

Document Version

Final published version

Licence

CC BY

Citation (APA)

Aynekulu Mersha, D. G., Lopuhaä, B., Gietema, H. A., Smit, J., Wind, A. E., Hulshof, A. M., Daenen, K., Khyali, J., von der Thüsen, J. H., & More Authors (2025). In situ pulmonary thrombosis and pulmonary embolus are distinct thrombotic phenotypes in critically ill patients with COVID-19 Acute Respiratory Distress Syndrome. *Journal of Thrombosis and Haemostasis*, 24(2), 662-671. <https://doi.org/10.1016/j.jtha.2025.10.021>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



An Active Learning Algorithm for Bidirectional Deterministic Finite Automata

Simon Dieck^(✉) and Sicco Verwer

Department of Software Technology, TU Delft, Delft, Netherlands
{S.Dieck,S.E.Verwer}@tudelft.nl

Abstract. In this paper, we present an L^* -style algorithm for actively learning a bidirectional deterministic finite automaton (biDFA) in polynomial time using three types of oracles. We show how the W -method for the equivalence oracle can be adapted to our algorithm and present a novel heuristic for choosing the orientation of states. With this algorithm, one can identify automata for a subset of the linear languages that includes but is not limited to the regular languages. Since the equivalence oracle is an important part of the algorithm, we also discuss complexity bounds for different versions of the language equivalence problem for biDFAs. These results, together with our algorithm, also prove complexity bounds for the biDFA minimisation problem. Finally, we provide an implementation of the algorithm and experimentally show its performance with different approximation heuristics.

1 Introduction

A deterministic finite automaton (DFA) is a versatile model that is used across a variety of different fields ranging from applications in biology for parsing DNA [18], in compilers as scanners [12], to cyber security for fingerprinting [9]. While there are several ways of obtaining DFAs, for many fields, it is useful to infer them by interacting with a system whose behaviour the DFA should replicate [7]. This process, normally referred to as active learning, ideally produces a DFA whose behaviour is equivalent to the system from which it was inferred. For example, one can use DFAs inferred from programs that were infected by malicious software to identify security vulnerabilities [2, 5, 6]. Similarly, one can analyse the behaviour of opaque models, like neural networks, by analysing a DFA surrogate model inferred from them instead [17].

One issue that all of these applications face is that DFAs are limited to recognising the regular languages. And while learning procedures for more expressive models exist, often based on a notion of substitutability, these algorithms are computationally significantly more expensive [4, 23].

In this context bidirectional deterministic automata (biDFAs) are interesting. One can think of these models as DFAs with two reading heads. At the beginning, one reading head is placed at the start of the string, and the other one is placed at

the end. In each step, the model decides on which reading head will read the next character. This way, a string is parsed from both ends by heads moving inwards. The process terminates if both heads meet. One can think of this meeting point as a “centre” of a string. For the deterministic version of these models, the decision of which reading head to move is also deterministic, and as such, every word has a fixed centre where the heads meet. Nondeterministic biDFAs can recognise all linear languages, while the deterministic version can recognise a subset of the linear languages that strictly includes the regular languages. Importantly, they can do this without requiring additional memory structures like, for example, pushdown automata [13, 15].

While not significantly more expressive than DFAs, the additional languages biDFA can recognise are such that normally require counting, like $a^n b^n$. They can also recognise languages like palindromes, which notably even deterministic pushdown automata cannot. Further, Dieck and Verwer have shown that on palindrome-like regular languages biDFAs can be exponentially smaller than DFAs [8]. Especially in the field of inferring models from software, one can expect such patterns. If software produces some output at the beginning and end of a recursive function, the output will have a palindrome-like structure. Accordingly, an inference algorithm for biDFAs has potential uses in many fields but is of special interest for software inference.

In this paper, we show that the L^* inference algorithm for DFA can be adapted to infer a biDFA. We show that if an oracle for state orientation, that is which reading head is used in a state, is given, this algorithm has the same runtime guarantees as L^* for DFAs. We also prove that the algorithm produces a biDFA that is correct and minimal with regard to a centre function defined by the oracle. Since L^* requires an equivalence oracle, we analyse the complexity of this problem, which is an open question for biDFAs, and prove results for a special case. The algorithm introduced in this work also proves that the biDFA minimisation problem is NP-complete if the biDFA recognises a regular language and in P if the centre function for the minimised model is equivalent to the centre function of the input model.

In practice, especially when inferring software, one has only access to membership queries. Accordingly, model-checking techniques such as the W -method [3] are often used to approximate an equivalence query heuristically [14]. In this work, we also prove that an adaptation of the W -method heuristic [3] works for biDFAs, giving similar guarantees. We also provide a heuristic that can be used for the state orientation oracle to decide which end of the string to read from in a given state.

Finally, we test the algorithm and heuristics experimentally. Here, we show that with the right setup for the heuristics, the algorithm consistently correctly identifies non-regular languages. This includes a palindrome language and even some that were chosen to be difficult for the heuristic. We also show that it often finds the globally minimal model on those languages.

The full version of this paper which includes all proofs is available together with the code used for the experiments at https://github.com/SimonDieck/biDFA_L_star.

