

# Recurrent Point Review Models

Kostadin Cvejovski<sup>\*‡</sup>, Ramsés J. Sánchez<sup>\*†</sup>, Bogdan Georgiev<sup>‡</sup>, Christian Bauckhage<sup>‡</sup> and César Ojeda<sup>§</sup>

<sup>\*</sup>Competence Center Machine Learning Rhine-Ruhr

<sup>†</sup>B-IT, University of Bonn, Bonn, Germany

<sup>‡</sup>Fraunhofer Center for Machine Learning and Fraunhofer IAIS, 53757 Sankt Augustin, Germany

<sup>§</sup>Berlin Center for Machine Learning and TU Berlin, 10587 Berlin, Germany

{kostadin.cvejovski, bogdan.georgiev, christian.bauckhage}@iais.fraunhofer.de,  
ojeda.marin@tu-berlin.de, sanchez@bit.uni-bonn.de

**Abstract**—Deep neural network models represent the state-of-the-art methodologies for natural language processing. Here we build on top of these methodologies to incorporate temporal information and model how review data changes with time. Specifically, we use the dynamic representations of recurrent point process models, which encode the history of how business or service reviews are received in time, to generate instantaneous language models with improved prediction capabilities. Simultaneously, our methodologies enhance the predictive power of our point process models by incorporating summarized review content representations. We provide recurrent network and temporal convolution solutions for modeling the review content. We deploy our methodologies in the context of recommender systems, effectively characterizing the change in preference and taste of users as time evolves. Source code is available at [1].

**Index Terms**—dynamic language models, marked point processes, recommender systems

## I. INTRODUCTION

Dynamic models of text aim at characterizing temporal changes in patterns of document generation. Most successful dynamic language models are Bayesian in nature, and lag behind state-of-the-art deep language models in terms of expressibility. A natural space to study some of the temporal aspects of language is that of the large review datasets found in e-commerce sites. The availability of millions of reviewed items, such as business or services, books or movies, whose reviews have been recorded in time scales of years, opens up the possibility to develop deep scalable models that can predict the change in taste and preference of users as time evolves. Originally, the interaction of users in these e-commerce sites were studied in the context of collaborative filtering, where the goal was to predict user ratings, based on user interaction metrics. Here we aim to look directly at the content of reviews as time evolves.

Customer reviews provide a rich and natural source of unstructured data which can be leverage to improve recommender system performance [2]. Indeed, reviews are effectively a form of recommendation. Recently, a variety of deep learning solutions for recommendation have profit from their ability to extract latent representations from review data, encoding rich information related to both users and items. Time represents yet another dimension of context, as user preference and item availability change with time – and indeed, causal and temporal relations have been known to improve

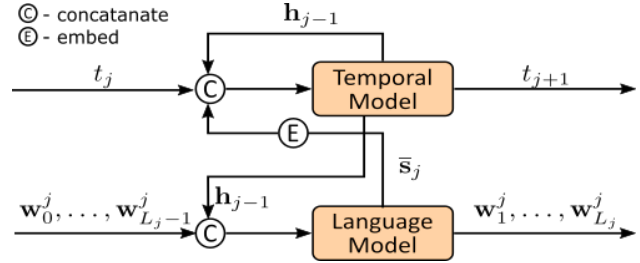


Fig. 1: Recurrent Point Review (RPR) Model. We treat the pair  $(t_j, (\mathbf{w}_0^j, \dots, \mathbf{w}_{L_j-1}^j))$  as input to our model. The hidden state of the temporal model  $\mathbf{h}_j$  captures the non-linear dependency between timing and text content from past reviews. On the other hand, the mean hidden state of the language model  $\bar{\mathbf{s}}_j$  holds information about the text content of the current review and is used for updating the global representation  $\mathbf{h}_j$ . See Table I for reference to our notation.

the performance of recommender systems [3] [4]. Despite this fact, recent natural language processing (NLP) methodologies for rating and reviews [5] lag behind at incorporating temporal structure in their language representations. In the present work we exploit recurrent neural network (RNN) models for point processes, and feed them neural representations of text, to characterize customer reviews. Our goal is to capture the changes in user taste and item importance during time, and to exploit those changes to better predict when are new reviews arriving, and what do they actually say. We summarize our contributions as follows:

- A text-augmented marked temporal point process model for review arrival times with better predictive performance,
- A hierarchical language model that leverages the global dynamic representations of the review arrival point process, effectively defining instantaneous language models with improved modelling performance.

We present the related work in Section II and introduce our model in Section III. The baseline models used for comparison in this paper are presented in Section IV. The experimental setup and results are presented in Section V. Finally, in Section VI we conclude and discuss future work.

## II. RELATED WORK

The dynamics of language is of fundamental importance in social sciences as a proxy for cultural evolution [6]. Complex system methods seek to understand the emergence of the encoding capabilities of language [7], and evolutionary approaches – following the Bayesian tradition from Phylogenetics, study the competition between grammar and syntax in the context of historical linguistics [8]. Closer to our line of work, research of online communities point to temporal linguistic changes as means to enforce community norms [9]. Our methodologies aim at studying similar systems in the e-commerce review context, wherein linguistic change is relevant in time scales of months and years.

The work on language dynamics from the machine learning community has mainly focused on the dynamics of word embeddings and topics. On the one hand, different embeddings, as e.g. word2vec [10], are trained in slices of temporal data and alignments methods are performed a posteriori [11], [12]. The probabilistic framing of word embeddings has, in contrast, allowed for a stochastic-process point of view of embedding evolution [13]. On the other hand, within the dynamic topic modelling approach, the parameters of models like Latent Dirichlet Allocation are defined to follow known stochastic processes in time [14], [15]. Likewise, self-exciting point processes in time have allowed for clustering of document streams [16], [17]. Lastly, while writing this paper we found [18], in which a RNN language model is conditioned on a global dynamic latent variable. In contrast to this work, our dynamic representations explicitly encode both timing and content of past reviews, and can capture non-Markovian dynamics.

Finally, within the recommender system realm deep neural networks models of review data for rating predictions use embedding representations, as well as convolutional neural networks [19]. They also provide characterization of review usefulness [20], use reviews for product description [21], and provide better representations for rating prediction [22]. The need to interact with the customer has also led to question answering solutions [23], [24]. Different from these works, we focus on the temporal aspects of review content.

