

REFERENCES

- [1] Y. Lu, S. Wu, Y.-W. Tai, and C.-K. Tang, “Image generation from sketch constraint using contextual gan,” in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 205–220.
- [2] Y. Li, Z. Gan, Y. Shen, J. Liu, Y. Cheng, Y. Wu, L. Carin, D. Carlson, and J. Gao, “Storygan: A sequential conditional gan for story visualization,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 6329–6338.
- [3] W. Liao, K. Hu, M. Y. Yang, and B. Rosenhahn, “Text to image generation with semantic-spatial aware gan,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 18 187–18 196.
- [4] N. Ruiz, Y. Li, V. Jampani, Y. Pritch, M. Rubinstein, and K. Aberman, “Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation,” *arXiv preprint arXiv:2208.12242*, 2022.
- [5] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen, “Hierarchical text-conditional image generation with clip latents,” *arXiv preprint arXiv:2204.06125*, 2022.
- [6] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, “High-resolution image synthesis with latent diffusion models,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 10 684–10 695.
- [7] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, “Analyzing and improving the image quality of stylegan,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 8110–8119.
- [8] V. Gallego, “Personalizing text-to-image generation via aesthetic gradients,” *arXiv preprint arXiv:2209.12330*, 2022.
- [9] A. Trammell, “Notes from the wargaming underground: Dungeons, dragons, and the history of games,” *Game Studies*, vol. 13, no. 2, 2013.
- [10] D. M. Ewalt, *Of dice and men: The story of Dungeons & Dragons and the people who play it*. Simon and Schuster, 2013.
- [11] J. P. Zagal and S. Deterding, “Definitions of “role-playing games”,” in *Role-Playing Game Studies*. Routledge, 2018, pp. 19–51.

- [12] W. J. White, J. Arjoranta, M. Hitchens, J. Peterson, E. Torner, and J. Walton, “Tabletop role-playing games,” in *Role-Playing Game Studies*. Routledge, 2018, pp. 63–86.
- [13] G. Gygax and D. Arneson, *dungeons & dragons*. Tactical Studies Rules Lake Geneva, WI, 1974, vol. 19.
- [14] O. Svan and A. Wuolo, “Emergent player-driven narrative in blades in the dark and dungeons & dragons: A comparative study,” 2021.
- [15] J. Bulmahn, G. Gygax, and D. Arneson, *Pathfinder roleplaying game: Core rulebook*. Paizo Publishing, LLC, 2013.
- [16] D. B. Staff, “Damp;d beyond,” Jan 2023. [Online]. Available: <https://www.dndbeyond.com/>
- [17] J. E. Holmes, “Dungeons & dragons basic set,” *Dungeons & Dragons*, p. 39, 1977.
- [18] S. Gu, D. Chen, J. Bao, F. Wen, B. Zhang, D. Chen, L. Yuan, and B. Guo, “Vector quantized diffusion model for text-to-image synthesis,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022*, pp. 10 696–10 706.
- [19] J. Yu, Y. Xu, J. Y. Koh, T. Luong, G. Baid, Z. Wang, V. Vasudevan, A. Ku, Y. Yang, B. K. Ayan *et al.*, “Scaling autoregressive models for content-rich text-to-image generation,” *arXiv preprint arXiv:2206.10789*, 2022.
- [20] A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. Chen, and I. Sutskever, “Zero-shot text-to-image generation,” in *International Conference on Machine Learning*. PMLR, 2021, pp. 8821–8831.
- [21] L. T. Pereira, “Procedural generation of dungeon maps, missions and rooms,” Ph.D. dissertation, Universidade de São Paulo, 2018.
- [22] E. Papagiannopoulou and G. Tsoumakas, “A review of keyphrase extraction,” *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, vol. 10, no. 2, p. e1339, 2020.
- [23] J. Jayasudha and M. Thilagu, “A survey on sentimental analysis of student reviews using natural language processing (nlp) and text mining,” in *Innovations in Intelligent Computing and Communication: First International Conference, ICHCC 2022, Bhubaneswar, Odisha, India, December 16-17, 2022, Proceedings*. Springer, 2023, pp. 365–378.