1.1 Related Work

Bidirectional automata, then called biautomata, were first introduced by Klíma and Polák [16]. While these models also read a string from both ends moving inwards and are almost equivalent to biDFAs, Klíma and Polák defined the transition function in such a way that these models could only recognise regular languages. This transition function was then changed to allow non-deterministic transitions by Holzer and Jakobi [13]. They analysed the language theoretical properties of these biautomata and showed that they could recognise all linear languages [13].

Jirásková and Klíma then adjusted the transition function of Holzer and Jakobi once again to be deterministic both in allowing only one transition per character and state and enforcing that all transitions from a single state need to read a character from the same end [15]. These models are equivalent to biDFAs since they implicitly introduce an orientation of the state and otherwise are fully deterministic. Jirásková and Klíma prove most language theoretical properties of these models in their work and show that the class of languages they can recognise is a proper subset of the linear languages and a proper superset of the regular languages [15]. Dieck and Verwer provide a Myhill-Nerode style characterisation of the class of languages recognised by these models [8].

The contributions of these works regarding the language equivalence and minimisation problem, together with our novel contributions are summarised in Table 1.

Table 1. Summary of some problems relating to biDFA. Results in bold are novel results from this paper. NP-hardness of the minimisation problem was proven by Dieck and Verwer [8] and the other known results by Jirásková and Klíma [15].

Given B_1 and B_2 , recognising L_1 and L_2 , defining centres c_1 and c_2	L_1 or L_2 is regular	$c_1 = c_2$	No restrictions
Emptiness of intersection	NL-complete	NL-complete	Undecidable
Language equivalence problem	NL-complete	In NL	Unknown
Minimisation problem	NP-complete	In P	NP-hard

Independently by Nagy very similar bidirectional models have been proposed in the context of DNA processing [18, 19]. Whether these models are equivalent has not been proven yet.

For the active learning algorithm, much of our work is based on the L^* algorithm introduced by Angluin [1]. However, for the counterexample processing we use the method first introduced by Rivest and Schapire [20,21].

2 Background

In this section, we will introduce some notation and results from other works that are used in this work. For notation, we will mostly refer to what was used by Dieck and Verwer [8] since we are focused on the deterministic version of these bidirectional models and will heavily use their characterisation. We assume familiarity with standard definitions and notations regarding formal languages and DFA.

For biDFAs we need to introduce the notion of a centre function c .

Definition 1 (Centre function). *A centre function $c : \Sigma^* \rightarrow \mathbb{N}$ maps a word $w \in \Sigma^*$ to the range $[0, |w|]$, where $|w|$ denotes the length of a word which is equal to the number of characters it contains.*

Given such a centre function, every word can be split into a left and a right half $w = w_{c,l} \cdot w_{c,r}$ s.t. if $w = \sigma_1 \sigma_2 \dots \sigma_{|w|}$ then $w_{c,l} = \sigma_1 \dots \sigma_{c(w)}$ and $w_{c,r} = \sigma_{c(w)+1} \dots \sigma_{|w|}$.

For the definition of a biDFA we use the one given by Dieck and Verwer [8]:

Definition 2 (biDFA). *A bidirectional deterministic finite automaton (biDFA) is a six-tuple $B = (Q_l, Q_r, \Sigma, \delta, q_\lambda, F)$, where $Q = Q_l \cup Q_r$, with $Q_l \cap Q_r = \emptyset$, forms a finite set of states, Σ is an alphabet, $\delta : Q \times \Sigma \rightarrow Q$ a transition function, $q_\lambda \in Q$ the initial state and $F \subseteq Q$ the set of accepting states.*

For this definition to work on words and not just characters, δ is extended to $\delta^* : Q \times \Sigma^* \rightarrow Q$ with the following definition:

$$\delta^*(q, (\sigma_1 \dots \sigma_n)) = \begin{cases} q & \text{if } \sigma_1 \dots \sigma_n = \lambda \\ \delta(q, (\sigma_1 \dots \sigma_n)) & \text{if } n = 1 \\ \delta^*(\delta(q, \sigma_1), (\sigma_2 \dots \sigma_n)) & \text{if } n > 1 \text{ and } q \in Q_l \\ \delta^*(\delta(q, \sigma_n), (\sigma_1 \dots \sigma_{n-1})) & \text{if } n > 1 \text{ and } q \in Q_r \end{cases}$$

The language L accepted by B can then be defined as $\{w \in \Sigma^* \mid \delta^*(q_\lambda, w) \in F\}$. One can think of a biDFA as parsing a word by parsing a left string from left to right and a right string from right to left. Whenever B is in a state in Q_l , δ^* will process the next character from the left. When in a state of Q_r , it will process the next character from the right. With this, one can use a biDFA B to define a centre function c_B . To obtain $c_B(w)$ one can count the number of states in Q_l that were encountered when parsing w with B , except for the final state in which the process terminates. This is exactly the number of characters that were read from the left end of the string.

