

DenseHMM: Learning Hidden Markov Models by Learning Dense Representations

Joachim Sicking^{1,2}, Maximilian Pintz^{1,3}, Maram Akila¹ and Tim Wirtz^{1,2}

¹Fraunhofer IAIS, Sankt Augustin, Germany

²Fraunhofer Center for Machine Learning, Sankt Augustin, Germany

³University of Bonn, Bonn, Germany

Keywords: Hidden Markov Model, Representation Learning, Gradient Descent Optimization.

Abstract: We propose *DenseHMM* – a modification of Hidden Markov Models (HMMs) that allows to learn dense representations of both the hidden states and the (discrete) observables. Compared to the standard HMM, transition probabilities are not atomic but composed of these representations via kernelization. Our approach enables constraint-free and gradient-based optimization. We propose two optimization schemes that make use of this: a modification of the Baum-Welch algorithm and a direct co-occurrence optimization. The latter one is highly scalable and comes empirically without loss of performance compared to standard HMMs. We show that the non-linearity of the kernelization is crucial for the expressiveness of the representations. The properties of the DenseHMM like learned co-occurrences and log-likelihoods are studied empirically on synthetic and biomedical datasets.

1 INTRODUCTION

Hidden Markov Models (Rabiner and Juang, 1986) have been a state-of-the-art approach for modelling sequential data for more than three decades (Hinton et al., 2012). Their success story is proven by a large number of applications ranging from natural language modelling (Chen and Goodman, 1999) and automatic speech recognition (Bahl et al., 1986; Varga and Moore, 1990) over financial services (Bhusari and Patil, 2016) to robotics (Fu et al., 2016). While still being used frequently, many more recent approaches are based on neural networks instead, like feed-forward neural networks (Schmidhuber, 2015), recurrent neural networks (Hochreiter and Schmidhuber, 1997) or spiking neural networks (Tavanaei et al., 2019). However, the recent breakthroughs in the field of neural networks (LeCun et al., 2015; Schmidhuber, 2015; Goodfellow et al., 2016; Minar and Naher, 2018; Deng et al., 2013; Paul et al., 2015; Wu et al., 2020) are accompanied by an almost equally big lack of their theoretical understanding. In contrast, HMMs come with a broad theoretical understanding, for instance of the parameter estimation (Yang et al., 2017), convergence (Wu, 1983; Dempster et al., 1977), consistency (Leroux, 1992) and short-term prediction performance (Sharan et al., 2018), despite of their

non-convex optimization landscape.

Different from HMMs, (neural) representation learning became more prominent only recently (Mikolov et al., 2013b; Pennington et al., 2014). From the very first day following their release, approaches like word2vec or Glove (Mikolov et al., 2013b; Pennington et al., 2014; Mikolov et al., 2013a; Le and Mikolov, 2014) that yield dense representations for state sequences, have significantly emphasized the value of pre-trained representations of discrete sequential data for downstream tasks (Kim, 2014; Wang et al., 2019). Since then those found application not only in language modeling (Mikolov et al., 2013c; Zhang et al., 2016b; Li and Yang, 2018; Almeida and Xexéo, 2019) but also in biology (Asgari and Mofrad, 2015; Zou et al., 2019), graph analysis (Perozzi et al., 2014; Grover and Leskovec, 2016) and even banking (Baldassini and Serrano, 2018). Similar approaches have received an overwhelming attention and became part of the respective state-of-the-art approaches.

Ever since, the quality of representation models increased steadily, driven especially by the natural language community. Recently, so-called transformer networks (Vaswani et al., 2017) were put forward, complex deep architectures that leverage attention mechanisms (Bahdanau et al., 2015; Kim et al., 2017). Their complexity and tremendously large amounts of com-

pute and training data lead again to remarkable improvements on a multitude of natural language processing (NLP) tasks (Devlin et al., 2019).

These developments are particularly driven by the question on *how to optimally embed discrete sequences into a continuous space*. However, many existing approaches identify optimality solely with performance and put less emphasis on aspects like conceptual simplicity and theoretical soundness. Intensified discussions on *well-understood* and therefore *trustworthy* machine learning (Saltzer and Schroeder, 1975; Dwork et al., 2012; Amodei et al., 2016; Gu et al., 2017; Varshney, 2019; Toreini et al., 2020; Brundage et al., 2020) indicate, however, that these latter aspects become more and more crucial or even mandatory for real-world learning systems. This holds especially true when representing biological structures and systems as derived insights may be used in downstream applications, e.g. of medical nature where physical harm can occur.

In light of this, we propose *DenseHMM* – a modification of Hidden Markov Models that allows to learn dense representations of both the hidden states and the discrete observables (Figure 1). Compared to the standard HMM, transition probabilities are not atomic but composed of these representations. Concretely, we contribute

- a parameter-efficient, non-linear matrix factorization for HMMs,
- two competitive approaches to optimize the resulting DenseHMM
- and an empirical study of its performance and properties on diverse datasets.

The rest of the work is organized as follows: first, we present related work on HMM parameter learning and matrix-factorization approaches in section 2. Next, DenseHMM and its optimization schemes are introduced in section 3. We study the effect of its softmax non-linearity and conduct empirical analyses and comparisons with standard HMMs in sections 4 and 5, respectively. A discussion in section 6 concludes our paper.

2 RELATED WORK

HMMs are generative models with Markov properties for sequences of either discrete or continuous observation symbols (Rabiner, 1989). They assume a number of non-observable (hidden) states that drive the dynamics of the generated sequences. If domain expertise allows to interpret these drivers, HMMs can be fully understood. This distinguishes HMMs from

sequence-modelling neural networks like long short-term memory networks (Hochreiter and Schmidhuber, 1997) and temporal convolutional networks (Bai et al., 2018). More recent latent variable models that keep the discrete structure of the latent space make use of, e.g., Indian buffet processes (Griffiths and Ghahramani, 2011). These allow to dynamically adapt the dimension of the latent space dependent on data complexity and thus afford more flexible modelling. While we stay in the HMM model class, we argue that our approach allows to extend or reduce the latent space in a more principled way compared to standard HMMs.

Various approaches exist to learn the parameters of hidden Markov models: A classical one is the Baum-Welch algorithm (Rabiner, 1989) that handles the complexity of the joint likelihood of hidden states and observables by introducing an iterative two-step procedure that makes use of the forward-backward algorithm (Rabiner and Juang, 1986). Another algorithm for (local) likelihood maximization is (Baldi and Chauvin, 1994). The authors of (Huang et al., 2018) study HMM learning on observation co-occurrences instead of observation sequences. Based on moments, i.e. co- and triple-occurrences, bounds on the empirical probabilities can be derived via spectral decomposition (Anandkumar et al., 2012). Approaches from Bayesian data analysis comprise Markov chain Monte Carlo (MCMC) and variational inference (VI). While MCMC can provide more stable and better results (Rybert Sipos, 2016), it traditionally suffers from poor scalability. A more scalable stochastic-gradient MCMC algorithm that tackles mini-batching of sequentially dependent data is (Ma et al., 2017). The same authors propose a stochastic VI (SVI) algorithm (Foti et al., 2014) that shares some technical details with (Ma et al., 2017). SVI for hierarchical Dirichlet process (HDP) HMMs is considered in (Zhang et al., 2016a). For our DenseHMM, we adapt two non-Bayesian procedures: the Baum-Welch algorithm (Rabiner, 1989) and direct co-occurrence optimization (Huang et al., 2018). The latter we optimize, solely for convenience, using a deep learning framework.¹ In (Tran et al., 2016) this idea was carried further, allowing different modifications to the original HMM context. Already earlier, combinations of HMMs and neural networks were employed, for instance for automatic speech recognition (Bouclard et al., 1992; Moon and Hwang, 1997; Trentin and Gori, 1999).

