



PhD Thesis

Casper Hansen

Representation Learning for Efficient and Effective Similarity Search and Recommendation

Advisors: Stephen Alstrup, Christina Lioma, Jakob Grue Simonsen

Handed in: April 28, 2021

This thesis has been submitted to the PhD School of The Faculty of Science, University of Copenhagen

Abstract

How data is represented and operationalized is critical for building computational solutions that are both effective and efficient. A common approach is to represent data objects as binary vectors, denoted *hash codes*, which require little storage and enable efficient similarity search through direct indexing into a hash table or through similarity computations in an appropriate space. Due to the limited expressibility of hash codes, compared to real-valued representations, a core open challenge is how to generate hash codes that well capture semantic content or latent properties using a small number of bits, while ensuring that the hash codes are distributed in a way that does not reduce their search efficiency. State of the art methods use representation learning for generating such hash codes, focusing on neural autoencoder architectures where semantics are encoded into the hash codes by learning to reconstruct the original inputs of the hash codes. This thesis addresses the above challenge and makes a number of contributions to representation learning that (i) improve effectiveness of hash codes through more expressive representations and a more effective similarity measure than the current state of the art, namely the Hamming distance, and (ii) improve efficiency of hash codes by learning representations that are especially suited to the choice of search method. The contributions are empirically validated on several tasks related to similarity search and recommendation.

Dansk Resumé

Hvordan data repræsenteres og operationaliseres er afgørende for opbygningen af effektive beregningsmodeller. En almindelig tilgang er at repræsentere dataobjekter som binære vektorer, betegnet hash-koder, der kræver lidt lagerplads og muliggør effektiv similaritetssøgning gennem direkte indeksering i en hash-tabel eller gennem similaritetsberegninger i et passende rum. På grund af den begrænsede ekspressibilitet af hash-koder, sammenlignet med flydende tal repræsentationer, så er en essentiel udfordring hvordan man genererer hash-koder, der repræsenterer semantisk indhold, eller latente egenskaber, godt ved hjælp af et lille antal bits, samtidig med at man sørger for, at hash-koder distribueres på en måde der ikke reducerer deres søgeeffektivitet. State of the art metoder bruger repræsentationslæring til generering af sådanne hash-koder med fokus på neurale autoencoderarkitekturer, hvor semantik er kodet i hash-koder ved at lære at rekonstruere de originale input af hash-koderne. Denne afhandling adresserer ovennævnte udfordring og bidrager med en række nye metoder til repræsentationslæring, der (i) forbedrer effektiviteten af hash-koder gennem mere ekspressive repræsentationer og et mere effektivt afstandsmål end den nuværende kendte teknik, nemlig Hamming distancen, og (ii) forbedre søgeeffektiviteten af hash-koder ved at lære repræsentationer, der er særligt velegnede til valget af søgemetode. Bidragene valideres empirisk på flere problemer relateret til similaritetssøgning og recommendation.

Acknowledgements

My PhD has been a wonderful academic and personal journey that would not have been possible without the support from family, friends, and colleagues. First, I would like to thank my academic advisors Christina, Jakob, and Stephen from whom I have without a doubt learned many valuable lessons. I have especially enjoyed my weekly meetings with Christina and Jakob, where we have had many insightful and entertaining discussions. I would also like to thank my office mates over the years, Dongsheng, Lucas, Niklas, and Stephan, for providing a great work environment with many joyous moments. I was fortunate enough to be able to do an internship at Spotify Research with many fantastic people, where I would especially like to thank Brian, Federico, Lucas, Mounia, and Rishabh. I had a great summer in London! I would also like to thank Benjamin and Rastin from Danmarks Nationalbank, and Esben and Jacob from the Municipality of Copenhagen, it was great discussing and collaborating with you. Next, I would like to thank all my friends (including many of the people mentioned above!) and family, especially David and Mikkel, who have been a steady source of support during the years and made sure I could always see the bigger picture. Lastly, and most importantly, I would like to deeply thank my twin brother Christian for always being there for me and being a great source of inspiration, food, and snacks.

Contents

Abstract	i
Dansk Resumé	ii
Acknowledgements	iii
1 Introduction	1
1.1 Research Outline	3
1.1.1 Document Similarity Search	3
1.1.2 Recommendation	6
1.2 Summary of Contributions	8
1.3 Future Work	9
1.4 List of Publications	12
2 Unsupervised Neural Generative Semantic Hashing	15
3 Unsupervised Semantic Hashing with Pairwise Reconstruction	26
4 Unsupervised Multi-Index Semantic Hashing	31
5 Content-aware Neural Hashing for Cold-start Recommendation	42
6 Projected Hamming Dissimilarity for Bit-Level Importance Coding in Collaborative Filtering	53
Bibliography	63

Chapter 1

Introduction

Similarity search and recommendation are two types of Information Retrieval (IR) tasks used for supporting users in filtering and exploring data, such as searching for interesting books or movies. To build solutions for such tasks, learning representations of information objects, such as words, documents, and items, is a core challenge in IR, with the aim of capturing key semantics and properties for retrieving the most relevant information.

Traditional IR text representations have focused on bag-of-words (BoW) representations with a fixed vocabulary (i.e., a limited set of words), based on various term weighting approaches, such as term-frequency [53], tf-idf [53], BM25 [51], and language models [50] (see [6, 43] for an overview on term weighting). In BoW vector representations, each dimension corresponds to a specific vocabulary word, which results in highly sparse vectors as the only non-zero dimensions are the vocabulary words occurring in a given text. Due to the sparsity, BoW representations are very efficient to use in practice through fast indexing structures for retrieval. However, such BoW representations suffer from the problem of *vocabulary mismatch* [15]. Vocabulary mismatch refers to the problem when different words or phrases are used to express the same meaning, which occurs for BoW representations as each word is only mapped to a single dimension [68].

One way to address the vocabulary mismatch problem is through *distributional semantics*, which is based on the distributional hypothesis stating that words occurring in the same context tend to have similar meanings [35]. In distributional semantics, text representations are typically represented as dense vectors of significantly fewer dimensions than the vocabulary size, where each word is distributed across all dimensions. Such vectors are typically denoted as embeddings, as they are obtained by embedding high dimensional data (or complex data objects generally) into a relatively low dimensional space in such a way that semantically similar data is placed close together in the embedding space¹. Early work on distributional semantics for representing documents was based on computing a matrix

¹Note that many types of data can be effectively embedded, not just words and documents, such that data objects close together in the embedding space correspond to objects with similar properties.

decomposition of the document-term co-occurrence matrix, e.g., using the seminal Latent Semantic Indexing (LSI) approach [10], such that similar documents are represented as similar embeddings. However, computing the decomposition of the co-occurrence matrix is computationally expensive on large-scale data due to the runtime complexity of the singular value decomposition typically used [54]. In a less computationally expensive way, embeddings have also been constructed on a word level, where methods such as word2vec [46] and GloVe [49] learn word embeddings based on word-word co-occurrences. Specifically, these methods aim at learning word embeddings that are able to predict a word given its context, or the context given a word, where the context is represented by n words before and after a given word (denoted as n -length context windows). The use of context windows can be seen in contrast to using all the words in a document as done by LSI. To create document embeddings, embeddings of the words within a document can be aggregated, e.g., by taking an average [39], a frequency-weighted average [3], or through max-pooling [69]. However, while such embeddings provide an effective encoding of document semantics, the cost associated with similarity computations (e.g., using the inner product), as well as the amount of storage required, is non-trivial and may hinder use in very large-scale settings. We will next describe a type of representation that supports faster similarity computations while simultaneously requiring significantly less storage.

Hash Code Representations

An efficient approach for representing documents (or more generally *data objects*) is through short binary vector representations denoted as *hash codes*. In contrast to real-valued embeddings, hash codes may require only a few bytes of storage, and support fast similarity computations through the Hamming distance, which can be computed in two CPU instructions counting the number of differing bits between two hash codes. Additionally, hash codes can be used for direct indexing into a hash table for finding similar documents with exact hash code matches, or by making appropriate bit flips for directly finding documents up to a specific Hamming distance to the reference (i.e., query) hash code. Locality-Sensitive Hashing (LSH) [16, 36] is a widely known family of hashing techniques, where highly similar (i.e., near-identical) documents have a high probability of hashing to the same hash code. However, LSH uses data-independent hashing functions for mapping similar documents to similar hash codes, such as random projections. Data-independent hashing functions have the limitation of not utilizing any training data in their construction, which limits their ability to encode semantics. For this reason, LSH techniques are less suited for computing the semantic similarity between documents.

In contrast, semantic hashing [52] methods are data-dependent and learn a mapping for transforming documents into semantically-aware hash codes, such that semantically similar documents have a short Hamming distance between their associated hash codes. Some of the first work on semantic hashing utilized quantized (i.e., binary) versions of familiar techniques, such as LSI [62] and spectral

clustering [60], and through two step approaches of first finding the optimal encoding and then learning to predict this as a classification problem [60, 63]. Recent work has focused on neural network approaches, specifically variational autoencoder architectures [7, 8, 13, 20, 22, 24, 56]. These models are trained with the objective of reconstructing the input representation (typically a tf-idf BoW representation) from the latent representation of the most internal layer in the autoencoder, which corresponds to the hash code. In this way, such models learn an encoding that captures the key semantics within the hash code, as it otherwise would be unable to reconstruct it well. Learned hash codes have been found to be effective not only in document similarity search, but also for representing users and items in recommender systems [23, 28, 38, 40, 65, 67, 70]. In this setting, user representations act as query objects for which the most relevant items should be retrieved. This is analogous to the document similarity search setting where documents act as both queries and items to be retrieved.

This thesis addresses the challenges of generating efficient hash codes that well capture semantic content, or latent properties, using a small number of bits. We make a number of contributions to representation learning that (i) improve effectiveness of hash codes through more expressive representations and a more effective similarity measure than the current state of the art, namely the Hamming distance, and (ii) improve efficiency of hash codes by learning representations that are especially suited to the choice of search method. The contributions are empirically validated on several tasks related to similarity search and recommendation.

1.1 Research Outline

This thesis is composed of five chapters, where each chapter is a published article (see Section 1.4 for the list of publications). These five chapters are grouped into the application areas of document similarity search (Section 1.1.1) and recommendation (Section 1.1.2).

This section provides an outline of each chapter of the thesis, where we state the main research questions, provide background information, and overview the main findings and observations.

1.1.1 Document Similarity Search

The task of document similarity search consists of searching a set of document objects by a given query object, such that those document objects most similar to the query are retrieved. In the setting of document similarity search using hash codes, similarity search is usually considered either a radius search or a k-nearest neighbour (kNN) search. In a radius search, all hash codes with a specified maximum distance to the query hash code are to be found, whereas for kNN the radius is incrementally increased until the k^{th} hash code has a distance equal to the search radius. In our work on document similarity search using semantic hashing, we fo-

cus on kNN search as per related work [7, 8, 13, 56], but note that the work can trivially be adapted for radius search.

Chapter 2. Unsupervised Neural Generative Semantic Hashing

Existing approaches for semantic hashing are typically unsupervised as document labels are most often not available for large data collections of a size where fast similarity search using hash codes is particularly useful. The approaches have evolved into three main groups. Firstly, being based on two-step procedures of first learning the optimal hash code encoding (based on low Hamming distances between semantically similar documents) and then learning to predict this to handle unseen documents [60, 63]. Secondly, variational autoencoder models with an input reconstruction loss using a post-processing quantization step for obtaining the hash codes [7, 8]. The quantization step consists of rounding each dimension based on its median value, such that the k^{th} bit of the hash code is 1 if the k^{th} dimension of the learned (real-valued) vector is larger than the median, and vice versa for setting a bit to 0. Thirdly, variational autoencoders using Bernoulli priors for learning hash codes in an end-to-end fashion [56], which improves hash code effectiveness due to reducing the quantization error. However, none of these approaches directly model the goal of ranking documents by their similarity to the query document, i.e., for maximizing the precision of a kNN search, but rather implicitly assume that the learned hash codes will enable this by focusing on encoding the document semantics well. This leads us to the first research question:

RQ1 To what extent can ranking become an organic part of learning semantic hash codes?

To answer this question, we first need to obtain relevance labels between a query document and the remaining documents. As we work in an unsupervised setting, this cannot be based on provided document tags or labels, but rather we take a weakly supervised direction by using an existing unsupervised semantic hashing approach for obtaining approximate top-K rankings of each document. Based on this, we extract ranking triplets for training, such that each sample consists of query document, a similar document, and a dissimilar document. With such ranking triplets, the aim is to learn hash codes that are better able to separate similar and dissimilar documents in the Hamming space. We propose Ranking-based Semantic Hashing (RBSH), a variational autoencoder with a traditional input reconstruction objective trained jointly with a hinge loss on the ranking triplets as the ranking objective. We find that the ranking objective has a beneficial regularizing effect, as the hash codes, especially short hash codes down to 8 bits, otherwise have a tendency to cluster and not sufficiently utilize the space, which reduces the kNN effectiveness. We experimentally find that RBSH hash codes significantly outperform state-of-the-art approaches, and most importantly, yield state-of-the-art performance while using 2-4x fewer bits than existing approaches.

Chapter 3. Unsupervised Semantic Hashing with Pairwise Reconstruction

The findings from **RQ1** suggest that incorporating neighbourhood information, in our case in the form of optimizing a joint ranking objective, can improve the generalization of the learned hash codes as observed by higher kNN effectiveness. This observation is also noted in early work [60, 63], where the cosine similarity between documents in their original space is used for constructing pairwise weights for learning to minimize a weighted Hamming distance between hash code pairs. In this way, hash codes from similar documents are forced to have a small Hamming distance, whereas dissimilar documents with a small or negative weight are ignored or forced further apart in the Hamming space. In more recent work, a variational autoencoder model, denoted NbrReg, with two reconstruction objectives (represented as two different decoders in the autoencoder) is proposed [7]. The first reconstruction objective is the typical input reconstruction, whereas the second aims at reconstructing all unique words occurring in a local neighbourhood around the input document (i.e., an aggregated neighbourhood document). Based on this line of work, we ask the following research question:

RQ2 To what extent can local semantic neighbourhoods be incorporated as an organic part of learning semantic hash codes?

We answer this question by proposing an extension to the input reconstruction objective shared among the variational autoencoder approaches [7, 8, 22, 56]. Similarly to our previously proposed RBSH (**RQ1**), we use an existing unsupervised semantic hashing approach to retrieve a set of similar documents to each query document. Using this set, we construct training pairs and train a variational autoencoder, named PairRec, to be able to reconstruct the query document from both hash codes (i.e., a pairwise reconstruction). In this way, more general hash codes are learned that not only encode their own semantics, but also those of similar documents, which we experimentally show to be more effective than existing state-of-the-art approaches. In contrast to NbrReg [7], we consider pairs of real documents, as opposed to the aggregated neighbourhood document used by NbrReg, which we argue may introduce a semantic shift as such documents cannot be encountered during retrieval. Furthermore, our approach uses only a single decoder, as the two decoders used by NbrReg increase the overall decoding complexity, which may be difficult to capture using the simple Hamming distance and hence unnecessarily reduce the kNN effectiveness.

Chapter 4. Unsupervised Multi-Index Semantic Hashing

RQ1 and **RQ2** address ways for learning more expressive hash codes that improve kNN effectiveness. However, both our and prior work have generally assumed that the hash codes will be efficient to use in practice, but not considered any concrete search method to explore whether any efficiency differences exist between the approaches. While real-time brute-force linear scans are possible using the Hamming

distance on large-scale data [55], significantly faster alternatives exist enabling sub-linear search time. One such alternative is multi-index hashing [47, 48], a method for performing exact radius and kNN search in the Hamming space. While hash codes can be used as direct addresses into a hash table, the number of such lookups scales exponentially with the radius (where the number of bits is the base), which can become infeasible for long hash codes or even moderate radii. To fix this, multi-index hashing splits each hash code into m disjoint substrings, and utilizes the pigeonhole principle to determine that if two hash codes are within radius r , then at least one substring exists where the Hamming distance between the substrings is at most $\lfloor \frac{r}{m} \rfloor$. Based on this, multi-index hashing builds a candidate set, significantly smaller than the entire set of hash codes, from which a linear scan is performed to determine the exact Hamming distances. This leads to our next research question:

RQ3 To what extent can key properties for multi-index hashing become an organic part of learning semantic hash codes?

To answer this question, we first identify two key properties for high multi-index efficiency. The first one is to reduce the number of documents per hash table lookup by reducing the number of false-positive candidates, i.e., those with a small substring distance, but high overall hash code distance. The second one is to distribute the hash codes sufficiently such that the distance to the k^{th} document is kept low to limit the exponential growth. We operationalize these properties into model-agnostic training objectives for training hash codes specifically designed for multi-index hashing in an end-to-end fashion, which we denote as Multi-Index Semantic Hashing (MISH). We experimentally find that state-of-the-art baselines are upwards of 33% slower than hash codes generated by MISH without being more effective.

1.1.2 Recommendation

Recommender systems are trained on data from user-item interactions, e.g., clicks or ratings, with the objective of learning to estimate user preferences to provide relevant recommendations. When users and items are represented as embeddings, such as hash codes, the recommendation task is highly similar to that of similarity search (Section 1.1.1), but with the modification that the user representation acts as the query in contrast to using an item as a query. In the following chapters, we focus on two different recommendation settings, depending on the availability of content information, but they have the commonality of generating user and item hash codes to be used for recommendation.

Chapter 5. Content-aware Neural Hashing for Cold-start Recommendation

Recommendation approaches based on collaborative filtering, content-based filtering, and their combinations have been well studied and shown to work well in

practice [1, 57]. Collaborative filtering learns user and item embeddings based on user-item interactions, such as implicit feedback (e.g., clicks) or explicit feedback (e.g., ratings). However, when new items appear (denoted *cold-start* items) no previous user-item interactions exist, hence collaborative filtering approaches are unable to learn such embeddings. Content-based filtering solves the cold-start problem by using item content information (or other item features) for recommending items similar to items a user has previously enjoyed. For efficiency reasons, a number of hashing-based approaches have also been proposed for the recommendation domain. These approaches have primarily focused on the collaborative filtering setting [38, 44, 64, 67, 70], but less so on content-aware approaches addressing the cold-start problem. Existing content-aware hashing approaches, DDL [66] and DCMF [41], learn to generate user and item hash codes for use in both standard and cold-start settings, however they both share the problem of generating item hash codes differently depending on whether the item is considered cold-start or not. Specifically, they both learn user and item hash codes in a typical collaborative filtering setting, but simultaneously learn separate item hash codes based on their content information, such that the distance between the two types of item hash codes is minimized. We argue that this is unnecessary and may limit generalizability, which leads to the next research question:

RQ4 How can item content information be used for generating hash codes in the same way for both standard and cold-start items to improve recommendation effectiveness?

To answer this question, we propose NeuHash-CF, a joint hashing model for generating user and item hash codes robust to the cold-start problem. Inspired by semantic hashing, we use the basic architecture of our previous work [22, 24] to construct an item component that learns hash codes based on content information alone. In contrast, the user component is based solely on a user ID as in the typical collaborative filtering setting, and the model is jointly optimized by learning to reconstruct the logged user-item ratings. Since the item hash codes are generated entirely based on content information, they are by default robust to the cold-start problem as long as new items share some similarity to some of the existing items. We experimentally evaluate NeuHash-CF hash codes against state-of-the-art baselines, including collaborative filtering and content-aware approaches, where we observe significant improvements in both cold-start and standard recommendation settings.

Chapter 6. Projected Hamming Dissimilarity for Bit-Level Importance Coding in Collaborative Filtering

In hashing-based learning, such as for documents (**RQ1-3**) or users and items (**RQ4**), the default way of measuring the similarity between two hash codes is through the Hamming distance, which is computed as the summation of the XOR operation between two hash codes. However, by definition, the Hamming distance

weighs each bit equally, which could be problematic when the importance of each bit’s underlying properties differ depending on the query. For example, in a collaborative filtering setting the user hash code represents the query, and depending on the user’s historic item interactions, it may be possible to infer that certain underlying properties are more important for the item ranking. While approaches have been proposed for assigning real-valued weights to certain substrings of bits in a hash code [14, 44], such a weighting has the problem of making the core similarity computation (e.g., Hamming distance) significantly slower, which limits its usage in large-scale settings where hashing-based solutions are most needed. This leads to the next research question:

RQ5 How can a similarity measure support hash code importance weighting without reducing efficiency?

To answer this question, we consider the collaborative filtering setting of user and item hash codes. We derive a new way of measuring the dissimilarity between a user (query) and item hash code with binary weighting of each dimension, corresponding to disabling specific bits. To this end, we consider a general definition of dissimilarity defined as the norm of the difference between the user representation and the item representation projected onto the user, which in the Euclidean space corresponds to the well known cosine distance. We show that when working in the Hamming space, this results in a novel projected Hamming dissimilarity, which by choice of projection, effectively allows an importance weighting of the item hash codes through the user hash code. Specifically, if the possible bit values are $\{-1, 1\}$, then all bit dimensions with -1 in the user hash code are ignored in the dissimilarity computation across all items, corresponding to a binary importance weighting of each bit dimension. We experimentally show that hash codes optimized for the projected Hamming dissimilarity lead to large gains in recommendation effectiveness compared to baselines using the Hamming distance, while requiring no additional storage and no reduction in efficiency compared to computing the Hamming distance.

1.2 Summary of Contributions

This thesis addresses challenges related to representation learning for enabling more efficient and effective similarity search and recommendation. The contributions focus on improving the representational power of learned hash codes, as well as learning to construct them in a way that enables more efficient similarity search.

- The first contribution is a method for retrieving semantically relevant documents based on learned hash codes, which are constructed in a way that encodes both document semantics and ranking information relative to other

hash codes. The ranking part ensures that the hash codes are directly designed for being ranked, rather than only encoding semantics, which improves retrieval effectiveness.

- The second contribution is an improved method of learning document hash codes, based on introducing local semantic neighbourhood information into the learning process. Specifically, hash codes from pairs of semantically similar documents are trained such that both hash codes can reconstruct the same original document, which effectively ensures that semantically similar documents are placed close together in the learned hash code space.
- The third contribution is a method for learning hash codes that enables more efficient retrieval by directly optimizing the hash code construction towards the search approach being used. In contrast to prior work that applies post-processing steps after training for improving efficiency, we find that learning it in a direct end-to-end trainable way enables even larger efficiency gains without compromising effectiveness.
- The fourth contribution is a method for learning user and item hash codes for recommendation, while being robust to the problem of cold-start items. This is achieved by using content information for generating item hash codes in a unified way for both standard and cold-start items, rather than distinguishing between them as done by prior work, which improves recommendation effectiveness.
- The fifth contribution is a new dissimilarity measure for comparing two hash codes, which enables a binary weighting of the hash codes, corresponding to disabling the bits deemed unimportant by one of the hash codes. Hash codes optimized for this dissimilarity, rather than the current state of the art, namely the Hamming distance, results in higher effectiveness in collaborative filtering experiments, while requiring no additional storage and no computational overhead compared to the Hamming distance.

1.3 Future Work

Our contributions towards more efficient and effective representations for similarity search and recommendation naturally have their limitations and possible directions of future work. Below we first discuss specific directions of future work for research presented in this thesis, and follow with a discussion of more general directions of future work.

New input representations

All text hashing approaches proposed in this thesis, as well as the baselines used for comparison, use a tf-idf weighted bag-of-words vector as the document rep-

resentation. As state-of-the-art approaches all use variational autoencoders, the motivation for using a fixed input vector is for using it in reconstruction when decoding the hash code as part of training the autoencoder. As an alternative, Doan and Reddy [12] recently explored two approaches using recurrent and convolutional neural networks, with the objective of being able to reconstruct the word embeddings of the input. They empirically found this to be more effective than the (tf-idf vector) baselines, but they did not compare against more recent state-of-the-art approaches, so the exact improvement is uncertain. However, their work does provide a first step towards exploring more expressive representations than a tf-idf vector, which is a natural next step for our work as well.

Exploring the projected Hamming dissimilarity and other functions in new settings

We proposed the projected Hamming dissimilarity as a way to compute weighted dissimilarities in the asymmetrical user-item setting, where the user hash code acts as a query used for searching among the item hash codes. While our work showed large effectiveness gains, a next step would be exploring it (as well as other potential functions) in a symmetrical item-item setting, such as document similarity search. A potential obstacle is that each item also acts as a query for similarity search, which may be problematic as the weighting in the projected Hamming dissimilarity is defined through the query, hence its dual purpose could hinder learning a good representation. However, a first step would be to simply learn two hash codes per item, one for the purpose of a query, and one for the normal purpose of an item. Note that this does not necessarily lead to increased storage requirements, as the similarity search often is performed on new (unseen) documents, which would require an encoding step for obtaining the hash code in any case.

Hashing-based representations in dynamic settings

This thesis has considered application areas where item (and user) hash codes could be generated without considering the need for any updates. While the hash code models should be retrained regularly, and new hash codes should be generated when new data arrives, this has not been a focus in our work. To this end, it would be interesting to investigate the effectiveness of hashing-based approaches for more dynamic settings, such as context-aware or feedback-aware recommender systems [5, 9, 21, 61], where the user representation should change often (e.g., after every few interactions) to better represent the (contextual or temporal) needs of the user. While more expressive real-valued representations might be needed for the final item ranking, the hash codes could prove useful for efficiently generating highly relevant candidate pools.

Representation interpretability

Interpretability is a core challenge for representation learning, as well as for machine learning as a whole. Investigating the relation between interpretable concepts and how embeddings are constructed, whether being real-valued or binary, is important for understanding what a model has learned, which may improve the trust we as humans assign to such models for use in real-world applications. For real-valued embeddings, one way to improve interpretability is through inducing sparsity when constructing the embeddings [58], such that the embedding of the document (or another type of object) can be described by a few non-zero dimensions, corresponding to some latent topics. Each topic, or combination of topics, would then be described by a set of documents from which a common theme could be extracted. However, for hash codes, this type of sparsity is not as straightforward to induce because hash codes are unable to represent the equivalent of a real-valued zero, as both possible binary values impact the distance computation. As little work has been done in hash code interpretability, it would be worth pursuing that direction in the future.

Representation pretraining

Learning representations that generalize well is another important challenge, but doing so may require large and varied amounts of appropriate training data. One way to solve this problem is through model pretraining, with one of the most successful examples being word embedding models (e.g., word2vec [46] and BERT [11]). These models are trained in a unsupervised fashion on massive collections of text, from which general language understanding is learned, such that the models later can be fine-tuned (or even applied directly) to downstream tasks. Utilizing, and further exploring, this idea of pretraining is highly attractive, as it allows new models to exploit the language understanding obtained through the pretraining for improving their generalization. Interestingly, such pretrained models have not yet been explored for the task of learning semantic hash codes for documents, neither through using existing word embedding models or (pre)training hash code models on the same massive collections of text as the original word embedding models.

1.4 List of Publications

The following publications are included as chapters of this thesis (* denotes equal contribution):

Chapter 2 Casper Hansen, Christian Hansen, Jakob Grue Simonsen, Stephen Alstrup, Christina Lioma (2019). Unsupervised Neural Generative Semantic Hashing. In SIGIR, pages 735-744. [22].

Chapter 3 Casper Hansen*, Christian Hansen*, Jakob Grue Simonsen, Stephen Alstrup, Christina Lioma (2020). Unsupervised Semantic Hashing with Pairwise Reconstruction. In SIGIR, pages 2009-2012. [24].

Chapter 4 Christian Hansen*, Casper Hansen*, Jakob Grue Simonsen, Stephen Alstrup, Christina Lioma (2021). Unsupervised Multi-Index Semantic Hashing. In WWW, in press. [20].

Chapter 5 Casper Hansen*, Christian Hansen*, Jakob Grue Simonsen, Stephen Alstrup, Christina Lioma (2020). Content-aware Neural Hashing for Cold-start Recommendation. In SIGIR, pages 971-980. [23].

Chapter 6 Christian Hansen*, Casper Hansen*, Jakob Grue Simonsen, Christina Lioma (2021). Projected Hamming Dissimilarity for Bit-Level Importance Coding in Collaborative Filtering. In WWW, in press. [28].

Furthermore, publications not part of this thesis have been written during the PhD studies on the following topics: (non hashing-based) recommender systems [21, 32, 34], computational approaches for COVID-19 prediction and retrieval [37, 42], fact-checking [4, 19, 25, 26, 27, 33], (non hashing-based) text representation and classification [18, 30, 45, 59], smart city analytics [17], and educational data mining [2, 29, 31]. These publications are listed below in reverse chronological order:

1. Dongsheng Wang*, Casper Hansen*, Lucas Chaves Lima, Christian Hansen, Maria Maistro, Jakob Grue Simonsen, Christina Lioma (2021). Multi-Head Self-Attention with Role-Guided Masks. In ECIR, pages 432-439. [59].
2. Christian Hansen, Rishabh Mehrotra, Casper Hansen, Brian Brost, Lucas Maystre, Mounia Lalmas (2021). Shifting Consumption towards Diverse Content on Music Streaming Platforms. In WSDM, pages 238-246. [34].
3. Espen Jimenez Solem, Tonny Studsgaard Petersen, Casper Hansen, Christian Hansen, et al. (2021). Developing and Validating COVID-19 Adverse Outcome Risk Prediction Models from a Bi-national European Cohort of 5594 Patients. In Scientific Reports 11 (1), pages 1-12. [37].
4. Casper Hansen, Christian Hansen, Lucas Maystre, Rishabh Mehrotra, Brian Brost, Federico Tomasi, Mounia Lalmas (2020). Contextual and Sequential User Embeddings for Large-Scale Music Recommendation. In RecSys, pages 53-62. [21].

5. Lucas Chaves Lima*, Casper Hansen*, Christian Hansen, Dongsheng Wang, Maria Maistro, Birger Larsen, Jakob Grue Simonsen, Christina Lioma (2021). Denmark's Participation in the Search Engine TREC COVID-19 Challenge: Lessons Learned about Searching for Precise Biomedical Scientific Information on COVID-19. In TREC COVID-19 Challenge. [42].
6. Casper Hansen, Christian Hansen, Jakob Grue Simonsen, Christina Lioma (2020). Fact Check-Worthiness Detection with Contrastive Ranking. In CLEF, pages 124-130. [27].
7. Christian Hansen, Casper Hansen, Jakob Grue Simonsen, Birger Larsen, Stephen Alstrup, Christina Lioma (2020). Factuality Checking in News Headlines with Eye Tracking. In SIGIR, pages 2013-2016. [33].
8. Isabelle Augenstein, Christina Lioma, Dongsheng Wang, Lucas Chaves Lima, Casper Hansen, Christian Hansen, Jakob Grue Simonsen (2019). MultiFC: A Real-World Multi-Domain Dataset for Evidence-Based Fact Checking of Claims. In EMNLP, pages 4685-4697. [4].
9. Casper Hansen, Christian Hansen, Jakob Grue Simonsen, Christina Lioma (2019). Neural Weakly Supervised Fact Check-Worthiness Detection with Contrastive Sampling-Based Ranking Loss. In CLEF-2019 Fact Checking Lab. [26].
10. Casper Hansen, Christian Hansen, Stephen Alstrup, Jakob Grue Simonsen, Christina Lioma (2019). Contextually Propagated Term Weights for Document Representation. In SIGIR, pages 897-900. [18].
11. Christian Hansen, Casper Hansen, Stephen Alstrup, Christina Lioma (2019). Modelling End-of-Session Actions in Educational Systems. In EDM, pages 306-311. [29].
12. Casper Hansen, Christian Hansen, Stephen Alstrup, Jakob Grue Simonsen, Christina Lioma (2019). Neural Check-Worthiness Ranking with Weak Supervision: Finding Sentences for Fact-Checking. In Companion Proceedings of WWW, pages 994-1000. [19].
13. Christian Hansen, Casper Hansen, Stephen Alstrup, Jakob Grue Simonsen, Christina Lioma (2019). Neural Speed Reading with Structural-Jump-LSTM. In ICLR. [30].
14. Christian Hansen, Casper Hansen, Jakob Grue Simonsen, Stephen Alstrup, Christina Lioma (2019). Modelling Sequential Music Track Skips Using a Multi-RNN Approach. In WSDM Cup. [32].
15. Rastin Matin, Casper Hansen, Christian Hansen, Pia Mølgaard (2019). Predicting Distresses using Deep Learning of Text Segments in Annual Reports. In Expert Systems With Applications (132), pages 199-208. [45].

16. Casper Hansen, Christian Hansen, Jakob Grue Simonsen, Christina Lioma (2018). The Copenhagen Team Participation in the Check-Worthiness Task of the Competition of Automatic Identification and Verification of Claims in Political Debates of the CLEF2018 CheckThat! Lab. In CLEF-2018 Fact Checking Lab. [25].
17. Casper Hansen, Christian Hansen, Stephen Alstrup, Christina Lioma (2017). Smart City Analytics: Ensemble-Learned Prediction of Citizen Home Care. In CIKM, pages 2095-2098. [17].
18. Stephen Alstrup, Casper Hansen, Christian Hansen, Niklas Hjuler, Stephan Lorenzen, Ninh Pham (2017). DABAI: A data driven project for e-Learning in Denmark. In ECEL, pages 18-24. [2].
19. Christian Hansen, Casper Hansen, Niklas Hjuler, Stephen Alstrup, Christina Lioma (2017). Sequence Modelling For Analysing Student Interaction with Educational Systems. In EDM, pages 232-237. [31].

Chapter 2

Unsupervised Neural Generative Semantic Hashing

Casper Hansen, Christian Hansen, Jakob Grue Simonsen, Stephen Alstrup, Christina Lioma (2019). Unsupervised Neural Generative Semantic Hashing. In SIGIR, pages 735-744. [22].

