

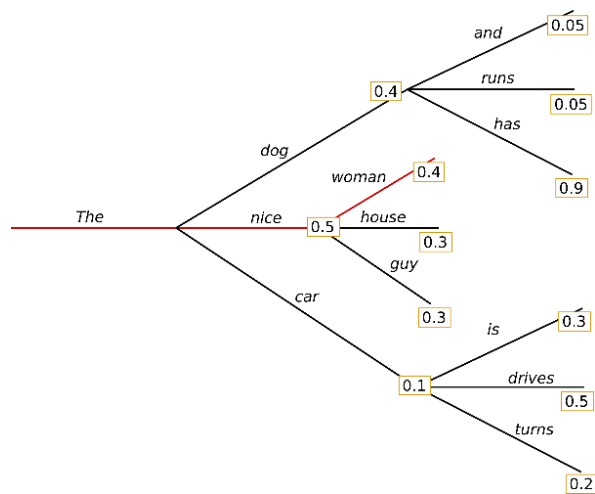
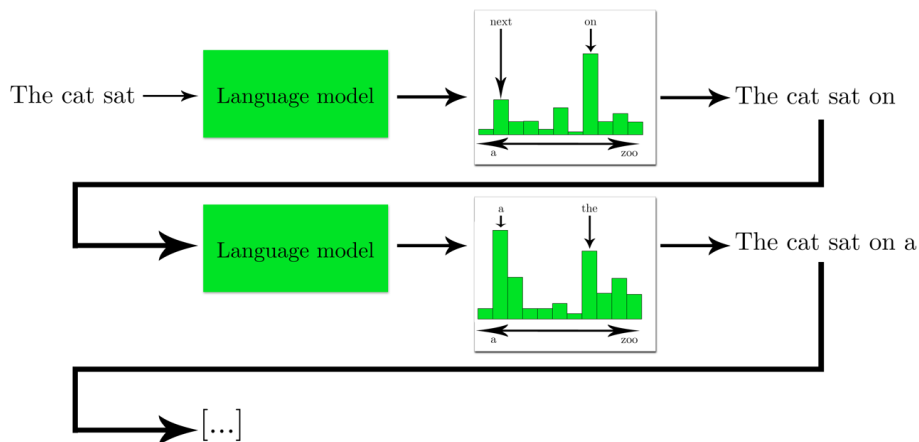
PPL-MCTS: Constrained Textual Generation Through Discriminator-Guided MCTS Decoding

Antoine Chaffin
Vincent Claveau
Ewa Kijak



Introduction

- 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**



- Few options to control the generation besides the **prompt**
- Adding some **constraints** is useful to control various aspects (writing style, emotion/polarity, detoxification, etc.)

Text
generation

I feel → Language model → I feel normal




Emotion: fear 🤖



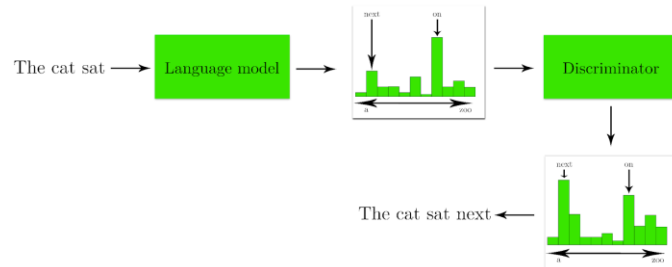
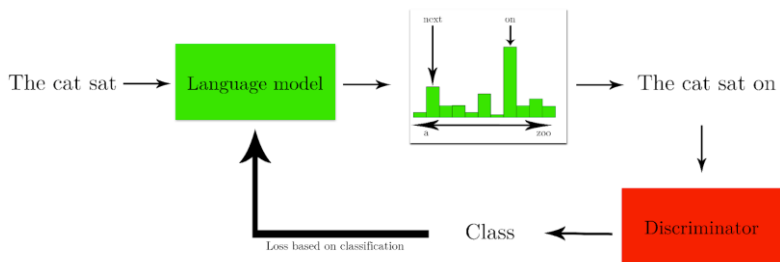
Constrained text
generation

