EURA NOVA

FACULTIES OF SCIENCE AND ENGINEERING

Master Thesis & PFE Offers 2021-2022

For the latest information, please visit https://research.euranova.eu.
20th November 2020

CONTENT

Eura Nova

Introduction

Eura Nova is a data-driven Belgian company founded in September 2008 and located in Brussels, Marseille, and Tunis. Our mission is simple: bring life to our customers' great ideas, by offering best-in-class services in data science, software engineering, and data architecture. To do so, we invest significantly in in-house expertise and state-of-the-art knowledge. In line with this course of action, we offer academic programs in collaboration with universities. These offers include bootcamps, master theses topics, research internships, and PhDs topics. See below for details.

Our Master Theses and PFE Offers

This document presents master theses and graduation project topics supervised by our research & development department. Each project is an opportunity to be actively involved in the development of solutions to address tomorrow's challenges in ICTs and to implement them today. The students will work in a dedicated international team of engineers with diverse expertise in machine learning, graph theory, artificial intelligence, high performance computing, etc. They will keep Eura Nova informed of the project advancement and share their ideas and challenges using the in-house knowledge management tool. We value continuous learning and teamwork. We love to have a good time together. For more information on our R&D activities, please visit our website at https://research.euranova.eu.

How To Apply

When you have gone through our master thesis offers, pick your favourite. Draft a short text, stating why you find it interesting and what you would do about it. Send us this statement, along with your CV at career@euranova.eu. If you are interested in working on a topic that is not in our range of offers, we would be delighted to hear your proposition and invite you to get in touch as well.

GridNet with Attention for Semantic Segmentation

Context

Image semantic segmentation is a popular field of research which consists in assigning a semantic label to each pixel of an image. Many methods have been proposed to tackle this problem, including GridNet [1], a little-known architecture. GridNet uses a grid-based approach where each row of grid processes features maps at different resolutions, and each column allows communication between lines. Due to its specific architecture, GridNet cannot use a pre-trained network and has to be trained from scratch. Despite this, it shows promising results.

Recently, a new technique, named Attention [2, 3], was developed by the deep learning community. Attention initially gave remarkable results in Natural Language Processing [4] and was soon adapted to image processing [5, 6]. The attention mechanism is based on a system of queries, keys and values, which allows the network to pay more attention to different input parts that may be far apart.

The objective of this thesis is to explore how the attention mechanism can be integrated into a GridNet. Especially, can the attention mechanism allow better communication between the resolutions (lines) of a GridNet? Can the attention mechanism capture better contextual information? And finally, can it facilitate the training? We expect that the use of the Attention mechanism will allow a GridNet to better capture contextual information, which will improve the long-range dependency and may also reduce the complexity of the model by reducing the training time and inference time. The study will be done on public datasets such as Cityscapes [7] and results will be compared with state-of-the-art methods [8].

Business opportunities

The exploration of the Attention mechanism carried out as part of this project will provide a good understanding of a very popular and cutting-edge technique. It is important for Eura Nova to know and master state-of-the-art techniques. In addition, it will set up a codebase useful for tasks related to image segmentation such as semantic segmentation, instances segmentation and panoptic segmentation.

Contribution

The objective is to implement and improve a GridNet using an attention mechanism. It will take a strong understanding of GridNet and Attention mechanisms as well as good creativity to combine them.

- Explore the state of the art of Attention for image processing.
- Read and understand GridNet architecture.
- Propose different approaches to use attention in a GridNet and implement them.
- Analyze the different results and explain the pros and cons.

Technologies/expertise to develop

- Keras and Tensorflow
- Semantic segmentation
- Attention mechanism

In practice

Belgium, early 2021 (5-6 months)

References

- [1] Fourure, D., Emonet, R., Fromont, E., Muselet, D., Tremeau, A., & Wolf, C. (2017). Residual conv-deconv grid network for semantic segmentation. arXiv preprint arXiv:1707.07958.
- [2] Bahdanau, D., Cho, K., & Bengio, Y. (2014). Neural machine translation by jointly learning to align and translate. arXiv preprint arXiv:1409.0473.
- [3] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. In Advances in neural information processing systems (pp. 5998-6008).
- [4] Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., ... & Agarwal, S. (2020). Language models are few-shot learners. arXiv preprint arXiv:2005.14165.
- [5] Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., & Zagoruyko, S. (2020). End-to-End Object Detection with Transformers. arXiv preprint arXiv:2005.12872.
- [6] Under double-blind review. LambdaNetworks: Modeling Long-Range Interactions Without Attention
- [7] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The Cityscapes Dataset for Semantic Urban Scene Understanding," in Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016.
- [8] Tao, A., Sapra, K., & Catanzaro, B. (2020). Hierarchical Multi-Scale Attention for Semantic Segmentation. arXiv preprint arXiv:2005.10821.

