

A Computational Analysis of the Dehumanisation of Migrants from Syria and Ukraine in Slovene News Media

Jaya Caporusso^{1,2*}
Boshko Koloski^{1,2*}

Damar Hoogland^{3*}
Matthew Purver^{1,5}

Mojca Brglez^{4*}
Senja Pollak¹

¹Jožef Stefan Institute ²Jožef Stefan International Postgraduate School ⁴University of Ljubljana
Jamova cesta 39, 1000 Ljubljana, Slovenia Ljubljana, Slovenia
jaya.caporusso96@gmail.com {boshko.koloski, senja.pollak}@ijs.si mojca.brglez@ff.uni-lj.si

³Newcastle University ⁵Queen Mary University of London
Newcastle upon Tyne, UK London, UK
d.hoogland2@newcastle.ac.uk m.purver@qmul.ac.uk

Abstract

Dehumanisation involves the perception and/or treatment of a social group's members as *less than human*. This phenomenon is rarely addressed with computational linguistic techniques. We adapt a recently proposed approach for English, making it easier to transfer to other languages and to evaluate, introducing a new sentiment resource, the use of zero-shot cross-lingual valence and arousal detection, and a new method for statistical significance testing. We then apply it to study attitudes to migration expressed in Slovene newspapers, to examine changes in the Slovene discourse on migration between the 2015-16 migration crisis following the war in Syria and the 2022-23 period following the war in Ukraine. We find that while this discourse became more negative and more intense over time, it is less dehumanising when specifically addressing Ukrainian migrants compared to others.

Keywords: news analysis, dehumanisation, migration, social bias

1. Introduction

Dehumanisation is the perception and/or treatment of a certain social group as if its members were *less than human* (Haslam and Stratemeyer, 2016). Negating a shared humanity can lead to strong ingroup-outgroup dynamics (Arcimaviciene and Baglama, 2018) and discriminatory behaviors (Utych, 2018). Dehumanisation and other social biases are reflected and perpetuated through language (Hovy and Prabhumoye, 2021) (i.e., *linguistic bias*). While linguistic bias can occur in any modality, this is particularly concerning for text news articles that reach a large audience.

During the last decade, migration towards the European Union (EU) has significantly increased on two occasions. In the wake of conflicts in Syria in 2015-16, around 1.3m people entered Europe, and around 7.3m since the Russian invasion of Ukraine in 2022 (Moise et al., 2023)¹. Although the former crisis was much smaller in scale than the latter, it was often represented as a threat

to European security (Prideaux de Lacy, 2023), while migration from Ukraine was presented in a more welcoming light (e.g., Dražanová and Geddes, 2022; Koppel and Jakobson, 2023; Tomczak-Boczek et al., 2023; Zawadzka-Paluckta, 2023). Such differences in the discourse on migration from Syria and Ukraine are often accredited to the higher perceived similarity between “us” (typically, EU member state citizens) and the Ukrainian people (Bayoumi, 2022; Paré, 2022).

Natural language processing (NLP) methods offer an efficient means to explore relevant phenomena including linguistic biases (e.g. Bolukbasi et al., 2016), and dehumanisation (Mendelsohn et al., 2020, discussed in detail below). These developments are further driven by pre-trained large language models (LLMs) such as BERT (Devlin et al., 2019), GPT-3 (Brown et al., 2020), and XLNet (Conneau and Lample, 2019). Trained on large datasets, LLMs can successfully capture language structure, and generalise well from few or no examples (Brown et al., 2020; Wei et al., 2022; Kojima et al., 2022), making them well-suited for knowledge transfer in low-resource settings (in our case, the Slovene language), if task-specific data from high-resourced languages is available.

Our objective is to characterise changes in Slovenian public attitudes towards migrants, as presented in news, between the Syria and Ukraine

* These authors contributed equally.

¹The people migrating during these two crises were not solely from Syria and Ukraine; we refer to the crises and periods by reference to the locations of their main triggering reasons (the wars in Syria and Ukraine) rather than making any assumptions about the origin or ethnic identity of the people affected.

migration crisis periods; and for the latter period, to describe differences in attitudes to Ukrainian and other migrants. To do so, we analyse a corpus of news articles published during these two periods using validated computational methods, here extended and adapted to Slovene. We expect to find the following trends:

- *H1a*) attitudes towards migrants became more positive/intense during the Ukraine period compared to the Syria period;
- *H1b*) dehumanising language was more prevalent in the Syria period than the Ukraine period;
- *H2a*) attitudes towards Ukrainian migrants were more positive/intense than those towards non-Ukrainian migrants; and
- *H2b*) dehumanising language was more prevalent in discourse about non-Ukrainian migrants than Ukrainian migrants.

Our main contributions are: 1) adapted computational techniques for analysing dehumanising discourse in Slovene, a less-resourced European language; 2) new public resources, including keyword lists for *moral disgust* and *vermin* concepts, and a Valence, Arousal and Dominance (VAD) sentiment lexicon for Slovene; 3) a new method using anchor vectors and the Kolmogorov–Smirnov test to measure significance of differences in cosine similarities between corpora; 4) an exploration of dehumanisation towards migrants during the Syria and Ukraine migration crises.

This paper is structured as follows. Section 2 addresses related work. Section 3 introduces our data and tools. Section 4 describes the methods and experimental setup. Results are presented in Sections 5 and 6, and discussed, along with conclusion and future work, in Section 7.

2. Related work

While social biases have been widely studied using NLP techniques (Liang et al., 2021), dehumanisation has only rarely been addressed (Wiegand et al., 2021), perhaps because it is challenging to measure it directly (He et al., 2022). However, Mendelsohn et al. (2020) presented an approach based on five characterizing dehumanisation components (Haslam, 2006): negative evaluation of the target group; denial of agency; moral disgust; likening the target group to something non-human; and psychological distancing and denial of subjectivity. They measured the first two using sentiment detection: employing a lexicon-based approach (Mohammad, 2018) and connotation frames (Rashkin et al., 2015), they measured valence and dominance, which, together

with arousal (i.e., VAD), are considered the three most important dimensions of sentiment (Osgood et al., 1957). Valence represents the continuum between pleasure and displeasure, arousal between engaging and non-engaging, and dominance between control and submission of the experiencer of an affective state (Russell and Mehrabian, 1977). To measure the second two elements of dehumanisation, moral disgust and metaphorisation through non-human concepts, Mendelsohn et al. (2020) employed distributional semantic methods: they measured the cosine similarity of the target group to the concepts of *moral disgust* and *vermin*, showing that this can capture interpretable patterns in the discourse on the LGBTQIA+ community in the New York Times; and compared different time periods by constructing and analysing a separate word embedding space for each one. In this study, we build on their work by modifying and applying it to the discourse on migration in Slovene.

