One-and-a-halfth order logic

Murdoch J Gabbay, Heriot Watt University, UK
Joint work with Aad Mathijssen, TU/e, the Netherlands

PPDP, Venice, Italy 11/7/2006

Thanks. . .

. . . to the Computer Science department of the University of Venice, for making possible a visit to Venice in early 2006 (photos on my webpage).

. . . to the Technical University of Eindhoven for supporting this work.

Recall: First-Order Logic with equality (FOL)

Fix countably infinitely many variable symbols $a, b, c, \ldots$. Let terms be:

$$t ::= a$$

(More interesting syntactic universes of terms are possible!)

Formulae or predicates are:

$$\phi ::= \bot \quad | \quad \phi \Rightarrow \phi \quad | \quad \forall a.\phi \quad | \quad t \approx t'.$$

Write $\equiv$ for syntactic identity (identify formulae up to α-equivalence).

Derivation

A context Φ and cocontext Ψ are finite and possibly empty sets of formulae. A judgement is a pair $\Phi \vdash \Psi$. Valid judgements:

$$(Axiom)\ \frac{}{\phi,\ \Phi \vdash \Psi,\ \phi} \qquad (\bot L)\ \frac{}{\bot,\ \Phi \vdash \Psi}$$

$$(\Rightarrow R)\ \frac{\phi,\ \Phi \vdash \Psi,\ \psi}{\Phi \vdash \Psi,\ \phi \Rightarrow \psi} \qquad (\Rightarrow L)\ \frac{\Phi \vdash \Psi,\ \phi \quad \psi,\ \Phi \vdash \Psi}{\phi \Rightarrow \psi,\ \Phi \vdash \Psi}$$

$$(\forall R)\ \frac{\Phi \vdash \Psi,\ \psi}{\Phi \vdash \Psi,\ \forall a.\psi}\ a \text{ fresh for } \Phi, \Psi \qquad (\forall L)\ \frac{\phi[a\mapsto t],\ \Phi \vdash \Psi}{\forall a.\phi,\ \Phi \vdash \Psi}$$

$$(\approx L)\ \frac{\phi[a\mapsto t],\ \Phi \vdash \Psi}{t \approx t',\ \phi[a\mapsto t'],\ \Phi \vdash \Psi} \qquad (\approx R)\ \frac{\Phi \vdash \Psi}{\Phi \vdash \Psi,\ t \approx t}$$

What is the status of this definition?

What are ϕ and ψ?

Meta-variables ranging over formulae.

What are t and a?

Meta-variables ranging over terms and variable symbols.

What is $\phi[a \mapsto t]$?

A meta-level operation defined which given a real predicate, variable symbol, and term, gives a predicate.

What is 'a fresh for Φ and Ψ'?

A meta-level condition only defined when given a real context and cocontext.

What does this definition serve to establish?

An entailment relation

$$\Phi \vdash \Psi.$$

Thanks to the predicate part of FOL this can be internalised as

$$\text{‘}\Phi^{\wedge} \Rightarrow \Psi^{\vee} \text{ holds’},$$

where $\{\phi_1, \ldots, \phi_n\}^{\wedge} \equiv \phi_1 \wedge \cdots \wedge \phi_n$ and $\{\}^{\wedge} = \top$, and $\{\phi_1, \ldots, \phi_n\}^{\vee} \equiv \phi_1 \vee \cdots \vee \phi_n$ and $\{\}^{\vee} = \bot$.

So FOL is a syntax, and a set of valid formulae.

We'll return to this later.

Proof-schema

Quite a lot of things happen in the meta-level in FOL. For example

$$\vdash \forall a.\forall b.\phi \Leftrightarrow \forall b.\forall a.\phi$$

is derivable for every value of the meta-variable ϕ:

$$\cfrac{\cfrac{\cfrac{\cfrac{\cfrac{}{\phi \vdash \phi}\,(Axiom)}{\forall b.\phi \vdash \phi}\,(\forall L)}{\forall a.\forall b.\phi \vdash \phi}\,(\forall L)}{\forall a.\forall b.\phi \vdash \forall a.\phi}\,(\forall R)}{\forall a.\forall b.\phi \vdash \forall b.\forall a.\phi}\,(\forall R)$$

Schema

However, the fact that this happens for all ϕ cannot be expressed in FOL.

