

Efficient Entity Candidate Generation for Low-Resource Languages

Alberto García-Durán, Akhil Arora, Robert West

EPFL, Switzerland

agaduran@gmail.com

{akhil.arora, robert.west}@epfl.ch

Abstract

Candidate generation is a crucial module in entity linking. It also plays a key role in multiple NLP tasks that have been proven to beneficially leverage knowledge bases. Nevertheless, it has often been overlooked in the monolingual English entity linking literature, as naïve approaches obtain very good performance. Unfortunately, the existing approaches for English cannot be successfully transferred to poorly resourced languages. This paper constitutes an in-depth analysis of the candidate generation problem in the context of cross-lingual entity linking with a focus on low-resource languages. Among other contributions, we point out limitations in the evaluation conducted in previous works. We introduce a characterization of queries into types based on their difficulty, which improves the interpretability of the performance of different methods. We also propose a *light-weight* and *simple* solution based on the construction of indexes whose design is motivated by more complex transfer learning based neural approaches. A thorough empirical analysis on 9 real-world datasets under 2 evaluation settings shows that our simple solution outperforms the state-of-the-art approach in terms of both quality and efficiency for almost all datasets and query types.

Keywords: Candidate generation, low-resource languages, evaluation pitfalls, efficiency, light-weight models

1. Introduction

Candidate generation (Shen et al., 2014; Zhou et al., 2020) is the task of retrieving a short list of plausible candidate entities from a knowledge base (KB) for a given mention. This module is usually a fundamental part of a more elaborate pipeline as that of entity linking (Guo et al., 2013), coreference resolution (Zhang et al., 2019) or question answering (Bao et al., 2016). With few exceptions (Gillick et al., 2019), it is present in almost all existing entity linking techniques, where the subsequent entity disambiguation technique only has to evaluate a small set of entities from the KB, instead of all of them. Therefore, in addition to obtaining a high recall—fraction of mentions where the true entity is present in the set of candidates—a desired characteristic of a candidate generator is to constitute a low computational overhead of the full pipeline. Cross-lingual entity linking (Tsai and Roth, 2016; Sil et al., 2018) associates mentions in non-English documents to entities from a KB—typically the English Wikipedia. This poses several challenges that can often be attributed to the amount of resources of the target language, especially if it has low resources. Candidate generation is deemed trivial in the English entity linking literature, as there exist resources (*i.e.* indexes/lookup tables (Spitkovsky and Chang, 2012)) capable of providing high candidate recall. However, this is not the case in the cross-lingual setting. This is specially true when dealing with low-resource languages (*e.g.*, karelian or turkmen), wherein, moreover, the number of works that tackle this important step is very scarce. Some recent approaches (Rijhwani et al., 2019; Zhou et al., 2020) are based on the idea of performing transfer learning from a pivot language that is lexically similar to the target language.

Contributions. In this paper we propose a *light-weight* solution, called pivot token indexes (PTI), to the candidate generation problem with a special focus on low-resource languages. PTI combines prior and posterior probabilities learned by aggregating information from the pivot and target language. These probabilities are constructed with tokens the mentions can be decomposed into. This simple solution keeps the efficiency of standard indexes—thereby possessing the desired property of a candidate generator—while being equipped with the sophistication of more complex solutions (Zhou et al., 2020). Most importantly, the main contribution of this work is an in-depth analysis of the candidate generation problem in the context of entity linking, to which we make several contributions: i) we establish that the (inherited) practice of ignoring mentions that are not linkable to the English Wikipedia leads to an unrealistic evaluation; ii) we perform a realistic evaluation and propose an efficient trick to increase the ceiling recall of *any* candidate generator; iii) we introduce three query types according to their difficulty based on the available training data; iv) we compare the performance of the different methods and establish differences across them based on the amount of resources of the pivot language and the query type. Moreover, our experimental section includes experiments on 9 datasets and two different learning settings (without and with some supervision from the target language). We also empirically analyze the complexity of PTI and conclude that it takes up much less memory, and requires significantly less training and inference time than the current state-of-the-art (Zhou et al., 2020).

Other applications of candidate generators. In general, tasks that can beneficially leverage the knowledge contained in KBs may benefit from better candidate

generators. Examples of such tasks are coreference resolution (Zhang et al., 2019) or question answering (Bao et al., 2016), to name but a few. These works retrieve facts from a KB that are related to a given mention. As in the case of entity linking, candidate generation may alleviate the computational cost of the subsequent modules by just keeping a reduced set of facts.

2. Background

Let m be a mention observed in a document, and let e be the ground truth entity associated to that mention. The task of a candidate generator \mathcal{CG}^K is to retrieve a list of K possible candidate entities such that $e \in \mathcal{CG}^K(m)$. The candidate generator is trained with mention entity pairs observed in a training corpus $\mathcal{C}_{\text{LANG}}$, where the mentions are written in a certain language LANG. At test time, candidate generators provide a score to all the entities that form the candidate space, and find the K entities with the largest scores—this is a linear operation (done with a max-heap structure) with respect to the candidate space.

A candidate generator must have good recall, and also be much more computationally efficient as compared to the more complex downstream model. See (Ling et al., 2015) for a detailed discussion. We use EN, PL and TL to indicate English, pivot and target language, respectively.

