Quine's Conjecture on Many-Sorted Logic

Thomas Barrett and Hans Halvorson

Dominik Ehrenfels
St Cross College

March 2, 2019

Many-sorted logic - Syntax

Non-logical vocabulary:

1. non-empty set of sort symbols $\sigma_1, \sigma_2, \ldots$
2. variables $x_1^{(\sigma_1)}, x_2^{(\sigma_1)}, \ldots x_1^{(\sigma_2)}, x_2^{(\sigma_2)}, \ldots$, indexed with a sort symbol
3. constant symbols $c^{(\sigma)}$, indexed with a sort symbol
4. predicate symbols P_i of arity $\sigma_{i_1} \times \ldots \times \sigma_{i_n}$
5. function symbols f_j of arity $\sigma_{j_1} \times \ldots \times \sigma_{j_n} \to \sigma_{j_{n+1}}$

Logical vocabulary:

1. connectives: $\neg, \vee, \ldots$
2. quantifiers $\forall_\sigma x$, $\exists_\sigma x$ for each sort symbol σ
3. equality sign $=$

Semantics

Structure $\mathcal{M}$:

1. A non-empty domain $\sigma^{\mathcal{M}}$ for each sort symbol σ. The domains of the sort symbols are pairwise disjoint ($\sigma_i^{\mathcal{M}} \cap \sigma_j^{\mathcal{M}} = \varnothing$ for $i \neq j$)
2. For any constant symbol c of sort σ, $c^{\mathcal{M}} \in \sigma^{\mathcal{M}}$
3. For any predicate P of arity $\sigma_{i_1} \times \ldots \times \sigma_{i_n}$, $P^{\mathcal{M}} \subseteq \sigma_{i_1}^{\mathcal{M}} \times \ldots \times \sigma_{i_n}^{\mathcal{M}}$
4. For function symbols analogous

$\mathcal{M} \vDash \forall_\sigma x \phi(x)$ iff $\mathcal{M} \vDash \phi[a]$ for all $a \in \sigma^{\mathcal{M}}$

Note that there are no quantifiers that range over multiple domains

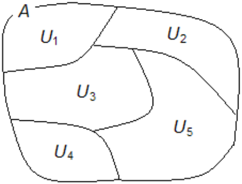

Quine's conjecture

Quine's conjecture

Every many-sorted theory is equivalent to a single-sorted theory

What is the precise sense of 'equivalent' here?

A standard notion of equivalence in the context of (single-sorted) FOL is *Definitional Equivalence*

Definitional Equivalence

Explicit definitions

Let Σ be a signature and let $P \notin \Sigma$, where P is some predicate symbol. An explicit definition of P in terms of Σ is a sentence

$$\forall \bar{x}(P\bar{x} \leftrightarrow \delta(\bar{x}))$$

where $\delta(\bar{x})$ is a Σ-formula

Constant symbols and function symbols can also be explicitly defined in a straightforward manner.

Definitional extensions

Let T be a theory in signature Σ, and let ψ_i, $i \in I$ be explicit definitions in terms of Σ of symbols not in Σ. Then $T \cup \{\psi_i | i \in I\}$ is a *definitional extension* of T.

Definitional equivalence

Two theories T_1 and T_2 of signature Σ_1 and Σ_2, respectively, are *definitionally equivalent* if they possess logically equivalent definitional extensions T_1^+ and T_2^+ of signature $\Sigma_1 \cup \Sigma_2$

An example

Let $\Sigma_1 = \{P\}$ and $\Sigma_2 = \{Q\}$ be signatures. Let T_1 be a Σ_1-theory and T_2 a Σ_2-theory:

$T_1 = \{\forall x Px\}$
$T_2 = \{\forall x \neg Qx\}$

Let $\delta \equiv \forall x(Qx \leftrightarrow \neg Px)$
Let $\delta' \equiv \forall x(Px \leftrightarrow \neg Qx)$

$T_1 \cup \{\delta\}$ and $T_2 \cup \{\delta'\}$ are definitional extensions of T_1 and T_2, respectively. They are furthermore logically equivalent. T_1 and T_2 are therefore definitionally equivalent

Generalising definitional equivalence - defining new sorts

So far, we have defined new constant symbols, predicate symbols, and function symbols via definitional extensions. But in the setting of many-sorted logic, we also encounter new sort symbols. How should we define these?

