

1 A Latent Confounders and Proxy

2 It is infeasible to estimate causal effects without any observed confounding information [12, 2]. With
3 unobserved confounders, a common approach is to introduce additional proxy variables [11, 1, 15].
4 For example, the socio-economic status of a patient is a confounder to the medication and outcome,
5 which is unobserved and hard to measure directly. But we can use known auxiliary information such
6 as the zip code and job type as a proxy to assess the confounder indirectly [8]. Note that directly
7 treating the proxy variables as ordinary confounders can induce bias [13, 4, 3]. Instead, recent work
8 [8, 16, 9, 10, 14] uses latent variable approaches. For example, Louizos et al. [8] use VAEs to infer
9 latent confounders from proxy information. Our work develops a latent variable causal framework for
10 controllable text generation. Due to the rich and abstract information in text, we introduce necessary
11 training objectives to avoid training collapse and encourage confounder balance.

12 B Experimental Details

13 **Model configuration** The VAE backbone of our model largely follows the architecture of the VAE
14 model in [6], except that our inference network (encoder) q_ϕ adopts the pretrained GPT-2 model
15 (same as the decoder $p_\theta(x|a, z)$) instead of BERT, in order to make sure both encoder q_ϕ and decoder
16 $p_\theta(x|a, z)$ have the same tokenization. We implement a to be either an all-zero or all-one vector
17 of dimension 50, and set the dimension of z to 718, so that concatenating a and z leads to a vector
18 of dimension 768, same as in [6]. We implement $p_\theta(c|z)$ as a single-layer MLP and $p_\theta(a|z)$ as a
19 two-layer MLP with the intermediate dimension the same as the input dimension (768).

20 **More details of BIOS dataset** For occupation which is the confounding factor, we subsample and
21 merge the occupations into two groups, i.e., $\{nurse, dietitian, paralegal, model, yoga\ teacher\}$, and
22 $\{rapper, DJ, surgeon, software\ engineer, composer\}$, which results in a correlation strength of 94%
23 between occupation and gender.

24 **Training confounding label classifier with data re-weighting** For the Conditional LM
25 (full) baseline for attribute-conditional generation, we first train a confounding label classifier
26 with the limited confounding (proxy) labels in the dataset. Due to the strong correlation between the
27 confounding factor and attribute (e.g., with a correlation strength of 90%), only a small fraction of
28 (e.g., 10%) instances have opposite confounding label and attribute label. We thus train the classifier
29 with data reweighting to reduce the bias. Specifically, we associate a weight of 0.9 to those instances
30 with opposite confounding and attribute labels, and a weight of 0.1 to other instances with the same
31 confounding and attribute labels. We tried other weights and obtained lower or similar classifier
32 accuracy.

33 **Human evaluation of text attribute transfer** Following previous work [e.g., 7], we conduct
34 comparison-based human evaluation for the output of different generation models. Specifically, for
35 each test instance, we present the outputs of two comparison models to the human rater, and ask the
36 human rater to rank which of the two outputs are better in terms of the goal of the task (i.e., accurately
37 rewriting the text to possess desired attribute and meanwhile preserving all other characteristics of the
38 original sentence). The human rater can also choose “no preference” if the two outputs are equally
39 good or bad. We asked three human annotators (who are graduate students and proficient English
40 speakers) to do the rating [7]. There were no potential participant risks. We evaluate on 50 test
41 cases for each pair of comparison models. Table 1 shows the results, which are consistent with the
42 observations from automatic evaluation.

	Ours better (%)	No preference (%)	Ours worse (%)
Hu et al. [5]	62	24	14
Ablation: Ours w/o cf - z/c	54	22	24

Table 1: Human evaluation of text attribute transfer on biased YELP. For example, the outputs of OURS are considered to be better than those of Hu et al. [5] on 62% test instances.

43 **Classifiers used in training and evaluation** We summarize the different classifiers used in training
44 and evaluation in the experiment to serve as a reference and avoid confusion.

- 45 • For training:
 - 46 – *attribute classifier* f (Eq.4) is used to train our causal model. The classifier is pretrained
 - 47 with the biased (attribute, text) training corpus.
 - 48 – *confounding label classifier* is used in the baseline Conditional LM (full) for
 - 49 attribute-conditional generation. The classifier is trained with the available confounding
 - 50 (proxy) labels with data re-weighting, as discussed above.
- 51 • For evaluation:
 - 52 – *evaluation attribute classifier* is used to evaluate the generation accuracy of desired
 - 53 attribute. As an evaluation metric, the classifier is obtained by training on additional
 - 54 *unbiased* (attribute, text) data.
 - 55 – *evaluation confounding label classifier* is used to evaluate the correlation of attribute
 - 56 and confounding (proxy) labels in the generation. Similarly, the classifier is obtained
 - 57 by training on additional large unbiased data.

58 C Generated Samples

CONDITIONAL LM (FULL)

$a = 0$ (sentiment negative)

this was the worst experience i 've ever had at a glazier .
 i even asked him if they could play on the tv channel .
 this was pretty fun the first time i went . "
 waited in line once but almost never reached the floor .
 if you are ever up in chandler , tony will stop by .

