

# Low-resource Deep Entity Resolution with Transfer and Active Learning

Jungo Kasai<sup>♡\*</sup>      Kun Qian<sup>♣</sup>      Sairam Gurajada<sup>♣</sup>  
Yunyaoli<sup>♡</sup>      Lucian Popa<sup>♣</sup>

<sup>♡</sup>Paul G. Allen School of Computer Science & Engineering, University of Washington

<sup>♣</sup>IBM Research – Almaden

`jkasai@cs.washington.edu`

`{qian.kun, Sairam.Gurajada}@ibm.com`

`{yunyaoli, lpopa}@us.ibm.com`

## Abstract

Entity resolution (ER) is the task of identifying different representations of the same real-world entities across databases. It is a key step for knowledge base creation and text mining. Recent adaptation of deep learning methods for ER mitigates the need for dataset-specific feature engineering by constructing distributed representations of entity records. While these methods achieve state-of-the-art performance over benchmark data, they require large amounts of labeled data, which are typically unavailable in realistic ER applications. In this paper, we develop a deep learning-based method that targets low-resource settings for ER through a novel combination of transfer learning and active learning. We design an architecture that allows us to learn a transferable model from a high-resource setting to a low-resource one. To further adapt to the target dataset, we incorporate active learning that carefully selects a few informative examples to fine-tune the transferred model. Empirical evaluation demonstrates that our method achieves comparable, if not better, performance compared to state-of-the-art learning-based methods while using an order of magnitude fewer labels.

## 1 Introduction

Entity Resolution (ER), also known as entity matching, record linkage (Fellegi and Sunter, 1969), reference reconciliation (Dong et al., 2005), and merge-purge (Hernández and Stolfo, 1995), identifies and links different representations of the same real-world entities. ER yields a unified and consistent view of data and serves as a crucial step in downstream applications, including knowledge base creation, text mining (Zhao et al., 2014), and social media analysis (Campbell

et al., 2016). For instance, seen in Table 1 are citation data records from two databases, DBLP and Google Scholar. If one intends to build a system that analyzes citation networks of publications, it is essential to recognize publication overlaps across the databases and to integrate the data records (Pasula et al., 2002).

Recent work demonstrated that deep learning (DL) models with distributed representations of words are viable alternatives to other machine learning algorithms, including support vector machines and decision trees, for performing ER (Ebraheem et al., 2018; Mudgal et al., 2018). The DL models provide a universal solution to ER across all kinds of datasets that alleviates the necessity of expensive feature engineering, in which a human designer explicitly defines matching functions for every single ER scenario. However, DL is well known to be data hungry; in fact, the DL models proposed in Ebraheem et al. (2018); Mudgal et al. (2018) achieve state-of-the-art performance by learning from thousands of labels.<sup>1</sup> Unfortunately, realistic ER tasks have limited access to labeled data and would require substantial labeling effort upfront, before the actual learning of the ER models. Creating a representative training set is especially challenging in ER problems due to the data distribution, which is heavily skewed towards negative pairs (i.e. non-matches) as opposed to positive pairs (i.e. matches).

This problem limits the applicability of DL methods in low-resource ER scenarios. Indeed, we will show in a later section that the performance of DL models degrades significantly as compared to other machine learning algorithms when only a limited amount of labeled data is available. To address this issue, we propose a DL-based method that combines transfer learning and

\*Work done during summer internship at IBM Research – Almaden.

<sup>1</sup>17k labels were used for the DBLP-Scholar scenario.

DBLP			
Authors	Title	Venue	Year
M Carey, D Dewitt, J Naughton, M Asgarian, P Brown, J Gehrke, D Shah	The Bucky Object-relational Benchmark (Experience Paper)	SIGMOD Conference	1997
A Netz, S Chaudhuri, J Bernhardt, U Fayyad	Integration of Data Mining with Database Technology	VLDB	2000
Google Scholar			
Authors	Title	Venue	Year
MJ Carey, DJ Dewitt, JF Naughton, M Asgarian, P	The Bucky Object Relational Benchmark	Proceedings of the SIGMOD Conference on Management of Data	NULL
A Netz, S Chaudhuri, J Bernhardt, U Fayyad	Integration of Data Mining and Relational Databases	Proc.	2000

Table 1: Data record examples from DBLP-Scholar (citation genre). The first records from DBLP and Google Scholar (red) refer to the same publication even though the information is not identical. The second ones (blue and brown) record different papers with the same authors and year.

active learning. We first develop a transfer learning methodology to leverage a few pre-existing scenarios with abundant labeled data, in order to use them in other settings of similar nature but with limited or no labeled data. More concretely, through a carefully crafted neural network architecture, we learn a transferable model from multiple source datasets with cumulatively abundant labeled data. Then we use active learning to identify informative examples from the target dataset to further adapt the transferred model to the target setting. This novel combination of transfer and active learning in ER settings enables us to learn a comparable or better performing DL model while using significantly fewer target dataset labels in comparison to state-of-the-art DL and even non-DL models. We also note that the two techniques are not dependent on each other. For example, one could skip transfer learning if no high-resource dataset is available and directly use active learning. Conversely, one could use transfer learning directly without active learning. We evaluate these cases in the experiments. Specifically, we make the following contributions:

- We propose a DL architecture for ER that learns attribute agnostic and transferable representations from multiple source datasets using dataset (domain) adaptation.
- To the best of our knowledge, we are the first to design an active learning algorithm for deep ER models. Our active learning algorithm searches for high-confidence examples and uncertain examples, which provide a guided way to improve the precision and recall of the transferred model to the target dataset.
- We perform extensive empirical evaluations over multiple benchmark datasets and demonstrate that our method outperforms state-of-

the-art learning-based models while using an order of magnitude fewer labels.