The discourse on migration frequently expresses an “us” and “them” dichotomy (Vezovnik, 2018; van Dijk, 2018; Chitrakar, 2020) and portrays migrants as a threat. Migration narratives are also known for their persistent use of mechanistic and animalistic metaphors, equating migrants to water, animals, or commodities (Taylor, 2021). Recent studies found *vermin* metaphors to be particularly dominant in anti-immigration online discourse (Šori and Vehovar, 2022).

3. Data and resources

We base our approach on that of Mendelsohn et al. (2020), but address some shortcomings—such as the lack of an arousal analysis and the use of only lexicon-based sentiment analysis—and extend their work in a number of ways. Specifically, we add arousal analysis and neural models for valence and arousal analysis (the latter aims to compensate for the possible shortcomings of the lexicon-based approach, as addressed in Section 3.2.4); introduce a novel inferential anchoring procedure allowing comparison of any two corpora with a shared vocabulary, without the need for vector spaces to be explicitly aligned or share parameters/dimensions; apply the framework to a less-resourced language and a new domain; and, via the novel inferential procedure and cross-lingual valence-arousal model, make it significantly easier to transfer to new datasets.

We first introduce the corpora used in Section 3.1. In Section 3.2, we explain the construction of the embeddings model, the construction of the term lists and their concept vectors, and we provide the details of the tools and models used for sentiment analysis.

3.1. Corpora

We use two corpora of Slovene news, each corresponding to a large-scale migration time period. Both corpora were obtained by one of Slovenia’s largest media monitoring companies, and constructed by selecting news articles from the online publications of 29 Slovene media outlets. The first corpus (C_{syr}) contains articles published following the war in Syria and the subsequent migration from August 2015 to April 2016, and the second corpus (C_{ukr}) contains articles published during the war in Ukraine from February 2022 to March 2023.

The corpora were constructed by selecting articles including the following migration-related keywords: *begunec**, *begunc**, *begunk**, *beguns**, *migracij**, *migrant**, *imigra**, *prebežni**, *pribežni**, *prebežni**, and *azil**. These are *unbiased*, almost synonymous terms corresponding to the concepts of *migrant* and *refugee* in English. They were taken from a larger list of migration-related keywords used in previous studies on Slovene (e.g. [Evkoski and Pollak, 2023](#)), with only the most general terms selected (avoiding, e.g., terms referring to specific nationalities and more loaded terms).

We report descriptive statistics in Table 1.

Statistic	C_{ukr}	C_{syr}
documents	8 470	8 556
sentences	311 185	338 759
paragraphs	137 164	132 934
total words	8 785 219	8 282 229
unique words	237 622	189 512
total lemmas	8 785 907	8 282 481
unique lemmas	100 895	77 927
words per doc.	1 037.22	968.00
words per sent.	28.23	24.45
word per par.	92.19	106.95

Table 1: Dataset statistics.

Note that the two corpora do not contain news only pertaining to migrants from Syria and Ukraine. In particular, in C_{ukr} , almost half (49.8%) of the articles do not contain mentions of Ukraine or Ukrainians. For this reason, we further split C_{ukr} into subcorpora of paragraphs (defined as text surrounded by /n/n) mentioning Ukraine (S_{ukr}) and paragraphs not mentioning Ukraine (S_{oth}), and our analyses include comparisons on the corpus and subcorpus levels.

3.2. Analysis Tools

The application of [Mendelsohn et al. \(2020\)](#)’s method requires the use of a word embedding model, vector representations of selected concepts, and a method to infer sentiment. For the latter, we investigate both lexicon- and classifier-based methods.

3.2.1. Static Embeddings

To build a static word embedding model, we first pre-process the corpora to remove titles, and segment text into paragraphs, following [Mendelsohn et al. \(2020\)](#). We then apply the CLASSLA tools for South Slavic languages ([Ljubešić and Dobrovoljc, 2019](#)) for tokenisation and lemmatisation, followed by training a Word2Vec model ([Mikolov et al., 2013](#)) for each (sub)corpus. To allow for direct comparison of vector spaces, we align the neighbourhoods of the individual models (following [Kim et al., 2014](#)) by initialising them from the unlemmatised kontekst.io pre-trained model for Slovene.² We therefore lemmatise the words in that vocabulary using the word-based LemmaGen3 lemmatiser ([Juršič et al., 2010](#)) and average the embeddings of any repeated lemmas. This reduces our vocabulary by c.58% from **572,261** word forms to **242,262** lemmas. We train distinct models on the sentences for each corpus C_{syr} and C_{ukr} . Next, we train a model for each subcorpus S_{ukr} and S_{oth} . Only sentences containing more than two words are considered. We set the min-count to 1 and train the model for 50 epochs.

3.2.2. Concept vector construction

We introduce three concept lists that are used to construct the concept embedding vectors, employed in the cosine similarity-based analyses described in Section 4.1.

Migrant terms To select the words forming the *migrant* vector (MV) representing the concept of migrant in the embedding space(s), we start from the list of search terms used to construct the corpora and derive their lemmas (e.g., the search term *migrant** can capture three lemmas, masculine noun *migrant*, feminine noun *migrantka*, and adjective *migrantski*). For our final concept list, we exclude feminine forms as these are rarely present in the corpora and nearly exclusively used in C_{ukr} . We also exclude migration-related adjectives and abstract nouns, because these are likely to refer to non-human migration, and can also be used in inherently dehumanising syntagms such as ‘*migrantski val*’ (migrant wave). The final list of the words for the MV is:

- (1) *migrant*, *imigrant*, *begunec*, *azilant*, *prebežnik*, *pribežnik*. (English: *migrant*, *immigrant*, *refugee*, *asylee*, *fugitive*, *escapee*).

