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 optimising 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 modelling approaches, combining unsupervised pre-training, sparse supervision, and RL optimised 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 fulfil each language task.

Introduction

Natural Language Processing (NLP) is a subset of Artificial Language (AI) that involves the use of machine learning in language whether be it text summarization, translation or conversational AI. Given its ever growing demand, it is a heavily researched field particularly with the emergence of Large Language Models (LLMs) with researchers trying to come up with newer ways to accommodate the demands by both academia and industry fellows. It is heavily guided by principles of machine learning of supervised and unsupervised learning. Supervised and Unsupervised learning are the building blocks of machine learning. Supervised learning makes use of labelled data while unsupervised learning uses unlabelled data. In NLP, examples of supervised learning could be seen in text classification whereas topic modelling is an example of unsupervised learning.

Traditionally, NLP has been dominated by supervised learning methods. This means that NLP models are trained on labelled data, i.e. data that has been explicitly labelled with the desired output. For example, a supervised machine translation model can be trained on a parallel sentence dataset, where every sentence in one language has been manually translated into the other.

However, supervised learning has certain limitations. First, it requires large amounts of labelled data, which can be expensive and time-consuming to collect. Second, supervised models can only be used to perform clearly labelled tasks. For example, the supervised machine translation model cannot be used to translate a new language pair that has not been seen before.

Unsupervised learning is an alternative to supervised learning. Unsupervised learning does not require labelled data. Instead, they learn from unlabeled data by identifying patterns and relationships in the data. This makes unsupervised learning more scalable than supervised learning as it does not require the collection of large amounts of labelled data. In addition, unsupervised models can be used to perform tasks that are not clearly labelled, such as identifying new language patterns or generating new words.

The Role of Reinforcement Learning in NLP

Reinforcement learning is a type of machine learning that allows agents to learn how to behave in an environment through trial and error. This makes it well-suited to tasks that require abstract reasoning, basic understanding of language, and common sense.

For example, reinforcement learning can be used to train chatbots to have more natural and engaging conversations. In the context of reinforcement learning, chatbots will be rewarded for generating relevant, informative, and engaging responses. Chatbots will also be penalised for generating irrelevant, inaccurate, or offensive responses. Over time, the chatbot will learn to generate more human responses and be more likely to be rewarded.

Reinforcement learning can also be used to train NLP models to perform tasks like machine translation and question answering. In these tasks, the NLP model will be rewarded for producing outputs that are accurate and consistent with the input. The model will also be penalised for producing incorrect or irrelevant output. Over time, the NLP model will learn to produce more accurate and relevant outputs that are more likely to be rewarded.

The Future of NLP: An Integrated Approach

The future of NLP is likely to be an integrated approach that combines unsupervised learning and reinforcement learning. Unsupervised learning can be used to train NLP models on large amounts of unlabeled data, while reinforcement learning can be used to refine these models for specific tasks. . This integrated approach will allow us to build NLP models that are more accurate, flexible, and capable of understanding and generating natural language at the human level.

In addition to unsupervised learning and reinforcement learning, there are other machine learning techniques that can be used to improve NLP. For example, transformation learning can be used to transfer knowledge from one NLP task to another. This can help reduce the amount of labelled data required to train NLP models.

Another promising technique is self-supervised learning. Self-supervised learning methods learn from unlabelled data by designing hypothetical tasks similar to the target task. For example, a self-monitoring language model can be trained to predict the next word in a sentence. This task is similar to the machine translation task, as both tasks require the model to understand the relationship between words. The future of NLP is bright and many exciting research directions are being explored. By combining different machine learning techniques, we are able to create NLP models that are more accurate, flexible, and can understand and generate natural language at the human level.

Conclusion

The future of NLP is likely to be an integrated approach that combines unsupervised learning and reinforcement learning. Unsupervised learning can be used to train NLP models on large amounts of unlabeled data, while reinforcement learning can be used to refine these models for specific tasks. . This integrated approach will allow us to build NLP models that are more accurate, flexible, and capable of understanding and generating natural language at the human level.

In addition to unsupervised learning and reinforcement learning, there are other machine learning techniques that can be used to improve NLP. For example, transformation learning can be used to transfer knowledge from one NLP task to another. This can help reduce the amount of labeled data required to train NLP models.

Another promising technique is self-monitoring learning. Self-supervised learning methods learn from unlabelled data by designing hypothetical tasks similar to the target task. For example, a self-monitoring language model can be trained to predict the next word in a sentence. This task is similar to the machine translation task, as both tasks require the model to understand the relationship between words. The future of NLP is bright and many exciting research directions are being explored. By combining different machine learning techniques, we are able to create NLP models that are more accurate, flexible, and can understand and generate natural language at the human level.

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*
- Is NLP supervised or unsupervised? (n.d.). Quora. Retrieved August 9, 2023, from https://www.quora.com/Is-NLP-supervised-or-unsupervised/
- IBM Cloud Education (2022, March 10). Supervised vs unsupervised learning: What's the difference? IBM Cloud Blog. https://www.ibm.com/blog/supervised-vs-unsupervised-learning/
- Supervised and Unsupervised Learning. (n.d.). JavaTpoint. Retrieved August 9, 2023, from https://www.javatpoint.com/difference-between-supervised-and-unsupervised-learning/
- Supervised and Unsupervised Learning. (n.d.). JavaTpoint. Retrieved August 9, 2023, from https://www.javatpoint.com/difference-between-supervised-and-unsupervised-learning/
- Bhatt, S. (2018, March 19). Reinforcement learning 101. Towards Data Science. https://towardsdatascience.com/reinforcement-learning-101-e24b50e1d292/