Extracting Explicit and Implicit Causal Relations from Sparse, Domain-Specific Texts

Ashwin Ittoo and Gosse Bouma

University of Groningen, 9747 AE
Groningen, The Netherlands
{r.a.ittoo,g.bouma}@rug.nl

Abstract. Various supervised algorithms for mining causal relations from large corpora exist. These algorithms have focused on relations explicitly expressed with causal verbs, e.g. "to cause". However, the challenges of extracting causal relations from domain-specific texts have been overlooked. Domain-specific texts are rife with causal relations that are implicitly expressed using verbal and non-verbal patterns, e.g. "reduce", "drop in", "due to". Also, readily-available resources to support supervised algorithms are inexistent in most domains. To address these challenges, we present a novel approach for causal relation extraction. Our approach is minimally-supervised, alleviating the need for annotated data. Also, it identifies both explicit and implicit causal relations. Evaluation results revealed that our technique achieves state-of-the-art performance in extracting causal relations from domain-specific, sparse texts. The results also indicate that many of the domain-specific relations were unclassifiable in existing taxonomies of causality.

Keywords: Relation exaction, Causal relations, Information extraction.

1 Introduction

Causal relations, between causes and effects, are a complex phenomenon, pervading all aspects of life. Causal relations are fundamental in many disciplines, including philosophy, psychology and linguistics. In Natural Language Processing (NLP), algorithms have been developed for discovering causal relations from large general-purpose [2,4,7,8,14] and bio-medical corpora [9]. These algorithms rely extensively on hand-coded knowledge (e.g. annotated corpora), and only extract explicit causal relations. Explicit relations are realized by explicit causal patterns, predominantly assumed to be causal verbs [4,7,14]. Causal verbs (e.g. *"induce"*) are synonymous with the verb "*to cause*". They establish a causal link between a distinct causal-agent (e.g. "rain"), and a distinct effect (e.g. "floods"), as in *"rain causes floods"*.

The discovery of causal relations from texts in other domains (e.g. business/corporate) has been largely overlooked despite numerous application opportunities. In Product Development/Customer Service, for instance, causal relations encode valuable operational knowledge that can be exploited for improving product quality. For example, in *"broken cable resulted in voltage loss"*, the causal relation between

R. Muñoz et al. (Eds.): NLDB 2011, LNCS 6716, pp. 52–63, 2011.

the cause "broken cable" and the effect "voltage loss", established by the pattern "*resulted in*", helps engineers during product diagnosis. Similarly, in *"new analog processor causes system shutdown"*, the causal relation between "new analog processor" and "system shutdown", realized by the pattern *"causes"*, provides business organizations with insights on customer dissatisfaction.

However, extracting causal relations from domain-specific texts poses numerous challenges to extant algorithms. A major difficulty in many domains is the absence of knowledge resources (e.g. annotated data), upon which traditional algorithms rely. In addition, current techniques are unable to detect implicit causal relations, which are rife in the English language. Implicit relations are realized by implicit causal patterns. These patterns do not have any (explicit) causal connotation. But they subtly bias the reader into associating certain events in the texts with causal-agents or effects [10]. Thus, implicit patterns have a causal valence, even though they are not synonymous with "*to cause*". We consider 3 main types of implicit causal relations. Relations of the first type, *T1*, are realized by resultative and instrumentative verbal patterns. These verbs, for e.g. "*increase*", "*reduce*", "*kill*", inherently specify (part of) the resulting situation, as in "*the temperature increased*". The second type of implicit causal relations, *T2*, involves patterns that make the causal-agents inseparable from the resulting situations [10]. Such patterns include *"mar (by)"*, "*plague(by)*". For example, in *"white spots mar the x-ray image"*, the causal-agent "white spots" is an integral component of the result "marred x-ray image". The last type of implicit causal relations, *T3*, involves non-verbal patterns, for e.g. the preposition *"due to"*, as in *"replaced camera due to horizontal calibration problem"*. Besides the difficulties posed by implicit patterns, existing algorithms are also unable to disambiguate ambiguous causal relations that involve polysemous patterns (e.g. "*result in*", "*lead to*"). These patterns express causality only in restricted contexts. For e.g., the pattern "*lead to*" establishes a causal relation in *"smoking leads to cancer"*, but not in *"path leads to garden"*.

