

Limitations of Autoregressive Models and Their Alternatives

Chu-Cheng Lin^{#*} Aaron Jaech^b Xin Li[#] Matthew R. Gormley^h Jason Eisner[#]

[#]Department of Computer Science, Johns Hopkins University

^bFacebook AI

^hMachine Learning Department, Carnegie Mellon University

{kitsing,lixints,jason}@cs.jhu.edu ajaech@fb.com mgormley@cs.cmu.edu

Abstract

Standard autoregressive language models perform only polynomial-time computation to compute the probability of the next symbol. While this is attractive, it means they cannot model distributions whose next-symbol probability is *hard* to compute. Indeed, they cannot even model them well enough to solve associated *easy* decision problems for which an engineer might want to consult a language model. These limitations apply no matter how much computation and data are used to train the model, unless the model is given access to oracle parameters that grow *superpolynomially* in sequence length.

Thus, simply training larger autoregressive language models is not a panacea for NLP. Alternatives include energy-based models (which give up efficient sampling) and latent-variable autoregressive models (which give up efficient scoring of a given string). Both are powerful enough to escape the above limitations.

1 Introduction

Sequence modeling is a core NLP problem. Many sequence models \tilde{p} are efficient at *scoring strings*: given a string \mathbf{x} , its score $\tilde{p}(\mathbf{x})$ can be computed in $O(\text{poly}(|\mathbf{x}|))$. For example, an RNN (Mikolov et al., 2011) scores \mathbf{x} in time $O(|\mathbf{x}|)$ while a Transformer (Vaswani et al., 2017) does so in time $O(|\mathbf{x}|^2)$. The score may be an unnormalized probability, and can be used to rank candidate strings.

Many sequence models also make it easy to compute marginal properties of \tilde{p} . They support efficient *sampling* of strings \mathbf{x} (which allows unbiased approximation of marginal expectations). And they support efficient computation of the *normalizing constant* $Z = \sum_{\mathbf{x}} \tilde{p}(\mathbf{x})$ (or simply guarantee $Z = 1$) for any value of the model parameters.

How about training? Briefly: If a sequence model can efficiently compute $\tilde{p}(\mathbf{x})$ (and its derivatives

*Part of this work was done at Facebook AI.

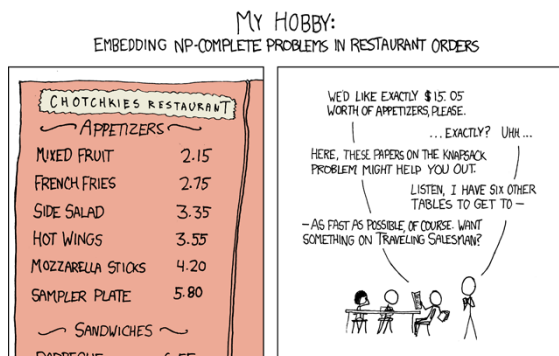


Figure 1: Valid answers to hard natural language inference problems can be hard to find (Munroe, 2009), but in many cases can be checked efficiently (e.g. the KNAPSACK problem here). Given a *large enough* parametric autoregressive model with *correct* parameters, we can efficiently solve *all* problem instances with input length n , and efficiently verify the solutions — but the required model size can grow superpolynomially in n . (This allows the model to store precomputed results that we can compare against at test time, in $O(n)$.) A main observation of this paper is that assuming $NP \not\subseteq P/\text{poly}$, then without such a superpolynomial growth in model size, autoregressive models cannot even be used to verify answers to some problems where polynomial-time verification algorithms exist.

with respect to model parameters), then it is efficient to compute parameter updates for noise-contrastive estimation (Gutmann and Hyvärinen, 2010; Gutmann and Hyvärinen, 2012) or score-matching (Hyvärinen, 2005). If sampling \mathbf{x} or computing Z (and its derivatives) is *also* efficient, then it is efficient to compute parameter updates for ordinary MLE training.

Finally, popular sequence models are *compact*. Usually a fixed-size model is used to score strings \mathbf{x} of all lengths. More generally, it might be reasonable to use an $O(\text{poly}(n))$ -sized parameter vector θ_n when \mathbf{x} has length n , at least if parameter vectors can be obtained (perhaps from an oracle) for all needed lengths. In this paper, we investigate what can and cannot be achieved with models that are compact in this sense. This setup allows us to discuss the asymptotic behavior of model families.

Standard autoregressive models have the form $p(\mathbf{x}) = \prod_t p(x_t | \mathbf{x}_{<t})^1$ where each factor is ef-

¹In this paper we use the shorthand $\mathbf{x}_{<t} \triangleq x_1 \dots x_{t-1}$.

Model family	Compact parameters?	Efficient scoring?	Efficient sampling and normalization?	Support can be ...
ELN/ELNCP: Autoregressive models (§3.1)	✓	✓	✓	<i>some but not all</i> $L \in \mathsf{P}$
EC/ECCP: Energy-based models (§4.1)	✓	✓	✗	<i>all</i> $L \in \mathsf{P}$ but <i>no</i> $L \in \mathsf{NPC}$
Lightly marginalized ELNCP: Latent-variable autoregressive models (§4.2)	✓	✗	✓	<i>all</i> $L \in \mathsf{NP}$
Lookup models (§4.3)	✗	✓	✓	<i>anything</i>

Table 1: A feature matrix of parametric model families discussed in this paper. Also see Figure 2 in the appendices.

ficient to compute from a fixed parameter vector. These models satisfy all three of the desiderata above. Through the use of flexible neural network architectures, standard autoregressive models have achieved stellar empirical results in many applications (Oord et al., 2016; Child et al., 2019; Zellers et al., 2019; Brown et al., 2020). However there are still tasks that they have not mastered: for example it is reported that they struggle at deep logical structure, even with the help of huge pretrained models (Wang et al., 2019a).

We point out that unfortunately, there are certain sequence distributions whose unnormalized string probabilities $\tilde{p}(\mathbf{x})$ are *easy* to compute, yet whose autoregressive factors $p(x_t \mid \mathbf{x}_{<t})$ are *hard* (NP-hard) to compute or even approximate. Thus, standard autoregressive models are *misspecified* for these distributions (cannot fit them). It does not help much to focus on strings of bounded length, or to enlarge the model: under the common complexity-theoretic assumption $\mathsf{NP} \not\subseteq \mathsf{P/poly}$, the parameter size $|\theta_n|$ must grow *superpolynomially* in n to model the probabilities of all strings of length up to n .

Indeed, one of our main findings is that there exist unweighted languages $L \in \mathsf{P}$ for which *no* standard autoregressive model has L as its support, i.e., assigns weight > 0 to just the strings $\mathbf{x} \in L$. This is downright depressing, considering the costs invested in training huge parametric autoregressive models (Bender et al., 2021): since $L \in \mathsf{P}$, it is trivial to build an efficient scoring function $\tilde{p}(\mathbf{x})$ with fixed parameters that has L as its support, just not an autoregressive one. The problem holds for *all* standard autoregressive models, regardless of how much computation and training data are used to learn the model parameters.

That is, for an NP-hard problem, scoring strings under a standard autoregressive model $p(\mathbf{x})$ cannot be used to *verify* witnesses. Nor can *finding* witnesses be solved by prompting such a model with a description of a problem instance and sampling a continuation of that string. Such problems are abundant in NLP: for example surface realization under Optimality Theory (Idsardi, 2006), decoding text from an AMR parse (Cai and Knight, 2013), phrase alignment between two sentences (DeNero and Klein, 2008), and in

general inference of propositional logic (Cook, 1971), which illustrates the NP-hardness of general natural language inference, as in Figure 1. In other words, our results imply that standard autoregressive models do not have the right structure to capture important linguistic regularities: e.g., that observed sequences were *in fact* constructed to be phonologically optimal, expressive of a semantic form, or logically coherent!

Our work is also relevant to autoregressive models of fixed-dimensional vectors, such as NADE (Uria et al., 2016). These models can be extended to arbitrary n -dimensional vectors by providing separate parameters θ_n for each n . However, our constructions imply that for some distributions, $|\theta_n|$ must grow superpolynomially in n , even though this would be not be necessary if the models were not autoregressive.

In the remainder of this paper, we formalize our three desiderata for sequence models. We formalize compact autoregressive models and describe some limitations on their expressiveness. We then show that it can help to choose an alternative model family that relaxes any one of the three desiderata (Table 1).

2 Background

2.1 Weighted languages

An **unweighted language** $L \subseteq V^*$ is a set of strings \mathbf{x} over a finite alphabet V . A **weighted language** \tilde{p} is a function $\tilde{p} : V^* \rightarrow \mathbb{R}_{\geq 0}$. It may be regarded as specifying an unweighted language $L = \text{support}(\tilde{p}) \triangleq \{\mathbf{x} : \tilde{p}(\mathbf{x}) \neq 0\}$ along with positive weights for the strings in L . We say that a weighted language \tilde{p} is **normalizable** if its **global normalizing constant** $Z \triangleq \sum_{\mathbf{x} \in V^*} \tilde{p}(\mathbf{x})$ is finite and strictly positive. When \tilde{p} is normalizable, $p(\mathbf{x}) \triangleq \tilde{p}(\mathbf{x})/Z$ is a probability distribution over L . A **distribution** is any weighted language whose global normalizing constant is 1.

Let $\hat{\mathbf{x}} \leq \mathbf{x}$ mean that $\hat{\mathbf{x}}$ is a prefix of $\mathbf{x} \in V^*$ (not necessarily a strict prefix). If \tilde{p} is normalizable, then $Z(\hat{\mathbf{x}}) \triangleq \sum_{\mathbf{x} \in V^* : \hat{\mathbf{x}} \leq \mathbf{x}} \tilde{p}(\mathbf{x})$ is $\leq Z$ for any $\hat{\mathbf{x}} \in V^*$, yielding a marginal **prefix probability** $Z(\hat{\mathbf{x}})/Z$. If the prefix $\hat{\mathbf{x}}$ has positive prefix probability, then it admits a **local conditional probability** $p(x \mid \hat{\mathbf{x}}) \triangleq Z(\hat{\mathbf{x}}x)/Z(\hat{\mathbf{x}})$ for each symbol $x \in V$, where the denominator is interpreted as a **local normalizing**

constant. This is the conditional probability that if a random string starts with the prefix $\hat{\mathbf{x}}$, the next symbol is x . There is also a probability $p(\$ | \hat{\mathbf{x}}) \triangleq 1 - \sum_{x \in V} p(x | \hat{\mathbf{x}}) = \tilde{p}(\hat{\mathbf{x}})/Z(\hat{\mathbf{x}}) \geq 0$ that the string ends immediately after $\hat{\mathbf{x}}$; the special symbol $\$ \notin V$ represents “end of string.”

2.2 Computation for weighted languages

We define a weighted language \tilde{p} to be **computable** if it is defined by a Turing machine (also called \tilde{p}) that maps any $\mathbf{x} \in V^*$ to $\tilde{p}(\mathbf{x}) \in \mathbb{Q}_{\geq 0}$ in finite time. The Turing machine does not have to compute Z .

While the computable weighted languages allow any computable function as \tilde{p} , most architectures for defining weighted languages (e.g., RNNs or Transformers) do only a bounded or linear amount of work per input symbol. As a result, they compute $\tilde{p}(\mathbf{x})$ in time $O(\text{poly}(|\mathbf{x}|))$. That is, $\tilde{p} \in \text{FP}$; we refer to such weighted languages as **efficiently computable (EC)**. This does not imply that the normalized version p is efficiently computable, since finding the denominator Z requires summing over all of V^* .

If we tried to construct the same normalized distribution p as in the previous paragraph using a standard autoregressive model, we would model it as a product of local conditional probabilities, $p(\mathbf{x}) = (\prod_{t=1}^{|\mathbf{x}|} p(x_t | \mathbf{x}_{<t}))p(\$ | \mathbf{x})$. Most such architectures again do only a bounded or linear amount of work per input symbol. Yet one suspects that this may not always be enough work to do the job: the local conditional probabilities of the original \tilde{p} are expensive to compute (unless \tilde{p} has some special structure making $Z(\hat{\mathbf{x}})$ tractable).

Indeed, the observation of this paper is that for some efficiently computable weighted languages \tilde{p} , the local conditional probabilities are expensive to compute or even to approximate well. More precisely, autoregressive models cannot fit the local conditional probabilities unless they are superpolynomial in either their runtime or in their number of parameters (where the parameters may be precomputed at training time). We now explain how to formalize these notions.

2.3 Non-uniform computation

In the machine learning approach to sequence modeling, we usually do not manually design the Turing machine behind \tilde{p} . Rather, we design a model M with *parameters* θ . M is a Turing machine that reads θ and outputs a specialized Turing machine $\tilde{p}_\theta \triangleq M(\theta)$ that can score strings \mathbf{x} and hence defines a weighted language. Without loss of generality, we will express

θ as a string in \mathbb{B}^* (where $\mathbb{B} \triangleq \{0, 1\}$). For each θ , we obtain a potentially different weighted language.

Strings vary in length, and accurate modeling of longer strings may sometimes require more complex computations with more parameters. For example, when V is a natural language alphabet, a recurrent neural network may require more hidden units to model sentences of the language rather than individual words, and even more units to model whole documents. To accommodate this, we allow an *infinite sequence* of parameter vectors, $\Theta = \{\theta_n \in \mathbb{B}^* \mid n \in \mathbb{N}\}$, which yields an infinite sequence of Turing machines $\{\tilde{p}_n \mid n \in \mathbb{N}\}$ via $\tilde{p}_n \triangleq M(\theta_n)$. We then define $\tilde{p}_\Theta(\mathbf{x}) \triangleq \tilde{p}_{|\mathbf{x}|}(\mathbf{x})$, so a string of length n is scored by the \tilde{p}_n machine. This is known as **non-uniform computation**. Of course, it is legal (and common) for all of the θ_n to be equal, or empty, but if desired, we can obtain more power by allowing the number of parameters to grow with n if needed.

We can now consider *how rapidly* the parametric and runtime complexity may grow.

- If $|\theta_n|$ is permitted to grow exponentially, then one can fit *any* weighted language \tilde{p} (even an uncomputable one).² Simply use θ_n to encode a trie with $O(|V|^{n+1})$ nodes that maps $\mathbf{x} \mapsto \tilde{p}(\mathbf{x})$ for any $|\mathbf{x}|$ of length n , and design M such that the Turing machine $\tilde{p}_n = M(\theta_n)$ has a (large) state transition table that mirrors the structure of this trie. The resulting collection of Turing machines $\{\tilde{p}_n \mid n \in \mathbb{N}\}$ can then compute $\tilde{p}(\mathbf{x})$ exactly for any \mathbf{x} , with only linear runtime $O(|\mathbf{x}|)$ (which is used to traverse the trie).
- Separately, if unbounded runtime is permitted for M , then one can exactly fit *any computable* weighted language \tilde{p} . Simply have M , when run on θ_n , *compute* and return the large trie-structured \tilde{p}_n that was mentioned above. In this case, M need not even use the parameters θ_n , except to determine n .
- Finally, if unbounded runtime is permitted for \tilde{p}_n , then again one can exactly fit *any computable* weighted language \tilde{p} . In this case, M trivially returns $\tilde{p}_n = \tilde{p}$ for all n .
- However, if the parameters Θ are “compact” in the sense that $|\theta_n|$ grows only as $O(\text{poly}(n))$, and also $\tilde{p}_n = M(\theta_n)$ is constructed by M in time $O(\text{poly}(n))$, and \tilde{p}_n scores any \mathbf{x} of length n in time $O(\text{poly}(n))$, then we say that the resulting weighted language \tilde{p} is **efficiently computable with compact parameters (ECCP)**.³ We refer

²See our remark on computability in Appendix A.

³Since we require M to run in polytime, it can only look at a polynomial-sized portion of θ_n . Hence it is not really crucial for

to M paired with a parameter space of possible compact values for Θ as an **ECCP model**.

Neural models of weighted languages are typically ECCP models. The construction and execution of the neural network \tilde{p}_n may perform a polynomial amount of total computation to score the string \mathbf{x} . This computation may involve parameters that were precomputed using any amount of effort (e.g., training on data) or even obtained from an oracle (they need not be computable). However, the exponentially many strings of length n must share a polynomial-size parameter vector θ_n , which prevents the solution given in the first bullet point above.

In practice one takes $\theta_n = \theta$ for all n and obtains $\theta \in \mathbb{R}^d$ by training. However, we do not consider whether such parameters are easy to estimate or even computable. We simply ask, for a given target language \tilde{p} , whether there *exists* a polynomially growing sequence Θ of “good” parameter vectors for any parametric model M . When not, there can be no scheme for estimating arbitrarily long finite prefixes of such a sequence. So for any polynomial f , any training scheme that purports to return a trained model of size $f(n)$ that works “well” for strings of length $\leq n$ must fail for large enough n — even if unlimited data, computation, and oracles are allowed at training time.

2.4 P, P/poly, and NP/poly

The phrase “efficiently computable with compact parameters” means that without access to those parameters, the ECCP weighted language may no longer be efficiently computable. Indeed, it need not be computable at all, if the parameter vectors store the outputs of some uncomputable function.

Our definitions above of EC and ECCP weighted languages are weighted generalizations of complexity classes P and P/poly, respectively,⁴ and their supports are always unweighted languages in P and P/poly, respectively. An unweighted language L is in P iff there is a deterministic Turing machine that decides in $O(\text{poly}(|\mathbf{x}|))$ time whether $\mathbf{x} \in L$. And an unweighted language L' is in P/poly iff⁵ there exist

the parameters θ_n^P to be compact, but we nonetheless include this intuitive condition, without loss of generality.

⁴Namely the nonnegative functions in FP and FP/poly.

⁵Our presentation of P/poly is a variant of [Arora and Barak \(2009, §6\)](#), in which inputs \mathbf{x} of length n are evaluated by a polytime function M that is given an advice string θ_n as an auxiliary argument. This corresponds to a neural architecture M that can consult trained parameters θ_n at runtime. We have replaced the standard call $M(\theta_n, \mathbf{x})$ with the “curried” expression $M(\theta_n)(\mathbf{x})$, which we still require to execute in polynomial total time. Here the intermediate result $M_n = M(\theta_n)$ corresponds to a

Turing machines $\{M_n : n \in \mathbb{N}\}$ such that M_n decides in $O(\text{poly}(n))$ time whether \mathbf{x} of length n is in L' , where each M_n can be constructed in $O(\text{poly}(n))$ time as $M(\theta_n)$, for some Turing machine M and some sequence of polynomially-sized **advice strings** $\Theta = \{\theta_n \mid n \in \mathbb{N}\}$ with $|\theta_n| \in O(\text{poly}(n))$. We define the language class NP/poly similarly to P/poly: the only difference is the family $\{M_n : n \in \mathbb{N}\}$ consists of *nondeterministic Turing machines*.

Naturally, $P \subseteq P/\text{poly}$. But P/poly is larger than P: it contains all sparse languages, regardless of their hardness — even sparse undecidable languages — as well as many dense languages. The extra power of P/poly comes from its access to compact advice strings that do not have to be recursively enumerable, let alone efficient to find. This corresponds to statistical modeling, where the trained model has a computationally efficient architecture plus access to parameters that might have taken a long time to find.

2.5 NP-completeness and SAT

NP-complete decision problems have solutions that are efficient to validate but inefficient to find (assuming $P \neq \text{NP}$). One of the most well-known NP-complete problems is the boolean satisfiability problem (SAT) ([Cook, 1971](#)). Given a boolean formula ϕ , SAT accepts ϕ iff ϕ can be satisfied by some value assignment. For example, the formula $(A_1 \vee \neg A_2 \vee A_3) \wedge (A_1 \vee \neg A_4)$ is in SAT, since there is a satisfying assignment $A_{1..4} = 1101$. We denote the number of satisfying assignments to ϕ as $\#(\phi)$.

It is widely believed that no NP-complete languages are in P/poly. Otherwise we would have all of $\text{NP} \subseteq P/\text{poly}$ and the polynomial hierarchy would collapse at the second level ([Karp and Lipton, 1980](#)).

