A formal approach for understanding the behavior of constrained language models

M. Iturbide ITURBIDE@ORT.EDU.UY
S. Yovine YOVINE@ORT.EDU.UY
M. Carrasco CARRASCO_M@ORT.EDU.UY
Facultad de Ingeniería, Universidad ORT Uruguay, Montevideo, Uruguay

Editors: J. Moerman, S. Verwer

Abstract

We study language models under explicit formal constraints by treating common operations, such as prompting, masking, temperature scaling, and vocabulary translation, as composable transformations, in particular those that preserve the set of generable sequences and/or its probability distribution. We focus on constraints that disable tokens, potentially leading to dead-ends during generation. Moreover, we formalize bounded-length and safe decoding procedures. The theoretical results enable the algorithmic approximation of the probability of subsets (properties) of the generable language, as well as active learning of automata representations, such as DFA o Moore machines. We evaluate the approach to analyze the behavior of several state-of-the-art large language models.

Keywords: Languages models. Constrained decoding. Safe generation. Formal languages.

1. Introduction

Large language models (LLMs) are increasingly deployed in operational pipelines where their textual outputs must satisfy rigid formal specifications, such as structured serialization formats, syntax-valid source code, database queries, or domain-specific notations. In practical applications, ensuring compliance with these structural requirements is handled at decoding time. This is typically achieved through a variety of runtime interventions, including prompt engineering, context-dependent temperature scaling, and token-level vocabulary masking that prunes the model's output distribution dynamically at each step Willsey et al. (2024); Koo et al. (2024); Beurer-Kellner et al. (2024); Geng et al. (2023).

Despite the widespread adoption of these decoding-time constraints, there is a lack of rigorous mathematical characterizations regarding how they affect properties of the generated language, such as language equivalence, probabilistic termination, and the occurrence of execution errors. This paper proposes a formal framework to address these challenges. Within this paradigm, decoding-time interventions are modeled cleanly as composable mathematical transformations. This algebraic abstraction allows us to achieve two core theoretical and practical objectives: (1) define, transform, compute, and compare languages and probability distributions induced under distinct, composable constraints; and (2) characterize and remove generation failures caused by reaching dead-ends with zero termination probability.

Our framework occupies a distinct and complementary niche relative to existing ones. It diverges from grammar-guided decoding libraries (e.g., Willsey et al. (2024); Ugare et al. (2025)), which enforce constraints during inference, and from probability analysis tools

© M. Iturbide, S. Yovine & M. Carrasco.

(e.g., Suresh et al. (2026)), which are restricted to safety properties. It is also conceptually orthogonal to white-box symbolic execution models of transformer architectures, such as RASP Weiss et al. (2021), which constructively describe the formal languages a network can theoretically compute.

We evaluate our approach empirically using state-of-the-art open-weight models, such as the Qwen3 family. The experimental results validate that our framework effectively allows understanding the effect of formally characterized decoding constraints, via the quantification of probabilities of desired properties, and the extraction of descriptive finite-state abstractions.

2. Language Models

In abstract terms, a *language model* (LM) is a partial function $\mathcal{L} : \Sigma^* \to \mathfrak{R}(\Sigma_\top)$ where: Σ^* is the free monoid of finite *strings* over a finite *alphabet* Σ, concatenation is the monoid operation and the *empty* string $\lambda \in \Sigma^*$ is the identity element; $\Sigma_\top \triangleq \Sigma \cup \{\top\}$, where $\top \notin \Sigma$ is a special symbol which represents *termination*; $\mathfrak{R}(\Sigma_\top)$ is the set of total functions from $\Sigma_\top$ to $\mathbb{R} \cup \{-\infty\}$, called the *logits*. A *probabilistic* language model (PLM) is a partial function from Σ^* to $\mathfrak{D}(\Sigma_\top)$, where $\mathfrak{D}(\Sigma_\top) \subset \mathfrak{R}_{\geq 0}(\Sigma_\top)$ is the set of probability distributions over $\Sigma_\top$. To apply vector operations such as $\oplus$ (component-wise addition), $\odot$ (component-wise multiplication), $\bullet$ (dot product), and $\cdot$ (multiplication by a scalar), elements of $\mathfrak{R}(\Sigma_\top)$ will be regarded as vectors, assuming an arbitrary but fixed (e.g., lexicographic) ordering. We write $\mathcal{L}(u) = \bot$ to denote that $\mathcal{L}$ is *undefined* in $u \in \Sigma^*$.

We call *transformation* any mechanism that modifies the domain or the image of an LM. Here, we formalize some useful transformations, such as *prompting*, *masking*, changing its vocabulary, etc. Unless explicitly stated otherwise, we assume all operations are *strict*, i.e., they are undefined wherever any argument is undefined.

Prompting. Given $v \in \Sigma^*$, *prompting* an LM $\mathcal{L}$ with v is defined by the following transformation: $\mathsf{prompt}[v](\mathcal{L})(u) \triangleq \mathcal{L}(vu)$, for all $u \in \Sigma^*$. We use $[\cdot]$ to explicitly denote that a transformation (likewise an LM) is parametric.

Normalization. $\mathsf{norm}(\mathcal{L})$ is the LM whose outputs are *normalized* using the L^1-norm $|\cdot|_1$: $\mathsf{norm}(\mathcal{L})(u) \triangleq |\mathcal{L}(u)|_1^{-1} \cdot \mathcal{L}(u)$ if $|\mathcal{L}(u)|_1^{-1} \neq 0$, or $\bot$ otherwise.

Softmax. The *exponential* transformation with *temperature* $\mathtt{t} : \Sigma^* \to \mathbb{R}_{\geq 0}$, is such that, for all $u \in \Sigma^*$, if $\mathtt{t}(u) > 0$, $\mathsf{exp}[\mathtt{t}](\mathcal{L})(u) \triangleq \exp(\mathtt{t}(u)^{-1} \cdot \mathcal{L}(u))$ with $x \cdot (-\infty) = -\infty$ and $\exp(-\infty) = 0$, otherwise $\mathsf{exp}[\mathtt{t}](\mathcal{L})(u) \triangleq \exp(\mathbb{1}_{\arg\max \mathcal{L}(u)} \cdot \mathcal{L}(u))$.[1] The transformation $\mathsf{softmax}[\mathtt{t}](\mathcal{L}) \triangleq \mathsf{norm}(\mathsf{exp}[\mathtt{t}](\mathcal{L}))$ gives a PLM.

Masking. An *additive mask* is an LM $\mathcal{S} : \Sigma^* \to \Sigma_\top \to \{0, -\infty\}$. Masking an LM $\mathcal{L}$ with $\mathcal{S}$ gives an LM $\mathcal{L} \oplus \mathcal{S}$, such that: $(\mathcal{L} \oplus \mathcal{S})(u) \triangleq \mathcal{L}(u) \oplus \mathcal{S}(u)$, for all $u \in \Sigma^*$, where $x + (-\infty) = -\infty$. Masking $\mathcal{L}$ with a *multiplicative mask* $\mathcal{M} : \Sigma^* \to \Sigma_\top \to \{0, 1\}$ gives an LM $\mathcal{L} \odot \mathcal{M}$ such that: $(\mathcal{L} \odot \mathcal{M})(u) \triangleq \mathcal{L}(u) \odot \mathcal{M}(u)$, for all $u \in \Sigma^*$.

Top-k masking (Fan et al. (2018)) is defined as follows. Given a threshold function $\mathtt{k} : \Sigma^* \to \mathbb{N}$, $\mathsf{kmask}^{\oplus}[\mathtt{k}]$ is the parametric additive mask: $\mathsf{kmask}^{\oplus}[\mathtt{k}](\mathcal{L})(u)(\sigma) \triangleq 0$ if

1. This definition is consistent with the usual interpretation of 0 temperature in current LLMs.

$u \leftarrow \lambda$ // $P_{\mathcal{L}}(\lambda) = 1$
while True **do** // `loop invariant:` $P_{\mathcal{L}}(u) > 0$
  if $\mathcal{L}(u) = \bot$ **then return** $\bot$
  $\sigma \sim \mathcal{L}(u)$ // $\mathcal{L}(u) \neq \bot$
  if $\sigma = \top$ **then return** u // $\mathcal{L}(u)(\top) > 0 \;\wedge\; P_{\mathcal{L}}^{\top}(u) > 0$
  $u \leftarrow u\sigma$ // $\mathcal{L}(u)(\sigma) > 0 \;\wedge\; P_{\mathcal{L}}(u\sigma) > 0$
end

Algorithm 1: String generation procedure $\mathtt{gen}(\mathcal{L})$.

$\sigma \in \mathsf{top}[\mathtt{k}(u)](\mathcal{L}(u))$, or $-\infty$ otherwise, where $\mathsf{top}[k](\mathbf{r}) \subseteq \Sigma_\top$ is the set of k symbols with highest value given by $\mathbf{r} \in \mathfrak{R}(\Sigma_\top)$, and $\mathsf{top}[k](\bot) = \emptyset$. Ties are resolved using the ordering of symbols. Then, $\mathsf{topk}^{\oplus}[\mathtt{k}](\mathcal{L}) \triangleq \mathcal{L} \oplus \mathsf{kmask}^{\oplus}[\mathtt{k}](\mathcal{L})$. Analogously, we define a multiplicative $\mathsf{kmask}^{\odot}[\mathtt{k}]$ mask and a transformation $\mathsf{topk}^{\odot}[\mathtt{k}](\mathcal{L}) \triangleq \mathcal{L} \odot \mathsf{kmask}^{\odot}[\mathtt{k}](\mathcal{L})$.

Similarly, given $\mathtt{p} : \Sigma^* \to [0, 1]$, *nucleus masking* (Holtzman et al. (2020)) is the transformation $\mathsf{topp}(\mathcal{L}) \triangleq \mathcal{L} \odot \mathsf{topp}[\mathtt{p}](\mathcal{L})$ where $\mathsf{topp}[\mathtt{p}](\mathcal{L})(u)(\sigma) \triangleq 1$ if $\sigma \in \mathsf{nucleus}[\mathtt{p}(u)](\mathcal{L}(u))$, or 0 otherwise. Here, $\mathsf{nucleus}[p](\mathbf{d}) \subseteq \Sigma_\top$ is the smallest set of symbols whose cumulative probability given by $\mathbf{d} \in \mathfrak{D}(\Sigma_\top)$ is at least p, with $\mathsf{nucleus}[p](\bot) \triangleq \emptyset$.

Also, given $\Theta \subset \Sigma$ a set of *allowed* symbols, the LM over Θ that behaves as $\mathcal{L}$ is defined as $\mathsf{restrict}[\Theta](\mathcal{L})(v)(\theta) \triangleq \mathcal{L}(v)(\theta)$ for all $v \in \Theta^*$ and $\theta \in \Theta$.