$a = 1$ (sentiment positive)

very good and long wait time .
 we loved our favorite harrah 's night ! "
 i would love to try this restaurant again when they open . "
 this place is great .
 everything you will find in this restaurant !

OURS

$a = 0$ (sentiment negative)

no , it 's obvious that they were overcooked .
 the seats were poorly done and basically sucked up .
 it was n't enough to ask us if it was okay .
 very disappointed with my food order yesterday .
 i declined to replace it tho they were bad .

$a = 1$ (sentiment positive)

great for a relaxed evening out .
 i 'm beyond impressed with the passion fruit and unbeatable service .
 it 's a true pleasure to meet andrew .
 jacksville became my go-to spot for dessert .
 thank you for the technique , i am quite impressed .

Table 2: Attribute-conditional generation trained on YELP dataset. CONDITIONAL LM (FULL) tends to generate non-restaurant reviews conditioning on $a = 0$, and restaurant reviews conditioning on $a = 1$.

OURS

$a = 0$ (sentiment negative) $\rightarrow a = 1$ (positive)

original: pick-up was just ok , but vehicle was filthy and had trash in it .

output: pick-up was pretty good , but atmosphere was just incredible and comfortable .

original: so that was nice but they served some sweet concoctions that made me sick .

output: so good that they served some sweet and flavorful cocktails that made me super happy .

original: similar to some of the other reviewers , the poutine was just that .

output similar to some of the other reviewers , the poutine was just perfect .

$a = 1$ (sentiment positive) $\rightarrow a = 0$ (negative)

original: the santa fe salad is awesome .

output: the santa fe salad is mediocre .

original: the employees were super helpful , friendly and attentive .

output: the employees were super rude , incompetent and unhelpful .

original: i love their eggs benedict and pancakes both are amazing !

output: i hate their eggs benedict and pancakes both are horrible .

Table 3: Text attribute transfer on the biased YELP dataset.

59 References

- 60 [1] J. D. Angrist and J.-S. Pischke. *Mostly harmless econometrics: An empiricist’s companion*.
61 Princeton university press, 2008.
- 62 [2] A. D’Amour. On multi-cause approaches to causal inference with unobserved confounding:
63 Two cautionary failure cases and a promising alternative. In *The 22nd International Conference*
64 *on Artificial Intelligence and Statistics*, pages 3478–3486. PMLR, 2019.
- 65 [3] W. A. Fuller. *Measurement error models*, volume 305. John Wiley & Sons, 2009.
- 66 [4] Z. Griliches and J. A. Hausman. Errors in variables in panel data. *Journal of econometrics*, 31
67 (1):93–118, 1986.
- 68 [5] Z. Hu, Z. Yang, X. Liang, R. Salakhutdinov, and E. Xing. Toward controlled generation of text.
69 In *International Conference on Machine Learning (ICML)*, 2017.
- 70 [6] C. Li, X. Gao, Y. Li, X. Li, B. Peng, Y. Zhang, and J. Gao. Optimus: Organizing sentences via
71 pre-trained modeling of a latent space. *arXiv preprint arXiv:2004.04092*, 2020.
- 72 [7] S. Lin, W. Wang, Z. Yang, X. Liang, F. F. Xu, E. Xing, and Z. Hu. Record-to-text generation
73 with style imitation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural*
74 *Language Processing: Findings*, pages 1589–1598, 2020.
- 75 [8] C. Louizos, U. Shalit, J. M. Mooij, D. Sontag, R. Zemel, and M. Welling. Causal effect
76 inference with deep latent-variable models. In *Advances in neural information processing*
77 *systems (NeurIPS)*, pages 6446–6456, 2017.
- 78 [9] D. Lu, C. Tao, J. Chen, F. Li, F. Guo, and L. Carin. Reconsidering generative objectives for
79 counterfactual reasoning. In *NeurIPS*, 2020.
- 80 [10] D. Madras, E. Creager, T. Pitassi, and R. Zemel. Fairness through causal awareness: Learning
81 causal latent-variable models for biased data. In *Proceedings of the Conference on Fairness,*
82 *Accountability, and Transparency*, pages 349–358, 2019.

- 83 [11] M. R. Montgomery, M. Gragnolati, K. A. Burke, and E. Paredes. Measuring living standards
84 with proxy variables. *Demography*, 37(2):155–174, 2000.
- 85 [12] J. Pearl. *Causality*. Cambridge university press, 2009.
- 86 [13] J. Pearl. On measurement bias in causal inference. *arXiv preprint arXiv:1203.3504*, 2012.
- 87 [14] R. Ranganath and A. Perotte. Multiple causal inference with latent confounding. *arXiv preprint*
88 *arXiv:1805.08273*, 2018.
- 89 [15] J. H. Stock, M. W. Watson, et al. *Introduction to econometrics*, volume 3. Pearson New York,
90 2012.
- 91 [16] D. Tran and D. M. Blei. Implicit causal models for genome-wide association studies. In
92 *International Conference on Learning Representations*, 2018.