

Progressive Adversarial Learning for Bootstrapping: A Case Study on Entity Set Expansion

Lingyong Yan^{1,3}, Xianpei Han^{1,2}, Le Sun^{1,2,*}

¹Chinese Information Processing Laboratory ²State Key Laboratory of Computer Science
Institute of Software, Chinese Academy of Sciences, Beijing, China

³University of Chinese Academy of Sciences, Beijing, China
{lingyong2014, xianpei, sunle}@iscas.ac.cn

Abstract

Bootstrapping has become the mainstream method for entity set expansion. Conventional bootstrapping methods mostly define the expansion boundary using seed-based distance metrics, which heavily depend on the quality of selected seeds and are hard to be adjusted due to the extremely sparse supervision. In this paper, we propose BootstrapGAN, a new learning method for bootstrapping which jointly models the bootstrapping process and the boundary learning process in a GAN framework. Specifically, the expansion boundaries of different bootstrapping iterations are learned via different discriminator networks; the bootstrapping network is the generator to generate new positive entities, and the discriminator networks identify the expansion boundaries by trying to distinguish the generated entities from known positive entities. By iteratively performing the above adversarial learning, the generator and the discriminators can reinforce each other and be progressively refined along the whole bootstrapping process. Experiments show that BootstrapGAN achieves the new state-of-the-art entity set expansion performance.

1 Introduction

Bootstrapping is a fundamental technique for entity set expansion (ESE). It starts from a few seed entities (e.g., {London, Beijing, Paris}) and iteratively extracts new entities in the target category (e.g., {Berlin, Moscow, Tokyo}) to expand the entity set, where new entities are often evaluated by their context similarities to seeds (e.g., sharing the same context pattern—“* is an important city”) (Riloff and Jones, 1999; Gupta and Manning, 2014; Yan et al., 2020a). During the above process, it is core to decide whether the new entities belong to the target category (within the expansion boundary) or not (outside the expansion boundary) (Shi et al., 2014; Gupta and Manning, 2014).

*Corresponding author.

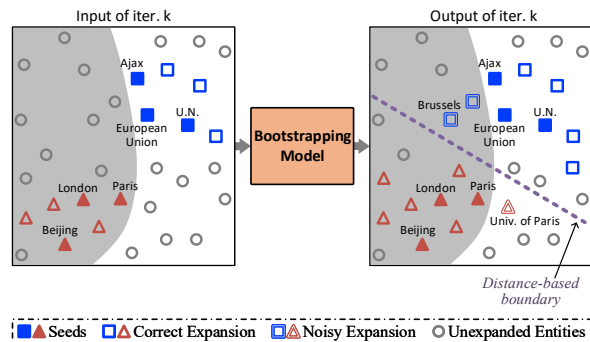


Figure 1: The expansion boundary problem of the bootstrapping technique. The areas in different background colors belong to different categories. Using the distance to positive entities will easily result in a bad expansion boundary at each iteration.

However, it is challenging to determine the expansion boundaries during the whole bootstrapping process, since only several seeds are used as the supervision at the beginning. Firstly, it is obviously not enough to define a good boundary using only several positive entities. For example, as shown in Figure 1, when only using several positive entities to learn distance-based boundaries, the boundaries are usually far from optimum, which in turn influences the quality of following bootstrapping iterations. Therefore, it is critical to enhancing the boundary learning with more supervision signals or prior knowledge (Thelen and Riloff, 2002; Curran et al., 2007). Secondly, bootstrapping is a dynamic process containing multiple iterations. Therefore, the boundary needs to be synchronously adjusted with the bootstrapping model, i.e., a good boundary should precisely restrict the current bootstrapping model from expanding negative entities.

Currently, most bootstrapping methods define expansion boundary using seed-based distance metrics, i.e., determining whether an entity should be expanded by comparing it with seeds. For instance, Riloff and Jones (1999); Gupta and Manning (2014); Batista et al. (2015) define the bound-

ary using pattern matching statistics or distributional similarities. Unfortunately, these heuristic metrics heavily depend on the selected seeds, making the boundary biased and unreliable (Curran et al., 2007; McIntosh and Curran, 2009). Although some studies extend them with extra constraints (Carlson et al., 2010) or manual participants (Berger et al., 2018), the requirement of expert knowledge makes them ad-hoc and inflexible. Some studies try to learn the distance metrics (Zupon et al., 2019; Yan et al., 2020a), but they still suffer from weak supervision. Furthermore, because the bootstrapping model and the boundary are mostly learned separately, it is hard for these methods to synchronously adjust the boundary when the bootstrapping model updates.

