

ERGAN: Generative Adversarial Networks for Entity Resolution

Anonymous authors

Abstract—Entity resolution targets at identifying records that represent the same real-world entity from one or more datasets. A major challenge in learning-based entity resolution is how to reduce the label cost for training. Due to the quadratic nature of record pair comparison, labeling is a costly task that requires a significant effort from human experts. However, without sufficient training data, a powerful machine learning model may be overfitting. This challenge is further aggravated when the underlying data distribution is highly imbalanced, which commonly occurs in entity resolution applications. Inspired by recent advances of generative adversarial network (GAN), in this paper, we propose a novel deep learning method, called ERGAN, to address the challenge. ERGAN consists of two key components: a label generator and a discriminator which are optimized alternatively through adversarial learning. To alleviate the issues of overfitting and highly imbalanced distribution, we design two novel modules for diversity and propagation, which can greatly improve the model generalization power. We theoretically prove that ERGAN can overcome the model collapse and convergence problems in the original GAN. We also conduct extensive experiments to empirically verify the labeling and learning efficiency of ERGAN. The experimental results show that ERGAN beats all state-of-the-art baselines, including unsupervised, semi-supervised, and unsupervised learning methods.

I. INTRODUCTION

Entity Resolution (ER) is an important and ubiquitous component of real-world applications in various fields, such as national census, health sector, crime and fraud detection, bibliographic statistics, and online shopping [8]. For example, if a national census agency wants to obtain the population growth of its country in a time period, ER is necessary to detect whether records from different time points refer to the same person, no matter whether he or she changed name (e.g. due to marriage), postal address and so on. Learning-based ER methods have been widely used in the past years. However, due to the quadratic nature of record pair comparison required by ER tasks [6], labeling is costly, time consuming, and highly imbalanced. This raises the difficulty of applying supervised learning methods for ER because supervised learning methods require a sufficient number of labeled instances for training, which is infeasible in many real-world applications. To reduce the labeling effort, a number of semi-supervised learning methods have been proposed [22], [38], [33].

Generally, semi-supervised learning aims to train a machine learning model with both labeled instances (called seeds) and unlabeled instances [42]. Several semi-supervised learning methods have been proposed based on a low-density separation assumption, i.e. there exists a low-density “boundary” so that instances belonging to different classes can be distinguished [2], [25]. However, such a boundary may not always exist

or can be clearly identified, especially when the number of labeled instances is small [19]. Some semi-supervised learning methods have utilized the idea of self-learning, which firstly trains a classifier using labeled instances, and then selects unlabeled instances with predicted labels to train a classifier iteratively [22]. Although promising, these methods often lead to the issue of overfitting when labeled instances in training are limited [29].

In this paper, we focus on tackling the following two challenges, which cannot be handled by the existing ER methods: (1) **the overfitting problem**; (2) **the imbalanced class problem**. The overfitting problem happens when the number of labeled instances is limited and a learning model is powerful enough to remember all the features of training instances. In such cases, the learning model can correctly predict the classes of seen instances with high certainty, but fail to predict the classes of unseen instances, thus losing the generalization ability. For the imbalanced class problem, it is due to the fact that the number of matches (record pairs referring to the same entity) is far less than the number of non-matches in ER tasks. For example, given two datasets with the number of records n and m , respectively, the largest number of matches between them is $\min(n, m)$ while the largest number of non-matches is $n*m$. Traditionally, blocking techniques can help to alleviate the imbalanced class problem by grouping potentially matched instances into the same cluster. However, selecting a blocking method also requires prior knowledge or sufficient training instances [37], [32], which is still hard to achieve under a very limited number of training instances.

Generative adversarial network (GAN) and its variants have recently emerged as a powerful deep learning technique for real-world applications across various domains such as image generation and natural language processing [17], [16]. Inspired by these advances, in this paper, we develop a novel generative adversarial network, called ERGAN, to solve the aforementioned challenges faced by ER applications. In ERGAN, there are two key components: (1) a *label generator* G that aims to generate pseudo labels for unlabeled instances, and (2) a *discriminator* D that aims to distinguish instances with pseudo labels from instances with real labels. The discriminator D is trained using not only a small number of instances with real labels but also a large number of instances with high-quality pseudo labels. However, the question arises: how to ensure the high-quality of pseudo labels generated for unlabeled instances? Unfortunately, the existing GAN and its variants cannot guarantee this when the number of instances with real

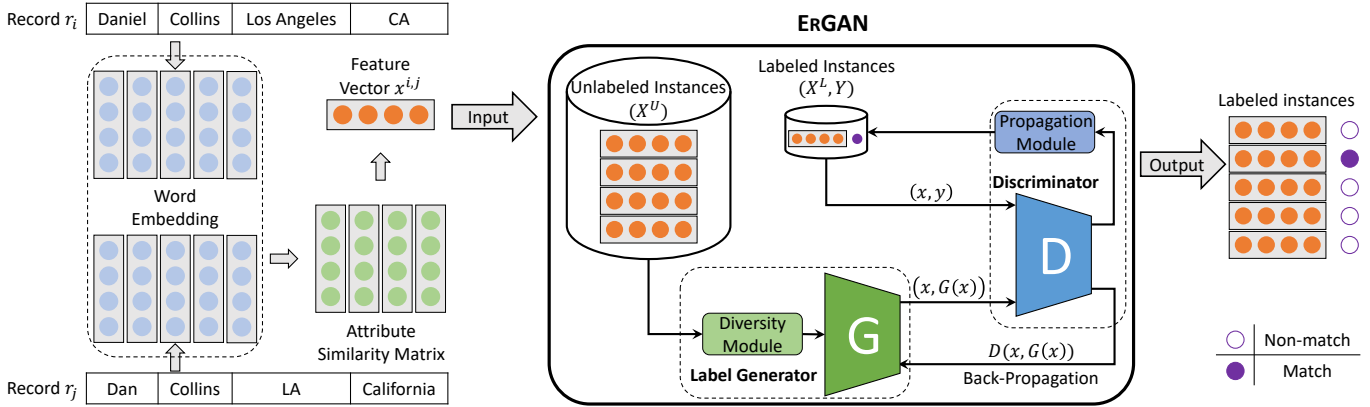


Fig. 1: **Overview of our framework.** Given a record pair r_i and r_j , each record is first represented as a matrix of vector embeddings using word embedding. Based on this, the record pair is transformed into an attribute similarity matrix, which leads to a feature vector (i.e., an unlabeled instance). Using only a limited number of labeled instances for training, ERGAN takes unlabeled instances as input and classifies them as being matches or non-matches (i.e. predicting their labels).

labels is limited. To address this question, our model ERGAN is designed to incorporate two modules: *diversity module* and *propagation module* into the label generator G and the discriminator D , respectively. The diversity module enriches the diversity of unlabeled instances during the sampling process, while the propagation module guarantees that only unlabeled instances with high-quality pseudo labels can be propagated into the training of the discriminator D . We formally prove that G and D converge to the equilibrium point, achieving the global optimality. Figure 1 shows an overview of our framework of using ERGAN for ER tasks.

