V	1–3 V	✗
	1–3 V	4–6 V	1–3 V	✗
	1–3 V	1–3 V	1–3 V	✗
	4–6 V	4–6 V	4–6 V	√

unconstrained; it may be mathematical, proximal, distal, narrow, or wide. Even if the inputs and outputs of an AND-gate are labelled with the numeral '0' or '1', such syntactic labelling is *still* representational content. Similarly, "[n]o physical, structural, or functional property decides", so Sprevak claims (2010, p. 269), whether *G* computes conjunction or disjunction. It is supposedly a difference in representational content.

One of the main arguments in support of the semantic individuative strategy concerns the semantic individuation of tasks. Lee (2018) summarises it very nicely; let us call it the 'task individuation' argument.

P1: Computations feature in explanations of task performance.

P2: Tasks are individuated semantically.

C: Hence, computation requires semantic individuation.

P1 may seem uncontentious at first, and, thus, may be accepted by computational mechanists. A closer inspection reveals that some proponents of mechanistic explanations may reject P1 as irrelevant to the case in point (Miłkowski, personal communication). For tasks, as such, are not the phenomena to be explained, but are often only experimental *effects*. In addition, "the phenomena we typically call 'effects' are incidental to the primary *explananda* of psychology" (Cummins, 2000, p. 140), namely cognitive capacities (e.g., learning capacity, the capacity for depth vision, and planning capacity). Given that mechanists would argue that scientific explanation is *phenomenon*-based, and task performance is only secondary, P1 should be rejected.[5]

Computational semanticists, such as Sprevak and Shagrir, however, endorse both P1 and P2. There are at least two good reasons for endorsing P2. The first one has just been discussed: computational I/O equivalence between systems like *G* and *H* can be easily defended (thereby, indirectly, also supporting the multiple realisability principle). The second reason is fending off Putnam- and Searle-like triviality arguments according to which every (complex enough) physical system computes every Turing-computable function (and there are infinitely many such functions!).

As stated above, the computational mechanist denies any appeal to semantic properties for computational individuation. Unlike the narrow individuative strategies proposed by Dewhurst and Coelho Mollo's, Piccinini accepts that *contextual*

---

[5] Piccinini also adds that the task individuation argument does not go through, if we reject the assumption that explanatia and their explananda must be individuated by the same properties (2015, p. 40).

Table 4 An electrical gate similar to *G* with the original low voltage range divided in two, and three types of corresponding arm movement: no movement (between 0 and 45 degrees), medium movement (between 45 and 90 degrees), and high movement (greater than 90 degrees)

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

factors can play a role in determining the computational identity of a physical system. However, in response to the task argument, he claims that "a (non-semantic) functional individuation of computational states is sufficient to determine which task is being performed by a mechanism, and hence which computation is explanatory in a context" (Piccinini, 2015, p. 43).

Shagrir agrees that the computational indeterminacy exhibited by *G*, for example, can indeed be settled by appealing to non-semantic functional properties. Without exceeding the boundaries of the encompassing system that contains *G*, one might appeal, say, to arm movement, as described by Table 3. If, as the table shows, the connected arm *only* moves when both inputs are within the high voltage range, then *G* may be said to compute conjunction (rather than disjunction). No semantic property needs to be invoked to determine *G*'s computational identity in this case.

To show that Piccinini's individuative strategy cannot deal with more intricate computational indeterminacies, Shagrir proposes a simple, yet clever, modification of the voltage ranges on which *G* operates. The resulting gate (see Table 4 below) now has *three* voltage ranges instead of *two*; it is a tri-stable system. Suppose that the low voltage range is (1–2 V) and (2–3 V). (Thus, grouping them both together gives us [1–3 V], as before.) Shagrir argues that this construction enables us to individuate movement as either (a) high movement only (i.e., when both inputs are within the high voltage range) or (b) medium movement *plus* high movement (i.e., in all possible input combinations except for {[1–2 V], [1–2 V]}). If we adopt the first option, the gate computes conjunction, but if we adopt the second, the gate computes disjunction. How can Piccinini's individuative strategy (as well as Coelho Mollo's) decide which functional kinds are relevant in identifying the computation that is actually performed? This challenge has remained unaddressed.

The upshot of the task individuation argument and this last example is that at least in some explanatory contexts, the functional and semantic tasks are not co-extensive. The last example illustrates that "arm movements by themselves do not suffice to determine the units of the computation, and hence, the computation itself"

(Harbecke & Shagrir, 2019). Thus, sometimes even "the system plus its immediate causal environment are not […] sufficient for fixing the actual computations performed by the system" (ibid). (Thus, Piccinini's short-arm strategy is insufficient in such cases.) Hence, the computational semanticist concludes that the computational individuation of, at least some, physical systems requires semantic constraints.

4 A new strategy: long-arm functional individuation of computation

The discrepancy between the mechanistic and semantic individuative strategies runs, *to some extent*, parallel with the disagreement between internalism and externalism about mental content, respectively.[6] An internalist claims that any cognitive phenomenon *P* supervenes *exclusively* on properties and processes inside an individual's head—or, more generally, body. For the only causally relevant factors are internal neurophysiological properties. How else can mental content exert any causal influence besides that which is manifested within the individual's neurophysiology? Thus, the internalist concludes that mental content is *narrow* in the sense that it supervenes on internal neurophysiological properties alone.

