

Learning Character-level Compositionality with Visual Features

Frederick Liu¹, Han Lu¹, Chieh Lo², Graham Neubig¹

¹Language Technology Institute

²Electrical and Computer Engineering

Carnegie Mellon University, Pittsburgh, PA 15213

{fliu1, hlu2, gneubig}@cs.cmu.edu

chiehl@andrew.cmu.edu

Abstract

Previous work has modeled the compositionality of words by creating character-level models of meaning, reducing problems of sparsity for rare words. However, in many writing systems compositionality has an effect even on the character-level: the meaning of a character is derived by the sum of its parts. In this paper, we model this effect by creating embeddings for characters based on their visual characteristics, creating an image for the character and running it through a convolutional neural network to produce a visual character embedding. Experiments on a text classification task demonstrate that such model allows for better processing of instances with rare characters in languages such as Chinese, Japanese, and Korean. Additionally, qualitative analyses demonstrate that our proposed model learns to focus on the parts of characters that carry semantic content, resulting in embeddings that are coherent in visual space.

1 Introduction

Compositionality—the fact that the meaning of a complex expression is determined by its structure and the meanings of its constituents—is a hallmark of every natural language (Frege and Austin, 1980; Szabó, 2010). Recently, neural models have provided a powerful tool for learning how to compose words together into a meaning representation of whole sentences for many downstream tasks. This is done using models of various levels of sophistication, from simpler bag-of-words (Iyyer et al., 2015) and linear recurrent neural network (RNN) models (Sutskever et al., 2014; Kiros et al., 2015), to more sophisticated models using tree-

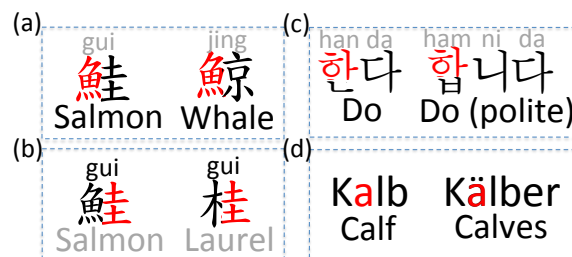


Figure 1: Examples of character-level compositionality in (a, b) Chinese, (c) Korean, and (d) German. The red part of the characters are shared, and affects the pronunciation (top) or meaning (bottom).

structured (Socher et al., 2013) or convolutional networks (Kalchbrenner et al., 2014).

In fact, a growing body of evidence shows that it is essential to look below the word-level and consider compositionality within words themselves. For example, several works have proposed models that represent words by composing together the characters into a representation of the word itself (Ling et al., 2015; Zhang et al., 2015; Dhingra et al., 2016). Additionally, for languages with productive word formation (such as agglutination and compounding), models calculating morphology-sensitive word representations have been found effective (Luong et al., 2013; Botha and Blunsom, 2014). These models help to learn more robust representations for *rare words* by exploiting morphological patterns, as opposed to models that operate purely on the lexical level as the atomic units.

For many languages, compositionality stops at the character-level: characters are atomic units of meaning or pronunciation in the language, and no further decomposition can be done.¹ However, for other languages, character-level compositionality, where a character’s meaning or pronunciation can

¹In English, for example, this is largely the case.

Lang	Geography	Sports	Arts	Military	Economics	Transportation
Chinese	32.4k	49.8k	50.4k	3.6k	82.5k	40.4k
Japanese	18.6k	82.7k	84.1k	81.6k	80.9k	91.8k
Korean	6k	580	5.74k	840	5.78k	1.68k
Lang	Medical	Education	Food	Religion	Agriculture	Electronics
Chinese	30.3k	66.2k	554	66.9k	89.5k	80.5k
Japanese	66.5k	86.7k	20.2k	98.1k	97.4k	1.08k
Korean	16.1k	4.71k	33	2.60k	1.51k	1.03k

Table 1: By-category statistics for the Wikipedia dataset. Note that Food is the abbreviation for “Food and Culture” and Religion is the abbreviation for “Religion and Belief”.

be derived from the sum of its parts, is very much a reality. Perhaps the most compelling example of compositionality of sub-character units can be found in logographic writing systems such as the Han and Kanji characters used in Chinese and Japanese, respectively.² As shown on the left side of Fig. 1, each part of a Chinese character (called a “radical”) potentially contributes to the meaning (i.e., Fig. 1(a)) or pronunciation (i.e., Fig. 1(b)) of the overall character. This is similar to how English characters combine into the meaning or pronunciation of an English word. Even in languages with phonemic orthographies, where each character corresponds to a pronunciation instead of a meaning, there are cases where composition occurs. Fig. 1(c) and (d) show the examples of Korean and German, respectively, where morphological inflection can cause single characters to make changes where some but not all of the component parts are shared.

