

# Technical Advisory Council Meeting

February 27, 2020

 THE **LINUX** FOUNDATION

 **LF AI**

# Antitrust Policy Notice

- › Linux Foundation meetings involve participation by industry competitors, and it is the intention of the Linux Foundation to conduct all of its activities in accordance with applicable antitrust and competition laws. It is therefore extremely important that attendees adhere to meeting agendas, and be aware of, and not participate in, any activities that are prohibited under applicable US state, federal or foreign antitrust and competition laws
- › Examples of types of actions that are prohibited at Linux Foundation meetings and in connection with Linux Foundation activities are described in the Linux Foundation Antitrust Policy available at <http://www.linuxfoundation.org/antitrust-policy>. If you have questions about these matters, please contact your company counsel, or if you are a member of the Linux Foundation, feel free to contact Andrew Updegrave of the firm of Gesmer Undergone LLP, which provides legal counsel to the Linux Foundation

# Recording of Calls

This is a reminder that TAC calls are recorded and available for viewing on the [TAC Wiki](#)

# Reminder: LF AI Useful Links

Web site: [lfai.foundation](https://lfai.foundation)  
Wiki: [wiki.lfai.foundation](https://wiki.lfai.foundation)  
GitHub: [github.com/lfai](https://github.com/lfai)  
Landscape: [landscape.lfai.foundation](https://landscape.lfai.foundation) or [l.lfai.foundation](https://l.lfai.foundation)  
Mail Lists: <https://lists.lfai.foundation>

LF AI Logos: <https://github.com/lfai/artwork/tree/master/lfai>

LF AI Presentation Template:

[https://drive.google.com/file/d/1eiDNJvXCqSZHT4Zk\\_-czASlz2GTBRZk2/view?usp=sharing](https://drive.google.com/file/d/1eiDNJvXCqSZHT4Zk_-czASlz2GTBRZk2/view?usp=sharing)

Events Page on LF AI Website: <https://lfai.foundation/events/>

Events Calendar on LF AI Wiki (subscribe available):

<https://wiki.lfai.foundation/pages/viewpage.action?pageId=12091544>

Event Wiki Pages: <https://wiki.lfai.foundation/display/DL/LF+AI+Foundation+Events>

# Agenda

- Roll Call
- Approval of Minutes
- TAC Chairperson Election
- Invited Talk: CodeReef
- LF AI Updates
- Upcoming TAC Meetings
- Open Discussion

# TAC Voting Members

Member	TAC Contact	TAC Email
Amdocs	Ofer Hermoni*	<a href="mailto:oferher@gmail.com">oferher@gmail.com</a>
AT&T	Jack Murray	<a href="mailto:jfm@research.att.com">jfm@research.att.com</a>
Baidu	Daxiang Dong	<a href="mailto:dongdaxiang@baidu.com">dongdaxiang@baidu.com</a>
Ericsson	Nimisha Radia	<a href="mailto:nimish.radia@ericsson.com">nimish.radia@ericsson.com</a>
Huawei	Huang Zhipeng	<a href="mailto:huangzhipeng@huawei.com">huangzhipeng@huawei.com</a>
Nokia	Pantelis Monogioudis	<a href="mailto:pantelis.monogioudis@nokia.com">pantelis.monogioudis@nokia.com</a>
Tech Mahindra	Nikunj Nirmal	<a href="mailto:nn006444@techmahindra.com">nn006444@techmahindra.com</a>
Tencent	Juniping Du	<a href="mailto:junipingdu@tencent.com">junipingdu@tencent.com</a>
Zilliz	Jun Gu	<a href="mailto:jun.gu@zilliz.com">jun.gu@zilliz.com</a>
ZTE	Wei Meng	<a href="mailto:meng.wei2@zte.com.cn">meng.wei2@zte.com.cn</a>
Acumos AI Project	Anwar Aftab	<a href="mailto:anwar@research.att.com">anwar@research.att.com</a>
Angel Project	Fitz Wang	<a href="mailto:fitzwang@tencent.com">fitzwang@tencent.com</a>
ONNX Project	Prasanth Pulavarthi	<a href="mailto:prasanth.pulavarthi@microsoft.com">prasanth.pulavarthi@microsoft.com</a>

\* TAC Chairperson

# Approval of Minutes

- › Draft minutes from the February 13th, meeting of the TAC were previously distributed to the TAC members
- › **Proposed Resolution:**
  - › That the minutes of the February 13th meeting of the Technical Advisory Council of the LF AI Foundation are hereby approved

# TAC Chairperson Election

- › The annual TAC Chairperson election will begin in April
- › Please begin to think about any interest in this leadership role
- › Nominations will kick off mid April, followed by voting, and results by end of April
  
- › General TAC and Chairperson details can be viewed within the LF AI Charter under section 7 [here](#) - A summary of the role:
  - › Represent the TAC as a voting member on the Governing Board (will attend monthly meeting)
  - › Lead TAC agenda and meetings with coordination among the TAC representatives and broader community
  - › Attend sync meetings with LF AI staff to discuss overall TAC activities and planning
  - › General representation of the TAC and the LF AI technical community



# Invited Talk: CodeReef

Grigori Fursin  
grigori.fursin@codereef.ai



# CK enables production-quality ML/AI/systems R&D with DevOps principles since 2015

[cKnowledge.org/partners](https://cKnowledge.org/partners)



CK complements and interconnects popular tools and services rather than substituting them



[github.com/ctuning/ck](https://github.com/ctuning/ck)

[CodeReef.ai/static/docs](https://CodeReef.ai/static/docs)

# Many groups are working to co-design efficient AI / ML / SW / HW stacks

## AI hardware

- All major vendors (Google, NVIDIA, ARM, Intel, IBM, Qualcomm, Apple, AMD ...)

## AI models

Many groups in academia & industry (Google, OpenAI, Microsoft, Facebook ...)

## AI software

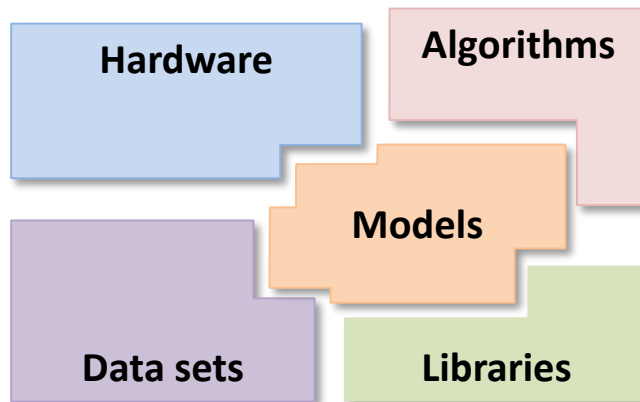
- AI frameworks (TensorFlow, MXNet, PyTorch, CNTK, Theano)
- AI libraries (cuDNN, libDNN, ArmCL, OpenBLAS)

## AI integration/services

- Cloud services (AWS, Google, Azure ...)



**Efficient AI/ML system must be very carefully co-designed**



**for various form factors**  
*(IoT, mobile, data centers)*



**while trading off multiple constraints**  
*(accuracy, speed, energy, size, costs)*

**and maximizing ROI**  
*(faster time to market, R&D sustainability, much better than all competitors)*



**Helping the society**

Healthcare  
Agriculture  
Finances  
Automotive  
Aerospace  
Meteorology  
Retail  
Robotics

...

# 90K+ AI / ML / SW / HW papers are published each year!

## AI hardware

- All major vendors (Google, NVIDIA, ARM, Intel, IBM, Qualcomm, Apple, AMD ...)

## AI models

- Many groups in academia & industry (Google, OpenAI, Microsoft, Facebook ...)

## AI software


- AI frameworks (TensorFlow, MXNet, PyTorch, CNTK, Theano)
- AI libraries (cuDNN, libDNN, ArmCL, OpenBLAS)

## AI integration/services

- Cloud services (AWS, Google, Azure ...)



## Numerous papers with ad-hoc code



**Numerous models, data sets, benchmarks, libraries and tools**

**Multiple competitions focusing mostly on accuracy (Kaggle, DawnBench)**

**A few benchmarks and competitions focusing on optimizing other metrics besides accuracy:**

*LPIRC, MLPerf*



## Helping the society

Healthcare  
Agriculture  
Finances  
Automotive  
Aerospace  
Meteorology  
Retail  
Robotics  
...

# 90K+ AI / ML / SW / HW papers are published each year!

## AI hardware

- All major vendors (Google, NVIDIA, ARM, Intel, IBM, Qualcomm, Apple, AMD ...)

## AI models

- Many groups in academia & industry (Google, OpenAI, Microsoft, Facebook ...)


## AI software

- AI frameworks (TensorFlow, MXNet, PyTorch, CNTK, Theano)
- AI libraries (cuDNN, libDNN, ArmCL, OpenBLAS)

## AI integration/services

- Cloud services (AWS, Google, Azure ...)

## Numerous papers with ad-hoc code



Can we now co-design efficient SW/HW/ML stacks and use them in production to support real-world applications?

*Non-proprietary code, datasets, benchmarks, libraries and tools*

Multiple competitions focusing mostly on accuracy (Kaggle, DawnBench)

A few benchmarks and competitions focusing on optimizing other metrics besides accuracy:

*LPIRC, MLPerf*

## Helping the society

Healthcare  
Agriculture  
Finances  
Automotive  
Aerospace  
Meteorology  
Retail  
Robotics

...



# Adoption of novel AI / ML techniques in production is extremely slow

## AI hardware

- All major vendors (Google, NVIDIA, ARM, Intel, IBM, Qualcomm, Apple, AMD ...)

## AI models

- Many groups in academia & industry (Google, OpenAI, Microsoft, Facebook ...)


## AI software

- AI frameworks (TensorFlow, MXNet, PyTorch, CNTK, Theano)
- AI libraries (cuDNN, libDNN, ArmCL, OpenBLAS)

## AI integration/services

- Cloud services (AWS, Google, Azure ...)

- Technological chaos: continuously changing algorithm/model/SW/HW stack




- Non-representative / outdated training sets
- No common experimental frameworks and established methodologies which can adapt to this chaos
- Numerous reproducibility issues
- Very little artifact reuse in 1000+ ML papers
- Very little tech. transfer from academia (toy examples and too many papers)
- Docker, Kubernetes and VM images hide the mess but do not solve above problems

**Public outcry about reproducibility and reusability crisis**

## Helping the society

- Healthcare
- Agriculture
- Finances
- Automotive
- Aerospace
- Meteorology
- Retail
- Robotics
- ...





