The Yoneda Lemma without category theory: algebra and applications

Vaughan Pratt
Stanford University

August 9, 2009

1 Algebraic formulation

1.1 Introduction

The Yoneda Lemma is ordinarily understood as a fundamental representation theorem of category theory. As such it can be stated as follows in terms of an object c of a locally small category $\mathcal{C}$, meaning one having a homfunctor $\mathcal{C}(-,-) : \mathcal{C}^{\mathrm{op}} \times \mathcal{C} \to \mathsf{Set}$ (i.e. small homsets), and a functor $F : \mathcal{C} \to \mathsf{Set}$ or presheaf.

Lemma 1 (Yoneda)**.** *The function* $\alpha : \mathrm{Hom}(C(c,-), F) \to F(c)$ *defined by* $\alpha(\tau : \mathcal{C}(c,-) \to F) = \tau_c(1_c)$ *is is a bijection natural in* c.

However it can just as well be considered a fundamental representation theorem of universal algebra, via the connection between the homfunctor of a category and free algebras for the theory represented by that category. This connection is not generally appreciated outside category theoretic circles, which this section endeavors to correct by presenting the relevant concepts from an algebraic perspective. The algebraically motivated notations $\mathcal{T}$, s, $\mathcal{A}$, h, h^*, and a^* below correspond to the respective notations $\mathcal{C}$, c, F, τ, $\alpha(\tau)$, and $\alpha^{-1}(a)$ (for $a \in F(c)$) above, with the free algebra $\mathcal{T}_s$ corresponding to the representable functor $\mathcal{C}(c,-) : \mathcal{C} \to \mathsf{Set}$, $f_u = T_{uf}$ to $\mathcal{C}(c,f)$, and $f^* = \mathcal{T}_f$ to $\mathcal{C}(f,-)$.

The remaining sections of the paper apply the Yoneda Lemma to a proposed simplification and generalization of algebra, illustrated with applications. Algebras and their homomorphisms have in effect been defined as respectively functors and natural transformations long before category theory came into existence. That is, we do not see any essential differences between the definitions, and category theorists are just as entitled to characterize algebra as merely another language for category theory as algebraists are to the converse claim.

Sections 2 and 3 on respectively dense and didense extensions of categories exploit the Yoneda Lemma to replace these conventional definitions of algebra/functor and homomorphism/natural transformation with definitions that are more abstract, primitive, and general. To the extent that algebras are fundamental to mathematics this makes the Yoneda Lemma a potentially insightful tool for the foundations of mathematics.

1.2 Homogeneous case

Clearly every element x of a set X is representable as the function $x^* : \{0\} \to X$ uniquely defined by $x^*(0) = x$. This is the discrete case of the representation of an element a of an algebraic structure $\mathcal{A}$ as the homomorphism $a^* : \mathcal{T}_1 \to \mathcal{A}$ uniquely defined by $a^*(\mathbf{x}) = a$ where $\mathcal{T}_1$ denotes the free algebra on one generator $\mathbf{x}$. For any homomorphism $h : \mathcal{T}_1 \to \mathcal{A}$ write h^* for the element $h(\mathbf{x})$ of $\mathcal{A}$. We then have $a^{**} = a$ and $h^{**} = h$, showing that this representation takes the form of a bijection between the underlying set A of $\mathcal{A}$ and the set $\mathrm{Hom}(\mathcal{T}_1, \mathcal{A})$ of homomorphisms from $\mathcal{T}_1$ to $\mathcal{A}$. Its existence is either an easily seen consequence of the definition of $\mathcal{T}_1$ as an algebra of unary terms $f(\mathbf{x})$ whose values are uniquely determined by the value of their common variable $\mathbf{x}$, or (part of) a definition in its own right of the notion of free algebra.

This representation is applicable to $\mathcal{T}$-algebras or models of an algebraic theory $\mathcal{T}$ defined by equations between terms formed from variables and operation symbols. Typical such structures include monoids, groups, vector spaces, Boolean algebras, lattices, and pointed sets.

For each of these the respective free $\mathcal{T}$-algebra $\mathcal{T}_1$ on one generator is, up to isomorphism, the monoid of natural numbers, the group of integers, the one-dimensional vector space over a given field, the four-element Boolean algebra $\{\mathbf{x}, \neg\mathbf{x}, 0, 1\}$, the singleton lattice $\{\mathbf{x}\}$, and the doubleton pointed set $\{\mathbf{x}, c\}$ (set with one constant c, as a constant unary operation). In all cases $\mathcal{T}_1$ can be understood as the unary abstract terms of the theory, obtainable from those concrete terms that are expressed with a single variable $\mathbf{x}$, modulo the congruence on terms (with respect to composition) generated by the equations of $\mathcal{T}$.

This representation of the elements of $\mathcal{A}$ extends to the unary operations f_A of $\mathcal{A}$ as follows. (This is the algebraic counterpart of naturality in c in the categorical formulation.) Taking $\mathcal{A}$ to be $\mathcal{T}_1$ in the first paragraph of this subsection yields $f^* : \mathcal{T}_1 \to \mathcal{T}_1$ as the homomorphism representing the element f, or $f(\mathbf{x})$, of $\mathcal{T}_1$. This permits the interpretation f_A of f in $\mathcal{A}$ to be recovered entirely from how homomorphisms compose as follows.

$$
\begin{aligned}
f_A(a) &= f_A(a^*(\mathbf{x})) && \text{(defn. of }^*\text{)} \\
&= a^*(f(\mathbf{x})) && (a^* \text{ is a homomorphism}) \\
&= a^*(f) && (f(\mathbf{x}) = f \in \mathcal{T}_1) \\
&= a^*(f^*(\mathbf{x})) && \text{(defn. of }^*\text{)} \\
&= (a^* f^*)(\mathbf{x}) && \text{(homomorphisms are composable)} \\
&= (a^* f^*)^* && \text{(defn. of }^*\text{)}
\end{aligned}
$$

That is, the left action of f_A on elements $a \in A$, as the element $f_A(a)$ of A, is represented by homomorphisms as the right action of f^* on a^*, or substitution of the element $f(\mathbf{x})$ of $\mathcal{T}_1$ for the argument of a^*, or in pictures the composite $\mathcal{T}_1 \xrightarrow{f^*} \mathcal{T}_1 \xrightarrow{a^*} \mathcal{A}$.

The beauty of this representation is that f^* is defined independently of $\mathcal{A}$: the same f^* can be used for every algebra $\mathcal{A}$, with the action of f on $a \in A$ determined by the choice of composite a^*f^* at $\mathcal{T}_1$, back at the depot so to speak.

This representation can be organized more crisply as an isomorphism of (left) modules. A module $(\mathcal{M}, X, \cdot)$ consists of a monoid $\mathcal{M} = (M, \circ, 1)$, a set X, and a scalar multiplication $\cdot : M \times X \to X$ satisfying $(m \circ n) \cdot x = m \cdot (n \cdot x)$ and $1 \cdot x = x$. The example above of the scalar multiplications of an arbitrary vector space nicely illustrates the notion of module, *but so do the other five examples.* For example any Boolean algebra $\mathcal{B}$ determines a module by taking X to be its underlying set B, $\mathcal{M}$ to be the monoid formed by the four unary Boolean terms under composition, and scalar multiplication to be application of the interpretation of those terms as the four unary operations of $\mathcal{B}$. Every $\mathcal{T}$-algebra $\mathcal{A} = (A, f_A, g_A, \ldots)$ is an expansion by non-unary operations of the module $(\mathcal{T}_1, A, \cdot)$ where $\mathcal{T}_1$ is the monoid of unary terms of $\mathcal{T}$ and $\cdot$ is application of these terms interpreted as operations of $\mathcal{A}$; call this the ***internal module*** of $\mathcal{A}$ formed by its unary operations. When all operations of $\mathcal{T}$ are unary a $\mathcal{T}$-algebra *is* in this way a module.

