

# The Open-Source LearnLib

## A Framework for Active Automata Learning

Malte Isberner<sup>1</sup>(✉), Falk Howar<sup>2</sup>, and Bernhard Steffen<sup>1</sup>

<sup>1</sup> TU Dortmund University, 44221 Dortmund, Germany  
{malte.isberner, steffen}@cs.tu-dortmund.de

<sup>2</sup> IPSSE/TU Clausthal, 38678  
Clausthal-Zellerfeld, Germany  
falk.howar@tu-clausthal.de



**Abstract.** In this paper, we present *LearnLib*, a library for active automata learning. The current, open-source version of *LearnLib* was completely rewritten from scratch, incorporating the lessons learned from the decade-spanning development process of the previous versions of *LearnLib*. Like its immediate predecessor, the open-source *LearnLib* is written in Java to enable a high degree of flexibility and extensibility, while at the same time providing a performance that allows for large-scale applications. Additionally, *LearnLib* provides facilities for visualizing the progress of learning algorithms in detail, thus complementing its applicability in research and industrial contexts with an educational aspect.

## 1 Introduction

Active automata learning, from its early beginnings almost thirty years ago [6], inspired a number of applications in quite a number of fields (see [19] for a survey). However, it took almost a decade for the software verification and testing community to recognize its value of being able to provide models of black-box systems for the plethora of model-based tools and techniques. More precisely, it was not until the seminal works of Peled *et al.* [36], employing automata learning to model check black-box systems, and Steffen *et al.* [18], who used it to automatically generate test cases for legacy computer-telephony integrated systems, that this use case of automata learning was discovered.

Since then, however, active automata learning has enjoyed quite a success story, having been used as a valuable tool in areas as diverse as automated GUI testing [13], fighting bot-nets [12], or tpestate analysis [5,41]. Most of these works, however, used their custom, one-off implementation of the well-known  $L^*$  learning algorithm [6], and hence invested relatively little effort for optimizations, or using a more sophisticated (but harder to implement and lesser-known) algorithm altogether.<sup>1</sup>

<sup>1</sup> An elaborate discussion on the theoretical aspects of active automata learning, as well as on the challenges that arise in practice, are outside the scope of this paper.

We refer the interested reader to [39] for an introduction focusing on these matters.

We started developing the *LearnLib*<sup>2</sup> library to provide researchers and practitioners with a reusable set of components to facilitate and promote the use of active automata learning, and to enable access to cutting-edge automata learning technology. From the beginnings of the development of *LearnLib*, started in 2003, until now, more than a decade has passed. In these years, many lessons were learned on what makes for a usable, efficient and practically feasible product that fulfills this goal (cf. [25,35,37]).

These lessons form the basis of the new *LearnLib* presented in this paper. The new *LearnLib* is not just an overhaul of the prior version, but completely re-written from scratch. It provides a higher level of abstraction and increased flexibility, while simultaneously being the fastest version of *LearnLib* to date (cf. Sect. 4). As a service to the community and to encourage contributions by and collaborations with other research groups, we decided to make *LearnLib* available under an open-source license (the *Apache License, version 2.0*<sup>3</sup>). In the remainder of this paper we highlight two aspects that we address with *LearnLib*.

**Advanced Features.** This is what we consider the strongest case for preferring a comprehensive automata learning framework such as *LearnLib* over a custom implementation. While implementing the original version of  $L^*$  is not a challenging task, the situation is different for more refined active learning algorithms, such as Rivest & Schapire’s [38], Kearns & Vazirani’s [30] or even the very recent TTT algorithm [28]. While we found these algorithms to consistently outperform  $L^*$ , the latter remains the most widely used. Also, several other advanced optimizations such as query parallelization or efficient query caches are typically neglected. Through *LearnLib*’s modular design, changing filters, algorithm parameters or even the whole algorithm is a matter of a few lines of code, yielding valuable insights on how different algorithms perform on certain input data. Many of these features rely on *AutomataLib*, the standalone finite-state machine library that was developed for *LearnLib*, which provides a rich toolbox of data structures and algorithms for finite-state machines. The design of *AutomataLib* is presented in Sect. 2, while Sect. 3 provides a more comprehensive overview of *LearnLib*’s feature set.

