



Semantic Change Detection for the Romanian Language

Ciprian-Octavian Truică, Victor Tudose and Elena-Simona Apostol

EasyChair preprints are intended for rapid dissemination of research results and are integrated with the rest of EasyChair.

September 6, 2023

Semantic Change Detection for the Romanian Language

Ciprian-Octavian Truică
University Politehnica of Buchares
Bucharest, Romania
ciprian.truica@upb.ro

Victor Tudose
University Politehnica of Buchares
Bucharest, Romania
victor.tudose@stud.acs.upb.ro

Elena-Simona Apostol
University Politehnica of Buchares
Bucharest, Romania
elena.apostol@upb.ro

Abstract—Automatic semantic change methods try to identify the changes that appear over time in the meaning of words by analyzing their usage in diachronic corpora. In this paper, we analyze different strategies to create static and contextual word embedding models, i.e., Word2Vec and ELMo, on real-world English and Romanian datasets. To test our pipeline and determine the performance of our models, we first evaluate both word embedding models on an English dataset (SEMEVAL-CCOHA). Afterward, we focus our experiments on a Romanian dataset, and we underline different aspects of semantic changes in this low-resource language, such as meaning acquisition and loss. The experimental results show that, depending on the corpus, the most important factors to consider are the choice of model and the distance to calculate a score for detecting semantic change.

Index Terms—semantic change, word embeddings, low resource language

I. INTRODUCTION

Language is in a continuous process of change that occurs permanently, language change being the phenomenon that drives language evolution, as a process of adaptation to the environment and the ways other speakers use the language [1], [2]. The various instances of language change are classified into different categories, such as regular phonetic changes, changes in word usage, and changes in the way words appear together, i.e., syntactic changes. Although it is usually a continuous process that follows regular patterns, very abrupt changes in the meanings of words can still occur, usually motivated by a change in the context a community lives in [7], [9], [15].

Semantic change, as a phenomenon permanently present in language evolution, is an important aspect that should be taken into account when working with historical data [3]. Historical linguists, lexical typologists, and other humanities and social science experts have studied the meaning of words and how it changes over time. The method used by these experts is known as "close reading", being the manual study of texts.

The growing availability of digitalized historical corpora has enabled larger-scale quantitative studies of language evolution and a rise in the usage of computational approaches, both automatic and semi-automatic, for computational linguistics tasks. New approaches to semantic change detection had appeared with the recent advances in Natural Language Processing such as word embedding techniques. The main approaches use either non-contextualized (static) word embeddings that generate one embedding per word regardless of the context [5], [14], [28]

and contextualized word embeddings that generate multiple contested dependent embeddings for the same word [17], [24].

The main objective of this work is to analyze the semantic change in Romanian. To achieve this, we:

- (1) develop an architecture for determining semantic change using two parallel corpora from different time periods to model static and contextual word embeddings;
- (2) determine the accuracy of our model using two lists of words, one with words that undergo semantic change and one with words for which their sense remains stable over the time periods;
- (3) explore how different metrics can be affected by the corpora employed in our analysis;
- (4) create a web application to visualize and analyze the models.

The rest of this paper is structured as follows: In Section II, we describe some of the most relevant current similar works. Section III introduces and presents our proposed solution. In Section IV, we briefly describe our datasets and experimental setup and analyze the results obtained in our experiments. Section V provides an in-depth discussion of our results. Section VI concludes our paper and hints at future research.

II. RELATED WORK

In the current literature, static and contextual word embedding models together with different distance metrics have been used to determine semantic change. The main idea is to create word embeddings for parallel corpora that span over different periods of time. Afterward, a score is calculated based on the distance between two vector representations from different periods of time of the same word. If the score is above a given threshold, then the word undergoes a semantic change.

Gong et al. (2020) [14] create condition-specific embeddings that consider spatial-temporal dimensions to model the words in a corpus and capture language evolution across both time and location. Bizzone et al. (2019) [5] use hyperbolic embeddings [19] to detect semantic changes that occur in domain-specific literature. Tsakalidis and Liakata (2020) [28] use deep neural networks, e.g., sequence-to-sequence (Seq2Seq) models, together with Word2Vec Skip-Gram model [18] to create new word representation. Semantic change is evaluated by calculating the average cosine similarity of word representations for different time periods. Wegmann et al. (2020) [30] use local

neighborhood [15] to determine semantic shifts using static word embeddings.

Recent studies have also explored the use of contextual word embedding and transformer models, e.g., ELMo (Embeddings from Language Models) [22] and BERT (Bidirectional Encoder Representations from Transformers) [8], in lexical semantic change tasks. These studies utilized contextualized word representations to measure semantic shifts of words over time and evaluated their approach on large diachronic English corpora as well as other high-resource languages [4], [13], [17], [24], e.g., Latin, Italian, German, and Swedish. The results demonstrated that their proposed models captured a range of synchronic and diachronic linguistic aspects, and outperformed baselines based on normalized frequency difference or cosine distance methods.

III. PROPOSED SOLUTION

Figure 1 presents our proposed architecture consisting of three main modules, i.e., Corpus, Embeddings, and Metric Modules. Our solution is designed to facilitate training multiple models on different corpora. After the model is trained, we store the results obtained by each evaluation metric as well as the ranked list of the words that undergo semantic changes. We implement our semantic change detection architecture using Python v3.7. The code is available on GitHub at <https://github.com/DS4AI-UPB/SemanticChange-RO>.

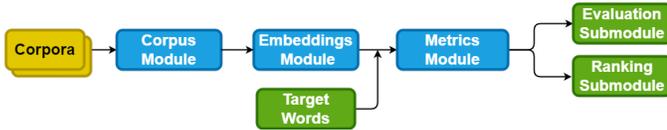


Fig. 1: The Proposed Architecture

a) *Corpus Module*: This module performs text preprocessing. Given a corpus of documents, we employ the following steps. (1) **Vocabulary extraction**: used by the Embeddings Module to extract the terms present in the corpora; and (2) **Time Interval Merging**: this takes a set of texts from different moments in time and a list of intervals, and merges the texts into new corpora that are obtained by merging the initial texts according to interval adherence.

b) *Embeddings Module*: This module receives a set of time interval split corpora and the vocabulary. For each corpus, we train static or contextual word embedding models, i.e., Word2Vec and ELMO, respectively.

