Cognitive Ontology in Terms of Cognitive Homology: The Role of Brain, Behavior, and Environment for Individuating Cognitive Categories

Beate Krickel, TU Berlin
Mariel Goddu, Stanford

How should scientists carve up cognition to generate good predictions, explanations, and models of cognition? This chapter argues that cognitive categories should be constructed the same way that biological categories are: in terms of *homology*. The chapter adapts a developmental account of trait identity from evolutionary-developmental biology to make sense of the notion of "cognitive homology." The consequence is that both brain structures and the organism's ongoing interactions with the environment are crucial for individuating cognitive homologies, and thus for cognitive ontology.

Keywords: cognitive ontology, homology, cognitive development, mechanisms

1 Introduction

The *cognitive ontology question* is how to carve up cognition so that the resulting carving (i.e., ontology, taxonomy, categorization) enables successful scientific predictions, explanations, and generalizations.[1] The issue with contemporary cognitive science, however, is that its categories are ill-defined and ambiguous (Poldrack et al., 2011; Zagaria et al., 2020) and not derived from rigorous scientific reasoning but from intuition, speculation, or "dreamed-up vocabulary" (Buzsáki, 2019, p. 6) (see also Vanderwolf, 2007). According to Buzsáki and Vanderwolf, William James introduced this vocabulary in 1890 and, indeed, it still broadly corresponds to the categories that guide cognitive scientific investigations today: "self-consciousness", "attention", "association", "perception", "memory", "imagination", "reasoning", "production of movement", "emotion" (James, 1890). Many theorists and researchers agree that these categories should be revised (e.g., Anderson, 2014; Cisek, 2019; De Brigard, 2017; Dewhurst, 2021; Figdor, 2018; Poldrack, 2010; Poldrack et al., 2011; Price & Friston, 2005). The challenge, then, is to find a new way (or ways) of categorizing cognitive capacities. This is what scientists and philosophers who engage with the cognitive ontology question aim to do.

There are two broad strategies for addressing this challenge (for a discussion of accounts corresponding to these two strategies see Francken et al., 2022; McCaffrey & Wright, 2022). *Incremental approaches* take existing cognitive categories and revise them. These approaches are hypothesis-driven and iterative: roughly, they hypothesize that some cognitive category X corresponds to a cognitive capacity, outline specific predictions that follow from this hypothesis, and conduct empirical research to test them. Then, they keep, revise, or reject X in light of the findings. The challenge for incremental approaches is that they are prone to inheriting the problems of the existing categories. *Radical approaches,* by contrast, reject the current categories altogether. They start from scratch, constructing new cognitive categories that might depart dramatically from the ones posited by contemporary cognitive science. Radical approaches face two main challenges: first, they must provide reasonable bases for deriving the novel categories, and second, they must justify why the resulting categories are indeed *cognitive* categories (rather than, for example, purely mathematical entities; see, e.g., Yeo et al. (2015)—and for a critical discussion along these lines see McCaffrey & Wright (2022, p. 446).

A key question in the cognitive ontology literature concerns the relationship between cognitive categories and neural categories (such as brain regions). One central issue is whether one should anticipate finding one-to-one correspondences between cognitive functions and neural structures, or whether more complex mappings, such as one-to-many, many-to-one, or many-to-many, might be all one can get (Viola, 2017). We will not address this issue here, as it is less relevant to the objectives of our paper. Another significant issue is whether or not the development of a new cognitive ontology should begin with the brain, or not. *Bottom-up approaches* to cognitive ontology seek to derive cognitive categories from neural categories, usually anatomically individuated brain regions. Radical approaches to cognitive ontology typically adopt this bottom-up perspective. In contrast, *top-down approaches* start with hypotheses about cognitive categories and then—if at all—examine how these cognitive categories are realized in the brain. Incremental approaches to cognitive ontology are usually top-down.

One feature that most approaches to cognitive ontology in the current literature share is that they tend to focus solely on *human* cognition (Figdor, 2022). This is a severe shortcoming. To the best of our scientific knowledge, cognition is a product of evolution (Cisek, 2019). An approach to cognitive ontology informed by basic considerations from evolutionary biology should assume that (some form of) cognition is found in many species.

In this chapter, we propose a new answer to the cognitive ontology question that takes the evolution of cognition into account. We suggest that cognitive carvings should be based on *cognitive homologies*, cognitive capacities that are "the same across species" by the standards of evolutionary biology. Still, our approach is incremental. We suggest investigating existing cognitive capacities to discover whether they are cognitive homologies. If presently distinct cognitive capacities are found to be homologous, revision is necessary: they collapse into one and the same capacity. If cognitive capacities that are currently considered to be the same turn out *not* to be homologous, revision is also necessary: they are different cognitive capacities. In that sense, our approach is top-down. Still, as we show, investigating the brain and neural mechanisms is crucial for determining cognitive homologies, and thus, for cognitive ontology.

More specifically, we aim to develop a novel account of cognitive homology by adopting a recent account of "trait identity" from evolutionary-developmental biology (DiFrisco, 2023; DiFrisco et al., 2020). This account supplements standard conceptions of homology by introducing "Character Identity Mechanisms" (ChIMs): developmental mechanisms that guarantee the stability of trait identity across phylogeny. Similarly, our account proposes to individuate cognitive capacities based on the developmental mechanisms that bring them about—what we will call *cognitive ChIMs*. As we will show, cognitive ChIMs are developmental mechanisms that crucially involve brain mechanisms as well as the organism's interactions with its environment.

Here, the goal is primarily to provide a proof of concept. We argue that a plausible notion of cognitive homology can be developed based on the notion of a *cognitive ChIM,* a concept which can then be used in answering the cognitive ontology question. Turning the idea into a full-fledged approach to cognitive ontology constitutes a novel interdisciplinary research program—of which this chapter can only be a starting point.

The paper proceeds as follows. Section 2 introduces the biological concept of homology and presents two conceptions of it: the *phylogenetic conception,* and the *ontogenetic conception* that will be crucial for our account of cognitive homology. In Section 3, we review two previous accounts of cognitive homology. We argue that they are unsuccessful mainly because they attempt to define a notion of structural similarity for entities that are characterized functionally. In Section 4, we argue that an account of cognitive homology that incorporates the ontogenetic conception of homology avoids the problems afflicting existing accounts. Section 5 concludes.

2 Homology in Biology

> Countless philosophy students have been taught that the definition of the heart as 'the organ whose purpose is to pump the blood' is a typical example of biological classification… [T]his picture "is utterly false…" …[I]n comparative anatomy the parts of organisms are primarily defined as *homologues*. (Griffiths, 2007, p. 644)

Contrary to a commonly held view among philosophers, biologists tend not to individuate biological entities—such as organs, tissues, and cells—in terms of what they do. Rather, they individuate them in terms of *homology*. Homologous traits (also called "homologues" or "homologies") share a kind of identity: "[w]hen two characters[2] or traits are homologous, they are fundamentally the same kind of character" (see also Fusco, 2022; DiFrisco, 2023, p. 1).

Homology is typically defined for morphological features (body parts) but need not be restricted to it (see Section 3). One example of a morphological homology is the 'pentadactyl forelimb', the general forelimb form of species that have five "fingers". Instances of this trait include the human arm, the cat's foreleg, the whale's flipper, and the bat's wing (see Figure 1).