²<https://kontekst.io/>

Moral disgust terms For the *moral disgust* vector (DV), we translate and further select (based on term-frequency analysis) terms identified by [Graham et al. \(2009\)](#). Because of the Covid-19 epidemic and the possible effects of its occurrence on the semantics of the models, we exclude disease-related terms. The final list of 76 terms for the DV includes, among others:

- (2) *skrunstvo, nečist, zamazanost, prostitut, grešnica, nezmeren* (English: *desecration, unclean, filthiness, prostitute, sinner, intemperate*)

The complete list of moral disgust terms is available in Appendix A.

Vermin terms Vermin metaphors are a prevalent feature of dehumanising, exclusionary, and racist discourse, and act as the dominant metaphor in offensive anti-immigrant comments ([Šori and Vehovar, 2022](#)). We thus also measure dehumanisation through this particular metaphor. To later construct a *vermin* concept vector (VV), we collect vermin-related terms by translating the list of terms used by [Mendelsohn et al. \(2020\)](#), based on previous metaphor studies. Our final list of terms is:

- (3) *golazen, žužek, roj, termit, parazit, zajedavec, glocavec, miš, vampir, kobilica, ščurek, gnida, uš, pršica, bolha, pijavka, podgana, krvoses, osa, škodljivec, mravlja, komar, žuželka* (English: *vermin, bug, swarm, termite, parasite, rodent, mouse, bloodsucker/vampire, locust, cockroach, louse egg, louse, mite, flea, leech, rat, bloodsucker, wasp, pest, ant, mosquito, insect*).

Concept vector construction Based on these concept lists, we build MV, DV, and VV by taking the average of individual word vectors weighted by their frequency in the corpora.

3.2.3. Slovene Valence, Arousal and Dominance Lexicon

For Slovene, there is no VAD lexicon comparable to the NRC English VAD ([Mohammad, 2018](#)) used by [Mendelsohn et al. \(2020\)](#). To replicate their valence and dominance analysis, we therefore adapt the English lexicon to Slovene. First, we take the Slovene part of the LiLaH lexicon ([Daelemans et al., 2020](#)), a manually validated translation of the NRC Emotion lexicon ([Mohammad and Turney, 2013, 2010](#)) containing c.14,000 words with binary values for positive/negative sentiment and 8 basic emotions; for these, we map the English VAD scores directly. We then extend this with 5,931 entries not present in LiLaH, translating

them using sloWNet ([Fišer, 2015](#)). If no mapping is found, we retain the translation in the machine-translated Slovenian version of the NRC-VAD lexicon. The final resource³ contains 19,998 entries with real-valued VAD scores and binary values of sentiment and emotion association.

3.2.4. Zero-shot Cross-Lingual VA Detection

While the lexicon-based approach above is likely to have high precision, it may have low recall, particularly given the transfer to Slovene. Furthermore, it does not capture contextual cues such as word sense, part of speech, and negation ([Mohammad, 2020](#)). We therefore also use a machine-learning-based approach; given the lack of relevant resources in Slovene, we derive this via cross-lingual transfer of an existing model. [Mendes and Martins \(2023\)](#) provide VAD models fine-tuned on 34 datasets from 18 languages (not including Slovene). They investigated three custom losses—Mean Square Error, CCCL, and robust loss—to leverage effective learning through the datasets and RoBERTa family models. The best-performing model was *XLM-Roberta-large* ([Conneau and Lample, 2019](#)). The model checkpoint was made available by the authors. We use HuggingFace ([Wolf et al., 2020](#)) to infer with the paragraph-level inputs to provide Valence and Arousal (VA). We apply this in a zero-shot cross-lingual transfer setting; although the fine-tuning for VA output used no Slovene data, the underlying multilingual language model includes Slovene.

4. Methods

We describe changes in dehumanising attitudes towards migrants expressed in different subcorpora using two methodologies, one based on Word2Vec vector space similarities (Section 4.1) and the other on sentiment analysis (Section 4.2).

4.1. Vector-based similarity analysis

We use our Word2Vec embeddings (see Section 3.2.1 above) to analyse the differences in the latent representation of the concept of migrant between the corpora.

4.1.1. Nearest neighbour analysis

In each (sub)corpus, we first extract the top k nearest neighbours (NN) for the MV (see Section 3.2.2), excluding words with the same root as the words used to construct MV. This allows for qualitative inspection of terms (see Section 5.1.1) and functions

³The lexicon is publicly available via CLARIN.SI at <http://hdl.handle.net/11356/1875>.

as input for the sentiment analysis of the corpus-specific NNs lists (see Section 4.2.2).

4.1.2. Similarity of migrants to moral disgust and vermin vectors

We compare the cosine similarities (CS) of the concept pairs MV-DV and MV-VV in each subcorpus, and perform a statistical test to assess the significance of the differences in distributions.

Moral disgust Following Mendelsohn et al. (2020), we assess whether migrants are described with greater or lesser degrees of moral disgust in two corpora by comparing the MV-DV similarity; we do this between the two corpora C_{ukr} and C_{syr} , and between the two subcorpora S_{ukr} and S_{oth} .

Dehumanising metaphors analysis Similarly, we assess the change in the use of dehumanising metaphorical language by comparing the MV-VV similarity; again we compare both C_{ukr} vs. C_{syr} and S_{ukr} vs. S_{oth} .

Anchoring Kolmogorov-Smirnov test for comparison of neighbourhoods To assess whether the difference between corpora in MV-DV or MV-VV similarity is statistically significant, we develop and apply the following novel anchoring procedure, which can be applied without relying on exact alignments between embedding spaces.

First, we take a selection S of 1000 random words w_i from the common vocabulary of the two corpora. Next, we use the MV and DV/VV as anchors v , denoted by v_{mv} and v_{dv} , and calculate their distance to each randomly selected word w_i of S as $d(w_i, v) = \cos(w_i, v)$, obtaining two vectors that represent each of the two anchors as their distance to each word in S : $a_{mv} = [d(w_1, v_{mv}), d(w_2, v_{mv}), \dots, d(w_N, v_{mv})]$, $a_{dv} = [d(w_1, v_{dv}), d(w_2, v_{dv}), \dots, d(w_N, v_{dv})]$. We then calculate the distance between these two anchor vectors as $\tilde{d} = a_{mv} - a_{dv}$.

