

Modeling and Measuring Short Text Similarities On the Multi-Dimensional Differences between German Poetry of Realism and Modernism

Anton Ehrmanntraut¹ 
Thora Hagen¹ 
Fotis Jannidis¹ 
Leonard Konle¹ 
Merten Kröncke² 
Simone Winko² 

1. Institut für Deutsche Philologie, Julius-Maximilians-Universität Würzburg , Würzburg, Germany.
2. Seminar für Deutsche Philologie, Georg-August-Universität Göttingen , Göttingen, Germany.

Citation

Anton Ehrmanntraut, Thora Hagen, Fotis Jannidis, Leonard Konle, Merten Kröncke, and Simone Winko (2022). "Modeling and Measuring Short Text Similarities. On the Multi-Dimensional Differences between German Poetry of Realism and Modernism". In: *Journal of Computational Literary Studies* 1 (1). [10.48694/jcls.116](https://doi.org/10.48694/jcls.116)

Date published 2022-12-08

Date accepted 2022-03-30

Date received 2021-12-23

Keywords

short text, similarity, poetry, modernism, realism

License

CC BY 4.0 

Reviewers

Anne-Sophie Bories 

Thomas Weitin 

Note

This paper has passed through the conference track of JCLS. In addition to being peer reviewed, it was presented and discussed at the 1st Annual Conference of Computational Literary Studies 2022 at the Technical University of Darmstadt.

Abstract. This study contributes to the ongoing discussion on how to operationalize text similarity for the purposes of computational literary studies by defining, justifying theoretically and employing a multi-dimensional text model. Additionally, we evaluate a set of strategies to implement this model for very short texts like poetry using a range of methods from weighted sparse vectors up to very recent neural sentence embeddings based on annotations of emotions, genre and similarity. And finally, we show the relevance of using such a complex text model by applying the best method to a research question about the development of early modernism in German poetry. While we can confirm some important hypotheses from literary studies, we are also able to differentiate or relativize others. In particular, our findings do not support the widely held thesis that the change from realism to modernism was a revolutionary 'rupture'.

1. Introduction

This paper pursues two equally important goals: First, to find a suitable state-of-the-art method to model and analyze text similarity for poetry, and second, to contribute to the field of literary studies by studying the transition from realist to modernist poetry using the concept of similarity. The perception of similarity between texts is the basis for the construction of many literary terms like genre, author, or, as in our case, period. Grouping texts according to these terms usually presupposes that these texts have something in common and that these groups can be distinguished via these commonalities from other texts. Though the concept of similarity is ubiquitous in the practice of literary studies it has seldom been analyzed explicitly. While there are several studies that focus on the political aspects of comparison and (dis)similarity analysis (e. g. Felski and Friedman 2013), these contributions generally have little direct connection to our methodological goals. More applicable are some studies by scholars in Comparative Studies who reflected on similarity as part of their discipline defining practice (e.g. Corbineau-Hoffmann 2013). Similar attempts to model text similarity beyond the aspect of content have also been undertaken in computational linguistics (e.g. Bär et al. 2011). So one of the major contributions of this paper is our attempt to bring

these discussion threads together. But while it is possible to discuss these dimensions on a very abstract level, it is not possible to evaluate them on the same level. When we talk about structural aspects of a text, we look at very different elements depending on the genre we look at: speaker, stage directions, *dramatis personae*, etc. for drama, or stanza, verse, rhyme, etc. for poetry. Therefore, in order to discuss the phenomenon not only theoretically, but also to be able to apply it practically – and that means, above all, to include an evaluation method – it is more productive to limit the task to one genre – in our case, poetry.

The second goal of our research is to provide a broad foundation for a literary history of the beginnings of modernism. In the last years, we assembled a corpus of German poetry consisting of poems from realist and modernist anthologies. We are analyzing this corpus under the perspective of whether we can contribute to the discussion about the transition from realism to (early) modernism. We are using these period terms, as is the custom nowadays in literary studies, as useful constructions. That means: On the one hand it is understood that real breaks and disruptions are very rare and that history can be better understood as an evolutionary, gradual process with many small changes at each step. On the other hand, we assume that this process is not happening at the same speed all the time and that many of the changes in one time segment show some commonalities. Specifically, we will use the concept of similarity to describe the changes between the texts from the different corpora.

We structure our paper as follows: In a theoretical section, we first develop a four-dimensional model of textual similarity for poetry ([section 2](#)). We then describe our corpora; mainly the digitized anthologies of the poetry of realism and early modernism mentioned above ([section 3](#)). A selection of these poems was previously manually annotated with a hierarchical system of emotion labels. Within the context of our work, a subset of this selection was then additionally annotated using the dimensions of similarity described in the theoretical section. The following section discusses how each of these four dimensions can be measured in poetry ([section 4](#)). Poetry presents specific computational challenges even for semantics, a relatively traditional dimension of similarity. The main issue poses the shortness of the texts. Semantic similarity, which is usually modeled by using weighted terms to locate a document in vector space, does not work reliably on short texts. Additionally, working with poetry entails having to adapt to its specific language. This includes a high percentage of figurative speech, which makes the analysis of semantic similarity especially difficult, and also a high percentage of archaic words and expressions. For each of the four dimensions of similarity, we discuss and evaluate different methods to measure them using our poetry corpus in a first step. Among our methods used are traditional sparse document vectors, short dense feature vectors, and dense document embeddings, created either by computing them from token vectors or by using the recently proposed approach for sentence embeddings. In a second step, we adapt the best-performing models to each of the four dimensions ([subsection 4.4](#)). In the last section, we employ our final models from this two-step approach to assess the degrees of similarity and difference between realist and modernist poetry ([section 5](#)). In particular, we take up three specific research questions from literary studies and discuss our results with respect to the predominant hypotheses within the field. These questions are:

1. How does naturalist poetry relate to realism and modernism?
2. How homogeneous are realist and modernist poems?
3. How revolutionary is early modernism?

In summary, this study contributes to the ongoing discussion on how to operationalize text for computational literary studies by defining, theoretically justifying, and employing a multi-dimensional model of similarity. Additionally, we evaluate a set of strategies to implement this model for poetry using a range of methods from weighted sparse vectors up to the recent neural sentence embeddings based on extensive annotations of emotions, genre, and similarity. And finally, we show the relevance of using such a complex text-based model by employing the best method to provide new input for the continued research on the development of early modernism in German poetry.

2. Theoretical Considerations

As far as we can see, in literary studies, text similarity has been discussed mainly by Comparative Studies, where the concept of ‘comparison’ has been closely linked to ‘similarity’ (e.g. Zelle 2005). There seems to be a consensus that comparison is only possible on the basis of similarity in some specific aspects. Though principally many different aspects have been and can be used to compare literature, some have been established as especially useful for the study of literature. Corbineau-Hoffmann (2013), for example, groups them under three headings:

- I. Content (1. theme, 2. motifs, 3. settings, 4. characters, 5. concepts)
- II. Text-organization (1. narrative/description, 2. poetry/prose, 3. style levels, 4. instances of speech, 5. discourse)
- III. History (1. influences, 2. epochs, 3. other arts, 4. sciences, 5. genre).

While the first two groups are aspects of a text, the last group refers to typical contexts, often established again by analyzing groups of texts. To avoid the recursive loop hidden here, we focus on the two first aspects, ‘content’ and ‘text-organization’. It is important to note that these are open lists. There are other interesting aspects, but the ones mentioned are often used when people compare literature. The terms grouped under ‘content’ can be seen as parts of text semantics in general. A text has a theme, or there are specific motifs in a text, but usually, the meaning of text is more than each of these, it encompasses all of them. The terms grouped under ‘text-organization’ on the other hand cover quite heterogeneous aspects – even if you substitute the more common ‘form’ for it. In our experience, especially the term ‘style’ is hard to subsume under the same dimension as other text-organizational aspects.

Semiotics and linguistics support this position as they also distinguish between form and style (Nöth 2008; Sandig 2006), and the three aspects – content, structure, and style – are also distinguished in one of the very few attempts in computational linguistics to model text similarity (Bär et al. 2015). We propose to add one dimension which can only be subsumed with difficulties under one of the three headings and which is usually

highly important, especially for literature and especially for poetry, which has been defined as the prototypical medium to express subjective feelings: emotion.¹

Content, Form, Style, and Emotion are the four dimensions of similarity which we will use to describe the relations between texts. From the perspective of this study, it is more useful to explicate the dimensions via operationalizations and examples rather than ‘exact’ definitions. To this end, the annotation guidelines (see section 3) list specific components that make up the dimensions. Content consists of components such as theme, character, or setting; form is operationalized primarily through stanza structure, meter, and rhyme; style, in contrast, refers to components such as register or metaphor; and for emotion, we consider, among other things, the extent to which emotions are represented and their polarity. In further studies, these components could be analyzed individually and be integrated into an even more complex model of text similarity.

The heterogeneity of the four dimensions will have a direct influence on the inter-annotator agreement and the performance of any machine learning model trained to detect these aspects automatically. From a theoretical perspective, it is unclear how the dimensions relate to each other, or in the language of statistics, how much they correlate. Winko (2003), for example, assigns the aspect ‘linguistic shaping of emotions’ via the aspect ‘presentation of emotions’ to what is called ‘style’ in our model, while she assigns it to content via the aspect ‘thematization of emotions’. From this perspective, a relatively high correlation of emotion with content and style is to be expected.

3. Corpus and Annotation

The corpus is a collection of anthologies of contemporary poetry from the two epochs ‘realism’ and ‘modernism’.² The collections contain poems that the anthologists, i.e. contemporary experts in poetry, consider to be particularly typical, outstanding, or representative among other aspects. From the large amount of poetry anthologies in both epochs, the corpus was compiled³ according to the following **criteria**: The collections contain **contemporary** poetry, have **no thematic restrictions**, and are all aimed at a **general audience** rather than a particular target group. The criteria minimize the risk that thematic constraints or specific addressee orientation could influence the poem selection as systematic factors. It is important to note that the corpus construction leads us to model ‘realism’ and ‘modernism’ from the perspective of the period of investigation. Ultimately, it is the anthologists who determine which texts are ‘realist’ and which are ‘modernist’, and their views do not always have to coincide with the perspective of today’s research. The corpus contains texts by both canonical and non-canonical authors. We call authors ‘canonical’ if they are frequently mentioned in recent literary histories. For early modernism, this applies to Stefan George, Hugo von Hofmannsthal, Arno Holz, Else Lasker-Schüler, and Rainer Maria Rilke.

1. Why emotion is a dimension of its own for the analysis of text is discussed in (Winko 2003).

2. Since this epoch is characterized by a multitude of literary trends, the more neutral label ‘turn of the century around 1900’ is preferred in literary studies. We choose the term ‘modernism’ because the anthologies we include claim to present modern poetry. In the following, ‘modernism’ always means ‘early modernism’, i.e. literature before expressionism.

3. For our corpus selection we used Günter Häntzschel’s comprehensive bibliography (Häntzschel 1991).

(1) Sub-corpus ‘Realism’: The first sub-corpus consists of 7 anthologies with German poems from the realist epoch: Avenarius 1882; Bern 1877; Kneschke 1865; Moltke 1882; Polko 1860; Prutz 1859; Willatzen 1875. The poems included in the anthologies cover the period under study, 1850 to 1880. Some of the anthologies, but especially Elise Polko’s widely distributed collection, also contain some poems written before the period of study; these have been excluded. This sub-corpus consists of 3,039 poems by a total of 484 different authors.

