

# Global Autoregressive Models for Data-Efficient Sequence Learning

**Tetiana Parshakova**  
Stanford University\*

tetianap@stanford.edu

**Jean-Marc Andreoli**  
Naver Labs Europe

{jean-marc.andreoli,marc.dymetman}@naverlabs.com

**Marc Dymetman**  
Naver Labs Europe

## Abstract

Standard autoregressive seq2seq models are easily trained by max-likelihood, but tend to show poor results under small-data conditions. We introduce a class of seq2seq models, GAMs (Global Autoregressive Models), which combine an autoregressive component with a log-linear component, allowing the use of global *a priori* features to compensate for lack of data. We train these models in two steps. In the first step, we obtain an *unnormalized* GAM that maximizes the likelihood of the data, but is improper for fast inference or evaluation. In the second step, we use this GAM to train (by distillation) a second autoregressive model that approximates the *normalized* distribution associated with the GAM, and can be used for fast inference and evaluation. Our experiments focus on language modelling under synthetic conditions and show a strong perplexity reduction of using the second autoregressive model over the standard one.

## 1 Introduction

Neural sequential text generation models have become the standard in NLP applications such as language modelling, NLG, machine translation. When enough data is available, these models can be trained end-to-end with impressive results. Generally, inference and training proceed in an auto-regressive manner, namely, the next decoded symbol is predicted by a *locally normalized* conditional distribution (the “softmax”). This has several advantages: (i) the probability of the sequence is already normalized, by the chain-rule over local decisions, (ii) max-likelihood (ML) training is easy, because the log-likelihood of the full sequence is simply the sum of local CE (cross-entropy) losses, (iii) exact sampling of full se-

\* Work conducted during an internship at NAVER Labs Europe.

quences from the model distribution is directly obtained through a sequence of local sampling decisions.

However, these autoregressive models (AMs) tend to suffer from a form of myopia. They have difficulty accounting for global properties of the predicted sequences, from overlooking certain aspects of the semantic input in NLG to duplicating linguistic material or producing “hallucinations” in MT, and generally through being unable to account for long-distance consistency requirements that would be obvious for a human reader.<sup>1</sup>

The main contributions of this paper are as follows.

*First*, we propose a hybrid seq2seq formalization, the *Global Autoregressive Model* (GAM), that combines a local autoregressive component with a global log-linear component, allowing the use of *a priori* features to compensate for the lack of training data. GAMs are related both to the class of Energy-Based Models (EBM) and to that of Exponential Families (EF), and inherit some important properties from those: an intimate relationship between training and sampling (EBM); the identity of empirical and model expectations at maximum-likelihood; convexity of log-likelihood (EF).

*Second*, we propose a training procedure in two steps. In the first step, we train through max-likelihood a GAM, which however is *unnormalized* and improper for fast inference or evaluation. In the second step, we use this GAM to train (by distillation) a second autoregressive model that approximates the *normalized* distribution associated with the GAM, and can be used for fast inference

<sup>1</sup>To borrow terminology from Reinforcement Learning (RL) (Sutton and Barto, 2018), such NLP models work by “imitation learning”, without any representation of “objectives” to be realized. While this defect can be mitigated in the presence of large training sets, it can become serious when this condition is not met.

and evaluation.

*Third*, we demonstrate the ability of GAMs to be data-efficient, namely, to exploit the original data better than a standard autoregressive model. In order to clarify the core techniques and issues, we design a simple class of synthetic data, consisting of random binary strings containing “motifs” (specific substrings) that we can manipulate in different ways. We show that, in limited data conditions, GAMs are able to exploit the features to obtain final autoregressive models that perform better than the original ones.

The remainder of the paper is structured as follows. In Section 2, we provide some background about autoregressive models, energy-based models, and log-linear models. In Section 3, we introduce GAMs. In section 4, we describe our focus on synthetic data. In Section 5, we explain our training procedure. In Section 6, we comment on related work. In Section 7, we describe our experiments. In Section 8, we provide an analysis of our results. We conclude with a discussion in Section 9. *Note that some additional explanations and experiments are provided in the Supplementary Material, indicated by [SM].*

## 2 Background

### 2.1 Autoregressive models (AM)

These are currently the standard for neural seq2seq processing, with such representatives as RNN/LSTMs (Hochreiter and Schmidhuber, 1997; Sutskever et al., 2014), ConvS2S (Gehring et al., 2017), Transformer (Vaswani et al., 2017)). Formally, they are defined through a distribution  $r_\eta(x|C)$ , where  $C$  is an input (aka Context, e.g. a source sentence in Machine Translation (MT)), and  $x$  is a target sequence (e.g. a target sentence in MT). We have:

$$r_\eta(x|C) \doteq \prod_i s_\eta(x_i|x_1, \dots, x_{i-1}, C),$$

where each  $s_\eta(x_i|x_1, \dots, x_{i-1}, C)$  is a normalized conditional probability over the next symbol of the sequence, computed by a neural network (NN) with parameters  $\eta$ . The local normalization of the incremental probabilities implies the overall normalization of the distribution  $r_\eta(x|C)$ , and consequently, the possibility of directly sampling from it and evaluating the likelihood of training sequences.

### 2.2 Energy-Based Models (EBM)

EBMs are a generic class of models, characterized by an energy function  $U_\eta(x|C)$  computed by a NN parametrized by  $\eta$  (LeCun et al., 2006). Equivalently, they can be seen as directly defining a potential (an unnormalized probability distribution)  $P_\eta(x|C) = e^{-U_\eta(x|C)}$ , and indirectly the normalized distribution  $p_\eta(x|C) = 1/Z_\eta(C) P_\eta(x|C)$ , with  $Z_\eta(C) = \sum_x P_\eta(x|C)$ . A fundamental property of these models is that, for maximum likelihood training, the SGD updates can be computed through the formula:<sup>2</sup>

$$\begin{aligned} \nabla_\eta \log p_\eta(x|C) &= \nabla_\eta \log P_\eta(x|C) \\ &\quad - E_{x \sim p_\eta(\cdot|C)} \nabla_\eta \log P_\eta(x|C), \end{aligned} \quad (1)$$

which, in principle, reduces the problem of *training* with unnormalized potentials to the problem of *sampling* from them.