## 2 Background and Related Work

### 2.1 Entity Resolution

Let  $D_1$  and  $D_2$  be two collections of entity records. The task of ER is to classify the entity record pair  $\langle e_1, e_2 \rangle$ ,  $\forall e_1 \in D_1, e_2 \in D_2$ , into a match or a non-match. This is accomplished by comparing entity record  $e_1$  to  $e_2$  on their corresponding attributes. In this paper, we assume records in  $D_1$  and  $D_2$  share the same schema (set of attributes). In cases where they have different attributes, one can use schema matching techniques (Rahm and Bernstein, 2001) to first align the schemas, followed by data exchange techniques (Fagin et al., 2009). Each attribute value is a sequence of words. Table 1 shows examples of data records from an ER scenario, DBLP-Scholar (Köpcke et al., 2010) from the citation genre and clearly depicts our assumption of datasets handled in this paper.

Since the entire Cartesian product  $D_1 \times D_2$  often becomes large and it is infeasible to run a high-recall classifier directly, we typically decompose the problem into two steps: *blocking* and *matching*. *Blocking* filters out obvious non-matches from the Cartesian product to obtain a candidate set. Attribute-level or record-level tf-idf and jaccard similarity can be used for blocking criteria. For example, in the DBLP-Scholar scenario, one blocking condition could be based on applying equality on “Year”. Hence, two publications in different years will be considered as obvious non-matches and filtered out from the candidate set. Then, the subsequent *matching* phase classifies the candidate set into matches and non-matches.

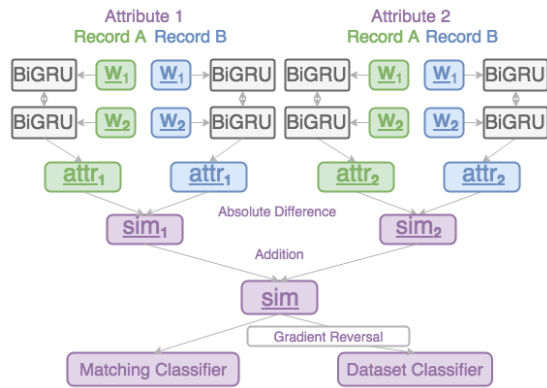


Figure 1: Deep ER model architecture with dataset adaptation via gradient reversal. Only two attributes are shown.  $W$ s indicate word vectors.

## 2.2 Learning-based Entity Resolution

As described above, after the blocking step, ER reduces to a binary classification task on candidate pairs of data records. Prior work has proposed learning-based methods that train classifiers on training data, such as support vector machines, naive bayes, and decision trees (Christen, 2008; Bilenko and Mooney, 2003). These learning-based methods first extract features for each record pair from the candidate set across attributes in the schema, and use them to train a binary classifier. The process of selecting appropriate classification features is often called feature engineering and it involves substantial human effort in each ER scenario. Recently, Ebraheem et al. (2018) and Mudgal et al. (2018) have proposed deep learning models that use distributed representations of entity record pairs for classification. These models benefit from distributed representations of words and learn complex features automatically without the need for dataset-specific feature engineering.

## 3 Deep ER Model Architecture

We describe the architecture of our DL model that classifies each record pair in the candidate set into a match or a non-match. As shown in Fig. 1, our model encompasses a sequence of steps that computes attribute representations, attribute similarity and finally the record similarity for each input pair  $\langle e_1, e_2 \rangle$ . A matching classifier uses the record similarity representation to classify the pair. For an extensive list of hyperparameters and training details we chose, see the appendix.

**Input Representations.** For each entity record pair  $\langle e_1, e_2 \rangle$ , we tokenize the attribute values and

vectorize the words by external word embeddings to obtain input representations ( $W$ s in Fig. 1). We use the 300 dimensional fastText embeddings (Bojanowski et al., 2017), which capture subword information by producing word vectors via character n-grams. This vectorization has the benefit of well representing out-of-vocabulary words (Bojanowski et al., 2017) that frequently appear in ER attributes. For instance, venue names *SIGMOD* and *ACL* are out of vocabulary in the publicly available GloVe vectors (Pennington et al., 2014), but we clearly need to distinguish them.

**Attribute Representations.** We build a universal bidirectional RNN on the word input representations of each attribute value and obtain attribute vectors ( $attr_1$  and  $attr_2$  in Fig. 1) by concatenating the last hidden units from both directions. Crucially, the universal RNN allows for transfer learning between datasets of different schemas without error-prone schema mapping. We found that gated recurrent units (GRUs, Cho et al. (2014)) yielded the best performance on the dev set as compared to simple recurrent neural networks (SRNNs, Elman (1990)) and Long Short-Term Memory networks (LSTMs, Hochreiter and Schmidhuber (1997)). We also found that using BiGRU with multiple layers did not help, and we will use one-layer BiGRUs with 150 hidden units throughout the experiments below.

**Attribute Similarity.** The resultant attribute representations are then used to compare attributes of each entity record pair. In particular, we compute the element-wise absolute difference between the two attribute vectors for each attribute and construct attribute similarity vectors ( $sim_1$  and  $sim_2$  in Fig. 1). We also considered other comparison mechanisms such as concatenation and element-wise multiplication, but we found that absolute difference performs the best in development, and we will report results from absolute difference.

**Record Similarity.** Given the attribute similarity vectors, we now combine those vectors to represent the similarity between the input entity record pair. Here, we take a simple but effective approach of adding all attribute similarity vectors ( $sim$  in Fig. 1). This way of combining vectors ensures that the final similarity vector is of the same dimensionality regardless of the number of attributes and facilitates transfer of all the subsequent parameters. For instance, the DBLP-Scholar and

Cora<sup>2</sup> datasets have four and eight attributes respectively, but the networks can share all weights and biases between the two. We also tried methods such as max pooling and average pooling, but none of them outperformed the simple addition method.

