

On the Representational Capacity of Neural Language Models with Chain-of-Thought Reasoning

Franz Nowak* Anej Svete* Alexandra Butoi Ryan Cotterell

{franz.nowak, anej.svete, alexandra.butoi, ryan.cotterell}@inf.ethz.ch

ETH zürich

Abstract

The performance of modern language models (LMs) has been improved by chain-of-thought (CoT) reasoning, i.e., the process of generating intermediate results that guide the model towards a final answer. A possible explanation for this improvement is that CoT reasoning extends an LM’s computational power, as RNNs and transformers with additional scratch space are known to be Turing complete. Comparing LMs to Turing machines, however, introduces a category error—Turing machines decide language membership, whereas LMs define *distributions* over strings. To bridge this gap, we formalize CoT reasoning in a probabilistic setting. We present several results on the representational capacity of recurrent and transformer LMs with CoT reasoning, showing that they can represent the same family of distributions over strings as *probabilistic* Turing machines.

 <https://github.com/rycolab/cot-lms>

1 Introduction

Motivated by how humans solve complex problems by noting down intermediary results, Wei et al. (2022b) introduced **chain-of-thought** (CoT) reasoning.¹ CoT reasoning helps language models (LMs) solve reasoning tasks by allowing them to store intermediary results in a scratch space that is *not* part of the final output. The empirical success of CoT reasoning has made it an established component of modern neural LMs (Nye et al., 2021; Wei et al., 2022a,b; Suzgun et al., 2022; Kojima et al., 2023; Wies et al., 2023)—seemingly overnight. Such empirical success motivates a thorough understanding of the abilities of CoT-augmented LMs. While existing theoretical treatments have shed light on some aspects of this framework, we are still far from a concrete theoretical understanding of CoT (Feng et al., 2023).

CoT reasoning can be interpreted as encouraging LMs to perform additional sequential computation

*Equal contribution.

¹We use the more general term *CoT reasoning* over the original term *CoT prompting* as prompting is just one way to elicit CoT reasoning (Wang and Zhou, 2024).




Figure 1: A schematic overview of our main results, showing which neural language models extended with CoT reasoning (left) correspond to which probabilistic models of computation (right).

steps while storing the result of that computation. Naturally, this suggests that CoT reasoning should be formalized in a manner that exploits our understanding of well-known, stateful models of computation. Theory of computation provides a natural toolbox for this avenue of exploration. Indeed, the internal configurations of many neural language models have been previously linked to intermediate steps in neural LMs (Pérez et al., 2021; Merrill and Sabharwal, 2024; Feng et al., 2023).

A **language model** is definitionally a distribution over strings from some fixed alphabet.² This definition is somewhat at odds with the way existing work has investigated the representational capacity of CoT-augmented LMs. Firstly, most results in this area concern unweighted language *recognition*, which is different from the inherently probabilistic task of language modeling. Secondly, Turing completeness results such as Pérez et al.’s (2021) construction—a seminal result showing the Turing completeness of transformer LMs—usually require *additional symbols* not present in the original alphabet to correctly simulate a Turing machine. Because the output alphabets of the transformer LM and the Turing machine it simulates fundamentally differ, it is difficult

²An alphabet is a finite, non-empty set.

to discuss their *equivalence*. This motivates a more fine-grained treatment of the representational capacity of neural LMs in which: (1) LMs are treated as *probabilistic* models of computation, i.e., they assign weights to strings rather than deciding language membership, and (2) language model equivalence is defined on the output string level, and the analysis takes into account the additional information required to be encoded in additional outputs to achieve the equivalence.

We address the first point by analyzing CoT-augmented LMs using *probabilistic* models of computation, which provide a convenient way of describing families of distributions over strings with standard models of sequential reasoning. To address the second point, we define a new type of relationship between language models, which we call **regular reducibility**. Intuitively, an LM is regularly reducible into another one if the strings generated by the latter are sufficiently simple transformations of those generated by the former. An instance of such a transformation is the deletion of the intermediary computation steps of CoT reasoning. Formalizing CoT reasoning in this framework, we find that it increases the computational power in both RNN and transformer LMs. Concretely, we find that CoT-augmented constant-precision RNN LMs are equivalent to probabilistic finite-state automata (PFSAs). This is in contrast to the constant-precision RNN LMs *without* CoT reasoning, which are equivalent to *deterministic* PFSAs (Svete and Cotterell, 2023). Additionally, we show how Turing-complete linear-precision RNN LMs can be thought of as performing CoT reasoning. Finally, we show that both linearly bounded precision RNN LMs and logarithmically bounded precision decoder-only transformer LMs with CoT-reasoning can simulate any probabilistic Turing machine. Taken together, our results frame CoT reasoning in pure language modeling terms and describe its intuitive and formal connections to probabilistic models of computation.

2 Preliminaries

In our analysis, we assume rational (rather than real) arithmetic for all definitions and computations.

2.1 Language Models

A **formal language** L is a subset of the Kleene closure Σ^* of some alphabet Σ . We call an element

\mathbf{y} of Σ^* a **string**. We denote the empty string by ε , and we assume that $\varepsilon \notin \Sigma$. A **discrete semimeasure** over Σ^* is a function $\mu: \Sigma^* \rightarrow [0, 1]$ such that $\sum_{\mathbf{y} \in \Sigma^*} \mu(\mathbf{y}) \leq 1$ (Bauwens, 2013; Icard, 2020). If the semimeasure of all strings sums to one, i.e., $\sum_{\mathbf{y} \in \Sigma^*} \mu(\mathbf{y}) = 1$, then μ is called a **probability measure**.³ Finally, for any alphabet Σ , we define the set $\Sigma_\varepsilon \stackrel{\text{def}}{=} \Sigma \cup \{\varepsilon\}$.

Definition 2.1. A *language model (LM)* p is a *semimeasure* over Σ^* . If p is a *probability measure*, it is called a **tight language model**.

Definition 2.2. Two LMs p and q are **weakly equivalent** if $p(\mathbf{y}) = q(\mathbf{y})$ for all $\mathbf{y} \in \Sigma^*$.

Most modern LMs are autoregressive, i.e., they define $p(\mathbf{y})$ through conditional distributions of the next symbol given the string produced so far and the probability of ending the string, i.e.,

$$p(\mathbf{y}) \stackrel{\text{def}}{=} p(\text{EOS} \mid \mathbf{y}) \prod_{t=1}^{|\mathbf{y}|} p(y_t \mid \mathbf{y}_{<t}). \quad (1)$$

Here, EOS denotes the special end-of-string symbol, which specifies that the generation of a string has halted. The inclusion of EOS allows (but does not guarantee) that a language model p , autoregressively, constitutes a probability measure over Σ^* (Du et al., 2023); a model defined as in Eq. (1) may sum to *less than* 1 in a pathological case. For any alphabet Σ , we define the set $\bar{\Sigma} \stackrel{\text{def}}{=} \Sigma \cup \{\text{EOS}\}$. The conditional probability distributions are usually defined based on *vectorial* representations of $\mathbf{y}_{<t}$ computed by a **language encoder** $\text{enc}: \Sigma^* \rightarrow \mathbb{R}^D$ (Chan et al., 2024).

Definition 2.3. A **representation-based LM** is any LM that can be written as an autoregressive language model (Eq. (1)) where the conditional distributions over the next symbol $\bar{y}_t \in \bar{\Sigma}$ are given by

$$p(\bar{y}_t \mid \mathbf{y}_{<t}) \stackrel{\text{def}}{=} \mathbf{f}(\mathbf{E} \text{enc}(\mathbf{y}_{<t}))_{\bar{y}_t}, \quad (2)$$

where $\text{enc}: \Sigma^* \rightarrow \mathbb{R}^D$ is a language encoder, $\mathbf{E} \in \mathbb{R}^{|\bar{\Sigma}| \times D}$ is an output matrix, and \mathbf{f} is a projection function.⁴

³This definition differs from Li and Vitányi’s (2008) who instead define semimeasures over *prefix strings*.

⁴A common choice for \mathbf{f} is the softmax. Since our analyses use rational arithmetic, we instead opt for sparsemax (Martins and Astudillo, 2016). However, all of our results can be extended to the use of the more common softmax function through the use of log activations and the extended real numbers $\mathbb{R} \cup \{-\infty, \infty\}$ (Svete et al., 2024).

2.2 Regular Language Models

Probabilistic finite-state automata are a well-understood probabilistic computational model.

Definition 2.4. A *probabilistic finite-state automaton* (PFSA) is a tuple $(\Sigma, Q, \lambda, \rho, \delta)$ where Σ is an alphabet, Q is a finite set of states, $\delta \subseteq Q \times \Sigma \times \mathbb{Q}_{\geq 0} \times Q$ is a finite set of weighted transitions where we write transitions $(q, y, w, q') \in \delta$ as $q \xrightarrow{y/w} q'$, and $\lambda, \rho: Q \rightarrow \mathbb{Q}_{\geq 0}$ are functions that assign each state its initial and final weight, respectively. Moreover, for all states $q \in Q$, δ, λ and ρ satisfy $\sum_{q \in Q} \lambda(q) = 1$, and $\sum_{q \xrightarrow{y/w} q' \in \delta} w + \rho(q) = 1$.

Next, we will define some basic concepts related to PFSA. A PFSA $\mathcal{A} = (\Sigma, Q, \lambda, \rho, \delta)$ is **deterministic** if $|\{q \mid \lambda(q) > 0\}| = 1$ and, for every $q \in Q, y \in \Sigma$, there is at most one $q' \in Q$ such that $q \xrightarrow{y/w} q' \in \delta$ with $w > 0$.⁵ Any state q where $\lambda(q) > 0$ is called an **initial state**, and if $\rho(q) > 0$, it is called a **final state**. A **path** π of length N is a sequence of subsequent transitions in \mathcal{A} , denoted as

$$q_1 \xrightarrow{y_1/w_1} q_2 \xrightarrow{y_2/w_2} q_3 \cdots q_N \xrightarrow{y_N/w_N} q_{N+1}. \quad (3)$$

The **yield** of a path is $\mathbf{s}(\pi) \stackrel{\text{def}}{=} y_1 \cdots y_N$. The **prefix weight** \tilde{w} of a path π is the product of the transition and initial weights, whereas the **weight** of a path additionally has the final weight multiplied in. In symbols, this means

$$\tilde{w}(\pi) \stackrel{\text{def}}{=} \prod_{n=0}^N w_n, \quad (4) \quad \mathbf{w}(\pi) \stackrel{\text{def}}{=} \prod_{n=0}^{N+1} w_n, \quad (5)$$

with $w_0 \stackrel{\text{def}}{=} \lambda(q_1)$ and $w_{N+1} \stackrel{\text{def}}{=} \rho(q_{N+1})$. We write $\Pi(\mathcal{A})$ for the set of all paths in \mathcal{A} and we write $\Pi(\mathcal{A}, \mathbf{y})$ for the set of all paths in \mathcal{A} with yield \mathbf{y} . The sum of weights of all paths that yield a certain string $\mathbf{y} \in \Sigma^*$ is called the **stringsum**, which we write as

$$\mathcal{A}(\mathbf{y}) \stackrel{\text{def}}{=} \sum_{\pi \in \Pi(\mathcal{A}, \mathbf{y})} \mathbf{w}(\pi). \quad (6)$$

The stringsum gives the probability of the string \mathbf{y} . This way, PFSA induce a particularly well-understood family of LMs.

⁵In this paper, we do *not* distinguish between a transition for a given symbol with weight $w = 0$ and the absence of a transition for that symbol. That is, we assume there always exists a transition $q \xrightarrow{y/w} q' \in \delta$ for any $q, q' \in Q$ and $y \in \Sigma$, albeit possibly with $w = 0$. Such a choice turns out to be useful in our technical exposition.

Definition 2.5. An LM p is a **regular LM** if there exists a PFSA \mathcal{A} whose induced language model $p_{\mathcal{A}}$ is weakly equivalent to p .

PFSA and non-determinism. Although PFSA share many properties with unweighted (boolean-weighted) finite-state automata, one important difference relates to determinization. In the unweighted case, the class of deterministic and non-deterministic FSAs are equivalent, i.e., any non-deterministic FSA has an equivalent deterministic FSA that accepts the same language. This result, however, does not hold for PFSA: There exist PFSA that admit no deterministic equivalent (Mohri, 1997; Buchsbaum et al., 2000), meaning that non-deterministic PFSA are strictly more expressive than deterministic ones.

Computing string probabilities under non-determinism. Autoregressive LMs (cf. Eq. (1)) and PFSA fall under the larger framework of models that specify probability distributions *implicitly* (Icard, 2020).⁶ However, in autoregressive neural LMs with fixed precision, only one sequence of computational actions can yield a particular string, meaning that they can only model deterministic weighted regular languages (Svete and Cotterell, 2023). On the other hand, non-deterministic PFSA compute string probabilities by *additively* combining the probabilities of all paths that yield the same string, and hence lie outside the grasp of such neural models.⁷ As we show later, CoT reasoning provides a principled way to overcome this limitation and allow fixed-precision neural models to simulate non-deterministic automata.

2.3 Regular Functions

We now define a finite-state machine that, in addition to scanning, also *outputs* strings.

Definition 2.6. A *finite-state transducer* (FST) is a 6-tuple $\mathcal{T} = (Q, \Sigma, \Xi, I, F, \delta)$, where Q is a finite set of **states**, Σ is an alphabet of **input symbols**, Ξ is an alphabet of **output symbols**, $I, F \subseteq Q$ are sets of initial and final states, respectively, and $\delta \subseteq Q \times \Sigma_{\varepsilon} \times \Xi_{\varepsilon} \times Q$ is a set of transitions.

Similar to PFSA transitions, we write FST transitions of the form $(q, x, y, q') \in \delta$ as $q \xrightarrow{x:y} q'$.

⁶Other examples of such models include hidden Markov models, which are equivalent to PFSA (Icard, 2020).

⁷Note that RNN LMs with linearly bounded precision can simulate non-deterministic PFSA in real-time by encoding a probability distribution over the PFSA's current state in the hidden state of the RNN (Svete et al., 2024).

Finally, we give the following definition which we will use to formalize CoT reasoning.

Definition 2.7. A *regular relation* is a relation $\phi \subseteq \Sigma^* \times \Delta^*$ that is representable by an FST. If ϕ is a (partial) function, it is called a *regular function*.

Like all relations, regular relations $\phi \subseteq \Sigma^* \times \Delta^*$ have a trivial inverse $\phi^{-1} \subseteq \Delta^* \times \Sigma^*$.

2.4 Turing Machines

We consider the following definition of a probabilistic Turing machine.

Definition 2.8. A *probabilistic Turing machine (PTM)* is a two-tape machine with a working tape and an output tape specified by the 6-tuple $\mathcal{M} = (Q, \Sigma, \Gamma, \delta, q_i, q_f)$, where Q is a finite set of states, Σ is an output alphabet, Γ is a tape alphabet including the blank symbol \sqcup , $q_i, q_f \in Q$ are the initial and final states, and $\delta \subseteq Q \times \Gamma \times \Sigma_\epsilon \times \{L, R\} \times \mathbb{Q}_{\geq 0} \times Q \times \Gamma$ is a rationally weighted transition relation. L and R signify the PTM head moving left (L) or right (R) on the tape after a transition. We write transitions $(q, \gamma, y, d, w, q', \gamma') \in \delta$ as $(q, \gamma) \xrightarrow{y, d/w} (q', \gamma')$. Moreover, we require that for any given $q \in Q, \gamma \in \Gamma$, the weights satisfy $\sum_{(q, \gamma) \xrightarrow{y, d/w} (q', \gamma') \in \delta} w = 1$.

A transition should be interpreted as follows: When \mathcal{M} in state q , reads γ on the working tape, writes y on the output tape, writes γ' on the working tape, move the head in direction d on the working tape. Each computation step randomly selects a transition according to its weight w .⁸

This definition of Turing machines straightforwardly induces a semimeasure over strings (Nowak et al., 2023, Remark 2.2). It is sometimes easier to prove claims about Turing machines through another equivalent machine (Hopcroft et al., 2001). The probabilistic two-stack pushdown automaton is one example (Nowak et al., 2023). See App. C.3 for an overview.

2.5 Recurrent Neural Language Models

Recurrent neural LMs are LMs whose conditional probabilities are given by an RNN.⁹

⁸For more details, see App. C.1.

⁹Throughout this paper, we will focus on Elman RNNs (Elman, 1990) as they are the easiest to analyze and a special case of more common networks, e.g., those based on long short-term memory (LSTM; Hochreiter and Schmidhuber, 1997) and gated recurrent units (GRUs; Cho et al., 2014).

Definition 2.9. An *Elman RNN* $\mathcal{R} = (\Sigma, \sigma, D, \mathbf{U}, \mathbf{V}, \mathbf{b}, \eta)$ is an RNN with the following hidden state recurrence:

$$\mathbf{h}_0 = \eta \quad (t = 0), \quad (7a)$$

$$\mathbf{h}_t = \alpha(\mathbf{U}\mathbf{h}_{t-1} + \mathbf{V}\mathbf{r}(y_t) + \mathbf{b}) \quad (t > 0), \quad (7b)$$

where $\mathbf{h}_t \in \mathbb{Q}^D$ is the state vector¹⁰ at time step t , $\eta \in \mathbb{Q}^D$ is an initialization parameter, $y_t \in \Sigma$ is the input symbol at time step t , $\mathbf{r}: \Sigma \rightarrow \mathbb{Q}^R$ is a symbol representation function, $\mathbf{U} \in \mathbb{Q}^{D \times D}$ and $\mathbf{V} \in \mathbb{Q}^{D \times R}$ are parameter matrices, $\mathbf{b} \in \mathbb{Q}^D$ is a bias vector, and $\alpha: \mathbb{Q}^D \rightarrow \mathbb{Q}^D$ is an element-wise, non-linear activation function.¹¹

Because \mathbf{h}_t hides the symbols consumed by the Elman RNN, we also use the evocative notation $\mathbf{h}(\mathbf{y})$ to denote the result of the application of Eq. (7b) over the string $\mathbf{y} = y_1 \cdots y_t$. The notation $\mathbf{h}(\mathbf{y})$ makes it clear that an RNN LM implicitly defines a language encoder.

Definition 2.10. A *representation-based LM* is called an *Elman LM* if its representation function is defined by the hidden state of an Elman RNN $\text{enc}(\mathbf{y}_{<t}) \stackrel{\text{def}}{=} \mathbf{h}(\mathbf{y}_{<t})$.

The most common choice for the projection function \mathbf{f} is the softmax, whose limitation is that it implies the LM has full support, i.e., an Elman LM with a softmax projection assigns positive probability to all strings in Σ^* . To construct non-full-support Elman LMs, we instead use the **sparsemax** (Martins and Astudillo, 2016):

$$\text{sparsemax}(\mathbf{x}) \stackrel{\text{def}}{=} \underset{\mathbf{z} \in \Delta^{N-1}}{\text{argmin}} \|\mathbf{z} - \mathbf{x}\|_2^2. \quad (8)$$

In contrast to the softmax function, sparsemax is the **identity** on Δ^{N-1} , i.e., we have $\text{sparsemax}(\mathbf{x}) = \mathbf{x}$ for $\mathbf{x} \in \Delta^{N-1}$.

2.6 Neural Networks and Numerical Precision

All implementations of LMs on modern hardware require representations to be fixed-precision floating-point numbers or arbitrary-precision (rational) numbers. In the case of arbitrary precision, an important consideration in processing strings $\mathbf{y} \in \Sigma^*$ is the number of bits required to store the representations and how the number of bits scales with the length of the string, $|\mathbf{y}|$. This motivates the following definition of precision.

¹⁰Throughout this paper all vectors are column vectors.

¹¹Common examples of α include the Heaviside function $H(x) \stackrel{\text{def}}{=} \mathbb{1}\{x > 0\}$, the sigmoid function $\sigma(x) \stackrel{\text{def}}{=} \frac{1}{1 + \exp(-x)}$, and the ReLU $(x) \stackrel{\text{def}}{=} \max(0, x)$.

