

Supporting Journalistic News Angles with Models and Analogies

Bjørnar Tessem

Dept. of Information Science and Media Studies
University of Bergen
Bergen, Norway
Email: bjornar.tessem@uib.no

Andreas L. Opdahl

Dept. of Information Science and Media Studies
University of Bergen
Bergen, Norway
Email: andreas.opdahl@uib.no

2019

Abstract

News angles are approaches to content presentation in journalism, where the journalist chooses which facts of an event to present. The News Angler project investigates how to computationally support the creation and selection of original news angles for a news event based on information from big data sources. At least two creative approaches are possible. One is to maintain a library of well-known news angles represented in a suitable modeling language, matching published reports on a current event to news angles in order to identify possible angles that have not yet been used. A second approach is not to represent news angles explicitly, instead matching the current event with previous events, and transferring angles from past to present reports by similarity and analogy. Both approaches are described and technologies needed to proceed in either direction are discussed.

1 Introduction

In the journalist's daily work they need to find the most interesting way of presenting a newsworthy event for the readers. A *news angle* is *how a journalist or other news worker makes an event (or situation) interesting for an audience* [1]. It is a general way to turn an event into a report (or news story) that will interest an audience. It may be obvious what this angle should be, but in many circumstances only a creative, original angle would warrant a new report on the same event. Simultaneously, tight deadlines and high productivity expectations (in the form of numbers of reports) may suppress natural creativity in search for good and original angles.

In the News Angler project we therefore aim at developing methods and tools to support journalists in their work, by making them able to find new ways of looking into an event. This will require computerised means of knowledge representation, collection, organisation, and reasoning to support the identification of novel, useful, and readable angles on an event, leveraging a big data repository organised as a knowledge graph of facts represented as RDF triples [2]. News Angler is a collaboration between a university research group and Wolftech Broadcast Solutions ¹, a developer of newsroom systems for the international market. Our collaboration aims to extend Wolftech’s commercial tool for newsrooms with support for creative journalism.

Knowing how to select a suitable angle on a current event is part of the human journalistic skill set, along with the ability to compose new angles by combining simpler ones. In this short paper, we ask *How can a newsroom tool automatically suggest original news angles on current or otherwise interesting events to journalists?* We attempt to answer this research question by exploring two alternative approaches: news-angle modelling and analogical reasoning (which can possibly also be used in combination).

News-angle modelling: The first approach is to maintain a library of news angles, which externalises the knowledge that is normally included in reports using each particular angle. If we are able to recognize which angles have been used already for a news event, a journalist can then be suggested to use one of the other alternatives, in particular if we are able to match the available knowledge about the event to another angle.

Analogical reasoning: A second approach is to forgo the news angle library, instead identifying original angles by identifying similarities or, even more challenging, analogies between a current and previous events. Furthermore, similarities can be identified among individual reports from two similar events, suggesting a similarity-based match. The unmatched reports of a similar event may then suggest a new news angle for the current event.

For a computer to support these approaches, it needs access to knowledge about news events: e.g., the people and other entities (animals, organisations) involved, their relations, and their attributes, as well as important sub-events they undergo and how they are temporally and causally related. The representation of such events is a central research topic for News Angler [1].

In section 2, we will explain the news angle concept and describe a typical use case. In sections 3 and 4 we describe the two approaches and how they can be realised. In the discussion section 5 we summarize the approaches, and compare the modeling and computational issues to be handled by the two approaches.

2 News Angles

In previous work [1], we have discussed the idea of news angles in the journalism literature and online sources. Examples are very general angles like the event itself (just report the facts of the newsworthy event), the influence of an event (the effect of an event on something else), proximity (a connection from the event to something near you), milestone (the event made something special be achieved), and many others. This is how an *event* becomes the nucleus of a *news report*: unless the event is so significant that it is newsworthy in itself,

¹<http://www.wolftech.no/>

it must be made newsworthy by the match with a good and hopefully original angle.

2.1 Use Case

In collaboration with Wolftech, we have already developed a first prototype, News Hunter, [3] that is able to harvest news-relevant information from the web, represent the information items semantically as facts in a knowledge graph, and make the information available to journalists that work on related news events. We are now re-implementing our prototype using proper big-data technologies to make it more scalable and augmenting it with support for news angles. Drawing on the prototype experiences, we have detailed the potential functionality of News Hunter through eleven use cases with extensions and variants. A particular important one is *What's my angle?*, which comprises the following steps:

1. A journalist types a working news report into the front end.
2. News Hunter lifts the working report and returns IRIs for named entities, concepts/topics/categories, relations, and sentiments in the report.
3. News Hunter retrieves angles that fit the working report.
4. News Hunter recommends the most suitable angles.
5. The front end makes recommendations to the journalist.

