Analogy and cognitive architecture: Two kinds of systematicity, one kind of (universal) construction

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

Abstract

Cognitive science recognizes two kinds of systematicity: (1) as the property where certain cognitive capacities imply certain other related cognitive capacities (Fodor & Pylyshyn, 1988); and (2) as the principle that analogical mappings based on collections of connected relations are preferred over relations in isolation (Gentner, 1983). These two kinds of systematicity were shown to derive from one type of (universal) construction (Phillips, 2014), using category theory (Mac Lane, 2000). Underlying both forms of systematicity is a kind of optimization. We provide an informal summary of this result, and suggest an extension to address other (semantic) aspects of analogy.

Keywords: analogy; cognitive architecture; systematicity; category theory; universal construction

Introduction

Cognitive science recognizes two kinds of systematicity. One kind of systematicity is the property of cognition where the capacity for certain cognitive abilities implies the capacity for certain other cognitive abilities, i.e. capacity is distributed around equivalence classes of cognitive abilities (Fodor & Pylyshyn, 1988). Another kind of systematicity is the preference for analogical mappings based on collections of connected relations over relations in isolation (Gentner, 1983). Whether these two kinds of systematicity are aspects of a deeper property of cognition is hitherto unknown.

In previous work (Phillips & Wilson, 2010, 2011, 2012, 2014), we explained Fodor and Pylyshyn's kind of systematicity using the category theory notion of *universal construction* (Mac Lane, 2000). We also showed that these two kinds of systematicity are derived from one type of universal construction as an optimization of cognitive resources (Phillips, 2014). We provide an informal summary of this result (after recalling the two notions of systematicity), and suggest extensions to address some other (semantic) aspects of analogy.

Systematicity (Fodor and Pylyshyn)

A remarkable property of human cognition is the distribution of cognitive capacities, where the capacity for certain cognitive abilities implies the capacity for certain other cognitive abilities. For example, suppose one is shown pairs of geometric shapes such as a square to the left of a triangle. If one has the ability to infer square as the left shape in the pair (square, triangle), then one also has the ability to infer triangle as the left shape in the pair (triangle, square), assuming that squares and triangles are individually recognizable. This property is generally referred to as *systematicity* (Fodor & Pylyshyn, 1988), and is characterized more broadly as having capacity c_1 *if and only if* c_2 (McLaughlin, 2009), i.e. as equivalence classes of cognitive capacities.

For cognitive science, a major challenge has been to provide a theory of cognitive architecture that *explains* systematicity: i.e. *why* certain cognitive capacities are distributed in a particular, non-arbitrary way (Fodor & Pylyshyn, 1988; Fodor & McLaughlin, 1990; Aizawa, 2003). Cognitive architecture is the collection of basic processes and their modes of composition that together provide the basis for a slew of cognitive abilities, from recognizing concrete physical objects to reasoning about abstract mathematical structures. The problem with the major theoretical frameworks, including classicism (symbols systems) and connectionism (neural networks), is that systematicity does not necessarily follow from the core principles and assumptions of their proposed theories. Although classical and connectionist theories are sufficiently general to derive models supporting the systematicity property, they are insufficiently specific to rule out models (derived from the same core principles and assumptions) that do not support systematicity (in those same cognitive domains). Further, *ad hoc* (essentially, arbitrary) assumptions are needed that take up the explanatory slack to exclude those classical or connectionist models that do not support systematicity, and so neither classical nor connectionist theory fully explains systematicity (Aizawa, 2003).

Systematicity (Gentner)

In the context of analogy, *systematicity* is the observation that when matching source and target domains of (relational) knowledge people match systems of (higher-order) relations in preference to isolated (lower-order) relations (Gentner, 1983). This observation is embodied as the *systematicity principle*, or assumption, in the *structure mapping theory* of analogy (Gentner, 1983). Structure mapping theory supposes that relational knowledge consists of a system of concepts arranged in tree-like structures. Three kinds of concepts are distinguished: *objects*, *attributes* and *predicates*. Object concepts are things like *John*; they are constants, which are nodes in a concept tree from which there are no more branches. Attributes and predicates are concepts that express some proposition about other concepts. An attribute expresses a proposition about a single concept, e.g., *Male*(*John*) expresses the proposition that *John is a male*, and its concept tree structure has *Male* (attribute) as the root concept node and John (object) at its only branch. Predicates are concepts that express propositions about two or more concepts, e.g., *Loves*(*John*, *Mary*) expresses the proposition that *John loves Mary*, and its concept tree structure has *Loves* (predicate) as the root node, *John* at its first branch and *Mary* at its second branch. Later, we consider objects, attributes and predicates as instances of relations of arity n: i.e. nullary ($n = 0$), unary ($n = 1$) and n-ary ($n \geq 2$) relations, respectively.

