Navigating Proof Search in Theorem Provers

A Brief Tour of ATP in Type Theory

Presented by Yanning Chen @ ProSE Seminar



¹https://www.youtube.com/watch?v=VAHyU6gHDdk



¹https://www.youtube.com/watch?v=VAHyU6gHDdk



¹https://www.youtube.com/watch?v=VAHyU6gHDdk



¹https://www.youtube.com/watch?v=VAHyU6gHDdk



1

¹https://www.youtube.com/watch?v=VAHyU6gHDdk

ATP, as a PL Problem

The good old C.H. Correspondence:

Logic Side	Programming Side
proposition	type
proof	term
proposition is true	type has an inhabitant
proposition is false	type does not have an inhabitant
proving a proposition	finding an inhabitant

ATP, as a PL Problem

The good old C.H. Correspondence:

Logic Side	Programming Side
proposition	type
proof	term
proposition is true	type has an inhabitant
proposition is false	type does not have an inhabitant
proving a proposition	finding an inhabitant

ATP is the **type-inhabitation problem**:

Given a proposition P, find a term of type P.

```
inductive Or (a b : Prop) : Prop
| inl : a \rightarrow a \lor b
| inr : b \rightarrow a \lor b
```

Or.inl : ∀ {a b : Prop}, a → a v b

Or.inr : $\forall \{a \ b : \Pr \}$, $b \rightarrow a \ v \ b$

Or.elim : \forall {a b c : Prop}, a v b \rightarrow (a \rightarrow c) \rightarrow (b \rightarrow c) \rightarrow c

Or.inl : $\forall \{a \ b : \Pr \}$, $a \rightarrow a \lor b$

Or.inr : $\forall \{a \ b : \Pr \}$, $b \rightarrow a \ v \ b$

Or.elim : \forall {a b c : Prop}, a v b \rightarrow (a \rightarrow c) \rightarrow (b \rightarrow c) \rightarrow c

Theorem: $\forall (p \ q : \operatorname{Prop}), P \lor Q \to Q \lor P$

example: \forall (P Q: Prop), P v Q \rightarrow Q v P := ?

Or.inl : $\forall \{a \ b : \Pr \}$, $a \rightarrow a \lor b$

Or.inr : $\forall \{a \ b : \Pr \}$, $b \rightarrow a \ v \ b$

Or.elim : \forall {a b c : Prop}, a v b \rightarrow (a \rightarrow c) \rightarrow (b \rightarrow c) \rightarrow c

Theorem: $\forall (p \ q : \operatorname{Prop}), P \lor Q \to Q \lor P$

example: \forall (P Q: Prop), P v Q \rightarrow Q v P := λ P Q hor => ?_

Or.inl : $\forall \{a \ b : \Pr \}$, $a \rightarrow a \lor b$

Or.inr : $\forall \{a \ b : \Pr \}$, $b \rightarrow a \ v \ b$

Or.elim : \forall {a b c : Prop}, a v b \rightarrow (a \rightarrow c) \rightarrow (b \rightarrow c) \rightarrow c

Theorem: $\forall (p \ q : \operatorname{Prop}), P \lor Q \to Q \lor P$

example: \forall (P Q: Prop), P v Q \rightarrow Q v P :=

λ P Q hor => Or.elim hor ?_ ?_

Or.inl : $\forall \{a \ b : \Pr \}$, $a \rightarrow a \lor b$

Or.inr : $\forall \{a \ b : \Pr \}$, $b \rightarrow a \ v \ b$

Or.elim : $\forall \{a \ b \ c : \ \mathsf{Prop}\}, a \ v \ b \rightarrow (a \rightarrow c) \rightarrow (b \rightarrow c) \rightarrow c$

Theorem: $\forall (p \ q : \operatorname{Prop}), P \lor Q \to Q \lor P$

Or.inl : $\forall \{a \ b : \Pr \}$, $a \rightarrow a \lor b$

Or.inr : $\forall \{a \ b : \Pr \}$, $b \rightarrow a \ v \ b$

Or.elim : $\forall \{a \ b \ c : \ \mathsf{Prop}\}, a \ v \ b \rightarrow (a \rightarrow c) \rightarrow (b \rightarrow c) \rightarrow c$

