Deep Neural Networks for Safety-Critical Applications: Vision and Open Problems

Daniel Casini, Alessandro Biondi, Giorgio Buttazzo

ReTiS Lab, Scuola Superiore Sant'Anna, Pisa, Italy

Sant'Anna
Scuola Universitaria Superiore Pisa

Real-Time Systems Laboratory

Motivations

1

Currently, many car manufacturers are tackling the race towards autonomous cars

2

Huge improvements in DNN accuracy for many tasks (e.g. image classification)

Motivations

1

Currently, many car manufacturers are tackling the race towards autonomous cars

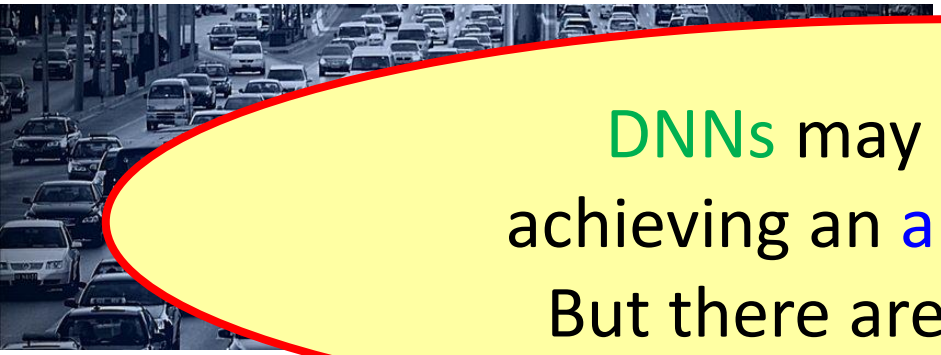

DNNs may be useful for achieving an autonomous car!
But there are many issues…

2

Huge improvements in DNN accuracy for many tasks (e.g. image classification)

Motivations

But not only autonomous driving…

- DNNs can be also adopted for other types of autonomous systems (e.g., robotics, industrial control)

Neural Networks

- Used to solve problems that are difficult to formalize by a set of rules.

Problems that depend on too many details are learned by direct experience

- Neural Networks imitate the way our brain works

Neural Networks

- Neural Networks consist of a set of neurons, often organized into layers
- Neurons are connected to each other by synaptic weights
- They are used for many different purposes, as speech recognition, image processing, wheather forecast, etc.

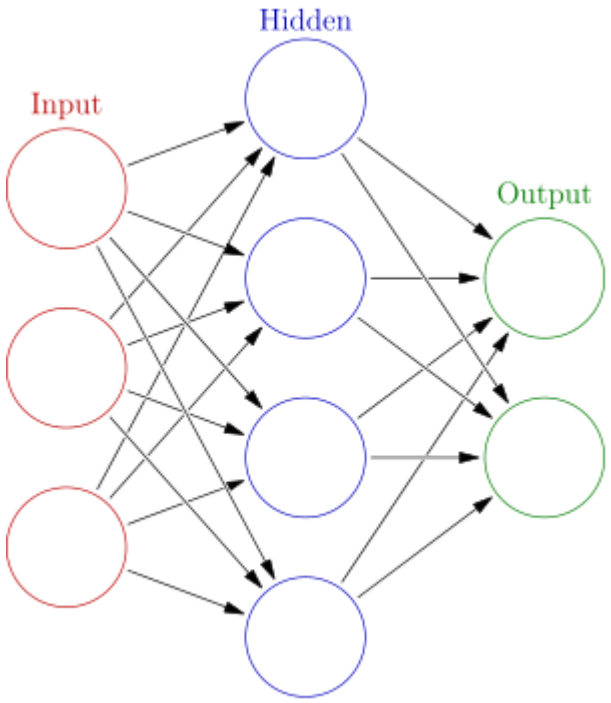

Supervised Learning

- They can learn "by examples" to associate input-output pairs

- Example: Before training

Input	Desired Output	Actual Output
	Red Semaphore	Car

Incorrect Classification!

Supervised Learning

- The training algorithm regulates the internal parameters (i.e., weights) of the network for producing the expected output

- Example: After training

Input	Desired Output	Actual Output
	Red Semaphore	Red Semaphore

Correct!

Image Recognition

- The ILSVRC Challenge is a competition held from 2010 in which networks compete in classifying objects from images to labels, with 1000 possible categories

Training set: 1.2 million images (1,000 categories)

Test set: 150,000 images

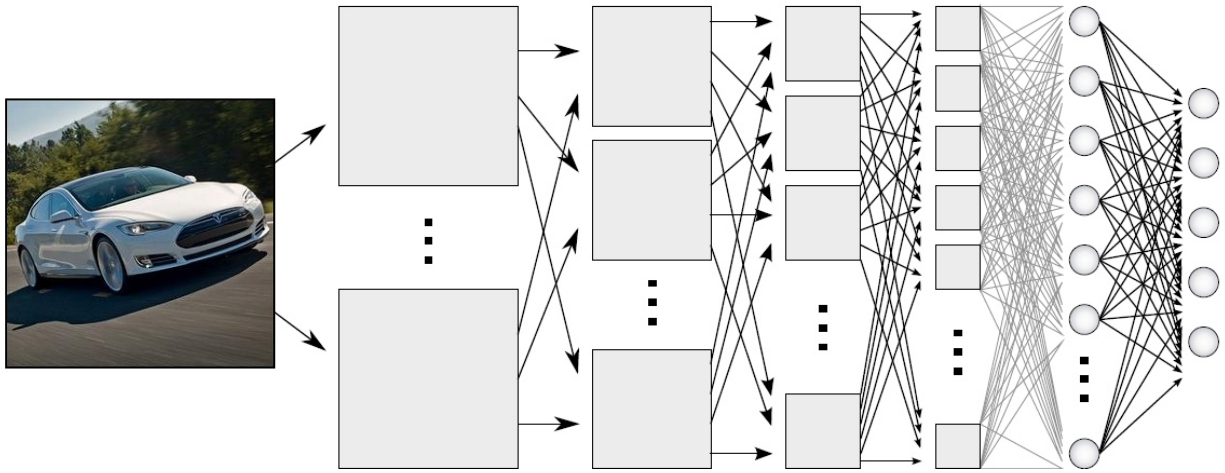

Are DNNs good enough?

