

Constructing large scale biomedical knowledge bases from scratch with rapid annotation of interpretable patterns

Julien Fauqueur*

BenevolentAI
4-8 Maple St, London
W1T 5HD

julien@benevolent.ai

Ashok Thillaisundaram*

BenevolentAI
4-8 Maple St, London
W1T 5HD

ashok@benevolent.ai

Theodosia Togia*

BenevolentAI
4-8 Maple St, London
W1T 5HD

sia@benevolent.ai

Abstract

Knowledge base construction is crucial for summarising, understanding and inferring relationships between biomedical entities. However, for many practical applications such as drug discovery, the scarcity of relevant facts (e.g. *gene X is therapeutic target for disease Y*) severely limits a domain expert's ability to create a usable knowledge base, either directly or by training a relation extraction model. In this paper, we present a simple and effective method of extracting new facts with a pre-specified binary relationship type from the biomedical literature, without requiring any training data or hand-crafted rules. Our system discovers, ranks and presents the most salient patterns to domain experts in an interpretable form. By marking patterns as compatible with the desired relationship type, experts indirectly batch-annotate candidate pairs whose relationship is expressed with such patterns in the literature. Even with a complete absence of seed data, experts are able to discover thousands of high-quality pairs with the desired relationship within minutes. When a small number of relevant pairs do exist - even when their relationship is more general (e.g. *gene X is biologically associated with disease Y*) than the relationship of interest - our system leverages them in order to i) learn a better ranking of the patterns to be annotated or ii) generate weakly labelled pairs in a fully automated manner. We evaluate our method both intrinsically and via a downstream knowledge

base completion task, and show that it is an effective way of constructing knowledge bases when few or no relevant facts are already available.

1 Introduction

In many important biomedical applications, experts seek to extract facts that are often complex and tied to particular tasks, hence data that are truly fit for purpose are scarce or simply non-existent. Even when only binary relations are sought, useful facts tend to be more specific (e.g. *mutation of gene X has a causal effect on disease Y in an animal model*) than associations typically found in widely available knowledge bases. Extracting facts with a pre-specified relationship type from the literature in the absence of training data often relies on handcrafted rules, which are laborious, ad-hoc and hardly reusable for other types of relations. Recent attempts to create relational data from scratch by denoising the output of multiple hand-written rules (Ratner et al., 2016) or by augmenting existing data through the induction of new black-box heuristics (Varma and Ré, 2018) are still dependent on ad-hoc human effort or pre-existing data. Our approach involves discovering and recommending, rather than prescribing, rules. Importantly, our rules are presented as text-like patterns whose meaning is transparent to human annotators, enabling integration of an automatic data generation (or augmentation) system with a domain expert feedback loop.

In this work, we make the following contributions:

- We propose a number of methods for extracting patterns from a sentence in which two eligible entities co-occur; different types of patterns have different trade-offs between expressive power and coverage.

*Equal contribution. Listing order is alphabetical. Theodosia proposed and co-ordinated the research project, built the early prototypes and contributed the different methods for extracting and lexicalising patterns. Ashok provided conceptual work on the metrics for ranking simplifications and for the intrinsic evaluation, developed the simplification extraction module, ran the experiments for the automated workflow (with all the parameter variations) and performed all the extrinsic evaluations. Julien was mainly responsible for the system architecture and workflow, the intrinsic evaluation (including interacting with the experts), handling negation and speculation and the clustering algorithm.

- We propose a simple method for presenting patterns in a readable way, enabling faster, more reliable human annotation
- For cases where a small number of seed pairs are already available, we propose a method which utilises these seed pairs to rank newly discovered patterns in terms of their compatibility with the existing data. The resulting patterns can be used with or without a human in the loop.

The rest of the paper is organised as follows. Section 2 describes some related work. Section 3 explains the relationship between patterns and labelling rules and presents some pattern types along with techniques for rendering them interpretable. Section 4 provides a high-level overview of the system and covers details of our different workflows (with and without seed data; with and without human feedback). Section 5 explains how we measure the system’s performance both intrinsically and via a downstream knowledge base completion task. In section 6, we report the details of our main experiments while in sections 7 and 8 we present some analysis along with further experiments. The paper ends with conclusions and proposals for further work in section 9.

2 Related work

The idea of extracting entity pairs by discovering textual patterns dates back to early work on bootstrapping for relation extraction with the DIPRE system (Brin, 1999). This system was designed to find co-occurrences of seed entity pairs of a known relationship type inside unlabelled text, then extract simple patterns (exact string matches) from these occurrences and use them to discover new entity pairs. Agichtein et al. (2000) introduced a pattern evaluation methodology based on the precision of a pattern on the set of entity pairs which had already been discovered; they also used the dot product between word vectors instead of an exact string match to allow for slight variations in text. Later work (Greenwood and Stevenson, 2006; Xu et al., 2007; Alfonseca et al., 2012) has proposed more sophisticated pattern extraction methods (based on dependency graphs or kernel methods on word vectors) and different pattern evaluation frameworks (document relevance scores).