To address these challenges, we develop and present a framework for automatically extracting high quality causal relations from domain-specific, sparse corpora. We implemented our methodology in a prototype as part of the DataFusion initiative[1], which aims at enhancing product quality and customer satisfaction using information extracted from corporate texts. The crux of our approach lies in acquiring a set of explicit and implicit causal patterns from Wikipedia, which we exploit as a knowledge-base. We then use these patterns to extract causal relations from domain-specific documents. Our strategy of applying the knowledge acquired from Wikipedia to specialized documents is based on domain-adaptation [3]. It circumvents the data sparsity issues posed by the domain-specific, corporate documents.

Our contributions are as follows. We present a minimally-supervised algorithm that extracts causal relations without relying on hand-coded knowledge. Also, our algorithm accurately disambiguates polysemous causal patterns, and discovers both explicit and implicit causal relations. In addition, we represent the extracted causal patterns as sophisticated syntactic structures, which overcome the shortcomings of traditional pattern representations based on surface-strings.

[1] DataFusion is a collaboration between academia and industry, sponsored by the Dutch Ministry of Economic Affairs.

Our experiments revealed that our approach achieved state-of-the-art performance in extracting causal relations from real-life, domain-specific and sparse documents. These results also demonstrate that Wikipedia can be effectively leveraged upon as a knowledge-base for extracting domain-specific causal relations. Furthermore, we found out that some of the identified domain-specific relations were not conclusively classifiable in Barriere's taxonomy of causality [1].

2 Related Work

Causality is a complex, non-primitive relation, which can be refined into more specialized sub-relations. A recent taxonomy of causal relations is that of Barriere [1], which distinguishes between *existential and influential causality*. *Existential causality* pertains to the creation, destruction, prevention and maintenance of events (or entities). *Influential causality* modifies features of events by increasing, decreasing or preserving their values.

Algorithms for extracting causal relations from texts adopt either a pattern-based or supervised-learning approach. The pattern-based algorithms in [8,9] extract text segments that match hand-crafted patterns, e.g. "*is the result of*". These techniques detect explicit causal relations from Wall Street Journal (WSJ) articles and from Medline abstracts with precisions of 25% and 76.8% respectively. Their corresponding recall scores are respectively 68% and 75.9%. In the supervised approach of [7], sentences in the L.A. Times corpus that contain explicit causal verbs (e.g. *"to cause"*) are manually annotated as positive or negative examples of causal relations. They are used to train a decision tree classifier, which detects new relations with a precision of 73.9% and recall of 88.7%. In [2], a support vector machine (SVM) trained over the manually annotated SemEval 2007 Task 4 corpus achieved an accuracy of 77.5% in identifying cause-effect noun pairs. SVMs are also used in [4], where they are trained over manually annotated texts of the WSJ, and detect causal relations with a precision of 24.4% and recall of 79.7%. Another supervised-learning approach is that of [14]. An SVM, trained on sub-graphs from annotated WSJ sentences, identifies causal relations with a precision of 26% and recall of 78%.

Existing algorithms have been applied solely to large general-purpose texts or to bio-medical documents. The discovery of causal relations from other domains (e.g. corporate documents) poses new challenges yet to be addressed. Readily-available resources (e.g. hand-crafted patterns and annotated data), which are extensively used by traditional algorithms, are inexistent in many domains. Also, domain-specific texts are sparse, and negatively impact the performance of relation extraction techniques [6]. In addition, these texts are rife with implicit causal relations, which are realized by implicit verbal and non-verbal causal patterns. Implicit patterns and relations are more complex and difficult to detect than their explicit counterparts [7], traditionally extracted by current algorithms. Furthermore, extant algorithms are unable to precisely disambiguate ambiguous causal relations, which are realized by polysemous patterns.

We address these challenges by developing a framework for extracting explicit and implicit causal relations from domain-specific texts in a minimally-supervised fashion. Our proposed approach is described in the next section.