The Water-Heat flow analogy example (see Gentner, 1983) illustrates the systematicity principle in analogical mapping. Suppose the following relational knowledge (concept trees):

1. *And*(*Contains*(*Vessel*, *Water*), *On-top-of*(*Lid*, *Vessel*)); and
2. *Cause*(*And*(*Puncture*(*Vessel*), *Contains*(*Vessel*, *Water*)), *Flows-from*(*Water*, *Vessel*)).

Suppose two corresponding trees by replacing the objects in the water flow domain with corresponding objects in the heat flow domain, as given by the following object concept mappings: *Vessel* → *House*, *Lid* → *Roof*, and *Water* → *Heat*. The systematicity principle as expressed in this example is the preference for mapping the second tree over the first, because the second tree involves a larger system of (higher-order) predicates than the first (Gentner, 1983).

Category theory: universal constructions

In this section, we provide an informal introduction to the category theory concepts needed to explain the two kinds of systematicity, i.e. universal constructions, which depend on the concepts of category and functor. Corresponding formal definitions and results are given in Phillips (2014). Deeper introductions to category theory can be found in many texts on the subject (e.g., Mac Lane, 2000; Simmons, 2011).

Category

A category is a collection of objects and arrows (“relations”) between objects with a composition operation for composing arrows in a way that satisfies certain axioms. We will consider the shapes example as part of a category whose objects are sets, arrows are functions between sets, and composition operation is function composition. Some of these objects (sets) are sets of perceptual, or conceptual states for corresponding shapes, and the arrows are functions (cognitive processes) transforming representational states.

Objects and arrows may be constructed from other objects and arrows. The set of shape concept pairs, for example, is constructed from the *Cartesian product* of the set of shape concepts, $S = \{\text{square}, \text{triangle}\}$, with itself: i.e. the set $S \times S = \{(\text{square}, \text{square}), (\text{square}, \text{triangle}), (\text{triangle}, \text{square}), \ldots\}$, which is another object. Elements are retrieved by two functions (projections): $\pi_1 : S \times S \to S$; $(\text{square}, \text{triangle}) \mapsto \text{square}$, $(\text{triangle}, \text{square}) \mapsto \text{triangle}$, etc.; and $\pi_2 : S \times S \to S$; $(\text{square}, \text{triangle}) \mapsto \text{triangle}$, $(\text{triangle}, \text{square}) \mapsto \text{square}$, etc. In general, the Cartesian product of sets A and B is the set $A \times B$ of all pairwise combinations of the elements taken from sets A and B, and two functions, $\pi_1 : A \times B \to A$ and $\pi_2 : A \times B \to B$, that return the first and second elements of each pair. A Cartesian product is a product in the category **Set**. More generally, in some category **C**, a product of objects A and B is an object P (also denoted $A \times B$) together with two arrows $\pi_1 : A \times B \to A$ and $\pi_2 : A \times B \to B$ such that certain universal conditions are met (see Universal construction).

Functor

Functors are to categories as arrows are to objects. They send objects and arrows in one category to (respectively) objects and arrows in another category. Functors can also be considered as a way of constructing categories from other categories. For example, the *product functor* constructs product objects and arrows from pairs of objects and arrows.

Universal construction

The intuition behind the formal notion of universal construction involves the idea of capturing the common component of a collection of entities (arrows). We can see this intuition in action from our shapes example. Observe that every pair of maps that extracts the first and second shape concept from shape images ($first_i$ and $second_i$) can be composed of a map sending each image to a pair of shape concepts in the Cartesian product set and the projections for extracting the first and second shape concepts from each pair of shape concepts. For example, the map $first_1 : \square\triangle \mapsto \text{square}$ is composed of the map $u_1 : \square\triangle \mapsto (\text{square}, \text{triangle})$ and the projection $\pi_1 : (\text{square}, \text{triangle}) \mapsto \text{square}$. The map $first_2 : \triangle\square \mapsto \text{triangle}$ is composed of the map $u_2 : \triangle\square \mapsto (\text{triangle}, \text{square})$ and π_1. Maps $first_1$ and $first_2$ share the common component map π_1. Similarly, maps that extract the second shape concept from each image, $second_i$, share the common component map $\pi_2 : (\text{square}, \text{square}) \mapsto \text{square}, (\text{square}, \text{triangle}) \mapsto \text{triangle}, (\text{triangle}, \text{square}) \mapsto \text{square}, (\text{triangle}, \text{triangle}) \mapsto \text{triangle}$. Universal constructions can also be thought of as a kind of optimization relative to the underlying functor, described next.

