

# A Deep Metric Learning Approach to Account Linking

Aleem Khan, Elizabeth Fleming, Noah Schofield, Marcus Bishop, and Nicholas Andrews

Human Language Technology Center of Excellence

Johns Hopkins University

{akhan141, eflemin9, nschofi1, marcus.bishop, noa}@jhu.edu

## Abstract

We consider the task of linking social media accounts that belong to the same author in an automated fashion on the basis of the content and metadata of their corresponding document streams. We focus on learning an embedding that maps variable-sized samples of user activity—ranging from single posts to entire months of activity—to a vector space, where samples by the same author map to nearby points. The approach does not require human-annotated data for training purposes, which allows us to leverage large amounts of social media content. The proposed model outperforms several competitive baselines under a novel evaluation framework modeled after established recognition benchmarks in other domains. Our method achieves high linking accuracy, even with small samples from accounts not seen at training time, a prerequisite for practical applications of the proposed linking framework.

## 1 Introduction

The scale and anonymity of social media pose systematic challenges for manual moderation efforts (Pennycook et al., 2020; Broniatowski et al., 2018). These challenges have motivated the development of automated methods to identify abusive content, such as Davidson et al. (2017), which considers automatically classifying hate speech, or Alvari et al. (2019), which deals with detecting violent extremism, both in the Twitter domain.

However, automatic moderation remains a difficult problem. Indeed, existing methods based on hand-constructed resources such as keyword lists may fail to adapt to novel trends (Corbett-Davies and Goel, 2018) whereas automatic methods based on statistics of large corpora may exhibit harmful biases (Caliskan et al., 2017). Additionally, individual posts may fail to contain sufficient information to reliably identify them as harmful.

This work considers *account-level* moderation. Specifically, we consider the problem of determining whether two *document streams* share the same author, based on *samples* from those streams rather than individual documents. This capability has numerous applications, such as detecting users attempting to circumvent account bans, identifying sockpuppet accounts, and detecting coordinated disinformation campaigns involving multiple authors controlling multiple accounts.

As a motivating application, we consider the enforcement of account bans on anonymous platforms, such as Reddit. Given a new account, the problem is to automatically identify whether it matches any previously banned account, which amounts to making binary decisions about whether pairs of accounts share the same author. Variations of this problem have been studied before. For example, Schwartz et al. (2013) learn a classifier to determine whether the author of a Twitter comment belongs to a small, closed set of authors. In contrast, we are interested in an open-world setting, requiring binary decisions about arbitrary pairs of accounts. This introduces a number of challenges.

First, any individual comment may be too short to serve as the basis for linking accounts. Figure 1a illustrates this empirically using a variation of our model, where embeddings of individual comments from the same account fail to coalesce, making it difficult to assert that an account has the same author as another account. See §4 for further experimental details. Therefore, we focus on *aggregating* information across contiguous sequences of documents. Figure 1b illustrates the impact of aggregation using our full model, where aggregations of contiguous sequences of documents from the same account exhibit an approximate convergence behavior as the number of documents aggregated increases. In fact, the motivating application above *requires* linking accounts on the basis of samples of widely varying sizes. Indeed, banned accounts

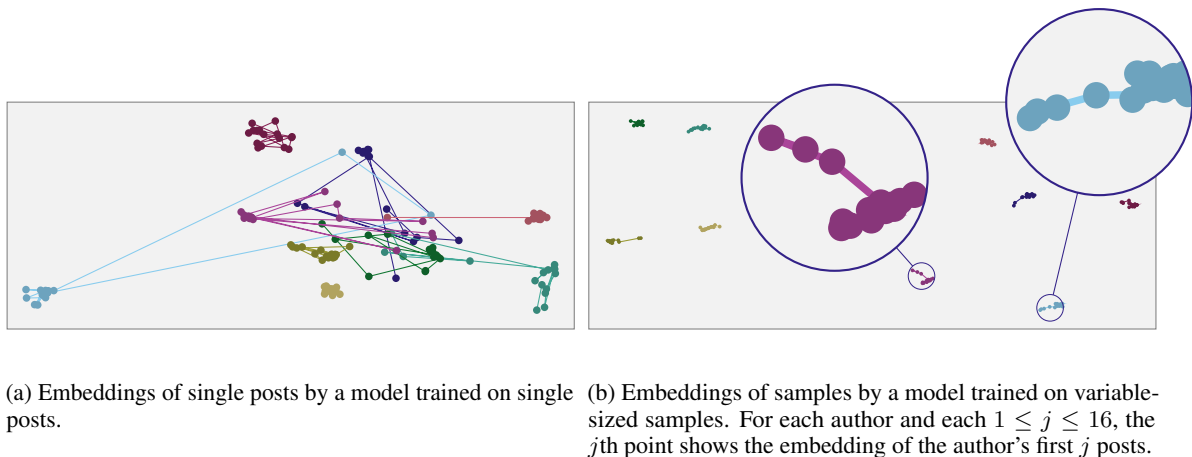


Figure 1: UMAP projections (McInnes et al., 2020) of embeddings of document streams. Each image shows the same 16 Reddit posts by 9 randomly chosen authors, connected chronologically. Posts by each author are shown in the same color in both figures.

typically have many documents, all of which we would like to consider, while new accounts generally have few documents, which nevertheless must be linked to banned accounts as quickly as possible to mitigate abusive behavior.

The second challenge is that of spurious associations. For example, over a short period of time, an author may discuss only a single, narrow topic. While a naïve model based on word statistics alone might be sufficient to link such an account to another, this approach would fail to generalize over longer periods of time due to topic drift. Both our training procedure and evaluation framework have been designed to ensure that our model learns the appropriate invariances to identify same-authorship, rather than being correct for the wrong reasons (McCoy et al., 2019). Namely, samples from an account are drawn from different time periods in each training iteration (see §2.3), while the evaluation data consists of posts by accounts not seen at training time and is future to all the training data (see §4.2).

