Opinion Mining Meets Decision Making: Towards Opinion Engineering

Klemens Schnattinger[1] and Heike Walterscheid[2]
[1]Computer Science, Baden-Wuerttemberg Cooperative State University Loerrach, Hangstrasse 46-50, 79539 Loerrach, Germany
[2]Economics, Baden-Wuerttemberg Cooperative State University Loerrach, Hangstrasse 46-50, 79539 Loerrach, Germany

Keywords: Opinion Mining, Decision Making, Natural Language Processing, Description Logics.

Abstract: We introduce a methodology for opinion mining based on recent approaches for natural language processing and machine learning. To select and rank the relevant opinions, decision making based on weighted description logics is introduced. Therefore, we propose an architecture called OMA (Opinion Mining Architecture) that integrates these approaches of our methodology in a common framework. First results of a study on opinion mining with OMA in the financial sector are presented.

1 INTRODUCTION

The work reported is part of the project OMA aiming at the development of an opinion mining and evaluation system for real-world domains. The methodological challenge is two-fold. The opinion mining task is that the textual sources must be pre-processed and analysed as well as the opinion evaluation task is to put the opinions in an order.

To address these problems, we concentrate on foundations of Natural Language Processing (NLP) in combination with machine Learning (ML) (Sun, Luo and Chen, 2017) and on Weighted Description Logics, an extension of "classical" Description Logics (DL) with utility theory for the calculation of quantitative preference relations (Acar et al., 2017). Hence, we combine these techniques in a common architecture, called the Opinion Mining Architecture (OMA). In addition, we present data from a first empirical evaluation of OMA. Qualitative measures are the subject of future research to focus more on validity and causality of sentiments.

2 OPINION MINING

2.1 Opinion and Opinion Mining

Usually, the term *opinion* is defined as *"the personal view that someone has about something"* (Dictionary, 2002). Formally, an opinion is defined as follows (Liu, 2012): $(e_i, a_{ij}, h_k, t_l, s_{ijkl})$, where e_i denotes the ith entity, a_{ij} the jth aspect of the ith entity, h_k the kth opinion holder, t_l the time when the opinion is expressed, s_{ijkl} the opinion towards the jth aspect of the ith entity from opinion holder h_k at time t_l.

For example, in "*The screen of this tablet is good*", the components e_i, a_{ij} and s_{ijkl} can be identified: *screen* is an aspect of the entity *tablet*. Additionally, a *positive sentiment* is expressed. h_k and t_l are not given, that is, the five components are not always necessary to express an opinion.

To perform opinion mining, machine learning approaches are meaningful. Classifiers are used and are trained with known texts to identify their sentiment orientation. For the task of identifying the opinion holder, detecting opinion expressions, and identifying the target or aspect of the opinion, corpora with annotated opinion or sentiment scores are necessary but difficult to get.

In contrast, lexicon approaches identify the sentiment of text purely without a training set according to given sentiment lexicons. A sentiment lexicon is a dictionary of sentiment words and phrases, contains a sentiment orientation and a strength for each sentiment entry, which is expressed through a sentiment score. Lexicons use less resources, because they don't use annotated corpora. In addition, such a sentiment lexicon can be

Schnattinger K. and Walterscheid H.
Opinion Mining Meets Decision Making: Towards Opinion Engineering.
DOI: 10.5220/0006576403340341
In *Proceedings of the 9th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management (KDIR 2017)*, pages 334-341
ISBN: 978-989-758-271-4

integrated into machine learning approaches. Thus, performance can be significantly increased.

2.2 Natural Language Processing

To perform opinion mining the reviewed texts must be pre-processed. For this purpose, the following processes are usually carried out for structuring the text and for extracting features:

Tokenization decomposes a sentence or document into tokens. Tokens represents words or phrases. For English or German, the decomposition of words is easy with spaces, but some additional expertise should be kept in mind, such as opinion phrases and named entities. Words, such as *"the"*, *"a"* only provide little information. Thus, tokenization must remove these words, which are called *stop words*.

POS tagging is a technique that analyses the lexical information of a word for determining their POS tag (e.g. *adjective* or *noun*). POS tagging is a so-called *sequential labelling problem*. Conditional Random Fields (CRFs) (Lafferty, McCallum and Pereira, 2001) and Markov models (Sutton and McCallum, 2011) are applied to this problem. On the one hand, adjectives can represent opinion words. On the other hand, entities and aspects of opinion mining can be expressed with nouns or combination of nouns.

Parsing is a technique that provides syntactic information. Among other things, it analyses the grammatical structures of a given sentence and generates a tree with the corresponding relationship of different so-called constituents as "a group of words treated by a syntactic rule as a unit" (Carnie, 2010). Unlike POS tagging, parsing determines richer structural information. It can be used especially for fine-grained opinion mining (Socher, Bauer and Manning, 2013).

2.3 Machine Learning

For opinion mining gaining features from texts is important. Thus, text features are discussed, including *n-gram features* with weighting schemes, *syntactic features* and *semantic features.*

An *n-gram* is a set of *n* adjacent items. Additionally, the number of times an item appears in the text is denoted. In opinion mining, double-digit weights of unigram and bigram are widely accepted. Instead of binary weights, other schemes can be used (Paltoglou and Thelwall, 2010).

Syntactic features include POS tags and syntactic information. These features either build up a feature space for machine learning approaches (Joshi and Penstein-Rosé, 2009), or generate rules for e.g. entities and aspects in fine-grained opinion mining (Gindl, Weichselbraun and Scharl, 2013).