ERGAN can alleviate the overfitting problem. This is because, instead of training a single classifier, ERGAN trains the label generator G and the discriminator D adversarially such that D improves G to generate pseudo labels for unlabeled instances and G regularizes D from overfitting through propagating unlabeled instances with high-quality pseudo labels. In essence, G serves as an adaptive regularization term acting on D during the minibatch training process. ERGAN can also alleviate the imbalanced class problem. We observe that, the key to solving the imbalanced class problem lies in how to effectively approximate the true distribution using limited instances with real labels. Driven by this observation, ERGAN employs the diversity module to sample instances diversely from different feature subspaces. This leads to a dual boosting: both improving label efficiency of selecting instances from the minority class, and improving learning efficiency of learning a fast approximation on the underlying data distribution.

Contributions. Our contributions are summarized as follows:

- **We propose a semi-supervised generative adversarial network, namely ERGAN, for entity resolution.** ERGAN has two key components: a label generator and a discriminator, which are optimized in an adversarial learning manner. As ERGAN is specifically designed for ER classification, it can be used jointly with any word embedding or string matching techniques for ER tasks.

- **We develop two integrated modules: diversity and propagation modules to improve the model generalization ability.** Consequently, even when only a very limited number of labeled instances are available, ERGAN can still effectively infer the true distribution of all labels in a semi-supervised manner.
- **We theoretically prove that ERGAN overcomes the model collapse and convergence problems in the original GAN.** Specifically, the following properties of ERGAN are proven: (1) The global optimality can be guaranteed; (2) The label generator G and the discriminator D converge to the equilibrium point; (3) The diversity module helps improve the generalization without compromising the global optimality and equilibrium; and (4) The mode collapse issue is alleviated in ERGAN.
- **We conduct extensive experiments to empirically verify the effectiveness of ERGAN.** The experimental results show that ERGAN outperforms all state-of-the-art baselines, including unsupervised, semi-supervised and full-supervised learning methods. Even when the label cost is very limited and the state-of-the-art baselines fail to predicate labels, ERGAN can still achieve reasonably good performance.

It is worthy to note that, although we only consider the use of ERGAN for entity resolution in this paper, the techniques of ERGAN for handling overfitting and imbalanced data can be much more widely applicable.

II. PROBLEM FORMULATION

Let R be an ER dataset consisting of a set of records where each $r \in R$ is associated with a number of attributes A . We use $r.a$ to refer to the value of an attribute $a \in A$ in a record r . Each record pair (r_i, r_j) in R corresponds to a feature vector $x^{(ij)} \in \mathbb{R}^m$ where each dimension $x_k^{(ij)}$ indicates a feature value, e.g., the textual similarity of values in an attribute $a_k \in A$, i.e., $x_k^{(ij)} = f(r_i.a_k, r_j.a_k)$, calculated by a function f .

Let $X = \{x^{(ij)} | (r_i, r_j) \subseteq R \times R\}$ be the set of all instances corresponding to record pairs in R and $Y = \{M, N\}$ be a label space where M and N refer to two labels *match* and *non-match*, respectively. There is a small subset $X^L \subseteq X$ of instances that are labeled, while the other instances in X are unlabeled, i.e., $X^U = X - X^L$. We assume $|X^L| \ll |X^U|$, i.e., X has a very limited number of labeled instances in X^L but a large number of unlabeled instances in X^U . We denote (X^L, Y) as a set of instances in X^L and their labels in Y , and $(x^L, y) \sim (X^L, Y)$ as a pair of instance $x^L \in X^L$ and its label $y \in Y$. Our task is to tackle the ER classification problem as formulated below.

Definition 1. Given a set X of instances with $X = X^L \cup X^U$ and $|X^L| \ll |X^U|$, and a label space $Y = \{M, N\}$, the **ER classification problem** is to learn a model Λ that can predict a label $\hat{y} \in Y$ for each unlabeled instance $x \in X^U$ w.r.t.

$$\max E(\Lambda) / |X^U| \quad (1)$$

$$\text{where } E(\Lambda) = \sum_{x \in X^U \wedge \hat{y}=y} 1.$$

Intuitively, $E(\Lambda)$ refers to the total number of unlabeled instances in X^U whose labels are correctly classified by Λ .

III. PROPOSED METHOD: ERGAN

Our proposed method ERGAN consists of two components: (1) a *label generator* G ; and (2) a *discriminator* D . Both G and D are differentiable functions.

A. Label Generator

In ERGAN, a label generator G can obtain instances from $p(X^U)$, but does not know about $p(Y)$ nor $p(X, Y)$. Nevertheless, we know that $p(X^U) \approx p(X)$ because $X^U \subseteq X$ and $|X^U|/|X|$ is close to 1. The goal of G is to learn a conditional distribution $p_g(Y|X^U) \approx p(Y|X^U)$, i.e., given an instance $x \sim p(X^U)$ as input, G generates a pseudo label \hat{y} for x . Ideally, the pseudo label \hat{y} generated for an instance x by G should be the same as the real label of x . To simulate the conditional distribution $p(Y|X^U)$, the label generator G receives feedback (i.e. gradients) from the discriminator D and is trained iteratively through backpropagation.

Diversity module. One major difference of our ERGAN from the original GAN and its variants such as CatGAN [35] is that we consider the diversity of instances in the minibatch sampling process. More specifically, for all instances in X , we partition them into a number of non-overlapping subspaces alike in certain features $\{X_1, \dots, X_b\}$ such that instances in the same subspace are more similar than those in different subspaces. Accordingly, labeled instances in X^L and unlabeled instances in X^U are partitioned into these b subspaces, i.e., $X_i^L = X_i \cap X^L$ and $X_i^U = X_i \cap X^U$.

Let $\mathbf{v} = (\mathbf{v}_1, \dots, \mathbf{v}_b)$ be a vector corresponding to b subspaces, where each $\mathbf{v}_i = (v_i^1, \dots, v_i^{n_i})^T \in [0, 1]^{n_i}$ and $n_i = |X_i^U|$. That is, each v_i^j ($1 \leq j \leq n_i$) is associated with an

instance in X_i^U . Then, a minibatch of m instances is selected from X^U according to the following objective function:

$$\text{maximize } \|\mathbf{v}\|_{2,1} \quad \text{s.t. } \sum_{i,j} v_i^j = m \quad (2)$$

where $\|\mathbf{v}\|_{2,1}$ is a $l_{2,1}$ -norm function defined as:

$$\|\mathbf{v}\|_{2,1} = \sum_{i=1}^b \|\mathbf{v}_i\|_2 = \sum_{i=1}^b \sqrt{\sum_{j=1}^{n_i} v_i^j{}^2} \quad (3)$$

Here, $\|\mathbf{v}_i\|_2$ is the l_2 -norm of \mathbf{v}_i . When $v_i^j = 1$, the instance in X_i^U corresponding to v_i^j is selected into the minibatch; otherwise, that instance is not selected. When the value of the $l_{2,1}$ -norm is small, instances are selected from a small number of subspaces in X^U and the diversity of instances is low. Conversely, when maximizing the $l_{2,1}$ -norm in Eq. 2, instances are selected from as many subspaces in X^U as possible and the diversity of instances is high.