Dieck and Verwer have shown, given a centre function c and a language L one can define an equivalence relation on Σ^* as follows [8]:

$$u \equiv_{c,L} v \text{ iff for all } m \in \Sigma^* : u_{c,l} \cdot m \cdot u_{c,r} \in L \iff v_{c,l} \cdot m \cdot v_{c,r} \in L$$

In words, two words are considered inter-equivalent w.r.t. c and L if and only if all words that can be produced from them by inserting a word at their centre behave the same with regards to membership in L . As this interjection comes up often, we introduce a shorthand $u \cdot_c m = u_{c,l} \cdot m \cdot u_{c,r}$, where “ \cdot_c ” is a left-associative operator. We will denote the equivalence classes for this by $[w]_{c,L}$ where $[w]_{c,L} = \{w' \in \Sigma^* \mid w' \equiv_{c,L} w\}$.

Finally, there is the important restriction for a centre function of “strict stability” which was introduced by Dieck and Verwer [8]:

Definition 3 (Strict stability). *A centre function c is considered stable if for all words $w \in \Sigma^*$ and all characters $s \in \Sigma$ $c(w) \leq c(w \cdot_c s) \leq c(w) + 1$.*

A stable centre function c is considered “strictly stable” for a language L if for all $u, v \in \Sigma^$ $u \equiv_{c,L} v$ implies that for all $s, s' \in \Sigma$: $c(u) = c(u \cdot_c s) \iff c(v) = c(v \cdot_c s')$*

In words, a strictly stable centre function shifts the centre by either 0 or 1 if a single character is inserted and two inter-equivalent words shift by the same amount. Since every word is inter-equivalent to itself, that also means that for a single word, the centre shifts in the same direction independent of which character is inserted. Also note that every centre function defined by a biDFA is strictly stable.

Given this, Dieck and Verwer prove a Myhill-Nerode style theorem whose corollary is important for this work [8]:

Corollary 1. *For a given strictly stable centre function c and a language L the number $|Q|$ of states of the minimal biDFA defining the same centre function c and recognising L is equal to the number of equivalence classes of the inter-equivalence w.r.t. c and L . There exists an isomorphism between such a minimal biDFA and the “canonical” minimal biDFA w.r.t. c and L .*

This corollary means that we can identify a biDFA that is minimal w.r.t. a given centre function if we can identify the corresponding equivalence classes.

3 The Algorithm

In this section, we will introduce the active learning algorithm for biDFA. The algorithm is in large parts an adaptation of the L^* algorithm introduced by Angluin [1] but uses a counterexample processing based on Rivest and Schapire [20, 21]. For the description of the algorithm, we will assume the existence of a teacher that acts as a membership and equivalence oracle, like the original L^* algorithm, in addition to a centre function oracle. For a teacher of a language $L \subseteq \Sigma^*$ the oracles are expected to have the following capabilities:

The membership oracle can answer whether a given word $w \in \Sigma^*$ is in L or not. The equivalence oracle can answer for a given biDFA B whether the

language recognised by B is equivalent to L . If it is not equivalent the oracle returns a counterexample. The centre function oracle for a fixed strictly stable centre function c will return for a given word $w \in \Sigma^*$ its centre $c(w)$.

We will discuss how each of the oracles can be realised in Sect. 3.2.

The algorithm at all times maintains two sets $A, E \subset \Sigma^*$, the set A of access words and the set E of extensions. Both are initialised with $A, E = \{\lambda\}$. In addition to these two sets it maintains two tables, D , the distinguishing table, with $|A|$ rows and $|E|$ columns and X , the extended table, with $|A||\Sigma|$ rows and $|E|$ columns. For $a \in A$ and $e \in E$, the distinguishing table has the entry $d_{a,e} = 1$ if $a \cdot_c e \in L$ and 0 otherwise. Similarly for $a \in A$, $s \in \Sigma$ and $e \in E$ the extended table has the entry $x_{a,s,e} = 1$ if $a \cdot_c s \cdot_c e \in L$ and 0 otherwise.

We consider A to be **complete** if for all $a' \in A$ and $s \in \Sigma$ there exists an $a \in A$ s.t. $x_{a',s,e} = d_{a,e}$ for all $e \in E$. If A is not **complete** we will add one $a' \cdot_c s$ to A for which for all $a \in A$ there exists an $e \in E$ s.t. $x_{a',s,e} \neq d_{a,e}$. One can think of A as the set of unique equivalence classes and, therefore, states that have already been identified. We know for all $a, a' \in A$ that $[a]_{c,L} \neq [a']_{c,L}$ since there exists an $e \in E \subset \Sigma^*$ s.t. $a \cdot_c e \in L$ and $a' \cdot_c e \notin L$ or vice versa, since $d_{a,e} \neq d_{a',e}$ for at least one $e \in E$. Since we want to maintain this property we will not add new access words outside of this routine to make A complete. This differs from the original L^* algorithm, which also adds new access words when obtaining a counterexample [1].

An important observation here is that we can make the queries to the membership oracle, that are necessary to fill the tables, even with only partial information on the centre function. If we know $c(a)$, we can form $a \cdot_c e$ for any $e \in E$. The same is true for $a \cdot_c s \cdot_c e$ if we know $c(a)$ and $c(a \cdot_c s)$. Further, since c is strictly stable, it suffices to know $c(a \cdot_c s)$ for one $s \in \Sigma$. This way, we can limit the number of times we need to call the centre function oracle. We can save all the information the algorithm needs to function by having a label $o(a)$ for all $a \in A$. This label will be “l” if we believe the state associated with a is in Q_l , and the centre function will shift forward by 1 upon inserting a single character and “r” otherwise. We can obtain this label by querying the centre function oracle for the centre of $a \cdot_c s$ for any $s \in \Sigma$ when a is added to A . Since $c(\lambda) = 0$ and all other access words are one-letter extensions of other access words, this is sufficient to build o starting from the initialisation.