But every $\mathcal{T}$-algebra $\mathcal{A}$ determines a second module $(\mathrm{End}(\mathcal{T}_1)^{\mathrm{op}}, \mathrm{Hom}(\mathcal{T}_1, \mathcal{A}), ;)$ where ; is converse composition (so $h; k = k \circ h = kh$), $\mathrm{End}(\mathcal{T}_1)^{\mathrm{op}}$ is the monoid of endomorphisms of $\mathcal{T}_1$ under converse composition, and $\mathrm{Hom}(\mathcal{T}_1, \mathcal{A}) = \{h : \mathcal{T}_1 \to \mathcal{A}\}$ as before. Call this the ***external module*** of $\mathcal{A}$ formed by its representing homomorphisms.

Lemma 2 (Yoneda, algebraic form)**.** *The internal and external modules of an algebraic structure are isomorphic.*

An alternative and convenient formulation that we shall sometimes use defines the external module as a right module $(\mathrm{End}(\mathcal{T}_1), \mathrm{Hom}(\mathcal{T}_1, \mathcal{A}), \circ)$ in terms of composition instead of its converse, distinguished from ordinary or left modules by writing the scalar multiplication conversely, as with $a^* \circ f^*$ or a^*f^*, so that the monoid now acts on the set from the right. The lemma then states that the external right module of $\mathcal{A}$ is dual to its internal left module.

In the case $\mathcal{A} = \mathcal{T}_1$ of the Yoneda Lemma, each element $f(\mathbf{x})$ of $\mathcal{T}_1$ is represented as the homomorphism $f^* : \mathcal{T}_1 \to \mathcal{T}_1$, a substitution mapping each element $g(\mathbf{x}) \in \mathcal{T}_1$ to $f^*(g(\mathbf{x})) = g(f(\mathbf{x}))$. That is, the monoid $\mathrm{End}(\mathcal{T}_1)$ of endomorphisms of $\mathcal{T}_1$ is isomorphic to ${\mathcal{T}_1}^{\mathrm{op}}$ as the opposite of the monoid of unary operations of $\mathcal{T}$.

If $\mathcal{T}$ is a commutative theory such as that of vector spaces, meaning that ${\mathcal{T}_1}^{\mathrm{op}}$ is isomorphic to $\mathcal{T}_1$ via $g(f(\mathbf{x})) = f(g(\mathbf{x}))$, nothing much changes in the passage from inside to outside. More striking is the case of a noncommutative theory such as that of Boolean algebras. For example the constant unary operation $0 \wedge \mathbf{x}$, as the element 0 of the free Boolean algebra $\mathcal{T}_1$ on $\mathbf{x}$, is represented by the endomorphism 0^* given by $0^*(f(\mathbf{x})) = f(0)$, which can be seen to be the retraction of $\mathcal{T}_1$ to its subboolean algebra $\{0, 1\}$ sending $\mathbf{x}$ to 0 and $\neg\mathbf{x}$ to 1.

This case of the Yoneda Lemma constitutes the Yoneda Embedding (thus far in the single-sorted case) because it embeds the monoid ${\mathcal{T}_1}^{\mathrm{op}}$ as a full one-object subcategory of the concrete category of $\mathcal{T}$-algebras and their homomorphisms.

1.3 Heterogeneous case

One generalization of this representation of the internal module of $\mathcal{A}$ by its external one would be to extend the notion of internal module to the n-ary operations of $\mathcal{A}$ for all n. The requisite bijection for doing so is that between the set A^n of n-tuples of elements of $\mathcal{A}$ and the set $\mathrm{Hom}(\mathcal{T}_n, \mathcal{A})$ of homomorphisms to $\mathcal{A}$ from the free algebra $\mathcal{T}_n$ on n generators $\mathbf{x}_1, \ldots, \mathbf{x}_n$. (The categorical counterpart of such algebras takes the form of functors that preserve products.)

The Yoneda Lemma treats a related generalization, namely to multisorted or heterogeneous unary algebras, or presheaves. These are the models $\mathcal{A}$ of a theory $\mathcal{T}$ having sorts $s, t, \ldots$, terms $f : s \to t$, and equations between terms of the same type $s \to t$ serving as axioms, in other words a category $\mathcal{T}$. The elements of $\mathcal{A}$ of sort s form a set A_s and each term $f : s \to t$ is interpreted in $\mathcal{A}$ as a function $f_A : A_s \to A_t$ satisfying the equations of the theory, in other words a functor $\mathcal{A} : \mathcal{T} \to \mathsf{Set}$.

As explained in subsection 1.5, the models of an algebraic theory form a full subcategory of its presheaf models, the latter being those for which A_{s^n} is not required to be A_s^n. In this sense the latter generalization subsumes the former, with the caveat that structural properties of the latter such as forming a topos [6] need not be preserved by subcategories. We can therefore assume henceforth, without loss of generality for algebraic theories (for our purposes), that we are working with the whole of $\mathcal{T}$. The single-sorted case was more constrained, obliging us to consider only the unary part $\mathcal{T}_1$ of a possibly larger theory $\mathcal{T}$. That we could reliably illustrate presheaves with Boolean algebras etc. is a consequence of the category of Boolean algebras forming a full subcategory of the category of all presheaves on the theory of Boolean algebras, with the Boolean homomorphisms $\mathcal{T}_1 \to \mathcal{A}$ in particular being the same thing as the presheaf homomorphisms.

In place of a single formal variable $\mathbf{x}$ serving as the generator of $\mathcal{T}_1$ we now need a separate variable $\mathbf{x}_s$ for each sort s. Each such variable generates its own free algebra $\mathcal{T}_s$ consisting of all abstract terms or polynomials of the form $f(\mathbf{x}_s)$, that is, all terms generated by $\mathbf{x}_s$. Following the notation A_t, T_{st} denotes the set of elements of $\mathcal{T}_s$ of sort t, namely the terms $f : s \to t$.

In the single-sorted case, terms lead a double life as the elements and the operations of $\mathcal{T}_1$. In the multisorted case, each term $f : s \to t$ exists as an *element* $f(\mathbf{x}_s) \in \mathcal{T}_s$ in only *one* free algebra, but it is interpreted as an *operation* in *every* free algebra $\mathcal{T}_u$, namely as the function $f_u : T_{us} \to T_{ut}$ mapping each term $g(\mathbf{x}_u)$ of sort s in $\mathcal{T}_u$ to the term $f(g(\mathbf{x}_u))$ of sort t. This "one-every" distinction did not arise in the single-sorted case.

Every element $a \in A_s$ is representable as the homomorphism $a^* : \mathcal{T}_s \to \mathcal{A}$ uniquely defined by $a^*(\mathbf{x}_s) = a$, by essentially the same reasoning as for the single-sorted case: evaluate the terms in $\mathcal{T}_s$. Applying this representation as for the single-sorted case, every element $f : s \to t$ in $\mathcal{T}_s$ is representable as the homomorphism $f^* : \mathcal{T}_t \to \mathcal{T}_s$ by taking $\mathcal{A} = \mathcal{T}_s$ and noting that f is of sort t.

We can understand f^* as the substitution of $f(\mathbf{x}_s)$ for $\mathbf{x}_t$ in each term $g(\mathbf{x}_t)$ of $\mathcal{T}_t$, producing $g(f(\mathbf{x}_s))$, a term of $\mathcal{T}_s$. The main difference from the single-sorted case is that this substitution acts as a homomorphism between possibly distinct free algebras. We can then represent $f_A(a)$ almost exactly as for the single-sorted case, namely as $\mathcal{T}_t \xrightarrow{f^*} \mathcal{T}_s \xrightarrow{a^*} \mathcal{A}$.

The Yoneda Lemma as we restated it algebraically in terms of modules was only for the single-sorted case. However we did not say so there because it can serve the multisorted case without changing its statement provided only that we broaden the definition of module to accommodate multisorted monoids. These are most efficiently organized as abstract categories (for which algebra offers no better organization) while continuing to call functors and natural transformations respectively (unary) algebras and homomorphisms.

