


Hammer for Coq: Automation for Dependent Type Theory

Łukasz Czajka¹ · Cezary Kaliszyk¹ 

Received: 30 March 2017 / Accepted: 20 February 2018 / Published online: 27 February 2018
© The Author(s) 2018. This article is an open access publication

Abstract Hammers provide most powerful general purpose automation for proof assistants based on HOL and set theory today. Despite the gaining popularity of the more advanced versions of type theory, such as those based on the Calculus of Inductive Constructions, the construction of hammers for such foundations has been hindered so far by the lack of translation and reconstruction components. In this paper, we present an architecture of a full hammer for dependent type theory together with its implementation for the Coq proof assistant. A key component of the hammer is a proposed translation from the Calculus of Inductive Constructions, with certain extensions introduced by Coq, to untyped first-order logic. The translation is “sufficiently” sound and complete to be of practical use for automated theorem provers. We also introduce a proof reconstruction mechanism based on an *eauto*-type algorithm combined with limited rewriting, congruence closure and some forward reasoning. The algorithm is able to re-prove in the Coq logic most of the theorems established by the ATPs. Together with machine-learning based selection of relevant premises this constitutes a full hammer system. The performance of the whole procedure is evaluated in a bootstrapping scenario emulating the development of the Coq standard library. For each theorem in the library only the previous theorems and proofs can be used. We show that 40.8% of the theorems can be proved in a push-button mode in about 40 s of real time on a 8-CPU system.

Keywords Hammer · Coq · Calculus of inductive constructions · Proof automation

✉ Cezary Kaliszyk
cezary.kaliszyk@uibk.ac.at

Łukasz Czajka
lukasz.czajka@uibk.ac.at

¹ University of Innsbruck, Innsbruck, Austria

1 Introduction

Interactive Theorem Proving (ITP) systems [44] become more important in certifying mathematical proofs and properties of software and hardware. A large part of the process of proof formalisation consists of providing justifications for smaller goals. Many of such goals would be considered trivial by mathematicians. Still, modern ITPs require users to spend an important part of the formalisation effort on such easy goals. The main points that constitute this effort are usually library search, minor transformations on the already proved theorems (such as reordering assumptions or reasoning modulo associativity-commutativity), as well as combining a small number of simple known lemmas.

ITP automation techniques are able to reduce this effort significantly. Automation techniques are most developed for systems that are based on somewhat simple logics, such as those based on first-order logic, higher-order logic, or the untyped foundations of ACL2. The strongest general purpose proof assistant automation technique is today provided by tools called “hammers” [17] which combine learning from previous proofs with translation of the problems to the logics of automated systems and reconstruction of the successfully found proofs. For many higher-order logic developments a third of the proofs can be proved by a hammer in push-button mode [15,52].

Even if the more advanced versions of type theory, as implemented by systems such as Agda [13], Coq [14], Lean [29], and Matita [5], are gaining popularity, there have been no hammers for such systems. This is because building such a tool requires a usable encoding, and a strong enough proof reconstruction.

A typical use of a hammer is to prove relatively simple goals using available lemmas. The problem is to find appropriate lemmas in a large collection of all accessible lemmas and combine them to prove the goal. An example of a goal solvable by our hammer, but not solvable by any standard Coq tactics, is the following.

```
forall (A : Type) (l1 l2 : list A) (x y1 y2 y3 : A),
  In x l1 \\/ In x l2 \\/ x = y1 \\/ In x (y2 :: y3 :: nil) ->
  In x (y1 :: (l1 ++ (y2 :: (l2 ++ (y3 :: nil)))))
```

The statement asserts that if x occurs in one of the lists $l1$, $l2$, or it is equal to $y1$, or it occurs in the list $y2 :: y3 :: nil$ consisting of the elements $y2$ and $y3$, then it occurs in the list

```
y1 :: (l1 ++ (y2 :: (l2 ++ (y3 :: nil))))
```

where $++$ denotes list concatenation and $::$ denotes the list cons operator. Eprover almost instantly finds a proof of this goal using six lemmas from the module `Lists.List` in the Coq standard library:

```
Lemma in_nil : forall (A : Type) (a : A), ~(In a nil).
Lemma in_inv : forall (A : Type) (a b : A) (l : list A),
  In b (a :: l) -> a = b \\/ In b l.
Lemma in_cons : forall (A : Type) (a b : A) (l : list A),
  In b l -> In b (a :: l).
Lemma in_or_app : forall (A : Type) (l m : list A) (a : A),
  In a l \\/ In a m -> In a (l ++ m).
Lemma app_comm_cons : forall (A : Type) (x y : list A) (a : A),
  a :: (x ++ y) = (a :: x) ++ y.
Lemma in_eq : forall (A : Type) (a : A) (l : list A), In a (a :: l).
```

The found ATP proof may be automatically reconstructed inside Coq.

The advantage of a hammer is that it is a general system not depending on any domain-specific knowledge. The hammer plugin may use all currently accessible lemmas, including those proven earlier in a given formalization, not only the lemmas from the standard library or other predefined libraries.

Contributions. In this paper we present a comprehensive hammer for the Calculus of Inductive Constructions together with an implementation for the Coq proof assistant. In particular:

- We introduce an encoding of the Calculus of Inductive Constructions, including the additional logical constructions introduced by the Coq system, in untyped first-order logic with equality.
- We implement the translation and evaluate it experimentally on the standard library of the Coq proof assistant showing that the encoding is sufficient for a hammer system for Coq: the success rates are comparable to those demonstrated by hammer systems for Isabelle/HOL and Mizar, while the dependencies used in the ATP proofs are most often sufficient to prove the original theorems.
- We present a proof reconstruction mechanism based on an `eauto`-type procedure combined with some forward reasoning, congruence closure and heuristic rewriting. Using this proof search procedure we are able to re-prove 44.5% of the problems in the Coq standard library, using the dependencies extracted from the ATP output.
- The three components are integrated in a plugin that offers a Coq automation tactic `hammer`. We show case studies how the tactic can help simplify certain existing Coq proofs and prove some lemmas not provable by standard tactics available in Coq.

Preliminary versions of the translation and reconstruction components for a hammer for Coq have been presented by us at HaTT 2016 [24]. Here, we improve both, as well as introduce the other required components creating a first whole hammer for a system based on the Calculus of Inductive Constructions.

The rest of this paper is structured as follows. In Sect. 2 we discuss existing hammers for other foundations, as well as existing automation techniques for variants of type theory including the Calculus of Constructions. In Sect. 3 we introduce CIC_0 , an approximation of the Calculus of Inductive Constructions which will serve as the intermediate representation for our translation. Section 4 discusses the adaptation of premise selection to CIC_0 . The two main contribution follow: the translation to untyped first-order logic (Sect. 5) and a mechanism for reconstructing in Coq the proofs found by the untyped first-order ATPs 6. The construction of the whole hammer and its evaluation is given in Sect. 7. Finally in Sect. 8 a number of case studies of the whole hammer is presented.

2 Related Work

A recent overview [17] discusses the three most developed hammer systems, large-theory premise selection, and the history of bridges between ITP and ATP systems. Here we briefly survey the architectures of the three existing hammers and their success rates on the various considered corpora, as well as discuss other related automation techniques for systems based on the Calculus of (Inductive) Constructions.

2.1 Existing Hammers

Hammers are proof assistant tools that employ external automated theorem provers (ATPs) in order to automatically find proofs of user given conjectures. Most developed hammers exist

for proof assistants based on higher-order logic (Sledgehammer [63] for Isabelle/HOL [74], HOLyHammer [52] for HOL Light [40] and HOL4 [67]) or dependently typed set theory (Mizar [55] for Mizar [10,73]). Less complete tools have been evaluated for ACL2 [46]. There are three main components of such hammer systems: premise selection, proof translation, and reconstruction.

Premise Selection is a module that given a user goal and a large fact library, predicts a smaller set of facts likely useful to prove that goal. It uses the statements and the proofs of the facts for this purpose. Heuristics that use recursive similarity include SInE [45] and the Meng-Paulson relevance filter [62], while the machine-learning based algorithms include sparse naive Bayes [70] and k -nearest neighbours (k-NN) [51]. More powerful machine learning algorithms perform significantly better on small benchmarks [1], but are today too slow to be of practical use in ITPs [34,58].

Translation (encoding) of the user given conjecture together with the selected lemmas to the logics and input formats of automated theorem provers (ATPs) is the focus of the second module. The target is usually first-order logic (FOL) in the TPTP format [68], as the majority of the most efficient ATPs today support this foundation and format. Translations have been developed separately for the different logics of the ITPs. An overview of the HOL translation used in Sledgehammer is given in [18]. An overview of the dependently-typed set theory of Mizar is given in [72]. The automated systems are in turn used to either find an ATP proof or just further narrow down the subset of lemmas to precisely those that are necessary in the proof (unsatisfiable core).

Finally, information obtained by the successful ATP runs can be used to re-prove the facts in the richer logic of the proof assistants. This is typically done in one of the following three ways. First, by a translation of the found ATP proof to the corresponding ITP proof script [9,64], where in some cases the script may be even simplified to a single automated tactic parametrised by the used premises. Second, by replaying the inference inside the proof assistant [20,50,64]. Third, by implementing verified ATPs [3], usually with the help of code reflection.

The general-purpose automation provided by the most advanced hammers is able to solve 40–50% of the top-level goals in various developments [17], as well as more than 70% of the user-visible subgoals [15].

2.2 Related Automation Techniques

The encodings of the logics of proof assistants based on the Calculus of Constructions and its extensions in first-order logic have so far covered only very limited fragments of the source logic [2,16,69]. Why3 [35] provides a translation from its own logic [33] (which is a subset of the Coq logic, including features like rank-1 polymorphism, algebraic data types, recursive functions and inductive predicates) to the format of various first-order provers (in fact Why3 has been initially used as a translation back-end for HOLyHammer).

Certain other components of a hammer have already been explored for Coq. For premise selection, we have evaluated the quality of machine learning advice [49] using custom implementations of Naive Bayes relevance filter, k-Nearest Neighbours, and syntactic similarity based on the Meng-Paulson algorithm [62]. *Coq Learning Tools* [59] provides a user interface extension that suggests to the user lemmas that are most likely useful in the current proof using the above algorithms as well as LDA. The suggestions of tactics which are likely to work for a given goal has been attempted in ML4PG [48], where the Coq Proof General [6] user interface has been linked with the machine learning framework Weka [41]. SEPIA [39] tries to infer automata based on existing proofs that are able to propose likely tactic sequences.

The already available HOL automation has been able to reconstruct the majority of the automatically found proofs using either internal proof search [43] or source-level reconstruction. The internal proof search mechanisms provided in Coq, such as the `firstorder` tactic [26], have been insufficient for this purpose so far: we will show this and discuss the proof search procedures of `firstorder` and `tauto` in Sect. 6. The `jp` tactic which integrates the intuitionistic first-order automated theorem prover JProver [66] into Coq does not achieve sufficient reconstruction rates either [24]. Matita's ordered paramodulation [7] is able to reconstruct many goals with up to two or three premises, and the congruence-closure based internal automation techniques in Lean [30] are also promising.

The SMTCoq [3] project has developed an approach to use external SAT and SMT solvers and verify their proof witnesses. Small checkers are implemented using reflection for parts of the SAT and SMT proof reconstruction, such as one for CNF computation and one for congruence closure. The procedure is able to handle Coq goals in the subset of the logic that corresponds to the logics of the input systems.

3 Type Theory Preliminaries

In this section we present our approximation CIC_0 of the Calculus of Inductive Constructions, i.e., of the logic of Coq. The system CIC_0 will be used as an intermediate step in the translation, as well as the level at which premise selection is performed. Note that CIC_0 is interesting as an intermediate step in the translation, but is not a sound type theory by itself (this will be discussed in Sect. 5.6). We assume the reader to be familiar with the Calculus of Constructions [22] and to have a working understanding of the type system of Coq [11, 25]. This section is intended to fix notation and to precisely define the syntax of the formalism we translate to first-order logic. The system CIC_0 is intended as a precise description of the syntax of our intermediate representation. It is a substantial fragment of the logic of Coq as presented in [25, Chapter 4], as well as of other systems based on the Calculus of Constructions. The features of Coq not represented in the formalism of CIC_0 are: modules and functors, coinductive types, primitive record projections, and universe constraints on Type.

The formalism of CIC_0 could be used as an export target for other proof assistants based on the Calculus of Inductive Constructions, e.g. for Matita or Lean. However, in CIC_0 , like in Coq, Matita and Lean, there is an explicit distinction between the universe of propositions `Prop` and the universe of sets `Set` or types `Type`. The efficiency of our translation depends on this distinction: propositions are translated directly to first-order formulas, while sets or types are represented by first-order terms. For proof assistants based on dependent type theories which do not make this distinction, e.g. Agda [13] and Idris [19], one would need a method to heuristically infer which types are to be regarded as propositions, in addition to possibly some adjustments to the formalism of CIC_0 .

The language of CIC_0 consists of terms and three forms of declarations. First, we present the possible forms of terms of CIC_0 together with a brief intuitive explanation of their meaning. The terms of CIC_0 are essentially simplified terms of Coq. Below by $t, s, u, \tau, \sigma, \rho, \kappa, \alpha, \beta$, etc., we denote terms of CIC_0 , by c, c', f, F , etc., we denote constants of CIC_0 , and by x, y, z , etc., we denote variables. We use \vec{t} for a sequence of terms $t_1 \dots t_n$ of an unspecified length n , and analogously for a sequence of variables \vec{x} . For instance, $s\vec{y}$ stands for $s y_1 \dots y_n$, where n is not important or implicit in the context. Analogously, we use $\lambda\vec{x} : \vec{t}.t$ for $\lambda x_1 : \tau_1. \lambda x_2 : \tau_2. \dots \lambda x_n : \tau_n. t$, with n implicit or unspecified.