Example 1 *A prompted top-k transformation of an LM $\mathcal{L} : \Sigma^* \to \mathfrak{R}(\Sigma_\top)$ is defined as:*

$$\textsc{Base}[v, \mathtt{k}] \triangleq \mathsf{topk}[\mathtt{k}]\left(\mathsf{prompt}[v](\mathcal{L})\right) \tag{1}$$

We would like to constrain Base *to obtain a PLM* Num *defined for the real numbers in $[0, 1)$ following the format given by the regular expression* $\mathtt{0.[0-9]*}$[2]*. Let $\Theta = \{\mathtt{.}, \mathtt{0}, \ldots, \mathtt{9}\} \subset \Sigma$:*

$$\textsc{Num}[v, \mathtt{k}, \mathtt{t}] \triangleq \mathsf{softmax}[\mathtt{t}]\big(\mathsf{restrict}[\Theta]\left(\textsc{Base}[v, \mathtt{k}] \oplus \textsc{NumMask}\right)\big) \tag{2}$$

where $\textsc{NumMask}(\lambda)(\mathtt{0}) = 0$, $\textsc{NumMask}(\mathtt{0})(\mathtt{.}) = 0$, $\textsc{NumMask}(\mathtt{0}.u)(\sigma) = 0$ *if $\sigma \in \Theta_\top \setminus \{\mathtt{.}\}$ for $u \in \{\mathtt{0}, \ldots, \mathtt{9}\}^+$, otherwise its value is $-\infty$.*

3. Language generation

The main purpose of a PLM $\mathcal{L} : \Sigma^* \to \mathfrak{D}(\Sigma_\top)$ is to be used as a *generative* device that produces strings that are continuations of a given string. W.l.o.g, prompting allows us to just focus on PLM that generate continuations of λ. The set of generable finite strings is the *language* of the PLM. Formally, $\mathit{lang}(\mathcal{L}) \triangleq \{u \in \Sigma^* \mid u = \mathtt{gen}(\mathcal{L})\}$ where $\mathtt{gen}(\mathcal{L})$ is the probabilistic program shown in Alg. 1. The following definitions are useful to analyze properties of $\mathtt{gen}$, such as the probability of termination.

$$P_{\mathcal{L}}(\lambda) \triangleq 1 \qquad P_{\mathcal{L}}(u\sigma) \triangleq P_{\mathcal{L}}(u) \cdot \mathcal{L}(u)(\sigma) \text{ if } \mathcal{L}(u) \neq \bot \text{ else } P_{\mathcal{L}}(u) \cdot 0 \tag{3}$$

$$P_{\mathcal{L}}^{\top}(u) \triangleq P_{\mathcal{L}}(u) \cdot \mathcal{L}(u)(\top) \text{ if } \mathcal{L}(u) \neq \bot \text{ else } P_{\mathcal{L}}(u) \cdot 0 \tag{4}$$

2. This example is borrowed from Willard and Louf (2023); Carrasco et al. (2024)

The relationship between $P_{\mathcal{L}}$ and $P_{\mathcal{L}}^{\top}$ with $\texttt{gen}(\mathcal{L})$ is shown as comments in Alg. 1. The loop invariant is $P_{\mathcal{L}}(u) > 0$ (line 2): it is true initially (line 1) and it is maintained by the assignment in line 6. Whenever $\texttt{gen}(\mathcal{L})$ returns $u \in \Sigma^*$, $P_{\mathcal{L}}^{\top}(u) > 0$ holds (line 5). Also, if $\mathcal{L}(u)(\top) > 0$, then $\top$ can be sampled in line 4. If that happens, u will be returned in line 5.

Proposition 1 *Let $\mathcal{L}$ be a PLM. For all $u \in \Sigma^*$, $u \in lang(\mathcal{L})$ if and only if $P_{\mathcal{L}}^{\top}(u) > 0$.*

Let $pref(\mathcal{L}) \triangleq \{u \in \Sigma^* \mid \mathbb{1}_{\mathcal{L}}(u) = 1\}$, where $\mathbb{1}_{\mathcal{L}}(u) = 1$ iff $P_{\mathcal{L}}(u) > 0$. Generators admit sampling from a PLM up to a maximum number $\texttt{n} \in \mathbb{N}$ of symbols. When the bound is reached before getting the terminal symbol, the sequence generated so far is returned.[3] Such behavior, not directly captured by the $\texttt{gen}$ (Alg. 1), can be modeled by the *bounded* generation function $\texttt{bndgen}[\texttt{n}]$:

$$\texttt{bnd}[\texttt{n}](\mathcal{L})(u) \triangleq \begin{cases} \mathcal{L}(u) & \text{if } |u| < \texttt{n} \\ \{\Sigma \mapsto 0, \top \mapsto 1\} & \text{otherwise} \end{cases} \quad (5) \qquad \texttt{bndgen}(\mathcal{L}, \texttt{n}) \triangleq \texttt{gen}(\texttt{bnd}[\texttt{n}](\mathcal{L})) \quad (6)$$

Corollary 2 *For all PLM $\mathcal{L}$, $pref(\mathcal{L}) = \bigcup_{\texttt{n} \in \mathbb{N}} lang(\texttt{bnd}[\texttt{n}](\mathcal{L}))$.*

Notice that $\texttt{gen}(\mathcal{L})$ terminates because either $\top$ is sampled (line 4) or $\mathcal{L}(u) = \bot$ (line 3). So, as long as during the execution of the loop there exists $\sigma \in \Sigma$ such that $\mathcal{L}(u)(\sigma) > 0$, $\texttt{gen}$ may not terminate. The probability of termination of $\texttt{gen}(\mathcal{L})$ and of $lang(\mathcal{L})$ is given by Prop. 3 based on $P_{\mathcal{L}}$ and $P_{\mathcal{L}}^{\top}$. The set of prefixes of w, including λ and w, is $\text{pref}(w) \subseteq \Sigma^*$.

Proposition 3 *For every PLM $\mathcal{L} : \Sigma^* \to \mathfrak{D}(\Sigma_{\top})$, there exists a unique Borel[4] probability measure $\boldsymbol{P}_{\mathcal{L}}$ in $\Sigma^* \cup \Sigma^{\omega} \cup \{\bot\}$ such that for all $u \in \Sigma^*$:*

$$P_{\mathcal{L}}(u) = \boldsymbol{P}_{\mathcal{L}}\{w \in \Sigma^* \cup \Sigma^{\omega} : u \in \text{pref}(w)\} \quad (7) \qquad P_{\mathcal{L}}^{\top}(u) = \boldsymbol{P}_{\mathcal{L}}\{u\} \quad (8)$$

$$\boldsymbol{P}_{\mathcal{L}}\{\bot\} = \sum_{u \in error(\mathcal{L})} P_{\mathcal{L}}(u) \text{ where } error(\mathcal{L}) \triangleq \{u \in \Sigma^* \mid \mathcal{L}(u) = \bot\}. \quad (9)$$

Corollary 4 *The following properties result from Prop. 1-3:*

1. $\boldsymbol{P}_{\mathcal{L}}(lang(\mathcal{L})) = \sum_{u \in lang(\mathcal{L})} P_{\mathcal{L}}^{\top}(u) = \sum_{u \in \Sigma^*} P_{\mathcal{L}}^{\top}(u) = \boldsymbol{P}_{\mathcal{L}}(\Sigma^*)$
2. $\texttt{gen}(\mathcal{L})$ *terminates with probability* $\boldsymbol{P}_{\mathcal{L}}(lang(\mathcal{L})) + \boldsymbol{P}_{\mathcal{L}}\{\bot\}$
3. $\texttt{gen}(\mathcal{L})$ *terminates with probability one if and only if* $\boldsymbol{P}_{\mathcal{L}}(\Sigma^{\omega}) = 0$

Example 2 *Let $\mathcal{L}$ in Exa. 1 be* GPT2*. Let $\mathcal{N}_1$ be obtained by instantiating* $\textsc{Num}[v, \texttt{k}, \texttt{t}]$ *with prompt v set to "An example of real number in the interval [0, 1) is: ", and for $u \in \Sigma^*$, $\texttt{k}(u) = 200$, and $\texttt{t}(u) = 1$. It is verified that $\mathbb{1}_{\mathcal{N}_1}(\texttt{0}.\sigma) = 1$, for all $\sigma \in \Sigma$, $\mathcal{N}_1(\texttt{0}.\sigma)(\top) = 1$, for $\sigma \in [\texttt{1}, \texttt{5}]$, otherwise $\mathcal{N}_1(\texttt{0}.\sigma) = \bot$, with $\boldsymbol{P}_{\mathcal{N}_1}\{\bot\} \approx 0.434$, $\boldsymbol{P}_{\mathcal{N}_1}(lang(\mathcal{N}_1)) \approx 0.566$, and $\boldsymbol{P}_{\mathcal{N}_1}(\Sigma^{\omega}) = 0$. Let $\mathcal{N}_2$ be obtained using the same prompt, with $\texttt{k}$ and $\texttt{t}$ be constants 200 and 0 (greedy generation), resp.: $lang(\mathcal{N}_2) = \{\texttt{0.5}\}$ and has probability equal to 1.*[5]

3. For instance, in Huggingface's `generate` function, $\texttt{n}$ corresponds to parameter `max_new_tokens`.
4. Proof in Appendix A. It extends results in Vidal et al. (2005); Dupont et al. (2005); Du et al. (2023)
5. See experiments in github public repository.

