A simple class of Kripke-style models in which logic and computation have equal standing

Murdoch J. Gabbay
Joint work with Michael J. Gabbay

Thanks to Alex Simpson

June 29, 2010

What logical quantifier is λ?

I spend a lot of time studying names and binders.

So let me ask you this:

> *What logical quantifier is λ-binding?*

On the face of it, this question is a category error. λ is not a logical quantifier; it creates functions, not truth-values.

Or does it do both?

I'll dive straight into the definitions. They aren't hard.

Frames

A frame F is a 3-tuple $(W, \bullet, H)$ of

- W a set of worlds,
- $\bullet$ an application function from $W \times W$ to $powerset(W)$, and
- $H \subseteq powerset(W)$.

This looks like an applicative structure, but it is not because $w' \bullet w$ is a set of worlds, not a world.

An intuition for frames

I can suggest an intuition for frames.

Think of worlds $w \in W$ as data, and sets of worlds as concepts (we identify a property with the set of things that have that property).

Then $\bullet$ takes two elements of data and returns a concept.

A knife applied to a loaf of bread is not just a physical process; it provides a sliced loaf. We do not run computations on data only to make more data; ultimately, we obtain meaning.

The intuition is: $\bullet$ takes two data inputs and outputs a meaning.

H identifies a collection of concepts as 'computable'. λ as a logical quantifier will quantify over these. More on this later.

λ-abstraction as a logical quantifier

As I mentioned, frames are not quite applicative structures. Given a frame $F = (W, \bullet, H)$ define an applicative structure on $powerset(W)$ by:

$$S' \bullet S = \bigcup \{ w' \bullet w \mid w' \in S,\ w \in S \} \quad \text{for } S' \subseteq W,\ S \subseteq W$$

This gives an applicative structure such that application commutes with arbitrary unions; $(\bigcup_i S_i) \bullet (\bigcup_j T_j) = \bigcup_{ij} (S_i \bullet T_j)$. More on this later.

Frames are also not quite Kripke structures, but we can view $\bullet$ as a ternary 'accessibility relation' R such that $Rw_1w_2w_3$ if and only if $w_3 \in w_1 \bullet w_2$.

It is possible to manipulate frames as a generalised kind of Kripke structure, and to import ideas from modal logic.

Syntax of λ-terms

Fix a countably infinite set of variables p, q, r.

λ-terms are defined inductively as usual by:

$$\mathtt{t} ::= p \mid \lambda p.\mathtt{t} \mid \mathtt{t}\cdot\mathtt{t}$$

λp binds in $\lambda p.\mathtt{t}$.

Frames denotation of λ-terms

Suppose $F = (W, \bullet, H)$ is a frame.

A valuation v to F maps variables p to sets $v(p) \subseteq W$ (variable$\mapsto$meaning).

A denotation $[\![\mathtt{t}]\!]^v$ is defined by:

$$\begin{aligned} [\![p]\!]^v &= v(p) \\ [\![\mathtt{t}\cdot\mathtt{s}]\!]^v &= [\![\mathtt{t}]\!]^v \bullet [\![\mathtt{s}]\!]^v \\ [\![\lambda p.\mathtt{t}]\!]^v &= \{w \mid w \bullet h \subseteq [\![\mathtt{t}]\!]^{v[p:=h]} \text{ for all } h \in H\} \end{aligned}$$

$[\![\lambda p.\mathtt{t}]\!]^v$ is the greatest set such that $[\![\lambda p.\mathtt{t}]\!]^v \bullet h \subseteq [\![\mathtt{t}]\!]^{v[p:=h]}$ for all $h \in H$. Alternative characterisation:

$$[\![\lambda p.\mathtt{t}]\!]^v = \bigcup \{S \mid S \bullet h \subseteq [\![\mathtt{t}]\!]^{v[p:=h]} \text{ for all } h \in H\}$$

β-reduction and η-expansion

Call a valuation adequate when every variable maps to an element of H.

Call a frame adequate if the denotation of every λ-term is in H for every adequate valuation, so $[\![\mathtt{t}]\!]^v \in H$. (Terminology from Henkin models.)

We can verify the following properties, for adequate valuations:

β-reduction	$[\![\lambda p.\mathtt{t}]\!]^v \bullet [\![\mathtt{s}]\!]^v \subseteq [\![\mathtt{t}[p{:=}\mathtt{s}]]\!]^v$	
η-expansion	$[\![\mathtt{t}]\!]^v \subseteq [\![\lambda p.(\mathtt{t}\cdot p)]\!]^v$	(p not free in $\mathtt{t}$)

Adequate frames are a sound model for β-reduction and η-expansion.

They are also complete. I sketch the proof shortly.

Logic

It is easy to interpret logic in frames.

Define a language $\mathcal{L}$ with sentences ϕ inductively by:

$$\phi \quad ::= \quad p, q, r \ldots \mid \phi \to \phi \mid \forall p.\, \phi \mid \phi \cdot \phi \mid \phi \rhd \phi \mid \bot$$

$\forall p$ binds in $\forall p.\, \phi$.