Objective function of G . After a minibatch of unlabeled instances is selected from X^U according to Eq. 2, the label generator G generates a pseudo label $G(x_i)$ for each unlabeled instance x_i in the minibatch. Then, $(x_i, G(x_i))$ is sent to the discriminator D . After receiving the gradient from D , G updates its parameters according to the following objective:

$$\mathcal{L}_G = \min_G \mathbb{E}_{x \sim p(X_i^U)} [\log(1 - D(x, G(x)))] \quad (4)$$

B. Discriminator

Unlike GAN, a discriminator D in our ERGAN does not know about the real distribution $p(X, Y)$. Instead, D has access only to a limited number of instances with real labels, i.e. (X^L, Y) . The goal of D is to distinguish whether a labeled instance $(x, G(x))$ is from the real distribution $p(X, Y)$, i.e., given a pair $(x, G(x))$ as input, D generates a scalar value in $[0, 1]$ to indicate the probability that $G(x)$ is the same as the real label y of x .

Propagation module. To achieve the above goal, as opposite to GAN and its variants in which the discriminator has the true distribution $p(X, Y)$, D in ERGAN is designed to approximate the true joint distribution $p(X, Y)$ progressively through a propagation module. The general principle of propagation is that, the more confident the pseudo label $G(x)$ of an instance x is the same as its real label y , the more likely such an instance is selected. Specifically, let $(X^t, G(X^t))$ denote all unlabeled instances with their pseudo labels at the t -th iteration of propagation. These instances are fed to D to obtain their scores $D(X^t, G(X^t))$ that indicates the probabilities of their pseudo labels being the same as their real labels. Based on the scores, a subset $\Delta X^t \subseteq X^t$ of instances is selected according to the following objective function:

$$\begin{aligned} \operatorname{argmax}_{\Delta X^t \subseteq X^t} \sum_{x \in \Delta X^t} D(x, G(x)) \\ \text{subject to } |\Delta X^t| = \gamma \end{aligned} \quad (5)$$

where γ is a hyper-parameter for the number of unlabeled instances being selected in the t -th iteration of propagation.

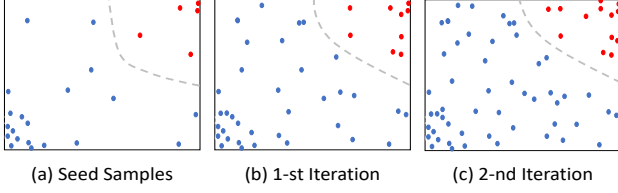


Fig. 2: **An illustration for propagation of ERGAN.** A boundary between two classes (red and blue) is learned through propagation.

Then, this subset of instances with their high-quality pseudo labels $(\Delta X^t, \hat{Y})$ is propagated into the set of labeled instances $(X^*, Y)^t$ to train D , i.e.,

- $(X^*, Y)^0 = (X^L, Y)$
- $(X^*, Y)^t = (X^*, Y)^{t-1} \cup (\Delta X^t, \hat{Y})$

Hence, at the t -th iteration of propagation, D has access to $(X^*, Y)^t$, which is a mixed set of labeled instances from X^L (with real labels) and unlabeled instances from X^U (with pseudo labels generated by G). The following holds:

$$(X^*, Y)^0 \subseteq (X^*, Y)^1 \subseteq \dots \subseteq (X^*, Y)^t \quad (6)$$

Figure 2 shows an example of the propagation in two iterations, where the grey dash line indicates a boundary between two classes (red and blue) and is learned through propagation.

Objective function of D . The objective function of D at the t -th iteration of propagation is defined as:

$$\mathcal{L}_D = \max_D \mathbb{E}_{x \sim p(X_i^U)} \log[1 - D(x, G(x))] + \lambda \mathbb{E}_{(x,y) \sim (X^*, Y)^t} \log[D(x, y)] \quad (7)$$

where λ refers to a weighted term. In the following, we will explain how unlabeled instances with their pseudo labels, i.e., (X^t, \hat{Y}) , is selected at the t -th iteration of propagation.

C. Algorithm Description

Our algorithm is described in Algorithm 1. It involves two key processes: batch training (Lines 3-8) and label propagation (Lines 9-12). In the batch training process, a minibatch is randomly sampled from unlabeled instances in X^U and G generates pseudo labels for samples in this minibatch (Lines 4-5). Then another minibatch is randomly sampled from instances in $(X^*, Y)^t$ (Line 6). After that, the discriminator D is trained using these two minibatches w.r.t. Eq. 7, while the label generator G is trained w.r.t. Eq. 4. The label propagation process only occurs after each batch training is finished. At this time, G and D reach (or closely approach) the equilibrium point and their parameters are fixed. In the label propagation process, G first generates pseudo labels for all unlabeled instances in X^t , and then D predicts scores for these instances. Based on these scores, a subset of instances $\Delta X^t \subseteq X^t$ with high-quality pseudo labels is propagated into the set of labeled instances $(X^*, Y)^t$ for D in the t -th iteration.

How to choose b , t and n ? In our algorithm, the number of subspaces b is decided based on the attributes in each dataset. Suppose that a dataset has four attributes, we first obtain the

median value for each attribute, and then partition instances into $4^2 = 16$ subspaces according to whether attribute values of each instance are above or below the median values of these four attributes [33]. Furthermore, n is a hyper-parameter referring to the number of iterations for converging G and D , and t is decided by the total number X^U of unlabeled instances and the number γ of instances being propagated in each iteration, i.e. $t = \lceil \frac{|X^U|}{\gamma} \rceil$.

Algorithm 1: Minibatch stochastic gradient descent and label propagation of ERGAN

Input: b subspaces in X ; X^U ; (X^L, Y) ;
Output: $(X^*, Y)^t$ where $X^* = X$

- 1 Initialize $t = 0$; $(X^*, Y)^0 = (X^L, Y)$;
 $X^0 = X^1 = X^U$
- 2 **while** $X^t \neq \emptyset$ **do**
- 3 **for** n iterations **do** // Batch training
- 4 Sample a minibatch $\{x_1, \dots, x_m\}$ from X^U
w.r.t. Eq. 2
- 5 Generate pseudo labels
 $\{(x_1, G(x_1)), \dots, (x_m, G(x_m))\}$
- 6 Sample a minibatch $\{(x_1^L, y_1), \dots, (x_m^L, y_m)\}$
from $(X^*, Y)^t$
- 7 Update the parameters of D w.r.t. Eq. 7
- 8 Update the parameters of G w.r.t. Eq. 4
- 9 $t = t + 1$ // Label propagation
- 10 Generate pseudo labels for X^t
- 11 Select $\Delta X^t \subseteq X^t$ for propagation w.r.t. Eq. 5
- 12 $(X^*, Y)^t = (X^*, Y)^{t-1} \cup (\Delta X^t, \hat{Y})$;
 $X^{t+1} = X^t - \Delta X^t$

IV. THEORETICAL ANALYSIS

We prove several nice theoretical properties of ERGAN.

Diversity. We first show that partitioning X into subspaces for diversity does not affect the learning goals of G and D .

Lemma 1. *If $p_g(Y|X_i^U)$ by G approximates $p(Y|X_i^U)$ for each feature subspace X_i^U , then $p_g(Y|X^U) \approx p(Y|X^U)$.*

Proof. X^U is partitioned into non-overlapping subspaces $\{X_i^U\}_{i=1}^b$. Thus, $p(X^U)$ is composed of b distributions $p(X_i^U)$ ($1 \leq i \leq b$). Further, Eq. 2 prefers instances from different subspaces, but within each subspace, instances are selected randomly. \square

Since unlabeled and labeled instances are sampled from the corresponding subspaces, we have the following lemma for D .