Two recent weak supervision techniques, Data Programming (Ratner et al., 2016) and the method

underlying the Snuba system (Varma and Ré, 2018) have attempted to combine the results of handcrafted rules and weak base classifiers respectively. Data Programming involves modelling the accuracy of ideally uncorrelated rules devised by domain experts, then combining their output into weak labels. Although this approach does not require any seed data, it does rely on handwritten rules, which are both time consuming and ad-hoc due to the lack of a data-driven mechanism for exploring the space of possible rules. Snuba learns black-box heuristics (parameters for different classifiers) given seed pairs of the desired relationship. This method avoids the need for manually composing rules, however, the rules it learns are not interpretable, which makes the pipeline harder to combine with an active learning step. Second, the system requires gold standard pairs. In contrast, while our system can leverage gold standard annotations, if available, in order to reduce the space of discovered rules, as well as tune the ranking of newly discovered patterns, it is entirely capable of starting without any gold data if ranking is heuristics-based (e.g. prioritisation by frequency) and a human assesses the quality of the highest coverage rules suggested. Our method does not preclude use within a data programming setup as a way of discovering labelling functions or within a system like Snuba, as a way of generating seed pairs. Another body of work, distant supervision (Verga et al., 2018; Lin et al., 2016) has been a recent popular way to extract relationships from weak labels, but does not give the user any control on the model performance.

A well known body of work, OpenIE (Banko et al., 2007; Fader et al., 2011; Mausam et al., 2012; Angeli et al., 2015) aims to extract patterns between entity mentions in sentences, thereby discovering new surface forms which can be clustered (Mohamed et al., 2011; Nakashole et al., 2012) in order to reveal new meaningful relationship types. In the biomedical domain, Percha and Altman (2018) attempt something similar by extracting and clustering dependency patterns between pairs of biomedical entities (e.g. chemical-gene, chemical-disease, gene-disease). Our work differs from these approaches in that we extract pairs for a pre-specified relationship type (either from scratch or by augmenting existing data written with specific guidelines), which is not guaranteed to correspond to a cluster of discovered sur-

face forms.

3 Extracting interpretable patterns

In a rule-based system, a rule, whether handwritten or discovered, can be described as a hypothetical proposition “*if P then Q*”, where *P* (the antecedent) is a set of conditions that may be true or false of the system’s input and *Q* (the consequent) is the system’s output. For instance, a standard rule-based relation extraction system can **i**) take as input a pair of entities (e.g. `TNF-GeneID:7124` and `Melanoma-MESH:D008545`) that are mentioned in the same piece of text, **ii**) test whether certain conditions are met (e.g. presence of lexical or syntactic features) and **iii**) output a label (e.g. 1: Therapeutic target, 0: Not therapeutic target.)

In this work, patterns are seen as the antecedents of rules that determine which label (consequent) should be assigned to some input (e.g. candidate pair + text that mentions it.) We aim to extract patterns that are expressive enough to allow a system or a domain expert to discriminate between the different labels available for an input but also generic enough to apply to a wide range of inputs. In this work, we have made the following simplifying assumptions:

1. Relationships are binary (i.e. hold between exactly two entities).
2. A pair of entities are candidates for relation extraction if they are mentioned simultaneously in the same sentence.
3. There is a one-to-many relationship between patterns and inputs. An input (i.e. sentence + entity pair) is described by a single pattern (although this pattern can be a boolean combination of other patterns) but one pattern can correspond to multiple inputs.
4. We can select patterns which are expressive enough to represent the relationship, so it is possible to classify the input from which a given pattern has been extracted by examining the pattern alone. However, the omitted part of the sentence may contain contextual information which specifies the condition when or where the relationship holds. Modeling such contextual information would be useful but is beyond the scope of this work. A consequence of this assumption is that it is possible to batch-annotate a group

of inputs that correspond to the same pattern by annotating the pattern itself.

Pattern interpretability An important consideration in this research is pattern interpretability, which could assist domain experts (who are not NLP experts) in exploring the space of labelling rule antecedents for a given relationship type in a given corpus. Hence, for each pattern, we construct what we call a pattern *lexicalisation*, that is converting a pattern to a readable text-like sequence.