A (heterogeneous) module $(\mathcal{C}, \mathcal{X}, \cdot)$ consists of a category $\mathcal{C}$, a family $\mathcal{X}$ of sets X_c indexed by objects c of $\mathcal{C}$, and a family of scalar multiplications $\cdot_{cd} : \mathcal{C}(c, d) \times X_c \to X_d$ doubly indexed by objects of $\mathcal{C}$, such that for each morphism $f : c \to d$ of $\mathcal{C}$ and $x \in X_c$, $f \cdot x \in X_d$ and satisfies $(gf) \cdot x = g \cdot (f \cdot x)$ and $1_c \cdot x = x$. The internal module of a $\mathcal{T}$-algebra $\mathcal{A}$ is then $(\mathcal{T}, |\mathcal{A}|, \cdot)$, where $|\mathcal{A}|$ denotes the family of underlying sets of $\mathcal{A}$ and $f \cdot a$ is interpreted as $f(a)$ constituting the rest of $\mathcal{A}$.

The Yoneda Lemma makes the internal module of $\mathcal{A}$ dual to the right module $(\mathcal{J}, |\widehat{\mathcal{J}}(F(-), \mathcal{A})|, \circ)$ of homomorphisms where $\mathcal{J} \cong \mathcal{T}^{\mathrm{op}}$ denotes the category of free $\mathcal{T}$-algebras $\mathcal{T}_s$ or representable functors $\mathcal{T}(s, -)$ along with their homomorphisms/natural transformations. Here $F : \mathcal{J} \to \widehat{\mathcal{J}}$ fully embeds $\mathcal{J}$ in the category $\widehat{\mathcal{J}}$ of all $\mathcal{T}$-algebras and their morphisms, which for the time being we can identify with $[\mathcal{T}, \mathsf{Set}]$, and $|\widehat{\mathcal{J}}(F(-), \mathcal{A})|$ denotes the family A_j of sets of homomorphisms from $F(j)$ to $\mathcal{A}$ as j ranges over the objects of $\mathcal{J}$. $\mathcal{J}$ can be viewed as a category of prototypical $\mathcal{T}$-algebras and $\widehat{\mathcal{J}}$ as its completion to all $\mathcal{T}$-algebras, a concept developed in more detail in the next section.

Our earlier definition of module can then be understood as the case when $\mathcal{C}$ has one object. Our algebraic version of the Yoneda Lemma in terms of modules applies equally to this general notion of module allowing multiple sorts.

In relating all this back to the categorical perspective sketched in the first section, one additional correspondence is that modules $(\mathcal{C}, \mathcal{X}, \cdot)$ as multisorted unary algebras are equivalent to functors $F : \mathcal{C} \to \mathsf{Set}$, with $F(c) = X_c$ and $F(f)(x) = f \cdot x$ for $x \in X_c$.

1.4 Directed reflexive multigraphs

This representation is nicely illustrated by the example of a directed reflexive multigraph $(V, E, \sigma, \tau, \iota)$, defined to consist of sets V and E of respectively vertices and edges, operations $\sigma, \tau : E \to V$ giving respectively the source $\sigma(e)$ and target $\tau(e)$ of each edge $e \in E$, and an operation $\iota : V \to E$ for which $\iota(v)$ is the distinguished self-loop at vertex v satisfying $\sigma\iota(v) = \tau\iota(v) = v$ for all $v \in V$. (This follows the usual practice of taking the signature of a $\mathcal{T}$-algebra to be only a sufficient basis from which to generate the remaining terms by identities and composition.) Without these two equations there would be infinitely many terms such as $\mathbf{e}$, $\iota\sigma$, $\iota\sigma\iota\sigma$, etc. With them there are just seven terms, organized as $T_{VV} = \{\mathbf{v}\}$, $T_{VE} = \{\iota\}$, $T_{EE} = \{\mathbf{e}, \iota\sigma, \iota\tau\}$, and $T_{EV} = \{\sigma, \tau\}$ where the generators $\mathbf{v}$ and $\mathbf{e}$ constitute the identity operations at respectively V and E.

We then have two free graphs on respective generators $\mathbf{v}$ and $\mathbf{e}$, namely $\mathcal{T}_V = (\{\mathbf{v}\}, \{\iota\}, \sigma_V, \tau_V, \iota_V)$ and $\mathcal{T}_E = (\{\sigma, \tau\}, \{\mathbf{e}, \iota\sigma, \iota\tau\}, \sigma_E, \tau_E, \iota_E)$. The graph $\mathcal{T}_V$ consists of one vertex $\mathbf{v}$ and one edge, the self-loop ι at $\mathbf{v}$. The operations of $\mathcal{T}_V$ satisfy $\sigma_V(\iota) = \tau_V(\iota) = \mathbf{v}$ and $\iota_V(\mathbf{v}) = \iota$.

The graph $\mathcal{T}_E$ consists of two vertices σ and ι and three edges: $\mathbf{e}$ from σ to τ, and self-loops $\iota\sigma$ and $\iota\tau$ at respectively σ and τ. The operations of $\mathcal{T}_E$ satisfy $\sigma_E(\mathbf{e}) = \sigma_E(\iota\sigma) = \sigma$, $\tau_E(\mathbf{e}) = \tau_E(\iota\tau) = \tau$, $\iota_E(\sigma) = \iota\sigma$, and $\iota_E(\tau) = \iota\tau$.

The Yoneda Lemma puts the seven terms of the theory in bijection with the seven homomorphisms of a category $\mathcal{J}$ of reflexive graphs whose objects are $\mathcal{T}_V$ and $\mathcal{T}_E$. A graph can then be represented simply as an object G with morphisms to G from V and E representing respectively the vertices and edges of G. The source and target vertices of each edge e are given by the respective composites $e^*\sigma^*, e^*\tau^* : \mathcal{T}_V \to G$. The self-loop at each vertex v is given by the composite $v^*\iota^* : \mathcal{T}_E \to G$.

This correspondence between operations and homomorphisms can give additional insight into a theory. For example the reflexivity provided by the operation ι is reflected in the retractibility of $\mathcal{T}_E$ to either vertex (as a self-loop). If we form $\mathcal{T}'$ from $\mathcal{T}$ by omitting ι, $\iota\sigma$ and $\iota\tau$ (graphs), $\mathcal{T}'_E$ is then no longer retractible because it contains no self-loop to receive $\mathbf{e}$. $\mathcal{T}'$ then has only four morphisms, whence so does its opposite $\mathcal{J}'$

Category theory offers yet more insight. Composing any reflexive graph G as a functor $G : \mathcal{T} \to \mathsf{Set}$ with the evident inclusion $K : \mathcal{T}' \to \mathcal{T}$ yields the underlying graph $U(G) = GK : \mathcal{T}' \to \mathsf{Set}$ of G, having the same vertices and edges but with the distinguished self-loops of G no longer differentiated from the other self-loops. For each homomorphism $h : G \to G'$ of $\widehat{\mathcal{J}}$, $U(h) = h$, making $U : \widehat{\mathcal{J}} \to \widehat{\mathcal{J}'}$ a faithful (but not full) functor, i.e. $\widehat{\mathcal{J}}$ is a subcategory of $\widehat{\mathcal{J}'}$. The new graphs in $\widehat{\mathcal{J}'}$, those not in the image of U, are those containing a vertex with no self-loop. Graph homomorphisms preserve self-loops whence there are no homomorphisms from G to H in $\widehat{\mathcal{J}'}$ when there exist self-loops in G but not H; in contrast $\widehat{\mathcal{J}}$ has no empty homsets. Ontology recapitulates phylogeny in the respective bases: the same sentence is true with the bases $\mathcal{J}$ and $\mathcal{J}'$ in place of $\widehat{\mathcal{J}}$ and $\widehat{\mathcal{J}'}$.

1.5 Applicability to algebras and coalgebras

Homogeneous algebras with non-unary operations such as $f : s^2 \to s$, as well as coalgebras with operations such as $f : s \to s + s$, can be understood as presheaves $A : \mathcal{T} \to \mathsf{Set}$ that preserve specified products and sums, e.g. $A_{s^2} = A_s^2$, $A_{s+s} = A_s + A_s$, etc. These form a full subcategory of the category $\widehat{\mathcal{J}}$ of all presheaves on $\mathcal{J}$, argued as follows.