(2) Sub-corpus ‘Modernism’: Of the 941 anthologies of German-language poetry published in first edition between 1885 and 1912 (Häntzschel 1991, 587–589), twelve anthologies meet the selection criteria: Arent 1885; Benzmann 1904; Bethge 1905; Bierbaum 1893, 1894; Bonsels et al. 1905; Federmann 1908; P. Friedrich 1911; Gemmel 1898; Huch 1911; Jacobowski 1899; Renner 1899; Tille 1896. They all claim to contain ‘modern poetry’. This sub-corpus consists of 2,882 poems by a total of 361 authors. We annotated 1,278 poems from both sub-corpora for emotion and thematic genre. Thematic genres such as love poetry or nature poetry provide information about the content of the poems.⁴ The annotated emotions are not the readers’ emotions, but rather the emotions expressed in the text itself. The annotators used a list of 40 discrete emotions which we categorized into 6 groups, inspired by the emotion hierarchy in Shaver et al. 1987: love, joy, agitation/surprise, anger, sadness, and fear. First, emotions and genres were annotated independently by two annotators, then they merged annotations manually into a consensus annotation. Their agreement before creating the consensus annotation, measured with γ (Mathet et al. 2015), was 0.6445 for individual emotions, 0.7491 for the emotion groups, and 0.69 Krippendorff’s alpha (Krippendorff 2011) for the thematic genres.

Additionally, we annotated the similarity of the poems.⁵ The task was not to annotate absolute similarities (“These two poems are not at all/a little/very similar”), but relative similarities (“Poem A is more similar to poem B than poem C”), which is much easier. For each triple of poems, the annotators had to judge for each similarity dimension (content, form, style, and emotion) and for a comprehensive ‘overall’ category whether the focus poem was more similar to the one on the left, to the one on the right, or equally (dis)similar to both. The annotation guidelines specify for each dimension which components should be taken into consideration, e.g. stanza structure, rhyme, meter, and text length in case of the formal dimension, and which of these aspects are typically most important. Nevertheless, the annotators ultimately had to weigh the components on a case-by-case basis, which required considerable literary expertise.

We annotated 470 triples, consisting of a total of 866 poems. One constraint for the selection of the triples was that the poem length had to be quite short due to technical prerequisites.⁶ In addition, we selected triples for which we expected a strong similarity of the middle text with either the left or right text based on formal features such as text length or previous annotations of thematic genres and emotions. This second constraint

4. As the annotation is still ongoing to cover more poems, the entire corpus and a detailed report on the annotation guidelines for emotions and genre will be published at a later date.

5. The annotation guidelines can be found here: https://github.com/cophi-wue/jcls2022-poem-similarity/blob/main/annotation_guidelines_text_similarity.pdf.

6. The length of poems is bound to a maximum of 124 sentence-piece tokens used as input for paraphrase-*xlm-r-multilingual-v1*.

Content	1	0.24	0.68	0.42	0.82
Form	0.24	1	0.22	0.43	0.48
Style	0.68	0.22	1	0.4	0.76
Emotion	0.42	0.43	0.4	1	0.66
Overall	0.82	0.48	0.76	0.66	1
	Content	Form	Style	Emotion	Overall

Figure 1: Pearson correlation in annotated dimensions (majority vote).

ensured that the annotators could deal with reasonably clear cases. There were 400 triples covering both constraints available in our annotated poems. 70 triples were additionally annotated without similarity expectations. Each triple was annotated by at least two people.

The agreement, measured with Krippendorff's alpha, was 0.53 for content, 0.68 for form, 0.44 for style, 0.32 for emotion, and 0.48 for overall. Possible reasons for the differences in agreement are that the dimensions with lower agreement are more dependent on interpretation or that the weighting of the components is more ambiguous in their case. An experiment showed, however, that three annotators who created a consensus annotation after annotating 60 triples were able to increase their agreement when annotating another 30 triples from 0.49 to 0.63 on content, from 0.48 to 0.69 on emotion, from 0.32 to 0.41 on style, and from 0.45 to 0.68 on overall (only the agreement on form deteriorated from 0.77 to 0.71, but still remained high). Since the creation of consensus annotations seems to significantly improve the annotation quality, we plan to create consensus annotations for all triples in the future. Until then, for the triples without consensus annotations, we will only use annotations that the majority of annotators agree with. In the evaluation of the following section, we also omit all annotations that the majority of annotators found that 'the middle text is equally (dis)similar to both the left and the right text'.⁷ That leaves us with 346 usable annotations for content, 388 for form, 331 for style, 359 for emotion, and 381 for overall, with every annotation stating that the middle text is more similar to either the left text or the right text.

Some of the similarity dimensions correlate strongly with each other, according to the annotations. The 'overall' dimension correlates most strongly, especially with content and style. This is understandable since the annotation of the 'overall' dimension is usually based on annotations of the other similarity dimensions. Another relevant correlation exists between content and style. The most independent dimension is form, whose correlation with the other dimensions is the weakest.

7. More precisely, for each triple and similarity dimension, we calculate the mode of the annotation results. We use 'The middle text is more similar to the left text' (from now on: 'left') as the final annotation if the mode is 'left', but also if it is 'left' and at the same time 'The middle text is equally (dis)similar to both the left and the right text'. The same is true in reverse for annotations on the right. All other annotations are discarded.

4. Dimensions of the Similarity of Poems

Measuring the similarity of poems along the dimensions discussed above poses two challenges: First, the shortness of the texts makes it difficult to apply well-established approaches with high reliability. Research in natural language processing has proposed a set of methods for the measurement of short text similarity (Prakoso et al. 2021) usually complementing the texts with other sources which compensate for the lack of information in the text themselves. But research on text similarity, in general, focuses on the ‘content’ aspect. So the second challenge lies in finding methods that can be used to model the other dimensions.

Overviews of the research on short text similarity classify the methods in four groups: string-based, corpus-based, knowledge-based, and hybrid-based (Gomaa and Fahmy 2013; Prakoso et al. 2021, 1): 1) string-based methods use only the word or character tokens to create a representation of the text. We use *tfidf* and *mfw*. 2) Knowledge-based methods use an external knowledge base like WordNet. Our two models *features-formal* and *features-emotional* can be seen as variants of this approach. 3) Corpus-based methods use an external corpus to create information-rich representations, nowadays usually word embeddings: *FastText*, *Glove*, *GBert*. Additionally, we experiment with document embeddings using different sentence embedding methods: *XLM-R*, *mpnet*, *MiniLM*, *cross-en-de-roberta*. The drawback of this approach is that we are limited to an input of 127 tokens, but it is reported to be the best representation for short texts. 4) Hybrid approaches, combining some of the strategies outlined above.⁸

The small number of poems we had annotated under the perspective of similarity made it inadvisable to use the typical finetuning approach. Instead, we opted for broader testing of how different text representations are able to mirror the results from our annotations, select the best representations, and then tweak the vector spaces with similarity learning based on our dimension annotations. So in sections 4.1 to 4.3, we introduce the different models we were able to use and their evaluation based on our similarity annotations. In subsection 4.4, we apply similarity learning to the best performing models.

4.1 Models

We evaluate the following embeddings, which can be roughly categorized into more simple baselines on the one hand, and some more complex embeddings derived from sophisticated deep-learning language models on the other. The baseline embeddings are defined as follows:

TFIDF-{1000,10000,20000}: Poems are represented by a vector, where the dimensions correspond to the 1000 etc. most frequent terms in our corpus. Each individual vector component is the relative term frequency of that term in the poem, weighted by the inverse document frequency.

MFW-{100,200,500,1000}: Defined like the embedding above, but the term frequencies are z-standardized for each term over all poems.

8. Bär et al. (2015) distinguish between compositional measures, which usually “compute pairwise word similarity between all words, and aggregate the resulting scores to an overall similarity score” (Bär et al. 2015, 5), and non-compositional measures, which project the texts into a shared space like the vector space model (Salton and McGill 1983). We concentrate here on the latter.

Features-Formal: Poems are represented by a vector of the following four formal features: stanza count, verse count, word count, average stanza length in verses – each z-standardized over all poems.

Features-Emotional: Each poem is embedded with a vector of its verse-level relative frequency of shaver emotions (see section 3). These emotions derive either from annotations or, if no annotations are available, predicted by a machine learning model.⁹

The following deep-learning embeddings are derived from pre-trained static type-based embeddings:

{FastText,GloVe}-{mean,median,meannorm,sif}: For each term in the poem (minus stopwords), we obtain the embedding vector for that term with FastText (trained on the German OSCAR corpus with $d = 1536$ as proposed by Ehrmantraut et al. 2021, resp. a GloVe model with $d = 300$ provided by Deepset¹⁰). Finally, on that set of vectors, we compute the arithmetic mean (resp. median, resp. meannorm (Ehrmantraut et al. 2021), resp. arithmetic mean weighted by smooth inverse frequency (Arora et al. 2017)) to obtain a single vector for a particular poem.

Similarly, the following embeddings are derived from the output of BERT, which generates vectors for each token, but also takes into consideration the textual context of the entire input sequence.

GBERT-lastlayer-{mean,median,meannorm}: For a particular poem, we plug in the tokenized poem into GBERT_{Base} (Chan et al. 2020), the currently best performing German BERT model. BERT then computes a contextualized output vector (i.e., the output of the last layer) for each token. We now aggregate all vectors by taking the arithmetic mean (resp. median, resp. meannorm), just like above. This results in a vector with 768 dimensions.

GBERT-alllayers-{mean,median,meannorm}: Defined just like above, except that we not only consider the final output vector, but the outputs of all layers. That is, for each token, we concatenate the input embedding with the 12 Transformer outputs to derive a vector with $d = 13 \times 768$. Then, like above, we aggregate this sequence of token vectors into a single vector for a particular poem.

In contrast, the following embeddings result from pre-trained language models following a Sentence-BERT-architecture, as proposed by Reimers and Gurevych 2019, 2020.

paraphrase-XLM-R: This model is a multilingual XLM-RoBERTa model that is fine-tuned to imitate Sentence-BERT-paraphrases, as described by Reimers and Gurevych 2020. We let paraphrase-XLM-R interpret our poems as sentences, which outputs a vector representation for each poem with $d = 768$. Note that in the case of paraphrase-XLM-R and all following Sentence-BERT models, fine-tuning was only performed for input sequences no longer than 126 SentencePiece tokens. Therefore, we also restrict our evaluation of these models to poems that are no longer than 126 SentencePiece tokens. (These are 29% of all poems in our corpus.)

9. The model achieves a performance of 0.73 (f1 score).

10. See: <https://www.deepset.ai/german-word-embeddings>.

paraphrase-mpnet, paraphrase-MiniLM, cross-en-de-roberta: Similarly, these are pre-trained Sentence-BERT models trained on a wide variety of sentence pair datasets and parallel multilingual data. Specifically, we use publicly available variants `paraphrase-multilingual-mpnet-base-v2`, `paraphrase-multilingual-MiniLM-L12-v2` provided by Reimers and Gurevych 2019, and `cross-en-de-roberta-sentence-transformer` provided by T-Systems online. Again, we use these on poems of length ≤ 126 SentencePiece tokens to obtain vector representations with $d = 768$.

