

The State of Lean

Leo de Moura
Senior Principal Applied Scientist, AWS
Chief Architect, Lean FRO

January 27, 2026



Lean is an open-source programming language and proof assistant that is transforming how we approach mathematics, software verification, and AI.

Lean and its tooling are implemented in Lean. Lean is very **extensible**.

LSP, Parser, Macro System, Elaborator, Type Checker, Tactic Framework, Proof automation, Compiler, Build System, Documentation Authoring Tool.

Lean has a **small trusted kernel**, proofs can be exported and independently checked.

The **Lean FRO** is a nonprofit dedicated to developing Lean.

Lean is impacting Mathematics

*"Lean enables large-scale collaboration by allowing mathematicians to break down complex proofs into smaller, verifiable components. This formalization process ensures the correctness of proofs and facilitates contributions from a broader community. **With Lean, we are beginning to see how AI can accelerate the formalization of mathematics, opening up new possibilities for research.**" — Terence Tao*

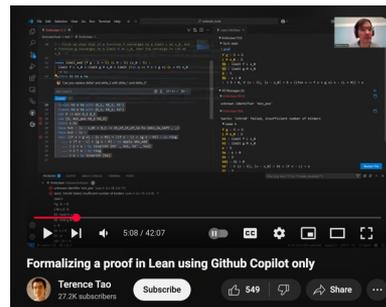
Liquid Tensor Experiment – Peter Scholze

The Equational Theories Project – Terence Tao

Fermat's Last Theorem – Kevin Buzzard

Carleson's Theorem – Floris van Doorn

Abc conjecture



nature

Explore content Journal information Publish with us Subscribe

nature > news > article

NEWS | 18 June 2021

Mathematicians welcome computer-assisted proof in 'grand unification' theory

Startups using Lean & AI



Math, Inc.



Logical Intelligence

...and more to come

Vibe Proving

Harmonic @HarmonicMath · Nov 29 🔄 ...

Mathematical superintelligence is coming, faster than you imagined

Vlad Tenev @vladtenev · Nov 29

We are on the cusp of a profound change in the field of mathematics. Vibe proving is here.

Aristotle from @HarmonicMath just proved Erdos Problem #124 in @leanprover, all by itself. This problem has been open for nearly 30 ... [Show more](#)

🗨️ 17 🔄 40 ❤️ 561 📊 89K 📌 📤

Logical Intelligence @logic_int 🔄 ...

Aleph prover just went BEAST MODE
4 math problems unsolved for 20+ years. Formal proofs in [Lean 4](#). Less than 48 hours. Under \$5k total.

- ✅ Binomial tail bounds conjecture (Telgarsky, 2009)
- ✅ Quantum gate lattice approximation (Greene & Damelin, 2015)*
- ✅ Erdős 124
- ✅ Erdős 481
- ✅ #1 on PutnamBench leaderboard

The era of AI mathematics is here.

Axiom reposted

Carina Hong @CarinaLHong · Dec 2 🔄 ...

AxiomProver solved Erdos problem #481 - took 5 hours

#124 simplified version took over 24 hours (oof) and was not as succinct as we'd like it to be

```

def test_of_lemma_0_of_lean_prover_01 : Prop :=
  -- lemma : k / a^n = 1 / 12 * M / c + 11
  -- From k / a^n = 1 / 12 * M / c + 11, get a^n / k = 240c / (k * 12 * M / c + 11)
  have h1 : (k / a^n) = 1 / 12 * M / c + 11 := by
    have h1 : (k / a^n) = 1 / 12 * M / c + 11 := by
      use (int_div_int_div)
    rw [h1]
  have h2 : 1 / 12 * M / c + 11 = 1 / 12 * M / c + 11 := by
    apply (int_div_int_div)
  simp only [int_div_int_div] at h2
  linarith
  
```

Harmonic @HarmonicMath · Nov 30 🔄 ...

Bartosz Naskręcki @nasqret · Nov 29

Aristotle by @HarmonicMath is acing group-theory puzzles.