Systematicity and universal constructions

Two kinds of systematicity are derived from universal constructions. For succinctness, Fodor and Pylyshyn’s kind of systematicity is termed *F-systematicity*, and Gentner’s kind of systematicity is termed *G-systematicity*.

F-systematicity (Fodor and Pylyshyn)

The shapes example is used for the explanation of F-systematicity based on universal constructions. In this example of F-systematicity, if one has the capacity to infer from $\square\triangle$ that the left shape is square, then one also has the capacity to infer from $\triangle\square$ that the left shape is triangle, and likewise for the right shape in each instance. Cognitive architecture is modeled in the category **Set** where objects are sets of cognitive representations, arrows are cognitive processes mapping representations, and the composition operator is function composition. For the specific shapes example, we have objects that are sets of representations of shape concepts (indicated by name, e.g., square) and images (indicated by symbol, e.g., $\square$), and arrows that are functions from representations to representations. For example, the set of shape concepts is the set $S = \{\text{square}, \text{triangle}\}$, the set containing the square-triangle image is the *singleton* (one-element) set $Z_1 = \{\square\triangle\}$, and the set containing the triangle-square image is the singleton set $Z_2 = \{\triangle\square\}$. (We also have sets $Z_3 = \{\square\square\}$ and $Z_4 = \{\triangle\triangle\}$.) The arrow representing the capacity to infer from $\square\triangle$ that the left shape is square is the function $\mathit{first}_1 : Z_1 \to S; \square\triangle \mapsto \text{square}$, and the arrow representing the other left-shape inferential capacity is $\mathit{first}_2 : Z_2 \to S; \triangle\square \mapsto \text{triangle}$. Likewise, we have arrows for right-shape inferential capacities: $\mathit{second}_1 : Z_1 \to S; \square\triangle \mapsto \text{triangle}$, and $\mathit{second}_2 : Z_2 \to S; \triangle\square \mapsto \text{square}$.

F-systematicity follows from the fact that in **Set** we also have the Cartesian product set of all pairwise combinations of elements of S, i.e. $S \times S = \{(\text{square}, \text{square}), (\text{square}, \text{triangle}), \dots\}$, and two functions (projections) that return the first and second elements of each pair, i.e. $\pi_1 : S \times S \to S; (\text{square}, \text{triangle}) \mapsto \text{square}$, etc., and $\pi_2 : S \times S \to S; (\text{square}, \text{triangle}) \mapsto \text{triangle}$, etc. Together, the Cartesian product and projections constitute the product construction $(S \times S, \pi_1, \pi_2)$, which is an instance of a universal construction. As a universal construction, for each set Z_i and each function $\mathit{first}_i : Z_i \to S$, there must exist a unique function $u_i : Z_i \to S \times S$ such that $\mathit{first}_i = \pi_1 \circ u_i$ and $\mathit{second}_i = \pi_2 \circ u_i$. Indeed, for first_1 and second_1, we have the function $u_1 : Z_1 \to S; \square\triangle \mapsto (\text{square}, \text{triangle})$, where $\mathit{first}_1(\square\triangle) = \pi_1((\text{square}, \text{triangle})) = \pi_1(u_1(\square\triangle)) = \pi_1 \circ u_1(\square\triangle)$. This function, u_1, is the only function that satisfies the equality $(\mathit{first}_1, \mathit{second}_1) = (\pi_1, \pi_2) \circ u_1$, as required by the definition of product. Likewise, $u_2 : Z_2 \to S; \triangle\square \mapsto (\text{triangle}, \text{square})$ is the only function satisfying $(\mathit{first}_2, \mathit{second}_2) = (\pi_1, \pi_2) \circ u_2$. The collection of objects and arrows modeling the shape capacities is given in the following *commutative diagram* (i.e. paths from the same start object to the same end object are equal, where one path has at least two arrows)—sets and arrows associated with the $\square\square$ and $\triangle\triangle$ cases are not shown; a dashed arrow indicates uniqueness:

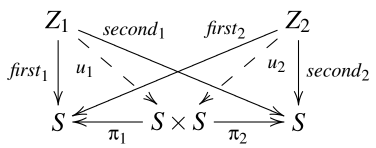

(1)

Systematicity is realized by common functions π_1 and π_2: the presence/absence of each arrow implies the presence/absence of each collection of inferential capacities.

