Geometric Resolution: A Proof Procedure Based on Finite Model Search

Hans de Nivelle and Jia Meng

Max Planck Institut für Informatik, Saarbrücken, Germany,

National ICT, Canberra, Australia,

Seattle, 19.08.2006

Introduction

We present a new calculus for first-order logic with equality.

We call the calculus geometric resolution, because it operates on a normal form, which is derived from geometric formulas. (this is a first-order fragment introduced by Thoralf Skolem)

We show that the calculus is sound and complete for first-order logic.

Motivation

- Try out something new.
- Avoid use of Herbrand's theorem, because unrestricted interpretations can be much more compact than Herbrand interpretations.
- Find general theorem proving strategies with good termination behaviour.
- Find theorem proving strategies that can deal better with partial functions, incompletely defined functions.

Definition: We assume an infinite set of variables $\mathcal{V}$.

A variable atom is an atom of one of the following two forms:

1. $p(v_1, \ldots, v_n)$ with $n \geq 0$ and $v_1, \ldots, v_n \in \mathcal{V}$.
2. $v_1 \not\approx v_2$ with $v_1, v_2 \in \mathcal{V}$.

Observe that:

- There are no positive equalities.
- There are no constants and no function symbols.

Definition: A geometric formula has form

$$\forall \overline{x}\ A_1(\overline{x}) \wedge \cdots \wedge A_p(\overline{x}) \wedge x_1 \not\approx x_1' \wedge \cdots \wedge x_q \not\approx x_q' \rightarrow Z(\overline{x}),$$

in which $x_1, x_1', \ldots, x_q, x_q' \in \overline{x} \subseteq \mathcal{V}$.

$Z(\overline{x})$ can have one of the following three forms:

1. The false constant $\bot$.
2. A disjunction of atoms $B_1(\overline{x}) \vee \cdots \vee B_r(\overline{x})$, with $r > 0$.
3. An existential formula of form $\exists y\ B(\overline{x}, y)$.

Types 1 and 2 overlap (if one would allow $r = 0$) but we prefer to distinguish the types. Geometric formulas of Type 1 are called lemmas. Formulas of Type 2 are called disjunctive. Formulas of Type 3 are called existential.

Example 1

We are interested in finding out whether $a \approx b,\ b \approx c \vdash\ a \approx c$.

We try to find a model for

$$a \approx b, \quad b \approx c, \quad a \not\approx c.$$

Resulting geometric formulas are:

$A(X) \wedge B(Y) \wedge X \not\approx Y \rightarrow \bot$,
$B(X) \wedge C(Y) \wedge X \not\approx Y \rightarrow \bot$,
$A(X) \wedge C(X) \rightarrow \bot$,

$\rightarrow \exists x\ A(x)$,
$\rightarrow \exists x\ B(x)$,
$\rightarrow \exists x\ C(x)$.

Example 2

What about $s(a) \approx a \vdash s(s(a)) \approx a$?

Try to find model for

$$s(a) \approx a, \quad s(s(a)) \not\approx a.$$

$A(X) \wedge S(X, Y) \wedge A(Y) \wedge X \not\approx Y \rightarrow \perp,$
$A(X) \wedge S(X, Y) \wedge S(Y, X) \rightarrow \perp,$

$\exists x\ A(x),$
$\forall x \exists y\ S(x, y).$

Example 3

$a \approx s(a), \quad p(a,a) \vee p(s(a),s(a)) \vdash p(a,a).$

Negation of goal:

$$a \approx s(a), \quad p(a,a) \vee p(s(a),s(a)), \quad \neg p(a,a).$$

$A(X) \wedge S(X,Y) \wedge X \not\approx Y \rightarrow \perp,$
$A(X) \wedge S(X,Y) \rightarrow p(X,X) \vee p(Y,Y),$
$A(X) \wedge p(X,X) \rightarrow \perp,$

$\exists x \; A(x),$
$\forall x \exists y \; S(x,y).$

After these examples, you might be willing to believe that:

Theorem:

Every set of first-order formulas can be translated into a set of geometric formulas, which is equisatisfiable.

The result (and the computation) can be linear in the size of the input.

- For each function symbol or constant f, introduce a new predicate symbol P_f, s.t. $\#P_f = \#f + 1$.
- for each new predicate symbol P_f, introduce a seriality axiom:

$$\forall \overline{x}\ \exists y\ P_f(\overline{x}, y).$$

- As long as F contains a functional term, let $f(x_1, \ldots, x_n)$ be a functional term with variable arguments.