4.2 Evaluation Setup

Evaluating the embeddings as described above requires us to formulate a task that probes each embedding space for its ability to represent certain dimensions of (dis-)similarity of poems via their distances in that particular embedding space, taking into consideration and comparing against the human annotations. As the embeddings define no particular distance function, we evaluate every embedding with each of the following three distance functions: Euclidean (L2), Manhattan (L1), and Cosine Distance.

We opted to replicate our prompts for the annotators by formulating a binary classification problem on a particular dimension of similarity, and checking if the model can replicate the majority vote. Note that these votes either take the value ‘annotated left’ or ‘annotated right’. Assume some embedding space and some distance function d fixed. For some annotated triple, let *left*, *anchor* and *right* denote the corresponding vectors in that embedding space. Now, we make the following prediction:

- Predict ‘annotated left’ if $d(\text{anchor}, \text{left}) < d(\text{anchor}, \text{right})$, i.e., *left* is closer to *anchor* than *right*.
- Otherwise, predict ‘annotated right’.

To compare the true majority annotations with the predicted ones, we use the *balanced accuracy* (arithmetic mean over the recall of both classes, cf. Grandini et al. 2020) as our metric. Note that the random ‘no skill’ classifier has a balanced accuracy score of 0.5.

We remark that variations on the above operationalization are possible as well, particularly if we do not omit cases where the majority of annotators chose ‘The middle text is equally (dis)similar to both the left and the right text’. However, while experimenting we observed that when including this third class ‘same’ in the operationalization, the balanced accuracy significantly drops. (We made the different balanced accuracies comparable by rescaling to the range $1/(1 - \#\text{classes})$ to 1, so that performance at random scoring is always at 0.) We suspect that this difference in performance is caused by the complexity of the triples that were labeled with ‘annotated same’: Human annotators agree on the features which make them classify a text as ‘more similar to the focus text’. But the label ‘same’ is given when neither of both comparison texts shows obvious similarities to the focus text, but that does not imply that the comparison texts have any features in common; they can be different to the focus text in very diverse ways.

In particular, we experimented with the following two variations of the original operationalization:

- (a) Probe whether the embedding can predict ‘annotated equally (dis)similar’ vs ‘an-

notated left or right' by evaluating $|d(anchor, left) - d(anchor, right)| < \epsilon$ against some optimal decision boundary ϵ .

(b) In a 3-class classification setup, probe whether the embedding space admits a classification using an optimal symmetric decision boundary ϵ . That is, predict 'annotated left' when $d(anchor, right) - d(anchor, left) > \epsilon$ (*left* is closer to *anchor* than *right* by at least ϵ). Symmetric, when $d(anchor, left) - d(anchor, right) > \epsilon$, predict 'annotated right'. And otherwise, when $|d(anchor, left) - d(anchor, right)| \leq \epsilon$, predict 'annotated equally (dis)similar'.

As outlined above, variant (a) is solved with lower balanced accuracy than the original operationalization throughout all embeddings and variant (b) with even lower accuracy.

4.3 Results

The results show for all dimensions except 'form' a clear increase with the complexity of text representation: Word Embeddings are better than sparse representations – with dynamic embeddings based on BERT showing a better performance than static embeddings – and sentence embeddings are better than word embeddings. The best sentence embedding is showing an acceptable performance, especially if the cosine is used. As almost all the strategies of text representation, which we applied here, have been developed with the main focus on the semantic aspect, it is not too surprising that the best model is the best in all dimensions. The one big exception is form. Using only a very small set of features is enough to match the annotations. Discussions with the annotators revealed that they usually based their decision on a very small set of observations.

The best model is *paraphrase-mpnet*. To evaluate all German sentence embedding models,¹¹ which are available at this moment, we use the about 9,000 sentences of the Sick dataset (Marelli et al. 2014) which we had translated into German with DeepL. Our results show *paraphrase-XLM-R* (correlation with human annotations: 0.82) slightly ahead of *paraphrase-mpnet* (0.8165), which is why we include these two models and the best model based on static word embeddings (*FastText-mean*) in the next step.

4.4 Similarity Learning

To adapt the text representations to the specific textual dimensions (content, form, style, and emotion), we additionally apply similarity learning. The goal of this step is to learn a transformation of the vectors presented in the previous chapter that allows for better reproduction of the annotation. We use a siamese neural network (Bromley et al. 1993) for this purpose, which we modeled following the maaten network structure from (Szubert et al. 2019). The base model consists of three dense layers (500, 500 and 2,000 neurons) each followed by a normalization activation function (see Klambauer et al. 2017) and dropout. The input for the network consists of our annotated poem triples.

Regardless of the original size of its vector representation, each poem is transformed into a space with 128 dimensions. The loss, and hence the optimization objective of the

11. The multilingual models in Huggingface's sentence transformers; see <https://huggingface.co/sentence-transformers>.

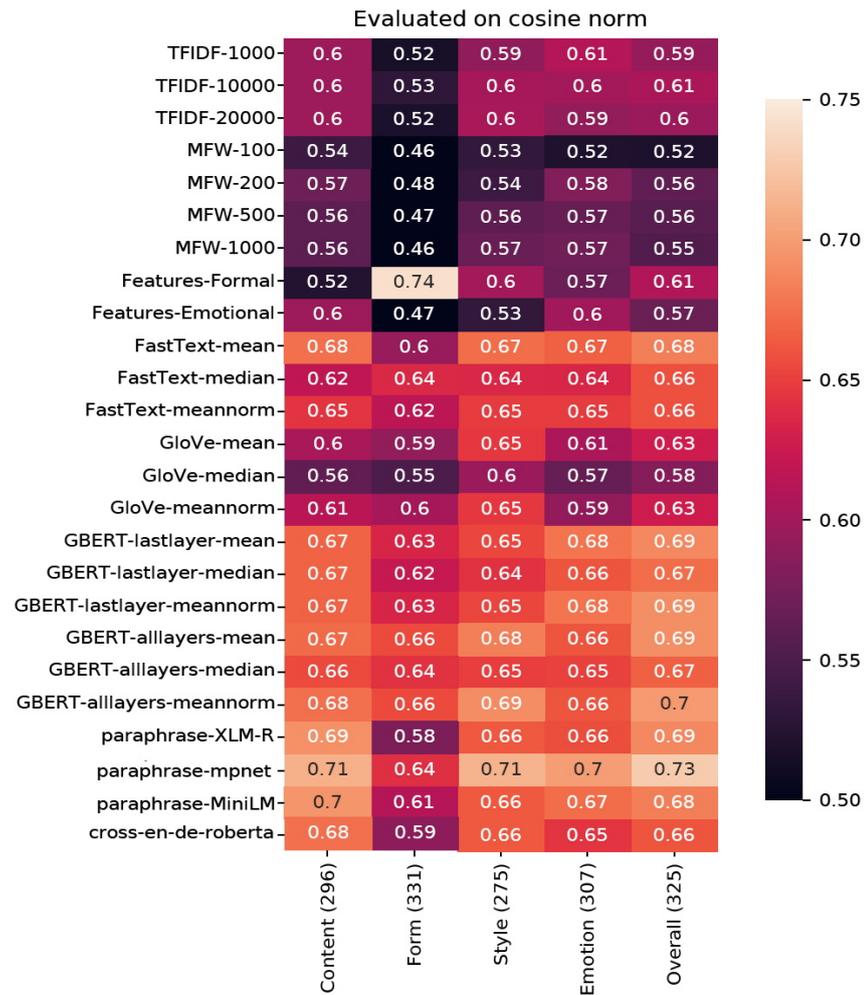


Figure 2: Balanced accuracy score for each model and dimension. Numbers on the x-axis indicate class support. For information on results with other distance metrics see [Figure 12](#) in [section 9](#).

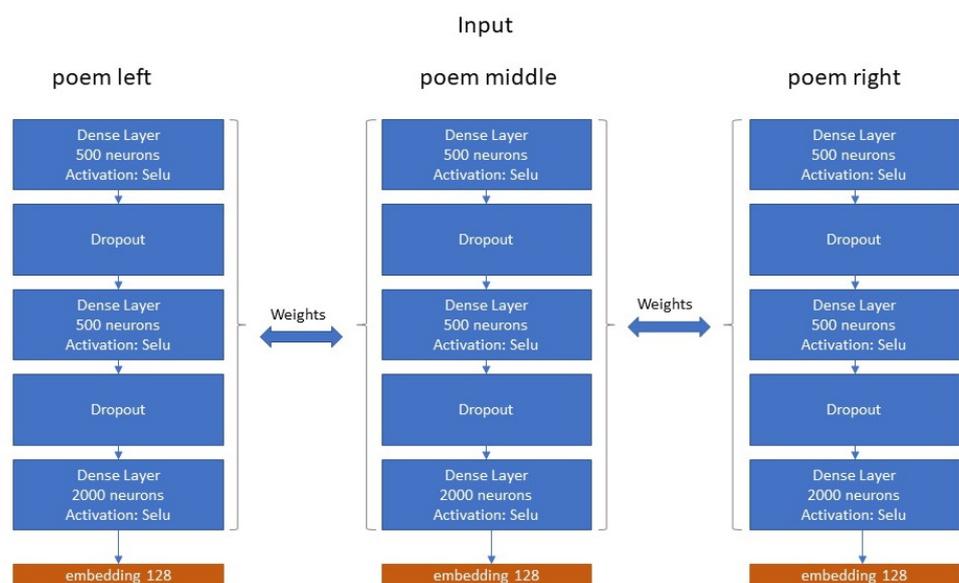


Figure 3: Architecture of the siamese neural network used for similarity learning.

Model	Content	Form	Style	Emotion	Overall
paraphrase-XLM-R	.69→.81	.58→.76	.66→.79	.66→.76	.69→.79
paraphrase-mpnet	.71→.75	.64→.68	.71→.71	.70→.74	.73→.74
FastText-mean	.66→.77	.59→.67	.65→.72	.66→.74	.66→.72
Formal-Features	-	.81→.81	-	-	-

Table 1: Similarity learning results (accuracy in 10-fold cross-validation). Format: best performance before similarity learning (see Figure 2) → performance afterwards.

network, is to maximize the difference between the focus text and the negative example while also minimizing the difference between the focus text and the positive example, i.e. the text which has been annotated as being more similar to the focus text. In short: in Euclidean distances ($\text{dist}(\text{anchor}, \text{negative}) - \text{dist}(\text{anchor}, \text{positive})$).¹² Learning rate decrease is bound to a *reduce on plateau* mechanism, which leads to strong performance gains compared to more common choices like constant or time-based decrease rates. The network’s performance is measured via the amount of correctly identified positive examples (accuracy, see Table 1).

4.5 Discussion

With our two-step approach, we are able to achieve good results for a complex task. It is probably open to discussion whether the restriction to 127 input tokens is acceptable compared to the small gain in performance. Future work will either improve on the input size or find a reliable way to compute representations for longer texts. Using one representation for three of the four aspects in the first step made us ask whether the representations after the second step are actually different. The correlations of distances (Figure 4) rate positive correlations between content and style, content and emotion, or style and emotion. In other words, the vector space was attuned to the specific dimension by similarity learning. The close relationship between ‘content’ and ‘overall’ was already noticed by the annotators.

