Teaching mathematics using Lean and controlled natural language

3 Patrick Massot 🖂 D

4 Université Paris-Saclay, France

5 Carnegie Mellon University, USA

6 — Abstract -

We present Verbose Lean, a library using the flexibility of the Lean programming language and proof assistant to teach undergrad mathematics students how to read and write traditional paper proofs. After explaining our main pedagogical goals, we explain how students use the library with varying levels of assistance to write proofs that are easy to transfer to paper because they look like natural language. We then describe how teachers can customize the student experience based on their specific pedagogical goals and constraints. Finally we describe some aspects of the implementation of the library, emphasizing how new aspects of the very recently released version 4 of Lean allow a lot of flexibility that could benefit many new creative uses of a proof assistant.

¹⁵ **2012 ACM Subject Classification** Human-centered computing \rightarrow Natural language interfaces; ¹⁶ Applied computing \rightarrow Interactive learning environments; Theory of computation \rightarrow Logic and ¹⁷ verification

18 Keywords and phrases mathematics teaching, proof assistant, controlled natural language

¹⁹ Digital Object Identifier 10.4230/LIPIcs.CVIT.2016.23

²⁰ Funding Patrick Massot: Partially funded by the Hoskinson center for formalized mathematics

Acknowledgements The library described here is the result of work spread over five years and 21 benefited from many interactions. Marie-Anne Poursat made it possible by trusting me to teach 22 using a proof assistant. Frédéric Bourgeois and Christine Paulin-Mohring later joined me in teaching 23 this course and discussing how to make it more useful. Our students also participated with their 24 patience, questions and enthusiasm. Very recently, Julien Narboux, Pierre Boutry and Evmorfia-Iro 25 Bartzia started using this library with students and made valuable comments. Simon Hudon started 26 my Lean meta-programming education. It was then continued by Mario Carneiro and Kyle Miller 27 who both also directly contributed crucial code that I would have been incapable of writing. Wojciech 28 Nawrocki helped me a lot with widgets, adding several features to his ProofWidgets library for 29 the needs of this project. More generarly many people from the Mathlib community contributed 30 indirectly by formalizing basic mathematics and creating tactics that we built on. I also benefited 31 from conversations about teaching using a proof assistant, particularly with Heather Macbeth, Jim 32 Portegies and Jelle Wemmenhove. Most of the work on the Lean 4 version was made possible 33 by Jeremy Avigad's invitation to benefit from some of Charles Hosinkson's generous donation to 34 CMU. And of course none of this would exist without the amazing work of Leonardo de Moura and 35 Sebastian Ullrich on Lean 4. 36

37 **1** Introduction

The transition from high-school mathematics to proof-based university mathematics is a well-known challenge for students. In the recent past, there have been several experiments using proof assistants to help students in this transition [2, 7, 12, 1, 17]. Here we really mean courses that consider the proof assistant only as an intermediate tool, not as a final goal. Note this tool is applicable to any kind of mathematics in principle but this account and the library it describes are biased towards elementary real analysis which is used in France as

© Patrick Massot; licensed under Creative Commons License CC-BY 4.0 42nd Conference on Very Important Topics (CVIT 2016). Editors: John Q. Open and Joan R. Access; Article No. 23; pp. 23:1–23:17 Leibniz International Proceedings in Informatics Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany ⁴⁴ the main introduction to rigorous proofs.

Verbose Lean¹ is a teaching library and meta library for the Lean programming language 45 and proof assistant. Lean is due mostly to Leonardo de Moura at Microsoft Research and 46 then Amazon Web Service and the Lean FRO^2 . This library has three main layers. The 47 bottom one is a set of tactics (i.e. proof producing programs) mimicking the granularity of 48 proofs on paper. The middle one is a controlled natural language syntax whose goal is to 49 ease the transition to paper proofs, at the cost of being slightly more difficult to write. The 50 third layer is made of mechanisms that help students to write proofs by suggesting potential 51 next steps. All three layers are heavily customizable, even without programming knowledge, 52 and everything exists in French and English and is translatable to other languages. 53

This paper is intended for mathematics teachers and for people who want to see examples of using Lean's flexibility. It is organized roughly in order of decreasing importance. We first explain our pedagogical goals, then describe the student experience, then the teacher experience before concluding with some aspects of the implementation. That last section is more technical but could be useful to anyone interested in what can be done with Lean.

⁵⁹ 2 Main pedagogical goals

The first main goal is to train students to have a crystal clear view of the current state of the proof: what is currently being proven, what are the current assumptions, which mathematical objects are fixed. The next pedagogical goal is to train students to automatically perform proof steps depending of the syntactic structure of the goal. For instance a direct proof of a universally quantified statement starts with fixing an object of the relevant type.

A basic requirement here is to make sure that every statement has a clear status: is 65 it something that is known to be true? Something that we assume? Something that will 66 be proven soon? We claim that this goal is much easier to achieve if there is a really clear 67 separation between stating, proving and using. For any given logical operator or quantifier, 68 there are syntactic rules that explain how to form a mathematical statement it and some 69 previously existing statements. Then there are rules that explains how to prove such a 70 statement. Finally there are rules about how to use such a statement. This distinction 71 seems obvious but in practice it is incredibly blurred, both by students and by professional 72 mathematicians. Of course the difference is that mathematicians know what they are doing 73 even when they are very sloppy about this distinction. One of the most frequent cases is 74 failing to distinguish between stating the existence of a mathematical object satisfying some 75 condition and fixing a witness. An example would be to say or write "since f is continuous 76 and ε is positive, there exists δ such that..." and then refer to some δ on the following 77 line. The other very common one is more tied to using symbols instead of text. It uses 78 the implication symbol as an abbreviation of the word "hence" or "therefore", and more 79 generally mixing using an implication and stating one. An example would be writing on 80 a blackboard "f is polynomial $\implies f$ is continuous" instead of "f is polynomial hence 81 continuous". This misuse of a symbol is extremely problematic since this alternative meaning 82 of the symbol turns every legitimate use into nonsense. For instance reading the definition of 83 convergence of a function f towards a number l at some point x_0 with this interpretation for 84 the implication symbol gives "for every positive ε , there exists some positive δ such that for 85 every x, $|x - x_0| < \delta$ hence $|f(x) - l| < \varepsilon$ ". A less ubiquitous but still common example is a 86

¹ https://github.com/PatrickMassot/verbose-lean4

² https://lean-fro.org/

confusion between stating a claim starting with a universal quantifier and beginning a proof
 of such a statement.

All those logical errors (or at least sloppy writing) also impact achieving a third pedagogical goal which is to train students to make a clear distinction between bound and free variables. This is of course very related to keeping track of what is fixed in the current state of the proof, but with the additional twist that some variable name can appear simultaneously as the name of a free variable and as a bound variable in some assumption or in the current target statement.