Theorem: $\forall (p \ q : \operatorname{Prop}), P \lor Q \to Q \lor P$

```
example: \forall (P Q: Prop), P v Q \rightarrow Q v P :=
```

```
\lambda P Q hor => Or.elim hor
```

```
(λ hp => Or.inr ?_)
```

 $(\lambda hq => 0r.inl ?_)$

Or.inl : $\forall \{a \ b : \Pr \}$, $a \rightarrow a \lor b$

Or.inr : $\forall \{a \ b : \Pr \}$, $b \rightarrow a \ v \ b$

Or.elim : $\forall \{a \ b \ c : \ \mathsf{Prop}\}, a \ v \ b \rightarrow (a \rightarrow c) \rightarrow (b \rightarrow c) \rightarrow c$

Theorem: $\forall (p \ q : \operatorname{Prop}), P \lor Q \to Q \lor P$

example: \forall (P Q: Prop), P v Q \rightarrow Q v P :=

 λ P Q hor => Or.elim hor

 $(\lambda hp => 0r.inr hp)$

 $(\lambda hq => 0r.inl hq)$

What are these "holes"?

```
example: \forall (P Q: Prop), P v Q \rightarrow Q v P :=

\lambda P Q hor => Or.elim h

(\lambda hp => Or.inr ?1)

(\lambda hq => Or.inl ?2)
```

What are these "holes"?

```
example: \forall (P Q: Prop), P v Q \rightarrow Q v P :=

\lambda P Q hor => Or.elim h

(\lambda hp => Or.inr ?1)

(\lambda hq => Or.inl ?2)
```

- Metavariable: A meta-level placeholder for an object-level term. It might have an assigned value (?1 = hp), partially assigned (?x [a ↦ 1]), or completely unknown.
- **Goal**: Any metavariable that's not completely assigned.

What are these "holes"?

```
example: \forall (P Q: Prop), P v Q \rightarrow Q v P :=

\lambda P Q hor => Or.elim h

(\lambda hp => Or.inr ?1)

(\lambda hq => Or.inl ?2)
```

$$\begin{array}{l} ?1: \forall (P\ Q: \operatorname{Prop}), h_{or}: P \lor Q, h_p: P \vdash P \\ ?2: \forall (P\ Q: \operatorname{Prop}), h_{or}: P \lor Q, h_q: Q \vdash Q \end{array}$$

Local Context: Available premises and the target type for this hole.

```
example : \forall (P Q: Prop), P v Q \rightarrow Q v P := by
  intro p q h
  cases h
  . case inl hp =>
    right
    exact hp
  . case inr hq =>
    left
    exact hq
```

 $\left(\vdash \forall (p,q:\operatorname{Prop}), p \lor q \to q \lor p \right)$

example: \forall (P Q: Prop), P v Q \rightarrow Q v P := ?_



example: \forall (P Q: Prop), P v Q \rightarrow Q v P := λ P => ?_



example: \forall (P Q: Prop), P v Q \rightarrow Q v P := λ P Q => ?_



example: \forall (P Q: Prop), P v Q \rightarrow Q v P := λ P Q hor => ?_





-⊳[



example:
$$\forall$$
 (P Q: Prop), P v Q \rightarrow Q v P :=
 λ P Q hor => Or.elim hor
(λ hp => Or.inr hp)
(λ hq => Or.inl hq)

Two Paradigms: term vs tactic

- **Term based**: giving a direct witness of type inhabitation
- **Tactic based**: refining the term by filling "holes" with meta-level instructions

(x: T, p: P x): ∃ x: T, P x

(x: T, p: P x): ∃ x: T, P x

theorem exists_nat_eq_3: 3 n: Nat, n = 3 := ...

```
(x: T, p: P x): ∃ x: T, P x
```

```
theorem exists_nat_eq_3: 3 n: Nat, n = 3 := by
refine (3, ?_)
rfl
```

(x: T, p: P x): ∃ x: T, P x

```
(x: T, p: P x): ∃ x: T, P x
```

```
theorem exists_nat_eq_3: ∃ n: Nat, n = 3 := by
refine (?_, ?_)
pick_goal 2
rfl
```

```
(x: T, p: P x): ∃ x: T, P x
```

```
theorem exists_nat_eq_3: ∃ n: Nat, n = 3 := by
refine (?_, ?_)
pick_goal 2
    rfl -- closes 1st goal automatically via unification
```

Intuition:

- 1. The difficulty of constructing a proof sometimes depends on the order one chooses to fill the holes.
- 2. Solving a goal might help solve another goal.

