

# GENERATING TEXT THROUGH ADVERSARIAL TRAINING USING SKIP-THOUGHT VECTORS

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

In the past few years, various advancements have been made in generative models owing to the formulation of Generative Adversarial Networks (GANs). GANs have been shown to perform exceedingly well on a wide variety of tasks pertaining to image generation and style transfer. In the field of Natural Language Processing, word embeddings such as word2vec and GloVe are state-of-the-art methods for applying neural network models on textual data. Attempts have been made for utilizing GANs with word embeddings for text generation. This work presents an approach to text generation using Skip-Thought sentence embeddings in conjunction with GANs based on gradient penalty functions and f-measures. The results of using sentence embeddings with GANs for generating text conditioned on input information are comparable to the approaches where word embeddings are used.

## 1 INTRODUCTION

Numerous efforts have been made in the field of natural language text generation for tasks such as sentiment analysis (Zhang et al. (2018)) and machine translation (Gangi & Federico; Qian et al. (2018)). Early techniques for generating text conditioned on some input information were template or rule-based engines, or probabilistic models such as n-gram. In recent times, state-of-the-art results on these tasks have been achieved by recurrent (Press et al.; Mikolov et al. (2010)) and convolutional neural network models trained for likelihood maximization. This work proposes an approach for text generation using Generative Adversarial Networks with Skip-Thought vectors.

GANs (Goodfellow et al. (2014)) are a class of neural networks that explicitly train a generator to produce high-quality samples by pitting against an adversarial discriminative model. GANs output differentiable values and hence the task of discrete text generation has to use vectors as differentiable inputs. This is achieved by training the GAN with sentence embedding vectors produced by Skip-Thought (Kiros et al. (2015)), a neural network model for learning fixed length representations of sentences.

## 2 RELATED WORKS

Deep neural network architectures have demonstrated strong results on natural language generation tasks (Xie (2017)). Recurrent neural networks using combinations of shared parameter matrices across time-steps (Sutskever et al. (2014); Mikolov et al. (2010); Cho et al. (2014)) with different gating mechanisms for easing optimization (Hochreiter & Schmidhuber (1997); Cho et al. (2014)) have found success in modeling natural language. Another approach is to use convolutional neural networks that reuse kernels across time-steps with attention mechanism to perform language generation tasks (Kalchbrenner et al. (2016; 2014)).

Supervised learning with deep neural networks in the framework of encoder-decoder models has become the state-of-the-art methods for approaching NLP problems (Young et al.). Stacked denoising autoencoder have been used for domain adaptation in classifying sentiments (Glorot et al. (2011)) and combinatory categorical autoencoders demonstrate learning the compositionality of sentences (Hermann & Blunsom (2013)). Recent text generation models use a wide variety of GANs such as gradient policy based sequence generation framework (Yu et al. (2016)) and an actor-critic conditional GAN to fill missing text conditioned on surrounding text (Fedus et al. (2018)) for performing

natural language generation tasks. Other architectures such as those proposed in Wang et al. (2017) with RNN and variational auto-encoder generator with CNN discriminator and in Guo et al. (2017) with leaky discriminator to guide generator through high-level extracted features have also shown great results.

### 3 SKIP-THOUGHT GENERATIVE ADVERSARIAL NETWORK (STGAN)

This section introduces Skip-Thought Generative Adversarial Network with a background on models that it is based on. The Skip-Thought model (Kiros et al. (2015)) induces embedding vectors for sentences present in training corpus. These vectors constitute the real distribution for the discriminator network. The generator network produces sentence vectors similar to those from the encoded real distribution. The generated vectors are sampled over training and decoded to produce sentences using a Skip-Thought decoder conditioned on the same text corpus.

#### 3.1 SKIP-THOUGHT VECTORS

Skip-Thought is an encoder-decoder framework with an unsupervised approach to train a generic, distributed sentence encoder. The encoder maps sentences sharing semantic and syntactic properties to similar vector representations and the decoder reconstructs the surrounding sentences of an encoded passage. The sentence encoding approach draws inspiration from the skip-gram model in producing vector representations using previous and next sentences.

The Skip-Thought model uses an RNN encoder with GRU activations (Chung et al. (2014)) and an RNN decoder with conditional GRU, the combination being identical to the RNN encoder-decoder of Cho et al. (2014) used in neural machine translation.