A *term* of CIC_0 has one of the following forms.

- c . A constant.
- x . A variable.
- ts . An application.
- $\lambda x : t.s$. A lambda-abstraction.
- $\Pi x : t.s$. A dependent product. If x does not occur free in s then we abbreviate $\Pi x : t.s$ by $t \rightarrow s$.
- $\text{case}(t, c, n, \vec{\lambda \vec{a}} : \vec{\alpha}.\lambda x : c \vec{p} \vec{a}.\tau, \vec{\lambda \vec{x}}_1 : \vec{\tau}_1.s_1, \dots, \vec{\lambda \vec{x}}_k : \vec{\tau}_k.s_k)$. A case expression. Here t is the term matched on, c is a constant such that

$$I_n(c : \gamma := c_1 : \gamma_1, \dots, c_k : \gamma_k)$$

is an inductive declaration in the global environment (see the definition of inductive declarations below for an explanation), the type of t has the form $c \vec{p} \vec{u}$, the integer n denotes the number of parameters (which is the length of \vec{p}), the type $\tau[\vec{u}/\vec{a}, t/x]$ is the return type, i.e., the type of the whole case expression, $\vec{a} \cap \text{FV}(\vec{p}) = \emptyset$, and $s_i[\vec{v}/\vec{x}_i]$ is the value of the case expression if the value of t is $c_i \vec{p} \vec{v}$.

- $\text{fix}(f_i, f_1 : t_1 := s_1, \dots, f_n : t_n := s_n)$. A mutually recursive fixpoint definition. The value of this is the function f_i (where $1 \leq i \leq n$) defined by s_i . The variables f_1, \dots, f_n may occur in s_1, \dots, s_n . All functions are required to be terminating.
- $\text{let}(x : t := s, u)$. A let-expression locally binding x of type t to s in u .
- $\text{cast}(t, \tau)$. A type cast: t is forced to have type τ .

We assume that the following special constants are among the constants of CIC_0 : Prop, Set, Type, \top , \perp , \forall , \exists , \wedge , \vee , \leftrightarrow , \neg , $=$. We usually write $\forall x : t.s$ and $\exists x : t.s$ instead of $\forall t(\lambda x : t.s)$ and $\exists t(\lambda x : t.s)$, respectively. For \wedge , \vee and \leftrightarrow we typically use infix notation. We usually write $t = s$ instead of $= \tau st$, omitting the type τ . The purpose of having the *logical primitives* $\top, \perp, \forall, \exists, \wedge, \vee, \leftrightarrow, \neg, =$ in CIC_0 is to be able to directly represent the Coq definitions of logical connectives. These primitives are used during the translation. We directly export the Coq definitions and inductive types which represent the logical connectives (the ones declared in the `Init.Logic` module), as well as equality, to the logical primitives of CIC_0 . In particular, `Init.Logic.all` is exported to \forall .

In CIC_0 the universe constraints on Type present in the Coq logic are lost. This is not dangerous in practice, because the ATPs are not strong enough to exploit the resulting inconsistency. Proofs of paradoxes present in Coq’s standard library are explicitly filtered-out by our plugin.

A *declaration* of CIC_0 has one of the following forms.

- A *definition* $c = t : \tau$. This is a definition of a constant c stating that c is (definitionally) equal to t and it has type τ .
- A *typing declaration* $c : \tau$. This is a declaration of a constant c stating that it has type τ .
- An *inductive declaration* $I_k(c : \tau := c_1 : \tau_1, \dots, c_n : \tau_n)$ of c of type τ with k parameters and n constructors c_1, \dots, c_n having types τ_1, \dots, τ_n respectively. We require $\tau \Downarrow \Pi \vec{y} : \vec{\sigma}.\Pi \vec{y}' : \vec{\sigma}'.s$ with $s \in \{\text{Prop}, \text{Set}, \text{Type}\}$ and $\tau_i \Downarrow \Pi \vec{y} : \vec{\sigma}.\vec{x}_i : \vec{\alpha}_i.c \vec{y} \vec{u}_i$ for $i = 1, \dots, n$, where the length of \vec{y} is k and $a \Downarrow b$ means that a evaluates to b . Usually, we omit the subscript k when irrelevant or clear from the context.

For instance, a polymorphic type of lists defined as an inductive type in Type with a single parameter of type Type may be represented by

$$I_1(\text{List} : \text{Type} \rightarrow \text{Type} := \\ \text{nil} : (\Pi A : \text{Type}.\text{List } A), \\ \text{cons} : (\Pi A : \text{Type}.A \rightarrow \text{List } A \rightarrow \text{List } A)).$$

Mutually inductive types may also be represented, because we do not require the names of inductive declarations to occur in any specific order. For instance, the inductive predicates `even` and `odd` may be represented by two inductive declarations

$$\begin{aligned}
 I_0(\text{even} : \text{nat} \rightarrow \text{Prop}) := & \\
 & \text{even}_0 : \text{even } 0, \\
 & \text{even}_S : \prod n : \text{nat}. \text{odd } n \rightarrow \text{even } (Sn). \\
 I_0(\text{odd} : \text{nat} \rightarrow \text{Prop}) := & \\
 & \text{odd}_S : \prod n : \text{nat}. \text{even } n \rightarrow \text{odd } (Sn).
 \end{aligned}$$

An *environment* of CIC_0 is a set of declarations. We assume an implicit global environment E . The environment E is assumed to contain appropriate typing declarations for the logical primitives. A CIC_0 *context* is a list of declarations of the form $x : t$ with t a term of CIC_0 and x the declared CIC_0 variable. We assume the variables declared in a context are pairwise disjoint. We denote environments by $E, E', \text{etc.}$, and contexts by $\Gamma, \Gamma', \text{etc.}$ We write $\Gamma, x : \tau$ to denote the context Γ with $x : \tau$ appended. We denote the empty context by $\langle \rangle$. A type judgement of CIC_0 has the form $\Gamma \vdash t : \tau$ where Γ is a context and t, τ are terms. If $\Gamma \vdash t : \tau$ and $\Gamma \vdash \tau : \sigma$ then we write $\Gamma \vdash t : \tau : \sigma$. A Γ -*proposition* is a term t such that $\Gamma \vdash t : \text{Prop}$. A Γ -*proof* is a term t such that $\Gamma \vdash t : \tau : \text{Prop}$ for some term τ .

The set $\text{FV}(t)$ of free variables of a term t is defined in the usual way. To save on notation we sometimes treat $\text{FV}(t)$ as a list. For a context Γ which includes declarations of all free variables of t , the free variable context $\text{FC}(\Gamma; t)$ of t is defined inductively:

- $\text{FC}(\langle \rangle; t) = \langle \rangle$,
- $\text{FC}(\Gamma, x : \tau; t) = \text{FC}(\Gamma; \lambda x : \tau. t), x : \tau$ if $x \in \text{FV}(t)$,
- $\text{FC}(\Gamma, x : \tau; t) = \text{FC}(\Gamma; t)$ if $x \notin \text{FV}(t)$.

If Γ includes declarations of all variables from a set of variables V , then we define $\text{FF}_\Gamma(V)$ to be the set of those $y \in V$ which are not Γ -proofs. Again, to save on notation we sometimes treat $\text{FF}_\Gamma(V)$ as a list.

Our translation encodes CIC_0 in untyped first-order logic with equality (FOL). We also implemented a straightforward information-forgetting export of Coq declarations into the syntax of CIC_0 . We describe the translation and the export in the next section.

In the translation of CIC_0 we need to perform (approximate) type checking to determine which terms are propositions (have type `Prop`), i.e. we need to check whether a given term t in a given context Γ has type `Prop`. For this purpose we implemented a specialised efficient procedure to do so. In fact, this procedure is slightly incomplete. The point here is to approximately identify which types are intended to represent propositions. In proof assistants or proof developments where types other than those of sort `Prop` are intended to represent propositions the procedure needs to be changed.

All CIC_0 terms we are interested in correspond to typable (and thus strongly normalizing) Coq terms, i.e., Coq terms are exported in a simple information-forgetting way to appropriate CIC_0 terms. We will assume that for any exported term there exists a type in logic of Coq, it is unique, and it is preserved under context extension. This assumption is not completely theoretically justified, but is useful in practice.

4 Premise Selection

The first component of a hammer preselects a subset of the accessible facts most likely to be useful in proving the user given goal. In this section we present the premise selection

algorithm proposed for a hammer for dependently typed theory. We reuse the two most successful filters used in HOLyHammer [52] and Sledgehammer [15] adapted to the CIC₀ representation of proof assistant knowledge. We first discuss the features and labels useful for that representation and further describe the k -NN and naive Bayes classifiers, which we used in our implementation.

4.1 Features and Labels

A simple possible characterization of statements in a proof assistant library is to use the sets of symbols that appear in these statements. It is possible to extend this set in many ways [56], including various kinds of structure of the statements, types, and normalizing variables (all variables will be replaced by a single symbol X). In the case of CIC₀, the constants are already both term constants and type constructors. We omit the basic logical constants, as they will not be useful for automated theorem provers which assume first-order logic. We further augment the set of features by inspecting the parse tree: constants and constant-variable pairs that share an edge in the parse tree give rise to a feature of the statement. We will denote such features of a theorem T by $F(T)$.

For each feature f we additionally compute a feature weight $w(f)$ that estimates the importance of the feature. Based on the HOLyHammer experiments with feature weights [54], we use TF-IDF [47] to compute feature weights. This ensures that rare features are more important than common ones.

Like in usual premise selection, the dependencies of theorems will constitute the labels for the learning algorithms. The dependencies for a theorem or definition T , which we will denote $D(T)$, are the constants occurring in the type of T or in the proof term (or the unfolding) of T . Note that these dependencies may not be complete, because in principle an ATP proof of T may need some additional information that in Coq is incorporated into type-checking but not used to build proof terms, e.g. definitions of constants, facts which are necessary to establish types of certain terms.

For example, consider the theorem $T = \text{Between.between_le}$ from the Coq standard library with the statement:

```
forall k l, between k l -> k <= l.
```

In the section where this theorem is declared there is the following variable declaration:

```
Variable P : nat -> Prop.
```

The features and dependencies of T are:

```
F(T) = {"Between.Between.between", "Between.Between.between-X",
        "Coq.Init.Datatypes.nat", "Coq.Init.Peano.le",
        "Coq.Init.Peano.le-X"}
```

```
D(T) = {"Between.Between.between", "Between.Between.between_ind",
        "Coq.Init.Datatypes.nat", "Coq.Init.Peano.le",
        "Coq.Init.Peano.le_S", "Coq.Init.Peano.le_n", "P"}
```

The $-X$ features correspond to constants applied to variables. Similarly, in more complex examples constant-constant applications (such as the successor of zero) give rise to such compound features.

4.2 *k*-Nearest Neighbors

The *k* nearest neighbors classifier (*k*-NN) finds a given number *k* of accessible facts which are most similar to the current goal. The distance for two statements *a*, *b* is defined by the function (higher values means more similar, τ_1 is a constant which gives more similar statements an additional advantage):

$$s(a, b) = \sum_{f \in F(a) \cap F(b)} w(f)^{\tau_1}$$

The dependencies of the selected facts will be used to estimate the relevance of all accessible facts. Given the set of the *k* nearest neighbors *N* together with their nearness values, the relevance of a visible fact *a* for the goal *g* is

$$\left(\tau_2 \sum_{b \in N | a \in D(b)} \frac{s(b, g)}{|D(b)|} \right) + \begin{cases} s(a, g) & \text{if } a \in N \\ 0 & \text{otherwise} \end{cases}$$

where τ_2 is a constant which gives more importance to the dependencies. We have used the values $\tau_1 = 6$ and $\tau_2 = 2.7$ in our implementation, which were found experimentally in our previous work [51].

There are two modifications of the standard *k*-NN algorithm. First, when deciding on the labels to predict based on the neighbors, we not only include the labels associated with the neighbors based on the training examples (this corresponds to past proofs) but also the neighbors themselves. This is because a theorem is in principle provable from itself in zero steps, and this information is not included in the training data. Furthermore, theorems that have been proved, but have not been used yet, would not be accessible to the algorithm without this modification.

Second, we do not use a fixed number *k*, instead we fix the number of facts with non-zero relevance that need to be predicted. We start with *k* = 1 and if not enough facts have been selected, we increase *k* iteratively. This allows creating ATP problems of proportionate complexity.

4.3 Sparse Naive Bayes

The sparse naive Bayes classifier estimates the relevance of a fact *a* for a goal *g* by the probability

$$P(a \text{ is used in the proof of } g)$$

Since the goal is only characterized by its features, the probability can be further estimated by:

$$P(a \text{ is used in a proof of } s \mid s \text{ has features } F(g))$$

where *s* is an arbitrary proved theorem, abstracting from the goal *g*.