In this paper, we investigate the feasibility of modeling the *compositionality of characters* in a way similar to how humans do: by visually observing the character and using the features of its shape to learn a representation encoding its meaning. Our method is relatively simple, and generalizable to a wide variety of languages: we first transform each character from its Unicode representation to a rendering of its shape as an image, then calculate a representation of the image using Convolutional Neural Networks (CNNs) (Cun et al., 1990). These features then serve as inputs to a down-stream processing task and trained in an end-to-end manner, which first calculates a loss function, then back-propagates the loss back to the CNN.

²Other prominent examples are largely for extinct languages: Egyptian hieroglyphics, Mayan glyphs, and Sumerian cuneiform scripts (Daniels and Bright, 1996).

As demonstrated by our motivating examples in Fig. 1, in logographic languages character-level semantic or phonetic similarity is often indicated by visual cues; we conjecture that CNNs can appropriately model these visual patterns. Consequently, characters with similar visual appearances will be biased to have similar embeddings, allowing our model to handle *rare characters* effectively, just as character-level models have been effective for rare words.

To evaluate our model’s ability to learn representations, particularly for rare characters, we perform experiments on a downstream task of classifying Wikipedia titles for three Asian languages: Chinese, Japanese, and Korean. We show that our proposed framework outperforms a baseline model that uses standard character embeddings for instances containing rare characters. A qualitative analysis of the characteristics of the learned embeddings of our model demonstrates that visually similar characters share similar embeddings. We also show that the learned representations are particularly effective under low-resource scenarios and *complementary* with standard character embeddings; combining the two representations through three different fusion methods (Snoek et al., 2005; Karpathy et al., 2014) leads to consistent improvements over the strongest baseline without visual features.

2 Dataset

Before delving into the details of our model, we first describe a dataset we constructed to examine the ability of our model to capture the compositional characteristics of characters. Specifically, the dataset must satisfy two desiderata: (1) it must be necessary to fully utilize each character in the input in order to achieve high accuracy, and (2) there must be enough regularity and com-

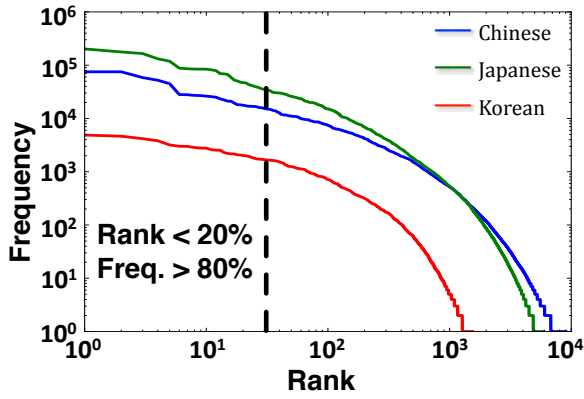


Figure 2: The character rank-frequency distribution of the corpora we considered in this paper. All three languages have a long-tail distribution.

positionality in the characters of the language. To satisfy these desiderata, we create a text classification dataset where the input is a Wikipedia article title in Chinese, Japanese, or Korean, and the output is the category to which the article belongs.³ This satisfies (1), because Wikipedia titles are short and thus each character in the title will be important to our decision about its category. It also satisfies (2), because Chinese, Japanese, and Korean have writing systems with large numbers of characters that decompose regularly as shown in Fig. 1. While this task in itself is novel, it is similar to previous work in named entity type inference using Wikipedia (Toral and Munoz, 2006; Kazama and Torisawa, 2007; Ratinov and Roth, 2009), which has proven useful for downstream named entity recognition systems.