F-systematicity, universal constructions and optimization A universal construction is also a kind of optimization in the sense that it consists of an object (from a collection of objects) that is "closest" (relative to the collection) to an object of interest. If we consider each cognitive capacity as a path from one set of cognitive states (input) to another set of cognitive states (output), then closest is interpreted in terms of path length, i.e. the number of component arrows between two objects. The following diagram illustrates this conception by comparing the paths associated with the set $S \times S$ and the paths associated with the set T, corresponding to architectures having universal and non-universal constructions (respectively):

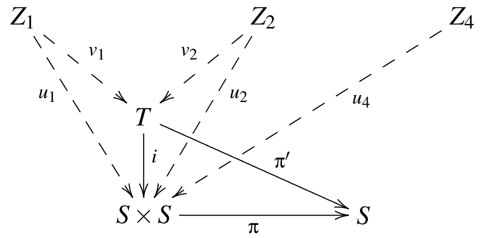

(2)

This diagram omits the functorial component of this universal construction and simplifies some objects and arrows to highlight the characterization of universal constructions as a kind of optimization. The omitted details are given in Phillips (2014). In particular, the projections, π_1 and π_2, are simply denoted here as the arrow $\pi : S \times S \to S$, which stands in for the arrow pair $(\pi_1, \pi_2) : (S \times S, S \times S) \to (S, S)$. Likewise, u_i and Z_i stands in for the arrow $(u_i, u_i) : (Z_i, Z_i) \to (S \times S, S \times S)$, and similarly for v_j and T. With the exception of S, π and π', these objects and arrows belong to the *image* of the diagonal functor, sending objects and arrows to pairs of objects and arrows. The image of a functor $F : \mathbf{C} \to \mathbf{D}$ is the collection of objects $F(A)$ and arrows $F(f)$.

Diagram 2 reveals the sense in which object $S \times S$ is closer to S than T. Arrow π' is composed of two arrows, the projection π and the injection i. Hence, the "distance" (number of component arrows) from T to S (two) is greater than the distance from $S \times S$ to S (one). All capacities $f_i : Z_i \to S$ are composed of π, only capacities f_1, f_2 and f_3 are composed of π'. Universal construction as closeness motivates a reconceptualization of analogy as structure approximation to provide a categorical treatment of G-systematicity, next.

G-systematicity (Gentner)

All categorical constructions, including universal constructions, reside in a category of some kind. So, the first step in providing a categorical account of G-systematicity is to recast source and target knowledge domains in terms of a suitable category. The second step is to show that G-systematicity derives from a universal construction in regard to that category. The explanation for G-systematicity also considers two cases of analogy: (1) a special case, where the source and target knowledge domains each consist of a single concept tree; and (2) the general case, where the source or target domains consist of multiple concept trees.

Special case: single pair of trees Each concept tree is considered as an object in some (to be specified) category. In the water flow knowledge domain, for example, the binary relation *Contains*(*Vessel*, *Water*) is represented by the tree $\langle \text{Contains}, (\langle \text{Vessel} \rangle, \langle \text{Water} \rangle) \rangle$, where the first component is the relational concept (*Contains*) and the second component is the list of trees representing the related concepts (*Vessel* and *Water*). An n-ary relation is represented by a tree consisting of an n-ary relational concept and a list of n concept trees.

Universal constructions are a kind of optimization. Optimization suggests analogical mapping as structure approximation. Approximation ordering for inductively defined data structures, such as lists and trees, affords a categorical treatment of recursion (Bird, 1998). Hence, approximation ordering over trees is considered as the arrows of this category of concept trees. In this context, approximation refers to partial knowledge about some concept tree. For example, suppose the contents of a vessel are unknown. A representation of this partial knowledge is the approximation tree $\langle \text{Contains}, (\langle \text{Vessel} \rangle, \bot) \rangle$, where the symbol $\bot$ indicates the unknown concept tree. Conversely, for example, suppose one does not know the source of a water leak. This situation is represented by the tree $\langle \text{Flows-from}, (\bot, \langle \text{Water} \rangle) \rangle$.