Unsupervised Neural Generative Semantic Hashing

Casper Hansen
University of Copenhagen
c.hansen@di.ku.dk

Christian Hansen
University of Copenhagen
chrh@di.ku.dk

Jakob Grue Simonsen
University of Copenhagen
simonsen@di.ku.dk

Stephen Alstrup
University of Copenhagen
s.alstrup@di.ku.dk

Christina Lioma
University of Copenhagen
c.lioma@di.ku.dk

ABSTRACT

Fast similarity search is a key component in large-scale information retrieval, where semantic hashing has become a popular strategy for representing documents as binary hash codes. Recent advances in this area have been obtained through neural network based models: generative models trained by learning to reconstruct the original documents. We present a novel unsupervised generative semantic hashing approach, *Ranking based Semantic Hashing* (RBSH) that consists of both a variational and a ranking based component. Similarly to variational autoencoders, the variational component is trained to reconstruct the original document conditioned on its generated hash code, and as in prior work, it only considers documents individually. The ranking component solves this limitation by incorporating inter-document similarity into the hash code generation, modelling document ranking through a hinge loss. To circumvent the need for labelled data to compute the hinge loss, we use a weak labeller and thus keep the approach fully unsupervised.

Extensive experimental evaluation on four publicly available datasets against traditional baselines and recent state-of-the-art methods for semantic hashing shows that RBSH significantly outperforms all other methods across all evaluated hash code lengths. In fact, RBSH hash codes are able to perform similarly to state-of-the-art hash codes while using 2-4x fewer bits.

KEYWORDS

Unsupervised semantic hashing, Deep learning, Generative model, Document ranking

ACM Reference Format:

Casper Hansen, Christian Hansen, Jakob Grue Simonsen, Stephen Alstrup, and Christina Lioma. 2019. Unsupervised Neural Generative Semantic Hashing. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '19)*, July 21–25, 2019, Paris, France. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3331184.3331255>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGIR '19, July 21–25, 2019, Paris, France

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6172-9/19/07...\$15.00

<https://doi.org/10.1145/3331184.3331255>

1 INTRODUCTION

The task of similarity search consists of querying a potentially massive collection to find the content most similar to a query. In Information Retrieval (IR), fast and precise similarity search is a vital part of large-scale retrieval [28], and has applications in content-based retrieval [14], collaborative filtering [13], and plagiarism detection [10, 26]. Processing large-scale data requires solutions that are both computationally efficient and highly effective, and that work in an unsupervised fashion (because manually labelling massive datasets is unfeasible). Semantic hashing [21] is a highly effective class of methods that encode the semantics of a document into a binary vector called a *hash code*, with the property that similar documents have a short Hamming distance between their codes, which is simply the number of differing bits in the codes as efficiently computed by the sum of the XOR operation. For short hash codes of down to a single byte, this provides a very fast way of performing similarity searches [34], while also reducing the storage requirement compared to full text documents.

Originally, work on semantic hashing focused on generating hash codes for a fixed collection [25], but more modern information needs require querying *unseen* documents for retrieving similar documents in the collection. Modern semantic hashing methods are based on machine learning techniques that, once trained, are able to produce the hash code based solely on the document alone. This can be done using techniques similar to Latent Semantic Indexing [33], spectral clustering [29], or two-step approaches of first creating an optimal encoding and then training a classifier to predict this [34]. Recent work has focused on deep learning based methods [4, 5, 23] to create a generative document model. However, none of the methods directly model the end goal of providing an effective similarity search, i.e., being able to accurately rank documents based on their hash codes, but rather just focus solely on generating document representations.

We present a novel unsupervised generative semantic hashing approach, *Ranking based Semantic Hashing* (RBSH) that combines the ideas of a variational autoencoder, via a so-called variational component, together with a ranking component that aims at directly modelling document similarity through the generated hash codes. The objective of the variational component is to maximize the document likelihood of its generated hash code, which is intractable to compute directly, so a variational lower bound is maximized instead. The variational component is modelled via neural networks and learns to sample the hash code from a Bernoulli distribution, thus allowing end-to-end trainability by avoiding a post-processing

step of binarizing the codes. The ranking component aims at learning to rank documents correctly based on their hash codes, and uses weak supervision through an unsupervised document similarity function to obtain pseudo rankings of the original documents, which circumvents the problem of lacking ground truth data in the unsupervised setting. Both components are optimized jointly in a combined neural model, which is designed such that the final model can be used to generate hash codes solely based on a new unseen document, without computing any similarities to documents in the collection. Extensive experimental evaluation on four publicly available datasets against baselines and state-of-the-art methods for semantic hashing, shows that RBSH outperforms all other methods significantly. Similarly to related work [4, 5, 23], the evaluation is performed as a similarity search of the most similar documents via the Hamming distance and measured using precision across hash codes of 8-128 bits. In fact, RBSH outperforms other methods to such a degree, that generally RBSH hash codes perform similarly to state-of-the-art hash codes while using 2-4x less bits, which corresponds to an effective storage reduction of a factor 2-4x.

In summary, we **contribute** a novel generative semantic hashing method, *Ranking based Semantic Hashing* (RBSH), that through weak supervision directly aims to correctly rank generated hash codes, by modelling their relation to weakly labelled similarities between documents in the original space. Experimentally this is shown to significantly outperform all state-of-the-art methods, and most importantly to yield state-of-the-art performance using 2-4x fewer bits than existing methods.

2 RELATED WORK

2.1 Semantic Hashing

Semantic hashing functions provide a way to transform documents to a low dimensional representation consisting of a sequence of bits. These compact bit vectors are an integral part of fast large-scale similarity search in information retrieval [28], as they allow efficient nearest neighbour look-ups using the Hamming distance. Locality Sensitive Hashing (LSH) [6] is a widely known data-independent hashing function with theoretically founded performance guarantees. However, it is general purpose and as such not designed for semantic hashing, hence it empirically performs worse than a broad range of semantic hashing functions [4, 5]. In comparison to LSH, semantic hashing methods employ machine learning based techniques to learn a *data-dependent* hashing function, which has also been denoted as learning to hash [28].

Spectral Hashing (SpH) [29] can be viewed as an extension of spectral clustering [17], and preserves a global similarity structure between documents by creating balanced bit vectors with uncorrelated bits. Laplacian co-hashing (LCH) [33] can be viewed as a version of binarized Latent Semantic Indexing (LSI) [7, 21] that directly optimizes the Hamming space as opposed to the traditional optimization of Latent Semantic Indexing. Thus, LCH aims at preserving document semantics, just as LSI traditionally does for text representations. Self-Taught Hashing (STH) [34] has the objective of preserving the local similarities between samples found via a k -nearest neighbour search. This is done through computing the bit vectors by considering document connectivity, however without learning document features. Thus, the objective of preserving

local similarities contrasts the global similarity preservation of SpH. Interestingly, the aim of our RBSH can be considered as the junction of the aims of STH and SpH: the variational component of RBSH enables the learning of local structures, while the ranking component ensures that the hash codes incorporate both local and global structure. Variational Deep Semantic Hashing (VDSH) [5] is a generative model that aims to improve upon STH by incorporating document features by preserving the semantics of each document using a neural autoencoder architecture, but without considering the neighbourhood around each document. The final bit vector is created using the median method [29] for binarization, which means the model is not end-to-end trainable. Chaidaroon et al. [4] propose a generative model with a similar architecture to VDSH, but in contrast incorporate an average document of the neighbouring documents found via BM25 [20] which can be seen as a type of weak supervision. The model learns to also reconstruct the average neighbourhood document in addition to the original document, which has similarities with STH in the sense that they both aim to preserve local semantic similarities. In contrast, RBSH directly models document similarities based on a weakly supervised ranking through a hinge loss, thus enabling the optimization of both local and global structure. Chaidaroon et al. [4] also propose a model that combines the average neighbourhood documents with the original document when generating the hash code. However this model is very computationally expensive in practice as it requires to find the top- k similar documents online at test time, while not outperforming their original model [4]. NASH [23] proposed an end-to-end trainable generative semantic hashing model that learns the final bit vector directly, without using a second step of binarizing the vectors once they have been generated. This binarization is discrete and thus not differentiable, so a straight-through estimator [2] is used when optimizing the model.

The related work described above has focused on unsupervised text hashing. Direct modelling of the hash code similarities as proposed in this paper has not been explored. For the case of *supervised* image hashing, some existing work has aimed at generating hash codes using ranking strategies from labelled data, e.g., based on linear hash functions [27] and convolutional neural networks [30, 36]. In contrast, our work develops a generative model and utilises weak supervision to circumvent the need for labelled data.

2.2 Weak Supervision

Weak supervision has showed strong results in the IR community [8, 9, 18, 32], by providing a solution for problems with small amounts of labelled data, but large amounts of unlabelled data. While none of these are applied in a problem domain similar to ours, they all show that increased performance can be achieved by utilizing weak labels. Zamani et al. [8] train a neural network end-to-end for ad-hoc retrieval. They empirically show that a neural model trained on weakly labelled data via BM25 is able to generalize and outperform BM25 itself. A similar approach is proposed by Nie et al. [18], who use a multi-level convolutional network architecture, allowing to better differentiate between the abstraction levels needed for different queries and documents. Zamani et al. [32] present a solution for the related problem of query performance

prediction, where multiple weak signals of clarity, commitment, and utility achieve state-of-the-art results.

3 RANKING BASED SEMANTIC HASHING

We first present an overview of our model, Ranking Based Semantic Hashing (RBSH), and then describe in detail the individual parts of the model. RBSH combines the principles of a variational autoencoder with a ranking component using weak supervision and is an unsupervised generative model. For document d , the variational component of RBSH learns a low dimensional binary vector representation $z \in \{0, 1\}^m$, called the hash code, where m is the number of bits in the code. RBSH learns an encoder and decoder function, modelled by neural networks, that are able to encode d to z and z to \hat{d} , respectively, where \hat{d} is an approximation of the original document d . The goal of the encoder-decoder architecture is to reconstruct the original document as well as possible via the hash code. Additionally, we incorporate a ranking component which aims to model the similarity between documents, such that the resulting hash codes are better suited for finding nearest neighbours. Specifically, during training RBSH takes document triplets, (d, d_1, d_2) , as inputs with estimated pairwise similarities, and through weak supervision attempts to correctly predict either d_1 or d_2 as being most similar to d . Training the model on inputs of various similarities (e.g., from the top 200 most similar documents) enables the model to learn both the local and global structure to be used in the hash code generation.

In summary, through the combination of the variational and ranking components the objective of RBSH is to be able to both reconstruct the original document as well as correctly rank the documents based on the produced hash codes. An overview of the model can be seen in Figure 1. In the sections below we describe the generative process of the variational component (Section 3.1), followed by the encoder function (Section 3.2), decoder function (Section 3.3), the ranking component (Section 3.4), and finally the combined model (Section 3.5).

3.1 Variational component

We assume each document d to be represented as a bag-of-words representation of vocabulary size V such that $d \in \mathbb{R}^V$. We denote the set of unique words in document d as \mathcal{W}_d . For each document we sample a binary semantic vector $z \sim p(z)$ where $p(z_i) = p_i^{z_i}(1 - p_i)^{1 - z_i}$, which allows the hash codes to be end-to-end trainable, similarly to Shen et al. [23]. For each bit, p_i corresponds to the probability of sampling a 1 at position i and $(1 - p_i)$ is the probability of sampling a 0. Thus, z is obtained by repeating a Bernoulli trial m times. Using the sampled semantic vector, we consider each word as $w_i \sim p(w_i|f(z))$ and define the document likelihood as follows:

$$p(d|z) = \prod_{j \in \mathcal{W}_d} p(w_j|f(z)) \quad (1)$$

that is, a simple product of word probabilities where the product iterates over all unique words in document d (denoted \mathcal{W}_d). In this setting $f(z)$ is a function that maps the hash code, z , to a latent vector useful for modelling word probabilities.

3.1.1 Variational loss. The first objective of our model is to maximize the document log likelihood:

$$\log p(d) = \log \int_{\{0,1\}^m} p(d|z)p(z)dz \quad (2)$$

However, due to the non-linearity of computing $p(w_j|f(z))$ from Equation 1 this computation is intractable and the variational lower bound [12] is maximized instead:

$$\log p(d) \geq E_Q[\log p(d|z)] - \text{KL}(Q(z|d)||p(z)) \quad (3)$$

where $Q(z|d)$ is a learned approximation of the posterior distribution $p(z|d)$, the computation of which we describe in Section 3.2, and KL is the Kullback-Leibler divergence. Writing this out using the document likelihood we obtain the model’s variational loss:

$$\mathcal{L}_{\text{var}} = E_Q \left[\sum_{j \in \mathcal{W}_d} \log p(w_j|f(z)) \right] - \text{KL}(Q(z|d)||p(z)) \quad (4)$$

where j iterates over all unique words in document d . The purpose of this loss is to maximize the document likelihood under our modelling assumptions, where the E_Q term can be considered the reconstruction loss. The KL divergence acts as a regularizer by penalizing large differences between the approximate posterior distribution and the Bernoulli distribution with equal probability of sampling 0 and 1 ($p = 0.5$), which can be computed in closed form as:

$$\text{KL}(Q(z|d)||p(z)) = Q(d) \log \frac{Q(d)}{p} + (1 - Q(d)) \log \frac{1 - Q(d)}{1 - p} \quad (5)$$

3.2 Encoder function

The approximate posterior distribution $Q(z|d)$ can be considered as the encoder function that transforms the original document representation into its hash code of m bits. We model this using a neural network that outputs the sampling probabilities used for the Bernoulli sampling of the hash code. First, we compute the representation used as input for computing the sampling probabilities:

$$v_1 = \text{ReLU}(W_a(d \odot E_{\text{imp}}) + b_a) \quad (6)$$

$$v_2 = \text{ReLU}(W_b v_1 + b_b) \quad (7)$$

where \odot corresponds to elementwise multiplication, W and b are weight matrices and bias vectors respectively, and E_{imp} is an *importance embedding* that learns a scalar for each word that is used to scale the word level values of the original document representation, and the same embedding is also used in the decoder function. The purpose of this embedding is to scale the original input such that unimportant words have less influence on the hash code generation. We transform the intermediate v_2 representation to a vector of the same size as the hash code, such that the i^{th} entry corresponds to the sampling probability for the i^{th} bit:

$$Q(d) = \sigma(W_m v_2 + b_m) \quad (8)$$

where W_m and b_m have the dimensions corresponding to the code length m , and σ is the sigmoid function used to enforce the values to be within the interval $[0, 1]$, i.e., the range of probability values. The final hash code can then be sampled from the Bernoulli distribution. In practice, this is estimated by a vector $\mu = [\mu_1, \mu_2, \dots, \mu_m]$ of




Figure 1: Model overview

values sampled uniformly at random from the interval $[0, 1]$ and computing each bit value of either 0 or 1 as:

$$z_i = \lceil Q(d)_i - \mu_i \rceil \quad (9)$$

Sampling μ uniformly at random corresponds to a stochastic strategy, as the same $Q(d)$ could result in different hash codes. The opposite deterministic strategy consists of fixing $\mu_i = 0.5$, such that the network always generates the same code for a given document. To encourage exploration during training, the stochastic strategy is chosen, while the deterministic is used for testing. To compute the gradient of the sampled z for back-propagation, we use a straight-through estimator [2].

3.3 Decoder function

The purpose of the decoder function is to reconstruct the original document d given the hash code z . This is computed using the document log likelihood (Equation 1) as the sum of word log probabilities:

$$\begin{aligned} \log p(d|z) &= \sum_{j \in \mathcal{W}_d} \log p(w_j | f(z)) \\ &= \sum_{j \in \mathcal{W}_d} \log \frac{e^{f(z)^T g(E_{\text{word}}(o_j \odot E_{\text{imp}})) + b_j}}{e^{\sum_{i \in \mathcal{W}_{\text{all}}} f(z)^T g(E_{\text{word}}(o_i \odot E_{\text{imp}})) + b_i}} \end{aligned} \quad (10)$$

where the sums iterate over all unique words in document d ; \odot corresponds to elementwise multiplication; o_j is a one-hot-vector with 1 in the j^{th} position and 0 everywhere else; E_{imp} is the same importance embedding as in the encoder function; E_{word} is a word embedding; b is a bias vector; \mathcal{W}_{all} contains all vocabulary words; and the g function will be detailed later. E_{word} is a mapping from a word to a *word embedding* space, such that $\log p(d|z)$ is maximized when the hash code is similar to most words in the document. To this end, the importance embedding assists in reducing the need to be similar to all words, as it learns to reduce the value of unimportant words. The word embedding E_{word} is made by learning a 300 dimensional embedding matrix, and $g(E_{\text{word}}(o_j \odot E_{\text{imp}}))$ corresponds to a transformation through a fully connected linear layer to fit the code length. The choice of 300 dimensions was made to be similar in size to standard GloVe and Word2vec word embeddings [16, 19]. This two-step embedding process was chosen to allow the model to learn a code length-independent embedding initially, such that the underlying word representation is not limited by the code length.

3.3.1 Reduce overfitting through noise injection. We inject noise into the hash code before decoding, which has been shown to reduce overfitting and to improve generalizability in generative models [3, 12, 24]. For semantic hashing applications, this corresponds to observing significantly more *artificial* documents with small perturbations, which is beneficial for reducing overfitting in the reconstruction step. To this end we choose a Gaussian noise model, which is traditionally done for variational autoencoders [12], such that $f(z)$ in Equation 10 is sampled as $f(z) \sim \mathcal{N}(z, \sigma^2 I)$ where I is the identity matrix and σ^2 is the variance. Instead of using a fixed variance, we employ variance annealing, where the variance is reduced over time towards 0. Variance annealing has previously been shown to improve performance for generative models in the image domain [3], as it reduces the uncertainty over time when the model confidence increases. However, the gradient estimate with this noise computation exhibits high variance [12], so we use the reparameterization trick to compute $f(z)$ as:

$$f(z; \sigma^2) = z + \epsilon \sigma, \quad \epsilon \sim \mathcal{N}(0, I) \quad (11)$$

which is based on a single source of normal distributed noise and results in a gradient estimate with lower variance [12].

3.4 Ranking component

The variational loss guides the model towards being able to reconstruct the original document from the hash code, but no hash code similarity is enforced between similar documents. We introduce a ranking component into the model, which ensures that similar documents have a small hash code distance between them. To enable the network to learn the correct document ranking we consider document triplets as inputs, (d, d_1, d_2) with corresponding pairwise similarities of $s_{\{d, d_1\}}$ and $s_{\{d, d_2\}}$. However, in the unsupervised setting we do not have a ground truth annotated ranking of the documents to extract the similarities. To this end, we generate pseudo pairwise similarities between the documents, such that weak supervision can be used to train the network in a supervised fashion.

3.4.1 Estimating pairwise similarities. For estimating pairwise similarities in our setting, one of many traditional ranking functions or document similarity functions could be employed. We assume such a function is chosen such that a similarity between d and d_1 can be computed.

For concreteness, in this paper we choose to compute document similarities using the hash codes generated by Self-Taught Hashing (STH) [34] as this has been shown to perform well for semantic hashing (see Section 4.5). Using the STH hash codes, document similarity is computed based on the Euclidean distance between two hash codes:

$$s_{\{d, d_1\}} = -\|z^{\text{STH}} - z_1^{\text{STH}}\|_2 \quad (12)$$

where z^{STH} corresponds to the STH hash code for document d , such that $s_{\{d, d_1\}}$ is highest when two documents are very similar. We use the k -nearest neighbour algorithm to find the top k most similar documents for each document.

3.4.2 Ranking loss. To train the ranking component we use a modified version of the hinge loss, as the hinge loss has previously been shown to work well for ranking with weak supervision [8]. We first define the following short-hand expressions:

$$\text{sign}_{d, d_1, d_2} = \text{sign}(s_{\{d, d_1\}} - s_{\{d, d_2\}}) \quad (13)$$

$$D_{d, d_1, d_2} = \|z - z_2\|_2^2 - \|z - z_1\|_2^2 \quad (14)$$

such that $\text{sign}_{d, d_1, d_2}$ corresponds to the sign of the estimated pairwise document similarities, and D_{d, d_1, d_2} is the difference between the squared Euclidean distance of the hash codes of the document pairs. Using this we can define our modified hinge loss as the following piece-wise function:

$$\mathcal{L}_{\text{rank}} = \begin{cases} \max(0, \epsilon - \text{sign}_{d, d_1, d_2} D_{d, d_1, d_2}) & s_{\{d, d_1\}} \neq s_{\{d, d_2\}} \\ |D_{d, d_1, d_2}| & \text{otherwise.} \end{cases} \quad (15)$$

where ϵ determines the margin of the hinge loss, which we fix to 1 to allow a small bitwise difference between hash codes of highly similar documents. Traditionally, the hinge loss consists only of the first part of the piece-wise function, but since the similarity estimates are based on distance computations on hash codes, some document pairs will have the same similarity. In that case the pairwise similarities are equal and the loss is simply the absolute value of D_{d, d_1, d_2} , as it should be close to 0.

3.5 Combining variational and ranking components

We train the variational and ranking components simultaneously by minimizing a combined weighted loss from Equation 4 and 15:

$$\mathcal{L} = \alpha \mathcal{L}_{\text{rank}} - E_Q \left[\sum_{j \in \mathcal{W}_d} \log p(w_j | f(z)) \right] + \beta \text{KL}(Q(z|d) || p(z)) \quad (16)$$

where j iterates over all unique words in document d , α is used to scale the ranking loss, β is used to scale the KL divergence of the variational loss, and we keep the unscaled version of the reconstruction part of the variational loss. During training we start with initial weight parameters of 0 and gradually increase the values in order to focus on just being able to reconstruct the input well.

	n	multi-class	num. classes	unique words
20news	18,846	No	20	52,447
TMC	28,596	Yes	22	18,196
Reuters	9,848	Yes	90	16,631
AGnews	127,598	No	4	32,154

Table 1: Dataset statistics

4 EXPERIMENTAL EVALUATION

4.1 Datasets

We use the four publicly available datasets summarized in Table 1. 1) *20 newsgroups*¹ is a dataset of posts from 20 different newsgroups. 2) *TMC*² is a dataset of NASA air traffic reports, where each report is labelled with multiple classes. 3) *Reuters21578*³ is a dataset of news documents from Reuters, where each document is labelled with one or more classes. The Reuters21578 dataset is subsampled such that documents are removed if none of their associated classes are among the 20 most frequent classes. This was done by Chaidaroon and Fang [5] and their subsampled dataset was used by Shen et al. [23]. 4) *AGnews* [35] contains news articles from 4 categories.

The datasets are commonly used in related work [4, 5, 23], but without full details of preprocessing. So, in the following we describe how we preprocess the data. We filter all documents in a dataset by removing *hapax legomena*, as well as words occurring in more than 90% of the documents. In addition, we apply stopword removal using the NLTK stopword list⁴, do not apply any stemming, and use TF-IDF [22] as the document representation.

For each dataset we make a training, validation, and testing split of 80%, 10%, and 10% of the data, respectively. In all experiments the training data is used to train an unsupervised model, the validation data is used for early stopping by monitoring when the validation loss starts increasing, and the results are reported on the testing data.

4.2 Performance metric

The purpose of generating binary hash codes (of equal length) is to use them to obtain fast similarity searches via the Hamming distance, i.e., computing the number of bits where they differ. If two documents are semantically similar, then the generated semantic hash codes should have small Hamming distance between them. To evaluate the effectiveness of a semantic hashing method we treat each testing document as a query and perform a k -nearest-neighbour (k NN) search using the Hamming distance on the hash codes. Similarly to previous work [4, 5, 23], we retrieve the 100 most similar documents and measure the performance on a specific test document as the precision among the 100 retrieved documents (Prec@100). The total performance for a semantic hashing method is then simply the average Prec@100 across all test documents. The used datasets are originally created for text classification, but we can define two documents to be similar if they share at least one class in their labelling, meaning that multiclass documents need not to be of exactly the same classes. This definition of similarity is also used by related work [4, 5, 23].

¹http://scikit-learn.org/0.19/datasets/twenty_newsgroups.html

²<https://catalog.data.gov/dataset/siam-2007-text-mining-competition-dataset>

³<http://www.nltk.org/book/ch02.html>

⁴https://www.nltk.org/nltk_data/

4.3 Baselines

We compare our method against traditional baselines and state-of-the-art semantic hashing methods used in related work as described in Section 2: Spectral Hashing (SpH) [29], Self-Taught Hashing (STH) [34], Laplacian co-hashing (LCH) [33], Variational Deep Semantic Hashing (VDSH) [5], NASH [23], and the neighbourhood recognition model (NbrReg) proposed by Chaidaroon et al. [4]. We tune the hyperparameters of these methods on the validation data as described in their original papers.

4.4 Tuning

For the encoder function (Section 3.2) we use two fully connected layers with 1000 nodes in each layer on all datasets. The network is trained using the ADAM optimizer [11]. We tune the learning rate from the set $\{0.001, 0.0005\}$, where 0.0005 was chosen consistently for 20news and 0.001 on the other datasets. To improve generalization we add Gaussian distributed noise to the hash code before reconstruction in the decoder function, where the variance of the sampled noise distribution is annealed over time. Initially we start with a variance of 1 and reduce it by 10^{-6} every iteration, which we choose conservatively to not reduce it too fast. For the ranking component we use STH [34] to obtain a ranking of the most similar documents for each document in the training and validation set, where we choose every 10^{th} document from the top 200 most similar documents. This choice was made to limit the number of triplets generated for each document as it scales quadratically in the number of similar documents to consider.

When combining the variational and ranking components of our model (Section 3.5), we added a weight parameter on the ranking loss and the KL divergence of the variational loss. We employ a strategy similar to variance annealing in this setting, however in these cases we start at an initial value and increase the weight parameters with every iteration. For the KL divergence we fix the start value at 0 and increase it by 10^{-5} with every iteration. For the ranking loss we tune the models by considering starting values from the set $\{0, 0.5, 1, 1.5\}$ and increase from the set $\{30000^{-1}, 300000^{-1}, 1500000^{-1}, 3000000^{-1}\}$. The code was implemented using the Tensorflow Python library [1] and the experiments were performed on Titan X GPUs.

4.5 Results

The experimental comparison between the methods is summarized in Table 2, where the methods are used to generate hash codes of length $m \in \{8, 16, 32, 64, 128\}$. We highlight the best performing method according to the Prec@100 metric on the testing data. We perform a paired two tailed t-test at the 0.05 level to test for statistical significance on the Prec@100 scores from each test document. We apply a Shapiro-Wilk test at the 0.05 level to test for normality, which is passed for all methods across all code lengths.

4.5.1 Baseline comparison. On all datasets and across all code lengths (number of bits) our proposed Ranking based Semantic Hashing (RBSH) method outperforms both traditional approaches (SpH, STH, and LCH) and more recent neural models (VDSH, NbrReg, and NASH). Generally, we observe a larger performance variation for the traditional methods depending on the dataset compared

to the neural approaches, which are more consistent in their relative performance across the datasets. For example, STH is among the top performing methods on Agnews, but performs among the worst on 20news. This highlights a possible strength of neural approaches for the task of semantic hashing.

Our RBSH consistently outperforms other methods to such a degree, that it generally allows to use hash codes with a factor of 2-4x fewer bits compared to state-of-the-art methods, while keeping the same performance. This provides a notable benefit on large-scale similarity searches, as computing the Hamming distance between two hash codes scales linearly with the code length. Thus, compared to prior work our RBSH enables both a large speed-up as well as a large storage reduction.

4.5.2 Performance versus hash code length. We next consider how performance scales with the hash code length. For all methods 128 bit codes perform better than 8 bit codes, but the performance of scaling from 8 to 128 bits varies. The performance of SpH and STH on Reuters peaks at 32 bit and reduces thereafter, and a similar trend is observed for VDSH on Agnews and TMC. This phenomenon has been observed in prior work [5, 23], and we posit that it is due to longer hash codes being able to more uniquely encode each document, thus resulting in a degree of overfitting. However, generally a longer hash code leads to better performance until the performance flattens after a certain code length, which for most methods happens at 32-64 bits.

4.5.3 Result differences compared to previous work. Comparing our experimental results to results reported in previous work [4, 5, 23], we observe some smaller differences most likely due to preprocessing. Previous work have not fully described the preprocessing steps used, thus to do a complete comparison we had to redo the preprocessing as detailed in Section 4.1.

On 20news and TMC the baseline performance scores we report in this paper are slightly larger for most hash code lengths. The vectorized (i.e., bag-of-words format) Reuters dataset released by the VDSH authors⁵, and also used in the NASH [23] paper, only consisted of 20 (unnamed) classes instead of the reported 90 classes, so these results are not directly comparable.

4.6 Effect of ranking component

To evaluate the influence of the ranking component in RBSH we perform an experiment where the weighting parameter of the ranking loss was set to 0 (thus removing it from the model), and report the results in Table 3. Generally, we observe that on all datasets across all hash code lengths, RBSH outperforms RBSH without the ranking component. However, it is interesting to consider the ranking component’s effect on performance as the hash code length increases. On all datasets we observe the largest improvement on 8 bit hash codes, but then on Reuters, Agnews, and TMC a relatively large performance increase happens that reduces the difference in performance. On 20news the performance difference is even larger at 16 bit than at 8 bit, but as the bit size increases the difference decreases until it is marginal. This highlights that one of the major strengths of RBSH, its performance using short hash codes, can be

⁵<https://github.com/unsuthee/VariationalDeepSemanticHashing/blob/master/dataset/reuters.tfidf.mat>

	20news					Agnews				
	8 bits	16 bits	32 bits	64 bits	128 bits	8 bits	16 bits	32 bits	64 bits	128 bits
SpH [29]	0.0820	0.1319	0.1696	0.2140	0.2435	0.3596	0.5127	0.5447	0.5265	0.5566
STH [34]	0.2695	0.4112	0.5001	0.5193	0.5119	0.6573	0.7909	0.8243	0.8377	0.8378
LCH [33]	0.1286	0.2268	0.4462	0.5752	0.6507	0.7353	0.7584	0.7654	0.7800	0.7879
VDSH [5]	0.3066	0.3746	0.4299	0.4403	0.4388	0.6418	0.6754	0.6845	0.6802	0.6714
NbrReg [4]	0.4267	0.5071	0.5517	0.5827	0.5857	0.4274	0.7213	0.7832	0.7988	0.7976
NASH [23]	0.3537	0.4609	0.5441	0.5913	0.6404	0.7207	0.7839	0.8049	0.8089	0.8142
RBSH	0.5190[▲]	0.6087[▲]	0.6385[▲]	0.6655[▲]	0.6668[▲]	0.8066[▲]	0.8288[▲]	0.8363[▲]	0.8393[▲]	0.8381[▲]
	Reuters					TMC				
	8 bits	16 bits	32 bits	64 bits	128 bits	8 bits	16 bits	32 bits	64 bits	128 bits
SpH [29]	0.4647	0.5250	0.6311	0.5985	0.5880	0.5976	0.6405	0.6701	0.6791	0.6842
STH [34]	0.6981	0.7555	0.8050	0.7984	0.7748	0.6787	0.7218	0.7695	0.7818	0.7797
LCH [33]	0.5619	0.6235	0.6587	0.6610	0.6586	0.6546	0.7028	0.7498	0.7817	0.7948
VDSH [5]	0.6371	0.6686	0.7063	0.7095	0.7129	0.6989	0.7300	0.7416	0.7310	0.7289
NbrReg [4]	0.5849	0.6794	0.6290	0.7273	0.7326	0.7000	0.7012	0.6747	0.7088	0.7862
NASH [23]	0.6202	0.7068	0.7644	0.7798	0.8041	0.6846	0.7323	0.7652	0.7935	0.8078
RBSH	0.7409[▲]	0.7740[▲]	0.8149[▲]	0.8120[▲]	0.8088[▲]	0.7620[▲]	0.7959[▲]	0.8138[▲]	0.8224[▲]	0.8193[▲]

Table 2: Prec@100 with varying bit size. Bold marks the highest score. ▲ shows statistically significant improvements with respect to the best baseline at the 0.05 level using a paired two tailed t-test. A Shapiro-Wilk test at the 0.05 level is used to test for normality.

	20news					Agnews				
	8 bits	16 bits	32 bits	64 bits	128 bits	8 bits	16 bits	32 bits	64 bits	128 bits
RBSH	0.5190	0.6087	0.6385	0.6655	0.6668	0.8066	0.8288	0.8363	0.8393	0.8381
RBSH w/o ranking	<u>0.4482</u>	0.5000	<u>0.6263</u>	<u>0.6641</u>	<u>0.6659</u>	<u>0.7986</u>	<u>0.8244</u>	<u>0.8344</u>	0.8332	0.8306
	Reuters					TMC				
	8 bits	16 bits	32 bits	64 bits	128 bits	8 bits	16 bits	32 bits	64 bits	128 bits
RBSH	0.7409	0.7740	0.8149	0.8120	0.8088	0.7620	0.7959	0.8138	0.8224	0.8193
RBSH w/o ranking	<u>0.7061</u>	<u>0.7701</u>	<u>0.8075</u>	<u>0.8099</u>	<u>0.8081</u>	<u>0.7310</u>	<u>0.7804</u>	<u>0.8040</u>	<u>0.8119</u>	<u>0.8172</u>

Table 3: Effect of including the ranking component. Prec@100 with varying bit size. Bold marks the highest score and underline marks a score better than the best baseline.

partly attributed to the ranking component. This is beneficial for the application of similarity search, as the Hamming distance scales linearly with the number of bits in the hash codes, and can thus provide a notable speed-up while obtaining a similar performance using fewer bits. Additionally, when comparing the performance of RBSH without the ranking component against the baselines in Table 2, then it obtains a better performance in 17 out of 20 cases, thus highlighting the performance of just the variational component.

To further investigate the ranking component effect, as well as RBSH in general, in Section 4.7 we consider word level differences in the learned importance embeddings, as well as relations between inverse document frequency (IDF) and the importance embedding weights for each word. In Section 4.8 we investigate what makes a word difficult to reconstruct (i.e., using the decoder function in Section 3.3), which is done by comparing the word level reconstruction log probabilities to both IDF and the learned importance embedding weights. Finally, in Section 4.9 we do a quantitative comparison of RBSH with and without the ranking component. The comparison is based on a t-SNE [15] dimensionality reduction of the hash codes, such that a visual inspection can be performed. In the following sections we consider 16 and 128 bit hash codes generated on 20news, as these provide the largest and one of the

smallest performance difference of RBSH with and without the ranking component, respectively.

