

An L^* Algorithm for Deterministic Weighted Regular Languages

Clemente Pasti Talu Karagöz Anej Svete

Franz Nowak Reda Boumasmoud Ryan Cotterell

{clemente.pasti, anej.svete, franz.nowak, ryan.cotterell}@inf.ethz.ch
reda.boumasmoud@math.ethz.ch, talukaragoz@gmail.com

ETH zürich

Abstract

Extracting finite state automata (FSAs) from black-box models offers a powerful approach to gaining interpretable insights into complex model behaviors. To support this pursuit, we present a weighted variant of Angluin’s (1987) L^* algorithm for learning FSAs. We stay faithful to the original formulation, devising a way to exactly learn deterministic weighted FSAs whose weights support division. Furthermore, we formulate the learning process in a manner that highlights the connection with FSA minimization, showing how L^* directly learns a minimal automaton for the target language.

 github.com/rycolab/weighted-angluin

1 Introduction

Learning formal languages from data is a classic problem in computer science. Unfortunately, learning only from positive examples is impossible, as shown by Gold (1978). However, by granting the learner access to more than just positive examples, Angluin (1987) introduced the *active* learning scheme L^* , where the learner interacts with an oracle by asking it queries. Concretely, Angluin’s (1987) L^* algorithm learns regular languages in the form of deterministic finite-state automata (DFSAs) from *membership* queries (analogous to asking for a ground truth label of a string in the training dataset) and *equivalence* queries (analogous to asking whether a hypothesis is correct).

Weighted formal languages, where strings are assigned weights such as probabilities or costs, naturally generalize membership-based (boolean) formal languages. Weighted languages, especially probabilistic languages, serve as a cornerstone in the conceptual framework of many NLP problems (Mohri, 1997). Their significance is twofold: first, in practical applications, where they underpin algorithms for tasks such as parsing (Goodman, 1996) and machine translation (Mohri, 1997), and second, as an analytical framework for better understanding

modern language models (Weiss et al., 2018; Jumelet and Zuidema, 2023; Nowak et al., 2024, *inter alia*). The practical applications in NLP, coupled with theoretical interests in formal language theory, have motivated the development of various weighted extensions of Angluin’s (1987) L^* . For instance, Weiss et al. (2019) describes a generalization of L^* that (approximately) learns a probabilistic DFSA by querying a neural language model to interpret it. Additionally, and less faithfully to the original L^* algorithm, multiple algorithms for learning *non-deterministic* weighted FSAs have been proposed (Bergadano and Varricchio, 1996; Beigel et al., 2000; Balle and Mohri, 2012; Balle et al., 2014; Daviaud and Johnson, 2024).

In this paper, we present a novel weighted generalization of the L^* algorithm that learns *semifield*-weighted *deterministic* FSAs with membership and equivalence queries. In contrast to other algorithms inspired by L^* , ours is a true generalization, i.e., we generalize Angluin’s (1987) original algorithm, resulting in a familiar procedure that, just like the original, learns a deterministic FSA exactly in a finite number of steps if the automata can be determinized.¹ Additionally, we loosen the requirement for field-weighted FSAs made by spectral algorithms, e.g., Balle and Mohri’s (2012) algorithm, that learn non-deterministic automata; our algorithm only requires a semifield-weighted FSA. Interestingly, our exposition further illuminates the connection between three important algorithms for weighted FSAs: minimization (Hopcroft and Ullman, 1979; Mohri, 1997), weight pushing (Mohri, 1997; Mohri and Riley, 2001) and L^* .

2 Weighted Regular Languages

Semirings and Semifields. A **monoid** is a 3-tuple $(\mathbb{K}, \bullet, \mathbf{1})$, where \mathbb{K} is a set, $\bullet : \mathbb{K} \times \mathbb{K} \rightarrow \mathbb{K}$ is an associative operator, and $\mathbf{1} \in \mathbb{K}$ is a distinguished identity element such that

¹All boolean-weighted FSA can be determinized, which is why Angluin’s (1987) L^* always halts.

$\mathbf{1} \bullet w = w \bullet \mathbf{1} = w$ for any $w \in \mathbb{K}$. A monoid is **commutative** if \bullet is a commutative operator. A **group** is a monoid $(\mathbb{K}, \bullet, \mathbf{1})$ where every element has an inverse: for every $w \in \mathbb{K}$ there exists a $w^{-1} \in \mathbb{K}$ such that $w \bullet w^{-1} = w^{-1} \bullet w = \mathbf{1}$; we use the notation w_1/w_2 as an alias for $w_2^{-1} \otimes w_1$. A **semiring** $(\mathbb{K}, \oplus, \otimes, \mathbf{0}, \mathbf{1})$ is a 5-tuple where $(\mathbb{K}, \oplus, \mathbf{0})$ is a commutative monoid, $(\mathbb{K}, \otimes, \mathbf{1})$ is a monoid, \otimes distributes over \oplus and $\mathbf{0}$ is an annihilator for \otimes . A **semifield** is a semiring $(\mathbb{K}, \oplus, \otimes, \mathbf{0}, \mathbf{1})$ where $(\mathbb{K} \setminus \{\mathbf{0}\}, \otimes, \mathbf{1})$ is a commutative group.

Strings and Languages. An **alphabet** Σ is a non-empty, finite set of symbols. A **string** is a finite sequence of symbols from an alphabet. We write \mathbf{xy} to denote the concatenation of the strings \mathbf{x} and \mathbf{y} . Let $\Sigma^{n+1} \stackrel{\text{def}}{=} \{\mathbf{ya} \mid \mathbf{y} \in \Sigma^n, a \in \Sigma\}$ and $\Sigma^0 \stackrel{\text{def}}{=} \{\varepsilon\}$, where ε is the empty string. The **Kleene closure** $\Sigma^* \stackrel{\text{def}}{=} \bigcup_{n=0}^{\infty} \Sigma^n$ of Σ is the set containing all strings made with symbols of Σ . We further introduce the set $\Sigma^{\leq k} = \bigcup_{n=0}^k \Sigma^n$. Given an alphabet Σ and a semiring $\langle \mathbb{K}, \oplus, \otimes, \mathbf{0}, \mathbf{1} \rangle$, a **weighted formal language** is a function $L: \Sigma^* \rightarrow \mathbb{K}$ that assigns weights $w \in \mathbb{K}$ to strings $\mathbf{y} \in \Sigma^*$. Unless differently specified, in this paper we will assume that all weighted languages are *semifield-weighted*.

Weighted Finite-state Automata. A **weighted finite-state automaton** (WFSA) \mathcal{A} over a semifield $\langle \mathbb{K}, \oplus, \otimes, \mathbf{0}, \mathbf{1} \rangle$ is a 5-tuple $(\Sigma, Q, \delta, \lambda, \rho)$ where Σ is an alphabet, Q is a finite set of states, δ is a set of weighted arcs rendered as $p \xrightarrow{a/w} q$ with $p, q \in Q$, $a \in \Sigma$, and $w \in \mathbb{K}$,² and $\lambda: Q \rightarrow \mathbb{K}$ and $\rho: Q \rightarrow \mathbb{K}$ are the initial and final weight function, respectively. A **path** π in \mathcal{A} is a finite sequence of contiguous arcs, denoted as

$$q_0 \xrightarrow{a_1/w_1} q_1, \dots, q_{N-1} \xrightarrow{a_N/w_N} q_N. \quad (1)$$

We call $i(\pi) = q_0$ the initial state of the path, and $f(\pi) = q_N$ the final state of the path. The **weight** of π is $w(\pi) = w_1 \otimes \dots \otimes w_N$ and its **yield** is $\sigma(\pi) = a_1 \dots a_N$. With $\Pi_{\mathcal{A}}$, we denote the set of all paths in \mathcal{A} , and with $\Pi_{\mathcal{A}}(\mathbf{p})$ the subset of all paths in \mathcal{A} with yield \mathbf{p} . We say that a WFSA $\mathcal{A} = (\Sigma, Q, \delta, \lambda, \rho)$ is **deterministic** (a WDFSA) if, for every $p \in Q$, $a \in \Sigma$, there is at most one $q \in Q$ such that $p \xrightarrow{a/w} q \in \delta$ with $w > 0$, and there is a single state q_I with $\lambda(q_I) \neq 0$. In such

²We do not consider ε -transitions. This is without loss of generality; any regular language can be represented by an ε -free automaton (Mohri, 2009, Theorem 7.1).

case, we refer to q_I as the **initial state**. Naturally, a WDFSA can have at most one path yielding a string $\mathbf{y} \in \Sigma^*$ from the initial state q_I .

Weighted Regular Languages. Every WFSA \mathcal{A} generates the weighted language

$$L_{\mathcal{A}}(\mathbf{p}) \stackrel{\text{def}}{=} \bigoplus_{\pi \in \Pi_{\mathcal{A}}(\mathbf{p})} \lambda(i(\pi)) \otimes w(\pi) \otimes \rho(f(\pi)) \quad (2)$$

for $\mathbf{p} \in \Sigma^*$. We define the set $\text{supp}(L) = \{\mathbf{p} \in \Sigma^* \mid L(\mathbf{p}) \neq \mathbf{0}\}$ to be the **support** of L . A weighted language is said to be **regular** if there exists a WFSA that generates it. If two WFSA generate the same language, they are said to be **equivalent**. Finally, a weighted regular language is said to be **deterministic** if there exists a WDFSA that generates it. In contrast to the boolean case, not every weighted regular language can be generated by a deterministic WFSA (Allauzen and Mohri, 2003), and, therefore, weighted deterministic regular languages are a strict subset of weighted regular languages. This distinction plays a critical role in our exposition—we develop a generalization of Angluin’s (1987) algorithm that learns weighted *deterministic* regular languages.