Trees are (partially) ordered by an *approximation order relation*, denoted $\sqsubseteq$. The expression $t \sqsubseteq r$ says that tree t is no better an approximation (expresses no more knowledge) than a tree r; or, in passive form, tree r is at least as good an approximation (expresses at least as much knowledge) as tree t. The definition of the specific approximation order relation for concept trees has two parts that formalize the following intuitions: (1) the concept tree $\bot$ is no better an approximation (expresses no more knowledge) than any tree t, and (2) recursively, an n-ary tree t is no better an approximation (expresses no more knowledge) than an n-ary tree r whenever the two trees express the same relational concept and each related tree t_i is no better an approximation than its corresponding related tree r_i. Formally, the approximation order relation for n-ary trees of arity 0 to N is defined by:

$$\bot \sqsubseteq t \quad \text{and}$$

$$\langle a, (t_i)_{i=1}^n \rangle \sqsubseteq \langle b, (r_i)_{i=1}^n \rangle \Leftrightarrow (a = b) \wedge \left(\bigwedge_{i=1}^{n} t_i \sqsubseteq r_i \right)$$

We have the following examples:

- $\bot \sqsubseteq \langle \text{Contains}, (\langle \text{Vessel} \rangle, \langle \text{Water} \rangle) \rangle$;
- $\langle \text{Contains}, (\langle \text{Vessel} \rangle, \bot) \rangle \sqsubseteq \langle \text{Contains}, (\langle \text{Vessel} \rangle, \langle \text{Water} \rangle) \rangle$; and
- $\langle \text{Contains}, (\bot, \langle \text{Water} \rangle) \rangle \sqsubseteq \langle \text{Contains}, (\langle \text{Vessel} \rangle, \langle \text{Water} \rangle) \rangle$.

There is no order relation between $\langle \text{Contains}, (\bot, \langle \text{Water} \rangle) \rangle$ and $\langle \text{Contains}, (\langle \text{Vessel} \rangle, \bot) \rangle$, for example.

The approximation order relation has three properties.

1. *Reflexivity* ($t \sqsubseteq t$): t no better approximates itself.
2. *Transitivity* ($s \sqsubseteq t \wedge t \sqsubseteq r \Rightarrow s \sqsubseteq r$): if s is no better an approximation than t, and t is no better an approximation than r then s is no better an approximation than r.
3. *Antisymmetry* ($t \sqsubseteq r \wedge r \sqsubseteq t \Rightarrow t = r$): if t is no better an approximation than r, and r is no better an approximation than t then t is the same as r.

Hence, the approximation relation is a *partial order*. The set of tree approximations, $T_\bot$, together with the partial order, $\sqsubseteq$, constitute a *partially ordered set*, or *poset*, denoted $(T_\bot, \sqsubseteq)$.

The composition operator, $\circ$, for partial orders is *conjunction*. The transitivity property of partial orders means that if we have (the corresponding arrows) $a \sqsubseteq b$ *and* $b \sqsubseteq c$ then we have $a \sqsubseteq c$, which satisfies the requirement for a category that there is an arrow for every pair of *composable* arrows. The reflexivity property of partial orders means that for every element (object) a we have (the corresponding arrow) $a \sqsubseteq a$, which satisfies the requirement for a category that every object has an identity arrow. The proof that this collection of objects, arrows and composition operator is a category follows immediately from the fact that $(T_\bot, \sqsubseteq)$ is a partially ordered set and every partially ordered set is a category.