A capacity limitation of EC/ECCP weighted languages naturally follows from this belief:⁶

Lemma 1. *For any $L \in P$, there exists an EC weighted language with support L . For any $L \in P/\text{poly}$, there exists an ECCP language with support L . But for any $L \in \text{NP-complete}$, there exists no ECCP language with support L (assuming $\text{NP} \not\subseteq P/\text{poly}$).*

In addition to not capturing the support of NP-complete languages, ECCP languages cannot help

trained runtime model for inputs of length n . Our Turing machines M_n have size polynomial in n (because they are constructed by M in polynomial time). They correspond to the polynomial-sized boolean circuits M_n that are used to evaluate inputs of length n under the classical definition of P/poly ([Ladner, 1975](#)). We exposed these intermediate results M_n only to observe in [§2.3](#) and [§4.3](#) that if we had allowed the M_n to grow exponentially, they would have been able to encode the answers in tries.

⁶All omitted proofs are in Appendix A.

solve other NP-hard problems, either. For example, many structured prediction problems in NLP can be formulated as $\operatorname{argmax}_{\mathbf{x}: \hat{\mathbf{x}} \leq \mathbf{x}} \tilde{p}(\mathbf{x})$: we are given a prefix $\hat{\mathbf{x}}$ as input and look for its optimal continuation under \tilde{p} . But if this problem is NP-hard for a particular \tilde{p} , then it is not in P/poly (assuming $\text{NP} \not\subseteq \text{P/poly}$), so it cannot be accomplished by any polytime algorithm that queries an ECCP model.

3 Autoregressive ECCP models (ELNCP models) have reduced capacity

In this section we formally define *autoregressive* ECCP languages, and prove that they have strictly less capacity than general ECCP languages or even just EC languages. Our proofs rely on the construction of a EC language \tilde{p} where computing the local conditional probabilities $p(x | \hat{\mathbf{x}})$ is NP-hard, so they cannot be computed with compact parameters, if $\text{NP} \not\subseteq \text{P/poly}$.

3.1 ELN and ELNCP models

Many parameter estimation techniques and inference methods specifically work with local conditional probabilities $p(x | \hat{\mathbf{x}})$. Thus, it is common to use parametric models where such quantities can be computed in time $O(\text{poly}(|\hat{\mathbf{x}}|))$ (given the parameters).⁷ These are the “standard autoregressive models” we discussed in §1. We say that the resulting distributions are **efficiently locally normalizable**, or **ELN**.

We may again generalize to allow the use of compact parameters. For any weighted language \tilde{p} , the Turing machine M^q **efficiently locally normalizes** \tilde{p} **with compact parameters** $\Theta^q = \{\theta_n^q | n \in \mathbb{N}\}$ if

- the parameter size $|\theta_n^q|$ grows only as $O(\text{poly}(n))$
- $M^q(\theta_n^q)$ returns a Turing machine q_n (similar to \tilde{p}_n in §2.3) in time $O(\text{poly}(n))$
- \tilde{p} is normalizable (so p exists)
- q_n maps $\hat{\mathbf{x}}x \mapsto p(x | \hat{\mathbf{x}})$ for all $x \in V \cup \{\$\}$ and all prefixes $\hat{\mathbf{x}} \in V^*$ with $|\hat{\mathbf{x}}| \leq n$ and $Z(\hat{\mathbf{x}}) > 0$

⁷An autoregressive model architecture generally defines $p(\mathbf{x})$ as an efficiently computable (§2.2) product of local conditional probabilities. However, the parametrization usually ensures only that $\sum_{x \in V} p_\theta(x | \hat{\mathbf{x}}) = 1$ for all prefixes $\hat{\mathbf{x}}$. Some parameter settings may give rise to **inconsistent** distributions where $Z \triangleq \sum_{\mathbf{x} \in V^*} p_\theta(\mathbf{x}) < 1$ because the generative process terminates with probability < 1 (Chen et al., 2018). In this case, the factors $p_\theta(x | \hat{\mathbf{x}})$ defined by the autoregressive model are not actually the conditional probabilities of the weighted language (as defined by §2.1). It is true that training θ with a likelihood objective does encourage finding a weighted language whose generative process always terminates (hence $Z = 1$), since this is the behavior observed in the training corpus (Chi and Geman, 1998; Chen et al., 2018; Welleck et al., 2020). In this paper, we are only concerned with such languages, since we require the *actual* conditional probabilities to be efficiently computable. Autoregressive models that do not sum to 1, whose normalized probabilities can be uncomputable, are not ruled out by our theorems in this section.

- q_n runs on those inputs $\hat{\mathbf{x}}x$ in time $O(\text{poly}(n))$

If there is M^q that efficiently locally normalizes a weighted language \tilde{p} with compact parameters Θ^q , we say \tilde{p} is **efficiently locally normalizable with compact parameters**, or **ELNCP**. Note that this is a property of the weighted language itself. In this case, it is not hard to see that $p = \tilde{p}/Z$ is ECCP:

Lemma 2. *An ELNCP model \tilde{p} is also ECCP. Likewise, an ELN model is also EC.*

If we define ELNCP *models* analogously to ECCP models, Lemma 2 means that locally normalized models do not provide any extra power. Their distributions can always be captured by globally normalized models (of an appropriate architecture that we used in the proof). But we will see in Theorem 1 that the converse is likely not true: provided that $\text{NP} \not\subseteq \text{P/poly}$, there are efficiently computable weighted languages that cannot be efficiently locally normalized, even with the help of compact parameters. That is, they are EC (hence ECCP), yet they are not ELNCP (hence not ELN).

3.2 ELNCP models cannot exactly capture all EC (or ECCP) distributions

We prove our claim by defining a certain weighted language \tilde{p} and reducing SAT to computing certain local conditional probabilities of \tilde{p} (as defined in §2.1). Each decision $\text{SAT}(\phi)$ (where ϕ ranges over formulas) corresponds to a particular local conditional probability, implying that there is no polytime scheme for computing all of these probabilities, even with polynomially sized advice strings (i.e., parameters).

Without loss of generality, we consider only formulae ϕ such that the set of variables mentioned at least once in ϕ is $\{A_1, \dots, A_j\}$ for some $j \in \mathbb{N}$; we use $|\phi|$ to denote the number of variables j in ϕ . We say that **a satisfies** ϕ if $\mathbf{a} \in \mathbb{B}^{|\phi|}$ and $(A_1 = a_1, \dots, A_{|\phi|} = a_{|\phi|})$ is a satisfying assignment. Finally, let boldface $\phi \in \mathbb{B}^*$ denote $\text{enc}(\phi)$ where enc is a prefix-free encoding function. We can now define the unweighted language $L = \{\phi \mathbf{a} | \phi \text{ is a formula and } \mathbf{a} \in \mathbb{B}^{|\phi|} \text{ and } \mathbf{a} \text{ satisfies } \phi\}$ over alphabet \mathbb{B} , which contains each possible SAT problem concatenated to each of its solutions.⁸

We now convert L to a weighted language \tilde{p} , defined by $\tilde{p}(\mathbf{x}) = \tilde{p}(\phi, \mathbf{a}) = (\frac{1}{3})^{|\mathbf{x}|+1}$ for $\mathbf{x} \in L$ (otherwise $\tilde{p}(\mathbf{x}) = 0$). \tilde{p} is normalizable since Z is both finite ($Z = \sum_{\mathbf{x} \in \mathbb{B}^*} \tilde{p}(\mathbf{x}) \leq \sum_{\mathbf{x} \in \mathbb{B}^*} (\frac{1}{3})^{|\mathbf{x}|+1} = 1$) and positive ($Z > 0$ because the example string in footnote 8 has weight > 0). The conditional distribution

⁸For example, L contains the string $\phi \mathbf{a}$ where $\phi = \text{enc}((A_1 \vee \neg A_2 \vee A_3) \wedge (A_1 \vee \neg A_4))$ and $\mathbf{a} = 1101$.

$p(\mathbf{a} \mid \phi)$ is uniform over the satisfying assignments \mathbf{a} of ϕ , as they all have the same length $|\phi|$.

\tilde{p} is efficiently computable, and so is $p = \tilde{p}/Z$.⁹ Yet deciding whether the local conditional probabilities of \tilde{p} are greater than 0 is NP-hard. In particular, we show that SAT can be reduced to deciding whether certain local probabilities are greater than 0, namely the ones that condition on prefixes $\hat{\mathbf{x}}$ that consist only of a formula: $\hat{\mathbf{x}} = \phi$ for some ϕ . This implies, assuming $\text{NP} \not\subseteq \text{P/poly}$, that no (M^q, Θ^q) can efficiently locally normalize \tilde{p} with compact parameters. Granted, the restriction of \tilde{p} to the finite set $\{\mathbf{x} \in \mathbb{B}^* : |\mathbf{x}| \leq n\}$ can be locally normalized by some polytime Turing machine q_n , using the same trie trick sketched in §2.3. But such tries have sizes growing exponentially in n , and it is not possible to produce a sequence of such machines, $\{q_n : n \in \mathbb{N}\}$, via a single master Turing machine M^q that runs in $O(\text{poly}(n))$ on θ_n^q . That is:

Theorem 1. *Assuming $\text{NP} \not\subseteq \text{P/poly}$, there exists an efficiently computable normalizable weighted language \tilde{p} that is not ELNCP.*

Proof sketch. Take \tilde{p} to be the weighted language we defined earlier in this section. \tilde{p} is clearly efficiently computable. We will show that if it is ELNCP via (M^q, Θ^q) , then the NP-complete problem SAT is in P/poly, contradicting the assumption. We must give a method for using (M^q, Θ^q) to decide SAT in polytime and with compact parameters Θ . Given ϕ , our method constructs a simple related formula ϕ' such that

- ϕ' has at least one satisfying assignment (so $Z(\phi') > 0$ and thus $p(1 \mid \phi')$ is defined)
- ϕ' has satisfying assignments with $A_1 = 1$ (i.e., $p(1 \mid \phi') > 0$) if and only if ϕ is satisfiable

Our construction also provides a polynomial function f such that $|\phi'|$ is guaranteed to be $\leq f(|\phi|)$. We now define Θ by $\theta_n = \theta_{f(n)}^q$ ($\forall n$). When our SAT algorithm with compact parameters Θ is given ϕ of length n , it can use the polynomial-size advice string θ_n to ask (M^q, Θ^q) in polynomial time for $p(1 \mid \phi')$. $\text{SAT}(\phi)$ returns true iff that probability is > 0 .¹⁰ \square

3.3 ELNCP models cannot even capture all EC (or ECCP) supports or rankings

We can strengthen Theorem 1 as follows:

Theorem 2. *Assuming $\text{NP} \not\subseteq \text{P/poly}$, there exists an efficiently computable normalizable weighted*

⁹Almost. This Z could be irrational, but at least it is computable to any desired precision. For any rational $\hat{Z} \approx Z$, we can say $\hat{p} = \tilde{p}/\hat{Z} \approx p$ is EC, via a Turing machine $M^{\hat{p}}$ that stores \hat{Z} . Further remarks on irrationality appear on page 14 in Appendix A.

¹⁰See also the remark on implications for seq2seq models following the proof in Appendix A.

language \tilde{p} where there is no ELNCP \tilde{q} such that $\text{support}(\tilde{p}) = \text{support}(\tilde{q})$.

Proof. Observe that for any two weighted languages \tilde{p} and \tilde{q} with the same support, $\forall \hat{\mathbf{x}} \in V^*$, $Z_{\tilde{p}}(\hat{\mathbf{x}}) > 0 \iff Z_{\tilde{q}}(\hat{\mathbf{x}}) > 0$ (where $Z_{\tilde{p}}$ and $Z_{\tilde{q}}$ return the prefix probabilities of \tilde{p} and \tilde{q} respectively). Thus, for any $\hat{\mathbf{x}}$ with $Z_{\tilde{p}}(\hat{\mathbf{x}}) > 0$, $p(1 \mid \hat{\mathbf{x}}) \triangleq Z_{\tilde{p}}(\hat{\mathbf{x}}1)/Z_{\tilde{p}}(\hat{\mathbf{x}})$ and $q(1 \mid \hat{\mathbf{x}}) \triangleq Z_{\tilde{q}}(\hat{\mathbf{x}}1)/Z_{\tilde{q}}(\hat{\mathbf{x}})$ are well-defined and $p(1 \mid \hat{\mathbf{x}}) > 0 \iff q(1 \mid \hat{\mathbf{x}}) > 0$. If \tilde{q} is ELNCP, then all such probabilities $q(1 \mid \hat{\mathbf{x}})$ can be computed in polytime with compact parameters, so it is likewise efficient to determine whether $p(1 \mid \hat{\mathbf{x}}) > 0$. But this cannot be the case when \tilde{p} is the weighted language used in the proof of Theorem 1, since that would suffice to establish that $\text{SAT} \in \text{P/poly}$, following the proof of that theorem. \square

To put this another way, there exists an unweighted language in P (namely $\text{support}(\tilde{p})$) that is not the support of any ELNCP distribution.

If they have different support, normalizable languages also differ in their ranking of strings:

Lemma 3. *Let \tilde{p}, \tilde{q} be normalizable weighted languages with $\text{support}(\tilde{p}) \neq \text{support}(\tilde{q})$. Then $\exists \mathbf{x}_1, \mathbf{x}_2 \in V^*$ such that $\tilde{p}(\mathbf{x}_1) < \tilde{p}(\mathbf{x}_2)$ but $\tilde{q}(\mathbf{x}_1) \geq \tilde{q}(\mathbf{x}_2)$.*

Therefore, no ELNCP \tilde{q} captures the string ranking of \tilde{p} from Theorem 2. And for some \tilde{p} , any ELNCP \tilde{q} misranks even string pairs of “similar” lengths:

Theorem 3. *Assuming $\text{NP} \not\subseteq \text{P/poly}$, there exists an efficiently computable normalizable weighted language \tilde{p} such that no ELNCP \tilde{q} with $\text{support}(\tilde{q}) \supseteq \text{support}(\tilde{p})$ has $\tilde{p}(\mathbf{x}_1) < \tilde{p}(\mathbf{x}_2) \implies \tilde{q}(\mathbf{x}_1) < \tilde{q}(\mathbf{x}_2)$ for all $\mathbf{x}_1, \mathbf{x}_2 \in V^*$. Indeed, any such \tilde{q} has a counterexample where $\tilde{p}(\mathbf{x}_1) = 0$. Moreover, there is a polynomial $f_{\tilde{q}} : \mathbb{N} \rightarrow \mathbb{N}$ such that a counterexample exists for every \mathbf{x}_1 such that $\tilde{p}(\mathbf{x}_1) = 0$ and $\tilde{q}(\mathbf{x}_1) > 0$, where the \mathbf{x}_2 in this counterexample always satisfies $|\mathbf{x}_2| \leq f_{\tilde{q}}(|\mathbf{x}_1|)$.*

Theorem 3 is relevant if one wishes to train a model \tilde{q} to rerank strings that are proposed by another method (e.g., beam search on \tilde{q} , or exact k -best decoding from a more tractable distribution). If the desired rankings are given by Theorem 3’s \tilde{p} , any smoothed¹¹ ELNCP model \tilde{q} will misrank some sets of candidate strings, even sets all of whose strings are “close” in length, by failing to rank an impossible string (\mathbf{x}_1 with $\tilde{p}(\mathbf{x}_1) = 0$) below a possible one (\mathbf{x}_2 with $\tilde{p}(\mathbf{x}_2) > 0$).

¹¹Smoothing is used to avoid ever incorrectly predicting 0 (a “false negative”) by ensuring $\text{support}(\tilde{q}) \supseteq \text{support}(\tilde{p})$. E.g., autoregressive language models often define $q(x \mid \hat{\mathbf{x}})$ using a softmax over $V \cup \{\$\}$, ensuring that $q(\mathbf{x}) > 0$ for all $\mathbf{x} \in V^*$.

3.4 ELNCP models cannot even approximate EC (or ECCP) distributions

Local probabilities cannot even be approximated well with compact parameters (if $\text{NP} \not\subseteq \text{P/poly}$).

Theorem 2 implies that there exists \tilde{p} whose local probabilities $p(x | \hat{\mathbf{x}})$ are not approximated by any ELNCP q to within any constant factor λ , since that would perfectly distinguish zeroes from non-zeroes and the resulting support sets would be equal.

Dropping the normalization requirement on the approximate local probabilities does not help. We say that M^q **efficiently approximately locally normalizes \tilde{p} with compact parameters Θ^q** if, for some $\lambda \geq 1$, it meets a relaxed version of the conditions from §3.1, where $q_n(\hat{\mathbf{x}}x)$ only needs to match $p(x | \hat{\mathbf{x}})$ within a factor of λ . (So possibly $\sum_{x \in V} q_n(\hat{\mathbf{x}}x) \neq 1$.) This is not possible for the language \tilde{p} from §3.2, since that would still allow a SAT algorithm to determine whether $p(1 | \phi') > 0$ in the proof of Theorem 1.

However, these demonstrations hinge on the difficulty of multiplicative approximation of zeroes — whereas real-world distributions may lack zeroes. Below we further show that it is hard even to approximate the *non-zero* local conditional probabilities (even with the additional help of randomness).

Theorem 4. *Assuming $\text{NP} \not\subseteq \text{P/poly}$, there exists an efficiently computable weighted language $\tilde{p} : V^* \rightarrow \mathbb{R}_{\geq 0}$ such that there is no (M^q, Θ^q) where $\Theta^q = \{\theta_n^q | n \in \mathbb{N}\}$ that satisfies all of the following properties (similar to §3.1):*

- the parameter size $|\theta_n^q|$ grows only as $O(\text{poly}(n))$
- $M^q(\theta_n^q)$ returns a probabilistic Turing machine q_n in time $O(\text{poly}(n))$
- there exists $\lambda \geq 1$ such that for each $x \in V \cup \{\$\}$ and $\hat{\mathbf{x}} \in V^*$ with $|\hat{\mathbf{x}}| \leq n$ and $p(x | \hat{\mathbf{x}}) > 0$, the probabilistic computation $q_n(\hat{\mathbf{x}}x)$ has probability $> 2/3$ of approximating $p(x | \hat{\mathbf{x}})$ to within a factor of λ (that is, $q_n(\hat{\mathbf{x}}x)/p(x | \hat{\mathbf{x}}) \in [1/\lambda, \lambda]$)
- q_n runs on those inputs $\hat{\mathbf{x}}x$ in time $O(\text{poly}(n))$

Moreover, the statement above still remains true with either of the following modifications:

- (a) the approximation guarantee is only required to hold for prefixes $\hat{\mathbf{x}}$ such that $\{\mathbf{x} : \hat{\mathbf{x}} \leq \mathbf{x}\}$ is finite (so that $p(x | \hat{\mathbf{x}})$ is computable by brute force)
- (b) $\text{support}(\tilde{p}) = V^*$

4 Alternative model families

We now discuss alternative families of sequence distributions that trade away efficiency or compactness in exchange for greater capacity, as shown in Table 1.

4.1 Energy-based models (EBMs)

Energy-based models (LeCun et al., 2006) of discrete sequences (Rosenfeld et al., 2001; Sandbank, 2008; Huang et al., 2018) traditionally refer to the EC models of §2.2. Only the unnormalized probabilities $\tilde{p}_\theta(\mathbf{x})$ are required to be efficiently computable. Lemmas 1 and 2 showed that this model family contains all ELN languages and can achieve any support in P. Theorem 1 shows that it also contains languages that are not ELN or even ELNCP: intuitively, the reason is that the sums $Z(\hat{\mathbf{x}})$ needed to compute the local normalizing constants (see §2.1) can be intractable.

If we generalize energy-based sequence models to include all ECCP models — that is, we allow non-uniform computation with compact parameters — then Lemmas 1 and 2 guarantee that they can capture all ELNCP languages and furthermore all languages in P/poly (though still not NP-complete languages).

Experiments on different parameterizations.