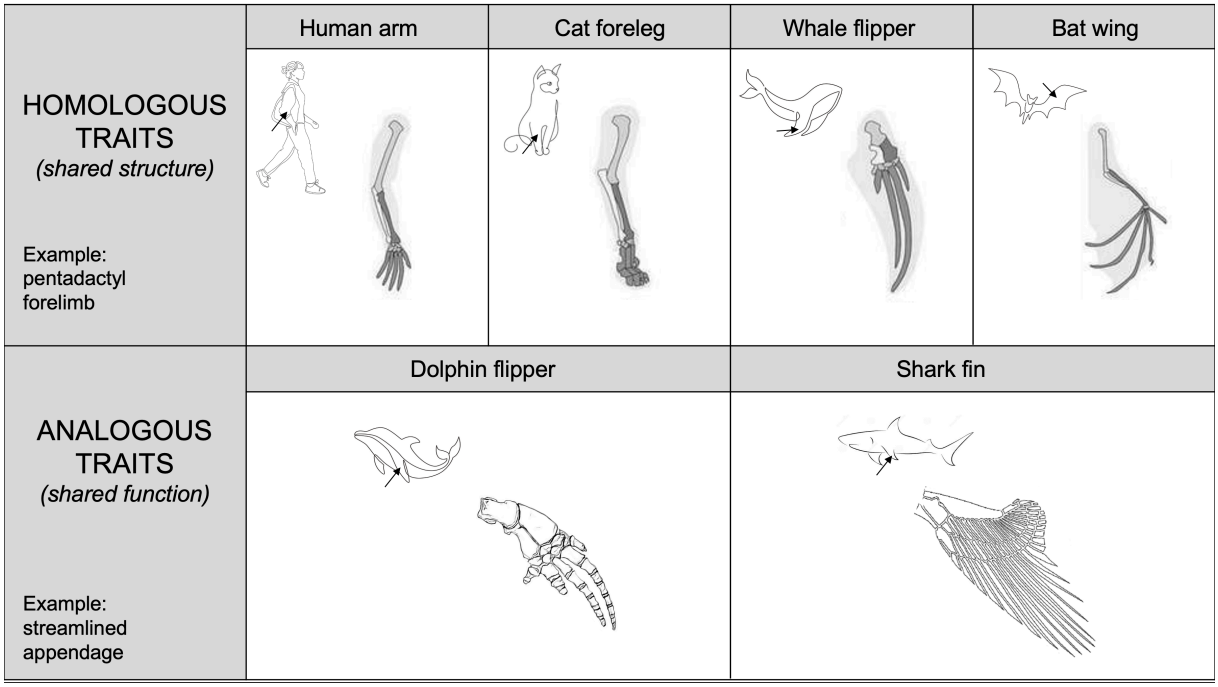

Figure 1. Examples of homologous versus analogous traits (own illustration).

These forms look different: they are different *phenotypes* or *phenotypical expressions* of the "pentadactyl forelimb" trait. They have different shapes and fulfill rather different functions—reaching, walking, swimming, and flying. However, the sameness relevant for homology of traits concerns their *structural* similarity rather than their outward appearance or function.

Structural similarity of morphological entities involves sameness of internal organization, positions relative to other structures, degree of complexity, and position within a series of trait transformations in evolution (García, 2010; Bergeron, 2021; DiFrisco, 2023).

The conceptual foil to homologous traits is *analogous* traits (also called "analogues" or "analogies"). In contrast to homologues, where the similarities are structural, analogues have merely functional similarities. Some species develop superficially similar traits because they adapted to similar environments. For example, the ancestors of dolphins and sharks faced adaptive pressures such that those individuals with streamlined appendages had a survival benefit. However, despite their common function (swimming) and apparent similarity of form (streamlined shape), the dolphin's flipper and the shark's fin are structurally *dis*similar (see Figure 1). These body parts are not instances of the same trait.

Biological sameness or trait-identity thus concerns *structural sameness* rather than *functional sameness*. Since the late 19th century (and already present in Darwin's work), a further crucial aspect of homology is *shared ancestry*. Species with homologous traits are descended from a common ancestor that had the trait in question and are linked to that ancestor by a continuous lineage of ancestors that all had that trait. Common descent and the continuous line of ancestors with the trait are assumed to be the cause and explanation for the structural similarities between them (Wagner, 1989; Griffiths, 2007). Call this concept of homology the *phylogenetic conception* (DiFrisco, 2023). The phylogenetic conception of homology is the standard view of homology.

Recently, the phylogenetic conception has become the target of various criticisms (for an overview see DiFrisco, 2023). One problem is that there is simply no continuity in the presence of a trait across generations: almost all traits need to develop anew in each individual. Another criticism is that the phylogenetic conception cannot address *serial homology,* the sameness of traits within the same individual organism. For example, the fore and hind legs of tetrapods are iterated versions of the same body part. Thus, they are homologous but do not have an obvious phylogenetic connection. A further problem for the phylogenetic conception of homology is that it does not provide a satisfactory explanation of what it means for traits to "maintain their identity" despite vast changes across evolution (Wagner, 1989; for details see DiFrisco, 2023). Finally, the phylogenetic conception is problematic for epistemic reasons. Shared ancestry and continuous inheritance are not directly accessible; rather, they must be inferred from fossil records or comparative anatomical or genetic analyses (DiFrisco, Love and Wagner, 2023).

A conception of homology that aims to address these problems may be called the *ontogenetic* or *developmental conception* of homology. This is seen as a supplement rather than an

alternative to the phylogenetic conception (see DiFrisco, 2023).[3] It is based on adopting the 'developmental homology principle':

> *Developmental homology principle (DHP)*: The identity and classification of morphological homologies is determined, at least in part, by the specific developmental factors that cause them in ontogeny. (DiFrisco, 2023, p. 4)

DHP provides a developmental basis for individuating and identifying homologous traits. Different accounts of ontogenetic homology make use of DHP in different ways. Here, we focus on an account due to DiFrisco, Wagner, and Love (2020) that centers on *character identity mechanisms* (ChIMs). The concept of a ChIM follows from two hypotheses:

> (1) there is a general, recognizable mechanistic architecture in development that explains the traceability of characters [traits], and
>
> (2) these mechanisms are themselves cohesive units that can be traced as homologues in evolution. (DiFrisco, Love and Wagner, 2020, p. 7)

Traceability of traits is a challenge for researchers: as outlined above, the expression of the same trait (homologue) in different species can look superficially different (compare a human arm and a bat wing). This gives rise to two questions. First: How can homologues be detected and identified? And second: In what sense can a homologue be said to maintain its identity (i.e., remain the "same trait"), even if it undergoes drastic changes?

According to DiFrisco et al., ChIMs provide answers to both questions because they are themselves traceable. They are highly evolutionarily conserved and resistant to change. This is because ChIMs have special characteristic features, as we will now explain.

ChIMs are a special type of biological mechanism. According to the so-called "new mechanists" (who DiFrisco et al. follow here), *mechanisms are entities and activities organized such that they bring about a phenomenon* (see e.g., Craver, 2007). Mechanisms are individuated based on their entities, activities, organizational features, and the phenomena they bring about.[4] For example, the mechanisms for neurotransmitter release consist of ions, ion channels, and vesicles (the entities) that open, close, diffuse, and move (the activities) at specific locations, times, and in specific orders (their organization) such that neurotransmitters are released into the synaptic cleft (the phenomenon). DiFrisco, Wagner, and Love (2020) provide a detailed discussion of ChIM individuation that is based on this general framework. For lack of space, we won't repeat these considerations here but refer the reader to their article.

ChIMs are *developmental* mechanisms. Many developmental mechanisms—such as those that govern the formation of particular feathers in birds or wing patterns in butterflies—"micro-manage" the expression of species-specific traits. That is, they are responsible for the

development of a particular *phenotypical expression of a trait* in a species. Changes to such mechanisms might cause slightly different feathers or wing patterns to develop. By contrast, ChIMs regulate *trait identity*. Changes to a ChIM might result in the total absence of feathers or wing patterns at all. One example of a ChIM is the mechanism involving *Ultrabithorax,* a gene that regulates hindwing development in insects. Evidence from "knockdown" experiments (in which *Ultrabithorax* is "switched off") show that this mechanism controls character identity rather than phenotype: the knockdown causes the development of a forewing in the hindwing location on the body without altering the size, shape or any other structural features of the developing structure (Wagner, 2007, p. 475; DiFrisco, Love and Wagner, 2020, p. 3). The gene *Ultrabithorax* thus determines whether a wing *is* a hindwing (trait identity) rather than, for example, the specific color, shape, or structure of the hindwing (trait expression).

