Semantics of FOL and λ-calculus using nominal techniques

Murdoch J. Gabbay

Samson@60, Oxford

Thanks to Luke Ong, Prakash Pananganden, and Dave Corne

May 30, 2013

Thank you Samson for bringing us all together.

Nominal techniques

Nominal techniques are obtained by taking names seriously; specifically, using Fraenkel-Mostowski sets.

So we assume a symmetric class of atoms or urelemente. Use a permutative convention that $a, b, c, \ldots$ range over distinct atoms.

Set theorists: assume we are in the hierarchy of Fraenkel–Mostowski sets. Category theorists: assume we are in the Schanuel Topos.

Every element has an atoms–permutation action $\pi{\cdot}x$, and a supporting set of atoms supp(x). Write $a\#x$ when $a \notin \mathrm{supp(x)}$.

These are our assumptions.

Nominal techniques

We should take names seriously because they are everywhere, and they are elementary. More elementary than ‘function’. More elementary than ‘first-order predicate’.

So we should be able to build first-order logic and the λ-calculus. Just out of names.

Given a set X we make a Boolean algebra out of its powerset $pset(X)$. Refining this, we get Stone duality.

So powersets give us propositional logic.

Surely, nominal powersets (sets with names and binding) should give us predicate logic (propositions with names and binding).

Nominal algebras

You can build algebras over nominal sets. Specifically [capasn 2006,capasn-jv 2008] you can specify a nominal algebra for substitution—σ-algebras:

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

You know these axioms, as lemmas: ‘if $x \notin fv(t)$ then $t[s/x] = t$’ is a lemma of syntax; $(\sigma\#)$ abstracts this. Need ‘nominal’ for the freshness side-condition.

Validity of $(\sigma\sigma)$ on concrete syntax is often called the substitution lemma.

Example σ-algebras

The axioms were specified in a paper on nominal algebra in 2006.

Examples of σ-algebras include:

- first-order syntax with substitution,
- λ-calculus terms with capture-avoiding substitution,
- non-syntactic models including Tarski-style valuation semantics and (amazingly)
- any cumulative hierarchy model of FM sets [stusun 2009].

The axioms are sound and complete for each of the classes of models above. So there are lots of models of these things, of greatly differing complexities.

Pick one, call it $\mathcal{X}$.

Nominal power sets

Given a σ-algebra $\mathcal{X}$ its nominal powerset $\mathit{powerset}(\mathcal{X})$ has a dual structure which I call an $\rotatebox[origin=c]{180}{σ}$-algebra (amgis-algebra):

$$\begin{array}{lrrcl}
(\rotatebox[origin=c]{180}{σ}\mathbf{id}) & & p[a\mapsfrom a] & = & p \\
(\rotatebox[origin=c]{180}{σ}\alpha) & b\#p \Rightarrow & p[u\mapsfrom a] & = & (b\ a)\cdot(p[u\mapsfrom b]) \\
(\rotatebox[origin=c]{180}{σ}\sigma) & a\#v \Rightarrow & p[v\mapsfrom b][u\mapsfrom a] & = & p[u[b\mapsto v]\mapsfrom a][v\mapsfrom b]
\end{array}$$

This looks a bit like a σ-algebra, but the axioms are all 'inside out'—as one might expect.

$$p[u\mapsfrom a] = \{x \in \mathcal{X} \mid x[a\mapsto u] \in p\} \qquad p \in \mathit{pset}(\mathcal{X})$$

So $(\rotatebox[origin=c]{180}{σ}\sigma)$ comes about since $x[a\mapsto u][b\mapsto v] \in p$ if and only if $x[b\mapsto v][a\mapsto u[b\mapsto v]] \in p$ by $(\sigma\sigma)$.

Nominal power sets

Take powersets again, and you get back a σ-algebra.

So if $\mathcal{X}$ is a σ-algebra, so is *powerset*(*powerset*($\mathcal{X}$)).

So far so good.

This is not as easy as I might make it sound. Here is the definition of the σ-action from the λ-calculus paper:

$$X[a{\mapsto}u] = \{p \mid \text{И}c.p[u{\mapsfrom}c] \in (c\ a)\cdot X\} \quad X \in pset(pset(\mathcal{X}))$$

Check out the 'nominal' stuff going on here: И, and permutation on sets of sets of sets.

But what you end up with in the end is a σ-algebra.

Nominal powersets

Powersets are lattices.

We can interpret $\bot$ and $\wedge$ in $\mathit{pset}(\mathit{pset}(\mathcal{X}))$ as the empty set and intersection.

We can interpret negation as complement.

Easy. Well known. Standard.

Fresh limits

New concept: nominal powersets have fresh-finite limits [nomspl 2012].

This means that given $X \in \mathit{pset}(\mathit{pset}(\mathcal{X}))$, $\forall a.X$ is the greatest subset of X such that $a\#\forall a.X$. So $\forall a.X$ is the $a\#$limit of $\{X\}$.

In the presence of the aforementioned σ-action, this coincides with the intersection of $X[a{\mapsto}u]$ for all u.

Equality

However, the definitions are not the same. The proofs work using the $a\#$limit characterisation, not the infinite-intersection characterisation.

