

Submission Number: 657

NewsViz: Extraction and Visualisation of Emotions from News Articles

Eva Hanser and Paul Mc Kevitt
School of Computing & Intelligent Systems
Faculty of Computing & Engineering
University of Ulster, Magee
Derry/Londonderry, BT48 7JL
Northern Ireland

E-mail: hanser-e@email.ulster.ac.uk, p.mckevitt@ulster.ac.uk

Abstract

News is commonly intended to be delivered in an objective, unbiased manner and therefore presented plainly and formally, even though its content often affects readers emotionally. The NewsViz system aims to enhance the news reading experience by integrating 30 seconds long Flash-animations into news article web pages depicting their content and emotional aspects. NewsViz interprets football news texts automatically and creates abstract 2D visualisations. The user interface enables animators to further refine animations. Here we focus on the emotion extraction component of NewsViz which facilitates subtle background visualisation. The emphasis of NewsViz lies on expression, impacting on the reader's understanding of the article and making it more memorable. NewsViz detects moods from news reports. The original text is part-of-speech tagged and adjectives and/or nouns, the word types conveying most emotional meaning, are filtered out and labelled with an emotion and intensity value. Subsequently reoccurring emotions are joined into longer lasting moods and matched with appropriate animation presets. Different linguistic analysis methods were tested on NewsViz: word-by-word, sentence based and incremental minimum threshold summarisation, to find a minimum number of occurrences of an emotion in forming a valid mood. NewsViz proved to be viable for the fixed domain of football news, grasping the overall moods and some more detailed emotions precisely. NewsViz introduces a novel approach to a universally applicable emotion scheme which offers an efficient technique to cater for the production of a large number of daily updated news stories. NewsViz fills the gap of lack of information for background or environment depiction encountered in similar applications. Further development may refine the detection of emotion shifts through summarisation with the full implementation of football and common linguistic knowledge. Future work will reveal whether NewsViz is feasible when extended to different domains.

Keywords: Natural Language Processing, Emotion Sensing, Automatic Visualisation, Text-to-Animation, NewsViz

1. Introduction

News reports are regarded as objective facts represented in a neutral and formal format: typically a static headline, a summarising paragraph with one image and the body text

eventually with one to three more images. Even though reporters find the content of news stories worth mentioning for emotional reasons, story brevity, scarce background information and poor combination of visual and verbal information hinders learning by viewers. In order to reach the audience emotionally, educate and entertain, emphasis on visual elements is important as they tend to be more memorable than verbal ones. The software prototype, NewsViz, creates animations from news articles. Abstract design elements show emotions conveyed in the stories. The main objective of NewsViz remains information provision, thus emotion extraction is universally applicable and without opinion bias. NewsViz is an efficient software tool for designers to be able to build daily updated animations. Input for NewsViz is natural language text. Multimodal systems automatically mapping text to visuals face challenges in interpreting human language which is variable, ambiguous, imprecise and relies on common knowledge between the communicators. Enabling a machine to understand a natural language text involves feeding the machine with grammatical structures, semantic relations and visual descriptions to be able to match suitable graphics.

2. Background and related research

Text-to-visual mapping relates to the area of natural language processing (NLP) and multimodal storytelling which attempt to enable computers to interpret and generate natural human language and mental images.

Text-to-visual mapping starts with linguistic analysis of the text. Despite variability, ambiguity and imprecision *syntactic analysis* tools achieve mostly reliable results. Part-of-speech tagger software tools identify grammatical types of words. For example, Qtag (Mason 2003) attaches a tag to each word labelling it as noun, verb, adjective or other.

Semantic interpretation and actual understanding of the meaning of a text is more difficult, because it depends largely on common sense knowledge. Common sense knowledge and mental images need to be structured, related through logical rules and entered into databases before computational text interpretation is possible. A commonly used tool for determining semantic relations between words is WordNet (Miller 1995), an extended dictionary specifying word relations such as similarity, part-of relations, hierarchy or manner. Sensing emotions from multimodal input has mainly been investigated with the objective of developing human-like agents. The football commentary system, Byrne (Binsted and Luke 1999), includes a commentator with emotions influenced by his personality and intentions. SOBA (Buitelaar et al. 2006) extracts information from soccer match reports, annotates relevant expressions (e.g. players, teams, goals.) and generates knowledge base entities. The collected football knowledge can set preconditions and context to consequently evaluate current events and assign appropriate emotions.