We now want to consider a particular universal construction in regard to this category. We consider an analogical mapping between two trees t and r as involving their greatest common approximation tree, i.e. the tree p that shares the greatest number of (higher-order) relational concepts, which is a product of concept trees in the category of trees and approximations. The definition of the *greatest common approximation* (*gca*) of two trees is motivated by the following considerations: (1) if either tree is the no approximation (no knowledge) tree, $\bot$, then their gca is also the no approximation tree; (2) if either tree represents a different relational concept then their gca is also the no knowledge tree; and (3) if both trees represent the same relational concept then their gca is that relational concept together with the gca of each pair of trees at the corresponding role of the relation. Formally, the gca for trees $t, r \in T_\bot$ is defined by:

$$\begin{aligned} gca(t, \bot) &= \bot \\ gca(\bot, r) &= \bot \\ gca(\langle a, (t_i)_{i=1}^m \rangle, \langle b, (r_j)_{j=1}^n \rangle) &= \bot \qquad a \neq b \\ gca(\langle a, (t_i)_{i=1}^n \rangle, \langle a, (r_i)_{i=1}^n \rangle) &= \langle a, (gca(t_i, r_i))_{i=1}^n \rangle \end{aligned}$$

Some examples of the gca of two trees follow:

- $gca(\langle \text{Contains}, (\langle \text{Vessel} \rangle, \langle \text{Water} \rangle) \rangle, \bot) = \bot$;
- $gca(\langle \text{Contains}, (\langle \text{Vessel} \rangle, \langle \text{Water} \rangle) \rangle, \langle \text{Contains}, (\langle \text{Vessel} \rangle, \langle \text{oil} \rangle) \rangle) = \langle \text{Contains}, (\langle \text{Vessel} \rangle, \bot) \rangle$;
- $gca(\langle \text{Contains}, (\langle \text{Vessel} \rangle, \langle \text{Water} \rangle) \rangle, \langle \text{Contains}, (\langle \text{House} \rangle, \langle \text{Heat} \rangle) \rangle) = \langle \text{Contains}, (\bot, \bot) \rangle$; and
- $gca(\langle \text{Contains}, (\langle \text{Vessel} \rangle, \langle \text{Water} \rangle) \rangle, \langle \text{Flows-from}, (\langle \text{Vessel} \rangle, \langle \text{Water} \rangle) \rangle) = \bot$.

As the last example illustrates, gca is the greatest common approximation tree, not the correspondence between two trees: Contains and Flows-from are relational concepts, not relational concept trees, hence the greatest common approximation tree is not $\langle \bot, (\langle \text{Vessel} \rangle, \langle \text{Water} \rangle) \rangle$. Having obtained the gca, a subsequent process can be employed to obtain correspondences between the other concepts.

The gca of trees t and r is their greatest lower bound p. The poset $(T_\perp, \sqsubseteq)$ is a category where each tree $t \in T_\perp$ is an object in the (poset as a) category $(T_\perp, \sqsubseteq)$. The product of two trees t and r in this category is their gca together with two approximation arrows: i.e. $(gca(t,r), \sqsubseteq_t, \sqsubseteq_r)$. The proof follows from the proof that the gca of two trees is their greatest lower bound, and that the greatest lower bound is a product in a poset considered as a category (see Phillips, 2014).

The product of trees t and r is the tree with the greatest number of connected higher-order relations in common to t and r (together with their approximation arrows). A product is a universal construction. Hence, a universal construction provides an explanation for G-systematicity.

General case: multiple pairs of trees In general, a source and target domain may consist of multiple trees, as illustrated in the Water-Heat flow example in the Introduction. The explanation generalizes to this situation, sketched here.

Suppose there are multiple candidate pairs of source-target trees. The pairs of source-target trees considered during an analogy constitute a list of tree pairs. Computing the gca of each pair gives a list of product trees. Since product trees are also trees, we can also define an ordering on them. In this case, the ordering is over tree size, rather than tree approximation. Again, we have a partially ordered set and hence a product of product trees as the greatest lower bound. Suppose size is a natural number, for example, indicating tree height which corresponds to the order of the root relation, where size of the unknown tree is zero, i.e., $size(\perp) = 0$. The set of natural numbers, $\mathbb{N}$, and the usual ordering on them, $\leq$, is the poset $(\mathbb{N}, \leq)$. This set is also a *totally ordered set* (i.e. every pair of elements in the set is ordered) and the product of two natural number objects x and y in the poset as a category $(\mathbb{N}, \leq)$ is the minimum of x and y: e.g., the *categorical* product $2 \times 3 = (2, \leq_2, \leq_3)$. In the case that we require the maximum size of two trees, we can work in the dual (*opposite*) category $(\mathbb{N}, \geq)$, whose product is the maximum of two numbers: e.g., in $(\mathbb{N}, \geq)$, the *categorical* product $2 \times 3 = (3, \geq_2, \geq_3)$. Equivalently, in $(\mathbb{N}, \leq)$, the maximum of two numbers is the dual, *coproduct* (universal) construction (denoted, $+$): e.g., the *categorical* coproduct $2 + 3 = (3, \leq_2, \leq_3)$. These two steps give us the largest common approximation tree for the Water-Heat flow analogy, which corresponds to G-systematicity.

