

# A Clique-based Graphical Approach to Detect Interpretable Adjectival Senses in Hungarian

Enikő Héja and Noémi Ligeti-Nagy

Hungarian Research Centre for Linguistics

Institute of Language Technology and Applied Sciences

Language Technology Research Group

[heja.eniko, ligeti-nagy.noemi]@nytud.hu

## Abstract

The present paper introduces an ongoing research which aims to detect interpretable adjectival senses from monolingual corpora applying an unsupervised WSI approach. According to our expectations the findings of our investigation are going to contribute to the work of lexicographers, linguists and also facilitate the creation of benchmarks with semantic information for the NLP community. For doing so, we set up four criteria to distinguish between senses. We experiment with a graphical approach to model our criteria and then perform a detailed, linguistically motivated manual evaluation of the results.

## 1 Introduction

The objective of this ongoing research is to model human intuition regarding meaning distinctions, and anchor it to observable data. Its importance is given by the fact that according to several authors (eg. Véronis, 2003; Kuti et al., 2010) human intuition on sense distinctions varies greatly among individuals, which in turn has a serious effect on lexicography, lexical semantics and NLP, as well. It goes without saying in the lexicographic community that in spite of scrupulous corpus-based investigations, monolingual dictionaries greatly vary with regard to their macro- and microstructure (Adamska-Sałaciak, 2006). The same problem arises in the field of NLP: the sense inventories or knowledge-bases exhibit a great variance regarding how fine-grained meaning distinctions they apply. Although lexical semantics in linguistics and word sense induction in NLP are widely studied fields (cf. Geeraerts, 2015; Amrami and Goldberg, 2019; Wiedemann et al., 2019), to our knowledge there is still no agreement on how the meaning space of a word should be partitioned to obtain well-motivated senses. For instance, Pustejovsky (1995, p. 32) introduces a very fine-grained meaning distinction asserting that “adjectives such as *good* have multiple

meanings depending on what they are modifying: *good car*, *good meal*, *good knife*”. However, he also adds that *good* may be conceived of merely “as a positive evaluation of the nominal head it is modifying.” Accordingly, the present experiment has two main objectives: first, we aim to come up with a definition that is able to provide necessary criteria to distinguish between senses. This definition should enable us to anchor meaning distinctions to not only a set of contexts, but conceptual categories as well. Secondly, we aim to model this definition via an unsupervised approach that is able to grasp this definition to minimize the role of human introspection in meaning distinction. We think that our approach is quite promising as one of the main drawbacks of unsupervised models is their poor interpretability, as pointed out by Camacho-Collados and Pilehvar (2018). On top of that, in their survey they tied graphical models to knowledge-based semantic representations, which implies that unsupervised graph-based WSI is underrepresented in the field.

The usual conception of meaning starts from meaning identity: the definition of synonymy (two expressions are synonymous iff they are interchangeable in every context preserving the original meaning) has a long tradition going back at least to Frege (1892), and all the senses that are not synonyms are considered to be different senses. The subsequent research tends to accept this chain of thoughts. However, in the present discussion we put it in the other way: as opposed to Frege and his followers, we do not give a definition for synonymy, but give one to distinguish between meanings. This choice is motivated by the fact that the notion of synonymy is intimately tied to truth-conditions, which are notoriously missing from pure distributional semantics. That is why it is so hard to detect true synonyms solely on distributional grounds. And indeed, automatically detected synonym-classes tend to cover also tight seman-

tic classes, such as names of nations, colors, even antonyms exhibiting very similar distributional behavior. Starting from the presupposition that attributive adjectives can be characterized in a rather simple feature space – constituted only by the following nouns – in the present research we confine ourselves to the investigation of the semantic properties of attributive adjectives. The paper is structured as follows: in section 2 our hypotheses are presented, section 3 describes our methodology, in section 4 we present our validation techniques, while section 5 focuses on the evaluation of our results. We conclude with a summary in section 6.