3 Methodology for Extracting Causal Relations

Our overall framework for mining causal relations from a domain-specific corpus is depicted in Figure 1 (dotted arrows correspond to inputs to the various phases, filled and solid arrows represent the outputs). To circumvent the data-sparsity issues of domain-specific texts, we first acquire a set of causal patterns from Wikipedia. We choose Wikipedia as a knowledge-base since its large size offers ample evidence for accurate statistical inferences. As it is a broad-coverage resource, it is also likely to contain the wide variety of explicit and implicit linguistic patterns that express causality. In addition, Wikipedia is readily-available, and has been successfully employed in NLP applications [12]. We transform Wikipedia's sentences into lexico-syntactic patterns in the Pattern Acquisition phase (Section 3.1). In Causal Pattern Extraction (Section 3.2), a minimally-supervised algorithm selects those patterns that encode causality. The harvested patterns are used during Causal Relation Extraction (Section 3.3) to discover causal relations from domain-specific, sparse documents. We refer to these documents as the *target texts*.

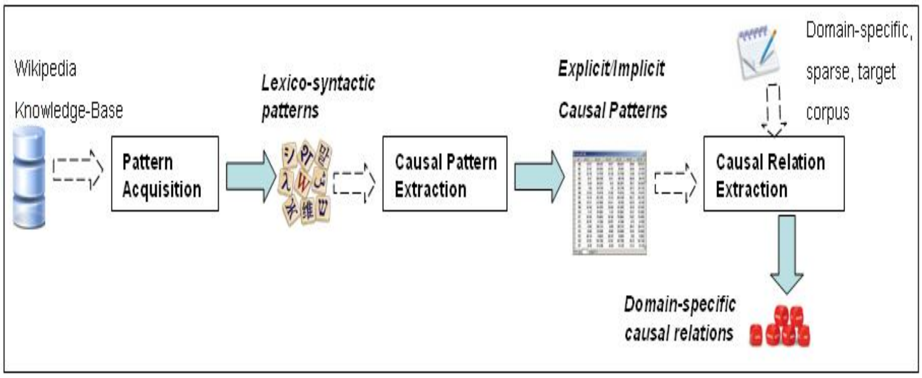

Fig. 1. Overall architecture of framework for extracting domain-specific causal relations

3.1 Pattern Acquisition

We syntactically parse Wikipedia's sentences, and represent the relation between each event-pair as the shortest path that connects the pair in the parse trees[2]. Such a path corresponds to a lexico-syntactic pattern. Sample (lexico-syntactic) patterns, the pairs they sub-categorize in Wikipedia, and the pair-pattern frequency are shown in the 4th, 3rd, 2nd and 1st columns of Figure 2. These patterns encode different semantic relations, as expressed in their corresponding Wikipedia sentences. The first 2 patterns are respectively inferred from sentences *S1=“hurricanes are the most severe climatic disturbance in this area and have been known to cause extensive damage”* and *S2*=“*hiv, the virus which causes aids, is transmitted through contact between blood*”. They express causal relations between the pairs “hurricane-damage” and “hiv-aids”. The 3rd

[2] At this stage, we only consider nominal events, corresponding to noun phrases.

pattern, derived from *S3="the poem consists of five stanzas written in terza rima*", expresses a part-whole relation between the pair "stanza-poem". ARG1 and ARG2 are generic placeholders, representing the pairs connected by the patterns.

```
14 | hurricane | damage | ARG1+nsubj < cause > dobj+ARG2
11 | hiv | aids | ARG1+nsubj < cause > dobj+ARG2
5  | stanza | poem | ARG2+nsubj < consist > prep+of+pobj+ARG1
```

Fig. 2. Lexico-syntactic patterns, pairs, and statistics from Wikipedia

Compared to the conventional surface-strings that existing algorithms employ to represent their patterns, our lexico-syntactic patterns neutralize word order and morphological variations. Thus, they alleviate the manual authoring of a large number of surface-patterns. For example, we derive a single general pattern to represent the relations in *S1* and *S2*. Our patterns also capture long range dependencies, regardless of the distance between related pairs and their positions in surface texts, for e.g. between "hurricane-damage" in *S1*.