- [24] S. Vijayarani, M. J. Ilamathi, M. Nithya *et al.*, “Preprocessing techniques for text mining-an overview,” *International Journal of Computer Science & Communication Networks*, vol. 5, no. 1, pp. 7–16, 2015.
- [25] B. Santana, R. Campos, E. Amorim, A. Jorge, P. Silvano, and S. Nunes, “A survey on narrative extraction from textual data,” *Artificial Intelligence Review*, pp. 1–43, 2023.
- [26] R. Srivastava, P. Singh, K. Rana, and V. Kumar, “A topic modeled unsupervised approach to single document extractive text summarization,” *Knowledge-Based Systems*, vol. 246, p. 108636, 2022.
- [27] S. Abdel-Salam and A. Rafea, “Performance study on extractive text summarization using bert models,” *Information*, vol. 13, no. 2, p. 67, 2022.
- [28] H. Y. Koh, J. Ju, M. Liu, and S. Pan, “An empirical survey on long document summarization: Datasets, models, and metrics,” *ACM computing surveys*, vol. 55, no. 8, pp. 1–35, 2022.
- [29] D. Yadav, J. Desai, and A. K. Yadav, “Automatic text summarization methods: A comprehensive review,” *arXiv preprint arXiv:2204.01849*, 2022.
- [30] D. O. Cajueiro, A. G. Nery, I. Tavares, M. K. De Melo, S. A. d. Reis, L. Weigang, and V. R. Celestino, “A comprehensive review of automatic text summarization techniques: method, data, evaluation and coding,” *arXiv preprint arXiv:2301.03403*, 2023.
- [31] F. S. Alotaibi and S. Sharma, “Distributed feature sets for document specific key-phrase extraction,” *Cybernetics and Systems*, pp. 1–29, 2022.
- [32] A. P. B. Veyseh, N. Meister, F. Derroncourt, and T. H. Nguyen, “Improving keyphrase extraction with data augmentation and information filtering,” *arXiv preprint arXiv:2209.04951*, 2022.
- [33] S. Beliga, “Keyword extraction: a review of methods and approaches,” *University of Rijeka, Department of Informatics, Rijeka*, vol. 1, no. 9, 2014.
- [34] S. Seo, S. Arik, J. Yoon, X. Zhang, K. Sohn, and T. Pfister, “Controlling neural networks with rule representations,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 11 196–11 207, 2021.
- [35] P. Liu, Y. Guo, F. Wang, and G. Li, “Chinese named entity recognition: The state of the art,” *Neurocomputing*, vol. 473, pp. 37–53, 2022.

- [36] L. Krasnov, I. Khokhlov, M. V. Fedorov, and S. Sosnin, “Transformer-based artificial neural networks for the conversion between chemical notations,” *Scientific Reports*, vol. 11, no. 1, pp. 1–10, 2021.
- [37] A. Conneau, K. Khandelwal, N. Goyal, V. Chaudhary, G. Wenzek, F. Guzmán, E. Grave, M. Ott, L. Zettlemoyer, and V. Stoyanov, “Unsupervised cross-lingual representation learning at scale,” *arXiv preprint arXiv:1911.02116*, 2019.
- [38] P. J. Chambon, C. Wu, J. M. Steinkamp, J. Adleberg, T. S. Cook, and C. P. Langlotz, “Automated deidentification of radiology reports combining transformer and “hide in plain sight” rule-based methods,” *Journal of the American Medical Informatics Association*, vol. 30, no. 2, pp. 318–328, 2023.
- [39] K. Clark, M.-T. Luong, Q. V. Le, and C. D. Manning, “Electra: Pre-training text encoders as discriminators rather than generators,” in *International Conference on Learning Representations*, 2020.
- [40] P. Wang, Z. Wang, X. Zhang, K. Wang, and Z. Feng, “Enhanced named entity recognition with semantic dependency,” in *PRICAI 2021: Trends in Artificial Intelligence: 18th Pacific Rim International Conference on Artificial Intelligence, PRICAI 2021, Hanoi, Vietnam, November 8–12, 2021, Proceedings, Part II 18*. Springer, 2021, pp. 287–298.
- [41] M. Benesty, “Ner algo benchmark: spacy, flair, m-bert and camembert on anonymizing french commercial legal cases,” *Towar. Data Sci*, 2019.
- [42] H. Shelar, G. Kaur, N. Heda, and P. Agrawal, “Named entity recognition approaches and their comparison for custom ner model,” *Science & Technology Libraries*, vol. 39, no. 3, pp. 324–337, 2020.
- [43] S. Naseer, M. M. Ghafoor, S. bin Khalid Alvi, A. Kiran, S. U. Rahmand, G. Murtazae, and G. Murtaza, “Named entity recognition (ner) in nlp techniques, tools accuracy and performance.” *Pakistan Journal of Multidisciplinary Research*, vol. 2, no. 2, pp. 293–308, 2021.
- [44] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” *arXiv preprint arXiv:1810.04805*, 2018.
- [45] J. Oppenlaender, “Prompt engineering for text-based generative art,” *arXiv preprint arXiv:2204.13988*, 2022.
- [46] S. W. McRoy, S. Channarukul, and S. S. Ali, “An augmented template-based approach to text realization,” *Natural Language Engineering*, vol. 9, no. 4, pp. 381–420, 2003.

