



# Distributed Vector Representations of Folksong Motifs

Aitor Arronte Alvarez<sup>1</sup> and Francisco Gómez-Martin<sup>2</sup>(✉)

<sup>1</sup> Center for Language and Technology,  
University of Hawaii at Manoa, Honolulu, USA  
arronte@hawaii.edu

<sup>2</sup> Applied Mathematics Department,  
Technical University of Madrid, Madrid, Spain  
fmartin@etsisi.upm.es

**Abstract.** This article presents a distributed vector representation model for learning folksong motifs. A skip-gram version of word2vec with negative sampling is used to represent high quality embeddings. Motifs from the Essen Folksong collection are compared based on their cosine similarity. A new evaluation method for testing the quality of the embeddings based on a melodic similarity task is presented to show how the vector space can represent complex contextual features, and how it can be utilized for the study of folksong variation.

**Keywords:** Folksong motifs · Melodic context · Motif embedding · Word2vec

## 1 Introduction

Vector representations of words have been widely used in Natural Language Processing (NLP) tasks [18]. Following the distributional hypothesis [5, 9], vector space models represent, or embed, words that are semantically related to each other closer in a continuous vector space [24]. A recent development in vector space models is word2vec [8, 13, 15], developed for learning high-quality word vectors from large corpora.

A neural network language model for learning word-embeddings was first proposed to learn a statistical language model and a word vector representation [1]. A simpler model using a neural net with a single hidden layer to learn word vector representations, and then train a language model was later developed [14]. Word2vec follows this simpler approach in two steps: first, continuous word vectors are learned using the simpler model [14], and then an n-gram is trained using these representations.

The relation between music and language has been studied in the cognitive science literature. Even though they are treated as different cognitive faculties, both share structural characteristics and generate similar expectations on the

listener [2]. NLP methods have been adapted and adopted in Music Information Retrieval (MIR) contexts [3, 4, 6]. More specifically, word2vec was used to model musical contexts in western classical music works [10], and for chord recommendations [11]. In both cases the music compositions studied were complex polyphonic works. The work presented in this article uses a much less data intensive material: monophonic songs.

Following the distributional hypothesis in semantics, the goal of this research is to adopt the skip-gram version of the word2vec model for the distributional representation of melodic units. Several melodic features such as contour, grouping, and small size motifs seem to be part of the so called ‘Statistical Music Universals’ [17, 19]. This sequential processing of melodic units may be related to the human capacity to group and comprehend motifs as units within a melodic context. Our hypothesis is that these units may relate to each other in a melody in similar ways as words do in sentences. If that is the case, the distributional hypothesis should hold true for folksong melodies.

In the following sections a description of the skip-gram version of word2vec to learn motifs from the Essen Folksong Collection [20] is presented. We will present different similarity measures to determine how melodic context can capture the similarity of folksong motifs.

## 2 Word2vec: Representing Folksong Motifs in a Distributed Vector Space

### 2.1 Word2vec Model

In the skip-gram version of the word2vec model, the goal is to find word embeddings that can predict the surrounding words of a target word in a sentence or document [15]. Formally, the model can be defined in the following terms: given a corpus  $W$  of words  $w$  and contexts  $c$ , the network tries to predict the surrounding words of a target in a context. The objective of the skip-gram is to maximize the following log probability:

$$\arg \max_{\theta} \prod_{w \in W} \left[ \prod_{c \in C} p(c \mid w; \theta) \right] \quad (1)$$

where  $p(c \mid w; \theta)$  is calculated by the softmax function:

$$p(c \mid w; \theta) = \frac{e^{v_c \cdot v_w}}{\sum_{c' \in C} e^{v_{c'} \cdot v_w}} \quad (2)$$

where  $v_c$  and  $v_w \in R^d$  are vector representations of  $v$  and  $c$ , and  $C$  is the set of all possible contexts. The set of parameters  $\theta$  is composed of  $v_{c_i}$ ,  $v_{w_i}$  for  $w \in W$ .

Since the term  $p(w; \theta)$  involves a summation over all possible contexts  $c'$  becomes computationally very intensive, and it is normally replaced with negative sampling [15]. This article uses this sampling technique.

The cosine similarity measure is used to determine the relatedness of two embeddings. The metric for a pair of words  $w_1$  and  $w_2$  can be defined as [22]:

$$\text{cos}(w_1, w_2) = \frac{\vec{w}_1 \cdot \vec{w}_2}{\|\vec{w}_1\| \|\vec{w}_2\|} \quad (3)$$

for all similarity computations in the embedding space, where  $\vec{w}$  is a real-valued vector embedding of word  $w$ .

