

---

# DEMONSTRATING CUSTOMIZABLE AND REUSABLE PIPELINES TO AUTOMATE AND REPRODUCE SYSML EXPERIMENTS

---

Grigori Fursin<sup>1 2</sup>

## ABSTRACT

Reproducing and comparing results from systems and machine learning papers is a very tedious, ad hoc and time-consuming process. I will demonstrate how to automate this process using open-source, portable, customizable and CLI-based Collective Knowledge workflows and pipelines developed by the community (CK). I will help participants run several real-world non-virtualized CK workflows from companies (General Motors, Arm), systems conferences (ASPLOS, CGO, Supercomputing), and MLPerf benchmark (MLP) to automate benchmarking and co-design of efficient software/hardware stacks for machine learning workloads. I hope that our approach will help authors reduce their effort when sharing reusable and extensible research artifacts while enabling artifact evaluators to automatically validate experimental results from published papers in a standard and portable way.

## 1 INTRODUCTION

Ten years ago I started collaborating with the community to automate and reproduce complex systems and machine learning experiments (Fursin, 2009). This experience helped us introduce and standardize validation of experimental results from published papers at PPOPP, CGO, Supercomputing and other systems conferences during the so-called Artifact Evaluation process (AE).

Voluntarily AE typically works as follows. Once a paper is accepted, the authors can submit a related artifact (code, data, models, results) to a special committee usually formed from senior graduate students and postdoctoral researchers. This committee attempts to reproduce results using submitted artifacts based on the standard ACM procedure we contributed to a few years ago (ACM): check if the artifact is permanently archived, follow its installation instructions and verify functional correctness, partially replicate results from the paper, assign reproducibility badges and so on.

The good news is that the number of papers participating in AE has dramatically increased over the years. For example, PPOPP'19 has set a new record with 20 papers passed through AE out of 29 papers accepted. Supercomputing'19 has even made artifact description appendices mandatory for all submitted papers. ACM also provided advanced article search based on reproducibility badges in the ACM Digital Library.

---

<sup>1</sup>cTuning foundation, France <sup>2</sup>dividiti, UK. Correspondence to: Grigori Fursin <Grigori.Fursin@cTuning.org>.

The bad news is that the lack of a common experimental framework, common research methodology, common formats and stable APIs places an increasing burden on evaluators to validate and compare a growing number of very complex artifacts. Furthermore, having too many papers with too many ad-hoc artifacts and Docker snapshots is almost as bad as not having any, since they cannot be easily reused, customized and built upon.

These problems are particularly noticeable in co-design projects developing efficient software and hardware stacks for emerging workloads including artificial intelligence and machine learning. Such stacks must be continuously benchmarked, optimized and compared across rapidly evolving machine learning models, data sets, frameworks, libraries, compilers and hardware architectures.

However, when overviewing more than 100 papers during artifact evaluation, I noticed that many of them use very similar experimental setups, scripts, pipelines, benchmarks, models, data sets and platforms. This motivated me to develop Collective Knowledge (CK), an open-source, extensible and CLI-based framework with a simple Python API to help researchers automate their repetitive tasks such as detecting software dependencies in a native, non-virtualized environment; installing missing packages; downloading data sets and models; compiling and running programs, implementing autotuning and co-design pipelines; crowdsourcing time-consuming experiments across computing resources provided by volunteers; automatically reproducing, plotting and validating experimental results, and so on (CK). Such approach helps a growing number of organizations apply DevOps principles to their research while integrating shared CK pipelines and components with their existing

projects ([USE](#)).

## 2 DEMONSTRATION

I will demonstrate how to prepare, run, visualize and reuse several CK pipelines for image classification from the 1st ACM ReQuEST tournament at ASPLOS'18 aiming to collaboratively co-design efficient SW/HW stacks for deep learning and reproduce results using the established AE methodology ([REQ](#)):

- **Edge:** TFLite with MobileNets on ARM-based Android devices - this CK pipeline will be reused in MLPerf ([MLP](#))
- **Cloud:** Intel Caffe with 8-bit ResNet-50 on AWS c5.18xlarge instances - this CK pipeline was validated by Amazon colleagues ([AIC](#))
- **FPGA:** TVM/MXNet with 8-bit ResNet-18 on an FPGA simulator

I will also show CK pipelines from several accepted SysML'19 papers which successfully passed Artifact Evaluation ([SAE, 2019](#)). The participants should be able to repeat all interactive steps on their own computers while learning CK concepts and interfaces. It can help future SysML authors share reusable pipelines with a common API to automate and accelerate validation of experimental results from their articles. The participants will need the following software to reproduce my demo ([CKI](#)):

- Linux or MacOS with bash (CK also supports Windows but several third-party tools used in my demo work only on Unix)
- Python 2.7 or 3.3+
- Git client
- CK CLI
- Any GCC, LLVM or ICC
- Internet access
- Optional: Android NDK and SDK to test ML pipelines on Android devices

## 3 SHORT BIOGRAPHY

I run the non-profit cTuning foundation where I work with the community to develop a methodology and open-source tools for collaborative, reproducible and reusable ML and systems research. I am also the architect of the "Collective Knowledge" technology ([CK](#)) to automate, crowdsource and reproduce experiments. Since 2014 I have helped initiate and standardize Artifact Evaluation at PPOPP, CGO and other systems conferences where I also introduced a unified Artifact Appendix.

## 4 PRIOR DEMO EXPERIENCE

My latest presentation and a live demo of the CK technology was at the O'Reilly AI conference ([AIC](#)) and FOSDEM'19 ([CKP, 2019](#)).

## REFERENCES

- (ACM) ACM Artifact Review and Badging Policy. <https://www.acm.org/publications/policies/artifact-review-badging>.
- (AE) Artifact Evaluation at computer systems conferences (methodology and automation tools). <http://cTuning.org/ae>.
- (AIC) Scaling deep learning on AWS using C5 instances with MXNet, TensorFlow, and BigDL: from the edge to the cloud. <https://conferences.oreilly.com/artificial-intelligence/ai-eu-2018/public/schedule/detail/71549>.
- (CK) Collective Knowledge: an open-source framework to automate, reproduce, reuse and crowdsource experiments at software, systems and machine learning conferences. <http://cKnowledge.org>.
- (CKI) Collective Knowledge installation guide. <https://github.com/ctuning/ck#installation>.
- (MLP) MLPerf: a broad ML benchmark suite for measuring performance of ML software frameworks, ML hardware accelerators, and ML cloud platforms. <https://mlperf.org>.
- (REQ) ReQuEST: open tournaments on collaborative, reproducible and Pareto-efficient software/hardware co-design of emerging workloads. <http://cKnowledge.org/request>.
- (USE) Collective Knowledge real-world use cases. <http://cKnowledge.org/partners>.
- (CKP) CK presentation and live demonstration at FOSDEM'19. [https://fosdem.org/2019/schedule/event/collective\\_knowledge](https://fosdem.org/2019/schedule/event/collective_knowledge), 2019.
- (SAE) Artifact Evaluation for SysML'19. <http://cTuning.org/ae/sysml2019.html>, 2019.
- Fursin, G. Collective Tuning Initiative: automating and accelerating development and optimization of computing systems. In *Proceedings of the GCC Developers' Summit*, June 2009. URL <https://hal.inria.fr/inria-00436029v2>.