Static Word Embeddings. We choose Word2Vec Skip-Gram with Negative Sampling (SGNS) as the static word embedding model based on the results of Hamilton et al. (2016b) [16]. Table I presents the parameters used for training the vectors.

We use two strategies to train our word embeddings. The first implementation, i.e., SGNS-OP, trains the vectors on the time interval split corpora in parallel and uses the orthogonal Procrustes method to align the word embeddings (Figure 2). The Procrustes method [27] states that given two matrices A and B we need to find a matrix Ω such as $R = \operatorname{argmin}_{\Omega} (\|\Omega \cdot$

TABLE I: Word2Vec SGNS model parameters

Word2Vec SGNS Parameters		
Name	Description	Value
min count	minimum number of word occurrences	1
vector size	dimensionality of vectors	100
window	max distance between current and predicted word	5
alpha	learning rate	0.025
negative	how many negative samples are used	5
ns exponent	correlation between frequency and negative sampling	1
epochs	number of training iterations	5

$A - B \|_F$), where $\Omega^T \cdot \Omega = I$ and $\|A\|_F = \sqrt{\sum_{i,j=1}^n |a_{i,j}|^2}$ called also the matrix norm or the Frobenius norm.

SGNS with Orthogonal Procrustes

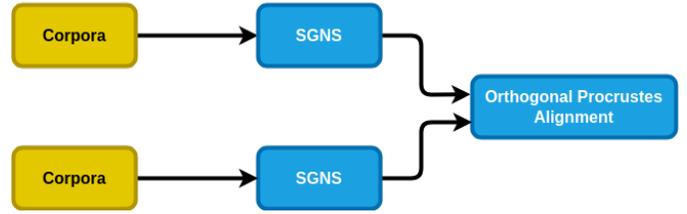


Fig. 2: SGNS-OP Architecture

The second strategy for training the static word embeddings, i.e., SGNS-WI, makes use of word injection. Using this method, we train the word embeddings on all the corpora at once but we tag some target words to obtain two different embeddings that correspond to the corpus the word originates from (Figure 3).

SGNS with Word Injection



Fig. 3: SGNS-WI Architecture

In our implementation, we use Gensim [23] to train both static word embeddings, i.e., SGNS-OP and SGNS-WI, respectively.

Contextual Word Embeddings. For the contextual word embeddings, we use AllenNLP [12] ELMo [22] implementation. One reason to use contextualized embeddings in the case of semantic change detection is to better model polysemy. Since the word's meaning is disambiguated by its surrounding, naturally, a way of encoding the immediate surrounding of a word could help with capturing the particular sense a word has in its respective context. The implementation of ELMo uses several modules. (1) **Dataset Reader Module**: A module for preparing the corpus that (a) extracts the tokens, and (b) creates a token index. (2) **Uncontextualised Embedding Module**: A module that creates the non-contextualized embeddings used by ELMo. This module uses a character-level Convolutional Neural Network that creates character embeddings of size 16. The character embeddings are then passed to a second layer that outputs non-contextualized word embeddings with a dimensionality of 256. (3) **Contextualizer**: A contextualizing

layer that extracts the context for the embeddings. We used a BiLSTM contextualizer which concatenates a forward contextualized embedding to a backward contextualized embedding. (4) **Bidirectional Language Model Loss Module**: Computes the loss for the ELMO model based on the sum of the log probabilities of all the sentences in the corpus.

We create a contextual word embedding model for each time window corpus. To analyze the semantic change, for each corpus used to create a model we take all the sentences that contain a target word and then extract the embeddings from these sentences using the ELMO models trained on each corpus. Thus, given two corpora A and B , we train a contextual word embedding mode for each: (1) ELMO-PREV trained on A , and (2) ELMO-POST trained on B . Using these embeddings we create 4 sentence embeddings, as follows: 1) E_1 are the embeddings obtained when applying the ELMO-PREV on the sentences $s_A \in A$ 2) E_2 are the embeddings obtained when applying the ELMO-PREV on the sentences $s_B \in B$ 3) E_3 are the embeddings obtained when applying the ELMO-POST on the sentences $s_A \in A$ 4) E_4 are the embeddings obtained when applying the ELMO-POST on the sentences $s_b \in A$. After obtaining these embeddings, we use different metrics to compare them and detect semantic changes. Figure 4 presents the architecture to train the ELMO-PREV and ELMO-POST models.

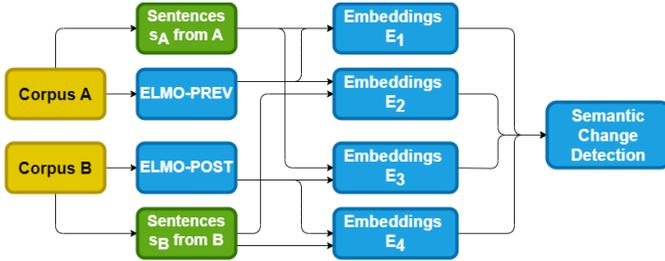


Fig. 4: Contextual Models Architecture

c) *Metrics Module*: We employ multiple metrics to detect different nuances of semantic change that words undergo over time [15], [25], [30], e.g., local neighborhoods metrics determine better cultural changes, while global displacement measurements tend to detect regular semantic changes. Thus, the metrics module receives the word embeddings and the target words for testing the semantic change and then computes a distance function to determine if the word’s meaning undergoes any changes over time.

For every corpus, we create associated lists of target words to test the models’ accuracy with words that undergo semantic change and words that kept their sense. We use the evaluation metrics with these lists to determine if we manage to detect correctly semantic changes by determining if the word maintains its ranking, thus, its meaning. Furthermore, when testing each model’s performance using the lists of target words, we also record the number of true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN). The evaluation metrics and the rankings are stored using

Evaluation Submodule and Ranking Submodule, respectively.

