To Beam Or Not To Beam: That is a Question of Cooperation for Language GANs

Thomas Scialom, Paul-Alexis Dray, Sylvain Lamprier, Benjamin Piwowarski, Jacopo Staiano

ReciTAL & Sorbonne Université, CNRS, LIP6, F-75005 Paris, France

An Important Problem

Language Generation: MLE, an imperfect training

⇒ Exposure Bias
(i.e. mismatch between training and inference)

Discriminators are very accurate:

⇒ Distinguish between human and machine texts with an accuracy > 90 [4, 3]

Two ways to leverage discriminators:

i) Inference: Cooperative Decoding Reranking generator's probabilities wrt the discriminator

- [3] used a BeamSearch
- [2] used Nucleus or Top-K Sampling

⇒ Both decoding suffer from the 'Left To Right Curse'

ii) Training: GANs

Discrete data implies using reinforcement learning (no gradient from the discriminator)

- Reinforcement Learning
- Sparse reward, unstable training

⇒ Existing language GANs are known to fall short [1]

Contributions

i) Coop-MCTS:
A new cooperative decoding mechanism beyond the left-to-right curse based on Monte Carlo Tree Search (MCTS).

ii) SelfGAN:
A new framework to propagate the discriminator signal in discrete GANs leveraging cooperative decoding mechanisms.

Beyond the left-to-right curse

Coop-MCTS: new cooperative decoding based on MCTS

- Policy Network: the generator
- Value Network: the discriminator

Conditioned Answer: *Super Bowl*
Context:

Super Bowl 50 was an American football game to determine the champion of the National Football League (NFL) for the 2015 season. The American Football Conference (AFC) champion Denver Broncos defeated the National Football Conference (NFC) champion Carolina Panthers 24â€"10 to earn their third Super Bowl title. The game was played on February 7, 2016, at Levi's Stadium in the San Francisco Bay Area at Santa Clara, California. As this was the 50th Super Bowl, the league emphasized the "golden anniversary" with various gold-themed initiatives, as well as temporarily suspending the tradition of naming each Super Bowl game with Roman numerals (under which the game would have been known as "Super Bowl L"), so that the logo could prominently feature the Arabic numerals

Step 01: *What*
⋮
Step 16: *What was the name of the game that would have been known as "Super Bowl*
Step 17: *How*
⋮
Step 46: *How is called the American football game that determines the NFL champion?*

Table 4: Progressive results obtained by our *Coop-MCTS* decoding method on Question Generation during a simulation. Until the 16th step, the generation is left-to-right. Then, the cooperation mechanism kicks in, allowing the model to safely abort this beam, by restarting a new question with *How*. We report the cross-attention weights on the input context for step 16 (red) and 17 (blue).

SelfGAN

A new GAN framework for discrete data

- Adversarial Sampling from a Cooperative Decoding
- The reward is infused directly to the generator
- Works with any Cooperative Decoding

Algorithm 1 *SelfGAN*

Input: a generator gen, a discriminator $discr$, and a cooperative decoding method $decod_{coop}$
for n epochs **do**
  for X, S_{ref} in training set **do** ▷ *Start Training*
    $S_{coop} \leftarrow decod_{coop}(X, gen, discr)$
    gen.train(srcs=X, tgts=S_{coop}) ▷ *Standard maximum likelihood but with S_{coop} as the target, and not S_{ref}*
    $discr$.train(srcs=X, human_exs= S_{ref}, machine_exs=S_{coop})

⇒ The signal from the discriminator is passed to the generator in a completely new way

Unconditional Text Generation

Model	T=0.5	T=1	T=2
MLE+Sample	0.42;0.29	0.31;0.11	0.18;0.07
ColdGAN+Sample	0.47;0.21	0.33;0.08	0.22;0.06
MLE+CoopMCTS	0.45;0.22	0.34;0.10	0.21;0.06
SelfGANCoopMCTS+CoopMCTS	0.48;0.20	0.37;0.09	0.24;0.05

Results on Unconditional Generation for samples realized at three different temperatures, in terms of BLEU Vs Self-BLEU (higher better;lower better)

Conditional NLG

Considered Task: Summarization (CNN/DM) and Question Generation (SQuAD)

