

Research



Cite this article: Phillips S. 2019 Sheaving—a universal construction for semantic compositionality. *Phil. Trans. R. Soc. B* **375**: 20190303.
<http://dx.doi.org/10.1098/rstb.2019.0303>

Accepted: 29 May 2019

One contribution of 16 to a theme issue 'Towards mechanistic models of meaning composition'.

Subject Areas:

cognition

Keywords:

compositionality, category theory, sheaf, sheaving, topology

Author for correspondence:

Steven Phillips
e-mail: steve@ni.aist.go.jp

Sheaving—a universal construction for semantic compositionality

Steven Phillips

National Institute of Advanced Industrial Science and Technology, Tsukuba 305-8566, Japan

SP, 0000-0002-5694-0670

Semantic compositionality—the way that meanings of complex entities obtain from meanings of constituent entities and their structural relations—is supposed to explain certain concomitant cognitive capacities, such as systematicity. Yet, cognitive scientists are divided on mechanisms for compositionality: e.g. a language of thought on one side versus a geometry of thought on the other. Category theory is a field of (meta)mathematics invented to bridge formal divides. We focus on sheaving—a construction at the nexus of algebra and geometry/topology, alluding to an integrative view, to sketch out a category theory perspective on the semantics of compositionality. Sheaving is a universal construction for making inferences from local knowledge, where meaning is grounded by the underlying topological space. Three examples illustrate how topology conveys meaning, in terms of the inclusion relations between the open sets that constitute the space, though the topology is not regarded as the only source of semantic information. In this sense, category (sheaf) theory provides a general framework for semantic compositionality.

This article is part of the theme issue 'Towards mechanistic models of meaning composition'.

1. Introduction

The way that representations and their meanings for complex entities obtain from the representations and meanings for the constituent entities and their structural relations is called *semantic compositionality*. Some form of compositionality is supposed to explain concomitant cognitive capacities, such as the *systematicity* of language [1] and thought [2], i.e. where possessing certain cognitive capacities implies possessing certain other (structurally related) cognitive capacities—an equivalence relation on cognitive abilities [3]—such as understanding the meaning of *John loves Mary* and *Mary loves John*. *Classical compositionality*, for example, supposedly explains systematicity by appealing to a *combinatorial syntax and semantics*, i.e. semantic relations between constituent entities are captured by syntactic relations between corresponding symbols, and cognitive processes that are sensitive to those syntactic/semantic relations [4]. So, to illustrate, the meaning of *John loves Mary* is captured by a symbol, JOHN, for the person *John*, a symbol, MARY, for the person *Mary*, and a symbol, LOVES, for the relation *loves* that are combined (juxtaposed) in a way that expresses the fact that *John* is the lover and *Mary* is the one loved, e.g. LOVES (JOHN, MARY). Classical compositionality can account for systematicity by assuming that structurally related capacities involve the same combinatorial syntactic process [4], e.g. lover is determined by the first argument position. Analogous accounts are supposed for other (non-classical) forms of compositionality, e.g. as used in connectionist (neural network) models, where symbols are replaced with vectors of neural activity and juxtapositioning with operations such as tensor product [5]. Though the claim that these forms are non-classical has been extensively debated [6], they generally suppose that semantic relations are captured by relations between corresponding representations.

Despite the efficacy of supposing a compositionality principle, cognitive scientists are generally divided on symbolic [4] versus non-symbolic [5]

mechanisms—a *language of thought* [2] on one side versus a *geometry of thought* [7] on the other—and their explanatory import [6]. The challenge is not just to explain how some form of compositionality accounts for properties such as systematicity, but *why* cognition is compositional in the first place [8].

Explaining the *why* versus *how* of systematicity was posed as a challenging problem for connectionist theories [4], and later shown to be also problematic for classical theory [6]. Problematically, while there are instances of compositionality that support a requisite systematicity property, there are also instances that do not support the same property. So, systematicity does not necessarily follow from core principles and assumptions of classical or connectionist theories. Auxiliary assumptions added to pick out just those instances of compositionality that support systematicity are *ad hoc* when they are unconnected to the theory's core principles and assumptions, cannot be confirmed independently of confirming the theory, and are motivated only by the need to fit the data, in which case, the theory fails to fully explain systematicity [6]. One recourse is to claim that the supposed counterexamples are not the 'canonical' forms of compositionality that classical theory takes as a core assumption [3]. Yet, it's unclear what characterizes canonicity, or why cognition is canonically compositional [9].

A *category theory* [10] approach to compositionality was introduced to address the *why* of systematicity [11]. Category theory is a field of (meta)mathematics invented to formally compare mathematical structures [12]. The core explanatory concept is *universal construction*, formalized as *universal morphism*, which is a way of comparing cognitive capacities modelled as compositions of maps—such constructions are characterized by a *universal mapping property* [13]: in regard to a collection of systematically related cognitive capacities, each map modelling a member capacity is composed of the map shared by all members and a map that is unique to that capacity. Hence, a universal morphism identifies an equivalence class of systematically related cognitive capacities. Such constructions are the 'best' one can do within a certain (categorical) context—every construction in that context 'leads to' a universal construction, so necessarily obtains via a recursive process [9].

An explanation for semantic compositionality must ultimately connect to the physical (neural) system that supports cognition. Classical theory assumes that symbols are supported by a neural system that implements the equivalent of memory registers, i.e. the *physical symbol system* hypothesis [14]. Connectionist theory makes this link more directly as the representations that supposedly support semantic compositionality are instantiated as neural activity for a network of (abstract) neurons. A categorical approach must also make this kind of connection. To this end, the current work focuses on another universal morphism, called *sheaving* [15] or *sheafification* [16], to sketch out a category theory perspective on the semantics of compositionality. Sheaving is a construction at the nexus of algebra and geometry/topology, which alludes to an integrative view. This view starts with a (*pre*)*sheaf* to model cognitive representations as data attached to a topological space [17]. As we shall see, the underlying topological space gives meaning to the data in terms of the relations between the open sets that constitute the topology.