## 2 Criteria for meaning distinction

In what follows, we describe the applied criteria, which were implemented in the next phase. Contrary to the usual procedure of definition, instead of searching an identity criteria to “give the necessary and sufficient conditions for  $a$  to be identical to  $b$  when  $a$  and  $b$  are  $K$ s” (cf. Carrara and Giaretta, 2004), we search for necessary and sufficient conditions to discriminate between  $a$  and  $b$ . That is, instead of modeling synonymy, we strive to grasp when the target word surely conveys different meanings on distributional grounds. For doing so, we introduce the notion of near-synonymy (cf. Ploux and Victorri, 1998) – a relaxed version of synonymy: two words are *near-synonyms* if they are interchangeable in a restricted set of contexts so that they preserve the meaning of the original sentence.<sup>1</sup> Moreover, in accordance with our original purpose (i.e. meaning distinction), we also consider the members of tight semantic classes to be near-synonyms, inasmuch various tight semantic classes denote different senses of a word, even though they do not preserve the truth value.<sup>2</sup>

According to our hypothesis two senses have to be differentiated iff:

1. There is (at least) one near-synonym for each sense of the adjective.
2. There is a set of context-nouns which form grammatical constructions with both the original adjective and with the near-synonym.
3. The two sets of context-nouns characterizing the different senses are non-overlapping sets.

<sup>1</sup>For instance, *finom* (‘fine’) and *lágy* (‘soft’) are synonyms before nouns related to music, such as the Hungarian counterparts of ‘music’, ‘rhythm’, ‘melody’, etc.

<sup>2</sup>For example, *fekete* (‘black’) may belong to two different near-synonymy sets: one containing surnames and the other containing names of colors.

4. The non-overlapping set of nouns form a semantic category “reflecting the sub-selectional properties of adjectives” (Pustejovsky, 1995).

Example 1 is intended to further illustrate the above criteria, using the automatically extracted two senses of the adjective *napfényes* (‘sunny’). As can be seen, there is a near-synonym for both senses: *napsütétes* (‘sunshiny’) for the first one and *napsütötte* (‘sunlit’) for the second one. The nouns listed below the adjectives are the ones that form grammatical constructions with the near-synonyms: *napfényes/napsütétes vasárnap* (‘sunny/sunshiny Sunday’), *napfényes/napsütétes nap* (‘sunny/sunshiny day’), etc., and *napfényes/napsütötte terület* (‘sunny/sunlit area’), *napfényes/napsütötte terasz* (‘sunny/sunlit terrace’), etc. However, the two sets of nouns do not overlap: there is no *napsütétes terasz* (‘sunshiny terrace’) or *napsütötte nap* (‘sunlit day’), and the same goes for all adjective-noun pairs where the noun comes from the context noun set of the other sense. Finally, the nouns that match the above criteria form a semantic category: time periods with the first sense, and areas, places with the second.

- (1) Sense 1: *napfényes* ‘sunny’, *napsütétes* ‘sunshiny’  
Nouns of sense 1: *vasárnap* ‘Sunday’, *nap* ‘day’  
Sense 2: *napfényes* ‘sunny’, *napsütötte* ‘sunlit’  
Nouns of sense 2: *terület* ‘area’, *sziget* ‘island’, *oldal* ‘side’, *terasz* ‘terrace’

We wish to examine to what extent the above conditions are necessary and sufficient to differentiate between meanings. For doing so, in Section 3 an unsupervised word sense induction experiment on Hungarian monolingual data will be described using cliques of target words and their contexts to retrieve senses. The workflow conceptually comprises two main stages: i) the detection of near-synonymy classes for a given adjective, ii) discriminating between the various meanings of the given adjective by the extraction of the relevant context nouns.

## 3 Method

Our methodology is based on Ah-Pine and Jacquet (2009), as far as meaning distinctions are mod-

eled via cliques. However, there are two main differences: first, instead of named entities we focused on adjectival meanings. As overproduction of cliques is much less pronounced in this case, clustering becomes an unnecessary step. However, the resulting cliques need to be validated in terms of the following nouns, possibly along with the sub-categorization patterns of the adjectives. Secondly, adjectives are represented with static dense embeddings instead of frequency based sparse vectors.

### 3.1 Input data

The adjectives of our interest were selected on the basis of the 180 million word Hungarian National Corpus (Váradi, 2002). Although the frequency list contains adjectives with various case suffixes, we took only nominative adjectives into consideration, presuming that the adjective is always in nominative in the Adj + Noun constructions.