## 2.2 Melodic Context and Motif Representation

We are interested in studying how word2vec can model melodic context using small musical motifs instead of words. In the present research context is understood as the sequential organization of melodic units that establish statistically relevant relationships with one another in a melodic segment.

Melodic similarity and classification methods depend strongly on melodic representation [23]. Motifs from the Essen folksong collection are represented by using strings. First, intervals are codified for each song by using Music21 [7] chromatic step values from the original Kern format, and encode interval direction with Boolean values (1 for ascending and 0 for descending). For instance, the string 21 represents an ascending major second, and the string 30 a descending minor third. Repeated notes are encoded as 00.

Once the entire folksong corpus is encoded using this scheme, motifs are extracted as multi-words [15]. A multi-word is then a concatenation of two or more intervals or durations that are found in a melody adjacent to each other. For example, an intervallic multi-word of size 3 30\_00\_21 represents a descending minor third, followed by a repeated note, and by an ascending major second.

The multi-word representation of motifs is obtained following these steps:

- From a corpus of intervals we create a vocabulary of multi-word  $M$  with multi-words  $mw$  of length 2. Only those  $mw$  that occur at least 10 times are kept, based on the quality of the results from ad-hoc queries.
- For each  $mw$  in  $M$  intervals in the corpus are substituted with their corresponding  $mw$ .

The same procedure is used for  $mw$  of size 3, with the only difference that the minimum number of occurrences of  $mw$  in a corpus is set to 5. The word2vec model is run based on the corpora created obtaining vector representations for all the motifs.

## 2.3 Evaluation Methods

Evaluation of Word Embeddings (WE) falls into two categories: intrinsic and extrinsic evaluation [22]. Intrinsic evaluation methods test for syntactic or semantic relationships between words using predefined queries. Then, methods are evaluated by aggregating correlation scores. Extrinsic evaluations are performed

by using WE as the input feature for another task, and then embeddings are evaluated based on the changes in the performance of that particular task.

This study concentrates on intrinsic evaluations, more specific on relatedness and analogy. Relatedness in WE is the cosine similarity between two words. Pairs of words should have higher correlation scores when compared with human annotated semantic similarity scores [22]. Analogical reasoning was first used for testing semantic relationships between pairs of words given specific phrases: given a term  $x$  and a term  $y$  so that  $x:y$  resembles a sample relationship  $i:j$  [13]. All these evaluation methods are language specific, and have not being adapted for MIR tasks.

Given the non-linguistic nature of music, and the difficulty of interpreting WE, more so when they represent melodic motifs, a new method is presented for evaluating Melodic Embeddings (ME) based on variations of motifs and similarity measures for those motifs in relation to a reference one. The method proceeds as follows:

1. For each multi-word  $mw_i$ , where  $i = 1, 2, \dots, l$  and  $l$  is the cardinality of the vocabulary  $M$  from corpus  $C$ , we compute  $\max(\cos(mw_i, mw_j))$  for all  $j$ , and obtain the most related multi-word  $mw_i^+$  of  $mw_i$ , so that  $mw_i : w_i^+$ , and an unrelated multi-word  $mw_i^-$ , where  $\cos(mw_i, mw_i^-) < h$ , where  $h$  is an acceptable similarity threshold.
2. Chose from  $C$  a melodic segment  $c$  and replace  $mw_i$  with  $mw_i^+$  and  $mw_i^-$ , obtaining a related  $c^+$  and an unrelated  $c^-$  melodic segments. This action is performed for all segments in  $C$ .
3. Obtain  $\text{sim}(c, c^+)$  and  $\text{sim}(c, c^-)$ , where  $\text{sim}()$  is a function that computes a measure of melodic similarity between pairs of melodic segments.

The idea behind this evaluation method is that, if vector representations of motifs are of good quality, when a motif  $mw_i$  is replaced with its most similar motif  $mw_i^+$  in a melodic segment  $c$  obtaining  $c^+$ , then a melodic similarity measure should indicate that segment  $c$  is more similar to  $c^+$  than to  $c^-$ .

