PhD in Information Technology and Electrical Engineering
Università degli Studi di Napoli Federico II

PhD Student: Franca Rocco di Torrepadula

Cycle: XXXVII

Training and Research Activities Report

Academic year: 2022-23 - PhD Year: Second

Tutor: prof. Nicola Mazzocca

Co-Tutor: prof. Sergio Di Martino

Date: October 23, 2023

1. Information:

- **PhD student: Franca Rocco di Torrepadula. PhD Cycle: XXXVII**
- **DR number: DR995856**
- **Date of birth: 26/06/1998**
- **Master Science degree: Computer Engineering**
  University: Università degli Studi di Napoli Federico II
- **Scholarship type: UNINA**
- **Tutor: Prof. Nicola Mazzocca**
- **Co-tutor: Prof. Sergio Di Martino**

2. Study and training activities:

Activity	Type[1]	Hours	Credits	Dates	Organizer	Certificate[2]
Complex network systems: introduction and open challenges	Seminar	2	0.4	17/11/22	SSSM	Y
Cybercrime and Information Warfare: National and International Actors	Seminar	2	0.4	18/11/22	Prof. Romano	Y
Privacy and Data Protection	Seminar	2	0.4	22/11/22	Prof. Romano	Y
Introduction to GAP9 for Autonomous Nano-Vehicles	Seminar	4.5	0.9	28/11/22	UNIBO	Y
IoT Data Analysis	Course	12	4	09-13-16-20-23-27/01/23	Dr. Della Corte	Y
Blockchain Day	Seminar	2	0.4	30/03/23	Napoli Blockchain	N
2023 Spring School in Transferable Skills	Doctoral School	8.5	2	24-25/05/23	Department of Pharmacy (Unina)	Y
20th International Symposium on Web and Wireless Geographical Information Systems	Seminar	8.5	1.7	12-13/06/23		Y
Ricerca e formazione	Seminar	6	1	22/09/23	CINI	Y

nella società della transizione digitale						
Semantic artifacts and multimedia knowledge graphs for bio-data integration	Course	10	2	11/09/23-03/10/23	Dr. Cristiano Russo	N

1) Courses, Seminar, Doctoral School, Research, Tutorship
2) Choose: Y or N

2.1. Study and training activities - credits earned

	Courses	Seminars	Research	Tutorship	Total
Bimonth 1	**0**	**2.1**	**7**	**0.04**	**9.14**
Bimonth 2	**4**	**0**	**6**	**0**	**10**
Bimonth 3	**0**	**0.4**	**9**	**0.48**	**9.88**
Bimonth 4	**2**	**1.7**	**6**	**0**	**9.7**
Bimonth 5	**0**	**0**	**10**	**0**	**10**
Bimonth 6	**2**	**1**	**7**	**0**	**10**
Total	**8**	**5.2**	**45**	**0.52**	**58.72**
Expected	**30 - 70**	**10 - 30**	**80 - 140**	**0 – 4.8**	

3. Research activity:

My research activity concerns the definition and the application of the intelligent systems to smart cities. In particular, during this second year, I mainly focused on the following aspects:

1) EdgeAI – Moving AI inference from cloud down to edge systems.

In recent years, complex and expensive techniques (*e.g.* based on Deep Neural Networks (DNNs)) are becoming more prevalent in the scientific literature related to the smart cities. However, the effectiveness of DNNs comes with expenses, in terms of both storage and computational costs, due to the large number of parameters and computations involved. For these reasons, such networks are typically deployed on cloud resources, implying that data must be moved from source locations to the cloud, thereby incurring possible latency, scalability, and privacy issues.

To mitigate the latter, DNN inference tasks are increasingly moving from cloud down to edge systems, often facing limited capabilities in terms of storage and computation. Thus, in the EdgeAI area several research efforts address the development of efficient DNNs, suitable for the inference on edge/embedded devices.

The *Knowledge Distillation* approach was proposed in 2015 as a model compression technique, with the aim of reducing the network structure (and thus both computational and storage costs) without significantly decreasing its performances. More in detail, a deep model (referred to as the *teacher*) is trained without strict requirements in terms of storage or computational costs. This training is typically performed offline, on a server or in the cloud. Subsequently, a smaller model (referred to as

the *student*) is designed to be more suitable for the deployment at the edge, namely by limiting the number of parameters or operations involved. Such a student is trained by exploiting (*i.e. distilling*) the knowledge acquired from the teacher, resulting in higher accuracy compared to training the student "from scratch". From a technical standpoint, applying KD involves redefining the loss function used during the training of the student, introducing also two novel hyperparameters, namely the *imitation factor* and the *temperature.*

In this context, there is still a lack of theoretical explanation of why and how the knowledge distillation works. Among several open issues, we focused on the so-called *temperature dilemma*. More in details, we deeply investigated the effect of different values of temperature in the KD, spotting a correlation with two other factors: the ability of the teacher w.r.t the task at hand, and the model capacity gap between the teacher and the student, related to the intrinsic capability of the student in following its teacher. On this basis, we propose a systematic methodologic called ***KYSA (Keep Your Student's Attention)*** to properly configure the temperature hyper-parameter, avoiding inefficient grid-search.