Semantic features are conjunctions which specifies negation, increase, and decrease of a sentiment. Negation turns the sentiment orientation into the opposite. Increase and decrease also influence the strength of sentiment, respectively, are useful for opinion mining (Taboada et al., 2011).

Opinion mining is usually divided into three levels: *document level*, *sentence level*, and *fine-grained level*. The task of the *document level opinion mining* determines sentiment orientation of an entire document. The objective of the document level opinion mining is identifying the s_{ijkl} in $(e_i, a_{ij}, h_k, t_l, s_{ijkl})$. Recent techniques for document level opinion mining are among others:

Supervised approaches: Usual classifiers in machine learning, such as a Naïve Bayes or Support Vector Machines, are used. The features considered are, among others, n-gram, POS tags, position information (Pang, Lee and Vaithyanathan, 2002) and semantic features (Kennedy and Inkpen, 2006).

Probabilistic generative models: Generative models such as joint sentiment topic model (Lin and He, 2009) are proposed which use a Markov chain.

Unsupervised lexicon-based approaches: Averaged sentiment orientation is used to suggest the overall sentiment orientation of an entire document (Turney, 2002). To improve the results e.g. discourse structure-based weighting scheme (Bhatia, Ji and Eisenstein, 2015) are proposed.

In opinion mining at the *sentence level*, sentiment orientation is determined for each sentence in the document. However, not all the detailed information of opinions is collected such as opinion target and opinion holder. For example, "*The screen of this tablet is good*." expresses a positive sentiment orientation to aspect "*screen*" of entity "*tablet*". Recent techniques for sentence level opinion mining are among others:

Supervised approaches: Again, Naïve Bayes classifiers are used to determine subjectivity of sentences (Yu and Hatzivassiloglou, 2003) and CRFs for the dependencies between sentences (Yang and Cardie, 2015).

Unsupervised approaches: For subjectivity classification in sentences graph-based (Pang and Lee, 2004), as well as lexicon-based approaches (Kim and Hovy, 2004) exists.

The problems with the *fine-grained level opinion mining* can't be traced with traditional classification techniques. Several variations are suggested including *aspect level opinion mining* (Cambria et al., 2013) that aims to discover aspects or entities of opinion mining and the corresponding sentiment orientation. Thus, it is split into two sub-tasks: *opinion target extraction* and *sentiment classification*. Recent techniques for fine-grained level opinion mining are:

Unsupervised approaches: Association mining algorithm for aspect detection and linguistic knowledge (Popescu, 2005) and part-whole patterns (Zhang et al., 2010) are considered. For aspects extraction (Qiu et al., 2009) propose *propagation algorithms.* Additionally, rule-based methods are also suitable (Gindl, Weichselbraun and Scharl, 2013).

Probabilistic generative models: For aspects detection (Brody and Elhadad, 2010) and sentiment detection (Lazaridou, Titov and Sporleder, 2013) so called *Latent Dirichlet Allocation (LDA)* topic models are adopted.

2.4 Comparative Opinion Mining

A comparative opinion is defined as a relationship of similarities or differences between two entities. Comparative opinion mining takes these entities and preferences of opinion holders into account. From comparative sentences, compared entities, comparative words and aspects can be extracted. For instance, in "*Tablet X's screen is better than tablet Y.*", "*tablet X*" and "*tablet Y*" are the compared entities, "*better*" is the comparative word and "*screen*" is the compared aspect. Because the word "*better"* expresses the preference, "*tablet X*" is preferred. However, many comparative words, e.g., "*longer*", express different positive or negative sentiment orientations in different contexts.

A rule-based method for this kind of sentence decomposes this problem into two sub-tasks (Jindal and Liu, 2006): *comparative sentence identification* and *comparative relation extraction*. Class Sequential Rules (CSRs) with class labels (i.e., "*comparative*" or "*noncomparative*") and Label Sequential Rules (LSRs) applied on comparative sentences help solving these tasks, respectively.

Another method divides comparative sentences into two categories: *opinionated comparatives* and *comparatives with context-dependent opinions* (Ganapathibhotla and Liu, 2008). In the first case, comparative words are used. In the second, external information is needed.

3 DECISION MAKING

3.1 Preference and Utility

Preferences are an important variable in the study of decisions such as in mathematical economics, social choice theory and opinion mining. To keep it simple in the beginning preferences will be "modelled as a binary relation over the set of choices" (Kaci, 2011). A set of choices for a rational agent as homo oeconomicus (Mill, 1836) which has the preference relation $\succ$ are named $\mathcal{C}$ and $c_1 \succcurlyeq c_2$ is read "c_1 *is at least as good as* c_2" where $c_1, c_2 \in \mathcal{C}$. Furthermore, at $\succcurlyeq$ is a complete, reflexive and transitive relation. There are two preference relations for $\succcurlyeq$:

- for any $c, c' \in \mathcal{C}$, $c \succ c'$ **iff**
  $c \succcurlyeq c'$ and $c' \not\succcurlyeq c$ (*Strict preference*)
  This is read: c *is better than* c'.
- for any $c, c' \in \mathcal{C}$, $c \sim c'$ **iff**
  $c \succcurlyeq c'$ and $c' \succcurlyeq c$ (*Indifference*)
  This is read: *the agent is indifferent between* c *and* c'.