To address the boundary learning problem, we propose a new learning method for bootstrapping—BootstrapGAN, which defines expansion boundaries via learnable discriminator networks, and jointly models the bootstrapping process and the boundary learning process in the generative adversarial networks (GANs) framework (Goodfellow et al., 2014):

(1) Instead of using unified seed-based distance metrics, we define the expansion boundaries of different bootstrapping iterations using different learnable discriminator networks, where each of them directly determines whether an entity belongs to the same category of seeds at each iteration. By defining boundaries using discriminator networks, our method is flexible to use different classifiers and learnable using different algorithms.

(2) At each bootstrapping iteration, by modeling the bootstrapping network as the generator and adversarially learning it with a discriminator network, our method can effectively resolve the sparse supervision problem for boundary learning. Specifically, at each bootstrapping iteration, the generator is trained to select the most confusing entities; the discriminator learns to determine the selected entities as negative instances, and previously expanded entities and seeds as positive instances. In this way, the generator and the discriminator can reinforce each other: the generator can enhance supervision signals for discriminator learning by selecting latent noisy entities, and the discriminator can influence the generator to select more indistinguishable entities. When reaching the generator-discriminator equilibrium, the discriminator finally learns a good expansion boundary that accurately identifies new

entities, and the bootstrapping network can expand new positive entities within the boundaries.

(3) By iteratively performing the above adversarial learning process, the bootstrapping network and the expansion boundaries are progressively refined along bootstrapping iterations. Specifically, we use a discriminator sequence containing multiple discriminators to progressively learn expansion boundaries for different bootstrapping iterations. And the bootstrapping network is also refined and restricted along the whole bootstrapping process by the current discriminator and previously learned discriminators.

We conduct experiments over two datasets, and our BootstrapGAN achieves the new state-of-the-art performance for entity set expansion.

2 Progressive Adversarial Learning for Bootstrapping

In this section, we introduce our boundary learning method for bootstrapping models—BootstrapGAN (see Figure 2), which contains a generator—the bootstrapping network that performs the bootstrapping process, and a set of discriminators that determine the expansion boundaries for different bootstrapping iterations. The bootstrapping network and the discriminator networks are progressively and adversarially trained during the bootstrapping process.

2.1 Generator: Bootstrapping Network

The generator is the bootstrapping model, which iteratively selects new entities to expand seed sets.

We adopt the recently proposed end-to-end bootstrapping network—BootstrapNet (Yan et al., 2020a) as the generator, which follows the encoder-decoder architecture:

Encoder The encoder is a multi-layer graph neural network (GNN) that encodes the context features around entities/patterns into their embeddings. And the encoder takes an entity-pattern bipartite graph as input to efficiently capture global evidence (i.e., the direct and multi-hop co-occurrences between entities and patterns). The bipartite graph is constructed from original datasets: entities and patterns are graph nodes; an entity and a pattern are linked if they co-occur.

Based on the above bipartite graph, each GNN layer aggregates information from node neighbors as follows:

$$v_i^l = \sigma(f(W^l v_i^{l-1}, \sum_{j \in N(i)} a_{i,j}^l W^l v_j^{l-1})) \quad (1)$$

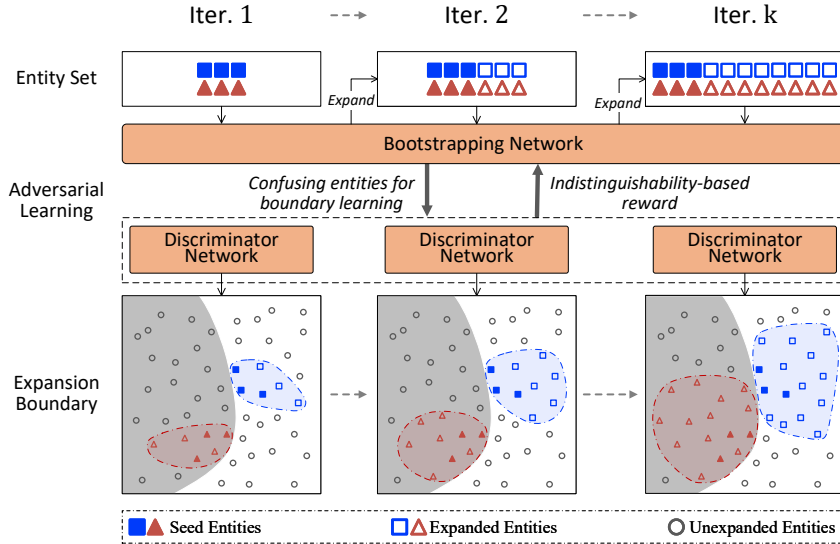


Figure 2: The overall framework of BootstrapGAN.

where v_i^l is node i 's embedding after layer l , $N(i)$ are i 's neighbors, W^l is the parameter matrix, $a_{i,j}^l$ is the attention-based weight, f is a linear sum function, and σ is the non-linear activation function.

