Topological nominal semantics

Murdoch J. Gabbay

Thanks to Dan Marsden and Luke Ong.

20 April 2018

Introduction

This talk is based on ideas from two papers:

- Representation and duality of the untyped lambda-calculus. Annals of Pure and Applied Logic.
  http://www.gabbay.org.uk/papers.html#repdul
  https://doi.org/10.1016/j.apal.2016.10.001
- Semantics out of context. Journal of the ACM.
  http://www.gabbay.org.uk/papers.html#semooc
  http://dx.doi.org/10.1145/2700819

There's well over a hundred pages of dense maths there. I regret this length, because they are trying to do something interesting.

I'll to try to communicate some sense of this.

Recall: Boolean algebras (BA)

A Boolean Algebra $\mathfrak{B}$ is an algebraic structure: a set, with operations $\bot$, $\neg$, $\wedge$ and reasonable equality axioms that say

- $\bot$ behaves like bottom,
- $\neg$ behaves like negation, and
- $\wedge$ behaves like conjunction.

We convert $\mathfrak{B}$ into a topological space $S(\mathfrak{B})$ as follows:

- Points of $S(\mathfrak{B})$ are ultrafilters $f \subseteq \mathfrak{B}$.
  (ultrafilter = maximal $\leq$-closed set with finite lower bounds; equivalently, BA homs to $\{\bot, \top\}$).
- The topology is generated by introducing $x^\bullet = \{f \mid x \in f\}$ for each $x \in \mathfrak{B}$ as clopen sets, and closing under topological operations.

Stone duality

Note the two maps:

$$x \in \mathfrak{B} \longmapsto x^\bullet = \{f \mid x \in f\} \in clopen(S(\mathfrak{B}))$$
$$C \in clopen(S(\mathfrak{B})) \longmapsto C_\bullet = \iota x.(x \in \bigcap C) \in \mathfrak{B}$$

The categories of Boolean Algebras and of Stone spaces (compact Hausdorff totally disconnected topological spaces) are dual.

- $\mathfrak{B}$ maps to $S(\mathfrak{B})$.
- A Stone space H maps to its lattice of clopen sets.

The maps above are units of the adjunction.

Thus are connected propositional logic, algebra, and topology.

Nominal semantics

My PhD of 2001 introduced nominal sets, also called sets with atoms and named sets.

These are simply 'sets with atoms'; definition follows shortly.

Nominal sets have applications including to inductive syntax with binding, to generalised finite automata, and to semantics of HoTT.

I believe nominal sets should be taught in mathematical foundational courses — because of these applications but also because of their applications to the foundations of logic and computation.

That's what my two papers cited above (and a few others) explore.

Nominal semantics

These are canonical foundations for logic and computation (note: names and binding feature prominently):

- First-order logic (FOL).
- The λ-calculus (λ-calculus).

The foundations of these systems were developed in a ZF universe.

Q. What happens if we develop foundations of logic and computation in a nominal sets universe, where names and binding are primitive, instead of in the ZF universe that they were historically developed in?

A. Interesting things.

We need some definitions.

Nominal sets definition

Fix a countably infinite set of atoms $\mathbb{A}$. Write $\Sigma_{\mathbb{A}}$ for the symmetry group of $\mathbb{A}$ (bijections π on $\mathbb{A}$; also called atoms-permutations).

A $\Sigma_{\mathbb{A}}$-set X is:

- An underlying set X, with
- a group action $\Sigma_{\mathbb{A}} \times X \rightarrow X$.

Example:

- $\mathbb{A}$ is a $\Sigma_{\mathbb{A}}$-set where $\pi{\cdot}a = \pi(a)$.
- Syntax of λ-terms over atoms as variable symbols is a $\Sigma_{\mathbb{A}}$-set where π acts pointwise on the atoms in a term. So $\pi{\cdot}(\lambda a.\lambda b.ab) = \lambda\pi(a).\lambda\pi(b).\pi(a)\pi(b)$.