Maximum-likelihood parameter estimation (MLE) can be expensive in an EBM because the likelihood formula involves the expensive summation $Z = \sum_{\mathbf{x} \in V^*} \tilde{p}_\theta(\mathbf{x})$. This forces us in practice to use alternative estimators that do not require computing normalized probabilities, such as noise-contrastive estimation (NCE) or score matching (§1), which are less statistically efficient. In pilot experiments we found that both RNN- and Transformer-based EBMs trained with NCE achieved *worse* held-out perplexity than comparable locally normalized models trained with MLE. This might be due to a capacity limitation of the specific globally normalized architectures (i.e., no parameters work well), or excess capacity (i.e., too many parameters work well on the finite sample), or statistical inefficiency of the estimator (the NCE objective on the finite sample, with the noise distribution we chose, does not distinguish among parameters as well as MLE does), or an optimization difficulty caused by local optima in the NCE optimization landscape.

Fortunately, it is possible to infuse a globally normalized architecture with the inductive bias of a locally normalized one, which empirically yields good results. **Residual energy-based models (REBMs)** (Bakhtin et al., 2021) are a simple hybrid architecture:

$$p_\theta(\mathbf{x}) \propto \tilde{p}_\theta(\mathbf{x}) \triangleq p_0(\mathbf{x}) \cdot \exp g_\theta(\mathbf{x})$$

This simply multiplies our previous weight by a new factor $p_0(\mathbf{x})$. The *base model* $p_0 : L \rightarrow (0, 1]$ is a locally normalized neural sequence model (ELN model) that was pretrained on the same distribution. $g_\theta : V^* \rightarrow \mathbb{R}$ is a learnable function (with parameters

θ) that is used to adjust p_0 , yielding a weighted language \tilde{p}_θ with the same support L . We implemented REBMs, again with NCE training, and evaluated them on two different neural architectures (GRU- and Transformer-based) and 3 datasets (WikiText (Merity et al., 2017), Yelp (Yelp), and RealNews (Zellers et al., 2019)). In each setting we tried, the RBM slightly but significantly improved the perplexity of the base model p_0 ($p < 0.05$).¹²

4.2 Latent-variable models

Autoregressive models have $Z = 1$ for any setting of the parameters (or at least any setting that guarantees consistency: see footnote 7). Clearly $Z = 1$ ensures that Z is both finite and tractable. Can we find a model family that retains this convenience (unlike EBMs), while still being expressive enough to have any non-empty language in P as support?

Autoregressive *latent-variable* models form such a family. As in directed graphical models, the use of latent variables provides a natural way to model partial observations of an underlying stochastic sequence of events. We will model an observed sequence \mathbf{x} of length n as a function of a latent string \mathbf{z} of length $O(\text{poly}(n))$. As in EBMs, the probability $p(\mathbf{x})$ can be computationally intractable, allowing these models to break the expressivity bottleneck of ordinary autoregressive models. However, the intractability no longer comes from exponentially many summands in the denominator Z , but rather from exponentially many summands in the numerator — namely, the summation over all latent \mathbf{z} that could have produced \mathbf{x} . Notice that as a result, even unnormalized string weights are now hard to compute, although once computed they are already normalized.

Formally, we define **marginalized** weighted languages. We say that \tilde{p} is a **marginalization** of the weighted language \tilde{r} if it can be expressed as $\tilde{p}(\mathbf{x}) = \sum_{\mathbf{z}: \mu(\mathbf{z})=\mathbf{x}} \tilde{r}(\mathbf{z})$, where $\mu : \mathcal{Z} \rightarrow V^*$ is some function (the **marginalization operator**). We say it is a **light marginalization** if $|\mathbf{z}| \in O(\text{poly}(|\mu(\mathbf{z})|))$ and μ runs in time $O(\text{poly}(|\mathbf{z}|))$.¹³ Typically $\mu(\mathbf{z})$ extracts a subsequence of \mathbf{z} ; it can be regarded as keeping the observed symbols while throwing away a

¹²We independently conceived of and implemented the RBM idea proposed in Bakhtin et al. (2021). Details of neural architecture choice, model parameter sizes, training regimen, and evaluation (Appendices B–D) differ between our work and theirs, which also reported positive empirical results (on different datasets). We regard the two independent positive findings as a strong indication that the RBM design is effective.

¹³WLOG, μ can be required to run in linear time $O(|\mathbf{z}|)$, as it does in our constructions below.

polynomially bounded number of latent symbols.

Light marginalizations of ELN distributions are a reasonable formalization of latent-variable autoregressive models. They are more powerful than ELN distributions, and even include some distributions that (by Lemma 1) are not even ELNCP or ECCP:

Theorem 5. *There exists a light marginalization p of an ELN distribution, such that $\text{support}(p)$ is an NP-complete language.*

Our proof of Theorem 5 relies on special structure of a certain NP-complete language (SAT) and does not evidently generalize to all languages in NP.

However, light marginalizations of ELNCP distributions are more powerful still,¹⁴ and can have any language $\in \mathsf{NP}$ or even NP/poly (§2.4) as support:

Theorem 6. *The following statements are equivalent for any nonempty $L \subseteq V^*$:*

- (a) $L \in \mathsf{NP/poly}$.
- (b) L is the support of a light marginalization of an ELNCP distribution.
- (c) L is the support of a light marginalization of an ECCP weighted language.

Theorems 5 and 6 make use of unrestricted latent-variable autoregressive models. There exist more practical restricted families of such models that admit tractable computation of $p(\mathbf{x})$ (Lafferty et al., 2001; Rastogi et al., 2016; Wu et al., 2018; Buys and Blunsom, 2018). Such models are EC (and indeed, typically ELN) — but this limits their expressivity, by Theorem 1. Both Lin et al. (2019) and Buys and Blunsom (2018) observed that such models yield worse empirical results than models that do not have tractable exact inference methods. The tractability requirement is dropped in “self-talk” (blixt, 2020; Gontier et al., 2020; Schwartz et al., 2020), where a neural autoregressive language model generates an analysis of the prefix $\hat{\mathbf{x}}$ via latent intermediate symbols before predicting the next output symbol.¹⁵

We remark that for autoregressive models, the *position* of the latent variables is significant. Marginalizing out latent variables at the *end* of the string adds no

¹⁴The capacity established by Theorem 6 does not need the full power of marginalization. We could similarly define light *maximizations* of ELNCP distributions, $\tilde{p}(\mathbf{x}) = \max_{\mathbf{z}: \mu(\mathbf{z})=\mathbf{x}} \tilde{r}(\mathbf{z})$. Replacing sum by max does not change the support.

¹⁵Here the marginal distribution of the next observed symbol can require superpolynomial time to compute (if $\#\mathsf{P} \neq \mathsf{FP}$, which follows from $\mathsf{NP} \not\subseteq \mathsf{P/poly}$). Theorem 1 could likewise be evaded by *other* autoregressive approaches that invest superpolynomial computation in predicting the next symbol (Graves, 2016). Each autoregressive step might explicitly invoke lookahead or reasoning algorithms, just as feed-forward network layers can invoke optimizers or solvers (Amos and Kolter, 2017; Wang et al., 2019b).

power. More precisely, if an ELNCP distribution is over strings \mathbf{z} of the form $\mathbf{x}\#\mathbf{y}$, then its marginalization via $\mu(\mathbf{x}\#\mathbf{y}) = \mathbf{x}$ can be expressed more simply as an ELNCP language. Thus, by Theorem 2, marginalizations of such distributions cannot have arbitrary NP languages as support. Our proofs of Theorems 5 and 6 instead use latent strings of the form $\mathbf{y}\#\mathbf{x}$, where all latent variables precede all observed ones (as in Kingma and Welling, 2014). (This simple design can always be used without loss of generality.) Trying to reorder those latent strings as $\mathbf{x}\#\mathbf{y}$ while preserving their weights would have yielded a non-ELNCP distribution $p(\mathbf{x}\#\mathbf{y})$ (because if it were ELNCP, then $p(\mathbf{x})$ would be ELNCP also, and we know from Lemma 1 that it cannot be for any distribution whose support is an NP-complete language).

How about lightly marginalizing ECCP languages instead of ELNCP ones? This cannot model any additional *unweighted* languages, by Theorem 6. But it may be able to model more probability distributions. One can easily construct a light marginalization p of an ECCP distribution such that $\#(\phi) = c_n \cdot p(\phi)$, where $\#(\phi)$ is the number of satisfying assignments of ϕ and the constant c_n depends only on $n = |\phi|$. We conjecture that this is not possible with lightly marginalized ELNCP distributions.

4.3 Lookup models

§2.3 noted that with exponential growth in stored parameters, it is possible to fit any weighted language up to length n , with local probabilities computed in only $O(n)$ time by lookup. Of course this rapidly becomes impractical as n increases, even if the amount of training data increases accordingly. However, there has been some recent movement toward storage-heavy models. Such models are typically semiparametric: they use a parametric neural model, such as an autoregressive model, together with an external knowledge base of text strings or factoids that are not memorized in the layer weights. The neural model generates queries against the knowledge base and combines their results. Examples include k NNLMs (Khandelwal et al., 2020) and semiparametric LMs (Yogatama et al., 2021). The knowledge base grows linearly with the training data rather than compressing the data into a smaller parameter vector. It is in fact a copy of the training data, indexed to allow fast lookup (Indyk and Motwani, 1998). (Preparing the index is much cheaper than neural network training.) Access to the large knowledge base may reduce the amount of computation needed to find the local conditional probabilities, much as in the trie construction of §2.3.

5 Related work

Chen et al. (2018) show that it is hard to map *RNN parameters* to properties of the resulting autoregressive weighted language. Our point is that any *given* autoregressive weighted language, if consistent, comes with an efficient algorithm to map a *string* to properties of that string (its prefix probability and next-symbol probabilities). Thus, given consistent RNN parameters, the resulting language is always one of those for which these problems are tractable — this tractability property of the language is *not* hard to determine.

In a Bayes network — which is really just an autoregressive model of fixed-length strings — approximate marginal inference is NP-hard (Roth, 1996). Assuming $\text{NP} \not\subseteq \text{P/poly}$ and the grid-minor hypothesis, Chandrasekaran et al. (2008, Theorem 5.6) further showed that for any infinite sequence of graphs G_1, G_2, \dots where G_n has treewidth n , there is no sequence of algorithms M_1, M_2, \dots such that M_n performs approximate marginal inference in time $O(\text{poly}(n))$ on graphical models of structure G_n . This remarkable negative result says that in *any* graph sequence of unbounded treewidth, approximating the normalizing constant for G_n given arbitrary parameters is hard (not $O(\text{poly}(n))$), even with advice strings. Our negative result (Theorem 4) focuses on *one particular* infinite weighted language, showing that approximating local conditional probabilities given an arbitrary length- n prefix is hard in the same way. (So this language cannot be captured by an RNN, even with advice strings.)

6 Conclusion and future work

While autoregressive models have several properties that appeal to practitioners, we have shown under common complexity-theoretic assumptions that they cannot match the expressivity of energy-based models — unless their runtime or parameter size grows superpolynomially in input length, or they are allowed to use latent variables, which add considerable power.

All model families we have discussed in this paper can be seen as making compromises between different desiderata (Table 1). Natural follow-up questions include ‘*Are there model families that win on all fronts?*’ ‘*What are other modeling desiderata?*’

While *some* languages $\in \text{P}$ cannot be supports of ELNCPs, we do not know if the same can be said for *most* languages $\in \text{P}$. This problem seems to be closely related to the average complexity of NP-complete languages, where most questions remain open (Levin, 1986; Bogdanov and Trevisan, 2006).

Acknowledgements

We thank the anonymous reviewers for their comments. We also thank our colleagues at Johns Hopkins University, Facebook, and Carnegie Mellon University for their comments on earlier versions of the manuscript. This material is based upon work at Johns Hopkins University supported by the National Science Foundation under Grant No. 1718846. It does not represent the views of Microsoft (where Dr. Eisner is also a paid employee, in an arrangement that has been reviewed and approved by the Johns Hopkins University in accordance with its conflict of interest policies).

References

- Brandon Amos and J. Zico Kolter. 2017. [OptNet: Differentiable optimization as a layer in neural networks](#). In *ICML*.
- Sanjeev Arora and Boaz Barak. 2009. *Computational Complexity: a Modern Approach*. Cambridge University Press.
- Anton Bakhtin, Yuntian Deng, Sam Gross, Myle Ott, Marc’Aurelio Ranzato, and Arthur Szlam. 2021. [Residual energy-based models for text generation](#). *JMLR*, 22(40):1–41.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. [On the dangers of stochastic parrots: Can language models be too big?](#) . In *FAccT*.
- Daniel Le Berre and Olivier Roussel. 2009. *SAT Competition 2009: Benchmark Submission Guidelines*. Accessed May 17, 2020.
- blixt. 2020. [Re: Teaching gpt-3 to identify nonsense](#). <https://news.ycombinator.com/item?id=23990902>. Online (accessed Oct 23, 2020).
- Andrej Bogdanov and Luca Trevisan. 2006. [Average-case complexity](#). *Foundations and Trends in Theoretical Computer Science*, 2(1):1–106.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#).
- Jan Buys and Phil Blunsom. 2018. [Neural syntactic generative models with exact marginalization](#). In *NAACL*.
- Shu Cai and Kevin Knight. 2013. [Smatch: an evaluation metric for semantic feature structures](#). In *ACL*.
- Venkat Chandrasekaran, Nathan Srebro, and Prahladh Harsha. 2008. [Complexity of inference in graphical models](#). In *UAI*.
- Yining Chen, Sorcha Gilroy, Andreas Maletti, Jonathan May, and Kevin Knight. 2018. [Recurrent neural networks as weighted language recognizers](#). In *NAACL*.
- Zhiyi Chi and Stuart Geman. 1998. [Estimation of probabilistic context-free grammars](#). *Computational Linguistics*, 24(2):299–305.
- Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. 2019. [Generating long sequences with sparse transformers](#). *ArXiv*, abs/1904.10509.
- Stephen A. Cook. 1971. [The complexity of theorem-proving procedures](#). In *STOC*.
- John DeNero and D. Klein. 2008. [The complexity of phrase alignment problems](#). In *ACL*.
- Nicolas Gontier, Koustuv Sinha, Siva Reddy, and C. Pal. 2020. [Measuring systematic generalization in neural proof generation with transformers](#). In *NeurIPS*.
- A. Graves. 2016. [Adaptive computation time for recurrent neural networks](#). *ArXiv*, abs/1603.08983.
- Michael U. Gutmann and Aapo Hyvärinen. 2010. [Noise-contrastive estimation: A new estimation principle for unnormalized statistical models](#). In *AISTATS*.
- Michael U. Gutmann and Aapo Hyvärinen. 2012. [Noise-contrastive estimation of unnormalized statistical models, with applications to natural image statistics](#). *JMLR*, 13(11):307–361.
- Y. Huang, A. Sethy, K. Audhkhasi, and B. Ramabhadran. 2018. [Whole sentence neural language models](#). In *ICASSP*.
- Aapo Hyvärinen. 2005. [Estimation of non-normalized statistical models by score matching](#). *JMLR*, 6(24):695–709.
- W. Idsardi. 2006. [A simple proof that optimality theory is computationally intractable](#). *Linguistic Inquiry*, 37:271–275.
- Piotr Indyk and Rajeev Motwani. 1998. [Approximate nearest neighbors: Towards removing the curse of dimensionality](#). In *STOC*.
- Richard M. Karp and Richard J. Lipton. 1980. [Some connections between nonuniform and uniform complexity classes](#). In *STOC*.

- Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. 2020. [Generalization through memorization: Nearest neighbor language models](#). In *ICLR*.
- Diederik P. Kingma and Max Welling. 2014. [Auto-encoding variational Bayes](#). In *ICLR*.
- Richard E. Ladner. 1975. [The circuit value problem is log space complete for P](#). *SIGACT News*, 7(1):18–20.
- John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. [Conditional random fields: Probabilistic models for segmenting and labeling sequence data](#). In *ICML*.
- Yann LeCun, Sumit Chopra, Raia Hadsell, Marc'Aurelio Ranzato, and Fu-Jie Huang. 2006. [A tutorial on energy-based learning](#). In *Predicting Structured Data*. MIT Press.
- Leonid A. Levin. 1986. [Average case complete problems](#). *SIAM Journal on Computing*, 15:285–286.
- Chu-Cheng Lin, Hao Zhu, Matthew R. Gormley, and Jason Eisner. 2019. [Neural finite-state transducers: Beyond rational relations](#). In *NAACL*.
- Zhuang Ma and Michael Collins. 2018. [Noise contrastive estimation and negative sampling for conditional models: Consistency and statistical efficiency](#). In *EMNLP*.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2017. [Pointer sentinel mixture models](#). *ArXiv*, abs/1609.07843.
- Tomas Mikolov, Stefan Kombrink, Anoop Deoras, Lukas Burget, and Jan Honza Cernocky. 2011. [RNNLM—Recurrent neural network language modeling toolkit](#). In *IEEE Automatic Speech Recognition and Understanding Workshop*.
- Randall Munroe. 2009. [My Hobby: Embedding NP-Complete Problems in Restaurant Orders](#). Online (accessed May 29, 2020).
- Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. 2016. [WaveNet: A generative model for raw audio](#). *ArXiv*, abs/1609.03499.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#).
- Pushpendre Rastogi, Ryan Cotterell, and Jason Eisner. 2016. [Weighting finite-state transductions with neural context](#). In *NAACL*.
- Ronald Rosenfeld, Stanley Chen, and Xiaojin Zhu. 2001. [Whole-sentence exponential language models: A vehicle for linguistic-statistical integration](#). *Computer Speech & Language*, 15(1):55–73.
- Dan Roth. 1996. [On the hardness of approximate reasoning](#). *Artificial Intelligence*, 82(1–2):273–302.
- Ben Sandbank. 2008. [Refining generative language models using discriminative learning](#). In *EMNLP*.
- Vered Shwartz, Peter West, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2020. [Unsupervised commonsense question answering with self-talk](#). In *EMNLP*.
- Hava T. Siegelmann and Eduardo D. Sontag. 1992. [On the computational power of neural nets](#). In *COLT*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. [Sequence to sequence learning with neural networks](#). In *NeurIPS*.
- Benigno Uribe, Marc-Alexandre Côté, Karol Gregor, Iain Murray, and Hugo Larochelle. 2016. [Neural autoregressive distribution estimation](#). *JMLR*, 17(1):7184–7220.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *NeurIPS*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019a. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *ICLR*.
- Po-Wei Wang, P. Donti, B. Wilder, and J. Z. Kolter. 2019b. [SATNet: Bridging deep learning and logical reasoning using a differentiable satisfiability solver](#). In *ICML*.
- Sean Welleck, Ilia Kulikov, Jaedeok Kim, Richard Yuanzhe Pang, and Kyunghyun Cho. 2020. [Consistency of a recurrent language model with respect to incomplete decoding](#). In *EMNLP*.
- Shijie Wu, Pamela Shapiro, and Ryan Cotterell. 2018. [Hard non-monotonic attention for character-level transduction](#). In *EMNLP*.
- Yelp. [Yelp open dataset](#). <https://www.yelp.com/dataset>.
- Dani Yogatama, Cyprien de Masson d’Autume, and Lingpeng Kong. 2021. [Adaptive semiparametric language models](#). *ArXiv*, abs/2102.02557.
- Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. 2019. [Defending against neural fake news](#). In *NeurIPS*.