The presentation of this work is primarily informal to facilitate an intuitive understanding of the approach. Connections to

formal details appear elsewhere [17], and deeper introductions to categories and sheaves appear in many textbooks on these topics [10,16,18,19]. We proceed with an example of a universal morphism that serves to illustrate the basic category theory concepts (§2) underlying the examples of sheaving given in the context of cognition (§3). This approach is discussed by comparison and contrast with classical notions of compositionality and possible neural mechanisms (§4). For convenience and to help ground concepts, some formal details appear in the appendix.

2. Categories and (universal) compositionality

We use *playing cards* as a running example of compositionality to bootstrap the needed category theory from the more familiar concepts of sets and functions. Each card has a *rank* (i.e. *two, three, ... , ten, jack, queen, king, ace*) and a *suit* (i.e. *spade, club, diamond, heart*). For example, *queen* and *heart* constitute the *queen of hearts*. The ranks can be represented by the set of symbols $Rank = \{2, 3, 4, 5, 6, 7, 8, 9, 10, J, Q, K, A\}$, the suits by the set of symbols $Suit = \{\spadesuit, \clubsuit, \diamondsuit, \heartsuit\}$ and the cards by the *Cartesian product* of those sets: $Card = Rank \times Suit = \{(2, \spadesuit), (2, \clubsuit), \dots\}$. For instance, the pair of symbols (Q, \heartsuit) represents the *queen of hearts*. This product also comes with two functions that retrieve the rank and suit of each card: e.g. $rk : (Q, \heartsuit) \mapsto Q$ and $st : (Q, \heartsuit) \mapsto \heartsuit$. Accordingly, sets and functions provide a basic set-theoretic model of *playing cards*.

Category theory starts with the formal concept of a *category* (definition A.1), which consists of a collection of entities, called *objects*, a collection of relations between objects, called *morphisms*, and an operation that takes two morphisms and returns a morphism, called *composition*. The archetypal category is **Set** (example A.2), the category of sets (objects) and functions (morphisms), with function composition as the composition operation (remark A.3). Hence, sets *Rank*, *Suit* and *Card* are objects and functions *rk* and *st* are morphisms in **Set**, constituting a *categorical product* (definition A.6), which is the Cartesian product for this category (example A.7). A deck of cards is modelled as a mapping of each face, signifying a *playing card*, to the corresponding symbol, e.g. a function $card : Face \rightarrow Card; Q^\heartsuit \mapsto (Q, \heartsuit)$. The mappings from faces to ranks and from faces to suits are given by compositions $faceRank = rk \circ card$ and $faceSuit = st \circ card$, respectively: e.g. $faceRank : Q^\heartsuit \mapsto Q$, which says that the rank of the card signified by the face Q^\heartsuit is *Q* (remark A.8). Thus, we have a category-theoretic model of the same *playing cards* concept.

Having introduced categories, we can now look at basic constructions and their relations. A *functor* (definition A.12) is a way of constructing, indexing, or identifying objects and morphisms. For example, the product functor (example A.14) constructs the set of cards from the sets of ranks and suits, i.e. $\Pi : (Rank, Suit) \mapsto Rank \times Suit$, and a constant functor identifies the set of cards (i.e. the functor that sends every set and function, in **Set**, to the set of cards, *Card*, and its identity function, 1_{Card}). Two functors are related by a *natural transformation* (definition A.15), and the optimal (or most efficient) transformation pertains to a *universal morphism* (definition A.17). For example, the transformation from the set of cards to their ranks and suits is the universal morphism $(Card, rs)$, where $rs = \{rk, st\}$. The transformation is efficient in that there are no more and no fewer mappings than needed to retrieve the rank and suit of every card.

Note that universal morphisms are *unique up to unique isomorphism* (remark A.19). So, constituents need not be ‘tokened’ in the classical sense. A characteristic of classical compositionality is that the symbols representing constituents are *tokened* (inscribed, or written out) whenever the representation of their complex host is tokened [4]. The symbol pair representation of cards is an example of tokening: for instance, the symbols for *queen*, Q , and *heart*, \heartsuit , are tokened whenever the symbol for *queen of hearts*, (Q, \heartsuit) , is tokened. In category theory, the product of two sets is conventionally given as the Cartesian product, but other products exist. For example, the cards can be represented as numbers, say from 1 to 52, provided the accompanying functions retrieve the requisite components. Being an isomorphic set is not sufficient, because one still needs the appropriate functions to recover the constituents—such isomorphisms are generally not unique (remark A.19).

3. Sheaving: bridging gaps in knowledge

Our categorical approach to semantic compositionality involves *presheaves/sheaves* (functors) and *sheaving* (natural transformation). A presheaf/sheaf (definitions A.20/A.21) models data attached to a topological space (definition A.4). A sheaf is a presheaf where the attached data are globally coherent, i.e. agree on overlapping regions. *Pullbacks* (definition A.9) express global coherency conditions (remark A.22). For **Set**, a pullback of f and g (example A.10) is a constrained product (remark A.11), which consists of only those pairs, (a, b) , whose components map to a common value (property): $f(a) = g(b)$. Hence, pullbacks pertain to non-local (global) properties. Sheaving is a universal morphism that constructs the ‘nearest’ sheaf from a given presheaf (remark A.23). This construction is likened to the *natural join* operation (example A.24) that extracts information from data stored locally in different tables of a relational database—say, the addresses of all people prescribed a particular medication, where contact and medical data are stored in separate tables. In this way, sheaving is a kind of relational inference: a way of bridging gaps in knowledge via meaning grounded in the underlying topological space.