**Performance.** The implementation of a learning algorithm comes with many performance pitfalls. Even though in most cases the time taken by the actual learning algorithm is an uncritical aspect (compared to the time spent in executing queries, which may involve, e.g., network communication), it should be kept as low as reasonably possible. Besides, an efficient management of data structures is necessary to enable learning of large-scale systems without running into out-of-memory conditions or experiencing huge performance slumps. In *LearnLib*, considerable effort was spent on efficient implementations while providing a conveniently high level of abstraction. This will be detailed in Sect. 4.

Finally, we conclude the paper by briefly discussing envisioned future work in Sect. 5.

<sup>2</sup> <http://www.learnlib.de>.

<sup>3</sup> <https://www.apache.org/licenses/LICENSE-2.0>.




Fig. 1. Architecture of *AutomataLib*

## 2 AutomataLib

One of the main architectural changes of the open-source *LearnLib* is that it uses a dedicated, stand-alone library for representing and manipulating automata, called *AutomataLib*.<sup>4</sup> While *AutomataLib* is formally independent of *LearnLib*, its development process is closely intertwined with the one of *LearnLib*. For this reason, *AutomataLib* mainly focuses on deterministic automata, even though selected classes of non-deterministic automata are supported as well (e.g., NFAs).

*AutomataLib* is divided into an *abstraction layer*, *automata implementations*, and *algorithms* (cf. Fig. 1). The abstraction layer comprises a set of *Java* interfaces to represent various types of automata and graphs, organized in a complex, fine-grained type hierarchy. Furthermore, these interfaces were designed in a generic fashion, to integrate existing, third-party automata implementations into *AutomataLib*'s interface landscape with as little effort and run-time overhead as possible. For instance, a proof-of-concept adapter for the BRICS automaton library<sup>5</sup> could be realized in as little as 20 lines of *Java* code.

Adapters like for the BRICS library form one part of the implementation layer. The other part are generic automaton implementations, e.g., for DFAs or Mealy machines, that provide good defaults for general setups, and are also used by most algorithms in *LearnLib* to store hypotheses.

Sample algorithms shipped with *AutomataLib* include minimization, equivalence testing, or visualization (via *GraphVIZ*'s<sup>6</sup> *dot* tool). The set of functionalities will be continuously extended, with a strong focus on functionality either directly required in *LearnLib*, or desirable in a typical automata learning application context.

An important aspect is that the algorithms operate solely on the abstraction layer, meaning that they are implementation agnostic: they can be used with a (wrapped) BRICS automaton as well as with other automaton implementations. Furthermore, the generic design enables a high degree of code reuse: the minimization (or equivalence checking) algorithm can be used for both DFA and Mealy machines, as it is designed to only require a *deterministic automaton*, instead of a concrete machine type (or even implementation).

<sup>4</sup> <http://www.automatalib.net/>.

<sup>5</sup> <http://www.brics.dk/automaton/>.

<sup>6</sup> <http://www.graphviz.org/>.


### 3 LearnLib

*LearnLib* provides a set of components to apply automata learning in practical settings, or to develop or analyze automata learning algorithms. These can be grouped into three main classes: learning algorithms, methods for finding counterexamples (so-called *Equivalence Queries*), and infrastructure components.

**Learning Algorithms.** *LearnLib* features a rich set of learning algorithms, covering the majority of algorithms which have been published (and many beyond that). Care was taken to develop the algorithms in a modular and parameterizable fashion, which allows us to use a single “base” algorithm to realize several algorithms described in the literature, e.g., by merely exchanging the involved counterexample analysis strategy. Perhaps the best example for this is the  $L^*$  algorithm [6], which can be configured to pose as **Maler & Pnueli’s** [31], **Rivest & Schapire’s** [38], or **Shahbaz’s** [26] algorithm, **Suffix1by1** [26], or variants thereof. Other base algorithms available in *LearnLib* are the **Observation Pack** [21] algorithm, **Kearns & Vazirani’s** [30] algorithm, the **DHC** [34] algorithm, and the **TTT** [28] algorithm. These, too, can be adapted in the way they handle counterexamples, e.g., by linear search, binary search (à la Rivest & Schapire), or exponential search [29]. With the exception of **DHC**, all these algorithms are available in both DFA and Mealy versions. Furthermore, *LearnLib* features the  $NL^*$  algorithm for learning NFAs [8].