3. Related Work

While candidate generation in monolingual (English) entity linking is deemed a trivial task given the existence of high-quality mention-entity lookup tables, most of other languages—specially low-resource languages—do not benefit from lookup tables of a similar quality. Therefore, contrary to (English) entity linking, candidate generation constitutes a challenge in the cross-lingual entity linking literature.

3.1. Wikipedia-based Indexes

This is the *de facto* standard for candidate generation in the entity linking literature (Sil and Florian, 2016; Yamada et al., 2017; Sil et al., 2018; Upadhyay et al., 2018a). An index of prior probabilities $P(\text{entity}|\text{mention})$ is constructed from a corpus. While this solution has been shown to provide a good recall for languages where the training corpus is large, for low-resource languages the training corpus may offer a very poor coverage in terms of both mentions (Fu et al., 2020) and entities.

3.2. Pivoting-based Embedding Approaches

Methods within this category leverage information of a pivot language that is lexically related to the target language: the candidate generator is trained with data from the pivot language, but then it is evaluated in the target language. These methods rely on a certain lexical overlap between pairs of related languages (*e.g.* Italian and Corsican) that allow to transfer learning between

the pivot—which is supposed to have more resources than the target language—and the target language. This is opposed to other techniques (Gollins and Sander-son, 2001) that rely on the existence of accurate machine translation models to map all the resources to a common language. Unfortunately, the performance of translation models is also largely affected by the scarce data that is available for low-resource languages. Similar ideas have been applied to other problems such as POS tagging (Fang and Cohn, 2017) or machine translation (Zoph et al., 2016).

The learning objective consists of maximizing a similarity function *sim* (the cosine similarity) between mention \mathbf{m} and entity e embedding pairs:

$$\mathcal{L}_{\text{PL}} = \sum_{(m,e) \in \mathcal{C}_{\text{PL}}} \max(0, 1 - \text{sim}(\mathbf{m}, \mathbf{e}^{\text{EN}^+}) + \text{sim}(\mathbf{m}, \mathbf{e}^{\text{EN}^-})), \quad (1)$$

where \mathcal{C}_{PL} is the collection of training mention entity pairs in the pivot language, and $\mathbf{m} = \text{Enc}(\text{tknr}(m))$ is the mention embedding obtained by applying an encoding function Enc to the output of a tokenizer $\text{tknr}(m)$. A similar architecture is used to obtain the entity embedding from its English name. The superscript $+$ or $-$ indicates whether the entity is the target or a randomly sampled one.

Previous works have explored different tokenizers (*e.g.* character-level (Rijhwani et al., 2019) or character n -grams (Zhou et al., 2020)) and encoding functions (*e.g.* LSTM- (Rijhwani et al., 2019) or CNN-based architectures (Zhang et al., 2015)). The best setup (Zhou et al., 2020) consists of a character n -gram tokenizer, followed by an encoder that maps each token to an embedding (via an embedding lookup table), sums them all up, and applies a nonlinear activation function to the aggregated representation. This approach, called CHARAGRAM (Wieting et al., 2016), has shown to also outperform other type of approaches based on indexes and translation (Pan et al., 2017).

Additional Related Work. There are related tasks such as transliteration (Upadhyay et al., 2018b) and bilingual lexical induction (Haghighi et al., 2008). The former is a harder task, as it consists of generating, instead of retrieving, English entity names. The latter relies to a great extent on the existence of large monolingual corpus in both languages (Vulic et al., 2020).

4. PIVOT TOKEN INDEXES

Intuitively, CHARAGRAM learns to associate tokens that co-occur between mentions and English entity names. The learned associations between tokens determine the similarity (score) between entity and mention embeddings. This approach has been shown to provide good recall but both memory requirements and training/inference time are considerable (an empirical analysis is in Section 6.1). Thus, it is arguable whether

CHARAGRAM fulfills a desired characteristic of a candidate generator, namely its computational overhead is low. Similarly to CHARAGRAM, our approach also learns associations between mentions and entities. Different to CHARAGRAM, the associations learned by PTI are meant to indicate the attachment strength between tokens and entities, and not equivalences between tokens. PTI consists of two indexes—then being high efficient—that are linearly combined to score entities. The indexes correspond to the prior probability that associates tokens to entities, and the posterior probability that associates entities to tokens. The goal of the posterior probability is to counterbalance the (low) score that unpopular entities obtain for popular tokens from the prior index. PTI has the following steps:

1. We apply a tokenizer tknr to the collection of mentions observed in the training corpus of the pivot language \mathcal{C}_{PL} . Each mention m , linked to a certain entity, is decomposed into a set of n_m tokens $\text{tknr}(m) = \{\text{token}_1, \text{token}_2, \dots, \text{token}_{n_m}\}$.
2. We construct prior $P(\text{entity}|\text{token})$ and posterior $P(\text{token}|\text{entity})$ token-based probabilities. Optionally, we can remove low probabilities by thresholding (Section 6.1) the indexes without an impact in the performance. The thresholding significantly reduces the memory requirements of PTI.
3. At test time, given a mention m , each entity ent obtains an score as follows:

$$\begin{aligned} & \sum_{\text{tok} \in \text{tknr}(m)} P(\text{ent}|\text{tok}) + \lambda P(\text{tok}|\text{ent}) \\ &= \sum_{\text{tok} \in \text{tknr}(m)} P(\text{ent}|\text{tok}) \left(1 + \lambda \frac{P(\text{tok})}{P(\text{ent})}\right), \quad (2) \end{aligned}$$

where λ is a weighting term. As argued, the posterior probability term may counterbalance the bias introduced by the prior probability towards the popular entities: a large value of λ will favor unpopular entities. While these scores do not have a probabilistic interpretation, one can easily normalize them so that they sum up to 1. As opposed to CHARAGRAM, PTI performs a search that only includes the entities that have a non-zero score; Then contributing towards a better efficiency.