An externalist, however, denies the thesis that *P* supervenes only on properties and processes in the head/body. One reason is that the mere distinction *between* the space inside the body and outside it seems ad hoc for the purpose of explaining the nature of many cognitive phenomena. The more traditional reason is provided by Putnam's Twin Earth (1975) and Burge's Arthritis (1986) thought experiments. Although my Doppelganger on Twin Earth is (supposedly) identical to me, when he thinks about water, his thought refers to XYZ, whereas my *identical* thought refers to H_2O on earth. Therefore, in such circumstances content must causally depend on the environment, thereby resulting in *wide* content. This completes our brief detour into the internalist-externalist debate.

The next subsection discusses Piccinini's short-arm strategy and Shagrir's semantic strategy, between which the long-arm individuative strategy is positioned. Subsection 4.2 defends the plausibility of the long-arm strategy. Subsection 4.3 distinguishes this newly proposed strategy from its short-arm counterpart. In Subsection 4.4, it is argued that the long-arm individuative strategy does not entail an appeal to semantic properties. Finally, Subsection 4.5 revisits Shagrir's indeterminacy challenge that is based on the tri-stable gate described in Table 4.

4.1 Wide, short-arm versus semantic individuation of computation.

In this subsection, we focus on Piccinini's and Shagrir's individuative strategies as representative of the two camps in the debate. One further clarification is needed, however, about the relation between the two axes of computational individuation before proceeding. According to Dewhurst's strategy, computational individuation is

---

[6] Nevertheless, despite their disagreement about computational individuation, computational semanticists and mechanists alike may be externalists about content.

narrow,[7] since it supervenes *only* on the internal states and/or properties of the computing system itself. Piccinini's strategy, on the other hand, is short-arm, yet wide. Therefore, short-arm functional individuation need not be narrow: computational individuation along these two axes may differ.

Let us, then, clarify the distinction between narrow versus wide individuation and short-arm versus long-arm individuation. Computational individuation is *narrow*, if it supervenes only on the internal states and/or properties of the computing system concerned; otherwise, it is *wide*. Computational individuation is *short-arm*, if the computational inputs and outputs have to be realised within the system itself; it is *long-arm* insofar as the inputs or outputs can also be realised *outside* the system itself.[8] Whilst every long-arm functional strategy is, hence, also wide, not every wide individuative strategy is long-arm. According to Piccinini's strategy, for example, the environment can *affect* which computation is realised, yet the inputs and outputs are realised within the computing system concerned. The long-arm individuative strategy developed henceforth is wide.

With this distinction in mind, we can now examine Piccinini's short-arm, wide individuative strategy. He claims that understanding the nature of wide individuation requires an epistemic distinction between functionally relevant and irrelevant properties of the physical computing system (Piccinini, 2015, pp. 139–140). Drawing this distinction, on his view, requires knowledge of (a) which of the system's properties are relevant to its computational inputs and outputs, and (b) how they are relevant to the computational explanandum. This knowledge, in turn, requires an understanding of the way(s) that the system interacts—via inputs and outputs—with the context in which it is embedded.[9] This may lead us to conclude that Piccinini, in fact, concedes that, at least in some cases, wide computational individuation amounts to *wide* (semantic) content. If that were so, this mechanistic individuative strategy would supposedly collapse into its semantic counterpart.

Piccinini, unsurprisingly, denies this consequence on the basis of two reasons (2015, pp. 140–141). First, he argues that the relevant functional properties are not *very* wide: they concern the interaction between the system and its immediate context via the system's input and output transducers. The resulting strategy is, thus, short-arm. In artificial computing systems, the boundaries may be drawn at the forces exerted on input devices, such as trackpads, and the outputs produced by output devices, such as the screen monitor or printer. In biological systems, the wideness of the relevant properties required for computational individuation "does not even reach into the organisms' environment; it only reaches sensory receptors and muscle fibers" (ibid).

---

[7] Another clear example of an internalist computational view is Tucker's. He argues that the environment, broadly construed, *cannot* affect the computational individuation of the system concerned (Tucker, 2018).

[8] I am grateful to an anonymous reviewer for insisting on teasing apart these two conceptual distinctions.

[9] Coelho Mollo would similarly argue that the functional decomposition of the computing system depends on a capacity of interest, and this capacity may often be determined in part by the context in which it is embedded.

The second reason provided is that the mechanistically relevant properties are to be identified under the empirical constraints set by the suitable natural science or engineering methods. The computational identity of the system concerned may be discovered and individuated without appealing to any semantic properties. We are, thus, left only with short-arm factors for computational individuation.

However, a computational mechanist need not limit herself only to short-arm factors. Understanding how mechanisms, including computational ones, actually function often requires to situate them in their operational context. Bechtel's citation nicely captures this idea.

> "The behavior of mechanisms is highly dependent on conditions in their environments, including any regularities that occur there. But these are not discovered by looking inside the mechanism to the parts and operations or how these are organized. They must be discovered by examining the environment in which the mechanism operates and employing tools appropriate for such inquiry" (Bechtel, 2009, p. 559)

On the other side of the spectrum, Shagrir advances a full-blown externalist individuative strategy. He claims that (a) whilst wide, short-arm individuation can indeed eliminate some computational indeterminacies, others remain (see Table 4 and the discussion in Sect. 3), and, hence, (b) one possible route is "going even more external, to the outside environment" (Shagrir, 2020, p. 4102). He rightly requires, however, that a functional strategy that extends all the way into the environment and resolves all cases of indeterminacy should be shown to be (a) plausible, and (b) preferable to a semantic individuation of computation. In the next subsection, we modestly take up only the first requirement[10]; the second must await another opportunity.

