



Deep Indexed Active Learning for Matching Heterogeneous Entity Representations

Arjit Jain
IIT Bombay
arjit@cse.iitb.ac.in

Sunita Sarawagi
IIT Bombay
sunita@iitb.ac.in

Prithviraj Sen
IBM Research, Almaden
senp@us.ibm.com

ABSTRACT

Given two large lists of records, the task in entity resolution (ER) is to find the pairs from the Cartesian product of the lists that correspond to the same real world entity. Typically, passive learning methods on such tasks require large amounts of labeled data to yield useful models. Active Learning is a promising approach for ER in low resource settings. However, the search space, to find informative samples for the user to label, grows quadratically for instance-pair tasks making active learning hard to scale. Previous works, in this setting, rely on hand-crafted predicates, pre-trained language model embeddings, or rule learning to prune away unlikely pairs from the Cartesian product. This blocking step can miss out on important regions in the product space leading to low recall. We propose DIAL, a scalable active learning approach that jointly learns embeddings to maximize recall for blocking and accuracy for matching blocked pairs. DIAL uses an Index-By-Committee framework, where each committee member learns representations based on powerful pre-trained transformer language models. We highlight surprising differences between the matcher and the blocker in the creation of the training data and the objective used to train their parameters. Experiments on five benchmark datasets and a multilingual record matching dataset show the effectiveness of our approach in terms of precision, recall and running time.

PVLDB Reference Format:

Arjit Jain, Sunita Sarawagi, and Prithviraj Sen. Deep Indexed Active Learning for Matching Heterogeneous Entity Representations. PVLDB, 15(1): 31 - 45, 2022.
doi:10.14778/3485450.3485455

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/ArjitJ/DIAL>.

1 INTRODUCTION

Entity resolution (ER) is a crucial task in data integration whose goal is to determine whether two mentions refer to the same real-world entity. With a history going back at least half a century (following Fellegi and Sunter [20]’s seminal work), the task goes by various names and formulations, with the most common one being: Given two sets R and S , for each pair of instances $(r, s) \in R \times S$ classify (r, s) as either being a match or a non-match. In essence, this is

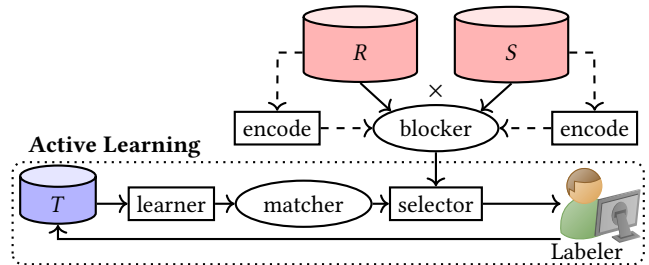


Figure 1: Traditional paired classification AL.

an instance of *paired classification* that requires learning a highly accurate binary-class classifier or *matcher*.

ER has a rich history of employing active learning (AL) [60] instead of supervised or passive learning which harbors some advantages such as incrementally adding labeled pairs instead of requiring voluminous labeled data up-front to train the matcher. A variety of previous works on ER [29, 39, 40, 54, 58] have utilized the AL workflow shown in Figure 1, with minor modifications. In each iteration, the learning algorithm (*learner*) learns a *matcher* (shown in an ellipse which we use to denote model components) from T , the labeled pairs collected from the (human) labeler so far, while the example selector (*selector*) chooses the most informative unlabeled pairs to acquire labels for. After including the new labels into T , the process repeats until we learn a matcher of sufficient quality. Popular choices for matcher includes support vector machines [58], random forests [39], and neural networks [29]. Popular choices for selector includes query-by-committee [21, 61] which has seen wide usage in ER [11, 40] and uncertainty sampling [29].

To efficiently pare down the number of unlabeled pairs in $R \times S$ that the selector needs to choose from, one usually employs a pre-specified *blocking* function. Commonly used blocking functions include string similarity measures (e.g. Jaccard similarity) to compare string representations of r and s , and keep only those pairs whose similarity exceeds a pre-determined threshold [39]. Konda et al. [31] recommend that the user acquire some domain knowledge about R, S so as to be able to specify an effective blocker. Even if domain knowledge is available, the user’s choice may still be sub-optimal. In some situations, it may even be impossible to acquire such knowledge, for example when one of R or S is in a language unfamiliar to the user (aka cross-lingual ER [38]). The other, possibly more disconcerting, conceptual issue with Figure 1 is that the blocker is removed from the matcher. To be clear, both matcher and blocker are paired classifiers but the requirements of them are different. While the matcher needs to provide high classification accuracy, the blocker only needs to efficiently identify matches while rejecting as many non-matches as possible (in other words,

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.
Proceedings of the VLDB Endowment, Vol. 15, No. 1 ISSN 2150-8097.
doi:10.14778/3485450.3485455

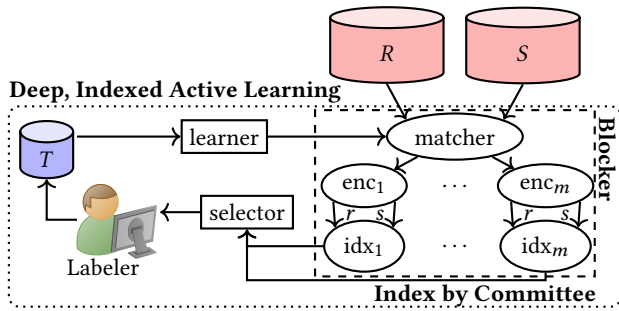


Figure 2: Deep, indexed AL with a committee of encoders.

high recall is desired). This implies that ideally, the blocker should be integrated into the AL feedback loop. As we obtain more labeled data, we expect both matcher and blocker to benefit instead of benefiting one and not the other as Figure 1 indicates. While there exist proposals to learn the blocking function automatically [8], these require copious amounts of labeled data up-front and thus it is not clear how to combine this with low-resource AL setting where such labeled data may not be available. Optionally, DeepER [17] encodes $r \in R$ and $s \in S$ into fixed-dimensional vectors or *encodings* $E(r)$ and $E(s)$, respectively, (see dashed edges in Figure 1). Similar pairs (r, s) may then be retrieved via nearest neighbor search implemented using locally sensitive hashing. However, even in this case the blocker is not integrated with the AL feedback loop.

Given the previous discussion, our goal is to propose a new AL approach for ER that satisfies the following desiderata:

- To ensure that both benefit from newly acquired labels, the blocker and matcher should be integrated.
- Slightly at logger heads with the previous property, we would also like to train the matcher and blocker with distinct loss functions since they need to satisfy distinct requirements.
- Need to achieve all this without adversely affecting scalability, since this is one of the purposes of the blocker in Figure 1.

Our proposed integrated matcher-blocker combination and new AL workflow is shown in Figure 2. Compared to Figure 1, the two most notable differences are 1) the blocker (dashed box) is now part of the AL feedback loop, and 2) the matcher is a component within the blocker. As base matcher, we use *transformer-based pre-trained language models* (TPLM) [16, 35] which have recently led to excellent ER accuracies in the passive (non-AL) setting [34]. Before we describe details of the proposed approach, please note that TPLMs can be invoked in two distinct modes. To obtain a prediction for a pair (r, s) , we invoke the matcher, a (fine-tuned) instance of TPLM, in *paired* mode where we concatenate string representations of r and s to obtain a joint representation. *Single* mode is where we input one of r or s 's string representation to obtain its encoding. To implement blocking, we invoke TPLM in single mode by first populating an index structure (FAISS [28]) with $E(r)$, $\forall r \in R$, followed by probing with $E(s)$ given $s \in S$ to retrieve potentially similar pairs for the selector to choose from. Since the blocker needs to attain high recall, we find that one encoding of r or s is not sufficient, improved recall can be achieved if we allow minor variations of $E(r)$ and $E(s)$. To this end, we allow

for multiple, distinct affine transformations of the base TPLM's last layer. This combination of TPLM, multiple encoders enc_1, \dots, enc_m and indices idx_1, \dots, idx_m , referred to as *index by committee* (IBC), is pictorially depicted in Figure 2. While previous works on AL and ER have attempted to learn committees instead of a single classifier [11, 39, 40, 58], none of them consider TPLMs, and none of them use committee for designing a high recall blocker.

A primary challenge of our integrated matcher-blocker system is training them simultaneously so that the blocker recalls all likely duplicates based on embedding similarity, and the matcher precisely separates duplicates from non-duplicates in paired mode. We show that the conventional classification loss on labeled examples that trains the matcher well, performs very poorly on the blocker. This led us to design a new contrastive training objective for the blocker that separates labeled duplicates from *random* non-duplicate pairs. In entity resolution tasks, random non-duplicates are generally much easier to separate out than the difficult non-duplicates selected during active learning [58]. While such difficult near-duplicates are essential for training a *precise matcher*, they interfere with the goals of co-embeddings duplicates during blocking. We observe dramatic drops in recall when blocker embeddings are trained with actively labeled non-duplicates, likewise dramatic drop in precision when the matcher's classifier is trained with random negatives. These findings were key to our jointly training the matcher and blocker in an active learning loop so as to match or even surpass the yield obtained by hand-designed rules on existing benchmarks. Additionally, on heterogeneous entity lists where existing methods relied on pre-trained embeddings, we obtained significantly higher recall with our method of training the blocker. Our contributions are:

- To the best of our knowledge, ours is the first active learning ER proposal to *integrate* the matcher with the blocker. With availability of more labels, improvements in the former directly benefits the latter.
- We design a novel method of learning record embeddings for the blocker using (1) a contrastive training objective and (2) random non-duplicates. This design choice is crucial to achieve high recall; and provides as much as 25 points increase compared to using matcher embeddings as-is.
- Learning a committee of encodings is in itself a novel contribution. To the best of our knowledge, we are not aware of any previous work that can learn a committee of TPLMs. By combining with indexing, leads to IBC, a novel, fast and effective example blocking technique.
- We evaluate the efficacy of our learned blocker by comparing with (1) hand-crafted blocking functions used in popular ER datasets, and (2) state of the art learned embeddings methods on a multilingual dataset and five ER datasets.
- DIAL provides an *absolute* improvement on the F1 scores by 6 – 20% on two product datasets, 4 – 10% on a bibliographic dataset, 40 – 55% on a textual dataset, and 5 – 18% on a multilingual dataset, over baseline approaches demonstrating the effectiveness of DIAL across various real world datasets. On some of these datasets DIAL produces even better recall than hand-tuned blocking functions without any external knowledge about the domain and with only limited number of judiciously chosen label.

2 PRELIMINARIES AND BACKGROUND

We formally state our problem and provide relevant background in this section.

2.1 Problem Statement

Given two large lists R and S of entities, our goal is to design an end to end system that can identify the subset DUPS of $R \times S$ that are duplicates across the two lists. R and S could be the same list, and the matchings could be many to many. Each entry $r \in R$ or $s \in S$ could consist of one or more attributes that are predominantly textual. In general, the attributes across the lists may not be aligned, and the space of their values may be incomparable. For example, list R may list product names and descriptions in German whereas list S may be in English. Our goal is to learn in an integrated active learning loop (1) a *blocker* to efficiently identify the subset CAND of $R \times S$ that are likely duplicates, and (2) a *matcher* to assign a final verdict of duplicate or not for each entity pair (r, s) in the filtered set CAND . We are given three types of resources: a transformer based pretrained language model (TPLM), a small seed labeled dataset T of duplicates and non-duplicate pairs, and a labeling budget B of getting human labels on pairs selected from $R \times S$ to augment T .