The winning network of 2017 (SENet), achieved an accuracy of 97.74%

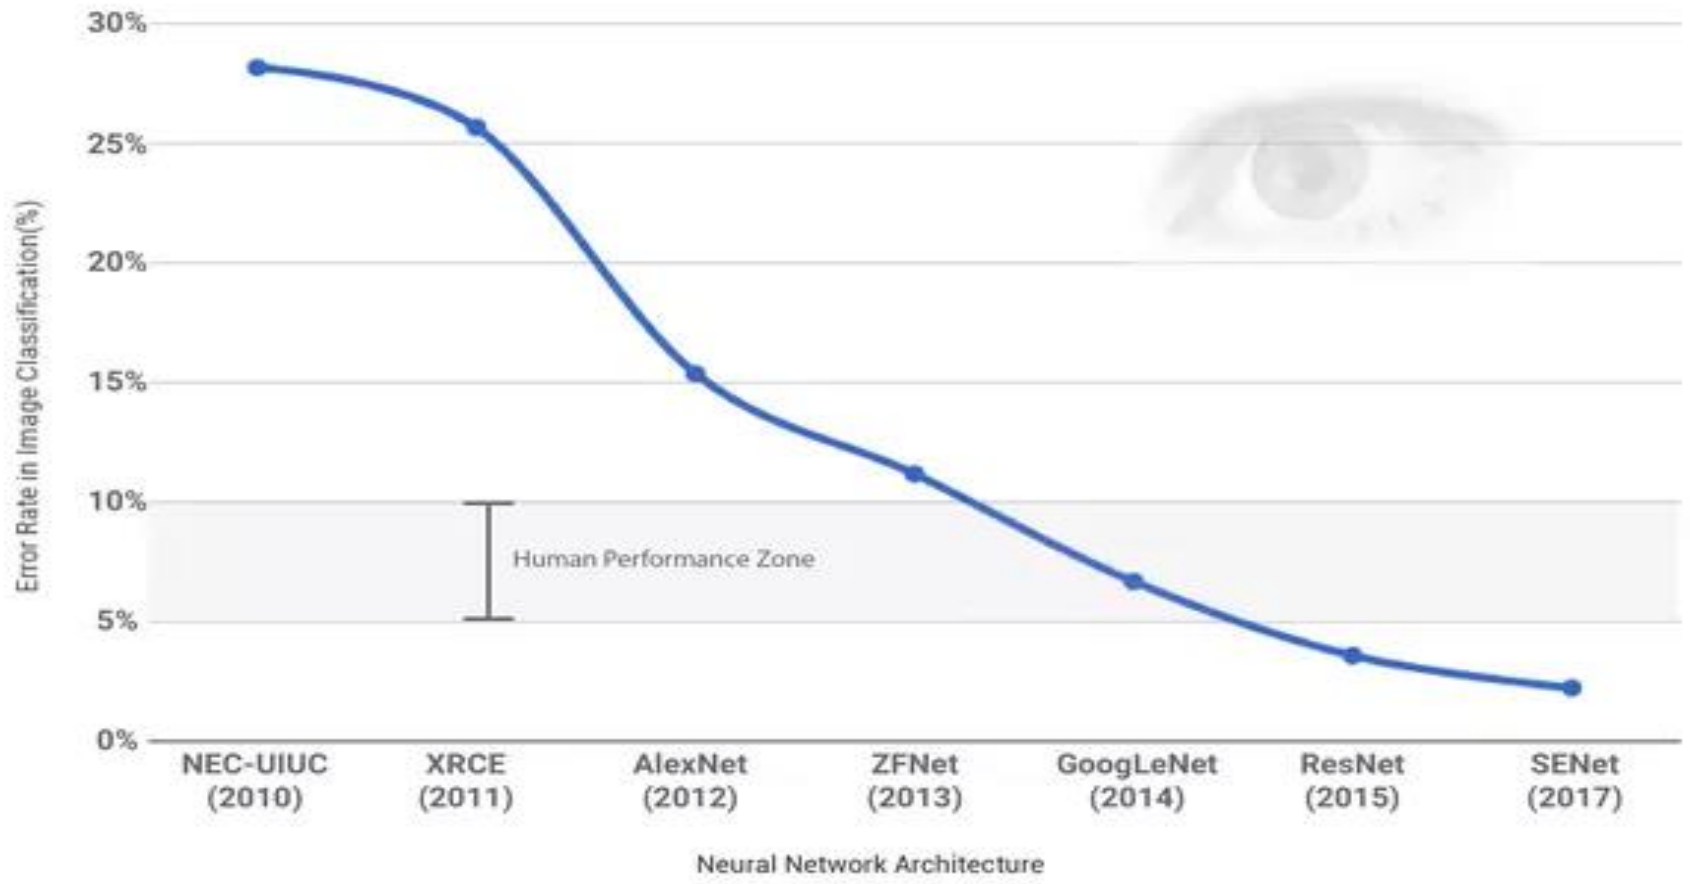

Source: http://blog.paralleldots.com/data-science/must-read-path-breaking-papers-about-image-classification/

Deep Neural Networks in Safety Critical Scenarios:

1. Certification Issues

Certification Issues

- Deep Neural Networks do not have a well-defined behavior

- Their results are difficult to be replicated (e.g., changing few pixels of an image may lead to different results)

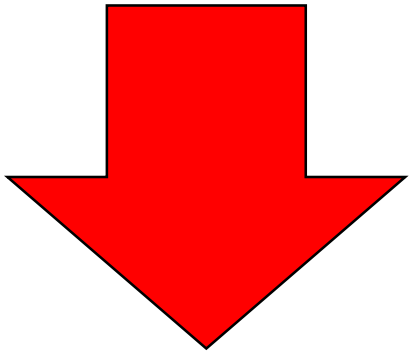

Huge problem for certification!

Certification Issues

- Deep Neural Networks do not have a well-defined behavior

- Their results are difficult to be replicated (e.g., changing few pixels of an image may lead to different results)

How to support DNNs in a safety-critical context to build a system that can be certified?

Huge problem for certification!

Hint of Solution

IDEA: Match each DNN with a corresponding algorithm based on hard computing (e.g., a convolution filter) to monitor their behavior and redirect the actuation to safe actions in case of detected misbehavior

