Computing Semantic Spectral Qualities using Neutral Semantic Subspace

Kazuki Otsuka
The Graduate School of Interdisciplinary Information Studies,
The University of Tokyo
7-3-1, Hongo, Bunkyo-ku, Tokyo 113-0033
kazuki-otsuka@g.ecc.u-tokyo.ac.jp

Abstract

A computational model is proposed to evaluate a non-Boolean proposition of the form "X is {A-like} to the extent σ" , where X is any concept, and {A-like} is an element of a set of ill-defined but intrinsic qualities of X. We call this type of evaluation Intrinsic Analysis (IA). To evaluate this form, a Neutral Semantic Subspace where σ becomes 0 is computed as a subspace in a high-dimensional word vector space using the geometric symmetry of opposing concepts. The results of three experiments show that this computational model (1) acts well in judging synonymy, a special case of Intrinsic Analysis (IA), (2) outperforms several existing methods in a sentiment analysis task, a subcategory of IA, and (3) is consistent with human intuition in its IA evaluation of the intrinsic descriptive meaning of a document.

1 Introduction

Is the Mona Lisa artistic? Most people would say yes, but is this really the case? Note that the term 'artistic' is used here instead of 'art'. The fact that the Mona Lisa belongs to the category of 'art' is judged as true or false. On the other hand, the judgment of whether the Mona Lisa is 'artistic' includes ambiguity and is not intrinsically evaluated with a binary classification. Now let us replace the 'Mona Lisa' with 'dresses' and the 'artistic' with the 'formal'. *Are dresses formal?* Mostly True. If we replace 'dresses' with 'shirts', it becomes more subtle. Shirts are formal to a certain extent and not formal to a certain extent. In any case, however, they are generally not as formal as dresses. Formality involves ambiguities that cannot be clearly defined.
In this paper, a computational model is proposed to a non-Boolean proposition of the form:

$$"X\ is\ \{\text{A-like}\}\ to\ the\ extent\ \sigma". \qquad (*)$$

where, X is any concept, and {A-like} is an element of a set of ill-defined but intrinsic qualities of X. This evaluation is contrasted with judgments of synonymy that have an 'is-a' type form, and is based on the assumption, as mentioned by Clark (1987), that absolute synonym do not exist. In other words, the proposition $(*)$ has the characteristic of "simultaneous decidability". We call this type of evaluation Intrinsic Analysis (IA). It is important to note that the degree to which concept X is {A-like} is not strictly defined, therefore dynamic and fuzzy. At the same time the intrinsic semantic property is possessed by the concept itself without reference to other concepts. {A-like} includes adjectives, hypernyms, synonyms and other relationships.
The goal of IA is to model a general structure of meaning by the spectrum of semantic qualities. Sentiment analysis can be viewed as a part of IA, where words related to sentiment are assigned to {A-like} in $(*)$. Topic modeling can also be regarded as a part of IA, where {A-like} in $(*)$ are words related to topics. The decision of synonymy in semantic similarity is a special case of $(*)$.
A concept similar to {A-like} has been defined by Cruse (2004) as *quality*. Difference of *quality* is roughly tested by whether we can say:

X is not A, X is B.

He listed the following as the typical set of ontological types:

THING / QUALITY / QUANTITY / PLACE / TIME / STATE / PROCESS / EVENT / ACTION / RELATION / MANNER

Is it modern? (time); Does it exist in nature? (place); Is it luxurious? (quality); Is it related to a political context? (relation); Is it formal? (manner). All of these qualities can be assigned to {A-like} in the IA form $(*)$. IA does not deal with the evaluation of sentences such as:

The man is my father.

Whether this statement is true or false, or how plausible it is, depends on a relation that is not intrinsic to the element "The man". It is sure that "The man" has a certain semantic fuzziness, but it is not intrinsic.

Let us consider a computational model that handles this IA-type problem. Firstly, the proposed computational model define a semantic quality using antisymmetry of opposing concepts. It relies on the assumption that opposing concepts have a symmetrical multi-dimensional structure that represents a spectral combinations of features of a meaning in a word vector space. Second, a subspace where $\sigma = 0$ in the proposition $(*)$ is generated. We call this subspace (i.e., a linear decision boundary separates the signs of σ) a Neutral Semantic Subspace, and the extent σ of X is given by the projection of X onto the orthogonal complement of this subspace. These technical details are described in Section 3.

Section 2 discusses the basic policies that are prerequisites for the realization of the IA model, based on existing research on semantic formalization and computational models of words. Section 3 describes the details of the theory of the computation model. In section 4, the results of three experiments show that this computational model (1) acts well in judging synonymy, a special case of IA, (2) outperforms several existing methods in a sentiment analysis task, a subcategory of IA, and (3) is consistent with human intuition in its IA evaluation of the intrinsic descriptive meaning of a document.

2 Related Work

The Formalization of Semantics. The question of what is the definition of ”meaning” and how it should be represented quantitatively is an open question. In Frege (1948)’s analytic philosophy, the reference and sense of a sign are distinguished from the associated idea which is subjective. However, by nature, subjective ambiguity is unavoidable in reasoning. Zadeh (1988) presented a logical system that tolerates ambiguity and allows statements to take continuous truth values. According to the quantum-probabilistic formalism developed by von Neumann (1955), an observable state variable is described by the spectrum of two or more states and therefore has ”simultaneous decidability”. The question of semantic representation seems to reduce to the problem of how the meanings of each concept are common, similar, or different. Thus, it is suggestive to consider the field of semantic similarity, which in recent years has moved toward greater rigor in the formalization of semantics.

Currently, semantic similarity is considered mainly for synonymy. Budanitsky and Hirst (2006) claim that similarity is only a special case of semantic relatedness and proposed the term *near-synonym*. Clark (1987), based on observations of the early stages in the acquisition of language, showed that many apparent synonyms are in fact not synonyms. For example, ’big’ and ’large’ are generally considered to be synonymous, both surely being close in the sense of taking high measurements, but simultaneously completely opposite in the sense of whether it is additive or non-additive. Therefore, the synonymy is regarded as merely the result of a particular biased situation described as a continuous spectrum of relationships. Thus, all related concepts may be simultaneously near and far in meaning, and the totality of qualities seems to determine the relationship and the definition.

From these perspectives, in this study, each concept is viewed as a set of semantic qualities, and the spectrum of the qualities defines the total meaning of the concept, or the relationships with other concepts. Each concept is in a vector space that represents semantic relatedness i.e., the superordinate concept of any linguistic category. These qualities are written as $\{$A-like$\}$ in the $(*)$ form of IA. Conversely, in theory, a concept X is represented as the union of all $\{$A-like$\}$ qualities in this model.

Computational Model of Words. According to the distributional hypothesis of Harris (1954), words sharing semantic relationships tend to occur in similar contexts. After this hypothesis was proposed, semantic vector space models learned from the distribution of words in large unlabeled text corpora have been successful in tasks including information retrieval. Latent Semantic Analysis (LSA) (Deerwester et al. 1990) generates a word-document co-occurrence matrix whose features are linear combinations of tf-idf features. Word2Vec (Mikolov et al. 2013) and GloVe (Pennington, Socher, and Manning 2014) efficiently learn statistical information from word-word co-occurrences on large datasets and estimate higher-dimensional continuous representations of words.