It can be extended by the *Fill in my angle* use case as follows:

6. The journalist selects a recommended angle.
7. The front end retrieves from the knowledge graph a subgraph of related named entities, concepts/topics/categories, dates and locations, relevant to the report and that fit the selected angle.
8. The front end presents the facts to the journalist.

This paper outlines two alternative solutions to these steps.

2.2 An example

As an example, consider an imagined news report, where a journalist in the local newspaper covering the Valderdal rural area, writes the following (in a more lively language): Football player Hans Hansen owns a cabin in Valderdal. Hans Hansen visits his cabin regularly. Hans Hansen plays on Vålerborg football team, which is coached by John Johnsen, and they will play the cup final the next weekend. An additional point is that Hans Hansen's mother Anna Hansen, who happens to live in Valderdal, recently married the coach John Johnsen, so that Hans Hansen now plays for his stepfather's team.

This report might instantiate two different news angles. The first is *Local Person in Sports Event*. Important entities in this news angle are the local person and, if a team sport is involved, the team of the local person. Furthermore, the relation of the player to the local area is important: in this case derived through the location of the cabin combined with the ownership of the cabin.

Additional facts about the sports event itself, with temporal and location attributes and relations, e.g., when and where the cup final is going to happen, are also in this news angle.

The second news angle is the *Possible Nepotism* angle, in which one powerful person (coach Johnsen) controls something of value (a place on the team), which a gaining person (player Hansen) attains access to. Furthermore, there must be a familiar or other private relation (stepson/stepfather) between the powerful and the gaining person.

2.3 Representing events

The News Hunter platform will maintain a dynamic knowledge graph that is continuously updated with published news articles/reports, social media messages and other potentially news-relevant information items, and enriched from relevant sources of background knowledge (for example from DBPedia² and Wikidata³). The knowledge graph will contain potentially newsworthy information as *event (sub-)graphs*, each comprising facts about entities, relations, and attributes related to the event.

An important research problem in News Angler is how to constrain event graphs to remain computationally feasible, without excluding context information that might suggest highly interesting and surprising news angles.

2.4 Representing news angles

As mentioned, we consider a news angle to be a general view that could match a particular event. The requirements for a match are that certain types of entities (for example agents such as people, organisations, teams...) are present and that there are certain kinds of relations between them, e.g., that particular types of events have happened or are happening.

For example, the nepotism angle can be described more precisely as follows:

```
Person(powerfulPerson)
Value(value)
controlsValue(powerfulPerson, value)
Person(gainingPerson)
privateRelation(powerfulPerson, gainingPerson)
accessesValue(gainingPerson, value)
```

2.5 Representing fabulas

To be able to work with news angles computationally we need to develop a firmer conceptualisation of how they match events in the knowledge graph. For this purpose, we borrow concepts from literature theory, as Gervás does in his discussion of computational storytelling [4]. In literature theory, a story consists of the *fabula*, which comprises facts about the event that the story (in our case: report) is about, and the *discourse*, which is how the facts in the fabula are presented (or reported). Because the knowledge graph may contain many more facts about entities, relations, and attributes than the journalist wants to present in a report, we use the term *fabula (sub-)graph* about that restricted part of an