The next teaching goal that we want to achieve with this technology is to train students to classify proof steps into safe reversible steps that require no initiative and risky irreversible steps that require some creativity. The alternation between the routine safe steps and the risky creative steps is a crucial aspect of mathematical proof creation which is not necessarily emphasized when presenting a proof on a blackboard.

On top of that there is a final goal which is to learn how to choose indirect strategies instead of simply following the syntax of the current target. Those includes stating and proving an intermediate fact, using a lemma instead of reproving everything, and the use of the excluded middle axiom, e.g. through proofs by contradiction or contraposition.

The usual interfaces of proof assistants have many qualities that help achieving those goals. We will use the Lean terminology, but that part of the discussion applies equally well to Coq for instance. Near the beginning of proof of sequential continuity from continuity, the proof assistant can display something like:

```
108
         f: \mathbb{R} \to \mathbb{R}
109
         u: \mathbb{N} \to \mathbb{R}
110
         \mathbf{x}_0: \mathbb{R}
111
         hu: u converges to x<sub>0</sub>
112
         hf: f is continuous at x<sub>0</sub>
113
         \varepsilon: \mathbb{R}
114
         \varepsilon_{pos: \varepsilon} > 0
115
         δ: R
116
         \delta pos: \delta > 0
117
         h\delta: \forall (x : \mathbb{R}), |x - x_0| \leq \delta \Rightarrow |f x - f x_0| \leq \varepsilon
118
         dash \exists N, \forall n \geq N, |(\texttt{f} \circ \texttt{u}) n - \texttt{f} x_0| \leq arepsilon
\frac{112}{128}
```

¹²¹ This display is called the *tactic state*. Most lines described mathematical objects that are ¹²² fixed in the proof and assumption. The last line shows the current goal.

This is already tremendously useful to students, and completely impossible to emulate 123 on a blackboard or in print. But there are also many challenges. The most obvious one 124 is the need to learn the syntax of the software. In order to progress in proofs, one need 125 to call *tactics* that are commands which usually do not look like mathematics but rather 126 like programming. This is not a crucial problem, especially with students who also learn 127 programming in other courses. But it does take time, so this issue prevents using a proof 128 assistant on the side of a regular course with no dedicated time. A much more serious issue 129 is that having a proof assistant syntax that is very different from paper proofs makes it a lot 130 harder to transfer proving skills to paper. This is true in the direct case of transcribing a 131 proof done on the computer but also in the longer run when students try to prove things on 132 paper without writing a formal proof first. 133

A second challenge is to set up the right kind of automation. Traditional paper proofs are very far from mentioning every lemma that is used. There can be some value to write a couple of proofs mentioning every lemma. But doing this systematically is counter-productive with respect to the goals listed above. An extreme example would be lemmas asserting the commutativity and associativity of addition and multiplication of numbers. More generally, things that are too obvious to be mentioned on paper should be done automatically by the proof assistant, whereas things that we want to see justified on paper should not. But this is an extremely vague and inconsistent specification. Powerful automation is also potentially problematic for students who have very slow computers with little memory, and it can make error reporting more complicated. But it is crucial for the success of that kind of teaching use of proof assistants.

A last challenge that is less universal is that some countries do not use English as their
 main language. This is especially relevant for very young students. In France in particular,
 having anything written in English raises a huge barrier for undergraduate students.

¹⁴⁸ **3** Using the library as a student

¹⁴⁹ 3.1 Tactic language

153

In this section we will show what Verbose Lean looks like from the point of view of student
 users. We will show several possible variations but it is probably unwise to show all those
 variations to students. The following is a typical exercise solution.

```
Exercise "Continuity implies sequential continuity"
154
         Given: (f : \mathbb{R} \to \mathbb{R}) (u : \mathbb{N} \to \mathbb{R}) (\mathrm{x}_0 : \mathbb{R})
155
         Assume: (hu : u converges to x_0) (hf : f is continuous at x_0)
156
         Conclusion: (f \circ u) converges to f x_0
157
     Proof:
158
         Let's prove that \forall \varepsilon > 0, \exists N, \forall n \ge N, |f(u n) - f x_0| \le \varepsilon
159
        Fix \varepsilon > 0
160
         By hf applied to \varepsilon using that \varepsilon > 0 we get \delta such that
161
            (\delta_{pos}: \delta > 0) and (\text{Hf}: \forall x, |x - x_0| \le \delta \Rightarrow |f x - f x_0| \le \varepsilon)
162
         By hu applied to \delta using that \delta > 0 we get N such that
163
           Hu : \forall n \geq N, |u n - x<sub>0</sub>| \leq \delta
164
        Let's prove that N works : \forall n \geq N, |f (u n) - f x<sub>0</sub>| \leq \varepsilon
165
        Fix n > N
166
         By Hf applied to u n it suffices to prove |u n - x_0| \leq \delta
167
         We conclude by Hu applied to n using that n \geq N
168
      QED
<del>1</del>98
```

The first difference with the default syntax of Lean is that the statement clearly distinguishes the data, the assumption and the conclusion. Then each line of the proof sounds like natural language, but it is actually as rigid as any programming language. Remember the goal of this language is not to be easier to write, it is to be easier to transfer to paper.

The first line is completely optional, it only unfolds a definition. The second line shows 175 how bounded quantifiers are handled by the library. The core logic of Lean does not involve 176 those quantifiers. Given a predicate P, say on real numbers, the statement $\forall \varepsilon > 0, P(\varepsilon)$ is a 177 notation for $\forall \varepsilon, \varepsilon > 0 \implies P(\varepsilon)$. In Verbose Lean, introducing a positive number can be 178 done in one step (of course it can also be done in two steps). This generates a name ε_{pos} 179 for the assumption that ε is positive. The next tactic, that spans the third and fourth lines 180 above, uses that positivity but in a declarative way. However it does use the name hf that 181 was assigned to the continuity assumption on f. A syntactic variant here would be to write 182 183

Since f is continuous at x_0 and $\varepsilon > 0$ we get δ such that $(\delta_{pos} : \delta > 0)$ and $(\text{Hf} : \forall x, |x - x_0| \le \delta \Rightarrow |f x - f x_0| \le \varepsilon)$

which only list claims and does not even explicitly mention ε before mentioning its positivity.

213

which use both the nameless approach and forward reasoning only (those two aspects are independent, we gathered them only to prevent a combinatorial explosion of examples). The nameless approach is not purely stylistic, it also involves some implicit reasoning. For instance, with the same assumptions, we could write since $\varepsilon \geq 0$ and the library would silently derive this from $\varepsilon > 0$.

The example used so far only uses two definitions and the rules of logic. It features both using and proving statements formed using quantifiers and implication. Stating a universally quantified claim on line one and starting its proof on line 2 are very distinct operations. The first one involves the quantifier symbol while the second one involves the word "Fix".