We will see in Section 6 that this simple approach has the following advantages as compared to CHARAGRAM: fewer parameters, less training time, less inference time, and better performance in almost all settings.

4.1. Joint learning with the target language

In practice, there is often some training data that is available for the target language. To incorporate this data into the construction of the prior and posterior

probabilities we modify step 2 by weighting the counts over the pivot language corpus with a factor α . Thus, the index of priors is computed as

$$\begin{aligned} P(\text{entity}|\text{token}) &= \\ &= \frac{\#\text{TL}(\text{entity}|\text{token}) + \alpha \#\text{PL}(\text{entity}|\text{token})}{\sum_{e \in \mathcal{E}} (\#\text{TL}(e|\text{token}) + \alpha \#\text{PL}(e|\text{token}))}, \end{aligned}$$

where $\#\text{LANG}(\cdot)$ is a shortcut that refers to the number of times the conditional argument is observed in the corpus $\mathcal{C}_{\text{LANG}}$. \mathcal{E} corresponds to the set of entities observed in the corpora. The same weighting is applied to compute the posterior probabilities. At test time, entities are scored following Eq. (2).

More complex versions of our approach PTI are in Appendix A.5. While these versions led to minor improvements, their complexity also increased.

5. Experimental Setup

All the resources required to reproduce the experiments are available at <https://github.com/epfl-dlab/pti-candgen>.

In order to comprehensively understand the performance of PTI we run experiments on a variety of publicly available datasets. The benchmarked methods are compared in terms of accuracy and complexity (number of parameters and training/inference time). The accuracy is measured by the recall, which is defined as the proportion of retrieved candidate lists containing the correct entity. More formally, the recall@ K is defined as

$$\text{recall}@K = \frac{\sum_{(m,e) \in \mathcal{T}} \delta(e \in \mathcal{CG}^K(m))}{|\mathcal{T}|} * 100$$

where \mathcal{T} is the set of mentions entity pairs in an evaluation set, and $\delta(\cdot)$ returns 1 if its argument is true, and 0 otherwise. As in previous work (Zhou et al., 2020) we set $K = 30$ for our main set of experiments. However, for a subset of the experiments we include a comparative performance for other values of K in Appendix A.4.

We compare the performance of the different methods in the two following learning settings:

- **Zero-shot.** We use zero resources—no training data—from the target language. Therefore, the candidate generator must fully rely on the available information for the pivot language. While it is possible to obtain a (small) training set of mention entity pairs for some target languages, in some other cases the target language is so poorly resourced that this is the only possible learning setting.
- **Joint learning with the target language.** Even for low-resource languages, there is often some training data for the target language that can be used in conjunction with the data that is available for the pivot language. We sometimes refer to this setting as supervised.




Figure 1: Query types according to their difficulty. The target and pivot language correspond to corsican (CO) and italian (IT), respectively. The backgrounds of the queries (*Mention : Ground Truth Entity*) indicate their type (EASY=green, MEDIUM=orange, and HARD=red).

We note that the aim of this section is to establish a detailed comparison of different candidate generation methods. As any progress in the candidate generator is independent from any improvement in the entity disambiguation module, the integration of the candidate generation in the entity linking pipeline is out of the scope of this work.

5.1. Query Types

As discussed in Section 5, despite the little resources that are available for the target languages, sometimes it is possible to obtain some training data. This training data will determine the difficulty of the queries contained in the test set. Given a mention query m that links to an entity e , we classify it according to our proposed categorization (see Figure 1 for a visual explanation):

- EASY. The mention m is observed with the entity e in the training data of the target language.
- MEDIUM. The entity e is observed in the training data of the target language but with mentions other than m .
- HARD. The entity e is never observed in the training data of the target language. To successfully link these mentions we fully rely on the pivot language. *In the zero-shot setting, all mention queries belong to this type.*

5.2. Baselines

- WIKIPRIORS. Prior probabilities (*i.e.* $P(\text{entity}|\text{mention})$) are estimated using a training corpus obtained from both a target and a pivot language, and constructs indexes mapping mentions to their co-occurring entities. This simple procedure resembles the approach followed by (Spitkovsky and Chang, 2012), where a lookup table was built for entities that have an English Wikipedia page. The strength of a mention-entity association is given by the estimated prior probabilities. In the supervised setting, it first generates candidates using the index constructed on the target language. More details are provided in Appendix A.2.
- CHARAGRAM (Zhou et al., 2020). This is the state-of-the-art on candidate generation for (low-resource) cross-lingual settings. It has shown to

significantly outperform other architectures such as (Rijhwani et al., 2019), and translation-based methods (Pan et al., 2017), among others. For the supervised setting, we extend CHARAGRAM to also incorporate training data from the target language by optimizing the loss $\mathcal{L}_{TL} + \mu\mathcal{L}_{PL}$. More details are provided in Section 3.2.