Definition 2.11. The *precision* $\psi(\mathbf{y})$ of a representation-based neural LM is the number of bits required to represent the entries of $\text{enc}(\mathbf{y})$:

$$\psi(\mathbf{y}) \stackrel{\text{def}}{=} \max_{d \in [D]} \min_{\substack{p, q \in \mathbb{N}, \\ \frac{p}{q} = \text{enc}(\mathbf{y})_d}} \lceil \log_2 p \rceil + \lceil \log_2 q \rceil. \quad (9)$$

We say that a representation-based LM is of

- **constant precision** if $\psi(\mathbf{y}) = \mathcal{O}(1)$, i.e., if $\psi(\mathbf{y}) \leq C$ for all $\mathbf{y} \in \Sigma^*$ and some $C \in \mathbb{R}$,
- **logarithmically bounded precision** if $\psi(\mathbf{y}) = \mathcal{O}(\log |\mathbf{y}|)$, i.e., if there exist $T_0 \in \mathbb{N}$ and $C \in \mathbb{R}$ such that for all $\mathbf{y} \in \Sigma^*$ with $|\mathbf{y}| \geq T_0$, $\psi(\mathbf{y}) \leq C \log_2 |\mathbf{y}|$,
- **linearly bounded precision** if $\psi(\mathbf{y}) = \mathcal{O}(|\mathbf{y}|)$, i.e., there exist $T_0 \in \mathbb{N}$ and $C \in \mathbb{R}$ such that for all $\mathbf{y} \in \Sigma^*$ with $|\mathbf{y}| \geq T_0$, $\psi(\mathbf{y}) \leq C|\mathbf{y}|$, and
- **unbounded precision** if $\psi(\mathbf{y})$ cannot be bounded by a function of $|\mathbf{y}|$.

The constructions in the existing literature range from constant to unbounded precision. The RNNs considered by Svete and Cotterell’s (2023) encoding of deterministic PFSAs, for example, results in an RNN that is of constant precision. Weiss et al. (2018), Merrill (2019) and Merrill et al. (2020) consider models with logarithmically bounded precision. In contrast, articles treating the Turing completeness of neural networks (Siegelmann and Sontag, 1992; Nowak et al., 2023) require unbounded precision to be able to represent the fact that a Turing machine may fail to halt. Naturally, our constructions of Turing complete LMs will also require unbounded precision.

2.7 Transformer Language Models

Transformer LMs compute the conditional distributions $p(\bar{y}_t | \mathbf{y}_{<t})$ by means of self-attention. Because transformer LMs necessitate the precise introduction of multiple sub-components, we postpone the full introduction to App. C.4 and only review the basics here. We borrow much of the notation and formalization from Svete and Cotterell (2024).

A transformer is a composition of multiple transformer **layers**, each of which implements the **attention mechanism**. The attention mechanism works as follows. It takes a **query** vector $\mathbf{q} \in \mathbb{R}^D$ and two matrices: The matrix $\mathbf{K} \in \mathbb{R}^{N \times D}$ of **keys** and the matrix $\mathbf{V} \in \mathbb{R}^{N \times D}$ of **values** and computes a weighted average of the value vectors based on

the compatibilities of the key vectors to the query vector, as scored by a scoring function f .

Attention weights are computed by normalizing the scores $f(\mathbf{q}, \mathbf{k}_1), \dots, f(\mathbf{q}, \mathbf{k}_t)$. The choice of the normalization function has concrete implications on representational capacity (Hao et al., 2022; Svete and Cotterell, 2024). We focus on the **hard attention** projection function.

Definition 2.12. *Hard attention* is computed with the hardmax projection function

$$\text{hardmax}(\mathbf{x})_d \stackrel{\text{def}}{=} \begin{cases} \frac{1}{m} & \text{if } d \in \text{argmax}(\mathbf{x}) \\ 0 & \text{otherwise,} \end{cases} \quad (10)$$

for $d \in [D]$, where $\mathbf{x} \in \mathbb{R}^D$ and $m \stackrel{\text{def}}{=} |\text{argmax}(\mathbf{x})|$ is the cardinality of the argmax set.

A transformer layer uses the attention mechanism followed by a position-wise MLP¹² to compute augmented representations \mathbf{z}_t of the input representations $\mathbf{X}_t = (\mathbf{x}_1^\top; \dots; \mathbf{x}_t^\top) \in \mathbb{R}^{t \times D}$. The query \mathbf{q}_t , the keys \mathbf{K}_t , and values \mathbf{V}_t are all transformations of the input representations \mathbf{X}_t . Initially, \mathbf{X}_t are computed by some static representation function of the symbols and their positions. Multiple transformer layers are stacked into a transformer, which computes the (deep) contextual representations of all symbols in the string. The contextual representation of the final symbol in the string is then used to define the representation function enc of a transformer LM.

3 CoT Reasoning and Weak Equivalence

In this section, we argue that CoT reasoning can be seen as a way of comparing the input language of a formal automaton with the output language of an autoregressive LM. We then introduce a formalization of CoT reasoning that allows us to characterize the representational capabilities of CoT-augmented LMs in terms of their weak equivalence to well-studied weighted automata.

3.1 Equivalence and Augmented Alphabets

Autoregressive LMs generate strings by outputting individual symbols $\bar{y} \in \bar{\Sigma}$ one at a time until EOS is generated. The resulting string \mathbf{y} is the output of the LM and the outputs generated in this manner implicitly define the distribution of the LM. Models studied in existing work on the representational capacity of neural LMs, however,

¹²For more details, see App. C.4

often do not (only) emit symbols from Σ . Rather, they also emit symbols from some *larger* alphabet Δ that includes symbols that encode additional information required to simulate a given formal model of computational. For example, the symbols generated by Pérez et al.’s (2021) transformer contain both the output symbols as well as (the changes to) the configuration of the Turing machine.

It is not difficult to see how generating strings from an augmented alphabet can be seen as a form of CoT reasoning. The outputs *not* intended to be a part of the final output can simply be seen as the result of the additional computational steps performed by the CoT-augmented LM. These outputs are later removed in post-processing and only the string with symbols from our target alphabet Σ remains. In this sense, Pérez et al.’s (2021) construction works with a form of CoT reasoning without explicitly mentioning it. This connection was made explicit in concurrent work by Merrill and Sabharwal (2024).

Outputting additional information is not regarded as an issue when the task is to simulate (unweighted) Turing machine runs on a given *input string* because what matters, in that case, is only whether the model accepts or rejects the input (Siegelmann and Sontag, 1992; Pérez et al., 2021, *inter alia*). However, when considering the LM induced by a PTM, we do care about the alphabet the distribution over strings is over. Thus, given an LM over Σ^* , it follows that neural LMs outputting additional information cannot define the same distribution, i.e., they cannot be weakly equivalent. Nevertheless, the ability to output additional information while generating a string seems natural and useful; it is the heart of CoT reasoning. And, indeed, models that can only output symbols that are part of the final string are restricted to doing real-time computation (Weiss et al., 2018; Nowak et al., 2023; Svete et al., 2024).

3.2 Regular Reducibility

We now formalize the notion of CoT reasoning through the following definition which allows a model to output additional information *not* considered part of the final output:


Definition 3.1. *An LM p over Δ is **regularly reducible** to a LM q over Σ if there exists a regular function $\phi: \Delta^* \rightarrow \Sigma^*$ such that $q \circ \phi$ is weakly equivalent to p .*

We can use the function ϕ to map the strings

sampled from an LM, ones additionally encoding the intermediary steps of computation, into the final output, i.e., strings $y \in \Sigma^*$.

Definition 3.2. *We say an alphabet Δ satisfies the **Σ -augmentation condition** for an alphabet Σ if $\Delta \subseteq \bar{\Sigma}_\varepsilon \times \Gamma_\varepsilon \setminus \{(\varepsilon, \varepsilon)\}$ for some alphabet Γ .*

This means if we care about strings from Σ^* and we have an LM over the closure of a Σ -augmented alphabet Δ , then each symbol the LM outputs encodes is either an output symbol from Σ , an intermediate computation symbol from Γ , or both.¹³ The computation symbols can then easily be removed by applying the per-symbol projection function $\phi((y, \gamma)) \stackrel{\text{def}}{=} y$, lifted to a non-length-increasing homomorphism over strings $\phi: \Delta^* \rightarrow \Sigma^*$ to the LM’s output element-wise.¹⁴ This function can be implemented by a simple single-state transducer (making it a regular function):



3.3 Properties of Regular Reducibility

To motivate our formalization, in this section, we show several properties of regular reducibility that will allow us to reason about the representational capacity of CoT-augmented LMs later in §3.4. Particularly, as we will see, there exists a strong connection between regular reducibility and the addition of *non-deterministic* choices in the model. To exemplify this concretely, the next theorem shows that regular reducibility allows us to model non-deterministic PFSA with deterministic ones, which, as discussed in §2.2, cannot be done with just deterministic PFSA in general.

Theorem 3.1. *Let $\mathcal{A} = (\Sigma, Q, \delta, \lambda, \rho)$ be a PFSA. Then, there exists a deterministic PFSA \mathcal{A}' over the alphabet $\Sigma \times Q$ and with the state space $\Sigma \times Q$ that is regularly reducible to \mathcal{A} .*

Proof. See App. E. ■

A similar connection exists for PPDAs (see App. E). Thm. 3.1 captures a crucial property of regular reducibility: It allows us to *simulate* non-determinism with a deterministic device. As

¹³While in practice the CoT and the final output come from the same alphabet, we use separate ones for ease of notation and analysis. This is without loss of generality; see App. D.

¹⁴A non-length-increasing homomorphism is a function $h: \Delta^* \rightarrow \Sigma^*$ where $h(\varepsilon) = \varepsilon$, $h(x) = y$ for $x \in \Delta$ and $y \in \Sigma_\varepsilon$, and that further satisfies $h(x_1x_2 \cdots x_T) = h(x_1)h(x_2) \cdots h(x_T)$.

An LM p generates strings from $\Delta^* \subseteq (\Sigma \times Q)^*$:

$$\mathbf{x}_1 = q_0, y_1, q_1, q_2, \dots, q_N, y_{T-1}, y_T \sim p$$

$$\mathbf{x}_2 = q'_0, q'_1, y_1, q_2, \dots, y_T, q'_{M-1}, q'_M \sim p$$

The function ϕ removes the CoT steps:

$$\phi(\mathbf{x}_1) = \phi(\mathbf{x}_2) = y_1, y_2, \dots, y_{T-1}, y_T \sim p^{\text{CoT}}$$

Figure 2: An LM p can generate strings from a state-augmented alphabet Δ . The intermediate outputs q_i allow it to condition on the previous states of the computation. By *post-hoc* removing the intermediate outputs with ϕ , we obtain a CoT-augmented LM p^{CoT} over Σ^* that generates more human-readable outputs.

exemplified in the cases of regular LMs, this can concretely increase the representational capacity of a model class, as in the case of Thm. 3.1. On the other hand, regular reducibility never *decreases* the representational capacity of a model class, since ϕ can always be set to the identity function. Due to the close connection between neural LMs and determinism (Svete and Cotterell, 2023), one could hope that a similar increase in capabilities could be achieved for neural LMs augmented with CoT reasoning as well. In §3, we show that this is indeed the case.

3.4 Regular Reducibility and CoT Reasoning

After introducing the notion of regular reducibility and presenting some of its core properties, we now use it to define CoT-augmented LMs.

Definition 3.3. Given alphabets Σ and Δ where Δ satisfies the Σ -augmentation condition, an LM p^{CoT} over Σ^* is called a **CoT-augmented LM** induced by LM p over Δ^* if p is regularly reducible from p^{CoT} . That is, there exists a regular function $\phi: \Delta^* \rightarrow \Sigma^*$ such that $p^{\text{CoT}} \stackrel{\text{def}}{=} p \circ \phi^{-1}$.

Def. 3.3 defines a CoT-augmented LM p^{CoT} through the LM p that generates strings from Δ^* and then applies the regular function ϕ to the generated strings to obtain strings from Σ^* . This allows p to output additional information while still defining a probability distribution over strings from Σ^* (see Fig. 2). The weight of a string $\mathbf{y} \in \Sigma^*$ under p^{CoT} is then computed as the sum over the weights of all strings in the preimage of ϕ , analogously to Eq. (6):

$$p^{\text{CoT}}(\mathbf{y}) = \sum_{\mathbf{x} \in \phi^{-1}(\mathbf{y})} p(\mathbf{x}). \quad (11)$$

This is the approach we take in §4.2.2.




Figure 3: If $\mathbf{y}_{\leq t}$ uniquely determines a subpath in the PFSA, the current state q_t can be uniquely determined *without* knowing the previous state (top). If, however, the substring $\mathbf{y}_{\leq t}$ can lead to multiple states, the relevant next-symbol distribution cannot be determined. Storing the *sampled* states as part of the output fixes this by keeping track of the sampled PFSA subpath (bottom).

4 CoT and Representational Capacity

We now connect the notions introduced in §3 to the representational capacity of neural LMs. We begin in §4.1 by showing how CoT reasoning endows autoregressive LMs with *non-determinism* by using scratch space to keep track of the current branch of execution. We then use similar principles to emulate probabilistic Turing machines with Elman RNNs (§4.2.1) and transformers (§4.2.2).

Neural LMs and formal models of computation.

Probabilistic models of computation are often presented as autoregressive *generators* of strings that implicitly define probability (semi)measures (Icard, 2020). This interpretation is particularly important when addressing non-determinism, because it allows for multiple executions (for example, in the case of PFSA, multiple different *paths*) generating the same string. Here, we treat neural LMs as autoregressive generators of strings and show that, by defining identical next-symbol distributions, they implicitly define the same semimeasures over strings as classical probabilistic models of computation.

4.1 Neural LMs and Regular LMs

We first discuss the connection between CoT-augmented neural LMs and PFSA.

4.1.1 Recurrent Neural LMs

Minsky’s (1954) construction provided one of the first connections between a neural network and a formal computational model. It showed that RNNs can emulate (deterministic) FSAs, where the RNN accepts string by activating a particular

neuron after reading the string. The relationship in the probabilistic case was explored by Svete and Cotterell (2023), who show the equivalence of pure Elman RNNs (without CoT reasoning) and *deterministic* PFSA. ¹⁵ This illustrates an important distinction between the deterministic and non-deterministic frameworks and thus a discrepancy between general PFSA and RNN LMs. Intuitively, the discrepancy comes from the fact that there is no non-determinism in the recurrence of an RNN, which is thus unable to capture the possibly non-deterministic decisions of the PFSA. CoT reasoning endows an RNN with exactly this non-determinism because it allows the RNN to sample and refer back to *trajectories* rather than only symbols by annotating each of its outputs with the current state of the PFSA. ¹⁶ Upon reading the previously randomly generated state of the automaton captured in the output, it can follow the randomly sampled generating trajectories.

The following two theorems show that RNN LMs with fixed precision and CoT reasoning are weakly equivalent to general PFSA. This requires showing the correspondence in both directions, i.e., that (1) the distribution induced by any PFSA can be generated by a constant-precision RNN LM with CoT, and (2) any CoT-augmented constant-precision RNN LM can be emulated by a PFSA.

Theorem 4.1. *For any regular LM, there exists a weakly equivalent CoT-augmented constant-precision Elman RNN LM.*

Proof intuition. We show how, using the RNN’s recurrence and output sampling step, we can implement the transition function of any PFSA. The RNN starts by sampling an output symbol containing an initial state q_0 according to the initial distribution λ without emitting a language symbol. This output gets fed back into the RNN at the next time step, allowing it to *read* the sampled state, and the next symbol–state pair is sampled according to the conditional distribution defined by the output state, as illustrated by Fig. 3. Finally, the states in the generated string are removed by a regular function, leaving only the language output. See App. F for the detailed proof. ■

¹⁵The relationship between RNN LMs and *non-deterministic* PFSA was explored by Svete et al. (2024), albeit with linearly bounded-precision RNNs.

¹⁶Recall that non-determinism means that multiple possible transitions to different states can yield the same symbol (§2.2).

Theorem 4.2. *For any constant-precision CoT-augmented Elman RNN LM, there exists a weakly equivalent PFSA.*

Proof. See App. F. ■

Thms. 4.1 and 4.2 establish the equivalence between CoT-augmented *constant-precision* Elman RNN LMs and general regular LMs. This illuminates the added representational capacity awarded by CoT reasoning: Storing the current FSA state in the output string and removing it later allows the model to handle non-determinism which is not possible without the additional information.

4.1.2 Transformer LMs

We now show an analogous claim to Thm. 4.1 for Transformer LMs with CoT. This is simply a stepping stone towards full probabilistic Turing completeness—a similar construction will then lead us to the full proof that transformer LMs with unbounded precision can simulate PTMs.

Theorem 4.3. *For any regular LM, there exists a weakly equivalent CoT-augmented constant-precision transformer LM.*

Proof sketch. We use the construction from Svete and Cotterell (2024) which shows how to encode n -gram LMs in a transformer. Here, the alphabet of the transformer is augmented with the states of the PFSA, i.e., the symbol at each position t contains not only an output symbol but also the current state of the PFSA at time t . Thereby each input symbol contains all the information required to compute the next-symbol probabilities (it is effectively a unigram LM). Because the next symbol probabilities in a PFSA only depend on the current state, the transformer does not need to use attention at all and can rely solely on the output of its residual connections. See App. F for the detailed proof. ■

4.2 Neural LMs and PTMs

We now extend the results from the previous section from representing the simple regular LM to expressing all enumerable semimeasures over strings by emulating probabilistic Turing machines. For this, in the RNN case, we require unbounded precision and ReLU activation functions rather than fixed precision and Heaviside activations.

4.2.1 Recurrent Neural LMs

First, note the following definition:

Definition 4.1. An autoregressive LM is called a *real-time LM* if it never outputs empty symbols (ϵ).

Nowak et al. (2023) show that RNN LMs with rationally valued activations that are not restricted to operating in real time are Turing complete and weakly equivalent to a subset of rationally weighted PTMs. Intuitively, not operating in real time gives an LM additional computation time while storing the results of its computations in the hidden states. This is a form of CoT reasoning since we could equivalently let the LM output additional symbols that do not count toward the final output.¹⁷ The ability to *erase* certain symbols, as done by our transducer ϕ , helps make the setup of Nowak et al. (2023) realistic. Importantly, beyond giving the LM more computation time, CoT also allows the LM to model non-determinism in PTMs.

Theorem 4.4. For every LM induced by a non-deterministic probabilistic Turing machine, there exists a weakly equivalent CoT-augmented RNN LM without unbounded precision.

Proof intuition. The proof follows Nowak et al.’s (2023) probabilistic version of the proof by Siegelmann and Sontag (1992), but extended to account for CoT reasoning. For the ability to simulate *all* PTMs rather than just a subset, we augment the output alphabet of the RNN with enough information about the current PTM configuration so that subsequent steps can uniquely identify the previous action taken. Concretely, the output alphabet Δ contains information about the current state, the symbol written on the working tape of the PTM, and the head action performed. This additional information is then removed by our regular function at the end, yielding only the output of the simulated PTM generated according to its probability. A detailed proof is presented in App. F. ■

4.2.2 Transformer LMs

The representational capacity of transformer LMs has received a lot of attention in the last few years. Pérez et al. (2021) established that the encoder–decoder variant of the architecture is Turing complete. Since then, concurrent work has shown non-probabilistic Turing completeness of LM-oriented decoder-only variants (Merrill and Sabharwal, 2024; Feng et al., 2023). We extend the work to the *probabilistic* case.

¹⁷On the other hand, some CoT-augmented LMs do operate in real time but over an extended alphabet; see §4.1.

Theorem 4.5. For any PTM-induced LM, there exists a weakly equivalent unbounded-precision CoT-augmented Transformer LMs.

Proof intuition. The proof follows Pérez et al. (2021) but is adapted to the probabilistic case. Again, the main idea is to augment the output alphabet with enough information about the current PTM configuration to reconstruct the probabilities of possible actions at each time step. This information and appropriate positional encodings are enough to recover the PTM’s current configuration and thus the next-action distribution, allowing us to construct a weakly equivalent transformer LM. A regular function then removes the additional information from the string. See App. F for details. ■

4.2.3 Weak Equivalence

Thms. 4.4 and 4.5 show that transformer and RNN LMs with CoT reasoning are at least as powerful as PTMs. Weak equivalence requires us to also prove the reverse of these two theorems, analogous to Thm. 4.2 for constant-precision Elman RNN LMs.

Theorem 4.6. For any rationally valued RNN LM and transformer LM with CoT reasoning, there exists a weakly equivalent PTM.

