Chu Spaces from the Representational Viewpoint

Vaughan Pratt

Department of Computer Science, Stanford University, Stanford, CA 94305-9045

Abstract

We give an elementary introduction to Chu spaces viewed as a set of strings all of the same length. This perspective dualizes the alternative view of Chu spaces as generalized topological spaces, and has the advantage of substituting the intuitions of formal language theory for those of topology.

1 Background

Chu spaces provide a simple, uniform, and well-structured approach to the representation of objects that may possess algebraic, relational, or other structure, and that can transform into one another in ways that respect that structure. Chu spaces are simple by virtue of being merely a rectangular array, with no further machinery. They are uniform in the sense that all transformable objects, whether sets, groups, Boolean algebras, vector spaces, or manifolds, are representable by Chu spaces within the same framework, and hence can coexist in a single typeless universe of mathematical objects. And they are well-structured in that this seemingly featureless universe in fact has a natural and rich structure given by Girard's linear logic [Gir87].

To climb up to this universe of structured objects, we use as our ladder a universe of unstructured or *discrete* objects, namely ordinary sets and the functions between them. Having done so, we then pull this ladder up after us by representing sets as discrete Chu spaces.

For all practical purposes almost any reasonable understanding of sets and functions will suffice for our development. But it will not hurt to say informally what parts of set theory we make essential use of. Nowhere shall we depend

[1] Supported by ONR under grant number N00014-92-J-1974.

on the existence of infinite sets, though of course if they do exist then one can manufacture infinite Chu spaces from them, essential for representing say the ring of integers or the field of reals.

What we do need is binary *cartesian product* $A \times B$, along with the *function space* $A \to B$, the set of all functions from the set A to the set B. We also need the Currying principle, namely the bijection between $(A \times B) \to C$ and $A \to (B \to C)$ which puts each f in the former into correspondence with f' in the latter via $f(a, b) = f'(a)(b)$. And, given two functions $f, g : A \to B$, we need to be able to form the subset of A on which f and g agree.

The history of Chu spaces is roughly as follows. The basic idea of representing duality as a contravariant pair of morphisms goes back to G. Mackey [Mac45]. [2] This idea was abstracted by M. Barr to Chu spaces enriched (in the sense of enriched category theory [Kel82]) in a symmetric monoidal category V and studied by his student Peter Chu [Bar79, App.]. The case $V = \mathbf{Set}$ of *ordinary* Chu spaces, the kind we treat here, was first studied in detail by Lafont and Streicher [LS91] under the rubric of games, and by Brown, Gurr, and de Paiva [BGdP91], and Blass [Bla95], who have treated a "lax-continuous" variant in which the adjointness condition defining continuity is relaxed from an equality to an inequality.

Our own interest in Chu spaces originated in their application to the representation of generalized event structures [GP93], but we have since found them also of interest as universal objects [Pra95,Pra96], broadening the denotational semantics of linear logic to a much larger, in fact universal, class of mathematical objects than previously associated with linear logic.

2 Representation

A Chu space resembles a formal language, in that it may be understood intuitively as a set of "words" over an alphabet Σ. An ordinary word w of length n can be defined as a function $w : \{1, 2, \ldots, n\} \to \Sigma$, having for its i-th symbol $w(i)$ or w_i.

Chu spaces modify this in two ways. First, a word is taken to be a function $w : X \to \Sigma$ where X is an arbitrary set, not necessarily an initial segment of the positive integers. We may then speak of X as the "length" of w. (Note that the number of words of "length" a five-element set is the same as the number of ordinary words of ordinary length 5.)

Second, the words of a given Chu space are all of the same length, i.e. all have

[2] Lawvere has recently suggested calling it the Chu-Mackey construction.

a given set X as their common domain. No restrictions are placed on either X or Σ, which may be empty, or initial segments of the positive integers, or sets of reals, or any other set. Likewise no restrictions are placed on the words, which may be any function from X to Σ.

This intuition is formalized as follows. A *Chu space over* Σ is a triple $\mathcal{A} = (A, r, X)$ consisting of sets A and X and a function $r : A \times X \to \Sigma$. We call A and X respectively the *carrier* and *cocarrier* of $\mathcal{A}$, their elements respectively *points* and *states*, and r the *interaction matrix*.

Words and dual words. The interaction matrix has *left* and *right transposes* $\hat{r} : A \to \Sigma^X$ and $\check{r} : X \to \Sigma^A$ satisfying $\hat{r}(a)(x) = r(a,x) = \check{r}(x)(a)$, which we may interpret as representations of A and X respectively. For each point $a \in A$, $\hat{r}(a)$ represents a as a function from X to Σ, i.e. of type Σ^X, namely a word over an alphabet Σ of length X in the above sense. Dually $\check{r}(x)$ represents state x as a function from A to Σ, similarly constituting a word of length A over the same alphabet Σ, which we shall refer to as a *dual* word of $\mathcal{A}$.

The Chu space whose words are the dual words of $\mathcal{A} = (A, r, X)$ is $(X, r^{\smile}, A)$ where $r^{\smile}(x, a) = r(a, x)$, called the *dual* of $\mathcal{A}$ and denoted $\mathcal{A}^{\perp}$.

When two points a, b are represented by the same word, i.e. when $\hat{r}(a) = \hat{r}(b)$, we call them *equivalent*, written $a \equiv b$, Dually, equivalent states, those satisfying $\check{r}(x) = \check{r}(y)$, are likewise indicated by $x \equiv y$.