²<https://wiki.dbpedia.org/>

³<http://wikidata.org>

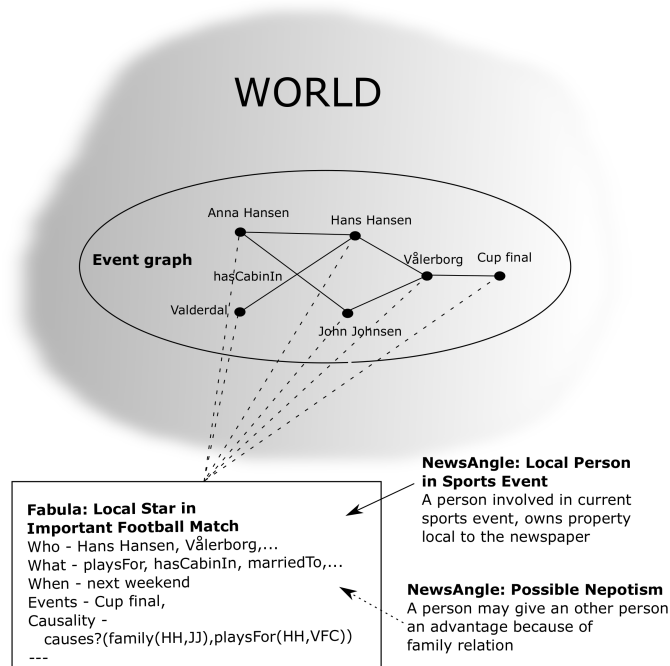


Figure 1: A conceptualisation of news angles with the world, an event graph, a fabula, and two news angles.

event (sub-)graph that forms the basis of a report, perhaps after having been extended by the journalist.

In Figure 1 we show the instantiated fabula “Local star Hans Hansen in the cup final”. The fabula consists of six entities — three human agents, an organisation, a location, and an event — as well as relations between them, all coming from the event graph. The construction of the fabula draws on a combination of requirements from the nepotism and local person news angles, and includes a subset of entities and relations in the event graph (like “has cabin in” or “plays for”). Notice, that already here we have identified general relations that are relevant for many events (like “family of”).

3 A Model-Based Approach

To represent the news-relevant world, its many events, and news angles precisely, we use OWL — the Web Ontology Language — to define useful terms and RDF — the Resource Description Framework — to represent the facts in our knowledge graph [2].

3.1 Harvesting and lifting news items

Following the earlier prototype [3], the News Hunter platform will continually harvest RSS messages, tweets and other kinds of news-relevant information

items from the web. Each item is then *lifted* into a small knowledge graph, an *item (sub-) graph*, using standard NLP techniques such as named entity, concept/topic/category, relation, and sentiment extraction [5, 6, 7]. Because they link the same named entities, concepts, and so on, the many item graphs will overlap to form a big connected knowledge graph. An ontology based on our earlier prototype, shown in Figure 2, provides a basis for this work.

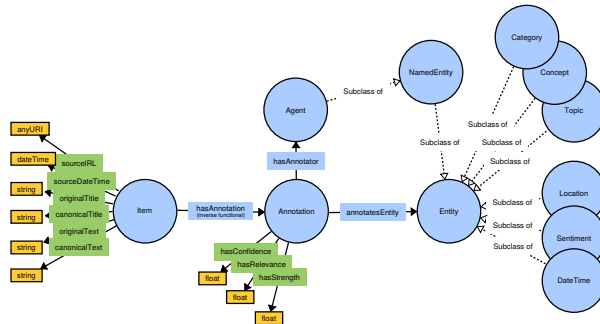


Figure 2: Ontology for representing news items and other texts.

Items represent the news-relevant texts themselves, which can for example be a: published article/report, tweet, Facebook message or report in preparation. An item has a title and a text, both in the original language and possibly translated into a canonical platform language (such as English). The results of semantic analysis are associated with the *Item* through *Annotations*, each having a confidence, relevance, and strength. Each annotation is in turn related to an *Entity* in the knowledge graph, which can be either a *NamedEntity*, *Concept/Topic/Category*, *Location*, *DateTime* or *Sentiment*. Because every *Entity* has an IRI, they can be enriched semantically with further facts (or RDF triples) taken from the Linked Open Data (LOD) cloud. Not shown in Figure 2 are *RelationAnnotations* that describe a semantic relationship between two annotation entities of the same item.

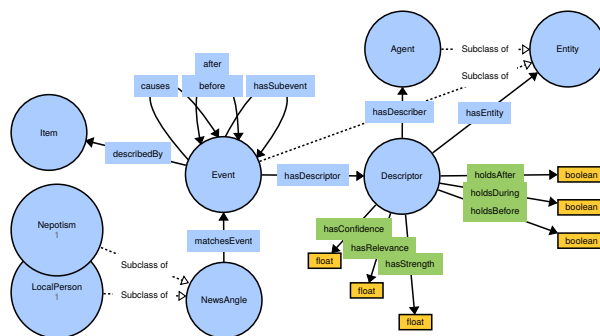


Figure 3: Ontology for representing events, their event graphs, and their matching news angles.