We give three examples of sheaving that pertain to cognition. The first example continues the introduction to category (sheaf) theory constructions via the familiar concept of *playing cards*. The second example involves visual feature binding [17] extended for triple conjunction search [20]. The third example involves a simple version of depth perception. Each example illustrates the different ways that meaning is conveyed by the relations between the open sets that constitute the topology.

(a) Playing cards

The *playing cards* example, introduced earlier, can be considered as a presheaf or sheaf on a topological space constituted by elements identifying the (feature) dimensions of rank and suit. For example, suppose the rank and suit dimensions are labelled as R and S , respectively. The set of dimension labels $D = \{R, S\}$ together with the topology $\{\emptyset, \{R\}, \{S\}, \{R, S\}\}$ constitute a *discrete topological space*, which consists of all subsets of labels and their inclusion relations (example A.5). And, the values of each card constitute the data attached to that space. For example, the *queen of hearts* and *two of spades* are represented

by the presheaf, $\mathcal{F}_{Q2} : D^{\text{op}} \rightarrow \mathbf{Set}$. In database terms, this presheaf can be regarded as a collection of tables whose attributes (headings) correspond to the open sets and rows correspond to the attached data, e.g. there is a two-column table whose attributes correspond to the open set $\{R, S\}$ that has two rows: one row for the *queen of hearts* and one row for the *two of spades* (example A.26). In sheaf theory terms, \mathcal{F}_{Q2} sends each open set to the set of functions on that set—each function maps the elements of the open set to the attached data—e.g. $\mathcal{F}_{Q2} : \{R, S\} \mapsto \{c_{QH}, c_{2S}\}$, where $c_{QH} : R \mapsto Q, S \mapsto \heartsuit$ and $c_{2S} : R \mapsto 2, S \mapsto \spadesuit$. The inclusions given by the topology are preserved as restrictions on functions, e.g. $\{R\} \subseteq \{R, S\}$ maps to the restriction $f|_R : c_{QH} \mapsto c_Q, c_{2S} \mapsto c_2$. Restriction corresponds to (database) projection of a table along the specified attribute(s).

Sheaving affords the systematic capacity to represent all cards (example A.27), but this capacity depends on the topology. To illustrate, suppose one knows the ranks and suits, i.e. there is a one-column table of 13 rows for ranks and a one-column table of four rows for suits. In this situation, sheaving simply constructs all pairwise combinations of ranks and suits, which is the sheaf $\mathcal{F}_{\text{card}}^+$. Thus, we have a systematic capacity to represent all 52 cards. One can think of sheaving as a kind of completion, or limit process—adding just enough rows to make a sheaf.

A contrasting scenario is where one knows some of the cards without knowing about constituents *rank* and *suit*: cards are understood as non-compositional entities. This situation is captured by the *indiscrete topology* (example A.5), i.e. $\{\emptyset, D\}$. Sheaving, in this case, does not add any rows to the table containing just the known (non-compositional) cards. Hence, one does not necessarily have a systematic capacity to represent all cards. Completion is trivial—the presheaf is a sheaf—because the topology does not consist of any other (non-empty) open sets.

This difference between sheaving with respect to a discrete versus indiscrete topological space was used to model the difference between generalization and lack of generalization observed with participants trained on cue-target maps [17]. The participants who failed to generalize were regarded as having learned the mappings from cues to targets—pairs of letters to coloured shapes—as mappings of non-compositional entities.

(b) Visual feature binding

Visual feature binding concerns the capacity to identify, say, a red square and a blue triangle, as opposed to a red triangle and a blue square based on globally coherent spatial information (location). This process is modelled as the sheaving of colour and shape location maps to obtain a colour-shape conjunction map that corresponds to objects observed in the visual field as needed to perform visual search [17]. Here, we show how this example of sheaving extends straightforwardly to triple conjunction search [20], i.e. where the target of search is identifiable by a triple of features, such as colour, orientation and (spatial) frequency.

In this example, we start with the set of dimension labels $D = \{C, O, F, L\}$ representing colour, orientation, frequency and location. The feature-location maps are data attached to the corresponding open sets of the topology: $\{\emptyset, \{L\}, \{C, L\}, \{O, L\}, \{F, L\}, \{C, O, L\}, \{C, F, L\}, \{O, F, L\}, \{C, O, F, L\}\}$. For instance, suppose there are four locations, i.e. $L = \{1, 2, 3, 4\}$, two colours $C = \{\text{red}, \text{blue}\}$, two orientations $O = \{\text{hor}, \text{ver}\}$

and two frequencies $F = \{low, high\}$. In database terms, the identification of the four objects, indexed by location, $(red, hor, low)_1, (red, hor, high)_2, (blue, hor, low)_3, (red, ver, low)_4$ obtains from the natural join of the corresponding feature-location tables: e.g. the join of the colour-location table with rows $(red)_1, (red)_2, (blue)_3, (red)_4$ and the orientation-location table with rows $(hor)_1, (hor)_2, (hor)_3, (ver)_4$ obtains the colour-orientation table with rows $(red, hor)_1, (red, hor)_2, (blue, hor)_3, (red, ver)_4$; and the join of colour-orientation with frequency-location $(low)_1, (high)_2, (low)_3, (low)_4$ obtains the four objects.