It is unclear to us, why the different embeddings show significantly different improvements in the second step (mean values): 0.126 for *paraphrase-XLM-R*, 0.026 for *paraphrase-mpnet*, and 0.08 for *FastText-mean*; on what factors does this capability for improvement depend? Which training data and training regime for the sentence embeddings enables the text representation to be adaptable to the text dimensions beyond content?

The results from Figure 2¹³ show that the best results are obtained using language models with transformer architecture and that they increase even more if those have previously been fine-tuned for sentence similarity. With the additional adaptation by similarity learning, we now perform a third tuning step of representations created this way. A next step would be instead of using the frozen output vectors of those networks, to include the network in the learning process and model the similarity learning as a fine-tuning step. Likewise, we should add another layer of pertaining before similarity learning and perform a domain adaptation (Gururangan et al. 2020) to our corpus.

¹². Triplet margin loss.

¹³. Motivated by feedback during the conference, we additionally tested the combination of paraphrase-XLM-R and Formal-Features. This variant leads to an accuracy of 0.78 for the form aspect. It improves the result of paraphrase-XLM-R slightly, but remains below the value achieved by Formal-Features alone.

Content	1	0.6	0.62	0.58	0.79
Form	0.6	1	0.66	0.47	0.64
Style	0.62	0.66	1	0.55	0.71
Emotion	0.58	0.47	0.55	1	0.64
Overall	0.79	0.64	0.71	0.64	1
	Content	Form	Style	Emotion	Overall

Figure 4: Pearson correlation of distances in vector space after similarity learning.

5. (Dis)similarity between the Poetry of Realism and the Poetry of Modernism

5.1 Hypotheses from Literary Studies

In the following, we continue a discussion in German literary studies about the relationship of poems of realism to those of early modernism and the special position of naturalistic poetry in this development. We hope to contribute to this discussion by a mix of explorative methods and hypothesis testing. To enable the latter, we will condense positions in the debate into three hypotheses related to this transformation.

Hypothesis 1: The poetry of naturalism, as represented in the anthology *Moderne Dichter-Charaktere*, is predominantly traditional rather than modernist. The question of where exactly naturalism can be located between realism and modernism has been debated many times. In this context, the anthology *Moderne Dichter-Charaktere*, which is part of our corpus, is considered central to naturalist poetry. The anthology's introductions emphatically assert the novelty and revolutionary character of the texts (especially Conradi 1885, I–III). Research, on the other hand, is mainly of the opinion that these statements are exaggerated and that the poetry of the anthology is, on the whole, traditional (Vietta 1992, 194; Fähnders 1998, 36–37; Sprengel 1998, 621; Austermühl 2000, 350–51; Lamping 2000, 145–146; Andreotti 2014, 17). However, some scholars, even if they consider the anthology as a whole to be traditional, argue that it was at least innovative in terms of *content* since new themes such as 'big cities' or 'social issues' were addressed (e.g. Fähnders 1998, 36–37).

Hypothesis 2: Modernist poetry is heterogeneous, that is, more heterogeneous than realist poetry. While the poetry of realism, or at least the mass-produced poetry of this period, is considered by researchers to be relatively homogeneous (e.g. Stockinger 2010, 88), modernist poetry is highly diverse, according to many scholars, given the simultaneity of a wide variety of literary movements (Anz 2007, 330–331; Becker and Kiesel 2007, 30; Fähnders 1998, IX, 4). But the hypothesis of modernist heterogeneity

has its limitations. For example, some researchers support the view that modernism is homogeneous at least insofar as it responds to the same social-cultural problems (Vietta 1992, 30–31.; Fähnders 1998, 9–10; Becker and Kiesel 2007, 30; for further statements on the homogeneity of modernist poetry see H. Friedrich 1992, 140–142; Lamping 2008, 13). One researcher, therefore, argues that the period around 1900 was characterized by a “homogeneity of the heterogeneous” (Fähnders 1998, 11). Despite these limitations, most scholars would probably agree that modernist poetry is at least more heterogeneous than the poetry of realism.

Hypothesis 3: There is a fundamental ‘rupture’ between modernist poetry and earlier, more traditional poetry. This view was already held by contemporary authors, critics, and anthologists, who spoke of a ‘revolution’ in poetry (as an example from the corpus anthologies see Bethge 1905, 13–14; cf. on contemporary statements H. Friedrich 1992, 141; Anz 2007, 333; Lamping 2012; Wieland 2019, 17). Many researchers also emphasize major differences between modernism and previous literary periods, often using the metaphor of ‘rupture’ (H. Friedrich 1992, 20; Kiesel 2004, 141–142; Frick 2007, 97–98; Goltschnigg 2007, 169; Lamping 2012; Andreotti 2014, 5; without this metaphor: Klinger 2002, 160; Lamping 2000, 140; Lamping 2008, 11, 13). But the ‘rupture’-thesis is also partly qualified. For example, it is emphasized that modernism still refers to traditions (even though it uses them in new ways) (Kiesel 2004, 142–143; Frick 2007, 98–99; Goltschnigg 2007, 169). Others argue that many relevant authors were located somewhere *between* realism and modernism or that they combined traditional as well as new elements, which implies a smoother transition between periods (see for C. F. Meyer Selbmann 1999, 149, 152; for Fontane (!) Selbmann 2007, 201; for Baudelaire, Rilke, Hofmannsthal, and Kafka Lamping 2012). Still, others relativize the novelty of modernism in general (Hiebel 2005, 27; Anz 2007, 333). Thus, hypothesis 3 is partly controversial in research.

It is possible to combine the aforementioned hypotheses in a visual model. The purpose of this model is threefold: it visually summarizes the research hypotheses, it relates the hypotheses to one another, and it demonstrates that all hypotheses about similarity and dissimilarity combined offer a fairly comprehensive interpretation of the transformation from realism to modernism, again underscoring the relevance of similarity as a category of analysis.

In the model, each point represents a poem. The greater the distances between the points, the more dissimilar the texts. The distances are not based on calculations but on a hermeneutic understanding of the research and are meant as rough approximations of general ideas. One can see that the distances within realism are smaller than the distances within modernism. It is also visible that there is a strong division between realism and modernism and that the naturalist poems tend to gravitate more towards realism than modernism.¹⁴

Admittedly, this model is not explicitly advocated in research. Only rarely does a single scholar state all the hypotheses that the model synthesizes. Like any model, it represents only a section of reality and neglects other aspects, such as the differentiation

14. Distances within naturalism should not be given any further significance; no research hypotheses were considered in this regard.

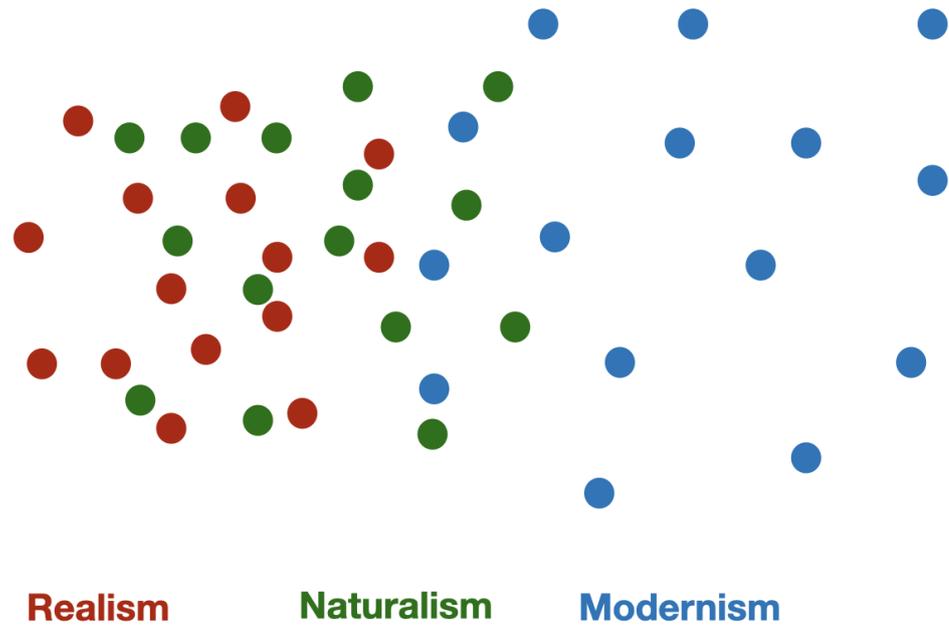


Figure 5: Model of the distances between poems of realism, naturalism and modernism according to research.

of individual dimensions of similarity, or synchronic and diachronic period-internal differentiations of, for example, individual authors, groups of authors, or literary movements. Some aspects of the model are, as explained, controversial in research, but it is all the more interesting to examine whether our results fit the model and the underlying hypotheses.

5.2 Results

For a first exploration of the (dis)similarities between realism and modernism, we project the poems into a two-dimensional space (Figure 6). Some similarities with the model derived from research (Figure 5) become apparent. In particular, a distinction between realism and modernism is evident, even though the separation is far from perfect since there are numerous overlaps between the two periods. Furthermore, it is consistent with the research model that the naturalist poems tend to stay within the realist spectrum and hardly enter ‘decidedly modernist’ areas. However, it is necessary to test the hypotheses from literary studies more precisely than just by explorative means.

Hypothesis 1: The poetry of naturalism, as represented in the anthology *Moderne Dichter-Charaktere*, is predominantly traditional rather than modernist To test the first hypothesis, we examined the similarity of the programmatically naturalistic anthology *Moderne Dichter-Charaktere* to the realism and modernism corpora. In addition, we measured the distances between the poems within the latter two corpora to be able to assess the comparative analyses more accurately.

The boxplots show that overall and for each individual dimension ‘content’, ‘form’, ‘style’, and ‘emotion’ the distances between naturalism and realism are smaller than the

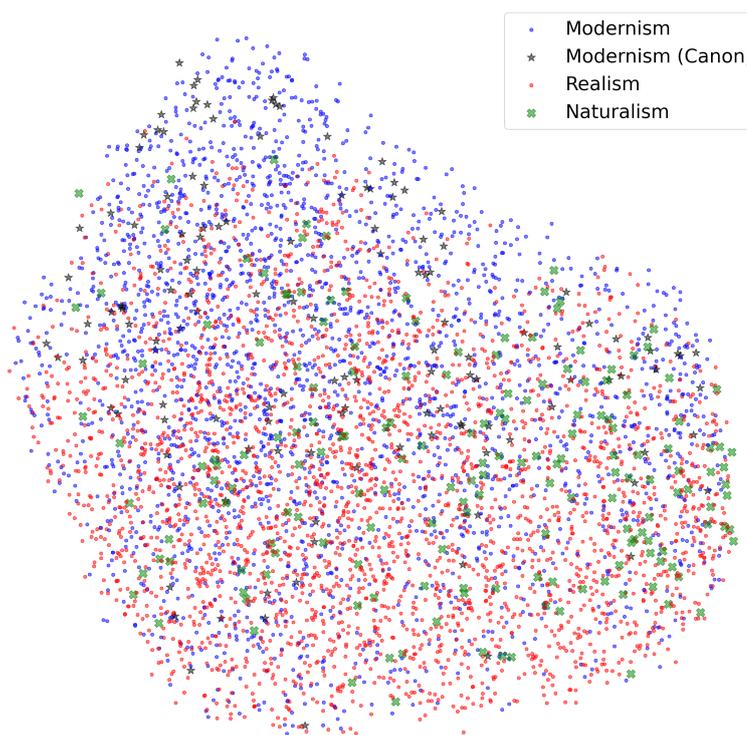


Figure 6: Poems embedded with both vanilla *GBERT-alllayers-meannorm* (see Figure 2) and *FastText-meannorm* transformed to reflect the aspect 'content' (see Table 1) projected in 2-dimensional space using UMAP (McInnes et al. 2018).

distances¹⁵ between naturalism and modernism. At the same time, the distance between the naturalism and realism corpus is larger than the distance between the poems within the realism corpus. Surprisingly, in the dimension 'content' no higher proximity to the modernism corpus is seen.