##### 3.1.1 SKIP-THOUGHT ARCHITECTURE

For a given sentence tuple  $(s_{i-1}, s_i, s_{i+1})$ , let  $w_i^t$  denote the  $t$ -th word for sentence  $s_i$ , and let  $\mathbf{x}_i^t$  denote its word embedding. The model has three components:

**Encoder.** Encoded vectors for a sentence  $s_i$  with  $N$  words  $w^1, w^{i+1}, \dots, w^n$  are computed by iterating over the following sequence of equations:

$$\begin{aligned} \mathbf{r}^t &= \sigma(\mathbf{W}_r \mathbf{x}^t + \mathbf{U}_r \mathbf{h}^{t-1}) \\ \mathbf{z}^t &= \sigma(\mathbf{W}_z \mathbf{x}^t + \mathbf{U}_z \mathbf{h}^{t-1}) \\ \tilde{\mathbf{h}}^t &= \tanh(\mathbf{W} \mathbf{x}^t + \mathbf{U}(\mathbf{r}^t \odot \mathbf{h}^{t-1})) \\ \mathbf{h}^t &= (\mathbf{1} - \mathbf{z}^t) \odot \mathbf{h}^{t-1} + \mathbf{z}^t \odot \tilde{\mathbf{h}}^t \end{aligned}$$

where  $\mathbf{h}_i^t$  is a hidden state at each time step and interpreted as a sequence of words  $w_i^1, \dots, w_i^n$ ,  $\mathbf{h}^t$  is the proposed state update at time  $t$ ,  $\mathbf{z}^t$  is the update gate and  $\mathbf{r}^t$  is the reset gate. Both update gates take values between zero and one.

**Decoder.** A neural language model conditioned on the encoder output  $\mathbf{h}_i$  serves as the decoder. Bias matrices  $C_z, C_r, C$  are introduced for the update gate, reset gate and hidden state computation by the encoder. Two decoders are used in parallel, one each for sentences  $s_i + 1$  and  $s_i - 1$ . The following equations are iterated over for decoding:

$$\begin{aligned} \mathbf{r}^t &= \sigma(\mathbf{W}_r^d \mathbf{x}^{t-1} + \mathbf{U}_r^d \mathbf{h}^{t-1} + C_r \mathbf{h}_i) \\ \mathbf{z}^t &= \sigma(\mathbf{W}_z^d \mathbf{x}^{t-1} + \mathbf{U}_z^d \mathbf{h}^{t-1} + C_z \mathbf{h}_i) \\ \tilde{\mathbf{h}}^t &= \tanh(\mathbf{W}^d \mathbf{x}^{t-1} + \mathbf{U}^d(\mathbf{r}^t \odot \mathbf{h}^{t-1}) + C \mathbf{h}_i) \\ \mathbf{h}_{i+1}^t &= (\mathbf{1} - \mathbf{z}^t) \odot \mathbf{h}^{t-1} + \mathbf{z}^t \odot \tilde{\mathbf{h}}^t \end{aligned}$$

**Objective.** For the same tuple of sentences, objective function is the sum of log-probabilities for the forward and backward sentences conditioned on the encoder representation:

$$\sum_t \log P(w_{i+1}^t | w_{i+1}^{<t}, \mathbf{h}_i) + \sum_t \log P(w_{i-1}^t | w_{i-1}^{<t}, \mathbf{h}_i)$$

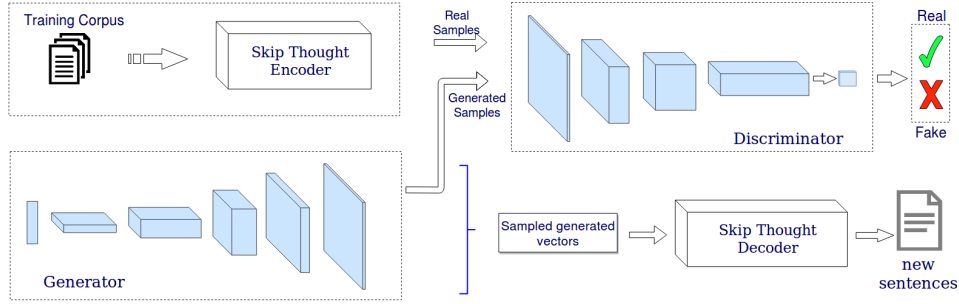


Figure 1: Model architecture: Skip-Thought Generative Adversarial Network

### 3.2 GENERATIVE ADVERSARIAL NETWORKS

Generative Adversarial Networks (Goodfellow et al. (2014)) are deep neural net architectures comprised of two networks, contesting with each other in a zero-sum game framework. For a given data, GANs can mimic learning the underlying distribution and generate artificial data samples similar to those from the real distribution. Generative Adversarial Networks consists of two players - a Generator and a Discriminator. The generator  $G$  tries to produce data close to the real distribution  $P(x)$  from some stochastic distribution  $P(z)$  termed as noise. The discriminator  $D$ 's objective is to differentiate between real and generated data  $G(z)$ .