Non-negative matrix factorization (NMF, (Lee and Seung, 1999)) splits a matrix into a pair of low-rank matrices with solely positive components. NMF for HMM learning is used, e.g., in (Lakshminarayanan and Raich, 2010; Cybenko and Crespi, 2011). In

¹See appendix C for details.

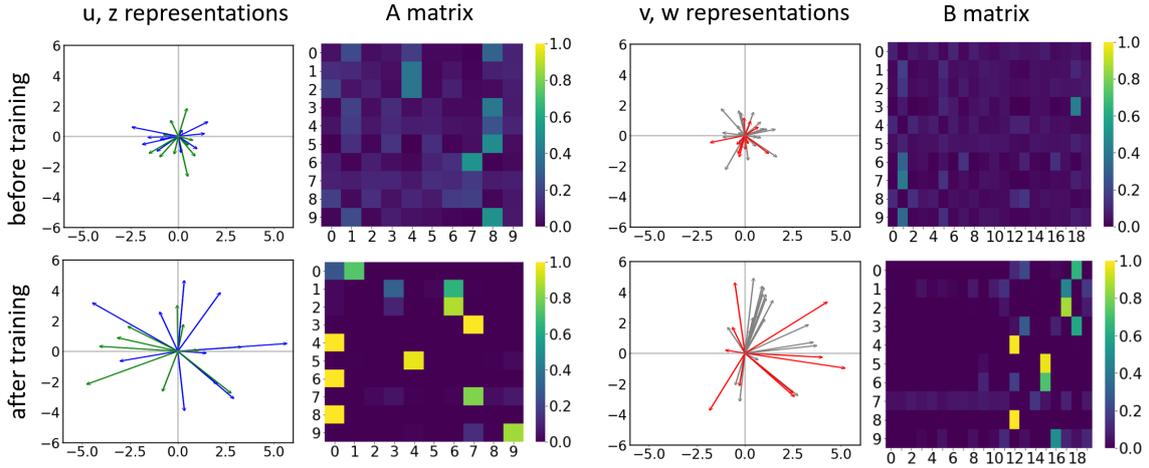


Figure 1: Exemplary DenseHMM to visualize the inner workings of our approach. All model components are shown before (top row) and after training (bottom row). The transition matrices \mathbf{A} (second column) and \mathbf{B} (fourth column) are learned by learning dense representations (first and third column). All representations are initialized by a standard Gaussian.

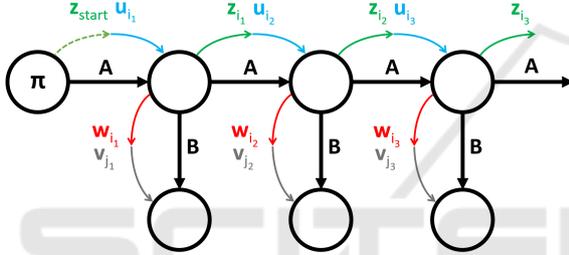


Figure 2: Structure of the DenseHMM. The HMM parameters \mathbf{A}, \mathbf{B} and π are composed of vector representations such that $\mathbf{A} = \mathbf{A}(\mathbf{U}, \mathbf{Z})$, $\mathbf{B} = \mathbf{B}(\mathbf{V}, \mathbf{W})$ and $\pi = \pi(\mathbf{U}, \mathbf{z}_{\text{start}})$.

contrast, we combine matrix factorization with a non-linear kernel function to ensure non-negativity and normalization of the HMM transition matrices. Further foci of recent work on HMMs are identifiability (Huang et al., 2018), i.e. uniqueness guarantees for an obtained model, and optimized priors over transition distributions (Qiao et al., 2015).

3 STRUCTURE AND OPTIMIZATION OF THE DENSE HMM

A HMM is defined by two time-discrete stochastic processes: $\{X_t\}_{t \in \mathbb{N}}$ is a Markov chain over hidden states $S = \{s_i\}_{i=1}^n$ and $\{Y_t\}_{t \in \mathbb{N}}$ is a process over observable states $O = \{o_i\}_{i=1}^m$. The central assumption of HMMs is that the probability to observe $Y_t = y_t$ depends only on the current state of the hidden process and the probability to find $X_t = x_t$ only on the the previous state of the hidden process, $X_{t-1} = x_{t-1}$ for all $t \in \mathbb{N}$. We denote the state-transition matrix as $\mathbf{A} \in \mathbb{R}^{n \times n}$

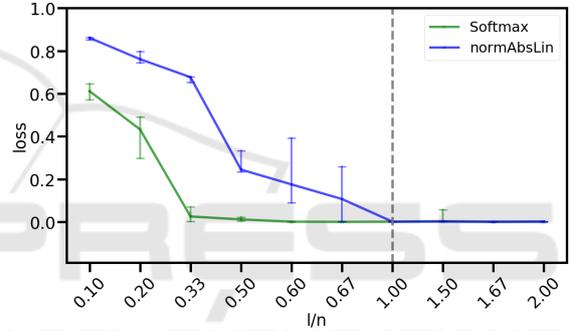


Figure 3: Approximation quality of non-linear matrix factorizations. The optimization errors (median, 25/75 percentile) of softmax (green) and normAbsLin (blue) matrices are shown over the ratio of representation length l and matrix size n . The vertical line indicates $l = n$.

with $a_{ij} = P(X_t = s_j | X_{t-1} = s_i)$, the emission matrix as $\mathbf{B} \in \mathbb{R}^{n \times m}$ with $b_{ij} = P(Y_t = o_j | X_t = s_i)$ and the initial state distribution as $\pi \in \mathbb{R}^n$ with $\pi_i = P(X_1 = s_i)$. A HMM is fully parametrized by $\lambda = (\mathbf{A}, \mathbf{B}, \pi)$.

HMMs can be seen as extensions of Markov chains (MCs) which are in turn closely related to word2vec embeddings. Let us elaborate on this: A MC has no hidden states and is defined by just one process $\{X_t\}_{t \in \mathbb{N}}$ over observables. The transition dynamics of the states of the MC is described by a transition matrix \mathbf{A} and an initial distribution π . Being in a given state s_t , the MC models conditional probabilities of the form $p(s_j | s_t)$. MCs are structurally similar to the approaches that learn word2vec representations, i.e. continuous bag of words and skip-gram (Mikolov et al., 2013b). Both models learn transitions between the words of a text corpus. Each word w_i of the vocabulary is represented by a learned dense vector u_i . The transition probabilities between words are recovered

from the scalar products of these vectors:

$$p(w_j | w_i) = \frac{\exp(\mathbf{u}_i \cdot \mathbf{v}_j)}{\sum_k \exp(\mathbf{u}_i \cdot \mathbf{v}_k)} \propto \exp(\mathbf{u}_i \cdot \mathbf{v}_j). \quad (1)$$

The learned word2vec representations are low-dimensional and context-based, i.e. they contain semantic information. This is in contrast to the trivial and high-dimensional one-hot (or bag-of-words) encodings.