Pattern types Simple patterns, which can potentially be combined with boolean operators, can be of different types. We illustrate some types of patterns used in our experiments through the following example sentences that include mentions of a gene-disease pair:

- (1) “We investigate the hypothesis that the knockdown of *BRAF* may affect *melanoma* progression.”
- (2) “The study did not record higher *NF-kb* activity in *cancer* patients.”

Below are some types of patterns, as well as their lexicalisations:

- **KEYWORDS:** words (e.g. ‘inhibiting’) or lemmas (e.g. ‘inhibit’) in the entire sentence or in the text between the entities. This pattern’s lexicalisation is, trivially, the word itself.
- **PATH:** shortest path between the two entity mentions in the dependency graph of the sentence. For instance, in example (1), the path could be `BRAF <-pobj- of <-prep- knockdown <-nsubj- affect -dojb-> progression -compound-> melanoma;` in example (2), the path could be `NF-kb <-compound- activity -prep-> in -pobj-> patients -compound-> cancer).` To lexicalise patterns of this type, we extract the nodes (i.e. words) from the path, arrange them as per their order in the sentence and replace the entity mentions by a symbol denoting simply their entity types. For instance, the first pattern becomes “knockdown of GENE affect DISEASE progression”. This pattern is used extensively in our experiments because it strikes a good balance between expressive power and coverage.

We call its lexicalisation a **simplification** because it is a text-like piece that simplifies a sentence by discarding all but the most essential information.

- **PATH_ROOT**: the root (word with no incoming edges) of the shortest path between the two entities (e.g. ‘affect’ and ‘activity’ in examples (1) and (2) respectively). The lexicalisation could be trivial (i.e. the root itself) or, alternatively, if this pattern is used in an AND boolean combination with the PATH pattern, the root can simply be highlighted (e.g. "knockdown of GENE **affect** DISEASE progression")
- **SENTENCE_ROOT**: the root of the dependency graph of the entire sentence (e.g. ‘investigate’ and ‘record’ in the examples above), which is often not the same as root of the path connecting the two entities. It can be lexicalised similarly to the pattern above.
- **PATH_BETWEEN_ROOTS**: the path between the root of the entire sentence and the root of the path between the two entities (e.g. `investigate -dobj-> hypothesis -acl-> affect and record -dobj-> activity` for examples (1) and (2) respectively). The pattern can be lexicalised as what we have called “simplification” (e.g. `investigate hypothesis affect, or, if AND-ed with the PATH pattern, all the words from both patterns can be merged and arranged as per their original order in the sentence, potentially with some highlighting to differentiate the two simpler patterns (e.g. "investigate hypothesis knockdown of GENE affect DISEASE progression")`)
- **SENTENCE_ROOT_DESCENDANTS**: the direct descendants of the SENTENCE_ROOT, for instance, ‘did’, ‘not’ and ‘activity’ in the example (1), because of the edges `did <-aux- record, not <-neg- record and record -dobj-> activity`. To lexicalise this pattern, we can extract the words and merge them with words of other patterns. Alternatively, we can devise some simpler sub-patterns, for instance, descendants with `aux`, that is auxiliary, edges, such as ‘may’, or descendants with `neg` edges such as ‘not’ and place them outside any simplification: "investigate hypothesis knockdown

of GENE affect DISEASE progression + hedging:[may]"

- **PATH_ROOT_DESCENDANTS**: the direct descendants of the root of the path between the entities (e.g. ‘may’ and ‘progression’ in example (1) because `may <-aux- affect and affect -dobj-> progression; ‘higher’ and ‘in’ in the example (2) because higher <-amod- activity and activity -prep-> in). Its lexicalisation can be the same as that of the previous pattern type.`

Other examples of patterns could be regular expressions or rules informed by an external biomedical ontology (e.g. `GENE is a Rhodopsin-like receptor`) or with lexical information from databases like WordNet (Miller, 1995) (e.g. for increasing pattern coverage leveraging synonyms or hypernyms of words in a pattern.)

It should be obvious that the more expressive a pattern becomes (for instance by AND-ing multiple other patterns), the less capable it is of subsuming many sentences. It is important to discover patterns with this trade-off in mind.

4 System overview

In this section, we will describe each step of our system, outlined in Figure 1.

4.1 Data preparation

Extracting named entities and patterns The first step is performing named entity recognition (NER) on the sentences in the corpus to enable us to identify all the sentences which contain entity pairs of interest. Our experiments are focused on gene-therapeutic target pairs, however, the system is designed to be agnostic to different types of entities and relationships between them. We then extract the desired patterns from each of these sentences, as described in section 3. For simplicity, we limited our experiments to sentences that contain exactly one gene-disease pair, however, extending the system to handle multiple pairs is straightforward.