A utility function u maps a choice to a real number representing the degree of request. The representation theorems formally are defined as follows (Fishburn, 1969):

Given the choices $c, c' \in \mathcal{C}$ a *utility function*, $u: \mathcal{C} \to \mathbb{R}$ represents

$$\begin{aligned} &\succcurlyeq \text{ if } \quad c \succcurlyeq c' \ \textbf{iff} \quad u(c) \geq u(c') \\ &\succ \text{ if } \quad c \succ c' \ \textbf{iff} \quad u(c) > u(c') \\ &\sim \text{ if } \quad c \sim c' \ \textbf{iff} \quad u(c) = u(c') \end{aligned}$$

For instance, if $u(high\_price) = 20$ and $u(low\_price) = 5$, this leads to $high\_price \succ low\_price$ since 5 < 20. This means, the choices low_price and $high\_price$ are values of a single attribute $price$ (Acar et al., 2017).

Normally, due to framing or irrationality e.g. decisions are more complex (Tversky and Kahneman, 1981). Therefore, choices are formalized as *values* or *elements of attributes*. For instance, if we will buy a car, not only the price will be of interest, but also its colour, and even more. Formally, the set of attributes is denoted by $\mathcal{X}$. Then, $X_i \in \mathcal{X}$ refer to a specific attribute in $\mathcal{X}$ where $i \in \{1, \ldots, |\mathcal{X}|\}$. With these preliminaries, we can formalize the set of choices made by the cartesian product over the set of attributes. This set of choices is denoted by Ω where $\Omega = X_1 \times \ldots \times X_n$. Now, the utility function u has been expanded: $u : \Omega \to \mathbb{R}$ is the (multi-attribute) utility function which represents $\succcurlyeq$ **iff**

$\forall (x_1, \ldots, x_n), (y_1, \ldots, y_n) \in \Omega,$

$$(x_1, \ldots, x_n) \succcurlyeq (y_1, \ldots, y_n) \ \textbf{iff}$$
$$u(x_1, \ldots, x_n) \geq u(y_1, \ldots, y_n)$$

The size of the Ω is $2^{|\mathcal{X}|}$, the assumption that u is *additive* helps to significantly reduce the complexity. A typical additive function is

$$u(x_1, \ldots, x_n) = u(x_1) + \ldots + (x_n) \quad (Additivity)$$

where $(x_1, \ldots, x_n) \in \Omega$.

Now, we can formulate an optimization task, namely that a rational agent should make the choice with the maximum utility:

$$Opt(\mathcal{C}) \coloneqq \arg\max_{c \in \mathcal{C}} u(c) \qquad (Optimal\ choice)$$

where $Opt(\mathcal{C})$ matches to maximal elements in $\mathcal{C}$ with respect to the utility function u (and therefore means w.r.t. the preference relation $\succcurlyeq$).

3.2 Description Logics

The signatures of description logics (Baader et al., 2003) can be given as a triple (N_C, N_R, N_I), where N_C denotes the set of atomic concepts, N_R the set of role names and N_I the set of atomic individuals.

We denote *concepts* or *classes* by C and D, *roles* by R and S, and *individuals* as a and b. Concept descriptions are defined in a common way from N_C as $\neg C$, $C \sqcap D$, and $C \sqcup D$ if C and D are concept descriptions. Further, $\exists r.C$ and $\forall r.C$ exist if $r \in N_R$ and C is a concept description. The top concept $\top$ is an abbreviation for $C \sqcup \neg C$ and $\bot$ for $\neg\top$.

For the semantic we need an interpretation for the presented syntax. An interpretation is a pair $\mathcal{I} \coloneqq (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ where the domain $\Delta^{\mathcal{I}}$ is a set that can't be empty, and $\cdot^{\mathcal{I}}$ is a so-called interpretation function. This function maps to every concept name C a set $C^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$ and to every role name R a binary relation $R^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$. The function also defines:

$$(\neg C)^{\mathcal{I}} \coloneqq \Delta^{\mathcal{I}} \setminus C^{\mathcal{I}} \qquad (C \sqcap D)^{\mathcal{I}} \coloneqq C^{\mathcal{I}} \cap D^{\mathcal{I}}$$
$$(C \sqcup D)^{\mathcal{I}} \coloneqq C^{\mathcal{I}} \cup D^{\mathcal{I}} \qquad (C \sqsubseteq D)^{\mathcal{I}} \coloneqq C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$$
$$(\exists r.C)^{\mathcal{I}} \coloneqq \{a \in \Delta^{\mathcal{I}} \mid \text{exists } b, (a,b) \in r^{\mathcal{I}}, b \in C^{\mathcal{I}}\}$$
$$(\forall r.C)^{\mathcal{I}} \coloneqq \{a \in \Delta^{\mathcal{I}} \mid \text{for all } b, (a,b) \in r^{\mathcal{I}} \rightarrow b \in C^{\mathcal{I}}\}$$

In DLs, we distinguish between terminological knowledge (so-called *TBox*) and assertional knowledge (so-called *ABox*). A TBox is a set of concept inclusions $C \sqsubseteq D$ which has the semantics $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$ and a concept definition is $C \equiv D$ if $C \sqsubseteq D$ and $D \sqsubseteq C$. An ABox is a set of concept assertions $C(a)$ where $a \in N_I$ and $C(a)^{\mathcal{I}} \coloneqq a^{\mathcal{I}} \in C^{\mathcal{I}}$, as well as role assertions $R(a,b)$ where $(a,b) \in N_I \times N_I$ and $R(a,b)^{\mathcal{I}} \coloneqq (a^{\mathcal{I}}, b^{\mathcal{I}}) \in R^{\mathcal{I}}$.

