

Aesop: White-Box Automation for Lean 4

Jannis Limperg
University of Munich (LMU)
jannis@limperg.de

BICMR and Istituto Grothendieck (online), 18. April
2024

Search Algorithm

Registering Rules

Built-In Rules

Debugging

Miscellaneous Features

Applications, Shortcomings and Work In Progress

Search Algorithm

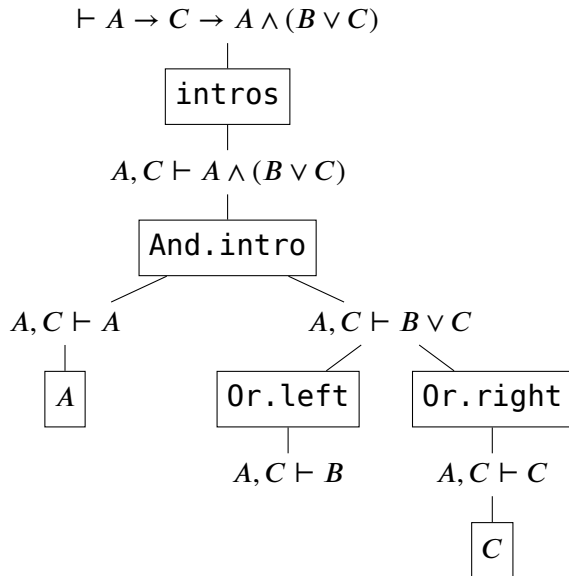
Rules

A *rule* is an arbitrary Lean tactic.

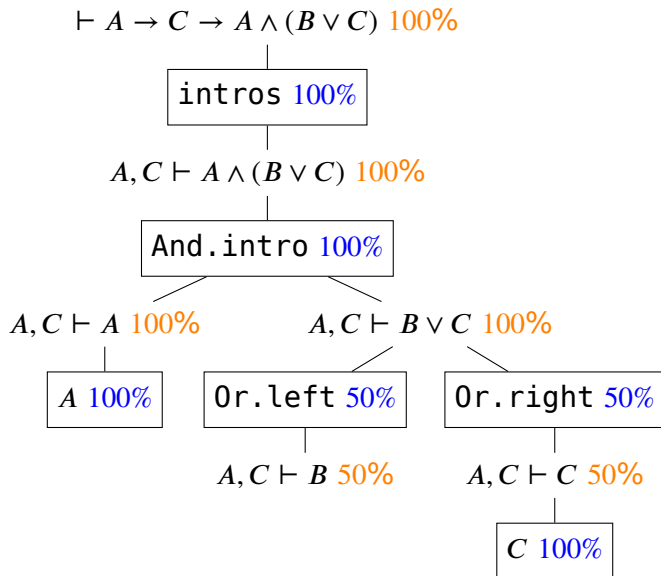
Aesop provides convenient syntax (*rule builders*) for creating rules from theorems.

Aesop always operates with a user-defined *rule set*.

Basic Tree Search



Best-First Tree Search



Safe Rules

- Run before unsafe rules
- If a safe rule succeeds on a goal G , no other rules are tried for G
- Integer penalty
- Treated as 100% success probability
- 🌟 Good for performance
- 😞 Users need to make sure that the rule really is safe

Examples

Safe rule: \wedge -introduction

$$\begin{array}{c} \Gamma \vdash A \wedge B \\ / \quad \backslash \\ \Gamma \vdash A \quad \Gamma \vdash B \end{array}$$

Unsafe rule: left \vee -introduction

$$\begin{array}{c} \Gamma \vdash A \vee B \\ | \\ \Gamma \vdash A \end{array}$$

When Is A Rule Safe?

A rule R is *logically safe* if it preserves provability:

For each goal G , if G is provable and R , applied to G , generates subgoals G_1, \dots, G_n , then G_1, \dots, G_n must still be provable.

A rule R is *relatively safe* if it preserves provability relative to a rule set S :

If a goal G is provable with rules from S and R , applied to G , generates subgoals G_1, \dots, G_n , then G_1, \dots, G_n must still be provable with rules from S .