Lemma 2. *If $p_d(X_i^U, Y)$ by D approximates $p(X_i^U, Y)$ for each feature subspace X_i^U , then $p_d(X^U, Y) \approx p(X^U, Y)$.*

By $|X^L| \ll |X^U|$, we have $p(X^U) \approx p(X)$. Accordingly, $p_g(Y|X^U) \approx p_g(Y|X)$, $p_d(X^U, Y) \approx p_d(X, Y)$, and $p(X^U, Y) \approx p(X, Y)$.

Global optimality. Here we use $p_d(X^*, Y)^t$ to refer to the distribution learned by D at the t -th iteration of propagation. We have the following lemma.

Lemma 3. *For the fixed label generator G , the optima D in the t -th iteration of propagation is:*

$$D_G^*(x, y) = \frac{p_d(X^*, Y)^t}{p_d(X^*, Y)^t + p_g(X, Y)^t} \quad (8)$$

Proof. At the t -th iteration of propagation, by Eq. 7, the training objective of D is to maximize:

$$\mathbb{E}_{x \sim p(X_i^U)} \log[(1 - D(x, G(x)))] + \lambda \mathbb{E}_{(x, y) \sim p_d(X^*, Y)^t} \log[D(x, y)]$$

Based on the "change of variable" technique [15], we have $p_g(s) = p_{X_i^U}(G^{-1}(y)) \frac{dG^{-1}(y)}{d(s)}$, where $s = (x, y)$ denotes the vector concatenated by x and y at the t -th iteration of propagation. The first part of the above formula equals to:

$$\begin{aligned} & \int_s p_{X_i^U}(G^{-1}(y)) \log(1 - D(s)) dG^{-1}(y) \\ &= \int_s p_g(s) \log[(1 - D(s))] d(s) \end{aligned}$$

Hence, we have an objective function that is the same as GAN [17]. This lemma is proven. \square

As a result, we have the following lemma.

Lemma 4. *The global minimum of the training criterion of G under the optimal D at the t -th iteration of propagation is achieved when $p_d(X^*, Y)^t = p_g(X, Y)^t$.*

Equilibrium. By Lemma 4, we have $p_g(X, Y)^t = p_d(X^*, Y)^t$. Thus, the equilibrium of the minimax game in ERGAN is achieved at each iteration of propagation. When the labels are propagated sufficiently, the real data distribution is well simulated. The following lemma states that D can approximate the real distribution $p(X, Y)$ when the number of propagation iterations is large.

Lemma 5. *$p_d(X^*, Y)^t$ approaches $p(X, Y)$ when t increases.*

Statistically, according to the *Central Limit Theorem*, the larger sample size the more likely an estimated distribution is close to the real distribution. This lemma corresponds to the central property of self-training based semi-supervised learning approaches [18], [36].

Mode collapse. The mode collapse problem occurs in GAN where the generator collapses to generate limited and non-sensical images. However, our approach ERGAN does not suffer from this problem for two reasons: (1) ERGAN takes unlabeled instances as input for both G and D . This allows D to distinguish $p_d(X^*, Y)^t$ from $p_g(X, Y)^t$, and D can prevent G to generate the same kind of pseudo labels for unlabeled instances, i.e., mode collapse. (2) In GAN, mode collapse only occurs when the discriminator is well trained but the generator is not optimal. In ERGAN, $(X^*, Y)^t$ only

TABLE I: **Characteristics of datasets.** The instances of these datasets are generated from their record pairs.

Dataset	#Attributes (A)	#Instances (X)	Imbalance Rate	#Subspaces (b)
Cora	4	837,865	1:49	16
DBLP-ACM	4/4	6,001,104	1:2,698	16
DBLP-Scholar	4/4	168,112,008	1:71,233	16
NCVoter	18/18	1,000,000	1:4,202	64

has limited instances with labels when t is small. Thus, D can hardly be optimal in early iterations of propagation. For later iterations, since G is trained iteratively based on its previous parameters, the mode collapse issue can also be avoided.

V. EXPERIMENTAL SETUP

We evaluate ERGAN to answer the following questions:

- Q1.** How does ERGAN perform in comparison with the state-of-the-art unsupervised, semi-supervised and fully supervised methods?
- Q2.** How do the design choices such as the diversity module, the propagation module, and GAN architecture affect performance of ERGAN?
- Q3.** To what degree ERGAN can work for classifying instances when the label cost is extremely limited (i.e. only a small number of instances with real labels)?

Datasets. We use four widely used benchmark datasets of ER tasks: (1) *Cora*¹ dataset contains bibliographic records of machine learning publications. (2) *DBLP-Scholar*¹ dataset contains bibliographic records from the DBLP and Google Scholar websites. (3) *DBLP-ACM* [24] dataset contains bibliographic records from the DBLP and ACM websites. (4) *North Carolina Voter Registration (NCVoter)*² dataset contains real-world voter registration information of people from North Carolina in the USA. Table I summarizes the characteristics of these four datasets. We can see that the datasets are highly imbalanced, i.e., the number of instances from the majority class (non-match) is much more than the number of those from the minority class (match).

Baselines. We compare ERGAN with the following baselines: (1) *Unsupervised methods: Two-Steps (2S)* is a widely-used unsupervised learning method proposed by Christen [6]. **Iterative Term-Entity Ranking and CliqueRank (ITER-CR)** is the state-of-the-art graph based unsupervised method for ER [40]. (2) *Semi-supervised methods: Semi-supervised Boosted Classifier (SBC)* is the state-of-the-art semi-supervised learning method with label propagation based on Adaboost classifier [30] proposed by Kejriwal and Miranker [22]. (3) *Fully supervised methods: Magellan* is a state-of-the-art open-source fully supervised learning-based ER solution designed by Konda et. al. [23]. We consider two supervised classifiers provided in Magellan: **Logistic Regression (LR)** and **Support Vector Machine (SVM)**. **eXtreme Gradient boosting**

¹Available from: <http://secondstring.sourceforge.net>

²Available from: <http://alt.ncsbe.gov/data/>

TABLE II: **Experimental results of f-measure with 60% training.** The results marked by * are taken from the original papers and the others are obtained by running the code provided by the authors.

Method	Datasets			
	Cora	DBLP-ACM	DBLP-Scholar	NCVoter
2S [6]	62.69	91.43	68.78	98.96
ITER-CR* [40]	89.00	—	—	—
SBC [22]	85.71	97.09	85.47	99.78
SVM [23]	88.95	97.19	85.71	98.48
LR [23]	80.25	95.56	83.84	99.37
XGBoost [5]	91.34	97.20	86.63	100
ERGAN	93.03	98.23	88.32	100
DM [27]	98.58	98.29	94.68	100
DTAL* [21]	98.68 \pm 0.26	98.45 \pm 0.22	92.94 \pm 0.47	—
ERGAN+WE	98.72 \pm 0.15	98.51 \pm 0.23	94.73 \pm 0.35	100

(**XGboost**) is a state-of-the-art fully supervised ensemble learning based method proposed by Chen and Guestrin [5]. **DeepMatcher (DM)** is a state-of-the-art deep learning based entity matching method for ER [27]. **Deep Transfer active learning (DTAL)** is the state-of-the-art active learning method which combines both transfer learning and active learning for handling ER tasks [21].