Finally, the numbers of banned and new accounts may be quite large, requiring a still larger number of pairwise comparisons. For this reason, our proposed approach to account linking consists of embedding variable-sized samples from document streams into a metric space whereby samples likely to have been composed by the same author map to nearby points. Under this embedding, comparisons between document streams amount to pairwise distance calculations, so our approach is highly scalable and amenable to various optimizations, such as approximate nearest neighbor methods.

Our primary contributions are the following:

- We provide a simple but effective data augmentation strategy which enables embedding variable-sized samples. In addition, we successfully train such an embedding on a large-scale dataset consisting of more than 300 million comments from 1 million distinct accounts using scalable losses.
- We propose a novel framework to assess account linking performance focused on challenging conditions and minimizing the impact of incidental authorship features, such as topic. In particular, we propose benchmark datasets as well as verification metrics tailored to our application.

Our code, data splits, and scripts to reproduce our experiments are available at <http://github.com/noa/naacl2021>.

## 2 Learning embeddings of document streams

We treat a document stream as a sequence of timestamped actions  $a_1, a_2, \dots, a_L$  where each  $a_i$  is a structure containing the data comprising an action. The possible contents of  $a_i$  are specific to the document stream, but include at least a timestamp  $t_i$  such that  $t_1 < t_2 < \dots < t_L$ .

In this work we focus on textual content published on social media platforms, although the approach would easily extend to allow  $a_1, a_2, \dots, a_L$  to contain other modalities such as images or video, which would be handled similarly. In addition, we

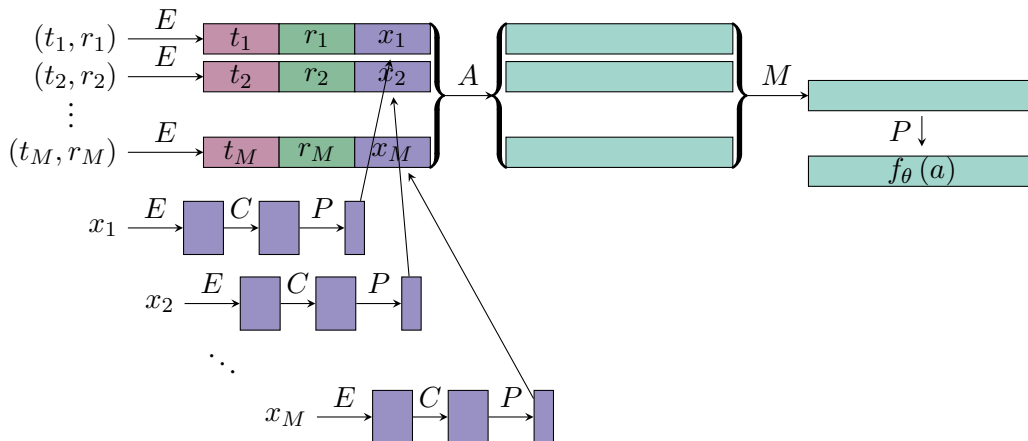


Figure 2: Illustration of the proposed model architecture. Each action  $a_i$  consists of text content  $x_i \in \mathbb{Z}^\ell$ , a subreddit feature  $r_i \in \mathbb{Z}$ , and is published at time  $t_i \in \{0, 1, \dots, 23\}$ . These elements are combined as shown using various embedding lookups  $E$ , one-dimensional convolutions  $C$ , and linear projections  $P$ , along with an attention mechanism  $A$  and a max pooling layer  $M$ .

also avail of certain categorical features contained in  $a_1, a_2, \dots, a_L$  such as hashtags or the subreddit to which a comment was posted.

A *sample* from a document stream is a contiguous subsequence of its actions. We introduce an embedding  $f_\theta$  in §2.1 mapping a *variable-sized* sample to a point in a vector space, such that the Euclidean distance between the embeddings of two samples quantifies the likelihood that they belong to the same author.

## 2.1 Architecture

We define an embedding  $f_\theta$  as follows. This embedding is illustrated in Figure 2. Consider a sample  $a = (a_1, a_2, \dots, a_M)$  where each action  $a_i$  consists of a subreddit feature  $r_i$ , a timestamp  $t_i$ , and text content  $x_i$ .<sup>1</sup> We encode  $t_i$  as the corresponding hour of the day and  $r_i$  by lookup in the list of 2048 most common subreddits, resulting in  $t_i \in \{0, 1, \dots, 23\}$  and  $r_i \in \{0, 1, \dots, 2048\}$ , where  $r_i = 2048$  when the subreddit is not among the top 2048. We encode the text feature using the SentencePiece unigram subword model (Kudo, 2018), resulting in  $x_i \in \mathbb{Z}^\ell$ , where the parameter  $\ell$  is defined in §2.2. Note that the chosen vocabulary size impacts the amount of content that can be encoded with  $\ell$  integers. Further details on the choice of text encoding are provided in Appendix C.

We replace each token of  $x_i$  with a correspond-

<sup>1</sup>In addition to  $r_i, t_i, x_i$  platforms like Reddit contain further metadata, such as the thread title and the submission and parent comments, all of which might help distinguish users because users read them and chose to respond. Incorporating these features would be interesting to explore in future work.

ing learned embedding in  $\mathbb{R}^N$  for all  $1 \leq i \leq M$ , resulting in  $M$  matrices in  $\mathbb{R}^{\ell \times N}$ . We apply one-dimensional convolutions of widths 2, 3, and 4 along the first axis of each, concatenate the convolved matrices along their second axes, max-pool along the first axis, and concatenate the results with learned embeddings of the corresponding subreddit features and one-hot encodings of the corresponding time features, resulting in  $M$  vectors, which we aggregate using dot-product attention. The resulting sequence of vectors is projected to a single vector through max-pooling followed by two fully-connected layers with bias, resulting in the encoding  $f_\theta(a) \in \mathbb{R}^D$  of the sample  $a$ .

## 2.2 Text Sampling

The variance in the lengths of documents poses computational challenges when aggregating large samples. Therefore we resort to truncating each document to a fixed number  $\ell$  of tokens, padding any documents containing fewer than  $\ell$  tokens. We take  $\ell = 32$  after observing that Reddit posts have an average length of approximately 43 tokens (see Appendix D).

