A category theory perspective on compositionality and (the development of) cognitive capacity

Steven Phillips (steve@ni.aist.go.jp)
Mathematical Neuroinformatics Group, National Institute of Advanced Industrial Science and Technology (AIST), Tsukuba, Ibaraki 305-8568 JAPAN

Abstract

A remarkable property of human cognition is the systematic co-occurrence of certain cognitive abilities. One challenge for cognitive science is to determine the (computational) principles that derive this property as part of a broader goal of establishing the foundations of cognitive architecture (i.e. the basic processes and modes of combination affording cognitive capacity) for a science of cognition. This paper continues a category theory approach to compositionality and cognitive capacity that posits universal construction (e.g., products) as a fundamental cognitive architectural component. As shown here, products can be modeled in other frameworks, thereby providing a link between an abstract computational principle and a concrete cognitive resource needed for particular capacities. For example, a network of weighted connections implementing a categorical product uses fewer resources when the number of task instances sharing a common product structure is greater than two; otherwise it is more economic to realize each instance independently. This cross-over may explain why human cognition is not always systematic: the cost of universal construction may not outweigh its expected gain.

Keywords: category theory; cognitive capacity; compositionality; systematicity, similarity; product; universal construction

Introduction

A remarkable property of human cognition is the systematic co-occurrence of certain cognitive abilities. For example, if one has the ability to infer square as the first shape in the pair (square, triangle), then one also has the ability to infer triangle as the first shape in the pair (triangle, square), assuming that squares and triangles are recognizable. This property is called *systematicity* (Fodor & Pylyshyn, 1988), and is characterized as having capacity c_1 *if and only if* c_2 (McLaughlin, 2009), i.e., as an equivalence class of cognitive capacities.

For cognitive science, a major challenge has been to explain *why* cognitive capacity is systematically distributed in a particular (non-arbitrary) way. The main theoretical frameworks propose some form of compositionality to explain systematic cognitive capacity. The *classical* (symbol system) explanation is that cognitive processes are sensitive to *grammatical* structure (Fodor & Pylyshyn, 1988). Thus, the two shape capacities (above) are inseparable because they involve a common process: say, $(p,q) \to p$, where p and q are symbols for squares and triangles, on the assumption of having component processes $\square, \triangle \to p, q$ for recognizing squares and triangles.[1] Similarly, a *connectionist* (neural network) account of capacity can make use of common processes in the form of activation units and weighted connections modeling cognitive processes that are sensitive to *spatial* structure (Phillips & Wilson, to appear). In this case, the two shape capacities factor through a common (sub)network of weighted connections that map (neuronal) vector representations of shape pairs to first shapes. Although classical and connectionist models can be constructed such that capacity c_1 if and only if c_2, systematicity doesn't *necessarily* follow from classical, or connectionist principles: one can also devise models from those principles such that c_1, but not c_2. Additional, *ad hoc* (i.e. arbitrary) assumptions are needed to exclude models that lack systematicity. So, classical and connectionist theories (principles) fail to fully explain systematicity (Aizawa, 2003).

Another approach to the systematicity problem (Phillips & Wilson, 2010, 2011, 2012) used a mathematical theory of structure, called *category theory* (Mac Lane, 2000). In category theory, a universal construction relates a collection of arrows (interpreted as cognitive processes) via a common *mediating* arrow (cognitive process) in a unique way. Hence, all capacities are indivisibly linked to this mediating arrow as the common component. The mediating arrow is associated with an equivalence class of systematically related cognitive capacities (Phillips & Wilson, 2011, Text S4).

This paper further explores a categorical approach to compositionality and cognitive capacity by looking at the relationship between universality and cognitive resources. First, our category theory approach to the systematicity problem, and related models, are recalled to motivate its continued use here. Then, the relationship between a specific universal construction (product) and the resources needed for implementation are examined: fewer resources are needed when the number of task instances sharing a product is greater than two; otherwise, it is cheaper to realize each task instance independently. The implications of this relationship for development (learning) are discussed in the final section.

Compositionality and universal constructions

Definitions of category and product are provided in this section. Other definitions are provided in Appendix A. Deeper introductions to category theory can be found in many books on the subject (see, e.g., Mac Lane, 2000; Simmons, 2011).