If $\mathcal{T}$ contains a product sort s^2 then it also contains projections $p, q : s^2 \to s$. Hence any interpretation of s and s^2 as respective sets A_s, A_{s^2} interprets p and q as functions $p_A, q_A : A_{s^2} \to A_s$. This allows A_{s^2} to be understood as a multiset of pairs (a, a') for $a, a' \in A_s$, with each such pair having κ occurrences for some cardinal $\kappa \geq 0$.

A homomorphism $h : \mathcal{A} \to \mathcal{B}$ associates to sorts s and s^2 respective functions $h_s : A_s \to B_s$ and $h_{s^2} : A_{s^2} \to B_{s^2}$ satisfying $h_s(p_A(a, a')) = p_B(h_{s^2}(a, a'))$ and $h_s(q_A(a, a')) = q_B(h_{s^2}(a, a'))$, or just $h(a, a') = (h(a), h(a'))$ without the subscripts. This holds independently of how many copies of (a, a') appear in A_{s^2}. Hence for any term f of $\mathcal{T}$, $h(f(a, a')) = f(h(a, a')) = f(h(a), h(a'))$, thereby meeting the usual requirement for homomorphisms as a consequence of including the projections in $\mathcal{T}$.

The intended models of a theory $\mathcal{T}$ with product sorts are those that respect products, namely those such that the product sorts are interpreted as sets of tuples containing $\kappa = 1$ copy of each tuple. These therefore form a subclass of all the presheaves for $\mathcal{T}$, along with all their homomorphisms by the preceding argument, that is, a full subcategory thereof.

In the homogeneous case we can write the free $\mathcal{T}$-algebra on n generators as just $\mathcal{T}_n$. The Yoneda Lemma then represents the n-tuples t of any $\mathcal{T}$-algebra $\mathcal{A}$ as homomorphisms $t^* : \mathcal{T}_n \to \mathcal{A}$ and n-ary operations $f_A : A^n \to A$ as homomorphisms $f^* : \mathcal{T}_1 \to \mathcal{T}_n$ acting on the representations t^* of n-tuples t on the right, namely as $t^* f^*$. This representation of n-tuples is sound for all presheaves $\mathcal{A}$, what is special about those $\mathcal{A}$ that preserve the designated powers s^n is that each n-tuple is represented by exactly one homomorphism; with arbitrary presheaves each n-tuple could be represented by any number of homomorphisms including none.

An analogous situation holds for coalgebraic theories with coproduct or sum types such as $s + t$, as might appear with an operation $f : s \to s + s$ (more generally $f : s \to t + u$ etc. for heterogeneous coalgebras). In this case we have inclusions $i, j : s \to s + s$ whose respective interpretations i_A, j_A in any presheaf $\mathcal{A}$ for that theory associate each element a of $\mathcal{A}_{s+s}$ to those elements mapped by either i_A or j_A to a. A presheaf is a coalgebra just when coproducts are interpreted as such, namely when every such a is associated to exactly one element of A_s, either via i_A or j_A, and the dual reasoning shows that these coalgebras form a full subcategory of the presheaf category.

While some properties of presheaf categories need not hold of their full subcategories, such as forming a topos, the notion of density treated in the next section does, being as applicable to algebra and coalgebra as it is to presheaves: every algebra or coalgebra on $\mathcal{J}$ arises as an object of some full dense extension of $\mathcal{J}$, including the category of all $\mathcal{T}$-algebras, or of all $\mathcal{T}$-coalgebras, whether or not $\mathcal{T} \cong \mathcal{J}^{\mathrm{op}}$ has non-unary operations other than projections and inclusions.

2 Density

2.1 Homomorphisms

The Yoneda Lemma as stated has nothing to say about the homomorphisms between $\mathcal{T}$-algebras, or natural transformations between the presheaves, other than those from objects of $\mathcal{J} \cong \mathcal{T}^{\mathrm{op}}$. A minor complication here that does not arise with the Yoneda Lemma as stated is that there may be too many homomorphisms between two algebras to form a (small) set (in category-theoretic language, the functor category $\mathsf{Set}^{\mathcal{T}}$ or $[\mathcal{T}, \mathsf{Set}]$ may not be locally small) because the algebras themselves have too many sorts to form a set, even though each A_s, as the homset $[\mathcal{T}, \mathsf{Set}](\mathcal{T}_s, \mathcal{A})$, is a set. This is easily addressed however with the further requirement, assumed henceforth, that $\mathcal{J}$ be essentially small, meaning equivalent to a small category (for example the subcategory of Set consisting of the finite sets is equivalent to a category whose objects are the natural numbers, even if Set has been constituted to have a proper class of singletons).

A more substantial reason is that there is no need to extend the representation of algebras by homomorphisms to the representation of homomorphisms between algebras because they already represent themselves.

However the Yoneda Lemma can be rephrased as a representation theorem for a more abstract notion of presheaf and morphisms thereof suggested by its original phrasing. Starting from $\mathcal{J}$ as an abstract category of prototypical structures and all their morphisms, define a presheaf as simply any object A of any full extension $\mathcal{C}$ of $\mathcal{J}$.

The representation theorem then takes the form that any functor $F : \mathcal{J} \to \mathcal{C}$ induces a representation of each object A of $\mathcal{C}$ as a $\mathcal{J}^{\mathrm{op}}$-algebra, namely the functor $\mathcal{C}(F(-), A) : \mathcal{J}^{\mathrm{op}} \to \mathsf{Set}$. And since the homfunctor is functorial in both arguments this extends to the representation of each morphism $h : A \to B$ of $\mathcal{C}$ as a homomorphism, namely the natural transformation $\mathcal{C}(F(-), h) : \mathcal{C}(F(-), A) \to \mathcal{C}(F(-), B)$. That is, F determines a functor $\widetilde{F}(A) = \mathcal{C}(F(-), A)$ of type $\widetilde{F} : \mathcal{C} \to [\mathcal{J}^{\mathrm{op}}, \mathsf{Set}]$, following the notation of [7][(3.38)].

Now a representation of (the objects and morphisms of) a category $\mathcal{C}$ by a category $\mathcal{D}$ is by definition a full and faithful functor from $\mathcal{C}$ to $\mathcal{D}$ [11]. This together with our understanding of $\mathcal{J}$ as itself a category of presheaves on $\mathcal{J}$ suggests the following notions and definitions.

A functor $F : \mathcal{J} \to \mathcal{C}$ is ***dense*** just when $\widetilde{F}$ is full and faithful, and is moreover ***complete*** when $\widetilde{F}$ is an equivalence. A ***category of presheaves on*** $\mathcal{J}$ is any dense full extension $\mathcal{C}$ of $\mathcal{J}$, where "full extension" means that $\mathcal{J}$ embeds fully in $\mathcal{C}$. A ***presheaf category*** on $\mathcal{J}$ is any complete full extension of $\mathcal{J}$.

All presheaf categories on $\mathcal{J}$ are equivalent, with the functor category $[\mathcal{J}^{\mathrm{op}}, \mathsf{Set}]$ among them, which for definiteness we continue to take as the definition of $\widehat{\mathcal{J}}$. Every category of presheaves on $\mathcal{J}$ is a full subcategory of $\widehat{\mathcal{J}}$, and (when defined as above) includes $\mathcal{J}$ as a full subcategory.

Dense functors originated with Isbell [5] who called them *left adequate*, and their opposite $F^{\mathrm{op}} : \mathcal{J}^{\mathrm{op}} \to \mathcal{C}^{\mathrm{op}}$ *right adequate*. The modern terminology, respectively dense and codense, is due to Gabriel and Ulmer [4] by analogy with density of the rationals in any extension thereof up to the reals as the completion of the rationals.

The above definition of dense functor, which we will call semantic density when so defined, makes the $\widetilde{F}$ representation of presheaf morphisms more an axiom than a theorem, having the character of defining a Boolean algebra to be any algebra isomorphic to a Boolean algebra of subsets of some set. One can then ask whether there exist natural characterizations of categories of presheaves for which it is a honest theorem that $\widetilde{F}$ is a representation.