Discussion and extensions

Psychologically, we can consider universal constructions as a kind of optimization of cognitive resources. In the context of cognitive capacity, the F-systematicity property affords the ecological benefit of not having to expend further resources for an already present (component) cognitive capacity (Phillips & Wilson, 2010). In the context of analogy, the G-systematicity property affords more correspondences between source and target knowledge domains (Gentner, 1983), and therefore greater opportunities to exploit knowledge in one domain for inferences in another. Hence, systematic cognitive capacity and analogical mapping of systems of higher-order relations are two aspects of one optimization principle.

How this category-theoretic level of analysis maps into a symbolic level, e.g., the *structure mapping engine* (Forbus, Gentner, & Law, 1995), or a connectionist level, e.g., DORA (Doumas, Hummel, & Sandhofer, 2008), of analysis is a further challenge. At the neural level, one approach that we have mentioned before (Phillips & Wilson, 2014) is to propose a suitable category of neural networks and arrows between networks. Networks are a kind of graph. The category of graphs and graph homomorphisms has products, so one possibility is a variation of this category that also has products.

A psychological interpretation of a categorical construction may also depend on the nature of the arrows and ambient category. The explanations for F-systematicity and G-systematicity involved one kind of universal construction (product), but two types of arrows and hence categories: functions between sets for F-systematicity, and order relations between trees or numbers for G-systematicity. Interpreting F-systematicity in terms of common processes mapping cognitive states to states is natural, but a similar interpretation for G-systematicity may appear less so, since the arrows appear to be comparisons (not transformations) between objects. Nonetheless, G-systematicity also has an interpretation in terms of common arrows: in a poset as a category, every comparison between objects z and t (and r), i.e. $z \leq t$, factors through the comparison of the least upper bound p of t and r, i.e. $p \leq t$, because the corresponding arrow $\leq_{zt} = \leq_{pt} \circ \leq_{zp}$ (i.e. z less-than-or-equal t whenever z less-than-or-equal p *and* p less-than-or-equal t). Hence, $\leq_{pt}$ is the common arrow underlying all comparisons with t, defining an equivalence class of capacities for comparisons with t (and likewise r).

"Bottom-up" components

There is a large literature on computational models of analogy for a broad range of phenomena (see Gentner & Forbus, 2010, for a review), and the category theoretical approach presented here is a modest step towards integrating properties of analogy with other components of cognition. One important aspect of analogy not addressed here is the role of the one-to-one correspondence principle that is a central feature of theories of analogy, such as structure mapping theory (Gentner, 1983). For example, the gca of *Causes*(*Loves*(*John*, *Mary*), *Kisses*(*John*, *Mary*)) and *Causes*(*Loves*(*Jane*, *Marcia*), *Kisses*(*Jane*, *Marcia*)) is the same as the gca of *Causes*(*Loves*(*John*, *Mary*), *Kisses*(*John*, *Mary*)) and *Causes*(*Loves*(*Jane*, *Marcia*), *Kisses*(*Susan*, *Tony*)), yet we may expect a preference for the first choice given that the repeating components (e.g., *John* as the lover and as the kisser) represent the same concept. One possibility is to include the dual notion of coproducts by considering each repetition as a single concept with more than one parent, i.e. by considering the structure as a lattice instead of a tree. In this case, matching is based on both top-down (product) and, dually, bottom-up (coproduct) universal constructions.

Semantic components

Another aspect of analogy not addressed here is the semantic relatedness of concepts, which is addressed in models of analogy such as LISA (Hummel & Holyoak, 1997) and DORA (Doumas et al., 2008) using semantic features. These models prefer *Loves*(*John*, *Mary*) to *Likes*(*Bill*, *Susan*) than *Fears*(*Peter*, *Beth*), because *Loves* and *Likes* share more features than *Loves* and *Fears* (see Hummel & Holyoak, 1997).