**Matching Classification.** We finally feed the similarity vector for the two records to a two-layer multilayer perceptron (MLP) with highway connections (Srivastava et al., 2015) and classify the pair into a match or a non-match (“Matching Classifier” in Fig. 1). The output from the final layer of the MLP is a two dimensional vector and we normalize it by the softmax function to obtain a probability distribution. We will discuss dataset adaptation for transfer learning in the next section.

**Training Objectives.** We train the networks to minimize the negative log-likelihood loss. We use the Adam optimization algorithm (Kingma and Ba, 2015) with batch size 16 and an initial learning rate of 0.001, and after each epoch we evaluate our model on the dev set. Training terminates after 20 epochs, and we choose the model that yields the best F1 score on the dev set and evaluate the model on the test data.

## 4 Deep Transfer Active Learning for ER

We introduce two orthogonal frameworks for our deep ER models in low resource settings: transfer and active learning. We also introduce the notion of *likely false positives* and *likely false negatives*, and provide a principled active labeling method in the context of deep ER models, which contributes to stable and high performance.

### 4.1 Adversarial Transfer Learning

The architecture described above allows for simple transfer learning: we can train all parameters in the network on source data and use them to classify a target dataset. However, this method of transfer learning can suffer from dataset-specific properties. For example, the author attribute in the DBLP-ACM dataset contains first names while that in the DBLP-Scholar dataset only has first initials. In such situations, it becomes crucial to construct network representations that are invariant with respect to idiosyncratic properties of datasets. To this end, we apply the technique of dataset (domain) adaptation developed in image recognition

<sup>2</sup><http://www.cs.umass.edu/mccallum/data/cora-refs.tar>

(Ganin and Lempitsky, 2015). In particular, we build a dataset classifier with the same architecture as the matching classifier (“Dataset Classifier” in Fig. 1) that predicts which dataset the input pair comes from. We replace the training objective by the sum of the negative log-likelihood losses from the two classifiers. We add a gradient reversal layer between the similarity vector and the dataset classifier so that the parameters in the dataset classifier are trained to predict the dataset while the rest of the network is trained to mislead the dataset classifier, thereby developing dataset-independent internal representations. Crucially, with dataset adaptation, we feed pairs from the target dataset as well as the source to the network. For the pairs from the target, we disregard the loss from the matching classifier.

### 4.2 Active Learning

Since labeling a large number of pairs for each ER scenario clearly does not scale, prior work in ER has adopted active learning as a more guided approach to select examples to label (Tejada et al., 2001; Sarawagi and Bhamidipaty, 2002; Arasu et al., 2010; de Freitas et al., 2010; Isele and Bizer, 2013; Qian et al., 2017).

Designing an effective active learning algorithm for deep ER models is particularly challenging because finding informative examples is very difficult (especially for positive examples due to the extremely low matching ratio in realistic ER tasks), and we need more than a handful of both negative and positive examples in order to tune a deep ER model with many parameters.

To address this issue, we design an iterative active learning algorithm (Algorithm 1) that searches for two different types of examples from unlabeled data in each iteration: (1) uncertain examples including likely false positives and likely false negatives, which will be labeled by human annotators; (2) high-confidence examples including high-confidence positives and high-confidence negatives. We will not label high-confidence examples and use predicted labels as a proxy. We will show below that those carefully selected examples serve different purposes.

Uncertain examples and high-confidence examples are characterized by the entropy of the conditional probability distribution given by the current model. Let  $K$  be the sampling size and the unlabeled dataset consisting of candidate record

pairs be  $D^U = \{x_i\}_{i=1}^N$ . Denote the probability that record pair  $x_i$  is a match according to the current model by  $p(x_i)$ . Then, the conditional entropy of the pair  $H(x_i)$  is computed by:

$$-p(x_i) \log p(x_i) - (1 - p(x_i)) \log(1 - p(x_i))$$

Uncertain examples and high-confidence examples are associated with high and low entropy.

Given this notion of uncertainty and high confidence, one can simply select record pairs with top  $K$  entropy as uncertain examples and those with bottom  $K$  entropy as high-confidence examples. Namely, take

$$\operatorname{argmax}_{D \subseteq D^U | |D|=K} \sum_{x \in D} H(x), \quad \operatorname{argmin}_{D \subseteq D^U | |D|=K} \sum_{x \in D} H(x)$$

as sets of uncertain and high-confidence examples respectively. However, these simple criteria can introduce an unintended bias toward a certain direction, resulting in unstable performance. For example, uncertain examples selected solely on the basis of entropy can sometimes contain substantially more negative examples than positive ones, leading the network to a solution with low recall. To address this instability problem, we propose a partition sampling mechanism. We first partition the unlabeled data  $D^U$  into two subsets:  $\overline{D}^U$  and  $\underline{D}^U$ , consisting of pairs that the model predicts as matches and non-matches respectively. Namely,  $\overline{D}^U = \{x \in D^U | p(x) \geq 0.5\}$ ,  $\underline{D}^U = \{x \in D^U | p(x) < 0.5\}$ .

Then, we pick top/bottom  $k = K/2$  examples from each subset with respect to entropy. Uncertain examples are now:

$$\operatorname{argmax}_{D \subseteq \overline{D}^U | |D|=k} \sum_{x \in D} H(x), \quad \operatorname{argmax}_{D \subseteq \underline{D}^U | |D|=k} \sum_{x \in D} H(x)$$

where the two criteria select *likely false positives* and *likely false negatives* respectively. Likely false positives and likely false negatives are useful for improving the precision and recall of ER models (Qian et al., 2017). However, the deep ER models do not have explicit features, and thus we use entropy to identify the two types of examples in contrast to the feature-based method used in Qian et al. (2017). High-confidence examples are identified by:

$$\operatorname{argmin}_{D \subseteq \overline{D}^U | |D|=k} \sum_{x \in D} H(x), \quad \operatorname{argmin}_{D \subseteq \underline{D}^U | |D|=k} \sum_{x \in D} H(x)$$

where the two criteria correspond to *high-confidence positives* and *high-confidence negatives* respectively. These sampling criteria equally partition uncertain examples and high-confidence examples into different categories. We will show that the partition mechanism contributes to stable and better performance in a later section.