Automated story visualisation systems deliver initial results for object and action depiction, as in WordsEye (Coyne and Sproat 2001), creating static 3D images from written descriptions.

Additionally, automated camera and character animation, interaction and speech synthesis is realised in CONFUCIUS (Ma 2006). The Unseen Video (Scheibel and Weinrother 2005), is a good example of abstract mood visualisation. Local weather data is automatically retrieved from news websites and influences the look and feel of the Flash animation through shapes, colours and images. The Story Picturing Engine (Joshi et al. 2004) visualises texts selecting and matching pictures and their annotations from image databases.

The work discussed in section 2 demonstrates that sufficient subsets of the English language can be mapped to computer understandable language for the visualisation of stories.

3 The NewsViz system

NewsViz takes online news articles as input and outputs animations reflecting the content of these news stories. NewsViz consists of three main components: the linguistic analysis, the animation composer and an interface for editing text and animations (Fig. 1). The linguistic component constructs three elements of the animation in different processes. The emotion extraction tool creates atmospheric background visuals, the action visualiser depicts people, objects and their actions and the audio creator selects music and sound effects. The composer synchronises the different outputs. Here, we focus on the emotion extraction component (Fig. 2) developed in Flash MX and Photoshop. Emotional aspects within the news story are identified and linked to appropriate presets of background animations.

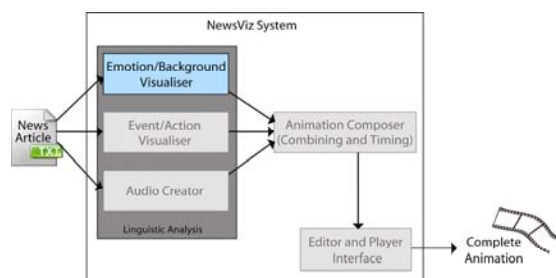


Fig. 1: NewsViz system architecture




Fig. 2: Emotion extraction component

3.1 Emotion extraction

The first step in processing the text is to define the grammatical type of all words. The part-of-speech tagger, Qtag (Mason 2003), attaches tags to nouns, verbs, adjectives and other parts of speech. The tagged text is sent on to the adjective and noun detector. Only these two types of words are selected for further processing because emotional meaning is mainly conveyed by them. Next, the emotion word selector checks the adjectives and nouns in the emotion dictionary and attaches emotion tags indicating their kind of emotion and intensity. The dictionary holds manually created emotion-indices and default intensity values of all affective words. Four emotions have been found relevant in relation to football matches –

happiness, sadness, tension and boredom. Words with a neutral emotion index do not describe football relevant emotions. To achieve a coherent course of emotion and animation, neutral phases are replaced by the previous mood with decreasing intensity. The list of emotion tagged words is handed to the emotion summariser. During the summarisation process subsequent emotions of the same type are combined to one longer-lasting mood. Each mood is labelled with its type, average intensity and display duration. With the 'word by word' summarisation method mood boundaries appear as soon as the emotion type of the next word differs. In order to reduce error and excessive mood swings, the minimum threshold method sets a minimum number of words required to represent a mood. Alternatively, the sentence based method assumes that one sentence conveys one idea and consequently one emotion. Hence, it calculates an average emotion for each sentence, before combining identical emotions. A chronological list of mood chunks is created.

3.2 Animation construction

The animation selection component loads the individual animation elements from the graphics database and combines them in a 30 seconds long animation. The graphics database contains prefabricated graphics sorted by an emotion index which are combined and adjusted according to mood intensities. Based on the weighted mood list, the emotion sequence order, the type of graphic element, its display duration, and the background colour are determined. The intensity value specifies the element size and the number of objects loaded. An emotion change causes the current animation elements to fade out and to load different elements. Animation examples are shown in Fig. 3.




