UMR IRISA

Institut de Recherche en Informatique et Systèmes Aléatoires

Which Discriminator for Cooperative Text Generation?

Antoine Chaffin, Thomas Scialom, Sylvain Lamprier, Jacopo Staiano, Benjamin Piwowarski, Ewa Kijak, Vincent Claveau

SORBONNE UNIVERSITÉ

MATAG

reciTAL.

UMR IRISA

Introduction

Language modeling

- Probability of the **next word given past ones**
- **Iteratively add tokens** to produce text
- Text generation can be seen as **tree exploration**
- **Root is the prompt**, each node correspond to its **parent sequence with an additional token**

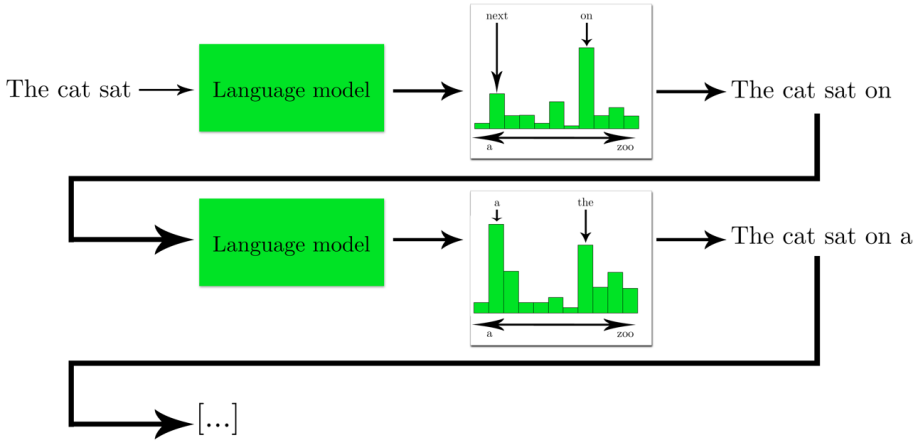

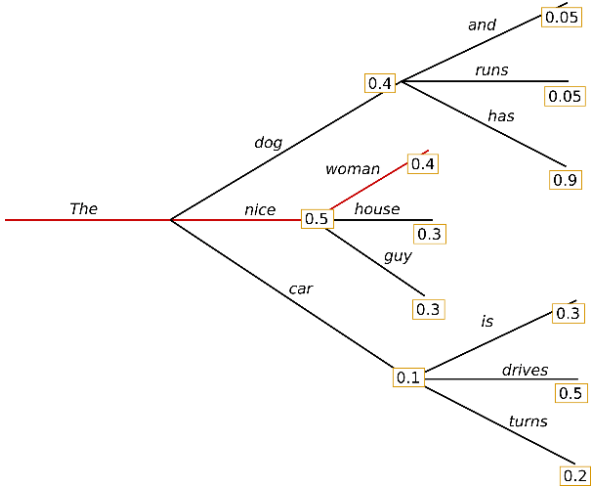

https://huggingface.co/blog/how-to-generate

Cooperative textual generation

- Adding some **constraints** is useful to control various aspects

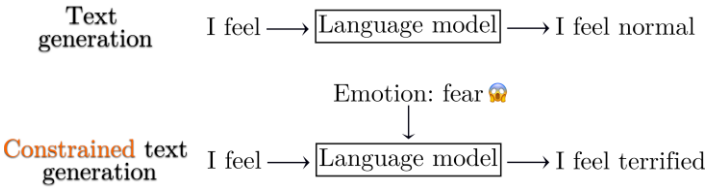

- Discriminators can be trained to detect if a text has the desired property
  - **Real/generated**, writing style, emotion/polarity, ...
- Can be used **to train** the language model (**adversarial** approach)