Decoder After encoding entities and patterns, the GRU-based decoder sequentially generates new entities as the expansions, where each GRU step refers to one bootstrapping iteration. Specifically, the hidden state of the decoder represents the semantics of the target category. At each GRU step, the last expanded entities are used as the inputs to update the hidden state, which models the process that newly expanded entities are added to the current set, and therefore the set semantics should be updated (The first step inputs are seeds); then, the generating probabilities of a new entity are calculated as follows¹:

$$P_k(j) = \frac{\exp(v_j M^\top h^k)}{\sum_{j'} \exp(v_{j'} M^\top h^k)} \quad (2)$$

where h^k is the hidden state at k -th GRU step, v_j is entity j 's embedding outputted by the encoder, j' is a candidate entity, M is the parameter matrix. And top- N new entities are expanded at each step.

2.2 Discriminator: Expansion Boundary

Given positive entities (i.e., seeds and expanded entities), the discriminator defines the expansion boundary of each bootstrapping iteration by identifying whether a new entity is positive (i.e., be-

longing to the same category as positive entities) or negative (otherwise).

Instead of using seed-based distance metrics (Riloff and Jones, 1999; Gupta and Manning, 2014), we take different categories of seeds into consideration, and design the discriminators to directly predict which category a new entity belongs to. The motivation comes from two aspects: (1) By enforcing the discriminator directly discriminating whether a new entity is positive to any category of seeds, the discriminator can essentially possess the category boundary and is flexible to leverage more supervision signals except for seeds; (2) According to the mutual exclusive assumption (Curran et al., 2007) (i.e., *most entities usually belong to only one category*), it is better to leverage different categories of seeds to alleviate noises and simultaneously learn their expansion boundaries.

Specifically, we set our discriminator a multi-class classifier, which contains a GNN followed by an MLP layer: The GNN module takes the entity-pattern bipartite graph as input, and encodes context features into entity embeddings as Eq. 1; The MLP layer followed by a softmax function outputs the entity's category probabilities, where each category refers to one kind of seed set. And a new entity is only regarded as positive to the category with the highest probability. Besides, we set the GNN module as 1-layer to avoid model overfitting.

2.3 Progressive Adversarial Learning

To learn the above generator and discriminator, we design the following progressive adversarial learn-

¹This probability function is different from the original version of BootstrapNet that leverages cosine similarities.

ing process: Before bootstrapping, we pre-train the generator for better convergence (**Pre-training**); At each bootstrapping iteration, the discriminator is used to learn the expansion boundaries of this iteration, and is adversarially trained with the generator to reinforce each other (**Local adversarial learning**). Along the whole bootstrapping process, we progressively refine the generator with multiple discriminators by iteratively performing the above local adversarial learning (**Global progressive refining**).

2.3.1 Pre-training

Many previous studies have suggested that pre-training is important for learning convergence in GANs (Li and Ye, 2018; Qin et al., 2018). This paper pre-trains the generator (i.e., the bootstrapping network), and uses the following two kinds of pre-training algorithms: (1) The multi-view learning algorithm (Yan et al., 2020a), where the generator is co-trained with an auxiliary network. (2) Self-supervised and supervised pre-training using external resources (Yan et al., 2020b). Note that, since the external resources are not always accessible, we use the first algorithm as our default setting and set the second one as an alternative.

2.3.2 Local Adversarial Learning

At each bootstrapping iteration, the discriminator and the generator are learned using the following adversarial goals: the generator tries to generate new positive entities; the discriminator should distinguish new entities from current positive entities.

However, it is difficult to adopt standard GAN settings for our method: (1) The discriminator is a multi-class classifier rather than a binary classifier. (2) The generator outputs discrete entities rather than continuous values. To address the above issues, we use a Shannon entropy-based objective that is consistent with the discriminator, and the policy gradient algorithm to optimize the generator.

Shannon entropy-based learning objective To make our GAN settings consistent with the multi-class discriminator, we modify the adversarial goals inspired by Springenberg (2016): The generator tries to generate new entities that are certainly predicted as the same category as known positive entities by the discriminator; The discriminator tries to be not fooled by certainly assigning categories to the known positive entities and keeping uncertain about the class assignment for newly generated entities.

Based on the new goals, we design a Shannon entropy-based learning objective, where the category assignment uncertainty is represented by the Shannon entropy. Formally, at bootstrapping iteration k , we use the following adversarial objective to learn the generator G and the discriminator D :

$$\begin{aligned} \min_G \max_D & - \mathbb{E}_{e \sim S^c \cup G_{<k}^c} [H(p_D(c|e))] \\ & + \lambda \mathbb{E}_{e \sim S^c \cup G_{<k}^c} [CE(c, p_D(c|e))] \quad (3) \\ & + \mathbb{E}_{e' \sim G_k^c} [H(p_D(c|e'))] \end{aligned}$$

where c is a target category, S^c is the corresponding seed set, $G_{<k}^c$ is the set of expanded entities before iteration k , entities in $S^c \cup G_{<k}^c$ are regarded as positive entities, G_k^c is the set of newly expanded entities at step k , $H(p_D(c|e))$ is the discriminator prediction entropy for e , $CE(\cdot)$ is the cross-entropy term to assign right classes for positive entities, and λ is a hyper-parameter (this paper sets $\lambda = 1$). The first two terms of Eq. 3 aim to maximize the class assignment probabilities (i.e., minimizing the uncertainty) of positive entities, and the third term aims to maximize the entropies (i.e., maximizing the uncertainty) of newly generated entities.

