Day - 5 : Overview of Common Tactics

Touseef Haider

May 16, 2025

Study Mathematics In Lean(MIL) Section 2.3

5.1. The intro Tactic: Introducing Hypotheses

The intro tactic (or intros for multiple introductions) is your go-to tool when your goal is an *implication* (like $P \rightarrow Q$) or a *universally quantified statement* (like $\forall x : T, P x$).

For Implication $P \rightarrow Q$

If your goal is $\exists P \rightarrow Q$ (read as "prove that P implies Q"), it means you need to show that *if* you assume P is true, then you can prove Q. The intro hP tactic does exactly this:

- 1. It takes the P from the left side of the ' \rightarrow '.
- 2. It adds hP: P to your list of hypotheses (your local context). So now you have P as an assumption named hP.
- 3. It changes your goal to $\dashv Q$. Now you need to prove Q, but you have the extra help of assuming P (via hP).

Example:

Suppose your Lean state is:

```
\begin{array}{cccc} 1 & 1 & \text{goal} \\ 2 & P & Q & : & Prop \\ 3 & \neg & P & \rightarrow & Q \end{array}
```

If you apply the tactic intro hP, the state becomes:

```
1 1 goal
2 P Q : Prop
3 hP : P
4 ⊣ Q
```

You've "introduced" the hypothesis P into your context.

For Universal Quantifiers ($\forall x : T, P x$):

If your goal is $\dashv \forall x : T$, $P \ge (read as "prove that for all x of type T, the property <math>P \ge holds"$), it means you need to show this for an *arbitrary* x of type T. The intro x tactic does this:

- 1. It introduces an arbitrary x of type T into your local context.
- 2. It changes your goal to \dashv P x. Now you need to prove P x for this arbitrary x.

Example:

Suppose your Lean state is:

```
\begin{array}{cccc} 1 & \text{goal} \\ 2 & P & : \text{Nat} \rightarrow \text{Prop} \\ 3 & \dashv & \forall n : \text{Nat, } P n \end{array}
```

If you apply the tactic intro n_val, the state becomes:

```
\begin{array}{cccc} 1 & \text{goal} \\ 2 & P : \text{Nat} \rightarrow \text{Prop} \\ 3 & \text{n_val} : \text{Nat} \\ 4 & \neg P \text{n_val} \end{array}
```

You've introduced an arbitrary natural number n_val and now must prove P holds for it.

Why is intro useful?

It allows you to peel off the "if..." parts of implications or the "for all..." parts of universal statements, turning them into assumptions you can use in later steps of your proof (often with **rw** or other tactics).

Think of it like this: If someone asks you to prove "If it's raining, then the ground is wet," your natural first step is to say, "Okay, *assume* it's raining..." That's what **intro** does!

Let's see a quick example in Lean code. Consider this simple (and perhaps obvious) theorem: If P is true, then P is true.

```
1 theorem if_P_then_P (P : Prop) : P \rightarrow P := by

2 -- What tactic would you use first here, based on what we just discussed?

3 -- And what would the goal become?
```

What do you think? Your initial goal is :

 $\begin{array}{cccc} 1 & \texttt{goal} \\ 2 & \texttt{P} & : & \texttt{Prop} \\ 3 & \dashv & \texttt{P} & \rightarrow & \texttt{P} \end{array}$

If you apply the tactic intro h (where h is the name you give to the new hypothesis, the Infoview will change to :

```
1 1 goal
2 P : Prop
3 h : P
4 ⊢ P
```

Now you need to prove P. Look at the current state:

- You need to prove P.
- You have a hypothesis h which is P.

How would you tell Lean, "Hey, the thing I need to prove is exactly this hypothesis h that I already have"?

5.2. The exact Tactic: Pointing to the Proof

The exact tactic is used when your current goal is exactly matched by one of your hypotheses or by a known lemma/theorem that directly states the goal.