## III. RECURRENT POINT REVIEW MODEL (RPR)

Rather intuitively, the generation process underlying the content of the reviews received by a business or service is intimately related to how fast and when these reviews are created (or received) in time. It is not hard to imagine, for example, reviews with similar content clustering and forming different kind of patterns in time. Based on this intuition, we develop a model that explicitly uses the text content of reviews to better predict the arrival times of new ones. Simultaneously, the model leverages the history of how reviews have arrived in time to yield *instantaneous* language models, i.e. language models which are informed by time. As a result we improve the text prediction capabilities of standard language models by learning representations which characterize the time evolution

Symbol	Description
$V$	vocabulary size
$H$	dimension of hidden state for temporal dynamics
$S$	dimension of hidden state for language model
$W$	word embedding dimension
$a$	item index (business)
$N_a$	number of reviews for business $a$
$M$	total number of businesses
$\mathbf{r}_j^a$	review of business $a$ at time step $j$
$\mathbf{x}_j^a$	text of review $\mathbf{r}_j^a$
$L_j$	number of words in the $j$ th review
$t_j^a$	timestamp of review $\mathbf{r}_j^a$
$\mathbf{w}_i^j$	$i$ th word embedding in the $j$ th review
$\mathbf{X}_j^a$	bag of words embedding for review $\mathbf{r}_j^a$
$\mathbf{h}_j^a$	hidden state of the temporal model for review $\mathbf{r}_j^a$
$\bar{\mathbf{s}}_j$	summary content representation for the $j$ th review

TABLE I: Summary of notation used throughout the paper.

of the review content describing the business or service in question.

Consider an item  $a$  (e.g. a business, service or movie) and assume that, since its opening to the public, it has received a collection of  $N_a$  reviews  $\mathbf{r}_j^a = \{(\mathbf{x}_j^a, t_j^a)\}_{j=1}^{N_a}$ , where  $t_j^a$  labels the creation time of review  $\mathbf{r}_j^a$  and  $\mathbf{x}_j^a = (\mathbf{w}_1^{a,j}, \dots, \mathbf{w}_{L_{a,j}}^{a,j})$  corresponds to its text<sup>1</sup>. Such a collection of reviews effectively defines a point process in time. Our main idea is to model these point processes as RNNs *in continuous time*, and feed them representations summarizing the content of past reviews. We then use the point process’ hidden representations, which encode the nonlinear relations between text and timing of past reviews, to predict both (i) how the reviews’ content of a given item changes with time, and (ii) when are new reviews going to arrive. The model thus consists of two interacting components: a neural point process model which leverages the information encoded in the review content, and a dynamic neural language model which uses the point process history. In what follows we dwell into the details of these two building blocks. Figure 1 summarizes the Recurrent Point Review model.

### A. Dynamic language models for review text content

Auto-regressive neural language models approximate the joint probability distribution over sequences of words with a product over conditional probabilities such that

$$p(\mathbf{x}) = \prod_{i=0}^L p_{\theta}(\mathbf{w}_i | \mathbf{w}_{<i}), \quad (1)$$

where  $\mathbf{x} = (\mathbf{w}_1, \dots, \mathbf{w}_L)$  labels the sequence of words in question, and  $p_{\theta}$  is a discrete probability distribution, parametrized by a neural network with parameter set  $\theta$ , which yields the probability of observing  $\mathbf{w}_i$  given the previously generated words  $\mathbf{w}_{<i}$  [25]. In order to capture how the review content of a given item changes with time, we assume the conditional probabilities in Eq. (1) additionally depend on a vector representation  $\mathbf{h}_j$ , which encodes the nonlinear relations between creation times and content of past reviews. Let us assume for the moment that  $\mathbf{h}_{j-1}$  is given. We shall explain how to

<sup>1</sup>In what follows we shall drop the index  $a$  over items.

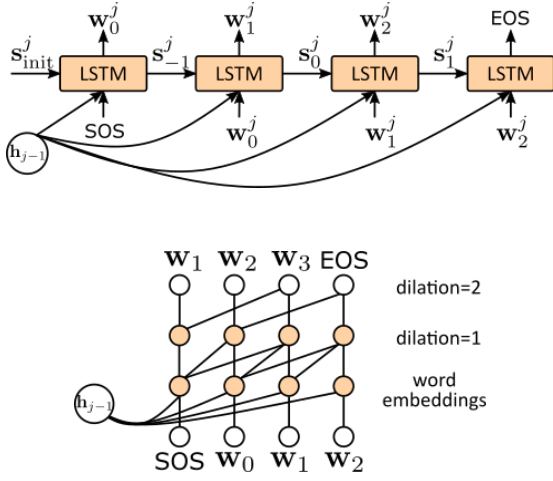


Fig. 2: Neural networks parametrizing our dynamic language models. *Upper figure*: LSTM network for which the hidden state  $s_i^j$  is updated recursively. *Lower Figure*: Dilated TCN, whose effective receptive size grows exponentially with each layer. The temporal vector representation  $\mathbf{h}_{j-1}$  is concatenated with each word embedding in both models. SOS and EOS label the star-of-sequence and end-of-sequence tokens. See Table I for reference to our notation.

calculate it in Section III-B below. Given  $\mathbf{h}_{j-1}$ , we can rewrite Eq. (1) for the  $j$ th review of an item thus

$$p(\mathbf{x}_j | \mathbf{h}_{j-1}) = \prod_{i=0}^{L_j} p_{\theta}(\mathbf{w}_i^j | \mathbf{w}_{<i}^j, \mathbf{h}_{j-1}). \quad (2)$$

In practice we define  $p_{\theta}$  as a categorical distribution over a vocabulary of size  $V$  whose class probabilities  $\pi_i^j \in \mathbb{R}^V$  are given by

$$\pi_i^j = \text{softmax}(\mathbf{W} \mathbf{s}_i^j), \quad \mathbf{s}_i^j = g_{\theta}(\mathbf{w}_{<i}^j, \mathbf{h}_{j-1}), \quad (3)$$

where the index “ $i$ ” runs from one to  $L_j$ ,  $\mathbf{W} \in \mathbb{R}^{V \times S}$  is a learnable weight matrix, and  $\mathbf{s}_i^j \in \mathbb{R}^S$  is the hidden representation of the  $i$ th word in review  $j$ . We shall model the function  $g_{\theta}$  with two standard neural network architectures, namely a Long short-term memory (LSTM) neural network and a Temporal Convolutional Network (TCN). We briefly revisit these two architectures below.