We index each sentence in a database along with the lexicalisation for its pattern (e.g. the ‘simplification’ for PATH or PATH_BETWEEN_ROOTS patterns) and the entity pair found. This allows us to easily query this database i) for all entity pairs that correspond to a pattern (which is now lexicalised

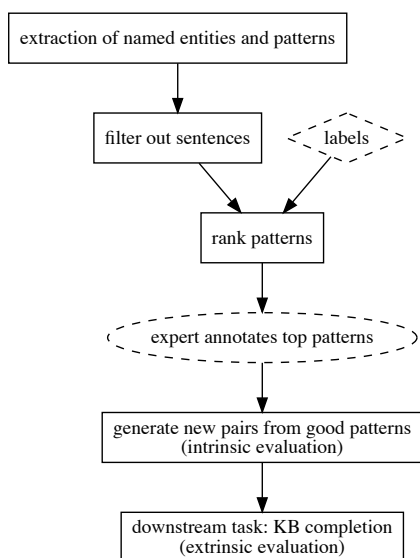


Figure 1: System overview. Diamond box is present only in workflows with seed labels available (i.e. “no expert but labels” and “expert with labels”), elliptical box is only present in workflows involving an expert (i.e. “expert - no labels” and “expert with labels”) and rectangular boxes are always present.

and stored as a string) or ii) for all patterns that correspond to an entity pair.

Filtering out sentences with negation and hedging. Since we are interested in inputs which unambiguously encode affirmations of facts about entities, we filter out any sentences which contain negation, speculation, or other forms of hedging. We adopt a conservative approach by excluding sentences which match specific instantiations of these pattern types: **i)** KEYWORDS (e.g. presence of terms such as “no”, “didn’t”, “doubt”, “speculate” etc. in the sentence); our list is modified from NegEx (Chapman et al., 2013), **ii)** SENTENCE_ROOT AND SENTENCE_ROOT_DESCENDANTS (e.g. “study we investigated”, which makes no statement of results), **iii)** PATH_ROOT AND PATH_ROOT_DESCENDANTS (e.g. “was used”, “was performed”), **iv)** path_between_roots (e.g. “found associated”) This filtering is applied at all stages in our system where sentences are used.

4.2 Ranking patterns

Below we describe methods for ranking and selecting top patterns in the presence or absence of domain expertise or labelled training data.

Baseline workflow: “no expert - no labels” In this workflow, we simply extract new pairs using

simplifications (from the PATH pattern type, but other types are also described in our experiments) that have a high enough (≥ 5) count of entity pairs.

Manual curation in the absence of any labelled training data: “expert - no labels” In this workflow, we have a domain expert (a biologist) available for manual curation but there is no labelled training data. It is not possible for a domain expert to annotate all simplifications; this would be too time-consuming. In such cases, active learning can be helpful in deciding which simplifications should be shown to the domain experts for manual curation to best improve the output of our system. The approach that we adopt here is simple but the system could be extended with more sophisticated active learning strategies. We rank the simplifications by their count of entity pairs; by this we mean the number of unique pairs contained in the sentences in our corpus which correspond to a given simplification (similar to section 4.2). We then show the top ranked simplifications (i.e. those with the greatest pair count) to our domain expert with a fixed number of random example sentences who then decides if a given simplification is an appropriate heuristic to extract new entity pairs from the corpus, by selecting one of three options “Yes”, “No”, “Maybe”.

Automated workflow: “no expert but labels”

For this workflow, a set of gold standard pairs exists as training data but we have no domain experts available. The sentence simplifications can be ranked using various metrics calculated against the gold standard training data. Each simplification is considered as a classifier: A given pair is ‘classified’ by the simplification as positive if the pair can be discovered using the simplification’s underlying rule in the corpus. Otherwise, it is classified as negative. The metrics we use to rank the simplifications are precision and recall. The gold standard pairs will form the positive pairs in our training data. To obtain negative pairs, we operate under the closed world assumption: any entity pair found in our corpus of sentences not present in our gold standard set is taken to be negative. This results in an imbalance in the sizes of positive and negative training data which skews the value of precision. To address this, we use a precision metric where the number of true positives and false positives are normalised

by the total number of positive and negative pairs respectively in our training data. For each simplification S we define true positives (TP_S) and false positives (FP_S) as the sets of correctly and incorrectly positive-labelled entity pairs respectively. Our variant of precision for a simplification S is then, $precision_S = \frac{|TP_S|/N_P}{|TP_S|/N_P + |FP_S|/N_N}$ where N_P and N_N are the number of positive and negative pairs respectively in the training data. With this metric, if a simplification classifies 10% of the positive pairs as true positives and 10% of the negative pairs as false positives then $precision_S = \frac{0.1}{0.1+0.1} = 0.5$. The metric utilises the percentage of each class instead of the absolute number of pairs, as would be the case for the standard precision metric. The definition of recall for a given simplification S is with respect to just the positive training data and is thus unaffected by an imbalance in the sizes.