- If you goal is \dashv G
- And you have a hypothesis h : G (meaning h is a proof of G)
- Or you know a lemma my_lemma : G
- Then exact h (or exact my_lemma) will solve the goal.

It's like telling Lean, "The proof you are looking for? It is this thing right here!". So for our if_P_then_P theorem, the complete proof is :

```
1 theorem if_P_then_P ( P : Prop ) : P \rightarrow P := by

2 intro hP -- Goal becomes \dashv P, with hP : P in hypotheses

3 exact hP -- Solves the goal using the hypothesis hP
```

You have now learned :

- intro : to assume the premise of an implication (or introduce an arbitrary variable for a \forall
- exact : To solve a goal when you have a hypothesis or lemma that is precisely what the goal states.

These two work together: intro to get hypotheses, and then other tactics (including exact or rw) to use those hypotheses to solve the new goal.

5.3. The apply Tactic : Matching Conclusions and Creating Subgoals

The apply tactic is similar to exact, but it is more flexible. You use apply when you have a hypothesis or lemma whose *conclusion* matches your current goal, but the hypothesis/lemma itself might have some prerequisites (premises).

How it works:

- You have a goal, say \dashv G.
- You have a hypothesis or lemma, say $h : P_1 \to P_2 \to \cdots \to P_n \to G$. (This mean if you can prove P_1, P_2, \cdots, P_n , then you get G).
- If you use tactic apply h, Lean sees that the conclusion of h (which is G) matches your goal G.
- Lean then says, "Okay, I can use h to prove G, provided you can give me proofs for all of h's premises."
- Your original goal \dashv G is replaced by n new subgoals : $\dashv P_1, \dashv P_2, \cdots, \dashv P_n$.

Contrast with exact:

- exact foo works if foo is directly a proof of your goal.
- apply foo works if foo is a function/implication that results in your goal. apply essentially does "modus ponens" in reverse : to prove G using $P \rightarrow G$, it asks you to prove P.

Example

Let's say you have the following state:

Your goal is R.

- You cannot exact hQR because hQR is $Q \rightarrow R$ not R.
- You cannot exact hPQ because hPQ is $P \rightarrow Q$, not R.

But notice that hQR concludes with R (i.e. $Q \rightarrow R$). So, you can apply hQR. If you type apply hQR, the state becomes :

1 goal 2 P Q R : Prop 3 hPQ : P \rightarrow Q 4 hQR : Q \rightarrow R 5 hP : P 6 \dashv Q -- New goal! This was the premise of hQR

Now your goal is Q.

• You can apply hPQ because hPQ concludes with Q (i.e. $P \rightarrow Q$). If you type apply hPQ, the state becomes:

1 1 goal 2 P Q R : Prop 3 hPQ : P \rightarrow Q 4 hQR : Q \rightarrow R 5 hP : P 6 \dashv P -- New goal! This was the premise of hPQ

Now your goal is P. And you have hP : P. So, how would you finish this proof? This chaining (apply . . . then apply . . . then often exact) is very common for proving implications.

Let's try to write out the full proof for the state above. Suppose the theorem was stated like this:

```
1 example (P Q R : Prop) (hPQ : P \rightarrow Q) (hQR : Q \rightarrow R) (hP : P) : R := by
2 -- How would you prove this step-by-step using apply and then the final tactic?
```

```
1 example (P Q R : Prop) (hPQ : P \rightarrow Q) (hQR : Q \rightarrow R) (hP : P) : R := by
     -- Initial state:
2
     -- P Q R : Prop
3
     -- hPQ : P \rightarrow Q
4
     -- hQR : Q \rightarrow R
     -- hP : P
6
     -- ⊣ R
7
8
     apply hQR -- The conclusion of hQR (Q 
ightarrow R) is R, matching the goal.
9
                  -- New goal becomes the premise of hQR.
10
11
     -- New state:
12
     -- P Q R : Prop
13
     -- hPQ : P \rightarrow Q
     -- hQR : Q \rightarrow R
14
     -- hP : P
     -- ⊣ Q
16
17
     apply hPQ -- The conclusion of hPQ (P 
ightarrow Q) is Q, matching the current goal.
18
                  -- New goal becomes the premise of hPQ.
19
     -- New state:
20
     -- P Q R : Prop
21
     -- hPQ : P \rightarrow Q
22
     -- hQR : Q \rightarrow R
23
     -- hP : P
24
     -- ⊣ P
25
26
     exact hP -- The hypothesis hP directly proves the current goal P.
28
     -- Proof complete! Infoview shows ''goals accomplished"
```

