Extracting Opinion Targets from Environmental Web Coverage and Social Media Streams

Albert Weichselbraun
Univ. of Applied Sciences Chur
Chur, Switzerland
albert.weichselbraun@htwchur.ch

Arno Scharl
MODUL University Vienna
Vienna, Austria
scharl@modul.ac.at

Stefan Gindl
MODUL University Vienna
Vienna, Austria
stefan.gindl@modul.ac.at

Abstract

Policy makers and environmental organizations have a keen interest in awareness building and the evolution of stakeholder opinions on environmental issues. Mere polarity detection, as provided by many existing methods, does not suffice to understand the emergence of collective awareness. Methods for extracting affective knowledge should be able to pinpoint opinion targets within a thread. Opinion target extraction provides a more accurate and fine-grained identification of opinions expressed in online media. This paper compares two different approaches for identifying potential opinion targets and applies them to comments from the YouTube video sharing platform. The first approach is based on statistical keyword analysis in conjunction with sentiment classification on the sentence level. The second approach uses dependency parsing to pinpoint the target of an opinionated term. A case study based on YouTube postings applies the developed methods and measures their ability to handle noisy input data from social media streams.

Index Terms

Opinion mining; sentiment analysis; opinion target extraction; keyword analysis; climate change.

1. Introduction

Methods to pinpoint and track opinion targets such as politicians, environmental groups or companies can guide and improve communication and public outreach activities. Organizations do not only want to know whether content is positive or negative - they need to know what is driving the discourse, how issues are being framed, to whom or to what the expressed opinion is directed, and the level of agreement among opinion holders.

The examples presented in this paper focus on climate change communication, where conflicting positions are common. Figure 1 shows a screenshot of the Media Watch on Climate Change [1], a news and social media aggregator available at www.ecoresearch.net/climate. Within the DecarboNet research project (www.decarbonet.eu), this publicly available platform is currently extended into a collective awareness platform, in close collaboration with NOAA Climate Program Office (www.climate.gov) and the World Wide Fund for Nature (www.wwf.ch).

Organized by WWF since 2007, the Earth Hour (www.earthour.org) is a unique opportunity to apply the technologies of the Media Watch on Climate Change, as it unites hundreds of millions of citizens, businesses and governments around the world to support one of the largest environmental events in history. The example query shown in Figure 1 analyzes Anglo-American news media coverage on the Earth Hour between March and June 2015, comparing its coverage with articles on the upcoming COP21 UN Climate Change Conference in Paris (www.cop21paris.org), and IPCC - the Intergovernmental Panel on Climate Change (www.ipcc.ch).

The system uses multiple coordinated view technology to synchronize several visual representations of the search results - including a trend chart, word tree, geographic map, tag cloud and keyword graph [2], [3]. The list of associations in the lower left corner reflects topics and entities associated with the query term "earth hour" - including locations such as the Eiffel Tower and the Golden Gate Bridge, and organizations such as the World Wide Fund for Nature (WWF). The color coding (red = negative; grey = neutral; green = positive) of the frequency counts indicates average sentiment. Without a clear identification of opinion targets, however, it remains unclear whether a negative or positive bias refers to the listed entity, or to associated events and topics.