1) *Long Short Term Memory Network*: One of the most common variants of recurrent neural networks, aimed at solving the vanishing gradient problem, is the long short-term memory network [26]. Let us define  $\hat{\mathbf{w}}_i^j = \text{concat}(\mathbf{w}_i^j, \mathbf{h}_{j-1})$ . The LSTM network recursively processes each element  $\hat{\mathbf{w}}_i^j$  in review  $j$  while updating its hidden state  $\mathbf{s}^j$  as follows

$$\begin{aligned} \mathbf{i}_t &= \sigma(\mathbf{W}_i^1 \hat{\mathbf{w}}_t + \mathbf{W}_i^2 \mathbf{s}_{t-1} + \mathbf{b}_i), \\ \mathbf{o}_t &= \sigma(\mathbf{W}_o^1 \hat{\mathbf{w}}_t + \mathbf{W}_o^2 \mathbf{s}_{t-1} + \mathbf{b}_o), \\ \mathbf{f}_t &= \sigma(\mathbf{W}_f^1 \hat{\mathbf{w}}_t + \mathbf{W}_f^2 \mathbf{s}_{t-1} + \mathbf{b}_f), \\ \mathbf{c}_t &= \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \tanh(\mathbf{W}_c^1 \hat{\mathbf{w}}_t + \mathbf{W}_c^2 \mathbf{s}_{t-1} + \mathbf{b}_c), \\ \mathbf{s}_t &= \mathbf{o}_t \odot \tanh(\mathbf{c}_t). \end{aligned} \quad (4)$$

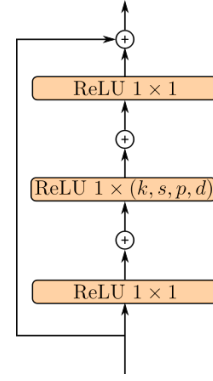


Fig. 3: Temporal Convolutional Network: Residual Block

Here  $\mathbf{i}_t$ ,  $\mathbf{o}_t$ ,  $\mathbf{f}_t$ ,  $\mathbf{c}_t$ , and  $\mathbf{s}_t$  are the input, output, forget, memory and hidden states of the LSTM, respectively; “ $t$ ” runs from one to  $L$ , the number of words in the  $j$ th review (note we have omitted  $j$  above);  $\sigma$  labels the ReLU nonlinearity and  $\odot$  labels element-wise multiplication. The upper sub-figure in Fig. 2 depicts the LSTM dynamic language model.

2) *Temporal Convolutional Network*: Convolutional neural networks have shown to be the workhorse for image processing tasks [27], [28], and NLP applications have profited from their rich representations of text [29]. To extend these model architectures for sequential modelling tasks, one must enforce two conditions [30]: (i) no information “leakage” from the future to the past, and (ii) the architecture should take a sequence of any length and map it to an output sequence of the same length. Both these conditions can be achieved with the correct padding, and we refer to [30] for reference.

*Dilation*: One would like the network to look very far into the past. This characteristic is accomplished by what is known as diluted convolutions [31]. Essentially, dilation boils down to skipping  $d - 1$  inputs at each step. More formally, for a filter  $f : \{0, \dots, k - 1\} \rightarrow \mathbb{R}$ , the dilated convolution operation  $g$  on the  $i$ th word of review  $j$  reads

$$g_{\theta}(\mathbf{w}_i^j) = \sum_{l=0}^{k-1} f_{\theta}(l) \cdot \mathbf{x}_{i-dl}^j, \quad (5)$$

where  $d$  is the dilation factor and  $k$  is the filter size. With dilation, the effective receptive size can grow exponentially<sup>2</sup>.

*Residual Block*: Our TCN model also uses residual connections [32] to speed up convergence. Specifically, we use a residual block similar to that of [33], [34]. The block is sketched in Fig. 3.

The sequence of word representation  $\mathbf{s}_i^j$  of Eq. (3) is thus given by

$$\mathbf{s}_i^j = g_{\theta_l}(\text{ReLU}(g_{\theta_{l-1}}(\dots \text{ReLU}(g_{\theta_1}(\hat{\mathbf{w}}_i^j)) \dots))), \quad (6)$$

where  $l$  is the total number of layers in the TCN,  $\hat{\mathbf{w}}_i^j = \text{concat}(\mathbf{w}_i^j, \mathbf{h}_{j-1})$  and  $g_{\theta}$  is defined in Eq. (5) above. The lower sub-figure in Fig. 2 shows the TCN dynamic language model.

<sup>2</sup>In practice, we double the dilation step at each layer, i.e.  $d_{l+1} = 2d_l$ .

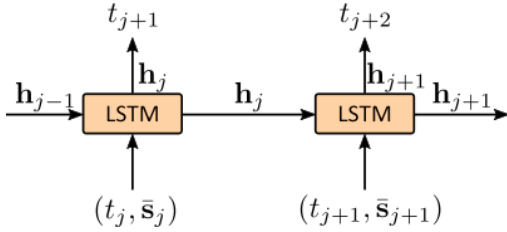


Fig. 4: Text-augmented Recurrent Point Process (RPP) model. We treat the pair  $(t_j, \bar{s}_j)$  as input to the temporal model. The hidden state  $\mathbf{h}_j$  captures the non-linear dependency between timing and content of the review arrival process. The summary representation of the language model  $\bar{s}_j$  holds information about the text content of the current review and is used for updating the global representation  $\mathbf{h}_j$ . See Table I for reference to our notation.