 $\langle x: T, p: P x \rangle: \exists x: T, P x$

```
theorem exists_nat_eq_3: ∃ n: Nat, n = 3 := by
refine (?_, ?_)
pick_goal 2
    rfl -- closes 1st goal automatically via unification
```

Two Representations: A view shift from *Search View* to *Presentation View*.

Three Views of Proofs¹

• **Presentation View**: Polished proof for presentation (natural language) or verification (machine readable). Contains "*magical*" values.

¹Aniva et al. (TACAS 2025). Pantograph: A Machine-to-Machine Interaction Interface for Advanced Theorem proving, ... in Lean 4. ²R. L. Morris. Motivated Proofs: What They Are, Why They Matter and How to Write Them
Three Views of Proofs¹

• **Presentation View**: Polished proof for presentation (natural language) or verification (machine readable). Contains "*magical*" values.

E.g. "Let
$$\varepsilon > 0$$
. Define $\delta := \min(2\varepsilon^2, \frac{\varepsilon}{5}, \frac{1}{3})$, ..."

¹Aniva et al. (TACAS 2025). Pantograph: A Machine-to-Machine Interaction Interface for Advanced Theorem proving, ... in Lean 4. ²R. L. Morris. Motivated Proofs: What They Are, Why They Matter and How to Write Them

Three Views of Proofs¹

• **Presentation View**: Polished proof for presentation (natural language) or verification (machine readable). Contains "*magical*" values.

E.g. "Let $\varepsilon > 0$. Define $\delta := \min(2\varepsilon^2, \frac{\varepsilon}{5}, \frac{1}{3})$, ..."

• **Search (proof) View**: The trajectory a prover (machine/human) follows, i.e. *Motivated Proof*².

"Let $\varepsilon > 0$. Let $\delta := ?1$. Using lemma ..., we know $\delta \leq \frac{\varepsilon}{5}$, ... Therefore we can set $\delta := \min(...)$."

¹Aniva et al. (TACAS 2025). Pantograph: A Machine-to-Machine Interaction Interface for Advanced Theorem proving, ... in Lean 4. ²R. L. Morris. Motivated Proofs: What They Are, Why They Matter and How to Write Them

Three Views of Proofs¹

- **Presentation View**: Polished proof for presentation (natural language) or verification (machine readable). Contains "*magical*" values.
- Search (proof) View: The trajectory a prover (machine/human) follows, i.e. *Motivated Proof*².
- **Kernel View**: The proof term itself. Might contain metavars (holes) if unfinished.

¹Aniva et al. (TACAS 2025). Pantograph: A Machine-to-Machine Interaction Interface for Advanced Theorem proving, ... in Lean 4. ²R. L. Morris. Motivated Proofs: What They Are, Why They Matter and How to Write Them

How humans prove

- 1. Prove a theory (either by constructing terms by hand or by using tactics) in the **Search View**.
 - Term based: a tree of *terms (with holes) attempted*
 - Tactic based: a tree of *tactic sequence applied*

How humans prove

- 1. Prove a theory (either by constructing terms by hand or by using tactics) in the **Search View**.
 - Term based: a tree of *terms (with holes) attempted*
 - Tactic based: a tree of *tactic sequence applied*
- 2. Close the proof
 - If term based: Kernel View proof
 - If tactic based:
 - 1. Polish the proof
 - 2. Get a **Presentation View** proof (→ **Kernel View**)

How humans prove



How machines prove: a historical survey of ATP

- 1. Brute-force term search (90s)
- 2. Hammer & SMT bridge (2000-)
- 3. Data-driven (201?-)

The most intuitive and straightforward way.

Theorem: $\forall (p \ q : \operatorname{Prop}), P \lor Q \to Q \lor P$



The most intuitive and straightforward way.

```
Theorem: \forall (p \ q : \operatorname{Prop}), P \lor Q \to Q \lor P
```

 $\lambda p q h \Rightarrow ?_$

The most intuitive and straightforward way.