## 3.2 Representations

### 3.2.1 Representation of adjectives

As opposed to Ah-Pine and Jacquet (2009), instead of count vectors we decided to use static word embeddings to represent adjectives. Our choice was motivated by Baroni et al. (2014), who presented a systematic comparison of traditional “context-counting” vectors (eg. Turney and Pantel, 2010; Clark, 2015) and the more recent “context-predicting” ones (eg. Bengio et al., 2003; Mikolov et al., 2013a) on a set of various standard lexical semantic benchmarks. Their findings show that the predictive models achieve an impressive overall performance, beating count vectors in all tasks. Therefore, a word2vec language model (Mikolov et al., 2013a,b) was trained on the first 999 file (21GB raw texts) of a Hungarian language corpus, the Webcorpus 2.0 (Nemeskey, 2020) containing the normalized version of the original texts, cc. 170M sentences<sup>3</sup>. 300-dimension vectors were trained using the Gensim Python package (Rehurek and Sojka, 2011) to perform CBoW training with a  $6k$  window size and a minimum frequency of 3. Since Hungarian is a highly inflective language and we trained embeddings on raw texts, this is not a pure bag-of-words model, as the abbreviation CBoW would imply. Our choice of input data

<sup>3</sup>As the evaluation of the model trained on the whole Webcorpus 2.0 (cc. 591.4M sentences) yielded only a slight improvement on the Hungarian translation of the Google Analogy Test Set (Makrai, 2015), this smaller model was used in our experiment.

was based on the presupposition that morphosyntactic information may contribute to the characterization of adjectival meanings. This hypothesis is in accordance with the findings of Novák and Novák (2018), who investigated the performance of various Hungarian static word embeddings in a word similarity task. Their experiment concludes that adjectival senses are best represented via embeddings trained on surface forms of words. Roughly 8,5M word forms were assigned embeddings as the result of our training. The trained LMs are available on GitHub: [https://github.com/nytud/w2v\\_models](https://github.com/nytud/w2v_models).

### 3.2.2 Representation of semantic similarity

In our graph-based representation of adjectives, vertex-labeled undirected graphs were generated. Vertices and their labels represent the adjectives, while the edges (or their lack) denote whether there is a semantic similarity relation between two adjectives (or not). This structure encodes some basic intuitions about meaning similarity:

- (1) ‘Undirectedness’ guarantees the symmetric nature of meaning similarity: if a meaning  $M$  is similar to meaning  $M'$ , then the reverse is also true.
- (2) Since every adjective is similar to itself, there is a self-loop at every node of the graph.

### 3.2.3 Representing near-synonyms as cliques

Meaning is grasped through the notion of near-synonymy. Following Ah-Pine and Jacquet (2009), near-synonyms which exhibit “very similar” distributional behavior, are grasped by cliques in the graph: that is, we search for those maximally connected subgraphs. Now the nodes in the clique represent a set of adjectives with “very similar” distributional behavior.

### 3.2.4 Representing meaning-discrimination as shared cliques

This approach, on the one hand, makes possible the detection of multiple near-synonymy classes comprising a common adjectival lexeme, where the corresponding cliques represent differing sense candidates. In addition, ideally, it also enables meaning discrimination based on explicit surface data, inasmuch all the resulting cliques are anchored to the contexts in which each element of the adjectival clique may occur.

Therefore, according to our hypothesis, an adjective has multiple meanings if it belongs to multiple

cliques, and the cliques are characterized by non-overlapping sets of context nouns.

### 3.3 Extraction of cliques

First a similarity matrix was created ( $A_{sim}$ ) containing adjectives as rows and columns. For doing so, a suitable similarity measure was applied to fill in the cells of  $A_{sim}$ . That is,  $A_{sim}(i, j) = sim(a_i, a_j)$ , where  $a_i$  and  $a_j$  denote the word2vec representations of adjectives from the selected vocabulary. The usual cosine similarity was calculated. That is:

$$sim_{cos}(v_1, v_2) = \frac{\mathbf{v}_1 \cdot \mathbf{v}_2}{\|v_1\| \|v_2\|} \quad (1)$$

In the second step  $A_{sim}$  similarity matrix was converted into an adjacency matrix  $A_a$  based on suitable cutting heuristics indicating whether the corresponding adjectives are semantically similar or not. Here a  $K$  cut-off parameter was set.

As a result, in this step  $A'_a$  symmetric square matrix is generated containing boolean values.  $A'_a$  adjacency matrix can be conceived of as a graph representation of the adjectives. Note that the use of cosine similarity guarantees that  $A'_a$  matrix is symmetric. Due to the reflexive nature of 'similarity' all the diagonal values of  $A'_a$  equal to 1.

In the last phase cliques were retrieved from the graph represented by the adjacency matrix to grasp adjectival near-synonymy classes.