And we sample the same size of newly generated entities as the positive entities to balance the above adversarial training process (We still select top- N entities for inference as Section 2.1).

Policy gradient learning for generator To optimize the generator that outputs discrete entities, we adopt the policy gradient algorithm. Specifically, we first rewrite the objective of the generator (the third term in Eq. 3) as maximizing the following function (denoted as L_G):

$$\begin{aligned} L_G & = \mathbb{E}_{e \sim G_k^c} [-H(p_D(c|e))] \\ & = \sum_{e \sim G_k^c} p_{G_k^c}(e) [-H(p_D(c|e))] \quad (4) \end{aligned}$$

where $p_{G_k^c}(e)$ is the expansion probability for entity e at step k , and e is a sampled discrete entity. We adopt the REINFORCE algorithm (Williams, 1992) to directly calculate L_G 's gradient $\nabla_{\theta} L_G$ as:

$$\begin{aligned} \nabla_{\theta} L_G & = \sum_{e' \sim G_k^c} p_{G_k^c}(e) \nabla_{\theta} \log p_{G_k^c}(e) R(e) \\ R(e) & = p_D(c|e) - b \quad (5) \end{aligned}$$

where $p_D(c|e)$ is the probability of e belonging to category c returned by the discriminator, $R(e)$ is the indistinguishability-base reward for generator

learning², b is the baseline value (This paper sets $b = \frac{1}{|C|}$, $|C|$ is the category number).

2.3.3 Global Progressive Refining

The local adversarial learning optimizes the generator and the discriminator at each bootstrapping iteration. This section describes how to refine them along the whole bootstrapping process—we call it global progressive refining.

One naive refining method is to iteratively perform the above local adversarial learning using one generator and one discriminator. However, this setting is not suitable for the dynamic bootstrapping process. Firstly, since the positive entities are iteratively expanded, the expansion boundaries at sibling iterations should also be slightly different. Therefore, it is necessary to use different discriminators for different iterations. Secondly, for the end-to-end bootstrapping network (Yan et al., 2020a), restricting the outputs of the current iteration will influence the outputs of previous iterations, but the naive refining method cannot continuously restrict the expansions of previous iterations to already learned boundaries.

Therefore, we propose a global progressive refining mechanism using a discriminator sequence containing multiple discriminators rather than one discriminator. Specifically:

(1). For each bootstrapping iteration, we use a unique discriminator to learn its expansion boundaries. That means for a total of K bootstrapping iterations, the discriminator sequence contains K different discriminators.

(2). At the k -th iteration, discriminator D_k is initialized by learned discriminator D_{k-1} ; then D_k and the generator G are trained using the local adversarial learning until coverage; finally, D_k can accurately define the expansion boundaries of iteration k and keeps fixed in the following iterations. Through the above process, we can progressively refine the expansion boundaries by iteratively fitting new discriminators from previously learned boundaries to new ones.

(3). At the k -th iteration, to restrict the generator’s previous expansion to the learned boundaries (possessed by $\{D_1, D_2, \dots, D_{k-1}\}$), we also use the learned discriminator D_i ($i \leq k$) to assign prediction probabilities as rewards for expanded

entities at iteration i . Finally, we replace the generator’s gradient calculated by Eq. 5 as:

$$\nabla_{\theta} L_G = \sum_{i=1}^k \sum_{e' \sim G_i^c} p_{G_i^c}(e) \nabla_{\theta} \log p_{G_i^c}(e) [p_{D_i}(c|e) - b] \quad (6)$$

where D_k is the discriminator to be learned at iteration k , and $\{D_1, \dots, D_{k-1}\}$ are already learned discriminators.

3 Experiments

3.1 Experimental Setup

Hyper-parameter	Value
Learning Rate	1e-4
Weight Decay	1e-3
Dropout Rate	0.1
Training Epoch per Iteration	10

Table 1: Main hyper-parameter settings.

Datasets The evaluation datasets we used are published by Zupon et al. (2019) and used by Yan et al. (2020a)—CoNLL and OntoNotes: The CoNLL contains 4 categories (5,522 entities), and the OntoNotes contains 11 categories (19,984 entities). For each category, 10 entities are used as the seeds, and all n -grams ($n \leq 4$) around candidate entities are defined as the context patterns.