1530-1605/16 $31.00 © 2016 IEEE
DOI 10.1109/HICSS.2016.133

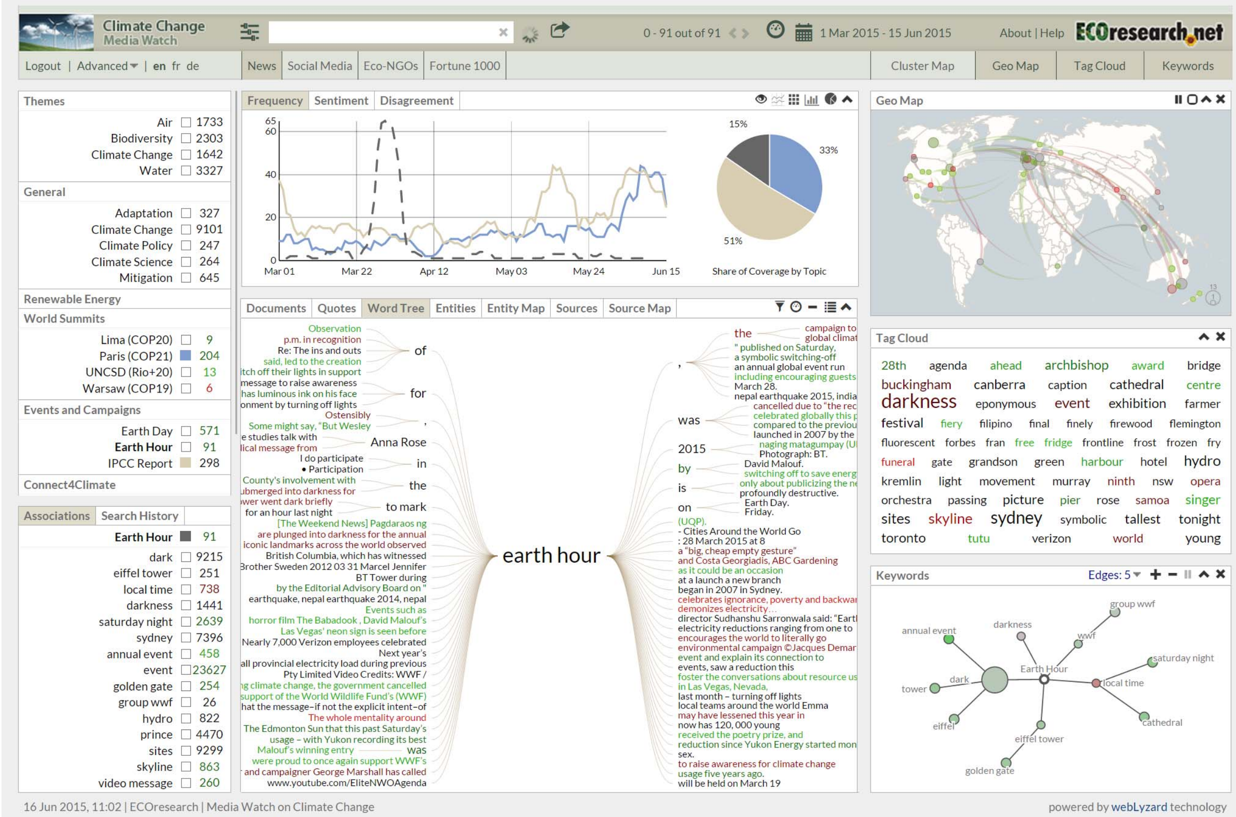

Figure 1. Screenshot of the Media Watch on Climate Change (www.ecoresearch.net/climate), analyzing the coverage about Earth Hour in Anglo-American news media outlets between March and June 2015

The approach presented in this paper addresses this shortcoming. It identifies opinions of social media users when formulating a message, and the target of this opinion. Without the ability to pinpoint opinion targets, sentiment analysis remains a statement about the polarity of textual content. While revealing the overall sentiment of a text is useful, identifying specific targets adds an important layer of knowledge. The author of a text might have a certain opinion about a topic in general, but assess subtopics or the role of specific stakeholders differently.

The ability to capture this distinction automatically has significant commercial potential as well, e.g. when analyzing consumer purchase decisions. A buyer of an electric car, for example, might feel good about the purchase, even when it turns out that the range in terms of kilometers per charge falls short of initial expectations (see dependency tree example of Section 3.2). Many organizations that provide products or services are interested in such fine-grained assessments, helping them to maximize customer value and identify features to improve in future revisions or new product lines. For communicators, the assessments can provide guidance in targeting a particular stakeholder group, and help align their messages to meet the informational requirements of citizens.

The remainder of this paper is organized as follows: Section 2 presents a general overview of statistical and syntactic approaches for opinion target extraction. Section 3 then describes two specific methods used in this paper. Section 4 elaborates on the corpora used for our experiments and presents a case study with the extracted terms from both methods based on YouTube postings. The paper closes with a summary and conclusion in Section 5.

2. Related Work

Opinion target extraction involves using statistical methods and linguistic rules. Statistical approaches determine keywords by comparing word distributions between a given text snippet and a larger text collection. Subsequently, a simple sentiment analysis algorithm could assign aggregated sentiment values to the extracted keywords. If such an approach relies on a straightforward mathematical model, it remains computationally feasible and scales well to large text corpora. Moreover, it is robust against typos and sloppy language usage; i.e., grammatical errors occurring in documents drafted by inexperienced authors or in colloquial social media postings. Linguistic approaches that consider the grammatical structure of a text tend to be more accurate, but computationally expensive - especially those that rely on advanced dependency parsing. They might also encounter difficulties when processing textual content containing errors or incomplete grammatical structures.

Traditionally, statistical approaches for opinion extraction have relied on comparing word frequencies between corpora by applying well-known measures such as the Log-Likelihood Ratio [4], Pointwise Mutual Information [5], Fisher's Exact Test, Pearson's Chi-Squared (χ^2) test, or the Dice Coefficient [6]. In recent years, these methods have been extended to use the distribution of information within the Web to determine the relatedness of concepts and relations.