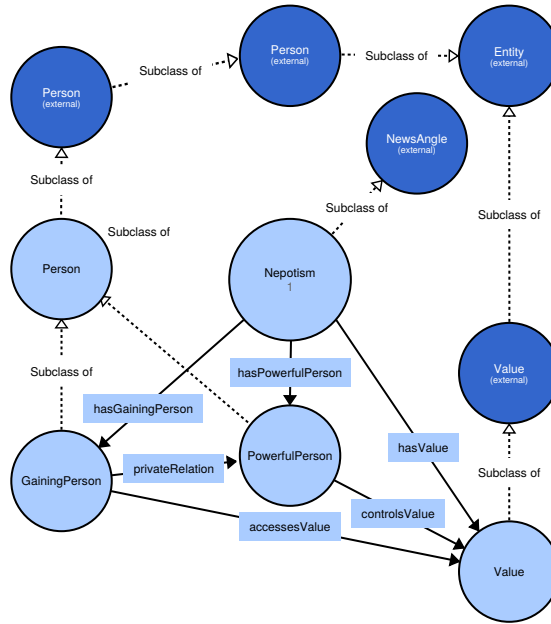


Figure 4: An OWL representation of the nepotism angle.

3.2 Detecting events

The individual item (sub-)graphs according to Figure 2, must be aggregated to form events, which are also described by knowledge graphs, called *event (sub-)graphs*. Because they are described using the same *NamedEntities*, *Concepts*, *Locations*, *DateTimes* so on, the many event graphs will also overlap and be part of the big connected knowledge graph. Aggregation serves to increase both the trustworthiness, completeness, and precision of the event graphs, which can be further enriched with data from linked-open data sources such as DBPedia⁴ and Wikidata⁵).

Figure 3 illustrates some of the most important concepts that we need in order to support events and match them with news angles. An *Event* is related to all the *Items* that describe it, and from which its event graph has been aggregated. The *Event* also has *Descriptors* that are similar to *Item Annotations*. Each descriptor relates the event to an *Entity* in the knowledge graph, which can be a named entity, concept/topic/category, location, date-time or sentiment as before. Descriptors can represent facts that hold either *before*, *during* or *after* the event. Not shown in Figure 3 are *RelationDescriptors*, similar to *RelationAnnotations*, that describe semantic relationships between pairs of entities in the same event graph. The particular analysis tool used to derive a descriptor is again represented as an *Agent*. Finally, an *Event* may partially match *NewsAngles*. In the figure we see two subclasses of *NewsAngle*, *Nepotism* and *LocalPerson*.

⁴<https://wiki.dbpedia.org/>

⁵<http://wikidata.org>

3.3 Detecting news angles

Detecting a *NewsAngle* on an *Event* now boils down to finding an angle whose restrictions on entities and relations match the entities and relations in the event graph, thus producing a fabula graph. For example, the *Local star Hans Hansen in cup final* fabula, would be an instance of both the *Nepotism* and the *LocalPerson* angles and inherit constraints from both.

The *Nepotism* angle with its constraints is illustrated in Figure 4, and involves a *PowerfulPerson*, a *Gainingperson*, and some entity of *Value*. They must stand in specific *familyRelations*, *controlsValue*, and *accessesValue* relations to one another, as explained in Section 2.4. The *NewsAngle* thus describes some constraints on its instances, which could be the presence of particular types of and relations between entities.

We plan to start by crafting news-angle ontologies by hand before exploring ways to automatically derive news angles from news knowledge graphs. In previous work [1], we have already collected news angle types from the journalism literature, generalising the angles that are available to the journalist.

Assuming we have such a library of news-angle ontologies describing the required entities and relations, the steps of suggesting original news angles can be summarized as: (1) For the whole collection of reports on an event, extract entities and relations existing in the collection, as well as extracting other relevant knowledge. This can be seen as constructing a representation of the event as an event graph. (2) For each report, identify the entities and relations presented in that particular report, i.e, identify the report’s fabula graph. (3) Identify all known news angles that match some report’s fabula graph. (4) Match *unused* news angles to the knowledge in the event graph. (5) Suggest instantiations of entities (including events) and relations in the unused news angles, i.e., a fabula, to the journalist.

4 An Analogy Approach

The model-based approach can be productive, but it also has some computational challenges which need to be solved: It relies on a library of ready-defined news angles (which also may be constraining), and also on semantic lifting of content from items into item graphs and aggregated event graphs.