Here we transfer the non-linear factorization approach of word2vec (eq. 1) to HMMs. This is done by composing \mathbf{A} , \mathbf{B} and $\boldsymbol{\pi}$ of dense vector representations such that

$$a_{ij} = a_{ij}(\mathbf{U}, \mathbf{Z}) = \frac{\exp(\mathbf{u}_j \cdot \mathbf{z}_i)}{\sum_{k \in [n]} \exp(\mathbf{u}_k \cdot \mathbf{z}_i)}, \quad (2a)$$

$$b_{ik} = b_{ik}(\mathbf{V}, \mathbf{W}) = \frac{\exp(\mathbf{v}_k \cdot \mathbf{w}_i)}{\sum_{k' \in [m]} \exp(\mathbf{v}_{k'} \cdot \mathbf{w}_i)}, \quad (2b)$$

$$\pi_i = \pi_i(\mathbf{U}, \mathbf{z}_{\text{start}}) = \frac{\exp(\mathbf{u}_i \cdot \mathbf{z}_{\text{start}})}{\sum_{k \in [n]} \exp(\mathbf{u}_k \cdot \mathbf{z}_{\text{start}})}, \quad (2c)$$

for $i, j \in [n]$ and $k \in [m]$. Let us motivate this transformation (Fig. 2) piece by piece: each representation vector corresponds to either a hidden state $(\mathbf{u}_i, \mathbf{w}_i, \mathbf{z}_i)$ or an observation (v_i) . The vector \mathbf{u}_i (\mathbf{z}_i) is the incoming (outgoing) representation of hidden state i along the (hidden) Markov chain. \mathbf{w}_i is the outgoing representation of hidden state i towards the observation symbols. These are described by the \mathbf{v}_i . All vectors are real-valued and of length l . \mathbf{A} and \mathbf{B} each depend on two kinds of representations instead of only one to enable non-symmetric transition matrices. Additionally, to choose \mathbf{A} independent of \mathbf{B} , as is typical for HMMs, we need \mathbf{w}_i as a third hidden representation. It is convenient to summarize all representation vectors of one kind in a matrix $(\mathbf{U}, \mathbf{V}, \mathbf{W}, \mathbf{Z})$.

A softmax kernel maps the scalar products of the representations onto the HMM parameters \mathbf{A} , \mathbf{B} and $\boldsymbol{\pi}$. Softmax maps to the simplex and thus ensures a_{ij} , b_{ij} , π_i to be in $[0, 1]$ as well as row-wise normalization of \mathbf{A} , \mathbf{B} and $\boldsymbol{\pi}$. While a different kernel function may be chosen, a strong non-linearity, such as in the softmax kernel, is essential to obtain matrices \mathbf{A} and \mathbf{B} with high ranks (compare experiments in section 4).

This non-linear kernelization enables constraint-free optimization which is a central property of our approach. We use this fact in two different ways: we derive a modified expectation-maximization (EM) scheme in section 3.1 and study an alternative to EM optimization that is based on co-occurrences in section 3.2.

3.1 EM Optimization: a Gradient-based M-step

We briefly recapitulate the EM-based Baum-Welch algorithm (Bishop, 2007) and adapt it to learn the proposed representations as part of the M-step:

Given a sequence \mathbf{o} of length $T \in \mathbb{N}$ over observations O , the Baum-Welch algorithm finds parameters $\boldsymbol{\lambda}$ that (locally) maximize the observation likelihoods. A latent distribution Q over the hidden states S is introduced such that the log-likelihood of the sequence decomposes as follows: $\mathcal{L}(Q, \boldsymbol{\lambda}) = \log P(\mathbf{o}, \boldsymbol{\lambda}) = \mathcal{L}(Q, \boldsymbol{\lambda}) + KL(Q || P(\cdot | \mathbf{o}, \boldsymbol{\lambda}))$. $KL(P || Q)$ denotes the Kullback-Leibler divergence from Q to P with P, Q being probability distributions and $\mathcal{L}(Q, \boldsymbol{\lambda}) = \sum_{\mathbf{x} \in S^T} Q(\mathbf{x}) \log [P(\mathbf{x}, \boldsymbol{\lambda}) / Q(\mathbf{x})]$. Starting from an initial guess for $\boldsymbol{\lambda}$, the algorithm alternates between two sub-procedures, the E- and M-step: In the E-step, the forward-backward algorithm (Bishop, 2007) is used to update $Q = P(\cdot | \mathbf{o}, \boldsymbol{\lambda})$, which maximizes $\mathcal{L}(Q, \boldsymbol{\lambda})$ for fixed $\boldsymbol{\lambda}$. The efficient computation of the conditional probabilities $\gamma_t(s, s') := P(X_{t-1} = s, X_t = s' | \mathbf{o})$ and $\gamma_t(s) := P(X_t = s | \mathbf{o})$ for $s, s' \in S$ is crucial for the E-step. In the M-step, the latent distribution Q is fixed and $\mathcal{L}(Q, \boldsymbol{\lambda})$ is maximized w.r.t. $\boldsymbol{\lambda}$ under normalization constraints. As the Kullback-Leibler divergence KL is set to zero in each E-step, the function to maximize in the M-step becomes

$$\begin{aligned} \mathcal{L}(Q, \boldsymbol{\lambda}) &= \sum_{\mathbf{x} \in S^T} Q(\mathbf{x}) \log \frac{P(\mathbf{x}, \boldsymbol{\lambda})}{Q(\mathbf{x})} \\ &= \sum_{\mathbf{x} \in S^T} P(\mathbf{x} | \boldsymbol{\lambda}^{\text{old}}) \log \frac{P(\mathbf{x}, \boldsymbol{\lambda})}{P(\mathbf{x} | \boldsymbol{\lambda}^{\text{old}})} \end{aligned}$$

with $\boldsymbol{\lambda}^{\text{old}}$ being the parameter obtained in the previous M-step. Applying the logarithm to $P(\mathbf{x}, \boldsymbol{\lambda})$, which has a product structure due to the Markov properties, splits the optimization objective into three summands. Each term depends on only one of the parameters \mathbf{A} , \mathbf{B} , $\boldsymbol{\pi}$:

$$\begin{aligned} \mathbf{A}^*, \mathbf{B}^*, \boldsymbol{\pi}^* &= \arg \max_{\mathbf{A}, \mathbf{B}, \boldsymbol{\pi}} \mathcal{L}(Q, \boldsymbol{\lambda}) \\ &= \arg \max_{\mathbf{A}, \mathbf{B}, \boldsymbol{\pi}} \mathcal{L}_1(Q, \mathbf{A}) + \mathcal{L}_2(Q, \mathbf{B}) + \mathcal{L}_3(Q, \boldsymbol{\pi}). \end{aligned}$$

Due to structural similarities between the three summands, we consider only \mathcal{L}_1 in the following. The treatment of \mathcal{L}_2 and \mathcal{L}_3 can be found in appendix A. For \mathcal{L}_1 , we have

$$\mathcal{L}_1(Q, \mathbf{A}) = \sum_{\mathbf{i} \in [n]^T} P(s_{\mathbf{i}} | \boldsymbol{\lambda}^{\text{old}}) \sum_{t=2}^T \log(a_{i_{t-1}, i_t})$$

with multi-index $s_{\mathbf{i}} = s_{i_1}, \dots, s_{i_T}$. The next step is to rewrite \mathcal{L}_1 in terms of γ_t and to use Lagrange multipliers

to ensure normalization. This gives the following part $\tilde{\mathcal{L}}_1$ of the full Lagrangian $\tilde{\mathcal{L}} = \tilde{\mathcal{L}}_1 + \tilde{\mathcal{L}}_2 + \tilde{\mathcal{L}}_3$:

$$\tilde{\mathcal{L}}_1 := \sum_{i,j \in [n]} \sum_{t=2}^T \gamma_t(s_i, s_j) \log a_{ij} + \sum_{i \in [n]} \varphi_i \left(1 - \sum_{j \in [n]} a_{ij} \right)$$

with Lagrange multipliers φ_i for each $i \in [n]$ to ensure that \mathbf{A} is a proper transition matrix.

To optimize DenseHMM, we leave the E-step unchanged and modify the M-step by applying the parametrization of $\boldsymbol{\lambda}$, i.e. equations (2a)-(2c), to the Lagrangian $\tilde{\mathcal{L}}$. Please note that we can drop all normalization constraints as they are explicitly enforced by the softmax function. This turns the original constrained optimization problem of the M-step into an unconstrained one, leading to a Lagrangian of the form $\tilde{\mathcal{L}}^{\text{dense}} = \tilde{\mathcal{L}}_1^{\text{dense}} + \tilde{\mathcal{L}}_2^{\text{dense}} + \tilde{\mathcal{L}}_3^{\text{dense}}$ with

$$\begin{aligned} \tilde{\mathcal{L}}_1^{\text{dense}} &= \sum_{i,j \in [n]} \sum_{t=2}^T \gamma_t(s_i, s_j) \mathbf{u}_j \cdot \mathbf{z}_i \\ &\quad - \sum_{i,j \in [n]} \sum_{t=2}^T \gamma_t(s_i, s_j) \log \sum_{k \in [n]} \exp(\mathbf{u}_k \cdot \mathbf{z}_i). \end{aligned}$$

We optimize $\tilde{\mathcal{L}}^{\text{dense}}$ with gradient-decent procedures such as SGD (Bottou, 2010) and Adam (Kingma and Ba, 2015).