Here is a complete formal proof of the popular Yu Tsumura 554 puzzle. What's nice is that the proof is very transparent, with easy-to-follow steps. It was generated in less than an hour without any hints. I am ... [Show more](#)



Lean in Software Verification

[Cedar](#) – Language for defining permissions as policies – AWS

[SampCert](#) – Verified Differential Privacy Primitives – AWS

[SymCrypt](#) – Verified Cryptography – Microsoft

[Neuron Compiler](#) – AWS

 | science

[Research areas](#) ▾ [Blog](#) [Publications](#) [Conferences](#) [Code and datasets](#) [Academia](#) ▾ [Careers](#)



AUTOMATED REASONING

How the Lean language
brings math to coding
and coding to math

Prospective: AI+CS as the Next Frontier is Imminent

Companies are starting to pivot their math-proven AIs towards program verification



Beyond math: Aristotle achieves SOTA 96.8% proof generation on VERINA: Benchmarking Verifiable Code Generation. You can read more about this performance on our engineering blog linked in bio

12:24 PM · Dec 3, 2025 · **1,860** Views



Ilya Sergey
@ilyasergey



The proof of the last remaining Lean theorem in our upcoming conference submission has now been completed with the help of AI.

The research community's perception of program verification is about to change irreversibly.

8:32 PM · Nov 13, 2025 · **15.8K** Views

harmonic.fun/news#blog-post-aristotle-tech-report

Spec auto-formalization + program synthesis + verification as a service



A Foundation for Computer Science in Lean

CSLib: A Mathlib for computer science.

Clark Barrett (Stanford University and Amazon)

Swarat Chaudhuri (Google DeepMind and UT Austin)

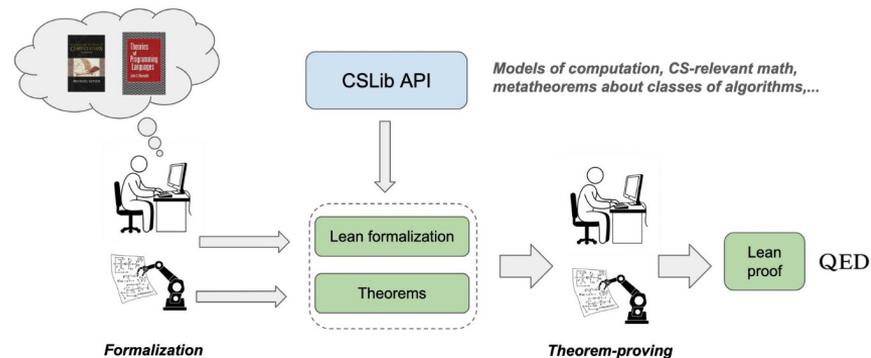
Jim Grundy (Amazon)

Pushmeet Kohli (Google DeepMind)

Leo de Moura (Amazon and Lean FRO)

Fabrizio Montesi (University of Southern Denmark)

Usage scenario: Adding to CS foundations



CSLib aims to be a foundation for **teaching**, **research**, and new **verification** efforts, including AI-assisted.



Lean FRO: Shaping the Future of Lean Development

The Lean Focused Research Organization (FRO) is a non-profit dedicated to Lean's development.

Founded in **August 2023**, the organization has 20 members.

We are very grateful for all philanthropic support we have received.

Its mission is to enhance critical areas: **scalability, proof automation, usability, and documentation.**

It must reach **self-sustainability in August 2028** and become the **Lean Foundation.**

25 releases and **+7,500 pull requests** merged in the main repository only since its launch in July 2023.

We run the FRO as a startup.

Proof Automation: grind



Why do we need proof automation?

"I thought AI would prove all theorems for us now."

AI at the IMO 2024

AlphaProof (Google DeepMind) achieved silver medal level using Lean.

AI at the IMO 2025

Google DeepMind and OpenAI achieved gold medal level using informal reasoning.

[ByteDance achieved silver* medal](#) using Lean. (*) [They reached gold after the competition.](#)

[Harmonic achieved gold medal](#) using Lean.

AI is playing the “Lean game”

The **moves** in this game are **tactics** from Automated Reasoning: good old proof automation.

Here are some “moves” played by AlphaProof:

```
simp_all[Finset.sum_range_id]
```

```
zify[*]at*
```

```
norm_num at*
```

```
nlinarith[(by norm_cast:(c:ℝ) >=A*(1-[ ])+[ ]+1), Int.floor_lex, Int.lt_floor_add_one x]
```

Even the most advanced AI relies on the same tactics we use every day.

By developing **better moves/tactics**, we enable even **more powerful AI**.



What is grind?

New proof automation (since Lean v4.22) developed by Kim Morrison and myself.

A proof-automation tactic **inspired by modern SMT solvers**. Think of it as a **virtual whiteboard**:

- Discovers new equalities, inequalities, etc.