Analogical reports about other newsworthy events may provide an alternative way of identifying new angles. The reports themselves can be seen as examples of angles that could be pursued by the journalist, and more or less advanced similarity measures for documents and document collections can be used to identify candidates for analogies, not by naming the news angles, but by presenting an example that could to some extent be transferred to the new report. This can be handled by different similarity measures, ranging from classic document similarity to structural similarity as found in several analogical reasoning techniques. The latter is an approach we also find in work by Maiden et al. [8].

Hence, the analogical approach is not constrained by a library of angles. And, while it can be implemented on top of knowledge graphs, it can also use similarity measures and analogy-reasoning techniques that work directly on item texts, thus removing the need for precise semantic lifting.

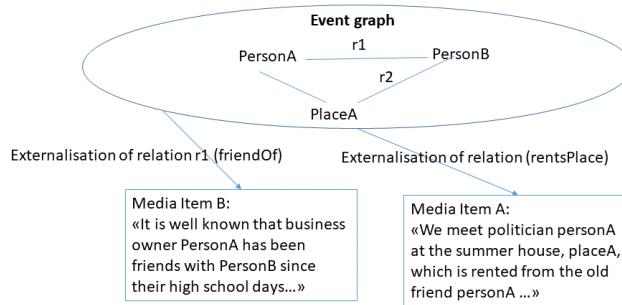


Figure 5: Externalisation of entities and relations

4.1 Assumptions

First, reiterate that any fabula may be considered to be a small knowledge graph. The graph contains entities e_i that are part of the event and also (binary) relations r_j between pairs of entities. These entities and relations are not necessarily known to the News Angler analogy machine but are *externalised* in news items m_k . See Figure 5.

Since the relations r_j may not be explicit and hard to identify from existing technologies, an alternative approach to detecting analogies may be to base a search on the idea that analogical or similar stories have similar externalisations.

4.2 Finding unmatched news items

Let us assume that we have a target (current) event t that we want to find a new angle for, as well as an *identified and analogical* base (previous) event b that already has n_b base news items $m_i \in B, i = 1, \dots, n_b$. We also may assume that t has n_t news items $m_j \in T, j = 1, \dots, n_t$. We must also assume that we have a similarity measure $sim(m_i, m_j) \in [0, 1]$ for each pair of items from the base and target events ($m_i \in B$ and $m_j \in T$). In terms of news items this similarity can be based on item type, text similarity, similarity among identified entities etc. As discussed later in the paper, the similarity measure is an abstraction that can be realised in many different ways, from traditional cosine similarity to analogical similarity.

Now, let A be a binary matrix with entries $a_{ij} = 1$ if there is a *match* between $m_i \in B$ and $m_j \in T$, otherwise 0. This matrix represents the total matching between items in the base and target events. There are several soft constraints that should be fulfilled for a matching to have high quality. These are first of all that we should have maximum similarity among matched items, secondly, items with low similarity should not be matched, even if they are the only candidates left for matching, and thirdly, items should normally not match more than one item. To handle the problem with low similarity we may subtract a constant c_l from all similarity values to ensure that all matched items have a similarity above the limit c_l . To ensure almost one-to-one matching we count the number of 1's in each row (cr_i) and each column (cc_j) of A . The penalty for having more than one 1 in a row or column is c_p . A lower c_p would allow for more freedom

in matches, higher c_p would lead to one-to-one matches.

A matching of presumed high quality is then found by maximizing the objective function

$$\begin{aligned}
 f(A) = & \sum_{i,j} a_{ij} (sim(m_i, m_j) - c_l) \\
 & - \sum_i max(0, cr_i - 1) \cdot c_p \\
 & - \sum_j max(0, cc_j - 1) \cdot c_p
 \end{aligned}$$

When we have found the solution A , there will be items in the base event which are unmatched, i.e., there are rows where all entries in A are 0. These items may represent a new angle. The most original candidate could be the item that has the lowest similarity to any existing item in the target, i.e., has the highest $o(m_i) = 1.0 - \max_j sim(m_i, m_j)$.

4.3 Ranking among many base events

Having established an algorithm for finding the candidates for a news angle in a known analogical base event b , and ranking them using an *originality*-function $o(m_i)$, it is possible to use a general ranking approach to finding the best candidate reports from all possible base events b_k . This requires that we are able to compute an analogical similarity among events based on their total collection of items and identified entities.