The two networks - generator and discriminator compete against each other in a zero-sum game. The minimax strategy dictates that each network plays optimally with the assumption that the other network is optimal. This leads to Nash equilibrium which is the point of convergence for GAN model.

**Objective.** Goodfellow et al. (2014) have formulated the minimax game for a generator  $G$ , discriminator  $D$  adversarial network with value function  $V(G, D)$  as:

$$\min_G \max_D V(D, G) = \mathbb{E}_{x \sim p_{data}(x)} [\log D(x)] + \mathbb{E}_{z \sim p_z(z)} [\log (1 - D(G(z)))]$$

### 3.3 MODEL ARCHITECTURE

The STGAN model uses a deep convolutional generative adversarial network, similar to the one used in Radford et al.. The generator network is updated twice for each discriminator network update to prevent fast convergence of the discriminator network.

The Skip-Thought encoder for the model encodes sentences with length less than 30 words using 2400 GRU units (Chung et al. (2014)) with word vector dimensionality of 620 to produce 4800-dimensional combine-skip vectors. The combine-skip vectors, with the first 2400 dimensions being uni-skip model and the last 2400 bi-skip model, are used as they have been found to be the best performing in the experiments<sup>1</sup>. The decoder uses greedy decoding taking argmax over softmax output distribution for given time-step which acts as input for next time-step. It reconstructs sentences conditioned on a sentence vector by randomly sampling from the predicted distributions with or without a preset beam width. Unknown tokens are not included in the vocabulary. A 620 dimensional RNN word embeddings is used with 1600 hidden GRU decoding units. Gradient clipping with Adam optimizer (Kingma & Ba (2014)) is used, with a batch size of 16 and maximum sentence length of 100 words for decoder.

### 3.4 IMPROVING TRAINING AND LOSS

The training process of a GAN is notably difficult as demonstrated by Salimans et al. (2016) and several improvement techniques such as batch normalization, feature matching, historical averaging (Salimans et al. (2016)) and unrolling GAN (Metz et al.) have been suggested for making the training more stable. Training the Skip-Thought GAN often results in mode dropping (Arjovsky &

<sup>1</sup><https://github.com/ryankiros/skip-thoughts/>

Table 1: BLEU-2, BLEU-3 and BLEU-4 metric scores for different models with (a) test set as reference, and (b) entire corpus as reference. **ST**: Skip-Thought, **GAN**: Generative Adversarial Network, **W**: Wasserstein **GP**: Gradient Penalty

MODEL	TEST SET	COMPLETE CORPUS REFERENCE		
	BLEU-3	BLEU-2	BLEU-3	BLEU-4
<b>STGAN</b>	0.521	<b>0.709</b>	0.564	0.525
<b>STGAN</b> (minibatch)	0.526	<b>0.745</b>	0.607	0.531
<b>STGAN-GP</b>	0.558	<b>0.791</b>	0.621	0.547
<b>STWGAN</b>	0.582	<b>0.833</b>	0.669	0.580
<b>STWGAN-GP</b>	<b>0.617</b>	<b>0.836</b>	<b>0.682</b>	<b>0.594</b>

Bottou; Srivastava et al.) with a parameter setting where it outputs a very narrow distribution of points. To overcome this, it uses minibatch discrimination by looking at an entire batch of samples and modeling the distance between a given sample and all the other samples present in that batch.

The minimax formulation for an optimal discriminator in a vanilla GAN is Jensen-Shannon Distance between the generated distribution and the real distribution. Arjovsky et al. (2017) used Wasserstein distance or earth mover’s distance to demonstrate how replacing distance measures can improve training loss for GAN. Gulrajani et al. (2017) have incorporated a gradient penalty regularizer in WGAN objective for discriminator’s loss function. The experiments in this work use the above f-measures to improve performance of Skip-Thought GAN on text generation.

## 4 RESULTS AND DISCUSSION

### 4.1 CONDITIONAL GENERATION OF SENTENCES.

GANs can be conditioned on data attributes to generate samples (Mirza & Osindero (2014); Radford et al.). In this experiment, both the generator and discriminator are conditioned on Skip-Thought encoded vectors (Kiros et al. (2015)). The encoder converts 70000 sentences from the BookCorpus dataset collected in Zhu et al. (2015) with a training/test/validation split of 5/1/1 into vectors used as real samples for discriminator. The decoded sentences are used to evaluate model performance under corpus level BLEU-2, BLEU-3 and BLEU-4 metrics (Papineni et al.), once using only test set as reference and then entire corpus as reference. Table 1 compares these results for different architectures that have been experimented with in this paper.