Whereas these models require only large documents as the training set, manually created lexical resources according to linguistic classification (Miller et al. 1990; Fellbaum 1998; Hill, Reichart, and Korhonen 2015) are used as a stand-alone evaluation set, or as external resources to be injected into and pre-trained vector models for their enrichment (Mrkšić et al. 2017; Wieting et al. 2015).

Surely the methods using lexical resources are accurate in some sense, and the injected vector models have shown state-of-the-art performance on some specific tasks. However, as discussed above, it is not practical to manually organize the semantic structure between concepts described with a myriad of combinations of continuous semantic qualities. Furthermore, words and concepts are being created and are evolving every day. Also, some semantic qualities that cannot be measured by the task may be lost due to optimization for a particular task, such as synonymy.

As the computational basis of IA model, using semantic relatedness is a reasonable choice since it encompasses subdivided complex semantic categories. And if we follow the distributional hypothesis, semantic relatedness is modeled by word co-occurrence. Therefore, in this work, in order to represent continuous spectral structures of meaning without using external lexical resources, Word2Vec and GloVe are used as the computational foundation for a co-occurrence-based semantic vector space.

3 The IA Model

3.1 Antisymmetric Concepts Matrix to Define a Semantic Quality

Suppose we wish to generate a mathematical model that evaluates the degree to which a given concept has an intended intrinsic semantic quality. From the word embedding matrix $E \in \mathbf{R}^{m,n}$, whose features are m and vocabularies are n, let us take $\frac{N}{2}$ opposing pairs that are representative for that semantic quality, and let S_0 be a new matrix whose column vectors are the $\frac{N}{2}$ selected pairs. $S_0 \in \mathbf{R}^{N\times m}$ is generally a full-rank rectangular matrix with $N < m$. To center the distribution in the vector space of S_0 at the center of the coordinates, let us consider a linear operator ψ that transforms S_0 such that $s_{ij} - \mu_i$ is the entry of matrix S after transformation, where s_{ij} is each entry of S_0 and μ_i is its row mean. The linear operator ψ is given by,

$$\psi : S_0 \mapsto S$$
$$s_{ij} - \mu_i$$
$$\psi = I - \frac{1}{N} J_N \tag{1}$$

where J_N is an N by N matrix ones and I is an N by N identity matrix.

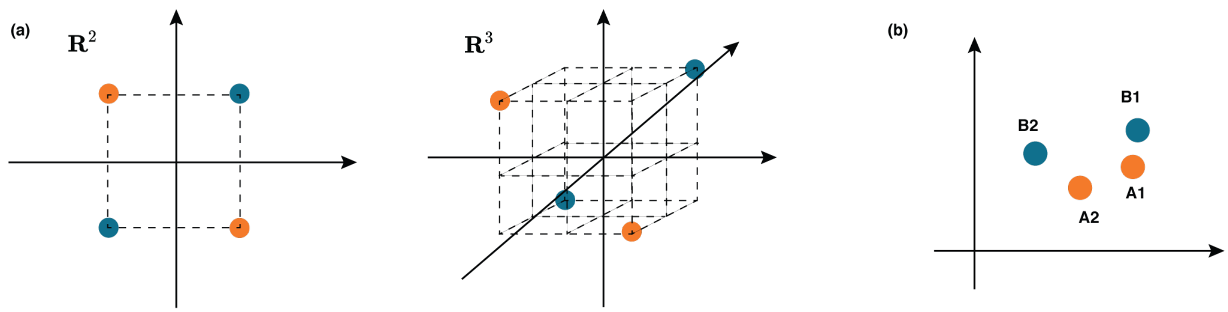

Figure 1: (a) Hypothetical example 1 - In the left figure, there is no hyperplane, i.e., straight line, in R^2 separating orange and green. In fact, the left figure is a compressed 3D space, and the vectors are placed as in the right figure. In this case, the hyperplane, i.e., plane, separating orange and green does exist. According to the experiments in Chapter 4, models with more than 10-dimensions tend to perform well. (b) Hypothetical Example 2 - The vector with the meaning close to A1 is A2. However, in cosine similarity, B1 is closer than A2. Nevertheless the meaning polarity of B1 is opposite to A1.

Here, let the transformed, centered matrix S be the *Anti-symmetric Qualities Matrix*.

3.2 Find a Neutral Semantic Subspace

When words with opposite qualities are distributed over a vector space, by its semantic linear regularity, a linear manifold that places these pairs of opposites symmetrically on one side and the other, can be considered a *Neutral Semantic Subspace*, where $\sigma = 0$ in $(*)$. For example, on a plane in $\mathbf{R}^2$, a line of $\mathbf{R}^1$ separates the semantically opposite pairs; in $\mathbf{R}^3$, a plane of $\mathbf{R}^2$ separates the opposite pairs as well symmetrically. Thus, in general, in a space of $\mathbf{R}^n$, the $\mathbf{R}^{n-1}$ linear manifold, that is hyperplane, separates opposite pairs symmetrically. Note that $n \geq 4$, since $n = (number\, of\, pairs\, chosen \times 2)$ and the at least 2 pairs are selected in the operation 3.1. Let us consider spanning the hyperplane by singular vectors of S. This is because the singular vectors especially corresponding to large singular values point in the direction of the spread of the distribution of antisymmetric pairs, so some of them can be assumed to be a boundary running through ideally the center of each pair. Therefore, we assume that the subspace spanned by $n-1$ of the n singular vectors of S is the expected linear manifold, that is the *Neutral Semantic Subspace*.

3.2.1 Singular Value Decomposition

To obtain a set of orthogonal basis vectors that reflect the spread of data in the antisymmetric pairs that constitute S and ideally run through the center of each pair, a (Reduced) Singular Value Decomposition is performed on S:

$$S = \sigma_1 \mathbf{u}_1 \mathbf{v}_1^T + \cdots + \sigma_N \mathbf{u}_N \mathbf{v}_N^T = U \Sigma V^T \tag{2}$$

where $\mathbf{v}$ represents the eigenvectors of $S^T S$, $\mathbf{u}$ represents the eigenvectors of SS^T, and the square roots of σ represent the non-zero eigenvalues of $A^T A$ and AA^T. $\mathbf{v}$ and $\mathbf{u}$ are called (left and right) singular vectors, respectively. The square roots of σ are called singular values of S. The dimensions of U, Σ, and V are

$$U \in \mathbf{R}^{N \times N}, \Sigma \in \mathbf{R}^{N \times N}, V \in \mathbf{R}^{m \times N} \tag{3}$$

U holds N singular vectors of dimension N, which are used in Section 3.2.3 to find the equations of the hyperplane, and V holds m singular vectors of dimension m, which are used for mapping a word vector on any S with m features to the N-dimensional space. As noted above, the singular vectors especially corresponding to large singular values of σ are assumed to span a subspace separating the pairs symmetrically.