We repeat this process for the two corpora to obtain two sets of distances between the anchors, $d_{corpus1}$ and $d_{corpus2}$. Finally, we apply the Kolmogorov-Smirnov test to assess if $d_{corpus1}$ and $d_{corpus2}$ originated from the same distribution or not, i.e., represent the semantic similarities between the MV and DV/VV as the same or not. We use a conventional $\alpha=0.05$ to draw inferences.

4.2. Sentiment analysis

We use two approaches to analyse the differences in the sentiment expressed in the corpora.

4.2.1. Lexicon and transformer approaches

We apply the lexicon and zero-shot multilingual VA models introduced in Section 3. We obtain sentiment scores expressing valence and arousal levels on a scale from 0 to 1 for each paragraph in two ways. Following Mendelsohn et al. (2020), we use our adapted Slovene VAD lexicon to calculate the score of a paragraph by taking the average score over words; here, for each valence and arousal. Due to the highly inflectional nature of the Slovene language, we use the lemmatised version of the corpus. We also employ the ML-based model presented in Section 3.2.4. In this approach, VA scores for each paragraph are predicted from the unlemmatised text. In the initial comparison of the two approaches, we exclude paragraphs with less than 20% coverage by the NRC lexicon and paragraphs of 15 or fewer words and 500 or more words. We also perform a qualitative analysis of the 20 paragraphs with the highest and lowest scores, to determine which method captures paragraph-level sentiment more successfully.

4.2.2. Hypothesis testing

Paragraph-level VA analysis To highlight the variations in the attitudes reflected in news reports across the two selected time periods, and between Ukrainian and non-Ukrainian migrants in the second period, we analyse valence and arousal on the paragraph level. We only include paragraphs between 15 and 500 words, with at least five unique words and at least one of the migrant terms (see Section 3.2.2). We compare the VA scores from the method that most accurately describes sentiment according to our qualitative analysis.

Nearest neighbour VAD analysis We also look at sentiment on a word level. In each subcorpus, we first extract the top k nearest neighbours (NNs) of the MV, disregarding words with the same root as the terms used to construct MV. We compare the 20 NNs of the MV across corpora both qualitatively and quantitatively—the latter, by comparing the sentiment scores (valence, arousal, and dominance) of the 500 NNs using the NRC lexicon.

Bayesian Hypothesis Testing We apply Bayesian Hypothesis Testing to assess the difference between distributions of VA scores of paragraphs and NN. For each comparison, we adopt normally distributed priors $\mathcal{N}(\mu = 0.5, \sigma^2 = 0.25)$. We assume that the standard deviation of the data is a half-normal distribution with $\sigma^2=0.25$ and model the data as truncated normal distributions since sentiment scores are defined in $[0, 1]$. We use Markov Chain Monte Carlo sampling, drawing

5,000 samples from the posterior distributions after an initial tuning phase of 1,000 samples. We assess modelled probabilities of the difference in means, effect sizes, and credible intervals (CI).

5. Results: Syria and Ukraine periods

In this section, we present our comparison of C_{ukr} and C_{syr} .

5.1. Vector-based similarity analysis

5.1.1. Nearest neighbours analysis

We qualitatively analyse the top 20 NNs of each of the C_{syr} and C_{ukr} MVs. The top NNs show a very similar pattern, including *človek*, *tujec*, *oseba*, *prišlek*, *prosilec* (English: *human*, *foreigner*, *person*, *newcomer*, *applicant*). However, while the NNs of MV in C_{syr} contain country names, unique human concepts appear closer to the C_{ukr} vector, such as *sirota*, *pacient*, *ilegalec*, *bolnik*, *družina* (English: *orphan*, *patient*, *illegal*, *patient*, *family*), but also a non-human concept, *žival* (*animal*).

5.1.2. Similarity of migrants to moral disgust and vermin vectors

Moral disgust The MV-DV CS in C_{syr} (-0.036) is lower than in C_{ukr} (0.033). This difference is statistically significant ($k=.072$, $p.011$), indicating that *migrant* becomes semantically closer to *moral disgust* in the Ukraine period, which implies a rising trend of migrant dehumanisation. This differs from the hypothesised direction (H1b).

Dehumanising metaphors analysis The MV-VV similarity in C_{ukr} (.100) is higher than in C_{syr} (.064). This difference is significant ($k=.122$, $p<.001$). Similarly to moral disgust, this indicates that the representation of migrants through vermin-related dehumanising metaphors increases in the Ukraine period, again contrary to hypothesis H1b.

5.2. Sentiment analysis

In Section 5.2.1, we describe the quantitative and qualitative comparisons of the two sentiment detection approaches. This allows us to select the approach to be used for sentiment analysis and statistical testing of results, as presented in Sections 5.2.2 and 6.2.2.

5.2.1. Lexicon and transformer approaches

Quantitative comparison of VA methods To compare the two sentiment detection approaches (VAD lexicon and pre-trained cross-lingual model),

we first compare the overall distributions of VA scores of paragraphs as assessed by the two approaches. As illustrated by Figure 1, we note that the valence score distributions obtained using the pre-trained model are wider than the distribution obtained using the lexicon. A similar distribution is observed for arousal scores. The difference in distributions indicates that the scores obtained by the model capture the sentiment expressed in paragraphs in a more fine-grained manner, while the lexicon-based scores all converge around some average score. This global view of VA distributions promotes the use of the pre-trained model over the use of the lexicon.

