

# A Curiously Effective Backtracking Strategy for Connection Tableaux

Michael Färber

## Abstract

Automated proof search with connection tableaux, such as implemented by Otten’s leanCoP prover, depends on backtracking for completeness. Otten’s restricted backtracking strategy loses completeness, yet for many problems, it significantly reduces the time required to find a proof. I introduce a new, less restricted backtracking strategy based on the notion of *exclusive cuts*. I implement the strategy in a new prover called *meanCoP* and show that it greatly improves upon the previous best strategy in leanCoP.

## Keywords

connection tableaux, backtracking, exclusive cut, REX, leanCoP, meanCoP

## 1. Introduction

Bibel’s connection method [1] is a proof search method similar to Andrews’s matings [2]. Compared to other proof search methods such as resolution, the connection method has several merits: It is goal-oriented, enabling natural conjecture-directed proof search. It can be used with relatively little effort for non-classical logics such as intuitionistic or modal logics [3, 4], and non-clausal search [5]. Finally, most connection calculi have only very few and simple rules, making it easy to certify proofs in proof assistants such as HOL Light [6, 7].

One of the most influential connection provers is Otten’s leanCoP [3]. Its outstanding ratio between code size and effectiveness has made it a frequently used vehicle to experiment with new search strategies. leanCoP uses bounded depth-first search together with iterative deepening to explore larger and larger potential proofs. As the proof search is not confluent, leanCoP employs backtracking to preserve completeness.

This article studies backtracking that guides connection proof search. In particular, a backtracking strategy deals with the question: When a literal  $L$  is solved with a proof  $P$ , which alternative proofs  $P'$  to solve  $L$  can proof search consider afterwards? A complete backtracking strategy does not impose any restriction on  $P'$ ; that is, it allows proof search to consider all different proofs  $P'$  for  $L$ . Otten showed that by restricting backtracking, the prover becomes significantly more effective for many problems as this reduces the search space, at the expense of losing completeness (section 3). His “restricted backtracking” strategy prevents exploring any alternative proof  $P'$  for a solved literal  $L$ . In this paper, I introduce a novel incomplete strategy called “less restricted backtracking”, which prevents exploring any alternative proof  $P'$  for a

---


AReCCa 2023: Automated Reasoning with Connection Calculi, 18 September 2023, Prague, Czech Republic

✉ michael.farber@gedenkt.at (M. Färber)

🆔 0000-0003-1634-9525 (M. Färber)



© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

solved literal  $L$  where  $P'$  starts with the same root step as  $P$  (section 4). In other words, for any literal  $L$ , unrestricted backtracking considers all proofs, restricted backtracking considers only a single proof, and less restricted backtracking considers only proofs with differing root steps.

I unexpectedly discovered less restricted backtracking upon implementing a new prover called *meanCoP* based on leanCoP (section 5). The new strategy improves upon the former best strategy in leanCoP dramatically (section 6).

## 2. Preliminaries

In this article, we will use classical first-order logic without equality. However, the techniques shown can be also applied to non-classical logics.

A term  $t$  is either a variable (denoted by  $x, y, z$ ) or the application of a constant (denoted by  $a, b, c$ ) to terms. An atom  $A$  is the application of a predicate (denoted by  $p, q, r$ ) to terms. Predicates and constants have associated fixed arities. A literal  $L$  is an atom  $A$  or its negation  $\neg A$ . The complement  $\bar{\phantom{x}}$  of a literal is defined such that  $\overline{\bar{A}} = A$  and  $\overline{\neg A} = \neg A$ . A term substitution  $\sigma$  is a mapping from variables to terms. Applying a substitution  $\sigma$  to a literal  $L$ , denoted as  $\sigma L$ , substitutes all variables of  $L$  with their mappings. Two literals  $L_1, L_2$  can be unified under a substitution  $\sigma$  if  $\sigma L_1 = \sigma L_2$ .

A formula in conjunctive normal form (CNF) is a conjunction ( $\wedge$ ) of disjunctions ( $\vee$ ) of literals. A clause is a set of literals, and a matrix is a set of clauses. We interpret a clause as the disjunction of its literals, and we interpret a matrix as the conjunction of its (interpreted) clauses. It is easy to see that for each formula in CNF, there is an equivalent matrix.

**Example 1.** Consider the formula

$$(p(x) \vee q(x)) \wedge (\neg p(y) \vee r(y)) \wedge \neg p(z) \wedge \neg r(a) \wedge \neg r(b) \wedge \neg q(c).$$

Its equivalent matrix is

$$M = \left[ \begin{array}{c} \left[ \begin{array}{c} p(x) \\ q(x) \end{array} \right] \left[ \begin{array}{c} \neg p(y) \\ r(y) \end{array} \right] \left[ \neg p(z) \right] \left[ \neg r(a) \right] \left[ \neg r(b) \right] \left[ \neg q(c) \right] \end{array} \right],$$

which we will use as running example throughout this paper.

In this paper, we treat proof search using the clausal connection tableaux calculus [8, 9].<sup>1</sup>


**Definition 1** (Connection Calculus). *The axiom and the rules of the clausal connection calculus are given in Figure 1. The words of the connection calculus are tuples  $\langle C, M, Path \rangle$ , where  $M$  is a matrix, and  $C$  and  $Path$  are sets of literals or  $\varepsilon$ .  $C$  is called the subgoal clause and  $Path$  is called the active path. In the calculus rules,  $\sigma$  is a global (or rigid) term substitution; that is, it is applied to the whole derivation.*

<sup>1</sup>Unlike this article, [8, 9] use disjunctive normal form (DNF) and check for validity of a formula. The two presentations are dual; in particular, the DNF of a formula is valid iff the CNF of its negation is unsatisfiable.

Axiom	$\frac{}{\{\}, M, Path} A$
Start	$\frac{C_2, M, \{\}}{\varepsilon, M, \varepsilon} S$ where $C_2$ is a copy of $C_1 \in M$
Reduction	$\frac{C, M, Path \cup \{L'\}}{C \cup \{L\}, M, Path \cup \{L'\}} R$ where $\sigma(L) = \sigma(\overline{L'})$
Extension	$\frac{C_2 \setminus \{L'\}, M, Path \cup \{L\} \quad C, M, Path}{C \cup \{L\}, M, Path} E$

where  $C_2$  is a copy of  $C_1 \in M$  and  $L' \in C_2$  with  $\sigma(L) = \sigma(\overline{L'})$

**Figure 1: Clausal connection calculus rules.**



**Figure 2: Graphical representation of proof search.**

An application of a proof rule is called a proof step. A derivation for  $\langle C, M, Path \rangle$  with the term substitution  $\sigma$ , in which all leaves are axioms, is called a *connection proof* for  $\langle C, M, Path \rangle$ . A connection proof for  $\langle \varepsilon, M, \varepsilon \rangle$  is called a connection proof for  $M$ .