3.2 Direct Optimization of Observation Co-occurrences: Gradient-based and Scalable

Inspired by (Huang et al., 2018), we investigate an alternative to the EM scheme: directly optimizing co-occurrence probabilities. The ground truth co-occurrences Ω^{gt} are obtained from training data \mathbf{o} by calculating the relative frequencies of subsequent pairs $(o_i(t), o_j(t+1)) \in \mathcal{O}^2$. If we know that \mathbf{o} is generated by a HMM with a stationary hidden process, we can easily compute Ω^{gt} analytically as follows: we summarize all co-occurrence probabilities $\Omega_{ij} = P(Y_t = o_i, Y_{t+1} = o_j)$ for $i, j \in [m]$ in a co-occurrence matrix $\Omega = \mathbf{B}^T \Theta \mathbf{B}$ with $\Theta_{kl} = P(X_t = s_k, X_{t+1} = s_l)$ for $k, l \in [n]$. We can further write $\Theta_{kl} = P(X_{t+1} = s_l | X_t = s_k) P(X_t = s_k) = A_{kl} \pi_k$ for $i, j \in [n]$ under the assumption that $\boldsymbol{\pi}$ is the stationary distribution of \mathbf{A} , i.e. $\pi_j = \sum_i A_{ij} \pi_i$ for all $i, j \in [n]$. Then, we obtain the co-occurrence probabilities

$$\Omega_{ij} = \sum_{k,l \in [n]} \pi_k b_{ki} a_{kl} b_{lj} \quad \text{for } i, j \in [m]. \quad (3)$$

Parametrizing the matrices \mathbf{A} and \mathbf{B} according to eq. (2a,2b) yields

$$\begin{aligned} &\Omega_{ij}^{\text{dense}}(\mathbf{U}, \mathbf{V}, \mathbf{W}, \mathbf{Z}) \\ &= \sum_{k,l \in [n]} \pi_k b_{ki}(\mathbf{V}, \mathbf{W}) a_{kl}(\mathbf{U}, \mathbf{Z}) b_{lj}(\mathbf{V}, \mathbf{W}) \end{aligned}$$

for $i, j \in [m]$. Please note that $\boldsymbol{\pi}$ is not parametrized here. Following our stationarity demand it is chosen as the eigenvector $\mathbf{v}_{\lambda=1}$ of \mathbf{A}^T . We minimize the squared distance between Ω^{dense} and Ω^{gt} w.r.t. the vector representations, i.e.

$$\arg \min_{\mathbf{U}, \mathbf{V}, \mathbf{W}, \mathbf{Z}} \|\Omega^{\text{gt}} - \Omega^{\text{dense}}(\mathbf{U}, \mathbf{V}, \mathbf{W}, \mathbf{Z})\|_F^2,$$

using gradient-decent procedures like SGD and Adam.

4 PROPERTIES OF THE DENSE HMM

To further motivate our approach, please note that a standard HMM with n hidden states and m observation symbols has $n^2 + n(m-1) - 1$ degrees of freedom (DOFs), whereas a DenseHMM with representation length l has $l(3n+m+1)$ DOFs. Therefore, a low-dimensional representation length l leads to DenseHMMs with less DOFs compared to a standard HMM for many values of n and m . A linear factorization with representation length $l < n$ leads to rank l for the matrices \mathbf{A} and \mathbf{B} , whereas a non-linear factorization can yield more expressive full rank matrices. This effect of non-linearities may be best understood with a simple toy example: assume a 2x2 matrix with co-linear columns: $[[1, 2], [2, 4]]$. Applying a softmax column-wise leads to a matrix $\propto [[e, e^2], C[e^2, e^4]]$ with linearly independent columns. More general, the softmax rescales and rotates each column of $\mathbf{U}\mathbf{Z}$ and $\mathbf{V}\mathbf{W}$ differently and (except for special cases) thus increases the matrix rank to full rank. It is worth to mention that any other kernel $k: \mathbb{R}^l \times \mathbb{R}^l \rightarrow \mathbb{R}^+$ could be used instead of \exp in the softmax to recover the HMM parameters from the representations, e.g. sigmoid, ReLU or RBFs.

As non-linear matrix factorization is a central building block of our approach, we compare the approximation quality of softmax with an appropriate linear factorization in the following setup: we generate a Dirichlet-distributed ground truth matrix $\mathbf{A}_{\text{gt}} \in \mathbb{R}^{n \times n}$ and approximate it (i) by $\tilde{\mathbf{A}} = \text{softmax}(\mathbf{U}\mathbf{Z})$ defined by $\text{softmax}(\mathbf{U}\mathbf{Z})_{ij} = \exp((\mathbf{U}\mathbf{Z})_{ij}) / \sum_k \exp((\mathbf{U}\mathbf{Z})_{ik})$ and (ii) by a normalized absolute matrix product $\tilde{\mathbf{A}} = \text{normAbsLin}(\mathbf{U}\mathbf{Z})$ defined by $\text{normAbsLin}(\mathbf{U}\mathbf{Z})_{ij} = |(\mathbf{U}\mathbf{Z})_{ij}| / \sum_k |(\mathbf{U}\mathbf{Z})_{ik}|$. Note that we report the resulting error $\|\tilde{\mathbf{A}} - \mathbf{A}_{\text{gt}}\|_F$ divided by $\|\mathbf{A}_{\text{gt}}\|_F$ to get comparable losses independent of the size of \mathbf{A}_{gt} . These optimizations are performed for matrix sizes $n = 3, 5, 10$, several representation lengths l and 10 different \mathbf{A}_{gt} for each (n, l) pair. Table 1 in appendix A provides all considered (n, l) pairs and detailed results. Fig. 3 shows that the softmax non-linearity yields closer approximations of

\mathbf{A}_{gt} compared to normAbsLin. Moreover, we observe on a qualitative level significantly faster convergence for softmax as l increases. For softmax, vector lengths $l \approx n/3$ suffice to closely fit \mathbf{A}_{gt} while the piece-wise linear normAbsLin requires $l = n$. This result is in accordance with our remarks above.

5 EMPIRICAL EVALUATION

We investigate the outlined optimization schemes w.r.t. obtained model quality and behaviors. We compare the following types of models:

- $\mathcal{H}_{\text{dense}}^{\text{EM}}$: a DenseHMM optimized with the EM optimization scheme (section 3.1),
- $\mathcal{H}_{\text{dense}}^{\text{direct}}$: a DenseHMM optimized with direct optimization of co-occurrences (section 3.2),
- $\mathcal{H}_{\text{stand}}$: a standard HMM optimized with the Baum-Welch algorithm (Rabiner, 1989).

These models all have the same number of hidden states n and observation symbols m . If a standard HMM and a DenseHMM use the same n and m , one of the models may have less DOFs than the other (cp. section 4). Therefore, we also consider the model $\mathcal{H}_{\text{stand}}^{\text{fair}}$, which is a standard HMM with a similar amount of DOFs compared to a given $\mathcal{H}_{\text{dense}}$ model. We denote the number of hidden states in $\mathcal{H}_{\text{stand}}^{\text{fair}}$ as n_{fair} , which is the positive solution of $n_{\text{fair}}^2 + n_{\text{fair}}(m-1) - 1 = l(3n+m+1)$ rounded to the nearest integer. Note that n_{fair} can get significantly larger than n for $l > n$ and therefore $\mathcal{H}_{\text{stand}}^{\text{fair}}$ is expected to outperform the other models in these cases.