For efficiency reasons the computation of the relevance of *a* is restricted to the features of *a* and the features that were ever present when *a* was used as a dependency. More formally, the *extended features* $\bar{F}(a)$ of *a* are:

$$\bar{F}(a) = F(a) \cup \bigcup_{a \in D(b)} F(b)$$

The probability can be thus estimated by the statements s which have the features $F(g)$ but do not have the features $\overline{F}(a) - F(g)$:

$$P(a \text{ is used in a proof of } s \mid F(a) \subseteq F(g) \wedge F(a) \text{ misses } \overline{F}(a) - F(g))$$

Assuming that the features are independent¹ the Bayes’s rule can be applied to transform the probability to the following product of probabilities:

$$\begin{aligned}
 &P(a \text{ is used in the proof of } s) \\
 &\cdot \prod_{f \in F(g) \cap \overline{F}(a)} P(s \text{ has feature } f \mid a \text{ is used in the proof of } s) \\
 &\cdot \prod_{f \in F(g) - \overline{F}(a)} P(s \text{ has feature } f \mid a \text{ is not used in the proof of } s) \\
 &\cdot \prod_{f \in \overline{F}(a) - F(g)} P(s \text{ does not have feature } f \mid a \text{ is used in the proof of } s)
 \end{aligned}$$

The expressions can be finally estimated:

$$\begin{aligned}
 P(a \text{ is used in a proof of } s) &= \frac{t(a)}{K} \\
 P(s \text{ has feature } f \mid a \text{ is used in the proof of } s) &= \frac{s(a, f)}{t(a)} \\
 P(s \text{ does not have feature } f \mid a \text{ is used in the proof of } s) &= 1 - \frac{s(a, f)}{t(a)}
 \end{aligned}$$

using two auxiliary functions that can be computed from the dependencies:

- $s(a, f)$ is the number of times a has been a dependency of a fact characterized by the feature f ;
- $t(a)$ is the number of times a has been a dependency;

as well as the number K of all theorems proved so far.

In our actual implementation we further introduce minor modifications to avoid any of the probabilities become zero and we estimate the logarithms of probabilities to avoid multiplying small numbers which might cause numerical instability. The classifier can finally estimate the relevance of all visible facts and return the requested number of them that are most likely to lead to a successful proof of the conjecture.

5 Translation

In this section we describe a translation of Coq goals through CIC_0 to untyped first-order logic with equality. The translation presented here is a significantly improved version of our translation presented at HaTT [24]. It has been made more complete, many optimisations have been introduced, and several mistakes have been eliminated.

The translation is neither sound nor complete. In particular, it assumes proof irrelevance (in the sense of erasing proof terms), it omits universe constraints on Type, and some information is lost in the export to CIC_0 . However, it is sound and complete “enough” to be practically

¹ There are many dependencies among the features, however considering such dependencies makes premise selection very slow and gives little improvement both when it comes to machine learning metrics and in practical hammer use [4].

usable by a hammer (just like the hammers for other systems, it works very well for essentially first-order logic goals and becomes much less effective with other features of the logics [17]). The limitations of the translation and further issues of the current approach are explained in more detail in Sects. 5.6 and 9. Some similar issues were handled in the context of code extraction in [60].

The translation proceeds in three phases. First, we export Coq goals to CIC_0 . Next we translate CIC_0 to first-order logic with equality. In the first-order language we assume a unary predicate P , a binary predicate T and a binary function symbol $@$. Usually, we write ts instead of $@(t, s)$. Intuitively, an atom of the form $P(t)$ asserts the provability of t , and $T(t, \tau)$ asserts that t has type τ . In the third phase we perform some optimisations on the generated FOL problem, e.g. replacing some terms of the form $P(cts)$ with $c(t, s)$.

A FOL *axiom* is a pair of a FOL formula and a constant (label). We translate CIC_0 to a set of FOL axioms. The labels are used to indicate which axioms are translations of which lemmas. When we do not mention the label of an axiom, then the label is not important.

5.1 Export of Coq data

The Coq declarations are exported in a straightforward way, translating Coq terms to corresponding terms of CIC_0 , possibly forgetting some information like e.g. universe constraints on Type. We implemented a Coq kernel plugin which exports the Coq kernel data structures. We briefly comment on several aspects of the export.

- Definitions are exported as CIC_0 definitions.
- Axioms are exported as CIC_0 typing declarations.
- Free variables (e.g. current hypotheses or variables from a currently open section) are exported as CIC_0 constants with appropriate typing declarations.
- Inductive types are exported as CIC_0 inductive declarations. Induction principles and recursor definitions are exported as separate CIC_0 definitions.
- Coinductive types are treated in the same way as inductive types, except that no induction principles or recursor definitions are exported for them.
- Mutual inductive types are exported separately for each constituent inductive type. See Sect. 3.
- The Coq construct `cofix` is exported to `fix` in CIC_0 with a special flag that affects the evaluation algorithm. We omitted this flag from the description of CIC_0 for the sake of simplicity.
- Modules and functors are not exported. Objects inside a module are exported with the name of the module prefixed to the name of the object.
- Universe constraints on Type are not exported. Proofs of paradoxes present in the standard library, e.g., Hurken's paradox, are explicitly filtered out and not exported.
- The following objects from the `Init.Logic` module are represented directly by the corresponding logical primitives of CIC_0 : `True`, `False`, `all`, `ex`, `and`, `or`, `iff`, `eq`. No other objects from the `Init.Logic` module are exported.
- Records are translated to inductive types already by Coq. Primitive record projections are not supported by our plugin.
- Existential metavariables are not exported. Currently it is not possible to use the hammer plugin when the proof state contains some uninstantiated existential metavariables.

The limitations of the translation, including these stemming from the incompleteness of the export as well as of the current architecture will be discussed in Sects. 5.6 and 9.

5.2 Translating Terms

The terms of CIC_0 are translated using three mutually recursively defined functions \mathcal{F} , \mathcal{G} and \mathcal{C} . The function \mathcal{F} encodes propositions as FOL formulas and is used for terms of CIC_0 having type Prop, i.e., for propositions of CIC_0 . The function \mathcal{G} encodes types as guards and is used for terms of CIC_0 which have type Type but not Prop. The function \mathcal{C} encodes CIC_0 terms as FOL terms. During the translation we add some fresh constants together with axioms (in FOL) specifying their meaning. Hence, strictly speaking, the codomain of each of the functions \mathcal{F} , \mathcal{G} and \mathcal{C} is the Cartesian product of the set of FOL formulas (or terms)—the desired encoding—and the powerset of the set of FOL formulas—the set of axioms added during the translation. However, it is more readable to describe the functions assuming a global mutable collection of FOL axioms.

Our translation assumes proof irrelevance. We use a fresh constant `prf` to represent an arbitrary proof object (of any inhabited proposition). For the sake of efficiency, CIC_0 propositions are translated directly to FOL formulas using the \mathcal{F} function. The CIC_0 types which are not propositions are translated to guards which essentially specify what it means for an object to have the given type. The formula $\mathcal{G}(t, \alpha)$ intuitively means “ t has type α ”. For instance, for a (closed) type $\tau = \Pi x : \alpha. \beta$ we have

$$\mathcal{G}(f, \tau) = \forall x. \mathcal{G}(x, \alpha) \rightarrow \mathcal{G}(fx, \beta)$$

So $\mathcal{G}(f, \tau)$ says that an object f has type $\tau = \Pi x : \alpha. \beta$ if for any object x of type α , the application fx has type β (in which x may occur free).

Below we give definitions of the functions \mathcal{F} , \mathcal{G} and \mathcal{C} . These functions are in fact parameterised by a CIC_0 context Γ , which we write as a subscript. In the description of the functions we implicitly assume that variable names are chosen appropriately so that no unexpected variable capture occurs. Also we assume an implicit global environment E . This environment is used for type checking. The typing declarations for CIC_0 logical primitives, as described in the previous section, are assumed to be present in E . During the translation also some new declarations are added to the environment. We assume all CIC_0 constants are also FOL constants, and analogously for variables. We use the notation $t_1 \approx_\Gamma t_2$ for $t_1 \leftrightarrow t_2$ if $\Gamma \vdash t_1 : \text{Prop}$, or for $t_1 = t_2$ if $\Gamma \not\vdash t_1 : \text{Prop}$.

The function \mathcal{F} encoding propositions as FOL formulas:

- If $\Gamma \vdash t : \text{Prop}$ then $\mathcal{F}_\Gamma(\Pi x : t. s) = \mathcal{F}_\Gamma(t) \rightarrow \mathcal{F}_{\Gamma, x:t}(s)$.
- If $\Gamma \not\vdash t : \text{Prop}$ then $\mathcal{F}_\Gamma(\Pi x : t. s) = \forall x. \mathcal{G}_\Gamma(x, t) \rightarrow \mathcal{F}_{\Gamma, x:t}(s)$.
- $\mathcal{F}_\Gamma(\forall x : t. s) = \forall x. \mathcal{G}_\Gamma(x, t) \rightarrow \mathcal{F}_{\Gamma, x:t}(s)$.
- $\mathcal{F}_\Gamma(\exists x : t. s) = \exists x. \mathcal{G}_\Gamma(x, t) \wedge \mathcal{F}_{\Gamma, x:t}(s)$.
- $\mathcal{F}_\Gamma(t \circ s) = \mathcal{F}_\Gamma(t) \circ \mathcal{F}_\Gamma(s)$ where $\circ \in \{\wedge, \vee, \leftrightarrow\}$.
- $\mathcal{F}_\Gamma(\neg t) = \neg \mathcal{F}_\Gamma(t)$.
- $\mathcal{F}_\Gamma(t = s) = (\mathcal{C}_\Gamma(t) = \mathcal{C}_\Gamma(s))$.
- Otherwise, if none of the above apply, $\mathcal{F}_\Gamma(t) = P(\mathcal{C}_\Gamma(t))$.

The function \mathcal{G} encoding types as guards:

- If $w = \Pi x : t. s$ and $\Gamma \vdash t : \text{Prop}$ then

$$\mathcal{G}_\Gamma(u, w) = \mathcal{F}_\Gamma(t) \rightarrow \mathcal{G}_{\Gamma, x:t}(u, s)$$

- If $w = \Pi x : t. s$ and $\Gamma \not\vdash t : \text{Prop}$ then $\mathcal{G}_\Gamma(u, w) = \forall x. \mathcal{G}_\Gamma(x, t) \rightarrow \mathcal{G}_{\Gamma, x:t}(ux, s)$.
- If w is not a product then $\mathcal{G}_\Gamma(u, w) = T(u, \mathcal{C}_\Gamma(w))$.

The function \mathcal{C} encoding terms as FOL terms:

- $\mathcal{C}_\Gamma(c) = c$ for a constant c ,
- $\mathcal{C}_\Gamma(x) = x$ for a variable x if x is not a Γ -proof,
- $\mathcal{C}_\Gamma(x) = \text{prf}$ for a variable x if x is a Γ -proof,
- $\mathcal{C}_\Gamma(ts)$ is equal to:
 - prf if $\mathcal{C}_\Gamma(t) = \text{prf}$,
 - $\mathcal{C}_\Gamma(t)$ if $\mathcal{C}_\Gamma(t) \neq \text{prf}$ but $\mathcal{C}_\Gamma(s) = \text{prf}$,
 - $\mathcal{C}_\Gamma(t)\mathcal{C}_\Gamma(s)$ otherwise.
- $\mathcal{C}_\Gamma(\Pi x : t.s) = R\vec{y}$ for a fresh constant F where $\vec{y} = \text{FF}_\Gamma(\text{FC}(\Gamma; \Pi x : t.s))$ and
 - if $\Gamma \vdash (\Pi x : t.s) : \text{Prop}$ then $\forall \vec{y}. P(F\vec{y}) \leftrightarrow \mathcal{F}_\Gamma(\Pi x : t.s)$ is a new axiom,
 - if $\Gamma \not\vdash (\Pi x : t.s) : \text{Prop}$ then $\forall \vec{y}z. T(z, F\vec{y}) \leftrightarrow \mathcal{G}_\Gamma(z, \Pi x : t.s)$ is a new axiom.
- $\mathcal{C}_\Gamma(\lambda \vec{x} : \vec{\tau}.t) = F\vec{y}_0$ for a fresh constant F where
 - t does not start with a lambda-abstraction any more,
 - $\Gamma, \vec{x} : \vec{\tau} \vdash t : \alpha$,
 - $\vec{y} : \vec{\rho} = \text{FC}(\Gamma; \lambda \vec{x} : \vec{\tau}.t)$,
 - $\vec{y}_0 = \text{FF}_\Gamma(\vec{y})$ and $\vec{x}_0 = \text{FF}_{\Gamma, \vec{x} : \vec{\tau}}(\vec{x})$,
 - the typing declaration $F : \Pi \vec{y} : \vec{\rho}. \Pi \vec{x} : \vec{\tau}. \alpha$ is added to the global environment E (before the recursive call to \mathcal{F}_Γ below),
 - the following is a new axiom:

$$\forall \vec{y}_0 \vec{x}_0. \mathcal{F}_{\Gamma, \vec{x} : \vec{\tau}}(F\vec{y}\vec{x} \approx_{\Gamma, \vec{x} : \vec{\tau}} t).$$

Note that the call to \mathcal{F} will remove those variable arguments to F which are $\Gamma, \vec{x} : \vec{\tau}$ -proofs. Hence, ultimately F will occur as $F\vec{y}_0\vec{x}_0$ in the above axiom.

- If t is a Γ -proof then

$$\mathcal{C}_\Gamma(\text{case}(t, c, n, \lambda \vec{a} : \vec{\alpha}. \lambda x : c\vec{p}\vec{a}. \tau, \lambda \vec{x}_1 : \vec{\tau}_1.s_1, \dots, \lambda \vec{x}_k : \vec{\tau}_k.s_k)) = C$$

for a fresh constant C .

