

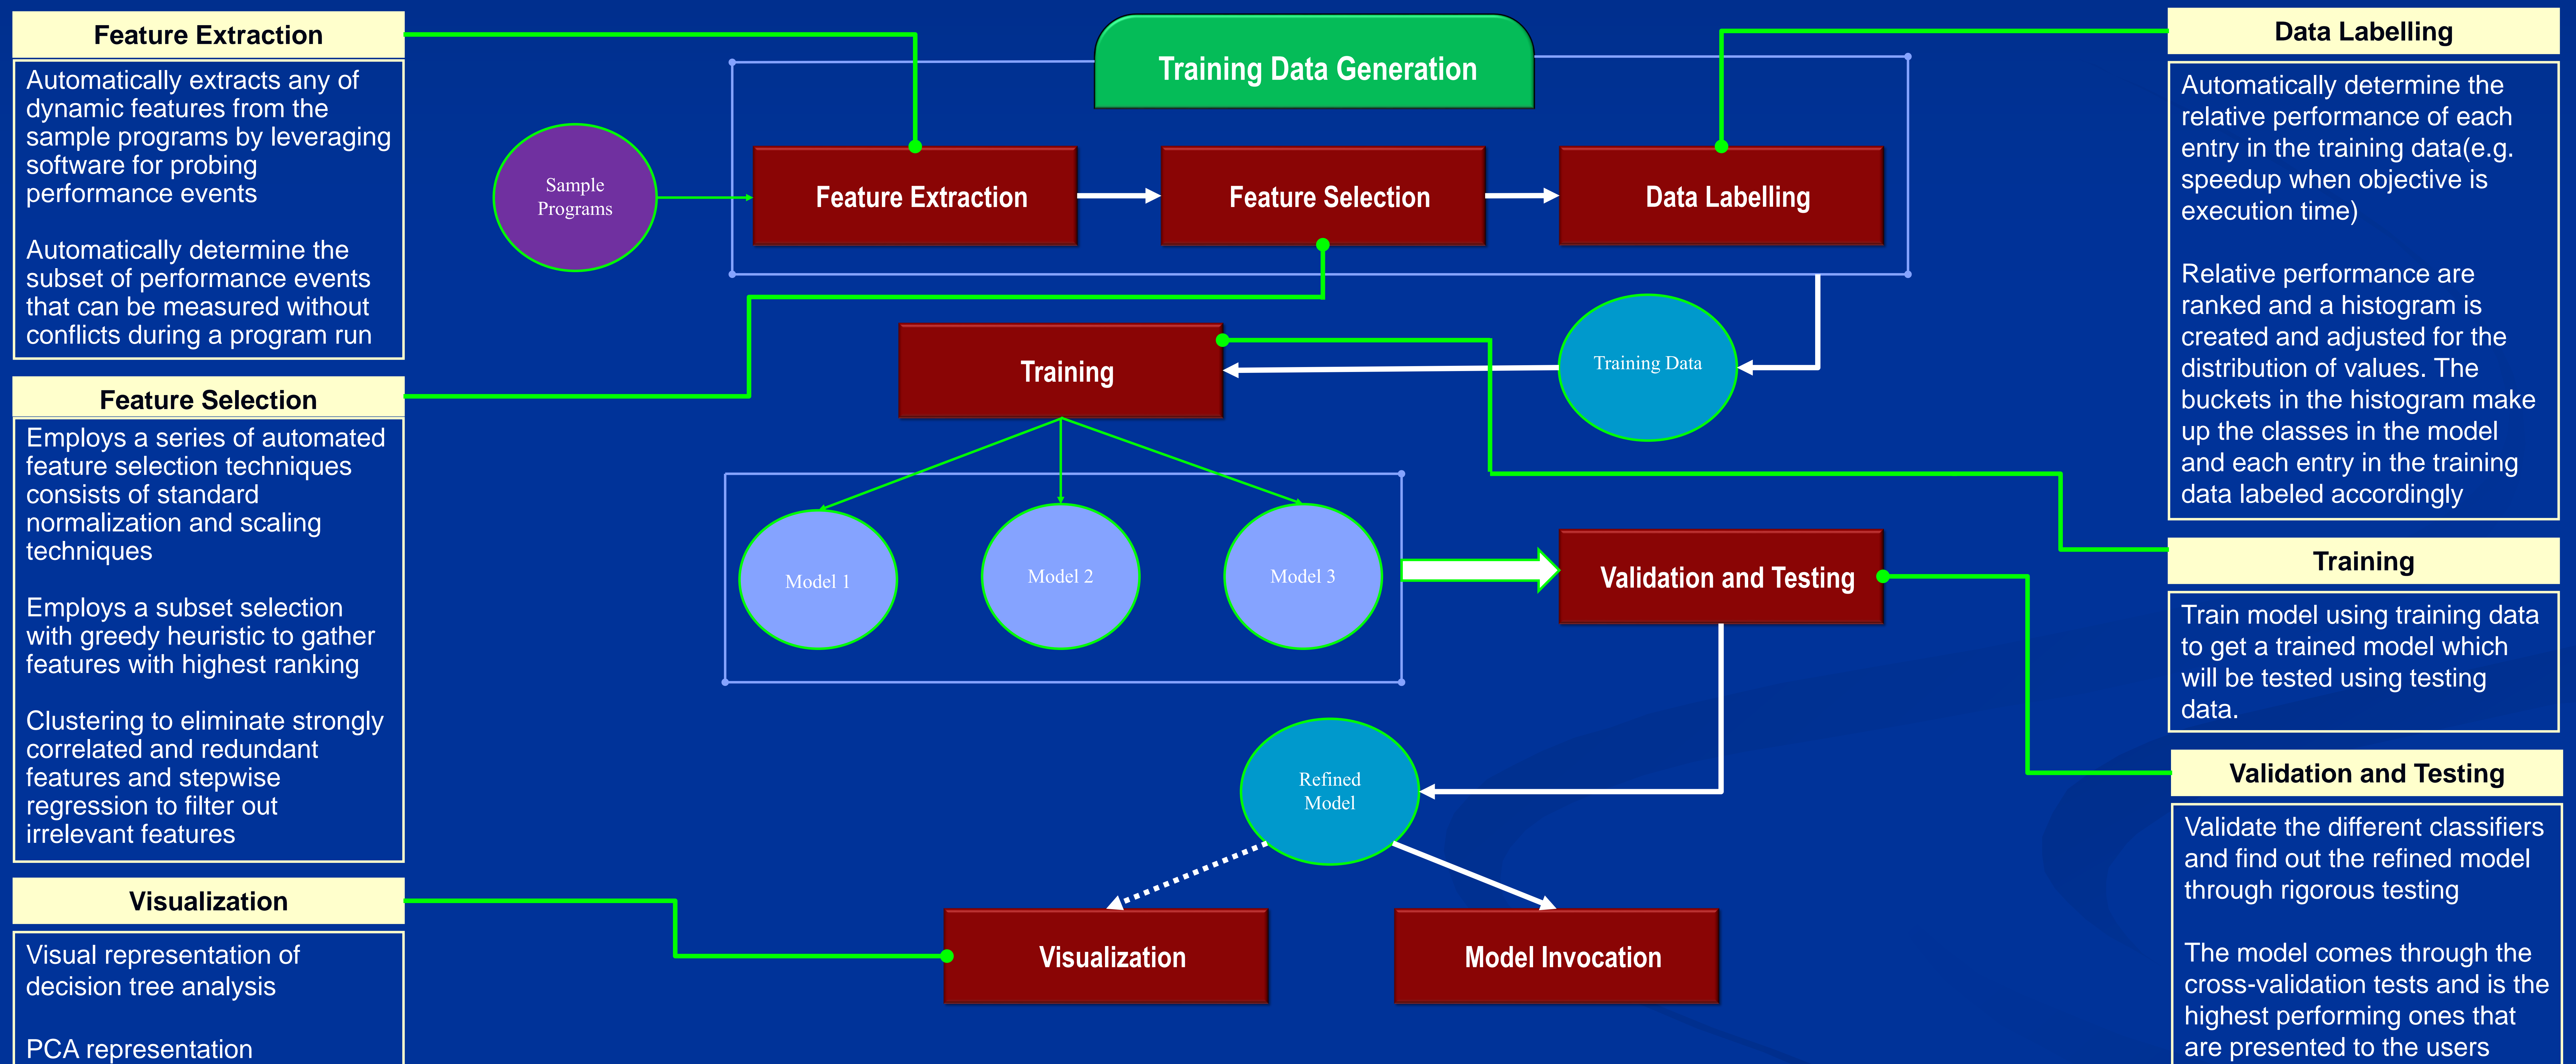
MLTUNE: A Tool-chain for Automating the Workflow of Machine-Learning Based Performance Tuning

