
HIGHLY FAST TEXT SEGMENTATION WITH PAIRWISE MARKOV CHAINS

Elie Azeraf*
Watson Department
IBM GSB France
elie.azeraf@ibm.com

Emmanuel Monfrini
SAMOVAR, Telecom SudParis
Institut Polytechnique de Paris

Emmanuel Vignon
Watson Department
IBM GSB France

Wojciech Pieczynski
SAMOVAR, Telecom SudParis
Institut Polytechnique de Paris

ABSTRACT

Natural Language Processing (NLP) models' current trend consists of using increasingly more extra-data to build the best models as possible. It implies more expensive computational costs and training time, difficulties for deployment, and worries about these models' carbon footprint reveal a critical problem in the future. Against this trend, our goal is to develop NLP models requiring no extra-data and minimizing training time. To do so, in this paper, we explore Markov chain models, Hidden Markov Chain (HMC) and Pairwise Markov Chain (PMC), for NLP segmentation tasks. We apply these models for three classic applications: POS Tagging, Named-Entity-Recognition, and Chunking. We develop an original method to adapt these models for text segmentation's specific challenges to obtain relevant performances with very short training and execution times. PMC achieves equivalent results to those obtained by Conditional Random Fields (CRF), one of the most applied models for these tasks when no extra-data are used. Moreover, PMC has training times 30 times shorter than the CRF ones, which validates this model given our objectives.

Keywords Chunking · Hidden Markov Chain · Named Entity Recognition · Pairwise Markov Chain · Part-Of-Speech tagging

1 Introduction

In the past ten years, developments of Deep Learning methods [1, 2] have enabled research in Natural Language Processing (NLP) to take an impressive leap. Some tasks, like Question Answering [3] or Sentiment Analysis [4], seemed unrealistic twenty years ago, and nowadays recent neural models achieved better scores than humans [5] [6] for these applications. This dynamic's main motivation is the direct applications of NLP models in the industry, with tasks such as Named-Entity-Recognition and mail classification. However, the cost of improving models' performances increases, and learning algorithms always need more data and power to be trained and make a prediction. We can produce models that achieve impressive scores, but deployment and climate change problems [7] raise some issues. Our aim in this paper is to initiate a reflection around light models by introducing a new Markov model design for text

*Elie Azeraf is also a member of SAMOVAR, Telecom SudParis, Institut Polytechnique de Paris

segmentation and comparing it with machine learning algorithms that have a reasonable carbon impact and execution time. Motivation about the choice of segmentation tasks is explained later in this paper.

The Hidden Markov Chains (HMC), introduced by Stratonovitch sixty years ago [8] [9] [10] [11] [12] [13], which model poor correlations, are widely used in machine learning, and have especially been applied for NLP segmentation tasks [14] [15] [16].

For over twenty years, HMCs have been strictly generalized to Pairwise Markov Chains (PMCs) [17], a family of models including HMCs. In PMCs, the “hidden” chain is not necessarily Markov, and the noise is modeled in a more correlated – and thus more informative – way. However, PMCs keep the same advantages as HMCs in regards to the hidden data estimation. In particular, the training and estimation tasks still have linear complexity. PMCs have especially been studied for image segmentation, with discrete hidden variables and continuous observations. It turns out that using PMCs instead of HMCs can divide the error rate by two [18] [19]. However, for specific reasons relative to language processing, which we will develop later, PMCs have never been applied for NLP tasks.

This paper explores the interest in using PMCs for three of the main text segmentation tasks: Part-Of-Speech (POS) Tagging, Chunking, and Named-Entity-Recognition (NER). The best methods for these tasks are based on Deep Learning models [20] [21]. However, to produce excellent scores, these models require a large amount of extra-data, which results in very long training time and difficulties in deploying them with classic architectures.

The paper is organized as follows. In the next section, we present and compare PMCs and HMCs, the bayesian segmentation methods, and the parameter estimation algorithm. The third section is devoted to the text segmentation tasks and an original way to adapt Markov chain models for these tasks while keeping fast training and relevant results. Experiments are presented in section four. The last section is devoted to conclusions and perspectives.

2 Markov Chain Models

Let $X_{1:T} = (X_1, \dots, X_T)$ and $Y_{1:T} = (Y_1, \dots, Y_T)$ be two discrete stochastic processes. For all t in $\{1, \dots, T\}$, X_t takes its values in $\Lambda_X = \{\lambda_1, \dots, \lambda_N\}$ and Y_t takes its values in $\Omega_Y = \{\omega_1, \dots, \omega_M\}$. Let us then consider the process $Z_{1:T} = (X_{1:T}, Y_{1:T})$.