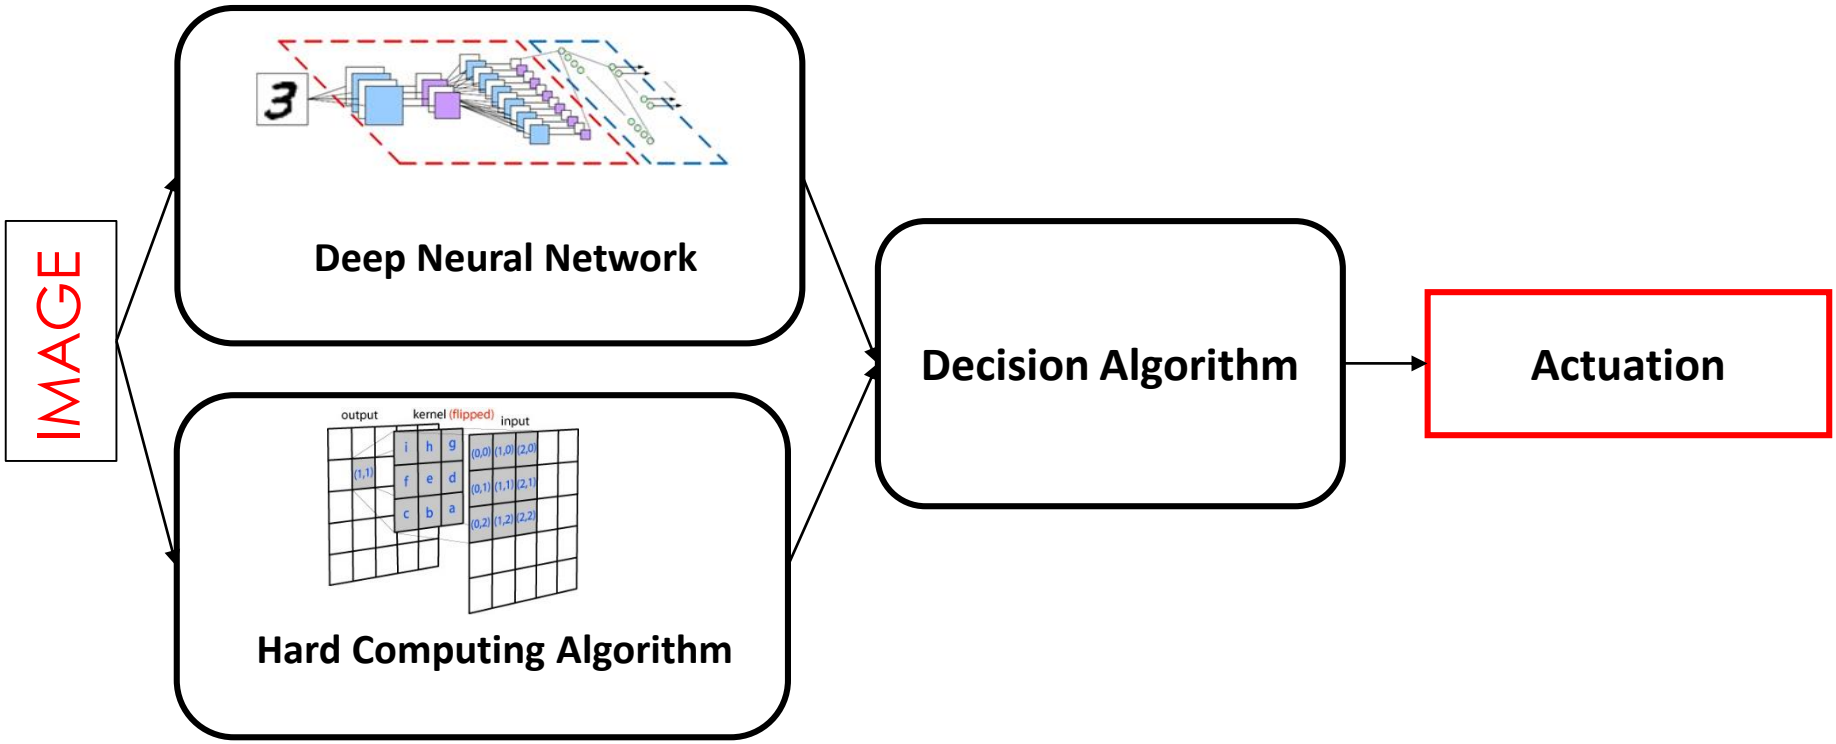

Example of safe action

IDEA: Match each DNN with a corresponding algorithm based on hard computing (e.g., a convolution filter) to monitor their behavior and redirect the actuation to safe actions in case of detected misbehavior

Example of safe action

IDEA: Match each DNN with a corresponding algorithm based on hard computing (e.g., a convolution filter) to monitor their behavior and redirect the actuation to safe actions in case of detected misbehavior

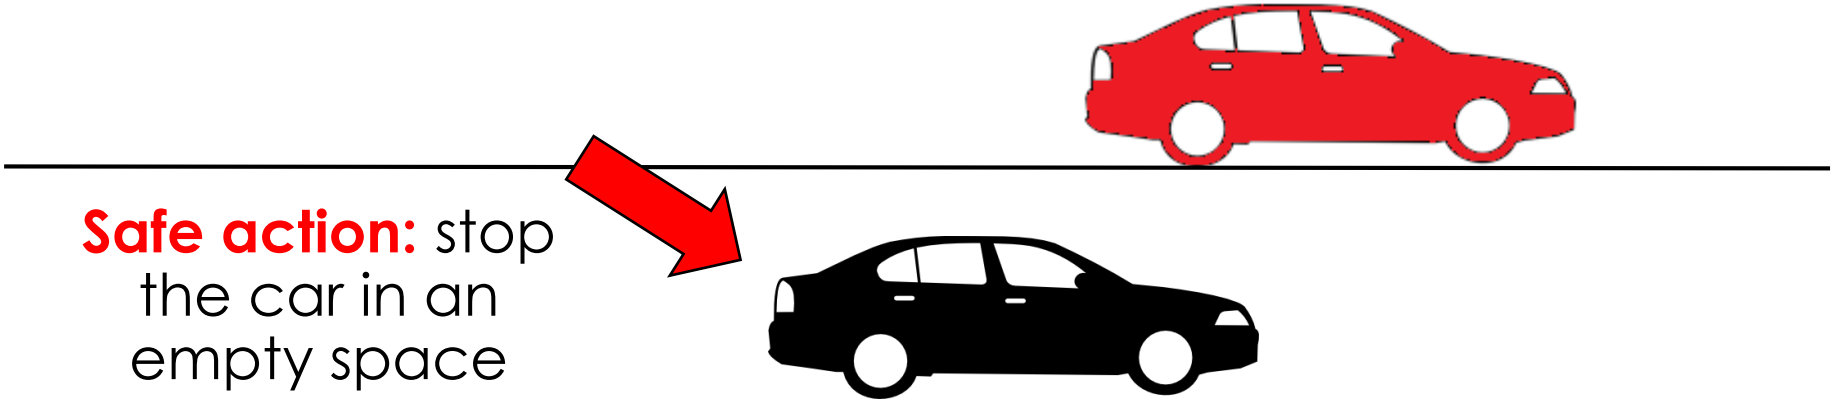

Safe action: stop the car in an empty space

Redundant Neural Networks

The average-case behavior can be improved by inserting redundant neural networks, based on different models or trained with a different algorithm.