The confusion between stating existence and extracting a witness is handled in a more subtle way. We could first state the existence using the symbol $\exists \delta$ and then have something like "Let us fix such a δ ". Nothing prevents a teacher from implementing this syntax, but the trick of saying we get δ such that is a very nice compromise which distinguishes fixing a witness from merely stating existence and which stays very light to read.

The confusion between claiming an implication and using it is handled very simply. First the implication symbol is used only when stating. Then both backward and forward uses of implication mention both the implication and its premise. This applies both to the style referring to assumption names and to the nameless style.

Let us now consider another example: proving the squeeze theorem.

```
214
      Example "The squeeze theorem."
215
        Given: (u v w : \mathbb{N} \to \mathbb{R}) (1 : \mathbb{R})
216
        Assume: (hu : u converges to 1) (hw : w converges to 1)
217
                    (h : \forall n, u n \leq v n) (h' : \forall n, v n \leq w n)
218
        Conclusion: v converges to 1
219
     Proof:
220
        Fix \varepsilon > 0
221
        Since u converges to 1 and \varepsilon > 0 we get N such that
222
           hN : \forall n \geq N, |u n - 1| \leq \varepsilon
223
        Since w converges to 1 and \varepsilon > 0 we get N' such that
224
           hN': \forall n > N', |w n - 1| < \varepsilon
225
        Let's prove that max N N' works : \forall n \geq max N N', |v n - 1| \leq \varepsilon
226
        Fix n \geq max N N'
227
        Since n \ge max N N' we get (hn : n \ge N) and (hn' : n \ge N')
228
        Since \forall n \geq N, |u n - 1| \leq \varepsilon and n \geq N we get
229
          (hNl : -\varepsilon \leq u n - 1) and (hNd : u n - 1 \leq \varepsilon)
230
        Since \forall n \geq N', |w n - 1| \leq \varepsilon and n \geq N' we get
231
            (hN'l : -\varepsilon \leq w n - 1) and (hN'd : w n - 1 \leq \varepsilon)
232
        Let's prove that |v n - 1| \leq \varepsilon
233
        Let's first prove that -\varepsilon \leq v n - 1
234
        Calc -\varepsilon \leq u n - 1 by assumption
235
                    \leq v n - l since u n \leq v n
236
        Let's now prove that v n - 1 < \varepsilon
237
        Calc v n - l \leq w n - l since v n \leq w n
238
                           \leq \varepsilon
                                           by assumption
239
      QED
240
241
```

The beginning of the proof uses the same tactics as our first example. New things start 242 with the line Since $n \ge max N N'$ we get (hn : $n \ge N$) and (hn' : $n \ge N'$). Here we 243 are using a lemma claiming that $n \ge \max(N, N')$ implies that $n \ge N$ and $n \ge N'$. But this 244 lemma is not mentioned explicitly. The tactic saw that the claim $n \ge \max(N, N')$ is not a 245 conjunction so it tried splitting it into the announced conclusions using so-called anonymous 246 fact splitting lemmas. Here anonymous refers to the fact that their names do not appear 247 in the proof script, but of course they actually do have names. The next two tactics (each 248 spanning two lines) use the exact same mechanism using an anonymous lemma that split an 249 inequality with shape $|x| \leq y$ into $-y \leq x$ and $x \leq y$. 250

The next tactic Let's prove that $|v n - 1| \leq \varepsilon$ is completely optional, it recalls what is the current goal since it was never explicitly spelled out and we just went through three tactics that created new facts without changing current goal. This tactic could also have been used right after Fix $n \geq \max N N'$.

The next line is something new: Let's first prove that $-\varepsilon \leq v n - 1$. This tactic can be used to start a conjunction proof. But here the current goal is not a conjunction, it is turned into a conjunction by a so-called anonymous goal splitting lemma, which happens to be the converse the anonymous fact splitting lemmas used before (but those are completely separate lemmas from the framework point of view).

This tactic does a bit more than applying the lemma and splitting the resulting conjunction. Indeed we want to force students to announce the second part of the conjunction when the first one is proven. So instead of showing directly the second goal, the tactic state displays: You need to announce: Let's now prove that $v n - 1 \leq \varepsilon$ and refuses³ any other tactic.

Returning to what happens during the proof of the first inequality, we see some com-265 putation introduced by the Calc word. This is based on the builtin Lean calc tactic, but 266 the justifications are specific to our library. The first one in the example is by assumption 267 which implicitly refer to the hNl assumption. A more explicit justification could be from hNl. 268 What comes after the word from could also contain the words applied to and using that 269 as in the third tactic of our first example. The next justification uses since which indicates 270 a nameless approach: we claim that $u n \leq v n$ without explaining why; the tactic has to 271 instantiate the assumption **h** to the free variable **n**. But there is more to it since this fact by 272 itself is not sufficient to justify that calculation step. The tactic has to secretely invoke a 273 lemma saying that $\forall x, y, z, x \leq y \implies x - z \leq y - z$. This is handled internally by calling 274 the gcongr tactic of Carneiro and Macbeth. 275

Those two examples illustrate the main mechanisms that we use to get students to develop proof skills that are easier to transfer to paper than using the native Lean tactics. Of course they do not exhaust the list of tactics provided by our library. In particular there are tactics that allow to prove things by case disjunctions, using contraposition or proof by contradiction, or using the axiom of choice.

281 3.2 Assisted modes

The above examples can all be typed in the editor (typically VSCode to avoid teaching at the same time how to use Lean and a powerful editor such as vim or emacs). Lean then answers by updating the proof state and displaying error messages if needed. But mastering

 $^{^{3}}$ The courage to make this mandatory came from seeing it in the Coq waterproof project.

289

the **help** tactics displays:

a lot of syntax is challenging, even if only one proof style is taught (for instance only the
variants that do not use assumption names). So Verbose Lean offers two levels of assistance.

The fist level is the help tactic that can be used inside the proof. For instance, if the current target is $(f \circ u)$ converges to $f x_0$ as at the beginning of our first example then

²⁹⁹ The above has two help messages and two suggestions. Clicking on a suggestion replaces the ³⁰⁰ help tactic with the suggestion.

In the same example, there is a local assumption named hu saying that u converges to x_0 . The answer to help hu is

```
303
     Help
304
       This assumption starts with the application of a definition.
305
        One can make it explicit with:
306
        We reformulate hu as \forall \varepsilon > 0, \exists N, \forall n \ge N, |u n - x_0| \le \varepsilon
307
       The assumption hu starts with "\forall \epsilon > 0, \exists N, \ldots"
308
        One can use it with:
309
        By hu applied to \varepsilon_0 using h\varepsilon_0 we get N such that
310
           (hN : \forall n \geq N, |un - x_0| \leq \varepsilon_0)
311
        where \varepsilon_0 is a real number and h\varepsilon_0 is a proof of the fact that \varepsilon_0 > 0
312
        The names N and hN can be chosen freely among available names.
313
```

Using this tactic with students suggests it already does a lot to tame the syntactic complexity of our controlled natural language. One could fear that students will constantly use it instead of analysing the goal or assumptions themselves, but this was not observed.