ChIMs have the power to determine trait identity because of their "mechanistic architecture" (DiFrisco, Love and Wagner, 2020, p. 2). They are characterized by their location within a causal topology (the "shape" of a causal process). Specifically, ChIMs are the knot in a bowtie-like causal structure (Figure). They occupy this position because they are causally necessary and non-redundant. ChIMs are causally necessary because there is no other way to turn the 'inputs' from the left bowtie wing (the 'upstream' developmental processes) into the 'outputs' on the right (the 'downstream' developmental processes) other than via the ChIM. They are causally non-redundant because they are not replaceable by any other causal mechanism: there is no back-up mechanism that can fill the position of the ChIM in case the ChIM fails. This is due to the complexity of ChIMs that is necessary to transform the ChIM's inputs into its outputs.

It is evolutionarily very unlikely that an organism develops a 'better-safe-than-sorry' system that instantiates such complexity twice.

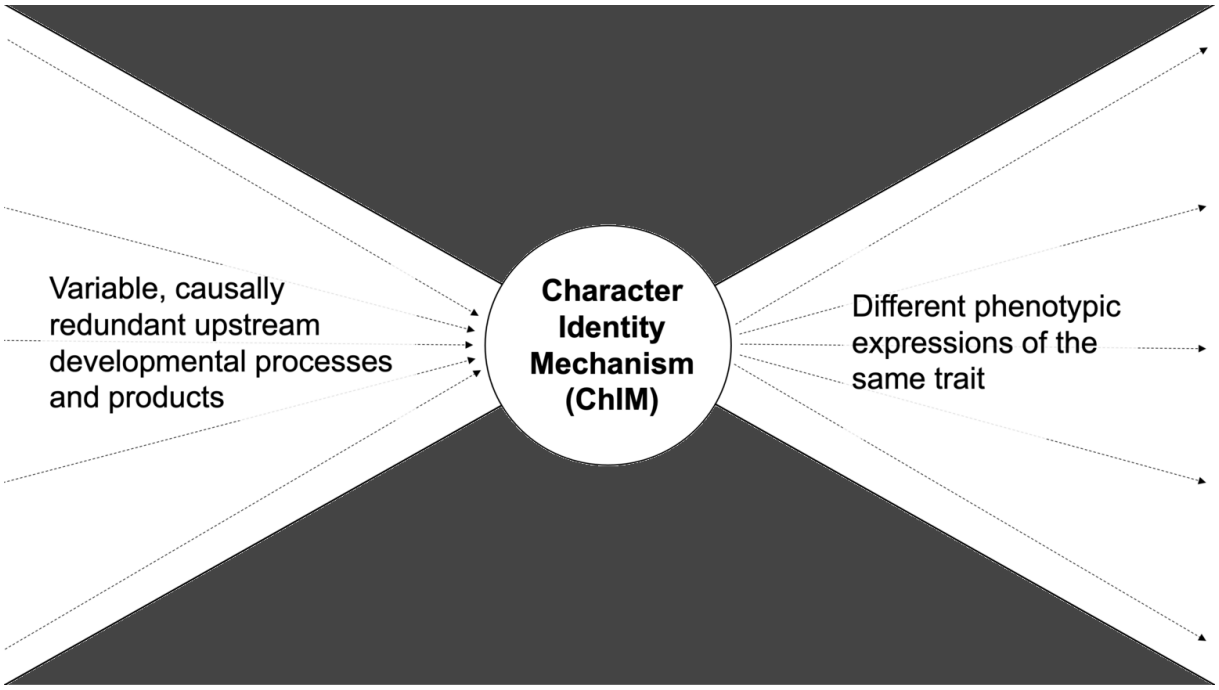

Figure 2. A Character Identity Mechanism (ChIM) is a developmental mechanism with a specific position in a causal process: it is the knot of a "bowtie" causal topology (in white on grey background). ChIMs are complex mechanisms that take a variety of variable, causally redundant upstream processes as inputs and transform them into a single output: the specific character that the ChIM produces. This output then typically undergoes subsequent developments. Due to their causal necessity and non-redundancy, ChIMs are highly conserved across evolution.

Due to these characteristic features, ChIMs can be seen as causing and explaining trait identity throughout the phylogenetic tree. Because ChIMs' necessity and complexity renders them less replaceable than other mechanisms, their "developmental outcomes… are subject to strong stabilizing selection." The result is that ChIMs "are more likely to be evolutionarily conserved than other developmental mechanisms" (DiFrisco, Love and Wagner, 2020, pp. 8–9). Individuals with a defective ChIM do not develop whatever trait the ChIM is causally responsible for, and are thus likely to die early and not reproduce. All living individuals who developed the trait also had the ChIM, which makes it very likely that they pass on the ChIM to their offspring. Figure presents the ChIM for the vertebrate paired limbs trait as an example.

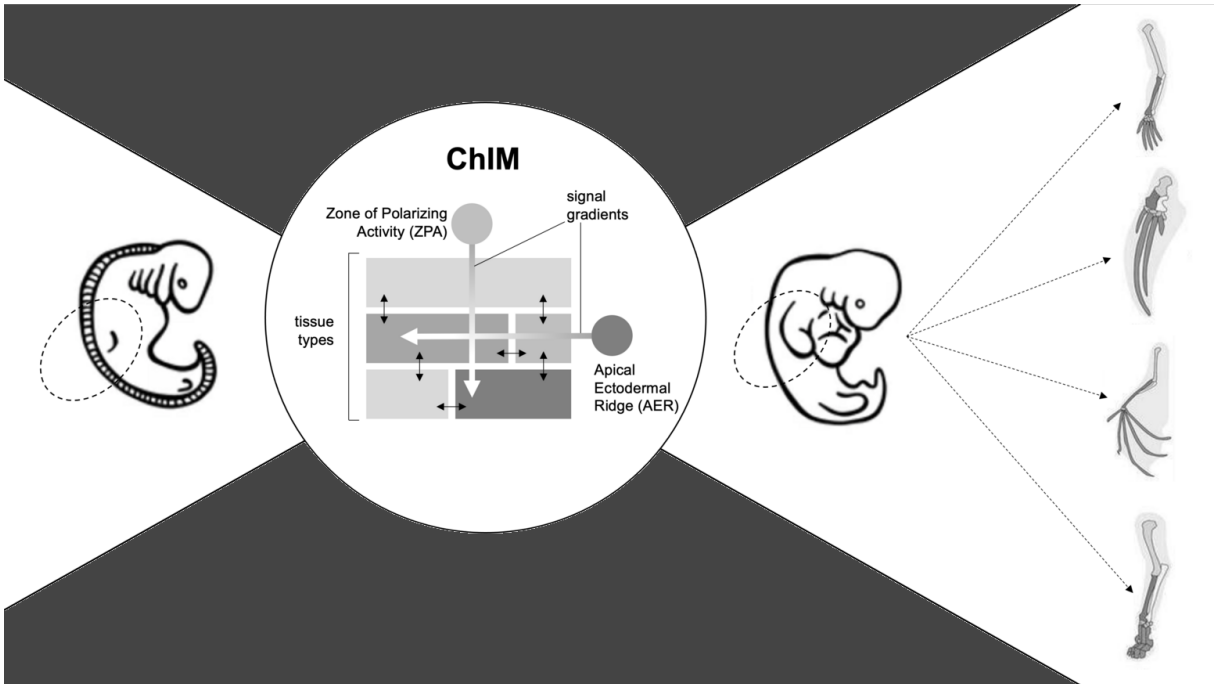

Figure 3 Schematic illustration (adapted from DiFrisco, Love and Wagner, 2020) of the "paired gnathostome appendage ChIM." The development of limb buds in vertebrate embryos precedes the development of paired appendages (e.g., fins and limbs). Bi-directional arrows schematically depict local, hormone-driven interactions between tissue types; limb development is orchestrated by four signaling centers throughout the limb bud. The two most well-known signaling centers are included here as circles with arrows indicating signal gradients. The ZPA establishes anterior-posterior organization of the limb, and the AER establishes proximal-distal organization. These signaling centers form a complex, causally non-redundant, and necessary mechanism that jointly determines limb identity. Their activities are interdependent and mutually reinforce each other (DiFrisco, Love and Wagner, 2020). Upstream of the ChIM (its causal 'input') is relatively undifferentiated, loosely organized mesodermal embryonic tissue; downstream of the ChIM (its causal 'output') is a structure with the identifiable tissue types and organization of a limb. This basic limb structure is then modified by further processes to determine a broad spectrum of possible limb phenotypes, such as the range of pentadactyl forelimb phenotypes discussed above.