Bibel proved soundness and completeness of the calculus: for any formula  $F$  in CNF, we have that  $F$  is unsatisfiable iff there is a connection proof for the matrix corresponding to  $F$  [1].

Proof search proceeds by constructing derivations from bottom to top. We can understand a derivation for  $\langle C, M, Path \rangle$  as an attempt to prove  $(M \wedge Path) \implies C$ , where we interpret  $Path$  as conjunction of its literals. By interpreting  $\varepsilon$  as empty set, a derivation for  $\langle \varepsilon, M, \varepsilon \rangle$  can be seen as a proof attempt of  $M \implies \perp$ .

We say that any reduction or extension step as in Figure 1 *connects*  $L$  to  $L'$ . We illustrate this by drawing an arrow from  $L$  to  $L'$  in the matrix. In this paper, we will only use extension steps in examples.

Let us walk through a *failed* proof search attempt for the matrix  $M$  from Example 1, and show its graphical representation as well as its resulting derivation in the calculus.

$$\begin{array}{c}
\frac{\frac{\frac{\{\}, M, \{p(x), r(y)\}}{\{r(y)\}, M, \{p(x)\}} \text{A}}{\{p(x), q(x)\}, M, \{\}} \text{S}}{\varepsilon, M, \varepsilon} \text{S} \quad \frac{\frac{\{\}, M, \{p(x)\}}{\{q(x)\}, M, \{\}} \text{E (1.1)}}{\vdots} \text{E (1.1)} \\
\text{E (1.2)}
\end{array}$$

**Figure 3: Derivation with  $\sigma(x) = \sigma(y) = a$ .**

**Example 2.** Consider matrix  $M$  from Example 1. Matrix (1) of Figure 2 illustrates a proof search attempt through  $M$ . We write the proof step  $n$  in matrix  $m$  as  $(m.n)$  and mark situations in which we are stuck with  $\downarrow$ . The proof search proceeds as follows: We first choose the first clause in  $M$  as start clause. This obliges us to connect both  $p(x)$  and  $q(x)$ . We start with  $p(x)$ , which we choose to connect in step (1.1) to  $\neg p(y)$ , setting  $\sigma(x) = y$ . This in turn obliges us to connect  $r(y)$ , which we choose to connect in step (1.2) to  $\neg r(a)$ , setting  $\sigma(x) = \sigma(y) = a$ . We are now left with the obligation to connect  $q(x)$ . However, at this point (1.3), we cannot connect  $q(x)$  to any literal due to  $\sigma$ . As we are stuck at this point, we mark this with  $\downarrow$ . Figure 3 shows a derivation for  $M$  that corresponds to that proof search. The extension steps in the derivation are labelled like the corresponding proof steps in matrix (1). The derivation is not a connection proof for  $M$ , because the leaf  $\langle \{q(x)\}, M, \{\} \rangle$  is not an axiom.

This example illustrates that search in connection tableaux is not confluent, i.e., we can end up with unprovable leaves in derivations for a matrix  $M$  although the formula corresponding to  $M$  is unsatisfiable. This makes it necessary to backtrack to previous states of derivations to obtain a complete proof search method. We will study two backtracking strategies in the next section.

### 3. Backtracking

In the failed proof attempt shown in Example 2, we frequently talked about *obligations* and *choices*. Fulfilling obligations assures soundness, and making alternative choices exhaustively assures completeness.

Otten's *unrestricted backtracking* strategy [3] is sound and complete. It makes choices until an obligation cannot be fulfilled. At this point, the strategy changes the most recent choice for which there is an untried alternative. The strategy succeeds if it fulfills all obligations, and fails if it runs out of alternatives.

**Example 3 (Unrestricted Backtracking).** Consider the proof search attempt in Example 2. The last choice we made in that example was to connect  $r(y)$  to  $\neg r(a)$  as part of step (1.2). We can make a different choice here, namely connect  $r(y)$  to  $\neg r(b)$ , which we perform in step (2.1). However, as it turns out, this will not help us once we have to deal with  $q(x)$  anew in step (2.2), for now we have  $\sigma(x) = \sigma(y) = b$ , which still does not permit a connection from  $q(x)$ . So we backtrack again, leading to the proof search shown in matrix (3). This time, the last choice was to connect  $r(y)$  to  $\neg r(b)$ , but now, we cannot find a different way to connect  $r(y)$ . So we look at our second to last choice, namely to connect  $p(x)$  to  $\neg p(y)$ . We can make an alternative choice here as step (3.1),

namely to connect  $p(x)$  to  $\neg p(z)$ . Now we are back once more to the dreaded  $q(x)$ , but finally, due to  $\sigma(x) = \sigma(z)$  not pointing to an actual term, we can connect  $q(x)$  to  $\neg q(c)$  as step (3.2). This concludes the proof, as we have no more obligations left at this point.

Otten's *restricted backtracking* strategy [10] is sound, but incomplete. However, it is often significantly more effective than the complete strategy. To define it, Otten introduces the property of *solvedness* on literals in a proof search.

**Definition 2** (Principal literal, solved literal). *When the reduction or extension rules are applied, the literal  $L$  (see Figure 1) is called the principal literal of the proof step. A reduction step solves a literal  $L$  iff  $L$  is its principal literal. An extension step  $S$  solves a literal  $L$  iff  $L$  is the principal literal of  $S$  and there is a proof for the left premise of  $S$ , i.e., there is a derivation for the left premise of  $S$  having only axioms as leaves.*

The restricted backtracking strategy works like the unrestricted one, with one exception: Once a literal is solved, restricted backtracking discards all choices to solve the literal differently.

**Example 4** (Restricted Backtracking). *Consider matrix (1) of Figure 2. Proof step (1.1) does not solve any literal, so at this point, proof search behaves like in Example 3. Proof step (1.2) solves  $r(y)$ , so at that point, alternative choices to solve  $r(y)$  are discarded. At the same time, step (1.2) also solves  $p(x)$ , so alternative choices to solve  $p(x)$  are discarded as well. In proof step (1.3), we note that  $q(x)$  cannot be connected. However, unlike in Example 3, we have no alternative choices left to backtrack to because they were discarded as a result of step (1.2). That means that the restricted backtracking strategy cannot find a proof once we commit to connecting  $p(x)$  to  $\neg p(y)$  as first step.*


## 4. Less Restricted Backtracking

Restricted backtracking can be decomposed into two cuts: cuts on reduction and cuts on extension steps. Kaliszyk already implemented these two cuts separately, but did not describe it, as they are usually most useful in conjunction. Here, the distinction arises naturally, as it allows to more succinctly describe a new backtracking strategy.