  Write $F = F[\ A[\ f(x_1, \ldots, x_n)\ ]\ ]$, where A is the smallest subformula that contains all occurrences of $f(x_1, \ldots, x_n)$.

  Replace

$$F[\ A[\ f(x_1, \ldots, x_n)\ ]\ ]$$

  by

$$F[\ \forall y\ (P_f(x_1, \ldots, x_n, y) \rightarrow A[y]\ )\ ].$$

Searching for a Model

Definition: An interpretation is a finite set of atoms, with arguments from a fixed, given set $\mathcal{E}$.

Equality is interpreted as object equality, therefore there are no disequality atoms in interpretations.

Examples of interpretations are

$$A(e_0), \quad S(e_0, e_1), \quad S(e_1, e_2), \quad B(e_2).$$

$$A(e_0), \quad B(e_1), \quad P(e_0, e_1, e_2), \quad Q(e_2, e_2, e_1).$$

A Naive Algorithm for Theorem Proving

Definition: Let I be an interpretation. We call geometric formula

$$\forall \overline{x}\ A_1(\overline{x}) \wedge \cdots \wedge A_p(\overline{x}) \wedge x_1 \not\approx x_1' \wedge \cdots \wedge x_q \not\approx x_q' \rightarrow Z(\overline{x})$$

applicable in I if there is a ground substitution Θ, s.t.

- All $A_i(\overline{x})\Theta$ are in I.
- For each $x_j \not\approx x_j'$, $x_j\Theta$ and $x_j'\Theta$ are distinct.

In addition $Z(\overline{x})\Theta$ has to be false in I.

1. If $Z(\overline{x})$ has form $\perp$, then $Z(\overline{x})\Theta$ is always false in I.
2. If $Z(\overline{x})$ has form $B_1(\overline{x}) \vee \cdots \vee B_r(\overline{x})$ then $Z(\overline{x})\Theta$ is false in I, if none of $B_j(\overline{x})\Theta$ is present in I.
3. If $Z(\overline{x})$ has form $\exists y\ B(\overline{x}, y)$ then $Z(\overline{x})\Theta$ is false in I if there is no element $e \in \mathcal{E}$, s.t. $(B(\overline{x}, y)\Theta)\ \{y := e\}$ is present in I.

Start with empty interpretation $I = \{ \}$.

- If there is no applicable rule, then I is a model.
- Otherwise, select a rule $\forall \overline{x}\ \Phi(\overline{x}) \rightarrow Z(\overline{x})$ that is applicable on I with substitution Θ.
  - If $Z(\overline{x})$ has form $\perp$, then backtrack.
  - If $Z(\overline{x})$ has form $B_1(\overline{x}) \vee \cdots \vee B_r(\overline{x})$, then backtrack through all of
  $$I \cup \{B_j(\overline{x})\Theta\}.$$
  - If $Z(\overline{x})$ has form $\exists y\ B(\overline{x}, y)$, then backtrack through
  $$I \cup \{ \ \ B(\overline{x}, y)\ \Theta \cdot \{x := e\} \ \ \},$$
  for each e that is present in I. In addition, try
  $$I \cup \{ \ \ B(\overline{x}, y)\ \Theta \cdot \{x := e'\} \ \ \}$$
  for a new element e' that is not present in I.

Remember the example

$A(X) \wedge B(Y) \wedge X \not\approx Y \rightarrow \bot, \quad B(X) \wedge C(Y) \wedge X \not\approx Y \rightarrow \bot,$
$A(X) \wedge C(X) \rightarrow \bot,$

$\rightarrow \exists x\ A(x), \quad \rightarrow \exists x\ B(x), \quad \rightarrow \exists x\ C(x).$

(empty interpretation),
$A(e_0)$,
$A(e_0),\ B(e_0)$,
$A(e_0),\ B(e_0),\ C(e_0)$,
$A(e_0),\ B(e_0),\ C(e_1)$,
$A(e_0),\ B(e_1)$.

