Affective Common Sense Knowledge Acquisition for Sentiment Analysis

Erik Cambria[1], **Yunqing Xia**[2], **Amir Hussain**[3]

[1]Temasek Laboratories, National University of Singapore, Singapore 117411, Singapore
E-mail: cambria@nus.edu.sg
[2]Department of Computer Science and Technology, Tsinghua University, Beijing 100084, P. R. China
E-mail: yqxia@tsinghua.edu.cn
[3]Department of Computing Science and Mathematics, University of Stirling, Stirling FK9 4LA, UK
E-mail: ahu@cs.stir.ac.uk

Abstract

Thanks to the advent of Web 2.0, the potential for opinion sharing today is unmatched in history. Making meaning out of the huge amount of unstructured information available online, however, is extremely difficult as web-contents, despite being perfectly suitable for human consumption, still remain hardly accessible to machines. To bridge the cognitive and affective gap between word-level natural language data and the concept-level sentiments conveyed by them, affective common sense knowledge is needed. In sentic computing, the general common sense knowledge contained in ConceptNet is usually exploited to spread affective information from selected affect seeds to other concepts. In this work, besides exploiting the emotional content of the Open Mind corpus, we also collect new affective common sense knowledge through label sequential rules, crowd sourcing, and games-with-a-purpose techniques. In particular, we develop Open Mind Common Sentics, an emotion-sensitive IUI that serves both as a platform for affective common sense acquisition and as a publicly available NLP tool for extracting the cognitive and affective information associated with short texts.

Keywords: knowledge acquisition, crowd sourcing, games with a purpose, natural language processing, sentic computing

1. Introduction

The Social Web has changed the ways people communicate, collaborate, and express their opinions. The potential for opinion sharing today is unmatched in history. Never before have so many knowledgeable people been connected by such a time and cost efficient and effective network.

The distillation of useful knowledge from the huge amount of unstructured information available online, however, is an extremely difficult task as today web-contents are perfectly suitable for human consumption but they remain hardly accessible to machines. The Web, in fact, mostly owes its success to the development of search engines like Google and Yahoo, which represent the starting point for information retrieval. Such engines, which base their searches on keyword-based algorithms relying on the textual representation of the web-page, are very good in retrieving texts, splitting them into parts, checking the spelling, counting their words. But when it comes to interpreting sentences and extracting useful information for users, their capabilities result still very limited.

Current attempts to perform automatic understanding of text, e.g., textual entailment and machine reading, still suffer from numerous problems including inconsistencies, synonymy, polysemy, and entity duplication, as they focus on a mere syntactical analysis of text. To bridge the cognitive and affective gap between word-level natural language data and the concept-level opinions and sentiments conveyed by them, we need intelligent systems able to learn new affective common sense knowledge and to perform reasoning on it, in order to semantically and affectively analyse natural language text. In human cognition, thinking and feeling are mutually present: emotions are often the product of our thoughts as well as our reflections are often the product of our affective states. Emotions, in fact, are intrinsically part of our mental activity and play a key role in decision-making processes: they are special states, shaped by natural selection, to adjust various aspects of our organism to make it better face particular situations, e.g., anger evolved for reaction, fear evolved for protection and affection evolved for reproduction (Minsky, 2006). For these reasons, we cannot prescind from emotions in the development of intelligent systems: if we want computers to be really intelligent, not just have the veneer of intelligence, we need to give them the ability to recognize, understand, and express emotions.

To this end, in this work we exploit games-with-a-purpose (GWAP) techniques, together with label sequential rules (LSR), and crowd-sourcing methods, to collect new affective common sense knowledge that we need, within sentic computing, for tasks such as social media marketing (Cambria et al., 2011), patient opinion mining (Cambria et al., 2012a), and affective resource design (Cambria et al., 2012b). In particular, the paper is organised as follows: Section 2 presents an overview on the use of games for knowledge acquisition, Section 3 explains the motivations for collecting affective common sense knowledge, Section 4, 5 and 6 illustrates the adopted LSR, crowd-sourcing, and GWAP techniques respectively, and Section 7 comprises concluding remarks and future directions.