4. Language and distribution preserving transformations

A transformation that preserves $\mathit{lang}(\mathcal{L})$ for every PLM $\mathcal{L}$ is said to be *language preserving*. It is *distribution preserving* if it preserves $\boldsymbol{P}_{\mathcal{L}}$ for every PLM $\mathcal{L}$. An LM $\mathcal{L}$ is called *error-absorbing* whenever for all $u \in \Sigma^*$, $\mathcal{L}(u) = \bot$ implies $\mathcal{L}(u\sigma) = \bot$ for all $\sigma \in \Sigma$. A PLM that satisfies Def. (10) is called a *zero-absorbing* PLM (ZAPLM):

$$\mathbb{1}_{\mathcal{L}}(u) = 0 \implies \mathcal{L}(u) = \bot, \text{ for all } u \in \Sigma^*. \tag{10}$$

PLM $\mathcal{N}_1$ in Exa. 2 is not error-absorbing since it is verified that $\mathcal{N}_1(\mathsf{0.7}) = \bot$ but $\mathcal{N}_1(\mathsf{0.72}) \neq \bot$. It is not a ZAPLM neither: $\mathcal{N}_1(\mathsf{0.1})(\mathsf{2}) = 0$, and so $\mathbb{1}_{\mathcal{N}_1}(\mathsf{0.12}) = 0$, but $\mathcal{N}_1(\mathsf{0.12}) \neq \bot$.[5]

Proposition 5 *Let $\mathcal{L}$ be a ZAPLM. Then, $\mathcal{L}$ is error-absorbing.*[6]

Let the multiplicative mask zam be defined as follows:

$$\mathsf{zam}[\mathcal{L}](u) \triangleq \mathbf{1} \text{ if } \mathbb{1}_{\mathcal{L}}(u) = 1 \text{ else } \mathbf{0} \tag{11}$$

where $\mathbf{0}$ and $\mathbf{1}$ are respectively the constant 0 and 1 vectors in $\mathfrak{R}_{\geq 0}(\Sigma_\top)$. Let

$$\mathsf{za}(\mathcal{L}) \triangleq \mathsf{norm}(\mathcal{L} \odot \mathsf{zam}[\mathcal{L}]). \tag{12}$$

Def. (11) and Def. (12), imply

$$\mathsf{za}(\mathcal{L})(u) = \mathcal{L}(u) \text{ if } \mathbb{1}_{\mathcal{L}}(u) = 1 \text{ else } \bot. \tag{13}$$

Proposition 6 *Transformation* za *is both language and distribution preserving.*[7]

For every PLM $\mathcal{L}$, Prop. 6 implies $\mathsf{za}(\mathcal{L})$ is a ZAPLM[7]. Furthermore, the transformation $\mathsf{accept}_{\bot}(\mathcal{L}) : \Sigma^* \to \{0, 1, \bot\}$ defined as:

$$\mathsf{accept}(\mathcal{L})(u) \triangleq 1 \text{ if } \mathsf{za}(\mathcal{L})(u)(\top) > 0 \text{ else } 0 \tag{14}$$

$$\mathsf{accept}_{\bot}(\mathcal{L})(u) \triangleq \bot \text{ if } \mathsf{za}(\mathcal{L})(u) = \bot \text{ else } \mathsf{accept}(\mathcal{L})(u) \tag{15}$$

is s.t. $\mathit{lang}(\mathcal{L}) = \{u \in \Sigma^* \mid \mathsf{accept}_{\bot}(\mathcal{L})(u) = 1\}$, and $\mathit{error}(\mathcal{L}) = \{u \in \Sigma^* \mid \mathsf{accept}_{\bot}(\mathcal{L})(u) = \bot\}$[7]. For a ZAPLM $\mathcal{L}$, $\mathit{error}(\mathcal{L})$ may be non-empty. Equivalently, $\mathtt{gen}(\mathcal{L})$ may return $\bot$. For this to happen, there must exist some $u \in \Sigma^*$ such that $\mathbb{1}_{\mathcal{L}}(u) = 1$ and $\mathcal{L}(u) = \bot$. Formally:

Proposition 7 *Let $\mathcal{L} : \Sigma^* \to \mathfrak{D}(\Sigma_\top)$. $\boldsymbol{P}_{\mathcal{L}}\{\bot\} \in \{0, 1\}$ if $\mathcal{L}$ is such that* [8]

$$\mathcal{L}(u) = \bot \overset{\triangle}{\implies} \mathbb{1}_{\mathcal{L}}(u) = 0, \text{ for all } u \in \Sigma^+. \tag{16}$$

The restriction to non-empty strings is a consequence of Def. (3) which implies $\mathbb{1}_{\mathcal{L}}(\lambda) = 1$. A PLM $\mathcal{L}$ can be transformed into a $\{0, 1\}$-error PLM by appropriately constraining it. Let

$$\mathsf{zo}(\mathcal{L}) \triangleq \mathsf{norm}(\mathcal{L} \odot \mathsf{zom}[\mathcal{L}]) \tag{17}$$

6. Proof in Appendix B.1.
7. Proof in Appendix B.2.
8. Proof in Appendix B.3.

where the mask $\mathsf{zom}[\mathcal{L}]$ is such that for all $u \in \Sigma^*$:

$$\mathsf{zom}[\mathcal{L}](u)(\sigma) \triangleq 0 \text{ if } \sigma \in \Sigma \wedge \mathbb{1}_{\mathcal{L}}(u\sigma) = 1 \wedge \mathcal{L}(u\sigma) = \bot \text{ else } 1. \tag{18}$$

The effect of zo is to *zero out* the probabilities of symbols that can be sampled in line 4 of `gen` and lead to returning $\bot$ in line 3, in the immediately following iteration (Alg. 1). For example, in $\mathcal{N}_1$ of Exa. 2, $\mathsf{zo}(\mathcal{N}_1)(\mathbf{0}.)(\sigma) = 0$ for $\sigma \notin [1, 5]$, and so $\boldsymbol{P}_{\mathsf{zo}(\mathcal{N}_1)}\{\bot\} = 0$. It follows that, unlike za, transformation zo is not distribution preserving. Nevertheless:

Proposition 8 zo *is language preserving.*[9]

In general, a single application of zo is not be enough to ensure Def. (16). Consider a PLM $\mathcal{L}$ over $\Sigma = \{a\}$ such that $\mathbb{1}_{\mathcal{L}}(aa) = 1$, $\mathcal{L}(aa) = \bot$, and $\mathcal{L}(a)(a) = 1$. Then, $\mathsf{zom}[\mathcal{L}](a)(a) = 0$, and so $\mathsf{zo}(\mathcal{L})(a) = \bot$, yet $\mathbb{1}_{\mathsf{zo}(\mathcal{L})}(a) = 1$ because $\mathsf{zo}(\mathcal{L})(u) = \mathcal{L}(u)$ for $u \in \{\lambda, a\}$. The following result states that being a fixpoint of zo is sufficient for Def. (16) to hold.

Proposition 9 *Let* $\mathcal{L} : \Sigma^* \to \mathfrak{D}(\Sigma_\top)$. *If* $\mathcal{L} = \mathsf{zo}(\mathcal{L})$, *then* $\mathcal{L}$ *satisfies Def.* (16).[10]

The set of of multiplicative masks (over $\Sigma_\top$), denoted $\mathbb{M}$, is a *complete lattice* where the minimum $O : u \mapsto \mathbf{0}$, the maximum $I : u \mapsto \mathbf{1}$, and $\mathcal{M}_1 \leq \mathcal{M}_2$ if and only if for all $u \in \Sigma^*$, $\mathcal{M}_1(u)(\sigma) \leq \mathcal{M}_2(u)(\sigma)$, for all $\sigma \in \Sigma_\top$. Let $\mathcal{M} \in \mathbb{M}$. We define:

$$\mathsf{m}[\mathcal{L}](\mathcal{M})(u)(\sigma) \triangleq \mathsf{em}[\mathcal{L}, \mathcal{M}](u\sigma) \cdot \mathcal{M}(u)(\sigma) \tag{19}$$

where

$$\mathsf{em}[\mathcal{L}, \mathcal{M}](u\sigma) \triangleq 0 \text{ if } \sigma \in \Sigma \wedge \mathbb{1}_{\mathsf{norm}(\mathcal{L} \odot \mathcal{M})}(u\sigma) = 1 \wedge \mathsf{norm}(\mathcal{L} \odot \mathcal{M})(u\sigma) = \bot \text{ else } 1. \tag{20}$$

Notice that $\mathsf{zom}[\mathcal{L}] = \mathsf{m}[\mathcal{L}](I)$. Moreover, $\mathsf{m}[\mathcal{L}]$ is *reductive*, that is, $\mathsf{m}[\mathcal{L}](\mathcal{M}) \leq \mathcal{M}$. Hence, $\{\mathcal{M}_i\}_{i \geq 0}$, with $\mathcal{M}_0 = I$, and $\mathcal{M}_{i+1} = \mathsf{m}[\mathcal{L}](\mathcal{M}_i)$, is a *descending chain*, that is, $\mathcal{M}_i \geq \mathcal{M}_{i+1}$. Let $\mathcal{M}_\infty$ be the *infimum* of the chain, which exists because $\mathbb{M}$ is a complete lattice.

Proposition 10 *Let* $\mathcal{L} : \Sigma^* \to \mathfrak{D}(\Sigma_\top)$. *Then,* $\mathsf{norm}(\mathcal{L} \odot \mathcal{M}_\infty)$ *is a fixpoint of* zo.[11]

Multiplicative mask $\mathcal{M}_\infty$ can be approximated via the auto-corrective generation procedure `safegen` (Alg. 2). It keeps a *persistent* multiplicative mask $\mathcal{M}$, initially set to I. Consider the case $\mathsf{norm}(\mathcal{L} \odot \mathcal{M})(u) = \bot$ (line 4). If $u = \lambda$, it means the PLM is 1-error, so `safegen` returns $\bot$. If $u = v\sigma$, it follows that u satisfies the if-condition in Def. (20). In this case, `safegen` first updates the mask by zeroing the error-enabling symbol (line 7), as in Def. (20), and then backtracks to continue generating from v (lines 8-9). When $\mathsf{norm}(\mathcal{L} \odot \mathcal{M})(u) \neq \bot$, `safegen` samples from $\mathsf{norm}(\mathcal{L} \odot \mathcal{M})(u)$. Recall $\top$ is never masked. Then, if the condition in line 4 holds, then $\mathcal{L}(u)(\top) = 0$, and so $u \notin \mathit{lang}(\mathcal{L})$. Also, if $\top$ is sampled in line 11, then $\mathcal{L}(u)(\top) > 0$, and so $u \in \mathit{lang}(\mathsf{norm}(\mathcal{L} \odot \mathcal{M}))$. Hence, by Prop. 8, $u \in \mathit{lang}(\mathcal{L})$.

Example 3 *Fig. 1 shows the effect of iterating* `safegen` *for* $\mathcal{L}$ *defined as the instantiation of* Num$[v, \mathtt{k}, \mathtt{t}]$ *with* Gpt2, v *as in Exa. 2,* $\mathtt{k}(u) = 310$, *and* $\mathtt{t}(u) = 1$. *Let* $\{\mathcal{L}_i\}_{i \geq 0}$ *be the sequence of models* $\mathcal{L}_i \triangleq \mathsf{norm}(\mathcal{L} \odot \mathcal{M}_i)$, *for iteration* i. $\boldsymbol{P}_{\mathcal{L}_i}\{\bot\}$ *decreases. This is because* $\{\mathcal{M}_i\}_{i \geq 0}$ *is a decreasing chain. Moreover, by (the proof of) Prop. 7 and the fact that* $\boldsymbol{P}_{\mathcal{L}}\{\bot\} < 1$, $\boldsymbol{P}_{\mathcal{L}_i}\{\bot\}$ *tends to 0. Besides,* $\boldsymbol{P}_{\mathcal{L}_i}(\mathit{lang}(\mathcal{L}_i))$ *increases and tends to 1.*[5]

9. Proof in Appendix B.4
10. Proof in Appendix B.5.
11. Proof in Appendix B.6.

static $\mathcal{M}(v) \leftarrow \mathbf{1}$ for all $v \in \Sigma^*$ `// persistent mask`
$u \leftarrow \lambda$
while True **do** `// loop invariant:` $\mathbb{1}_{\mathsf{norm}(\mathcal{L} \odot \mathcal{M})}(u) = 1$
  if $\mathsf{norm}(\mathcal{L} \odot \mathcal{M})(u) = \bot$ **then**
    if $u = \lambda$ **then return** $\bot$ `//` $\boldsymbol{P}_{\mathsf{norm}(\mathcal{L} \odot \mathcal{M})}\{\bot\} = 1$
    let $u = v\sigma$
      $\mathcal{M}(v)(\sigma) \leftarrow 0$ `// apply Def.` (19)
      $u \leftarrow v$ `//` $\mathbb{1}_{\mathsf{norm}(\mathcal{L} \odot \mathcal{M})}(v) = 1$
    continue `// backtrack`
  end
  $\sigma \sim \mathsf{norm}(\mathcal{L} \odot \mathcal{M})(u)$ `//` $\mathsf{norm}(\mathcal{L} \odot \mathcal{M})(u) \neq \bot$
  if $\sigma = \top$ **then return** u `//` $u \in \mathit{lang}(\mathsf{norm}(\mathcal{L} \odot \mathcal{M}))$
  $u \leftarrow u\sigma$ `//` $\mathbb{1}_{\mathsf{norm}(\mathcal{L} \odot \mathcal{M})}(u) = 1$
end

Algorithm 2: Auto-corrective string generation procedure `safegen`($\mathcal{L}$).