2.1 Dataset Collection

As the labels we would like to predict, we use 12 different main categories from the Wikipedia web page: Geography, Sports, Arts, Military, Economics, Transportation, Health Science, Education, Food Culture, Religion and Belief, Agriculture and Electronics. Wikipedia has a hierarchical structure, where each of these main categories has a number of subcategories, and each subcategory has its own subcategories, etc. We traverse this hierarchical structure, adding each main category tag to all of its descendants in this subcategory tree structure. In the case that a particular article is the descendant of multiple main categories, we favor the main category that minimizes the depth of the

³The link to the dataset and the crawling scripts – https://github.com/frederick0329/Wikipedia_title_dataset

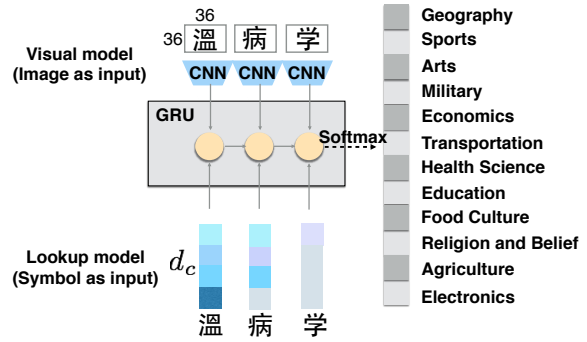


Figure 3: An illustration of two models, our proposed VISUAL model at the top and the baseline LOOKUP model at the bottom using the same RNN architecture. A string of characters (e.g. “温病学”), each converted into a 36x36 image, serves as input of our VISUAL model. d_c is the dimension of the character embedding for the LOOKUP model.

article in the tree (e.g., if an article is two steps away from Sports and three steps away from Arts, it will receive the “Sports” label). We also perform some rudimentary filtering, removing pages that match the regular expression “.*:.*”, which catches special pages such as “title:agriculture”.

2.2 Statistics

For Chinese, Japanese, and Korean, respectively, the number of articles is 593k/810k/46.6k, and the average length and standard deviation of the title is $6.25 \pm 3.96/8.60 \pm 5.58/6.10 \pm 3.71$. As shown in Fig. 2, the character rank-frequency distributions of all three languages follows the 80/20 rule (Newman, 2005) (i.e., top 20% ranked characters that appear more than 80% of total frequencies), demonstrating that the characters in these languages belong to a long tail distribution.

We further split the dataset into training, validation, and testing sets with a 6:2:2 ratio. The category distribution for each language can be seen in Tab. 1. Chinese has two varieties of characters, traditional and simplified, and the dataset is a mix of the two. Hence, we transform this dataset into two separate sets, one completely simplified and the other completely traditional using the Chinese text converter provided with Mac OS.

3 Model

Our overall model for the classification task follows the encoder model by Sutskever et al. (2014).

Layer#	3-layer CNN Configuration
1	Spatial Convolution (3, 3) \rightarrow 32
2	ReLU
3	MaxPool (2, 2)
4	Spatial Convolution (3, 3) \rightarrow 32
5	ReLU
6	MaxPool (2, 2)
7	Spatial Convolution (3, 3) \rightarrow 32
8	ReLU
9	Linear (800, 128)
10	ReLU
11	Linear (128, 128)
12	ReLU

Table 2: Architecture of the CNN used in the experiments. All the convolutional layers have 3×3 filters.

We calculate character representations, use a RNN to combine the character representations into a sentence representation, and then add a softmax layer after that to predict the probability for each class. As shown in Fig. 2.1, the baseline model, which we call it the LOOKUP model, calculates the representation for each character by looking it up in a character embedding matrix. Our proposed model, the VISUAL model instead learns the representation of each character from its visual appearance via CNN.

LOOKUP model Given a character vocabulary C , for the LOOKUP model as in the bottom part of Fig. 2.1, the input to the network is a stream of characters c_1, c_2, \dots, c_N , where $c_n \in C$. Each character is represented by a 1-of- $|C|$ (one-hot) encoding. This one-hot vector is then multiplied by the lookup matrix $T_C \in \mathbf{R}^{|C| \times d_c}$, where d_c is the dimension of the character embedding. The randomly initialized character embeddings were optimized with classification loss.

VISUAL model The proposed method aims to learn a representation that includes image information, allowing for better parameter sharing among characters, particularly characters that are less common. Different from the LOOKUP model, each character is first transformed into a 36-by-36 image based on its Unicode encoding as shown in the upper part of Fig. 2.1. We then pass the image through a CNN to get the embedding for the image. The parameters for the CNN are learned through backpropagation from the classification

loss. Because we are training embeddings based on this classification loss, we expect that the CNN will focus on parts of the image that contain semantic information useful for category classification, a hypothesis that we examine in the experiments (see Section 5.5).