- [47] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, “Language models are few-shot learners,” *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [48] A. Radford, K. Narasimhan, T. Salimans, I. Sutskever *et al.*, “Improving language understanding by generative pre-training,” 2018.
- [49] S. Michellmann, M. Kumar, K. A. Norman, and M. Toneva, “Large language models can segment narrative events similarly to humans,” Jan 2023. [Online]. Available: <https://arxiv.org/abs/2301.10297>
- [50] T. Zhang, V. Kishore, F. Wu, K. Q. Weinberger, and Y. Artzi, “Bertscore: Evaluating text generation with bert,” *arXiv preprint arXiv:1904.09675*, 2019.
- [51] K. Wang, C. Gou, Y. Duan, Y. Lin, X. Zheng, and F.-Y. Wang, “Generative adversarial networks: introduction and outlook,” *IEEE/CAA Journal of Automatica Sinica*, vol. 4, no. 4, pp. 588–598, 2017.
- [52] R. Wei, C. Garcia, A. El-Sayed, V. Peterson, and A. Mahmood, “Variations in variational autoencoders—a comparative evaluation,” *Ieee Access*, vol. 8, pp. 153 651–153 670, 2020.
- [53] L. Yang, Z. Zhang, Y. Song, S. Hong, R. Xu, Y. Zhao, Y. Shao, W. Zhang, B. Cui, and M.-H. Yang, “Diffusion models: A comprehensive survey of methods and applications,” *arXiv preprint arXiv:2209.00796*, 2022.
- [54] H. Huang, P. S. Yu, and C. Wang, “An introduction to image synthesis with generative adversarial nets,” *arXiv preprint arXiv:1803.04469*, 2018.
- [55] L. Wang, W. Chen, W. Yang, F. Bi, and F. R. Yu, “A state-of-the-art review on image synthesis with generative adversarial networks,” *IEEE Access*, vol. 8, pp. 63 514–63 537, 2020.
- [56] Z. Pan, W. Yu, X. Yi, A. Khan, F. Yuan, and Y. Zheng, “Recent progress on generative adversarial networks (gans): A survey,” *IEEE access*, vol. 7, pp. 36 322–36 333, 2019.
- [57] Y. Liao, K. Schwarz, L. Mescheder, and A. Geiger, “Towards unsupervised learning of generative models for 3d controllable image synthesis,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 5871–5880.

- [58] S. Reed, Z. Akata, X. Yan, L. Logeswaran, B. Schiele, and H. Lee, “Generative adversarial text to image synthesis,” in *International conference on machine learning*. PMLR, 2016, pp. 1060–1069.
- [59] P. Dhariwal and A. Nichol, “Diffusion models beat gans on image synthesis,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 8780–8794, 2021.
- [60] L. Xu, J. S. Ren, C. Liu, and J. Jia, “Deep convolutional neural network for image deconvolution,” *Advances in neural information processing systems*, vol. 27, 2014.
- [61] R. Durall, M. Keuper, and J. Keuper, “Watch your up-convolution: Cnn based generative deep neural networks are failing to reproduce spectral distributions,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 7890–7899.
- [62] A. A. Heydari and A. Mehmood, “Srvae: super resolution using variational autoencoders,” in *Pattern Recognition and Tracking XXXI*, vol. 11400. SPIE, 2020, pp. 87–100.
- [63] J. Ho, C. Saharia, W. Chan, D. J. Fleet, M. Norouzi, and T. Salimans, “Cascaded diffusion models for high fidelity image generation.” *J. Mach. Learn. Res.*, vol. 23, no. 47, pp. 1–33, 2022.
- [64] A. Nichol, P. Dhariwal, A. Ramesh, P. Shyam, P. Mishkin, B. McGrew, I. Sutskever, and M. Chen, “Glide: Towards photorealistic image generation and editing with text-guided diffusion models,” *arXiv preprint arXiv:2112.10741*, 2021.
- [65] G. Kim, T. Kwon, and J. C. Ye, “Diffusionclip: Text-guided diffusion models for robust image manipulation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 2426–2435.
- [66] X. Liu, D. H. Park, S. Azadi, G. Zhang, A. Chopikyan, Y. Hu, H. Shi, A. Rohrbach, and T. Darrell, “More control for free! image synthesis with semantic diffusion guidance,” in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2023, pp. 289–299.
- [67] Z. Lyu, X. Xu, C. Yang, D. Lin, and B. Dai, “Accelerating diffusion models via early stop of the diffusion process,” *arXiv preprint arXiv:2205.12524*, 2022.
- [68] W. Wang, J. Bao, W. Zhou, D. Chen, D. Chen, L. Yuan, and H. Li, “Semantic image synthesis via diffusion models,” *arXiv preprint arXiv:2207.00050*, 2022.