Nominal sets definition

Note that if X is a $\Sigma_{\mathbb{A}}$-set then so is *power*(X) the powerset of X, by the pointwise action

$$\pi \cdot X = \{\pi \cdot x \mid x \in X\} \quad \text{where} \quad X \in \mathit{power}(\mathsf{X}).$$

Example if $a, b, c \in \mathbb{A}$ then:

$$\begin{aligned}
\pi \cdot \{a, b, c\} &= \{\pi(a), \pi(b), \pi(c)\} \\
\pi \cdot (\mathbb{A} \setminus \{a, b, c\}) &= \mathbb{A} \setminus \{\pi(a), \pi(b), \pi(c)\} \\
\pi \cdot \big\{\{a, b, c\}, \mathbb{A} \setminus \{a, b, c\}\big\} &= \\
&\big\{\{\pi(a), \pi(b), \pi(c)\}\ , \mathbb{A} \setminus \{\pi(a), \pi(b), \pi(c)\}\big\}
\end{aligned}$$

Nominal sets definition

If X is a $\Sigma_{\mathbb{A}}$-set and $x \in$ X and $X \subseteq_{fin}$ X (finite subset) then define:

$$\begin{aligned} fix(x) &= \{\pi{\in}\Sigma_{\mathbb{A}} \mid \pi{\cdot}x = x\} \\ stab(X) &= \textstyle\bigcap_{x\in X} fix(x) = \{\pi{\in}\Sigma_{\mathbb{A}} \mid \forall x{\in}X.\pi{\cdot}x = x\} \end{aligned}$$

Easy to check that if $A \subseteq_{fin} \mathbb{A}$ then:

$$stab(A) = \{\pi{\in}\Sigma_{\mathbb{A}} \mid \forall a{\in}A.\pi(a) = a\}.$$

Define $A\$x$ and say $A \subseteq \mathbb{A}$ supports $x \in$ X as follows:

$$\begin{array}{lll} A\$x & \text{when} & stab(A) \subseteq fix(x) \\ & \Leftrightarrow & \forall\pi{\in}\Sigma_{\mathbb{A}}.\big(\forall a{\in}A.\pi(a) = a\big) \Rightarrow \pi{\cdot}x = x. \end{array}$$

Nominal sets definition

A nominal set X is a $\Sigma_{\mathbb{A}}$-set such that every $x \in X$ has a unique least finite supporting set $supp(x) \subseteq_{fin} \mathbb{A}$.

Examples:

- $\mathbb{A}$ is a nominal set. If $a \in \mathbb{A}$ then $\{a\}\$a$.
- Finite sets of atoms are a nominal set. If $A \subseteq_{fin} \mathbb{A}$ then $A\$A$.
- Cofinite sets of atoms (complements of finite sets of atoms) are a nominal set. If $A \subseteq_{fin} \mathbb{A}$ then $A\$\, \mathbb{A} \setminus A$.
- Syntax is a nominal set (where variables symbols are atoms). t is supported by the atoms mentioned in t (free or bound).
- Syntax quotiented by α-equivalence is a nominal set (where variables symbols are atoms).
  $[t]_\alpha$ is supported by $fv(t)$. (The α-equivalence class $[t]_\alpha$ is supported by the free variables of t.)
- If X is a nominal set then $power_{fs}(X)$ the set of $X \subseteq X$ with finite support under the permutation action, is a nominal set.

Freshness

If $a \in \mathbb{A}$ define $a\#x$ and say a is fresh for x by:

$$a\#x \quad \text{when} \quad a \notin \text{supp}(x).$$

- If X is finite sets of atoms then $a\#X$ when $a \notin X$.
- If X is cofinite sets of atoms then $a\#X$ when $a \in X$.
- If X is syntax quotiented by α-equivalence then $a\#[t]_\alpha$ when $a \notin fv(t)$ (a is not free in t).

This concludes the definitions (for the moment).