4.2 Long-arm individuative strategy

To defend the plausibility of a long-arm, functional individuative strategy, consider a toy, but realistic, example of a shared physical subsystem S in rodents. Suppose that S receives two inputs: one from the hypothalamus in the form of orexin hormone (which is involved in the sleep and wake cycle in addition to energy balance), and another from the visual system. Orexin signals hunger when the organism's blood glucose levels are low, prompting the organism to search for food. The visual input, specifying the contours and textures of a visual object, signals the presence of an object, which *is likely to* be edible, within a visible distance from the organism. S produces a single output signal that is sent to the motor cortex, and when this output signal exceeds a certain threshold, it functions as a seek-food command. Such a computational description of S is mechanistic. But absent further constraints S might still be indeterminate between an AND- and an OR-characterisation.

---

[10] To be fair, even showing that this strategy addresses *all* cases of computational indeterminacies is a task that exceeds the scope of this article.

Why is that? Depending on the specific organism and its interaction with the environment, it may be the case that only when both inputs are "positive" (i.e., the orexin input exceeds a certain threshold and so does the visual input), S sends a seek-food command as a "positive" output. But if one of the inputs (or both) is "negative" (i.e., the relevant input threshold is not reached), then the output signal does not exceed the relevant threshold, and therefore, no seek-food command is sent. This description is consistent with conjunction. However, it is likewise plausible that organisms of another related species would seek food even when only one input is "positive". That is, if either blood glucose levels are low (thereby secreting orexin above a certain threshold) or a target object is identified within sight, the organism may forage for food (in response to the seek-food motor output of S). In this scenario, the computational description of S is consistent with (inclusive) disjunction.[11] Roughly the same S can be used to compute two different mathematical functions (more on the 'roughly' qualification in footnote 13).

The specific organism-environment interaction can, and often does, play a role in fixing the computational identity of S. In this sense, computation is not equated with the standard "offline" view of Turing machines (Wells, 1998),[12] but rather is understood ecologically—as inclusive of both S and its surrounding environment. The computational identity of S can be fixed by the biological function of S (cf. Coelho Mollo, 2019). Nevertheless, the contextual factors that are relevant to determining its biological function may extend beyond the organism's sensory receptors and muscle fibers—as Piccinini suggests. For that reason, "wide mechanistic explanations can [and should] be used by all researchers interested in the interaction of cognitive systems with their environments" (Miłkowski et al., 2018).

Suppose that S is a subsystem in the hopping mouse. The foraging behaviour of any species depends on the location and consumption of available resources, securing and storing these resources, existing competition with conspecifics and other species, and the risk of predation. It is quite plausible that, on average, a positive energy budget by the hopping mouse is expected only when both inputs to S are "positive". S may likely perform conjunction in the hopping mouse.

Another rodent, however, such as the golden hamster, might exhibit a different behaviour, if it were equipped with a similar subsystem.[13] Why? Because the amount of food hoarded by this organism increases significantly when food becomes

---

[11] An *exclusive* disjunction (XOR) interpretation under these circumstances is implausible, as it entails that when the organism is hungry and sees food-like object, it does not reach out to grab that food-like object.

[12] Wells advances an alternative conception of Turing machines that is arguably more consistent with theories of cognitive architecture. On this conception, "[the] memory/tape mechanism [of the Turing machine], is hypothesized to exist in the external environment. Consequently, cognitive computation is a process of organism-environment interaction" (Wells, 1998, p. 271).

[13] Despite possible minute differences between S in the golden hamster and in the hopping mouse, what matters here are the input–output relations and the connectivity between S and the relevant upstream/downstream subsystems. Thus, even if to qualify as a "positive" input to S, the orexin threshold is slightly higher, say, in the mouse (as compared to the hamster), this difference is not functionally important. Such differences may manifest even between different mice of the *same species*. For similar reasons, we do not doubt that the hypothalamus as a neuroendocrine organ exists in both the hamster and the mouse despite any physical differences between them.

available after being in short supply. Nevertheless, the amount of food this hamster typically consumes remains unchanged from pre-fast levels (Buckley et al., 2007). Such behaviour does nothing to decrease the hamster's appetite, and will likely continue until the food stored in its cheek pouches is actually chewed and swallowed, thereby resulting in an increased blood glucose level, and a decrease in the secretion of orexin. The hamster will, hence, forage for food when the relevant visual input to S is "positive", even if its low blood glucose level does not result in the secretion of orexin above the required threshold. Similarly, the hamster may also forage for food when its blood glucose level drops (and orexin is secreted above the relevant threshold as a "positive" input to S) without receiving the relevant visual input.[14] If so, S in the golden hamster may likely compute disjunction.