In terms of universal morphisms, sheaving involves pullbacks (remark A.22). For instance, the colour-orientation map obtains from the pullback of the projections of the colour-location (CL) and orientation (OL) maps onto location: $\pi_2: CL \rightarrow L$ and $\pi_2: OL \rightarrow L$ to obtain the colour-orientation map, denoted $C \times_L O$, and its projections. Thus, triple conjunction obtains from two pullbacks: $(C \times_L O) \times_L F$.

The topology in this example conveys a different (relational) meaning from the meanings conveyed by the discrete and indiscrete topologies. Each topology induces a corresponding order over the elements of the underlying space, called the *specialization (pre)order* (remark A.28): $C \leq L, O \leq L, F \leq L$ for the current example, which says that colour, orientation and frequency specialize location; conversely, location is a general (global) property of the data (object features) attached to the topological space. By contrast, the discrete topology in the cards example has the corresponding order $R \leq R, S \leq S$, which says that neither dimension is a specialization of the other. In other words, the dimensions are independent; sheaving is effectively a Cartesian product of the sets of values on those dimensions (example A.27). The preorder corresponding to the indiscrete topology in the cards example has $R \leq S$ and $S \leq R$, which says that the dimensions are specializations of each other, i.e. effectively the same dimension (remark A.28). Thus, topology plays a significant role in our approach to semantic compositionality.

(c) Depth perception

Binocular vision can be used to infer (triangulate) location of a target object using lines of sight and relative eye positions. This computation can be achieved as an instance of sheaving, using simple geometry. Suppose the position of the target object is $(x, y) \in P$ and the angles of the eyes (lines of sight) to the target are λ and ρ for the left and right eyes, respectively. Left and right lines of sight specify position as functions of distance from the eyes, $l \in L$ and $r \in R$, parameterized by angle:

- $left_\lambda: l \mapsto l(\cos \lambda, \sin \lambda)$, and
- $right_\rho: r \mapsto r(\cos \rho, \sin \rho)$.

The position of the target is the intersection of the two lines of sight, which is the pullback of $left_\lambda$ and $right_\rho$. This pullback is equivalent to the pullback of projections $\pi_2: LP \rightarrow P$ and $\pi_2: RP \rightarrow P$, where LP is the relation $\{(l, left_\lambda(l)) | l \in L\}$ and RP is the relation $\{(r, right_\rho(r)) | r \in R\}$. Suppose the set of distance and position labels $\{L, R, P\}$ is given the topology $\{\emptyset, \{P\}, \{L, P\}, \{R, P\}, \{L, R, P\}\}$. This topology is analogous to the previous (binding) example, so it conveys similar meaning. Sheaving on this space computes the pullback of the projections (remark A.22). Hence, the location of the target object is obtained as an instance of sheaving.

4. Discussion

Semantic compositionality concerns the way that representations and the entities they stand in for correspond in some systematic, structurally consistent manner. Our sheaf theory approach regards this correspondence as data attached to a topological space (presheaf/sheaf), where the shape (topology) of the underlying space conveys meaning to the representations. Shape is determined by the open sets and its structure is preserved by restrictions of the data, either locally (presheaf), or in a systematic, globally coherent manner (sheaf). Systematicity is afforded by a universal construction (sheaving). Sheaving infers non-local information from locally sourced knowledge to construct the nearest sheaf by gluing together data that agree on the overlapping regions (global coherency). Three examples were given: (1) inferring the ranks and suits of every card, given ranks and suits of some cards, (2) inferring the binding of features to objects given the binding of features to locations and (3) inferring object location given binocular line of sight. In each case, local knowledge is extended (composed) to infer non-local information, and this form of compositionality depends on the topology.

Note that there are two senses in which sheaving spans a formal divide. There is a ‘vertical’ sense in that presheaves are maps that preserve spatial relations (inclusions) as algebraic relations (restrictions). We limited ourselves to the simplest case where attached data were sets. In general, other categories can be used, such as categories of partially ordered sets, or groups. And there is a ‘horizontal’ sense in that data attached to open sets are glued together to construct data attached to a larger open set. These two senses arise because functors are maps between categories, whereas natural transformations (sheavings) are maps between functors.

This sheaf theory approach can be compared/contrasted with classical approaches to compositionality. Classical compositionality, in comparison, says that representations of complex entities are given by representations of their constituent entities so that the semantic relations between constituents are preserved by syntactic relations between corresponding symbolic representations. Functors preserve structure. So, classical and categorical approaches are similar to the extent that classical structures are category-like. Classical theory assumes symbolic representations are instantiated on some physical system, e.g. memory registers (or, slots), hence classical systems are sometimes called physical symbol systems [14]. Given a set of registers, one can impose the discrete topological space, in which the instantiated symbols are data attached to that space, thus realizing a presheaf. In this way, classical compositionality can be seen as an instance of categorical compositionality. By contrast, however, functoriality is only one part of the categorical approach to compositionality presented here. Presheaves and sheaves are functors, but only presheaves that are sheaves satisfy the global coherency conditions.

As noted elsewhere [17], pullbacks are reminiscent of symbolic connectionist models, LISA [21] and DORA [22]. The idea is that (relational) entities are represented via connections to corresponding neurons representing the constituent entities (fillers) and their roles in the relation based on shared semantic information represented by a common pool of neurons. Neurons representing related entities that have shared semantic features tend to bind together. Similarly, the pullback of morphisms $f: A \rightarrow C$ and $g: B \rightarrow C$ is a generalized intersection

of A and B constrained by C . In terms of those models, objects A and B pertain to roles and fillers, C to semantic features, and the pullback object to relational binding. This correspondence is suggestive of a way to connect sheaving to neural network models. Neurons are topologically organized and their activities are the attached data.

