

Exploring the potential of Large Language Models for personalized recommendations

Encadrant-e/Supervisor : Matthieu Labeau, IDS

Nombre d'étudiant-es minimum dans chaque instance de ce projet/ minimum number of students per group : 2

Nombre d'étudiant-es maximum dans chaque instance de ce projet/ maximum number of students per group : 3

Combien d'instances de ce projet proposez-vous ? / How many groups for this project ?
1

Tags : NLP, Large Language Models, Personalization, Recommendation

1 Contexte/Context

Recently, Large Language Models (LLMs) have successively and quickly shown new capabilities, among which great performance at zero-shot and few-shot learning [2], in-context learning, and alignment on user preferences [1]. Generative models are nowadays used with prompts, where tasks can be described as a given textual input x , which the model can use to generate an output text y . Then, a way to specialize the model in doing new tasks is usually to provide an *instruction* x^i to the model; and an expected output (written by a human annotator) y^i . Another possibility which accelerates annotation is to provide the annotator with several possible outputs and ask them to choose a *winner* and *loser*, resulting in a sample (x^i, y_w^i, y_l^i) . While diversity and agreement among annotators when creating data is a difficult issue in the field, we would like here to investigate the potential of a LLM to produce personalized outputs given a set of samples associated to a user u : $S_u = \{(x_u^1, y_u^1), (x_u^2, y_u^2), \dots, (x_u^n, y_u^n)\}$. In particular, we are interested in the capacity of a LLM to use reviews written by a user to predict how they will rate new content. We would like to compare this to on-the-shelf recommender systems.

2 Attendus du projet/Expectations

The first objective of the project is to get familiarized with a Natural Language Processing pipeline. This includes finding appropriate data, defining an experimental methodology and evaluation metrics, and using the appropriate pre-existing python libraries for data pre-processing and implementing the model. The second objective is to be able to measure a *zero-shot* (*i.e.*, with no subsequent training of the model) recommendation performance with the LLM, and to compare it with a simple recommender system. If time allows, we will try to continue training the model (*fine-tune* it) to improve its performance.

References

- [1] Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback, 2022.
- [2] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. 2019.