3 Minimal Deterministic Automata

We now introduce the notion of **right language equivalence** between two strings. It provides the basis for active learning (Biermann and Feldman, 1972; Angluin, 1987) and minimization algorithms (Hopcroft and Ullman, 1979; Mohri, 1997), and with it, our weighted extension of L^* . We begin by defining a notion of equivalence between two languages L_1 and L_2 with weights on the same semifield $\langle \mathbb{K}, \oplus, \otimes, \mathbf{0}, \mathbf{1} \rangle$. We write $L_1 \equiv L_2$ if there exists a $k \in \mathbb{K} \setminus \{\mathbf{0}\}$ such that $L_1(\mathbf{p}) = k \otimes L_2(\mathbf{p})$ for every $\mathbf{p} \in \Sigma^*$.³ Next, given a weighted language L and $\mathbf{y} \in \Sigma^*$, we define \mathbf{y} ’s **right language** $\mathbf{y}^{-1}L(\mathbf{p}) \stackrel{\text{def}}{=} L(\mathbf{yp})$ for $\mathbf{p} \in \Sigma^*$. The definition of right language naturally induces on Σ^* the **right language equivalence relation** $\mathbf{y} \sim_L \mathbf{x} \iff \mathbf{y}^{-1}L \equiv \mathbf{x}^{-1}L$, $\mathbf{x}, \mathbf{y} \in \Sigma^*$ (Mohri, 1997).⁴ The core intuition behind both minimization (Hopcroft and Ullman, 1979; Mohri, 1997) and active learning algorithms (Angluin, 1987) is that of building a DFSA whose states correspond to the right language equivalence classes of L .

³ \equiv is an equivalence relation; see App. A.

⁴The fact that \sim_L is an equivalence relation follows directly from the fact that \equiv is an equivalence relation.

Minimality. We say that a WDFSA is **minimal** if no other equivalent WDFSA has fewer states (Mohri, 1997). The Myhill–Nerode theorem binds right-equivalence classes of a regular language to the states of the minimal DFSA (Hopcroft and Ullman, 1979). We state its weighted version:

Theorem 1 (Myhill–Nerode). *Let \mathcal{A} be a semifield-weighted DFSA and $L_{\mathcal{A}}$ its weighted language. Then, $\sim_{L_{\mathcal{A}}}$ induces a finite number of equivalence classes on Σ^* , which equals the number of states of a minimal automaton for $L_{\mathcal{A}}$.*

Proof. See App. C.1. ■

In a DFSA \mathcal{A} , every state q can be uniquely identified as $q_{\mathbf{p}}$,⁵ where \mathbf{p} is the string that is read by traversing the non-zero-weight path from the initial state to q . Therefore, it is straightforward to extend the notion of right language equivalence from strings to states of \mathcal{A} as $q_{\mathbf{p}} \sim_{L_{\mathcal{A}}} p_{\mathbf{y}} \iff \mathbf{p} \sim_{L_{\mathcal{A}}} \mathbf{y}$. Minimization algorithms for unweighted DFSA are based on merging states that are right-equivalent (Hopcroft and Ullman, 1979). This strategy can be adapted to WDFSA after *canonicalizing* the distribution of the weights over the WDFSA transitions with weight pushing (Mohri, 1997).⁶

4 A Weighted L^* Algorithm

We now present our weighted L^* algorithm. Let $\langle \mathbb{K}, \oplus, \otimes, \mathbf{0}, \mathbf{1} \rangle$ be a semifield and $L^* : \Sigma^* \rightarrow \mathbb{K}$ a deterministic regular language. As Angluin (1987), we assume that we have access to an **oracle** that can answer the following queries about L^* :

- (1) **Membership query:** What is the weight $L^*(\mathbf{p})$ of the string $\mathbf{p} \in \Sigma^*$?
- (2) **Equivalence query:** Does an hypothesis automaton \mathcal{H} generate L^* ? If it does *not*, the oracle provides a **counterexample**, which is a string \mathbf{t} such that $L_{\mathcal{H}}(\mathbf{t}) \neq L^*(\mathbf{t})$.

Our algorithm reduces to Angluin’s (1987) in the case of the boolean semifield.

⁵More precisely, this holds for DFSA where every state is accessible—a condition we will assume throughout the paper without loss of generality.

⁶In the weighted case, equivalent automata with the same topology but different weight distribution along the transitions and the initial and final weights exist. Weight pushing (Mohri, 1997) re-distributes the weights while preserving the language generated by the automaton. It can be used to obtain a canonical weight distribution, which facilitates, for instance, WDFSA equivalence testing. In essence, our algorithm learns a minimal WDFSA of a weighted regular language.

4.1 The Hankel Matrix

Let L be a weighted language as above. L ’s **Hankel matrix** \mathbf{H} is a bi-infinite matrix indexed by elements of Σ^* defined entry-wise as $\mathbf{H}(\mathbf{p}, \mathbf{s}) \stackrel{\text{def}}{=} L(\mathbf{p}\mathbf{s})$ for $\mathbf{p}, \mathbf{s} \in \Sigma^*$. Additionally, for every $\mathbf{p} \in \Sigma^*$, we define the function $\mathbf{H}_{\mathbf{p}} : \Sigma^* \rightarrow \mathbb{K}$ as $\mathbf{H}_{\mathbf{p}}(\mathbf{s}) \stackrel{\text{def}}{=} \mathbf{H}(\mathbf{p}, \mathbf{s}) = \mathbf{p}^{-1}L(\mathbf{s})$ for $\mathbf{s} \in \Sigma^*$.

The L^* algorithm learns from an *empirical* Hankel matrix. Let $\tilde{\mathbf{P}} \subseteq \Sigma^*$ be a prefix-closed set of prefixes and let $\tilde{\mathbf{S}} \subseteq \Sigma^*$ be a suffix-closed set of suffixes.⁷ The **empirical Hankel matrix** $\tilde{\mathbf{H}}$ is a matrix of size $|\tilde{\mathbf{P}} \circ \Sigma^{\leq 1}| \times |\tilde{\mathbf{S}}|$ defined entry-wise as $\tilde{\mathbf{H}}(\mathbf{p}, \mathbf{s}) \stackrel{\text{def}}{=} L(\mathbf{p}\mathbf{s})$ for $\mathbf{p} \in \tilde{\mathbf{P}} \circ \Sigma^{\leq 1}, \mathbf{s} \in \tilde{\mathbf{S}}$. Analogously to the Hankel matrix, we define the function $\tilde{\mathbf{H}}_{\mathbf{p}} : \tilde{\mathbf{S}} \rightarrow \mathbb{K}$ as $\tilde{\mathbf{H}}_{\mathbf{p}}(\mathbf{s}) \stackrel{\text{def}}{=} \tilde{\mathbf{H}}(\mathbf{p}, \mathbf{s})$ for every $\mathbf{s} \in \tilde{\mathbf{S}}$. Note that $\tilde{\mathbf{H}}_{\mathbf{p}}$ is a restriction of $\mathbf{H}_{\mathbf{p}}$ to a finite domain. We define the equivalence relation \equiv on the set $\{\tilde{\mathbf{H}}_{\mathbf{p}} \mid \mathbf{p} \in \tilde{\mathbf{P}} \circ \Sigma^{\leq 1}\}$ as we did for weighted languages (§3). We say that $\tilde{\mathbf{H}}$ is:

- (1) **closed** if, for every $\mathbf{p} \in \tilde{\mathbf{P}}$ and $a \in \Sigma$, there exists $\mathbf{p}' \in \tilde{\mathbf{P}}$ such that $\tilde{\mathbf{H}}_{\mathbf{p}a} \equiv \tilde{\mathbf{H}}_{\mathbf{p}'}$;
- (2) **consistent** if, for every $\mathbf{p}, \mathbf{p}' \in \tilde{\mathbf{P}}$ and $a \in \Sigma$ we have that $\tilde{\mathbf{H}}_{\mathbf{p}} \equiv \tilde{\mathbf{H}}_{\mathbf{p}'} \implies \tilde{\mathbf{H}}_{\mathbf{p}a} \equiv \tilde{\mathbf{H}}_{\mathbf{p}'a}$.

Given a closed and consistent empirical Hankel matrix $\tilde{\mathbf{H}} : \tilde{\mathbf{P}} \circ \Sigma^{\leq 1} \times \tilde{\mathbf{S}} \rightarrow \mathbb{K}$, we define the following equivalence relation on $\tilde{\mathbf{P}} \circ \Sigma^{\leq 1}$:

$$\mathbf{p} \sim_{\tilde{\mathbf{H}}} \mathbf{p}' \iff \tilde{\mathbf{H}}_{\mathbf{p}} \equiv \tilde{\mathbf{H}}_{\mathbf{p}'}. \quad (3)$$

We note the similarity to the right language equivalence relation on Σ^* induced by a weighted language (cf. §3). Because $\tilde{\mathbf{P}}$ is finite by assumption, the relation $\sim_{\tilde{\mathbf{H}}}$ induces a finite number of equivalence classes $\tilde{\mathbf{P}} \circ \Sigma^{\leq 1} / \sim_{\tilde{\mathbf{H}}}$. We denote each equivalence class as $[\mathbf{p}] = \{\mathbf{p}' \in \tilde{\mathbf{P}} \circ \Sigma^{\leq 1} \mid \mathbf{p}' \sim_{\tilde{\mathbf{H}}} \mathbf{p}\}$.