- Writes facts on the board and merges equivalent terms

- Multiple engines cooperate on the same workspace

Cooperating Engines:

- Congruence closure; E-matching; Constraint propagation; Guided case analysis

- Satellite theory solvers (linear integer arithmetic, commutative rings, linear arithmetic, AC)

Supports dependent types, type-class system, and dependent pattern matching

Produces ordinary Lean proof terms for every fact.



What grind is NOT

Not designed for combinatorially explosive search spaces:

- Large-n pigeonhole instances

- Graph-coloring reductions

- High-order N-queens boards

- 200-variable Sudoku with Boolean constraints

Why? These require thousands/millions of case-splits that overwhelm grind's branching search

Key takeaway: grind excels at cooperative reasoning with multiple engines, but struggles with brute-force combinatorial problems.

For massive case-analysis, use `bv_decide`



grind: Design Principles

Native to **Dependent Type Theory**: No translation to first-order or higher-order logic needed.

Solves trivial goals automatically.

Fast startup time: No server startup, no external tool dependencies, no translations

Rich **diagnostics**: When it fails, it tells you why.

Configurable via Type Classes and annotations.

Provide **grind?** similarly to `bv_decide?` and `aesop?`

Interactive mode for full control

Stdlib and Mathlib pre-annotated.



grind: Architecture

Preprocessing: normalization, canonicalization, extracting nested proofs, hash-consing, ...

Internalization: process of converting Lean expressions into solver's internal data-structures.

E-graph: congruence closure, E-matching, constraint propagation.

Satellite Solvers: linear integer arithmetic, commutative rings, linear arithmetic, AC, orders, etc.

grind: Model-based theory solvers

For linear arithmetic (linarith) and linear integer arithmetic (lia).

linarith is parametrized by a Module over the integers. It supports preorders, partial orders, and linear orders.

"I'm interested in developing some API for linearly ordered vector spaces, in order to easily handle manipulations of asymptotic orders" – Terence Tao on the Lean Zulip

```
example {R} [OrderedVectorSpace R] (x y z : R)
  : x ≤ 2•y → y < z → x < 2•z := by
  grind -- 🦄
```

OrderedVectorSpace implements IntModule, LinearOrder, IntModule.IsOrdered.



grind: Model-based theory solvers

lia is parametrized by the ToInt type class used to embed types such as Int32, BitVec 64 into the integers

```
/--  
The embedding into the integers takes addition to addition, wrapped into the range interval.  
-/  
class ToInt.Add (α : Type u) [Add α] (I : outParam IntInterval) [ToInt α I] where  
  /-- The embedding takes addition to addition, wrapped into the range interval. -/  
  toInt_add : ∀ x y : α, toInt (x + y) = I.wrap (toInt x + toInt y)  
  
/--  
The embedding into the integers is monotone.  
-/  
class ToInt.LE (α : Type u) [LE α] (I : outParam IntInterval) [ToInt α I] where  
  /-- The embedding is monotone with respect to `≤`. -/  
  le_iff : ∀ x y : α, x ≤ y ↔ toInt x ≤ toInt y
```

grind: Model-based theory solvers

```
example (x y : Int) :  
  27 ≤ 11*x + 13*y → 11*x + 13*y ≤ 45 →  
  -10 ≤ 7*x - 9*y → 7*x - 9*y ≤ 4 → False := by  
grind
```

```
example (a b c : UInt32) :  
  -a + -c > 1 →  
  a + 2*b = 0 →  
  -c + 2*b = 0 → False := by  
grind
```

```
example (a : Fin 4) : 1 < a → a ≠ 2 → a ≠ 3 → False := by grind
```

grind: commutative rings and fields

Support for commutative rings and fields uses Grobner basis.

Parametrized by the type classes: CommRing, CommSemiring, NoNatZeroDivisors, Field, AddRightCancel, and IsCharP

```
example {α} [CommRing α] (a b c : α)
  : a + b + c = 3 →
    a^2 + b^2 + c^2 = 5 →
    a^3 + b^3 + c^3 = 7 →
    a^4 + b^4 + c^4 = 9 := by
  grind
```

```
example [Field R] (a : R) : (2 * a)^-1 = a^-1 / 2 := by grind
```

```
example [Field R] (a : R) : (2 : R) ≠ 0 → 1 / a + 1 / (2 * a) = 3 / (2 * a) := by grind
```

```
example [Field R] [IsCharP R 0] (a : R) : 1 / a + 1 / (2 * a) = 3 / (2 * a) := by grind
```

```
example (x y : BitVec 16) : x^2*y = 1 → x*y^2 = y → y*x = 1 := by grind
```

Associative (commutative, idempotent) operators & neutral elements

This is not a toy: it encodes a real algebraic constraint derived from relations among diagrams in a pivotal tensor category.