4.7 Investigation of the importance embedding

We posit that the ranking component in RBSH enables the model to better differentiate between the importance of individual words when reconstructing the original document. If we consider the decoder function in Equation 10, then it is maximized when the hash code is similar to most of the importance weighted words, which in the case of equally important words would correspond to a word embedding average. However, if the hash code is short, e.g., 8 bits, then similar documents have a tendency to hash to exactly the same code, as the space of possible codes are considerably smaller than at e.g., 128 bits. This leads to worse generalizability observed on unseen documents when using short hash codes, but the ranking component enables the model to better prioritize which words are the most important. Figure 2 compares the learned importance embedding weights for 16 and 128 bit codes on 20news with and without the ranking component. For 16 bit codes we observe that RBSH without the ranking component tends to estimate a higher importance for most words, and especially for words with an RBSH importance over 0.6. This observation could be explained by the




Figure 2: Visualization of the learned importance embedding for each word with and without using the ranking component of RBSH. The plot is made on 20news with 16 and 128 bit hash codes, and the black diagonal line corresponds to equal importance weights.




Figure 3: Visualization of the learned importance embedding for each word compared to the inverse document frequency (IDF). The plot is made on 20news with 16 and 128 bit hash codes.

ranking component acting as a regularizer, by enabling a direct modelling of which words are important for correctly ranking documents as opposed to just reconstruction. However, as the code length increases this becomes less important as more bits are available to encode more of the occurring words in a document, which is observed from the importance embedding comparison for 128 bits, where the over estimation is only marginal.

Figure 3 shows the importance embedding weights compared to the inverse document frequency (IDF) of each word. For both 16 and 128 bits we observe a similar trend of words with high importance weight that also have a high IDF; however words with a high IDF do not necessarily have a high importance weight. When we consider low importance weights, then the corresponding IDF is more evenly distributed, especially for 16 bit hash codes. For 128 bit we observe that lower importance weights are more often associated with a low IDF. These observations suggest that the model learns to emphasize rare words, as well as words of various rarity that the model deems important for both reconstruction and ranking.

4.8 Investigation of the difficulty of word reconstruction

To better understand what makes a word difficult to reconstruct we study the word level reconstruction log probabilities, i.e., each summand in Equation 10, where a 0 value represents a word that is always possible to reconstruct while a smaller value corresponds to a word more difficult to reconstruct. Figure 4 compares the word




Figure 4: Comparison of the word level reconstruction log probability compared to each word’s inverse document frequency (IDF). The plot is made on 20news with 16 and 128 bit hash codes.




Figure 5: Comparison of the word level reconstruction log probability compared to each word’s learned importance weighting. The plot is made on 20news with 16 and 128 bit hash codes.

level reconstruction log probabilities to each word’s IDF for 16 and 128 bit hash codes. There is no notable difference between the plots, which both show that the model prioritizes being able to reconstruct rare words, while focusing less on words occurring often. This follows our intuition of an ideal semantic representation, as words with a high IDF are usually more informative than those with a low IDF.

Figure 5 shows a comparison similar to above, where the word level reconstruction log probabilities are plotted against the learned importance embedding weights. For both 16 and 128 bit hash codes we observe that words that are difficult to reconstruct (i.e., have a low log probability) are associated with a low importance weight. Words with a low reconstruction log probability are also associated with a low IDF. This shows that the model chooses to ignore often occurring words with low importance weight. When considering words with a reconstruction log probability close to 0, then in the case of 16 bit hash codes the corresponding important weights are very evenly distributed in the entire range. In the case of 128 bit hash codes we observe that words the model reconstructs best have importance weights in the slightly higher end of the spectrum, however for lower log probabilities the two hash code lengths behave similarly. This shows that the model is able to reconstruct many words well irrespectively of their learned importance weight, but words with a high importance weight are always able to be reconstructed well.

4.9 Hash code visualization

In Section 4.7 we argued that the ranking component of RBSH enables the model to better prioritize important words for short hash codes, by directly modelling which words were relevant for




Figure 6: t-SNE [15] visualization of the 16 and 128 bit hash codes from our RBSH with and with the ranking component. 20news was used as the dataset and the same color coding for class labels is used across the plots.

ranking the documents. To further study this we perform a qualitative visualization using t-SNE [15] of 16 and 128 bit hash codes on 20news (see Figure 6), where we do the visualization for RBSH with and without the ranking component. For 16 bit hash codes we observe that RBSH without the ranking component most often creates very tight clusters of documents, corresponding to the fact that many of the produced hash codes are identical. When the ranking component is included the produced hash codes are more varied. This creates larger, more general clusters of similar documents. This leads to better generalizability as the space is better utilized, such that unseen documents are less likely to hash into unknown regions, which would result in poor retrieval performance. When considering the 128 bit hash codes for RBSH with and without the ranking component, we observe that they are highly similar, which was also expected as the Prec@100 performance was almost identical.

5 CONCLUSION

We presented a novel method for unsupervised semantic hashing, *Ranking based Semantic Hashing* (RBSH), which consists of a variational and ranking component. The variational component has similarities with variational autoencoders and learns to encode a input document to a binary hash code, while still being able to reconstruct the original document well. The ranking component is trained on document triplets and learns to correctly rank the documents based on their generated hash codes. To circumvent the need of labelled data, we utilize a weak labeller to estimate the

rankings, and then employ weak supervision to train the model in a supervised fashion. These two components enable the model to encode both local and global structure into the hash code. Experimental results on four publicly available datasets showed that RBSH is able to significantly outperform state-of-the-art semantic hashing methods to such a degree, that RBSH hash codes generally perform similarly to other state-of-the-art hash codes, while using 2-4x fewer bits. This means that RBSH can maintain state-of-the-art performance while allowing a direct storage reduction of a factor 2-4x. Further analysis showed that the ranking component provided performance increases on all code lengths, but especially improved the performance on hash codes of 8-16 bits. Generally, the model analysis also highlighted RBSH’s ability to estimate the importance of rare words for better hash encoding, and that it prioritizes the encoding of rare informative words in its hash code.

Future work includes incorporating multiple weak labellers when generating the hash code ranking, which under certain independence assumptions has been theoretically shown to improve performance of weak supervision [31]. Additionally, it could be interesting to investigate the effect of more expressive encoding functions, such as recurrent or convolutional neural networks, that have been used for image hashing [30, 36].

ACKNOWLEDGMENTS

Partly funded by Innovationsfonden DK, DABAI (5153-00004A).

REFERENCES

- [1] Martin Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. 2016. Tensorflow: a system for large-scale machine learning. In *OSDI*, Vol. 16. 265–283.
- [2] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. 2013. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432* (2013).
- [3] Andrew Brock, Jeff Donahue, and Karen Simonyan. 2019. Large Scale GAN Training for High Fidelity Natural Image Synthesis. In *International Conference on Learning Representations*.
- [4] Suthee Chaidaroon, Travis Ebesu, and Yi Fang. 2018. Deep Semantic Text Hashing with Weak Supervision. *ACM SIGIR Conference on Research and Development in Information Retrieval*, 1109–1112.
- [5] Suthee Chaidaroon and Yi Fang. 2017. Variational deep semantic hashing for text documents. In *ACM SIGIR Conference on Research and Development in Information Retrieval*. 75–84.
- [6] Mayur Datar, Nicole Immorlica, Piotr Indyk, and Vahab S Mirrokni. 2004. Locality-sensitive hashing scheme based on p-stable distributions. In *Proceedings of the twentieth annual symposium on Computational geometry*. ACM, 253–262.
- [7] Scott Deerwester, Susan T Dumais, George W Furnas, Thomas K Landauer, and Richard Harshman. 1990. Indexing by latent semantic analysis. *Journal of the American society for information science* 41, 6 (1990), 391–407.
- [8] Mostafa Dehghani, Hamed Zamani, Aliaksei Severyn, Jaap Kamps, and W Bruce Croft. 2017. Neural ranking models with weak supervision. In *ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 65–74.
- [9] Casper Hansen, Christian Hansen, Stephen Alstrup, Jakob Grue Simonsen, and Christina Lioma. 2019. Neural Check-Worthiness Ranking with Weak Supervision: Finding Sentences for Fact-Checking. In *Companion Proceedings of the 2019 World Wide Web Conference*.
- [10] Monika Henzinger. 2006. Finding near-duplicate web pages: a large-scale evaluation of algorithms. In *ACM SIGIR Conference on Research and Development in Information Retrieval*. 284–291.
- [11] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *International Conference on Learning Representations*.
- [12] Diederik P Kingma and Max Welling. 2014. Auto-encoding variational bayes. In *International Conference on Learning Representations*.
- [13] Yehuda Koren. 2008. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *ACM SIGKDD international conference on Knowledge discovery and data mining*. 426–434.
- [14] Michael S Lew, Nicu Sebe, Chabane Djeraba, and Ramesh Jain. 2006. Content-based multimedia information retrieval: State of the art and challenges. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* 2, 1 (2006), 1–19.
- [15] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9, Nov (2008), 2579–2605.
- [16] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*. 3111–3119.
- [17] Andrew Y Ng, Michael I Jordan, and Yair Weiss. 2002. On spectral clustering: Analysis and an algorithm. In *Advances in neural information processing systems*. 849–856.
- [18] Yifan Nie, Alessandro Sordoni, and Jian-Yun Nie. 2018. Multi-level abstraction convolutional model with weak supervision for information retrieval. In *ACM SIGIR Conference on Research and Development in Information Retrieval*. 985–988.
- [19] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global Vectors for Word Representation. In *Empirical Methods in Natural Language Processing (EMNLP)*. 1532–1543.
- [20] Stephen E Robertson, Steve Walker, Susan Jones, Micheline M Hancock-Beaulieu, Mike Gattford, et al. 1995. Okapi at TREC-3. *Nist Special Publication Sp 109* (1995), 109.
- [21] Ruslan Salakhutdinov and Geoffrey Hinton. 2009. Semantic hashing. *International Journal of Approximate Reasoning* 50, 7 (2009), 969–978.
- [22] Gerard Salton and Christopher Buckley. 1988. Term-weighting approaches in automatic text retrieval. *Information processing & management* 24, 5 (1988), 513–523.
- [23] Dinghan Shen, Qinliang Su, Paidamoyo Chapfuwa, Wenlin Wang, Guoyin Wang, Ricardo Henao, and Lawrence Carin. 2018. NASH: Toward End-to-End Neural Architecture for Generative Semantic Hashing. In *Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, 2041–2050.
- [24] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. 2015. Learning structured output representation using deep conditional generative models. In *Advances in Neural Information Processing Systems*. 3483–3491.
- [25] Benno Stein. 2007. Principles of hash-based text retrieval. In *ACM SIGIR Conference on Research and Development in Information Retrieval*. 527–534.
- [26] Benno Stein, Sven Meyer zu Eissen, and Martin Potthast. 2007. Strategies for retrieving plagiarized documents. In *ACM SIGIR Conference on Research and Development in Information Retrieval*. 825–826.
- [27] Jun Wang, Wei Liu, Andy X Sun, and Yu-Gang Jiang. 2013. Learning hash codes with listwise supervision. In *Proceedings of the IEEE International Conference on Computer Vision*. 3032–3039.
- [28] Jingdong Wang, Ting Zhang, Nicu Sebe, Heng Tao Shen, et al. 2018. A survey on learning to hash. *IEEE transactions on pattern analysis and machine intelligence* 40, 4 (2018), 769–790.
- [29] Yair Weiss, Antonio Torralba, and Rob Fergus. 2009. Spectral hashing. In *Advances in neural information processing systems*. 1753–1760.
- [30] Ting Yao, Fuchen Long, Tao Mei, and Yong Rui. 2016. Deep Semantic-Preserving and Ranking-Based Hashing for Image Retrieval. In *IJCAI*. 3931–3937.
- [31] Hamed Zamani and W Bruce Croft. 2018. On the theory of weak supervision for information retrieval. In *Proceedings of the 2018 ACM SIGIR International Conference on Theory of Information Retrieval*. ACM, 147–154.
- [32] Hamed Zamani, W. Bruce Croft, and J. Shane Culpepper. 2018. Neural Query Performance Prediction Using Weak Supervision from Multiple Signals. In *ACM SIGIR Conference on Research and Development in Information Retrieval*. 105–114.
- [33] Dell Zhang, Jun Wang, Deng Cai, and Jinsong Lu. 2010. Laplacian co-hashing of terms and documents. In *European Conference on Information Retrieval*. Springer, 577–580.
- [34] Dell Zhang, Jun Wang, Deng Cai, and Jinsong Lu. 2010. Self-taught hashing for fast similarity search. In *ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 18–25.
- [35] Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *Advances in neural information processing systems*. 649–657.
- [36] Fang Zhao, Yongzhen Huang, Liang Wang, and Tieniu Tan. 2015. Deep semantic ranking based hashing for multi-label image retrieval. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1556–1564.

Chapter 3

Unsupervised Semantic Hashing with Pairwise Reconstruction

Casper Hansen*, Christian Hansen*, Jakob Grue Simonsen, Stephen Alstrup, Christina Lioma (2020). Unsupervised Semantic Hashing with Pairwise Reconstruction. In SIGIR, pages 2009-2012. [24]. * denotes equal contribution.

Unsupervised Semantic Hashing with Pairwise Reconstruction

Casper Hansen*
University of Copenhagen
c.hansen@di.ku.dk

Christian Hansen*
University of Copenhagen
chrh@di.ku.dk

Jakob Grue Simonsen
University of Copenhagen
simonsen@di.ku.dk

Stephen Alstrup
University of Copenhagen
s.alstrup@di.ku.dk

Christina Lioma
University of Copenhagen
c.lioma@di.ku.dk

ABSTRACT

Semantic Hashing is a popular family of methods for efficient similarity search in large-scale datasets. In Semantic Hashing, documents are encoded as short binary vectors (i.e., hash codes), such that semantic similarity can be efficiently computed using the Hamming distance. Recent state-of-the-art approaches have utilized weak supervision to train better performing hashing models. Inspired by this, we present Semantic Hashing with Pairwise Reconstruction (PairRec), which is a discrete variational autoencoder based hashing model. PairRec first encodes weakly supervised training pairs (a query document and a semantically similar document) into two hash codes, and then learns to reconstruct the same query document from both of these hash codes (i.e., pairwise reconstruction). This pairwise reconstruction enables our model to encode local neighbourhood structures within the hash code directly through the decoder. We experimentally compare PairRec to traditional and state-of-the-art approaches, and obtain significant performance improvements in the task of document similarity search.

KEYWORDS

Semantic Hashing; Variational; Pairwise Reconstruction

ACM Reference Format:

Casper Hansen, Christian Hansen, Jakob Grue Simonsen, Stephen Alstrup, and Christina Lioma. 2020. Unsupervised Semantic Hashing with Pairwise Reconstruction. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '20)*, July 25–30, 2020, Virtual Event, China. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3397271.3401220>

1 INTRODUCTION

Document similarity search is a core information retrieval task, where semantically similar documents are retrieved based on a query document. Large-scale retrieval requires methods that are both effective and efficient, and that can—ideally—be trained in an

unsupervised fashion due to the high cost associated with labeling massive data collections. To this end, Semantic Hashing [11] methods learn to transform objects (e.g., text documents) into short binary vector representations, which are called *hash codes*. The semantic similarity between two documents can then be computed using the Hamming distance, i.e., the sum of differing bits between two hash codes, which can be implemented highly efficiently on hardware due to operating on fixed-length bit strings (real-time retrieval among a billion hash codes [12]). Hash codes are typically the same length as a machine word (32 or 64 bits), thus the storage cost for large document collections is relatively low.

The state-of-the-art on unsupervised semantic hashing uses *weak supervision* in different ways to learn hash codes that better encode the structure of local neighbourhoods around each document. NbrReg [2] used BM25 to associate each document with an aggregation of the most similar neighbourhood documents, where two different decoders are trained to reconstruct the document hash code to both the original *and* aggregated neighbourhood document. However, we argue that using multiple different decoders on a single hash code is ineffective, since each decoder will attempt to enforce (potentially) different semantics, which may harm generalization of the hash code. Additionally, an aggregated neighbourhood document is not a *real* document encountered during retrieval, which means that learning from it can introduce further semantic shift. Recently, RBSH [7] proposed to use weak supervision for incorporating a ranking objective in the model, with the aim of improving the hash codes performance in document ranking tasks. However, RBSH uses two weakly (positively and negatively) labeled documents to generate a ranking triplet, each of which is obtained from a noisy relevance estimate, which may lead to larger inaccuracies when combined.

To address the above problems, we propose to use weak supervision to extract the top-K most similar documents to a given query document, which are split into K pairs, each consisting of the query document and a top-K document. Using an end-to-end discrete variational autoencoder architecture, each document within a pair is encoded to a hash code, and through a single decoder they are both trained in an unsupervised fashion to be able to reconstruct the query document (i.e., they are *pairwise* reconstructed to obtain the query document). In contrast to NbrReg [2], our PairRec aims at learning a more generalizable decoding through a single decoder used on pairs of (non-aggregated) documents, as opposed to using different decoders as done in NbrReg. In contrast to RBSH [7], our PairRec is only based on a single weakly labeled document per sample, thus aiming at reducing the inaccuracies originating from comparing noisy relevance estimates for ranking in RBSH.

*Both authors share the first authorship

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGIR '20, July 25–30, 2020, Virtual Event, China

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8016-4/20/07... \$15.00

<https://doi.org/10.1145/3397271.3401220>




Figure 1: PairRec model overview.

In summary, we **contribute** a novel weakly supervised semantic hashing approach named PairRec, based on our concept of pairwise reconstruction for encoding local neighbourhood structures within the hash code. We experimentally evaluate the effectiveness of PairRec against traditional and state-of-the-art semantic hashing approaches, and show that PairRec obtains significant improvements in the task of document similarity search. In fact, PairRec hash codes generally perform similar or better than the state-of-the-art while using 2-4x fewer bits.

2 RELATED WORK

Early work on semantic hashing used techniques adopted from spectral clustering [14], encapsulating global similarity structures, and later local similarity structures between neighbours found using k-nearest neighbour [16]. Following the popularity of deep learning, VDSH [3] was proposed as a neural model enabling complex encoding of documents, that aimed to learn more descriptive hash codes. Inspired by the benefit of weak supervision in related domains [4, 6], NbrReg [2] was proposed for incorporating aggregated neighbourhood documents in the hash code decoder, for the purpose of incorporating local similarity structure. However, these methods do not learn the hash code in an end-to-end fashion, since they rely on a post-processing rounding stage. To this end, NASH [13] was proposed as an end-to-end trainable variational autoencoder, where bits were sampled according to a learned sample probability vector from a Bernoulli distribution. As a step towards more expressive document encoding, BMSH [5] utilized a Bernoulli mixture prior generative model, but was only able to outperform a simple version of the NASH model, and not consistently outperform the proposed full version. Lastly, RBSH [7] was the first semantic hashing approach that utilized a ranking objective in the model (through sampling semantically similar documents [9]), thus enabling the hash codes to combine both local and global structures for improved retrieval performance. RBSH was able to significantly outperform existing state-of-the-art semantic hashing approaches. Recently, semantic hashing has also been successfully applied to the task of cold-start collaborative filtering, where recent advances enabled a better semantic representation of the items [8].

3 PAIRWISE RECONSTRUCTION BASED HASHING

Pairwise reconstruction based hashing (PairRec) is a discrete variational autoencoder with a pairwise reconstruction loss. Given a

document d , PairRec generates an m -bit hash code $z \in \{0, 1\}^m$ for d , such that two semantically similar documents have low Hamming distance. Specifically, z is sampled by repeating m consecutive Bernoulli trials based on learned sampling probabilities. Given a similarity function, PairRec is trained on pairs of semantically similar documents, and learns to encode local document neighbourhood structures by training to reconstruct one of the documents from both hash codes (i.e., pairwise reconstruction). We first cover the model architecture and then the pairwise reconstruction loss function. Figure 1 shows a model overview.

To compute the hash code z , we let the document likelihood be conditioned on z and define the conditional document likelihood as a product over word probabilities:

$$p(d|z) = \prod_{j \in \mathcal{W}_d} p(w_j|z) \quad (1)$$

where \mathcal{W}_d denotes the set of all unique words in document d . Based on this, the document log likelihood can be found as:

$$\log p(d) = \log \sum_{z \in \{0,1\}^m} p(d|z)p(z) \quad (2)$$

where $p(z)$ is the hash code prior of a Bernoulli distribution with equal probability of sampling 0 and 1. However, maximizing $\log p(d)$ is intractable in practice [10], so instead we maximise the variational lower bound:

$$\log p(d) \geq E_{Q(\cdot|d)} [\log p(d|z)] - \text{KL}(Q(z|d)||p(z)) \quad (3)$$

where $Q(z|d)$ is a learned approximation of the posterior distribution, and KL is the Kullback-Leibler divergence, which has a closed form solution for Bernoulli distributions [13]. Next, we cover our model’s encoder ($Q(z|d)$) and decoder ($p(d|z)$), and subsequently specify the pairwise reconstruction loss.

3.1 Encoder

The approximate posterior $Q(z|d)$ is computed using a feedforward network with two hidden layers with ReLU activations, and a final output layer using a sigmoid activation to get the sampling probability for each bit:

$$Q(z|d) = \text{FF}_\sigma(\text{FF}_{\text{ReLU}}(\text{FF}_{\text{ReLU}}(d \odot e_{\text{imp}}))) \quad (4)$$

where FF denotes a single feed forward layer, \odot is elementwise multiplication, and e_{imp} is a learned word level importance [7]. During training, the bits are Bernoulli sampled according to their sampling probabilities, while the most probable bits are chosen greedily for evaluation. This enables exploration during training, and a deterministic evaluation output. As the sampling is non differentiable, the straight through estimator is used to do back propagation through the sampling [1].

3.2 Decoder

The decoder should reconstruct the original document d . Previous work has shown a single linear projection works well [7, 13] because the hash codes are used for (linear) Hamming distance computations. We compute the word probabilities by a softmax, where the logit for a single word is given by:

$$\text{logit}(w|z) = f(z)^T (E_{\text{word}}(I(w) \odot e_{\text{imp}})) + b_w \quad (5)$$

Table 1: Dataset statistics

	documents	multi-class	classes	unique words
TMC	28,596	Yes	22	18,196
reuters	9,848	Yes	90	16,631
agnews	127,598	No	4	32,154

where $f(z)$ is a noise infused hash code, E_{word} is a word embedding learned during training, $I(w)$ is a one-hot encoding of word w , e_{imp} is the word level importance also used in the encoder, and b_w is a word level bias term. The noise infusion is done by adding Gaussian noise with zero mean and variance σ^2 to the hash code, resulting in lower variance for the gradient estimates [10]. We apply variance annealing to reduce the variance over time while training the model. Thus, the conditional document log likelihood is given by:

$$\log p(d|z) = \sum_{j \in \mathcal{W}_d} \log \frac{e^{\text{logit}(w_j|z)}}{e^{\sum_{i \in \mathcal{W}_{\text{all}}} \text{logit}(w_i|z)}} \quad (6)$$

where \mathcal{W}_{all} is the set of unique words over all documents.

3.3 Pairwise Reconstruction

PairRec assumes access to some similarity function, which given a document d can be used to obtain a set of the K most similar documents \mathcal{D}_d^K . A training pair (d, d_+) is constructed from the document d and a single document sampled from the set, i.e., $d_+ \in \mathcal{D}_d^K$. Using the variational lower bound from Eq. 3, the pairwise reconstruction loss for the pair is defined as:

$$\begin{aligned} \mathcal{L}_{\text{PairRec}} = & -E_{Q(\cdot|d)}[\log p(d|z)] + \beta \text{KL}(Q(z|d)||p(z)) \\ & -E_{Q(\cdot|d_+)}[\log p(d|z_+)] + \beta \text{KL}(Q(z_+|d_+)||p(z_+)) \end{aligned} \quad (7)$$

Note that this is a negation of the variational lower bound because the loss needs to be minimized. The loss consists of two parts: (i) the first part is an ordinary variational lower bound for document d ; (ii) in the second variational lower bound, document d_+ is used in the encoding, while the decoding is of document d . This transfers local neighbourhood structure from the document space into the Hamming space, since z_+ needs to be able to reconstruct the original d . Lastly, the KL divergence is weighed by a tuneable parameter.

4 EXPERIMENTAL EVALUATION

We use 3 publicly available datasets commonly used in related work [3, 7, 13] consisting of TMC, reuters, and agnews (see Table 1). *TMC*¹ is a multi-class dataset of air traffic reports. *reuters*² is a multi-class dataset of news, and filtered such that a document is removed if none of its labels are among the 20 most frequent labels (similarly done by [3, 7, 13]). Lastly, *agnews* [17] is a single-class dataset of news.

We use the preprocessed data provided in [7], where TF-IDF is used as the document representation and words occurring only once are removed, as well as words occurring in more than 90% of the documents. The datasets were split into training, validation, and testing (80%/10%/10%). We use the validation loss to determine when to stop training a model (using early stopping with a patience of 5 epochs).

¹<https://catalog.data.gov/dataset/siam-2007-text-mining-competition-dataset>

²<http://www.nltk.org/book/ch02.html>

4.1 Baselines and Tuning

We compare our PairRec against traditional post-processing rounding approaches (SpH [14], STH [16], and LCH [15]), neural post-processing rounding approaches (VDSH [3] and NbrReg [2]), and neural end-to-end approaches (NASH [13] and RBSH [7]). NbrReg and RBSH both make use of weak supervision as discussed in Section 1. The baselines are tuned as described in their original papers.

In PairRec³, we tune the number of hidden units in each encoder layers across {500, 1000}, and the number of top K reconstruction pairs across {1, 2, 5, 10, 25, 50, 100, 150, 200}. For obtaining the reconstruction pairs, we generate 64 bit STH [16] hash codes and retrieve the top K most semantically similar documents (STH was also used by RBSH [7]). For the KL divergence, we tune β from {0, 0.01, 0.1}. Note that when $\beta = 0$ is chosen, it corresponds to removing the regularizing KL divergence from the loss. For the variance annealing, we use an initial value of 1 and reduce it by 10^{-6} every iteration (as done in [7]). Lastly, we use the Adam optimizer with a learning rate of 0.0005.

4.2 Evaluation Setup

Following related work [2, 3, 7, 13], we evaluate the semantic hashing approaches based on their top 100 retrieval performance using Prec@100 based on the Hamming distance. Given a query document, we define a retrieved document to be relevant if it shares at least one label with the query document (to ensure that we can accommodate the multi-class datasets, where each document may have one or more associated labels).

4.3 Results

We generate hash codes of {8, 16, 32, 64, 128} bits and report Prec@100 in Table 2. The best performing method for each dataset and bit size is highlighted in bold, and statistically significant improvements (0.05 level) using a two tailed paired t-test are indicated by \blacktriangle .

Our PairRec method consistently outperforms all the traditional and state-of-the-art approaches across all datasets on all bit sizes. RBSH, which also utilizes weak supervision for generating ranking triplets, consistently obtains the second best scores, indicating the benefit of weak supervision for semantic hashing. While NbrReg also makes use of weak supervision (for creating aggregated neighbourhood documents), it performs worse than both NASH and RBSH, but generally better than VDSH, to which its architecture is most similar to. The absolute Prec@100 increases depend on dataset and bit size, but overall PairRec improves state-of-the-art by 1-4%, which correspondingly enables PairRec hash codes to generally perform better or similar to state-of-the-art hash codes with 2-4x more bits.

4.4 Impact of Pairwise Reconstruction

The primary novelty of PairRec is the introduction of pairwise reconstruction. We study the impact of (i) the performance gain obtained by the pairwise reconstruction, and (ii) the performance variance across a varying number of document pairs.

³We make our code available at <https://github.com/casperhansen/PairRec>.

Table 2: Prec@100 with different bit sizes. Bold numbers highlights the highest scores, and [▲] represents statistically significant improvements over RBSH (the best baseline) at the 0.05 level using a two tailed paired t-test.

	Agnews					Reuters					TMC				
	8	16	32	64	128	8	16	32	64	128	8	16	32	64	128
SpH [14]	.3596	.5127	.5447	.5265	.5566	.4647	.5250	.6311	.5985	.5880	.5976	.6405	.6701	.6791	.6842
STH [16]	.6573	.7909	.8243	.8377	.8378	.6981	.7555	.8050	.7984	.7748	.6787	.7218	.7695	.7818	.7797
LCH [15]	.7353	.7584	.7654	.7800	.7879	.5619	.6235	.6587	.6610	.6586	.6546	.7028	.7498	.7817	.7948
VDSH [3]	.6418	.6754	.6845	.6802	.6714	.6371	.6686	.7063	.7095	.7129	.6989	.7300	.7416	.7310	.7289
NbrReg [2]	.4274	.7213	.7832	.7988	.7976	.5849	.6794	.6290	.7273	.7326	.7000	.7012	.6747	.7088	.7862
NASH [13]	.7207	.7839	.8049	.8089	.8142	.6202	.7068	.7644	.7798	.8041	.6846	.7323	.7652	.7935	.8078
RBSH [7]	.8066	.8288	.8363	.8393	.8381	.7409	.7740	.8149	.8120	.8088	.7620	.7959	.8138	.8224	.8193
PairRec (ours)	.8119[▲]	.8354[▲]	.8452[▲]	.8492[▲]	.8498[▲]	.7502[▲]	.8028[▲]	.8268[▲]	.8329[▲]	.8468[▲]	.7656[▲]	.7991[▲]	.8239[▲]	.8280[▲]	.8303[▲]




Figure 2: PairRec with and without pairwise reconstruction.




Figure 3: 64 bit PairRec while varying the top K.

Performance gain by pairwise reconstruction. We compute the Prec@100 with and without the pairwise reconstruction and plot the scores in Figure 2. The largest improvements occur for 64-128 bit on the reuters dataset, but across all datasets and bit sizes, pairwise reconstruction obtains consistent improvements. In comparison, the original RBSH paper [7] also did an ablation with and without weak supervision, but found their improvements to be primarily isolated to 8-16 bits. This further highlights the benefit of using a single weakly supervised document, rather than combining multiple sources for generating ranking triplets as done in RBSH.

Performance variance across number of pairs. We now investigate the impact of the choice of the number of pairs. We fix the bit size to 64 and plot the Prec@100 for all datasets using $\{0, 1, 5, 10, 25, 50, 100, 150, 200\}$ pairs, where 0 corresponds to no pairwise reconstruction. The optimal values for agnews, reuters, and TMC are 100, 25, and 25, respectively. Interestingly, Prec@100 drops after 25 pairs on reuters, which most likely is due to a combination of its small dataset size and high number of classes, corresponding to pairs from top 50 and above no longer being sufficiently semantically similar to the original document. In contrast, for TMC and agnews, we observe no significant performance drop as the number of pairs is increased. In all cases, we note that the optimal value of pairs is also identified by the model parameter configuration with the minimum loss.

5 CONCLUSION

Inspired by recent advances in semantic hashing using weak supervision, we presented a novel semantic hashing approach with

pairwise reconstruction (PairRec). PairRec is a discrete variational autoencoder trained on semantically similar document pairs (obtained through weak supervision), where the model is trained such that the hash codes from both pairwise documents reconstruct the same document. We denote this type of reconstruction as *pairwise reconstruction*; it enables PairRec to encode local neighbourhood structures within the hash code. In an experimental comparison, PairRec was shown to consistently outperform existing state-of-the-art semantic hashing approaches. These improvements generally enable PairRec hash codes to use 2-4x fewer bits than state-of-the-art hash codes while achieving the same or better retrieval performance.

REFERENCES

- [1] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. 2013. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432* (2013).
- [2] Suthee Chaidaroon, Travis Ebesu, and Yi Fang. 2018. Deep Semantic Text Hashing with Weak Supervision. *SIGIR*, 1109–1112.
- [3] Suthee Chaidaroon and Yi Fang. 2017. Variational deep semantic hashing for text documents. In *SIGIR*, 75–84.
- [4] Mostafa Dehghani, Hamed Zamani, Aliaksei Severyn, Jaap Kamps, and W Bruce Croft. 2017. Neural ranking models with weak supervision. In *SIGIR*. ACM, 65–74.
- [5] Wei Dong, Qinliang Su, Dinghan Shen, and Changyou Chen. 2019. Document Hashing with Mixture-Prior Generative Models. In *EMNLP*. 5226–5235.
- [6] Casper Hansen, Christian Hansen, Stephen Alstrup, Jakob Grue Simonsen, and Christina Lioma. 2019. Neural Check-Worthiness Ranking with Weak Supervision: Finding Sentences for Fact-Checking. In *Companion Proceedings of WWW*. 994–1000.
- [7] Casper Hansen, Christian Hansen, Jakob Grue Simonsen, Stephen Alstrup, and Christina Lioma. 2019. Unsupervised Neural Generative Semantic Hashing. In *SIGIR*. 735–744.
- [8] Casper Hansen, Christian Hansen, Jakob Grue Simonsen, Stephen Alstrup, and Christina Lioma. 2020. Content-aware Neural Hashing for Cold-start Recommendation. In *SIGIR*, in press.
- [9] Casper Hansen, Christian Hansen, Jakob Grue Simonsen, and Christina Lioma. 2019. Neural weakly supervised fact check-worthiness detection with contrastive sampling-based ranking loss. In *CLEF-2019 CheckThat! Lab*.
- [10] Diederik P Kingma and Max Welling. 2014. Auto-encoding variational bayes. In *ICLR*.
- [11] Ruslan Salakhutdinov and Geoffrey Hinton. 2009. Semantic hashing. *International Journal of Approximate Reasoning* 50, 7 (2009), 969–978.
- [12] Ying Shan, Jian Jiao, Jie Zhu, and JC Mao. 2018. Recurrent binary embedding for gpu-enabled exhaustive retrieval from billion-scale semantic vectors. In *KDD*. 2170–2179.
- [13] Dinghan Shen, Qinliang Su, Paidamoyo Chapfuwa, Wenlin Wang, Guoyin Wang, Ricardo Henao, and Lawrence Carin. 2018. NASH: Toward End-to-End Neural Architecture for Generative Semantic Hashing. In *ACL*. 2041–2050.
- [14] Yair Weiss, Antonio Torralba, and Rob Fergus. 2009. Spectral hashing. In *NeurIPS*. 1753–1760.
- [15] Dell Zhang, Jun Wang, Deng Cai, and Jinsong Lu. 2010. Laplacian co-hashing of terms and documents. In *ECIR*. Springer, 577–580.
- [16] Dell Zhang, Jun Wang, Deng Cai, and Jinsong Lu. 2010. Self-taught hashing for fast similarity search. In *SIGIR*. ACM, 18–25.
- [17] Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *NeurIPS*. 649–657.