I will now distinguish *inclusive* and *exclusive* cuts. An inclusive cut discards all alternatives to solve a literal, whereas an exclusive cut discards all alternatives to solve a literal, *except for derivations starting with a different proof step*. Otten's restricted backtracking strategy shown in section 3 uses inclusive cuts on both reduction and extension steps. To the best of my knowledge, exclusive cuts have not been researched before.

**Example 5** (Exclusive Cut). *We will, for the last time, revisit the proof search in Figure 2. After proof step (1.2), exclusive cut discards alternative ways to solve  $p(x)$ , except for derivations starting with different extension steps. As a result, after being stuck at step (1.3) with  $q(x)$ , we can backtrack unlike in Example 4, namely to the proof search in matrix (3), because it solves  $p(x)$  in step (3.1) starting with a different extension step. From there, proof search behaves again like in Example 3, solving the problem in step (3.2) after connecting  $q(x)$ .*

To sum up the outcomes of different backtracking strategies on the proof search in Figure 2: The complete strategy (Example 3) solves the problem, going through all stages from (1) to (3).



**Figure 4: Effect of different cuts on the tree of alternatives.**

The exclusive cut (Example 5) also solves the problem, but takes one stage less, going through only (1) and (3). The inclusive cut (Example 4) fails after (1). For this example, exclusive cut therefore is the most efficient strategy.

For reduction steps, an exclusive cut is equivalent to no cut, so I distinguish between inclusive and exclusive cut only for extension steps. I abbreviate (inclusive) cut on reduction steps as R and inclusive and exclusive cut on extension steps as EI and EX, respectively. Otten’s restricted backtracking strategy can be described as a combination of R and EI, written as REI.

Figure 4 visualises the alternatives that are cut once a literal is solved. In each of the trees, the left child of the root is a proof step  $S$  that solves a literal, the children of the left child are alternatives to proof steps that are descendants of  $S$ , and the right child is the alternative to  $S$ . Reduction and extension steps are marked as R and E, respectively, and alternatives are marked as “?”. Both the R and EI cut are inclusive cuts because the right child is cut, and the EX cut is exclusive because the right child is preserved. Both EI and EX cuts eliminate all alternatives below the proof step.


The seemingly small difference between inclusive and exclusive cut has a large impact on the effectiveness of the prover. We will see this in the evaluation (section 6).

## 5. Implementation

The connection prover leanCoP is compactly implemented in Prolog as a recursive predicate `prove` that takes  $C$  and  $Path$  as parameters, using a helper predicate `lit` that models  $M$  [3]. leanCoP implements restricted backtracking using Prolog’s built-in cut operator [10]. Kaliszyk has reimplemented leanCoP using stacks for backtracking [11].

Based on Kaliszyk’s stack-based implementation, I implemented a connection prover called *meanCoP* in Rust.<sup>2</sup> Like C++, Rust favours zero-overhead abstractions [12], making it a suitable candidate for the development of high-performance automated theorem provers. I use functional programming for preprocessing and imperative programming for the prover loop. I follow Kaliszyk’s implementation for the prover loop, but I use dynamic instead of static arrays in order to allow for arbitrarily sized stacks, terms, etc. The prover loop does not use Rust’s standard library and can be therefore compiled to targets such as WASM, which can be used to create websites with an embedded prover that is run locally in a web browser. Furthermore, meanCoP

<sup>2</sup>The name meanCoP abbreviates “more efficient, albeit non-lean connection prover”. The source code of meanCoP is available at <https://github.com/01mf02/cop-rs>. I evaluated revision 884aea4 compiled with Rust 1.49.



**Figure 5: Effect of different cuts on the stack of alternatives.**

contains a tiny proof checker that is run before outputting a proof. This is useful to assure that the prover is sound.

meanCoP supports most major features of leanCoP, such as conjecture-directed search, regularity, and lemmas [3]. Furthermore, meanCoP supports the R, EI, and EX cuts (section 4). By default, meanCoP uses (inclusive) cut on lemma steps, which does not hamper completeness, as lemma steps do not impact the substitution.

Kaliszyk uses a stack of alternatives to keep track of proof steps to backtrack to. This allows for a compact implementation of inclusive and exclusive cuts. Figure 5 shows the effect of inclusive and exclusive cut on the stack of alternatives. Figure 5a shows the initial situation of the stack after a literal was solved with a proof step whose alternative is  $a_n$ . Above  $a_n$  are alternatives to proof steps added after  $a_n$ , and below  $a_n$  are alternatives to proof steps added before  $a_n$ . Using no cut does not change the stack at this point. Both exclusive and inclusive cut eliminate all alternatives added after  $a_n$ , but the exclusive cut (Figure 5b) keeps one alternative more than the inclusive cut (Figure 5c), namely  $a_n$ . Using inclusive instead of exclusive cut amounts to truncating the stack of alternatives to length  $n$  instead of length  $n - 1$  once a literal was solved.

## 6. Evaluation

I evaluate the performance of meanCoP (section 5) and other provers on several first-order problem datasets.<sup>3</sup> For every dataset and prover, I measure the number of problems solved by the prover in a given time. All evaluated connection provers use a single strategy with conjecture-directed search and non-definitional (i.e. standard or naive) translation into CNF [13, 10], unless specified otherwise. I use the same hardware, the same timeout, and the same datasets as in my previous evaluation of connection provers together with Kaliszyk and Urban [14]. I will compare the results in this paper with those of the previous evaluation.

I use a 48-core server with AMD Opteron 6174 2.2GHz CPUs, 320 GB RAM, and 0.5 MB L2 cache per CPU. Each problem is always assigned one CPU. I run every prover with a timeout of 10 seconds per problem.

I use several first-order logic datasets for evaluation, with statistics given in Table 1:

- TPTP [15] is a large benchmark for automated theorem provers. It is used in CASC [16].

<sup>3</sup>The evaluation results are available in more detail at <http://cl-informatik.uibk.ac.at/~mfaerber/arecca-2023.html>.