We also experimented with a more complicated text sampling strategy, namely sampling contiguous segments of  $\ell$  tokens from each post uniformly at random during training. While this approach leverages all available textual information by affording slightly different samples of each post in each iteration of training, we found it to yield similar results and also complicates the comparison to our primary baseline model, which uses the prefix

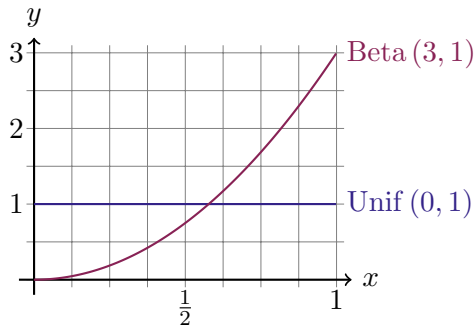


Figure 3: In contrast with choosing sample sizes uniformly at random, sampling according to a left skewed Beta distribution results in more samples of sizes closer to the specified maximum.

method described above.

### 2.3 Sample selection

During training we randomly select document stream samples of sizes varying between  $R = 1$  and  $S = 16$ , which we regard as hyperparameters of the model. To select a sample from the stream  $a_1, a_2, \dots, a_L$  we first choose its length  $M = R + \lceil x(S - R) \rceil$  where  $x \sim \text{Beta}(3, 1)$  and take the sample  $a_i, a_{i+1}, \dots, a_{i+M-1}$  where  $1 \leq i \leq L - M + 1$  is chosen uniformly at random. Selecting  $M$  according to  $\text{Beta}(3, 1)$  provides an expected sample size closer to  $S$  than to  $R$ , a trade-off that allows the model to quickly learn features of a document stream by exposing it to larger samples most of the time, while still maintaining the flexibility to handle samples of varying sizes. Indeed, the latter is critical in the evaluation described in §4.3, which requires linking large samples to small samples. The density function of  $\text{Beta}(3, 1)$  is shown in Figure 3 together with that of the uniform distribution  $\text{Unif}(0, 1)$  for comparison. We explore the benefits of  $\text{Beta}(3, 1)$  and other related distributions in §4.5.

### 2.4 Scalable deep metric learning losses

Deep metric learning methods aim to embed observations into a low-dimensional space such that instances from the same class map to nearby points under a chosen metric, such as Euclidean distance. In our setting, we take the instances to be document stream samples and the classes to be the corresponding accounts, which serve as proxies for latent authorship. Therefore, training the mapping  $f_\theta$  defined in §2.1 using metric learning affords an embedding under which samples by the same author map to nearby points.

Recent work in deep metric learning has introduced a number of training objectives with state of the art performance on computer vision tasks (Kim et al., 2020; Wang et al., 2019). Unfortunately, many of these objectives scale linearly with the number  $K$  of classes considered due to a costly linear projection onto  $\mathbb{R}^K$ . Note that because account names in effect provide labels for the corresponding document streams, we may use raw social media content to fit our model directly, availing of a virtually unlimited source of data. We stipulate that the ability to exploit larger amounts of data may be more important than per-example efficiency, and therefore consider the classical *triplet loss* (Schroff et al., 2015) in our experiments, whose complexity does not depend on  $K$ . In particular, we use semi-hard negative mining with a fixed margin penalty. We also consider the *top-k* loss recently proposed by Lu et al. (2019), which optimizes precision-at- $k$  as follows. Given targets ranked by similarity to a query, *top-k* arranges for as many matches as possible to be among the top  $k$  ranked targets. It accomplishes this by penalizing only those targets that would need to move the smallest amount in order to maximize the number of matching targets among the top  $k$ . Like triplet loss, *top-k* also uses an additive margin penalty to separate classes. See Appendix A for further experimental details on both loss functions.

## 3 Related work

The separate but related problem of *closed-world* author attribution has received considerable attention. For example, the PAN 2019 challenge (Daelemans et al., 2019) employed a closed-world setting with a small number of authors that are the same at training and test time. That task also considered longer documents, obviating the need for aggregating evidence of authorship across multiple documents.

Generic text embedding methods such as the universal sentence encoder (Cer et al., 2018) and BERT (Devlin et al., 2019) are fit using auxiliary tasks, such as conditional language modeling. In the case of BERT, this is usually followed by supervised fine-tuning for a downstream task of interest. In this work, we are interested in learning representations that are immediately useful for our account linking task. However, because a large corpus of task-specific training data may be collected without human supervision, the benefits of generative pre-



training are diminished in our setting. Indeed, the parameters of the text encoding are learned from a random initialization in all our experiments. Our approach is further distinguished from generic embedding methods by featuring a *multi-document* embedding, mapping a sequence of documents to a single vector, where each document may consist of both text and metadata.

The most closely related prior work is the Invariant User Representation (IUR) proposed by Andrews and Bishop (2019), whose approach is broadly similar to ours, but only considers samples of a fixed size. Our approach may be viewed as a generalization of that work in support of the account linking task. In addition we use a simpler dot product attention mechanism and introduce the use of scalable metric learning losses in §2.4, which enable us to train our model on an order of magnitude more data than previously considered. We validate these improvements in §4.2 using the ranking task proposed by Andrews and Bishop (2019). We also adapt IUR to serve as a baseline in our primary linking task in §4.3.

We believe that our treatment of account linking as a pairwise recognition task between document stream samples and our proposed general-purpose evaluation protocols are both novel. However, in prior work, there have been several platform-specific approaches to account linking. For example Silvestri et al. (2015) explore a heuristic approach to linking accounts across social media platforms. Separately, on platforms with rich social network information, graph-matching methods have been explored (Fan, 2012). Our focus is on content-based account linking, which is more general than prior methods we are aware of. Some other related but distinct problems include detecting deceptive accounts (Van Der Walt and Eloff, 2018) and authorship classification of short messages (Ishihara, 2011).