- If t is not a Γ -proof then

$$\mathcal{C}_\Gamma(\text{case}(t, c, n, \lambda \vec{a} : \vec{\alpha}. \lambda x : c\vec{p}\vec{a}. \tau, \lambda \vec{x}_1 : \vec{\tau}_1.s_1, \dots, \lambda \vec{x}_k : \vec{\tau}_k.s_k)) = F\vec{y}_0$$

for a fresh constant F where

- $I(c : \gamma := c_1 : \gamma_1, \dots, c_k : \gamma_k) \in E$,
- $\vec{y} : \vec{\rho} = \text{FC}(\Gamma; \text{case}(t, c, n, \lambda \vec{a} : \vec{\alpha}. \lambda x : c\vec{p}\vec{a}. \tau, \lambda \vec{x}_1 : \vec{\tau}_1.s_1, \dots, \lambda \vec{x}_k : \vec{\tau}_k.s_k))$,
- $\vec{y}_0 = \text{FF}_\Gamma(\vec{y})$,
- $\vec{y}_1 : \vec{\rho}_1 = \text{FC}(\Gamma; t)$,
- $\Gamma \vdash t : c\vec{p}\vec{u}$ for some terms \vec{u} ,
- the declaration $F : \Pi \vec{y} : \vec{\rho}. \tau[\vec{u}/\vec{a}, t/x]$ is added to the global environment E ,
- the following is a new axiom:

$$\begin{aligned} \forall \vec{y}_0. \text{guards}_{\vec{y}_1 : \vec{\rho}_1}(\mathcal{F}_\Gamma ((\exists \vec{x}_1 : \vec{\tau}_1. t = c_1\vec{p}\vec{x}_1 \wedge F\vec{y} \approx_{\Gamma, \vec{x}_1 : \vec{\tau}_1} s_1) \\ \vee \dots \\ \vee (\exists \vec{x}_k : \vec{\tau}_k. t = c_k\vec{p}\vec{x}_k \wedge F\vec{y} \approx_{\Gamma, \vec{x}_k : \vec{\tau}_k} s_k))) \end{aligned}$$

where for a FOL formula φ and a context Γ we define $\text{guards}_\Gamma(\varphi)$ inductively as follows:

- $\text{guards}_{\emptyset}(\varphi) = \varphi$,
- $\text{guards}_{\Gamma, x : \tau}(\varphi) = \text{guards}_\Gamma(\mathcal{F}_\Gamma(\tau) \rightarrow \varphi)$ if $\Gamma \vdash \tau : \text{Prop}$,

- $\text{guards}_{\Gamma, x: \tau}(\varphi) = \text{guards}_{\Gamma}(\mathcal{G}_{\Gamma}(x, \tau) \rightarrow \varphi)$ if $\Gamma \not\vdash \tau : \text{Prop}$.
- $\mathcal{C}_{\Gamma}(\text{fix}(f_j, f_1 : \tau_1 := t_1, \dots, f_n : \tau_n := t_n)) = F_j \vec{y}_0$ where
 - $\vec{y} : \vec{\alpha} = \text{FC}(\Gamma; \text{fix}(f_j, f_1 : \tau_1 := t_1, \dots, f_n : \tau_n := t_n))$,
 - $\vec{y}_0 = \text{FF}_{\Gamma}(\vec{y})$,
 - F_1, \dots, F_n are fresh constants,
 - for $i = 1, \dots, n$ the typing declarations $F_i : \Pi \vec{y} : \vec{\alpha}. \tau_i$ are added to the global environment E ,
 - for $i = 1, \dots, n$ the following are new axioms:

$$\forall \vec{y}_0. \mathcal{F}_{\Gamma}(F_i \vec{y} \approx_{\Gamma} t_i [F_1 \vec{y} / f_1, \dots, F_n \vec{y} / f_n]).$$
- $\mathcal{C}_{\Gamma}(\text{let}(x : \tau := t, s)) = \mathcal{C}_{\Gamma}(s [F \vec{y}_0 / x])$ for a fresh constant F where
 - $\vec{y} : \vec{\alpha} = \text{FC}(\Gamma; t \tau)$,
 - $\vec{y}_0 = \text{FF}_{\Gamma}(\vec{y})$,
 - $\sigma = \Pi \vec{y} : \vec{\alpha}. \tau$,
 - the definition $F = (\lambda \vec{y} : \vec{\alpha}. t) : \sigma$ is added to the global environment E (before the recursive call to \mathcal{C}_{Γ} above),
 - if $\not\vdash \sigma : \text{Prop}$ then $\forall \vec{y}_0. F \vec{y}_0 = \mathcal{C}_{\Gamma}(t)$ is a new axiom.
- $\mathcal{C}_{\Gamma}(\text{cast}(\text{prf}, \tau)) = \text{prf}$.
- If $t \neq \text{prf}$ then $\mathcal{C}_{\Gamma}(\text{cast}(t, \tau)) = F \vec{y}_0$ for a fresh constant F where
 - $\vec{y} : \vec{\alpha} = \text{FC}(\Gamma; t \tau)$,
 - $\vec{y}_0 = \text{FF}_{\Gamma}(\vec{y})$,
 - $\sigma = \Pi \vec{y} : \vec{\alpha}. \tau$,
 - the definition $F = (\lambda \vec{y} : \vec{\alpha}. t) : \sigma$ is added to the global environment E ,
 - if $\not\vdash \sigma : \text{Prop}$ then $\forall \vec{y}_0. F \vec{y}_0 = \mathcal{C}_{\Gamma}(t)$ is a new axiom.

Example 1 A CIC_0 proposition

$$t = \Pi x : N. \Pi f : \alpha \rightarrow N \rightarrow N. \Pi q : \alpha. f q x = x$$

in the context

$$\Gamma = N : \text{Type}, \alpha : \text{Prop}$$

is translated to

$$\mathcal{F}_{\Gamma}(t) = \forall x. T(x, N) \rightarrow \forall f. (P(\alpha) \rightarrow \forall y. T(y, N) \rightarrow T(fy, N)) \rightarrow P(\alpha) \rightarrow fx = x.$$

In practice, checking the conditions $\Gamma \vdash t : \text{Prop}$ is performed by our specialised approximate proposition-checking algorithm. Checking whether a term t is a Γ -proof occurs in two cases.

1. t is the term matched on in a `case`-expression `case(t, c, ...)`. Then there is an inductive declaration $I_n(c : \gamma := \dots)$ in the global environment. We check if the normal form of γ has target `Prop`.
2. $t = x$ is a variable. Then we check if the type assigned to x by the context Γ is a proposition.

We write $\varphi(\sigma)$ to denote that a FOL formula φ has σ as a subformula. Then $\varphi(\sigma')$ denotes the formula φ with σ replaced by σ' . We use an analogous notation when σ is a FOL term instead of a formula.

Note that each new axiom defining a constant F intended to replace (“lift-out”) a λ -abstraction, a case expression or a fixpoint definition has the form

$$\forall \vec{x}. \varphi(F\vec{x} = t)$$

or

$$\forall \vec{x}. \varphi(P(F\vec{x}) \leftrightarrow \psi).$$

We will call each such axiom the *lifting axiom for F* . For lambda abstractions, this is equivalent to lambda-lifting, which is a common technique used by hammers for HOL and Mizar. In CIC_0 however other kinds of terms do bind variables (for example `case` and `fix`) and lifting axioms need to be created for such terms as well.

5.3 Translating Declarations

Declarations of CIC_0 are encoded as FOL axioms. As before, a global CIC_0 environment E is assumed. During the translation of a declaration the functions \mathcal{F} , \mathcal{G} and \mathcal{C} from the previous subsection are used. These functions may themselves add some FOL axioms, which are then also included in the result of the translation of the declaration. We proceed to describe the translation for each of the three forms of CIC_0 declarations. Whenever we write \mathcal{F} , \mathcal{G} , \mathcal{C} without subscript, the empty context $\langle \rangle$ is assumed as the subscript.

A definition $c = t : \tau$ is translated as follows.

- If $\vdash \tau : \text{Prop}$ then add $\mathcal{F}(\tau)$ as a new axiom with label c .
- If $\not\vdash \tau : \text{Prop}$ then
 - add $\mathcal{G}(c, \tau)$ as a new axiom,
 - if $\tau = \text{Prop}$ then add $c \leftrightarrow \mathcal{F}(t)$ as a new axiom with label c ,
 - if $\tau = \text{Set}$ or $\tau = \text{Type}$ then add $\forall f. cf \leftrightarrow \mathcal{G}(f, t)$ as a new axiom with label c ,
 - if $\tau \notin \{\text{Prop}, \text{Set}, \text{Type}\}$ then add $c = \mathcal{C}(t)$ as a new axiom with label c .

A typing declaration $c : \tau$ is translated as follows.

- If $\vdash \tau : \text{Prop}$ then add $\mathcal{F}(\tau)$ as a new axiom with label c .
- If $\not\vdash \tau : \text{Prop}$ then add $\mathcal{G}(c, \tau)$ as a new axiom with label c .

An inductive declaration $I(c : \tau := c_1 : \tau_1, \dots, c_n : \tau_n)$ is translated as follows, where $\tau \Downarrow \Pi \vec{p} : \vec{\beta}. \Pi \vec{y} : \vec{\gamma}. s$ and $s \in \{\text{Prop}, \text{Set}, \text{Type}\}$ and $\vec{\beta}$ are the types of the parameters of the inductive type and $\tau_i \Downarrow \Pi \vec{p} : \vec{\beta}. \Pi \vec{x}_i : \vec{\alpha}_i. c \vec{p} \vec{t}_i$ and the length of \vec{y} and each \vec{t}_i is m .

- Translate the typing declaration $c : \tau$.
- Translate each typing declaration $c_i : \tau_i$ for $i = 1, \dots, n$.
- If $s \neq \text{Prop}$ then for each $i = 1, \dots, n$ add the following injectivity axiom:

$$\mathcal{F}(\forall \vec{x}_i : \vec{\alpha}_i. \forall \vec{x}'_i : \vec{\alpha}'_i. c_i \vec{x}_i = c_i \vec{x}'_i \rightarrow x_{i,1} = x'_{i,1} \wedge \dots \wedge x_{i,k_i} = x'_{i,k_i})$$

where $\vec{\alpha}'_i = \vec{\alpha}_i[\vec{x}'_i / \vec{x}_i]$.

- If $s \neq \text{Prop}$ then for each $i, j = 1, \dots, n$ with $i \neq j$ add the following discrimination axiom:

$$\mathcal{F}(\forall \vec{x}_i : \vec{\alpha}_i. \forall \vec{x}_j : \vec{\alpha}_j. c_i \vec{x}_i \neq c_j \vec{x}_j).$$

- If $s \neq \text{Prop}$ then add the following inversion axiom:

$$\begin{aligned} \mathcal{F}(\forall \vec{p} : \vec{\beta}. \forall \vec{y} : \vec{\gamma}. \forall z : c \vec{p} \vec{y} . & (\exists \vec{x}_1 : \vec{\alpha}_1. z = c_1 \vec{p} \vec{x}_1 \wedge y_1 = t_{1,1} \wedge \dots \wedge y_m = t_{1,m}) \\ & \vee \dots \\ & \vee (\exists \vec{x}_n : \vec{\alpha}_n. z = c_n \vec{p} \vec{x}_n \wedge y_1 = t_{n,1} \wedge \dots \wedge y_m = t_{n,m})). \end{aligned}$$

- If $s = \text{Prop}$ then add the following inversion axiom:

$$\mathcal{F}(\forall \vec{p} : \vec{\beta}. \forall \vec{y} : \vec{\gamma}. c \vec{p} \vec{y} \rightarrow ((\exists \vec{x}_1 : \vec{\alpha}_1. y_1 = t_{1,1} \wedge \dots \wedge y_m = t_{1,m}) \vee \dots \vee (\exists \vec{x}_n : \vec{\alpha}_n. y_1 = t_{n,1} \wedge \dots \wedge y_m = t_{n,m})))$$

5.4 Translating Problems

A CIC_0 problem consists of a set of assumptions which are CIC_0 declarations, and a conjecture which is a CIC_0 proposition. A CIC_0 problem is translated to a FOL problem by translating the assumptions to FOL axioms in the way described in the previous subsection, and translating the conjecture t to a FOL conjecture $\mathcal{F}(t)$. New declarations added to the environment during the translation are *not* translated. For every CIC_0 problem the following FOL axioms are added to the result of the translation:

- $T(\text{Prop}, \text{Type}), T(\text{Set}, \text{Type}), T(\text{Type}, \text{Type}),$
- $\forall y. T(y, \text{Set}) \rightarrow T(y, \text{Type}).$

5.5 Optimisations

We perform the following optimisations on the generated FOL problems, in the given order. Below, by an occurrence of a term t (in the FOL problem) we mean an occurrence of t in the set of FOL formulas comprising the given FOL problem.

- We recursively simplify the lifting axioms for the constants encoding λ -abstractions, case expressions and fixpoint definitions. For any lifting axiom A for a constant F , if A has the form

$$\forall \vec{x}. \varphi(F \vec{x} = G \vec{x})$$

such that G has a lifting axiom B

$$\forall \vec{x} \forall \vec{y}. \psi(G \vec{x} \vec{y} = t)$$

and either $\varphi(\square) = \square$ or \vec{y} is empty, then we replace the axiom A by

$$\forall \vec{x}. \varphi(\forall \vec{y}. \psi(F \vec{x} \vec{y} = t))$$

and we remove the axiom B and replace all occurrences of G by F . When in the lifting axioms A and B we have logical equivalence \leftrightarrow instead of equality $=$, then we adjust the replacement of A appropriately, using \leftrightarrow instead of $=$. We repeat applying this optimisation as long as possible.

- For a constant c , we replace any occurrence of $T(s, ct_1 \dots t_n)$ by $c_T(t_1, \dots, t_n, s)$ where c_T is a new function symbol of arity $n + 1$. We then also add a new axiom:

$$\forall x_1 \dots x_n y. c_T(x_1, \dots, x_n, y) \leftrightarrow T(y, cx_1 \dots x_n).$$

Note that after performing this replacement the predicate T may still occur in the FOL problem, e.g., a term $T(s, xt_1 \dots t_n)$ may occur. This optimisation is useful, because it simplifies the FOL terms and replaces the T predicate with a specialised predicate for a constant. This makes it easier for the ATPs to handle the problem.