### B. Temporal model for review creation: Recurrent Point Process (RPP) model

The collection of reviews received by a business or service effectively defines a point process in time. Let us consider a point process with compact support  $\mathcal{S} \subset \mathbb{R}$ . Formally, we write the likelihood of a new arrival (i.e. a new review)  $\mathbf{r}_{j+1}$  as an inhomogeneous Poisson process between reviews, conditioned on the history  $\mathcal{H}_j \equiv \{\mathbf{r}_1, \dots, \mathbf{r}_j\}$ <sup>3</sup> [35]. For the one-dimensional processes we are concerned here, the conditional likelihood function reads

$$f^*(t) = \lambda^*(t) \exp \left\{ \int_{t_j}^t \lambda^*(t') dt' \right\}, \quad (7)$$

where  $\lambda^*$  is (locally) integrable and is known as the intensity function of the point process. Following [36], [37], we define the functional dependence of the intensity function to be given by a RNN with hidden state  $\mathbf{h}_j \in \mathbb{R}^H$ , where an exponential function guarantees that the intensity is *non-negative*

$$\lambda^*(t) = \exp \{ \mathbf{v}^t \cdot \mathbf{h}_j + w^t (t - t_j) + b^t \}. \quad (8)$$

Here the vector  $\mathbf{v}^t \in \mathbb{R}^d$  and the scalars  $w^t$  and  $b^t$  are trainable variables.

Now, we can modify the original RPP model of [37] and augment  $\mathbf{h}_j$  to explicitly encode the nonlinear relations between content and timing of past reviews. Consider the summary representation  $\bar{s}_j$  for review  $\mathbf{r}_j$  defined as the average

$$\bar{s}_j = \frac{1}{L_j} \sum_i^{L_j} \mathbf{s}_i^j, \quad (9)$$

where  $\mathbf{s}_i^j$  is given by Eq. (3). The update equation for the hidden representation  $\mathbf{h}_j$  can then be written as

$$\mathbf{h}_j = g_\theta(t_j, \bar{s}_j, \mathbf{h}_{j-1}), \quad (10)$$

where  $t_j$  labels the creation time of review  $\mathbf{r}_j$  and  $\theta$  denotes the network's parameters. In practice,  $g_\theta$  is implemented via

<sup>3</sup>a.k.a. filtration.

an LSTM network with update equations similar to those in Eq. (4). To summarize, we use the summary representations  $\bar{s}_j$  of the review content as marks in the recurrent marked temporal point process [38]. Figure 4 sketches our text-augmented RPP model.

Inserting Eq. (8) into Eq. (7) and integrating over time immediately yields the likelihood  $f^*$  as a function of  $\mathbf{h}_j$ , that is

$$f^*(t|\mathbf{h}_j) = \exp \left\{ \mathbf{v}^t \cdot \mathbf{h}_j + w^t (t - t_j) + b^t + \frac{1}{w^t} \exp \{ \mathbf{v}^t \cdot \mathbf{h}_j + b^t \} - \frac{1}{w^t} \exp \{ \mathbf{v}^t \cdot \mathbf{h}_j + w^t (t - t_j) + b^t \} \right\}. \quad (11)$$

*Prediction and sampling.* In order to use the RPP model for both prediction and sampling we require  $p(T|H_j)$ , the probability that the next review arrives at time  $T$  given the previous history until the arrival of  $\mathbf{r}_j$ . In the Appendix we calculate the latter and give expressions for both the average time of the next arrival and the inverse of the cumulative function of  $p(T|H_j)$ , which is required to sample the next review arrival time via inverse transform sampling.

### C. Recurrent point review model log-likelihood

We shall train our model using maximum likelihood. The complete log-likelihood of the RPP model can be written as

$$\mathcal{L} = \sum_{a=1}^M \sum_{j=1}^{N_a} \left\{ \log f^*(\delta_{j+1}^a | \mathbf{h}_j^a) + \log p(x_j^a | \mathbf{h}_{j-1}^a) \right\}, \quad (12)$$

where  $f^*$  is the conditional likelihood function of the RPP model defined in Eq. (11);  $\delta_{j+1}^a \equiv t_{j+1}^a - t_j^a$  denotes the inter-review time for item  $a$ ; and  $p(\mathbf{x}_j^a | \mathbf{h}_{j-1}^a)$  labels the joint probability distribution over the sequence of words in  $\mathbf{r}_j^a$ , as defined in Eq. (2).

## IV. BASELINE MODELS

In this section we present a set of dynamic review models to be used as baselines [39] in the experiment section below. These baselines will help us test the relevance of some of the properties of the RPP model, as well as of some of the hypothesis implicitly made while defining it. We shall replace the auto-regressive neural language models (LM) of Section III-A with a Bag of Words (BoW) Neural Text Model. This allows us to test, for example, whether the review word order (i.e. its grammar and distributed semantics) helps in modelling the arrival times of new reviews, or if it is enough to simply know the presence of some key words to successfully perform this task. We also consider replacing the RPP model of Section III-B with a simpler LSTM model. Such a modification helps us test e.g. whether the point-process nature of how reviews are received in time (e.g. its self-exciting properties) really plays a role when modelling the evolution of review content.

In total we consider four combinations of dynamic review models: (i) RPP+(BoW) in which we replaced the autoregressive LM with a BoW model, (ii) t-LSTM+(RNN) and (iii) t-LSTM+(CNN), for which we replaced the RPP model with

Dataset	#reviews	#items	#users	#sentences	#words
Yelp-Shopping	227K	12K	174K	1.6M	11.5M

TABLE II: Yelp shopping category dataset statistics.

a simple LSTM, and (iv) t-LSTM+(BoW), for which both language and RPP models were replaced. Below we describe the BoW and t-LSTM models in detail.

#### A. Bag of Words (BoW) Neural Text Model

We start by transforming the text of each review into a *bag of words* (BoW) representation  $\mathbf{X}_j \in \mathbb{R}^V$ , where  $V$  is the vocabulary size [40]. To model the text component of reviews we assume the words in review  $\mathbf{r}_{j+1}$  are generated independently, conditioned on  $\mathbf{h}_j$ , a vector representation describing the review creation dynamics. Note that  $\mathbf{h}_j$  can come from either the RPP model of Section III-B or the t-LSTM model of Section IV-B. Following [41] we write the conditional probability of generating the  $i$ th word  $\mathbf{w}_i^{j+1}$  of the  $(j+1)$ th review as

$$p_\theta(\mathbf{w}_i^{j+1}|\mathbf{h}_j) = \frac{\exp\{-z(\mathbf{w}_i^j, \mathbf{h}_j)\}}{\sum_{k=1}^V \exp\{-z(\mathbf{w}_k^j, \mathbf{h}_j)\}}, \quad (13)$$

$$z(\mathbf{w}_i^j, \mathbf{h}_j) = -\mathbf{h}_j^T \mathbf{R} \mathbf{w}_i^j - \mathbf{b} \mathbf{w}_i^j, \quad (14)$$

where  $\mathbf{R} \in \mathbb{R}^{H \times V}$  and  $\mathbf{b} \in \mathbb{R}^V$  are trainable parameters, and  $\mathbf{w}_i^j$  is the one-hot representation of the word at position  $i$  of review  $j$ . We train this model via maximum likelihood.