Proof. RNN LMs define enumerable semimeasures (Nowak et al., 2023, App. G), which can always be expressed by a PTM (Icard, 2020, Thm. 3). Following the same reasoning, transformer LMs define enumerable semimeasures as well; the probability of a string y is defined as the sum probabilities of all runs of the transformer which result in the output y , of which there are countably many. Finally, there is only a countable number of ways a regular function can transform any string, so RNN LMs and Transformer LMs with CoT still define enumerable semimeasures. ■

5 Conclusion

Many modern language models have been shown to perform better with CoT reasoning. Despite its empirical success, CoT reasoning has yet to be well understood formally. Recent theoretical work in this area has analyzed the Turing completeness of CoT-augmented LMs, failing to account for the inherently probabilistic nature of LMs. We fix this mismatch in the current literature by introducing a novel formalization of CoT reasoning.

Limitations

Our constructions simulating PFSAs rely on fixed-precision arithmetic in line with the finite-memory nature of such automata. In contrast, the Turing-complete models all rely on *unbounded* precision with respect to the length of the string—either to be able to encode arbitrarily large stacks in the hidden state in the case of RNN LMs or to be able to encode positional information in the case of transformer LMs. This inevitably results from the unbounded number of computational steps per emitted symbol by a Turing machine and is distinctly different from the scaling with respect to the number of *computational steps*. In the transformer case, the precision scales logarithmically with the number of steps (since we use the same positional encodings as Pérez et al. (2021)), which is standard for theoretical investigations of transformer models (Yao et al., 2021; Merrill et al., 2022; Merrill and Sabharwal, 2023). In the case of RNNs, the precision scales linearly with the computation sequence length, due to the encoding of a stack as a rational number, where each entry on the stack occupies a single digit. While in line with previous theoretical work, these assumptions are unrealistic in practice since neural LMs normally use fixed precision floating point arithmetic.

We do not use layer normalization with transformers—this is done for simplicity—but note that layer normalization has been found to increase the representational capacity of transformers in some cases (Chiang and Cholak, 2022; Merrill and Sabharwal, 2024).

Moreover, we only prove theoretically the equivalence between CoT-augmented LMs and formal models of computation. That is, we do not give a training algorithm to elicit the emergence of specific Turing complete or non-deterministic regular automata in neural LMs, and we make no claims about the training efficiency or even the feasibility of training neural LMs for this purpose.

Ethics Statement

Since this work deals with purely theoretical properties of neural language models, we do not foresee any ethical issues arising from this work.

References

Steven Abney, David McAllester, and Fernando Pereira. 1999. *Relating probabilistic grammars and automata*.

In Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics, pages 542–549, College Park, Maryland, USA. Association for Computational Linguistics.

Y. Bar-Hillel, M. Perles, and E. Shamir. 1961. *On formal properties of simple phrase structure grammars*. *Zeitschrift für Phonetik, Sprachwissenschaft und Kommunikationsforschung*, 14:143–172.

Bruno Bauwens. 2013. *Upper semicomputable sumtests for lower semicomputable semimeasures*. *arXiv preprint arXiv:1312.1718*.

Satwik Bhattamishra, Arkil Patel, and Navin Goyal. 2020. *On the computational power of transformers and its implications in sequence modeling*. *In Proceedings of the 24th Conference on Computational Natural Language Learning*, pages 455–475, Online. Association for Computational Linguistics.

Adam L. Buchsbaum, Raffaele Giancarlo, and Jeffrey R. Westbrook. 2000. *On the determinization of weighted finite automata*. *SIAM Journal on Computing*, 30(5):1502–1531.

Robin Chan, Reda Boumasmoud, Anej Svete, Yuxin Ren, Qipeng Guo, Zhijing Jin, Shauli Ravfogel, Mrinmaya Sachan, Bernhard Schölkopf, Mennatallah El-Assady, and Ryan Cotterell. 2024. *On affine homotopy between language encoders*. *arXiv preprint arXiv:2406.02329*.

David Chiang and Peter Cholak. 2022. *Overcoming a theoretical limitation of self-attention*. *In Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7654–7664, Dublin, Ireland. Association for Computational Linguistics.

Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. *On the properties of neural machine translation: Encoder–decoder approaches*. *In Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, pages 103–111, Doha, Qatar. Association for Computational Linguistics.

Stephen Chung and Hava Siegelmann. 2021. *Turing completeness of bounded-precision recurrent neural networks*. *In Advances in Neural Information Processing Systems*, volume 34, pages 28431–28441. Curran Associates, Inc.

Li Du, Lucas Torroba Hennigen, Tiago Pimentel, Clara Meister, Jason Eisner, and Ryan Cotterell. 2023. *A measure-theoretic characterization of tight language models*. *In Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9744–9770, Toronto, Canada. Association for Computational Linguistics.

Jeffrey L. Elman. 1990. *Finding structure in time*. *Cognitive Science*, 14(2):179–211.

- Guhao Feng, Bohang Zhang, Yuntian Gu, Haotian Ye, Di He, and Liwei Wang. 2023. [Towards revealing the mystery behind chain of thought: A theoretical perspective](#). *arXiv preprint arXiv:2305.15408*.
- Michael Hahn. 2020. [Theoretical limitations of self-attention in neural sequence models](#). *Transactions of the Association for Computational Linguistics*, 8:156–171.
- Yiding Hao, Dana Angluin, and Robert Frank. 2022. [Formal language recognition by hard attention transformers: Perspectives from circuit complexity](#). *Transactions of the Association for Computational Linguistics*, 10:800–810.
- Yiding Hao, William Merrill, Dana Angluin, Robert Frank, Noah Amsel, Andrew Benz, and Simon Mendelsohn. 2018. [Context-free transductions with neural stacks](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 306–315, Brussels, Belgium. Association for Computational Linguistics.
- John Hewitt, Michael Hahn, Surya Ganguli, Percy Liang, and Christopher D. Manning. 2020. [RNNs can generate bounded hierarchical languages with optimal memory](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1978–2010, Online. Association for Computational Linguistics.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural Computation*, 9(8):1735–1780.
- John E. Hopcroft, Rajeev Motwani, and Jeffrey D. Ullman. 2001. *Introduction to Automata Theory, Languages, and Computation*, 3 edition. Pearson.
- Thomas F. Icard. 2020. [Calibrating generative models: The probabilistic Chomsky–Schützenberger hierarchy](#). *Journal of Mathematical Psychology*, 95:102308.
- S. C. Kleene. 1956. [Representation of events in nerve nets and finite automata](#). In C. E. Shannon and J. McCarthy, editors, *Automata Studies. (AM-34), Volume 34*, pages 3–42. Princeton University Press, Princeton.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2023. [Large language models are zero-shot reasoners](#). *arXiv preprint arXiv:2205.11916*.
- Samuel A. Korsky and Robert C. Berwick. 2019. [On the computational power of RNNs](#). *arXiv preprint arXiv:1906.06349*.
- Ming Li and Paul M.B. Vitányi. 2008. *An Introduction to Kolmogorov Complexity and Its Applications*, 3 edition. Springer Publishing Company, Incorporated.
- André F. T. Martins and Ramón F. Astudillo. 2016. [From softmax to sparsemax: A sparse model of attention and multi-label classification](#). In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48, ICML’16*, page 1614–1623.
- Warren S. McCulloch and Walter Pitts. 1943. [A logical calculus of the ideas immanent in nervous activity](#). *The bulletin of mathematical biophysics*, 5(4):115–133.
- 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.
- William Merrill and Ashish Sabharwal. 2023. [The parallelism tradeoff: Limitations of log-precision transformers](#). *Transactions of the Association for Computational Linguistics*, 11:531–545.
- William Merrill and Ashish Sabharwal. 2024. [The expressive power of transformers with chain of thought](#). In *The Twelfth International Conference on Learning Representations*.
- William Merrill, Ashish Sabharwal, and Noah A. Smith. 2022. [Saturated transformers are constant-depth threshold circuits](#). *Transactions of the Association for Computational Linguistics*, 10:843–856.
- William Merrill and Nikolaos Tsilivis. 2022. [Extracting finite automata from RNNs using state merging](#). *arXiv preprint arXiv:2201.12451*.
- William Merrill, Gail Weiss, Yoav Goldberg, Roy Schwartz, Noah A. Smith, and Eran Yahav. 2020. [A formal hierarchy of RNN architectures](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 443–459, Online. Association for Computational Linguistics.
- Marvin Lee Minsky. 1954. *Neural Nets and the Brain Model Problem*. Ph.D. thesis, Princeton University.
- 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, Fernando Pereira, and Michael Riley. 2008. *Speech Recognition with Weighted Finite-State Transducers*, pages 559–584. Springer Berlin Heidelberg, Berlin, Heidelberg.
- 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.

- Maxwell Nye, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin, David Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, Charles Sutton, and Augustus Odena. 2021. [Show your work: Scratchpads for intermediate computation with language models](#). *arXiv preprint arXiv:2112.00114*.
- C.H. Papadimitriou. 1994. *Computational Complexity*. Theoretical computer science. Addison-Wesley.
- Clemente Pasti, Andreas Opedal, Tiago Pimentel, Tim Vieira, Jason Eisner, and Ryan Cotterell. 2023. [On the intersection of context-free and regular languages](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 737–749, Dubrovnik, Croatia. Association for Computational Linguistics.
- Fernando C. N. Pereira and Michael D. Riley. 1997. [Speech recognition by composition of weighted finite automata](#). In *Finite-State Language Processing*. The MIT Press.
- Jorge Pérez, Pablo Barceló, and Javier Marinkovic. 2021. [Attention is turing-complete](#). *Journal of Machine Learning Research*, 22(75):1–35.
- Hava T. Siegelmann and Eduardo D. Sontag. 1992. [On the computational power of neural nets](#). In *Proceedings of the Fifth Annual Workshop on Computational Learning Theory, COLT '92*, page 440–449, New York, NY, USA. Association for Computing Machinery.
- Michael Sipser. 2013. *Introduction to the Theory of Computation*, 3 edition. Cengage Learning.
- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V. Le, Ed H. Chi, Denny Zhou, and Jason Wei. 2022. [Challenging BIG-Bench tasks and whether chain-of-thought can solve them](#). *arXiv preprint arXiv:2210.09261*.
- Anej Svete and Ryan Cotterell. 2023. [Recurrent neural language models as probabilistic finite-state automata](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 8069–8086, Singapore. Association for Computational Linguistics.
- Anej Svete and Ryan Cotterell. 2024. [Transformers can represent \$n\$ -gram language models](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics*, Mexico City. Association for Computational Linguistics.
- Anej Svete, Franz Nowak, Anisha Mohamed Sahabdeen, and Ryan Cotterell. 2024. [Lower bounds on the expressivity of recurrent neural language models](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics*, Mexico City. Association for Computational Linguistics.
- Johannes Von Oswald, Eyvind Niklasson, Ettore Randazzo, João Sacramento, Alexander Mordvintsev, Andrey Zhmoginov, and Max Vladymyrov. 2023. [Transformers learn in-context by gradient descent](#). In *Proceedings of the 40th International Conference on Machine Learning, ICML'23*. JMLR.org.
- Xuezhi Wang and Denny Zhou. 2024. [Chain-of-thought reasoning without prompting](#). *arXiv preprint arXiv:2402.10200*.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022a. [Emergent abilities of large language models](#). *arXiv preprint arXiv:2206.07682*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. 2022b. [Chain-of-thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837. Curran Associates, Inc.
- Gail Weiss, Yoav Goldberg, and Eran Yahav. 2018. [On the practical computational power of finite precision RNNs for language recognition](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 740–745, Melbourne, Australia. Association for Computational Linguistics.
- Noam Wies, Yoav Levine, and Amnon Shashua. 2023. [Sub-task decomposition enables learning in sequence to sequence tasks](#). *arXiv preprint arXiv:2204.02892*.
- Shunyu Yao, Binghui Peng, Christos Papadimitriou, and Karthik Narasimhan. 2021. [Self-attention networks can process bounded hierarchical languages](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3770–3785, Online. Association for Computational Linguistics.

A Discussion

At a high level, our claims characterize neural architectures endowed with CoT reasoning in terms of well-understood *probabilistic* models of computation, giving the framework of CoT reasoning a novel probabilistic perspective. This allows the reconciliation of theoretical results about neural LMs with the general probabilistic language modeling framework. In doing so, we put existing results about the representational capacity of LMs into perspective and show how they can be thought of as performing “CoT reasoning in disguise.” Perhaps surprisingly, the inclusion of CoT reasoning steps results in a natural inclusion of non-determinism in otherwise deterministic neural LMs, as exemplified by Thm. 4.1. Outputting the states as part of the generated string—whose generation is inherently non-deterministic—allows a neural LM to keep track of the trajectory of the current execution. This suggests that CoT reasoning might provide an interesting avenue for exploring the non-deterministic representational capacity of neural LMs. Note that this means CoT reasoning can endow neural LMs with higher expressivity since, e.g., non-deterministic PFSA are strictly more expressive than deterministic ones. Similarly, our result that RNNs or transformer LMs with CoT reasoning can simulate PTMs rather than regular Turing machines is important as it allows *probabilistic* computation using such neural LMs. This means one can sample from them multiple times to assess the certainty of string inclusion in a language. Furthermore, the problems LMs could solve efficiently can be described by different complexity classes such as BPP, ZPP, etc. instead of P.¹⁸

B Related Work

Turing completeness of RNNs. Plenty of existing work has investigated the representational capacity of RNNs, both as recognizers as well as LMs (e.g., McCulloch and Pitts, 1943; Kleene, 1956; Siegelmann and Sontag, 1992; Hao et al., 2018; Korsky and Berwick, 2019; Merrill, 2019; Merrill et al., 2020; Hewitt et al., 2020; Chung and Siegelmann, 2021; Merrill et al., 2022; Merrill and Tsilivis, 2022; Svete and Cotterell, 2023; Nowak et al., 2023, *inter alia*). Most relevant to our work, Siegelmann and Sontag (1992); Chung and Siegelmann (2021) show how RNNs with unbounded precision and unbounded computation time can simulate Turing machines and discuss the implications. Nowak et al. (2023) extend this to the probabilistic setting, showing that RNN LMs can emulate certain PTMs, but require the LMs to be able to perform non-emitting steps.

Turing completeness of transformers. Pérez et al. (2021) show the Turing completeness of hard attention encoder–decoder transformer by encoding the configuration of the Turing machine in the output of the transformer. Bhattamishra et al. (2020) provide a different perspective on Turing completeness of the architecture by showing that transformers can simulate RNNs. Turing completeness is in this sense a simple consequence of the Turing completeness of RNNs. Bhattamishra et al.’s (2020) construction, however, relies on emitting symbols from an (uncountably) infinite set (rational-valued vectors), in contradiction to the requirement of the LM working over a *finite* alphabet. In this sense, they make use of the so-called *regression* transformer setup, which does not lend itself well to the discrete language modeling (Von Oswald et al., 2023). In concurrent work, Merrill and Sabharwal (2024) adapt Pérez et al.’s (2021) result to the CoT setting, connecting Turing completeness of transformer decoders to CoT reasoning similar to our work. In contrast to our work, they make statements about the model’s ability to decide language membership (simulating a Turing machine that accepts or rejects an input) and not to represent probabilistic languages. In that sense, the transformer they construct is not a language model. Besides being probabilistic, our construction also enables the analysis of any autoregressive LM architecture, and we focus on RNN and transformer LMs. Feng et al. (2023) similarly show that CoT transformers can perform arithmetic expressions and dynamic programming by generating intermediate results, backing these claims up with experimental evidence, but stopping short of showing Turing completeness. Du et al. (2023) provide a related result, showing that transformer LMs with continuous transformation functions are *tight*. This might at first glance contradict the results on Turing completeness, since the latter might require the model to run indefinitely, resulting in a non-tight model. However,

¹⁸For a comparison of these complexity classes, see e.g. Papadimitriou (1994).

the discrepancy is resolved by noting that the tightness result crucially relies on the use of the softmax normalization function over the reals \mathbb{R} (without $\pm\infty$) in Eq. (2). This, together with the encoding function of the transformer, results in a non-diminishing probability of generating EOS and thus in a tight model. Our use of the sparsemax function sidesteps this issue by allowing us to assign EOS probability 0.

C Additional Preliminaries

In our construction, we assume that strings are prefixed with the `beginning-of-string` symbol BOS. This is to give a base case for recurrent definitions such as Def. 2.1. This is just for ease of notation in our proofs. Note that instead of $p(\bar{y} \mid \text{BOS})$, we could also equivalently write $p(\bar{y} \mid \varepsilon)$.

C.1 Turing Machines

Under our definition, a probabilistic Turing machine has two tapes.¹⁹ The first is the **processing tape** on which symbols from the tape alphabet can be read and written. The second is a write-only **output tape** for symbols of the output alphabet. In the beginning, the processing tape contains the designated \perp symbol in the leftmost cell while all other cells contain the blank symbol \sqcup . The output tape is empty at the beginning. Starting in the initial state q_ι , at each time step t , a transition is sampled out of all available transitions for the given state q and the current working tape symbol γ , and then applied. The sampling happens according to the transition probability w (recall that the transition weights w are non-negative and sum to 1 for each pair of state q tape symbol γ). We write transitions as $(q, \gamma) \xrightarrow{y, d/w} (q', \gamma')$, where $q, q' \in Q, \gamma, \gamma' \in \Gamma, y \in \Sigma_\varepsilon, w \in \mathbb{Q}_{\geq 0}$, and $d \in \{L, R\}$, with the following interpretation. When the machine is in state q and its head is reading γ on the working tape, it moves to state q' , writes γ' to the working tape, writes y to the output tape if $y \in \Sigma$ or nothing if $y = \varepsilon$, and it moves the head on the working tape by one symbol in the direction d , i.e., to the left (L), or to the right (R). As with any non-deterministic machine, the above definition naturally gives rise to the notion of a tree of possible computations (Sipser, 2013, p. 48). A **branch** π of the computation tree of a PTM is a sequence of consecutive transitions

$$(q_\iota, \perp) \xrightarrow{y_1, d_1/w_1} (q_2, \gamma'_1), (q_2, \gamma_2) \xrightarrow{y_2, d_2/w_2} (q_3, \gamma'_2), \dots (q_N, \gamma_N) \xrightarrow{y_N, d_N/w_N} (q_{N+1}, \gamma'_N) \quad (12)$$

We say that a branch is **accepting** if the branch reaches the final state q_φ .²⁰ The **yield** of an accepting computation branch²¹ is the sequence of symbols $\mathbf{y} \in \Sigma^*$ written on the output tape at that point in the computation, i.e., a concatenation of the (non- ε) symbols y_1, \dots, y_N , where $N = |\pi|$ is the number of transitions in the branch. The **weight** of an accepting branch is the product of the weights of its transitions, i.e.,

$$w(\pi) \stackrel{\text{def}}{=} \prod_{n=0}^N w_n. \quad (13)$$

We denote the branches that yield a given string \mathbf{y} by $\Pi(\mathcal{M}, \mathbf{y})$. The sum of weights of all branches that yield a certain string $\mathbf{y} \in \Sigma^*$ is the **stringsum** of that string, defined as

$$\mathcal{M}(\mathbf{y}) \stackrel{\text{def}}{=} \sum_{\pi \in \Pi(\mathcal{M}, \mathbf{y})} w(\pi). \quad (14)$$

This definition gives rise to a semimeasure over strings whose sum over all possible strings is exactly the halting probability²² of the PTM, i.e., the probability that starting from the initial state q_ι , \mathcal{M} reaches a final state q_φ

$$\sum_{\mathbf{y} \in \Sigma^*} \mathcal{M}(\mathbf{y}) = \mathbb{P}(\mathcal{M} \text{ halts}). \quad (15)$$

¹⁹Adding another tape does not increase the computational power of a Turing machine (Sipser, 2013, Ch. 3).

²⁰The final state has no outgoing transitions.

²¹We only consider branches that end in a final state when discussing the yield and weight of branches. This means we only take into account finite instances of computation.

²²The halting probability is the probability that the execution of the PTM will end in a halting state after finitely many steps.