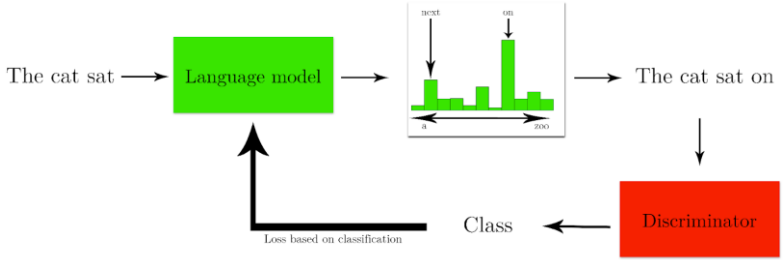

- But also to **guide the decoding (cooperative** approach)

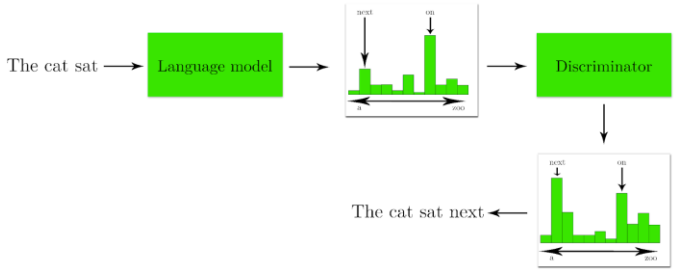

Monte Carlo Tree Search

- Heuristic based iterative algorithm that use randomness to solve deterministic problems with a **too large search space**
- Compromise between **exploiting good sequences** and **exploring promising ones**
- Score of a node is defined by children's (simulation)
  - **Short-term decisions to optimize a long-term result**

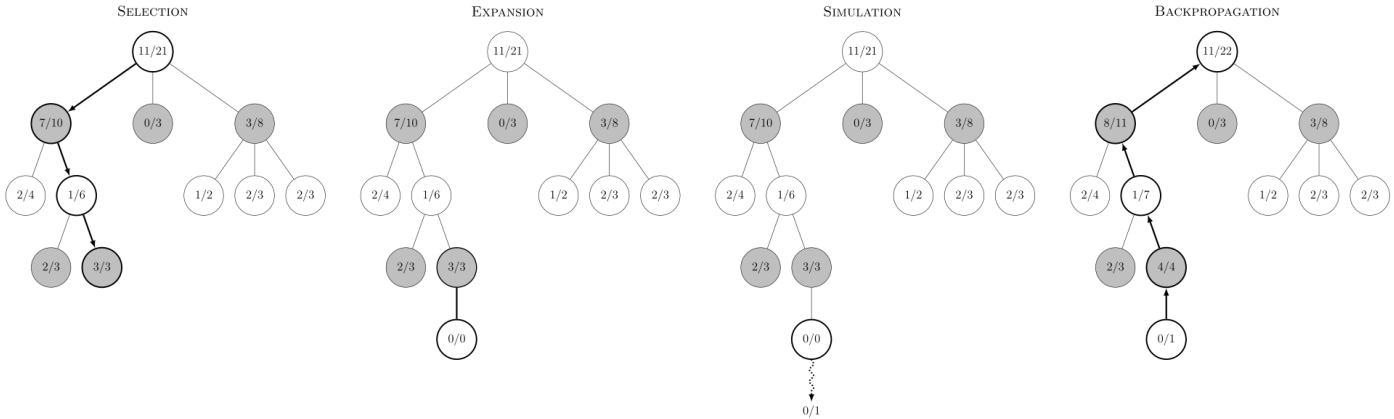

https://en.wikipedia.org/wiki/Monte_Carlo_tree_search

PPL-MCTS

- Monte Carlo Tree Search (MCTS) as a **cooperative decoding strategy** achieves **state-of-the-art results in numerous applications**
  - Take **short-term decisions to optimize a long-term result**, offering a long-view on the generation process
  - Explore promising branches guided by scores of the discriminator