I feel → Language model → I feel terrified

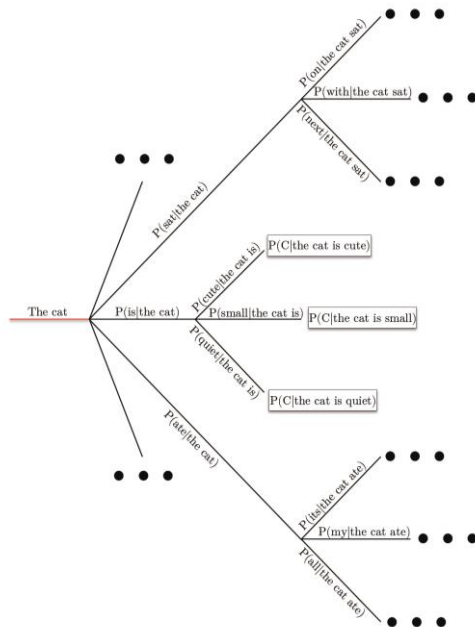
- Language models (LMs) tuning
 - Train and store **one model for each constraint**
 - **Very costly** when even possible for very large LM
- Class-conditional language models
 - Add a **control code** before texts
 - **Training/tuning for any new additional constraint**
- Discriminator-guided generation
 - Change the LM distribution based on **scores from a discriminator**



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

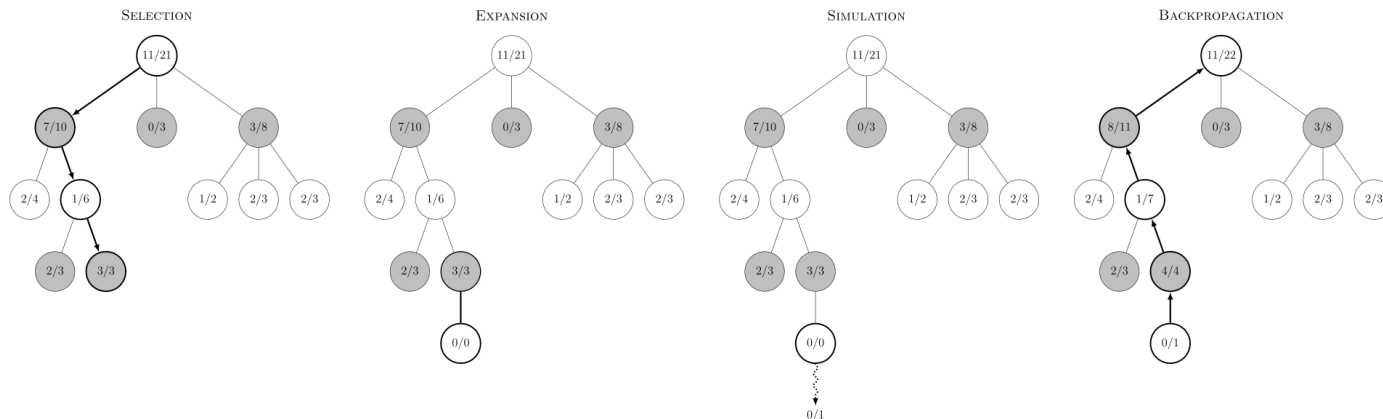


- Generation is **iterative**, one token is produced after the other
 - Previous work focus on the next token to emit (lack of **long-term vision**)
- Words meaning are **context dependent**
 - This book is **awfully**
 - Good/bad ?
 - This made me **cry**
 - Joy/Sadness ?
- **Short term** decisions might not be optimal in the **long run**

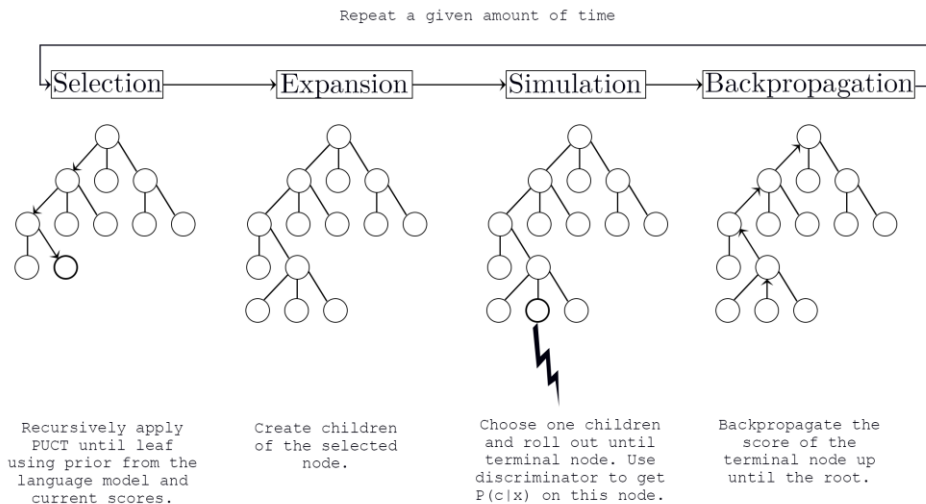


Method

- 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**



- Monte Carlo Tree Search (MCTS) properties:
 - 1. Long-term vision:** scores the next token using finished sequences (rollout)
 - 2. Efficient:** exploration of sub-optimal paths has an upper bound
 - 3. Modular:** outputs a solution according to the computational budget
 - 4. Plug and play:** can be used with any LM and discriminator without any tuning