### 4.2 LANGUAGE GENERATION.

Language generation is done on a dataset comprising simple English sentences referred to as CMU-SE<sup>2</sup> in Rajeswar et al. (2017). The CMU-SE dataset consists of 44,016 sentences with a vocabulary of 3,122 words. For encoding, the vectors are extracted in batches of sentences having the same length. The samples represent how mode collapse is manifested when using least-squares distance (Mao et al. (2016)) f-measure without minibatch discrimination. Table 2(a) contains sentences generated from STGAN using least-squares distance (Mao et al. (2016)) in which there was no mode collapse observed, while 2(b) contains examples wherein it is observed. Table 2(c) shows generated sentences using gradient penalty regularizer(GAN-GP). Table 2(d) has samples generated from STGAN when using Wasserstein distance f-measure as WGAN in Arjovsky et al. (2017)) and 2(e) contains samples when using a gradient penalty regularizer term as WGAN-GP from (Gulrajani et al. (2017)).

<sup>2</sup><https://github.com/clab/sp2016.11-731/tree/master/hw4/data>

Table 2: Sample sentences generated from training on CMU-SE Dataset; mode collapse is overcome by using minibatch discrimination. Formation of sentences further improved by changing f-measure to Wasserstein distance along with gradient penalty regularizer.

<b>Mode collapse</b>	<ol style="list-style-type: none"> <li>1. it ?</li> <li>2. it ?</li> <li>3. it ?</li> <li>4. it ? how would it ?</li> <li>5. it ? how would it ?</li> </ol>
<b>With minibatch discrimination</b>	<ol style="list-style-type: none"> <li>1. it a bottle ?</li> <li>2. a glass bottle ?</li> <li>3. a glass bottle it ?</li> <li>4. it my hand a bottle ?</li> <li>5. the phone my hand it</li> </ol>
<b>With gradient penalty</b>	<ol style="list-style-type: none"> <li>1. battery is eighteen percent um ?</li> <li>2. what fine are cash please ?</li> <li>3. youre gonna go over the t- house .</li> <li>4. do you have a nice store around here?</li> <li>5. open this flight number six zero one.</li> </ol>
<b>Wasserstein STGAN</b>	<ol style="list-style-type: none"> <li>1. we have new year s holidays, always.</li> <li>2. here you can nt see your suitcase ,</li> <li>3. please show me how much is a transfer?</li> <li>4. i had a police take watch out of my wallet .</li> <li>5. here i collect my telephone card and telephone number</li> </ol>
<b>Wasserstein STGAN with gradient penalty</b>	<ol style="list-style-type: none"> <li>1. my passport and a letter card with my card , please</li> <li>2. here on my telephone, mr. kimuras registration cards address.</li> <li>3. i can nt see some shopping happened .</li> <li>4. get him my camera found a person s my watch .</li> <li>5. delta airlines flight six zero two from six p.m. to miami, please?</li> </ol>

### 4.3 FURTHER WORK

Another performance metric that can be computed for this setup has been described in Rajeswar et al. (2017) which is a parallel work to this. Simple CFG<sup>3</sup> and more complex ones like Penn Treebank CFG generate samples (Eisner & Smith (2008)) which are used as input to GAN and the model is evaluated by computing the diversity and accuracy of generated samples conforming to the given CFG.

Skip-Thought sentence embeddings can be used to generate images with GANs conditioned on text vectors for text-to-image conversion tasks like those achieved in Reed et al.; Bodnar (2018). These embeddings have also been used to Models like neural-storyteller<sup>4</sup> which use these sentence embeddings can be experimented with generative adversarial networks to generate unique samples.

### ACKNOWLEDGMENTS

Daniel Ricks for Penseur: An interface for the Sent2Vec encoder and training code from the "Skip-Thought Vectors" paper.

Tae-hoon Kim for DCGAN-tensorflow: A tensorflow implementation of "Deep Convolutional Generative Adversarial Networks".

Dept of CSIS, BITS Hyderabad for providing Nvidia Titan XP GPUs for experimentation with code.