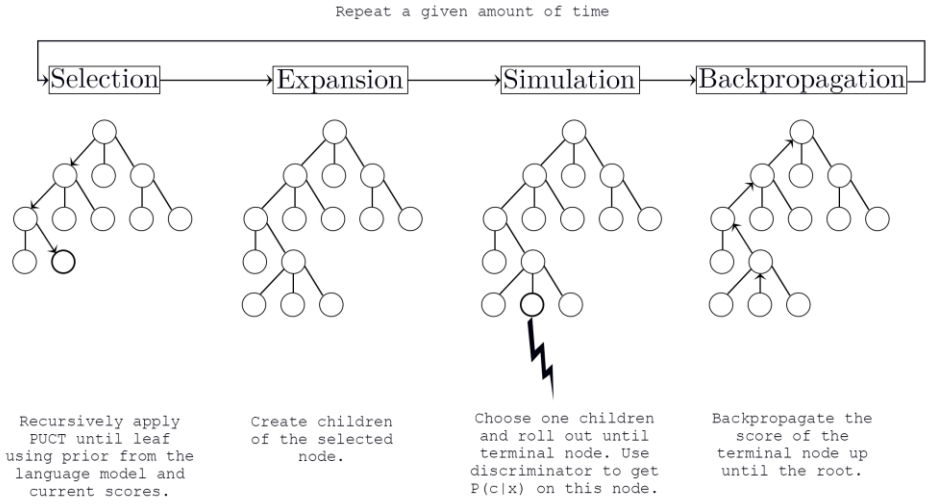

Multiple type of discriminators

- **Bidirectional vs. Unidirectional**

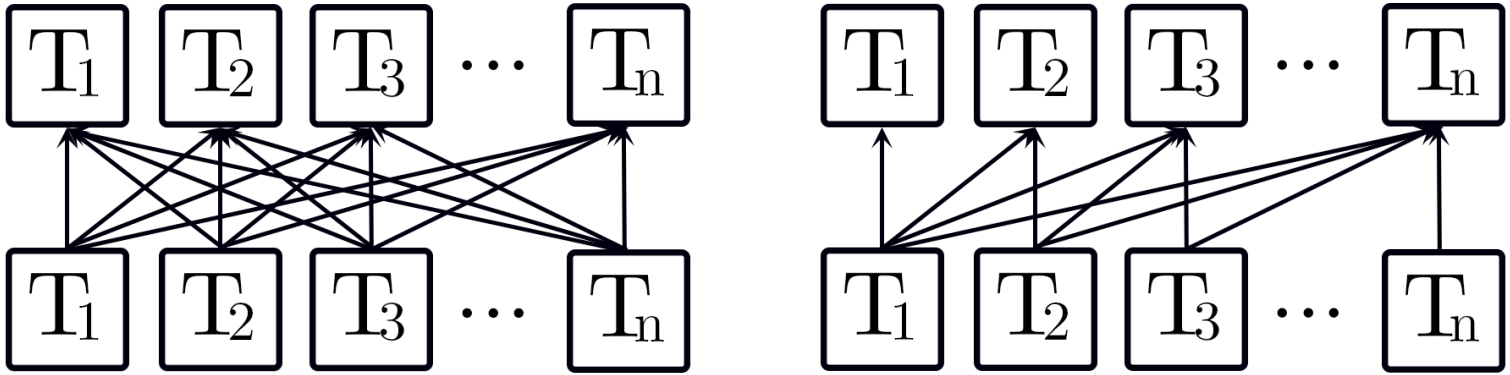

- Unidirectional attention only require to compute attention score on the additional token **(t against t^2 at step t)**

Multiple type of discriminators

- **Generative Discriminators (GeDi)[1]**
  - Leverage **Class-Conditionnal Language Models** to get discrimination scores for the whole vocabulary
  - Get every score in **|C|** forward passes against **|V|** for the standard case (**|V| >> |C|**)