In the following we will consider only a *coherent* TBox $\mathcal{T}$. This means that all concepts in $\mathcal{T}$ are satisfiable. The usual interpretation function is used for the notion satisfiable (Baader et al., 2003) and write $\mathcal{T} \vDash C$. We say that an ABox $\mathcal{A}$ *entails* an assertion α (and write $\mathcal{A} \vDash \alpha$), if every model of $\mathcal{A}$ also satisfies α. An ABox $\mathcal{A}$ is called *consistent* with a TBox $\mathcal{T}$ if there exists an interpretation $\mathcal{I}$ that satisfies $\mathcal{T}$ and $\mathcal{A}$. We then call the pair $\mathcal{K} \coloneqq \langle \mathcal{T}, \mathcal{A} \rangle$ a *knowledge base*. Further, $\mathcal{K}$ is satisfiable if $\mathcal{A}$ is consistent w.r.t. $\mathcal{T}$. In the remainder, we will use the *instance check*. Thus, for a knowledge base $\mathcal{K}$ and an assertion α, one can check whether $\mathcal{A} \vDash \alpha$ holds.

A *concrete domain* $\mathcal{D}$ is defined as a pair $(\Delta^{\mathcal{D}}, pred(\mathcal{D}))$. $\Delta^{\mathcal{D}}$ is the domain of $\mathcal{D}$ and *pred*($\mathcal{D}$) is the set of predicate names of $\mathcal{D}$. The following assumptions have been applied: $\Delta^{\mathcal{I}} \cap \Delta^{\mathcal{D}} = \emptyset$ and for each $P \in pred(\mathcal{D})$ with arity n there is $P^{\mathcal{D}} \subseteq (\Delta^{\mathcal{D}})^n$. According to (Baader et al., 2003), *functional roles* are denoted with lower case letters, for example with r. In description logics with concrete domains, N_R is partitioned into a set of functional roles and one of ordinary roles. A role r is *functional* if for every $(x,y) \in r$ and $(w,z) \in r$ it is necessary that $x = w \Rightarrow y = z$. Functional roles are explained as partial functions from $\Delta^{\mathcal{I}}$ to $\Delta^{\mathcal{I}} \times \Delta^{\mathcal{D}}$. A concrete domain is closed under negation (denoted by $\bar{P}$). For this reason, a logical formula can be calculated which are in the so-called negation normal form (NNF). A formula is in NNF when the negation operators are only used between atomic statements.

3.3 Weighted Description Logics

We will introduce an ontological approach to decision making. This approach can be considered as a generic framework, the so-called DL *decision base* (Acar et al., 2017). We use an *a priori* preference relation over attributes (called the *ontological classes*). Thereby, an *a posteriori* preference relation over choices (called *ontological individuals*) can be derived. Formally, a priori utility function U over $\mathcal{X}$ (the set of attributes) is defined ($U: \mathcal{X} \rightarrow \mathbb{R}$). Additionally, a utility function u defined over choices, which uses logical entailment, extends the utility function U to the subset of attributes. The utility function u was used because a choice was defined as an individual and its outcome as a set of concepts. Another reason is that U can take various forms, e.g., $max, mean$. Modelling attributes has two steps:

1. Each attribute is modelled by a concept.
2. For every value of an attribute a new (sub)concept has been introduced.

For instance, if *colour* is an attribute to be modelled, it is simply represented by the concept $Colour$ (i.e., $Colour \in \mathcal{X}$). A colour can be regarded as a value, as if it were a concept of its own. If *blue* is a value of the attribute *colour*, the attribute set $\mathcal{X}$ is simply extended by adding the concept *Blue*, as a sub-concept of $Colour$. It should be noted, that an axiom has been introduced to guarantee the disjointedness. (e.g. $Red \sqsubseteq \neg Blue$) and that this procedure results in a binary term vector for $\mathcal{X}$, because an individual *c* (as a choice) is either a member of the concept $\mathcal{X}$ or not.

Given a total preference relation (i.e., $\succcurlyeq_{\mathcal{X}}$) over an ordered set of not necessarily atomic attributes $\mathcal{X}$, and a function $U: \mathcal{X} \to \mathbb{R}$ that represents $\succcurlyeq$ (i.e., $U(X_1) \geq U(X_2)$ **iff** $X_1 \succcurlyeq_{\mathcal{X}} X_2$ for $X_1, X_2 \in \mathcal{X}$). The function U asigns an a priori weight to each concept $X \in \mathcal{X}$. Therefore, one can say, that "U makes the description logic weighted". The utility of a concept $X \in \mathcal{X}$ is denoted by $U(X)$. The following applies: The greater the utility of an attribute the more the attribute is preferable. Furthermore, the attribute set $\mathcal{X}$ can be divided into two subsets:

- *desirable* denotes the set of attributes with non-negative weights, denoted $\mathcal{X}^+$, and
- *undesirable* $\mathcal{X}^-$, i.e., $X \in \mathcal{X}$ **iff** $U(X) \geq 0$ and $\mathcal{X} = \mathcal{X}^+ \cup \mathcal{X}^-$ with $\mathcal{X}^+ \cap \mathcal{X}^- = \emptyset$

This means that any attribute that is not in $\mathcal{X}^+$ (not desirable) must lie in $\mathcal{X}^-$ and is therefore undesirable. In addition, it should be noted that an attribute with weight zero can be interpreted as desirable with no utility.