4.2 The Empirical Hankel Automaton

We now introduce the conversion of an empirical Hankel matrix into a WDFSA, the **empirical Hankel automaton** $\mathcal{H} = (\Sigma, Q, \delta, \lambda, \rho)$. We first detail how to construct \mathcal{H} ’s states, transitions, and initial and final weights:

- (1) **States.** We define the states Q to be in a one-to-one correspondence with $\tilde{\mathbf{P}} / \sim_{\tilde{\mathbf{H}}}$: Each state $q_{[\mathbf{p}]} \in Q$ corresponds to an equivalence class $[\mathbf{p}] \in \tilde{\mathbf{P}} / \sim_{\tilde{\mathbf{H}}}$. Here, $\mathbf{p} \in \tilde{\mathbf{P}}$ is a *representative* of the class, which we **assume to be fixed** throughout the construction of \mathcal{H} . We

⁷A set of strings is prefix-closed (suffix-closed) if it contains all prefixes (suffixes) of each of its elements.

also define the map $r : \tilde{P} \circ \Sigma^{\leq 1} \rightarrow \tilde{P}$ that assigns to each prefix the representative of its equivalence class; $r(p') = p \iff p' \in [p]$.

- (2) **Transitions.** For every state $q_{[p]} \in Q$, and every symbol $a \in \Sigma$, let the transition $q_{[p]} \xrightarrow{a/w} q_{[r(pa)]}$ be in δ , where

$$w \stackrel{\text{def}}{=} \frac{\bigoplus_{s \in \tilde{S}} \tilde{\mathbf{H}}(pa, s)}{\bigoplus_{s \in \tilde{S}} \tilde{\mathbf{H}}(p, s)} \quad (4)$$

if $\bigoplus_{s \in \tilde{S}} \tilde{\mathbf{H}}(p, s)$ is non-zero, and $w \stackrel{\text{def}}{=} \mathbf{0}$ else.

- (3) **Initial weight.** For every state $q_{[p]} \in Q$, we define the initial weighting function as

$$\lambda(q_p) \stackrel{\text{def}}{=} \begin{cases} \bigoplus_{s \in \tilde{S}} \tilde{\mathbf{H}}(\varepsilon, s) & p = \varepsilon \\ \mathbf{0} & p \neq \varepsilon \end{cases} \quad (5)$$

- (4) **Final weights.** For every state $q_{[p]}$, we define the final weight

$$\rho(q_{[p]}) \stackrel{\text{def}}{=} \frac{\tilde{\mathbf{H}}(p, \varepsilon)}{\bigoplus_{s \in \tilde{S}} \tilde{\mathbf{H}}(p, s)} \quad (6)$$

if $\bigoplus_{s \in \tilde{S}} \tilde{\mathbf{H}}(p, s)$ is non-zero, and $\rho(q_{[p]}) \stackrel{\text{def}}{=} \mathbf{0}$ else.

Theorem 2 (The empirical Hankel Automaton). *Let $\tilde{\mathbf{H}} : \tilde{P} \circ \Sigma^{\leq 1} \times \tilde{S} \rightarrow \mathbb{K}$ be a closed and consistent empirical Hankel matrix and let \mathcal{H} be the empirical Hankel automaton induced by $\tilde{\mathbf{H}}$. Then:*

- (1) \mathcal{H} is a well-defined WDFSA.
- (2) $L_{\mathcal{H}}(ps) = \tilde{\mathbf{H}}(p, s)$ for all $p \in \tilde{P}$ and $s \in \tilde{S}$, meaning that \mathcal{H} is consistent with $\tilde{\mathbf{H}}$.
- (3) \mathcal{H} is minimal.

Proof. See App. C.2. ■

4.3 The Learning Algorithm

Our weighted \mathbf{L}^* algorithm, with its main loop detailed in Alg. 1, employs the subroutines outlined in Alg. 2.

Initialization. The prefix and suffix sets \tilde{P} and \tilde{S} are initialized as $\{\varepsilon\}$ and the empirical Hankel matrix to the zero matrix.

Handling inconsistencies. The subroutine MAKECONSISTENT in (Line 7, Alg. 1, and Alg. 2) looks for rows $p, p' \in \tilde{P}$ that make $\tilde{\mathbf{H}}$ non-consistent: $\tilde{\mathbf{H}}_p \equiv \tilde{\mathbf{H}}_{p'}$, but $\tilde{\mathbf{H}}_{pa} \not\equiv \tilde{\mathbf{H}}_{p'a}$ for some $a \in \Sigma$. To find the column(s) that make the relation $\tilde{\mathbf{H}}_{pa} \equiv \tilde{\mathbf{H}}_{p'a}$ not true, we normalize over the sum of row entries and compare each entry pair-wise

Algorithm 1 The Weighted \mathbf{L}^* algorithm. Initially, the empirical Hankel matrix $\tilde{\mathbf{H}}$ is set to the zero matrix and the sets \tilde{P}, \tilde{S} to the empty string ε .

```

1. def  $\mathbf{L}^*$  ( $\mathbb{O}$ ):
2.   while true :
3.     while true :
4.       if  $\tilde{\mathbf{H}}$  is not consistent :
5.         MAKECONSISTENT( $\mathbb{O}, \tilde{\mathbf{H}}$ )
6.       else if  $\tilde{\mathbf{H}}$  is not closed :
7.         MAKECLOSED( $\mathbb{O}, \tilde{\mathbf{H}}$ )
8.       else : break
9.      $\mathcal{H} \leftarrow$  MAKEAUTOMATON( $\tilde{\mathbf{H}}$ )
10.    if EQUIVALENT( $\mathbb{O}, \mathcal{H}$ ) : return  $\mathcal{H}$ 
11.    else :
12.       $p \leftarrow$  COUNTEREXAMPLE( $\mathbb{O}, \mathcal{H}$ )
13.      for  $t = 1$  to  $|p| + 1$  :
14.         $\tilde{P} \leftarrow \tilde{P} \cup \{p_{<t}\}$ 
15.      COMPLETE( $\mathbb{O}, \tilde{\mathbf{H}}$ )

```

(Alg. 2, Line 7).⁸ If a column indexed by $s \in \tilde{S}$ is found to make the empirical Hankel matrix not consistent, as is added to \tilde{S} . This results in the new equivalence classes $[p]$ and $[p']$ because $\tilde{\mathbf{H}}_p$ and $\tilde{\mathbf{H}}_{p'}$ do not match anymore on the column indexed by as . See Lemma 5 in App. C for more details.

Closing $\tilde{\mathbf{H}}$. MAKECLOSED (Alg. 1, Line 7; Alg. 2) adds to \tilde{P} the missing prefixes required to make the empirical Hankel matrix closed. It searches for $p \in \tilde{P}, a \in \Sigma$ such that $\tilde{\mathbf{H}}_{pa} \not\equiv \tilde{\mathbf{H}}_{p'}$ for every $p' \in \tilde{P}$, and adds pa to \tilde{P} . This results in the new equivalence class $[pa]$. See Lemma 5 in App. C for more details.

Filling out $\tilde{\mathbf{H}}$. Finally, COMPLETE fills the empty entries of $\tilde{\mathbf{H}}$ by asking membership queries of the oracle.

Handling inconsistencies, closing $\tilde{\mathbf{H}}$, and filling $\tilde{\mathbf{H}}$ is carried out by the inner while loop (Lines 3 to 8) of Alg. 1, which continues until $\tilde{\mathbf{H}}$ is both closed and consistent.

Generating the hypothesis automaton. When $\tilde{\mathbf{H}}$ is closed and consistent, Alg. 1 generates the hypothesis automaton \mathcal{H} (Line 9) and submits an equivalence query to the oracle (Line 10). If the oracle answers positively, Alg. 1 halts and returns \mathcal{H} . Otherwise, the oracle provides a counterexam-

⁸Indeed, $\tilde{\mathbf{H}}_{pa} \equiv \tilde{\mathbf{H}}_{p'a}$ iff $\tilde{\mathbf{H}}_{pa}(s) = k \otimes \tilde{\mathbf{H}}_{p'a}(s)$ for every $s \in \tilde{S}$ and for some $k \in \mathbb{K} \setminus \{0\}$. Therefore $\tilde{\mathbf{H}}_{pa} \equiv \tilde{\mathbf{H}}_{p'a}$ entails that $\frac{\tilde{\mathbf{H}}(pa, s)}{\bigoplus_{s' \in \tilde{S}} \tilde{\mathbf{H}}(pa, s')} = \frac{k \otimes \tilde{\mathbf{H}}(p'a, s)}{k \otimes \bigoplus_{s' \in \tilde{S}} \tilde{\mathbf{H}}(p'a, s')}$ for every $s \in \tilde{S}$.