2.2 Pre-trained Language Models

Transformer based pretrained LMs (TPLM) such as BERT [16] and RoBERTa [35] have been shown to transfer remarkably well to many different tasks and domains. The input to the transformer is a sequence of tokens. The transformer uses multiple layers of self-attention to output for each token a fixed dimensional contextual embedding. The hundreds of million parameters used in a transformer are pre-trained using large amounts of unlabeled text corpus e.g. Wikipedia. This results in assigning each word an embedding that captures its semantics in the context of the current sentence. These highly contextual embeddings have been found useful in a number of downstream NLP tasks. In ER they have been shown to lead to robustness to spelling mistakes, and abbreviations, and provide state of the art performance on "dirty" datasets [10, 34]. A standard approach to use these models in a new task is to add task specific layers on top of the transformer and fine-tune using a task specific objective. There are two common modes to fine-tune a transformer for a pairwise classification task required in ER.

2.2.1 Paired mode. In this mode the transformer is fed a concatenation of the tokens of the two records as follows:

$$[\text{CLS}], r_1 \dots r_n, [\text{SEP}], s_1 \dots s_m, [\text{SEP}] \quad (1)$$

where r_1, \dots, r_n denote tokens of record r , $s_1 \dots s_m$ denote tokens of s , CLS denotes a special start token and SEP denotes a special separator token. The last layer of the transformer assigns fixed d dimensional contextual embeddings to all $m + n + 3$ tokens. The contextual embedding of the [CLS] token is treated as an embedding $E(r, s)$ of the pair. This embedding is used to classify the pair as duplicates or not via additional light-weight layers. This is the mode we use for the matcher since the learned attention across tokens in the records can focus on distinguishing words. Consider an example of a pair of records describing two different editions of the same book. An embedding based model can have a hard time trying to distinguish these two instances, however, a transformer model

can by aligning the attention between the tokens corresponding to book edition between the two instances. Other examples include the price attribute in a products dataset, and house number in a postal addresses dataset.

2.2.2 Single mode. The above paired mode is not practical to invoke on every (r, s) in the Cartesian product $R \times S$. A second way is to first separately encode each record. For a record x in R or S we obtain its embedding from the TPLM by first feeding to the transformer:

$$[\text{CLS}], x_1 \dots x_n [\text{SEP}] \quad (2)$$

where x_1, \dots, x_n denote the tokens in record x . We obtain fixed d dimensional contextual embeddings $E(x_1), \dots, E(x_n)$ from the TPLM. We then define the embedding of the record x as the mean of its token embeddings.

$$E(x) = \frac{1}{n} \sum_{i=1}^n E(x_i) \quad (3)$$

For a pair of records (r, s) we separately compute embeddings $E(r)$ and $E(s)$ and decide on whether they are duplicate or not based only on these fixed embeddings. A well-known example is SentenceBERT [56] whose classifier takes as input the concatenation of the embedding $E(r)$ of r , embedding $E(s)$ of s , and the absolute element-wise difference between the two embeddings $|E(r) - E(s)|$ and adds a linear layer above it. After training with appropriate labeled data, these embeddings can be used for efficient nearest neighbour search to retrieve likely duplicates. We will harness the single mode for the design of our blocker.

2.3 Example selection for ER

2.3.1 Query-by-Committee via Bootstrap. Query-by-Committee (QBC) [21, 61] has a rich history of application in ER going back to ALIAS [58]. We review 1) the bootstrap-based classifier-agnostic approach towards building a committee [40], followed by 2) selecting examples for labeling using said committee. While there exist many techniques to build a committee of classifiers given the same labeled data, most of these are specifically designed for certain classifiers, e.g., randomizing the choice of the feature to split on while adding a node in the decision tree is a specific technique to learn a committee of decision trees [58]. Mozafari et al. propose *bootstrap* as a way to build a committee that is *agnostic* to the classifier being used. Given labeled data T , bootstrap creates multiple versions T_1, \dots, T_m by sampling from T with replacement so that each T_i contains the same number of pairs as T . Subsequently, one may use T_i to train a member of the committee by using it as training data. Given an unlabeled pair (r, s) , one may then compute the variance in its predicted label as:

$$\text{var}(r, s) = \frac{\#\text{match}(r, s)}{m} \left(1 - \frac{\#\text{match}(r, s)}{m} \right)$$

where $\#\text{match}(r, s)$ denotes the number of committee members predicting (r, s) to be a duplicate out of the m -sized committee. Pairs with higher variance are selected for labeling.

2.3.2 Uncertainty Sampling. Besides variance, other metrics are also available to measure the uncertainty of the prediction for (r, s) . These may be used independent of the committee, especially when

the classifier produces prediction probabilities besides the label. DTAL [29] uses (conditional) *entropy*:

$$H(p) = -p \log p - (1 - p) \log(1 - p) \quad (4)$$

where p denotes $\Pr(y = \text{match} | (r, s))$, the predicted probability of (r, s) being a match.

2.3.3 High Confidence Sampling with Partition. Besides entropy, DTAL [29] also proposes High Confidence Sampling with Partition. They divide the candidate set into two subsets consisting of pairs that are predicted as positives, and negatives respectively by the matcher. From both these sets they choose an equal amount of most confident and least confident pairs, based on their entropy, giving four sets, $p_{hc}, p_{lc}, n_{hc}, n_{lc}$ representing high and low confidence positives, and high and low confidence negatives respectively. They query the user to label p_{lc} and n_{lc} , but they do NOT query the user to label p_{hc} and n_{hc} . Instead, they directly add them to the labeled positives and negatives, i.e. $T_p \leftarrow T_p \cup p_{hc}$ and $T_n \leftarrow T_n \cup n_{hc}$

2.3.4 BADGE. In a batch active learning setup, BADGE [5] tries to combine uncertainty and diversity for example selection by computing hallucinated gradient embeddings. Given a neural network classifier $f(x; \theta)$, with weights θ_0 , and a query point x from the candidate set, BADGE calculates \hat{y} , the most likely label for x according to the class probabilities output by f . It then uses \hat{y} to compute the gradient embedding

$$g_x = \frac{\partial}{\partial \theta_{\text{out}}} \ell(f(x; \theta), \hat{y}) \Big|_{\theta=\theta_0}$$

where θ_{out} refers to the parameters of the output layer, and ℓ is a loss function, usually taken to be the standard cross entropy loss. Notice that the magnitude of these gradient embeddings can be used as a proxy for uncertainty, as confident samples will have lower gradient magnitudes. To incorporate diversity, examples to query the user are selected using the k-MEANS++ [4] seeding algorithm on the set $\{g_x : x \in \text{CAND}\}$.

3 DIAL

DIAL starts with an initial set of labeled pairs T of duplicates and non-duplicates and iteratively collects B more labeled pairs in an active learning loop. In each iteration of the loop, it performs the following steps: (1) trains a Matcher model that given a pair of records can assign a probability of the pair being duplicate, (2) trains a Blocker model to encode records in R or S so that duplicates are close, (3) performs an indexed nearest neighbor search over the encodings to filter a candidate set $\text{CAND} \subset R \times S$ of likely duplicate pairs, (4) selects a subset SEL of CAND using uncertainty assignments from Matcher, (5) collects user’s duplicate or not labels on pairs in SEL and augments T . At the end of the loop, all pairs in the candidate set predicted duplicates by the Matcher are returned as the duplicate set. This labeling loop differs from earlier AL-based ER systems in one crucial way. While existing systems assume a fixed candidate set CAND under a user-provided or pre-trained blocking function, we propose to learn a Blocker and adaptively create candidates CAND within the AL loop. Our challenge then is how to perform this step while ensuring that our learned Blocker can match hand-crafted rules in terms of recall, and do that without enumerating the Cartesian product $R \times S$.

Our matcher and blocker are integrated and both leverage TPLMs. We present the design of the main modules of DIAL. An overview appears in Figure 3.

3.1 Matcher

For each record pair (r, s) the matcher needs to assign a probability $\Pr(y = 1 | (r, s))$ of the pair being a duplicate. The matcher uses the transformer in the paired mode described in Section 2.2.1 to get a joint embedding $E(r, s) \in R^d$ of (r, s) . Let Θ denote all the parameters of the transformer. These embeddings are converted into a probability of the pair being duplicates using additional neural layers $F_W : R^d \mapsto R$:

$$\Pr(y = 1 | (r, s)) = (1 + \exp(-F_W(E(r, s))))^{-1} \quad (5)$$

where W denote parameters of the matcher specific layers to be learned along with parameters Θ of the transformer. In our case, F_W comprised of a linear layer, followed by a tanh activation, followed by another linear layer to get a single scalar score which is then converted into a probability using the above sigmoid function. During training, the initial values of parameters Θ are from the TPLM whereas W, b take random values. All three sets of parameters are optimized using the standard cross entropy loss on the labeled training set T .

$$\begin{aligned} \min_{\Theta, W} & \sum_{(r^i, s^i) \in T_p} \log(1 + \exp(-F_W(E_{\Theta}(r^i, s^i)))) \\ & + \sum_{(r^i, s^i) \in T_n} \log(1 + \exp(F_W(E_{\Theta}(r^i, s^i)))) \end{aligned} \quad (6)$$

where T_p denotes the duplicate pairs in T and $T_n = T - T_p$ denotes the non-duplicates. In the above we put the subscript Θ on the embeddings to denote that the transformer parameters are further fine-tuned to achieve the matcher’s goal of assigning probability close to 1 to the duplicates and close to 0 to the non-duplicates. See Step 1 in Figure 3 for a summary of this part.

3.2 Blocker

Here our goal is to obtain embeddings of each record in R and S so we can retrieve likely duplicates via nearest neighbor search. Existing methods for getting such embeddings is to use the transformer in *single* mode as outlined in Section 2.2.2, either as-is or with further fine-tuning using T as in SentenceBERT [34, 56]. We will show in Section 4.4 that both methods perform surprisingly poorly in retrieving duplicates. The blocker in DIAL makes three important design choices that jointly provide significant gains over existing methods. We outline each of these next.

3.2.1 Index by Committee of Embeddings (IBC). Our blocker assigns a committee of N different embeddings to a record in R or S . Traditionally in AL, committees (Section 2.3.1) are used to assign uncertainty values during example selection. Here, we propose to use multiple embeddings for a different goal of casting a wider net so that all likely duplicates are covered in *any one* of the N embeddings.