Chapter 4

Unsupervised Multi-Index Semantic Hashing

Christian Hansen*, Casper Hansen*, Jakob Grue Simonsen, Stephen Alstrup, Christina Lioma (2021). Unsupervised Multi-Index Semantic Hashing. In WWW, in press. [20]. * denotes equal contribution.

Unsupervised Multi-Index Semantic Hashing

Christian Hansen*
University of Copenhagen
chrh@di.ku.dk

Casper Hansen*
University of Copenhagen
c.hansen@di.ku.dk

Jakob Grue Simonsen
University of Copenhagen
simonsen@di.ku.dk

Stephen Alstrup
University of Copenhagen
s.alstrup@di.ku.dk

Christina Lioma
University of Copenhagen
c.lioma@di.ku.dk

ABSTRACT

Semantic hashing represents documents as compact binary vectors (hash codes) and allows both efficient and effective similarity search in large-scale information retrieval. The state of the art has primarily focused on learning hash codes that improve similarity search effectiveness, while assuming a brute-force linear scan strategy for searching over all the hash codes, even though much faster alternatives exist. One such alternative is multi-index hashing, an approach that constructs a smaller candidate set to search over, which depending on the distribution of the hash codes can lead to sub-linear search time. In this work, we propose Multi-Index Semantic Hashing (MISH), an unsupervised hashing model that learns hash codes that are both effective and highly efficient by being optimized for multi-index hashing. We derive novel training objectives, which enable to learn hash codes that reduce the candidate sets produced by multi-index hashing, while being end-to-end trainable. In fact, our proposed training objectives are model agnostic, i.e., not tied to how the hash codes are generated specifically in MISH, and are straight-forward to include in existing and future semantic hashing models. We experimentally compare MISH to state-of-the-art semantic hashing baselines in the task of document similarity search. We find that even though multi-index hashing also improves the efficiency of the baselines compared to a linear scan, they are still upwards of 33% slower than MISH, while MISH is still able to obtain state-of-the-art effectiveness.

KEYWORDS

Semantic hashing; multi-index hashing; similarity search

ACM Reference Format:

Christian Hansen, Casper Hansen, Jakob Grue Simonsen, Stephen Alstrup, and Christina Lioma. 2021. Unsupervised Multi-Index Semantic Hashing. In *Proceedings of the Web Conference 2021 (WWW '21)*, April 19–23, 2021, Ljubljana, Slovenia. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3442381.3450014>

*Both authors share the first authorship.

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW '21, April 19–23, 2021, Ljubljana, Slovenia

© 2021 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC-BY 4.0 License.

ACM ISBN 978-1-4503-8312-7/21/04.

<https://doi.org/10.1145/3442381.3450014>

1 INTRODUCTION

Similarity search is a fundamental information retrieval task that aims at finding items similar to a given query. Efficient and effective similarity search is essential for a multitude of retrieval tasks such as collaborative filtering, content-based retrieval, and document search [28, 29]. Semantic Hashing [25] methods enable very efficient search by learning to represent documents (or other types of data objects) as compact bit vectors called *hash codes*, where the Hamming distance is used as the distance metric between hash codes. In this setting, similarity search is expressed as either radius search (finding all hash codes with a specified maximum Hamming distance), or as k-nearest neighbour (kNN) search by incrementally increasing the search radius until the Hamming distance to the k^{th} document is equal to the search radius. Early work on semantic hashing [31, 32] was inspired by techniques similar to spectral clustering [20] and latent semantic indexing [5], whereas modern approaches use deep learning techniques, typically unsupervised autoencoder architectures where hash codes are optimized by learning to reconstruct their original document representations [2, 3, 6, 8, 10, 27]. This line of work has led to extensive improvements of the effectiveness of document similarity search, but has had a lesser focus on efficiency, as it uses a brute-force linear scan of all the hash codes. While the highly efficient Hamming distance does enable large-scale search using linear scans [26], significantly faster alternatives exist. Multi-index hashing [7, 22, 23] is such an alternative, which, depending on a query hash code, can enable sub-linear search time by constructing a smaller set of candidate hash codes to search over. Hash codes can be used extremely efficiently as direct indices into a hash table for finding exact matches, however when doing radius search the number of such hash table lookups grows exponentially. Multi-index hashing is based on the observation that by splitting the hash codes into m substrings and building a hash table per substring, the exponential growth of the number of lookups can be significantly reduced. In practice, the efficiency of multi-index hashing is heavily dependent on the distribution of the hash codes, and most particularly their substrings, which affects the size of the constructed candidate set for a given query hash code. However, no existing semantic hashing methods consider this aspect, which we experimentally verify limits their efficiency.

To address the above efficiency problem, we **contribute** Multi-Index Semantic Hashing (MISH), an unsupervised semantic hashing model that generates hash codes that are both effective and highly efficient through being optimized for multi-index hashing. We identify two key hash code properties for improving multi-index




Figure 1: Method comparison relative to our proposed MISH, averaged over all used datasets. We plot each method as a point regarding its mean speedup of multi-index hashing compared to a linear scan, as well as mean prec@100, relative to MISH.

hashing efficiency, related to limiting the size of the candidate set produced by multi-index hashing. We operationalize these into two novel model agnostic training objectives that effectively reduce the number of hash codes per hash table lookup, while also limiting the necessary search radius for kNN search. These new objectives are fully differentiable and enable training MISH in an end-to-end fashion, and thus enable learning hash codes highly suited for multi-index hashing. We experimentally compare MISH to state-of-the-art semantic hashing baselines in the task of document similarity search. We evaluate their efficiency by comparing the speedup obtained by multi-index hashing over a linear scan, and find that on average the baselines are upwards of 33% slower than MISH. Even though MISH enables large efficiency gains, it is still able to obtain state-of-the-art effectiveness—summarized in Figure 1—where we plot the mean multi-index hashing speedup and mean prec@100 for each method relative to our MISH. Furthermore, we find that MISH can be tuned to enable even larger efficiency improvements at the cost of a slight reduction in effectiveness.

2 RELATED WORK

The problem of nearest neighbour search, also known as similarity search or proximity search, aims at finding the nearest item to a given query using a given distance measure. An efficient way of doing this is by using compact bit vectors (hash codes) that have low storage requirements and enable fast search through the use of the Hamming distance. Locality Sensitive Hashing (LSH) [4] is well-known type of hashing methods with strong theoretical guarantees [4, 15, 29]. However, these types of methods are *data-independent*, and thus unable to capture the semantics of a document (or other types of data objects such as images). In contrast, semantic hashing [25] based methods are *data-dependent*, and aim to learn hash codes such that nearest neighbour search on the hash codes leads to a search result similar to nearest neighbour search on the original document space [29].

2.1 Semantic hashing

Early work on semantic hashing focused on Spectral Hashing (SpH) [31], an approach inspired by spectral clustering [20] that aims at

learning hash codes that preserve global similarity structures of the original documents. Laplacian co-hashing (LCH) [32] learns hash codes representing document semantics through a decomposition similar to latent semantic indexing [5]. Similarly to SpH, Graph Hashing [17] uses a graph representation for learning to capture the underlying global structure. Self-Taught Hashing (STH) [33] contrasts prior work by learning to preserve the *local* structures between samples identified by an initial kNN search in the original document space. The prior work above has primarily been solved as relaxed optimization problems, while later work has utilized deep learning for better capturing document semantics. Variational Deep Semantic Hashing (VDSH) [3], the first work in this direction, uses a more complex encoding of documents through a variational autoencoder architecture (this has since become the primary architecture in subsequent work). Similarly to STH, the authors of VDSH later expanded their model (now named NbrReg) [2] by including a loss function forcing the hash codes to be able to reconstruct unique words occurring in both the document and its neighbours in the original document space (found using BM25 [24]). While both VDSH and NbrReg improved effectiveness, they share the problem of using a post-hoc rounding of learned real-valued vectors, rather than learning the hash codes end-to-end. To fix this, NASH [27] proposed to learn the hash codes end-to-end through learning to sample the bits according to a Bernoulli distribution, which reduced the quantization errors compared to a rounding approach. Based on the same principle, BMSH [6] uses a Bernoulli mixture prior, but only manages to outperform a simple version of NASH, rather than consistently outperform the full NASH model. Similarly to NbrReg and STH, recent state-of-the-art approaches have incorporated neighbourhood knowledge: RBSH [8] incorporates a ranking-based objective, while PairRec [10] uses a pairwise reconstruction loss. The pairwise reconstruction loss is also used in our proposed MISH, and forces two hash codes of semantically similar documents to be able to reconstruct the unique words occurring in both documents, thus directly enabling the encoding of neighbourhood information into the hash codes.

2.2 Semantic hashing efficiency

The semantic hashing approaches above have led to substantial improvements in effectiveness, but they all use a brute-force linear scan for doing similarity search. While this is fast due to the high efficiency of the Hamming distance, hash codes were originally developed to be used as direct indices into a hash table [21, 25, 31], as to avoid a linear scan on a dataset of potential massive size. Through a hash table, finding exact hash code matches only requires a single lookup, but when varying the search radius in a similarity search (e.g., for performing kNN search) it leads to an exponentially increasing number of lookups. To fix this, multi-index hashing [7, 22, 23] has been explored, which enables sub-linear search time by building hash tables on substrings of the original hash codes (see Section 3.2 for a detailed description), and has been used for fast kNN search in hashing-based approaches related to collaborative filtering [9, 11, 12, 16, 34], knowledge graph search [30], and video hashing [35]. However, none of these approaches optimize the hash codes towards improving their multi-index hashing efficiency, but rather simply apply it on already learned hash codes. In contrast, our

proposed MISH is designed to directly learn hash codes suited for multi-index hashing in an end-to-end fashion, which significantly improves efficiency.

3 PRELIMINARIES

3.1 Hamming distance

Given a document $d \in \mathcal{D}$, let $z_d \in \{-1, 1\}^n$ be its associated bit string of n bits, called a *hash code*. The Hamming distance between two hash codes is defined as the number of differing bits between the codes:

$$d_H(z_d, z_{d'}) = \sum_{i=1}^n 1_{z_{d,i} \neq z_{d',i}} = \text{SUM}(z_d \text{ XOR } z_{d'}) \quad (1)$$

where the summation can be computed efficiently due to the *popcnt* instruction that counts the number of bits set to one within a machine word. Due to the efficiency of the Hamming distance, representing documents as hash codes enables both efficient radius search (retrieving all hash codes with a maximum Hamming distance of r to a given hash code), as well as kNN search. Specifically, kNN is performed through radius search by incrementally increasing the search radius up until the distance to the k^{th} hash code is equal to the search radius.

3.2 Multi-index hashing

Norouzi et al. [22, 23] propose a multi-index hashing strategy for performing exact radius and kNN search in the Hamming space. The aim of multi-index hashing is to build a candidate set C of hash codes, significantly smaller than the full document collection, $|C| \ll |\mathcal{D}|$. Given a query hash code z , it then suffices to compute the Hamming distances between z and the hash codes within C , rather than every hash code in the full collection \mathcal{D} . If the size of C is sufficiently small, and can be constructed efficiently, this leads to a sub-linear runtime compared to computing all possible Hamming distances.

Multi-index hashing is an efficient and easy to implement algorithm for building the candidate set C , where each code $z \in \mathcal{D}$ is split into m disjoint substrings, $z = [z^1, z^2, \dots, z^m]^1$. It now follows by the pigeonhole principle that if two codes z and z' are within radius r , i.e., $d_H(z, z') \leq r$, there exists at least one substring where the distance between the two codes is at most $r^* = \lfloor \frac{r}{m} \rfloor$. More specifically, if we assume some arbitrary fixed ordering of the substrings, and write the search radius as $r = r^*m + a$, $a < m$, one substring will have distance at most r^* in the first $a + 1$ substrings, or distance $r^* - 1$ in the remaining $m - a - 1$ substrings. Thus, the candidate set can be constructed by finding all hash codes where the distance within a substring is at most r^* for the first $a + 1$ substrings, or at most $r^* - 1$ for the remaining $m - a - 1$ substrings. For ease of notation, we will denote the substring search radius for substring i as r_i^* .

3.2.1 Efficient candidate set construction. Multi-index hashing uses hash tables to construct the candidate set efficiently. It constructs m hash tables, one for each substring, where the integer value of the substring is used as a key into the hash table, which then maps to all

¹For ease of notation we will assume the substrings have the same length, but it is not a requirement.

documents containing the same substring. Given substring z^i with $\frac{n}{m}$ bits, finding exact matches would require only a single lookup, but the number of lookups for radius search scales exponentially with the substring radius r_i^* as $\sum_{r'=0}^{r_i^*} (\frac{n}{m})^{r'}$. However, the exponential growth is significantly suppressed through fixing $m > 1$, which reduces both the base (through the substring length) and exponent (through the substring search radius), thus making it feasible to run in practice.

Performing radius search on the hash codes can then be done in a straight-forward fashion, by searching within each substring using their associated hash table, and then taking the union over the documents found for each substring search. Note that for kNN search, incrementally increasing the search radius from r to $r + 1$ only changes the search radius within a single substring, and this procedure can therefore be done very efficiently by building the candidate set incrementally as r is increased.

3.2.2 Hash code properties for efficient multi-index hashing. The efficiency of multi-index hashing depends heavily on the properties of the hash codes and how they are distributed in the Hamming space. The computational cost is dominated by the cost of sorting the candidate set according to the Hamming distance to the query hash code, such that the largest speedups are obtained when the candidate set size is small. Focusing on kNN search, the candidate set size is controlled by two factors:

Documents per hash table lookup Given a query hash code, the hash codes should be distributed such that the documents added to the candidate set are likely to appear among the top k documents with the least Hamming distance to the query hash code. To achieve this, the hash codes should be generated such that two hash codes with a low substring Hamming distance also have a low Hamming distance between the entire hash codes.

Search radius for kNN Given a query hash code, the search radius for kNN search is determined by the Hamming distance to the k^{th} closest hash code, which is unknown at query time. Since the number of hash table lookups increases exponentially with the substring search radius, the hash codes should be distributed such that the Hamming distance to the k^{th} document is kept low to limit the exponential growth (corresponding to substring distance less than 2).

Based on these factors it follows that different sets of hash codes for the same dataset can potentially have highly varying search *efficiency* without necessarily affecting the search *effectiveness* of the hash codes. To ensure that learned hash codes enable both efficient *and* effective search, the learning procedure must reflect both of these as part of the training objective. In the next section, we present how such codes can be learned for semantic hashing.

4 MULTI-INDEX SEMANTIC HASHING

We present Multi-Index Semantic Hashing (MISH), a semantic hashing model for unsupervised semantic hashing, which learns to generate hash codes that enable *both* effective and efficient search through being optimized for multi-index hashing. For a document $d \in \mathcal{D}$, MISH learns to generate an n -bit hash code $z_d \in \{-1, 1\}^n$ that represents its semantics, such that two semantically similar

documents have a low Hamming distance between them. MISH consists of a component learning to encode document semantics, and two novel components that ensure the learned hash codes are well suited for multi-index hashing (see Section 3.2), based on the hash code properties discussed in Section 3.2.2. These two components are in fact model agnostic, i.e., not tied to how the hash codes are encoded, so they are straight-forward to include in future work for improving efficiency. Below is an overview of the components:

Semantic encoding MISH is based on a variational autoencoder architecture, where the encoder learns to generate the semantic hash code according to repeating n Bernoulli trials, while a decoder learns to be able to reconstruct the original document from the generated hash code. We choose to use the pairwise reconstruction loss proposed in the state-of-the-art PairRec [10] model to ensure that document semantics are well captured within the hash codes.

Reducing the number of documents per hash table lookup

Given a query hash code z_q , during training we sample another hash code z_s with a low Hamming distance within one of its substrings, but a high Hamming distance using the full hash code, which can be considered a false positive match. We derive an objective that maximises the Hamming distance between the particular substrings of z_q and z_s , thus effectively pushing the substrings of z_q and z_s apart, reducing the number of such false positive matches in the hash table lookup.

Control the search radius for kNN Given a query hash code z_q , during training we sample a hash code z_r with $d_H(z_q, z_r) = r$, where r is the Hamming distance at the top k^{th} position in a kNN search from z_q . In case r is too large, which leads to a large number of hash table lookups, r is reduced through an objective that minimizes the Hamming distance between z_q and z_r , thus effectively pushing the top k hash codes closer together.

In the following sections we present each component individually, and describe how they are jointly optimized for learning the hash codes in an end-to-end fashion.

4.1 Semantic encoding

We use a variational autoencoder to learn a document encoder that generates a hash code capturing the document’s semantics, as well as encoding the local neighbourhood structure of encoded document. This is done by training the codes such that a hash code z should be able to reconstruct not only the original document d , but also documents in the neighborhood of d defined by an appropriate similarity function.

To learn the hash codes, we compute the log likelihood of document $d \in \mathcal{D}$ conditioned on its code z as a sum of word likelihoods, which needs to be maximized:

$$\log p(d|z) = \sum_{j \in \mathcal{W}_d} \log p(w_j|z) \quad (2)$$

where $p(z)$ is sampled by repeating n Bernoulli trials and \mathcal{W}_d is the set of unique words in document d . However, due to the size of the Hamming space, the above is intractable to compute in practice, so

the variational lower bound [14] is maximized instead:

$$\log p(d) \geq E_{Q(\cdot|d)} [\log p(d|z)] - \text{KL}(Q(z|d)||p(z)) \quad (3)$$

where $Q(z|d)$ is a learned approximation of $p(z)$ that functions as the decoder, and KL is the Kullback-Leibler divergence. In the text below, we first describe the encoder ($Q(z|d)$), then the decoder ($p(d|z)$), and lastly the loss function.

4.1.1 Encoder. The encoder is a feed forward network, with two hidden layers using ReLU activation units, followed by a final output layer using a sigmoid activation function, to get the bitwise sampling probabilities:

$$Q(z|d) = \text{FF}_\sigma(\text{FF}_{\text{ReLU}}(\text{FF}_{\text{ReLU}}(d \odot e_{\text{imp}}))) \quad (4)$$

where FF denotes a feed forward layer, and e_{imp} is a learned word level importance embedding. The purpose of the importance embedding is to scale each word of the document representation, such that unimportant words have less influence on generating the hash codes. During training, the bits are sampled according to the bitwise sampling probabilities, while being chosen deterministically for evaluation (choosing the most probable bit value without sampling). To make the sampling differentiable, we employ the straight-through estimator [1].

4.1.2 Decoder. The decoder, $\log p(d|z)$, is defined as maximizing the log likelihood of each word in document d :

$$\log p(d|z) = \sum_{j \in \mathcal{W}_d} \log \frac{e^{\text{logit}(w_j|z)}}{e^{\sum_{i \in \mathcal{W}_{\text{all}}} \text{logit}(w_i|z)}} \quad (5)$$

where $\text{logit}(w|z)$ is the logit for word w_j and \mathcal{W}_{all} are all the words in the corpus. The logit for each word is computed as:

$$\text{logit}(w|z) = f(z)^T (E_{\text{word}}(I(w) \odot e_{\text{imp}})) + b_w \quad (6)$$

where $f(z)$ is a noise-infused hash code with added Gaussian noise (zero mean and a parameterized variance σ^2), which is annealed during training and results in lower variance for the gradient estimates [14]. E_{word} is a word embedding learned during training, $I(w)$ denotes a one-hot encoding of word w , and b_w is a bias term.

4.1.3 Semantic encoder loss. To make the loss function aware of the local neighbourhood structure around a given document, we use pairwise reconstruction as proposed by Hansen et al. [10]. To this end, we use a similarity function independent of the learned hash codes to compute a set of the p most semantically similar documents in the neighbourhood around document d , denoted as \mathcal{N}_d^p . For each $d_+ \in \mathcal{N}_d^p$, we construct (d_q, d_+) with corresponding hash codes (z_d, z_+) and define the loss function based on the variational lower bound from Eq. 3 as:

$$\begin{aligned} \mathcal{L}_{\text{semantic}} = & -E_{Q(\cdot|d_q)} [\log p(d_q|z_q)] + \beta \text{KL}(Q(z_q|d_q)||p(z_q)) \\ & -E_{Q(\cdot|d_+)} [\log p(d_+|z_+)] + \beta \text{KL}(Q(z_+|d_+)||p(z_+)) \end{aligned} \quad (7)$$

As the hash codes, z_q and z_+ both have to reconstruct document d_q (known as *pairwise reconstruction*) the hash codes are forced to not only encode their associated document, but also the local neighbourhood \mathcal{N}_d^p as a whole.

4.2 Reduce the number of documents per hash table lookup

The candidate set estimated by multi-index hashing can be reduced by limiting the number of documents added by each hash table lookup. Specifically, we are interested in limiting the number of false positive matches, i.e., candidate documents added due to a low substring Hamming distance, but where the Hamming distance on the full hash code is above the search radius. Given z_q , a substring i , and the top k search radii r and r_i^* , we sample a hash code z_s as follows:

$$z_s = \operatorname{argmax}_{z_j} d_H(z_q, z_j) \cdot 1_{[d_H(z_q^i, z_j^i) \leq r_i^*]} \cdot 1_{[d_H(z_q, z_j) > r]} \quad (8)$$

which corresponds to sampling the hash code with the largest Hamming distance that has a substring Hamming distance below r_i^* and is outside the r -ball centered on z_q (expressed via $1_{[d_H(z_q, z_j) > r]}$). By sampling hash codes with the largest value of $d_H(z_q, z_s)$, z_s is unlikely to be within top k , but would still appear in the candidate set due to the low substring Hamming distance. Based on the sampling of such hash codes, we can derive an objective that maximizes the Hamming distance within the substring as long as z_s appears in the candidate set, which is expressed in the following loss function:

$$\mathcal{L}_{\text{false-positive}} = -d_H(z_q^i, z_s^i) \quad (9)$$

In case one or both indicator functions in Eq. 8 are always 0 for a given query, this loss is simply set to 0.

4.2.1 Finding the pair (z_q, z_s). While training the network, the hash codes potentially change in each iteration, hence we need to continuously sample the pair (z_q, z_s) during training. As recomputing every hash code for every batch is computationally expensive, we employ a memory module that is continuously updated with the generated hash codes in addition to the associated document ids. We denote this memory module as M . The memory size (i.e., the number of hash codes to keep in M) is denoted by s_{mem} , and the memory module is updated using a first-in first-out (FIFO) strategy. Due to the compact representation of the hash codes and document ids, we fix s_{mem} to be the size of the training set².

Using the sampling requirements from Eq. 8, z_s can now be obtained from the memory module. However, since the memory module may contain outdated hash codes due to model updates in previous training iterations, the validity of the pair (z_q, z_s) must be ensured. This is done by recomputing z_s (based on the document features obtained through the stored document id) on the current model parameters, and verifying whether it is still valid according to the sampling requirements (if not, the loss is set to 0). Note that the search radii r_i^* and r for top k retrieval also needs to be estimated based on the memory module. However, since z_s is sampled as the hash code with the largest Hamming distance to z_q , smaller deviations from the true radii are not problematic, as the worst-case outcome simply is that the already far apart (z_q, z_s) pair is pushed slightly further apart than necessary.

²For truly massive-scale datasets, or due to specific hardware constraints, the memory size could be fixed to a number less than the training set size.

Table 1: Dataset statistics.

	documents	multi-class	classes	unique words
TMC	28,596	Yes	22	18,196
reuters	9,848	Yes	90	16,631
agnews	127,598	No	4	32,154

4.3 Control search radius

In Section 4.2 we detailed how to reduce the number of false positive documents per hash table lookup, while we now focus on how to reduce the number of such lookups. Given a query hash code z_q , we aim to control the search radius r to limit the exponential increase in the number of hash table lookups, which happens when the substring search radius $r_i^* > 1$, $i \in \{1, \dots, m\}$ (see Section 3.2.1), corresponding to $r > 2m - 1$. To this end, we compute r for the query based on the memory module, and sample a hash code z_r with $d_H(z_q, z_r) = r$, resulting in the hash code pair (z_q, z_r). To reduce the number of lookups, we define a loss function that minimizes the Hamming distance of the pair:

$$\mathcal{L}_{\text{radius}} = d_H(z_q, z_r) \cdot 1_{[r > 2m-1]} \quad (10)$$

where the indicator function ensures that the Hamming distance is only minimized in cases where the search radius is too large. Similarly to sampling z_s for reducing the number of documents per hash table lookup (Eq. 8), z_r may be outdated in the memory module, but is recomputed and it is verified whether its radius is still equal to r (otherwise the loss is set to 0).

4.4 Combined loss function

MISH is trained in an end-to-end fashion by jointly optimizing the semantic loss (Eq. 7), reducing the number of false positive documents per hash table lookup (Eq. 9), and controlling the number of such lookups (Eq. 10) as follows:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{semantic}} + \alpha_1 \mathcal{L}_{\text{false-positive}} + \alpha_2 \mathcal{L}_{\text{radius}} \quad (11)$$

where the hyperparameter weights, α_1 and α_2 , control the trade-off between the semantic encoding and tuning the hash codes towards more efficient multi-index hashing search. However, optimizing both effectiveness and efficiency are not necessarily mutually exclusive because any permutation of the hash code bits provides the same effectiveness, but some permutations result in better multi-index hashing efficiency. Lastly, observe that $\mathcal{L}_{\text{false-positive}}$ and $\mathcal{L}_{\text{radius}}$ are model agnostic, as neither are tied to how the hash codes are generated, and can thus easily be incorporated in any semantic hashing model for improving efficiency.

5 EXPERIMENTAL EVALUATION

5.1 Datasets

We evaluate MISH on well-known and publicly available datasets used in related work [3, 8, 10, 27] and summarized in Table 1: (1) *TMC* consists of multi-class air traffic reports; (2) *Agnews* consists of single-class news documents; and (3) *reuters* consists of multi-class news documents, where a filtering is applied that removes documents if none of its labels occur among the top 20 most frequent labels (as done in [3, 8, 10, 27]).

We use the preprocessed data provided by Hansen et al. [10], where documents are represented using TF-IDF vectors, where words occurring only once, or in more than 90% of the documents, are removed. We use the provided data splits, which split the datasets into training (80%), validation (10%), and testing (10%), where the validation loss is used for early stopping.

5.2 Evaluation setup

We evaluate the hash codes in the task of document similarity search (using kNN search), where we evaluate both effectiveness and efficiency. Each document is represented as a hash code, such that document similarity search can be performed using the Hamming distance between two hash codes. For evaluation purposes, we denote a document to be relevant (i.e., similar) to a query document, if the documents share at least one label, meaning that in multi-class datasets two documents do not need to share all labels.

For effectiveness we follow related work [3, 8, 10, 27] by considering the retrieval performance as measured by precision@100. However, existing work computes the scores based on random tie splitting, which is problematic for hash codes as ties occur often due to the limited number of $n+1$ different Hamming distances for n -bit hash codes. Instead, we compute the tie-aware precision@100 metric [19] corresponding to the average-case retrieval performance. Additionally, due to the large number of possible ties, we also compute the *worst-case* retrieval performance by fixing ties such that irrelevant documents appear before relevant ones before computing precision@100.

For efficiency we measure the runtime for performing top-100 retrieval on the training set, where each test document acts as a query document once. We perform a linear scan, i.e., brute-force computation of all Hamming distances, as well as multi-index hashing based search [23]. We use the linear scan and multi-index implementation made available by Norouzi et al. [23]³, where we follow the practical recommendation of splitting hash codes into substrings of 16 bits for multi-index hashing⁴. We repeat all timing experiments 100 times, and report the median speedup of using multi-index hashing over the linear scan. Note that the runtime of multi-index hashing naturally varies between the methods used for generating hash code, whereas the linear scan time is the same independent of the used method. All timing experiments were performed on an Intel Core i9-9940X@3.30 GHz.

5.3 Baselines

We compare our proposed MISH to non-neural approaches with post-processing rounding of document vectors to obtain hash codes (STH [33] and LCH [32]), neural approaches with post-processing rounding (NbrReg [2]), and neural end-to-end approaches that incorporate the rounding as part of the model (NASH [27], RBSH [8], and PairRec [10]). All baselines are tuned following the original papers.

In contrast to our MISH, existing semantic hashing baselines have focused on maximizing effectiveness, while simply assumed retrieval is done using a brute-force linear scan, rather than faster alternatives such as multi-index hashing. However, how bits are

assigned into substrings impacts multi-index hashing efficiency, as the candidate set size may be larger than necessary. To this end, we include the greedy substring optimization (GSO) heuristic proposed by Norouzi et al. [23], which greedily assigns bits to substrings as to minimize the correlation between the bits within each substring.

5.4 Tuning

To tune MISH⁵ we fix the number of hidden units in the encoder to 1000, and vary the number of documents in the pairwise reconstruction neighbourhood (\mathcal{N}_d^p) from {10, 25, 50, 100}, where both 10 and 25 worked well for reuters and TMC, whereas 100 was consistently chosen for agnews. Similarly to Hansen et al. [10], \mathcal{N}_d^p is constructed based on retrieving the top p most semantically similar documents based on 64 bit STH hash codes. For the KL-divergence, we tune β from {0, 0.01}, where 0 was chosen most often, thus effectively removing the KL term in those cases. For the variance annealing in the noise-infused hash codes, we fix the initial value to 1 and reduce by 10^{-6} every iteration (as per [8, 10]). For the combined loss, we tune α_1 from {1, 3, 5, 7} and α_2 from {0.01, 0.05, 0.1}, where $\alpha_1 = 3$ and $\alpha_2 = 0.01$ was chosen for reuters and TMC, and $\alpha_1 = 7$ and $\alpha_2 = 0.05$ for agnews. As we focus on learning hash codes that maintain state-of-the-art effectiveness, while improving efficiency, we choose to use only the semantic loss ($\mathcal{L}_{\text{semantic}}$) on the validation set for model selection and early stopping, rather than the weighted total loss. Lastly, MISH is optimized using the Adam optimizer [13] with a learning rate from {0.001, 0.005}, where 0.005 was chosen for reuters and TMC, and 0.001 for agnews.

5.5 Results

The experimental results are summarized for effectiveness in Table 2 and efficiency in Table 3. In Table 2, the highest worst-case and average-case scores per column are highlighted in bold, and the second highest are underlined. In Table 3, the largest and second largest speedups (independent of applying the greedy substring optimization (GSO)) are highlighted in bold and underlined, respectively. Additionally, we report the linear scan time per document as a point of reference for the speedups. In both tables, \blacktriangle represents statistically significant improvements over the second best method at the 0.05 level using a two-tailed paired t-test.

5.5.1 Retrieval effectiveness. Table 2 shows the effectiveness measured by worst-case and average-case prec@100 across the datasets using 32 and 64 bit hash codes (corresponding to the typical machine word sizes). Across all methods, we observe a larger gain in worst-case prec@100 when increasing the number of bits in the hash codes, compared to the average-case prec@100, where only smaller increases are obtained. Thus, increasing the number of bits is beneficial when worst-case performance is important no matter the chosen method. The increase in worst-case prec@100 happens because the documents are being spread out more in the Hamming space as the number of bits are increased, which reduce the number of Hamming distance ties.

Our MISH method obtains the best results for worst-case prec@100 in all cases, and additionally it obtains the best average-case prec@100 for reuters and agnews. For TMC, PairRec obtains marginally higher

³<https://github.com/norouzi/mih/>

⁴<https://github.com/norouzi/mih/blob/master/RUN.sh>

⁵We make our code publicly available at <https://github.com/Varyn/MISH>

Table 2: Worst case and average case precision@100. The highest precision is highlighted in bold, and the second highest is underlined. [▲] represents statistically significant improvements over the second best method at the 0.05 level using a two tailed paired t-test.

Prec@100	Reuters				TMC				Agnews			
	32 bits		64 bits		32 bits		64 bits		32 bits		64 bits	
	Worst	Average	Worst	Average	Worst	Average	Worst	Average	Worst	Average	Worst	Average
LCH	0.5995	0.6616	0.6283	0.6613	0.6658	0.7510	0.7421	0.7817	0.6822	0.7599	0.7423	0.7775
STH	0.7730	0.8046	0.7803	0.7968	0.6858	0.7693	0.7481	0.7816	0.6823	0.8237	0.7931	0.8374
NbrReg	0.5785	0.6329	0.6327	0.6616	0.3272	0.6648	0.5862	0.6827	0.7274	0.7914	0.7535	0.7928
NASH	0.7330	0.7737	0.7767	0.7967	0.6845	0.7709	0.7535	0.7953	0.7059	0.8018	0.7748	0.8107
RBSH	0.7809	0.8110	0.8011	0.8182	<u>0.7459</u>	0.8107	0.7852	0.8158	<u>0.7797</u>	0.8347	0.8053	0.8317
PairRec	<u>0.7812</u>	<u>0.8218</u>	<u>0.8087</u>	<u>0.8316</u>	0.7320	0.8187	<u>0.7922</u>	0.8288	<u>0.7700</u>	<u>0.8348</u>	<u>0.8114</u>	<u>0.8407</u>
MISH	0.7965	0.8286	0.8248	0.8377	0.7608[▲]	<u>0.8156</u>	0.7931	<u>0.8261</u>	0.7818	0.8375	0.8116	0.8419

Table 3: Speedup of multi-index hashing over a brute-force linear scan, as well as linear scan time per document. Greedy substring optimizing (GSO) [23] corresponds to the correlation-based post-hoc heuristic, and Default corresponds to using the hash codes as is. The largest speedup is highlighted in bold (independent of post-hoc fix), and the second largest is underlined. [▲] represents statistically significant improvements (100 repeated timing experiments) over the second best method at the 0.05 level using a two tailed paired t-test.