Theorem: $\forall (p \ q : \operatorname{Prop}), P \lor Q \rightarrow Q \lor P$

λ p q h => 0r.inl ?_ 🤤

The most intuitive and straightforward way.

Theorem: $\forall (p \ q : \operatorname{Prop}), P \lor Q \rightarrow Q \lor P$

λ p q h => Or.inr ?_ 🤤

The most intuitive and straightforward way.

```
Theorem: \forall (p \ q : \operatorname{Prop}), P \lor Q \rightarrow Q \lor P
```

 $\lambda p q h \Rightarrow 0r.elim ?_?_$

The most intuitive and straightforward way.

Theorem: $\forall (p \ q : \operatorname{Prop}), P \lor Q \to Q \lor P$

 $\lambda p q z \Rightarrow 0r.elim (\lambda h \Rightarrow ?_) (\lambda h \Rightarrow ?_)$

The most intuitive and straightforward way.

Theorem: $\forall (p \ q : \operatorname{Prop}), P \lor Q \to Q \lor P$

 $\lambda p q z \Rightarrow 0r.elim (\lambda h \Rightarrow 0r.inr h) (\lambda h \Rightarrow 0r.inl h)$

The most intuitive and straightforward way.

Theorem: $\forall (p \ q : \operatorname{Prop}), P \lor Q \to Q \lor P$

 $\lambda p q z \Rightarrow 0r.elim (\lambda h \Rightarrow 0r.inr h) (\lambda h \Rightarrow 0r.inl h)$

Mimics how humans prove by hand, sans the intuition and experience.

The most intuitive and straightforward way.

Theorem: $\forall (p \ q : \operatorname{Prop}), P \lor Q \to Q \lor P$

 $\lambda p q z \Rightarrow 0r.elim (\lambda h \Rightarrow 0r.inr h) (\lambda h \Rightarrow 0r.inl h)$

Mimics how humans prove by hand, sans the intuition and experience.



The most intuitive and straightforward way.

```
Theorem: \forall (p \ q : \operatorname{Prop}), P \lor Q \rightarrow Q \lor P
```

 $\lambda p q z \Rightarrow 0r.elim (\lambda h \Rightarrow 0r.inr h) (\lambda h \Rightarrow 0r.inl h)$

Mimics how humans prove by hand, sans the intuition and experience.



Problem: huge search space



 Cut out a decidable fragment (e.g. QF_uf)



- Cut out a decidable fragment (e.g. QF_uf)
- 2. Find out related theorems



RE - General ATP

- Cut out a decidable fragment (e.g. QF_uf)
- 2. Find out related theorems
- 3. Send them to a solver



RE - General ATP

- Cut out a decidable fragment (e.g. QF_uf)
- 2. Find out related theorems
- 3. Send them to a solver
- 4. Reconstruct the proof



RE - General ATP

- Cut out a decidable fragment (e.g. QF_uf)
- 2. Find out related theorems
- 3. Send them to a solver
- 4. Reconstruct the proof





RE - General ATP

- Cut out a decidable fragment (e.g. QF_uf)
- 2. Find out related theorems
- 3. Send them to a solver
- 4. Reconstruct the proof



The data-driven way (The magic blackbox 😨)



The data-driven way (The magic blackbox 😨)



Challenges

• **Premise Selection**: Which constructor/lemma to use? A huge headache: induction!

¹Limperg and Halkjær. (CPP 2023). Aesop: White-Box Best-First Proof Search for Lean.

Challenges

- **Premise Selection**: Which constructor/lemma to use? A huge headache: induction!
- Search Space Pruning: Do not revisit the same state.

¹Limperg and Halkjær. (CPP 2023). Aesop: White-Box Best-First Proof Search for Lean.

Challenges

- **Premise Selection**: Which constructor/lemma to use? A huge headache: induction!
- **Search Space Pruning**: Do not revisit the same state.
- Metavar Coupling¹ (Selection): Which goal to solve first?

¹Limperg and Halkjær. (CPP 2023). Aesop: White-Box Best-First Proof Search for Lean.