- [69] Y. Shih, S. Paris, C. Barnes, W. T. Freeman, and F. Durand, “Style transfer for headshot portraits,” 2014.
- [70] Y. Deng, F. Tang, W. Dong, C. Ma, X. Pan, L. Wang, and C. Xu, “Stytr2: Image style transfer with transformers,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 11 326–11 336.
- [71] T. Karras, S. Laine, and T. Aila, “A style-based generator architecture for generative adversarial networks,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 4401–4410.
- [72] Y. Li, C. Fang, J. Yang, Z. Wang, X. Lu, and M.-H. Yang, “Universal style transfer via feature transforms,” *Advances in neural information processing systems*, vol. 30, 2017.
- [73] G. Parmar, K. K. Singh, R. Zhang, Y. Li, J. Lu, and J.-Y. Zhu, “Zero-shot image-to-image translation,” *arXiv preprint arXiv:2302.03027*, 2023.
- [74] T. Wang, T. Zhang, B. Zhang, H. Ouyang, D. Chen, Q. Chen, and F. Wen, “Pretraining is all you need for image-to-image translation,” *arXiv preprint arXiv:2205.12952*, 2022.
- [75] S. Yang, L. Jiang, Z. Liu, and C. C. Loy, “Unsupervised image-to-image translation with generative prior,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 18 332–18 341.
- [76] C. Saharia, W. Chan, H. Chang, C. Lee, J. Ho, T. Salimans, D. Fleet, and M. Norouzi, “Palette: Image-to-image diffusion models,” in *ACM SIGGRAPH 2022 Conference Proceedings*, 2022, pp. 1–10.
- [77] A. Obukhov and M. Krasnyanskiy, “Quality assessment method for gan based on modified metrics inception score and fréchet inception distance,” in *Software Engineering Perspectives in Intelligent Systems: Proceedings of 4th Computational Methods in Systems and Software 2020, Vol. 1 4*. Springer, 2020, pp. 102–114.
- [78] M. Kokare, B. Chatterji, and P. Biswas, “Comparison of similarity metrics for texture image retrieval,” in *TENCON 2003. Conference on convergent technologies for Asia-Pacific region*, vol. 2. IEEE, 2003, pp. 571–575.
- [79] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, “Image quality assessment: from error visibility to structural similarity,” *IEEE transactions on image processing*, vol. 13, no. 4, pp. 600–612, 2004.