### 3.4 Retrieving context nouns

In this phase adjectival cliques are validated by retrieving the set of nouns they may co-occur with. According to our expectation, different senses of an adjective are characterized by the different sets of nouns they co-occur with. These non-overlapping sets provide explicit information on the context of meaning discrimination. A characteristic set of nouns is found as follows:

1. We collect all the nouns an adjective co-occurs with; we do this for all adjectives in a clique. This step was performed on the basis of a 91.4 million-token subcorpus of the Hungarian Gigaword Corpus (Oravecz et al., 2014) compiled specifically for the present experiment. During the compilation process we aimed at preserving the original proportion of the genres, thus, every domain of HGC was included in the new corpus: newspapers, literature, scientific, official, personal and spoken language. Accordingly, our corpus was made up of

30.5m, 6.5 m, 11.6m, 8.8m, 28m and 6.6m tokens, respectively.

2. We compute the intersection of the above sets: those are the nouns that co-occur with each adjective of a clique. If at least one such noun exist for a clique, then we consider the given clique as a potential meaning candidate.

3. We repeat step 1 and step 2 for each clique a given adjective belongs to. This results in a set of nouns for each clique.

4. Finally, we take these sets and omit the intersections: we keep only the nouns for a clique which are exclusive to the given clique; they do not appear in the sets of the other cliques. Example (2) shows the cliques of the adjective *cinikus* 'cynical'. The nouns listed below the cliques are those shared by all members of the clique. Nouns in bold are the ones specific to the clique. These are the nouns indicating the specific meanings, therefore, we kept them for further evaluation.

(2) *cinikus* 'cynical'

Clique 1: *ostoba* 'silly', *cinikus* 'cynical', *demagóg* 'demagogic'

Nouns: *dolog* 'thing', *kérdés* 'question', *lépés* 'move', *mód* 'way', *szöveg* 'text'

Clique 2: *ostoba* 'silly', *cinikus* 'cynical', *arcátlan* 'impudent'

Nouns: *dolog* 'thing', ***ember*** 'person', *kérdés* 'question', *lépés* 'move', *mód* 'way'

Our presumption is that the resulting sets of nouns are the ones specific to the given cliques: they capture the given sense of the adjective that is shared among the other adjectives of the clique.

### 3.5 Evaluation

Finally, the results were evaluated according to different parameter settings. Since, to our knowledge, there is no similar database available for Hungarian, a qualitative evaluation was performed.

The main objective of the evaluation phase was twofold. On the one hand we aimed to verify our basic hypothesis, according to which the proposed techniques are able to provide a solid methodological background to discriminate between meanings. On the other hand, we also had the intention to catalogue the automatically retrieved adjectival senses with their salient context nouns and their perceived semantic categories, if possible. For doing so, first a coarse-grained evaluation was performed focusing on the main semantic properties of the automatically retrieved adjectival cliques. This was

followed by a fine-grained evaluation phase where we concentrated on the context nouns.

### 3.6 Parameter setting

Three parameters were identified as having a serious impact on the results.

- (i) The frequency of adjectives in the Hungarian National Corpus
- (ii) The  $K$  cut-off parameter
- (iii) The minimum frequency count of the nouns in the clique-validation step

#### *The frequency of the adjectives*

This parameter had to be taken into account to ensure that the word2vec representations were trained on sufficient amount of data.

#### *The impact of $K$ cut-off value*

Interestingly, we found that the value of the  $K$  cut-off parameter has a serious impact not only on the number of the resulting cliques but also on the semantic field to which they belong to. For instance, in the case of adjectives occurring at least 200 times,  $K = 0.9$  yielded only a handful of results: only 8 adjectives were assigned to more than one clique and only two cliques were validated by nouns. The retrieved cliques refer to numbers, months and days exclusively, therefore, they are not very interesting from a sense discrimination perspective. On the other hand, with the same parameter settings, but with a lower similarity cut-off value ( $K = 0.7$ )<sup>4</sup> we had 187 different adjectives belonging to multiple cliques, where all cliques are validated and discriminated by at least one following noun. Setting  $K$  to 0.7 resulted in 3847 single nodes and 1085 node pairs with one edge leaving only 1110 adjectives to possibly belong to multiple cliques. The high proportion of single nodes clearly implies that the  $K$  cut-off value should be set to a lower value.