New sorts are definable from old sorts via four constructions: We can introduce product sorts, coproduct sorts, subsorts, and quotient sorts

Product sort

A product sort can be thought of as the Cartesian product of two sorts

Example: Let $\sigma_1^{\mathcal{M}} = \{a, b\}$ and $\sigma_2^{\mathcal{M}} = \{c\}$

Then $\sigma_P^{\mathcal{M}^+} = \{\langle a, c\rangle, \langle b, c\rangle\}$

Formally: The product sort σ of sorts σ_1 and σ_2 is defined by

$$\forall_{\sigma_1} x \forall_{\sigma_2} y \exists^{=1}_{\sigma} z (\pi_1(z) = x \wedge \pi_2(z) = y)$$

Here, the π_i are **new** function symbols of arity $\sigma \to \sigma_i$. Think of them as projections.

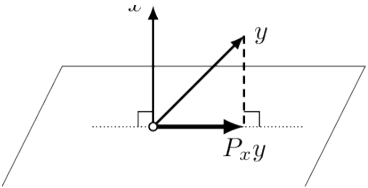

Coproduct sort

A coproduct sort can be thought of as the disjoint union of two sorts

Example: Let $\sigma_1^{\mathcal{M}} = \{a, b\}$ and $\sigma_2^{\mathcal{M}} = \{\alpha, \beta\}$

Then $\sigma_C^{\mathcal{M}^+} = \{\langle a, 1\rangle, \langle b, 1\rangle, \langle \alpha, 2\rangle, \langle \beta, 2\rangle\}$

Formally: The coproduct sort σ of sorts σ_1 and σ_2 is defined by

$$\forall_\sigma z(\exists^{=1}_{\sigma_1} x(\rho_1(x) = z) \vee \exists^{=1}_{\sigma_2} y(\rho_2(y) = z)) \wedge \forall_{\sigma_1} x \forall_{\sigma_2} y(\rho_1(x) \neq \rho_2(y))$$

Here, the ρ_i are **new** function symbols of arity $\sigma_i \to \sigma$. Think of them as equipping each element of σ_i with an index i

Subsort

Think of a subsort of σ as a copy of a definable subset of $\sigma^{\mathcal{M}}$

Example: Let $\sigma^{\mathcal{M}} = \{a, b, c\}$ and let $P^{\mathcal{M}} = \{a, b\}$. We can then define a subsort σ_S of σ that is a copy of $P^{\mathcal{M}}$, i.e. $\sigma_S^{\mathcal{M}} = \{a', b'\}$

Formally: A subsort σ of a sort σ_1 is defined by

$$\forall_{\sigma_1} x(\phi(x) \leftrightarrow \exists_\sigma z(h(z) = x)) \wedge \forall_\sigma y \forall_\sigma z(h(y) = h(z) \to y = z)$$

Here, $\phi(x)$ is an **old** formula which defines the subset of σ_1 we want to copy. h is a **new** function symbol of arity $\sigma \to \sigma_1$. Think of h as a bijection between σ and its copy.

Note that we cannot allow the domain of σ to be empty. $\exists_{\sigma_1} x \phi(x)$ must therefore hold. This is called the admissibility condition for the subsort σ

Quotient sort

The elements of a quotient sort σ_Q of σ are the equivalence classes of elements of σ with respect to some equivalence relation $\phi(x_1, x_2)$ on σ

Example: Let $\sigma^{\mathcal{M}} = \{Mark, John, Rachel, Mary\}$
Let $\phi(x_1, x_2)$ describe the equivalence relation $'x_1$ is the same gender as x_2'

$[Mark]_\phi = \{Mark, John\}$
$[Rachel]_\phi = \{Rachel, Mary\}$

$\sigma_Q^{\mathcal{M}^+} = \{[Mark]_\phi, [Rachel]_\phi\}$

Formally: A quotient sort σ of a sort σ_1 is defined by

$$\forall_{\sigma_1} x \forall_{\sigma_1} y (\epsilon(x) = \epsilon(y) \leftrightarrow \phi(x, y)) \wedge \forall_{\sigma} z \exists_{\sigma_1} x (\epsilon(x) = z)$$

Here, ϵ is a **new** function symbol of arity $\sigma_1 \to \sigma$. ϵ maps every element of σ_1 to its equivalence class.