2.2 Presheaves as colimits

One natural notion of category of presheaves on $\mathcal{J}$ is as a category $\mathcal{C}$ fully embedding $\mathcal{J}$ such that every object of $\mathcal{C}$ arises as a colimit of a functor to $\mathcal{C}$ that factors through F. As a degenerate case in point, consider how one might build the category Set as $\widehat{\mathbf{1}}$ by starting just from the base of Set, namely the category $\mathbf{1}$ whose one object is to be understood as the singleton 1. It is natural to express the set with 3 elements up to isomorphism as the sum or coproduct $1 + 1 + 1$, as an object of a suitable extension $\mathcal{C}$ of $\mathbf{1}$, the extension being formulated as a functor $G : \mathbf{1} \to \mathcal{C}$. We can formalize this as the colimit of a (necessarily constant) functor $\mathbf{3} \xrightarrow{F} \mathbf{1} \xrightarrow{G} \mathcal{C}$, where $\mathbf{3}$ denotes the 3-object discrete category.

This approach yields not only the set 3 itself but also the morphisms from 3 to any other object X of $\mathcal{C}$, via the adjunction $[\mathbf{3}, \mathcal{C}](GF, \Delta(X)) \cong \mathcal{C}(\mathrm{Colim}(GF), X)$ defining the notion of colimit, thinking of the left side of the adjunction as the definiens of colimit and the right side as the definiendum. Thus if $X = 5$ (the 5-element set) then $\Delta(X) = (5, 5, 5)$, $GF = (1, 1, 1)$, and the functor category $[\mathbf{3}, \mathcal{C}]$ therefore has $5^3 = 125$ natural transformations from GF to $\Delta(X)$, which the adjunction then obliges its right side to have as well. In this way $\widehat{\mathbf{1}}$ acquires not only all the objects but all the morphisms of Set.

The natural generalization to any $\widehat{\mathcal{J}}$ is obtained along these lines as the completion of $\mathcal{J}$ under colimits, with the comma category $(F \downarrow c)$ in place of $\mathbf{3}$ when defining c as a colimit, as per equation (1) of [10][§X.6]. For $\mathcal{J}$ as the basis for graphs as in subsection 1.4, a graph can be understood as a sum of copies of $\mathbf{v}$ and $\mathbf{e}$ with some identifications made between source and target vertices specified by the comma category.

The category $\widehat{\mathcal{J}}$ of "all" presheaves on $\mathcal{J}$ can then be taken to be some completion under colimits of $\mathcal{J}$, meaning a category $\mathcal{C}$ every extension of which by colimits of functors that factor through $\mathcal{J}$ is equivalent to $\mathcal{C}$.

2.3 Syntactic density

Defining presheaves as colimits makes the proposition that $\widetilde{F}$ is a representation an honest representation theorem instead of a mere definition. However abstraction promises both generality and simplicity. Whether or not the colimit approach can be considered a general characterization of presheaves and their morphisms, one would have to be thoroughly wedded to the categorical point of view to call it a simple one.

The following approach to density is developed in terms solely of categories and their extensions, with the latter understood as a more elementary notion than functor. The approach can be understood as simply a translation into more elementary language of the semantic definition of density.

Call category $\mathcal{D}$ an ***extension*** of category $\mathcal{C}$ just when $\mathcal{C}$ is a subcategory of $\mathcal{D}$. This is usually understood in category theory to mean that there exists a faithful functor $F : \mathcal{C} \to \mathcal{D}$. However if we impose the stronger condition that $F(c) = c$ for all objects and morphisms c in C, call this ***set-theoretic extension***, the kind normally understood in algebra, we can define the notion of extension of a category $\mathcal{C}$ before that of functor, namely as a category formed by adjoining new objects and morphisms to $\mathcal{C}$ and specifying how they compose, with no modification to the old objects and morphisms of $\mathcal{C}$. In particular functoriality $F(gf) = F(g)F(f)$ and $F(1) = 1$ is automatic for functors constituted as set-theoretic extensions.

A ***full*** extension of $\mathcal{C}$ is an extension that leaves all existing homsets of $\mathcal{C}$ unchanged, in the sense that it adds no new morphisms to them. More generally, for a specified class H of homsets of $\mathcal{C}$ an extension is called ***full on*** H when it leaves the homsets in H unchanged. An extension of $\mathcal{C}$ that adjoins only morphisms (no new objects) is called an extension of $\mathcal{C}$ ***by morphisms***. An extension $\mathcal{D}$ of $\mathcal{C}$ is called $\mathcal{C}$***-extensional*** (an unfortunate clash of terminology) when for all morphisms $f, g : d \to e$ in $\mathcal{D}$, if for all objects c in $\mathcal{C}$ and all morphisms $a : c \to d$ in $\mathcal{D}$ we have $fa = ga$, then $f = g$. Lastly an extension of an extension $\mathcal{D}$ of $\mathcal{C}$ is called $\mathcal{C}$***-full*** when it is full (as an extension of $\mathcal{D}$) on all homsets of $\mathcal{D}$ from $\mathcal{C}$. (The parameter $\mathcal{C}$ in "$\mathcal{C}$-full" and "$\mathcal{C}$-extensional" permits these notions to be applied unambiguously to extensions of extensions of $\mathcal{C}$.)

Call a full extension $\mathcal{C}$ of $\mathcal{J}$ ***dense*** when

(i) $\mathcal{C}$ is $\mathcal{J}$-extensional, and

(ii) every $\mathcal{J}$-full $\mathcal{J}$-extensional extension of $\mathcal{C}$ by morphisms is full. That is, $\mathcal{C}$ cannot be further extended without either extending some homset from $\mathcal{J}$, or violating $\mathcal{J}$-extensionality, or adjoining new objects.

By ***syntactic density*** we shall mean density defined in this way.

More operationally, a dense extension of $\mathcal{J}$ is one that can be constructed in two stages. At stage one, take any full extension of $\mathcal{J}$ such that the homsets from the new objects are empty. This creates all the algebras of the extension, along with just those homomorphisms required by the Yoneda Lemma. This extension is automatically $\mathcal{J}$-extensional because if $a, b : j \to A$ (the only new morphisms) are distinct then $a1_j = a \neq b = b1_j$, that is, a and b map the element 1_j of j to distinct elements of A.

At stage two, further extend the result by morphisms, by saturating the hitherto empty homsets from the new objects (those not in $\mathcal{J}$). That is, adjoin as many morphisms to those homsets as possible while preserving $\mathcal{J}$-extensionality to obtain the remaining homomorphisms. The homsets from $\mathcal{J}$ introduced at the first stage determine the objects of the extension and therefore must not be touched during the second stage, which adjoins only those morphisms between objects that don't participate in the Yoneda representation of the objects.

Theorem 3. *A full extension $\mathcal{C}$ of $\mathcal{J}$ is syntactically dense if and only if it is semantically dense.*

Proof. (If) Semantic density implies that $\mathcal{C}$ is a full subcategory of $[\mathcal{J}^{\mathrm{op}}, \mathsf{Set}]$, which along with having $\mathcal{J}$ as a full subcategory is then easily seen to satisfy the conditions of syntactic density.

(Only if) Any full extension of $\mathcal{J}$ is interpreted via $\widetilde{F}$ as having presheaves on $\mathcal{J}$ as objects, with every morphism interpreted as some homomorphism. It therefore remains to verify that each homomorphism is represented exactly once. Condition (i) of the definition of semantic density supplies the upper bound, condition (ii) the lower. □

For our purposes therefore the semantic-syntactic distinction for density can henceforth be dropped. In this way we obtain the notion of a category of presheaves without the notions of functor or natural transformation.

One might argue that surely algebras and homomorphisms accomplished this long ago, having been invented well before category theory. But properly understood these are merely synonyms for functors and natural transformations. Syntactic density involves only abstract categories and their extensions, without bringing in the usual notions of algebra and homomorphism thereof, however named. Defining these notions in terms of syntactic density is therefore more abstract and primitive than the more conventional approaches.

