Temperature Is All You Need

Massimo Caccia*, Lucas Caccia*, William Fedus,
Hugo Larochelle, Joelle Pineau, Laurent Charlin

MILA, Universite de Montreal, RLlab, McGill University & Google Brain

September 1, 2018

tl;dr

tl;dr

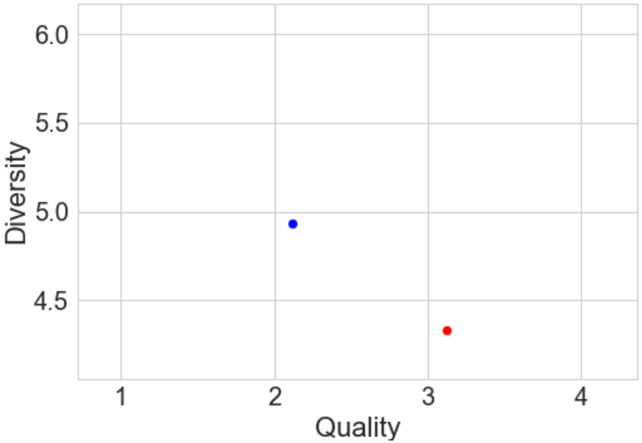

(lower is better for both metrics)

tl;dr

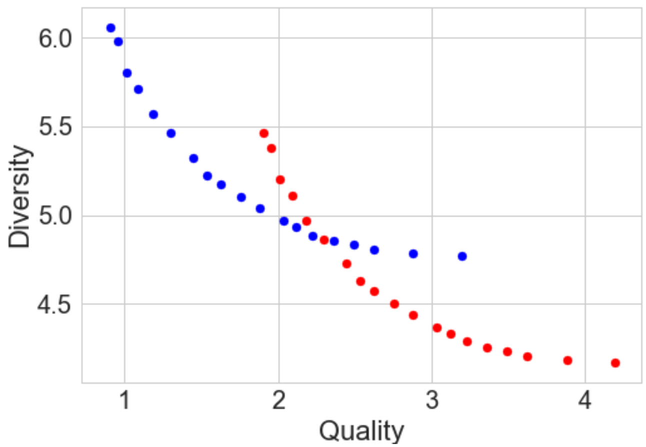

tl;dr

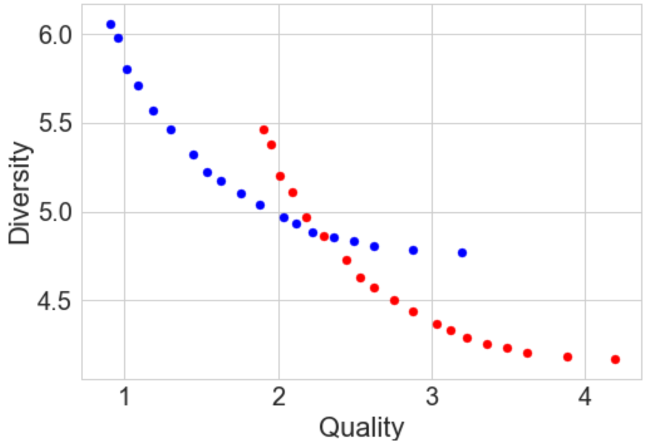

- The most effective way to evaluate NLG models is to compare them in quality/diversity space w.r.t multiple temperatures

tl;dr

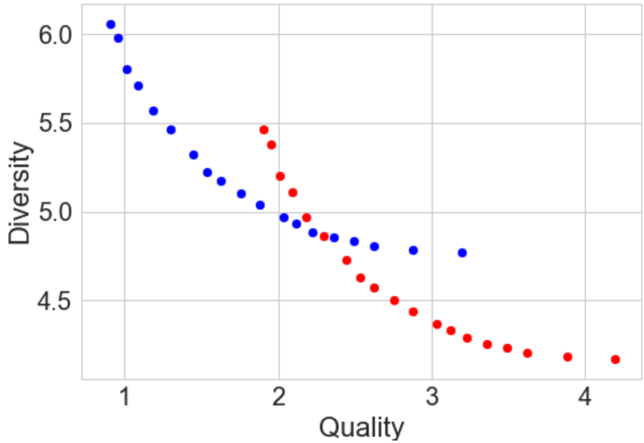

- The most effective way to evaluate NLG models is to compare them in quality/diversity space w.r.t multiple temperatures
- Maximum-likelihood training is superior (for now) to Textual GANs

MLE trained models generate poor samples

(1) Researchers are expected to comment on where a scheme is sold , but it is no longer this big name at this point .
(2) We know you ' re going to build the kind of home you ' re going to be expecting it can give us a better understanding of what ground test we ' re on this year , he explained .