3.2.2 Projection Operator

In considering a manifold on an N-dimensional space, let us prepare an operator that projects an arbitrary m-dimensional matrix (or vector) X onto an N-dimensional space, so that words in S represented by an m-dimensional feature space can be handled in the N-dimensional space.
$V' = \{\mathbf{v}_i\}$, leaving only r singular vectors from V, is a linear map that maps X to an r-dimensional space. Let V' be the projection operator ϕ and the transformed matrix (or vector) be X'.

$$\phi : XV' \mapsto X' \in \mathbf{R}^{n \times r}, \ \forall X \in \mathbf{R}^{n \times m}$$
$$V' \in \mathbf{R}^{m \times r} = \{\mathbf{v}_i\} \qquad (i = \{1, \ldots, N\}) \tag{4}$$

Note that the basis of the transformed matrix X' is $\{\mathbf{v}_i\}$ $(i = 1, \ldots, N)$. This allows any word vector $x \in \mathbf{R}^m$ of S or E to be projected onto a N-dimensional space with $\{\mathbf{v}_i\}$ (or $\{\mathbf{u}_i\}$) as the orthonormal bases.

$$\forall \mathbf{x} \in \mathbf{R}^m \quad \phi(\mathbf{x}) = \mathbf{x}' \in \mathbf{R}^N \tag{5}$$

3.2.3 Compute Neutral Semantic Subspace

Hyperplane on n-dimensional space $\mathbf{R}^n$ $\mathcal{H}$ is gehnerally represented by

$$\mathcal{H}: \ e_1 x_1 + e_2 x_2 + \cdots + e_n x_n = 0 \tag{6}$$

where,

$$\mathbf{x} = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix}$$

The coefficients $e_1, \ldots, e_n$ can be obtained as the coefficients of each term of the cross product of $n-1$ linearly independent vectors $\mathbf{v}_1, \ldots, \mathbf{v}_{n-1} \in \mathbf{R}^n$. That is, hyperplane $\mathcal{H}$ is equal to

$$\mathcal{H} : e_1 x_1 + e_2 x_2 + \cdots + e_n x_n = \mathbf{v}_1 \times \mathbf{v}_2 \times \cdots \times \mathbf{v}_{n-1}$$

$$
= \begin{vmatrix} v_{1,1} & v_{1,2} & \cdots & v_{1,n} \\ \vdots & \vdots & & \vdots \\ v_{n-1,1} & v_{n-1,2} & \cdots & v_{n-1,n} \\ x_1 & x_2 & \cdots & x_n \end{vmatrix}
$$

$$
\begin{aligned}
&= \det \begin{bmatrix} v_{1,1} & \cdots & v_{1,n} \\ \vdots & & \vdots \\ v_{n-1,1} & \cdots & v_{n-1,n} \end{bmatrix} x_1 \\
&- \det \begin{bmatrix} v_{1,1} & v_{1,3} & \cdots & u_{1,n} \\ \vdots & \vdots & & \vdots \\ v_{n-1,1} & v_{n-1,3} & \cdots & v_{n-1,n} \end{bmatrix} x_2 \\
&+ \cdots \\
&+ (-1)^n \det \begin{bmatrix} v_{1,1} & \cdots & v_{1,n-1} \\ \vdots & & \vdots \\ v_{n-1,1} & \cdots & v_{n-1,n-1} \end{bmatrix} x_N
\end{aligned} \tag{7}
$$

In the case of an N-dimensional space, $N-1$ linearly independent basis vectors $\{\mathbf{v}_i\}$, $(i = \{1, \ldots, N-1\})$ are needed. Now let us use the N singular vectors of S obtained in 3.2.1 as this basis vectors. More specifically, $\tilde{U} \in \mathbf{R}^{N \times N-1}$ is used, that is the N by N matrix U with one arbitrary singular vector removed.

Note that the index $N-1$ in $\mathbf{u}_{N-1}$ simply indicates the number of vectors $\mathbf{u}$ contained in $\tilde{U}$, and does not mean that the N-th singular vector is removed from U. The same applies to the following.

$$
\tilde{U} = \begin{bmatrix} \mathbf{u}_1 & \ldots & \mathbf{u}_{N-1} \end{bmatrix} \in R^{N \times N-1} \tag{8}
$$

Using $\tilde{U}$, equation (7) follows

$$
\mathcal{H} : e_1x_1 + e_2x_2 + \cdots + e_Nx_N = \mathbf{u}_1, \mathbf{u}_2 \times \cdots \times \mathbf{u}_{N-1}
$$

$$
= \begin{vmatrix} u_{1,1} & u_{1,2} & \cdots & u_{1,N} \\ \vdots & \vdots & & \vdots \\ u_{N-1,1} & u_{N-1,2} & \cdots & u_{N-1,N} \\ x_1 & x_2 & \cdots & x_N \end{vmatrix} \tag{9}
$$

Each coefficient e_i is computed by

$$
e_i = (-1)^i \, det \begin{bmatrix} u_{1,1} & \cdots & u_{1,i-1} & u_{1,i+1} & \cdots & u_{1,N} \\ \vdots & & \vdots & \vdots & & \vdots \\ u_{N-1,1} & \cdots & u_{N-1,i-1} & u_{N-1,i+1} & \cdots & u_{N-1,N} \end{bmatrix} \tag{10}
$$

As shown in equation (6), the coefficient vector $\mathbf{e}$ of $\mathcal{H}$ is orthogonal to any vector $\mathbf{x}$ on $\mathcal{H}$. In other words, if we denote the 1-dimensional subspace by $\mathcal{V}$ with a single vector $\mathbf{e}$ as basis, then this $\mathcal{V}$ is the orthogonal complement of the Neutral Semantic Subspace $\mathcal{H}$, where $\sigma = 0$ in $(*)$.

$$
\forall \mathbf{x} \in \mathcal{H}, \;\; \mathbf{e} \in \mathcal{V}, \;\; \mathcal{V} = \mathcal{H}^{\perp}
$$

Thus, the inner product of $\mathbf{e}$ and any vector $\mathbf{x}$

$$
\mathbf{e}^T\mathbf{x} \tag{11}
$$

is equal to the projective position of the word $\mathbf{x}$ onto the orthogonal complement space $\mathbf{e}$ of $\mathcal{H}$. That is, the magnitude of this value indicates how far away from the semantically neutral subspace the word $\mathbf{x}$ is, and its sign indicates to which side of $\mathcal{V}$, separated by the semantically neutral space $\mathcal{H}$, $\mathbf{x}$ belongs.

If $\mathbf{x}$ satisfies

$$
\mathbf{e}^T\mathbf{x} = 0 \tag{12}
$$

, then $\mathbf{x}$ lies on $\mathcal{H}$. In other words, $\mathbf{x}$ is neutral in the meaning defined here. Note that $\mathbf{e}$ is assumed to be normalized to norm 1.