Some nice example theorems:

- If $t \approx t'$ then $\phi[a \mapsto t] \Leftrightarrow \phi[a \mapsto t']$.
- If $a \notin fv(\phi)$ then $\vdash (\forall a.\phi) \Leftrightarrow \phi$.
- $\forall a.\forall b.\phi$ if and only if $\forall b.\forall a.\phi$.

Normally you might go to higher-order logic to express universal properties ranging over all predicates. However, unification up to β-equality is undecidable, the models get more complex, and there are other prices for the convenience.

One-and-a-half proof schema

One-and-a-halfth order logic applies nominal terms to represent the meta-level.

Take (nominal) term-formers $\approx$, $\forall$, $\Rightarrow$, $\bot$, and sub.

Read these as 'equals', 'forall', 'implies', 'false' (or 'bot'), and 'substitute'.

Terms are $t ::= a$ and formulae or predicates are:

$$\phi ::= P \mid \bot \mid P \Rightarrow P \mid \forall [a] P \mid t \approx t' \mid P[a \mapsto t]$$

Here we write $P[a \mapsto t]$ for $\mathsf{sub}([a]P, t)$.

Sugar

Write

- $\neg\phi$ for $\phi \Rightarrow \bot$,
- $\phi \wedge \phi'$ for $\neg(\phi \Rightarrow \neg\phi')$,
- $\phi \Leftrightarrow \phi'$ for $(\phi \Rightarrow \phi') \wedge (\phi' \Rightarrow \phi)$,
- $\phi \vee \phi'$ for $(\neg\phi) \Rightarrow \phi'$,
- $\top$ for $\bot \Rightarrow \bot$.

Write Φ, Ψ for contexts, which are finite sets of formulae.

Let a primitive freshness assertion be $a\#P$, read it as 'a does not occur in P'. Write Δ for a freshness context, a finite set of primitive freshness assertions.

Sequent derivation rules

$$\frac{}{\phi,\ \Phi \vdash_{\triangle} \Psi,\ \phi}\,(Axiom) \qquad \frac{}{\bot,\ \Phi \vdash_{\triangle} \Psi}\,(\bot L)$$

$$\frac{\Phi \vdash_{\triangle} \Psi,\ \phi \qquad \psi,\ \Phi \vdash_{\triangle} \Psi}{\phi \Rightarrow \psi,\ \Phi \vdash_{\triangle} \Psi}\,(\Rightarrow L) \qquad \frac{\phi,\ \Phi \vdash_{\triangle} \Psi,\ \psi}{\Phi \vdash_{\triangle} \Psi,\ \phi \Rightarrow \psi}\,(\Rightarrow R)$$

$$\frac{\phi',\ \Phi \vdash_{\triangle} \Psi \qquad \Delta \vdash_{\mathsf{SUB}} \phi' = \phi[a \mapsto t]}{\forall[a]\phi,\ \Phi \vdash_{\triangle} \Psi}\,(\forall L)$$

$$\frac{\Phi \vdash_{\triangle} \Psi,\ \psi \qquad \Delta \vdash a\#\Phi, \Psi}{\Phi \vdash_{\triangle} \Psi,\ \forall[a]\psi}\,(\forall R)$$

Em. . . just a few more sequent derivation rules

$$\frac{}{\Phi \vdash \Psi,\ t \approx t}\ (\approx R)$$

$$\frac{\phi',\ \Phi \vdash \Psi \qquad \Delta \models_{\mathsf{SUB}} \phi' = \phi''[a{\mapsto}t'] \qquad \Delta \models_{\mathsf{SUB}} \phi = \phi''[a{\mapsto}t]}{t'{\approx}t,\ \phi,\ \Phi \vdash_{\Delta} \Psi}\ (\approx L)$$

$$\frac{\phi',\ \Phi \vdash_{\Delta} \Psi \qquad \Delta \models_{\mathsf{SUB}} \phi' = \phi}{\phi,\ \Phi \vdash_{\Delta} \Psi}\ (StructL)$$

$$\frac{\Phi \vdash_{\Delta} \Psi,\ \psi' \qquad \Delta \models_{\mathsf{SUB}} \psi' = \psi}{\Phi \vdash_{\Delta} \Psi,\ \psi}\ (StructR)$$