Our study takes place in the case of latent variables, with an observed realization $y_{1:T}$ of $Y_{1:T}$, and the corresponding hidden realization $x_{1:T}$ of $X_{1:T}$ having to be estimated.

In the following, in order to simplify the notation, for all t in $\{1, \dots, T\}$, events $\{X_t = x_t\}$, $\{Y_t = y_t\}$ and $\{Z_t = z_t\}$ will be denoted by $\{x_t\}$, $\{y_t\}$ and $\{z_t\}$, respectively.

2.1 HMC - PMC

Starting from the most correlated shape of the couple $(X_{1:T}, Y_{1:T})$, we can say that $Z_{1:T}$ is an HMC if it is possible to write its probability law as:

$$p(z_{1:T}) = p(x_1)p(y_1|x_1)p(x_2|x_1)p(y_2|x_2) \dots p(x_T|x_{T-1})p(y_T|x_T) \quad (1)$$

Moreover, for all $\lambda_i, \lambda_j \in \Lambda_X$ and $\omega_k \in \Omega_Y$, we consider homogeneous HMCs defined with the following parameters:

- Initial probability $\pi = (\pi(1), \dots, \pi(N))$ with $\pi(i) = p(x_1 = \lambda_i)$
- Transition matrix $A = \{a_i(j)\}$ with $\forall t \in \{1, \dots, T-1\}, a_i(j) = p(x_{t+1} = \lambda_j | x_t = \lambda_i)$
- Emission matrix $B = \{b_i(k)\}$ with $\forall t \in \{1, \dots, T\}, b_i(k) = p(y_t = \omega_k | x_t = \lambda_i)$

An oriented dependency graph of the HMC is given in figure 1.

We can note that the dependency graph is almost reduced to the minimum to link dependent couple of sequential data, and yet HMC is known to be a robust model. PMCs, which we will detail just after, allow to introduce more connections to the graph, as shown in figure 2.

Starting from the most correlated shape of the couple $(x_{1:T}, y_{1:T})$, we can say that $z_{1:T}$ is a PMC if $z_{1:T}$ is a Markov chain, which is equivalent to admitting that its distribution is of the form:

$$p(z_{1:T}) = p(z_1)p(z_2|z_1)\dots p(z_T|z_{T-1}) \quad (2)$$




Figure 1: Probabilistic graphical model of HMC




Figure 2: Probabilistic graphical model of PMC

which can be rewritten, without loss of generality,

$$\begin{aligned}
 p(z_{1:T}) = & p(x_1)p(y_1|x_1)p(x_2|x_1, y_1)p(y_2|x_2, x_1, y_1) \\
 & \dots p(x_T|x_{T-1}, y_{T-1})p(y_T|x_T, x_{T-1}, y_{T-1})
 \end{aligned} \tag{3}$$

When comparing (1) and (3) we can see that a PMC is a HMC [17] if and only if $\forall t \in \{1, \dots, T-1\}$,

- $p(x_{t+1}|x_t, y_t) = p(x_{t+1}|x_t)$ and
- $p(y_{t+1}|x_{t+1}, x_t, y_t) = p(y_{t+1}|x_{t+1})$.

From a graphical point of view, we can materialize on figures 1 and 2 how PMCs are more general than HMCs. In particular, the markovian assumption of $x_{1:T}$ made for HMCs is not supposed for PMCs, but the determinist bayesian inference is still possible.

In this paper, we consider homogeneous PMCs defined with three sets of parameters, $\forall \lambda_i, \lambda_j \in \Lambda_X, \forall \omega_k, \omega_l \in \Omega_Y, \forall t \in \{1, \dots, T-1\}$:

- Initial probability matrix $\Pi^{PMC} = \{\pi^{PMC}(i, k)\}$ with $\pi^{PMC}(i, k) = p(x_1 = \lambda_i, y_1 = \omega_k)$;
- Transition matrix $A^{PMC} = \{a_{i,k}^{PMC}(j)\}$ with: $a_{i,k}^{PMC}(j) = p(x_{t+1} = \lambda_j | x_t = \lambda_i, y_t = \omega_k)$;
- Emission matrix $B^{PMC} = \{b_{i,j,k}^{PMC}(l)\}$: $b_{i,j,k}^{PMC}(l) = p(y_{t+1} = \omega_l | x_t = \lambda_i, x_{t+1} = \lambda_j, y_t = \omega_k)$.