(backtracking complete)

An example with disjunction:

$\to \exists x \; A(x),$
$A(X) \to B(X) \vee C(X), \quad A(X) \wedge B(X) \to \bot, \quad C(X) \to \bot.$

$$
\begin{aligned}
&\text{(empty interpretation)},\\
&A(e_0),\\
&A(e_0),\; B(e_0),\\
&A(e_0),\; C(e_0).
\end{aligned}
$$

(backtracking complete)

Evaluation of Model Search

- A clever implementation of naive model search performs better than I expected.
- Much depends on the selection strategy. (i.e. which applicable rule is expanded first)
- But, of course, this algorithm will never be seriously competitive.

How to improve?

$\Rightarrow$ Avoid work being redone, add learning.

Model Search with Learning

The search algorithm backtracks using an interpretation I. It maintains a set of geometric formulas $\mathcal{G}$.

Consider a recursive implementation **search**$(I, \mathcal{G})$. The improved version has the following invariant:

At every time when it returns (including returns from recursive calls) :

Either I has been extended to a complete model (no rules in $\mathcal{G}$ are applicable),
or there is a rule of form $\forall \overline{x}\ \Phi(\overline{x}) \rightarrow \bot$ in $\mathcal{G}$, that is applicable in I.

The improved algorithm **search**$(I, \mathcal{G})$ has the following structure:

- Either I is a model, or we can find a rule $\forall \overline{x}\ \Phi(\overline{x}) \rightarrow Z(\overline{x})$ that is applicable with substitution Θ.
- The algorithm successively tries interpretations $I \cup \{A_1\}, \ldots, I \cup \{A_r\}$.
- If none of them resulted in a model, we have for each $I \cup \{A_j\}$ an applicable rule of form $\forall \overline{y}_j\ \Phi_j(\overline{y}_j) \rightarrow \bot$.
- What we need is a calculus that allows to make an inference from $\forall \overline{x}\ \Phi(\overline{x}) \rightarrow Z(\overline{x})$ and

$$\forall \overline{y}_1\ \Phi_1(\overline{y}_1) \rightarrow \bot, \ldots, \forall \overline{y}_r\ \Phi_r(\overline{y}_r) \rightarrow \bot.$$

The result must have form $\forall \overline{z}\ \Psi(\overline{z}) \rightarrow \bot$ and be applicable in I.

Rules for Lemma Learning

A complete calculus can be obtained by the following three rules:

- Instantiation (followed by merging)
- Disjunction resolution.
- Existential resolution.

Lemma Factoring:

Let $\lambda =$

$$\forall \overline{x}\ A_1(\overline{x}) \wedge A_2(\overline{x}) \wedge \cdots \wedge A_p(\overline{x}) \wedge x_1 \not\approx x_1' \wedge \cdots \wedge x_q \not\approx x_q' \rightarrow \bot,$$

be a lemma. Let Σ be a substitution of form $\{y := y'\}$. Then the following lemma is a factor of λ:

$$\forall \overline{x}\Sigma\ \ A_1(\overline{x}\Sigma) \wedge \cdots \wedge A_p(\overline{x}\Sigma) \wedge\ x_1\Sigma \not\approx x_1'\Sigma\ \wedge \cdots \wedge\ x_q\Sigma \not\approx x_q'\Sigma\ \rightarrow \bot.$$

Disjunction Resolution:

Let $\rho =$

$$\forall \overline{x}\ \Phi(\overline{x}) \rightarrow B_1(\overline{x}) \vee \cdots \vee B_q(\overline{x})$$

be a disjunctive formula.

Let $\lambda =$

$$\forall \overline{y}\ D_1(\overline{y}) \wedge \cdots \wedge D_r(\overline{y}) \wedge y_1 \not\approx y_1' \wedge \cdots \wedge y_s \not\approx y_s' \rightarrow \perp,$$

be a lemma, s.t. $B_1(\overline{x})$ and $D_1(\overline{y})$ are unifiable. Then the following formula is a disjunction resolvent of ρ and λ:

$$\forall\ \ \overline{x}\Sigma\ \ \overline{y}\Sigma\ \ \Phi(\overline{x})\Sigma \wedge$$

$$D_2(\overline{y})\Sigma \wedge \cdots \wedge D_r(\overline{y})\Sigma\ \wedge\ y_1\Sigma \not\approx y_1'\Sigma\ \wedge \cdots \wedge\ y_s\Sigma \not\approx y_s'\Sigma\ \rightarrow$$

$$B_2(\overline{x})\Sigma \vee \cdots \vee B_q(\overline{x})\Sigma.$$

Existential Resolution:

Let $\rho =$

$$\forall \overline{x}\ \Phi(\overline{x}) \rightarrow \exists y\ B(\overline{x}, y)$$

be an existential formula.