Parametrized by the type classes: Associative, Commutative, IdempotentOp, Lawfullidentity.

Coming soon: AC E-matching.

```
example {α} (f : α → α) (op : α → α → α) [Associative op] [Commutative op] (a b : α)
  : op (f (op a b)) b = op b (f (op b a)) := by
  grind only
```

```
example {α} (bar : α → α) (op : α → α → α) [Associative op] [IdempotentOp op]
  (a b c d e f x y w : α)
  : op d (op x c) = op a b →
    op e (op f (op y w)) = op (op d a) (op b c) →
    bar (op d (op x c)) = bar (op e (op f (op y w))) := by
  grind only
```

```
example (a b c : Nat) : min a (max b c) = min (max c b) a := by
  grind -lia only
```

```
example (a b c : Nat) : min a (max b (max c 0)) = min (max c b) a := by
  grind -lia only
```

grind: E-matching

E-matching is a heuristic for instantiating theorems. It is used in many SMT solvers.

It is matching modulo equalities.

```
@[grind =] theorem fg {x} : f (g x) = x := by
  unfold f g; omega
```

```
example {a b c} : f a = b → a = g c → b = c := by
  grind
```

```
-- Whenever `grind` sees `cos` or `sin`, it adds `(cos x)^2 + (sin x)^2 = 1` to the whiteboard.
-- That's a polynomial, so it is sent to the Grobner basis module.
-- And we can prove equalities modulo that relation!
```

```
example {x} : (cos x + sin x)^2 = 2 * cos x * sin x + 1 := by
  grind
```



grind: Extensibility

You can configure grind using type classes.

You can annotate theorems and definitions with the `[grind]` attributes.

```
@[grind =]
theorem getElem?_cons {a l i} : (a :: l)[i]? = if i = 0 then some a else l[i-1]? := by
  cases i <|> simp
```

```
@[grind →]
theorem getElem_of_getElem? {a i a} {l : List a} : l[i]? = some a → ∃ h : i < l.length, l[i] = a :=
  getElem?_eq_some_iff.mp
```

That said, grind is implemented in Lean, and you can extend its implementation using Lean itself.

No need to learn another programming language, or how to create shared objects.



grind: Extensibility - constraints

You can associate constraints to theorem instantiation patterns.

```
theorem shiftLeft_add (m n : Nat) :  $\forall k, m \lll (n + k) = (m \lll n) \lll k$   
  | 0 => rfl  
  | k + 1 => by simp [ $\leftarrow$  Nat.add_assoc, shiftLeft_add _ _ k, shiftLeft_succ]  
  
grind_pattern shiftLeft_add => m  $\lll (n + k)$  where  
  m  $\neq 0$ 
```

```
theorem div_pow_of_pos (a n : Nat) :  $n > 0 \rightarrow a \mid a ^ n :=$  by...  
  
grind_pattern div_pow_of_pos => a ^ n where  
  is_value a  
  guard n > 0
```

```
protected theorem dvd_mul_left_of_dvd {a b : Nat} (h : a  $\mid$  b) (c : Nat) : a  $\mid$  c * b :=...  
  
grind_pattern Nat.dvd_mul_left_of_dvd => a  $\mid$  b, c * b where  
  guard a  $\mid$  b
```



grind: Extensibility - solvers

You can plugin your own solver. We implemented all built-in solvers using the plugin API.

```
/-- State for all associative operators detected by `grind`. -/  
structure State where  
  /--  
  Structures/operators detected.  
  We expect to find a small number of associative operators in a given goal.  
  Thus, using `Array` is fine here.  
  -/  
  structs : Array Struct := {}  
  /--  
  Mapping from operators to its "operator id". We cache failures using `none`.  
  `opIdOf[op]` is `some id`, then `id < structs.size`. -/  
  opIdOf : PHashMap ExprPtr (Option Nat) := {}  
  /--  
  Mapping from expressions/terms to their structure ids.  
  Recall that term may be the argument of different operators. -/  
  exprToOpIds : PHashMap ExprPtr (List Nat) := {}  
  steps := 0  
  deriving Inhabited  
  
builtin_initialize acExt : SolverExtension State ← registerSolverExtension (return {})
```