**Table 1**  
**Evaluation datasets and the number of contained first-order problems.**

Dataset	TPTP	bushy	chainy	Miz40	FS-top
Problems	7492	2078	2078	32524	27111

**Table 2**  
**Number of solved problems.**

Cut	TPTP	bushy	chainy	Miz40	FS-top
None	1731	546	208	9247	4038
R	1857	644	252	12965	4447
EI	1984	724	333	13853	4249
EX	2056	820	268	15507	4758
REI	1988	730	<b>341</b>	13562	4267
REX	<b>2126</b>	<b>850</b>	294	<b>16135</b>	<b>4994</b>

Its problems are based on different logics and come from various domains. I use the nonclausal first-order problems (files matching `*+?.p`) of TPTP 6.3.0.

- MPTP2078 [17] contains 2078 problems exported from the Mizar Mathematical Library. It comes in the two flavours “bushy” and “chainy”: In the “chainy” dataset, every problem contains all facts stated before the problem, whereas in the “bushy” dataset, every problem contains only the Mizar premises required to prove the problem.
- Miz40 contains the problems from the Mizar library for which at least one ATP proof has been found using one of the 14 combinations of provers and premise selection methods considered in [18]. The problems are translated to untyped first-order logic using the MPTP infrastructure [19]. The problems are minimised using ATP-based minimisation, i.e., re-running the ATP only with the set of proof-needed axioms until this set no longer becomes smaller. This typically leads to even better axiom pruning and ATP-easier problems than in the Mizar-based pruning used for the “bushy” version above.
- FS-top is a translation to first-order logic of the top-level HOL Light theorems of the Flyspeck project, which finished in 2014 a formal proof of the Kepler conjecture [20].

## 6.1. Comparison of meanCoP strategies

The first part of the evaluation studies the impact of different combinations of cuts on the number of problems solved by meanCoP.

Table 2 shows the number of problems solved by strategy. For all datasets, the complete strategy without any cut solves the least problems. Among the previously implemented cuts, namely R, EI, and REI, REI solves the most problems, except on the dataset FS-top, where R prevails. Adding cut on reduction (R) to any strategy increases the number of solved problems, except for the Miz40 dataset. The strategies with exclusive cut on extension steps (EX, REX)



**Table 3**  
**Union of solved problems.**

Cut	TPTP	bushy	chainy	Miz40	FS-top
Any but (R)EX	2277	834	391	15885	5185
REX and REI	2297	907	389	16803	5191
Any	2387	927	402	17437	5560

outperform those with inclusive cut (EI, REI) on all datasets except for the chainy one. I explore the reason for this in [subsection 6.2](#).

On most datasets, the strategies using exclusive cut bring an impressive improvement of the prover power. The REX strategy increases the number of solved problems compared to the REI strategy by 16.4% for bushy, 17.0% for FS-top (12.3% if we compare with the R cut), and 19.0% for Miz40. Remarkably, on TPTP, the improvement turns out much smaller with only 6.9%.

Table 3 shows the union of problems solved by a portfolio of strategies. The first row shows the problems solved by any of the four previously used cut strategies, including the unrestricted backtracking strategy without cut, but also combinations of cut on reduction and inclusive cut on extension, while excluding our new exclusive cut. Comparing the first row with the REX results from Table 2, we see that the new REX strategy solves single-handedly more problems than a union of four strategies on the bushy and Miz40 datasets, which is quite noteworthy. The second row shows the problems solved by any of the two most powerful strategies, including the REX strategy that uses exclusive cut. This combination is better than the combination of all previous cut strategies in the first row on all datasets except for chainy, where it is only two problems behind. Combining all strategies (row 3) clearly boosts the number of solved problems compared to the previously available strategies (row 1), namely 11.2% for bushy, 4.8% for TPTP, 9.8% for Miz40, and 7.2% for FS-top.

In conclusion, the new strategies do not only prove more problems, but the problems they solve are also sufficiently complementary from the problems solved by previously available strategies. This makes the new strategies attractive in portfolio modes.

## 6.2. Proof analysis

I compare the proofs of the complete, REX, and REI strategies, similar to Otten’s comparison of the complete and REI strategies [10, sec. 4.2].

There are two indicators for the quality of a cut strategy C2 with respect to a more complete cut strategy C1: which percentage of C1 proofs C2 finds, and how many more inferences C1 takes to find these proofs. When two problems are solved identically by C2 and C1, the additional backtracking done by C1 is superfluous for the proof. The fewer inferences C2 takes to find identical proofs, the more likely it is that C2 also finds proofs which are out of reach for C1 in a given time limit.

Table 4 shows the number of problems for which two strategies find identical proofs. To understand how these numbers emerge, let us consider the chainy dataset, comparing the complete (C1 = None) and the REX (C2 = REX) strategies. Here, Table 2 shows us that C1 solves

**Table 4**  
**Percentage of problems solved by C1 that are identically solved by C2.**

C1	C2	TPTP	bushy	chainy	Miz40	FS-top
None	REX	84.5	66.5	89.4	77.8	81.0
None	REI	68.3	46.7	57.7	54.2	67.4
REX	REI	63.3	40.8	59.2	50.1	66.6

**Table 5**  
**Ratio between sum of inferences taken by C1 and inferences taken by C2, for problems identically solved by C1 and C2.**

C1	C2	TPTP	bushy	chainy	Miz40	FS-top
None	REX	4.4	37.0	9.9	37.4	19.8
None	REI	4.2	55.4	32.0	54.6	28.8
REX	REI	3.3	4.0	<b>8.4</b>	2.4	2.2

**Table 6**  
**Average of ratios between inferences taken by C1 and inferences taken by C2, for problems identically solved by C1 and C2.**

C1	C2	TPTP	bushy	chainy	Miz40	FS-top
None	REX	7.3	22.1	5.9	62.8	12.5
None	REI	41.5	28.1	23.7	78.9	32.0
REX	REI	<b>39.6</b>	2.9	<b>6.9</b>	3.7	2.8

208 and C2 solves 294 chainy problems. C2 finds for 186 of the 208 problems solved by C1 the same proof as C1, which amounts to the 89.4% given in Table 4. Note that C2 solves 203 of the 208 problems that C1 solves, which means that for 17 problems, it finds different proofs than C1.

Of all proofs found by the complete strategy, REX finds between 89.4% (chainy) to 66.5% (bushy), whereas REI finds only between 68.3% (TPTP) and 46.7% (bushy). REI also finds only between 66.6% (FS-top) and 40.8% (bushy) of the proofs found by REX. This shows that there are significantly fewer proofs requiring unrestricted backtracking (no cut) than proofs requiring backtracking that replaces root steps (REX).