We start with the d -dimensional embeddings $E(x)$ obtained from the Matcher-trained Transformer operating in single mode as described in Equation 3. Then we create a committee of N different

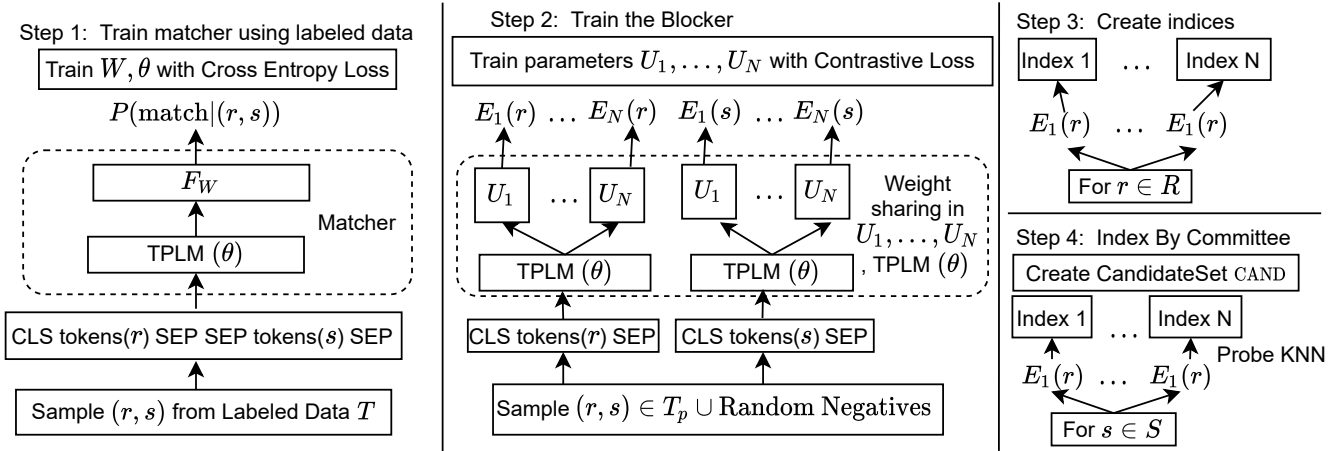


Figure 3: Outline of the proposed system. DIAL integrates TPLM based matcher and blocker models. In each iteration of the active learning loop, it performs the following steps: trains a Matcher model that given a pair of records can assign a probability of the pair being duplicate, trains a Blocker model to independently encode records in R or S , and performs an indexed nearest neighbor search over the encodings to filter a candidate set $\text{CAND} \in R \times S$ of likely duplicate pairs. Candidate set CAND is used by the selector to obtain a subset of samples to be labeled by the user.

light-weight layers to produce a set of N d -dimensional embeddings: E_1, \dots, E_N . Each committee member k , first chooses a fixed random mask $M_k \in \{0, 1\}^d$ to retain only a random fraction p of the initial embeddings $E(x)$. This step is inspired by the choice of random attribute selection in random forests [9]. Then a linear layer transforms the masked embeddings via learned parameters to obtain the k -th embedding vector $E_k(x)$ as:

$$E_k(x) = \tanh(U_k(M_k \odot E(x), 1)) \quad (7)$$

where $U_k \in R^{d(d+1)}$ denote the learned parameters used to obtain the k -th embedding vector of record x . The transformer parameters Θ used to compute $E(x)$ are not trained by the blocker. We next describe how we train the U_k parameters.

3.2.2 Choice of Training data. One subtle problem we encountered is using the labeled data T collected via AL to train the blocker. The negatives in T are mostly *near-duplicates* and were chosen by AL because they were hard to separate from duplicates. While such hard negatives are extremely useful for learning a precise matcher as several previous AL work have shown [29, 58], they are detrimental to learning good embeddings for blocking where the goal is high recall rather than high-precision. Embeddings trained to separate the similar non-duplicates T_n from the actual duplicates T_p , might also throw the unseen duplicates apart. We therefore create easier non-duplicates in the following way:

Given a set D_p of b duplicates in a training batch, we randomly sample a set $\text{rand}(R)$ of b records from R and an independent random set $\text{rand}(S)$ of b records from S . We then obtain embeddings $E(x)$ for all records in $\text{rand}(R), \text{rand}(S)$, and each record in D_p the duplicate pairs. Now each committee randomly shuffles the set of records in $\text{rand}(R), \text{rand}(S)$ and obtains a random set of b non-duplicate pairs $(r_1, s_1) \dots (r_b, s_b)$ by concatenating the shuffled lists. Further, for each duplicate pair (r_p, s_p) in the training batch D_p we obtain further non-duplicates as $(r_p, s_i), (r_i, s_p)$ for $i = 1 \dots b$.

3.2.3 Choice of Training Objective. Given the set of duplicates and non-duplicates, a default training objective would be impose a binary classification loss to separate them as is done for the matcher in Eq 6. However, again considering the differing goals of the two systems, we propose a different contrastive training objective that jointly separates a duplicate from all non-duplicates. The contrastive loss requires a similarity function $\text{sim}(u, v)$ between any two embedding vectors u, v . The training objective of the k -th committee member is then

$$\max_{U_k} \sum_{(r_p, s_p) \in T_p} \log \left[\frac{s(r_p, s_p)}{s(r_p, s_p) + \sum_{i=1}^b (s(r_i, s_p) + s(r_p, s_i) + s(r_i, s_i))} \right] \quad (8)$$

where $s(r, s) = e^{\text{sim}(E_k(r), E_k(s))}$

We use the negative squared ℓ_2 distance as a similarity function. Scaled cosine similarity is another good choice. The only requirement is that we should be able to retrieve nearest neighbour efficiently using that similarity function.

Step 2 in Figure 3 summarizes the training of the blocker. Notice the differences in the input, the training objective, and the training dataset with the training of the matcher in step 1.

3.3 Overall Algorithm

Algorithm 1 outlines the pseudo-code of DIAL. In each round of Active Learning, DIAL first trains the TPLM parameters Θ , and parameters W of the matcher specific layer F_W with the binary classification objective (Equation 6) on the labeled data T . It then freezes the weights of parameters Θ , and creates a committee where each member implements an embedding layer as described in Section 3.2 (Equation 7). To train the committee, it samples duplicate pairs from the labeled data T_p , creates random negative pairs (r, s) where

$r \in \text{rand}(R)$ and $s \in \text{rand}(S)$, and individually obtains the transformer representations for each of these. Every committee member computes individual embeddings for each of these instances, and is trained using the contrastive objective (Equation (8)). After training the committee, each member creates an index on the embeddings of instances in R , and queries this index to get the k nearest neighbours for each instance in S . The closest pairs across all members are used to construct the set CAND , which are fed to an active learning instance selector to select the most informative B pairs to be labeled by the user. DIAL is agnostic to the specific selection algorithm and we present results with many existing selection algorithms in Section 4.7. Our default is uncertainty sampling (Eq 4). Figure 3 highlights the main operations performed by DIAL in an active learning round, and clearly describes the data flow.

Algorithm 1 Pseudo-Code of the proposed system DIAL.

Require: TPLM with parameters Θ , Lists R and S , Seed Labeled Data T , CAND Size, Labeling Budget per round B , Committee Size N , Number of neighbours k

- 1: **for each** round of Active Learning **do**
- 2: **▷ Train the matcher**
- 3: Find Θ, W that minimize Eq. (6) using T
- 4: **▷ Create committee: each member k , has trainable parameters U_k , computes embedding $E_k(x)$ using Eq. (7)**
- 5: **▷ Train the embeddings**
- 6: **for each** committee member k **do**
- 7: Find U_k that maximize Eq. (8) using T_p & Random Negatives (See Sec 3.2.2)
- 8: **end for**
- 9: **▷ Retrieving Pairs**
- 10: Create Indexes IDX_i for each committee member i
- 11: **for each** r in R **do**
- 12: Compute TPLM embedding $E(r)$
- 13: **for each** committee member k **do**
- 14: Add $E_k(r)$ to IDX_k
- 15: **end for**
- 16: **end for**
- 17: Create list RP to store Retrieved Pairs
- 18: RP = []
- 19: **for each** s in S **do**
- 20: Compute TPLM embedding $E(s)$
- 21: **for each** committee member c **do**
- 22: Add k nearest neighbours of $E_c(s)$ in IDX_c to RP
- 23: **end for**
- 24: **end for**
- 25: Create CAND containing the closest pairs from RP
- 26: Select B pairs from CAND & query user labels. (See Sec 4.7)
- 27: Update T with the newly labeled data
- 28: **end for**

4 EXPERIMENTS

We present an extensive comparison of DIAL with existing methods based on hand-crafted predicates, learned embeddings, and existing meta-blocking methods. We also present a detailed ablation study and analyze DIAL’s running time.

Table 1: Statistics reporting the scale of the datasets used to evaluate DIAL.

| Dataset | $ R $ | $ S $ | $ \text{DUPS} $ | $ \frac{\text{DUPS}}{R \times S} $ | $ \mathcal{D}_{\text{test}} $ |
|----------------|-------|-------|-----------------|------------------------------------|-------------------------------|
| Walmart-Amazon | 2554 | 22074 | 1154 | $\sim 2e-5$ | 2049 |
| Amazon-Google | 1363 | 3226 | 1300 | $\sim 3e-4$ | 2293 |
| DBLP-ACM | 2616 | 2294 | 2224 | $\sim 3e-4$ | 2473 |
| DBLP-Scholar | 2616 | 64263 | 5347 | $\sim 3e-5$ | 5742 |
| Abt-Buy | 1081 | 1092 | 1097 | $\sim 1e-3$ | 1916 |
| MultiLingual | 100k | 100k | 100k | $\sim 1e-5$ | 2000 |

4.1 Datasets

We validate our approach on five widely used real world datasets from DeepMatcher [41], ER Benchmark [32] and the Magellan data repository [15] as summarized in Table 1. Walmart-Amazon, Amazon-Google and Abt-Buy are product datasets, whereas DBLP-ACM and DBLP-Scholar are citation datasets. Abt-Buy is a textual dataset, whereas the other four are structured datasets. To use Walmart-Amazon and Amazon-Google as structured datasets, we follow the schema used by DeepMatcher [41]. As we motivated in Section 2.1, there may be scenarios where the elements of lists R and S are incomparable, making rule based blocking methods infeasible. To make a case for our method for such settings, we also evaluate DIAL against baseline approaches on a multilingual dataset [26]. Section 4.5 provides more information on the dataset, as well as describes the corresponding experiments and results.

Evaluation Metrics. We are interested in three questions to evaluate our matcher and blocker system:

- Recall of the Blocker: What fraction of the duplicates DUPS are retrieved in CAND .
- Overall F1 score on unseen test pairs: How accurately can our system classify unseen pairs from a test set, $\mathcal{D}_{\text{test}}$, into duplicates and non duplicates. The overall system predicts a record pair to be a duplicate only if the record pair is retrieved in CAND , and the matcher assigns a probability greater than 0.5 of the pair being a duplicate.
- Overall F1 score on all pairs: How accurately can our system find all duplicate pairs from the set of all possible pairs in the data? We compare the gold list of all duplicates in the data, to pairs that our system predicts to be duplicates.

The test dataset, $\mathcal{D}_{\text{test}}$, is the same dataset used to evaluate DeepMatcher. Hence, the test set evaluation metric gives us a way to compare our system with other approaches that may or may not be using active learning. However, the evaluation on all pairs is more aligned with the practical utility of any EM system.

4.2 Implementation Details

Compute. We implemented all the systems, and experiments, in PyTorch 1.6 [52], and used transformers library by huggingface [68]. All experiments were conducted on a machine with 64 2.10GHz Intel Xeon Silver 4216 CPUs with 1007GB RAM and a single NVIDIA Titan Xp 12 GB GPU with CUDA 10.2 running Ubuntu 18.04. To retrieve nearest neighbours we use the Facebook AI Similarity Search (FAISS) [28] library.

Model Architectures. We use the pre-trained RoBERTa model as our base transformer. The RoBERTa model builds on BERT, but with a careful selection of training sensitive hyperparameters like learning rate, and batch size. The RoBERTa model was pre-trained on five English corpora of 160GB total size, 10 times that used for BERT. We use 6 layers out of the 12 layered uncased RoBERTa base model. We use 12 attention heads, with 768 dimensional hidden vectors, and limit the number of input tokens to 512. The paired classifier, on top of the base RoBERTa model, is the default classification head used in RoBERTa based models, consisting of two dropout layers with dropout probability 0.1, a fully connected layer with a tanh activation, and a softmax classifier layer. Unless stated otherwise, DIAL uses a committee of size $N = 3$, with each member using a masking probability $p = 0.5$.

Optimization. We use the AdamW (Adam with Weight Decay) optimizer [36], with a learning rate of $3e - 5$ for the base transformer, and $1e - 3$ for the embedding and classifier layers. We use a linear learning rate schedule with no warm-up steps. The choice of optimizer parameters, and learning rate schedule, was based on previous works [10, 34], and the standard choices for using RoBERTa models for classification tasks. We did not tune these hyper-parameters. The mini-batch size is 16. The number of epochs is 20 for the matcher and 200 for the blocker.