For now, let us assume that we are able to compute the similarity $sim(b_k, t)$ for all base events b_k and the target event t . Further, assume that we for each b_k is able to compute an optimal item matching A^k with t . All unmatched items in the various base events will now be candidates for being suggested as an angle. To rank all these candidate items, we may use the similarity of the event they belong to, as well as the originality-measure $o(m_i)$ suggested above. A function used for ranking could be $score(m_i^k) = sim(b_k, t) \cdot o(m_i^k)$. Notice that it is possible to swap the originality measure o for any measure indicating some requirements for the new report item, for example by using click counts for previous reports, or other established knowledge about the journalist and the media users.

4.4 Similarity measures

The discussion on matching above is based on undefined similarity measures for documents. In her classical work on analogy, Gentner [9] claimed that analogies can only be detected based on the structural similarity of object relations in two models, and together with Falkenhainer et al., she proposed an analogical matching approach based on structural similarity, the structure mapping engine [10]. An alternative constraint-satisfaction approach to analogical matching is based on structure, semantics, and pragmatics [11].

Analogical similarity is useful as soon as we have fabulas for every report, but initially we will have to use other approaches. Classical document similarity measures could be a good starting point to establish a baseline for later research

in the project. Traditional information retrieval techniques like cosine similarity, combined with tf-idf are obvious candidates [12].

The use of word embeddings is a recent technology that enables document similarity at a level closer to graph similarities, and thus also may be more of the analogical type than standard cosine similarity. Pre-computed and possibly multi-lingual word embeddings generated from neural networks such as word2vec [13, 14] and GloVe [15] combine precise semantics with efficient similarity calculation as scalar products of dense word vectors. doc2vec [16] captures the semantics of longer texts as similarly comparable vectors, but can only be used for existing collections of aggregated or important texts because new vectors must be learned per document.

Taking it a step further, in the ConceptNet project⁶, Speer et al. [17] describe how word embeddings generated by neural networks can be improved by aligning them with embeddings generated for the same words from knowledge graphs. Algorithms have also been proposed that generate embeddings for nodes and subgraphs of graphs. Such graph-embedding approaches are potentially useful for News Angler, but only, as for analogical reasoning strategies, when we are able to represent news items and events as graphs. Only a few of the proposals around graph embeddings address RDF specifically [18, 19], and none of them support incremental node and subgraph embeddings.

Calculating similarities between pairs of items is the inner loop of this algorithm. They should therefore be stored for reuse by this and other future similarity-based analyses. It is also likely that similarities should be pre-calculated when news items are lifted into the knowledge graph, as part of categorisation and clustering.

5 Discussion

We have presented two different approaches to suggesting news angles to the journalist, a news-angle modelling approach, and an analogical reasoning approach. In addition, we suggest several similarity models that may be used for analogical reasoning: traditional document similarity and tf-idf; word embeddings representations of item texts; knowledge-graph embeddings representations of items; and pure analogical reasoning based on the structure-mapping engine.

Evaluating the different approaches will have to take two different directions. One will be to see if analogical approaches to news angles or a model-based approach will be more successful with practitioners. This type of evaluations can be conducted as empirical investigations in collaboration with our industrial partner Wolftech, and will inform the modeling of news angles as well as the support functionality itself.

As for evaluating the similarity approaches, we need to choose a rich corpus of news articles/reports, grouping them into stories, and selecting additional items (tweets, Wikipedia) of relevance for each. We aim to build such a corpus by selecting content from, for example, the NOW corpus⁷.

Many of the ideas we present here are based on an ability to lift item content into knowledge graphs. How to extract entities like persons and other entities

⁶<http://www.conceptnet.io/>

⁷<https://corpus.byu.edu/now/>

from such content is a problem where well established methods exist [5, 6]. However, News Angler is very much based on being able to extract relations among entities in texts [7]. This is a research topic of relevance in many domains. For example, it is an issue in automated literature-based discovery as applied in medicine [20]. Anyhow, relation extraction is still an immature field of research, where we need to contribute in order to enable news angle suggestions.

6 Conclusion

This paper has presented ongoing work in the News Angler project. Focus has been on supporting creative journalism with use of two different approaches to identifying alternative news angles, one model-based and another analogy-based. A conceptual understanding of the domain is developed, the fundamentals of the two approaches for selecting news angles are presented and future research topics are identified. In particular there is a need for developing algorithms for mining relations in documents, a library of news angles, and strategies for evaluating the approaches selected.