Once again, there is an admissibility condition: $\phi(x, y)$ must be an equivalence relation, i.e. reflexive, symmetric, transitive

Morita equivalence

We are now in a position to define our new notion of Generalised Definitional Equivalence, or *Morita* Equivalence

Morita extensions

Let $\Sigma \subset \Sigma^+$ be signatures and T a Σ-theory. A Morita extension T^+ of T is a Σ^+-theory

$$T \cup \{\delta_s | s \in \Sigma^+ - \Sigma\}$$

For which it holds that

1. δ_s is an explicit definition of s
2. If α_s is an admissibility condition for s, then $T \vDash \alpha_s$

Morita equivalence

Let T_1 be a Σ_1-theory and T_2 a Σ_2-theory. T_1 and T_2 are Morita equivalent if there are theories $T_1^1, \ldots T_1^m$ and $T_2^1, \ldots T_2^n$ such that

1. T_1^{i+1} is a Morita extension of T_i^i for $0 \leq i \leq m-1$
2. T_2^{i+1} is a Morita extension of T_1^i for $0 \leq i \leq n-1$
3. T_1^m and T_2^n are logically equivalent

Why are multiple steps upwards needed (unlike for definitional equivalence)?
Answer: We can construct new sorts from complex sorts, which in turn are constructed from more basic sorts

Example

The following two theories are Morita equivalent:

$T_1 = \{\exists^{=1}_{\sigma_1} x(x = x)\}$ and $T_2 = \{\exists^{=2}_{\sigma_2} y(y = y)\}, \exists^{=1}_{\sigma_2} yPy\}$

To show this, we need to define the symbols of $\Sigma_1 = \{\sigma_1\}$ in terms of the symbols of $\Sigma_2 = \{\sigma_2, P\}$, and vice versa. Then we can build a common Morita extension T^+ of T_1 and T_2

The domain of σ_1 in any model $\mathcal{M}$ of T_1 has exactly one element, e.g. $\sigma_1^{\mathcal{M}} = \{a\}$. To construct a domain for σ_2 out of this, we need to turn this one element into two.

$\Rightarrow$ σ_2 must be defined as the coproduct of σ_1 with itself

$$\sigma_2^{\mathcal{M}^+} = \{\langle a, 1\rangle, \langle a, 2\rangle\}$$

We also need to define $P \in \Sigma_2$ in terms of Σ_1. For this, we can just define $P^{\mathcal{M}^+} = \{\langle a, 1\rangle\}$, i.e. the first element of the $\sigma_2^{\mathcal{M}^+}$ we just constructed

Now we need to define σ_1 in terms of Σ_2. $\sigma_1^{\mathcal{M}^+}$ needs to have exactly one element. We just saw that $P^{\mathcal{M}^+}$ has exactly one element. So let's define σ_1 as a copy of $P^{\mathcal{M}^+}$, i.e. as a subsort of σ_2. For instance, $\sigma_1^{\mathcal{M}^+} = \{a\}$

Now let's do all of this in the syntax!

The signature Σ^+ of our common Morita extension will be $\Sigma^+ = \{\sigma_1, \sigma_2, P, \rho_1, \rho_2\}$. ρ_1 and ρ_2 are function symbols of arity $\sigma_1 \to \sigma_2$ which we need for the definitions of the product sort and the subsort. Let's define σ_2 and P:

$$\begin{aligned}
\delta_{\sigma_2} \equiv & \forall_{\sigma_2} z(\exists^{=1}_{\sigma_1} x(\rho_1(x) = z) \vee \exists^{=1}_{\sigma_1} x(\rho_2(x) = z)) \\
& \wedge \forall_{\sigma_1} x \forall_{\sigma_1} y(\rho_1(x) \neq \rho_2(y)) \\
\delta_P \equiv & \forall_{\sigma_2} z(Pz \leftrightarrow \exists_{\sigma_1} x(z = \rho_1(x)))
\end{aligned}$$

And let's define σ_1:

$$\begin{aligned}
\delta_{\sigma_1} \equiv & \forall_{\sigma_2} z(Pz \leftrightarrow \exists_{\sigma_1} x(z = \rho_1(x))) \\
& \wedge \forall_{\sigma_1} x \forall_{\sigma_1} y(\rho_1(x) = \rho_1(y) \to x = y)
\end{aligned}$$