Table 1: MLE samples lack in global coherence

MLE trained models generate poor samples

(1) Researchers are expected to comment on where a scheme is sold , but it is no longer this big name at this point .
(2) We know you ’ re going to build the kind of home you ’ re going to be expecting it can give us a better understanding of what ground test we ’ re on this year , he explained .

Table 1: MLE samples lack in global coherence

It is hypothesized that this is in-part because of the well-known issues of exposure bias [1], [2].

GAN generate great samples (in images)

GAN generate great samples (in images)

Figure 1: GAN generated images from [3]

GAN visualized

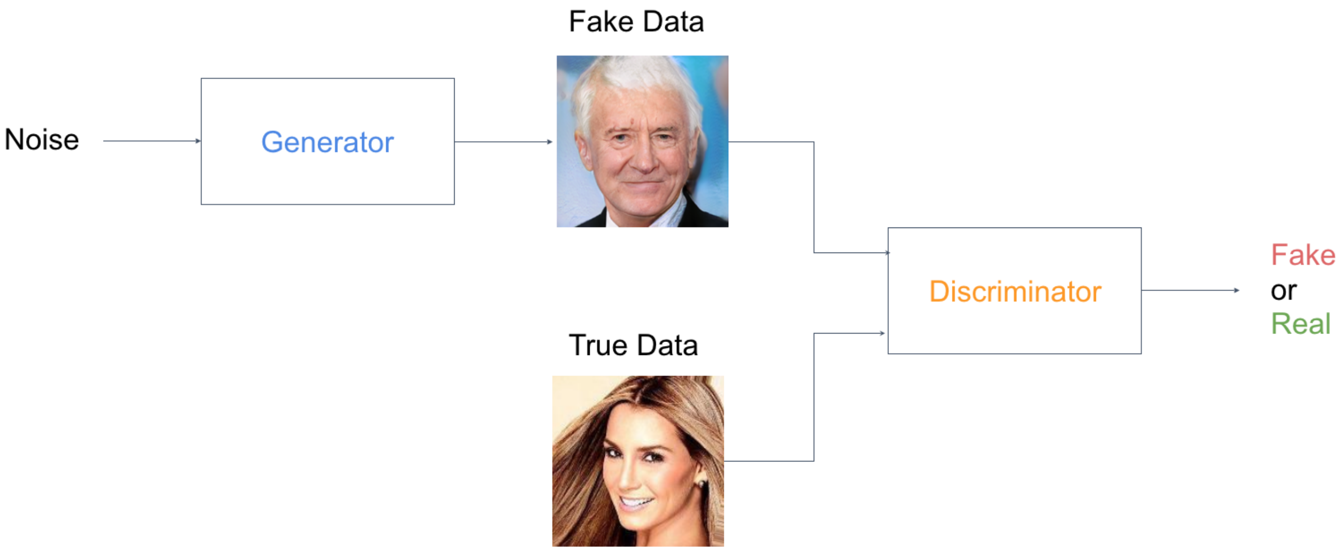

Figure 2: GAN [4] Framework visualized

non-differentiable in output space

Difficult to backpropagate through discrete variables

non-differentiable in output space

Difficult to backpropagate through discrete variables

- Straight-through estimator [5]
- Gumbel softmax [6]
- Reinforcement Learning (in particular REINFORCE [7])

Recent advancements in Sequential GANs

Recent advancements in Sequential GANs

SeqGAN [8]: First sequential GAN (trained with REINFORCE)

SeqGAN [8]: First sequential GAN (trained with REINFORCE)

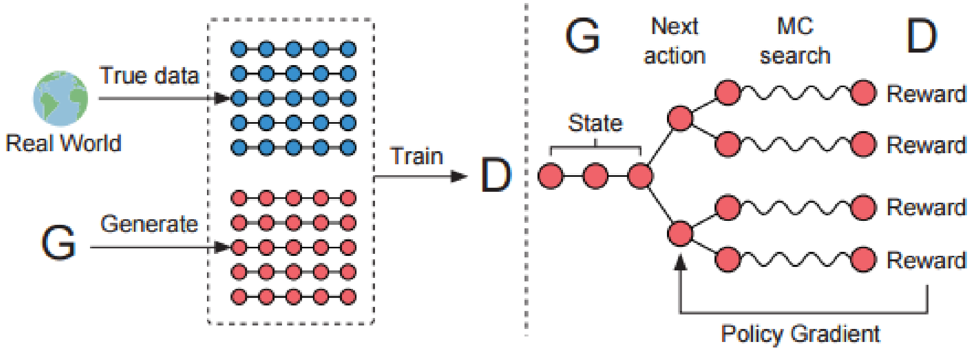

Figure 3: A cartoon of SeqGAN

Recent advancements in Sequential GANs

Recent advancements in Sequential GANs

- MaliGAN [9]: rescale the reward to alleviate the vanishing gradient

Recent advancements in Sequential GANs

- MaliGAN [9]: rescale the reward to alleviate the vanishing gradient
- RankGAN [10]: replace binary classification with ranking score