The vertebrate limb ChIM has the hallmark features described by DiFrisco et al.: it is a developmental mechanism that is causally necessary and nonredundant for the transformation of 'upstream' developmental inputs (here, pluripotent mesenchymal cells in early embryos) into 'downstream' developmental outputs (here, various possible phenotypical expressions of a limb in different species). The ChIM's central role in generating the limb trait explains the structural similarities across species. By contrast, analogous traits—such as the streamlined appendages of sharks and dolphins mentioned above—are generated by different mechanisms and are, thus, not of the same type.

To summarize: biological traits are carved up in terms of homology. The *phylogenetic conception of homology* holds that homologies share structural similarities that can be traced back to a common ancestor. The *ontogenetic conception of homology* supplements this

approach to solve several problems afflicting a purely phylogenetic conception of homology. A recent ontogenetic account introduces the concept of Character Identity Mechanisms (ChIMs): causally necessary, non-redundant developmental mechanisms that are responsible for trait identity. ChIMs are conserved across evolution. On this account, two traits are the same in different species if only if they are produced by the same ChIM.

Now that we have a better grasp of homology and its role in trait identification in biology, we are prepared to ask: Can this concept be applied to cognitive capacities?

3 Cognitive homology – Existing accounts

The idea of cognitive homology might seem like a non-starter. Cognitive capacities are characterized by their functional profiles (e.g., "memory is the capacity to recall information that was acquired in the past") rather than their structural features. Given that homology concerns the sameness of structure rather than function (see Section 2), the idea that functionally characterized entities might be homologous seems to be a contradiction in terms (García, 2010; Love, 2007). How are we to make sense of the idea of cognitive homologies?

In this section, we present two existing accounts of cognitive homology—i.e., Claudia García's (2010) and Vincent Bergeron's (2021) — that try to make sense of cognitive homology by explaining how functionally defined entities can still have structural features. The discussion of these accounts illustrates the challenges that arise for accounts of cognitive homology that seek to base the relevant notion of "homology" on structural features of functional entities.

Both Bergeron and García begin their analyses with the already-established notion of "behavioral homology". Homologous behaviors are the same across species due to common ancestry. Tail wagging behaviors in different species of birds and web weaving behaviors in different species of spiders are examples of behavioral homologies. The notion of sameness relevant for behavioral homology is the same as for morphological homology: it is *structural* sameness that matters (García, 2010, p. 126). Behaviors can have similar internal organizations, relative positions in behavioral sequences, and degrees of complexity. Thus, behavioral homology makes plausible notions of "structure" and "structural similarity" that are not restricted to material structure. This makes it an apparently promising starting point for cognitive homology.

Despite their common starting point, García and Bergeron wind up with different accounts of cognitive homology. The crucial difference lies in the structural features they identify relative to which structural similarity, and thus homology, can be spelled out.

Garcia's account characterizes cognitive homology as a particular type of what she calls "functional homology". Functional homology, according to her account, can hold between *cognitive systems* such as neural networks. Two cognitive systems S1 and S2 are functionally homologous if and only if two conditions hold:

> i. S1 and S2 perform the same function, albeit possibly in different ways, i.e., whether differently implemented (…) or subsystemically differently performed (…)[5], and
> ii. S1 and S2 have a common ancestor. (García, 2010, p. 128)

On García's account, the evidence that two cognitive systems are indeed functionally homologous and the criteria for structural similarity between functions is the same as in morphological homology (see Section 2):

- S1 and S2 have the same position in a causal topology of two larger systems T1 and T2 that perform the same function,
- the subsystemic operations of S1 and S2 have the same complexity and the functions of S1 and S2 have the same specificity, and
- the functions have the same place in an evolutionary series of traits.

In other words, García's account consists in a direct translation of the phylogenetic concept of homology to cognitive systems—but without referring to material structure (such as brain structures). Her account allows that the material implementation of the functions can differ between homologous functions.

García's account presents several challenges. One key issue is the ambiguity surrounding condition ii, which states that "S1 and S2 have a common ancestor." Since S1 and S2 are cognitive systems, this must mean that the *organisms implementing* these systems share a common ancestor. However, this alone does not establish homology. Tracing back far enough in evolutionary history, every organism shares a common ancestor with every other organism. Therefore, an important addition to García's condition is that the common ancestor must have already possessed a precursor version of S1 and S2 (call it "S0"), which must have performed the same function as S1 and S2. This requirement poses a significant epistemic challenge. Unlike morphological structures, cognitive functions do not leave fossil records. Even if a common ancestor with S0 can be identified, determining whether S0 performed the same cognitive function as S1 and S2 remains elusive. One potential counterargument might suggest

that, if we have already established that S1 and S2 perform the same function and that a common ancestor had S0, it is highly improbable that S0 did not perform a similar function to S1 and S2.

A second issue arises with condition i. and the criteria for structural sameness of function. If we require that S1 and S2 share the same relative position within their systems, we risk falling into an infinite regress. According to García, the relative positions of S1 and S2 are determined relative to larger systems, T1 and T2, of which S1 and S2 are subsystems, where T1 and T2 *perform the same function*. Thus, to establish structural sameness between the functions of S1 and S2, we need to assess the functions of T1 and T2 and determine whether they are the same. However, under the current proposal, this is the case only if T1 and T2 occupy the same relative position within yet further larger systems. This leads to a regress that can only stop when we reach a system that is not a subsystem of any larger system, or when T1 and T2 are subsystems of the same larger system (e.g., sharing the same ecological niche or ecosystem). In such cases, the present account cannot determine whether the functions are structurally the same. Eliminating the requirement for relative position does not resolve the problem, as the remaining criteria are insufficient to fully define structural sameness of functions. Surely, a more detailed analysis of these criteria is needed, but that is beyond the scope of this chapter.

Bergeron's account avoids the problems of García's account because it takes the underlying neural structures to be essential for determining cognitive homology and does not rely on an account of structural similarities of functions. On Bergeron's account, structural sameness is defined in terms of sameness of material, i.e., brain structure. Because there is already plenty of evidence for homologous brain structures in extant species, it is in principle possible to use comparative anatomical or genetic analyses to determine whether two species that share a brain structure shared a common ancestor that already had a version of it. Cognitive homologies, then, according to Bergeron, are the "cognitive workings" of homologous brain structures.

> [T]he notion of cognitive homology can be understood as the cognitive workings of a homologous (brain) structure, regardless of how these cognitive workings are being used in different species—the same cognitive workings in different species under every variety of cognitive uses, where "the same cognitive workings" means "the cognitive workings of a homologous brain structure". (Bergeron, 2021, p. 36)

Bergeron adopts a distinction between "cognitive use" and "cognitive working". To illustrate, consider the difference between the use and the working of a forelimb. The *use* of a forelimb can be different across species—for example, in cats they are for running, in whales they are for swimming, and in bats they are for flying. Still, the forelimbs' *workings* are the same: there

are similar bones, muscles, and ligaments in similar relative positions that move and interact in similar ways. "Cognitive workings" of brain structures, for Bergeron, are "the basic cognitive contribution [the brain structure] makes to the various cognitive functions it participates in" (ibid. 35), or the basic contributions a brain structure makes to its different uses. According to Bergeron, the basic cognitive contributions of a brain structure can be described in different ways—such as in psychological, computational, or neural terms—as well as in a 'domain-specific' or 'domain-neutral' way. The preferred description depends on the experimental, theoretical, or clinical contexts. So, for Bergeron, the recipe for identifying cognitive homologies seems to be: 1) identify a brain structure that is homologous in different species; and then 2) determine its "basic cognitive contribution"—that is, what the brain structure contributes to its various cognitive uses in different species.