# Many great tools, data sets and models to help researchers ...

## Applications

- Meteorology
- Health
- Robotics
- Automotive
- Economics
- Physics
- Astronomy
- Education

## Programs

- Image classification
- Object detection
- Natural Language processing
- Text processing
- Video processing
- Personal assistant

## AI/ML

### frameworks

- TensorFlow
- PyTorch
- MXNet
- Caffe
- MCT (CNTK)
- Keras
- Kubeflow
- AutoML
- SageMaker
- Apache Spark

## Scientific tools

- MATLAB
- Scilab
- Simulink
- LabVIEW
- Gnuplot
- LaTeX
- Ipython

## Build tools

- Make
- Cmake
- SCons
- Bazel
- Gradle
- Ninja

## Languages

- C++
- C#
- C
- Go
- PHP
- Fortran
- Java
- Python

## Compilers

- LLVM
- GCC
- Intel
- PGI
- TVM
- CUDA

## DevOps tools

- Git
- Jenkins
- Docker
- Kubernetes
- Singularity

## Package managers

- Anaconda
- Go
- Npm
- Pip
- Sbt
- dpkg
- Spack
- EasyBuild

## Libraries

- SciPy
- TFLite
- OpenBLAS
- MAGMA
- cuDNN
- cuFFT
- ArmNN
- CLBlast
- gemmlowp
- Boost
- HDF5
- MPI
- OpenCV
- Protobuf

## OS

- Linux
- MacOS
- BSD
- Windows
- Android

## Shells

- bash
- sh
- csh
- ksh
- Windows shell

## Benchmarks

- SPEC
- EEMBC
- HPCG
- LINPACK
- cBench
- MLPerf

## Hardware

- CPU
- GPU
- TPU / NN
- DSP
- FPGA
- Quantum
- Simulators
- Interconnects

## Knowledge sharing

- ArXiv
- ACM DL
- IEEE DL
- GitHub
- Zenodo
- FigShare
- Web pages

## Platforms

- HPC
- Desktops
- IoT
- Mobile
- Cloud services

## Databases / experiments

- MySQL
- PostgreSQL
- MongoDB
- CouchDB
- Text files
- JSON files
- XLS files

## Web services

- GitHub
- GitLab
- BitBucket
- Travis
- JupyterHub
- Codelabs
- SageMaker

## Workload managers



- MPI
- SLURM
- PBS
- FLUX

## Datasets

- ImageNet
- KITTI
- COCO
- MiDataSets
- Human Cell Atlas
- 1000 Genomes
- Earth models
- OpenStreetMap

## Models

- GoogleNet
- AlexNet
- VGG
- ResNet
- MobileNets
- SSD
- SqueezeNet
- DeepSpeech



# ... but it's not easy to connect them together into reproducible AI / ML workflows

## Applications

- Meteorology
- Health
- Robotics
- Automotive
- Economics
- Physics
- Astronomy
- Education

## Programs

- Image classification
- Object detection
- Natural Language processing
- Text processing
- Video processing
- Personal assistant

## AI/ML frameworks

- TensorFlow
- PyTorch
- MXNet
- Caffe
- MCT (CNTK)
- Keras
- Kubeflow
- AutoML
- SageMaker
- Apache Spark

## Scientific tools

- MATLAB
- Scilab
- Simulink
- LabVIEW
- Gnuplot
- LaTeX
- Ipython

## Models

- GoogleNet
- AlexNet
- VGG
- ResNet
- MobileNets
- SSD
- SqueezeNet
- DeepSpeech

## Build tools

- Make
- Cmake
- SCons

## Languages

- C++
- C#
- C
- Go

## Compilers

- LLVM
- GCC
- Intel
- PGI

## DevOps tools

- Git
- Jenkins
- Docker

## Package managers

- Anaconda
- Go
- Nnm

## Libraries

- SciPy
- TFLite
- OpenBLAS
- MAGMA
- cuDNN
- cuFFT
- ArmNN
- CLBlast
- gemmlowp
- Boost
- HDF5
- MPI
- OpenCV
- Protobuf

## OS

- Linux
- MacOS
- BSD
- Windows
- Android

## Shells

- bash
- sh
- csh
- ksh
- Windows shell

## Benchmarks

- SPEC
- EEMBC
- HPCG
- LINPACK
- cBench
- MLPerf

## Hardware

- CPU
- GPU
- TPU / NN
- DSP
- FPGA
- Quantum
- Simulators
- Interconnects

## Knowledge sharing

- ArXiv
- ACM DL
- IEEE DL
- GitHub
- Zenodo
- FigShare
- Web pages

## managers

- MiDataSets
- Human Cell Atlas
- 1000 Genomes
- Earth models
- OpenStreetMap

- MPI
- SLURM
- PBS
- FLUX

- GitLab
- BitBucket
- Travis
- JupyterHub
- Codelabs
- SageMaker

- PostgreSQL
- MongoDB
- CouchDB
- Text files
- JSON files
- XLS files




© Elnur / Adobe Stock

# What I've noticed when reproducing papers at ACM and IEEE conferences since 2014

Authors develop their own ad-hoc scripts to do exactly the same "actions" across nearly all artifact submissions:

- Detect target hardware properties
- Detect software dependencies
- Install missing packages (code/data)
- Build code; run experiments
- Plot graphs and validate results
- Generate papers



Ad-hoc scripts to compile and run a program or a benchmark

- image corner detection
- matmul OpenCL
- compression
- neural network CUDA

*Have some common meta: which datasets can use, how to compile, CMD, ...*

Ad-hoc scripts to install packages or set up environment for code and data deps on a given platform

- GCC V8.1
- LLVM V7.0
- Intel Compilers 2017

*Have some common meta: compilation, linking and optimization flags*

Ad-hoc dirs for data sets with some ad-hoc scripts to find them, extract features, etc

- image-jpeg-0001
- bzip2-0006
- txt-0012
- video-raw-1280x1024

*Have some (common) meta: filename, size, width, height, colors, ...*

Ad-hoc dirs and scripts to record and analyze experiments

- cvs speedups
- txt hardware counters
- xls table with graph

*Have some common meta: features, characteristics, optimizations*

**Result**

# What I've noticed when reproducing papers at ACM and IEEE conferences since 2014

Authors develop their own ad-hoc scripts to do exactly the same "actions" across nearly all artifact submissions:

- Detect ta
- Detect so
- Install mi
- Build cod
- Plot and
- Generate

Ad-hoc scripts to compile and

image corner detection

Have some

## The reason I started developing the open-source Collective Knowledge framework (CK) was

to help researchers share their artifacts (code, data sets, models, scripts, experiments, papers) as reusable, portable and customizable packages and workflows with a simple Python API, CLI and JSON meta description.

I needed CK to support reproducibility initiatives and help the community crowd-benchmark AI/ML techniques across diverse software, hardware, data sets and models, collaboratively reproduce and compare results, and help companies move new techniques to production while supporting any technology and legacy code

**cKnowledge.org**

and analyze experiments

xls table with graph

characteristics, optimizations

**Result**

# CK concept : create, share and reuse ML/system “actions” as Python API, CLI and JSON

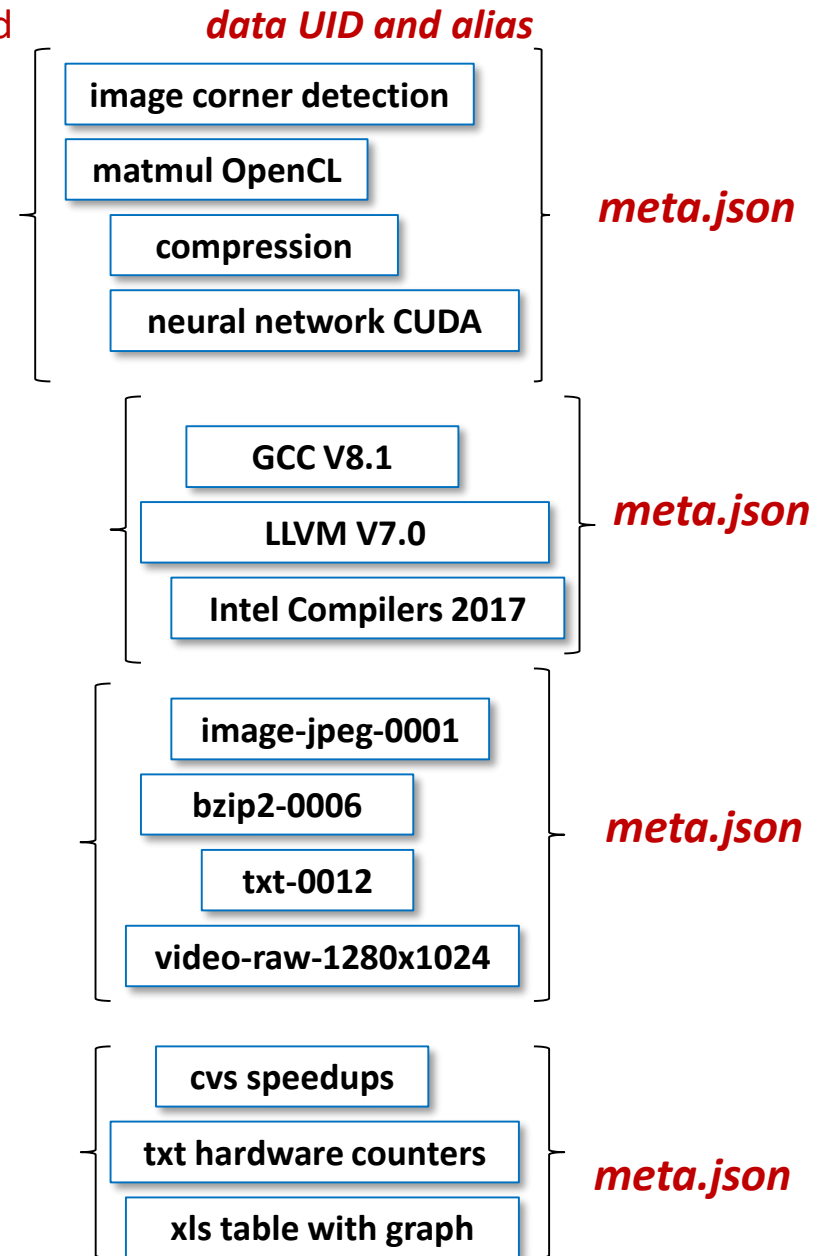
Such approach helped us bring standard DevOps practices and Continuous Integration to research and experimentation!