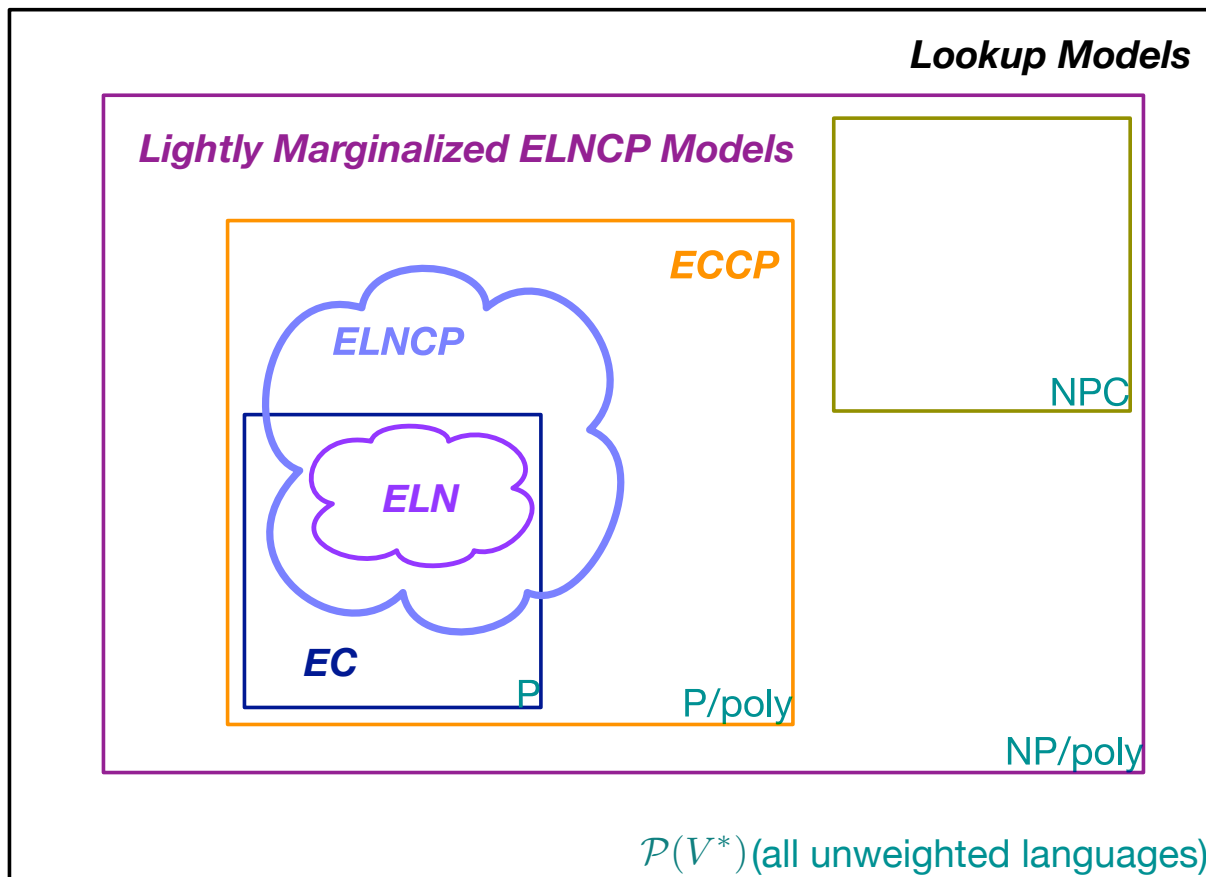


Figure 2: The space of unweighted languages. We assume in this diagram that $NP \not\subseteq P/poly$. Each rectangular outline corresponds to a complexity class (named in its lower right corner) and encloses the languages whose decision problems fall into that class. Each ***bold-italic label*** (colored to match its shape outline) corresponds to a model family and encloses the languages that can be expressed as the *support* of some *weighted* language in that family. All induced partitions in the figure are non-empty sets: shape A properly encloses shape B if and only if language class A is a strict superset of language class B. As mentioned in Table 1, standard autoregressive models (ELN models) have support languages that fall in P (Theorem 2). ELNCP models (§3.1) extend ELN models by allowing the parameter size to grow polynomially in string length, allowing them to capture both more languages inside P (Lemma A.4) and languages outside P (including undecidable but sparse languages) that can be characterized autoregressively with the help of these compact parameters. All of those languages belong in the class P/poly. Theorem 2 establishes that energy-based (EC) and ECCP models go strictly further than ELN and ELNCP models, respectively (Theorem 2): they correspond to the *entire* classes P and P/poly (Lemma 1). However, even ECCP does not capture any NP-complete languages under our assumption $NP \not\subseteq P/poly$. Allowing a polynomial number of latent symbols extends the power further still: lightly marginalized ELNCP or ECCP distributions cover exactly the languages $\in NP/poly$ (Theorem 6). Finally, if we were to drop the requirement that the parameters Θ must be compact, we could store lookup tries to model any weighted language (§4.3).

A Proofs

Lemma 1. *For any $L \in P$, there exists an EC weighted language with support L . For any $L \in P/\text{poly}$, there exists an ECCP language with support L . But for any $L \in \text{NP-complete}$, there exists no ECCP language with support L (assuming $\text{NP} \not\subseteq P/\text{poly}$).*

This simple lemma relates our classes EC and ECCP of *weighted* languages to the complexity classes P and P/poly of their supports, which are *unweighted* formal languages (§2). It holds because computing a string’s weight can be made as easy as determining whether that weight is nonzero (if we set the weights in a simple way), but is certainly no easier. We spell out the trivial proof to help the reader gain familiarity with the formalism.

Proof. Given L , define a weighted language \tilde{p} with support L by $\tilde{p}(\mathbf{x}) = 1$ if $\mathbf{x} \in L$ and $\tilde{p}(\mathbf{x}) = 0$ otherwise.

If $L \in P$, then clearly \tilde{p} is EC since the return value of 1 or 0 can be determined in polytime.

If $L \in P/\text{poly}$, L can be described as a tuple (M, Θ) following our characterization in §2.4. It is easy to show that \tilde{p} is ECCP, using the same polynomially-sized advice strings Θ . We simply construct $M^{\tilde{p}}$ such that $M^{\tilde{p}}(\theta_n)$ returns 1 or 0 on input \mathbf{x} according to whether $M(\theta_n)$ accepts or rejects \mathbf{x} . Both $M^{\tilde{p}}(\theta_n)$ and $M^{\tilde{p}}(\theta_n)(\mathbf{x})$ are computed in time $O(\text{poly}(n))$ if $|\mathbf{x}| = n$. (The technical construction is that $M^{\tilde{p}}$ simulates the operation of M on the input θ_n to obtain the description of the Turing machine $M_n = M(\theta_n)$, and then outputs a slightly modified version of this description that will write 1 or 0 on an output tape.)

For the second half of the lemma, we use the reverse construction. Suppose \tilde{p} is an ECCP weighted language with support L . \tilde{p} can be characterized by a tuple $(M^{\tilde{p}}, \Theta)$. It is easy to show that $L \in P/\text{poly}$, using the same polynomially-sized advice strings Θ . We simply construct M such that $M(\theta_n)$ accepts \mathbf{x} iff $M^{\tilde{p}}(\theta_n)(\mathbf{x}) > 0$. Then by the assumption, $L \notin \text{NP-complete}$. \square

Lemma 2. *An ELNCP model \tilde{p} is also ECCP. Likewise, an ELN model is also EC.*

Proof. Let \tilde{p} be an ELNCP language. Let M^q efficiently locally normalize \tilde{p} with compact parameters $\Theta^q = \{\theta_n^q \mid n \in \mathbb{N}\}$. It is simple to define a Turing machine M^r that maps each parameter string θ_n^q to a Turing machine r_n , where $r_n(\mathbf{x})$ simply computes $\left(\prod_{t=1}^n q_n(x_t \mid \mathbf{x}_{<t})\right) \cdot q_n(\$ \mid \mathbf{x})$. Then for all \mathbf{x} of

length n , $r_n(\mathbf{x}) = \left(\prod_{t=1}^n p(x_t \mid \mathbf{x}_{<t})\right) \cdot p(\$ \mid \mathbf{x})$, by the definition of local normalization, and thus $r_n(\mathbf{x}) = p(\mathbf{x})$.

M^r can be constructed by incorporating the definition of M^q , so that $r_n = M^r(\theta_n^q)$ can include $q_n = M^q(\theta_n^q)$ as a subroutine. This allows r_n to query q_n for local conditional probabilities and multiply them together.

- Since M^q runs in polytime, it is straightforward for this construction to ensure that M^r runs in polytime as well.
- Since $q_n(\cdot \mid \hat{\mathbf{x}}) \in O(\text{poly}(n))$, this construction can ensure that r_n runs in polytime as well.
- We were given that $|\theta_n^q| \in O(\text{poly}(n))$ (compact parameters).

Since p is the weighted language defined by (M^r, Θ^q) , and M^r and Θ^q have the properties just discussed, we see that p is efficiently computable with compact parameters.

In the case where \tilde{p} is more strongly known to be ELN (the parameters Θ^q are not needed), a simplification of this argument shows that it is EC. \square

Theorem 1. *Assuming $\text{NP} \not\subseteq P/\text{poly}$, there exists an efficiently computable normalizable weighted language \tilde{p} that is not ELNCP.*

Proof. The proof was sketched in §3.2. Here we fill in the details.

The unweighted language \tilde{p} defined in that section is efficiently computable via the following simple algorithm that outputs $\tilde{p}(\mathbf{x})$ given $\mathbf{x} \in \mathbb{B}^*$. If \mathbf{x} has a prefix that encodes a formula ϕ , and the remainder of \mathbf{x} is a satisfying assignment \mathbf{a} to the variables of ϕ , then return $(\frac{1}{3})^{|\mathbf{x}|+1}$. Otherwise return 0. This algorithm can be made to run in polynomial time because whether an assignment satisfies a formula can be determined in polynomial time (a fact that is standardly used to establish that $\text{SAT} \in \text{NP}$).

Given a formula ϕ with variables A_1, \dots, A_j , we define $\phi' = (\neg A_1 \wedge \neg A_2 \wedge \dots \wedge \neg A_j \wedge \neg A_{j+1}) \vee (A_1 \wedge \text{Shift}(\phi))$, where $\text{Shift}(\phi)$ is a version of ϕ in which A_i has been renamed to A_{i+1} for all $1 \leq i \leq j$. It is obvious that ϕ' and p have the properties stated in the proof sketch. The strings in L that begin with ϕ' are precisely the strings of the form $\phi' \mathbf{a}'$ where \mathbf{a}' is a satisfying assignment of ϕ' — which happen just when $\mathbf{a}' = \mathbf{0}^{j+1}$ or $\mathbf{a}' = \mathbf{1} \mathbf{a}$ where \mathbf{a} is a satisfying assignment of ϕ . At least one string in L begins with ϕ' , namely $\phi' \mathbf{0}^{j+1}$, so $Z(\phi') > 0$. Moreover, $Z(\phi' \mathbf{1}) > 0$ iff ϕ has any satisfying assignments. Therefore the

local probability $p(1 \mid \phi') = Z(\phi'1) / Z(\phi')$ is defined (see §2.1), and is > 0 iff $\text{SAT}(\phi)$.

Notice that the formal problem used in the proof is a version of SAT whose inputs are encoded using the same prefix-free encoding function enc that was used by our definition of L in §3.2. We must choose this encoding function to be concise in the sense that $\phi \triangleq \text{enc}(\phi)$ can be converted to and from the conventional encoding of ϕ in polynomial time. This ensures that our version of SAT is \leq_m^P -interreducible with the conventional version and hence NP-complete. It also ensures that there is a polynomial function f such that $|\phi'| \leq f(|\phi|)$, as required by the proof sketch, since there is a polynomial-time function that maps $\phi \rightarrow \phi \rightarrow \phi' \rightarrow \phi'$ and the output length of this function is bounded by its runtime. This is needed to show that our version of SAT is in P/poly.

Specifically, to show that the existence of (M^q, Θ^q) implies $\text{SAT} \in \text{P/poly}$, we use it to construct an appropriate pair (M, Θ) such that $(M(\theta_n))(\phi) = \text{SAT}(\phi)$ if $|\phi| = n$. As mentioned in the proof sketch, we define Θ by $\theta_n = \theta_{f(n)}^q$, and observe that $|\theta_n| \in O(\text{poly}(n))$ (thanks to compactness of the parameters Θ^q and the fact that f is polynomially bounded). Finally, define $M(\theta_n)$ to be a Turing machine that maps its input ϕ of length n to ϕ' of length $\leq f(n)$, then calls $M^q(\theta_n) = M^q(\theta_{f(n)}^q)$ on $\phi'1$ to obtain $p(1 \mid \phi')$, and returns true or false according to whether $p(1 \mid \phi') > 0$. Computing ϕ' takes time polynomial in n (thanks to the properties of enc). Constructing $M^q(\theta_{f(n)})$ and calling it on ϕ' each take time polynomial in n (thanks to the properties of f and M^q). \square

Remark on conditional models. While we focus on modeling *joint* sequence probabilities in this work, we note that in many applications it often suffices to just model *conditional* probabilities (Sutskever et al., 2014). Unfortunately, our proof of Theorem 1 above implies that ELNCPs do not make good conditional models either: specifically, there exists ϕ such that deciding whether $p(1 \mid \phi) > 0$ is NP-hard, and thus beyond ELNCP's capability.

Remark on irrationality. In our definitions of ECCP and ELNCP languages, we implicitly assumed that the Turing machines that return weights or probabilities would write them in full on the output tape, presumably as the ratio of two integers. Such a Turing machine can only return rational numbers.

But then our formulation of Theorem 1 allows another proof. We could construct \tilde{p} such that the local

conditional probabilities $p(x \mid \hat{\mathbf{x}}) \triangleq Z(\hat{\mathbf{x}}x) / Z(\hat{\mathbf{x}})$ are sometimes irrational. In this case, they cannot be output exactly by a Turing machine, implying that \tilde{p} is not ELNCP. However, this proof exposes only a trivial weakness of ELNCPs, namely the fact that they can only define distributions whose local marginal probabilities are rational.

We can correct this weakness by formulating ELNCP languages slightly differently. A real number is said to be **computable** if it can be output by a Turing machine to any desired precision. That Turing machine takes an extra input b which specifies the number of bits of precision of the output. Similarly, our definitions of ECCP and ELNCP can be modified so that their respective Turing machines \tilde{p}_n and q_n take this form, are allowed to run in time $O(\text{poly}(n + b))$, and have access to the respective parameter vectors Θ_{n+b}^p and Θ_{n+b}^q . Since some of our results concern the ability to distinguish zero from small values (arbitrarily small in the case of Lemma A.4), our modified definitions also require \tilde{p}_n and q_n to output a bit indicating whether the output is *exactly* zero. For simplicity, we suppressed these technical details from our exposition.

Relatedly, in §4.3, we claimed that lookup models can fit any weighted language up to length n . This is not strictly true if the weights can be irrational. A more precise statement is that for any weighted language \tilde{p} , there is a lookup model that maps (\mathbf{x}, b) to the first b bits of $\tilde{p}(\mathbf{x})$. Indeed, this holds even when $\tilde{p}(\mathbf{x})$ is uncomputable.

Remark on computability. In §2.1 we claimed that any weighted language \tilde{p} that has a finite and strictly positive Z can be normalized as $p(\mathbf{x}) = \tilde{p}(\mathbf{x}) / Z$. However, Z may be uncomputable: that is, there is no algorithm that takes number of bits of precision b as input, and outputs an approximation of Z within b bits of precision. Therefore, even if \tilde{p} is computable, p may have weights that are not merely irrational but even *uncomputable*. An example appears in the proof of Lemma A.4 below. Weighted language classes (e.g. ELNCP) that only model normalized languages will not be able to model such languages, simply because the partition function is uncomputable.

However, our proof of Theorem 1 does not rely on this issue, because the \tilde{p} that it exhibits happens to have a computable Z . For any b , Z may be computed to b bits of precision as the explicit sum $\sum_{\mathbf{x}:|\mathbf{x}|\leq N} \tilde{p}(\mathbf{x})$ for a certain large N that depends on b .

Remark on RNNs. Our proof of Theorem 1 showed that our problematic language \tilde{p} is efficiently computable (though not by any locally normalized architecture with compact parameters). Because this paper is in part a response to popular neural architectures, we now show that \tilde{p} can in fact be computed efficiently by a recurrent neural network (RNN) with compact parameters. Thus, this is an example where a simple globally normalized RNN parameterization is fundamentally more efficient (in runtime or parameters) than any locally normalized parameterization of any architecture (RNN, Transformer, etc.).

Since we showed that \tilde{p} is efficiently computable, the existence of an RNN implementation is established in some sense by the ability of finite rational-weighted RNNs to simulate Turing machines (Siegelmann and Sontag, 1992), as well as an extension to Chen et al. (2018, Thm. 11) to a family of RNNs, where each RNN instance also takes some formula encoding as input. However, it is straightforward to give a concrete construction, for each $n \in \mathbb{N}$, for a simple RNN that maps each string $\mathbf{x} \in \mathbb{B}^n$ to $\tilde{p}(\mathbf{x})$. Here $\tilde{p}(\mathbf{x})$ will be either $(\frac{1}{3})^{n+1}$ or 0, according to whether \mathbf{x} has the form $\phi\mathbf{a}$ where ϕ encodes a 3-CNF-SAT formula ϕ that is satisfied by \mathbf{a} .¹⁶ The basic idea is that ϕ has $j \leq n$ variables, so there are only $O(n^3)$ possible 3-CNF clauses. The RNN allocates one hidden unit to each of these. When reading $\phi\mathbf{a}$, each clause encountered in ϕ causes the corresponding hidden unit to turn on, and then each literal encountered in \mathbf{a} turns off the hidden units for all clauses that would be satisfied by that literal. If any hidden units remain on after \mathbf{x} has been fully read, then ϕ was not satisfied by \mathbf{a} , and the RNN’s final output unit should return 0. Otherwise it should return $(\frac{1}{3})^{n+1}$, which is constant for this RNN. To obtain digital behaviors such as turning hidden units on and off, it is most convenient to use ramp activation functions for the hidden units and the final output unit, rather than sigmoid activation functions. Note that our use of a separate RNN M_n^{RNN} for each input length n is an example of using more hidden units for larger problems, a key idea that we introduced in §2.3 in order to look at asymptotic behavior. The RNN’s parameter sequence $\Theta^{\text{RNN}} = \{\theta_n^{\text{RNN}} \mid n \in \mathbb{N}\}$ is obviously

¹⁶The restriction to 3-CNF-SAT formulas is convenient, but makes this a slightly different definition of L and \tilde{p} than we used in the proofs above. Those proofs can be adjusted to show that this \tilde{p} , too, cannot be efficiently locally normalized with compact parameters. The only change is that in the construction of Theorem 1, ϕ' must be converted to 3-CNF. The proof then obtains its contradiction by showing that 3-CNF-SAT \in P/poly (which suffices since 3-CNF-SAT is also NP-complete).

compact, as θ_n^{RNN} only has to store the input length n . With our alphabet \mathbb{B} for \tilde{p} , $|\theta_n^{\text{RNN}}| \in O(\log n)$.

Lemma 3. *Let \tilde{p}, \tilde{q} be normalizable weighted languages with $\text{support}(\tilde{p}) \neq \text{support}(\tilde{q})$. Then $\exists \mathbf{x}_1, \mathbf{x}_2 \in V^*$ such that $\tilde{p}(\mathbf{x}_1) < \tilde{p}(\mathbf{x}_2)$ but $\tilde{q}(\mathbf{x}_1) \geq \tilde{q}(\mathbf{x}_2)$.*

Proof. Suppose that the claim is false, i.e., \tilde{p} and \tilde{q} have the same ranking of strings. Then the minimum-weight strings under \tilde{p} must also be minimum-weight under \tilde{q} . WLOG, there exists $\mathbf{x} \in V^*$ with $\tilde{p}(\mathbf{x}) = 0$ and $\tilde{q}(\mathbf{x}) = c > 0$. Then $c > 0$ is the minimum weight of strings in \tilde{q} . But this is not possible for a normalizable language \tilde{q} , since it means that $Z_{\tilde{q}} \triangleq \sum_{\mathbf{x}' \in V^*} q(\mathbf{x}') \geq \sum_{\mathbf{x}' \in V^*} c$ diverges. \square

Theorem 3. *Assuming $\text{NP} \not\subseteq \text{P/poly}$, there exists an efficiently computable normalizable weighted language \tilde{p} such that no ELNCP \tilde{q} with $\text{support}(\tilde{q}) \supseteq \text{support}(\tilde{p})$ has $\tilde{p}(\mathbf{x}_1) < \tilde{p}(\mathbf{x}_2) \Rightarrow \tilde{q}(\mathbf{x}_1) < \tilde{q}(\mathbf{x}_2)$ for all $\mathbf{x}_1, \mathbf{x}_2 \in V^*$. Indeed, any such \tilde{q} has a counterexample where $\tilde{p}(\mathbf{x}_1) = 0$. Moreover, there is a polynomial $f_{\tilde{q}} : \mathbb{N} \rightarrow \mathbb{N}$ such that a counterexample exists for every \mathbf{x}_1 such that $\tilde{p}(\mathbf{x}_1) = 0$ and $\tilde{q}(\mathbf{x}_1) > 0$, where the \mathbf{x}_2 in this counterexample always satisfies $|\mathbf{x}_2| \leq f_{\tilde{q}}(|\mathbf{x}_1|)$.*

Our proof of Theorem 3 below is based on an observation that any rational number a takes $\theta(\log a)$ bits to store, which implies that any polytime algorithm that assigns probability $q(\mathbf{x})$ to string \mathbf{x} must have $|\log q(\mathbf{x})|$ bounded below some polynomial function of $|\mathbf{x}|$. In other words, the negative quantity $\log q(\mathbf{x})$ is bounded above some polynomial function. Then we show there exists a family of \mathbf{x}_2 ’s, such that for every \mathbf{x}_1 there exists \mathbf{x}_2 , where $q(\mathbf{x}_1) > 0$, $q(\mathbf{x}_1) > q(\mathbf{x}_2)$ because \mathbf{x}_2 is longer, but $0 = \tilde{p}(\mathbf{x}_1) < \tilde{p}(\mathbf{x}_2)$ following Theorem 2 and Lemma 3.