5.3. Datasets

As previous works (Rijhwani et al., 2019; Zhou et al., 2020), we create our datasets from Wikipedia. We choose Wikipedia as, contrary to other datasets (*e.g.* DARPA LORELEI (Strassel and Tracey, 2016)), it can be accessed without any restriction, and, importantly, its crowdsourcing spirit guarantees a realistic setting.

We select 9 pairs of languages, where the target language is always (extremely) low-resourced, and the pivot language can be categorized as either high, medium, or low-resourced. The distinction regarding the amount of resources of the pivot language has never done in previous works, but it helps to analyze the performance of the different methods. The 9 language pairs (TARGET-PIVOT) are Corsican-Italian (CO-IT), Limburgish-Dutch (LI-NL), Bavarian-German (BAR-DE), Karelian-Finnish (OLO-FI), Macedonian-Bulgarian (MK-BG), Javanese-Indonesian (JV-ID), Turkmen-Azerbaijani (TK-AZ), Marathi-Hindi (MR-HI) and Faroese-Icelandic (FO-IS). Appendix A.1 contains information about the dataset statistics (Table 4), additional preprocessing details and detailed information about the selection of the pivot languages.

Except for cases when zero-shot is the only possible learning setting (CO-IT, OLO-FI and TK-AZ) we create the validation and test sets by sampling a maximum of 1,000 unique queries of each query type—in total it amounts to a maximum of 3,000 queries; and the training data is adapted accordingly. The same validation and test sets are used in the zero-shot setting. For target languages that can only be evaluated in the zero-shot setting, we equally distribute the available mention entity pairs into validation and test. These sets never contain duplicates.

5.4. Towards a Realistic Evaluation

A common practice in the (English) entity linking literature (Yamada et al., 2017; Ganea and Hofmann, 2017) consists of ignoring mentions whose correct entity does not exist in the considered KB—typically the English Wikipedia which contains over 5 million entities. This is not a major problem in these models, as this has been shown (Liu et al., 2019) to remove a small percentage (typically less than 5%) of the data. Previous works (Rijhwani et al., 2019; Zhou et al., 2020) on candidate generation for cross-lingual entity linking have adopted the same practice. Indeed, cross-lingual entity linking is defined as the task of grounding mentions in non-English documents to entries in the English Wikipedia

Table 1: Micro recall@30 of different approaches. The upper and lower block show the performance of the methods in the zero-shot and supervised setting, respectively. For the pairs CO - IT, OLO - FI and TK - AZ, where there is no training data in the target language, the supervised setting amounts to the zero-shot one. The best performance for each target language and setting is in **bold**.

Target Lang.	Pivot Language								
	High Resource			Medium Resource			Low Resource		
	CO - IT	LI - NL	BAR - DE	OLO - FI	MK - BG	JV - ID	TK - AZ	MR - HI	FO - IS
ZERO-SHOT									
WIKIPRIORS	47	51	63	45	31	74	48	30	29
CHARAGRAM	51	51	50	48	47	60	67	48	49
PTI	63	70	76	72	49	79	69	40	47
JOINT LEARNING WITH THE TARGET LANGUAGE									
WIKIPRIORS	47	68	77	45	59	82	48	58	58
CHARAGRAM	51	72	73	48	69	72	67	66	74
PTI	63	84	88	72	71	87	69	70	79

Table 2: Performance breakdown in terms of recall@30 for the case when the models are trained with some supervision from the target language—except for CO - IT, OLO - FI and TK - AZ, which correspond to the zero-shot setting. E, M and H are shortcuts that refer to EASY, MEDIUM and HARD queries, respectively. The best performance for each target language and query type is shown in **bold**.

Target Lang.	Query Type	Pivot Language																				
		High Resource						Medium Resource						Low Resource								
		CO - IT	LI - NL		BAR - DE		OLO - FI	MK - BG		JV - ID		TK - AZ	MR - HI		FO - IS							
		H	E	M	H	E	M	H	H	E	M	H	E	M	H	H	E	M	H			
WIKIPRIORS		47	100	49	57	100	56	74	45	99	53	26	100	79	68	48	99	61	15	100	38	35
CHARAGRAM		51	83	72	61	84	75	62	48	86	74	46	82	70	61	67	82	75	42	90	80	54
PTI		63	98	80	73	99	84	82	72	92	78	44	99	86	78	69	98	82	30	100	93	45

(Tsai and Roth, 2016; Upadhyay et al., 2018a). However, we will show in Section 6.1 that this practice is problematic in the cross-lingual setting, as we would have to ignore a very large percentage of data.

For some methods such as CHARAGRAM is easy to expand the candidate space by considering entities in Wikipedias other than English. While one might evaluate all entities (more than 20 million), we discovered one much less costly alternative. We experimentally observe that expanding the candidate space with the entities observed in the training data of the pivot language significantly improves the entity coverage of the entities observed in the test data of the target language (more in Section 6.1). With the goal of performing a realistic evaluation we simply evaluate all mentions.