Python module  
“program”  
with functions:  
**compile and run**

Python module  
“soft”  
with function:  
**detect**

Python module  
“dataset”  
with function:  
**extract\_features**

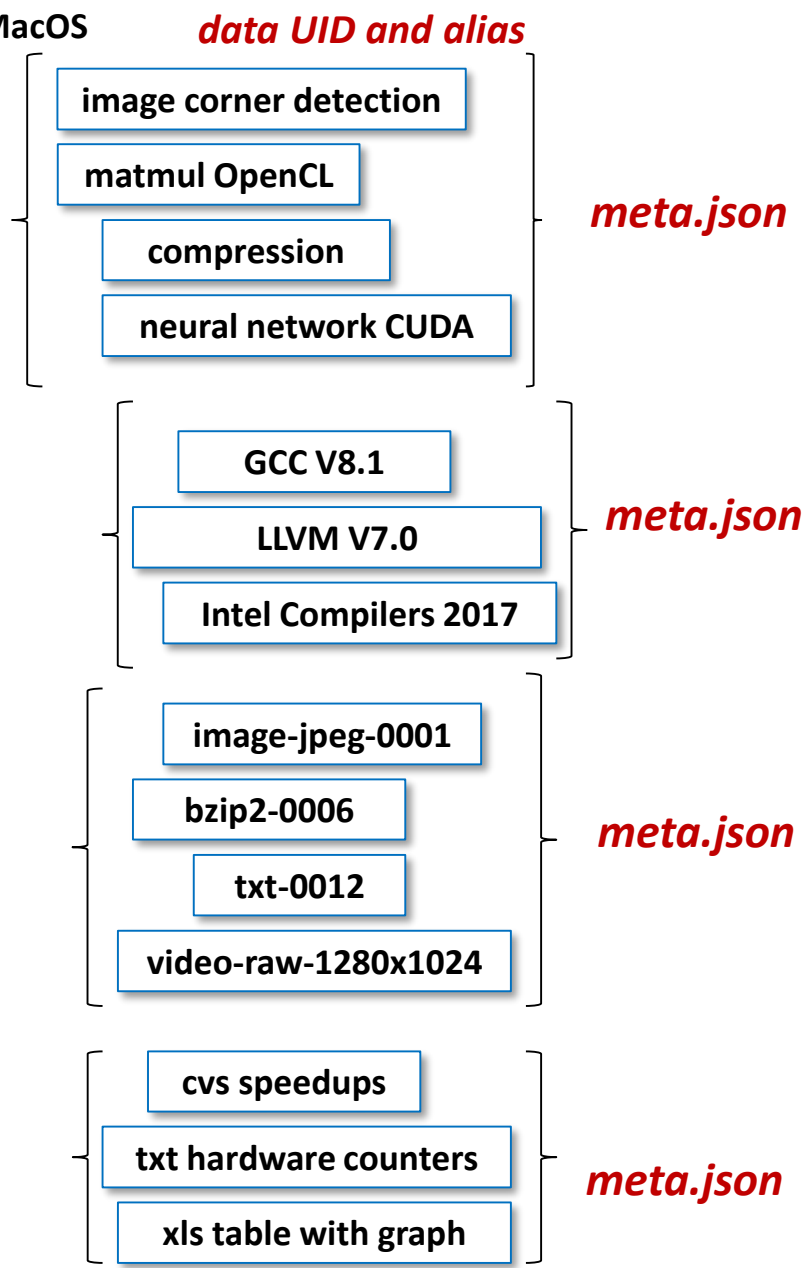
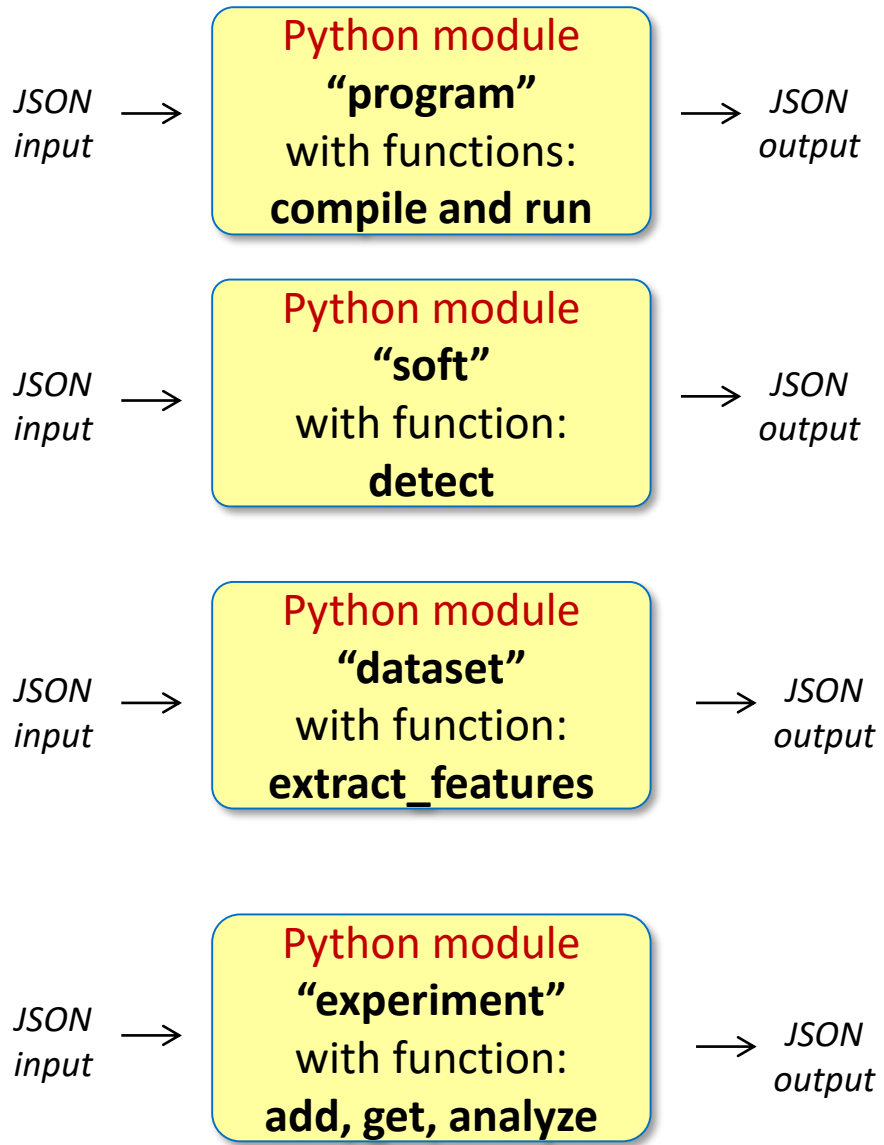
Python module  
“experiment”  
with function:  
**add, get, analyze**



# CK framework: simple CLI to create and access APIs (very portable - minimal dependencies)

CK: small python module (~200Kb); any python and git; Linux; Win; MacOS

```
$ pip install ck  
$ ck {function} {module name}:{data name} @input.json
```

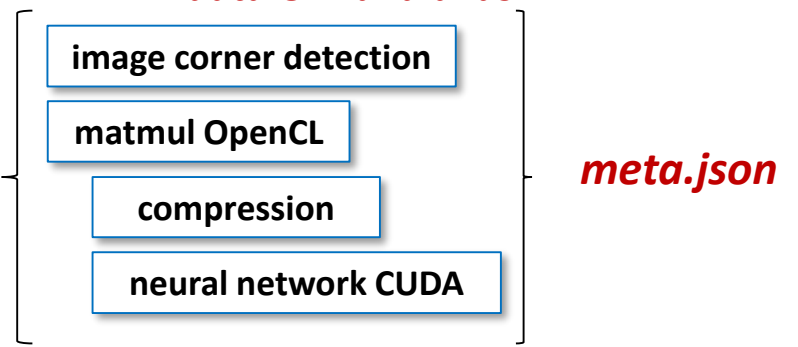
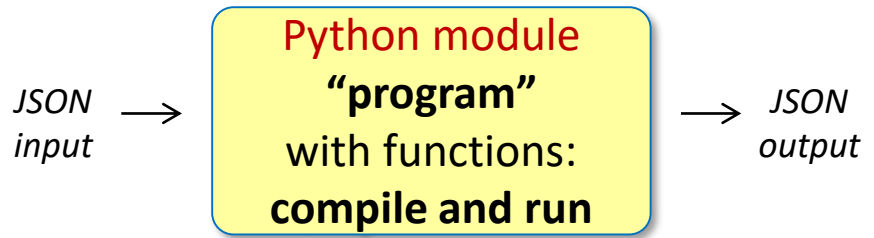


# CK framework: simple CLI to create and access APIs (very portable - minimal dependencies)

CK: small python module (~200Kb); any python and git; Linux; Win; MacOS

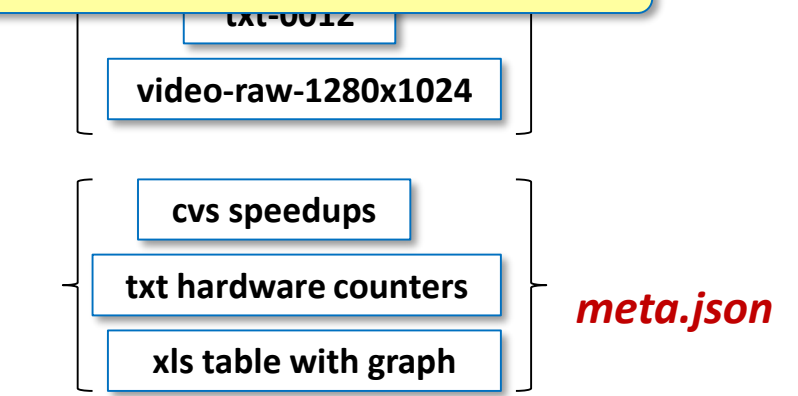

```
$ pip install ck  
$ ck {function} {module name}:{data name} @input.json
```