Normalisation Rules

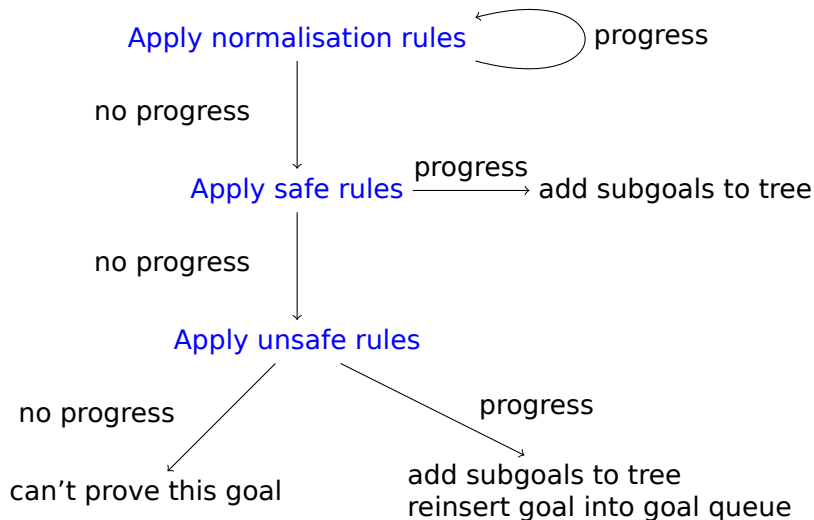
- Run before safe rules
- Integer penalty
- Treated as 100% success probability
- May produce only one subgoal
- Run in a fixpoint loop, i.e. until no normalisation rule succeeds any more
- 🧙 Can establish invariants for other rules
- 😞 Typically run multiple times on every goal

Example

\wedge -elimination

$$\frac{\Gamma, h : A \wedge B \vdash T}{\Gamma, h_1 : A, h_2 : B \vdash T}$$

Summary: Aesop's Search Algorithm



Registering Rules

Registering Rules

Globally

```
@[aesop unsafe 100%]  
theorem And.intro : A → B → A ∧ B
```

Locally

```
aesop (add 100% And.intro)
```

Safe rules

```
@[aesop safe 10]  
theorem And.intro : A → B → A ∧ B
```

Rule Builders

A rule builder turns a declaration into an Aesop rule.

In the examples so far, we have implicitly used a default builder.

Aesop currently provides 7 rule builders.

Apply Builder

```
@[aesop safe apply 10]  
theorem And.intro : A → B → A ∧ B
```

Builds a rule that runs **apply** `And.intro`.

$$\begin{array}{c} \Gamma \vdash A \wedge B \\ / \quad \backslash \\ \Gamma \vdash A \quad \Gamma \vdash B \end{array}$$

Constructors Builder

```
@[aesop 50% constructors]
inductive Or (A B : Prop) where
| left  : A → Or A B
| right : B → Or A B
```

Equivalent to one **apply** rule for each constructor.

Cases Builder

```
@[aesop safe cases]
inductive Or (A B : Prop) where
| left  : A → Or A B
| right : B → Or A B
```

Builds a rule that runs **cases** on any hypothesis of type `Or A B`.

Forward Builder

```
@[aesop safe forward]
```

```
theorem pos_of_min_pos :  $\forall \{x y : \mathbb{N}\},$ 
```

```
  0 < min x y  $\rightarrow$ 
```

```
  0 < x  $\wedge$  0 < y
```

$$\Gamma, x y : \mathbb{N}, h : 0 < \min x y \vdash T$$
$$\Gamma, x y : \mathbb{N}, h : 0 < \min x y, h_1 : 0 < x \wedge 0 < y \vdash T$$

Destruct Builder

```
@[aesop safe destruct]
```

```
theorem pos_of_min_pos :  $\forall \{x y : \mathbb{N}\},$ 
```

```
   $0 < \min x y \rightarrow$ 
```

```
   $0 < x \wedge 0 < y$ 
```

$$\Gamma, x y : \mathbb{N}, h : 0 < \min x y \vdash T$$
$$\Gamma, x y : \mathbb{N}, h : 0 < x \wedge 0 < y \vdash T$$

Simp Builder

Aesop runs `simp_all` as a built-in normalisation rule with penalty 0.