Algorithm 2 Subroutines of Alg. 1.

```
1. def MAKECONSISTENT( $\mathbb{O}, \tilde{\mathbf{H}}$ ):
2.   for  $\langle p, p' \rangle \in \tilde{\mathbb{P}} \times \tilde{\mathbb{P}}$ :
3.     if  $\tilde{\mathbf{H}}_p \equiv \tilde{\mathbf{H}}_{p'}$ :
4.       for  $\langle a, s \rangle \in \Sigma \times \tilde{\mathbb{S}}$ :
5.          $Z_{pa} \leftarrow \bigoplus_{s' \in \tilde{\mathbb{S}}} \tilde{\mathbf{H}}(pa, s')$ 
6.          $Z_{p'a} \leftarrow \bigoplus_{s' \in \tilde{\mathbb{S}}} \tilde{\mathbf{H}}(p'a, s')$ 
7.         if  $\tilde{\mathbf{H}}(pa, s) / Z_{pa} \neq \tilde{\mathbf{H}}(p'a, s) / Z_{p'a}$ :
8.            $\tilde{\mathbb{S}} \leftarrow \tilde{\mathbb{S}} \cup \{as\}$ 
9.   COMPLETE( $\mathbb{O}, \tilde{\mathbf{H}}$ )
10. def MAKECLOSED( $\mathbb{O}, \tilde{\mathbf{H}}$ ):
11.   for  $\langle p, a \rangle \in \tilde{\mathbb{P}} \times \Sigma$ :
12.     if  $\exists p' \in \tilde{\mathbb{P}}$  s.t.  $\tilde{\mathbf{H}}_{pa} \neq \tilde{\mathbf{H}}_{p'}$ :
13.        $\tilde{\mathbb{P}} \leftarrow \tilde{\mathbb{P}} \cup \{pa\}$ 
14.   COMPLETE( $\mathbb{O}, \tilde{\mathbf{H}}$ )
15. def COMPLETE( $\mathbb{O}, \tilde{\mathbf{H}}$ ):
16.   for  $p \in \tilde{\mathbb{P}} \circ \Sigma^{\leq 1}$ :
17.     for  $s \in \tilde{\mathbb{S}}$ :
18.        $\tilde{\mathbf{H}}(p, s) \leftarrow \text{MEMBERSHIP}(\mathbb{O}, ps)$ 
```

ple t , which is added to $\tilde{\mathbb{P}}$ along with its prefixes. $\tilde{\mathbf{H}}$ is then updated through membership queries (Lines 12 to 15). The algorithm continues until $\tilde{\mathbf{H}}$ is closed and consistent again.

The correctness of Alg. 1 is shown by the following theorem.

Theorem 3. *Let $\langle \mathbb{K}, \oplus, \otimes, \mathbf{0}, \mathbf{1} \rangle$ be a semifield and Σ an alphabet. Let \mathbb{O} be an oracle for the deterministic regular language $L^\star : \Sigma^* \rightarrow \mathbb{K}$, whose minimal WDFSA has N states. Then, Alg. 1 returns a minimal WDFSA \mathcal{H} generating L^\star in time $\mathcal{O}(N^4 M^2 |\Sigma|)$, where M is the length of the longest counterexample provided by \mathbb{O} .*

Proof. See App. C.3. ■

Relationship Angluin’s (1987) algorithm. We strive to keep the weighted L^* algorithm as faithful as possible to Angluin’s (1987). This brings with it the advantages of familiarity and ease of analysis. A closer look at the pseudocode in Alg. 1 and the proofs of Thms. 2 and 3 reveals that our algorithm should be familiar to anyone acquainted with the unweighted version: The learner uses the same learning loop, procedures of closing and completing the observation table, the same set of calls to the oracle, and an analogous construction of the hypothesis automaton.

L^* and Minimization. Building an automaton whose states correspond to right language equivalence classes (cf. §3) provides the basis for both minimization (Hopcroft and Ullman, 1979; Mohri, 1997) as well as the L^* algorithm. Both algorithms also rely on some sort of weight normalization to identify equivalent states—weight pushing (Mohri, 1997) normalizes the outgoing transition weights in minimization, while the L^* algorithm normalizes the rows of the empirical Hankel matrix to compute the transition weights (Eqs. (4) and (6)). The key distinction between the two algorithms lies in their utilization of the discovered equivalence classes: minimization *merges* equivalent states, whereas L^* incrementally *adds* states to the automaton as it identifies new equivalence classes until termination.

5 Conclusion

We introduce a weighted L^* algorithm, an oracle-based algorithm for learning weighted regular languages, building upon the paradigm pioneered by Angluin (1987). While similar methods have been proposed before, our method is novel in that it learns an *exact* deterministic WFSAs, akin to the original Angluin’s (1987) unweighted version. We highlight language model analysis as a possible application, given that language models describe probability distributions over strings (Icard, 2020; Nowak et al., 2023), i.e., weighted languages.

Limitations

One of the limitations of weighted L^* is that it requires an oracle capable of answering membership and equivalence queries. However, in the case we want to use L^* to study a language model, this is the ideal setting, as we can use the language model itself as the oracle (Weiss et al., 2018; Okudono et al., 2019; Weiss et al., 2019). Another limitation to the applications of our work is that not every language model is efficiently representable as a finite-state machine. For instance, Merrill (2019) shows that LSTMs are strictly more powerful than FSAs. Therefore, in practice, one may have to use a simplified abstraction of the model one aims to learn (Weiss et al., 2021), inevitably reducing the model’s expressivity. Lastly, we note that L^* is not capable of learning *non-deterministic* regular languages, which, in the weighted case, can be a

strict subset of weighted regular languages.⁹

Acknowledgements

Ryan Cotterell acknowledges support from the Swiss National Science Foundation (SNSF) as part of the “The Nuts and Bolts of Language Models” project. Anej Svete is supported by the ETH AI Center Doctoral Fellowship.

References

- Cyril Allauzen and Mehryar Mohri. 2003. [Efficient algorithms for testing the twins property](#). *Journal of Automata, Languages and Combinatorics*, 8(2):117–144.
- Dana Angluin. 1987. [Learning regular sets from queries and counterexamples](#). *Information and Computation*, 75(2):87–106.
- Borja Balle, Xavier Carreras, Franco M. Luque, and Ariadna Quattoni. 2014. [Spectral learning of weighted automata](#). *Machine Learning*, 96(1):33–63.
- Borja Balle and Mehryar Mohri. 2012. [Spectral learning of general weighted automata via constrained matrix completion](#). In *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc.
- Amos Beimel, Francesco Bergadano, Nader H. Bshouty, Eyal Kushilevitz, and Stefano Varricchio. 2000. [Learning functions represented as multiplicity automata](#). *J. ACM*, 47(3):506–530.
- Francesco Bergadano and Stefano Varricchio. 1996. [Learning behaviors of automata from multiplicity and equivalence queries](#). *SIAM Journal on Computing*, 25(6):1268–1280.
- A. W. Biermann and J. A. Feldman. 1972. [On the synthesis of finite-state machines from samples of their behavior](#). *IEEE Transactions on Computers*, C-21(6):592–597.
- Laure Daviaud and Marianne Johnson. 2024. [Feasibility of learning weighted automata on a semiring](#). *Preprint*, arXiv:2309.07806.
- E Mark Gold. 1978. [Complexity of automaton identification from given data](#). *Information and Control*, 37(3):302–320.
- Joshua Goodman. 1996. [Parsing algorithms and metrics](#). In *34th Annual Meeting of the Association for Computational Linguistics*, pages 177–183, Santa Cruz, California, USA. Association for Computational Linguistics.
- John E. Hopcroft and Jeffrey D. Ullman. 1979. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley Publishing Company, Inc.
- Thomas F. Icard. 2020. [Calibrating generative models: The probabilistic chomsky–schützenberger hierarchy](#). *Journal of Mathematical Psychology*, 95:102308.
- Jaap Jumelet and Willem Zuidema. 2023. [Transparency at the source: Evaluating and interpreting language models with access to the true distribution](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 4354–4369, Singapore. Association for Computational Linguistics.
- William Merrill. 2019. [Sequential neural networks as automata](#). In *Proceedings of the Workshop on Deep Learning and Formal Languages: Building Bridges*, pages 1–13, Florence. Association for Computational Linguistics.
- Mehryar Mohri. 1997. [Finite-state transducers in language and speech processing](#). *Computational Linguistics*, 23(2):269–311.
- Mehryar Mohri. 2009. *Weighted Automata Algorithms*, pages 213–254. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Mehryar Mohri and Michael Riley. 2001. [A weight pushing algorithm for large vocabulary speech recognition](#). In *Proc. 7th European Conference on Speech Communication and Technology (Eurospeech 2001)*, pages 1603–1606.
- Franz Nowak, Anej Svete, Alexandra Butoi, and Ryan Cotterell. 2024. [On the representational capacity of neural language models with chain-of-thought reasoning](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Bangkok, Thailand. Association for Computational Linguistics.
- Franz Nowak, Anej Svete, Li Du, and Ryan Cotterell. 2023. [On the representational capacity of recurrent neural language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7011–7034, Singapore. Association for Computational Linguistics.
- Takamasa Okudono, Masaki Waga, Taro Sekiyama, and Ichiro Hasuo. 2019. [Weighted automata extraction from recurrent neural networks via regression on state spaces](#). *Preprint*, arXiv:1904.02931.
- Gail Weiss, Yoav Goldberg, and Eran Yahav. 2018. [Extracting automata from recurrent neural networks using queries and counterexamples](#). In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 5247–5256. PMLR.
- Gail Weiss, Yoav Goldberg, and Eran Yahav. 2019. [Learning deterministic weighted automata with queries and counterexamples](#). In *Advances in Neural Information Processing Systems*, volume 32.