Neutral Semantic Subspace $\mathcal{H}$ divides the N-dimensional space R^N into two. Let us call one side of the partition A and the other side B. The following equation determines which side of A or B a vector $\mathbf{x}$ lies on.

$$
\begin{cases} \mathbf{e}^T\mathbf{x} > 0 : \textit{On the A side divided by} \mathcal{H} \\ \mathbf{e}^T\mathbf{x} < 0 : \textit{On the B side divided by} \mathcal{H} \end{cases} \tag{13}
$$

The further away from $\mathcal{H}$ toward A, i.e., the larger $\mathbf{e}^T\mathbf{x}$ is in the positive direction, the stronger the degree of meaning A is, and the further away from B, i.e., the larger $\mathbf{e}^T\mathbf{x}$ is in the negative direction, the stronger the degree of meaning B is.

There are N ways to choose $\mathbf{u}_1, \ldots, \mathbf{u}_{N-1}$ in equation (8). Thus, this choice allows us to find $\mathcal{H}$ in N ways. The method to obtain the most reasonable Neutral Semantic Subspace out of N ways of $\mathcal{H}$ is defined as follows. Let $\mathbf{x}_A$ be a vector of words that belong to semantic category A in the Antisymmetric Quailties Matrix S and $\mathbf{x}_B$ be the vector of words in semantic category B. If the subspace $\mathcal{H}$ spanned by a certain $n-1$ basis vectors separates them correctly, then $\mathbf{e}^T\mathbf{x}_A > 0$ holds for $\mathbf{x}_A$ and $\mathbf{e}^T\mathbf{x}_B < 0$ holds for $\mathbf{x}_B$, so the desired $\mathcal{H}$ is assumed to maximize the ratio of correctly classified words out of N number of words. That is, the subspace that maximizes the following α is the Neutral Semantic Subspace.

$$
\alpha = \frac{\sum_{i=0}^{N/2} \frac{\mathbf{e}^T\mathbf{x}_{A,i}}{\|\mathbf{e}^T\mathbf{x}_{A,i}\|} + \sum_{i=0}^{N/2} \frac{-\mathbf{e}^T\mathbf{x}_{B,i}}{\|\mathbf{e}^T\mathbf{x}_{B,i}\|}}{N} \tag{14}
$$

3.3 Intrinsic Quality Analysis

3.3.1 Projection onto the Complement of Neutral Semantic Subspace

In relation to the semantic quality defined by S, let us infer to which side an arbitrary concept belongs and to what degree it belongs. As shown in Section 3.2.3, the result of the inference is equal to the projection of the word $\mathbf{x}$ onto $\mathbf{e}$, where $\mathbf{e}$ is in the orthogonal complement of $\mathcal{H}$. That is, when we denote any word vector in the word embedding matrix E by $\mathbf{d}$, by using the coefficient vector $\mathbf{e}$ of the Neutral Semantic Subspace $\mathcal{H}$. the evaluation result is obtained as the following value.

$$
\mathbf{e}^T\mathbf{d}' \tag{15}
$$

Let us call this value *Intrinsic Quality*, i.e., σ in $(*)$. Here, $\mathbf{d}'$ is a map $\mathbf{R}^m \mapsto \mathbf{R}^N$ using the projection operator ϕ defined in equation (5).

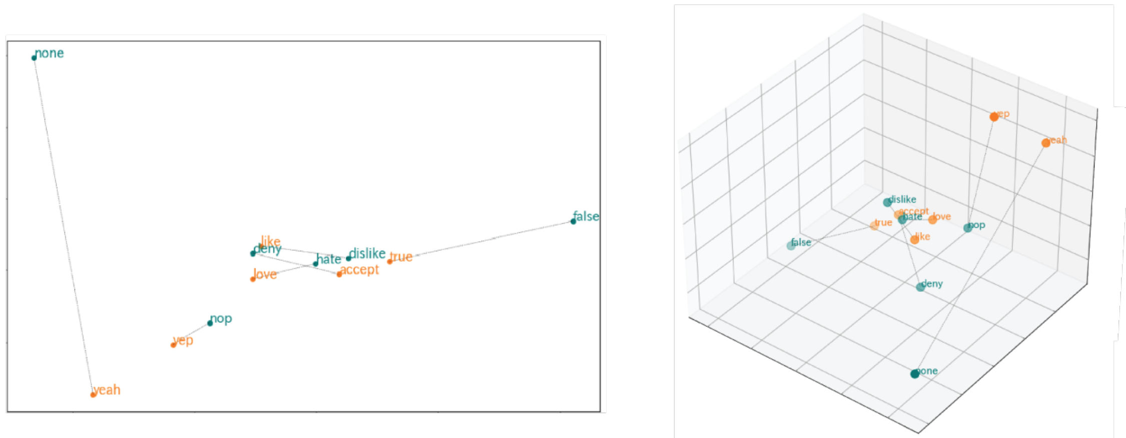

Figure 2: Six pairs used as the test set in the yes/no examination, plotted in 2- and 3-dimensional low-dimensional space. This figure shows that there is no line or plane, i.e., linear subspace, that separates all conflicting pairs, and thus a multidimensional hyperplane is required.

$$\forall \mathbf{d} \in E, \;\; \phi(\mathbf{d}) = \mathbf{d}' \in R^{N \times 1} \tag{16}$$

As described in section 3.2.3, the sign of this value indicates to which side of the space divided by the Neutral Semantic Subspace $\mathcal{H}$ the word $\mathbf{d}$ belongs. And its absolute value is equal to how far away from $\mathcal{H}$ the word is, therefore indicating the degree to which the word $\mathbf{d}$ belongs to the meaning of A or B.

$$\begin{cases} \mathbf{e}^T\mathbf{d}' > 0 : \textit{the word } \mathbf{d} \textit{ has the meaning of } A \textit{ and its degree is } \mathbf{e}^T\ \mathbf{d}' \\ \mathbf{e}^T\mathbf{d}' = 0 : \textit{the word } \mathbf{d} \textit{ has neigher } A \textit{ nor } B \textit{ meaning} \\ \mathbf{e}^T\mathbf{d}' < 0 : \textit{the word } \mathbf{d} \textit{ has the meaning of } B \textit{ and its degree is } \mathbf{e}^T\ \mathbf{d}' \end{cases} \tag{17}$$

3.3.2 Documents and Multiple Quality Estimation

When we consider a document as a set of l words $\{\mathbf{d}_i\}$, the evaluation of the document is given by the following value.

$$\sum_{i=0}^{l} \mathbf{e}^T \mathbf{d}'_i \tag{18}$$

We then represent the document by a matrix D whose columns are $\mathbf{d}_1, \ldots, \mathbf{d}_l$. With projection operator ϕ defined in equation (5), D is mapped to N-dimensional space by a single matrix operation. Let us denote this result D'.

$$D = \begin{bmatrix} - & \mathbf{d}_1^T & - \\ & \vdots & \\ - & \mathbf{d}_l^T & - \end{bmatrix} \in \mathbf{R}^{l \times m}$$

$$\phi : D \in \mathbf{R}^{l \times m} \mapsto D' \in \mathbf{R}^{l \times N}$$

The computation to simultaneously perform multiple quality estimations of a document is given by

$$D'\mathcal{M} \tag{19}$$

where M is a matrix whose columns are coefficient vectors $\{\mathbf{e}_j\}$ of multiple models generated for k semantic qualities, and the sum of each entry in row i of $D'M$ is the estimation value of the quality j of the document.