3.2 Causal Pattern Extraction

We develop a minimally-supervised algorithm for determining which of the patterns are causal. Unlike traditional algorithms, it does not require annotated training data. It is initialized with cause-effect pairs (e.g. "hiv-aids"), called seeds. Our algorithm then starts by identifying patterns in Wikipedia that connect these pairs (seeds). The reliability, *r(p)*, of a pattern, *p*, is computed with equation (1) [13]. It measures the association strength between *p* and pairs, *e*, weighted by the pairs' reliability, *r(e)*. Initially, *r(e)*=1 for the seeds. In (1), *E* refers to the set of pairs, and *pmi(e,p)* is the point-wise mutual information [5] between pattern *p* (e.g. "*cause*") and pair *e=x-y* (e.g. "hiv-aids")

$$r(p) = \frac{\sum_{e \in E} \left(\frac{pmi(e,p)}{max_{pmi}} \times r(e) \right)}{|E|} \tag{1}$$

Then, we select the top-*k* most reliable patterns, and identify other pairs that they connect in Wikipedia. The pairs' reliability is estimated using equation (2). In the next iteration, we select the top-*m* most reliable pairs, and extract new patterns that connect them.

$$r(e) = \frac{\sum_{p \in P} \left(\frac{pmi(e,p)}{max_{pmi}} \times r(p) \right)}{|P|} + purity(e) \tag{2}$$

The reliability, *r(e)*, of a pair, *e*, is defined by two components. In the first component, which is analogous to the pattern reliability, |*P*| is a set of patterns. The second component determines the pair's *purity* in instantiating a causal relation. It stems from the Latent Relation hypothesis that pairs which co-occur with similar patterns instantiate similar semantic relations [15]. Thus, if a pair, $e=x\text{-}y$, is connected by a pattern, p_{ref}, which also connects our initial seeds, then e instantiates causality. As p_{ref}, we choose the pattern "*caused by*" since it explicitly and unambiguously expresses causality, and specifies the causal link between a distinct causal-agent and an effect. Also, *"caused by"* was found to co-occur with all our seeds. The *purity* of a pair $e=x\text{-}y$ (e.g. "rain-flooding") is then calculated as its probability of being sub-categorized by p_{ref}. This is obtained by querying the Yahoo! search engine[3] with $q=$ "$y\ p_{ref}\ x$" (e.g. *"flooding caused by rain"*), and determining the fraction of the top-50 search results that contains phrase q in their summaries. In this way, invalid pairs, for e.g. "trail-summit", identified by ambiguous causal patterns, for e.g. "*lead to*" (*"trail leads to summit"*), are assigned lower purity values. This is because queries like $q=$*"summit caused by trail"* (formed by the pair "trail-summit" and p_{ref}) are incoherent, and are unlikely to return any search results. The overall reliability scores of these invalid pairs will then be smaller, and they will be discarded. Otherwise, the invalid pairs will be selected in the next iteration, and incorrect patterns connecting them, for e.g. "*pass through*" (*"trail passes though summit"*), will be extracted, degrading our performance. Conversely, valid pairs will be awarded much higher purity values, increasing their reliability scores. They are selected in the next iteration, enabling our algorithm to extract both explicit and implicit causal patterns that connect them.

Our recursive procedure of learning new patterns from pairs, and vice-versa is repeated until a suitable number, t, of causal patterns have been harvested. Parameter values for k, m, and t will be defined during Experimental Evaluation (Section 4.2). Figure 3 shows 2 example patterns extracted by our minimally-supervised algorithm from Wikipedia. ARG1 and ARG2 respectively denote the cause and effect events.

```
ARG1+nsubj < induce > dobj+ARG2
ARG2+nsubj < increase > prep+by+pobj+ARG1
```

Fig. 3. Example of causal patterns learnt from Wikipedia

The 1st pattern explicitly indicates causality with a causal verb, namely "*induce*", as in "*a non-zero current induces a magnetic field by Ampere's law*". The 2nd pattern implicitly expresses causality with a resultative/instrumentative verb, viz. *"increase (by)"*, as in "*the population of the state of Nebraska was increased by positive birthrates*". In this sentence, "*increase (by)*" biases the reader into ascribing the causal-agent role to "positive birthrates", as in "*positive birthrates caused an increase in the population"*.