Baselines We compare BootstrapGAN with the following baselines:

(1) Bootstrapping methods using heuristic seed-based distance metrics, including statistical metric—**Gupta** (Gupta and Manning, 2014), and lookahead search-based method—**LTB** (Yan et al., 2019);

(2) Bootstrapping methods using weakly-supervised learned boundaries, including custom embedding-based method—**Emboot** (Zupon et al., 2019), and end-to-end bootstrapping model learned by multi-view—**BootstrapNet** (Yan et al., 2020a), end-to-end bootstrapping model pre-trained using external datasets—**GBN** (Yan et al., 2020b).

For BootstrapGAN, we report the results of its two versions: BootstrapGAN, which uses the multi-view learning algorithm for pre-training; BootstrapGAN(ext), which uses external datasets for pre-training like Yan et al. (2020b).

Evaluation Metrics Following Zupon et al. (2019), this paper uses the precision-throughput curves to compare all methods. For further precise evaluation, we also report the precision@ K values (P@K, i.e., the precision at expansion step K).

²Maximizing the probability of one class still equals maximizing the minus entropy (i.e., indistinguishability), thus we use the probability for efficiency.




Figure 3: The precision-throughput curves on CoNLL and OntoNotes.

Method	CoNLL				OntoNotes			
	P@5	P@10	P@20	Mean	P@5	P@10	P@20	Mean
Gupta	70.4	63.4	61.3	65.0	-	-	-	-
LTB	78.5	71.0	62.2	70.6	42.2	36.6	32.3	37.0
Emboot	71.3	68.8	62.3	67.5	28.4	24.8	23.7	25.6
BootstrapNet	92.0	88.3	80.8	87.0	60.2	52.1	43.1	51.8
GBN	97.0	95.3	91.5	94.6	62.2	55.6	47.7	55.2
BootstrapGAN	98.7(±0.5)	94.8(±0.4)	86.4(±0.9)	93.3	63.0(±0.7)	57.1(±0.4)	48.9(±0.5)	56.3
- pre-training	98.0(±0.4)	94.5(±0.6)	87.1(±0.5)	93.2	54.8(±1.8)	49.1(±1.3)	44.0(±1.3)	49.3
BootstrapGAN(ext)	98.0(±0.5)	96.4(±0.5)	91.8(±0.9)	95.4	68.5(±0.7)	60.3(±0.5)	50.7(±0.4)	59.8

Table 2: The P@K values (%) of different bootstrapping models.

And we run our method for 10 repetitive training pieces and report the mean values of P@K as well as the standard deviations.

Implementation We implement the BootstrapGAN using the PyTorch (Paszke et al., 2019) with the PyTorch Geometric extension (Fey and Lenssen, 2019), and run it on a single Nvidia TiTan RTX GPU. And we use Adam (Kingma and Ba, 2015) and Rmsprop (Tieleman and Hinton, 2012) to respectively optimize the generator and the discriminators. Main hyperparameters are shown in Table 1. Our code is released at <https://www.github.com/lingyongyan/BootstrapGAN>.

3.2 Overall Evaluation Results

The precision-throughput curves of all methods are shown in Figure 3, and P@K values are also shown in Table 2. We can observe that:

(1) **Adversarial learning can effectively learn good expansion boundaries for bootstrapping models.** Comparing to all baselines without external resource pre-training (i.e., Gupta, LTB, Emboot, and BootstrapNet), BootstrapGAN achieves significant improvements (All p-values of t-test

evaluation are less than 0.01), and the precision-throughput curves of BootstrapGAN are the most smooth ones. That means more correct entities and less noisy entities are expanded at each iteration. It verifies that the learned expansion boundaries of BootstrapGAN contain fewer noisy entities than other methods, and therefore are the better boundaries. Besides, comparing to the baseline model using external resources for pre-training (i.e., GBN), the external resource pre-trained version—BootstrapGAN(ext) also outperforms it.

(2) **Progressive adversarial learning is complementary with self-supervised and supervised pre-training, and combining them can achieve the new state-of-the-art performance.** Comparing to the original BootstrapGAN, BootstrapGAN(ext), which combines self-supervised and supervised pre-training, achieves further improvements: On CoNLL, the P@10 and P@20 values achieve 1.6% and 5.4% improvements; On OntoNotes, the P@10 and P@20 values achieve 3.2% and 1.8% improvements.

(3) **The end-to-end bootstrapping paradigm outperforms other bootstrapping methods.** Comparing to other methods, the end-to-end learning methods (i.e., BootstrapNet, GBN and

Refining Strategy	CoNLL			OntoNotes		
	P@5	P@10	P@20	P@5	P@10	P@20
BootstrapGAN	98.7	94.8	86.4	63.0	57.1	48.9
- refining	93.1	83.0	73.1	56.4	50.7	43.2
- g-refining	95.2	92.6	87.0	63.0	56.4	48.5

Table 3: Performance comparison of BootstrapGAN with different refining mechanisms.