⁹Precisely, this depends on the specific semifield we are considering. The boolean semifield reduces to the unweighted case thus every regular language is deterministic. Conversely, Allauzen and Mohri (2003) provide an example of a language in the real semifield which is not deterministic.

Zarah Weiss, Xiaobin Chen, and Detmar Meurers. 2021. Using broad linguistic complexity modeling for cross-lingual readability assessment. In *Proceedings of the 10th Workshop on NLP for Computer Assisted Language Learning*, pages 38–54, Online. LiU Electronic Press.

A Language Equivalence

Lemma 1. Let L_1 and L_2 be two languages with weights on the same semifield $\langle \mathbb{K}, \oplus, \otimes, \mathbf{0}, \mathbf{1} \rangle$. Further, let us write $L_1 \equiv L_2$ iff $\text{supp}(L_1) = \text{supp}(L_2)$ and there exists a $k \in \mathbb{K} \setminus \{\mathbf{0}\}$ such that $L_1(\mathbf{p}) = k \otimes L_2(\mathbf{p})$ for every $\mathbf{p} \in \text{supp}(L_1)$. Then \equiv is an equivalence relation.

Proof. We show that \equiv is:

- (1) symmetric, since if $\text{supp}(L_1) = \text{supp}(L_2)$ and there exists $k \in \mathbb{K} \setminus \{\mathbf{0}\}$ such that $L_1(\mathbf{p}) = k \otimes L_2(\mathbf{p})$ for every $\mathbf{p} \in \text{supp}(L_1)$, then $\text{supp}(L_2) = \text{supp}(L_1)$ and we can write $L_2(\mathbf{p}) = k^{-1} \otimes L_1(\mathbf{p})$, for every $\mathbf{p} \in \text{supp}(L_2)$.
- (2) transitive, since if $\text{supp}(L_1) = \text{supp}(L_2)$ and $L_1(\mathbf{p}) = k_1 \otimes L_2(\mathbf{p})$ for every $\mathbf{p} \in \text{supp}(L_1)$ and for some $k_1 \in \mathbb{K} \setminus \{\mathbf{0}\}$; and $\text{supp}(L_2) = \text{supp}(L_3)$, $L_2(\mathbf{p}) = k_2 \otimes L_3(\mathbf{p})$ for every $\mathbf{p} \in \text{supp}(L_2)$ and for some $k_2 \in \mathbb{K} \setminus \{\mathbf{0}\}$ then $\text{supp}(L_1) = \text{supp}(L_3)$ and $L_1(\mathbf{p}) = k_1 \otimes k_2 \otimes L_3(\mathbf{p})$.
- (3) reflexive, since $\text{supp}(L_1) = \text{supp}(L_1)$ and $L_1(\mathbf{p}) = \mathbf{1} \otimes L_1(\mathbf{p})$.

Note that if $\langle \mathbb{K}, \oplus, \otimes, \mathbf{0}, \mathbf{1} \rangle$ is a generic semiring that does not satisfy the semifield axioms, then \equiv may not be symmetric, since an inverse for k may not be defined. \blacksquare

B Learning WDFSAs in a Semifields of Fractions

The algorithm presented in the main text learns a weighted deterministic finite-state automaton (WDFSA) over a *semifield* $\langle \mathbb{K}, \oplus, \otimes, \mathbf{0}, \mathbf{1} \rangle$. This is done mainly to allow for a more concise and approachable presentation of the algorithm. In this section, we present another framework that allows for learning WDFSAs over semirings that can be lifted into a *semifield of fractions*.

Definition 1. A *semi-integral domain* is a semiring $\langle \mathbb{K}, \oplus, \otimes, \mathbf{0}, \mathbf{1} \rangle$ where $(\mathbb{K}, \otimes, \mathbf{1})$ is a commutative *cancellative monoid*: for all $x, y \in \mathbb{K}$, if $x \otimes y = \mathbf{0}$, then $x = \mathbf{0}$ or $y = \mathbf{0}$.

A semi-integral domain allows us to define a **semifield of fractions** as follows.

Definition 2. Let $\langle \mathbb{K}, \oplus, \otimes, \mathbf{0}, \mathbf{1} \rangle$ be a semi-integral domain. Its *semifield of fractions* is the tuple $\langle \mathbb{S}, \oplus, \otimes, \mathbf{0}, \mathbf{1} \rangle$ where

$$\mathbb{S} = \left\{ \frac{x}{y} \mid x, y \in \mathbb{K}, y \neq \mathbf{0} \right\} \quad (7)$$

equipped with the operations

$$\frac{x_1}{y_1} \oplus \frac{x_2}{y_2} \stackrel{\text{def}}{=} \frac{x_1 \otimes y_2 \oplus x_2 \otimes y_1}{y_1 \otimes y_2} \quad (8a)$$

$$\frac{x_1}{y_1} \otimes \frac{x_2}{y_2} \stackrel{\text{def}}{=} \frac{x_1 \otimes x_2}{y_1 \otimes y_2} \quad (8b)$$

$$\mathbf{0} \stackrel{\text{def}}{=} \frac{\mathbf{0}}{\mathbf{1}} \quad (8c)$$

$$\mathbf{1} \stackrel{\text{def}}{=} \frac{\mathbf{1}}{\mathbf{1}} \quad (8d)$$

where \oplus and \otimes are the operations of the semi-integral domain.

This definition induces a natural equivalence relation on the semifield of fractions:

$$\frac{x_1}{y_1} \sim \frac{x_2}{y_2} \iff x_1 \otimes y_2 = x_2 \otimes y_1. \quad (9)$$

We can then define the **equivalence class** of $\frac{x}{y}$ as $[\frac{x}{y}] = \left\{ \frac{x'}{y'} \mid \frac{x'}{y'} \sim \frac{x}{y} \right\}$. The **canonical form** of a fraction $\frac{x}{y}$ is the equivalence class $[\frac{x}{y}]$.

Alg. 1 generalizes over a semifield of fractions in the following sense.

1. We lift the weights of the semifield-weighted learned language to the semifield of fractions. This simply entails interpreting the string weights $L^\star(\mathbf{y}) \in \mathbb{K}$ as $\frac{L^\star(\mathbf{y})}{\mathbf{1}}$.
2. The WDFSA is then learned as in the semifield case, but with the weights in the semifield of fractions.

The initial, final, and transition weights of the automaton learned in this manner naturally fall in the semifield of fractions, and thus do not necessarily correspond to any specific element of the original semiring. However, the string weights defined by the learned automaton—and therefore its language—will be the same as in the ground-truth *semiring-weighted* WFSAs.

Theorem 4. *Let L^\star be a weighted regular language over a semi-integral domain $\langle \mathbb{K}, \oplus, \otimes, \mathbf{0}, \mathbf{1} \rangle$. Let \mathcal{H} be the WDFSA learned by Alg. 1 over the corresponding semifield of fractions over $\mathbb{K}, \langle \mathbb{S}, \oplus, \otimes, \mathbf{0}, \mathbf{1} \rangle$. Then, $L_{\mathcal{H}}(\mathbf{y}) \sim \frac{L^\star(\mathbf{y})}{\mathbf{1}}$ for all $\mathbf{y} \in \Sigma^*$.*

Proof. The semifield of fractions is a semifield. Furthermore, since lifting the string weights does not affect the determinism of the language, the lifted language L^\star satisfies the requirements for Alg. 1, which then correctly learns the field-of-fractions-weighted language: $L_{\mathcal{H}}(\mathbf{y}) \sim \frac{L^\star(\mathbf{y})}{\mathbf{1}}$. ■

C Proofs

C.1 Proof of Thm. 1

Theorem 1 (Myhill–Nerode). *Let \mathcal{A} be a semifield-weighted DFSA and $L_{\mathcal{A}}$ its weighted language. Then, $\sim_{L_{\mathcal{A}}}$ induces a finite number of equivalence classes on Σ^* , which equals the number of states of a minimal automaton for $L_{\mathcal{A}}$.*

Proof. Let $\mathcal{A} = (\Sigma, Q, \delta, \lambda, \rho)$ be a semifield-weighted DFSA. Since \mathcal{A} is deterministic, there is at most one initial state q_I . Every string $s \in \Sigma^*$ can be mapped to at most one non-zero-weight path π_s starting in q_I and yielding s . Set $q_s = n(\pi_s)$.