3.3 Causal Relation Extraction

We use the reliable causal patterns harvested from Wikipedia to extract relations from a domain-specific, sparse (target) corpus. A domain-specific causal relation is made up of a causal pattern and the events that it connects in the target corpus. We consider both nominal events like noun phrases (e.g. "loose connection"), and verbal events

[3] Using the Yahoo! Boss API, http://developer.yahoo.com/search/boss/boss_guide/

like verb phrases (e.g. "replacing the cables"). Figure 4 illustrates a domain-specific causal relation, extracted from the sentence *S4*="*replacing the cables generated intermittent x-rays"* in the target corpus.

```
<pattern = "generate" , cause="replace cable"  , effect="intermittent x-ray">
```

Fig. 4. Domain-specific causal relation triples

4 Experimental Evaluation

This section describes experiments to evaluate the performance of our approach in extracting causal relations from domain-specific, sparse texts.

4.1 Pattern Acquisition

Using the Stanford parser [11], we syntactically parsed the sentences of the English Wikipedia collection [16] (August 2007 dump, around 500 million words). We identified 2,176,922 distinct lexico-syntactic patterns that connected 6,798,235 distinct event-pairs. Sample patterns, event-pairs they sub-categorized and the pair-pattern co-occurrence frequencies were shown in Figure 2.

4.2 Causal Pattern Extraction

To identify patterns that express causality, we implemented a minimally-supervised algorithm. It was initialized with 20 seeds. Seeds were cause-effect pairs that unambiguously instantiated causal relations (e.g. "bomb-explosion") and that occurred at least 3 times in Wikipedia. The 1st iteration of our algorithm extracted 10 most reliable patterns connecting the seeds. Then, 100 other pairs that were also connected by these patterns were extracted. In subsequent iterations, we identified 5 additional patterns (with the previously extracted pairs), and 20 additional pairs (with the previously extracted patterns). The values of parameters k and m (Section 3.2) were therefore $k=|P|+5$ and $m=|E|+20$, where $|P|$ and $|E|$ are respectively the number of previously collected patterns and pairs.

The recursive process of learning new pairs from patterns, and vice-versa was repeated until we observed a drop in the quality of the harvested patterns (e.g. when non-causal patterns were extracted). The performance peaked in the 16th iteration, where we harvested t=81 causal patterns[4] from Wikipedia. Examples are in Table 1. Fifteen of the patterns were causal verbs that explicitly indicated causality, for e.g. "*induce*" (column T0). Resultative/instrumentative verbs, for e.g. "*increase*", which implicitly expressed causality, accounted for 50 patterns (column T1). Fourteen non-verbal, implicit causal patterns were also found (column T3). They included nominalizations of resultative/instrumentative verbs, for e.g. "*increase in*"; adjectives, for e.g. "*responsible for*"; and prepositions, for e.g. "*due to*". Implicit causal patterns,

[4] Precision was roughly estimated as 79%.

Table 1. Explicit and implicit causal patterns extracted from Wikipedia

T0 (15/81=18.5%)	T1 (50/81 =61.7%)	T2 (2/81=2.5%)	T3 (14/81=17.3%)
cause (by, of)	affect (with, by)	mar (by)	drop (in)
induce	decrease (by, to)	plague (by)	due to
lead to	increase (by, to)		increase (in)
result in	inflict (by, on)		rise (in)
spark	limit (by, to)		source of
trigger	prevent (by, to)		responsible for

which made the causal-agent inseparable from the result (effect), were least frequent. Only 2 such patterns, viz. "*mar (by)*" and "*plague (by)*", were detected (column T2). We can deduce from these results that causality in the English Wikipedia corpus is more commonly expressed by implicit causal patterns, particularly resultative/instrumentative verbs, than by explicit causal verbs.