We are now going to analyse how much different strategies reduce the search space. For this, we will compare the number of inferences taken by two strategies when they find the same proofs. Given two strategies C1 and C2, we can construct an  $I$  as follows: if C1 and C2 found the same proof  $p$  for a problem, then  $(p, n_1, n_2) \in I$ , where  $n_1$  and  $n_2$  are the number of inferences taken by C1 and C2. Table 5 shows the ratio of the sum of all inferences, calculated by

$$\frac{\sum_{(p, n_1, n_2) \in I} n_1}{\sum_{(p, n_1, n_2) \in I} n_2},$$

and Table 6 shows the average of the ratios of inferences, calculated by

$$\sum_{(p,n_1,n_2) \in I} \frac{n_1}{n_2} \div |I|.$$

In both cases, the higher a value, the more C2 reduces the search space with respect to C1.

Let us look at Table 5. For example, on the Miz40 dataset, for all problems identically solved by REX and the complete strategy, the sum of inferences by the complete strategy is 37.4 times the sum of inferences by the REX strategy. This is the highest ratio for REX with respect to the complete strategy. REX also achieves on Miz40 the largest increase of solved problems compared to the complete strategy (+74.5%). Conversely, on TPTP, where REX shows the smallest inference ratio (4.4), REX also least improves the number of solved problems (+22.8%). On most datasets, the ratios between REI and REX are significantly smaller than the ratios between REX/REI and the complete strategy; for example, on the Miz40 dataset, REI reduces inferences compared to REX only by 2.4, whereas REX and REI reduce inferences compared to the complete strategy by 37.4 and 54.6. This indicates that REX and REI are much “closer” to each other than to the complete strategy. Notable exceptions are the chainy and TPTP datasets, where the ratios between the complete strategy and REX are quite similar to the ratios between REX and REI.<sup>4</sup> Interestingly, these are the datasets where REX proves fewer (chainy) or only few more (TPTP) problems than REI. On datasets like chainy that contain many problems with unusually many axioms, implying a larger explosion of the search space, a more aggressive cut such as REI turns out to be beneficial. In general, we can observe that REX yields the best results when REX greatly reduces the search space with respect to the complete strategy and REI slightly reduces the search space with respect to REX.

In summary, REX is successful because it conserves a considerable amount of existing proofs, while sufficiently reducing the number of inferences in order to find new proofs.

### 6.3. Comparison with other leanCoP implementations

I evaluate meanCoP and two other implementations of leanCoP, namely leanCoP 2.1 using the Prolog compiler ECLiPSe 5.10, and fleanCoP, which is a reimplementaion of leanCoP in OCaml using streams [14]. All evaluated connection provers in this section use a single strategy with conjecture-directed search, non-definitional translation into CNF, and restricted backtracking, i.e. REI.<sup>5</sup> Care is taken that leanCoP-REI, fleanCoP-REI, and meanCoP-REI perform the same inferences.

Table 7 shows the runtime of different leanCoP implementations on sets of problems solved by the original leanCoP. The meanCoP prover is between 27.7 (FS-top) and 6.5 (TPTP) times faster than the original leanCoP and between 11.1 (Miz40) and 2.4 (TPTP) times faster than its OCaml reimplementaion using streams.

Table 8 shows that the higher performance of meanCoP translates to a vastly increased number of proven problems. The largest improvement can be seen on the chainy dataset, where

<sup>4</sup>This can be also seen in Table 6, where the ratio between REX and REI on TPTP is unusually high and clearly exceeds the ratio between REX and the complete strategy.

<sup>5</sup>This amounts to running meanCoP with `--conj --cuts rei`, leanCoP with `SET=' [nodef, conj, cut] '`, and fleanCoP with `-schedule 0 -nodefcnf`.

**Table 7****Prover runtime in seconds for problems solved by leanCoP-REI.**

Prover	TPTP	bushy	chainy	Miz40	FS-top
leanCoP-REI	1299.7	461.9	319.1	9308.7	2451.6
fleanCoP-REI	488.1	190.9	69.8	3845.6	657.2
meanCoP-REI	200.0	17.3	29.0	347.9	88.5

**Table 8****Number of solved problems for different leanCoP implementations.**

Prover	TPTP	bushy	chainy	Miz40	FS-top
leanCoP-REI	1673	606	182	11243	3664
fleanCoP-REI	1859	670	289	12204	3980
meanCoP-REI	1988	730	341	13562	4267

leanCoP, fleanCoP and meanCoP (all using the REI strategy) prove 182, 289 (+58.8%), and 341 (+87.4% compared to leanCoP and +18.0% compared to fleanCoP) problems, respectively.

#### 6.4. Comparison with other provers

I compare meanCoP with several non-connection provers that I previously evaluated in joint work with Kaliszyk and Urban [14, table 2]. In particular, I evaluate Vampire 4.0 [21] and E 2.0 [22], which performed best in the first-order category of CASC-J8 [16]. Vampire and E are written in C++ and C, respectively, implement the superposition calculus, and perform premise selection with SInE [23]. Furthermore, Vampire integrates several SAT solvers [24], and E automatically determines proof search settings for a given problem. I run E with `--auto-schedule` and Vampire with `--mode casc`. In addition, I evaluate the ATP Metis 2.3 (release 20171005) [25]: It implements the ordered paramodulation calculus (having inference rules for equality just like the superposition calculus), but is considerably smaller than Vampire and E and is implemented in Standard ML.

I also evaluate two versions of leanCoP 2.1: First, I evaluate leanCoP 2.1 with strategy scheduling, which will be simply called “leanCoP” in this section. Running leanCoP with a timeout of 10 seconds runs about 10 different search strategies for one second each. Second, I evaluate the first strategy in the strategy schedule of leanCoP 2.1 which searches using restricted backtracking until a path limit of 7, then switches to a complete search with unrestricted backtracking. I call this strategy “leanCoP-CC7”. Unlike all other evaluated connection provers with a single strategy, leanCoP-CC7 does not use conjecture-directed search; furthermore, it uses definitional translation into CNF for the conjecture of the input problem. Finally, I evaluate an adapted version of leanCoP-CC7 with conjecture-directed search and with non-definitional translation into CNF. I call this strategy “leanCoP-NCCC7”.<sup>6</sup>

<sup>6</sup>leanCoP-CC7 and leanCoP-NCCC7 amount to running leanCoP with `SET='[cut,comp(7)]'` and `SET='[nodef,conj,cut,comp(7)]'`, respectively.