2.2 Bayesian segmentation

There are two Bayesian methods to estimate the realization $\hat{x}_{1:T}$ of the hidden chain: the Marginales Posterior Mode (MPM) and the Maximum A Posteriori (MAP).

The MPM estimator is given by :

$$\begin{aligned}
 \hat{x}_{MPM} = & (\hat{x}_1, \dots, \hat{x}_T) \text{ with } \forall t \in \{1, \dots, T\}, \\
 p(\hat{x}_t|y_{1:T}) = & \sup_{\lambda_i \in \Lambda_X} p(x_t = \lambda_i | y_{1:T})
 \end{aligned}$$

and the MAP estimator is :

$$\hat{x}_{MAP} = \arg \max_{x_{1:T} \in (\Lambda_X)^T} p(x_{1:T} | y_{1:T})$$

MPM estimator is given with Forward-Backward algorithm [9] [12], while MAP is given by the Viterbi one [22]. We tested both of them, and MPM gives slightly better results than MAP in term of accuracy for the tasks we consider, so we only present results for MPM.

We thus have to compute $\forall t \in \{1, \dots, T\}, \forall \lambda_i \in \Lambda_X, p(x_t = \lambda_i | y_{1:T})$ to maximize posterior marginals with the most probable state at every time $t \in \{1, \dots, T\}$. We present the Forward-Backward algorithm in the case of PMC which generalizes the classical HMC one.

$\forall t \in \{1, \dots, T\}, \forall \lambda_i \in \Lambda_X$:

$$p(x_t = \lambda_i | y_{1:T}) = \frac{\alpha_t(i)\beta_t(i)}{\sum_{j \in \Lambda_X} \alpha_t(j)\beta_t(j)}$$

where $\forall t \in \{1, \dots, T\}, \forall \lambda_i \in \Lambda_X$:

$$\begin{aligned} \alpha_t(i) &= p(y_{1:t}, x_t = \lambda_i) \\ \beta_t(i) &= p(y_{t+1:T} | x_t = \lambda_i, y_t) \end{aligned}$$

$\forall \lambda_i \in \Lambda_X, \forall t \in \{1, \dots, T\}$, $\alpha_t(i)$ and $\beta_t(i)$ can still be computed thanks to forward and backward recursions:

- with $y_1 = \omega_k$,

$$\alpha_1(i) = \pi^{PMC}(i, k)$$

- $\forall t \in \{1, \dots, T-1\}$, with $y_t = \omega_k, y_{t+1} = \omega_l$,

$$\alpha_{t+1}(i) = \sum_{\lambda_j \in \Lambda_X} a_{j,i}^{PMC}(k) b_{j,i,k}^{PMC}(l) \alpha_t(j)$$

and

- $\beta_T(i) = 1$
- $\forall t \in \{1, \dots, T-1\}$ with $y_t = \omega_k, y_{t+1} = \omega_l$,

$$\beta_t(i) = \sum_{\lambda_j \in \Lambda_X} a_{i,j}^{PMC}(k) b_{i,j,k}^{PMC}(l) \beta_{t+1}(j)$$

One can normalize forward and backward probabilities at every step, which allows them to avoid underflowed computational problems without modifying the results. This algorithm can be executed with matrix computation, allowing a highly fast execution.

2.3 Parameter estimation

We estimate the parameters with maximum likelihood estimator [23] [14]. It consists in computing the empirical frequencies of the probabilities of interest. We have, $\forall \lambda_i, \lambda_j \in \Lambda_X, \forall \omega_k, \omega_l \in \Omega_Y$,

$$\begin{aligned} \hat{\pi}(i) &= \frac{N_i^0}{L}, \quad \hat{a}_i(j) = \frac{N_{i,j}}{N_i}, \quad \hat{b}_i(k) = \frac{M_{i,k}}{N_i} \\ \text{and } \hat{\pi}^{PMC}(i, k) &= \frac{N_{i,k}^0}{L}, \quad \hat{a}_{i,k}^{PMC}(j) = \frac{N_{i,k,j}}{M_{i,k}}, \\ \hat{b}_{i,j,k}^{PMC}(l) &= \frac{N_{i,k,j,l}}{N_{i,k,j}} \end{aligned}$$