$T_1 \cup \{\delta_{\sigma_2}, \delta_P\}$ is a theory in the target signature $\Sigma^+ = \{\sigma_1, \sigma_2, P, \rho_1, \rho_2\}$. We have reached our common Morita extension, starting from T_1

But $T_2 \cup \{\delta_{\sigma_1}\}$ is in the signature $\Sigma' = \{\sigma_1, \sigma_2, P, \rho_1\}$. We have not yet reached the common Morita extension from T_2 since we have not yet defined ρ_2. We need to extend $T_2 \cup \{\delta_{\sigma_1}\}$ once more to reach the common Morita extension. Let's add

$$\delta_{\rho_2} \equiv \forall_{\sigma_1} x \forall_{\sigma_2} y(\rho_2(x) = y \leftrightarrow \rho_1(x) \neq y)$$

One can verify that $T_1 \cup \{\delta_{\sigma_2}, \delta_P\}$ is logically equivalent to $T_2 \cup \{\delta_{\sigma_1}, \delta_{\rho_2}\}$. We have therefore found our common Morita extension of T_1 and T_2.

What do its models look like? Here is one, call it $\mathcal{M}^+$, based on our earlier constructions:

- $\sigma_1^{\mathcal{M}^+} = \{a\}$
- $\sigma_2^{\mathcal{M}^+} = \{\langle a, 1\rangle, \langle a, 2\rangle\}$
- $P^{\mathcal{M}^+} = \{\langle a, 1\rangle\}$
- $\rho_1^{\mathcal{M}^+} = \langle a, \langle a, 1\rangle\rangle$
- $\rho_2^{\mathcal{M}^+} = \langle a, \langle a, 2\rangle\rangle$

This brings us to two important notions: *expansions* and *reducts* of structures

Expansions and reducts

Let $\Sigma \subset \Sigma'$ be signatures. Let $\mathcal{M}$ be a Σ'-structure. The unique Σ-structure that agrees with $\mathcal{M}$ on the interpretation of every symbol in Σ is called the *reduct* of $\mathcal{M}$ to Σ. We write $\mathcal{M}|_\Sigma$ for the reduct.

If a Σ-structure $\mathcal{M}$ is the reduct of some Σ'-structure $\mathcal{M}^+$, then $\mathcal{M}^+$ is called an *expansion* of $\mathcal{M}$. Expansions are in general not unique

Recall the model $\mathcal{M}^+$ of the common Morita extension of T_1 and T_2 in the previous example. Its reducts to $\Sigma_1 = \{\sigma_1\}$ and $\Sigma_2 = \{\sigma_2, P\}$, respectively are

- $\mathcal{M}^+|_{\Sigma_1}$: $\sigma_1^{\mathcal{M}^+|_{\Sigma_1}} = \{a\}$
- $\mathcal{M}^+|_{\Sigma_2}$: $\sigma_2^{\mathcal{M}^+|_{\Sigma_2}} = \{\langle a, 1\rangle, \langle a, 2\rangle\}$, $P^{\mathcal{M}^+|_{\Sigma_2}} = \{\langle a, 1\rangle\}$

Note that $\mathcal{M}^+|_{\Sigma_1}$ and $\mathcal{M}^+|_{\Sigma_2}$ are models of T_1 and T_2, respectively. In general, the following holds

Proposition. If T_1 and T_2 are Morita equivalent theories of signature Σ_1 and Σ_2, respectively, and T^+ some common Morita extension, then every model $\mathcal{M}$ of T_1 can be expanded into a model $\mathcal{M}^+$ of T^+, such that $\mathcal{M}^+|_{\Sigma_2}$ is a model of T_2, and vice versa.

Returning to Quine's conjecture

Quine's conjecture (precise version)

Every many-sorted theory is Morita equivalent to a single-sorted theory

It turns out that Quine's conjecture is false. I will now sketch the proof of this claim. We need some (very basic) category theory

Category theory - basic notions

A category consists of

1. objects $a, b, c, \dots$
2. arrows between objects $f : a \to b, g : c \to a, \dots$
3. in particular, an identity arrow $1_a : a \to a$ for every object a

Arrows compose: If $f : a \to b$ and $g : b \to c$ are arrows, then there exists an arrow $g \circ f : a \to c$

Equivalence of categories

A mapping $F : A \to B$ between categories A, B which maps objects and arrows of A to objects and arrows of B, respectively, and has the following properties

1. $F(f : a \to b) = Ff : Fa \to Fb$
2. $F(g \circ h) = Fg \circ Fh$
3. $F(1_a) = 1_{Fa}$

is called a *functor.*

A functor is called an equivalence of categories if it is full, faithful, and essentially surjective. If $F : A \to B$ is an equivalence of categories, then there exists an equivalence of categories $G : B \to A$

A functor $F : A \to B$ is **full** if for all arrows $g : Fa_1 \to Fa_2$ there exists an arrow $f : a_1 \to a_2$ such that $g = Ff$, for all a_1 and a_2 from A.