(Why? Intuitively, $\bigwedge_u X[a{\mapsto}u]$ depends on the size of the set of u, whereas an $a\#$limit does not. Discuss here shades of ($\forall$**R**) rule.)

So we have $\bot$, $\wedge$, $\neg$, and $\forall$. In nominal powersets, we can interpret first-order logic.

Equality

We can go further and interpret $u = v$ as

$$\{p \in pset(\mathcal{X}) \mid \text{И}c.p[u \leftarrow\!\!\dashv c] = p[v \leftarrow\!\!\dashv c]\}.$$

И is the new-quantifier meaning ‘for some/any fresh c’.

So we have equality too. In nominal powersets, we can interpret first-order logic with equality!

References

I've simplified—a lot!

See "Stone duality for first-order logic" [stodfo 2011] and see "Semantics out of context" http://arxiv.org/abs/1305.6291 (submitted).

Warning: the papers are 32 and 56 pages long respectively. It's meaty stuff.

But what we get out of them is a comprehensive account of first-order logic in nominal sets: as a nominal algebra, a nominal lattice, and as a topological (a Stone) space, along with soundness, completeness, duality results, and translations of traditional Tarski and Herbrand models to the nominal framework.

The ideas are simple enough and are drawn directly from studying $\mathit{pset}(\mathit{pset}(\mathcal{X}))$ for a σ-algebra $\mathcal{X}$, as outlined. All we have done is take names seriously and use powersets.

λ-calculus

A similar story, only harder because λ is harder. See http://arxiv.org/abs/1305.5968 (86 pages!).

Assume that $\mathcal{X}$ has a combination action

$$\circ : \mathit{pset}(\mathcal{X}) \times \mathit{pset}(\mathcal{X}) \to \mathit{pset}(\mathit{pset}(\mathcal{X}))$$

(that's an odd type). Also assume that atoms are a subset of $\mathcal{X}$.

The combination action acts pointwise to give a binary application function on sets of sets $\mathit{pset}(\mathit{pset}(\mathcal{X}))^2 \to \mathit{pset}(\mathit{pset}(\mathcal{X}))$. This has a right adjoint $\multimap$.

λ-calculus

So given X and Y we can form $X \bullet Y$ and $Y \mathbin{-\!\bullet} X$, and $X \bullet Y \subseteq Z$ if and only if $X \subseteq Y \mathbin{-\!\bullet} Z$.

Then $\lambda a.X$ can be identified with $\forall a.(\partial a \mathbin{-\!\bullet} X)$, where $\partial a = \{p \mid a \in p\}$.

β-reduction and η-expansion emerge from the adjoint properties of $\bullet$ and $\mathbin{-\!\bullet}$:

- $(\partial a \mathbin{-\!\bullet} X) \bullet \partial a \subseteq X$ leads to β-reduction, and
- $X \subseteq \partial a \mathbin{-\!\bullet} (X \bullet \partial a)$ leads to η-expansion.

A flavour: the notion of filter used in both papers

A filter in $\mathcal{D}$ is a nonempty subset $p \subseteq |\mathcal{D}|$ (which need not have finite support) such that:

1. $\bot \notin p$.
2. If $x \in p$ and $x \leq x'$ then $x' \in p$.
3. If $x \in p$ and $x' \in p$ then $x \wedge x' \in p$.
4. If $\reflectbox{N}b.(b\ a){\cdot}x \in p$ then $\forall a.x \in p$.

A flavour: why the nominal models are not just ordinary models

Tarski-style models of FOL can be converted into corresponding nominal structures. I.e. valuation-based models have natural nominal and σ-algebra structure.

These models are complete; they have all limits, because the standard poset of Booleans $\{\bot, \top\}$ is complete and the Tarski denotation of a predicate ϕ is a function from valuations to $\{\bot, \top\}$.

The nominal models give ϕ a semantic in a nominal Stone space. Open and open compact sets are not closed under arbitrary intersections (i.e. do not have all limits).

They only have fresh-finite limits. Precise characterisation of FOL.

Why are the papers so difficult?

You've got to set up the axiomatisation, which is non-trivial because the 'nominal' aspects of the quantifiers have not been explored before.

You have to set up nominal algebras; we can't assume the reader knows them.

Then you find the right notion of filter and topology, then prove duality. Duality results are difficult (and addictive).

Soundness is fairly easy but completeness is not a straightforward generalisation of the non-nominal case. More hard work.

Plus, all these things are interacting with one another like crazy. A change on page 50 can lead to changes on pages 10 to 80 (and often does).

In the case of λ it's even more delicate, because $Y {-\!\!\bullet} X$ is negative in Y. The proofs get really tight.

What is this good for?

I probably don't have to sell this so hard to this audience. You are probably better-qualified to answer this question than me.

Interpret predicates and λ-terms as open sets. Interpret $\forall$ as a literal intersection, and also as a fresh-finite limit.

The models are absolute (no valuations).

Duality results for both FOL and the λ-calculus and representations in nominal powersets (sketched in this talk).

Direct derivation of first-order logic with equality and the λ-calculus just from atoms and powersets. Great personal satisfaction.