Querying search engines such as Google, Yahoo! or Bing yields statistics on the distribution of keywords on the Web, that provide a rough approximation of the real distribution of that information [7]. Sanchez and Moreno [8] use these statistics together with the Pointwise Mutual Information (PMI) measure to acquire labeled relations and terms for ontology learning. They assure domain relatedness of new terms by computing their semantic association to domain keywords. Wong et al. [9] use the mutual information between constituents of terms to guide term simplifications for ontology learning. Weichselbraun et al. [10] query term statistics from Delicious and Yahoo! for refining domain ontologies. They use PMI to asses the quality of domain terms and to replace *weak concepts* with more adequate concepts retrieved from social media sources.

Aggarwal et al. [11] use lexical semantic analysis to derive online brand positions. They use the Google API to obtain search results for brands and associated brand positions, and then apply statistical methods such as PMI to determine how strongly Procter & Gamble detergent brands are associated with adjectives describing brand positions such as "Fresh", "Color", "Baby", etc.

Syntactic methods, by contrast, typically invoke stronger linguistic pre-processing. One way to tackle the problem of opinion target extraction is automatic semantic role labeling. Such an approach yields acceptable results, but requires the integration of other strategies such as anaphora resolution [12].

Syntactic methods are also often combined with machine learning. Jakob and Gurevych [13] extract opinion targets on multiple domains using conditional random fields. Their approach exploits several features, such as simple tokens, but also part-of-speech and dependency parsing. Nakagawa et al. [14] apply a similar approach using Conditional Random Fields and dependency parsing to Japanese and English sentences. Sayeed et al. [15] connect a-priori sentiment terms with their targets using syntactic relationships, derived from suffix-tree data structures. They use crowd-sourcing to overcome data sparseness, a problem common in this research area. Qiu et al [16] define grammar rules applied to the dependencies of terms to identify opinion targets. Their approach propagates the value from opinion-bearing words to their targets. After identifying targets, their algorithm connects them with additional terms within the sentence, given that target and new term have a dependency relation specified in a predefined set of relations. Thus, a freshly identified target can transfer its sentiment value onto other terms. The terms identified in this second step can either be new targets or unknown sentiment terms, ready for inclusion into a sentiment lexicon.

The approach presented in this paper uses a subset of the rules compiled in [16]. Rules for double-propagation, i.e. the bidirectional transfer of sentiment values onto targets and back to unknown sentiment terms, are excluded since it is outside the scope of this paper that focuses on target identification.

The aim of this work is a comparison of a statistical and syntactical method to provide insights into their respective strengths and weaknesses. The syntactical approach is limited to single nouns, while the statistical approach can also identify phrases but lacks the accuracy of the former. Future research will develop strategies for improving the versatility of these methods by combining them, i.e. labeling a sequence of terms as related using the statistical method and subsequently applying the syntactical method to exactly pinpoint targeted phrases.

3. Methodology

This paper compares statistical and linguistic approaches to opinion target extraction. This comparison gives insight on whether the higher accuracy of a syntactical approach legitimates the acceptance of increased time consumption as compared to a computationally less intensive statistical approach. A corpus compiled from the comments of YouTube videos is the basis for the evaluation. The comments have been extracted from more than 5900 environmental videos covering the topics *climate change, greenhouse effect,*

Table 1. Contingency table, potential bigram AB.

	A	$\neg A$
B	$n(AB)$	$n(\neg AB)$
$\neg B$	$n(A\neg B)$	$n(\neg A\neg B)$

greenhouse gas, global warming, global dimming, nuclear winter and *global cooling*. A comparison of the most frequent terms from the syntactic approach and the terms extracted with the statistical approach allow assessing the meaningfulness of the extracted terms.

3.1. Statistical Approach

The statistical approach for extracting associations consists of two major steps: (i) a collocation identification to extract idiomatic phrases, narrow collocations, and fixed phrases [17]; (ii) a significance test to compare the frequency distribution of all tokens (= unigrams and the identified collocations) and identify tokens that are significantly over-represented as potential associations. The following paragraphs outline these two steps in greater detail.

3.1.1. Collocation Identification. Pedersen et al. [18] distinguish the following groups of methods for identifying collocations: (i) methods based on Mutual Information such as the Log-Likelihood Ratio [4], true and pointwise Mutual Information [5], and Pearson-Stirling [19], (ii) Fisher's Exact Test, Pearson's Chi-Squared (χ^2), and the Dice Coefficient [6].

We use the contingency table in Table 1 to determine whether a word sequence AB such as "climate change" represents a valid collocation: The letters A and B refer to the corresponding words, the negated variable $\neg A$ indicates every possible word with the exception of A. Therefore, $n(AB)$ indicates the number of bigrams that contain the sequence AB, $n(A\neg B)$ the number of bigrams that start with A but do not contain B in the second position, and $n(\neg A\neg B)$ the count of bigrams that do not contain an A at the first position and no B on the second position.

