Cross-lingual Word Embeddings in Hyperbolic Space

Anonymous ACL submission

Abstract

Cross-lingual word embeddings can be applied to several natural language processing applications across multiple languages. Unlike prior works that use word embeddings based on the Euclidean space, this short paper presents a simple and effective cross-lingual Word2Vec model that adapts to the Poincaré ball model of hyperbolic space to learn unsupervised cross-lingual word representations from a German-English parallel corpus. It has been shown that hyperbolic embeddings can capture and preserve hierarchical relationships. We evaluate the model on both hypernymy and analogy tasks. The proposed model achieves comparable performance with the vanilla Word2Vec model on the cross-lingual analogy task, the hypernymy task shows that the cross-lingual Poincaré Word2Vec model can capture latent hierarchical structure from free text across languages, which are absent from the Euclidean-based Word2Vec representations. Our results show that by preserving the latent hierarchical information, hyperbolic spaces can offer better representations for cross-lingual embeddings.

1 Introduction

In Natural Language Processing (NLP), cross-lingual word embeddings refer to the representations of words from two or more languages in a joint feature space. Prior works have demonstrated the use of these continuous representations in a variety of NLP tasks such as information retrieval (Zoph et al., 2016), semantic textual similarity (Cer et al., 2017), knowledge transfer (Gu et al., 2018), lexical analysis (Dong and De Melo, 2018), plagiarism detection (Alzahrani and Aljuaid, 2020), etc. across different languages.

Natural language data possesses latent tree-like hierarchies in linguistic ontologies (e.g., hypernyms, hyponyms) (Dhingra et al., 2018; Astefanoaei and Collignon, 2020) such as the taxonomy of WordNet (Miller, 1998) for a language. From the statistics of word co-occurrence in training text, word embeddings models in Euclidean space can capture associations of words and their semantic relatedness. However, they fail to capture asymmetric word relations, including the latent hierarchical structure of words such as *specificity* (Dhingra et al., 2018). For example, 'bulldog' is more specific than 'dog.' The use of non-Euclidean spaces has recently been advocated as alternatives to the conventional Euclidean space to infer latent hierarchy from the language data (Nickel and Kiela, 2017, 2018; Dhingra et al., 2018; Tifrea et al., 2018). Learning cross-lingual hierarchies such as cross-lingual types-sub types and hypernyms-hyponyms, is useful for tasks like cross-lingual lexical entailment, textual entailment, machine translation, etc. (Vulić et al., 2019).

This paper builds upon previous work in monolingual hyperbolic Word2Vec[1] modeling from Tifrea et al. (2018) by learning cross-lingual hyperbolic embeddings from a parallel corpus, As a first step, we adopt the German-English parallel corpus from Wołk and Marasek (2014). We summarize the main contributions as follows: (1) To the best of our knowledge, we are the first to attempt at learning cross-lingual embeddings of natural language data using non-Euclidean geometry; (2) we evaluate the hyperbolic embeddings on cross-lingual HyperLex hypernym task to evaluate its performance in learning latent hierarchies from free text and how a word's specificity correlates to its embedding's norm. We also compare the hyperbolic Word2Vec embeddings with the vanilla Word2Vec embeddings in the cross lingual analogy task.

2 Related Work

2.1 Cross-lingual Word Embeddings

Cross-lingual word representations have been a subject of extensive research (Upadhyay et al., 2016;

[1] The hyperbolic Word2Vec model is not described in Tifrea et al. (2018)'s paper, but available in the corresponding codebase

Ruder et al., 2019). Recent advances in the field can be grouped into unsupervised, supervised, and joint learning algorithms. Unsupervised models (Lample et al., 2017; Artetxe and Schwenk, 2019; Chen et al., 2018) exploit existing monolingual word embeddings, followed by various cross-lingual alignment procedures. Supervised models (Mikolov et al., 2013; Smith et al., 2017; Grave et al., 2018) learn a mapping function from a source embedding space to the target embedding space based on different objective criteria. Joint learning models (Coulmance et al., 2015; Josifoski et al., 2019; Sabet et al., 2019; Lachraf et al., 2019) use parallel corpora to train bilingual embeddings in the same space jointly. This work adopts the settings similar to the joint learning model for embedding alignments by Lachraf et al. (2019).