2. Games with a Purpose

The Casual Games Association[1] reports more than 200 million casual gamers worldwide this year. People play games for different reasons, e.g., to relax, to be entertained, for the need of competition and to be thrilled (PopCap, 2010).

[1]`http://casualgamesassociation.org`

Additionally, they want to be challenged, both on a mental and on a skill-based level. Such army of gamers could be exploited for performing tasks that are relatively easy to complete by humans, but computationally rather infeasible to solve (von Ahn et al., 2003). The idea is to integrate such tasks as goal of games (von Ahn, 2006) by producing a win-win situation where people have fun playing games while actually doing something useful. The nature of these games, in fact, focuses on exploiting player inputs to both create meaningful data and provide a funnier game experience (Thaler et al., 2011).
Such a human-based computational power can be exploited for tasks such as video annotation, e.g., OntoTube (Siorpaes and Hepp, 2008), PopVideo (Fig. 1), Yahoo's Videotaggame (van Zwol et al., 2008), and Waisd (Addis et al., 2010), in which two players have to timely agree on a set of tags about the same streaming YouTube[2] video.
Similarly, in ESP game (von Ahn and Dabbish, 2004) and Google Image Labeler (before being discontinued last September) players have to consensually guess content objects or properties of random images by simultaneously typing what they see. Other games for image annotation include Matchin (Hacker and von Ahn, 2009), which focuses on image perceived quality by asking players to pairwise choose the picture they like better, Phetch (von Ahn et al., 2006a), a game that collects explanatory descriptions of images in order to improve accessibility of the Web for the visually impaired by letting a player describe an image and others retrieve it using an image search engine, Peekaboom (von Ahn et al., 2006c), which focuses on locating objects within images by letting a player reveal specific parts of an image in order for the other to guess the correct object name, Squigl, in which players have to spot objects in images previously annotated within ESP Game, and Picture This, which asks players to choose, among a set of images, the one that best suits the given query.
Among games for image annotation, there are also games for streamlining the robustness evaluation of CAPTCHAs, namely: Magic Bullet (Yan and Yu, 2009), a team game in which players need to agree on the meaning of CAPTCHAs, and TagCaptcha (Morrison et al., 2009), where players are asked to quickly describe CAPTCHA images with one word each. GWAPs are also exploited to automatically tag music with deeper semantic labels. HerdIt (Barrington et al., 2009), for example, asks players have to accomplish different tasks related to the song they are listening to, while in Tagatune (Law et al., 2007) two players have to listen to an audio file and describe to the other what they are hearing, in order for him/her to decide whether the game has played the same soundtrack to both or not.
Several games have also been designed for text annotation. Verbosity (von Ahn et al., 2006b), for example, is a real time quiz game for collecting common sense facts. In the game, two players take different roles at different times: a narrator, who has to describe a word using templates, and a guesser, who has to guess such word in the shortest time possible. Sentiment Quiz, instead, gathers information about the polarity associated to words.

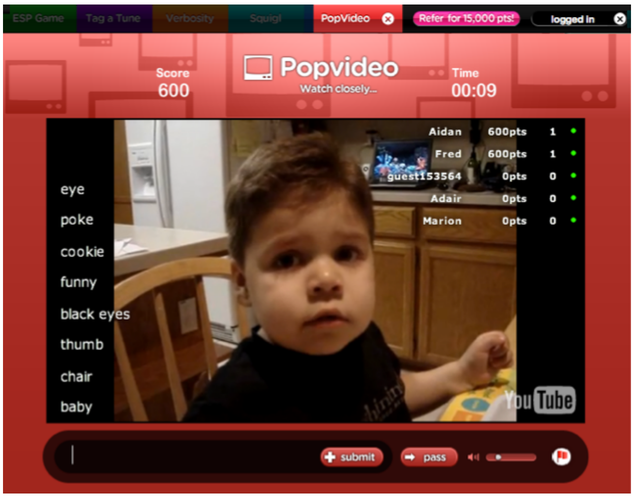

Figure 1: A screenshot of PopVideo. Two or more players are shown the same video and have to timely describe the objects that appear in the video in order earn points.