Proof. Let \tilde{p} be the efficiently computable language defined in §3.2. And let q be the normalized language of \tilde{q} . We know there exists $\mathbf{x}_1 \in V^*$ such that $\tilde{p}(\mathbf{x}_1) = 0$ but $q(\mathbf{x}_1) > 0$, following the proof of Theorem 2. Since it takes polynomial time to output $\log q(\mathbf{x}_1)$ — \mathbf{x}_1 ’s weight in log scale, $|\log q(\mathbf{x}_1)|$ is bounded below some polynomial function of $|\mathbf{x}_1|$. Namely,

$$\exists a > 0, \exists s > 0 \text{ such that } |\log q(\mathbf{x}_1)| < a|\mathbf{x}_1|^s.$$

And since $\forall \mathbf{x}_1 \in L$, $\log q(\mathbf{x}_1) < 0$,

$$\exists a > 0, \exists s > 0 \text{ such that } \log q(\mathbf{x}_1) > -a|\mathbf{x}_1|^s.$$

We will refer to the condition above as ‘the \mathbf{x}_1 condition’ in the rest of the proof.

We now focus on the other string $\mathbf{x}_2 \in \text{support}(\tilde{p})$. Since \tilde{p} is an infinite weighted language, $\text{support}(\tilde{p})$ contains infinitely many strings. We claim that there is a family of (infinite) strings $L_2 = \{\mathbf{x}_2^{(n)} \mid n \in \mathbb{N}\}$ that is both a subset of $\text{support}(\tilde{p})$ and a subset of $\text{support}(\tilde{q})$, where $\forall n \in \mathbb{N}$,

1. $\tilde{p}(\mathbf{x}_2^{(n)}) > 0$, and
2. $-\infty < \log q(\mathbf{x}_2^{(n)}) < -|\mathbf{x}_2^{(n)}| \log 2$, and
3. $n < |\mathbf{x}_2^{(n)}| < |\mathbf{x}_2^{(n+1)}|$, and
4. $\exists b > 0, \exists t > 0, |\mathbf{x}_2^{(n)}| < bn^t$.

We describe how we identify L_2 , and prove the four claims above (which we refer to as ‘claims regarding L_2 ’ in the remainder of the proof), as follows. Let $\phi^{(n)}$ be a formula of the form $(A_1 \vee \neg A_1 \vee A_1) \wedge (A_2 \vee \neg A_2 \vee A_2) \dots (A_n \vee \neg A_n \vee A_n)$. And let $\boldsymbol{\phi}^{(n)} = \text{enc}(\phi^{(n)})$ be the encoding of $\phi^{(n)}$ (defined in §3.2). Since $\phi^{(n)}$ accepts any assignment to $A_1 \dots A_n$, $\forall \mathbf{a} \in \mathbb{B}^n, \tilde{p}(\boldsymbol{\phi}^{(n)}.\mathbf{a}) > 0$ (which checks our first claim $\forall n \in \mathbb{N}, \tilde{p}(\mathbf{x}_2^{(n)}) > 0$). Under our assumption that $\text{support}(\tilde{p}) \subseteq \text{support}(\tilde{q})$, we know $\forall \mathbf{a} \in \mathbb{B}^n, q(\boldsymbol{\phi}^{(n)}.\mathbf{a}) > 0$ as well. We define $\mathbf{x}_2^{(n)} \triangleq \boldsymbol{\phi}^{(n)}.\mathbf{a}_{\min}^{(n)}$, where $\mathbf{a}_{\min}^{(n)} = \text{argmin}_{\mathbf{a} \in \mathbb{B}^n} q(\boldsymbol{\phi}^{(n)}.\mathbf{a})$. Now we consider the decay rate of $\log q(\mathbf{x}_2^{(n)})$. Regardless of the choice of q , $q(\mathbf{x}_2^{(n)}) = q(\boldsymbol{\phi}^{(n)}.\mathbf{a}_{\min}^{(n)}) \leq q(\boldsymbol{\phi}^{(n)}.\frac{1}{2^n})$. Since $\log q(\boldsymbol{\phi}^{(n)}) \leq 0$, $\log q(\mathbf{x}_2^{(n)}) \leq -n \log 2 < -|\mathbf{x}_2^{(n)}| \log 2$. This checks our second claim $\forall n \in \mathbb{N}, -\infty < \log q(\mathbf{x}_2^{(n)}) < -|\mathbf{x}_2^{(n)}| \log 2$. Moreover, $\forall n \in \mathbb{N}, |\boldsymbol{\phi}^{(n)}| \leq |\boldsymbol{\phi}^{(n+1)}|$ (because ϕ_{n+1} has strictly more variables than ϕ_n), and $|\mathbf{a}_{\min}^{(n)}| + 1 = n + 1 = |\mathbf{a}_{\min}^{(n+1)}|$. Therefore $n < |\mathbf{x}_2^{(n)}| = |\boldsymbol{\phi}^{(n)}| + |\mathbf{a}_{\min}^{(n)}| < |\boldsymbol{\phi}^{(n+1)}| + |\mathbf{a}_{\min}^{(n+1)}| = |\mathbf{x}_2^{(n+1)}|$, checking our third claim. Finally, we note that $\forall n \in \mathbb{N}, \mathbf{x}_2^{(n)}$ can be generated in $O(\text{poly}(n))$, which implies that $\exists b > 0, \exists t > 0, |\mathbf{x}_2^{(n)}| < bn^t$, checking our fourth claim.

Since L_2 is a family of strictly increasing strings (third claim regarding L_2) and longer strings have lower upperbounds to their weights under q (second claim regarding L_2), together with the \mathbf{x}_1 condition, we can identify a subset of L_2 whose strings have lower weights under q than \mathbf{x}_1 . Specifically, we know that there is a polynomial function f such that $\forall \mathbf{x}_2 \in L_2, q(\mathbf{x}_1) > q(\mathbf{x}_2)$ if $|\mathbf{x}_2| > f(|\mathbf{x}_1|)$. Let these strings form a subset $L_2^{(|\mathbf{x}_1|)} \subseteq L_2$. Now we show that there exists a polynomial function of $|\mathbf{x}_1|$, such that there is a non-empty subset of any such $L_2^{(|\mathbf{x}_1|)}$, whose string lengths are bounded below that function as well. We know such function exists, since we know

$\log q(\mathbf{x}_1)$ is bounded *above* $-a|\mathbf{x}_1|^s$, and that $\forall \mathbf{x}_2^{(n)} \in L_2, \log q(\mathbf{x}_2^{(n)})$ is bounded *below* some polynomial function of bn^t (the second and fourth claims regarding L_2). Therefore there exists a family of strings $L_2^{(|\mathbf{x}_1|)'} = \{\mathbf{x}_2^{(n(|\mathbf{x}_1|))} \mid |\mathbf{x}_1| \in \mathbb{N} \cup \{0\}\} \subseteq L_2^{(|\mathbf{x}_1|)}$ where $n : \mathbb{N} \cup \{0\} \rightarrow \mathbb{N}$ is a polynomial function, such that by combining the \mathbf{x}_1 condition and our second claim regarding L_2 , we have $\forall \mathbf{x}_1 \in V^*, \log q(\mathbf{x}_1) > -a|\mathbf{x}_1|^s > -b \log 2(n(|\mathbf{x}_1|))^t > \log q(\mathbf{x}_2^{(n(|\mathbf{x}_1|))})$. From our fourth claim regarding L_2 , we know $\forall \mathbf{x}_1 \in V^*, \forall \mathbf{x}_2 \in L_2^{(|\mathbf{x}_1|)'}, |\mathbf{x}_2| \leq f_{\tilde{q}}(|\mathbf{x}_1|) = b(n(|\mathbf{x}_1|))^t$, which is what our theorem claims. Also from Theorem 2 and Lemma 3 we know $\exists \mathbf{x}_1 \in V^*$ such that $\tilde{p}(\mathbf{x}_1) = 0, \tilde{q}(\mathbf{x}_1) > 0$. And from our first claim regarding L_2 we know $\forall \mathbf{x}_2 \in L_2^{(|\mathbf{x}_1|)'}, \tilde{p}(\mathbf{x}_2) > 0$. Therefore given any \mathbf{x}_1 , we can find a counterexample $\mathbf{x}_2^{(n(|\mathbf{x}_1|))} \in L_2^{(|\mathbf{x}_1|)'}$. \square

Lemma A.1. *The first part of Theorem 4 (without the modifications (a) and (b)).*

We first prove the first part of Theorem 4 (which is restated in full below). In this case we will use a distribution \tilde{p} that does not have support V^* (so it does not prove modification (b)).

Proof. We take \tilde{p} to be the weighted language that was defined in §3.2, which was already shown to be efficiently computable. Suppose (M^q, Θ^q, λ) is a counterexample to Lemma A.1. Choose integer $k \geq 1$ in a manner (dependent only on λ) to be described at the end of the proof.

Suppose we would like to answer SAT where ϕ is a formula with variables A_1, \dots, A_j . Define $\phi' = (\neg A_1 \wedge \neg A_2 \wedge \dots \wedge \neg A_j \wedge \neg A_{j+1} \wedge \dots \wedge \neg A_{j+k}) \vee (A_1 \wedge \text{Shift}(\phi))$. Note that ϕ' augments ϕ with k additional variables, namely A_1 and A_{j+2}, \dots, A_{j+k} . For $k = 1$, this is the same construction as in the proof of Theorem 1. Let $n = |\phi'|$ and note that n is polynomial in the size of ϕ (holding k constant).

The strings in $L = \text{support}(\tilde{p})$ that begin with ϕ' are precisely the strings of the form $\phi' \mathbf{a}'$ where \mathbf{a}' is a satisfying assignment of ϕ' . This is achieved precisely when $\mathbf{a}' = \mathbf{0}^{j+k}$ or $\mathbf{a}' = \mathbf{1a}\vec{b}$ where \mathbf{a} is a satisfying assignment of ϕ and $\vec{b} \in \mathbb{B}^{k-1}$.

By our definition of \tilde{p} , all strings in L that begin with ϕ' have equal weight under \tilde{p} . Call this weight w .¹⁷ Clearly $Z(\phi' \mathbf{0}) = w$, and $Z(\phi' \mathbf{1}) = w \cdot 2^{k-1}$. (number of satisfying assignments of ϕ).

¹⁷Specifically, each such string has length $n + j + k$, so \tilde{p} gives it a weight of $w = (\frac{1}{3})^{n+j+k+1}$.

Recall that $p(\mathbf{0} \mid \phi') = Z(\phi'\mathbf{0})/(Z(\phi'\mathbf{0}) + Z(\phi'\mathbf{1}))$. Let us abbreviate this quantity by p . It follows from the previous paragraph that if ϕ is unsatisfiable, then $p = 1$, but if ϕ is satisfiable, then $p \leq 1/(1+2^{k-1})$. By hypothesis, p is approximated (with error probability $< 1/3$) by the possibly random quantity $(M^q(\theta_{|\phi'|}^q))(\phi'\mathbf{0})$, which we abbreviate by q , to within a factor of λ . That is, $p \in [q/\lambda, \lambda q]$. By choosing k large enough¹⁸ such that $[q/\lambda, \lambda q]$ cannot contain both 1 and $1/(1+2^{k-1})$, we can use q to determine whether $p = 1$ or $p \leq 1/(1+2^{k-1})$. This allows us to determine $\text{SAT}(\phi)$ in polynomial time with error probability $< 1/3$, since by hypothesis q is computable in polynomial time with compact parameters. This shows that $\text{SAT} \in \text{BPP/poly} = \text{P/poly}$, implying $\text{NP} \subseteq \text{P/poly}$, contrary to our assumption. (BPP/poly is similar to P/poly but allows M^q to be a bounded-error probabilistic Turing machine.) \square

Theorem 4. *Assuming $\text{NP} \not\subseteq \text{P/poly}$, there exists an efficiently computable weighted language $\tilde{p} : V^* \rightarrow \mathbb{R}_{\geq 0}$ such that there is no (M^q, Θ^q) where $\Theta^q = \{\theta_n^q \mid n \in \mathbb{N}\}$ that satisfies all of the following properties (similar to §3.1):*

- *the parameter size $|\theta_n^q|$ grows only as $O(\text{poly}(n))$*
- *$M^q(\theta_n^q)$ returns a probabilistic Turing machine q_n in time $O(\text{poly}(n))$*
- *there exists $\lambda \geq 1$ such that for each $x \in V \cup \{\$\}$ and $\hat{\mathbf{x}} \in V^*$ with $|\hat{\mathbf{x}}| \leq n$ and $p(x \mid \hat{\mathbf{x}}) > 0$, the probabilistic computation $q_n(\hat{\mathbf{x}}x)$ has probability $> 2/3$ of approximating $p(x \mid \hat{\mathbf{x}})$ to within a factor of λ (that is, $q_n(\hat{\mathbf{x}}x)/p(x \mid \hat{\mathbf{x}}) \in [1/\lambda, \lambda]$)*
- *q_n runs on those inputs $\hat{\mathbf{x}}x$ in time $O(\text{poly}(n))$*

Moreover, the statement above still remains true with either of the following modifications:

- (a) *the approximation guarantee is only required to hold for prefixes $\hat{\mathbf{x}}$ such that $\{\mathbf{x} : \hat{\mathbf{x}} \leq \mathbf{x}\}$ is finite (so that $p(x \mid \hat{\mathbf{x}})$ is computable by brute force)*
- (b) *support(\tilde{p}) = V^**

Proof. It remains to show that the statement remains true with modification (a) and with modification (b). For (a), the proof of Lemma A.1 suffices, since it reduces SAT to approximate local probability queries of the stated form. That is, the true local probabilities $p(x \mid \hat{\mathbf{x}})$ that can be computed with finite summations, thanks to the structure of our example language \tilde{p} , which guarantees that the prefix $\hat{\mathbf{x}}$ can only continue with suffixes of a fixed length that is easily determined from $\hat{\mathbf{x}}$.

¹⁸It suffices to ensure that $1 + 2^{k-1} > \lambda^2$, so take any $k > 1 + \log_2(\lambda^2 - 1)$.

For modification (b), again let $V = \mathbb{B} = \{\mathbf{0}, \mathbf{1}\}$. Choose some $\epsilon > 0$ (any choice will do), and let

$$\tilde{p}_1(\mathbf{x}) = \begin{cases} (\frac{1}{3})^{|\mathbf{x}+1|} & \text{if } \mathbf{x} = \phi\mathbf{a} \text{ where } \phi = \text{enc}(\phi) \\ & \text{and } \mathbf{a} \text{ satisfies } \phi \\ 0 & \text{otherwise} \end{cases}$$

$$\tilde{p}_2(\mathbf{x}) = (\frac{1}{9})^{|\mathbf{x}+1|} > 0$$

$$\tilde{p}(\mathbf{x}) = \tilde{p}_1(\mathbf{x}) + \epsilon \cdot \tilde{p}_2(\mathbf{x})$$

We use Z_1 , Z_2 , and Z respectively to denote normalizing constants of these three weighted languages. Note that \tilde{p}_1 is the weighted language that was previously used in the proofs of Theorem 1 and Lemma A.1. Our new \tilde{p} is intended to be very similar while satisfying the additional condition (b). It is easy to show that \tilde{p} is efficiently computable, much as we showed for \tilde{p}_1 in Theorem 1. Also, \tilde{p} is normalizable, since $Z = Z_1 + \epsilon \cdot Z_2$, where $Z_1 \leq (\frac{1}{3})/(1 - \frac{2}{3}) = 1$ and $Z_2 = (\frac{1}{9})/(1 - \frac{2}{9}) = \frac{1}{7}$ are both finite.

The proof proceeds as in Lemma A.1, with ϕ' constructed from ϕ as before. Recall that ϕ has j variables, ϕ' has $j + k$ variables, and $|\phi'| = n$. We may assume WLOG that the encoding function enc is such that an encoded formula always has at least as many bits as the number of variables in the formula, so $n \geq j + k$.

Notice that $Z_1(\phi')$ sums over the satisfying assignments of ϕ' , and there may be as few as one of these (if ϕ is unsatisfiable). By contrast, $Z_2(\phi')$ sums over an infinite number of continuations with positive probability. The faster decay rate of $\frac{1}{9}$ in \tilde{p}_2 was chosen to keep $Z_2(\phi')$ small relative to $Z_1(\phi')$ despite this. Specifically,

$$Z_1(\phi'\mathbf{0}) = (\frac{1}{3})^{n+j+k+1}$$

$$Z_1(\phi'\mathbf{1}) = (\frac{1}{3})^{n+j+k+1} \cdot 2^{k-1} \cdot (\# \text{ of satisfying assignments of } \phi)$$

$$Z_2(\phi'\mathbf{0}) = (\frac{1}{9})^n \cdot \frac{1}{9} \cdot (\frac{1}{9}/(1 - \frac{2}{9}))$$

$$= \frac{1}{7} \cdot (\frac{1}{3})^{2(n+1)}$$

$$< \frac{1}{7} \cdot Z_1(\phi'\mathbf{0})$$

(because $2(n+1) > n + j + k + 1$)

$$Z_2(\phi'\mathbf{1}) = Z_2(\phi'\mathbf{0})$$

As in the proof of Lemma A.1, we will show that $p(\mathbf{0} \mid \phi')$ is much larger when ϕ is unsatisfiable. Recall that $Z(\hat{\mathbf{x}}) = Z_1(\hat{\mathbf{x}}) + \epsilon \cdot Z_2(\hat{\mathbf{x}})$. When ϕ has

zero satisfying assignments,

$$\begin{aligned} p(\mathbf{0} \mid \phi') &= \frac{Z(\phi'\mathbf{0})}{Z(\phi'\mathbf{0}) + Z(\phi'\mathbf{1})} \\ &= \frac{Z(\phi'\mathbf{0})}{Z_1(\phi'\mathbf{0}) + \epsilon \cdot Z_2(\phi'\mathbf{0}) + \epsilon \cdot Z_2(\phi'\mathbf{1})} \\ &> \frac{Z(\phi'\mathbf{0})}{Z_1(\phi'\mathbf{0}) + 2 \cdot \frac{\epsilon}{7} \cdot Z_1(\phi'\mathbf{0})} \end{aligned}$$

whereas if ϕ has at least one satisfying assignment, then

$$\begin{aligned} p(\mathbf{0} \mid \phi') &= \frac{Z(\phi'\mathbf{0})}{Z(\phi'\mathbf{0}) + Z(\phi'\mathbf{1})} \\ &< \frac{Z(\phi'\mathbf{0})}{Z_1(\phi'\mathbf{0}) + Z_1(\phi'\mathbf{1})} \\ &\leq \frac{Z(\phi'\mathbf{0})}{Z_1(\phi'\mathbf{0}) + 2^{k-1} Z_1(\phi'\mathbf{0})} \end{aligned}$$

This rewrites both probabilities in terms of $Z(\phi'\mathbf{0})$ quantities, which do not depend on the number of satisfying assignments. So now we can see that the first probability is at least $(1 + 2^{k-1}) / (1 + \frac{2\epsilon}{7})$ times as large as the second probability. Choose k large enough¹⁹ such that $[q/\lambda, \lambda q]$ cannot contain both probabilities, and complete the proof as in Lemma A.1. \square

Theorem 5. *There exists a light marginalization p of an ELN distribution, such that $\text{support}(p)$ is an NP-complete language.*

Proof. We will construct p such that $\text{support}(p)$ is the NP-complete language SAT of all satisfiable boolean formulas. The idea is to construct an ELN distribution r that can autoregressively generate any assignment \mathbf{a} followed by any formula ϕ that is satisfied by \mathbf{a} . Thus, if we delete the \mathbf{a} prefixes, the support consists of exactly the satisfiable formulas ϕ (or more precisely, their encodings ϕ).