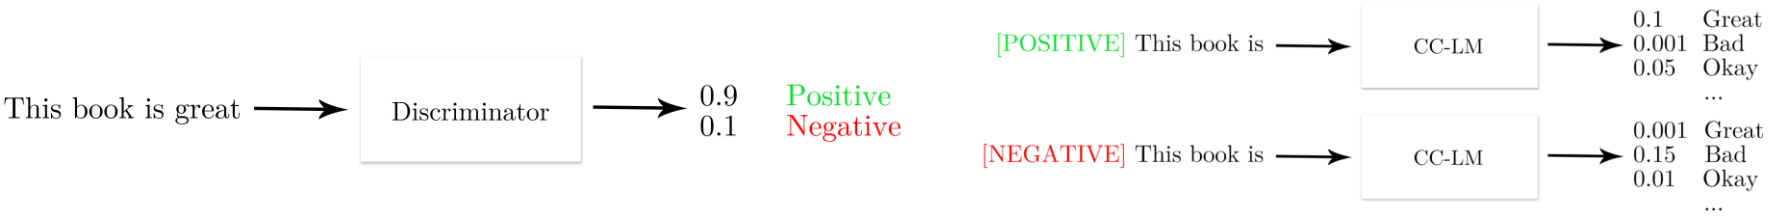

$$\mathrm{P}(\text{positive} \mid \text{This book is great}) \propto \frac{\mathrm{P}(\text{This book is great} \mid \text{positive})}{\mathrm{P}(\text{This book is great} \mid \text{positive}) + \mathrm{P}(\text{This book is great} \mid \text{negative})}$$

[1] Ben Krause, Akhilesh Deepak Gotmare, Bryan McCann, Nitish Shirish Keskar, Shafiq Joty, Richard Socher, Nazneen Fatema Rajani. **GeDi: Generative Discriminator Guided Sequence Generation**

UMR IRISA

Experiments

Datasets

Review polarity 😍😡

amazon_polarity

[POSITIVE] Stuning even for the non-gamer. This sound track was beautiful! It paints the senery in your mind so well I would recomend it even to people who hate vid. game music! I have played the game Chrono Cross but out of all of the games I have ever played it has the best music! It backs away from crude keyboarding and takes a fresher step with grate guitars and soulful orchestras. It would impress anyone who cares to listen! ^_^

News topic 💼🌎🎖️💻

ag_news

[BUSINESS] Carlyle Looks Toward Commercial Aerospace (Reuters) Reuters - Private investment firm Carlyle Group,\which has a reputation for making well-timed and occasionally\controversial plays in the defense industry, has quietly placed\its bets on another part of the market.

Discrimination strength

- Main desired property: **informative output with restricted input** to guide the language model during the generation process