Category theory and cognition

Category theory starts with a definition of a category, which is a collection of objects and relations between objects, called arrows (or morphisms, or maps). The category theory approach to cognition presented here regards a cognitive archi-

[1] Systematicity pertains to "molecular" not "atomic" capacities, hence that an architecture permits the recognition of squares, but not triangles, is not a counterexample to having the systematicity property (Fodor & Pylyshyn, 1988).

tecture (i.e., the basic component processes and the ways in which those processes are combined) as a category (of possibly other categories), where objects are interpreted as components of the architecture and arrows are relations between those components. For instance, an object may be a set of cognitive representational states and an arrow may be a function between two sets of states (objects) that transforms cognitive representations (states).

Definition (Category). A *category* $\mathbf{C}$ consists of:

- a class of objects $|\mathbf{C}| = (A, B, \dots)$;
- for each pair of objects A and B in $\mathbf{C}$, a set $\mathbf{C}(A, B)$ of morphisms (also called arrows, or maps) from A to B, where each morphism $f : A \to B$ has A as its domain and B as its codomain, including the *identity* morphism $1_A : A \to A$ for each object A; and
- a composition operation, denoted "$\circ$", of morphisms $f : A \to B$ and $g : B \to C$, written $g \circ f : A \to C$ that satisfies the laws of:
  - *identity*, where $f \circ 1_A = f = 1_B \circ f$, $\forall f : A \to B$; and
  - *associativity*, where $h \circ (g \circ f) = (h \circ g) \circ f$, $\forall f : A \to B$, $g : B \to C$, $h : C \to D$. □

An example of a category is **Set**, which has sets for objects and total functions for arrows, where the identities are the identity functions and composition is function composition. **Set** is used here to model cognitive architecture as a collection of sets of representational states and cognitive processes as functions sending states to states, where identities are "do nothing" functions.

Most models of cognition support some kind of compositionality that affords the representation of a pair of entities from which the constituent entity representations are recoverable. In category theory, a basic kind of (universal) construction of a pair of objects from which the component objects are recoverable is called a *product*.[2]

Definition (Product). In a category $\mathbf{C}$, a *product* of two objects A and B is an object P (also denoted $A \times B$) together with two morphisms $p_1 : P \to A$ and $p_2 : P \to B$, such that for every object $Z \in |\mathbf{C}|$ and every pair of morphisms $f : Z \to A$ and $g : Z \to B$, there is a unique morphism $u : Z \to P$ (also denoted $\langle f, g \rangle$, since it is determined by f and g), such that $f = p_1 \circ u$ and $g = p_2 \circ u$, as indicated in commutative diagram[3]

$$\begin{tikzcd} & Z \arrow[dl, "f"'] \arrow[d, dashed, "\langle f,g \rangle"] \arrow[dr, "g"] & \\ A & A \times B \arrow[l, "p_1"] \arrow[r, "p_2"'] & B \end{tikzcd} \tag{1}$$

□

In **Set**, the *Cartesian product* of sets A and B is the set $A \times B = \{(a,b) | a \in A, b \in B\}$ together with functions $p_1 : (a,b) \mapsto a$ and $p_2 : (a,b) \mapsto b$ recovering the first and second elements of each pair is a product. Accordingly, for any set Z and functions f and g there is a unique function, called the *product function*, $\langle f, g \rangle : z \mapsto (f(z), g(z))$, that maps each element $z \in Z$ to the pair of elements $(f(z), g(z))$ with the desired elements $f(z) \in A$ and $g(z) \in Z$ recoverable via p_1 and p_2. Projections p_1 and p_2 are part of this universal construction; they are the same functions employed for every Z, f and g. Cartesian product is an instance of categorical compositionality, and a concrete example of deriving systematicity with regard to the shape-pairs example (see next).

Categorical compositionality and systematicity

A challenge for cognitive science is to explain why cognitive capacity is organized in a particular way: e.g., why is it that if one can infer square as the first shape in the pair (square, triangle), then one can infer triangle as the first shape in the pair (triangle, square), given that they can recognize squares and triangles; yet, the capacity to infer squares and triangles is independent of the capacity to count? The general, category-theoretic claim is that underlying every collection of systematically related cognitive capacities is a universal construction of some kind. Each member of a collection of systematically related capacities is realized by two cognitive components: a common component and a unique capacity-specific component.[4] The presence or absence of the common arrow implies the presence or absence of the entire collection of capacities. For the shapes example, the common arrow(s) corresponds to a capacity to retrieve the first and second elements of a pair, and the unique arrow corresponds to recognition of specific shape instances. By contrast, the capacity to count involves a different kind of universal construction (Phillips & Wilson, 2012), and hence a different mediating arrow.