## 4 Experiments

We conduct evaluations on the two primary tasks illustrated in Figure 4. First, our ranking evaluation described in §4.2 is motivated by information retrieval needs. Although ranking is not the focus of this paper, it provides an assessment of the quality of the learned embedding in terms of similarity judgements and facilitates comparison with the baseline model IUR. In addition, we use the ranking evaluation to monitor training using de-

Model	Features	Loss	MRR	R@4	R@8
Prop	TPS	Top- <i>k</i>	<b>0.637</b>	<b>0.709</b>	<b>0.765</b>
Prop	TPS	Triplet	0.634	0.702	0.762
Prop	TP	Top- <i>k</i>	0.450	<b>0.522</b>	<b>0.595</b>
Prop	TP	Triplet	<b>0.452</b>	0.520	0.591
Prop	T	Triplet	0.372	0.439	0.512
IUR	TPS	Arcface	0.520	0.590	0.650
IUR	T	Arcface	0.200	0.240	0.290

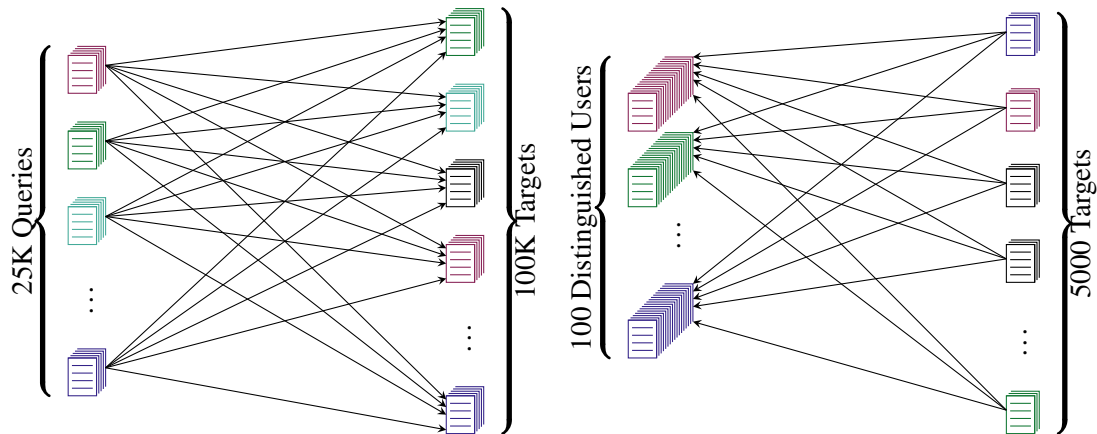
Table 1: Ranking results for the proposed model (Prop) and the baseline (IUR), both trained and evaluated using various combinations of text content (T), publication time (P), and subreddit (S). IUR results are reported from Andrews and Bishop (2019).

velopment data disjoint from the test data used for the final evaluation. Second, we introduce an account linking evaluation framework in §4.3 inspired by similar evaluations used for speaker recognition (Doddington et al., 2000; Van Leeuwen and Brümmer, 2007). Both evaluations involve setting up two sets of samples as described below, the *queries* and the *targets*. For each query, there is exactly one target drawn from the same document stream. Roughly speaking, both evaluations involve matching targets with their corresponding queries.

### 4.1 MUD: a Web-scale training dataset

Reddit is currently one of the most popular social media platforms, where anonymous users interact primarily by posting *comments* to discussion threads. Together with its text content, each comment is labeled by its publication time and the subreddit to which it was posted, a categorical feature roughly indicating its topic.

We construct a dataset consisting of 300 million Reddit posts from 1 million users published over an entire year to be used to train our proposed model. This *Million User Dataset* (MUD) consists of all posts by authors who published at least 100 and at most 1000 posts between July 2015 and June 2016, where the lower bound ensures a sufficiently long history from which to sample, and the upper bound is intended to reduce the impact of bot and spam accounts. We obtained the data by drawing from the existing Pushshift Reddit corpus (Baumgartner et al., 2020). Some further statistics of MUD are shown in Table 8.



(a) Ranking evaluation: for each query, the targets are ranked by similarity to the query. (b) Linking evaluation: for each target, the queries are determined to match each query or not.

Figure 4: Illustrations of our two primary evaluation frameworks. Document stream samples are shown as stacks of documents of sizes reflecting the corresponding sample sizes.

## 4.2 Ranking evaluation

As shown in Figure 4a, the ranking experiment consists of ranking the targets by similarity to each query. For compatibility, we mimic the experimental setup from Andrews and Bishop (2019), which proposes *separate* sets of queries and targets to be used for training and testing. We adopt the training split for validation and the testing split for evaluation, although we *train* our model on MUD (see §4.1). We select hyperparameters based on dev split performance (see Appendix A). The test split consists of samples, each of size exactly 16, although we train the proposed model using samples from MUD of *varying* sizes as described in §2.3. Note that the posts comprising MUD *precede* those of both IUR splits in publication time, ensuring that our training data is disjoint from IUR’s test data. Of the 111,396 authors contributing to the test split, 69,275 or 62% contribute to the IUR training split. In contrast, MUD has only 39,529 users in common with the test split, a significantly smaller overlap than IUR. In principle, the increase in novel users at test time puts the proposed model at a disadvantage because it places more importance on generalization to novel users.

We report recall-at- $k$  ( $R@k$ ) and mean reciprocal rank (MRR), calculated exactly as in Andrews and Bishop (2019). MRR is the expected value of the reciprocal of the position of the correct target in the ranked list.  $R@k$  is the probability that the unique target composed by the same author as a given query appears in the top  $k$  ranked results. We

limit ourselves to  $R@4$  and  $R@8$  as proxies for the “first page” of search results returned to a user issuing a query.

The results of this evaluation calculated with the test split are shown in Table 1. Note that the full version of the proposed model significantly outperforms the previously published state-of-the-art.<sup>2</sup> We conclude that although the angular margin loss used by Andrews and Bishop (2019) is considered state-of-the-art, the simpler triplet loss outperforms it, most likely because it admits the use of a considerably larger dataset. We remark that the models trained with top- $k$  performed only slightly better than those trained with triplet loss, an observation consistent with recent findings that when matching experimental conditions, the choice of ranking loss is less important than previously believed (Muscgrave et al., 2020).