Especially in situations where there is not much time allocated to the use of a proof assistant, one can use a more assisted mode where proofs can be assembled at least partly through clicking. In this interaction mode, students click on expressions in the tactic state and get tactics suggestions in return. This subsumes the help tactic: when clicking on the full target or on a full assumption, the suggestions that given are the same that appear in the help command (assuming the default configuration is used). But one can also click on multiple assumptions, or on sub-expressions inside the target or inside an assumption.

For instance the example that proved sequential continuity from continuity can be done entirely by clicking. Clicking on the full goal suggests the first two lines of the proof. Then one needs to specialize the continuity assumption to the positive ε that was just introduced. This is done by clicking on the assumption and then clicking on ε . With this selection in the tactic state, one gets the following suggestions:

```
 \begin{array}{l} \overset{330}{331} & \cdot \text{ By hf applied to } \varepsilon \text{ using that } \varepsilon > 0 \text{ we get } \delta \text{ such that } (\delta_{pos} : \delta > 0) \\ & (h\delta : \forall (x : \mathbb{R}), |x - x_0| \leq \delta \Rightarrow |f x - f x_0| \leq \varepsilon) \\ \overset{333}{333} & \cdot \text{ By hf applied to } \varepsilon / 2 \text{ using that } \varepsilon / 2 > 0 \text{ we get } \delta \text{ such that } (\delta_{pos} : \delta \\ & > 0) (h\delta : \forall (x : \mathbb{R}), |x - x_0| \leq \delta \Rightarrow |f x - f x_0| \leq \varepsilon / 2) \\ \end{array}
```

Those two suggestions have the same shape but use either ε or $\varepsilon/2$ since specializing to half a given number is a very common move in elementary analysis and the default configuration is biased towards this kind of mathematics. In the above example, the library does not check that it will be able to automatically prove the positivity side condition that appears after using that. This lets students judge

³⁴¹ the different suggestions.

357

358

³⁴² **4** Using the library as a teacher

343 4.1 Basic configuration

In this section we explain various mechanisms used for configuration which do not require
programming expertise from teachers (of course a lot more is possible with programming).
The goal is not to document every configuration possibility – since this is not a manual –
but to show the configuration mechanisms that we use. We also show how this configuration
depends on specific pedagogical goals, students expertise and time constraints.

The first decision to make is how much automation, if any, is desired when implicitly 349 using lemmas. As explained in the previous section, there are two kinds of such lemmas, 350 depending on whether they split a given fact or the current goal. For instance lemmas in 351 the second category can be configured using configureGoalSplitingLemmas Iff.into 352 Subset.antisymm. Listing a lot of lemmas in these commands could become very tedious. 353 So we allow defining lemmas lists and referring to them in the configuration commands. We 354 also pre-define some lists. Defining a list named MyList which contains the pre-defined list 355 LogicIntros and the lemma proving set equality from double inclusion is done using 356

```
AnonymousGoalSplittingLemmasList MyList := LogicIntros Subset.antisymm
```

Those commands and all configuration commands are meant to be "hidden" in the teacher file. When changes are needed in the middle of an exercise file, one can use a macro. For instance macro "switchConf" : command => `(configureGoalSplittingLemmas x) would allow to simply to write switchConf between two exercises to avoid a long distracting line. Note that teachers will probably want to tell students about the list of anonymous lemmas, at least informally, to avoid creating confusion about what needs to be justified.

Tactics can also have configuration flags. For instance, say the goal is $\neg \exists x:\mathbb{Q}, x^2 = 2$. By default, starting the proof with Assume for contradiction $H : \exists x : \mathbb{Q}, x^2 = 2$ will lead to the error message: "The goal is a negation, there is no point in proving it by contradiction. You can directly assume $\exists x : \mathbb{Q}, x^2 = 2$ ". Teachers who fully embrace the confusion between a direct proof of a negation and a proof by contradiction, or simply need to focus on another battle, can use allowProvingNegationsByContradiction.

The next thing to configure is the assisted modes (help tactic and suggestion widget). 372 First there are commands to disable those modes for teachers who want students to write 373 everything by hand (in an exam setting, one can simply delete the relevant file for extra 374 safety). Assuming they are enabled, many aspects are configurable. Each help message comes 375 from a function and one can configure the available functions using the same kind of lists 376 as with anonymous lemmas. For instance, in a basic course which progressively introduces 377 different kinds of reasoning, one can disable messages suggesting a proof by contradiction in 378 the beginning. One can also modify existing help functions with no programming knowledge 379 by copy-pasting and editing only the text. 380

If the current lecture focusses on students knowing definitions then one can completely disable the help that unfolds definitions. One can also control in detail which definitions participate in unfolding suggestions. This has to be an opt-in mechanism to avoid having suggestions unfolding fundamental definitions such as the definition of real numbers or even the definition of addition of natural numbers. For instance the teacher file could contain: 336 configureUnfoldableDefs continuous seq_limit assuming the teacher library defined 387 continuous and seq_limit.

The suggestion widget is also fully configurable. Really changing its behavior requires programming. But an easy tweak is to change how to produce data from the selection. We saw that selecting a real number ε and a universally quantified assumption h does not only propose to specialize h to ε but also to $\varepsilon/2$. The configuration for this can look like:

```
392
    dataProvider mkSelf a := a
393
    dataProvider mkHalf a := a/2
394
    dataProvider mkMin a b := min a b
395
     dataProvider mkMax a b := max a b
396
    DataProviderList CommonProviders := mkSelf mkMin mkMax
397
398
     configureDataProviders {
399
       \mathbb{R} : [CommmonProviders, mkHalf],
400
       \mathbb{N}:
           [CommmonProviders] }
481
```

In addition to a declaration list CommonProviders analogous to the one we use for 403 anonymous lemmas, there are two new kinds of micro–DSLs (domain specific languages) 404 here. First we define four "functions" with the dataProvider command which features no 405 type information at all. Those are purely syntactic objects. They only participate in creating 406 the widget suggestions on the syntactic level. Writing meaningless functions there would 407 of course create trouble when accepting suggestions. Finally there is a JSON-like syntax in 408 the configureDataProviders that registers data providers for different types. The goal of 409 those DSL is to allow configuring this even for teachers who basically know nothing about 410 Lean, maybe not even enough to write a function that can perform an algebraic operation 411 either on natural numbers or on reals. 412

All the configuration options mentioned so far are specific to the Verbose library, but of course they come on top of the usual Lean configuration. A lot of the flexibility offered by Lean out of the box is very relevant to the kind of teaching targeted by our library.