The nature of sheaves depends on the nature of the data and the underlying topology. The examples of sheaves presented here are relatively simple. Sheaf theory has applications in other areas that may be adaptable to cognition. For example, a sheaf theory approach to sensor fusion [23] suggests applications to the psychology of perception. Human probability judgments that violate classical probability laws motivate quantum probability theory for cognition [24]. The close connection between sheaf theory and contextuality effects in quantum physics [25] suggests that our sheaving approach to semantic compositionality may also be applicable to quantum-like compositionality effects [17]. In these applications, the data are measurements, or probabilities [23,25].

One important direction for further work is modelling the development of the underlying topological space. Our examples illustrate how different topologies ground relational information differently. However, we have not considered how these topological spaces are obtained. Sheaf theory methods in applied topology [26] may be useful here, where the underlying topological space is inferred from data.

The importance of the underlying topology is another way that the sheaving approach goes beyond classical and artificial neural network approaches to compositionality. In this paper, we focused on the universal morphism aspect of sheaves and sheaving, because universal morphisms were argued to play a crucial role in explaining systematicity [9,11], which is a cognitive property motivating compositionality principles [8]. Yet, the topological aspect of sheaving is also crucial. Any set of registers or neurons can be given a topology. The deeper question is *why* one topology arises over another. Discrete and indiscrete topologies were asserted for an application of sheaving [17] because they are two extremes obtained from universal morphisms. So, their determination accords with the general universal construction principle [9,11]. Determination of other topologies will depend on other constraints. For instance, the physical (geometrical) relations between sensors ground triangulation of object location. This view of semantics differs from the classical view, which regards the computational (psychological) level as supported by, but independent of the specific physical (implementational) level—just as a programming language is supported by, but independent of a specific computer.

Topology captures order, and order is implicit even in the productive (recursive) aspects of cognition, e.g. level within a tree hierarchy. We have not dealt with productivity, as it purportedly implies recursion in language [27]. Category theory also provides general constructions for recursion [28], and these methods have been applied to some aspects of cognition [9]. Topology is not regarded as the only source of semantic information. So, in this sense, category (sheaf) theory provides a general framework for semantic compositionality.

Data accessibility. This article does not contain any additional data.

Competing interests. I declare I have no competing interests.

Funding. This work was supported by a Japanese Society for the Promotion of Science Grant-in-aid [16KT0025].

Appendix A. Basic theory

Conceptual introductions to the formal concepts provided in this appendix can be found in [23,29,30], see also in [17]. Deeper introductions to the category theory concepts can be found in [13,16,19] and sheaf theory concepts in [16,18]. Specific results are referenced where they appear in the appendix.

Definition A.1 (Category). A category \mathbf{C} consists of a collection of *objects*, $\mathcal{O}(\mathbf{C}) = \{A, B, \dots\}$, a collection of *morphisms*, $\mathcal{M}(\mathbf{C}) = \{f, g, \dots\}$ —a morphism written in full as $f: A \rightarrow B$ indicates object A as the *domain* and object B as the *codomain* of f —including for each object $A \in \mathcal{O}(\mathbf{C})$ the *identity morphism* $1_A: A \rightarrow A$, and a *composition* operation, \circ , that sends each pair of *compatible* morphisms $f: A \rightarrow B$ and $g: B \rightarrow C$ (i.e. the codomain of f is the domain of g) to the *composite* morphism $g \circ f: A \rightarrow C$, that together satisfy the laws of:

- *identity*: $f \circ 1_A = f = 1_B \circ f$ for every $f \in \mathcal{M}(\mathbf{C})$, and
- *associativity*: $h \circ (g \circ f) = (h \circ g) \circ f$ for every triple of compatible morphisms $f, g, h \in \mathcal{M}(\mathbf{C})$.

Example A.2 (Set). The collection of sets (objects) and the collection of functions between sets (morphisms) form a category, denoted **Set**. The composition operation is function composition and the identity morphisms are the identity functions.

Remark A.3. Morphisms may be transformations, or relations between objects that may have their own internal structure, in which case morphisms typically preserve that structure. A category may have zero or more morphisms from an object A to an object B . In **Set**, the number of functions $A \rightarrow B$ is n^m , where m (n) is the size of A (B). Composition need not return a ‘new’ morphism: e.g. $f \circ 1_A = f$ (composition with an identity); $f \circ f = f$ (self-composition of an *idempotent* function).

Definition A.4 (Topological Space). A *topological space* (X, T) consists of a set X and a collection T of subsets U of X , called the *open sets* of T , such that the empty set (\emptyset) and X are open sets, and arbitrary unions and finite intersections of open sets are open sets. The set T is called the *topology* of X .


Example A.5 (Topological space). A topological space is a category of open sets (objects) and inclusions (morphisms)—there is just one morphism $V \rightarrow U$ if $V \subseteq U$. Composition is by transitivity: $V \subseteq U \wedge W \subseteq V \Rightarrow W \subseteq U$. The discrete topology on X is the set of all subsets of X ; the indiscrete topology on X is $\{\emptyset, X\}$.

Definition A.6 (Product). A *product* of objects A and B , in a category \mathbf{C} , is an object P (also written $A \times B$) together with a pair of morphisms $\pi_1: P \rightarrow A$ and $\pi_2: P \rightarrow B$ such that for every object Z and morphisms $f: Z \rightarrow A$ and $g: Z \rightarrow B$ there exists a unique morphism $u: Z \rightarrow P$ such that $f = \pi_1 \circ u$ and $g = \pi_2 \circ u$. Morphism u is also denoted $\langle f, g \rangle$, as it is uniquely given by f and g .