Experiments

- Two tasks: **polarity** 😍😡 and **emotion** 😡😞😄🙌💖
- Two languages: **French** and **English** 🇫🇷🇬🇧

amazon_polarity

[POSITIVE] Stunning 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! ^_^

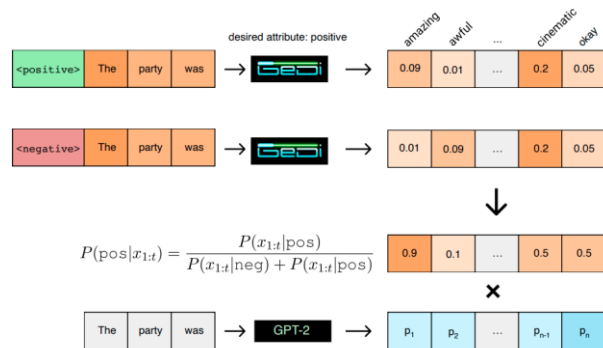

CLS (FLUE)

[POSITIVE] Robert Downey Jr en héros de Marvel? Ca apparait de prime abord complètement improbable et après avoir vu le film on se dit que personne d'autre n'aurait pu jouer le rôle d'Iron Man. En effet le film évite les clichés à la testostérone habituel des supers héros, ajoute une très bonne dose d'humour et de glamour, et propose un scénario original (sauf la fin que je trouve bidon). Au final Iron Man est un très bon film d'action qui parvient à renouveler le genre et où les effets spéciaux supportent très bien l'histoire (et non l'inverse). Bref un très bon moment en perspective à découvrir!

Emotion

[SADNESS] ive been feeling a little burdened lately wasnt sure why that was

- Class-conditional language models (**CC-LMs**)^[1]
- **Re-ranking based** baselines
- **PPLM**^[2]: use discriminator scores to update LM hidden states
- Generative Discriminators (**GeDi**)^[3]: leverage CC-LMs to get classification scores over the whole vocabulary



[1] Nitish Shirish Keskar, Bryan McCann, Lav R. Varshney, Caiming Xiong, Richard Socher. **CTRL: A Conditional Transformer Language Model for Controllable Generation**

[2] Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, Rosanne Liu. **Plug and Play Language Models: A Simple Approach to Controlled Text Generation**

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

- 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 

- PPL-MCTS yields **state-of-the-art results on both tasks and languages**

- Matching GeDi performance that **rely on tuned CC-LMs**



- **Human evaluation (50 reviews per method)** to support automatic metric results
 1. **Polarity:** Rate from 1 to 5 how well the text corresponds to the desired label 
 2. **Readability:** Rate from 1 to 5 how well the text is written 
- **Confirm automatic metrics results**

Generation method	Polarity	Readability
GeDi	4,46 ± 0,08*	4,19 ± 0,28*
PPL-MCTS	4,43 ± 0,12*	4,05 ± 0,23*
PPLM	3,74 ± 0,08	3,12 ± 0,19
Sampling - Argmax	4,00 ± 0,11	2,83 ± 0,33

Table 2: Results of the human evaluation on the CLS dataset (averaged over 3 annotators). * indicates statistically significant ($p \leq 1\%$) improvement against PPLM.

- PPL-MCTS yields **state-of-the-art results** on constrained textual generation **without tuning the language model**
- Powerful **depth exploration** that offers a **long-term view**
 - **Combination** with **Generative Discriminators** for **width exploration**
 - **Adaptative rollout size**
- **Cooperative generation** is very promising to **tweak the LM distribution**
 - Generative Cooperative Networks for Natural Language Generation^[1]
 - Which Discriminator for Cooperative Text Generation?^[2]
- **Code based on Hugging Face transformer library available on Github**

[1] Sylvain Lamprier, Thomas Scialom, Antoine Chaffin, Vincent Claveau, Ewa Kijak, Jacopo Staiano, Benjamin Piwowarski. **Generative Cooperative Networks for Natural Language Generation**

[2] Antoine Chaffin, Thomas Scialom, Sylvain Lamprier, Jacopo Staiano, Benjamin Piwowarski, Ewa Kijak, Vincent Claveau. **Which Discriminator for Cooperative Text Generation?**

**Thank you for your attention !
Any question ?**



antoine.chaffin@irisa.fr  @antoine_chaffin

Institut de Recherche en Informatique et Systèmes Aléatoires

- Fixed number of token in the random simulation rather than until end of sequence
 - Rollout is very useful **up to a given number of tokens**
 - **Adaptative** rollout size: rollout until discriminator score reach a threshold




Figure 4: Accuracy according to the roll-out size; CLS dataset