It asks its players to evaluate random words on a five grade scale, from very negative over neutral to very positive. Another approach to collecting common sense knowledge is the FACTory Game (Lenat and Guha, 1989), published by Cycorp. FACTory randomly chooses facts from Cyc and presents them to players, in order for them to guess whether a statement is true, false, or does not make sense. A variant of the FACTory game is the Concept Game on Facebook (Herdagdelen and Baroni, 2010), which collects common sense knowledge by proposing random assertions to users in a slot machine fashion and asking them to decide whether this assertion is meaningful or not (Fig. 2).
Page Hunt (Ma et al., 2009) is a GWAP for the annotation of web-sites. It allows to index web sites and hence to improve the search index of a search engine (Microsoft Bing). The player gets assigned a random web-site and is asked to describe it with keywords. The game then shows players the top five page hits for the entered keywords and they are rewarded depending on how high ranked the previously assigned web page is in the result set.
Other GWAPs engage players in building ontologies. OntoPronto (Siorpaes and Hepp, 2008), for example, is a quiz game for vocabulary building that attempts to build a huge domain ontology from Wikipedia[3] articles. This is achieved by mapping random articles to the most specific class of the Proton ontology using the *subClassOf* relationship.
Virtual Pet Game (Kuo et al., 2009) aims to construct a semantic network that encodes common sense knowledge. The game is built on top of PPT, a popular Chinese bulletin board system that is accessible through a terminal interface. Each player owns a pet, which they should take care of by asking and answering questions. The pet in this game is just a substitute for other players, who receive such questions and answers, and have to respond or validate them. Rapport Game (Kuo et al., 2009), similarly to Virtual Pet Game, exploits player labour for constructing a semantic network that encodes common sense knowledge.

[2] `http://youtube.com`

[3] `http://wikipedia.org`

Figure 2: A screenshot of Concept Game, a turn-based single player game taking advantage of the Facebook platform for collecting random pieces of common sense knowledge.

Rapport Game, however, is built on top of Facebook and uses direct interaction between players. An interesting game for the creation of formal domain ontologies from Linked Open Data is Guess What?! (Markotschi and Volker, 2010). Given a seed concept, a player has to find a matching URI in DBpedia, Freebase and OpenCyc. The resulting labels/URIs are analyzed by simple NLP tools in order to identify expressions that can be translated into logical operators and break down complex descriptions into small fragments. The game starts with the most general fragment and, at each round, a more specific fragment is connected to it through a logical operator, with players having to guess the concept described by it.

There are GWAPs that try to align ontologies. Wordhunger, for example, is a web-based application that maps WordNet synsets to Freebase. Each game round consists of a WordNet term and up to three suggested possible Freebase articles, among which players have to select the most fitting (or pass or select 'no match'). SpotTheLink is a two player game focusing on the alignment of random concepts from the DBpedia Ontology to the Proton upper ontology. Each player has to select Proton concepts that are either the same as or more specific than a randomly selected DBpedia concept. The data generated by SpotTheLink is a SKOS mapping between the concepts of the two input ontologies.

Based on Wikipedia, there are three Wikiracing game, The Wiki Game, Wikispeedia and WikipediaMaze, where the objective is to find connections between two Wikipedia articles by clicking links within the text. WikipediaGame and Wikispedia focus on completing the race faster and with fewer clicks than other players. In WikipediaMaze, instead, players are allowed to create races for each other and are incentivised to create and play races by earning badges.

3. Sentic Computing

Sentic computing (Cambria and Hussain, 2012) is a multidisciplinary approach to opinion mining and sentiment analysis at the crossroads between affective computing and common sense computing, which exploits both computer and social sciences to better recognize, interpret and process opinions and sentiments over the Web.