Biplab Kumar Saha
Texas State University
San Marcos, TX
b_s183@txstate.edu


Saami Rahman
Texas State University
San Marcos, TX
saami.rahman@txstate.edu

Apan Qasem
Texas State University
San Marcos, TX
apan@txstate.edu

Recent interest in machine learning-based techniques have produced several excellent models for performance optimization and characterization of application behavior. Generally, these models are sensitive to architectural parameters and hence, are most effective when trained on the target platform. Training of these models, however, is a fairly involved process and requires knowledge of machine learning and statistics, which some users may lack. This poster presents MLTUNE, a tool-chain that automates the workflow of developing machine learning algorithms for performance tuning and enables performance engineers to build their own learning models. Leveraging existing open-source software the tool-chain provides automated mechanisms for sample generation, dynamic feature extraction, feature selection, data labelling, validation and model selection. The poster highlights the key features of the design of MLTUNE, which sacrifices sophistication in favor of generalization and automation. The system's applicability is demonstrated with a model for predicting profitable thread affinity configuration for parallel workloads.



Experimental Results



Current Status

Most component tools have been implemented and version 0 is available for download on github.

We are currently working on visualization



Acknowledgements

This research funded by NSF award CNS 1253292

We thank John Mellor-Crummey and Nathan Tallent for their help with the HPCToolkit interface