When no two distinct points are equivalent, i.e. $\hat{r} : A \to \Sigma^X$ is injective, we call $\hat{r}$ a *faithful* representation of A, and say that $\mathcal{A}$ is *separable*. Dually when $\check{r} : Y \to \Sigma^A$ is injective we call it a faithful representation of X, and say that $\mathcal{A}$ is *extensional*. A Chu space that is both separate and extensional is called *biextensional*.

The usual notion of a formal language as a set of words all distinct corresponds to the property of separability. Nonseparable Chu spaces may be understood as multisets of words, allowing the same word to occur more than once in the language. For our purposes the identity of points and states is determined by their representations, and for this reason the biextensional Chu spaces will be the ones we shall be mainly concerned with.

The alphabet Σ itself forms a language consisting of all words of length 1 over alphabet Σ. This makes it a Chu space, which we denote $\perp = (\Sigma, \pi_1, 1)$ where $1 = \{0\}$ and $\pi_1 : \Sigma \times 1 \to \Sigma$ is projection on the first coordinate, satisfying $\pi_1(i, 0) = i$ for all $i \in \Sigma$. Its dual, $\perp^{\perp}$, consists of a single word containing one occurrence of every symbol in Σ. Its role is as the discrete singleton, denoted by 1.

The discrete empty language, denoted 0, is $(\emptyset, !, 1)$. The inconsistent empty language, $(\emptyset, !, \emptyset)$, plays no important role and needs no name.

3 Transformation

We now consider how Chu spaces transform. As one would expect, a function $f : \mathcal{A} \to \mathcal{B}$ between Chu spaces sends each word $a \in \mathcal{A}$ to some word $f(a) \in \mathcal{B}$. But if some Chu spaces are to have nontrivial structure, not all functions will preserve that structure. Those that do preserve it we shall call *continuous*. What we shall define however is not the notion of structure but of continuity. Later we shall define and defend a suitable notion of structure, and show that, among all functions between Chu spaces, the continuous ones are exactly those preserving that structure.

Our basic example of a continuous function will be any *projection* from $\mathcal{A} = (A, r, X)$ to $\perp$. A projection is defined as any dual word of $\mathcal{A}$, that is, a function $\check{r}(x)$ from $\mathcal{A}$ to Σ for some $x \in X$.

By way of motivation we give a preliminary definition of continuous function.

Continuity 1. The continuous functions are the largest class such that

(i) every continuous function to $\perp$ is a projection, and

(ii) the composition of two continuous functions is continuous.

This is a third order characterization of continuity, being phrased in terms of classes of functions which themselves are second-order entities. We now give a second-order definition of continuity and prove its equivalence to the above.

Continuity 2. $f : \mathcal{A} \to \mathcal{B}$ is continuous just when for every projection $\pi : \mathcal{B} \to \perp$ of $\mathcal{B}$, πf is a projection of $\mathcal{A}$.

Proposition 1 *Continuity-1 and continuity-2 are equivalent.*

PROOF. We first show that continuous-2 implies continuous-1. For this it suffices to show that the class of continuous-2 functions meets both 1(i) and 1(ii). Observe first that the identity function on $\perp$ is a projection of $\perp$ (the only projection in fact). Now if $f : \mathcal{A} \to \perp$ is continuous-2 then the composition of the identity on $\perp$ with f, namely f itself, must be be a projection, whence 1(i) is satisfied by continuity-2.

For 1(ii), let $\mathcal{A} \xrightarrow{f} \mathcal{B} \xrightarrow{g} \mathcal{C}$ be the composition of two continuous-2 functions. For any projection $\pi : \mathcal{C} \to \perp$ of $\mathcal{C}$, πg must be a projection of $\mathcal{B}$, but then

$\pi g f$ must be a projection of $\mathcal{A}$. Hence gf is continuous-2 and so continuity-2 satisfies 1(ii). This completes one direction.

We now show that continuous-1 implies continuous-2. For a contradiction, let $f : \mathcal{A} \to \mathcal{B}$ be any function that is continuous-1 but not continuous-2. By the latter there must exist a projection $\pi : \mathcal{B} \to \perp$ such that πf is not a projection of $\mathcal{A}$. Hence both f and π are continuous-1, but then their composition πf cannot simultaneously satisfy 1(i) and 1(ii).

We remark in passing that the continuous maps to $\perp$ are exactly the projections. This yields another view of a Chu space, namely as the dual of the space consisting of the continuous functions to $\perp$. The shorter the words of a space, the fewer the continuous functions to $\perp$.

Now the operation of composing with f defines a function from the dual words of $\mathcal{B} = (B, s, Y)$ to those of $\mathcal{A} = (A, r, X)$. When $\mathcal{A}$ is extensional this determines a function $g : Y \to X$ such that for each $y \in Y$, $g(y)$ indexes the projection of $\mathcal{A}$ that equals $\check{s}(y)f$, i.e. $\lambda a.s(f(a), y)$. The projection of $\mathcal{A}$ that it must equal is $\lambda a.r(a, g(y))$. The continuity condition can now be stated in first order terms (no quantification over functions or predicates) as the equation

$$s(f(a), y) = r(a, g(y)).$$

We call this equation the *adjointness condition* on account of its resemblance to adjoint relationships in linear algebra and categorical adjunctions. The condition may be understood loosely as saying that g is a form of inverse of f, more precisely an adjoint. We call g the *adjoint of f*.