In particular, sentic computing involves the use of AI and Semantic Web techniques, for knowledge representation and inference; mathematics, for carrying out tasks such as graph mining and multi-dimensionality reduction; linguistics, for discourse analysis and pragmatics; psychology, for cognitive and affective modelling; sociology, for understanding social network dynamics and social influence; finally ethics, for understanding related issues about the nature of mind and the creation of emotional machines. Unlike statistical classification, which generally requires large inputs and thus cannot appraise texts with satisfactory granularity, sentic computing enables the analysis of documents not only on the page or paragraph level but also on the sentence and clause level.

This is possible thanks to an affective common sense knowledge base built upon ConceptNet (Havasi et al., 2007) and WordNet-Affect (Strapparava and Valitutti, 2004), which provides the cognitive and affective information associated to concepts extracted from opinionated text by means of a semantic parser. Affective common sense knowledge consists of information that people usually take for granted and, hence, normally leave unstated. Affective common sense, in fact, is not a kind of knowledge that we can find in Wikipedia but it consists in all the basic relationships among words, concepts, phrases, emotions and thoughts that allow people to communicate with each other and face everyday life problems. Therefore, in this work, we collect such kind of knowledge through label sequential rules (LSR), crowd sourcing, and GWAP techniques.

4. Sentic Patterns

Human emotions and their modelling are increasingly understood to be a crucial aspect in the development of intelligent systems (Picard, 1997; Minsky, 2006; Cambria and Hussain, 2012). Emotions are a basic part of human communication and have therefore to be taken into account for the development of more effective interfaces for human-machine communication such as chat systems, e-house, e-learning, e-health or emphatic voice boxes.

The Open Mind Common Sense project has been collecting general common sense knowledge from volunteers over the Web since 2000. Part of such knowledge contains affective information, e.g., "a gift is for celebrating a birthday" or "making a mistake causes embarrassment", as common sense encompasses, among many other aspects, also knowledge about the emotional facets of typical everyday life.

However, the amount of affective information contained in the Open Mind corpus is still very limited, as relationships such as ArisesEmotion, MakesFeel or AffectivelyRelated are missing from the set of properties that are currently used for collecting pieces of common sense knowledge from the public. Since computers now have the ability to search vast amounts of data in little time, the use of a search engine to collect the affective information we need is pretty tempting. To this end, we use different lexical patterns, which we call sentic patterns, for extracting affective information from the Web. We build such patterns using label sequential rules (LSR), which are generated from sequential patterns in data mining (Liu et al., 2005). A rule is of the form $X \rightarrow Y$, where Y is a sequence and X is a sequence produced from

Y by replacing some of its items with wildcards, denoted by a '*', which can match any item. During the learning process, each segment is converted to a sequence. Each sequence element is a word, which is represented by both the word itself and its POS tag in a set. In the training data, all concepts are manually labelled and replaced by the label *\$concept*. A concept can be expressed with a noun (NN), adjective (JJ), verb (VB) or adverb (RB). The labels and their POS tags used in mining LSRs are {*\$concept*, NN}, {*\$concept*, JJ}, {*\$concept*, VB} and {*\$concept*, RB}, where *\$concept* denotes a concept to be extracted. For example, the sentence segment "chocolate makes me feel happy" is turned into the sequence <{chocolate, NN}{make (me|you) feel, VB}{happy, JJ}>. After labeling, it becomes <{*\$concept*, NN}{make (me|you) feel, VB}{happy, JJ}>. All the resulting sequences are then used to mine LSRs. A typical rule, for example, is <{*, NN}{put (me|you) on, VB}{cloud nine, NN}> → <{*\$concept*, NN}{put (me|you) on, VB}{cloud nine, NN}> confidence = 80%, where the confidence is the conditional probability, Pr($Y|X$), which measures the accuracy of the rule. Concept extraction is performed by matching the patterns with each sentence segment in a new web-page to extract affective information about concepts contained AffectNet. That is, the words in the sentence segment that both match *\$concept* in a pattern and an AffectNet concept are extracted. In the pattern match, only the right-hand side of each rule is used. In rule generation, both the right- and the left-hand sides are needed to compute the conditional probability or confidence.