- For each occurrence of a constant c with $n > 0$ arguments, i.e., each occurrence $ct_1 \dots t_n$ where $n > 0$ is maximal (there are no further arguments), we replace this occurrence with $c^n(t_1, \dots, t_n)$ where c^n is a new n -ary function symbol. We then also add a new axiom:

- $\forall x_1 \dots x_n. P(c^n(x_1, \dots, x_n)) \leftrightarrow P(cx_1 \dots x_n)$ if (after replacement of all such occurrences) all terms of the form $c^n(t_1, \dots, t_n)$ occur only as arguments of the predicate P , i.e., occur only in $P(c^n(t_1, \dots, t_n))$.
- $\forall x_1 \dots x_n. c^n(x_1, \dots, x_n) = cx_1 \dots x_n$ otherwise.

This optimisation is similar to the optimisation originally described by Meng and Paulson in [61, Section 2.7].

- For any constant c and $n > 0$, if all terms of the form $c^n(t_1, \dots, t_n)$ occur only as arguments of P , then replace each occurrence of a term of the form $P(c^n(t_1, \dots, t_n))$ by $c^n(t_1, \dots, t_n)$.

5.6 Properties of the Translation

In this section we briefly comment on the theoretical aspects of the translation. Further limitations of the whole approach will be mentioned in Sect. 9. The translation is neither sound nor complete. The lack of soundness is caused e.g. by the fact that we forget universe constraints on Type, the assumption of proof irrelevance, and the combination of omitting type guards for lifted-out lambda-abstractions with translating Coq equality to FOL equality. However, our experimental evaluation indicates that the translation is both sound and complete “enough” to be practically usable. Also, a “core” version of our translation is sound. A soundness proof and a more detailed discussion of the theoretical properties of a core version of our translation may be found in [27].

Note that e.g. in the axiom added for lifted-out lambda-abstractions

$$\forall \vec{y}_0 \vec{x}_0. \mathcal{F}_{\Gamma, \vec{x}; \vec{\tau}}(F \vec{y} \vec{x} \approx_{\Gamma, \vec{x}; \vec{\tau}} t)$$

we do not generate type guards for the free (\vec{y}_0) or bound (\vec{x}_0) variables of the lambda-expression. In practice, omitting these guards slightly improves the success rate of the ATPs without significantly affecting the reconstruction success rate. We conjecture that, ignoring other unsound features of the translation, omitting these guards is sound provided that the inductive Coq equality type `eq` is *not* translated to FOL equality. Note also that it is not sound (and our translation does not do it) to omit guards for the free variables of the term matched on in the `case` construct, even if Coq equality is not translated to FOL equality. For example, assume $I_0(c : \text{Set} := c_0 : c)$ is in the global environment. With the guards omitted, for the `case`-expression `case(x, c, 0, c, c_0)` we would add an axiom

$$\forall x. x = c_0 \wedge Fx = c_0$$

with F a fresh first-order constant. This obviously leads to an inconsistency by substituting for x two distinct constants c_1, c_2 such that $c_1 \neq c_2$ is provable.

In our translation we map Coq equality to FOL equality which is not sound in combination with omitting the guards for free variables. In particular, if a CIC₀ problem contains a functional extensionality axiom then the generated FOL problem may be inconsistent, and in contrast to the inconsistencies that may result from omitting certain universe constraints, this inconsistency may be “easy enough” for the ATPs to derive. Our plugin has an option to turn on guard generation for free variables. See also [27, Section 6].

6 Proof Reconstruction

In this section we will discuss a number of existing Coq internal automation mechanisms that could be useful for proof reconstruction and finally introduce our combined proof reconstruction tactic.

The tactic `firstorder` is based on an extension of the contraction-free sequent calculus LJT of Dyckhoff [32] to first-order intuitionistic logic with inductive definitions [26]. A decision procedure for intuitionistic propositional logic based on the system LJT is implemented in the tactic `tauto`. The tactic `firstorder` does not take into account many features of Coq outside of first-order logic. In particular, it does not fully axiomatise equality.

In general, the tactics based on extensions of LJT do mostly forward reasoning, i.e., they predominantly manipulate the hypotheses in the context to finally obtain the goal. Our approach is based more on an `auto`-type proof search which does mostly backward Prolog-style reasoning—modifying the goal by applying hypotheses from the context. The core of our search procedure may be seen as an extension of the Ben-Yelles algorithm [21,42] to first-order intuitionistic logic with all connectives [71,75]. It is closely related to searching for η -long normal forms [12,31]. Our implementation extends this core idea with various heuristics. We augment the proof search procedure with the use of existential metavariables like in `eauto`, a looping check, some limited forward reasoning, the use of the `congruence` tactic, and heuristic rewriting using equational hypotheses.

It is important to note that while the external ATPs we employ are classical and the translation assumes proof irrelevance, the proof reconstruction phase does not assume any additional axioms. We re-prove the theorems in the intuitionistic logic of Coq, effectively using the output of the ATPs merely as hints for our hand-crafted proof search procedure. Therefore, if the ATP proof is inherently classical then proof reconstruction will fail. Currently, the only information from ATP runs we use is a list of lemmas needed by the ATP to prove the theorem (these are added to the context) and a list of constant definitions used in the ATP proof (we try unfolding these constants and no others).

Another thing to note is that we do not use the information contained in the Coq standard library during reconstruction. This would not make sense for our evaluation of the reconstruction mechanism, since we try to re-prove the theorems from the Coq standard library. In particular, we do not use any preexisting hint databases available in Coq, not even the core database (for the evaluation we use the `auto` and `eauto` tactics with the `nocore` option, but in the final version of the reconstruction tactics we also use `auto` without this option). Also, we do not use any domain-specific decision procedures available as Coq tactics, e.g., `field`, `ring` or `omega`. Including such techniques in HOLyHammer did allow fast solving of many simple arithmetic problems [53].

We now describe a simplification of our proof search procedure. We will treat the current proof state as a collection of judgements of the form $\Gamma \vdash G$ and describe the rules as manipulating a single such judgement. In a judgement $\Gamma \vdash G$ the term G is the *goal* and Γ is the *context* which is a list of *hypothesis* declarations of the form $H : A$. We use an informal notation for Coq terms similar to how they are displayed by Coq. For instance, by $\forall x : A, B$ we denote a dependent product. We write $\forall x, B$ when the type of x is not essential. Note that in $\forall x, B$ the variable x may be a proposition, so $\forall x, B$ may actually represent a logical implication $A \rightarrow B$ if A is the omitted type of x which itself has type `Prop` and x does not occur in B . To avoid confusion with $=$ used to denote the equality inductive predicate in Coq, we use \equiv as a metalevel symbol to denote identity of Coq terms. We use the notation $\Gamma; H : A$ to denote Γ with $H : A$ inserted at some fixed position. By $\Gamma, H : A$ we denote the

context Γ with $H : A$ appended. We omit the hypothesis name H when irrelevant. By $C[t]$ we denote an occurrence of a term t in a term context C .

The proof search procedure applies the rules from Fig. 1. An application of a rule of the form

$$\frac{\Gamma_1 \vdash G_1 \quad \dots \quad \Gamma_n \vdash G_n}{\Gamma \vdash G}$$

replaces a judgement $\Gamma \vdash G$ in the current proof state by the judgements $\Gamma_1 \vdash G_1, \dots, \Gamma_n \vdash G_n$. The notation $\text{tac}[\Gamma \vdash G]$ (resp. $\text{tac}(A)[\Gamma \vdash G]$) in a rule premise means applying the Coq tactic tac (with argument A) to the judgement $\Gamma \vdash G$ and making the judgements (subgoals) generated by the tactic be the premises of the rule. In a rule of the form e.g.

$$\frac{\Gamma; A' \vdash G}{\Gamma; A \vdash G}$$

the position in Γ at which A is inserted is implicitly assumed to be the same as the position at which A' is inserted.

In Fig. 1 the variables $?e_i, ?e$ denote fresh existential metavariables of appropriate types. These metavariables need to be instantiated later by Coq’s unification algorithm. In the rules (orsplit) and (exsimpl) the types of x_1, \dots, x_n are assumed not to be propositions. In the rule (exinst) the types of x_1, \dots, x_k are not propositions and either $k = n$ or the type of x_{k+1} is a proposition. In the rule (orinst) the x_{i_1}, \dots, x_{i_m} are all those among x_1, \dots, x_n for which T_{i_1}, \dots, T_{i_m} are not propositions; and the index k ranges over all $k \in \{1, \dots, n\} \setminus \{i_1, \dots, i_m\}$ (so that each T_k is a proposition)—all judgements for any such k are premises of the rule, not just a single one. Moreover, in these rules for any term T by T' we denote $T[?e_{i_1}/x_{i_1}, \dots, ?e_{i_m}/x_{i_m}]$, and $T_{j_1}, \dots, T_{j_m:k}$ are those among T_1, \dots, T_k which are propositions. In the (apply) and (invert) rules P is an atomic proposition, i.e., a proposition which is not a dependent product, an existential, a disjunction or a conjunction. In the (destruct) rule T is not a proposition.

The tactic `yapply` in rule (apply) works like `eapply` except that instead of simply unifying the goal with the target of the hypothesis, it tries unification modulo some simple equational reasoning. The idea of the `yapply` tactic is broadly similar to the smart matching of Matita [8], but our implementation is more heuristic and not based on superposition.

The tactic `yrewrite` in rule (rewrite) uses Coq’s tactic `erewrite` to try to rewrite the hypothesis in the goal. If it fails to rewrite it directed from left to right, then it tries the other direction.

The rules in Fig. 1 are divided into groups. The rules in each group are either applied with backtracking (marked by (b) in the figure), i.e., if applying one of the rules in the group to a judgement $\Gamma \vdash G$ does not ultimately succeed in finishing the proof then another of the rules in the group is tried on $\Gamma \vdash G$; or they are applied eagerly without backtracking (marked by (e) in the figure). There are also restrictions on when the rules in a given group may be applied. The rules in the group “Leaf tactics” must close a proof tree branch, i.e., they are applied only when they generate zero premises. The rules in the group “Final splitting” are applied only before the “leaf tactics”. The rules in the groups “Splitting”, “Hypothesis simplification” and “Introduction” are applied whenever possible. The rules in the group “Proof search” constitute the main part of the proof search procedure. They are applied only when none of the rules in the groups “Splitting”, “Hypothesis simplification” and “Introduction” can be applied. The rules in the group “Initial proof search” may only be applied after an application of (intro) followed by some applications of the rules in the “Splitting” and “Hypothesis simplification”

$\frac{\text{eauto}[\Gamma \vdash G]}{\Gamma \vdash G}$	$\frac{\text{congruence}[\Gamma \vdash G]}{\Gamma \vdash G}$	$\frac{\text{constructor}[\Gamma \vdash G]}{\Gamma \vdash G}$	$\frac{\text{easy}[\Gamma \vdash G]}{\Gamma \vdash G}$
Leaf tactics (b)			
<hr/>			
$\frac{\Gamma; A[?e_1/x_1, \dots, ?e_n/x_n] \vdash G \quad \Gamma; B[?e_1/x_1, \dots, ?e_n/x_n] \vdash G}{\Gamma; \forall x_1 \dots x_n, A \vee B \vdash G}$ (orsplit)			
$\frac{\Gamma; A[?e_1/x_1, \dots, ?e_n/x_n] \vdash G \quad y \text{ fresh}}{\Gamma; \forall x_1 \dots x_n, \exists y, A \vdash G}$ (exsimpl)			
Final splitting (e)			
<hr/>			
$\frac{\text{destruct}(t)[\Gamma \vdash C[\text{match } t \text{ with } b]]}{\Gamma \vdash C[\text{match } t \text{ with } b]}$		$\frac{\text{destruct}(t)[\Gamma; C[\text{match } t \text{ with } b] \vdash G]}{\Gamma; C[\text{match } t \text{ with } b] \vdash G}$	
$\frac{\Gamma \vdash G_1 \quad \Gamma \vdash G_2}{\Gamma \vdash G_1 \wedge G_2}$			
Splitting (e)			
<hr/>			
$\frac{\Gamma; A; B \vdash G}{\Gamma; A; A \rightarrow B \vdash G}$		$\frac{\Gamma; \forall x_1 \dots x_n, A; \forall x_1 \dots x_n, B \vdash G}{\Gamma; \forall x_1 \dots x_n, A \wedge B \vdash G}$	
$\frac{\Gamma; \forall x_1 \dots x_n, A \rightarrow B \rightarrow C \vdash G}{\Gamma; \forall x_1 \dots x_n, A \wedge B \rightarrow C \vdash G}$		$\frac{\Gamma; \forall x_1 \dots x_n, A \rightarrow C; \forall x_1 \dots x_n, B \rightarrow C \vdash G}{\Gamma; \forall x_1 \dots x_n, A \vee B \rightarrow C \vdash G}$	
$\frac{\Gamma; A \vdash G \quad x \text{ fresh}}{\Gamma; \exists x, A \vdash G}$		$\frac{\Gamma; A \vdash G}{\Gamma; A; A \vdash G}$	
Hypothesis simplification (e)			
<hr/>			
$\frac{\Gamma; x : A \vdash B}{\Gamma \vdash \forall x : A, B}$ (intro)			
Introduction (e)			
<hr/>			
$\frac{}{\Gamma; \text{False} \vdash G}$		$\frac{\Gamma \vdash G[?e/x]}{\Gamma \vdash \exists x, G}$	
$\frac{}{\Gamma; G \vdash G}$		$\frac{\Gamma \vdash G_1}{\Gamma \vdash G_1 \vee G_2}$	
$\frac{}{\Gamma \vdash G_1 \vee G_2}$		$\frac{\Gamma \vdash G_2}{\Gamma \vdash G_1 \vee G_2}$	
$\frac{\text{yapply}(H)[\Gamma; H : A \vdash P]}{\Gamma; H : A \vdash P}$ (apply)		$\frac{\text{yrewrite}(H)[\Gamma; H : \forall x_1 \dots x_n, A = B \vdash G]}{\Gamma; H : \forall x_1 \dots x_n, A = B \vdash G}$ (rewrite)	
$\frac{\Gamma; \forall x_{k+1} \dots x_n, \exists x, A[?e_1/x_1, \dots, ?e_k/x_k] \vdash G}{\Gamma; \forall x_1 \dots x_k \dots x_n, \exists x, A \vdash G}$ (exinst)			
$\frac{\Gamma; T'_{j_1}, \dots, T'_{j_{m:n}}, A'_1 \vdash G \quad \Gamma; T'_{j_1}, \dots, T'_{j_{m:n}}, A'_2 \vdash G \quad \Gamma; T'_{j_1}, \dots, T'_{j_{m:k-1}} \vdash T'_k}{\Gamma; \forall (x_1 : T_1) \dots (x_n : T_n), A_1 \vee A_2 \vdash G}$ (orinst)			
Proof search (b)			
<hr/>			
$\frac{\Gamma; \forall x_1 \dots x_n, \text{False} \vdash \text{False}}{\Gamma; \forall x_1 \dots x_n, \text{False} \vdash G}$			
$\frac{\text{inversion}(H)[\Gamma; H : P \vdash G]}{\Gamma; H : P \vdash G}$ (invert)		$\frac{\text{destruct}(x)[\Gamma; x : T \vdash G]}{\Gamma; x : T \vdash G}$ (destruct)	
Initial proof search (b)			
<hr/>			