Example derivations

$$\dfrac{\dfrac{\forall[a]\forall[b]X \vdash X \qquad a\#\forall[a]\forall[b]X}{\forall[a]\forall[b]X \vdash \forall[a]X}(\forall R) \qquad b\#\forall[a]\forall[b]X}{\forall[a]\forall[b]X \vdash \forall[b]\forall[a]X}(\forall R)$$

$$\dfrac{\dfrac{\dfrac{}{X \vdash X}(Axiom) \qquad \vdash_{\mathsf{SUB}}\ X = X[b\mapsto b]}{\forall[b]X \vdash X}(\forall L) \qquad \vdash_{\mathsf{SUB}}\ \forall[b]X = (\forall[b]X)[a\mapsto a]}{\forall[a]\forall[b]X \vdash X}(\forall L)$$

Semantics in FOL: "For all ϕ, $\forall a.\forall b.\phi \vdash \forall b.\forall a.\phi$."

Freshness part of the derivation

$$\cfrac{\cfrac{\cfrac{\cfrac{}{b\#[b]X}\,(\#[]a)}{b\#\forall[b]X}\,(\#\mathsf{f})}{b\#[a]\forall[b]X}\,(\#[]a)}{b\#\forall[a]\forall[b]X}\,(\#\mathsf{f})$$

Another example derivation

$$\dfrac{\dfrac{}{X[a{\mapsto}T'] \vdash X[a{\mapsto}T']}\,(Axiom) \quad \begin{array}{c} \models_{\mathsf{SUB}}\ X[a{\mapsto}a][a{\mapsto}T'] = X[a{\mapsto}T'], \\ \models_{\mathsf{SUB}}\ X[a{\mapsto}a][a{\mapsto}T] = X[a{\mapsto}T] \end{array}}{T' \approx T,\ X[a{\mapsto}T] \vdash X[a{\mapsto}T']}\,(\approx L)$$

Semantics in FOL:

"For all t and t' and ϕ, $\quad t' \approx t,\ \phi[a{\mapsto}t] \vdash \phi[a{\mapsto}t']$."

One more example derivation

$$\dfrac{\dfrac{}{X \vdash_{a\#X} X}(Axiom) \qquad a\#X \vdash a\#X}{X \vdash_{a\#X} \forall[a]X}(\forall R)$$

Semantics in FOL:

"For all ϕ and a, if $a \notin fv(\phi)$ then $\phi \vdash \forall a.\phi$."

A nice theorem:

$$\frac{\Phi \vdash_{\Delta} \Psi, \phi \qquad \phi, \Phi \vdash_{\Delta} \Psi}{\Phi \vdash_{\Delta} \Psi} (Cut)$$

Theorem (cut-elimination): Cut is eliminable.

The cut-elimination procedure is almost standard — but this is cut-elimination in the presence of unknown formulae.

Since the cut-elimination procedure is normally written parametrically over those formulae, this is no surprise really. However, the meta-level reasoning about substitution and α-equivalence is now all completely explicit on the nominal terms.

Another nice theorem:

Say a nominal term is closed when it mentions no unknowns. So a is closed but X is not.

Theorem: First-order logic (and its derivations) correspond to sequents of closed terms (and their derivations); term-for-term up to $\vdash_{\mathsf{SUB}}$, and proof-rule by proof-rule (up to $(Struct)$).

Recall that FOL is just valid formulae

$$P \Rightarrow Q \Rightarrow P \qquad (P \Rightarrow Q) \Rightarrow (Q \Rightarrow R) \Rightarrow (P \Rightarrow R)$$

$$\neg\neg P \Rightarrow P \qquad \forall[a](P \wedge Q) \Leftrightarrow \forall[a]P \wedge \forall[a]Q$$

$$\bot \Rightarrow P \qquad a\#P \vdash \forall[a](P \Rightarrow Q) \Leftrightarrow P \Rightarrow \forall[a]Q$$

$$T \approx T \qquad \forall[a]P \Rightarrow P[a \mapsto T]$$

$$U \approx T \wedge P[a \mapsto T] \Rightarrow P[a \mapsto U]$$

This plus modus ponens gives the same valid formulae as the sequent system (but no proof-theory!).