Models of categorical compositionality

The relationship to the above categorical theory of compositionality and systematicity to specific (classical, or connectionist) models is analogous to a relationship in categorical algebra, where a model of (say) a theory of groups is a functor (a *structure-preserving* map, see Appendix A) from a category representing the theory and a category capturing the semantics of the model, e.g., **Set**, where the group is based on elements of a set (Lawvere, 1963). A (classical or connectionist) model of categorical compositionality and therefore systematicity is a functor from a universal construction (as a category) to a particular concrete category of cognitive representations and processes.[5]

Returning to the shape example, suppose we model cognitive architecture in the category **Set** where objects are sets of cognitive representations and arrows are cognitive processes

[2]A closely related construction involving pairs of objects is called *coproduct*, which is obtained by reversing the directions of the arrows in the definition of product (Mac Lane, 2000).

[3]I.e. where all pairs of paths (one or both having more than one arrow) from the same start object to the same finish object are equal.

[4]See Appendix A for the general form of universal construction.

[5]Functors compose, so a model at one level may be construed as a theory (functor image as a category) relative to a model of that theory at another level (functor from it into another category).

mapping representations. Consider the collection of objects and arrows (identities not shown) in Diagram 1 as a category (theory). A functor (model) of this category into **Set** is indicated by commutative diagram

$$\begin{tikzcd}
Z \arrow[d, "fst"'] \arrow[drr, "snd"] \arrow[dr, dashed, "u"] & & Z' \arrow[dll, "fst'"'] \arrow[dl, dashed, "u'"'] \arrow[d, "snd'"] \\
S & S \times S \arrow[l, "p_1"] \arrow[r, "p_2"'] & S
\end{tikzcd} \tag{2}$$

where $S = \{\text{square}, \text{triangle}\}$ is a set of elements representing square and triangle, $S \times S = \{(\text{square}, \text{triangle}), \dots\}$ is the set of all pairwise combinations of square and triangle, $Z = \{\square\triangle\}$ is a singleton set containing the image of a square to the left of a triangle, $Z' = \{\triangle\square\}$ is a singleton set containing the image of a triangle to the left of a square, $fst : \square\triangle \mapsto \text{square}$ maps the image of a square and triangle to square, $snd : \square\triangle \mapsto \text{triangle}$ maps the image of a square and triangle to triangle, and $u = \langle fst, snd\rangle : \square\triangle \mapsto (\text{square}, \text{triangle})$. Functions fst' and snd' are defined similarly. Systematicity is realized by common mediating functions (projections) $p_1 : (\text{square}, \text{triangle}) \mapsto \text{square}, \dots$ and $p_2 : (\text{square}, \text{triangle}) \mapsto \text{triangle}, \dots$. This model can be regarded as classical where each element is a symbol, and the constituents of the complex symbols for each pair are tokened whenever each pair is tokened—classical compositionality.

Another model is obtained from the category **Vec** of vector spaces for objects and linear functions for arrows. In the case of the objects and arrows in Diagram 2, suppose S is a vector space over the field of real numbers containing vectors $\vec{s}$ and $\vec{t}$ representing square and triangle, respectively. Likewise, suppose Z and Z' are vector spaces containing vectors $\vec{st}$ and $\vec{ts}$ representing the images $\square\triangle$ and $\triangle\square$, respectively. A product object in **Vec** is a product vector space. Hence, $S \times S$ is a product vector space of vector space S with itself containing all pairwise products of vectors $\vec{s}$ and $\vec{t}$, e.g., $\vec{s}\,\vec{t}$ and $\vec{t}\,\vec{s}$, etc. In this model, the arrows are linear functions mapping vectors in one space to vectors in another space. Specifically, $fst : \vec{st} \mapsto \vec{s}$, $\langle fst, snd\rangle : \vec{st} \mapsto \vec{s}\,\vec{t}$, $p_1 : \vec{s}\,\vec{t} \mapsto \vec{s}$ and $p_2 : \vec{s}\,\vec{t} \mapsto \vec{t}$ (similarly for the other arrows). If we choose particular bases for these vector spaces, then we have a category of coordinate spaces, and the linear functions are realized as matrices, where the identities are identity matrices and composition is matrix multiplication. Further, if we identify each basis vector with a neuron, hence each object is a collection of real-valued neurons, and each matrix as a matrix of weighted connections between neurons in different objects, then we have a connectionist network realizing this model. Systematicity is realized by common linear functions p_1 and p_2. Note that there is no requirement that constituent vectors be tokened whenever their host product vectors are tokened (e.g., the coordinates specifying $\vec{s}$ are not necessarily a subset of the coordinates specifying $\vec{s}\,\vec{t}$)—connectionist (functional) compositionality.