#### *The effect of the frequency count of the following noun*

The minimum frequency count of the validating nouns ( $Freq_n$ ) also had to be taken into consideration. Two settings were tested ( $Freq_{ADJ} = 200, K = 0.7$ ). In the first setting a clique was considered valid if there was at least 1 noun occurring

<sup>4</sup>Our findings meet with the results of Veremyev et al. (2019). They constructed semantic networks based on word2vec representations of words with various thresholds and found that the threshold 0.7 resulted in the smallest, most compact cliques (largest clique size equaled to 245 and to 14, for the threshold 0.5 and 0.7, respectively).

at least 5 times with every element of the clique ( $Freq_n \geq 5$ ). Validating only a handful of cliques, this threshold value was deemed to be too high. To keep the coverage as high as possible, the value of  $Freq_n$  was set to 2. This change clearly improved the coverage, yielding 446 adjectives belonging to multiple cliques – out of the 6042 adjectives occurring at least 200 times in our input corpus with a word2vec representation.

In the rest of this section the results of the qualitative evaluation of these cliques will be presented ( $Freq_{ADJ} = 200, K = 0.7, Freq_n = 2$ ).

## 4 Relevant senses

In the present section we introduce some linguistic consideration that had to be taken into account during the evaluation phase to detect distinct classes of attributive modification.

### 4.1 Productivity

Distinct meanings may come from different sources. It is common to differentiate between collocational and more productive uses of an expression. In the course of the present research productivity is interpreted as a scale. On the one end of this scale there are collocations where both the adjective and the noun are fixed. In this case the meaning of the construction is yielded in a fully non-compositional way: neither component can be substituted with a near-synonym preserving the original meaning of the expression (eg. *fehér zaj* 'white noise' or *fekete doboz* 'black box').

Albeit collocations are possible sources of additional meanings, we are more interested in 'semi-compositional' constructions in the present WSI task, where compositionality operates on a restricted set of adjectives or nouns. For example, *fehér/szürke/fekete gazdaság* (literally 'white/gray/black economy')<sup>5</sup> are not considered collocations in the strict sense, since the restricted set of colors denotes a new dimension of meaning in the context of the noun *gazdaság* ('economy') (i.e. the extent to which a sector of economy is monitored and taxed). That is, one step further from collocations on the 'productivity scale' more interesting instances emerge, for example, *ékes* ('ornate') means *tipikus* ('typical') before a restricted

<sup>5</sup>Here, as opposed to the meaning of the English expression ('health related goods and services'), the Hungarian counterpart of 'white economy' refers to the monitored and taxed sectors of economy.

set of nouns (*példa* 'example' and *képviselő* 'representative').

## 4.2 Subcategorization

And indeed, the most interesting cases are those where the nouns form one or more semantic classes allowing the adjectives in the cliques to be synonyms in those semantically restricted contexts. In these cases the adjective subcategorizes the subsequent nouns (cf. Pustejovsky, 1995). For example, the different meanings of *könnyű* ('easy'), *komoly* ('serious'), *szép* ('nice'), *éles* ('sharp'), *finom* ('fine, delicate'), all can be discriminated on the basis of a set of synonym adjectives along with their semantically constrained nominal contexts. For example *könnyű* ('easy') has different meanings in the context of nouns referring to physical objects ('a lightweight bag'), nouns referring to clothes ('a light clothing'), foods ('a light lunch'), and before nouns like 'answer', 'task', 'solution' ('an easy answer/task/solution').

The size of the semantically constrained nominal sets may vary: on the other end of the scale there are really productive uses of adjectives that are still important for our purposes. For instance, the retrieved cliques imply that *vidám* 'merry' and *szomorú* 'sad' have different meanings when modifying nouns denoting humans and when modifying nouns referring to time periods. According to the cliques, we can say both *szomorú* [*időszak, év, nap*] ('sad [period, year, day]') and *gyászos* [*időszak, év, nap*] ('mournful [period, year, day]') but there is neither *bánatos* [*időszak, év, nap*] ('sorrowful [period, year, day]'), nor *gyászos* [*lány, ember*] ('mournful [girl, human]').

- (3) Clique 1: *szomorú* 'sad', *gyászos* 'mournful'  
Nouns: *időszak* 'period', *year* 'év', *nap* 'day'  
Clique 2: *szomorú* 'sad', *bánatos* 'sorrowful'  
Nouns: *lány* 'girl', *ember* 'human'

The adjective *vidám* 'merry' exhibits rather similar behavior to *szomorú* 'sad' from this perspective.