Manual curation with labelled data: “expert with labels” For this workflow, both domain experts and labelled training data are available to us. We improve on our methodology in the “expert - no labels” workflow by making use of the metrics discussed in the “no expert but labels” workflow which are calculated using the labelled training data. As we want to maximise the number and precision of new pairs extracted, we keep only simplifications with recall and precision above certain respective thresholds and present them to domain experts ranked by pair count to ensure they see the most impactful simplifications first.

4.3 Generating new pairs

All previous stages aim at generating a list of good simplifications. We now have a collection of rules which can be used to extract new entity pairs from the corpus. Any simplification selected as useful implies that all entity pairs recovered from the corpus using this rule can be added as positive examples to the dataset. With the selected simplifications, we can batch-annotate thousands of sentences, and hence pairs, with minimal effort. We simply query our database for all new pairs which are found in a sentence expressing any of our selected simplifications.

Clustering simplifications We found that many simplifications can be very similar up to a few characters. We create clusters of quasi-identical simplifications, and use them i) to enforce diver-

sity in the selection of simplifications for the user to annotate, by picking only one simplification per cluster and, ii) to safely extend the selection of positive simplification to other simplification in the cluster. We create clusters of simplifications by detecting connected components in a graph where the nodes are the simplifications and the edges are between simplifications which are at a maximum Levenshtein distance of 2. This allows us to be invariant to plural forms, upper/lower case, to short words like in/of etc. Note that some (not all) of these variations could be captured with a lemmatiser. Example of a distance 2 cluster:

```
{GENE effects on DISEASE,
GENE effect on DISEASE,
GENE effects in DISEASE}
```

With a distance of 2, we typically increase the number of positive simplifications by 50%, which significantly increases the recall on new pairs.

5 Evaluation

We implement two evaluation frameworks. The first is an intrinsic evaluation of the quality of the new extracted pairs. The second is extrinsic; we consider how the inclusion of the new pairs discovered by our system affects the performance of a downstream knowledge base completion task.*

5.1 Intrinsic evaluation

Pair-level Our aim in this subsection is to construct an intrinsic evaluation framework which can directly measure the quality of the discovered pairs. We do this by holding out a fraction of the gold standard positive pairs and the negative pairs (under the closed world assumption) to be used as a test set. The remaining fraction is used as training data. We evaluate our system by measuring its recall, specificity (true negative rate), precision, and F-score against this test set. In more detail, the new pairs discovered by our selected simplifications are taken to be the positive pairs predicted by our system. The overlap between these new pairs and the positive test set are the true positives (TP) while the overlap with the negative test set are the false positives (FP). Recall and specificity take their standard definitions. Again, we consider a precision score which is normalised to correct

*We consider the second type of evaluation extrinsic because knowledge base completion aims to recover latent relationships, whereas knowledge base construction, which the system is built for, is limited to extracting pairs from the literature.

for the imbalance in numbers between positive and negative pairs, $precision = \frac{|TP|/N_P}{|TP|/N_P + |FP|/N_N}$, where N_P and N_N are the numbers of positive and negative pairs respectively present in the test set (as described in section 4.2). We take the F-score to be the harmonic mean of this precision variant and recall.

Simplification-level For the manual workflows, we also consider the expert annotations while assessing the quality of the simplifications. We report MSP , the manual simplification precision, based on N_{Yes} , N_{No} and N_{Maybe} , the number of simplifications that the expert has annotated as “Yes”, “No” and “Maybe”. $MSP = \frac{N_{Yes}}{N_{Yes} + N_{No} + N_{Maybe}}$. We expect MSP to be as high as possible.

Extrinsic evaluation via knowledge base completion The setup for our extrinsic evaluation framework is straightforward and intuitive. The initial gold standard set of positive pairs is split into training and test data. A graph completion model is then trained using the training data and evaluated to determine whether it can predict the existence of the pairs in the test data. To determine whether our knowledge base construction system can add value, we use the new pairs found from our system to augment the training data for the graph completion model, and observe whether this improves its performance against the test set. We use ComplEx (Trouillon et al., 2017), a well-established tensor factorisation model, as our knowledge base completion model. We provide standard information retrieval metrics to quantify the performance of the graph completion model. These are the precision, $P(k)$, and recall, $R(k)$, calculated for the top k predictions along with the mean average precision (mAP). For gene-disease entity pairs, for example, $mAP = \frac{1}{N_d} \sum_d AveP$, where the sum is over the diseases d with N_d being the total number of diseases, and $AveP = \sum_k P(k) (R(k) - R(k - 1))$ with $P(k)$ and $R(k)$ as defined above.