*data UID and alias*



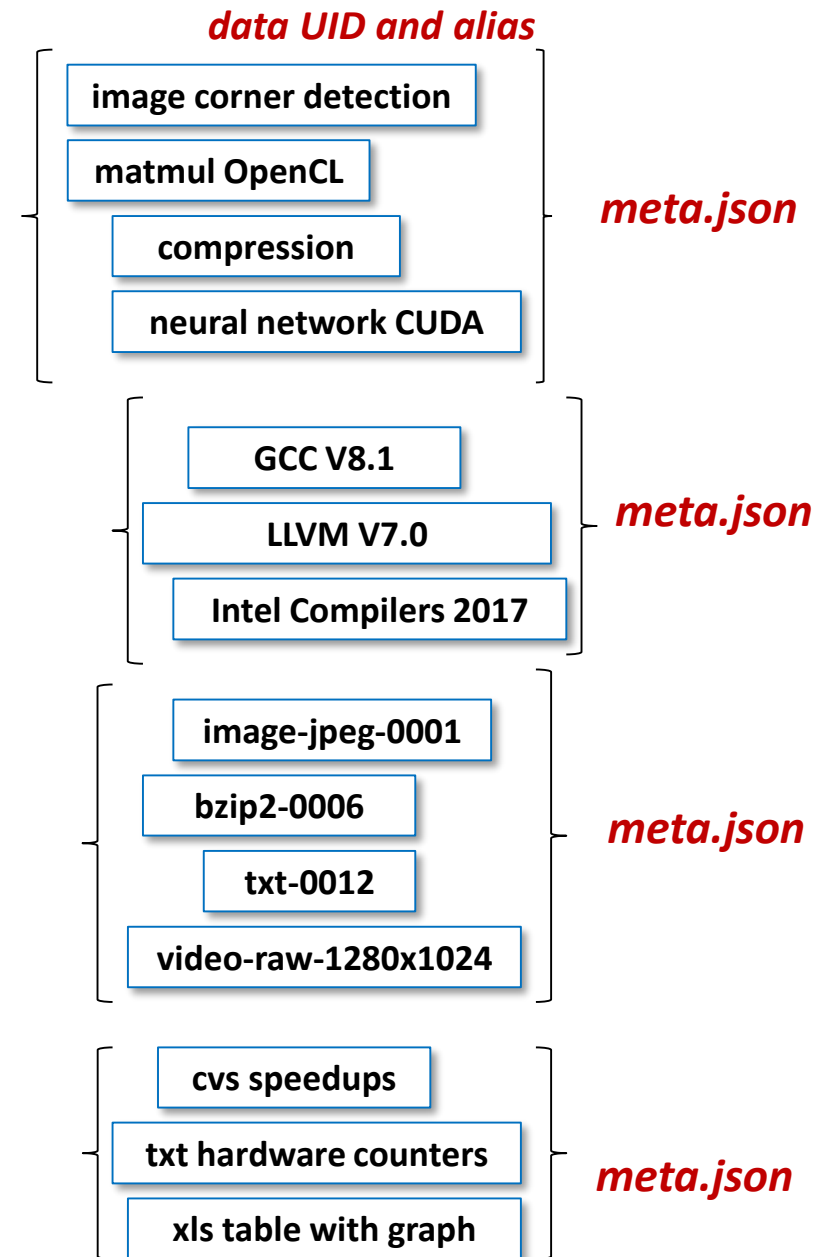
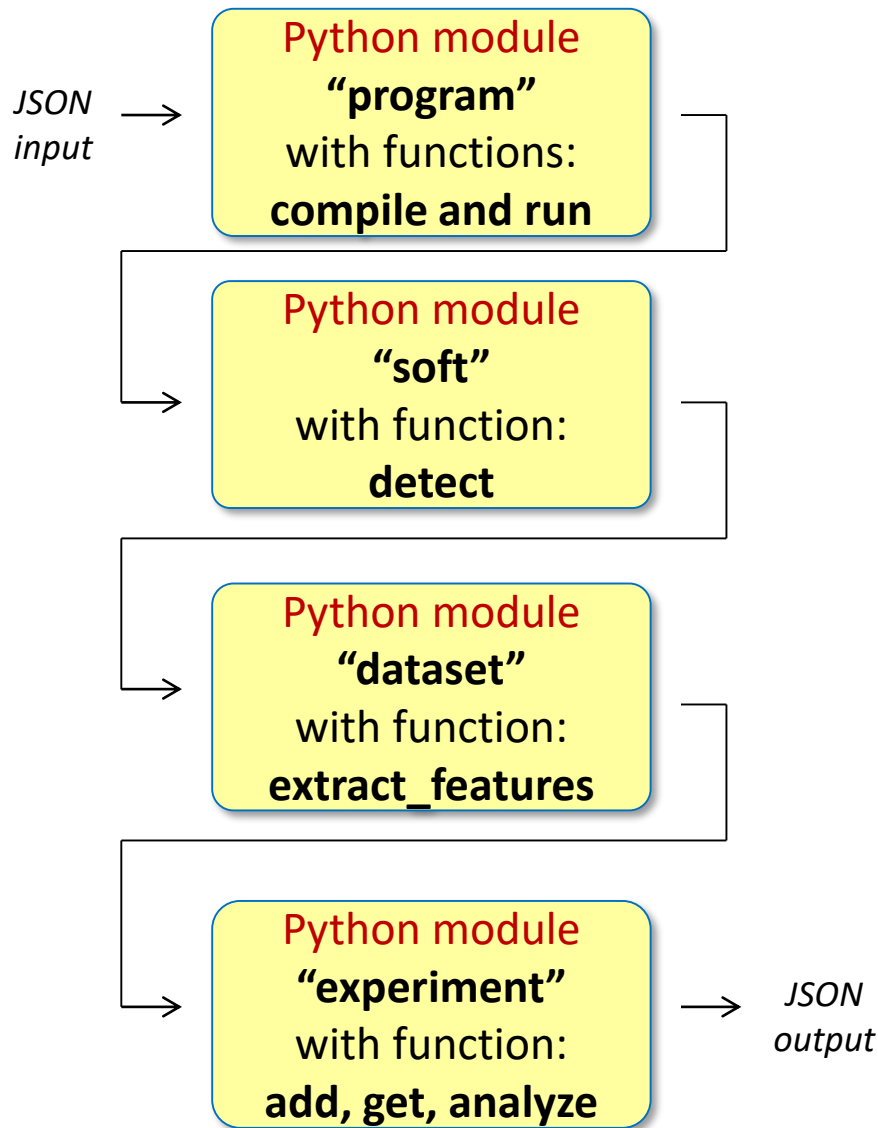
**600+ shared open-source automation actions from the past ACM, IEEE and ArXiv papers:**

[CodeReef.ai/portal/c/cr-action](http://CodeReef.ai/portal/c/cr-action)



# CK concept: assemble customizable workflows with JSON input and output

CK workflows describe dependencies on other CK packages using simple tags (code, data sets, models, frameworks, etc)







## 1) Describing different operating systems

```
$ ck pull repo:ck-env  
$ ck ls os  
$ ck load os:linux-64 --min
```

Needed to support real use cases from our partners:

<https://cKnowledge.org/partners>

## 2) Detecting and unifying information about platforms

```
$ ck detect platform --help  
$ ck detect platform --out=json  
$ ck load os:linux-64 --min
```

## 3) Detecting installed “software” (both code and data):

```
$ ck search soft --tags=dataset  
$ ck detect soft:compiler.llvm  
$ ck show env --tags=llvm
```

250+ software detection plugins:

[codereef.ai/portal/c/soft](https://codereef.ai/portal/c/soft)

## 4) Installing missing packages (both code and data): front-end to EasyBuild, Spack, scons ...

```
$ ck search package --tags=model  
$ ck install compiler:compiler-llvm-7.0.0-universal  
$ ck show env --tags=llvm  
$ ck virtual env --tags=llvm,v7.0.0
```

600+ shared packages:

[codereef.ai/portal/c/package](https://codereef.ai/portal/c/package)

We can now adapt to any native environment and use containers to make stable snapshots

[github.com/ctuning/ck/wiki/Portable-workflows](https://github.com/ctuning/ck/wiki/Portable-workflows)

# We then started developing CK workflows for portable and customizable AI / ML benchmarking

We developed a universal program workflow to compile, run, profile and autotune AI/ML applications across diverse models, data sets and platforms, validate results, record experiments, share and reproduce them, and report discrepancies

```
$ ck pull repo:ck-crowdtuning
```

```
$ ck ls program  
$ ck ls dataset
```

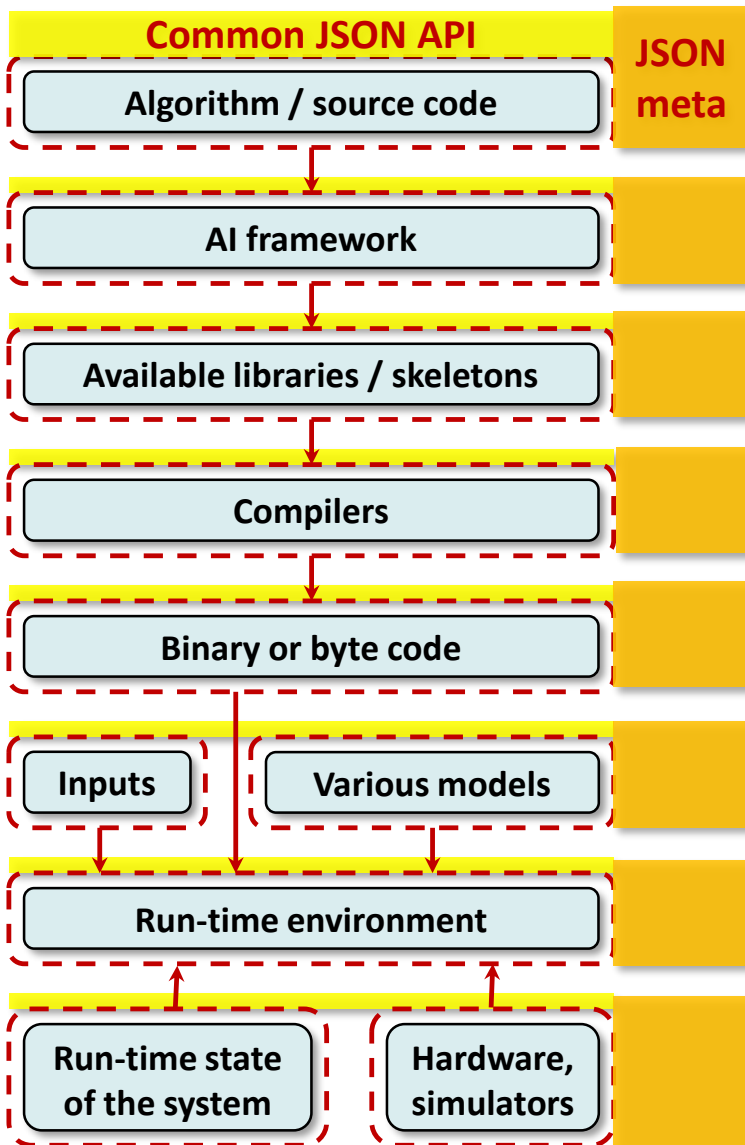
```
$ ck load program:cbench-automotive-susan --min  
$ ck compile program:cbench-automotive-susan --fast
```

```
$ ck run program:cbench-automotive-susan
```

```
$ ck autotune program:cbench-automotive-susan
```

```
$ ck crowdtune program:cbench-automotive-susan
```

```
$ ck replay experiment
```