A stronger similarity between naturalist and modernist poems would have been expected based on the literary-historical theses we have mentioned above. As expected, the analyses support the thesis that the naturalism corpus is more similar to the realism corpus than to the modernist corpus. However, a more detailed look shows that differences can be found in the individual dimensions. This could indicate that the naturalistic poems probably do not use the same means as realistic poems. What exactly these differences are should be investigated in a further study. However, equating naturalist with realist poetry falls short in any case since the internal distance in the realism corpus is smaller than that in the comparison between naturalism and realism. It should be emphasized that we have studied the effect only for the anthology *Moderne Dichter-Charaktere* and only using its short poems, as stated above. Further study would have to take into account that the modernism corpus also contains some naturalistic poems.

15. We tested for significance and all differences are highly significant. To make sure this is not solely an effect of the large sample size we randomly selected 100 texts, but the differences stay significant. New guidelines usually recommend complementing p-values with effect size. In our case this is not easy to apply, because the measure is not grounded in an intuitively comprehensible unit.

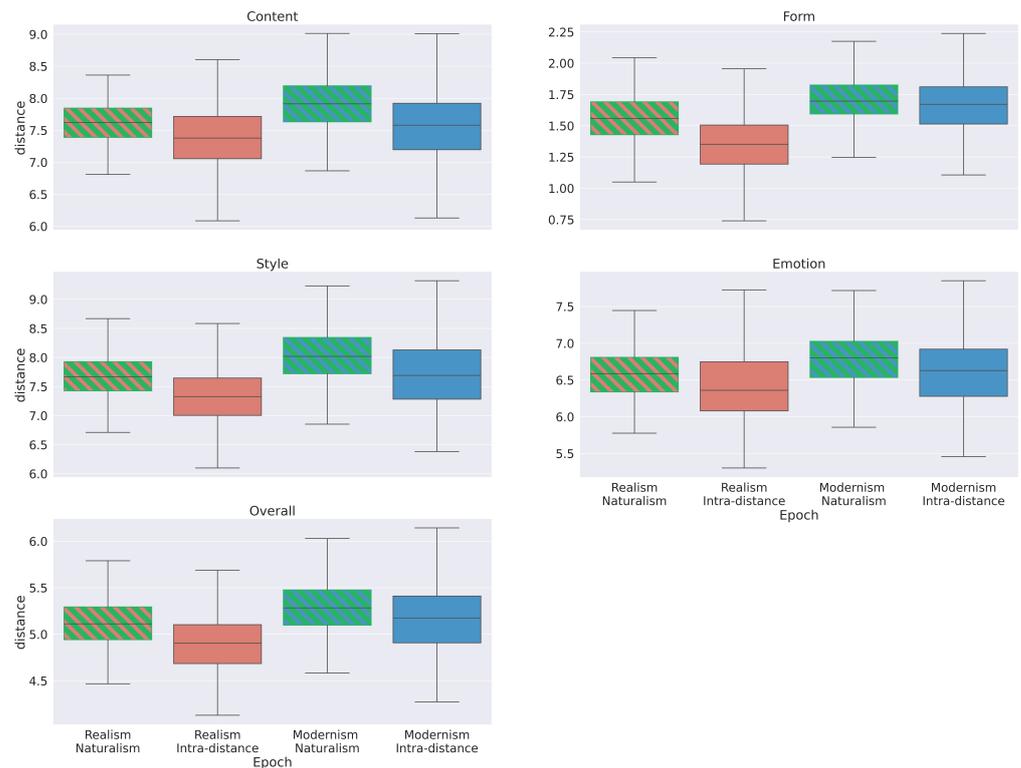


Figure 7: Distances between poems from Realism/Naturalism and Modernism/Naturalism and poems within Realism and Modernism. Distances in ‘content’, ‘style’, ‘emotion’ and ‘overall’ are measured in the space of *paraphrase-XLM-R* embeddings transformed via similarity learning (see section 4.4). Distances in ‘form’ are measured in the Feature-Form embedding space (see section 4.1). Each boxplot represents pairwise euclidean distances of 2,000 samples with a size of 20 poems.

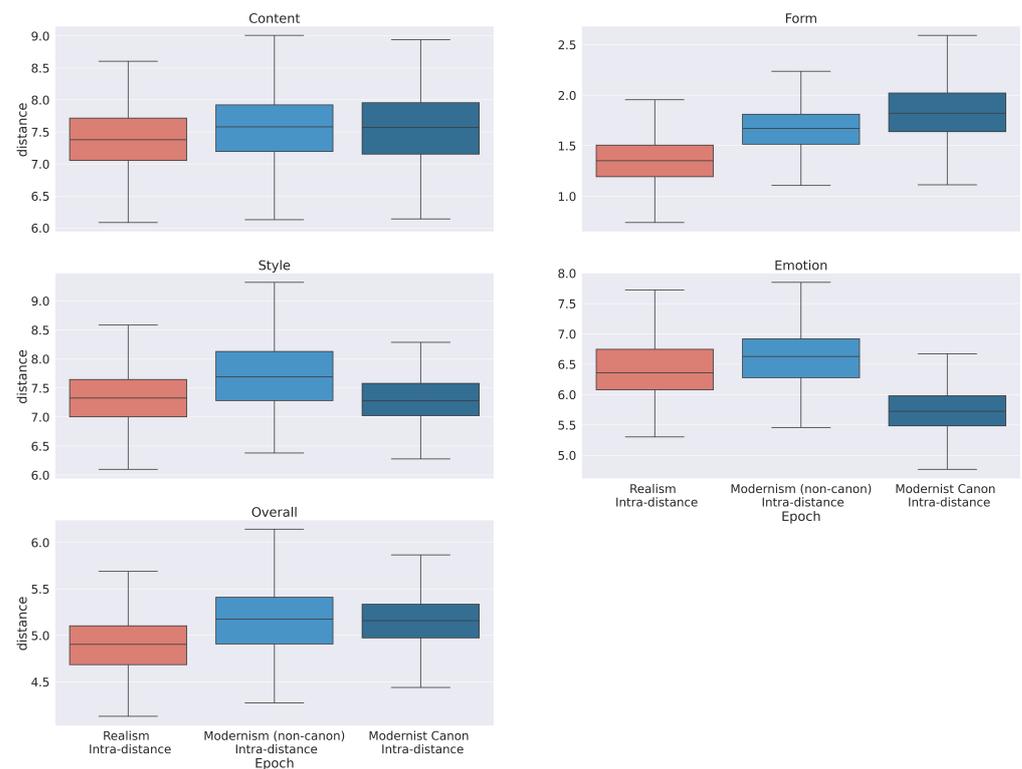


Figure 8: Distances within poems from Realism, Modernism and canonic Modernism. Distances in ‘content’, ‘style’, ‘emotion’ and ‘overall’ are measured in the space of *paraphrase-XLM-R* embeddings transformed via similarity learning (see section 4.4). Distances in ‘form’ are measured in the Feature-Form embedding space (see section 4.1). Each boxplot represents pairwise euclidean distances of 2,000 samples with a size of 20 poems.

Hypothesis 2: Modernist poetry is heterogeneous, that is, more heterogeneous than realist poetry. From now on, when we compare realism with modernism, we no longer include the naturalist poems in our calculations and visualizations, since we have seen that naturalism is located somewhere between realism and modernism. However, we now distinguish in modernist texts canonical and non-canonical authors in order to point out some peculiarities of the canonical poems.¹⁶

To test hypothesis 2, we compare the distances within realism with those within modernism (Figure 8). In all dimensions, the distances within modernism are greater than in realism, most clearly in the dimension 'form'. Thus, the hypothesis that modernist poetry is more heterogeneous than realist poetry can be confirmed by our data. However, the differences in heterogeneity are mostly small and should not be overemphasized. Modernist poems by canonical authors are slightly more heterogeneous than non-canonical poems regarding the dimension 'form'. Otherwise, the canonical poems are not characterized by greater distances among themselves than non-canonical modernist poems. On the contrary, the distances for the dimensions of style and especially emotion are much smaller within the canonical texts than within the non-canonical modernist poems. All in all, the canonical texts are no more heterogeneous than the non-canonical ones. This is surprising, since one might have expected a particularly high degree of individuality and thus heterogeneity in the canon. In any case, it must be kept in mind that the subcorpus of canonical modernist poems is very small (58 poems, 5 authors), which limits the validity of the results. Further research is needed here.

Hypothesis 3: There is a fundamental 'rupture' between modernist poetry and earlier, more traditional poetry. Some difficulties arise in testing hypothesis 3. It is not clear what exactly is meant by 'rupture' and how we should measure it. One possibility is to assume that the term 'rupture' in hypothesis 3 denotes a certain kind of literary change, namely a change that (a) is particularly large compared to other changes between periods (e.g. between romanticism and realism) and that (b) occurs abruptly, i.e. in a very short period of time. However, we do not have data on other changes between literary periods and we cannot analyse whether the shift between realism and modernism was a continuous, decades-long process or a matter of a few years.

So while we are not able to test hypothesis 3 directly, we can at least share some observations that are likely to be related. In particular, we will compare the distances between realism and modernism with distances within realism. If the distances between realism and modernism are greater than within realism, it can at least be said that modernism is different from realism.

In all dimensions, the distances between realism and modernism are larger than the distances within realism. However, the differences are not enormous. Moreover, the two-dimensional plot above (Figure 6) shows that modernist poems appear not only outside realism, but often within the realist spectrum as well. While these observations cannot falsify hypothesis 3 directly, they certainly do not confirm the hypothesis either. If anything, our results suggest that the notion of a fundamental 'rupture' between realism and modernism might be exaggerated.

16. In our study, in accordance with German literary history, Stefan George (6 poems), Hugo von Hofmannsthal (6 poems), Arno Holz (19 poems), Else Lasker-Schüler (3 poems), and Rainer Maria Rilke (24 poems) represent canonical modernism.