C.2 Pushdown Automata

Definition C.1. A *probabilistic pushdown automaton* (PPDA) is a tuple $(Q, \Sigma, \Gamma, \delta, (q_\iota, S), (q_\varphi, \varepsilon))$ where Q is a finite set of states, Σ is the input alphabet, Γ is the stack alphabet, $\delta \subseteq Q \times \Gamma \times \Sigma_\varepsilon \times \mathbb{Q}_{\geq 0} \times Q \times \Gamma^*$ is a finite set of weighted transitions, and (q_ι, S) and (q_φ, ε) are called the initial and final configuration, respectively. Moreover, for all states $q \in Q$ and stack symbols $\gamma \in \Gamma$, δ satisfies

$$\sum_{(q, \gamma) \xrightarrow{y/w} (q', \gamma) \in \delta} w = 1.$$

Our definition of PPDA is similar to that of [Abney et al. \(1999\)](#), but, unlike theirs, our machine must end its computation in the final configuration rather than just empty the stack. We write transitions $(q, \gamma, y, w, q', \gamma) \in \delta$ as $(q, \gamma) \xrightarrow{y/w} (q', \gamma)$ which represent a move with weight w from state q to q' while scanning or outputting y , popping the symbol γ from the stack and pushing the sequence of symbols γ .

Definition C.2. A *configuration* $(q, \gamma) \in Q \times \Gamma^*$ is a pair containing the current state and the current contents of the stack.

A PPDA $\mathcal{P} = (Q, \Sigma, \Gamma, \delta, (q_\iota, S), (q_\varphi, \varepsilon))$ is called **deterministic** if for every $q \in Q$, $\gamma \in \Gamma$ and $y \in \Sigma_\varepsilon$, there is at most one transition $(q, \gamma) \xrightarrow{y/w} q', \gamma$ with $w > 0$. Additionally, if there is a transition $q, \gamma \xrightarrow{y/w} q', \gamma$ such that $y \in \Sigma$, then there is no transition $q, \gamma \xrightarrow{\varepsilon/w'} q'', \gamma'$ with $w' > 0$. Intuitively, there exists at most one next move given any configuration in a deterministic PPDA. A **run** π of a PPDA is a sequence of configurations and transitions,

$$(q_0, \gamma_0), \tau_1, (q_1, \gamma_1), \dots, \tau_N, (q_N, \gamma_N), \quad (16)$$

where (q_n, γ_n) is the configuration reached by taking transition τ_n from configuration (q_{n-1}, γ_{n-1}) for any $n \in [1, N]$. If $(q_0, \gamma_0) = (q_\iota, S)$ and $(q_N, \gamma_N) = (q_\varphi, \varepsilon)$, we call π **accepting**. The **yield** of an accepting run is $\mathbf{s}(\pi) \stackrel{\text{def}}{=} y_1 \cdots y_N$, where y_n is the symbol scanned by τ_n . The **weight** of an accepting run, $\mathbf{w}(\pi)$, is the product of the weight of its transitions,

$$\mathbf{w}(\pi) \stackrel{\text{def}}{=} \prod_{n=1}^N w_n, \quad (17)$$

where w_n is the weight of transition τ_n . We write $\Pi(\mathcal{P}, \mathbf{y})$ for the set of all accepting runs with yield \mathbf{y} . The sum of the weights of all accepting runs of some PPDA \mathcal{P} that yield a certain string $\mathbf{y} \in \Sigma^*$ is called the **stringsum** and is defined as

$$\mathcal{P}(\mathbf{y}) \stackrel{\text{def}}{=} \sum_{\pi \in \Pi(\mathcal{P}, \mathbf{y})} \mathbf{w}(\pi). \quad (18)$$

C.3 Two-stack Pushdown Automata

Definition C.3. A *two-stack pushdown automaton* (2PDA) is a machine specified by the 6-tuple $\mathcal{P} = (Q, \Sigma, \Gamma, \delta, q_\iota, q_\varphi)$, where Q is a finite set of states, Σ is an alphabet of input symbols, Γ is an alphabet of stack symbols, including the bottom-of-stack symbol \perp , $\delta: Q \times \Gamma \times \Sigma_\varepsilon \times \mathbb{Q}_{\geq 0} \times Q \times \Gamma_\varepsilon^4$ is a finite set of rationally-weighted transitions, and q_ι and q_φ are the initial and the final state, respectively.

We use a definition similar to that of [Nowak et al.'s \(2023\)](#), which assumes without loss of generality that transitions are determined by the current state and the top symbol of the first stack only. See [Nowak et al. \(2023, Appendix B\)](#) for a proof. We write transitions as $q \xrightarrow[\gamma_2 \rightarrow \gamma_4]{\gamma, y, \gamma_1 \rightarrow \gamma_3 / w} q'$, which represent a move with weight w from state q , with the top of the first stack γ , to state q' , while scanning or outputting y , popping γ_1 and γ_2 from the first and second stack, respectively, and pushing γ_3 and γ_4 to the first and second stack, respectively. A 2PDA is called **probabilistic** if, for every state q and top symbol γ of the first stack, the weights of the transitions define a probability distribution, i.e.,

$$\sum_{q \xrightarrow[\gamma_2 \rightarrow \gamma_4]{\gamma, y, \gamma_1 \rightarrow \gamma_3 / w} q' \in \delta} w = 1. \quad (19)$$

Symbol	Type	Meaning
$[N]$	$\subset \mathbb{N}$	The set $\{1, \dots, N\}$ for $N \in \mathbb{N}$.
$\Sigma, \underline{\Sigma}, \overline{\Sigma}$	alphabet	Σ is a set of symbols, $\underline{\Sigma} \stackrel{\text{def}}{=} \Sigma \cup \{\text{BOS}\}$, $\overline{\Sigma} \stackrel{\text{def}}{=} \Sigma \cup \{\text{EOS}\}$
$y, \underline{y}, \overline{y}$	$\in \Sigma$	A symbol, element of $\Sigma, \underline{\Sigma}$, or $\overline{\Sigma}$.
\mathbf{y}	$\in \Sigma^*$	A string over Σ .
\mathbf{y}_j^i	$\in \Sigma^*$	A substring of \mathbf{y} , a string.
$\llbracket y \rrbracket$	$\in \{0, 1\}^{ \Sigma }$	One-hot encoding of the symbol $y \in \Sigma$.
D	$\in \mathbb{N}$	Size of the contextual representations in the transformer or RNN.
Δ^{N-1}	$\subseteq \mathbb{R}^N$	The $N - 1$ -dimensional probability simplex.
f	$\mathbb{R}^D \times \mathbb{R}^D \rightarrow \mathbb{R}$	A scoring function.
\mathbf{f}	$\mathbb{R}^N \rightarrow \Delta^{N-1}$	A normalization function.
Q, K, V, O	$\mathbb{R}^D \rightarrow \mathbb{R}^D$	The query, key, value, and output functions.
F	$\mathbb{R}^D \rightarrow \mathbb{R}^D$	The final transformer LM transformation function.
enc	$\Sigma^* \rightarrow \mathbb{R}^D$	The string representation function.
\mathbf{r}_Π	$\underline{\Sigma} \times \mathbb{N} \rightarrow \mathbb{R}^D$	The position-augmented representation function.
L	$\in \mathbb{N}$	Number of layers.
H	$\in \mathbb{N}$	Number of heads.
\mathcal{H}	$\mathbb{R}^{HD} \rightarrow \mathbb{R}^D$	The head combining function.
$(\cdot; \dots; \cdot)$		Vertical concatenation operator of vectors or matrices.

Table 1: A summary of the notation used in the paper.

As before, we define a **run** in the 2PDA as a sequence of consecutive transitions. A run is called **accepting** if it ends in a final configuration. The yield of an accepting run is the sequence of symbols $y \in \Sigma$ that the 2PDA scans (or outputs) during the run. And the **weight** of an accepting run, $w(\boldsymbol{\pi})$, is the product of the weight of its transitions,

$$w(\boldsymbol{\pi}) \stackrel{\text{def}}{=} \prod_{n=1}^N w_n, \quad (20)$$

where w_n is the weight of transition τ_n . Finally, the **stringsum** of string \mathbf{y} is defined as the sum of the weights of all accepting runs that yield \mathbf{y} , i.e., $\Pi(\mathcal{P}, \mathbf{y})$:

$$\mathcal{P}(\mathbf{y}) \stackrel{\text{def}}{=} \sum_{\boldsymbol{\pi} \in \Pi(\mathcal{P}, \mathbf{y})} w(\boldsymbol{\pi}). \quad (21)$$

C.4 Transformer Language Models

Transformer LMs are LMs whose conditional distributions $p(\overline{y}_t \mid \mathbf{y}_{<t})$ are computed by a *transformer*. A transformer is a composition of multiple transformer *layers*, each of which implements the *attention mechanism*. We give definitions of these building blocks in what follows. Our formalization and notation closely follows [Svete and Cotterell \(2024\)](#).

Notation. We use bold unitalicized letters such as $\mathbf{x} \in \mathbb{R}^D$ to denote real-valued vectors and italicized letters $x_j \in \mathbb{R}$ for their entries. Capital bold letters such as $\mathbf{X} \in \mathbb{R}^{N \times D}$ denote matrices. All vectors are *column* vectors unless transposed. We define the vertical stacking operator $(\cdot; \dots; \cdot)$, which denotes the vertical concatenation of the D -dimensional *column* vectors $\mathbf{x}_1, \dots, \mathbf{x}_N$ into a ND -dimensional vector $(\mathbf{x}_1; \dots; \mathbf{x}_N) \in \mathbb{R}^{ND}$ and the concatenation of the D -dimensional *row* vectors $\mathbf{x}_1^\top, \dots, \mathbf{x}_N^\top$ into a matrix $\mathbf{X} \in \mathbb{R}^{N \times D}$ with N rows and D columns. Given the matrix $\mathbf{X} = (\mathbf{x}_1^\top; \dots; \mathbf{x}_N^\top)$, we write $\mathbf{X}_n = (\mathbf{x}_1^\top; \dots; \mathbf{x}_n^\top)$ for the submatrix composed of the first n rows. We call a function $f: \mathbb{R}^D \times \mathbb{R}^D \rightarrow \mathbb{R}$ whose purpose is to evaluate the compatibility of two vectors a **scoring function**. A **normalization function** $\mathbf{f}: \mathbb{R}^N \rightarrow \Delta^{N-1}$ maps vectors in \mathbb{R}^N to N probabilities. Here, $\Delta^{N-1} \stackrel{\text{def}}{=} \left\{ \mathbf{x} \in [0, 1]^N \mid \sum_{n=1}^N x_n = 1 \right\}$ is the $N - 1$ -dimensional probability simplex. This notation is summarized in Tab. 1.

The Attention Mechanism. The attention mechanism works as follows. It takes a **query** vector $\mathbf{q} \in \mathbb{R}^D$ and two matrices: The matrix $\mathbf{K} \in \mathbb{R}^{N \times D}$ of **keys** and the matrix $\mathbf{V} \in \mathbb{R}^{N \times D}$ of **values** and computes a weighted average of the value vectors based on the compatibilities of the key vectors to the query vector, as scored by a scoring function f . A formal definition is given below.

Definition C.4 (Attention Mechanism). The *attention mechanism* $\text{Att}: \mathbb{R}^D \times \mathbb{R}^{N \times D} \times \mathbb{R}^{N \times D} \rightarrow \mathbb{R}^D$ is defined as

$$\text{Att}(\mathbf{q}, \mathbf{K}, \mathbf{V}) \stackrel{\text{def}}{=} \sum_{n=1}^N s_n \mathbf{v}_n \quad (22)$$

where $\mathbf{q} \in \mathbb{R}^D$ be a query vector and let $\mathbf{K} = (\mathbf{k}_1^\top; \dots; \mathbf{k}_N^\top) \in \mathbb{R}^{N \times D}$ and $\mathbf{V} = (\mathbf{v}_1^\top; \dots; \mathbf{v}_N^\top) \in \mathbb{R}^{N \times D}$ be matrices of keys and values, respectively, and

$$\mathbf{s} \stackrel{\text{def}}{=} \mathbf{f}(f(\mathbf{q}, \mathbf{k}_1), \dots, f(\mathbf{q}, \mathbf{k}_N)) \quad (23)$$

is the vector of normalized scores between the query \mathbf{q} and the keys in \mathbf{K} , f is a scoring function and \mathbf{f} is a normalization function.

Attention types. Attention weights are computed by normalizing the scores $f(\mathbf{q}, \mathbf{k}_1), \dots, f(\mathbf{q}, \mathbf{k}_L)$. The choice of the projection function \mathbf{f} determines the type of attention and has concrete implications on representational capacity (Hao et al., 2022). We focus on the **hard attention** projection function.

Definition C.5. *Hard attention* is computed with the hardmax projection function:

$$\text{hardmax}(\mathbf{x})_d \stackrel{\text{def}}{=} \begin{cases} \frac{1}{m} & \text{if } d \in \text{argmax}(\mathbf{x}) \\ 0 & \text{otherwise} \end{cases} \quad (24)$$

for $d \in [D]$, where $\mathbf{x} \in \mathbb{R}^D$ and $m \stackrel{\text{def}}{=} |\text{argmax}(\mathbf{x})|$ is the cardinality of the argmax set.

Definition C.6. A *multi-layer perceptron* (MLP) $F: \mathbb{R}^D \rightarrow \mathbb{R}^D$ is a function defined as the composition of elementary functions $\mathbf{f}_1, \dots, \mathbf{f}_L$

$$F(\mathbf{x}) \stackrel{\text{def}}{=} \mathbf{f}_L \circ \mathbf{f}_{L-1} \circ \dots \circ \mathbf{f}_1(\mathbf{x}), \quad (25)$$

where each function \mathbf{f}_ℓ for $\ell \in [L]$ is defined as

$$\mathbf{f}_\ell(\mathbf{x}) \stackrel{\text{def}}{=} \beta(\mathbf{W}_\ell \mathbf{x} + \mathbf{b}_\ell) \quad \ell \in [L-1] \quad (26a)$$

$$\mathbf{f}_L(\mathbf{x}) \stackrel{\text{def}}{=} \mathbf{W}_L \mathbf{x} + \mathbf{b}_L, \quad (26b)$$

where $\mathbf{W}_\ell \in \mathbb{R}^{D \times D}$ is a square weight matrix specific to layer ℓ , $\mathbf{b}_\ell \in \mathbb{R}^D$ is a bias vector, and β is an element-wise non-linear activation function. The function \mathbf{f}_1 is called the **input layer**, the function \mathbf{f}_L is called the **output layer**, and the function \mathbf{f}_ℓ for $\ell = 2, \dots, L-1$ are called **hidden layers**.²³

The Transformer Architecture. A transformer layer uses the attention mechanism to compute augmented representations $\mathbf{z}_t = \text{Att}(\mathbf{q}_t, \mathbf{K}_t, \mathbf{V}_t)$ of the input representations $\mathbf{X}_t = (\mathbf{x}_1; \dots; \mathbf{x}_t)$. The query \mathbf{q}_t , the keys \mathbf{K}_t , and values \mathbf{V}_t are all transformations of the input representations \mathbf{X}_t .

Definition C.7. Given query, key, value, and **output** functions $Q, K, V, O: \mathbb{R}^D \rightarrow \mathbb{R}^D$, a **transformer layer** is a function $\mathcal{L}: \mathbb{R}^{T \times D} \rightarrow \mathbb{R}^{T \times D}$ that computes

$$\mathcal{L}(\mathbf{x}_1^\top; \dots; \mathbf{x}_T^\top) = (\mathbf{z}_1^\top; \dots; \mathbf{z}_T^\top) \in \mathbb{R}^{T \times D} \quad (27)$$

for $t \in [T]$ where

$$\mathbf{a}_t \stackrel{\text{def}}{=} \text{Att}(\mathbf{q}_t, \mathbf{K}_t, \mathbf{V}_t) + \mathbf{x}_t \in \mathbb{R}^D \quad (28a)$$

$$\mathbf{z}_t \stackrel{\text{def}}{=} O(\mathbf{a}_t) + \mathbf{a}_t \in \mathbb{R}^D. \quad (28b)$$

²³Note that we refer to MLPs by the number of hidden layers, e.g., a one-layer-MLP is an MLP with one *hidden* layer.

Here, we define

$$\mathbf{q}_t \stackrel{\text{def}}{=} Q(\mathbf{x}_t) \in \mathbb{R}^D \quad (29a)$$

$$\mathbf{K}_t \stackrel{\text{def}}{=} \left(K(\mathbf{x}_1)^\top; \dots; K(\mathbf{x}_t)^\top \right) \in \mathbb{R}^{t \times D} \quad (29b)$$

$$\mathbf{V}_t \stackrel{\text{def}}{=} \left(V(\mathbf{x}_1)^\top; \dots; K(\mathbf{x}_t)^\top \right) \in \mathbb{R}^{t \times D}. \quad (29c)$$

Note: For simplicity, we do not include layer normalization.

The functions Q, K, V are usually implemented as linear transformations and O as an MLP. Some theoretical literature, however, also considers more general function classes, e.g., all smooth functions (Hahn, 2020).

Without further modification, the transformations applied by the transformer layer are position-invariant, which necessitates the addition of explicit positional information.

Definition C.8. A symbol **representation function** is a function $\mathbf{r}: \underline{\Sigma} \rightarrow \mathbb{R}^{D_r}$ and a **positional encoding** is a function $\Pi: \mathbb{N} \rightarrow \mathbb{R}^{D_\Pi}$. A **position-augmented representation function** $\mathbf{r}_\Pi: \underline{\Sigma} \times \mathbb{N} \rightarrow \mathbb{R}^D$ (with $D = D_r + D_\Pi$) is defined as

$$\mathbf{r}_\Pi \left(\underline{y}_t, t \right) \stackrel{\text{def}}{=} \left(\mathbf{r}(\underline{y}_t); \Pi(t) \right). \quad (30)$$

Definition C.9. A **static encoding** \mathcal{R} is a function $\mathcal{R}: \Sigma^T \rightarrow \mathbb{R}^{T \times D}$ defined for any $T \in \mathbb{N}$ as

$$\mathcal{R}(\mathbf{y}) \stackrel{\text{def}}{=} \left(\mathbf{r}_\Pi(y_1, 1)^\top; \dots; \mathbf{r}_\Pi(y_T, T)^\top \right). \quad (31)$$

Multiple transformer layers are stacked into a transformer, which computes the (deep) contextual representations of all symbols in the string.

Definition C.10. For $L \in \mathbb{N}$, an L -layer **transformer** \mathcal{T} is defined as

$$\mathcal{T}(\mathcal{R}) \stackrel{\text{def}}{=} \mathcal{L}_L \circ \dots \circ \mathcal{L}_1 \circ \mathcal{R}, \quad (32)$$

where \mathcal{L}_ℓ for $\ell \in [L]$ are transformer layers and \mathcal{R} is a static encoding.

A transformer computes the contextual representations of the symbols $\mathbf{y} = y_1 \dots y_T$ as

$$\left(\mathbf{z}_1^\top; \dots; \mathbf{z}_T^\top \right) \stackrel{\text{def}}{=} \left(\mathbf{x}_1^{L\top}; \dots; \mathbf{x}_T^{L\top} \right) \stackrel{\text{def}}{=} \mathcal{T}(\mathcal{R})(\mathbf{y}). \quad (33)$$

If \mathcal{R} is clear from the context or arbitrary, we will omit it as an argument to \mathcal{T} and just write $\mathcal{T}(\mathbf{y})$.

Definition C.11. Given a transformer \mathcal{T} , a final representation transformation function $F: \mathbb{R}^D \rightarrow \mathbb{R}^D$, and a string $\mathbf{y} \in \Sigma^*$ with $|\mathbf{y}| = T$, we define the **encoding function** enc as

$$\text{enc}(\mathbf{y}) \stackrel{\text{def}}{=} F(\mathbf{z}_T) \quad (34)$$

where \mathbf{z}_T is the representation of the T^{th} symbol in \mathbf{y} computed by \mathcal{T} , i.e., $(\mathbf{z}_1^\top; \dots; \mathbf{z}_T^\top) = \mathcal{T}(\mathbf{y})$.

Transformer Language Models. So far, we have only defined how the transformer architecture can be used to compute the contextual representations of the symbols. To complete the definition, we define a transformer *language model* as follows.

Definition C.12. A **transformer LM** $p_{\mathcal{T}}$ is the representation-based autoregressive LM with the representation function enc from Eq. (34). That is, $p_{\mathcal{T}}$ defines the conditional probability distributions

$$p_{\mathcal{T}}(y_t \mid \mathbf{y}_{<t}) \stackrel{\text{def}}{=} \text{sparsemax}(\mathbf{E} \text{enc}(\mathbf{y}_{<t}))_{y_t}. \quad (35)$$

C.5 Weighted Finite-state Transducers

Definition C.13. A (rational-)weighted finite-state transducer (WFST) is the tuple $(\Sigma, \Xi, Q, \delta, \lambda, \rho)$ where Σ and Ξ are the input and output alphabets, respectively, Q is a finite set of states, $\delta \subseteq Q \times \Sigma_\varepsilon \times \Xi_\varepsilon \times \mathbb{Q}_{\geq 0} \times Q$ is a finite set of weighted transitions where we write transitions $(q, y, x, w, q') \in \delta$ as $q \xrightarrow{y:x/w} q'$, and $\lambda, \rho: Q \rightarrow \mathbb{Q}_{\geq 0}$ are functions that assign each state its initial and final weight, respectively.