BootstrapGAN, BootstrapGAN(ext)) can achieve obviously higher performance. And comparing to the BootstrapNet/GBN, BootstrapGAN/BootstrapGAN(ext) can further achieve noticeable improvements, especially on the more complex dataset–OntoNotes.

3.3 Detail Analysis

Effect of pre-training strategies. To analyze the effects of pre-training, we compare the performance of BootstrapGAN using different pre-training settings (see Table 2): **BootstrapGAN**, and BootstrapGAN without pre-training (**- pre-train**). And we can see that: pre-training is an effective way to improve bootstrapping performance in some tasks. Without the pre-training, the BootstrapGAN’s performance on OntoNotes substantially drops—all mean P@K values decrease at least 4.9%. This may be because complex datasets (e.g., the OntoNotes) usually contain massive amounts of entities, and the search space of the bootstrapping network is extremely large, which makes it hard to converge to the optimum without appropriate pre-training.

Effect of global progressive refining. To analyze the effects of global progressive refining, we conduct the comparison experiments with different refining mechanisms (see Table 3): original settings using global progressive refining (**BootstrapGAN**); performing local adversarial learning without refining, i.e., only seeds are taken as positive entities, all expanded entities from different iterations are taken as negative ones in Eq. 3 (**-refining**); performing refining using the naive refining mechanism rather than our global progressive refining (**- g-refining**). From Table 3, we can see that:

(1) Refining is useful when performing adversarial learning for bootstrapping. Without the refining mechanism, the BootstrapGAN performance sharply drops on both datasets (All P@K values decrease by at least 5.6%).

(2) Our global progressive refining mechanism is very suitable for BootstrapGAN learning. By




Figure 4: Mean and standard deviation bands of P@K values of BootstrapGAN across different learning iterations.

replacing the global progressive refining with the naive mechanism, we can see that most BootstrapGAN performance results decrease, especially on P@5 and P@10. This verifies our observation that the expansion of previous iterations can be influenced when adversarially learning for later iterations. And our global progressive refining can well alleviate the influence, and therefore a better refining mechanism.

Stability of adversarial learning. To analyze the stability of our adversarial learning method, we report the P@K values of BootstrapGAN at different iterations (see Figure 4). We can see that: (1) Our adversarial learning method can coverage quickly. At around the 10th bootstrapping iteration, the performance of BootstrapGAN reaches a reasonable level. (2) Our adversarial learning method is stable. On both datasets, most P@K values steadily increase with more training iterations, and the standard deviations of most P@K values progressively decrease. Those can verify the stability of our learning algorithm (Although some P@K values decrease a little from iteration 10 to iteration 20, we still consider our algorithm stable since the differences are slight enough to be omitted).

Iter.	BootstrapNet	BootstrapGAN
1	the States, <i>Asia Development Corp.</i> , <i>Atlantic</i> , <i>North China Area Army</i> , <i>Continent</i> , <i>the Gulf of Mexico</i> , <i>Mediterranean</i> , <i>Scandinavia</i> , <i>the East Coast</i> , <i>Bank of China</i>	Mexico, Poland, Romania, Denmark, Moscow, The Netherlands, Vienna, Hungary, Greece, Bulgaria
10	Indonesia, Somalia, <i>Northern German</i> , <i>the Convention on Trade in Endangered Species</i> , ..., <i>the Western Hemisphere</i> , <i>the Asia Pacific region</i>	Melbourne, Havana, Tajikstan, Lausanne, Tehran, Abidjan, Thailand, Bahrain, Aqaba, <i>the Shaanxi International Exhibition Center</i>
20	the Republic of Iraq, <i>the United Nations World Human Rights Convention</i> , <i>Arab - Israelis</i> , ..., <i>Budget Group</i> , <i>Sino - Kirghizian</i>	Adelaide, Rangoon, Cologne, Madrid, Phnom Penh, Jinan, Karachi, Palermo, <i>Baghdad Airport</i> , <i>the North Pole</i>

Table 4: The examples of expanded GPEs using BootstrapNet and BootstrapGAN (Entities in red are noisy entities).

Examples for learned expansion boundaries.

To intuitively show the quality of learned expansion boundaries by BootstrapGAN, we show a typical case of different expanded entities for GPE (geopolitical entities) on the OntoNotes using BootstrapNet and BootstrapGAN (see Table 4)³. And we can see that BootstrapGAN can expand more correct entities, and most of them are tightly related to the GPE semantics; while the expansion boundaries of BootstrapNet contain many noisy entities at the very beginning and tend to introduce more noises at later iterations. This further verifies the importance of expansion boundary learning and BootstrapGAN’s effectiveness.