2.2 Hyperbolic Word Embeddings

Hyperbolic spaces offer a continuous representation for embedding tree-like structures with arbitrarily low distortion (Sala et al., 2018; Chami et al., 2020). Word embeddings in hyperbolic spaces have been applied to diverse NLP applications such as text classification (Zhu et al., 2020), learning taxonomy (Astefanoaei and Collignon, 2020), and concept hierarchy (Le et al., 2019). By using hyperbolic space these applications were able to outperform their euclidean counterparts by exploiting the benefits of hierarchical structure of the text data with high quality embedding which capture similarity and generality of concept together enforce transitivity of the is-a-relations in a smaller embedding space (Le et al., 2019). Some recent work use supervised models (Nickel and Kiela, 2017, 2018; Ganea et al., 2018) that require external information on word relations such as WordNet or ConceptNet in addition to free text corpora to learn word and sentence embeddings in the hyperbolic space. Nickel and Kiela (2017) consider a non-parametric method to learn hierarchical representation from a lookup table for symbolic data. Ganea et al. (2018) propose a supervised method to learn embeddings for an acyclic graph structure of words. Unsupervised word embedding models (Leimeister and Wilson, 2018; Dhingra et al., 2018; Tifrea et al., 2018) which can directly learn from text corpora have been recently applied in the hyperbolic spaces. Leimeister and Wilson (2018) employ the skip-gram with negative sampling architecture of the Word2Vec model for learning word embeddings from free text. Dhingra et al. (2018) present a two-step model to embed a co-occurrence graph of words and map the output of the encoder to the Poincaré ball using the algorithm from Nickel and Kiela (2017). Tifrea et al. (2018) remodel the GloVe algorithm to learn unsupervised word representation in hyperbolic spaces.

3 Methodology

3.1 Hyperbolic Space

Hyperbolic space in Riemannian geometry is a homogeneous space of constant negative curvature with special geometric properties. Hyperbolic space can endow infinite trees to have nearly isometric embeddings. We embed words using the Poincaré ball model of the hyperbolic space.

The Poincaré Ball. The Poincaré ball model $\mathcal{B}^n$ of n-dimensional hyperbolic geometry is a manifold equipped with a Riemannian metric g^B. Formally, an n-dimensional Poincaré unit ball is defined as $(\mathcal{B}^n, g^B)$ and the metric g^B is conformal to the Euclidean metric g^E as $g^B = {\lambda_x}^2 . g^E$. Where $\lambda_x = \frac{2}{1-||x||^2}$, $x \in \mathcal{B}^n$, and $||.||$ stands for the Euclidean norm. Notably, the hyperbolic distance $d_{\mathcal{B}^n}$ between n-dimensional points $(x, y) \in \mathcal{B}^n$ in the Poincaré ball is defined as:

$$d_{\mathcal{B}^n}(x, y) = \operatorname{arcosh}\left(1 + 2\frac{||x-y||^2}{(1-||x||^2)(1-||y||^2)}\right) \quad (1)$$

where $\operatorname{arcosh}(w) = \ln(w + \sqrt{w^2 - 1})$ is the inverse of hyperbolic cosine function. Using ambient Euclidean geometry, the geodesic distance between points (x, y) can be induced using Equation (1) as $d_{\mathcal{B}^n}(x, y) = \operatorname{arcosh}\left(1 + \frac{1}{2}\lambda_x\lambda_y||x-y||^2\right)$. This indicates that the distance changes evenly w.r.t. $||x||$ and $||y||$, which is a key point to learning continuous representation for hierarchical structures (Chen et al., 2020).