To measure intervallic similarity, sequences are evaluated using the mean absolute difference in intervals (*diffint*) [16]. Since this study deals with equal-length sequences, note sequences are evaluated with city block distance (*citydist*) [21], and for duration-weighted pitch sequences correlation distance (*corrdist*) [12]. In order to compute distance measures based on note sequences, a vector of pitches represented as numerical MIDI values is used.

## 2.4 Evaluating Motif Embeddings

A sample of 2000 melodic segments is randomly selected from the European subcollection from the Essen folksong corpus. Multi-word embeddings of size 2 and 3 are obtained using the skip-gram version of word2vec with context size of 5 and vector dimension of 150. We measure melodic similarity using *diffint*, *citydist*, and *corrdist* for related and unrelated multi-word melodic segments using the method presented in Sect. 2.3, and compare their means.

Wilcoxon rank sum test is performed on related and unrelated melodic segments for all similarity measures, resulting on significant differences in means for all measures ( $p\text{-value}<0.01$ ). Ad-hoc queries of intervallic motif embeddings of size 2 show similarity between motifs based on the context. For instance, Fig. 1 shows similar motifs from *mw* of size 2 (transposed to C), and Fig. 2, shows melodic examples where those motifs are present in similar melodic contexts: all three fragments contain the target motif, either 00\_20 or 20\_00 preceded by a melodic unison and followed by an ascending major second.




Fig. 1. Similar intervallic motifs from *mw* of size 2

Joseph Lieber Joseph, Germany



Die Koenigstochter Stander, Denmark



Der Brautmoeerder, Germany



Fig. 2. Fragments of European folksongs with similar intervallic motifs colored in red (Color figure online)

Next, closely related and unrelated melodic segments variations from a reference segment using the procedure described in Sect. 2.3 are computed. We compare the similarity between a reference melodic segment with its most related variation and the same reference segment with a close variation, and with a non related (or distant) variation. The cosine similarity for multi-words of size 2 and

**Table 1.** Euclidean distance between similarity scores

Results			
Measure	ref_var_ref_close_var	ref_var_ref_distant_var	mw_size
diffint	<b>6.231</b>	7.853	2
citydist	<b>8.836</b>	11.213	2
corrdist	<b>0.503</b>	2.378	2
diffint	<b>2.163</b>	4.782	3
citydist	<b>4.556</b>	7.181	3
corrdist	<b>0.713</b>	4.044	3

3 is used to select closely related and unrelated motifs. We utilize the Euclidean distance for comparing the average similarity scores of the 2000 segments and all the variants described.

The results in Table 1 show that the distance of the similarity scores between the reference segments and their variations, and the reference segments and closely related variants (*ref\_var\_ref\_close\_var*) yield better results than when we compare the reference segments and their variants, with the reference segments with distantly related variants (*ref\_var\_ref\_distant\_var*).

Overall, the results of the motif embeddings show that vector representations of folksong motifs capture contextual melodic features. Query results show how motifs can be modeled with the skip-gram version of the word2vec from monophonic contexts. One of the advantages of this method is that motifs can be easily modeled in a complete unsupervised manner given a context, and they can be retrieved using the cosine distance. At the same time, with large corpora the algorithm tends to discover multiple motifs, some of which may be irrelevant for the musicological analysis.

### 3 Conclusions

Word2vec has been used to model complex Western polyphonic classical music [10]. In this article the skip-gram version of word2vec is used to learn rich representations of monophonic motifs from the Essen folksong collection. The proposed approach shows how motifs from folksongs can be learned from a large corpus and compared with each other using the cosine similarity. This approach can be very useful for the musicological study of folksong variation using small melodic units such as motifs. It also shows, how word2vec is able to capture and model melodic contexts from monophonic songs. Future work should concentrate on the filtering of motifs based on different musicological criteria, to avoid a combinatorial explosion and to select relevant motifs for the musical analysis.

The evaluation of WE is an important research topic in the NLP literature [22]. In this article a novel computational method for evaluating the quality of motif embeddings is proposed. The approach presented shows how the model

captures different degrees of motif similarity. This evaluation method can be very useful for studying the similarity of melodic segments based on motifs and their related variants. Future work in this area should include a cognitive similarity evaluation task performed by human participants to test the quality of the embeddings.