This `simp_all` call uses the default `simp` set plus an Aesop-specific `simp` set.

The `simp` builder adds an equation or proposition to this Aesop-specific set.

Tactic Builder

```
aesop (add safe (by norm_num; done))
```

```
add_aesop_rules safe (by norm_num; done)
```

Requirement

If the tactic does not change the goal, it should fail.

Built-In Rules

Logic: \wedge

\wedge -introduction (safe, penalty 101)

$$\begin{array}{c} \Gamma \vdash A \wedge B \\ / \quad \backslash \\ \Gamma \vdash A \quad \Gamma \vdash B \end{array}$$

\wedge -elimination (norm, penalty 0)

$$\begin{array}{c} \Gamma, h : A \wedge B \vdash T \\ | \\ \Gamma, h_1 : A, h_2 : B \vdash T \end{array}$$

Similar for Prod, PProd, MProd

Logic: \vee

\vee -introduction (unsafe, success probability 50%)

$$\begin{array}{c} \Gamma \vdash A \vee B \\ | \\ \Gamma \vdash A \end{array}$$

$$\begin{array}{c} \Gamma \vdash A \vee B \\ | \\ \Gamma \vdash B \end{array}$$

\vee -elimination (safe, penalty 100)

$$\begin{array}{ccc} & \Gamma, h : A \vee B \vdash T & \\ & \swarrow \quad \searrow & \\ \Gamma, h : A \vdash T & & \Gamma, h : B \vdash T \end{array}$$

Similar for Sum, PSum

Logic: \forall and \rightarrow

\forall -introduction (norm, penalty -100)

Run the intros tactic

\forall -elimination (unsafe, success probability 75%)

$$\begin{array}{c} \Gamma, h : \forall(x_1 : A_1) \dots (x_n : A_n), B \vdash B \\ \quad \swarrow \quad \quad \quad \downarrow \quad \quad \quad \searrow \\ \Gamma, h \vdash A_1 \quad \dots \quad \Gamma, h \vdash A_n \end{array}$$

Logic: \exists

\exists -introduction (unsafe, success probability 30%)

$$\begin{array}{c} \Gamma \vdash \exists x, P x \\ | \\ \Gamma \vdash P ?x \end{array}$$

\exists -elimination (norm, penalty 0)

$$\begin{array}{c} \Gamma, h : \exists x : A, P x \vdash T \\ | \\ \Gamma, x : A, h : P x \vdash T \end{array}$$

Similar for Subtype, Sigma, PSigma

Logic: \leftrightarrow

\leftrightarrow -introduction (safe, penalty 100)

$$\begin{array}{c} \Gamma \vdash A \leftrightarrow B \\ \swarrow \quad \searrow \\ \Gamma \vdash A \rightarrow B \quad \Gamma \vdash B \rightarrow A \end{array}$$

\leftrightarrow hypotheses

A hypothesis of type $A \leftrightarrow B$ is treated as the equation $A = B$ by the simplifier and the substitution rule.

Logic: T

T-introduction¹ (safe, penalty 0)

$$\Gamma \vdash T$$

↓
✓

Similar for Unit, PUnit.

¹T = True

Logic: \perp

The simplifier already solves goals with an assumption $h : \perp$.²

We add destruct rules to conclude \perp from Empty and PEmpty.

² $\perp = \text{False}$

Logic: \neg

\neg -introduction

$$\frac{\Gamma \vdash \neg A}{\Gamma, h : A \vdash \perp}$$

Negated hypotheses

Given a hypothesis of type $\neg A$, the simplifier replaces A with \perp everywhere in the goal.

Logic: Simplifier

The simplifier performs limited logical reasoning. If A and B are propositions:

- With assumption $h : A$: rewrite $A = \top$
- With assumption $h : \neg A$: rewrite $A = \perp$
- $(\top \wedge \perp) = \perp$
- $(\top \wedge \top) = \top$
- $(\top \rightarrow A) = A$
- etc.