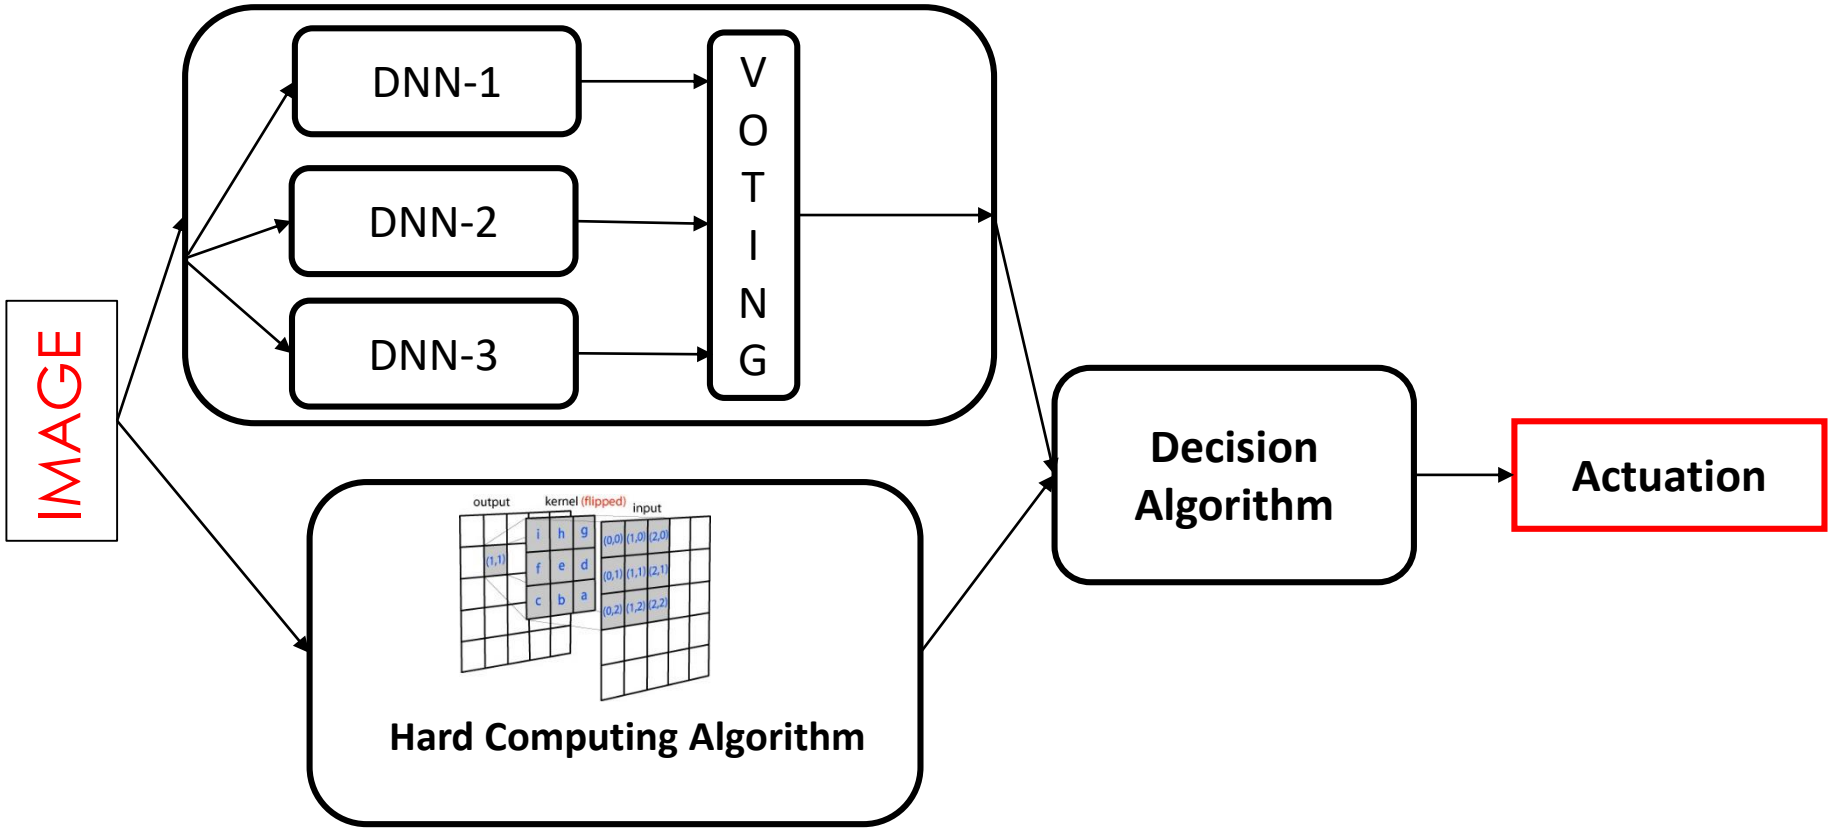

Redundant Neural Networks

The average-case behavior can be improved by inserting redundant neural networks, based on different models or trained with a different algorithm.

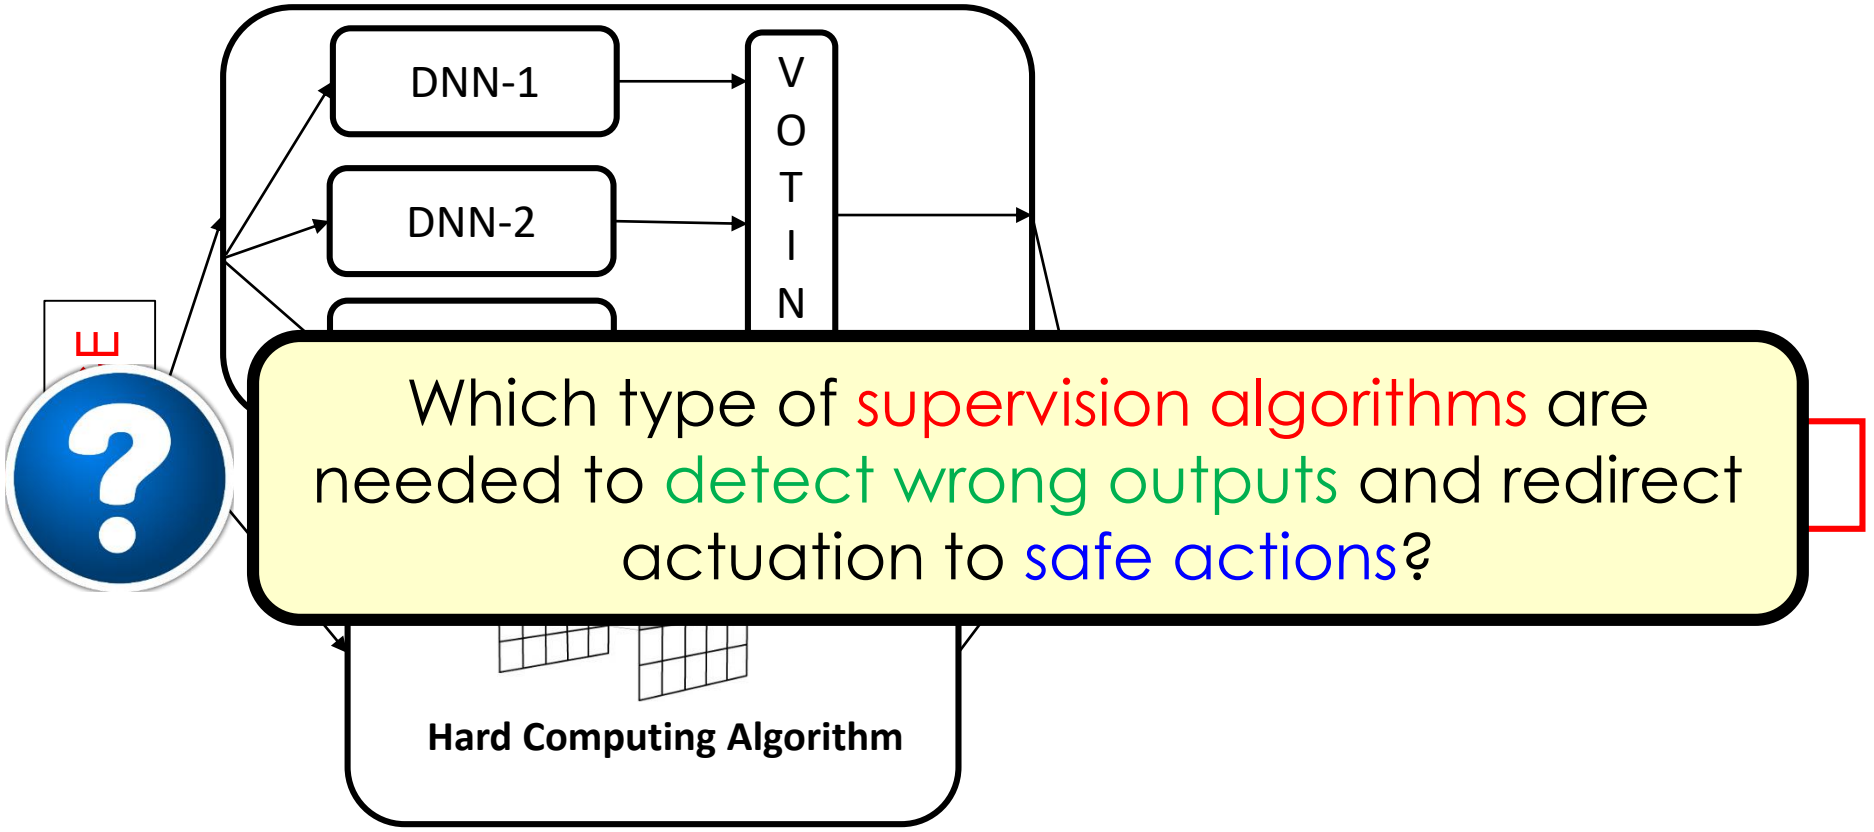

Which type of supervision algorithms are needed to detect wrong outputs and redirect actuation to safe actions?