Fig. 3: Animations for Boredom (green), Happiness (yellow), Sadness (blue) and Tension (red)

3.3 User interface

NewsViz users load or type news stories into the text editor. The options menu offers different emotion extraction and mood summarisation methods. By pressing the 'run' button the visualisation can be watched in the preview window. The text processing runs invisibly in the background. If the user is satisfied he can save the animation. If the user prefers to alter the animation manually, he has the option to edit the original text or the animation elements frame by frame. Fig. 4 shows the user interface with animation player. The final animations are integrated at the top of the news article's internet page (Fig. 5).




Fig. 4: NewsViz User Interface



Fig. 5: Animation integrated into website

3.4 Evaluation and testing

NewsViz was tested on four football match reports from the BBC and FIFA websites. NewsViz emotion output was evaluated against two forms of human interpretation of these news articles, a brief outline of the prevailing moods naming three to five emotions per article and a more fine grained interpretation assigning one (or two) emotions to each sentence. Three types of emotion extraction error were distinguished, falsely detected, missing and very close, but not exact, emotions. On average, the best results are achieved when extracting both, adjective and nouns, with the word by word method, but this output is too fine grained for visualisation. Summarisation methods need improvement. Overall the outcome of NewsViz is satisfactory and demonstrates that news texts are suitable for emotion extraction. The different sensations of the described football matches were distinguishable.

4 Conclusion and future work

NewsViz extracts emotion-bearing words from online football news reports based on an extended dictionary with emotion-indices assigned to each entry. The extracted emotions are processed and illustrated in abstract background animations. Results from initial testing demonstrate that this automated process has satisfactory performance. Technologically, NewsViz is viable for the fixed domain of football reports and offers a sound basis for more affective text-to-visual mapping. Future work will aim to improve the semantic processing of emotions. This involves the integration of rules for common and linguistic knowledge. Linguistic knowledge identifies emotions in context applying language rules to emotion interpretation, i.e. it solves negation by inverting emotions. With a dependency parser expressions of related words can be found and their average emotion determined. Domain-specific knowledge (e.g. football) provides background information including match statistics, players' and teams' names, team colours and league tables. It also accommodates game rules or match situations with their emotional consequences. The mood list is refined through moods discovered with common knowledge and football facts which set pre-conditions and context representing long-term moods influencing current event-based emotions. The emotion database could be extended through the WordNet-Affect dictionary (Strapparava and Valitutti

2004). NewsViz enriches standard news websites with attractive and informative animations and can track emotional aspects of people's views on world events. NewsViz brings news reported on the internet closer to readers, making it more easily understood and memorised which is much appreciated by online users overloaded with information. NewsViz assists animation designers in the production of daily updated visualisations creating initial scenes.

References

Binsted, K. and Luke, S. "Character Design for Soccer Commentary". *Lecture Notes in Computer Science. RoboCup-98: Robot Soccer World Cup II*. Springer-Verlag, London. 1604. (1999): 22-33.

Buitelaar, P., Eigner, T., Gulrajani, G., Schutz, A., Siegel, M., Weber, N., Cimiano, P., Ladwig, G., Mantel, M., Zhu, H. "Generating and Visualizing a Soccer Knowledge Base". *Proceedings of the EAACL06 Demo Session*. 4/2006. (2006): 123-126.

Coyne, B. and Sproat, R. "WordsEye: an automatic text-to-scene conversion system". *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*. ACM Press, Los Angeles. (2001): 487-496.

Joshi, D. and Wang, J. Z. and Li, J. "The Story Picturing Engine: Finding Elite Images to Illustrate a Story Using Mutual Reinforcement". *Proceedings of the 6th ACM SIGMM International Workshop on Multimedia Information Retrieval*. ACM Press, New York. (2004): 119-126.

Ma, M. "Automatic Conversion of Natural Language to 3D Animation". *Ph.D. Thesis*, School of Computing and Intelligent Systems, University of Ulster. (2006).

Mason, O. "[Qtag](#)". (2003). May 2009.

Miller, G. A. "WordNet: a lexical database for English". *Communications of the ACM*, 38(11). (1995): 39-41.

Scheibel, D. and Weinrother, F. "[The Unseen Video](#)". (2005). May 2009.

Strapparava, C. and Valitutti, A. "WordNet-Affect: an affective extension of WordNet". *Proceedings of the 4th International Conference on Language Resources and Evaluation (LREC 2004)*. 4. (2004): 1083-1086.