The astute reader may reasonably object at this point and claim that this toy example can be explained by a short-arm mechanistic strategy; this, however, is not necessarily so. The main reason for that is that the visual inputs to S in both the hopping mouse and the golden hamster in response to seeds should, on average, be produced by seeds, and not by any light reflected from seed-like objects. This is part of the standard consumer teleosemantic story (Millikan, 1993). That is, the subsystem S has the adapted proper function of searching for food in that environment. Given the particular environmental conditions (internal: glucose blood level, and external: availability of seeds), S has the adapted proper function of producing a seek-food signal. It also has the derived proper function of enabling the mouse and the hamster to survive in their environment by reaching out to the observed seed-like object or even by seeking yet-unseen seeds (in the hamster).[15]

The take-home message is that S has two different proper functions depending on how the computation of S is affected by and contributes to the organism-environment interaction. In the hopping mouse, S has the function of triggering a seek-food motor command if and only if both inputs are "positive" (i.e., to compute conjunction on the inputs). In the golden hamster, however, S has the function of triggering a seek-food motor command if at least one input is "positive" (i.e., to compute disjunction on the inputs).[16] These evolutionary functions should have been performed often enough in the evolutionary past of the respective species to have been selected. The same S can perform two different computations depending on the relevant context. The proposed long-arm individuative strategy can tell the two apart.

---

[14] It is probably for that reason, that vets often recommend to make home-grown hamsters work hard for their meals and hide food pellets or seeds inside paper bags or cardboard tubes.

[15] A seed-like object with similar surface properties of a seed may be further discriminated by the rodent's main olfactory system, which influences its foraging behaviour and food preferences.

[16] The consumer/producer versions of teleosemantics diverge on how the content of a representation is fixed by either how the representation is consumed or produced. But, importantly, adopting Neander's producer teleosemantics (2017) may just push the proposed computational individuative strategy all the way to full-blown externalism. This is because Neander's 'response function'—a causal-functional notion—already combines function and information. Information is understood, though, fundamentally as causation (roughly, A carries information about B if A is caused by B). Thus, response functions may determine the content of sensory-perceptual representations all by themselves.

4.3 The long-arm and wide, short-arm strategies are similar, but different

Whilst the long-arm strategy is indeed similar in many respects to the wide, short-arm counterpart, they are importantly different. In many cases, the inputs to the computational process may be internal to the organism; in those cases, the long-arm strategy, too, need not look any further for computational individuation. However, in the rodents example, the kinds of objects that cause S to compute are part of the environment of the organism, and hence, the inputs are long-arm. Likewise, the kinds of outputs resulting from the computation of S may affect the environment, and hence, S' outputs may also be long-arm. The computational states of organisms sometimes reach out into the world of things and their adaptive value for the organism.[17] Computational inputs in such cases are "identified with events in the external environment" of S and not simply with "more internal [events] such as the stimulation of [the] organism's sensory organs" (Harman, 1990, p. 32).

Accordingly, the computational identity of S, in our case, extends to the visual objects to which the organism responds adaptively, rather than stopping at its physical boundaries. In principle, the same retinal input caused by a seed-like object may be caused by some kind of illusion or a malfunction in the visual system. The ecological approach to computation proposed here need not stop at the sensory receptors of the organism—as in Piccinini's short-arm strategy. The relevant interaction of the organism with its environment (e.g., with the edible object spotted) is part of the factors contributing to fixing the computational identity of S. That is because the organism and relevant aspects of the environment are co-designed and interconnected.

The long-arm functional strategy should be seen as more encompassing than the wide, short-arm counterpart. When short-arm factors suffice for determining the computational identity of the system, we need look no further than these factors. On Piccinini's strategy, S performs both inclusive disjunction and conjunction, yet its teleological function is to compute only *one* of them. Its teleological function can supposedly be characterised through the rodent's immediate context (Piccinini, 2015, p. 43). However, it is unclear *how immediate* the 'immediate context' should be. If the hamster travels to great distances in search of food, does it count as an immediate context? Moreover, according to Piccinini, a teleological function is determined by the stable contribution by a trait of organisms—belonging to a biological population—to an objective goal of those organisms (2015, p. 108). This non-selectionist notion raises some complications that exceed the scope

---

[17] The implication of the '*sometimes*' qualification with respect to the long-arm individuative strategy is a pluralistic approach to computational individuation along the lines advocated by Lee (2018) as discussed elsewhere (Fresco, forthcoming).

of this article. On the long-arm strategy, computational individuation *may* include wide contextual factors: *S* may turn out to compute distinct (mathematical) functions depending on the *wider* context in which it is embedded.

4.4 The long-arm strategy need not be semantic

A consequence of adopting a long-arm, functional individuative strategy, however, need not amount to a full-blown externalist strategy.[18] Why? Whether a physical state *p* having some proper function suffices for *p* to also represent some feature, event or object in the environment depends on the relevant theory of representation that one adopts. If, for example, one is very liberal about what it takes for physical states of an organism to represent (Millikan, 1989),[19] then the proposed long-arm, functional individuative strategy entails that system *S* in the rodents does not only have the proper function of yielding seek-food commands, but it also represents the presence of food (or some such). On Artiga's proposal (2021), for instance, *S* may count as a genuinely representational detector, since appealing to the representational content of the states of *S* contributes to the explanation of the rodents' behaviour under the relevant circumstances.[20]

However, if one adopts a more restrictive view of representation, then the long-arm, functional individuative strategy need not amount to a full-blown externalist strategy. For one thing, on Lloyd's account of representation (1989), for *p* to qualify as a representational vehicle it need not only yield some behavioural output (e.g., a seek-food command in our rodents). *p* also has to depend (in terms of conditional probability), via multiple channels, on the *simultaneous conjunction* of *multiple* events (e.g., receptor events) all responding to the same, single environmental condition *C* (e.g., the presence of the observed seed-like object). Thus, the seek-food command event is insufficient on its own to qualify as a representation of *C*: another simultaneous receptor event in the rodent must be triggered—via another channel—in response to *C*. (And, even then, it will be the conjunction of these events that may qualify as a representation of *C*.)