This includes the whole parsing and elaboration pipelines. For instance Verbose Lean overwrites the notation for implication to use the double arrow symbol that is normally used in mathematics instead of a single arrow. The examples in this paper also use an infix notation for continuity and limits as in u tendsto x.

Overriding notation is not only about having a nice output. It also help mitigating 420 unwanted side-effects of using type theory. For instance say we want to use the sequence 421 of real numbers $n \mapsto 1/(n+1)$. Using its default configuration, Mathlib may need help 422 to understand that 1/(n+1) is a real number and not a natural number. The correct 423 interpretation can be enforced using a type ascription such as $1/(n+1;\mathbb{R})$. But this is 424 distracting for students in the provided code, and failing to use such ascription in their own 425 code can lead to inscrutable error messages. In this case it is much easier to override the 426 meaning of the division symbol to always mean division or real numbers. One can also use 427 a custom notation for function abstraction specialized to sequences of real numbers. For 428 instance one can ensure seq $n \mapsto \ldots$ gets expanded to fun $n : \mathbb{N} \mapsto (\ldots : \mathbb{R})$. 429

Note there is no way to completely avoid type ambiguities in the input without type ascriptions. We have seen students feeling the need to state as an intermediate fact something like 0 < 1. Here there is no way Lean can guess whether this is meant as an inequality of real numbers or of natural numbers. And there is no way students can be aware of this issue without discussing the subtle status of the "inclusion" of \mathbb{N} into \mathbb{R} . Lean will interpret the above statement as an inequality of natural numbers and our tactics will happily prove it. But then using this intermediate fact will fail if the intended meaning was an inequality of
real numbers. Note that Lean has an option, namely pp.numeralTypes, to always decorate
literal numbers such as 0 or 1 when it displays them. This helps making to above problem
easier to spot, but it does not fix the input issue and does not avoid discussing the subtlety.

440 4.2 Translating to a new language

447

The English language can be a huge barrier for undergraduate students. One can also imagine teachers who want to use a dialect of English. Verbose Lean comes with an English version and a French version. Adding a new language can be done by imitation without any knowledge of Lean programming. The process is to copy the English folder of the Verbose library and replace English words. In case of doubt, comparing the French and English versions can show the required modifications. For instance we see⁴ in the English version:

```
declare syntax cat maybeApplied
448
    syntax term : maybeApplied
449
    syntax term "applied to " term : maybeApplied
450
    syntax term "applied to " term " using " term : maybeApplied
451
    syntax term "applied to " term " using that " term : maybeApplied
452
453
    def maybeAppliedToTerm : TSyntax 'maybeApplied \rightarrow MetaM (TSyntax 'term)
454
    | `(maybeApplied| $e:term) => pure e
455
    \'(maybeApplied| $e:term applied to $x:term) => '($e $x)
456
      `(maybeApplied| $e:term applied to $x:term using $y) => `($e $x $y)
457
    | `(maybeApplied| $e:term applied to $x:term using that $y) =>
458
        `($e $x (strongAssumption% $y))
459
    | _ => pure default
460
461
    elab "We" " conclude by " e:maybeApplied : tactic => do
462
      concludeTac (\leftarrow maybeAppliedToTerm e)
463
```

We will comment more on this code in the next section. Our point here is that we see a lot of mysterious things but understanding them is not required to write the French version:

```
467
    declare_syntax_cat maybeAppliedFR
468
    syntax term : maybeAppliedFR
469
    syntax term "appliqué à " term : maybeAppliedFR
470
    syntax term "appliqué à " term " en utilisant " term : maybeAppliedFR
471
    syntax term "appliqué à " term " en utilisant que " term : maybeAppliedFR
472
473
    def maybeAppliedFRToTerm : TSyntax 'maybeAppliedFR \rightarrow MetaM Term
474
    | `(maybeAppliedFR| $e:term) => pure e
475
     `(maybeAppliedFR| $e:term appliqué à $x:term) => `($e $x)
476
    | `(maybeAppliedFR| $e:term appliqué à $x:term en utilisant $y) => `($e $x $
477
        y)
478
    | `(maybeAppliedFR| $e:term appliqué à $x:term en utilisant que $y) =>
479
        `($e $x (strongAssumption% $y))
480
        => pure default
481
482
    elab "On" " conclut par " e:maybeAppliedFR : tactic => do
483
      concludeTac ( maybeAppliedFRToTerm e)
484
485
```

 $^{^4}$ The actual code has some more cases that were removed here for conciseness.

486 **5** Some implementation mechanisms

The previous sections have been all about the pedagogical choices of the library, about how they can be tweaked by teachers and how students can use them. We now switch gears and turn to the question of the Lean mechanisms that allow all this. Many of those mechanisms are specific to Lean 4, the new family of versions of Lean that was officially released for the first time in September 2023 and puts flexibility of use in the center [5].

The flexibility of Lean as a proof assistant rests on two main pillars. The first one is that Lean is also a programming language and that almost all of Lean is implemented in Lean. This allows in particular to offer the possibility of overriding a lot of what Lean is doing, even pretty deep down, but we don't really use that directly. However we certainly use Lean as a programming language here, so we need a very quick introduction to that.

Lean is a pure functional programming language. So, fundamentally, it never does 497 anything but defining and applying functions. However Lean makes extensive use of the 498 monad pattern together with an extremely sophisticated notation system that can make it 499 look a lot like imperative programming, but without the bad surprises [15]. An important 500 inspiration is Haskell here. For our purposes, one can think of a monad M as a programming 501 environment with a well specified state that can be read or written during program execution, 502 a well specified way it can fail or not, and a well specified way of interacting or not with the 503 outside world (such interactions could include reading files or printing things for instance). 504 For any type α , the type M α is the type of programs that can do all this and return an 505 element of type α (when they don't throw an exception if M includes the exception throwing 506 capability). Running such a program requires passing in an initial state and actually also 507 return the new state if the state includes writable parts. 508

An important example is the **CoreM** monad defined by Lean itself. It allows to describe and run programs that have read and write access to all definitions in scope, read only access to options, can fail by throwing certain kinds of exceptions, and can interact with the outside world (of course it has a more precise definition, we only give the flavor here).