We claim that

$$\begin{aligned} \mathbf{x}^{-1}L_{\mathcal{A}} \equiv \mathbf{y}^{-1}L_{\mathcal{A}} &\iff^1 \forall z \in \Sigma^*, \mathbf{x}z \in \text{supp}(L_{\mathcal{A}}) \text{ iff} \\ &\mathbf{y}z \in \text{supp}(L_{\mathcal{A}}) \iff^2 (q_{\mathbf{x}} = q_{\mathbf{y}} \text{ and } \mathbf{w}(\pi_{\mathbf{x}}) \otimes \mathbf{w}(\pi_{\mathbf{z}}) \neq 0) \text{ or } (\mathbf{w}(\pi_{\mathbf{x}}) = \mathbf{w}(\pi_{\mathbf{y}}) = 0) \end{aligned} \quad (10)$$

\implies)^{1&2} Both are straightforward. \Leftarrow) If $q_{\mathbf{x}} = q_{\mathbf{y}}$. Accordingly, for any $z \in \Sigma^*$, we have

$$\begin{aligned} \mathbf{x}^{-1}L_{\mathcal{A}}(z) &= L_{\mathcal{A}}(\mathbf{x}z) \\ &= \lambda(q_I) \otimes \mathbf{w}(\pi_{\mathbf{x}}) \otimes \mathbf{w}(\pi_{\mathbf{z}}) \otimes \rho(q_z) \\ &= \frac{\mathbf{w}(\pi_{\mathbf{x}})}{\mathbf{w}(\pi_{\mathbf{y}})} \otimes \mathbf{y}^{-1}L_{\mathcal{A}}(z). \end{aligned}$$

And this shows that $\mathbf{x}^{-1}L_{\mathcal{A}} \equiv \mathbf{y}^{-1}L_{\mathcal{A}}$. In conclusion, the claim proves that $|\mathcal{A}/\sim_{L_{\mathcal{A}}}| \leq |Q|$. This shows that the equivalence classes induced by $L_{\mathcal{A}}$ are finite. The proof that their number is equal to $|Q^{\mathcal{M}}|$, where $Q^{\mathcal{M}}$ is the set of states of a minimal automaton for $L_{\mathcal{A}}$ is done by construction, for which we refer the reader to (Mohri, 1997, section 3.7). ■

C.2 Proof of Thm. 2

Theorem 2 (The empirical Hankel Automaton). *Let $\tilde{\mathbf{H}}: \tilde{\mathbb{P}} \circ \Sigma^{\leq 1} \times \tilde{\mathbb{S}} \rightarrow \mathbb{K}$ be a closed and consistent empirical Hankel matrix and let \mathcal{H} be the empirical Hankel automaton induced by $\tilde{\mathbf{H}}$. Then:*

- (1) \mathcal{H} is a well-defined WDFSA.
- (2) $L_{\mathcal{H}}(\mathbf{p}\mathbf{s}) = \tilde{\mathbf{H}}(\mathbf{p}, \mathbf{s})$ for all $\mathbf{p} \in \tilde{\mathbb{P}}$ and $\mathbf{s} \in \tilde{\mathbb{S}}$, meaning that \mathcal{H} is consistent with $\tilde{\mathbf{H}}$.
- (3) \mathcal{H} is minimal.

Proof. Let $\tilde{\mathbf{H}}: \Sigma^* \times \Sigma^* \rightarrow \mathbb{K}$ be a closed and consistent empirical Hankel matrix and let \mathcal{H} be $\tilde{\mathbf{H}}$'s induced WDFSA.

- (1) The set of states Q is a well-defined finite set since, by assumption, the prefixes \tilde{P} are a finite set and $\sim_{\tilde{H}}$ is a well-defined equivalence relation on \tilde{P} . The set of transitions (Eq. (4)) is well-defined since for any $q_{[p]}$ and every symbol $a \in \Sigma$: *i*) $q_{[r(pa)]}$ is uniquely defined *ii*) $q_{[r(pa)]}$ is in Q since the empirical Hankel matrix is closed. Regarding the initial and final states, we just note that since \tilde{P} and \tilde{S} are respectively prefix and suffix closed, they always contain the empty string ε .

Determinism: To see that \mathcal{H} is deterministic, note that:

- (a) There is at most one state $q_{[\varepsilon]}$ with non-zero initial weight (Eq. (5)).
 (b) Following Eq. (4), for every state $q_{[p]}$ and for every input symbol $a \in \Sigma$, there is exactly one state $q_{[r(pa)]}$ such that the transition $q_{[p]} \xrightarrow{a/w} q_{[r(pa)]}$ is in the set of transitions δ . We further, consistency ensures that the next state is independent of the specific choice of representatives for the equivalence classes.

- (2) We begin by showing the following lemmata.

Lemma 2. *Let \tilde{H} be a closed and consistent empirical Hankel matrix, with rows $\tilde{P} \circ \Sigma^{\leq 1}$ and columns \tilde{S} . Then, for every two $p, p' \in \tilde{P}$ such that $\tilde{H}(p, s) = k \otimes \tilde{H}(p', s) \forall s \in \tilde{S}$, we also have that $\tilde{H}(pa, s) = k \otimes \tilde{H}(p'a, s)$, $\forall s \in \tilde{S}$ and for $a \in \Sigma$.*

Proof. Since the table is consistent, we know that if for two $p, p' \in \tilde{P}$, $\tilde{H}(p, s) = k_1 \otimes \tilde{H}(p', s)$ for every $s \in \tilde{S}$ and for some $k_1 \in \mathbb{K} \setminus \{0\}$, then $\tilde{H}(pa, s) = k_2 \otimes \tilde{H}(p'a, s)$ for $s \in \tilde{S}$, $a \in \Sigma$ and for some $k_2 \in \mathbb{K} \setminus \{0\}$. Now, since \tilde{S} is suffix closed, we can chose $s \in \tilde{S}$ such that we can write $s = as'$ for some s' in \tilde{S} , and therefore:

$$\tilde{H}(p, s) = k_1 \otimes \tilde{H}(p', s) \quad (11a, \text{By hypothesis})$$

$$\tilde{H}(p, as') = k_1 \otimes \tilde{H}(p', as') \quad (11b, \text{Rewriting } s \text{ as } as')$$

$$\tilde{H}(pa, s') = k_1 \otimes \tilde{H}(p'a, s') \quad (11c, \text{Skew diagonals of the Hankel matrix are constant})$$

and hence $k_1 = k_2$. ■

Lemma 3. *Let \tilde{H} be a closed and consistent empirical Hankel matrix with rows $\tilde{P} \circ \Sigma^{\leq 1}$ and columns \tilde{S} , and let \mathcal{H} be the empirical hankel automaton generated from \tilde{H} . Then for every $p \in \tilde{P}$ let us denote with π_p the path in \mathcal{H} with initial state $q_{[\varepsilon]}$ and yield p . The following equation holds:*

$$\lambda(q_{[\varepsilon]}) \otimes w(\pi_p) = \bigoplus_{s \in \tilde{S}} \tilde{H}(p, s) \quad (12)$$

Proof. We proceed by induction on the string length.

Base Case. Let $p = \varepsilon$. Then

$$\lambda(q_{[\varepsilon]}) = \bigoplus_{s \in \tilde{S}} \tilde{H}(\varepsilon, s) \quad (13)$$

by the definition of initial weight.

Induction Step. Recall that \tilde{P} is prefix-closed. Let $p \in \tilde{P}$ be a string with length $i \geq 1$ so that we can write $p = p_{<i}a$ for some $a \in \Sigma$ and some $p \in \tilde{P}$. Then our induction hypothesis is that:

$$\lambda(q_{[\varepsilon]}) \otimes w(\pi_{p_{<i}}) = \bigoplus_{s \in \tilde{S}} \tilde{H}(p_{<i}, s) \quad (14)$$

And then we can write

$$\lambda(q_{[\varepsilon]}) \otimes w(\boldsymbol{\pi}_{\mathbf{p}}) = \lambda(q_{[\varepsilon]}) \otimes w(\boldsymbol{\pi}_{\mathbf{p}_{<i}}) \otimes \frac{\bigoplus_{s' \in \tilde{\mathcal{S}}} \tilde{\mathbf{H}}(r(\mathbf{p}_{<i})a, \mathbf{s})}{\bigoplus_{s \in \tilde{\mathcal{S}}} \tilde{\mathbf{H}}(r(\mathbf{p}_{<i}), \mathbf{s})} \quad (15a)$$

$$= \bigoplus_{s \in \tilde{\mathcal{S}}} \tilde{\mathbf{H}}(\mathbf{p}_{<i}, \mathbf{s}) \otimes \frac{\bigoplus_{s' \in \tilde{\mathcal{S}}} \tilde{\mathbf{H}}(r(\mathbf{p}_{<i})a, \mathbf{s})}{\bigoplus_{s \in \tilde{\mathcal{S}}} \tilde{\mathbf{H}}(r(\mathbf{p}_{<i}), \mathbf{s})} \quad (15b, \text{Induction Hypothesis})$$

$$= \bigoplus_{s \in \tilde{\mathcal{S}}} \tilde{\mathbf{H}}(\mathbf{p}_{<i}, \mathbf{s}) \otimes \frac{k \otimes \left(\bigoplus_{s' \in \tilde{\mathcal{S}}} \tilde{\mathbf{H}}(\mathbf{p}_{<i}a, \mathbf{s}) \right)}{k \otimes \left(\bigoplus_{s \in \tilde{\mathcal{S}}} \tilde{\mathbf{H}}(\mathbf{p}_{<i}, \mathbf{s}) \right)} \quad (15c, \text{Lemma 2})$$

$$= \bigoplus_{s \in \tilde{\mathcal{S}}} \tilde{\mathbf{H}}(\mathbf{p}_{<i}, \mathbf{s}) \otimes \frac{\bigoplus_{s' \in \tilde{\mathcal{S}}} \tilde{\mathbf{H}}(\mathbf{p}_{<i}a, \mathbf{s})}{\bigoplus_{s \in \tilde{\mathcal{S}}} \tilde{\mathbf{H}}(\mathbf{p}_{<i}, \mathbf{s})} \quad (15d, \text{Lemma 2})$$

$$= \bigoplus_{s' \in \tilde{\mathcal{S}}} \tilde{\mathbf{H}}(\mathbf{p}_{<i}a, \mathbf{s}) \quad (15e)$$

$$= \bigoplus_{s' \in \tilde{\mathcal{S}}} \tilde{\mathbf{H}}(\mathbf{p}, \mathbf{s}) \quad (15f)$$