On Sterelny's account of representation (1995), the physical states of *S* need not qualify as representations either, because simple control systems need not amount to being representational of the very events or features that they control. What is distinctive about *representational* states is the use of *diverse* stimuli for tracking the *same* distal stimulus. Thus, a distal feature or object that is only detected through a single cue does not qualify as a proper representation. "[T]here [has] to be sufficient

---

[18] Dewhurst indeed raises this objection against Piccinini's short-arm mechanistic strategy. Since it is not clear, so he claims, how Piccinini's strategy avoids the risk of being equated with a semantic theory of computation. For "once we have teleological functions we are not far from having a full-blown teleosemantic theory of representation" (Dewhurst, 2016, p. 796). Coelho Mollo's individuative strategy—discussed in Subsection 2.2—similarly appeals to teleological function, but denies even narrow content, such as logical properties.

[19] For Millikan, a representation simply requires that the organism (or a consumer subsystem) can fulfil its task normally when the producer (such as *S* in the case of our rodents) goes into a state that correlates with a given environmental condition (e.g., the existence of seeds in the proximal environment).

[20] Artiga, in fact, also requires that such a representational description of *S* accommodate misrepresentation and guide research on the rodents' behaviour in useful ways.

Table 5 An adaptation of Shagrir's tri-stable gate described in Table 4 with three types of corresponding movement pattern: no movement, medium movement, and high movement. The first input is received in the form of orexin hormone from the monkey's hypothalamus. The second input is received from the monkey's visual system. The output is a specific motor command

Input-channel 1	Input-channel 2	Movement pattern
Low Orexin Level	Few Leaves and No Fruits	None
Low Orexin Level	Rich in Leaves and Few Fruits	Medium
Intermediate Orexin Level	Few Leaves and No Fruits	Medium
Low Orexin Level	Few Leaves and Many Fruits	Medium
Normal Orexin Level	Few Leaves and No Fruits	Medium
Intermediate Orexin Level	Rich in Leaves and Few Fruits	Medium
Intermediate Orexin Level	Few Leaves and Many Fruits	Medium
Normal Orexin Level	Rich in Leaves and Few Fruits	Medium
Normal Orexin Level	Few Leaves and Many Fruits	High

variety in proximal routes, and sufficient stability of distal sources, for the organism's adaptive reaction to the environmental feature to be robust" (Sterelny, 1995, pp. 261–262). Thus, the states of *S* tracking features of the relevant stimulus (i.e., the seed-like object) need not be representational, because *S* detects the stimulus through a single cue only. It is very plausible, of course, that other subsystems in the rodent are used to track that stimulus via another kind of cue (e.g., specific odour molecules). That extra requirement, however, is not part of the long-arm, functional individuative strategy.

Finally, drawing on Burge's account of representation (2010), Schulte (2015) argues that constancy mechanisms are needed in addition to the function of a given system to track some environmental feature. A perceptual constancy for a property *F* is, roughly, the capacity to generate stable *F*-representations despite possible variations in the sensory stimulation. Perceptual constancies mark the distinction between mere sensitivity and perception, and, hence, also the contrast between proximal stimulus detection and representational capacities. "[T]heir presence in a sensory system is necessary and sufficient for the system's being a perceptual system" (Burge, 2010, p. 413). Although the rodents plausibly have constancy mechanisms that enable them to perceive properties of seed-like objects as constant—despite changes in distance, illumination, or viewing angles, such mechanisms are not part of the computational identity of *S* according to the long-arm, functional strategy.

The long-arm functional ingredient does bring this computational individuative strategy closer to the semantic one, but it certainly does not, and need not, go quite all the way there.

4.5 Shagrir's indeterminacy challenge revisited

Before concluding this section, we offer a partial reply to Shagrir's indeterminacy challenge, which is raised in the form of the tri-stable gate, discussed in Sect. 3. His challenge is based on a simple limbic movement output (i.e., none, medium,

or high), but we adapt it to tree-climbing movement pattens. Let us consider the foraging behaviour of the mantled howler monkey. Depending on its natural habitat (Hopkins, 2016), the daily diet of this monkey typically comprises mostly leaves (~54%), fruit (~40%), flowers (~5%), and petioles (~1%). The highest concentrations of mature fruits can normally be found at the tree's canopy, since that part receives plenty of sunlight and rain.

Table 5 describes three possible—yet highly simplified—tree-climbing behaviours of that monkey. Let us assume that low orexin levels—as the first input—signal normal glucose blood level, whereas normal orexin levels signal low glucose blood level, thereby suggesting that the monkey is very hungry. The second input—from the visual system—varies between *few* leaves and *no* fruits (i.e., a very poor nutrition source) and *few* leaves and *many* fruits (i.e., a highly rich nutrition source). Nutrition-wise, the canopy contains many fruits, but requires greater energetic expenditure (the monkey has to climb higher up the tree). In between, the tree layers are rich in leaves and contain few scattered fruits. This example should clearly illustrate that the long-arm functional strategy caters for cases in which the inputs (see input channel 2 in Table 5) and/or outputs are not physically realised in the organism, but rather in its environment. Let us now connect the dots to see how Shagrir's challenge might be met.