To be more precise, we will have $\text{support}(r)$ be the language $L = \{1^j \mathbf{0} \mathbf{a} \phi \mid j \geq 0 \text{ and } \mathbf{a} \in \mathbb{B}^j \text{ and } \phi \text{ is a formula satisfied by } \mathbf{a}\}$. This is defined similarly to the support language L in §3.2, but with the order of ϕ and \mathbf{a} crucially swapped: r will now generate the “solution” \mathbf{a} *before* the “problem” ϕ . We have prepended to \mathbf{a} a unary encoding of j , the number of variables, so that it is clear where \mathbf{a} ends and ϕ begins. The marginalization operator μ maps $1^j \mathbf{0} \mathbf{a} \phi$ to ϕ , which deletes the first $2j + 1$ symbols.

As in §3.2, we require ϕ to use all of the variables A_1, \dots, A_j (and only those variables), implying that

¹⁹It suffices to ensure that $(1 + 2^{k-1}) / (1 + \frac{2\epsilon}{7}) > \lambda^2$, so take any $k > 1 + \log_2(\lambda^2 \cdot (1 + \frac{2\epsilon}{7}) - 1)$.

$|\phi| \geq j$. This ensures that marginalizing over the $2j + 1$ latent symbols is only light marginalization since $2j + 1 + |\phi| \in O(\text{poly}(|\phi|))$. For convenience, we will also require ϕ to be a CNF formula. These requirements shrink $\text{support}(p)$ but do not affect its NP-completeness.

The remaining challenge is to construct an autoregressive distribution r whose support is L . We can think of this distribution as describing an efficient procedure for randomly generating a string from left to right so that the procedure terminates with probability 1,²⁰ runs in time that is polynomial in the length of the resulting string (so that it is ELN), has positive probability of producing any string in L , and has zero probability of producing any string not in L . Below we give such a procedure.²¹

1. First, the procedure generates a string of the form $1^j \mathbf{0}$. At each step, it chooses uniformly from $\{\mathbf{0}, 1\}$ until it generates a $\mathbf{0}$. Let j be the number of 1 symbols generated so far. For example, if $111\mathbf{0}$ was generated, then $j = 3$.
2. Next, the procedure generates \mathbf{a} as a sequence of j random symbols from $\{\mathbf{0}, 1\}$, again making a uniform draw at each step. In our example, it might generate $\mathbf{010}$.
3. Finally, the procedure must generate the encoding ϕ of a random CNF formula ϕ that is satisfied by \mathbf{a} , such as $(A_2 \vee \neg A_3 \vee \neg A_2 \vee A_2) \wedge (\neg A_1)$ in our example. This involves generating a random sequence of 0 or more satisfied clauses connected by \wedge . At each step, the procedure decides whether to generate a new clause or end the formula. The probability of generating a new clause is ordinarily $1/2$. However, this probability is 1 if the previous clauses do not yet mention all the variables A_1, \dots, A_j .

How does it generate each satisfied disjunctive clause? This involves generating a sequence of literals connected by \vee , at least one of which must be true. At each step of this subroutine, it uniformly chooses an integer $i \in [1, j]$, and

²⁰It is clear that the procedure below terminates in finite time almost surely (i.e., with probability 1), so that r is in fact a consistent probability distribution over the finite strings V^* (see footnote 7). Phase 1 almost surely terminates after a finite number of 1’s. Phase 2 always terminates. Phase 3 almost surely terminates after a finite number of clauses, and each clause almost surely terminates after a finite number of literals.

²¹Our presentation here makes use of an infinite alphabet that includes symbols such as A_i and $\neg A_i$ for all $i \in \mathbb{N}_{>0}$, as well as symbols such as $\mathbf{0}, 1, \wedge, \vee$. We implicitly invoke some prefix-free encoding scheme to translate each symbol into a fixed string over a finite alphabet V . We note that such translation can be done in $O(\text{poly}(j))$.

then flips a fair coin to decide whether to add the literal A_i or $\neg A_i$ to the current clause. If the clause is now satisfied by \mathbf{a} (i.e., at least one of the literals is true), it then flips another fair coin to decide whether to end the clause.

r is ELN because there exists a Turing machine which computes from input $\hat{\mathbf{x}}$ —in time $O(\text{poly}(|\hat{\mathbf{x}}|))$ —the probability that the next symbol generated after the prefix $\hat{\mathbf{x}}$ would be x , under the above procedure. \square

Theorem 6. *The following statements are equivalent for any nonempty $L \subseteq V^*$:*

- (a) $L \in \text{NP/poly}$.
- (b) L is the support of a light marginalization of an ELNCP distribution.
- (c) L is the support of a light marginalization of an ECCP weighted language.

Proof. (b) implies (c) since any ELNCP distribution is an ECCP weighted language (Lemma 2). (c) implies (a) by Lemma A.2 below. Finally, (a) implies (b) by Lemma A.3 below. \square

Lemma A.2. *For any ECCP weighted language \tilde{r} , if \tilde{p} is a light marginalization of \tilde{r} , then $\text{support}(\tilde{p}) \in \text{NP/poly}$.*

Notice that this lemma concerns the class NP/poly, not P/poly (see §2.4). The proof is straightforward.

Proof. Suppose \tilde{r} is ECCP via $(M^{\tilde{r}}, \theta^{\tilde{r}})$, and μ is the marginalization operator such that $\tilde{p}(\mathbf{x}) = \sum_{\mathbf{z}: \mu(\mathbf{z})=\mathbf{x}} \tilde{r}(\mathbf{z})$. By the light marginalization assumption, there is a polynomial f such that $|\mathbf{z}| \leq f(|\mu(\mathbf{z})|)$.

To prove $\text{support}(\tilde{p}) \in \text{NP/poly}$, we must show that there exists (M, Θ) such that for all $n \geq 0$, a *nondeterministic* Turing machine M_n can be constructed as $M(\theta_n)$ in time $O(\text{poly}(n))$, which can in turn decide in time $O(\text{poly}(n))$ whether $\tilde{p}(\mathbf{x}) > 0$ for any \mathbf{x} with $|\mathbf{x}| = n$.

Deciding $\tilde{p}(\mathbf{x}) > 0$ means deciding whether $(\exists \mathbf{z} \in V^*) \mu(\mathbf{z}) = \mathbf{x}$ and $\tilde{r}(\mathbf{z}) > 0$. But if $|\mathbf{x}| = n$, the first condition $\mu(\mathbf{z}) = \mathbf{x}$ implies $|\mathbf{z}| \leq f(|\mu(\mathbf{z})|) = f(|\mathbf{x}|) = f(n)$. Thus, we need M_n to nondeterministically check only the \mathbf{z} of length up to $f(n)$ to see whether $\mu(\mathbf{z}) = \mathbf{x}$ and $\tilde{r}(\mathbf{z}) > 0$.

How can M_n check a string \mathbf{z} of length m ? It can decide the first condition $\mu(\mathbf{z}) = \mathbf{x}$ in time $O(\text{poly}(m))$, since the marginalization operator μ is a polytime function. To decide the second condition $\tilde{r}(\mathbf{z}) > 0$, it must construct the (deterministic) Turing machine $M^{\tilde{r}}(\theta_m^{\tilde{r}})$ and then apply it to \mathbf{z} to obtain $\tilde{r}(\mathbf{z})$:

since \tilde{r} is ECCP, both steps take time $O(\text{poly}(m)) = O(\text{poly}(f(n))) \subseteq O(\text{poly}(n))$ as required.

However, this means that $M_n = M(\theta_n)$ must have access to the parameter vectors $\theta_m^{\tilde{r}}$ for all $m \leq f(n)$. We therefore make θ_n include this collection of parameter vectors. Each $|\theta_m^{\tilde{r}}| \in O(\text{poly}(m)) \subseteq O(\text{poly}(n))$ since \tilde{r} is ECCP. So $|\theta_n| \in O(\text{poly}(n))$ as required. \square

Lemma A.3. *For any $L \in \text{NP/poly}$, there exists a light marginalization p of an ELNCP distribution, such that $\text{support}(p) = L$.*

In this proof we will demonstrate how an ELNCP distribution r can correspond to a left-to-right stochastic string generation process such that when generation terminates, we have a string whose suffix is guaranteed to be $\in L$.

Our generative story is inspired by rejection sampling, a widely used Monte Carlo method for sampling from intractable distributions. In standard rejection sampling schemes, we first sample a string from a tractable distribution, then toss a coin to decide whether we keep the sample. If the sample is not in the support of a target distribution, we reject unconditionally. Otherwise, there is a nonzero probability that we accept the sample, then halt. If we reject a sample, we try again, until we accept (and halt). Rejection sampling offers a general framework for exploiting randomness to (approximately) solve computational problems. But it is not guaranteed to halt in finite time, which poses as a problem if we require light marginalization—we would like a guarantee that all accepted strings are relatively short to the answer it contains. We therefore make use of the polysize parameter vectors of ELNCP languages to store certain ‘example strings’ that are guaranteed to be $\in L$. Wherever ordinary rejection sampling would reject a string and try generating another, we switch to accepting a previously stored example string of an appropriate length. This places a lot of probability mass on a small number of example strings, whereas rejection sampling effectively throws away this mass and renormalizes. Unlike rejection sampling, it distorts the local probabilities, but it does preserve the support.

At a high level, r is a distribution over strings that record traces of the generative story we describe above. Formally they take the form $\mathbf{z} = \mathbf{abcde}$, where \mathbf{a} , \mathbf{b} , \mathbf{c} and \mathbf{e} are themselves strings. \mathbf{a} represents string length in L : we will accept a string $\in L$ that is at least as long as $|\mathbf{a}| - 1$. We then sample a string \mathbf{b} from a proposal distribution over $V^{|\mathbf{a}|-1}$. \mathbf{c} is an encoding

of a polysize prefix to a computation path from the Turing machine that decides on strings of length $|\mathbf{b}|$: since $L \in \text{NP/poly}$, $M_{|\mathbf{b}|}$ (following the notation established in §2.4) must halt in $\leq f_L(|\mathbf{b}|)$ steps on an input string of length $|\mathbf{b}|$, where $f_L : \mathbb{N} \rightarrow \mathbb{N}$ is a polynomial function. c is a convenience variable that indicates the outcome of rejection sampling ($c = \mathbf{0}$ for accepting \mathbf{b} , and $c = 1$ for rejecting \mathbf{b}). We always reject \mathbf{b} if it is not $\in L$: that is, $p_z(c = \mathbf{0}) > 0 \iff \mathbf{b} \in L$. Finally, \mathbf{d} is a copy of \mathbf{b} iff $\mathbf{b} \in L$ is accepted in the previous rejection sampling (indicated by $c = \mathbf{0}$). If \mathbf{b} was not accepted (indicated by $c = 1$), we let \mathbf{d} be some example string memorized in the ELNCP weighted language r 's parameter vectors that is also known to be $\in L$, and of length $\geq |\mathbf{b}|$. And after the generation of \mathbf{d} , we stop. \mathbf{d} is guaranteed to be $\in L$.

Proof. WLOG we assume L uses the Boolean alphabet \mathbb{B} . We discuss two cases of L . If L is finite, then there exists a polytime Turing machine M that accepts $\hat{\mathbf{x}} \in \mathbb{B}^*$ iff $\exists \mathbf{x} \in L, \hat{\mathbf{x}} \leq \mathbf{x}$. M encodes a finite trie, where $\forall \hat{\mathbf{x}} \in \mathbb{B}^*$, we can decide if $\hat{\mathbf{x}}$ is a valid prefix of some string $\in L$ in time $O(|\hat{\mathbf{x}}|)$ by following a path in the encoded trie. Given such an M , we can define an ELN r that computes any $r(x | \hat{\mathbf{x}})$ where the prefix probability $Z_r(\hat{\mathbf{x}}) > 0$ in $O(\text{poly}(|\hat{\mathbf{x}}|))$. Let the marginalization operator $\mu(\mathbf{x}) = \mathbf{x}$, r lightly marginalizes itself and has support L .

Now we discuss the case where L is infinite. First, we describe the set of ‘example strings’ formally. We define $\mathcal{L}_0 = \{\mathbf{x}_0^{(n)} : n \in \mathbb{N}\}$ to be a subset of L where the strings $\mathbf{x}_0^{(n)}$, speaking at a high level, is a ‘specimen’ n -th shortest string from L . And $|\mathbf{x}_0^{(n)}| \geq n - 1$. More formally: $\forall m \in \mathbb{N} \cup \{0\}$ such that $\exists \mathbf{x} \in L$ where $|\mathbf{x}| = m, \exists n \in \mathbb{N}$ such that $\mathbf{x}_0^{(n)} \in \mathcal{L}_0, \mathbf{x}_0^{(n)} \in L$, and $|\mathbf{x}_0^{(n)}| = m$. We also define that $\forall \mathbf{x}_0^{(s)} \in \mathcal{L}_0, \forall \mathbf{x}_0^{(t)} \in \mathcal{L}_0, |\mathbf{x}_0^{(s)}| > |\mathbf{x}_0^{(t)}| \iff s > t$. To identify a unique \mathcal{L}_0 , we arbitrarily define that $\forall \mathbf{x}_0^{(n)} \in \mathcal{L}_0, \mathbf{x}_0^{(n)}$ has the lowest lexicographical order among all $\mathbf{x} \in \{\mathbf{x} \in L \mid |\mathbf{x}| = |\mathbf{x}_0^{(n)}|\}$.

We define a weighted language $r(\mathbf{z}) = \prod_{t=0}^{|\mathbf{z}|-1} r(\mathbf{z}_t | \mathbf{z}_{<t})$, where $\mathbf{z}_{|z|+1} \triangleq \$$, to have support \mathcal{Z} , which contains strings of the form $\mathbf{z} = \mathbf{abcde}$. Below we (re)use t as an index variable to scan through strings $\{\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}, \mathbf{e}\}$. \mathbf{a} is of form $\mathbf{0}^*1$. We define $\mathbf{a} = \mathbf{0} \dots \mathbf{0}1 \in \mathcal{A}$, and $r(a_t = 0 \mid \mathbf{a}_{<t}) = 1/2$. We define $\mathbf{b} \in \mathbb{B}^{|\mathbf{a}|-1}$, and $r(b_t = 0 \mid \mathbf{ab}_{<t}) = r(b_t = 1 \mid \mathbf{ab}_{<t}) = 1/2$ iff $|\mathbf{b}_{<t}| < |\mathbf{a}| - 1$. \mathbf{c} is some encoding of choices made during $M_{|\mathbf{b}|}$'s execution

on \mathbf{b} . Specifically, \mathbf{c} is a list of $\lceil \log w(|\mathbf{b}|) \rceil$ -bit (also *word width* later) integers (*words*) $c_1 c_2 \dots$, where $w : \mathbb{N} \rightarrow \mathbb{R}_{>1}$ is a polynomial function, $c_t \in \mathbb{N}$ means that at timestep t , the nondeterministic Turing machine $M_{|\mathbf{b}|}$ takes the c_t -th entry in the lexicographically sorted transition relation table that has at most $O(\text{poly}(|\mathbf{b}|))$ entries, on an input string \mathbf{b} . We define $r(c \mid \mathbf{abc}_{<t}) \propto 1 \iff \mathbf{c}_{<t}c$ is a prefix to some valid transition of $M_{|\mathbf{b}|}$ on input $|\mathbf{b}|$, and $r(c \mid \mathbf{abc}_{<t}) = 0$ otherwise. d is a single bit, with $r(d = \mathbf{0} \mid \mathbf{abc}) \propto \mathbb{I}(\mathbf{c}$ ends at an ACCEPT state) and $r(c = 1 \mid \mathbf{ab}) \propto 1$. Finally, $r(e_t \mid \mathbf{abc}1\mathbf{e}_{<t}) = 1 \iff \mathbf{e}_{<t}e_t \leq \mathbf{x}_0^{|\mathbf{a}|-1}\$, and $r(e_t \mid \mathbf{abc}\mathbf{0}\mathbf{e}_{<t}) = 1 \iff \mathbf{e}_{<t}e_t \leq \mathbf{b}\$$. The marginalization operator is defined as $\mu(\mathbf{abcde}) = \mathbf{e}$.$

First we check the claim that $p(\mathbf{x}) > 0 \iff \mathbf{x} \in L$. We check both strings $\in L$ and strings $\notin L$: if $\mathbf{x} \in L$, there must exist a transition string of $M_{|\mathbf{x}|}$ such that $M_{|\mathbf{x}|}$ accepts \mathbf{x} in at most $f_L(|\mathbf{x}|)$ steps, by the definition of NP/poly. Therefore, $\exists \mathbf{a} \in \mathbb{B}^*, |\mathbf{a}| = |\mathbf{x}| + 1, \exists \mathbf{c} \in \mathbb{B}^*, |\mathbf{c}| \in O(\text{poly}(|\mathbf{x}|))$ such that $\mathbf{z} = \mathbf{axc}\mathbf{0}\mathbf{c} \in \mathcal{Z}$ and $p(\mathbf{x}) \geq r(\mathbf{z}) > 0$. On the other hand, if $\mathbf{x} \notin L$, and if there were $\mathbf{z} \in \mathcal{Z}$ such that $\mu(\mathbf{z}) = \mathbf{x}$, $p(\mathbf{x}) \geq r(\mathbf{z}) > 0$, there must be strings of form $\mathbf{z} = \{\mathbf{0}, 1\}^* \mathbf{x} \in \mathcal{Z}$ where $\mu(\mathbf{z}) = \mathbf{x}$. We first show that $\nexists \mathbf{z} = \{\mathbf{0}, 1\}^* \mathbf{0}\mathbf{x} \in \mathcal{Z}$ where $\mu(\mathbf{z}) = \mathbf{x}$. If there were such a string, $\exists \mathbf{a} \in \mathcal{A}, \mathbf{c} \in \mathbb{B}^*$ such that $\mathbf{z} = \mathbf{axc}\mathbf{0}\mathbf{x} \in \mathcal{Z}$ where \mathbf{c} encoded an accepting transition string of \mathbf{x} . Since given $\mathbf{x} \notin L$, all computation of $M_{|\mathbf{x}|}$ must reject \mathbf{x} in at most $f_L(|\mathbf{x}|)$ steps by definition, we have a contradiction. We then show $\nexists \mathbf{z} = \{\mathbf{0}, 1\}^* 1\mathbf{x} \in \mathcal{Z}$ where $\mu(\mathbf{z}) = \mathbf{x}$ either: because of the preceding 1 before suffix \mathbf{x} of \mathbf{z} , \mathbf{x} must be $\in \mathcal{L}_0$. Since $\mathcal{L}_0 \subseteq L$, this contradicts our assumption $\mathbf{x} \notin L$. Finally, $\{\mathbf{z} \mid \mathbf{z} = \{\mathbf{0}, 1\}^* \mathbf{x} \in \mathcal{Z}, \mathbf{x} = \mu(\mathbf{z})\}$ is an empty set, therefore $p(\mathbf{x}) = \sum_{\mathbf{z}: \mu(\mathbf{z})=\mathbf{x}} r(\mathbf{z}) = 0$.