In addition, Figure 5 shows the results of the evaluation performed after every hour of training.<sup>3</sup> Note that after only six hours of training the full model outperforms the baseline. Figure 5 also shows the learning curve for an ablation of our model that eliminates the subreddit feature. We observe that this ablated model performs almost as well as the full-featured baseline, which suggests that the proposed approach may be effective in domains where only text and timestamps are

<sup>2</sup>A paired sign test of the differences in ranking between IUR and the proposed model is significant at the  $p < 10^{-15}$  level.

<sup>3</sup>Although we generated Figure 5 *post-hoc* using test data, we did not use test data for model or hyperparameter selection.

available.

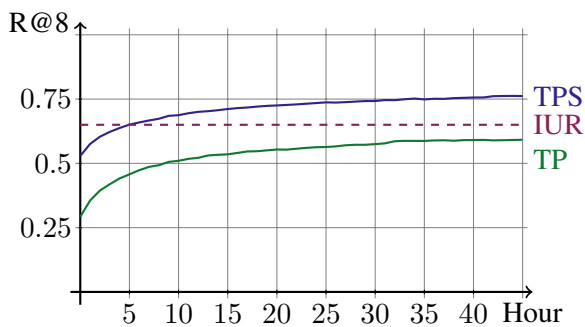


Figure 5: Test results after every hour of training for the proposed model using combinations of text (T), publication time (P), and subreddit (S) features as shown. The baseline result 0.65 shown as a dashed line is reported from [Andrews and Bishop \(2019\)](#) and corresponds with a model availing of all three features.

### 4.3 A new framework for account linking

While the ranking experiments in §4.2 were designed to measure the quality of the learned embedding, they do not directly measure task performance: moderation applications require *decisions* rather than *rankings*. To this end, we propose an account linking benchmark modeled after the problem of enforcing account bans, in which a fixed number of accounts are linked against novel accounts at test-time. Compared to the ranking experiments, the key difference is that we introduce a *distinguished* subset of authors from which we have accumulated a significant number of previously published documents to serve as queries. The procedure is illustrated in [Figure 4b](#).

Because the subreddit feature serves as a proxy for topic, restricting to a single subreddit results in a more challenging problem by increasing the likelihood that the comments considered deal with similar topics. To this end, we repeat the following procedure for each of the five most popular subreddits. Each result of the experiment reported in [Tables 2 and 9](#) is the average over the five subreddits of the corresponding results calculated using those subreddits individually.

Given a specified subreddit, we first randomly select 100 distinguished accounts, each publishing at least 100 posts to that subreddit in November 2016. The queries in the experiment consist of the 100 most recently published posts to the subreddit by each of the distinguished accounts in November 2016. In addition, the distinguished ac-

counts must have published at least 16 posts to the subreddit between December 2016 and May 2017 to serve as the corresponding targets, as described below.

Next we randomly select 4900 accounts distinct from the distinguished accounts, each publishing at least 16 posts to the subreddit between December 2016 and May 2017. The targets in the experiment consist of the 4 most recently published posts to the subreddit by each of the 5000 accounts.

**Performance metrics.** For every query and target, each model considered returns a score, with smaller scores associated with a higher likelihood that the query and the target have the same author. For example, the proposed model returns the distance between their embeddings under the model. A decision rule to predict an author match is obtained by thresholding this score with respect to a chosen *operating point*. In production settings, one adjusts the operating point to obtain acceptable rates of false positives and false negatives. In our running application of ban enforcement, these types of errors correspond respectively with mistakenly banning an innocent user and failing to ban a new account of a banned user. Because the severity of these types errors are different, we consider the *detection cost function*

$$C_{\text{det}} = \pi C^- P^- + (1 - \pi) C^+ P^+$$

proposed by [Van Leeuwen and Brümmer \(2007\)](#), where  $P^-$  and  $P^+$  are empirical probabilities of false negatives and false positives,  $C^-$  and  $C^+$  are the costs of false negatives and false positives, and  $\pi$  is the *a priori* probability of a match. We take  $\pi = 0.05$  and we set  $C^- = 1$  and  $C^+ = 2$ , reflecting our presumption that banning an innocent account is more severe than failing to recognize a banned user. Our choices of  $C^-$  and  $C^+$  are only meant to reflect the asymmetric nature of the problem, although in practice these costs would be highly platform-specific.

We report the minimum value of  $C_{\text{det}}$  over all operating points (minDCF) and the value of  $P^+$  at the operating point for which  $P^- = P^+$ , also known as the *equal error rate* (EER).

**Baseline models.** We compare the proposed method with three baselines. First we consider TF-IDF vector representations of the concatenated text content of a sample, which are compared using cosine similarity. Next, we consider universal sentence encodings ([Cer et al., 2018](#)), which

Model	Training Length	EER	minDCF
TF-IDF	–	0.341	0.971
Universal	–	0.363	0.981
IUR	16	0.247	0.999
TP	1–8	0.169	0.848
TP	1–16	0.132	0.792

Table 2: Linking evaluation, averaged over 5 subreddits. The proposed model was trained using triplet loss. Smaller scores are better for both metrics.

are compared using angular distance. We experimented with two versions of this baseline, namely embedding the concatenation of the text content of the documents in a sample, and averaging the embeddings of the individual documents. Since we found the concatenated version to perform better, we only report on this variation. Finally, we consider IUR (Andrews and Bishop, 2019). Because this model only embeds samples of size 16, we pad samples containing fewer than 16 posts. To handle samples containing more than 16 posts, we organize the sample into contiguous groups of at most 16 posts, apply the embedding to each group, and average the embeddings.

**Results.** Table 2 compares the linking performance of the three baseline models along with two of variations of the proposed model arising from varying the sizes of the training samples. Note that both variations of the proposed model outperform the baselines. A further variation on this experiment is reported in Table 9 in which the queries are drawn from the training dataset, better reflecting the context of the motivating example.

Figure 6 shows the effect of the sizes of the targets on linking performance. Note that performance rapidly improves with larger samples, relative to the baselines. This trend is promising for our motivating application of ban enforcement, where it is desirable to recognize banned users as early as possible. Figure 7 shows receiver operator curves (ROC), which plot false positive rates against true positive rates as the operating points vary.