grind: Extensibility - solvers

After you declare your solver extension. You implement your internalizer, propagators, and equality handlers.

```
def processNewDiseq (a b : Expr) : GoalM Unit := withExprs a b do
  let ea ← asACExpr a
  let lhs ← norm ea
  let eb ← asACExpr b
  let rhs ← norm eb
  { lhs, rhs, h := .core a b ea eb : DiseqCnstr }.assert
```

```
builtin_initialize
acExt.setMethods
  (internalize := AC.internalize)
  (newEq := AC.processNewEq)
  (newDiseq := AC.processNewDiseq)
  (check := AC.check)
  (checkInv := AC.checkInvariants)
```



grind: Tooling

How to maintain annotations in a huge libraries with more than 2M lines of code?

```
-- `BitVec.msb_signExtend` is reasonable at 22.  
#guard_msgs in  
#grind_lint inspect (min := 25) BitVec.msb_signExtend  
  
/ -! Check BitVec namespace: -/  
  
#guard_msgs in  
#grind_lint check (min := 20) in BitVec
```

Can AI/users control grind? Yes, grind interactive mode.

```
example {x : R} (f : R → Nat)
  : max 3 (4 * f ((cos x + sin x)^2)) ≠ 2 + f (2 * cos x * sin x + 1) := by
  grind =>
  use [Nat.max_def, trig_identity]
  ring
  cases_next
```

Available since v4.25.0.

Higher level of abstraction.

Foundation for grind?

Prediction: AI will learn how to use grind interactive mode in a few months.

grind?

```
theorem mem_insert (m : IndexMap α β) (a a' : α) (b : β) :
  ⚡ a' ∈ m.insert a b ↔ a' = a ∨ a' ∈ m := by
  grind?
```

Try these:

```
[apply] grind only [= mem_indices_of_mem, insert, =_
HashMap.contains_iff_mem, = getElem?_neg, = getElem?_pos,
= HashMap.contains_insert, #4ed2, #ffdf, #10d8, #2688]
[apply] grind only [= mem_indices_of_mem, insert, =_
HashMap.contains_iff_mem, = getElem?_neg, = getElem?_pos,
= HashMap.contains_insert]
[apply] grind ⇒
  instantiate only [= mem_indices_of_mem, insert]
  instantiate only [= _ HashMap.contains_iff_mem, = getElem?_neg, =
getElem?_pos]
  cases #4ed2
  · cases #ffdf
    · instantiate only
    · instantiate only
      instantiate only [= HashMap.contains_insert]
  · cases #10d8
    · cases #2688
      · instantiate only
      · instantiate only
        instantiate only [= HashMap.contains_insert]
    · cases #ffdf
      · instantiate only
      · instantiate only
        instantiate only [= HashMap.contains_insert]
```

grind?

```
theorem mem_insert (m : IndexMap  $\alpha$   $\beta$ ) (a a' :  $\alpha$ ) (b :  $\beta$ ) :  
  a'  $\in$  m.insert a b  $\leftrightarrow$  a' = a  $\vee$  a'  $\in$  m := by  
  grind only [= mem_indices_of_mem, insert, =_ HashMap.contains_iff_mem, = getElem?_neg,  
    = getElem?_pos, = HashMap.contains_insert]
```

grind?

```
theorem mem_insert (m : IndexMap  $\alpha$   $\beta$ ) (a a' :  $\alpha$ ) (b :  $\beta$ ) :
  a'  $\in$  m.insert a b  $\leftrightarrow$  a' = a  $\vee$  a'  $\in$  m := by
  grind =>
  instantiate only [= mem_indices_of_mem, insert]
  instantiate only [= _ HashMap.contains_iff_mem, = getElem?_neg, = getElem?_pos]
  cases #4ed2
  · cases #ffdf
    · instantiate only
    · instantiate only
      instantiate only [= HashMap.contains_insert]
  · cases #10d8
    · cases #2688
      · instantiate only
      · instantiate only
        instantiate only [= HashMap.contains_insert]
    · cases #ffdf
      · instantiate only
      · instantiate only
        instantiate only [= HashMap.contains_insert]
```

+suggestions: grind's best friend

```

/-- A product set is included in a product set if and only factors are included, or a factor
first set is empty. -/
theorem prod_subset_prod_iff : s ×s t ⊆ s1 ×s t1 ↔ s ⊆ s1 ∧ t ⊆ t1 ∨ s = ∅ ∨ t = ∅ := by
  grind +suggestions