- [80] J. Wu, G. Shi, W. Lin, A. Liu, and F. Qi, “Just noticeable difference estimation for images with free-energy principle,” *IEEE Transactions on Multimedia*, vol. 15, no. 7, pp. 1705–1710, 2013.
- [81] T. Frese, C. A. Bouman, and J. P. Allebach, “Methodology for designing image similarity metrics based on human visual system models,” in *Human Vision and Electronic Imaging II*, vol. 3016. SPIE, 1997, pp. 472–483.
- [82] S. Behdenna, F. Barigou, and G. Belalem, “Document level sentiment analysis: a survey,” *EAI Endorsed Transactions on Context-aware Systems and Applications*, vol. 4, no. 13, pp. e2–e2, 2018.
- [83] E. Tomadaki and A. Salway, “Matching verb attributes for cross-document event co-reference,” in *Procs. Interdisciplinary Workshop on the Identification and Representation of Verb Features and Verb Classes*, 2005, pp. 127–132.
- [84] X. Han and L. Wang, “A novel document-level relation extraction method based on bert and entity information,” *IEEE Access*, vol. 8, pp. 96 912–96 919, 2020.
- [85] C. Girardi, M. Speranza, R. Sprugnoli, and S. Tonelli, “Cromer: A tool for cross-document event and entity coreference,” in *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14)*. LREC, 2014, pp. 3204–3208.
- [86] Y. Yao, D. Ye, P. Li, X. Han, Y. Lin, Z. Liu, Z. Liu, L. Huang, J. Zhou, and M. Sun, “Docred: A large-scale document-level relation extraction dataset,” *arXiv preprint arXiv:1906.06127*, 2019.
- [87] A. Cherian and A. Sullivan, “Sem-gan: Semantically-consistent image-to-image translation,” in *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 2019, pp. 1797–1806.
- [88] A. Maharana, D. Hannan, and M. Bansal, “Storydall-e: Adapting pretrained text-to-image transformers for story continuation,” in *European Conference on Computer Vision*. Springer, 2022, pp. 70–87.
- [89] T. Gupta, D. Schwenk, A. Farhadi, D. Hoiem, and A. Kembhavi, “Imagine this! scripts to compositions to videos,” in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 598–613.
- [90] L. Anne Hendricks, O. Wang, E. Shechtman, J. Sivic, T. Darrell, and B. Russell, “Localizing moments in video with natural language,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 5803–5812.

- [91] H. Jeong, G. Kwon, and J. C. Ye, “Zero-shot generation of coherent storybook from plain text story using diffusion models,” *arXiv preprint arXiv:2302.03900*, 2023.
- [92] T. Rahman, H.-Y. Lee, J. Ren, S. Tulyakov, S. Mahajan, and L. Sigal, “Make-a-story: Visual memory conditioned consistent story generation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 2493–2502.
- [93] X. Pan, P. Qin, Y. Li, H. Xue, and W. Chen, “Synthesizing coherent story with auto-regressive latent diffusion models,” *arXiv preprint arXiv:2211.10950*, 2022.
- [94] J. Cho, A. Zala, and M. Bansal, “Dall-eval: Probing the reasoning skills and social biases of text-to-image generation models,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 3043–3054.
- [95] P. Esser, R. Rombach, and B. Ommer, “Taming transformers for high-resolution image synthesis,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 12 873–12 883.
- [96] J. Liang, W. Pei, and F. Lu, “CpGAN: full-spectrum content-parsing generative adversarial networks for text-to-image synthesis,” *arXiv preprint arXiv:1912.08562*, 2019.
- [97] A. Maharana, D. Hannan, and M. Bansal, “Improving generation and evaluation of visual stories via semantic consistency,” *arXiv preprint arXiv:2105.10026*, 2021.
- [98] K. Frans, L. Soros, and O. Witkowski, “Clipdraw: Exploring text-to-drawing synthesis through language-image encoders,” *NeurIPS*, vol. 35, pp. 5207–5218, 2022.
- [99] T.-C. Wang, M.-Y. Liu, J.-Y. Zhu, A. Tao, J. Kautz, and B. Catanzaro, “High-resolution image synthesis and semantic manipulation with conditional GANs,” in *CVPR*, 2018, pp. 8798–8807.
- [100] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, “Lora: Low-rank adaptation of large language models,” *arXiv preprint arXiv:2106.09685*, 2021.
- [101] L. Zhang and M. Agrawala, “Adding conditional control to text-to-image diffusion models,” *arXiv preprint arXiv:2302.05543*, 2023.
- [102] T. Ceccherini-Silberstein, M. Coornaert, T. Ceccherini-Silberstein, and M. Coornaert, *Cellular automata*. Springer, 2010.