As mentioned above, a *choice* is an individual $c \in N_I$. $\mathcal{C}$ denotes the finite set of choices. To determine a preference relation (*a posteriori*) over $\mathcal{C}$ (i.e., $\succcurlyeq_{\mathcal{C}}$), which respects $\succcurlyeq_{\mathcal{X}}$, a utility function $u(c) \in \mathbb{R}$ is introduced. $u(c)$ indicates *the utility of a choice* c relative to the attribute set $\mathcal{X}$. Also, a utility function U over attributes as an aggregator is introduced. For simplicity, the symbol $\succcurlyeq$ is used for both choices and attributes whenever it is evident from the context.

The σ-utility is a particular u and is defined as $u_\sigma(c) \coloneqq \sum\{U(X) \mid X \in \mathcal{X} \text{ and } \mathcal{K} \vDash X(c)\}$ and is called the *sigma utility of a choice* $c \in \mathcal{C}$. u_σ triggers a preference relation over $\mathcal{C}$ i.e., $u_\sigma(c_1) \geq u_\sigma(c_2)$ **iff** $c_1 \succcurlyeq c_2$. Each choice corresponds to a set of attributes, which is logically *entailed* e.g., $\mathcal{K} \vDash X(c)$. Due to the criterion Additivity, each selection *c* corresponds to a result.

Putting things (DL, U and u) together, a generic *UBox* (so-called *Utility Box*) is defined as a pair $\mathcal{U} := (u_\sigma, U)$, where U is a utility function over $\mathcal{X}$ and u is the utility function over $\mathcal{C}$. Also, a *decision base* can be defined as a triple $D = (\mathcal{K}, \mathcal{C}, \mathcal{U})$ where $\mathcal{K} \coloneqq \langle \mathcal{T}, \mathcal{A} \rangle$ is a consistent knowledge base, $\mathcal{T}$ is a TBox and $\mathcal{A}$ is an ABox, $\mathcal{C} \subseteq N_I$ is the set of choices, and $\mathcal{U} = (u, U)$ is an UBox. Note: $\mathcal{K}$ provides assertional information about the choices and terminological information about the agent ability to reason over choices.

Example:
We want to buy a tablet computer. Two alternatives are considered, which fit the original purpose. The buyer's decision base ($\mathcal{T}, \mathcal{A}$, choices $\mathcal{C} = \{tab_1, tab_2\}$, and attributes mentioned in $\mathcal{U}$) are given. The language uses discrete domains. The domain $tablet$ is used and $\Delta^{tablet} \coloneqq \Delta^{€} \cup \Delta^{g}$ with $\Delta^{€} \cap \Delta^{g} = \emptyset$ and $pred(Tablet) \coloneqq pred(€) \cup pred(g)$. The partition $\Delta^{€}$ of domain Δ^{tablet} is

$\Delta^{€} := \{i\text{ €} \mid i \in \mathbb{S} \subset \mathbb{Q}\}$ and

$pred(€) \coloneqq \{<_{€}, >_{€}, \leq_{€}, \geq_{€}, =_{€}, \neq_{€}\}$ with

$(<_{€})^{€}(x, y) = \{(x, y) \in \Delta^{€} \times \Delta^{€} \mid$ $i, j \in \mathbb{S}$ with $x := i€$ and $y\ j€$ such that $i < j\}$.

Further predicates are defined similar. Note: $pred(€)$ is closed under negation. This means that we can invert the predicates in an obvious way like $\overline{<_{€}}\,(x, y) = \geq_{€}(x, y)$. The other partition is defined as follows: $\Delta^{g} := \{i\ g \mid i \in \mathbb{S}^{+}\{0\}\}$. The remaining predicate names and functional roles are also defined (basic predicate names and functional roles like $Tablet^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$ are not given here):

$\mathcal{T} = \{InexpensiveTablet \sqsubseteq Tablet,$
$ExpensiveTablet \sqsubseteq Tablet,$
$ExpensiveTablet \sqcap InexpensiveTablet \sqsubseteq \bot,$
$\exists hasPrice.\leq_{700\text{ €}} \equiv InexpensiveTablet,$
$UpperClassTablet \sqsubseteq Tablet,$
$Convertable \sqsubseteq UpperClassTablet,$
$LightTablet \sqsubseteq Tablet,$
$\exists hasWeight.\leq_{900\ g} \sqsubseteq LightTablet,$
$\forall hasKeyboard.Keyboard \equiv Convertable,$
$Keyboard \sqsubseteq Device, Device \sqcap Tablet \sqsubseteq \bot\}$

$\mathcal{A} = \{UpperClassTablet(tab_1),$
$hasPrice(tab_1, 769\text{ €}), hasWeight(tab_1, 710\text{ g}),$
$Tablet(tab_2), hasPrice(tab_2, 629\text{ €}),$
$hasWeight(tab_2, 1250\text{ g}), Keyboard(keyb_1),$
$hasKeyboard(tab_2, keyb_1)\}$

$\mathcal{U} = \{(InexpensiveTablet, 50),$
$(UpperClassTablet, 30), (LightTablet, 40),$
$(\exists hasKeyboard.Keyboard, 60)\}$

Considering $\mathcal{U}$ the agent is more interested in a tablet with a keyboard than in an upper class or inexpensive tablet. The utilities can be calculated by $u_\sigma(tab_1) = 30 + 40 = 70$ and $u_\sigma(tab_2) = 50 + 30 + 60 = 140$. Thus, $tab_2 \succ tab_1$.