---

### Algorithm 1 Deep Transfer Active Learning

---

**Require:**

Unlabeled data  $D^U$ , sampling size  $K$ , batch size  $B$ , max. iteration number  $T$ , max. number of epochs  $I$ .

**Ensure:**

Denote the deep ER parameters and the set of labeled examples by  $\mathcal{W}$  and  $D^L$  respectively.  $\text{Update}(\mathcal{W}, D^L, B)$  denotes a parameter update function that optimizes the negative log-likelihood of the labeled data  $D^L$  with batch size  $B$ . Set  $k = K/2$ .

- 1: Initialize  $\mathcal{W}$  via transfer learning. Initialize also  $D^L = \emptyset$
  - 2: **for**  $t \in \{1, 2, \dots, T\}$  **do**
  - 3:   Select  $k$  likely false positives and  $k$  likely false negatives from  $D^U$  and remove them from  $D^U$ . Label those examples and add them to  $D^L$ .
  - 4:   Select  $k$  high-confidence positives and  $k$  high-confidence negatives from  $D^U$  and add them with positive and negative labels to  $D^L$ .
  - 5:   **for**  $t \in \{1, 2, \dots, I\}$  **do**
  - 6:      $\mathcal{W} \leftarrow \text{Update}(\mathcal{W}, D^L, B)$
  - 7:     Run deep ER model on  $D^L$  with  $\mathcal{W}$  and get the F1 score.
  - 8:     **if** the F1 score improves **then**
  - 9:        $\mathcal{W}_{best} \leftarrow \mathcal{W}$
  - 10:    **end if**
  - 11:   **end for**
  - 12:    $\mathcal{W} \leftarrow \mathcal{W}_{best}$
  - 13: **end for**
  - 14: **return**  $\mathcal{W}$
- 

High-confidence examples prevent the network from overfitting to selected uncertain examples (Wang et al., 2017). Moreover, they can give the DL model more labeled data without actual manual effort. Note that we avoid using any entropy level thresholds to select examples, and instead fix the number of examples. In contrast, the active learning framework for neural network image recognition in Wang et al. (2017) uses entropy thresholds. Such thresholds necessitate fine-tuning for each target dataset: Wang et al. (2017) use different thresholds for different image recognition datasets. However, since we do not have sufficient labeled data for the target in low-resource ER problems, the necessity of fine-tuning thresholds would undermine the applicability of the active learning framework.

dataset	genre	size	matches	attr
DBLP-ACM	citation	12,363	2,220	4
DBLP-Scholar	citation	28,707	5,347	4
Cora	citation	50,000	3,969	8
Fodors-Zagats	restaurant	946	110	6
Zomato-Yelp	restaurant	894	214	4
Amazon-Google	software	11,460	1,167	3

Table 2: Post-blocking statistics of the ER datasets we used. (attr denotes the number of attributes.)

## 5 Experiments

### 5.1 Experimental Setup

For all datasets, we first conduct blocking to reduce the Cartesian product to a candidate set. Then, we randomly split the candidate set into training, development, and test data with a ratio of 3:1:1. For the datasets used in Mudgal et al. (2018) (DBLP-ACM, DBLP-Scholar, Fodors-Zagats, and Amazon-Google), we adopted the same feature-based blocking strategies and random splits to ensure comparability with the state-of-the-art method. The candidate set of Cora was obtained by randomly sampling 50,000 pairs from the result of the jaccard similarity-based blocking strategy described in Wang et al. (2011). The candidate set of Zomato-Yelp was taken from Das et al. (2016).<sup>3</sup> All dataset statistics are given in Table 2. For evaluation, we compute precision, recall, and F1 score on the test sets. In the active learning experiments, we hold out the test sets *a priori* and sample solely from the training data to ensure fair comparison with non-active learning methods. The sampling size  $K$  for active learning is 20. As preprocessing, we tokenize with NLTK (Bird et al., 2009) and lowercase all attribute values. For every configuration, we run experiments with 5 random initializations and report the average. Our DL models are all implemented using the publicly available deepmatcher library.<sup>4</sup>

### 5.2 Baselines

We establish baselines using a state-of-the-art learning-based ER package, Magellan (Konda et al., 2016). We experimented with the following 6 learning algorithms: Decision Tree, SVM, Ran-

<sup>3</sup>We constructed Zomato-Yelp by merging Restaurants 1 and 2, which are available in Das et al. (2016). Though the two datasets share the same source, their schemas slightly differ: Restaurants 1 has an address attribute that contains zip code, while Restaurants 2 has a zip code attribute and an address attribute. We put a null value for the zip code attribute in Restaurants 1 and avoid merging errors.

<sup>4</sup><https://github.com/anhaidgroup/deepmatcher>

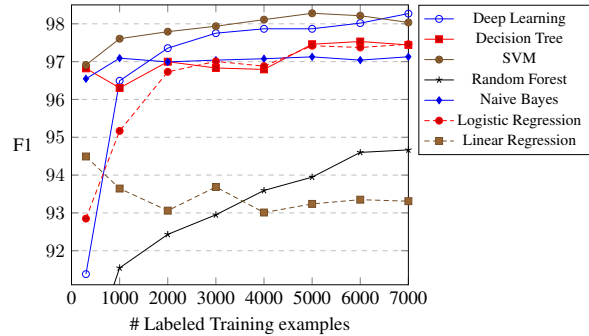


Figure 2: Performance vs. data size (DBLP-ACM).

dom Forest, Naive Bayes, Logistic Regression, and Linear Regression. We use the same feature set as in Mudgal et al. (2018). See the appendix for extensive lists of features chosen.