5.5. Training Details

The simplicity of our approach PTI is also reflected in its number of hyperparameters. The hyperparameter α , which controls the contribution of the pivot language in the construction of the prior and posterior probabilities, is validated among the values $\{0.1, 0.4, 1\}$. The hyperparameter λ , which controls the importance of the posterior probability, is validated among the values $\{0.2, 0.4, 1\}$. As in CHARAGRAM, the tokenizer `tknr` applied to mentions returns character n -grams with $n \in \{2, 3, 4, 5\}$. The tokenizer counts each symbol as a character regardless of the script of the lan-

guage, although more sophisticated mechanisms are also possible (*e.g.* by converting strings to international phonetic alphabet (IPA) symbols). For CHARAGRAM we replicate the exact same setup as reported in (Zhou et al., 2020). The hyperparameter μ of the supervised setting is validated among the values $\{0.01, 0.1, 0.2\}$. More details are provided in Appendix A.2. The hyperparameters α and μ are required by both PTI and CHARAGRAM, respectively, only in the supervised setting. For all cases, the validation metric is micro recall@30.

6. Results

Table 1 depicts a comparison of the methods in terms of micro recall@30. Similar conclusions are drawn for both learning settings. Overall, our approach PTI clearly outperforms CHARAGRAM in those target languages where the pivot language is high-resourced. For pivot languages that are medium-resourced, while PTI is the best performing technique CHARAGRAM performs very similarly in one target language (MK). This tendency is not as clear when the pivot language is low-resourced, where CHARAGRAM shows a clear superior performance in the target language MR in the zero-shot learning setting. We see that all models largely benefit from the existence of training data in the target language.




Figure 2: Ceiling recall for all query types and target languages. In the experiments reported in Section 5, the candidate space of PTI is as follows: PL and PL+TL in the zero-shot and supervised setting, respectively. Whereas for CHARAGRAM is as follows: EN+PL and EN+PL+TL in the zero-shot and supervised setting, respectively.

By the definition of the query types introduced in Section 5.1, it follows that all queries are of type HARD in the zero-shot setting. However, this is not the case when the models are also trained with data from the target language. We show the performance breakdown in Table 2. The conclusions are clear: PTI always outperforms CHARAGRAM for queries of type EASY and MEDIUM, and it is only beaten for queries of type HARD when the pivot language has low resources, with the exception of AZ. We note that while in this work we have created the evaluation splits so that all query types are equally important, the distribution of the query types depends on the target evaluation data. For instance, splitting Wikipedia articles into training, validation and test in a proportion of 70%, 15% and 15%, respectively, leads to a distribution of query types largely dominated by EASY ones.

The performance of WIKIPRIORS is also dependent on the amount of resources of the pivot language, showing a performance that is comparable or even better than CHARAGRAM when the pivot language is high- and medium-resourced (see Table 1). The more detailed analysis depicted in Table 2 shows that WIKIPRIORS has a very heterogeneous performance with respect to the different query types in the supervised setting. As expected, it obtains a perfect performance for EASY queries, but its performance tends to be poor, with some few exceptions, for MEDIUM and HARD queries.




Figure 3: Recall@30 of PTI (left) and CHARAGRAM (right) for different amount of training data. For simplicity, the candidate space of CHARAGRAM is always formed as if all training data was observed.

6.1. Analysis

Entity Coverage As argued in Section 5.4, it is unrealistic to simply ignore those mentions that are not in the considered KB, specially if this translates into a performance that would not resemble that of the realistic evaluation that we propose. We proceed to analyze the ceiling recall—which depends on the entity coverage—of the methods in the realistic setting, where no mention is excluded.

Figure 2 shows the ceiling recall of PTI and CHARAGRAM for the different target languages and query types. The ceiling recall of CHARAGRAM is lower than 30% in almost all cases when the candidate space is given by the English Wikipedia (EN) considered in previous work (Rijhwani et al., 2019; Zhou et al., 2020). This showcases that previous work were limiting, to a great extent, the mention queries that were

evaluated. This is not only unrealistic, but also a notable limiting factor in performance. The ceiling recall notably increases when the candidate space also includes the entities that are observed in the training data of the pivot language. We hypothesize that this is because the pivot language is not only lexically similar to the target language, but also their respective Wikipedias exhibit similar topical interests. Moreover, the ceiling for queries of type EASY and MEDIUM is also increased when the candidate space is also formed by the entities observed in the training data of the target language. The latter corresponds to the supervised setting; and surprisingly, as opposed to PTI, CHARAGRAM does not reach a ceiling recall of 1 for queries of type EASY and MEDIUM. The reason is that there are a fair amount of entities in the respective Wikipedias that do not have an English name (*e.g.* Q61068957 or Q15842755). This is the strength and weakness of CHARAGRAM: its candidate space can include any entity as long as its English title is available. On the other side, PTI is language-agnostic to the representation of the entity, but its entity coverage is limited to the entities observed in the pivot language (PL), or in both the pivot and target language (PL+TL). It is for this reason that PTI is beaten in queries of type HARD when the pivot language is low-resourced: it fully relies on the entity coverage of the pivot language, which is low, to successfully answer these queries. This is illustrated in Figure 2h and 2i, where CHARAGRAM shows a higher ceiling recall for HARD mentions due to the additional entity coverage provided by the English (EN) Wikipedia. Future work should address this limitation of PTI (*e.g.* by incorporating more pivot languages). However, in all other situations where the ceiling recall is higher or comparable to that of CHARAGRAM, PTI shows the best performance.