#### B. t-LSTM model for review creation

The inter-review time  $\delta_{j+1} \equiv t_{j+1} - t_j$  can be modelled as the mean of an exponential distribution with parameter  $\lambda_\theta(\mathbf{h}_j)$ . The function  $\lambda_\theta$  is modelled by a multilayer perceptron and the hidden state encoding the history of review arrivals is given by

$$\mathbf{h}_j = f_\theta(t_j, \alpha_j, \mathbf{h}_{j-1}), \quad (15)$$

where  $f_\theta$  is implemented by an LSTM network and  $\alpha_j$  labels a text representation which can either be the summary representation  $\bar{\mathbf{s}}_j$  from the language models of Section III-A or the BoW representation  $\mathbf{X}_j$  of Section IV-A.

	mean	std	max	50%	95%	99%
reviews	17.57	29.99	1 147	10	53	123
sentences	7.32	5.93	99	6	18	31
words	50.41	46.40	493	36	136	238

TABLE III: Review statistics.

## V. EXPERIMENTS AND RESULTS

### A. Data set

For performance validation of the models we choose the Yelp19 dataset.<sup>4</sup> We take all reviews for businesses that are labeled with the *shopping* parent category from *01 Jan 2016* to *30 Nov 2018*. The creation time of a review is defined as the difference in days between the original timestamp and *01*

<sup>4</sup><https://www.yelp.com/dataset>

*Jan 2016*. Next, we group reviews by business. All businesses with less than 5 reviews are removed. The sequential language models use the raw text from the reviews changed into lower case. In contrast, we convert the text from each review into a BoW vector of size 2000 [40] for the BoW models. Preprocessing scripts can be found at [1].

The result from the preprocessing is summarized in Tables II and III. Note that the average number of reviews per business is 17.57, and 50% of the businesses have less or equal to 10 reviews. The average length of a review is 50.41 sentences. This information gives insights helpful in the hyperparameter search during training. Finally, Figure 5 shows the number of posts per business distribution (Fig. 5a) and the *log* inter-arrival distribution (Fig. 5b), as well as the number of sentences (Fig. 5c) and number of words (Fig. 5d) per review distributions.

### B. Training and Model Configuration

*Training.* The RPP models training on the objective Eq. (12) has shown to be challenging. This is due to the fact that each model component (RPP and dynamic LM) use the latent representation from the other ( $\bar{\mathbf{s}}_j$  and  $\mathbf{h}_j$ , respectively). Since we are using very flexible models, we have to be careful not to end up in a situation wherein one of the models is significantly much better than the other. Such a situation would cause one, or maybe both, of the models to ignore the latent embedding from the other one. To prevent such a situation, the frequency of gradient updates for each model should be tuned. Bear in mind that the RPP model is trained with *backpropagation through time* (BPTT) [42], [43], as to avoid the inherent RNN problems – like vanishing or exploding gradient. Therefore, our propose solution is to run two forward steps for each BPTT window. In the first step, we freeze the parameters of the RPP model, and only update the parameters of the dynamic LM at each time step in the window, while storing the summary content representations. For the second forward run we freeze the parameters of the LM, and use the stored  $\bar{\mathbf{s}}_j$  as input to compute the gradients of the RPP model. This time we can perform gradient updates at the end of the BPTT window. Using one optimizer to train both models showed to be unstable.<sup>5</sup>

*Model Configuration.* In our experiments, we randomly split each dataset into two parts: training set (80%) and test set (20%). We use grid search for hyper-parameters tuning. The temporal component of all dynamic review models have hidden state size of  $H = 128$ , and we project  $\bar{\mathbf{s}}_j$  into a 16-dimensional space before concatenation. The LM-RNN model has a hidden dimension  $S$  of 1024. We set the convolution filter to size  $k = 3$  for the LM-CNN model, and choose to use 6 residual blocks with dilation [1, 2, 4, 8, 16, 32], respectively. Weight normalization [45] is used between the convolution layers. The vocabulary size is  $V = 3000$  for all sequential language models and we use the GloVe word embeddings [46], with dimension  $W = 300$ .

<sup>5</sup>This is due to the fact that optimizers like Adam [44] keep information about the dynamics of training. Using the same optimizer for the two models will lead to instabilities since the models have different learning dynamics.



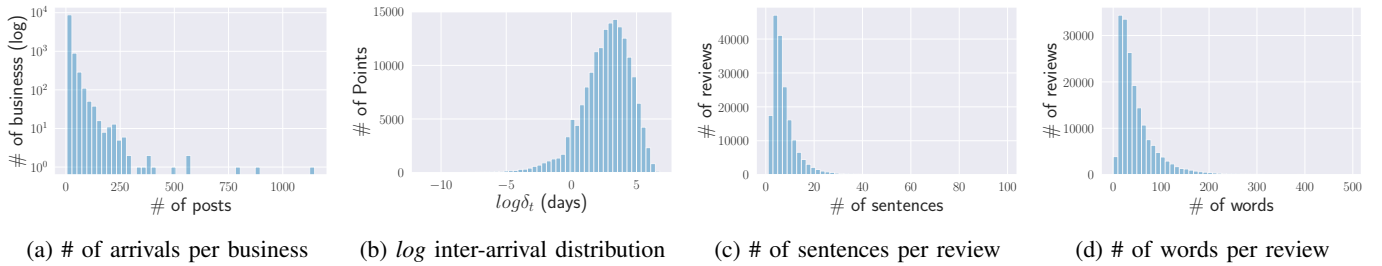


Fig. 5: Empirical distribution for different aspects of the data.