Active Learning. We conduct 10 rounds of active learning, with a labeling budget of $B = 128$ samples per round. We start with an initial labeled seed set containing $|T_p| = 64$ positive and $|T_n| = 64$ negative pairs. These pairs were sampled at random from the benchmarked training splits of the datasets. All results are averaged over three such randomly constructed labeled seed sets. The default value of the candidate set size is $|CAND| = 3 \cdot |S|$ where $|S|$ denotes the size of the second list. The number of nearest neighbours retrieved is $k = 3$. The size of List S, shown in Table 1, is very small for the Abt-Buy dataset, hence we use a candidate set size of $CAND = 20 \cdot |S|$, and $k = 20$ for this dataset. We retrieve the nearest neighbours based on the ℓ_2 distance. We do not warm start the model parameters between active learning rounds, i.e. after each round M is re-initialized with the pre-trained weights of the TPLM.

Unless stated otherwise, all systems use uncertainty sampling to select examples from the candidate set. In all our experiments, we exclude the pairs in $\mathcal{D}_{test} \cap CAND$ from the process of selecting examples to query the labeler.

4.3 Methods Compared

We compare DIAL with four baseline approaches of blocking while using a TPLM-based matcher in an active learning loop:

- **PairedFixed** uses a non adaptive blocking strategy, where the candidate set is created by conducting a similarity search on the embeddings obtained as is from the pre-trained TPLM, i.e. no task specific finetuning is employed
- **PairedAdapt** uses the embeddings from the TPLM as it gets finetuned by the matcher in paired mode as described in Section 3.1. However, the candidate set is created in a similar manner as PairedFixed, i.e. a similarity search on the embeddings obtained from the TPLM in the single mode.

- **SentenceBERT** finetunes the TPLM and a SentenceBERT-like classifier on the labeled data T to obtain embeddings conducive for similarity search. To keep comparisons uniform, even though the method is called SentenceBERT, we use the same RoBERTa transformer in all methods. This method is also what is called the Advanced Blocking method in DITTO [34] except that we learn it in an Active Learning setup much like DIAL.
- **Rules** depends on hand-crafted rules to perform blocking. These exist only for the five benchmark datasets and not for the multilingual dataset. These five benchmarks already provide pairs after pre-blocking with human-designed rules, so we did not create our own rules and instead define all pairs in these pre-blocked datasets as the candidate set for this method.

All baselines use a TPLM based matcher, similar to DIAL.

We further compare DIAL with three well-established non-TPLM based methods. [40] conducted an exhaustive experimental study to compare various active learning methods for entity resolution on several real-world datasets and found that random forests with learner-aware QBC, described in Section 2.3.1, perform remarkably well. We compare DIAL with a Random Forest learner implemented as an ensemble of 20 decision trees using QBC via bootstrap [40].

JedAI [47, 51] is another recent open-source toolkit for Entity Resolution. JedAI offers highly scalable implementations of end-to-end schema-based and schema-agnostic pipelines including meta-blocking techniques. Schema-based workflows rely on similarity joins, whereas schema-agnostic workflows leverage all attribute values to extract overlapping blocks. We compare DIAL with the *best configuration* [47] of both workflows, as found through Grid Search on each dataset using the gold list of duplicates DUPs.

4.4 Overall Results

Figure 4 plots the progressive F1 scores obtained by overall system of baseline methods and DIAL on the unseen test dataset as described in Section 4.1. The x-axis denotes the increasing number of example pairs in T as active learning progresses. In Table 2 we show the efficacy of each system at the end of AL in retrieving all duplicate pairs. Here we also show precision, recall and running time. We find that DIAL provides significant gains in F1 over baselines methods both at each stage of AL and at the end of AL.

Table 2 shows that DIAL produces the best F1 scores on all the product datasets while performing close to the best on the citation datasets. With respect to recall, we note that DIAL’s recall is often close and in some cases, perhaps surprisingly, *exceeds* Rules’. The intent behind Rules was to perform blocking which in turn, calls for recall. So it is quite surprising that on datasets such as Walmart-Amazon and Abt-Buy, without any external knowledge about the domain and with only limited number of judiciously chosen labels, DIAL produces even better recall than hand-tuned blocking functions. Figure 5 provides a more detailed view of this phenomenon by showing the recall of the candidate set CAND at each stage of AL. Here we see that the recall offered by DIAL’s blocker is significantly higher than other methods. Note the recall of PairedFixed and Rules does not change since the candidate set remains fixed. In most cases PairedAdapt’s F1 is better than PairedFixed indicating

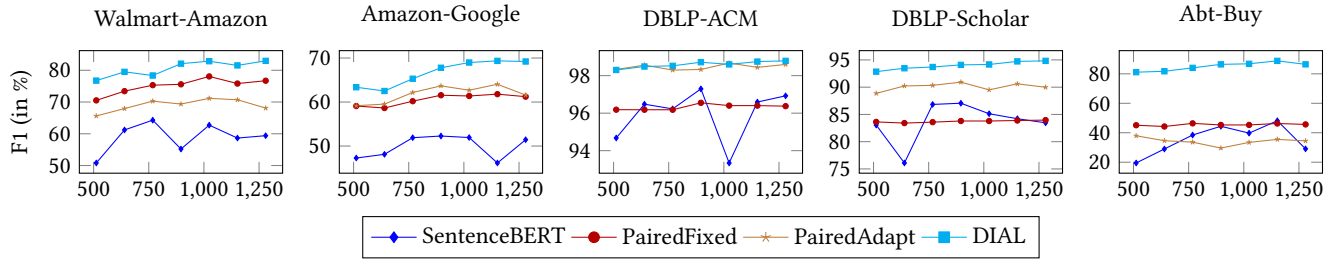


Figure 4: Comparison of DIAL with baseline approaches with respect to F1 on a fixed test-set against increasing number of instances selected by active learning. In all cases, DIAL provides significant gains over existing methods.

Table 2: Comparison of DIAL with baseline approaches with respect to Precision, Recall, and F1 evaluated on all pairs at the end of the AL loop. DIAL achieves high recall and consequently high F1 scores. The RT column denotes time in seconds to find All duplicate pairs and includes both blocking and matching time.

| Method | Walmart-Amazon | | | | Amazon-Google | | | | DBLP-ACM | | | | DBLP-Scholar | | | | Abt-Buy | | | |
|-----------------------|----------------|------|------|------|---------------|------|------|-----|----------|------|------|------|--------------|------|------|-----|---------|------|------|-----|
| | P | R | F1 | RT | P | R | F1 | RT | P | R | F1 | RT | P | R | F1 | RT | P | R | F1 | RT |
| Non TPLM based | | | | | | | | | | | | | | | | | | | | |
| Random Forest | 96.5 | 63.0 | 76.2 | 1.1 | 84.7 | 54.6 | 66.3 | 1.1 | 99.0 | 99.1 | 99.0 | 1.3 | 97.2 | 96.3 | 96.7 | 2.7 | 83.9 | 52.4 | 64.4 | 0.9 |
| JedAI:Schema-based | 82.9 | 55.2 | 66.3 | 0.5 | 66.3 | 42.3 | 51.7 | 0.5 | 97.8 | 93.2 | 95.4 | 0.6 | 95.3 | 77.5 | 85.5 | 14 | 88.4 | 43.8 | 58.5 | 0.4 |
| JedAI:Schema-agnostic | 59.0 | 75.3 | 66.2 | 5.3 | 57.6 | 64.1 | 60.7 | 4.5 | 99.3 | 99.2 | 99.3 | 1.3 | 94.6 | 94.9 | 94.7 | 30 | 94.9 | 85.6 | 90.0 | 1.1 |
| TPLM based | | | | | | | | | | | | | | | | | | | | |
| SentenceBERT | 87.1 | 43.9 | 58.0 | 87.6 | 73.2 | 38.5 | 50.4 | 7.9 | 99.3 | 94.3 | 96.7 | 15.5 | 97.0 | 74.4 | 84.2 | 255 | 87.6 | 20.3 | 32.6 | 42 |
| PairedFixed | 96.6 | 71.2 | 82.0 | 87.6 | 94.9 | 52.1 | 67.2 | 7.9 | 99.6 | 93.6 | 96.5 | 15.5 | 98.5 | 74.2 | 84.6 | 255 | 97.9 | 33.0 | 49.3 | 42 |
| PairedAdapt | 96.3 | 61.2 | 74.4 | 87.6 | 91.6 | 58.3 | 71.1 | 7.9 | 99.7 | 98.0 | 98.8 | 15.5 | 98.2 | 85.8 | 91.6 | 255 | 97.6 | 23.4 | 37.7 | 42 |
| Rules | 93.7 | 77.3 | 84.7 | 9.2 | 85.4 | 75.2 | 79.9 | 5.6 | 99.4 | 99.2 | 99.3 | 15.1 | 96.3 | 98.0 | 97.1 | 26 | 96.3 | 87.2 | 91.6 | 15 |
| DIAL | 94.9 | 85.2 | 89.8 | 88.3 | 87.4 | 77.4 | 82.1 | 8.0 | 99.6 | 98.6 | 99.1 | 15.6 | 97.5 | 96.1 | 96.8 | 257 | 97.8 | 87.4 | 92.3 | 42 |

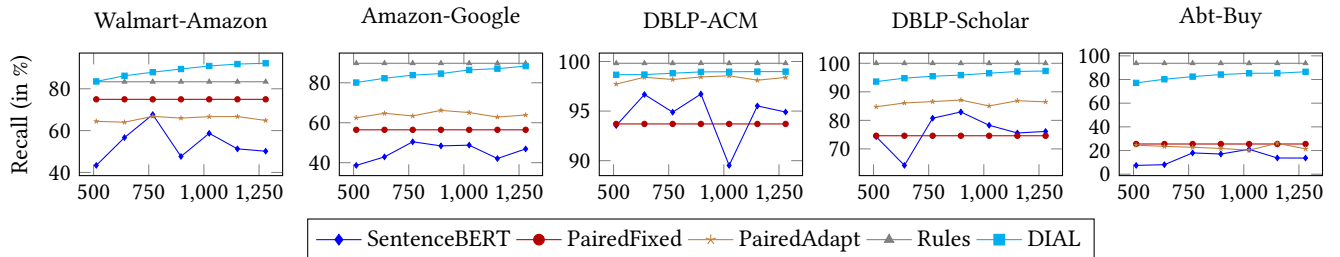


Figure 5: Recall on CAND against increasing number of instances selected by active learning. In all cases, DIAL provides significant gains over baseline methods and is able to achieve recall at par with hand crafted rule based blocking.

that fine-tuning the transformer parameters with the task specific training data T is helpful. The recall of SentenceBERT is worse than PairedAdapt perhaps because the SentenceBert network architecture, choice of training data, and training objective are not effective in co-embedding duplicates. The finding on the poor performance of SentenceBERT is significant because DITTO [34], a recent state-of-the-art ER system proposed to use SentenceBERT as its advanced blocking strategy on their large internal dataset.

In terms of running time, we observe that all deep-learning (TPLM-based) methods are between one and two orders of magnitude slower than pre-deep learning methods in the first three rows. However, given the substantial gains in accuracy that TPLM-based methods provide, an end-user may be willing to invest in the extra running time.