In more detail, the specific structure of the CNN that we utilize consists of three convolution layers where each convolution layer is followed by the max pooling and ReLU nonlinear activation layers. The configurations of each layer are listed in Tab. 2. The output vector for the image embeddings also has size d_c which is the same as the LOOKUP model.

Encoder and Classifier For both the LOOKUP and the VISUAL models, we adopt an RNN encoder using Gated Recurrent Units (GRUs) (Chung et al., 2014). Each of the GRU units processes the character embeddings sequentially. At the end of the sequence, the incremental GRU computation results in a hidden state e embedding the sentence. The encoded sentence embedding is passed through a linear layer whose output is the same size as the number of classes. We use a softmax layer to compute the posterior class probabilities:

$$P(y = j|e) = \frac{\exp(w_j^T e + b_j)}{\sum_{i=1}^L \exp(w_i^T e + b_i)} \quad (1)$$

To train the model, we use cross-entropy loss between predicted and true targets:

$$J = \frac{1}{B} \sum_{i=1}^B \sum_{j=1}^L -t_{i,j} \log(p_{i,j}) \quad (2)$$

where $t_{i,j} \in \{0, 1\}$ represents the ground truth label of the j -th class in the i -th Wikipedia page title. B is the batch size and L is the number of categories.

4 Fusion-based Models

One thing to note is that the LOOKUP and the VISUAL models have their own advantages. The LOOKUP model learns embedding that captures the semantics of each character symbol without sharing information with each other. In contrast, the proposed VISUAL model directly learns embedding from visual information, which naturally shares information between visually similar characters. This characteristic gives the VISUAL

Lookup/Visual	100%	50%	12.5%
zh.trad	0.55/0.54	0.53/0.50	0.48/0.47
zh.simp	0.55/0.54	0.53/0.52	0.48/0.46
ja	0.42/0.39	0.47/0.45	0.44/0.41
ko	0.47/0.42	0.44/0.39	0.37/0.36

Table 3: The classification results of the LOOKUP / VISUAL models for different percentages of full training size.

model the ability to generalize better to rare characters, but also has the potential disadvantage of introducing noise for characters with similar appearances but different meanings.

With the complementary nature of these two models in mind, we further combine the two embeddings to achieve better performances. We adopt three fusion schemes, early fusion, late fusion (described by Snoek et al. (2005) and Karpathy et al. (2014)), and fallback fusion, a method specific to this paper.

Early Fusion Early fusion works by concatenating the two varieties of embeddings before feeding them into the RNN. In order to ensure that the dimensions of the RNN are the same after concatenation, the concatenated vector is fed through a hidden layer to reduce the size from $2 \times d_c$ to d_c . The whole model is then fine-tuned with training data.

Late Fusion Instead of learning a joint representation like early fusion, late fusion averages the model predictions. Specifically, it takes the output of the softmax layers from both models and averages the probabilities to create a final distribution used to make the prediction.

Fallback Fusion Our final fallback fusion method hypothesizes that our VISUAL model does better with instances which contain more rare characters. First, in order to quantify the overall rareness of an instance consisting of multiple characters, we calculate the average training set frequency of the characters therein. The fallback fusion method uses the VISUAL model to predict testing instances with average character frequency below or equal to a threshold (here we use 0.0 frequency as cutoff, which means all characters in the instance do not appear in the training set), and uses the LOOKUP model to predict the rest of the instances.

5 Experiments and Results

In this section, we compare our proposed VISUAL model with the baseline LOOKUP model through three different sets of experiments. First, we examine whether our model is capable of classifying text and achieving similar performance as the baseline model. Next, we examine the hypothesis that our model will outperform the baseline model when dealing with low frequency characters. Finally, we examine the fusion methods described in Section 4.

5.1 Experimental Configurations

The dimension of the embeddings and batch size for both models are set to $d_c = 128$ and $B = 400$, respectively. We build our proposed model using Torch (Collobert et al., 2002), and use Adam (Kingma and Ba, 2014) with a learning rate $\eta = 0.001$ for stochastic optimization. The length of each instance is cut off or padded to 10 characters for batch training.

5.2 Comparison with the Baseline Model

In this experiment, we examine whether our VISUAL model achieves similar performance with the baseline LOOKUP model in classification accuracy.