Recall

```
theorem exists_nat_eq_3: ∃ n: Nat, n = 3 := by
refine (?_, ?_)
pick_goal 2
    rfl -- closes 1st goal automatically via unification
```

Recall

Recall

theorem exists_nat_eq_3:
$$\exists$$
 n: Nat, n = 3 :=
(?1 -- 3 by $n = 3 \sim 3 = 3$, rfl)

Recall

theorem exists_nat_eq_3:
$$\exists$$
 n: Nat, n = 3 :=
(?1 -- 3 by $n = 3 \sim 3 = 3$, rfl)

There's a dependency between **?1** and **?2**¹. So, the algorithm must be able to ...

¹Note that while the term "metavar coupling" was first introduced by aesop, the concept itself is not new. It's a natural consequence of dependent unification.

Recall

theorem exists_nat_eq_3: \exists n: Nat, n = 3 := (?1 -- 3 by $n = 3 \sim 3 = 3$, rfl)

There's a dependency between ?1 and ?2¹.

So, the algorithm must be able to ...

1. keep track of the dependencies and update them during unification.

¹Note that while the term "metavar coupling" was first introduced by aesop, the concept itself is not new. It's a natural consequence of dependent unification.

Recall

theorem exists_nat_eq_3: 3 n: Nat, n = 3 :=

(?1 -- 3 by $n = 3 \sim 3 = 3$, rfl)

There's a dependency between ?1 and ?2¹.

So, the algorithm must be able to ...

- 1. keep track of the dependencies and update them during unification.
- 2. choose wisely which metavariable to solve first.

¹Note that while the term "metavar coupling" was first introduced by aesop, the concept itself is not new. It's a natural consequence of dependent unification.

Case Study

- **Canonical**: Brute-force search, done smartly
- Hammer: When type theory based reasoning meets solvers
- **DSP** (**LLM**): The **panacea**, but at what cost?

Case Study: Canonical

Conventional term-based ID-DFS, with a special type theory.
Conventional term-based ID-DFS, with a special type theory.

Core idea: maintain β -normal η -long form, even through substitution, for free.

 $\begin{array}{ll} (\lambda x.C \ x) \ 1 & \bigotimes \beta \text{-reducible} \\ f: \tau_1 \to \tau_2 & \bigotimes \eta \text{-expandable} \\ \lambda x.fx: \tau_1 \to \tau_2 & \checkmark \beta \text{-normal } \eta \text{-long} \\ (fg)h & \bigotimes \eta \text{-expandable: } (\lambda x.fgx)h \text{, i.e. no partial app} \end{array}$

Maintain β -normal η -long form, even through substitution, for free.

Why? It solves the state pruning problem.

All $\alpha\beta\eta$ -equiv terms are syntactically equal! ($\beta n\eta l + ln$) A hashset of seen terms is enough. No normalization.

Maintain β -normal η -long form, even through substitution, for free.

Why? It (partly) simplifies the premise selection problem.

E.g. h ?1 with local context $C_1:a \to T, C_2:a \to b \to T, D_1:a \to U \vdash T$

The only two viable candidates are

- $C_1: h \ (C_1 \ ? \ 2)$, and
- C_2 : $h(C_1 ? 2 ? 3)$

Maintain β -normal η -long form, even through substitution, for free.

Why? It (partly) solves the premise selection problem.

E.g. id ?1 ?2 with local context $C_1 : a \to T, C_2 : a \to b \to T, D_1 : a \to U \vdash ?t$

We have three candidates.

Maintain β -normal η -long form, even through substitution, for free.

Why? It (partly) solves the premise selection problem.

E.g. id ?1 ?2 with local context $C_1 : a \to T, C_2 : a \to b \to T, D_1 : a \to U \vdash ?t$

If, we are not enforcing $\beta n\eta l$ form, it can be a lambda! (We always have n + 1 choices)

Worse, id $(\Pi_{x:a} ? 1) ? 2 ? 3!$

Maintain β -normal η -long form, even through substitution, for free.

Why? It (partly) solves the premise selection problem.

E.g. id ?1 ?2 with local context $C_1 : a \to T, C_2 : a \to b \to T, D_1 : a \to U \vdash ?t$

But now that we are enforcing $\beta n\eta l$, it's impossible for a metavar to be a lambda:

$$?1:\tau_1 \rightarrow \tau_2 \Longrightarrow \lambda x:\tau_1.(?1:\tau_2)$$

Maintain β -normal η -long form, even through substitution, for free.