6 Main experiments and results

6.1 Datasets

For all the following experiments, our data was drawn from the following datasets: DisGeNET (Pinerio et al., 2016) and Comparative Toxicogenomics Database (CTD) (Davis et al., 2018). CTD

contains two relation types: ‘marker/mechanism’ and ‘therapeutic’. We use both the entire CTD dataset and the subset of therapeutic gene-disease pairs which we refer to as CTD therapeutic.

The datasets above are first restricted to human genes and then to the gene-disease pairs which appear in our corpus of sentences; this corpus consists of sentences from PubMed articles which have been restricted, for simplicity, to sentences which contain just one gene-disease pair each. With these restrictions in effect, the CTD dataset has 8828 gene-disease pairs, CTD therapeutic has 169 pairs, and Disgenet has 33844 pairs.

6.2 Intrinsic evaluation results

In table 1, we report the pair-level metrics (see section 5.1) for our three proposed workflows and a baseline (see section 4). We also report the expert-based metric MSP (see section 5.1) for the two manual workflows. The CTD therapeutic dataset was the most suitable dataset for this evaluation because **i**) it is very relevant to crucial domains of application such as drug discovery, and **ii**) its small size makes it a good candidate for expansion. In each session, the expert annotated 200 simplifications accompanied by 20 sentences. It took the expert about 3 hours to annotate the first session, which is a rapid way to generate thousands of new pairs from scratch.

We find that our three main proposed workflows (‘expert - with labels’, ‘expert - no labels’, and the fully automated ‘no expert but labels’) all discover a significant number of new gene-disease therapeutic pairs. As confirmed by both pair-level and user-based metrics, incorporating the use of domain expert’s time and the use of labelled data results in higher precision at the expense of recall.

6.3 Extrinsic evaluation results

In table 2, we list the results of the downstream knowledge base completion task for the fully automated workflow and the baseline. We compare the performance of our knowledge base completion model when trained with just the initial seed training data versus the seed training data augmented with the new pairs discovered by our fully automated workflow (and baseline workflow).

The addition of new pairs from the fully automated workflow gives us a higher mean average precision (mAP) than with just the seed dataset. We obtain a higher precision (for the top 100 and top 1000 predictions) while maintaining the same

Selection method	MSP	New pairs	Recall	Specificity	Precision	F-score
expert with labels	0.315	8875	0.286	0.976	0.923	0.436
expert - no labels	0.265	9560	0.250	0.975	0.908	0.392
no expert but labels	-	30006	0.679	0.920	0.894	0.772
no expert - no label (baseline)	-	59913	0.774	0.842	0.830	0.801

Table 1: Intrinsic evaluation results for the main experiments on the CTD therapeutic dataset. This was carried out with a train/valid/test split of 0.4/0.1/0.5, and precision threshold of 0.6 for the ‘expert with labels’ and ‘no expert but labels’ workflows. MSP is our “manual simplification precision” metric. The precision and F-scores reported here are normalised as described in the section 4.2.

level of recall. For the baseline workflow, mAP is higher but with lower precision (for the top 100 and 1000 predictions respectively).

7 Top simplifications

In table 3, we show the simplifications with the highest count of Disease-Gene pairs in our whole corpus (after the sentence filtering), which have been annotated by the expert as “Yes” or “No”, for the CTD therapeutic dataset. While “Yes” and “No” patterns look similar, we can clearly see differences in language. The “No” annotations look unspecific while the “Yes” ones express the target has a therapeutic effect on the disease.

8 Further experiments

We performed several other experiments using our fully automated workflow to evaluate the quality of the new pairs discovered as we varied our experiment parameters.

We consider three dimensions of variation: varying the precision threshold for selecting simplifications, varying the size of the seed training set, and varying the expressiveness of the simplification (for example, by including the SENTENCE_ROOT or restricting to simplifications with at least a specified number of words).

The intrinsic evaluation results for these experiments are listed in tables 4, 5, and 6. In all cases, as we make our system more selective either by raising the precision threshold, by starting with fewer seeds pairs, or by restricting to more informative simplifications, we unsurprisingly obtain higher precision at the expense of lower recall.

The extrinsic evaluation framework is less sensitive to these changes but improvements were observed (without any noticeable trend) for all these parameter changes.