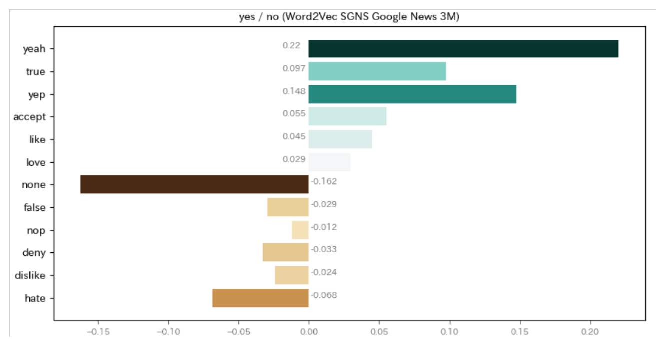

Figure 3: Intrinsic Quality computed using the IA model for six pairs of yes/no implied words.

$$\mathcal{M} = \begin{bmatrix} \mathbf{e}_1 & \ldots & \mathbf{e}_k \end{bmatrix} \in \mathbf{R}^{N \times k}$$

4 Experiments

To test the basic performance of the IA model, firstly a simple Yes/No word classification is performed. Second, a baseline comparison on a Twitter sentiment classification task using SemEval-2013 is conducted. Third, some semantic qualities of concepts described with words (film plot summaries in this case) are estimated.

4.1 Yes/No Classification

First, to verify the basic performance, the test that classifies words with simple yes and no meanings was performed. These synonyms generally fail to be classified because they are close in cosine distance (Tang et al. 2014). Three pairs of six words were selected as words similar to yes and no: set A (yes, absolutely, positive) and set B (no, never, negative) as antisymmetric pairs to generate a model. These words were then extracted from the column vectors of the embedding matrix to create an Antisymmetric Concepts Matrix,

word pairs used to calculate models	embedding	Macro F1	recall	precision	F1
A(good, nice, awesome)	Word2Vec	68.65	74.75	80.60	77.56
B(bad, ugly, horrible)	GloVe	70.13	81.53	78.19	78.83
A(good, nice, awesome, **happy**)	Word2Vec	73.38	94.43	75.19	83.72
B(bad, ugly, horrible, **sad**)	GloVe	73.82	81.60	**82.16**	81.88
A(good, nice, awesome, happy, **love**, **fun**)	Word2Vec	73.03	96.81	74.00	83.88
B(bad, ugly, horrible, sad, **hate**, **boring**)	GloVe	**75.00**	**99.12**	74.68	85.18
A(good, nice, **beautiful**, awesome, happy, love, fun)	Word2Vec	74.11	97.28	74.67	84.49
B(bad, **unpleasant**, ugly, horrible, sad, hate, boring)	GloVe	**77.71**	90.36	**81.06**	**85.46**

Table 1: Performance results of positive/negative classification on the SemEval-2013 test dataset for different word pairs and embeddings used to compute the Neutral Semantic Subspace.

	Positive	Negative	Neutral
Train	2,642	994	3,4346
Test	1,570	601	1,639
Dev	408	219	493

Table 2: Statistics of the SemEval-2013 Twitter sentiment analysis dataset.

Method	Macro F1
DistSuper + unigram	61.74
DistSuper + uni/bi/tri-gram	63.84
SVM + unigram	74.50
SVM + uni/bi/tri-gram	75.06
RAE	75.12
IA model 14-dim + GloVe	**77.71**

Table 3: Comparison of Macro F1 with existing methods for positive/negative classification of tweets.

and a 6-dimensional Neutral Semantic Subspace was computed based on the method described in Chapter 3. The test sets for A and B were selected as A' (yeah, true, yep, accept, like, love) and B' (none, false, nop, deny, dislike, hate), respectively, and the Intrinsic Quality defined by equation (15) was obtained for each word.

Figure 2 shows the 2D and 3D mappings of the word vectors in the six test sets A' and B'. The word embedding matrix used was Word2Vec SGNS trained on the 3 million word Google News dataset. Figure 3 shows the results of the Intrinsic Quality for the same test set. The results show that yeah/none, which have almost equal meanings to yes/no, are the best related. All other words belonging to A' have positive values, and all words belonging to B' have negative values. Similar results were obtained using GloVe 1.2 trained on the 2.2 million word Common Crawl dataset under the same conditions.

The results of the Intrinsic Quality values vary depending on how the antisymmetric pairs are chosen. The choice of the best pairs is currently left to heuristics. In other words, the choice of base terms can be flexible depending on the purpose.

	Label positive	Label negative	Total
Predict positive	1,331	331	1,642
Predict negative	142	248	390
Total	1,473	559	2,032

Table 4: Classification results for the model with the best Macro F1 among 8 models showed in Table2

4.2 SemEval-2013 Sentiment Analysis Task

Next, experiments on the Twitter sentiment classification benchmark dataset of SemEval-2013 (Nakov et al. 2013) were conducted. Since the purpose of the IA model is to evaluate the degree of positivity/negativity of σ in $(*)$ as a continuous value, the investigation was conducted using a 2-class classification of Positive/Negative, excluding Neutral. Thus, the test set that contains 2,032 Positive/Negative texts was used. A model was generated by selecting a few antisymmetric pairs that were close to the Positive/Negative meaning, and the Intrinsic Quality of each Tweet text was calculated according to equation (18). Words that were not included in the vocabulary of the embedding matrix were ignored. 174 stop words were removed in advance. Table 1 shows the results of Macro F1 and other performance indicators of classification performance on the SemEval-2013 test data.

The same pretrained Word2Vec and GloVe models as in 4.1 were used for comparison. In both cases, performance was improved by increasing the number of antisymmetric pairs, but beyond seven pairs, it became difficult to find additional words that would improve performance further. This does not imply that there is no possibility to improve performance beyond a $8 \times 2 = 16$ dimensional model, although it does mean that simply increasing the number of pairs does not always produce better results. It suggests that the optimal design will vary depending on how the pairs are chosen and the nature of the target task. Further discussions of the optimal selection method require additional research. Finally, Macro F1 reached 77.71% with the model generated from GloVe using 7 pairs. This exceeds the performance of DistSuper, SVM and RAE (Tang et al. 2014). Table 3 com-

Semantic	1		2		3	
War	Sin City	**-1.49**	Monster-in-Law	-1.43	Sideways	-1.15
Peace	50 First Dates	**0.36**	Coach Carter	0.34	The Longest Yard	0.32
Modern	Napoleon Dynamite	**1.53**	S.W.A.T	1.51	Coach Carter	1.29
Classic	Bewitched	**-0.48**	Finding Neverland	-0.4	Ray	-0.14
Art	-	-	-	-	-	-
Pop	Ray	**1.35**	50 First Dates	1.26	The Hitchhiker's Guide to the Galaxy	1.25

Table 5: The films with the highest 3 values for each definition of quality based on the analysis of their plot summaries.