### 2.3 Log-Linear Models / Exponential Families

Log-Linear models (Jebara, 2013) are the conditional version of Exponential Families (Jordan, 2010). The general form of a log-linear model (for the discrete case) is as follows:

$$p_\lambda(x|C) = 1/Z_\lambda(C) \mu(x; C) e^{\langle \lambda(C), \phi(x; C) \rangle},$$

with  $Z_\lambda(C) = \sum_x \mu(x; C) e^{\langle \lambda(C), \phi(x; C) \rangle}$ . Here  $\phi(x; C)$  is a vector of predefined real features of the pair  $(x, C)$ , which is combined by scalar product with a real vector of weights  $\lambda(C)$  of the same dimension;  $\mu(x; C)$  is an arbitrary “base measure”, which is fixed. These models, which allow to introduce prior knowledge through features and have nice formal properties (see below), were mainstream in NLP before the revival of neural approaches.

## 3 Proposal: GAMs

We now define *Global Autoregressive Models* (GAMs). These are hybrid seq2seq models that exploit both local autoregressive properties as well as global properties of the full target sequence. A GAM is an unnormalized distribution  $P_\eta(x|C)$  over sequences  $x$ , parametrized by a vector  $\eta = \eta_1 \oplus \eta_2$ :

$$P_\eta(x|C) = r_{\eta_1}(x|C) \cdot e^{\langle \lambda_{\eta_2}(C), \phi(x; C) \rangle}. \quad (2)$$

<sup>2</sup>See (LeCun et al., 2006, p. 15), and [SM] for a derivation.

Here  $r_{\eta_1}(x|C)$  is an autoregressive seq2seq model for generating  $x$  from input  $C$ , parametrized by  $\eta_2$ ;  $\phi(x; C)$  is a vector of predefined real features of the pair  $(x, C)$ , which is combined by a scalar product with a real vector  $\lambda_{\eta_2}(C)$  of the same dimension, computed over the input  $C$  by a network parametrized by  $\eta_2$ . The normalized distribution associated with the GAM is  $p_\eta(x|C) = \frac{P_\eta(x|C)}{Z_\eta(C)}$ , where  $Z_\eta(C) = \sum_x P_\eta(x|C)$ .

GAMs appear promising for the following reasons:

- Features  $\phi(x; C)$  provide a simple way to draw attention of the model to potentially useful aspects that may be difficult for the AM component to discover on its own from limited data.
- GAMs are an instance of EBMs, where the potential  $P_\eta(x|C)$  is the product of the an AM potential  $r_{\eta_1}(x|C)$  with a “log-linear” potential  $e^{\langle \lambda_{\eta_2}(C), \phi(x; C) \rangle}$ . Here the gradient relative to the log-linear part takes the especially simple form:

$$\begin{aligned} \nabla_{\eta_2} \log p_\eta(x|C) &= \phi(x; C) \\ &- E_{x \sim p_\eta(\cdot|C)} \phi(x; C). \end{aligned} \quad (3)$$

- Log-linear models, on their own, while great at expressing prior knowledge, are not as good as AM models at discovering unforeseen regularities in the data. Also, they are typically problematic to train from a log-likelihood perspective, because sampling from them is often unfeasible. GAMs address the first issue through the  $r$  component, and alleviate the second issue by permitting the use of  $r$  as a powerful “proposal” (aka “surrogate”) distribution in importance sampling and related approaches, as we will see.

## 4 Experimental focus

While the motivation for GAMs ultimately lies in practical NLP applications such as those evoked earlier, in this paper we aim to understand some of their capabilities and training techniques in simple and controllable conditions. We focus on the unconditional (i.e. language modelling) case, and on synthetic data. Our setup is as follows:

- We consider an underlying process  $p_{true}$  that generates *binary sequences* according to a well-defined and flexible process. In this paper we use PFSAs (Probabilistic Finite State Automata) to impose the presence or absence of

sub-strings (“motifs”) anywhere in the generated data, exploiting the intersection properties of automata.

- Due to the dynamic programming properties of PFSAs, it is possible to compute the true entropy  $H(p_{true}) = -\sum_x p_{true}(x) \log p_{true}(x)$  of the process (see [SM]), as well as other quantities (Partition Functions, Mean sequence length); it is also possible to generate training ( $D$ ), validation ( $V$ ), and test data ( $T$ ) in arbitrary quantities.
- We employ an unconditional GAM of the simple form:

$$\begin{aligned} p_\lambda(x) &\doteq \frac{P_\lambda(x)}{Z_\lambda}, \text{ with } Z_\lambda \doteq \sum_x P_\lambda(x) \text{ and} \\ P_\lambda(x) &\doteq r(x) \cdot e^{\langle \lambda, \phi(x) \rangle}, \end{aligned} \quad (4)$$

where  $r$  is trained on  $D$  and then kept fixed, and where  $\lambda$  is then trained on top of  $r$ , also on  $D$ .

It should be noted that with  $r$  fixed in this way, this formulation exactly corresponds to the definition of an exponential family (Jordan, 2010), with  $r$  as *base measure*. In such models, we have two important properties: (i) the log-likelihood of the data is *convex* relative to the parameters  $\lambda$ , and thus a local maximum is also global; (ii) the max-likelihood value  $\lambda^*$  has the property that the model expectation  $E_{x \sim p_{\lambda^*}(\cdot)} \phi(x)$  is equal to the empirical expectation  $|D|^{-1} \sum_{x \in D} \phi(x)$  (“Moment Matching” property of exponential families).

