

BRIDGING HMMs AND RNNs THROUGH ARCHITECTURAL TRANSFORMATIONS

Anonymous authors

Paper under double-blind review

ABSTRACT

A distinct commonality between HMMs and RNNs is that they both learn hidden representations for sequential data. In addition, it has been noted that the backward computation of the Baum-Welch algorithm for HMMs is a special case of the back-propagation algorithm used for neural networks (Eisner (2016)). Do these observations suggest that, despite their many apparent differences, HMMs are a special case of RNNs? In this paper, we investigate a series of architectural transformations between HMMs and RNNs, both through theoretical derivations and empirical hybridization, to answer this question. In particular, we investigate three key design factors—*independence assumptions between the hidden states and the observation, the placement of softmax, and the use of non-linearity*—in order to pin down their empirical effects. We present a comprehensive empirical study to provide insights on the interplay between expressivity and interpretability with respect to language modeling and parts-of-speech induction.

1 INTRODUCTION

Sequence is a common structure among many forms of naturally occurring data, including speech, text, video, and DNA. As such, sequence modeling has long been a core research problem across several fields of machine learning and AI. By far the most widely used approach for decades is the Hidden Markov Models of Baum & Eagon (1967); Jelinek et al. (1975), which assumes a sequence of *discrete* latent variables to generate a sequence of observed variables. When the latent variables are unobserved, unsupervised training of HMMs can be performed via the Baum-Welch algorithm (which, in turn, is based on the forward-backward algorithm), as a special case of Expectation-Maximization (EM) (Dempster et al. (1977)). Importantly, the discrete nature of the latent variables has the benefit of interpretability, as they recover contextual clustering of the output variables.

In contrast, Recurrent Neural Networks (RNNs), introduced later in the form of Jordan (1986) and Elman (1990) networks, assume *continuous* latent representations. Notably, unlike the hidden states of HMMs, there is no probabilistic interpretation of the hidden states of RNNs, regardless of their many different architectural variants (e.g. LSTMs of Hochreiter & Schmidhuber (1997), GRUs of Cho et al. (2014) and RANs of Lee et al. (2017)).

Despite their many apparent differences, both HMMs and RNNs model hidden representations for sequential data. At the heart of both models are: a state at time t , a transition function $f : h_{t-1} \rightarrow h_t$ in latent space, and an emission function $g : h_t \rightarrow x_t$. In addition, it has been noted that the backward computation in the Baum-Welch algorithm is a special case of back-propagation for neural networks (Eisner (2016)). Therefore, a natural question arises as to the fundamental relationship between HMMs and RNNs. Might HMMs be a special case of RNNs?

In this paper, we investigate a series of architectural transformations between HMMs and RNNs—both through theoretical derivations and empirical hybridization. In particular, we demonstrate that the forward marginal inference for an HMM—accumulating forward probabilities to compute the marginal emission and hidden state distributions at each time step—can be reformulated as equations for computing an RNN cell. In addition, we investigate three key design factors—*independence assumptions between the hidden states and the observation, the placement of softmax, and the use of non-linearity*—in order to pin down their empirical effects.

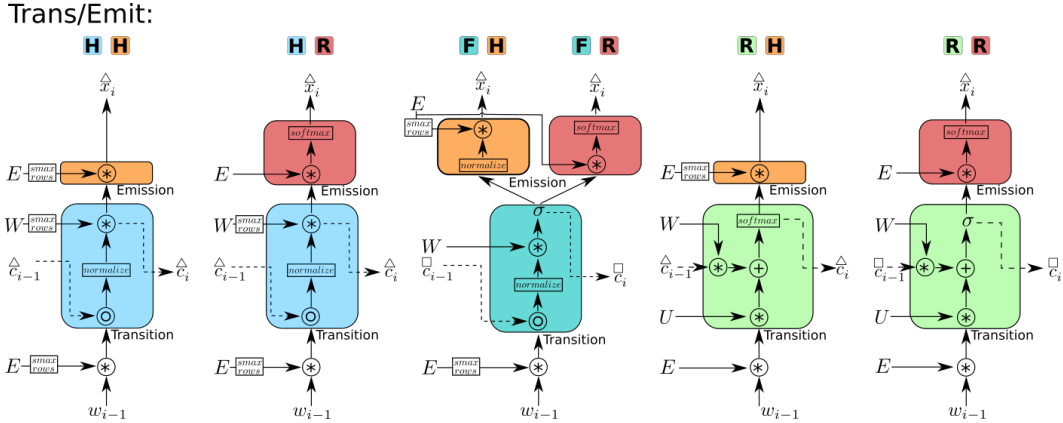


Figure 1: Above each of the models we indicate the type of transition and emission cells used. **H** for HMM, **R** for RNN/Elman and **F** is a novel Fusion defined in §3.3. It is particularly important to understanding this work to track when a vector is a distribution (resides in a simplex) versus in the unit cube (e.g. after a sigmoid non-linearity). These cases are indicated by \hat{c}_i and \square_{c_i} , respectively.