**Equivalence Tests and Finding Counterexamples.** Once a learning algorithm converges to a stable hypothesis, a *counterexample* is needed to ensure further progress. In the context of active learning, the process of searching for a counterexample is also referred to as an *equivalence query*. “Perfect” equivalence queries are possible only when a model of the target system is available. In this case, *LearnLib* uses Hopcroft and Karp’s near-linear equivalence checking algorithm [4, 20] available through *AutomataLib*. In black-box scenarios, equivalence queries can be approximated using conformance tests. *AutomataLib* provides implementations of the W-method [14] and the Wp-method [16], two of the few conformance tests that can find missing states. Often, the cheapest and fastest way of approximating equivalence queries is searching for counterexamples directly: *LearnLib* implements a random walk (only for Mealy machines), randomized generation of tests, and exhaustive generation of test inputs (up to a certain depth).

**Infrastructure.** The third class of components that come with *LearnLib* provide useful infrastructure functionality such as a logging facility, an import/export mechanism to store and load hypotheses, or utilities for gathering statistics. An important component for many practical applications are (optimizing) *filters*, which pre-process the queries posed by the learning algorithm. A universally useful example of such a filter is a *cache filter* [32], eliminating duplicate queries that most algorithms pose. Other examples include a parallelization component that distributes queries across multiple workers [22], a mechanism for reusing system states to reduce the number of resets [7], and for prefix-closed systems [32].



**Fig. 2.** Performance comparison between the new *LearnLib* and *libalf*. Left: run-time of the classic  $L^*$  algorithm on a series of randomly generated automata with state counts between 10 and 1000. Right: run-time of five comparable algorithms from *LearnLib* and *libalf* on a DFA with 500 states.

For a learning algorithm to work in practice, some interface to the system under learning (SUL) needs to be available. While this is generally specific to the SUL itself, *LearnLib* provides *SUL adapters* for typical cases, e.g., *Java* classes, web-services, or processes that are interfaced with via standard I/O.

## 4 Evaluation

We are aware of two other open-source automata learning libraries that provide implementations of textbook algorithms, complemented by own developments:

**libalf**<sup>7</sup>. The *Automata Learning Framework* [9], was developed primarily at the RWTH Aachen. It is available under LGPLv3 and written in C++. Its active development seems to have ceased; the last version was released in April 2011.

**AIDE**<sup>8</sup>. The *Automata-Identification Engine*, under active development, is available under the open-source license LGPLv2.1 and written in C#.

The ambitions behind *LearnLib* go further: It is specifically designed to easily compose new custom learning algorithms on the basis of components for counterexample analysis, approximations of equivalence queries, as well as connectors to real life systems. Moreover, *LearnLib* provides a variety of underlying data structures, and various means for visualizing the algorithm and its statistics. This does not only facilitate the construction of highly performant custom solutions, but also provides a deeper understanding of the algorithms' characteristics. The latter has been essential, e.g., for designing the TTT algorithm [28], which almost uniformly outperforms all the previous algorithms.

**Performance.** As we have mentioned earlier, the open-source *LearnLib* is the fastest version of *LearnLib* to date, and moreover the fastest automata learning

<sup>7</sup> <http://libalf.informatik.rwth-aachen.de/>.

<sup>8</sup> <http://aide.codeplex.com/>.

implementation that we are aware of. We have conducted a preliminary performance evaluation, comparing the new *LearnLib* to *libalf* and the old, closed-source version of *LearnLib* (which we will refer to as *JLearn* in order to avoid confusion). A visualization of some of the results comparing *LearnLib* and *libalf* is shown in Fig. 2. It can be clearly seen that in the considered setting, *LearnLib* is more than an order of magnitude faster than *libalf* (even though the former is implemented in Java while the latter is implemented in C++). More importantly, the gap grows with the size of the system to be learned. In our experiments, the open-source *LearnLib* also outperformed *JLearn* on a similar scale. More detailed performance data can be found on the *LearnLib* website.<sup>9</sup>