Without functors there is no $\widetilde{F}$ and hence no representation theorem. As soon as the homfunctor is introduced however we obtain both $\widetilde{F}$ and the representation of presheaves and their morphisms as $\mathcal{J}^{\mathrm{op}}$-algebras and their homomorphisms.

One benefit of this more elementary definition of density is that it generalizes quite easily to the notion of didensity defined in the next section. In contrast the corresponding generalization of $\widetilde{F}$ will be seen to be somewhat awkward.

3 Communes

3.1 Didensity

Thus far we have encountered morphisms in only one of the two argument positions of the homfunctor. In general, for any pair of morphisms $f : c' \to c$ and $h : d \to d'$ of a category $\mathcal{C}$, the homfunctor maps them to the function $\mathcal{C}(f, h) : \mathcal{C}(c, d) \to \mathcal{C}(c', d')$, which in turn maps each morphism $g : c \to d$ to $hgf : c' \to d'$. In this section we exploit the full generality of the homfunctor in a notion of commune (the suggested name in algebraic contexts) or disheaf (as its categorical counterpart making the connection with presheaf), as a common generalization of the notions of presheaf and Chu space [1, 2] (a construct that subsumes point set topology) that is particularly easy to define in terms of a notion of didense extension.

In place of a single small category $\mathcal{J}$ as the base we take two small categories $\mathcal{J}$ and $\mathcal{L}$. Elements of an object D (for disheaf) are represented as morphisms $a : j \to D$ as before, but now we also allow dual elements or *states* of D, represented as morphisms $x : D \to \ell$ for objects ℓ of $\mathcal{L}$. And since this necessitates a composite $xa : j \to \ell$, we fix homsets $K_{j\ell}$ supplying the values permitted for each such xa, along with their composites with morphisms of $\mathcal{J}$ at one end and $\mathcal{L}$ at the other.

This structure K, including $\mathcal{J}$ and $\mathcal{L}$ but excluding D, is called variously a bimodule, profunctor, or distributor. It can be understood as a functor $K : \mathcal{J}^{\mathrm{op}} \times \mathcal{L} \to \mathsf{Set}$, also notated $K : \mathcal{L} \nrightarrow \mathcal{J}$ when referred to as a profunctor. It can be understood equivalently as a category $\mathcal{K}$ formed from the category $\mathcal{J}+\mathcal{L}$ by adjoining the elements of $K(j,\ell)$ as morphisms from j to ℓ for objects j of $\mathcal{J}$ and ℓ of $\mathcal{L}$, with composites involving each element $k \in K(j,\ell)$ defined for each $f : i \to j$ in $\mathcal{J}$ as $kf = K(f,\ell)(k)$ (via the typing $K(f,\ell) : K(j,\ell) \to K(i,\ell)$) and for each $g : \ell \to m$ in $\mathcal{L}$ as $gk = K(j,g)(k)$ (via the typing $K(j,g) : K(j,\ell) \to K(j,m)$). The homsets from $\mathcal{L}$ to $\mathcal{J}$ are left empty.

Let $\mathcal{C}$ be a set-theoretic extension of the category $\mathcal{K}$. The notions of full, and full on certain homsets, for extensions are as before. Call the extension ***K-extensional*** when for all morphisms $f, g : c \to d$ in $\mathcal{C}$, if for all objects j of $\mathcal{J}$ and ℓ of $\mathcal{L}$ and for all morphisms $a : j \to c$ and $x : d \to \ell$ in $\mathcal{C}$ we have $xfa = xga$, then $f = g$. Call it ***K-full*** when it is full on homsets from $\mathcal{J}$ and homsets to $\mathcal{L}$.

Let $\mathcal{C}$ be any K-full extension of $\mathcal{K}$, that is, full on all homsets of $\mathcal{K}$ save those from $\mathcal{L}$ to $\mathcal{J}$. Call $\mathcal{C}$ ***didense*** when

(i) $\mathcal{C}$ is K-extensional, and

(ii) every K-full K-extensional extension of $\mathcal{C}$ by morphisms is full.

A ***category of communes*** on $\mathcal{K}$ is any K-full didense extension of $\mathcal{K}$. The terminology "on $\mathcal{K}$" is taken as implying that $\mathcal{K}$ is included as a K-full subcategory.

As with presheaves, merely belonging to the same category $\mathcal{C}$ as $\mathcal{K}$ is enough to qualify an object for interpretation as a commune, whose elements are the morphisms to it from $\mathcal{J}$ and whose states are the morphisms from it to $\mathcal{L}$, with the remaining structure given by the composites of elements and states with each other and with the morphisms of $\mathcal{J}$ and $\mathcal{L}$. From this perspective communes are as simple and natural a notion as presheaves.

One might expect such a generalization to have a straightforward conventional formulation generalizing the notion of module $(\mathcal{C}, \mathcal{X}, \cdot)$, or presheaf $A : \mathcal{C} \to \mathsf{Set}$. Following that approach naively for communes requires attention to ten components, namely the doubly-indexed family $K_{j\ell}$ of sets, the category $\mathcal{J}$, the composites of K with $\mathcal{J}$, the category $\mathcal{L}$, the composites of $\mathcal{L}$ with K, the sets of elements of presheaves on $\mathcal{J}$, their composites with $\mathcal{J}$, the sets of states of dual presheaves on $\mathcal{L}$, the composites of $\mathcal{L}$ with states, and the composites of states with elements.

With further organization all this can be packaged more neatly to define a commune on the profunctor $K : \mathcal{L} \nrightarrow \mathcal{J}$ as a triple (A, X, ρ) where $A : 1 \nrightarrow \mathcal{J}$ and $X : \mathcal{L} \nrightarrow 1$ are profunctors and $\rho : AX \to K$ is a natural transformation from the composite $AX : \mathcal{L} \nrightarrow \mathcal{J}$ to K. Without going into detail about composition of profunctors in general it suffices to say that the composite AX in this special case is a functor $\mathcal{J}^{\mathrm{op}} \times \mathcal{L} \to \mathsf{Set}$ satisfying $AX(j,\ell) = A(j) \times X(\ell)$ and $AX(f,h)(a,x) = (A(f)(a), X(h)(x))$. Following the tradition in algebra of identifying algebras with their representation as functors (construed sufficiently broadly as to allow the tradition to predate the invention of category theory per se), we can identify communes with their representation as triples (A, X, ρ) while continuing to regard abstract didense K-full extensions of a bimodule $\mathcal{K}$

as categories of communes understood more abstractly.

This organization brings out the relationship between communes on a bimodule K and Chu spaces (A, X, r) on a set K defined as sets A and X and a function $r : A \times X \to K$. One might call communes the categorification of Chu spaces, with the caveat that although Chu spaces form a self-dual category, communes do not unless $\mathcal{K}$ is self-dual. Chu spaces arise as the special case $\mathcal{J} = \mathcal{L} = \mathbf{1}$, making the domain of $K : \mathcal{J}^{\mathrm{op}} \times \mathcal{L} \to \mathsf{Set}$ a singleton and hence K merely a set, namely the set of morphisms from the object of $\mathcal{J}$ to that of $\mathcal{L}$.

We write $\widehat{\mathcal{K}}$ for the category of communes on $\mathcal{K}$ by analogy with $\widehat{\mathcal{J}}$ for the category of presheaves on $\mathcal{J}$.

For presheaves we called the functor $F : \mathcal{J} \to \mathcal{C}$ semantically dense when the functor $\widetilde{F} : \mathcal{C} \to \widehat{\mathcal{J}}$ defined by $\widetilde{F}(A)(j) = \mathcal{C}(F(j), A)$ is full and faithful. Taking didensity as defined above to be syntactic didensity, for communes we call the functor $F : \mathcal{K} \to \mathcal{C}$ ***semantically didense*** when the functor $\widetilde{F} : \mathcal{C} \to \widehat{\mathcal{K}}$ defined by $\widetilde{F}(D)(j, \ell) = \mu_{F(j)AF(\ell)} : \mathcal{C}(A, F(\ell)) \times \mathcal{C}(F(j), A) \to \mathcal{C}(F(j), F(\ell))$ is full and faithful, where $\mu_{cde} : \mathcal{C}(d, e) \times \mathcal{C}(c, d) \to \mathcal{C}(c, e)$ is composition in $\mathcal{C}$. When F is K-full and faithful, the equivalence of syntactic and semantic didensity can be argued similarly to the proof of Theorem 3.