Why? It enforces static arity, enabling usage of efficient data structures.

Always full app: { head: Term, args: Array Term }
Partial app: { lhs: Term, rhs: { lhs: Term, rhs: ... }}

Case Study: Canonical, and its Achilles' Heel

Induction is a big problem.

Case Study: Canonical, and its Achilles' Heel

Induction is a big problem.

```
Nat.rec : {motive : Nat → Type} →
motive Nat.zero → ((n : Nat) → motive n → motive n.succ) →
(t : Nat) → motive t
```

motive is universally quantified. It's a candidate for every hole.

Outsourcing the proof search to an external solver (e.g. Z3, Vampire).

 Export subgoals to external solvers, often making dependent types opaque

- 1. Export subgoals to external solvers, often making dependent types opaque
- 2. Filter out a list of likely related theorems and send them to the solver

- 1. Export subgoals to external solvers, often making dependent types opaque
- 2. Filter out a list of likely related theorems and send them to the solver
- 3. Replaying external proofs back to the type theory via reconstruction

	Hammer	ATP
Automation	push-button	varies
Expressiveness	limited to decidable fragment	full support for DT and HoL
Trust	longer trust chain	kernel-checked directly
Performance	highly optimized	slow in large libs or proofs
Visibility	hidden in the solver	full trace available

Outsourcing the proof search to an external solver (e.g. Z3, Vampire).

	Hammer	ATP
Automation	push-button	varies
Expressiveness	limited to decidable fragment	full support for DT and HoL
Trust	longer trust chain	kernel-checked directly
Performance	highly optimized	slow in large libs or proofs
Visibility	hidden in the solver	full trace available

Trade-off between expressiveness and automation.

Outsourcing the proof search to an external solver (e.g. Z3, Vampire).

Challenges:

1. Premise Selection: depending on other solutions (adhoc heuristics, reinforcement learning, etc.)

Outsourcing the proof search to an external solver (e.g. Z3, Vampire).

Challenges:

- 1. Premise Selection: depending on other solutions (adhoc heuristics, reinforcement learning, etc.)
- 2. Induction: no support (recursors are dependently typed)

The good:

1. It works magically! ($\sim 50\%$ success rate on some datasets)

The good:

- 1. It works magically! ($\sim 50\%$ success rate on some datasets)
- 2. Somewhat explainable: the informal proof sketch is human-readable.

The ugly:

• Limited to tactic based approach due to the nature of DSP, thus

The ugly:

- Limited to tactic based approach due to the nature of DSP, thus
- Need to tackle with tactic selection, a slightly different problem with premise selection, and

The ugly:

- Limited to tactic based approach due to the nature of DSP, thus
- Need to tackle with tactic selection, a slightly different problem with premise selection, and
- Need to learn how to parametrize the tactic.

The bad:

1. Neural networks are not explainable.

The bad:

- 1. Neural networks are not explainable.
- 2. Insufficient dataset: most datasets are synthetic, and the realworld datasets are not large enough.

The bad:

- 1. Neural networks are not explainable.
- 2. Insufficient dataset: most datasets are synthetic, and the realworld datasets are not large enough.
- 3. Unreliable benchmarking: data leak problem
 - Avoidable: training set pollution
 - Unavoidable: prior knowledge from natural language pretraining

Best of all three worlds

1. Symbolic approach with a solver heart Regain the power of domain specific tactics.

Best of all three worlds

- 1. Symbolic approach with a solver heart Regain the power of domain specific tactics.
- Symbolic approach with a neuro heart (Domain-specific) learning based premise/metavar selection.

Best of all three worlds

- 1. Symbolic approach with a solver heart Regain the power of domain specific tactics.
- Symbolic approach with a neuro heart (Domain-specific) learning based premise/metavar selection.
- 3. Neuro approach with a symbolic heart Explainable & reliable logic core.

Open problems

1. Premise Selection

Open problems

- 1. Premise Selection
- 2. Lemma Discovery: reuse proofs, generalizing theorems Might help with the induction headache.

Acknowledgements

Leni Aniva, for her brilliant idea of the three views of proof, insights on neuro-symbolic methods, and also her help on the diagrams.

Tesla Zhang, for his clear explanation of the Canonical theorem prover.

Alexander Chichigin, for his attentive review and constructive criticism.