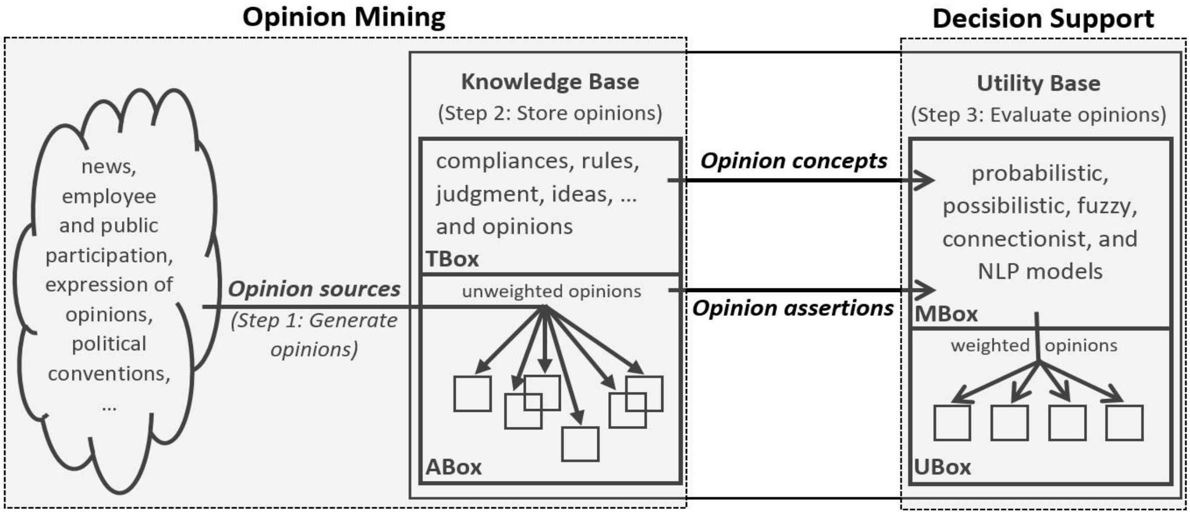

Figure 1: The Opinion Mining Architecture OMA.

4 OMA - THE OPINION MINING ARCHITECTURE

The Opinion Mining Architecture (OMA) we propose is strongly basing on the approaches of natural language processing and machine learning presented in Section 0 as well as on decision making with weighted description logics presented in section 3. We separate OMA into opinion mining (pre-processing, analysing texts, filtering out opinions) and decision support (evaluating extracted opinions) according to SYNDIKATE (Hahn and Schnattinger, 1997). OMA serves the generation of opinions from texts like news, employee and public participation, expressions of opinions, political conversations, etc. (see step 1 in Figure 1). The representation of the underlying domain (TBox) as well as the opinions expressed as assertions (ABox) use a description logic model (see step 2 in Figure 1). The TBox contains concepts which represents artefacts like compliance, rule, judgment, idea, sentiment, opinion, etc. The ABox contains assertions. In terms of content, it consists of opinions that are extracted from the sources of text. Whenever an opinion is stored in the ABox, different types of machine learning and natural language processing models carried out an evaluation (see step 3 in Figure 1). These models are presented in a so-called MBox (*methodology box*). The evaluation provides a ranking of the opinions according to their utility. These weighted opinions are stored in the UBox (see section 3.3). Note: Not every opinion can be weighted and therefore does not appear in the UBox.

In view of OMA architecture, we intend to build a model for opinion mining in various domains such as sentiment mining for the financial sector. The results of a first attempt to determine sentiments for Deutsche Bank, Commerzbank, Volksbank and Sparkasse during the introduction of account management fees in spring 2017 has shown that OMA can deliver conclusive results. Starting from measured sentiment score for each of these banks, the sentiment scores for those banks fell, which have announced the introduction of a fee for account management in April 2017. As you can see in Figure 2 sentiment scores for the Sparkasse and Volksbank ran relatively uniformly from January to March 2017. In April, the score declined due to the announcement of account management fees. One month later in May, after first account fees were reported on the account statement, the score fell significantly. For Deutsche Bank and Commerzbank such behavior couldn’t be observed, since these banks charge account fees for a long time already.

Interestingly, this result could have been achieved by the fact that a supervised learning method had to be used to improve the results of the score calculation in addition to the pre-processing techniques of NLP, such as stop word lists and tokenization. Therefore, we used a Naïve Bayes classifier at document level and trained him with several hundred tweets. To select the right tweets, we use a bag-of-word model with unigrams. As a technological platform, we used OpenNLP.