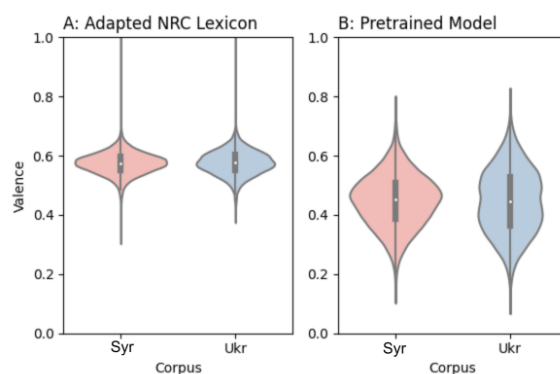


Figure 1: Distributions of valence scores for C_{syr} and C_{ukr} according to the lexicon approach (A) and the transformer model (B).

Qualitative analysis of VA methods To evaluate the valence scores obtained by the two methods, we also manually evaluate 20 paragraphs per subcorpus, including the top-10 with the highest valence and the top-10 with the lowest valence for each of the two approaches. The evaluation is conducted by a Slovene native speaker. They address two aspects of sentiment. In the first step, they assess whether the analysed paragraph presents a positive, negative, or neutral attitude towards migrants (aspect-based sentiment); in the second step, whether the overall sentiment of the paragraph is positive, negative, or neutral (general sentiment).

Overall, the qualitative evaluation of the highest-valenced (i.e., most positive) paragraphs indicates that both approaches perform well (specifically, paragraphs with a positive valence towards migrants using the lexicon approach: 14/20; paragraphs with a positive valence towards migrants using the crosslingual XLMRoberta approach: 20/20). A common theme in these paragraphs is the expression of support, empathy, and solidarity towards migrants, as shown in Exam-

ple 1.

Example 1: *To so torej obrazi ljudi, prostovoljcev, ki nesebično pomagajo beguncem iz dneva v dan in jim s tem vivajo upanje v nov in boljši jutri. 'So these are the faces of people, volunteers, who selflessly help refugees day by day, giving them hope for a new and better tomorrow.'*

In the qualitative analysis of paragraphs with the lowest valence, the picture is a little less clear. Only 9 out of 20 and 7 out of 20 paragraphs for the lexicon-based and model-based methods, respectively, are actually negative towards migrants. In these, a prominent common theme is crimes committed by migrants, including passages that depict migrants in an animalistic manner, arguing for their lack of respect for property, cleanliness, and order.

On the other hand, many of the most negative paragraphs do not necessarily communicate a negative, dehumanising attitude towards migrants. In 3 out of 20 and 7 out of 20 for the lexicon-based and model-based approach, respectively, a common theme is the bad conditions and poor treatment of migrants, and the causes of migration which use negatively-valenced words such as *vojna, slabo, izgubiti, trpljenje* (English: *war, bad, to lose, suffering*). While the topics or events described by the paragraphs are indeed negative, they still communicate a positive attitude towards migrants, accompanied by expressions of support, empathy, and solidarity towards them.

Although the manual sentiment annotation revealed that migrants are not necessarily negatively evaluated in negatively valenced paragraphs, we find a general trend concerning the dehumanising treatment of migrants. Specifically, while the language is not used to directly dehumanise migrants, it often describes the inhumane and degrading conditions they experience. Our qualitative analysis shows that neither resource accurately captures sentiment expressed *towards* migrants; however, the neural-model-based approach better captures general sentiment by accounting for the wider context. For these reasons, we use the model predictions in all our subsequent analyses.

5.2.2. Hypothesis testing

Paragraph-level VA analysis Hypothesis H1a predicts that news about migrants is more positive and more intense in the Ukraine period compared to the Syria period, meaning higher valence and arousal scores in C_{ukr} than in C_{syr} . This prediction is not borne out with regard to valence: by a small margin of .002, valence in C_{ukr} is lower (mean=.446, sd=.108) than in C_{syr} (mean=.449, sd=.089). We find weak evidence that this difference reflects a true population difference, with

a probability of .02, a CI close to zero (-.005, -.000), and a posterior Cohen's d of -.024 (CI: -.048, -.001). However, arousal is higher in C_{ukr} (mean=.480, sd=.054) than in C_{syr} (mean=.466, sd=.056) by .013. We find strong evidence that this reflects a robust population difference with a probability of 1.00 (CI: .012, .015) and a posterior Cohen's d of .245 (CI: .223, .267).

Nearest neighbours VAD analysis We compare the sentiment scores between the 500 NNs of the MV in C_{ukr} and C_{syr} corpora. The NRC lexicon provides only limited coverage of the NNs lists; for the C_{ukr} corpus, only 22.8% of the words are in the NRC lexicon, and for C_{syr} , only 26.4%.

- **Valence** - The NNs of the C_{ukr} MV are higher in valence (mean=.145, sd=.073) than those of the C_{syr} MV (mean=.134, sd=.065) by .010. We find no evidence that this reflects a true population difference, with a high probability of .73, but the CI for this parameter straddling zero (-.018, .032; Posterior Cohen's d: .097 with CI: -.234, .409)
- **Arousal** - The NNs of the C_{ukr} MV (mean=.125, sd=.046) are higher in arousal than those of the C_{syr} MV (mean=.107, sd=.046) by .018. We find evidence that this reflects a true population difference, with a high probability 1.00 (CI: .006, .033) and a posterior Cohen's d: .399 (CI: .132, .666)
- **Dominance** - The NNs of the C_{ukr} MV (mean=.134, sd=.055) are higher in dominance than those of the C_{syr} MV (mean=.121, sd=.054) by .013. We find no evidence that this reflects a true population difference, with a high probability at .95 but a CI that straddles zero (-.002, .030) and a posterior Cohen's d: .239 (CI: -.040, .512).

6. Results: Ukraine sub-corpora

In this section, we present our comparison of the subcorpora of news articles produced during the Ukrainian migration crisis—including, respectively, articles mentioning (S_{ukr}) and not mentioning (S_{oth}) Ukraine.

6.1. Vector-based similarity analysis

6.1.1. Nearest neighbours analysis

The top 10 NNs of the two MVs in S_{ukr} and S_{oth} have slightly different orders of similarity. We first ensure that each subcorpus corresponds to a different nationality group by verifying that *Ukrainian* appears only in the S_{ukr} MV neighbourhood and *Kurd* appears only in the S_{oth} MV neighbourhood.