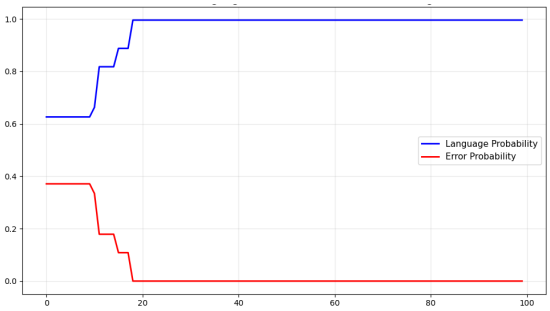


Figure 1: Effect of iterating `safegen`.

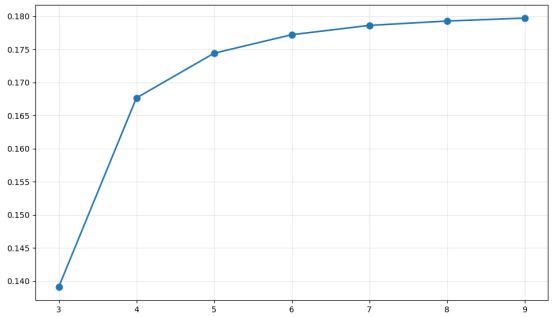


Figure 2: Probability of $L_{[0,0.2)}$.

5. Applications

Given a PLM $\mathcal{L}$, the probability of $\boldsymbol{P}_{\mathcal{L}}(L)$ of a language $L \subseteq \Sigma^*$ can be approximated with the help of the BFS-exploration procedure `explore` (Alg. 3). Starting with λ, it keeps a queue of strings which remain to be analyzed. Basically, it explores a *trie* whose nodes are strings, and there is an arc from u to $u\sigma$ only if $\mathcal{L}(u)(\sigma) > 0$ and $u\sigma$ is required for the analysis. The purpose of the `callback` function is twofold: it performs the analysis of interest and decides whether to expand the string.

$q \leftarrow \{\lambda\}$
while $q \neq \emptyset$ **do**
  $u \leftarrow q.pop()$
  $\mathit{expand} \leftarrow$ `callback`$(\mathcal{L}, u)$
  if $\mathit{expand} = 1$ **then foreach** $\sigma \in \Sigma$ s.t. $\mathcal{L}(u)(\sigma) > 0$ **do** $q \leftarrow q.push(u\sigma)$
end

Algorithm 3: `explore`($\mathcal{L}$, `callback`).

static $p_\top \leftarrow 0$ `// language probability`
static $p_\bot \leftarrow 0$ `// error probability`
$expand \leftarrow 0$ if $\mathcal{L}(u) = \bot$ else 1
if $expand = 0$ **then** $p_\bot \leftarrow p_\bot + P_{\mathcal{L}}(u)$ **else** $p_\top \leftarrow p_\top + P^{\top}_{\mathcal{L}}(u)$
return $expand$

Algorithm 4: Callback for computing language and error probability.

Alg. 4 shows the `callback` function that accumulates the probabilities of strings that belong to either $lang(\mathcal{L})$ (in $p_\top$) or $error(\mathcal{L})$ (in $p_\bot$). Furthermore, it decides to expand a string u provided $\mathcal{L}(u) \neq \bot$. This callback has been applied to obtain the curves depicted in Fig. 1, by calling `explore` after each iteration of `safegen`.

Consider the set $L_{[0,0.2)}$ of strings corresponding to real numbers in $[0, 0.2)$. Fig. 2 shows the cumulated probability by length of strings in $L_{[0,0.2)}$ which are generable by $\texttt{safegen}(\mathcal{L})$, i.e., $lang(\mathcal{L}) \cap L_{[0,0.2)}$, for $\mathcal{L}$ defined in Exa. 3, where overall $\boldsymbol{P}_{\mathcal{L}}(L_{[0,0.2)}) \approx 0.18$.[5]

Besides computing probabilities of sets of strings, this formalization enables learning automata representations of languages through *membership* (MQ) and *equivalence* (EQ) queries. This can be done by defining $\text{MQ}_{\mathcal{L}}$ as either accept or $\mathsf{accept}_\bot$, depending on whether we want to characterize $lang(\mathcal{L})$ alone or $error(\mathcal{L})$ as well. Fig. 3 depicts the automaton representation of $\mathsf{accept}_\bot(\mathcal{L})$ for PLM $\mathcal{L}$ in Exa. 3. This Moore machine has been learned with an implementation of Angluin's L* using strings up to length 5 for EQ.

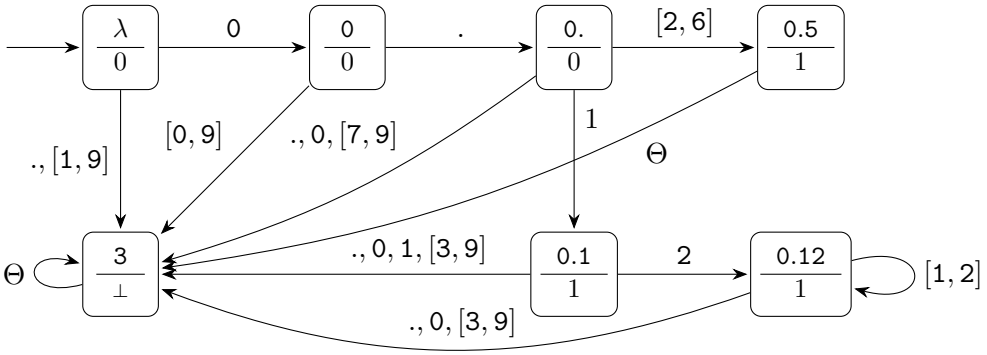


Figure 3: Learned Moore machine of $\mathsf{accept}_\bot(\mathcal{L})$ in Exa. 3

6. Case study

The goal of this case study is to analyze to what extent a real-world open-weight LLM can generate numbers with some property stated as instruction, and formatted according to the DFA Number with alphabet $\Theta = \{-, ., \mathrm{e}, 0, \ldots, 9\}$ shown in Fig. 4. This format accepts numbers to be written in both decimal and scientific e notation.[12] For instance, 21.76 and 2.176e1 are accepted, as well as non-normalized writings such as 21.76e0 or 0.2176e2, etc.[13]

12. For simplicity, the grammar only accepts e, instead of E or even both.
13. In *normalized* scientific e notation, the mantissa is a single non-zero digit, e.g., 2.176e1.

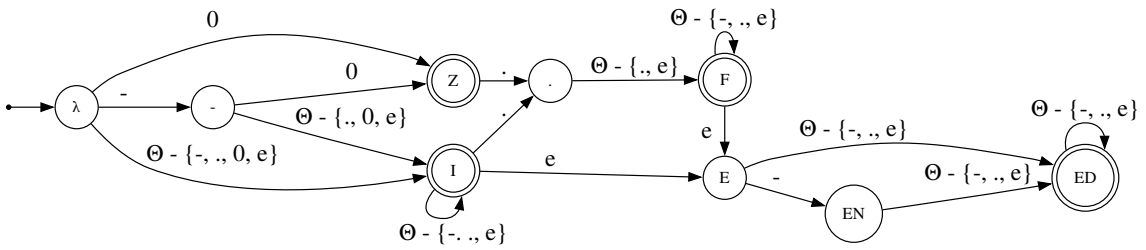


Figure 4: DFA Number for numbers. State labels are merely informative. The initial state is the one labeled λ. Double-circles depict accepting states. Multiple outgoing transitions with the same target sate are grouped.

6.1. Modeling

Let $\mathcal{L}$ be an LLM over Σ. We define the *model under analysis* MUA parameterized transformation as follows:

$$\mathrm{MUA}[v, \mathtt{k}, \mathtt{p}, \mathtt{t}](\mathcal{L}) \triangleq (\mathsf{za} \circ \mathsf{norm} \circ \mathsf{topp}[\mathtt{p}] \circ \mathsf{softmax}[\mathtt{t}] \circ \mathsf{topk}[\mathtt{k}]^{\oplus} \circ \mathsf{prompt}[v])(\mathcal{L}) \tag{21}$$

Prompt v is the instruction given to the LLM, parameters $\mathtt{t}$, $\mathtt{k}$, and $\mathtt{p}$ constrain the distribution over Σ, and DFA Number acts as a decoding contraint. Since $\mathrm{MUA}(\mathcal{L})$ is a PLM over Σ and Number is a DFA over Θ, a *tokenizer* is required to *encode* strings in Θ^* into strings in Σ^*. Formally, we define a tokenizer as an injective function $\mathtt{enc} : \Theta^* \to \Sigma^*$. Notice that $\mathtt{enc}$ is bijective over its image. Let $\mathtt{enc}^{-1}$ denote its inverse image.

Let $G \subseteq \Theta^*$ be a language of interest for the analysis. We define the *target* languages, over Σ and Θ, as follows:

$$\mathsf{Target}_{\Sigma}(G) \triangleq \mathit{lang}(\mathrm{MUA}(\mathcal{L})) \cap \mathtt{enc}(G) \tag{22}$$

$$\mathsf{Target}_{\Theta}(G) \triangleq \mathtt{enc}^{-1}(\mathsf{Target}_{\Sigma}(G)) \tag{23}$$

6.2. Experiments

As LLMs $\mathcal{L}$ we use four models of the Qwen3 family, namely, the hybrid 4B model and the mixture-of-experts model 30B-A3B, as well as their latest versions 4B-2507 and 30B-A3B-2507, all in instruct mode. The Qwen3 family models have default (or suggested) values for parameters controlling token distribution, namely, $\mathtt{p} = 0.8$, $\mathtt{k} = 20$, and $\mathtt{t} = 0.7$.