The vector spaces and linear functions in this **Vec**-based model of compositionality contain many more vectors and mappings than may be needed as a model of cognitive capacity. A further connectionist refinement of the model may be to add connections between neurons within an object, e.g., as in a Hopfield network (Hopfield, 1982), so as to restrict the number of representational states to just those needed (in the shape example, two vectors representing square and triangle for the object S, and four vectors representing each pairwise combination of square and triangle for the product object $S \times S$). In this case, we have networks (not just collections of neurons) for objects, whence we need some kind of structure-preserving map for arrows satisfying the usual axioms to be a category. However, not just any combination of objects and arrows constitutes a product (see next). One possibility, for further work, is to extend the category **Gph** of graphs and graph homomorphisms, which has products,[6] by considering a connectionist network as a graph with additional structure.[7]

(Non-)Universality and cognitive capacity

Understanding the relationship between universal constructions, systematic cognitive capacity and cognitive resources requires a specific model of categorical compositionality that gives meaning to a concept of cognitive resource. An example of compositionality that does not involve a universal construction follows. Suppose a **Set**-based model of the shape-related capacities, shown in Diagram 3, that consists of an object $T = \{t_1, t_2\}$ and arrows $g_1 : Z \to T; \square\triangle \mapsto t_1$, $h_1 : T \to S; t_1 \mapsto \text{square}, t_2 \mapsto \text{triangle}$ and $h_2 : T \to S; t_1 \mapsto \text{triangle}, t_2 \mapsto \text{square}$, where composition $h_1 \circ g_1$ correctly infers square as the first object from image $\square\triangle$ and $h_2 \circ g_1$ correctly infers triangle as the second object from image $\square\triangle$. Likewise, suppose arrow $g_2 : Z' \to T; \triangle\square \mapsto t_2$, where composition $h_1 \circ g_2$ correctly infers triangle as the first object from image $\triangle\square$ and $h_2 \circ g_2$ correctly infers square as the second object from image $\triangle\square$. The object T together with arrows h_1 and h_2 do not constitute a product, because for object $Z'' = \{\square\square\}$ and arrows $fst'' : Z'' \to S; \square\square \mapsto \text{square}$ and $snd'' : Z'' \to S; \square\square \mapsto \text{square}$ there does not exist an arrow $u'' : Z'' \to T$ such that $h_1 \circ u'' = fst''$ *and* $h_2 \circ u'' = snd''$. This arrangement does not support the systematically related capacities of inferring square as the first and second shapes from the image $\square\square$. (A similar lack of systematicity also applies for image $\triangle\triangle$, etc.)

$$\begin{tikzcd}
Z \arrow[d, "fst"', "g_1"] \arrow[drr, "snd"] \arrow[dr, dashed] & & Z' \arrow[dll, "fst'"'] \arrow[dl, dashed] \arrow[d, "g_2"', "snd'"] \\
S & T \arrow[l, "h_1"] \arrow[r, "h_2"'] & S
\end{tikzcd} \tag{3}$$

In general, for a network that dedicates one unit for each element of each set, the number of units required to represent

[6]A categorical product of two unlabeled graphs G and H is the Cartesian product of their nodes with edges between pairs of nodes just in case their is an edge between the first node of each pair in G and an edge between the second node of each pair in H.

[7]See Ehresmann and Vanbremeersch (2007) for a category theory description of neurons and systems, and Healy et al. (2009) for an application to neural network modeling.

a categorical product of sets A and B is the size of set A times the size of set B. This additional cost in resources to represent a universal construction raises the question of why a cognitive system would incur such an expense. As shown in the next section, implementing a categorical product becomes cheaper when the number of task instances sharing a common product structure exceeds a certain number. That number depends on the relative costs of representing the task-common and task-specific components.