We compute the Log-Likelihood Ratio (G) to determine how significantly the bigram counts $n(ij)$ with $i \in \{A, \neg A\}$ and $j \in \{B, \neg B\}$ deviate from the expected counts $m(ij)$ under the hypothesis of independence between the words A and B [20], [4].

$$G^2 = 2 \sum_{\substack{i\in\{A,\neg A\}\\ j\in\{B,\neg B\}}} n(ij) log \frac{n(ij)}{m(ij)} \qquad (1)$$

Adapting the equation for trigram collocations ABC yields formula 2.

$$G^2 = 2 \sum_{\substack{i\in\{A,\neg A\}\\ j\in\{B,\neg B\}\\ k\in\{C,\neg C\}}} n(ijk) log \frac{n(ijk)}{m(ijk)} \qquad (2)$$

The algorithm extracts bigrams and trigrams that exceed a threshold significance ($G^* = 2.0$) and includes them in the list of tokens considered for the subsequent extraction process of associations.

3.1.2. Extraction of Potential Opinion Targets. The keyword extraction component [21] identifies relevant features by comparing the token frequency distribution in a target corpus of YouTube comments on videos covering climate change and related environmental issues with a reference distribution that was obtained by assembling YouTube comments on general political issues. It is important to note that such keywords are not necessarily opinion targets, but rather candidate terms for such targets.

We identify tokens with a significant deviation between their expected $m(i)$ and observed $n(i)$ counts in the reference corpus by using the Chi-square test with Yates' correction for continuity

$$\chi^2_{Yates} = \sum_{i=1}^{N} \frac{(|n(i) - m(i)| - 0.5)^2}{m(i)} \qquad (3)$$

and consider overrepresented ($m(i) > n(i)$) terms as potential associations. The component then extracts these tokens since they tend to be specific to the discussions in the target corpus.

3.2. Syntactic Approach

The syntactic approach uses the NLP processing pipeline of the Media Watch on Climate Change to split the corpus into sentences, subsequently parses each sentence using the Stanford parser (http://nlp.stanford.edu/software/lex-parser.shtml) and applies the following rules [16] to the output of the parser:

$$O \rightarrow O - Dep \rightarrow T, \qquad (4)$$

s.t. $O \in \{O\}$, $O - Dep \in$ {MR}, POS(T) $\in$ {N, NN, NNP}, and

$$O \rightarrow O - Dep \rightarrow H \leftarrow T - Dep \leftarrow T, \qquad (5)$$

s.t. $O \in \{O\}$, $O/T - Dep \in$ {MR}, $POS(T) \in$ {N, NN, NNP}; $MR =$ {advmod, amod, rcmod, nsubj, s, obj, obj2, desc, nn}.

O represents the set of opinionated terms, *T* is their target and must be a noun, i.e. $POS(\text{T}) \in \{\text{N, NN, NNP}\}$. *MR* is the set of modifying relations, e.g. "amod" means "adjective modifier". The first rule is a single propagation from sentiment terms to their noun targets. For instance, in the sentence "The phone has a good screen" the term "screen" (*T*) receives positive sentiment from "good" (*O*). The second rule transfers the sentiment value a target has received from an a-priori sentiment term to another target within the same sentence. In the same example sentence as before it connects the "good screen" with "phone" and transfers the positive sentiment onto the latter. The approach is limited to these two rules, because the remaining rules support double-propagation, i.e. the identification of previously unkown sentiment terms, which was outside the scope of this work. The application of the rule set, initially created for Minipar in [16], required renaming some of the dependency relations to fit them to the output the Stanford parser produces. Parsing and propagating the charges from sentiment terms to their targets is a time-consuming task, but has a clear advantage compared to the statistical approach. The latter can only assign the overall sentiment value of the total text snippet, i.e. the sentence, paragraph or even document, to the keywords extracted from it. The former, on the other hand, is able to exactly pinpoint the target of a sentiment term via its syntactic relation. Syntactic extraction allows several targets within one and the same sentence, each having a different sentiment value. Figure 2 shows the dependency tree for a sentence discussing the *Tesla Roadster* with the targets "driver" and "ranges". Both targets are identified via their dependency relations to the sentiment terms "longest" and "nervous" and have corresponding, differing sentiment values.

The syntactic approach suffers from the frequent use of noisy language in YouTube comments - i.e., grammar and spelling mistakes, abbreviations and acronyms, Web lingo, etc. Parsing mistakes introduced by a defective sentence structure strongly affect the precision of finding targets. Thus, the cleanliness of the data set strongly influences the outcome. Nevertheless, our results in Section 4 demonstrate that, given a reasonable corpus size, even the limited linguistic quality of YouTube comments suffices for extracting useful and intuitive opinion targets.