Combining Generative Adversarial Networks and Attention Mechanism for Multi-View Representation Learning

Context

In the real world, multi-view or multi-modal data are quite common. Indeed, a sample can have different representations depending on its source, on the sensor that captured it, or even on the applied method generating its features. For instance, if we consider social media content or e-commerce websites, most of the items are represented by an image and by some description text. A second example can be related to healthcare applications where patients can be described by their symptoms, the results from tests, and data coming from sensors monitoring their vitals.

Several techniques have been used to handle multi-modal or multi-view datasets. The objective is to learn the relationships between the different views that refer to the same item and embed them in a common latent space. The learnt common latent manifold can then be used to solve several machine learning tasks such as clustering, classification, etc.

Recently, generative adversarial networks (GANs) [1] have gained popularity in the field of multi-modal embedding thanks to their limited need for input data, their ability to generate fake data that are very similar to the real ones, as well as their ability to handle missing data or partial views. For instance, in [2], high-level features extracted from each view are embedded and fused to give a unique representation of the given sample. The proposed framework called GP-MVC is even able to handle missing views by the means of the cycle consistency between the generators.

The main challenge with the fusion approach in the context of multi-view representation learning is to soft-select relevant features of the views by assigning learnable weights. If we consider the context of action recognition using RGB and Depth images, some actions can be more detectable within the first sensor while others may appear clearer with the second type of images [4].

The objective of this thesis is to explore how the attention mechanism [4] can lead to discovering meaningful weights/scores and provide complementary information and thus enhance multi-view representation as shown in [5] [6] [7]. The main idea is to integrate one or more attention networks within a GAN-based architecture to better represent multimodal/view data.

Business opportunities

The exploration of Attention mechanism carried out as part of this project will provide a good understanding of a very popular and cutting-edge technology. This thesis might lead to a base implementation that would be used for several tasks related to multimodal/view data handling for machine learning applications.

Contribution

The objective is to integrate the attention mechanism within a GAN-based architecture for multiview representation learning.

- Explore the state of the art of multimodal/view representation learning.
- Explore the state of the art of Attention and self-attention mechanisms.
- Read and understand the GP-MVC architecture and code.
- Propose different approaches to integrate attention networks with GP-MVC and implement them.
- Conduct a comparative analysis of the results.

Technologies/expertise to develop

- Encoders/decoders, GANs.
- Tensorflow, Keras and Pytorch.
- Multiview representation learning.
- Attention mechanism.

In practice

Tunis, early 2021 (5-6 months)

References

- [1] Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A. C., and Bengio, Y. Generative adversarial nets. In NIPS, 2014
- [2] Wang, Q., Ding, Z., Tao, Z., Gao, Q., & Fu, Y. (2020). Generative Partial Multi-View Clustering. arXiv preprint arXiv:2003.13088
- [3] Wang, L., Ding, Z., Tao, Z., Liu, Y., & Fu, Y. (2019). Generative multi-view human action recognition. In Proceedings of the IEEE International Conference on Computer Vision (pp. 6212-6221)
- [4] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. In Advances in neural information processing systems (pp. 5998-6008).
- [5] Yuan, Y., Xun, G., Ma, F., Wang, Y., Du, N., Jia, K., ... & Zhang, A. (2018, November). Muvan: A multi-view attention network for multivariate temporal data. In 2018 IEEE International Conference on Data Mining (ICDM) (pp. 717-726). IEEE.
- [6] Liu, J. W., Ding, X. H., Lu, R. K., & Xionglin, L. U. O. (2020, August). Self-attention Multi-view Representation Learning with Diversity-promoting Complementarity. In 2020 Chinese Control And Decision Conference (CCDC) (pp. 3972-3978). IEEE
- [7] Li, X., Wang, C., Tan, J., Zeng, X., Ou, D., Ou, D., & Zheng, B. (2020, April). Adversarial Multimodal Representation Learning for Click-Through Rate Prediction. In Proceedings of The Web Conference 2020 (pp. 827-836).

Implementing a learning agent to optimize Big Data infrastructures

Context

Cloud platforms provide two main advantages to end-users and cloud operators: flexibility and cost efficiency. They allow the users, from one side, to pay only for the resources used and the cloud providers, from the other side, to share their resources between multiple users and workloads in a way that allows them to reduce cost.