4.3 Causal Relation Extraction

The domain-specific (target) corpus from which we extracted causal relations contained 32,545 English documents (1.1 million words). The documents described customer complaints and engineers' repair actions on professional medical equipment. They were linguistically pre-processed (e.g. to derive syntactic information) prior to their analysis for relation extraction.

Out of the 81 causal patterns from Wikipedia, 72 were found to connect nominal and verbal events in the target corpus, yielding a total of 9,550 domain-specific causal relations. Examples are presented in Table 2. The 3rd column shows the causal patterns that realized these relations. The 2nd column gives the percentage of the 9,550 relations that contained these patterns. The domain-specific causal relations are

Table 2. Domain-specific causal relations from target corpus

Id	Freq (%)	Causal Pattern	Linguistic realization
T1	55	*destroy*	*"short-circuit in brake wiring destroyed the power supply"*
		prevent	*"message box prevented viewer from starting"*
		exceed	*"breaker voltage exceeded allowable limit"*
		reduce	*"the radiation output was reduced"*
T0	23	*cause (by)*	*"gray lines caused by magnetic influence"*
		induce	*"bad cable extension might have induced the motion problem"*
T3	21	*due to*	*"replacement of geometry connection cable due to wear and tear"*
		drop in	*"there was a slight drop in the voltage"*
T2	1	*mar*	*"cluttered options mars console menu"*

depicted in the 4th column as their linguistic manifestations in the target corpus. The 1st column is an identifier for each pattern/relation group.

The most frequent patterns, participating in around 55% of the extracted relations, were resultative/instrumentative verbs (e.g. *"destroy"*, *"exceed"*), which implicitly expressed causality. The high frequency of these patterns in our target corpus could be attributed to their common use in describing product failures (e.g. *"short-circuit in brake wiring destroyed the power supply"*), and in reporting observations on product behavior (e.g. *"breaker voltage exceeded allowable limit"*). We observed that these relations had optional causal-agents, and that they could not be conclusively classified in Barriere's taxonomy of causality [1]. When the causal-agents were specified, as in *"[short-circuit in brake wiring$_{causal\text{-}agent}$] destroyed the power supply"*, the relations indicated the creation or destruction of events. Thus, according to Barriere's taxonomy, they established *existential causality*. However, when their causal-agents were unspecified, they expressed changes in magnitude, as in *"breaker voltage exceeded allowable limit"*. Then, these relations established *influential causality* based on Barriere's taxonomy. In the domain of Product Development/Customer Service (PD-CS), these relations provide useful information for product quality improvement. The next most frequent patterns, appearing in around 23% of the extracted relations, were explicit causal verbs (e.g. "*cause by*", *"induce"*). These relations always specified a distinct causal-agent and its effect, as in *"[gray lines$_{effect}$] caused by [magnetic influence$_{causal\text{-}agent}$]"*. They established *existential causality* since they were realized by verbs synonymous with "*to cause (to exist)*". In the PD-CS domain, these relations can be exploited to facilitate the diagnosis procedure of engineers. Around 21% of the relations were realized by non-verbal implicit causal patterns, such as noun phrases (e.g. *"drop in"*) and prepositions (e.g. *"due to"*). These relations were not conclusively classifiable in Barriere's taxonomy. When they were realized by noun phrases, they described unexplained phenomena with unknown causes. Hence, their causal-agents were often unspecified, as in *"there was a slight drop in the voltage"*. The relations then established *influential causality*. Conversely, when they were realized by prepositions, they established *existential causality* since they always specified a causal-agent that brought a resulting situation (effect) into existence, as in *"[replacement of geometry connection cable$_{effect}$] due to [wear and tear$_{causal\text{-}agent}$]"*. In the PD-CS domain, these relations provide pertinent information on customer dissatisfaction and on repair actions of engineers. The rarest patterns, appearing in less than 1% of the relations, were those that implicitly expressed causality by making the causal-agent inseparable from the result (e.g. *"mar", "plague"*). For example, in "*cluttered options mars console menu*", *[cluttered options$_{causal\text{-}agent}$]* is an integral part of the *[marred console menu$_{effect}$]*. These relations could not be unambiguously classified in Barriere's taxonomy. They established neither *existential* nor *influential causality*.[5]

Nine (out of 81) patterns from Wikipedia were not found in the target corpus. They included explicit causal verbs, for e.g. "*spark*"; implicit resultative/instrumentative verbs, for e.g. "*end*" and "*outstrip*"; and non-verbal expressions for e.g. "*growth*". These patterns were unlikely to occur in our corpus of customer complaints and of engineers' repair actions.