4 Related Work

Bootstrapping Bootstrapping is a widely used technique for information extraction (Riloff, 1996; Ravichandran and Hovy, 2002; Yoshida et al., 2010; Angeli et al., 2015; Saha et al., 2017), and also benefits many other NLP tasks, like question answering (Ravichandran and Hovy, 2002), named entity translation (Lee and Hwang, 2013), knowledge base population (Angeli et al., 2015), etc. To address the expansion boundary problem, most early methods (Riloff, 1996; Riloff and Jones, 1999) heuristically decide boundaries using pattern-matching statistics, but often result in a rapid quality degrading, which is known as the semantic drifting (Curran et al., 2007). To reduce semantic drifting, some studies leverage external resources or constraints, e.g., mutual exclusive constraints (Yan-garber et al., 2002; Thelen and Riloff, 2002; Curran et al., 2007; Carlson et al., 2010), lexical and statistical features (Gupta and Manning, 2014), look-ahead feedbacks (Yan et al., 2019), manually defined patterns (Zhang et al., 2020). However, those heuristic constraints are usually not flexible due

to their requirement for expert efforts. In contrast, recent studies focus on learning the distance metrics to determine boundaries using weak supervision (Gupta and Manning, 2015; Berger et al., 2018; Zupon et al., 2019; Yan et al., 2020a). For example, Yan et al. (2020a) propose an end-to-end bootstrapping network learned by multi-view learning, and extend it by self-supervised and supervised pre-training (Yan et al., 2020b). However, these methods usually learn a loose boundary using sparse supervision. Furthermore, these methods’ boundary learning process and model learning process are usually separately performed and therefore fail to be adjusted synchronously.

Adversarial Learning in NLP Adversarial learning (Goodfellow et al., 2014) is widely applied in NLP. For example, in sequential generation tasks, GAN is mainly used to alleviate the problem of lacking explicitly defined criteria (Yu et al., 2017; Lin et al., 2017; Yang et al., 2018). GAN has also been used in weakly supervised information extraction to identify informative instances and filter out noises (Qin et al., 2018; Wang et al., 2019), which inspires our method.

5 Conclusion

Due to very sparse supervision and the dynamic nature, one fundamental challenge of bootstrapping is how to learn precise expansion boundaries. In this paper, we propose an effective learning method for bootstrapping—BootstrapGAN, which defines expansion boundaries via learnable discriminator networks and jointly models the bootstrapping process and the boundary learning process in the GANs framework. Experimental results show that, by adversarially learning and progressively refining the bootstrapping network and the discriminator networks, our method achieves the new state-of-the-art performance. In the future, we plan to leverage extra knowledge (e.g., knowledge graph) to improve bootstrapping learning.

³The seeds are {Washington, New York, the United States, Russia, Iran, Hong Kong, France, London, California, China}.

Acknowledgments

This work is supported by the National Natural Science Foundation of China under Grants no. U1936207 and 61772505, Beijing Academy of Artificial Intelligence (BAAI2019QN0502), and in part by the Youth Innovation Promotion Association CAS(2018141).

References

- Gabor Angeli, Victor Zhong, Danqi Chen, Arun Tejasvi Chaganty, Jason Bolton, Melvin Jose Johnson Premkumar, Panupong Pasupat, Sonal Gupta, and Christopher D Manning. 2015. Bootstrapped self training for knowledge base population. In *TAC*.
- David S. Batista, Bruno Martins, and Mário J. Silva. 2015. Semi-supervised bootstrapping of relationship extractors with distributional semantics. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 499–504, Lisbon, Portugal. Association for Computational Linguistics.
- Matthew Berger, Ajay Nagesh, Joshua Levine, Mihai Surdeanu, and Helen Zhang. 2018. Visual supervision in bootstrapped information extraction. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2043–2053, Brussels, Belgium. Association for Computational Linguistics.
- Andrew Carlson, Justin Betteridge, Richard C. Wang, Estevam R. Hruschka Jr., and Tom M. Mitchell. 2010. Coupled semi-supervised learning for information extraction. In *Proceedings of the Third International Conference on Web Search and Web Data Mining*, pages 101–110. ACM.
- James R. Curran, Tara Murphy, and Bernhard Scholz. 2007. Minimising semantic drift with mutual exclusion bootstrapping. In *PAFLING*, pages 172–180.
- Matthias Fey and Jan E. Lenssen. 2019. Fast graph representation learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*.
- Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014*, pages 2672–2680.
- Sonal Gupta and Christopher Manning. 2014. Improved pattern learning for bootstrapped entity extraction. In *Proceedings of the Eighteenth Conference on Computational Natural Language Learning*, pages 98–108, Ann Arbor, Michigan. Association for Computational Linguistics.
- Sonal Gupta and Christopher D. Manning. 2015. Distributed representations of words to guide bootstrapped entity classifiers. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1215–1220, Denver, Colorado. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations*.
- Taesung Lee and Seung-won Hwang. 2013. Bootstrapping entity translation on weakly comparable corpora. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 631–640, Sofia, Bulgaria. Association for Computational Linguistics.
- Yan Li and Jieping Ye. 2018. Learning adversarial networks for semi-supervised text classification via policy gradient. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1715–1723. ACM.
- Kevin Lin, Dianqi Li, Xiaodong He, Ming-Ting Sun, and Zhengyou Zhang. 2017. Adversarial ranking for language generation. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017*, pages 3155–3165.
- Tara McIntosh and James R. Curran. 2009. Reducing semantic drift with bagging and distributional similarity. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 396–404, Suntec, Singapore. Association for Computational Linguistics.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019*, pages 8024–8035.
- Pengda Qin, Weiran Xu, and William Yang Wang. 2018. DSGAN: Generative adversarial training for distant supervision relation extraction. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 496–505, Melbourne, Australia. Association for Computational Linguistics.
- Deepak Ravichandran and Eduard Hovy. 2002. Learning surface text patterns for a question answering system. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*,