Generator / Decoder	QG B4	QG R1	QG RL	QG Base	QG Base+	Sum. B4	Sum. R1	Sum. RL	Sum. Base	Sum. Base+
MLE										
BeamSearch [26]	19,7	45,2	41,1	15%	15%	15,9	42,3	40,4	9%	8%
DAS_{local} [32]	19,9	45,2	41,1	28%	19%	16,6	43,8	40,9	17%	11%
DAS_{global} [7]	20,0	45,2	41,2	20%	17%	16,2	44,1	41,9	12%	9%
Coop-MCTS	19.8	45,3	41,5	33%	21%	16,3	42,5	40,6	20%	12%
ColdGAN										
BeamSearch [31]	19.9	45,2	41,4	26%	17.9%	16,3	42,8	40,7	15%	10%
DAS_{local}	19,8	45,3	41,1	31%	20%	15,9	42,5	42,0	19%	11%
DAS_{global}	20,2	45,6	41,5	26%	18%	16,6	44,6	41,2	16%	10%
Coop-MCTS	19,9	45,4	41,2	39%	22%	15,9	44,2	41,2	23%	12%
$SelfGAN_{DAS_{loc}}$										
BeamSearch	20,2	45,4	41,6	27%	21%	16,9	44,2	42,5	16%	11%
DAS_{local}	20,5	45,5	41,7	30%	23%	16,9	44,4	41,9	18%	13%
DAS_{global}	20,1	45,4	41,7	33%	20%	16,6	44,0	42,3	19%	11%
Coop-MCTS	20,4	45,5	41,8	39%	23%	16,4	43,8	42,8	23%	13%
$SelfGAN_{DAS_{glob}}$										
BeamSearch	20,4	45,5	41,7	24%	19%	16,9	43,0	41,5	14%	11%
DAS_{local}	19,9	45,4	41,3	32%	22%	15,9	42,7	40,6	18%	12%
DAS_{global}	20,7	45,6	41,9	29%	20%	17,0	43,7	42,6	17%	11%
Coop-MCTS	20,0	45,3	41,4	40%	24%	16,1	43,4	42,3	23%	13%
$SelfGAN_{Coop\text{-}MCTS}$										
BeamSearch	20,5	46,6	42,6	34%	21%	17,0	42,8	41,5	20%	13%
DAS_{local}	20,6	46,7	41,7	42%	24%	16,6	43,7	42,8	25%	13%
DAS_{global}	20,5	46,6	41,7	39%	21%	16,5	42,8	40,9	23%	12%
Coop-MCTS	**21,1**	**48,9**	**44,7**	40%	**26%**	**17,5**	**43,5**	**42,3**	23%	**15%**

Table 1: Results of our experiments on QG (left) and Summarization (right). For each generator, we report the results with the four different decoders. The reported metrics correspond to BLEU4 (B4), ROUGE-1 (R1), ROUGE-L (RL) and the discriminators Base and Base+ as described in Section 5.3. For Base and Base+ the scores correspond to the probability of being human, so higher is better for all the metrics. For *SelfGAN*$_{MCTS}$, we experimented with 5 different seeds and the standard deviation is always inferior to 0.1 for BLEU4 and ROUGE, and inferior to 0.5% for Base and Base+.

Colinearity of the gradients

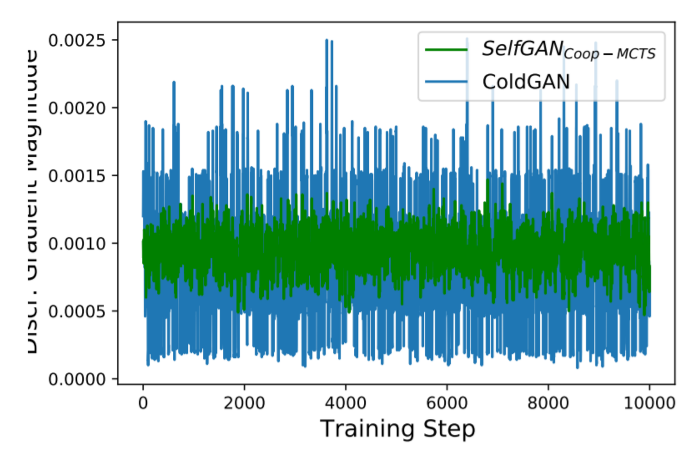

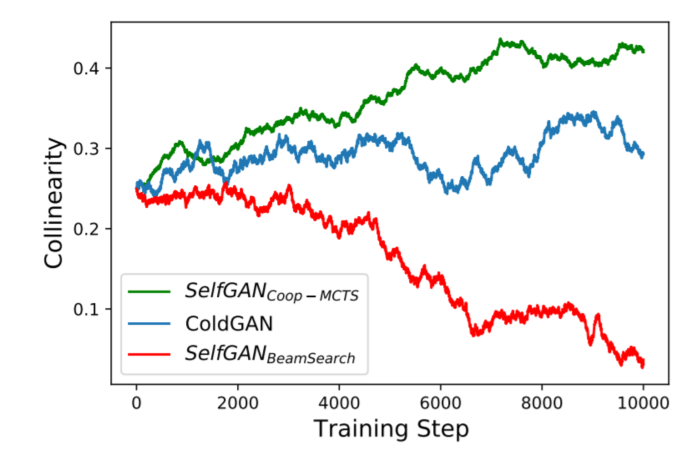

Left: Moving Average of the magnitude of the *discriminators* gradients during training. Right: colinearity of the *generators* gradients between the sampled texts and their corresponding human reference for *SelfGAN*$_{\text{Coop-MCTS}}$, ColdGAN and *SelfGAN*$_{\text{BeamSearch}}$. Both on Summarization.

References

[1] Massimo Caccia et al. "Language GANs Falling Short". In: *International Conference on Learning Representations*. 2020.

[2] Yuntian Deng et al. "Residual energy-based models for text generation". In: *arXiv preprint arXiv:2004.11714* (2020).

[3] Thomas Scialom et al. "Discriminative Adversarial Search for Abstractive Summarization". In: *arXiv preprint arXiv:2002.10375* (2020).

[4] Rowan Zellers et al. "Defending against neural fake news". In: *Advances in Neural Information Processing Systems*. 2019, pp. 9051–9062.