Although we obtained the adjointness condition from our original definition of continuity by assuming that $\mathcal{A}$ was extensional, the condition itself does not make any use of that assumption. In fact pairs of functions satisfying this condition define the most basic notion of morphism of Chu spaces. The continuous functions $f : \mathcal{A} \to \mathcal{B}$ can then be defined as those functions $f : A \to B$ such that there exists $g : Y \to X$ making (f, g) an adjoint pair. This version of the definition of continuity makes no assumption about either separability or extensionality of either $\mathcal{A}$ or $\mathcal{B}$.

Adjoint pairs $\mathcal{A} \xrightarrow{(f,g)} \mathcal{B} \xrightarrow{(f',g')} \mathcal{C}$, where $\mathcal{A} = (A, r, X)$, $\mathcal{B} = (B, s, Y)$, and $\mathcal{C} = (C, t, Z)$, compose via $(f', g')(f, g) = (f'f, gg')$. That this is itself an adjoint pair follows from $t(f'f(a), z) = s(f(a), g'(z)) = r(a, gg'(z))$. Hence Chu spaces over Σ and their adjoint pairs form a category, denoted $\mathbf{Chu}_\Sigma$. We denote by $\mathbf{chu}_\Sigma$, pronounced "little chu," the subcategory of $\mathbf{Chu}_\Sigma$ whose objects are the biextensional Chu spaces over Σ and whose morphisms are all adjoint pairs between them (i.e. a full subcategory of $\mathbf{Chu}_\Sigma$).

Now the adjointness condition, despite being first-order, is a little bit magical, and for this reason we started out with higher-order definitions that did not contain a magic formula and hence were better motivated. One advantage of the adjointness condition besides its elementary nature is that it demonstrates the symmetry of continuity with respect to transposition or duality: the dual (g, f) of an adjoint pair (f, g) from $\mathcal{A}$ to $\mathcal{B}$ is itself an adjoint pair from $\mathcal{B}^\perp$ to $\mathcal{A}^\perp$, its adjointness condition being

$$r^{\smile}(g(y), a) = s^{\smile}(y, f(a)).$$

That is, the dual of a continuous function f is the adjoint of f whose existence the continuity of f requires. This dualizability is not at all apparent from either of our first two definitions of continuity.

We now give a definition of continuity that combines the best features of both the non-magical definitions and the adjoint-pair definition. We exploit the representational aspect of Chu spaces in such a way that duality can be integral to the definition of continuity, yet without pulling any formulas out of a hat.

Lift the representation $\hat{s} : B \to \Sigma^Y$ of points of $\mathcal{B} = (B, s, Y)$ to a representation of $f : \mathcal{A} \to \mathcal{B}$, simply by forming the composition $\hat{s}f : A \to \Sigma^Y$. This represents f pointwise in terms of the representation $\hat{s}(f(a))$ of each point in the image of f. But $\hat{s}f$ is the left transpose of a function $\varphi : A \times Y \to \Sigma$, namely $\varphi(a, y) = \hat{s}(f(a)) = s(f(a), y)$, which we can view as a Chu space $\mathcal{F} = (A, \varphi, Y)$ representing f,

We define a function $f : \mathcal{A} \to \mathcal{B}$ to be continuous when the dual $\mathcal{F}^\perp$ of its representation $\mathcal{F}$ represents a function from $\mathcal{B}^\perp$ to $\mathcal{A}^\perp$.

That is, for some function $g : Y \to X$, the left transpose of $\varphi^{\smile}$ must be $\check{r}g : Y \to \Sigma^A$ satisfying $(\check{r}g)(y)(a) = r(a, g(y))$. So our dualization requirement becomes $s(f(a), y) = r(a, g(y))$. But this is exactly the adjointness condition and hence is equivalent to the other definitions.

To these four definitions of continuity we may add a fifth in the case $\Sigma = 2 = \{0, 1\}$: a function $f : \mathcal{A} \to \mathcal{B}$ is continuous when the inverse image of each dual word of $\mathcal{B}$, viewed as a subset of B, is a dual word of $\mathcal{A}$. This is the standard definition of continuity from point-set topology, where our dual words play the role of open sets. But this is easily seen to be just a restatement of Continuity-2.

Concreteness. Let $U(\mathcal{A})$ denote the underlying set A of $\mathcal{A} = (A, r, X)$, and for an adjoint pair (f, g) let $U(f, g)$ denote f. Then U is a functor from $\mathbf{Chu}_\Sigma$ to $\mathbf{Set}$.

Now f need not determine g uniquely. In particular if both r and s are all-zero matrices, every pair (f, g) of functions between (A, r, X) and (B, s, Y) is trivially an adjoint pair. Hence U is not a faithful functor, whence $\mathbf{Chu}_\Sigma$ with this choice of forgetful functor is not concrete.

When $\mathcal{A}$ is extensional and $f : A \to B$ is continuous from $\mathcal{A}$ to $\mathcal{B}$, the adjoint $g : Y \to X$ of f making (f, g) an adjoint pair is uniquely determined. Hence the restriction of U to the extensional Chu spaces in $\mathbf{Chu}_\Sigma$ is a faithful functor, making that subcategory of $\mathbf{Chu}_\Sigma$ a concrete category.

Now the dual of an extensional Chu space while separable is not extensional. The dual of a biextensional space however is biextensional. We therefore have two self-dual categories, big **Chu** and little **chu**, only the latter of which however is concrete.

4 Inherited Structure

From now on we shall treat only biextensional Chu spaces. Most of what we say here could be applied to general Chu spaces provided we substitute “a” for “the” and $\equiv$ (same representation) for $=$. However this complicates the story without any real gain. It would be like studying preordered sets instead of partially ordered sets, where the notion of *the* least upper bound of a subset must be generalized to that of a clique of least upper bounds, with no essential advantage to the elementary theory.