You have beautifully demonstrated how apply works backward from the goal, using implications, and how exact can then seal the deal when a premise is directly available as a hypothesis. This pattern of applying implications and then proving the new sub goals (perhaps with more apply, rw, exact, or intros) is a cornerstone of proving in Lean.

5.4. The simp Tactic : The Simplifier

The simp tactic is one of Lean's most powerful and frequently used automation tools. Think of it as a skilled assistant that tries to simplify your goal (or hypotheses) using a predefined set of rewrite rules, definitions, and basic computational steps.

What it does:

- simp takes your current goal.
- It repeatedly applies lemmas that are tagged with the @[simp] attribute (these are called "simp lemmas"). These lemmas are equalities that are generally considered "simplifying" like n + 0 = n, 0 + n = n, ¬(¬P) ↔ P, etc.
- It unfolds definitions that are marked as @[reducible] or sometimes just basic definitions.
- It performs computations (like 2 + 2 becomes 4).
- It continues this process until no more such simplifications can be applied.

How to use it:

- simp: Tries to simplify the current goal.
- simp at h : Tries to simplify hypothesis h.
- simp at * : Tries to simplify all hypotheses and the goal.
- simp [lemma1, lemma2] : Simplifies using the default simp set plus lemma1 and lemma2 (even if they are not @[simp] lemmas).
- simp only [lemma1, lemma2] : Simplifies using only lemma1 and lemma2 (and not the default simp set). This gives you more control.
- simp [-lemma_to_remove] : Simplifies using the default simp set but *excludes* lemma_to_remove.

When is it useful?

- Cleaning up expressions: If your goal has things like x+0 or true && P. simp can often clean these up to x or P.
- Unfolding definitions: if a definition is obscuring the structure of your goal, simp might unfold it to reveal something easier to work with.
- Applying many basic rewrites : Instead of rw [lemma1], rw [lemma2], rw [lemma3], sometimes a single simp can do the job if those lemmas are part of its simp set.
- After intro or cases (another tactic we haven't covered yet, for breaking down inductive types), the goal might become messy, and simp can tidy it up.

Example

Suppose your goal is \dashv (x+ 0) + (y * 1) = x + y. Lemmas like Nat.add_zero (n : Nat) : n + 0 = n and Nat.mul_one (n: Nat) : n * 1 = n are typically @[simp] lemmas.

If you apply simp, Lean would:

- See x + 0 and rewrite it to x (using Nat.add_zero). Goal becomes x + (y * 1) = x + y.
- See y + 1 and rewrite it to y (using Nat.mul_one). Goal becomes x + y = x + y
- This is now rfl, and simp is often smart enough to solve such goals by reflexivity it it reaches that state. So, it might directly show "goal accomplished".

A Word of Caution (and Power!):

simp is powerful, but it can sometimes feel like a "black box". It might do a lot of steps, and if it does not solve the goal, the new goal might look quite different, and it might not be obvious how it got there. Let's try a small example theorem where simp can be very helpful :

```
import Mathlib.Data.Nat.Basic -- For Nat properties
theorem simp_example (P : Prop) (n m : Nat) (h : P) : (true P) ∧ (n + 0 = m) → m = n := by
-- What do you think simp will do to the goal if we apply it after an intro?
```

If you were to first use intro h_imp for the main implication, your goal would become m = n and you would have a hypothesis h_imp : (true $\land P$) \land (n + 0 = m). What do you think simp at h_imp would do to this hypothesis h_imp? Think about true $\land P$ and n + 0.