Boolean algebras in a nominal context

Consider a Boolean algebra $\mathfrak{B}$ in nominal sets (i.e. a Boolean algebra whose underlying set happens to also be a nominal set).

To what extent might $\mathfrak{B}$ already model FOL?

If $x \in \mathfrak{B}$, we can define

$$\forall a.x = \bigvee \{x' \leq x \mid a \# x'\},$$

if this bound exists in $\mathfrak{B}$.

$\forall a.x$, if it exists, is the greatest element below x for which a is fresh.

Quantifiers in nominal BAs

If $x \in \mathfrak{B}$ then $\forall a.x$ is the greatest a-fresh lower bound for x.

To see this is natural, rewrite in natural deduction style:

$$\frac{x' \leq x \quad a\#x'}{x' \leq \forall a.x}$$

Compare with the forall-right intro-rule:

$$\frac{\Gamma \vdash \phi \quad (a \notin fv(\Gamma))}{\Gamma \vdash \forall a.\phi.}$$

They are the same!

So just the shift from ZF to nominal sets, lets us naturally express $\forall$-intro.

The research programme of my two papers

Our research programme is now as follows:

- Run the 'ultrafilters and clopen sets' idea for algebras for first-order logic and the λ-calculus . . . in nominal sets.
- Let topological structure tell us what the algebraic structure needs to be, and vice versa.

I did this. Huge amounts of structure just dumped themselves into my lap . . . whence the two big papers.

It was a spigot. The challenge was to sort through it and present it as accessibly as possible.

Substitution / σ-action

We can build a simple account of FOL just with what we already have, but it is better if we include substitution.

After all, FOL and λ-calculus have two name-related structures:

- binders $\forall$ and λ, and
- substitution $[a{\mapsto}u]$, which I will call a σ-action.

Substitution is not normally described as a binding structure, but it is, because of this α-equivalence property:

$$t[a{\mapsto}u] = ((b\ a)\cdot t)[b{\mapsto}u] \quad \text{usually written} \quad t[u/a] = t[b/a][u/b].$$

Here $(b\ a)$ is the permutation mapping b to a and vice versa, and $[b/a]$ is the renaming mapping a to b.

σ-algebra / terms

A termlike σ-algebra is a nominal set T along with operations

- $\partial : \mathbb{A} \to \mathsf{T}$ and
- $\mathsf{T} \times \mathbb{A} \times \mathsf{T} \to \mathsf{T}$ written $x[a{\mapsto}u]$,

satisfying the following equations:

$$
\begin{array}{lrcl}
(\sigma\#) & a\#x \Rightarrow \quad x[a{\mapsto}u] & = & x \\
(\sigma\alpha) & b\#x \Rightarrow \quad x[a{\mapsto}u] & = & ((b\ a)\cdot x)[b{\mapsto}u] \\
(\sigma\sigma) & a\#v \Rightarrow x[a{\mapsto}u][b{\mapsto}v] & = & x[b{\mapsto}v][a{\mapsto}u[b{\mapsto}v]] \\
(\sigma\mathbf{id}) & x[a{\mapsto}\partial a] & = & x \\
(\sigma\mathbf{a}) & \partial a[a{\mapsto}u] & = & u
\end{array}
$$

This is a nominal algebra that distils the essential features of a substitution action (can be made precise: see [capasn-jv]).

σ-algebra / terms

If you don't like the axioms above, just think of the canonical examples:

- Terms of FOL. $[a \mapsto u]$ is substitution and $\partial : \mathbb{A} \hookrightarrow \mathit{Terms}$ is just the variable symbol a considered as a term.
- Terms of λ-calculus. As for terms of FOL.
- $\mathbb{A}$ itself. $[a \mapsto b]$ is the natural replacement map and $\partial : \mathbb{A} \to \mathbb{A}$ is the identity.