6.2.1. Computing probabilities

In this experiment, we instantiate $\mathcal{L}$ in MUA with Qwen3-4B, and use the default configuration parameters. For this, we use Alg. 3, with callback functions set to cut expanding at length equal to 4. We started with prompt "Generate a number between -1 and 1. Only reply with the number".[14] This defines the language of interest

$$G_{[-1,1]} \triangleq \mathsf{Number} \cap \{x \in \Theta^* \mid -1 \leq x \leq 1\} \tag{24}$$

14. Instruction-tuned LLMs like to respond with a large explanation before generating the actual number. By adding "Only ..." the model is more willing to generate the number right away.

Since $\boldsymbol{P}$ is defined over Σ, probabilities are computed for $\mathsf{Target}_\Sigma(G_{[-1,1]})$. We found $\boldsymbol{P}(\mathsf{Target}_\Sigma(G_{[-1,1]})) \approx 0.5$, and also discovered that the probability of negative numbers is 0, meaning the model does not generate numbers below 0.

Prompt "Generate a positive integer. Only reply with the number" gives the property $G_{(0,\infty)}$. Its probability is approximately 0.78, and with probability 0.22 the model under analysis generates (non-integer) numbers containing ".".

6.2.2. Learning

Here, we are interested in learning the target language $\mathsf{Target}_\Sigma(\mathsf{Number})$, for several values of `t`, `k`, and `p`. For the sake of readability, we reference it as just Target. As prompt, we use the character string "Generate a real number between 0 and 1 in scientific e notation. Only reply with the number". Let $G_{[0,1]}$ be the language that corresponds to the prompt.[15] It can be represented by a minimal DFA obtained by removing state $-$ (minus) and some appropriate arcs from DFA Number in Fig. 4. As DFA learning algorithm, we use an implementation of Bounded-L^* based on a probably approximately correct (PAC) equivalence Mayr and Yovine (2018).

Experiment 1 Here, we set generation parameters `t` and `p` to 1, and learn DFA for each LLM, with different values of `k`. PAC parameters ε and δ are both set to 0.001. Every equivalence query (EQ) randomly samples half of the sequences in Θ^*, and half in the language of Number.[16] Hence, we can safely use as bound the number of states of Number, because any DFA with more than 9 states could not be equivalent to neither Number nor $G_{[0,1]}$. Besides, we use a time bound of 500 seconds.

	Qwen3-4B			Qwen3-4B-2507			Qwen3-30B-A3B			Qwen3-30B-A3B-2507		
Top-k	States	Time(s)	Target-EQ	States	Time(s)	Target-EQ	States	Time(s)	Target-EQ	States	Time(s)	Target-EQ
20	>9	14.3	False	>9	15.5	False	>9	174.8	False	>9	150.8	False
200	>9	19.4	False	>9	15.0	False	>9	143.0	False	>9	187.4	False
2000	>9	18.6	False	>9	26.5	False	8	500.0	False	>9	120.9	False
20000	>9	20.8	False	>9	17.1	False	4	500.0	False	>9	142.1	False
40000	>9	13.5	False	>9	35.7	False	6	500.0	False	8	500.0	False
80000	>9	24.2	False	>9	28.1	False	>9	150.4	False	7	294.4	True
120000	>9	194.0	False	>9	29.0	False	>9	243.9	False	7	295.8	True
140000	>9	131.5	False	8	48.2	True	7	308.3	True	8	252.8	True
None	8	42.7	True	8	39.5	True	7	301.3	True	8	243.6	True
None				9	418.8	True						

Table 1: Experiments across models with `t`=1 and `p`=1.

Table. 1 shows the number of states of the learned DFAs, the running times of the learning algorithm, and whether the resulting hypothesis is equivalent to the target language, for several values of `k`. "None" means the whole alphabet of Qwen3 family (∼150K) is used. The results reveal that Qwen3-4B versions are clearly not equivalent to neither $G_{[0,1]}$ nor Number for values of `k` below 120K, because the number of states grew larger than 9, the size of the minimal DFAs. For `k` less than 80K, none of the fourth model is equivalent to

15. Formally speaking, it should be the rational numbers between 0 and 1.
16. It is worth noticing that sampling only from Θ^* did not result in learning meaningful DFA because most of the sampled sequences did not belong to the language of Number.

	Qwen3-4B			Qwen3-4B-2507			Qwen3-30B-A3B			Qwen3-30B-A3B-2507		
Top-k	States	Time(s)	PAC-EQ	States	Time(s)	PAC-EQ	States	Time(s)	PAC-EQ	States	Time(s)	PAC-EQ
20	>8	33.7	False	>8	28.8	False	>8	293.7	False	>8	286.3	False
200	>8	33.9	False	>8	29.1	False	>8	354.6	False	8	283.6	True
2000	8	46.3	False	8	42.1	False	8	473.2	True	8	285.3	True
20000	8	53.5	True	8	45.2	True	8	473.3	True	8	285.7	True
40000	8	55.7	True	8	47.1	True	8	473.3	True	8	287.2	True
80000	8	58.9	True	8	49.7	True	8	477.0	True	8	289.2	True
120000	8	62.2	True	8	52.6	True	8	479.9	True	8	292.7	True
140000	8	64.1	True	8	54.3	True	8	483.3	True	8	292.3	True
None	8	52.7	True	8	44.4	True	8	471.4	True	8	284.1	True

Table 2: Experiments for prefix language across models with $\mathtt{t} = 1$ and $\mathtt{p} = 1$.

	Qwen3-4B			Qwen3-4B-2507			Qwen3-30B-A3B			Qwen3-30B-A3B-2507		
Top-k	States	Time(s)	PAC-EQ	States	Time(s)	PAC-EQ	States	Time(s)	PAC-EQ	States	Time(s)	PAC-EQ
20.0	3	500.0	False	4	16.7	True	4	109.2	True	4	104.7	True
200.0	7	94.1	True	4	15.8	True	4	107.9	True	4	105.2	True
2000.0	7	96.0	True	4	16.0	True	4	107.4	True	4	104.9	True
20000.0	7	98.6	True	4	16.9	True	4	108.3	True	4	103.3	True
40000.0	7	102.2	True	4	17.2	True	4	108.1	True	4	104.4	True
80000.0	7	109.0	True	4	18.7	True	4	108.8	True	4	105.6	True
120000.0	7	116.6	True	4	20.1	True	4	110.2	True	4	106.8	True
140000.0	7	120.2	True	4	20.3	True	4	110.4	True	4	107.6	True
None	7	111.1	True	4	18.9	True	4	109.3	True	4	105.9	True

Table 3: Experiments for prefix language across models with default parameters.

Target. The cases where a hypothesis is found to be PAC-equivalent to Target, it turns out that it is checked not to be equivalent to Target, because it accepts strings not in Number.

We performed additional experiments with Qwen3-4B-2507 with the entire alphabet, lowering the values of ϵ and δ to 0.0001 (last line of Table 1). In this case, the learned hypothesis is verified to actually be equivalent to Target. This concludes that the limitation here is the PAC precision. The decision of not using a more precise PAC-EQ is the significant increase in time for learning in all the models (expected roughly 10x increase as suggested by the last experiment). Making it prohibitively expensive in some cases.

Experiment 2 To analyze the impact of the default configuration on the generated language, we applied the learning algorithm just for Qwen3-30B-A3B-2507, setting `p` and `t` to the recommended values, while varying `k` starting at 20. PAC parameters were set to 0.001 as before. All experiments finished in about 105 secs with a positive EQ. In all cases, the learned hypothesis consisted in the empty language. It turns out that, attempts to query the instantiated MUA with these parameters, resulted in a quite small set of outputs (less than 100). In other words, the actual target language is a finite set of very low cardinality. Therefore, it becomes virtually impossible to get a positive sequence by sampling from Θ^* and Number using PAC parameters equal to 0.001. Notwithstanding, we made another interesting observation: the logits of the terminal token, more often than not were very small. As a consequence, the combined effect of top-k masking, and the application of a temperature less than 1 that concentrates the distribution, and 0.8-nucleus sampling makes the terminal symbol rarely be in the support of the token distribution during generation.

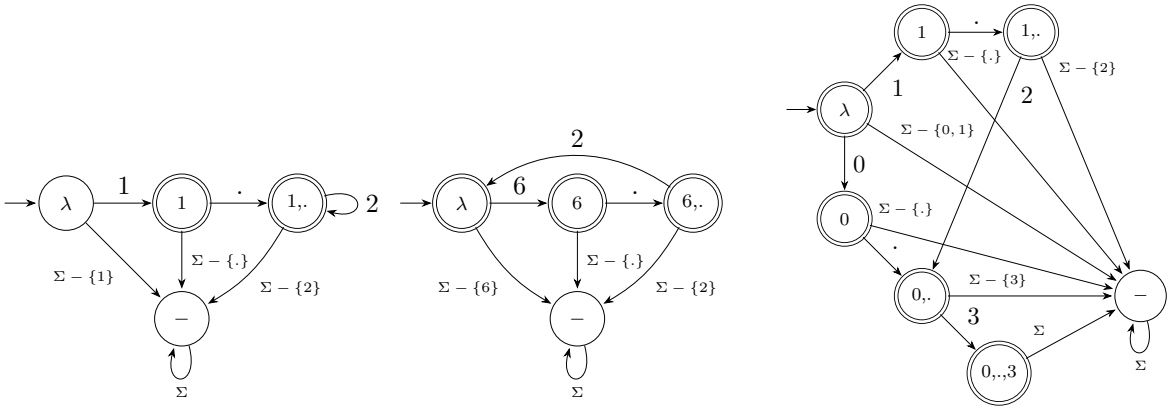


Figure 5: (left) Qwen3-4B-2507 (center) Qwen3-30B-A3B/2507 (right) Qwen3-4B

Now, we can define the target language as the prefix language of MUA as follows:

$$\mathsf{PrefTarget}_{\Sigma} \triangleq \mathit{pref}(\mathrm{MUA}(\mathcal{L})) \cap \mathtt{enc}(\mathsf{PrefNumber}) \tag{25}$$

$$\mathsf{PrefTarget}_{\Theta} \triangleq \mathtt{enc}^{-1}(\mathsf{PrefTarget}_{\Sigma}) \tag{26}$$

where PrefNumber is obtained from Number in Fig. 4 by setting all states to be accepting, and removing state ED since it is equivalent to EN to make the DFA minimal (8 states).