Up to this point we have dealt with A, X, and Σ as unstructured sets, but with the promise that by representing the elements of A as words we would somehow impose structure on A.

In the next section we shall define a general but extremely austere notion of structure. The austerity takes some getting used to, so by way of motivation we first consider in this section a less general but more familiar notion of structure which we shall show is preserved by continuous functions. We illustrate this with three examples, pointed sets, posets, and semilattices.

The class of all relational structures of a given signature is closed under both arbitrary product and substructure. Hence if we equip the unstructured set Σ with a relational structure, chosen completely arbitrarily, then this structure is inherited by Σ^X, and thence by any subset $\mathcal{A}$ thereof.

For example if we order the set $2 = \{0, 1\}$ so that $0 \leq 1$, then this order is inherited by every Chu space over $\mathcal{A}$ in the usual way: $a \leq b$ just when for all $x \in X$, $r_a(x) \leq r_b(x)$. This is just the ordinary inclusion order on 2^X when understood as the power set of X.

What we shall show is that continuous functions preserve the structure inherited from Σ. What is remarkable about this is that *the notion of continuity was defined without reference to that structure,* regardless of how rich or combinatorially complex that structure might be.

Let us begin with the simple case of declaring one element of Σ, say 0, to be a constant (definable as a unary relation holding just at 0), making Σ a pointed set. Pointed sets transform like ordinary sets except for the requirement that the constant in the source be mapped to the constant in the target. Now consider Σ^X for any set X. The induced constant in Σ^X is the constant word $00\ldots0$. This is as in universal algebra, where the n-th direct power of an algebra with a constant contains that constant as a constant n-tuple.

Now suppose $00\ldots0$ appears in $\mathcal{A}$, call it $\mathbf{0}$. Then every projection of $\mathcal{A}$ onto Σ sends $\mathbf{0}$ to 0. Let $f : \mathcal{A} \to \mathcal{B}$ be continuous. Then every projection of $\mathcal{B}$ to Σ must send $f(\mathbf{0})$ to 0, or we would have a projection of $\mathcal{B}$ whose composition with f fails to be a projection of $\mathcal{A}$. Hence $f(\mathbf{0})$ must itself be the constantly zero word.

But this behavior of f must hold even if we do not equip Σ with any structure, because the definition of continuous function is independent of any structure we might assign to Σ. That is, constant words must be preserved by continuous functions: not only must $\mathbf{0}$ always be mapped to $\mathbf{0}$ (of the appropriate length) but $\mathbf{1}$ must be mapped to $\mathbf{1}$ and so on.

We now pass from Σ as a structure with one or more constants to the above example $\Sigma = 2$ ordered by $0 \leq 1$. Any Chu space (A, r, X) over this alphabet, viewed as a subset $A \subseteq 2^X$ of the power set of X, becomes a set of subsets of X, with the inherited order being the ordinary inclusion order.

Proposition 2 *All continuous functions between two Chu spaces over 2 are monotone with respect to the inclusion order.*

PROOF. Suppose to the contrary that $a \leq b$ in $\mathcal{A}$ but the continuous function $f : \mathcal{A} \to \mathcal{B}$ is such that $f(a) \not\leq f(b)$. Hence for some projection $\pi : \mathcal{B} \to \Sigma$, $\pi(f(a)) = 1$ but $\pi(f(b)) = 0$. But $\pi \circ f$ must be a projection of $\mathcal{A}$, whose just-observed behavior contradicts $a \leq b$.

For our third and last example, furnish $\Sigma = 2$ with the usual join or disjunction operation $\vee : 2^2 \to 2$. This induces a partial join operation on $\mathcal{A} \subseteq \Sigma^X$ defined as the join (bitwise OR) of words of $\mathcal{A}$. If for any two words a, b of $\mathcal{A}$ that join is itself a word of $\mathcal{A}$, call it $a \vee b$, then we say that $\vee$ is defined at (a, b).

Proposition 3 *Whenever $a \vee b$ is defined in $\mathcal{A}$ and $f : \mathcal{A} \to \mathcal{B}$ is continuous, then $f(a) \vee f(b)$ is defined in $\mathcal{B}$ and equals $f(a \vee b)$.*

PROOF. Suppose not. Then there must exist a projection $\pi : \mathcal{B} \to \Sigma$ with $\pi(f(a \vee b)) \neq \pi(f(a)) \vee \pi(f(b))$. But as before, $\pi \circ f$ must be a projection of $\mathcal{A}$, giving us a position (element $x \in X$) at which the alleged $a \vee b$ fails to be the join of a and b after all.

Common to these proofs is the existence of a refuting state of $\mathcal{B}$, which is turned into a refuting state of $\mathcal{A}$ by composing it with f. This technique generalizes in the now obvious way to any relational structure we might impose on Σ.

Now one might deduce from all this that in order to preserve every possible relational structure we can think of for Σ, continuous functions must be terribly constrained. In some situations this is indeed the case: for example if $\mathcal{A}$ contains a constant word and $\mathcal{B}$ does not contain the same constant, there is *no* continuous function from $\mathcal{A}$ to $\mathcal{B}$.

