



CAVIA: Fast Context Adaptation via Meta-Learning

Luisa Zintgraf



Kyriacos Shiarlis



Vitaly Kurin



Katja Hofmann




Shimon Whiteson




Meta-Learning for Fast Adaptation:

- Learn how to map x to y , on new tasks, fast and with little data.



Meta-Learning for Fast Adaptation:

- Learn how to map x to y , on new tasks, fast and with little data.



- MAML (Finn et al. 2017): learns model initialisation st. new task can be learned within a few gradient updates

CAVIA: Fast **C**ontext **A**daptation **via** Meta-Learning

Less prone to overfitting
compared to MAML

Interpretable

Easy to parallelise

Many tasks /
benchmarks only require
task identification

Updating all model parameters at test time isn't necessary!


- Many tasks and current benchmarks only require **task identification**.
- Many parameters + few data points can lead to **overfitting**.

Updating all model parameters at test time isn't necessary!


- Many tasks and current benchmarks only require **task identification**.
- Many parameters + few data points can lead to **overfitting**.

Solution: Adapt only part of the network!

MAML (Finn et al. 2017)



CAVIA




Context parameters:

Task-specific *input parameter vector*.

Updated at test time via gradient descent.

Represents task embedding.




Context parameters:

Task-specific *input parameter vector*..

Updated at test time via gradient descent.

Represents task embedding.





Network parameters:

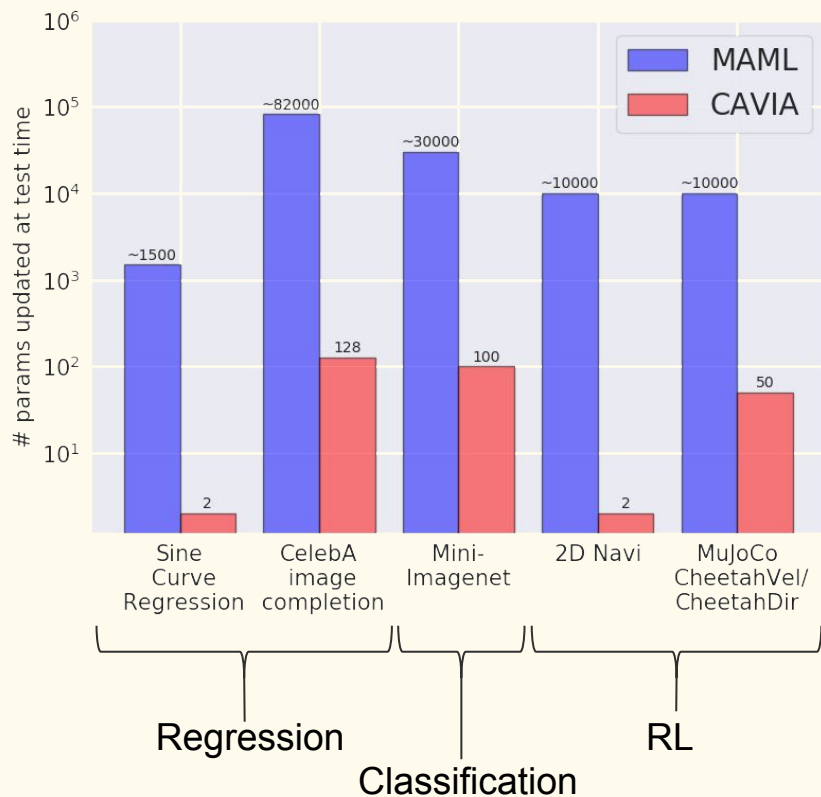
Shared across tasks.

Meta-trained, fixed at test time.