3.2 Hyperbolic Cross-lingual Word Embedding

We first adopt the mono-lingual hyperbolic word embedding from a model defined in the work by Tifrea et al. (2018). We extend it to cross-lingual hyperbolic word embedding by using parallel text corpora input to capture word relationsships through bilingual word co-occurrence statistics. Tifrea et al. (2018) added a hyperparameter function h on the distance between word and context pairs in the hyperbolic Word2Vec's objective function. Hence, the effective distance function in the objective function becomes $h(d_{\mathcal{B}^n}(x, y))$.

Hyperbolic word embeddings have shown to embed general words near the origin and specific words towards the edges – we attempt to exploit this property to identify latent hierarchies and in hypernym evaluation task by using the Poincaré norms of the words to determine their hierarchy as words with higher norm will be more specific, i.e., lower in hierarchy (Nickel and Kiela, 2017; Dhingra et al., 2018; Linzhuo et al., 2020). We evaluate the hyperbolic model on the cross-lingual analogy task to compare it with its Euclidean counterpart.

3.3 Cross-lingual Alignment

To train the cross-lingual Word2Vec model in the hyperbolic space, we perform a pre-processing step of word-to-word alignment as defined by Lachraf et al. (2019) using parallel sentences from a bilingual parallel corpus. We generate word-to-word alignment by matching the indices of tokens from both languages in parallel sentences.

3.4 Evaluation Methodology

Hypernymy Evaluation. We perform hypernymy evaluation to assess performance of the proposed model based on learning the latent hierarchical structure from free text. In the hypernymy evaluation task, given a word pair (u, v), we evaluate $is\text{-}a(u, v)$ i.e., to what degree u is of type v.

For English, German and cross-lingual German-English hypernymy evaluation, we use the HyperLex benchmark Vulić et al. (2017, 2019), which contains word pairs (u, v) and a corresponding degree to which u is of type v i.e. the $is\text{-}a$ score. This score has been obtained by human annotators, scored by the degree of typicality and semantic category membership (Vulić et al., 2017). For example, in the HyperLex dataset, $is\text{-}a(chemistry, science) = 6.00$ and $is\text{-}a(chemistry, knife) = 0.50$ as chemistry is a type of science but not a type of knife.

To generate the $is\text{-}a$ score we follow the same approach as used by Nickel and Kiela (2017):

$$is\text{-}a(u, v) = -(1+\alpha(||v||-||u||))d_{\mathcal{B}^n}(u, v) \quad (2)$$

The evaluation is performed by calculating the Spearman correlation between the ground-truth score and the predicted score. Note that our model is not trained on any hypernymy detection task but tries to learn latent hierarchy from free text.

Word	Closest Children
Species	arten, gattung, subspecies, unterfamilie
Physics	astrophysik, astrophysics, mechanik
Molekülen	atomen, protonen, elektronen, ionen
Orchestra	symphony, philharmonic, concerto
Regierung	governments, regierungen, bundesregierung

Table 1: For a given word in the left column, this table shows the top closest children using a 100Dim with bias hyperbolic Word2Vec model. Note that the children consist of both English and German words.

	HyperLex		
Hyperbolic Model	**English**	**German**	**Cross**
	en	de	de-en
100D	0.166	0.130	0.150
100D w/ bias	0.175	0.104	0.162
120D w/ bias	**0.192**	0.120	**0.179**
300D w/ bias	0.183	**0.125**	0.155

Table 2: Spearman correlations from different hyperbolic Word2Vec models on the English, German and German-English HyperLex dataset for hypernymy evaluation. Best results are in bold.

Cross-lingual Analogy Evaluation. The analogy evaluation task is one of the standard intrinsic evaluations for word embeddings. In cross-lingual analogy evaluation task, given a word pair (w_1, w_2) in one language, and a word w_3 in the other language, the goal is to predict the word w_4^* such that w_4^* is related to w_3 same way w_2 is related to w_1. For example, as prince (w_1) is to princess (w_2), prinz (w_3; German equivalent for *prince*) is to prinzessin (w_4^*; German equivalent for *princess*). For evaluating cross-lingual analogy for the German and English language, we use the cross-lingual analogy dataset provided by Brychcín et al. (2018).