To make a fair comparison, we follow the default parameters suggested in the original papers of the baselines. Note that DM uses the imbalance rate as a hyper-parameter which is normally unknown in real-life applications. Both DTAL and DM are deep learning based methods in which FastText [4] is used for learning word embeddings. For other baselines, we use 2-gram Jaccard similarity for textual comparison.

To compare with the baselines that use word embeddings, we use **ERGAN+WE** to refer to the model of ERGAN augmented with word embeddings for attribute values. In our ablation study, we use **ERGAN-D** and **ERGAN-P** to refer to a model being obtained by removing the diversity and propagation modules from ERGAN, respectively, and **ERNN** a model in which the GAN architecture (i.e. G and D are trained alternatively) is replaced by a single multi-layer perceptron for semi-supervised learning with the diversity module. We set $\lambda = 1$, $m \leq 100$, and $\gamma = |X^*|$ at each iteration of propagation. Our models use the same word embedding and similarity comparison techniques as the baselines.

Measures. We use the following widely used measures in ER tasks for performance evaluation [8]: (1) *Recall* (R) is the fraction of true matches being predicted by the classifier among the total number of true matches; (2) *Precision* (P) is the fraction of true matches over all matches being predicted by the classifier; (3) *F-measure* (FM) is the harmonic mean of recall and precision, i.e. $FM = \frac{2 * R * P}{R + P}$. Due to the existence of highly imbalanced classes in ER datasets, F-measure is preferred rather than accuracy in our experiments.

VI. EXPERIMENTAL RESULTS

In this section, we discuss the results of our experiments to answer the aforementioned questions.

A. Performance Comparison

Task 1. We conduct an experiment to evaluate how our methods perform against the baselines. Following the previous work for the supervised methods DM [27] and DTAL [21], we split the datasets with 60% for training and the rest for testing.

Table II shows the results of the experiment, where the last three methods DM, DTAL and ERGAN+WE are deep-learning methods which use word embeddings for attribute values and the other methods use Jaccard similarity for comparing attribute values (without using word embeddings). We can see that, the unsupervised method 2S performs the worst among all the methods. However, the other unsupervised method ITER-CR performs better than SBC, SVM and LR due to its ability to leverage graph based structure. Compared with the fully supervised methods, the semi-supervised method SBC performs better than LR, comparably with SVM, but worse than XGBoost. Our method ERGAN performs better than any non-deep-learning method, but worse than the deep-learning methods with word embedding, i.e., DM, DTAL and ERGAN+WE. Nonetheless, our method ERGAN+WE outperforms all the baseline, including two deep-learning methods DM and DTAL, over all databases they have the results.

Observation 1. *With sufficient training data, ERGAN+WE performs better than ERGAN due to the power of word embedding for records. ERGAN+WE has superior performance against all the baselines consistently.*

Task 2. To understand how performance may change when the amount of training data is reduced, we conduct an experiment on the supervised methods by splitting the datasets into 20% for training and 80% for testing. DTAL is excluded in this experiment because the original paper does not have results in this setting and also no code is available.

Figure 3 shows the results of precision, recall and f-measure over four datasets. LR performs the worst except for precision on DBLP-ACM. The performance of the semi-supervised method SBC is better than LR and SVM, and comparable with XGBoost w.r.t. f-measure. This is because SBC takes the advantage of label propagation but its classifier Adaboost is not as powerful as XGBoost. For the deep-learning methods, the performance of DM drops significantly on DBLP-ACM and DBLP-Scholar and even worse than ERGAN which is different from the case of 60% training in Table II. ERGAN+WE still performs the best on all datasets w.r.t. all the three measures. Overall, compared with Table II, the performance of SVM, LR and DM is affected considerably on one or more datasets, whereas SBC, XGBoost, ERGAN and ERGAN+WE remain comparable performance.

Observation 2. *With reduced but still considerable training data, the performance of ERGAN and ERGAN+WE remains strong and consistent. This is because ERGAN and ERGAN+WE benefit from its adversarial learning architecture and diversity and propagation modules.*

Task 3. To study performance under a limited number of instances with real labels, we further conduct an experiment

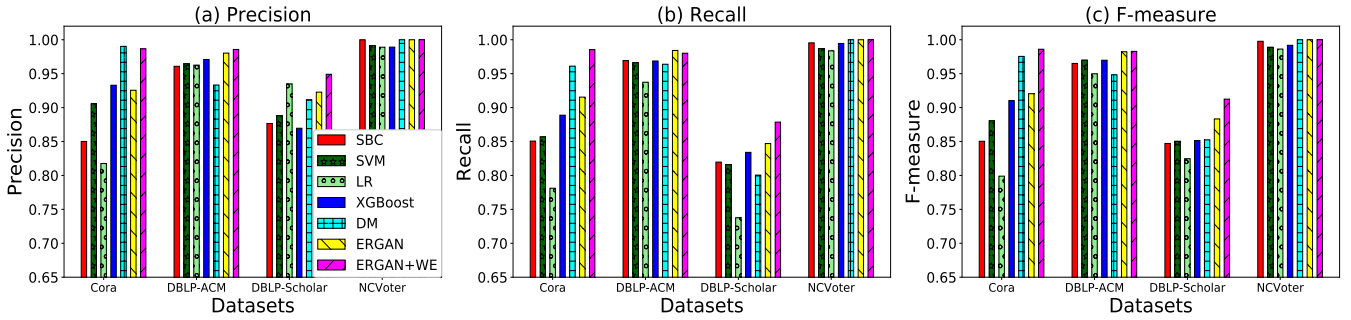


Fig. 3: Experimental results of precision, recall and f-measure with 20% training.

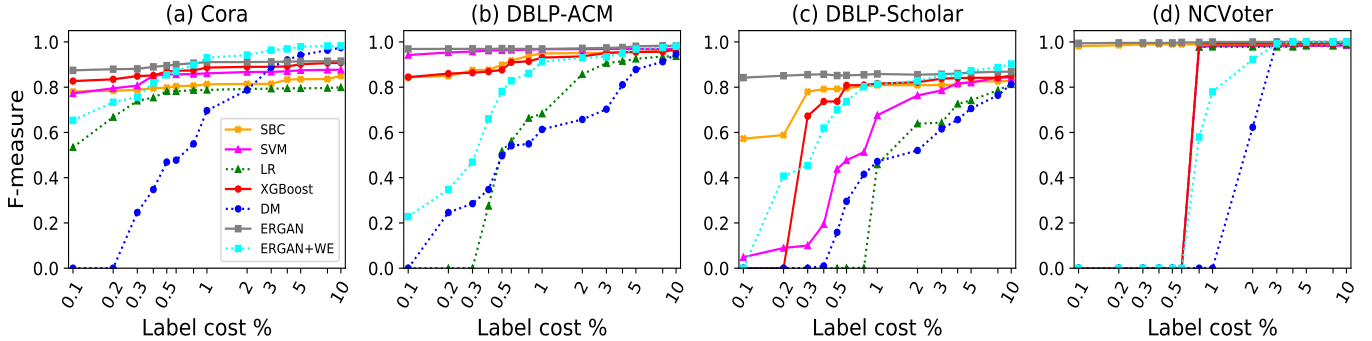


Fig. 4: Experimental results of f-measure with 0.1% - 10% training.

using only a small percentage for training, ranging from 0.1% to 10% of the datasets, and the rest for testing.