σ-algebra future research

We can fine-tune these axioms:

$$
\begin{array}{llrcl}
(\sigma\#) & a\#x \Rightarrow & x[a\mapsto u] & = & x \\
(\sigma\alpha) & b\#x \Rightarrow & x[a\mapsto u] & = & ((b\ a)\cdot x)[b\mapsto u] \\
(\sigma\sigma) & a\#v \Rightarrow & x[a\mapsto u][b\mapsto v] & = & x[b\mapsto v][a\mapsto u[b\mapsto v]] \\
(\sigma\mathbf{id}) & & x[a\mapsto \partial a] & = & x \\
(\sigma\mathbf{a}) & & \partial a[a\mapsto u] & = & u
\end{array}
$$

- Drop $(\sigma\alpha)$ to obtain a kind of 'fusion' operation: without $(\sigma\alpha)$, $[a\mapsto u]$ binds u to a but leaves the 'port' a open for further bindings.
- Drop $(\sigma\#)$ to permit communication on any port.
- Drop ∂, $(\sigma\mathbf{a})$, and $(\sigma\mathbf{id})$ if e.g. we only want to substitute for closed terms.

σ-algebra / predicates

A σ-algebra is:

- a nominal set X along with
- a σ-action $\mathsf{X} \times \mathbb{A} \times \mathsf{T} \rightarrow \mathsf{X}$ over a termlike σ-algebra T.

Canonical example:

- FOL predicates with their substitution action over terms, as in $\phi[a \mapsto u]$.

Dualise to σ^*-algebra

Consider a σ-algebra X over a termlike σ-algebra T.

What structure does the powerset *power*(X) have?

We obtain an algebra dual to σ-algebras; I call it amgis-algebra in my papers; here I will call it a σ^*-algebra. It is axiomatisable:

$$\begin{array}{ll} (\sigma^*\sigma) & a\#v \Rightarrow X[b{\mapsto}v]^*[a{\mapsto}u]^* = X[a{\mapsto}u[b{\mapsto}v]]^*[b{\mapsto}v]^* \\ (\sigma^*\mathbf{id}) & X[a{\mapsto}\partial a]^* = X \end{array}$$

This is the structure that *power*(X) naturally inherits from X, via the dualising action

$$X[a{\mapsto}u]^* = \{x \in \mathsf{X} \mid x[a{\mapsto}u] \in X\}.$$

Dualise to σ^*-algebra

Note that we consider the full powerset $\mathit{power}(\mathsf{X})$, not the finitely-supported powerset $\mathit{power}_{\mathit{fs}}(\mathsf{X})$.

This is important: intuitively $X \in \mathit{power}(\mathsf{X})$ may be a point (e.g. a maximally consistent theory). It need not be finitely supported.

(For specific applications, points may be subject to further coherence conditions; e.g. $\leq$-closure if X is a poset. A lot of technical action is here.)

In summary:

- If X is a σ-algebra with $(\sigma\#)$, $(\sigma\alpha)$, $(\sigma\sigma)$, $(\sigma\mathbf{id})$, and $(\sigma\mathbf{a})$, then
- $\mathit{power}(\mathsf{X})$ the set of all subsets of X is naturally a σ^*-algebra with $(\sigma^*\sigma)$ and $(\sigma^*\mathbf{id})$, where $x \in X[a{\mapsto}u]^* \Leftrightarrow x[a{\mapsto}u] \in X$.

Obviously we're going to do this again, moving from a σ^*-algebra $Y \in \mathsf{Y}$ to $\mathcal{X} \in \mathit{power}_{\mathit{fs}}(\mathsf{X})$.

But give me a moment first:

Are you still with me?

Let's orient ourselves.

Imagine you're Tarski: you're in ZF and you want to give semantics to FOL over (for simplicity) the trivial term language $\mathbb{A}$. Then you do the following:

- Choose a domain of individuals I.
- Call $\rho \in \mathbb{A} \to \mathsf{I}$ valuations.
- Build semantics in the type

$$(\mathbb{A} \to \mathsf{I}) \to \{\bot, \top\}.$$

This only seems simple because it's familiar! It's not simple at all. This is simpler:

- Build semantics in $\mathit{power}_{\mathit{fs}}(\mathit{power}(\mathbb{A}))$.

