



# Make E Smart Again (Short Paper)

Zarathustra Amadeus Goertzel<sup>(✉)</sup>

Czech Technical University in Prague, Prague, Czechia  
zariuq@gmail.com

**Abstract.** In this work in progress, we demonstrate a new use-case for the ENIGMA system. The ENIGMA system using the XGBoost implementation of gradient boosted decision trees has demonstrated high capability to learn to guide the E theorem prover’s inferences in real-time. Here, we strip E to the bare bones: we replace the KBO term ordering with an identity relation as the minimal possible ordering, disable literal selection, and replace evolved strategies with a simple combination of the clause weight and FIFO (first in first out) clause evaluation functions. We experimentally demonstrate that ENIGMA can learn to guide E as well as the smart, evolved strategies even without these standard automated theorem prover functionalities. To this end, we experiment with XGBoost’s meta-parameters over a dozen loops.

## 1 Introduction: Making E Stupid and Then Smart Again

State-of-the-art saturation-based automated theorem provers (ATPs) for first-order logic (FOL), such as E and Vampire [12], employ the *given clause algorithm* [13], translating the input FOL problem  $T \cup \{-C\}$  (background theory and negated conjecture) into a refutationally equivalent set of clauses. The search for a contradiction is performed maintaining sets of *processed* ( $P$ ) and *unprocessed* ( $U$ ) clauses (the *proof state*  $\Pi$ ). The algorithm repeatedly selects a *given clause*  $g$  from  $U$ , moves  $g$  to  $P$ , and extends  $U$  with all clauses inferred with  $g$  and  $P$ . This process continues until a contradiction is found,  $U$  becomes empty, or a resource limit is reached.

Historically, *term ordering*, together with *literal selection*, is used to guarantee the completeness of the proof search [1] and to “tame the growth of the search space and help steer proof search” [5]. Term ordering ensures that rewriting happens in only one direction, toward smaller terms. Literal selection limits the inferences done with each given clause  $g$  to the selected literals, which slows down the growth of the search space and reduces redundant inferences.

E includes a *strategy* language of *clause evaluation functions*, made up of weight and priority functions, to heuristically guide the proof search. In this

---

Supported by the ERC Consolidator grant no. 649043 AI4REASON and by the Czech project AI&Reasoning CZ.02.1.01/0.0/0.0/15\_003/ 0000466 and the European Regional Development Fund.

work, I use two algorithmically invented [6,7] strategies, E1 and E2<sup>1</sup>, that use many sophisticated clause evaluation functions, the Knuth-Bendix ordering (KBO6), literal selection, and other E heuristics.

The ENIGMA [4,8–10] system with the XGBoost [2] implementation of gradient boosted decision trees has recently demonstrated high capability to learn to guide the E [14] theorem prover’s inferences in real-time. ENIGMA uses the XGBoost model as a clause evaluation function to recommend clauses for selection based on clause and conjecture features. In particular, after several proving and learning iterations, its performance on the 57880 problems from the Mizar40 [11] benchmark improved by 70% (= 25397/14933) [10] over the strategy E1 used for the initial proving phase.

In this work, E is stripped to the bare bones by disabling term ordering and literal selection. KBO6 is replaced with an identity relation as the minimal possible ordering (called IDEN – an addition to E<sup>2</sup>). While this frees E to do inferences in any order, E can no longer perform rewriting inferences. The strategy E1 is replaced with the simple combination of the clause weight and FIFO (first in first out) evaluation functions. E is thus practically reduced to a basic superposition prover, without advanced heuristics, rewriting, or completeness guarantees. We call this strategy E0:

```
--definitional-cnf=24 --prefer-initial-clauses -tIDEN
--restrict-literal-comparisons -WNoSelection
-H'(5*Clauseweight(ConstPrio,1,1,1),1*FIFOweight(ConstPrio))'
```

E0 solves only 3872 of the Mizar40 problems in 5 s compared to 14526 for E1. The first research question is the extent to which ENIGMA with this basic prover can learn ATP guidance completely on its own. The second is to what extent ENIGMA’s learning can be boosted with data from strong strategies and models. That is, I explore how smart machine learning can become in this *zero-strategy* setting. The more general related question is to what extent can machine learning replace the sophisticated human-invented theorem-proving body of wisdom used in today’s ATPs for restricting advanced proof calculi.