### 5.3 Results and Discussions

**Model Performance and Data Size.** Seen in Fig. 2 is F1 performance of different models with varying data size on DBLP-ACM. The DL model improves dramatically as the data size increases and achieves the best performance among the 7 models when 7000 training examples are available. In contrast, the other models suffer much less from data scarcity with an exception of Random Forest. We observed similar patterns in DBLP-Scholar and Cora. These results confirm our hypothesis that deep ER models are data-hungry and require a lot of labeled data to perform well.

**Transfer Learning.** Table 3 shows results from our transfer learning framework when used in isolation (i.e., without active learning, which we will discuss shortly). Our dataset adaptation method substantially ameliorates performance when the target is DBLP-Scholar (from 41.03 to 53.84 F1 points) or Cora (from 38.3 to 43.13 F1 points) and achieves the same level of performance on DBLP-ACM. Transfer learning with our dataset adaptation technique achieves a certain level of performance without any target labels, but we still observe high variance in performance (e.g. 6.21 standard deviation in DBLP-Scholar) and a huge discrepancy between transfer learning and training directly on the target dataset. To build a reliable and stable ER model, a certain amount of target labels may be necessary, which leads us to apply our active learning framework.

**Active Learning.** Fig. 3 shows results from our active learning as well as the 7 algorithms trained on labeled examples of corresponding size that are

Method	Target	DBLP-ACM			DBLP-Scholar			Cora		
		Prec	Recall	F1	Prec	Recall	F1	Prec	Recall	F1
Train on Source		86.98	98.38	92.32 $\pm$ 1.15	73.41	43.20	41.03 $\pm$ 6.33	92.54	24.22	38.30 $\pm$ 3.77
+Adaptation		88.71	96.21	92.31 $\pm$ 1.36	88.06	39.03	<b>53.84<math>\pm</math>6.21</b>	40.64	52.16	<b>43.13<math>\pm</math>3.62</b>
Train on Target		98.30	98.60	98.45 $\pm$ 0.22	92.72	93.08	92.94 $\pm$ 0.47	98.01	99.37	98.68 $\pm$ 0.26
Mudgal et al. (2018)		—	—	98.4	—	—	93.3	—	—	—

Table 3: Transfer learning results (citation genre). We report standard deviations of the F1 scores. For each target dataset, the source is given by the other two datasets (e.g., the source for DBLP-ACM is DBLP-Scholar and Cora.)

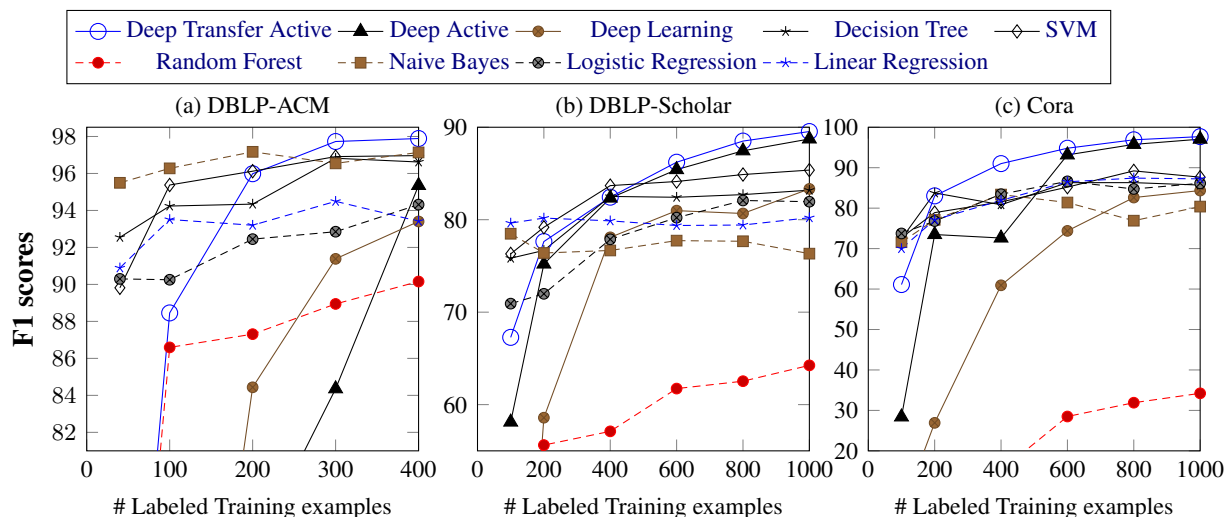


Figure 3: Low-resource performances on different datasets.

randomly sampled.<sup>5</sup> Deep transfer active learning (DTAL) initializes the network parameters by transfer learning whereas deep active learning (DAL) starts with a random initialization. We can observe that DTAL models remedy the data scarcity problem as compared to DL models with random sampling in all three datasets. DAL can achieve competitive performance to DTAL at the expense of faster convergence.

Seen in Table 4 is performance comparison of different algorithms in low-resource and high-resource settings. (We only show the SVM results since SVM performed best in each configuration among the 6 non-DL algorithms.) First, deep transfer active learning (DTAL) achieves the best performance in the low-resource setting of each dataset. In particular, DTAL outperforms the others to the greatest degree in Cora (97.68 F1 points) probably because Cora is the most complex dataset with 8 attributes in the schema. Non-DL algorithms require many interaction features, which lead to data sparsity. Deep active learning (DAL) also outperforms SVM and yields comparable performance to DTAL. However, the standard deviations in performance of DAL are substantially higher than those of DTAL (e.g. 4.15

vs. 0.33 in DBLP-ACM), suggesting that transfer learning provides useful initializations for active learning to achieve stable performance.