Given a **complete** A , we can now construct a biDFA by creating a state $q_a \in Q_l$ for all $a \in A$ with $o(a) = \text{“l”}$ and a state $q_{a'} \in Q_r$ for all $a' \in A$ with $o(a') = \text{“r”}$. Since A is **complete** for every $a' \in A$ and $s \in \Sigma$ there exists a unique $a \in A$ s.t. $d_{a,e} = x_{a',s,e}$ for all $e \in E$. We will call this a $x(a', s)$. With this we can define $\delta(q_a, s) = q_{x(a,s)}$. Further, we will set q_λ to be the state created from $\lambda \in A$ and $F = \{q_a | d_{a,\lambda} = 1\}$. This fully defines a biDFA since we assume Σ is given.

Once a biDFA B is built from a **complete** set A we can use the equivalence oracle. If the oracle certifies that B recognises L , we terminate returning B . Otherwise, we process the counterexample w given by the oracle. Here, we use a method based on the work of Rivest and Schapire [20, 21]. Since w is a coun-

terexample, meaning either B accepts it, while it is not in L or vice versa, we can use it to identify a new extension that distinguishes a word $a' \cdot_c s$ from every row in the distinguishing table. There must exist a pre- and inter- and suffix of w , say w_p, w_i and w_s with $w_p \cdot w_i \cdot w_s = w$ and $c(w_p w_s) = |w_p|$ s.t. $\delta^*(q_\lambda, w_p w_s) = q_a$ but $[w_p w_s]_{c,L} \neq [a]_{c,L}$ and w.l.o.g. $w_p w_s \cdot_c w_i \notin L$ and $a \cdot_c w_i \in L$. If this was not the case, B would also reject w . Let $w_p w_s$ be the shortest of such pre- and suffixes. So if w.l.o.g. $w'_p \cdot \sigma \cdot w_s = w_p w_s$ for some $\sigma \in \Sigma$, and $\delta^*(q_\lambda, w'_p w_s) = q_{a'}$, then $x_{a', \sigma, w_i} \neq d_{a, w_i}$. But since the row associated with a is unique and otherwise equal with the row associated with $a', \sigma w_i$ distinguishes $a' \cdot \sigma$ from every row in the distinguishing table, which was our claim. Finally, observe that we can identify w_i since for some pre- and suffix $w_p w_s$ with $\delta(q_\lambda, w_p w_s) = q_a$ if $[w_p w_s]_{c,L} = [a]_{c,L}$ then $w_p w_s \cdot_c w_i \in L \iff a \cdot_c w_i \in L$. It, therefore, suffices to process w character by character and substituting the pre- and suffix already processed with the access trace of the state that was reached. A membership query is asked before and after the substitution. The first time the queries differ before and after substitution, the shortest pre- and suffixes are identified. This can also be used to improve the query complexity of this search by using binary search to find the shortest pre- and suffix from which predictions begin to differ, which is the result obtained by Rivest and Schapire [20, 21].

3.1 Correctness and Running Time

While many of the key arguments have already been made in the previous section, we formalise the correctness of the presented algorithm with the following theorem:

Theorem 1. *Given a teacher T for a language L and a strictly stable centre function c , the presented algorithm will return a biDFA that is isomorphic to the canonical biDFA given c and L .*

Theorem 1 shows that if a biDFA is returned by the algorithm, then it will be the minimal one w.r.t. c and L . However, it gives no guarantee for the algorithm terminating and how long it will take for the algorithm to terminate. For the L^* -style algorithms, they are often evaluated on how many calls to the membership and equivalence oracle they make [1, 7]. We will do the same but also include the number of calls to the centre function oracle.

Theorem 2. *Given a teacher T for a language L and a centre function c the presented algorithm will identify the minimal biDFA $B = (Q_l, Q_r, \Sigma, \delta, q_\lambda, F)$ after asking at most $O(|Q|)$ equivalence queries, $O(|Q|)$ centre function queries and $O(|Q|^2|\Sigma| + |Q|\log_2(|x|))$ membership queries, where $Q = Q_l \cup Q_r$ and x is the longest counterexample.*

Essentially, we retain the same running time guarantees as modern L^* -style algorithms used for learning a DFA while learning a biDFA.

3.2 Heuristics

While the version of the algorithm relying on oracles is interesting from a theoretical perspective for many applications where DFA active learning is normally used, namely learning surrogate models for more complex models, these oracles are normally not available. In this section, we will discuss how each of the oracles can be implemented.

First, note that the membership oracle does not interact with a biDFA and, therefore, is independent of the model we are trying to learn. The implementation of the membership oracle just relies on what model is used as a teacher. When learning from a more complex model, this normally involves running the complex model with the queried input to obtain the output [14].