Resources and (non-)universal constructions

A simple calculation for the resources needed to realize a sequence of task instances with a common product structure using non-universal versus universal constructions reveals a cross-over point. Suppose a sequence of tasks such that each task $i \in \{1,2,3,\ldots,n\}$ consists of two maps $f_i : Z_i \to A$ and $g_i : Z_i \to B$. Consider a cognitive system without the capacity for constructing products. To realize n such tasks, a cognitive system must implement $2n$ functions (i.e. two functions for each task instance). Consider a second architecture with the capacity to construct a product $(A \times B, p_1, p_2)$. In this case, the n tasks are realized by the two projections $p_1 : A \times B \to A$ and $p_2 : A \times B \to B$, and the n unique functions $u_i : Z_i \to A \times B$, totalling $n+2$ functions. Thus, the advantage of constructing products is obtained when $n > 2$, i.e. when the cognitive system must have a capacity for three or more task instances sharing the same product structure.

The precise relationship between cognitive resource and systematic cognitive capacity will depend on the underlying model and the kind of universal construction. The general form of a universal construction (given in Appendix A) with respect to a functor F is a pair (A, ϕ) such that each capacity $f : F(Z) \to Y$ is composed of a common component $\phi : F(A) \to Y$ and a unique component $F(u) : F(Z) \to F(A)$ such that $f = \phi \circ F(u)$. Hence, the number of task instances (n) at which fewer resources are deployed to realize capacity via a universal construction will depend on the relative costs of realizing components ϕ, $F(u)$ and f. The benefit of universal constructions is more pronounced when the most expensive part of realizing f is with ϕ, since the ϕ construction is only required once, whereas one $F(u)$ component is required for each task instance. In general, computing with universal constructions over n tasks becomes cheaper when

$$n > \frac{cost(\phi)}{cost(f) - cost(F(u))} \tag{4}$$

Hence, the benefit of universal construction will be obtained from fewer task instances when the cost of constructing each unique component $F(u)$ is low compared to each f. Otherwise, the benefits will only start to accrue after realizing many task instances.

Discussion

Category theory generalizes the classical and connectionist notions of compositionality. Classical compositionality is the idea that representations of the constituents of complex entities are tokened, in a consistent way, whenever the representations of their complex host entities are tokened (Fodor & Pylyshyn, 1988). Symbol systems are the paradigmatic classical cognitive architecture, e.g., where the symbols representing constituents *John*, *Mary*, and *loves* appear in the symbolic representation of complex entity *John loves Mary*. Categorical compositionality generalizes this idea from symbols to arrows, and from tokening to arrow composition, while specializing the kinds of compositions to those that constitute universal constructions. This specialization is crucial to avoid otherwise arbitrary assumptions over which modes of tokening capture systematicity.

Connectionist compositionality is the idea that representations of complex entities are a function of the representations of their constituent entities without necessarily being tokened whenever the complex host representation is tokened (van Gelder, 1990). Category theory generalizes this idea from functions between sets to arrows between objects with additional internal structure beyond set membership and spatial structure (e.g., a graph, or a group), while specializing to functors (which appear in the definition of universal construction) that preserve that internal structure. This specialization is crucial, since although a functor is a kind of generalized function mapping objects *and* arrows, not all generalized functions are functors (i.e., preserve identity and associativity). Again, specialization is crucial to avoid arbitrary assumptions over which generalized functions capture systematicity.

Systematic capacity and cognitive development

Models of universal constructions provide a functorial link from an abstract computational principle to a concrete cognitive resource. A connectionist example illustrated this point as a capacity-resource tradeoff. This non-universal situation may arise during cognitive development when the cognitive system does not have sufficient cognitive/neural resources to compute products, as suggested by the example in **Set**. Consistent with this idea, a category theory approach to the analysis of reasoning in young children explained developmental differences in terms of the capacity to compute (co)products (Phillips, Wilson, & Halford, 2009). The analysis was based on empirical evidence from multi-paradigm studies where children from different age groups were tested on various reasoning tasks under *easy* and *hard* conditions (see Halford, 1993): a repeatedly observed result (see, e.g., Andrews & Halford, 2002; Andrews, Halford, Bunch, Bowden, & Jones, 2003; Andrews, Halford, Murphy, & Knox, 2009) was that if a child (typically older than five years) was significantly above chance for both easy and hard conditions on one reasoning task (e.g., *transitive inference*), then that child was significantly above chance for easy and hard conditions on another reasoning task (e.g., *class inclusion*); conversely, if a child (typically younger than five years) was significantly above chance for the easy, but not the hard condition on one task, then they were significantly above chance for the easy, but not the hard condition on another task. Analysis of the easy and hard conditions for seven reasoning tasks used to test children in such multi-paradigm studies showed that all hard conditions involved some kind of binary (co)product, in contrast to the easy conditions that did not, or what could also be called a unary (co)product[8] (Phillips et al., 2009).[9]