**Applications.** The performance data demonstrates that *LearnLib* provides a robust basis for fast and scalable active automata learning solutions. Consequently, in its ten years of continued development, *LearnLib* has been used in a number of research and industry projects, of which we briefly present some of the more recent ones. A more complete list can be found on the *LearnLib* homepage. *LearnLib* has been used to infer models of smart card readers [11] and of bank cards [3]. The models were used to verify security properties of these systems. In [2, 15], models of communication protocols are inferred using *LearnLib*. The models are used to verify the conformance of protocol implementations to the corresponding specifications. At TU Dortmund, *LearnLib* has been used in an industry project [40] to generate models of a web application. The models were used to test regressions in the user interface and in the business processes of this application. The authors of [33] propose a method for generating checking circuits for functions implemented in FPGAs. The method uses models of the functions that are inferred with *LearnLib*. *LearnLib* is also used in other tools: PSYCO [17, 23] is a tool for generating precise interfaces of software components developed at CMU and NASA Ames. The tool combines concolic execution and active automata learning (i.e., *LearnLib*). Tomte, developed at the Radboud University of Nijmegen [1] leverages regular inference algorithms provided by *LearnLib* to infer richer classes of models by simultaneously inferring sophisticated abstractions (or “mappers”).

## 5 Conclusion

In this paper we have presented *LearnLib*, a versatile open-source library of active automata learning algorithms. *LearnLib* is unique in its modular design, which has furthered the development of new learning algorithms (e.g., the TTT algorithm [28]) and tools (e.g., Tomte [1] and PSYCO [17, 23]).

While in many aspects the open-source *LearnLib* by far surpasses the capabilities of the previous version, there are two major features which have yet to be ported. The first is *LearnLib Studio* (cf. [35]), a graphical user interface for *LearnLib*, and the second is an extension for learning *Register Automata*. An extension for learning Register Automata with the theory of equality only was

<sup>9</sup> <http://learnlib.de/features/performance>.

available upon request for the old *LearnLib* in binary form [24,27]. We are currently working on a generalized approach [10], which will be included in the open-source release.

## References

1. Aarts, F., Heidarian, F., Kuppens, H., Olsen, P., Vaandrager, F.: Automata learning through counterexample guided abstraction refinement. In: Giannakopoulou, D., Méry, D. (eds.) FM 2012. LNCS, vol. 7436, pp. 10–27. Springer, Heidelberg (2012)
2. Aarts, F., Jonsson, B., Uijen, J., Vaandrager, F.W.: Generating models of infinite-state communication protocols using regular inference with abstraction. *Form. Meth. Syst. Des.* **46**(1), 1–41 (2015)
3. Aarts, F., De Ruiter, J., Poll, E.: Formal models of bank cards for free. In: 2013 IEEE Sixth International Conference on Software Testing, Verification and Validation, Workshops Proceedings, pp. 461–468, Luxembourg, 18–22 Mar 2013
4. Almeida, M., Moreira, N., Reis, R.: Testing the equivalence of regular languages. In: Proceedings Eleventh International Workshop on Descriptive Complexity of Formal Systems, DCFS 2009, pp. 47–57, Magdeburg, Germany, 6–9 Jul 2009. <http://dx.doi.org/10.4204/EPTCS.3.4>
5. Alur, R., Cerný, P., Madhusudan, P., Nam, W.: Synthesis of interface specifications for java classes. In: Palsberg, J., Abadi, M. (eds.) Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2005, pp. 98–109. ACM, Long Beach, California, USA, 12–14 Jan 2005. <http://doi.acm.org/10.1145/1040305.1040314>
6. Angluin, D.: Learning regular sets from queries and counterexamples. *Inf. Comput.* **75**(2), 87–106 (1987)
7. Bauer, O., Neubauer, J., Steffen, B., Howar, F.: Reusing system states by active learning algorithms. In: Moschitti, A., Scandariato, R. (eds.) ETERNALS 2011. CCIS, vol. 255, pp. 61–78. Springer, Heidelberg (2012)
8. Bollig, B., Habermehl, P., Kern, C., Leucker, M.: Angluin-style learning of NFA. In: Proceedings IJCAI 2009, pp. 1004–1009. IJCAI 2009, San Francisco, CA, USA (2009)
9. Bollig, B., Katoen, J.-P., Kern, C., Leucker, M., Neider, D., Piegdon, D.R.: `libalf`: The automata learning framework. In: Touili, T., Cook, B., Jackson, P. (eds.) CAV 2010. LNCS, vol. 6174, pp. 360–364. Springer, Heidelberg (2010)
10. Cassel, S., Howar, F., Jonsson, B., Steffen, B.: Learning extended finite state machines. In: Giannakopoulou, D., Salaün, G. (eds.) SEFM 2014. LNCS, vol. 8702, pp. 250–264. Springer, Heidelberg (2014)
11. Chalupar, G., Peherstorfer, S., Poll, E., De Ruiter, J.: Automated reverse engineering using lego. In: 8th USENIX Workshop on Offensive Technologies, WOOT 2014, San Diego, CA, USA, 19 Aug 2014
12. Cho, C.Y., Babić, D., Shin, R., Song, D.: Inference and analysis of formal models of botnet command and control protocols. In: Proceedings CCS 2010, pp. 426–440, ACM, Chicago, Illinois, USA (2010)
13. Choi, W., Necula, G., Sen, K.: Guided gui testing of android apps with minimal restart and approximate learning. In: Proceedings of the 2013 ACM SIGPLAN International Conference on Object Oriented Programming Systems Languages & Applications, pp. 623–640. OOPSLA 2013, ACM, New York, NY, USA (2013). <http://doi.acm.org/10.1145/2509136.2509552>

