The Future Remains Unsupervised

Victor Ashioya

Deep Learning Indaba

Abstract

Recent advances in natural language processing have been primarily driven by unsupervised learning from huge datasets. Although transformative, unsupervised models are limited to pattern matching and lack language as a means of achieving goals and goals. Achieving human-level cross-domain language understanding and generation requires research into reinforcement learning as a complement to complex and interactive language tasks.

This position paper argues that unsupervised learning should continue to be a priority in open language learning, but supplemented, where necessary, with reinforcement learning tailored to the task at hand. By designing languages as action sequences and optimizing models to achieve specific goals, RL can harness the power of unsupervised representation learning. It may be uniquely suited to address current limitations around abstract reasoning, grounded language understanding and common-sense knowledge that constrain broader real-world application of models like BERT and GPT-3.

Continued progress in self-supervised learning and world knowledge will enable more sophisticated unsupervised language models, but interactive learning with feedback is key to human-level language competence for open-domain conversation and problem-solving. Researchers should seek integrated modeling approaches, combining unsupervised pre-training, sparse supervision, and RL optimized for purpose, rather than framing progress as a choice between ML paradigms.

With a recognition of the strengths and limits of pattern matching versus interactive optimization, and how they may interact at different levels of language abstraction, NLP can achieve AI systems that understand, generate and reason about language with human-level versatility, open-domain breadth and purposeful skill. The future of human-level NLP is an integrated, interactive one that remains grounded in unsupervised learning but complements it where needed to fulfill each language task.

Key Points

- Unsupervised learning is a key driver of recent advances in NLP, but it has limitations.
- Reinforcement learning can complement unsupervised learning to address these limitations.
- Integrated modeling approaches that combine unsupervised pre-training, sparse supervision, and RL are needed to achieve human-level NLP.
- The future of NLP is an integrated, interactive one that remains grounded in unsupervised learning but complements it where needed.

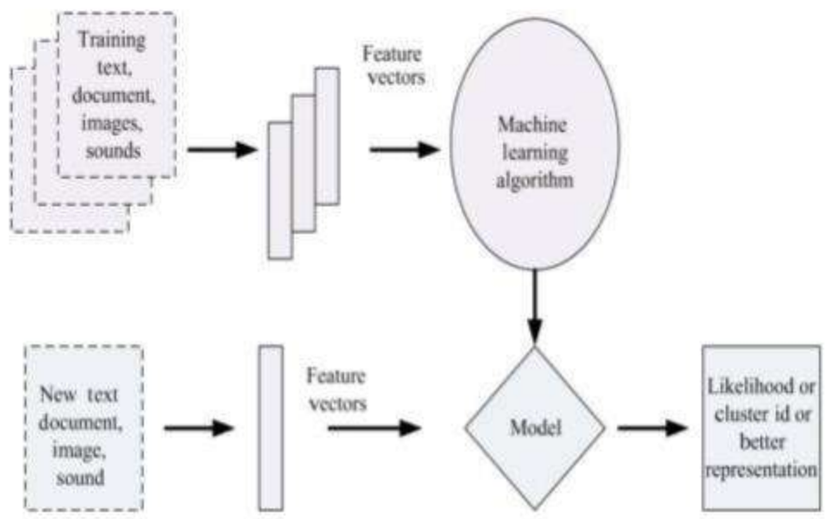

Figure 1: Unsupervised Learning

References

- Unsupervised NLP methods and intuitions behind working with unstructured texts. (2020, October 26). Towards Data Science. Retrieved July 15, 2023, from https://towardsdatascience.com/unsupervised-nlp-methods-and-intuitions-behind-working-with-unstructured-texts-b84aa4d286da
- Unsupervised NLP: How I learned to love the data. (2020, April 28). ODSC. Medium. Retrieved July 15, 2023, from https://odsc.medium.com/unsupervised-nlp-how-i-learned-to-love-the-data-1dde7dc4a3c1
- Unsupervised and supervised NLP approach. (n.d.). Aisera. Retrieved July 15, 2023, from https://aisera.com/blog/unsupervised-and-supervised-nlp-approach/
- 5 reasons why large language models (LLMs) like ChatGPT use reinforcement learning instead of supervised learning for finetuning. (2023, May 5). MarkTechPost. Retrieved July 15, 2023, from https://www.marktechpost.com/2023/03/05/5-reasons-why-large-language-models-llms-like-chatgpt-use-reinforcement-learning-instead-of-supervised-learning-for-finetuning/
- RLHF: Reinforcement Learning for Human-like Language Understanding. (2022, December 9). Hugging Face. Retrieved July 15, 2023, from https://huggingface.co/blog/rlhf
- The rise of the unsupervised learning-based chatbot models. (2020, September 24). SearchUnify. Retrieved July 15, 2023, from https://www.searchunify.com/blog/the-rise-of-the-unsupervised-learning-based-chatbot-models/
- Unsupervised machine learning. (n.d.). In ScienceDirect Topics: Biochemistry, Genetics, and Molecular Biology. *Date accessed :July 15, 2023*

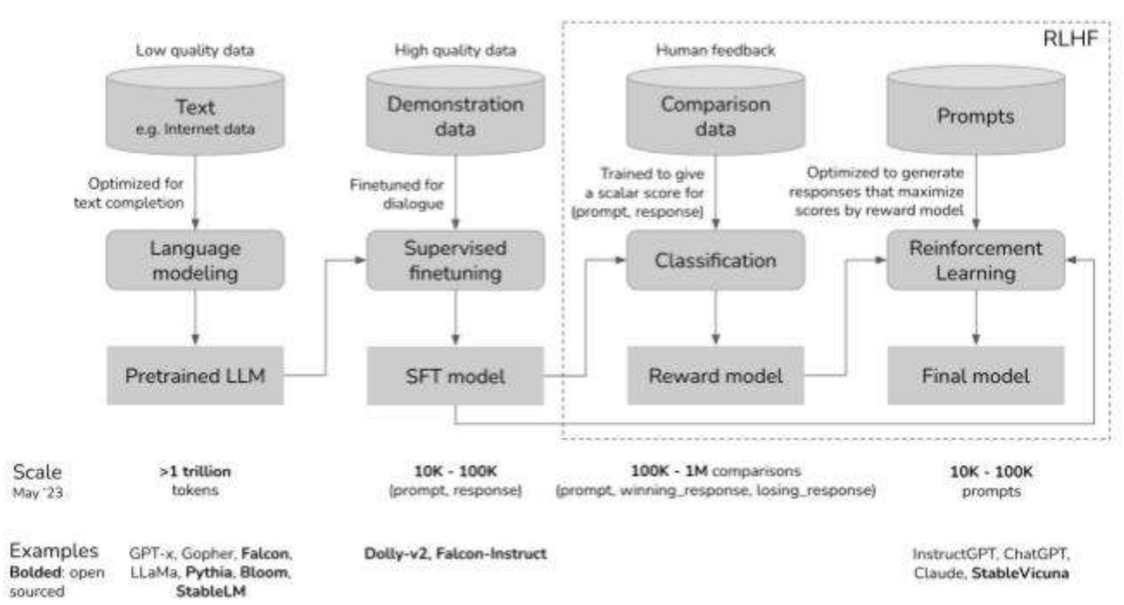

Figure 2: Reinforcement Learning from Human Feedback