DENSO

Crafting the Core

Research Challenges and Opportunities towards Safe Autonomous Driving

Chih-Hong Cheng

Corporate R&D

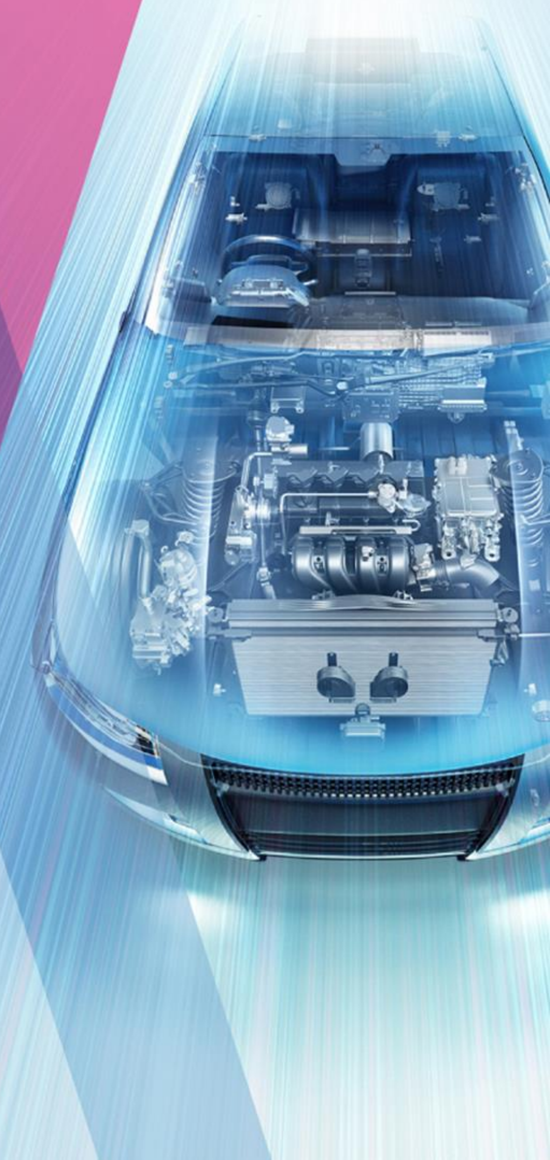

Before I start

- Some of the contents may be of my personal view and shall not be viewed as official statement of DENSO
- Some presented work are based on prior results during my tenure at fortiss

Source: www.fortiss.org

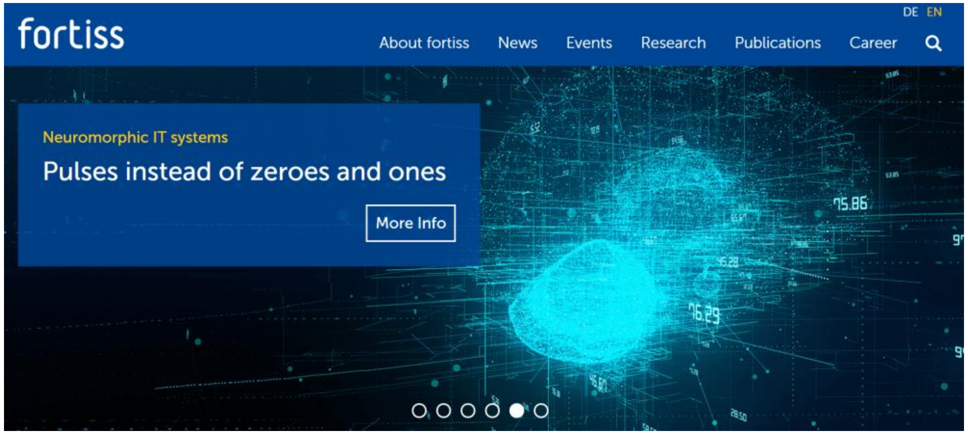

The race towards automated driving continues

OEM	BMW, VW, Toyota, …
Classical Tier-1 suppliers	ZF, Continental, Bosch, DENSO, …
New-comers	Google, Intel, Nvidia, …
Startups	Uber, Lyft, Zoox, Five.AI, Oxbotica, …

Safe autonomy is the destination

Source: Ytoutube (abc news)

There is a gap between running demos and safe products

But it's a money burning business

Starsky Robotics

Source: Medium

TEAM COMPANY NEWS WE'RE HIRING

The End of Starsky Robotics

Stefan Seltz-Axmacher Follow
Mar 19 · 9 min read

In 2015, I got obsessed with the idea of driverless trucks and started Starsky Robotics. In 2016, we became the first street-legal vehicle to be paid to do real work without a person behind the wheel. In 2018, we became the first street-legal truck to do a fully unmanned run, albeit on a closed road. In 2019, our truck became the first fully-unmanned truck to drive on a live highway.

And in 2020, we're shutting down.