Communes are a generalization of a notion due to Isbell and called by Lawvere the ***Isbell envelope*** $E(\mathcal{C})$ of a category $\mathcal{C}$. $E(\mathcal{C})$ is the special case of a category of communes where the base has the form of a homfunctor $\mathcal{C}^{\mathrm{op}} \times \mathcal{C} \to \mathsf{Set}$, equivalently the identity profunctor $1_C : \mathcal{C} \nrightarrow \mathcal{C}$. An object D of the Isbell envelope can be understood as a commune whose elements are morphisms from objects of $\mathcal{C}$ to D and whose states are morphisms from D to objects of $\mathcal{C}$. Conversely the commune category $\widehat{\mathcal{K}}$ can be obtained from $E(\check{K})$ as the full subcategory of $E(\mathcal{K})$ consisting of those objects having no elements from $\mathcal{L}$ and no states to $\mathcal{J}$.

Call a commune D ***extensional*** when as a one-object full extension of $\mathcal{K}$ (i.e. $\mathcal{K}$ together with the elements and states of D) it forms a J-extensional extension of $\mathcal{J}$. That is, for all objects ℓ of $\mathcal{L}$ and all $x, y : D \to \ell$, if for all objects j of $\mathcal{J}$ and all elements $a : j \to D$ we have $xa = ya$, then $x = y$. (This agrees with the terminology for an extensional Chu space [2].) Call D ***consistent*** when for every ℓ in $\mathcal{L}$ there exists a state $x : D \to \ell$. Call D ***discrete*** when it is maximally extensional, meaning that it is extensional and no new state can be adjoined without violating extensionality.

3.2 Categories of graphs

Subsection 1.4 illustrated presheaves with two bases $\mathcal{J}$ and $\mathcal{J}'$ for respectively reflexive and general graphs. Here we give some extensions of the latter to profunctors $K : \mathcal{L} \nrightarrow \mathcal{J}'$ for various choices of $\mathcal{L}$ each understood itself as a category of graphs playing the role of the Sierpinski space in this generalization of point set topology. The benefit of $\mathcal{J}'$ over $\mathcal{J}$ in these examples is that the absence of self-loops (more generally idempotents) enhances combinatorial expressivity, in contrast to the Sierpinski space which has little to offer combinatorics. In the terminology of Grothendieck, combinatorics is better conducted in a *gros topos* than a *petit* one.

These bimodules can be presented as dense full extensions of either $\mathcal{J}'$ or $\mathcal{L}$. That is, the extensions produce respectively the objects of either $\mathcal{L}$ or $\mathcal{J}'$. We take the latter as being more convenient for these graph examples.

Call a graph ***uniform*** when for any pair u, v of vertices all paths from u to v have the same length. For example the Hasse diagram of any finite modular lattice is uniform. We represent uniform graphs as consistent extensional communes on the following bimodule $\mathcal{K}$.

Take $\mathcal{L}$ to be the monoid $\mathcal{Z} = (\mathbb{Z}, +, 0)$ of integers. Form the $\mathcal{Z}$-module $(\mathcal{Z}, \mathbb{Z}, +)$ extending $\mathcal{Z}$. Form $\mathcal{J}'$ by taking V and E to be two such extensions of $\mathcal{Z}$, with $\sigma^* : V \to E$ acting as the identity on the integers and $\tau^* : V \to E$ as predecessor. This defines $\check{K}$ to consist of $\mathcal{J}'$ having two objects V and E and two morphisms $\sigma^*, \tau^* : V \to E$, $\mathcal{L}$ with one object Z for which $\mathrm{End}(Z)$ is $\mathcal{Z}$ as a monoid, and $K(V, Z) = K(E, Z) = \mathbb{Z}$ supplying respectively the vertices and edges of the graph. Just as each of V and E forms a right $\mathbb{Z}$-module in $\mathcal{K}$, so does Z form a left $\mathcal{J}'$-module in $\check{K}$, namely a graph with the property that for every integer i, edge i runs from vertex i to vertex $i + 1$.

A consistent extensional commune on this bimodule $\mathcal{K}$ then consists of a graph G and a nonempty set X of states expressible as distinct graph homomorphisms $x : G \to \mathcal{Z}$, which we can understand as a painting of G with integers such that every path in G is painted with consecutive integers. Since X is nonempty (the consistency requirement) this forces G to be a uniform graph. The presence in $\mathcal{L}$ of all translations of $\mathcal{Z}$ forces G to be discrete.

Thinking of $\mathcal{L}$ here as the graph of the successor function on $\mathbb{Z}$, replacing it by the graph of the binary relation $i < j$ yields communes that are acyclic graphs. In place of translations we take the corresponding morphisms of $\mathcal{L}$ to be strict monotone functions on $\mathbb{Z}$, those satisfying $f(i) < f(j)$ when $i < j$. For finite communes we obtain all acyclic graphs, namely as the discrete consistent communes on this bimodule. (Discreteness is needed here because the morphisms of $\mathcal{L}$ must be dilations; if we start with some painting $x : G \to Z$ which interprets edges as very large steps through the integers, applying dilations to these does not force more conservative paintings, the omission of which would impute unintended structure to the model, whence the discreteness requirement.)

For larger communes some acyclic graphs are not so representable, for example the graph of the less-than relation on the rationals, which contains infinitely many vertices between any two distinct vertices, which cannot be labeled with integers from any finite interval of integers. This can be dealt with for graphs of size up to any given cardinal by taking $\mathcal{L}$ to be a suitable large linear order.

This raises the question, does there exist a category of discrete consistent communes equivalent to the category of all acyclic graphs? Allowing $\mathcal{L}$ to be some partial order seems not to help.

3.3 Ontology of properties and qualia

Three long-standing problems of philosophy are, in decreasing order of seniority, Cartesian dualism, the nature of properties or attributes, and the existence of qualia.

The problem of Cartesian dualism concerns the coexistence of mind and body and in particular the nature of their interaction. Descartes proposed in 1647 that the universe consists of three components, God, mind, and body, with the latter two independent except in man where they interact by some as yet unknown mechanism, possibly the recently discovered pituitary gland for which no role had yet been found. (The recent house arrest of Galileo may have had something to do with the inclusion of God in this analysis.) Such sharply defined Theories of Everything being something of a novelty in those days, philosophers struggled with the question of how mind and body could interact for a century before giving up and rejecting one of them as illusory, with the more scientifically inclined preferring to dismiss mind as not existing in the universe the way body does.

The notion of type such as cat or dog is readily understood in terms of its extension in any given universe, namely the set of all entities of that type. The notion of property has proved harder to pin down. C. Swoyer's comprehensive (28,000-word) Stanford Encyclopedia of Philosophy article on the subject of properties [12] lists a number of possible answers to the primary question, what is a property, all of which raise to one degree or another secondary questions about the number of properties, their identity, their modal status, and their epistemic status. The Wikipedia article "Property (philosophy)" on the other hand merely states baldly that properties differ from classes by lacking any concept of extensionality. This has both the succinctness and falsifiability of a scientific hypothesis.

In 1929 C.I.Lewis, an early contributor to modal logic, wrote *Mind and the World Order: Outline of a Theory of Knowledge* [9] in which he summarized his thinking about qualia as entities bridging the physically observable (as measured by scientific instruments) and the psychologically observable (as the sensations reported by human observers). Philosophers have since divided themselves into qualiaphiles such as Edmond Wright, editor of *The Case for Qualia* [13], a just-published score of qualia-friendly essays, and qualiaphobes such as Daniel Dennett [3] who maintain that the concept is incoherent.

Communes are a new mathematical construct that provide a common solution to all three problems by giving a way of thinking about them. Since communes are well-defined, this allows the questions to be formulated more sharply as, how faithfully do communes capture the notions of mind, property, and quale? Communes also suggests novel ways of defining and organizing those notions so as to make them more consistent both individually and in combination with each other.