## References

1. Bengio, Y., Ducharme, R., Vincent, P., Jauvin, C.: A neural probabilistic language model. *J. Mach. Learn. Res.* **3**(Feb), 1137–1155 (2003)
2. Besson, M., Schön, D.: Comparison between language and music. *Ann. New York Acad. Sci.* **930**(1), 232–258 (2001)
3. Boom, C.D., et al.: Large-scale user modeling with recurrent neural networks for music discovery on multiple time scales. *Multimed. Tools Appl.* **77**, 15385–15407 (2017)
4. Boulanger-Lewandowski, N., Bengio, Y., Vincent, P.: Modeling temporal dependencies in high-dimensional sequences: application to polyphonic music generation and transcription. arXiv preprint [arXiv:1206.6392](https://arxiv.org/abs/1206.6392) (2012)
5. Clark, S.: Vector space models of lexical meaning. In: Lappin, S., Fox, C. (eds.) *The Handbook of Contemporary Semantic Theory*, pp. 463–472. Wiley-Blackwell, Hoboken (2015)
6. Conklin, D., Witten, I.H.: Multiple viewpoint systems for music prediction. *J. New Music Res.* **24**(1), 51–73 (1995)
7. Cuthbert, M.S., Ariza, C.: Music21: A toolkit for computer-aided musicology and symbolic music data. In: ISMIR. Utrecht, The Netherlands (2010)
8. Goldberg, Y., Levy, O.: word2vec explained: deriving mikolov et al.’s negative-sampling word-embedding method. arXiv preprint [arXiv:1402.3722](https://arxiv.org/abs/1402.3722) (2014)
9. Harris, Z.S.: Distributional structure. *Word* **10**(2–3), 146–162 (1954)
10. Herremans, D., Chuan, C.H.: Modeling musical context with word2vec. arXiv preprint [arXiv:1706.09088](https://arxiv.org/abs/1706.09088) (2017)
11. Huang, C.Z.A., Duvenaud, D., Gajos, K.Z.: Chordripple: recommending chords to help novice composers go beyond the ordinary. In: *Proceedings of the 21st International Conference on Intelligent User Interfaces*, pp. 241–250. ACM, Sonoma (2016)
12. Janssen, B., van Kranenburg, P., Volk, A.: Finding occurrences of melodic segments in folk songs employing symbolic similarity measures. *J. New Music Res.* **46**(2), 118–134 (2017)
13. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. arXiv preprint [arXiv:1301.3781](https://arxiv.org/abs/1301.3781) (2013)
14. Mikolov, T., Kopecky, J., Burget, L., Glembek, O., et al.: Neural network based language models for highly inflective languages. In: *2009 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2009*, pp. 4725–4728. IEEE, Taipei (2009)
15. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. In: *Advances in Neural Information Processing Systems*, pp. 3111–3119. Lake Tahoe, Nevada (2013)
16. Müllensiefen, D., Frieler, K., et al.: Cognitive adequacy in the measurement of melodic similarity: algorithmic vs. human judgments. *Comput. Musicology* **13**(2003), 147–176 (2004)

17. Nettl, B.: An ethnomusicologist contemplates universals in musical sound and musical culture. In: Brown, S., Nils, L., Wallin, B.M. (eds.) *The Origins of Music*, pp. 463–472. MIT Press, Cambridge (2000)
18. Rumelhart, D.E., Hinton, G.E., Williams, R.J.: Learning representations by back-propagating errors. *Nature* **323**(6088), 533 (1986)
19. Savage, P.E., Brown, S., Sakai, E., Currie, T.E.: Statistical universals reveal the structures and functions of human music. *Proc. National Acad. Sci.* **112**(29), 8987–8992 (2015)
20. Schaffrath, H., Huron, D.: The essen folksong collection in the humdrum kern format. Technical report, Center for Computer Assisted Research in the Humanities, Menlo Park, CA, USA (1995)
21. Scherrer, D.K., Scherrer, P.H.: An experiment in the computer measurement of melodic variation in folksong. *J. Am. Folklore* **84**(332), 230–241 (1971)
22. Schnabel, T., Labutov, I., Mimno, D., Joachims, T.: Evaluation methods for unsupervised word embeddings. In: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pp. 298–307. Lisbon, Portugal (2015)
23. Toivainen, P., Eerola, T.: A computational model of melodic similarity based on multiple representations and self-organizing maps. In: *Proceedings of the seventh international conference on music perception and cognition*, Sydney. Causal Productions, Adelaide, pp. 236–239 (2002)
24. Turney, P.D., Pantel, P.: From frequency to meaning: vector space models of semantics. *J. Artif. Intell. Res.* **37**, 141–188 (2010)