Let $\lambda =$

$$\forall \overline{z}\ v \quad \Psi(\overline{z}) \wedge B(\overline{z}, v) \wedge v \not\approx z_1 \wedge \cdots \wedge v \not\approx z_s \rightarrow \bot,$$

be a lemma, s.t. $B(\overline{x}, y)$ and $B(\overline{z}, v)$ are unifiable and $v \notin \overline{z}$. Then the following formula is an existential resolvent of ρ and λ:

$$\forall\ \overline{x}\Sigma \quad \overline{z}\Sigma \quad \Phi(\overline{x})\Sigma \wedge \Psi(\overline{z})\Sigma \rightarrow B(\overline{z}, z_1)\Sigma \vee \cdots \vee B(\overline{z}, z_s)\Sigma.$$

Providing some Evidence

Suppose we have $I = \; p(e_0), \;\; q(e_0).$

Assume that the applicable rule is:

$$p(X) \rightarrow r(X) \vee s(X).$$

Assume that $p(e_0), \;\; q(e_0), \;\; r(e_0)$ has applicable rule

$$r(X) \rightarrow \perp.$$

Assume that $p(e_0), \;\; q(e_0), \;\; s(e_0)$ has applicable rule

$$q(X) \wedge s(X) \rightarrow \perp.$$

By disjunction resolution, one can obtain:

$$p(X) \wedge q(X) \rightarrow \perp.$$

Existential Resolution

The simplest form of existential resolution is:

From

$$p(X, Y) \rightarrow \exists z \; q(X, Y, z)$$

and

$$q(X, Y, Z) \wedge r(X, Y) \rightarrow \bot$$

derive

$$p(X, Y) \wedge r(X, Y) \rightarrow \bot.$$

Existential Resolution (2)

Now suppose we have

$$p(X, Y) \to \exists z\ q(X, Y, z)$$

and

$$q(X, Y, Z) \wedge Z \not\approx X \wedge r(X, Y) \to \bot.$$

The second rule refutes almost all possible choices for Z, except the case where $Z \approx X$.

Therefore, we must keep this possibility in the conclusion:

$$p(X, Y) \wedge r(X, Y) \to q(X, Y, X).$$

Existential Resolution (3)

Similarly,

$$p(X, Y) \rightarrow \exists z\ q(X, Y, z)$$

and

$$q(X, Y, Z) \wedge Z \not\approx X \wedge Z \not\approx Y \wedge r(X, Y) \rightarrow \perp$$

result in

$$p(X, Y) \wedge r(X, Y) \rightarrow q(X, Y, X) \vee q(X, Y, Y).$$

Providing Evidence for Existential Resolution

Suppose that we have $I = p(e_0)$.

Assume that the applicable rule is $\rightarrow \exists y\ q(y)$.

Assume that $p(e_0),\ \ q(e_0)$ has applicable rule

$$p(X) \wedge q(X) \rightarrow \bot.$$

Assume that $p(e_0),\ \ q(e_1)$ has applicable rule

$$p(X) \wedge q(Y) \wedge X \not\approx Y \rightarrow \bot.$$

Existential resolution gives

$$p(X) \rightarrow q(X).$$

Disjunction resolution results in

$$p(X) \rightarrow \bot.$$

Theorem: What I did in the examples, can always be done.

We have an implementation of this calculus, which is called geo. it took part in this year's CASC. It solved:

FOF: 73 out of 150,
CNF: 45 out of 150,
SAT: 51 out of 100,
UEQ: 2 out of 100.

This is not bad for a first time, but there is still some work to do.

Conclusions, Future Work

- We gave a new calculus, which is somewhat similar to resolution, and which is refutationally complete for first-order logic.
- Since the algorithm provides an implicit completeness proof, this calculus could be used for saturation-based theorem proving.
- But we do not recommand this: The calculus is intended to be used in combination with the model search algorithm.
- In the implementation, understand which lemmas should be forgotten. Find good heuristics. Develope an intuition of how it searches, and what the proofs mean.
- Extend calculus? (theories, well-behaved infinite models)