But in some situations *every* function from $\mathcal{A}$ to $\mathcal{B}$ can be continuous. This happens in particular when $\mathcal{A}$ is the dual of the Chu space Σ^X consisting of all words of length X. In this case a projection from $\mathcal{A}$ to Σ is just a word of Σ^X, of length X. But Σ^X contains *every* possible word of length X. Hence the test as to whether $\pi \circ f$ is a projection of $\mathcal{A}$ must always succeed, whence every function from $\mathcal{A}$ to any Chu space whatsoever will be continuous. Such an $\mathcal{A}$ therefore behaves like a pure set, an object devoid of structure, which we call a *discrete* Chu space. This is the Chu counterpart of a discrete topological space.

This way of constructing discrete spaces ensures that no matter what relational structure we equip Σ with, and hence Σ^X, all of this structure may be eliminated by shrinking $A \subset \Sigma^X$ enough, namely to size double-logarithmic (with base $|\Sigma|$) in the size of Σ^X (since $\mathcal{A}$ is discrete when X is Σ^A).

So just how much relational structure *can* we equip Σ with? The answer is that it suffices to take all possible relations of all possible arities on Σ, including infinite arities. For arity X (as a set of variables in preference to just a number) these form the set 2^{Σ^X} of all X-ary relations on Σ. One set X of each cardinality suffices. Most of these relations will be obtainable by composition from a very small basis. When $\Sigma = 2$, X-ary relations on 2 for finite X are just X-ary Boolean operations, for which a sufficient basis is implication and the constant 0. More generally, with Σ still 2, the maximal structure with which we may equip Σ^X is that of a complete atomic Boolean algebra. If we ignore questions of concreteness, this remains true for larger Σ [Pra95].

5 Abstract Structure

In this section we shall identify the states of any extensional Chu space (A, r, X) with the dual words representing them. Thus instead of an extensional Chu space being (A, r, X) with $\check{r} : X \to \Sigma^A$ being injective, we have just (A, X) with $r(a, x)$ being defined implicitly as $x(a)$ (application of $x : A \to \Sigma$ to $a \in A$).

Define a *property* of a Chu space (A, X) to be any superset of X that is a subset of Σ^A. The discrete Chu space Σ^A has only the one property, namely itself, which we identify with the vacuous property *true*. In general the properties of $\mathcal{A}$ form the set $2^{(\Sigma^A - X)}$, namely all ways of adding new columns to $\mathcal{A}$.

Given two sets A and B, a function $f : A \to B$ induces a map $\hat{f} : 2^{\Sigma^A} \to 2^{\Sigma^B}$ of properties sending the property $Z \subseteq \Sigma^A$ to the property $\{h \in \Sigma^B \mid h \circ f \ \in \ Z\}$. This is true independently of the choice of $X \subseteq \Sigma^A$ making A into a Chu space (A, X). When $\hat{f}$ sends every property of $\mathcal{A}$ to a property of $\mathcal{B}$, i.e. when $Z \supseteq$ implies $\hat{f}(Z) \supseteq Y$, we call f a *homomorphism* of Chu spaces. The following is easily seen.

Proposition 4 *A function is a homomorphism of Chu spaces if and only if it is continuous.*

PROOF. If f is a homomorphism from (A, X) to (B, Y) then it sends the property X to a superset of Y. But this is equivalent to saying that for every column $y \in Y$, $y \circ f$ is in X, whence f is continuous.

Conversely, suppose f is continuous. Let $Z \supseteq X$ be a property of $\mathcal{A}$. We wish to show that $\{y \mid y \circ f \in Z\}$ is a superset of Y. But this is just the requirement that every $y \in Y$ satisfies $y \circ f \in Z$. Since f is continuous we have the stronger property that $y \circ f \in X \subseteq Z$.

This notion of property is quite abstract, and it will therefore be helpful to develop some intuition about it to demonstrate its relationship to more conventional notions of property.

To begin with, its generality notwithstanding, this notion of property is internal to the Chu space being described, since it is defined in terms of a fixed carrier A and alphabet Σ. We cannot use it to express properties that refer to other Chu spaces, such as the property of being the smallest Chu space meeting some condition.

Some intuition for the scope of this concept of property can be built up as follows. We call a property *atomic* when it excludes exactly one state. The intersection of two atomic properties excludes both their respective states and thus corresponds to conjunction, yielding a compound (non-atomic) property. By allowing arbitrary conjunction we can express any property of a Chu space $\mathcal{A}$ as a conjunction of its atomic properties; in particular $\mathcal{A}$ itself can be expressed as the conjunction of all its atomic properties, namely those states absent from $\mathcal{A}$.

When $\Sigma = 2$, a natural presentation of an atomic property is as $\Gamma \vdash \Delta$ where Γ is a list of the points projected by the property to 1 and Δ lists the remaining points, those projected to 0. If we interpret points as propositions that may be either true (1) or false (0), and take Γ to be the conjunction of its members and Δ the disjunction, then $\Gamma \vdash \Delta$ is the logical expression of this atomic property, since it is false in the state excluded by that property and true in all other states.

This language generalizes to larger alphabets by providing a region for each letter. With three letters one could write something like $\Gamma_0 \vdash \Gamma_1 \vdash \Gamma_2$, where every letter appears in exactly one Γ_i, and describes the state in which each letter in Γ_i has value i.

Every property of a Chu space $\mathcal{A}$ can be expressed as a conjunction of atomic properties of $\mathcal{A}$, which therefore can serve as the constants of a description language for $\mathcal{A}$. But if we insist on atomicity for the constants of our language, it may be unnecessarily rich.

Consider a Chu space having $|\Sigma|$ atomic properties whose excluded states differ from each other only at one element a, which is assigned a different letter of Σ in each state. Instead of listing all $|\Sigma|$ of these properties, we could simply take one of them and drop a from it, which would say the same thing. More generally if the space had Σ^n atomic properties which collectively were independent of n points in the same way, we could again simply take one of those properties and drop any mention of all n of those points.