Figure 4 shows the experimental results. ERGAN performs best among all the methods over all the datasets when training data is below 1%. ERGAN+WE performs poorly in this range. However, the performance of ERGAN+WE increases rapidly with increasing training data and exceeds all the other methods on all the datasets when training data reaches 10%. LR and DM have a similar trend as ERGAN+WE, but perform significantly worse. The semi-supervised method SBC performs better than ERGAN+WE only when training data is small, i.e. below 0.2% for Cora and below 0.9% for DBLP-ACM, DBLP-Scholar and NCVoter. The performance of SVM and XGBoost varies in datasets, i.e., perform well on Cora and DBLP-ACM, but badly on DBLP-Scholar and NCVoter. This demonstrates that the performance of SVM and XGBoost is sensitive to the imbalance rate of a dataset, and they fail to handle imbalanced data when no sufficient training data is available. For NCVoter, due to a clear boundary existing between matches and non-matches in the underlying distribution, the performance of all the methods that perform poorly for small training data can be dramatically improved after using 0.6% or more training data. In general, we may conclude that, compared with the cases of 60% and 20% training in Table II and Figure 3, the performance gain of the methods with word embedding against the methods without word embedding does not exist anymore. Instead, the methods with word embedding performs worse than most of the methods without word embedding when training data is small, i.e., below 1%.

Observation 3. When decreasing training data, ERGAN+WE gradually performs worse than ERGAN+WE. This is because, ERGAN+WE transforms instances into a high dimensional space through word embedding and thus requires much more labels in training than ERGAN.

B. Ablation Analysis

We conduct an ablation study to evaluate the effects of the key components of ERGAN, including the adversarial learning architecture, the diversity module and the propagation module, under different label costs, ranging from 0.1% to 60% for training. The results are presented in Table III. We observe that the performance of all the methods ERNN, ERGAN-D, ERGAN-P and ERGAN become stable and gradually converge when the label cost increases, e.g. in the case of 60% training. Nonetheless, ERGAN performs the best among all the methods, and the performance of the other methods varies in different datasets. In the following, we will discuss how each key component of ERGAN may affect the performance.

Adversarial learning architecture. The performance of ERNN generally lies in between ERGAN-D and ERGAN-P, and significantly worse than ERGAN. This indicates that the use of adversarial learning architecture by ERGAN helps to improve the performance, particularly when training data is limited, e.g., for 0.1% training, ERGAN improves around 3% on Cora and more than 8% on DBLP-ACM upon ERNN.

Diversity module. In Table III, the results of ERGAN-D are the worst among all the methods over all the datasets. This indicates that diverse instances are more informative for model

TABLE III: Experimental results of f-measure with 0.1%, 1%, 20% and 60% training for ablation analysis.

Datasets	Cora				DBLP-ACM				DBLP-Scholar				NCVoter			
	0.1%	1%	20%	60%	0.1%	1%	20%	60%	0.1%	1%	20%	60%	0.1%	1%	20%	60%
ERNN	84.46	90.67	91.43	92.78	88.05	95.68	98.20	98.22	82.76	83.17	86.71	87.73	99.39	100	100	100
ERGAN-D	79.87	85.14	91.27	92.97	0	93.30	97.16	98.21	0	78.85	83.43	88.29	0	99.58	100	100
ERGAN-P	85.18	90.76	91.42	93.03	92.67	95.96	98.21	98.23	83.43	85.34	86.55	88.32	99.39	99.79	100	100
ERGAN	87.45	91.07	91.54	93.03	96.89	96.93	98.22	98.23	84.23	85.85	86.86	88.32	99.45	100	100	100

TABLE IV: Experimental results of f-measure under extremely limited real-labeled instances. The methods SVM, LR, XGBoost, DM, ERGAN-D, and ERGAN+WE have the f-measure value 0 in all these settings and are thus excluded from the table.

Dataset	Label Cost (#Instances)	Methods			
		SBC	ERNN	ERGAN-P	ERGAN
Cora	50	0	0.6648	0.7358	0.7735
	100	0	0.7684	0.7960	0.8083
	200	0.303	0.7742	0.8156	0.8314
	500	0.7629	0.8234	0.8493	0.8691
DBLP-ACM	50	0	0.6694	0.8492	0.8869
	100	0	0.7261	0.9143	0.9673
	200	0	0.7463	0.9151	0.9656
	500	0	0.8013	0.9174	0.9681
DBLP-Scholar	50	0	0.0043	0.6777	0.7760
	100	0	0.0536	0.7335	0.8045
	200	0	0.6869	0.7869	0.8124
	500	0	0.7903	0.8256	0.8372
NCVoter	50	0	0.3603	0.6389	0.7192
	100	0	0.8202	0.9091	0.9532
	200	0	0.9289	0.9431	0.9583
	500	0	0.9724	0.9740	0.9740

training, which can improve the label efficiency. Specifically, with 0.1% training, ERGAN-D fails to work (i.e., f-measure value is 0) on three datasets except for Cora. This is because ERGAN-D lacks the diversity module and can only randomly select instances for training. As a result, all training instances are selected from the majority class (non-matches), and accordingly no matched instance can be classified correctly by ERGAN-D, i.e. all the instances are classified as non-matches. Since datasets in ER applications are usually highly imbalanced, training data without diversity may hardly contain instances from the minority class (matches) when labels are limited, thus leading to poor performance.

Propagation module. Table III shows that ERGAN-P generally has better performance than ERNN and ERGAN-D, and thus it may affect the performance of ERGAN least compared with the other two key components: the adversarial learning architecture and the diversity module, especially when the label cost is small, e.g. 0.1% and 1% training. Additionally, when the label cost is 60%, the performance of ERGAN-P and ERGAN is the same. This is because instances with real labels in 60% training data can provide sufficient information for learning, and the propagation of instances with pseudo labels becomes unnecessary.

Observation 4. In ERGAN, all the three key components, i.e., the adversarial learning architecture, the diversity module and the propagation module, are necessary, each serving as an

integral part of the entire framework.

C. Extremeness Test

In the previous experiments, ERGAN has demonstrated strong and consistent performance when training data is reduced from 60% to 0.1%. However, a question left is: what are the minimum label costs required by ERGAN to achieve reasonable performance? To answer this, we conduct an experiment under extremely limited label cost, ranging from 50 to 500 instances with real labels. Table IV shows the results of our experiment. Note that we exclude the results of the methods SVM, LR, XGBoost, DM, ERNN, ERGAN-D and ERGAN+WE from this table because they fail to work when the label costs are below 500, i.e., f-measure value is 0 in all the settings in Table IV.