A long-arm functional individuation of this tri-stable system plausibly lends itself to an OR description, since both climbing patterns (i.e., medium and high) qualify as movement. It is only when the monkey's glucose blood levels are normal (i.e., it is not hungry) and the tree has few leaves and no fruits, will the monkey avoid climbing the tree. Otherwise, it certainly makes little sense according to optimal foraging theory. The monkey has to exert energy, for very little reward (the nutrition value of *that* tree is very low), when it is not even hungry. On the other hand, when its orexin levels are normal—that is, it is very hungry—climbing to the canopy carries a big reward in the form of nutrient-rich fruits. According to Shagrir's suggestion, then, the underlying biological system computes disjunction.

Nevertheless, the reason this analysis may be deemed only a partial reply is that one combination of inputs—namely, low orexin level *plus* leaves and many fruits—is less plausible. For that means, according to our construction, that although the monkey is not hungry, it will perform a *medium* climb to consume many fruits. One possibility is that the monkey must climb to the canopy, which means a *high*, rather than *medium*, movement pattern. Moreover, the monkey does so whilst exerting excessive energy when it is not even hungry. Another possibility is that the tree might be rich in fruits at a lower layer; that would make the relevant environment such that a medium movement pattern *is* plausible. Either way, whilst Table 4 covers the space of all logical possibilities with respect to Shagrir's original construction, it is an *idealisation*. Any given *ecological* setting, however, will seldom fit so neatly with *all* logical possibilities. The monkey's foraging behaviour would depend on various factors besides just the two specified inputs (including its group behaviour, the risk of predation, past learning history, and other hormones flowing in its body). Besides, it is probabilistic, rather than deterministic, and so it might be the case that the less plausible combination of inputs above is simply assigned a very low probability.

The long-arm functional individuative strategy proposed above has to be further regimented, and include an explanation of how it also applies to artificial computing

systems. At the very least, though, the present analysis renders this strategy plausible and a possible competitor to the wide, short-arm mechanistic strategy, on the one hand, and the full-blown externalist strategy, on the other hand. In the next section, the role of teleology in the long-arm strategy is further elaborated.

5 The long-arm functional strategy is (not) question-begging

An important objection concerns the explanatory work performed by the teleofunctional description of the computing system concerned (*S*). According to the long-arm functional strategy, the teleological function determines whether the computational identity of *S* is conjunction or disjunction. However, so the objection continues, whether *S* computes either conjunction or disjunction is predefined by the teleofunctional description so as to fit the computed function. In other words, the teleofunctional description of *S* is *tailored* to *fit* the computational identity of *S*. If so, then the arguments above based on the hopping mouse, golden hamster, and mantled howler monkey beg the question.

It should be stated, right off the bat, that the examples above are indeed tailor-made to show the biological plausibility of selective pressures that lead to the exploitation of one mathematical function (e.g., conjunction) over another (e.g., disjunction). Resource availability (e.g., presence of sufficient food), biological factors (e.g., presence of predators) and environmental conditions may result in the same *type* of system being used in to compute conjunction in one context (e.g., in the hopping mouse) and disjunction in another (e.g., in the golden hamster). However, none of the considerations appealed to above are devoid of biological plausibility. These simple examples can be replaced with more elaborate scientific case studies (e.g., contrastive studies of OR-based negative regulatory and AND-based positive regulatory networks in bacteria and yeast, which respond to intra- and extracellular inputs, in Mittenthal and Zou (2011)). More expository work would be required to lead to the *same* moral.

The long-arm functional strategy differs from other mechanistic strategies in the explanatory role played by the teleological function in computational individuation. Notably, both Piccinini (2015, p. 121) and Coelho Mollo (2017, p. 3487) would view the basic teleological function of *S* as processing medium-independent vehicles (e.g., digits) according to a rule—which is sensitive only to the vehicles' degrees of freedom. This *function to compute* is the one that is relevant to the computational identity of *S*. Thus, in neither the mouse nor the hamster is the '*seek-food*' function of *S* constitutive of the conjunction (mouse) or disjunction (hamster) function that *S* computes. According to the long-arm functional strategy, in each of these rodents, *S* indeterminately computes both functions. The differential selective pressure exerted on the mouse and the hamster is responsible for one actually computing conjunction and the other disjunction.

To drive the message home, let us imagine that system *S* is surgically removed from the mouse and carefully transplanted in the hamster and vice versa. Does it follow that now the hamster will compute conjunction and the mouse disjunction despite the rodent's 'task requirements' not having changed? The answer depends

on the precise notion of (biological) function the individuative strategy appeals to. A standard etiological theory of function may be committed to that logical implication, since the function *S* has depends entirely on its *phylogenetic history*. A complete answer must await another opportunity.

Yet, a selectionist theory of function need not be confined to *natural* selection. A function may result from the effect that has contributed to the differential persistence of a trait in the organism (Garson & Papineau, 2019). The shaping of behaviour through trial-and-error learning and the corresponding change of the underlying neural structures could count as a function-bestowing selection process. This idea can be explained by appeal to selection in the Pricean sense (Price, 1995). Thus understood, selection encompasses not only natural selection, but also *sample* selection[21] that may contribute to adaptation via variation and selective retention without replication and multiplication (Fresco et al., 2018). A second alternative could be to understand function relative to a technical notion of *design* that refers to neither the design history nor the organism's phylogeny. In brief, a function may be defined as "the contribution of a part of a system with design to such a capacity of the system to which, possibly among other parts, type-fixed components contribute" (Krohs, 2009, p. 80). This alternative has the primary facie explanatory advantage of catering for human-designed computing systems, too.