One can argue that DTAL performs best in the low-resource scenario, but the other algorithms can also boost their low-resource performance by active learning. While there are many approaches to active learning on feature-based (non-DL) ER (e.g. Bellare et al. (2012); Qian et al. (2017)) that yield strong performance under certain condition, it requires further research to quantify how these methods perform with varying datasets, genres, and blocking functions. It should be noted, however, that in DBLP-Scholar and Cora, DTAL in the low-resource setting even significantly outperforms SVM (and the other 5 algorithms) in the high-resource scenario. These results imply that DTAL would significantly outperform SVM with active learning in the low-resource setting since the performance with the full training data with labels serves as an upper bound. Moreover, we can observe that DTAL with a limited amount of data (less than 6% of training data in all datasets), performs comparably to DL models with full training data. Therefore, we have demonstrated that a deep ER system with our transfer and active learning frameworks can provide a stable and reliable solu-

<sup>5</sup>We average the results over 5 random samplings.

Dataset	Method	Train Size	F1
DBLP-ACM	DTAL	400	<b>97.89</b> $\pm$ 0.33
	DAL	400	95.35 $\pm$ 4.15
	DL	400	93.40 $\pm$ 2.61
	SVM	400	96.97 $\pm$ 0.69
	DL	7,417	98.45 $\pm$ 0.22
	SVM	7,417	98.35 $\pm$ 0.14
DBLP-Scholar	DTAL	1000	<b>89.54</b> $\pm$ 0.39
	DAL	1000	88.76 $\pm$ 0.76
	DL	1000	83.33 $\pm$ 1.26
	SVM	1000	85.36 $\pm$ 0.32
	DL	17,223	92.94 $\pm$ 0.47
	SVM	17,223	88.56 $\pm$ 0.46
Cora	DTAL	1000	<b>97.68</b> $\pm$ 0.39
	DAL	1000	97.05 $\pm$ 0.64
	DL	1000	84.35 $\pm$ 4.25
	SVM	1000	87.66 $\pm$ 3.15
	DL	30,000	98.68 $\pm$ 0.26
	SVM	30,000	95.39 $\pm$ 0.31

Table 4: Low-resource (shaded) and high-resource (full training data) performance comparison. DTAL, DAL, and DL denote deep transfer active learning, deep active learning, and deep learning (random sampling).

tion to entity resolution with low annotation effort.

**Other Genre Results.** We present results from the *restaurant* and *software* genres.<sup>6</sup> Shown in Table 5 are results of transfer and active learning from Zomato-Yelp to Fodors-Zagats. Similarly to our extensive experiments in the citation genre, the dataset adaptation technique facilitates transfer learning significantly, and only 100 active learning labels are needed to achieve the same performance as the model trained with all target labels (894 labels). Fig. 4 shows low-resource performance in the software genre. The relative performance among the 6 non-DL approaches differs to a great degree as the best non-DL model is now logistic regression, but deep active learning outperforms the rest with 1200 labeled examples (10.4% of training data). These results illustrate that our low-resource frameworks are effective in other genres as well.

**Active Learning Sampling Strategies.** As discussed in a previous section, we adopted high-confidence sampling and a partition mechanism for our active learning. Here we analyze the effect of the two methods. Table 6 shows deep transfer active learning performance in DBLP-ACM with varying sampling strategies. We can observe that high-confidence sampling and the partition mech-

<sup>6</sup>We intend to apply our approaches to more genres, but unfortunately we lack large publicly available ER datasets in other genres than citation. Applications to non-English languages are also of interest. We leave this for future.

Method	Prec	Recall	F1
Train on Src	100.00	6.37	11.76 $\pm$ 6.84
+Adaptation	95.33	57.27	70.13 $\pm$ 19.89
+100 active labels	100.00	100.00	100.00 $\pm$ 0.00
Train on Tgt	100.00	100.00	100.00 $\pm$ 0.00
Mudgal et al. (2018)	–	–	100

Table 5: Transfer and active learning results in the restaurant genre. The target and source datasets are Fodors-Zagats and Zomato-Yelp respectively.

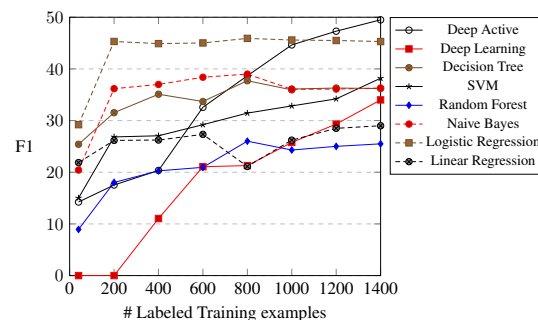


Figure 4: Low-resource performance (software genre).

anism contribute to high and stable performance as well as good precision-recall balance. Notice that there is a huge jump in recall by adding partition while precision stays the same (row 4 to row 3). This is due to the fact that the partition mechanism succeeds in finding more false negatives. The breakdown of labeled examples (Table 7) shows that is indeed the case. It is noteworthy that the partition mechanism lowers the ratio of misclassified examples (FP+FN) in the labeled sample set because partitioning encourages us to choose likely false negatives more aggressively, yet false negatives tend to be more challenging to find in entity resolution due to the skewness toward the negative (Qian et al., 2017). We observed similar patterns in DBLP-Scholar and Cora.