```



```

/-- A product set is included in a product set if and only factors are included, or a factor
first set is empty. -/
theorem prod_subset_prod_iff : s ×s t ⊆ s1 ×s t1 ↔ s ⊆ s1 ∧ t ⊆ t1 ∨ s = ∅ ∨ t = ∅ := by
  grind? +suggestions

```



grind finds out for you exactly with theorems are relevant

```

/-- A product set is included in a product set if and only factors are included, or a factor
first set is empty. -/
theorem prod_subset_prod_iff : s ×s t ⊆ s1 ×s t1 ↔ s ⊆ s1 ∧ t ⊆ t1 ∨ s = ∅ ∨ t = ∅ := by
  grind only [prod_eq_empty_iff, prod_subset_iff, prod_mono, = subset_def, mem_empty_iff_fa
  mem_prod, empty_subset]

```



"if-normalization" challenge by Leino, Merz, and Shankar

```
def normalize (assign : Std.HashMap Nat Bool) : IfExpr → IfExpr
| lit b => lit b
| var v =>
  match assign[v]? with
  | none => var v
  | some b => lit b
| ite (lit true) t _ => normalize assign t
| ite (lit false) _ e => normalize assign e
| ite (ite a b c) t e => normalize assign (ite a (ite b t e) (ite c t e))
| ite (var v) t e =>
  match assign[v]? with
  | none =>
    let t' := normalize (assign.insert v true) t
    let e' := normalize (assign.insert v false) e
    if t' = e' then t' else ite (var v) t' e'
  | some b => normalize assign (ite (lit b) t e)
termination_by e => e.normSize

-- We tell `grind` to unfold our definitions above.
attribute [local grind] normalized hasNestedIf hasConstantIf hasRedundantIf disjoint vars eval List.disjoint

theorem normalize_spec (assign : Std.HashMap Nat Bool) (e : IfExpr) :
  (normalize assign e).normalized
  ∧ (∀ f, (normalize assign e).eval f = e.eval fun w => assign[w]?.getD (f w))
  ∧ ∀ (v : Nat), v ∈ vars (normalize assign e) → ¬ v ∈ assign := by
  fun_induction normalize with grind
```



"if-normalization" challenge by Leino, Merz, and Shankar

Interactive tactic suggestion tool: the `try?` tactic

It tries many different tactics, guesses induction principle, and is **extensible**

```
✓ theorem normalize_spec (assign : Std.HashMap Nat Bool) (e : IfExpr) :  
  (normalize assign e).normalized  
  ∧ (∀ f, (normalize assign e).eval f = e.eval fun w => assign[w]?.getD (f w))  
  ⚙️ ∧ ∀ (v : Nat), v ∈ vars (normalize assign e) → ¬ v ∈ assign := by  
  try?
```

Lean Infoview ×

▼ Suggestions

Try these:

- `fun_induction normalize <=> grind`
- `fun_induction normalize <=> grind only [vars, normalized, disjoint, =_ Std.HashMap.contains_iff_mem, =_ List.contains_iff_mem, List.contains_eq_mem, hasNestedIf, hasConstantIf, hasRedundantIf, List.elem_nil, eval, cases Or, List.contains_cons, List.eq_or_mem_of_mem_cons, Option.getD_none, List.mem_cons_of_mem, getElem?_pos, getElem?_neg, Option.getD_some, = Std.HashMap.mem_insert, = Std.HashMap.getElem?_insert, = Std.HashMap.getElem_insert, = Std.HashMap.contains_insert, =_ List.cons_append, = List.append_assoc, = List.contains_append, List.nil_append, List.disjoint, List.append_nil, = List.cons_append, =_ List.append_assoc, → List.eq_nil_of_append_eq_nil, List.mem_append]`



grind: Initial reactions



Markus de Medeiros Jul 22nd at 1:43 PM

I keep being surprised by how many nuisance goals `grind` is able to solve. Props to everyone who worked on it!



+1 1



You and Kim Morrison, Oliver Nash

AUG 8



Oliver Nash EDITED

8:27 AM

I was just singing `grind`'s praises in the Mathlib community meeting and highlighted my favourite example was [#27372](#) (which massively golfs some of my work).