4 Experiments & Results

4.1 Dataset

This paper uses the Wikipedia corpus of parallel sentences extracted by Wołk and Marasek (2014) to train the model. The dataset is accessed through OPUS (Tiedemann, 2012). The corpus consists of ~2.5 million parallel aligned German-English sentence pairs with 43.5 million German tokens and 58.4 million English tokens.

"music"			"art"			"film"			"chemistry"		
Word	Count	Norm	Word	Count	Norm	Word	Count	Norm	Word	Count	Norm
music	33167	0.607	art	28551	0.606	film	61682	0.606	chemistry	3165	0.628
musik	10637	0.608	arts	13888	0.623	films	7185	0.607	chemie	2530	0.629
musical	6585	0.612	design	11558	0.624	drama	4948	0.617	chemiker	908	0.620
musicians	1955	0.628	skulptur	480	0.632	comedy	3937	0.630	chemischen	628	0.647
filmmusik	278	0.640	kunstgalerie	102	0.665	stummfilm	179	0.648	organischen	344	0.651

Table 3: Words in order of increasing hyperbolic norm which are related to the word indicated in the top row along with their counts in the corpus. General words have a lower norm and specific words have a higher norm.

Model Type	Dim	Bias term	Accuracy
Vanilla	20D	✗	16.8
Poincaré	20D	✗	20.5
Vanilla	40D	✗	25.4
Poincaré	40D	✗	26.5
Vanilla	80D	✗	30.8
Poincaré	80D	✗	28.7
Vanilla	180D	✓	36.1
Poincaré	180D	✓	29.3

Table 4: Accuracy on the cross-lingual analogy task.

4.2 Experimental Settings

We reference Tifrea et al. (2018)'s Poincaré Word2Vec implementation[2] and extended it to learn cross-lingual word embeddings. We set the minimum frequency of words in the vocabulary to 100, and a window size of 5. The models use Negative-Log-Likelihood loss. The non-hyperbolic vanilla Word2Vec uses Stochastic Gradient Descent optimizer, whereas hyperbolic Word2Vec uses Weighted Full Riemannian Stochastic Gradient Descent optimizer (Bonnabel, 2013). For hyperbolic embeddings, the hyperparameter h is set to $cosh^2(x)$. During the analogy evaluation, we use the cosine distance instead of Poincaré distance for hyperbolic models. We use the hypernymysuite[3] for hypernymy evaluation (Roller et al., 2018).

4.3 Evaluation Results

Hypernymy Evaluation. We present the top closest children of selected words in Table 1. As described in Section 3.2, the closest children are calculated by finding the target word's (t) nearest neighbours (N) and extracting the neighbour $n \in N$ such that $||n||_p > ||t||_p$, where $||.||_p$ is the Poincaré norm. We observe that the model is able to find the hyponyms of the words using the closest children across languages. For example, the children of 'Physics' are its subtypes – 'astrophysik' (astrophysics), 'astrophysics', 'mechanik' (mechanics), and 'biophysics'.

Table 2 reports the results on the hypernymy evaluation task. Although the models were not trained on hypernymy tasks, we observe that they could still learn some latent hierarchies from the free text across languages. Word pairs with out-of-vocabulary words were ignored during evaluation.

Table 3 shows lists of related words in order of increasing hyperbolic norm and specificity, similar to Dhingra et al. (2018)'s evaluation. We show counts of these words in the corpus. Higher the count, more generic the word, and has a smaller hyperbolic norm. The Spearman correlation between $1/f$, where f is the frequency of a word in the corpus, and its embedding's hyperbolic norm is 0.747 using a 300D w/bias Poincaré model.

Cross-lingual Analogy Evaluation. Table 4 reports the results on the cross-lingual analogy task. We observe that for 20D models, hyperbolic model outperformed the vanilla model. For higher dimension models, hyperbolic Word2Vec performed on par with its Euclidean counterpart. Similar to hypernymy evaluation, analogy pairs with out-of-vocabulary words were ignored during evaluation.