The mathematical function that *S* would actually compute following the surgical exchange ultimately depends on the ontogeny of the individual rodents. Whether or not *S* in the mouse will compute disjunction depends on the new functional role assumed by *S* in response to the bodily and environmental conditions (and likewise for the hamster). It might even be the case that *S* will no longer compute either conjunction or disjunction. Why? Because the mouse, for example, lacks the hamster's cheek pouches, and is, hence, incapable of storing too much food in its mouth without digesting it. Accordingly, it might make less biological sense for the mouse to forage for food in the absence of suitable visual stimuli. If that were indeed the case, then the selected Boolean function in the mouse would be neither conjunction nor disjunction. In the hamster, on the other hand, *S* might assume the same functional role as before and thus compute disjunction. The end result can, in principle, be empirically tested. Nevertheless, the long-arm functional strategy is not committed to the computed functions being equally interchanged with *S* being exchanged between the mouse and the hamster.

6 Conclusion

This article examines two main approaches to computational individuation: mechanistic and semantic. Some computational mechanists advocate a very narrow individuative strategy that relies only on physical properties, whereas other appeal to contextual factors. Nevertheless, even mechanists—who accept that contextual

---

[21] This is a process of selecting a subset from a set based on some value criterion. The selection of radio stations with the turning of a dial is an example Price used to describe this kind of process.

factors may play a role in determining the computational identity of a physical system—insist that the functional individuation is rather narrow (i.e., it does not exceed the boundaries of the physical system). The computational semanticist agrees that narrow contextual factors may settle some cases of indeterminacy, but argues that some important cases remain unanswered. She, therefore, claims that, at least in those cases, one must appeal to semantic content for computational individuation. The long-arm, functional individuative strategy proposed above may address such open cases of computational indeterminacy without adopting full-blown externalism. As such, it opens up the possibility of an intermediate between these two opposing positions.

References

Anderson, B. L. (2015). Can Computational Goals Inform Theories of Vision? *Topics in Cognitive Science, 7*(2), 274–286. https://doi.org/10.1111/tops.12136

Artiga, M. (2021). Strong liberal representationalism. *Phenomenology and the Cognitive Sciences*. https://doi.org/10.1007/s11097-020-09720-z

Bechtel, W. (2009). Looking down, around, and up: Mechanistic explanation in psychology. *Philosophical Psychology, 22*(5), 543–564. https://doi.org/10.1080/09515080903238948

Bickle, J. (2015). Marr and Reductionism. *Topics in Cognitive Science, 7*(2), 299–311. https://doi.org/10.1111/tops.12134

Bishop, J. M. (2009). A cognitive computation fallacy? Cognition, computations and panpsychism. *Cognitive Computation, 1*(3), 221–233. https://doi.org/10.1007/s12559-009-9019-6

Blokpoel, M. (2018). Sculpting Computational-Level Models. *Topics in Cognitive Science, 10*(3), 641–648. https://doi.org/10.1111/tops.12282

Buckley, C. A., Schneider, J. E., & Cundall, D. (2007). Kinematic analysis of an appetitive food-handling behavior: The functional morphology of Syrian hamster cheek pouches. *Journal of Experimental Biology, 210*(17), 3096–3106. https://doi.org/10.1242/jeb.003210

Buller, D. J. (1993). Confirmation and the computational paradigm (or: Why do you think they call itartificial intelligence?). *Minds and Machines, 3*(2), 155–181. https://doi.org/10.1007/BF00975530

Burge, T. (1986). Individualism and Psychology. *The Philosophical Review, 95*(1), 3–45. https://doi.org/10.2307/2185131

Burge, T. (2010). Origins of Objectivity. *Oxford University Press*. https://doi.org/10.1093/acprof:oso/9780199581405.001.0001

Coelho Mollo, D. (2017). Functional individuation, mechanistic implementation: The proper way of seeing the mechanistic view of concrete computation. *Synthese*. https://doi.org/10.1007/s11229-017-1380-5

Coelho Mollo, D. (2019). Are There Teleological Functions to Compute? *Philosophy of Science, 86*(3), 431–452. https://doi.org/10.1086/703554

Cummins, R. (2000). How does it work?" versus" what are the laws?": Two conceptions of psychological explanation. In F. C. Keil & R. A. Wilson (Eds.), *Explanation and cognition* (pp. 117–144). MIT Press.

Dewhurst, J. (2016). Physical computation: A mechanistic account. *Philosophical Psychology, 29*(5), 795–797. https://doi.org/10.1080/09515089.2016.1150450

Dewhurst, J. (2018). Individuation without Representation. *The British Journal for the Philosophy of Science, 69*(1), 103–116. https://doi.org/10.1093/bjps/axw018

Fresco, N. (forthcoming). How context can determine the identity of physical computation. In M. Hemmo, S. Ioannidis, O. Shenker, and G. Vishne (Eds.), *Levels of reality in science and philosophy*.