Are you still with me?

We perceive $(\mathbb{A} \to \mathsf{I}) \to \{\bot, \top\}$ as a natural place for FOL semantics only because we've been brainwashed by our ZF universe.

To be clear, I'm not saying Tarski semantics is bad.

On the contrary it's very good. I'm just pointing out that

$$\mathit{power}_{\mathit{fs}}(\mathit{power}(\mathbb{A}))$$

— in words: the nominal powerset of the powerset of atoms — is incredibly natural.

The NEW quantifier

Write $\text{И} b.\phi(b)$ for the following assertion:

- $\{b \in \mathbb{A} \mid \phi(b)\}$ is cofinite, or equivalently
- $\{b \in \mathbb{A} \mid \neg\phi(b)\}$ is finite.

The natural dualisation of a σ^*-algebra Y to a σ-algebra is this:

$$Y \in \mathcal{X}[a \mapsto u] \Leftrightarrow Y[a \mapsto u]^* \in \mathcal{X}.$$

But this doesn't give us $(\sigma\alpha)$. Recall:

$$\begin{array}{ll} (\sigma\#) & a\#x \Rightarrow x[a \mapsto u] = x \\ (\sigma\alpha) & b\#x \Rightarrow x[a \mapsto u] = ((b\ a)\cdot x)[b \mapsto u] \end{array}$$

Obtaining σ from σ^*

We need a cleverer definition:

$$Y \in \mathcal{X}[a\mapsto u] \quad \Leftrightarrow \quad \text{И}b.(Y[b\mapsto u]^* \in (b\ a)\cdot\mathcal{X}).$$

This builds in that $b\#\mathcal{X}$ implies $\mathcal{X}[a\mapsto u] = ((b\ a)\cdot\mathcal{X})[b\mapsto u]$.

Then given a σ^*-algebra Y, the set of $\mathcal{X} \subseteq$ Y such that

1. $\mathcal{X}$ is finitely supported (to allow the definition above to make sense so we can have $(\sigma\alpha)$), and
2. If $a\#\mathcal{X}$ then $\mathcal{X}[a\mapsto u] = \mathcal{X}$ (to get $(\sigma\#)$)

naturally forms a σ-algebra.

If you're worried that conditions 1 and 2 above seem artifical; actually they're just right, and I did the legwork to prove it by establishing the topological duality result.

($\forall$**L**) and ($\forall$**R**)

What does $\forall$ look like in $power_{fs}(power(\mathbb{A}))$?

By definition,

$$\forall a.\mathcal{X} = \bigcup\{\mathcal{X}' \subseteq \mathcal{X} \mid a\#\mathcal{X}'\}.$$

We already remarked that this builds in $\forall$-right-intro, in the sense that by construction

$$\frac{\mathcal{X}' \subseteq \mathcal{X} \quad a\#\mathcal{X}'}{\mathcal{X}' \subseteq \forall a.\mathcal{X}} \quad \text{which is patently this} \quad \frac{\Gamma \vdash \phi \quad (a \notin fv(\Gamma))}{\Gamma \vdash \forall a.\phi}.$$

What about ($\forall$**R**)? Let's look at the right elim rule:

$$\frac{\Gamma \vdash \forall a.\phi}{\Gamma \vdash \phi[a{:=}s]}$$

($\forall$**L**) and ($\forall$**R**)

Now by condition 2 above, if $a\#\mathcal{X}$ then $\mathcal{X}[a\mapsto u] = \mathcal{X}$. Also by construction $a\#\forall a.\mathcal{X}$.