It should be clear how to interpret morphisms $a : j \to D$ as individuals of type j in universe D. In particular the homset $\widehat{\mathcal{K}}(j, D)$ represents the extension of type j in universe D as a set, while morphisms $f^* : j \to i$ contravariantly represent operations $f : i \to j$ acting on individuals.

We propose to interpret objects ℓ of $\mathcal{L}$ as properties, states $x : D \to \ell$ as local states specific to an abstract or idealized observer of property ℓ, and morphisms $d : \ell \to m$ as dependencies between properties, for example hue and saturation may depend on color but not brightness. The set $\widehat{\mathcal{K}}(D, \ell)$ consists of the *possible* states of abstract observer ℓ. At any one time that observer will be in the state resulting from observing all the individuals of D.

The meaning of $xa : j \to \ell$ is the value of property ℓ of individual a of type j in ℓ-specific state x. This value is drawn from the set $K(j, \ell)$ of possible values of property ℓ for individuals of type j, for example the height of a building, the color of a cat, etc.

The global state of a universe D is a state vector indexed by properties. In relational databases a row of a relation, or unit record, would correspond to the values of some such state vector as applied to a particular individual a. Updates to the database, say of attribute ℓ, are understand as replacement of the current state $x : D \to \ell$ with $x' : D \to \ell$, entailing possible changes to values of other attributes as required by the dependencies.

We identify individuals as the inhabitants of the physical world, its bodies, and states as the alternatives of the mental world, its minds. To be of two minds about something is to be torn between two alternative states. Qualia as the elements of $K(j, \ell)$ arise at the boundary of these physical and mental entities, which would seem to come about as close to capturing the schizophrenic mind-body nature of Lewis's qualia as any simple mathematical theory is likely to get. Philosophers may well have a somewhat different notion in mind, say with a stronger psychological element than a naive mathematical model can convey. However communes at least provide a coherent notion of qualia as mediating the covariant and contravariant notions of respectively individual and state, thereby overcoming the basic objection that qualia are incoherent in the sense that there is no consistent plausible notion of them.

3.4 Implications for modal logic

Global states of a universe as state vectors of a commune indexed by properties give a notion of possible world similar to that of a Kripke structure. One difference from Kripke structures is the concurrent nature of the local states. The abstract observers of the properties can be understood as observing concurrently, and independently to the extent permitted by the dependencies. In the absence of dependencies the observers behave independently and asynchronously: one observer might make few observations in the time that another makes many. Dependencies create (more or less) synchronous correlations between properties.

In modal logic, the Barcan formula $\Diamond\exists x\Phi(x) \to \exists x\Diamond\Phi(x)$ asserts that if it is possible for a certain entity (namely one satisfying an arbitrary formula $\Phi(x)$, for example "x is Santa Claus,") to exist, then there exists an entity which can acquire that property. In other words, the notion of "possible" permits changes to existing entities but not the creation of new ones. The converse Barcan formula $\exists x\Diamond\Phi(x) \to \Diamond\exists x\Phi(x)$ asserts the persistence of entities: it is not possible for entities to vanish from the universe.

Within a commune individuals persist independently of choice of state: the properties of an individual may change but not the individual itself, providing an interpretation of Kripke's notion of rigid designator [8]. Hence possibility within a universe of alternative states does not allow for creation or annihilation of individuals, nor even for their identification (as with the realization that the morning and evening star are both Venus) or resolution into multiple individuals (as with the realization that someone who seemed to get around very quickly is actually a pair of identical twins). For the fixed-universe notion of possibility both Barcan formulas hold, and can even be strengthened suitably to reflect the impossibility of identification or resolution.

If however possibility admits the transformation of universes by homomorphisms, then although individuals cannot be annihilated or resolved, they can be created and identified. The Barcan formula then fails, although its converse still holds. With yet more general notions of possibility, for example via opposite homomorphisms, one can refute both formulas. Communes therefore provide a framework in which one can eat one's cake and have it too as far as the Barcan formulas are concerned.

3.5 Implications for evolution

We certainly perceive our universe in terms of types and properties. But why? The naive answer is because that's how our universe, and surely every universe, happens to be organized.

A more objective answer might be that brains evolved those notions because they were simpler than the alternatives. If evolution is a billion-year experiment to determine what features of self-reproducing organisms work best in any given environment, it is going to stumble over simple mechanisms long before it invents complicated ones.

We certainly think in terms of types and properties. At the other extreme it seems unlikely that single-cell organisms do so. Somewhere in between, animal brains made that distinction. If there is a simple way of structuring the distinction, for example along the lines of communes, its simplicity lends plausibility to the theory that our ancestors at some point stumbled on the concept and found it very useful.

To the extent that density is a simpler notion than either functor or natural transformation, it stands a better chance of being stumbled on first by evolving brains. And since presheaves and communes are even simpler, being merely objects in the same category as respectively some $\mathcal{J}$ and some $\mathcal{K}$, these objects may well have been found much more easily than the morphisms.

Any competing theory of thought needs a comparably simple structure if it is to be a plausible candidate for chance discovery in the limited time available for the evolution of modern thought processes.

4 Acknowledgments

Although Bill Lawvere had pointed me at Isbell's papers in connection with left and right adequacy at Category Theory 2004 in Vancouver where I first spoke about communes (during which I was introduced to bimodules by Robert Seely), I first learned about Lawvere's term "Isbell envelope" $E(\mathcal{C})$ for that concept much more recently from Ross Street, which Ross defined for me in terms of left Kan extensions. Richard Wood then pointed out the definition of objects

of $E(\mathcal{C})$ as triples (A, X, ρ) where $A : \mathbf{1} \nrightarrow \mathcal{C}$ and $X : \mathcal{C} \nrightarrow \mathbf{1}$ are profunctors and $\rho : AX \to 1_C$ is a natural transformation thereof, and Jeff Egger independently pointed out that communes could be accommodated via the generalization to $A : \mathbf{1} \nrightarrow \mathcal{J}$ and $X : \mathcal{L} \nrightarrow \mathbf{1}$ with $\rho : AX \to K$ for $K : \mathcal{L} \nrightarrow \mathcal{J}$. Many thanks to Mike Barr for the concept of Chu spaces, without which I would never have thought of communes.

References

[1] M. Barr. ∗-*Autonomous categories*, volume 752 of *Lecture Notes in Mathematics*. Springer-Verlag, 1979.

[2] M. Barr. The Chu construction. *Theory and Applications of Categories*, 2(2):17–35, 1996.

[3] D.C. Dennett. *Consciousness Explained*. Little, Brown, and Company, 1991.

[4] P. Gabriel and F. Ulmer. *Lokal präsentierbare Kategorien*, volume 221 of *Lecture Notes in Mathematics*. Springer-Verlag, 1971.

[5] J.R. Isbell. Adequate subcategories. *Illinois J. Math.*, 4:541–552, 1960.

[6] P.T. Johnstone. *Sketches of an Elephant: A Topos Theory Compendium*. Oxford Science Publications, 2002.

[7] G.M. Kelly. *Basic Concepts of Enriched Category Theory: London Math. Soc. Lecture Notes*. 64. Cambridge University Press, 1982.

[8] S. Kripke. Identity and necessity. In S.P. Schwartz, editor, *Naming, Necessity, and Natural Kinds*, pages 66–101. Cornell University Press, Ithaca and London, 1977.

[9] C.I. Lewis. *Mind and the World Order: Outline of a Theory of Knowledge*. Scribner's Sons, 1929.

[10] S. Mac Lane. *Categories for the Working Mathematician*. Springer-Verlag, 1971.

[11] A. Pultr and V. Trnková. *Combinatorial, Algebraic and Topological Representations of Groups, Semigroups, and Categories*. North-Holland, 1980.

[12] C. Swoyer. Properties. In *Stanford Encyclopedia of Philosophy*. Stanford University, 2000.

[13] E. Wright. *The case for qualia*. MIT Press, Boston, 2008.