We use two standard measures to assess model quality: the co-occurrence mean absolute deviation (MAD) and the normalized negative log-likelihood (NLL). The MAD between two co-occurrence matrices Ω^{gt} and Ω^{model} is defined as $1/m^2 \sum_{i,j \in [m]} |\Omega_{ij}^{\text{model}} - \Omega_{ij}^{\text{gt}}|$. We compute both Ω^{gt} and Ω^{model} based on sufficiently long sampled sequences (more details in appendix C). In the case of synthetically generated ground truth sequences, we compute Ω^{gt} analytically instead. In addition, we take a look at the negative log-likelihood of the ground truth test sequences $\{\mathbf{o}_i^{\text{test}}\}$ under the model, i.e. $\text{NLL} = -\sum_i \log P(\mathbf{o}_i^{\text{test}}; \lambda)$. We conduct experiments with $n \in \{3, 5, 10\}$ and different representation lengths l for each n . For each (n, l) combination, we run 10 experiments with different train-test splits. We evaluate the median and 25/75 percentiles of the co-occurrence MADs and the normalized NLLs for each of the four models (see appendix C for details).

In the following, we consider synthetically generated data as well as two real-world datasets: amino acid sequences from the RCSB PDB dataset (Berman et al., 2000) and part-of-speech tag sequences of biomedical text (Smith et al., 2004), referred to as the MedPost dataset.

Synthetic Sequences. We sample training and test ground truth sequences from a standard HMM \mathcal{H}_{syn} that is constructed as follows: Each row of the transition matrices \mathbf{A} and \mathbf{B} is drawn from a Dirichlet distribution $\text{Dir}(\alpha)$, where all entries in α are set to a fixed value $\alpha = 0.1$. The initial state distribution π is set to the normalized eigenvector $\mathbf{v}_{\lambda=1}$ of \mathbf{A}^T . This renders \mathcal{H}_{syn} stationary and allows a simple analytical calculation of Ω^{gt} according to eq. 3. For both training and testing, we sample 10 sequences, each of length 200 with $m = n$ emission symbols. Figure 4 left shows our evaluation w.r.t. co-occurrence MADs. Note that the performance of $\mathcal{H}_{\text{stand}}$ changes slightly with l for fixed n as training is performed on different sequences for every (n, l) pair. This is because the sequences are re-drawn from \mathcal{H}_{syn} for every experiment. The standard HMMs and $\mathcal{H}_{\text{dense}}^{\text{direct}}$ perform similarly, with $\mathcal{H}_{\text{dense}}^{\text{direct}}$ performing slightly better throughout the experiments and especially for $n = 3$. $\mathcal{H}_{\text{dense}}^{\text{EM}}$ shows a higher MAD than the other models. The good performance of $\mathcal{H}_{\text{dense}}^{\text{direct}}$ may be explained by the fact that it optimizes a function similar to co-occurrence MADs, whereas the other models aim to optimize negative log-likelihoods. The results in Figure 4 right show that the DenseHMMs reach comparable NLLs, although the standard HMMs perform slightly better in this metric.

Proteins. The RCSB PDB dataset (Berman et al., 2000) consists of 512,145 amino acid sequences from which we only take the first 1,024. After applying preprocessing (described in further detail in appendix C), we randomly shuffle the sequences and split train and test data 50:50 for each experiment. Figure 5 left shows the results of our evaluation w.r.t. co-occurrence MADs. $\mathcal{H}_{\text{dense}}^{\text{EM}}$ performs slightly worse than both $\mathcal{H}_{\text{stand}}$ and $\mathcal{H}_{\text{stand}}^{\text{fair}}$. We observe that $\mathcal{H}_{\text{dense}}^{\text{direct}}$ yields the best results. While the co-occurrence MADs of $\mathcal{H}_{\text{dense}}^{\text{EM}}$, $\mathcal{H}_{\text{stand}}$ and $\mathcal{H}_{\text{stand}}^{\text{fair}}$ stay roughly constant throughout different experiments, $\mathcal{H}_{\text{dense}}^{\text{direct}}$ can utilize larger n and l to further decrease co-occurrence MADs. As can be seen in Figure 5 right, all models achieve almost identical normalized NLLs throughout the experiments. The results suggest that model size has only a minor impact on normalized NLL performance for the protein dataset.

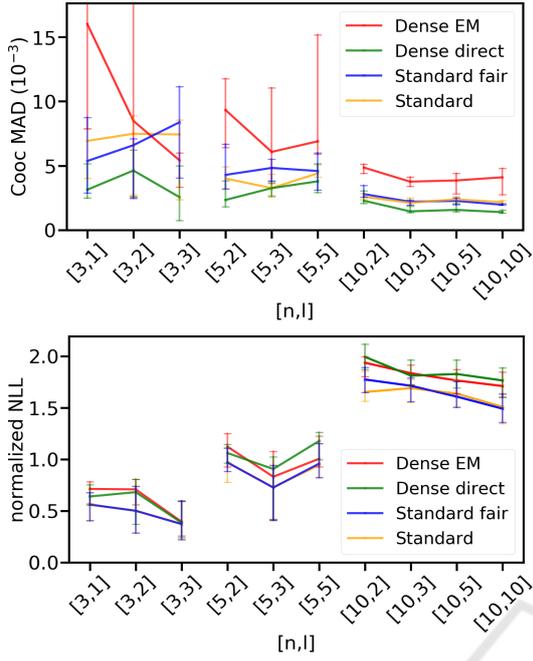


Figure 4: Co-occurrence mean absolute deviation (left) and normalized negative log-likelihood (right) of the models $\mathcal{H}_{\text{dense}}^{\text{EM}}$, $\mathcal{H}_{\text{dense}}^{\text{direct}}$, $\mathcal{H}_{\text{stand}}^{\text{fair}}$, $\mathcal{H}_{\text{stand}}$ on synthetically generated sequences evaluated for multiple combinations of n and l .

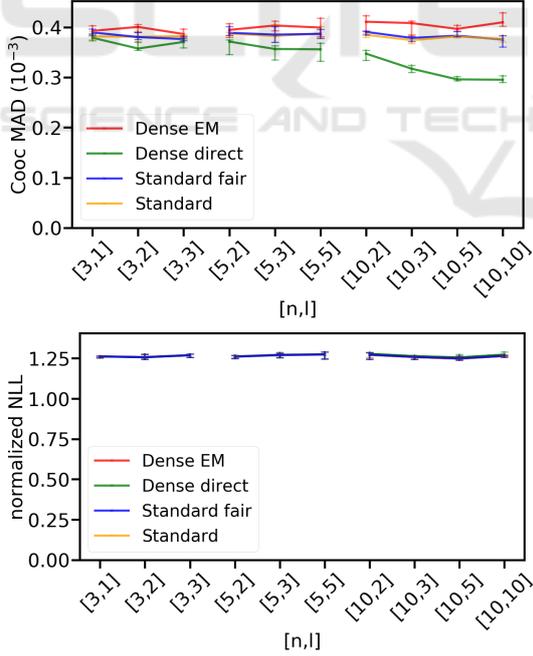


Figure 5: Co-occurrence mean absolute deviation (left) and normalized negative log-likelihood (right) of the models $\mathcal{H}_{\text{dense}}^{\text{EM}}$, $\mathcal{H}_{\text{dense}}^{\text{direct}}$, $\mathcal{H}_{\text{stand}}^{\text{fair}}$, $\mathcal{H}_{\text{stand}}$ on amino acid sequences evaluated for multiple combinations of n and l .