Speedup	Reuters				TMC				Agnews			
	32 bits		64 bits		32 bits		64 bits		32 bits		64 bits	
	Default	GSO	Default	GSO	Default	GSO	Default	GSO	Default	GSO	Default	GSO
LCH	2.5182	0.3508	0.8937	0.5603	7.2427	1.5879	2.1837	3.0767	18.2709	25.6365	3.7311	5.1354
STH	2.9374	0.2783	1.0116	0.5695	7.0382	2.2211	2.1947	2.4791	14.6936	29.2779	4.4515	<u>9.5676</u>
NbrReg	<u>6.3841</u>	0.7190	<u>4.9589</u>	0.8299	7.1497	1.7346	2.8800	<u>5.2506</u>	21.7102	23.7698	6.8518	7.5279
NASH	5.6356	0.6037	4.5869	0.8417	9.4069	1.5725	4.3819	4.3886	20.7673	23.1443	4.8355	5.3849
RBSH	4.4342	0.2965	1.7083	0.8051	7.0831	1.9708	2.7957	2.9146	25.9186	27.1164	4.5036	4.7345
PairRec	5.4296	0.3721	3.0770	0.9433	<u>10.9118</u>	1.7063	5.1129	4.9614	29.7880	<u>33.4096</u>	7.1765	7.8676
MISH	7.0698[▲]	1.1216	5.4466[▲]	1.0282	14.6296[▲]	2.1923	8.8696[▲]	6.1645	44.0151[▲]	35.9177	13.6756[▲]	12.5788
Linear scan time	0.000070 s		0.000069 s		0.000111 s		0.000109 s		0.000432 s		0.000407 s	

average-case prec@100 compared to MISH, but PairRec is in most cases the second best method. In general, the difference in effectiveness between the best and second best performing method is relatively small, and we only obtain statistically significant improvements for worst-case prec@100 at 32 bits for TMC. As both MISH and the state-of-the-art PairRec are based on the same semantic loss (Eq. 7), it was to be expected that MISH would not significantly improve effectiveness over PairRec. However, it is important to notice that including the two additional losses in MISH, for improving efficiency, did not negatively impact effectiveness either. In fact, the additional losses have a regularizing effect that reduces the number of Hamming distance ties (hence improving the worst-case performance), as the losses force the hash codes to be better spread in the Hamming space.

Table 4 shows the average percentage decreases in prec@100 (for both worst-case and average-case) compared to the best worst-case and average-case scores, respectively. For a given method, a decrease of 0% corresponds to that method always being the best performing method across all datasets. We observe that on average, MISH outperforms the other methods in both cases, with a noticeable better average worst-case effectiveness. Additionally, it

Table 4: Average decrease in worst-case and average-case prec@100 compared to the best scores per dataset. An average decrease of 0% corresponds to obtaining the best performance across all datasets.

	32 bits		64 bits	
	Δ Worst	Δ Average	Δ Worst	Δ Average
LCH	-16.65%	-12.57%	-12.93%	-11.46%
STH	-8.51%	-3.53%	-4.45%	-3.71%
NbrReg	-30.44%	-15.98%	-18.84%	-14.83%
NASH	-9.24%	-5.58%	-5.12%	-4.21%
RBSH	-1.39%	-1.15%	-1.55%	-1.71%
PairRec	-2.40%	-0.38%	-0.70%	-0.29%
MISH	0.00%	-0.13%	0.00%	-0.11%

can be seen that the baselines generally have larger worst-case decreases compared to the average-case decreases, which shows that the baselines broadly cluster the hash codes more, thus resulting in larger number of Hamming distance ties.

5.5.2 Retrieval efficiency. Table 3 shows the relative speedup compared to a linear scan of the hash codes produced by each method,




Figure 2: Speedup of multi-index hashing over linear scan relative to MISH averaged across all datasets and bit configurations.

with and without greedy substring optimization (GSO) [23], together with the linear scan time per document. The linear scan time per document is slightly lower for 64 bit hash codes compared to 32 bit, due to our machine using a 64 bit operating system. In addition, due to the sequential access pattern in a linear scan, the larger memory requirement of 64 bit hash codes does not increase the scan time. Overall, we observe that all methods achieve a higher speedup for 32 bits compared to the speedup at 64 bits, caused by an increase in the search radius for top 100 retrieval, which leads to a larger number of hash table lookups, thus increasing the candidate set of multi-index hashing. Furthermore, as expected the speedup increases as the dataset size increases, because the relative size of the candidate set decreases compared to the entire set of hash codes.

Table 3 shows that MISH is significantly faster than the baselines on all datasets and number of bits. The speedup obtained by the baselines are largely varied across the datasets and number of bits. No baseline can perform consistently well in all settings. GSO leads to larger speedups for the baselines on agnews and for 64 bit hash codes on TMC (except for PairRec), but worse on reuters and 32 bit hash codes on TMC. For MISH, applying GSO always decreases the speedup, which is caused by GSO changing the substring structure learned during training through our proposed loss functions (Eq. 9 and Eq. 10), to a less optimal one. This highlights that GSO is not always beneficial and, more importantly, the limitation of making a post-hoc change to the structure of the hash codes. In contrast, MISH directly optimizes the desirable hash code properties for multi-index hashing (see Section 3.2.2) in an end-to-end fashion, which significantly improves multi-index hashing efficiency.

Figure 2 shows the average relative speedup, across all datasets and number of bits, compared to MISH. Generally the baselines can be clustered in two groups of similar average efficiency (LCH, RBSH, STH) and (NASH, NbrReg, PairRec). Interestingly, RBSH is the only neural method among the least efficient methods, even though PairRec and RBSH share the same underlying neural architecture with the difference being that RBSH uses a pairwise ranking loss and PairRec uses a pairwise reconstruction loss. This shows that it is problematic to assume the produced hash codes by a given method will be efficient, without directly optimizing them as done in MISH, as even small changes in the model may greatly influence the efficiency.




Figure 3: Hyperparameter impact on speedup and prec@100: α_1 reduces the number of documents per hash table lookup, and α_2 controls the number of such lookups.

5.6 Efficiency and effectiveness impact of α_1 and α_2

To maintain state-of-the-art effectiveness, model selection was done based only on the semantic loss (Eq. 7) on the validation set, but model training is naturally still done on the weighted total loss (Eq. 11). We now investigate the impact of the total loss weights (α_1 and α_2) on the efficiency and effectiveness of MISH, where α_1 reduces the number of documents per hash table lookup, and α_2 controls the number of such lookups. To this end, we report the speedup and average-case prec@100 for as two grid plots with α_1 and α_2 as the axes, exemplified for 64 bit hash codes on agnews.

The grid plots can be seen in Figure 3, where the top left corner ($\alpha_1 = 0$, $\alpha_2 = 0$) corresponds to the PairRec baseline [10]. For the speedup plot, we observe a clear trend that higher values of both α_1 and α_2 improve efficiency, but α_1 has the largest impact. This is expected since α_1 directly reduces the number of documents per hash table lookup, reducing the candidate set across all queries, while α_2 affects a smaller subset of queries who exhibit a large search radius. For the prec@100 plot, we observe that higher values of α_1 and α_2 reduce effectiveness, which highlights the possible trade-off when tuning the hyperparameters. However, the area with the largest prec@100 scores is relatively large, thus enabling a large efficiency improvement without compromising effectiveness.

5.7 Distribution of candidate set sizes

The computational cost of multi-index hashing is dominated by the cost of sorting the candidate set according to the Hamming distances to the query hash code. We now investigate the distribution of the candidate set size per hash code query for each method, as visualized in Figure 4 using symmetrical violin plots⁶. For the cases where GSO leads to improved speedups, we compute the candidate set based on the hash codes after applying GSO. Across all datasets and bit sizes, we observe that MISH is able to better concentrate the density towards a lower number of candidates compared to the baselines, thus explaining the large speedup improvements. By observing the larger candidate sizes for 64 bit hash codes compared to 32 bit hash code, we can directly see the reason for the speedup gap between the two hash code sizes reported in Table 3. Furthermore, the baselines exhibiting poor speedups are primarily due to long density tails or a more uniform distribution of candidate set sizes.

⁶A violin plot is a combination of a boxplot and a symmetrical density plot that shows the full distribution of the data.

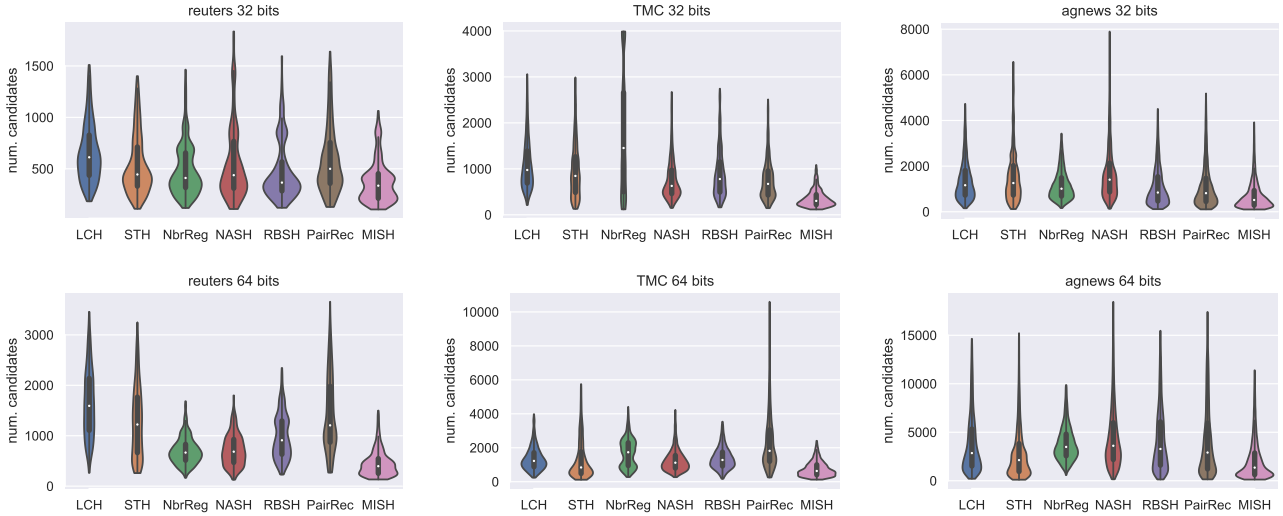


Figure 4: Violin plots (combined density and boxplot) of the number of document candidates found using multi-index hashing for top-100 retrieval.

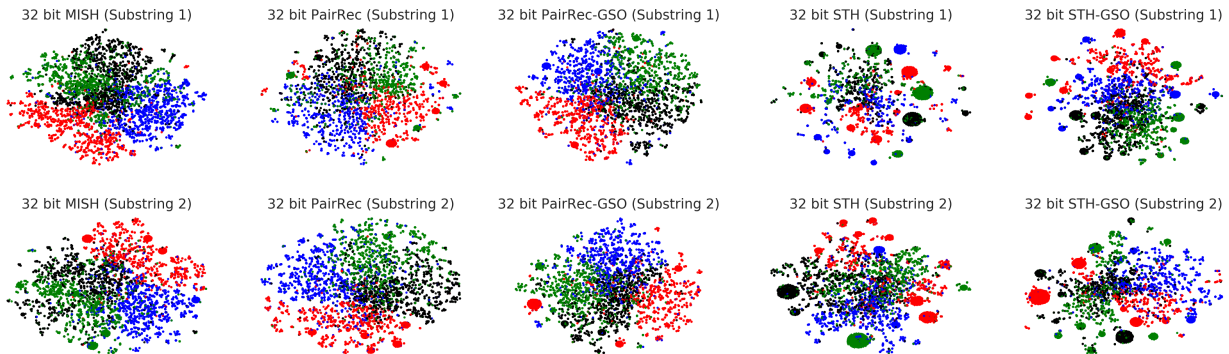


Figure 5: t-SNE [18] visualization of the two substring of MISH, PairRec, and STH 32 bit hash codes on agnews. Each color represents one of 4 different classes. Greedy substring optimization (GSO) [23] is applied on PairRec and STH as it improved their efficiency for multi-index hashing.

While this naturally leads to worse overall efficiency, it also has the potential problem of large query time variance, which may limit their application in extremely time-constrained use cases.

5.8 Substring-level visualization

To further investigate the differences in how the methods distribute the bits within the hash codes, we choose to visualize the individual substrings. We consider 32 bit hash codes from the test set of agnews, as it only contains two substrings, while agnews is a single-class dataset, which makes it easier to visualize if hash codes of the same class are clustered. We consider MISH, PairRec, and STH as representative methods, as it also enables visualizing the impact of GSO on the two baselines. PairRec is chosen due to being the second best method on both effectiveness and efficiency for 32 bit hash codes on agnews, while STH represents a method where GSO

greatly improves (doubles) its speedup (see Table 3). Figure 5 shows a two-dimensional t-SNE [18] visualization, where it is important to keep in mind that each plot contains the same number of points, such that a plot appearing more sparse (more distance between points) corresponds to more hash codes being highly similar. When comparing MISH and PairRec, we observe that they do appear similar, but MISH is slightly less sparse, meaning the hash codes are better spread throughout the Hamming space. For STH, we observe that prior to applying GSO, the hash codes are tightly clustered and highly sparse (especially in the first substring), but GSO is able to redistribute the bits such that the substrings better utilize the space. This redistribution reduces the amount of false positive candidates found in multi-index hashing, thus leading to the large observed speedup.

5.9 Conclusion

We presented Multi-Index Semantic Hashing (MISH), an unsupervised semantic hashing model that learns hash codes well suited for multi-index hashing [23], which enables highly efficient document similarity search. Compared to a brute-force linear scan over all the hash codes, multi-index hashing constructs a smaller candidate set to search over, which can provide sub-linear search time. We identify key hash code properties that affect the size of the candidate set, and use them to derive two novel objectives that enable MISH to learn hash codes that results in smaller candidate sets when using multi-index hashing. Our objectives are model agnostic, i.e., not tied to how the hash codes are generated specifically in MISH, which means they are straight-forward to incorporate in existing and future semantic hashing models. We experimentally compared MISH to state-of-the-art semantic hashing baselines in the task of document similarity search, where we evaluated both efficiency and effectiveness. While multi-index hashing also improves the efficiency of the baseline hash codes compared to a linear scan, they are still upwards of 33% slower than our proposed MISH. Interestingly, these large efficiency gains of MISH can be obtained without reducing effectiveness, as MISH is still able to obtain state-of-the-art effectiveness, but we do find that even further efficiency improvements can be obtained, but at the cost of an effectiveness reduction. In future work, we plan to explore supervised versions of MISH, specifically the impact of expanding our proposed efficiency objectives with label information, which could decrease the number of irrelevant documents in the candidate sets.

REFERENCES

- [1] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. 2013. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432* (2013).
- [2] Suthae Chaidaroon, Travis Ebesu, and Yi Fang. 2018. Deep Semantic Text Hashing with Weak Supervision. *SIGIR*, 1109–1112.
- [3] Suthae Chaidaroon and Yi Fang. 2017. Variational deep semantic hashing for text documents. In *SIGIR*. 75–84.
- [4] Mayur Datar, Nicole Immorlica, Piotr Indyk, and Vahab S Mirrokni. 2004. Locality-sensitive hashing scheme based on p-stable distributions. In *Proceedings of the twentieth annual symposium on Computational geometry*. ACM, 253–262.
- [5] Scott Deerwester, Susan T Dumais, George W Furnas, Thomas K Landauer, and Richard Harshman. 1990. Indexing by latent semantic analysis. *Journal of the American society for information science* 41, 6 (1990), 391–407.
- [6] Wei Dong, Qinliang Su, Dinghan Shen, and Changyou Chen. 2019. Document Hashing with Mixture-Prior Generative Models. In *EMNLP*. 5226–5235.
- [7] Dan Greene, Michal Parnas, and Frances Yao. 1994. Multi-index hashing for information retrieval. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*. IEEE, 722–731.
- [8] Casper Hansen, Christian Hansen, Jakob Grue Simonsen, Stephen Alstrup, and Christina Lioma. 2019. Unsupervised Neural Generative Semantic Hashing. In *SIGIR*. 735–744.
- [9] Casper Hansen, Christian Hansen, Jakob Grue Simonsen, Stephen Alstrup, and Christina Lioma. 2020. Content-aware Neural Hashing for Cold-start Recommendation. In *SIGIR*. 971–980.
- [10] Casper Hansen, Christian Hansen, Jakob Grue Simonsen, Stephen Alstrup, and Christina Lioma. 2020. Unsupervised Semantic Hashing with Pairwise Reconstruction. In *SIGIR*. 2009–2012.
- [11] Christian Hansen, Casper Hansen, Jakob Grue Simonsen, and Christina Lioma. 2021. Projected Hamming Dissimilarity for Bit-Level Importance Coding in Collaborative Filtering. In *Proceedings of The Web Conference 2021*. In print.
- [12] Wang-Cheng Kang and Julian McAuley. 2019. Candidate Generation with Binary Codes for Large-Scale Top-N Recommendation. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 1523–1532.
- [13] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *ICLR*.
- [14] Diederik P Kingma and Max Welling. 2014. Auto-encoding variational bayes. In *ICLR*.
- [15] Jure Leskovec, Anand Rajaraman, and Jeffrey David Ullman. 2020. *Mining of massive data sets*. Cambridge university press.
- [16] Defu Lian, Xing Xie, and Enhong Chen. 2019. Discrete matrix factorization and extension for fast item recommendation. *IEEE Transactions on Knowledge and Data Engineering* (2019).
- [17] Wei Liu, Jun Wang, Sanjiv Kumar, and Shih-Fu Chang. 2011. Hashing with graphs. In *ICML*.
- [18] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9, Nov (2008), 2579–2605.
- [19] Frank McSherry and Marc Najork. 2008. Computing Information Retrieval Performance Measures Efficiently in the Presence of Tied Scores. In *Proceedings of the IR Research, 30th European Conference on Advances in Information Retrieval*. 414–421.
- [20] Andrew Y Ng, Michael I Jordan, and Yair Weiss. 2002. On spectral clustering: Analysis and an algorithm. In *NeurIPS*. 849–856.
- [21] Mohammad Norouzi and David J Fleet. 2011. Minimal loss hashing for compact binary codes. In *ICML*.
- [22] Mohammad Norouzi, Ali Punjani, and David J Fleet. 2012. Fast search in hamming space with multi-index hashing. In *2012 IEEE conference on computer vision and pattern recognition*. IEEE, 3108–3115.
- [23] Mohammad Norouzi, Ali Punjani, and David J Fleet. 2013. Fast exact search in hamming space with multi-index hashing. *IEEE transactions on pattern analysis and machine intelligence* 36, 6 (2013), 1107–1119.
- [24] Stephen E Robertson, Steve Walker, Susan Jones, Michelle M Hancock-Beaulieu, Mike Gafford, et al. 1995. Okapi at TREC-3. *Nist Special Publication Sp 109* (1995), 109.
- [25] Ruslan Salakhutdinov and Geoffrey Hinton. 2009. Semantic hashing. *International Journal of Approximate Reasoning* 50, 7 (2009), 969–978.
- [26] Ying Shan, Jian Jiao, Jie Zhu, and JC Mao. 2018. Recurrent binary embedding for gpu-enabled exhaustive retrieval from billion-scale semantic vectors. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2170–2179.
- [27] Dinghan Shen, Qinliang Su, Paidamoyo Chapfuwa, Wenlin Wang, Guoyin Wang, Ricardo Henao, and Lawrence Carin. 2018. NASH: Toward End-to-End Neural Architecture for Generative Semantic Hashing. In *ACL*. 2041–2050.
- [28] Jun Wang, Wei Liu, Sanjiv Kumar, and Shih-Fu Chang. 2015. Learning to hash for indexing big data—A survey. *Proc. IEEE* 104, 1 (2015), 34–57.
- [29] Jingdong Wang, Ting Zhang, Nicu Sebe, and Heng Tao Shen. 2018. A survey on learning to hash. *IEEE transactions on pattern analysis and machine intelligence* 40, 4 (2018), 769–790.
- [30] Meng Wang, Haomin Shen, Sen Wang, Lina Yao, Yinlin Jiang, Guilin Qi, and Yang Chen. 2019. Learning to Hash for Efficient Search Over Incomplete Knowledge Graphs. In *2019 IEEE International Conference on Data Mining (ICDM)*. IEEE, 1360–1365.
- [31] Yair Weiss, Antonio Torralba, and Rob Fergus. 2009. Spectral hashing. In *NeurIPS*. 1753–1760.
- [32] Dell Zhang, Jun Wang, Deng Cai, and Jinsong Lu. 2010. Laplacian co-hashing of terms and documents. In *ECIR*. Springer, 577–580.
- [33] Dell Zhang, Jun Wang, Deng Cai, and Jinsong Lu. 2010. Self-taught hashing for fast similarity search. In *SIGIR*. ACM, 18–25.
- [34] Hanwang Zhang, Fumin Shen, Wei Liu, Xiangnan He, Huanbo Luan, and Tat-Seng Chua. 2016. Discrete collaborative filtering. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*. 325–334.
- [35] Hanwang Zhang, Meng Wang, Richang Hong, and Tat-Seng Chua. 2016. Play and rewind: Optimizing binary representations of videos by self-supervised temporal hashing. In *Proceedings of the 24th ACM international conference on Multimedia*. 781–790.

Chapter 5

Content-aware Neural Hashing for Cold-start Recommendation

Casper Hansen*, Christian Hansen*, Jakob Grue Simonsen, Stephen Alstrup, Christina Lioma (2020). Content-aware Neural Hashing for Cold-start Recommendation. In SIGIR, pages 971-980. [23]. * denotes equal contribution.

Content-aware Neural Hashing for Cold-start Recommendation

Casper Hansen*
University of Copenhagen
c.hansen@di.ku.dk

Christian Hansen*
University of Copenhagen
chrh@di.ku.dk

Jakob Grue Simonsen
University of Copenhagen
simonsen@di.ku.dk

Stephen Alstrup
University of Copenhagen
s.alstrup@di.ku.dk

Christina Lioma
University of Copenhagen
c.lioma@di.ku.dk

ABSTRACT

Content-aware recommendation approaches are essential for providing meaningful recommendations for *new* (i.e., *cold-start*) items in a recommender system. We present a content-aware neural hashing-based collaborative filtering approach (NeuHash-CF), which generates binary hash codes for users and items, such that the highly efficient Hamming distance can be used for estimating user-item relevance. NeuHash-CF is modelled as an autoencoder architecture, consisting of two joint hashing components for generating user and item hash codes. Inspired from semantic hashing, the item hashing component generates a hash code directly from an item's content information (i.e., it generates cold-start and seen item hash codes in the same manner). This contrasts existing state-of-the-art models, which treat the two item cases separately. The user hash codes are generated directly based on user id, through learning a user embedding matrix. We show experimentally that NeuHash-CF significantly outperforms state-of-the-art baselines by up to 12% NDCG and 13% MRR in cold-start recommendation settings, and up to 4% in both NDCG and MRR in standard settings where all items are present while training. Our approach uses 2-4x shorter hash codes, while obtaining the same or better performance compared to the state of the art, thus consequently also enabling a notable storage reduction.

KEYWORDS

Hashing; Cold-start Recommendation; Collaborative Filtering; Content-Aware Recommendation; Autoencoders

ACM Reference Format:

Casper Hansen, Christian Hansen, Jakob Grue Simonsen, Stephen Alstrup, and Christina Lioma. 2020. Content-aware Neural Hashing for Cold-start Recommendation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '20)*, July 25–30, 2020, Virtual Event, China. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3397271.3401060>

*Both authors share the first authorship

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGIR '20, July 25–30, 2020, Virtual Event, China

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8016-4/20/07... \$15.00

<https://doi.org/10.1145/3397271.3401060>

1 INTRODUCTION

Personalizing recommendations is a key factor in successful recommender systems, and thus is of large industrial and academic interest. Challenges arise both with regards to efficiency and effectiveness, especially for large-scale systems with tens to hundreds of millions of items and users.

Recommendation approaches based on collaborative filtering (CF), content-based filtering, and their combinations have been investigated extensively (see the surveys in [2, 31]), with CF based systems being one of the major methods in this area. CF based systems learn directly from either implicit (e.g., clicks) or explicit feedback (e.g., ratings), where matrix factorization approaches have traditionally worked well [5, 20]. CF learns m -dimensional user and item representations based on a factorization of the interaction-matrix between users and items, e.g., based on their click or rating history, such that the inner product can be used for computing user-item relevance. However, in the case of new unseen items (i.e., *cold-start* items), standard CF methods are unable to learn meaningful representations, and thus cannot recommend those items (and similarly for cold-start users). To handle these cases, content-aware approaches are used when additional content information is available, such as textual descriptions, and have been shown to improve upon standard CF based methods [21].

In large-scale recommendation settings, providing top-K recommendations among all existing items using an inner product is computationally costly, and thus provides a practical obstacle in employing these systems at scale. Hashing-based approaches solve this by generating *binary* user and item hash codes, such that user-item relevance can be computed using the Hamming distance (i.e., the number of bit positions where two bit strings are different). The Hamming distance has a highly efficient hardware-level implementation, and has been shown to allow for real-time retrieval among a billion items [29]. Early work on hashing-based collaborative filtering systems [17, 40, 41] learned real-valued user and item representations, which were then in a later step discretized into binary hash codes. Further work focuses on end-to-end approaches, which improve upon the two-stage approaches by reducing the discretizing error by optimizing the hash codes directly [24, 37]. Recent content-aware hashing-based approaches [22, 39] have been shown to perform well in both standard and cold-start settings, however they share the common problem of generating cold-start item hash codes differently from standard items, which we claim is unnecessary and limits their generalizability in cold-start settings.

We present a novel neural approach for content-aware hashing-based collaborative filtering (NeuHash-CF) robust to cold-start recommendation problems. NeuHash-CF consists of two joint hashing components for generating user and item hashing codes, which are connected in a variational autoencoder architecture. Inspired by semantic hashing [28], the item hashing component learns to directly map an item’s content information to a hash code, while maximizing its ability to reconstruct the original content information input. The user hash codes are generated directly based on the user’s id through learning a user embedding matrix, and are jointly optimized with the item hash codes to optimize the log likelihood of observing each user-item rating in the training data. Through this end-to-end trainable architecture, all item hash codes are generated in the same way, independently of whether they are seen or not during training. We experimentally compare our NeuHash-CF to state-of-the-art baselines, where we obtain significant performance improvements in cold-start recommendation settings by up to 12% NDCG and 13% MRR, and up to 4% in standard recommendation settings. Our NeuHash-CF approach uses 2-4x fewer bits, while obtaining the same or better performance than the state of the art, and notable storage reductions.

In summary, we **contribute** a novel content-aware hashing-based collaborative filtering approach (NeuHash-CF), which in contrast to existing state-of-the-art approaches generates item hash codes in a unified way (not distinguishing between standard and cold-start items).

2 RELATED WORK

The seminal work of Das et al. [9] used a Locality-Sensitive Hashing [10] scheme, called Min-Hashing, for efficiently searching Google News, where a Jaccard measure for item-sharing between users was used to generate item and user hash codes. Following this, Karatzoglou et al. [17] used matrix factorization to learn real-valued latent user and item representations, which were then mapped to binary codes using random projections. Inspired by this, Zhou and Zha [41] applied iterative quantization [11] as a way of rotating and binarizing the real-valued latent representations, which had originally been proposed for efficient hashing-based image retrieval. However, since the magnitude of the original real-valued representations are lost in the quantization, the Hamming distance between two hash codes might not correspond to the original relevance (inner product of real-valued vectors) of an item to a user. To solve this, Zhang et al. [40] imposed a constant norm constraint on the real-valued representations followed by a separate quantization.

Each of the above approaches led to improved recommendation performance, however, they can all be considered two-stage approaches, where the quantization is done as a post-processing step, rather than being part of the hash code learning procedure. Furthermore, post-processing quantization approaches have been shown to lead to large quantization errors [37], leading to the investigation of approaches learning the hash codes directly.

Next, we review (1) hashing-based approaches for recommendation with explicit feedback; (2) content-aware hashing-based recommendation approaches designed for the cold-start setting of item recommendation; and (3) the related domain of semantic hashing, which our approach is partly inspired from.

2.1 Learning to Hash Directly

Discrete Collaborative Filtering (DCF) [37] was the first approach towards learning item and user hash codes directly, rather than through a two-step approach. DCF is based on a matrix factorization formulation with additional constraints enforcing the discreteness of the generated hash codes. DCF further investigated balanced and de-correlation constraints to improve generalization by better utilizing the Hamming space. Inspired by DCF, Zhang et al. [38] proposed Discrete Personalized Ranking (DPR) as a method designed for collaborative filtering with implicit feedback (in contrast to explicit feedback in the DCF case). DPR optimized a ranking objective through AUC and regularized the hash codes using both balance and de-correlation constraints similar to DCF. While these and previous two-stage approaches have led to highly efficient and improved recommendations, they are still inherently constrained by the limited representational ability of binary codes (in contrast to real-valued representations). To this end, Compositional Coding for Collaborative Filtering (CCCF) [24] was proposed as a hybrid approach between discrete and real-valued representations. CCCF considers each hash code as consisting of a number of blocks, each of which is associated with a learned real-valued scalar weight. The block weights are used for computing a *weighted* Hamming distance, following the intuition that not all parts of an item hash code are equally relevant for all users. While this hybrid approach led to improved performance, it has a significant storage overhead (due to each hash code’s block weights) and computational runtime increase, due to the weighted Hamming distance, compared to the efficient hardware-supported Hamming distance.

2.2 Content-aware Hashing

A common problem for collaborative filtering approaches, both binary and real-valued, is the cold-start setting, where a number of items have not yet been seen by users. In this setting, approaches based solely on traditional collaborative filtering cannot generate representations for the new items. Inspired by DCF, Discrete Content-aware Matrix Factorization (DCMF) [22] was the first hashing-based approach that also handled the cold-start setting. DCMF optimizes a multi-objective loss function, which most importantly learns hash codes directly for minimizing the squared rating error. Secondly, it also learns a latent representation for each content feature (e.g., each word in the content vocabulary), which is multiplied by the content features to approximate the learned hash codes, such that this can be used for generating hash codes in a cold-start setting. DCMF uses an alternating optimization strategy and, similarly to DCF, includes constraints enforcing bit balancing and de-correlation. Another approach, Discrete Deep Learning (DDL) [39] learns hash codes similarly to DCMF, through an alternating optimization strategy solving a relaxed optimization problem. However, instead of learning latent representations for each content feature to solve the cold-start problem, they train a deep belief network [16] to approximate the already learned hash codes based on the content features. This is a problem as described below.

DCMF and DDL both primarily learn hash codes not designed for cold-start settings, but then as a sub-objective learn how to map content features to new compatible hash codes for the cold-start setting. In practice, this is problematic as it corresponds to learning

cold-start item hash codes based on previously learned hash codes from standard items, which we claim is unnecessary and limits their generalizability in cold-start settings. In contrast, our proposed NeuHash-CF approach does not distinguish between the settings for generating item hash codes, but rather always bases the item hash codes on the content features through a variational autoencoder architecture. As such, our approach can learn a better mapping from content features to hash code, since it is learned directly, as opposed to learning it in two steps by approximating the existing hash codes that have already been generated.

2.3 Semantic Hashing

The related area of Semantic Hashing [28] aims to map objects (e.g., images or text) to hash codes, such that similar objects have a short Hamming distance between them. Early work focused on two-step approaches based on learning real-valued latent representations followed by a rounding stage [34–36]. Recent work has primarily used autoencoder-based approaches, either with a secondary rounding step [7], or through direct optimization of binary codes using Bernoulli sampling and straight-through estimators for back-propagation during training [13, 14, 30]. We draw inspiration from the latter approaches in the design of the item hashing component of our approach, as substantial performance gains have previously been observed in the semantic hashing literature over rounding-based approaches.

3 HASHING-BASED COLLABORATIVE FILTERING

Collaborative Filtering learns real-valued latent user and item representations, such that the inner product between a user u and item i corresponds to the item’s relevance to that specific user, where the ground truth is denoted as a user-item rating $R_{u,i}$. Hashing-based collaborative filtering learns *hash codes*, corresponding to *binary* latent representations, for users and items. We denote m -bit user and item hash codes as $z_u \in \{-1, 1\}^m$ and $z_i \in \{-1, 1\}^m$, respectively. For estimating an item’s relevance to a specific user in the hashing setting, the Hamming distance is computed as opposed to the inner product, as:

$$H(z_u, z_i) = \sum_{j=1}^m 1_{[z_u^{(j)} \neq z_i^{(j)}]} = \text{SUM}(z_u \text{ XOR } z_i) \quad (1)$$

Thus, the Hamming distance corresponds to summing the differing bits between the codes, which can be implemented very efficiently using hardware-level bit operations through the bitwise XOR and *popcount* operations. The relation between the inner product and Hamming distance of hash codes is simply:

$$z_u^T z_i = m - 2H(z_u, z_i) \quad (2)$$

meaning it is trivial to replace real-valued user and item representations with learned hash codes in an existing recommender system.

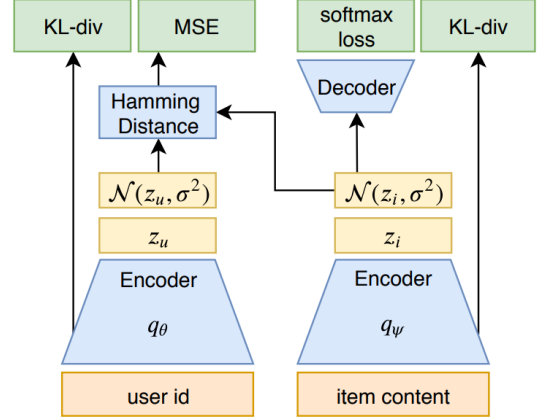


Figure 1: NeuHash-CF model overview.