Fresco, N., Ginsburg, S., & Jablonka, E. (2018). Functional Information: A Graded Taxonomy of Difference Makers. *Review of Philosophy and Psychology*. https://doi.org/10.1007/s13164-018-0410-7

Fresco, N., & Miłkowski, M. (2019). Mechanistic Computational Individuation without Biting the Bullet. *The British Journal for the Philosophy of Science*. https://doi.org/10.1093/bjps/axz005

Fresco, N., Copeland, B. J., & Wolf, M. J. (2021). The indeterminacy of computation. *Synthese*. https://doi.org/10.1007/s11229-021-03352-9

Gabbiani, F., Krapp, H. G., Koch, C., & Laurent, G. (2002). Multiplicative computation in a visual neuron sensitive to looming. *Nature, 420*(6913), 320–324. https://doi.org/10.1038/nature01190

Garson, J., & Papineau, D. (2019). Teleosemantics, selection and novel contents. *Biology & Philosophy*. https://doi.org/10.1007/s10539-019-9689-8

Harbecke, J., & Shagrir, O. (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

Hardcastle, V. G., & Hardcastle, K. (2015). Marr's Levels Revisited: Understanding How Brains Break. *Topics in Cognitive Science, 7*(2), 259–273. https://doi.org/10.1111/tops.12130

Harman, G. (1990). The Intrinsic Quality of Experience. *Philosophical Perspectives, 4*, 31–52. https://doi.org/10.2307/2214186

Hopkins, M. E. (2016). Mantled howler monkey spatial foraging decisions reflect spatial and temporal knowledge of resource distributions. *Animal Cognition, 19*(2), 387–403. https://doi.org/10.1007/s10071-015-0941-6

Illari, P. M., & Williamson, J. (2012). What is a mechanism? Thinking about mechanisms across the sciences. *European Journal for Philosophy of Science, 2*(1), 119–135. https://doi.org/10.1007/s13194-011-0038-2

Jones, P. W., & Gabbiani, F. (2012). Logarithmic Compression of Sensory Signals within the Dendritic Tree of a Collision-Sensitive Neuron. *Journal of Neuroscience, 32*(14), 4923–4934. https://doi.org/10.1523/JNEUROSCI.5777-11.2012

Krohs, U. (2009). Functions as based on a concept of general design. *Synthese, 166*(1), 69–89. https://doi.org/10.1007/s11229-007-9258-6

Lee, J. (2018). Mechanisms, Wide Functions, and Content: Towards a Computational Pluralism. *The British Journal for the Philosophy of Science*. https://doi.org/10.1093/bjps/axy061

Lloyd, D. E. (1989). *Simple minds*. The MIT Press.

Love, B. C. (2015). The Algorithmic Level Is the Bridge Between Computation and Brain. *Topics in Cognitive Science, 7*(2), 230–242. https://doi.org/10.1111/tops.12131

Miłkowski, M. (2016). Computation and Multiple Realizability. In V. C. Müller (Ed.), *Fundamental Issues of Artificial Intelligence* (pp. 29–41). Springer. https://doi.org/10.1007/978-3-319-26485-1_3

Miłkowski, M. (2017). The False Dichotomy Between Causal Realization and Semantic Computation. *Hybris, 38*, 1–21.

Miłkowski, M., Clowes, R., Rucińska, Z., Przegalińska, A., Zawidzki, T., Krueger, J., Gies, A., McGann, M., Afeltowicz, Ł, Wachowski, W., Stjernberg, F., Loughlin, V., & Hohol, M. (2018). From Wide Cognition to Mechanisms: A Silent Revolution. *Frontiers in Psychology, 9*, 2393. https://doi.org/10.3389/fpsyg.2018.02393

Millikan, R. G. (1989). Biosemantics. *The Journal of Philosophy, 86*(6), 281–297. https://doi.org/10.2307/2027123

Millikan, R. G. (1993). *White Queen psychology and other essays for Alice*. MIT Press.

Mittenthal, J. E., & Zou, L. (2011). To signal a conjunction of many inputs negative regulation is likely. *Mathematical Biosciences, 231*(1), 69–75. https://doi.org/10.1016/j.mbs.2011.02.001

Neander, K. (2017). *A mark of the mental: In defense of informational teleosemantics*. The MIT Press.

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

Price, G. R. (1995). The nature of selection. *Journal of Theoretical Biology, 175*(3), 389–396. https://doi.org/10.1006/jtbi.1995.0149

Putnam, H. (1975). The Meaning of "Meaning." In K. Gunderson (Ed.), *Language, mind, and knowledge* (pp. 131–193). University of Minnesota Press.

Schulte, P. (2015). Perceptual representations: A teleosemantic answer to the breadth-of-application problem. *Biology & Philosophy, 30*(1), 119–136. https://doi.org/10.1007/s10539-013-9390-2

Shagrir, O. (2001). Content, computation and externalism. *Mind, 110*(438), 369–400. https://doi.org/10.1093/mind/110.438.369

Shagrir, O. (2020). In defense of the semantic view of computation. *Synthese, 197*(9), 4083–4108. https://doi.org/10.1007/s11229-018-01921-z

Sprevak, M. (2010). Computation, individuation, and the received view on representation. *Studies in History and Philosophy of Science Part A, 41*(3), 260–270. https://doi.org/10.1016/j.shpsa.2010.07.008

Sterelny, K. (1995). Basic Minds. *Philosophical Perspectives, 9*, 251. https://doi.org/10.2307/2214221

Stinson, C. (2016). Mechanisms in psychology: Ripping nature at its seams. *Synthese, 193*(5), 1585–1614. https://doi.org/10.1007/s11229-015-0871-5

Tucker, C. (2018). How to Explain Miscomputation. *Philosophers'. Imprint, 18*(24), 1–17.

Wells, A. J. (1998). Turing's Analysis of Computation and Theories of Cognitive Architecture. *Cognitive Science, 22*(3), 269–294. https://doi.org/10.1207/s15516709cog2203_1