Fig. 1 Simplified proof search rules

groups. They are applied only if none of the rules in the groups “Splitting”, “Hypothesis simplification” and “Introduction” can be applied.

The above description is only a readable approximation of what is actually implemented. Some further heuristics are used and more complex restrictions are put on what rules may be applied when. In particular, some loop checking (checking whether a judgement repeats) is implemented, the number of times a hypothesis may be used for rewriting is limited, and we also use heuristic rewriting in hypotheses and heuristic instantiation of universal hypotheses. Some heuristics we use are inspired by the `crush` tactic of Adam Chlipala [23].

As mentioned before, our proof search procedure could be seen as an extension of a search for η -long normal forms for first-order intuitionistic logic using a Ben-Yelles-type algorithm [71, 75]. As such it would be complete for the fragment of type theory “corresponding to” first-order logic, barring two simplifications we introduced to make it more practical. For the sake of efficiency, we do not backtrack on instantiations of existential metavariables solved by unification, and the rules (`exinst`) and (`orinst`) are not general enough. These cause incompleteness even for the first-order fragment, but this incompleteness does not seem to matter much in practice. The usual reasons why proof reconstruction fails is that either the proof is inherently classical, too deep, or uses too much rewriting which cannot be easily handled by our rewriting heuristics. It is left for future work to integrate rewriting into our proof search procedure in a more principled way.

The proof reconstruction phase in the `hammer` tactic uses a number of tactics derived from the procedure described above, with different depth limits, a bit different heuristics and rule application restrictions; plus a few other tactics, including Coq’s `intuition`, `simpl`, `subst`, and heuristic constant unfolding. Various reconstruction tactics are tried in order with a time limit for each, until one of them succeeds (or none succeed—then the proof cannot be reconstructed).

It is important to note that no time limits are supposed to be present in the final proof scripts. The CoqHammer plugin shows which of the tactics succeeded, and the user is supposed to copy this tactic, replacing the `hammer` tactic invocation. The final reconstruction tactic does not rely on any time limits or make any calls to external ATPs. Its results are therefore completely reproducible on different machines, in contrast to the main `hammer` tactic itself.

7 Integrated Hammer and Evaluation

In this section we present the technique used to select the combination of strategies included in the integrated hammer and present an evaluation of the components as well as the final offered strategy.

The evaluation in this section will perform a push-button re-proving of Coq problems without using their proofs. In order for the evaluation of the system to be fair, we need ensure that no information from a proof is used in its re-proving, as well as that the actual strategy that is used by the whole system has been developed without the knowledge of the proofs being evaluated.

The system will be evaluated on the problems generated from all theorems in the Coq standard library of Coq version 8.5 (a version of the plugin works with Coq 8.6 and 8.7 as well). The problems were generated from the source code of the library, counting as theorems all definitions (introduced with any of `Lemma`, `Theorem`, `Corollary`, `Fact`, `Instance`, etc.) that were followed by the `Proof` keyword. The source code of the library was then modified to insert a hook to our hammer plugin after each `Proof` keyword. The

plugin tries to re-prove the theorem using the Coq theorems accessible at the point when the statement of the theorem is introduced, using the three phases of premise selection, ATP invocation and proof reconstruction as described above.

This simulates how a hammer would be used in the development of the Coq standard library. In particular, when trying to re-prove a given theorem we use only the objects accessible in the Coq kernel at the moment the theorem statement is encountered by Coq. Of course, neither the re-proved theorem itself nor any theorems or definitions that depend on it are used. The number of problems obtained by automatically analysing the Coq standard library source code in the way described above is 9276. This differs significantly from the number of problems reported in [24]. There the theorems in the Coq standard library were extracted from objects of type Prop in the Coq kernel. Because of how the Coq module system works, there may be many Coq kernel objects corresponding to one definition in a source file (this is the case e.g. when using the `Include` command).

Furthermore, the problems are divided in a training set consisting of about 10% of the problems in the standard library and a validation set containing the remaining 90% of the problems. The training set is used to find a set of complementary strategies. Just like for the hammers for higher-order logic based systems and for Mizar a single best combination of the premise-selection algorithm, number of selected premises, and ATP run for a longer time is much weaker than running a few such combinations even for a shorter time. Contrary to existing hammer constructions [52,55], we decided to include the reconstruction mechanism among the considered strategy parameters since generally reconstruction rates are lower and it could happen that proofs originating from a particular prover and number of premises would be too hard to reconstruct.

In our evaluation we used the following ATPs: E Prover version 1.9 [65], Vampire version 4.0 [57] and Z3 version 4.0 [28]. The evaluation was performed on a 48-core server with 2.2GHz AMD Opteron CPUs and 320GB RAM. Each problem was always assigned one CPU core. The two considered premise selection algorithms were asked for an ordering of premises, and all powers of two between 16 and 1024 were considered. Finally we considered both `firstorder` and `hrecon` reconstruction. Having evaluated all combinations of premise selection algorithms we ordered them in a greedy sequence: each following strategy is the one that adds most to the current selection of strategies. The first 14 strategies in the greedy sequence are presented in Table 1. The column “Solved” indicates the number of problems that were successfully solved by the given ATP with the given premise selection method and a given number of premises, and they could be reconstructed by the proof reconstruction procedure described in Sect. 6. The ATPs were run with a time limit of 30 s. The maximum time limit for a single reconstruction tactic was 10 s, depending on the tactic, as described in Sect. 6. No time limit was placed on the premise selection phase, however for goals with largest number of available premises the time does not exceed 0.5 s for either of the considered algorithms. The first strategy that includes `firstorder` appears only on twelfth position in the greedy sequence and is therefore not used as part of the hammer. We show cumulative success rates to display the progress in the greedy sequence.

The results of the hammer strategies including the premise selection are very good in comparison with the results on the dependencies. Evaluating the translation with `hrecon` reconstruction is presented in Table 2. The results are significantly worse, mainly for two reasons. First, some dependencies are missing due to our way of recording them which does not take into account the delta-conversion. Secondly, the dependencies in proof terms often were added by automated tactics and are difficult to use for the ATPs. It is sometimes easier for the ATPs to actually prove the theorem from other lemmas in the library than from the original dependencies.

Table 1 Success rates of the strategies on the training set in the greedy sequence order

Prover	Selection	Premises	Reconstruction	Solved (%)	Solved
Vampire	k-NN	1024	Hrecon	30.778	285
Z3	k-NN	128	Hrecon	37.473	347
E-Prover	k-NN	1024	Hrecon	39.741	368
Vampire	k-NN	64	Hrecon	40.929	379
Z3	n. Bayes	32	Hrecon	41.469	384
Z3	n. Bayes	512	Hrecon	42.009	389
Z3	n. Bayes	128	Hrecon	42.549	394
E-Prover	n. Bayes	256	Hrecon	43.089	399
Z3	n. Bayes	16	Hrecon	43.521	403
E-Prover	n. Bayes	1024	Hrecon	43.952	407
Vampire	n. Bayes	256	Hrecon	44.276	410
Z3	k-NN	64	Hrecon	44.492	412
Vampire	k-NN	512	Hrecon	44.708	414
E-Prover	k-NN	512	Firstorder	44.924	416
total				46.112	427

Table 2 Prover results on the dependencies

Prover	Solved (%)	Solved
Vampire	24.749	2292
Z3	23.961	2219
E-Prover	23.162	2145
Total	26.747	2477

Table 3 The success rate of of the combination of strategies on the validation set

Prover	Selection	Premises	Reconstruction	Solved (%)	Solved
Vampire	k-NN	1024	Hrecon	28.816	2673
E-Prover	k-NN	1024	Hrecon	25.593	2374
Vampire	k-NN	64	Hrecon	25.367	2353
Z3	n. Bayes	128	Hrecon	24.299	2254
Z3	k-NN	128	Hrecon	24.127	2238
Z3	n. Bayes	512	Hrecon	23.243	2156
Z3	n. Bayes	32	Hrecon	19.028	1765
E-Prover	n. Bayes	256	Hrecon	17.497	1623
Total				40.815	3786

Given the common hardware configuration of computers today, we consider as the integrated system a combination of eight complementary strategies. The final results of the hammer including reconstruction on the validation set are presented in Table 3.

8 Case Studies

The intended use of a hammer is to prove relatively simple goals using available lemmas. The main problem a hammer system tries to solve is that of finding appropriate lemmas in a large collection and combining them to prove the goal. The advantage of a hammer over specialised domain-specific tactics is that it is a general system not depending on any domain knowledge. The hammer plugin may use all currently accessible lemmas, which includes lemmas proven earlier in a given formalization, not only the lemmas from the standard library or other predefined libraries.

It sometimes happens that the ATPs find proofs with fewer dependencies than the proofs in the standard library. One example is the Coq lemma `isometric_rotation`:

```
Lemma isometric_rotation : forall x1 y1 x2 y2 theta : R,
  dist_euc x1 y1 x2 y2 =
  dist_euc (xr x1 y1 theta) (yr x1 y1 theta)
  (xr x2 y2 theta) (yr x2 y2 theta).
```

Its current proof in the Coq standard library uses 6 auxiliary facts and is performed using the following 7 line script:

```
unfold dist_euc; intros; apply Rsqr_inj;
  [ apply sqrt_positivity; apply Rplus_le_le_0_compat
    | apply sqrt_positivity; apply Rplus_le_le_0_compat
    | repeat rewrite Rsqr_sqrt;
      [ apply isometric_rotation_0
        | apply Rplus_le_le_0_compat
        | apply Rplus_le_le_0_compat ] ]; apply Rle_0_sqr
```

Multiple ATPs found a shorter proof which uses only two of the dependencies: the definition of euclidean distance and the lemma `isometric_rotation_0`. This suggests that the proof using the injectivity of square root is a detour, and indeed it is possible to write a much simpler valid Coq proof of the lemma using just the two facts used by the ATPs:

```
unfold dist_euc; intros;
  rewrite (isometric_rotation_0 _ _ _ _ theta); reflexivity.
```

The proof may also be reconstructed from the found dependencies inside Coq. This is also the case for all other examples presented in this section.

Also for some theorems the ATPs found proofs which use premises not present in the dependencies extracted from the proof of the theorems in the standard library. An example is the lemma `le_double` from `Reals.ArithProp`:

```
forall m n : nat, 2 * m <= 2 * n -> m <= n.
```

The proof of this lemma in the standard library uses 6 auxiliary lemmas and is performed by the following proof script (two lemmas not visible in the script were added by the tactic `prove_sup0`):

```
intros; apply INR_le.
assert (H1 := le_INR _ _ H).
do 2 rewrite mult_INR in H1.
apply Rmult_le_reg_l with (INR 2).
replace (INR 2) with 2; [ prove_sup0 | reflexivity ].
assumption.
```


ATPs found a proof of `le_double` using only 3 lemmas: `Arith.PeanoNat.Nat.le_0_1`, `Arith.Mult.mult_S_le_reg_1` and `Init.Peano.le_n`. None of these lemmas appear among the original dependencies.

Another example of hammer usage is a proof of the following fact:

```
forall m n k : nat, m * n + k = k + n * m.
```

This cannot be proven using the `omega` tactic because of the presence of multiplication. The tactic invocations `eauto` with `arith` or `firstorder` with `arith` do not work either. The hammer tool finds a proof using two lemmas from `Arith.PeanoNat.Nat:add_comm` and `mul_comm`.

A similar example is the goal

```
forall n : nat, 3 * 3 ^ n = 3 ^ (n + 1).
```

This goal cannot be solved using standard Coq tactics, including the tactic `omega`. Z3 with 128 preselected premises found a proof using the following lemmas from `Arith.PeanoNat.Nat:add_succ_r`, `le_0_1`, `pow_succ_r`, `add_0_r`. The proof may be reconstructed using `hexhaustive 0` or `hyelles 5` tactic invocations.