Part-of-Speech Sequences. The MedPost dataset (Smith et al., 2004) consists of 5,700 sentences. Each

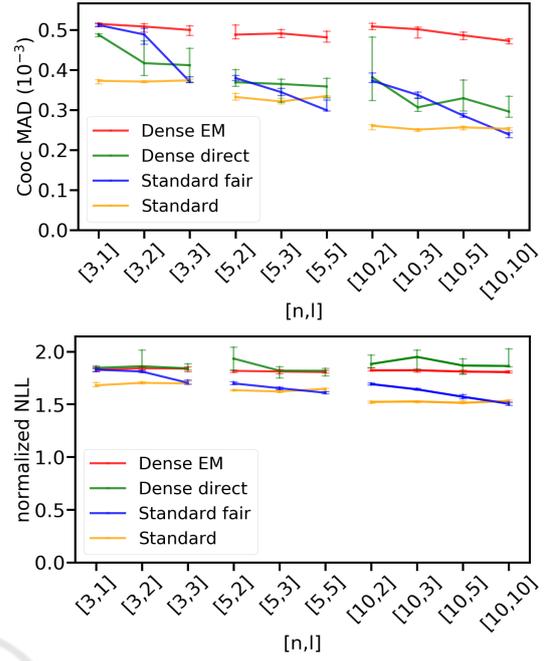


Figure 6: Co-occurrence mean absolute deviation (left) and normalized negative log-likelihood (right) of the models $\mathcal{H}_{\text{dense}}^{\text{EM}}$, $\mathcal{H}_{\text{dense}}^{\text{direct}}$, $\mathcal{H}_{\text{stand}}^{\text{fair}}$, $\mathcal{H}_{\text{stand}}$ on part-of-speech tag sequences (Medpost) evaluated for multiple combinations of n and l .

sentence consists of words that are tagged with one of 60 part-of-speech items. Sequences of part-of-speech tags are considered such that each sequence corresponds to one sentence. We apply preprocessing similar to the protein dataset (more details in appendix C). The sequences are randomly shuffled and train and test data is split 50:50 for each experiment. Figure 6 left shows performance of our models in terms of co-occurrence MADs. We see that the performance of the standard HMM models as well as $\mathcal{H}_{\text{dense}}^{\text{direct}}$ increases with increasing n . The number of hidden states seems to be a major driver of performance. Accordingly, $\mathcal{H}_{\text{stand}}^{\text{fair}}$ improves with growing $n_{\text{fair}}(l) \propto l$. Plus, we have $n_{\text{fair}} > n$ for $l \approx n$ which fully explains why $\mathcal{H}_{\text{stand}}^{\text{fair}}$ is the best performing model in these cases. Overall, $\mathcal{H}_{\text{dense}}^{\text{direct}}$ performs competitive to $\mathcal{H}_{\text{stand}}^{\text{fair}}$, especially for the practically more relevant cases with $l < n$. Similar to the other datasets, $\mathcal{H}_{\text{dense}}^{\text{EM}}$ performs worse than the other models and is barely affected by increasing n . Both DenseHMM models have increasing performance for increasing l . Normalized NLLs (Figure 6 right) are best for the standard HMM models. Both DenseHMM models achieve similar normalized NLLs, which are slightly worse than the ones achieved by the standard HMM models.

6 DISCUSSION

We learn hidden Markov models by learning dense, real-valued vector representations for its hidden and observable states. The involved softmax non-linearity enables to learn high-rank transition matrices, i.e. prevents that the matrix ranks are immediately determined by the chosen (in most cases) low-dimensional representation length. We successfully optimize our models in two different ways and find direct co-occurrence optimization to yield competitive results compared to the standard HMM. This optimization technique requires only one gradient descent procedure and no iterative multi-step schemes. It is highly scalable with training data size and also with model size - as it is implemented in a modern deep-learning framework. The optimization is stable and does neither require fine-tuning of learning rate nor of the representation initializations. We release our full tensorflow code to foster active use of DenseHMM in the community.

We leave it to future work to adapt DenseHMMs to HMMs with continuous emissions and study variants of DenseHMM with fewer kinds of learnable representations. First experiments with DenseHMMs that learn only Z and V lead to almost comparable model quality. From a practitioner's viewpoint, it is worth to investigate how DenseHMM and the learned representations perform on downstream tasks. Using the MedPost dataset, one could consider part-of-speech labeling of word sequences via the Viterbi algorithm after identifying the hidden states of the model with a pre-defined set of ground truth tags. For the protein dataset, a comparison with LSTM-based and BERT embeddings (Bepler and Berger, 2019; Min et al., 2019) could help to understand similarities and differences resulting from modern representation learning techniques. An analysis of geometrical properties of the learned representations seems promising for systems with $l \gg 1$ as $\exp(l)$ many vectors can be almost-orthogonal in \mathbb{R}^l . The V representations are a natural choice for such a study as they directly correspond to observation items. The integration of Bayesian optimization techniques like MCMC and VI with DenseHMM is another research avenue.

ACKNOWLEDGEMENTS

The research of J. Sicking and T. Wirtz was funded by the Fraunhofer Center for Machine Learning within the Fraunhofer Cluster for Cognitive Internet Technologies. The work of M. Pintz and M. Akila was funded by the German Federal Ministry of Education and Research, ML2R - no. 01S18038B. The authors

thank Jasmin Brandt for fruitful discussions.

REFERENCES

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al. (2016). Tensorflow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation*, pages 265–283.
- Almeida, F. and Xexéo, G. (2019). Word embeddings: A survey. *arXiv preprint arXiv:1901.09069*.
- Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., and Mané, D. (2016). Concrete problems in AI safety. *arXiv preprint arXiv:1606.06565*.
- Anandkumar, A., Hsu, D. J., and Kakade, S. M. (2012). A method of moments for mixture models and hidden Markov models. In Mannor, S., Srebro, N., and Williamson, R. C., editors, *COLT 2012 - The 25th Annual Conference on Learning Theory, June 25-27, 2012, Edinburgh, Scotland*, volume 23 of *JMLR Proceedings*, pages 33.1–33.34. JMLR.org.
- Asgari, E. and Mofrad, M. R. (2015). Continuous distributed representation of biological sequences for deep proteomics and genomics. *PLoS One*, 10(11):e0141287.
- Bahdanau, D., Cho, K., and Bengio, Y. (2015). Neural machine translation by jointly learning to align and translate. In Bengio, Y. and LeCun, Y., editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Bahl, L., Brown, P., De Souza, P., and Mercer, R. (1986). Maximum mutual information estimation of hidden Markov model parameters for speech recognition. In *ICASSP'86. IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 11, pages 49–52. IEEE.
- Bai, S., Kolter, J. Z., and Koltun, V. (2018). An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271*.
- Baldassini, L. and Serrano, J. A. R. (2018). client2vec: towards systematic baselines for banking applications. *arXiv preprint arXiv:1802.04198*.
- Baldi, P. and Chauvin, Y. (1994). Smooth on-line learning algorithms for hidden Markov models. *Neural Computation*, 6(2):307–318.
- Bepler, T. and Berger, B. (2019). Learning protein sequence embeddings using information from structure. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*.
- Berman, H. M., Westbrook, J., Feng, Z., Gilliland, G., Bhat, T. N., Weissig, H., Shindyalov, I. N., and Bourne, P. E. (2000). The protein data bank. *Nucleic acids research*, 28(1):235–242.
- Bhusari, V. and Patil, S. (2016). Study of hidden Markov model in credit card fraudulent detection. In *2016 World Conference on Futuristic Trends in Research and Innovation for Social Welfare (Startup Conclave)*,