- (4) Clique 1: *vidám* 'merry', *derűs* 'bright'  
Nouns: *perc* 'minute', *nap* 'day', *hétvége* 'weekend'  
Clique 2: *vidám* 'merry', *jókedvű* 'cheerful'  
Nouns: *fiú* 'boy', *delfin* 'dolphin'

As opposed to humans (and dolphins), periods of time cannot be *jókedvű*, and in tandem with this, *derűs fiú* and *derűs delfin* are not well-formed constructions in Hungarian<sup>6</sup>.

<sup>6</sup>Interestingly, this is a well-known example in the lexical

## 5 Evaluation

### 5.1 Coarse-grained classification of adjectival cliques

#### *Tight semantic classes*

One problem we had to face during the evaluation phase is that not all adjectives were equally relevant from a meaning discrimination perspective. For example, dates and measures did not exhibit any interesting properties in most cases, even if they were assigned to multiple cliques. Instead, adjectives from these tight semantic classes tended to belong to multiple cliques with the very same meaning. According to our hypothesis, due to their varying sizes and varying distances between the elements, the adjectives belonging to tight semantic classes cannot be grouped into one clique in a coherent way, no matter what the parameter setting is. Another reason to disregard adjectives from tight semantic classes is that their lexical meaning seems to be rather straightforward not allowing for polysemy, except for a handful of more complex ones (eg. *fekete* 'black', *fehér* 'white', *szürke* 'gray'). For instance, *hétfői* ('of.Monday') was grouped under two different cliques:

- (5) Clique 1: *hétfői* 'of.Monday', *pénteki* 'of.Friday', *szombati* 'of.Saturday', *vasárnapi* 'of.Sunday'

Clique 2: *hétfői* 'of.Monday', *tegnapi* 'of.yesterday', *keddi* 'of.Tuesday', *csütörtöki* 'of.Thursday', *szerdai* 'of.Wednesday', *szombati* 'of.Saturday', *pénteki* 'of.Friday'

In the case of numerals, dates, names of colors, units of measurements and various national currencies the nouns did not supply enough evidence to accept the meaning discrimination indicated by the cliques.

#### *Named entities*

Another class of adjectives was made up of named entities, primarily countries, cities and surnames. In spite of the rather striking results, they were not considered in the present investigation, since our main focus is on lexical meaning here, while the clique-membership of NEs tend to reflect factual knowledge rather than lexical meaning. For instance, *egri* (related to the city of Eger) was assigned to two cliques [*egri, soproni, veszprémi*] (related to the cities of Eger, Sopron and Veszprém,

semantic research concerning English. As Pustejovsky (1995, p. 48) notes "[...] *sad* and *happy* are able to predicate of both individuals [...] as well as event denoting nouns".

respectively) indicating viticultural areas, whereas the other clique [*egri, esztergomi*] (related to the cities of Eger and Esztergom, respectively) are referring to archdioceses.

One interesting finding of the manual evaluation was that the 6*k* window size word2vec representation was rather efficient in the detection of tight semantic classes and cliques of named entities: out of the 446 adjectives 99 belonged to some types of named entities, 28 adjectives were terms of measurements, while 11 adjectives assigned to at least two cliques referred to numerals.

#### *Emotive intensifiers*

We found that emotive intensifiers tend to group in cliques not conveying separate meanings. For example:

- (6) Clique 1: *borzalmas* 'terrible', *iszonyatos* 'terrific', *rettenetes* 'awful'  
Nouns: *szenvedés* 'suffering', *kép* 'picture', *körülmény* 'circumstance'
- Clique 2: *borzalmas* 'terrible', *félelmetes* 'dreadful', *rettenetes* 'awful', *szörnyű* 'horrible'  
Nouns: *látvány* 'spectacle', *nap* 'day', *érzés* 'feeling'
- Clique 3: *borzalmas* 'terrible', *borzasztó* 'terrifying', *rettenetes* 'awful', *szörnyű* 'horrible', *rémes* 'fearful'  
Nouns: *emlék* 'memory', *élmény* 'experience'

While the cliques imply that negative emotive intensifiers form a coherent semantic class among adjectives, neither the cliques nor the following nouns do not supply enough evidence to discriminate between the meaning of cliques.

*nagy* 'great'

The adjective *nagy* ('great') and related notions, such as *óriási* ('huge'), *hatalmas* ('large'), etc, are posing another problem: here the abstraction step is quite easy to make along the various dimensions, therefore, in this case, lumping the sub-meanings indicated by the cliques may be a motivated choice. For example, *óriási* belongs to two different cliques characterized by plenty of nouns:

- (7) Clique 1: *óriási* 'huge', *nagy* 'great', *hatalmas* 'large'  
Nouns: *mosoly* 'smile', *oroszlán* 'lion', *roham* 'attack', *piramis* 'pyramid', etc.