While in Sect. 4 we will discuss how the language equivalence problem, and therefore oracle, can be solved if the teacher is a biDFA or DFA, this problem cannot be solved exactly for arbitrary models. There exist methods that heuristically predict equivalence while using only membership queries. The two most popular ones are the W -method [3] and random walks [14]. The W -method generates the set of all words of length less than a given parameter m , $W = \{u \in \Sigma^* \mid |u| \leq m\}$ and then asks membership queries of both the teacher and the hypothesized model for all words $a \cdot u \cdot e$ for all $a \in A$, $u \in W$ and $e \in E$. It compares the output on both queries. When learning DFA, if all query pairs give the same output, it guarantees that the models either recognise the same language or that the DFA would require at least m more states to recognise the same language [3]. We can easily adjust this method to work with biDFAs by querying for all words $a \cdot_c u \cdot_c e$ for all $a \in A$, $u \in W$, $e \in E$ instead. One can prove that such an adaptation retains the same guarantees by showing that the shortest distinguishing sequence being longer than m implies that at least m states need to be split into at least two states.

Lastly, we will discuss how the centre function oracle can be implemented. Unless $P=NP$, a centre function oracle, that runs in polynomial time but results in the globally minimal biDFA will not be possible even if a biDFA is given as a teacher since that would result in a P-time algorithm for the minimisation problem with a regular language, which we show in Sect. 4 to be an NP-complete problem. In the more realistic setting where the teacher is some opaque model, which we can only interact with through membership queries and the centre function is no longer given, it is, therefore, unrealistic to expect a heuristic that leads to a globally good solution. Accordingly, we propose a greedy heuristic that attempts to minimise the immediate number of new states that would appear as a result of fixing the centre function. This heuristic is shown in Algorithm 1.

A simple summary of the heuristic is that it chooses the direction in such a way that the new rows that will be added to the extended table have the maximum overlap with already existing rows in the distinguishing table. If a newly added row in the extended table overlaps with an existing one the distinguishing table one can think of this as an indication that likely no new state will be introduced by this row. The heuristic therefore chooses the direction of a state in such a way that locally the number of states is greedily minimised. If no direction is

better, it flips a coin. One can think of this as looking one step into the future since only single-character extensions are considered.

```

input : A word  $a^* = a \cdot_c s'$ , a teacher  $T$ , a set of access words  $A$ , a set of
distinguishing words  $E$ , a partial centre function  $o$  and a
distinguishing table  $d$ 
output: An orientation  $o(a^*)$ 
1  $l \leftarrow 0$ ; // An overlap counter for assuming  $o(a^*) = 'l'$ 
2  $r \leftarrow 0$ ; // An overlap counter for assuming  $o(a^*) = 'r'$ 
3 for all  $s \in \Sigma$  do
4   Assume  $o(a^*) = 'l'$ ;
5    $x_{a^*,s,e}^l = \text{MembershipQuery}(a^* \cdot_c s \cdot_c e)$ ;
6   if  $\exists a \in A : \forall e \in E : d_{a,e} = x_{a^*,s,e}^l$  then
7      $l \leftarrow l + 1$ ;
8   Assume  $o(a^*) = 'r'$ ;
9    $x_{a^*,s,e}^r = \text{MembershipQuery}(a^* \cdot_c s \cdot_c e)$ ;
10  if  $\exists a \in A : \forall e \in E : d_{a,e} = x_{a^*,s,e}^r$  then
11     $r \leftarrow r + 1$ ;
12 if  $l > r$  then
13   return  $o(a^*) = 'l'$ 
14 else if  $r > l$  then
15   return  $o(a^*) = 'r'$ 
16 else
17   return  $o(a^*) = 'l'$  or  $o(a^*) = 'r'$  with probability 0.5 each

```

Algorithm 1: Heuristic for choosing state orientation

One could also extend the heuristic by running the algorithm for n steps and choosing the direction that minimised the number of states after those steps. But since each step requires choosing a direction again, this would require making $2^n |A| |E| |\Sigma|$ membership queries. This quickly becomes infeasible but might be worth considering for very small values of n .

As given, there are instances where the heuristic of Algorithm 1 will result in the algorithm taking exponential time in expectation if we assume E always contains a distinguishing sequence if it exists. For example $L = \{a, b\}^n (c\{a, b\}^{10})^n$ over the alphabet $\Sigma = \{a, b, c\}$ is problematic for the heuristic. Since every time the algorithm chooses a state to read from the left without choosing right-oriented states for the next 11 successors, the number of minimal states of a model which respect this centre function will increase by 12.

However, while this example gives a good idea that the heuristic performs worse the longer and more generic the pumpable sequences are, it is not very useful for the actual application of the algorithm. This is due to the assumption that E contains a distinguishing sequence if it exists. Generally, the heuristic approach relies on an accurate prediction of inter-equivalence. This is also the

main motivation for using the counterexample processing of Rivest and Schapire [20, 21]. If we did not have the invariant that entries in A represent unique states then in a heuristic setting we would need to guarantee that entries in A that cannot be distinguished by E get assigned the same orientation. However, after adding new entries to E one might discover that those states are not equivalent but it would be no longer possible to assign a new orientation to one of the states. By setting up the algorithm as we did we delay orientation assignments as long as possible. This should lead to more accurate assessments of overlap by the heuristic since E will contain more entries. For our experiments, we wanted to further improve this overlap assessment by initialising E to contain several words. We achieved this in two different ways. One was to simply add all words up to a certain length to E . The other method was to run the algorithm for some iterations and then restart it while retaining E between restarts. We call this second method **Reset and Remember**.