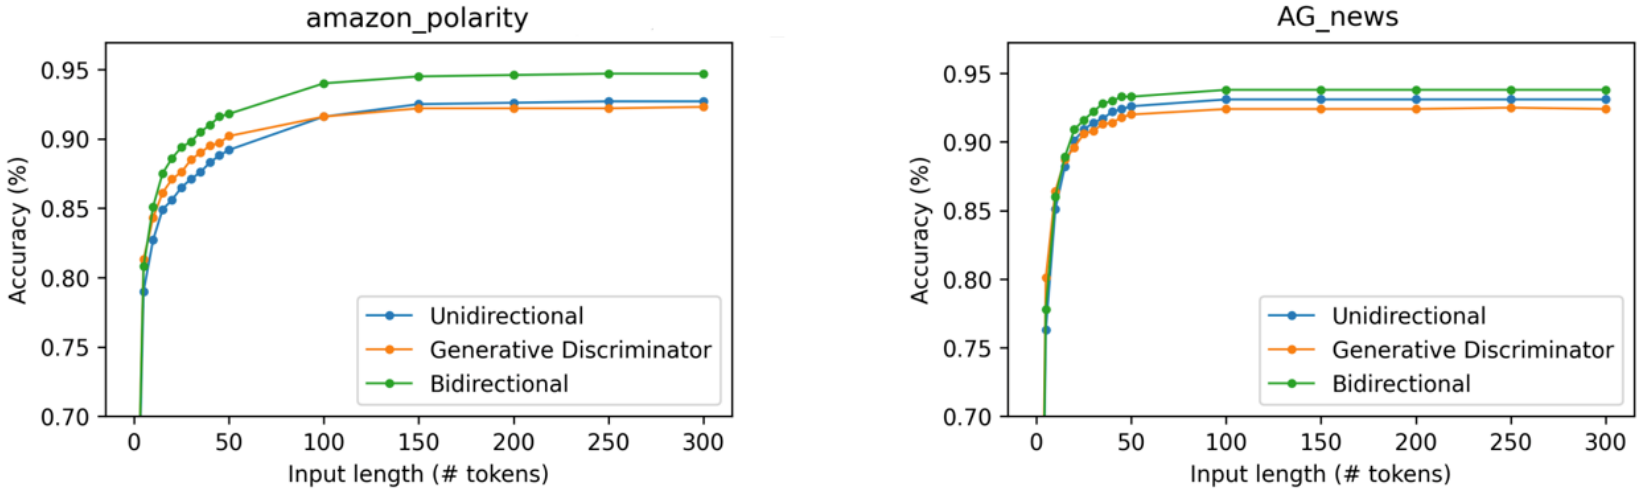

Figure 1: Accuracy (%) of the different type of discriminators w.r.t. the input length (# tokens)

Impact on the generation quality

- Does these small differences in accuracy **reflect on resulting samples**?
- Automatic metrics
  1. **Accuracy**: samples belong to the target class 🎯
  2. **Perplexity**: samples are well written ✍️
  3. **Self-BLEU**: there is enough diversity across samples 📚

	amazon_polarity			AG_news		
Value	Accuracy ↑	5 - Self-BLEU ↓	Oracle perplexity ↓	Accuracy ↑	5 - Self-BLEU ↓	Oracle perplexity ↓
$p(x)$	70.8	0.652	10.49	86.6	0.306	29.08
Bidirectional	96.0*	0.531*	12.25	94.8*	0.319	29.13
Unidirectional	93.0*	0.528*	11.98	93.4	0.313	29.99
Unidirectional (100 its)	93.6*	0.522*	10.73	94.6*	0.323	30.92
Generative discriminator	84.4	0.576	11.92	91.8	0.321	29.43

Table 1: Performance of MCTS w.r.t. the metric to optimize on amazon_polarity (left) and AG_news (right) datasets. ∗ indicates statistically significant improvement against Generative Discriminator. Note that no model demonstrated significant improvement over unidirectional discriminator.

Computational gain

- Cached hidden states allow **linear speed gain**
  - Make cooperative decoding **tractable for long sequences**