**Table 9**  
**Number of solved problems by different provers.**

Prover	TPTP	bushy	chainy	Miz40	FS-top
Vampire	4404	1253	656	30341	6358
E	3664	1167	287	26003	7382
Metis	1376	500	75	18519	3537
leanCoP-CC7	1749	635	154	13121	3892
leanCoP-NCCC7	1752	651	188	11637	4188
leanCoP	1917	673	196	13636	4373
meanCoP-REI	1988	730	<b>341</b>	13562	4267
meanCoP-REX	2126	850	294	16135	4994


**Table 10**  
**Number of problems solved by meanCoP-REX, but not by another prover.**

Prover	TPTP	bushy	chainy	Miz40	FS-top
Vampire	33	19	21	88	1038
E	190	65	92	1803	720
Metis	972	390	222	3111	2463
leanCoP	383	207	104	3394	908
Any above	15	12	16	15	176

Table 9 shows the results: Vampire proves most problems on all datasets except for FS-top, where E prevails. On the chainy dataset, meanCoP proves more problems than E, which is likely due to the conjecture-directed search. Metis proves the fewest problems, except on the Miz40 dataset, where it proves more problems than any connection prover, but less than Vampire and E. leanCoP-CC7 proves more problems than Metis on all datasets except for Miz40, but proves fewer problems than leanCoP (with strategy scheduling) on all datasets. meanCoP-REI proves more problems than leanCoP on all datasets but Miz40 and FS-top, and meanCoP-REX proves more problems than leanCoP on all datasets.

Table 10 shows for several provers  $P$  how many problems meanCoP-REX can solve that were not solved by  $P$ . For example, it shows that meanCoP-REX proves 1038 FS-top problems that were not solved by Vampire, which solves 6358 problems in total. The last line in Table 10 shows the number of problems that meanCoP-REX solves which no other prover in the table solves.

Figure 6 shows for several provers the number of problems on the bushy dataset proved up to a certain time. meanCoP-REX proves considerably more problems than Vampire in the first 10 milliseconds, namely 428 versus a single one. However, Vampire catches up after about 50 milliseconds, leaving all other provers behind. leanCoP and leanCoP-REI solve their first problem about 50 milliseconds after any other prover, which is due to the relatively high start-up time caused by the compilation of the prover at each run. After 50 milliseconds, the order between the provers remains stable. The curves for the provers without strategy scheduling (leanCoP-REI, meanCoP-REI, meanCoP-REX) flatten with time, whereas the curves for Vampire



**Figure 6: Evolution of number of solved bushy problems over time.**

and leanCoP shows several “bumps” due to strategy scheduling. The average time used to solve a problem is 0.53 seconds for meanCoP-REI, 0.64 seconds for meanCoP-REX, 0.76 seconds for leanCoP-REI, 0.87 seconds for leanCoP, and 0.99 seconds for Vampire.

## 7. Related Work

The MaLeCoP prover by Urban et al. [26] and the FEMaLeCoP prover by Kaliszyk and Urban [27] were among the first to use machine learning to guide connection proof search. These provers order the applicable extension steps in prover states by Naive Bayesian probabilities that are inferred from previous proofs. Like leanCoP, they use depth-first search, iterative deepening, and backtracking, which makes such provers likely to benefit from advances in backtracking strategies as presented in this work.

Other connection provers have moved away more from leanCoP’s traditional backtracking-based search. I developed monteCoP in joint work with Kaliszyk and Urban [14], Kaliszyk et al. developed rlCoP [28], and Olšák et al. developed follow-up work to rlCoP [29]. All these provers use machine-learned policies to explore the search space, with Monte Carlo Tree Search taking the role that backtracking plays in leanCoP. For that reason, such provers can probably not directly profit from this work.

We evaluate FEMaLeCoP and monteCoP on the bushy dataset, using 60 seconds timeout, definitional clausification and the REI strategy. Comparing the non-learning with the learning versions of the provers, the increase in number of solved problems is from 563 to 601 for monteCoP (+6.7%) and from 577 to 592 for FEMaLeCoP (+2.6%) [14, table 8], thus far below the increase of 16.4% gained in the current work.

Kaliszyk et al. evaluate rlCoP on the Miz40 dataset, where it proves 16108 problems after 10 iterations of training. Although I evaluate meanCoP on the same dataset, where meanCoP proves 16134 problems, it is unfortunately difficult to compare the results for two reasons: First,

Kaliszyk et al. limit the number of inferences instead of the time allotted to the prover. Second, most inferences performed by rlCoP end up in prover states that are not actually explored, due to not being chosen by Monte Carlo Tree Search.

Another line of work extends connection provers with native support for equality. Rawson's lazyCoP is a connection prover based on Paskevich's connection tableaux calculus with lazy paramodulation [30, 31]. It supports first-order logic with equality. Given that lazyCoP does not use backtracking to control the search, it seems unlikely that exclusive cut could be integrated in this system.

Otten's ileanCoP for intuitionistic logic [3] and MleanCoP for modal logic [4], as well as nanoCoP for nonclausal proof search [5], could all integrate exclusive cut seamlessly.

## 8. Conclusion

I introduced a new kind of cut on extension steps called exclusive cut, which discards all alternatives to solve a literal, except for derivations starting with a different extension step. I implemented the described techniques in a new prover called meanCoP. Evaluating meanCoP on several first-order problem datasets yielded that a combination of cut on reduction steps and exclusive cut on extension steps (REX) improves the number of solved problems compared to the previous best strategy by up to 19%.

## Acknowledgements

I am grateful to the anonymous CADE, TABLEAUX, and AReCCa reviewers as well as to Jasmin Blanchette, Mathias Fleury, Cezary Kaliszyk, and Petar Vukmirović for their comments on drafts of this paper. This work has been supported by the Schrödinger grant (J 4386) of the Austrian Science Fund (FWF).