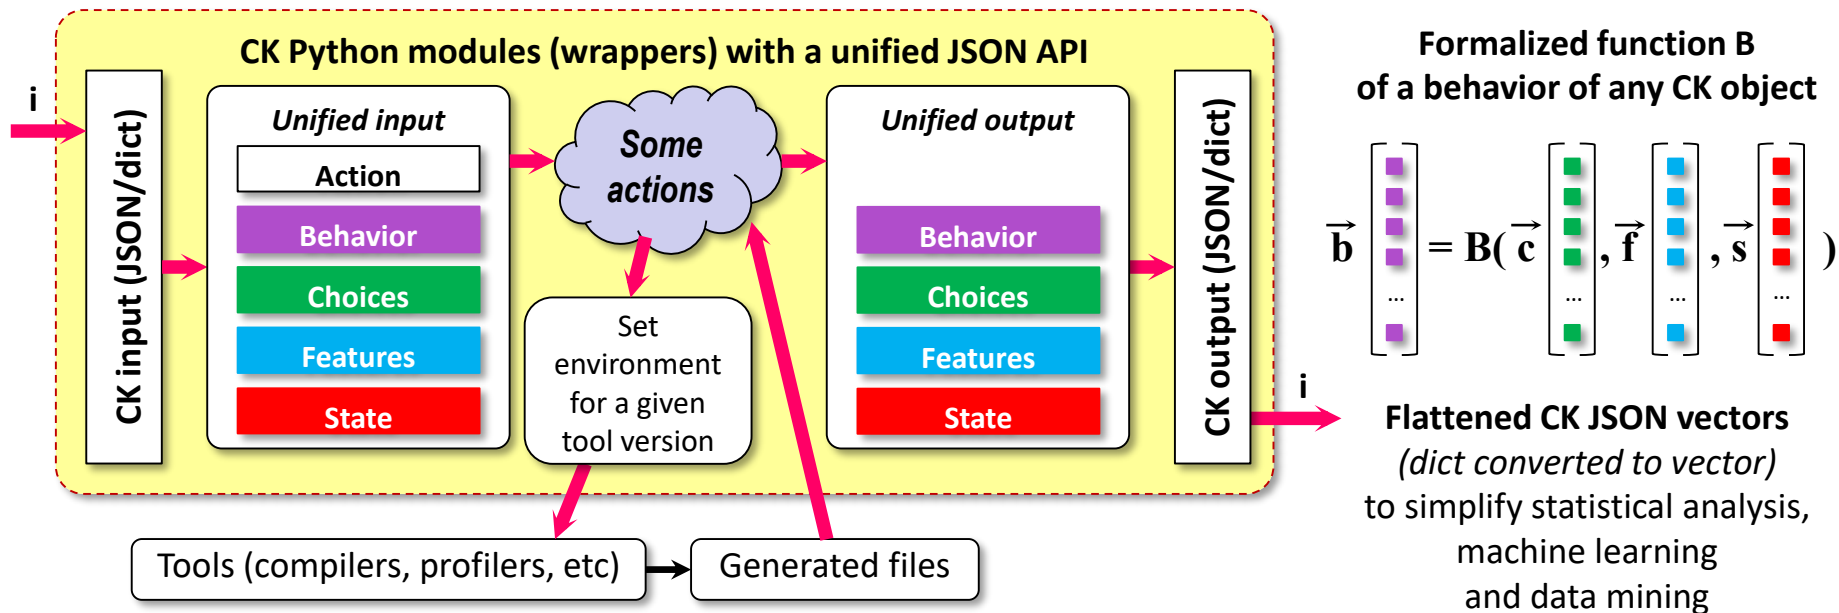


**CK workflows describe dependencies on CK soft detection plugins and packages to automatically adapt to a given platform and environment**

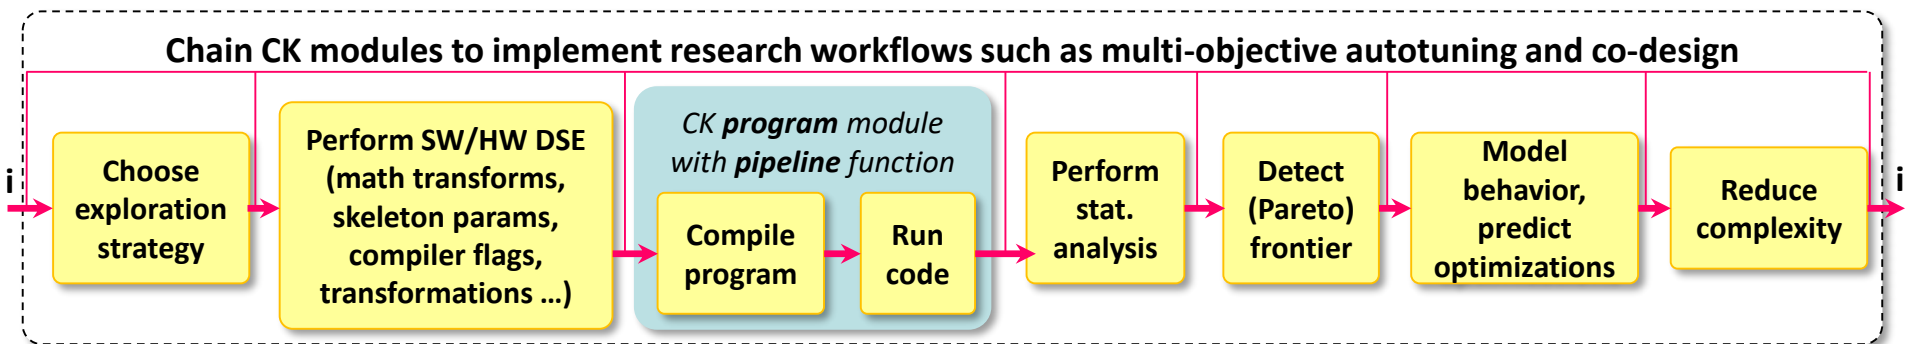
[CodeReef.ai/portal/c/cr-solution/demo-obj-detection-coco-tf-cpu-benchmark-linux-portable-workflows/#dependencies](https://CodeReef.ai/portal/c/cr-solution/demo-obj-detection-coco-tf-cpu-benchmark-linux-portable-workflows/#dependencies)

# Customizable CK workflows can be also used to autotune the whole AI/ML/SW/HW stack!

First expose coarse grain high-level choices, features, system state and behavior characteristics via CK APIs



Then automate crowd-benchmarking and optimization across diverse models, datasets and platforms



Keep best species (AI/SW/HW choices); model behavior; predict better optimizations and designs

[CodeReef.ai/portal/c/cr-result](https://CodeReef.ai/portal/c/cr-result)

# Gradually expose more optimizations and characteristics via JSON files

Autotuning and machine learning specification:

**CK flattened JSON key**

##characteristics#execution\_times@1

```
{  
  "characteristics":{  
    "execution times": ["10.3","10.1","13.3"],  
    "code size": "131938", ...},  
  "choices":{  
    "os":"linux", "os version":"2.6.32-5-amd64",  
    "compiler":"gcc", "compiler version":"4.6.3",  
    "compiler_flags":"-O3 -fno-if-conversion",  
    "platform":{"processor":"intel xeon e5520",  
      "l2":"8192", ...}, ...},  
  "features":{  
    "semantic features": {"number_of_bb": "24", ...},  
    "hardware counters": {"cpi": "1.4" ...}, ... }  
  "state":{  
    "frequency":"2.27", ...}  
}
```

↓

```
"flattened_json_key":{  
  "type": "text"|"integer" | "float" | "dict" | "list" | "uid",  
  "characteristic": "yes" | "no",  
  "feature": "yes" | "no",  
  "state": "yes" | "no",  
  "has_choice": "yes" | "no",  
  "choices": [ list of strings if categorical choice],  
  "explore_start": "start number if numerical range",  
  "explore_stop": "stop number if numerical range",  
  "explore_step": "step if numerical range",  
  "can_be_omitted": "yes" | "no"  
  ...  
}
```

# We were invited to organize reproducible AI/ML optimization tournaments (2018)

## AI hardware

- All major vendors (Google, NVIDIA, ARM, Intel, IBM, Qualcomm, Apple, AMD ...)

## AI models

Many groups in academia & industry (Google, OpenAI, Microsoft, Facebook ...)

## AI software

- AI frameworks (TensorFlow, MXNet, PyTorch, CNTK, Theano)
- AI libraries (cuDNN, libDNN, ArmCL, OpenBLAS)

## AI integration/services

- Cloud services (AWS, Google, Azure ...)

[cKnowledge.org/request](http://cKnowledge.org/request)

Finding the most efficient AI/SW/HW stacks across diverse models, data sets and platforms via open competitions, share them as reusable CK components and visualize on a public scoreboard


## Organizers (A-Z)

**Luis Ceze**, University of Washington  
**Natalie Enright Jerger**, University of Toronto  
**Babak Falsafi**, EPFL  
**Grigori Fursin**, cTuning foundation  
**Anton Lokhmotov**, dividiti  
**Thierry Moreau**, University of Washington  
**Adrian Sampson**, Cornell University  
**Phillip Stanley Marbell**, University of Cambridge

## Real use-cases

Healthcare  
Agriculture  
Finances  
Automotive  
Aerospace  
Meteorology  
Retail  
Robotics  
...

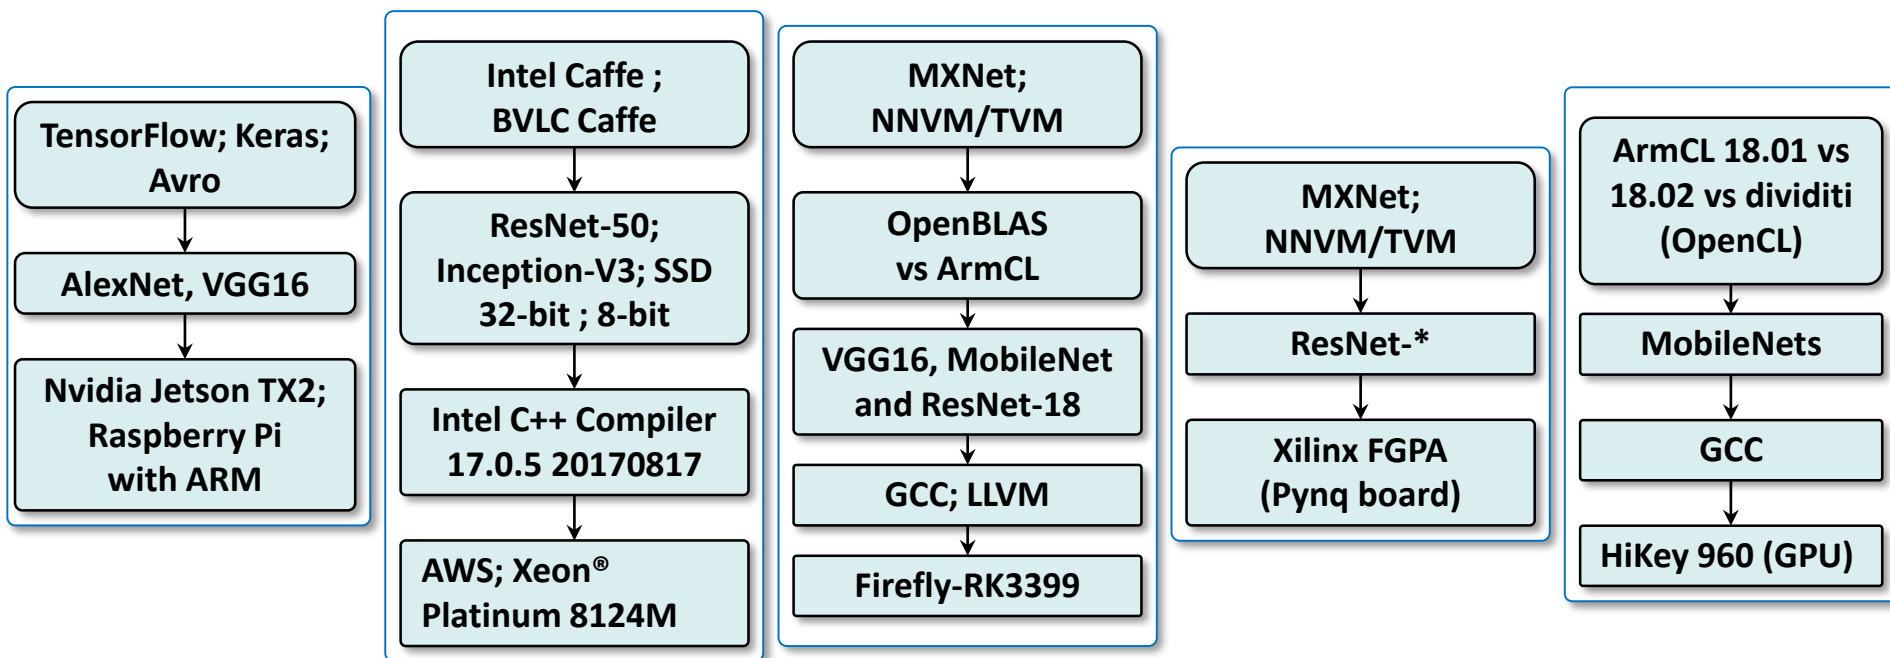
## CK platform & CodeReef



Interdisciplinary community

# We reproduced results from published ML papers and shared portable ML workflows!

8 intentions to submit and 5 submitted image classification workflows with unified Artifact Appendices



Public validation at [github.com/ctuning/ck-request-asplos18-results](https://github.com/ctuning/ck-request-asplos18-results) via GitHub issues.