Such patterns yield some useful data, however, they are not good enough for our purposes for three reasons. Firstly, most of the affective common sense knowledge we are trying to collect is so obvious that no one has bothered to record it. Secondly, there exists incorrect knowledge on the Web (for example, the query "plastic can love" returns 33,200 results on Google, while "plastic cannot love" returns just a few links). Thirdly, the text on the Web is unstructured and turning it into a directly useful format is a non-trivial task. For these reasons, we mainly rely on crowd-sourcing techniques.

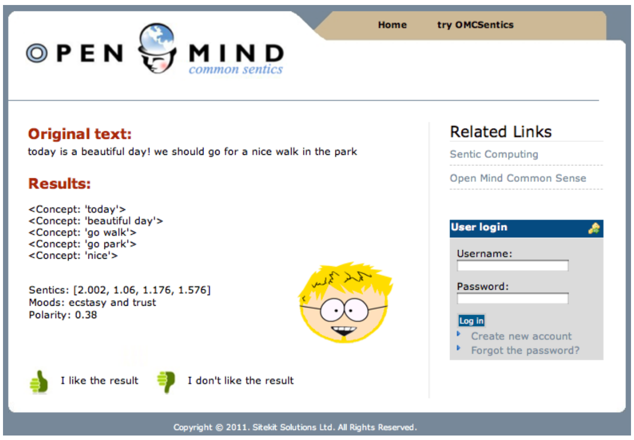

Figure 3: A typical output on Open Mind Common Sentics contains a list of extracted concepts, their valence and polarity, a list of sentics and a polarity value.

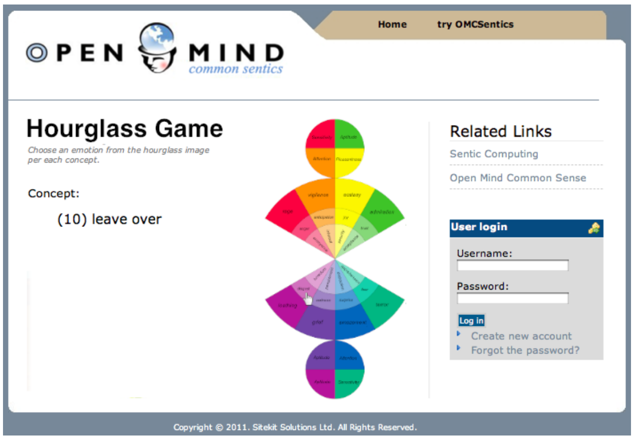

Figure 4: A screenshot of the Hourglass Game. The GWAP aims to collect affective common sense knowledge from the general public by engaging users in a speed game.

5. Hourglass Game

Distributed online knowledge acquisition projects have become quite popular in the past years. Examples include Freebase[4], with its 1,450 concepts, WikiTaxonomy, counting 127,000 concepts, YAGO[5], with 149,162 instances, NELL[6], containing 959,654 beliefs, ProBase[7], Microsoft's universal probabilistic ontology, and the different projects associated with the Open Mind Initiative, e.g., OMCS, Open Mind Word Expert (Mihalcea, 2003), an active learning system that aims to create large annotated corpora, and Open Mind Indoor Common Sense (Gupta et al., 2004), which aims to develop intelligent mobile robots for use in home and office environments.

In a similar fashion to the Open Mind family of distributed knowledge capture projects, Open Mind Common Sentics[8] aims to collect affective common sense knowledge for sentiment analysis (Fig. 3). Whereas previous approaches have mainly relied on paid experts or unpaid volunteers, we put much stronger emphasis on creating a system that is appealing to a large audience of people, regardless of whether or not they are interested in contributing to AI.