Deep Neural Networks in Safety Critical Scenarios:

2. Security and Isolation

Security and Isolation

- A DNN is a complex software, exposed to security threats

- What if an attacker exploits the weakness of a DNN to take control of the steering system?

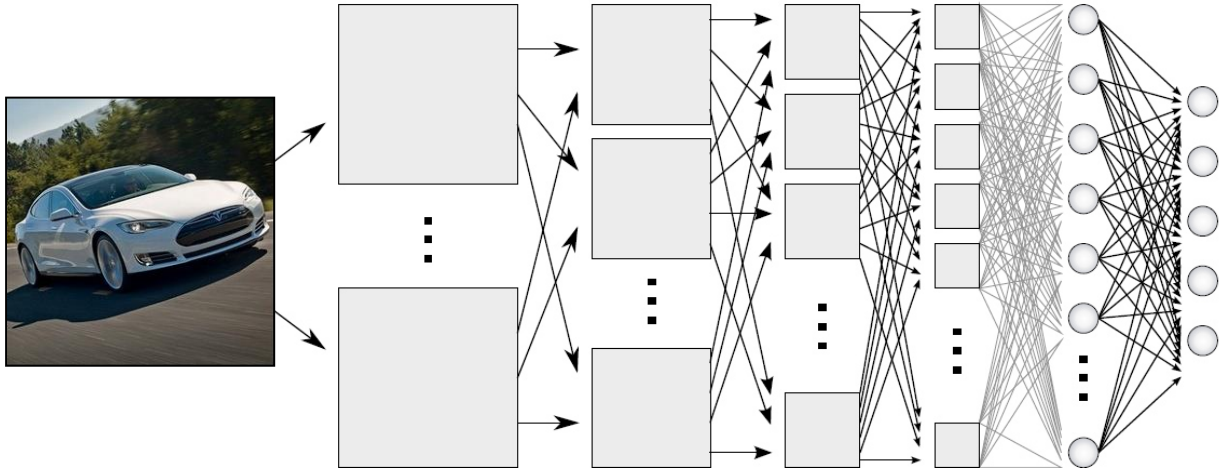

Security and Isolation

- A DNN is a complex software, exposed to security threats

- What if an attacker exploits the weakness of a DNN to take control of the steering system?

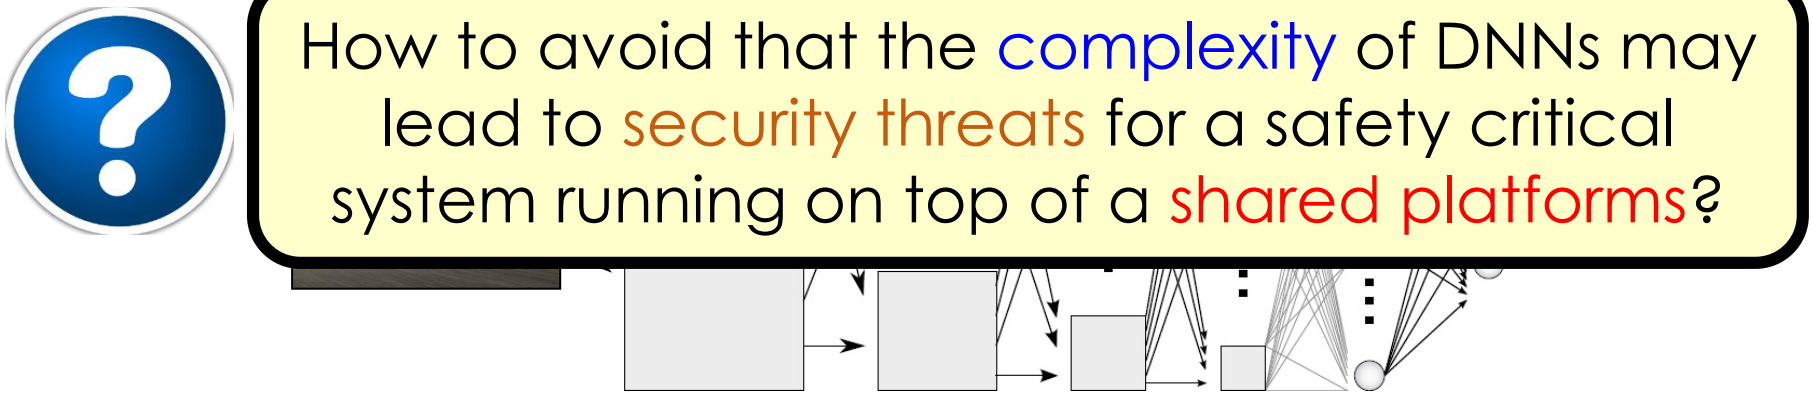

How to avoid that the complexity of DNNs may lead to security threats for a safety critical system running on top of a shared platforms?

A hypervisor-based solution

IDEA: Divide a multicore heterogeneous platform in two domains

Prone to attacks and malfunctioning

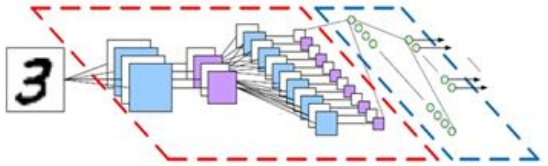

Deep Neural Networks

TensorFlow™ Caffe

Safety-critical

AUTOSAR

Hypervisor

Multicore Heterogeneous Platform

A hypervisor-based solution

IDEA: Divide a multicore heterogeneous platform in two domains

Prone to attacks and malfunctioning

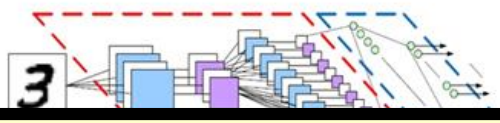

Which mechanisms have to be provided to allow them interacting while running on different OSes?