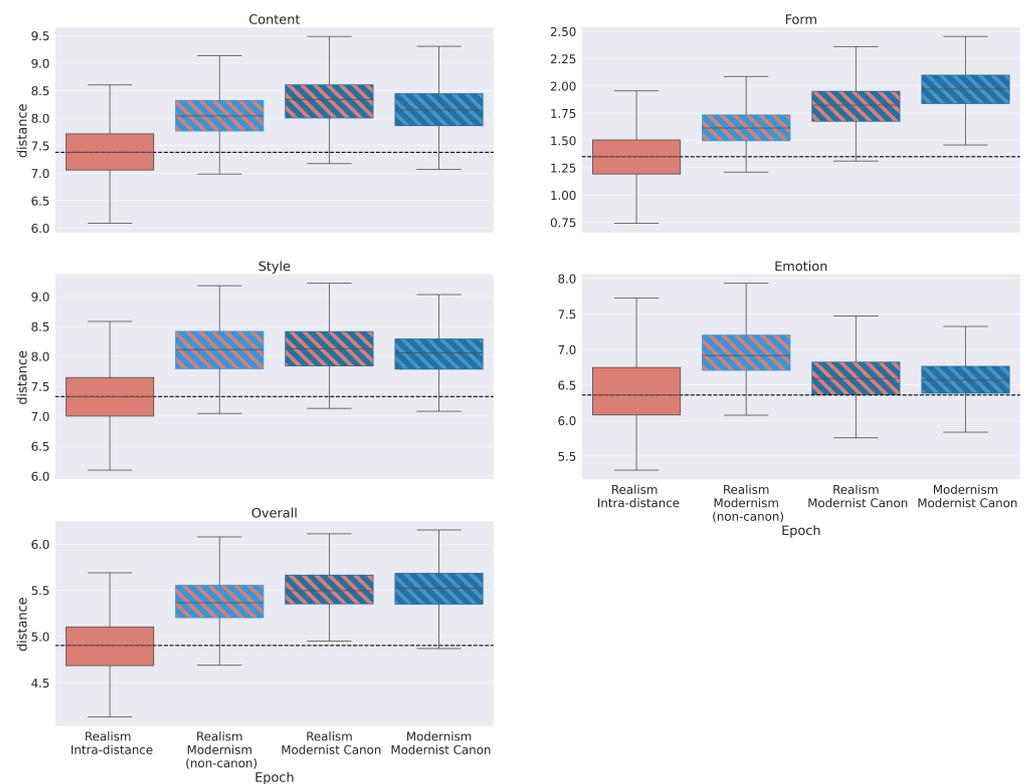


Figure 9: Distances within poems from Realism and between Realism/non-canon Modernism, Realism/canonic Modernism and non-canon Modernism/canonic Modernism. Distances in 'content', 'style', 'emotion' and 'overall' are measured in the space of *paraphrase-XLM-R* embeddings transformed via similarity learning (see section 4.4). Distances in 'form' are measured in the Feature-Form embedding space (see section 4.1). Each boxplot represents pairwise euclidean distances of 2,000 samples with a size of 20 poems.

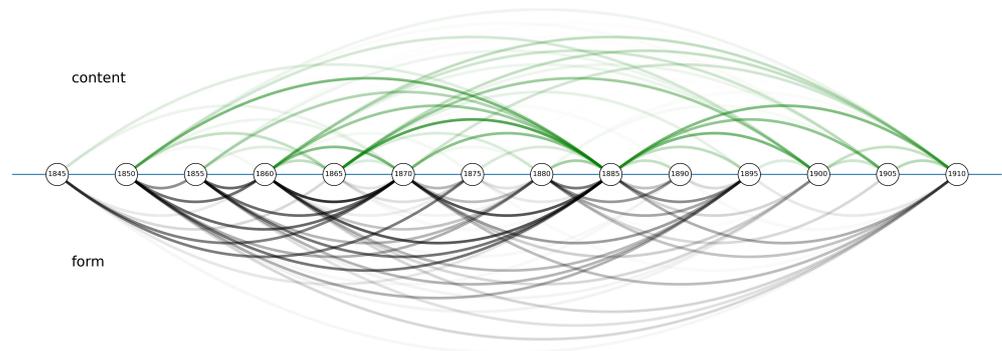


Figure 10: Graph timelines for the ‘content’ and ‘form’ dimensions based on the mean pairwise similarities of 30 poems, sampled for each 5-year time span, based on the similarity-adapted vectors of *paraphrase-XLM-R* (content) and the formal feature vectors (form). See [Figure 13](#) in [section 9](#) for a larger version of this figure.

One could assume that researchers use the metaphor of ‘rupture’ because they focus on other, namely canonical texts. The distance between realism and canonical modernism is indeed larger than the distance between realism and non-canonical modernism regarding the form, and at least a tiny bit larger for the dimensions ‘content’ and ‘overall’. But in terms of style, canonical modernism is no further from realism than non-canonical modernism, and in regards to emotion, the distance between realism and canonical modernism is even smaller than between realism and non-canonical modernism. Thus, our results do not show that the distances between canonical modernism and realism are systematically larger than between non-canonical modernism and realism. The idea that the canonical texts set a trend that the non-canonical texts follow, just not as decisively, cannot be confirmed.

One might expect the canonical modernist poems to be at least closer to the non-canonical modernist texts than to the realist poems, but this is not true either, according to our data: The distances from canonical modernist poems to realist texts on the one hand and to non-canonical modernist texts on the other hand do not differ significantly. In the case of the dimension ‘form’, the canonical modernist poems are even closer to the realist ones than to the non-canonical modernist ones.

The results for the canon are counter-intuitive and call for further research. Again, our observations may have something to do with the fact that our subcorpus of canonical texts is very small and that we only analyze short poems.

To further explore the differences between modernist and realist poetry in our vector space, we constructed a timeline from a graph network. The network was created using all pairwise distances (or similarities more precisely) between the document vectors. For all dimensions except ‘form’, the distances are based on the vectors of *paraphrase-XLM-R*, after the adaptation with similarity learning. For ‘form’, only the formal feature vector similarities were used. All distances were standardized per dimension to lie between 0 and 1 (due to the different metrics used to determine the vector distances).

Each node in the graphs represents a span of 5 years (i.e. 1865 for the span 1863-1867). The edge between two year slices is depicted by the mean distances of a sample of 30 poems – if less than 30 poems were available, poems were drawn multiple times. The alpha of one edge between two years visualizes the degree of their similarity based on

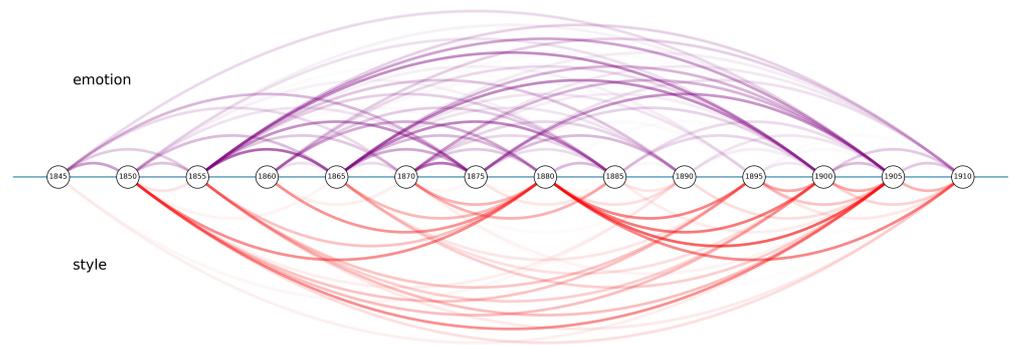


Figure 11: Graph timelines for the ‘emotion’ and ‘style’ dimensions based on the mean pairwise similarities of 30 poems, sampled for each 5-year time span, based on the similarity-adapted vectors of *paraphrase-XLM-R*. See Figure 14 in section 9 for a larger version of this figure.

the chosen poems. We only used poems where the corresponding years were manually checked and corrected by us if necessary. This amounted to 321 poems between 1845 and 1911 specifically.

From this visualization which is based not on the assignment of the poems to a period by the editors of the anthologies, but on the publication date of the poems, we can make some observations. In terms of form, we can surmise from the timeline that realist poems are more similar to each other and thus more homogenous than modernist poems are (coinciding with our findings from hypothesis 2). Additionally, the further the nodes are away from realism, the weaker the similarity becomes, implying that later modernist poems become even more estranged from the form of realist poems. The networks for content and style seem similar: both suggest a kind of split between the epochs, hinting at the possibility that modernist and realist poetry have a higher inter-than extra-epochal similarity (coinciding with our findings from hypothesis 3). The timelines could potentially not only help with identifying whether a rupture between the epochs exists or not but also when exactly such a rupture occurs. While ‘style’ shows its split around 1880, the split for ‘content’ appears to be at around 1885, implying that the change from realism to modernism first became apparent in style and then in content. For ‘emotion’, we cannot discover any kind of pattern in the timeline, suggesting that emotions thematized or expressed in the poems might contribute to a continuity between the two epochs.

In summary, we were able to confirm some important hypotheses from literary studies, while differentiating or relativizing others. Our data supports the view that naturalist poetry is closer to realism than to modernism; however, simply equating naturalist and realist poetry would not be appropriate. We showed that modernist poetry is indeed more heterogeneous than realist poetry, even though the differences are limited. Finally, our findings suggest that the change from realism to modernism was an evolutionary transition rather than a revolutionary disruption. The results encourage increased attention in literary history to processes of gradual, limited change, rather than thinking only in terms of either stasis or rupture.

The assumptions made in this section are still only based on exploratory visualizations and comparatively little data. Subsequent research could expand this subcorpus of year-annotated poems (most importantly including longer poems as already mentioned)

while further research questions could investigate these assumptions, e.g. whether the rupture between the epochs could have happened at slightly different points in time for different dimensions or whether ‘form’ really is the most suitable dimension to measure homogeneity and heterogeneity within realism and modernism for example.

In a recent article (Underwood and So 2021) discuss the question of whether the mapping of cultural artifacts to some spatial representation is not ‘distorting’ them, whether cultural relationships obey a spatial logic at all. Their experiments show that even if we have some seemingly convincing arguments against this kind of mapping, we accumulate more and more empirical evidence that it works very often astonishingly well. Our paper adds to this evidence: Textual representations in high-dimensional space seem well-suited to express even complex text models though more empirical work may expose its shortcomings in the future. In the meantime, we hope our approach can be used to reevaluate our understanding of the fundamental concept of similarity, not only in Computational Literary Studies.

6. Data Availability

Data can be found here: <https://github.com/cophi-wue/jcls2022-poem-similarity>

7. Software Availability

Software can be found here: <https://github.com/cophi-wue/jcls2022-poem-similarity>

8. Acknowledgements

This work was funded by the Deutsche Forschungsgemeinschaft as part of the SPP 2207 Computational Literary Studies in the project *The beginnings of modern poetry – Modeling literary history with text similarities*.