WFSTs are thus a special class of (weighted) finite-state automata that operate over two alphabets. Just like PFSA are a generalization of unweighted FSAs, WFSTs represent a more general class of machines than unweighted FSTs, such as the ones used in the definition of regular reducibility (cf. Def. 3.1). The weighted versions will prove useful when talking about various aspects of equivalence with regular reducibility. As a special case of weighted finite-state automata, WFSTs compute weights of their inputs—in this case, weights of *pairs* of strings (the inputs and their outputs):

$$\mathcal{T}(\mathbf{y}, \mathbf{x}) \stackrel{\text{def}}{=} \sum_{\pi \in \Pi(\mathcal{T}, (\mathbf{y}, \mathbf{x}))} w(\pi). \quad (36)$$

Of interest are also marginal sums of the inputs, i.e.,

$$\mathcal{T}(\mathbf{y}) \stackrel{\text{def}}{=} \sum_{\mathbf{x} \in \Xi^*} \mathcal{T}(\mathbf{y}, \mathbf{x}). \quad (37)$$

C.5.1 Operations on WFSTs

We will use WFSTs as building blocks in the exposition of regular reducibility (cf. §3.2). In this subsection, we outline some operations on the computational models that will be particularly useful.

Composition. The **composition** of two WFSTs results in a WFST that, similarly to function composition, maps the strings with the first WFST, passes them to the second one, and returns the (weight of the) output of the second transducer (Pereira and Riley, 1997). More formally, given the WFSTs $\mathcal{T}_1 = (\Sigma, \Upsilon, Q, \delta, \lambda, \rho)$ and $\mathcal{T}_2 = (\Upsilon, \Xi, Q, \delta, \lambda, \rho)$, their composition $\mathcal{T}_2 \circ \mathcal{T}_1$ computes

$$(\mathcal{T}_2 \circ \mathcal{T}_1)(\mathbf{y}, \mathbf{x}) \stackrel{\text{def}}{=} \sum_{z \in \Upsilon^*} \mathcal{T}_2(z, \mathbf{y}) \cdot \mathcal{T}_1(\mathbf{x}, z). \quad (38)$$

Intuitively, Eq. (38) computes the weight of all the possible ways of mapping \mathbf{y} to \mathbf{x} by first mapping \mathbf{y} to some $z \in \Upsilon^*$ and then mapping that z into \mathbf{x} . The WFST $(\mathcal{T}_2 \circ \mathcal{T}_1)$ can be computed from \mathcal{T}_1 and \mathcal{T}_2 in time $\mathcal{O}((|Q_1| + |\delta_1|)(|Q_2| + |\delta_2|))$ (Mohri, 2009).

Projecting a WFST. Any WFST can be **projected** onto an (input or output) weighted finite-state automaton (WFSA)²⁴ that computes the cumulative weights of individual input or output strings (Mohri et al., 2008). Given a WFST $\mathcal{T} = (\Sigma, \Xi, Q, \delta, \lambda, \rho)$, its input projection is the automaton $\mathcal{A}_{\mathcal{I}} = (\Sigma, Q, \delta_{\mathcal{I}}, \lambda, \rho)$ that computes

$$\mathcal{A}_{\mathcal{I}}(\mathbf{y}) \stackrel{\text{def}}{=} \sum_{\mathbf{x} \in \Xi^*} \mathcal{T}(\mathbf{y}, \mathbf{x}). \quad (39)$$

The output projection automaton is defined analogously. While it might not be immediately obvious that Eq. (39) represents the computations of a WFSA, $\mathcal{A}_{\mathcal{I}}$ can be easily constructed by ignoring the output labels on the transitions (and additively merging any transitions that become identical after the removal of the output labels). That is:

$$\delta_{\mathcal{I}} = \left\{ q \xrightarrow{y/\sum_{q' \in \delta} w} q' \mid q, q' \in Q, y \in \Sigma_\varepsilon \right\}. \quad (40)$$

²⁴A weighted finite-state automaton is a generalization of a probabilistic finite-state automaton where the weights do not have to form probability distributions. We also extend the definition of WFSA to allow ε -transitions, i.e. with symbols from Σ_ε . Note that every such WFSA can be converted into a weakly equivalent WFSA without ε -transitions.

Lifting a PFSA. **Lifting** transforms a given PFSA into a trivial WFST that implements the identity function $\mathbf{y} \mapsto \mathbf{y}$. Concretely, given the PFSA $\mathcal{A} = (\Sigma, Q, \lambda, \rho, \delta)$, its lifted WFST $\mathcal{T}_{\mathcal{A}} = (\Sigma, \Sigma, Q, \delta_{\mathcal{T}}, \lambda, \rho)$ defines the transitions

$$\delta_{\mathcal{T}} \stackrel{\text{def}}{=} \left\{ q \xrightarrow{y:y/w} q' \mid q \xrightarrow{y/w} q' \in \delta \right\}. \quad (41)$$

It is easy to see that $\mathcal{T}_{\mathcal{A}}$ computes

$$\mathcal{T}_{\mathcal{A}}(\mathbf{y}, \mathbf{y}') = \mathbb{1} \{ \mathbf{y} = \mathbf{y}' \} \mathcal{A}(\mathbf{y}). \quad (42)$$

Lifting is useful when one wants to compose a PFSA with a WFST, as is the case when discussing regular reducibility.

Inverting a WFST. The **inverted** WFST of some WFST $\mathcal{T} = (\Sigma, \Xi, Q, \delta, \lambda, \rho)$ is the WSFT $\mathcal{T}^{-1} = (\Xi, \Sigma, Q, \delta^{-1}, \lambda, \rho)$ that maps the outputs of \mathcal{T} to their inputs. δ^{-1} is defined as

$$\delta^{-1} \stackrel{\text{def}}{=} \left\{ q \xrightarrow{x:y/w} q \mid q \xrightarrow{y:x/w} q \in \delta \right\}, \quad (43)$$

i.e., it is composed of transitions from δ with their input–output labels flipped.

D Separation of Alphabets

In our analysis, we use a different alphabet for the chain of thought (Γ) than for the final output (Σ) to clarify the distinction between output that encodes automata configurations vs the output that constitutes the weighted output language of that automaton. However, we can also define a regular function for the case where $\Gamma = \Sigma$. For this, we choose a specific symbol in Σ , e.g., X , whose first occurrence signifies the boundary between the chain of thought and the language output. Then, ϕ can be defined as:

$$\phi(\mathbf{y} X \mathbf{z}) \stackrel{\text{def}}{=} \mathbf{z} \quad (44)$$

where $\mathbf{y} \in \Sigma^* \setminus \{X\}$, and $\mathbf{z} \in \Sigma^*$. This can be easily modeled by the following simple two-state transducer:




Figure 4: A simple WFST \mathcal{T} implementing ϕ .

Note that in this case, the chain of thought has to be confined in its entirety as one string, with the entire final output coming afterward. This simply means that the simulated Turing machine first performs all the computations it needs on the working tape, and then writes the final string to the output tape in real time. The proof of Thm. 4.4 can easily be adapted to work in this setting by only tagging the ε symbols in Γ , e.g. by symbols a and b and defining ϕ to remove them until the occurrence of another symbol, e.g. c .

E Proofs: Regular Reducibility

E.1 Nondeterminism in PFSA

Theorem 3.1. *Let $\mathcal{A} = (\Sigma, Q, \delta, \lambda, \rho)$ be a PFSA. Then, there exists a deterministic PFSA \mathcal{A}' over the alphabet $\Sigma \times Q$ and with the state space $\Sigma \times Q$ that is regularly reducible to \mathcal{A} .*

Proof. Given the PFSA $\mathcal{A} = (\Sigma, Q, \lambda, \rho, \delta)$, we want to show the existence of a *deterministic* PFSA \mathcal{A}' over the alphabet $\Sigma \times Q$ that is regularly reducible to \mathcal{A} . We construct the deterministic PFSA $\mathcal{A}' = (\Sigma \times Q, \Sigma \times Q, \delta', \lambda', \rho')$ as follows.

$$\delta' \stackrel{\text{def}}{=} \left\{ (y, q) \xrightarrow{(y', q')/w} (y', q') \mid q \xrightarrow{y'/w} q' \in \delta, y \in \Sigma \right\} \quad (45a)$$

$$\lambda'(y, q) \stackrel{\text{def}}{=} \frac{\lambda(q)}{|\Sigma|}, \quad \forall y \in \Sigma, q \in Q \quad (45b)$$

$$\rho'(y, q) \stackrel{\text{def}}{=} \rho(q), \quad \forall y \in \Sigma, q \in Q \quad (45c)$$

We first prove that \mathcal{A}' is probabilistic and deterministic.

- **Probabilistic:** We compute

$$\sum_{(y, q) \in \Sigma \times Q} \lambda'(y, q) = \sum_{(y, q) \in \Sigma \times Q} \frac{\lambda(q)}{|\Sigma|} \quad (46a)$$

$$= \sum_{y \in \Sigma} \sum_{q \in Q} \frac{\lambda(q)}{|\Sigma|} \quad (46b)$$

$$= \sum_{q \in Q} \lambda(q) \quad (46c)$$

$$= 1 \quad (46d)$$

and, for any $(y, q) \in \Sigma \times Q$,

$$\sum_{(y, q) \xrightarrow{(y', q')/w} (y', q') \in \delta'} w + \rho'(y, q) = \sum_{(y, q) \xrightarrow{(y', q')/w} (y', q') \in \delta'} w + \rho(q) \quad (47a)$$

$$= \sum_{q \xrightarrow{w/y'} y' \in \delta} w + \rho(q) \quad (47b)$$

$$= 1 \quad (47c)$$

- **Deterministic:** Let $(y, q) \in \Sigma \times Q$ be a state of \mathcal{A}' and (y', q') a symbol in its alphabet. By definition of δ' , (y', q') uniquely determines the target state, which is identical to the symbol— (y', q') .

We now show that \mathcal{A}' is indeed regularly reducible to \mathcal{A} . The alphabet $\Sigma \times Q$ clearly satisfies the Σ -augmentation condition. We define the regular function $\phi(y, q) \stackrel{\text{def}}{=} y$, an instance of the general function described in §3.2. We will show that \mathcal{A}' is weakly equivalent to $\mathcal{A} \circ \phi$, or, equivalently, that $\mathcal{A}' \circ \phi^{-1}$ is weakly equivalent to \mathcal{A} .

Since \mathcal{A}' is an *acceptor* of strings in $(\Sigma \times Q)^*$, we first transform it into weighted finite-state *transducer* $\mathcal{T}_{\mathcal{A}'}$ implementing the mapping²⁵

$$\mathcal{T}_{\mathcal{A}'}(\mathbf{x}, \mathbf{x}') \stackrel{\text{def}}{=} \mathbb{1} \{ \mathbf{x} = \mathbf{x}' \} \mathcal{A}(\mathbf{x}) \quad (48)$$

for $\mathbf{x}, \mathbf{x}' \in \Delta^*$. This is a simple WFST with transitions $\delta_{\mathcal{T}_{\mathcal{A}'}} \stackrel{\text{def}}{=} \left\{ q \xrightarrow{(y, q):(y, q)/w} q' \mid q \xrightarrow{(y, q)/w} q' \in \delta' \right\}$.

$\mathcal{T}_{\mathcal{A}'}$ can then be composed with the weighed version of the FST implementing ϕ^{-1} , $\mathcal{T}_{\phi^{-1}}$:

$$\mathcal{T}_{\phi^{-1}}(\mathbf{y}, \mathbf{x}) \stackrel{\text{def}}{=} \mathbb{1} \{ \mathbf{x} \in \phi^{-1}(\mathbf{y}) \}. \quad (49)$$

²⁵See App. C.5 for a definition of weighted finite-state transducers and the operations on them.

$\mathcal{T}_{\mathcal{A}'} \circ \mathcal{T}_{\phi^{-1}}$ then computes, for $\mathbf{y} \in \Sigma^*$, $\mathbf{x}' \in \Delta^*$,

$$\left(\mathcal{T}_{\mathcal{A}'} \circ \mathcal{T}_{\phi^{-1}}\right)(\mathbf{y}, \mathbf{x}') = \sum_{\mathbf{x} \in \Delta^*} \mathcal{T}_{\phi^{-1}}(\mathbf{y}, \mathbf{x}) \cdot \mathcal{T}_{\mathcal{A}'}(\mathbf{x}, \mathbf{x}') \quad (50a, \text{Eq. 15.4 in Pereira and Riley (1997).})$$

$$= \sum_{\mathbf{x} \in \Delta^*} \mathbb{1}\{\mathbf{x} \in \phi^{-1}(\mathbf{y})\} \cdot \mathcal{T}_{\mathcal{A}'}(\mathbf{x}, \mathbf{x}') \quad (50b)$$

$$= \sum_{\mathbf{x} \in \phi^{-1}(\mathbf{y})} \mathcal{T}_{\mathcal{A}'}(\mathbf{x}, \mathbf{x}') \quad (50c)$$

$$= \sum_{\mathbf{x} \in \phi^{-1}(\mathbf{y})} \mathbb{1}\{\mathbf{x} = \mathbf{x}'\} \mathcal{A}(\mathbf{x}) \quad (50d)$$

$$= \mathbb{1}\{\mathbf{x}' \in \phi^{-1}(\mathbf{y})\} \mathcal{A}(\mathbf{x}') \quad (50e)$$

We then turn $\mathcal{T}_{\mathcal{A}'} \circ \mathcal{T}_{\phi^{-1}}$ into a PFSA in the standard way by projecting the WFST onto the transition input labels (Mohri et al., 2008). This results in a PFSA that assigns the string $\mathbf{y} \in \Sigma^*$ the probability that equals the sum of all possible mappings of \mathbf{y} to any $\mathbf{x}' \in \Delta^*$:

$$\left(\mathcal{T}_{\mathcal{A}'} \circ \mathcal{T}_{\phi^{-1}}\right)(\mathbf{y}) \stackrel{\text{def}}{=} \sum_{\mathbf{x}' \in \Delta^*} \left(\mathcal{T}_{\mathcal{A}'} \circ \mathcal{T}_{\phi^{-1}}\right)(\mathbf{y}, \mathbf{x}') \quad (51a)$$

$$= \sum_{\mathbf{x}' \in \Delta^*} \mathbb{1}\{\mathbf{x}' \in \phi^{-1}(\mathbf{y})\} \mathcal{A}(\mathbf{x}') \quad (51b)$$

$$= \sum_{\mathbf{x}' \in \phi^{-1}(\mathbf{y})} \mathcal{A}(\mathbf{x}') \quad (51c)$$

$$= \mathcal{A}'(\phi^{-1}(\mathbf{y})) \quad (51d)$$

$$= (\mathcal{A}' \circ \phi^{-1})(\mathbf{y}) \quad (51e)$$

$$= \mathcal{A}(\mathbf{y}) \quad (51f)$$

This shows that $\mathcal{A}' \circ \phi^{-1}$ is weakly equivalent to \mathcal{A} , which finishes the proof. \blacksquare

E.2 Nondeterminism in PPDA's

Theorem E.1. *Let $\mathcal{P} = (Q, \Sigma, \Gamma, \delta, (q_l, S), (q_\varphi, \varepsilon))$ be a PPDA. Then, there exists a deterministic PPDA \mathcal{P}' over the alphabet $(\Sigma_\varepsilon) \times Q \times \Gamma$ with the state space $(\Sigma_\varepsilon) \times Q \times \Gamma$ that is regularly reducible to \mathcal{P} .*

Proof. Given a PPDA $\mathcal{P} = (Q, \Sigma, \Gamma, \delta, (q_l, S), (q_\varphi, \varepsilon))$, we prove the existence of a deterministic PPDA \mathcal{P}' over the alphabet $\Sigma_\varepsilon \times Q \times \Gamma$ that is regularly reducible to \mathcal{P} .

We construct the deterministic PPDA $\mathcal{P}' = (\Sigma_\varepsilon \times Q \times \Gamma, \Sigma_\varepsilon \times Q \times \Gamma, \Gamma, \delta', (q_l, S), (q_\varphi, \varepsilon))$ with the set of transitions

$$\delta' \stackrel{\text{def}}{=} \{(y, q, \gamma), \gamma' \xrightarrow{(y', q', \gamma')/w} (y', q', \gamma'), \gamma \mid q, \gamma' \xrightarrow{y'/w} q', \gamma \in \delta, y \in \Sigma_\varepsilon, \gamma \in \Gamma\}. \quad (52)$$

- **Probabilistic:** For any $(y, q, \gamma) \in \Sigma_\varepsilon \times Q \times \Gamma$ and $\gamma' \in \Gamma$,

$$\sum_{(y, q, \gamma), \gamma' \xrightarrow{(y', q', \gamma')/w} (y', q', \gamma'), \gamma} w = \sum_{q, \gamma' \xrightarrow{y'/w} q', \gamma} w \quad (53a)$$

$$= 1 \quad (53b)$$

- **Deterministic:** Just like in the finite-state case, let $(y, q, \gamma) \in \Sigma_\varepsilon \times Q \times \Gamma$ be a state, $\gamma' \in \Gamma$ a stack symbol and $(y', q', \gamma') \in \Sigma_\varepsilon \times Q \times \Gamma$ an input symbol. Then (y', q', γ') uniquely determines the next configuration.

We now show that \mathcal{P}' is regularly reducible to \mathcal{P} . More precisely, we define the regular function $\phi(y, q, \gamma) \stackrel{\text{def}}{=} y$ and show that $\mathcal{P}' \circ \phi^{-1}$ is weakly equivalent to \mathcal{P} .

Let \mathcal{G}' be a CFG²⁶ that is equivalent to \mathcal{P}' (see Hopcroft et al., 2001, Section 6.3.2 for the construction). Using grammars, rather than pushdown automata, will simplify the proof as there are well-known constructions (e.g., Bar-Hillel et al. (1961), Pasti et al. (2023)) for composing CFGs and FSTs. We therefore prove that $\mathcal{G}' \circ \phi^{-1}$ is weakly equivalent to \mathcal{P} .

Recall that ϕ^{-1} can be implemented as an FST $\mathcal{T}_{\phi^{-1}}$ defined as

$$\mathcal{T}_{\phi^{-1}}(\mathbf{y}, \mathbf{x}) \stackrel{\text{def}}{=} \mathbb{1}\{\mathbf{x} \in \phi^{-1}(\mathbf{y})\}. \quad (54)$$

Adapting Corollary 1 from Pasti et al. (2023) to the FST case, we get

$$\left(\mathcal{G}' \circ \mathcal{T}_{\phi^{-1}}\right)(\mathbf{y}, \mathbf{x}) = \mathcal{G}'(\mathbf{x}) \cdot \mathcal{T}_{\phi^{-1}}(\mathbf{y}, \mathbf{x}) \quad (55a)$$

$$= \mathcal{G}'(\mathbf{x}) \cdot \mathbb{1}\{\mathbf{x} \in \phi^{-1}(\mathbf{y})\}. \quad (55b)$$

Then, just like in the regular case, we compute $\left(\mathcal{G}' \circ \mathcal{T}_{\phi^{-1}}\right)(\mathbf{y})$ by summing over all possible $\mathbf{x} \in \Sigma^*$ values:

$$\left(\mathcal{G}' \circ \mathcal{T}_{\phi^{-1}}\right)(\mathbf{y}) \stackrel{\text{def}}{=} \sum_{\mathbf{x} \in \Sigma^*} \left(\mathcal{G}' \circ \mathcal{T}_{\phi^{-1}}\right)(\mathbf{y}, \mathbf{x}) \quad (56a)$$

$$= \sum_{\mathbf{x} \in \Sigma^*} \mathcal{G}'(\mathbf{x}) \mathbb{1}\{\mathbf{x} \in \phi^{-1}(\mathbf{y})\} \quad (56b)$$

$$= \sum_{\mathbf{x} \in \phi^{-1}(\mathbf{y})} \mathcal{G}'(\mathbf{x}) \quad (56c)$$

$$= \sum_{\mathbf{x} \in \phi^{-1}(\mathbf{y})} \mathcal{P}'(\mathbf{x}) \quad (56d)$$

$$= \mathcal{P}'(\phi^{-1}(\mathbf{y})) \quad (56e)$$

$$= (\mathcal{P}' \circ \phi^{-1})(\mathbf{y}) \quad (56f)$$

$$= \mathcal{P}(\mathbf{y}). \quad (56g)$$

This shows that $\mathcal{G}' \circ \mathcal{T}_{\phi^{-1}}$ is weakly equivalent to \mathcal{P} , thus $\mathcal{P}' \circ \mathcal{T}_{\phi^{-1}}$ is weakly equivalent to \mathcal{P} . This concludes the proof. \blacksquare

F Proofs: Representational Capacity of Neural LMs

F.1 Finite-state Language Models

Theorem 4.1. *For any regular LM, there exists a weakly equivalent CoT-augmented constant-precision Elman RNN LM.*

Proof. Let $\mathcal{A} = (\Sigma, Q, \lambda, \rho, \delta)$ be a PFSA. We divide the definition of a weakly equivalent Elman LM $\mathcal{R} = (\Sigma \times Q, D, \mathbf{U}, \mathbf{V}, \mathbf{b}, \boldsymbol{\eta})$ with the output matrix \mathbf{E} into multiple steps.²⁷

Note: The paper assumes a BOS-padding of the input strings. Since CoT-augmented LMs work over a potentially larger alphabet, the padding symbol has to be changed accordingly. Particularly, the CoT-augmented RNN will work over the alphabet $\Sigma \times Q$. For simplicity, assume that any input string is padded by the symbol (BOS, q_0) for some arbitrary (but fixed) state $q_0 \in Q$.

