Procesamiento de Lenguaje Natural: Embeddings y Modelos Generativos

Dr. Luciano Del Corro (Profesor Adjunto)

Programa:

El hilo conductor de la materia es recorrer la evolución de los modelos de lenguaje, culminando con GPT-4 y los avances del último tiempo.
La intención es brindar un panorama completo del estado del arte de los modelos de lenguaje tanto para NLU como NLG, combinando teoría con aplicaciones prácticas y explorando las posibilidades y desafíos que se nos presentan actualmente.

Temario:

- Unidad 1: Introducción a *Deep Learning* para NLP
- Unidad 2: Embeddings: codificando el lenguaje
- Unidad 3: Introducción a Modelos de Lenguaje: Ngrams y primeros modelos neuronales
- Unidad 4: ELMo: LSTMs y Redes Neuronales Recurrentes
- Unidad 5: BERT: Atención, *Transformers*, *Encoders* y *Decoder*
- Unidad 6: GPT: *Pre-training* y modelos Generativos
- Unidad 7: ChatGPT: RLHF, DPO e Instruction Tuning
- Unidad 8: *Information Retrieval* y RAG
- Unidad 9: *Model Distillation* y *Specialization*
- Unidad 10: Modelos Multimodales
- Unidad 11: Agentes
- Unidad 12: Prompt Engineering y NLP en la industria
- Unidad 13: State of the Art y ¿ahora qué? *Reward Modelling*, *Model Adaptation*, *Distillation*

Bibliografía:

- A Primer on Neural Network Models for Natural Language Processing, Goldberg, [1510.00726] A Primer on Neural Network Models for Natural Language Processing (arxiv.org)
- Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems, Aurélien Géron, 3ra edición
- Dive into Deep Learning, https://d2l.ai
- Deep Learning, Goodfellow, et. al, Deep Learning (deeplearningbook.org)
- Mikolov et. al., Efficient Estimation of Word Representations in Vector Space, 2013.
- Pennington et. al., GloVe: Global Vectors for Word Representation, 2014. (GloVe: Global Vectors for Word Representation (stanford.edu))
- Bojanowski et. al., Enriching Word Vectors with Subword Information, 2017.
- Schnabel et. al., Evaluation methods for unsupervised word embeddings, 2015.

- Speech and Language Processing. Daniel Jurafsky & James H. Martin, Capítulo 3 (https://web.stanford.edu/~jurafsky/slp3/ed3book_jan72023.pdf)
- Speech and Language Processing. Daniel Jurafsky & James H. Martin, Capítulo 7
- Neural net language models, Yoshua Bengio, 2008
- Deep contextualized word representations, 2018
- Contextual String Embeddings for Sequence Labeling, Akbik, 2018
- Neural Architectures for Named Entity Recognition, Lampe et al. 2016
- Stanford Neural Machine Translation Systems for Spoken Language Domains, Luong et al., 2015
- A Primer on Neural Network Models for Natural Language Processing, Goldberg, 2015
- Hierarchical Attention Networks for Document Classification, Yang et al., 2016
- Attention Is All You Need, Vaswani et. al. 2018
- BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding et. al., 2019
- RoBERTa: A Robustly Optimized BERT Pretraining Approach Liu et. al, 2019
- Improving Language Understanding by Generative Pre-Training, Radford et. al., 2019
- DeBERTa: Decoding-enhanced BERT with Disentangled Attention, He et. al, 2020
- ELECTRA: Pre-training Text Encoders as Discriminators Rather Than Generators, Clark et. al. 2020
- Parameter-Efficient Transfer Learning for NLP, Housby et al., 2019
- LoRA: Low-Rank Adaptation of Large Language Models, Hu et. al., 2021
- Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer, Raffel et al., 2020
- PaLM: Scaling Language Modeling with Pathways, Chowdhery et. al., 2022
- PaLM 2 Technical Report , 2023
- Language Models are Unsupervised Multitask Learners, Radford et. al.,2019
- Language Models are Few-Shot Learners, Brown wt. Al., 2020
- Training language models to follow instructions with human feedback, Ouyang et. al., 2022
- The Flan Collection: Designing Data and Methods for Effective Instruction Tuning, Longpre et. al., 2023
- Instruction Tuning with GPT-4, Peng et. al, 2023
- Finetuned Language Models are Zero-shot Learners, Wei et. al., 2022
- Llama 2: Open Foundation and Fine-Tuned Chat Models, Touvron et. al, 2023
- Direct Preference Optimization: Your Language Model is Secretly a Reward Model, Rafailov et. al, 2023
- GPT-4 Technical Report, 2023, OpenAI
- Sparks of Artificial General Intelligence: Early experiments with GPT-4, Bubeck et. al., 2023