Recent advancements in Sequential GANs

- MaliGAN [9]: rescale the reward to alleviate the vanishing gradient
- RankGAN [10]: replace binary classification with ranking score
- LeakGAN [11]: Discriminator *leaks* information to the Generator

Recent advancements in Sequential GANs

- MaliGAN [9]: rescale the reward to alleviate the vanishing gradient
- RankGAN [10]: replace binary classification with ranking score
- LeakGAN [11]: Discriminator *leaks* information to the Generator
- TextGAN [12]: MMD loss + latent regularization

Recent advancements in Sequential GANs

- MaliGAN [9]: rescale the reward to alleviate the vanishing gradient
- RankGAN [10]: replace binary classification with ranking score
- LeakGAN [11]: Discriminator *leaks* information to the Generator
- TextGAN [12]: MMD loss + latent regularization
- a lot more: DPGAN, GSGAN, IRLGAN, etc

flawed evaluation protocol

flawed evaluation protocol

- BLEU [13] is an n-gram overlap metric most popular in MT

flawed evaluation protocol

- BLEU [13] is an n-gram overlap metric most popular in MT
- To advertise the Textual GANs, *corpus-level* BLEU was designed

flawed evaluation protocol

- BLEU [13] is an n-gram overlap metric most popular in MT
- To advertise the Textual GANs, *corpus-level* BLEU was designed
- *corpus-level* BLEU (alone) is really bad

quick fix

Self-BLEU [14]: BLEU score between generated sentences

quick fix

Self-BLEU [14]: BLEU score between generated sentences

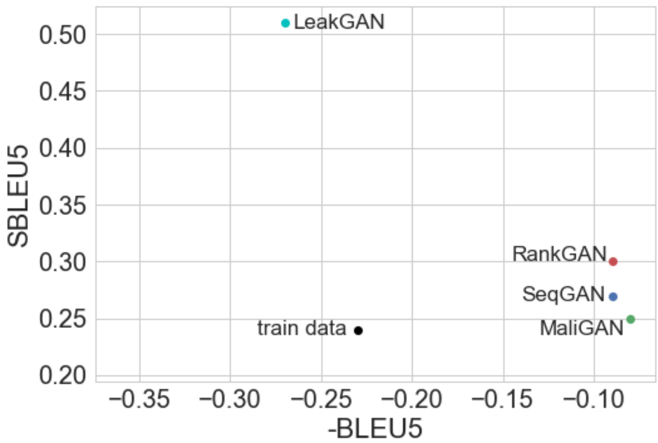

Figure 4: Results taken from [15]. Lower is better for both metrics.

temperature tuning

temperature tuning

$$G(y_t|y_{1:t-1}) = \text{softmax}(o_t \cdot W/\alpha).$$

o_t: generator's pre-logits activation at t
W: word embedding matrix
α: temperature parameter

temperature tuning

α	
2	(1) If you go at watch crucial characters putting awareness in Washington , forget there are now unique developments organized personally then why charge . (2) Front wants zero house blood number places than above spin 5 provide school projects which youth particularly teenager temporary dollars plenty of investors enjoy headed Japan about if federal assets own , at 41 .
0.7	(1) The other witnesses are believed to have been injured , the police said in a statement , adding that there was no immediate threat to any other witnesses . (2) The company ' s net income fell to 5 . 29 billion , or 2 cents per share , on the same period last year .
0	(1) The company ' s shares rose 1 . 5 percent to 1 . 81 percent , the highest since the end of the year . (2) The company ' s shares rose 1 . 5 percent to 1 . 81 percent , the highest since the end of the year .

Table 2: Effect of varying the temperature of the softmax layer in an autoregressive language model