4.5 Multilingual Dataset

The multilingual dataset that we use is from [26]. The dataset, originally proposed for machine translation of structured data, consists

Table 3: Precision, Recall, and F1 evaluated on all pairs on the Multilingual dataset at the end of the 10 AL rounds. DIAL achieves higher almost 7.3 points higher F1 compared to existing practice of solving this task.

| Method | P | R | F1 |
|-------------|------|------|------|
| PairedFixed | 81.2 | 56.8 | 66.9 |
| PairedAdapt | 94.8 | 31.6 | 47.4 |
| DIAL | 92.2 | 62.3 | 74.3 |

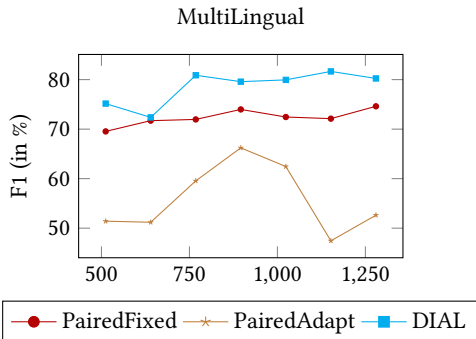


Figure 6: Comparison of DIAL with baselines on progressive F1 scores on a fixed test-set. DIAL consistently outperforms baseline methods.

of accurately-aligned parallel XML files in multiple languages. For our experiments, we use the English-Deutsch subset. Concretely, in our setup, each element of list R is a string in English which can contain HTML/XML tags, and similarly each element of list S is a string in German which can contain HTML/XML tags. As a result of the parallel alignment in data, we have $|\text{DUPS}| = |R| = |S|$.

For the multilingual dataset, we use 6 layers out of the 12 layered uncased multilingual BERT base model. This model was pre-trained on 104 languages from the Wikipedia dataset using Masked Language Modelling and Next Sentence Prediction [16]. Apart from changing the base transformer all other implementation details, including the architectures of the classifiers remain the same as that for the earlier five benchmark datasets.

We now describe the construction of the labeled seed set. We use a pre-trained 12 layered uncased multilingual BERT base model and create an index on the embedding of each $r \in R$. Then, we query $k = 3$ nearest neighbours in this index for the embedding of each $s \in S$. Using the gold list of duplicates DUPS , we divide these retrieved pairs into duplicates and non-duplicates. A random sample of 64 duplicate, and 64 non-duplicate pairs, from these sets respectively, is then chosen to create the labeled seed set. The test set is constructed in a similar manner, except the index is created, and probed, on the elements of the dev split of the dataset. The multilingual BERT model, as mentioned above, is pre-trained on 104 different languages and hence learns an extremely strong prior. Moreover, the dataset that we use consists mostly of natural language text, as opposed to the deepmatcher datasets involving product or bibliographical data. These two key differences from the

Table 4: Comparing labeled negatives with random negatives to train the committee embeddings in DIAL after 10 rounds of AL. Note the 12-25 points jump in recall with Random negatives on product datasets.

| Negatives | W-A | A-G | D-A | D-S | A-B |
|----------------------|-------|-------|-------|-------|-------|
| Recall of CAND | | | | | |
| Labeled | 80.94 | 76.54 | 99.02 | 93.47 | 66.45 |
| Random | 92.20 | 88.36 | 98.98 | 97.30 | 92.50 |
| Test Evaluation | | | | | |
| Labeled | 75.47 | 67.93 | 98.75 | 93.32 | 69.74 |
| Random | 82.97 | 69.21 | 98.79 | 94.83 | 88.81 |
| All Pairs Evaluation | | | | | |
| Labeled | 85.36 | 78.78 | 99.14 | 95.49 | 78.12 |
| Random | 89.80 | 82.07 | 99.13 | 96.81 | 92.31 |

previous setup influence the decision to fine-tune the TPLM, i.e. we find that freezing the TPLM parameters leads to slightly better F1 scores.

Progressive F1 scores calculated on the test data can be found in Figure 6. Table 3 compares DIAL against the PairedFixed and PairedAdapt baselines on All-Pairs F1 scores calculated after 10 active learning rounds. We notice that on both evaluation measures, DIAL outperforms baselines significantly. Compared to indexing the transformer embeddings as-is, DIAL achieves more than 7 percent points increase in F1!

4.6 Ablation Study

We next present a detailed ablation study to evaluate the impact of the many design decisions we made in the design of DIAL.

4.6.1 Choice of Training Data. To validate the intuition presented in Section 3.2.2, we compare DIAL, where the blocker is trained to drive apart embeddings of “easy” non-duplicates, to where the blocker is trained to separate the “difficult” labeled negatives ($T - T_p$) chosen during DIAL’s AL loop. Table 4 evaluates, on all three metrics, the two systems. We observe that Random Negatives achieves higher recall on CAND providing absolute gains of 12 – 25% over Labeled negatives on the product datasets! This subsequently results in much better F1 scores on both evaluation measures, compared to Labeled Negatives. We note here that Random Negatives while significantly improving recall of blocker, can be detrimental to precision if used to train the Matcher. On product datasets, a matcher trained with random negatives suffers a loss of 30 – 60% in precision compared to labeled negatives.

4.6.2 Choice of Training Objective. Once we have established that random negatives are more effective than labeled negatives, we evaluate the objective function to train the blocker for maximizing recall. We compare our Contrastive objective, defined in Equation 8, with two other objectives:

Classification objective as used in SentenceBert to separate duplicates from non-duplicates using cross entropy (Eq 6)

Table 5: Evaluation of DIAL with different objectives to train the committee embeddings after 10 rounds of Active Learning. Contrastive objective consistently outperforms Classification and Triplet objectives.

| Objective | W-A | A-G | D-A | D-S | A-B |
|----------------------|--------------|--------------|--------------|--------------|--------------|
| Test Evaluation | | | | | |
| Classification | 79.63 | 67.40 | 98.75 | 93.28 | 70.90 |
| Triplet | 80.94 | 68.71 | 98.79 | 94.38 | 87.21 |
| Contrastive | 82.97 | 69.21 | 98.79 | 94.83 | 88.81 |
| All Pairs Evaluation | | | | | |
| Classification | 84.88 | 79.17 | 99.05 | 95.15 | 76.03 |
| Triplet | 87.72 | 81.04 | 99.06 | 96.48 | 91.95 |
| Contrastive | 89.80 | 82.07 | 99.13 | 96.81 | 92.31 |

Table 6: Evaluation of DIAL with increasing candidate set size after 10 rounds of Active Learning.

| CAND | W-A | A-G | D-A | D-S | A-B |
|----------------------|-------|-------|-------|-------|-------|
| Recall | | | | | |
| Small | 55.78 | 79.31 | 98.98 | 92.55 | 71.92 |
| Medium | 92.20 | 88.36 | 98.98 | 97.30 | 86.54 |
| Large | 94.60 | 89.90 | 99.09 | 97.85 | 92.50 |
| All Pairs Evaluation | | | | | |
| Small | 70.19 | 80.09 | 99.08 | 95.01 | 82.68 |
| Medium | 89.80 | 82.07 | 99.13 | 96.81 | 90.49 |
| Large | 90.80 | 81.41 | 99.19 | 97.00 | 92.31 |

Triplet Used in [64] for product matching with TPLM, a triplet loss is computed on examples that are positive and negative with respect to an anchor. This loss penalizes the model if the anchor is farther away from the positive than the negative example. The TripletObjective is expressed as

$$\text{TripletObjective} = \max(d(r_p, s_p) - d(r_p, s_r) + \text{margin}, 0) \\ + \max(d(s_p, r_p) - d(s_p, r_r) + \text{margin}, 0)$$

We use the euclidean distance metric d , and set the margin to be 1. However, unlike [64], we do not perform hard negative mining.

Table 5 reports the F1 scores on Test and All pairs evaluations at the end of the active learning loop, for the three different training objectives used to train the blocker. We see that the Contrastive consistently outperforms Classification and Triplet objectives. The similarity between instance embeddings of the positive (and negative) pairs is maximized (and minimized) explicitly in contrastive and triplet objectives, whereas this is implicit in classification. The contrastive objective is able to leverage multiple random negatives as opposed to triplet which only uses 2, one for each instance as an anchor.

4.6.3 Choice of Candidate Size. The size of Candidate set |CAND|, is an important factor that influences the overall recall of the system. A small candidate set can lead to low recall, and a large candidate

Table 7: Evaluation of DIAL with increasing committee size (N) after 10 rounds of Active Learning.

| N | W-A | A-G | D-A | D-S | A-B |
|----------------------|--------------|--------------|--------------|--------------|--------------|
| Test Evaluation | | | | | |
| 1 | 83.16 | 68.62 | 98.52 | 94.38 | 88.56 |
| 3 | 82.97 | 69.21 | 98.79 | 94.83 | 88.81 |
| 5 | 83.51 | 70.85 | 98.71 | 94.76 | 88.31 |
| All Pairs Evaluation | | | | | |
| 1 | 89.85 | 80.82 | 99.20 | 96.21 | 92.22 |
| 3 | 89.80 | 82.07 | 99.13 | 96.81 | 92.31 |
| 5 | 90.19 | 82.14 | 99.10 | 96.66 | 92.79 |

set can inadvertently lead to low precision. Table 6 compares DIAL with different candidate set sizes. Small corresponds to $|\text{CAND}| = 3 \cdot |\text{DUPS}|$. Medium and Large correspond to $|\text{CAND}| = 10 \cdot |S|$ and $20 \cdot |S|$ for Abt-Buy and $|\text{CAND}| = 3 \cdot |S|$ and $5 \cdot |S|$ respectively for all other datasets. On average, All Pairs Evaluation is maximized for Large |CAND|.

4.6.4 Impact of Committee size in our blocker. Table 7 evaluates DIAL with different committee sizes N . As we motivate in Section 3.2.1, the committee is introduced to improve recall with the intuition that as opposed to one embedding it is less likely that a duplicate pair is missed by a committee of different embeddings. We find that on average, having multiple members improves performance as compared to a single member. An immediate question that then arises is, what is the cost of introducing an additional member in the committee? In Section 4.8 we provide a running time analysis varying the committee size. We show that DIAL is optimized to efficiently handle large committee sizes.

4.7 Selection Strategies

Unless stated otherwise, we have used Uncertainty Sampling as the example selection strategy for active learning. However, DIAL is agnostic to the choice of selection strategy. In this section, we compare different example selection strategies with DIAL. We implement the following methods

- **Random:** The naive baseline of choosing samples at random from the candidate set
- **Greedy:** Selecting the most similar pairs from the candidate set. We use the negative ℓ_2 distance as a similarity metric
- **Partition:** As explained in Section 2.3, High Confidence Sampling with Partition is not strictly an Active Learning selection strategy since it assumes labels not provided by a human labeler. Hence, to use a similar method as [29] in our setup, we implement two selection strategies. Partition-2 queries the user to label p_{1c} and n_{1c} , and Partition-4 queries the user to label $p_{hc}, p_{1c}, n_{hc}, n_{1c}$.
- **Query By Committee:** Select pairs from the candidate set which achieve the highest disagreement in a committee of classifiers. If member k of a committee of size N assigns a probability $\text{Pr}_k(y = 1|(r, s))$ to a pair (r, s) of being a

Table 8: Comparison of DIAL with different example selection strategies on F1 scores evaluated on all pairs after 10 rounds of active learning. DIAL is agnostic to the choice of selection strategy, and hence can operate with many different methods used in the active learning literature.

| Method | W-A | A-G | D-A | D-S | A-B |
|-------------|------|------|------|------|------|
| Random | 58.8 | 63.0 | 97.8 | 89.5 | 78.2 |
| Greedy | 78.2 | 74.9 | 90.0 | 77.9 | 79.9 |
| Partition-2 | 90.7 | 82.2 | 99.1 | 96.8 | 93.2 |
| Partition-4 | 85.4 | 74.5 | 99.0 | 95.0 | 90.6 |
| QBC | 79.1 | 75.2 | 98.8 | 94.6 | 83.9 |
| BADGE | 90.5 | 82.8 | 99.1 | 96.8 | 92.5 |
| Uncertainty | 89.8 | 82.1 | 99.1 | 96.8 | 92.3 |

duplicate pair, then the disagreement is measured as

$$H\left(\frac{1}{N}\sum_{k=1}^N \Pr(y = 1|(r, s))\right)$$

where $H(x)$ is given by Equation 4. Note that, this is a “soft” measure of disagreement, as opposed to the the hard disagreement defined in Section 2.3.