14. Chow, T.S.: Testing software design modeled by finite-state machines. *IEEE Trans. Softw. Eng.* **4**(3), 178–187 (1978)
15. Fiterău-Broștean, P., Janssen, R., Vaandrager, F.: Learning fragments of the TCP network protocol. In: Lang, F., Flammini, F. (eds.) *FMICS 2014*. LNCS, vol. 8718, pp. 78–93. Springer, Heidelberg (2014)
16. Fujiwara, S., Von Bochmann, G., Khendek, F., Amalou, M., Ghedamsi, A.: Test selection based on finite state models. *IEEE Trans. Softw. Eng.* **17**(6), 591–603 (1991)
17. Giannakopoulou, D., Rakamarić, Z., Raman, V.: Symbolic learning of component interfaces. In: Miné, A., Schmidt, D. (eds.) *SAS 2012*. LNCS, vol. 7460, pp. 248–264. Springer, Heidelberg (2012)
18. Hagerer, A., Hungar, H.: Model generation by moderated regular extrapolation. In: Kutsche, R.-D., Weber, H. (eds.) *FASE 2002*. LNCS, vol. 2306, p. 80. Springer, Heidelberg (2002)
19. De la Higuera, C.: A bibliographical study of grammatical inference. *Pattern Recogn.* **38**(9), 1332–1348 (2005). <http://dx.doi.org/10.1016/j.patcog.2005.01.003>
20. Hopcroft, J., Karp, R.: A linear algorithm for testing equivalence of finite automata. Technical report 0, Department of Computer Science, Cornell U, Dec 1971
21. Howar, F.: Active learning of interface programs. Ph.D. thesis, TU Dortmund University (2012). <http://dx.doi.org/2003/29486>
22. Howar, F., Bauer, O., Merten, M., Steffen, B., Margaria, T.: The teachers' crowd: the impact of distributed oracles on active automata learning. In: Hähnle, R., Knoop, J., Margaria, T., Schreiner, D., Steffen, B. (eds.) *ISOLA 2011 Workshops 2011*. CCIS, vol. 336, pp. 232–247. Springer, Heidelberg (2012)
23. Howar, F., Giannakopoulou, D., Rakamarić, Z.: Hybrid learning: interface generation through static, dynamic, and symbolic analysis. In: *Proceedings of the International Symposium on Software Testing and Analysis (ISSTA)*, pp. 268–279, ACM (2013)
24. Howar, F., Steffen, B., Jonsson, B., Cassel, S.: Inferring canonical register automata. In: Kuncak, V., Rybalchenko, A. (eds.) *VMCAI 2012*. LNCS, vol. 7148, pp. 251–266. Springer, Heidelberg (2012)
25. Hungar, H., Niese, O., Steffen, B.: Domain-specific optimization in automata learning. In: Hunt Jr, W.A., Somenzi, F. (eds.) *CAV 2003*. LNCS, vol. 2725, pp. 315–327. Springer, Heidelberg (2003)
26. Irfan, M.N., Oriat, C., Groz, R.: Angluin style finite state machine inference with non-optimal counterexamples. In: *1st International Workshop on Model Inference In Testing* (2010)
27. Isberner, M., Howar, F., Steffen, B.: Learning register automata: from languages to program structures. *Mach. Learn.* **96**(1–2), 65–98 (2014). <http://dx.doi.org/10.1007/s10994-013-5419-7>
28. Isberner, M., Howar, F., Steffen, B.: The TTT algorithm: a redundancy-free approach to active automata learning. In: Bonakdarpour, B., Smolka, S.A. (eds.) *RV 2014*. LNCS, vol. 8734, pp. 307–322. Springer, Heidelberg (2014)
29. Isberner, M., Steffen, B.: An abstract framework for counterexample analysis in active automata learning. In: Clark, A., Kanazawa, M., Yoshinaka, R. (eds.) *Proceedings of the 12th International Conference on Grammatical Inference, ICGI 2014*, Kyoto, Japan, 17–19 Sep 2014. *JMLR Proceedings*, vol. 34, pp. 79–93, <http://JMLR.org> (2014). <http://jmlr.org/proceedings/papers/v34/isberner14a.html>
30. Kearns, M.J., Vazirani, U.V.: *An Introduction to Computational Learning Theory*. MIT Press, Cambridge (1994)