On top of this CoreM monad sits the MetaM monad which mainly adds read and write 513 access to the meta-variable context. A meta-variable is a place-holder that can be used 514 in particular in a partially constructed proof. For instance at the very beginning of an 515 interactive proof of a lemma, the full proof is a single meta-variable. By itself a meta-variable 516 only stores a unique identifier. The meta-variable context is a data structure holding in 517 particular for each meta-variable its expected type (that would be the conclusion of the 518 lemma in our example) and its local context (that would be the assumptions of the lemma). 519 On top of MetaM sits the TacticM monad which describes tactics, with additional access to 520 all the relevant goals. 521

The second pillar of flexibility is the existence of typed concrete syntax objects as first 522 class citizens [14, 13]. Here an important source of inspiration is the family of LISP languages, 523 especially modern incarnations such as Racket. This is crucial for us. First it is crucial to our 524 translation system that the syntax of tactics is clearly separated from their implementation. 525 It also allows the assisted modes to provide suggestions that are guaranteed to be syntactically 526 correct, because they produce syntax objects that are then printed as strings. This is seen in 527 the snippets above that use syntax quotations such as '(maybeApplied| \$e:term). Syntax 528 objects can also be turned into other syntax objects, either by macros such as the one showed 529 in Section 4 or by functions in the library. 530

In order to explain this, we will follow part of the journey that allowed us to write in an earlier example: By hu applied to δ using that $\delta > 0$ we get N such that Hu :

```
<sup>533</sup> \forall n \geq N, |u n - x_0| \leq \delta. The place where the corresponding syntax is hooked to the
<sup>534</sup> tactic implementation is
```

```
signed state =>
signed by the state =>
s
```

The **elab** command is a shortcut that allows to define a syntax in the first line and immediately 539 assign it some meaning in the second line. On the first line we see two literal strings and two 540 variables e and news. Those variables are also syntax objects with some syntax categories 541 that are defined in Verbose Lean. They are based in the crucial syntax category term which 542 is used for syntax representing Lean expressions. The (simplified) definition of maybeApplied 543 representing functions that may be applied to arguments, and then a function turning such 544 syntax objects into terms syntax objects have been seen in the first code snippet in Section 4.2. 545 The first line of that snippet registers our syntax category and the next four lines describe 546 four ways to build a syntax object in this syntax category. Those four ways are very simple 547 and combine terms and literal words. 548

Now we can move to the second line of our elab command that calls the actual tactic. The do keyword starts a monadic program. Here it is a TacticM Unit program, i.e. a program in the TacticM monad that returns nothing – its only purpose is to manipulate the state of this monad. The type of the three involved functions are

We will ignore the third function. Note that the return type of the second one does not match the type of the first input to the first function. We need a term and not only a program computing a term in the MetaM monad. Fortunately, program in this monad can be used in the TacticM monad which extends it. And the arrow in (\leftarrow maybeAppliedToTerm e) tells Lean to run the maybeAppliedToTerm e program and feed the result at this spot.

The definition of the maybeAppliedToTerm function was shown in Section 4.2. It is of 563 course defined by pattern matching on the four possibilities to create a maybeApplied syntax 564 object (plus a wild card possibility that is required because the syntax category could in 565 principle be extended after the definition of this function). Note that the last interesting case 566 also uses the strongAssumption% \$y macro that expands to (by strongAssumption : \$y) 567 where strongAssumption is one of our tactics. Of course we could have used the expanded 568 version here but the macro is used in several other places. Hence this example is a library 569 function turning some syntax objects into other syntax objects using pattern matching and a 570 macro that also does such a transformation. 571

The result type is not directly TSyntax 'term but a program in the MetaM monad. This is because the quotation mechanism includes hygiene guarantees, a mechanism preventing accidental name capture and requiring some information from the surrounding context.

Note that one could inline this function into the obtainTac function. But this would make the latter into a language dependent function. As we saw in Section 4.2, Verbose Lean contains a French maybeAppliedFR syntax category and a function to turn syntax objects in this category into terms. It then uses the exact same obtainTac function from the language agnostic part of the library. Hopefully this already illustrate how we put to good use the monadic meta-programming framework of Lean and, crucially, its treatment of syntax objects. Syntax objects are also used to implement all the little DSLs we saw earlier.

We now want to discuss part of the implementation of obtainTac. Its first task is to turn the term it got as its first argument into a Lean expression, i.e. an abstract syntax

Patrick Massot

⁵⁸⁴ object whose type Expr is an inductive type whose main constructors correspond to the

 $_{585}$ fundamental operations of lambda calculus: function application, function abstraction...

This process is possible in the TacticM monad which has access to all definitions in scope as well as to the local context of the proof where this tactic is used.

Then obtainTac first tries to decompose the type of this expression. In our example this 588 type is $\exists N, \forall n > N, |u n - x_0| < \delta$ which can indeed be decomposed as a witness N 589 and its property. This job of obtainTac is a trivial wrapper around a standard Lean tactic. 590 More interestingly, when such a decomposition does not make sense, the tactic will try to 591 apply an anonymous splitting lemma. We saw already how to configure the lemmas that are 592 tried. Now we want to discuss how this configuration is stored and updated. We saw that 593 programs in the CoreM monad have access to existing declarations. More generally they have 594 access to the so-called environment that stores declarations but also a lot more information. 595 Lean allows to declare environment extensions storing user information for later use. In the 596 case at hand, the stored information is simply a list of lists of declaration names. But our 597 library also uses more interesting extensions for assisted modes. 598

The multilingual support for help and suggestions is based on a multilingual function dispatch framework by Mario Carneiro. Multilingual functions are first register it using the **register_endpoint** command. This gives a function that can immediately be used to define other functions. But running those functions requires implementing the endpoint in the current language, which is **en** by default but can be changed using **setLang**. For instance:

```
604
    /-- Multilingual hello function. -/
605
    register_endpoint hello : CoreM String
606
607
     /-- Greeting function refering to our endpoint before any implementation
608
        is defined. -/
609
    def greet (name : String) : CoreM String :=
610
      return (~ hello) ++ " " ++ name
611
612
    #eval greet "Patrick" -- throws error: no implementation of hello found
613
        for language en
614
615
    implement_endpoint (lang := en) hello : CoreM String := return "Hello"
616
617
    implement_endpoint (lang := fr) hello : CoreM String := return "Bonjour"
618
619
    #eval greet "Patrick" -- returns "Hello Patrick"
620
621
    setLang fr
622
623
    #eval greet "Patrick" -- returns "Bonjour Patrick"
624
625
```

Note that above example creates three declarations: hello, hello.en and hello.fr, but only the first one is explicitly used. The implementations in this example are silly since they do no perform anything inside the CoreM monad, they simply return a value without reading or writing any CoreM state. But the hello function itself, which is created by the register_endpoint command, crucially uses CoreM to fetch the relevant information from the environment after reading the current language setting.

The way this is achieved illustrates an important point about Lean's flexibility. Lean as a proof assistant has very strong soundness guarantees, and the whole proof checking is handled by its type system. This translates to the default behavior of Lean as a programming language. We saw that being a pure functional programming language does not prevent ⁶³⁶ us from accessing state and having side-effects. One simply has to be honest about it by
⁶³⁷ announcing in which monad we are working. Lean also allows to throw away type safety
⁶³⁸ as long as we clearly announce it. And of course functions which do that cannot appear in
⁶³⁹ proofs (they can participate in creating proofs, but can't appear in the end result).

