

## Transfer Learning for Performance Analysis of Machine Learning Systems

Pooyan Jamshidi, Assistant Professor, Computer Science and Engineering Department,  
University of South Carolina

**Abstract.** A wide range of modern machine learning systems (e.g., autonomous systems, big data analytics, robotics, deep neural architectures) are built configurable. These systems offer a rich space for adaptation to different domains and tasks. Developers and users often need to reason about the performance of such systems, making tradeoffs to change specific quality attributes or detecting performance anomalies. For instance, developers of image recognition mobile apps are not only interested in learning which deep neural architectures are accurate enough to classify their images correctly, but also which architectures consume the least power on the mobile devices on which they are deployed. Recent research has focused on models built from performance measurements obtained by instrumenting the system [1]. However, the fundamental problem is that the learning techniques for building a reliable performance model do not scale well, simply because the configuration space is exponentially large that is impossible to exhaustively explore. For example, it will take over 60 years to explore the whole configuration space of a system with 25 binary options.

In this talk, I will start motivating the configuration space explosion problem based on my previous experience with large-scale big data systems in industry. I will then present our transfer learning solution to tackle the scalability challenge: instead of taking the measurements from the real system, we learn the performance model using samples from cheap sources, such as simulators that approximate the performance of the real system, with a fair fidelity and at a low cost [1, 2]. Results show that despite the high cost of measurement on the real system, learning performance models can become surprisingly cheap as long as certain properties are reused across environments. In the second half of the talk, I will present empirical evidence, which lays a foundation for a theory explaining why and when transfer learning works by showing the similarities of performance behavior across environments. I will present observations of environmental changes' impacts (such as changes to hardware, workload, and software versions) for a selected set of configurable systems from different domains to identify the key elements that can be exploited for transfer learning. These observations demonstrate a promising path for building efficient, reliable, and dependable machine learning systems. Finally, I will share my research vision for the next five years and outline my immediate plans to further explore the opportunities of transfer learning.

**Keywords:** Machine Learning, Artificial Intelligence, Transfer Learning.

### References:

1. Jamshidi, P. and Casale, G., 2016, September. An uncertainty-aware approach to optimal configuration of stream processing systems. In *2016 IEEE 24th International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS)* (pp. 39-48). IEEE.
2. Jamshidi, P., Siegmund, N., Velez, M., Kästner, C., Patel, A. and Agarwal, Y., 2017, October. Transfer learning for performance modeling of configurable systems: An exploratory analysis. In *Proceedings of the 32nd IEEE/ACM International Conference on Automated Software Engineering*(pp. 497-508). IEEE Press.

**Bio:** *Dr. Pooyan Jamshidi* is an Assistant Professor at the University of South Carolina. Prior to his current position, he was a research associate at Carnegie Mellon University and Imperial

College London, where he primarily worked on transfer learning for performance analyses of highly-configurable systems including robotics and big data systems. He holds a PhD from Dublin City University. Jamshid's general research interests are at the intersection of software engineering, systems, and machine learning, and his focus is primarily in the areas of distributed machine learning systems. Jamshidi has spent 7 years in the software industry before his PhD. Dr. Jamshidi can be reached at [pjamshid@cse.sc.edu](mailto:pjamshid@cse.sc.edu)