Master internship in Natural Language Processing and Machine Learning: Text generation with disentangled semantic and syntactic representations

Caio Corro

- **Duration:** 5-6 months, during the year 2020
- **Location:** LIMSI, Orsay (south of Paris)
- **Supervisor:** Caio Corro
- **Team:** Spoken Language Processing / Traitement Automatique de la Parole
- **Contact:** caio.corro@limsi.fr

1 Context

This internship will focus on text generation with deep generative models, in particular Variational Auto-Encoders (VAEs) [1, 2]. The goal is to study how we can build a generative model for text generation where the semantic and syntactic representations are disentangled [3]. That is, we aim to generate a sentence through the following process:

- sample z: a latent variable encoding a meaning,
- sample z': a latent variable encoding a surface structure (i.e. how the meaning is expressed),
- sample x from $p(x|z, z')$: a sentence conditioned on its meaning and syntactic structure.

This kind of models could be used for sentence simplification, paraphrasing or generating diverse text responses [4, 5]. Previous work in the literature has explored models where z' is encoded as a discrete combinatorial structure [6, 7]. However, these methods require annotation of linguistic structures to be available during training and they may not be suitable for large scale learning as they are computationally expensive.

Therefore, we aim to focus on techniques closer to the ones developed in computer vision where both the semantic and syntactic representations are encoded in a fixed size continuous latent space and learned in a fully unsupervised setting. To this end, the successful candidate will explore generative losses for disentangled representation learning and propose neural architectures specifically developed for text generation in this setting.

2 Missions

- review the literature on learning disentangled latent space with VAEs;
- reproduce the experiments from [3, 8] with a transformer architecture instead of recurrent networks;
- explore VAE losses for learning disantangled representations;
- propose transformer architectures that isolates structural information from semantic information (e.g. distance, see Section 3 in [9]);

References

[1] Diederik P Kingma and Max Welling. *Auto-Encoding Variational Bayes.*

[2] Danilo Jimenez Rezende et al. *Stochastic backpropagation and approximate inference in deep generative models.*

[3] Mingda Chen et al. *A Multi-Task Approach for Disentangling Syntax and Semantics in Sentence Representations.*

[4] Yizhe Zhang et al. *Generating Informative and Diverse Conversational Responses via Adversarial Information Maximization.*

[5] Danial Alihosseini et al. *Jointly Measuring Diversity and Quality in Text Generation Models.*

[6] Pengcheng Yin et al. *StructVAE: Tree-structured Latent Variable Models for Semi-supervised Semantic Parsing.*

[7] Caio Corro and Ivan Titov *Differentiable Perturb-and-Parse: Semi-Supervised Parsing with a Structured Variational Autoencoder*.

[8] Tom Pelsmaeker and Wilker Aziz *Effective Estimation of Deep Generative Language Models*.

[9] Nikita Kitaev and Dan Klein. *Constituency Parsing with a Self-Attentive Encoder*.