## 6 Further Related Work

Transfer learning has proven successful in fields such as computer vision and natural language processing, where networks for a target task is pretrained on a source task with plenty of training data (e.g. image classification (Donahue et al., 2014) and language modeling (Peters et al., 2018)). In this work, we developed a transfer learning framework for a deep ER model. Concurrent work (Thirumuruganathan et al., 2018) to ours has also proposed transfer learning on top of the features from distributed representations, but they focused on classical machine learning classifiers (e.g., logistic regression, SVMs, decision trees, random forests) and they did not con-



Sampling Method	Prec	Recall	F1
High-Confidence	93.32	97.21	95.19 $\pm$ 2.21
Partition	96.14	97.12	96.61 $\pm$ 0.57
High-Conf.+Part.	97.63	97.84	<b>97.73</b> $\pm$ 0.43
Top K Entropy	96.16	89.64	92.07 $\pm$ 9.73

Table 6: Low-resource performance (300 labeled examples) of different sampling strategies (DBLP-ACM).

Method	FP	TP	FN	TN
Part	79.6 <sub>5.9</sub>	70.4 <sub>5.9</sub>	59.2 <sub>5.6</sub>	90.8 <sub>5.6</sub>
W/o Part	101.6 <sub>7.7</sub>	57.4 <sub>15.9</sub>	41.6 <sub>4.4</sub>	99.4 <sub>22.5</sub>

Table 7: Breakdown of 300 labeled samples (uncertain samples) from deep transfer active learning in DBLP-ACM. Part, FP, TP, FN, and TN denote the partition mechanism, false positives, true positives, false negatives, and true negatives respectively.

sider active learning. Their distributed representations are computed in a “bag-of-words” fashion, which can make applications to textual attributes more challenging (Mudgal et al., 2018). Moreover, their method breaks attribute boundaries for tuple representations in contrast to our approach that computes a similarity vector for each attribute in an attribute-agnostic manner. In a complex ER scenario, each entity record is represented by a large number of attributes, and comparing tuples as a single string can be infeasible. Other prior work also proposed a transfer learning framework for linear model-based learners in ER (Negahban et al., 2012).

## 7 Conclusion

We presented transfer learning and active learning frameworks for entity resolution with deep learning and demonstrated that our models can achieve competitive, if not better, performance as compared to state-of-the-art learning-based methods while only using an order of magnitude less labeled data. Although our transfer learning alone did not suffice to construct a reliable and stable entity resolution system, it contributed to faster convergence and stable performance when used together with active learning. These results serve as further support for the claim that deep learning can provide a unified data integration method for downstream NLP tasks. Our frameworks of transfer and active learning for deep learning models are potentially applicable to low-resource settings beyond entity resolution.

## Acknowledgments

We thank Sidharth Mudgal for assistance with the DeepMatcher/Magellan libraries and replicating experiments. We also thank Vamsi Meduri, Phoebe Mulcaire, and the anonymous reviewers for their helpful feedback. JK was supported by travel grants from the Masason Foundation fellowship.

## References

- Arvind Arasu, Michaela Götz, and Raghav Kaushik. 2010. [On active learning of record matching packages](#). In *Proc. of SIGMOD*.
- Kedar Bellare, Suresh Iyengar, Aditya G. Parameswaran, and Vibhor Rastogi. 2012. [Active sampling for entity matching](#). In *Proc. of KDD*.
- Mikhail Bilenko and Raymond J. Mooney. 2003. [Adaptive duplicate detection using learnable string similarity measures](#). In *Proc. of KDD*, pages 39–48, New York, NY, USA. ACM.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural Language Processing with Python*. OReilly Media.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. [Enriching word vectors with subword information](#). *TACL*, 5:135–146.
- William M. Campbell, Lin Li, Charlie K. Dagli, Joel Acevedo-Aviles, K. Geyer, Joseph P. Campbell, and C. Priebe. 2016. [Cross-domain entity resolution in social media](#). In *Proc. of SocialNLP*.
- Kyunghyun Cho, Bart van Merriënboer, aglar Gülehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. [Learning phrase representations using RNN encoder-decoder for statistical machine translation](#). In *Proc. of EMNLP*.
- Peter Christen. 2008. [Febri: A freely available record linkage system with a graphical user interface](#). In *Proc. of HDKM*, pages 17–25, Darlinghurst, Australia, Australia. Australian Computer Society, Inc.
- Sanjib Das, AnHai Doan, Paul Suganthan G. C., Chaitanya Gokhale, and Pradap Konda. 2016. The Magellan data repository. <https://sites.google.com/site/anhaidgroup/projects/data>.
- Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, and Trevor Darrell. 2014. [DeCAF: A deep convolutional activation feature for generic visual recognition](#). In *ICML*, volume 32 of *Proceedings of Machine Learning Research*, pages 647–655, Beijing, China.