## References

- [1] W. Bibel, Automated theorem proving, 2nd Edition, Artificial intelligence, Vieweg, 1987. URL: <https://www.worldcat.org/oclc/16641802>.
- [2] P. B. Andrews, Theorem proving via general matings, J. ACM 28 (1981) 193–214. URL: <https://doi.org/10.1145/322248.322249>. doi:10.1145/322248.322249.
- [3] J. Otten, leanCoP 2.0 and ileanCoP 1.2: High performance lean theorem proving in classical and intuitionistic logic (system descriptions), in: A. Armando, P. Baumgartner, G. Dowek (Eds.), Automated Reasoning, 4th International Joint Conference, IJCAR 2008, Sydney, Australia, August 12-15, 2008, Proceedings, volume 5195 of *Lecture Notes in Computer Science*, Springer, 2008, pp. 283–291. URL: [https://doi.org/10.1007/978-3-540-71070-7\\_23](https://doi.org/10.1007/978-3-540-71070-7_23). doi:10.1007/978-3-540-71070-7\_23.
- [4] J. Otten, MleanCoP: A connection prover for first-order modal logic, in: S. Demri, D. Kapur, C. Weidenbach (Eds.), Automated Reasoning - 7th International Joint Conference, IJCAR 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna,

- Austria, July 19-22, 2014. Proceedings, volume 8562 of *Lecture Notes in Computer Science*, Springer, 2014, pp. 269–276. URL: [https://doi.org/10.1007/978-3-319-08587-6\\_20](https://doi.org/10.1007/978-3-319-08587-6_20). doi:10.1007/978-3-319-08587-6\_20.
- [5] J. Otten, nanocop: Natural non-clausal theorem proving, in: C. Sierra (Ed.), Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017, ijcai.org, 2017, pp. 4924–4928. URL: <https://doi.org/10.24963/ijcai.2017/695>. doi:10.24963/ijcai.2017/695.
- [6] C. Kaliszyk, J. Urban, J. Vyskočil, Certified connection tableaux proofs for HOL Light and TPTP, in: X. Leroy, A. Tiu (Eds.), Proceedings of the 2015 Conference on Certified Programs and Proofs, CPP 2015, Mumbai, India, January 15-17, 2015, ACM, 2015, pp. 59–66. URL: <https://doi.org/10.1145/2676724.2693176>. doi:10.1145/2676724.2693176.
- [7] M. Färber, C. Kaliszyk, Certification of nonclausal connection tableaux proofs, in: S. Cerrito, A. Popescu (Eds.), Automated Reasoning with Analytic Tableaux and Related Methods - 28th International Conference, TABLEUX 2019, London, UK, September 3-5, 2019, Proceedings, volume 11714 of *Lecture Notes in Computer Science*, Springer, 2019, pp. 21–38. URL: [https://doi.org/10.1007/978-3-030-29026-9\\_2](https://doi.org/10.1007/978-3-030-29026-9_2). doi:10.1007/978-3-030-29026-9\_2.
- [8] R. Letz, G. Stenz, Model elimination and connection tableau procedures, in: J. A. Robinson, A. Voronkov (Eds.), Handbook of Automated Reasoning (in 2 volumes), Elsevier and MIT Press, 2001, pp. 2015–2114. URL: <https://doi.org/10.1016/b978-044450813-3/50030-8>. doi:10.1016/b978-044450813-3/50030-8.
- [9] J. Otten, W. Bibel, leanCoP: lean connection-based theorem proving, *J. Symb. Comput.* 36 (2003) 139–161. URL: [https://doi.org/10.1016/S0747-7171\(03\)00037-3](https://doi.org/10.1016/S0747-7171(03)00037-3). doi:10.1016/S0747-7171(03)00037-3.
- [10] J. Otten, Restricting backtracking in connection calculi, *AI Commun.* 23 (2010) 159–182. URL: <https://doi.org/10.3233/AIC-2010-0464>. doi:10.3233/AIC-2010-0464.
- [11] C. Kaliszyk, Efficient low-level connection tableaux, in: H. de Nivelle (Ed.), Automated Reasoning with Analytic Tableaux and Related Methods - 24th International Conference, TABLEUX 2015, Wrocław, Poland, September 21-24, 2015. Proceedings, volume 9323 of *Lecture Notes in Computer Science*, Springer, 2015, pp. 102–111. URL: [https://doi.org/10.1007/978-3-319-24312-2\\_8](https://doi.org/10.1007/978-3-319-24312-2_8). doi:10.1007/978-3-319-24312-2\_8.
- [12] B. Stroustrup, Foundations of C++, in: H. Seidl (Ed.), Programming Languages and Systems - 21st European Symposium on Programming, ESOP 2012, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2012, Tallinn, Estonia, March 24 - April 1, 2012. Proceedings, volume 7211 of *Lecture Notes in Computer Science*, Springer, 2012, pp. 1–25. URL: [https://doi.org/10.1007/978-3-642-28869-2\\_1](https://doi.org/10.1007/978-3-642-28869-2_1). doi:10.1007/978-3-642-28869-2\_1.
- [13] D. A. Plaisted, S. Greenbaum, A structure-preserving clause form translation, *J. Symb. Comput.* 2 (1986) 293–304. URL: [https://doi.org/10.1016/S0747-7171\(86\)80028-1](https://doi.org/10.1016/S0747-7171(86)80028-1). doi:10.1016/S0747-7171(86)80028-1.
- [14] M. Färber, C. Kaliszyk, J. Urban, Machine learning guidance for connection tableaux, *J. Autom. Reason.* 65 (2021) 287–320. URL: <https://doi.org/10.1007/s10817-020-09576-7>. doi:10.1007/s10817-020-09576-7.
- [15] G. Sutcliffe, The TPTP problem library and associated infrastructure - from CNF to



