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

1. Constrained textual generation

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

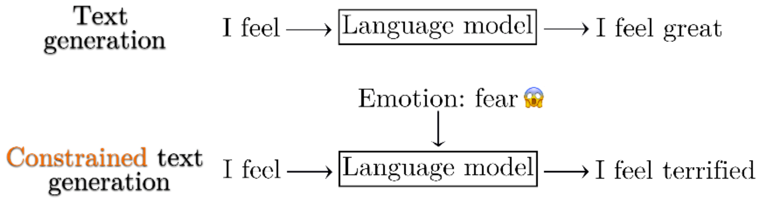

Language models (LM) tuning

- Train and store **one model for each constraint**
- **Very costly** when even possible for very large LM (e.g. GPT-3)

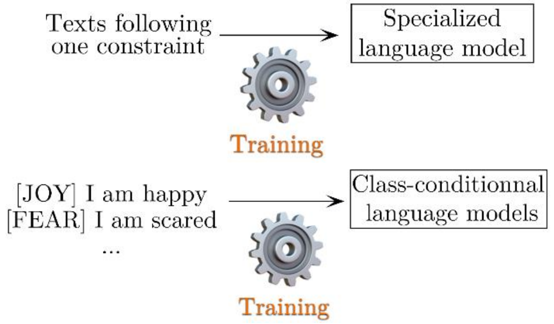

- Class-conditional language models (CC-LMs) [1]
  - Add a control code before texts
  - **Training/tuning for any new additional constraint**

Discriminator-guided generation

- Change the LM distribution based on scores from a discriminator
- Plug and Play Language Models (PPLM) [2]
  - Shift hidden states using discriminator's gradient
  - Require **direct access to the LM** (not compatible with GPT-3 API)
- Generative Discriminator Guided Sequence Generation (GeDi) [3]
  - Exploits **CC-LMs as discriminators** to lower the classification cost

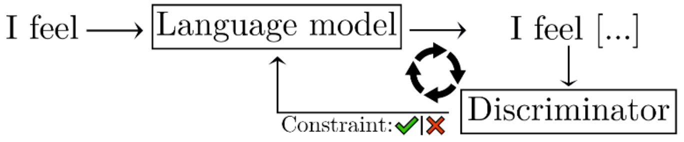

2. PPL-MCTS

- Previous works lack of **long-term vision**
  - Meaning of words are **context depend**
- Short-term decisions to optimize a long-term result
  - Tree exploration similar to game setups

Monte Carlo Tree Search (MCTS)

- Iterative algorithm that finds solutions in a space **too large to be exhaustively searched**
- 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 classifier without any tuning

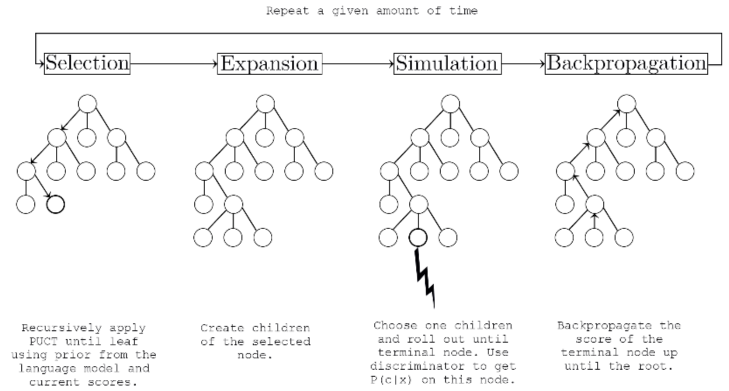

3. Results

- Two tasks: polarity 😍😡 and emotion 😠😥😃😱💏
- Two languages: 🇫🇷🇬🇧
- 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 📕📗📘
- Human evaluation to support automatic metric results
- PPL-MCTS yields **state-of-the-art results on both tasks and languages**
- Rollout is very useful **up to a given number of tokens**

Generation method	amazon_polarity Acc. ↑	5 - Self-BLEU ↓	Oracle pplty ↓	emotion Acc. ↑	5 - Self BLEU ↓	Oracle pplty ↓	CLS Acc. ↑	5 - Self BLEU ↓	Oracle pplty ↓
CC-LM - Classloss	0.82	0.79	**2.56***,†	**0.89***	0.65†	3.72*,†	0.89*	**0.04***,†	50.6
CC-LM	0.91	0.71	3.21†	0.52	**0.13***,†	11.1	0.66	0.06*,†	31.5
GeDi - Classloss	0.96*	0.6*	5.16	0.88*	0.68	5.57*	0.94*	0.4	7.99*
GeDi	0.96*	0.6*	5.16	0.54	0.52†	4.09*,†	0.83*	0.31†	11.9
PPLM	0.89	0.66	2.84†	0.67	0.19†	7.31	0.79	0.23†	8.3
PPL-MCTS	0.97*	0.63*	5.69	0.84*	0.37†	4.82*,†	0.89*	0.54	4.98*,†

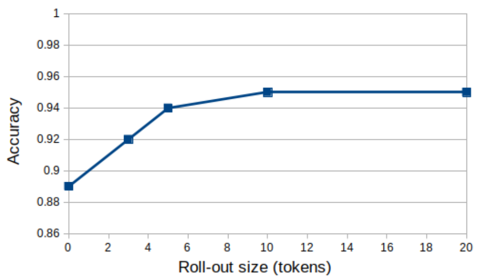

4. Conclusion

- PPL-MCTS shows that **depth search is helpful for constraint generation**
- The extra cost of the classifier **still limit the search in width**
- **Avenues of research:**
  1. Merge GeDi width and PPL-MCTS depth search
  2. Trade-off between accuracy and perplexity
  3. Adaptative rollout size
- Code available on Github

Antoine Chaffin
IMATAG, IRISA, France
Vincent Claveau
CNRS, IRISA, France
Ewa Kijak
Univ. Rennes, IRISA, France

Contacts
antoine.chaffin@irisa.fr
vincent.claveau@irisa.fr
ewa.kijak@irisa.fr

References

[1] **CTRL: A Conditional Transformer Language Model for Controllable Generation.**
Nitish Shirish Keskar, Bryan McCann, Lav R. Varshney, Caiming Xiong, Richard Socher
[2] **Plug and Play Language Models: A Simple Approach to Controlled Text Generation**.
Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, Rosanne Liu
[3] GeDi: Generative Discriminator Guided Sequence Generation.
Ben Krause, Akhilesh Deepak Gotmare, Bryan McCann, Nitish Shirish Keskar, Shafiq Joty, Richard Socher, Nazneen Fatema Rajani