- We are specially interested in the relative *data-efficiency* of the GAM compared to the AM  $r$ : namely the ability of the GAM to recover a lower perplexity approximation of  $p_{true}$  than  $r$ , especially in small training-set conditions.

## 5 Training procedure

### 5.1 Two-stage training

We consider a two-stage training procedure (see Fig. 1).




Figure 1: Two-stage training. At the end of the process, we compare the perplexities of  $r$  and  $\pi_\theta$  on test data:  $CE(T, r)$  vs.  $CE(T, \pi_\theta)$ .

**Training-1** This consists in training the model  $P_\lambda$  on  $D$ . This is done by first training  $r$  on  $D$  in the standard way (by cross-entropy) and then by training  $\lambda$  by SGD with the formula (adapted from (3)):

$$\nabla_\lambda \log p_\lambda(x) = \phi(x) - E_{x \sim p_\lambda(\cdot)} \phi(x). \quad (5)$$

The main difficulty then consists in computing an estimate of the model moments  $E_{x \sim p_\lambda(\cdot)} \phi(x)$ . In our experiments, we compare two Monte-Carlo approaches (Robert and Casella, 2005) for addressing this problem: (i) *Rejection Sampling* (rs), using  $r$  as the proposal distribution and (ii) *Self-Normalized Importance Sampling* (snis) (Owen, 2017; Y. Bengio and J. S. Senecal, 2008), also using  $r$  as the proposal.

Rejection sampling is performed as follows. We use  $r(x)$  as the proposal, and  $P_\lambda(x) = r(x) e^{\lambda \cdot \phi(x)}$  as the unnormalized target distribution; for any specific  $\lambda$ , because our features are bounded between 0 and 1, we can easily upper-bound the ratio  $\frac{P_\lambda(x)}{r(x)} = e^{\lambda \cdot \phi(x)}$  by a number  $\beta$ ; we then sample  $x$  from  $r$ , compute the ratio  $\rho(x) = \frac{P_\lambda(x)}{\beta r(x)} \leq 1$ , and accept  $x$  with probability  $\rho(x)$ . The accepted samples are unbiased samples from  $p_\lambda(x)$  and can be used to estimate model moments.

Snis also uses the proposal distribution  $r$ , but does not require an upper-bound, and is directly oriented towards the computation of expectations. In this case, we sample a number of points  $x_1, \dots, x_N$  from  $r$ , compute ‘‘importance ratios’’  $w(x_i) = \frac{P_\lambda(x_i)}{r(x_i)}$ , and estimate  $E_{x \sim p_\lambda(\cdot)} \phi(x)$  through  $\hat{E} = \frac{\sum_i w(x_i) \phi(x_i)}{\sum_i w(x_i)}$ . The estimate is biased for a given  $N$ , but consistent (that is, it converges to the true  $E$  for  $N \rightarrow \infty$ ).

**Training-2** While Training-1 results in a well-defined model  $P_\lambda(x)$ , which may fit the data closely in principle, we should not conclude

that  $P_\lambda(x)$  is convenient to use for inference — namely, in language modeling, efficiently sampling from its normalized version  $p_\lambda(x)$ ; as seriously, because of the partition factor  $Z_\lambda$ , it is also not obvious to *evaluate* the perplexity of  $P_\lambda(x)$  on test data. In order to do both, one approach consists in using a *distillation* technique (Hinton et al., 2015), where, during training, one expends generous time towards producing a set of samples from  $P_\lambda$ , for instance by Monte-Carlo (e.g. Rejection Sampling) techniques, and where this set (which may be arbitrarily larger than the original  $D$ ) is in turn used to train a new autoregressive model  $\pi_\theta(x)$ , which can then be used directly for sampling or for computing data likelihood. This is the approach that we use in our current experiments, again using the original  $r(x)$  as a proposal distribution.

## 5.2 Cyclical training

In the case of small  $|D|$ , the proposal distribution  $r$  is weak and as a result the distillation process, based on rejection sampling, can be slow. To address this issue, we also consider a cyclical training regime that updates the proposal distribution after distilling each batch of samples, with the intention of reducing the rejection rate. Once the process of distillation is finished, we use the aggregated samples to train the final  $\pi_\theta$ . The two-stage training procedure is a variant of the cyclical one, with a fixed proposal (see Algorithm 1 for more details).

## 6 Related Work

(Hoang et al., 2018), working in a NMT context, have a similar motivation to ours. They first train an autoregressive seq2seq model (Transformer in their case) on bilingual data, then attempt to control global properties of the generated sequences through the introduction of a priori features. They interpolate the training of the autoregressive model with training of a Moment Matching component which tries to equate the features expectations of the model with those of the data. Contrarily to our approach, they do not directly try to maximize likelihood in an integrated model.

(Andor et al., 2016) consider transition-based neural networks, and contrast local to global normalization of decision sequences, showing how the global approach avoids the *label bias* problem in such tasks as tagging or parsing. They



