

Semi-supervised Text Regression with Conditional Generative Adversarial Networks

Tao Li

Purdue University
taoli@purdue.edu

Xudong Liu

ObEN, Inc.
xudong@oben.com

Shihan Su

California Institute of Technology
ssu@caltech.edu

Abstract—Enormous online textual information provides intriguing opportunities for understandings of social and economic semantics. In this paper, we propose a novel text regression model based on a conditional generative adversarial network (GAN), with an attempt to associate textual data and social outcomes in a semi-supervised manner. Besides promising potential of predicting capabilities, our superiorities are twofold: (i) the model works with unbalanced datasets of limited labelled data, which align with real-world scenarios; and (ii) predictions are obtained by an end-to-end framework, without explicitly selecting high-level representations. Finally we point out related datasets for experiments and future research directions.

I. INTRODUCTION

With millions of textual information uploaded every day, the Internet embeds tremendous data of social and economic phenomena, and have attracted consistent interests not only from sociologists and economists but also statisticians and computer scientists. For example, [1] forecasted movie revenues using online reviews; based on social media data, [2] monitored flu pandemic and [3] predicted election results.

To our best knowledge, the concept of text regression was first introduced by [4] who described it as: given a piece of text, predict a real-world continuous quantity associated with the text's meaning. They applied a linear model to estimate financial risks by using financial reports directly and claimed a significant outperformance compared to previous methods. Subsequently, several linear text regression models were proposed; to name a few: [5]–[7].

Although easy for interpretation and implementation, linear models rely heavily on specific selections of high-level textual representations and fail to properly capture complicated distributions. Recent successes of deep neural networks in the field of computer vision (e.g., [8] and [9]) encourage researchers to discover their potential in natural language processing. Unlike image synthesis, using deep networks for natural language generation (NLG) is notoriously difficult [10], as the feature space of a sentence is discrete and thereby discontinuous and non-differentiable. [11] attacked this issue by using one-hot vectors obtained from softmax function for backpropagation. [12] used ranking scores instead of real/fake prediction for the objective function of the discriminator.

Our idea of using GANs for text regression was inspired by recent advances in NLG (e.g., [13] and [14]). We further shift the focus from realistic language synthesis to the generation of adversarial samples from a LSTM [15], who competes against

a discriminator for regression (see Figure 1). The performance of our model is guaranteed by deep neural networks' power of capturing complicated distributions especially when obtained in an adversarial manner. The capability of training with limited supervision also facilitates promising future applications.

The rest of the paper is organized as follows: in Section II we discuss existing text regression techniques and previous works in semi-supervised learning with GANs; the model is detailed in Section III; we conclude the paper in Section IV by future works.

II. RELATED WORK

A. Text Regression

Previous attempts at text regression mainly focused on linear models. [4] adopted a support vector regression (SVR) [16] in financial reports to predict the volatility of stock returns, a widely used measure of financial risk, and reported a significant outperformance compared to state-of-the-arts. To correlate movies' online reviews and corresponding revenues, [1] extracted high-level features of textual reviews and incorporated them into a elastic net model [17]. [3] exploited a multi-task learning scheme that leverages textual data with user profiles for voting intention prediction. As mentioned earlier, linear models sometimes are oversimplified and fail to properly capture real-world scenarios. [18] proposed the first non-linear model, a deep convolutional neural network, for text regression which surpassed previous state of the art even with limited supervision.

B. Semi-supervised Learning

Semi-supervised learning tackles the problem of learning a mapping between data and labels when only a small subset of labels are available. Earlier approaches of generative models with semi-supervised learning consider Gaussian mixture models [19] and non-parametric density models [20], but suffer from limitations of scalability and inference accuracy. Recently [21] addresses this problem by developing stochastic variational inference algorithms for joint optimization of model and variational parameters.

Since generative adversarial networks (GANs) has been shown to be promising in generating realistic images [22], several approaches have been proposed to use GANs in semi-supervised learning. [23] extends the discriminator (D) to be a K class classifier with objective function to minimize