Static Word Embeddings Metrics. For the static word embedding models, given two n -dimensional static word embeddings $\mathbf{a} \in \mathbb{R}^n$ and $\mathbf{b} \in \mathbb{R}^n$ for the same word w , we use the following distances:

Euclidean distance computes the global displacement between the embeddings, indicating a change in the semantic properties: $d(\mathbf{a}, \mathbf{b}) = \sqrt{\sum_{i=1}^n (\mathbf{a}_i - \mathbf{b}_i)^2}$

Manhattan distance measures the global distance between two embeddings taking into account the dimensionality and number of vector components: $d(\mathbf{a}, \mathbf{b}) = \sum_{i=1}^n |\mathbf{a}_i - \mathbf{b}_i|$.

Canberra distance is a weighted measure of the Manhattan distance that proved it provides good results when used for detecting semantic change [4], [13] $d(\mathbf{a}, \mathbf{b}) = \sum_{i=1}^n \frac{|\mathbf{a}_i - \mathbf{b}_i|}{|\mathbf{a}_i| + |\mathbf{b}_i|}$

Cosine distance determines the dissimilarity between the embeddings to detect semantic change: $d(\mathbf{a}, \mathbf{b}) = 1 - \frac{\sum_{i=1}^n \mathbf{a}_i \cdot \mathbf{b}_i}{\sqrt{\sum_{i=1}^n \mathbf{a}_i^2} \sqrt{\sum_{i=1}^n \mathbf{b}_i^2}}$

Bray-Curtis distance measures the dissimilarity between two word embeddings that considers their values relative magnitudes: $d(\mathbf{a}, \mathbf{b}) = \sum_{i=1}^n \frac{|\mathbf{a}_i - \mathbf{b}_i|}{|\mathbf{a}_i + \mathbf{b}_i|}$.

Correlation distance: computes the correlation between two word embeddings: $d(\mathbf{a}, \mathbf{b}) = 1 - \frac{(\mathbf{a} - \bar{\mathbf{a}}) \cdot (\mathbf{b} - \bar{\mathbf{b}})}{|\mathbf{a} - \bar{\mathbf{a}}|_2 |\mathbf{b} - \bar{\mathbf{b}}|_2}$, where $\bar{\mathbf{a}}$ and $\bar{\mathbf{b}}$ are the mean of the elements of \mathbf{a} and \mathbf{b} , respectively, and \cdot is used to denote the dot product between two vectors.

Contextual Word Embeddings Metrics. For measuring the performance of the contextual word embeddings models for the semantic change detection task, we need to consider that we are analyzing multiple word embeddings for the same word in the same period of time in that are context dependent. Thus, we must consider different set-based methods that group together the word embeddings, i.e., Average Pairwise Distances, Jensen-Shannon divergence, and Cluster Count based on Affinity Propagation.

Average Pairwise Distances (APD) are a class of distance functions that compute a given distance $d(\cdot)$ for every possible pair of vectors $\langle \mathbf{a}, \mathbf{b} \rangle$ from two sets of vectors A and B [13]: $APD(A, B, d) = \frac{1}{|A| \cdot |B|} \sum_{\mathbf{a} \in A, \mathbf{b} \in B} |d(\mathbf{a}, \mathbf{b})|$. We compute the Average Pairwise Distances using the following distances: Euclidean, Manhattan, Canberra, and Cosine.

Jensen-Shannon divergence is a method for measuring similarity between two probability distributions P and Q based on Kullback-Leibler divergence that always outputs a finite value: $JSD(P, Q) = \frac{1}{2}(D(P, M) + D(M, Q))$, where M is the average of the two probability distributions $M = \frac{1}{2}(P + Q)$ and $D(A, B)$ is the Kullback-Leibler divergence.

Cluster Count obtains similar vector clusters by using the Affinity Propagation algorithm [11].

d) *Visualization and exploration*: To visualize and explore the results of our models, we created a web application using Flask. The application contains the following pages created dynamically using Jinja2: (1) Index: a page that lists the selected models. (2) Results: a page for visualizing the results of the different models (Figures 5 and 6). (3) Select: a page for selecting models for comparison. (4) Compare: a page

that presents the comparison results for the selected models (Figures 7a, and 7b).

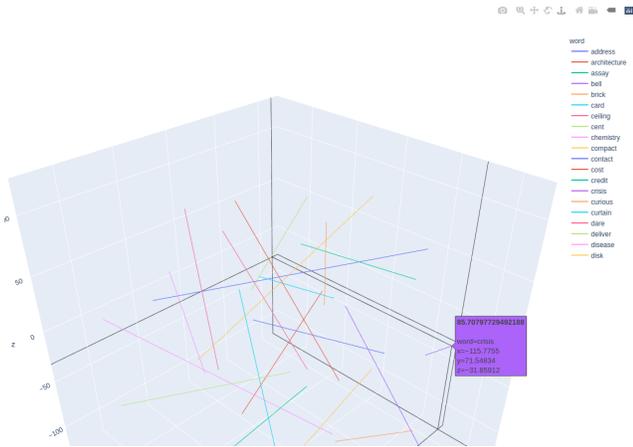


Fig. 5: Visualization of semantic change using 3D scatter plot where embeddings, that are projected into the 3D scape, are united with a line if they represent the same word

The projection from the embedding space, which has at least 100 dimensions, to the 3D space is done using the t-distributed stochastic neighbor embedding (t-SNE) method [29]. The t-SNE algorithm is composed of two stages:

- 1) The algorithm creates a probability distribution over pairs of data entities, where very similar entities are assigned a higher probability and dissimilar objects are assigned a lower probability
- 2) The algorithm creates a set of points in a low dimensionality space, creates a similar probability distribution and minimizes the Kullback-Leibler divergence between the two distributions concerning the location of points in space

Name	Language	Task
sgns_wl_rodica_transylvania_19vs20_21	ro	Transylvanian corpora.one before the 19th century and one after the 20th century

Metric	Score	True positives	True negatives	False positives	False negatives
euclidean	0.875	7	0	0	1
cosine	0.75	6	0	0	2
canberra	0.875	7	0	0	1
manhattan	0.875	7	0	0	1
braycurtis	0.875	7	0	0	1
correlation	0.75	6	0	0	2

Fig. 6: Results page presenting the selected model and the obtained results

The graphics are created using the Plotly plotting framework. We choose this framework since the graphics could be directly embedded into the web pages and they would be interactive. The graphics can be manipulated in a variety of ways (Figure 7b), e.g., selecting/deselecting metrics for visualization, selecting/deselecting the model for analysis, etc. Furthermore, these graphs can be saved.