And competition is fierce

MAD MONEY

Robotaxis will be available as soon as 2022, self-driving tech supplier Mobileye CEO says

PUBLISHED FRI, JAN 10 2020·6:48 PM EST

Tyler Clifford
@_TYLERTHETYLER_

SHARE

KEY POINTS

- "Robotaxi is not that far away. We are targeting early 2022," Mobileye CEO Amnon Shashua told CNBC.
- Self-driving cars will start with fleet operators before general use due to regulatory and cost constraints "that you cannot put on a consumer," he said in a "Mad Money" interview.
- "If more cars will be autonomous, more lives would be saved. A computer will do a better job than a human, eventually," he said.

BEYOND THE VALLEY
GET YOUR TECH INSIG
ACROSS THE

Source: CNBC

GOOGLE \ TECH \ TRANSPORTATION

Waymo tells riders that 'completely driverless' vehicles are on the way

'You can enjoy having the car to yourself'

By Andrew J. Hawkins | @andyjayhawk | Oct 10, 2019, 11:42am EDT

SHARE

Source: www.theverge.com

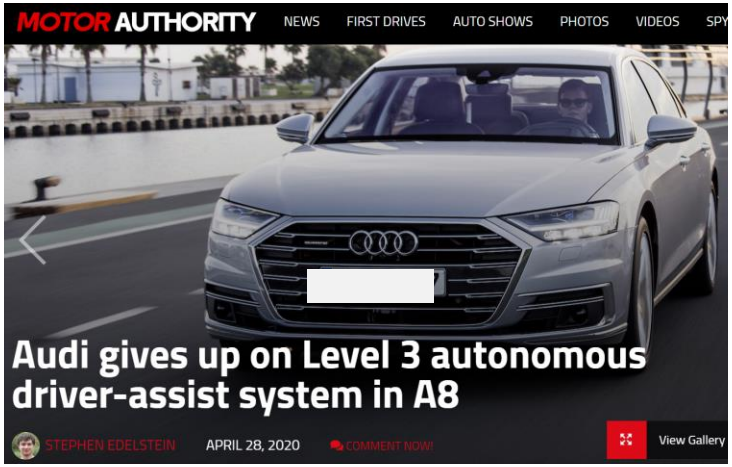

Source: MOTOR AUTHORITY

Safe autonomy is the destination

Source: Ytoutube (abc news)

We may close this gap quicker by scientific-driven methods (e.g., from "miles" to "intelligent miles")

Opportunities

Engineering tool provider and
component provider for chasing competitors

- Use total solution development as a learning process to validate the concept, but no need to be perfect in the solution

/* Selling hardware & EDA tools */

Agenda

- Background
- **DNN safety in automated driving**
- Concluding remarks

Methodology

Data quality

Robust training

Formal Verification

Runtime monitoring

Why can my DNN go wrong?

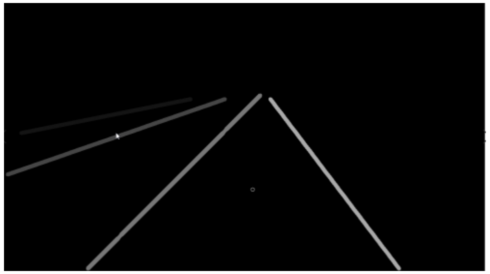

Average-case vs worst-case mindset

Very hard question

Being lazy in data collection
(Garbage-in-garbage-out)

...

Surprises in Operating Design Domain (ODD)

...

Maybe we should think systematically

Specification, data collection & labelling	Architectural design & training	testing /generalization	Operation