Our work is supported by several earlier works such as Wessels & Omlin (2000) and Wu et al. (2016) that have also noted the connection between RNNs and HMMs (see §7 for more detailed discussion). Our contribution is to provide the first thorough theoretical investigation into the model variants, carefully controlling for every design choices, along with comprehensive empirical analysis over the spectrum of possible hybridization between HMMs and RNNs.

We find that the key elements to better performance of the HMMs are the use of a sigmoid instead of softmax linearity in the recurrent cell, and the use of an unnormalized output distribution matrix in the emission computation. On the other hand, multiplicative integration of the previous hidden state and input embedding, and intermediate normalizations in the cell computation are less consequential. We also find that HMM outperforms other RNNs variants for unsupervised prediction of the next POS tag, demonstrating the advantages of discrete bottlenecks for increased interpretability.

The rest of the paper is structured as follows. First, we present in §2 the derivation of HMM marginal inference as a special case of RNN computation. Next in §3, we explore a gradual transformation of HMMs into RNNs. In §4, we present the reverse transformation of Elman RNNs back to HMMs. Finally, building on these continua, we provide empirical analysis in §5 and §6 to pin point the empirical effects of varying design choices over the possible hybridization between HMMs and RNNs. We discuss related work in §7 and conclude in §8.

2 FORMULATING HMMs AS RECURRENT NEURAL NETWORKS

We start by defining HMMs as sequence models, together with the forward-backward algorithm which is used for inference. Then we show that, by rewriting the forward algorithm, the computation can be viewed as updating a hidden state at each time step by feeding the previous word prediction, and then computing the next word distribution, similar to the way RNNs are structured. The resulting architecture corresponds to the first cell in Figure 1.

2.1 MODEL DEFINITION

Let $\mathbf{x}^{(1:n)} = \{x^{(1)}, \dots, x^{(n)}\}$ be a sequence of random variables, where each x is drawn from a vocabulary \mathbb{V} of size v , and an instance x is represented as an integer w or a one-hot vector $e^{(w)}$, where w corresponds to an index in \mathbb{V} .¹ We also define a corresponding sequence of hidden variables $\mathbf{h}^{(1:n)} = \{h^{(1)}, \dots, h^{(n)}\}$, where $h \in \{1, 2, \dots, m\}$. The distribution $P(\mathbf{x})$ is defined by

¹We use the recommended notation following Goodfellow et al. (2016). Sequences are notated as \mathbf{x} or w .

marginalizing over \mathbf{h} , and factorizes as follows:

$$P(\mathbf{x}) = \sum_{\mathbf{h}} P(\mathbf{x}, \mathbf{h}) = \sum_{\mathbf{h}} P(\mathbf{h}^{(1)}) p(\mathbf{x}^{(1)} | \mathbf{h}^{(1)}) \prod_{i=2}^n P(\mathbf{h}^{(i)} | \mathbf{h}^{(i-1)}) P(\mathbf{x}^{(i)} | \mathbf{h}^{(i)}) \quad (1)$$

We define the hidden state distribution, referred to as the *transition* distribution, as

$$P(\mathbf{h}^{(i)} | \mathbf{h}^{(i-1)} = l) = \text{softmax}(\mathbf{W}_{l,:} + \mathbf{b}), \quad \mathbf{W} \in \mathbb{R}^{m \times m}, \mathbf{b} \in \mathbb{R}^m \quad (2)$$

$$P(\mathbf{h}^{(1)}) = \text{softmax}\left(\sum_l \mathbf{W}_{l,:}^\top + \mathbf{b}\right), \quad (3)$$

and the *emission* (output) distribution as

$$p(\mathbf{x}^{(i)} | \mathbf{h}^{(i)} = k) = \text{softmax}(\mathbf{E}_{k,:} + \mathbf{d}), \quad \mathbf{E} \in \mathbb{R}^{m \times v}, \mathbf{d} \in \mathbb{R}^v. \quad (4)$$

2.2 INFERENCE

Inference for HMMs (marginalizing over the hidden states to compute the observed sequence probabilities) is performed with the forward-backward algorithm. The backward algorithm is equivalent to automatically differentiating the forward algorithm Eisner (2016). Therefore, while traditional HMM implementations had to implement both the forward and backward algorithm, and train the model with the EM algorithm, we only implement the forward algorithm in standard deep learning software, and perform end-to-end minibatched SGD training, efficiently parallelized on the GPU.