9. Author Contributions

Anton Ehrmantraut: Software, Writing – original draft

Thora Hagen: Visualization, Writing – original draft

Fotis Jannidis: Conceptualization, Supervision, Writing – original draft, Funding acquisition

Leonard Konle: Formal Analysis, Software, Writing – original draft

Merten Kröncke: Data curation, Methodology, Writing – original draft

Simone Winko: Data Curation, Conceptualization, Supervision, Writing – original draft, Funding acquisition

References

- Andreotti, Mario (2014). *Die Struktur der modernen Literatur. Neue Formen und Techniken des Schreibens: Erzählprosa und Lyrik*. 5th ed. Haupt Verlag.
- Anz, Thomas (2007). "Thesen zur expressionistischen Moderne". In: *Literarische Moderne. Begriff und Phänomen*. Ed. by Sabina Becker and Helmuth Kiesel. De Gruyter, 329–346.
- Arent, Wilhelm, ed. (1885). *Moderne Dichter-Charaktere*. Kanzlah.
- Arora, Sanjeev, Yingyu Liang, and Tengyu Ma (2017). "A Simple but Tough-to-Beat Baseline for Sentence Embeddings". In: *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. <https://openreview.net/pdf?id=SyK00v5xx> (visited on 11/08/2022).
- Austermühl, Elke (2000). "Lyrik der Jahrhundertwende". In: *Naturalismus, Fin de siècle, Expressionismus, 1890–1918*. Ed. by York-Gothart Mix. Hansers Sozialgeschichte der deutschen Literatur vom 16. Jahrhundert bis zur Gegenwart 7. Carl Hanser, 350–366.
- Avenarius, Ferdinand, ed. (1882). *Deutsche Lyrik der Gegenwart seit 1850. Eine Anthologie mit biographischen und bibliographischen Notizen. Aus den Quellen*. Ehlermann.
- Bär, Daniel, Torsten Zesch, and Iryna Gurevych (2011). "A Reflective View on Text Similarity". In: *Proceedings of Recent Advances in Natural Language Processing*, 515–520. <https://aclanthology.org/R11-1071> (visited on 04/22/2022).
- (2015). *Composing Measures for Computing Text Similarity*. <https://tuprints.ulb.tu-darmstadt.de/4342/> (visited on 04/22/2022).
- Becker, Sabina and Helmuth Kiesel (2007). "Literarische Moderne. Begriff und Phänomen". In: *Literarische Moderne. Begriff und Phänomen*. Ed. by Sabina Becker and Helmuth Kiesel. De Gruyter, 9–36.
- Benzmann, Hans, ed. (1904). *Moderne deutsche Lyrik*. Reclam.
- Bern, Maximilian, ed. (1877). *Deutsche Lyrik seit Goethes Tode*. Reclam.
- Bethge, Hans, ed. (1905). *Deutsche Lyrik seit Liliencron*. Hesse.
- Bierbaum, Otto Julius, ed. (1893). *Moderner Musenalmanach auf das Jahr 1893*. Albert.
- ed. (1894). *Moderner Musenalmanach auf das Jahr 1894*. Albert.
- Bonsels, Waldemar, Hans Brandenburg, Bernd Isemann, and Will Vesper, eds. (1905). *Die Erde*. Bonsels.
- Bromley, James, Isabelle Guyon, Yann LeCun, Eduard Sackinger, and Roopak Shah (1993). "Signature Verification Using a 'Siamese' Time Delay Neural Network". In: *Proceedings of the 6th International Conference on Neural Information Processing Systems*, 737–744.
- Chan, Branden, Stefan Schweter, and Timo Möller (2020). "German's Next Language Model". In: *arXiv preprint*. [10.48550/arxiv.2010.10906](https://arxiv.org/abs/10.48550/arxiv.2010.10906).
- Conradi, Hermann (1885). "Unser Credo". In: *Moderne Dichter-Charaktere*. Ed. by Wilhelm Arent. Wilhelm Friedrich, I–IV.
- Corbineau-Hoffmann, Angelika (2013). *Einführung in die Komparatistik*. 3rd ed. Erich Schmidt.
- Ehrmanntraut, Anton, Thora Hagen, Leonard Konle, and Fotis Jannidis (2021). "Type- and Token-based Word Embeddings in the Digital Humanities". In: *Proceedings of the Conference on Computational Humanities Research, CHR2021*, 16–38. http://ceur-ws.org/Vol-2989/long_paper35.pdf (visited on 04/22/2022).
- Fähnders, Walter (1998). *Avantgarde und Moderne 1890-1933*. Lehrbuch Germanistik. J. B. Metzler.

- Federmann, Herta, ed. (1908). *Der Schatzbehälter*. Steinicke & Lehmkuhl.
- Felski, Rita and Susan Stanford Friedman, eds. (2013). *Comparison*. Baltimore: John Hopkins University Press.
- Frick, Werner (2007). "Avantgarde und longue durée. Überlegungen zum Traditionsverbrauch der klassischen Moderne". In: *Literarische Moderne. Begriff und Phänomen*. Ed. by Sabine Becker and Helmuth Kiesel. De Gruyter, 97–112.
- Friedrich, Hugo (1992). *Die Struktur der modernen Lyrik. Von der Mitte des neunzehnten bis zur Mitte des zwanzigsten Jahrhunderts*. Rowohlt.
- Friedrich, Paul, ed. (1911). *Neuland. Ein Buch jüngstdeutscher Lyrik*. Borngräber.
- Gemmel, Ludwig, ed. (1898). *Die Perlenschnur. Eine Anthologie moderner Lyrik*. Schuster & Loeffler.
- Goltschnigg, Dietmar (2007). "Traditionszusammenhänge der österreichischen Moderne (am Beispiel der Heine- und Büchner-Rezeption)". In: *Literarische Moderne. Begriff und Phänomen*. Ed. by Sabine Becker and Helmuth Kiesel. De Gruyter, 169–180.
- Gomaa, Wael and Aly Fahmy (2013). "A Survey of Text Similarity Approaches". In: *International Journal of Computer Applications* (68), 13–18. [10.5120/11638-7118](https://doi.org/10.5120/11638-7118).
- Grandini, Margherita, Enrico Bagli, and Giorgio Visani (2020). "Metrics for Multi-Class Classification: an Overview". In: *arXiv preprint*. [10.48550/arxiv.2008.05756](https://arxiv.org/abs/10.48550/arxiv.2008.05756).
- Gururangan, Suchin, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith (2020). "Don't Stop Pretraining: Adapt Language Models to Domains and Tasks". In: *arXiv preprint*. [10.48550/arXiv.2004.10964](https://arxiv.org/abs/10.48550/arXiv.2004.10964).
- Häntzschel, Günter (1991). *Bibliographie der deutschsprachigen Lyrikanthologien 1840-1914*. K. G. Saur.
- Hiebel, Hans H. (2005). *Das Spektrum der modernen Poesie. Teil 1. 1900–1945*. Königshausen & Neumann.
- Huch, Margarethe, ed. (1911). *Frauenlyrik der Gegenwart*. Eckardt.
- Jacobowski, Ludwig, ed. (1899). *Neue Lieder der besten neueren Dichter für's Volk*. Liemann.
- Kiesel, Helmuth (2004). *Geschichte der literarischen Moderne. Sprache – Ästhetik – Dichtung im zwanzigsten Jahrhundert*. C.H. Beck.
- Klambauer, Günter, Thomas Unterthiner, Andreas Mayr, and Sepp Hochreiter (2017). "Self-Normalizing Neural Networks". In: *arXiv preprint*. [10.48550/arXiv.1706.02515](https://arxiv.org/abs/10.48550/arXiv.1706.02515).
- Klinger, Cornelia (2002). "Modern/Moderne/Modernismus". In: *Ästhetische Grundbegriffe. Bd. 4*. Ed. by Karlheinz Barck, Martin Fontius, Dieter Schlenstedt, and Friedrich Wolfzettel. J.B. Metzler, 121–167.
- Kneschke, Emil, ed. (1865). *Anthologie deutscher Lyriker seit 1850*. Lorck.
- Krippendorff, Klaus (2011). *Computing Krippendorff's Alpha-Reliability*. https://repository.upenn.edu/asc_papers/43 (visited on 04/22/2022).
- Lamping, Dieter (2000). *Das lyrische Gedicht. Definitionen zu Theorie und Geschichte der Gattung*. 3rd ed. Vandenhoeck & Ruprecht.
- (2008). *Moderne Lyrik*. Vandenhoeck & Ruprecht.
- (2012). *Klassiker der Moderne. Über die Kanonisierung moderner Literatur*. <https://literaturkritik.de/id/16853> (visited on 04/22/2022).
- Marelli, Marco, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, and Roberto Zamparelli (2014). "A SICK cure for the evaluation of compositional distributional semantic models". In: *Proceedings of the Ninth International Conference*

- on Language Resources and Evaluation (LREC'14). European Language Resources Association (ELRA), 216–223. http://www.lrec-conf.org/proceedings/lrec2014/pdf/363_Paper.pdf (visited on 04/22/2022).
- Mathet, Yann, Antoine Widlöcher, and Jean-Philippe Métivier (2015). “The Unified and Holistic Method Gamma (γ) for Inter-Annotator Agreement Measure and Alignment”. In: *Computational Linguistics* 3 (41), 437–479. [10.1162/COLI_a_00227](https://doi.org/10.1162/COLI_a_00227).
- McInnes, Leland, John Healy, Nathaniel Saul, and Lukas Großberger (2018). “UMAP: Uniform Manifold Approximation and Projection”. In: *The Journal of Open Source Software* 29 (3). [10.21105/joss.00861](https://doi.org/10.21105/joss.00861).
- Moltke, Max, ed. (1882). *Neuer deutscher Parnaß. Silberblicke aus der Lyrik unserer Tage*. Rühle.
- Nöth, Winfried (2008). “Stil als Zeichen”. In: *Rhetoric and Stylistics. Handbooks of Linguistics and Communication Science*. Ed. by Ulla Fix, Andreas Gardt, and Joachim Knappe. De Gruyter, 1178–1196.
- Polko, Elise, ed. (1860). *Dichtergrüße. Neuere deutsche Lyrik*. Amelang.
- Prakoso, Dimas Wibisono, Asad Abdi, and Chintan Amrit (Mar. 2021). “Short text similarity measurement methods: a review”. In: *Soft Computing* 6 (25), 4699–4723. [10.1007/s00500-020-05479-2](https://doi.org/10.1007/s00500-020-05479-2).
- Prutz, Robert, ed. (1859). *Deutsche Dichter der Gegenwart. Ein lyrisches Album*. Kober & Markgraf.
- Reimers, Nils and Iryna Gurevych (2019). “Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks”. In: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics, 3980–3990. [10.18653/v1/D19-1410](https://doi.org/10.18653/v1/D19-1410).
- (2020). “Making Monolingual Sentence Embeddings Multilingual using Knowledge Distillation”. In: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, 4512–4525. [10.18653/v1/2020.emnlp-main.365](https://doi.org/10.18653/v1/2020.emnlp-main.365).
- Renner, August, ed. (1899). *Das lyrische Wien. Eine moderne Lese*. Georg Szeliński.
- Salton, Gerard and Michael J. McGill (1983). *Introduction to Modern Information Retrieval*. McGraw-Hill.
- Sandig, Barbara (2006). *Textstilistik des Deutschen*. De Gruyter.
- Selbmann, Rolf (1999). *Die simulierte Wirklichkeit. Zur Lyrik des Realismus*. Aisthesis.
- (2007). “Die Lyrik des Realismus”. In: *Realismus. Epoche – Autoren – Werke*. Ed. by Christian Begemann. Wissenschaftliche Buchgesellschaft, 189–206.
- Shaver, Phillip, Judith Schwartz, Donald Kirson, and Cary O'Connor (1987). “Emotion Knowledge. Further Exploration of a Prototype Approach”. In: *Journal of Personality and Social Psychology* 6 (52), 1061–1086. [10.1037/0022-3514.52.6.1061](https://doi.org/10.1037/0022-3514.52.6.1061).
- Sprengel, Peter (1998). *Geschichte der deutschsprachigen Literatur 1870-1900. Von der Reichsgründung bis zur Jahrhundertwende*. Geschichte der deutschen Literatur von den Anfängen bis zur Gegenwart 9.1. C.H. Beck.
- Stockinger, Claudia (2010). *Das 19. Jahrhundert. Zeitalter des Realismus*. Akademie Verlag.
- Szubert, Benjamin, Jennifer E. Cole, Claudia Monaco, and Ignat Drozdov (2019). “Structure-preserving visualisation of high dimensional single-cell datasets”. In: *Scientific Reports* 1 (9). [10.1038/s41598-019-45301-0](https://doi.org/10.1038/s41598-019-45301-0).
- Tille, Alexander, ed. (1896). *Deutsche Lyrik von Heute und Morgen*. Neumann.