## 2 Experiments

We evaluate ENIGMA with the basic strategy, E0, in several scenarios and over two datasets of different sizes. All experiments are run with 5 s per problem<sup>3 4</sup>.

<sup>1</sup> Strategies E1 and E2 are displayed in the appendix.

<sup>2</sup> The E version used in this paper can be found at <https://github.com/zariuq/eprover/tree/identity-order>, and the library for running ENIGMA with E can be found at <https://github.com/zariuq/enigmatic>.


<sup>3</sup> As a rule of thumb, E solves most problems within a few seconds or *not for a very long time*.

<sup>4</sup> All the experiments are run on the same hardware unless otherwise specified: Intel(R) Xeon(R) Gold 6140 CPU @ 2.30GHz with 188GB RAM.

ENIGMA has so far been used in two ways: *coop* combines the learned model with some standard E strategy equally (50:50) while *solo* only uses the learned model for choosing the given clauses. The best results have been achieved by MaLAREa-style [16] looping: that is, an ENIGMA model is trained and run with E (loop 0), then the resulting data are added to the initial training data and a new ENIGMA model is trained (loop 1).

In this work, ENIGMA trains with both *solo* and *coop* data. I present results from *solo* runs because they represent the most minimal setting.

### 2.1 Small Data (2000 Problems)



The E evaluations and XGBoost training can take a long time on the full Mizar40 dataset, so 2000 randomly sampled problems are used to test meta-parameters on. Each XGBoost model consists of  $T$  decision trees of depth  $D$ , the most important training meta-parameters in addition to the learning rate ( $\eta = 0.2$ ). In previous work with ENIGMA,  $T$  and  $D$  were fixed for all loops of learning. Here we try to vary the values of  $T$  and  $D$  during 16 loops. Let  $S_{D,T}$  denote the experiment with specific  $T$  and  $D$ . Of the many protocols tested, the following are included in the plot of solved problems (above): *Fives* ( $S_{5,100}$ ), *Nines* ( $S_{9,100}$ ), *Thirteens* ( $S_{13,200}$ ), *Sixteens* ( $S_{16,100}$ ).

We also experiment with adaptively setting the meta-parameters as the number of training examples increases according to the following protocols:

- *Inc* ( $S_{[3,33],100}$ ) increases  $D$  by 2 from 3 to 33 and keeps  $T = 100$  fixed.
- *32\_inc* ( $S_{32,[50,250]}$ ) fixes  $D = 32$  and gradually increases  $T$  from 50 to 250.
- *Inc2* ( $S_{[3,33],*}$ ) gradually decreases  $T$  from 150 to 50, varying the value intuitively<sup>5</sup>.

<sup>5</sup> Precise details of intuitively set parameters can be seen in the appendix.

- $Inc3(S_{[3,33],[50,250]})$  aims to be more systematic and steps  $T$  from 50 to 250.
- $Dec3(S_{[3,33],[250,50]})$  decreases  $T$  from 250 to 50.

At the 16th loop  $Inc$ 's performance is best, solving 299 problems, doubling the performance of E0 (152). However  $Inc2$  and  $Inc3$  solve 298 problems and  $32\_inc$  solves 291 problems. The conclusion is that simple protocols work well so long as  $T$  or  $D$  is incremented adaptively rather than fixed.

## 2.2 Big Data (57880 Problems)



These experiments are done on the large benchmark of 57880 Mizar40 [11] problems from the MPTP dataset [15]. E1 and E2 are two strong E strategies that solve 14526 and 12788 problems.