Second/Higher-order logic

In Higher-Order Logic (HOL), propositions have a type o and $\forall_\sigma$ is a constant with type $(\sigma \longrightarrow o) \longrightarrow o$, write just $\forall$ or $\forall : (\sigma \longrightarrow o) \longrightarrow o$.

Then the two judgements express the same idea:

'$\forall\lambda f.\big(\forall\lambda a.\forall\lambda b.fab \Leftrightarrow \forall\lambda b.\forall\lambda a.fab\big)$ is valid'

'$\forall[a]\forall[b]P \Leftrightarrow \forall[b]\forall[a]P$ is valid'.

f has function type. If $a : \sigma$ and $b : \tau$ then $f : \sigma \rightarrow \tau \rightarrow o$ and 'fab is P'.

Second/Higher-order logic

Similarly:

- ‘If $t \approx t'$ then $\phi[a\mapsto t] \Leftrightarrow \phi[a\mapsto t']$’ becomes

$$t \approx t' \vdash \forall\lambda f.\ \big(ft \Leftrightarrow ft'\big).$$

  in HOL.

  Note the types: f has function type and if $t : \sigma$ then $f : \sigma \rightarrow o$ and $\forall : ((\sigma \rightarrow o) \rightarrow o) \rightarrow o$.

- ‘If $a \notin fv(\phi)$ then $\vdash \forall a.\phi \Leftrightarrow \phi$’ is not expressible in HOL.

Relation to HOL

One-and-a-halfth-order logic is not fully higher-order. We can write

$$X \vdash Y$$

meaning in FOL "For all formulae ϕ and ψ, $\phi \vdash \psi$."

In HOL we can write this as $\vdash \forall\phi.\forall\psi.(\phi \Rightarrow \psi)$.

However we can also write $\vdash \forall\psi.\big((\forall\phi.\phi) \Rightarrow \psi\big)$.

This is not possible in one-and-a-halfth-order logic: $(\forall[X]X) \vdash Y$ is not syntax.

Relation to HOL

Not direct since we can express $a\#t$ and HOL cannot, but HOL can quantify over predicates to the left of an implication.

Also, suppose $X : o$ and $t : \mathbb{T}$.

$X[a \mapsto t]$ corresponds to ft and so X corresponds to f where $f : \mathbb{T} \to o$.

But $X[a \mapsto t][a' \mapsto t']$ corresponds to $f'tt'$ and X corresponds to $f' : \mathbb{T} \to \mathbb{T} \to o$.

But $X[a' \mapsto t'][a \mapsto t]$ corresponds to $f't't$ and X corresponds to $f' : \mathbb{T} \to \mathbb{T} \to o$.

Similarly $X[a \mapsto t][a' \mapsto t'][a'' \mapsto t'']\ldots$

Relation to HOL

This is type raising.

In one-and-a-halfth-order logic, X remains at o and the universal quantification implicit in the use of X allows this one symbol to represent a function of arbitrary arity — just like the meta-variable ϕ, which we cn write under substitutions $[a{\mapsto}t]$, $[a{\mapsto}t][a'{\mapsto}t]$, and so on, as we please.

Conclusions

Nominal Terms have ‘weak’ object-level variable symbols (atoms) with primitive facilities for abstraction and α-renaming and ‘strong’ meta-level variable symbols (variables or unknowns).

We can use this to axiomatise/build sequent systems for logic-with-binding, like first-order logic.

Conclusions

Sequent and axiomatic presentations systems are possible:

‘$\Phi \vdash_{\triangle} \Psi$ is derivable’

translating to ‘$\Delta \vdash (\Phi^{\wedge} \Rightarrow \Psi^{\vee})$ is valid’.

($^{\wedge}$ means ‘put $\wedge$ between the elements of Φ’, similarly for $^{\vee}$).

Conclusions

We get extra. E.g. one-and-a-halfth order logic has predicate unknowns; thus enabling us to reason universally on predicates in a new way.