This relationship between universality and cognitive resource may shed light on why human cognition is not always systematic: the initial (short-term) cost of realizing a universal construction does not outweigh the expected (long-term) gain. Put in more familiar terms, when faced with a problem one often has two choices: devise a quick and simple solution that works for the current situation only, or a general-purpose solution that takes more time to develop. This choice will depend on how likely the same problem will reappear in other contexts, and whether the cost of developing a general-purpose solution will be repaid in subsequent savings. From a developmental perspective, this choice may be forced: the necessity of having a solution that works at least for the currently presented cases may preclude development of a general-purpose alternative. Hence, children may first go through a period of non-systematicity before attaining some forms of systematic cognitive capacity.

If non-systematicity is forced by the immediate demands on having a functioning system, then how do universal constructions (systematic capacity) eventually develop? The development of universal constructions may be driven (in part) by limited cognitive resources and the subsequent need to simultaneously realize multiple instances of tasks sharing a common component.

Further work is needed to explore the implications of the relationship between compositionality, systematic cognitive capacity and cognitive resource. The category-theoretic analysis presented here identified the essence of this relationship.

Acknowledgments

This work was supported by a Japanese Society for the Promotion of Science Grant-in-aid (Kakkenhi: 22300092).

References

Aizawa, K. (2003). *The systematicity arguments*. New York: Kluwer Academic.

Andrews, G., & Halford, G. S. (2002). A cognitive complexity metric applied to cognitive development. *Cognitive Psychology*, *45*(2), 153–219.

Andrews, G., Halford, G. S., Bunch, K. M., Bowden, D., & Jones, T. (2003). Theory of mind and relational complexity. *Child Development*, *74*(5), 1476–1499.

Andrews, G., Halford, G. S., Murphy, K., & Knox, K. (2009). Integration of weight and distance information in young children: The role of relational complexity. *Cognitive Development*, *24*(1), 49–60.

Ehresmann, A. C., & Vanbremeersch, J.-P. (2007). *Memory evolutive systems: Hierarchy, emergence, cognition* (Vol. 4). Oxford, UK: Elsevier.

Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, *28*(1–2), 3–71.

Halford, G. S. (1993). *Children's understanding: The development of mental models*. Hillsdale, NJ: Lawrence Erlbaum.

Healy, M. J., Olinger, R. D., Young, R. J., Taylor, S. E., Caudell, T. P., & Larson, K. W. (2009). Applying category theory to improve the performance of a neural architecture. *Neurocomputing*, *72*(13–17), 3158–3173.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, *79*(8), 2554–2558.

Lawvere, F. W. (1963). Functorial semantics of algebraic theories. *Proceedings of the National Academy of Sciences*, *50*(5), 869–872.

Mac Lane, S. (2000). *Categories for the working mathematician* (2nd ed.). New York, NY: Springer.

McLaughlin, B. P. (2009). Systematicity redux. *Synthese*, *170*, 251–274.

Phillips, S., Takeda, Y., & Singh, A. (2012). Visual feature integration indicated by phase-locked frontal-parietal EEG signals. *PLoS ONE*, *7*(3), e32502.

Phillips, S., & Wilson, W. H. (2010). Categorial compositionality: A category theory explanation for the systematicity of human cognition. *PLoS Computational Biology*, *6*(7), e1000858.

Phillips, S., & Wilson, W. H. (2011). Categorial compositionality II: Universal constructions and a general theory of (quasi-)systematicity in human cognition. *PLoS Computational Biology*, *7*(8), e1002102.

Phillips, S., & Wilson, W. H. (2012). Categorial compositionality III: F-(co)algebras and the systematicity of recursive capacities in human cognition. *PLoS ONE*, *7*(4), e35028.