CAVIA

Fast **C**ontext **A**daptation **via** Meta-Learning

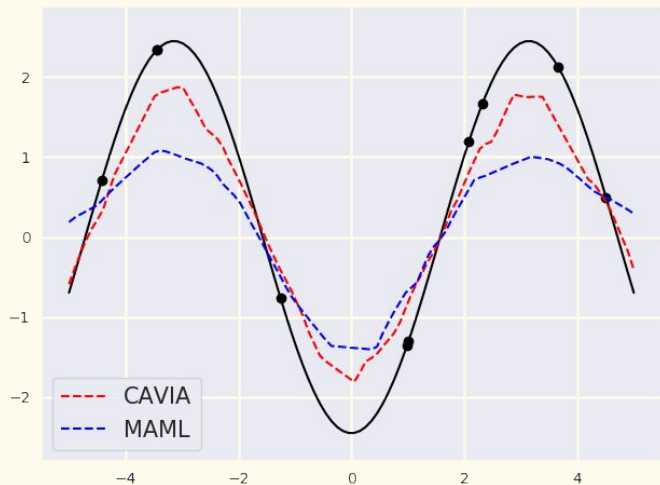




Many tasks / benchmarks only require **task identification**

CAVIA: Sine Curve Experiments

Task defined by amplitude + phase



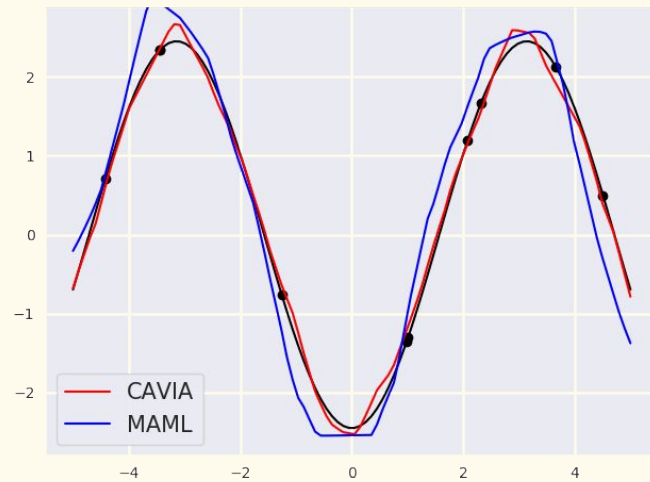
Before update

MAML: ~ 1500
params



CAVIA: 2 params

Gradient update




After update

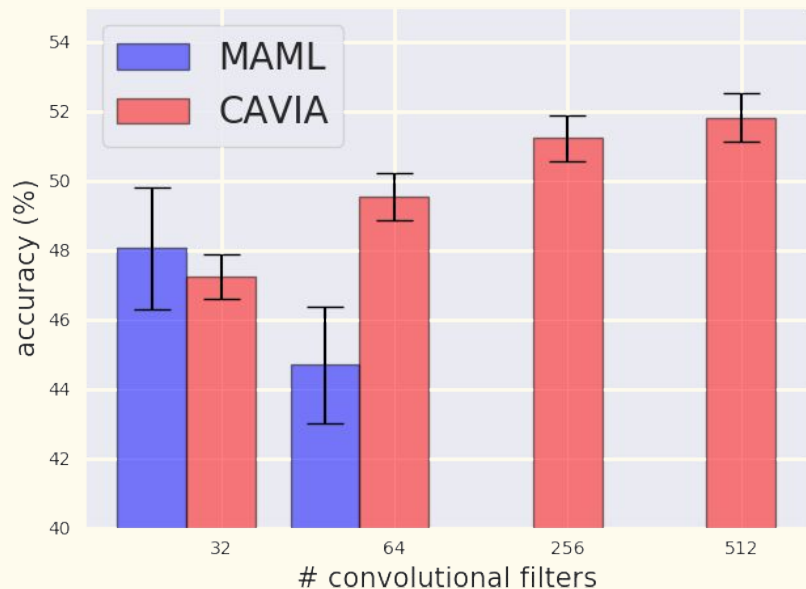
Could be re-used for
related / auxiliary tasks,
or distributed systems

Could be re-used for
related / auxiliary tasks,
or distributed systems

Context parameters
are interpretable



CAVIA: Mini-Imagenet Experiments



Model parameters: > 30,000

Context parameters: 100

Less prone to overfitting
compared to MAML

CAVIA: Fast Context Adaptation via Meta-Learning

Code: github.com/lmzintgraf/cavia
Blog: whirl.cs.ox.ac.uk/blog/cavia
Contact: luisa.zintgraf@cs.ox.ac.uk
[@luisa_zintgraf](https://twitter.com/luisa_zintgraf)



Luisa Zintgraf



Kyriacos Shiarlis



Vitaly Kurin



Katja Hofmann



Shimon Whiteson