---

**Algorithm 1 Training**


---

```

1: function TRAIN( $D, V, T, ft, DsSize, tReg, mode$ )
2:    $r \leftarrow$  TRAINRNN( $D, V, optAdam$ ) ▷ initialize and then train RNN
3:    $P_\lambda \leftarrow$  TRAINGAM( $r, D, V, tReg, ft$ ) ▷ train  $\lambda$  for a given proposal  $r$ 
4:   if mode = 'two_stage' then ▷ Training-2: distill in one step
5:      $\tilde{D}, \tilde{V}, \_ \leftarrow$  DISTILLBATCH( $P_\lambda, DsSize$ )
6:   else if mode = 'cyclic' then ▷ Cyclic-training: distill in several steps
7:      $\tilde{D} \leftarrow \{\}; \tilde{V} \leftarrow \{\}; flag_\lambda \leftarrow$  False
8:     while  $|\tilde{D}| < DsSize$  do ▷ proceed to the distillation process
9:        $\tilde{D}_B, \tilde{V}_B, acceptRate \leftarrow$  DISTILLBATCH( $P_\lambda, bSize$ ) ▷ acceptRate - acceptance rate of  $rs$  during distillation
10:       $\tilde{D}.insert(\tilde{D}_B); \tilde{V}.insert(\tilde{V}_B)$ 
11:      if not  $flag_\lambda$  then
12:         $r \leftarrow$  SINGLEUPDATERNN( $r, \tilde{D}_B, optAdam$ ) ▷ improve proposal  $r$ 
13:         $P_\lambda \leftarrow$  TRAINGAM( $r, D, V, tReg, ft$ ) ▷ train  $\lambda$  for a given proposal  $r$ 
14:         $flag_\lambda \leftarrow$  EARLYSTOPPING $_d(acceptRate)$  ▷ check if acceptance rate has stopped improving
15:         $\tilde{D}.insert(D); \tilde{V}.insert(V)$  ▷ add true data to the distilled one
16:       $\pi_\theta \leftarrow$  TRAINRNN( $\tilde{D}, \tilde{V}, optAdam$ )
17:      return  $\pi_\theta$ 
18: function TRAINGAM( $P_\lambda, D, V, tReg, ft$ ) ▷ Training-1
19:    $\alpha_0 \leftarrow$  10 ▷ initial learning rate
20:   target_mom  $\leftarrow$  GETMOMENTS( $D, V, ft$ ) ▷ empirical moments of the given dataset
21:   while not EARLYSTOPPING( $\ell_1\_mom$ ) do ▷ check if  $\ell_1\_mom$  has stopped improving
22:     model_mom  $\leftarrow$   $[0] \times |ft|$  ▷ accumulate the model's moments
23:      $\alpha_t \leftarrow \frac{\alpha_0}{1 + \#epoch}$ 
24:     for  $b \in$  range( $\#updatesPerEpoch$ ) do
25:       mean_mom  $\leftarrow$  GETMOMENTSGAM( $P_\lambda, D, V, tReg, ft$ ) ▷ use  $rs$  or  $snis$  to estimate  $E_{x \sim P_\lambda(\cdot)} \phi(x)$ 
26:       model_mom  $\leftarrow$  (model_mom + mean_mom / ( $b - 1$ ))  $\cdot \frac{b-1}{b}$  ▷ moving average
27:        $\nabla_\lambda \leftarrow$  target_mom - mean_mom ▷ use Eq. 5 to compute gradients
28:        $\lambda \leftarrow \lambda + \alpha_t \cdot \nabla_\lambda$ 
29:        $\ell_1\_mom \leftarrow$   $\|target\_mom - model\_mom\|_1$ 
30:   return  $P_\lambda$ 

```

---

focus on inference as maximization, e.g. finding the best sequence of tags for a sequence of words, and consistent with that objective, their training procedure exploits a beam-search approximation. By contrast, our focus is on inference as sampling in a language modelling perspective, on the complementarity between auto-regressive models and log-linear models, and on the relations between training and sampling in energy-based models.

## 7 Experiments

We conduct a series of experiments on synthetic data to illustrate our approach.

### 7.1 Synthetic data

To assess the impact of GAMs, we focus on distributions  $p_{true}(x)$  that are likely to be well approximated by the AM  $r(x)$  in the presence of large data. The first class of distributions is obtained through a PFSA that filters binary strings of fixed length  $n = 30$ , 0's and 1's being equally probable (white-noise strings), through the condition that they contain a specific substring ("motif") anywhere; here the relative frequency of sequences containing the motif among all sequences varies from  $\sim 0.01$  (shorter motifs  $|m| = 10$ ) to  $\sim 0.001$  (longer motifs  $|m| = 14$ ).

We also consider mixtures of two PFSA (motif/anti-motif): the first (with mixture prob.

0.9) produces white-noise strings containing the motif and the second (with mixture prob. 0.1) strings excluding the motif.

From these processes we produce a training set  $D$ , of size  $|D|$  varying between  $5 \cdot 10^2$  and  $2 \cdot 10^4$ , a validation set  $V$  of size  $0.25 \cdot |D|$  (but never smaller than  $5 \cdot 10^2$  or bigger than  $2 \cdot 10^3$ ) and a test set  $T$  of fixed size  $5 \cdot 10^3$ .

### 7.2 Features

In a real world scenario, prior knowledge about the true process will involve, along with predictive features, a number of noisy and useless features. By training the  $\lambda$  parameters to match the empirical moments, the GAM will learn to distinguish between these types. In order to simulate this situation we consider feature vectors over our artificial data that involve both types.

With  $x$  the full string and  $m$  the fixed motif used in constructing the training data, we consider variations among the 7 binary features in the set  $F$ :

$$F = \{m, m_{+0}, m_{/2}, d_0, d_1, d_2, d_3\},$$

where  $m = 0$  iff the motif  $m$  appears in  $x$ ,  $m_{+0} = 0$  iff the motif followed by a zero ("super-motif") appears in  $x$ ,  $m_{/2} = 0$  iff an initial section of the motif ("sub-motif", roughly half the size of  $m$ ) appears in  $x$ . These three features are chosen because they have some correlation with