4. Case Study on YouTube Coverage

YouTube is a good source for retrieving comments on political and educational videos on climate change and related environmental issues. Due to the emotional nature of the social media discourse, comments from YouTube Videos often contain positive or negative statements in conjunction with corresponding opinion targets. Querying the YouTube Data API for seven climate change terms (climate change, global dimming, nuclear winter, global cooling, greenhouse effect, greenhouse gas, global warming) taken from the climate change ontology constructed by Liu et al. [22] yielded a total of 5990 videos (4325 of which contained comments). Cleanup and pre-processing steps obtained video comments, used a rule-based algorithm for sentence splitting, identified and removed duplicates based on their MD5 hash sums, and provided part-of-speech tagging for the extracted sentences. This process yielded a total of 505,226 video comments.

A YouTube search for the search terms "Barack Obama" and "Mitt Romney" yielded a reference corpus of 1808 relevant videos and after the removal of duplicate sentences a total of 349,739 comments. This corpus provided the reference distribution for computing the expected counts ($m(i)$) in Equation 3 that is necessary for the statistical approach.

Table 2 shows the five most frequent terms extracted by the statistical and syntactic approach from each of the seven subcorpora, for both positive and negative sentiment. The statistical approach uses a reference distribution that has been computed based on the YouTube coverage of the US 2012 Elections for eliminating generic terms. This makes it less likely to extract generic and, therefore, irrelevant artifacts introduced by the social media platform. The syntactic approach, in contrast, does not yet eliminate such terms and, therefore, contains artifacts such as "thing" and "song". It is also currently limited to unigrams – leading to results such as *al*, which represents the first name of *Al Gore*.

To assess opinion target quality, we compared the most frequent opinion targets computed by (i) applying the *statistical keyword extraction algorithm* to single comments, and aggregating the results; and (ii) the *syntactic approach*. We combine Web statistics retrieved from Google search results and pointwise mutual information (PMI) for assessing the strength of the semantic relation between the YouTube query term and the retrieved opinion target (Equation 6).

$$PMI(\text{target, seed concept}) = log\frac{n(\text{target, seed concept})}{n(\text{target}) \cdot n(\text{seed concept})} \quad (6)$$

An approximation of the distribution of documents containing both, the opinion target and the seed term (or one of them) was obtained from the number of search results for corresponding Google queries.

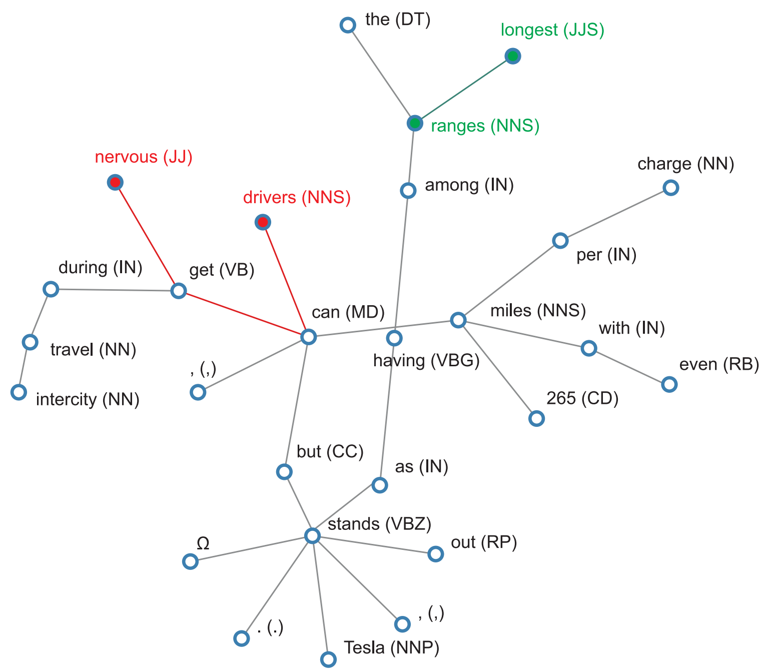

Figure 2. Multi-target sentiment propagation within a single sentence.

The notion behind this Web metric is that relevant opinion targets are assumed to show a strong association to the seed concept, in user-generated content from social media *and* in the Web corpus. Figure 3 visualizes how strongly the top-ranked opinion topics obtained with keyword analysis and opinion target extraction relate to the YouTube query terms according to this measure.

The graphs reveal that "mod" (upper right spiderweb graph) – which is also an abbreviation for a unit of measurement – has only a relatively small association with global cooling and global dimming according to the PMI Web metric. "Job" (lower right graph), in contrast, is a highly relevant opinion target for Web documents indexed by Google that discuss climate change and global warming. The introduced Web metric, therefore, provides means to refine the final selection of opinion targets by removing uncommon targets, i.e. opinion targets with a low relation strength in the general Web corpus.