#### 4.4 Embeddings of variable-sized samples

Our experiments in §4.3 show that linking samples from newly created accounts to those of distinguished authors is more successful when using as much historical data from the distinguished authors

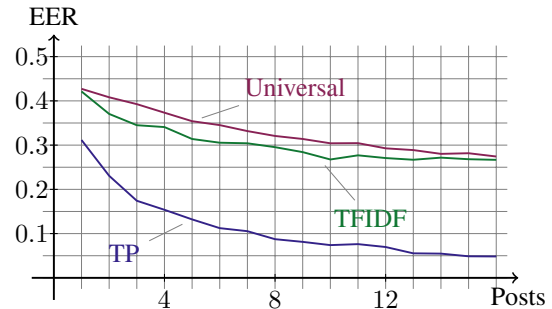


Figure 6: Equal error rate (EER) as the length of the target samples varies. Results are averaged over 5 subreddits. Smaller is better for EER. IUR is omitted as it does not admit variable-sized samples.

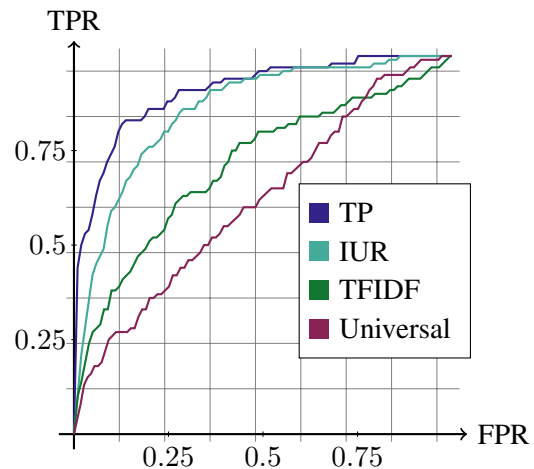


Figure 7: Receiver operator curves of models evaluated on a single subreddit and averaged over the five most popular subreddits. A point on a ROC near the northwest corner corresponds with an operating point with low FPR and high TPR, so a model is better if its ROC stretches further to the northwest.

as possible. However, computational constraints typically inhibit embedding full account histories during training. Instead, in §4.3 we embed large samples of distinguished authors’ histories using models trained on samples of sizes up to a maximum tractable length  $S$ . We take  $S = 16$  in our experiments as described in §2.3. Here, we examine the ability of a model trained on samples of sizes at most  $S$  to generalize to samples of sizes greater than  $S$ .

We also compare to a further baseline that averages single-post embeddings produced by a variation of the proposed model trained on single posts, which we denote by Avg. This is in contrast with the proposed model, which aggregates embeddings of multiple posts using an attention mechanism.

Table 3 shows the ranking performance of a num-



ber of variations of the proposed model trained with triplet loss and all features (TPS) on samples of fixed or varying sizes as specified. These results demonstrate that a model trained on variable-sized samples appears to generalize well to much longer samples. We observe substantially better performance compared to the simple averaging baseline, and only a slight decrease in performance compared to the fixed length models as the evaluation sample size increases beyond the lengths seen at training time.

Train Size	Test Size	Model	MRR	R@4
1	8	Avg	0.205	0.243
1-8	8	Prop	0.349	0.411
8	8	Prop	0.358	0.420
1	12	Avg	0.282	0.330
1-8	12	Prop	0.497	0.568
12	12	Prop	0.516	0.589
1	16	Avg	0.347	0.397
1-8	16	Prop	0.603	0.672
16	16	Prop	0.632	0.700

Table 3: Performance of the proposed model (Prop) and a simple averaging baseline (Avg), both trained and evaluated on fixed or variable-sized samples as shown.

#### 4.5 Selecting the distribution of sizes

As mentioned in §2.3, we use Beta (3, 1) to select sample sizes during training. We hypothesize that a negatively skewed distribution tends to improve training efficiency by supplying longer samples most of the time, while retaining the ability to handle shorter samples. To evaluate this claim, we investigate several distributions of varying degrees of negative skew.

Table 4 shows the ranking performance of variations of the proposed model trained on samples of sizes varying between 1 and 16 posts and evaluated on samples of size 16. These models differ only in the distribution used to select sample sizes. Indeed, the negatively skewed distributions *do* improve ranking performance over the uniform distribution, although the choice of negatively skewed distribution appears to be mostly immaterial.

Model	Skew	MRR	R@4
Unif (0, 1)	0	0.604	0.673
Beta (2, 1)	-0.566	0.626	0.692
tPois (16)	-0.786	0.632	0.703
Beta (3, 1)	-0.861	0.633	0.701
Beta (4, 1)	-1.049	0.633	0.704

Table 4: Ranking evaluations of models trained on samples of sizes selected with various distributions. tPois (16) is derived from the Poisson distribution with mean 16 by truncating its support to  $\{1, 2, \dots, 16\}$ .

## 5 Future work

This work motivates a number of interesting research questions. First, the proposed model makes use of publication times, but only avails of the hour of the day. It would be interesting to examine continuous-time variants of our encoder that incorporate relative time differences between actions when aggregating their embeddings, in light of the fact that patterns of user activity might be highly discriminative. For example, bots and spammers typically post at certain times of day and with particular frequencies. Separately, the proposed data augmentation methods we use to handle variable-sized samples may also be applicable in other settings, such as multi-document summarization (Liu and Lapata, 2019). Finally, the scores we use to determine author matches could be *calibrated*, providing confidence estimates associated with the account linking decisions.

To our knowledge, this work is the first to demonstrate the *feasibility* of a general-purpose account linking framework at web scale. Indeed, Figure 6 shows that performance improves as the size of the target increases, suggesting a speed-accuracy trade-off that can be tuned for different application settings. Expanding on an idea above, if confidence estimates were available, they could be used to inform the necessary sample sizes to achieve an acceptable level of risk.

Finally, we note that the generality of the proposed approach make it potentially applicable to a wide range of applications, including source code attribution (Burrows et al., 2009; Yang et al., 2017; Kalgutkar et al., 2019), plagiarism detection (Potthast et al., 2010; Meuschke et al., 2018; Foltýnek et al., 2019), and authorship attribution in collaborative documents (Flöck and Acosta, 2014; Dauber et al., 2017).