3.1 Content-aware Neural Hashing-based Collaborative Filtering (NeuHash-CF)

We first give an overview of our model, Content-aware Neural Hashing-based Collaborative Filtering (NeuHash-CF), and then detail its components. NeuHash-CF consists of two joint components for generating user and item hash codes. The item hashing component learns to derive item hash codes directly from the content features associated with each item. The item hashing component has two optimization objectives: (1) to maximize the likelihood of the observed user-item ratings, and (2) the unsupervised objective of reconstructing the original content features. Through this design, all item hash codes are based on content features, thus directly generating hash codes usable for both standard and cold-start recommendation settings. This contrasts existing state-of-the-art models [22, 39] that separate how standard and cold-start item hash codes are generated. Through this choice, NeuHash-CF can generate higher quality cold-start item hash codes, but also improve the representational power of already observed items by better incorporating content features.

The user hashing component learns user hash codes, located within the same Hamming space as the item hash codes, by maximizing the likelihood of the observed user-item ratings, which is a shared objective with the item hashing component. Maximizing the likelihood of the observed user-item ratings influences the model optimization in relation to both user and item hash codes, while the unsupervised feature reconstruction loss of the item hashing component is focused only on the item hash codes. The aim of this objective combination is to ensure that the hash code distances enforce user-item relevance, but also that items with similar content have similar hash codes.

Next, we describe the architecture of our variational autoencoder (Section 3.2), followed by how users and items are encoded into hash codes (Section 3.3), decoded for obtaining a target value (Section 3.4), and lastly the formulation of the final loss function (Section 3.5). We provide a visual overview of our model in Figure 1.

3.2 Variational Autoencoder Architecture

We propose a variational autoencoder architecture for generating user and item hash codes, where we initially define the likelihood

functions of each user and item as:

$$p(u) = \prod_{i \in \mathbb{I}_u} p(R_{u,i}) \quad (3)$$

$$p(i) = p(c_i) + \prod_{u \in \mathbb{U}_i} p(R_{u,i}) \quad (4)$$

where \mathbb{I}_u is the set of all items rated by user u , \mathbb{U}_i is the set of all users who have rated item i , and $p(c_i)$ is the probability of observing the content of item i . We denote as $c_i \in \mathbb{R}$ the n -dimensional content feature vector (a bag-of-words representation) associated with each item, and denote the non-zero entries as \mathbb{W}_{c_i} . Thus, we can define the content likelihood similar to Eq. 3 and 4:

$$p(c_i) = \prod_{w \in \mathbb{W}_{c_i}} p(w). \quad (5)$$

In order to maximize the likelihood of the users and items, we need to maximize the likelihood of the observed ratings, $p(R_{u,i})$, as well as the word probabilities $p(w)$. Since they must be maximized based on the generated hash codes, we assume that $p(R_{u,i})$ is conditioned on both z_u and z_i , and that $p(w)$ is conditioned on z_i . For ease of derivation, we choose to maximize the log likelihood instead of the raw likelihoods, such that the log likelihood of the observed ratings and item content can be computed as:

$$\log p(R_{u,i}) = \log \sum_{z_i, z_u \in \{-1, 1\}^m} p(R_{u,i} | z_i, z_u) p(z_i) p(z_u) \quad (6)$$

$$\log p(c_i) = \log \sum_{z_i \in \{-1, 1\}^m} p(c_i | z_i) p(z_i) \quad (7)$$

where the hash codes are sampled by repeating m consecutive Bernoulli trials, which as a prior is assumed to have equal probability of sampling either 1 or -1. Thus, $p(z_i)$ and $p(z_u)$ can be computed simply as:

$$p(z) = \prod_{j=1}^m p^{\delta_j} (1-p)^{1-\delta_j}, \quad \delta_j = 1_{[z^{(j)} > 0]} \quad (8)$$

where $z^{(j)}$ is the j 'th bit of a hash code (either user or item), and where we set $p = 0.5$ for equal sampling probability of 1 and -1. However, optimizing the log likelihoods directly is intractable, so instead we maximize their variational lower bounds [19]:

$$\log p(R_{u,i}) \geq E_{q_{\psi}, q_{\theta}} \left[\log p(R_{u,i} | z_i, z_u) - \text{KL}(q_{\psi}(z_i | i) || p(z_i)) - \text{KL}(q_{\theta}(z_u | u) || p(z_u)) \right] \quad (9)$$

$$\log p(c_i) \geq E_{q_{\psi}} \left[\log p(c_i | z_i) - \text{KL}(q_{\psi}(z_i | c_i) || p(z_i)) \right] \quad (10)$$

where $q_{\psi}(z_i | i)$ and $q_{\theta}(z_u | u)$ are *learned* approximate posterior probability distributions (see Section 3.3), and KL is the Kullback-Leibler divergence. Intuitively, the conditional log likelihood within the expectation term can be considered a reconstruction term, which represents how well either the observed ratings or item content can be decoded from the hash codes (see Section 3.4). The KL divergence can be considered as a regularization term, by punishing large deviations from the Bernoulli distribution with equal sampling

probability of 1 and -1, which is computed analytically as:

$$\begin{aligned} \text{KL}(q_{\psi}(z_i | c_i) || p(z_i)) &= q_{\psi}(c_i) \log \frac{q_{\psi}(c_i)}{p} \\ &+ (1 - q_{\psi}(c_i)) \log \frac{1 - q_{\psi}(c_i)}{p} \end{aligned} \quad (11)$$

with $p = 0.5$ for equal sampling probability. The KL divergence is computed similarly for the user hash codes using θ . Next we describe how to compute the learned approximate posterior probability distributions.

3.3 Encoder Functions

The learned approximate posterior distributions q_{ψ} and q_{θ} can be considered encoder functions for items and users, respectively, and are both modeled through a neural network formulation. Their objective is to transform users and items into m bit hash codes.

3.3.1 Item encoding. An item i is encoded based on its content c_i through multiple layers to obtain sampling probabilities for generating the hash code:

$$l_1 = \text{ReLU}(W_1(c_i \odot w_{\text{imp}}) + b_1) \quad (12)$$

$$l_2 = \text{ReLU}(W_2 l_1 + b_2) \quad (13)$$

where W and b are learned weights and biases, \odot is elementwise multiplication, and w_{imp} is a learned importance weight for scaling the content words, which has been used similarly for semantic hashing [13]. Next, we obtain the sampling probabilities by transforming the last layer, l_2 , into an m -dimensional vector:

$$q_{\psi}(c_i) = \sigma(W_3 l_2 + b_3) \quad (14)$$

where σ is the sigmoid function to scale the output between 0 and 1, and ψ is the set of parameters used for the item encoding. We can now sample the item hash code from a Bernoulli distribution, which can be computed for each bit as:

$$z_i^{(j)} = 2 \lceil q_{\psi}(i)^{(j)} - \mu^{(j)} \rceil - 1 \quad (15)$$

where $\mu \in [0, 1]^m$ is an m -dimensional vector with uniformly sampled values. The model is trained using randomly sampled μ vectors, since it encourages model exploration because the same item may be represented as multiple different hash codes during training. However, to produce a deterministic output for testing once the model is trained, we fix each value within μ to 0.5 instead of a randomly sampled value.

3.3.2 User encoding. The user hash codes are learned similarly to the item hash codes, however, since we do not have a user feature vector, the hash codes are learned using only the user id. Thus, the sampling probabilities are learned as:

$$q_{\theta}(u) = \sigma(E_{\text{user}} 1_u) \quad (16)$$

where $E_{\text{user}} \in \mathbb{R}^{|U| \times m}$ is the learned user embedding, and 1_u is a one-hot encoding of user u . Following the same approach as the item encoding, we can sample the user hash code based on $q_{\theta}(u)$ for each bit as:

$$z_u^{(j)} = 2 \lceil q_{\theta}(u)^{(j)} - \mu^{(j)} \rceil - 1 \quad (17)$$

where θ is the set of parameters for user encoding. During training and testing, we use the same sampling strategy as for the item

encoding. For both users and items, we use a straight-through estimator [4] for computing the gradients for backpropagation through the sampled hash codes.

3.4 Decoder Functions

3.4.1 User-item rating decoding. The first decoding step aims to reconstruct the original user-item rating $R_{u,i}$, which corresponds to computing the conditional log likelihood of Eq. 9, i.e., $\log p(R_{u,i}|z_i, z_u)$. We first transform the user-item rating into the same range as the inner product between the hash codes:

$$\hat{R}_{u,i} = 2m \frac{R_{u,i}}{\text{max rating}} - m \quad (18)$$

Similarly to [23, 27], we assume the ratings are Gaussian distributed around their true mean for each rating value, such that we can compute the conditional log likelihood as:

$$\log p(R_{u,i}|z_i, z_u) = \log \mathcal{N}(\hat{R}_{u,i} - z_i^T z_u, \sigma^2) \quad (19)$$

where the variance σ^2 is constant, thus providing an equal weighting of all ratings. However, the exact value of the variance is irrelevant, since maximizing Eq. 19 corresponds to simply minimizing the squared error (MSE) of the mean term, i.e., $\hat{R}_{u,i} - z_i^T z_u$. Thus, maximizing the log likelihood is equivalent to minimizing the MSE, as similarly done in related work [22, 37, 39]. Lastly, note that due to the equivalence between the inner product and the Hamming distance (see Eq. 2), this directly optimizes the hash codes for the Hamming distance.

3.4.2 Item content decoding. The secondary decoding step aims to reconstruct the original content features given the generated item hash code in Eq. 10, i.e., $\log p(c_i|z_i)$. We compute this as the summation of word log likelihoods (based on Eq. 5) using a softmax:

$$\log p(c_i|z_i) = \sum_{w \in \mathbb{W}_{c_i}} \log \frac{e^{z_i^T (E_{\text{word}}(1_w \odot w_{\text{imp}})) + b_w}}{e^{\sum_{w' \in \mathbb{W}} z_i^T (E_{\text{word}}(1_{w'} \odot w_{\text{imp}})) + b_{w'}}} \quad (20)$$

where 1_w is a one-hot encoding for word w , \mathbb{W} is the set of all vocabulary words of the content feature vectors, $E_{\text{word}} \in \mathbb{R}^{|\mathbb{W}| \times m}$ is a learned word embedding, b_w is a word-level bias term, and the learned importance weight w_{imp} is the same as in Eq. 12. This softmax expression is maximized when the item hash codes are able to decode the original content words.

3.4.3 Noise infusion for robustness. Previous work on semantic hashing has shown that infusing random noise into the hash codes before decoding increases robustness, and leads to more generalizable hash codes [6, 13, 30]. Thus, we apply a Gaussian noise to both user and item hash codes before decoding:

$$\text{noise}(z, \sigma^2) = z + \epsilon \sigma^2, \quad \epsilon \sim \mathcal{N}(0, I) \quad (21)$$

where variance annealing is used for decreasing the initial value of σ^2 in each training iteration.

3.5 Combined Loss Function

NeuHash-CF can be trained in an end-to-end fashion by maximising the combination of the variational lower bounds from Eq. 9 and 10, corresponding to the following loss:

$$\mathcal{L} = \mathcal{L}_{\text{rating}} + \alpha \mathcal{L}_{\text{content}} \quad (22)$$

Table 1: Dataset statistics after preprocessing such that each user has at least rated 20 items, and each item has at least been rated by 20 users.

Dataset	#users	#items	#ratings	sparsity
Yelp	27,147	20,266	1,293,247	99.765%
Amazon	35,736	38,121	1,960,674	99.856%

where $\mathcal{L}_{\text{rating}}$ corresponds to the lower bound in Eq. 9, $\mathcal{L}_{\text{content}}$ corresponds to the lower bound in Eq. 10, and α is a tunable hyperparameter to control the importance of decoding the item content.

4 EXPERIMENTAL EVALUATION

4.1 Datasets

We evaluate our approach on well-known and publicly available datasets with explicit feedback, where we follow the same preprocessing as related work [22, 33, 39] as described in the following. We disallow users to have rated the same item multiple times and use only the last rating in these cases. Due to the very high sparsity of these types of datasets, we apply a filtering to densify the dataset. We remove users who have rated fewer than 20 items, as well items that have been rated by fewer than 20 users. Since the removal of either a user or item may violate the density requirements, we apply the filtering iteratively until all users and items satisfy the requirement. The datasets are described below and summarized in Table 1:

Yelp is from the Yelp Challenge¹, which consists of user ratings and textual reviews on locations such as hotels, restaurants, and shopping centers. User ratings range between 1 (worst) to 5 (best), and most ratings are associated with a textual review.

Amazon [15] is from a collection of book reviews from Amazon². Similarly to Yelp, each user rates a number of books between 1 to 5, and most are accompanied by a textual review as well.

Similarly to related work [22, 33, 39], to obtain content information related to each item, we use the textual reviews (when available) by users for an item. We filter stop words and aggregate all textual reviews for each item into a single large text, and compute the TF-IDF bag-of-words representations, where the top 8000 unique words are kept as the content vocabulary. We apply this preprocessing step separately on each dataset, thus resulting in two different vocabularies.

4.2 Experimental Design

Following Wang and Blei [33], we use two types of recommendations settings: 1) in-matrix regression for estimating the relevance of known items with existing ratings, and 2) out-of-matrix regression for estimating the relevance of cold-start items. Both of these recommendation types lead to different evaluation setups as described next.

¹<https://www.yelp.com/dataset/challenge>

²<http://jmcauley.ucsd.edu/data/amazon/>

4.2.1 In-matrix regression. In-matrix regression can be considered the standard setup of all items (and users) being known at all times, and thus corresponds to the setting solvable by standard collaborative filtering. We split each user’s items into a training and testing set using a 50/50 split, and use 15% of the training set as a validation set for hyper parameter tuning.

4.2.2 Out-of-matrix regression. Out-of-matrix regression is also known as a cold-start setting, where new items are to be recommended. In comparison to in-matrix regression, this task cannot be solved by standard collaborative filtering. We sort all items by their number of ratings, and then proportionally split them 50/50 into a training and testing set, such that each set has approximately the same number of items with similar number of ratings. Similarly to the in-matrix regression setting, we use 15% of the training items as a validation set for hyper parameter tuning.

4.3 Evaluation Metrics

We evaluate the effectiveness of our approach and the baselines as a ranking task with the aim of placing the most relevant (i.e., highest rated) items at the top of a ranked list. As detailed in Section 4.2, each user has a number of rated items, such that the ranked list is produced by sorting each user’s testing items by their Hamming distance between the user and item hash codes. To measure the quality of the ranked list, we use Normalized Discounted Cumulative Gain (NDCG), which incorporates both ranking precision and the position of ratings. Secondly, we are interested in the first position of the item with the highest rating, as this ideally should be in the top. To this end, we compute the Mean Reciprocal Rank (MRR) of the highest ranked item with the highest given rating from the user’s list of testing items.

4.4 Baselines

We compare NeuHash-CF against existing state-of-the-art content-aware hashing-based recommendation approaches, as well as hashing-based approaches that are not content-aware to highlight the benefit of including content:

DCMF Discrete Content-aware Matrix Factorization [22]³ is a content-aware matrix factorization technique, which is discretized and optimized through solving multiple mixed-integer subproblems. Similarly to our approach, its primary objective is to minimize the squared error between the rating and estimated rating based on the Hamming distance. It also learns a latent representation for each word in the text associated to each item, which is used for generating hash codes for cold-start items.

DDL Discrete Deep Learning [39]⁴ also uses an alternating optimizing strategy for solving multiple mixed-integer subproblems, where the primary objective is a mean squared error loss. In contrast to DCMF, DDL uses a deep belief network for generating cold-start item hash codes, which is trained by learning to map the content of known items into their hash codes generated in the first part of the approach.

DCF Discrete Collaborative Filtering [37]⁵ can be considered the predecessor to DCMF, but is not content-aware, which was the primary novelty of DCMF.

NeuHash-CF/no.C We include a version of our NeuHash-CF that is not content-aware, which is done by simply learning item hash codes similarly to user hash codes, thus not including any content features.

For both DCMF and DDL, hash codes for cold-start items are seen as a secondary objective, as they are generated differently from non-cold-start item hash codes. In contrast, our NeuHash-CF treats all items identically as all item hash codes are generated based on content features alone.

To provide a comparison to non-hashing based approaches, which are notably more computationally expensive for making recommendations (see Section 4.7), we also include the following baselines:

FM Factorization Machines [26] works on a concatenated n -dimensional vector of the one-hot encoded user id, one-hot encoder item id, and the content features. It learns latent vectors, as well as scalar weights and biases for each of the n dimensions. FM estimates the user-item relevance by computing a weighted sum of all non-zero entries and all interactions between non-zero entries of the concatenated vector. This results in a large amount of inner product computations and a large storage cost associated with the latent representations and scalars. We use the FastFM implementation [3]⁶.

MF Matrix Factorization [20] is a classic non-content-aware collaborative filtering approach, which learns real-valued item and user latent vectors, such that the inner product corresponds to the user-item relevance. MF is similar to a special case of FM without any feature interactions.

4.5 Tuning

For training our NeuHash-CF approach, we use the Adam [18] optimizer with learning rates selected from $\{0.0005, 0.0001\}$ and batch sizes from $\{500, 1000, 2000\}$, where 0.0005 and 2000 were consistently chosen. We also tune the number of encoder layers from $\{1, 2, 3\}$ and the number of neurons in each from $\{500, 1000, 2000\}$; most runs had the optimal validation performance with 2 layers and 1000 neurons. To improve robustness of the codes we added Gaussian noise before decoding the hash codes, where the variance was initially set to 1 and decreased by 0.01% after every batch. Lastly, we tune α in Eq. 22 from $\{0.001, 0.01, 0.1\}$, where 0.001 was consistently chosen. The code⁷ is written in TensorFlow [1]. For all baselines, we tune the hyper parameters on the validation set as described in the original papers.

4.6 Results

The experimental comparison is summarized in Table 2 and 3 for NDCG@ $\{2, 6, 10\}$ and MRR, respectively. The tables are split into in-matrix and out-of-matrix evaluation settings for both datasets, and the methods can be categorized into groups: (1) content-aware

³<https://github.com/DefuLian/recsys/tree/master/alg/discrete/dcmf>

⁴<https://github.com/yixianqianzy/ddl>

⁵<https://github.com/hanwangzhang/Discrete-Collaborative-Filtering>

⁶<https://github.com/ibayer/fastFM>

⁷We make the code publicly available at <https://github.com/casperhansen/NeuHash-CF>

Table 2: NDCG@k scores on in-matrix and out-of-matrix settings for the Amazon and Yelp datasets. Bold numbers represent the best hashing-based approach and statistically significant results compared to the best hashing-based baseline per column are marked with a star. Dashed lines correspond to not content-aware approaches in out-of-matrix setting.

NDCG	Yelp (in-matrix)									Yelp (out-of-matrix)								
	16 dim.			32 dim.			64 dim.			16 dim.			32 dim.			64 dim.		
	@2	@6	@10	@2	@6	@10	@2	@6	@10	@2	@6	@10	@2	@6	@10	@2	@6	@10
NeuHash-CF	.662*	.701*	.752*	.681*	.718*	.766*	.697*	.731*	.776*	.646*	.694*	.747*	.687*	.725*	.772*	.702*	.737*	.780*
DCMF	.642	.678	.733	.655	.691	.743	.670	.701	.752	.611	.647	.703	.617	.655	.709	.626	.664	.717
DDL	.636	.674	.729	.651	.686	.739	.664	.698	.749	.575	.615	.673	.579	.622	.681	.612	.646	.700
NeuHash-CF/no.C	.634	.672	.727	.655	.689	.741	.666	.699	.749	-	-	-	-	-	-	-	-	-
DCF	.639	.676	.730	.649	.685	.738	.671	.700	.750	-	-	-	-	-	-	-	-	-
MF (real-valued)	.755*	.763*	.800*	.755*	.763*	.800*	.755*	.763*	.800*	-	-	-	-	-	-	-	-	-
FM (real-valued)	.754*	.763*	.801*	.750*	.760*	.798*	.744*	.755*	.794*	.731*	.750*	.789*	.724*	.744*	.785*	.719*	.740*	.781*

NDCG	Amazon (in-matrix)									Amazon (out-of-matrix)								
	16 dim.			32 dim.			64 dim.			16 dim.			32 dim.			64 dim.		
	@2	@6	@10	@2	@6	@10	@2	@6	@10	@2	@6	@10	@2	@6	@10	@2	@6	@10
NeuHash-CF	.759*	.777*	.810*	.780*	.798*	.827*	.786*	.803*	.831*	.758*	.778*	.809*	.769*	.788*	.818*	.787*	.804*	.831*
DCMF	.749	.767	.800	.761	.777	.810	.773	.788	.818	.727	.748	.782	.729	.749	.784	.733	.752	.786
DDL	.734	.755	.791	.748	.768	.802	.762	.779	.811	.704	.728	.766	.705	.729	.767	.705	.727	.766
NeuHash-CF/no.C	.748	.768	.802	.760	.776	.808	.771	.785	.816	-	-	-	-	-	-	-	-	-
DCF	.745	.767	.802	.759	.776	.809	.774	.787	.818	-	-	-	-	-	-	-	-	-
MF (real-valued)	.824*	.826*	.848*	.824*	.826*	.848*	.824*	.826*	.848*	-	-	-	-	-	-	-	-	-
FM (real-valued)	.821*	.822*	.845*	.817*	.819*	.843*	.813*	.816*	.841*	.792*	.800*	.827*	.785*	.793*	.821*	.780*	.790*	.819*

Table 3: MRR scores in both in-matrix and out-of-matrix settings. Bold numbers represent the best hashing-based approach and statistically significant results compared to the best hashing-based baseline per column are marked with a star. Dashed lines correspond to not content-aware approaches in out-of-matrix setting.

MRR	Yelp (in-matrix)			Yelp (out-of-matrix)			Amazon (in-matrix)			Amazon (out-of-matrix)		
	16 dim.	32 dim.	64 dim.	16 dim.	32 dim.	64 dim.	16 dim.	32 dim.	64 dim.	16 dim.	32 dim.	64 dim.
NeuHash-CF	.646*	.668*	.687*	.628*	.674*	.692*	.749*	.770*	.779*	.750*	.764*	.782*
DCMF	.629	.644	.660	.598	.604	.612	.738	.753	.767	.719	.721	.726
DDL	.620	.638	.651	.557	.562	.604	.721	.741	.753	.696	.694	.694
NeuHash-CF/no.C	.621	.642	.656	-	-	-	.737	.752	.764	-	-	-
DCF	.626	.636	.664	-	-	-	.736	.751	.769	-	-	-
MF (real-valued)	.767*	.767*	.767*	-	-	-	.826*	.826*	.826*	-	-	-
FM (real-valued)	.761*	.756*	.750*	.730*	.722*	.717*	.824*	.821*	.815*	.792*	.784*	.780*

(NeuHash-CF, DCMF, DDL), (2) not content-aware (NeuHash-CF/no.C, DCF), (3) real-valued not content-aware (MF), and (4) real-valued content-aware (FM). For all methods, we compute hash codes (or latent representations for MF and FM) of length $m \in \{16, 32, 64\}$. We use a two-tailed paired t-test for statistical significance testing against the best performing hashing-based baseline. Statistically significant improvements, at the 0.05 level, over the best performing hashing-based baseline per column are marked with a star (*), and the best performing hashing-based approach is shown in bold.

4.6.1 In-matrix regression. In the in-matrix setting, where all items have been rated in the training data, our NeuHash-CF significantly outperforms all hashing-based baselines. On Yelp, we observe improvements in NDCG by up to 0.03, corresponding to a 4.3% improvement. On Amazon, we observe improvements in NDCG by up to 0.02, corresponding to a 2.7% improvement. Similar improvements are noted on both datasets for MRR (1.6-4.1% improvements). On all datasets and across the evaluated dimensions, NeuHash-CF performs similarly or better than state-of-the-art hashing-based

approaches while using 2-4 times fewer bits, thus providing both a significant performance increase as well as a 2-4 times storage reduction. Interestingly, the performance gap between existing content-aware and not content-aware approaches is relatively small. When considering the relative performance increase of our NeuHash-CF with and without content features, we see the benefit of basing the item hash codes directly on the content. DCMF and DDL both utilize the content features for handling cold-start items, but not to the same degree for the in-matrix items, which we argue explains the primary performance increase observed for NeuHash-CF, since NeuHash-CF/no.C performs similarly to the baselines.

We also include MF and FM as real-valued baselines to better gauge the discretization gap. As expected, the real-valued approaches outperform the hashing-based approaches, however as the number of bit increases the performance difference decreases. This is to be expected, since real-valued approaches reach faster a potential representational limit, where more dimensions would not positively impact the ranking performance. In fact, for FM we observe a marginal performance drop when increasing its number of

latent dimensions, thus indicating that it is overfitting. In contrast, MF keeps the same performance (differing on far out decimals) independently of its number of latent dimensions.

4.6.2 Out-of-matrix regression. We now consider the out-of-matrix setting, corresponding to recommending cold-start items. NeuHash-CF significantly outperforms the existing state-of-the-art hashing-based baselines even more than for the in-matrix setting. On Yelp, we observe the smallest NDCG increase for 16 bit at 0.035, which is however doubled in most cases for 32 and 64 bits, corresponding to improvements of up to 12.1% gain over state-of-the-art baselines. We observe a similar trend on Amazon, where the lowest improvement of 0.027 NDCG is observed at 16 bits, but increasing the number of bits leads to consistently larger improvements of up to 7.4%. These results are also consistent with MRR, where increasing the number of bits provides increasingly larger performance increases between +5 and +13.1% on Yelp and between +4.3 and +7.7% on Amazon. In all cases, the performance of NeuHash-CF on 16 bits is even better than the best baseline at 64 bits, thus verifying the high quality of the hash codes generated by NeuHash-CF.

For the real-valued FM baseline, we observe that it outperforms ours and existing baselines at 16 and 32 dimensions, however at 64 dimensions NeuHash-CF outperforms FM on Amazon for $\text{NDCG}@6, 10$ (across all dimensions). When we consider Yelp, NeuHash-CF obtains a $\text{NDCG}@10$ within 0.01 of FM, but worse on the other NDCG cut offs and on MRR.

4.6.3 Out-of-matrix regression with limited training data. To evaluate how the content-aware approaches generalize to the cold-start setting depending on the number of training items, we furthermore create smaller versions of the 50/50 out-of-matrix split used previously. In addition to using 50% of the data for the training set, we consider splits using 10%, 20%, 30%, and 40% as well. In all out-of-matrix settings the validation and testing sets are identical to be able to compare the impact of the training size. The results can be seen in Table 4 for 32 bit hash codes and 32 latent dimensions in FM. Similarly to before, NeuHash-CF outperforms the hashing-based baselines in all cases with similar gains as observed previously. Most approaches, except DDL on Amazon, obtain the lowest performance using 10% of the data, and more training items generally improve the performance, although at 30-50% the pace of improvement slows down significantly. This indicates that the methods have observed close to sufficiently many training items and increasing the amount may not lead to better generalizability of the cold-start hash codes. Interestingly, NeuHash-CF obtains the largest improvement going from 10% to 50% on both NDCG and MRR, indicating that it generalizes better than the baselines. In contrast, DDL does not improve on Amazon by including more training items, which indicates that its ability to generalize to cold-start items is rather limited.

4.7 Computational Efficiency

To study the high efficiency of using hash codes in comparison to real-valued vectors, we consider a setup of 100,000 users and 1,000-1,000,000 items. We randomly generate hash codes and real-valued vectors and measure the time taken to compute all Hamming distances (or inner products) from each user to all items, resulting in




Figure 2: Computation time for all Hamming distances and inner products for 100,000 users and up to 1,000,000 items.

a total 10^8 - 10^{11} computations. We use a machine with a 64 bit instruction set⁸, and hence generate hash codes and vectors of length 64. We report the average runtime over 10 repetitions in Figure 2, and observe a speed up of a factor 40-50 for the Hamming distance, highlighting the efficiency benefit of hashing-based approaches. For FM, its dominating cost is its large number of inner product computations, which scales quadratically in the number of non-zero content features for a given item, thus making it highly intractable in large-scale settings.

4.8 Impact of Average Item Popularity per User

We now look at how different user characteristics impact the performance of the methods. We first compute the average item popularity of each user’s list of rated items, and then order the users in ascending order of that average. An item’s popularity is computed as the number of users who have rated that specific item, and thus the average item popularity of a user is representative of their attraction to popular content. Figure 3 plots the $\text{NDCG}@10$ for 32 dimensional representations using a mean-smoothing window size of 1000 (i.e., each shown value is averaged based on the values within a window of 1000 users). Generally, all methods perform better for users who have a high average item popularity, where for Yelp we see a $\text{NDCG}@10$ difference of up to 0.25 from the lowest to highest average popularity (0.2 for Amazon). This observation can be explained by highly popular items occurring more times in the training data, such that they have a better learned representation. Additionally, the hashing-based approaches have a larger performance difference, compared to the real-valued MF and FM, which is especially due to their lower relative performance for users with a very low average item popularity (left side of plots). In the out-of-matrix setting the same trend is observed, however with our NeuHash-CF performing highly similarly to FM when excluding the users with the lowest average item popularity. We hypothesize that users with a low average item popularity have a more specialized preference, thus benefitting more from the higher representational power of real-valued representations.

4.9 Impact of Number of Items per User

We now consider how the number of items each user has rated impacts performance. We order users by their number of rated items and plot $\text{NDCG}@10$ for 32 bit hash codes. Figure 4 plots this in the

⁸We used an Intel Xeon CPU E5-2670

Table 4: NDCG@10 and MRR scores for 32 dimensional representations in varying cold-start scenarios with 10-50% of the items used for training. Bold numbers represent the best hashing-based approach and statistically significant results compared to the best hashing-based baseline in each column are marked with a star.

NDCG	Yelp (out-of-matrix)										Amazon (out-of-matrix)									
	10%		20%		30%		40%		50%		10%		20%		30%		40%		50%	
	@10	MRR	@10	MRR	@10	MRR	@10	MRR	@10	MRR	@10	MRR	@10	MRR	@10	MRR	@10	MRR	@10	MRR
NeuHash-CF	.730*	.603*	.750*	.634*	.769*	.666*	.771*	.668*	.772*	.674*	.794*	.727*	.812*	.753*	.817*	.761*	.818*	.763*	.818*	.764*
DCMF	.688	.572	.693	.578	.704	.593	.710	.602	.709	.604	.774	.710	.778	.712	.781	.717	.784	.720	.784	.721
DDL	.678	.556	.681	.562	.687	.572	.684	.571	.681	.562	.770	.713	.766	.689	.767	.700	.765	.693	.767	.694
FM (real-valued)	.766*	.688*	.776*	.707*	.778*	.712*	.786*	.724*	.785*	.722*	.806*	.759*	.813*	.771*	.817*	.775*	.823*	.786*	.821*	.784*




Figure 3: Impact of the average item popularity per user on NDCG@10 for 32 bit hash codes.




Figure 4: Impact of the number of items per user on NDCG@10 for 32 bit hash codes.

same way as in Figure 3. Generally across all methods, we observe that performance initially increases, but then drops once a user has rated close to 100 items, depending on the dataset. While the hashing-based approaches keep steadily dropping in performance, MF and FM do so at a slower pace and even increase for users with the highest number of rated items in the in-matrix setting. The plots clearly show that the largest performance difference, between the real-valued and hashing-based approaches, is for the group of users with a high number of rated items, corresponding to users with potentially the highest diversity of interests. In this setting, the limited representational power of hash codes, as opposed to real-valued representations, may not be sufficient to encode users with largely varied interests. We observe very similar trends for the out-of-matrix setting for cold-start items, although the performance gap between our NeuHash-CF and the real-valued approaches is almost entirely located among the users with a high number of rated items.

5 CONCLUSION

We presented content-aware neural hashing for collaborative filtering (NeuHash-CF), a novel hashing-based recommendation approach, which is robust to cold-start recommendation problems (i.e., the setting where the items to be recommended have not been rated previously). NeuHash-CF is a neural approach that consists of two joint components for generating user and item hash codes.

The user hash codes are learned from an embedding based procedure using only the user’s id, whereas the item hash codes are learned directly from associated content features (e.g., a textual item description). This contrasts existing state-of-the-art content-aware hashing-based methods [22, 39], which generate item hash codes differently depending on whether they are cold-start items or not. NeuHash-CF is formulated as a variational autoencoder architecture, where both user and item hash codes are sampled from learned Bernoulli distributions to enforce end-to-end trainability. We presented a comprehensive experimental evaluation of NeuHash-CF in both standard and cold-start settings, where NeuHash-CF outperformed state-of-the-art approaches by up to 12% NDCG and 13% MRR in cold-start recommendation (up to 4% in both NDCG and MRR in standard recommendation settings). In fact, the ranking performance of NeuHash-CF on 16 bit hash codes is better than that of 32-64 bit state-of-the-art hash codes, thus resulting in both a significant effectiveness increase, but also in a 2-4x storage reduction. Analysis of our results showed that the largest performance difference between hashing-based and real-valued approaches occurs for users interested in the least popular items, and for the group of users with the highest number of rated items. Future work includes extending the architecture to accept richer item and user representations, such as [8, 12, 25, 32].

