
Transfer learning of Kernel configuration performance across Linux versions

Hugo Martin*¹

¹Institut de Recherche en Informatique et Systèmes Aléatoires (IRISA) – Université de Rennes 1, Université de Rennes, Institut National des Sciences Appliquées - Rennes, Institut National des Sciences Appliquées, Université de Bretagne Sud, École normale supérieure - Rennes, Institut National de Recherche en Informatique et en Automatique, CentraleSupélec : UMR6074, Centre National de la Recherche Scientifique, IMT Atlantique Bretagne-Pays de la Loire, Institut Mines-Télécom [Paris] – Avenue du général Leclerc Campus de Beaulieu 35042 RENNES CEDEX, France

Résumé

With large scale and complex configurable systems, it is hard for users to choose the right combination of options (i.e., configurations) in order to obtain the wanted trade off between functionality and performance goals such as speed or size. Machine learning can help in relating these goals to the configurable system options, and thus, predict the effect of options on the outcome, typically after a costly training step. However, many configurable systems evolve at such a rapid pace that it is impractical to retrain a new model from scratch for each new version. In this paper, we propose a new method to enable transfer learning of performance predictions among versions of the same configurable system. Taking the extreme case of the Linux kernel with its 14.500 configuration options, we first investigate how performance predictions of kernel size degrade over successive versions. We show that the direct reuse of an accurate prediction model from 2017 quickly becomes inaccurate when Linux evolves, up to a 32% mean error by August 2020. We thus propose a new approach for transfer evolution-aware model shifting (tEAMS). It leverages the structure of a configurable system to transfer an initial predictive model towards its future versions with a minimal amount of extra processing for each version. We show that outperforms state of the art approaches and using tEAMS on the history of Linux kernels from 4.13 to 5.8, spanning over a period of 3 years, the model accuracy stays low and roughly constant over time.

*Intervenant