where $N_{i,k,j,l}$ is the number of occurrences of the pattern $(x_t = \lambda_i, y_t = \omega_k, x_{t+1} = \lambda_j, y_{t+1} = \omega_l)$ in the L chains of the training set. Then, $N_{i,k,j} = \sum_{\omega_l \in \Omega_Y} N_{i,k,j,l}$, $N_{i,j} = \sum_{\omega_k \in \Omega_Y} N_{i,k,j}$, $M_{i,k} = \sum_{\lambda_j \in \Lambda_X} N_{i,k,j}$ and $N_i = \sum_{\lambda_j \in \Lambda_X} N_{i,j}$. Finally, N_i^0 and $N_{i,k}^0$ are respectively the number of times $x_1 = \lambda_i$ and $z_1 = (\lambda_i, \omega_k)$ in the L chains of the training set.

When $\text{card}(\Lambda_X)$ or $\text{card}(\Omega_Y)$ are huge, some patterns may not be observed in the training set, which implies that the corresponding estimation is zero. It is the case for NLP with Ω_Y , the space of possible written words. To minimize this problem, especially in PMC, we accept to partially “downgrade” PMC to HMC when necessary. This original process is represented in figure 3. It consists of approximating the forward and backward probabilities of the PMC by the HMC ones when those of PMC equal 0.

Moreover, it is essential to note that “online” learning is fast and easy in those models. If new sentences enrich the training set, updating parameters only requires us to “add” information to already determined values without retraining the model on the complete dataset.




Figure 3: Graphical model of a partially “downgraded” PMC - The couple of observations (y_2, y_3) and (y_5, y_6) never appeared in this order in the training set

3 Text segmentation tasks with Markov Chain models

Let us now have a few words on the three text segmentation tasks that we are interested in: POS Tagging, Chunking, and NER. For those three tasks, we observe the words of some sentences for which we have to find the specific labels.

POS tagging consists in labeling each word with its grammatical function as verb, determinant, adjective... For example, $(y_1, y_2, \dots, y_{12}) = (\text{John, likes, the, blue, house, at, the, end, of, the, street, .})$ and $(x_1, x_2, \dots, x_{12}) = (\text{Noun, Verb, Det, Adj, Noun, Prep, Det, Noun, Prep, Det, Noun, Punct})$. The performance of POS Tagging is evaluated with the accuracy error.

NER consists in discriminating the “entities” from the words of sentences. The “entities” can be a person (PER), a localization (LOC), or an organization (ORG). The set of entities depends on the use case. For example, $(\text{John, works, at, IBM, in, Paris, .})$ can have the following entities $(\text{PER, O, O, ORG, O, LOC, O})$ where O stands for “no entity”. The set of entities depends on the use case, for example, finding the name of protein or DNA in medical data as in [24]. The performance for NER is evaluated with F_1 score [25].

Chunking consists of segmenting a sentence into its sub-constituents, such as noun (NP), verb (VP), and prepositional phrases (PP). For example, $(\text{We, saw, the, yellow, dog, .})$ has the chunks: $(\text{NP, VP, NP, NP, NP, O})$. Like NER, the performance is evaluated with the F_1 score. More details about these tasks can be found in the NLTK book [26].

We took these three tasks as they can be the basis of more complex ones. For example, with text classification when the corpus is small, with limited architecture avoiding the deployment of a heavy deep learning model. This case can happen, for example, with email classification, with models having to process hundreds or thousands of emails per second with a maximum RAM power of 4Gb, given a relatively small training set. One way to handle this problem is to construct a text segmentation model, then another one using the predicted labels, to make a prediction. The three label types we consider, POS tag, chunk tag, and entity, are particularly useful to construct this type of architecture, motivating the choice of these tasks.

However, When working in the context of text segmentation, the cardinal of Ω_Y is very large, and managing unknown patterns (not in the train set) is then a real challenge. Although the Markov properties of our processes partially help to tackle this problem, it is not enough. PMC and HMC have to be adapted to consider this problem. Our goal is also to keep a very short training time. First of all, we can use the “downgrading” process from PMC to HMC described in Section II.B to alleviate this problem when PMC is used.