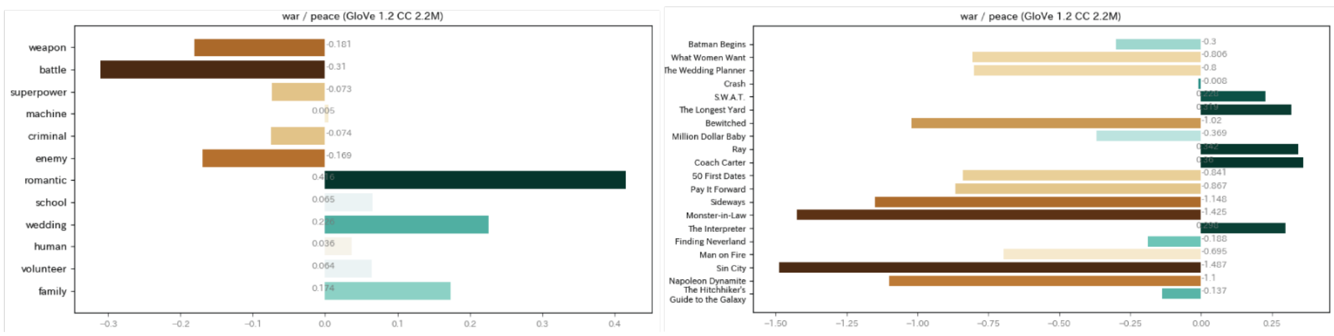

Figure 4: An Intrinsic Quality {(-war/+peace)-like} computed for an arbitrary selected test set of concepts (left) and for the 20 film plot summaries (right) using the IA model.

pares Macro F1 score with each method on the Positive/Negative classification task to SemEval-2013 using the scores presented in (Tang et al. 2014).

The crucial difference between other models and this model is that, while other models use the same kind of data for training as they use for inference and therefore are optimized for the task, this method does not use any of the same kind of data for training as it uses for inference.

4.3 Intrinsic Analysis of Concepts Described with Words

Finally, an evaluation experiment were conducted on plot summaries of several films to analyze Intrinsic Qualities of concepts described with words. Here, experiments were conducted on three defined meanings using the 20 films with the highest number of ratings in the Netflix Prize Data Set (Funk 2006). All the plot summaries are listed in Technical Appendices.

Three meanings were defined for this experiment: (1) war/-peace, (2) modern/classic, (3) art/pop. (1) war/peace was modeled by the word group A (war, fight, violence) and B (peace, love, romance). (2) modern/classic was modeled by the group A (futuristic, innovative, modern) and B (classical, old-fashioned, traditional). (3) art/pop was modeled by the group A (art, documentary, academic) and B (entertainment, pop, commercial).

Table 7 shows the top three for each semantic definition and their Intrinsic Qualities. As a whole, the results are consistent with what most people are likely to imagine. Interestingly, not a single one of the 20 works was judged to be "art-like".

5 Conclusion

This study presented one possibility for a general structure of meaning and a computation model to analyze that structure. In this model, a concept was represented as a spectrum of continuous qualities or a union set of them with simultaneous decidability. A non-Boolean propositional form that evaluates individual semantic {A-like} properties was called IA (Intrinsic Analysis), and the IA model was presented as a computational model for it. Since there are a myriad of elements that make up the meaning of concepts, this paper only discussed and verified the basic ideas and demonstrates the potential of the computation model. The question of what meaning is in the quantization of human thought and reasoning is an essential one that has not yet been clearly answered. It is believed that the exploration of this question will eventually open up a new relationship between humans and machines.

References

Barsalou, L. W. 1989. *Intraconcept similarity and its implications for interconcept similarity*, 76–121. Cambridge University Press.

Blei, D. M.; Ng, A. Y.; and Jordan, M. I. 2003. Latent dirichlet allocation. *J. Mach. Learn. Res.*, 3: 993–1022.

Budanitsky, A.; and Hirst, G. 2006. Evaluating WordNet-based Measures of Lexical Semantic Relatedness. *Computational Linguistics*, 32(1): 13–47.

Clark. 1987. The Principle of contrast : A constraint on language acquisition. *Mechanisms of language acquisition : The 20th annual Carnegie Symposium on Cognition*.

Cohen, D. W. 1989. *An Introduction to Hilbert Space and Quantum Logic*. 1. New York, NY: Springer New York, NY.

Cruse, D. A. 2004. *Meaning in language: An introduction to semantics and pragmatics*. Oxford: Oxford University Press.

Dagan, I.; and Itai, A. 1994. Word Sense Disambiguation Using a Second Language Monolingual Corpus. *Comput. Linguist.*, 20(4): 563–596.

Deerwester, S.; Dumais, S.; Furnas, G.; Landauer, T.; and Harshman, R. 1990. Indexing by latent semantic analysis. *Journal of the American Society for Information Science 41*, 391–407.

Edmonds, P.; and Hirst, G. 2002. Near-Synonymy and Lexical Choice. *Computational Linguistics*, 28(2): 105–144.

Fellbaum, C. D. 1998. A Semantic Network of English: The Mother of All WordNets. *Computers and the Humanities*, 32: 209–220.

Frege, G. 1948. Sense and Reference. *Philosophical Review*, 57(3): 209–230.

Frege, G. 1951. On Concept and Object. *Mind*, 60(238): 168–180.

Funk, S. 2006. Netflix update: Try this at home.

Griffiths, T. L.; Steyvers, M.; and Tenenbaum, J. B. 2007. Topics in semantic representation. *Psychological Review*, 114(2): 211–244.

Harris, Z. 1954. Distributional structure. *Word*, 10(2-3): 146–162.

Hill, F.; Reichart, R.; and Korhonen, A. 2015. SimLex-999: Evaluating Semantic Models With (Genuine) Similarity Estimation. *Computational Linguistics*, 41(4): 665–695.

Mikolov, T.; Chen, K.; Corrado, G.; and Dean, J. 2013. Efficient Estimation of Word Representations in Vector Space. In Bengio, Y.; and LeCun, Y., eds., *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*.

Miller, G. A.; Beckwith, R.; Fellbaum, C.; Gross, D.; and Miller, K. J. 1990. Introduction to WordNet: An On-line Lexical Database*. *International Journal of Lexicography*, 3(4): 235–244.

Morris, J.; and Hirst, G. 2004. Non-Classical Lexical Semantic Relations. In *Proceedings of the Computational Lexical Semantics Workshop at HLT-NAACL 2004*, 46–51. Boston, Massachusetts, USA: Association for Computational Linguistics.

Mrkšić, N.; Vulić, I.; Ó Séaghdha, D.; Leviant, I.; Reichart, R.; Gašić, M.; Korhonen, A.; and Young, S. 2017. Semantic Specialization of Distributional Word Vector Spaces using Monolingual and Cross-Lingual Constraints. *Transactions of the Association for Computational Linguistics*, 5: 309–324.

Nakov, P.; Rosenthal, S.; Kozareva, Z.; Stoyanov, V.; Ritter, A.; and Wilson, T. 2013. SemEval-2013 Task 2: Sentiment Analysis in Twitter. In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, 312–320. Atlanta, Georgia, USA: Association for Computational Linguistics.

Pennington, J.; Socher, R.; and Manning, C. 2014. GloVe: Global Vectors for Word Representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1532–1543. Doha, Qatar: Association for Computational Linguistics.