9 Conclusions and further work

We have presented a simple and effective method for knowledge base construction when the desired relational data are scarce or absent. We have demonstrated its effectiveness via **i)** classification metrics on a held-out test set, **ii)** human evaluation and **iii)** performance on a downstream knowledge base completion task. We further show that in the presence of a small set of data, it is possible to control the quality of the pairs discovered, by introducing stricter precision thresholds when ranking patterns. Our method could in principle be extended in order to: **1)** handle higher-order (e.g. ternary) relations between tuples, as opposed to pairs (for instance using dependency subgraphs that connect more than two entities cooccurring in a sentence), **2)** discover explicit negative examples of a binary relation instead of simply positive examples, **3)** train sentence-level relation extraction systems, **4)** collect and utilise continuous, rather than discrete annotations for each pattern (e.g. annotators could indicate the percentage of correct example sentences that correspond to a pattern displayed) as part of a more sophisticated active learning strategy, **5)** extract patterns from a semantic representation (Banarescu et al., 2013) and, finally, **6)** map patterns to a vector space using a distributional representation (e.g. defined by their neighbouring words in sentences) and cluster them for an optimal balance between expressive power and coverage.

Acknowledgments

We are very grateful to Nathan Patel for his engineering support and to Alex de Giorgio for his thorough feedback and domain expertise. We would also like to thank many of our colleagues working on drug discovery and link prediction for insightful conversations, as well as Felix Kruger for proofreading the final version of this paper.

Selection method	New pairs	MAP	Precision	Recall
seed dataset only	-	0.0414	0.0179 / 0.0179	1.0 / 1.0
no expert but labels	30006	0.0545	0.0192 / 0.0192	1.0 / 1.0
no expert - no label (baseline)	59913	0.1885	0.01 / 0.0015	0.6019 / 0.9208

Table 2: Extrinsic evaluation results for the CTD therapeutic dataset. The experiment parameters are the same as those given in table 1. Precision figures are given as ‘top 100 / top 1000’ and similarly for recall.

Pairs	“Yes” simplif.	Pairs	“No” simplif.
3345	role of GENE in DISEASE	6629	GENE DISEASE
839	GENE plays in DISEASE	4110	DISEASE GENE
648	GENE involved in DISEASE	3350	GENE and DISEASE
321	GENE target in DISEASE	2370	GENE in DISEASE
318	GENE target for DISEASE	2333	DISEASE and GENE
289	GENE mice develop DISEASE	1228	GENE DISEASE cells
279	DISEASE caused by mutations in GENE	904	DISEASE of GENE
276	GENE gene for DISEASE	879	DISEASE in GENE
273	role of GENE in development of DISEASE	638	DISEASE in GENE mice
		572	role for GENE in DISEASE
237	GENE promotes DISEASE	528	GENE in DISEASE patients

Table 3: Top 10 simplifications for CTD Therapeutic annotated “Yes” (left) and “No” (right) by the expert.

Dataset	Thres.	New pairs	R	S	P	F
CTD	0.8	29592	0.297	0.918	0.783	0.430
CTD	0.4	50329	0.379	0.863	0.735	0.500
DG	0.8	17441	0.180	0.947	0.773	0.292
DG	0.4	45446	0.314	0.867	0.703	0.434

Table 4: Intrinsic evaluation results (Recall, Specificity, Precision and F-score) on CTD and DisGeNET (DG) as we vary the precision threshold for the ‘no expert but labels’ workflow. Experiments are done with a train/valid/test split of 0.8/0.1/0.1 and we restrict to simplifications with at least 5 words to ensure that they are reasonably expressive.

References

- Eugene Agichtein and Luis Gravano. 2000. [Snowball: Extracting relations from large plain-text collections](#). In *Proceedings of the Fifth ACM Conference on Digital Libraries*, DL ’00, pages 85–94, New York, NY, USA. ACM.
- Enrique Alfonseca, Katja Filippova, Jean-Yves Delort, and Guillermo Garrido. 2012. [Pattern learning for relation extraction with a hierarchical topic model](#). In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Short Papers - Volume 2*, ACL ’12, pages 54–59, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Gabor Angeli, Melvin Jose Johnson Premkumar, and Christopher D. Manning. 2015. [Leveraging linguistic structure for open domain information extraction](#). In *ACL (1)*, pages 344–354. The Association for Computer Linguistics.
- Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. [Abstract meaning representation](#)

Train/val/test	New pairs	R	S	P	F
0.8/0.1/0.1	29592	0.297	0.918	0.783	0.430
0.5/0.1/0.4	25539	0.274	0.930	0.798	0.408
0.2/0.1/0.7	18268	0.225	0.950	0.818	0.352

Table 5: Intrinsic evaluation results (Recall, Specificity, Precision and F-score) for CTD as we vary size of the seed training data for the ‘no expert but labels’ workflow. Experiments are done with a precision threshold of 0.8 and we restrict to simplifications with at least 5 words to ensure that they are reasonably specific.