Impact of the amount of training data We train both PTI and CHARAGRAM with increasing amount of training data, and computed micro recall@30 for several target languages. This is shown in Figure 3. The performance of PTI increases with the more data up to a point—around 10 million. One reason is that the entity coverage of PTI is expected to increase with the amount of training data. Another reason is that the constructed indexes become more accurate. On the other side, CHARAGRAM barely improves beyond 80,000 data points¹. CHARAGRAM needs relatively little data to learn equivalences between n -grams co-occurring in mentions and English entity names, but gets stuck very quickly and does not leverage all the statistical information contained in the training corpora.

Runtime and Memory The goal of a candidate generator is to reduce the computational cost of a more complex subsequent model. Therefore, besides recall, complexity must also be a factor that drives the design of a candidate generator. We perform a detailed anal-

¹Authors of CHARAGRAM confirmed in an email they used, at most, this amount of data for training

Table 3: Recall@30 in the high-resource setting.

Lang.	IT		NL		DE	
	E	M	E	M	E	M
WIKIPRIORS	95	32	98	30	98	34
CHARAGRAM	50	45	57	47	39	38
PTI	79	49	87	53	77	54

ysis of the complexity in Appendix A.3. The conclusion of such analysis is that in average, PTI requires 20 times less memory than CHARAGRAM, and it is 300 and 7 times faster than CHARAGRAM in terms of training and inference time, respectively.

6.2. Can PTI also become the *de-facto* choice for high-resource languages?

We analyze whether PTI may compete with WIKIPRIORS in the most standard candidate generation setting: candidates have to be generated for mentions of a high-resource language by only exploiting its own training corpus. We build this experiment motivated by the recent findings disclosed in the recent work by *Fu et al.* (Fu et al., 2020), where authors found that even high-resources languages, such as Spanish, are also limited by their mention coverage. This limitation relates to the queries of type MEDIUM. *Fu et al.* partially circumvent the limitation with information coming from search engine query logs.

Following the same protocol as in Section 5.3 we create 1,000 queries of type EASY and MEDIUM for the three high-resource languages used in Section 5 as pivot languages: Italian (IT), Dutch (NL) and German (DE). We do not include queries of type HARD as they cannot be successfully completed. The results, depicted in Table 3, indicate that PTI always obtains the best performance for queries of type MEDIUM, followed by CHARAGRAM. On the other side, for queries of type EASY standard lookup tables (*i.e.* WIKIPRIORS) exhibit the best performance. We conclude that for the standard setting, WIKIPRIORS must be used as the primary technique, falling back to PTI only if WIKIPRIORS fails to retrieve any candidate.

7. Conclusions

We perform an in-depth analysis of the candidate generation problem. We show that the inherited practice of ignoring mentions where the target entity is not in the English Wikipedia is unrealistic in a cross-lingual setting. We alleviate this problem with an efficient solution that consists of using the pivot and target language to expand the candidate space. We also contribute with a categorization of the mention queries according to their difficulty. Finally, we propose a *light-weight* approach, called PTI, that outperforms the current state-of-the-art in almost all target languages and query types.

Acknowledgements

This project was partly funded by the Swiss National Science Foundation (grant 200021_185043), the European Union (TAILOR, grant 952215), and the Microsoft Swiss Joint Research Center. We also gratefully acknowledge generous gifts from Facebook and Google supporting West’s lab.