Tang, D.; Wei, F.; Yang, N.; Zhou, M.; Liu, T.; and Qin, B. 2014. Learning Sentiment-Specific Word Embedding for Twitter Sentiment Classification. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1555–1565. Baltimore, Maryland: Association for Computational Linguistics.

von Neumann, J. 1955. *1932, Mathematische Grundlagen der Quantenmechanik, Berlin: Springer-Verlag; English translation: Mathematical Foundations of Quantum Mechanics*. Princeton: Princeton University Press.

Weeds, J.; and Weir, D. 2003. A General Framework for Distributional Similarity. In *Proceedings of the 2003 Conference on Empirical Methods in Natural Language Processing*, 81–88.

Wieting, J.; Bansal, M.; Gimpel, K.; and Livescu, K. 2015. From Paraphrase Database to Compositional Paraphrase Model and Back. *Transactions of the Association for Computational Linguistics*, 3: 345–358.

Wilce, A. 2021. Quantum Logic and Probability Theory. In Zalta, E. N., ed., *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Fall 2021 edition.

Zadeh, L. 1971. Quantitative fuzzy semantics. *Information Sciences*, 3(2): 159–176.

Zadeh, L. 1988. Fuzzy logic. *Computer*, 21(4): 83–93.

A Technical Appendices

A.1 List of Stopwords Used in All Experiments

a, about, above, after, again, against, all, am, an, and, any, are, arent, as, at, be, because, been, before, being, below, between, both, but, by, cant, cannot, could, couldnt, did, didnt, do, does, doesnt, doing, dont, down, during, each, few, for, from, further, had, hadnt, has, hasnt, have, havent, having, he, hed, hell, hes, her, here, heres, hers, herself, him, himself, his, how, hows, i, id, ill, im, ive, if, in, into, is, isnt, it, its, its, itself, lets, me, more, most, mustnt, my, myself, no, nor, not, of, off, on, once, only, or, other, ought, our, ours, ourselves, out, over, own, same, shant, she, shed, shell, shes, should, shouldnt, so, some, such, than, that, thats, the, their, theirs, them, themselves, then, there, theres, these, they, theyd, theyll, theyre, theyve, this, those, through, to, too, under, until, up, very, was, wasnt, we, wed, well, were, weve, were, werent, what, whats, when, whens, where, wheres, which, while, who, whos, whom, why, whys, with, wont, would, wouldnt, you, youd, youll, youre, youve, your, yours, yourself, yourselves

A.2 SemEval-2013 Sentiment Analysis Task Results of 8 Models Discribed in Chapter 4.2

Antisymmetric concepts used to compute a model 1:
A (good, nice, awesome)
B (bad, ugly, horrible)

1-1. Word2Vec

	Label positive	Label negative	Total
Predict positive	**1,101**	265	1,366
Predict negative	372	**294**	666
Total	1473	559	2,032

1-2. GloVe

	Label positive	Label negative	Total
Predict positive	**1,201**	335	1,536
Predict negative	272	**224**	496
Total	1,473	559	2,032

Antisymmetric concepts used to compute a model 2:
A (good, nice, awesome, happy)
B (bad, ugly, horrible, sad)

2-1. Word2Vec

	Label positive	Label negative	Total
Predict positive	**1,201**	335	1,536
Predict negative	272	**224**	496
Total	1,473	559	2,032

2-2. GloVe

	Label positive	Label negative	Total
Predict positive	**1,202**	261	1,463
Predict negative	271	**298**	569
Total	1,473	559	2,032

Antisymmetric concepts used to compute a model 3:
A (good, nice, awesome, happy, love, fun)
B (bad, ugly, horrible, sad, hate, boring)

3-1. Word2Vec

	Label positive	Label negative	Total
Predict positive	**1,426**	501	1,927
Predict negative	47	**58**	105
Total	1,473	559	2,032

3-2. GloVe

	Label positive	Label negative	Total
Predict positive	**1,460**	495	1,955
Predict negative	13	**64**	77
Total	1,473	559	2,032

Antisymmetric concepts used to compute a model 4:
A (good, nice, beautiful, awesome, happy, love, fun)
B (bad, unpleasant, ugly, horrible, sad, hate, boring)

4-1. Word2Vec

	Label positive	Label negative	Total
Predict positive	**1,433**	486	1,919
Predict negative	40	**73**	113
Total	1,473	559	2,032

4-2. GloVe

	Label positive	Label negative	Total
Predict positive	**1,331**	311	1,642
Predict negative	142	**248**	390
Total	1,473	559	2,032

A.3 Yes/No Classification Results

Word2Vec SGNS Google News 3M

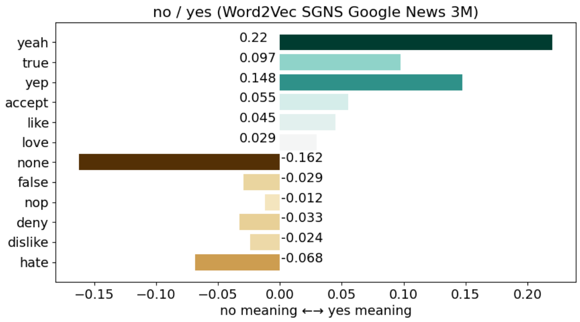

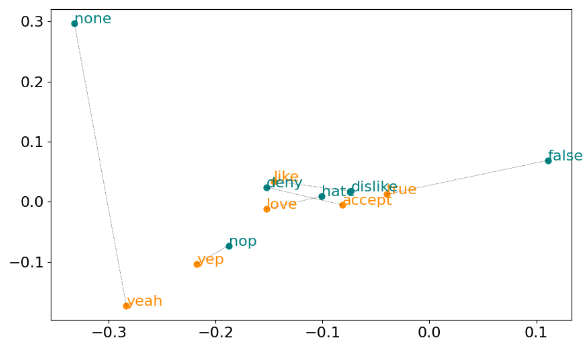

GloVe 1.2 CC 2.2M

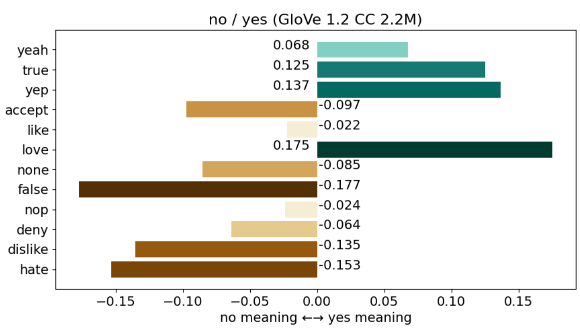

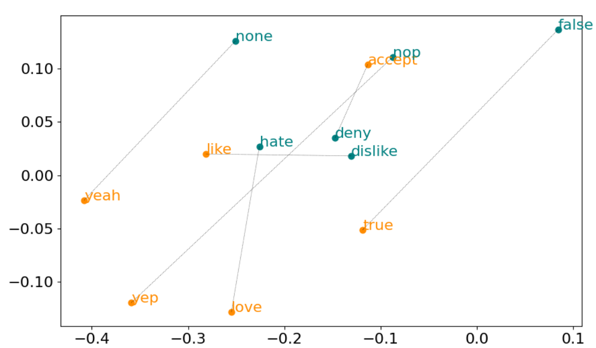

A.4 Titles and Plot Summaries of Films used in Chapter 4.3

1. Batman Begins: As a toxic threat endangers a corrupt city, Bruce Wayne finds himself at odds with a assassins and forced to battle more than his own demons.