We use Adam [44] to optimize the language model components of our models, with learning rate 0.0002 and  $\beta_1 = 0.9$ . The optimization of the temporal component was done with Adadelta [47], initialized with a learning rate of 1. The size of the BPTT window for the RPP model is 20, since the average number of arrivals per business is 17.57. We limit the length of each review to be of 80 tokens. All methods are implemented using PyTorch v1.3<sup>6</sup>. Source code for all models can be found at [1].

### C. Results

Given an item (business) of interest, the RPR model predicts both the arrival time of its next review and the probability of the word sequences within that review. To quantitatively evaluate the performance of the model on these tasks, we compute the root-mean-squared error (RMSE)  $\varepsilon$  on the inter-review times, and the perplexity per word  $\mathcal{P}$  with respect to the review content. Specifically, we define the RMSE as

$$\varepsilon = \sqrt{\frac{\sum_a^M \sum_j^{N_a} |\delta_j^a - \langle \tilde{\delta}_j^a \rangle|}{MN_a}}, \quad (16)$$

where  $\langle \tilde{\delta}_j^a \rangle$  is the mean predicted inter-review time<sup>7</sup> and  $\delta_j^a$  is its empirical value,  $N_a$  is the number of reviews for business  $a$  and  $M$  the total number of businesses. Likewise we define the perplexity per word as

$$\mathcal{P} = \exp \left\{ - \frac{\sum_a^M \sum_j^{N_a} \sum_i^{L_{a,j}} \log p(\mathbf{w}_i^{a,j} | \mathbf{w}_{<i}^{a,j}, \mathbf{h}_{j-1}^a)}{MN_a L_{a,j}} \right\}, \quad (17)$$

where  $\mathbf{w}_i^{a,j}$  and  $L_{a,j}$  label the  $i$ th word and the number of words, respectively, in the  $j$ th review of business  $a$ .

To set yardsticks by which to judge the performance of our models, let us first analyse the review arrival process and the word distribution of the review content separately, that is, as though they were independent. Consider the time series defined by the review creation times  $\{t_j^a\}_{j=1}^{N_a}$ . We model the inter-review time of these series with (i) the RPP model of section III-B and (ii) the t-LSTM model of section IV-B, excluding in both cases any representation ( $\bar{\mathbf{s}}_j$  or  $\mathbf{X}_j$ ) encoding

information from the reviews' content. The RMSE obtained via these models is shown in the first two rows of Table IV. Note how the RPP model slightly outperforms the LSTM model. Likewise, consider the aggregated set of all review content received by all the businesses we study. We compute the perplexity per word of all reviews with respect to two language model, one parametrized by a LSTM and the other by a TCN. Let us refer to these models as LM-RNN and LM-TCN. Their perplexity is shown in the third and fourth rows of Table IV. Note that for this data set the RNN performs better than the TCN.

With these values as yardsticks we are now in position to discuss our results. Let us start by evaluating how well our dynamic review models predict the arrival time of new reviews. The results are shown in the second column of Table IV. The first thing one can conclude from these results is that, as expected, review content does provide useful information for predicting when are the next reviews going to arrive. Even using a simple text representation as that of a BoW improves the RMSE values, as compared to the yardsticks. This makes sense intuitively, for one can imagine that reviews with similar content clusters in time. The last four rows of column 2 in Table IV show our results for models whose content representation comes from auto-regressive LMs. All these instances beat the BoW models, which hints at the importance of accounting for word order. Finally our RPR models (with either RNN or TCN) improve by about 25% the RMSE values as compared to the yardsticks, and about 19% as compared to the baselines. This last comparison is surprising, given that RPP and t-LSTM models with no review content information yield similar RMSE values. Thus a marked point process model which uses the content summary representations from auto-regressive LMs as marks describes significantly better the empirical review arrival process, as compared to e.g. a stand alone RPP.

Having seen that the review content information helps when modeling the review arrival process for the businesses we study, we can now ask whether the opposite flow of information is also useful, i.e. does the dynamics of review arrival help in modeling the content of reviews? The last column in Table IV shows the perplexity per word computed over all reviews received by a collection of businesses (i.e. aggregated over time). Low perplexity values indicate the model is able to predict well the sentences in a given review.

<sup>6</sup><https://pytorch.org/>

<sup>7</sup>To compute  $\langle \tilde{\delta}_j^a \rangle$  via the RPP model we sample  $t_j$  1000 times using Eq. (21). In the t-LSTM model  $\langle \tilde{\delta}_j^a \rangle$  is obtained directly from  $\lambda_\theta$ .

Model	RMSE ↓	R <sup>2</sup> ↑	Perplexity ↓
t-LSTM	96.8813	0.1788	-
RPP	96.3794	0.1873	-
LM-RNN	-	-	32.09
LM-TCN	-	-	32.81
t-LSTM+(BoW)	95.3414	0.2046	519.90*
RPP+(BoW)	92.3850	0.2533	511.32*
t-LSTM+(RNN)	89.9555	0.2920	29.86
t-LSTM+(TCN)	89.8053	0.2944	31.74
RPR+(RNN)	<b>71.8748</b>	<b>0.5480</b>	<b>29.61</b>
RPR+(TCN)	74.0945	0.5197	31.48

TABLE IV: Model performance on RMSE, R<sup>2</sup> and perplexity per word. The RPR models significantly outperforms all baselines in all metrics. The perplexity for the BoW models corresponds to the predictive perplexity as defined in [14].

Let us start with the BoW models. By definition, these models do not take into account the word sequence in a review  $\mathbf{r}_j$ , and can only predict the distribution of words in the next review  $\mathbf{r}_{j+1}$ . Accordingly, their perplexity values corresponds to the predicted perplexity, as defined in [14]. Despite not being able to directly compare the BoW models with the LMs, we can see that RPP+(BoW) performs better than t-LSTM+(BoW). Let us move on now to our dynamic review models with summary content representation  $\mathbf{s}_i$ . All models yield lower perplexity values than the yardsticks. Therefore all models are able to successfully use the vector representation  $\mathbf{h}_j$  encoding the review arrival dynamics. We can conclude that its inclusion does make sequence words prediction easier. Note, in particular, that the RPR+(RNN) model outperforms all baselines. Finally, we have also computed the instantaneous perplexity, or perplexity per review, and noticed that the RPR models consistently yield better values than all baselines.