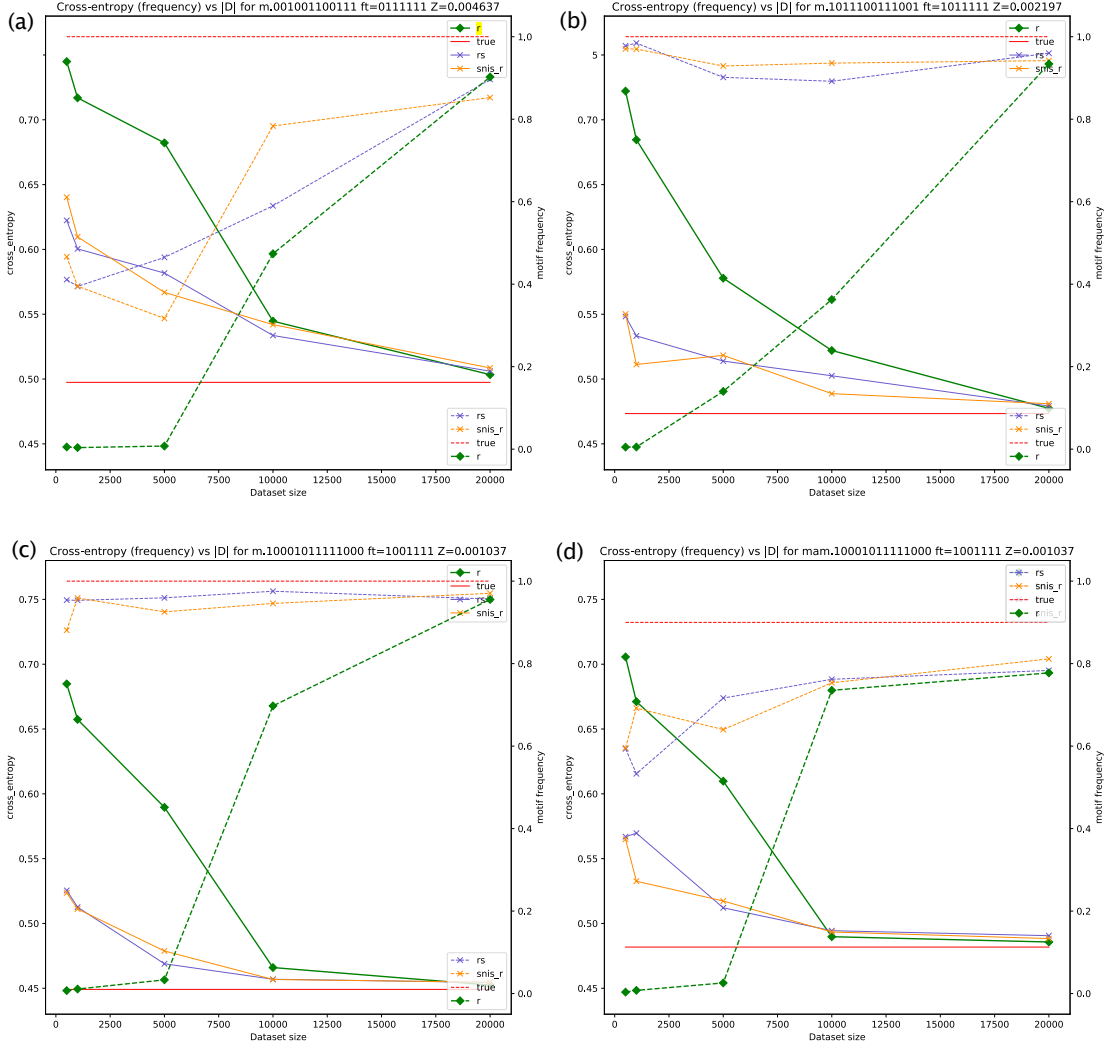


Figure 2: Cross-entropy in nats per character and frequency of sampling motif, depending on  $|D|$ . Two-stage Training. Features  $d_0, d_1, d_2, d_3$  are on for all panels ( $ft_{[4:7]} = \{1111\}$ ). Panel (a): pure  $D$ , features  $m_{+0}$  (super-motif) and  $m_{/2}$  (sub-motif) on; (b): pure  $D$ ,  $m$  (motif) and  $m_{/2}$  (sub-motif) on; (c) pure  $D$ ,  $m$  on; (d) mixture  $D, m$  on. The plain lines represent cross-entropy, the dashed lines motif frequency.

the process for generating the training data. By contrast, the four remaining features are “distractors”:  $d_0 = 0$  iff  $x$  begins with a 0,  $d_1 = 0$  (resp.  $d_2 = 0, d_3 = 0$ ) iff a certain random, but fixed, string of similar length to  $m$  (resp. of larger length, of smaller length) appears in  $x$ . We test different configurations of these features for training  $\lambda$ , and document the use/non-use of features with a bit-vector  $ft$  of length  $|F|$ , for instance  $ft = 0111111$  means that all features are exploited, apart from  $m$ .<sup>3</sup>

<sup>3</sup>In the experiments reported here, one of the provided features,  $m$ , is a detector of the motif actually present in the data generating process, an extreme form of prior knowledge used to illustrate the technique. In general, milder forms of useful prior features can be provided. A simple formal example is to consider one real-valued (non binary) feature for the length, and one for the square of the length, an experiment

## 7.3 Implementation aspects

### 7.3.1 Autoregressive models

The AMs are implemented in PyTorch<sup>4</sup> (Paszke et al., 2017) using a 2-layered LSTM (Hochreiter and Schmidhuber, 1997) with hidden-state size 200. The input is presented through one-hot encodings over the vocabulary  $V = \{0, 1, \langle \text{EOS} \rangle\}$ . These LSTMs are optimized with Adam (Kingma and Ba, 2014), with learning rate  $\alpha = 0.001$ , and

that we did recently but do not report here; by matching the data expectations of these two additional features, the model is able to represent the mean and variance of length in the data. Here the prior knowledge provided to the model just tells it to be attentive to the distribution of length, a much weaker form of prior knowledge than telling it to be attentive to a specific motif.

<sup>4</sup><https://github.com/parshakova/GAMS-for-Data-Efficient-Learning>

with early stopping (patience = 20) over a validation set.

### 7.3.2 Training: Two-Stage and Cyclical

The implementation is described in (Algorithm 1). Here we provide some additional details.