8. Bibliographical References

- Bao, J., Duan, N., Yan, Z., Zhou, M., and Zhao, T. (2016). Constraint-based question answering with knowledge graph. In *COLING*.
- Fang, M. and Cohn, T. (2017). Model transfer for tagging low-resource languages using a bilingual dictionary. In *ACL*.
- Fu, X., Shi, W., Zhao, Z., Yu, X., and Roth, D. (2020). Design challenges for low-resource cross-lingual entity linking. In *EMNLP*.
- Ganea, O.-E. and Hofmann, T. (2017). Deep joint entity disambiguation with local neural attention. In *EMNLP*.
- Gillick, D., Kulkarni, S., Lansing, L., Presta, A., Baldrige, J., Ie, E., and García-Olano, D. (2019). Learning dense representations for entity retrieval. In *CoNLL*.
- Gollins, T. and Sanderson, M. (2001). Improving cross language retrieval with triangulated translation. In *SIGIR*, pages 90–95.
- Guo, Y., Qin, B., Li, Y., Liu, T., and Li, S. (2013). Improving candidate generation for entity linking. In *International Conference on Application of Natural Language to Information Systems*, pages 225–236. Springer.
- Haghighi, A., Liang, P., Berg-Kirkpatrick, T., and Klein, D. (2008). Learning bilingual lexicons from monolingual corpora. In *ACL*.
- Levin, L. S., Littell, P., Mortensen, D. R., Lin, K., Kairis, K., and Turner, C. (2017). Uriel and lang2vec: Representing languages as typological, geographical, and phylogenetic vectors. In *EACL*.
- Ling, X., Singh, S., and Weld, D. S. (2015). Design challenges for entity linking. *TACL*, 3:315–328.
- Liu, C., Li, F., Sun, X., and Han, H. (2019). Attention-based joint entity linking with entity embedding. *Information*, 10:46.
- Ngo, L. and Wand, M. P. (2004). Smoothing with mixed model software. *Journal of Statistical Software*, 9(1).
- Pan, X., Zhang, B., May, J., Nothman, J., Knight, K., and Ji, H. (2017). Cross-lingual name tagging and linking for 282 languages. In *ACL*.
- Rijhwani, S., Xie, J., Neubig, G., and Carbonell, J. G. (2019). Zero-shot neural transfer for cross-lingual entity linking. In *AAAI*.
- Shen, W., Wang, J., and Han, J. (2014). Entity linking with a knowledge base: Issues, techniques, and solutions. *IEEE Transactions on Knowledge and Data Engineering*, 27(2):443–460.
- Sil, A. and Florian, R. (2016). One for all: Towards language independent named entity linking. In *ACL*.
- Sil, A., Kundu, G., Florian, R., and Hamza, W. (2018). Neural cross-lingual entity linking. In *AAAI*.
- Spitkovsky, V. I. and Chang, A. X. (2012). A cross-lingual dictionary for english wikipedia concepts. In *LREC*.
- Strassel, S. and Tracey, J. (2016). Lorelei language packs: Data, tools, and resources for technology development in low resource languages. In *LREC*.
- Tsai, C.-T. and Roth, D. (2016). Cross-lingual wikification using multilingual embeddings. In *HLT-NAACL*.
- Upadhyay, S., Gupta, N., and Roth, D. (2018a). Joint multilingual supervision for cross-lingual entity linking. In *EMNLP*.
- Upadhyay, S., Kodner, J., and Roth, D. (2018b). Bootstrapping transliteration with constrained discovery for low-resource languages. In *EMNLP*.
- Vulic, I., Ruder, S., and Søgaard, A. (2020). Are all good word vector spaces isomorphic? In *EMNLP*, volume abs/2004.04070.
- Wieting, J., Bansal, M., Gimpel, K., and Livescu, K. (2016). Charagram: Embedding words and sentences via character n-grams. In *EMNLP*.
- Yamada, I., Shindo, H., Takeda, H., and Takefuji, Y. (2017). Learning distributed representations of texts and entities from knowledge base. *Transactions of the Association for Computational Linguistics*, 5:397–411.
- Zhang, X., Zhao, J. J., and LeCun, Y. (2015). Character-level convolutional networks for text classification. In *NIPS*.
- Zhang, H., Song, Y., Song, Y., and Yu, D. (2019). Knowledge-aware pronoun coreference resolution. In *ACL*.
- Zhou, S., Rijhawani, S., Wieting, J., Carbonell, J. G., and Neubig, G. (2020). Improving candidate generation for low-resource cross-lingual entity linking. *TACL*, 8:109–124.
- Zoph, B., Yuret, D., May, J., and Knight, K. (2016). Transfer learning for low-resource neural machine translation. In *EMNLP*.

A. Supplementary Material

A.1. Dataset Information

We preprocess the Wikipedia dumps for each of the target and pivot languages. We apply possible redirects to each Wikipedia entity, and map the (redirected) Wikipedia entity to the corresponding Wikidata entity² (e.g. Q21 is the Wikidata identifier for the Wikipedia entity ENGLAND). We do not use aliases (Zhou et al., 2020) to enrich the data. Nevertheless its inclusion will improve the performance of *any* technique.

Table 4: Pivot languages used in this work categorized according to their amount of resources: high-, medium- and low-resource are in the upper, medium and lower block, respectively.

Lang.	WP Size	#Mention-Entity pairs.
IT (Italian)	1.5M	24M
NL (Dutch)	2M	16M
DE (German)	2.5M	44M
FI (Finnish)	500k	5.7M
BG (Bulgarian)	250k	3.4M
ID (Indonesian)	500k	4.4M
AZ (Azerbaijani)	150k	1.1M
HI (Hindi)	140k	700k
IS (Icelandic)	50k	400k

The transfer learning paradigm of applying a model, learned in a pivot language, to a target language has been explored in different natural language problems. Nevertheless, such works have never reached a consensus on how to select a pivot language that is closely related to a target language. For those target languages that have been used in previous work (e.g. JV or MR) we select the same pivot languages as in these works. For all others we rely on LANG2VEC (Levin et al., 2017) to choose appropriate pivot languages.

A.2. More Details about WIKIPRIORS and CHARAGRAM

WIKIPRIORS is based on the candidate generator used in previous work (Tsai and Roth, 2016; Upadhyay et al., 2018a). It fallbacks to the pivot language when the index of the target language does not retrieve K candidates. On the contrary, for the zero-shot setting it only relies on the index constructed using the pivot language. Furthermore, for both zero-shot and supervised settings, it constructs indexes that associate the constituent words of mention strings with entities using prior probabilities $P(\text{entity}|\text{word}), \forall \text{word} \in \text{mention}$. These indexes are used to retrieve candidates if the indexes on mention strings provide less than K candidates. Previous work (Zhou et al., 2020) has shown that sometimes it is competitive as compared to CHARAGRAM.

We replicate the exact same setup for CHARAGRAM as reported in (Zhou et al., 2020): the tokenizer `tknr`,

²This last step is not really necessary, but we prefer a language-agnostic representation of the entities.

which is applied to mentions and (English) entity names, returns character n -grams with $n \in \{2, 3, 4, 5\}$. The embedding size is 300. We train the model with stochastic gradient descent (SGD) with batch size 64, and a learning rate of 0.1. We stop training if the micro recall@30 on the validation set does not increase for 50 epochs, and the maximum number of training epochs is set to 200.

Table 5: Micro recall@30 after applying thresholding. In parenthesis the percentage of entries that are removed from the indexes built by PTI.