<sup>3</sup><http://www.cs.jhu.edu/~jason/465/hw-grammar/extra-grammars/holygrail>

<sup>4</sup><https://github.com/ryankiros/neural-storyteller>

## REFERENCES

- M. Arjovsky and L. Bottou. Towards Principled Methods for Training Generative Adversarial Networks. *ArXiv e-prints*.
- M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein GAN. *ArXiv e-prints*, January 2017.
- Cristian Bodnar. Text to image synthesis using generative adversarial networks. *CoRR*, abs/1805.00676, 2018.
- Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the properties of neural machine translation: Encoder–decoder approaches. In *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*. Association for Computational Linguistics, 2014.
- Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.
- Jason Eisner and Noah A Smith. Competitive grammar writing. In *Proceedings of the Third Workshop on Issues in Teaching Computational Linguistics*, pp. 97–105. Association for Computational Linguistics, 2008.
- W. Fedus, I. Goodfellow, and A. M. Dai. MaskGAN: Better Text Generation via Filling in the----- *ArXiv e-prints*, January 2018.
- Mattia Antonino Di Gangi and Marcello Federico. Deep neural machine translation with weakly-recurrent units. *CoRR*, abs/1805.04185.
- Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Domain adaptation for large-scale sentiment classification: A deep learning approach. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pp. 513–520, 2011.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems 27*, pp. 2672–2680. Curran Associates, Inc., 2014.
- Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *Advances in Neural Information Processing Systems 30*, pp. 5767–5777. 2017.
- Jiaxian Guo, Sidi Lu, Han Cai, Weinan Zhang, Yong Yu, and Jun Wang. Long text generation via adversarial training with leaked information. *CoRR*, abs/1709.08624, 2017.
- Karl Moritz Hermann and Phil Blunsom. The role of syntax in vector space models of compositional semantics. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pp. 894–904, 2013.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8), November 1997.
- Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. A convolutional neural network for modelling sentences. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2014.
- Nal Kalchbrenner, Lasse Espeholt, Karen Simonyan, Aaron van den Oord, Alex Graves, and Koray Kavukcuoglu. Neural machine translation in linear time. *CoRR*, abs/1610.10099, 2016.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014.
- Ryan Kiros, Yukun Zhu, Ruslan Salakhutdinov, Richard S Zemel, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. Skip-thought vectors. *arXiv preprint arXiv:1506.06726*, 2015.
- Xudong Mao, Qing Li, Haoran Xie, Raymond Y. K. Lau, and Zhen Wang. Multi-class generative adversarial networks with the L2 loss function. *CoRR*, 2016.

- Luke Metz, Ben Poole, David Pfau, and Jascha Sohl-Dickstein. Unrolled generative adversarial networks. *CoRR*, abs/1611.02163.
- Tomas Mikolov, Martin Karafiát, Lukás Burget, Jan Cernocký, and Sanjeev Khudanpur. Recurrent neural network based language model. In *INTERSPEECH*, 2010.
- Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *CoRR*, abs/1411.1784, 2014.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: A method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, ACL '02.
- Ofir Press, Amir Bar, Ben Bogin, Jonathan Berant, and Lior Wolf. Language generation with recurrent generative adversarial networks without pre-training. *CoRR*, abs/1706.01399.
- Xin Qian, Ziyi Zhong, and Jieli Zhou. Multimodal machine translation with reinforcement learning. *CoRR*, abs/1805.02356, 2018.
- Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *CoRR*, abs/1511.06434.
- Sai Rajeswar, Sandeep Subramanian, Francis Dutil, Christopher Joseph Pal, and Aaron C. Courville. Adversarial generation of natural language. *CoRR*, abs/1705.10929, 2017.
- Scott E. Reed, Zeynep Akata, Xincheng Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text to image synthesis. *CoRR*, abs/1605.05396.
- Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, NIPS'16, 2016.
- A. Srivastava, L. Valkov, C. Russell, M. U. Gutmann, and C. Sutton. VEEGAN: Reducing Mode Collapse in GANs using Implicit Variational Learning. *ArXiv e-prints*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. Sequence to sequence learning with neural networks. *CoRR*, abs/1409.3215, 2014.
- Heng Wang, Zengchang Qin, and Tao Wan. Text generation based on generative adversarial nets with latent variable. *CoRR*, abs/1712.00170, 2017.
- Ziang Xie. Neural text generation: A practical guide. *arXiv preprint arXiv:1711.09534*, 2017.
- Tom Young, Devamanyu Hazarika, Soujanya Poria, and Erik Cambria. Recent trends in deep learning based natural language processing. *CoRR*, abs/1708.02709.
- Lantao Yu, Weinan Zhang, Jun Wang, and Yong Yu. Seqgan: Sequence generative adversarial nets with policy gradient. *CoRR*, abs/1609.05473, 2016.
- Lei Zhang, Shuai Wang, and Bing Liu. Deep learning for sentiment analysis : A survey. *CoRR*, abs/1801.07883, 2018.
- Yukun Zhu, Ryan Kiros, Richard S. Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. 2015.