Bergeron's account can only be successful if we assume that homologous brain regions always make the same basic cognitive contribution. It is unclear why this should be the case, i.e., whether it should be impossible that homologous brain regions make different basic contributions to the various cognitive uses that they are involved in. Indeed, there seems to be empirical evidence that homologous brain structures can change their basic cognitive contribution. For example, small structural differences in the internal organization of *nucleus laminaris* (a homologous brainstem auditory region) between barn owls and chickens led to a change in the computational contribution of that brain region allowing owls to detect interaural time differences with greater accuracy (Striedter, 2002, p. 241). Now, one might think that the computational contribution just is the basic cognitive contribution of *nucleus laminaris* which, thus, seems to have changed in owls compared to chicken. Now, Bergeron might argue that the basic cognitive contribution of nucleus laminaris indeed did *not* change in owls—if individuated in a sufficiently coarse-grained way—the basic cognitive contribution of nucleus laminaris is "detecting interaural time differences", not the specific computational contribution that it makes—and that is the still the same in owls and in chicken. This, however, hints at the second problem for Bergeron's account.

The second issue with Bergeron's account is that he does not provide a generalized account of what "basic contributions" might be and how they are to be identified. He provides two examples: the basic cognitive contributions of the amygdala and Broca's area. He suggests that the amygdala's fundamental role is (roughly) "relevance detection," while Broca's area is responsible for (roughly) "translating information into complex motor actions." He argues that there is substantial evidence indicating that structures homologous to the human amygdala and Broca's area exist in non-human animals, and that these structures fulfill the same basic

cognitive functions. However, a critical question remains: How can we determine that "relevance detection" and "translating information into complex motor action" are indeed the basic cognitive contributions of the amygdala and Broca's area?

Bergeron's approach seems to be similar to Price and Friston's (2005) strategy for determining the function of a brain structure. According to Price and Friston, scientists should ascribe functions to brain regions that explain "all patterns of activation" of that brain region (Price and Friston, 2005, p. 268). To do so, they claim, scientists should look at all tasks that elicit activity in the brain region of interest. Then, scientists should identify what these tasks have in common when that brain region is most active. As an example, Price and Friston set out to analyze the function of the *left posterior lateral fusiform*. They conclude it is not—as many researchers have argued—primarily a "visual word form area," an area "sensitive to the visual attributes of animals," or a "tactile-visual region". Rather, they argue, it is only on the assumption that the left posterior lateral fusiform "integrates sensory cues with motor output" that one can explain all its activations during the different tasks.

In the context of Bergeron's account, one could apply a similar strategy. One could compare the activity of the same (i.e., homologous) brain region in members of different species across similar sets of tasks and ask which basic cognitive contribution should be ascribed to the homologous brain region such that the region's activity is best explained.

The problem that arises for Bergeron's account of cognitive homology, then, is the same problem that has been raised in the context of Price and Friston's account of brain function. Recall that our overall goal is to answer the cognitive ontology question: how to carve up the cognitive domain in ways that promote robust scientific prediction, explanation, and generalization. Even if it were adequate to ascribe the basic contribution of "sensory-motor integration" to a brain region, this ascription is not "cognitively interesting" (Klein, 2012, p. 955). First, it is rather difficult to find a brain region that does *not* perform sensory-motor integration (ibid.). Second, categories like "sensory-motor integration" are too broad to increase the predictive and explanatory power of a scientific model or theory of cognition. Such broad function ascriptions are useless for research and theory building—yet this is exactly what we aim to improve by re-thinking cognitive ontology.

Take Bergeron's example of the amygdala. Bergeron writes that "the amygdala can be described as a relevance detector, a structure that contributes to increasing vigilance and attention based on stimulus saliency, ambiguity and unpredictability" (Bergeron, 2021, p. 9). However, the amygdala is also involved in tasks that do not obviously fall under that description. For example, it is also involved in memory consolidation. Ferry et al. (1999)

conclude that the amygdala (or rather one of its nuclei) "is part of a neuromodulatory system that regulates the strength of memories in relation to their emotional significance". Here, the amygdala increases neither vigilance nor attention, and it does not react to external stimuli at all. Surely, one could now try to find a more general description that covers all scenarios in which the amygdala is activated. For example, one could argue that the amygdala's basic cognitive contribution is that it "orchestrates the allocation of processing resources (sensorimotor, cognitive, emotional, endocrine) to a given stimulus based on its relevance, or value, to the organism in a particular context" (Bergeron, 2021, p. 10). To capture the role of the amygdala on memory consolidation, one would have to add something like "to an external stimulus *or internal representation*".

This function description, however, again seems too vague and general to be relevant for a cognitive ontology, similar to "sensory-motor integration". Many brain structures orchestrate the allocation of processing resources to a given stimulus or internal representation based on their relevance or value to the organism in a particular context. For example, the prefrontal cortex (PFC) plays a crucial role in evaluating the relevance of stimuli and making decisions based on the agent's goals. It is heavily involved in orchestrating processing resources (Miller & Cohen, 2001). The basal ganglia are essential for decision-making processes and for selecting actions based on reward and value. The basal ganglia work closely with the PFC to help prioritize which actions or thoughts should be pursued (O'Reilly & Frank, 2006). While the primary sensory cortices are focused on processing specific types of sensory information, they are also involved in allocating processing resources to stimuli that are deemed relevant or valuable (Moore & Zirnsak, 2017). This is often modulated by top-down influences from regions like the PFC. The thalamus acts as a relay station, filtering and prioritizing sensory information before it reaches the cortex, thus playing a role in the allocation of processing resources (Saalmann & Kastner, 2011). The anterior cingulate cortex (ACC) is involved in error detection, conflict monitoring, and assessing the value of different actions, which contributes to the allocation of cognitive and emotional resources (Bush et al., 2000). The hypothalamus is involved in orchestrating the body's response to internal states (e.g., hunger, stress) and ensuring that resources are allocated to maintain homeostasis, which includes both endocrine and autonomic functions (Saper & Lowell, 2014). Thus, the overall process of allocating resources based on relevance and value is a general, distributed function of the brain rather than being isolated to specific regions.

A broader issue with Bergeron's account reflects a common concern about all similarly radical, bottom-up approaches to cognitive ontology: it begins with predefined brain

delineations (in this case, based on the homology of brain regions) and then attempts to identify the functions of these regions. Various authors have expressed different critiques of such approaches. Beyond the challenge of ensuring that the resulting categories remain cognitively meaningful, some have argued that the idea of finding one-to-one mappings between cognitive entities and neural entities is illusory (Viola, 2017). If at all, such one-to-one mappings can only be established when accepting that the functions of brain regions are context-dependent (Klein, 2012). Additionally, it has been suggested that brain regions may not be the appropriate neural entities for deriving a cognitive ontology. Instead, cognitive functions might be better understood in terms of neural networks (Klein, 2012), neural mechanisms (Francken et al., 2022; Krickel, 2024), or neural processes (Viola & Zanin, 2017). Notably, neural mechanisms and processes may not be easily decomposable into brain regions as defined by homology.

Both García's and Bergeron's accounts follow the same general strategy to make sense of the notion of cognitive homology. The core of these accounts are different ways of cashing out the notion of "structural sameness" for cognitive entities. García attempts to locate structural sameness in the functional organization of cognitive systems and runs into a regress. Bergeron reduces the structural sameness of "cognitive workings" to structural sameness of brain regions, and thereby runs the risk of missing the cognitive.

In the next section, we introduce an account of cognitive homology that is not based on structural sameness of cognitive entities, and thus avoids the problems facing these accounts.