5 Conclusion and Future Work

This work adapts a monolingual hyperbolic Word2Vec model and extend to cross-lingual embeddings. We observe that the hyperbolic Word2Vec embeddings are competent on cross-lingual analogy task. The hypernymy evaluation show that it also captures some latent hierarchies across languages without being trained on a hypernymy task. Future work will include extrinsic evaluation of hyperbolic cross-lingual word embeddings on downstream tasks such as machine translation, cross-lingual textual entailment detection, cross-lingual taxonomy learning, etc.

[2] https://github.com/alex-tifrea/poincare_glove
[3] https://github.com/facebookresearch/hypernymysuite

References

Salha Alzahrani and Hanan Aljuaid. 2020. Identifying cross-lingual plagiarism using rich semantic features and deep neural networks: A study on arabic-english plagiarism cases. *Journal of King Saud University-Computer and Information Sciences*.

Mikel Artetxe and Holger Schwenk. 2019. Margin-based parallel corpus mining with multilingual sentence embeddings. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3197–3203.

Maria Astefanoaei and Nicolas Collignon. 2020. Hyperbolic embeddings for music taxonomy. In *Proceedings of the 1st Workshop on NLP for Music and Audio (NLP4MusA)*, pages 38–42.

Silvere Bonnabel. 2013. Stochastic gradient descent on riemannian manifolds. *IEEE Transactions on Automatic Control*, 58(9):2217–2229.

Tomáš Brychcín, Stephen Eugene Taylor, and Lukáš Svoboda. 2018. Cross-lingual word analogies using linear transformations between semantic spaces.

Daniel Cer, Mona Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. 2017. Semeval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14.

Ines Chami, Albert Gu, Vaggos Chatziafratis, and Christopher Ré. 2020. From trees to continuous embeddings and back: Hyperbolic hierarchical clustering. *arXiv preprint arXiv:2010.00402*.

Boli Chen, Xin Huang, Lin Xiao, Zixin Cai, and Liping Jing. 2020. Hyperbolic interaction model for hierarchical multi-label classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 7496–7503.

Xilun Chen, Yu Sun, Ben Athiwaratkun, Claire Cardie, and Kilian Weinberger. 2018. Adversarial deep averaging networks for cross-lingual sentiment classification. *Transactions of the Association for Computational Linguistics*, 6:557–570.

Jocelyn Coulmance, Jean-Marc Marty, Guillaume Wenzek, and Amine Benhalloum. 2015. Trans-gram, fast cross-lingual word-embeddings. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1109–1113.

Bhuwan Dhingra, Christopher Shallue, Mohammad Norouzi, Andrew Dai, and George Dahl. 2018. Embedding text in hyperbolic spaces. In *Proceedings of the Twelfth Workshop on Graph-Based Methods for Natural Language Processing (TextGraphs-12)*, pages 59–69.

Xin Dong and Gerard De Melo. 2018. Cross-lingual propagation for deep sentiment analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.

Octavian Ganea, Gary Bécigneul, and Thomas Hofmann. 2018. Hyperbolic entailment cones for learning hierarchical embeddings. In *International Conference on Machine Learning*, pages 1646–1655. PMLR.

Édouard Grave, Piotr Bojanowski, Prakhar Gupta, Armand Joulin, and Tomáš Mikolov. 2018. Learning word vectors for 157 languages. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*.

Jiatao Gu, Hany Hassan, Jacob Devlin, and Victor OK Li. 2018. Universal neural machine translation for extremely low resource languages. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 344–354.

Martin Josifoski, Ivan S Paskov, Hristo S Paskov, Martin Jaggi, and Robert West. 2019. Crosslingual document embedding as reduced-rank ridge regression. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, pages 744–752.

Raki Lachraf, Youcef Ayachi, Ahmed Abdelali, Didier Schwab, et al. 2019. Arbengvec: Arabic-english cross-lingual word embedding model. In *Proceedings of the Fourth Arabic Natural Language Processing Workshop*, pages 40–48.

Guillaume Lample, Alexis Conneau, Ludovic Denoyer, and Marc'Aurelio Ranzato. 2017. Unsupervised machine translation using monolingual corpora only. *arXiv preprint arXiv:1711.00043*.