- **BADGE:** Described in Section 2.3. For a record pair $(r, s) \in \text{CAND}$, the input x is the joint encoding of r and s . The most likely label \hat{y} is calculated based on the class probabilities output by $F_W(E(x))$. The loss used to calculate the gradient embedding is the standard cross entropy loss.

Figure 7 shows All-Pair F1 at each step of active learning and Table 8 reports the All-Pair F1 scores after 10 active learning rounds on each of the 5 datasets using different selection strategies. We find that Partition-2 and BADGE provide gains over plain uncertainty sampling, as well as beat all other strategies by a high margin establishing their effectiveness for active learning.

4.8 Running time

Table 9 reports the time required by the different operations of DIAL in the 10^{th} round of active learning on all datasets. We emphasize here that matcher and committee training times are cumulative training times, i.e. we measure the time taken to train on all data labeled so far. We notice that the committee training time is comparable to matcher training time, despite the fact that the committee is trained for 10x more epochs than the matcher. The testing time of DIAL with different committee sizes is reported in Table 10. We notice that as the committee size is increased from $N = 1$ to 10, the corresponding testing time increases by less than 5% establishing the scalability of Index-By-Committee.

5 RELATED WORK

DIAL lies in the intersection of four distinct research areas: deep learning, entity resolution, active learning, and blocking for ER. In this section, we provide further details supporting the discussion presented in Section 1 and review work from the broader literature.

Table 9: Time taken, in seconds, by the different operations of DIAL in the 10^{th} round of active learning

| Operation | W-A | A-G | D-A | D-S | A-B |
|----------------------|-------|-------|-------|-------|-------|
| Train Matcher | 109.8 | 71.5 | 147.0 | 110.1 | 161.9 |
| Train Committee | 102.0 | 132.2 | 141.2 | 145.7 | 35.3 |
| Indexing & Retrieval | 1.8 | 0.4 | 0.5 | 4.8 | 0.2 |
| Selection | 73.0 | 6.0 | 8.9 | 221.9 | 34.71 |

Table 10: Testing Times (in seconds) of DIAL with different committee sizes

| Method | W-A | A-G | D-A | D-S | A-B |
|-------------------|------|-----|------|-------|------|
| DIAL ($N = 1$) | 87.6 | 7.9 | 15.5 | 254.8 | 41.8 |
| DIAL ($N = 3$) | 88.3 | 8.0 | 15.6 | 256.7 | 42.0 |
| DIAL ($N = 10$) | 90.8 | 8.2 | 15.8 | 263.1 | 42.0 |

5.1 Active Learning for Entity Resolution

Over the years, a number of works have applied active learning to ER using a variety of (paired) classifiers including support vector machines, decision trees [58, 63], explainable ER rules [3, 54, 55]. However, most of these assume that the blocking function is known. In fact, some of the aforementioned works that attempt to learn rules [54, 55] ask the user to mark every possible blocker in the input feature space. DIAL attempts to improve upon these approaches by not only *learning* the blocker but also via the use of a more powerful paired classifier (pre-trained language models). Meduri et al. [39] provide an in-depth comparison of various matchers and example selectors but neither consider TPLMs nor address how to learn a blocker. While HEALER [11] attempts to improve upon Mozafari et al. [40]’s committee-based approach to ER by including different kinds of matchers, it does not consider neural networks or TPLMs in its heterogeneous committee.

Alongside AL for ER, another line of work attempts to solve ER by crowd-sourcing [66]. The drawbacks of this approach include that no model is learned (neither matcher nor blocker) thus incurring monetary costs needed to pay the crowd each time we are faced with new data to deduplicate. Distinct from DIAL’s, their focus is more towards correcting the labels obtained from the crowd (may not constitute experts) [23] and most efficient interface to elicit most labels at least cost [65].

5.2 Deep Learning for Entity Resolution

Deep Learning has been used to tackle various aspects of the Entity Resolution task including blocking [17, 71], and matching [10, 34, 41, 64, 67, 72], and detecting variations [19] which are duplicates on a given set of base attributes but differ on other attributes. The example used in the Section 2.2.1 of a pair of records describing two different editions of the same book is an example of a variation. We refer the reader to [6] for an extensive survey on deep learning for entity matching. Deep learning methods for ER can be broadly classified into methods which operate on separate embeddings of instances r and s [17, 34, 71], and those which operate on the joint embedding of the record pair (r, s) [10, 34, 41, 43, 64]. While

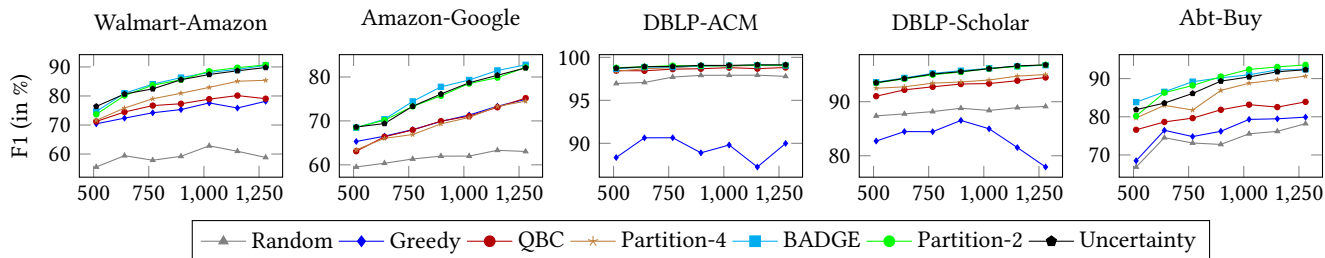


Figure 7: Comparison of DIAL with different selection methods on All Pairs Evaluation against increasing number of instances selected by active learning

joint embeddings provide more information useful for ER, DIAL shows there is a place for both, i.e., jointly embedding the pair and embedding them separately for use in blocking. DeepMatcher [41] proposed one of the first neural network architectures for ER, which was improved upon by DITTO [34]. Neither of these consider blocking nor active learning. In an effort to tackle ER in low-resource settings such as scarcely available labeled data, DTAL [29] proposes learning a neural network via active learning with uncertainty sampling along with partitioning but does not consider TPLMs and neither addresses learning a blocker. DIAL improves upon DTAL by learning an integrated matcher and blocker where the matcher is a more powerful TPLM, and DITTO’s advanced blocking in the active learning setting as shown via our experiments.

5.3 Active Learning for Deep Learning

Perhaps the closest work to our setting is [42], which also considers TPLMs for active learning on pairwise classification tasks. They use a similar architecture as [25], i.e. a TPLM invoked and trained in single mode to retrieve similar embeddings. Key differences from DIAL are 1) They do not use random negatives, 2) They do not consider separate models for matching and blocking, 3) They do not create a committee of multiple embeddings.

At the intersection of committee based methods for active learning, and deep learning, lies [7] which creates a committee of Convolutional Neural Networks (CNN) based classifiers for Active Learning in Image Classification. The area of Deep Active Learning is rapidly growing with exciting works like BALD [22], Loss Prediction [69], and Batch Aware methods like BatchBALD [30], BADGE [5] and [59, 73]. A comprehensive survey of deep active learning methods can be found in [57]. As stated earlier, most of these are compatible for use as example selectors in DIAL.

5.4 Blocking

Besides hand-coded blocking functions, earlier methods for blocking relied on unsupervised clustering [37] and passive learning with labeled data required up-front [8]. The latter uses red-blue set cover to learn an effective blocking function but its need for labeled data makes it ineffective in settings that call for active learning. While other approaches for blocking are available, a number of these utilize unsupervised learning [1, 24, 53].

Token Blocking [45] uses tokens from every attribute value as blocking keys, and records with common tokens are put in one block.

This yield high recall at the cost of low precision. Several methods have been proposed to deal with redundant and superfluous comparisons [2, 12, 27, 37, 44, 45, 48]. Meta-Blocking [46] operates on redundancy-positive block collections where the number of shared blocks indicate likelihood of matching. A blocking graph is created from the given redundancy-positive block collection, and is pruned using matching likelihoods [14, 18, 46, 48, 49, 62, 70].

AutoBlock [71] assumes knowledge of strong attributes (e.g., UPC code for grocery products), that may be used to produce labeled data for learning the blocking function. DIAL does not make any such assumptions and can work with heterogeneous lists. Both AutoBlock and DeepER [17] use Locality Sensitive Hashing (LSH) for retrieval, and DITTO uses similarity search by blocked matrix multiplication [1]. In contrast, DIAL uses FAISS [28], a highly optimized library for k-selection which relies on product quantization for fast asymmetric distance computations.

Another related task is training a retrieval system for entity linkage [25]. Key similarities with our blocker model include fine-tuning the TPLM in the single mode, and using random negatives to train the TPLM. This work differs from our work in that it does not perform active learning. We refer the reader to [12, 13, 33, 50] for an extensive survey on blocking.

6 CONCLUSION

In this work we present DIAL, a scalable active learning system with an integrated matcher-blocker combination. As opposed to most works in ER, DIAL learns the blocker in addition to the matcher. Furthermore, the blocker and matcher are integrated in a way so that improvements in one can benefit the other. We show that our approach leads to improved recall during blocking and improved matching via the use of transformer-based pre-trained language models. We successfully train a committee on top of powerful TPLM representations, and use it to perform Index-by-Committee, a novel and efficient example retrieval technique. Our experimental results on 5 real world datasets show that DIAL outperforms baseline methods by a large margin while also requiring minimal human involvement. We showcase our approach by reporting the effectiveness of DIAL on a multilingual dataset where hand-coding a blocking function may not be possible due to the different languages involved.