For the hypothesis h_{imp} : (true $\land P$) \land (n + 0 = m):

- simp would look at true \land P. Since true \land P is equivalent to P (often via simp lemma like true_and P : true \land P \iff P), it simplifies this part to P.
- simp would look at n + 0 = m. Since n + 0 simplifies to n (often via a simp lemma like Nat.add_zero n : n + 0 = n), this part of the equality becomes n=m.

```
So, after simp at h_imp, the hypothesis h_imp would effectively become P \land (n = m). Let's see how we could complete the proof of the simp_example theorem:
```

```
1 import Mathlib.Data.Nat.Basic -- For Nat properties
2
  theorem simp_example (P : Prop) (n m : Nat) (_h_given_P : P) : (true \land P) \land (n + 0 = m) \rightarrow m
3
       = n := by
                           -- h_imp : (true \land P) \land (n + 0 = m)
    intro h_imp
                           -- Goal: \dashv m = n
6
    -- Simplify the hypothesis h_imp
\overline{7}
    simp at h_imp --
8
                           -- After this, h_imp effectively means P is true AND n = m.
9
                           -- In Lean 4, h_imp will likely become: h_imp : P \land n = m
11
    -- Now h_imp is 'P \land (n = m)'. We need to use the 'n = m' part.
    -- We can extract the parts of the conjunction.
13
    -- A common way is to use the 'cases' tactic or dot notation if available.
14
    -- Let's assume h_imp is now 'P \wedge n = m'. We can get 'n = m' from it.
    -- If h_imp is h_imp : P \land n = m
16
    -- We can get 'h_eq : n = m' from 'h_imp.right' (or 'h_imp.2' in some versions)
17
    -- or by using 'cases h_imp with hP_internal h_eq_actual'
18
    cases h_imp with hp_from_h_imp h_eq_from_h_imp
19
                           -- Now we have:
20
                           -- hp_from_h_imp : P
21
                           -- h_eq_from_h_imp : n = m
22
23
    -- Our goal is \dashv m = n. We have h_eq_from_h_imp : n = m.
24
25
    -- We need to show m = n. This is the *symmetric* of n = m.
    -- We can rewrite or use 'Eq.symm'.
26
    rw [h_eq_from_h_imp] -- This rewrites n to m in the goal, making it m = m
27
                           -- Or, if it rewrites m to n, it becomes n = n.
28
                           -- Let's assume it rewrites n in the goal if possible, or m.
29
                           -- If goal is m = n and h_eq_from_h_imp is n = m.
30
                           -- 'rw [h_eq_from_h_imp]' might try to replace n with m in the goal.
31
                           -- This would make goal 'm = m'.
32
                           -- Or, 'rw [ \leftarrow h_eq_from_h_imp]' would replace m with n, making it 'n
33
       = n'.
34
    -- A more direct way if you have 'h_eq_from_h_imp : n = m' to prove 'm = n' is:
    exact Eq.symm h_eq_from_h_imp -- 'Eq.symm' turns 'n=m' into 'm=n'
35
36
    -- So, a cleaner finish after 'simp at h_imp' which results in 'h_imp : P \wedge n = m':
37
    -- exact (h_imp.right).symm -- Assuming .right gives n=m
38
39
```

```
40 -- Let's try the cases approach which is common:

41 -- (after 'simp at h_imp' resulted in 'h_imp : P \land n = m')

42 -- cases h_imp with unused_P_part actual_equality -- 'actual_equality' is 'n = m'

43 -- exact actual_equality.symm -- This uses the .symm property of an equality.
```