These considerations lead to the notion of an *axiomatization* as a set of properties, not necessarily atomic, whose conjunction is the property X of being the Chu space (A, X).

We can take this notion a step further by defining a *language* for a set A to consist of a set of possible properties of A. Each such property is some subset of Σ^X. The purpose of such a language is to specify Chu spaces in a more conventional way than simply by giving the whole matrix, namely as the conjunction of the available properties.

For example, take $\Sigma = 2$ and regard states as subsets of A. We can specify those Chu spaces representing a partial ordering of the set A by taking the language to consist of one property for each pair (a, b) in A^2. The property associated with (a, b) would exclude all states containing a but not b, and hence express the relationship $a \leq b$. A partial order can now be defined on A by forming the conjunction of those $a \leq b$ properties sufficient to express that order.

The advantage of this particular language is that each property $a \leq b$ can be named by naming just a and b, which for large A requires many fewer bits of information than needed to identify even one state let alone the many needed to identify the whole partial order.

For another example, still with $\Sigma = 2$, take the language to consist of one property for each triple (a, b, c), namely the property whose states are all those satisfying $a \vee b = c$. With this language we can equip A with the structure of a join-semilattice by listing one such property for each pair a, b in A, choosing c to be whatever the join of a and b happens to be in that semilattice.

If we list only a subset of those properties then we will have specified a partial join-semilattice, one for which the join operation is defined only for some pairs. Continuous functions from such a structure will preserve only those joins that exist. In particular we can specify an arbitrary partial order by listing only those triples (a, b, c) such that a and b are comparable in that order, in which case c will always be the larger of a and b, being their improper join. Continuous functions from such a structure will then preserve only the partial order, that is, they will simply be monotone functions. Even if two incomparable elements happen to have a least upper bound in that partial order, continuous functions need not preserve that least upper bound, i.e. they will not recognize least upper bounds as being "part of the signature."

6 Representing general structures with Chu spaces

A natural question to ask is, how general are the ideas of the preceding section? We have given two answers to this question elsewhere, one measuring the generality of Chu spaces in terms of arbitrary relational structures and their homomorphisms, of which the foregoing are examples, the other in terms of arbitrary small categories, for which a forgetful functor may or may not be given.

For all these situations we have shown [Pra93,Pra96], that the objects in question can be represented by Chu spaces in such a way that the continuous functions between the representing Chu spaces are exactly the homomorphisms of the corresponding represented objects. More precisely, in each case the category of interest embeds fully and concretely in the category of Chu spaces over an appropriate alphabet.

Proposition 5 *The relational structures* (A, R) *of arity* n *embed fully and concretely in* $\mathbf{chu}_{2^n}$*.*

PROOF. (Outline) We first give the representation of (A, R) in the case when R consists of exactly one n-tuple $t \in A^n$. Take (A, r, X) to be the Chu space such that $X \subseteq (2^A)^n$ consists of those n-tuples $(x1, \ldots, x_n)$ of subsets $x_i \subseteq A$ such that there exists $i \leq n$ for which $t_i \in x_i$, and $r(a, x) = \{i | a \in x_i\}$.

We then generalize this representation to arbitrary R by intersecting the X's that arise as above for each n-tuple in R, and restricting r to A times that intersection. One can then show that a function from (A, R) to (B, S) is a homomorphism if and only if it is a continuous function between the respective representing Chu spaces [Pra93].

The construction generalizes to multiple relations by combining them into one super-relation with arity the sum of its constituent arities. It further generalizes to multiple sorts by combining them into one sort with disjoint union and indicating the sorts with one unary relation per sort.

Structures so representable include groups, transforming standardly by group homomorphisms ($|\Sigma| = 8$ because the group multiplication relation $ab = c$ is ternary), directed graphs (or binary relations) transforming by graph homomorphisms ($|\Sigma| = 4$), small categories transforming by functors ($|\Sigma| = 16$ because we need a ternary relation for composition and also a unary relation to mark the identities), and so on.

The corresponding theorem for categories is as follows.

Proposition 6 *Every small category* C *embeds fully and concretely in the category of Chu spaces over the set* $ar(C)$ *of arrows of* C*.*

(In the absence of an explicitly given forgetful functor $U : C \to \mathbf{Set}$ we take that given by the Yoneda embedding. On the one hand this is the standard way to make any small category concrete, on the other it is not a very economical way so the result is not as useful as the theorem following this.)

PROOF. Represent each object b of C as the Chu space $F(b)$ whose points are the morphisms $f : a \to b$ from any object a of C, whose states are the morphisms $h : b \to c$ to any object c of C, and for which $r(f, h)$ is defined as the composite $f; h$ $(= h \circ f)$. Represent each morphism $g : b \to b'$ of C as the continuous function $F(g) = \lambda f.f;g$ (which maps points of $F(b)$ to points of $F(b'))$, with adjoint $\lambda h.g;h$ (which maps states of $F(b')$ to states of $F(b)$).

That this is indeed the adjoint follows immediately from $(f;g);h = f;(g;h)$, associativity of composition. That F is full and faithful is a nice exercise (or see [Pra96]). That F is concrete (commutes with the forgetful functor to **Set**) holds when we take as the forgetful functor for small categories the functor that sends each object to the set of all morphisms to that object, which is just the covariant half of our Chu representation.