Hypervisor

Multicore Heterogeneous Platform

Deep Neural Networks in Safety Critical Scenarios:

3. Predictability

Predictability

KEY ISSUE: Guaranteeing that a real-time workload composed of DNNs is schedulable

- Focus on the inference phase only

- A DNN is composed of a pipeline of layers, where each one implements an operation

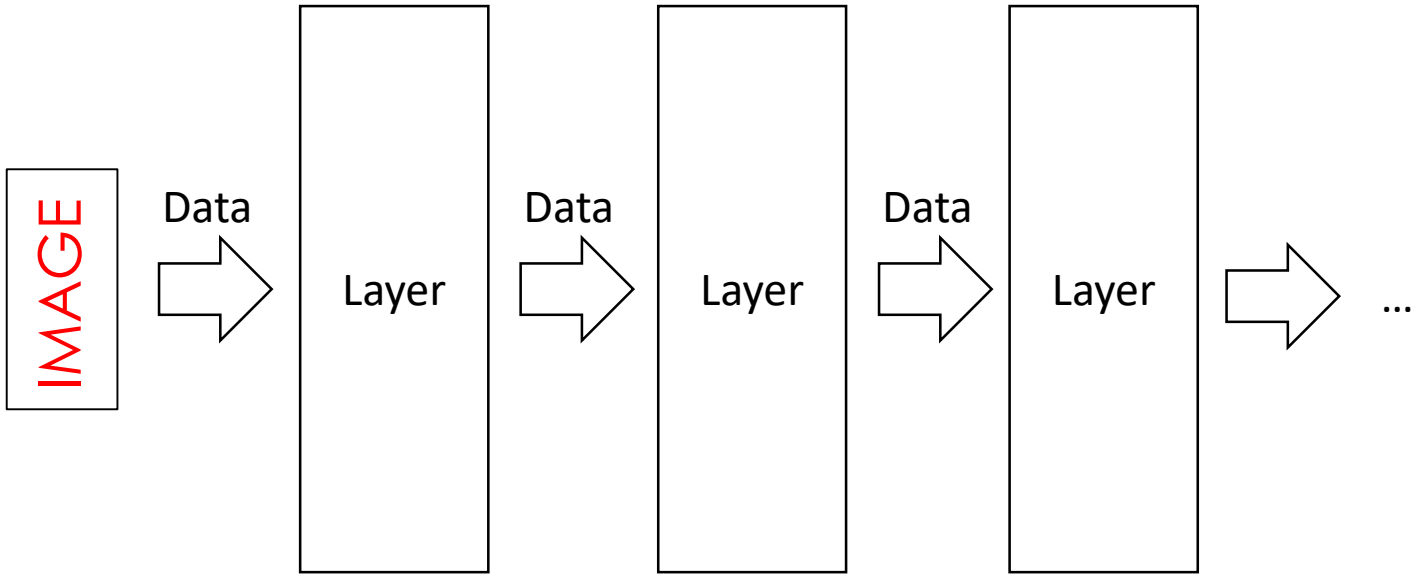

Predictability

- Many inference frameworks furtherly parallelize each layer

- The resulting computational activity can be represented by a Direct Acyclic Graph (DAG)

- A properly defined task model should also account for tensors (i.e., memory) exchanged among nodes

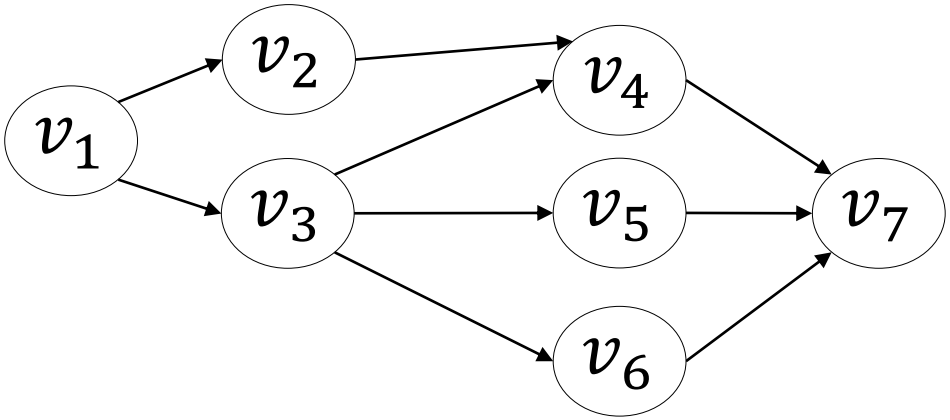

Predictability

- Many inference frameworks furtherly parallelize each layer
- The resulting computational activity can be represented by a Direct Acyclic Graph (DAG)
- A properly defined
  exchanged among
  nodes

Which is a suitable task model to describe and analyze the temporal behavior of DNNs?

Heterogeneous Platforms

- Timing analysis should also account for the heterogeneity of the underlying hardware platform
- DNN execution on FPGA is not yet fully supported by inference engines

  Dynamic partial reconfiguration can be exploited for accelerating complex layers

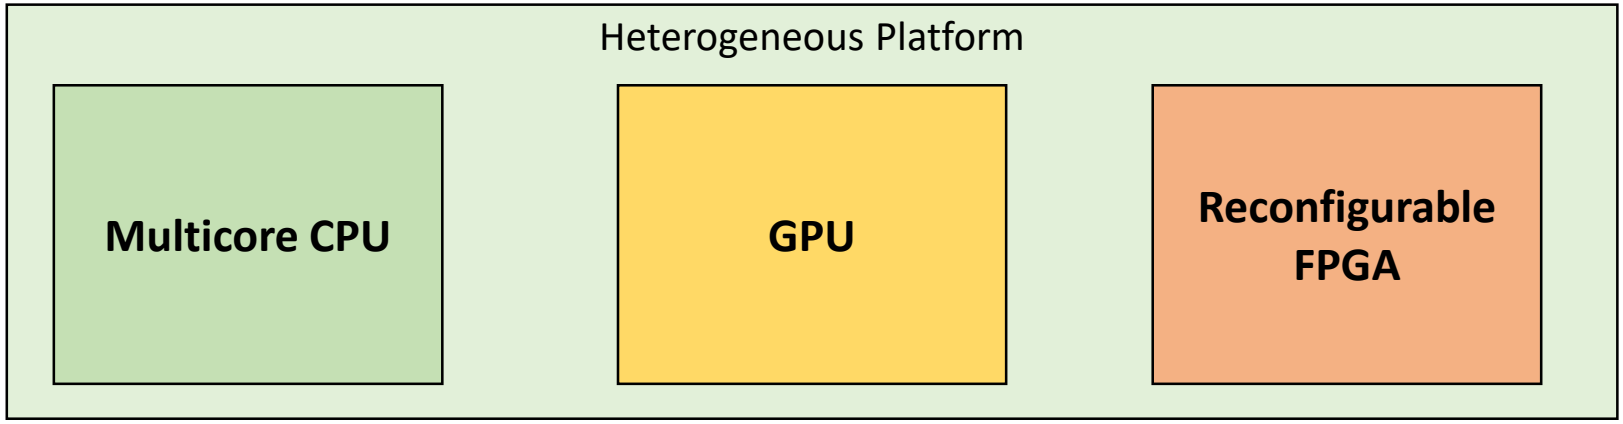

Heterogeneous Platforms

- Recently, ad hoc application specific integrated circuits have been recently produced (e.g., the Tensor Processing Unit by Google)

- They can be included in commercial heterogeneous platforms soon

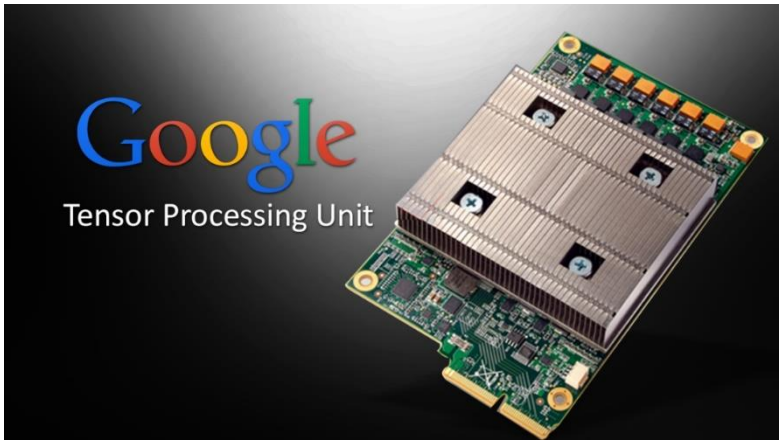

Heterogeneous Platforms

- Recently, ad hoc application specific integrated circuits have been recently produced (e.g., the Tensor Processing Unit by Google)

How to account for novel (highly heterogenous) computing platforms?