The next example of a goal solvable by the hammer involves operations on lists.

```
forall {A} (x : A) l1 l2 (P : A -> Prop),
  In x (l1 ++ l2) -> (forall y, In y l1 -> P y) ->
  (forall y, In y l2 -> P y) ->
  P x.
```

This goal cannot be solved (in reasonable time) using either `eauto` with `datatypes` or `firstorder` with `datatypes`. The hammer solves this goal using just one lemma: `Lists.List.in_app_iff`.

A similar example is

```
forall {A} (y1 y2 y3 : A) l l' z, In z l \ / In z l' ->
  In z (y1 :: y2 :: l ++ y3 :: l').
```

This goal cannot be solved using standard Coq tactics. Eprover with 512 preselected premises found a proof using two lemmas from `Lists.List:in_cons` and `in_or_app`.

The hammer is currently not capable of reasoning by induction, except in some very simple cases. Here is an example of a goal where induction is needed.

```
forall (A : Type) (P : A -> Prop) (a : A) (l l' : list A),
  List.Forall P l /\ List.Forall P l' /\ P a ->
  List.Forall P (l ++ a :: l').
```

This goal can be solved neither by standard Coq tactics nor by the hammer. However, it suffices to issue the `ltac` command `induction 1` and the hammer can solve the resulting two subgoals, none of which could be solved by standard Coq tactics. The subgoal for induction base is:

```
A : Type
P : A -> Prop
a : A
=====
forall l' : list A, Forall P nil \ / Forall P l' /\ P a ->
  Forall P (nil ++ a :: l')
```

The hammer solves this goal using the lemma `Forall_cons` from `Lists.List` and the definition of `++` (`Datatypes.app`). The subgoal for the induction step is:

```
A : Type
P : A -> Prop
a, a0 : A
l : list A
IH1 : forall l' : list A, Forall P l /\ Forall P l' /\ P a ->
      Forall P (l ++ a :: l')
=====
forall l' : list A, Forall P (a0 :: l) /\ Forall P l' /\ P a ->
      Forall P ((a0 :: l) ++ a :: l')
```

The hammer solves this goal using the lemma `Forall_cons`, the inductive hypothesis (IH1) and the definition of `++`. Note that to reconstruct the ATP proof for this goal it is crucial that our reconstruction tactics can do inversion on inductive predicates in the context.

9 Limitations

In this section we briefly discuss the limitations of the current implementation of the `CoqHammer` tool. We also compare the hammer with the automation tactics already available in `Coq`.

The intended use of a hammer is to prove relatively simple goals using accessible lemmas. Currently, the hammer works best with lemmas from the `Coq` standard library. Testing with other libraries has been as yet very limited and the hammer tool may need some adjustments to achieve comparable success rates.

The hammer works best when the goal and the needed lemmas are “close to” first-order logic, as some more sophisticated features of the `Coq` logic are not translated adequately. In particular, when dependent types are heavily used in a development then the effectiveness of the hammer tool is limited. Specifically, case analysis over inhabitants of small propositional inductive types is not translated properly, and the fact that in `Coq` all inhabitants of `Prop` are also inhabitants of `Type` is not accounted for.

A small propositional inductive type is an inductive type in `Prop` having just one constructor and whose arguments are all non-informative (e.g. propositional). In `Coq` it is possible to perform case analysis over an inhabitant of a small propositional inductive type. This is frequently done when dealing with data structures where dependent types are heavily exploited to capture the data structure invariants. Currently, all such pattern matches are translated to a fresh constant about which nothing is assumed. Therefore, the ATPs will fail to find a proof, except for trivial tautologies.

In `Coq` all propositions (inhabitants of `Prop`) are also types (inhabitants of `Type`). Therefore, type formers expecting types as arguments may sometimes be fed with propositions. For instance, one can use the pair type former as if it was a conjunction. Our translation heavily relies on the possibility of detecting whether a subterm is a proposition or not, in order to translate it to a FOL formula or a FOL term. The currently followed approach to proposition detection is relatively simplistic. For example, the pair type former should be translated to four different definitions, one taking in input two propositions, etc. Currently, only one definition is generated (the one with both arguments being of type `Type`).

In the context of code extraction the above two problems and some similar issues were handled in Pierre Letouzey’s Ph.D. thesis [60]. In [60] `Coq` terms are translated into an intermediate language where propositions are either removed from the terms or turned into unit types when used as types. It may be worthwhile to investigate if our translation could

be factorized reusing the intermediate representation from [60]. If successful, this would be a better approach.

We leave it for future work to increase effectiveness of the hammer on a broader fragment of dependent type theory. In this regard our hammer is similar to hammers for proof assistants based on classical higher-order logic, which are less successful when the goal or the lemmas make heavy use of higher-order features.

The success of the `hammer` tactic is not guaranteed to be reproducible, because it relies on external ATPs and uses time limits during proof reconstruction. Indeed, small changes in the statement of the goal or a change of hardware may change the behaviour of the hammer. However, once a proof has been found and successfully reconstructed the user should replace the `hammer` tactic with an appropriate reconstruction tactic shown by the hammer in the response window. This reconstruction tactic does not depend on any time limits or external ATPs, so its success is independent of the current machine.

In comparison to the hammer, domain-specific decision procedures, e.g., the `omega` tactic, are generally faster and more consistently reliable for the goals they can solve. On the other hand, the proof terms generated by the `hammer` tactic are typically smaller and contain fewer dependencies which are more human-readable.

An advantage of Coq proof-search tactics like `auto`, `eauto` or `firstorder` is that they can be configured by the user by means of hint databases. However, they are in general much weaker than the hammer. The idea of a hammer is to be a strong general-purpose tactic not requiring much configuration by the user.

10 Conclusions and Future Work

We have developed a first whole hammer system for intuitionistic type theory. This involved proposing an approximation of the Calculus of Inductive Constructions, adapting premise selection to this foundation, developing a translation mechanism to untyped-first order logic, and proposing reconstruction mechanisms for the proofs found by the ATPs. We have implemented the hammer as a plugin for the Coq proof assistant and evaluated it on all the proofs in its standard library. The source code of the plugin for Coq versions 8.5, 8.6 and 8.7, as well as all the experiments are available at: <http://cl-informatik.uibk.ac.at/cek/coqhammer/>

The hammer is able to re-prove completely automatically 40.8% of the standard library proofs on a 8-CPU system in about 40 s. This success rate is already comparable to that offered by the first generations of hammer systems for HOL and Mizar and can already offer a huge saving of human work.

To our knowledge this is the first translation which is usable by hammers. Strictly speaking, our translation is neither sound nor complete. However, our experiments suggest that the encoding is “sound enough” to be usable and that it is particularly good for goals close to first-order logic. Moreover, a “core” version of the translation is in fact sound [27].

There are many ways how the proposed work can be extended. First, the reconstruction mechanism currently is able to re-prove only 85.2% (4215 out of 4841) of the proofs founds by the ATPs, which is lower than that in other systems. The premise selection algorithms are not as precise as those involving machine learning algorithms tailored for particular logics. In particular, for similar size parts of the libraries almost the same premise selection algorithms used in HOLyHammer [52] or Isabelle/MaSh on parts of the Isabelle/HOL library [15], require on average 200–300 best premises to cover the dependencies, whereas in the Coq standard library on average 499–530 best premises are required.

The core of the hammer—the translation to FOL—could be improved to make use of more knowledge available in the prover in order to offer a higher success rate. It could also be modified to make it more effective on developments heavily using dependent types, and to more properly handle the advanced features of the Coq logic, possibly basing on some of the ideas in [60]. Finally, the dependencies extracted from the Coq proof terms do miss information used implicitly by the kernel, and are therefore not as precise as those offered in HOL-based systems.

In our work we have focused on the Coq standard library. Evaluations on a proof assistant standard library were common in many hammer comparisons, however this is rarely the level at which users are actually working, and looking at more advanced Coq libraries could give interesting insights for all components of a hammer. Since we focused on the standard library during development, it is likely that the effectiveness of the hammer is lower on libraries not similar to the standard library.

In particular, the Mathematical Components Library based on SSReflect [37] would be a particularly interesting example, as it heavily relies on unification hints to guide Coq automation. It has been used for example in the proofs of the four color theorem [38] and the odd order theorem [36]. On a few manually evaluated examples, the success rate is currently quite low. It remains to be seen, whether a hammer can provide useful automation also for such developments, and how the currently provided translation could be optimized, to account for the more common use of dependent types. Lastly, we would like to extend the work to other systems based on variants of CIC and other interesting foundations, including Matita, Agda, and Idris.

Acknowledgements Open access funding provided by Austrian Science Fund (FWF). We thank the organisers of the First Coq Coding Sprint, especially Yves Bertot, for the help with implementing Coq export plugins. We wish to thank Thibault Gauthier for the first version of the Coq exported data, as well as Claudio Sacerdoti-Coen for improvements to the exported data and fruitful discussions on Coq proof reconstruction. This work has been supported by the Austrian Science Fund (FWF) Grant P26201 and European Research Council (ERC) Grant No. 714034 *SMART*.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

1. Alemi, A.A., Chollet, F., Irving, G., Szegedy, C., Urban, J.: DeepMath—Deep sequence models for premise selection. In: Lee, D.D., Sugiyama, M., Luxburg, U.V., Guyon, I., Garnett, R. (eds.) *Advances in Neural Information Processing Systems (NIPS 2016)*, pp. 2235–2243 (2016)
2. Abel, A., Coquand, T., Norell, U.: Connecting a logical framework to a first-order logic prover. In: Gramlich, B. (ed.) *Frontiers of Combining Systems (FroCoS 2005)*, Volume 3717 of LNCS, pp. 285–301. Springer, New York (2005)
3. Armand, M., Faure, G., Grégoire, B., Keller, C., Théry, L., Werner, B.: A modular integration of SAT/SMT solvers to Coq through proof witnesses. In: Jouannaud, J., Shao, Z. (eds.) *Certified Programs and Proofs (CPP 2011)*, Volume 7086 of LNCS, pp. 135–150. Springer, New York (2011)
4. Alama, J., Heskens, T., Kühlwein, D., Tsiptsivadze, E., Urban, J.: Premise selection for mathematics by corpus analysis and kernel methods. *J. Autom. Reason.* **52**(2), 191–213 (2014)
5. Asperti, A., Ricciotti, W., Coen, C.: Sacerdoti: Matita tutorial. *J. Formaliz. Reason.* **7**(2), 91–199 (2014)
6. Aspinall, D.: Proof general: a generic tool for proof development. In: Graf, S., Schwartzbach, M.I. (eds.) *Tools and Algorithms for Construction and Analysis of Systems*, 6th International Conference, TACAS 2000, volume 1785 of LNCS, pp. 38–42. Springer, New York (2000)