Then we check whether p is lightly marginalized under r . Note that $\forall \mathbf{z} = \mathbf{abcde} \in \mathcal{Z}, |\mathbf{a}| \leq |\mathbf{e}| + 1$ (because $|\mathbf{x}_0^{(|\mathbf{a}|)}| = |\mathbf{d}| \geq |\mathbf{a}| - 1, |\mathbf{b}| = |\mathbf{a}| - 1 \leq |\mathbf{e}|, |\mathbf{c}| \leq f_L(|\mathbf{b}|) \cdot \lceil \log w(|\mathbf{b}|) \rceil \in O(\text{poly}(|\mathbf{b}|)) \subseteq O(\text{poly}(|\mathbf{e}|)), |d| = 1$. Therefore $\forall \mathbf{e} \in L, \mu(\mathbf{z}) = \mathbf{e} \implies |\mathbf{z}| \in O(\text{poly}(|\mathbf{e}|)) = O(|\mu(\mathbf{z})|)$. Therefore p is lightly marginalized under r .

Finally, we check whether r is an ELNCP language: first we check whether r defines a distribution over \mathcal{Z} . We observe that given any $\mathbf{a} \in \mathcal{A}$, the conditional probability that \mathbf{a} has a finitely long suffix under r is 1: specifically, $\forall \mathbf{a} \in \mathcal{A}, r(\mathbf{z}) = 0$ if $\mathbf{z} \in \{\mathbf{z} \in \mathbb{B}^* \mid \mathbf{a} \leq \mathbf{z}, |\mathbf{z}| > 3|\mathbf{x}_0^{(|\mathbf{a}|-1)}| + f_L(|\mathbf{x}_0^{(|\mathbf{a}|-1)}|)\}$.

$\lceil \log w(|\mathbf{x}_0^{(|\mathbf{a}|-1)}|) \rceil + 2 \}$. Therefore

$$\begin{aligned} r(|\mathbf{x}| = \infty) &= \lim_{n \rightarrow \infty} r(|\mathbf{x}| > n) \\ &= \lim_{n \rightarrow \infty} r(|\mathbf{a}| > n) \\ &= \lim_{n \rightarrow \infty} 1/2^n \\ &= 0. \end{aligned}$$

And therefore r defines a distribution over \mathcal{Z} .

We then check if there exists a family of parameter vectors such that on string \mathbf{z} , computing $r(z_t \mid \mathbf{z}_{<t}, t \in [1 \dots |\mathbf{z}| + 1])$ uses $O(\text{poly}(|\mathbf{z}|))$ time and a parameter vector of size $O(\text{poly}(|\mathbf{z}|))$. We define a family of parameter vectors $\{\theta_n^r : n \in \mathbb{N}\}$ where θ_n^r encodes both all example strings that a string of length n may use $\in \mathcal{L}_0$: $\mathcal{L}_0^{(n)} \triangleq \{\mathbf{x}_0^{(m)} \mid \mathbf{x}_0^{(m)} \in \mathcal{L}_0, n \geq |\mathbf{x}_0^{(m)}|\}$, and advice strings for input length $\leq n$: $\Theta^{(n)} \triangleq \{\theta_m \mid \theta_m \in \Theta, n \geq m\}$ of the NP/poly language L . Since $\forall \mathbf{x}_0^{(m)} \in \mathcal{L}_0^{(n)}, n \geq |\mathbf{x}_0^{(m)}|$, we have $|\mathcal{L}_0^{(n)}| \in O(\text{poly}(n))$. Therefore $\mathcal{L}_0^{(n)}$ can be stored using size $O(\text{poly}(n))$. As for the advice strings $\Theta^{(n)}$, again $\forall \theta_m \in \Theta^{(n)}, |\theta_m| \in O(\text{poly}(m)) \subseteq O(\text{poly}(n))$ can be stored using size $O(\text{poly}(n))$. Therefore $|\theta_{|\mathbf{z}|}^r| \in O(\text{poly}(|\mathbf{z}|))$.

For every $\hat{\mathbf{z}}$ such that the prefix probability $Z_r(\hat{\mathbf{z}}) > 0$, we can compute $r(\cdot \mid \hat{\mathbf{z}})$ as follows:

1. If $\hat{\mathbf{z}}$ is a proper prefix to some $\mathbf{a} \in \mathcal{A}$ (decidable in $O(|\hat{\mathbf{z}}|)$), $r(\cdot \mid \hat{\mathbf{z}})$ can be computed in $O(|\hat{\mathbf{z}}|)$. Otherwise, extract \mathbf{a} from $\hat{\mathbf{z}}$ and proceed to subsequent checks.
2. If we fall through the previous check, check if $\hat{\mathbf{z}}$ is a proper prefix to some \mathbf{ab} , where $\mathbf{b} \in \mathbb{B}^{|\mathbf{a}|-1}$. If $\hat{\mathbf{z}}$ is a proper prefix to \mathbf{ab} , $r(\cdot \mid \hat{\mathbf{z}}) = 1/2$. Otherwise, extract \mathbf{b} from $\hat{\mathbf{z}}$ and proceed to subsequent checks.
3. Let $\hat{\mathbf{c}} = \hat{\mathbf{z}}_{>|\mathbf{a}||\mathbf{b}|}$. If we fall through the previous checks, check if $\hat{\mathbf{z}}$ is a proper prefix to some \mathbf{abc} , where \mathbf{c} is a transition string of $M_{|\mathbf{b}|}$ on input string \mathbf{b} . $\hat{\mathbf{z}}$ is a proper prefix to some \mathbf{abc} iff $\hat{\mathbf{c}}$ is a prefix to valid transitions of $M_{|\mathbf{b}|}$ on input string \mathbf{b} that do not contain end in ACCEPT or REJECT states. Such a check can be done in $O(\text{poly}(|\mathbf{b}|)) \subseteq O(\text{poly}(|\hat{\mathbf{z}}|))$ (first we build $M_{|\mathbf{b}|}$ from advice string $\theta_{|\mathbf{b}|}$ as described in §2.4, then we proceed to check words on $\hat{\mathbf{c}}$). If $\hat{\mathbf{c}}$ is a proper prefix to some accepting/rejecting \mathbf{c} , we compute $r(\mathbf{c} \mid \hat{\mathbf{z}}) \propto \mathbb{I}(\hat{\mathbf{c}}\mathbf{c}$ is a prefix to some accepting/rejecting $\mathbf{c})$, $\forall \mathbf{c} \in \mathbb{B}$. $r(\mathbf{c} \mid \hat{\mathbf{z}})$ can be computed in $O(\text{poly}(|\hat{\mathbf{z}}|))$. Otherwise (if $\hat{\mathbf{z}}$ is *not* a proper prefix to some \mathbf{abc} — meaning that $\hat{\mathbf{z}}$ contains an accepting/rejecting \mathbf{c} substring),

extract \mathbf{c} from $\hat{\mathbf{z}}$ and proceed to subsequent checks.

4. If we fall through the previous checks, check if $\hat{\mathbf{z}}$ has prefix \mathbf{abc} but is a proper prefix to \mathbf{abcd} . The check follows easily from the previous step. If $\hat{\mathbf{z}}$ is a proper prefix to \mathbf{abcd} in this manner, we let $r(d = \mathbf{0} \mid \mathbf{abc}) \propto \mathbb{I}(\mathbf{c}$ ends in an ACCEPT state) and $r(d = \mathbf{1} \mid \mathbf{abc}) \propto 1$, which can be computed in $O(\text{poly}(|\hat{\mathbf{z}}|))$. Otherwise, we extract d from $\hat{\mathbf{z}}$.
5. If we fall through the previous checks, check if $\hat{\mathbf{z}}$ has prefix \mathbf{abcde} but is a proper prefix to \mathbf{abcde} . This can be checked easily by looking at d : if $c = 1, p_z(z \mid \hat{\mathbf{z}}) = 1 \iff \hat{\mathbf{z}}z \leq \mathbf{abc1x}_0^{(|\mathbf{a}|-1)}$, where $\mathbf{x}_0^{(|\mathbf{a}|-1)}$ can be looked up in $\theta_{|\hat{\mathbf{z}}|}^r$ in polytime. And $r(z \mid \hat{\mathbf{z}}) = 1 \iff \hat{\mathbf{z}}z \leq \mathbf{abc0b}$, which can also be computed in polytime.
6. Finally, if we fall through the previous checks, if $\hat{\mathbf{z}} = \mathbf{abcde}$, $r(\$ \mid \hat{\mathbf{z}}) = 1$. Otherwise $r(\$ \mid \hat{\mathbf{z}}) = 0$.

For every valid prefix $\hat{\mathbf{z}}$, we can compute $r(\cdot \mid \hat{\mathbf{z}})$ in $O(\text{poly}(|\hat{\mathbf{z}}|)) \subseteq O(\text{poly}(|\mathbf{z}|))$. Since the parameter vector $\theta_{|\mathbf{z}|}^r \in O(\text{poly}(|\mathbf{z}|))$ as previously described, p_z is ELNCP. \square

Lemma A.4. *The set $\{\tilde{p} : \tilde{p} \in \text{EC}, \tilde{p} \in \text{ELNCP}, \tilde{p} \notin \text{ELN}\}$ is not empty.*

We include this lemma to justify why this region is drawn as non-empty in Figure 2. To prove it, we first construct a weighted language \tilde{p} that is known to be in both EC and ELNCP. We then show if this $\tilde{p} \in \text{ELN}$, we would have an algorithm that solves the halting problem of deterministic Turing machines, which is impossible.

Proof. Given any unweighted language $L \subseteq \mathbb{B}^*$, we can define a normalizable weighted language \tilde{p} with support L by $\tilde{p}(\mathbf{x}) = 1/3^{|\mathbf{x}|+1}$ for $\mathbf{x} \in L$ and $\tilde{p}(\mathbf{x}) = 0$ otherwise. Moreover, if $L \in \text{P}$, then $\tilde{p} \in \text{EC}$.

For our purposes, we take L to consist of all strings of the form $\mathbf{x}^{(1)}\mathbf{x}^{(2)}$ such that $\mathbf{x}^{(1)} = \text{enc}(M)$ is a prefix-free encoding of some deterministic Turing machine M that takes an empty input tape, and $\mathbf{x}^{(2)}$ is a halting execution trace of M on the empty tape, represented as a sequence of states of M that begins with an initial state and ends at a HALT state. Note that any deterministic TM $\mathbf{x}^{(1)}$ can be paired with at most one halting execution trace $\mathbf{x}^{(2)}$, and cannot be paired with any $\mathbf{x}^{(2)}$ if it does not halt.

Clearly $L \in \mathbb{P}$: given $\mathbf{x} \in \mathbb{B}^*$, we can decide whether $\mathbf{x} \in L$ by first checking if \mathbf{x} contains well-formed strings $\mathbf{x}^{(1)}$ and $\mathbf{x}^{(2)}$. Then we build M from $\mathbf{x}^{(1)}$, then check the state transitions in $\mathbf{x}^{(2)}$ on M step-by-step. All these operations will take $O(\text{poly}(|\mathbf{x}|))$ in time. We conclude that the \tilde{p} derived from L is EC.

However, \tilde{p} is not necessarily ELNCP as desired. To fix this, we modify the construction to use a weighted language \tilde{p}' with *sparse* support L' . We will again be able to show that \tilde{p}' is EC. To show that \tilde{p}' is also ELNCP, we rely on the sparsity of L' , meaning that $\text{prefixes}(L') \triangleq \{\hat{\mathbf{x}}' : (\exists \mathbf{x}' \in L') \hat{\mathbf{x}}' \leq \mathbf{x}'\}$ contains at most $O(\text{poly}(n))$ strings $\hat{\mathbf{x}}'$ of length $\leq n + 1$. Thus, we can use $\Theta_n^{\mathbb{Q}}$ to store all of those strings $\hat{\mathbf{x}}'$ in polynomial space, along with their $Z(\hat{\mathbf{x}}')$ values.²² Notice that all strings $\hat{\mathbf{x}}' \notin \text{prefixes}(L')$ have $Z(\hat{\mathbf{x}}') = 0$, so they need not be stored. Now for any $\hat{\mathbf{x}}'$ of length $\leq n$, a Turing machine that consults $\Theta_n^{\mathbb{Q}}$ can compute $q(x | \hat{\mathbf{x}}') = Z_{\tilde{p}'(\hat{\mathbf{x}}'x)} / Z_{\tilde{p}'(\hat{\mathbf{x}}')}$ in time $O(\text{poly}(n))$ as desired.

We may define \tilde{p}' as follows. Let $\text{sparsify}(\mathbf{x})$ be a version of $\hat{\mathbf{x}}$ with many extra \emptyset symbols inserted: specifically, it inserts 2^t copies of \emptyset immediately before the t^{th} bit of \mathbf{x} , for all $1 \leq t \leq |\mathbf{x}|$. We construct \tilde{p}' so that $\tilde{p}'(\text{sparsify}(\mathbf{x})) = \tilde{p}(\mathbf{x})$. Specifically, let $L' \triangleq \text{sparsify}(L)$. The inverse function $\text{sparsify}^{-1}(\mathbf{x}')$ is defined on exactly $\mathbf{x}' \in L'$, and is unique when defined. For all $\mathbf{x}' \in \mathbb{B}^*$, let $\tilde{p}'(\mathbf{x}') \triangleq \tilde{p}(\text{sparsify}^{-1}(\mathbf{x}'))$ if $\text{sparsify}^{-1}(\mathbf{x}')$ is defined, and $\tilde{p}'(\mathbf{x}') \triangleq 0$ otherwise. This can be computed in polytime, so \tilde{p}' is EC. Also, its support L' is sparse as claimed, so \tilde{p}' is ELNCP.

Finally, we claim \tilde{p}' is not ELN. Given any Turing machine M that takes an empty input tape, M halts iff $\text{enc}(M) \in \text{prefixes}(L)$ iff $\text{sparsify}(\text{enc}(M)) \in \text{prefixes}(L')$ iff $q(x | \hat{\mathbf{x}}') > 0$ for all $\hat{\mathbf{x}}', x$ such that $\hat{\mathbf{x}}'x \leq \text{sparsify}(\text{enc}(M))$. But this would be decidable if \tilde{p}' were ELN as defined in §3.1, since then we would have a Turing machine to compute the local conditional probabilities $q(x | \hat{\mathbf{x}}')$. \square

B Implementation details of REBMs

B.1 Modeling finite subsets of infinite languages

The experiments of this paper are conducted on datasets where we only observe strings that are finitely long. We use the notation $L_{\leq T} = \{\mathbf{x} | \mathbf{x} \in L, |\mathbf{x}| \leq T\}$ for the subset of an infinite language L

²²More precisely, the first b bits of $Z(\hat{\mathbf{x}}') \leq 1$ may be stored in $\Theta_{n+b}^{\mathbb{Q}}$, when ELNCP is defined as explained in our ‘‘Remark on irrationality’’ above.

that contains all strings that are most T symbols long. Specific values of T of datasets used in experiments are listed in Appendix D.1.

B.2 Design of base models p_0

p_0 can be any distribution over $L_{\leq T}$ ²³ provided that we can sample from it, and evaluate $p_0(\mathbf{x}), \forall \mathbf{x} \in L_{\leq T}$, both in $O(\text{poly}(|\mathbf{x}|))$. In this work, we experiment with two designs of p_0 : GRU- and Transformer-based locally normalized language models. GRU-based models are used in WikiText and Yelp experiments. The GRU-based p_0 ’s are parametrized with 2-layer GRUs with 500 hidden units, and word embeddings of dimension size 500.

As for Transformer-based p_0 ’s, we make use of Grover models (Zellers et al., 2019), which effectively are GPT-2 models trained on the aforementioned REALNEWS dataset. In this work, we experiment with the ‘base’ variant of public available weights, which are 12-layered Transformers, with 12 heads, and 768 hidden units.

B.3 Design of discriminators g_θ

We formulate $g_\theta(\mathbf{x})$ as a summation of scores at positions $1 \dots |\mathbf{x}|$, passed through an activation function f :

$$g_\theta(\mathbf{x}) = f \left(\sum_{i=1}^{|\mathbf{x}|} g_t(\mathbf{x}; \theta) \right). \quad (1)$$

To verify whether lower-bounding g_θ would help with learning, as we discuss in §4.1, we experiment with two variants of f :

- **tanh**: $f(x) = 2 \cdot \tanh(x)$
- **softplus**: $f(x) = -\log(1 + \exp(x + s))$

The former one is bounded between $(-2, 2)$, while the second one has range $(-\infty, 0)$. The offset term s in the softplus activation function determines initial values of Z_θ . In this paper we set $s = 20$.

The design of $g_t(\mathbf{x}; \theta)$ follows their base model counterparts: we use Bi-GRU discriminators for GRU base models; and bi-directional Transformer discriminators for Transformer ones. For GRUs $g_t(\mathbf{x}; \theta) = \mathbf{h}_t \cdot x_t$, For Transformers $g_t(\mathbf{x}; \theta) = \sum \mathbf{h}_t$ where \mathbf{h}_t are the hidden states at time step t . In both cases, the discriminators have access to information of the whole sequence \mathbf{x} at any timestep: the Bi-GRU discriminators achieve this through the bi-directional RNNs, and the Transformers through the attention mechanism without directional masking.

²³Note that since p_0 does not have support over L , it has to assign $p(\$ | \mathbf{x}_{1..T}) = 1$, which is generally not an issue.

B.4 Training procedure

As we note in §4.1, MLE-based training methods are generally not feasible for globally normalized models. We therefore opt to train our model using the ranking variant of noise contrastive estimation (NCE) (Ma and Collins, 2018), which does not require samples from p_0 and has a simple form for residual LMs. Using p_0 as a *noise distribution*, NCE training requires minimizing the following single-sequence loss, in expectation over the true distribution p :

$$\mathcal{L}_{\text{NCE}}(\theta, \mathbf{x}, p_0, K) = -\log \frac{\tilde{p}_\theta(\mathbf{x})}{p_0(\mathbf{x})}, \quad (2)$$

where $\mathbf{x}^{(0)} \triangleq \mathbf{x}$, $\tilde{p}_\theta(\mathbf{x}) \triangleq \frac{\tilde{p}_\theta(\mathbf{x})}{p_0(\mathbf{x})}$, and $\mathbf{x}^{(1)} \dots \mathbf{x}^{(K)} \sim p_0$. Since $\tilde{p}_\theta(\mathbf{x}) = p_0(\mathbf{x}) \cdot \exp g_\theta(\mathbf{x})$, we have $\frac{\tilde{p}_\theta}{p_0}(\mathbf{x}) = \exp g_\theta(\mathbf{x})$. The NCE minimization objective (2) now reduces to the simple form

$$\begin{aligned} \mathcal{L}_{\text{NCE}}(\theta, \mathbf{x}, p_0, K) &= -g_\theta(\mathbf{x}) \\ &+ \log(\exp g_\theta(\mathbf{x}) + \sum_{k=1}^K \exp g_\theta(\mathbf{x}^{(k)})). \end{aligned} \quad (3)$$

Notice that minimizing the expected loss with stochastic gradient descent methods \mathcal{L}_{NCE} defined in equation (3) requires only evaluating sequence probabilities under g_θ , and tuning its parameters, but not the base model p_0 . We only need to generate the noise samples $\{\mathbf{x}^{(k)} \sim q \mid k \in [K]\}$ from p_0 . This way we do not need to backpropagate through parameters of the base model p_0 , which can speed up training considerably when p_0 is backed by a huge network. In fact, the training of g_θ can be completely agnostic to the design of p_0 , allowing for the application of finetuning any locally normalized p_0 .