Example A.7 (Product). A product of A and B , in **Set**, is the Cartesian product: the set $A \times B = \{(a, b) \mid a \in A, b \in B\}$ and projections $\pi_1: (a, b) \mapsto a$ and $\pi_2: (a, b) \mapsto b$.

Remark A.8. The function $u: Z \rightarrow P$ need not be a one-to-one correspondence (bijection). For instance, the rules of a game may stipulate that certain cards are duplicated or withheld, so a deck may contain more or less than 52 cards, i.e. the map from faces to cards, $card: Face \rightarrow Card$, is onto (surjection) or into (injection).

Definition A.9 (Pullback). A *pullback* of morphisms $f: A \rightarrow C$ and $g: B \rightarrow C$, in a category \mathbf{C} , is an object P (also written $A \times_C B$) together with a pair of morphisms $\pi_1: P \rightarrow A$ and $\pi_2: P \rightarrow B$ such that for every object Z and morphisms $z_1: Z \rightarrow A$ and $z_2: Z \rightarrow B$ there exists a unique morphism $u: Z \rightarrow P$ such that diagram



commutes: paths (where one path has at least two arrows) with the same start/end object are equal, e.g. $f \circ \pi_1 = g \circ \pi_2$. The dashed arrow indicates uniqueness.

Example A.10 (Pullback). A pullback of functions $f: A \rightarrow C$ and $g: B \rightarrow C$, in \mathbf{Set} , is the subset of pairs of elements $\{(a, b) \in A \times B \mid f(a) = g(b)\}$ and functions π_1, π_2 .

Remark A.11. A pullback is a (generalized) product constrained by f and g . A product of A and B is equivalently a pullback of $f: A \rightarrow 1$ and $g: B \rightarrow 1$, where 1 is *terminal*: an object such that for every object X , in \mathbf{C} , there exists a unique morphism from X to 1 . In \mathbf{Set} , a terminal is any singleton set, hence $f(a) = g(b)$ for all $a \in A$ and $b \in B$. Thus, a product is effectively an ‘unconstrained’ pullback.

Definition A.12 (Functor). A *functor* is a ‘structure-preserving’ map from a category \mathbf{C} to a category \mathbf{D} , written $F: \mathbf{C} \rightarrow \mathbf{D}$, sending each object A and morphism $f: A \rightarrow B$ in \mathbf{C} to the object $F(A)$ and the morphism $F(f): F(A) \rightarrow F(B)$ in \mathbf{D} (respectively) that satisfies the laws of:

- *identity*: $F(1_A) = 1_{F(A)}$ for every object $A \in \mathcal{O}(\mathbf{C})$, and
- *compositionality*: $F(g \circ_C f) = F(g) \circ_D F(f)$ for every pair of compatible morphisms $f, g \in \mathcal{M}(\mathbf{C})$.

Remark A.13. A functor written $F: \mathbf{C}^{\text{op}} \rightarrow \mathbf{D}$ indicates that the directions of the morphisms in \mathbf{C} are reversed.

Example A.14 (Diagonal, product). The *diagonal functor* $\Delta: A \mapsto (A, A), f \mapsto (f, f)$ sends each object and each morphism to their pairs. The *product functor* $\Pi: (A, B) \mapsto A \times B, (f, g) \mapsto f \times g$ sends pairs of objects/morphisms to their products.

Definition A.15 (Natural transformation). Let $F, G: \mathbf{C} \rightarrow \mathbf{D}$ be functors. A *natural transformation* $\eta: F \rightarrow G$ is a family of \mathbf{D} -morphisms $\{\eta_A: F(A) \rightarrow G(A) \mid A \in \mathcal{O}(\mathbf{C})\}$ such that $G(f) \circ \eta_A = \eta_B \circ F(f)$ for every morphism $f: A \rightarrow B$ in \mathbf{C} .

Example A.16 (Projection). Projection, $\pi = \{\pi_1, \pi_2\}$, is a natural transformation.

Definition A.17 (Universal morphism). A *universal morphism* from functor $F: \mathbf{C} \rightarrow \mathbf{D}$ to object Y in \mathbf{D} is a pair (B, ψ) consisting of an object B in \mathbf{C} and a morphism $\psi: F(B) \rightarrow Y$ in \mathbf{D} such that for every object X in \mathbf{C} and every morphism $g: F(X) \rightarrow Y$ in \mathbf{D} there exists a unique morphism $u: X \rightarrow B$ in \mathbf{C} such that $g = \psi \circ F(u)$.

Example A.18 (Products, pullbacks). A product of A and B is a universal morphism $(A \times B, \pi)$ from the diagonal functor, Δ , to the pair of objects (A, B) , where $\pi = (\pi_1, \pi_2)$. A pullback of morphisms $f: A \rightarrow C$ and $g: B \rightarrow C$ is a universal morphism $(A \times_C B, \pi)$ from the (generalized) diagonal functor [10] to the pair of morphisms (f, g) .

Remark A.19. Universal morphisms (e.g. products/pullbacks) are *unique up to unique isomorphism*: e.g. $(B \times A, \pi')$, where $\pi' = (\pi_2, \pi_1)$ is also a product of A and B , but there is just one isomorphism $(A \times B, \pi) \cong (B \times A, \pi')$ making the associated diagram commute. An isomorphism $A \times B \cong B \times A$ is generally not unique. ($f: X \rightarrow Y$ is an isomorphism if it has a *left/right inverse*, $g: Y \rightarrow X$, i.e. $f \circ g = 1_Y$ and $g \circ f = 1_X$.)