■

Lemma 4. Let $\tilde{\mathbf{H}}$ be a closed and consistent empirical Hankel matrix with rows $\tilde{\mathcal{P}} \circ \Sigma^{\leq 1}$ and columns $\tilde{\mathcal{S}}$, and let \mathcal{H} be the empirical hankel automaton generated from $\tilde{\mathbf{H}}$. Then for every $\mathbf{p} \in \tilde{\mathcal{P}}$ and $\mathbf{s} \in \tilde{\mathcal{S}}$, let us denote with $\boldsymbol{\pi}_{\mathbf{p}\mathbf{s}}$ the path in \mathcal{H} with initial state $q_{[\varepsilon]}$ and yield $\mathbf{p}\mathbf{s}$. The following equation holds:

$$\lambda(q_{[\varepsilon]}) \otimes w(\boldsymbol{\pi}_{\mathbf{p}}) \otimes \rho(n(\boldsymbol{\pi}_{\mathbf{p}\mathbf{s}})) = \tilde{\mathbf{H}}(\mathbf{p}, \mathbf{s}) \quad (16)$$

where we recall that $\rho(n(\boldsymbol{\pi}_{\mathbf{p}\mathbf{s}}))$ denotes the final state of $\boldsymbol{\pi}_{\mathbf{p}\mathbf{s}}$.

Proof. **Base Case:** for every $\mathbf{p} \in \tilde{\mathcal{P}}$ we can write.

$$\lambda(q_{[\varepsilon]}) \otimes w(\boldsymbol{\pi}_{\mathbf{p}}) \otimes \rho(q_{r(\mathbf{p})}) = \lambda(q_{[\varepsilon]}) \otimes w(\boldsymbol{\pi}_{\mathbf{p}}) \otimes \frac{\tilde{\mathbf{H}}(r(\mathbf{p}), \varepsilon)}{\bigoplus_{s \in \tilde{\mathcal{S}}} \tilde{\mathbf{H}}(r(\mathbf{p}), \mathbf{s})} \quad (17a)$$

$$= \lambda(q_{[\varepsilon]}) \otimes w(\boldsymbol{\pi}_{\mathbf{p}}) \otimes \frac{\tilde{\mathbf{H}}(\mathbf{p}, \varepsilon)}{\bigoplus_{s \in \tilde{\mathcal{S}}} \tilde{\mathbf{H}}(\mathbf{p}, \mathbf{s})} \quad (17b, \text{Since } \mathbf{p} \sim_{\tilde{\mathcal{H}}} r(\mathbf{p}):)$$

$$= \bigoplus_{s \in \tilde{\mathcal{S}}} \tilde{\mathbf{H}}(\mathbf{p}, \mathbf{s}) \otimes \frac{\tilde{\mathbf{H}}(\mathbf{p}, \varepsilon)}{\bigoplus_{s \in \tilde{\mathcal{S}}} \tilde{\mathbf{H}}(\mathbf{p}, \mathbf{s})} \quad (17c, \text{By Lemma 3:})$$

$$= \tilde{\mathbf{H}}(\mathbf{p}, \varepsilon) \quad (17d)$$

Induction Step: Assume that:

$$\lambda(q_{[\varepsilon]}) \otimes w(\boldsymbol{\pi}_{\mathbf{p}\mathbf{s}'}) \otimes \rho(n(\boldsymbol{\pi}_{\mathbf{p}\mathbf{s}'})) = \tilde{\mathbf{H}}(\mathbf{p}, \mathbf{s}') \quad (18)$$

$\forall \mathbf{p} \in \tilde{\mathcal{P}}$ and for all $\mathbf{s}' \in \tilde{\mathcal{S}}$, such that $|\mathbf{s}'| < k$. Now, let us assume that $\mathbf{s} \in \tilde{\mathcal{S}}$ is a suffix with $|\mathbf{s}| = k$;

since \tilde{S} is suffix closed, we can always write $s = as'$ for some $a \in \Sigma$ and $s' \in \tilde{S}$. Then, we write:

$$\lambda(q_{[\varepsilon]}) \otimes w(\pi_{ps}) \otimes \rho(n(\pi_{ps})) \quad (19a)$$

$$= \lambda(q_{[\varepsilon]}) \otimes w(\pi_{pas'}) \otimes \rho(n(\pi_{pas'})) \quad (19b)$$

$$= \lambda(q_{[\varepsilon]}) \otimes w(\pi_{pa}) \otimes w(\pi_{s'}) \otimes \rho(n(\pi_{pas'})) \quad (19c, \text{ where } \pi_{s'} \text{ is the path with initial state } q_{[r(pa)]} \text{ and yield } s')$$

$$= \frac{\lambda(q_{[\varepsilon]}) \otimes w(\pi_{pa})}{\lambda(q_{[\varepsilon]}) \otimes w(\pi_{r(pa)})} \otimes w(\pi_{r(pa)}) \otimes w(\pi_{s'}) \otimes \rho(n(\pi_{pas'})) \quad (19d, \text{ Since } \tilde{H} \text{ is closed})$$

$$= \frac{\lambda(q_{[\varepsilon]}) \otimes w(\pi_{pa})}{\lambda(q_{[\varepsilon]}) \otimes w(\pi_{r(pa)})} \otimes \tilde{H}(r(pa), s') \quad (19e, \text{ By induction hypothesis})$$

$$= \frac{\bigoplus_{s_i \in \tilde{S}} \tilde{H}(pa, s_i)}{\bigoplus_{s_i \in \tilde{S}} \tilde{H}(r(pa), s_i)} \otimes \tilde{H}(r(pa), s') \quad (19f, \text{ By Lemma 3})$$

$$= \frac{\tilde{H}(pa, s')}{\tilde{H}(r(pa), s')} \otimes \tilde{H}(r(pa), s') \quad (19g, \text{ Since } pa \in [r(pa)])$$

$$= \tilde{H}(pa, s') = \tilde{H}(p, as') \quad (19h, \text{ The empirical Hankel matrix is Skew diagonal})$$

■

This lemma entails the proof of the second point of the theorem.

- (3) By 2) of this theorem, we know that $\tilde{H}(p, s) = L_{\mathcal{H}}(ps)$ for every $p \in \tilde{P}$ and $s \in \tilde{S}$. Hence $\tilde{H}_p : \tilde{S} \rightarrow \mathbb{K}$ is a restriction of $L_{\mathcal{H}} : \Sigma^* \rightarrow \mathbb{K}$ from Σ^* to \tilde{S} , and following the definition of the relations $\sim_{\tilde{H}}$ and $\sim_{L_{\mathcal{H}}}$, we have that $|\tilde{P}/\sim_{\tilde{H}}| \leq |\Sigma^*/\sim_{L_{\mathcal{H}}}|$.

Now, let us denote by N the number of states of \mathcal{H} , which by construction we know is equal to $|\tilde{P}/\sim_{\tilde{H}}|$. Therefore we can write $|\Sigma/\sim_{L_{\mathcal{H}}}| \geq |\tilde{P}/\sim_{\tilde{H}}| = N$. Then by the Myhill–Nerode theorem (Thm. 1), we have the number of states of the minimal automaton for $L_{\mathcal{H}}$ is at least N . But \mathcal{H} has N states and hence \mathcal{H} is minimal.

■

C.3 Proof of Thm. 3

In order to prove Thm. 3, we first give the following sequence of lemmata. Throughout the section, L^{\star} denotes an unknown deterministic semifield-weighted regular language. Further, we write \mathcal{A}^{\star} to denote an arbitrary, but fixed, minimal WFSAs generating L^{\star} with N states.

Lemma 5 (Equivalence classes increase #1). *Let \tilde{P} be a prefix-closed set of prefixes and \tilde{S} a suffix-closed set of suffixes. Let \tilde{H} be the empirical Hankel matrix of L^{\star} , whose rows and columns are indexed with $\tilde{P} \circ \Sigma^{\leq 1}$ and \tilde{S} respectively. Then, the following statements hold:*

- (1) *If \tilde{H} is not consistent, Alg. 1 adds a suffix to \tilde{S} and the number of equivalence classes of \tilde{P} increases.*
- (2) *If \tilde{H} is not closed, Alg. 1 adds a prefix to \tilde{P} and the number of equivalence classes of \tilde{P} increases.*

Proof. We prove each of the statements in turn.