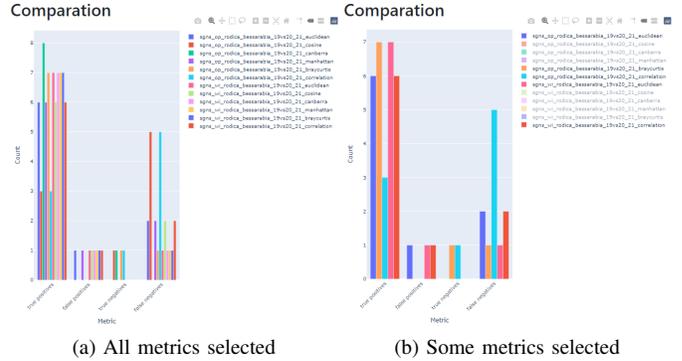


Fig. 7: Metrics comparison graph

IV. EXPERIMENTAL RESULTS

In this section, we present our datasets and experimental setup, and we analyze the obtained results.

A. Datasets details

a) *SEMEVAL-CCOHA (SemEval Clean Corpus of Historical American English)*: The first set of experiments is conducted on the SEMEVAL-CCOHA corpora [26]. We use these corpora to test the different word embeddings and evaluate their performance before moving to the Romanian corpus.

b) *RODICA (Romanian Diachronic Corpus with Annotations)*: RODICA is a collection of 153 Romanian language new articles collected from the four historical regions of current-day Romania: Wallachia, Transylvania, Moldavia, and Bessarabia. The news articles were published in the local press, spanning over a time period that covers the half of the 19th century until the early part of the 21st century. Table II presents the distribution of texts by region and by date.

TABLE II: RODICA documents distribution

Region	19th century	20th century	21st century	Total
Bessarabia	4	63	0	67
Moldavia	11	3	6	20
Transylvania	37	0	7	44
Wallachia	5	60	0	19
All	58	63	32	153

B. Experimental setup

For the SEMEVAL-CCOHA corpus, we use the textual data as provided in [26]. For RODICA, we create 4 corpora for each of the historical regions: (1) RODICA-BS, the text from the region of Bessarabia; (2) RODICA-MD, the text from the region of Moldavia; (3) RODICA-TR, the text from the region of Transylvania; (4) RODICA-WL, the text from the region of Wallachia. Before training the word embeddings and to mitigate against the small number of text in the dataset written in the 21st century, we split each dataset into 2 time periods: before and after 1900. We obtain the following pairs of corpora that are used together in the tests.

We use The Catalogue of Semantic Shifts [31] for Romanina to obtain a list that contains only words that changed their meaning. We also use Swadesh [20] to create a second list with of words that tend to change their meaning rather slowly in a language at most of their levels, semantic and phonetic, being mainly used to asses when the split between two languages occurred in their ancestry. For these 2 lists, we keep only the words that appear in both time periods for a given region. We use the evaluation metrics with these lists to determine if we manage to detect correctly semantic changes by determining if the word maintains its ranking, and thus its meaning, in the list of words created using Swadesh.

For the experiments, we train the following models: (1) SGNS-OP: the representation model formed with Skip-Gram negative sampling and Orthogonal Procrustes, and (2) SGNS-WI: the representation model formed with Skip-Gram negative sampling and Word Injection, (3) ELMO-PREV: an ELMO model trained on the diachronic corpus with text from the first time period; (4) ELMO-POST: an ELMO model trained on the diachronic corpus with text from the second time period.

C. Results on SEMEVAL-CCOHA

Static word embeddings. The first set of experiments uses SGNS-OP and SGNS-WI on the SEMEVAL-CCOHA. We observe that both approaches obtain similar results (Table III).

TABLE III: SEMEVAL-CCOHA results (Abbreviations: TP - True positive, TN - True Negative, FP - False Positive, FN - False Negative)

Model	Metric	Score	TP	TN	FP	FN
SGNS-OP	Euclidean	0.73	39	0	2	12
	Manhattan	0.73	39	0	2	12
	Canberra	0.90	48	0	2	3
	Cosine	0.56	29	1	1	22
	Bray Curtis	0.86	46	0	2	5
	Correlation	0.56	29	1	1	22
SGNS-WI	Euclidean	0.72	38	0	2	13
	Manhattan	0.72	38	0	2	13
	Canberra	0.89	47	0	2	4
	Cosine	0.57	29	1	1	22
	Bray Curtis	0.85	45	0	2	6
	Correlation	0.57	29	1	1	22

For both models, by intersecting the various metrics, we can observe that Euclidean and Manhattan distance have the highest correlation. Another strong correlation can be observed between the Consine and The Bray-Cutis distances. Table IV presents the words that undergo semantic change detected by the SGNS-OP model w.r.t. each metric, while Table V presents the words that undergo semantic change detected by the SGNS-WI model w.r.t. each metric.

Contextual word embeddings. Table VI presents the results when using contextual word embeddings. The best performing average pairwise distance is APD Cosine followed by Jensen–Shannon divergence for both ELMO-PREV and ELMO-POST on SEMEVAL-CCOHA.