One suggestion (sketched here) is to regard the concept tree nodes as indices (pointers) to semantic feature sets. In this case, a concept tree structure has two components: (1) a tree of indices, τ, and (2) an indexed set of semantic feature sets, $\sigma = \{s_i\}_{i \in I_\tau}$, where I_τ is the set of indices given by τ, together denoted as the pair (τ, σ). (Indices may be specified as the paths from root to node.) Approximation is defined componentwise: i.e. $t \sqsubseteq t' \Leftrightarrow \tau \sqsubseteq \tau' \wedge (\bigwedge_{i \in I_\tau \cap I_{\tau'}} (\sigma_i \subseteq \sigma'_i))$. The empty set, $\emptyset$, now takes on the "no knowledge" role of the empty tree, $\perp$. A product in a category of sets and inclusions is set intersection. Thus, a product of concept trees (τ, σ) and (τ', σ') is $(gca(\tau, \tau'), (\sigma_i \cap \sigma'_i)_{i \in I_\tau \cap I_{\tau'}})$.

For multiple product trees, we need a way of comparing tree sizes. Let $m(p) = (m_\tau(p), m_\sigma(p))$ be a measure of product tree size in terms of the sizes of each component. Partial order over product tree size is defined as $(m_\tau(p_1), m_\sigma(p_1)) \leq (m_\tau(p_2), m_\sigma(p_2)) \Leftrightarrow (m_\tau(p_1) \leq m_\tau(p_2)) \wedge (m_\sigma(p_2) \leq m_\sigma(p_2))$. Returning to the Loves-Likes example, the tree structures for the two alternatives (i.e. *Likes*(*Bill*, *Susan*) and *Fears*(*Peter*, *Beth*)) match the source to the same degree, but the size of the intersection for semantic feature sets for *Loves* and *Likes* will be greater than for *Loves* and *Fears*, hence *Likes*(*Bill*, *Susan*) is preferred. To consider the partial order over product trees as a total order, we need some way of integrating the sizes of each component into a single measure. One possibility is a function that accords with the MAX principle (Goldstone, Medin, & Gentner, 1991), where similarity judgement is based on the maximum of relational and featural similarity measures.

Summary and further work

F-systematicity and G-systematicity are two sides of the same coin; two aspects of a common principle, universal construction, which is a form of optimization. Induction is closely related to analogy, and both are modeled by DORA (Doumas et al., 2008). Further work considers induction as universal construction. For example, one approach (Schwering, Krumnack, Kuhnberger, & Gust, 2009), which treats analogy as induction, has its roots in category theory as a kind of universal construction called the *least general generalization* (Plotkin, 1970). Comparing these categorical methods may lead to a deeper understanding of cognition as structural optimization.

References

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

Bird, R. (1998). *Introduction to functional programming using Haskell* (2nd ed.). Harlow, England: Prentice Hall.

Doumas, L. A., Hummel, J. E., & Sandhofer, C. M. (2008). A theory of the discovery and predication of relational concepts. *Psychological Review*, *115*(1), 1–43.

Fodor, J. A., & McLaughlin, B. P. (1990). Connectionism and the problem of systematicity: Why Smolensky's solution doesn't work. *Cognition*, *35*(2), 183–204.

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

Forbus, K. D., Gentner, D., & Law, K. (1995). A model of similarity-based retrieval. *Cognitive Science*, *19*(2), 141–205.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, *7*(2), 47–59.

Gentner, D., & Forbus, K. D. (2010). Computational models of analogy. *Wiley Interdisciplinary Reviews: Cognitive Science*, *2*(3), 266–276.

Goldstone, R. L., Medin, D., & Gentner, D. (1991). Relational similarity and the non-independence of features in similarity judgements. *Cognitive Psychology*, *23*(2), 222–264.

Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, *104*(3), 427–466.

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. (2014). Analogy, cognitive architecture and universal construction: a tale of two systematicities. *PLoS ONE*, *9*(2), e89152.

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. (2014). A category theory explanation for systematicity: Universal constructions. In P. Calvo & J. Symons (Eds.), *The architecture of cognition: Rethinking Fodor and Pylyshyn's systematicity challenge* (pp. 227–249). Cambridge, MA: MIT Press.

Plotkin, G. D. (1970). A note on inductive generalization. *Machine Learning*, *5*(1), 153–163.

Schwering, A., Krumnack, U., Kuhnberger, K. U., & Gust, H. (2009). Syntactic principles of heuristic-driven theory projection. *Cognitive Systems Research*, *10*(3), 251–269.

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