4 Cognitive Ontology in terms of Cognitive ChIMs

In this section, we apply DiFrisco, Wagner, and Love's notion of a Character Identity Mechanism to cognitive capacities to develop a novel account of cognitive homology. The core idea is this: If we find the same ChIM producing a cognitive capacity in different species, we can infer that the cognitive capacities are the same. But this raises an important question: what are ChIMs in the cognitive domain?

To introduce the idea of *cognitive ChIMs*, let us first consider an intuitive example. Suppose that you hypothesize that humans' cognitive capacity to form grammatically correct sentences in language is the same as songbirds' cognitive capacity to form correct song patterns (Berwick et al., 2011; Figdor, 2022, p. 11). The capacities seem to share some similarities: for example, both speech and birdsongs have specific, recurring patterns, and both are used in communicative interactions between conspecifics. However, human syntactic capacities and songbird syntactic capacities are also quite different. Human syntax is famously generative: it

allows for productivity (an infinite set of grammatically correct sentences can be formed) and compositionality (the elements of a sentence can be recombined to form novel sentences that get their meanings from the meaning of their elements). By comparison, birdsong syntax is far more limited. Furthermore, these cognitive capacities manifest in a different range of situations for the different species: whereas human language (and thus syntax) occurs in a wide variety of social interactions, birdsong is primarily more narrowly involved in specific behavioral contexts, such as mating, deterring rivals, or defending territory (Berwick et al., 2011). So, are the cognitive capacities underlying human syntax and birdsong syntax different phenotypic expressions of the same cognitive trait (e.g., a capacity to produce regularly patterned stimuli for the purpose of communication), or are they different?

According to our account, this question is to be addressed by investigating the mechanisms by which the capacities *develop* in each species. If and only if they are produced by the same, trait-determining, causally necessary and non-redundant developmental mechanism—that is, the same cognitive ChIM—then they are instances of the same cognitive capacity.

In DiFrisco, Love, and Wagner's original account (2020; see Section 2), ChIMs underlie the development of morphological traits, and thus largely make their contributions during embryogenesis. In contrast, the mechanisms underlying the development of *cognitive* traits are largely manifest during the organisms' postnatal life and often involve learning. Learning involves robust changes in the neural and behavioral setup of an organism. Thus, cognitive ChIMs will involve neural mechanisms as well as the organisms' patterns of interactions with their respective ecological niches (Quartz and Sejnowski, 1997; Sirois and Karmiloff-Smith, 2009; Anderson, 2014). For example, the typical development of cognitive capacities underlying human syntax and birdsong syntax both depend on what is called *experience-expectant* input from the environment during early life (Greenough, Black and Wallace, 1987), such as exposure to adult speech or birdsong (Louder et al., 2024). They also depend on infants' and birds' own attempted production and practice of language and songs (Johnson and Whitney, 2005; Tchernichovski and Marcus, 2014). Furthermore, it has been found that specific neurons in the auditory forebrain, the caudomedial nidopallium (NCM), are crucial for song learning (Louder et al., 2024). It could be demonstrated that there are compelling computational and physiological similarities between corresponding neuronal cell types in the mammalian cortex (Spool et al., 2021). Thus, cognitive ChIMs involve not only organism-internal neural entities and activities, but also species-typical interactions between organisms and their environments.

Determining whether the ChIMs that give rise to the cognitive capacities for linguistic and birdsong syntax are the same would thus involve two steps. First, you must identify the

developmental mechanisms in each species on the neural and the behavioral level. Here, this means that you must identify the set of causally necessary and nonredundant biological endowments and learning experiences that cause humans and birds, respectively, to develop their species-specific capacity to produce regularly patterned stimuli for the purpose of communication (talking/singing). Second, you must compare these mechanisms to determine whether they are instances of the same ChIM. That is, you must determine whether the developmental mechanisms that control the development of the capacity (trait identity) in both species have a common mechanistic architecture.

These considerations, along with the general characterizations of ChIMs outlined in Section 2 above, suggest three desiderata for an account of cognitive ChIMs:

(i) The notion of a cognitive ChIM must take into account that cognitive development rests on neural change.

(ii) The notion of a cognitive ChIM must account for the fact that many cognitive capacities develop through experience-expectant inputs from early, species-typical interactions with the environment that crucially involve the organism's own *behavior*.

(iii) Cognitive ChIMs must be *specific enough* to provide a convincing and useful categorization of cognitive capacities, while at the same time being *general enough* to be possibly present in different species and to abstract away from the phenotypic specificities of the different expressions of that capacity.

Having sketched this preliminary outline of our ontogenetic approach to cognitive homology, we now turn to a more specific possible candidate for a cognitive ChIM as a proof of concept. Here, we consider the developmental mechanism underlying *episodic memory*. Contemporary cognitive science characterizes episodic memory as the capacity to recall specific experiences along with details about the time and place of their occurrence (Tulving, 1993).

Cognitive psychologists Arthur M. Glenberg and Justin Hayes (2016) propose an account for a common mechanism by which episodic memory develops in multiple species. The immediate goal of Glenberg and Hayes' account is to provide an explanation for *infantile amnesia*, the phenomenon that humans cannot remember the first few years of their lives. Interestingly, there is evidence that infantile amnesia occurs in other species, such as monkeys and rodents, as well. Glenberg and Hayes' explanation of infantile amnesia essentially proposes a common developmental mechanism for episodic memory in different species. Our main

purpose here is to show how their account could (in principle) be construed as an account of an "episodic memory ChIM."

The starting point for Glenberg and Hayes' account is that species with infantile amnesia have another feature in common: young juveniles do not self-locomote, but rather are carried around by their caregivers. Based on this observation, Glenberg and Hayes propose that infantile amnesia results from the fact that "the pre-locomotor infant does not have the opportunity to tune his hippocampal place and grid cells [brain regions associated with spatial representation] to the environment" (ibid., 3): in order for this to happen, self-locomotion is necessary. Because episodic memory involves memory for places, an organism with an "untuned" hippocampus will lack the capacity to form such memories. Once a juvenile begins to locomote, "[r]eciprocal connections between the hippocampus and sensorimotor and emotional cortices" develop (ibid., 2).

The explanation of infantile amnesia that Hayes and Glenberg propose relies on the assumption that the cognitive capacity "episodic memory," which is shared by multiple species, is acquired via a specific developmental mechanism that seems to have the role of a ChIM. Assuming species-typical neonate brain anatomy and physiology (i.e., being born with a species-typical hippocampus and sensorimotor cortex), the development of episodic memory is driven by the development of self-locomotion. This proposed mechanism is the same in different species: the development of self-locomotion tunes their hippocampal place and grid cells to regularities in the environment. This mechanism also essentially involves the organism's brain, behavior, and the environment.

A caveat to this 'proof of concept' example is that it is incomplete: the "tuning" of hippocampal place and grid cells to regularities in the environment via self-locomotion is likely only necessary rather than sufficient for episodic memory to develop. In addition, the "development of self-locomotion" and "tuning" would also need to be spelled out in far more detail. Finally, it is also very likely that the "tuning" mechanism described is not specific to episodic memory but rather necessary for the development of other cognitive capacities. As mentioned above, at least further reciprocal connections between the hippocampus and sensorimotor and emotional cortices must be established.

However, despite these caveats, the example suggests that there are indeed plausible candidates for cognitive ChIMs. That is: there are developmental mechanisms that underlie the production of cognitive capacities and are shared across species. The example also illustrates the way that a cognitive ontology based on a notion of cognitive homology grounded in ChIMs can be put to good scientific work: the common developmental mechanism both accounts for

the common cognitive capacity (episodic memory), and it also explains another phenomenon of interest (infantile amnesia) in a way that generalizes to all altricial species (not just humans).

Providing full details for a general account that satisfies the two desiderata outlined above must be left for future work. However, this example suggests that an ontogenetic account of cognitive homology based on cognitive ChIMs is indeed plausible.