Definition A.20 (Presheaf). Let (X, T) be a topological space. A *presheaf* is a functor $\mathcal{F}: T^{\text{op}} \rightarrow \mathbf{Set}$ such that for each open set U in T there is a set $\mathcal{F}(U)$ of elements, called the *sections* over U , and for each inclusion $V \subseteq U$ in T there is a morphism, called a *restriction morphism*, $f|_V: \mathcal{F}(U) \rightarrow \mathcal{F}(V)$ that satisfies:

- *identity*: for each open set U in T , the restriction morphism $f|_U: \mathcal{F}(U) \rightarrow \mathcal{F}(U)$ is the identity morphism $1_{\mathcal{F}(U)}$, and
- *compositionality*: for each triple of open sets U, V, W in T , if $W \subseteq V \subseteq U$, then $g|_W \circ f|_V = (g \circ f)|_W$.

Definition A.21 (Sheaf). A sheaf is a presheaf $\mathcal{F}: T^{\text{op}} \rightarrow \mathbf{Set}$ that satisfies:

- *gluing* (existence): if $\{U_i\}_{i \in I}$ is an open cover of an open set $U \in T$, and if for each $i \in I$ a section $s_i \in \mathcal{F}(U_i)$ is given such that for each pair (U_i, U_j) of covering sets $s_i|_{U_i \cap U_j} = s_j|_{U_i \cap U_j}$, i.e. s_i and s_j agree on overlap, then there is a section $s \in \mathcal{F}(U)$ such that $s|_{U_i} = s_i$ for each i , and
- *locality* (uniqueness): if $\{U_i\}_{i \in I}$ is an open cover of an open set $U \in T$, and if $s, t \in \mathcal{F}(U)$ such that $s|_{U_i} = t|_{U_i}$ for each U_i , then $s = t$.

Remark A.22. The sheaf conditions are indicated by commutative diagram



where the right diamond indicates the pullback of $f|_{U \cap V}$ and $g|_{U \cap V}$; or equivalently, by the *equalizer* of pairs of morphisms $\prod_i \mathcal{F}(U_i) \rightrightarrows \prod_{i,j} \mathcal{F}(U_i \cap U_j)$, where the products iterate over the open sets [16].

(a)
*

(b)
c ₂
c ₁
C

(c)	
a ₂	c ₁
a ₁	c ₁
A	C

(d)	
c ₂	b ₂
c ₁	b ₁
C	B

(e)		
a ₁	c ₁	b ₁
A	C	B

(f)		
a ₂	c ₁	b ₁
a ₁	c ₁	b ₁
A	C	B

Figure 1. Relational tables for a presheaf (a–e) and its nearest sheaf (a–d, f).

Remark A.23. *Sheaving* [18] constructs the ‘nearest’ (smallest) sheaf \mathcal{F}^+ from presheaf \mathcal{F} , i.e. the universal morphism (\mathcal{F}^+, θ) , where $\theta: \mathcal{F} \rightarrow \mathcal{F}^+$.

Example A.24 (Natural join). A perspicuous illustration of sheaving pertains to the *natural join* of relational tables [25], which involves joining tables at their shared attribute values. Attributes (headings) constitute open sets of the topological space and rows constitute the sections attached to the open sets. Restrictions are just projections along the specified attributes. Relational tables illustrate sheaving (figure 1) where table headings correspond to the open sets of the topology $\{\emptyset, \{C\}, \{A, C\}, \{B, C\}, \{A, B, C\}\}$. Tables (a–e) constitute a presheaf, where (a) is the empty table. The natural join of tables (c) and (d) constructs table (f), which constitutes the sheaf, i.e. tables (a–d, f).

Remark A.25. For the empty set, $\mathcal{F}: \emptyset \mapsto 1$, i.e. the terminal for **Set**, written $\{*\}$, as the element’s name is unimportant. Restriction $\mathcal{F}(U) \rightarrow 1$, for inclusion $\emptyset \subseteq U$, is guaranteed to exist, since 1 is terminal. The attached data is an empty function.

Example A.26 (Two cards). The *queen of hearts* and *two of spades* are represented by the presheaf \mathcal{F}_{Q2} , which is expressed as a collection of relational tables (figure 2).

Example A.27 (Deck of cards). Knowing each rank and suit corresponds to a presheaf, $\mathcal{F}_{\text{card}}$, that sends the open sets $\{R\}$ and $\{S\}$ to the sets of all ranks and all suits, respectively. Sheaving computes the pullback of restrictions $\mathcal{F}_{\text{card}}(\{R\}) \rightarrow 1$ and $\mathcal{F}_{\text{card}}(\{S\}) \rightarrow 1$, since $\{R\} \cap \{S\} = \emptyset$, which

(a)
*

(b)
Q
2
R

(c)	
Q	♥
2	♠
R	S

(d)
♥
♠
S

Figure 2. Relational tables for the presheaf \mathcal{F}_{Q2} .

is equivalent to the product $\mathcal{F}_{\text{card}}(\{R\}) \times \mathcal{F}_{\text{card}}(\{S\})$. Thus, sheaving constructs the sheaf $\mathcal{F}_{\text{card}}^+$ representing the deck of all 52 cards.