In Table IV, the results of SBC are almost 0 except for the cases when the label cost is 200 and 500 on Cora. This shows that SBC as a semi-supervised method can perform better than other fully supervised methods SVM, LR, XGBoost and DM under extremely limited label cost. Moreover, the performance of SBC is affected by the imbalance rate, and SBC fails to perform when the imbalance rate of a dataset is high.

We also notice that ERGAN can perform reasonably well on all the datasets even with 50 labels, and achieve good performance using only 500 labels. Moreover, ERNN, ERGAN-P and ERGAN have results in all the setting, but ERGAN-D does not have results in any of these settings. This is due to the diversity module, which can effectively select balanced training data to maximize the use of labels under extremely limited label cost. ERGAN-P performs better than ERNN on all the datasets. It indicates that performance may be harmed rather than helped by propagation if the quality of pseudo labels being propagated is not guaranteed. It is worthy to note that, for the cases of 50 and 100 labels on DBLP-Scholar dataset, we find that the reason why ERNN has very low feature values is because of high recall values (i.e., 0.749 for 50 labels and 0.777 for 100 labels) but low precision values (i.e., 0.002 for 50 labels and 0.028 for 100 labels). It means ERNN incorrectly classifies many non-matches as matches (i.e., false positives) and then propagates them into training data, leading to poor performance. With the increase in the label cost, this phenomenon is alleviated.

Observation 5. ERGAN can achieve good performance even with extremely limited labels. This is because the label generator G and the discriminator D are trained adversarially in ERGAN such that G uses the diversity module to balance the selection of instances from different classes and D propagates instances with high-quality pseudo labels into training.

VII. RELATED WORK

A. Entity Resolution

Entity Resolution (ER) has been extensively studied for decades since it was first reported in 1946 [11], [8], [9]. Traditionally, an ER task is performed through two stages (1) *blocking*, and (2) *matching*. Blocking aims to reduce the search space for record pair comparison which measures similarity of record pairs. Matching aims to determine whether record pairs refer the same real-world entity, and classification is the core problem in this stage. In this work, we focus on the classification problem of ER tasks.

Generally, two kinds of classification approaches are widely used in ER: rule-based and learning-based. Rule-based classification approaches [39], [13], often involve hand-crafted rules that associate with certain thresholds for defining similarity or dissimilarity of records [7], [34]. Learning-based approaches usually adopt a learning model to classify whether two records refer to the same entity. In the literature, there are three main categories: (1) Supervised learning approaches, which train a model to fit labeled instances so that it can predict unlabeled instances. The most recent work is *Magellan* [23], which considered learning models including *Decision Tree*, *Random Forest* and *Support Vector Machine* (SVM). Some work also studied ensemble learning approaches for ER to build a strong learned based on a set of weak learners [14]. A widely used ensemble classifier is *extreme gradient boosting* (XGBoost) [5], which uses the sparsity-aware algorithm and the weighted quantile sketch for approximate learning. (2) Unsupervised learning approaches, which take no real labeled instances, and assign labels to instances based on prior knowledge [3], [20]. A standard approach for ER is called *two-steps* (2S) [6]. It first labels a number (e.g. 10 percents of a dataset) of the most similar and dissimilar record pairs, respectively, and then trains a SVM in the second step. A most recent work in this line was proposed by Jurek et. al. [20], which considered both ensemble learning and automatic self-learning for classification based on training labels which are automatically generated based on different similarity measure schemes. A recent unsupervised approach for ER is proposed by Zhang et al. [40] which has two components: Iterative Term-Entity Ranking (ITER) and CliqueRank for record graph construction. (3) Semi-supervised learning approaches, which sit between supervised and unsupervised learning in that they take both limited real labeled and unlabeled instances for training. The state-of-the-art semi-supervised learning approach is an ensemble learning based approach using ensemble classifier Adaboost [30] for label prediction based on seed instances that have real labels.

B. Generative Adversarial Network

Generative adversarial network (GAN) was proposed by Goodfellow et. al. [17]. The key idea of GAN is that two networks, a *generator* and a *discriminator*, play a minimax game so that they converge gradually to an optimal solution. The generator aims to generate fake instances to "fool" the discriminator by simulating the distribution of real instances,

while the discriminator targets to distinguish fake instances (generated by the generator) from real instances. Due to the success of GAN in generating realistic images, a large number of studies have extended GAN to dealing with various tasks such as instance classification with semi-supervised learning [35], [31], labeled instance generation [26] and label generation [10]. Various techniques have also been proposed to improve GAN's performance by alleviating the mode collapse and convergence problems [31], [1].

Although GAN-based techniques are exploding, they cannot be directly used in solving ER tasks for three reasons: (1) ER datasets are often highly imbalanced, which aggravates the need of sufficient labeled training data, and may cause the mode collapse problem during the training process; (2) Most of the GAN-based approaches, including the ones designed for semi-supervised learning [35], have not considered the case of training with a extremely limited number of real labeled instances; (3) Traditionally, the generator in GANs is designed to generate new instances; however, for ER tasks, classifying all unlabeled instances is the ultimate goal. Until now, there is no existing GAN-based approach that can address all these problems. In this work, we build ERGAN to fill in this gap.

C. Deep Learning for Entity Resolution

In recent years, motivated by the success of deep learning techniques on tasks in computer vision, natural language processing, etc. [16], several attempts have been made to design deep learning solutions for ER tasks [27], [12]. Ebraheem et al. proposed DeepER, which uses bi-directional Recurrent Neural Networks (RNNs) with Long Short Term Memory (LSTM) units to learn a distributed representation for each record [12]. Mudgal et al. studied how to use deep learning techniques developed in natural language processing to handle the problems of attribute embedding, attribute summarization and attribute comparison [27]. A recent work proposed by Nie et al. [28] uses an align-compare-aggregate framework for a token level sequence-to-sequence ER which aims to solve the heterogeneous and dirty data problems. To deal with limited labels, several approaches take the advantages of transfer learning technique [41], [21]. However, the well-known limitations of transfer learning are that it needs a pre-trained model before applying to a target task, and a prior assumption on the correlation between the source and target tasks is also required, which restrict its practical applicability for ER problems in real-world applications.

To the best of our knowledge, this work is the first to explore the potentials of using GAN techniques to build powerful ER classifiers without prior knowledge. Our proposed method ERGAN can be synthesized with other deep learning techniques, including word embedding and transfer learning, and used as a key building block of an end-to-end deep learning solution.

VIII. CONCLUSION

In this paper, we have proposed a novel method, called ERGAN, to solve the ER classification problem with very limited labeled instances. ERGAN incorporates the diversity

of instances into sampling, prior to training the models. ERGAN consists of a label generator G to generate pseudo labels for unlabeled instances, and a discriminator D to distinguish instances with pseudo labels from instances with real labels. This method can be extended with word embedding for handling attribute values, leading to an enhanced method, called ERGAN+WE. We have formally proven that even with a limited number of instances with real labels, D and G in ERGAN can converge to the true distribution of instances and their labels. Our experimental results show that the performance of our methods beats all the baselines.