3.3 Experiments

Setup. For experiments, we implemented the algorithm in C++ and ran it on several simple languages as a proof of concept. For each language, the teacher was given in the form of a biDFA. However, the algorithm only had access to it in the form of membership queries to simulate the more common use case. For the equivalence query, we implemented the W -method, and for the centre function oracle, we implemented Algorithm 1. Additionally, we used three methods to initialise the algorithm, one **standard** method where $A = E = \{\lambda\}$ and one **μ -warm start** method, where E is initialised as $\{s \in \Sigma^* \mid |s| \leq \mu\}$. Finally, we implemented the **Reset and Remember** method described at the end of Sect. 3.2. For this, we reset the algorithm for the k th time if a model with at least 2^{1+k} states was discovered. The algorithm terminated if it predicted the smallest so far known model twice in a row. For all methods, we also set a limit to forcibly terminate once the distinguishing table A contained more than 100 rows. The algorithm was used 10 times on each language for each experiment, and we recorded the average runtime, the number of times the forced termination occurred, as well as the maximum, minimum, and average number of states of the final model, excluding the ones that were forcibly terminated.

The languages we tested on where $L_1 = a^n b^n$, $L_2 = a^n ccb^n$, $L_3 = a^n (bb)^n$ as some simple linear languages that give an idea of how the algorithm performs on these simple variations of a two-sided pump pattern. We also ran experiments on a regular language $L_4 = (abbb)^n$, which illustrates some of the pitfalls of the algorithm. To showcase the strength of the algorithm, we used the palindrome language over $\Sigma = \{a, b, c\}$, which is $L_5 = ww^{-1}$ a string followed by its reversed string. Finally, we chose two languages designed to challenge our proposed heuristic of Algorithm 1. The chosen languages try to avoid revealing local information immediately. These languages are $L_6 = ((\{a, b\})^2)^n (b\{a, b\}^5)^n$ and $L_7 = (\{a, b\}^2 a \{a, b\}^2)^n (\{a, b\}^2 b \{a, b\}^2)^n$.

Experimental Results. An informative subset of the experimental results can be found in Table 2 while the full experimental results are available in the extended version of the paper¹.

Table 2. Experimental results for different initialisation methods.

	$m = 3, 7, 8$ $\mu = 0$ “RR”=No	$m = 1, 5, 6$ $\mu = 2$ “R and R”=No	$m = 3, 7, 8$ $\mu = 0$ “R and R”=Yes	$m = 1, 5, 6$ $\mu = 2$ “R and R”=Yes
L_1	#MinDisc	6/10	10/10	10/10
	#ForcedTerm	0/10	0/10	0/10
	Average #states	5.8	3	3
L_2	#MinDisc	5/10	10/10	10/10
	#ForcedTerm	0/10	0/10	0/10
	Average #states	7.1	5	5
L_3	#MinDisc	6/10	5/10	10/10
	#ForcedTerm	0/10	0/10	0/10
	Average #states	7.3	5.1	4
L_4	#MinDisc	1/10	1/10	1/10
	#ForcedTerm	0/10	0/10	0/10
	Average #states	8.8	7.4	9.0
L_5	#MinDisc	2/10	10/10	10/10
	#ForcedTerm	8/10	0/10	0/10
	Average #states	5	5	5
L_6	#MinDisc	0/10	1/10	0/10
	#ForcedTerm	10/10	3/10	8/10
	Average #states	_*	14.4	18
L_7	#MinDisc	2/10	4/10	6/10
	#ForcedTerm	0/10	0/10	0/10
	Average #states	15.8	12.9	14.2

A initial observation of the experiments was that one can reduce the value of m by the value of μ since the W method compares outputs on all strings $a \cdot_c u \cdot_c e$ for all $a \in A, u \in W$ and $e \in E$. Normally E initially contains only λ . If strings of length μ are added, the W -method can find μ longer counterexamples during the first equivalence query.

In the experiments, both the μ -warm start method and **Reset and Remember** method were successful in making the heuristic notably more performant. Without it, the algorithm essentially always failed on languages L_5 and

¹ https://github.com/SimonDieck/biDFA_L_star/blob/main/CIAA_2025_biDFA_active_learning_camera_ready_full_version.pdf.

L_6 and struggled with L_2 . Combining the two methods showed some improvement in the average number of states. For running time, the μ -**warm start** was fastest, followed by the combined method with **Reset and Remember** being the slowest.

When run with one of the initialisation methods, the algorithm almost always finds a minimal size model for our simple languages L_1 and L_2 . At the same time, only **Reset and Remember** was very consistent on L_3 . It also often finds these minimal size models for these simple languages when run without adding extra words to E . On the palindrome language L_5 however, it seems the initialisation methods are necessary. Without them, the model struggled to find the minimal size model and even often encountered the forced termination. Nevertheless, if one of the initialisation methods was used, the algorithm consistently found a minimal-size model for the palindrome language. For L_4 , the regular language, one can however observe that the algorithm rarely finds the minimal size model with only the combined initialisation performing well. The minimal size model orients all states in the same direction, and as soon as one non-sink state is oriented differently the number of required states is almost doubled. However, the heuristic cannot obtain this information by looking only one step into the future and thus struggles to find a minimal-size model. Surprisingly though the algorithm performed well on L_6 and L_7 which were designed to exploit the same problems of the centre function heuristic. While the algorithm, even with the initialisation methods, still sometimes hit forced termination on those languages it often found a small model that recognised those languages and sometimes even a minimal-size model.

