

# Learning to Identify Useful Lemmas from Failure\*

Michael Rawson<sup>1</sup>, Christoph Wernhard<sup>2</sup>, and Zsolt Zombori<sup>3,4</sup>

<sup>1</sup> TU Wien, Austria [michael@rawsons.uk](mailto:michael@rawsons.uk)

<sup>2</sup> University of Potsdam, Germany [info@christophwernhard.com](mailto:info@christophwernhard.com)

<sup>3</sup> Alfréd Rényi Institute of Mathematics, Hungary [zombori@renyi.hu](mailto:zombori@renyi.hu)

<sup>4</sup> Eötvös Loránd University, Budapest, Hungary

**Introduction** We investigate learning to identify useful lemmas for ATP, where usefulness is defined in terms of 1) reducing proof search and 2) shortening the length of the overall proof. How can ATP performance be improved by the generation and selection of useful lemmas? In particular, we raise the question of what one can learn from a failed proof attempt. We present a proposal about how to learn from failed proof attempts of a single problem. This is in sharp contrast with previous works that rely on a large corpus of problems and aim to improve performance based on success obtained with easier problems. By way of motivation, we argue that human mathematicians learn from failed attempts as well.

**Restriction to a class of problems with accessible and simple proof structures** Interested in novel techniques, we work with a restricted class of first-order problems, *condensed detachment* (CD) problems [9, 10, 7], due to Carew A. Meredith [5]. Inference steps can be characterized by detachment (modus ponens) combined with unification. Proof structures are particularly simple and accessible: full binary trees, or terms with a binary function symbol  $D$ , which we call  $D$ -terms. Constants in these terms label axioms. As examples of  $D$ -terms consider  $1$ , a constant representing a use of the axiom labeled by  $1$ ;  $D(1, 1)$ , representing a detachment step applied to axiom  $1$  as major and minor premise; or  $D(1, D(1, 1))$ , representing a proof with two detachment steps. These proof terms are closely related to proof structures of the connection method [3, 4].

**Proof search and data extraction** We rely on the theorem prover SGCD [8] which performs proof search by structure enumeration of binary trees (interwoven with formula unification), until a suitable  $D$ -term is found. Enumeration can be axiom-driven, i.e. starting from axiom set  $As$ , producing  $D$ -terms that represent complete proofs of unit lemmas. We can also enumerate goal-driven, starting from conjecture  $C$  and obtaining partial proof trees of  $C$ . In practice we interleave goal-driven and axiom-driven phases. Using the idea of Hindsight Experience Replay [1], we can “pretend” that both sorts of failed proof attempts are in fact successful: In the axiom-driven case, we change the goal conjecture to the one that was actually proven and in the goal-driven case, we change the axioms to include whatever is needed to complete the proof. Hence, we end up enumerating complete proof trees of “some” problems. We note that the idea of Hindsight Experience Replay has already been applied to theorem proving in [2] in the context of training a policy model to guide saturation-style forward reasoning.

Given a proof tree  $P'$  of some formula  $C'$  from axiom set  $As'$ , any connected subgraph  $S'$  of  $P'$  can be considered as the proof of a lemma candidate  $L$ . If  $S'$  is a full tree, it proves a unit lemma, which is the formula associated with its root. Otherwise, it proves a Horn clause, whose head is the root formula of  $S'$  and whose body consists of the open leaves of  $S'$ . If we

---

\*Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project-ID 457292495, by the North-German Supercomputing Alliance (HLRN), by the ERC grant CoG ARTIST 101002685, by the Hungarian National Excellence Grant 2018-1.2.1-NKP-00008, the Hungarian Artificial Intelligence National Laboratory Program (RRF-2.3.1-21-2022-00004), the ELTE TKP 2021-NKTA-62 funding scheme and the COST action CA20111.

can measure how useful lemma  $L$  is for proving  $C'$  given axioms  $As'$ , this can serve as a useful training signal for a guidance model. For the utility measure  $U$ , there are easy-to-compute logical candidates, such as the compression in tree size, tree height, DAG size etc. A more refined measure is obtained if we reprove  $C'$  with the lemma  $L$  added to the axioms  $As'$  and observe how the number of inference steps changes.<sup>1</sup> This is slower to compute, but takes into account the particularities of the prover, hence providing more focused guidance. In practice, we find that the best performance is obtained by reproving and then normalising the inference step reduction into  $[-1, 1]$ , where  $-1$  means that the problem could not be solved within the original inference limit and  $1$  is assigned to the lemma that yields the greatest speedup. We end up with tuples  $\langle C', As', L, U \rangle$  to learn from.

**Interwoven proof search and training** During the proof search of conjecture  $C$  from axiom set  $As$ , we keep track of all produced proof trees  $P'$  and collect  $\langle C', As', L, U \rangle$  tuples, forming a training dataset  $D$ . Once proof search is over, we fit a neural lemma selector to  $D$ . This neural model  $M(\text{conjecture}, \text{axioms}, \text{lemma})$  predicts the utility of the input lemma for proving the conjecture from the axioms. Model  $M$  is next evaluated on all collected lemmas, along with the original conjecture and axioms, i.e. we compute pairs

$$\{\langle L, U \rangle \mid \langle \_, \_, L, \_ \rangle \in D, U = M(C, As, L)\}$$

Lemmas with the top  $k$  utilities are selected, where  $k$  is a hyperparameter to be tuned. The selected lemmas are added to the problem as axioms<sup>2</sup> and we can start proof search again.

**Experiments** We have performed experiments on training a unit lemma selector from successful proof attempts, which is explained in [6]. A recent addition to this work is extracting training data from failed attempts, for which we now present preliminary results.