- pages 41–47, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Ellen Riloff. 1996. Automatically Generating Extraction Patterns from Untagged Text. In *AAAI*, pages 1044–1049.
- Ellen Riloff and Rosie Jones. 1999. Learning dictionaries for information extraction by multi-level bootstrapping. In *AAAI/IAAI*, pages 474–479.
- Swarnadeep Saha, Harinder Pal, and Mausam. 2017. [Bootstrapping for numerical open IE](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 317–323, Vancouver, Canada. Association for Computational Linguistics.
- Bei Shi, Zhenzhong Zhang, Le Sun, and Xianpei Han. 2014. [A probabilistic co-bootstrapping method for entity set expansion](#). In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 2280–2290, Dublin, Ireland. Dublin City University and Association for Computational Linguistics.
- Jost Tobias Springenberg. 2016. [Unsupervised and semi-supervised learning with categorical generative adversarial networks](#). In *4th International Conference on Learning Representations*.
- Michael Thelen and Ellen Riloff. 2002. [A bootstrapping method for learning semantic lexicons using extraction pattern contexts](#). In *Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing*, pages 214–221. Association for Computational Linguistics.
- Tijmen Tieleman and Geoffrey Hinton. 2012. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural networks for machine learning*, 4(2):26–31.
- Xiaozhi Wang, Xu Han, Zhiyuan Liu, Maosong Sun, and Peng Li. 2019. [Adversarial training for weakly supervised event detection](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 998–1008, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ronald J. Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. In *Reinforcement Learning*, pages 5–32.
- Lingyong Yan, Xianpei Han, Ben He, and Le Sun. 2020a. [End-to-End Bootstrapping Neural Network for Entity Set Expansion](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 9402–9409.
- Lingyong Yan, Xianpei Han, Ben He, and Le Sun. 2020b. [Global bootstrapping neural network for entity set expansion](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3705–3714, Online. Association for Computational Linguistics.
- Lingyong Yan, Xianpei Han, Le Sun, and Ben He. 2019. [Learning to bootstrap for entity set expansion](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 292–301, Hong Kong, China. Association for Computational Linguistics.
- Zhen Yang, Wei Chen, Feng Wang, and Bo Xu. 2018. [Improving neural machine translation with conditional sequence generative adversarial nets](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1346–1355, New Orleans, Louisiana. Association for Computational Linguistics.
- Roman Yangarber, Winston Lin, and Ralph Grishman. 2002. [Unsupervised learning of generalized names](#). In *COLING 2002: The 19th International Conference on Computational Linguistics*.
- Minoru Yoshida, Masaki Ikeda, Shingo Ono, Issei Sato, and Hiroshi Nakagawa. 2010. [Person name disambiguation by bootstrapping](#). In *Proceeding of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 10–17. ACM.
- Lantao Yu, Weinan Zhang, Jun Wang, and Yong Yu. 2017. [Seqgan: Sequence generative adversarial nets with policy gradient](#). In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, pages 2852–2858. AAAI Press.
- Yunyi Zhang, Jiaming Shen, Jingbo Shang, and Jiawei Han. 2020. [Empower entity set expansion via language model probing](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8151–8160, Online. Association for Computational Linguistics.
- Andrew Zupon, Maria Alexeeva, Marco Valenzuela-Escárcega, Ajay Nagesh, and Mihai Surdeanu. 2019. [Lightly-supervised representation learning with global interpretability](#). In *Proceedings of the Third Workshop on Structured Prediction for NLP*, pages 18–28, Minneapolis, Minnesota. Association for Computational Linguistics.