- **Experiment 1** is done with  $D = 9$  and  $T = 200$  and uses our previously trained model that allowed us to solve 25562 problems when cooperating with E1 in our previous experiments [10]. This strong model, which hashes the features into  $32768 (2^{15})$  buckets [3, Sect. 3.4], is used with E0 now.
- **Experiment 2**'s parameters were intuitively toggled during the looping as in  $Inc3$ , and a feature size of  $2^{16}$  is used. Exp. 2 uses training data from E1 and E2 for additional guidance up to the 4th loop (and then stops including them in the training data based on the assumption they may confuse learning).
- **Experiment 3** sets  $T$  and  $D$  according to protocol  $Inc3$ . Exp. 3 only learns from E run with E0 and trains on the GPU, which requires the feature size to be reduced to 256.
- **Experiment 4** mimics Exp. 3 but uses E1 and E2 data for training (up to the 4th loop).
- **Experiment 5** further tests boosting with data from an E0 ENIGMA model that proved 9759 problems and an E1 that proved 21542 problems.

Tree depth is intuitively varied among 32, 512, and 1000, the number of trees is varied among 2, 100, 200, and 32. The feature vector size starts at  $2^{14}$  and is decreased to allow the data to fit on the RAM, down to 32 ( $= 2^5$ ).

As seen in the figure, the strong model does not help much in guiding E without ordering or selection in Exp. 1. Exp. 2 learns gradually and catches up with Exp. 1, but seems to plateau around 10,000. Surprisingly the pure Exp. 3 learns fast with the small feature size, but plateaus and drops in performance (perhaps due to overfitting). Exp. 4 indicates that guidance is useful and surpasses E2 with 13805 in round 13. Exp. 5 solves 15990 problems, showing that ENIGMA can take E0 beyond the smart strategies with appropriate parameters and boosting. This is a great improvement over the 3872 problems solved by E0.

### 3 Conclusion

ENIGMA can learn to guide the E prover effectively even without smart strategies and term orderings. The models confer a 256% increase over the naive E0 after 13 rounds of the proving/learning loop, and even trained without guidance data, a 121% increase.

The experiments indicate that machine learning can be used to fully control an ATP’s guidance, learning to replace orderings, heuristic strategies, and deal with the increase in generated clauses without literal selection. However the combination of ENIGMA and standard ATP heuristics still significantly outperforms ENIGMA alone.

Given the large symmetry-breaking impact of these methods in classical ATP, future work includes, e.g., training the guidance in such a way that redundant (symmetric) inferences are not done by the trained model once it has committed to a certain path. This probably means equipping the learning with more history and knowledge of the proof state in the saturation-style setting. ENIGMAWatch [4] may aid with symmetry breaking by focusing the proof search on particular proof paths. Additional work is needed to isolate the factors in Exp. 5’s performance, and determine the most effective boosting methods in addition to increasing  $D$  and  $T$  with training loops. Ablation studies should be done to discover the impact of term ordering and literal selection individually on E and ENIGMA’s performance. Perhaps term ordering alone is sufficient to train good ENIGMA models.

Running ENIGMA without term ordering and other restrictions is important because it may allow us to combine training data from different strategies, and it may allow ENIGMA to find novel proofs.

**Acknowledgments.** The research topic was proposed by Jan Jakubuv and Josef Urban, and further discussed with them, Martin Suda, and Thomas Tan. I also thank the AITP’20 anonymous referees for their comments on the first extended abstract of this work.