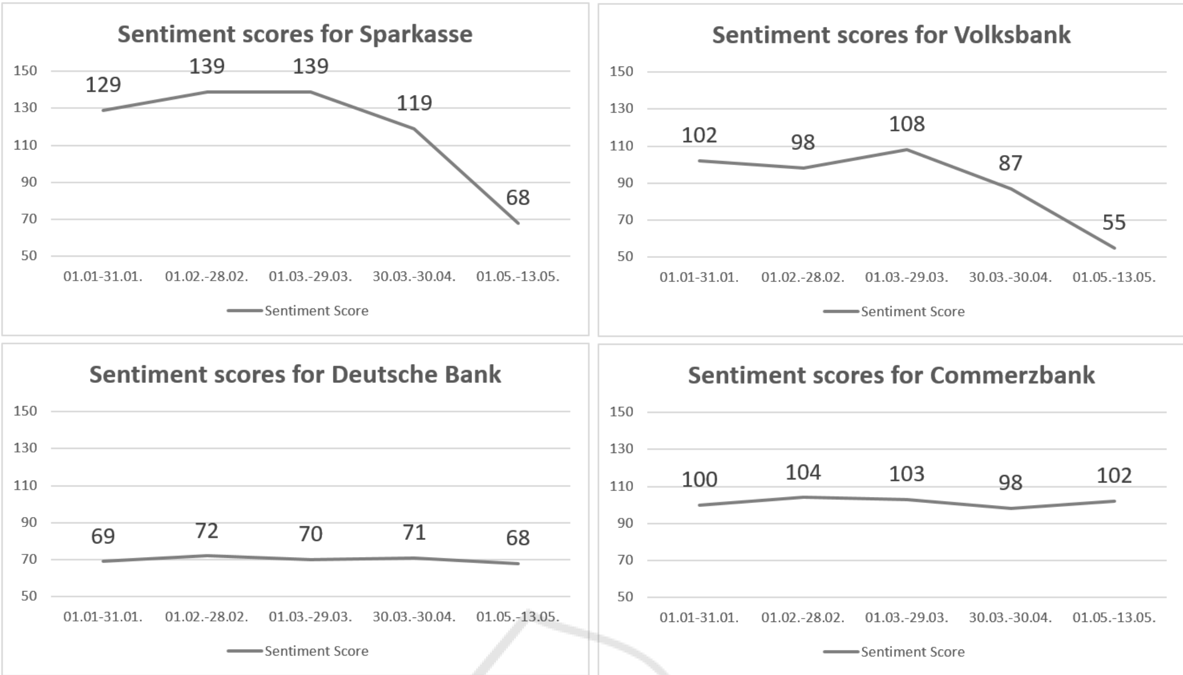

Figure 2: Sentiment scores for Sparkasse, Volksbank, Deutsche Bank and Commerzbank from January to May 2017.

5 CONCLUSION & FUTURE WORK

We have presented a methodology for opinion mining together with decision making based on machine learning, natural language processing methods for emerging opinions and weighted description logics. We were also able to present an initial evaluation showing that OMA can deliver good results.

May the approaches of opinion mining depend on specific domains, the principles underlying the ordering of opinions are to be generalized. Nevertheless, as weighted assertions are ubiquitous, one may easily envisage assertions with other content, e.g. data from IoT devices that provide incorrect values due to electronic fluctuations. The extension of OMA to data from IoT is also part for our project. From a formal perspective, we will introduce the methods mentioned in Section 2.3, such as supervised approaches with semantic features to get more information about the opinions causal nexus. Finally, we want to compare these approaches in a comprehensive evaluation and make recommendations for one or the other approach.

ACKNOWLEDGEMENTS

We would like to thank our colleague Prof. Dr. Jürgen Schenk for cooperation in the project OMA. Klemens and Heike are winners of the Dr. Karl Helmut Eberle Foundation's award on the study “Digitization and Knowledge Transformation”.

REFERENCES

Acar, E., Fink, M., Meilicke, C., Thome, C. and Stuckenschmidt, H. (2017) 'Multi-attribute Decision Making with Weighted Description Logics', *IFCoLog: Journal of Logics and its Applications, 4*, pp. 1973-1995.

Baader, F., McGuinness, D., Narci, D. and Patel-Schneider, P. (2003) *The Description Logic Handbook: Theory, Imlementation, and Applications*, Cambridge University Press.

Bhatia, P., Ji, Y. and Eisenstein, J. (2015) 'Better Document-level Sentiment Analysis from RST Discourse Parsing', EMNLP'15: Proc. 2015 Conf. on Empirical Methods in Natural Language Processing, 2212-2218.

Brody, S. and Elhadad, N. (2010) 'An Unsupervised Aspect-Sentiment Model for Online Reviews', HLT-NAACL'10: Human Language Technologies: The 2010 Annual Conf. North American Chapter of the Association for Computational Linguistics, 804–812.

Cambria, E., Schuller, B., Liu, B., Wang, H. and Havasi, C. (2013) 'Statistical Approaches to Concept-Level Sentiment Analysis', *IEEE Intelligent Systems 28*, pp. 6-9.

Carnie, A. (2010) *Constituent Structure*, Second Edition edition, Oxford University Press.

Dictionary (2002) *The American Heritage® Idioms Dictionary*, [Online], Available: www.dictionary.com/browse/opinion [1 Jul 2017].

Fishburn, P. (1969) *Utility Theory for Decision Making*, New York, London, Sydney, Toronto: John Wiley & Sons, Inc.

Ganapathibhotla, M. and Liu, B. (2008) 'Mining Opinions in Comparative Sentences', COLING'08: Proc. 22nd Intl. Conf. on Computational Linguistics, 241–248.

Gindl, S., Weichselbraun, A. and Scharl, A. (2013) 'Rule-based Opinion Target and Aspect Extraction to Acquire Affective Knowledge', WWW'13: Proc. 22nd Intl. Conf. on World Wide Web, New York, NW, 557-564.