Let $\mathbf{w} = \{w^{(1)}, \dots, w^{(n)}\}$ be the observed sequence, and $\mathbf{w}^{(i)}$ the one-hot representation of $w^{(i)}$. The forward probabilities \mathbf{a} are defined recurrently (i.e., sequentially recursively) as

$$a_k^{(i)} = P(\mathbf{h}^{(i)} = k, \mathbf{x}^{(1:i)} = \mathbf{w}^{(1:i)}), \quad (5)$$

$$= P(\mathbf{x}^{(i)} = w^{(i)} | \mathbf{h}^{(i)} = k) \sum_{l=1}^m a_l^{(i-1)} P(\mathbf{h}^{(i)} = k | \mathbf{h}^{(i-1)} = l). \quad (6)$$

This can be rewritten by defining

$$\mathbf{c}^{(i)} = P(\mathbf{h}^{(i)} | \mathbf{x}^{(1:i-1)} = \mathbf{w}^{(1:i-1)}), \quad (7)$$

$$\mathbf{s}^{(i)} = P(\mathbf{h}^{(i)} | \mathbf{x}^{(1:i)} = \mathbf{w}^{(1:i)}), \quad (8)$$

$$x^{(i)} = P(\mathbf{x}^{(i)} = w^{(i)} | \mathbf{x}^{(1:i-1)} = \mathbf{w}^{(1:i)}), \quad (9)$$

and substituting \mathbf{a} , so that equation 6 is rewritten as (left below) or if expressed directly in terms of the parameters used to define the distributions with vectorized computations (right below):

$$c_k^{(i)} = \sum_{l=1}^m s_l^{(i-1)} P(\mathbf{h}^{(i)} = k | \mathbf{h}^{(i-1)} = l), \quad \mathbf{c}^{(i)} = \text{softmax}_{\text{rows}}(\mathbf{W})^\top \mathbf{s}^{(i-1)}, \quad (10)$$

$$\mathbf{e}^{(i)} = \text{softmax}_{\text{rows}}(\mathbf{E}) \mathbf{w}^{(i)}, \quad (11)$$

$$x^{(i)} = \sum_{k=1}^m P(\mathbf{x}^{(i)} = w^{(i)} | \mathbf{h}^{(i)} = k) c_k^{(i)}, \quad x^{(i)} = \mathbf{e}^{(i)\top} \mathbf{c}^{(i)}, \quad (12)$$

$$s_k^{(i)} = \frac{1}{x^{(i)}} P(\mathbf{x}^{(i)} = w^{(i)} | \mathbf{h}^{(i)} = k) c_k^{(i)}. \quad \mathbf{s}^{(i)} = \frac{1}{x^{(i)}} \mathbf{e}^{(i)} \circ \mathbf{c}^{(i)}. \quad (13)$$

Here $\mathbf{w}^{(i)}$ used as a one-hot vector, and the bias vectors \mathbf{b} and \mathbf{d} are omitted for clarity. Note that the computation of $\mathbf{s}^{(i)}$ can be delayed until time step $i + 1$. The computation step can therefore be

rewritten to let \mathbf{c} be the recurrent vector (equivalent logspace formulations presented on the right):²

$$\mathbf{e}^{(i-1)} = \text{softmax}_{\text{rows}}(\mathbf{E})\mathbf{w}^{(i-1)}, \quad = \text{logsoftmax}_{\text{rows}}(\mathbf{E})\mathbf{w}^{(i-1)}, \quad (14)$$

$$\mathbf{s}^{(i-1)} = \text{normalize}(\mathbf{e}^{(i-1)} \circ \mathbf{c}^{(i-1)}), \quad = \text{softmax}(\mathbf{e}^{(i-1)} + \mathbf{c}^{(i-1)}), \quad (15)$$

$$\mathbf{c}^{(i)} = \text{softmax}_{\text{rows}}(\mathbf{W})^\top \mathbf{s}^{(i-1)}, \quad = \text{log}(\text{softmax}_{\text{rows}}(\mathbf{W})^\top \mathbf{s}^{(i-1)}), \quad (16)$$

$$\mathbf{e}^{(i)} = \text{softmax}_{\text{rows}}(\mathbf{E})\mathbf{w}^{(i)}, \quad = \text{logsoftmax}_{\text{rows}}(\mathbf{E})\mathbf{w}^{(i)}, \quad (17)$$

$$\mathbf{x}^{(i)} = \mathbf{e}^{(i)\top} \mathbf{c}^{(i)}, \quad = \text{logsumexp}(\mathbf{e}^{(i)} + \mathbf{c}^{(i)}). \quad (18)$$

This can be viewed as a step of a recurrent neural network with tied input and output embeddings: Equation 14 embeds the previous prediction, equations 15 and 16, the *transition* step, updates the hidden state \mathbf{c} , corresponding to the cell of a RNN, and equations 17 and 18, the *emission* step, computes the output next word probability.

