



UNIVERSITA' DEGLI STUDI - L'AQUILA

Dipartimento di Ingegneria e Scienze dell'Informazione e Matematica
Corso di Laurea in Informatica



L'Aquila, June 8th 2020

A thesis on the following topic is available "**Machine learning-driven Context-aware service discovery for Microservice architectures**"

Thesis Degree: Master Thesis in Computer Science

Prior Knowledge: Basic understanding of Microservice architectures, Basic to intermediate knowledge of Machine Learning, Knowledge of dockers and deployment technologies is a plus

In collaboration with: Mauro Caporuscio, Associate Professor, Linnaeus University, Sweden and Karthik Vaidhyanathan, Ph.D. student, Gran Sasso Science Institute

Starting Point: You may read the thesis of Marco De Toma available here: <https://tinyurl.com/yd3pz2s7>

Description: Service Discovery is achieved in traditional systems using technologies like Apache Zookeeper, Netflix Eureka, etc. These approaches often use a greedy strategy for selecting the instance in which they select the instance that best satisfies the properties requested by the client/applications. However, these approaches do not take into account the context of the client/application or that of the running instances. They also do not take into account the expected QoS of the instances. Towards this direction, we have developed a machine learning-driven context-aware service discovery mechanism that makes use of deep neural networks and reinforcement learning techniques to select the best instance during the process of service discovery

Context: Microservice-based architecture (MSA) has emerged as one of the popular styles for architecting large scale systems. The major reasons for this trend can be attributed to the out-of-the-box features provided by these architectures such as scalability, heterogeneity, etc. This is especially true given the fact that companies like Netflix, Amazon, Uber, etc have their entire systems developed based on MSA consisting of thousands of microservices that are replicated across hundreds of instances to handle the scale. MSA based systems make use of a mechanism known as Service Discovery to allow applications/clients to locate and discover running instances of different microservices and it is accomplished with the help of a service registry.

Objective: Goal of this work is threefold 1) Extend the existing approach to support QoS trade-offs and large-scale systems; 2) Experiment the approach by integrating with a diverse range of ML algorithms; 3) Explore the possibility of using transfer learning for making the approach robust across different classes of systems.

Prof. Henry Muccini