## VI. CONCLUSION AND FUTURE WORK

In this work we introduced neural dynamic language models of text for review data. We are able to leverage dynamic representations of point process models in language modelling tasks, and augment the point processes with text representations. We showed that our approach improves performance on both content and arrival times prediction, as well as opens the door for *dynamic generative language models*. Future work includes the implementation of attention mechanisms, as well as the inclusion of neural factorization machines aimed at predicting ratings values.

## VII. APPENDIX

We start by denoting with  $p(T|\mathcal{H}_j)$  the probability that the next point arrives at  $T$  given the previous history until the arrival of  $\mathbf{r}_j$  — we require  $p(T|\mathcal{H}_j)$  for both prediction and sampling. First, notice that the probability of no review arriving between  $t_j$  and  $t_j + \tau$  can be obtained as an integral over  $p(T|\mathcal{H}_j)$ , say

$$\exp \left\{ - \int_{t_j}^{t_j + \tau} \lambda(t) dt \right\} = \int_{\tau}^{\infty} p(T|\mathcal{H}_j) dT \equiv G(\tau),$$

with  $p(T|\mathcal{H}_j) = -\frac{dG(T)}{dT}$ , where we used the Poisson distribution for zero arrivals in the first expression. Solving for  $G(\tau)$  we find

$$G(\tau) = \exp \left\{ -e^{\alpha_j} \frac{1}{w^t} \left( e^{w^t \tau} - 1 \right) \right\}, \quad (18)$$

with  $\alpha_j = \mathbf{v}^t \mathbf{h}_j + b^t$ . The average time of the next arrival is then given by

$$\mathbb{E}[T] = \int_0^{\infty} p(T|\mathcal{H}_j) T dT = \int_0^{\infty} G(T) dT. \quad (19)$$

Finally, in order to sample the next arrival time one can use inverse transform sampling on  $P(T)$ . To this end one requires the inverse of the cumulative function of  $P(T)$ . We calculate the cumulative function thus

$$\begin{aligned} F[p(T|\mathcal{H}_j)] &= \int_0^{\tau} p(T|\mathcal{H}_j) dT \\ &= - \int_0^{\tau} \frac{dG(T)}{dT} dT = G(0) - G(\tau) \end{aligned} \quad (20)$$

whose inverse function then follows

$$\begin{aligned} F^{-1}[p(T|\mathcal{H}_j)](y) &= \frac{1}{w^t} \left( -\alpha_j + \log \left\{ w^t \left( \log \left\{ \frac{-1}{y-1} \right\} \right. \right. \right. \\ &\quad \left. \left. \left. + \frac{e^{\alpha_j}}{w^t} \right) \right\} \right). \end{aligned} \quad (21)$$

## ACKNOWLEDGMENT

The authors of this work were supported by the Fraunhofer Research Center for Machine Learning (RCML), within the Fraunhofer Cluster of Excellence Cognitive Internet Technologies (CCIT), and by the Competence Center for Machine Learning Rhine Ruhr (ML2R) which is funded by the Federal Ministry of Education and Research of Germany (grant no. 01—S18038A). Part of the work was also funded by the German Ministry for Education and Research as BIFOLD-Berlin Institute for the Foundations of Learning and Data (ref. 01IS18025A and ref 01IS18037A). We gratefully acknowledge this support.

## REFERENCES

- [1] “Source code,” <https://github.com/evejoski/recurrent-point-review-models>.
- [2] D. Liu, J. Li, B. Du, J. Chang, and R. Gao, “Daml: Dual attention mutual learning between ratings and reviews for item recommendation,” in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. ACM, 2019, pp. 344–352.
- [3] C.-Y. Wu, A. Ahmed, A. Beutel, A. J. Smola, and H. Jing, “Recurrent recommender networks,” in *Proceedings of the tenth ACM international conference on web search and data mining*. ACM, 2017, pp. 495–503.
- [4] N. Sachdeva, G. Manco, E. Ritacco, and V. Pudi, “Sequential variational autoencoders for collaborative filtering,” in *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*. ACM, 2019, pp. 600–608.
- [5] L. Zheng, V. Noroozi, and P. S. Yu, “Joint deep modeling of users and items using reviews for recommendation,” in *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*. ACM, 2017, pp. 425–434.
- [6] S. Kirby, M. Dowman, and T. L. Griffiths, “Innateness and culture in the evolution of language,” *Proceedings of the National Academy of Sciences*, vol. 104, no. 12, pp. 5241–5245, 2007.