After the call somebody suggested I highlight it to you both for your enjoyment :)



You

SEP 7



Fabrizio Montesi

11:28 AM

Testing a bundled definition of `Bisimulation`, and holy cow does `grind` shine. With the right annotations, it managed to prove that bisimilarity is a bisimulation.

```
def Bisimilarity (lts : Lts State Label) : Bisimulation lts :=
  rel s1 s2 := ∃ r : Bisimulation lts, r s1 s2
  is_bisimulation := by grind
```



Chris Henson, Shreyas Srinivas, Kim Morrison

```
- refine (λ _ _ _ _ ha, haj, hb, hbj, hc, hcj, hd, hdj, ?_, ?_, ?_, ?_, ?_)
- <|> rw [mem_insert] at * <|> try rintro rfl
- · obtain (rfl | ha) := ha
- · obtain (rfl | hb) := hb
- · exact hw.isPathGraph3Compl.fst_ne_snd rfl
- · exact hw.fst_notMem_right hb
- · obtain (rfl | hb) := hb
- · exact hw.snd_notMem_left ha
- · exact haj <| hw <| mem_inter_of_mem ha hb
- · obtain (rfl | ha) := ha
- · obtain (rfl | hd) := hd
- · exact hw.isPathGraph3Compl.ne_fst rfl
- · exact hw.fst_notMem_right hd
- · obtain (rfl | hd) := hd
- · exact hw.notMem_left ha
- · exact haj <| hw <| mem_inter_of_mem ha hd
- · obtain (rfl | hb) := hb
- · obtain (rfl | hc) := hc
- · exact hw.isPathGraph3Compl.ne_snd rfl
- · exact hw.snd_notMem_left hc
- · obtain (rfl | hc) := hc
- · exact hw.notMem_right hb
- · exact hbj <| hw <| mem_inter_of_mem hc hb
- · intro hat
- obtain (rfl | ha) := ha
- · exact hw.fst_notMem_right hat
- · exact haj <| hw <| mem_inter_of_mem ha hat
- · intro hbs
- obtain (rfl | hb) := hb
- · exact hw.snd_notMem_left hbs
- · exact hbj <| hw <| mem_inter_of_mem hbs hb
+ exact (λ _ _ _ _ ha, haj, hb, hbj, hc, hcj, hd, hdj, by grind)
```

Comment on line R312



Yael Dillies 19 minutes ago

Collaborator ...

Wow! 🤩



grind: Initial reactions



Terence Tao

@tao@mathstodon.xyz

In contrast, AI chatbots are usually tuned to avoid a "failure mode" as much as possible, at the expense of increasing the occurrence of "intermediate modes" where the chatbot response looks potentially useful, and invites further interaction from the user, but is not exactly providing what the user wants, and could contain hallucinations or some fundamental misunderstanding of the task that would take significant effort to uncover. Paradoxically, such tools may become significantly more useful if they simply reported that they were unable to provide a high quality answer to a query in such cases.

A comparison may be drawn with the increasingly advanced, but stringently verified, "tactics" used in a modern proof assistant such as Lean. I have been experimenting recently with the new tactic ``grind`` in Lean, which is a powerful tool (inspired more by "good old-fashioned AI" such as satisfiability modulo theories (SMT) solvers, than modern data-driven AI) to try to close complex proof goals if all the tools needed to do so are already provided in the proof environment; roughly speaking, this corresponds to proofs that can be obtained by "expanding everything out and trying all obvious combinations of the hypotheses". When I apply ``grind`` to a given subgoal, it can report a success within seconds, closing that subgoal in a Lean-verified fashion and allowing me to move on to the next subgoal. But, importantly, when this does not work, I quickly get a "``grind` failed" message, in which case I simply delete `grind` from the code and proceed by a more pedestrian sequence of lower level tactics. (2/3)`

Challenge

AI systems are writing impressive Lean proofs, but **they often bypass powerful automation** tactics, **generating long manual proofs that a single tactic call could replace.**

- **The bottleneck is annotations:** metadata that tells **symbolic automation** how to use theorems.
- Most mathematicians can write proofs but can't write good annotations.
- **We propose training AI to predict annotations.**

simp, aesop, and grind become much more powerful after theorems are properly annotated.

If we had perfect annotations:

Proof synthesis becomes dramatically easier (predict grind instead of 20 tactic steps)

AI-generated proofs become shorter, more robust, more maintainable

RL signals: success rate when reproofing existing libraries, grind diagnostics, #grind_lint

Symbolic Simulation Framework (SymM)

SymM is a new framework for implementing verification condition generators (mvcgen, aeneas, velvet, strata) in Lean.