REFERENCES

- [1] Firas Abuzaid, Geet Sethi, Peter Bailis, and Matei Zaharia. 2019. To Index or Not to Index: Optimizing Exact Maximum Inner Product Search. In *35th IEEE International Conference on Data Engineering, ICDE 2019, Macao, China, April 8-11, 2019*. IEEE, 1250–1261. <https://doi.org/10.1109/ICDE.2019.00114>
- [2] Akiko N. Aizawa and Keizo Oyama. 2005. A Fast Linkage Detection Scheme for Multi-Source Information Integration. In *2005 International Workshop on Challenges in Web Information Retrieval and Integration (WIRI 2005), 8-9 April 2005, Tokyo, Japan*. IEEE Computer Society, 30–39. <https://doi.org/10.1109/WIRI.2005.2>
- [3] Arvind Arasu, Michaela Götz, and Raghav Kaushik. 2010. On Active Learning of Record Matching Packages. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data (Indianapolis, Indiana, USA) (SIGMOD '10)*. Association for Computing Machinery, New York, NY, USA, 783–794. <https://doi.org/10.1145/1807167.1807252>
- [4] David Arthur and Sergei Vassilvitskii. 2007. K-Means++: The Advantages of Careful Seeding. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms (New Orleans, Louisiana) (SODA '07)*. Society for Industrial and Applied Mathematics, USA, 1027–1035.
- [5] Jordan T. Ash, Chicheng Zhang, Akshay Krishnamurthy, John Langford, and Alekh Agarwal. 2020. Deep Batch Active Learning by Diverse, Uncertain Gradient Lower Bounds. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net. <https://openreview.net/forum?id=ryghZJBKPS>
- [6] Nils Barlaug and Jon Atle Gulla. 2021. Neural Networks for Entity Matching: A Survey. *ACM Trans. Knowl. Discov. Data* 15, 3 (2021), 52:1–52:37. <https://doi.org/10.1145/3442200>
- [7] William H. Beluch, Tim Genewein, Andreas Nürnberger, and Jan M. Köhler. 2018. The Power of Ensembles for Active Learning in Image Classification. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*. Computer Vision Foundation / IEEE Computer Society, 9368–9377. <https://doi.org/10.1109/CVPR.2018.00976>
- [8] Mikhail Bilenko, Beena Kamath, and Raymond J. Mooney. 2006. Adaptive Blocking: Learning to Scale Up Record Linkage. In *Proceedings of the 6th IEEE International Conference on Data Mining (ICDM 2006), 18-22 December 2006, Hong Kong, China*. IEEE Computer Society, 87–96. <https://doi.org/10.1109/ICDM.2006.13>
- [9] Leo Breiman. 2001. Random Forests. *Mach. Learn.* 45, 1 (Oct. 2001), 5–32. <https://doi.org/10.1023/A:1010933404324>
- [10] Ursin Brunner and Kurt Stockinger. 2020. Entity Matching with Transformer Architectures - A Step Forward in Data Integration. In *Proceedings of the 23rd International Conference on Extending Database Technology, EDBT 2020, Copenhagen, Denmark, March 30 - April 02, 2020*. Angela Bonifati, Yongluan Zhou, Marcos Antonio Vaz Salles, Alexander Böhm, Dan Olteanu, George H. L. Fletcher, Arijit Khan, and Bin Yang (Eds.). OpenProceedings.org, 463–473. <https://doi.org/10.5441/002/edbt.2020.58>
- [11] Xiao Chen, Yinlong Xu, David Briones, Gabriel Campero Durand, Roman Zoun, and Gunter Saake. 2019. Heterogeneous Committee-Based Active Learning for Entity Resolution (HeALER). In *Advances in Databases and Information Systems, Tatjana Welzer, Johann Eder, Vili Podgorelec, and Aida Kamišalić Latifić (Eds.)*. Springer International Publishing, Cham, 69–85.
- [12] P. Christen. 2012. A Survey of Indexing Techniques for Scalable Record Linkage and Deduplication. *IEEE Transactions on Knowledge and Data Engineering* 24, 9 (2012), 1537–1555. <https://doi.org/10.1109/TKDE.2011.127>
- [13] Vassilis Christophides, Vasilis Efthymiou, Themis Palpanas, George Papadakis, and Kostas Stefanidis. 2020. An Overview of End-to-End Entity Resolution for Big Data. *ACM Comput. Surv.* 53, 6, Article 127 (Dec. 2020), 42 pages. <https://doi.org/10.1145/3418896>
- [14] Guilherme dal Bianco, Marcos André Gonçalves, and Denio Duarte. 2018. BLOSS: Effective meta-blocking with almost no effort. *Information Systems* 75 (2018), 75–89. <https://doi.org/10.1016/j.is.2018.02.005>
- [15] Sanjib Das, AnHai Doan, Paul Suganthan G. C., Chaitanya Gokhale, Pradap Konda, Yash Govind, and Derek Paulsen. 2021. The Magellan Data Repository. <https://sites.google.com/site/anhaidgroup/useful-stuff/the-magellan-data-repository>.
- [16] 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)*. Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186. <https://doi.org/10.18653/v1/N19-1423>
- [17] Muhammad Ebraheem, Saravanan Thirumuruganathan, Shafiq Joty, Mourad Ouzzani, and Nan Tang. 2018. Distributed Representations of Tuples for Entity Resolution. *Proc. VLDB Endow.* 11, 11 (July 2018), 1454–1467. <https://doi.org/10.14778/3236187.3236198>
- [18] Vasilis Efthymiou, George Papadakis, Kostas Stefanidis, and Vassilis Christophides. 2019. MinoanER: Schema-Agnostic, Non-Iterative, Massively Parallel Resolution of Web Entities. In *Advances in Database Technology - 22nd International Conference on Extending Database Technology, EDBT 2019, Lisbon, Portugal, March 26-29, 2019*, Melanie Herschel, Helena Galhardas, Berthold Reinwald, Irini Fundulaki, Carsten Binnig, and Zoi Kaoudi (Eds.). OpenProceedings.org, 373–384. <https://doi.org/10.5441/002/edbt.2019.33>
- [19] Varun Embar, Bunyamin Sisman, Hao Wei, Xin Luna Dong, Christos Faloutsos, and Lise Getoor. 2020. Contrastive Entity Linkage: Mining Variational Attributes from Large Catalogs for Entity Linkage. In *Conference on Automated Knowledge Base Construction, AKBC 2020, Virtual, June 22-24, 2020*, Dipanjan Das, Hannaneh Hajishirzi, Andrew McCallum, and Sameer Singh (Eds.). <https://doi.org/10.24432/C5W88R>
- [20] Ivan P. Fellegi and Alan B. Sunter. 1969. A Theory for Record Linkage. *J. Amer. Statist. Assoc.* 64, 328 (1969), 1183–1210. <https://doi.org/10.1080/01621459.1969.10501049> arXiv:<https://www.tandfonline.com/doi/pdf/10.1080/01621459.1969.10501049>
- [21] Yoav Freund, H. Sebastian Seung, Eli Shamir, and Naftali Tishby. 1997. Selective Sampling Using the Query by Committee Algorithm. *Mach. Learn.* 28, 2-3 (1997), 133–168. <https://doi.org/10.1023/A:1007330508534>
- [22] Yarín Gal, Riashat Islam, and Zoubin Ghahramani. 2017. Deep Bayesian Active Learning with Image Data. In *Proceedings of the 34th International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Doina Precup and Yee Whye Teh (Eds.), Vol. 70. PMLR, 1183–1192. <http://proceedings.mlr.press/v70/gal17a.html>
- [23] Sainyam Galhotra, Donatella Firmani, Barna Saha, and Divesh Srivastava. 2018. Robust Entity Resolution Using Random Graphs. In *Proceedings of the 2018 International Conference on Management of Data (Houston, TX, USA) (SIGMOD '18)*. Association for Computing Machinery, New York, NY, USA, 3–18. <https://doi.org/10.1145/3183713.3183755>
- [24] Sainyam Galhotra, Donatella Firmani, Barna Saha, and Divesh Srivastava. 2021. Efficient and effective ER with progressive blocking. *The VLDB Journal* 30, 4 (Mar 2021), 537–557. <https://doi.org/10.1007/s00778-021-00656-7>
- [25] Daniel Gillick, Sayali Kulkarni, Larry Lansing, Alessandro Presta, Jason Baldrige, Eugene Ie, and Diego Garcia-Olano. 2019. Learning Dense Representations for Entity Retrieval. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*. Association for Computational Linguistics, Hong Kong, China, 528–537. <https://doi.org/10.18653/v1/K19-1049>
- [26] Kazuma Hashimoto, Raffaella Buschiazio, James Bradbury, Teresa Marshall, Richard Socher, and Caiming Xiong. 2019. A High-Quality Multilingual Dataset for Structured Documentation Translation. In *Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers)*. Association for Computational Linguistics, Florence, Italy, 116–127. <https://doi.org/10.18653/v1/W19-5212>
- [27] Mauricio A. Hernández and Salvatore J. Stolfo. 1995. The Merge/Purge Problem for Large Databases. In *Proceedings of the 1995 ACM SIGMOD International Conference on Management of Data (San Jose, California, USA) (SIGMOD '95)*. Association for Computing Machinery, New York, NY, USA, 127–138. <https://doi.org/10.1145/223784.223807>
- [28] Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2021. Billion-Scale Similarity Search with GPUs. *IEEE Trans. Big Data* 7, 3 (2021), 535–547. <https://doi.org/10.1109/TBDATA.2019.2921572>
- [29] Jungo Kasai, Kun Qian, Sairam Gurajada, Yunyao Li, and Lucian Popa. 2019. Low-resource Deep Entity Resolution with Transfer and Active Learning. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28-August 2, 2019, Volume 1: Long Papers*, Anna Korhonen, David R. Traum, and Lluís Màrquez (Eds.). Association for Computational Linguistics, 5851–5861. <https://doi.org/10.18653/v1/p19-1586>
- [30] Andreas Kirsch, Joost van Amersfoort, and Yarín Gal. 2019. BatchBALD: Efficient and Diverse Batch Acquisition for Deep Bayesian Active Learning. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett (Eds.). 7024–7035. <https://proceedings.neurips.cc/paper/2019/hash/95323660ed2124450caaac2c46b5ed90-Abstract.html>
- [31] Pradap Konda, Sanjib Das, Paul Suganthan G. C., AnHai Doan, Adel Ardalan, Jeffrey R. Ballard, Han Li, Fatemah Panahi, Haojun Zhang, Jeff Naughton, Shishir Prasad, Ganesh Krishnan, Rohit Deep, and Vijay Raghavendra. 2016. Magellan: Toward Building Entity Matching Management Systems. *Proc. VLDB Endow.* 9, 12 (Aug. 2016), 1197–1208. <https://doi.org/10.14778/2994509.2994535>
- [32] Hanna Köpcke, Andreas Thor, and Erhard Rahm. 2010. Evaluation of Entity Resolution Approaches on Real-World Match Problems. *Proc. VLDB Endow.* 3, 1-2 (Sept. 2010), 484–493. <https://doi.org/10.14778/1920841.1920904>
- [33] Hanna Köpcke and Erhard Rahm. 2010. Frameworks for entity matching: A comparison. *Data & Knowledge Engineering* 69, 2 (2010), 197–210. <https://doi.org/10.1016/j.datak.2009.10.003>
- [34] Yuliang Li, Jinfeng Li, Yoshihiko Suhara, AnHai Doan, and Wang-Chiew Tan. 2020. Deep Entity Matching with Pre-Trained Language Models. *Proc. VLDB Endow.* 14, 1 (Sept. 2020), 50–60. <https://doi.org/10.14778/3421424.3421431>
- [35] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *CoRR abs/1907.11692* (2019). arXiv:1907.11692 <http://arxiv.org/abs/1907.11692>