It is a fact that by construction $[a\mapsto u]$ is monotone, and since $\forall a.\mathcal{X} \subseteq \mathcal{X}$, we have

$$\forall a.\mathcal{X} = (\forall a.\mathcal{X})[a\mapsto u] \subseteq \mathcal{X}[a\mapsto u].$$

This is patently $\forall$-elim. So:

1. condition 1 above = $(\sigma\alpha)$ = $\forall$-intro, and
2. condition 2 above = $(\sigma\#)$ = $\forall$-elim!

Equality

We get a semantics for equality for free. Define

$$a{=}a' \quad = \quad \{X \in power(\mathbb{A}) \mid \text{И}b.X[b{\mapsto}a]^* = X[b{\mapsto}a']^*\}.$$

Run the calculations and you'll see the FOL equality rules drop out.

Summary (FOL)

We obtain the following schema for FOL semantics, where X is any σ-algebra:

$$\mathsf{X} \quad \to \quad \mathit{power}(\mathsf{X}) \quad \to \quad \mathit{power}_{\mathit{fs}}(\mathit{power}(\mathsf{X})).$$

This schematic view can be refined to a topological duality.

There is scope to tweak and experiment. We could consider other logics; I suspect it will prove interesting to vary the structure of X. For example:

- We could drop $(\sigma\#)$. That would simplify the development (but 'substitution' wouldn't be complete for substitution).
- We could drop $(\sigma\mathbf{id})$.
- We could drop $(\sigma\alpha)$.

λ-calculus

The case of the λ-calculus is more complicated, but I can give an overview.

Start from a termlike σ-algebra T that is a magma: it has an additional binary operation

$$\bullet : \mathsf{T} \times \mathsf{T} \to \mathsf{T}.$$

This should commute with the σ-action but no other special axioms are required; in particular we do not require S and K combinators.

We follow the overall schema

$$\mathsf{T} \quad \to \quad \mathit{power}(\mathsf{T}) \quad \to \quad \mathit{power}_{fs}(\mathit{power}(\mathsf{T}))$$

and study the structure in $\mathit{power}_{fs}(\mathit{power}(\mathsf{T}))$.

λ-calculus

Given $\partial a \in \mathsf{T}$ (remember T is termlike so $\partial \in \mathbb{A} \hookrightarrow \mathsf{T}$) we can form

$$(\partial a)^\bullet = \{X \in \mathit{powerset}(\mathsf{T}) \mid a \in X\}.$$

Also we can extend $\bullet$ up the sets hierarchy in a natural way.

We define a right adjoint $-\!\!\bullet$ to $\bullet$ by

$$\mathcal{Z} \subseteq \mathcal{Y} -\!\!\bullet \mathcal{X} \Leftrightarrow \mathcal{Z} \bullet \mathcal{Y} \subseteq \mathcal{X}.$$

Then λ is defined by

$$\lambda a.\mathcal{X} = \forall a.((\partial a)^\bullet -\!\!\bullet \mathcal{X}).$$

[repdul, Notation 10.2.1]

Unpacking the symbols, this gives us η-expansion and β-reduction [repdul, Proposition 10.2.4].

Conclusions

This is a different approach to semantics than Tarski semantics. In terms of elegance I feel it has a lot going for it.

The papers are long because we have 80 years of foundations to develop, literally in an alternative universe. There's a lot of ground to cover.

For me, a sufficient motivation for doing this is the sheer beauty of

$$\mathsf{X} \quad \rightarrow \quad \mathit{power}(\mathsf{X}) \quad \rightarrow \quad \mathit{power}_{\mathit{fs}}(\mathit{power}(\mathsf{X})),$$

and the pleasure of watching logical and computational structure emerge from the double powerset iteration.

Conclusions

But already, we go strictly beyond what is possible in ZF:

- It's quite beautiful.
- Duality results obtained as described; can't do that in ZF.
- The nominal semantics is canonical in the sense that we only insist on precisely the limits we need to model FOL or λ-calculus.
  Tarski semantics gives 'too many limits', in the sense that $(\mathbb{A} \to \mathsf{I}) \to \{\bot, \top\}$ has all limits, and not just the limits representing predicates.