All validated papers are published in the ACM DL  
with **portable, customizable and reusable CK components and workflows:**

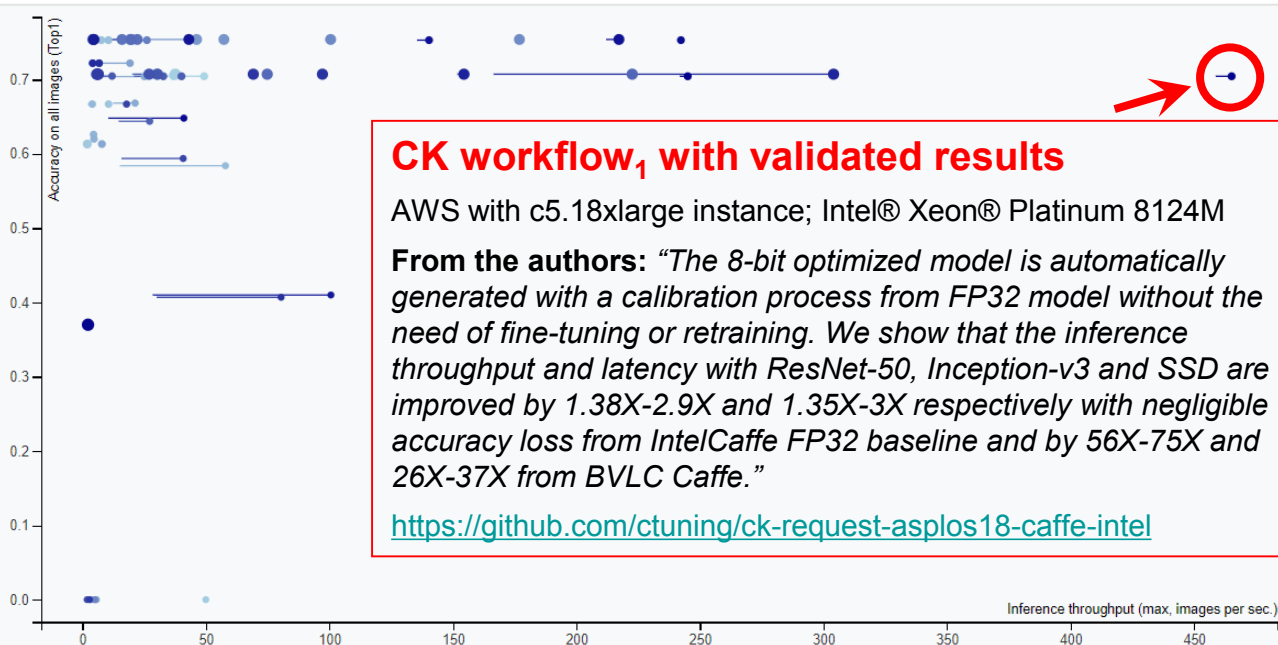
[dl.acm.org/citation.cfm?doid=3229762](https://dl.acm.org/citation.cfm?doid=3229762)

See ACM ReQuEST report: [portalparts.acm.org/3230000/3229762/fm/frontmatter.pdf](https://portalparts.acm.org/3230000/3229762/fm/frontmatter.pdf)

# All results are also available at online scoreboards

Multi-objective results for all AI/SW/HW stacks are presented on a live scoreboard and become available for public comparison and further customization, optimization and reuse!

ReQuEST @ ASPLOS'18 tournament (Pareto-efficient image classification)



Plot dimension X

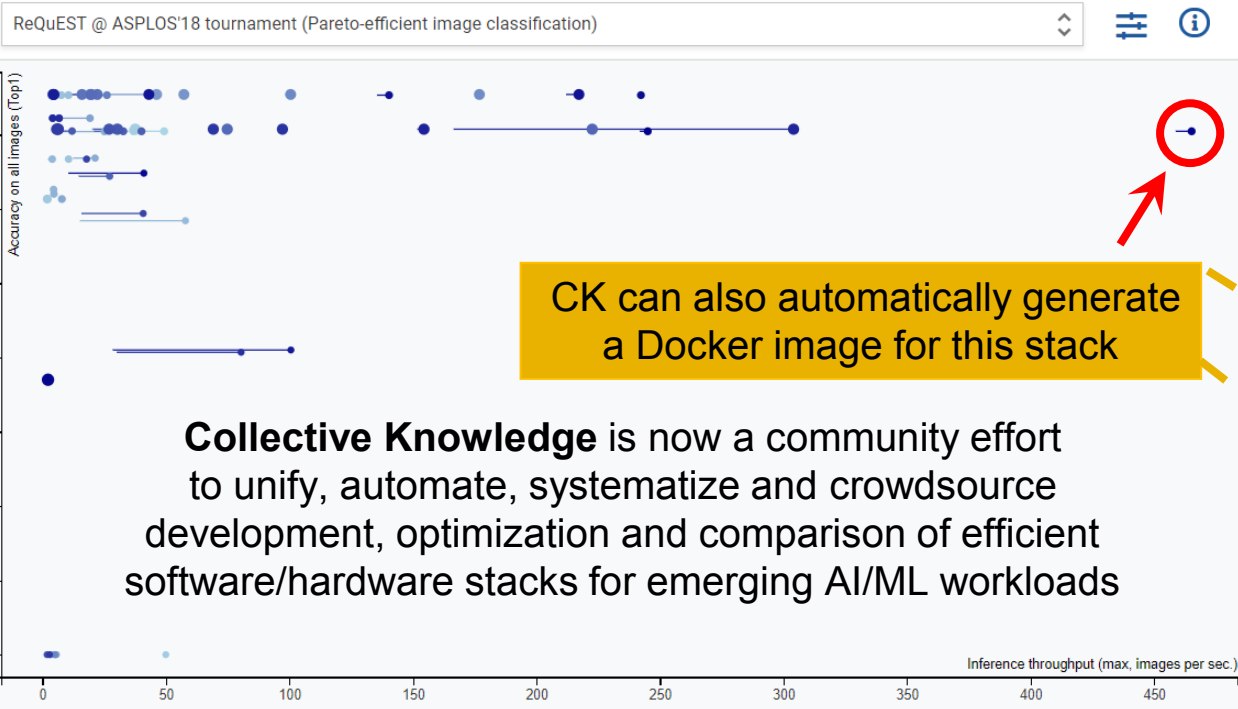
Inference throughput (max, images per sec.)
Experiment number
Prediction time per 1 image (min, sec.)
Inference latency for 1 image (min, sec.)
<b>Inference throughput (max, images per sec.)</b>
Accuracy on all images (Top1)
Accuracy on all images (Top5)
Model size (B)
Platform peak power (W)
Platform price (\$)
Usage cost (\$)
Platform species
Model species
Model precision
Dataset species
Device frequency (MHz)
CPU frequency (MHz)
GPU frequency (MHz)
Batch size

We are not announcing a single winner! We show all multi-dimensional results at [CodeReef.ai/portal/c/cr-result/pareto-efficient-ai-co-design-tournament-request-acm-asplos-2018](https://CodeReef.ai/portal/c/cr-result/pareto-efficient-ai-co-design-tournament-request-acm-asplos-2018) and let the users select best ML/SW/HW stacks depending on multiple constraints for their production use!



# Other companies managed to reproduce these results and started using CK

Multi-objective results for all AI/SW/HW stacks are presented on a live scoreboard and become available for public comparison and further customization, optimization and reuse!



**CK assists**  
**AWS market place**  
with collaboratively  
optimized AI/ML stacks



Accelerate technology transfer: companies can validate published techniques in their production environment using shared CK workflows!

We made a joint presentation with Amazon at O'Reilly AI conference (October 2018)

# General Motors uses CK to select the most efficient SW/HW stacks

Collaboratively optimizing deep learning via Collective Knowledge



OBJECT	FOUND	EXPECTED	FALSE POSITIVES	PRECISION	RECALL
	8	0	8	0	0
	0	0	0	1	1
	0	0	0	1	1

**MODE**  
Object detection

**ENGINE**  
TensorFlow library (prebuilt, cpu)

**MODEL**  
TensorFlow model - SqueezeDet (SqueezeDet)

**IMAGE SOURCE**  
KITTI Drive 0009

**IMAGES PER SECOND**  
**1.19**

**AVERAGE PRECISION**  
**0.67**

Stop


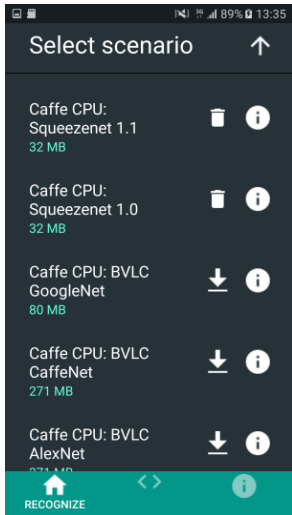
**Performance, accuracy, power consumption practically never match official reports!**

CK workflows and automation helped GM evaluate numerous models, datasets, frameworks and libraries to find the most efficient SW/HW stacks for object detection across Nvidia, AMD, ARM and Intel platforms (CUDA, OpenCL, OpenMP ...)