- pages 1–4. IEEE.
- Bishop, C. M. (2007). *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer, 1 edition.
- Bottou, L. (2010). Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010*, pages 177–186. Springer.
- Bourlard, H., Morgan, N., and Renals, S. (1992). Neural nets and hidden Markov models: Review and generalizations. *Speech Communication*, 11(2-3):237–246.
- Brundage, M., Avin, S., Wang, J., Belfield, H., Krueger, G., Hadfield, G., Khlaaf, H., Yang, J., Toner, H., Fong, R., et al. (2020). Toward trustworthy AI development: Mechanisms for supporting verifiable claims. *arXiv preprint arXiv:2004.07213*.
- Chen, S. F. and Goodman, J. (1999). An empirical study of smoothing techniques for language modeling. *Computer Speech & Language*, 13(4):359–394.
- Cybenko, G. and Crespi, V. (2011). Learning hidden Markov models using nonnegative matrix factorization. *IEEE Transactions on Information Theory*, 57(6):3963–3970.
- Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1):1–22.
- Deng, L., Li, J., Huang, J.-T., Yao, K., Yu, D., Seide, F., Seltzer, M., Zweig, G., He, X., Williams, J., et al. (2013). Recent advances in deep learning for speech research at Microsoft. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 8604–8608. IEEE.
- Devlin, J., Chang, M., Lee, K., and Toutanova, K. (2019). BERT: pre-training of deep bidirectional transformers for language understanding. In Burstein, J., Doran, C., and Solorio, T., editors, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Dwork, C., Hardt, M., Pitassi, T., Reingold, O., and Zemel, R. (2012). Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, pages 214–226.
- Foti, N., Xu, J., Laird, D., and Fox, E. (2014). Stochastic variational inference for hidden Markov models. In *Advances in Neural Information Processing Systems*, pages 3599–3607.
- Fu, R., Wang, H., and Zhao, W. (2016). Dynamic driver fatigue detection using hidden Markov model in real driving condition. *Expert Systems with Applications*, 63:397–411.
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep learning*. MIT press.
- Griffiths, T. L. and Ghahramani, Z. (2011). The Indian buffet process: An introduction and review. *Journal of Machine Learning Research*, 12(Apr):1185–1224.
- Grover, A. and Leskovec, J. (2016). node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 855–864. ACM.
- Gu, T., Dolan-Gavitt, B., and Garg, S. (2017). Badnets: Identifying vulnerabilities in the machine learning model supply chain. *arXiv preprint arXiv:1708.06733*.
- Hinton, G., Deng, L., Yu, D., Dahl, G. E., Mohamed, A.-r., Jaitly, N., Senior, A., Vanhoucke, V., Nguyen, P., Sainath, T. N., et al. (2012). Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal Processing Magazine*, 29(6):82–97.
- Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8):1735–1780.
- Huang, K., Fu, X., and Sidiropoulos, N. (2018). Learning hidden Markov models from pairwise co-occurrences with application to topic modeling. In *Proceedings of the 35th International Conference on Machine Learning*, pages 2068–2077. PMLR.
- Kim, Y. (2014). Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1746–1751, Doha, Qatar. Association for Computational Linguistics.
- Kim, Y., Denton, C., Hoang, L., and Rush, A. M. (2017). Structured attention networks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Kingma, D. P. and Ba, J. (2015). Adam: A method for stochastic optimization. In Bengio, Y. and LeCun, Y., editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Lakshminarayanan, B. and Raich, R. (2010). Non-negative matrix factorization for parameter estimation in hidden Markov models. In *2010 IEEE International Workshop on Machine Learning for Signal Processing*, pages 89–94. IEEE.
- Le, Q. and Mikolov, T. (2014). Distributed representations of sentences and documents. In *International Conference on Machine Learning*, pages 1188–1196.
- LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature*, 521(7553):436–444.
- Lee, D. D. and Seung, H. S. (1999). Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791.
- Leroux, B. G. (1992). Maximum-likelihood estimation for hidden Markov models. *Stochastic Processes and their Applications*, 40(1):127–143.
- Li, Y. and Yang, T. (2018). Word embedding for understanding natural language: a survey. In *Guide to Big Data Applications*, pages 83–104. Springer.
- Ma, Y.-A., Foti, N. J., and Fox, E. B. (2017). Stochastic gradient MCMC methods for hidden Markov models. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2265–2274. JMLR. org.

- Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013a). Efficient estimation of word representations in vector space. In Bengio, Y. and LeCun, Y., editors, *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. (2013b). Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems*, pages 3111–3119.
- Mikolov, T., Yih, W.-t., and Zweig, G. (2013c). Linguistic regularities in continuous space word representations. In *Proceedings of the 2013 Conference of the North American Chapter of the Association For Computational Linguistics: Human Language Technologies*, pages 746–751.
- Min, S., Park, S., Kim, S., Choi, H.-S., and Yoon, S. (2019). Pre-training of deep bidirectional protein sequence representations with structural information. *Neural Information Processing Systems, Workshop on Learning Meaningful Representations of Life*.
- Minar, M. R. and Naher, J. (2018). Recent advances in deep learning: An overview. *arXiv preprint arXiv:1807.08169*.
- Moon, S. and Hwang, J.-N. (1997). Robust speech recognition based on joint model and feature space optimization of hidden Markov models. *IEEE Transactions on Neural Networks*, 8(2):194–204.
- Paul, S., Singh, L., et al. (2015). A review on advances in deep learning. In *2015 IEEE Workshop on Computational Intelligence: Theories, Applications and Future Directions (WCI)*, pages 1–6. IEEE.
- Pennington, J., Socher, R., and Manning, C. (2014). Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.
- Perozzi, B., Al-Rfou, R., and Skiena, S. (2014). Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 701–710. ACM.
- Qiao, M., Bian, W., Da Xu, R. Y., and Tao, D. (2015). Diversified hidden Markov models for sequential labeling. *IEEE Transactions on Knowledge and Data Engineering*, 27(11):2947–2960.
- Rabiner, L. R. (1989). A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286.
- Rabiner, L. R. and Juang, B.-H. (1986). An introduction to hidden Markov models. *IEEE ASSP Magazine*, 3(1):4–16.
- Rybert Sipos, I. (2016). Parallel stratified MCMC sampling of AR-HMMs for stochastic time series prediction. In *Proceedings of the 4th Stochastic Modeling Techniques and Data Analysis International Conference with Demographics Workshop (SMTDA 2016)*. Valletta, Malta: University of Malta, pages 361–364.
- Saltzer, J. H. and Schroeder, M. D. (1975). The protection of information in computer systems. *Proceedings of the IEEE*, 63(9):1278–1308.
- Schmidhuber, J. (2015). Deep learning in neural networks: An overview. *Neural Networks*, 61:85–117.
- Sharan, V., Kakade, S., Liang, P., and Valiant, G. (2018). Prediction with a short memory. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1074–1087. ACM.
- Smith, L., Rindfleisch, T., and Wilbur, W. J. (2004). MedPost: a part-of-speech tagger for bioMedical text. *Bioinformatics*, 20(14):2320–2321.
- Tavanaei, A., Ghodrati, M., Kheradpisheh, S. R., Masquelier, T., and Maida, A. (2019). Deep learning in spiking neural networks. *Neural Networks*, 111:47–63.
- Toreini, E., Aitken, M., Coopamootoo, K., Elliott, K., Zelaya, C. G., and van Moorsel, A. (2020). The relationship between trust in AI and trustworthy machine learning technologies. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 272–283.
- Tran, K. M., Bisk, Y., Vaswani, A., Marcu, D., and Knight, K. (2016). Unsupervised neural hidden Markov models. In *SPNLP@EMNLP*, pages 63–71.
- Trentin, E. and Gori, M. (1999). Combining neural networks and hidden Markov models for speech recognition. *Neural Nets WIRN VIETRI-98*, pages 63–79.
- Varga, A. and Moore, R. K. (1990). Hidden Markov model decomposition of speech and noise. In *International Conference on Acoustics, Speech, and Signal Processing*, pages 845–848. IEEE.
- Varshney, K. R. (2019). Trustworthy machine learning and artificial intelligence. *XRDS: Crossroads, The ACM Magazine for Students*, 25(3):26–29.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008.
- Wang, B., Wang, A., Chen, F., Wang, Y., and Kuo, C.-C. J. (2019). Evaluating word embedding models: Methods and experimental results. *APSIPA transactions on signal and information processing*, 8.
- Wu, C. J. (1983). On the convergence properties of the EM algorithm. *The Annals of Statistics*, pages 95–103.
- Wu, X., Sahoo, D., and Hoi, S. C. (2020). Recent advances in deep learning for object detection. *Neurocomputing*.
- Yang, F., Balakrishnan, S., and Wainwright, M. J. (2017). Statistical and computational guarantees for the Baum-Welch algorithm. *The Journal of Machine Learning Research*, 18(1):4528–4580.
- Zhang, A., Gultekin, S., and Paisley, J. (2016a). Stochastic variational inference for the HDP-HMM. In *Artificial Intelligence and Statistics*, pages 800–808.
- Zhang, Y., Rahman, M. M., Braylan, A., Dang, B., Chang, H., Kim, H., McNamara, Q., Angert, A., Banner, E., Khetan, V., McDonnell, T., Nguyen, A. T., Xu, D., Wallace, B. C., and Lease, M. (2016b). Neural information retrieval: A literature review. *CoRR*, abs/1611.06792.
- Zou, Q., Xing, P., Wei, L., and Liu, B. (2019). Gene2vec: gene subsequence embedding for prediction of mammalian N6-methyladenosine sites from mRNA. *RNA*, 25(2):205–218.