- Underwood, Ted and Richard Jean So (2021). "Can We Map Culture?" In: *Journal of Cultural Analytics* 3 (6), 32–51. [10.22148/001c.24911](#).
- Vietta, Silvio (1992). *Die literarische Moderne. Eine problemgeschichtliche Darstellung der deutschsprachigen Literatur von Hölderlin bis Thomas Bernhard*. J.B. Metzler.
- Wieland, Klaus (2019). "Die deutschsprachige Lyrik der Frühen Moderne (1890-1930)". In: *Recherches Germaniques* 14, 5–27. [10.4000/rg.976](#).
- Willatzen, Peter Johann, ed. (1875). *Blüthenzweige deutscher Lyrik nach Goethe. Eine Anthologie*. Kühnemann.
- Winko, Simone (2003). *Kodierte Gefühle. Zu einer Poetik der Emotionen in lyrischen und poetologischen Texten um 1900*. Erich Schmidt.
- Zelle, Carsten (2005). "Komparatistik und 'comparatio' – der Vergleich in der Vergleichenden Literaturwissenschaft: Skizze einer Bestandsaufnahme". In: *Komparatistik. Jahrbuch der Deutschen Gesellschaft für Allgemeine und Vergleichende Literaturwissenschaft*, 13–33.

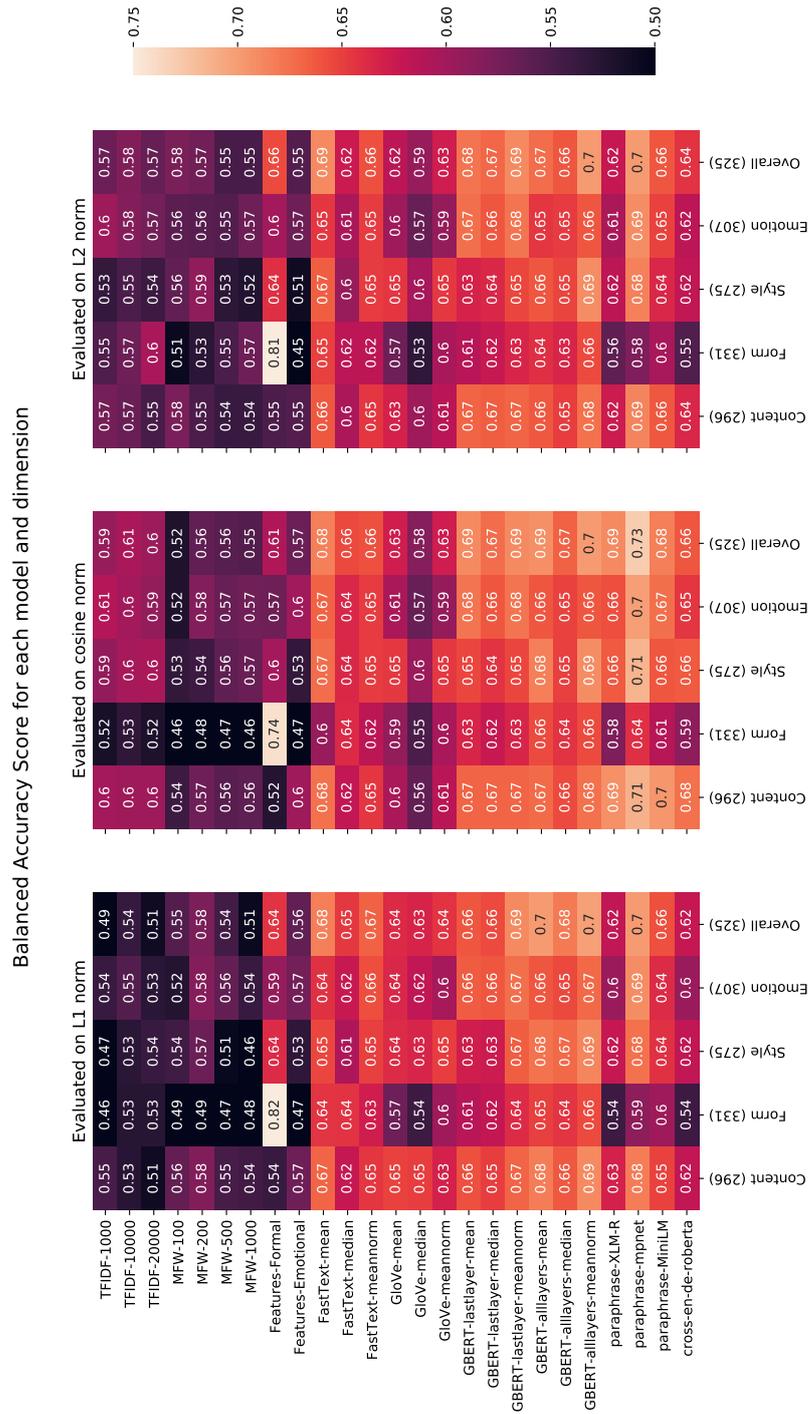


Figure 12: Balanced Accuracy Score for each model and dimension and three distance metrics: L1, Cosine, and L2. Number on the x-axis indicate class support.

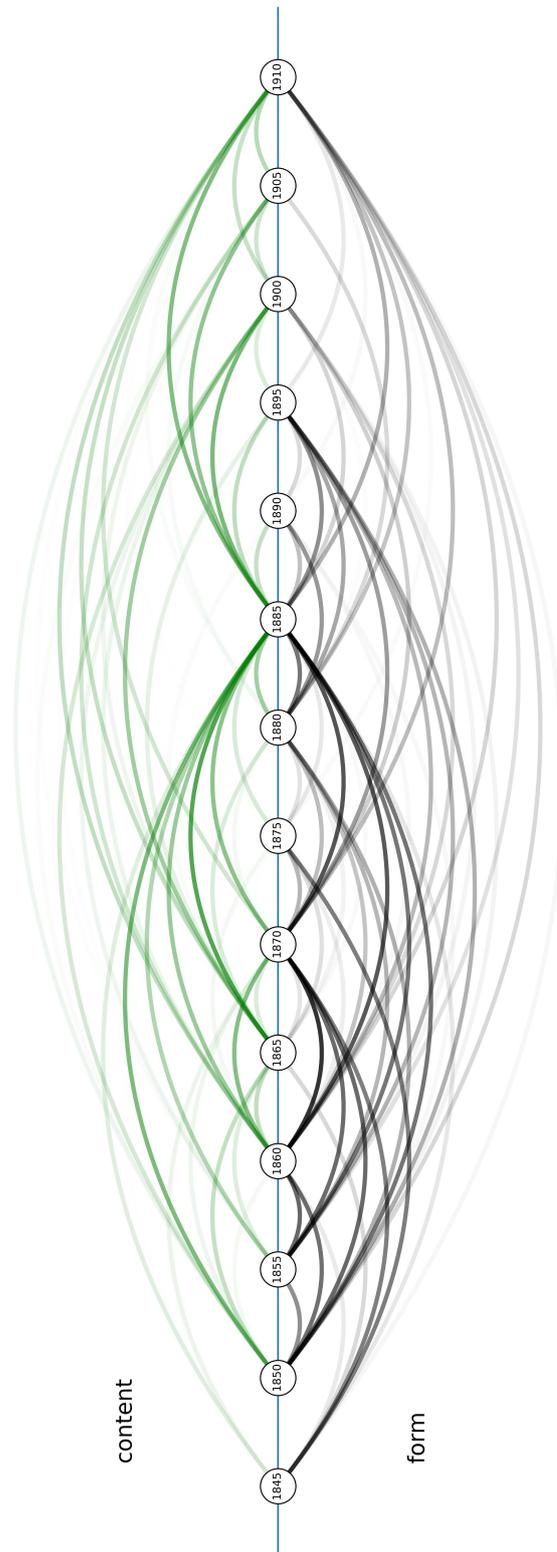


Figure 13: Graph timelines for the ‘content’ and ‘form’ dimensions based on the mean pairwise similarities of 30 poems, sampled for each 5-year time span, based on the similarity-adapted vectors of *paraphrase-XLM-R* (content) and the formal feature vectors (form).

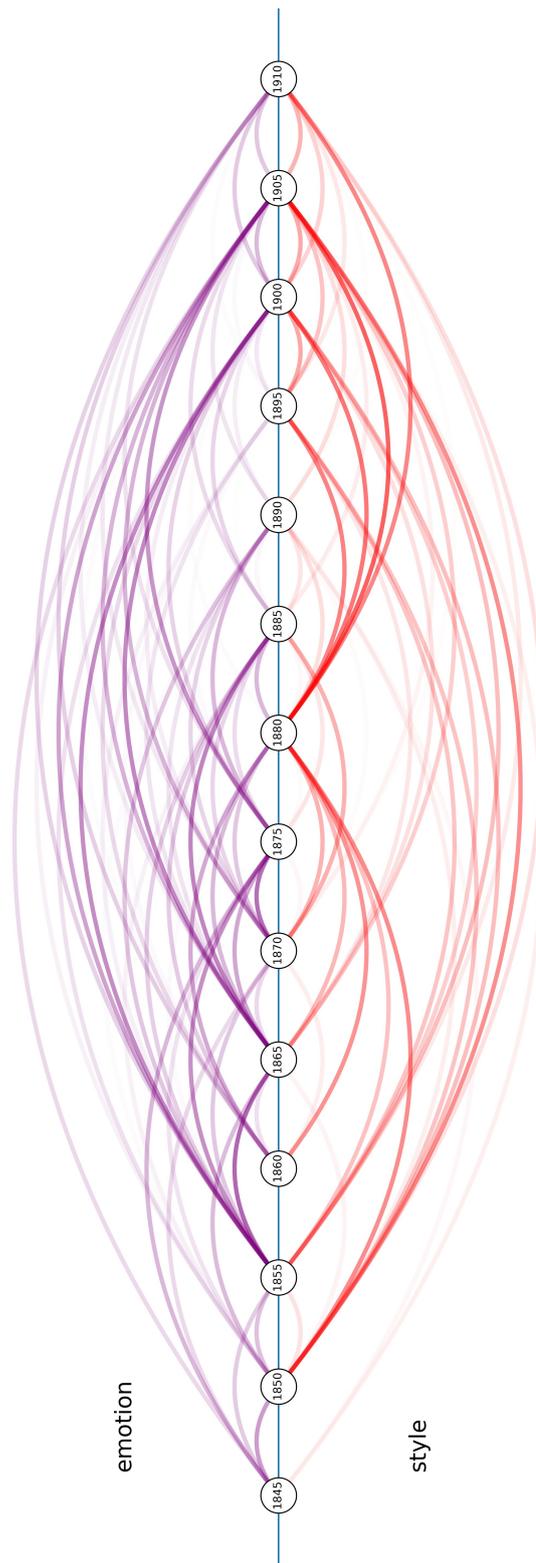


Figure 14: Graph timelines for the ‘emotion’ and ‘style’ dimensions based on the mean pairwise similarities of 30 poems, sampled for each 5-year time span, based on the similarity-adapted vectors of *paraphrase-XLM-R*.