GM presentation about using CK: [www.youtube.com/watch?v=1ldgVZ64hEI](https://www.youtube.com/watch?v=1ldgVZ64hEI)

# CK workflows were used to crowdsource AI/ML benchmarking across Android devices

## Continuously collect statistics, bugs and misclassifications



The number of distinct participated platforms: **800+**

The number of distinct CPUs: **260+**

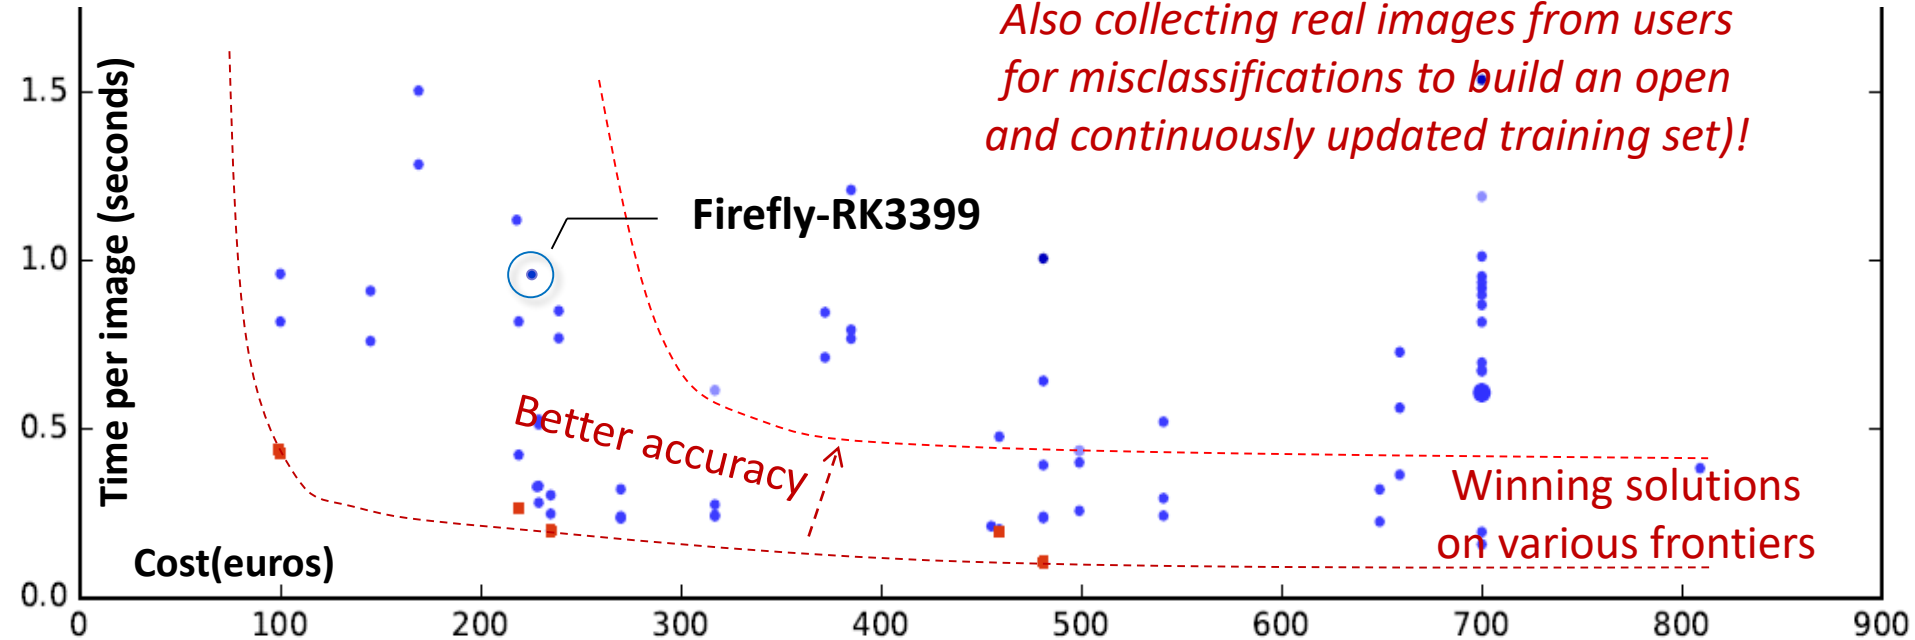
The number of distinct GPUs: **110+**

The number of distinct OS: **280+**

Power range: **1-10W**

No need for a dedicated and expensive cloud –  
volunteers help us validate research ideas  
similar to SETI@HOME

*Also collecting real images from users  
for misclassifications to build an open  
and continuously updated training set!*



## CK helped to automate MLPerf.org inference submissions

“**MobileNets**: Efficient Convolutional Neural Networks for Mobile Vision Applications” (Andrew G. Howard et al., 2017, <https://arxiv.org/abs/1704.04861>):

- Parameterised CNN family using depthwise separable convolutions.
- Channel multiplier: 1.00, 0.75, 0.50, 0.25 - marker shape (see below).
- Input image resolution: 224, 192, 160, 128 - marker size.

[CodeReef.ai/portal/search/?q="SOTA+MLPerf"](https://code-reef.ai/portal/search/?q=)



A broad ML benchmark suite for measuring performance of ML software frameworks, ML hardware accelerators, and ML cloud platforms.

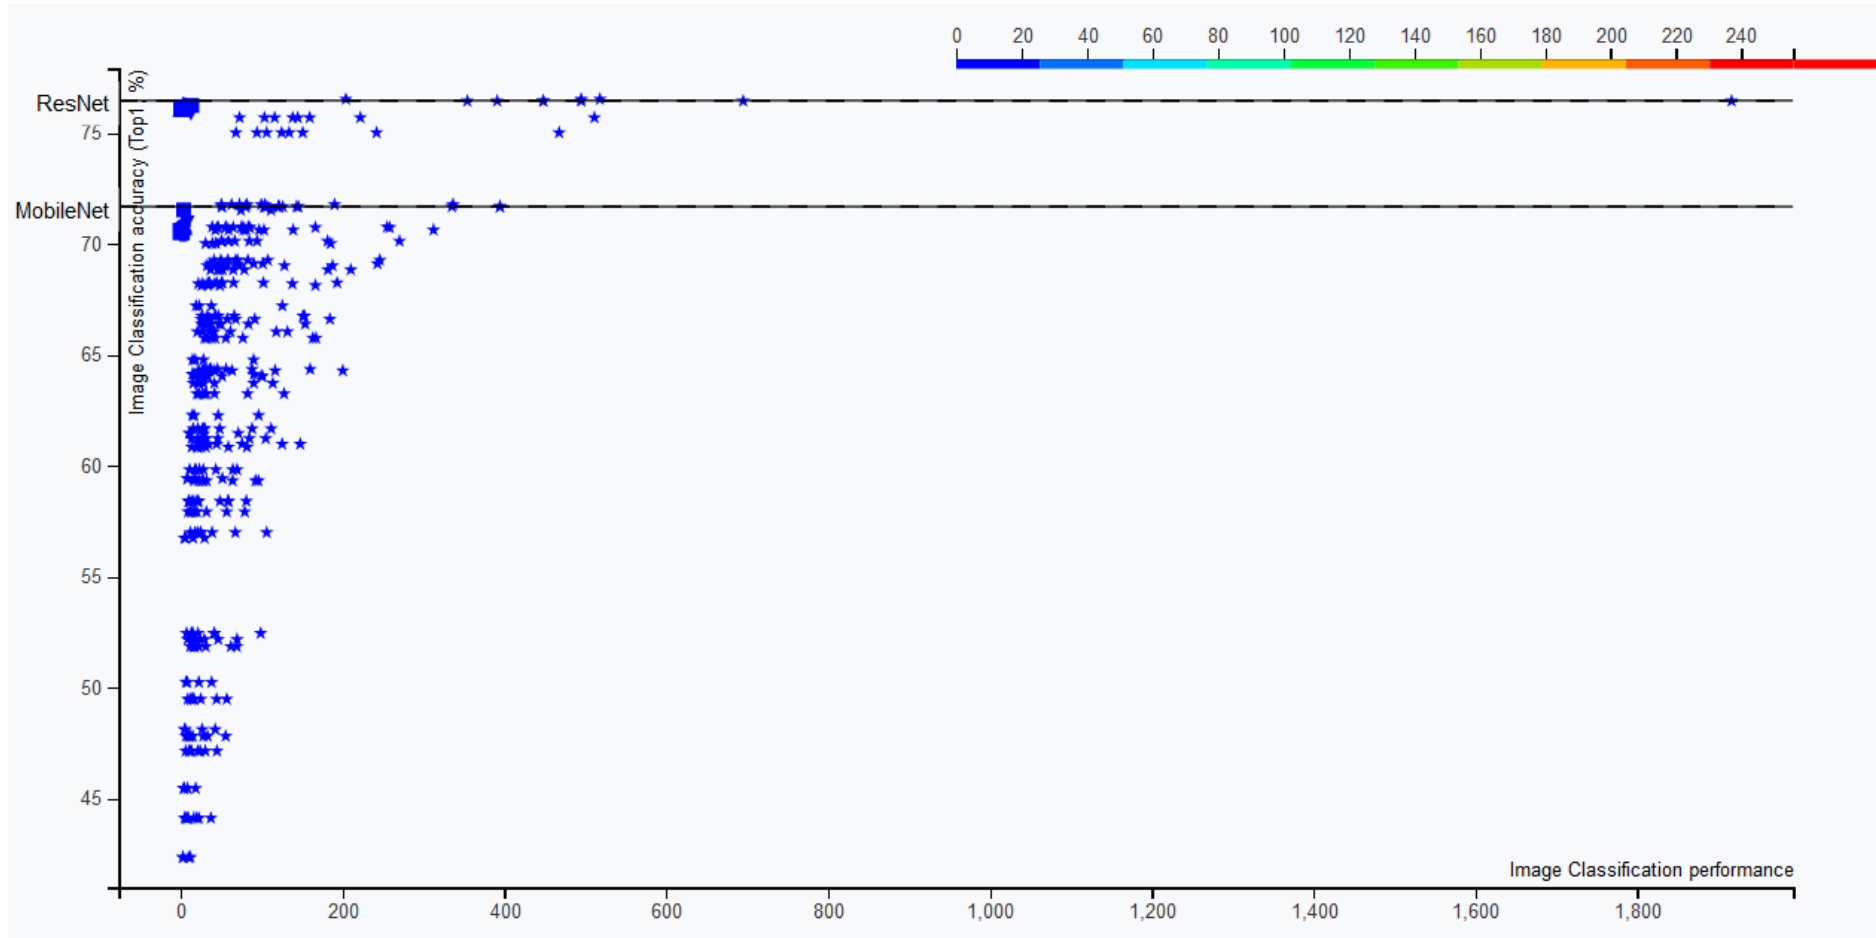
It is even possible to test how object detection from MLPerf works live:

[CodeReef.ai/portal/c/cr-solution/demo-obj-detection-coco-tf-cpu-webcam-linux-azure](https://code-reef.ai/portal/c/cr-solution/demo-obj-detection-coco-tf-cpu-webcam-linux-azure)

# Use CK to autotune MobileNets across diverse devices for MLPerf submissions

The [MLPerf](#) consortium has recently released over 500 validated [inference benchmarking results](#) from 14 organizations (including DellEMC, Nvidia, Google, Intel, Alibaba, Habana) measuring how fast and how well a pre-trained computer system can classify images, detect objects, and translate sentences.

Over 400 of these results were automated with the CK framework and CK workflows.