Experiment 3 Tables 2 and 3 show the results of learning $\mathsf{PrefTarget}_{\Theta}$, for the same models and configurations as in the previous experiments. Clearly, for $\mathtt{t} = 1$ and $\mathtt{p} = 1$, all DFA with more than 8 states are not equivalent to PrefNumber. Those with 8 states have been checked not to be equivalent neither. Figure 5 shows the DFA obtained with the default configuration for all language models. Overall, no one is equivalent to PrefNumber. The learned DFA are consistent with the most common strings over 3,000 generations: 1.234... and 0.345... for Qwen3-4B, 1.23e-1 (unique sequence) for Qwen3-4B-2507, 6.23..., 6.28..., and 6.32... for Qwen3-30B-A3B, and 6.28e-1 (unique sequence) for Qwen-30B-A3B-2507.

7. Conclusions

We introduced a formal framework for analyzing constrained language models by viewing decoding-time operations as composable transformations of language models. By making partiality, termination, and error behavior explicit, we identified a class of zero-absorbing models that admit sound abstraction into acceptors, preserving the induced language and the probability distribution of strings. We also formalized bounded-length generation typically used by state-of-the-art LLMs. Moreover, we analyzed a safe generation procedure that avoids blocking during decoding, in line with backtracking mechanisms proposed for constraint decoding. The framework enables automata-theoretic analysis and trie-based exploration to compute probabilities of sets of strings (properties).

Acknowledgments

The research that led to the results presented in this publication received funding from Agencia Nacional de Investigación e Innovación (ANII) under grants POS_FMV_2023_1_1012218 and FMV_1_2023_1_175864, and PEDECIBA, Uruguay.

References

Luca Beurer-Kellner, Marc Fischer, and Martin Vechev. Prompting is programming: A system for constrained language model generation. In *Proceedings of the ACM on Programming Languages (PACMPL)*, volume 8. ACM, 2024.

Matías Carrasco, Franz Mayr, Sergio Yovine, Johny Kidd, Martín Iturbide, Juan Pedro da Silva, and Alejo Garat. Analyzing constrained llm through pdfa-learning. LearnAut 2024, 2024. URL https://arxiv.org/abs/2406.08269.

L. Du, L. Torroba Hennigen, T. Pimentel, C. Meister, J. Eisner, and R. Cotterell. A measure-theoretic characterization of tight language models. In *61st ACL*, pages 9744–9770, July 2023.

Pierre Dupont, François Denis, and Yann Esposito. Links between probabilistic automata and hidden markov models: probability distributions, learning models and induction algorithms. *Pattern recognition*, 38(9):1349–1371, 2005.

Angela Fan, Mike Lewis, and Yann Dauphin. Hierarchical neural story generation. In Iryna Gurevych and Yusuke Miyao, editors, *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 889–898, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-1082. URL https://aclanthology.org/P18-1082/.

Xinyu Geng, Xiaojun Wan, Lin Wang, Min Zhang, and Jie Jiang. Grammar-constrained decoding for llms with lookahead backtracking. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, page 10932–10952. ACL, 2023.

Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration, 2020. URL https://arxiv.org/abs/1904.09751.

Terry Koo, Frederick Liu, and Luheng He. Automata-based constraints for language model decoding, 2024. URL https://arxiv.org/abs/2407.08103.

Franz Mayr and Sergio Yovine. Regular inference on artificial neuralnetworks. In *MAKE*, pages 350–369, Cham, 2018. Springer International Publishing.

P. C. Shields. *The ergodic theory of discrete sample paths*. AMS, 1996.

Tarun Suresh, Nalin Wadhwa, Debangshu Banerjee, and Gagandeep Singh. Beaver: An efficient deterministic llm verifier, 2026. URL https://arxiv.org/abs/2512.05439.

Shubham Ugare, Rohan Gumaste, Tarun Suresh, Gagandeep Singh, and Sasa Misailovic. Itergen: Iterative semantic-aware structured llm generation with backtracking, 2025. URL https://arxiv.org/abs/2410.07295.

E. Vidal, F. Thollard, C. de la Higuera, F. Casacuberta, and R.C. Carrasco. Probabilistic finite-state machines - part i. *IEEE PAMI*, 27(7):1013–1025, 2005.

Gail Weiss, Yoav Goldberg, and Eran Yahav. Thinking like transformers. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 11080–11090. PMLR, 18–24 Jul 2021. URL https://proceedings.mlr.press/v139/weiss21a.html.

B. T. Willard and R. Louf. Efficient guided generation for LLMs. *Preprint arXiv:2307.09702*, 2023.

Max Willsey, Willard Brandon, Dan Reid, and Clark Barrett. Outlines: Fast and deterministic linear-time constrained decoding. In *International Conference on Machine Learning (ICML) Workshops*, 2024.

Appendix A. Proof of Proposition 3

Proof We first enlarge the alphabet to accommodate both the termination and error symbols. Let Σ be the base alphabet and let

$$\Sigma_{\text{ext}} \triangleq \Sigma \cup \{\top, \bot\}.$$

Let Σ^* denote the set of finite words over Σ. Recall that a PLM $\mathcal{L}$ is a (possibly partial) function $\mathcal{L} : \Sigma^* \to \mathfrak{D}(\Sigma_\top)$ with an additional undefined value equal to $\bot$.

We extend $\mathcal{L}$ to a *total* function

$$\mathcal{L}_{\text{ext}} : \Sigma^*_{\text{ext}} \to \mathfrak{D}(\Sigma_{\text{ext}})$$

as follows. For $w \in \Sigma^*_{\text{ext}}$, set

$$\mathcal{L}_{\text{ext}}(w) = \begin{cases} \mathcal{L}(w) & \text{if } w \in \Sigma^* \text{ and } \mathcal{L}(w) \in \mathfrak{D}(\Sigma_\top), \\ \delta_\bot & \text{if } w \in \Sigma^* \text{ and } \mathcal{L}(w) = \bot, \\ \delta_\top & \text{if } w \in (u\top)\Sigma^*_{\text{ext}} \text{ for some } u \in \Sigma^*, \\ \delta_\bot & \text{if } w \in (u\bot)\Sigma^*_{\text{ext}} \text{ for some } u \in \Sigma^*. \end{cases} \tag{27}$$

In the above definition, if $\mathcal{L}(w) \in \mathfrak{D}(\Sigma_\top)$, we regard it as a distribution over Σ_{ext} by extending it with zero mass on $\bot$.

Notice that all possible cases are considered in (27). Indeed, let $w \in \Sigma^*_{\text{ext}}$, there are two disjoint possibilities:

1. w contains no $\top$ and no $\bot$. Then all symbols of w lie in Σ, so $w \in \Sigma^*$. Since $\mathcal{L}$ is defined only on Σ^*, exactly one of the following holds:
   - $\mathcal{L}(w) \in \mathfrak{D}(\Sigma_\top)$ (first case in (27)),

- $\mathcal{L}(w) = \perp$ (second case in (27)).

2. w contains at least one $\top$ or $\perp$. Then we can write $w = u\sigma v$, where σ is the first symbol of w that belongs to $\{\top, \perp\}$, and so $u \in \Sigma^*$ and $v \in \Sigma^*_{\text{ext}}$. Thus w satisfies exactly one of:

   - $w \in (u\top)\Sigma^*_{\text{ext}}$ if and only if $\sigma = \top$ (third case in (27)),
   - $w \in (u\perp)\Sigma^*_{\text{ext}}$ if and only if $\sigma = \perp$ (fourth case in (27)).

For each $k \geq 1$ we define the finite–dimensional distribution $\boldsymbol{P}_k : \Sigma^k_{\text{ext}} \to [0,1]$ by

$$\boldsymbol{P}_k\{w\} \triangleq \prod_{i=1}^{k} \mathcal{L}_{\text{ext}}(w_{<i})(\sigma_i),$$

where $w = \sigma_1 \cdots \sigma_k \in \Sigma^k_{\text{ext}}$ and $w_{<i} := \sigma_1 \cdots \sigma_{i-1}$ with the convention $w_{<1} = \lambda$.

We now check that $\{\boldsymbol{P}_k\}_{k\geq 1}$ is a consistent family of finite–dimensional distributions. For $w \in \Sigma^k_{\text{ext}}$,

$$\begin{aligned}
\sum_{\sigma_{k+1}\in\Sigma_{\text{ext}}} \boldsymbol{P}_{k+1}\{w\sigma_{k+1}\} &= \sum_{\sigma_{k+1}\in\Sigma_{\text{ext}}} \left[\prod_{i=1}^{k} \mathcal{L}_{\text{ext}}(w_{<i})(\sigma_i)\right] \mathcal{L}_{\text{ext}}(w)(\sigma_{k+1}) \\
&= \left[\prod_{i=1}^{k} \mathcal{L}_{\text{ext}}(w_{<i})(\sigma_i)\right] \sum_{\sigma_{k+1}\in\Sigma_{\text{ext}}} \mathcal{L}_{\text{ext}}(w)(\sigma_{k+1}) \\
&= \prod_{i=1}^{k} \mathcal{L}_{\text{ext}}(w_{<i})(\sigma_i) = \boldsymbol{P}_k\{w\},
\end{aligned}$$

since $\mathcal{L}_{\text{ext}}(w)$ is a probability distribution on Σ_{ext}. Thus $\{\boldsymbol{P}_k\}$ is consistent.

By Kolmogorov's Extension Theorem (see Shields (1996), Thm. I.1.2), there exists a unique probability measure $\widehat{\boldsymbol{P}}$ on $\Sigma^\omega_{\text{ext}}$ whose restriction to cylinder sets

$$C(a_1, \ldots, a_k) \triangleq \{(\sigma_i)_{i\geq 1} \in \Sigma^\omega_{\text{ext}} : \sigma_i = a_i \text{ for } 1 \leq i \leq k\}$$

satisfies

$$\widehat{\boldsymbol{P}}\left(C(a_1, \ldots, a_k)\right) = \boldsymbol{P}_k\{a_1 \cdots a_k\}.$$

Notice that

$$C(a_1, \ldots, a_k) = \left\{x \in \Sigma^\omega_{\text{ext}} : a_1 \cdots a_k \in \text{pref}(x)\right\}.$$

$\widehat{\boldsymbol{P}}$ concentrates on absorbing sequences. Define

$$A_\top \triangleq \{x \in \Sigma^\omega_{\text{ext}} : \exists\, u \in \Sigma^*,\ x = u\top^*\} \cup \Sigma^\omega, \tag{28}$$

$$A_\perp \triangleq \{x \in \Sigma^\omega_{\text{ext}} : \exists\, u \in \Sigma^*,\ x = u\perp^*\} \cup \Sigma^\omega. \tag{29}$$

Let $A \triangleq A_\top \cup A_\perp$.