We can now compare this formulation against the definition of a Elman RNN with tied embeddings and a sigmoid non-linearity. These equations correspond to the first and last cells in Figure 1. The Elman RNN has the same parameters, except for an additional input matrix $\mathbf{U} \in \mathbb{R}^{m \times m}$.

$$\mathbf{e}^{(i-1)} = \mathbf{E}\mathbf{w}^{(i-1)}, \quad (19)$$

$$\mathbf{c}^{(i)} = \sigma(\mathbf{W}\mathbf{c}^{(i-1)} + \mathbf{U}\mathbf{e}^{(i-1)}), \quad (20)$$

$$\mathbf{x}^{(i)} = \text{softmax}(\mathbf{E}\mathbf{c}^{(i)})\mathbf{w}^{(i)}. \quad (21)$$

3 TRANSFORMING AN HMM TOWARDS AN RNN

Having established the relation between HMMs and RNNs, we propose a number of models that we hypothesize have intermediate expressiveness between HMMs and RNNs. The architecture transformations can be seen in the first 3 cells in Figure 1. We will evaluate these model variants empirically, and also investigate their interpretability.

3.1 CONDITIONING TRANSITION PROBABILITY ON PREVIOUS WORD

By relaxing the independence assumption of the HMM transition probability distribution we can increase the expressiveness of the HMM “cell” by modelling more complex interactions between the fed word and the hidden state.

Tensor-based feeding:

Following Tran et al. (2016) we define the transition distribution as

$$P(\mathbf{h}^{(i)} | \mathbf{h}^{(i-1)} = l, \mathbf{x}^{(i-1)} = w) = \text{softmax}(\mathbf{W}_{l,:} \mathbf{e}^{(i-1)} + \mathbf{B}_{l,:}), \quad (22)$$

where $\mathbf{W} \in \mathbb{R}^{m \times m \times m}$, $\mathbf{B} \in \mathbb{R}^{m \times m}$.

Addition-based feeding:

As the tensor-based methods increases the number of parameters considerably, we also propose an additive version:

$$P(\mathbf{h}^{(i)} | \mathbf{h}^{(i-1)} = l, \mathbf{x}^{(i-1)} = w) = \text{softmax}(\mathbf{W}_{l,:} + \mathbf{U}\mathbf{e}^{(i-1)} + \mathbf{b}), \quad (23)$$

where $\mathbf{W} \in \mathbb{R}^{m \times m}$, $\mathbf{U} \in \mathbb{R}^{m \times m}$, $\mathbf{b} \in \mathbb{R}^m$.

Gating-based feeding:

Finally we propose a more expressive model where interaction is controlled via a gating mechanism and the feeding step uses unnormalized embeddings (this does not violate the HMM factorization):

$$\mathbf{e}'^{(i-1)} = \mathbf{E}\mathbf{w}^{(i-1)}, \quad (24)$$

$$\mathbf{f}^i = \sigma(\mathbf{U}\mathbf{e}'^{(i-1)} + \mathbf{b}), \quad (25)$$

$$P(\mathbf{h}^{(i)} | \mathbf{h}^{(i-1)} = l, \mathbf{x}^{(i-1)} = w) = \text{softmax}(\mathbf{W}_{l,:} \circ \mathbf{f}^{(i)}), \quad (26)$$

²where $\text{normalize}(\mathbf{y}) = \frac{\mathbf{y}}{\sum_i y_i}$.

where $\mathbf{U} \in \mathbb{R}^{m \times m}$, $\mathbf{b} \in \mathbb{R}^m$, $\mathbf{W} \in \mathbb{R}^{m \times m}$.

3.2 DELAYED SOFTMAXES

Another way to make HMMs more expressive is to relax their independence assumptions through delaying when vectors are normalized to probability distributions by applying the softmax function.

Delayed transition softmax

The computation of the recurrent vector $\mathbf{c}^{(i)} = P(\mathbf{h}^{(i)} | \mathbf{x}^{(1:i-1)})$ is replaced with

$$\mathbf{c}^{(i)} = \text{softmax}(\mathbf{W}\mathbf{s}^{(i-1)}). \quad (27)$$

Both \mathbf{c} and \mathbf{s} are still valid probability distributions, but the independence assumption in the distribution over $\mathbf{h}^{(i)}$ no longer holds.