The first assumption we made regarding contextualized embeddings is that they can more accurately represent the

TABLE IV: Semantic changes detected using SGNS-OP on SEMEVAL-CCOHA

Euclidean	Manhattan	Canberra	Cosine	Bray-Curtis	Correlation
justice	justice	justice	lift	disease	lift
federal	federal	leaf	disease	lift	disease
address	address	maid	drug	drug	drug
maybe	maybe	disease	leaf	justice	leaf
maid	maid	lift	justice	leaf	justice
lift	lift	drug	brick	maid	maid
leaf	gain	domain	maid	excuse	brick
gain	leaf	excuse	disorder	domain	disorder
disorder	drug	disorder	excuse	brick	excuse
drug	disease	cost	domain	disorder	domain

TABLE V: Semantic changes detected using SGNS-WI on SEMEVAL-CCOHA

Euclidean	Manhattan	Canberra	Cosine	Bray-Curtis	Correlation
honour	disease	disease	disk	maid	disk
disease	honour	honour	credit	assay	credit
deliver	deliver	disk	family	dining	family
address	address	gallery	federal	disk	federal
excuse	excuse	doubt	martial	context	martial
lip	lip	lip	disorder	credit	disease
federal	federal	federal	disease	architecture	disorder
gallery	gallery	family	maid	martial	maid
energy	energy	cost	architecture	family	honour
family	family	extracellular	honour	compact	dining

TABLE VI: Results obtained for ELMO-PREV and ELMO-POST on SEMEVAL-CCOHA

Model	Metric	Score	TP	TN	FP	FN
ELMO-PREV	APD Euclidean	0.15	8	7	0	83
	APD Manhattan	0.15	8	7	0	83
	APD Canberra	0.15	8	7	0	83
	APD Cosine	0.78	75	1	6	16
	Jensen–Shannon divergence	0.55	53	1	6	38
	Cluster Count	0.52	48	3	4	43
ELMO-POST	APD Euclidean	0.20	13	7	0	78
	APD Manhattan	0.20	13	7	0	78
	APD Canberra	0.20	13	7	0	78
	APD Cosine	0.78	75	1	6	16
	Jensen–Shannon divergence	0.50	48	1	6	43
	Cluster Count	0.39	34	4	3	57

various nuances in the meaning of a word. Thus, we choose for experiments some words that were listed in the Catalogue of Semantic Shifts as Polysemy (acquisition of new meaning by a word). For example, the word "*ceiling*" (with the meaning "*the hard upper part of the inside of the mouth*" in addition to its original sense) is detected by both ELMO models as semantic change, while the SGNS-OP and SGNS-WI models do not detect it. This result is most likely a false positive since the word did not acquire this sense in the second corpus.

When ranking the semantic changes (Tables VII and VIII), we observe that the contextual word embedding models detect several semantic changes that are not listed in the Catalogue of Semantic Shifts, e.g., *orange*. At first, we might think this captures the fact that at some point in time, the word for the orange (the fruit) also began to mean the color of the fruit. But the earliest recorded usage of this word as color comes from 1512 and dates before this year are not included in SEMEVAL-CCOHA.

Analyzing only the words that both models detected that have undergone semantic change, we obtain different results. We stored the true positives for every metric of both models (Table VI). The best performing metric for both ELMO models

is Cosine which also manages to detect all the words from the other metrics. Furthermore, the lists of semantic changes detected with APD Cosine for ELMO-PREV and ELMO-POST are identical. We also observe using the Cluster Count metric that with ELMO we obtain more regular embeddings [10].

TABLE VII: Semantic changes detected using ELMO-PREV on SEMEVAL-CCOHA

Euclidean	Manhattan	Canberra	Cosine	Bray-Curtis	Correlation
disorder	virtual	disorder	virtual	disorder	virtual
disk	strategy	disk	strategy	disk	strategy
ceiling	spine	ceiling	spine	ceiling	spine
negligence	replacement	negligence	replacement	negligence	replacement
upwards	program	upwards	program	upwards	program
assay	participate	assay	participate	prestige	participate
prestige	node	prestige	node	assay	node
optical	mirror	optical	mirror	optical	mirror
virtual	jet	virtual	jet	virtual	jet
strategy	infection	strategy	infection	strategy	infection

TABLE VIII: Semantic changes detected using ELMO-POST on SEMEVAL-CCOHA

Euclidean	Manhattan	Canberra	Cosine	Bray-Curtis	Correlation
disk	woman	disorder	structure	disk	lip
disorder	wine	disk	orange	disorder	credit
negligence	west	negligence	mirror	negligence	address
jet	virus	ceiling	disease	jet	bell
ceiling	virtual	jet	cent	ceiling	cost
node	vector	node	crisis	node	scene
virtual	vaccine	optical	gain	virtual	justice
prestige	user	replacement	sleep	prestige	west
replacement	upwards	prestige	lift	replacement	gain
upwards	trap	virtual	drug	upwards	parent

Table VI presents some of the words detected by the contextual word embedding models that undergo semantic change.

TABLE IX: SEMEVAL-CCOHA: semantic change examples

Word	Change
lift	this is an example of the emergence of a new meaning for a word as a response to the evolution of technology. In the second corpora, the word lift is used in the context "all this have been reporting to the management by the lift boy" where the term "lift" refers to the elevator, while it refers to the verb "to lift" for all the other instances as in the context "the angel shall stand upon the sea and the earth and lift up his hand to heaven"
bell	this word is used in the first corpus with the meanings of "bell"(musical instrument), "bell of cowslip"(a flower), and a surname. In the second corpus the meaning of "bell of cowslip" is not encountered.
leaf	In the first corpus leaf is used both in its literal sense as a part of the plant as in the context "tree with rough leaf", and in a metaphoric use in expressions such as "say the fiend and he shakes like a leaf when cast" and "the leaf of the book once". In the second corpus, it is used only in its denominative sense
mirror	this is most likely a false positive since the newly acquired meaning of the word mirror from the Catalogue of Semantic Shift is the meaning of "a copy of an internet page to be used if the original is down".

Although the cosine distance performs better with the contextual embedding models, non-contextualized embeddings performed better. Our findings are similar to the findings of Basile et al. (2020) [4]. Thus, because of the small size of the RODICA corpus and a very small performance increase for the cosine distance, we choose not to perform experiments with ELMO in the RODICA corpus.