[cknowledge.org/dashboard](https://cknowledge.org/dashboard)

[mlperf.org](https://mlperf.org)

## Current state of CK


- CK is a promising open-source technology already used by companies and universities in production but there is still a lot to be improved!

downloads 140k pypi package 1.12.2 python 2.7 | 3.4+ License BSD 3-Clause

- Current major issues preventing further adoption:
  - CLI and JSON meta is not user friendly (similar to Git)
  - Distributed nature of CK makes it difficult to understand who is using CK and ensure the stability/testing of workflows
  - Lack of an open portal to exchange stable components (similar to PyPi)
  - Lack of automatic testing of all components and workflows
- Currently supported by my non-profit cTuning foundation (cTuning.org) but our resources are very limited
- Want to attract more companies to improve our open-source technology together

## An open CodeReef.ai portal to solve above issues

Based on the feedback from the CK users I am now working on an open and free portal to solve above problems, continue supporting our open-source developments, and enable practical and portable MLOps



[CodeReef.ai/portal](https://CodeReef.ai/portal)  
[CodeReef.ai/static/docs](https://CodeReef.ai/static/docs)

Reproduced papers with portable CK workflows:

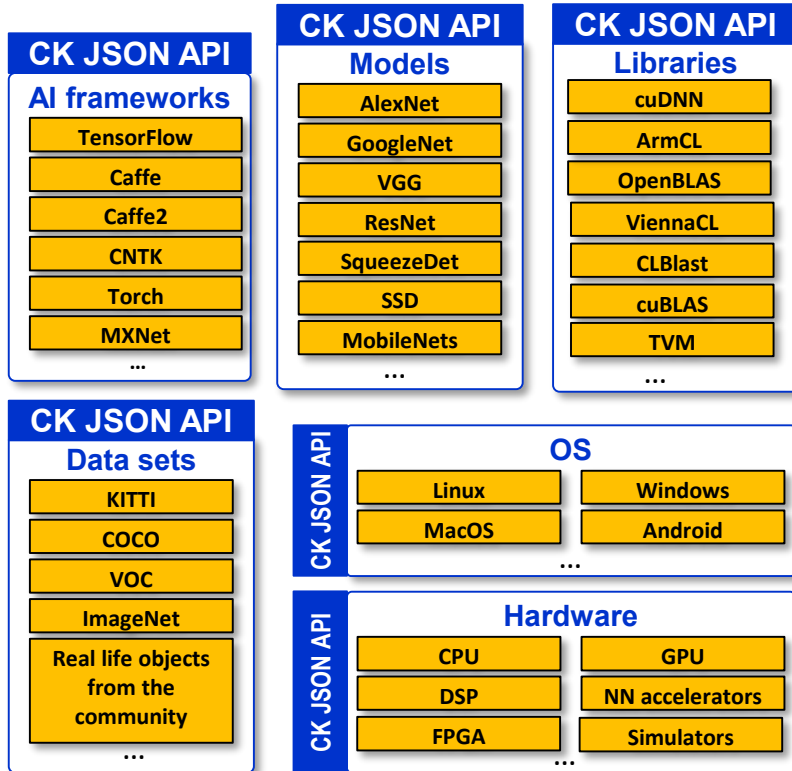
[CodeReef.ai/portal/search/?q=\"portable-workflow-ck\"](https://CodeReef.ai/portal/search/?q=\)

We will present CodeReef+CK for MLOps at the following events:

- NeurIPS MLSys'20 (Austin), [mlsys.org](https://mlsys.org)
- AI hardware summit (Munich), [www.kisacoresearch.com/events/ai-hardware-summit-europe](https://www.kisacoresearch.com/events/ai-hardware-summit-europe)
- ACM ASPLOS (Lausanne), [asplos-conference.org](https://asplos-conference.org)
- IEEE ISPASS (Boston), [www.ispass.org/ispass2020](https://www.ispass.org/ispass2020)

# CodeReef.ai: an open platform for portable AI/ML benchmarking and MLOps

Repository of customizable, portable and reusable research components with CK API



Share complete workflows along with published papers to automate artifact evaluation and help the community build upon prior work

Crowdsource experiments, benchmarking and co-design with the help of volunteers across diverse models, data sets and platforms




Present best results, AI/ML workflows and components on a live scoreboard for fair comparison and reuse

Help researchers learn ML/AI techniques, quickly prototype new ones, validate them in practice with companies, and even contribute back new research components

Customizable CK workflows for real-world AI/ML tasks

Assemble scenarios such as image classification as LEGO™



Collaboratively benchmark and co-design SW/HW stacks, select the most efficient AI/ML solutions, reduce development costs, and increase productivity

[CodeReef.ai/portal](https://CodeReef.ai/portal)



Thank you for your interest!

Contact

[Grigori.Fursin@cTuning.org](mailto:Grigori.Fursin@cTuning.org) or [Grigori.Fursin@CodeReef.ai](mailto:Grigori.Fursin@CodeReef.ai)

# LF AI Updates

# Outreach Committee

# Recent Announcements

1. **New Member Welcome - RStudio + inwinSTACK + ISSIP:**  
<https://lfai.foundation/blog/2020/01/30/lf-ai-foundation-new-member-welcome-rstudio-inwinstack-issip/>
2. **Sparklyr New Project:**  
<https://lfai.foundation/blog/2020/01/29/sparklyr-joins-lf-ai-as-its-newest-incubation-project/>
3. **LF AI 2019 Year in Review:**  
<https://lfai.foundation/blog/2020/01/22/lf-ai-2019-year-in-review/>
4. **Horovod 0.19.0 Release:**  
<https://lfai.foundation/blog/2020/01/14/horovod-version-0-19-0-now-available/>
5. **Angel Graduation:**  
<https://lfai.foundation/blog/2019/12/19/lf-ai-foundation-announces-graduation-of-angel-project/>
6. **Zilliz new premier member:**  
<https://lfai.foundation/blog/2019/12/17/lf-ai-welcomes-zilliz/>
7. **LF AI Day Shanghai Summary:**  
<https://lfai.foundation/blog/2019/12/05/thank-you-lf-ai-day-shanghai-summary/>
8. **LF AI Receives Contribution Award from CAAI:**  
<https://lfai.foundation/blog/2019/11/27/lf-ai-receives-best-contribution-award/>
9. **Acumos Clio Release:**  
<https://lfai.foundation/press-release/2019/11/26/lf-ai-delivers-acumos-ai-clio-release/>
10. **Pyro 1.0.0 Release:** <https://lfai.foundation/blog/2019/11/18/pyro-1-0-has-arrived/>

# Upcoming Announcements

## March:

Acumos and Angel Collaboration  
Marquez New Project  
Milvus New Project  
Angel Roadmap + Project Participation Invite  
LF AI 2 Year Anniversary  
ITU Global Challenge on AI/ML in 5G Networks Promotion

## April:

Q2 New Member Announcements  
IEEE Infrastructure Conference - LF AI Participation  
Open Networking & Edge Summit

# LF AI Events

**Event Leads Needed - Please Encourage Member Participation Across Your Organizations**

Events Page on LF AI Website:

<https://lfai.foundation/events/>

Events Calendar on LF AI Wiki (subscribe available):

<https://wiki.lfai.foundation/pages/viewpage.action?pageId=12091544>

2020 Events Wiki:

<https://wiki.lfai.foundation/display/DL/2020+Events>

**The [Outreach Committee](#) and LF AI Community need to identify Event Leads to drive these events as the Foundation is not resources/budgeted to manage them all. Please encourage participation within your organizations. LF AI provides [event planning templates/resources](#) and is available to help with guidance and questions along the way. Please email the Jessica Kim (Outreach Committee Chair) or Jacqueline/Ibrahim to discuss participation.**

# Call to Participate in Ongoing Efforts

# Trusted AI

- › **Leadership:**  
Animesh Singh (IBM), Souad Ouali (Orange), and Jeff Cao (Tencent)
- › **Goal:** Create policies, guidelines, tooling and use cases by industry
- › **Github:**  
<https://github.com/lfai/trusted-ai>
- › **Wiki:**  
<https://wiki.lfai.foundation/display/DL/Trusted+AI+Committee>
- › **To participate:**  
<https://lists.lfai.foundation/g/trustedai-committee/>
- › **Next call:** Bi-weekly on Thursdays at 7am PT, subscribe to group calendar on wiki  
<https://wiki.lfai.foundation/pages/viewpage.action?pageId=12091895>

# ML Workflow

- › **Leadership:**  
Ofer Hermoni
- › **Goal:**  
Define an ML Workflow and promote cross project integration
- › **Wiki:**  
<https://wiki.lfai.foundation/display/DL/ML+Workflow+Committee>
- › **To participate:**  
<https://lists.lfai.foundation/g/mlworkflow-committee>
- › **Next call:** Bi-weekly on Thursdays at 7:00 am PT, subscribe to group calendar on wiki  
<https://wiki.lfai.foundation/pages/viewpage.action?pageId=18481242>

# Upcoming TAC Meetings



# Upcoming TAC Meetings

**March 12:** Incubation Project Proposal + TAC Vote: NNStreamer from Samsung

**March 26:** TBD

# TAC Meeting Details

- › To subscribe to the TAC Group Calendar, visit the wiki: <https://wiki.lfai.foundation/x/XQB2>
- › Join from PC, Mac, Linux, iOS or Android: <https://zoom.us/j/430697670>
- › Or iPhone one-tap:
  - › US: +16465588656,,430697670# or +16699006833,,430697670#
- › Or Telephone:
  - › Dial(for higher quality, dial a number based on your current location):
  - › US: +1 646 558 8656 or +1 669 900 6833 or +1 855 880 1246 (Toll Free) or +1 877 369 0926 (Toll Free)
- › Meeting ID: 430 697 670
- › International numbers available: <https://zoom.us/u/achYtcw7uN>

# Open Discussion

# Legal Notices

- › The Linux Foundation, The Linux Foundation logos, and other marks that may be used herein are owned by The Linux Foundation or its affiliated entities, and are subject to The Linux Foundation's Trademark Usage Policy at <https://www.linuxfoundation.org/trademark-usage>, as may be modified from time to time.
- › Linux is a registered trademark of Linus Torvalds. Please see the Linux Mark Institute's trademark usage page at <https://lmi.linuxfoundation.org> for details regarding use of this trademark.
- › Some marks that may be used herein are owned by projects operating as separately incorporated entities managed by The Linux Foundation, and have their own trademarks, policies and usage guidelines.
- › TWITTER, TWEET, RETWEET and the Twitter logo are trademarks of Twitter, Inc. or its affiliates.
- › Facebook and the "f" logo are trademarks of Facebook or its affiliates.
- › LinkedIn, the LinkedIn logo, the IN logo and InMail are registered trademarks or trademarks of LinkedIn Corporation and its affiliates in the United States and/or other countries.
- › YouTube and the YouTube icon are trademarks of YouTube or its affiliates.
- › All other trademarks are the property of their respective owners. Use of such marks herein does not represent affiliation with or authorization, sponsorship or approval by such owners unless otherwise expressly specified.
- › The Linux Foundation is subject to other policies, including without limitation its Privacy Policy at <https://www.linuxfoundation.org/privacy> and its Antitrust Policy at <https://www.linuxfoundation.org/antitrust-policy>, each as may be modified from time to time. More information about The Linux Foundation's policies is available at <https://www.linuxfoundation.org>.
- › Please email [legal@linuxfoundation.org](mailto:legal@linuxfoundation.org) with any questions about The Linux Foundation's policies or the notices set forth on this slide.