- Xin Dong, Alon Halevy, and Jayant Madhavan. 2005. [Reference reconciliation in complex information spaces](#). In *Proc. of SIGMOD*, pages 85–96, New York, NY, USA. ACM.
- Muhammad Ebraheem, Saravanan Thirumuranathan, Shafiq Joty, Mourad Ouzzani, and Nan Tang. 2018. [Distributed representations of tuples for entity resolution](#). *VLDB*.
- Jeffrey L. Elman. 1990. [Finding structure in time](#). *Cognitive Science*, 14:179–211.
- Ronald Fagin, Laura M. Haas, Mauricio A. Hernández, Renée J. Miller, Lucian Popa, and Yannis Velegrakis. 2009. [Clio: Schema mapping creation and data exchange](#). In *Conceptual Modeling: Foundations and Applications*.
- Ivan P. Fellegi and Alan B. Sunter. 1969. [A theory for record linkage](#). *JASA*.
- Junio de Freitas, Gisele Lobo Pappa, Altigran Soares da Silva, Marcos André Goncalves, Edleno Silva de Moura, Adriano Veloso, Alberto H. F. Laender, and Moisés G. de Carvalho. 2010. [Active learning genetic programming for record deduplication](#). *IEEE Congress on Evolutionary Computation*, pages 1–8.
- Yaroslav Ganin and Victor Lempitsky. 2015. [Unsupervised domain adaptation by backpropagation](#). In *Proc. of ICML*, volume 37 of *Proceedings of Machine Learning Research*, pages 1180–1189, Lille, France. PMLR.
- Mauricio A. Hernández and Salvatore J. Stolfo. 1995. [The merge/purge problem for large databases](#). In *Proc. of SIGMOD*, pages 127–138, New York, NY, USA. ACM.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural Computation*, 9:1735–1780.
- Robert Isele and Christian Bizer. 2013. [Active learning of expressive linkage rules using genetic programming](#). *J. Web Sem.*, 23:2–15.
- Diederik P. Kingma and Jimmy Lei Ba. 2015. [ADAM: A Method for Stochastic Optimization](#). In *ICLR*.
- Pradap Konda, Sanjib Das, C. PaulSuganthanG., An-Hai Doan, Adel Ardalan, Jeffrey R. Ballard, Han Li, Fatemah Panahi, Haojun Zhang, Jeffrey F. Naughton, Shishir Prasad, Ganesh Krishnan, Rohit Deep, and Vijay Raghavendra. 2016. [Magellan: Toward building entity matching management systems](#). *VLDB*, 9:1197–1208.
- Hanna Köpcke, Andreas Thor, and Erhard Rahm. 2010. [Evaluation of entity resolution approaches on real-world match problems](#). *VLDB*, pages 484–493.
- Sidharth Mudgal, Han Li, Theodoros Rekatsinas, An-Hai Doan, Youngchoon Park, Ganesh Krishnan, Rohit Deep, Esteban Arcaute, and Vijay Raghavendra. 2018. [Deep learning for entity matching: A design space exploration](#). In *Proc. of SIGMOD*, pages 19–34, New York, NY, USA. ACM.
- Sahand N. Negahban, Benjamin I. P. Rubinstein, and Jim Gemmell. 2012. [Scaling multiple-source entity resolution using statistically efficient transfer learning](#). In *Proc. of CIKM*.
- Hanna M. Pasula, Bhaskara Marthi, Brian Milch, Stuart J. Russell, and Ilya Shpitser. 2002. [Identity uncertainty and citation matching](#). In *Proc. of NeurIS*.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. [GloVe: Global vectors for word representation](#). In *Proc. of EMNLP*.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke S. Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proc. of NAACL*.
- Kun Qian, Lucian Popa, and Prithviraj Sen. 2017. [Active learning for large-scale entity resolution](#). In *CIKM*, pages 1379–1388, New York, NY, USA. ACM.
- Erhard Rahm and Philip A. Bernstein. 2001. [A survey of approaches to automatic schema matching](#). *VLDB*, 10:334–350.
- Sunita Sarawagi and Anuradha Bhamidipaty. 2002. [Interactive deduplication using active learning](#). In *Proc. of KDD*.
- Rupesh Kumar Srivastava, Klaus Greff, and Jürgen Schmidhuber. 2015. [Training very deep networks](#). In *Proc. of NeurIS*.
- Sheila Tejada, Craig A. Knoblock, and Steven Minton. 2001. [Learning object identification rules for information integration](#). *Inf. Syst.*, 26:607–633.
- Saravanan Thirumuranathan, Shameem Puthiya Parambath, Mourad Ouzzani, Nan Tang, and Shafiq R. Joty. 2018. [Reuse and adaptation for entity resolution through transfer learning](#). arXiv:1809.11084.
- Jiannan Wang, Guoliang Li, Jeffrey Xu Yu, and Jianhua Feng. 2011. [Entity matching: How similar is similar](#). *VLDB*, 4(10):622–633.
- Keze Wang, Dongyu Zhang, Ya Li, Ruimao Zhang, and Liang Lin. 2017. [Cost-effective active learning for deep image classification](#). *IEEE Trans. Cir. and Sys. for Video Technol.*, 27(12):2591–2600.
- Xin Zhao, Yuexin Wu, Hongfei Yan, and Xiaoming Li. 2014. [Group based self training for e-commerce product record linkage](#). In *Proc. of COLING*, pages 1311–1321, Dublin, Ireland. Dublin City University and Association for Computational Linguistics.

## A Appendices

### A.1 Deep ER Hyperparameters

Seen in Table 8 is a list of hyperparameters for our deep entity resolution models. We use the same hyperparameters regardless of scenario and dataset. We initialize the 300 dimensional word embeddings by the character-based pre-trained fastText vectors publicly available.<sup>7</sup>

Input Representations	
Word embedding size	300
Input dropout rate	0.0
Word-level BiGRU	
GRU size	150
# GRU layers	1
Final output	concat
Similarity Representations	
Attr. sim.	absolute diff.
Record sim.	sum
Matching Classification	
# MLP layers	2
# MLP size	300
# MLP activation	relu
Highway Connection	Yes
Domain Classification (Adversarial)	
# MLP layers	2
# MLP size	300
# MLP activation	relu
Highway Connection	Yes
Training	
Objective	cross-entropy
Batch size	16
# Epochs	20
Adam (Kingma and Ba, 2015) lr rate	0.001
Adam $\beta_1$	0.9
Adam $\beta_2$	0.999

Table 8: Deep ER hyperparameters.

### A.2 Non-DL Learning Algorithms

Magellan (Konda et al., 2016) is an open-source package that provides state-of-the-art learning-based algorithms for ER.<sup>8</sup> We use the package to run the following 6 learning algorithms for baselines: Decision Tree, SVM, Random Forest, Naive Bayes, Logistic Regression, and Linear Regression. For each attribute in the schema, we apply the following similarity functions: q-gram jaccard, cosine distance, Levenshtein distance, Levenshtein similarity, Monge-Elkan measure, and exact matching.

<sup>7</sup><https://github.com/facebookresearch/fastText>

<sup>8</sup><https://sites.google.com/site/anhaidgroup/projects/magellan>