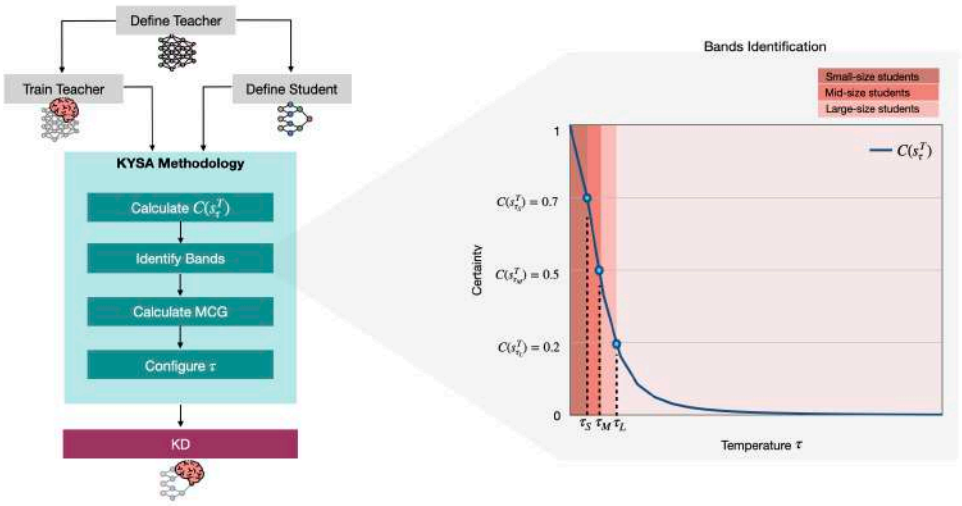

To briefly summarize how the methodology works, the main idea is that higher values of temperature force the student in considering more advanced hints from the teacher. Such hints can not be exploited at the same way by all the students, it depends on the model capacity gap: the higher the model capacity gap, the lower the ability of understanding the teacher. Thus, considering both the ability of the teacher and the model capacity gap, determining the level of complexity that is comprehensible to the student, and subsequently configuring the optimal temperature value, is a straightforward process.
Our approach has been tested and validated across different datasets and network structures. Additionally, the resulting networks have been deployed on diverse edge platforms, ranging from the STM32F4 to the GAPuino board, also in collaboration with *ETH Zürich* and the *University of Bologna*. This deployment has demonstrated a significant reduction in both storage and computational costs.

2) Predictive techniques for public transport optimization

Thanks to a collaboration with the *Hitachi Rail* company, I had the opportunity to delve into one of the key components of smart cities, namely *Intelligent Public Transportation Systems (IPTSs),* participating to the development of a real IPTS realized by Hitachi in a major Italian city.

In the context of public mobility, we conducted a Systematic Mapping Study (SMS) to investigate the research trends and gaps on passenger demand prediction. Indeed, despite the plethora and heterogeneity of researches carried out on this problem, to the best of our knowledge, there were no secondary studies attempting to synthesize the state of the art on this topic. We found out that such a problem is typically framed as a short-term prediction task. Longer-term predictions are still scarcely investigated, despite the potential utility of such insights for transportation companies in areas like service planning and fleet management. Moreover, most of the researches are mainly focused on proposing sophisticated and expensive prediction techniques, often leveraging Graph Convolutional Neural Networks (GCNNs) and/or Recurrent Neural Networks (RNNs). Finally, our SMS highlights lack of comparisons of the effectiveness of the research proposals. This is probably due to the scarcity of open datasets, as well as the lack of attention from the scientific community towards experiment reproducibility.

Moreover, within a collaboration with the *University College Dublin* and the *Hitachi Rail* company, we conducted a deep investigation on the generalizability and transferability of predictive methods in different mobility scenarios. We applied the same methods on both the city of Dublin, Ireland, and Genova, Italy, considering the bus journey prediction time as case study. We choose such cities since they present a number of differences in terms of urban plan, road network topology, and bus network structure.
The main conclusion of this study is that while scientific literature is increasingly adopting more complex and expensive techniques, these are not always necessary. It depends on the context. For instance, on Genova data, it is possible to achieve the same accuracy as in Dublin with a much simpler technique. To the best of our knowledge, the literature lacks studies of this type. This is mainly because the datasets used are typically small-scale (do not consider entire cities) and private.

Finally, given such a lack of public datasets in the IPTS field, within the collaboration with the *Hitachi Rail* company, we realized a simulator of the City of Genova, meant to generate synthetic data. To this purpose, we exploited mobility data (summarized in Origin-Destination (OD) matrices), the GTFS files (related to the scheduling and the organization of public transports), and the road network of the city (obtained through OpenStreetMap). The simulator was calibrated by comparing the simulated data with real-world ones, acquired through sensors installed on the buses of Genova.
The resulting simulator have been used for the generation of a synthetic but realistic dataset that can be exploit as input to AI-based predictive models. Moreover, it can be integrated as part of an urban digital twin.