A very simple version of the trick used in the multilingual framework is implemented 640 in the example below. The runFunctionOn takes a string and a natural number, searches 641 the environment for a declaration whose name is that string, then forcefully tells the Lean 642 type system that this declaration is a function from natural numbers to natural numbers 643 and applies it to the given number. The result is in CoreM \mathbb{N} rather than \mathbb{N} since it needs 644 access to the environment to search for the relevant declaration and it could fail to find it 645 so it needs to be able to throw errors – this is what happen in the second example below. 646 But the new piece is the **unsafe** signpost in front of **def**. Indeed the declaration could be 647 found but not with type $\mathbb{N} \to \mathbb{N}$, bringing us into undefined behavior territory. This is what 648 happen in the third example below where a function concatenating strings is found. 649

```
650
     unsafe def runFunctionOn (function : String) (a : \mathbb{N}) : CoreM \mathbb{N} := do
651
       let myFun \leftarrow evalConst (\mathbb{N} \rightarrow \mathbb{N}) (Name.mkSimple function)
652
       return myFun a
653
654
     def foo (a : ℕ) := 2*a + 1
655
656
     #eval runFunctionOn "foo" 1 -- returns 3
657
658
     #eval runFunctionOn "baz" 1 -- fails with error message: unknown
659
660
         declaration baz
661
     def bar (a : String) := a ++ a
662
663
     #eval runFunctionOn "bar" 1 -- crashes Lean
664
```

The moral is that programming in Lean allows to do very unsafe things that completely 666 bypass the guarantees offered by the type system, but this must be clearly flagged and 667 cannot be used as a proof (explaining how soundness is protected is beyond the scope of this 668 discussion). Note that our actual multilingual dispatch is much more careful and checks that 669 types match before calling functions so that teachers don't crash Lean when they make a 670 type mistake in the implementation of a new help message. For performance reasons we 671 don't want to perform this check at every function call, so we also use an extension that 672 keeps track of a list of endpoints and type-checked implementations. 673

Although we insisted on our use of concrete syntax objects, we also use abstract ones, 674 with type Expr. in the tactic backends. We even have a custom version VExpr with many 675 more constructors that are useful when analysing goals and assumptions in assisted modes. 676 For instance bounded quantifiers have dedicated constructors. We have a function parsing 677 an Expr into a VExpr, hence factoring out work that many help functions would need to do. 678 The type of help functions that analyse the goal is MVarId \rightarrow VExpr \rightarrow SuggestionM Unit 679 where MVarId is used to indicate the relevant goal and the SuggestionM monad accumulates 680 suggestions while providing access to the MetaM monad. Such functions are registered as 681 part of the configuration by teachers, together with a pattern indicating (coarsely) which 682 kind of goal they comment on. Calling the help tactic uses a discrimination tree to quickly 683 locate functions with the relevant pattern and then check whether they are active in the 684 configuration. This use of discrimination trees is not necessary until someone implements 685 thousands of help functions, but the infrastructure is provided by Lean so it is free. 686

Patrick Massot

The last piece of Lean infrastructure that we want to comment on is the framework that 687 allows us to build the suggestion widget. There are quite a few layers here. Lean implements 688 the Language Server Protocol (LSP) with many extensions related to the so called info view 689 which gather the tactic state display, various messages and user-defined widgets. Deep down, 690 widgets are Javascript modules that export a React component that is displayed by the 691 VSCode extension, can access information from Lean and modify the current document. 692 However the ProofWidgets library [9] offers a powerful abstraction that allows us to ignore 693 Javascript. In particular it features a JSX-like DSL as well as React components written in 694 Lean and having a Lean interface. As a result, Verbose Lean does not contain a single line of 695 actual Javascript or HTML. For instance the loop printing the suggestions is: 696

```
697
    for (linkText, newCode, range?) in suggs do
698
       let p : MakeEditLinkProps := .ofReplaceRange doc.meta
699
          \langle \texttt{params.pos}, \texttt{params.pos} \rangle (<code>ppAndIndentNewLine curIndent newCode) range?</code>
700
       children := children.push
701
         style={json% {"margin-bottom": "1rem"}}>
702
            <MakeEditLink
703
                edit={p.edit} newSelection?={p.newSelection?} title?={p.title?}>
704
              {.text linkText}
705
706
            </MakeEditLink>
         787
```

The li tag is directly turned into a HTML list item, whereas MakeEditLink refers to a ProofWidgets component. This component is rendered as a HTML link which, upon getting clicked, edits the proof script. All this is fully type-correct Lean code, with real-time typechecking and the expected editor support (for instance ctrl-clicking on MakeEditLink jumps to the relevant declaration).

The tactic state natively allows to select names or sub-expressions in the local context or 714 the goal. It records this information as an array where each element inhabits an inductive type 715 Lean.SubExpr.GoalLocation having a constructor for each kind of selection, for instance 716 a constructor for en element of the local context, one for a sub-expression in the type of 717 such an element, etc. This layout is not convenient for our purposes so we introduce another 718 datatype SelectionInfo which gather the same information by type of selections. We 719 also have many functions querying this information. Each suggestion provider has type 720 SelectionInfo \rightarrow MVarId \rightarrow WidgetM Unit analogous to the help function type. 721

6 Related work

722

Both the dream of using proof assistants for teaching and the work on alternative interfaces
are very common. However most teaching uses focus on computer science, logic or discrete
mathematics, or even on proof assistants for themselves.

One notable exception is the work of Heather Macbeth at Fordham university [8]. However her course is more focused on computations and less on reasoning, so the need for a controlled natural language is less pressing. What is common to both contexts is the need for automation that is adapted to the level of details expected from students. And indeed some of our tactics rely on tactics develop for Macbeth's course.

Even more relevant is the comparison with the Coq-Waterproof project [16] that was developed independently and shares many goals with Verbose Lean. Discussing with its authors led to several improvements in our work. One thing that we still do not have is a nice custom text editor to mix rendered comments and interactive exercises. On the other hand, we do benefit from using a very flexible proof assistant that easily allows syntactic ⁷³⁶ freedom and interactive interfaces, as explained in this paper. The resulting proof scripts are
 ⁷³⁷ closer to paper proofs and the user interaction model is richer.

Also very relevant and interesting is the Diproche system [3, 4]. Its focus on controlled natural language is even greater than in Verbose Lean. Proofs are sequences of assertions in a more flexible language. Those assertions are sent to an automated prover that complains if it cannot justify a step. The main downside is that the proof structure is much less clear.

On the topic of alternative interfaces, there are again many attempts that seem practicable only for pure logic. For instance this is the case of the Actema project [6] which proposes a drag and drop interface that is partly a more graphical version of our suggestion widget and can interface with Coq. However it seems difficult to integrate with computations as in our squeeze theorem example (which was chosen as the simplest example involving computations). And it very explicitly targets leaving no written trace at all, hence has a very different goal.