D. Results on RODICA

For this set of experiments, we train both models on each of the 4 regional datasets. We assess the models' performance by looking at how high in the ranking the words that did not change its semantic properties are positioned, i.e., *aici* (*here*)

and *acolo* (*there*). Note: in the following tables, orange marks the word(s) that did not undergo any semantic change.

a) *Bessarabia Region.*: For this dataset, we observe that SGNS-OP performs slightly better than SGNS-WI (Tables X and XI). We have the same results when using SGNS-OP for both Cosine and Correlation distances, e.g., the word *aici* is ranked 5 out of 10 words in both cases. Furthermore, we observe a strong correlation (i.e., the ranking lists are similar to identical) among the results obtained with the following distance pairs: < Euclidean, Manhattan >, < Cosine, Correlation >, and < Canberra, Bray-Curtis >. Table XII discusses some of the semantic changes detected.

TABLE X: SGNS-OP: RODICA-BS Semantic changes

Euclidean	Manhattan	Canberra	Cosine	Bray-Curtis	Correlation
mare	mare	grâu	grâu	grâu	grâu
aici	aici	vară	vară	vară	vară
grâu	grâu	aici	mie	aici	mie
mic	mic	mare	mic	mare	mic
vară	vară	mie	aici	mic	aici
mie	mie	mic	cale	cale	cale
inimă	inimă	inimă	inimă	inimă	inimă
cale	cale	cale	parte	cale	parte
parte	parte	parte	mare	parte	mare

TABLE XI: SGNS-WI: RODICA-BS Semantic changes

Euclidean	Manhattan	Canberra	Cosine	Bray-Curtis	Correlation
mare	mare	aici	mare	aici	mare
aici	aici	mare	aici	mare	aici
cale	cale	cale	cale	cale	cale
parte	parte	vară	parte	vară	parte
inimă	inimă	inimă	inimă	grâu	inimă
vară	vară	parte	mie	inimă	mie
grâu	grâu	grâu	vară	parte	vară
mie	mie	mie	mic	mie	mic
mic	mic	mic	grâu	mic	grâu

TABLE XII: RODICA-BS: semantic change examples

Word	Change
cale	this word is used in the before 1900 corpus with the sense "to find adequate" in the context "Prea Sfântul Sinod nu va găsi cu cale a-i da arătată Eparhie" (The Holy Synod won't find it adequate to give him an Eparchy), in the after 1900 corpus is used with several meanings, some of them being "method" "pe cale de economii aspre" (with harsh economic methods)
vară	this is most likely a false positive given by a difference in the number of usages across corpora

b) *Moldavia Region.*: The Moldavia corpora are somewhat balanced being made out of 11 texts for the period before 1900 and 9 texts for the period after 1900. The best results are obtained using the SGNS-OP model together with the Cosine and Correlation distances (Tables XIII and XIV), which also obtain the same rankings. Table XV discusses some of the detected semantic changes.

c) *Transylvania Region.*: We observe that SGNS-WI does not manage to correctly identify semantic change as the words that do not change their sense are ranked last (Table XVI). In contrast, SNGS-OP obtains good results managing to determine correctly semantic changes (Table XVII). A particularity of this corpora is the high occurrence of non-standard spelling in the before 1900 corpus, containing diacritics originating from an earlier form of the Romanian Latin Alphabet and non-standard phonetic realizations, e.g., "*dobendî*" instead of

TABLE XIII: SGNS-OP: RODICA-MD Semantic changes

Euclidean	Manhattan	Canberra	Cosine	Bray-Curtis	Correlation
parte	parte	grâu	grâu	grâu	grâu
aici	aici	aici	drum	aici	drum
mare	mare	țara	mie	țara	mie
țara	țara	parte	deschis	drum	deschis
acolo	acolo	drum	lume	parte	lume
drum	drum	mare	aici	mare	aici
grâu	grâu	mie	țara	mie	țara
mie	mie	acolo	acolo	acolo	acolo
deschis	deschis	deschis	mare	deschis	mare
lume	lume	lume	parte	lume	parte

TABLE XIV: SGNS-WI: RODICA-BS Semantic changes

Euclidean	Manhattan	Canberra	Cosine	Bray-Curtis	Correlation
parte	parte	acolo	parte	acolo	parte
aici	aici	aici	aici	parte	aici
acolo	acolo	parte	acolo	aici	acolo
mare	mare	cale	mare	cale	mare
cale	cale	mare	greu	mare	greu
față	față	greu	cale	față	cale
greu	greu	întreg	întreg	întreg	întreg
întreg	întreg	față	față	greu	față

TABLE XV: RODICA-MD: semantic change examples

Word	Change
grâu	this words appear to change due to a difference in the number of its usage across corpora, having four usages in the first corpus and only one in the second corpus
deschis	this word is used in with the sense "open" in the corpus with texts after 1900, while is used in a more figurative manner in the corpus with texts before 1900: "pocnetul deschis" (the loud noise)

the correct standard Contemporary Romanian word "dobânda". Table XVIII discusses some of the detected semantic changes.

TABLE XVI: SGNS-OP: RODICA-TR Semantic changes

Euclidean	Manhattan	Canberra	Cosine	Bray-Curtis	Correlation
mare	mare	față	întreg	întreg	întreg
față	față	întreg	față	față	față
întreg	întreg	greu	cale	greu	cale
greu	greu	cale	greu	cale	greu
cale	cale	mare	mare	mare	mare
parte	parte	acolo	aici	acolo	aici
acolo	acolo	parte	acolo	parte	acolo
aici	aici	aici	parte	aici	parte

TABLE XVII: SGNS-WI: RODICA-TR Semantic changes

Euclidean	Manhattan	Canberra	Cosine	Bray-Curtis	Correlation
parte	parte	acolo	parte	acolo	parte
aici	aici	aici	aici	parte	aici
acolo	acolo	parte	acolo	aici	acolo
mare	mare	cale	mare	cale	mare
cale	cale	mare	greu	mare	greu
față	față	greu	cale	față	cale
greu	greu	întreg	întreg	întreg	întreg
întreg	întreg	față	față	greu	față