²⁶Similar to the notation used for PPDAs, we use $\mathcal{G}(\mathbf{y})$ to denote the weight of the string \mathbf{y} under the grammar \mathcal{G} .

²⁷Notice that \mathcal{R} is CoT-augmented by construction, as it works over the alphabet $\Delta \stackrel{\text{def}}{=} \Sigma \times Q$. Δ also satisfies the Σ -augmentation condition.

Hidden states. Let $D = |\Sigma_\varepsilon| |Q|$. We will represent the input symbols with the one-hot representation function $\llbracket \cdot, \cdot \rrbracket$ that computes $\llbracket \underline{y}, q \rrbracket \in \{0, 1\}^{|\Sigma_\varepsilon| |Q|}$.²⁸

$$\llbracket \underline{y}, q \rrbracket_{y,q} \stackrel{\text{def}}{=} 1 \quad (57)$$

while other entries of the vector are zero.²⁹ We can define $\boldsymbol{\eta} \stackrel{\text{def}}{=} \mathbf{0}_D$.

Recurrence. Conveniently, the CoT-augmented RNN will store the current PFSA state in its output “symbol”. Therefore, its next hidden state \mathbf{h}_t will not, in fact, depend on the previous one (\mathbf{h}_{t-1}); \mathbf{h}_{t-1} will only be used to compute the next-symbol–state output distribution through Eq. (2). The RNN therefore only has to appropriately incorporate the input information. With this in mind, we define the following Elman RNN parameters \mathbf{U} , \mathbf{V} , and \mathbf{b} :

$$\mathbf{U} \stackrel{\text{def}}{=} \mathbf{0}_D \quad \mathbf{U} \in \mathbb{R}^{D \times D} \quad (58a)$$

$$\mathbf{V} \stackrel{\text{def}}{=} \mathbf{I}_D \quad \mathbf{V} \in \mathbb{R}^{D \times D} \quad (58b)$$

$$\mathbf{b} \stackrel{\text{def}}{=} \mathbf{0}_D \quad \mathbf{b} \in \mathbb{R}^D \quad (58c)$$

where $\mathbf{0}_D$ is a $D \times D$ -dimensional matrix of zeros, \mathbf{I}_D is the D -dimensional identity matrix, and $\mathbf{0}_D$ is a D -dimensional vector of zeros. Then, we define the output matrix $\mathbf{E} \in \mathbb{R}^{|\Sigma_\varepsilon| |Q| \times D}$ as

$$E_{(y',q'),(y,q)} \stackrel{\text{def}}{=} p(q', y' | q) \quad q, q' \in Q, \quad y \in \Sigma_\varepsilon, y' \in \Sigma_\varepsilon \quad (59a)$$

$$E_{(\varepsilon,q),(\text{BOS},q_0)} \stackrel{\text{def}}{=} \lambda(q) \quad q \in Q \quad (59b)$$

$$E_{(\text{EOS},q),(\varepsilon,q)} \stackrel{\text{def}}{=} \rho(q) \quad q \in Q, \quad y \in \Sigma_\varepsilon. \quad (59c)$$

Computation of string probabilities. We now show that \mathcal{R} with \mathbf{E} defines the same conditional distributions as \mathcal{A} , meaning that it implicitly defines the same probability distribution over strings. As mentioned above, we assume that the input string is prefixed with the (BOS, q_0) symbol. Let the previously generated output symbol $\in \Delta$ be (\underline{y}, q) (in the first step, the “generated” pair is (BOS, q_0)). We get that

$$\mathbf{h}_t = \mathbf{H}(\mathbf{U}\mathbf{h}_{t-1} + \mathbf{V}\llbracket \underline{y}, q \rrbracket + \mathbf{b}) \quad (60a)$$

$$= \mathbf{H}(\mathbf{0}_D\mathbf{h}_{t-1} + \mathbf{V}\llbracket \underline{y}, q \rrbracket + \mathbf{0}_D) \quad (60b)$$

$$= \mathbf{H}(\mathbf{V}\llbracket \underline{y}, q \rrbracket) \quad (60c)$$

$$= \llbracket \underline{y}, q \rrbracket \quad (60d)$$

The one-hot encoding $\llbracket \underline{y}, q \rrbracket$ is then used to “index” the appropriate column of the output matrix \mathbf{E} as $\mathbf{E}\mathbf{h}_t$, which contains the probabilities of the *next* symbol–state pairs defined by the PFSA, as per Eq. (59a):

$$(\mathbf{E}\mathbf{h}_t)_{(q',\bar{y}')} = E_{(\bar{y}',q'),(y,q)} = \begin{cases} \lambda(q') & \text{if } \underline{y} = \text{BOS}, q = q_0, \bar{y}' = \varepsilon \\ p(q', \bar{y}' | q) & \text{otherwise} \end{cases}. \quad (61)$$

Passing $\mathbf{E}\mathbf{h}_t$ through the normalization function, the next state and symbol are sampled and passed as input to the RNN, repeating the update step. Since the Elman RNN following these dynamics by constructions generates paths (sequences of states in the output symbols) with the same probabilities as the input automaton, it enumerates all accepting paths of any string $\mathbf{y} \in \Sigma^*$ with the same probabilities as the original PFSA. Applying the transformation ϕ to the produced outputs and summing over sequences that yield the same string will thus result in strings sampled from the PFSA. This finishes the proof. ■

Theorem 4.2. *For any constant-precision CoT-augmented Elman RNN LM, there exists a weakly equivalent PFSA.*

²⁸Throughout the paper, we index vectors and matrices directly with set elements rather than integer values, meaning each index corresponds to exactly one element from the given set.

²⁹If any of the two arguments are empty, its corresponding component is also zero.

Proof. Let p^{CoT} be a CoT Heaviside Elman RNN LM over the alphabet Σ with the regular function $\phi: \Delta^* \rightarrow \Sigma^*$ for an Σ -augmented alphabet Δ . Let p be the underlying RNN LM over Δ^* , that is,

$$p^{\text{CoT}} = p \circ \phi^{-1} \quad \text{and} \quad p = p^{\text{CoT}} \circ \phi$$

by Def. 3.3. We want to show the existence of a possibly non-deterministic PFSA \mathcal{A}^{CoT} that is weakly equivalent to p^{CoT} . We write $p \cong p'$ to mean p is weakly equivalent to p' (cf. Def. 2.2).

By Svete and Cotterell (2023, Lem. 4.1), there exists a DPFSA $\mathcal{A} = (\Delta, Q, \delta, \lambda, \rho)$ over the augmented alphabet Δ weakly equivalent to p . This means that

$$\mathcal{A} \cong p \cong p^{\text{CoT}} \circ \phi \quad (62)$$

and thus

$$p^{\text{CoT}} \cong \mathcal{A} \circ \phi^{-1}. \quad (63)$$

Just like in the proof of Thm. 3.1, $\mathcal{A} \circ \phi^{-1}$ can be computed as a composition of (lifted) WFSTs that is then projected onto the input component, resulting in a PFSA weakly equivalent to p^{CoT} . This finishes the proof. ■

Theorem 4.3. *For any regular LM, there exists a weakly equivalent CoT-augmented constant-precision transformer LM.*

Before we proceed to the full proof of Thm. 4.3, we first show the following simple but useful lemma.

Lemma F.1. *Define the following linear transformation functions:*

$$V(\mathbf{x}) \stackrel{\text{def}}{=} \mathbf{0}, \quad (64a)$$

$$O(\mathbf{x}) \stackrel{\text{def}}{=} \mathbf{0} \quad (64b)$$

for all $\mathbf{x} \in \mathbb{R}^D$. A transformer layer \mathcal{L} with the parameters V and O and residual connections implements the identity function irrespective of the parameters Q, K , and f :

$$\mathcal{L}(\mathbf{x}) = \mathbf{x} \quad (65)$$

for all $\mathbf{x} \in \mathbb{R}^D$.

Proof. By definition of a transformer layer (cf. Eqs. (27), (28a) and (28b)), \mathcal{L} computes

$$\mathbf{a} = \text{Att}(\mathbf{q}, \mathbf{K}, \mathbf{V}) + \mathbf{x} = \sum_{n=1}^N s_n \mathbf{v}_n + \mathbf{x} = \sum_{n=1}^N s_n \mathbf{0} + \mathbf{x} = \mathbf{x} \quad (66a)$$

$$\mathcal{L}(\mathbf{x}) = O(\mathbf{a}) + \mathbf{a} = \mathbf{0} + \mathbf{a} = \mathbf{a} = \mathbf{x}. \quad (66b)$$

■

This allows us to show Thm. 4.3.

Proof. Let $\mathcal{A} = (\Sigma, Q, \lambda, \rho, \delta)$ be a PFSA. Like the CoT-augmented RNN in Thm. 4.2, the CoT-augmented transformer LM will work over the alphabet $\Delta \stackrel{\text{def}}{=} \underline{\Sigma}_\epsilon \times Q$. It will start by generating a random initial state according to the initial-state distribution λ upon reading the designated padding symbol (BOS, q_0). Then, at step t of generation, it will generate the next symbol–state pair by sampling from the next-transition distribution $p(q_t, y_t | q_{t-1})$ given the state stored in the previously generated output tuple.

More formally, define $\text{BOS}' \stackrel{\text{def}}{=} (\text{BOS}, q_0)$ for some arbitrary but fixed $q_0 \in Q$. Further, define the static representation function

$$\mathcal{R}(\underline{y}, q, t) \stackrel{\text{def}}{=} \llbracket \underline{y}, q \rrbracket \in \{0, 1\}^{|\underline{\Sigma}_\epsilon| |Q|} \quad (67)$$

as in Eq. (57). Let \mathcal{L} be the identity transformer layer from Lemma F.1 and \mathcal{T} the transformer with the single layer \mathcal{L} . Then, by Lemma F.1,

$$\mathcal{T}(\mathcal{R})((\text{BOS}, q_0), (y_1, q_1), \dots, (y_{t-1}, q_{t-1})) = \left(\llbracket \text{BOS}, q_0 \rrbracket^\top; \llbracket y_1, q_1 \rrbracket^\top; \dots; \llbracket y_{t-1}, q_{t-1} \rrbracket^\top \right) \quad (68)$$

for any t . Defining the transformation $F(\mathbf{x}) \stackrel{\text{def}}{=} \mathbf{x}$, it therefore holds that

$$\text{enc}((\text{BOS}, q_0), (y_1, q_1), \dots, (y_{t-1}, q_{t-1})) = F(\mathbf{x}_{t-1}^1) = \llbracket y_{t-1}, q_{t-1} \rrbracket. \quad (69)$$

In words, $\text{enc}(\mathbf{y}_{<t})$ contains the current symbol–state pair of the PFSA. The CoT-augmented transformer LM can therefore sample the next symbol–state pair from the conditional distribution defined by the PFSA by setting the values of the output matrix \mathbf{E} as in Thm. 4.1:

$$E_{(y', q'), (y, q)} \stackrel{\text{def}}{=} p(q', y' | q) \quad q, q' \in Q, \quad y, y' \in \Sigma \quad (70a)$$

$$E_{(\varepsilon, q'), (\text{BOS}, q_0)} \stackrel{\text{def}}{=} \lambda(q) \quad q \in Q \quad (70b)$$

$$E_{(\text{EOS}, q), (y, q)} \stackrel{\text{def}}{=} \rho(q) \quad q \in Q, y \in \Sigma_\varepsilon. \quad (70c)$$

Clearly, the identical conditional probabilities defined by the CoT-augmented transformer \mathcal{T} result in strings over $\Sigma_\varepsilon \times Q$ sampled with probabilities equal to the path probabilities from the PFSA, as in Thm. 4.1. Applying the transformation ϕ to the produced outputs will thus result in strings sampled from the PFSA, giving us a CoT-augmented transformer LM weakly equivalent to \mathcal{A} . Lastly, since all the representations in the model are position-invariant, the constructed transformer is of constant precision. ■

F.2 Probabilistic Turing Machine Language Models

RNN LMs with CoT reasoning. Next, we show that, for every PTM, its induced LM can be encoded in an unbounded precision CoT RNN LM.

Theorem 4.4. *For every LM induced by a non-deterministic probabilistic Turing machine, there exists a weakly equivalent CoT-augmented RNN LM without unbounded precision.*

Proof. We rely mainly on existing results by Nowak et al. (2023) but use the additional information afforded by the augmented output alphabet, resulting in a simpler construction and interpretation. By Nowak et al. (2023, Prop. 3.1 and Thm. 3.1), the LM families induced by PTMs and probabilistic 2PDA are weakly equivalent, so it suffices to show that a CoT RNN LM can encode any given 2PDA.

We first reiterate their main result together with its condition.

Definition F.2. *A 2PDA \mathcal{P} is called Σ -deterministic if, for any current state q any top symbol on its first stack γ and any output symbol from Σ_ε , there is at most one transition with non-zero weight.*

Theorem F.3. *Nowak et al. (2023, Thm 3.2) Every Σ -deterministic probabilistic 2PDA can be encoded in an RNN LM that can output empty tokens ε .*

Nowak et al. (2023, Thm. 3.2) shows that an RNN LM with output alphabet Σ_ε can simulate any probabilistic 2PDA over Σ if it is Σ -deterministic. Let \mathcal{P} be an arbitrary probabilistic 2PDA, with $\mathcal{P} = (Q, \Sigma, \Gamma, \delta, q_L, q_\varphi)$. Now define another 2PDA $\mathcal{P}' = (Q, \Delta, \Gamma, \delta', q_L, q_\varphi)$, which generates outputs over the extended alphabet $\Delta \stackrel{\text{def}}{=} Q \times \Gamma_\varepsilon^4 \times \Sigma_\varepsilon$. Define the transitions, add the following transitions to δ' in \mathcal{P}' :

$$\delta' \stackrel{\text{def}}{=} \left\{ q \xrightarrow[\gamma_2 \rightarrow \gamma_4]{\gamma, (q, \gamma_1, \gamma_2, \gamma_3, \gamma_4, y), \gamma_1 \rightarrow \gamma_3/w} q' \mid q \xrightarrow[\gamma_2 \rightarrow \gamma_4]{\gamma, y, \gamma_1 \rightarrow \gamma_3/w} q' \in \delta \right\} \quad (71)$$

The constructed \mathcal{P}' satisfies the required Σ -determinism condition, meaning that we can construct a weakly equivalent RNN LM with ε outputs. Finally, define the following regular function $\phi: \Delta_\varepsilon \rightarrow \Sigma_\varepsilon$ that removes all additional information post-hoc:

$$\phi(x) \stackrel{\text{def}}{=} \begin{cases} y & \text{if } x = (q, \gamma_1, \gamma_2, \gamma_3, \gamma_4, y) \\ \varepsilon & \text{if } x = \varepsilon \end{cases} \quad (72)$$

The RNN LM therefore directly simulates runs from the 2PDA \mathcal{P}' . Applying ϕ to the outputs of the RNN LM results in strings in Σ^* . Since the transformation groups all execution runs that output $\mathbf{y} \in \Sigma^*$ and the resulting CoT-augmented LM p^\bullet (cf. Def. 3.3) sums over them, we are left with a weakly equivalent LM. ■

Theorem 4.5. *For any PTM-induced LM, there exists a weakly equivalent unbounded-precision CoT-augmented Transformer LMs.*

Proof. Here, we adapt the construction by Pérez et al. (2021) to the decoder-only case with prompt length (denoted by n in Pérez et al. (2021)) zero since we only care about *generating* strings from the CoT-augmented transformer LM. Moreover, rather than implementing a deterministic transition function of a Turing machine (which is simulated by a particular layer of their transformer), we implement the probabilistic transition function of a rational-valued PTM in the *sampling* step of the transformer LM. As in our previous constructions, this step endows the model with non-determinism.

At a high level, we will construct a CoT-augmented transformer LM over Σ^* that will output symbols from the augmented alphabet $\Delta \stackrel{\text{def}}{=} Q \times \Gamma \times \overline{\Sigma}_\varepsilon \times A \times A$. Here, $A \stackrel{\text{def}}{=} \{-1, 0, 1\}$, where we will for conciseness identify the action LEFT with -1 and the action RIGHT with 1 . Note that we, like Pérez et al. (2021), *do not* allow the action NOOP in our construction, but we include the value 0 in the action set as it will be useful for the starting conditions of the constructed transformer. As explained later, such an alphabet Δ contains enough information for the transformer LM to be able to reconstruct the configuration of the PTM at every time step, and thus match its conditional probability distributions.

Notation. Let $\mathcal{M} = (Q, \Sigma, \Gamma, \delta_{\mathcal{M}}, q_t, q_\varphi)$ be a QPTM and p the LM over Σ^* it induces. In the following, we denote with $q_t \in Q$ the state of the PTM, with $v_t \in \Gamma$ the symbol written to the working tape, with a_t the action performed, with s_t the symbol read, and with y_t the symbol written to the output tape, all at time step t .

High-level idea of the construction. Before formally describing the components of the CoT-augmented transformer LM in separate lemmata, we give a high-level overview of the construction. The two-layer (single-head) transformer will, when computing $\text{enc}(\mathbf{y}_{<t}^\bullet)$ for $\mathbf{y}_{<t}^\bullet \in \Delta^*$, perform the following computations, very similar to those described by Pérez et al. (2021):

1. **Input representations** (\mathcal{R} ; Lemma F.5): The input representation function represents the input symbols $y^\bullet \in \Delta$ with a multi-hot encoding (one that contains an individual one-hot encoding of each of the components $q_t, v_{t-1}, y_{t-1}, a_{t-1}, a_{t-2}$) of the form^{30,31}

$$\mathbf{r}_{\Pi}(y^\bullet, t) = \begin{pmatrix} \llbracket q_t \rrbracket \\ \llbracket v_{t-1} \rrbracket \\ \llbracket a_{t-1} \rrbracket \\ a_{t-1}, a_{t-2} \\ 0, 0 \\ 0, \mathbf{0}_{|\Gamma|} \\ 1, t+1, \frac{1}{t+1}, \frac{1}{(t+1)^2} \end{pmatrix} \quad (73)$$

for $t \in \mathbb{N}_{\geq 0}$. The representations also include additional components (the “empty” zero values in the vector, explained below) used for processing and the positional encoding of the time step, $1, t+1, \frac{1}{t+1}, \frac{1}{(t+1)^2}$. This component is explained in more detail in Lemma F.5.

2. **Layer 1** (\mathcal{L}_1 ; Lemma F.6): The first layer uses the information about the actions performed at all previous time steps (contained in the input static representations of the CoT-augmented symbols) to

³⁰In the following, green color denotes the components that are added or computed by each of the described components.