- [7] M. A. Nowak and D. C. Krakauer, "The evolution of language," *Proceedings of the National Academy of Sciences*, vol. 96, no. 14, pp. 8028–8033, 1999.
- [8] S. J. Greenhill, C.-H. Wu, X. Hua, M. Dunn, S. C. Levinson, and R. D. Gray, "Evolutionary dynamics of language systems," *Proceedings of the National Academy of Sciences*, vol. 114, no. 42, pp. E8822–E8829, 2017.
- [9] C. Danescu-Niculescu-Mizil, R. West, D. Jurafsky, J. Leskovec, and C. Potts, "No country for old members: User lifecycle and linguistic change in online communities," in *Proceedings of the 22nd international conference on World Wide Web*. ACM, 2013, pp. 307–318.
- [10] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in *Advances in neural information processing systems*, 2013, pp. 3111–3119.
- [11] W. L. Hamilton, J. Leskovec, and D. Jurafsky, "Diachronic word embeddings reveal statistical laws of semantic change," *arXiv preprint arXiv:1605.09096*, 2016.
- [12] V. Kulkarni, R. Al-Rfou, B. Perozzi, and S. Skiena, "Statistically significant detection of linguistic change," in *Proceedings of the 24th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 2015, pp. 625–635.
- [13] M. Rudolph and D. Blei, "Dynamic embeddings for language evolution," in *Proceedings of the 2018 World Wide Web Conference*. International World Wide Web Conferences Steering Committee, 2018, pp. 1003–1011.
- [14] C. Wang, D. Blei, and D. Heckerman, "Continuous time dynamic topic models," *arXiv preprint arXiv:1206.3298*, 2012.
- [15] D. Yogatama, C. Wang, B. R. Routledge, N. A. Smith, and E. P. Xing, "Dynamic language models for streaming text," *Transactions of the Association for Computational Linguistics*, vol. 2, pp. 181–192, 2014.
- [16] N. Du, M. Farajtabar, A. Ahmed, A. J. Smola, and L. Song, "Dirichlet-hawkes processes with applications to clustering continuous-time document streams," in *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2015, pp. 219–228.
- [17] X. He, T. Rekatsinas, J. Foulds, L. Getoor, and Y. Liu, "Hawkestopic: A joint model for network inference and topic modeling from text-based cascades," in *International conference on machine learning*, 2015, pp. 871–880.
- [18] E. Delasalles, S. Lamprier, and L. Denoyer, "Dynamic neural language models," in *Neural Information Processing*, T. Gedeon, K. W. Wong, and M. Lee, Eds. Cham: Springer International Publishing, 2019, pp. 282–294.
- [19] R. Catherine and W. Cohen, "Transnets: Learning to transform for recommendation," in *Proceedings of the Eleventh ACM Conference on Recommender Systems*. ACM, 2017, pp. 288–296.
- [20] M. Fan, C. Feng, L. Guo, M. Sun, and P. Li, "Product-aware helpfulness prediction of online reviews," in *The World Wide Web Conference*. ACM, 2019, pp. 2715–2721.
- [21] S. Novgorodov, G. Elad, I. Guy, and K. Radinsky, "Generating product descriptions from user reviews," in *The World Wide Web Conference*. ACM, 2019, pp. 1354–1364.
- [22] B. Esmaili, H. Huang, B. Wallace, and J.-W. van de Meent, "Structured neural topic models for reviews," in *The 22nd International Conference on Artificial Intelligence and Statistics*, 2019, pp. 3429–3439.
- [23] S. Chen, C. Li, F. Ji, W. Zhou, and H. Chen, "Driven answer generation for product-related questions in e-commerce," in *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*. ACM, 2019, pp. 411–419.
- [24] Q. Yu and W. Lam, "Aware answer prediction for product-related questions incorporating aspects," in *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*. ACM, 2018, pp. 691–699.
- [25] T. Mikolov, M. Karafiát, L. Burget, J. Cernocký, and S. Khudanpur, "Recurrent neural network based language model," in *INTERSPEECH*, T. Kobayashi, K. Hirose, and S. Nakamura, Eds. ISCA, 2010, pp. 1045–1048.
- [26] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, pp. 1735–80, 12 1997.
- [27] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner *et al.*, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [28] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *nature*, vol. 521, no. 7553, p. 436, 2015.
- [29] M. Denil, A. Demiraj, N. Kalchbrenner, P. Blunsom, and N. de Freitas, "Modelling, visualising and summarising documents with a single convolutional neural network," *arXiv preprint arXiv:1406.3830*, 2014.
- [30] S. Bai, J. Z. Kolter, and V. Koltun, "An empirical evaluation of generic convolutional and recurrent networks for sequence modeling," *arXiv preprint arXiv:1803.01271*, 2018.
- [31] F. Yu and V. Koltun, "Multi-scale context aggregation by dilated convolutions," *arXiv preprint arXiv:1511.07122*, 2015.
- [32] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016, pp. 770–778.
- [33] N. Kalchbrenner, L. Espeholt, K. Simonyan, A. van den Oord, A. Graves, and K. Kavukcuoglu, "Neural machine translation in linear time," 2016. [Online]. Available: <https://arxiv.org/abs/1610.10099>
- [34] Z. Yang, Z. Hu, R. Salakhutdinov, and T. Berg-Kirkpatrick, "Improved variational autoencoders for text modeling using dilated convolutions," in *Proceedings of the 34th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, D. Precup and Y. W. Teh, Eds., vol. 70. International Convention Centre, Sydney, Australia: PMLR, 06–11 Aug 2017, pp. 3881–3890.
- [35] D. J. Daley and D. Vere-Jones, *An introduction to the theory of point processes: volume II: general theory and structure*. Springer Science & Business Media, 2007.
- [36] H. Mei and J. M. Eisner, "The neural hawkes process: A neurally self-modulating multivariate point process," in *Advances in Neural Information Processing Systems*, 2017, pp. 6738–6748.
- [37] N. Du, H. Dai, R. Trivedi, U. Upadhyay, M. Gomez-Rodriguez, and L. Song, "Recurrent marked temporal point processes: Embedding event history to vector," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2016, pp. 1555–1564.
- [38] —, "Recurrent marked temporal point processes: Embedding event history to vector," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2016, pp. 1555–1564.
- [39] K. Cvejoski, R. J. Sanchez, B. Georgiev, J. Schuecker, C. Bauckhage, and C. Ojeda, "Recurrent point processes for dynamic review models," in *Workshop on Interactive and Conversational Recommendation Systems at AAAI*, 2020.
- [40] G. E. Hinton and R. R. Salakhutdinov, "Replicated softmax: an undirected topic model," in *Advances in Neural Information Processing Systems 22*, Y. Bengio, D. Schuurmans, J. D. Lafferty, C. K. I. Williams, and A. Culotta, Eds. Curran Associates, Inc., 2009, pp. 1607–1614. [Online]. Available: <http://papers.nips.cc/paper/3856-replicated-softmax-an-undirected-topic-model.pdf>
- [41] Y. Miao, L. Yu, and P. Blunsom, "Neural variational inference for text processing," in *International conference on machine learning*, 2016, pp. 1727–1736.
- [42] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning internal representations by error propagation," California Univ San Diego La Jolla Inst for Cognitive Science, Tech. Rep., 1985.
- [43] P. J. Werbos *et al.*, "Backpropagation through time: what it does and how to do it," *Proceedings of the IEEE*, vol. 78, no. 10, pp. 1550–1560, 1990.
- [44] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [45] T. Salimans and D. P. Kingma, "Weight normalization: A simple reparameterization to accelerate training of deep neural networks," in *Advances in Neural Information Processing Systems*, 2016, pp. 901–909.
- [46] J. Pennington, R. Socher, and C. D. Manning, "Glove: Global vectors for word representation," in *Empirical Methods in Natural Language Processing (EMNLP)*, 2014, pp. 1532–1543. [Online]. Available: <http://www.aclweb.org/anthology/D14-1162>
- [47] M. D. Zeiler, "Adadelata: an adaptive learning rate method," *arXiv preprint arXiv:1212.5701*, 2012.