However, a bad utilization of the cloud facilities can greatly affect cost effectiveness. For example, a recent study [1] showed that, for a large production cluster at Twitter, managed with Mesos, the gap between resource utilization and resource reservation for the aggregate CPU and memory capacity is quite wide. The achieved utilization is 25-35% and 40% for aggregate CPU and memory respectively, while reserved resources exceed 75% and 60% of available capacity for CPU and memory respectively.

Another study [2] showed that server utilization on Amazon EC2 is estimated between 3.6% and 16.9% only if it's not subject to optimization.

Properly identifying the best configuration for different applications with low search cost is a challenging task for many reasons. First, the relationship between the resources and the execution time is not linear, so increasing resources does not always increase the performance of the application. However, it can increase the execution cost. Second, the best cloud configuration depends on the type of workload or the resource requirements (CPU-bound, memory-bound, disk-bound) and its internal behaviour.

When using a cloud service, it is important to ensure the performance is at its best while reducing the related cost, and energy consumption. Different tuning methods have been proposed in the literature. They are classified into five main groups: rule-based, simulation-based, cost-based, search-based and learning-based. The most recent ones are the search-based and learning-based methods. The former makes use of search algorithms to find the best configuration [3][4], while the latter mainly relies on the machine learning algorithms to find the best configuration [5][6].

We are interested in the implementation of an intelligent agent that is capable of scaling up and down the resources in a way to match the need of the big data application in use. The agent should also be able to dynamically adjust to a different workload.

Business opportunities

Optimizing the cloud resources is a common need among cloud users. In a pay-as-you-go billing model, it is important for organizations to pay for the performance they need in accordance with the applications in use. Implementing such an agent can be a good opportunity for Eura Nova, whose goal is to accompany its clients in their digital transformation journey.

Contribution

The internship will focus on exploring the reinforcement learning techniques to achieve the performance optimization goals. In this internship, the student will:

- Explore the state of the art regarding the automatic tuning of cloud configurations for big data applications.
- Gather the requirements and constraints to address
- Perform the RL formulation of the problem
- Select the best family of RL methods to implement the agent
- Implement selected algorithm(s) and perform the required experiments to prove the efficiency of the agent

The results of this internship will be shared with the community as a talk, a workshop, an article published in a journal or conference, a blog post, or a contribution to an open-source project.

Technologies/expertise to develop

- Understanding of the big data environment and the relationship between the application's performance and the allocated resources
- Machine learning algorithms and Gym for the RL environment
- Docker and Kubernetes for testing and experiments
- Gitlab/git: for development and versioning
- If needed, deep learning models are developed using PyTorch or TensorFlow

In practice

Tunis, early 2021 (5-6 months)

References

[1] Christina Delimitrou and Christos Kozyrakis. Quasar: Resource-Efficient and QoS-Aware Cluster Management. SIGARCH Comput. Archit. News, 42(1):127–144, feb 2014.

[2] H. Liu. A Measurement Study of Server Utilization in Public Clouds. In 2011 IEEE Ninth International Conference on Dependable, Autonomic and Secure Computing, pages 435–442, dec 2011.

[3] Omid Alipourfard, Hongqiang Harry Liu, Jianshu Chen, Shivaram Venkataraman, Minlan Yu, and Ming Zhang. CherryPick: Adaptively Unearthing the Best Cloud Configurations for Big Data Analytics. In And and Aditya Akella Jon Howell, editors, 14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17), pages 469–482. USENIX Association, 2017.

[4] Yuqing Zhu, Jianxun Liu, Mengying Guo, Yungang Bao, Wenlong Ma, Zhuoyue Liu, Kunpeng Song, and Yingchun Yang. BestConfig: Tapping the performance potential of systems via automatic configuration tuning. In SoCC 2017 - Proceedings of the 2017 Symposium on Cloud Computing, pages 338–350. Association for Computing Machinery, Inc, sep 2017.

[5] Constantinos Bitsakos, Ioannis Konstantinou, and Nectarios Koziris. DERP: A deep reinforcement learning cloud system for elastic resource provisioning. In Proceedings of the International Conference on Cloud Computing Technology and Science, CloudCom, volume 2018-Decem, pages 21–29, 2018.

[6] Zixia Liu, Hong Zhang, Bingbing Rao, and Liqiang Wang. A Reinforcement Learning Based Resource Management Approach for Time-critical Workloads in Distributed Computing environment. Proceedings - 2018 IEEE International Conference on Big Data, Big Data 2018, pages 252–261, 2019.

Completing Missing data for Multimodal Recommendation Engines

Context

Recommendation systems have become a critical building block of current information systems, particularly in industries such as e-commerce and social network platforms. These systems have steadily evolved over the last years by relying on increasingly accurate models and capturing more complex multimodal data.