TABLE XVIII: RODICA-TR: semantic change examples

Word	Change
greu	in the before 1900 corpus the word is used as an adjective in contexts such as "drumul acelu greu al publicității" (that hard way of publicity), while in the after 1900 corpus is used more as an adverb in contexts such as "este greu de crezut"(it's hardly believable)
vară	this is most likely a false positive given by a difference in the number of its usage across the corpora

d) Wallachia Region.: Both models do not obtain satisfactory results on the RODICA-WL. The RODICA-WL corpus with texts after the 1900 does not contain any of the Swadesh words. Thus, we cannot use the same ranking method for our analysis. Thus, we test for semantic change the following

3 words: *lună* (moon), *țară* (county), and *parte* (part). We observe that the models determine that the word "parte" has the highest semantic change regardless of the metric employed (Table XIX). Table XX presents some of the detected semantic changes.

TABLE XIX: RODICA-WL Semantic changes

Model	Metric	Word 1	Word 2	Word 3
SGNS-WI	Euclidean	parte	lună	țară
	Manhattan	parte	lună	țară
	Cosine	parte	lună	țară
	Canberra	parte	lună	țară
	Bray-Curtis	parte	lună	țară
SGNS-OP	Euclidean	parte	țară	lună
	Manhattan	parte	țară	lună
	Cosine	parte	lună	țară
	Canberra	parte	țară	lună
	Bray-Curtis	parte	țară	lună
	Correlation	parte	lună	țară

TABLE XX: RODICA-WL: semantic change examples

Word	Change
parte	in the after 1900 corpus is used as "a part of a country" as in the context "în această parte de țară", while in the before 1900 corpus it is used more with the sense of "part of something", e.g., <i>parte a lucrării pentru articolul viitor</i> (part of the work for the future article), or as "to take part in", e.g., <i>boierii luau parte la serviciul divin</i> (the boyars took part in the divine service)
lună	in both corpora the word means "month", as we can see in contexts like "aproape o lună de zile" in the second corpus, and contexts such as "Apare de 2 ori pe lună"
țară	this is again a word whose sense did not change, being used in both corpora with the sense of <i>country</i>

V. DISCUSSIONS

The cosine distance has the overall worst performance on the SEMEVAL-CCOHA corpus regardless of the model, while Canberra has the best performance in detecting semantic change (Table III). When considering the same distance metric, both models perform similarly. We can conclude that the vector space model, i.e., Orthogonal Procrustes and Word Injection, have very little influence over the final word representation for the task of semantic change.

Assessing semantic change for the Romanian Corpora is a non-trivial task since Romanian is usually a low-resource language. Moreover, we can observe that the Basarabian corpus contains bilingual texts. Words and expressions from the Russian language appear on many occasions directly in the texts written with latin script, e.g., *za dva leia* which is the Russian phrases *за два лея* (*for two lei*). These occurrences tend to be rare enough and the performances of the model do not seem to be affected in any significant measure.

Spelling reforms and slight phonetic changes can also affect the performance of the models and are non-trivial problems to overcome when dealing with historic corpora. An example would be that the Romanian word *religie* could be found in the archive with various spellings corresponding to archaic phonemic realizations such as *religiune* or *relighie*. Furthermore, named entities for places also suffer from spelling changes or different naming depending on the region. This leads to the same concept being treated as two separate ones. For

example, *Elisabethgrad* is the name used in the Transylvania corpus for the city *Kropyvnytskyi* or the city of *Kyiv* appears with the outdated spelling *Kiew*.

VI. CONCLUSIONS

In this paper, we propose a new architecture for semantic change detection on low-resource languages. For this architecture, we design three strategies to train our word embeddings. The first two strategies, i.e., SGNS-OP and SGNS-WI, use static word embeddings. The third strategy uses ELMo, a contextual word embedding model. The proposed static and contextual models are trained on the SEMEVAL-CCOHA English dataset. The obtained results act as a baseline for choosing the models to train the Romanian language dataset RODICA.

The experimental results on the English corpus show that the two static word embedding models performs very similarly, while the contextual word embedding model performs poorly in comparison. For the RODICA corpus, we train the static word embeddings and obtain good results for detecting semantic change by using the Word Injection vector space representation.

In future work, we aim to find better ways to handle historical corpora by testing more models and architectures that could accommodate, especially for low-resource languages, the various diachronic differences in orthography. We will also consider other embedding techniques such as FastText [6], GloVe [21], or Hyperbolic Embeddings [19] for the non-contextualized embeddings and BERT [8] for contextualized embeddings. Also, we will research additional metrics that make better use of contextualized embeddings (e.g., ELMo and BERT).