5. Conclusion

Tracking the emergence of collective awareness among environmental stakeholders imposes challenges on existing opinion extraction techniques. This paper addresses these challenges by going beyond simple polarity detection and presenting statistical and syntactic extraction as two alternative approaches to identifying opinion targets within a corpus. The strength of the statistical approach lies in its robustness and time efficiency. It is able to determine meaningful keywords even if the corpus is not written in accurate language. This is a clear advantage for processing Web documents not written by professional authors. A disadvantage is the lack of granularity. Multiple targets with differing sentiment values occurring in one and the same sentence cannot be distinguished properly. The algorithm calculates an overall sentiment value for the whole sentence and applies it to both targets. In the worst case scenario, both targets get a neutral value assigned, when sentiment terms in the sentence cancel each other out. The syntactic approach, on the other hand, is more time-consuming and faces difficulties when processing noisy text. However, it is able to discern multiple targets in a single sentence and can assign different sentiment values to them.

The examples used in this paper stem from the Media Watch on Climate Change (MWCC), a news and social media aggregator that is publicly available at www.ecoresearch.net/climate. Several research initiatives (see acknowledgment below) are currently extending the underlying Web intelligence platform (www.weblyzard.com). The Radar Media Criticism project develops linguistic methods for opinion target extraction and sentiment analysis. The ASAP-FP7.eu Project increases the scalability of methods to analyze and visualize big data archives [23], which are applied