Delayed emission softmax

A further transformation is to delay the emission softmax until after multiplication with the hidden vector. This effectively replaces the HMM's emission computation with that of the RNN:

$$x^{(i)} = \text{softmax}(\mathbf{E}\mathbf{c}^{(i)})\mathbf{w}^{(i)}. \quad (28)$$

This formulation breaks the independence assumption that the output distribution is only conditioned on the hidden state assignment. Instead it can be viewed as taking the expectation over the (unnormalized) embeddings with respect to the state distribution \mathbf{c} , then softmaxed ($\mathbf{H}\mathbf{R}$ in Fig 1).

3.3 SIGMOID NON-LINEARITY

We can go further towards RNNs and replace the softmax in the transition by a sigmoid non-linearity. The sigmoid is placed in the same position as the delayed softmax. The recurrent state \mathbf{c} is no longer a distribution so the output has to be renormalized so the emission still computes a distribution:

$$\mathbf{c}^{(i)} = \text{sigmoid}(\mathbf{W}\mathbf{s}^{(i-1)}), \quad (29)$$

$$x^{(i)} = \mathbf{e}^{(i)\top} \text{normalize}(\mathbf{c}^{(i)}). \quad (30)$$

This model could also be combined with a delayed emission softmax - which we'll see makes it closer to an Elman RNN. This model is indicated as \mathbf{F} for fusion in Figure 1

4 TRANSFORMING AN RNN TOWARDS AN HMM

Analogously to making the HMM more similar to Elman RNNs, we can make Elman networks more similar to HMMs. Examples of these transformations can be seen in the last 2 cells in Figure 1.

4.1 HMM EMISSION

First, we use the Elman cell with an HMM emission. This requires the hidden state be a distribution, thus we consider two options. One is to replace the sigmoid non-linearity with the softmax function:

$$\mathbf{c}^{(i)} = \text{softmax}(\mathbf{W}\mathbf{c}^{(i-1)} + \mathbf{U}\mathbf{e}^{(i-1)}) \quad (31)$$

$$x^{(i)} = (\text{softmax}(\mathbf{E})\mathbf{w}^{(i)})^\top \mathbf{c}^{(i)}. \quad (32)$$

This model is depicted as $\mathbf{R}\mathbf{H}$ in Figure 1. The second formulation is to keep the sigmoid non-linearity, but normalize the hidden state output in the emission computation:

$$\mathbf{c}^{(i)} = \sigma(\mathbf{W}\mathbf{c}^{(i-1)} + \mathbf{U}\mathbf{e}^{(i-1)}) \quad (33)$$

$$x^{(i)} = (\text{softmax}(\mathbf{E})\mathbf{w}^{(i)})^\top \text{normalize}(\mathbf{c}^{(i)}). \quad (34)$$

4.2 MULTIPLICATIVE INTEGRATION

In the HMM cell, the integration of the previous recurrent state and the input embedding is modelled through an element-wise product instead of adding affine transformations of the two vectors. We can modify the Elman cell to do a similar multiplicative integration:³

$$\mathbf{c}^{(i)} = \sigma((\mathbf{W}\mathbf{c}^{(i-1)}) \circ (\mathbf{U}\mathbf{e}^{(i-1)})) \quad (35)$$

Or, using a single transformation matrix:

$$\mathbf{c}^{(i)} = \sigma(\mathbf{W}(\mathbf{c}^{(i-1)} \circ \mathbf{e}^{(i-1)})) \quad (36)$$

4.3 SOFTMAX NON-LINEARITY

Finally, and most extreme, we experiment with replacing the sigmoid non-linearity with a softmax:

$$\mathbf{c}^{(i)} = \text{softmax}(\mathbf{W}\mathbf{c}^{(i-1)} + \mathbf{U}\mathbf{e}^{(i-1)}) \quad (37)$$

And a more flexible variant, where the softmax is applied only to compute the emission distribution, while the sigmoid non-linearity is still applied to recurrent state:

$$\mathbf{c}^{(i)} = (\mathbf{W}\sigma(\mathbf{c}^{(i-1)}) + \mathbf{U}\mathbf{e}^{(i-1)}) \quad (38)$$

$$x^{(i)} = \text{softmax}(\mathbf{E}\text{softmax}(\mathbf{c}^{(i)})\mathbf{w}^{(i)}). \quad (39)$$

5 LANGUAGE MODELING EXPERIMENTS

Our formulations investigate a series of small architectural changes to HMMs and Elman cells. In particular, these changes raise questions about the expressivity and importance of (1) normalization within the recurrence and (2) independence assumptions during emission. In this section, we analyze the effects of these changes quantitatively via a standard language modeling benchmark.

5.1 SETUP

We follow the standard PTB language modeling setup Chelba & Jelinek (1998); Mikolov et al. (2011). We work with one-layer models to enable a direct comparison between RNNs and HMMs and a budget of 10 million parameters (typically corresponding to hidden state sizes of around 900). Models are trained with batched backpropagation through time (35 steps). Input and output embeddings are tied in all models.