31. Maler, O., Pnueli, A.: On the learnability of infinitary regular sets. *Inf. Comput.* **118**(2), 316–326 (1995)
32. Margaria, T., Raffelt, H., Steffen, B.: Knowledge-based relevance filtering for efficient system-level test-based model generation. *Innov. Syst. Softw. Eng.* **1**(2), 147–156 (2005)
33. Matuova, L., Kastil, J., Kotásek, Z.: Automatic construction of on-line checking circuits based on finite automata. In: 17th Euromicro Conference on Digital System Design, DSD 2014, pp. 326–332, Verona, Italy, 27–29 Aug 2014
34. Merten, M., Howar, F., Steffen, B., Margaria, T.: Automata learning with on-the-fly direct hypothesis construction. In: Hähnle, R., Knoop, J., Margaria, T., Schreiner, D., Steffen, B. (eds.) *ISoLA 2011 Workshops 2011*. CCIS, vol. 336, pp. 248–260. Springer, Heidelberg (2012)
35. Merten, M., Steffen, B., Howar, F., Margaria, T.: Next generation LearnLib. In: Abdulla, P.A., Leino, K.R.M. (eds.) *TACAS 2011*. LNCS, vol. 6605, pp. 220–223. Springer, Heidelberg (2011)
36. Peled, D., Vardi, M.Y., Yannakakis, M.: Black box checking. In: Wu, J., Chanson, S.T., Gao, Q. (eds.) *Proceedings FORTE 1999*, pp. 225–240, Kluwer Academic (1999)
37. Raffelt, H., Steffen, B., Berg, T., Margaria, T.: LearnLib: a framework for extrapolating behavioral models. *Int. J. Softw. Tools Technol. Transf.* **11**(5), 393–407 (2009)
38. Rivest, R.L., Schapire, R.E.: Inference of finite automata using homing sequences. *Inf. Comput.* **103**(2), 299–347 (1993)
39. Steffen, B., Howar, F., Merten, M.: Introduction to active automata learning from a practical perspective. In: Bernardo, M., Issarny, V. (eds.) *SFM 2011*. LNCS, vol. 6659, pp. 256–296. Springer, Heidelberg (2011)
40. Windmüller, S., Neubauer, J., Steffen, B., Howar, F., Bauer, O.: Active continuous quality control. In: *CBSE*, pp. 111–120 (2013)
41. Xiao, H., Sun, J., Liu, Y., Lin, S., Sun, C.: Tzuyu: learning stateful tpestates. In: Denney, E., Bultan, T., Zeller, A. (eds.) *2013 28th IEEE/ACM International Conference on Automated Software Engineering, ASE 2013*, pp. 432–442, IEEE, Silicon Valley, CA, USA, 11–15 Nov 2013. <http://dx.doi.org/10.1109/ASE.2013.6693101>