Given the same discriminator g_θ , the difference of KL-divergence between the true model p and residual language models $\tilde{p}'_\theta(\mathbf{x}) = p'_0(\mathbf{x}) \cdot \exp g_\theta(\mathbf{x})$, and the KL-divergence between the true model and $\tilde{p}''_\theta(\mathbf{x}) = p''_0(\mathbf{x}) \cdot \exp g_\theta(\mathbf{x})$, defined with base models p'_0 and p''_0 respectively, can be written as

$$\begin{aligned} \text{KL}[p||p'_\theta] - \text{KL}[p||p''_\theta] &= \text{KL}[p||p'_0] - \text{KL}[p||p''_0] + \log \frac{Z'}{Z''}, \end{aligned} \quad (4)$$

where $Z' = \mathbb{E}_{\mathbf{x} \sim p'_0}[\exp g_\theta(\mathbf{x})]$, and Z'' is similarly defined with p''_0 . As a direct result of equation (4), we can see that finding p''_0 where $\text{KL}[p||p''_0] < \text{KL}[p||p'_0]$ implies improvement in $\text{KL}[p||p'_\theta]$ over $\text{KL}[p||p_\theta]$, under mild conditions:

Lemma 4. *If $\exists k > 0$ such that $\frac{\mathbb{E}_{\mathbf{x} \sim p'_0}[\exp g_\theta(\mathbf{x})]}{\mathbb{E}_{\mathbf{x} \sim p''_0}[\exp g_\theta(\mathbf{x})]} > \exp(-k)$ and $\text{KL}[p||p'_0] - \text{KL}[p||p''_0] > k$ then $\text{KL}[p||p'_\theta] > \text{KL}[p||p''_\theta]$.*

Proof.

$$\begin{aligned} \text{KL}[p||p'_\theta] - \text{KL}[p||p''_\theta] &= \mathbb{E}_{\mathbf{x} \sim p} [\log p''_\theta(\mathbf{x}) - \log p'_\theta(\mathbf{x})] \\ &= \mathbb{E}_{\mathbf{x} \sim p} \left[\log \frac{p''_0(\mathbf{x}) \exp g_\theta(\mathbf{x})}{\sum_{\mathbf{x}' \in L_{\leq T}} p''_0(\mathbf{x}') \exp g_\theta(\mathbf{x}')} \right. \\ &\quad \left. - \log \frac{p'_0(\mathbf{x}) \exp g_\theta(\mathbf{x})}{\sum_{\mathbf{x}' \in L_{\leq T}} p'_0(\mathbf{x}') \exp g_\theta(\mathbf{x}')} \right] \\ &= \mathbb{E}_{\mathbf{x} \sim p} \left[\log \frac{p''_0(\mathbf{x}) \exp g_\theta(\mathbf{x})}{\mathbb{E}_{\mathbf{x}' \sim p''_0}[\exp g_\theta(\mathbf{x}')]}} \right. \\ &\quad \left. - \log \frac{p'_0(\mathbf{x}) \exp g_\theta(\mathbf{x})}{\mathbb{E}_{\mathbf{x}' \sim p'_0}[\exp g_\theta(\mathbf{x}')]}} \right] \\ &= \mathbb{E}_{\mathbf{x} \sim p} [\log p''_0(\mathbf{x}) - \log p'_0(\mathbf{x})] \\ &\quad + \mathbb{E}_{\mathbf{x} \sim p} \left[\log \mathbb{E}_{\mathbf{x}' \sim p'_0}[\exp g_\theta(\mathbf{x}')] - \log \mathbb{E}_{\mathbf{x}' \sim p''_0}[\exp g_\theta(\mathbf{x}')] \right] \\ &= \text{KL}[p||p'_0] - \text{KL}[p||p''_0] \\ &\quad + \log \frac{\mathbb{E}_{\mathbf{x}' \sim p'_0}[\exp g_\theta(\mathbf{x}')]}}{\mathbb{E}_{\mathbf{x}' \sim p''_0}[\exp g_\theta(\mathbf{x}')]}}. \end{aligned} \quad (5)$$

Plugging assumptions $\frac{\mathbb{E}_{\mathbf{x} \sim p'_0}[\exp g_\theta(\mathbf{x})]}{\mathbb{E}_{\mathbf{x} \sim p''_0}[\exp g_\theta(\mathbf{x})]} > \exp(-k)$ and $\text{KL}[p||p'_0] - \text{KL}[p||p''_0] > k$ into equation (5), $\text{KL}[p||p'_\theta] - \text{KL}[p||p''_\theta] > 0$. \square

Lemma 4 suggests a training strategy that we first train the base model p_0 , then finetune g_θ : under a roughly uniform g_θ (e.g. when θ is newly initialized), $\mathbb{E}_{\mathbf{x} \sim p'_0}[\exp g_\theta]/\mathbb{E}_{\mathbf{x} \sim p''_0}[\exp g_\theta] \approx \exp(0)$; so improvements on the inclusive KL-divergence of base model $\text{KL}[p||p_0]$ will mostly translate to improvement in $\text{KL}[p||\tilde{p}_\theta]$. Optimizing the base model (i.e. finding p''_0 such that $\text{KL}[p||p''_0] < \text{KL}[p||p'_0]$) is much easier than directly minimizing $\text{KL}[p||p'_\theta]$: the former can be done by minimizing empirical cross entropy, which is computationally efficient, while the latter involves an intractable partition function $\sum_{\mathbf{x} \in L_{\leq T}} \tilde{p}'_\theta(\mathbf{x})$.

Pseudocode for fine-tuning g_θ is listed in Algorithm 1.

B.5 Computing normalized probabilities

The unnormalized probability $\tilde{p}_\theta(\mathbf{x})$ (in equation (1)) can be evaluated easily, and should suffice for (re)ranking purposes (e.g. for ASR and MT applications). However, the *normalized* probability

Algorithm 1: Pseudocode for training g_θ

Input:

- Training/validation corpora $\mathcal{D}_{\{\text{train}, \text{dev}\}}$
- base model $p_0 : L_{\leq T} \rightarrow [0, 1]$
- initial parameter vector $\theta_0 \in \mathbb{B}^d$
- noise sample size $K \in \mathbb{N}$

Output: unnormalized residual language model $\tilde{q}_\theta : L_{\leq T} \rightarrow [0, 1]$ $\theta \leftarrow \theta_0;$ /* \mathcal{L}_{NCE} is defined in equation (3) */**while** $\sum_{\mathbf{x} \in \mathcal{D}_{\text{dev}}} \mathcal{L}_{\text{NCE}}(\theta, \mathbf{x}, p_0, K)$ is still decreasing **do** **foreach** $\mathbf{x} \in \text{shuffle}(\mathcal{D}_{\text{train}})$ **do** $\nabla_\theta \mathcal{L}_{\text{NCE}} = \nabla_\theta \mathcal{L}_{\text{NCE}}(\theta, \mathbf{x}, p_0, K);$ $\theta \leftarrow \text{update-gradient}(\theta, \nabla_\theta \mathcal{L}_{\text{NCE}});$ **end****end****return** $\mathbf{x} \mapsto p_0(\mathbf{x}) + \exp g_\theta(\mathbf{x});$

$q_\theta(\mathbf{x}) \triangleq \frac{\tilde{p}_\theta(\mathbf{x})}{\sum_{\mathbf{x}} \tilde{p}_\theta(\mathbf{x})}$ does require computing the partition function Z_θ . An unbiased importance sampling estimate of $\sum_{\mathbf{x} \in L_{\leq T}} \tilde{p}_\theta(\mathbf{x})$ is

$$\begin{aligned} Z_\theta &= \sum_{\mathbf{x} \in L_{\leq T}} \tilde{p}_\theta(\mathbf{x}) \\ &= \sum_{\mathbf{x} \in L_{\leq T}} p_0(\mathbf{x}) \exp g_\theta(\mathbf{x}) \\ &= \mathbb{E}_{\mathbf{x} \sim p_0} [\exp g_\theta(\mathbf{x})] \\ &\approx \sum_{m=1}^M \frac{\exp g_\theta(\mathbf{x}^{(m)})}{M} = \hat{Z}_{\theta M}, \end{aligned} \quad (6)$$

where $\mathbf{x}^{(1)} \dots \mathbf{x}^{(M)} \sim q_0$.

C Comparison between REBMs and autoregressive models

We evaluate the effectiveness of REBMs on two different neural architectures (GRU- and Transformer-based) and 3 datasets: WikiText (Merity et al., 2017), Yelp (Yelp), and RealNews (Zellers et al., 2019), on the task of modeling sequence probabilities. An REBM \tilde{p}_θ has two components, g_θ and p_0 , and we would like to see how \tilde{p}_θ competes against p_0 itself. We do not further tune p_0 while training p_θ . As a fair comparison, we also see how p'_0 compares against p_0 , where p'_0 is simply a version of p_0 that has been trained as many additional epochs as were used to train p_θ .

p_0 models are pretrained on moderately large corpora (in GRU cases) or a very large corpus (in the

Transformer case).²⁴ We compare residual energy-based models \tilde{p}_θ to further-fine-tuned base models p'_0 , on conservatively estimated (at the low end of 95% confidence interval) token perplexity and bootstrap-sampled log likelihood improvements. The results are in Table 2. Residual energy-based models show consistent perplexity improvement compared to p'_0 that are trained on the same data using the same maximum numbers of iterations. Although the improvement in log-likelihood of p_θ over p_0 is modest (especially for RealNews experiments, where p_0 is a very strong baseline), we verify that these improvements are all statistically significant ($p < 0.05$) using bootstrapped test datasets.

We experiment with different designs of the discriminator g_θ , evaluating the effectiveness of bounding g_θ and varying its number of parameters. We find that in Transformer-based experiments, bounding g_θ considerably helps with performance; but the opposite happens for GRU-based models. We speculate that this is due to the base models' performance: the Transformer base models have high parameter count and were trained on a lot of data; and the true distribution p likely is relatively similar to p_0 , and benefits from a small hypothesis space — even though we don't know if the at-most- ϵ error assumption in §4.1 holds. On the other hand our GRU-based p_0 has neither the capacity, nor the huge amount of training data. As a result, the unbounded variant g_θ (and q_θ) may end up learning a better approximation of p .

D Experimental details

D.1 Datasets

Residual language model experiments are conducted on these datasets:

- **Segmented WikiText:** we take the standard WikiText-2 corpus (Merity et al., 2017), and segment it into sequences at new line breaks. We discard all empty lines, and any line that starts with the '=' token. In effect, we obtain sequences that are mostly entire paragraphs. We also only keep lines that are shorter than 800 tokens after BPE tokenization. Because of our preprocessing, Segmented WikiText loses much interparagraph context information, and doesn't have the 'simple' header sequences that were in the original WikiText corpus, and is much

²⁴In the Transformer case we simply take p_0 to be the Grover (Zellers et al., 2019) pretrained language model, which is based on the GPT-2 (Radford et al., 2019) architecture and performs competitively on news article generation.

Experiment (Architecture)	Model	Best configuration	log likelihood improvement (95% CI)	perplexity improvement
RealNews (Transformer)	p_θ	4-layer, tanh	$(-0.18, -0.13), \mu = -0.15$.03%
RealNews (Transformer)	p'_0	N/A	N/A	.00%
WikiText (GRU)	p_θ	1-layer/500, softplus	$(-1.85, -1.54), \mu = -1.69$	1.44%
WikiText (GRU)	p'_0	N/A	N/A	.50%
Yelp (GRU)	p_θ	2-layer/500, softplus	$(-1.89, -1.67), \mu = -1.80$	1.82%
Yelp (GRU)	p'_0	N/A	N/A	.49%

Table 2: Residual energy-based model \tilde{p}_θ improvements over autoregressive base models p_0 . The perplexity numbers are per-token, and log likelihood improvements are per sequence (in nats). We only report each dataset’s best model (according to validation data) in this table. See Appendix D for experimental details.

harder to language-model.

- **Yelp**: the Yelp dataset (Yelp) contains business reviews. As in Segmented WikiText, We keep reviews shorter than 800 tokens.
- **REALNEWS**: we make use of the standard REALNEWS corpus comes from (Zellers et al., 2019), which contains news articles that are up to 1, 024 tokens long.

In all experiments we tokenize with BPE tokenizers derived from the GPT-2 language models: the GRU models use Huggingface’s implementation²⁵ and the Transformers use Grover’s²⁶. Number of sequences in preprocessed datasets are listed in Table 3.

	Train	Dev	Test
RealNews	3, 855	1, 533	6, 158
WikiText	18, 519	878	2, 183
Yelp	10, 951	9, 964	994

Table 3: Number of sequences in preprocessed datasets (for training and tuning the discriminators g_θ , and evaluation).

D.2 Pretraining base models p_0

We use a pretrained Grover model as the base model in RealNews experiments. For GRU-based experiments, we train base models on WikiText and Yelp datasets using separate training and validation splits than those of the discriminator g_θ (Table 4). The base models are periodically (every 1, 000 iterations) evaluated on the validation split for early stopping, where we stop if there is no improvement on validation perplexity for 10 consecutive evaluations. The base models q_θ achieve 113.98 for Segmented WikiText, and 110.89 in test set perplexity, respectively. Note that these base models are further fine-tuned on additional datasets in our comparison against residual language models.

²⁵<https://github.com/huggingface/transformers>

²⁶<https://github.com/rowanz/grover>

	Train	Dev
WikiText	17, 556	1, 841
Yelp	9, 954	1, 000

Table 4: Number of sequences in preprocessed datasets (for training and tuning the base model q). Note that we do not train our own base models for RealNews, but use one of the pretrained models provided by (Zellers et al., 2019).

D.3 Metrics

We evaluate the relative performance of residual language models against autoregressive models (i.e. fine-tuned base models) on two metrics, log likelihood and perplexity improvement, which are approximated as follows:

- **Log likelihood improvement**: since p, p_θ and q_0 are all distributions over $L_{\leq T}$, we can quantitatively evaluate their difference in log likelihood. We measure the difference between $\text{KL}[p||p_\theta]$ and $\text{KL}[p||p_0]$:²⁷

$$\begin{aligned}
& \text{KL}[p||p_\theta] - \text{KL}[p||p_0] \\
&= \mathbb{E}_{\mathbf{x} \sim p} [\log p_\theta(\mathbf{x}) - \log p_0(\mathbf{x})] \\
&= \mathbb{E}_{\mathbf{x} \sim p} [\log \tilde{p}_\theta(\mathbf{x}) - \log p_0(\mathbf{x})] - \log Z_\theta \\
&= \mathbb{E}_{\mathbf{x} \sim p} [g_\theta(\mathbf{x})] - \log Z_\theta \\
&\approx \frac{\sum_{\mathbf{x} \in \mathcal{D}_{\text{test}}} g_\theta(\mathbf{x})}{|\mathcal{D}_{\text{test}}|} - \log \hat{Z}_{\theta M}, \tag{7}
\end{aligned}$$

where $\hat{Z}_{\theta M}$ is estimated using equation (6). A negative value of log likelihood difference indicates that \tilde{q}_θ approximates p better than p_0 in terms of KL-divergence.

- **Perplexity improvement**: perplexity is a common language modeling metric. Following

²⁷Note that p_0 here is the base model component of \tilde{p}_θ . While comparing between residual language models and autoregressive models, we also finetune p_0 on additional data to get a new model q'_0 , which has different parameters than p_0 .

(Rosenfeld et al., 2001), we compute

$$\begin{aligned} & \text{perplexity improvement of } p_\theta \\ &= \frac{\exp \frac{|\mathcal{D}| \log \hat{Z}_{\theta M}}{w(\mathcal{D}_{\text{test}})}}{\exp \frac{\sum_{\mathbf{x} \in \mathcal{D}_{\text{test}}} g_\theta(\mathbf{x})}{w(\mathcal{D}_{\text{test}})}}, \end{aligned} \quad (8)$$

where $w(\mathcal{D})$ is the total token count of dataset \mathcal{D} , and $|\mathcal{D}|$ is the number of sequences of \mathcal{D} . $\hat{Z}_{\theta M}$ is ecomputed Appendix B.5

Both evaluation metrics involve estimating the partition function with $\hat{Z}_{\theta M}$. For the perplexity improvement metric, we obtain 32 estimates of $\hat{Z}_{\theta M}$ ²⁸, which are normally distributed, and compute equation (8) using $\hat{Z}_{\theta M}$ the conservative end of a 95% confidence level. To account for variance in our test datasets, we further make use of bootstrapping estimation for log likelihood improvement: we bootstrap-sample 1,000 subsamples for each test dataset, and compute equation (7) for each datapoint in the Cartesian product ($1,000 \times 32$ in total). We then report results at the 2.5% and 97.5% percentiles.

D.4 Hyperparameters

Transformer experiments. We train our models on 64 GPUs across 8 nodes, with a total batch size of $64 \times 8 \times 2 = 1,024$, and with 1 noise sequence ($K = 1$ in Appendix B.4) per batch. We use an initial learning rate of $5e - 5$. The rest of the hyperparameters largely follow settings in (Zellers et al., 2019). Optimization is done with the Grover implementation of AdaFactor.

GRU experiments. We train our models on 8 GPUs on a single node, with a total batch size of $8 \times 2 = 16$, and with 25 noise sequences ($K = 25$ in Appendix B.4) per batch. We have an initial learning rate of $1e - 4$. Upon no improvement on validation data, we half the learning rate, with patience = 1. The model parameters are l_2 regularized with a coefficient of $1e - 5$. We also apply dropout regularization with $p = 0.5$. Optimization is done with PyTorch-supplied Adam.

D.5 Configurations

We study the effects of these configurations:

- **Bounding g_θ :** we note in §4.1 that with the strong hypothesis that the base model p_0 has bounded error, g_θ will have a bounded range, and leads to a much smaller hypothesis space. In this work we experiment with both bounded and unbounded g_θ 's, with ranges $(-\infty, 0)$ and

$(-2, 2)$ respectively. More details can be found in Appendix B.3.

- **Model capability of g_θ :** we hypothesize that the expressiveness of g_θ does not need to be as rich as the parametrization of p_0 , since g_θ essentially only has to tell whether the sequence \mathbf{x} comes from p or p_0 . For the GRU + WikiText experiments, we experiment with $\{1, 2\}$ -layer GRU models of g_θ . For 1-layer models, we additionally experiment with a setup that has only 250 hidden units. For the Transformers/RealNews dataset, we experiment with $\{12, 4\}$ -layer Transformer models.

D.6 Log likelihood improvements under different configurations

We also see in Table 5 that using tanh as the activation function f does better than softplus for Transformers; but performs very poorly for GRUs. We also observe degeneracy problems. We speculate that our Transformer-based base models q_θ have already learned a good approximation of the true distribution; and limiting the model capacity of g_θ in exchange of smaller variance results in a favorable trade-off, and vice versa for GRUs. Regarding discriminator capability: we see that performance is not sensitive to model size. Our best Transformers run actually is from the smaller-model runs. And the 1-layer 500-unit GRU models achieve best performance. Overall, results in Table 5 suggests that performance is sensitive to the choice of model configuration.

²⁸We set $M = 512$ in this paper.

Model Size	Activation	log likelihood improvement	
		95% CI	μ
RealNews (Transformers)			
12-layer	softplus	(-0.13, 0.08)	-0.09
12-layer	tanh	(-0.14, -0.10)	-0.12
4-layer	softplus	(-0.15, 2.62)	-0.02
4-layer	tanh	(-0.18, -0.13)	-0.16
WikiText (GRUs)			
2-layer / 500	tanh	(-0.00, 0.00)	-0.00
2-layer / 500	softplus	(-1.32, -0.85)	-1.18
1-layer / 500	tanh	(-0.79, -0.64)	-0.71
1-layer / 500	softplus	(-1.85, -1.54)	-1.69
1-layer / 250	tanh	(-0.02, 0.02)	-0.00
1-layer / 250	softplus	(-1.85, -1.46)	-1.67
Yelp (GRUs)			
2-layer / 500	tanh	(-0.03, 0.01)	-0.02
2-layer / 500	softplus	(-1.89, -1.67)	-1.80
1-layer / 500	tanh	(-0.65, -0.57)	-0.61
1-layer / 500	softplus	(-2.62, -2.03)	-2.43
1-layer / 250	tanh	(-0.00, 0.00)	-0.00
1-layer / 250	softplus	(-2.25, -1.99)	-2.13

Table 5: Comparison of different configurations.