This really new, because $a\#X$ is not expressible using other techniques (to our knowledge); not in full generality for a completely unkown X.

Conclusions

For some further work, how about. . .

- Two-and-a-halfth-order logic, where you can abstract P, and a predicate can assert freshness properties $a\#P$ of its own unknowns?
- Implementation and automation?
- Semantics (aside from in FOL)?

Note that this work is based on nominal algebra, a theory of algebraic equality of nominal terms. Watch this space.

The end

Axioms of substitution $\vdash^{\mathsf{SUB}}$

Write $t \vdash^{\mathsf{SUB}}_{\Delta} u$ when $t = u$ is derivable from assumptions Δ using the following axioms:

$$(f{\mapsto}) \quad \mathsf{f}(u_1, \ldots, u_n)[a{\mapsto}t] = \mathsf{f}(u_1[a{\mapsto}t], \ldots, u_n[a{\mapsto}t])$$

$$([b]{\mapsto}) \quad b\#t \Rightarrow ([b]u)[a{\mapsto}t] = [b](u[a{\mapsto}t])$$

$$(var{\mapsto}) \quad a[a{\mapsto}t] = t$$

$$(u{\mapsto}) \quad a\#u \Rightarrow u[a{\mapsto}t] = u$$

$$(ren{\mapsto}) \quad b\#u \Rightarrow u[a{\mapsto}b] = (b\ a) \cdot u$$

$$(perm) \quad a, b\#t \Rightarrow (a\ b) \cdot t = t$$

SUB

Equality is decidable in the theory of substitution (philosophically interesting fact, that!).

I do not know if unification is decidable.

The axioms above are sound and complete for a Herbrand-style model.

The cateogory of all models of substitution is cartesian-closed. Very interesting programming and logic principles; one-and-a-halfth-order logic is one creature inhabiting this new and wonderful universe.

There is much more out there.

Permutation action

$$\pi \cdot a \equiv \pi(a) \qquad \pi \cdot (\pi' X) \equiv (\pi \circ \pi')X$$

$$\pi \cdot [a]t \equiv [\pi(a)](\pi t)$$

$$\pi \cdot \mathsf{f}(t_1, \ldots, t_n) \equiv \mathsf{f}(\pi t_1, \ldots, \pi t_n)$$

Nominal Terms

Nominal terms are a syntax inductively generated by

$$t ::= a \mid \pi X \mid [a]t \mid \mathsf{f}(t,\ldots,t).$$

Here:

- We fix $a, b, c, \ldots \in \mathbb{A}$ a countably infinite set of atoms.
- We fix $X, Y, Z, \ldots \in \mathbb{V}$ a countably infinite set of unknowns (disjoint from the atoms; everything's disjoint).
- We fix f, g, ... some term-formers.
- Call $[a]t$ an abstraction.

Nominal Terms

$$t ::= a \quad | \quad \pi X \quad | \quad [a]t \quad | \quad \mathsf{f}(t, \ldots, t).$$

π is a permutation. A permutation is a finitely supported bijection on $\mathbb{A}$.

Finitely supported means:

$\pi(a) = a$ for all $a \in \mathbb{A}$ except for a finite set of atoms.

Nominal Terms

For example permutations are:

$$(a\ b\ c) \quad \text{and} \quad \textbf{Id}$$

(a to b to c to a, and the identity function). Permutations are not:

$$(a_1\ a_2)(a_3\ a_4)\ldots$$

for $\mathbb{A} = \{a_1, a_2, \ldots\}$.

Freshness assertions $a\#t$

Read $a\#X$ as 'a does not occur in X', or 'a is fresh for X'.

Then we can characterise α-equivalence as:

$$b\#X \Rightarrow [b](b\ a)X = [a]X.$$

For the moment I'm just telling you that this is the case.

Call a pair $a\#t$ a freshness assertion. If $t \equiv X$ call it primitive.