---

[5] To some extent, they can be treated as influential causality.

To evaluate the performance of our approach, we randomly selected 3000 of the extracted causal relations, equally distributed across the groups T0, T1, and T3 of Table 2 (i.e. 1000 explicit relations with causal verbs, 1000 implicit relations involving resultative/instrumentative verbs, and 1000 implicit relations with non-verbal expressions). Relations in the group T2 were omitted as they were too sparse. The evaluation set of 3000 relations was manually inspected by human judges, and 2295 were deemed to correctly express causality (*true_positive*). The remaining 705 relations were incorrect (*false_positive*). They were realized by causal patterns, for e.g. *"due to"*, which did not connect valid events in the target corpus, as in "*screen due to arrive today*". Using equation (3), we then estimated the precision of our approach as 76.5%. To determine the recall, a gold-standard of 500 valid causal relations was manually constructed from a subset of the target corpus. The sub-corpus was then automatically analyzed by our approach. We detected 410 of the gold-standard relations (*true_positive*), but failed to discover remaining 90 (*false_negative*). The recall was computed as 82% using equation (4).

$$\text{precision} = \frac{\text{true\_positive}}{\text{true\_positive} + \text{false\_positive}} \tag{3}$$

$$\text{recall} = \frac{\text{true\_positive}}{\text{true\_positive} + \text{false\_negative}} \tag{4}$$

Our scores of precision (76.5%) and recall (82%) compare favorably with those reported by other state-of-the-art algorithms [2,4,7,8,9,14]. These latter techniques, however, extensively relied on manually-crafted knowledge (e.g. annotated data). Also, they focused solely on detecting causal relations that were explicitly realized by causal verbs in large corpora. Our approach, on the other hand, extracts both explicit and implicit causal relations, expressed by verbal and non-verbal causal patterns, from sparser texts. In addition, it is minimally-supervised, and alleviates the need for hand-coded knowledge, which is expensive to generate in most domains.

We conducted another set of experiments to illustrate the significance of our event-pair *purity* measure (Section 3.2) in extracting the most reliable causal patterns and relations. We re-implemented our minimally-supervised algorithm such that the reliability, *r(e)*, of an event-pair, *e*, was estimated with only the 1st component of equation (2). That is, the *purity* estimation was omitted. The pattern reliability measure in equation (1) was unchanged. Our algorithm was initialized with seeds, and was ran over the Wikipedia collection as before. We observed that invalid pairs, for e.g. "street-exit", which were connected by many ambiguous patterns, for e.g. *"lead to", "result in" ("street leads to exit", "street results in exit")*, were awarded higher reliability scores than valid ones. Subsequently, these invalid pairs were selected, and other incorrect patterns that connected them, for e.g. *"at" ("street at exit")*, were in turn extracted. The quality of the harvested patterns deteriorated in much earlier iterations. Optimal performance was achieved in the 7th iteration, where 34 causal patterns were identified[6]. We also found out that the new implementation failed to discover many of the implicit causal patterns that had been detected by our original algorithm. We used

[6] Precision was roughly estimated as 65%.

the 34 causal patterns discovered by the new implementation from Wikipedia to extract domain-specific causal relations from the target corpus. The precision and recall were calculated by equations (3) and (4) as 65.7% and 68.3% respectively. These scores indicate a drop of performance compared to our original algorithm, and suggest that the new implementation is less accurate in extracting causal patterns and relations. The results reveal that our event-pair *purity* measure for disambiguating polysemous patterns is crucial to reliably harvest the wide range of patterns participating in both explicit and implicit causal relations.