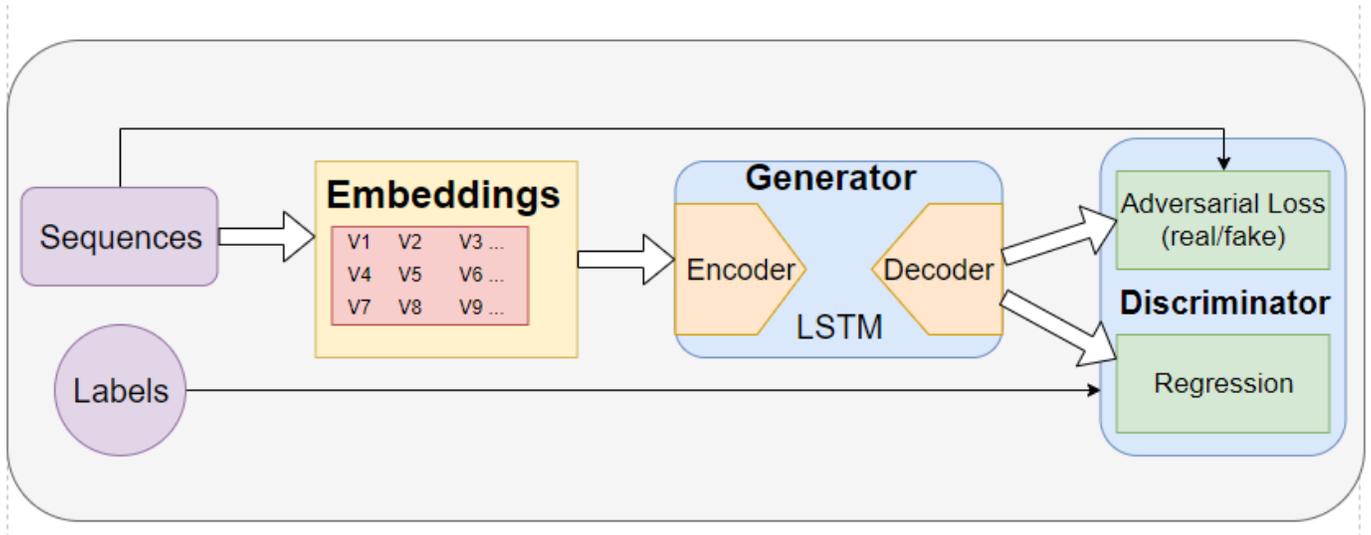


Fig. 1. Architecture of the TR-GAN model.

prediction certainty on generated images, while generator aims for maximize the same objective. [24] augments the K class discriminator to include a $K + 1$ label as fake for the generated images. These work have shown that incorporating adversarial objectives can make the learning of classifier robust and data efficient. While previous works mainly focus on classification setting, in our work, we extend the GAN based semi-supervised learning to regression task.

III. THE TR-GAN MODEL

In this section, we detail the conditional generative adversarial network for text regression in a semi-supervised setting (TR-GAN). We first introduce the word embedding method.

A. Word Embedding

Word embedding method learns a high dimension representation for each word, thereby incorporate semantic information that cannot be captured by the single token. In our work, we adopt a pretrained word embedding model for each word in the text input. Then each document can be represented by a $D \times N$ matrix, where D is the number of words in the document and N is the dimension of the word embedding in the pretrained model.

B. Model Architecture

As illustrated in Figure 1, the network architecture is a conditional GAN with a generator and a discriminator. A long short-term memory network (LSTM) [15] is deployed as the generator for natural languages. As the embedding is fed into LSTM, the generator is a LSTM-based sentence decoder. The discriminator is a convolutional neural network (CNN) [25], where several residual blocks [26] are followed by batch normalization with *ReLU* as the activate function. Subsequently, two fully connected layers are finalized for adversarial learning and the regression task.

The objective function adopt mean absolute error (MAE) for regression tasks and adversarial loss for sequence generation. Not only can this model generate realistic sentences through the optimized generator but the discriminator is also trained as a regression model for multiple prediction tasks (e.g., auto sales prediction, public opinion tracking, and even epidemiological surveillance from social media [27]), which are of great interest to a wide range of stakeholders.

IV. FUTURE WORK

We are excited about the idea of using GANs for text regression. Given the nature of the TR-GAN model, it is not challenging to find an experimental dataset; for example, [28] collected 50,000 textual comments below YouTube videos, among which 20,000 are labelled by state-of-the-art algorithms and 1,000 are labelled manually. We also are interested to see how the generated languages look like, given that existing literatures of using GANs for NLG merely report original experimental results but instead numerical metrics.

ACKNOWLEDGMENTS

We thank Hao Peng and Kantapon Kaewtip for insightful discussions. The idea of this work originally came out during discussions of [29] and [30].

REFERENCES

- [1] M. Joshi, D. Das, K. Gimpel, and N. A. Smith, "Movie reviews and revenues: An experiment in text regression," in *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, 2010, pp. 293–296.
- [2] V. Lamos and N. Cristianini, "Tracking the flu pandemic by monitoring the social web," in *Cognitive Information Processing (CIP), 2010 2nd International Workshop on*. IEEE, 2010, pp. 411–416.
- [3] V. Lamos, D. Preojuic-Pietro, and T. Cohn, "A user-centric model of voting intention from social media," in *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, vol. 1, 2013, pp. 993–1003.