REFERENCES

- [1] Martin Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. 2016. Tensorflow: A system for large-scale machine learning. In *USENIX Symposium on Operating Systems Design and Implementation (OSDI)*. 265–283.
- [2] Gediminas Adomavicius and Alexander Tuzhilin. 2005. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge & Data Engineering* 6 (2005), 734–749.
- [3] Immanuel Bayer. 2016. fastFM: A Library for Factorization Machines. *Journal of Machine Learning Research* 17, 184 (2016), 1–5.
- [4] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. 2013. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432* (2013).
- [5] James Bennett, Stan Lanning, et al. 2007. The netflix prize. In *KDD cup and workshop*, Vol. 2007. 35.
- [6] Suthesh Chaidaroon, Travis Ebesu, and Yi Fang. 2018. Deep Semantic Text Hashing with Weak Supervision. In *International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1109–1112.
- [7] Suthesh Chaidaroon and Yi Fang. 2017. Variational deep semantic hashing for text documents. In *International ACM SIGIR Conference on Research and Development in Information Retrieval*. 75–84.
- [8] Felipe Soares Da Costa and Peter Dolog. 2019. Collective embedding for neural context-aware recommender systems. In *ACM Conference on Recommender Systems*. 201–209.
- [9] Abhinandan S Das, Mayur Datar, Ashutosh Garg, and Shyam Rajaram. 2007. Google news personalization: scalable online collaborative filtering. In *ACM Conference on World Wide Web*. 271–280.
- [10] Aristides Gionis, Piotr Indyk, Rajeev Motwani, et al. 1999. Similarity search in high dimensions via hashing. In *Vldb*, Vol. 99. 518–529.
- [11] Yunchao Gong, Svetlana Lazebnik, Albert Gordo, and Florent Perronnin. 2012. Iterative quantization: A procrustean approach to learning binary codes for large-scale image retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35, 12 (2012), 2916–2929.
- [12] Casper Hansen, Christian Hansen, Stephen Alstrup, Jakob Grue Simonsen, and Christina Lioma. 2019. Contextually Propagated Term Weights for Document Representation. In *International ACM SIGIR Conference on Research and Development in Information Retrieval*. 897–900.
- [13] Casper Hansen, Christian Hansen, Jakob Grue Simonsen, Stephen Alstrup, and Christina Lioma. 2019. Unsupervised Neural Generative Semantic Hashing. In *International ACM SIGIR Conference on Research and Development in Information Retrieval*. 735–744.
- [14] Casper Hansen, Christian Hansen, Jakob Grue Simonsen, Stephen Alstrup, and Christina Lioma. 2020. Unsupervised Semantic Hashing with Pairwise Reconstruction. In *International ACM SIGIR Conference on Research and Development in Information Retrieval*. in press.
- [15] Ruining He and Julian McAuley. 2016. Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering. In *ACM Conference on World Wide Web*. 507–517.
- [16] Geoffrey E Hinton, Simon Osindero, and Yee-Whye Teh. 2006. A fast learning algorithm for deep belief nets. *Neural computation* 18, 7 (2006), 1527–1554.
- [17] Alexandros Karatzoglou, Alex Smola, and Markus Weimer. 2010. Collaborative filtering on a budget. In *International Conference on Artificial Intelligence and Statistics*. 389–396.
- [18] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*.
- [19] Diederik P Kingma and Max Welling. 2014. Auto-encoding variational bayes. In *International Conference on Learning Representations*.
- [20] Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix Factorization Techniques for Recommender Systems. *Computer* 42, 8 (Aug. 2009), 30–37.
- [21] Defu Lian, Yong Ge, Fuzheng Zhang, Nicholas Jing Yuan, Xing Xie, Tao Zhou, and Yong Rui. 2015. Content-aware collaborative filtering for location recommendation based on human mobility data. In *IEEE international conference on data mining*. 261–270.
- [22] Defu Lian, Rui Liu, Yong Ge, Kai Zheng, Xing Xie, and Longbing Cao. 2017. Discrete Content-aware Matrix Factorization. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 325–334.
- [23] Dawen Liang, Rahul G Krishnan, Matthew D Hoffman, and Tony Jebara. 2018. Variational autoencoders for collaborative filtering. In *ACM Conference on World Wide Web*. 689–698.
- [24] Chenghao Liu, Tao Lu, Xin Wang, Zhiyong Cheng, Jianling Sun, and Steven C.H. Hoi. 2019. Compositional Coding for Collaborative Filtering. In *International ACM SIGIR Conference on Research and Development in Information Retrieval*. 145–154.
- [25] Ahmed Rashed, Josif Grabocka, and Lars Schmidt-Thieme. 2019. Attribute-aware non-linear co-embeddings of graph features. In *ACM Conference on Recommender Systems*. 314–321.
- [26] Steffen Rendle. 2010. Factorization machines. In *International Conference on Data Mining*. IEEE, 995–1000.
- [27] Noveen Sachdeva, Giuseppe Manco, Ettore Ritacco, and Vikram Pudi. 2019. Sequential Variational Autoencoders for Collaborative Filtering. In *Proceedings of the ACM International Conference on Web Search and Data Mining*. 600–608.
- [28] Ruslan Salakhutdinov and Geoffrey Hinton. 2009. Semantic hashing. *International Journal of Approximate Reasoning* 50, 7 (2009), 969–978.
- [29] Ying Shan, Jie Zhu, JC Mao, et al. 2018. Recurrent binary embedding for gpu-enabled exhaustive retrieval from billion-scale semantic vectors. In *ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2170–2179.
- [30] Dinghan Shen, Qinliang Su, Paidamoyo Chapfuwa, Wenlin Wang, Guoyin Wang, Ricardo Henao, and Lawrence Carin. 2018. NASH: Toward End-to-End Neural Architecture for Generative Semantic Hashing. In *Annual Meeting of the Association for Computational Linguistics*. 2041–2050.
- [31] Yue Shi, Martha Larson, and Alan Hanjalic. 2014. Collaborative filtering beyond the user-item matrix: A survey of the state of the art and future challenges. *ACM Computing Surveys (CSUR)* 47, 1 (2014), 3.
- [32] Benyou Wang, Donghao Zhao, Christina Lioma, Qiuchi Li, Peng Zhang, and Jakob Grue Simonsen. 2020. Encoding word order in complex embeddings. In *International Conference on Learning Representations*.
- [33] Chong Wang and David M Blei. 2011. Collaborative topic modeling for recommending scientific articles. In *ACM SIGKDD international conference on Knowledge discovery and data mining*. 448–456.
- [34] Yair Weiss, Antonio Torralba, and Rob Fergus. 2009. Spectral hashing. In *Advances in neural information processing systems*. 1753–1760.
- [35] Dell Zhang, Jun Wang, Deng Cai, and Jinsong Lu. 2010. Laplacian co-hashing of terms and documents. In *European Conference on Information Retrieval*. Springer, 577–580.
- [36] Dell Zhang, Jun Wang, Deng Cai, and Jinsong Lu. 2010. Self-taught hashing for fast similarity search. In *International ACM SIGIR Conference on Research and Development in Information Retrieval*. 18–25.
- [37] Hanwang Zhang, Fumin Shen, Wei Liu, Xiangnan He, Huanbo Luan, and Tat-Seng Chua. 2016. Discrete collaborative filtering. In *International ACM SIGIR Conference on Research and Development in Information Retrieval*. 325–334.
- [38] Yan Zhang, Defu Lian, and Guowu Yang. 2017. Discrete personalized ranking for fast collaborative filtering from implicit feedback. In *AAAI Conference on Artificial Intelligence*. 1669–1675.
- [39] Yan Zhang, Hongzhi Yin, Zi Huang, Xingzhong Du, Guowu Yang, and Defu Lian. 2018. Discrete Deep Learning for Fast Content-Aware Recommendation. In *ACM International Conference on Web Search and Data Mining*. 717–726.
- [40] Zhiwei Zhang, Qifan Wang, Lingyun Ruan, and Luo Si. 2014. Preference preserving hashing for efficient recommendation. In *International ACM SIGIR Conference on Research and Development in Information Retrieval*. 183–192.
- [41] Ke Zhou and Hongyuan Zha. 2012. Learning binary codes for collaborative filtering. In *ACM SIGKDD international conference on Knowledge discovery and data mining*. 498–506.

Chapter 6

Projected Hamming Dissimilarity for Bit-Level Importance Coding in Collaborative Filtering

Christian Hansen*, Casper Hansen*, Jakob Grue Simonsen, Christina Lioma (2021).
Projected Hamming Dissimilarity for Bit-Level Importance Coding in Collaborative
Filtering. In WWW, in press. [28]. * denotes equal contribution.

Projected Hamming Dissimilarity for Bit-Level Importance Coding in Collaborative Filtering

Christian Hansen*
University of Copenhagen
chrh@di.ku.dk

Jakob Grue Simonsen
University of Copenhagen
simonsen@di.ku.dk

Casper Hansen*
University of Copenhagen
c.hansen@di.ku.dk

Christina Lioma
University of Copenhagen
c.lioma@di.ku.dk

ABSTRACT

When reasoning about tasks that involve large amounts of data, a common approach is to represent data items as objects in the Hamming space where operations can be done efficiently and effectively. Object similarity can then be computed by learning binary representations (hash codes) of the objects and computing their Hamming distance. While this is highly efficient, each bit dimension is equally weighted, which means that potentially discriminative information of the data is lost. A more expressive alternative is to use real-valued vector representations and compute their inner product; this allows varying the weight of each dimension but is many magnitudes slower. To fix this, we derive a new way of measuring the dissimilarity between two objects in the Hamming space *with* binary weighting of each dimension (i.e., disabling bits): we consider a field-agnostic dissimilarity that projects the vector of one object onto the vector of the other. When working in the Hamming space, this results in a novel projected Hamming dissimilarity, which by choice of projection, effectively allows a binary importance weighting of the hash code of one object through the hash code of the other. We propose a variational hashing model for learning hash codes optimized for this projected Hamming dissimilarity, and experimentally evaluate it in collaborative filtering experiments. The resultant hash codes lead to effectiveness gains of up to +7% in NDCG and +14% in MRR compared to state-of-the-art hashing-based collaborative filtering baselines, while requiring no additional storage and no computational overhead compared to using the Hamming distance.

KEYWORDS

importance coding; hash codes; collaborative filtering

ACM Reference Format:

Christian Hansen, Casper Hansen, Jakob Grue Simonsen, and Christina Lioma. 2021. Projected Hamming Dissimilarity for Bit-Level Importance Coding in Collaborative Filtering. In *Proceedings of the Web Conference 2021 (WWW '21)*, April 19–23, 2021, Ljubljana, Slovenia. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3442381.3450011>

*Both authors share the first authorship.

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW '21, April 19–23, 2021, Ljubljana, Slovenia

© 2021 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC-BY 4.0 License.

ACM ISBN 978-1-4503-8312-7/21/04.

<https://doi.org/10.1145/3442381.3450011>

1 INTRODUCTION

Hashing-based learning aims to find short binary representations of data objects (called *hash codes*) that allow effective and efficient computations. For tasks involving multiple interacting objects, hash codes must be chosen carefully to ensure that certain properties of the codes (e.g., their inner products or mutual distance) carry information relevant to the task. For example, in hashing-based collaborative filtering, binary user and item representations must be learned so that the distance between them reflects how much user u likes item i . Currently, the most efficient way of doing this is to compute their Hamming distance, which is the sum of differing bits between two hash codes. However, by definition, the Hamming distance weighs each bit equally. This is a problem because the importance of the underlying properties encoded by each bit may differ. An alternative is to use real-valued vectors and compute their inner product, which allows varying the weight of each dimension, and thus enables a dimension-specific importance weighting not possible using the Hamming distance for binary codes. However, binary codes allow large storage reductions compared to floating point representations, while the Hamming distance enables massively faster computations compared to the inner product (e.g., real-time brute-force search in a billion items [23]). Motivated by this, we ask: can we introduce bit-level importance weighting on binary representations without compromising efficiency?

We reason that, by the definition of the inner product, the distance between two vectors u and i should be identical to the difference in length between vector u and vector i 's projection on u . This observation can be exploited by using a vector space where projections and lengths can be computed several magnitudes more efficiently than in Euclidean space. We show that performing the exact same projection and length computations in the Hamming vector space over \mathbb{Z}_2 results in a novel projected Hamming dissimilarity, which corresponds to traditional measures used in real-valued representations. By choice of projection in the Hamming space, the projected Hamming dissimilarity effectively allows a bit-level binary weighting (corresponding to disabling bits) of i 's hash code via the hash code of u , but without decreasing efficiency compared to the Hamming distance. We propose a variational hashing model for learning hash codes optimized for our projected Hamming dissimilarity, which we experimentally evaluate in collaborative filtering experiments. Compared to state-of-the-art baselines using the Hamming distance, we observe effectiveness gains of up to +7% in NDCG and +14% in MRR, while also significantly improving the

convergence rate during model training compared to optimizing with the Hamming distance.

In summary, we **contribute** the first alternative to the Hamming distance for hashing-based learning that allows bit-level importance coding, while requiring no additional storage and no computational overhead. We make our code publicly available.¹

2 RELATED WORK

Matrix factorization is one of the most popular collaborative filtering methods [14], but to reduce storage requirements and speed up computation, hashing-based collaborative filtering has been researched. For hashing-based methods, the users and items are represented as binary hash codes (as opposed to real-valued vectors), which traditionally have used the highly efficient Hamming distance (as opposed to the inner product) for computing user-item similarities. In the following, we review the literature on both binary representation learning and importance coding of such binary representations.

2.1 Learning binary representations (hash codes)

Early hashing-based collaborative filtering methods include two stages: First, real-valued representations of the data (vectors) are learned, and then the real-valued vectors are transformed into binary hash codes. Zhang et al. [34] use matrix factorization initially, followed by a binary quantization of rounding the real-valued vectors, while ensuring that the hash code is preference preserving with respect to observed properties of the data using their Constant Feature Norm constraint. Zhou and Zha [35] and Liu et al. [19] both explore binary quantization strategies based on orthogonal rotations of the real-valued vectors, which share similarities with Spectral Clustering [28]. However, the two-stage approaches often suffer from large quantization errors [17, 30], because the hash codes are not learned directly, but rather based on different quantization procedures. More recently, hash codes are learned directly: this has been done using relaxed integer optimization while enforcing bit balancing and decorrelation constraints [30]; and (ii) using an autoencoder to learn the codes in an end-to-end manner [7]. The latter approach is most similar to the variational hashing model proposed in our work for optimizing hash codes for our projected Hamming dissimilarity, but their work is designed for cold-start recommendation based on generalizing hash codes from item descriptions. In contrast, our hashing model is designed to work based purely on user-item ratings, without assuming any additional knowledge of users or items.

2.2 Importance coding of binary representations

Approaches for coding importance into hash codes have been studied for the Hamming distance in several applications. Zhang et al. [32] present an image ranking approach that uses a weighted Hamming distance, where bit-level real-valued weights are learned based on both the discriminative power of each bit across all hash codes,

but also dynamically computed based on the hash code used for querying. The bit-level weights are multiplied on each differing bit between two hash codes, such that the distance is computed as the sum of the weights. Different ways of defining bit-level weights have been explored based on fixed weights per bit [27], fixed weights based on byte-level block differences between hash codes [5], and query-adaptive weights [12, 31]. While fixed weights enable faster retrieval than dynamic weights, they all share the same limitation of being significantly less efficient than the Hamming distance, because they can no longer be expressed using highly efficient Boolean hardware-level operations. Furthermore, in addition to the increased storage requirement due to the weights, transferring the weights to the lowest level of memory (i.e., the register) adds additional computational overhead compared to the Hamming distance.

More recent work addresses the problem that hash codes have reduced representational power compared to real-valued vectors, but increasing the hash code dimensionality to match the amount of bits used in the real-valued case hurts model generalization [17]. An alternative, in the task of collaborative filtering, is Compositional Coding for Collaborative Filtering (CCCF) [17], which is a broadly similar method to learning compositional codes for (word) embedding compression [3, 25]. CCCF is a hybrid of hash codes and real-valued weights: each hash code is split into k blocks of r bits each, and each block is associated with a real-valued scalar indicating the *weight* of the block. The distance between two CCCF hash codes is then computed as a weighted sum of the Hamming distances of the individual blocks, where each weight is the product of each block’s weight. The problem with this approach is that each block requires an individual Hamming distance computation, as well as floating point multiplications of the block weights. In fact, the CCCF block construction no longer allows for highly efficient Boolean operations because the distance computation is weighted by each block’s weight.

In contrast to the above approaches, our projected Hamming dissimilarity can exploit the same highly efficient Boolean operations as the Hamming distance, while enabling a bit-level binary weighting on hash codes without reducing efficiency.

3 BIT-LEVEL IMPORTANCE CODING IN HASH CODES

3.1 Preliminaries

Given two data objects u and i , let $z_u \in \{-1, 1\}^m$ and $z_i \in \{-1, 1\}^m$ denote their hash codes, where m is the number of bits in the hash code, which is typically chosen to fit into a machine word. The Hamming distance d_H between z_u and z_i is defined as:

$$d_H(z_u, z_i) = \sum_{j=1}^m 1_{[z_u^{(j)} \neq z_i^{(j)}]} = \text{SUM}(z_u \text{ XOR } z_i) \quad (1)$$

which can be computed very efficiently due to the Boolean operations on the word level, and the SUM which is computed using the *popcnt* bit string instruction. Because Hamming distance is integer-valued, the Hamming distances between several pairs of objects can be linear-time sorted using e.g. radix sort (Hamming distances must be elements of $[0, m]$, hence they are bounded) in ascending

¹The code is available at <https://github.com/casperhansen/Projected-Hamming-Dissimilarity>

order to create a ranked list, allowing for very fast object similarity computation in data-heavy tasks like information retrieval [23].

The above definition clearly shows the efficiency strength of the Hamming distance, but also its weakness in weighting each bit-dimension equally, even though the underlying data properties encoded by some bits may be more discriminative than those of others. In contrast, representations using real-valued vectors and the inner product allow varying the weight of each dimension, and thus enable a dimension-specific importance weighting not possible using the Hamming distance for binary codes. Next, we show how we derive a new way of measuring the dissimilarity between u and i in the Hamming space *with* binary weighting of each dimension.

3.2 Projected Hamming dissimilarity for bit-level importance coding

Let V be a vector space over any field F , and let $P_{\vec{u}}(\cdot)$ be a projection operator on V , i.e., for each fixed $\vec{u}, \vec{i} \in V$, $P_{\vec{u}}(\cdot) : V \rightarrow V$ is a linear map such that

$$P_{\vec{u}}(P_{\vec{u}}(\vec{i})) = P_{\vec{u}}(\vec{i}). \quad (2)$$

In the following, we consider an asymmetric relationship between two objects \vec{u} and \vec{i} , such that \vec{u} is considered a query object used for searching among data items denoted as \vec{i} . We consider both query and data items as elements of V . Intuitively, each dimension of V corresponds to a property of potential importance to a query (e.g., classical music); the projection of each query on the dimension corresponds to the strength of importance, and the projection of each item on the dimension corresponds to how much the item scores on that dimension.

Let $\|\cdot\| : V \rightarrow \mathbb{R}$ be a norm on V ; we define the *dissimilarity* between \vec{u} and \vec{i} , denoted $\delta(\vec{u}, \vec{i})$, as the norm of the projection of \vec{i} on \vec{u} :

$$\delta(\vec{u}, \vec{i}) = \left\| \vec{u} - P_{\vec{u}}(\vec{i}) \right\| \quad (3)$$

Similarly to Eq. 1, the more different u and i are, the higher their dissimilarity $\delta(\vec{u}, \vec{i})$ should be. The dissimilarity is a natural concept: several existing notions of distance or similarity can be seen as special cases of Eq. 3. For example, in the standard Euclidean space the often-used cosine distance², $1 - \cos(\vec{i}, \vec{u})$, is one instance of Eq. 3 if we assume unit length vectors.

In hashing-based search, we are particularly interested in the binary vector space $V = \{-1, 1\}^m$, with bitwise addition and scalar multiplication over the field $F = \mathbb{Z}_2$, where the projection operator is given by:

$$P_{z_u}(z_i) = z_u \text{ AND } z_i \quad (4)$$

for $z_u, z_i \in V$, i.e., masking the item hash code z_i by the query hash code z_u . Due to working in the Hamming space, the norm $\|\cdot\|$ is chosen as the Hamming norm (sometimes also called the zero norm or L_0). Using this, we obtain the projected Hamming dissimilarity δ_H^P , defined as:

$$\delta_H^P(z_u, z_i) = \left\| z_u - P_{z_u}(z_i) \right\| = \text{SUM}(z_u \text{ XOR } \underbrace{(z_u \text{ AND } z_i)}_{\text{projection}}) \quad (5)$$

²We use the conventional naming of the cosine distance, even though it is not a proper distance as it does not satisfy the triangle inequality.

While having a similar formulation as the Hamming distance (see Eq. 1), the projection of the item hash code z_i onto the query hash code z_u in Eq. 5 allows a binary importance weighting of z_i , which corresponds to disabling unimportant bits as defined by the query hash code z_u (corresponding to the bit-dimensions where the query hash code is -1). We consider bits to be disabled since a -1 bit in z_u leads to *all* item hash codes also having a -1 in that bit after the projection. Note that due to the choice of projection, the projected Hamming dissimilarity is asymmetric (i.e., in general $\delta_H^P(z_u, z_i) \neq \delta_H^P(z_i, z_u)$), whereas the Hamming distance is symmetric.

Compared to the Hamming distance, the projected Hamming dissimilarity fundamentally changes the purpose of the query hash code: instead of each dimension encoding a positive or negative preference for a property, it now encodes which properties of the item are important to the query (-1's from the query hash code are copied to the item due to the AND operation). Thus, this formulation can produce query-specific item representations while still only using a single code for each query and item respectively.

3.2.1 Speeding up the projected Hamming dissimilarity. The projected Hamming dissimilarity in Eq. 5 requires one additional Boolean operation compared to the Hamming distance. However, because the item codes are static once learned, we can reduce the time complexity to the same as the Hamming distance by observing that:

$$\begin{aligned} \delta_H^P(z_u, z_i) &= \text{SUM}(z_u \text{ XOR } (z_u \text{ AND } z_i)) \\ &= \text{SUM}(z_u \text{ AND } (\text{NOT } z_i)) \end{aligned} \quad (6)$$

where $(\text{NOT } z_i)$ can be precomputed and stored instead of the original item hash codes, thus requiring no additional storage and the same number of Boolean operations as the Hamming distance.

Next, we present how the projected Hamming dissimilarity can be used for learning hash codes in collaborative filtering, where, notationally, a *user* takes the role of the query, and *items* are to be recommended based on their relevance to the user.

4 PROJECTED HAMMING DISSIMILARITY IN COLLABORATIVE FILTERING

We propose a variational hashing model for collaborative filtering that learns user and item hash codes optimized for our projected Hamming dissimilarity. To derive a variational framework for hashing-based collaborative filtering, we define the likelihood of a user u as the product over the likelihoods of the observed user specified ratings:

$$p(u) = \prod_{i \in I_u} p(R_{u,i}), \quad p(i) = \prod_{u \in U_i} p(R_{u,i}) \quad (7)$$

where I_u is the set of items rated by user u , and U_i is the set of users who have rated item i . This formulation enforces a dual symmetric effect of users being defined by all their rated items, and items being defined by the ratings provided by all the users. To maximize the likelihood of all observed items and users, we need to maximize the likelihood of the observed ratings $p(R_{u,i})$. Note that instead of maximizing the raw likelihood, we consider the log-likelihood to derive the objective below. We assume that the likelihood of a rating, $p(R_{u,i})$, is conditioned on two latent vectors: a user hash code z_u , and an item hash code z_i . We sample the hash code of

a user or item by performing m Bernoulli trials, which as a prior is assumed to have equal probability of sampling -1 and 1 ($p(z_i)$ and $p(z_u)$ below). This yields the following log-likelihood to be maximized:

$$\log p(R_{u,i}) = \log \left(\sum_{z_i, z_u \in \{-1,1\}^m} p(R_{u,i}|z_u, z_i) p(z_i) p(z_u) \right) \quad (8)$$

However, this is intractable to compute, hence we derive a lower bound. First, we define the learned approximate posterior distributions for the user and item hash codes as $q_\phi(z_i|i)$ and $q_\psi(z_u|u)$, respectively, where ψ and ϕ are the distribution parameters. Next, we multiply and divide with the approximate posterior distributions, rewrite to an expectation, and finally apply Jensen's inequality to obtain a lower bound on the log-likelihood:

$$\begin{aligned} \log p(R_{u,i}) \geq \mathbb{E}_{q_\phi(z_i|i), q_\psi(z_u|u)} \left[\log [p(R_{u,i}|z_u, z_i)] \right. \\ \left. + \log p(z_i) - \log q_\phi(z_i|i) + \log p(z_u) - \log q_\psi(z_u|u) \right] \quad (9) \end{aligned}$$

Since z_i and z_u will be sampled independently, then $q_\phi(z_i|i)$ and $q_\psi(z_u|u)$ are independent and we can rewrite to the variational lower bound:

$$\begin{aligned} \log p(R_{u,i}) \geq \mathbb{E}_{q_\phi(z_i|i), q_\psi(z_u|u)} \left[\log [p(R_{u,i}|z_u, z_i)] \right] \\ - \text{KL}(q_\phi(z_i|i) || p(z_i)) - \text{KL}(q_\psi(z_u|u) || p(z_u)) \quad (10) \end{aligned}$$

where $\text{KL}(\cdot||\cdot)$ is the Kullback-Leibler divergence. Thus, to maximize the expected log-likelihood of the observed rating, we need to maximize the conditional log-likelihood of the rating, while minimizing the KL divergence between the approximate posterior and prior distribution of the two latent vectors. Maximizing the expected conditional log-likelihood can be considered as a reconstruction term of the model, while the KL divergence can be considered as a regularizer.

Next we present the computation of the approximate posterior distributions $q_\phi(z_i|i)$ and $q_\psi(z_u|u)$ (Section 4.1) and the conditional log-likelihood of the rating $p(R_{u,i}|z_u, z_i)$ (Section 4.2).

4.1 Computing the approximate posterior distributions

The approximate posterior distributions can be seen as two encoder functions modelled as embedding layers in a neural network. Each encoder function maps either a user or an item to a hash code. We present below the derivation of the encoder function for the user (the item function is computed in the same way). The probability of the j 'th bit is given by:

$$q_\psi^{(j)}(z_u|u) = \sigma(E_u^{(j)}) \quad (11)$$

where $E_u^{(j)}$ is the j 'th entry in a learned real-valued embedding E for user u , and σ is the sigmoid function. The j 'th bit is then given by:

$$z_u^{(j)} = 2 \lceil \sigma(E_u^{(j)}) - \mu^{(j)} \rceil - 1 \quad (12)$$

where $\mu^{(j)}$ is either chosen stochastically by sampling $\mu^{(j)}$ from a uniform distribution in the interval $[0,1]$, or chosen deterministically to be 0.5 (deterministic choice allows to obtain fixed hash

codes for later evaluation). As the sampling is non-differentiable, a straight-through estimator [1] is used for backpropagation.

4.2 Computing the conditional log-likelihood

The conditional log-likelihood can be considered a reconstruction of the rating, given the user and item hash codes. Similarly to [16], we model the observed ratings as a ground truth rating with additive Gaussian distributed noise, which is then discretized to the observed categorical rating. The conditional log-likelihood can then be computed as:

$$p(R_{u,i}|z_u, z_i) = \mathcal{N}(R_{u,i} - f(z_u, z_i), \sigma^2) \quad (13)$$

where $f(z_u, z_i)$ is a function that reconstructs the rating given the user and item hash codes. Maximising the log-likelihood $\log p(R_{u,i}|z_u, z_i)$ corresponds to minimising the mean squared error (MSE) between $R_{u,i}$ and $f(z_u, z_i)$, which is done for training the model. Existing work on hashing-based collaborative filtering [17, 30] also employs a MSE objective, and thus makes the same Gaussian distribution assumption as done here.

We define the reconstruction function to be our proposed projected Hamming dissimilarity:

$$f(z_u, z_i) = g(\delta_H^P(z_u, z_i)) \quad (14)$$

where g is a fixed affine transformation that maps the interval of the projected Hamming dissimilarity to the interval of the ratings, such that the minimum and maximum of the dissimilarity correspond to the minimum and maximum of the ratings. It should be noted that while variational autoencoders are generative models, we do not explicitly utilize this in the model, as we are primarily concerned with the reconstruction of the observed ratings. This is standard in the related domain of semantic hashing [6, 8, 9, 24].

5 EXPERIMENTAL EVALUATION

We evaluate the effectiveness and efficiency of the projected Hamming dissimilarity for bit-level importance coding in hash codes in collaborative filtering experiments, where the task is to recommend relevant items to users. Items and users are represented as learned hash codes, and user-item relevance is approximated by operations (such as the Hamming distance or the projected Hamming dissimilarity) on those hash codes.

5.1 Datasets

We use 4 publicly available datasets commonly used in prior work [15, 17, 18, 20, 30, 33] and summarized in the bottom of Table 1. The datasets comprise two movie rating datasets, Movielens 1M (ML-1M)³ and Movielens 10M (ML-10M)⁴; a Yelp dataset with ratings of e.g., restaurant and shopping malls⁵; and a book rating dataset from Amazon⁶ [10]. Following Rendle et al. [20], we remove users and items with less than 10 ratings. Following Zhang et al. [30], for each user 50% of the ratings are used for testing, 42.5% for training, and the last 7.5% for validation. If a user rates an item multiple times, only the first rating is kept.

³<https://grouplens.org/datasets/movielens/1m/>

⁴<https://grouplens.org/datasets/movielens/10m/>

⁵<https://www.yelp.com/dataset/challenge>

⁶<http://jmcauley.ucsd.edu/data/amazon/>

5.2 Evaluation metrics

We measure the effectiveness with the mean Normalised Discounted Cumulative Gain (NDCG) [11], which is often used to evaluate recommender systems with non-binary ratings (or relevance values). We use the average NDCG at cutoffs {5, 10} over all users and report the average for each cutoff value. We also compute the Reciprocal Rank (RR) of the highest rated item per user, averaged over all users, which represents how well the approaches are at getting a highly relevant item to the top of the ranked list:

$$\text{DCG}@k = \sum_{i=1}^k \frac{2^{\text{rel}_i} - 1}{\log_2(i + 1)}, \quad \text{NDCG}@k = \frac{\text{DCG}@k}{\text{DCG}@k^{(\text{opt})}}, \quad \text{RR} = \frac{1}{\text{rank}} \quad (15)$$

where rel_i is the relevance of the item in position i of the ranked list of items, $\text{DCG}@k^{(\text{opt})}$ is the $\text{DCG}@k$ of the optimal ranking, and rank is the position of the first highest rated item for a user. These measures are averaged across all users, and following the standard notation, the mean reciprocal rank is denoted as MRR.

5.3 Baselines

We learn hash codes optimized for the projected Hamming dissimilarity by incorporating it in the variational hashing model as described in Section 4 (denoted VH_{PHD}). We compare this to standard and state-of-the-art baselines (listed below) that use Hamming distance in different ways. For reference, we also include real-valued Matrix Factorization (MF)⁷ [14]. MF is not directly comparable to the hashing approaches (it learns latent real-valued vectors for users and items, instead of hash codes, and computes their inner product), but we include it as a point of reference to a real-valued collaborative filtering baseline. In fact, MF has been shown to be both more efficient and highly competitive in effectiveness compared to neural real-valued collaborative filtering approaches [21]. For a fair comparison of MF to the hashing approaches, we set the latent dimension in MF to be the same as the number of bits used in the hashing approaches. We experiment with hash code lengths of 32 and 64 bits, which correspond to the common machine word sizes. The hashing baselines are:

- **DCF**⁸ [30] learns user and item hash codes through binary matrix factorization solved as a relaxed integer problem, while enforcing bit balancing and decorrelation constraints. The codes can be used directly for computing the Hamming distance.
- **CCCF**⁹ [17] learns hash codes of k blocks, where each block has r bits. A floating point weight is given to each block for computing user-item similarities as a weighted sum of block-level Hamming distances. In [17], the floating point weights are mistakenly not counted towards the amount of bits used, thus leading to an unfair advantage. For a fair comparison, we count each floating point weight only as 16 bits (rather than the typical 32 or 64 bits used for single or double precision, respectively).
- **MFmean** and **MFmedian** are based on matrix factorization (MF), but use each dimension’s mean or median for doing the

⁷Included as baseline in the CCCF repository <https://github.com/3140102441/CCCF>

⁸<https://github.com/hanwangzhang/Discrete-Collaborative-Filtering>

⁹<https://github.com/3140102441/CCCF>

binary quantization to bits [29]. Similar to DCF, these codes can be used directly for computing the Hamming distance. We include these to highlight the large quantization loss occurring when the hash codes are not learned directly.

- **VH** is the same variational hashing model that we use for learning hash codes to be used with the projected Hamming dissimilarity, but here the codes are learned using the Hamming distance.

5.4 Tuning

All hyper parameters for the baselines are tuned using the same set of possible values as in the original papers. For CCCF, we use block sizes of {8, 16, 32, 64} and each floating point weight counts for 16 bits. We try all possible combinations that fit within the bit budget, and if a single block is chosen, then the weight is not included in the bit calculation. Using a Titan X GPU, we train our variational hashing model from Section 4 using the Adam optimizer [13], and tune the learning rate from the set {0.005, **0.001**, 0.0005} and the batch size from the set {100, 200, **400**, 800} (best values in bold). As noise injection has been found beneficial to reduce over-fitting in variational neural models [26], we add Gaussian noise to the ratings during training; we initially set the variance of the Gaussian to 1 and reduce by a factor of $1 - 10^{-4}$ in every training iteration.

5.5 Effectiveness results

Table 1 reports the effectiveness results measured with NDCG and MRR, where the highest NDCG and MRR per column among the hashing-based baselines is shown in **bold**. Results statistically significantly better than the other Hamming distance baselines per column, using a paired two tailed t-test at the 0.05 level and Bonferroni correction, are indicated by an asterisk *. The Amazon results for CCCF are not included because the released implementation requires >128GB of RAM on this dataset due to the larger amount of items and users.

There are 4 main findings in Table 1: (1) Hash codes optimized for the projected Hamming distance (VH_{PHD}) outperform all hashing baselines at all times. (2) The gains of VH_{PHD} are larger for MRR than for NDCG, which means that the bit-level importance coding impacts the very top of the ranking list (i.e., the recommendations that matter the most). (3) The best hashing baselines (CCCF, DCF, and VH) have overall similar scores, which indicates a potential ceiling in effectiveness with standard Hamming distance on hash codes. (4) MF with real-valued vectors (i.e., no hash codes) using the inner product outperforms all the hashing approaches, which is to be expected as the representational power of 32/64 floating point numbers is notably higher than that of 32/64 bits. However, VH_{PHD} bridges a large part of the effectiveness gap between the hashing baselines and MF, such that the NDCG differences between VH_{PHD} and MF in 9 out of 16 cases are below 0.01, while the MRR differences in 4 out of 8 cases are close to 0.01.