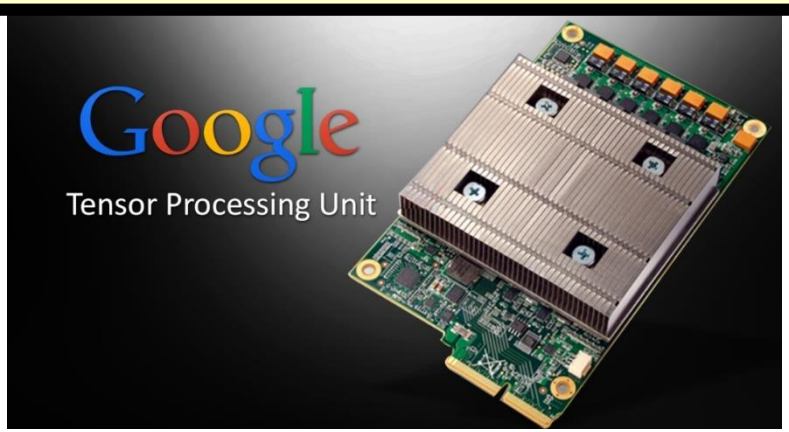

Inference Engines

- DNNs are typically executed by means of inference engines
  - Inference engines can affect the execution of DNNs

Caffe

TensorFlow™

Google

theano

torch

Inference Engines

- DNNs are typically executed by means of inference engines
  - Inference engines can affect the execution of DNNs

How to account for inference engines that affect the DNN's execution?

Caffe

TensorFlow™

theano

torch

How Tensorflow works on CPUs?

- TensorFlow assigns ready nodes to threads of a thread pool (of globally-scheduled threads)

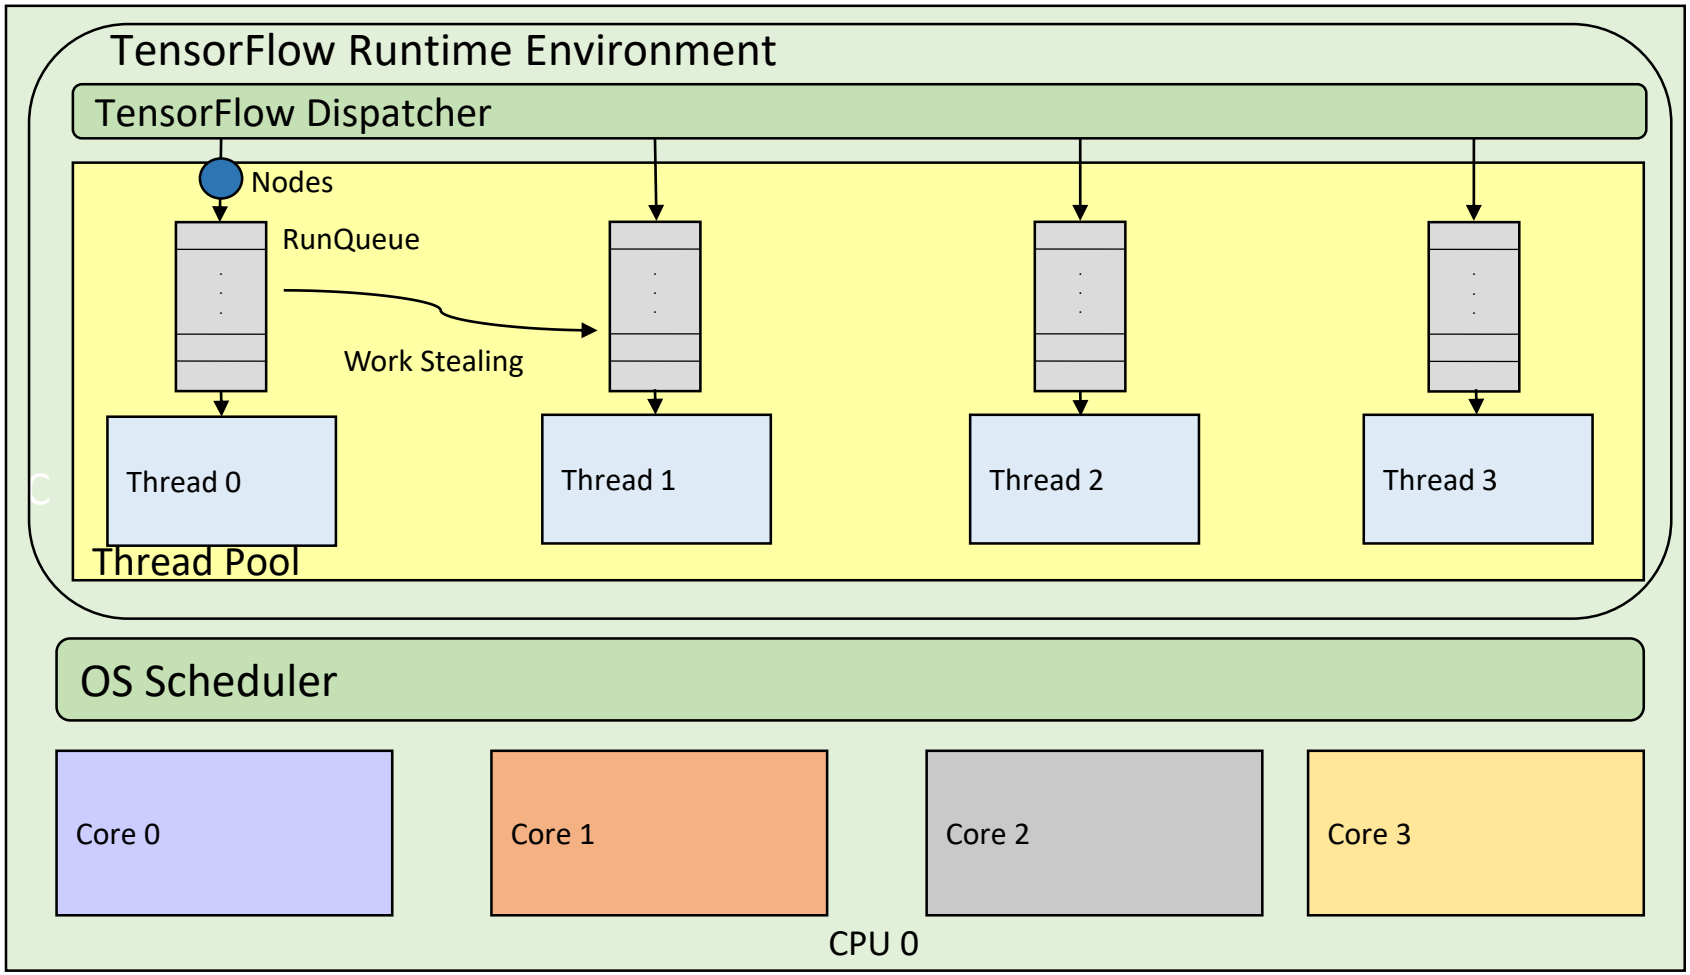

How Tensorflow works on CPUs?