**Training-1** For training  $P_\lambda(x)$  we test two regimes in Eq. 5, namely *rs* and *snis*; in both cases, we first train  $r(x)$  on the whatever  $D$  is available, and use it as the proposal distribution. During *rs*, we compute the model’s expectation over 10 accepted samples, update the  $\lambda$ ’s according to (5), and iterate. During *snis*, we keep a buffer of the last  $5 \cdot 10^4$  samples from  $r(x)$  to compute the weighted average of the feature moments. For the training of  $\lambda$ ’s, we use a basic SGD optimization with learning rate  $\alpha(\text{\#epoch}) = \frac{\alpha_0}{1+\text{\#epoch}}$ ,  $\alpha_0 = 10$ . To assess the quality of  $P_\lambda(x)$  for early stopping during training, we use the distance between the empirical and model moments:

$$\ell_{1\text{-mom}} = \left\| \frac{1}{|D|} \sum_{d \in D} \phi(d) - E_{x \sim p_\lambda(\cdot)} \phi(x) \right\|_1. \quad (6)$$

**Training-2 and Cyclical Training** When distilling from  $P_\lambda$  in Training-2, we use a single proposal  $r$ , and systematically produce a distilled dataset of size  $D_{\text{Size}} = 2 \cdot 10^4$ , which corresponds to the highest value of  $|D|$  among those considered for training  $r$ . In Cyclical Training, the distillation process is performed in several stages, with an evolving  $r$  for improving the rejection rate.

## 8 Results

### 8.1 Cross-entropy comparison

We conduct experiments to compare the cross-entropy (measured in nats) between the initial AM  $r(x)$  relative to the test set  $T$  and the final AM  $\pi_\theta(x)$  also relative to  $T$ ; we vary the size of  $|D| \in \{0.5, 1, 5, 10, 20\} \cdot 10^3$ , the regimes (tReg) for Training-1 (*rs* or *snis*), the features employed, the rarity of the motifs. Figure 2 depicts the resulting curves at the end of the two-stage training (plain lines).

Here we show only a few experiments (a more extensive set is provided in the [SM]).

We observe that, for a small dataset size  $|D|$ , there is a big gap between the CE of  $r(x)$  and the CE of  $\pi_\theta(x)$ . As  $|D|$  increases, these cross-entropies become closer to one another, but a large gap persists for  $|D| = 5000$ .

We note that the presence of the “fully-predictive” feature  $m$  results in a  $\pi_\theta(x)$  that has CE very close to the theoretical entropy, even in low  $|D|$  regimes, where  $r$  on its own is very weak.<sup>5</sup> Thus, not only is the distilled AM much better than the initial AM, but this is an indication that  $P_\lambda$  itself (for which the cross-entropy is more difficult to compute exactly) is a good approximation of the true process.

By contrast, if the  $m$  feature is absent, then, while  $\pi_\theta$  is still better than  $r$  in low  $|D|$  regimes, it cannot reach the theoretical entropy in such regimes, because features such as  $m_{0+}$  and  $m_{/2}$  can only partially model the data. With large  $|D|$ , on the other hand,  $r$  on itself does a good job at predicting the data, and  $P_\lambda$  adds little on top of its  $r$  component.

Finally, we note that the two regimes for training  $P_\lambda(x)$ , *rs* and *snis*, result in  $\pi_\theta$ ’s with similar accuracies.

We also observe that with a good performance of  $\pi_\theta(x)$ , the moments of motif feature on the distilled dataset are close to the true ones (see [SM] Figure 4, 5, 7).

These trends are consistent across the experiments with different motifs, as can be checked in Table 3 and with the additional plots in the [SM].

### 8.2 Motif frequencies

In order to assess the *predictive* properties of obtained AMs, we also compare the frequency of motifs in strings sampled from  $r$  and from  $\pi_\theta$  ( $2 \cdot 10^3$  samples in total). From Figure 2 we see that when vary  $|D|$ , the frequency of motifs (dashed lines) is aligned with the CE performance. Namely,  $\pi_\theta$  produces a higher fraction of strings with motif than  $r$  when  $|D|$  is small ( $|D| \in \{0.5, 1, 5\} \cdot 10^3$ ).

**Detailed illustration** To provide more intuition, we provide an illustration from one experiment in Table 1.

### 8.3 Mixture $D_{mam}$ vs pure $D_m$

In our experiments, the strings in  $D_{mam}$  (motif-anti-motif) contain a motif with  $p = 0.9$ . However, if not all of the samples in  $D_{mam}$  contain the

<sup>5</sup>The CE of a model relative to the true underlying process (approximated by the test set  $T$ ) can never be below the entropy of this process, due to the KL-divergence being non-negative.

1	<i>true</i>	101 <b>10001011111000</b> 1000001001001
2	<i>r</i>	011111000010111110001110001011
3	$\pi_\theta$	111010 <b>10001011111000</b> 0111111100
4	<i>ft</i>	$[m, -, -, d_0, d_1, d_2, d_3]$
5	$\lambda$ 's	$[-10.1, -, -, -0.15, -0.06, 0.0, -0.14]$
6	mom <i>true</i>	$[0.0, -, -, 0.47, 0.99, 1.0, 0.91]$
7	mom <i>r</i>	$[0.95, -, -, 0.53, 0.99, 1.0, 0.91]$
8	mom $\pi_\theta$	$[0.0006, -, -, 0.43, 0.99, 0.99, 0.91]$
9	CEs	<i>true</i> : 0.45, <i>r</i> : 0.56, $\pi_\theta$ : 0.47
10	motif freqs	<i>true</i> : 1.0, <i>r</i> : 0.045, $\pi_\theta$ : 0.959