The fundamental aim of Open Mind Common Sentics, in fact, is to transform as much as possible the activity of entering knowledge into an enjoyable interactive process. To this end, the system adopts a two-fold strategy: crowd sourcing, that is challenge volunteers over the Web through mood-spotting and fill-in-the blank questions, in the same wake as Open Mind Commons (Speer, 2007), and GWAPs, that is engage users through online games, in the same wake as Verbosity and ESP game. In particular, the mood-spotting questions consist in asking users to select an emoticon according the overall affect they can infer from a given sentence. The fill-in-the blank questions, in turn, are sentences to be completed such as "opening a Christmas gift makes feel ___".

[4] http://freebase.com
[5] http://mpi-inf.mpg.de/yago-naga/yago
[6] http://rtw.ml.cmu.edu/rtw
[7] http://research.microsoft.com/probase
[8] http://omcsentics.labs.sitekit.net

Sometimes, one or more taboo affect concepts are shown in the game window. Such concepts are entries that have been validated a sufficient amount of times; hence they are not valid input any more (in order to collect synonyms or alternatives of a given affective common sense concept). As for the GWAPs, we developed the Hourglass Game (Fig. 4), a speed game consisting in selecting from the Hourglass model the sentic level that is most likely associated with a given affective concept. Players earn points not only according to accuracy but also quickness in clicking on the right area of the Hourglass. The game is quite engaging, although very simple, and players like to challenge each other to see who has higher emotional quotient (EQ) but users are not too keen on playing more than once. What is lacking from most of crowd-sourcing and GWAP techniques, in fact, is stickiness.

GWAPs can be fun to play for a relatively short period of time but then players are not too much keen on returning. In other words, GWAPs generally have a pretty low sticky factor. The sticky factor is defined as the amount of daily active users (DAUs) of an application divided by the number of monthly active users (MAUs). MAU is the most-quoted measure of a game's size, but it is effective only to discuss size or reach, not engagement.

DAU, in turn, can be a very valuable number as it relates how much activity your game is seeing on a daily basis, but it falls into the same trap as MAU in that it does not discriminate between retention and acquisition. The single-most important metric for engagement is stickiness, i.e., DAU/MAU, which allows to more accurately calculate repeat visits and average knowledge acquired per user (AKAPU). A key for driving the sticky factor, besides great game play, is the ability of the application to prompt users to reach out to their friends, e.g., via stories and pictures about their gameplay.

6. Sentic Pet

To this end, we are developing Sentic Pet (Fig. 5), a massively multi-player online (MMO) game in which players have to raise and take care of their own pets. Unlike old-style tamagotchi games, in Sentic Pet, raising and caring pets is not about cleaning, feeding and petting them but rather training them, both at mental and skill level, by playing mini-GWAPs. Targeting players of a wide age range (10 to 50 year old), the game should appeal anyone who enjoyed and enjoys PetVille[9] or FarmVille[10].

Players start from level 1 with their pet being a baby born having very little affective common sense knowledge. The game involves balancing two main activities: training the pet and testing its skills by challenging other players. Training does not involve simply teaching the pet new knowledge, but also refining acquired knowledge. Challenges can be taken both at mental and skill level, which involve different kinds of activities. At mental level, for example, pets can be challenged according to different modalities, e.g., affective vocabulary learning (in the same wake of Verbosity) or affective meaning of images (in the same wake of ESP game).

[9]http://petville.com
[10]http://farmville.com

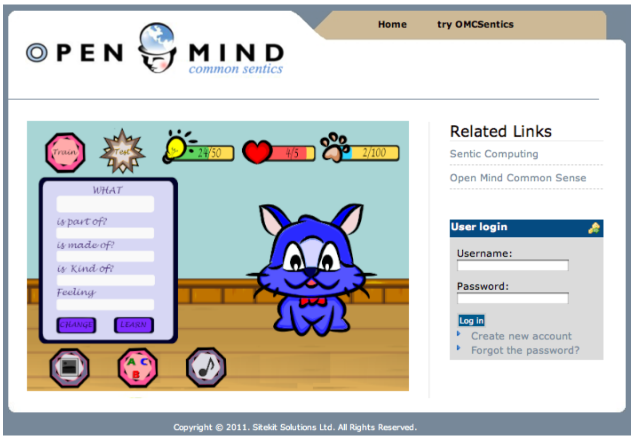

Figure 5: Sentic Pet. Icons in the up-right corner specify the abilities of the pet. Icons in the up-left corner allow players to train their pets according to different modalities.