Min length	SENTENCE ROOT	New pairs	R	S	P	F
5	No	29592	0.297	0.918	0.783	0.430
5	Yes	13284	0.174	0.963	0.824	0.288
7	No	5313	0.066	0.986	0.824	0.122
7	Yes	1814	0.028	0.995	0.856	0.055

Table 6: Intrinsic evaluation results (Recall, Specificity, Precision and F-score) for CTD as we vary simplification expressive power (minimum length) for the ‘no expert but labels’ workflow. Experiments are done with a train/valid/test split of 0.8/0.1/0.1 and a precision threshold of 0.8.

for [sembanking](#). In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 178–186, Sofia, Bulgaria. Association for Computational Linguistics.

- Michele Banko, Michael J. Cafarella, Stephen Soderland, Matt Broadhead, and Oren Etzioni. 2007. [Open information extraction from the web](#). In *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, IJCAI’07, pages 2670–2676, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Sergey Brin. 1999. [Extracting patterns and relations from the world wide web](#). In *Selected Papers from the International Workshop on The World Wide Web and Databases*, WebDB ’98, pages 172–183, London, UK, UK. Springer-Verlag.
- WW Chapman, D Hillert, S Velupillai, M Kvist, M Skeppstedt, BE Chapman, M Conway, M Tharp, DL Mowery, and L. Deleger. 2013. Extending the negex lexicon for multiple languages. In *Proceedings of the 14th World Congress on Medical & Health Informatics (MEDINFO)*.
- AP Davis, CJ Grondin, RJ Johnson, D Sciaky, R McMorran, J Wieggers, TC Wieggers, and CJ Mattingly.

2018. The comparative toxicogenomics database: update 2019. *Nucleic Acids Res.*
- Anthony Fader, Stephen Soderland, and Oren Etzioni. 2011. [Identifying relations for open information extraction](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing, EMNLP '11*, pages 1535–1545, Stroudsburg, PA, USA. Association for Computational Linguistics.
- MA Greenwood and Mark Stevenson. 2006. [Improving semi-supervised acquisition of relation extraction patterns](#). In *Proceedings of the Workshop on Information Extraction Beyond The Document, IEBeyondDoc '06*, pages 29–35, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Yankai Lin, Zhiyuan Liu, Shiqi Shen, Huanbo Luan, and Maosong Sun. 2016. [Neural relation extraction with selective attention over instances](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2124–2133. Association for Computational Linguistics.
- Mausam, Michael Schmitz, Robert Bart, Stephen Soderland, and Oren Etzioni. 2012. [Open language learning for information extraction](#). In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning, EMNLP-CoNLL '12*, pages 523–534, Stroudsburg, PA, USA. Association for Computational Linguistics.
- George Miller. 1995. Wordnet A lexical database for English. *Communications of ACM*, 38(11):39–41.
- Thahir P Mohamed, Estevam R Hruschka Jr, and Tom M Mitchell. 2011. Discovering relations between noun categories. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 1447–1455. Association for Computational Linguistics.
- Ndapandula Nakashole, Gerhard Weikum, and Fabian M. Suchanek. 2012. [Discovering and exploring relations on the web](#). *PVLDB*, 5(12):1982–1985.
- Bethany Percha and Russ B. Altman. 2018. [A global network of biomedical relationships derived from text](#). *Bioinformatics*, 34(15):2614–2624.
- Janet Pinero, Alex Bravo, Nuria Queralt-Rosinach, Alba Gutierrez-Sacristan, Jordi Deu-Pons, Emilio Centeno, Javier Garcia-Garcia, Ferran Sanz, and Laura I. Furlong. 2016. [Disgenet: a comprehensive platform integrating information on human disease-associated genes and variants](#). *Nucl. Acids Res.*
- Alexander Ratner, Christopher De Sa, Sen Wu, Daniel Selsam, and Christopher Ré. 2016. [Data Programming: Creating Large Training Sets, Quickly](#). *arXiv e-prints*, page arXiv:1605.07723.
- Theo Trouillon, Christopher R Dance, Johannes Welbl, Sebastian Riedel, Eric Gaussier, and Guillaume Bouchard. 2017. [Knowledge graph completion via complex tensor factorization](#). *JMLR*.
- Paroma Varma and Christopher Ré. 2018. [Snuba: Automating weak supervision to label training data](#). *Proc. VLDB Endow.*, 12(3):223–236.
- P Verga, E Strubell, and A McCallum. 2018. Simultaneously self-attending to all mentions for full-abstract biological relation extraction. In *North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Feiyu Xu, Hans Uszkoreit, and Hong Li. 2007. A seed-driven bottom-up machine learning framework for extracting relations of various complexity. In *ACL*, pages 584–591.