The results in Tab. 3 show that the baseline model performs 1-2% better across four datasets; this is due to the fact that the LOOKUP model can directly learn character embeddings that capture the semantics of each character symbol for frequent characters. In contrast, the VISUAL model learns embeddings from visual information, which constraints characters that has similar appearance to have similar embeddings. This is an advantage for rare characters, but a disadvantage for high frequency characters because being similar in appearance does not always lead to similar semantics.

To demonstrate that this is in fact the case, besides looking at the overall classification accuracy, we also examine the performance on classifying low frequency instances which are sorted according to the average training set frequency of the characters therein. Tab. 4 and Fig. 4 both show that our model performs better in the 100 lowest frequency instances (the intersection point of the two models). More specifically, take Fig. 4(a) as example, the solid (proposed) line is higher than the dashed (baseline) line up to 10^2 , indicating that the proposed model outperforms the baseline for the




Figure 4: Experiments on different training sizes for four different datasets. More specifically, we consider three different training data size percentages (TPs) (100%, 50%, and 12.5%) and four datasets: (a) traditional Chinese, (b) simplified Chinese, (c) Japanese, and (d) Korean. We calculate the accumulated number of correctly predicted instances for the VISUAL model (solid lines) and the LOOKUP model (dashed lines). This figure is a log-log plot, where x-axis shows rarity (rarest to the left), y-axis shows cumulative correctly classified instances up to this rank; a perfect classifier will result in a diagonal line.

first 100 instances. Lines depart the x-axis when the model classifies its first instance correctly, and the LOOKUP model did not correctly classify any of the first 80 rarest instances, resulting in it crossing later than the proposed model. This confirms that the VISUAL model can share visual information among characters and help to classify low frequency instances.

For training time, visual features take significantly more time, as expected. VISUAL is 30x slower than LOOKUP, although they are equivalent at test time. For space, images of Chinese characters took 36MB to store for 8985 characters.

5.3 Experiments on Different Training Sizes

In our second experiment, we consider two smaller training sizes (i.e., 50% and 12.5% of the full training size) indicated by green and red lines in Fig. 4. We performed this experiment under the hypothesis that because the proposed method was more robust to infrequent characters, the proposed model may perform better in low-resourced scenarios. If this is the case, the intersection point of the two models will shift right because of the increase of the number of instances with low average character frequency.

Lookup/Visual	100	1000	10000
zh_trad	0.22/ 0.49	0.35/0.35	0.40 /0.39
zh_simp	0.25/ 0.53	0.39 /0.37	0.41 /0.40
ja	0.30/ 0.35	0.45 /0.41	0.44 /0.41
ko	0.44 /0.33	0.44 /0.33	0.48 /0.42

Table 4: Classification results for the LOOKUP / VISUAL of the k lowest frequency instances across four datasets. The 100 lowest frequency instances for traditional and simplified Chinese and Korean were both significant (p -value < 0.05). Those for Japanese were not (p -value = 0.13); likely because there was less variety than Chinese and more data than Korean.

As we can see in Fig. 4, the intersection point for 100% training data lies between the intersection point for 50% training data and 12.5%. This disagrees with our hypothesis; this is likely because while the number of low-frequency characters increases, smaller amounts of data also adversely impact the ability of CNN to learn useful visual features, and thus there is not a clear gain nor loss when using the proposed method.

As a more extreme test of the ability of our proposed framework to deal with the unseen char-

	zh_trad	zh_simp	ja	ko
Lookup	0.5503	0.5543	0.4914	0.4765
Visual	0.5434	0.5403	0.4775	0.4207
early	0.5520	0.5546	0.4896	0.4796
late	0.5658	0.5685	0.5029	0.4869
fall	0.5507	0.5547	0.4914	0.4766

Table 5: Experiment results for three different fusion methods across 4 datasets. The late fusion model was better (p-value < 0.001) across four datasets.

acters in the test set, we use traditional Chinese as our training data and simplified Chinese as our testing data. The model was able to achieve around 40% classification accuracy when we use the full training set, compared to 55%, which is achieved by the model trained on simplified Chinese. This result demonstrates that the model is able to transfer between similar scripts, similarly to how most Chinese speakers can guess the meaning of the text, even if it is written in the other script.

5.4 Experiment on Different Fusion Methods

Results of different fusion methods can be found in Tab. 5. The results show that late fusion gives the best performance among all the fusion schemes combining the LOOKUP model and the proposed VISUAL model. Early fusion achieves small improvements for all languages except Japanese, where it displays a slight drop. Unsurprisingly, fallback fusion performs better than the LOOKUP model and the VISUAL model alone, since it directly targets the weakness of the LOOKUP model (e.g., rare characters) and replaces the results with the VISUAL model. These results show that simple integration, no matter which schemes we use, is beneficial, demonstrating that both methods are capturing complementary information.