- [36] Ilya Loshchilov and Frank Hutter. 2019. Decoupled Weight Decay Regularization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net. <https://openreview.net/forum?id=Bkg6RiCqY7>
- [37] Andrew McCallum, Kamal Nigam, and Lyle H. Ungar. 2000. Efficient Clustering of High-Dimensional Data Sets with Application to Reference Matching. In *Proceedings of the Sixth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Boston, Massachusetts, USA) (KDD '00). Association for Computing Machinery, New York, NY, USA, 169–178. <https://doi.org/10.1145/347090.347123>
- [38] Paul McNamee, James Mayfield, Dawn Lawrie, Douglas Oard, and David Doermann. 2011. Cross-Language Entity Linking. In *Proceedings of 5th International Joint Conference on Natural Language Processing*. Asian Federation of Natural Language Processing, Chiang Mai, Thailand, 255–263. <https://www.aclweb.org/anthology/I11-1029>
- [39] Venkata Vamsikrishna Meduri, Lucian Popa, Prithviraj Sen, and Mohamed Sarwat. 2020. A Comprehensive Benchmark Framework for Active Learning Methods in Entity Matching. In *Proceedings of the 2020 International Conference on Management of Data, SIGMOD Conference 2020, online conference [Portland, OR, USA], June 14-19, 2020*, David Maier, Rachel Pottinger, AnHai Doan, Wang-Chiew Tan, Abdussalam Alawini, and Hung Q. Ngo (Eds.). ACM, 1133–1147. <https://doi.org/10.1145/3318464.3380597>
- [40] Barzan Mozafari, Purna Sarkar, Michael Franklin, Michael Jordan, and Samuel Madden. 2014. Scaling up Crowd-Sourcing to Very Large Datasets: A Case for Active Learning. *Proc. VLDB Endow.* 8, 2 (Oct. 2014), 125–136. <https://doi.org/10.14778/2735471.2735474>
- [41] Sidharth Mudgal, Han Li, Theodoros Rekatsinas, AnHai Doan, Youngchoon Park, Ganesh Krishnan, Rohit Deep, Esteban Arcaute, and Vijay Raghavendra. 2018. Deep Learning for Entity Matching: A Design Space Exploration. In *Proceedings of the 2018 International Conference on Management of Data* (Houston, TX, USA) (SIGMOD '18). Association for Computing Machinery, New York, NY, USA, 19–34. <https://doi.org/10.1145/3183713.3196926>
- [42] Stephen Mussmann, Robin Jia, and Percy Liang. 2020. On the Importance of Adaptive Data Collection for Extremely Imbalanced Pairwise Tasks. In *Findings of the Association for Computational Linguistics: EMNLP 2020*. Association for Computational Linguistics, Online, 3400–3413. <https://doi.org/10.18653/v1/2020.findings-emnlp.305>
- [43] Hao Nie, Xianpei Han, Ben He, Le Sun, Bo Chen, Wei Zhang, Suhui Wu, and Hao Kong. 2019. Deep Sequence-to-Sequence Entity Matching for Heterogeneous Entity Resolution. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management* (Beijing, China) (CIKM '19). Association for Computing Machinery, New York, NY, USA, 629–638. <https://doi.org/10.1145/3357384.3358018>
- [44] George Papadakis, George Alexiou, George Papastefanatos, and Georgia Koutrika. 2015. Schema-Agnostic vs Schema-Based Configurations for Blocking Methods on Homogeneous Data. *Proc. VLDB Endow.* 9, 4 (Dec. 2015), 312–323. <https://doi.org/10.14778/2856318.2856326>
- [45] George Papadakis, Ekaterini Ioannou, Themis Palpanas, Claudia Niederée, and Wolfgang Nejdl. 2013. A Blocking Framework for Entity Resolution in Highly Heterogeneous Information Spaces. *IEEE Transactions on Knowledge and Data Engineering* 25, 12 (2013), 2665–2682. <https://doi.org/10.1109/TKDE.2012.150>
- [46] George Papadakis, Georgia Koutrika, Themis Palpanas, and Wolfgang Nejdl. 2014. Meta-Blocking: Taking Entity Resolution to the Next Level. *IEEE Transactions on Knowledge and Data Engineering* 26, 8 (2014), 1946–1960. <https://doi.org/10.1109/TKDE.2013.54>
- [47] George Papadakis, George Mandilaras, Luca Gagliardelli, Giovanni Simonini, Emmanouil Thanos, George Giannakopoulos, Sonia Bergamaschi, Themis Palpanas, and Manolis Koubarakis. 2020. Three-dimensional Entity Resolution with JedAI. *Information Systems* 93 (2020), 101565. <https://doi.org/10.1016/j.is.2020.101565>
- [48] George Papadakis, George Papastefanatos, and Georgia Koutrika. 2014. Supervised Meta-Blocking. *Proc. VLDB Endow.* 7, 14 (Oct. 2014), 1929–1940. <https://doi.org/10.14778/2733085.2733098>
- [49] George Papadakis, George Papastefanatos, Themis Palpanas, and Manolis Koubarakis. 2016. Scaling Entity Resolution to Large, Heterogeneous Data with Enhanced Meta-blocking. In *Proceedings of the 19th International Conference on Extending Database Technology, EDBT 2016, Bordeaux, France, March 15-16, 2016, Bordeaux, France, March 15-16, 2016*, Evaggelia Pitoura, Sofian Maabout, Georgia Koutrika, Amélie Marian, Letizia Tanca, Ioana Manolescu, and Kostas Stefanidis (Eds.). OpenProceedings.org, 221–232. <https://doi.org/10.5441/002/edbt.2016.22>
- [50] George Papadakis, Dimitrios Skoutas, Emmanouil Thanos, and Themis Palpanas. 2020. Blocking and Filtering Techniques for Entity Resolution: A Survey. *ACM Comput. Surv.* 53, 2, Article 31 (March 2020), 42 pages. <https://doi.org/10.1145/3377455>
- [51] George Papadakis, Leonidas Tsekouras, Emmanouil Thanos, George Giannakopoulos, Themis Palpanas, and Manolis Koubarakis. 2018. The Return of jedAI: End-to-end Entity Resolution for Structured and Semi-structured Data. *PVLDB* 11, 12 (2018), 1950–1953. <https://doi.org/10.14778/3229863.3236232>
- [52] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems* 32, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.). Curran Associates, Inc., 8024–8035. <http://papers.nips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>
- [53] Anna Primpeli, Christian Bizer, and Margret Keuper. 2020. Unsupervised Bootstrapping of Active Learning for Entity Resolution. In *The Semantic Web*, Andreas Harth, Sabrina Kirrane, Axel-Cyrille Ngonga Ngomo, Heiko Paulheim, Anisa Rula, Anna Lisa Gentile, Peter Haase, and Michael Cochez (Eds.). Springer International Publishing, Cham, 215–231.
- [54] Kun Qian, Lucian Popa, and Prithviraj Sen. 2017. Active Learning for Large-Scale Entity Resolution. In *Proceedings of the 2017 ACM Conference on Information and Knowledge Management, CIKM 2017, Singapore, November 06 - 10, 2017*, Ee-Peng Lim, Marianne Winslett, Mark Sanderson, Ada Wai-Chee Fu, Jimeng Sun, J. Shane Culpepper, Eric Lo, Joyce C. Ho, Debora Donato, Rakesh Agrawal, Yu Zheng, Carlos Castillo, Aixun Sun, Vincent S. Tseng, and Chenliang Li (Eds.). ACM, 1379–1388. <https://doi.org/10.1145/3132847.3132949>
- [55] Kun Qian, Lucian Popa, and Prithviraj Sen. 2019. SystemER: A Human-in-the-loop System for Explainable Entity Resolution. *Proc. VLDB Endow.* 12, 12 (2019), 1794–1797. <https://doi.org/10.14778/3352063.3352068>
- [56] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. 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 2019, Hong Kong, China, November 3-7, 2019*, Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (Eds.). Association for Computational Linguistics, 3980–3990. <https://doi.org/10.18653/v1/D19-1410>
- [57] Pengzhen Ren, Yun Xiao, Xiaojun Chang, Po-Yao Huang, Zhihui Li, Xiaojiang Chen, and Xin Wang. 2020. A Survey of Deep Active Learning. *CoRR* abs/2009.00236 (2020). arXiv:2009.00236 <https://arxiv.org/abs/2009.00236>
- [58] Sunita Sarawagi and Anuradha Bhamidipaty. 2002. Interactive Deduplication Using Active Learning. In *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Edmonton, Alberta, Canada) (KDD '02). Association for Computing Machinery, New York, NY, USA, 269–278. <https://doi.org/10.1145/775047.775087>
- [59] Ozan Sener and Silvio Savarese. 2018. Active Learning for Convolutional Neural Networks: A Core-Set Approach. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net. <https://openreview.net/forum?id=H1a1uk-RW>
- [60] Burr Settles. 2009. *Active Learning Literature Survey*. Technical Report.
- [61] H. Sebastian Seung, Manfred Opper, and Haim Sompolinsky. 1992. Query by Committee. In *Proceedings of the Fifth Annual ACM Conference on Computational Learning Theory, COLT 1992, Pittsburgh, PA, USA, July 27-29, 1992*, David Haussler (Ed.). ACM, 287–294. <https://doi.org/10.1145/130385.130417>
- [62] Giovanni Simonini, Sonia Bergamaschi, and H. V. Jagadish. 2016. BLAST: A Loosely Schema-Aware Meta-Blocking Approach for Entity Resolution. *Proc. VLDB Endow.* 9, 12 (Aug. 2016), 1173–1184. <https://doi.org/10.14778/2994509.2994533>
- [63] Sheila Tejada, Craig A Knoblock, and Steven Minton. 2001. Learning object identification rules for information integration. *Information Systems* 26, 8 (2001), 607–633.
- [64] Janusz Tracz, Piotr Iwo Wójcik, Kalina Jasinska-Kobus, Riccardo Belluzzo, Robert Mroczkowski, and Ireneusz Gawlik. 2020. BERT-based similarity learning for product matching. In *Proceedings of Workshop on Natural Language Processing in E-Commerce*. Association for Computational Linguistics, Barcelona, Spain, 66–75. <https://www.aclweb.org/anthology/2020.ecomnlp-1.7>
- [65] Vasilis Verroios, Hector Garcia-Molina, and Yannis Papakonstantinou. 2017. Waldo: An Adaptive Human Interface for Crowd Entity Resolution. In *Proceedings of the 2017 ACM International Conference on Management of Data* (Chicago, Illinois, USA) (SIGMOD '17). Association for Computing Machinery, New York, NY, USA, 1133–1148. <https://doi.org/10.1145/3035918.3035931>
- [66] Jiannan Wang, Tim Kraska, Michael J. Franklin, and Jianhua Feng. 2012. CrowdER: Crowdsourcing Entity Resolution. *Proc. VLDB Endow.* 5, 11 (July 2012), 1483–1494. <https://doi.org/10.14778/2350229.2350263>
- [67] Zhengyang Wang, Bunyamin Sisman, Hao Wei, Xin Luna Dong, and Shuiwang Ji. 2020. CorDEL: A Contrastive Deep Learning Approach for Entity Linkage. In *20th IEEE International Conference on Data Mining, ICDM 2020, Sorrento, Italy, November 17-20, 2020*, Claudia Plant, Haixun Wang, Alfredo Cuzzocrea, Carlo Zaniolo, and Xindong Wu (Eds.). IEEE, 1322–1327. <https://doi.org/10.1109/ICDM50108.2020.00171>
- [68] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu,

- Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. Transformers: State-of-the-Art Natural Language Processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Association for Computational Linguistics, Online, 38–45. <https://doi.org/10.18653/v1/2020.emnlp-demos.6>
- [69] Donggeun Yoo and In So Kweon. 2019. Learning Loss for Active Learning. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*. Computer Vision Foundation / IEEE, 93–102. <https://doi.org/10.1109/CVPR.2019.00018>
- [70] Fulin Zhang, Zhipeng Gao, and Kun Niu. 2017. A pruning algorithm for Meta-blocking based on cumulative weight. *Journal of Physics: Conference Series* 887 (aug 2017), 012058. <https://doi.org/10.1088/1742-6596/887/1/012058>
- [71] Wei Zhang, Hao Wei, Bunyamin Sisman, Xin Luna Dong, Christos Faloutsos, and Davd Page. 2020. AutoBlock: A Hands-off Blocking Framework for Entity Matching. In *Proceedings of the 13th International Conference on Web Search and Data Mining (Houston, TX, USA) (WSDM '20)*. Association for Computing Machinery, New York, NY, USA, 744–752. <https://doi.org/10.1145/3336191.3371813>
- [72] Chen Zhao and Yeye He. 2019. Auto-EM: End-to-End Fuzzy Entity-Matching Using Pre-Trained Deep Models and Transfer Learning. In *The World Wide Web Conference (San Francisco, CA, USA) (WWW '19)*. Association for Computing Machinery, New York, NY, USA, 2413–2424. <https://doi.org/10.1145/3308558.3313578>
- [73] Fedor Zhdanov. 2019. Diverse mini-batch Active Learning. *CoRR* abs/1901.05954 (2019). arXiv:1901.05954 <http://arxiv.org/abs/1901.05954>