Pets can level-up according to the combination of IQ (light bulb icon) and EQ (heart icon) points earned playing the mini-GWAPs. Data are validated by majority and reputation, that is, the confidence score with which a piece of affective common sense knowledge is saved into the knowledge base depends on how many players validated it and on the expertise level of these.

7. Conclusion and Future Work

Affective common sense knowledge consists in all the basic relationships among words, concepts, phrases, emotions, and thoughts that allow people to communicate with each other and face everyday life problems. It is not a kind of knowledge that we can easily find in text documents and web pages as people usually take it for granted and, hence, normally leave it unstated. However, it is a kind of information that can be extremely useful for tasks such as sentiment analysis.

In this work, we aim to collect affective common sense knowledge through label sequential rules (LSR), crowd sourcing, and GWAP techniques. Open Mind Common Sentics, in particular, is an example of an emerging class of games that can be considered 'human algorithms', since humans act as processing nodes for problems that computers cannot yet solve. By providing an incentive for players, we gain a large quantity of computing power that can be harnessed for multiple applications.

In the future, we plan to further develop and exploit Sentic Pet and the GWAPs related to it in order to gather and refine more and more pieces of affective common sense knowledge that can be shared among the different games and, hence, improve their stickiness. Constructing a comprehensive affective common sense database, finally, will be extremely beneficial for many research and business communities working in fields such as social data mining, human-computer interaction, social media marketing, and patient-centred health-care, and we believe Open Mind Common Sentics can be highly effective in doing so.

8. References

Matthew Addis, Laurent Boch, Walter Allasia, Francesco Gallo, Werner Bailer, and Richard Wright. 2010. 100 million hours of audiovisual content: Digital preservation and access in the prestoprime project. In *Digital Preservation Interoperability Framework Symposium*, Dresden.

Luke Barrington, Damien O'Malley, Douglas Turnbull, and Gert Lanckriet. 2009. User-centered design of a social game to tag music. In *ACM SIGKDD*, pages 7–10, Paris.

Erik Cambria and Amir Hussain. 2012. *Sentic Computing: Techniques, Tools, and Applications*. Springer, Berlin Heidelberg.

Erik Cambria, Marco Grassi, Amir Hussain, and Catherine Havasi. 2011. Sentic computing for social media marketing. *Multimedia Tools and Applications*, http://dx.doi.org/10.1007/s11042-011-0815-0.

Erik Cambria, Tim Benson, Chris Eckl, and Amir Hussain. 2012a. Sentic PROMs: Application of sentic computing to the development of a novel unified framework for measuring healthcare quality. *Expert Systems with Applications*, http://dx.doi.org/10.1016/j.eswa.2012.02.120.

Erik Cambria, Catherine Havasi, and Amir Hussain. 2012b. SenticNet 2: A semantic and affective resource for opinion mining and sentiment analysis. In *FLAIRS*, Marco Island.

Rakesh Gupta, Mykel Kochenderfer, Deborah Mcguinness, and George Ferguson. 2004. Common sense data acquisition for indoor mobile robots. In *AAAI*, pages 605–610, San Jose.

Severin Hacker and Luis von Ahn. 2009. Matchin: Eliciting user preferences with an online game. In *CHI*, Boston.

Catherine Havasi, Robert Speer, and Jason Alonso. 2007. ConceptNet 3: A flexible, multilingual semantic network for common sense knowledge. In *RANLP*, Borovets.

Amac Herdagdelen and Marco Baroni. 2010. The concept game: Better commonsene knowledge extraction by combining text mining and game with a purpose. In *AAAI CSK*, Arlington.