5.5 Visualization of Character Embeddings

Finally, we qualitatively examine what is learned by our proposed model in two ways. First, we visualize which parts of the image are most important to the VISUAL model’s embedding calculation. Second, we show the 6-nearest neighbor results for characters using both the LOOKUP and the VISUAL embeddings.



Figure 5: Examples of how much each part of the character contributes to its embedding (the darker the more). Two characters are shown per radical to emphasize that characters with same radical have similar patterns.

Emphasis of the VISUAL Model In order to delve deeper into what the VISUAL model has learned, we measure a modified version of the occlusion sensitivity proposed by Zeiler and Fergus (2014) by masking the original character image in four ways, and examine the importance of each part of the character to the model’s calculated representations. Specifically, we leave only the upper half, bottom half, left half, or right half of the image, and mask the remainder with white pixels since Chinese characters are usually formed by combining two radicals vertically or horizontally. We run these four images forward through the CNN part of the model and calculate the L_2 distance between the masked image embeddings with the full image embedding. The larger the distance, the more the masked part of the character contributes to the original embedding. The contribution of each part (e.g. the L_2 distance) is represented as a heat map, and then it is normalized to adjust the opacity of the character strokes for better visualization. The value of each corner of the heatmap is calculated by adding the two L_2 distances that contribute to this corner.

The visualization is shown in Fig. 5. The meaning of each Chinese character in English is shown below the Chinese character. The opacity of the character strokes represent how much the corresponding parts contribute to the original embedding (the darker the more). In general, the darker part of the character is related to its semantics. For example, “金” means gold in Chinese, which is

of what radicals are included in which characters (Li et al., 2015; Shi et al., 2015; Yin et al., 2016). The motivation of this method is similar to ours, but is only applicable to Chinese, in contrast to the method in this paper, which works on any language for which we can render text.

7 Conclusion and Future Work

In this paper, we proposed a new framework that utilizes appearance of characters, convolutional neural networks, recurrent neural networks to learn embeddings that are compositional in the component parts of the characters. More specifically, we collected a Wikipedia dataset, which consists of short titles of three different languages and satisfies the compositionality in the characters of the language. Next, we proposed an end-to-end model that learns visual embeddings for characters using CNN and showed that the features extracted from the CNN include both visual and semantic information. Furthermore, we showed that our VISUAL model outperforms the LOOKUP baseline model in low frequency instances. Additionally, by examining the character embeddings visually, we found that our VISUAL model is able to learn visually related embeddings.

In summary, we tackled the problem of rare characters by using embeddings learned from images. In the future, we hope to further generalize this method to other tasks such as pronunciation estimation, which can take advantage of the fact that pronunciation information is encoded in parts of the characters as demonstrated in Fig. 1, or machine translation, which could benefit from a wholistic view that considers both semantics and pronunciation. We also hope to apply the model to other languages with complicated compositional writing systems, potentially including historical texts such as hieroglyphics or cuneiform.

Acknowledgments

We thank Taylor Berg-Kirkpatrick, Adhiguna Kuncoro, Chen-Hsuan Lin, Wei-Cheng Chang, Wei-Ning Hsu and the anonymous reviewers for their enlightening comments and feedbacks.