Quality/Diversity space w.r.t temperatures

Quality/Diversity space w.r.t temperatures

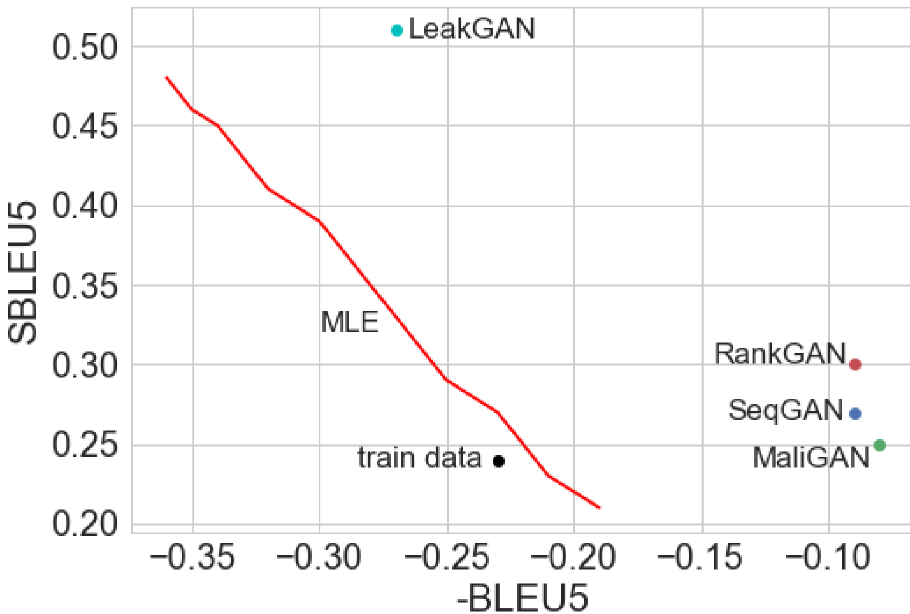

New (global) metrics in town

New (global) metrics in town

- Language Model score (quality): likelihood of the generated sentences under a LM

New (global) metrics in town

- Language Model score (quality): likelihood of the generated sentences under a LM
- Reverse LM score [16] (diversity + quality)

Quality/Diversity space w.r.t temperatures (global)

Quality/Diversity space w.r.t temperatures (global)

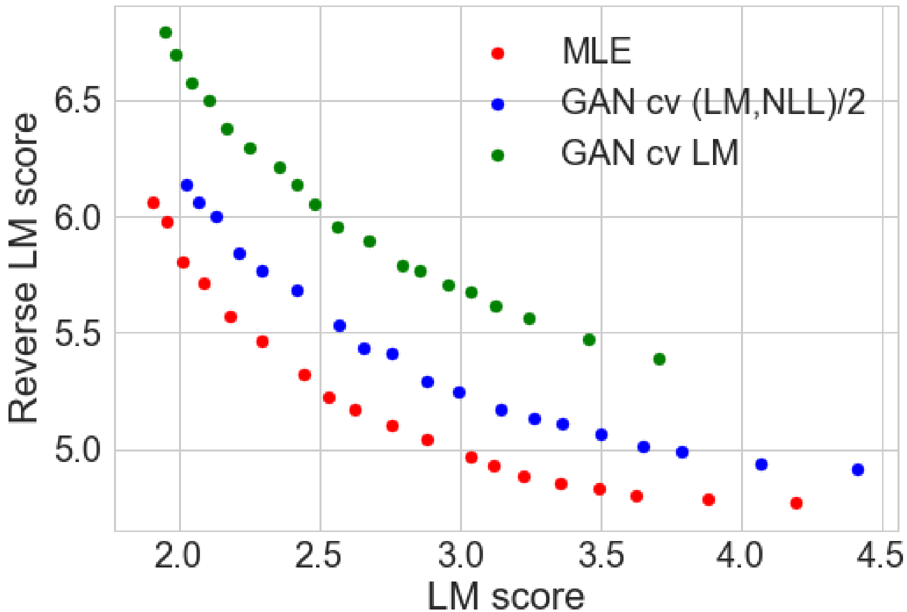

Figure 5: lower is better for both metrics

Conclusion

Conclusion

- most effective way to evaluate NLG models is in quality/diversity space w.r.t. multiples temperatures

Conclusion

- most effective way to evaluate NLG models is in quality/diversity space w.r.t. multiples temperatures
- MLE training is superior (for now) to adversarial training

Future Work

Future Work

- What about adversarial latent space regularization ?

Future Work

- What about adversarial latent space regularization ?
- Temperature Is All You Need for Image Generation ?

References

Marc'Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba.
Sequence level training with recurrent neural networks.
arXiv preprint arXiv:1511.06732, 2015.

Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer.
Scheduled sampling for sequence prediction with recurrent neural networks.
In *Advances in Neural Information Processing Systems*, pages 1171–1179, 2015.

Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen.
Progressive growing of gans for improved quality, stability, and variation.
arXiv preprint arXiv:1710.10196, 2017.

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*, pages 2672–2680, 2014.

Yoshua Bengio, Nicholas Léonard, and Aaron Courville.
Estimating or propagating gradients through stochastic neurons for conditional computation.
arXiv preprint arXiv:1308.3432, 2013.

Eric Jang, Shixiang Gu, and Ben Poole.
Categorical reparameterization with gumbel-softmax.
arXiv preprint arXiv:1611.01144, 2016.

Ronald J Williams.
Simple statistical gradient-following algorithms for connectionist reinforcement learning.
Machine learning, 8(3-4):229–256, 1992.

Lantao Yu, Weinan Zhang, Jun Wang, and Yong Yu.
Seqgan: Sequence generative adversarial nets with policy gradient.
2017.