4. Research products:

[J1] Di Martino, S., Landolfi, E., Mazzocca, N., **Rocco di Torrepadula, F.,** & Starace, L. L. L. (2023). *A visual-based toolkit to support mobility data analytics.* Published on *Expert Systems with Applications (ESWA)*. Indexed in Scopus and ISI Web of Science.

[J2] Cilardo, A., Maisto, V., Mazzocca, N., **Rocco Di Torrepadula, F**. (2023). An approach to the systematic characterization of multitask accelerated CNN inference in edge MPSoCs. Published on *ACM Transactions on Embedded Computing Systems (ACM TECS).* Indexed in Scopus and ISI Web of Science.

[J3] Rocco Di Torrepadula, F., Napolitano E. V., Di Martino S., Mazzocca N., *Data-Driven Public Transportation Demand Prediction: A Systematic Mapping Study.* Submitted to the *IEEE Transactions on Intelligent Transportation Systems (T-ITS).* Under the second round of review.

[J4] Rocco Di Torrepadula, F., Cilardo, A., Mazzocca N. *Keep Your Student's Attention: A Methodological Approach for Configuring the Temperature in Knowledge Distillation.* Submitted to the *IEEE Transactions on Pattern Analysis and Machine Intelligence (T-PAMI).*

[J5] Rocco Di Torrepadula, F., Maisto, V., Cilardo, A., Mazzocca. *Energy-Efficient DNNs via Knowledge Distillation and FPGA technologies.* Submitted to the *IEEE Transactions on Sustainable Computing (T-SUSC)*

[C1] Dunne, L., **Rocco Di Torrepadula, F**., Di Martino, S., McArdle, G., & Nardone, D. (2023, June). *Bus Journey Time Prediction with Machine Learning: An Empirical Experience in Two Cities*. In *International Symposium on Web and Wireless Geographical Information Systems (W2GIS2023)*, pp. 105-120. Part of Lecture Notes in Computer Science.

[C2] Di Martino, S., Mazzocca, N., **Rocco Di Torrepadula, F**., & Starace, L. L. L. (2023, June). *Mobility Data Analytics with KNOT: The KNime mObility Toolkit.*In *International Symposium on Web and Wireless Geographical Information Systems(W2GIS2023)*, pp. *95-104*. Part of Lecture Notes in Computer Science.

[C3] Rocco Di Torrepadula, F., Russo, D., Di Martino S., Mazzocca N., Sannino, P. *Using SUMO towards Proactive Public Mobility: Some Lessons Learned.* Accepted for the *1st ACM SIGSPATIAL Workshop on Sustainable Mobility (SuMob 2023).*

Best presentation award at *International Symposium on Web and Wireless Geographical Information Systems (2023)* for the presentation of the paper *Bus Journey Time Prediction with Machine Learning: An Empirical Experience in Two Cities.*

5. Conferences and seminars attended

International Symposium on Web and Wireless Geographical Information Systems. W2GIS2023. Quebec City, Canada. 12-13/06/2023. I presented the paper "Bus Journey Time Prediction with Machine Learning: An Empirical Experience in Two Cities".

6. Periods abroad and/or in international research institutions

Abroad research period at the *University College of Dublin,* Ireland, under the supervision of Prof. Gavin McArdle.
The research activities carried out in this period were focused on the topic of public transport analytics, in particular bus scheduling and passenger load estimation, by applying several novel spatiotemporal data analysis techniques.

The abroad research period took place from January 24th, 2023, to February 4th, 2023.

7. Tutorship

Computer System Design course: support and tutorship on Motorola 6800 programming and simulation on ASIM/ASIM Tool 10/03/23 (3 hours) and 17/03/23 (3 hours); Intel 82C59A Priority Interrupt Controller programming and Intel 8237A DMA controller programming 28/04/23 (2 hours). Total 8 hours.

Architettura dei sistemi digitali course: support and tutorship on the design of the Robertson's Multiplier on FPGA-based boards 13/12/22 (1 hour). Total 1 hour.

Risk Assessment course: support and tutorship on Fault Tree Analysis 19/04/23 (2 hours) and 20/04/23 (2 hours). Total 4 hours.

8. Plan for year three

For the next year, I plan to apply the proposed KYSA methodology to the ITS field. Indeed, the machine learning models applied so far in the IPTS I am studying are deployed in the cloud, thus incurring in the privacy and latency issues previous discussed. The application of the KYSA methodology would allow to run the inference tasks directly at the edge, avoiding latency/privacy issues and improving the performance of the IPTS.
The next step is moving the training directly at the edge (*i.e.* online learning). This is one of the topic of our collaboration with the research group of Prof. Luca Benini (*ETH Zürich, Università di Bologna*). Within this collaboration, I plan to investigate the integration of the Knowledge Distillation approach in the online learning framework, as a technique for effectively initializing the model, one of the main open problem for online learning.
Finally, I plan to do a research period abroad at the *Leibniz Universität* of Hannover, Germany, under the supervision of Prof. Wolfgang Nejdl and Dr. Marco Fisichella. Here, I will explore the application of the federated learning approach (a particular case of online learning) in the ITS field.

The draft title of my thesis is "EdgeAI for Smart Cities".