Moreover, for HMC, we look for extra information in the unknown observed words. Given a word ω , We introduce the following functions:

- $u(\omega) = 1$ if the first letter of ω is up, 0 otherwise
- $h(\omega) = 1$ if ω has a hyphen, 0 otherwise
- $f(\omega) = 1$ if ω is the first word of the sentence, 0 otherwise
- $d(\omega) = 1$ if ω has a digit, 0 otherwise
- $s_m(\omega) =$ the suffix of length m of ω

Then, for $\omega_k \in \Omega_Y$ unknown, and $\forall \lambda_i \in \Lambda_X$, we approximate $b_{i,k}$ by

$$b_i(k) \approx p(u(\omega_k), h(\omega_k), f(\omega_k), d(\omega_k), s_3(\omega_k) | x_t = \lambda_i)$$

which is also estimated with empirical frequencies. If $s_3(\omega_k)$ is unknown in the training set, we use $s_2(\omega_k)$ or otherwise $s_1(\omega_k)$ or finally $s_0(\omega_k)$. This original approximation method allows us to improve the segmentation of unknown words while keeping a very short training time with PMC and HMC. The choice of the different features depends of the language, these ones are selected for English, and one can add any features according to language characteristics.

It is important to note that the choice of the features is crucial as we cannot use arbitrary features with Markov chain models [23] [27] [28]. It preserves the model from a “second level” of unknown patterns, and it avoids having our new approximation of $b_i(k)$ to be equal to 0 too many times.

4 Experiments

To calibrate the performance of our models, we compare their performances to benchmark models with no extra-data: Maximum Entropy Markov Models [27], Recurrent Neural Networks [29] [30], Long-Short-Term-Memory (LSTM) network [31], Gated Recurrent Unit [32], Conditional Random Field (CRF) [33], and the BiLSTM-CRF [34]. We present the results of the best one with appreciable training and execution time, the CRF, having the same features described in section 3, to which we add the suffix of length 4 and prefixes of length 1 to 4. Comparing results to this model is particularly relevant, as it is one of the most applied models for these tasks when no extra-data are used. The experiments are done with python. We code our own Markov Chain models, and we use the CRF Suite [35] library for the CRF models.

Table 1: HMC, CRF and PMC for POS Tagging with error rates, KW stands for Known Words, and UW for Unknown Words

	HMC	CRF	PMC
CoNLL 2000	2.96%	2.47%	2.32%
CoNLL 2000 KW	1.94%	1.72%	1.27%
CoNLL 2000 UW	16.54%	12.47%	16.41%
CoNLL 2003	5.29%	4.28%	4.71%
CoNLL 2003 KW	4.03%	3.21%	3.40%
CoNLL 2003 UW	15.30%	12.79%	15.16%
UD ENGLISH	8.13%	6.95%	7.16%
UD ENGLISH KW	6.07%	5.46%	5.00%
UD ENGLISH UW	30.73%	23.25%	30.78%

Table 2: HMC, CRF and PMC for NER with F_1 scores, KW stands for Known Words, and UW for Unknown Words

	HMC	CRF	PMC
CoNLL 2003	78.44	79.16	79.52
CoNLL 2003 KW	87.07	87.41	88.47
CoNLL 2003 UW	56.67	58.67	56.87

Table 3: HMC, CRF and PMC for Chunking with F_1 scores, KW stands for Known Words, and UW for Unknown Words

	HMC	CRF	PMC
CoNLL 2000	92.72	93.05	94.49
CoNLL 2000 KW	93.18	93.35	95.09
CoNLL 2000 UW	87.45	89.63	87.58
CoNLL 2003	94.30	94.79	95.61
CoNLL 2003 KW	94.65	94.92	96.17
CoNLL 2003 UW	91.95	93.88	91.85

We use reference datasets for every tasks: CoNLL 2000 [36] for Chunking, UD English [37] for POS Tagging, and CoNLL 2003 [38] for NER ². In addition, we take advantage of having enough data to use CoNLL 2000 and CoNLL

²All these datasets are freely available. CoNLL 2003 upon request at <https://www.clips.uantwerpen.be/conll2003/ner/>, UD English is available on <https://universaldependencies.org/#language->, and CoNLL 2000 with the NLTK [26] library with python.

Table 4: Training times of the different models on CPU with 8 Gb of RAM

	HMC	CRF	PMC
CoNLL 2000 POS	1.3s	140s	5s
CoNLL 2000 CHUNK	1.3s	140s	5s
CoNLL 2003 POS	1.5s	180s	5.5s
CoNLL 2003 NER	1.5s	260s	5.5s
CoNLL 2003 CHUNK	1.5s	160s	5.5s
UD ENGLISH POS	1.7s	185s	6.3s

2003 for POS Tagging, and CoNLL 2003 for chunking. We use universal tagset [39] for each experiment of POS Tagging.