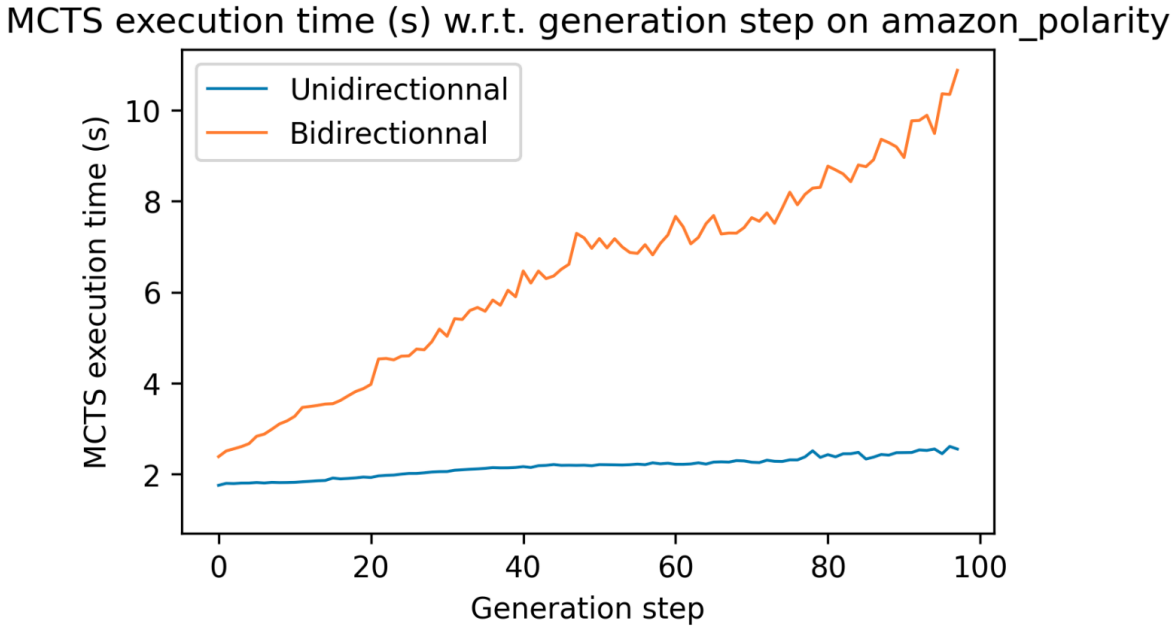

Computational gain

- Exploration is **deeper than wider**
  - Generative discriminators are **more costly for MCTS working points**

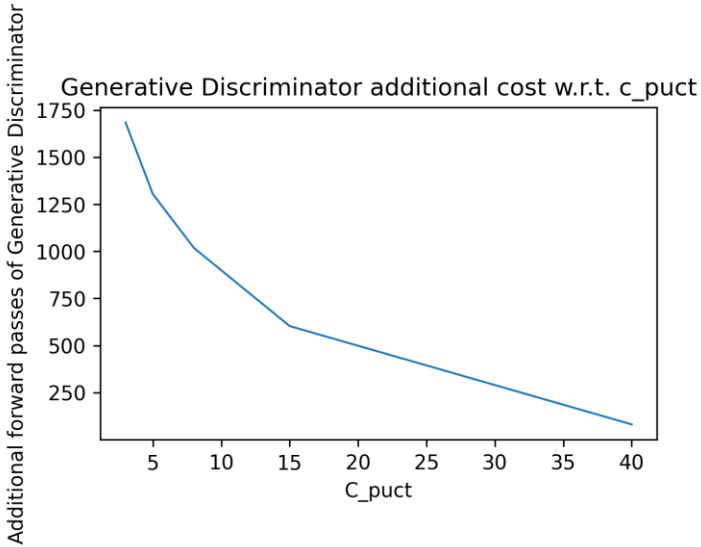

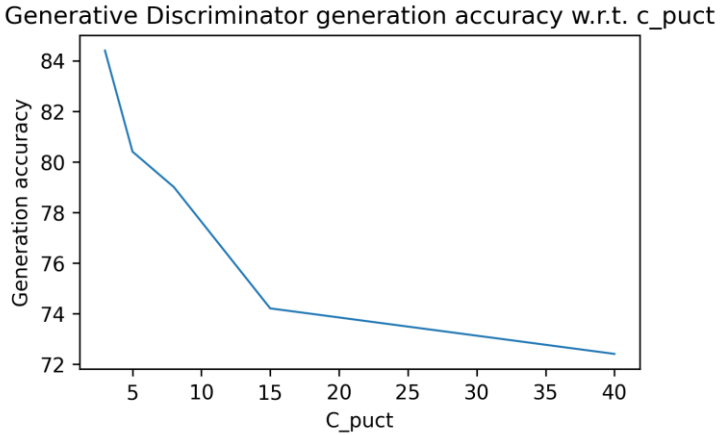

Conclusion

- Standard bidirectional attention discriminators **are justified for accuracy-critical tasks**
- For cooperative generation, unidirectional models produce **very similar results**
  - While offering **an huge speed-up and allowing to scale**
- Generative Discriminators seems interesting at first glance but offers a **less informative signal**
  - Show really useful with **adapted methods that exploit width exploration**
- *« Discriminators »* may **not be transformers**
  - Boolean logic, vocabulary constraint, human evaluation, other heuristics, …
- **Code based on Hugging Face transformer library available on Github**

Thank you for your attention !
Any question ?

antoine.chaffin@irisa.fr @antoine_chaffin

Institut de Recherche en Informatique et Systèmes Aléatoires

Plug And Play Language – Monte Carlo Tree Search

- Heuristic based iterative algorithm that use randomness to solve deterministic problems with a **too large search space**
- Compromise between **exploiting good sequences** and **exploring promising ones**
- Score of a node is defined by children's (simulation): **short-term decisions to optimize a long-term result**