4 Language Theoretical Properties

Language Equivalence. The algorithm that is presented in this work needs access to an equivalence oracle that produces a counterexample if languages are not equivalent. However, the complexity of the problem of deciding whether two arbitrary biDFA recognise the same language is currently not known. We however conjecture that the problem is at least in co-NP. Jirásková and Klíma have shown that the problem is in NL if one of the languages is regular [15]. For the non-regular case we believe that for a biDFA to correctly identify a two-sided pump-pattern like $uv^nw x^n y$, where both v and x are non-empty it will have to place the centre in between the copies of v and x . If this placement can be limited to at most some constant it would force two biDFAs whose languages have a significant overlap to define similar centre functions. If one can prove and combine these results it should be possible to limit the size of the smallest counterexample.

A simpler version of the language equality problem we can consider, even if the defined languages are non-regular, is if two biDFA define the same centre function.

Lemma 1. *Given two biDFAs $B_1 = (Q_l, Q_r, \Sigma, \delta, q_\lambda, F)$ and $B_2 = (Q'_l, Q'_r, \Sigma, \delta', q'_\lambda, F')$ which recognise L_1 and L_2 and define the centre functions c_1 and c_2 respectively. If $c_1 = c_2$ determining whether $L_1 \cap L_2 = \emptyset$ is NL-complete.*

The core idea of the proof is that once the centre function is fixed one can construct an intersection biDFA which respects the same centre function.

Corollary 2. *The language equality problem for two biDFAs B_1 and B_2 defining L_1 and L_2 respectively, that both define the same centre function c is in NL.*

To prove the corollary one can use the standard construction of $L_1 \cap \overline{L_2} = \emptyset$ and $\overline{L_1} \cap L_2 = \emptyset$ implies $L_1 = L_2$.

Minimisation. The algorithm shown in the previous section also answers some questions about the complexity of the biDFA minimisation problem. One can define this problem as given a biDFA B which recognises L and a constant k to decide whether there exists a biDFA B^* that also recognises L with $|Q^*| \leq k$.

Theorem 3 (Minimisation complexity). *Given a membership and equivalence oracle for L there exists a nondeterministic polynomial time algorithm that solves the biDFA minimisation problem which requires a polynomial number of calls to the membership oracle and a linear number of calls to the equivalence oracle.*

Proof. Note that there are only two possible outcomes whenever the algorithm calls the centre function oracle. A nondeterministic algorithm can guess the outcome and otherwise run the algorithm as described. The algorithm will return “yes” if the final biDFA has at most k states and “no” otherwise. Theorems 1 and 2 guarantee that the returned algorithm recognises L and that this procedure had a nondeterministic polynomial running time.

Since the input biDFA to the minimisation problem can be given as the teacher, and therefore the membership query can be solved in linear time, the complexity of the minimisation problem depends on the complexity of the language equivalence problem. This leads us to the following corollary:

Corollary 3. *Given a biDFA B recognising the language L and defining the centre function c and a constant k the minimisation problem is:*

1. In P when restricted to finding a biDFA that also defines c .
2. NP-complete if L is regular.

Proof. 1. follows from Theorem 3 and Corollary 2 as well as the observation that a biDFA can be used as an oracle for its own centre function.
 2. follows from Theorem 3 and Dieck and Verwer’s proof that the problem is NP-hard [8] as well as Jirásková and Klíma’s proof that the language equivalence problem is NL-complete if L is regular [15].

5 Discussion

In this work, we have introduced an active learning algorithm for biDFAs and some proofs about the complexity of the language equivalence problem for biDFAs. Together this also answered some open questions about the complexity of the biDFA minimisation problem. When combined with the heuristics we provided for the oracle required by the algorithm, we believe this work to be an important step to move bidirectional automata from a theoretical novelty towards something that can be used in practice. Indeed for DFA applications where the memory footprint is of significant importance, one could use the NP algorithm for minimisation introduced in this work. For the more common reverse engineering applications in security where DFA are currently learned one could experiment with the algorithm introduced in this work. For this, we expect that there is also significant room for improving the heuristics we provided. Especially the centre function heuristic we proposed, while working well on toy examples, provides no theoretical guarantees. Since beyond their superior expressiveness biDFA are also interesting because they can be potentially exponentially smaller than DFA when recognising regular languages [8], it would be desirable to find a centre function heuristic that works well on regular languages.

Another issue that keeps these results from being applicable to many real-life scenarios is that often not DFA are learned but Mealy machines [7, 14]. Nevertheless, with the theoretical foundations provided in this work, we expect the same algorithm can be used to learn bidirectional Mealy machines, similarly to how L^* is often used to learn them [14].

If one could show that the language equivalence problem is indeed in co-NP it would likely also imply that biDFAs are learnable in the limit [10]. If that is the case it would also be interesting if passive learning algorithms [22] or exact algorithms [11], that learn from a fixed sample could be adapted to learn biDFAs instead.