However, current recommendation systems suffer from data sparsity, in both the interaction data and the content data. In the case of multimodal data, some modalities might be incomplete for a subset of items which prevents an accurate recommendation for those items.

To address the sparsity issue, a plethora of models were proposed. In particular interest for this project are the Generative Adversarial Networks (GAN) [1]. GANs have recently gained popularity in the recommendation field as they can contribute to the generation of missing data[2]. The generated missing data will then be used by a multimodal recommendation engine to derive more accurate recommendations[3].

Business opportunities

RecoGAN is a multimodal recommendation engine that integrates different data types to offer personalized items recommendation to users. RecoGAN solves the data sparsity problem, and therefore benefits a wide range of recommendation applications such as e-commerce, retail and streaming platforms and enables them to improve their user experience and make more targeted and personalized recommendations.

Contribution

The purpose of this project is to integrate the generation of the missing modalities in RecoGAN, a multimodal recommendation engine based on autoencoders and generative adversarial networks developed at Eura Nova, to improve the accuracy of predictions.

In this project, you will:

- Study the state of the art in generating the missing data and modalities.
- Study the state-of-the-art variational autoencoder and GANs for the task of generating missing modalities.
- Design an ML model to generate the missing data and integrate it to the RecoGAN framework.
- Evaluate the results of the different approaches and quantify the improvement brought by the data completion.

Technologies/expertise to develop

- Encoders, Decoders, VAE, and GAN.
- Tensorflow, Keras, and Pytorch.
- Multimodal representation learning.
- Generative models.

In practice

Tunis, early 2021 (5-6 months)

References

[1] Goodfellow, Ian, et al. "Generative adversarial nets." *Advances in neural information processing systems*. 2014.
[2] Gao, Min, et al. "Recommender Systems Based on Generative Adversarial Networks: A Problem-Driven Perspective." *arXiv preprint arXiv:2003.02474* (2020).
[3] Mirza, Mehdi, and Simon Osindero. "Conditional generative adversarial nets." *arXiv preprint arXiv:1411.1784* (2014).

Multimodal Learning with Graph Neural Networks

Context

Graphs are a fundamental structure that provides an intuitive abstraction for modelling and analysing complex and highly interconnected data. A multitude of emerging problems can be represented using graph models and addressed using graph algorithms. Domains such as social networks, transportation, and biological networks are naturally modelled as graphs. Advanced applications such as detection of influential people, communities in a social network, or drug discovery are efficiently solved using graph algorithms.

With the emergence of machine learning, many research initiatives were led to extend machine learning pipelines with graph data. Traditional ML is mainly applied on data with an underlying Euclidean or grid-like structure. However, graph data is in a non-Euclidean space. Therefore, a multitude of methods such as Graph Neural Networks were introduced to adapt ML models to non-Euclidean structured data to preserve the topological information that otherwise could be lost in transforming the data into a Euclidean input.

In this master thesis, we propose to study different graph neural network architectures, such as graph autoencoders, graph convolutional networks, and graph attention networks etc. The purpose is to choose an architecture that best captures the graph topology and to integrate it within a multimodal framework in-development by Eura Nova.

Business opportunities

Graphs have different business applications ranging from social networks where they can be used to uncover structural information such as communities and influencers, to technological networks, where graphs make it possible to show the central nodes of a telecommunications network infrastructure and optimize the routing and load balancing across the infrastructure, or in biology for modelling metabolic pathways, genetic regulations, and protein interactions. The results of this project could substantially improve the accuracy of current ML algorithms by enriching them with topological information.

Contribution

The purpose of this project is to integrate the generation of the missing modalities in RecoGAN, a multimodal recommendation engine based on autoencoders and generative adversarial networks developed at Eura Nova, to improve the accuracy of predictions.

In this project, you will:

- Study the state of the art on graph neural networks
- Assess the performance of each architecture on a set of ML applications missing modalities
- Integrate the best architecture within our multimodal ML framework
- Experiment and discuss the results

Technologies/expertise to develop

- Graphs, GNN, GCN ...
- Tensorflow, Keras and Pytorch.
- Multimodal representation learning.

In practice

Tunis, early 2021 (5-6 months)

References

[1] Scarselli, Franco, et al. "The graph neural network model." *IEEE Transactions on Neural Networks* 20.1 (2008): 61-80.
[2] Wu, Zonghan, et al. "A comprehensive survey on graph neural networks." *IEEE Transactions on Neural Networks and Learning Systems* (2020).
[3] Cai, Hongyun, Vincent W. Zheng, and Kevin Chen-Chuan Chang. "A comprehensive survey of graph embedding: Problems, techniques, and applications." *IEEE Transactions on Knowledge and Data Engineering* 30.9 (2018): 1616-1637.