## A Strategies

Strategy E1 is:

```
--definitional-cnf=24 --split-aggressive --simul-paramod
--forward-context-sr --destructive-er-aggressive --destructive-er
--prefer-initial-clauses -tKBO -winvfreqrank -c1 -Ginvfreq -F1
--delete-bad-limit=150000000 -WSelectMaxLComplexAvoidPosPred
-H' (1*ConjectureTermPrefixWeight(DeferSOS,1,3,0.1,5,0.0.1,1,4),
1*ConjectureTermPrefixWeight(DeferSOS,1,3,0.5,100,0,0.2,0.2,4),
1*Refinedweight(PreferWatchlist,4,300,4,4,0.7),
1*RelevanceLevelWeight2(PreferProcessed,0,1,2,1,1,1,200,200,2.5,9999.9,9999.9),
1*StaggeredWeight(DeferSOS,1),
1*SymbolTypeweight(DeferSOS,18,7,-2,5,9999.9,2,1.5),
2*Clauseweight(PreferWatchlist,20,9999,4),
2*ConjectureSymbolWeight(DeferSOS,9999,20,50,-1,50,3,3,0.5),
2*StaggeredWeight(DeferSOS,2))'
```

Strategy E2 is:

```
--definitional-cnf=24 --split-aggressive --split-reuse-defs
--simul-paramod --forward-context-sr --destructive-er-aggressive
--destructive-er --prefer-initial-clauses -tKBO -winvfreqrank
-c1 -Ginvfreq -F1 --delete-bad-limit=150000000
-WSelectMaxLComplexAvoidPosPred -H' (
3*ConjectureRelativeSymbolWeight(PreferUnitGroundGoals,0.1,100,100,50,100,0.3,1.5,1.5),
4*FIFOweight(PreferNonGoals),
5*RelevanceLevelWeight2(ConstPrio,1,0,2,1,50,-2,-2,100,0.2,3,4))'
```

## B Additional Protocol Details

In this section I include the details for *intuitively toggled* protocols. Protocol *Inc2* is as follows:

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Depth	3	5	7	9	11	13	15	17	19	21	23	25	27	29	31	33
Trees	150	150	150	100	100	100	75	50	75	100	150	75	100	150	75	100

The protocol for Exp. 2 is as follows:

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21
D	4	5	6	7	8	9	10	11	12	13	14	15	16	16	32	9	16	32	64	24	25	32
T	50	150	160	170	180	190	200	200	200	200	210	220	225	225	225	300	300	225	150	250	250	250

The protocol for Exp. 5 requires some explanation. The motivation is to see how far E0 can be taken, even if the methods are too CPU-intensive for a thorough grid search.

Exp. 2 and Exp. 4 demonstrate the utility of boosting. Thus to create better boosting data I trained ENIGMA for 10 loops with strategies E1 through E12 and used this as boosting data for the first 4 of 10 loops of training. In addition to training E0, and in the spirit of ablation studies, I also trained ENIGMA models for E0 with KBO ordering (and no literal selection) and for E0 with KBO ordering and restricted literal comparisons. The motivation is that these versions may serve as a bridge between standard E and the basic E0.

Then I used these results to boost an ENIGMA model in loop 0, and trained based on this for 10 loops, proving 9759 problems.

Finally this data and the data from a loop 3 ENIGMA model trained with E1 is used to boost E0 with the following meta-parameters:

	0	1	2	3	4	5	6	7	8	9	10	11
Depth	512	512	32	1000	32	1000	32	1000	32	1000	1000	100
Trees	2	2	100	100	200	100	200	32	300	32	32	32
Feature size	16384	8192	4096	28	4096	28	4096	32	2048	64	32	128