5 Conclusion

This chapter developed a proof of concept for a novel account of cognitive homology based on an ontogenetic account of homology from evolutionary-developmental biology. While several previous authors have attempted to shed light on cognitive homology, our account avoids their pitfalls by incorporating insights about *cognitive development.* We propose that cognitive capacities should be individuated by the causally necessary, non-redundant developmental mechanisms that give rise to them—cognitive Character Identity Mechanisms (cognitive ChIMs). These mechanisms crucially involve not only brain structures, but also behavior and environmental interactions.

We began with the *cognitive ontology question*: How should cognitive scientists "carve up" the cognitive domain to generate robust predictions, explanations, and generalizations? Following a growing literature, we suggested that a human-centered cognitive ontology based on outdated folk psychological categories is untenable. We assume that a solution to the cognitive ontology question must incorporate the fact that cognition is the product of evolution. As such, we argue that cognitive categories should, like biological kinds, be individuated in terms of homology: two traits *x* and *y* are the same if and only if *x* and *y* are homologous.

After showing that two existing accounts of cognitive homology run into trouble (due to efforts to characterize *structural similarities* in entities that are functionally individuated), we argued that an account of cognitive homology grounded in ontogenetic mechanisms has the potential to avoid these problems. According to our account, cognitive capacities should be grounded in, and individuated by, cognitive developmental mechanisms: in particular, two cognitive capacities are the same if they are produced by the same cognitive ChIM.

Our account of cognitive ChIMs suggests that, unlike the account of (morphological) ChIMs on which our account is based (DiFrisco et al., 2020), cognitive ChIMs consist in interactions between organisms' brains, behaviors, and environments. Consequently, while investigating the brain is crucial for cognitive ontology, it cannot be the only consideration. On our account, cognitive homologies (and thus cognitive ontology) cannot be constructed by mapping brain

regions to cognitive capacities. Rather, we must identify developmental mechanisms for cognitive capacities that situate the brain in the context of behavior and environment.

References

Anderson, M. L. (2014). *After Phrenology - Neural Reuse and the Interactive Brain*. MIT Press.

Bergeron, V. (2021). Carving the mind at its homologous joints. *Biology & Philosophy*, *36*(4), 36. https://doi.org/10.1007/s10539-021-09812-3

Berwick, R. C., Okanoya, K., Beckers, G. J. L., & Bolhuis, J. J. (2011). Songs to syntax: The linguistics of birdsong. *Trends in Cognitive Sciences*, *15*(3), 113–121. https://doi.org/10.1016/j.tics.2011.01.002

Bush, G., Luu, P., & Posner, M. I. (2000). Cognitive and emotional influences in anterior cingulate cortex. *Trends in Cognitive Sciences*, *4*(6), 215–222. https://doi.org/10.1016/S1364-6613(00)01483-2

Buzsáki, G. (2019). *The Brain from Inside Out*. Oxford University Press. https://doi.org/10.1093/oso/9780190905385.001.0001

Cisek, P. (2019). Resynthesizing behavior through phylogenetic refinement. *Attention, Perception, and Psychophysics*, *81*(7), 2265–2287. https://doi.org/10.3758/s13414-019-01760-1

Craver, C. F. (2007). *Explaining the brain: mechanisms and the mosaic unity of neuroscience*. Oxford University Press.

De Brigard, F. (2017). Cognitive systems and the changing brain. *Philosophical Explorations*, *20*(2), 224–241. https://doi.org/10.1080/13869795.2017.1312503

Dewhurst, J. (2021). Folk Psychological and Neurocognitive Ontologies. In *Neural Mechanisms* (pp. 311–334). https://doi.org/10.1007/978-3-030-54092-0_14

DiFrisco, J. (2023). Toward a Theory of Homology: Development and the De-Coupling of Morphological and Molecular Evolution. *The British Journal for the Philosophy of Science*, *December*, 000–000. https://doi.org/10.1086/714959

DiFrisco, J., Love, A. C., & Wagner, G. P. (2020). Character identity mechanisms: a conceptual model for comparative-mechanistic biology. *Biology & Philosophy*, *35*(4), 44. https://doi.org/10.1007/s10539-020-09762-2

DiFrisco, J., Love, A. C., & Wagner, G. P. (2023). The hierarchical basis of serial homology and evolutionary novelty. *Journal of Morphology*, *284*(1), 1–18. https://doi.org/10.1002/jmor.21531

Ferry, B., Roozendaal, B., & McGaugh, J. L. (1999). Role of norepinephrine in mediating stress hormone regulation of long-term memory storage: A critical involvement of the amygdala. *Biological Psychiatry*, *46*(9), 1140–1152. https://doi.org/10.1016/S0006-3223(99)00157-2

Figdor, C. (2018). *Pieces of Mind - The Proper Domain of Psychological Predicates*. Oxford University Press.

Figdor, C. (2022). What could cognition be, if not human cognition?: Individuating cognitive abilities in the light of evolution. *Biology and Philosophy*, *37*(6), 1–21. https://doi.org/10.1007/s10539-022-09880-z

Francken, J. C., Slors, M., & Craver, C. F. (2022). Cognitive ontology and the search for neural mechanisms: three foundational problems. *Synthese*, *200*(5), 378. https://doi.org/10.1007/s11229-022-03701-2

Fusco, G. (2022). Serial Homology. *Biological Theory*, *17*(2), 114–119. https://doi.org/10.1007/s13752-021-00395-6

García, C. L. (2010). Functional Homology and Functional Variation in Evolutionary Cognitive Science. *Biological Theory*, *5*(2), 124–135. https://doi.org/10.1162/BIOT_a_00036

Glenberg, A. M., & Hayes, J. (2016). Contribution of embodiment to solving the riddle of infantile amnesia. *Frontiers in Psychology*, *7*(JAN), 1–6. https://doi.org/10.3389/fpsyg.2016.00010

Greenough, W. T., Black, J. E., & Wallace, C. S. (1987). Experience and brain development. *Child Development*, *58*(3), 539–559. https://doi.org/10.1111/j.1467-8624.1987.tb01400.x

Griffiths, P. E. (2007). The phenomena of homology. *Biology and Philosophy*, *22*(5), 643–658. https://doi.org/10.1007/s10539-007-9090-x

James, W. (1890). *The principles of psychology*. Henry Holt and Co. https://doi.org/10.1037/10538-000

Johnson, F., & Whitney, O. (2005). Singing-driven gene expression in the developing songbird brain. *Physiology & Behavior*, *86*(3), 390–398. https://doi.org/10.1016/J.PHYSBEH.2005.08.009

Khalidi, M. A. (2023). *Cognitive Ontology*. Cambridge University Press. https://doi.org/10.1017/9781009223645

Klein, C. (2012). Cognitive ontology and region- versus network-oriented analyses. *Philosophy of Science*, *79*(5), 952–960. https://doi.org/10.1086/667843

Krickel, B. (2024). The New Mechanistic Approach and Cognitive Ontology — Or : What Role do (Neural) Mechanisms Play in Cognitive Ontology ? *Minds and Machines*, 1–19. https://doi.org/10.1007/s11023-024-09679-9

Louder, M. I. M., Kuroda, M., Taniguchi, D., Komorowska-Müller, J. A., Morohashi, Y., Takahashi, M., Sánchez-Valpuesta, M., Wada, K., Okada, Y., Hioki, H., & Yazaki-Sugiyama, Y. (2024). Transient sensorimotor projections in the developmental song learning period. *Cell Reports*, *43*(5), 114196. https://doi.org/10.1016/j.celrep.2024.114196

Love, A. C. (2007). Functional homology and homology of function: Biological concepts and philosophical consequences. *Biology and Philosophy*, *22*(5), 691–708. https://doi.org/10.1007/s10539-007-9093-7

McCaffrey, J., & Wright, J. (2022). 14 Neuroscience and Cognitive Ontology: A Case for Pluralism. In F. De Brigard & W. Sinnott-Armstrong (Eds.), *Neuroscience and Philosophy* (pp. 427–466). MIT Press.