Freshness derivation rules (formally)

$$\frac{}{a\#b}\,(\#ab) \qquad \frac{a\#t_1 \quad \cdots \quad a\#t_n}{a\#\mathsf{f}(t_1,\ldots,t_n)}\,(\#f)$$

$$\frac{}{a\#[a]t}\,(\#[]a) \qquad \frac{a\#t}{a\#[b]t}\,(\#[]b) \qquad \frac{\pi^{-1}(a)\#X}{a\#\pi X}\,(\#X)$$

Core equality derivation rules (formally)

$$\frac{}{t = t}\,(refl) \qquad \frac{t = u}{u = t}\,(symm) \qquad \frac{t = u \quad u = v}{t = v}\,(tran)$$

$$\frac{t = u}{C[t] = C[u]}\,(cong) \qquad \frac{a\#t \quad b\#t}{(a\ b) \cdot t = t}\,(perm)$$

For example

$$\dfrac{\dfrac{}{a\#b}\,(\#ab) \qquad \dfrac{}{a\#b}\,(\#ab)}{[a]a = [b]b}\,(perm)$$

$$(b\ a) \cdot [b]b \equiv [a]a$$

$$\dfrac{\dfrac{b\#X}{b\#[a]X}\,(\#[]a) \qquad \dfrac{}{a\#[a]X}\,(\#[]a)}{[b](b\ a)X = [a]X}\,(perm)$$

$$(b\ a) \cdot [a]X \equiv [b](b\ a)X$$

Here $\equiv$ is syntactic identity.

Axioms

A freshness context Δ is a finite set of primitive freshness assertions.

An axiom $\Delta \vdash t = u$ is a pair of a freshness context and an equality assertion. If Δ is empty write it just $t = u$.

We can use axioms to enrich provable equality, which currently stands at some generalisation of α-equivalence.

Theory of λ-calculus LAM

$$(\lambda[a]Y)X = Y[a \mapsto X].$$

(Assume suitable term-formers λ, app and sugar.)

As an axiom, we instantiate Y and X to 'any term' when we enrich equality, generating a family of equalities for each instantiation (and each context). Thus, Y and X do represent 'any term', with universal quantification at top level. Instantiation is direct replacement of an unknown by a term (no capture avoidance).

Theory of first-order logic FOL

$$P \Rightarrow Q \Rightarrow P \;=\; \top \qquad (P \Rightarrow Q) \Rightarrow (Q \Rightarrow R) \Rightarrow (P \Rightarrow R) \;=\; \top$$

$$\neg\neg P \Rightarrow P \;=\; \top \qquad \forall[a](P \wedge Q) \Leftrightarrow \forall[a]P \wedge \forall[a]Q \;=\; \top$$

$$\bot \Rightarrow P \;=\; \top \qquad a\#P \;\vdash\; \forall[a](P \Rightarrow Q) \Leftrightarrow (P \Rightarrow \forall[a]Q) \;=\; \top$$

$$T \approx T \;=\; \top \qquad \forall[a]P \Rightarrow P[a{\mapsto}T] \;=\; \top$$

$$U \approx T \wedge P[a{\mapsto}T] \Rightarrow P[a{\mapsto}U] \;=\; \top$$

(Assume suitable term-formers $\approx, \forall, \Rightarrow, \bot$ and sugar.)

The '$= \top$' bit just converts a predicate into a nominal algebra judgement.

Theory of substitution SUB

$$\mathsf{f}(X_1,\ldots,X_n)[a\mapsto T] = \mathsf{f}(X_1[a\mapsto T],\ldots,X_n[a\mapsto T])$$

$$b\#T \vdash ([b]X)[a\mapsto T] = [b](X[a\mapsto T])$$

$$a[a\mapsto T] = T$$

$$a\#X \vdash X[a\mapsto T] = X$$

$$b\#X \vdash X[a\mapsto b] = (b\ a)X$$

Picture of what we have done

- $\equiv$ is syntactic identity $\qquad [a]a \not\equiv [b]b$
- $=$ (no axioms) is α-equivalence $\qquad b\#X \vdash [b](b\ a)X = [a]x$
- $=_{\mathsf{SUB}}$ is substitution $\qquad b\#Y \vdash Y[b{\mapsto}X] = Y$
- $=_{\mathsf{LAM}}$ is $\alpha\beta$-equivalence $\qquad (\lambda[a]a)b = b$
- $=_{\mathsf{FOL}}$ is logical equivalence $\qquad (\forall[a](a \approx a)) = \top$

All these theories are really very interesting beasts.