A functor $F : A \to B$ is **faithful** if $Ff = Fg$ implies $f = g$ for all arrows $f : a_1 \to a_2$ and $g : a_1 \to a_2$ in A

[Note that a faithful functor may map $f : a_1 \to a_2$ and $h : a_2 \to a_1$ to the same arrow in B however]

A functor $F : A \to B$ is **essentially surjective** if for every object b in B there is an object a in A such that Fa is isomorphic to b.

[Here, isomorphic means that there are arrows $g : Fa \to b$ and $g^{-1} : b \to Fa$ in B such that $g \circ g^{-1} = 1_b$ and $g^{-1} \circ g = 1_{Fa}$]

The category $Mod(T)$

Let T be a theory. The category $Mod(T)$ has for its objects the models of T. The arrows between the models are elementary embeddings

Elementary embeddings

Let Σ be a signature and let $\mathcal{M}$ and $\mathcal{M}'$ be Σ-structures

An *elementary embedding* $h : \mathcal{M} \to \mathcal{M}'$ is a family of injective maps $h_\sigma : \sigma^{\mathcal{M}} \to \sigma^{\mathcal{M}'}$, for $\sigma \in \Sigma$, with the following property

- $\mathcal{M} \models \phi[a_1, \ldots a_n]$ iff $\mathcal{M}' \models \phi[h_{\sigma_1}(a_1), \ldots h_{\sigma_n}(a_n)]$ for all Σ-formulae $\phi(x_1, \ldots x_n)$ and any elements $a_1 \in \sigma_1^{\mathcal{M}}, \ldots, a_n \in \sigma_n^{\mathcal{M}}$

If every h_σ is surjective, h is called an *isomorphism*

An isomorphism $h : \mathcal{M} \to \mathcal{M}$ is called an *automorphism*

Categorical equivalence - a simple example

Let $\Sigma_1 = \{P\}$ and $\Sigma_2 = \{Q\}$ be signatures, with P and Q unary predicate symbols.
Let T_1 and T_2 be a Σ_1- and a Σ_2-theory, respectively.

$$T_1 = \{\exists! x(x = x) \wedge \forall x Px\}$$

$$T_2 = \{\exists! x(x = x) \wedge \forall x \neg Qx\}$$

Every model $\mathcal{M}$ of T_1 looks like this:

1. $|\mathcal{M}| = \{a\}$, for some object a
2. $P^{\mathcal{M}} = \{a\}$

The models of T_2 also have singleton domains, but $Q^{\mathcal{M}} = \varnothing$

To show that T_1 and T_2 are categorically equivalent, we have to look at the categories $Mod(T_i)$

The models of T_1 will be the objects in the category $Mod(T_1)$.

What about the arrows in $Mod(T_1)$, i.e. the elementary embeddings between the models of T_1?
For any two models $\mathcal{M}$ and $\mathcal{M}'$ of T_1, there exists a unique function $f : |\mathcal{M}| \to |\mathcal{M}'|$
f is evidently an isomorphism, and hence an elementary embedding.

$\Rightarrow$ We have identified all the arrows in $Mod(T_1)$

All the same goes for $Mod(T_2)$

We now want to show that there exists an equivalence of categories between $Mod(T_1)$ and $Mod(T_2)$. We will show that

$$\begin{aligned} F : Mod(T_1) &\to Mod(T_2) \\ \mathcal{M} &\mapsto \mathcal{M}[Q] \\ f &\mapsto \tilde{f} \end{aligned}$$

is such an equivalence.