Started after the Lean@Google hackathon.

The insight: Don't optimize expensive operations: avoid them entirely. Pattern matching in simplification was dominated by definitional equality checks; we redesigned to use fast syntactic matching instead.

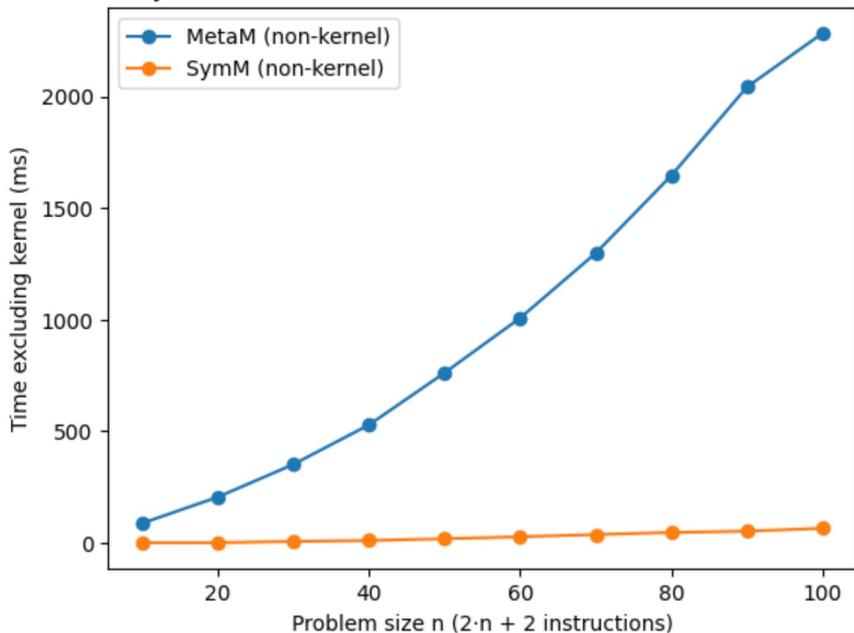
Maximally shared terms: Store terms so identical subterms share the same memory. Equality becomes pointer comparison, and caching becomes trivial.

Two-phase matching: First pass is pure syntactic matching (fast). Only fall back to unification when you hit binders or metavariables, which is rare in practice.

Reuse grind state: Most verification conditions share the same prefix. Avoid duplicate work.

Symbolic Simulation Framework (SymM)

Symbolic simulation benchmark (kernel time excluded)



Symbolic simulation benchmark — tactic time only

Problem size n corresponds to a program with $4 \cdot n$ monadic actions.

n	MetaM tactic (ms)	SymM tactic (ms)	Speedup
10	82.10	11.37	~7.2x
20	176.21	17.71	~9.9x
30	306.47	25.39	~12.1x
40	509.52	34.53	~14.7x
50	689.19	43.51	~15.8x
60	905.86	52.47	~17.3x
70	1172.31	62.50	~18.8x
80	1448.48	70.65	~20.5x
90	1787.15	80.89	~22.1x
100	2128.12	90.77	~23.5x

To follow the progress: <https://github.com/leanprover/lean4/issues?q=is%3Apr+author%3Aleodemoura>

Conclusion

Lean is an **efficient programming language** and **proof assistant** – no compromise

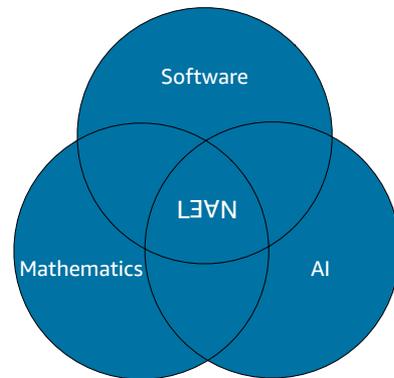
Lean is very **extensible** and is implemented in Lean: the ultimate extensibility test.

Proofs are machine-checkable, maintainable, and auditable.

The Lean FRO is shipping: module system, new compiler, grind, more to come.

Mathlib is changing how math is done.

Our goal: make verified software the norm, not the exception.



Thank You

<https://leanprover.zulipchat.com/>

x: @leanprover

LinkedIn: Lean FRO

Mastodon: @leanprover@functional.cafe

#leanlang, #leanprover

<https://www.lean-lang.org/>