The results are presented in Table I for POS Tagging, Table II for NER, and Table III for Chunking. Training times on CPU with 8 Gb of RAM are presented in Table IV. HMC has the lowest training times, but its performances are significantly worse than those of the other models. The difference can achieve 20%. As HMC is a particular case of PMC, these results were expected. They illustrate the interest in using the PMC model rather than the HMC one. The most relevant results are those regarding the comparison between PMC and CRF. They have equivalent scores, with PMC having better results in 4 of the 6 experiments. In general, PMC is better for known words and has worse results for unknown ones. The main difference between the two models rests on the time required to train them. PMC is about 30 times faster to train than CRF, and adding new data does not entail the retraining of models. About execution time, we have the same observation, with PMC about ten times faster. It is our work’s central goal: to have the best and fastest segmentation model as possible with no extra-data, and PMC seems to achieve this objective.

5 Conclusion

PMC has good performances with no extra-data for these tasks, with equivalent results to the CRF. The main advantage of PMC is its training and execution time, much faster for text segmentation. It confirms the interest of using PMC to build an extra-light model for text segmentation, which can then be used for most complex tasks without being restrictive for deployment.

We are conscious that our models are not as competitive as the best deep learning ones, especially for NER. These latter use a large amount of data to construct word embeddings [40] [41] [42], which become the input of models, like BiLSTM-CRF, for example, to achieve state-of-the-art results [43]. On the one hand, using these embeddings is not possible with PMC with the Viterbi or the Forward-Backward algorithms, as this model cannot use observations with arbitrary features. On the other hand, these neural models are heavy and impossible to deploy without a significant configuration.

The training time, the execution time, and the carbon footprint of PMC are significantly lower, which is our project’s main objective. As a perspective, PMC has been extended with the Triplet Markov Chain (TMC) model [44] [45] [46]. We can apply this extension to observe the possible improvements, especially for NER, while keeping Markov models’ relevant properties.

References

- [1] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- [2] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.
- [3] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. Squad: 100,000+ questions for machine comprehension of text. *arXiv preprint arXiv:1606.05250*, 2016.
- [4] Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. Learning Word Vectors for Sentiment Analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150, Portland, Oregon, USA, June 2011. Association for Computational Linguistics.