Matthew Le, Stephen Roller, Laetitia Papaxanthos, Douwe Kiela, and Maximilian Nickel. 2019. Inferring concept hierarchies from text corpora via hyperbolic embeddings. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3231–3241.

Matthias Leimeister and Benjamin J Wilson. 2018. Skip-gram word embeddings in hyperbolic space. *arXiv preprint arXiv:1809.01498*.

Li Linzhuo, Wu Lingfei, and Evans James. 2020. Social centralization and semantic collapse: Hyperbolic embeddings of networks and text. *Poetics*, 78:101428.

Tomas Mikolov, Quoc V Le, and Ilya Sutskever. 2013. Exploiting similarities among languages for machine translation. *arXiv preprint arXiv:1309.4168*.

George A Miller. 1998. *WordNet: An electronic lexical database*. MIT press.

Maximillian Nickel and Douwe Kiela. 2017. Poincaré embeddings for learning hierarchical representations. *Advances in Neural Information Processing Systems*, 30:6338–6347.

Maximillian Nickel and Douwe Kiela. 2018. Learning continuous hierarchies in the lorentz model of hyperbolic geometry. In *International Conference on Machine Learning*, pages 3779–3788. PMLR.

Stephen Roller, Douwe Kiela, and Maximilian Nickel. 2018. Hearst patterns revisited: Automatic hypernym detection from large text corpora. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics.

Sebastian Ruder, Ivan Vulić, and Anders Søgaard. 2019. A survey of cross-lingual word embedding models. *Journal of Artificial Intelligence Research*, 65:569–631.

Ali Sabet, Prakhar Gupta, Jean-Baptiste Cordonnier, Robert West, and Martin Jaggi. 2019. Robust cross-lingual embeddings from parallel sentences. *arXiv preprint arXiv:1912.12481*.

Frederic Sala, Chris De Sa, Albert Gu, and Christopher Ré. 2018. Representation tradeoffs for hyperbolic embeddings. In *International conference on machine learning*, pages 4460–4469. PMLR.

Samuel L Smith, David HP Turban, Steven Hamblin, and Nils Y Hammerla. 2017. Offline bilingual word vectors, orthogonal transformations and the inverted softmax. *arXiv preprint arXiv:1702.03859*.

Jörg Tiedemann. 2012. Parallel data, tools and interfaces in opus. In *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC'12)*, Istanbul, Turkey. European Language Resources Association (ELRA).

Alexandru Tifrea, Gary Becigneul, and Octavian-Eugen Ganea. 2018. Poincare glove: Hyperbolic word embeddings. In *International Conference on Learning Representations*.

Shyam Upadhyay, Manaal Faruqui, Chris Dyer, and Dan Roth. 2016. Cross-lingual models of word embeddings: An empirical comparison. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1661–1670.

Ivan Vulić, Daniela Gerz, Douwe Kiela, Felix Hill, and Anna Korhonen. 2017. HyperLex: A large-scale evaluation of graded lexical entailment. *Computational Linguistics*, 43(4):781–835.

Ivan Vulić, Simone Paolo Ponzetto, and Goran Glavaš. 2019. Multilingual and cross-lingual graded lexical entailment. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4963–4974, Florence, Italy. Association for Computational Linguistics.

Krzysztof Wołk and Krzysztof Marasek. 2014. Building subject-aligned comparable corpora and mining it for truly parallel sentence pairs. *Procedia Technology*, 18:126–132. International workshop on Innovations in Information and Communication Science and Technology, IICST 2014, 3-5 September 2014, Warsaw, Poland.

Yudong Zhu, Di Zhou, Jinghui Xiao, Xin Jiang, Xiao Chen, and Qun Liu. 2020. Hypertext: Endowing fasttext with hyperbolic geometry. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 1166–1171.

Barret Zoph, Deniz Yuret, Jonathan May, and Kevin Knight. 2016. Transfer learning for low-resource neural machine translation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1568–1575.