Miller, E. K., & Cohen, J. D. (2001). An Integrative Theory of Prefrontal Cortex Function. *Annual Review of Neuroscience*, *24*(1), 167–202. https://doi.org/10.1146/annurev.neuro.24.1.167

Moore, T., & Zirnsak, M. (2017). Neural Mechanisms of Selective Visual Attention. *Annual Review of Psychology*, *68*, 47–72. https://doi.org/10.1146/ANNUREV-PSYCH-122414-033400

O'Reilly, R. C., & Frank, M. J. (2006). Making working memory work: A computational model of learning in the prefrontal cortex and basal ganglia. *Neural Computation*, *18*(2), 283–328. https://doi.org/10.1162/089976606775093909

Pampush, J. D., & Daegling, D. J. (2016). The enduring puzzle of the human chin. *Evolutionary Anthropology*, *25*(1), 20–35. https://doi.org/10.1002/evan.21471

Poldrack, R. A. (2010). Mapping mental function to brain Structure: How can cognitive Neuroimaging Succeed? *Perspectives on Psychological Science*, *5*(6), 753–761. https://doi.org/10.1177/1745691610388777

Poldrack, R. A., Kittur, A., Kalar, D., Miller, E., Seppa, C., Gil, Y., Stott Parker, D., Sabb, F. W., & Bilder, R. M. (2011). The cognitive atlas: Toward a knowledge foundation for cognitive neuroscience. *Frontiers in Neuroinformatics*, *5*(September), 1–11. https://doi.org/10.3389/fninf.2011.00017

Price, C. J., & Friston, K. J. (2005). Functional ontologies for cognition: The systematic definition of structure and function. *Cognitive Neuropsychology*, *22*(3–4), 262–275. https://doi.org/10.1080/02643290442000095

Quartz, S. R., & Sejnowski, T. J. (1997). The neural basis of cognitive development: A constructivist manifesto. *Behavioral and Brain Sciences*, *20*(4), 537–596. https://doi.org/10.1017/S0140525X97001581

Saalmann, Y. B., & Kastner, S. (2011). Cognitive and Perceptual Functions of the Visual Thalamus. *Neuron*, *71*(2), 209–223. https://doi.org/10.1016/J.NEURON.2011.06.027

Saper, C. B., & Lowell, B. B. (2014). The hypothalamus. *Current Biology : CB*, *24*(23), R1111-6. https://doi.org/10.1016/J.CUB.2014.10.023

Sirois, S., & Karmiloff-Smith, A. (2009). Ontogenetic Development Matters. In *Cognitive Biology* (pp. 321–334). The MIT Press. https://doi.org/10.7551/mitpress/9780262012935.003.0293

Spool, J. A., Macedo-Lima, M., Scarpa, G., Morohashi, Y., Yazaki-Sugiyama, Y., & Remage-Healey, L. (2021). Genetically identified neurons in avian auditory pallium mirror core principles of their mammalian counterparts. *Current Biology*, *31*(13), 2831-2843.e6. https://doi.org/10.1016/J.CUB.2021.04.039/ATTACHMENT/D39BC6B3-6587-4FEC-8676-70BEB6B01E7B/MMC2.MP4

Striedter, G. F. (2002). Brain homology and function: An uneasy alliance. *Brain Research Bulletin*, *57*(3–4), 239–242. https://doi.org/10.1016/S0361-9230(01)00692-X

Tchernichovski, O., & Marcus, G. (2014). Vocal learning beyond imitation: mechanisms of adaptive vocal development in songbirds and human infants. *Current Opinion in Neurobiology*, *28*, 42–47. https://doi.org/10.1016/J.CONB.2014.06.002

Tulving, E. (1993). What Is Episodic Memory? *Psychological Science*, *2*(3), 67–70.

Vanderwolf, C. H. (2007). The evolving brain: The mind and the neural control of behavior. *The Evolving Brain: The Mind and the Neural Control of Behavior*, 1–104. https://doi.org/10.1007/978-0-387-34230-6/COVER

Viola, M. (2017). Carving Mind at Brain's Joints. The Debate on Cognitive Ontology. *Phenomenology and Mind*, *0*(12), 162–172. https://doi.org/10.13128/Phe

Viola, M., & Zanin, E. (2017). The standard ontological framework of cognitive neuroscience: Some lessons from Broca's area. *Philosophical Psychology*, *30*(7), 945–969. https://doi.org/10.1080/09515089.2017.1322193

Wagner, G. P. (1989). The Biological Homology Concept. *Annual Review of Ecology and Systematics*, *20*, 51–69. https://www.jstor.org/stable/2097084

Wagner, G. P. (2007). The developmental genetics of homology. *Nature Reviews Genetics*, *8*(6), 473–479. https://doi.org/10.1038/nrg2099

Weiskopf, D. A., & Adams, F. (2015). *An introduction to the philosophy of psychology*. Cambridge University Press.

Yeo, B. T. T., Krienen, F. M., Eickhoff, S. B., Yaakub, S. N., Fox, P. T., Buckner, R. L., Asplund, C. L., & Chee, M. W. L. (2015). Functional Specialization and Flexibility in Human Association Cortex. *Cerebral Cortex*, *25*(10), 3654–3672. https://doi.org/10.1093/cercor/bhu217

Zagaria, A., Ando', A., & Zennaro, A. (2020). Psychology: a Giant with Feet of Clay. *Integrative Psychological & Behavioral Science*, *54*(3), 521–562. https://doi.org/10.1007/S12124-020-09524-5

[1] The cognitive ontology question has a metaphysical and an instrumental (epistemological, or philosophy of science) reading. Understood metaphysically, the cognitive ontology question asks how to "carve cognition at its joints", identifying natural kinds that exist independently of our interests. On an epistemic or instrumental reading, the cognitive ontology question asks how to develop categories that best serve our interests, such as increasing the explanatory and predictive power of scientific models and theories. These readings might differ more in emphasis than in content: for example, the explanatory and predictive power of our models and theories will likely increase depending on how well they capture the actual causal structure of the world (Khalidi, 2023). Here, our primary aim is to contribute to the epistemic or

instrumental project: we approach the cognitive ontology question as a means for improving the explanations, predictions, models, and theories of the cognitive sciences while being indeterminate on whether we thereby "cut nature at its joints".

[2] The notion of a "character" in evolutionary biology is notoriously difficult to define. Roughly, a character is a trait, structure, property, or process that can be "seen" by evolution. The intuition behind this concept can be illustrated with help of an example. Scientists disagree on whether the chin—a forward pointing part of the bone structure below the lower lip—is a character. If chins have been selected for due to an evolutionary benefit, such as resistance against mechanical pressures to the mouth region or a sexual selection advantage, chins are a character. Alternatively, the chin might be simply a side effect of other evolutionarily selected features, such as a shortening of the breadth of the dental arch (Pampush & Daegling, 2016)—in which case it is not a character.

[3] We do not aim to contribute to the debate between proponents of the phylogenetic conception and proponents of the ontogenetic conception. Instead, we follow DiFrisco (2023)'s convincing argument that a satisfying account of homology will likely incorporate both conceptions. For present purposes, what is important is that the ontogenetic conception provides an explanation of trait identity that is objective, rooted in causal mechanisms, and allows for detection and tracing of traits in practice.

[4] When individuating mechanisms, some degree of vagueness cannot be avoided. For example, two mechanism tokens that produce neurotransmitter release might still be considered the "same type" of mechanism if they involve different amounts of sodium and potassium ions, and they might be considered different mechanisms if one involves ion-channels and the other does not (since this would entail different "entities", "activities", and "organization", even if the "phenomenon" of neurotransmitter release were shared; see Section 2). In between the two extremes might lie various mechanism tokens for which questions of "sameness" are less clear. In practice, such vagueness is often resolved based on researchers' explanatory interests.

[5] García also refers to a distinction between "structured" and "unstructured" systems. A system is structured, according to García, if it is "constituted by functional subsystems", whereas an unstructured system is not so constituted (García, 2010, p. 128). We will ignore this distinction here as cognitive systems are plausibly always or even necessarily "structured" (see, e.g., Weiskopf & Adams, 2015, p. 39).