Clique 2: *óriási* 'huge', *komoly* 'serious'

Nouns: *kaland* 'adventure', *konkurencia* 'concurrence', *kérdés* 'question', *lemaradás* 'lag', *marketing* 'marketing', *infláció* 'inflation', etc.

However, although *komoly* ('serious') cannot be used as a synonym of 'huge' before the elements of the first clique (eg. *komoly mosoly* 'a serious smile'  $\neq$  *óriási mosoly* 'a huge smile' and *komoly oroszlán* 'a serious lion'  $\neq$  *óriási oroszlán* 'a giant lion'), someone may claim that – in certain contexts at least – *óriási* and *komoly* conveys the same meaning at a certain level of abstraction. We confine ourselves only to make a notice on this phenomenon in the present paper and do not want to take a definite stance on this question.

## 5.2 Fine-grained evaluation of cliques

After excluding the irrelevant cases (cc. 240 adjectives altogether), a detailed evaluation took place aiming to create an adjectival database, where each sense is well-motivated and is characterized by the set of the context nouns. We also investigated whether these nouns can help humans to form concepts. For doing so, we went through on the resulting cliques manually. Maximum five context nouns were included into our database and we strove to select the salient context nouns for the given sense. We followed the procedure below:

1. The word2vec representations of the context nouns were used. They were generated as described in subsection 3.1.
2. The noun vectors were clustered using a hierarchic agglomerative algorithm to find subcategorization patterns.

For instance, we had *mindennapi* ('common') assigned to two cliques: dendograms in Figure 1 and Figure 2 depict the clusters of the context nouns. On the one hand, the respective near-synonyms are rather enlightening with regards to the two senses of the adjective, one of them being 'normal' or 'ordinary' while the other referring to regular, everyday activities. Based on the figures we can conclude that for example language-related things, such as *szóhasználat* ('word usage'), *nyelvhasználat* ('language use') are rather common or ordinary things than periodical ones; while *gyakorlás* ('practice') or *testmozgás* ('exercise') are regular, everyday activities and not necessarily common or ordinary ones. Therefore, the branches




Figure 1: The clusters of the context nouns of the adjectival clique [*mindennapi* 'common', *hétköznapi* 'normal']




Figure 2: The clusters of the context nouns of the adjectival clique [*mindennapi* 'common', *mindennapos* 'everyday']

of the dendrogram indicate the semantic classes of nouns the adjectival senses subcategorize.

As a result, out of the 446 adjectives with the given parameter setting, 53 adjectives were assigned to multiple cliques: to 118 cliques altogether. The list is available on GitHub: <https://github.com/nytud/HuWiC>. The qualitative evaluation yielded surprisingly insightful results in many cases, which may be not accessed with an introspective or even with a corpus-based methodology. Therefore, in spite of the low coverage we think that the research discussed here definitively worth pursuing in the future.

## 6 Conclusion and future work

The present paper describes an ongoing research, which intends to apply an unsupervised WSI approach to detect interpretable senses from monolingual corpora to contribute to the work of lexicogra-

phers, linguists and facilitate the creation of related benchmarks for the NLP community. For doing so, we came up with 4 necessary criteria to distinguish between senses, which were implemented in the next step. Finally, a detailed evaluation of the sense distinctions was performed yielding the conclusion that although the coverage definitively needs to be improved, in many cases the attained senses were surprisingly insightful supplying interpretable and intuitively not obvious sense distinction. However, during the evaluation it turned out that belonging to multiple near-synonymy classes is only a necessary but not sufficient condition for meaning discrimination, as adjectives may have collocate nouns or subcategorize multiple sets of nouns in a single clique (see the case of *könnyű* 'easy' in subsection 4.2). Since this method does not rely on any external knowledge base, it should be suitable for any low- or medium-resourced language.