APPENDIX

A Full Lagrangians of Standard HMM and DenseHMM

The full Lagrangian of the standard HMM model in the M-step reads

$$\begin{aligned}\bar{\mathcal{L}} &= \bar{\mathcal{L}}_1 + \bar{\mathcal{L}}_2 + \bar{\mathcal{L}}_3 \\ &= \sum_{i,j \in [n]} \sum_{t=2}^T \gamma_t(s_i, s_j) \log a_{ij} + \sum_{i \in [n]} \varphi_i \left(1 - \sum_{j \in [n]} a_{ij}\right) \\ &\quad + \sum_{i \in [n]} \sum_{t=1}^T \gamma_t(s_i) \log b_{i, j_{o_t}} + \sum_{i \in [n]} \varepsilon_i \left(1 - \sum_{j \in [m]} b_{ij}\right) \\ &\quad + \sum_{i \in [n]} \gamma_1(s_i) \log \pi_i + \bar{\varphi} \left(1 - \sum_{i \in [n]} \pi_i\right),\end{aligned}$$

where j_{o_t} describes the index of the observation observed at time t and $\bar{\varphi}, \varepsilon_i$ are Lagrange multipliers. Applying the transformations $\mathbf{A} = \mathbf{A}(\mathbf{U}, \mathbf{Z})$, $\mathbf{B} = \mathbf{B}(\mathbf{V}, \mathbf{W})$ and $\boldsymbol{\pi} = \boldsymbol{\pi}(\mathbf{U}, \mathbf{z}_{\text{start}})$ yields the full Lagrangian of the DenseHMM:

$$\begin{aligned}\bar{\mathcal{L}}^{\text{dense}} &= \bar{\mathcal{L}}_1^{\text{dense}} + \bar{\mathcal{L}}_2^{\text{dense}} + \bar{\mathcal{L}}_3^{\text{dense}} \\ &= \sum_{i,j \in [n]} \sum_{t=2}^T \gamma_t(s_i, s_j) \mathbf{u}_j \cdot \mathbf{z}_i \\ &\quad - \sum_{i,j \in [n]} \sum_{t=2}^T \gamma_t(s_i, s_j) \log \sum_{k \in [n]} \exp(\mathbf{u}_k \cdot \mathbf{z}_i) \\ &\quad + \sum_{i \in [n]} \sum_{t=1}^T \gamma_t(s_i) \mathbf{v}_{j_{o_t}} \cdot \mathbf{w}_i \\ &\quad - \sum_{i \in [n]} \sum_{t=1}^T \gamma_t(s_i) \log \sum_{j \in [m]} \exp(\mathbf{v}_j \cdot \mathbf{w}_i) \\ &\quad + \sum_{i \in [n]} \gamma_1(s_i) \mathbf{u}_i \cdot \mathbf{z}_{\text{start}} \\ &\quad - \sum_{i \in [n]} \gamma_1(s_i) \log \sum_{j \in [n]} \exp(\mathbf{u}_j \cdot \mathbf{z}_{\text{start}}).\end{aligned}$$

B Non-linear A-Matrix Factorization

All matrix sizes n and representation lengths l that contribute to the visualized l/n ratios in Figure 3 are shown in Table 1.

C Implementation Details and Data Preprocessing

Implementation Details. The backbone of our implementation is the library `hmmlearn`² that provides functions to optimize and score HMMs. The optimization schemes for the DenseHMM models $\mathcal{H}_{\text{dense}}^{\text{EM}}$ and $\mathcal{H}_{\text{dense}}^{\text{direct}}$ are implemented in tensorflow (Abadi et al., 2016). Both models use `tf.train.AdamOptimizer` with a fixed learning rate for optimization. At this point we note that experiments done with other optimizers such that `tf.train.GradientDescentOptimizer` lead to similar results in the evaluation. The representations are initialized using a standard isotropic Gaussian distribution. NLL values are normalized by the number of test sequences and by the maximum test sequence length.

Hardware Used. All experiments are conducted on a Intel(R) Xeon(R) Silver 4116 CPU with 2.10GHz and a NVidia Tesla V100.

Protein Dataset Preprocessing. The first 1,024 sequences of the RCSB PDB dataset have 22 unique symbols. We cut each sequence after a length of 512. Note that less than 4.9% of the 1,024 sequences exceed that length. Additionally, we collect the symbols of lowest frequency that together make up less than 0.2% of all symbols in the sequences and map them onto one residual symbol. This reduces the number of unique symbols in the sequences from 22 to 19.

Part-of-speech Sequences Preprocessing. We take 1,000 sequences from the Medpost dataset (from `tag_mb.ioc`) and cut them after a length of 40 which affects less than 15% of all sequences. We also collect the tags of lowest frequency that together make up less than 1% of all tags in the sequences and map them onto one residual tag. This reduces the number of tag items from 60 to 42.

Calculation of Ω^{gt} and Ω^{Model} . The co-occurrence matrices Ω^{model} and Ω^{gt} used to calculate the co-occurrence MADs in section 5 are estimated by counting subsequent pairs of observation symbols $(o_i(t), o_j(t+1)) \in \mathcal{O}^2$. For real-world data, Ω^{gt} is estimated based on the test data ground truth sequences. Equally long sequences sampled from the trained model are used to estimate Ω^{model} . In case of synthetic data, Ω^{gt} is calculated analytically (eq. 3) instead.

²<https://github.com/hmmlearn/hmmlearn>

Table 1: Approximation errors (median with 25/75 percentile) of normAbsLin-based and softmax-based matrix factorizations for different matrix sizes n and representation lengths l .

n	l	median (25/75 percentile) of loss($\tilde{\mathbf{A}}, \mathbf{A}_{gt}$)	
		$\tilde{\mathbf{A}} = \text{normAbsLin}(\mathbf{UZ})$	$\tilde{\mathbf{A}} = \text{softmax}(\mathbf{UZ})$
3	1	0.678 (0.652/0.696)	0.048 (0.004/0.110)
3	2	0.162 (0.002/0.270)	0.001 (0.000/0.001)
3	3	0.001 (0.001/0.001)	0.001 (0.000/0.001)
3	5	0.001 (0.001/0.001)	0.001 (0.001/0.001)
5	1	0.769 (0.745/0.827)	0.453 (0.321/0.505)
5	3	0.346 (0.093/0.396)	0.001 (0.001/0.003)
5	5	0.001 (0.001/0.002)	0.001 (0.001/0.001)
5	10	0.001 (0.001/0.002)	0.002 (0.001/0.003)
10	1	0.862 (0.851/0.868)	0.616 (0.581/0.645)
10	5	0.310 (0.235/0.345)	0.012 (0.005/0.028)
10	10	0.002 (0.002/0.002)	0.003 (0.002/0.005)
10	15	0.002 (0.002/0.002)	0.003 (0.003/0.043)