- (1) If the empirical Hankel matrix is not consistent, MAKECONSISTENT (Alg. 1, Line 5) finds two prefixes $p, p' \in \tilde{P}$ such that $\tilde{H}_p \equiv \tilde{H}_{p'}$ but $\tilde{H}_p(as) \not\equiv \tilde{H}_{p'}(as)$ for some $a \in \Sigma$ and $s \in \tilde{S}$, and adds as to \tilde{S} . After adding the suffix s to \tilde{S} , we have that $\tilde{H}_p \not\equiv \tilde{H}_{p'}$, and therefore an equivalence class has been divided in two.
- (2) If \tilde{H} is not closed, MAKECLOSED (Alg. 1, Line 7) finds $p \in \tilde{P}$ and $a \in \Sigma$ such that $\tilde{H}_{pa} \not\equiv \tilde{H}_{p'}$ for every $p' \in \tilde{P}$ and adds pa to \tilde{P} . Since there is no $p' \neq pa$ in \tilde{P} , such that $p' \sim_{\tilde{H}} pa$, it follows that a new equivalence class $[pa]$ is added to $\tilde{P}/\sim_{\tilde{H}}$.

■

Lemma 6 (Equivalence classes increase #2). *Let \tilde{P} be a prefix-closed set of prefixes and \tilde{S} a suffix-closed set of suffixes. Let \tilde{H} be the empirical Hankel matrix of L^\star , whose rows and columns are indexed with $\tilde{P} \circ \Sigma^{\leq 1}$ and \tilde{S} respectively. Then, each time the oracle replies with a counterexample t , the number of equivalence classes of \tilde{P} increases.*

Proof. When we add the counterexample t and its prefixes to \tilde{P} (Alg. 1, Line 13) three outcomes are possible:

- (1) The new empirical Hankel matrix \tilde{H}' is not closed and/or consistent. In this case, new rows or columns are added, which, as shown by Lemma 5, increases the number of equivalence classes in \tilde{P} .
- (2) The new empirical Hankel matrix \tilde{H}' is closed and consistent and \tilde{P} has at least one more equivalence class than in \tilde{H} .
- (3) The new empirical Hankel matrix \tilde{H}' is closed and consistent and \tilde{P} has the same number of equivalence classes as in \tilde{H} . This implies that $t_{[1:i]} \in [p_i]$ for $i = 1, \dots, |t|$, and for some $p_i \in \tilde{P}$, that was already in \tilde{P} before the counterexample and its prefixes were added. Since the choice of representatives for the construction of the hypothesis automaton is fixed, this entails that the automaton \mathcal{H}' generated from \tilde{H}' is equal to the automaton \mathcal{H} generated from \tilde{H} . However, this implies that $L_{\mathcal{H}} = L_{\mathcal{H}'}$ and therefore t was not a counterexample.

Therefore, both in case 1) and 2), the number of equivalence classes of \tilde{P} increases, and case 3) cannot occur, as it would contradict the hypothesis that t is a counterexample. ■

Lemma 7 (Upper bound on the number of equivalence classes). *Let L be a semifield-weighted deterministic regular language, and let \tilde{H} be an empirical Hankel matrix with prefixes \tilde{P} and suffixes \tilde{S} , containing observations from L . Then the number of equivalence classes induced by \sim_L on Σ^* upper bounds the number of equivalence classes induced by $\sim_{\tilde{H}}$ on \tilde{P} .*

Proof. For every $p_1, p_2 \in \tilde{P}$, we show that $p_1 \sim_L p_2 \implies p_1 \sim_{\tilde{H}} p_2$. Indeed $p_1 \sim_L p_2$ means that $p_1^{-1}L \equiv p_2^{-1}L$ and $\text{supp}(p_1^{-1}L) = \text{supp}(p_2^{-1}L)$. Following the definition of the relation \equiv between weighted languages (cf. §3), we then have that:

$$\text{supp}(\tilde{H}_{p_1}) = \text{supp}(p_1^{-1}L) \cap \tilde{S} \quad \text{since } \tilde{H}_{p_1} : \tilde{S} \rightarrow \mathbb{K} \text{ and } \tilde{H}_{p_1} = p_1^{-1}L(s), \forall s \in \tilde{S} \quad (20a)$$

$$= \text{supp}(p_2^{-1}L) \cap \tilde{S} \quad \text{since } \text{supp}(p_1^{-1}L) = \text{supp}(p_2^{-1}L) \quad (20b)$$

$$= \text{supp}(\tilde{H}_{p_2}) \quad (20c)$$

and further, $\forall s \in \text{supp}(\tilde{H}_{p_1})$

$$\tilde{H}_{p_1}(s) = p_1^{-1}L(s) \quad (21a)$$

$$= k \otimes p_2^{-1}L(s) \text{ since } p_1^{-1}L(s) = k \otimes p_2^{-1}L(s) \text{ for some } k \in \mathbb{K} \setminus \{0\}, \forall s \in \text{supp}(p_1^{-1}L) \quad (21b)$$

$$= k \otimes \tilde{H}_{p_2}(s) \quad (21c)$$

For a fixed $k \in \mathbb{K} \setminus \{0\}$. Therefore $\tilde{H}_{p_1} \equiv \tilde{H}_{p_2}$ and $p_1 \sim_{\tilde{H}} p_2$. Since $p_1 \sim_L p_2 \implies p_1 \sim_{\tilde{H}} p_2$, it follows that $|\tilde{P} / \sim_{\tilde{H}}| \leq |\Sigma^* / \sim_L|$ ■

Theorem 3. *Let $\langle \mathbb{K}, \oplus, \otimes, \mathbf{0}, \mathbf{1} \rangle$ be a semifield and Σ an alphabet. Let \mathcal{O} be an oracle for the deterministic regular language $L^\star : \Sigma^* \rightarrow \mathbb{K}$, whose minimal WDFSA has N states. Then, Alg. 1 returns a minimal WDFSA \mathcal{H} generating L^\star in time $\mathcal{O}(N^4 M^2 |\Sigma|)$, where M is the length of the longest counterexample provided by \mathcal{O} .*

Proof. We split the remainder of the proof into the proof of termination and the runtime analysis.

Termination. We will show that, at any iteration of the loop on Line 2 in Alg. 1, the algorithm either terminates or gets strictly closer to termination. Concretely, at each execution of the loop, three events can occur:

- (1) The empirical Hankel matrix $\tilde{\mathbf{H}}$ is not consistent and/or closed, resulting in the execution of the inner while loop (Alg. 1, Line 3).
- (2) The oracle \mathbb{O} replies positively to the equivalence query positively on Line 10, halting the execution.
- (3) The oracle \mathbb{O} replies with a counterexample.

Clearly, events (2) and (3) do not occur at the same time. By Lemmata 5 and 6, we know that (1) and (3) result an increase of the number of equivalence classes of $\tilde{\mathbb{P}}$.

Assume that after some number of iterations of the outer while loop, the prefixes $\tilde{\mathbb{P}}$ are partitioned into N equivalence classes and that the algorithm has not terminated yet. In the next iteration, events (1) and (3) would increase the number of equivalence classes of $\tilde{\mathbb{P}}$ to at least $N + 1$. Let \mathcal{H} be the empirical Hankel automaton constructed during this iteration of the algorithm and let $L_{\mathcal{H}}$ be its language. By Thm. 2, we know that the \mathcal{H} is consistent with the empirical Hankel matrix $\tilde{\mathbf{H}}$, meaning that $\sim_{\tilde{\mathbf{H}}}$ and $\sim_{L_{\mathcal{H}}}$ induce isomorphic equivalence classes. Now, by Lemma 7 we know that

$$N + 1 \leq |\Sigma^* / \sim_{L_{\mathcal{H}}}| = |\tilde{\mathbb{P}} / \sim_{\tilde{\mathbf{H}}}| \leq |\Sigma^* / L^*| = N. \quad (23)$$

Another execution therefore brings us to a contradiction with the Myhill–Nerode theorem (Thm. 1) since the number of states of the minimal automaton for L^* is N . To summarize, since event (3) results in a contradiction, event (2) must occur, meaning that the algorithm halts and return an automaton generating L^* after at most N iterations of the outer loop. The minimality of the returned automaton follows from Thm. 2.

Note that the same argument can be applied to show that the inner while loop (Alg. 1, Line 3) can be executed at most N times throughout the entire execution of L^* , since by Lemma 5 at each iteration, the number of equivalence classes increases.

Runtime Analysis. First, we analyze the sizes of $\tilde{\mathbb{P}}$ and $\tilde{\mathbb{S}}$. The size of $\tilde{\mathbb{P}}$ is upper bounded by $|\tilde{\mathbb{P}}| \leq N + MN$: The empirical Hankel matrix can be found to be not closed at most N times—and in such case one string is added to $\tilde{\mathbb{P}}$ —and the oracle can reply with a counterexample at most N times—and in any such case, at most M strings are added to the matrix. The size of $\tilde{\mathbb{S}}$ is at most N , since the empirical Hankel matrix can be found to be not consistent at most N times.

Now, we know that the outer while loop of Alg. 1 (Line 2) can be repeated at most N times, and similarly the inner while loop (Line 3) can be repeated at most N times during the entire execution of the outer loop. Of all the operations performed during the execution of the outer loop, the most expensive is MAKECLOSED, as it requires a double pass over $\tilde{\mathbb{P}}$ to check if the table is consistent. Therefore the total runtime is:

$$\mathcal{O}\left(\left((N + MN)^2 |\Sigma| N\right) N\right) = \mathcal{O}\left(N^4 M^2 |\Sigma|\right) \quad (24)$$

Where $\mathcal{O}\left((N + MN)^2 |\Sigma| N\right)$ is the cost of MAKECLOSED, and N is the maximum number of calls to MAKECLOSED. ■