Models are optimized with a grid search over optimizer parameters for two strategies: SGD⁴ and AMSProp. AMSProp is based on the optimization setup proposed by Melis et al. (2017).⁵

5.2 RESULTS

We see from the results in Table 1 (also depicted in Figure 2) that the HMM models perform significantly worse than the Elman network, as expected. Interestingly, many of the HMM variants that

³This is related to the architecture proposed by Wu et al. (2016), which shows that models with multiplicative integration can obtain competitive performance.

⁴Choose a learning rate from {5, 10, 20} and decayed by 4.0 after each epoch where validation did not improve. Batch size was 20, and embedding parameters are initialized from the uniform distribution in the range (-0.1, 0.1). These mostly follow the settings of https://github.com/pytorch/examples/tree/master/word_language_model.

⁵We use AMSGrad Reddi et al. (2018) (instead of Adam Kingma & Ba (2014)) with $\beta_1 = 0$, making it more similar to RMSProp. We pick an initial learning rate from {0.001, 0.002}, and a l_2 weight decay rate from {0, $1e-4$, $1e-5$, $1e-6$, $1e-7$ }. Batch size is 32, embedding parameters are initialized from the uniform distribution in the range (-0, 8, 0.8).

Model	dev ppl	Model	dev ppl
HMM		Elman	
-	284.59	- softmax, HMM emission	313.84
- Tensor feeding	288.15	- HMM emission	312.63
- Addition feeding	288.62	- softmax non-linearity	207.95
- Gated feeding	243.51	- normalize before emit	225.36
- Delayed transition softmax	284.59	- multiplicative (single matrix)	107.45
- Delayed emission softmax	287.00	- multiplicative	100.71
- Delayed transition and emission softmax	293.72	-	87.27
- Sigmoid non-linearity	240.91		
- Sigmoid non-linearity, delayed emission softmax	142.31	LSTM	80.61

Table 1: Language Modeling Perplexity for our baseline and transformed models.




Figure 2: This plot shows how perplexities change under our transformations, and which lead the models to converge and pass each other.

in principle have more expressivity or weaker independence assumptions do not perform better than the vanilla HMM. This includes delaying the transition or emission softmax, and most of the feeding models. The exception is the gated feeding model, which does substantially better, showing that gating is an effective way of incorporating more context into the transition matrix. Using a sigmoid non-linearity before the output of the HMM cell (instead of a softmax) does improve performance (by 44 ppl), and combining that with delaying the emission softmax gives a substantial improvement (almost another 100 ppl), making it much closer to some of the RNN variants.

We also evaluate variants of Elman RNNs: Just replacing the sigmoid non-linearity with the softmax function leads to a substantial drop in performance (120 ppl), although it still performs better than the HMM variants where the recurrent state is a distribution. Another way to investigate the effect of the softmax is to normalize the hidden state output just before applying the emission function, while keeping the sigmoid non-linearity: This performs somewhat worse than the softmax non-linearity, which indicates that it is significant whether the input to the emission function is normalized or softmaxed before multiplying with the (emission) embedding matrix. As a comparison for how much the softmax non-linearity acts as a bottleneck, a neural bigram model outperforms these approaches, obtaining 177 validation perplexity on this same setup.

Replacing the RNN emission function with that of an HMM leads to even worse performance than the HMM: Using a softmax non-linearity or a sigmoid followed by normalization does not make a significant difference. Using multiplicative integration leads to only a small drop in performance compared to a vanilla Elman RNN, and doing so with a single transformation matrix (making it comparable to what an RNN is doing) leads to only a small further drop. In contrast, preliminary experiments showed that the second transformation matrix is crucial in the performance of the vanilla Elman network.

In our experimental setup an LSTM performs only slightly better than the Elman network (80 vs 87 perplexity). While more extensive hyperparameter tuning Melis et al. (2017) or more sophisticated optimization and regularization techniques Merity et al. (2017) would likely improve performance, that is not the goal of this evaluation.

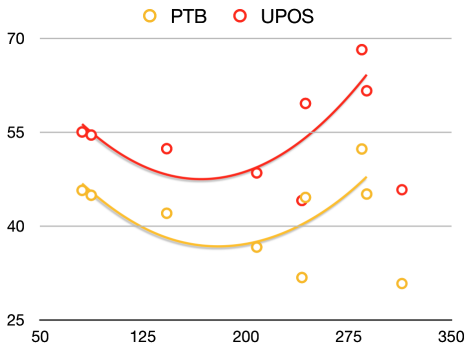


Figure 3: Tagging accuracies (right) are plotted against perplexities from Table 1. We see a somewhat quadratic relationship.