Table 1: Illustration. Setting is from Fig. 2, panel (c):  $n=30$ , motif = 1000101111000 (always present in  $D$ ),  $ft = 1001111$ ,  $|D| = 5000$ ,  $rs$  used for Training-1. Lines 1,2,3 show one example from *true*, *r*,  $\pi_\theta$  respectively; with training set of size 5000, *r* is only able to generate the motif a fraction of the time (0.045, see line 10), but is better able to generate some submotifs (underlined);  $\pi_\theta$  generates the motif frequently (0.959), as illustrated on line 3. With the features from *ft* (line 4), Training-1 produces a  $P_\lambda$  with first feature  $\lambda_m$  strongly negative (line 5), meaning that  $P_\lambda$  strongly penalizes the absence of the motif; the “distractor” features  $d_0, d_1, d_2, d_3$  get a weight close to 0, meaning that they have little predictive power in combination with feature  $m$ . It is visible from lines 6,7,8 that  $\pi_\theta$  is much better able to approximate the true feature expectations than *r* [features expectations (aka moments) under *r* (resp.  $\pi_\theta$ ):  $E_{x \sim r(\cdot)} \phi(x)$  (resp.  $E_{x \sim \pi_\theta(\cdot)} \phi(x)$ )] Finally (line 9), the CE of  $\pi_\theta$  relative to the test set is close to the true entropy of the process, while that of *r* is much further away.

$ D $	$m; \frac{\text{mtf\_freq}_{rs}}{\text{mtf\_freq}_{snis}}$	$m; \frac{\text{CE}(rs)}{\text{CE}(snis)}$	$m; \frac{\text{time}(rs)}{\text{time}(snis)}$	mam; $\frac{\text{mtf\_freq}_{rs}}{\text{mtf\_freq}_{snis}}$	mam; $\frac{\text{CE}(rs)}{\text{CE}(snis)}$	mam; $\frac{\text{time}(rs)}{\text{time}(snis)}$
500	0.998	0.967	2.92	0.997	1.003	4.7
1000	1.009	0.973	2.038	0.77	1.07	3.638
5000	0.995	0.967	0.756	1.12	0.99	1.365
10000	1.134	0.956	1.514	1.011	1.002	1.005
20000	1.497	0.961	0.938	0.965	1.005	0.975

Table 2: Comparison of the time for Training-1 in *rs* and *snis*; for motif 1000101111000;  $ft = 1011111$ ;  $H(p_{true}) = 0.449$  with pure  $D$  (*m*) and  $ft = 1001111$ ;  $H(p_{true}) = 0.482$  with mixture of motif-anti-motif  $D$  (*mam*).

tReg	$ D $	$m; \frac{\text{CE}(T,r)}{\text{CE}(T,\pi_\theta)}$	$m; \frac{\text{CE}(T,\pi_\theta)}{H(p_{true})}$	$m; \frac{\text{mtf\_freq}(\pi_\theta)}{\text{mtf\_freq}(r)}$	mam; $\frac{\text{CE}(T,r)}{\text{CE}(T,\pi_\theta)}$	mam; $\frac{\text{CE}(T,\pi_\theta)}{H(p_{true})}$	mam; $\frac{\text{mtf\_freq}(\pi_\theta)}{\text{mtf\_freq}(r)}$
<i>rs</i>	500	$1.24 \pm 0.07$	$1.19 \pm 0.07$	[32.0, 392.0]	$1.23 \pm 0.03$	$1.16 \pm 0.03$	[59.26, 433.33]
<i>rs</i>	1000	$1.24 \pm 0.07$	$1.16 \pm 0.07$	[23.87, 653.33]	$1.21 \pm 0.03$	$1.14 \pm 0.03$	[26.29, 233.33]
<i>rs</i>	5000	$1.18 \pm 0.08$	$1.09 \pm 0.05$	[3.59, 206.67]	$1.16 \pm 0.05$	$1.08 \pm 0.04$	[7.32, 130.0]
<i>rs</i>	10000	$1.08 \pm 0.1$	$1.04 \pm 0.02$	[0.89, 196.0]	$1.02 \pm 0.03$	$1.04 \pm 0.03$	[1.0, 4.97]
<i>rs</i>	20000	$0.99 \pm 0.01$	$1.02 \pm 0.01$	[0.81, 1.76]	$0.99 \pm 0.0$	$1.02 \pm 0.0$	[0.85, 1.04]

Table 3: Overall statistics: for  $D_m, motif \in \{10001010001, 01011101101, 001001100111, 1011100111001, 10001011111000\}$ ,  $ft \in \{1001111, 1011111, 0111111\}$  and  $D_{mam}, motif \in \{01011101101, 001001100111, 1011100111001, 100010100011, 10001011111000\}$ ,  $ft \in \{1001111\}$ .

motif, then the motif feature itself is not fully predictive. It can be seen in panel (d) of Figure 2 that the  $\pi_\theta$  achieved with  $P_\lambda$  trained on mixture  $D_{mam}$

has consistent behaviour with the results obtained on the pure  $D_m$  of panels (a,b,c).



## 8.4 Regimes in Training-1

For training GAM we consider two methods, *snis* and *rs*. As described in the previous sections, their impact on  $P_\lambda$  leads to  $\pi_\theta$ 's that have similar CE's and motif frequencies. Despite such resemblance in terms of accuracy, these two methods differ in terms of speed (see Table 2). Namely, when  $r$  is close to white noise due to small  $|D|$ , then for the rare events *rs* rejects most samples not containing the motif due to the effect of the log linear term and negative value of the component  $\lambda_m$  corresponding to the  $m$  feature, while *snis* is able to exploit all samples. Despite being faster than *rs*, *snis* remains competitive in terms of CE.

## 8.5 Cyclical vs two-stage training

We conducted a small experiment to compare the performance of cyclical training with two-stage training in terms of speed and accuracy for a fixed motif  $m$  and features  $ft$  (see [SM] Table 4, Figure 3). We observed that CEs of the obtained  $\pi_\theta$ 's were about the same for different values of  $|D|$  and Training-1 regimes. On the other hand, there was no systematic improvement in the training speed of one method over the other.