Here, $\mathcal{M}[Q]$ is the model of T_2 with $|\mathcal{M}[Q]| = |\mathcal{M}|$. And $\tilde{f}$ is the unique arrow $\tilde{f} : \mathcal{M}[Q] \to \mathcal{M}'[Q]$, if f is the unique arrow $f : \mathcal{M} \to \mathcal{M}'$

F is an equivalence of categories:

F is full: Let g be the unique arrow from $\mathcal{M}[Q]$ to $\mathcal{M}'[Q]$. $g = Ff$ for f the unique arrow $f : \mathcal{M} \to \mathcal{M}'$

F is faithful: immediate from the uniqueness of f

F is essentially surjective: T_2 is categorical, therefore **all** of its models are isomorphic

Quine's conjecture is false

Proposition. Let $\Sigma = \{\sigma_1, \sigma_2, \dots\}$ be a signature with infinitely many sort symbols. The Σ-theory $T = \varnothing$ is not Morita equivalent to any single-sorted theory.

Intuitive justification: The theory T says that everything is either σ_1 or σ_2 or..., but that is not expressible in FOL without an infinite disjunction

Proposition. Let $\Sigma = \{\sigma_1, \sigma_2, ...\}$ be a signature with infinitely many sort symbols. The Σ-theory $T = \varnothing$ is not Morita equivalent to any single-sorted theory.

Sketch of proof: Barrett and Halvorson's proof is a proof by contradiction. Assume that there is some single-sorted theory T' in signature Σ' that is Morita equivalent to T. Call their common Morita extension T^+. Remember that every model $\mathcal{M}$ of T can be expanded into a model $\mathcal{M}^+$ of T^+ such that $\mathcal{M}^+|_{\Sigma'}$ is a model of T'. It can be shown that there exists an equivalence of categories $F : Mod(T) \to Mod(T')$ that maps every model $\mathcal{M}$ of T to the corresponding model $\mathcal{M}^+|_{\Sigma'}$ of T'.

Sketch of proof (cont.): Barrett and Halvorson now consider a specific model $\mathcal{M}$ of T. It has two important properties: (1) $\sigma_i^{\mathcal{M}}$ is finite for every sort symbol $\sigma_i \in \Sigma$, and (2) $\mathcal{M}$ possesses infinitely many automorphisms.

They now consider what the corresponding $\mathcal{M}^+|_{\Sigma'}$ must look like. $\mathcal{M}^+|_{\Sigma'}$ is a Σ'-structure, and Σ' only contains a single sort symbol, call it α. Since T and T' are by assumption Morita equivalent, $\alpha^{\mathcal{M}^+|_{\Sigma'}}$ must in some way be constructed from the finite domains of the sort symbols in Σ via the product-, coproduct-, subsort-, and quotient-operations. But this means that $\alpha^{\mathcal{M}^+|_{\Sigma'}}$ is itself finite, since these operations all produce finite sets when applied to finite sets.

Sketch of proof (cont.): Hence, $\mathcal{M}^+|_{\Sigma'}$ possesses at most finitely many automorphisms. But then the functor F, which maps $\mathcal{M}$ to $\mathcal{M}^+|_{\Sigma'}$ cannot be a faithful functor, since it maps the infinitely many automorphisms of $\mathcal{M}$ onto the finitely many automorphisms of $\mathcal{M}^+|_{\Sigma'}$. Contradiction.

A weakened version of Quine's conjecture holds true:

Quine's conjecture (weakened)

If Σ contains only finitely many sort symbols, then every Σ-theory is Morita equivalent to some single-sorted theory

The proof is too long to go over the details. But I will sketch how, given a finitely-sorted theory, one finds a Morita equivalent single-sorted theory.

Say we are given a theory T in a signature Σ with finitely many sort symbols $\sigma_1, ...\sigma_n$. In the single-sorted theory T' we want to find, we will represent these by unary predicate symbols $Q_{\sigma_1}, ... Q_{\sigma_n}$.

To mimic the semantics of the sort symbols, we add the following axioms to T':

- $\exists_\sigma x Q_{\sigma_i} x$ for every σ_i in Σ
- $\forall_\sigma x(Q_{\sigma_1} x \vee ... \vee Q_{\sigma_n} x)$
- $\forall_\sigma x(Q_{\sigma_i} x \rightarrow (\neg Q_{\sigma_1} x \wedge ... \wedge \neg Q_{\sigma_{i-1}} x \wedge \neg Q_{\sigma_{i+1}} x \wedge ... \wedge \neg Q_{\sigma_n} x))$

One can then find a translation of the sentences of T into the single-sorted language. For details see the paper.