- [5] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [6] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*, 2019.
- [7] Emma Strubell, Ananya Ganesh, and Andrew McCallum. Energy and policy considerations for deep learning in NLP. *arXiv preprint arXiv:1906.02243*, 2019.
- [8] Ruslan Leont'evich Stratonovich. Conditional Markov processes. In *Non-linear transformations of stochastic processes*, pages 427–453. Elsevier, 1965.
- [9] Leonard E Baum and Ted Petrie. Statistical inference for probabilistic functions of finite state Markov chains. *The annals of mathematical statistics*, 37(6):1554–1563, 1966.
- [10] Olivier Cappé, Eric Moulines, and Tobias Rydén. *Inference in hidden Markov models*. Springer Science & Business Media, 2006.
- [11] Lawrence Rabiner and B Juang. An introduction to hidden Markov models. *IEEE ASSP Magazine*, 3(1):4–16, 1986.
- [12] Lawrence R Rabiner. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.
- [13] Yariv Ephraim and Neri Merhav. Hidden Markov processes. *IEEE Transactions on information theory*, 48(6):1518–1569, 2002.
- [14] Thorsten Brants. TnT: a statistical part-of-speech tagger. In *Proceedings of the sixth conference on Applied natural language processing*, pages 224–231. Association for Computational Linguistics, 2000.
- [15] Sudha Morwal, Nusrat Jahan, and Deepti Chopra. Named entity recognition using hidden markov model (hmm). *International Journal on Natural Language Computing (IJNLC)*, 1(4):15–23, 2012.
- [16] Asif Ekbal, S Mondal, and Sivaji Bandyopadhyay. Pos tagging using hmm and rule-based chunking. *The Proceedings of SPSAL*, 8(1):25–28, 2007.
- [17] Wojciech Pieczynski. Pairwise Markov Chains. *IEEE Trans. Pattern Anal. Mach. Intell.*, 25(5):634–639, May 2003.
- [18] Stéphane Derrode and Wojciech Pieczynski. Signal and image segmentation using pairwise Markov chains. *IEEE Transactions on Signal Processing*, 52(9):2477–2489, Sep. 2004.
- [19] Ivan Gorynin, Hugo Gangloff, Emmanuel Monfrini, and Wojciech Pieczynski. Assessing the segmentation performance of pairwise and triplet Markov models. *Signal Processing*, 145, 12 2017.
- [20] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding. In *Advances in neural information processing systems*, pages 5753–5763, 2019.
- [21] Alan Akbik, Duncan Blythe, and Roland Vollgraf. Contextual String Embeddings for Sequence Labeling. In *COLING 2018, 27th International Conference on Computational Linguistics*, pages 1638–1649, 2018.
- [22] Andrew Viterbi. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE transactions on Information Theory*, 13(2):260–269, 1967.
- [23] Daniel Jurafsky and James H. Martin. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*.
- [24] Tomoko Ohta, Yuka Tateisi, Jin-Dong Kim, Hideki Mima, and Junichi Tsujii. The GENIA corpus: An annotated research abstract corpus in molecular biology domain. In *Proceedings of the second international conference on Human Language Technology Research*, pages 82–86, 2002.
- [25] Leon Derczynski. Complementarity, f-score, and nlp evaluation. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 261–266, 2016.
- [26] Edward Loper and Steven Bird. NLTK: the natural language toolkit. *arXiv preprint cs/0205028*, 2002.
- [27] Andrew McCallum, Dayne Freitag, and Fernando CN Pereira. Maximum Entropy Markov Models for Information Extraction and Segmentation. In *Icml*, volume 17, pages 591–598, 2000.
- [28] Charles Sutton and Andrew McCallum. An introduction to conditional random fields for relational learning. *Introduction to statistical relational learning*, 2:93–128.
- [29] Rafal Jozefowicz, Wojciech Zaremba, and Ilya Sutskever. An empirical exploration of recurrent network architectures. In *International conference on machine learning*, pages 2342–2350, 2015.

- [30] Zachary C Lipton, John Berkowitz, and Charles Elkan. A critical review of recurrent neural networks for sequence learning. *arXiv preprint arXiv:1506.00019*, 2015.
- [31] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [32] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.
- [33] John Lafferty, Andrew McCallum, and Fernando CN Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. 2001.
- [34] Zhiheng Huang, Wei Xu, and Kai Yu. Bidirectional LSTM-CRF models for sequence tagging. *arXiv preprint arXiv:1508.01991*, 2015.
- [35] Naoaki Okazaki. CRFsuite: a fast implementation of Conditional Random Fields (CRFs), 2007.
- [36] Erik F Sang and Sabine Buchholz. Introduction to the CoNLL-2000 shared task: Chunking. *arXiv preprint cs/0009008*, 2000.
- [37] Joakim Nivre, Marie-Catherine De Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajic, Christopher D Manning, Ryan McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, et al. Universal dependencies v1: A multilingual treebank collection. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 1659–1666, 2016.
- [38] Erik F Sang and Fien De Meulder. Introduction to the conll-2003 shared task: Language-independent named entity recognition. *arXiv preprint cs/0306050*, 2003.
- [39] Slav Petrov, Dipanjan Das, and Ryan McDonald. A universal part-of-speech tagset. *arXiv preprint arXiv:1104.2086*, 2011.
- [40] Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543, 2014.
- [41] Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146, 2017.
- [42] Alan Akbik, Tanja Bergmann, and Roland Vollgraf. Pooled contextualized embeddings for named entity recognition. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 724–728, 2019.
- [43] Alan Akbik, Tanja Bergmann, Duncan Blythe, Kashif Rasul, Stefan Schweter, and Roland Vollgraf. FLAIR: An Easy-to-Use Framework for State-of-the-Art NLP. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 54–59, 2019.
- [44] Wojciech Pieczynski. Chames de Markov triplet. *Comptes Rendus Mathématique*, 335(3):275–278, 2002.
- [45] Shou Chen and Xiangqian Jiang. Modeling Repayment Behavior of Consumer Loan in Portfolio across Business Cycle: A Triplet Markov Model Approach. *Complexity*, 2020, 2020.
- [46] Haoyu Li, Stéphane Derrode, and Wojciech Pieczynski. An adaptive and on-line imu-based locomotion activity classification method using a triplet markov model. *Neurocomputing*, 362:94–105, 2019.