Logic: Completeness

In practice, these rules solve most 'purely logical' goals.

However, they are incomplete for first-order and even propositional logic.

Equality

Simplifier (norm, penalty 0)

Run the `simp_all` tactic as described previously

Reflexivity (safe, penalty 0)

Run the `rfl` tactic

Substitution (norm, penalty -50)

Run the `subst` tactic on any hypothesis of type $x = t$ or $t = x$ where x is a variable.

This substitutes t for x everywhere in the goal and removes the now-redundant hypothesis.

Case Splitting

Split target (safe, penalty 100)

Runs the **split** tactic.

This tactic looks for **if-then-else** and **match** expressions in the target and performs case splits on their discriminines.

Split hypotheses (safe, penalty 1000)

Ditto, but we look for case splits in hypotheses.

Extensionality

Extensionality (unsafe, success probability 80%)

Run the ext tactic.

This exhaustively applies extensionality lemmas to an equational goal. E.g.:

$$\begin{array}{c} \Gamma, pq : A \times B \vdash p = q \\ | \\ \Gamma, pq : A \times B \vdash p.1 = q.1 \wedge p.2 = q.2 \end{array}$$

Debugging

Leftover Goals

When Aesop fails to solve a goal, it reports the goals that remain after safe rules have been exhaustively applied.

This helps to check whether the safe rules do what you think they should.

Proof Generation

```
theorem last_cons {a :  $\alpha$ } {l : List  $\alpha$ } (h : l  $\neq$  nil) :  
  last (a :: l) (cons_ne_nil a l) = last l h := by  
  aesop? (add 1% cases List)
```

aesop? generates a proof script:

```
intro h  
unhygienic aesop_cases l  
· simp_all only [last, ne_eq]  
  split  
· simp_all only [last]
```

One click replaces **aesop?** with the generated proof.

Tracing

```
set_option trace.aesop true in  
aesop
```

- Lists the rules that Aesop (tried to) run and the resulting goals
- For other trace options see autocompletion for `trace.aesop`.

Miscellaneous Features

Custom Rule Sets

```
-- RuleSet.lean  
  
-- Must be in a different file for technical reasons  
declare_aesop_rule_sets [Foo]
```

```
import RuleSet  
  
@[aesop 100% (rule_sets [Foo])]  
theorem foo
```

Used for domain-specific automation, e.g. continuity, measurability, ...

Metavariables

Aesop supports rules that generate metavariables:

```
example {a b c d : Nat} (h1 : a < b)
  (h2 : a < c) (h3 : c < d) : a < d := by
apply Nat.lt_trans
-- ⊢ a < ?x
-- ⊢ ?x < d

example {a b c d : Nat} (h1 : a < b)
  (h2 : a < c) (h3 : c < d) : a < d := by
aesop (add 1% Nat.lt_trans)
```

- 🤖 Aesop's search algorithm is conjectured complete even in the presence of metavariables
- 😞 This is very expensive

Applications, Shortcomings and Work In Progress

Applications

- General-purpose automation
- Domain-specific solvers: continuity, measurability, `aesop_cat`, ...
- Domain-specific 'goal preprocessors' with **aesop**?
- Tree search backend for LLMs: Peiyang Song, Kaiyu Yang, and Anima Anandkumar. "Towards Large Language Models as Copilots for Theorem Proving in Lean". In: *Workshop on Mathematical Reasoning and AI*. 2023. URL: <https://mathai2023.github.io/papers/4.pdf>

Shortcomings and Work In Progress

- The built-in logical rules do not deal well with all-quantified hypotheses and sometimes negation
- The default rule set is currently missing many rules
- Repeated calls to **simp** during normalisation are bad for performance
- Tactic scripts generated by **aesop?** are not particularly idiomatic (WIP)
- **aesop?** only works if Aesop solves the goal (WIP)
- Nontrivial sets of forward rules are very slow (WIP)