- [103] T. Archer, “Procedurally generating terrain,” in *44th annual midwest instruction and computing symposium, Duluth*, 2011, pp. 378–393.
- [104] S. G. Kobourov, “Spring embedders and force directed graph drawing algorithms,” *arXiv preprint arXiv:1201.3011*, 2012.
- [105] R. J. Trudeau, *Introduction to graph theory*. Courier Corporation, 2013.
- [106] Q. Wu, G. Bansal, J. Zhang, Y. Wu, S. Zhang, E. Zhu, B. Li, L. Jiang, X. Zhang, and C. Wang, “Autogen: Enabling next-gen llm applications via multi-agent conversation framework,” *arXiv preprint arXiv:2308.08155*, 2023.
- [107] A. Peiris and N. de Silva, “Synthesis and evaluation of a domain-specific large data set for dungeons & dragons,” *arXiv preprint arXiv:2212.09080*, 2022.
- [108] —, “Shade: Semantic hypernym annotator for domain-specific entities-dungeons and dragons domain use case,” *IEEE*, pp. 1–6, 2023.
- [109] P. Zhou, A. Zhu, J. Hu, J. Pujara, X. Ren, C. Callison-Burch, Y. Choi, and P. Ammanabrolu, “An ai dungeon master’s guide: Learning to converse and guide with intents and theory-of-mind in dungeons and dragons,” *arXiv preprint arXiv:2212.10060*, 2022.
- [110] T. K. Sang and D. M. EF, “F., 2003. introduction to the conll-2003 shared task: Language-independent named entity recognition,” in *Proceedings of the Conference on Computational Natural Language Learning. Edmonton, Canada*, pp. 142–147.
- [111] R. Weischedel *et al.*, “Ontonotes release 5.0 ldc2013t19,” Web Download, Philadelphia, 2013.
- [112] K. Sugathadasa, B. Ayesha, N. de Silva, A. S. Perera, V. Jayawardana, D. Lakmal, and M. Perera, “Synergistic union of word2vec and lexicon for domain specific semantic similarity,” in *2017 IEEE international conference on industrial and information systems (ICIIS)*. IEEE, 2017, pp. 1–6.
- [113] N. de Silva, D. Dou, and J. Huang, “Discovering inconsistencies in pubmed abstracts through ontology-based information extraction,” in *Proceedings of the 8th ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics*, 2017, pp. 362–371.
- [114] N. de Silva and D. Dou, “Semantic oppositeness assisted deep contextual modeling for automatic rumor detection in social networks,” in *Proceedings of the 16th Conference of the European Chapter of the Association for*

- Computational Linguistics: Main Volume*, P. Merlo, J. Tiedemann, and R. Tsarfaty, Eds. Online: Association for Computational Linguistics, Apr. 2021, pp. 405–415. [Online]. Available: <https://aclanthology.org/2021.eacl-main.31>
- [115] Y. Yao, S. Huang, W. Wang, L. Dong, and F. Wei, “Adapt-and-distill: Developing small, fast and effective pretrained language models for domains,” *arXiv preprint arXiv:2106.13474*, 2021.
- [116] B. Mohit, “Named entity recognition,” *Natural language processing of semitic languages*, pp. 221–245, 2014.
- [117] R. Baker and C. Perkins, “Lost mine of phandelver,” *Wizards of the Coast*, 2014.
- [118] W. Baur, S. Winter, and A. Winter, *Hoard of the Dragon Queen*. Wizards of the Coast, 2014.
- [119] —, *Rise of Tiamat*. Wizards of the Coast, 2014.
- [120] C. Perkins, T. Hickman, and L. Hickman, *Curse of Strahd*. Wizards of the Coast, 2016.
- [121] C. Perkins, W. Doyle, and S. Winter, *Tomb of Annihilation*. Wizards of the Coast, 2017.
- [122] C. Perkins, G. Barber, B. Benham *et al.*, *Candlekeep Mysteries*. Wizards of the Coast, 2021.
- [123] S. Allan, W. Doyle, A. Levitch, and C. Perkins, *The Wild Beyond the Witchlight*. Wizards of the Coast, 2021.
- [124] T. L. Scao, A. Fan, C. Akiki, E. Pavlick, S. Ilić, D. Hesslow, R. Castagné, A. S. Luccioni, F. Yvon, M. Gallé *et al.*, “Bloom: A 176b-parameter open-access multilingual language model,” *arXiv preprint arXiv:2211.05100*, 2022.
- [125] X. Geng and H. Liu, “Openllama: An open reproduction of llama,” May 2023. [Online]. Available: https://github.com/openlm-research/open_llama
- [126] Databricks, “Dolly: The first open source, commercially viable instruction-tuned llm,” *Databricks Blog*, April 2023. [Online]. Available: <https://www.databricks.com/blog/2023/04/12/dolly-first-open-commercially-viable-instruction-tuned-llm>
- [127] S. Tedeschi, V. Maiorca, N. Campolungo, F. Cecconi, and R. Navigli, “WikiNEuRal: Combined neural and knowledge-based silver data creation for multilingual NER,” in *Findings of the Association for Computational Linguistics: EMNLP 2021*, 2021, pp. 2521–2533.