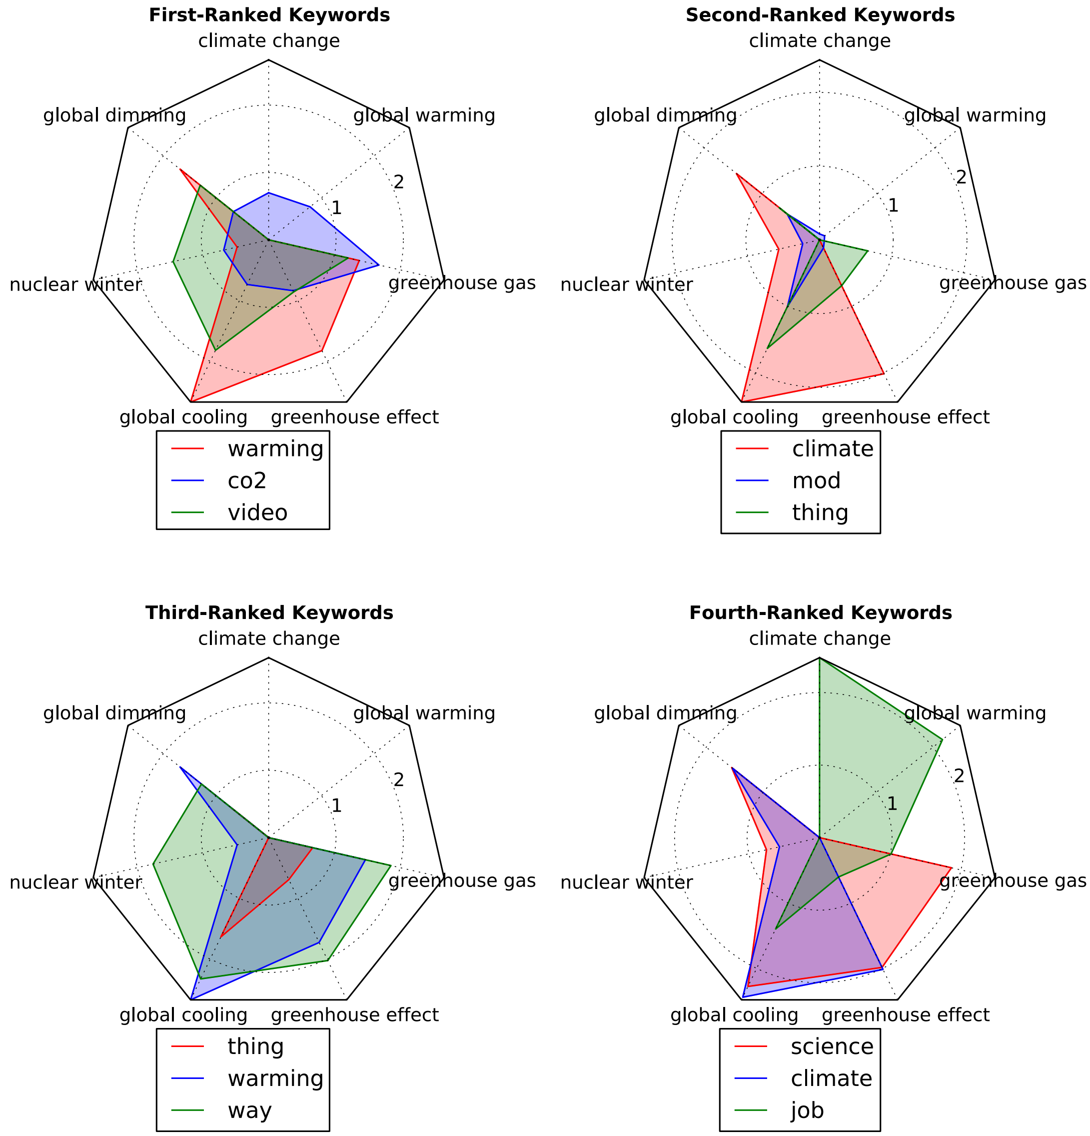

Figure 3. PMI between the extracted opinion targets and the seven YouTube query terms. The green (red) areas represent positive (negative) opinion targets suggested by the syntactic approach, blue graphs refer to targets obtained by the statistical approach.

by the uComp.eu project to identify patterns across factual and affective knowledge [24]. The DecarboNet.eu Project builds on such patterns to investigate information diffusion processes and shed light on the emergence of collective awareness.

Future work will integrate both approaches for improved accuracy, multi-term phrase detection (syntactical approach), and the correct processing of spelling mistakes and other orthographic variations (statistical approach). Sentiment analysis tools will benefit from such an integration. The statistical approach serves for the determination of topics in larger areas of text. It summarizes sentiment expressed towards these topics. To allow for an in-depth analysis of the text, a syntactic approach can then be invoked to obtain more detailed target statistics - including the sentiment expressed towards them.

Table 2. **The eight most frequently identified targets for the statistical (keywords) and syntactic approach (postive & negative terms) in each YouTube corpus.**

Corpus	Keywords	Positive terms	Negative terms
Climate change	climate, co2, warming, carbon dioxide, temperature	thing, way, job, science, work	warming, climate, thing, science, change
Global dimming	chemtrails, spraying, trails, clouds, chem	thing, way, song, job, al	warming, thing, climate, al, science
Nuclear winter	mod, al, fallout, companion, jerry	al, song, thing, mod, way	al, thing, metal, game, warming
Global cooling	co2, cooling, warming, ice age, temperature	thing, way, job, work, song	warming, climate, thing, science, al
Greenhouse effect	my hair, co2, hair, eminem, your hair	thing, way, job, science, work	warming, climate, thing, science, change
Greenhouse gas	co2, meat, vegan, methane, dioxide	thing, way, idea, energy, job	warming, thing, climate, gas, science
Global warming	co2, warming, agw, temperature, global warming	thing, way, song, al, job	warming, thing, climate, al, science

Acknowledgment

DecarboNet.eu and ASAP-FP7.eu have received funding by the European Union's 7th Framework Program for research, technology development and demonstration under (Grant Agreement No. 610829 and 619706); the Radar Media Criticism Switzerland (RMCS) project by the Swiss National Science Foundation (No. 150327); uComp by EPSRC EP/K017896/1, FWF 1097-N23 and ANR-12-CHRI-0003-03 in the framework of the CHIST-ERA ERA-NET program line.

References

[1] A. Scharl, A. Hubmann-Haidvogel, A. Weichselbraun, H.-P. Lang, and M. Sabou, "Media Watch on Climate Change – Visual Analytics for Aggregating and Managing Environmental Knowledge from Online Sources," in *Proceedings of the 46th Hawaii International Conference on System Sciences (HICSS-46)*. Maui, Hawaii: IEEE Computer Society Press, January 2013, pp. 955–964.

[2] A. Hubmann-Haidvogel, A. Scharl, and A. Weichselbraun, "Multiple Coordinated Views for Searching and Navigating Web Content Repositories," *Information Sciences*, vol. 179, no. 12, pp. 1813–1821, 2009.

[3] D. Fischl and A. Scharl, "Metadata Enriched Visualization of Keywords in Context," in *6th ACM SIGCHI Symposium on Engineering Interactive Computing Systems*. Rome, Italy: ACM Press, 2014, forthcoming.

[4] T. Dunning, "Accurate Methods for the Statistics of Surprise and Coincidence," *Computational Linguistics*, vol. 19, no. 1, pp. 61–74, 1993.

[5] K. W. Church and P. Hanks, "Word Association Norms, Mutual Information, and Lexicography," *Computational Linguistics*, vol. 16, no. 1, pp. 22–29, 1990.