## References

- H. Alvari, S. Sarkar, and P. Shakarian. 2019. Detection of violent extremists in social media. In *2019 2nd International Conference on Data Intelligence and Security (ICDIS)*, pages 43–47.
- Nicholas Andrews and Marcus Bishop. 2019. [Learning invariant representations of social media users](#). 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 1684–1695, Hong Kong, China. Association for Computational Linguistics.
- Jason Baumgartner, Savvas Zannettou, Brian Keegan, Megan Squire, and Jeremy Blackburn. 2020. [The pushshift reddit dataset](#).
- David A Broniatowski, Amelia M Jamison, SiHua Qi, Lulwah AlKulaib, Tao Chen, Adrian Benton, Sandra C Quinn, and Mark Dredze. 2018. Weaponized health communication: Twitter bots and russian trolls amplify the vaccine debate. *American journal of public health*, 108(10):1378–1384.
- Steven Burrows, Alexandra L Uitdenbogerd, and Andrew Turpin. 2009. Application of information retrieval techniques for source code authorship attribution. In *International Conference on Database Systems for Advanced Applications*, pages 699–713. Springer.
- Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334):183–186.
- Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St. John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, Brian Strope, and Ray Kurzweil. 2018. [Universal sentence encoder for English](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 169–174, Brussels, Belgium. Association for Computational Linguistics.
- Sam Corbett-Davies and Sharad Goel. 2018. The measure and mismeasure of fairness: A critical review of fair machine learning. *arXiv preprint arXiv:1808.00023*.
- Walter Daelemans, Mike Kestemont, Enrique Manjavacas, Martin Potthast, Francisco Rangel, Paolo Rosso, Günther Specht, Efsthathios Stamatatos, Benno Stein, Michael Tschuggnall, Matti Wiegmann, and Eva Zangerle. 2019. [Overview of PAN 2019: Author Profiling, Celebrity Profiling, Cross-domain Authorship Attribution and Style Change Detection](#). In *10th International Conference of the CLEF Association (CLEF 2019)*. Springer.
- Edwin Dauber, Rebekah Overdorf, and Rachel Greenstadt. 2017. Stylometric authorship attribution of collaborative documents. In *International Conference on Cyber Security Cryptography and Machine Learning*, pages 115–135. Springer.
- Thomas Davidson, Dana Warmusley, Michael Macy, and Ingmar Weber. 2017. Automated hate speech detection and the problem of offensive language. In *Eleventh International AAAI Conference on Web and Social Media*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [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, Minneapolis, Minnesota. Association for Computational Linguistics.
- George R Doddington, Mark A Przybocki, Alvin F Martin, and Douglas A Reynolds. 2000. The nist speaker recognition evaluation—overview, methodology, systems, results, perspective. *Speech communication*, 31(2-3):225–254.
- Wenfei Fan. 2012. Graph pattern matching revised for social network analysis. In *Proceedings of the 15th International Conference on Database Theory*, pages 8–21.
- Fabian Flöck and Maribel Acosta. 2014. Wikiwho: Precise and efficient attribution of authorship of revisioned content. In *Proceedings of the 23rd international conference on World wide web*, pages 843–854.
- Tomáš Foltýnek, Norman Meuschke, and Bela Gipp. 2019. Academic plagiarism detection: a systematic literature review. *ACM Computing Surveys (CSUR)*, 52(6):1–42.
- Shunichi Ishihara. 2011. A forensic authorship classification in sms messages: A likelihood ratio based approach using n-gram. In *Proceedings of the Australasian Language Technology Association Workshop 2011*, pages 47–56.
- Vaibhavi Kalgutkar, Ratinder Kaur, Hugo Gonzalez, Natalia Stakhanova, and Alina Matyukhina. 2019. Code authorship attribution: Methods and challenges. *ACM Computing Surveys (CSUR)*, 52(1):1–36.
- Sungyeon Kim, Dongwon Kim, Minsu Cho, and Suha Kwak. 2020. [Proxy anchor loss for deep metric learning](#).
- Taku Kudo. 2018. [Subword regularization: Improving neural network translation models with multiple subword candidates](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 66–75, Melbourne, Australia. Association for Computational Linguistics.

- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. Albert: A lite bert for self-supervised learning of language representations. In *International Conference on Learning Representations*.
- Yang Liu and Mirella Lapata. 2019. [Hierarchical transformers for multi-document summarization](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5070–5081, Florence, Italy. Association for Computational Linguistics.
- Jing Lu, Chaofan Xu, Wei Zhang, Ling-Yu Duan, and Tao Mei. 2019. Sampling wisely: Deep image embedding by top-k precision optimization. In *The IEEE International Conference on Computer Vision (ICCV)*.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. [Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy. Association for Computational Linguistics.
- Leland McInnes, John Healy, and James Melville. 2020. [Umap: Uniform manifold approximation and projection for dimension reduction](#).
- Norman Meuschke, Vincent Stange, Moritz Schubotz, and Bela Gipp. 2018. Hyplag: a hybrid approach to academic plagiarism detection. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 1321–1324.
- Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory Diamos, Erich Elsen, David Garcia, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, and Hao Wu. 2017. [Mixed precision training](#).
- Kevin Musgrave, Serge Belongie, and Ser-Nam Lim. 2020. A metric learning reality check. *arXiv preprint arXiv:2003.08505*.
- Gordon Pennycook, Jonathon McPhetres, Yunhao Zhang, Jackson G Lu, and David G Rand. 2020. Fighting covid-19 misinformation on social media: Experimental evidence for a scalable accuracy-nudge intervention. *Psychological science*, 31(7):770–780.
- Martin Potthast, Benno Stein, Alberto Barrón-Cedeño, and Paolo Rosso. 2010. An evaluation framework for plagiarism detection. In *Coling 2010: Posters*, pages 997–1005.
- Florian Schroff, Dmitry Kalenichenko, and James Philbin. 2015. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823.
- Roy Schwartz, Oren Tsur, Ari Rappoport, and Moshe Koppel. 2013. Authorship attribution of micro-messages. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1880–1891.
- Giuseppe Silvestri, Jie Yang, Alessandro Bozzon, and Andrea Tagarelli. 2015. Linking accounts across social networks: the case of stackoverflow, github and twitter. In *KDWeb*, pages 41–52.
- Estée Van Der Walt and Jan Eloff. 2018. Using machine learning to detect fake identities: Bots vs humans. *IEEE Access*, 6:6540–6549.
- David A Van Leeuwen and Niko Brümmer. 2007. An introduction to application-independent evaluation of speaker recognition systems. In *Speaker classification I*, pages 330–353. Springer.
- Xun Wang, Haozhi Zhang, Weilin Huang, and Matthew R. Scott. 2019. [Cross-batch memory for embedding learning](#).
- Xinyu Yang, Guoai Xu, Qi Li, Yanhui Guo, and Miao Zhang. 2017. Authorship attribution of source code by using back propagation neural network based on particle swarm optimization. *PloS one*, 12(11):e0187204.