Phillips, S., & Wilson, W. H. (to appear). A category theory explanation for systematicity: Universal constructions. In P. Calvo & J. Symons (Eds.), *Systematicity and cognitive architecture* (p. xx). Cambridge, MA: MIT Press.

Phillips, S., Wilson, W. H., & Halford, G. S. (2009). What do Transitive Inference and Class Inclusion have in common? Categorical (co)products and cognitive development. *PLoS Computational Biology*, *5*(12), e1000599.

Simmons, H. (2011). *An introduction to category theory*. New York, NY: Cambridge University Press.

van Gelder, T. (1990). Compositionality: A connectionist

[8] A unary product object is just A, which is isomorphic to the product of A with a terminal object (i.e., $A \times 1 \cong A$). In the dual case, A is isomorphic to the coproduct of A, denoted $+$, with an *initial object*, denoted 0 (i.e., $A + 0 \cong A$).

[9] One possible neural resource involved in categorical products is *synchrony*: analysis of EEG *phase-locking values* (PLV) during visual search for a target identifiable on one, two, or three feature dimensions, corresponding to a unary, binary, and ternary product (respectively), revealed significantly positive PLV-arity slopes for (frontal, parietal) electrode pairs (Phillips, Takeda, & Singh, 2012).

variation on a classical theme. *Cognitive Science*, *14*, 355–384.

Appendix A

Definition (Functor). A *functor* $F : \mathbf{C} \to \mathbf{D}$ is a map from category $\mathbf{C}$ to category $\mathbf{D}$ that associates each object A in $\mathbf{C}$ an object $F(A)$ in $\mathbf{D}$; and each map $f : A \to B$ in $\mathbf{C}$ a map $F(f) : F(A) \to F(B)$ in $\mathbf{D}$, such that $F(1_A) = 1_{F(A)}$ for each object A in $\mathbf{C}$; and $F(g \circ_{\mathbf{C}} f) = F(g) \circ_{\mathbf{D}} F(f)$ for all maps $f : A \to B$ and $g : B \to C$, where $\circ_{\mathbf{C}}$ and $\circ_{\mathbf{D}}$ are compositions in categories $\mathbf{C}$ and $\mathbf{D}$. □

Two examples of functors, used in the construction of products, are the *diagonal* and *product* functors. The diagonal functor $\Delta : \mathbf{C} \to \mathbf{C} \times \mathbf{C}; A \mapsto (A,A), f \mapsto (f,f)$ sends objects and arrows to pairs of objects and arrows. The product functor $\Pi : \mathbf{C} \times \mathbf{C} \to \mathbf{C}; (A,B) \mapsto A \times B, (f,g) \mapsto f \times g$ send pairs of objects and arrows to their respective products.

Definition (Universal morphism/construction). Given a functor $F : \mathbf{A} \to \mathbf{C}$ and an object $Y \in |\mathbf{C}|$, a *universal morphism* from F to Y is a pair (A, ϕ) where A is an object in $\mathbf{A}$, and ϕ is a morphism in $\mathbf{C}$, such that for every object $Z \in |\mathbf{A}|$ and every morphism $f : F(Z) \to Y$, there exists a unique morphism $u : Z \to A$, such that $\phi \circ F(u) = f$, as indicated in commutative diagram

$$\begin{tikzcd}
Z \arrow[d, dashed, "u"'] & F(Z) \arrow[d, dashed, "F(u)"'] \arrow[dr, "f"] & \\
A & F(A) \arrow[r, "\phi"'] & Y
\end{tikzcd} \tag{5}$$

A *universal construction* is a universal morphism, or its dual a *couniversal morphism*, obtained by reversing the directions of the arrows in the definition of universal morphism. □

A product $(A \times B, (p_1, p_2))$ is an instance of a universal morphism, hence a universal construction, and indicated in commutative diagram

$$\begin{tikzcd}
Z \arrow[d, dashed, "\langle f,g \rangle"'] & (Z,Z) \arrow[d, dashed, "(\langle f,g \rangle,\langle f,g \rangle)"'] \arrow[dr, "(f,g)"] & \\
A \times B & (A \times B, A \times B) \arrow[r, "(p_1,p_2)"'] & (A,B)
\end{tikzcd} \tag{6}$$

where functor $F : \mathbf{A} \to \mathbf{C}$ in the definition of universal morphism is the diagonal functor, and product object $A \times B$ is obtained by application of the product functor to object (A,B). Compare Diagram 6 with Diagram 1.