Remark A.28. A topological space, (X, T) , induces a specialization preorder on the elements of the underlying set, X . Two elements $x, y \in X$ are comparable, $x \leq y$, if x is an element of the *closure* of y , i.e. the intersection of all *closed sets* containing y —if U is an open set of T , then the complement of U (i.e. the set of elements in X that are not in U) is a closed set. In the cards example, the indiscrete topology has closed sets \emptyset and $\{R, S\}$. The closure of R and the closure of S are the same set, $\{R, S\}$. Hence, the preorder has $R \leq S$ and $S \leq R$. Open sets specify closeness. Accordingly, the open set $\{R, S\}$ says that R and S are close to each other, but not preferentially so, since there are no other open sets. The open sets of a discrete topology are also the closed sets. So, in the discrete case, R and S are not comparable, since R is not in the closure of S , i.e. $\{S\}$, and S is not in the closure of R , i.e. $\{R\}$. Note that an element is always comparable to itself, $x \leq x$, because any topology T on X must contain X as an open set of T (by definition).

References

- Chomsky N. 1980 *Rules and representations*. New York, NY: Columbia University Press.
- Fodor JA. 1975 *The language of thought*. New York, NY: Crowell.
- McLaughlin BP. 2009 Systematicity redux. *Synthese* **170**, 251–274. (doi:10.1007/s11229-009-9582-0)
- Fodor JA, Pylyshyn ZW. 1988 Connectionism and cognitive architecture: a critical analysis. *Cognition* **28**, 3–71. (doi:10.1016/0010-0277(88)90031-5)
- Smolensky P. 1990 Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artif. Intell.* **46**, 159–216. (doi:10.1016/0004-3702(90)90007-M)
- Aizawa K. 2003 *The systematicity arguments*. Studies in Mind and Brain 1. New York, NY: Kluwer Academic.
- Gardenfors P. 2000 *Conceptual spaces: the geometry of thought*. New York, NY: MIT Press.
- Szabo ZG. 2017 Compositionality. In *The Stanford Encyclopedia of Philosophy* (ed. EN Zalta). Stanford University, CA: Metaphysics Research Lab. See <https://plato.stanford.edu/archives/sum2017/entries/compositionality/>.
- Phillips S, Wilson WH. 2016 Systematicity and a categorical theory of cognitive architecture: universal construction in context. *Front. Psychol.* **7**, 1139. (doi:10.3389/fpsyg.2016.01139)
- Mac Lane S. 1998 *Categories for the working mathematician*, 2nd edn. Graduate Texts in Mathematics. New York, NY: Springer.
- Phillips S, Wilson WH. 2010 Categorical compositionality: a category theory explanation for the systematicity of human cognition. *PLoS Comput. Biol.* **6**, e1000858. (doi:10.1371/journal.pcbi.1000858)
- Eilenberg S, Mac Lane S. 1945 General theory of natural equivalences. *Trans. Am. Math. Soc.* **58**, 231–294. (doi:10.1090/S0002-9947-1945-0013131-6)
- Awodey S. 2010 *Category theory*, 2nd edn. Oxford Logic Guides. New York, NY: Oxford University Press.
- Newell A. 1980 Physical symbol systems. *Cogn. Sci.* **4**, 135–183. (doi:10.1207/s15516709cog0402_2)
- Vialar T. 2016 *Handbook of mathematics*. Norderstedt, Germany: HDBoM.
- Mac Lane S, Moerdijk I. 1992 *Sheaves in geometry and logic: a first introduction to topos theory*. New York, NY: Springer.
- Phillips S. 2018 Going beyond the data as the patching (sheaving) of local knowledge. *Front. Psychol.* **9**, 1926. (doi:10.3389/fpsyg.2018.01926)
- Hartshorne R. 1977 *Algebraic geometry*, vol. 52. Graduate Texts in Mathematics. New York, NY: Springer.

19. Leinster T. 2014 *Basic category theory*, vol. 143. Cambridge Studies in Advanced Mathematics. Cambridge, UK: Cambridge University Press.
20. Wolfe JM, Cave KR, Franzel SL. 1989 Guided search: an alternative to the feature integration model for visual search. *J. Exp. Psychol.: Human Percept. Perform.* **15**, 419–433. (doi:10.1037/0096-1523.15.3.419)
21. Hummel JE, Holyoak KJ. 1997 Distributed representations of structure: a theory of analogical access and mapping. *Psychol. Rev.* **104**, 427–466. (doi:10.1037/0033-295X.104.3.427)
22. Dumas LAA, Hummel JE, Sandhofer CM. 2008 A theory of the discovery and predication of relational concepts. *Psychol. Rev.* **115**, 1–43. (doi:10.1037/0033-295X.115.1.1)
23. Robinson M. 2017 Sheaves are the canonical data structure for sensor integration. *Inf. Fusion* **36**, 208–224. (doi:10.1016/j.inffus.2016.12.002)
24. Busemeyer JR, Bruza PD. 2012 *Quantum models of cognition and decision*. Cambridge, UK: Cambridge University Press.
25. Abramsky S, Brandenburger A. 2011 The sheaf-theoretic structure of non-locality and contextuality. *New J. Phys.* **13**, 113036. (doi:10.1088/1367-2630/13/11/113036)
26. Ghrist R. 2014 *Elementary applied topology*. Philadelphia, PA: Robert Ghrist.
27. Bolhuis JJ, Beckers GJL, Huybregts MAC, Berwick RC, Everaert MBH. 2018 Meaningful syntactic structure in songbird vocalizations? *PLoS Biol.* **16**, e2005157. (doi:10.1371/journal.pbio.2005157)
28. Hinze R, Wu N. 2016 Unifying structured recursion schemes: an extended study. *J. Funct. Program.* **26**, e1. (doi:10.1017/S0956796815000258)
29. Lawvere FW, Schanuel SH. 1997 *Conceptual mathematics: a first introduction to categories*. Foundations of Computing. Cambridge, UK: Cambridge University Press.
30. Simmons H. 2011 *An introduction to category theory*. New York, NY: Cambridge University Press.