Some formal syntax

Define a denotation $[\![\phi]\!]^v$ by:

$$
\begin{aligned}
[\![\bot]\!]^v &= \varnothing \\
[\![\phi \to \psi]\!]^v &= (W \backslash [\![\phi]\!]^v) \cup [\![\psi]\!]^v \\
[\![\phi \cdot \psi]\!]^v &= [\![\phi]\!]^v \bullet [\![\psi]\!]^v \\
[\![\phi \rhd \psi]\!]^v &= \{w \mid w \bullet [\![\phi]\!]^v \subseteq [\![\psi]\!]^v\} \\
[\![\forall p.\, \phi]\!]^v &= \bigcap_{h \in H} [\![\phi]\!]^{v[p \mapsto h]}
\end{aligned}
$$

We can now answer the question which which I opened the talk: λ is $\forall$ combined with $\rhd$.

λ as a quantifier

Define a quantifier y by:

$$\mathsf{y}p.\phi \stackrel{def}{=} \forall p.(p \rhd \phi)$$

y is the quantifier that corresponds to λ-abstraction.

Under this correspondence,

reduction $\sim$ *logical implication* $\sim$ *subset inclusion.*

λ-terms as truth-values

- Q. What is the negation of a λ-term?
- A. Its sets complement.
- Q. What is the conjunction of two λ-terms?
- A. Their sets intersection.

Note that the H does not have to be closed under sets complement and intersection. The negation of a λ-term does not have to be computational.

A dual to λ-quantification

You can do some nice things with this semantics.

For example, define a quantifier *match* by:

$$\mathit{match}\, p.\phi \stackrel{def}{=} \forall p.(\phi \rhd p)$$

Matching has dual properties to y. With this definition:

$$[\![(\mathit{match}\, p.\phi)\cdot(\phi[p{:=}\psi])]\!]^v = [\![\psi]\!]^v \quad ([\![\psi]\!]^v \in H)$$

See the paper for more details. There's an example at the end of these slides if we have time.

Completeness

The semantics briefly sketched above is complete: $\mathtt{s} \longrightarrow \mathtt{t}$ if and only if $[\![\mathtt{s}]\!]^{\nu} \subseteq [\![\mathtt{t}]\!]^{\nu}$ for all adequate frames and valuations.

Here's an intuition of how the proof works. Construct a frame whose worlds are sets $w_{\mathtt{t}} = \{\mathtt{s} \mid \mathtt{t} \longrightarrow \mathtt{s}\}$.

Computable sets of worlds in H are of the form $\{w_{\mathtt{t}} \mid \mathtt{s} \in w_{\mathtt{t}}\}$.

This is adequate by construction. Perhaps you can see that if you push through the calculations, you might get completeness. It works.

β-reduction and η-expansion

We suggest an intuition why models should satisfy β-reduction and η-expansion.

Both lose information:

- β-reduction loses information by performing a substitution; a non-invertible operation.
- η-expansion $\mathtt{t} \longrightarrow \lambda p.(\mathtt{t} \cdot p)$ loses intensional information in $\mathtt{t}$.

Neil Ghani wrote on this with Barry Jay in "the virtues of η-expansion".

It is possible to refine models to eliminate η-expansion, at some cost in complexity. You need a model with two accessibility relations $\bullet^{app} \subseteq \bullet^{lam}$.

Conclusions

We can tweak frames in other ways.

We can include $\beta\eta$-equality rather than reduction/expansion. The conditions can be viewed as lattice equalities on $\mathit{powerset}(W)$ or as geometric conditions on the accessibility relation $\bullet$ in the style of modal logic.

Conclusions

The title of the paper is carefully worded.

> *"A simple class of Kripke-style models in which logic and computation have equal standing."*

I don't think we're identifying logic and computation. We are providing models in which logic and computation exist side-by-side.

Where logic and computation intersect, they are equal: e.g subset inclusion models both reduction and entailment; universal quantification can be used to build y which models λ.

But note that the universal quantification is over H, not over $\mathit{powerset}(W)$. Note also that we cannot just take the transitive symmetric closure of subset inclusion to model the transitive symmetric closure of reduction, since everything is a subset of the top element W.

Computational behaviour happens in H.

Questions

Is there a frame in which H has the structure of a lattice or even a boolean algebra, so e.g. the intersection of two λ-terms is itself a λ-term?

What is the topological 'pointless' version of frames?

Can we apply the modal logic flavour of frames to interpret programming constructs, e.g. by using worlds to model notions of location?

An example of a construction using worlds as locations

An 'if. . . then. . . else' λ-operator: $\lambda_C p\begin{cases}\mathtt{t}_1\\ \mathtt{t}_2\end{cases}$ where

$(\lambda_C p\begin{cases}\mathtt{t}_1\\ \mathtt{t}_2\end{cases})\cdot\mathtt{s}$ reduces to $\mathtt{t}_1[p/\mathtt{s}]$ if C, and $\mathtt{t}_2[p/\mathtt{s}]$ otherwise.

The condition is on the function:

$$\forall p.\,\big((\chi \to (p \rhd \phi_1)) \wedge (\neg\chi \to (p \rhd \phi_2))\big)$$

The condition is on the argument:

$$\forall p.\,\big(((p \wedge \chi) \rhd \phi_1) \wedge ((p \wedge \neg\chi) \rhd \phi_2)\big)$$

The condition is on the application:

$$\forall p.\,\big(p \rhd ((\chi \to \phi_1) \wedge (\neg\chi \to \phi_2))\big)$$