## A Hyperparameter selection

Top- $k$  loss involves a number of hyperparameters. In addition to  $k$  and the margin penalty  $m$ , one must also select the number  $n_+$  of targets from each class presented in every batch. We found the values of  $k, m, n_+$  suggested by (Lu et al., 2019) to not perform well in our setting. We therefore conducted a small grid search for these values, selecting the optimal configuration based on validation scores. We considered  $k \in \{4, 8\}$ ,  $n_+ \in \{4, 16\}$ , and  $m \in \{0.05, 0.40\}$ , resulting in  $k = 4$ ,  $n_+ = 8$ , and  $m = 0.25$ .

In experiments with the triplet loss, we use a fixed margin  $m = 0.2$ . We arrived at this value through the grid search illustrated in Table 5. We do not use dropout regularization, as previous work has shown dropout can degrade performance when training with larger datasets (Lan et al., 2019).

Margin	MRR	R@8
0.1	0.59	0.73
0.2	<b>0.62</b>	<b>0.75</b>
0.3	0.59	0.74
0.4	0.61	0.74
0.5	0.60	0.73
0.6	0.59	0.73
0.8	0.58	0.72

Table 5: Validation ranking results for various triplet loss margins.

## B Implementation details

**Mixed precision.** We use mixed precision training (Micikevicius et al., 2017). This reduces the GPU memory consumed at training time by about half through the use of half precision floats, enables faster forward and backward pass computations, and allows for a larger batch size.

**Multi-GPU training.** Using 2 V100 GPUs to train the model significantly speeds up the process by quadrupling the effective batch size. Our models have an average training time of around 72 hours.

**Simplified text encoding.** We limit our convolutional text encoders to windows of 2, 3, and 4 subwords, excluding the largest window of 5 used by other models. Surprisingly, this did not impact ranking performance, which suggests that a small receptive field is sufficient for purposes of comparing authorship. As further support for this claim,

we also experimented with larger receptive fields than 5, which *reduced* ranking performance.

**Model Hyper-parameters.** Our reported models are trained with an embedding dimension of  $D = 1024$ , 512 convolutional filters, and an attention mechanism producing outputs of dimension 512. Additionally the subword and subreddit embedding dimensions are both  $N = 512$ .

**Model Parameters** Table 6 shows the numbers of parameters of the proposed model when using various combinations of text content (T), publication time (P), and subreddit (S).

Model	Parameters
T	44.0M
TP	44.1M
TPS	45.5M

Table 6: Numbers of parameters in various trained models.

## C Text encoding

We consider two methods to encode raw text content into integer arrays, namely taking the integer values of the corresponding UTF-8 encoded bytes directly, and using the SentencePiece unigram subword model (Kudo, 2018). SentencePiece tokenizes select character groupings according to a pretrained vocabulary of a specified size, which can be orders of magnitudes larger than the vocabulary of size  $2^8$  used by the byte encoding. For experiments conducted in §4, we used SentencePiece with a vocabulary size of  $2^{16}$ .

The potentially large disparity in vocabulary size between encoding methods can result in text encoded as integer arrays of significantly different lengths. In light of the need to truncate these arrays at training, we hypothesize that subword encoding with a large vocabulary results in better model performance as more textual information is captured after truncation.

To evaluate this claim, we fit additional SentencePiece subword models with vocabulary sizes  $2^{12}$  and  $2^{14}$  on text content from MUD. Table 7 provides ranking performance of variations of the proposed model trained with only the text feature on samples of fixed sized 16 and evaluated on samples of fixed sized 16. We remark that increases in vocabulary size *do* correspond with increases



Model	Size	MRR	R@4	R@8
SP	$2^{16}$	0.372	0.439	0.512
SP	$2^{14}$	0.352	0.412	0.483
SP	$2^{12}$	0.307	0.360	0.433
Byte	$2^8$	0.105	0.124	0.163

Table 7: Ranking performance of models using various subword vocabulary sizes. All models use Sentence-Piece (SP) or the byte encoding (Byte) described in [Appendix C](#).

in ranking performance. We found that increasing the vocabulary size beyond  $2^{16}$  did *not* increase performance further.

## D Further statistics about MUD

[Table 8](#) shows some further details of the MUD dataset introduced in [§4.1](#).

Number of users contributing	1,071,477
Number of posts	321,659,421
Mean post length	42.5 tokens
Mean number of posts contributed by a user	300.2
Mean number of subreddits accessed by a user	22.1
Mean number of months a user was active	9.9
Percentage of posts containing more than 64 tokens	17.37%

Table 8: Some statistics of the Million User Dataset.

## E Further experiments

We repeat the linking experiment from [§4.3](#) using queries that are observed at training time. We expect this to improve performance, something which is confirmed in [Table 9](#).

Model	Training Length	EER	minDCF
TF-IDF	–	0.362	0.967
Universal	–	0.371	0.989
IUR	16	0.223	0.951
TP	1–8	0.145	0.791
TP	1–16	0.136	0.765

Table 9: Same experiment as [Table 2](#), but with queries *observed* at training time.