A very common and often more robust way in Lean 4 after simp at h_{imp} (if h_{imp} becomes $P \land n = m$) is:

```
1 theorem simp_example_lean4 (P : Prop) (n m : Nat) (_hP_given : P) : (true \land P) \land (n + 0 = m
      ) \rightarrow m = n := bv
    intro h_imp
    simp (config := { contextual := true }) [Nat.add_zero, and_true] at h_imp -- contextual
      simp might be more powerful
4
    -- Assuming h_imp is now P \land n = m
    -- We can use dot notation to access parts of a conjunction or an equality:
5
    rw [h_imp.right] -- If h_imp.right is n = m, this will rewrite the goal m = n to m = m if
6
     n is found, or n=n if m is found.
                       -- Or more directly, if h_imp.right is n=m
    rfl
7
    -- exact h_imp.right.symm -- This should work
8
9
    -- Let's assume 'simp at h_imp' simplifies it effectively to give us 'n = m'.
10
    -- One way 'simp' can be very powerful is using 'simp [*] at *'.
    -- This uses all hypotheses to simplify each other and the goal.
12
13
    -- Consider this alternative approach using that:
```

Let's show a more direct proof using **simp**'s power:

```
import Mathlib.Data.Nat.Basic

theorem simp_example_direct (P : Prop) (n m : Nat) (_h_given_P : P) : (true ∧ P) ∧ (n + 0 =
m) → m = n := by
intro h_imp -- h_imp : (true ∧ P) ∧ (n + 0 = m)
-- Goal: ⊣ m = n
simp at h_imp -- h_imp becomes P ∧ n = m
exact h_imp.right.symm -- h_imp.right is n = m, so h_imp.right.symm is m = n
```

This version is clean and shows simp simplifying the hypothesis, and then we use the parts of the simplified hypothesis. The .right gets the second part of the \land (which is n = m), and .symm gives the symmetric equality m=n, which exactly matches the goal.

The key takeaway for simp is that it is a powerful tool for automatically applying many known simplifying rewrite rules.

Now you have rfl, rw, intro, exact, apply, and simp. That is a very solid set of basic tactics.

5.5. When to Use Which Tactic : A Quick Summary

rfl

- When : Your goal is an equality A=B (or an equivalence $P \iff Q$) where A and B (or P and Q) are *definitionally* equal. Lean can see they are the same just by unfolding definitions.
- Example : $\dashv x + 0 = x$ (if add_zero is definitional), $\dashv 2 + 2 = 4$.

rw [lemma_or_hyp]

- When : Your goal contains a subexpression that matches one side of an equality stated in lemma_or_hyp, and you want to replace it with the other side.
- Use rw [< lemma_or_hyp] to rewrite from right-to-left.
- Example: Goal \dashv (a + b) +c = (b + a) +c, use rw [Nat.add_comm a b] to change a+b to b+a.

intro h (or intro h1 h2 ...):

- When : Your goal is an implication $P \rightarrow Q$ or a universal qualification $\forall x : T, P x$.
- Effect: Adds P (as h : P) or x : T to your hypotheses and changes the goal to Q or P x.
- Example: Goal \dashv P \rightarrow (Q \rightarrow R), intros hP hQ gives hypotheses hP : P, hQ : Q and goal \dashv R.

exact h (or exact lemma_name):

- When : Your current goal \dashv G is *exactly* what hypothesis h states (i.e. , h : G) or what lemma_name states.
- Example : Goal ⊢ P, hypothesis hP ; P. Use exact hP.

apply h(or apply lemma_name):

- When : The conclusion of hypothesis h (or lemma_name) matches your goal, but h itself might have premises.
- Effect : Solves the current goal using h and creates new subgoals for each premise of h.
- Example : Goal \dashv R, hypothesis hQR : Q \rightarrow R. Use apply hQR. New goal becomes \dashv Q.

simp

- When : You want to simplify the goal or hypotheses using a collection of known rewrite rules (like x + 0 = x, true $\land P \iff P$, etc.) and definitions.
- Can be used as simp, simp at h, simp [*] at * (uses hypotheses to simplify too).
- Example : goal \dashv (x + 0) + (y * 1) = z, simp might change it to \dashv x + y =z.

This summary should give a good starting point for choosing your tactics. Often, a proof involves a sequence of these.