- [128] J. Baptiste, “roberta-large-ner-english: model fine-tuned from roberta-large for ner task,” <https://huggingface.co/Jean-Baptiste/roberta-large-ner-english>, Jan 2022, (Accessed on 05/10/2023).
- [129] F. Souza, R. Nogueira, and R. Lotufo, “Portuguese named entity recognition using bert-crf,” *arXiv preprint arXiv:1909.10649*, 2019.
- [130] A. Akbik, D. Blythe, and R. Vollgraf, “Contextual string embeddings for sequence labeling,” in *COLING 2018, 27th International Conference on Computational Linguistics*, 2018, pp. 1638–1649.
- [131] M. Honnibal and I. Montani, “spacy 2: Natural language understanding with bloom embeddings, convolutional neural networks and incremental parsing,” *To appear*, 2017.
- [132] T. V. Nguyen, T. Nguyen, V. S. Ta, P. Q. N. Minh, V. X. Son, N. M. Hieu, N. P. Thai, V. T. Hung, P. M. Quang, N. T. M. Huyen, and H. Li, “Trankit: A light-weight transformer-based toolkit for multilingual natural language processing,” in *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics*, 2021.
- [133] G. R. Terrell and D. W. Scott, “Variable kernel density estimation,” *The Annals of Statistics*, pp. 1236–1265, 1992.
- [134] G. Weerasundara and N. de Silva, “Comparative analysis of named entity recognition in the dungeons and dragons domain,” 2023.
- [135] J. Ho, A. Jain, and P. Abbeel, “Denoising diffusion probabilistic models,” *Advances in neural information processing systems*, vol. 33, pp. 6840–6851, 2020.
- [136] N. Ruiz, Y. Li, V. Jampani, Y. Pritch, M. Rubinstein, and K. Aberman, “Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 22 500–22 510.
- [137] J. Crawford, D. Dillon, B. Petrisor, F. W. Schneider, and E. Teague, *Tasha’s Cauldron of Everything*. Wizards of the Coast Publishing, 2020. [Online]. Available: <https://books.google.lk/books?id=WARPEAAAQBAJ>
- [138] J. Li, D. Li, C. Xiong, and S. Hoi, “Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation,” in *International Conference on Machine Learning*. PMLR, 2022, pp. 12 888–12 900.

- [139] G. Jocher, A. Chaurasia, A. Stoken, J. Borovec, Y. Kwon, K. Michael, J. Fang, Z. Yifu, C. Wong, D. Montes *et al.*, “ultralytics/yolov5: v7. 0-yolov5 sota realtime instance segmentation,” *Zenodo*, 2022.
- [140] O. Avrahami, A. Hertz, Y. Vinker, M. Arar, S. Fruchter, O. Fried, D. Cohen-Or, and D. Lischinski, “The chosen one: Consistent characters in text-to-image diffusion models,” *arXiv preprint arXiv:2311.10093*, 2023.
- [141] M. Oquab, T. Darcet, T. Moutakanni, H. Vo, M. Szafraniec, V. Khalidov, P. Fernandez, D. Haziza, F. Massa, A. El-Nouby *et al.*, “Dinov2: Learning robust visual features without supervision,” *arXiv preprint arXiv:2304.07193*, 2023.
- [142] D. Arthur and S. Vassilvitskii, “K-means++ the advantages of careful seeding,” in *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, 2007, pp. 1027–1035.
- [143] M. Zermani, M.-C. Larabi, and C. Fernandez-Maloigne, “A comprehensive assessment of the structural similarity index,” *Signal Processing: Image Communication*, vol. 99, p. 116336, 2021.
- [144] H. Liu, C. Li, Q. Wu, and Y. J. Lee, “Visual instruction tuning,” *arXiv preprint arXiv:2304.08485*, 2023.
- [145] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, “Gans trained by a two time-scale update rule converge to a local nash equilibrium,” *NeurIPS*, vol. 30, 2017.
- [146] T. Brooks, A. Holynski, and A. A. Efros, “Instructpix2pix: Learning to follow image editing instructions,” *arXiv preprint arXiv:2211.09800*, 2022.