## 9 Discussion

The basic idea behind GAMs is very simple. First, we extend the representational power of the autoregressive model  $r$  by multiplying by a log-linear potential, obtaining an unnormalized model  $P_\lambda$  (Training-1). Then we try to “project” this extended representation again to an autoregressive model  $\pi_\theta$  (Training-2). Our results showed that, under favorable prior knowledge conditions, the final  $\pi_\theta$  was able to perform as well, when trained on small data, as the standard  $r$ , trained on large data. During our experiments, we noticed that training  $P_\lambda$  was actually easier than training  $\pi_\theta$  from it. Intuitively, the small number of parameters to be fitted in the log-linear model requires less work and fewer data than the training of an autoregressive component.<sup>6</sup>

<sup>6</sup>At a deeper level, there are extreme situations where the  $P_\lambda$  obtained at the end of Training-1 can perfectly represent the true process, but where no autoregressive model can actually fit  $P_\lambda$ : one way to obtain such situations consists in generating binary strings that satisfy a certain cryptographic predicate, associated with a specific feature; the importance of this feature can be easily detected through Training-1, but an autoregressive model has no chance of generalizing from distilled or true data, even in large quantities.

It is interesting to relate our study to certain aspects of Reinforcement Learning (RL).

First, consider Training-2. There, we have a “score”  $P_\lambda$  that we are trying to approximate through an autoregressive model  $\pi_\theta$ , which is basically a sequential “policy”. The main difference with RL is that we are not trying to find a policy that *maximizes* the score (which would be a bad idea for language modelling, as it would tend to concentrate the mass on a few sequences), but one that approximates  $P_\lambda$  in a *distributional* sense; our current distillation technique is only one way to approach this problem, but other techniques more in the spirit of RL are possible, a direction that we leave for future work.

Second, consider Training-1. Our approach, consisting in suggesting to the model a number of prior features, might look too easy and suspicious. But notice that in RL, one would typically directly provide to the model an externally defined *reward*, a very strong form of prior knowledge. Here, instead, we “only” indicate to the models which features it might attend to, and Training-1 then determines the “reward”  $P_\lambda$  through max-likelihood, a milder form of prior knowledge, more respectful for what the data has to say.<sup>7</sup>

## Acknowledgements

Thanks to Matthias Gallé and Ioan Calapodescu for comments on a previous version of this paper and to the anonymous reviewers for their detailed reading and feedback.

<sup>7</sup>We could say that while Training-2 addresses a question directly related to Reinforcement Learning, Training-1 addresses one related to *Inverse* Reinforcement Learning (Russell, 1998; Ng and Russell, 2000): it derives a reward from training evidence rather than imposing it externally.

## References

- Daniel Andor, Chris Alberti, David Weiss, Aliaksei Severyn, Alessandro Presta, Kuzman Ganchev, Slav Petrov, and Michael Collins. 2016. [Globally Normalized Transition-Based Neural Networks](#).
- Marc G. Bellemare, Will Dabney, and Rémi Munos. 2017. [A Distributional Perspective on Reinforcement Learning](#). *arXiv:1707.06887 [cs, stat]*. ArXiv: 1707.06887.
- Rafael C. Carrasco. 1997. Accurate computation of the relative entropy between stochastic regular grammars. *Theoretical Informatics and Applications*, 31:437–444.
- Corinna Cortes, Mehryar Mohri, Ashish Rastogi, and Michael Riley. 2008. [On the computation of the relative entropy of probabilistic automata](#). *Int. J. Found. Comput. Sci.*, 19(1):219–242.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. 2017. [Convolutional sequence to sequence learning](#). *CoRR*. Cite arxiv:1705.03122.
- Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. 2015. [Distilling the knowledge in a neural network](#). *CoRR*, abs/1503.02531.
- Cong Duy Vu Hoang, Ioan Calapodescu, and Marc Dymetman. 2018. [Moment Matching Training for Neural Machine Translation: A Preliminary Study](#).
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Tony Jebara. 2013. [Log-Linear Models, Logistic Regression and Conditional Random Fields](#).
- Michael I. Jordan. 2010. [Chapter 8 The exponential family : Basics](#).
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Yann LeCun, Sumit Chopra, Raia Hadsell, Marc’Aurelio Ranzato, and Fu Jie Huang. 2006. A Tutorial on Energy-Based Learning. *Predicting Structured Data*, pages 191–246.
- Andrew Y. Ng and Stuart J. Russell. 2000. [Algorithms for inverse reinforcement learning](#). In *Proceedings of the Seventeenth International Conference on Machine Learning, ICML ’00*, pages 663–670, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Art Owen. 2017. [Adaptive Importance Sampling \(slides\)](#).
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in PyTorch. In *NIPS Autodiff Workshop*.
- Christian P. Robert and George Casella. 2005. *Monte Carlo Statistical Methods (Springer Texts in Statistics)*. Springer-Verlag, Berlin, Heidelberg.
- Stuart Russell. 1998. [Learning agents for uncertain environments \(extended abstract\)](#). In *Proceedings of the Eleventh Annual Conference on Computational Learning Theory, COLT’ 98*, pages 101–103, New York, NY, USA. ACM.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. [Sequence to sequence learning with neural networks](#). In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 3104–3112.
- Richard S. Sutton and Andrew G. Barto. 2018. *Reinforcement Learning: An Introduction*, second edition. The MIT Press.
- 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 *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 6000–6010.
- Y. Bengio and J. S. Senecal. 2008. [Adaptive Importance Sampling to Accelerate Training of a Neural Probabilistic Language Model](#). *Ieee Transactions on Neural Networks*, 19(4):713–722.