Its complement A^c is the set of sequences in which $\top$ or $\bot$ is followed at some point by a different symbol. More precisely,

$$A^c \;=\; \bigcup_{k=1}^{\infty} B_k^{(\top)} \;\cup\; \bigcup_{k=1}^{\infty} B_k^{(\bot)},$$

where

$$B_k^{(\top)} \triangleq \{x \in \Sigma_{\text{ext}}^{\omega} : x_k = \top \text{ and } x_{k+1} \neq \top\},$$
$$B_k^{(\bot)} \triangleq \{x \in \Sigma_{\text{ext}}^{\omega} : x_k = \bot \text{ and } x_{k+1} \neq \bot\}.$$

Consider $B_k^{(\top)}$. It is a finite disjoint union of cylinders of the form

$$C_{u,\sigma} \triangleq \{x \in \Sigma_{\text{ext}}^{\omega} : u\top\sigma \in \text{pref}(x)\},$$

with $u \in \Sigma_{\text{ext}}^{k-1}$ and $\sigma \in \Sigma_{\text{ext}} \setminus \{\top\}$. Thus

$$\widehat{\boldsymbol{P}}\left(B_k^{(\top)}\right) = \sum_{u,\sigma} \widehat{\boldsymbol{P}}(C_{u,\sigma}) = \sum_{u,\sigma} \boldsymbol{P}_{k+1}\{u\top\sigma\}.$$

By the definition of $\boldsymbol{P}_{k+1}$ and of $\mathcal{L}_{\text{ext}}$, once $\top$ is generated at position k we have $\mathcal{L}_{\text{ext}}(u\top) = \delta_{\top}$, and hence $\mathcal{L}_{\text{ext}}(u\top)(\sigma) = 0$ for all $\sigma \neq \top$. Therefore each summand is zero, and $\widehat{\boldsymbol{P}}\left(B_k^{(\top)}\right) = 0$. The same argument applies to $B_k^{(\bot)}$: $\widehat{\boldsymbol{P}}\left(B_k^{(\bot)}\right) = 0$ for all $k \geq 1$. By the union bound,

$$\widehat{\boldsymbol{P}}(A^c) \leq \sum_{k=1}^{\infty} \widehat{\boldsymbol{P}}\left(B_k^{(\top)}\right) + \sum_{k=1}^{\infty} \widehat{\boldsymbol{P}}\left(B_k^{(\bot)}\right) = 0,$$

so $\widehat{\boldsymbol{P}}(A) = 1$.

Definition of $\boldsymbol{P}_{\mathcal{L}}$. Define a measurable map $\phi : A \to \Sigma^* \cup \Sigma^{\omega} \cup \{\bot\}$ by:

1. If $x \in \Sigma^{\omega}$, set $\phi(x) = x$.
2. Otherwise let $k(x)$ be the first index where $x_k \in \{\top, \bot\}$.
   If $x_{k(x)} = \top$, set $\phi(x) = x_1 \cdots x_{k(x)-1}$.
   If $x_{k(x)} = \bot$, set $\phi(x) = \bot$.

We then define the desired probability measure on $\Sigma^* \cup \Sigma^{\omega} \cup \{\bot\}$. as the pushforward $\boldsymbol{P}_{\mathcal{L}} \triangleq \widehat{\boldsymbol{P}}_{\mathcal{L}} \circ \phi^{-1}$. Since $\widehat{\boldsymbol{P}}(A) = 1$, this is indeed a probability measure.

By construction, $\boldsymbol{P}_{\mathcal{L}}$ is Borel in the sense of the statement: its σ-algebra is generated by the cylinder sets coming from prefixes in Σ^* or from finite restrictions of infinite sequences in Σ^{ω}, together with the singleton $\{\bot\}$.

Prefix probabilities. Let $u \in \Sigma^*$ with length $k \geq 1$. The event

$$E_u \triangleq \{w \in \Sigma^* \cup \Sigma^\omega : u \in \mathrm{pref}(w)\}$$

corresponds under ϕ^{-1} to the cylinder

$$C_k \triangleq \{x \in \Sigma_{\mathrm{ext}}^\omega : u \in \mathrm{pref}(x)\},$$

because any sequence in A whose image under ϕ has u as prefix must itself start with u (once $\top$ or $\bot$ appears, it never disappears).

Hence, using $\widehat{\boldsymbol{P}}(A) = 1$,

$$\boldsymbol{P}_{\mathcal{L}}(E_u) = \widehat{\boldsymbol{P}}\left(\phi^{-1}(E_u)\right) = \widehat{\boldsymbol{P}}(C_k) = \boldsymbol{P}_k\{u\} = \prod_{i=1}^{k} \mathcal{L}_{\mathrm{ext}}(u_{<i})(u_i) = \prod_{i=1}^{k} \mathcal{L}(u_{<i})(u_i).$$

By definition of the prefix-probability function $P_{\mathcal{L}}$ in (3), this last expression equals $P_{\mathcal{L}}(u)$. Therefore,

$$\boldsymbol{P}_{\mathcal{L}}\{w \in \Sigma^* \cup \Sigma^\omega : u \in \mathrm{pref}(w)\} = P_{\mathcal{L}}(u),$$

which is exactly (7) when $u \neq \lambda$.

For $u = \lambda$, the event $\{w : \lambda \in \mathrm{pref}(w)\}$ is the whole space $\Sigma^* \cup \Sigma^\omega \cup \{\bot\}$, and $\boldsymbol{P}_{\mathcal{L}}$ assigns probability 1 to it, which equals $P_{\mathcal{L}}(\lambda)$ by definition.

Probability of a finite word. Let now $u \in \Sigma^*$ be fixed. Under ϕ, the point u corresponds to the unique sequence $u\top^\omega \in A$, namely: $\phi^{-1}\{u\} = \{u\top^\omega\}$. Indeed, if the first absorbing symbol in $x \in A$ is $\top$, then $\phi(x)$ is exactly the finite prefix before that $\top$, and no sequence of A containing $\bot$ in its tail is mapped to a finite word.

Moreover,

$$\{u\top^\omega\} = \bigcap_{n \geq 1} C_{u,n}, \qquad C_{u,n} \triangleq \{x \in \Sigma_{\mathrm{ext}}^\omega : u\top^n \in \mathrm{pref}(x)\}.$$

By continuity from above of $\widehat{\boldsymbol{P}}$,

$$\boldsymbol{P}_{\mathcal{L}}\{u\} = \widehat{\boldsymbol{P}}\left(\{u\top^\omega\}\right) = \lim_{n\to\infty} \widehat{\boldsymbol{P}}(C_{u,n}) = \lim_{n\to\infty} \boldsymbol{P}_{|u|+n}\{u\top^n\}.$$

From the definition of $\boldsymbol{P}_k$ and $\mathcal{L}_{\mathrm{ext}}$,

$$\boldsymbol{P}_{|u|+n}\{u\top^n\} = \left[\prod_{i=1}^{|u|} \mathcal{L}_{\mathrm{ext}}(u_{<i})(u_i)\right] \mathcal{L}_{\mathrm{ext}}(u)(\top) \prod_{j=1}^{n-1} \delta_\top(\top), \qquad n \geq 1.$$

If $\mathcal{L}(u) \in \mathfrak{D}(\Sigma_\top)$, then $\mathcal{L}_{\mathrm{ext}}(u)(\top) = \mathcal{L}(u)(\top)$. If $\mathcal{L}(u) = \bot$, then $\mathcal{L}_{\mathrm{ext}}(u) = \delta_\bot$ and $\mathcal{L}_{\mathrm{ext}}(u)(\top) = 0$. In both cases, by definition (4), the value above equals $P_{\mathcal{L}}^{\top}(u)$. Hence

$$\boldsymbol{P}_{\mathcal{L}}\{u\} = P_{\mathcal{L}}^{\top}(u),$$

which is exactly (8).

Probability of $\perp$. Finally, we compute the probability of $\perp$. Recall that

$$error(\mathcal{L}) = \{u \in \Sigma^* : \mathcal{L}(u) = \perp\}.$$

By definition of ϕ

$$\phi^{-1}\{\perp\} = \bigcup_{u \in error(\mathcal{L})} \{u\,\perp^\omega\},$$

and this union is disjoint since each $u\perp^\omega$ has a different finite prefix before the first occurrence of $\perp$. Thus, by countable additivity

$$\boldsymbol{P}_{\mathcal{L}}\{\perp\} = \widehat{\boldsymbol{P}}\left(\phi^{-1}\{\perp\}\right) = \sum_{u \in error(\mathcal{L})} \widehat{\boldsymbol{P}}\{u\,\perp^\omega\}.$$

As before

$$\{u\perp^\omega\} = \bigcap_{n \geq 1} C_{u,n}, \qquad C_{u,n} \triangleq \{x \in \Sigma^\omega_{\text{ext}} : u\perp^n \in \text{pref}(x)\}$$

and

$$\widehat{\boldsymbol{P}}\{u\perp^\omega\} = \lim_{n \to \infty} \boldsymbol{P}_{|u|+n}\{u\perp^n\}.$$

By the definition of $\mathcal{L}_{\text{ext}}$ on $error(\mathcal{L})$, once we reach such a u we have $\mathcal{L}_{\text{ext}}(u) = \delta_\perp$, so

$$\boldsymbol{P}_{|u|+n}\{u\perp^n\} = \left[\prod_{i=1}^{|u|} \mathcal{L}_{\text{ext}}(u_{<i})(u_i)\right] \mathcal{L}_{\text{ext}}(u)(\perp) \prod_{j=1}^{n-1} \delta_\perp(\perp) = \prod_{i=1}^{|u|} \mathcal{L}(u_{<i})(u_i).$$

By the definition of $P_{\mathcal{L}}$, this last expression equals $P_{\mathcal{L}}(u)$. Therefore

$$\boldsymbol{P}_{\mathcal{L}}\{\perp\} = \sum_{u \in error(\mathcal{L})} P_{\mathcal{L}}(u),$$

which is (9).

This completes the proof. ■

Appendix B. Proofs of Section 4

Recall that $\mathbb{1}_{\mathcal{L}}(u) = 1$ iff $P_{\mathcal{L}}(u) > 0$.