## References

- Arleta Adamska-Sałaciak. 2006. *Meaning and the Bilingual Dictionary. The Case of English and Polish. (Polish Studies in English Language and Literature 18)*. Peter Lang, Frankfurt am Main.
- Julien Ah-Pine and Guillaume Jacquet. 2009. [Clique-Based Clustering for Improving Named Entity Recognition Systems](#). In *Proceedings of the 12th Conference of the European Chapter of the ACL (EACL 2009)*, pages 51–59, Athens, Greece. Association for Computational Linguistics.
- Asaf Amrami and Yoav Goldberg. 2019. Towards better substitution-based word sense induction. *ArXiv*, abs/1905.12598.
- Marco Baroni, Georgiana Dinu, and Germán Kruszewski. 2014. [Don't count, predict! a systematic comparison of context-counting vs. context-predicting semantic vectors](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 238–247, Baltimore, Maryland. Association for Computational Linguistics.
- Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Janvin. 2003. A neural probabilistic language model. *J. Mach. Learn. Res.*, 3:1137–1155.
- Jose Camacho-Collados and Mohammad Taher Pilehvar. 2018. [From word to sense embeddings: A survey on vector representations of meaning](#).
- Massimiliano Carrara and Pierdaniele Giaretta. 2004. [The many facets of identity criteria](#). *Dialectica*, 58(2):221–232.
- Stephen Clark. 2015. [Vector space models of lexical meaning](#). In *The Handbook of Contemporary Semantic Theory*, chapter 16, pages 493–522. John Wiley & Sons, Ltd.
- G. Frege. 1892. [Über Sinn und Bedeutung](#). In Mark Textor, editor, *Funktion - Begriff - Bedeutung*, volume 4 of *Sammlung Philosophie*. Vandenhoeck & Ruprecht, Göttingen.
- Dirk Geeraerts. 2015. [Lexical semantics](#). *International Encyclopedia of the Social & Behavioral Sciences*, pages 273–295.
- Judit Kuti, Enikő Héja, and Bálint Sass. 2010. [Sense disambiguation - 'Ambiguous sensation'?](#) In *Proceedings of the Seventh Conference on International Language Resources and Evaluation (LREC 2010)*, page 23–30, la Valetta, Malta.
- Márton Makrai. 2015. [Comparison of distributed language models on medium-resourced languages](#). In *XI. Magyar Számítógépes Nyelvészeti Konferencia (MSZNY 2015)*.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. [Efficient Estimation of Word Representations in Vector Space](#).
- Tomas Mikolov, Ilya Sutskever, Kai Chen, G.s Corrado, and Jeffrey Dean. 2013b. [Distributed Representations of Words and Phrases and their Compositionality](#). *Advances in Neural Information Processing Systems*, 26.
- Dávid Márk Nemeskey. 2020. *Natural Language Processing Methods for Language Modeling*. Ph.D. thesis, Eötvös Loránd University.
- Attila Novák and Borbála Novák. 2018. [Pos, ana and lem: Word embeddings built from annotated corpora perform better](#). In *Computational Linguistics and Intelligent Text Processing: 17th International Conference, (CICLing 2018)*, Hanoi, Vietnam. Springer International Publishing, Cham.
- Csaba Oravecz, Tamás Váradi, and Bálint Sass. 2014. [The Hungarian Gigaword Corpus](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 1719–1723, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Sabine Ploux and Bernard Victorri. 1998. [Construction d'espaces sémantiques à l'aide de dictionnaires de synonymes](#). *Traitement automatique des langues*, 1(39):146–162.
- James Pustejovsky. 1995. *The Generative Lexicon*. MIT Press, Cambridge, MA.
- Radim Rehurek and Petr Sojka. 2011. [Gensim—python framework for vector space modelling](#). *NLP Centre, Faculty of Informatics, Masaryk University, Brno, Czech Republic*, 3(2).
- Peter D. Turney and Patrick Pantel. 2010. [From frequency to meaning: Vector space models of semantics](#). *J. Artif. Int. Res.*, 37(1):141–188.
- Alexander Veremyev, Alexander Semenov, Eduardo L. Pasiliao, and Vladimir Boginski. 2019. [Graph-based exploration and clustering analysis of semantic spaces](#). *Applied Network Science*, 4:1–26.
- Tamás Váradi. 2002. [The Hungarian National Corpus](#). In *Proceedings of the Second International Conference on Language Resources and Evaluation, Las Palmas*, pages 385–389.
- Jean Véronis. 2003. [Sense tagging: does it make sense?](#) In *Corpus Linguistics by the Lune: a festschrift for Geoffrey Leech*, Frankfurt. Peter Lang.
- Gregor Wiedemann, Steffen Remus, Avi Chawla, and Chris Biemann. 2019. [Does BERT make any sense? interpretable word sense disambiguation with contextualized embeddings](#). *CoRR*, abs/1909.10430.