5 Conclusion

In this paper, we have described an approach for mining causal relations from texts. The novelty in our technique lies in the use of Wikipedia as a knowledge-base. We proposed a minimally-supervised algorithm, which unlike previous techniques, alleviates the need for manually-annotated training data. Our algorithm employs sophisticated statistical analyses for disambiguating and extracting both explicit and implicit causal patterns from Wikipedia. The causal patterns from Wikipedia are then used to detect causal relations from a domain-specific corpus of corporate documents. This strategy, of applying the knowledge acquired from one domain (Wikipedia) to another domain (corporate documents), is based on previous studies in domain-adaptation. It overcomes the data-sparsity issues that domain-specific texts pose to relation extraction techniques. Evaluations results on real-life data reveal that our approach achieves state-of-the-art performance in discovering explicit and implicit causal relations from domain-specific, sparse texts. As future work, we are investigating whether the performance can be improved by using certain keywords that indicate causality, for e.g. modal verbs (“can”, “will”, “may”) and subordinating conjunctions (“after”, “as”, “because”). We are also interested in question-answering systems based on causal relations for answering “why” questions.

Acknowledgment. This work is being carried out as part of the DataFusion project, sponsored by the Netherlands Ministry of Economic Affairs, Agriculture and Innovation, under the IOP-IPCR program.

References

1. Barriere, C.: Hierarchical refinement and representation of the causal relation. Terminology 8(1), 91–111 (2002)
2. Beamer, B., Bhat, S., Chee, B., Fister, A., Rozovskaya, A., Girju, R.: UIUC: A knowledge-rich approach to identifying semantic relations between nominals. In: 4th International Workshop on Semantic Evaluations, pp. 386–389 (2007)
3. Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., Vaughan, J.W.: A theory of learning from different domains. Machine Learning 79(1), 151–175 (2010)
4. Bethard, S., Corvey, W., Klingenstein, S., Martin, J.H.: Building a corpus of temporal-causal structure. In: 6th International Language Resources and Evaluation Conference (2008)
5. Church, K.W., Hanks, P.: Word association norms, mutual information, and lexicography. Computational Linguistics 16(1), 22–29 (1990)

6. Cimiano, P., Pivk, A., Schmidt-Thieme, L., Staab, S.: Learning taxonomic relations from heterogeneous sources of evidence. In: Buitelaar, P., Cimiano, P., Magnini, B. (eds.) Ontology Learning From Text: Methods, Evaluation, and Applications. IOS Press, Amsterdam (2005)
7. Girju, R.: Automatic detection of causal relations for question answering. In: ACL 2003 Workshop on Multilingual Summarization and Question Answering, pp. 76–83 (2003)
8. Khoo, C.S.G., Kornfilt, J., Oddy, R.N., Myaeng, S.H.: Automatic extraction of cause-effect information from newspaper text without knowledge-based inferencing. Literary and Linguistic Computing 13(4), 177–186 (1998)
9. Khoo, C.S.G., Chan, S., Niu, Y.: Extracting causal knowledge from a medical database using graphical patterns. In: 38th Annual Meeting on Association for Computational Linguistics, pp. 336–343 (2000)
10. Khoo, C., Chan, S., Niu, Y.: The semantics of relationships: an interdisciplinary perspective. Kluwer Academic, Dordrecht (2002)
11. Klein, D., Manning, C.D.: Accurate unlexicalized Parsing. In: 41st Annual Meetingon Association for Computational Linguistics, pp. 423–430 (2003)
12. Medelyan, O., Milne, D., Legg, C., Witten, I.H.: Mining meaning from Wikipedia. International Journal of Human-Computer Studies 67(9), 716–754 (2009)
13. Pantel, P., Pennacchiotti, M.: Espresso: Leveraging Generic Patterns for Automatically Harvesting Semantic Relations. In: COLING/ACL 2006, pp. 113–120 (2006)
14. Rink, B., Bejan, C., Harabagiu, S.: Learning Textual Graph Patterns to Detect Causal Event Relations. In: 23rd Florida Artificial Intelligence Research Society International Conference (2010)
15. Turney, P.D.: The latent relation mapping engine: algorithm and experiments. Journal of Artificial Intelligence Research 33, 615–655 (2008)
16. Wikipedia, XML format (August 2007), `http://ilps.science.uva.nl/WikiXML`