A category admitting a full embedding of any small category is called *universal* [PT80,Trn66,HL69]. Unlike all previous universal categories however, Chu is *concretely* universal in the sense that the embedding preserves the carrier. This is the case both for this representation of objects of an arbitrary small category and for the preceding representation of relational structures (which normally form large categories).

A variant of this categorical embedding holds for small concrete categories (C, U) where $U : C \to \mathbf{Set}$ is a given faithful functor. The advantage of this theorem over the preceding one is that it is concrete with respect to the given forgetful functor, not with respect to one we make up for the occasion. For it to go through we require one minor additional condition on (C, U), namely that if $U(a) = \emptyset$ then for all objects b, there must exist a morphism $a \to b$, necessarily just one since U is faithful.

Denote by $Elt(C, U)$ the disjoint union of the $U(a)$'s over all objects a of C, itself a set since C is small. This is the set of all occurrences of elements in the underlying sets of the objects of C.

Proposition 7 *Every small honest concrete category (C, U) embeds fully, faithfully, and concretely in $\mathbf{chu}_{Elt(C,U)}$.*

PROOF. (Outline) Represent object a of C by the Chu space $(U(a), r, X)$ where X is the set of all morphisms $x : a \to b$ of C with domain a, and for each element u of $U(a)$, $r(u, x) = U(x)(u)$, the image of element u under the representation of x. It is then straightforward to show that this representation is faithful, full, and concrete.

7 Linear Logic

Chu spaces and their continuous functions form a *-autonomous category [Bar79], one that is self-dual and symmetric monoidal closed. This places them squarely in the domain of discourse of Girard's linear logic [Bar91,Gir87]. In this section we define the multiplicative fragment of linear logic from the representational viewpoint, in which $\mathcal{A} \multimap \mathcal{B}$ in particular has a strikingly natural definition.

In ordinary logic the logical connectives denote operations on truth values. Linear logic however makes the most sense when understood as a categorical logic, whose connectives are functors rather than mere operations.

The language of multiplicative linear logic consists of propositional variables $P, Q, \ldots$ and logical connectives $A^\perp$, $A \multimap B$, $A \otimes B$, and $A \parr B$. In the following, $A, B, \ldots$ denote arbitrary formulas.

In the Chu space interpretation of linear logic, all formulas are interpreted as Chu spaces.

The first operation of linear logic is linear negation, $A^\perp$. This is interpreted simply as the dual $\mathcal{A}^\perp$ of the Chu space $\mathcal{A}$ interpreting A.

The operation $A \multimap B$ is interpreted as the Chu space consisting of the representations of the continuous functions from the interpretation $\mathcal{A} = (A, r, X)$ of A to the interpretation $\mathcal{B} = (B, s, Y)$ of B. Recall from section 3 the Chu space $\mathcal{F}_f = (A, \varphi_f, Y)$ representing $f : \mathcal{A} \to \mathcal{B}$. Let $\mathcal{B}^{\mathcal{A}}$ denote the set of continuous functions from $\mathcal{A}$ to $\mathcal{B}$. Then $\mathcal{A} \multimap \mathcal{B}$ is defined as the Chu space $(\mathcal{B}^{\mathcal{A}}, r, A \times Y)$ consisting of the representations of those functions viewed as words of length $A \times Y$, where $r(f, (a, y)) = \varphi_f(a, y) = s(f(a), y)$.

This defines the behavior of $\multimap$ on Chu spaces. Its behavior on morphisms is defined so as to give this operation the characteristics of an internal hom-functor. Given continuous functions $f : \mathcal{A}' \to \mathcal{A}$ and $h : \mathcal{B} \to \mathcal{B}'$, we take $f \multimap h$ to be the continuous function from $\mathcal{A} \multimap \mathcal{B}$ to $\mathcal{A}' \multimap \mathcal{B}'$ defined by $\lambda g.hgf$, sending the Chu transform $g : \mathcal{A} \to \mathcal{B}$ (where $\mathcal{A} \multimap \mathcal{B} = (\mathcal{B}^{\mathcal{A}}, t, A \times Y)$) to the composite $hgf : \mathcal{A}' \multimap \mathcal{B}'$ (where $\mathcal{A}' \multimap \mathcal{B}' = (\mathcal{B}'^{\mathcal{A}}, t', A' \times Y')$). Since Chu transforms are closed under composition hgf is a Chu transform, whence $f \multimap h$ is well-defined.

It remains to show that $f \multimap h$ is continuous. Evidently its adjoint should map $(a', y') \in A' \times Y'$ to $(f(a'), h^\perp(y')) \in A \times Y$ where $h^\perp : Y' \to Y$ is the adjoint of h. Instantiating the adjointness condition calls for the equality of $t'(\lambda g.hgf(g), (a', y'))$ (namely $t(hgf, (a', y'))$) and $t(g, (f(a'), h^\perp(y'))$. The definition of t gives the left hand side as $S'(hgf(a'), y')$ and the right hand side as $S(g(f(a')), h^\perp(y'))$, but these are equal because h is a Chu transform, and we are done.

Having defined $A \multimap B$ as a functor it is now easy to define $A \otimes B$, namely as

$(A \multimap B^{\perp})^{\perp}$, and $A \parr B$, as $A^{\perp} \multimap B$.

In [Pra97] we make a start on the project of proving full completeness of linear logic interpreted over Chu spaces, by showing that for formulas of multiplicative linear logic having at most two occurrences of each variable, cut-free proof-nets of those formulas are in bijective correspondence with the dinatural elements of the corresponding functors. This shows that, up to this case, dinaturality exactly captures proof on the semantic side of the mathematical ledger.