From a theoretical perspective, there are also some open questions left. Most importantly, whether the language equality problem is indeed as we conjectured in co-NP for the general case. Together with Theorem 3 this would imply that minimisation is also in co-NP.

For two biDFA that define the same centre function we have proven membership in NL for the equivalence problem and membership in P for the minimisation problem but completeness remains an open question.

In conclusion, we believe that bidirectional automata can be useful in real-life applications and to that end, we have in this work provided a foundation upon which such methods can be built and with which initial experiments can be conducted.

References

1. Angluin, D.: Learning regular sets from queries and counterexamples. *Inf. Comput.* **75**(2), 87–106 (1987)

2. Cho, C.Y., Babić, D., Poosankam, P., Chen, K.Z., Wu, E.X.J., Song, D.: Mace: model-inference-assisted concolic exploration for protocol and vulnerability discovery. In: 20th USENIX Security Symposium (USENIX Security 2011) (2011)
3. Chow, T.S.: Testing software design modeled by finite-state machines. *IEEE Trans. Softw. Eng.* **4**(3), 178–187 (1978)
4. Clark, A., Eyraud, R.: Polynomial identification in the limit of substitutable context-free languages. *J. Mach. Learn. Res.* **8**(8), 1–21 (2007)
5. Cook, J.E., Wolf, A.L.: Discovering models of software processes from event-based data. *ACM Trans. Softw. Eng. Methodol. (TOSEM)* **7**(3), 215–249 (1998)
6. Cui, W., Kannan, J., Wang, H.J.: Discoverer: automatic protocol reverse engineering from network traces. In: USENIX Security Symposium, pp. 1–14 (2007)
7. De la Higuera, C.: *Grammatical Inference: Learning Automata and Grammars*. Cambridge University Press, Cambridge (2010)
8. Simon Dieck and Sizzo Verwer. On bidirectional deterministic finite automata. In: Fazekas, S.Z. (ed) *Implementation and Application of Automata*. CIAA 2024. LNCS, vol. 15015, pp. 109–123. Springer, Cham (2024). https://doi.org/10.1007/978-3-031-71112-1_8
9. Ficara, D., Giordano, S., Procissi, G., Vitucci, F., Antichi, G., Di Pietro, A.: An improved DFA for fast regular expression matching. *ACM SIGCOMM Comput. Commun. Rev.* **38**(5), 29–40 (2008)
10. Mark Gold, E.: Language identification in the limit. *Inf. Control* **10**(5), 447–474 (1967)
11. Heule, M.J.H., Verwer, S.: Exact DFA identification using SAT solvers. In: Semper, J.M., García, P. (eds.) *ICGI 2010*. LNCS (LNAI), vol. 6339, pp. 66–79. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-15488-1_7
12. Hoe, A.V., Sethi, R., Ullman, J.D.: *Compilers—Principles, Techniques, and Tools*. Pearson Addison Wesley Longman (1986)
13. Holzer, M., Jakobi, S.: Minimization and characterizations for biautomata. *Fund. Inform.* **136**(1–2), 113–137 (2015)
14. Isberner, M., Howar, F., Steffen, B.: The open-source LearnLib. In: Kroening, D., Păsăreanu, C.S. (eds.) *CAV 2015*. LNCS, vol. 9206, pp. 487–495. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-21690-4_32
15. Jirásková, G., Klíma, O.: On linear languages recognized by deterministic biautomata. *Inf. Comput.* **286**, 104778 (2022)
16. Klíma, O., Polák, L.: On biautomata. *RAIRO-Theoret. Inform. App.* **46**(4), 573–592 (2012)
17. Muškardin, E., Aichernig, B.K., Pill, I., Tappler, M.: Learning finite state models from recurrent neural networks. In: ter Beek, M.H., Monahan, R. (eds.) *integrated formal methods*. IFM 2022. LNCS, vol. 13274, pp. 229–248. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-07727-2_13
18. Nagy, B.: On a hierarchy of $5' \rightarrow 3'$ sensing Watson-crick finite automata languages. *J. Log. Comput.* **23**(4), 855–872 (2013)
19. Nagy, B., Parchami, S.: On deterministic sensing $5' \rightarrow 3'$ Watson-crick finite automata: a full hierarchy in 2detlin. *Acta Informatica* **58**, 153–175 (2021)
20. Rivest, R.L., Schapire, R.E.: Inference of finite automata using homing sequences. *Inf. Comput.* **103**(2), 299–347 (1993). <https://www.sciencedirect.com/science/article/pii/S0890540183710217>, <https://doi.org/10.1006/inco.1993.1021>
21. Rivest, R.L., Schapire, R.E.: Inference of finite automata using homing sequences. In: *Proceedings of the Twenty-first Annual ACM Symposium on Theory of Computing*, pp. 411–420 (1989)

22. Verwer, S., Hammerschmidt, C.A.: Flexfringe: a passive automaton learning package. In: 2017 IEEE International Conference on Software Maintenance and Evolution (ICSME), pp. 638–642. IEEE (2017)
23. Yoshinaka, R.: Efficient learning of multiple context-free languages with multi-dimensional substitutability from positive data. *Theoret. Comput. Sci.* **412**(19), 1821–1831 (2011)