In our first experiment, we fit separate models for each problem based on the prover’s failed attempt on the given problem. The model is then used to select lemmas for a second round of proof search. We use a set of 312 CD problems extracted or derived from the TPTP library and allow a single strategy for SGCD (`provecd_sgcd_s1`) with 10 seconds per problem time limit. As Table 1 shows, the base prover solves 176 problems, to which 11 problems are added thanks to the added lemmas. This is 6% improvement.

Learn from	Iter 0	Iter 1	Total
failure	176	187	187 (+11)

Table 1: Fitting a lemma selector model for each failed proof attempt

Next, we train a single model to guide lemma selection from all proof attempts on the entire problem set. We compare three scenarios: 1) learning only from successful attempts, 2) learning only from failed attempts and 3) learning from both success and from failure. Based on a set of 411 CD problems, using SGCD with strategy `provecd_sgcd_s1` and 10 s time limit, Table 2 shows that learning brings significantly more improvement when learning from failure (15%), compared to learning from success (12%). It is important to emphasize that this is not because the training signal from failed attempts would be more valuable, but rather because there is much more data that can be extracted from failure than from success.

<sup>1</sup>*Number of inference steps* refers to the underlying Prolog engine, not a calculus. It provides a measure similar to the time required for proving, but independent from hardware or system load fluctuations.

<sup>2</sup>If the prover has some special mechanism for handling lemmas, they need not be treated as axioms.

Learn from	Iteration					Total
	0	1	2	3	4	
success	199	203	206	216	205	222 (+23)
failure	199	211	219	209	205	229 (+30)
both	199	212	207	<b>223</b>	200	<b>230 (+31)</b>

Table 2: Fitting a single lemma selector model for the entire training set. Single strategy mode.

In our last experiment, we use a portfolio of 9 SGCD strategies instead of a single one. Table 3 shows that the base prover is stronger and the improvement due to learning is somewhat less (11% vs 15%), but it still remains significant.

Learn from	Iteration					Total
	0	1	2	3	4	
both	236	257	246	249	244	263 (+27)

Table 3: Fitting a single lemma selector model for the entire training set. Portfolio mode.

**Conclusion** Our work has two key takeaway messages. First, identifying useful lemmas can greatly aid automated proof search. Second, our preliminary results suggest that a lot of useful training signal can be extracted from failed proof attempts. This latter is in sharp contrast with most existing learning assisted provers (and in fact with most reinforcement learning algorithms), which focus on learning from successful attempts. We hope that future work will uncover the full potential of these observations.

## References

- [1] Andrychowicz, M., Wolski, F., Ray, A., Schneider, J., Fong, R., Welinder, P., McGrew, B., Tobin, J., Pieter Abbeel, O., Zaremba, W.: Hindsight experience replay. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 30. Curran Associates, Inc. (2017), [https://proceedings.neurips.cc/paper\\_files/paper/2017/file/453fadbd8a1a3af50a9df4df899537b5-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2017/file/453fadbd8a1a3af50a9df4df899537b5-Paper.pdf)
- [2] Aygün, E., Anand, A., Orseau, L., Glorot, X., Mcaleer, S.M., Firoiu, V., Zhang, L.M., Precup, D., Mourad, S.: Proving theorems using incremental learning and hindsight experience replay. In: Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., Sabato, S. (eds.) *Proceedings of the 39th International Conference on Machine Learning*. *Proceedings of Machine Learning Research*, vol. 162, pp. 1198–1210. PMLR (17–23 Jul 2022), <https://proceedings.mlr.press/v162/aygun22a.html>
- [3] Bibel, W.: *Automated Theorem Proving*. Vieweg, Braunschweig (1982). <https://doi.org/10.1007/978-3-322-90102-6>, second edition 1987
- [4] Bibel, W., Otten, J.: From Schütte’s formal systems to modern automated deduction. In: Kahle, R., Rathjen, M. (eds.) *The Legacy of Kurt Schütte*, chap. 13, pp. 215–249. Springer (2020). [https://doi.org/10.1007/978-3-030-49424-7\\_13](https://doi.org/10.1007/978-3-030-49424-7_13)
- [5] Prior, A.N.: Logicians at play; or Syll, Simp and Hilbert. *Australasian Journal of Philosophy* **34**(3), 182–192 (1956). <https://doi.org/10.1080/00048405685200181>

- [6] Rawson, M., Wernhard, C., Zombori, Z., Bibel, W.: Lemmas: Generation, selection, application. In: Ramanayake, R., Urban, J. (eds.) TABLEAUX 2023. LNAI, vol. 14278, pp. 153–174 (2023). [https://doi.org/10.1007/978-3-031-43513-3\\_9](https://doi.org/10.1007/978-3-031-43513-3_9)
- [7] Ulrich, D.: A legacy recalled and a tradition continued. *J. Autom. Reasoning* **27**(2), 97–122 (2001). <https://doi.org/10.1023/A:1010683508225>
- [8] Wernhard, C.: CD Tools — Condensed detachment and structure generating theorem proving (system description). *CoRR* **abs/2207.08453** (2023). <https://doi.org/10.48550/ARXIV.2207.08453>
- [9] Wernhard, C., Bibel, W.: Learning from Łukasiewicz and Meredith: Investigations into proof structures. In: Platzer, A., Sutcliffe, G. (eds.) CADE 28. LNCS (LNAI), vol. 12699, pp. 58–75. Springer (2021). [https://doi.org/10.1007/978-3-030-79876-5\\_4](https://doi.org/10.1007/978-3-030-79876-5_4)
- [10] Wernhard, C., Bibel, W.: Investigations into proof structures. *CoRR* **abs/2304.12827** (2023). <https://doi.org/10.48550/ARXIV.2304.12827>, submitted