Second, while the first two NNs of MV in S_{ukr} are *človek* (*human*) and *prosilec* (*applicant*), the terms closer to MV in S_{oth} are *priseljensec* (*immigrant, settler*) and *tujec* (*foreigner*), implying that media frames Ukrainian migrants as less foreign or alien compared to other nationalities. Moreover, NNs of MV in S_{ukr} present a higher number of human roles (e.g., *otrok*, *sirota*, *učenec*, *študent*; English: *child, orphan, pupil, citizen, student*), while NNs of MV in S_{oth} include more impersonal roles *potnik*, *civilist*, *ilegalec* (*passenger/traveler, civilian, illegal*). However, the top 10 NNs of the S_{ukr} MV do include a very dehumanising term: *žival* (*animal*).

6.1.2. Similarity of migrants to moral disgust and vermin vectors

Moral disgust The MV-DV CS in S_{oth} (.068) is larger than in S_{ukr} (.038). This difference is significant ($k=.095$, $p<.001$). This result indicates that news articles published during the period of the war in Ukraine communicate less moral disgust when they address Ukrainian migrants compared to when they address migrants of other nationalities—as hypothesised (H2b).

Dehumanising metaphors analysis The MV-VV CS in S_{oth} (.056) is smaller than in S_{ukr} (.162). This is a non-significant difference ($k=.05$, $p=.164$), pointing in the opposite direction from what was hypothesised (H2b). Namely, we found no evidence to support the hypothesis that Ukrainian migrants are less associated to dehumanising metaphors than migrants of other nationalities.

6.2. Sentiment analysis

6.2.1. Lexicon and zero-shot approaches

Based on our findings in 5.2.1, we use the pre-trained transformer model to analyse sentiment in S_{ukr} and S_{oth} and omit the lexicon-based results.

6.2.2. Hypothesis testing

Paragraph-level VA analysis Within C_{ukr} , we find that the paragraphs in S_{ukr} have a higher valence (mean=.469, sd=.108) than in S_{oth} (mean=.432, sd=.106) by .036. We also find strong evidence that this reflects a population difference, with the probability of the two means being different at 1.00 (CI: .032, .040) and a posterior Cohen’s d of .335 (CI: .295, .373). The results support our hypothesis H2a, i.e., that attitudes towards Ukrainian migrants are more positive and intense than those towards other nationalities.

Nearest neighbours VAD analysis We compare the valence scores between the 500 NNs of

the MV in S_{ukr} and S_{oth} . As in section 5.2.2, the lexicon provides only limited coverage: only 34.4% of the words from S_{ukr} and 19.6% of the words from S_{oth} are present in the NRC lexicon. Our analyses show that the NNs of S_{ukr} MV are higher in valence and dominance and lower in arousal than those of S_{oth} MV. However, as in Section 5.2.2, our statistical tests find no notable difference between the two sub-corpora in any of the three sentiment dimensions.

7. Discussion and Conclusion

We extend the framework of Mendelsohn et al. (2020) and apply it to the investigation of dehumanisation of migrants in Slovene news articles in the periods of the 2015-16 and the 2022-23 migration crises. Specifically, we extend their work in the following innovative ways: adding arousal analysis to the original valence-only approach; using neural models for valence and arousal (instead of solely lexicon-based approaches); introducing a novel inferential anchoring procedure allowing comparison of any two corpora with shared vocabulary, without the need for vector spaces to be explicitly aligned or share parameters/dimensions; applying the method to a less-resourced language and a new domain; and, via the novel inferential procedure and cross-lingual valence-arousal model, making the method significantly easier to transfer to new datasets.

Our analysis of linguistic correlates of dehumanisation in the news articles from Syria (C_{syrr}) and Ukraine (C_{ukr}) periods show the following. Concerning hypothesis H1a (“attitudes towards migrants became more positive/intense during the Ukraine period compared to the Syria period”), contrary to our expectations, valence appears to be significantly higher in C_{syrr} ; this is supported by the paragraph-level but not by the NNs valence analysis, as the latter did not show any significant differences across corpora. However, the arousal appears to be higher for C_{ukr} , in line with the part of H1a concerning intensity; this is supported by both the paragraph-level and the NNs arousal analysis. We interpret this pattern as tentative evidence that attitudes towards migrants expressed in news articles have become less positive and more intense during the Ukrainian period.

When looking closer at C_{ukr} , we observe that valence and arousal are both significantly higher in paragraphs that mention Ukraine (S_{ukr}) than in paragraphs that do not (S_{oth}), as put forward by our hypothesis H2a (“attitudes towards Ukrainian migrants were more positive/intense than those towards non-Ukrainian migrants”).

We also observe that the mean valence of S_{oth} is lower than that of C_{syrr} (.432 and .449 respec-

tively), which is somewhat contrary to the conclusions of Moise et al. (2023): namely, that positive attitudes to migrants from Ukraine “spill over” to attitudes to migrants from other origins.

The hypotheses H1b) (“*dehumanising language was more prevalent in the Syria period compared to the Ukraine period*”) and H2b) (“*dehumanising language was more prevalent in discourse about non-Ukrainian migrants than Ukrainian migrants*”) are analysed from three perspectives. First, in terms of **Denial of agency**, the NNs dominance analysis does not find any statistically significant difference in any of our corpora, meaning that we do not detect any difference in the degree to which the agency of migrants was denied. However, as Mendelsohn et al. (2020) also pointed out, this measure of denial of agency is limited in that it does not capture sentence-level information about whose agency is being denied/affirmed. This method may thus be insufficient to detect this aspect of dehumanising language use. Next, in terms of **Moral disgust**, contrary to H1b), the concept of migrant appears to be closer to the concept of moral disgust in C_{ukr} . However, our analyses on the Ukrainian period sub-corpora confirm H2b), as the concept of migrant is closer to the concept of moral disgust in S_{oth} . Concerning **Dehumanising metaphor**, contrarily to our expectations (H1b), the concept of vermin is closer to the concept of migrant in C_{ukr} . The higher level of dehumanisation in the later analysed period might relate to increased migratory movements and a context of a general crisis of the EU (Bello, 2022). At the same time, Schmidt-Catran and Czymara (2023) argue that higher migratory influxes are not related to a more negative perception of migrants, unlike exclusionary discourses by political elites that are influencing negative attitudes. The results concerning S_{oth} and S_{ukr} , although not statistically significant, showed a tendency in the direction opposite to what was formulated in H2b, suggesting that the concept of vermin is closer to migrants from Ukraine.