2. What Women Want: When a sexist advertising executive is suddenly able to hear women's thoughts, he's not nearly as charming as he thinks he is.

3. The Wedding Planner": Wedding planner Mary Fiore is saved from an accident by the man of her dreams – only to discover that he happens to be her latest client's fiance.

4. Crash: In post-Sept. 11 Los Angeles, tensions erupt when the lives of people from all walks of life converge during a 36-hour period

5. S.W.A.T.: A veteran cop is tasked with drafting and training and training a special weapons and tactics team, who soon find themselves up against an international criminal.

6. The Longest Yard: While doing time, a professional quarterback persuades a fellow convict and former coach to prepare a group of inmates for a game against the guards.

7. Bewitched: Isabel Bigelow seems to be the perfect Samantha to star in a remake of the 1960s sitcom 'Bewitched' – but no one knows she really is a witch!

8. Million Dollar Baby: When a cantankerous trainer mentors a persistent amateur boxer determined to go pro, deep- seated emotions become their strongest opponents.

9. Ray: From a satire to a psychological thriller, four short stories from celebrated auteur and writer Satyajit Ray are adapted for the screen in this series.

10. Coach Carter: When he takes over as Richmond High School's new basketball coach, Ken Carter demands that players show up academically as well as athletically.

11. 50 First Dates: After falling for an art teacher with short-term memory loss, a veterinarian finds he must win her over again every single day.

12. Pay It Forward: A young boy attempts to make the world a better place after his teacher gives him that chance.

13. Sideways: Two men reaching middle age with not much to show but disappointment embark on a week-long road trip through California's wine country, just as one is about to take a trip down the aisle.

14. Monster-in-Law: The love life of Charlotte is reduced to an endless string of disastrous blind dates, until she

meets the perfect man, Kevin. Unfortunately, his merciless mother will do anything to destroy their relationship.

15. The Interpreter: An interpreter for the United Nations finds herself in danger after she overhears an assassination plot and turns to a skeptical federal agent for help.

16. Finding Neverland: The story of Sir J.M. Barrie's friendship with a family who inspired him to create Peter Pan.

17. Man on Fire: A jaded ex-CIA operative reluctantly accepts a job as the body guard for a 10-year-old girl in Mexico City and will stop at nothing when she's kidnapped.

18. Sin City: A busy couple tries to give their love life a boost by taking an impromptu weekend trip only to find their relationship tested in unexpected ways.

19. Napoleon Dynamite: A listless and alienated teenager decides to help his new friend win the class presidency in their small western high school, while he must deal with his bizarre family life back home.

20. The Hitchhiker's Guide to the Galaxy, Mere seconds before the Earth is to be demolished by an alien construction crew, journeyman Arthur Dent is swept off the planet by his friend Ford Prefect, a researcher penning a new edition of "The Hitchhiker's Guide to the Galaxy".

A.5 Results of Intrinsic Analysis Described in Chapter 4.3

1. War/Peace

1-1. Related concepts

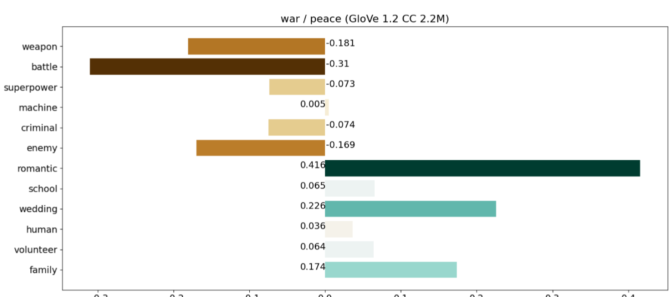

1-2. Concepts described with words (film plot summaries)

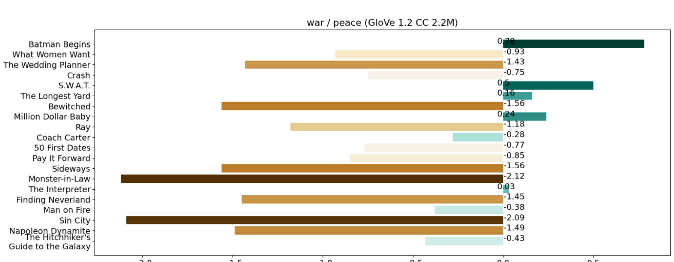

2. Modern/Classic

2-1. Related concepts

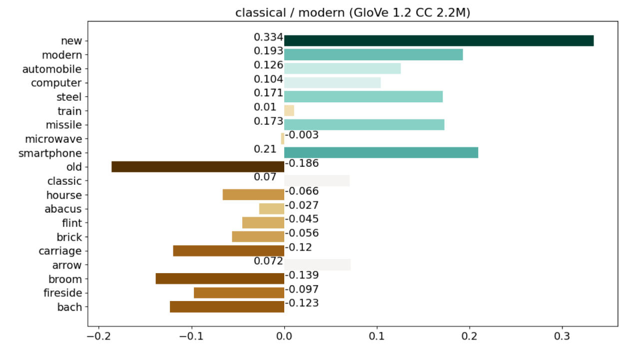

2-2. Concepts described with words (film plot summaries)

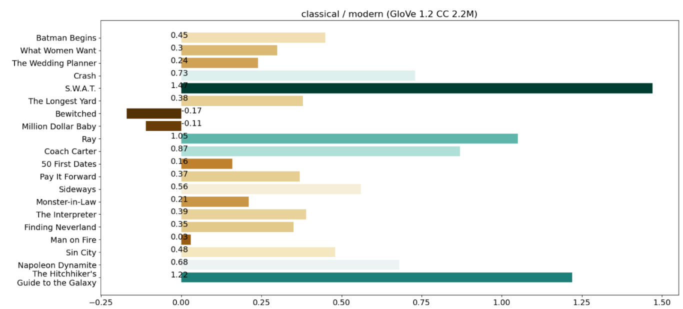

3. Art/Pop

3-1. Related concepts

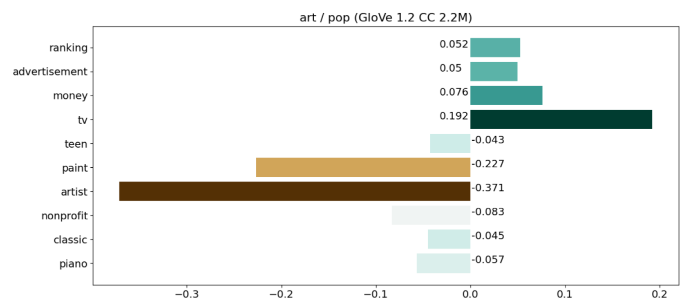

3-2. Concepts described with words (film plot summaries)