B.1. Proposition 5

Proof Let $u \in \Sigma^*$ be such that $\mathcal{L}(u) = \perp$. Let $\sigma \in \Sigma$. Def. (3) implies $\boldsymbol{P}_{\mathcal{L}}(u\sigma) = 0$. Then, by definition of $\mathbb{1}_{\mathcal{L}}$, we have $\mathbb{1}_{\mathcal{L}}(u\sigma) = 0$. Hence, Def. (10) implies $\mathcal{L}(u\sigma) = \perp$. ■

B.2. Proposition 6

Proof Let $\mathcal{L} : \Sigma^* \to \mathfrak{D}(\Sigma_\top)$ be a PLM. By Prop. 11 and Prop. 12, $\boldsymbol{P}_{\mathcal{L}} = \boldsymbol{P}_{\mathsf{za}(\mathcal{L})}$, and so, by Prop. 1, $\mathit{lang}(\mathcal{L}) = \mathit{lang}(\mathsf{za}(\mathcal{L}))$. ■

Proposition 11 *Let $\mathcal{L} : \Sigma^* \to \mathfrak{D}(\Sigma_\top)$. For all $u \in \Sigma^*$, $P_{\mathsf{za}(\mathcal{L})}(u) = P_{\mathcal{L}}(u)$.*

Proof By induction on $u \in \Sigma^*$. The case $u = \lambda$ is immediate by Def. (3). Let $u = w\sigma$. By I.H., $P_{\mathsf{za}(\mathcal{L})}(w) = P_{\mathcal{L}}(w)$. Suppose $\mathbb{1}_{\mathcal{L}}(w) = 0$. Then, $P_{\mathcal{L}}(w) = 0$, and so, by I.H., $P_{\mathsf{za}(\mathcal{L})}(w) = 0$. Then, by Def. (3), $P_{\mathsf{za}(\mathcal{L})}(w\sigma) = 0 = P_{\mathcal{L}}(w\sigma)$. Now, suppose $\mathbb{1}_{\mathcal{L}}(w) = 1$. Thus, by Eq. (13), $\mathsf{za}(\mathcal{L})(w) = \mathcal{L}(w)$. If $\mathcal{L}(w) = \bot$, then also $\mathsf{za}(\mathcal{L})(w) = \bot$. So, by Def. (3), $P_{\mathsf{za}(\mathcal{L})}(w\sigma) = P_{\mathcal{L}}(w\sigma) = 0$. If $\mathcal{L}(w) \neq \bot$, then also $\mathsf{za}(\mathcal{L})(w) \neq \bot$. Hence, $P_{\mathsf{za}(\mathcal{L})}(w\sigma) = P_{\mathsf{za}(\mathcal{L})}(w) \cdot \mathsf{za}(\mathcal{L})(w)(\sigma) = P_{\mathsf{za}(\mathcal{L})}(w) \cdot \mathcal{L}(w)(\sigma) = P_{\mathcal{L}}(w) \cdot \mathcal{L}(w)(\sigma) = P_{\mathcal{L}}(w\sigma)$. ■

Proposition 12 *Let $\mathcal{L} : \Sigma^* \to \mathfrak{D}(\Sigma_\top)$. For all $u \in \Sigma^*$, $P^{\top}_{\mathsf{za}(\mathcal{L})}(u) = P^{\top}_{\mathcal{L}}(u)$.*

Proof Let $u \in \Sigma^*$. By Prop. 11, $P_{\mathsf{za}(\mathcal{L})}(u) = P_{\mathcal{L}}(u)$. Then, by Def. (4), it follows that $P^{\top}_{\mathsf{za}(\mathcal{L})}(u) = P_{\mathcal{L}}(u) \cdot \mathsf{za}(\mathcal{L})(u)(\top)$. If $\mathbb{1}_{\mathcal{L}}(u) = 0$, then $P_{\mathcal{L}}(u) = 0$. If $\mathbb{1}_{\mathcal{L}}(u) = 1$, then by Def. (12), $\mathsf{za}(\mathcal{L})(u) = \mathcal{L}(u)$. Hence, $P^{\top}_{\mathsf{za}(\mathcal{L})}(u) = P^{\top}_{\mathcal{L}}(u)$. ■

Corollary 13 *For all PLM $\mathcal{L}$, $\mathsf{za}(\mathcal{L})$ is a ZAPLM.*

Proof Suppose $\mathbb{1}_{\mathsf{za}(\mathcal{L})}(u) = 0$. Then, by Prop. 11, $\mathbb{1}_{\mathcal{L}}(u) = 0$, and so, by Eq. (13), $\mathsf{za}(\mathcal{L})(u) = \bot$. Therefore, $\mathsf{za}(\mathcal{L})$ satisfies Def. (10). ■

Corollary 14 *For any ZAPLM $\mathcal{L} : \Sigma^* \to \mathfrak{D}(\Sigma_\top)$, $\mathit{lang}(\mathcal{L}) = \{u \in \Sigma^* \mid \mathsf{accept}_\bot(\mathcal{L})(u) = 1\}$, and $\mathit{error}(\mathcal{L}) = \{u \in \Sigma^* \mid \mathsf{accept}_\bot(\mathcal{L})(u) = \bot\}$.*

Proof By Def. (10), $\mathcal{L}(u)(\top) > 0$ implies $\mathbb{1}_{\mathcal{L}}(u) = 1$. Then, $P^{\top}_{\mathcal{L}}(u) > 0$, and so $u \in \mathit{lang}(\mathcal{L})$. Conversely, $u \in \mathit{lang}(\mathcal{L})$ then $P^{\top}_{\mathcal{L}}(u) > 0$, and so $\mathcal{L}(u)(\top) > 0$. The proof for $\mathit{error}(\mathcal{L})$ is similar. The result holds for every PLM $\mathcal{L}$ because $\mathsf{za}(\mathcal{L})$ is a ZAPLM. ■

B.3. Proposition 7

Proof By Def. (16) and Def. (3), for all $u \in \mathit{error}(\mathcal{L})$ such that $u \neq \lambda$, $P_{\mathcal{L}}(u) = 0$. Then, if $\lambda \in \mathit{error}(\mathcal{L})$, by Prop. 3(9), $\boldsymbol{P}_{\mathcal{L}}\{\bot\} = P_{\mathcal{L}}(\lambda) = 1$, otherwise $\boldsymbol{P}_{\mathcal{L}}\{\bot\} = 0$. ■

B.4. Proposition 8

Proof Prop. 15 implies $lang(\mathcal{L}) \subseteq lang(\mathsf{zo}(\mathcal{L}))$. Let $u \in lang(\mathsf{zo}(\mathcal{L}))$. So, $P^{\top}_{\mathsf{zo}(\mathcal{L})}(u) > 0$. Then, by Def. (17) and Def. (18), $\mathcal{L}(u)(\top) > 0$, and $\mathcal{L}(v)(\sigma) > 0$ for all $v\sigma \in \mathrm{pref}(u)$. Therefore, $P_{\mathcal{L}}(u) > 0$, and also $P^{\top}_{\mathcal{L}}(u) > 0$. Hence, $u \in lang(\mathcal{L})$. ■

Proposition 15 *Let $\mathcal{L} : \Sigma^* \to \mathfrak{D}(\Sigma_\top)$. For all $u \in lang(\mathcal{L})$, $P^{\top}_{\mathsf{zo}(\mathcal{L})}(u) > 0$.*

Proof Let $u \in lang(\mathcal{L})$. By Prop. 1, $P^{\top}_{\mathcal{L}}(u) > 0$. So, by Def. (18), $\mathsf{zom}[\mathcal{L}](u)(\top) = 1$, and by Def. (17), $\mathsf{zo}(\mathcal{L})(u)(\top) > 0$. Let $u = \lambda$. Then, Def. (3) implies $P_{\mathcal{L}}(\lambda) = 1 = P_{\mathsf{zo}(\mathcal{L})}(\lambda)$. Now, let $u = v\sigma$. By hypothesis $\mathcal{L}(v\sigma) \neq \bot$ and also $\mathcal{L}(v)(\sigma) > 0$. Then, Def. (18) implies $\mathsf{zom}[\mathcal{L}](u)(\sigma) = 1$, and so Def. (17) implies $\mathsf{zo}(\mathcal{L})(v)(\sigma) > 0$. Suppose, by contradiction, that $P_{\mathsf{zo}(\mathcal{L})}(v) = 0$. Then, there exists $w\sigma' \in \mathrm{pref}(v)$ such that $P_{\mathsf{zo}(\mathcal{L})}(w) > 0$ and $P_{\mathsf{zo}(\mathcal{L})}(w\sigma') = 0$. So, $\mathsf{zo}(\mathcal{L})(w)(\sigma') = 0$ or $\mathsf{zo}(\mathcal{L})(w) = \bot$. Since by hypothesis $\mathcal{L}(w)(\sigma') > 0$, Def. (17) implies $\mathsf{zom}[\mathcal{L}](w)(\sigma') = 0$, and so by Def. (18), $\mathcal{L}(w\sigma') = \bot$. Then, $u \notin lang(\mathcal{L})$, which contradicts the hypothesis. Therefore, $P_{\mathsf{zo}(\mathcal{L})}(v)$ must be positive. Then, $P_{\mathsf{zo}(\mathcal{L})}(v\sigma) > 0$. Hence, $P^{\top}_{\mathsf{zo}(\mathcal{L})}(u) > 0$, for all $u \in lang(\mathcal{L})$, since $\mathsf{zom}[\mathcal{L}](u)(\top) = 1$. ■

B.5. Proposition 9

Proof Let $u\sigma \in \Sigma^+$ with $\mathcal{L}(u\sigma) = \bot$. Suppose $\mathbb{1}_{\mathcal{L}}(u\sigma) = 1$. By Def. (18), $\mathsf{zom}[\mathcal{L}](u)(\sigma) = 0$, and so by Def. (17), $\mathsf{zo}(\mathcal{L})(u)(\sigma) = 0$. Then, by hypothesis, $\mathcal{L}(u)(\sigma) = 0$. Therefore, $P_{\mathcal{L}}(u\sigma) = 0$, and so $\mathbb{1}_{\mathcal{L}}(u\sigma) = 0$, which is a contradiction. Hence, $\mathcal{L}$ satisfies Def. (16). ■

B.6. Proposition 10

Proof Since $[\Sigma_\top \to \{0,1\}]$ is finite, for every $u \in \Sigma^*$, there exists i_u, such that $\mathcal{M}_{i_u+1}(u) = \mathcal{M}_{i_u}(u)$. Then,

$$\mathsf{m}[\mathcal{L}](\mathcal{M}_\infty)(u) = \mathsf{m}[\mathcal{L}](\mathcal{M}_{i_u})(u) = \mathcal{M}_{i_u+1}(u) = \mathcal{M}_{i_u}(u) = \mathcal{M}_\infty(u).$$

So, $\mathcal{M}_\infty$ is a fixpoint of $\mathsf{m}[\mathcal{L}]$. Moreover,

$$\begin{aligned}\mathsf{zo}(\mathsf{norm}(\mathcal{L} \odot \mathcal{M}_\infty)) &= \mathsf{norm}(\mathsf{norm}(\mathcal{L} \odot \mathcal{M}_\infty) \odot \mathsf{zom}[\mathcal{L}]) \\ &= \mathsf{norm}(\mathsf{norm}(\mathcal{L} \odot \mathcal{M}_\infty) \odot \mathsf{m}[\mathcal{L}](I)) \\ &= \mathsf{norm}(\mathsf{norm}(\mathcal{L} \odot \mathcal{M}_\infty)) = \mathsf{norm}(\mathcal{L} \odot \mathcal{M}_\infty).\end{aligned}$$

Therefore, $\mathsf{norm}(\mathcal{L} \odot \mathcal{M}_\infty)$ is a fixpoint of zo. ■