[6] F. Smadja, "Retrieving Collocations from Text: Xtract," *Computational Linguistics*, vol. 19, no. 1, pp. 143–177, 1993.

[7] R. Cilibrasi and P. M. B. Vitanyi, "Automatic Meaning Discovery Using Google," December 2004.

[8] D. Sánchez and A. Moreno, "Learning Non-taxonomic Relationships from Web Documents for Domain Ontology Construction," *Data & Knowledge Engineering*, vol. 64, no. 3, pp. 600–623, 2008.

[9] W. Wong, W. Liu, and M. Bennamoun, "Acquiring Semantic Relations Using the Web for Constructing Lightweight Ontologies," in *PAKDD '09: Proceedings of the 13th Pacific-Asia Conference on Advances in Knowledge Discovery and Data Mining*. Berlin, Heidelberg: Springer-Verlag, 2009, pp. 266–277.

[10] A. Weichselbraun, G. Wohlgenannt, and A. Scharl, "Applying Optimal Stopping Theory to Improve the Performance of Ontology Refinement Methods," in *Proceedings of the 44th Hawaii International Conference on System Sciences (HICSS-44)*. Maui, Hawaii: IEEE Computer Society Press, January 2011, pp. 1–10.

[11] P. Aggarwal, R. Vaidyanathan, and A. Venkatesh, "Using Lexical Semantic Analysis to Derive Online Brand Positions: An Application to Retail Marketing Research," *Journal of Retailing*, vol. 85, no. 2, pp. 145 – 158, 2009.

[12] J. Ruppenhofer, S. Somasundaran, and J. Wiebe, "Finding the Sources and Targets of Subjective Expressions," in *Proceedings of the Sixth International Language Resources and Evaluation (LREC)*, 2008.

[13] N. Jakob and I. Gurevych, "Extracting Opinion Targets in a Single- and Cross-Domain Setting with Conditional Random Fields," in *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP 2010)*, 2010, pp. 1035–1045.

[14] T. Nakagawa, K. Inui, and S. Kurohashi, "Dependency Tree-based Sentiment Classification Using CRFs with Hidden Variables," in *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics (HLT)*, 2010, pp. 786–794.

[15] A. B. Sayeed, J. Boyd-Graber, B. Rusk, and A. Weinberg, "Grammatical Structures for Word-Level Sentiment Detection," in *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL HLT)*, 2012, pp. 667–676.

[16] G. Qiu, B. Liu, J. Bu, and C. Chen, "Opinion Word Expansion and Target Extraction through Double Propagation," *Computational Linguistics*, vol. 37, no. 1, pp. 9–27, 2011.

[17] J. Wermter and U. Hahn, "Collocation Extraction Based on Modifiability Statistics," in *Proceedings of the 20th international conference on Computational Linguistics (COLING 2004)*, Stroudsburg, PA, USA, 2004.

[18] T. Pedersen, S. Banerjee, B. T. McInnes, S. Kohli, M. Joshi, and Y. Liu, "The N-gram Statistics Package (NSP): A Flexible Tool for Identifying N-grams, Collocations, and Word Associations," in *Proceedings of the Workshop on Multiword Expressions: from Parsing and Generation to the Real World (MWE)*. Stroudsburg, PA, USA: Association for Computational Linguistics, 2011, pp. 131–133.

[19] K. W. Church, "Empirical Estimates of Adaptation: The Chance of Two Noriegas is Closer to p/2 than p^2," in *Proceedings of the 18th conference on Computational linguistics (COLING 2000)*, Stroudsburg, PA, USA, 2000, pp. 180–186.

[20] E. Agirre and P. G. Edmonds, *Word Sense Disambiguation: Algorithms and Applications*, ser. Text, Speech and Language Technology. Springer, 2006, no. 33.

[21] A. Scharl and A. Weichselbraun, "An Automated Approach to Investigating the Online Media Coverage of US Presidential Elections," *Journal of Information Technology & Politics*, vol. 5, no. 1, pp. 121–132, 2008.

[22] W. Liu, A. Weichselbraun, A. Scharl, and E. Chang, "Semi-automatic Ontology Extension Using Spreading Activation," *Journal of Universal Knowledge Management*, vol. 0, no. 1, pp. 50–58, 2005.

[23] A. Scharl, D. Herring, W. Rafelsberger, A. Hubmann-Haidvogel, R. Kamolov, D. Fischl, M. Föls, and A. . Weichselbraun, "Semantic Systems and Visual Tools to Support Environmental Communication," *IEEE Systems Journal*, 2015, forthcoming (Accepted 31 July 2015).

[24] A. Weichselbraun, S. Gindl, and A. Scharl, "Enriching Semantic Knowledge Bases for Opinion Mining in Big Data Applications," *Knowledge-Based Systems*, 2014.