5.5.1 Impact of user difficulty. Given MF as the best performing method, we consider each user’s MF performance to be an indicator of difficulty for modeling that particular user. To see how this type of user difficulty impacts recommendation performance, we sort all users (per dataset) increasingly according to their 64-dimensional MF NDCG@10 scores (x axis), and plot the average NDCG@10

Table 1: NDCG@k and MRR scores. * marks statistically significant gains over the other Hamming distance baselines per column using Bonferroni correction. $\Delta\%$ shows the gain of VH_{PHD} over the best hashing-based baseline per column.

32 bit/dim.	Yelp			ML-1M			ML-10M			Amazon		
	NDCG@5	NDCG@10	MRR	NDCG@5	NDCG@10	MRR	NDCG@5	NDCG@10	MRR	NDCG@5	NDCG@10	MRR
Hamming distance												
CCCF	.7286	.8000	.6250	.6867	.7110	.6493	.5491	.5987	.5683	-	-	-
DCF	.7412	.8095	.6368	.6791	.7092	.6382	.5645	.6120	.5843	.8256	.8714	.7759
MFmean	.6912	.7712	.5815	.5631	.5950	.5053	.4111	.4688	.4271	.7899	.8452	.7342
MFmedian	.6935	.7734	.5769	.5631	.5952	.5085	.4082	.4665	.4225	.7899	.8452	.7343
VH	.7467	.8132	.6473	.6851	.7123	.6419	.5669	.6157	.5815	.8254	.8712	.7758
Projected Hamming dissimilarity												
VH_{PHD}	.8036*	.8547*	.7406*	.7135*	.7360*	.6940*	.5939*	.6358*	.6235*	.8479*	.8877*	.8062*
$\Delta\%$	+7.6%	+5.1%	+14.4%	+3.9%	+3.3%	+6.9%	+4.8%	+3.3%	+6.7%	+2.7%	+1.9%	+3.9%
Inner product												
MF	.8071*	.8573*	.7513*	.7352*	.7502*	.7370*	.6029*	.6427*	.6385*	.8586*	.8954*	.8262*
64 bit/dim.												
Hamming distance												
CCCF	.7371	.8060	.6329	.7016	.7259	.6716	.5645	.6134	.5837	-	-	-
DCF	.7497	.8155	.6574	.7049	.7285	.6766	.5865	.6316	.6088	.8299	.8747	.7825
MFmean	.6912	.7712	.5810	.5666	.5981	.5172	.4104	.4675	.4257	.7902	.8458	.7340
MFmedian	.6954	.7752	.5780	.5649	.5966	.5105	.4113	.4679	.4270	.7902	.8457	.7334
VH	.7537	.8185	.6561	.7103	.7338	.6759	.5860	.6328	.6013	.8300	.8746	.7828
Projected Hamming dissimilarity												
VH_{PHD}	.8075*	.8577*	.7540*	.7267*	.7459*	.7136*	.6034*	.6427*	.6373*	.8521*	.8908*	.8147*
$\Delta\%$	+7.1%	+4.8%	+14.7%	+2.3%	+1.7%	+5.5%	+2.9%	+1.6%	+4.7%	+2.7%	+1.8%	+4.1%
Inner product												
MF	.8096*	.8591*	.7573*	.7424*	.7552*	.7439*	.6188*	.6562*	.6594*	.8586*	.8954*	.8259*
Dataset properties	22,087 users 14,873 items 602,517 ratings 0.18% density			6,040 users 3,260 items 998,539 ratings 5.07% density			69,878 users 9,708 items 9,995,471 ratings 1.47% density			158,650 users 128,939 items 4,701,968 ratings 0.02% density		

score per user smoothed by averaging the 500 nearest users (y axis). We do this for the three best Hamming distance baselines (CCCF, DCF, and VH), VH_{PHD} , and MF, which can be seen in Figure 1.

We observe that VH_{PHD} outperforms all Hamming distance baselines, showing that the projected Hamming dissimilarity is robust across users. Note that, for the 20,000 users with the lowest MF NDCG@10 on ML-10M, all hashing-based methods outperform MF, highlighting that MF is not always consistently better than hashing-based alternatives (despite allowing for much richer (real-valued versus binary) data representations). In addition, for Yelp and Amazon, VH_{PHD} obtains near-identical performance as MF for a majority of the users. Interestingly, on Amazon the hashing-based approaches generally perform worse than MF on the [80000, 120000] user interval. We argue that this observation is due to those users not being expressed well by the limited representational power of hash codes using the Hamming distance, compared to real-valued vectors using the inner product, but the projected Hamming dissimilarity reduces a large part of this gap.

5.5.2 Impact of the number of items per user. We investigate the impact of a user’s activity as defined by their number of rated items

and plot the NDCG@10 scores per user for 64 bit hash codes and 64-dimensional vectors for MF (See Figure 2).

Generally for all methods, we observe higher performance for users with few rated items, and the performance drops as the number of rated items increases. On all datasets except Yelp, the hashing-based baselines perform similar to the real-valued MF initially, but the performance gap occurs as the number of rated items increases, especially for the users with the highest number of rated items for ML-10M and Amazon. The hash codes may perform worse on users with a high number of rated items due to their limited representational power (compared to a real-valued vector), however, using the projected Hamming dissimilarity in our VH_{PHD} enables to reduce this gap significantly. In fact, our VH_{PHD} performs almost identically to MF on Yelp and most of the users on ML-10M and Amazon, except those with the highest number of rated items.

5.5.3 Impact of the average item popularity per user. We now consider how a user’s average item popularity impacts performance. We denote an item’s popularity as the fraction of users who have rated the item, such that a user’s average item popularity can vary between 0.0 to 1.0, where 1.0 corresponds to only having rated items all other users have rated as well. We order users by their




Figure 1: Users are ordered by NDCG@10 for MF and the user-level performances are plotted.




Figure 2: Users are ordered by their number of rated items and the user-level NDCG@10 are plotted.




Figure 3: Users are ordered by their average item popularity and the user-level NDCG@10 are plotted.

average item popularity and plot the NDCG@10 scores per user for 64 bit hash codes and 64-dimensional vectors for MF (See Figure 3).

Generally for all methods, users with a high average item popularity obtain the highest performance scores, whereas users with a lower average item popularity tend to be harder to model. This can be explained by popular items appearing often during training, thus making it easier to learn high quality item representations (both binary and real-valued) matching the properties of those items. Interestingly, on Amazon and partly Yelp, the performance increases for users with the lowest average item popularity. We argue this is primarily due to the high sparsity of those datasets, meaning that some items are rated by very few users, thus making it possible to fit to a small set of user types. Similarly to the analysis on the number of rated items in Section 5.5.2, we overall observe a similar trend that our VH_{PHD} significantly reduces the gap between the existing hashing-based methods and the real-valued MF.

5.5.4 Stochastic or deterministic hash codes. We investigate the effect of the sampling strategy for hash codes (see Eq. 12) during training and evaluation. The sampling can either be deterministic ($\mu^{(j)} = 0.5$) or stochastic ($\mu^{(j)}$ is sampled uniformly at random

from $[0, 1]$), and does not have to be the same for training and evaluation. Figure 4 shows the performance for the four configurations of stochastic sampling or deterministic output across all datasets. We observe that stochastic training with deterministic evaluation consistently performs the best, while deterministic training and deterministic evaluation perform second best. As expected, stochastic sampling at evaluation performs significantly worse than the deterministic option (even more so when trained deterministically), as every item has a small probability of being sampled such that it has a small distance to a user, even though it has a low rating (and vice versa for highly rated items).

5.6 Efficiency results

We measure the efficiency of the projected Hamming dissimilarity in terms of convergence rate (when integrated into the variational hashing model) and runtime overhead compared to the Hamming distance.

5.6.1 Convergence rate. Figure 5 shows the convergence rate for the variational hashing model using either the Hamming distance or the projected Hamming dissimilarity for 64 bit hash codes. We




Figure 4: NDCG@10 of VHPHD when varying whether 64 bit hash codes are sampled stochastically or deterministically.

see that training with the projected Hamming dissimilarity significantly improves the convergence rate compared to the model with the Hamming distance. The time to run a single batch is the same for both the Hamming distance and the projected Hamming dissimilarity (approximately 0.007 seconds for a batch size of 400), from which we conclude that using and optimizing for the projected Hamming dissimilarity not only improves NDCG and MRR, but also drastically reduces training time. We argue that the masking done in the projected Hamming dissimilarity makes the hash codes easier to learn: During training, any update in the item hash code that makes a bit flip will change the distance to all users. However, for the projected Hamming dissimilarity, an update to an item only influences the projected Hamming dissimilarity to a user if the user’s corresponding bit is 1 (as opposed to -1). Thus, the Hamming distance has a global reach for each bit, compared to a more localised reach for the projected Hamming dissimilarity, which effectively makes the hash codes easier to learn.

5.6.2 Runtime analysis. The projected Hamming dissimilarity has the same number of Boolean operations as the Hamming distance (see Eq. 1 and 6), due to exploiting that the negation of the item hash codes in the projected Hamming dissimilarity can be precomputed and stored instead of the original item hash codes (i.e., requiring no additional storage). We now verify the actual runtime of both the Hamming distance and the projected Hamming dissimilarity, when considering a fixed set of hash codes. Note that this is a comparison of generated hash codes, thus all approaches using the Hamming distance (or projected Hamming dissimilarity) would have the same runtime in the following experiment. We implement both the Hamming distance and projected Hamming dissimilarity efficiently in C on a machine with a 64 bit instruction set. A test environment was made with 100M randomized 64 bit hash codes, where we measure the time taken to compute 100M Hamming distances and projected Hamming dissimilarities (averaged over 1000 repeated runs). All experiments were run on a single thread¹⁰, with all hash codes loaded in RAM. The source code was compiled with the highest optimization level utilizing all optimization flags applicable to the hardware.

As reported in Table 2, the mean experiment time was 0.07401 seconds using both the Hamming distance the projected Hamming dissimilarity. Thus, the projected Hamming dissimilarity add no

Table 2: Runtime in seconds and runtime overhead compared to the Hamming distance for 100M computations.

	Runtime (s)	Runtime overhead
Hamming distance	0.07401	-
Projected Hamming dissimilarity	0.07401	+0.0%
Inner product	4.71414	+6269.6%

computational overhead, but still allows learning hash codes enabling significant effectiveness improvements as reported in Table 1. Finally, Table 2 also reports the runtime of computing the inner product of floating point vectors of length 64: the computation time is 4.71414 seconds, thus being significantly slower than the Hamming space operations.

6 CONCLUSION

We presented the projected Hamming dissimilarity, which allows bit-level binary importance weighting (i.e., disabling bits), to produce hash codes that accurately represent dissimilarity between data objects and allow for very efficient subsequent processing. Next, we proposed a variational hashing model for learning hash codes to be optimized for the projected Hamming dissimilarity, and experimentally evaluated it in collaborative filtering experiments. Compared to state-of-the-art hashing-based baselines, we obtained effectiveness improvements of up to +7% in NDCG and +14% in MRR, across 4 widely used datasets. These gains come at no additional cost in storage or recommendation time, as the projected Hamming distance has the same extremely fast computation time as the Hamming distance. Compared to the Hamming distance, we further find that model optimization using the projected Hamming dissimilarity significantly improves the convergence rate, thus speeding up model training.

In future work, we plan to investigate the projected Hamming dissimilarity, and possible adaptations of it, in symmetric retrieval settings consisting of item-item similarity search, as opposed to asymmetric user-item search explored in this work. One such example is document similarity search, which in the hashing setting is known as Semantic Hashing [22], where current work has focused on using the Hamming distance for measuring document similarities [2, 4, 6, 8, 9, 24].

¹⁰We used a Intel Core i9-9940X @ 3.30GHz and had 128GB RAM available.




Figure 5: Convergence rate optimizing projected Hamming dissimilarity or Hamming distance.

REFERENCES

- [1] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. 2013. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432* (2013).
- [2] Suthee Chaidaroon and Yi Fang. 2017. Variational Deep Semantic Hashing for Text Documents. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval*. 75–84.
- [3] Ting Chen, Martin Renqiang Min, and Yizhou Sun. 2018. Learning K-way D-dimensional Discrete Codes for Compact Embedding Representations. In *Proceedings of the International Conference on Machine Learning*. 853–862.
- [4] Wei Dong, Qinliang Su, Dinghan Shen, and Changyou Chen. 2019. Document Hashing with Mixture-Prior Generative Models. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 5226–5235.
- [5] Bin Fan, Qingqun Kong, Xiaotong Yuan, Zhiheng Wang, and Chunhong Pan. 2013. Learning weighted Hamming distance for binary descriptors. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2395–2399.
- [6] Casper Hansen, Christian Hansen, Jakob Grue Simonsen, Stephen Alstrup, and Christina Lioma. 2019. Unsupervised Neural Generative Semantic Hashing. In *Proceedings of International ACM SIGIR Conference on Research and Development in Information Retrieval*. 735–744.
- [7] Casper Hansen, Christian Hansen, Jakob Grue Simonsen, Stephen Alstrup, and Christina Lioma. 2020. Content-aware Neural Hashing for Cold-start Recommendation. In *Proceedings of International ACM SIGIR Conference on Research and Development in Information Retrieval*. 971–980.
- [8] Casper Hansen, Christian Hansen, Jakob Grue Simonsen, Stephen Alstrup, and Christina Lioma. 2020. Unsupervised Semantic Hashing with Pairwise Reconstruction. In *SIGIR*. 2009–2012.
- [9] Christian Hansen, Casper Hansen, Jakob Grue Simonsen, Stephen Alstrup, and Christina Lioma. 2021. Unsupervised Multi-Index Semantic Hashing. In *Proceedings of The Web Conference 2021*. In print.
- [10] Ruining He and Julian McAuley. 2016. Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering. In *Proceedings of the International Conference on World Wide Web*. 507–517.
- [11] Kalervo Järvelin and Jaana Kekäläinen. 2000. IR evaluation methods for retrieving highly relevant documents. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval*. 41–48.
- [12] Tianxu Ji, Xianglong Liu, Cheng Deng, Lei Huang, and Bo Lang. 2014. Query-adaptive hash code ranking for fast nearest neighbor search. In *Proceedings of the 22nd ACM international conference on Multimedia*. 1005–1008.
- [13] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations*.
- [14] Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix Factorization Techniques for Recommender Systems. *Computer* 42, 8 (Aug. 2009), 30–37.
- [15] Defu Lian, Rui Liu, Yong Ge, Kai Zheng, Xing Xie, and Longbing Cao. 2017. Discrete Content-aware Matrix Factorization. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 325–334.
- [16] Dawen Liang, Rahul G Krishnan, Matthew D Hoffman, and Tony Jebara. 2018. Variational autoencoders for collaborative filtering. In *Proceedings of the 2018 World Wide Web Conference*. 689–698.
- [17] Chenghao Liu, Tao Lu, Xin Wang, Zhiyong Cheng, Jianling Sun, and Steven C.H. Hoi. 2019. Compositional Coding for Collaborative Filtering. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval*. 145–154.
- [18] Han Liu, Xiangnan He, Fuli Feng, Liqiang Nie, Rui Liu, and Hanwang Zhang. 2018. Discrete Factorization Machines for Fast Feature-based Recommendation. In *Proceedings of the International Joint Conference on Artificial Intelligence*. 3449–3455.
- [19] Xianglong Liu, Junfeng He, Cheng Deng, and Bo Lang. 2014. Collaborative hashing. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2139–2146.
- [20] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian personalized ranking from implicit feedback. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*. 452–461.
- [21] Steffen Rendle, Walid Krichene, Li Zhang, and John Anderson. 2020. Neural Collaborative Filtering vs. Matrix Factorization Revisited. In *Fourteenth ACM Conference on Recommender Systems (RecSys ’20)*. 240–248.
- [22] Ruslan Salakhutdinov and Geoffrey Hinton. 2009. Semantic hashing. *International Journal of Approximate Reasoning* 50, 7 (2009), 969–978.
- [23] Ying Shan, Jie Zhu, JC Mao, et al. 2018. Recurrent binary embedding for gpu-enabled exhaustive retrieval from billion-scale semantic vectors. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2170–2179.
- [24] Dinghan Shen, Qinliang Su, Paidamoyo Chapfuwa, Wenlin Wang, Guoyin Wang, Ricardo Henao, and Lawrence Carin. 2018. NASH: Toward End-to-End Neural Architecture for Generative Semantic Hashing. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*. 2041–2050.
- [25] Raphael Shu and Hideki Nakayama. 2018. Compressing Word Embeddings via Deep Compositional Code Learning. In *Proceedings of the International Conference on Learning Representations*.
- [26] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. 2015. Learning structured output representation using deep conditional generative models. In *Advances in Neural Information Processing Systems*. 3483–3491.
- [27] Qifan Wang, Dan Zhang, and Luo Si. 2013. Weighted hashing for fast large scale similarity search. In *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*. 1185–1188.
- [28] Stella X. Yu and Jianbo Shi. 2003. Multiclass Spectral Clustering. In *Proceedings of the IEEE International Conference on Computer Vision - Volume 2*. 313–319.
- [29] Dell Zhang, Jun Wang, Deng Cai, and Jinsong Lu. 2010. Self-taught Hashing for Fast Similarity Search. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval*. 18–25.
- [30] Hanwang Zhang, Fumin Shen, Wei Liu, Xiangnan He, Huanbo Luan, and Tat-Seng Chua. 2016. Discrete collaborative filtering. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval*. 325–334.
- [31] Jian Zhang and Yuxin Peng. 2018. Query-adaptive image retrieval by deep-weighted hashing. *IEEE Transactions on Multimedia* 20, 9 (2018), 2400–2414.
- [32] Lei Zhang, Yongdong Zhang, Jinhua Tang, Ke Lu, and Qi Tian. 2013. Binary code ranking with weighted hamming distance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 1586–1593.
- [33] Yan Zhang, Defu Lian, and Guowu Yang. 2017. Discrete personalized ranking for fast collaborative filtering from implicit feedback. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*.
- [34] Zhiwei Zhang, Qifan Wang, Lingyun Ruan, and Luo Si. 2014. Preference Preserving Hashing for Efficient Recommendation. In *International ACM SIGIR Conference on Research and Development in Information Retrieval*. 183–192.
- [35] Ke Zhou and Hongyuan Zha. 2012. Learning Binary Codes for Collaborative Filtering. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 498–506.

Bibliography

- [1] Gediminas Adomavicius and Alexander Tuzhilin. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE transactions on knowledge and data engineering*, 17(6):734–749, 2005.
- [2] Stephen Alstrup, Casper Hansen, Christian Hansen, Niklas Hjuler, Stephan Lorenzen, and Ninh Pham. Dabai: A data driven project for e-learning in denmark. In *ECEL17 - Proceedings of the 16th European Conference on e-Learning*, 2017.
- [3] Sanjeev Arora, Yingyu Liang, and Tengyu Ma. A simple but tough-to-beat baseline for sentence embeddings. In *ICLR*, 2017.
- [4] Isabelle Augenstein, Christina Lioma, Dongsheng Wang, Lucas Chaves Lima, Casper Hansen, Christian Hansen, and Jakob Grue Simonsen. MultiFC: A real-world multi-domain dataset for evidence-based fact checking of claims. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4685–4697, Hong Kong, China, 2019. Association for Computational Linguistics.
- [5] Alex Beutel, Paul Covington, Sagar Jain, Can Xu, Jia Li, Vince Gatto, and Ed H Chi. Latent cross: Making use of context in recurrent recommender systems. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, pages 46–54, 2018.
- [6] Roi Blanco and Christina Lioma. Graph-based term weighting for information retrieval. *Information retrieval*, 15(1):54–92, 2012.
- [7] Suthee Chaidaroon, Travis Ebesu, and Yi Fang. Deep semantic text hashing with weak supervision. In *SIGIR*, pages 1109–1112, 2018.
- [8] Suthee Chaidaroon and Yi Fang. Variational deep semantic hashing for text documents. In *SIGIR*, pages 75–84, 2017.

- [9] Hanjun Dai, Yichen Wang, Rakshit Trivedi, and Le Song. Deep coevolutionary network: Embedding user and item features for recommendation. *arXiv preprint arXiv:1609.03675*, 2016.
- [10] Scott Deerwester, Susan T Dumais, George W Furnas, Thomas K Landauer, and Richard Harshman. Indexing by latent semantic analysis. *Journal of the American society for information science*, 41(6):391–407, 1990.
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. 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 4171–4186, 2019.
- [12] Khoa D Doan and Chandan K Reddy. Efficient implicit unsupervised text hashing using adversarial autoencoder. In *Proceedings of The Web Conference 2020*, pages 684–694, 2020.
- [13] Wei Dong, Qinliang Su, Dinghan Shen, and Changyou Chen. Document hashing with mixture-prior generative models. In *EMNLP*, pages 5226–5235, 2019.
- [14] Bin Fan, Qingqun Kong, Xiaotong Yuan, Zhiheng Wang, and Chunhong Pan. Learning weighted hamming distance for binary descriptors. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 2395–2399. IEEE, 2013.
- [15] George W. Furnas, Thomas K. Landauer, Louis M. Gomez, and Susan T. Dumais. The vocabulary problem in human-system communication. *Communications of the ACM*, 30(11):964–971, 1987.
- [16] Aristides Gionis, Piotr Indyk, Rajeev Motwani, et al. Similarity search in high dimensions via hashing. In *VLDB*, pages 518–529, 1999.
- [17] Casper Hansen, Christian Hansen, Stephen Alstrup, and Christina Lioma. Smart city analytics: Ensemble-learned prediction of citizen home care. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, CIKM 2017, Singapore, November 06 - 10, 2017*, pages 2095–2098. ACM, 2017.
- [18] Casper Hansen, Christian Hansen, Stephen Alstrup, Jakob Grue Simonsen, and Christina Lioma. Contextually propagated term weights for document representation. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2019, Paris, France, July 21-25, 2019*, pages 897–900. ACM, 2019.

- [19] Casper Hansen, Christian Hansen, Stephen Alstrup, Jakob Grue Simonsen, and Christina Lioma. Neural check-worthiness ranking with weak supervision: Finding sentences for fact-checking. In *Companion of The 2019 World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13-17, 2019*, pages 994–1000. ACM, 2019.
- [20] Casper Hansen, Christian Hansen, Stephen Alstrup, Jakob Grue Simonsen, and Christina Lioma. Unsupervised multi-index semantic hashing. In *The 2021 World Wide Web Conference, WWW 2021, Ljubljana, Slovenia, April 19-23, 2021*, page in press. ACM, 2021.
- [21] Casper Hansen, Christian Hansen, Lucas Maystre, Rishabh Mehrotra, Brian Brost, Federico Tomasi, and Mounia Lalmas. Contextual and sequential user embeddings for large-scale music recommendation. In *RecSys 2020: Fourteenth ACM Conference on Recommender Systems, Virtual Event, Brazil, September 22-26, 2020*, pages 53–62. ACM, 2020.
- [22] Casper Hansen, Christian Hansen, Jakob Grue Simonsen, Stephen Alstrup, and Christina Lioma. Unsupervised neural generative semantic hashing. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2019, Paris, France, July 21-25, 2019*, pages 735–744. ACM, 2019.
- [23] Casper Hansen, Christian Hansen, Jakob Grue Simonsen, Stephen Alstrup, and Christina Lioma. Content-aware neural hashing for cold-start recommendation. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, July 25-30, 2020*, pages 971–980. ACM, 2020.
- [24] Casper Hansen, Christian Hansen, Jakob Grue Simonsen, Stephen Alstrup, and Christina Lioma. Unsupervised semantic hashing with pairwise reconstruction. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, July 25-30, 2020*, pages 2009–2012. ACM, 2020.
- [25] Casper Hansen, Christian Hansen, Jakob Grue Simonsen, and Christina Lioma. The copenhagen team participation in the check-worthiness task of the competition of automatic identification and verification of claims in political debates of the clef-2018 fact checking lab. In *CLEF-2018 CheckThat! Lab*, 2018.
- [26] Casper Hansen, Christian Hansen, Jakob Grue Simonsen, and Christina Lioma. Neural weakly supervised fact check-worthiness detection with contrastive sampling-based ranking loss. In *CLEF-2019 CheckThat! Lab*, 2019.

- [27] Casper Hansen, Christian Hansen, Jakob Grue Simonsen, and Christina Lioma. Fact check-worthiness detection with contrastive ranking. In *Experimental IR Meets Multilinguality, Multimodality, and Interaction - 11th International Conference of the CLEF Association, CLEF 2020, Thessaloniki, Greece, September 22-25, 2020, Proceedings*, volume 12260 of *Lecture Notes in Computer Science*, pages 124–130. Springer, 2020.
- [28] Casper Hansen, Christian Hansen, Jakob Grue Simonsen, and Christina Lioma. Projected hamming dissimilarity for bit-level importance coding in collaborative filtering. In *The 2021 World Wide Web Conference, WWW 2021, Ljubljana, Slovenia, April 19-23, 2021*, page in press. ACM, 2021.
- [29] Christian Hansen, Casper Hansen, Stephen Alstrup, and Christina Lioma. Modelling end-of-session actions in educational systems. In *Proceedings of the 12th International Conference on Educational Data Mining, EDM 2019, Montréal, Canada, July 2-5, 2019*. International Educational Data Mining Society (IEDMS), 2019.
- [30] Christian Hansen, Casper Hansen, Stephen Alstrup, Jakob Grue Simonsen, and Christina Lioma. Neural speed reading with structural-jump-lstm. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*, 2019.
- [31] Christian Hansen, Casper Hansen, Niklas Hjuler, Stephen Alstrup, and Christina Lioma. Sequence modelling for analysing student interaction with educational systems. In *Proceedings of the 10th International Conference on Educational Data Mining, EDM 2017, Wuhan, Hubei, China, June 25-28, 2017*. International Educational Data Mining Society (IEDMS), 2017.
- [32] Christian Hansen, Casper Hansen, Jakob Grue Simonsen, Stephen Alstrup, and Christina Lioma. Modelling sequential music track skips using a multi-rnn approach. In *WSDM Cup*. ACM, 2019.
- [33] Christian Hansen, Casper Hansen, Jakob Grue Simonsen, Birger Larsen, Stephen Alstrup, and Christina Lioma. Factuality checking in news headlines with eye tracking. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, July 25-30, 2020*, pages 2013–2016. ACM, 2020.
- [34] Christian Hansen, Rishabh Mehrotra, Casper Hansen, Brian Brost, Lucas Maystre, and Mounia Lalmas. Shifting consumption towards diverse content on music streaming platforms. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining, WSDM '21*, page 238–246. Association for Computing Machinery, 2021.
- [35] Zellig S Harris. Distributional structure. *Word*, 10(2-3):146–162, 1954.

- [36] Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: towards removing the curse of dimensionality. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 604–613, 1998.
- [37] Espen Jimenez-Solem, Tonny S Petersen, Casper Hansen, Christian Hansen, Christina Lioma, Christian Igel, Wouter Boomsma, Oswin Krause, Stephan Lorenzen, Raghavendra Selvan, et al. Developing and validating covid-19 adverse outcome risk prediction models from a bi-national european cohort of 5594 patients. *Scientific reports*, 11(1):1–12, 2021.
- [38] Alexandros Karatzoglou, Alex Smola, and Markus Weimer. Collaborative filtering on a budget. In *International Conference on Artificial Intelligence and Statistics*, pages 389–396, 2010.
- [39] Guy Lev, Benjamin Klein, and Lior Wolf. In defense of word embedding for generic text representation. In *International conference on applications of natural language to information systems*, pages 35–50. Springer, 2015.
- [40] Defu Lian, Rui Liu, Yong Ge, Kai Zheng, Xing Xie, and Longbing Cao. Discrete content-aware matrix factorization. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 325–334, 2017.
- [41] Defu Lian, Rui Liu, Yong Ge, Kai Zheng, Xing Xie, and Longbing Cao. Discrete content-aware matrix factorization. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 325–334, 2017.
- [42] Lucas Chaves Lima, Casper Hansen, Christian Hansen, Dongsheng Wang, Maria Maistro, Birger Larsen, Jakob Grue Simonsen, and Christina Lioma. Denmark’s participation in the search engine trec covid-19 challenge: Lessons learned about searching for precise biomedical scientific information on covid-19. In *TREC COVID-19 Challenge*, 2021.
- [43] Christina Lioma and Roi Blanco. Part of speech based term weighting for information retrieval. In *European Conference on Information Retrieval*, pages 412–423. Springer, 2009.
- [44] Chenghao Liu, Tao Lu, Xin Wang, Zhiyong Cheng, Jianling Sun, and Steven CH Hoi. Compositional coding for collaborative filtering. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 145–154, 2019.
- [45] Rastin Matin, Casper Hansen, Christian Hansen, and Pia Mølgaard. Predicting distresses using deep learning of text segments in annual reports. *Expert Systems with Applications*, 132:199–208, 2019.

- [46] Tomás Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. In *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*, 2013.
- [47] Mohammad Norouzi, Ali Punjani, and David J Fleet. Fast search in hamming space with multi-index hashing. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3108–3115. IEEE, 2012.
- [48] Mohammad Norouzi, Ali Punjani, and David J Fleet. Fast exact search in hamming space with multi-index hashing. *IEEE transactions on pattern analysis and machine intelligence*, 36(6):1107–1119, 2013.
- [49] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, 2014.
- [50] Jay M Ponte and W Bruce Croft. A language modeling approach to information retrieval. In *Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval*, pages 275–281, 1998.
- [51] Stephen E Robertson, Steve Walker, Susan Jones, Micheline M Hancock-Beaulieu, Mike Gatford, et al. Okapi at trec-3. *Nist Special Publication Sp*, 109:109, 1995.
- [52] Ruslan Salakhutdinov and Geoffrey Hinton. Semantic hashing. *International Journal of Approximate Reasoning*, 50(7):969–978, 2009.
- [53] Gerard Salton and Christopher Buckley. Term-weighting approaches in automatic text retrieval. *Information processing & management*, 24(5):513–523, 1988.
- [54] Hinrich Schütze, Christopher D Manning, and Prabhakar Raghavan. *Introduction to information retrieval*, volume 39. Cambridge University Press Cambridge, 2008.
- [55] Ying Shan, Jian Jiao, Jie Zhu, and JC Mao. Recurrent binary embedding for gpu-enabled exhaustive retrieval from billion-scale semantic vectors. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2170–2179, 2018.
- [56] Dinghan Shen, Qinliang Su, Paidamoyo Chapfuwa, Wenlin Wang, Guoyin Wang, Ricardo Henao, and Lawrence Carin. Nash: Toward end-to-end neural architecture for generative semantic hashing. In *ACL*, pages 2041–2050, 2018.

- [57] Yue Shi, Martha Larson, and Alan Hanjalic. Collaborative filtering beyond the user-item matrix: A survey of the state of the art and future challenges. *ACM Computing Surveys (CSUR)*, 47(1):1–45, 2014.
- [58] Anant Subramanian, Danish Pruthi, Harsh Jhamtani, Taylor Berg-Kirkpatrick, and Eduard Hovy. Spine: Sparse interpretable neural embeddings. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018.
- [59] Dongsheng Wang, Casper Hansen, Lucas Chaves Lima, Christian Hansen, Maria Maistro, Jakob Grue Simonsen, and Christina Lioma. Multi-head self-attention with role-guided masks. In *Proceedings of the 43rd European Conference on Information Retrieval Research*, pages 432–439, 2021.
- [60] Yair Weiss, Antonio Torralba, and Rob Fergus. Spectral hashing. In *Proceedings of the 21st International Conference on Neural Information Processing Systems, NIPS’08*, page 1753–1760. Curran Associates Inc., 2008.
- [61] Chao-Yuan Wu, Amr Ahmed, Alex Beutel, Alexander J Smola, and How Jing. Recurrent recommender networks. In *Proceedings of the tenth ACM international conference on web search and data mining*, pages 495–503, 2017.
- [62] Dell Zhang, Jun Wang, Deng Cai, and Jinsong Lu. Laplacian co-hashing of terms and documents. In *European Conference on Information Retrieval*, pages 577–580. Springer, 2010.
- [63] Dell Zhang, Jun Wang, Deng Cai, and Jinsong Lu. Self-taught hashing for fast similarity search. In *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval*, pages 18–25, 2010.
- [64] Hanwang Zhang, Fumin Shen, Wei Liu, Xiangnan He, Huanbo Luan, and Tat-Seng Chua. Discrete collaborative filtering. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, pages 325–334, 2016.
- [65] Yan Zhang, Hongzhi Yin, Zi Huang, Xingzhong Du, Guowu Yang, and Defu Lian. Discrete deep learning for fast content-aware recommendation. In *ACM International Conference on Web Search and Data Mining*, pages 717–726, 2018.
- [66] Yan Zhang, Hongzhi Yin, Zi Huang, Xingzhong Du, Guowu Yang, and Defu Lian. Discrete deep learning for fast content-aware recommendation. In *Proceedings of the eleventh ACM international conference on web search and data mining*, pages 717–726, 2018.

- [67] Zhiwei Zhang, Qifan Wang, Lingyun Ruan, and Luo Si. Preference preserving hashing for efficient recommendation. In *International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 183–192, 2014.
- [68] Le Zhao and Jamie Callan. Term necessity prediction. In *Proceedings of the 19th ACM international conference on Information and knowledge management*, pages 259–268, 2010.
- [69] V. Zhelezniak, A. Savkov, April Shen, Francesco Moramarco, Jack Flann, and N. Hammerla. Don’t settle for average, go for the max: Fuzzy sets and max-pooled word vectors. In *ICLR*, 2019.
- [70] Ke Zhou and Hongyuan Zha. Learning binary codes for collaborative filtering. In *ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 498–506, 2012.