- TH0, TPTP v6.4.0, *J. Autom. Reason.* 59 (2017) 483–502. URL: <https://doi.org/10.1007/s10817-017-9407-7>. doi:10.1007/s10817-017-9407-7.
- [16] G. Sutcliffe, The 8th IJCAR automated theorem proving system competition - CASC-J8, *AI Commun.* 29 (2016) 607–619. URL: <https://doi.org/10.3233/AIC-160709>. doi:10.3233/AIC-160709.
- [17] J. Alama, T. Heskes, D. Kühlwein, E. Tsvitsivadze, J. Urban, Premise selection for mathematics by corpus analysis and kernel methods, *J. Autom. Reason.* 52 (2014) 191–213. URL: <https://doi.org/10.1007/s10817-013-9286-5>. doi:10.1007/s10817-013-9286-5.
- [18] C. Kaliszyk, J. Urban, MizAR 40 for MizAR 40, *J. Autom. Reason.* 55 (2015) 245–256. URL: <https://doi.org/10.1007/s10817-015-9330-8>. doi:10.1007/s10817-015-9330-8.
- [19] J. Urban, MPTP - motivation, implementation, first experiments, *J. Autom. Reason.* 33 (2004) 319–339. URL: <https://doi.org/10.1007/s10817-004-6245-1>. doi:10.1007/s10817-004-6245-1.
- [20] T. C. Hales, M. Adams, G. Bauer, D. T. Dang, J. Harrison, T. L. Hoang, C. Kaliszyk, V. Magron, S. McLaughlin, T. T. Nguyen, T. Q. Nguyen, T. Nipkow, S. Obua, J. Pleso, J. Rute, A. Solovyev, A. H. T. Ta, T. N. Tran, D. T. Trieu, J. Urban, K. K. Vu, R. Zumkeller, A formal proof of the Kepler conjecture, *Forum of Mathematics, Pi* 5 (2017). doi:10.1017/fmp.2017.1.
- [21] L. Kovács, A. Voronkov, First-order theorem proving and Vampire, in: N. Sharygina, H. Veith (Eds.), *Computer Aided Verification - 25th International Conference, CAV 2013, Saint Petersburg, Russia, July 13-19, 2013. Proceedings*, volume 8044 of *Lecture Notes in Computer Science*, Springer, 2013, pp. 1–35. URL: [https://doi.org/10.1007/978-3-642-39799-8\\_1](https://doi.org/10.1007/978-3-642-39799-8_1). doi:10.1007/978-3-642-39799-8\_1.
- [22] S. Schulz, System description: E 1.8, in: K. L. McMillan, A. Middeldorp, A. Voronkov (Eds.), *Logic for Programming, Artificial Intelligence, and Reasoning - 19th International Conference, LPAR-19, Stellenbosch, South Africa, December 14-19, 2013. Proceedings*, volume 8312 of *Lecture Notes in Computer Science*, Springer, 2013, pp. 735–743. URL: [https://doi.org/10.1007/978-3-642-45221-5\\_49](https://doi.org/10.1007/978-3-642-45221-5_49). doi:10.1007/978-3-642-45221-5\_49.
- [23] K. Hoder, A. Voronkov, Sine qua non for large theory reasoning, in: N. Bjørner, V. Sofronie-Stokkermans (Eds.), *Automated Deduction - CADE-23 - 23rd International Conference on Automated Deduction, Wrocław, Poland, July 31 - August 5, 2011. Proceedings*, volume 6803 of *Lecture Notes in Computer Science*, Springer, 2011, pp. 299–314. URL: [https://doi.org/10.1007/978-3-642-22438-6\\_23](https://doi.org/10.1007/978-3-642-22438-6_23). doi:10.1007/978-3-642-22438-6\_23.
- [24] A. Biere, I. Dragan, L. Kovács, A. Voronkov, Experimenting with SAT solvers in Vampire, in: A. F. Gelbukh, F. Castro-Espinoza, S. N. Galicia-Haro (Eds.), *Human-Inspired Computing and Its Applications - 13th Mexican International Conference on Artificial Intelligence, MICAI 2014, Tuxtla Gutiérrez, Mexico, November 16-22, 2014. Proceedings, Part I*, volume 8856 of *Lecture Notes in Computer Science*, Springer, 2014, pp. 431–442. URL: [https://doi.org/10.1007/978-3-319-13647-9\\_39](https://doi.org/10.1007/978-3-319-13647-9_39). doi:10.1007/978-3-319-13647-9\_39.
- [25] J. Hurd, First-order proof tactics in higher-order logic theorem provers, in: M. Archer, B. D. Vito, C. Muñoz (Eds.), *Design and Application of Strategies/Tactics in Higher Order Logics (STRATA)*, number NASA/CP-2003-212448 in *NASA Technical Reports*, 2003, pp. 56–68. URL: <http://www.gilith.com/research/papers>.
- [26] J. Urban, J. Vyskočil, P. Štěpánek, MaLeCoP machine learning connection prover, in: K. Brunnler, G. Metcalfe (Eds.), *Automated Reasoning with Analytic Tableaux and Re-*

- lated Methods - 20th International Conference, TABLEAUX 2011, Bern, Switzerland, July 4-8, 2011. Proceedings, volume 6793 of *Lecture Notes in Computer Science*, Springer, 2011, pp. 263–277. URL: [https://doi.org/10.1007/978-3-642-22119-4\\_21](https://doi.org/10.1007/978-3-642-22119-4_21). doi:10.1007/978-3-642-22119-4\_21.
- [27] C. Kaliszyk, J. Urban, FEMaLeCoP: Fairly efficient machine learning connection prover, in: M. Davis, A. Fehnker, A. McIver, A. Voronkov (Eds.), *Logic for Programming, Artificial Intelligence, and Reasoning - 20th International Conference, LPAR-20 2015*, Suva, Fiji, November 24-28, 2015, Proceedings, volume 9450 of *Lecture Notes in Computer Science*, Springer, 2015, pp. 88–96. URL: [https://doi.org/10.1007/978-3-662-48899-7\\_7](https://doi.org/10.1007/978-3-662-48899-7_7). doi:10.1007/978-3-662-48899-7\_7.
- [28] C. Kaliszyk, J. Urban, H. Michalewski, M. Olśák, Reinforcement learning of theorem proving, in: S. Bengio, H. M. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, R. Garnett (Eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018*, December 3-8, 2018, Montréal, Canada, 2018, pp. 8836–8847. URL: <http://papers.nips.cc/paper/8098-reinforcement-learning-of-theorem-proving>.
- [29] M. Olśák, C. Kaliszyk, J. Urban, Property invariant embedding for automated reasoning, in: G. D. Giacomo, A. Catalá, B. Dilkina, M. Milano, S. Barro, A. Bugarín, J. Lang (Eds.), *ECAI 2020 - 24th European Conference on Artificial Intelligence, 29 August-8 September 2020, Santiago de Compostela, Spain, August 29 - September 8, 2020 - Including 10th Conference on Prestigious Applications of Artificial Intelligence (PAIS 2020)*, volume 325 of *Frontiers in Artificial Intelligence and Applications*, IOS Press, 2020, pp. 1395–1402. URL: <https://doi.org/10.3233/FAIA200244>. doi:10.3233/FAIA200244.
- [30] G. Sutcliffe, Proceedings of the 10th IJCAR ATP system competition (CASC-J10), 2020. URL: <http://www.tptp.org/CASC/J10/Proceedings.pdf>.
- [31] A. Paskevich, Connection tableaux with lazy paramodulation, *J. Autom. Reason.* 40 (2008) 179–194. URL: <https://doi.org/10.1007/s10817-007-9089-7>. doi:10.1007/s10817-007-9089-7.