Acknowledgement

The News Angler project is funded by the Norwegian Research Council’s IKT-PLUSS programme as project 275872.

References

- [1] Marc Gallofré Ocaña, Lars Nyre, Andreas Lothe Opdahl, Bjørnar Tessem, Christoph Trattner, and Csaba Veres. Towards a Big Data Platform for News Angles. In *Proceedings of the 4th Norwegian Big Data Symposium (NOBIDS 2018)*, volume 2316, pages 17–29. CEUR Workshop Proceedings, November 2018.
- [2] Dean Allemang and James Hendler. *Semantic web for the working ontologist: effective modeling in RDFS and OWL*. Elsevier, 2011.
- [3] Arne Berven, Ole A. Christensen, Sindre Moldeklev, Andreas L. Opdahl, and Kjell J. Villanger. News hunter: Building and mining knowledge graphs for newsroom systems. In *NOKOBIT*, volume 26, 2018.
- [4] Pablo Gervás. Computational Approaches to Storytelling and Creativity. *AI Magazine*, 30(3):49–62, July 2009.
- [5] David Nadeau and Satoshi Sekine. A survey of named entity recognition and classification. *Linguisticae Investigationes*, 30(1):3–26, August 2007.
- [6] Vikas Yadav and Steven Bethard. A survey on recent advances in named entity recognition from deep learning models. In *Proc. 27th COLING, Santa Fe, New Mexico, August 20-26*, pages 2145–2158, 2018.
- [7] Wilson Wong, Wei Liu, and Mohammed Bennamoun. Ontology learning from text: A look back and into the future. *ACM Computing Surveys*, 44(4):1–36, August 2012.

- [8] Neil Maiden, Konstantinos Zachos, Amanda Brown, George Brock, Lars Nyre, Aleksander Nygård Tonheim, Dimitris Apsotolou, and Jeremy Evans. Making the news: Digital creativity support for journalists. In *Proc. 2018 CHI Conf. on Human Factors in Computing Systems*, pages 475:1–475:11, New York, NY, USA, 2018. ACM.
- [9] Dedre Gentner. Structure-mapping - A theoretical framework for analogy. *Cognitive Science*, 7(2):155–170, 1983.
- [10] Brian Falkenhainer, Kenneth D. Forbus, and Dedre Gentner. The structure-mapping engine: Algorithm and examples. *Artif. Intell.*, 41(1):1–63, 1989.
- [11] Keith J. Holyoak and Paul Thagard. Analogical mapping by constraint satisfaction. *Cognitive Science*, 13(3):295–355, 1989.
- [12] Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütze. *Introduction to Information Retrieval*. Cambridge Univ. Press, New York, 2008.
- [13] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient Estimation of Word Representations in Vector Space. *arXiv:1301.3781 [cs]*, January 2013. arXiv: 1301.3781.
- [14] Xin Rong. word2vec Parameter Learning Explained. *arXiv:1411.2738 [cs]*, November 2014. arXiv: 1411.2738.
- [15] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. Glove: Global Vectors for Word Representation. In *Proc. EMNLP 2014, October 25-29, 2014, Doha, Qatar*, pages 1532–1543, 2014.
- [16] Quoc V. Le and Tomas Mikolov. Distributed representations of sentences and documents. *CoRR*, abs/1405.4053, 2014.
- [17] Robert Speer, Joshua Chin, and Catherine Havasi. Conceptnet 5.5: An open multilingual graph of general knowledge. In *Proc. 21st AAAI, February 4-9, San Francisco, USA.*, pages 4444–4451, 2017.
- [18] Petar Ristoski and Heiko Paulheim. Rdf2vec: RDF graph embeddings for data mining. In *Proc. 15th ISWC 2016 — The Semantic Web, Kobe, Japan, October 17-21*, pages 498–514, 2016.
- [19] Jessica Rosati, Petar Ristoski, Tommaso Di Noia, Renato De Leone, and Heiko Paulheim. RDF graph embeddings for content-based recommender systems. In *Proc. 3rd W’shop New Trends in Content-Based Recommender Systems (@RecSys 2016), Boston, MA, September 16*, pages 23–30, 2016.
- [20] Sam Henry and Bridget T. McInnes. Literature based discovery: Models, methods, and trends. *Journal of Biomedical Informatics*, 74:20 – 32, 2017.