7. Asperti, A., Tassi, E.: Higher order proof reconstruction from paramodulation-based refutations: the unit equality case. In: Kauer, M., Kerber, M., Miner, R., Windsteiger, W. (eds.) *Mathematical Knowledge Management (MKM 2007)*, Volume 4573 of LNCS, pp. 146–160. Springer, New York (2007)
8. Asperti, A., Tassi, E.: Smart matching. In: *Intelligent Computer Mathematics, 10th International Conference, AISC 2010, 17th Symposium, Calculemus 2010, and 9th International Conference, MKM 2010, Paris, France, July 5–10, 2010. Proceedings*, pp. 263–277 (2010)
9. Blanchette, J.C., Böhme, S., Fleury, M., Smolka, S.J., Steckermeier, A.: Semi-intelligible Isar proofs from machine-generated proofs. *J. Autom. Reason.* (2015)
10. Bancerek, G., Byliński, C., Grabowski, A., Kornitowicz, A., Matuszewski, R., Naumowicz, A., Pąk, K., Urban, J.: Mizar: State-of-the-art and beyond. In: *Intelligent Computer Mathematics—International Conference, CICM 2015, Washington, DC, USA, July 13–17, 2015. Proceedings*, pp. 261–279 (2015)
11. Bertot, Y., Castéran, P.: *Interactive Theorem Proving and Program Development: Coq’Art: The Calculus of Inductive Constructions*. Springer, New York (2004)
12. Broda, S., Damas, L.: On long normal inhabitants of a type. *J. Log. Comput.* **15**(3), 353–390 (2005)
13. Bove, A., Dybjer, P., Norell, U.: A brief overview of Agda—A functional language with dependent types. In: Berghofer, S., Nipkow, T., Urban, C., Wenzel, M. (eds.) *Theorem Proving in Higher Order Logics (TPHOLs 2009)*, Volume 5674 of LNCS, pp. 73–78. Springer, New York (2009)
14. Bertot, Y.: A short presentation of Coq. In: Mohamed, O.A., Muñoz, C.A., Tahar, S. (eds.) *Theorem Proving in Higher Order Logics (TPHOLs 2008)*, Volume 5170 of LNCS, pp. 12–16. Springer, New York (2008)
15. Blanchette, J.C., Greenaway, D., Kaliszky, C., Kühlwein, D., Urban, J.: A learning-based fact selector for Isabelle/HOL. *J. Autom. Reason.* **57**(3), 219–244 (2016)
16. Bezem, M., Hendriks, D., de Nivelle, H.: Automated proof construction in type theory using resolution. *J. Autom. Reason.* **29**(3–4), 253–275 (2002)
17. Blanchette, J.C., Kaliszky, C., Paulson, L.C., Urban, J.: Hammering towards QED. *J. Formaliz. Reason.* **9**(1), 101–148 (2016)
18. Blanchette, J.C.: *Automatic Proofs and Refutations for Higher-Order Logic*. PhD thesis, Technische Universität München (2012). <http://www21.in.tum.de/~blanchet/phdthesis.pdf>
19. Brady, E.: Idris, a general-purpose dependently typed programming language: design and implementation. *J. Funct. Program.* **23**(5), 552–593 (2013)
20. Böhme, S., Weber, T.: Fast LCF-style proof reconstruction for Z3. In: Kaufmann, M., Paulson, L. (eds.) *Interactive Theorem Proving (ITP 2010)*, Volume 6172 of LNCS, pp. 179–194. Springer, New York (2010)
21. Ben-Yelles, C.: *Type-assignment in the lambda-calculus: syntax and semantics*. Ph.D. thesis, Mathematics Department, University of Wales, Swansea, UK (1979)
22. Coquand, T., Huet, G.P.: The calculus of constructions. *Inf. Comput.* **76**(2/3), 95–120 (1988)
23. Chlipala, A.: *Certified Programming with Dependent Types—A Pragmatic Introduction to the Coq Proof Assistant*. MIT Press, Cambridge (2013)
24. Czajka, Ł., Kaliszky, C.: Goal translation for a hammer for Coq (extended abstract). In: Blanchette, J.C., Kaliszky, C. (eds.) *First International Workshop on Hammers for Type Theories (HaTT 2016)*, Volume 210 of EPTCS, pp. 13–20 (2016)
25. Coq Development Team: *The Coq proof assistant reference manual* (2016). Version 8.6
26. Corbineau, P.: First-order reasoning in the calculus of inductive constructions. In: Berardi, S., Coppo, M., Damiani, F. (eds.) *Types for Proofs and Programs (TYPES 2003)*, Volume 3085 of LNCS, pp. 162–177. Springer, New York (2003)
27. Czajka, Ł.: A shallow embedding of pure type systems into first-order logic. Submitted. (2016). <http://www.mimuw.edu.pl/~lukaszcz/emb.pdf>
28. de Moura, L.M., Bjørner, N.: Z3: An efficient SMT solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) *TACAS 2008*, Volume 4963 of LNCS, pp. 337–340. Springer, New York (2008)
29. de Moura, L.M., Kong, S., Avigad, J., van Doorn, F., von Raumer, J.: The Lean theorem prover. In: Felty, A.P., Middeldorp, A. (eds.) *International Conference on Automated Deduction (CADE 2015)*, Volume 9195 of LNCS, pp. 378–388. Springer, New York (2015)
30. de Moura, L., Selsam, D.: Congruence closure in intensional type theory. In: Olivetti, N., Tiwari, A. (eds.) *International Joint Conference on Automated Reasoning, IJCAR 2016*, Volume 9706 of LNCS. Springer, New York (2016)
31. Dowek, G.: A complete proof synthesis method for the cube of type systems. *J. Log. Comput.* **3**(3), 287–315 (1993)
32. Dyckhoff, R.: Contraction-free sequent calculi for intuitionistic logic. *J. Symb. Log.* **57**(3), 795–807 (1992)
33. Filliâtre, J.-C.: One logic to use them all. In: Bonacina, M.P. (ed.) *International Conference on Automated Deduction (CADE 2013)*, Volume 7898 of LNCS, pp. 1–20. Springer, New York (2013)

34. Färber, M., Kaliszyk, C.: Random forests for premise selection. In: Lutz, C., Ranise, S. (eds.) *Frontiers of Combining Systems (FroCoS 2015)*, Volume 9322 of LNCS, pp. 325–340 (2015)
35. Filliâtre, J.-C., Paskевич, A.: Why3—Where programs meet provers. In: Felleisen, M., Gardner, P. (eds.) *European Symposium on Programming (ESOP 2013)*, Volume 7792 of LNCS, pp. 125–128. Springer, New York (2013)
36. Gonthier, G., Asperti, A., Avigad, J., Bertot, Y., Cohen, C., Garillot, F., Roux, S.L., Mahboubi, A., O'Connor, R., Biha, S.O., Pasca, I., Rideau, L., Solovyev, A., Tassi, E., Théry, L.: A machine-checked proof of the odd order theorem. In: Blazy, S., Paulin-Mohring, C., Pichardie, D. (eds.) *Interactive Theorem Proving (ITP 2013)*, Volume 7998 of LNCS, pp. 163–179. Springer, New York (2013)
37. Gonthier, G., Mahboubi, A.: An introduction to small scale reflection in Coq. *J. Formaliz. Reason.* **3**(2), 95–152 (2010)
38. Gonthier, G.: The four colour theorem: Engineering of a formal proof. In: Kapur, D. (ed.) *ASCM*, Volume 5081 of LNCS, pp. 333. Springer, New York (2007)
39. Gransden, T., Walkinshaw, N., Raman, R.: SEPIA: search for proofs using inferred automata. In: Felty, A.P., Middeldorp, A. (eds.) *International Conference on Automated Deduction (CADE 2015)*, Volume 9195 of LNCS, pp. 246–255. Springer, New York (2015)
40. Harrison, J.: HOL light: an overview. In: Berghofer, S., Nipkow, T., Urban, C., Wenzel, M. (eds.) *Theorem Proving in Higher Order Logics (TPHOLs 2009)*, Volume 5674 of LNCS, pp. 60–66. Springer, New York (2009)
41. Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., Witten, I.H.: The Weka data mining software: an update. *SIGKDD Explor. Newsl.* **11**(1), 10–18 (2009)
42. Hindley, J.R.: *Basic Simple Type Theory*, Volume 42 of Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, Cambridge (1997)
43. Hurd, J.: First-order proof tactics in higher-order logic theorem provers. In: Archer, M., Vito, B.D., Muñoz, C. (eds.) *Design and Application of Strategies/Tactics in Higher Order Logics (STRATA 2003)*, Number NASA/CP-2003-212448 in NASA Technical Reports, pp. 56–68 (2003)
44. Harrison, J., Urban, J., Wiedijk, F.: History of interactive theorem proving. In: Siekmann, J. (ed.) *Handbook of the History of Logic vol 9 (Computational Logic)*, pp. 135–214. Elsevier, Amsterdam (2014)
45. Hoder, K., Voronkov, A.: Sine qua non for large theory reasoning. In: Bjørner, N., Sofronie-Stokkermans, V. (eds.) *23rd International Conference on Automated Deduction (CADE 2011)*, Volume 6803 of LNCS, pp. 299–314. Springer, New York (2011)
46. Joosten, S., Kaliszyk, C., Urban, J.: Initial experiments with TPTP-style automated theorem provers on ACL2 problems. In: Verbeek, F., Schmaltz, J. (eds.) *ACL2 Theorem Prover and Its Applications (ACL2 2014)*, Volume 152 of EPTCS, pp. 77–85 (2014)
47. Jones, K.S.: A statistical interpretation of term specificity and its application in retrieval. *J. Doc.* **28**, 11–21 (1972)
48. Komendantskaya, E., Heras, J., Grov, G.: Machine learning in Proof General: Interfacing interfaces. In: Kaliszyk, C., Lüth, C. (eds.) *User Interfaces for Theorem (UITP 2012)*, Volume 118 of EPTCS, pp. 15–41 (2013)
49. Kaliszyk, C., Mamane, L., Urban, J.: Machine learning of Coq proof guidance: First experiments. In: Kutsia, T., Voronkov, A. (eds.) *Symbolic Computation in Software Science (SCSS 2014)*, Volume 30 of EPiC, pp. 27–34. EasyChair (2014)
50. Kaliszyk, C., Urban, J.: PRoCH: Proof reconstruction for HOL Light. In: Bonacina, M.P. (ed.) *International Conference on Automated Deduction (CADE 2013)*, Volume 7898 of LNCS, pp. 267–274. Springer, New York (2013)
51. Kaliszyk, C., Urban, J.: Stronger automation for Flyspeck by feature weighting and strategy evolution. In: Blanchette, J.C., Urban, J. (eds.) *Proof Exchange for Theorem Proving (PxTP 2013)*, Volume 14 of EPiC, pp. 87–95. EasyChair (2013)
52. Kaliszyk, C., Urban, J.: Learning-assisted automated reasoning with Flyspeck. *J. Autom. Reason.* **53**(2), 173–213 (2014)
53. Kaliszyk, C., Urban, J.: HOL(y)Hammer: online ATP service for HOL light. *Math. Comput. Sci.* **9**(1), 5–22 (2015)
54. Kaliszyk, C., Urban, J.: Learning-assisted theorem proving with millions of lemmas. *J. Symb. Comput.* **69**, 109–128 (2015)
55. Kaliszyk, C., Urban, J.: MizAR 40 for Mizar 40. *J. Autom. Reason.* **55**(3), 245–256 (2015)
56. Kaliszyk, C., Urban, J., Vyskočil, J.: Efficient semantic features for automated reasoning over large theories. In: Yang, Q., Wooldridge, M. (eds.) *International Joint Conference on Artificial Intelligence (IJCAI 2015)*, pp. 3084–3090. AAAI Press, Palo Alto (2015)
57. Kovács, L., Voronkov, A.: First-order theorem proving and Vampire. In: Sharygina, N., Veith, H. (eds.) *Computer-Aided Verification (CAV 2013)*, Volume 8044 of LNCS, pp. 1–35. Springer, New York (2013)

58. Kühlwein, D., van Laarhoven, T., Tsvitsivadze, E., Urban, J., Heskes, T.: Overview and evaluation of premise selection techniques for large theory mathematics. In: Gramlich, B., Miller, D., Sattler, U. (eds.) International Joint Conference on Automated Reasoning (IJCAR 2012), volume 7364 of LNCS, pp. 378–392. Springer, New York (2012)
59. Laurent, J.: Suggesting relevant lemmas by learning from successful proofs. Technical report, École normale supérieure (2016). Internship Report
60. Letouzey, P.: Programmation fonctionnelle certifiée : L'extraction de programmes dans l'assistant Coq. (Certified functional programming : Program extraction within Coq proof assistant). PhD thesis, University of Paris-Sud, Orsay, France, (2004)
61. Meng, J., Paulson, L.C.: Translating higher-order clauses to first-order clauses. *J. Autom. Reason.* **40**(1), 35–60 (2008)
62. Meng, J., Paulson, L.C.: Lightweight relevance filtering for machine-generated resolution problems. *J. Appl. Log.* **7**(1), 41–57 (2009)
63. Paulson, L.C., Blanchette, J.: Three years of experience with Sledgehammer, a practical link between automated and interactive theorem provers. In: 8th IWIL (2010)
64. Paulson, L.C., Susanto, K.W.: Source-level proof reconstruction for interactive theorem proving. In: Schneider, K., Brandt, J. (eds.) Theorem Proving in Higher Order Logics (TPHOLs 2007), Volume 4732 of LNCS, pp. 232–245. Springer, New York (2007)
65. Schulz, S.: System description: E 1.8. In: McMillan, K.L., Middeldorp, A., Voronkov, A. (eds.) Logic for Programming, Artificial Intelligence (LPAR 2013), Volume 8312 of LNCS, pp. 735–743. Springer, New York (2013)
66. Schmitt, S., Lorigo, L., Kreitz, C., Nogin, A.: Jprover : Integrating connection-based theorem proving into interactive proof assistants. In: Goré, R., Leitsch, A., Nipkow, T. (eds.) Automated Reasoning, First International Joint Conference, IJCAR 2001, Siena, Italy, June 18–23, 2001, Proceedings, Volume 2083 of Lecture Notes in Computer Science, pp. 421–426. Springer, New York (2001)
67. Slind, K., Norrish, M.: A brief overview of HOL4. In: Mohamed, O.A., Muñoz, C., Tahar, S. (eds.) TPHOLs 2008, Volume 5170 of LNCS, pp. 28–32. Springer, New York (2008)
68. Sutcliffe, G.: The TPTP world-infrastructure for automated reasoning. In: Clarke, E., Voronkov, A. (eds.) LPAR-16, Number 6355 in LNAI, pp. 1–12. Springer, New York (2010)
69. Tamm, T., Smith, J.M.: Optimized encodings of fragments of type theory in first-order logic. *J. Log. Comput.* **8**(6), 713–744 (1998)
70. Urban, J.: MPTP—motivation, implementation. *First Exp. J. Autom. Reason.* **33**(3–4), 319–339 (2004)
71. Urzyczyn, P.: Intuitionistic games: determinacy, completeness, and normalization. *Stud. Log.* **104**(5), 957–1001 (2016)
72. Urban, J., Sutcliffe, G.: Automated reasoning and presentation support for formalizing mathematics in Mizar. In: Autexier, S., Calmet, J., Delahaye, D., Ion, P.D.F., Rideau, L., Rioboo, R., Sexton, A.P. (eds.) Intelligent Computer Mathematics (CICM 2010), Volume 6167 of LNCS, pp. 132–146 (2010)
73. Wiedijk, F.: Mizar's soft type system. In: Theorem Proving in Higher Order Logics, 20th International Conference, TPHOLs 2007, Kaiserslautern, Germany, September 10–13, 2007, Proceedings, pp. 383–399 (2007)
74. Wenzel, M., Paulson, L.C., Nipkow, T.: The Isabelle framework. In: Mohamed, O.A., Muñoz, C.A., Tahar, S. (eds.) Theorem Proving in Higher Order Logics (TPHOLs 2008), Volume 5170 of LNCS, pp. 33–38. Springer, New York (2008)
75. Zielenkiewicz, M., Schubert, A.: Automata theory approach to predicate intuitionistic logic. In: Logic-Based Program Synthesis and Transformation—26th International Symposium, LOPSTR 2016, Revised Selected Papers, pp. 345–360 (2016)