- TensorFlow assigns ready nodes to threads of a thread pool (of globally-sch

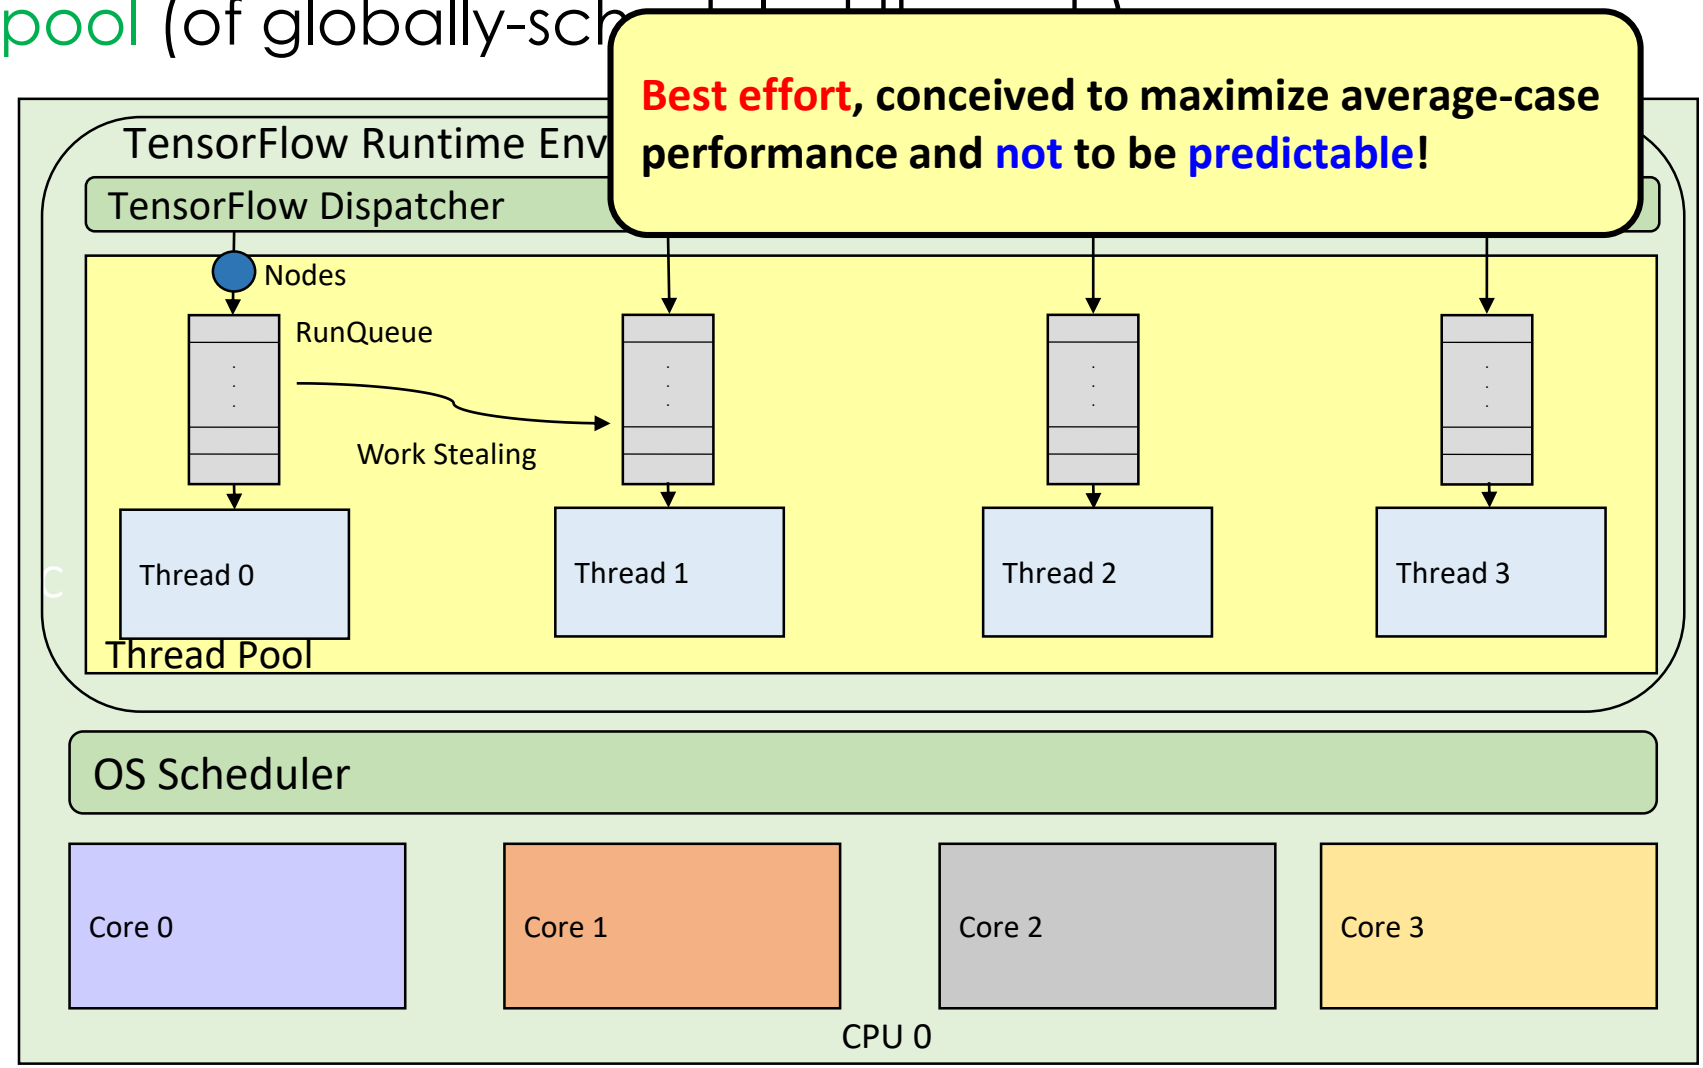

What we are doing?

Adding a layer for predictability in **TensorFlow**

- Extracting a computation model from DNNs, deriving precedence constraints, computation time of each node, memory exchanged, etc.
- Providing a **predictable** scheduling layer of nodes also aware of memory accesses
- Designing a partitioning scheme that considers producer-consumer relationships among nodes, for improving cache coherency
- Development of analysis techniques to assess schedulability of neural networks

Summary and conclusions

- Deep Neural Networks represent a promising technique for enacting autonomous driving, but…
  - their adoption in safety-critical scenarios presents many issues
- We focused on:
  - Certificability
  - Security and Isolation
  - Predictability

Summary and conclusions

- Deep Neural Networks represent a promising technique for enacting autonomous driving, but…

There is still a lot of work to do…

- We focused on:
  - Certificability
  - Security an
  - Predictabili

Let's start!

Thank you!

Daniel Casini
daniel.casini@sssup.it