Model	PTB	UPOS
HMM		
-	52.36	68.23
- Tensor feeding	45.16	61.66
- Gated feeding	44.62	59.64
- Sigmoid non-linearity	31.82	44.13
- Sigmoid non-linearity with delayed emission softmax	42.09	52.41
Elman		
- softmax, HMM emission	30.86	45.85
- softmax non-linearity	36.68	48.54
-	44.97	54.59
LSTM		
-	45.75	55.08

Table 2: Tagging accuracies for several representative models. Accuracy is calculated by converting $p(w)$ to $p(t)$ according to WSJ tag distributions.

6 SYNTACTIC EVALUATION

A strength of HMM bottlenecks is forcing the model to produce an interpretable hidden representation. A classic example of this property is part-of-speech tag induction. It is therefore natural to ask whether changes in the architecture of our models correlate with their ability to discover syntactic properties. We evaluate this by analyzing the models implicitly predicted tag distribution at each time step. Specifically, while no model is likely to predict the correct next word, we assume the HMMs errors will preserve basic tag-tag patterns of the language, and that this may not be true for RNNs. We test this by computing the accuracy of predicting the tag of the word in the sequence out of the next word distribution. None of the models were trained to perform this task.

First, we compute a tag distribution $p(t|w)$ for every word in the training portion of the Penn Treebank. Next, we multiply this value by the model’s $p(w) = \hat{x}_i$, and sum across the vocabulary. This provides us the model’s distribution over tags at the given time $p(t)_i$. We compare the most likely marginal tag against the ground truth to compute a tagging accuracy. This evaluation rewards models which place their emission probability mass predominantly on words of the correct part-of-speech. We compute this metric across both the full PTB tagset and the universal tags of Petrov et al. (2012).

The HMM allows for Viterbi decoding which allows us to compute $p(t|\max_{\text{dim}}(c_i))$. The more distributed the models’ representations are, the more the tag distribution given the max dimension will differ from the complete marginal. For HMMs with distributional hidden states the maximum dimension provided the best performance. In contrast, Elman models perform best when conditioned on the full hidden state. Results are shown in Table 2 and plotted against perplexity in Figure 3.⁶

7 RELATED WORK

Recently, a number of recent papers have identified variants of gated RNNs which are simpler than LSTMs but perform competitively or satisfy properties that LSTMs lack. Foerster et al. (2017) proposed RNNs without recurrent non-linearities to improve interpretability. Balduzzi & Ghifary (2016) proposed gated RNN variants with type constraints. Peng et al. (2018) identified a class of RNNs called rational recurrences, in which the hidden states can be computed by WFSA.

Another strand of recent work proposed neural models that learn discrete, interpretable structure: Yang et al. (2017) introduced a mixture of softmax model where the output distribution is conditioned on discrete latent variable. Shen et al. (2017) proposed a language model that jointly learns unsupervised syntactic (tree) structure, while Tran et al. (2016) used neural hidden Markov models for Part-of-Speech induction. Wiseman et al. (2018) and Wang et al. (2017) proposed models for segmental structure over sequences, while neural transduction models with discrete latent alignments have also been proposed Yu et al. (2016).

⁶For simplicity, our trend-lines ignore the Elman + softmax, HMM Emission as an outlier (bottom right).

8 CONCLUSION

In this work, we presented a theoretical and empirical investigation into the model variants over the spectrum of possible hybridization between HMMs and RNNs. By carefully controlling for every design choices, we provide new insights into several factors including independence assumptions, the placement of softmax, and the use of nonlinearity and how these choices influence the interplay between expressiveness and interpretability. Comprehensive empirical results demonstrate that the key elements to better performance of the HMM are the use of a sigmoid instead of softmax linearity in the recurrent cell, and the use of an unnormalized output distribution matrix in the emission computation. Multiplicative integration of the previous hidden state and input embedding, and intermediate normalizations in the cell computation are less consequential. We also find that HMM outperforms other RNNs variants in a next POS tag prediction task, which demonstrates the advantages of models with discrete bottlenecks in increased interpretability.