In conclusion, our results show that while news discourse seems to dehumanise migrants more and more, it does so in a selective way. While the general trend of greater dehumanisation holds for migrants in general, Ukrainian migrants are dehumanised to a lesser extent than other migrants. This confirms previous studies’ findings (e.g., Dražanová and Geddes 2022), supporting that we perceive and treat Ukrainians differently than other migrants due to their higher perceived similarity to “us” (Bayoumi, 2022; Paré, 2022).

In future work, we plan to investigate target-based sentiment analysis. Additionally, we aim to investigate differences in the discourse around different terms; in the context of immigration, for

example, Zawadzka-Palucktau (2023) focused on the terms *refugee* and (*im*)*migrant*, finding the latter “less legitimising”. Finally, we will apply the dehumanisation analysis framework to investigate the differences between different newspaper sources.

Data Availability

The resources, i.e., the list of terms used to build concept vectors and the Slovenian emotion and VAD lexicon, are made available, respectively, in Appendix A and on CLARIN.SI.

Acknowledgements

We acknowledge the financial support from the Slovenian Research Agency (ARRS) core research programs Knowledge Technologies (P2-0103) and ‘Slovene Language – Basic, Contrastive, and Applied Studies’ (P6-0215), and from the projects CANDAS (Computer-assisted multilingual news discourse analysis with contextual embeddings, No. J6-2581), SOVRAG (Hate speech in contemporary conceptualizations of nationalism, racism, gender and migration, No. J5-3102), and EMMA (Slovenian Research and Innovation Agency research project Embeddings-based techniques for Media Monitoring Applications, No. L2-50070). The work of DH was partially supported by Northern Bridge (UK AHRC); MP by Sodestream (UK EPSRC EP/S033564/1); and ARRS supported the work of MB within the national research programme ‘Slovene Language – Basic, Contrastive, and Applied Studies’ (P6-0215), and BK by the Young Researcher Grant (PR-12394).

Limitations

This paper addresses the dehumanisation of migrants in Slovenian news media using a computational, quantitative methodology. A limitation of our work lies in our reliance on translated English lexis for the study of dehumanisation, which may not perfectly match the context in Slovene. A more comprehensive evaluation of dehumanising elements in language would necessitate the application of qualitative approaches from critical discourse analysis, examining the wider context in texts and discourses. In turn, this could inform new applications of perhaps more suitable quantitative methods. Furthermore, our sentiment detection approaches did not specifically evaluate sentiment towards the target group. This aspect could in the future be addressed with aspect-/target-based sentiment analysis tools. Finally, the transformer VAD models did not contain any Slovene training

data, and therefore their accuracy could be sub-optimal.

8. Bibliographical References

- Liudmila Arcimaviciene and Sercan Hamza Baglama. 2018. Migration, metaphor and myth in media representations: The ideological dichotomy of “them” and “us”. *Sage Open*, 8(2):2158244018768657.
- Moustafa Bayoumi. 2022. They are ‘civilised’ and ‘look like us’: the racist coverage of ukraine. *The Guardian*, 2.
- Valeria Bello. 2022. Prejudice and cuts to public health and education: A migration crisis or a crisis of the european welfare state and its socio-political values? *Societies*, 12(2):51.
- Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. 2016. [Man is to computer programmer as woman is to homemaker? debiasing word embeddings](#). In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS’20*, Red Hook, NY, USA. Curran Associates Inc.
- Rok Chitrakar. 2020. [Threat perception in online anti-migrant speech: a slovene case study](#). *Slovenščina 2.0: empirical, applied and interdisciplinary research*.
- Alexis Conneau and Guillaume Lample. 2019. [Cross-lingual language model pretraining](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Lenka Dražanová and Andrew Geddes. 2022. Attitudes towards ukrainian refugees and governmental responses in 8 european countries. In *EU Responses to the Large-Scale Refugee Displacement from Ukraine: An Analysis on the Temporary Protection Directive and Its Implications for the Future EU Asylum Policy*, pages 135–147. European University Institute.
- Bojan Evkoski and Senja Pollak. 2023. [XAI in Computational Linguistics: Understanding Political Leanings in the Slovenian Parliament](#), pages 56–61.
- Jesse Graham, Jonathan Haidt, and Brian A Nosek. 2009. Liberals and conservatives rely on different sets of moral foundations. *Journal of personality and social psychology*, 96(5):1029.
- Nick Haslam. 2006. Dehumanization: An integrative review. *Personality and social psychology review*, 10(3):252–264.
- Nick Haslam and Michelle Stratemeyer. 2016. Recent research on dehumanization. *Current Opinion in Psychology*, 11:25–29.
- Yulan He, Heng Ji, Yang Liu, Sujian Li, Chia-Hui Chang, Soujanya Poria, Chenghua Lin, Wray Buntine, Maria Liakata, Hanqi Yan, et al. 2022. Findings of the association for computational linguistics: Aacl-ijcnlp 2022. In *Findings of the Association for Computational Linguistics: AACL-IJCNLP 2022*.
- Dirk Hovy and Shrimai Prabhumoye. 2021. Five sources of bias in natural language processing. *Language and Linguistics Compass*, 15(8):e12432.
- Matjaz Juršič, Igor Mozetic, Tomaz Erjavec, and Nada Lavrac. 2010. Lemmagen: Multilingual lemmatisation with induced ripple-down rules. *Journal of Universal Computer Science*, 16(9):1190–1214.
- Yoon Kim, Yi-I Chiu, Kentaro Hanaki, Darshan Hegde, and Slav Petrov. 2014. [Temporal analysis of language through neural language models](#). In *Proceedings of the ACL 2014 Workshop on Language Technologies and Computational Social Science*, pages 61–65, Baltimore, MD, USA. Association for Computational Linguistics.