Target - Pivot	Thresholding Value		
	0	1e-2	0.1
CO - IT	63 (0%)	62 (72%)	48 (98%)
LI - NL	84 (0%)	83 (65%)	73 (98%)
BAR - DE	88 (0%)	87 (75%)	70 (99%)
OLO - FI	72 (0%)	71 (68%)	55 (97%)
MK - BG	71 (0%)	69 (66%)	55 (95%)
JV - ID	87 (0%)	87 (64%)	79 (96%)
TK - AZ	69 (0%)	69 (56%)	64 (92%)
MR - HI	70 (0%)	69 (62%)	58 (94%)
FO - IS	79 (0%)	79 (53%)	73 (90%)

A.3. Complexity of PTI and CHARAGRAM

All the experiments were done using code written in Python on an Intel(R) Xeon(R) E5-2680 24-core machine with 2.50GHz CPU, and 256 GB RAM running Linux Ubuntu 16.04. Experiments for CHARAGRAM are run in a TitanX GPU. We do not make use of any parallelism in the experiments.

As previously discussed, the candidate generator must be computational efficient. For this reason we compare the memory and time complexity of PTI to that of CHARAGRAM. We experimentally observe that it is possible to reduce significantly the amount of non-zero prior and posterior probabilities in PTI by thresholding (see Table 5), while having very little impact on the performance. These experiments correspond to the best performing learning setting. We perform further analysis of PTI with a thresholding value of 0.01, and compare it to CHARAGRAM (trained with 80,000 data points) in terms of number of parameters, training and inference time. The inference time corresponds to the time needed to evaluate all mention queries—3,000 queries for almost all target languages. All target languages are evaluated in the best performing learning setting. A comparison is shown in Table 6. The findings are summarized below:

- **Number of parameters.** While the number of parameters remains more or less constant for CHARAGRAM, PTI shows its number of parameters is proportional to the amount of resources of the pivot language.
- **Training time.** The training time of PTI corresponds to the time that takes to construct the prior and posterior probabilities, which is proportional to the

Target - Pivot	Memory		Time			
	#Parameters		Training (minutes)		Inference (seconds)	
	PTI	CHARAGRAM	PTI	CHARAGRAM	PTI	CHARAGRAM
CO - IT	45M	200M	10	180	45	85
LI - NL	38M	220M	7	495	60	120
BAR - DE	80M	225M	20	645	150	150
OLO - FI	18M	180M	3	420	12	70
MK - BG	10M	225M	1	770	30	120
JV - ID	11M	200M	2	480	20	120
TK - AZ	4M	160M	<1	250	3	60
MR - HI	3.5M	150M	<1	680	10	110
FO - IS	2.2M	120M	<1	480	6	115

Table 6: Comparison between PTI (with a thresholding value of 0.01) and CHARAGRAM (trained with 80,000 data points as in (Zhou et al., 2020)) in terms of number of parameters, training and inference time.

amount of resources of the pivot language. The optimization of the learning objective done by CHARAGRAM is significantly much slower.

- **Inference time.** PTI uses sparse matrices to represent $P(\text{entity}|\text{token})$ and $P(\text{token}|\text{entity})$. At test time, scores are simply computed by aggregating rows of such matrices. Moreover, the search performed by PTI to find the K largest values only includes those candidates with non-zero scores. These characteristics make often PTI more efficient than CHARAGRAM during inference.

A.4. Recall@ K for different values of K

Previous works on candidate generation have set a value of 30 for K . However, when the candidate generator is followed by an entity disambiguation technique, there are other values that are typically used by these methods such as (approximately) 10 (Ganea and Hofmann, 2017) and 20 (Upadhyay et al., 2018a). Figure 4 shows the recall of PTI and CHARAGRAM for these other values. Except for CO-IT, OLO-FI and TK-AZ, where zero-shot is the only possible learning setting, all other metrics are evaluated in the supervised setting. The comparative performance between the methods is similar for all values of K .

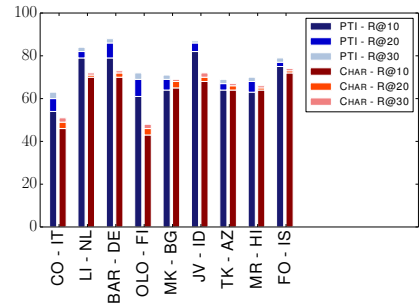


Figure 4: Recall@ K ($R@K$) of PTI and CHARAGRAM (CHAR) for several values of K .

A.5. Other versions of PTI

Wildcard Tokens The tokenizer extends the tokens returned for a given mention with wildcard tokens. Wildcard tokens includes placeholders represented by an asterisk, which can be interpreted as any character. We observed some minor improvements in recall for some target languages, but this notably increased the amount of tokens and, consequently, the memory requirements of the approach.

Probability-level Fusion For the supervised learning setting, we explored another approach to integrate the information from both the target and pivot language. Independent prior and posterior probabilities are computed for both languages, and linearly combined with a weighting factor. Finally, the probabilities are normalized. We did not observe any improvement in performance with this fusion technique.

Probability Smoothing We perform one additional step in the previous version. We use additive smoothing (Ngo and Wand, 2004) (also known as Laplace smoothing) in the computation of the prior and posterior probabilities for the pivot language. These smoothed probabilities account for possible discrepancies between the target and pivot language. We sometimes observed a minor improvement for some query types, but this comes at the cost of validating one more hyperparameter (the smoothing factor).