³¹Notice that the symbol $y \in \Sigma$ is not part of the internal representation, since it is not required for the simulation of the PTM. It is only used to construct the final output string stored in the output tape.

compute the locations of \mathcal{M} 's head at each time step. This results in the internal representations of the form

$$\mathbf{x}_t^1 = \begin{pmatrix} \llbracket q_t \rrbracket \\ \llbracket v_{t-1} \rrbracket \\ \llbracket a_{t-1} \rrbracket \\ a_{t-1}, a_{t-2} \\ \frac{c(t)}{t+1}, \frac{c(t-1)}{t+1} \\ 0, \mathbf{0}_{|\Gamma|} \\ 1, t+1, \frac{1}{t+1}, \frac{1}{(t+1)^2} \end{pmatrix} \quad (74)$$

for $t \in \mathbb{N}_{\geq 0}$, where $c(t)$ denotes the position of \mathcal{M} 's head at time t . This component is identical to the *second* layer of the transformer from Pérez et al. (2021) and is explained in more detail in Lemma F.6.

3. **Layer 2** (\mathcal{L}_2 ; Lemma F.7): Uses the \mathcal{L}_1 -computed information about the head locations at each time step to (almost) compute s_t —the symbol read by the head of the PTM at time step t . This results in the internal representations of the form

$$\mathbf{x}_t^2 = \begin{pmatrix} \llbracket q_t \rrbracket \\ \llbracket v_{t-1} \rrbracket \\ \llbracket a_{t-1} \rrbracket \\ a_{t-1}, a_{t-2} \\ \frac{c(t)}{t+1}, \frac{c(t-1)}{t+1} \\ \ell(t), \llbracket v_{\ell(t)} \rrbracket \\ 1, t+1, \frac{1}{t+1}, \frac{1}{(t+1)^2} \end{pmatrix} \quad (75)$$

for $t \in \mathbb{N}_{\geq 0}$, where $\ell(t)$ denotes the last time step when \mathcal{M} 's head wrote to $c(t)$. This component is analogous to the *third* layer of the transformer from Pérez et al. (2021) and is explained in more detail in Lemma F.7.

4. **Output function** (F ; Lemma F.8): The function F uses the information computed by the two layers of the transformer to compute the one-hot encoding of the current configuration of the PTM. In particular, this includes \mathcal{M} 's current state (q_t), the symbol read by \mathcal{M} 's head (s_t), and, for reasons that we explain shortly, the action performed by \mathcal{M} at the previous time step (a_{t-1}). This results in the representation

$$\text{enc}(\mathbf{y}_{<t}^{\bullet}) = \llbracket q_t, s_t, a_{t-1} \rrbracket, \quad (76)$$

which is used for sampling the next transition of the PTM. This component resembles the output function of Pérez et al. (2021) and is explained in more detail in Lemma F.8.

5. **Sampling** (Lemma F.9): The representation $\llbracket q_t, s_t, a_{t-1} \rrbracket$ is used to index the output matrix \mathbf{E} that contains the conditional probabilities $p(\cdot \mid q_t, s_t)$. This can be used to sample the next augmented symbol $\llbracket q^{t+1}, v_t, \bar{y}_t, a_t, a_{t-1} \rrbracket$ (here, the last component, a_{t-1} , equals the “input” to the sampling step). This component is explained in more detail in Lemma F.9.

The idea of the construction—iteratively computing and storing the modifications to the PTM configuration in the generated string \mathbf{y}^{\bullet} —is therefore identical to Pérez et al.'s (2021) one. In particular, the two components that perform the bulk of the simulation—Layer 1 and Layer 2—are identical to the components from Pérez et al. (2021). We describe and show the correctness of the components in the lemmata in the rest of the section. The correctness of the construction follows from the correctness of the components.

Weak equivalence. To define a CoT-augmented LM with such a transformer, we can define a transducer that projects the outputs in $(Q \times \Gamma \times \bar{\Sigma}_\varepsilon \times A \times A)^*$ onto Σ^* in the standard way by retaining only the outputs in $\bar{\Sigma}_\varepsilon$. This collapses all the executions of the transformer LM (which simulates the executions of the PTM with the same probabilities, as per Lemma F.9) yielding the same string in Σ^* , resulting in a CoT-augmented transformer LM weakly equivalent to p .

Precision. The constructed transformer computes and stores position-dependent values at several points during the computation. The precision required for the representation of these values grows logarithmically with the number of computational steps. However, since the number of computational steps performed by a PTM generating a string is potentially unbounded in the length of the string (Hopcroft et al., 2001, p. 339), the transformer’s precision is subsequently unbounded as well (Nowak et al., 2023). ■

We begin with a general lemma about the disjunction of one-hot encodings.

Lemma F.4. Let $\mathcal{S}_1, \dots, \mathcal{S}_n$ be finite sets and let $\llbracket s_1, \dots, s_n \rrbracket \in \{0, 1\}^{|\mathcal{S}_1| \cdots |\mathcal{S}_n|}$ denote the one-hot encoding of the tuple $(s_1, \dots, s_n) \in \mathcal{S}_1 \times \cdots \times \mathcal{S}_n$. Define the matrices

$$\mathbf{W}^{\mathcal{S}_i} \in \{0, 1\}^{|\mathcal{S}_i| \times |\mathcal{S}_1| \cdots |\mathcal{S}_n|} \quad (77)$$

element-wise as

$$\mathbf{W}_{s, (s_1, \dots, s_{i-1}, s, s_{i+1}, \dots, s_n)}^{\mathcal{S}_i} \stackrel{\text{def}}{=} 1 \quad \text{for all } s_j \in \mathcal{S}_j, j = 1, \dots, n, j \neq i. \quad (78)$$

Then, it holds that

$$\mathbf{W}^{\mathcal{S}_i} \llbracket s_1, \dots, s_n \rrbracket = \llbracket s_i \rrbracket. \quad (79)$$

Proof. Eq. (78) can equivalently be written as

$$\mathbf{W}_{s, (s_1, \dots, s_{i-1}, s, s_{i+1}, \dots, s_n)}^{\mathcal{S}_i} \stackrel{\text{def}}{=} \mathbb{1} \{s_i = s\}. \quad (80)$$

Indexing the elements of $\mathbf{W}^{\mathcal{S}_i} \llbracket s_1, \dots, s_n \rrbracket$ directly with $s \in \mathcal{S}_i$, we then compute

$$\left(\mathbf{W}^{\mathcal{S}_i} \llbracket s_1, \dots, s_n \rrbracket \right)_s = \sum_{\substack{s'_j \in \mathcal{S}_j \\ j=1, \dots, n}} \mathbf{W}_{s, (s'_1, \dots, s'_n)}^{\mathcal{S}_i} \llbracket s_1, \dots, s_n \rrbracket_{s'_1, \dots, s'_n} \quad (81a)$$

$$= \sum_{\substack{s'_j \in \mathcal{S}_j \\ j=1, \dots, n}} \mathbb{1} \{s'_i = s\} \llbracket s_1, \dots, s_n \rrbracket_{s'_1, \dots, s'_n} \quad (81b)$$

$$= \sum_{\substack{s'_j \in \mathcal{S}_j \\ j=1, \dots, n}} \mathbb{1} \{s'_i = s\} \mathbb{1} \{s'_1 = s_1, \dots, s'_n = s_n\} \quad (81c)$$

$$= \sum_{s'_i \in \mathcal{S}_i} \mathbb{1} \{s'_i = s\} \underbrace{\sum_{\substack{s'_j \in \mathcal{S}_j \\ j=1, \dots, n, j \neq i}} \mathbb{1} \{s'_1 = s_1, \dots, s'_n = s_n\}}_{=1} \quad (81d)$$

$$= \sum_{s'_i \in \mathcal{S}_i} \mathbb{1} \{s'_i = s\}, \quad (81e)$$

which is the definition of the elements of $\llbracket s_i \rrbracket$. The equality in Eq. (81d) follows from the fact that the summand is non-zero exactly when $s'_j = s_j$ for all $j \neq i$. ■

Lemma F.5 (Input representations). Define the following static representation function of the CoT-augmented symbols $y^{\bullet} \stackrel{\text{def}}{=} (q, v, \bar{y}, a', a) \in \Delta$:

$$\mathbf{r}_{\Pi} \left(y^{\bullet}, t \right) \stackrel{\text{def}}{=} \begin{pmatrix} \mathbf{W} \llbracket q, v, \bar{y}, a', a \rrbracket \\ 0, 0 \\ 0, \mathbf{0}_{|\Gamma|} \\ 1, t + 1, \frac{1}{t+1}, \frac{1}{(t+1)^2} \end{pmatrix} \in \mathbb{R}^D \quad (82)$$

where $\llbracket q, v, \bar{y}, a', a \rrbracket \in \{0, 1\}^{|\mathcal{Q}| |\Gamma| \lceil \bar{\Sigma}_\varepsilon \rceil |A| |A|}$ denotes the one-hot encoding of the tuple (q, v, \bar{y}, a', a) ,

$$\mathbf{W} \stackrel{\text{def}}{=} (\mathbf{W}^{\mathcal{Q}}; \mathbf{W}^{\Gamma}; \mathbf{W}^A; \mathbf{A}'; \mathbf{A}) \in \mathbb{R}^{D \times (|\mathcal{Q}| |\Gamma| \lceil \bar{\Sigma}_\varepsilon \rceil |A| |A|)} \quad (83)$$

for matrices $\mathbf{W}^Q, \mathbf{W}^\Gamma, \mathbf{W}^A$ from Lemma F.4, $\mathbf{A}', \mathbf{A} \in \mathbb{R}^{1 \times D}$ defined as³²

$$\mathbf{A}'_{1,(q,v,\bar{y},a',a)} \stackrel{\text{def}}{=} a', \quad q \in Q, v \in \Gamma, \bar{y} \in \bar{\Sigma}_\varepsilon, a, a' \in A \quad (84a)$$

$$\mathbf{A}_{1,(q,v,\bar{y},a',a)} \stackrel{\text{def}}{=} a, \quad q \in Q, v \in \Gamma, \bar{y} \in \bar{\Sigma}_\varepsilon, a, a' \in A, \quad (84b)$$

and $D \stackrel{\text{def}}{=} [|Q| + |\Gamma| + |A| + 2] + 2 + [1 + |\Gamma|] + 4$. Then, it holds that

$$\mathbf{r}_\Pi(y^{\text{BOS}}, t) = \begin{pmatrix} \llbracket q_t \rrbracket \\ \llbracket v \rrbracket \\ \llbracket a \rrbracket \\ a', a \\ 0, 0 \\ 0, \mathbf{0}_{|\Gamma|} \\ 1, t + 1, \frac{1}{t+1}, \frac{1}{(t+1)^2} \end{pmatrix} \quad (85)$$

Proof. The first part (computing the one-hot encodings of \sim_δ, v , and a') follows from the construction of the matrices $\mathbf{W}^Q, \mathbf{W}^\Gamma$, and \mathbf{W}^A from Lemma F.4. The second part (computing the values a' and a with the matrices \mathbf{A}' and \mathbf{A}) follows from the definition of the matrices \mathbf{A} and \mathbf{A}' : By construction, the values of a' and a will be copied from the one-hot encodings into the resulting entry of the vector. ■

We now describe the two layers of the transformer that compute the quantities required to be able to determine the current configuration of the PTM at each time step. Here, we heavily rely on the construction by Pérez et al. (2021).

Let $\mathcal{M} = (Q, \Sigma, \Gamma, \delta_{\mathcal{M}}, q_i, q_\varphi)$ be a rationally-weighted PTM. Let $t \in \mathbb{N}$ and define

$$c(t) \stackrel{\text{def}}{=} \sum_{j=0}^{t-1} a_j \quad (86)$$

as the position of the head of the PTM at time t . Further, define the set

$$L(t) \stackrel{\text{def}}{=} \{j \mid c(j) = c(t), j \in [t - 1]\}. \quad (87)$$

$L(t)$ contains the time steps at which the PTM \mathcal{M} so far visited (and thus wrote to) the tape cell read by \mathcal{M} at time t . Then, define $\ell(t)$ as

$$\ell(t) \stackrel{\text{def}}{=} \begin{cases} \max L(t) & \text{if } |L(t)| > 0 \\ t & \text{otherwise.} \end{cases} \quad (88)$$

In words, $\ell(t)$ denotes the time step at which the PTM \mathcal{M} last visited (and thus wrote on) the tape cell read by \mathcal{M} at time t if this cell was visited yet. Otherwise, $\ell(t)$ equals t .

Now, define the BOS symbol over the augmented alphabet Γ as

$$\text{BOS}^{\text{BOS}} \stackrel{\text{def}}{=} (q_0, \perp, \text{BOS}, 0, 0) \in Q \times \Gamma \times \bar{\Sigma}_\varepsilon \times A \times A. \quad (89)$$

Let $q_i, q_1, \dots, q_T, \perp, v_1, \dots, v_T, \text{BOS}, y_1, \dots, y_T$, and a_0, a_1, \dots, a_T be the sequences of states visited, symbols written on the processing tape, symbols written on the output tape, and actions performed by \mathcal{M} in the first T steps. Define $a_{-1} = a_0 \stackrel{\text{def}}{=} 0$,

$$\mathbf{x}_t^0 \stackrel{\text{def}}{=} \begin{cases} \mathbf{r}_\Pi(\text{BOS}^{\text{BOS}}, t) & \text{if } t = 0 \\ \mathbf{r}_\Pi((q_t, v_t, y_t, a_{t-1}, a_{t-2}), t) & \text{otherwise,} \end{cases} \quad (90)$$

³²Recall that we can identify the actions of the PTM with the integers $-1, 0$, and 1 .

and

$$\mathbf{X}^0 \stackrel{\text{def}}{=} (\mathbf{x}_0^0; \mathbf{x}_1^0; \cdots; \mathbf{x}_T^0). \quad (91)$$

These will be the inputs to the transformer layers and thus,³³

$$\mathcal{R}(\mathbf{y}_{<t}) \stackrel{\text{def}}{=} (\mathbf{x}_0^0; \mathbf{x}_1^0; \cdots; \mathbf{x}_{t-1}^0). \quad (92)$$

With this, we can describe and prove the correctness of the first layer of the transformer we are building.

Lemma F.6 (Layer 1). *There exists a transformer layer \mathcal{L}_1 that, given the inputs \mathbf{X}^0 , computes the values $\frac{c(t)}{t+1}$ and $\frac{c(t-1)}{t+1}$ for all $t = 1, \dots, T$. More precisely, denoting*

$$\mathbf{X}^1 \stackrel{\text{def}}{=} (\mathbf{x}_0^1; \mathbf{x}_1^1; \cdots; \mathbf{x}_T^1) = \mathcal{L}_1(\mathbf{X}^0), \quad (93)$$

it holds that \mathbf{x}_t^1 contains the entries containing the values $\frac{c(t)}{t+1}$ and $\frac{c(t-1)}{t+1}$:

$$\mathbf{x}_t^1 = \begin{pmatrix} \llbracket q_t \rrbracket \\ \llbracket v_{t-1} \rrbracket \\ \llbracket a_{t-1} \rrbracket \\ a_{t-1}, a_{t-2} \\ \frac{c(t)}{t+1}, \frac{c(t-1)}{t+1} \\ 0, \mathbf{0}_{|\Gamma|} \\ 1, t+1, \frac{1}{t+1}, \frac{1}{(t+1)^2} \end{pmatrix} \quad (94)$$

Proof. This follows from Pérez et al. (2021, Lemma 9), which we summarize here.³⁴ Concretely, \mathcal{L}_1 is implemented by a transformer layer with trivial (zero-valued) query and key transformations and a value transformation V that copies the values of the actions a_{t-1} and a_{t-2} to the entry that will hold the values $\frac{c(t)}{t+1}$ and $\frac{c(t-1)}{t+1}$. Since the current head location $c(t)$ is simply the sum of those values (cf. Eq. (86)), attending to all previous positions results in Eq. (94).

Formally, we define

$$\mathbf{Q}^1 \stackrel{\text{def}}{=} \mathbf{0}_{D \times D}, \quad \mathbf{K}^1 \stackrel{\text{def}}{=} \mathbf{0}_{D \times D}, \quad (95a)$$

$$f(\mathbf{q}, \mathbf{k}) \stackrel{\text{def}}{=} \langle \mathbf{q}, \mathbf{k} \rangle \quad (95b)$$

$$\mathbf{V}_{n',:}^1 = \mathbf{V}_{n,:}^1 = \begin{pmatrix} \mathbf{0}_{|Q|} \\ \mathbf{0}_{|\Gamma|} \\ \mathbf{0}_{|A|} \\ 0, 0 \\ 1, 1 \\ 0, \mathbf{0}_{|\Gamma|} \\ 0, 0, 0, 0 \end{pmatrix}^\top \quad (95c)$$

$$\mathbf{O}^1 = \mathbf{0}_{D \times D} \quad (95d)$$

Here, n' and n refer to the indices of the rows at which the values $\frac{c(t)}{t+1}$ and $\frac{c(t-1)}{t+1}$ will be stored (that is, the rows below the values of a_{t-1} and a_{t-2}). All other rows of \mathbf{V}^1 are zero.

Then, for $t \in \mathbb{N}_{\geq 0}$, it holds that

$$f(\mathbf{q}_t, \mathbf{k}_j) = f(Q(\mathbf{x}_t^0), K(\mathbf{x}_j^0)) = f(\mathbf{Q}^1 \mathbf{x}_t^0, \mathbf{K}^1 \mathbf{x}_j^0) = \langle \mathbf{0}, \mathbf{0} \rangle = 0 \quad (96)$$

³³Note that since the transformer over the augmented alphabet is deterministic, this representation is unique.