## References

1. Bachmair, L., Ganzinger, H.: Rewrite-based equational theorem proving with selection and simplification. *J. Log. Comput.* **3**(4), 217–247 (1994)
2. Chen, T., Guestrin, C.: XGBoost: a scalable tree boosting system. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2016)*, pp. 785–794. ACM, New York (2016)
3. Chvalovský, K., Jakubův, J., Suda, M., Urban, J.: ENIGMA-NG: efficient neural and gradient-boosted inference guidance for E. In: Fontaine, P. (ed.) *CADE 2019*. LNCS (LNAI), vol. 11716, pp. 197–215. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-29436-6\\_12](https://doi.org/10.1007/978-3-030-29436-6_12)
4. Goertzel, Z., Jakubův, J., Urban, J.: ENIGMAWatch: proofWatch meets ENIGMA. In: Cerrito, S., Popescu, A. (eds.) *TABLEAUX 2019*. LNCS (LNAI), vol. 11714, pp. 374–388. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-29026-9\\_21](https://doi.org/10.1007/978-3-030-29026-9_21)
5. Hoder, K., Reger, G., Suda, M., Voronkov, A.: Selecting the selection. In: Olivetti, N., Tiwari, A. (eds.) *IJCAR 2016*. LNCS (LNAI), vol. 9706, pp. 313–329. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-40229-1\\_22](https://doi.org/10.1007/978-3-319-40229-1_22)
6. Jakubův, J., Urban, J.: Hierarchical invention of theorem proving strategies. *AI Commun.* **31**(3), 237–250 (2018)
7. Jakubův, J., Urban, J.: Extending E prover with similarity based clause selection strategies. In: Kohlhase, M., Johansson, M., Miller, B., de Moura, L., Tompa, F. (eds.) *CICM 2016*. LNCS (LNAI), vol. 9791, pp. 151–156. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-42547-4\\_11](https://doi.org/10.1007/978-3-319-42547-4_11)
8. Jakubův, J., Urban, J.: ENIGMA: efficient learning-based inference guiding machine. In: Geuvers, H., England, M., Hasan, O., Rabe, F., Teschke, O. (eds.) *CICM 2017*. LNCS (LNAI), vol. 10383, pp. 292–302. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-62075-6\\_20](https://doi.org/10.1007/978-3-319-62075-6_20)

9. Jakubův, J., Urban, J.: Enhancing ENIGMA given clause guidance. In: Rabe, F., Farmer, W.M., Passmore, G.O., Youssef, A. (eds.) CICM 2018. LNCS (LNAI), vol. 11006, pp. 118–124. Springer, Cham (2018). [https://doi.org/10.1007/978-3-319-96812-4\\_11](https://doi.org/10.1007/978-3-319-96812-4_11)
10. Jakubuv, J., Urban, J.: Hammering Mizar by learning clause guidance. In: Harrison, J., O’Leary, J., Tolmach, A. (eds.) 10th International Conference on Interactive Theorem Proving, (ITP 2019) of LIPIcs, 9–12 September 2019, Portland, OR, USA, vol. 141, pp. 34:1–34:8. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2019)
11. Kaliszyk, C., Urban, J.: MizAR 40 for Mizar 40. *J. Autom. Reasoning* **55**(3), 245–256 (2015). <https://doi.org/10.1007/s10817-015-9330-8>
12. Kovács, L., Voronkov, A.: First-order theorem proving and VAMPIRE. In: Sharygina, N., Veith, H. (eds.) CAV 2013. LNCS, vol. 8044, pp. 1–35. Springer, Heidelberg (2013). [https://doi.org/10.1007/978-3-642-39799-8\\_1](https://doi.org/10.1007/978-3-642-39799-8_1)
13. Overbeek, R.A.: A new class of automated theorem-proving algorithms. *J. ACM* **21**(2), 191–200 (1974)
14. Schulz, S., Cruanes, S., Vukmirović, P.: Faster, higher, stronger: E 2.3. In: Fontaine, P. (ed.) CADE 2019. LNCS (LNAI), vol. 11716, pp. 495–507. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-29436-6\\_29](https://doi.org/10.1007/978-3-030-29436-6_29)
15. Urban, J.: MPTP 0.2: Design, implementation, and initial experiments. *J. Autom. Reasoning* **37**(1–2), 21–43 (2006)
16. Urban, J., Sutcliffe, G., Pudlák, P., Vyskočil, J.: MaLAREa SG1 - machine learner for automated reasoning with semantic guidance. In: Armando, A., Baumgartner, P., Dowek, G. (eds.) IJCAR 2008. LNCS (LNAI), vol. 5195, pp. 441–456. Springer, Heidelberg (2008). [https://doi.org/10.1007/978-3-540-71070-7\\_37](https://doi.org/10.1007/978-3-540-71070-7_37)