- Takeshi Kojima, Shixiang (Shane) Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. In *Advances in Neural Information Processing Systems*, volume 35, pages 22199–22213.
- Katrina Koppel and Mari-Liis Jakobson. 2023. Who is the worst migrant? migrant hierarchies in populist radical-right rhetoric in estonia. In *Anxieties of Migration and Integration in Turbulent Times*, pages 225–241. Springer International Publishing Cham.
- Paul Pu Liang, Chiyu Wu, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2021. Towards understanding and mitigating social biases in language models. In *International Conference on Machine Learning*, pages 6565–6576. PMLR.
- Nikola Ljubešić and Kaja Dobrovoljc. 2019. [What does neural bring? analysing improvements in morphosyntactic annotation and lemmatisation of Slovenian, Croatian and Serbian](#). In *Proceedings of the 7th Workshop on Balto-Slavic Natural Language Processing*, pages 29–34, Florence, Italy. Association for Computational Linguistics.
- Julia Mendelsohn, Yulia Tsvetkov, and Dan Jurafsky. 2020. A framework for the computational linguistic analysis of dehumanization. *Frontiers in artificial intelligence*, 3:55.
- Gonçalo A. Mendes and Bruno Martins. 2023. Quantifying valence and arousal in text with multilingual pre-trained transformers. In *ECIR 2023: Advances in Information Retrieval*, pages 84–100, Cham. Springer Nature Switzerland.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Saif M. Mohammad. 2020. [Practical and ethical considerations in the effective use of emotion and sentiment lexicons](#).
- Alexandru D. Moise, James Dennison, and Hanspeter Kriesi. 2023. [European attitudes to refugees after the russian invasion of ukraine](#). *West European Politics*, 0(0):1–26.
- C.E. Osgood, G.J. Suci, and P.H. Tannenbaum. 1957. *The Measurement of Meaning*. Illini Books, IB47. University of Illinois Press.
- Céline Paré. 2022. Selective solidarity? racialized othering in european migration politics. *Amsterdam Review of European Affairs*, 1(1):42–61.
- Jade Marie Prideaux de Lacy. 2023. "the white-washing of europe": A comparative analysis of migration policy towards the middle east and ukraine, as a reflection of european identity politics.
- Hannah Rashkin, Sameer Singh, and Yejin Choi. 2015. Connotation frames: A data-driven investigation. *arXiv preprint arXiv:1506.02739*.
- James A Russell and Albert Mehrabian. 1977. [Evidence for a three-factor theory of emotions](#). *Journal of Research in Personality*, 11(3):273–294.
- Alexander W Schmidt-Catran and Christian S Czyszara. 2023. Political elite discourses polarize attitudes toward immigration along ideological lines. a comparative longitudinal analysis of europe in the twenty-first century. *Journal of Ethnic and Migration Studies*, 49(1):85–109.
- Charlotte Taylor. 2021. [Metaphors of migration over time](#). *Discourse & Society*, 32(4):463–481.
- Justyna Tomczak-Boczko, Klaudia Gołębiowska, and Maciej Górny. 2023. Who is a 'true refugee'? polish political discourse in 2021–2022. *Discourse Studies*, page 14614456231187488.
- Stephen M Utych. 2018. How dehumanization influences attitudes toward immigrants. *Political Research Quarterly*, 71(2):440–452.
- Teun A. van Dijk. 2018. *Discourse and Migration*, pages 227–245. Springer International Publishing, Cham.
- Andreja Vezovnik. 2018. [Securitizing migration in slovenia: A discourse analysis of the slovenian refugee situation](#). *Journal of Immigrant & Refugee Studies*, 16:39 – 56.
- Jason Wei, Maarten Paul Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew Mingbo Dai, and Quoc V. Le. 2022. [Finetuned language models are zero-shot learners](#).
- Michael Wiegand, Josef Ruppenhofer, and Elisabeth Eder. 2021. Implicitly abusive language—what does it actually look like and why are we not getting there? In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 576–587.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick

von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

Natalia Zawadzka-Palucka. 2023. Ukrainian refugees in Polish press. *Discourse & Communication*, 17(1):96–111.

Iztok Šori and Vasja Vehovar. 2022. [Reported user-generated online hate speech: The ‘ecosystem’, frames, and ideologies](#). *Social Sciences*, 11(8).

A. List of Moral Disgust Terms

skrunstvo, nečist, zamazanost, prostitut, grešnica, nezmeren, zloba, klatež, svetoskrunski, izkoriščevalski, razuzdanost, opolzek, beda, izprijen, perverznej, opolzkost, zamazan, nespodoben, odbijajoč, gnus, brezbožen, vlačuga, razvraten, kužen, profana, prostaški, grešnik, cenenost, svetoskrunskost, izprijenost, kurba, umazanija, grešiti, profan, brezobziren, oskruniti, gnusen, ogaben, prešuštnik, grešen, nečistost, nalezljiv, perverzen, grešenje, skrunjen, ogabnost, oskrunjen, beden, razsipen, nalezljivost, razvrat, umazan, prostaškost, bogokleten, razuzdan, greh, odvraten, okuženost, omadeževanost, kužnost, odvratnost, omadeževan, nespodobnost, profanost, bogokletnost, brezbožnost, prostitucija, cenen, prešuštvo, prešuštnica

9. Language Resource References

Walter Daelemans, Darja Fišer, Jasmin Franza, Denis Kranjčič, Jens Lemmens, Nikola Ljubešić, Ilija Markov, and Damjan Popič. 2020. [The LiLaH emotion lexicon of Croatian, Dutch and Slovene](#). Slovenian language resource repository CLARIN.SI.

Darja Fišer. 2015. [Semantic lexicon of Slovene sloWNet 3.1](#). Slovenian language resource repository CLARIN.SI.

Saif Mohammad. 2018. [Obtaining reliable human ratings of valence, arousal, and dominance for 20,000 English words](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 174–184, Melbourne, Australia. Association for Computational Linguistics.

Saif Mohammad and Peter Turney. 2010. [Emotions evoked by common words and phrases: Using Mechanical Turk to create an emotion lexicon](#). In *Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, pages 26–34, Los Angeles, CA. Association for Computational Linguistics.

Saif M. Mohammad and Peter D. Turney. 2013. Crowdsourcing a word-emotion association lexicon. *Computational Intelligence*, 29(3):436–465.