Also in the same category but explicitly targeting teaching very young university students is the d∀∃duction project [11]. It features a graphical interface based on selecting subexpressions and clicking buttons, with a completely invisible Lean backend. Again there is no written trace and it seems difficult for students to transfer their new skills to paper.

Edukera [10] is another point and click interface that produces a written text. It is a web interface based on Coq. Its first main drawback is that teachers cannot write exercises or configure anything, they simply have access to a fixed set of exercises. In addition, there is no possibility to directly input text. The interface is only based on clicks and the text is purely on the output side. Also the development of Edukera stopped in 2018.

As far as we know, none of those very nice projects have multilingual support except for
Edukera. Most of them are only in English, Diproche is only in German and d∀∃duction is
only in French.

760 7 Future work

Although some version of this library has been used for five years in University Paris-Saclay 761 at Orsay, it is very much a work in progress, especially since the switch to Lean 4 made it 762 a lot more flexible. The suggestion widget in particular has not been used with students 763 yet, and is bound to need refinements and extensions. Error reporting is also a never-ending 764 work in progress. Each new interface or piece of automation requires more care in case of 765 incorrect input, and students always find new ways to trigger unexpected error messages. 766 On the multi-lingual side, one short-term goal is to make it easier to create variants of an 767 existing language. Another project is to offer more exercises that are ready to use or modify 768 for teachers. Existing exercises are not yet all ported to Lean 4, and new ones should be 769 created in different fields of elementary mathematics. 770

On the pedagogical side, there are plan to rigorously assess the benefits of using this library with the APPAM⁵ team which includes specialists in education sciences.

773 — References

Jeremy Avigad. Learning Logic and Proof with an Interactive Theorem Prover, pages 277–290.
 Springer International Publishing, Cham, 2019. doi:10.1007/978-3-030-28483-1_13.

Evmorfia-Iro Bartzia, Emmanuel Beffara, Antoine Meyer, and Julien Narboux. Underlying
 theories of proof assistants and potential impact on the teaching and learning of proof.

⁵ https://appam.icube.unistra.fr/

In 12th International Workshop on Theorem proving components for Educational software,
 Rome, Italy, July 2023. Julien Narboux and Walther Neuper and Pedro Quaresma. URL:
 https://hal.science/hal-04227823.

- Merlin Carl. Number theory and axiomatic geometry in the diproche system. In Pedro Quaresma, Walther Neuper, and João Marcos, editors, Proceedings 9th International Workshop on Theorem Proving Components for Educational Software, ThEdu@IJCAR 2020, Paris, France, 29th June 2020, volume 328 of EPTCS, pages 56–78, 2020. doi:10.4204/EPTCS.328.4.
- Merlin Carl, Hinrich Lorenzen, and Michael Schmitz. Natural language proof checking in introduction to proof classes - first experiences with diproche. In João Marcos, Walther Neuper, and Pedro Quaresma, editors, Proceedings 10th International Workshop on Theorem Proving Components for Educational Software, ThEdu@CADE 2021, (Remote) Carnegie Mellon University, Pittsburgh, PA, United States, 11 July 2021, volume 354 of EPTCS, pages 59–70, 2021. doi:10.4204/EPTCS.354.5.
- Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming language. In André Platzer and Geoff Sutcliffe, editors, Automated Deduction - CADE 28
 28th International Conference on Automated Deduction, Virtual Event, July 12-15, 2021, Proceedings, volume 12699 of Lecture Notes in Computer Science, pages 625–635. Springer, 2021. doi:10.1007/978-3-030-79876-5_37.
- Pablo Donato, Pierre-Yves Strub, and Benjamin Werner. A drag-and-drop proof tactic. In Andrei Popescu and Steve Zdancewic, editors, CPP '22: 11th ACM SIGPLAN International Conference on Certified Programs and Proofs, Philadelphia, PA, USA, January 17 - 18, 2022, pages 197–209. ACM, 2022. doi:10.1145/3497775.3503692.
- Marie Kerjean, Frédéric Le Roux, Patrick Massot, Micaela Mayero, Zoé Mesnil, Simon Modeste,
 Julien Narboux, and Pierre Rousselin. Utilisation des assistants de preuves pour l'enseignement
 en L1. In *Gazette de la SMF*, volume 174, Octobre 2022.
- 803 8 Heather Macbeth. The mechanics of proof. https://hrmacbeth.github.io/math2001/.
- Wojciech Nawrocki, Edward W. Ayers, and Gabriel Ebner. An extensible user interface for Lean 4. In Adam Naumowicz and René Thiemann, editors, 14th International Conference on Interactive Theorem Proving, ITP 2023, July 31 to August 4, 2023, Bialystok, Poland, volume 268 of LIPIcs, pages 24:1-24:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.
 URL: https://doi.org/10.4230/LIPIcs.ITP.2023.24, doi:10.4230/LIPICS.ITP.2023.24.
- 809 10 Benoît Rognier. Edukera. https://edukera.com/.
- $11 \quad {\rm Fr\'ed\'eric\ Le\ Roux.\ D} \forall \exists duction.\ https://perso.imj-prg.fr/frederic-leroux/d \forall \exists duction/.$
- Athina Thoma and Paola Iannone. Learning about proof with the theorem prover Lean: the
 abundant numbers task. International Journal of Research in Undergraduate Mathematics
 Education, 8(1):64–93, Apr 2022. doi:10.1007/s40753-021-00140-1.
- Sebastian Ullrich. An Extensible Theorem Proving Frontend. PhD thesis, Karlsruhe Institute of Technology, Germany, 2023. URL: https://nbn-resolving.org/urn:nbn:de:101:
 1-2023080204582480933072.
- Sebastian Ullrich and Leonardo de Moura. Beyond notations: Hygienic macro expansion
 for theorem proving languages. Log. Methods Comput. Sci., 18(2), 2022. URL: https:
 //doi.org/10.46298/lmcs-18(2:1)2022, doi:10.46298/LMCS-18(2:1)2022.
- Sebastian Ullrich and Leonardo de Moura. 'do' unchained: embracing local imperativity in a
 purely functional language (functional pearl). *Proc. ACM Program. Lang.*, 6(ICFP):512–539,
 2022. doi:10.1145/3547640.
- I6 Jelle Wemmenhove, Thijs Beurskens, Sean McCarren, Jan Moraal, David Tuin, and Jim
 Portegies. Waterproof: educational software for learning how to write mathematical proofs,
 2022. arXiv:arXiv:2211.13513.
- Xiaoheng Yan and Gila Hanna. Using the Lean interactive theorem prover in undergraduate mathematics. International Journal of Mathematical Education in Science and Technology, 0(0):1-15, 2023. arXiv:https://doi.org/10.1080/0020739X.2023.2227191, doi:10.1080/ 0020739X.2023.2227191.