References

- Jan A Botha and Phil Blunsom. 2014. Compositional morphology for word representations and language modelling. In *ICML*. pages 1899–1907.
- Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*.
- Ronan Collobert, Samy Bengio, and Johnny Marthoz. 2002. Torch: A modular machine learning software library.
- Y. Le Cun, B. Boser, J. S. Denker, R. E. Howard, W. Hubbard, L. D. Jackel, and D. Henderson. 1990. Advances in neural information processing systems 2. pages 396–404.
- Peter T Daniels and William Bright. 1996. *The world's writing systems*. Oxford University Press.
- Hal Daumé and Jagadeesh Jagarlamudi. 2011. Domain adaptation for machine translation by mining unseen words. In *ACL-HLT*. pages 407–412.
- Yuntian Deng, Anssi Kanervisto, and Alexander M. Rush. 2016. What you get is what you see: A visual markup decompiler. *arXiv preprint arXiv:1609.04938*.
- Bhuwan Dhingra, Zhong Zhou, Dylan Fitzpatrick, Michael Muehl, and William W Cohen. 2016. Tweet2vec: Character-based distributed representations for social media. *ACL*.
- Gottlob Frege and John Langshaw Austin. 1980. *The foundations of arithmetic: A logico-mathematical enquiry into the concept of number*. Northwestern University Press.
- Dan Gillick, Cliff Brunk, Oriol Vinyals, and Amarnag Subramanya. 2015. Multilingual language processing from bytes. *arXiv preprint arXiv:1512.00103*.
- Nizar Habash. 2008. Four techniques for online handling of out-of-vocabulary words in Arabic-English statistical machine translation. In *HLT-Short*. pages 57–60.
- Mohit Iyyer, Varun Manjunatha, and Jordan L Boyd-Graber. 2015. Deep unordered composition rivals syntactic methods for text classification. In *ACL*.
- Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. 2014. A convolutional neural network for modelling sentences. *ACL* pages 655–665.
- Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. 2014. Large-scale video classification with convolutional neural networks. In *CVPR*. pages 1725–1732.
- Jun'ichi Kazama and Kentaro Torisawa. 2007. Exploiting Wikipedia as external knowledge for named entity recognition. In *EMNLP-CoNLL*. pages 698–707.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *EMNLP*. pages 1746–1751.

- Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Ryan Kiros, Yukun Zhu, Ruslan R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Skip-thought vectors. In *NIPS*. pages 3294–3302.
- Yanran Li, Wenjie Li, Fei Sun, and Sujian Li. 2015. Component-enhanced chinese character embeddings. *EMNLP* pages 829–834.
- Wang Ling, Chris Dyer, Alan W Black, Isabel Trancoso, Ramon Fernandez, Silvio Amir, Luis Marujo, and Tiago Luis. 2015. Finding function in form: Compositional character models for open vocabulary word representation. In *EMNLP*. pages 1520–1530.
- Minh-Thang Luong and Christopher D Manning. 2016. Achieving open vocabulary neural machine translation with hybrid word-character models. *ACL* pages 1054–1063.
- Thang Luong, Richard Socher, and Christopher Manning. 2013. Better word representations with recursive neural networks for morphology. In *CoNLL*. pages 104–113.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *NIPS*. pages 3111–3119.
- Mej Newman. 2005. Power laws, Pareto distributions and Zipf’s law. *CONTEMP PHYS* pages 323–351.
- Lev Ratinov and Dan Roth. 2009. Design challenges and misconceptions in named entity recognition. In *CoNLL*. pages 147–155.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2015. Neural machine translation of rare words with subword units. *ACL* pages 1715–1725.
- Xinlei Shi, Junjie Zhai, Xudong Yang, Zehua Xie, and Chao Liu. 2015. Radical embedding: Delving deeper to chinese radicals. In *ACL*. pages 594–598.
- Cees GM Snoek, Marcel Worring, and Arnold WM Smeulders. 2005. Early versus late fusion in semantic video analysis. In *ACM MM*. pages 399–402.
- Richard Socher, Alex Perelygin, Jean Y Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *EMNLP*. pages 1631–1642.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *NIPS*. pages 3104–3112.
- Zoltán Gendler Szabó. 2010. Compositionality. *Stanford encyclopedia of philosophy*.
- Antonio Toral and Rafael Munoz. 2006. A proposal to automatically build and maintain gazetteers for named entity recognition by using wikipedia. In *EACL*. pages 56–61.
- Jiang Wang, Yi Yang, Junhua Mao, Zhiheng Huang, Chang Huang, and Wei Xu. 2016. Cnn-rnn: A unified framework for multi-label image classification. In *CVPR*. pages 2285–2294.
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron C Courville, Ruslan Salakhutdinov, Richard S Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*.
- Rongchao Yin, Quan Wang, Rui Li, Peng Li, and Bin Wang. 2016. Multi-granularity chinese word embedding. *EMNLP* pages 981–986.
- Matthew D Zeiler and Rob Fergus. 2014. Visualizing and understanding convolutional networks. In *ECCV*. Springer, pages 818–833.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *NIPS*. pages 649–657.
- Yuke Zhu, Oliver Groth, Michael Bernstein, and Li Fei-Fei. 2016. Visual7w: Grounded question answering in images. In *CVPR*. pages 4995–5004.