In more recent work this summer with Gordon Plotkin we have shown that dinaturality is not sufficient for formulas with four occurrences. In particular there are at least four dinaturals from $A \multimap A$ to itself. However by strengthening the naturality condition to logicality, i.e. invariance under logical relations, we believe we can extend full completeness to the whole of MLL.

8 Conclusion

We have presented Chu spaces, as $A \times X$ matrices, from a perspective that emphasizes their rows as representing their elements. Elsewhere we have developed the analogy with topology by emphasizing their columns, which can be understood as generalized open sets.

The advantage of the row perspective is that topology is not as widely appreciated as it could be, with the result that most people find it easier to think about the elements of a set than about the open sets of a topological space. Not only are open sets relatively unfamiliar, but they also transform strangely, namely backwards.

Functions $f : \mathcal{A} \to \mathcal{B}$ are naturally presented in terms of elements, namely by listing the values of $f(a)$ for all elements $a \in A$. This is then an A-indexed list of values of $\mathcal{B}$ each represented in $\mathcal{B}$ as a word of length Y. A function is therefore representable by a word of length $A \times Y$, and the space $\mathcal{A} \multimap \mathcal{B}$ of all continuous functions from $\mathcal{A}$ to $\mathcal{B}$ is represented by a Chu space of width $A \times Y$. This is *all* the information we need to give about $\mathcal{A} \multimap \mathcal{B}$ in order to equip it with exactly the right structure for it to soundly interpret linear implication.

One topic we did not treat is the Stone gamut, described in detail elsewhere [Pra95]. The Stone gamut coordinatizes transformational mathematics, which we understand as dealing with transformable objects and its associated logic of transformation, defined in terms of (di)natural transformations. There are two dimensions. In the horizontal direction range from the discrete, namely sets, to the coherent, namely complete atomic Boolean algebras or dual sets. In the vertical direction objects are classified according to the Σ by which they are represented.

This picture locates sets at one "edge" of the mathematical universe, and dual sets at the other, with all other structures in between. One can think of sets as made up of discrete atoms or particles (in the abstract sense of being dual to waves rather than in the more concrete "particle zoo" sense of particle physics) and dual sets as consisting of coherent waves. In the middle are the *square* Chu spaces, *mens sana in corpore sano*, such as finite-dimensional vector spaces, Hilbert spaces (suitably transforming), complete semilattices (join or meet is immaterial), locally compact Abelian groups, and finite chains with bottom.

Abstracts of the above and other papers by the author and colleagues, together with links to their postscript versions, may be found on the web site http://boole.stanford.edu/chuguide.html.

References

[Bar79] M. Barr. *∗-Autonomous categories*, volume 752 of *Lecture Notes in Mathematics*. Springer-Verlag, 1979.

[Bar91] M. Barr. ∗-Autonomous categories and linear logic. *Math Structures in Comp. Sci.*, 1(2):159–178, 1991.

[BGdP91] C. Brown, D. Gurr, and V. de Paiva. A linear specification language for Petri nets. Technical Report DAIMI PB-363, Computer Science Department, Aarhus University, October 1991.

[Bla95] Andreas Blass. A category arising in linear logic, complexity theory, and set theory. In J.-Y. Girard, Y. Lafont, and L. Regnier, editors, *Advances in Linear Logic*, pages 61–81, Ithaca, NY, June 1995. Cambridge University Press.

[Gir87] J.-Y. Girard. Linear logic. *Theoretical Computer Science*, 50:1–102, 1987.

[GP93] V. Gupta and V.R. Pratt. Gates accept concurrent behavior. In *Proc. 34th Ann. IEEE Symp. on Foundations of Comp. Sci.*, pages 62–71, November 1993.

[HL69] Z. Hedrlín and J. Lambek. How comprehensive is the category of semigroups. *J. Algebra*, 11:195–212, 1969.

[Kel82] G.M. Kelly. *Basic Concepts of Enriched Category Theory: London Math. Soc. Lecture Notes*. 64. Cambridge University Press, 1982.

[LS91] Y. Lafont and T. Streicher. Games semantics for linear logic. In *Proc. 6th Annual IEEE Symp. on Logic in Computer Science*, pages 43–49, Amsterdam, July 1991.

[Mac45] G. Mackey. On infinite dimensional vector spaces. *Trans. Amer. Math. Soc.*, 57:155–207, 1945.

[Pra93] V.R. Pratt. The second calculus of binary relations. In *Proceedings of MFCS'93*, volume 711 of *Lecture Notes in Computer Science*, pages 142–155, Gdańsk, Poland, 1993. Springer-Verlag.

[Pra95] V.R. Pratt. The Stone gamut: A coordinatization of mathematics. In *Logic in Computer Science*, pages 444–454. IEEE Computer Society, June 1995.

[Pra96] V.R. Pratt. Broadening the denotational semantics of linear logic. In *Special Issue on Linear Logic 96*, volume 3 of *ENTCS (Electronic Notes of Theoretical Computer Science)*, Tokyo, 1996.

[Pra97] V.R. Pratt. Towards full completeness of the linear logic of chu spaces. In *Electronic Notes in Theoretical Computer Science*, volume 6, Pittsburgh, 1997. URL: http://www.elsevier.nl/locate/entcs/volume6.html, 18 pages.

[PT80] A. Pultr and V. Trnková. *Combinatorial, Algebraic and Topological Representations of Groups, Semigroups, and Categories*. North-Holland, 1980.

[Trn66] V. Trnková. Universal categories. *Comment. Math. Univ.Carolinae*, 7:143–206, 1966.