YenLing Kuo, JongChuan Lee, Kaiyang Chiang, Rex Wang, Edward Shen, ChengWei Chan, and Jane Yung-Jen Hu. 2009. Community-based game design: Experiments on social games for commonsense data collection. In *ACM SIGKDD*, pages 15–22, Paris.

Edith Law, Luis von Ahn, Roger Dannenberg, and Mike Crawford. 2007. Tagatune: A game for music and sound annotation. In *International Conference on Music Information Retrieval*, Vienna.

Douglas Lenat and Ramachandra Guha. 1989. *Building Large Knowledge-Based Systems: Representation and Inference in the Cyc Project*. Addison-Wesley, Boston.

Bing Liu, Minqing Hu, and Junsheng Cheng. 2005. Opinion observer: Analyzing and comparing opinions on the web. In *WWW*, Chiba.

Hao Ma, Raman Chandrasekar, Chris Quirk, and Abhishek Gupta. 2009. Page hunt: Improving search engines using human computation games. In *SIGIR*, pages 746–747, Boston.

Thomas Markotschi and Johanna Volker. 2010. GuessWhat?! - Human intelligence for mining linked data. In *EKAW*, Lisbon.

Rada Mihalcea. 2003. Open mind word expert: Creating large annotated data collections with web users' help. In *Workshop on Linguistically Interpreted Corpora*, Budapest.

Marvin Minsky. 2006. *The Emotion Machine: Commonsense Thinking, Artificial Intelligence, and the Future of the Human Mind*. Simon & Schuster, New York.

Donn Morrison, Stephane Maillet, and Eric Bruno. 2009. Tagcaptcha: Annotating images with captchas. In *ACM SIGKDD*, pages 44–45, Paris.

Rosalind Picard. 1997. *Affective Computing*. The MIT Press, Boston.

PopCap. 2010. 2010 social gaming research. Technical report, Pop Cap Games.

Katharina Siorpaes and Martin Hepp. 2008. Ontogame: Weaving the semantic web by online games. In *ESWC*, pages 751–766, Tenerife.

Robert Speer. 2007. Open Mind Commons: An inquisitive approach to learning common sense. In *Workshop on Common Sense and Interactive Applications*, Honolulu.

Carlo Strapparava and Alessandro Valitutti. 2004. WordNet-Affect: An affective extension of WordNet. In *LREC*, Lisbon.

Stefan Thaler, Katharina Siorpaes, Elena Simperl, and Christian Hofer. 2011. A survey on games for knowledge acquisition. Technical report, Semantic Technology Institute.

Roelof van Zwol, Lluis Garcia, Georgina Ramirez, Borkur Sigurbjornsson, and Marcos Labad. 2008. Video tag game. In *WWW*, Beijing.

Luis von Ahn and Laura Dabbish. 2004. Labeling images with a computer game. In *CHI*, pages 319–326, Vienna.

Luis von Ahn, Manuel Blum, Nicholas Hopper, and John Langford. 2003. Captcha: Using hard ai problems for security. In ELi Biham, editor, *EUROCRYPT*, Lecture Notes in Computer Science, pages 294–311. Springer, Berlin Heidelberg.

Luis von Ahn, Shiry Ginosar, Mihir Kedia, Ruoran Liu, and Manuel Blum. 2006a. Improving accessibility of the web with a computer game. In *CHI*, pages 79–82, Quebec.

Luis von Ahn, Mihir Kedia, and Manuel Blum. 2006b. Verbosity: A game for collecting common sense facts. In *CHI*, pages 75–78, Quebec.

Luis von Ahn, Ruoran Liu, and Manuel Blum. 2006c. Peekaboom: A game for locating objects in images. In *CHI*, pages 55–64.

Luis von Ahn. 2006. Games with a purpose. *IEEE Computer Magazine*, 6:92–94.

Jeff Yan and Su-Yang Yu. 2009. Magic bullet: a dual-purpose computer game. In *ACM SIGKDD*, pages 32–33, Paris.