REFERENCES

- David Balduzzi and Muhammad Ghifary. Strongly-typed recurrent neural networks. *CoRR*, abs/1602.02218, 2016. URL <http://arxiv.org/abs/1602.02218>.
- Leonard E. Baum and J. A. Eagon. An inequality with applications to statistical estimation for probabilistic functions of markov processes and to a model for ecology. *Bulletin of the American Mathematical Society*, 73(3):360 – 363, 1967.
- Ciprian Chelba and Frederick Jelinek. Exploiting syntactic structure for language modeling. In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics, Volume 1*, pp. 225–231, Montreal, Quebec, Canada, August 1998. Association for Computational Linguistics. doi: 10.3115/980845.980882. URL <http://www.aclweb.org/anthology/P98-1035>.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1724–1734, Doha, Qatar, October 2014. Association for Computational Linguistics.
- A Dempster, N Laird, and D Rubin. Maximum likelihood from incomplete data via the em algorithm. 01 1977.
- Jason Eisner. Inside-outside and forward-backward algorithms are just backprop (tutorial paper). In *Proceedings of the Workshop on Structured Prediction for NLP*, pp. 1–17, Austin, TX, November 2016. Association for Computational Linguistics. URL <http://aclweb.org/anthology/W16-5901>.
- J Elman. Finding structure in time. 14(2):179–211, 06 1990.
- Jakob N Foerster, Justin Gilmer, Jascha Sohl-Dickstein, Jan Chorowski, and David Sussillo. Input switched affine networks: An rnn architecture designed for interpretability. In *International Conference on Machine Learning*, pp. 1136–1145, 2017.
- Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. *Deep learning*, volume 1. MIT Press, 2016.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8): 1735–1780, 1997.
- Frederick Jelinek, Lalit R. Bahl, and Robert L Mercer. Design of a linguistic statistical decoder for the recognition of continuous speech. *IEEE Transactions on Information Theory*, 21(3):250–256, May 1975.
- Michael I Jordan. Serial order: A parallel distributed processing approach. Technical report, University of California, San Diego, 1986.

- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kenton Lee, Omer Levy, and Luke S Zettlemoyer. Recurrent additive networks. *arXiv preprint arXiv:1705.07393*, 05 2017.
- Gábor Melis, Chris Dyer, and Phil Blunsom. On the state of the art of evaluation in neural language models. *arXiv preprint arXiv:1707.05589*, 2017.
- Stephen Merity, Nitish Shirish Keskar, and Richard Socher. Regularizing and optimizing LSTM language models. *CoRR*, abs/1708.02182, 2017. URL <http://arxiv.org/abs/1708.02182>.
- Tomáš Mikolov, Anoop Deoras, Stefan Kombrink, Lukáš Burget, and Jan Černocký. Empirical evaluation and combination of advanced language modeling techniques. In *Twelfth Annual Conference of the International Speech Communication Association*, 2011.
- Hao Peng, Roy Schwartz, Sam Thomson, and Noah A. Smith. Rational recurrences. In *Proc. of EMNLP*, 2018.
- Slav Petrov, Dipanjan Das, and Ryan McDonald. A universal part-of-speech tagset. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC-2012)*, pp. 2089–2096, Istanbul, Turkey, 05 2012.
- Sashank J Reddi, Satyen Kale, and Sanjiv Kumar. On the convergence of adam and beyond. 2018.
- Yikang Shen, Zhouhan Lin, Chin-Wei Huang, and Aaron C. Courville. Neural language modeling by jointly learning syntax and lexicon. *CoRR*, abs/1711.02013, 2017. URL <http://arxiv.org/abs/1711.02013>.
- Ke M. Tran, Yonatan Bisk, Ashish Vaswani, Daniel Marcu, and Kevin Knight. Unsupervised neural hidden markov models. In *Proceedings of the Workshop on Structured Prediction for NLP*, pp. 63–71, Austin, TX, November 2016. Association for Computational Linguistics. URL <http://aclweb.org/anthology/W16-5907>.
- Chong Wang, Yining Wang, Po-Sen Huang, Abdelrahman Mohamed, Dengyong Zhou, and Li Deng. Sequence modeling via segmentations. *arXiv preprint arXiv:1702.07463*, 2017.
- T Wessels and Christian W Omlin. Refining hidden markov models with recurrent neural networks. In *Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks*, volume 2, pp. 271–276. IEEE, 2000.
- Sam Wiseman, Stuart M Schieber, and Alexander M Rush. Learning neural templates for text generation. *arXiv preprint arXiv:1808.10122*, 08 2018.
- Yuhuai Wu, Saizheng Zhang, Ying Zhang, Yoshua Bengio, and Ruslan R Salakhutdinov. On multiplicative integration with recurrent neural networks. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 29*, pp. 2856–2864. Curran Associates, Inc., 2016. URL <http://papers.nips.cc/paper/6215-on-multiplicative-integration-with-recurrent-neural-networks.pdf>.
- Zhilin Yang, Zihang Dai, Ruslan Salakhutdinov, and William W Cohen. Breaking the softmax bottleneck: A high-rank rnn language model. *arXiv preprint arXiv:1711.03953*, 2017.
- Lei Yu, Jan Buys, and Phil Blunsom. Online segment to segment neural transduction. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 1307–1316, Austin, Texas, November 2016. Association for Computational Linguistics. URL <https://aclweb.org/anthology/D16-1138>.