³⁴More precisely, since their construction stores the actions a_t and a_{t-1} (the former is possible because the action is not sampled but deterministically computed based on the configuration of the PTM), their layer computes the values $\frac{c(t)}{t+1}$ and $\frac{c(t-1)}{t+1}$. As we show later in Lemma F.7, this does not affect the correctness of the construction.

for all $j \leq t$, resulting in $\mathbf{s}_t = \text{hardmax}(\mathbf{0}) = \frac{1}{t+1} \mathbf{1}_t$, and

$$V(\mathbf{x}_j^0) = \begin{pmatrix} \mathbf{0}_{|Q|} \\ \mathbf{0}_{|\Gamma|} \\ \mathbf{0}_{|A|} \\ 0, 0 \\ a_{j-1}, a_{j-2} \\ 0, \mathbf{0}_{|\Gamma|} \\ 0, 0, 0, 0 \end{pmatrix}. \quad (97)$$

This results in

$$\text{Att}(\mathbf{q}_t, \mathbf{K}_t, \mathbf{V}_t) = \sum_{j=0}^t s_j V(\mathbf{x}_j^0) \quad (98a)$$

$$= \sum_{j=0}^t \frac{1}{t+1} V(\mathbf{x}_j^0) \quad (98b)$$

$$= \frac{1}{t+1} \sum_{j=0}^t V(\mathbf{x}_j^0) \quad (98c)$$

$$= \frac{1}{t+1} \sum_{j=0}^t \begin{pmatrix} \mathbf{0}_{|Q|} \\ \mathbf{0}_{|\Gamma|} \\ \mathbf{0}_{|A|} \\ 0, 0 \\ a_{j-1}, a_{j-2} \\ 0, \mathbf{0}_{|\Gamma|} \\ 0, 0, 0, 0 \end{pmatrix} \quad (98d)$$

$$= \frac{1}{t+1} \begin{pmatrix} \mathbf{0}_{|Q|} \\ \mathbf{0}_{|\Gamma|} \\ \mathbf{0}_{|A|} \\ 0, 0 \\ c(t), c(t-1) \\ 0, \mathbf{0}_{|\Gamma|} \\ 0, 0, 0, 0 \end{pmatrix} \quad (98e)$$

$$= \begin{pmatrix} \mathbf{0}_{|Q|} \\ \mathbf{0}_{|\Gamma|} \\ \mathbf{0}_{|A|} \\ 0, 0 \\ \frac{c(t)}{t+1}, \frac{c(t-1)}{t+1} \\ 0, \mathbf{0}_{|\Gamma|} \\ 0, 0, 0, 0 \end{pmatrix}. \quad (98f)$$

Furthermore, we have that

$$\mathbf{a}_t = \text{Att}(\mathbf{q}_t, \mathbf{K}_t, \mathbf{V}_t) + \mathbf{x}_t^0 = \begin{pmatrix} \mathbf{0}_{|Q|} \\ \mathbf{0}_{|\Gamma|} \\ \mathbf{0}_{|A|} \\ 0, 0 \\ \frac{c(t)}{t+1}, \frac{c(t-1)}{t+1} \\ 0, \mathbf{0}_{|\Gamma|} \\ 0, 0, 0, 0 \end{pmatrix} + \begin{pmatrix} \llbracket q_t \rrbracket \\ \llbracket v_{t-1} \rrbracket \\ \llbracket a_{t-1} \rrbracket \\ a_{t-1}, a_{t-2} \\ 0, 0 \\ 0, \mathbf{0}_{|\Gamma|} \\ 1, t+1, \frac{1}{t+1}, \frac{1}{(t+1)^2} \end{pmatrix} = \begin{pmatrix} \llbracket q_t \rrbracket \\ \llbracket v_{t-1} \rrbracket \\ \llbracket a_{t-1} \rrbracket \\ a_{t-1}, a_{t-2} \\ \frac{c(t)}{t+1}, \frac{c(t-1)}{t+1} \\ 0, \mathbf{0}_{|\Gamma|} \\ 1, t+1, \frac{1}{t+1}, \frac{1}{(t+1)^2} \end{pmatrix} \quad (99)$$

$$\mathbf{x}_t^1 = O(\mathbf{a}_t) + \mathbf{a}_t = \mathbf{0}_D + \begin{pmatrix} \llbracket q_t \rrbracket \\ \llbracket v_{t-1} \rrbracket \\ \llbracket a_{t-1} \rrbracket \\ a_{t-1}, a_{t-2} \\ \frac{c(t)}{t+1}, \frac{c(t-1)}{t+1} \\ 0, \mathbf{0}_{|\Gamma|} \\ 1, t+1, \frac{1}{t+1}, \frac{1}{(t+1)^2} \end{pmatrix} = \begin{pmatrix} \llbracket q_t \rrbracket \\ \llbracket v_{t-1} \rrbracket \\ \llbracket a_{t-1} \rrbracket \\ a_{t-1}, a_{t-2} \\ \frac{c(t)}{t+1}, \frac{c(t-1)}{t+1} \\ 0, \mathbf{0}_{|\Gamma|} \\ 1, t+1, \frac{1}{t+1}, \frac{1}{(t+1)^2} \end{pmatrix}, \quad (100)$$

which is what we needed to show. ■

Lemma F.7 (Layer 2). *There exists a transformer layer \mathcal{L}_2 that, given the outputs \mathbf{X}^1 of \mathcal{L}_1 from Lemma F.6, computes the values $\ell(t)$ and $\llbracket v_{\ell(t)} \rrbracket$ for all $t = 1, \dots, T$. More precisely, denoting*

$$\mathbf{X}^2 \stackrel{\text{def}}{=} (\mathbf{x}_0^2; \mathbf{x}_1^2; \dots; \mathbf{x}_T^2) = \mathcal{L}_2(\mathbf{X}^1), \quad (101)$$

it holds that \mathbf{x}_t^2 contains the entries containing the values $\ell(t) + 1$ and $\llbracket v_{\ell(t)} \rrbracket$:

$$\mathbf{x}_t^2 = \begin{pmatrix} \llbracket q_t \rrbracket \\ \llbracket v_{t-1} \rrbracket \\ \llbracket a_{t-1} \rrbracket \\ a_{t-1}, a_{t-2} \\ \frac{c(t)}{t+1}, \frac{c(t-1)}{t+1} \\ \ell(t) + 1, \llbracket v_{\ell(t)} \rrbracket \\ 1, t+1, \frac{1}{t+1}, \frac{1}{(t+1)^2} \end{pmatrix}. \quad (102)$$

Proof. This follows from Pérez et al. (2021, Lemma 10). The idea of the construction is for the self-attention mechanism at time t to attend to exactly the entry from time step $\ell(t) + 1$ (i.e., to compute the query and key vectors such that $\arg\max(\mathbf{s}_t) = \ell(t) + 1$) since that entry will contain the information about the symbol written at the time step before—at $\ell(t)$. Then, the values $\ell(t)$ and $v_{\ell(t)}$ are obtained by copying the corresponding values of the positional encoding and the written symbol from that time step.

Formally, we define

$$\mathbf{Q}^2 \stackrel{\text{def}}{=} \begin{pmatrix} \dots & \underbrace{\frac{c(t)}{t+1}} & \dots & \underbrace{\frac{1}{t+1}} & \underbrace{\frac{1}{(t+1)^2}} & \dots \\ & 1 & & 1 & & \\ & & & & \frac{1}{3} & \end{pmatrix} \quad (103a)$$

$$\mathbf{K}^2 \stackrel{\text{def}}{=} \begin{pmatrix} \dots & \underbrace{\frac{c(t-1)}{t+1}} & \dots & \underbrace{\frac{1}{t+1}} & \underbrace{\frac{1}{(t+1)^2}} & \dots \\ & -1 & & 1 & & \\ & & & & \frac{1}{3} & \end{pmatrix} \quad (103b)$$

$$f(\mathbf{q}, \mathbf{k}) \stackrel{\text{def}}{=} -|\langle \mathbf{q}, \mathbf{k} \rangle| \quad (103c)$$

$$\mathbf{V}^2 \stackrel{\text{def}}{=} \begin{pmatrix} \dots & \underbrace{\llbracket v_{t-1} \rrbracket} & \dots & \underbrace{t+1} & \dots \\ & \mathbf{I}_{|\Gamma|} & & 1 & \\ & & & & \ell(t) + 1 \\ & & & & \llbracket v_{\ell(t)} \rrbracket \\ & & & & \vdots \end{pmatrix} \quad (103d)$$

$$\mathbf{O}^1 = \mathbf{0}_{D \times D} \quad (103e)$$

Pérez et al. (2021, Lemma 10) show that, given the parameters above, the output of the scoring function f is maximized at the entry $\ell(t) \in \{0, \dots, t\}$.³⁵ Since \mathbf{V}^2 copies the second value from the positional encoding of the symbol at time step $\ell(t)$, which is $\ell(t) + 1$, this entry appears in the output of the attention mechanism. Furthermore, \mathbf{V}^2 also copies the value $\llbracket v_{\ell(t)} \rrbracket$ from the same entry. This, together with the zero-valued function O and residual connections, results in Eq. (102). ■

Lemma F.8 (Correctness of the output function). *There exists an MLP F that, given the outputs \mathbf{X}^2 of \mathcal{L}_2 , computes the one-hot encoding of the current configuration of the PTM. More concretely, it holds that*

$$F(\mathbf{x}_t^2) = \llbracket q_t, s_t, a_{t-1} \rrbracket. \quad (104)$$

Proof. Here, we define a function F similar to that of Pérez et al. (2021, Lemma 11), but with an additional layer that handles the addition of the \perp symbol not handled by Pérez et al. (2021). The logic nonetheless remains the same: F receives the output of \mathcal{L}_2 of the form

$$\mathbf{x}_t^2 = \begin{pmatrix} \llbracket q_t \rrbracket \\ \llbracket v_{t-1} \rrbracket \\ \llbracket a_{t-1} \rrbracket \\ a_{t-1}, a_{t-2} \\ \frac{c(t)}{t+1}, \frac{c(t-1)}{t+1} \\ \ell(t) + 1, \llbracket v_{\ell(t)} \rrbracket \\ 1, t + 1, \frac{1}{t+1}, \frac{1}{(t+1)^2} \end{pmatrix}. \quad (105)$$

and (1) copies the value of q_t , (2) copies the value of a_{t-1} (3) compares the value of $\ell(t)$ to t to determine whether $s_t = \perp$ or $s_t = v_{\ell(t)}$, (4) compares the value of t to 0 to determine whether $s_t = \perp$ or $s_t = v_{\ell(t)}$.

³⁵More precisely, their construction results in the maximum being at $\ell(t + 1)$, since layer 1 in their construction, in contrast to ours, computes the values $\frac{c(t+1)}{t+1}$ and $\frac{c(t)}{t+1}$, shifting all computations by one step. This is the consequence of the aforementioned difference between their deterministic and our probabilistic framework.

$$\mathbf{W}^2 = \begin{pmatrix} \overbrace{1:C_1} & \overbrace{C_1:C_2} & \overbrace{C_2:C_3} & \overbrace{C_3:C_4} & \overbrace{C_4:C_5} & \overbrace{C_5+1} & \overbrace{C_5+2} \\ \mathbf{I}_{|Q|} & & & & & & \\ & \mathbf{I}_{|\Gamma|} & & & & & \\ & & \mathbf{I}_{|A|} & & & & \\ & & & \mathbf{I}_{|\Gamma|} & & & \\ & & & & \mathbf{I}_{|\Gamma|} & & \\ & & & & & 1 & -1 \\ & & & & & & 1 \end{pmatrix} \quad (112) \quad \mathbf{b}^2 = \begin{pmatrix} \mathbf{0}_{|Q|} \\ \mathbf{0}_{|\Gamma|} \\ \mathbf{0}_{|A|} \\ \mathbf{0}_{|\Gamma|} \\ \mathbf{0}_{|\Gamma|} \\ 0 \\ 0 \end{pmatrix} \quad (113)$$

$$\mathbf{W}^3 = \begin{pmatrix} \overbrace{1:C_1} & \overbrace{C_1:C_2} & \overbrace{C_2:C_3} & \overbrace{C_3:C_4} & \overbrace{C_4:C_5} & \overbrace{C_5+1} & \overbrace{C_5+2} \\ \mathbf{I}_{|Q|} & & & & & & \\ & \mathbf{I}_{|\Gamma|} & & & & -\mathbf{1}_{|\Gamma|} & -\mathbf{1}_{|\Gamma|} \\ & & \mathbf{I}_{|A|} & & & \mathbf{1}_{|\Gamma|} & \\ & & & \mathbf{I}_{|\Gamma|} & & & \\ & & & & \mathbf{I}_{|\Gamma|} & & \mathbf{1}_{|\Gamma|} \end{pmatrix} \quad (114) \quad \mathbf{b}^3 = \begin{pmatrix} \mathbf{0}_{|Q|} \\ \mathbf{0}_{|\Gamma|} \\ \mathbf{0}_{|A|} \\ -\mathbf{1}_{|\Gamma|} \\ -\mathbf{1}_{|\Gamma|} \end{pmatrix} \quad (115)$$

$$\mathbf{W}^4 = \begin{pmatrix} \overbrace{1:C_1} & \overbrace{C_1:C_2} & \overbrace{C_2:C_3} & \overbrace{C_3:C_4} & \overbrace{C_4:C_5} \\ \mathbf{I}_{|Q|} & & & & \\ & \mathbf{I}_{|\Gamma|} & & & \\ & & \mathbf{I}_{|A|} & & \\ & & & \mathbf{I}_{|\Gamma|} & \\ & & & & \mathbf{I}_{|\Gamma|} \end{pmatrix} \quad (116) \quad \mathbf{b}^4 = \begin{pmatrix} \mathbf{0}_{|Q|} \\ \mathbf{0}_{|\Gamma|} \\ \mathbf{0}_{|A|} \end{pmatrix} \quad (117)$$

Lastly, we set \mathbf{W}^5 and \mathbf{b}^5 such that the the function $\mathbf{x} \mapsto \text{ReLU}(\mathbf{W}^5 \mathbf{x} + \mathbf{b}^5)$ computes the one-hot encoding of the three input one-hot encodings $\llbracket q_t \rrbracket$, $\llbracket s_t \rrbracket$, $\llbracket a_{t-1} \rrbracket$ by implementing the logic AND operations, as in [Svete and Cotterell \(2024, Lemma B.1\)](#).

By construction, we get that

$$\mathbf{y}^1 = \text{ReLU}(\mathbf{W}^1 \mathbf{x}_t^2 + \mathbf{b}^1) \quad (118a)$$

$$= \begin{pmatrix} \llbracket q_t \rrbracket \\ \llbracket v_{t-1} \rrbracket \\ \llbracket a_{t-1} \rrbracket \\ \llbracket \perp \rrbracket \\ \llbracket \perp \rrbracket \\ \text{ReLU}(\ell(t) + 1 + 1 - (t + 1)) \\ \text{ReLU}(2 - (t + 1)) \end{pmatrix} \quad (118b)$$

$$= \begin{pmatrix} \llbracket q_t \rrbracket \\ \llbracket v_{t-1} \rrbracket \\ \llbracket a_{t-1} \rrbracket \\ \llbracket \perp \rrbracket \\ \llbracket \perp \rrbracket \\ \text{ReLU}(\ell(t) + 1 - t) \\ \text{ReLU}(1 - t) \end{pmatrix} \quad (118c)$$

$$= \begin{pmatrix} \llbracket q_t \rrbracket \\ \llbracket v_{t-1} \rrbracket \\ \llbracket a_{t-1} \rrbracket \\ \llbracket \perp \rrbracket \\ \llbracket \perp \rrbracket \\ \mathbb{1} \{ \ell(t) \geq t \} \\ \mathbb{1} \{ t < 1 \} \end{pmatrix} \quad (118d)$$

$$= \begin{pmatrix} \llbracket q_t \rrbracket \\ \llbracket v_{t-1} \rrbracket \\ \llbracket a_{t-1} \rrbracket \\ \llbracket \perp \rrbracket \\ \llbracket \perp \rrbracket \\ \mathbb{1} \{ \ell(t) = t \} \\ \mathbb{1} \{ t = 0 \} \end{pmatrix} \quad (118e)$$

$$\mathbf{y}^2 = \text{ReLU}(\mathbf{W}^2 \mathbf{y}^1 + \mathbf{b}^2) \quad (119a)$$

$$= \begin{pmatrix} \llbracket q_t \rrbracket \\ \llbracket v_{t-1} \rrbracket \\ \llbracket a_{t-1} \rrbracket \\ \llbracket \perp \rrbracket \\ \llbracket \perp \rrbracket \\ \text{ReLU}(\mathbb{1} \{ \ell(t) = t \} - \mathbb{1} \{ t = 0 \}) \\ \mathbb{1} \{ t = 0 \} \end{pmatrix} \quad (119b)$$

$$\quad (119c)$$

$$= \begin{pmatrix} \llbracket q_t \rrbracket \\ \llbracket v_{t-1} \rrbracket \\ \llbracket a_{t-1} \rrbracket \\ \llbracket \perp \rrbracket \\ \llbracket \perp \rrbracket \\ \mathbb{1} \{ \mathbb{1} \{ \ell(t) = t \} \&\mathbb{1} \{ t > 0 \} \} \\ \mathbb{1} \{ t = 0 \} \end{pmatrix} \quad (119d)$$

$$\mathbf{y}^3 = \text{ReLU}(\mathbf{W}^3 \mathbf{y}^2 + \mathbf{b}^3) \quad (120a)$$

$$= \begin{pmatrix} \llbracket q_t \rrbracket \\ \text{ReLU}(\llbracket v_{t-1} \rrbracket - (\mathbb{1}\{\ell(t) = t \wedge t > 0\} + \mathbb{1}\{t = 0\}) \mathbf{1}_{|\Gamma|}) \\ \text{ReLU}(\llbracket \sqcup \rrbracket + \mathbb{1}\{\ell(t) = t \wedge t > 0\} \mathbf{1}_{|\Gamma|} - \mathbf{1}_{|\Gamma|}) \\ \text{ReLU}(\llbracket \perp \rrbracket + \mathbb{1}\{t = 0\} \mathbf{1}_{|\Gamma|} - \mathbf{1}_{|\Gamma|}) \end{pmatrix} \quad (120b)$$

$$= \begin{pmatrix} \llbracket q_t \rrbracket \\ \mathbb{1}\{\neg(\ell(t) = t \wedge t > 0) \wedge \neg(t = 0)\} \llbracket v_{t-1} \rrbracket \\ \llbracket a_{t-1} \rrbracket \\ \mathbb{1}\{\ell(t) = t \wedge t > 0\} \llbracket \sqcup \rrbracket \\ \mathbb{1}\{t = 0\} \llbracket \perp \rrbracket \end{pmatrix} \quad (120c)$$

$$= \begin{pmatrix} \llbracket q_t \rrbracket \\ \mathbb{1}\{(\ell(t) < t \vee t = 0) \wedge t > 0\} \llbracket v_{t-1} \rrbracket \\ \llbracket a_{t-1} \rrbracket \\ \mathbb{1}\{\ell(t) = t \wedge t > 0\} \llbracket \sqcup \rrbracket \\ \mathbb{1}\{t = 0\} \llbracket \perp \rrbracket \end{pmatrix} \quad (120d)$$

$$= \begin{pmatrix} \llbracket q_t \rrbracket \\ \mathbb{1}\{\ell(t) < t \wedge t > 0\} \llbracket v_{t-1} \rrbracket \\ \llbracket a_{t-1} \rrbracket \\ \mathbb{1}\{\ell(t) = t \wedge t > 0\} \llbracket \sqcup \rrbracket \\ \mathbb{1}\{t = 0\} \llbracket \perp \rrbracket \end{pmatrix} \quad (120e)$$

$$\mathbf{y}^4 = \text{ReLU}(\mathbf{W}^4 \mathbf{y}^3 + \mathbf{b}^4) \quad (121a)$$

$$= \begin{pmatrix} \llbracket q_t \rrbracket \\ \mathbb{1}\{\ell(t) < t \wedge t > 0\} \llbracket v_{t-1} \rrbracket + \mathbb{1}\{\ell(t) = t \wedge t > 0\} \llbracket \sqcup \rrbracket + \mathbb{1}\{t = 0\} \llbracket \perp \rrbracket \\ \llbracket a_{t-1} \rrbracket \end{pmatrix} \quad (121b)$$

$$\stackrel{\text{def}}{=} \begin{pmatrix} \llbracket q_t \rrbracket \\ \llbracket w_t \rrbracket \\ \llbracket a_{t-1} \rrbracket \end{pmatrix} \quad (121c)$$

Since the three logical expressions in Eq. (121b) are complementary, it holds that the second component of \mathbf{y}^4 holds either the value of v_{t-1} , \sqcup , or \perp depending on the value of $\ell(t)$ and t . This is exactly the symbol that will be read by the PTM at time step t , i.e., s_t : If $\ell(t) < t$ (which means that cell $c(t)$ has been visited before), then $s_t = v_{t-1}$; if $\ell(t) = t$ and $t > 0$ (which means that cell $c(t)$ has not been visited before, and $t > 0$), then $s_t = \sqcup$; and if $t = 0$ (meaning that the PTM just started executing and is still reading the initial symbol \perp), then $s_t = \perp$. By the construction of \mathbf{W}^5 and \mathbf{b}^5 , we get that $\mathbf{y}^5 = \llbracket q_t, s_t, a_{t-1} \rrbracket$. ■

Lemma F.9 (Correctness of the sampling step). *Define the output matrix $\mathbf{E} \in \mathbb{R}^{|Q||\Gamma| \times |\bar{\Sigma}_\varepsilon| |A| |A| \times |Q||\Gamma| |A|}$ as*

$$E_{(q', v, \bar{y}, a', a), (q, s, a)} \stackrel{\text{def}}{=} \begin{cases} p(q', v, \bar{y}, a', a \mid q, s) & \text{if } \bar{y} \neq \text{EOS} \\ p(q_\varphi, v, \bar{y}, a', a \mid q, s) & \text{otherwise} \end{cases} \quad (122)$$

for $q, q' \in Q$, $v \in \Gamma$, $\bar{y} \in \bar{\Sigma}_\varepsilon$, $a \in A$, $a' \in \{-1, 1\}$, and $s \in \Gamma$. Then, it holds that

$$\mathbf{E} F(\mathbf{x}_t^2) = p(\cdot \mid q_t, s_t). \quad (123)$$

Proof. Follows directly from Lemma F.8 and the construction of \mathbf{E} : By Lemma F.8, we have that

$$F(\mathbf{x}_t^2) = \llbracket q_t, s_t, a_{t-1} \rrbracket. \quad (124)$$

By construction of \mathbf{E} , we have that

$$(\mathbf{E} F(\mathbf{x}_t^2))_{(q', v, \bar{y}, a', a)} = (\mathbf{E} \llbracket q_t, s_t, a_{t-1} \rrbracket)_{(q', v, \bar{y}, a', a)} \quad (125a)$$

$$= \mathbf{E}_{(q_t, s_t, a_{t-1}), (q', v, \bar{y}, a', a)} \quad (125b)$$

$$= \begin{cases} p(q', v, \bar{y}, a', a | q, s) & \text{if } \bar{y} \neq \text{EOS} \\ p(q_\varphi, v, \bar{y}, a', a | q, s) & \text{otherwise} \end{cases} \quad (125c)$$

■