REFERENCES

- [1] F. Armaselu, E.-S. Apostol, C. Chiacros, A. F. Khan, C. Liebeskind, B. McGillivray, C.-O. Truică, and G. V. Oleškevičienė, “Tracing semantic change with multilingual l1od and diachronic word embeddings,” in *International Scientific Interdisciplinary Conference LLOD approaches for language data research and management*, 2022, pp. 41–42.
- [2] F. Armaselu, E.-S. Apostol, A. F. Khan, C. Liebeskind, B. McGillivray, C.-O. Truică, A. Utka, G. V. Oleškevičienė, and M. van Erp, “LL(o)d and NLP perspectives on semantic change for humanities research,” *Semantic Web*, vol. 13, no. 6, pp. 1051–1080, sep 2022.
- [3] F. Armaselu, E.-S. Apostol, A. F. Khan, C. Liebeskind, B. McGillivray, C.-O. Truică, and G. Valūnaitė Oleškevičienė, “Historiae, history of socio-cultural transformation as linguistic data science. a humanities use case,” in *Conference on Language, Data and Knowledge*, 2021, pp. 34:1–13.
- [4] P. Basile, A. Caputo, T. Caselli, P. Cassotti, and R. Varvara, “Diacr-ita@evalita2020: overview of the evalita2020 diachroniclexical semantics (diacr-ita) task,” in *Evaluation Campaign of Natural Language Processing and Speech Tools for Italian*, 2020.
- [5] Y. Bizzoni, M. Mosbach, D. Klakow, and S. Degaetano-Ortlieb, “Some steps towards the generation of diachronic WordNets,” in *Nordic Conference on Computational Linguistics*, 2019, pp. 55–64.
- [6] P. Bojanowski, E. Grave, A. Joulin, and T. Mikolov, “Enriching word vectors with subword information,” *Transactions of the Association for Computational Linguistics*, vol. 5, pp. 135–146, 2017.
- [7] C.-G. Chiru, C.-O. Truică, E.-S. Apostol, and A. Ionescu, “Improving WordNet using word embeddings,” in *International Symposium on Symbolic and Numeric Algorithms for Scientific Computing*, 2021.
- [8] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *Conference of the North American Chapter of the Association for Computational Linguistics*, 2019, pp. 4171–4186.
- [9] H. Dubossarsky, D. Weinshall, and E. Grossman, “Outta control: Laws of semantic change and inherent biases in word representation models,” in *Conference on Empirical Methods in Natural Language Processing*, 2017, pp. 1136–1145.
- [10] K. Ethayarajh, “How contextual are contextualized word representations? comparing the geometry of bert, elmo, and gpt-2 embeddings,” *arXiv preprint arXiv:1909.00512*, 2019.
- [11] B. J. Frey and D. Dueck, “Clustering by passing messages between data points,” *Science*, vol. 315, no. 5814, pp. 972–976, 2007.
- [12] M. Gardner, J. Grus, M. Neumann, O. Taffjord, P. Dasigi, N. F. Liu, M. Peters, M. Schmitz, and L. Zettlemoyer, “AllenNLP: A deep semantic natural language processing platform,” in *Workshop for NLP Open Source Software*, 2018, pp. 1–6.
- [13] M. Giulianelli, M. Del Tredici, and R. Fernández, “Analysing lexical semantic change with contextualised word representations,” in *Annual Meeting of the Association for Computational Linguistics*, 2020, pp. 3960–3973.
- [14] H. Gong, S. Bhat, and P. Viswanath, “Enriching word embeddings with temporal and spatial information,” in *Conference on Computational Natural Language Learning*, 2020, pp. 1–11.
- [15] W. L. Hamilton, J. Leskovec, and D. Jurafsky, “Cultural shift or linguistic drift? comparing two computational measures of semantic change,” in *Conference on Empirical Methods in Natural Language Processing*, 2016, pp. 2116–2121.
- [16] —, “Diachronic word embeddings reveal statistical laws of semantic change,” in *Annual Meeting of the Association for Computational Linguistics*, 2016, pp. 1489–1501.
- [17] V. Kanjirang, S. Mitrovic, A. Antonucci, and F. Rinaldi, “SST-BERT at semeval-2020 task 1: Semantic shift tracing by clustering in bert-based embedding spaces,” in *Workshop on Semantic Evaluation*, A. Herbelot, X. Zhu, A. Palmer, N. Schneider, J. May, and E. Shutova, Eds., 2020, pp. 214–221.
- [18] T. Mikolov, K. Chen, G. Corrado, and J. Dean, “Efficient estimation of word representations in vector space,” in *International Conference on Learning Representations*, 2013, pp. 1–12.
- [19] M. Nickel and D. Kiela, “Poincaré embeddings for learning hierarchical representations,” in *International Conference on Neural Information Processing Systems*, 2017, pp. 6341–6350.
- [20] P. Novotná and V. Blažek, “Glottochronology and its application to the balto-slavic languages,” *Baltistica*, vol. 42, no. 2, pp. 185–210, 2007.
- [21] J. Pennington, R. Socher, and C. Manning, “GloVe: Global vectors for word representation,” in *Conference on Empirical Methods in Natural Language Processing*, 2014, pp. 1532–1543.
- [22] M. E. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer, “Deep contextualized word representations,” in *Conference of the North American Chapter of the Association for Computational Linguistics*, 2018, pp. 2227–2237.
- [23] R. Řehůřek and P. Sojka, “Software Framework for Topic Modelling with Large Corpora,” in *Workshop on New Challenges for NLP Frameworks*, 2010, pp. 45–50.
- [24] J. Rodina, Y. Trofimova, A. Kutuzov, and E. Artemova, “Elmo and bert in semantic change detection for russian,” in *International Conference on Analysis of Images, Social Networks and Texts*, 2020, pp. 175–186.
- [25] H. Rubenstein and J. B. Goodenough, “Contextual correlates of synonymy,” *Communications of the ACM*, vol. 8, no. 10, pp. 627–633, 1965.
- [26] D. Schlechtweg, B. McGillivray, S. Hengchen, H. Dubossarsky, and N. Tahmasebi, “SemEval-2020 task 1: Unsupervised lexical semantic change detection,” in *Workshop on Semantic Evaluation*, 2020, pp. 1–23.
- [27] P. H. Schönemann, “A generalized solution of the orthogonal procrustes problem,” *Psychometrika*, vol. 31, no. 1, pp. 1–10, 1966.
- [28] A. Tsakalidis and M. Liakata, “Sequential modelling of the evolution of word representations for semantic change detection,” in *Conference on Empirical Methods in Natural Language Processing*, 2020, pp. 8485–8497.
- [29] L. Van der Maaten and G. Hinton, “Visualizing data using t-sne,” *Journal of machine learning research*, vol. 9, no. 11, 2008.
- [30] A. Wegmann, F. Lemmerich, and M. Strohmaier, “Detecting different forms of semantic shift in word embeddings via paradigmatic and syntagmatic association changes,” in *International Semantic Web Conference*, 2020, pp. 619–635.
- [31] A. A. Zalizniak, M. Bulakh, D. Ganenkov, I. Gruntov, T. Maisak, and M. Russo, “The catalogue of semantic shifts as a database for lexical semantic typology,” *Linguistics*, vol. 50, no. 3, pp. 633–669, 2012.