REFERENCES

- [1] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International Conference on Machine Learning (ICML)*, pages 214–223, 2017.
- [2] Sugato Basu, Arindam Banerjee, and Raymond Mooney. Semi-supervised clustering by seeding. In *International Conference on Machine Learning (ICML)*, 2002.
- [3] Indrajit Bhattacharya and Lise Getoor. A latent dirichlet model for unsupervised entity resolution. In *SIAM International Conference on Data Mining*, pages 47–58, 2006.
- [4] Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146, 2017.
- [5] Tianqi Chen and Carlos Guestrin. Xgboost: a scalable tree boosting system. In *international conference on Knowledge Discovery and Data Mining (SIGKDD)*, 2016.
- [6] Peter Christen. Automatic record linkage using seeded nearest neighbour and support vector machine classification. In *international conference on Knowledge Discovery and Data mining (SIGKDD)*, 2008.
- [7] Peter Christen. Development and user experiences of an open source data cleaning, deduplication and record linkage system. *ACM SIGKDD Explorations Newsletter*, 11(1):39–48, 2009.
- [8] Peter Christen. *Data matching: concepts and techniques for record linkage, entity resolution, and duplicate detection*. Springer Science & Business Media, 2012.
- [9] Vassilis Christophides, Vasilis Efthymiou, Themis Palpanas, George Papadakis, and Kostas Stefanidis. End-to-end entity resolution for big data: A survey. *arXiv preprint:1905.06397*, 2019.
- [10] Yue Deng, KaWai Chen, Yilin Shen, and Hongxia Jin. Adversarial active learning for sequences labeling and generation. In *IJCAI*, pages 4012–4018, 2018.
- [11] Halbert L Dunn. Record linkage. *American Journal of Public Health and the Nations Health*, 36(12):1412–1416, 1946.
- [12] Muhammad Ebraheem, Saravanan Thirumuruganathan, Shafiq Joty, Mourad Ouzzani, and Nan Tang. Distributed representations of tuples for entity resolution. *Proceedings of the VLDB Endowment*, 11(11):1454–1467, 2018.
- [13] Wenfei Fan, Xibei Jia, Jianzhong Li, and Shuai Ma. Reasoning about record matching rules. *Proceedings of the VLDB Endowment*, 2(1):407–418, 2009.
- [14] Yoav Freund. Boosting a weak learning algorithm by majority. *Information and computation*, 121(2):256–285, 1995.
- [15] Crispin W Gardiner et al. *Handbook of stochastic methods*, volume 3. springer Berlin, 1985.
- [16] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.
- [17] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems (NeurIPS)*, pages 2672–2680, 2014.
- [18] Yves Grandvalet and Yoshua Bengio. Semi-supervised learning by entropy minimization. In *Advances in neural information processing systems (NeurIPS)*, pages 529–536, 2005.
- [19] Anil K Jain. Data clustering: 50 years beyond k-means. *Pattern recognition letters*, 31(8):651–666, 2010.
- [20] Anna Jurek, Jun Hong, Yuan Chi, and Weiru Liu. A novel ensemble learning approach to unsupervised record linkage. *Information Systems*, 71:40–54, 2017.
- [21] Jungo Kasai, Kun Qian, Sairam Gurajada, Yunyao Li, and Lucian Popa. Low-resource deep entity resolution with transfer and active learning. *arXiv preprint arXiv:1906.08042*, 2019.
- [22] Mayank Kejriwal and Daniel P Miranker. Semi-supervised instance matching using boosted classifiers. In *European Semantic Web Conference*, pages 388–402. Springer, 2015.
- [23] Pradap Konda, Sanjib Das, AnHai Doan, Adel Ardalani, Jeffrey R Ballard, Han Li, Fatemah Panahi, Haojun Zhang, Jeff Naughton, and Shishir Prasad. Magellan: toward building entity matching management systems over data science stacks. *Proceedings of the VLDB Endowment*, 9(13):1581–1584, 2016.
- [24] Hanna Köpcke, Andreas Thor, and Erhard Rahm. Evaluation of entity resolution approaches on real-world match problems. *VLDB Endowment*, 3(1-2):484–493, 2010.
- [25] Xiang Li, Yao Wu, Martin Ester, Ben Kao, Xin Wang, and Yudian Zheng. Semi-supervised clustering in attributed heterogeneous information networks. In *International Conference on World Wide Web (WWW)*, pages 1621–1629, 2017.
- [26] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint:1411.1784*, 2014.
- [27] Sidharth Mudgal, Han Li, Theodoros Rekatsinas, AnHai Doan, Youngchoon Park, Ganesh Krishnan, Rohit Deep, Esteban Arcaute, and Vijay Raghavendra. Deep learning for entity matching: A design space exploration. In *Proceedings of the 2018 International Conference on Management of Data*, pages 19–34. ACM, 2018.
- [28] Hao Nie, Xianpei Han, Ben He, Le Sun, Bo Chen, Wei Zhang, Suhui Wu, and Hao Kong. Deep sequence-to-sequence entity matching for heterogeneous entity resolution. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, pages 629–638, 2019.
- [29] Dorian Pyle. *Data preparation for data mining*. morgan kaufmann, 1999.
- [30] Gunnar Rätsch, Takashi Onoda, and K-R Müller. Soft margins for adaboost. *Machine learning*, 42(3):287–320, 2001.
- [31] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Advances in neural information processing systems (NeurIPS)*, pages 2234–2242, 2016.
- [32] Jingyu Shao and Wang Qing. Active blocking scheme learning for entity resolution. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD)*, 2018.
- [33] Jingyu Shao, Qing Wang, and Fangbing Liu. Learning to sample: an active learning framework. In *International Conference on Data Mining (ICDM)*, 2019.
- [34] Rohit Singh, Venkata Vamsikrishna Meduri, Ahmed Elmagarmid, Samuel Madden, Paolo Papotti, Jorge-Arnulfo Quiané-Ruiz, Armando Solar-Lezama, and Nan Tang. Synthesizing entity matching rules by examples. *VLDB Endowment*, 11(2):189–202, 2017.
- [35] Jost Tobias Springenberg. Unsupervised and semi-supervised learning with categorical generative adversarial networks. *arXiv preprint:1511.06390*, 2015.
- [36] Isaac Triguero, Salvador García, and Francisco Herrera. Self-labeled techniques for semi-supervised learning: taxonomy, software and empirical study. *Knowledge and Information systems*, 42(2):245–284, 2015.
- [37] Qing Wang, Mingyuan Cui, and Huizhi Liang. Semantic-aware blocking for entity resolution. *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, 28(1):166–180, 2016.
- [38] Qing Wang, Dinusha Vatsalan, and Peter Christen. Efficient interactive training selection for large-scale entity resolution. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 562–573. Springer, 2015.
- [39] Steven Euijong Whang and Hector Garcia-Molina. Entity resolution with evolving rules. 2010.
- [40] Dongxiang Zhang, Long Guo, Xiangnan He, Jie Shao, Sai Wu, and Heng Tao Shen. A graph-theoretic fusion framework for unsupervised entity resolution. In *2018 IEEE 34th International Conference on Data Engineering (ICDE)*, pages 713–724. IEEE, 2018.
- [41] Chen Zhao and Yeye He. Auto-em: End-to-end fuzzy entity-matching using pre-trained deep models and transfer learning. In *The World Wide Web Conference*, pages 2413–2424, 2019.
- [42] Xiaojin Jerry Zhu. Semi-supervised learning literature survey. Technical report, University of Wisconsin-Madison Department of Computer Sciences, 2005.