Hahn, U. and Schnattinger, K. (1997) 'Deep Knowledge Discovery from Natural Language Texts', KDD'97: Proc. 3rd Intl. Conf. on Knowledge Discovery and Data Mining, Newport Beach, CA, 175-178.

Jindal, N. and Liu, B. (2006) 'Mining Comparative Sentences and Relations', AAAI'06: Proc. 21st National Conf. on Artificial Intelligence, 1331–1336.

Joshi, M. and Penstein-Rosé, C. (2009) 'Generalizing Dependency Features for Opinion Mining', ACL'09: Proc. 47th Annual Meeting of the Association for Computational Linguistics and the 4th Intl. Joint Conf. on Natural Language Processing, 313-316.

Kaci, S. (2011) *Working with Preferences. Less is More*, Heidelberg, Dordrecht, London, New York: Springer.

Kennedy, A. and Inkpen, D. (2006) 'Sentiment Classification of Movie Reviews Using Contextual Valence Shifters', *Computational Intelligence 22*, pp. 110-125.

Kim, S.-M. and Hovy, (2004) 'Determining the Sentiment of Opinions', COLING'04: Proc. 20th Intl. Conf. on Computational Linguistics, 1367–1373.

Lafferty, J., McCallum, A. and Pereira, F. (2001) 'Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data', ICML'01: Proc. 8th Intl. Conf. on Machine Learning, 282-289.

Lazaridou, A., Titov, I. and Sporleder, C. (2013) 'A Bayesian Model for Joint Unsupervised Induction of Sentiment, Aspect and Discourse Representations', ACL'13: Proc. 51st Annual Meeting of the Association for Computational Linguistics, 1630–1639.

Lin, C. and He, Y. (2009) 'Joint Sentiment/Topic Model for Sentiment Analysis', CIKM'09: Proc. 8th ACM Conf. on Information and Knowledge Management, New York, NY, 375-384.

Liu, B. (2012) *Sentiment Analysis and Opinion Mining*, Morgan & Claypool Publishers.

Mill, J.S. (1836) *On the Definition of political economy and on the method of philosophical investigation in that science*, London and Westminster.

Paltoglou, G. and Thelwall, M. (2010) 'A Study of Information Retrieval Weighting Schemes for Sentiment Analysis', ACL'10: Proc. 48th Annual Meeting of the Association for Computational Linguistics, 1386-1395.

Pang, B. and Lee, L. (2004) 'A Sentimental Education: Sentiment Analysis Using Subjectivity Summarization Based on Minimum Cuts', ACL'04: Proc. 42nd Annual Meeting of Association for Computational Linguistics, 271–278.

Pang, B., Lee, L. and Vaithyanathan, S. (2002) 'Thumbs up? Sentiment Classification using Machine Learning Techniques', EMNLP'02: Proc. Conf. on Empirical Methods in Natural Language Processing, 79–86.

Popescu, A.-M. (2005) 'Extracting Product Features and Opinions from Reviews', HLT'05: Proc. Conf. on Human Language Technology and Empirical Methods in Natural Language Processing, 339–346.

Qiu, G., Liu, B., Bu, J. and Chen, C. (2009) 'Expanding Domain Sentiment Lexicon through Double Propagation', IJCAI'09: Proc. 21st Intl. Joint Conf. on Artificial Intelligence, 1199–1204.

Socher, R., Bauer, J. and Manning, C. (2013) 'Parsing with Compositional Vector Grammars', ACL'13: Proc. 51st Annual Meeting on Association for Computational Linguistics, 455-465.

Sun, S., Luo, C. and Chen, J. (2017) 'A Review of Natural Language Processing Techniques for Opinion Mining Systems, Volume 36', *Information Fusion*, pp. 10-25.

Sutton, C. and McCallum, A. (2011) 'An Introduction to Conditional Random Fields', *Foundations and Trends in Machine Learning 4*, pp. 267-373.

Taboada, M., Brooke, J., Tofiloski, M., Voll, K. and Stede, M. (2011) 'Lexicon-Based Methods for Sentiment Analysis', *Computational Linguistics 37*, pp. 267-307.

Turney, P. (2002) 'Thumbs Up or Thumbs Down? Semantic Orientation Applied to Unsupervised Classification of Reviews', ACL'02: Proc. 40th Annual Meeting on Association for Computational Linguistics, 417-424.

Tversky, A. and Kahneman, D. (1981) 'The Framing of Decisions and the Psychology of Choice', *Science, Vol. 211, No. 4481*, pp. 453-458.

Yang, B. and Cardie, C. (2015) 'Context-aware Learning for Sentence-level Sentiment Analysis with Posterior Regularization', ACL'14: Proc. 42nd Annual Meeting of the Association for Computational Linguistics, 325-336.

Yu, H. and Hatzivassiloglou, V. (2003) 'Towards Answering Opinion Questions: Separating Facts from Opinions and Identifying the Polarity of Opinion Sentences', EMNLP'03: Proc. 2003 Conf. on Empirical Methods in Natural Language Processing, 129-136.

Zhang, L., Liu, B., Lim, S. and O'Brien-Strain, E. (2010) 'Extracting and Ranking Product Features in Opinion Documents', COLING'10: Proc. 23rd Intl. Conf. on Computational Linguistics, 1462–1470.