- [4] S. Kogan, D. Levin, B. R. Routledge, J. S. Sagi, and N. A. Smith, "Predicting risk from financial reports with regression," in *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, 2009, pp. 272–280.
- [5] S. Volkova, G. Coppersmith, and B. Van Durme, "Inferring user political preferences from streaming communications," in *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, vol. 1, 2014, pp. 186–196.
- [6] V. Lampos, N. Aletras, D. Preoțiuc-Pietro, and T. Cohn, "Predicting and characterising user impact on twitter," in *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, 2014, pp. 405–413.
- [7] D. Preoțiuc-Pietro, V. Lampos, and N. Aletras, "An analysis of the user occupational class through twitter content," in *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, vol. 1, 2015, pp. 1754–1764.
- [8] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*. Ieee, 2009, pp. 248–255.
- [9] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. P. Aitken, A. Tejani, J. Totz, Z. Wang *et al.*, "Photo-realistic single image super-resolution using a generative adversarial network," in *CVPR*, vol. 2, no. 3, 2017, p. 4.
- [10] T. Li, K. Fu, M. Choi, X. Liu, and Y. Chen, "Toward robust and efficient training of generative adversarial networks with bayesian approximation," in *the Approximation Theory and Machine Learning Conference*, 2018.
- [11] M. J. Kusner and J. M. Hernández-Lobato, "Gans for sequences of discrete elements with the gumbel-softmax distribution," *arXiv preprint arXiv:1611.04051*, 2016.
- [12] K. Lin, D. Li, X. He, Z. Zhang, and M.-T. Sun, "Adversarial ranking for language generation," in *Advances in Neural Information Processing Systems*, 2017, pp. 3155–3165.
- [13] L. Yu, W. Zhang, J. Wang, and Y. Yu, "Seqgan: Sequence generative adversarial nets with policy gradient," in *AAAI*, 2017, pp. 2852–2858.
- [14] J. Li, W. Monroe, T. Shi, S. Jean, A. Ritter, and D. Jurafsky, "Adversarial learning for neural dialogue generation," *arXiv preprint arXiv:1701.06547*, 2017.
- [15] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [24] A. Odena, "Semi-supervised learning with generative adversarial networks," *arXiv preprint arXiv:1606.01583*, 2016.
- [16] H. Drucker, C. J. Burges, L. Kaufman, A. J. Smola, and V. Vapnik, "Support vector regression machines," in *Advances in neural information processing systems*, 1997, pp. 155–161.
- [17] H. Zou and T. Hastie, "Regularization and variable selection via the elastic net," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 67, no. 2, pp. 301–320, 2005.
- [18] Z. Bitvai and T. Cohn, "Non-linear text regression with a deep convolutional neural network," in *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, vol. 2, 2015, pp. 180–185.
- [19] X. Zhu, "Semi-supervised learning literature survey," *Computer Science, University of Wisconsin-Madison*, vol. 2, no. 3, p. 4, 2006.
- [20] C. Kemp, T. L. Griffiths, S. Stromsten, and J. B. Tenenbaum, "Semi-supervised learning with trees," in *Advances in neural information processing systems*, 2004, pp. 257–264.
- [21] D. P. Kingma, S. Mohamed, D. J. Rezende, and M. Welling, "Semi-supervised learning with deep generative models," in *Advances in Neural Information Processing Systems*, 2014, pp. 3581–3589.
- [22] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [23] J. T. Springenberg, "Unsupervised and semi-supervised learning with categorical generative adversarial networks," *arXiv preprint arXiv:1511.06390*, 2015.
- [25] N. Kalchbrenner, E. Grefenstette, and P. Blunsom, "A convolutional neural network for modelling sentences," *arXiv preprint arXiv:1404.2188*, 2014.
- [26] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [27] T. Li, M. Choi, Y. Guo, and L. Lin, "Opinion mining at scale: A case study of the first self-driving car fatality," in *2018 IEEE International Conference on Big Data (Poster)*, 2018.
- [28] T. Li, L. Lin, M. Choi, K. Fu, S. Gong, and J. Wang, "YouTube AV 50K: an annotated corpus for comments in autonomous vehicles," in *the 13th International Joint Symposium on Artificial Intelligence and Natural Language Processing*, 2018.
- [29] T. Li, K. Kaewtip, J. Feng, and L. Lin, "IVAS: Facilitating safe and comfortable driving with intelligent vehicle audio systems," in *2018 IEEE International Conference on Big Data (Poster)*, 2018.
- [30] C. Wang, S. Gong, A. Zhou, T. Li, and S. Peeta, "Cooperative adaptive cruise control for connected autonomous vehicles by factoring communication-related constraints," *arXiv preprint arXiv:1807.07232*, 2018.