GSN for DNN safety argumentation

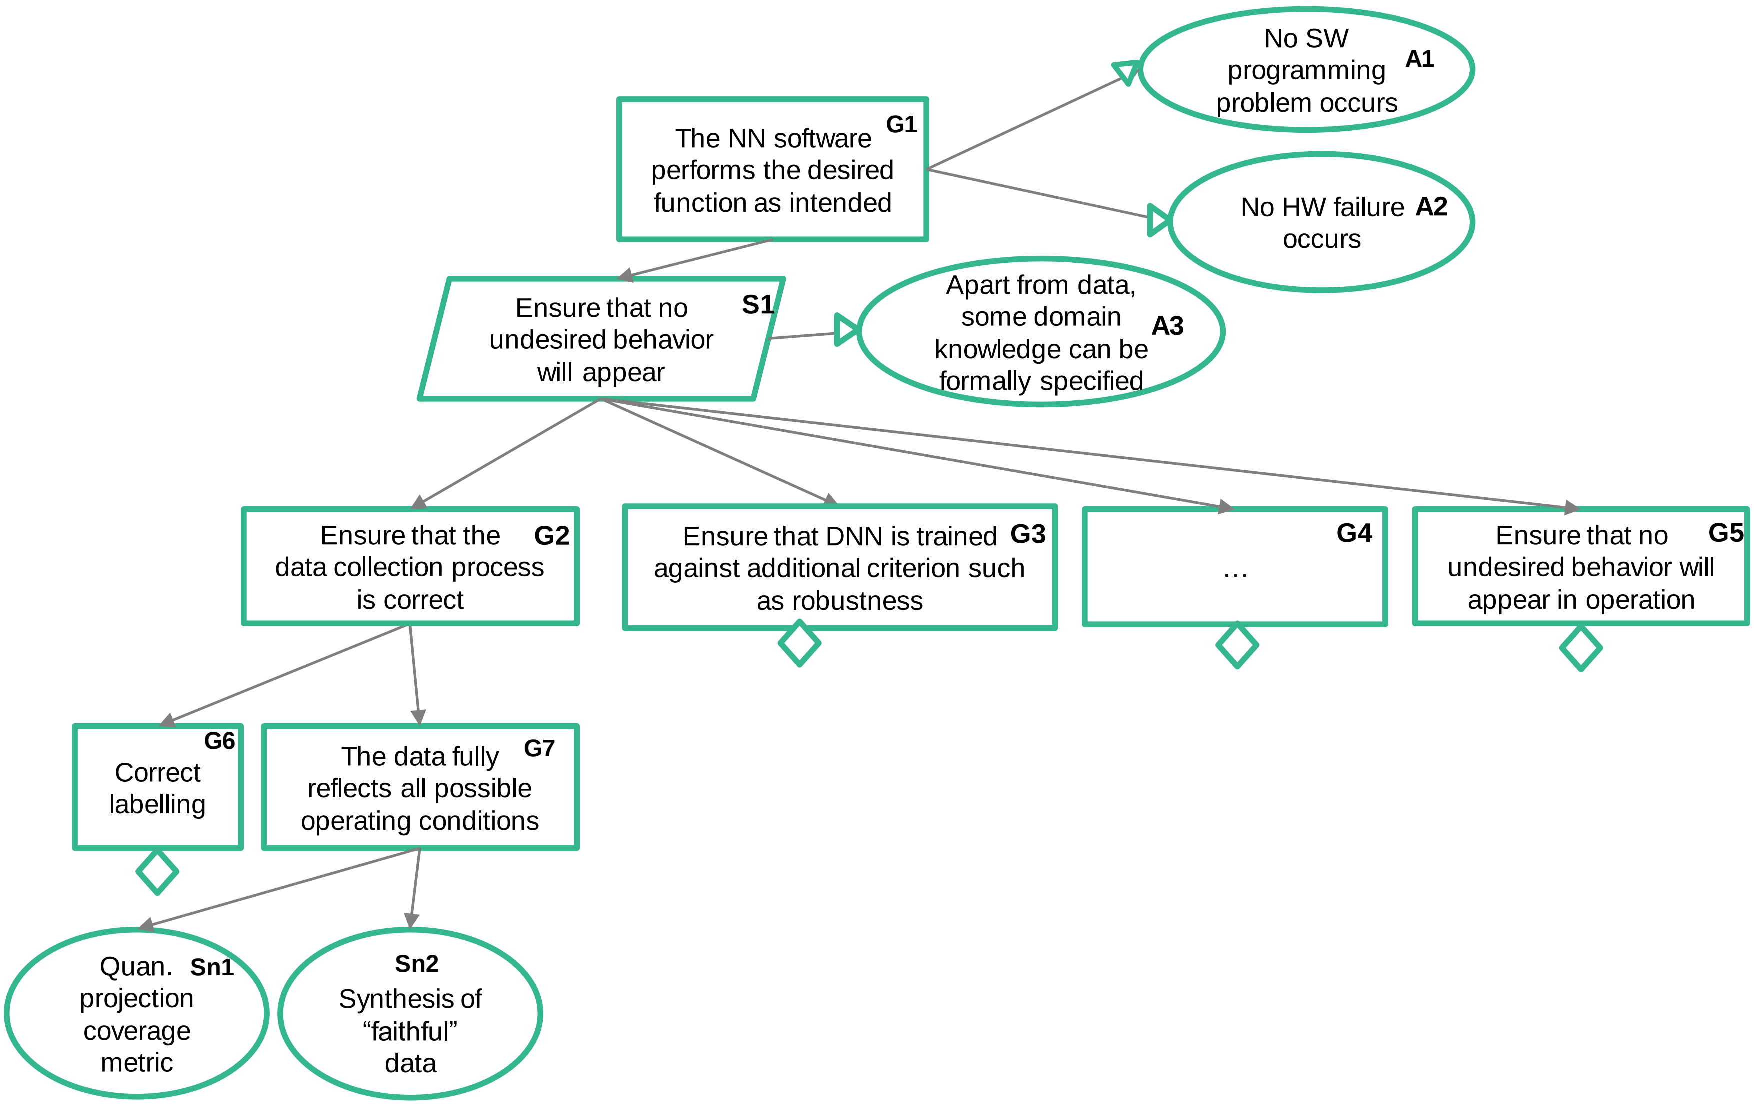

Addressing the DNN Safety via a Structured Approach

Systematically decompose problems into subproblems

Use scientific methods to provide elegant solutions (as evidences) to these problems

- Great battlefield for AI/ML/Safety/SE/FM researchers

Limitations

1. Currently, everyone (research institute, company, certification body) wants to have his “own” GSN
   - This is of course a waste of efforts
   - Also, it makes sense to focus on “what to be addressed”, and leave the “how” part open for creativity until best practice is out

2. GSN is nothing logical

GSN for DNN safety argumentation

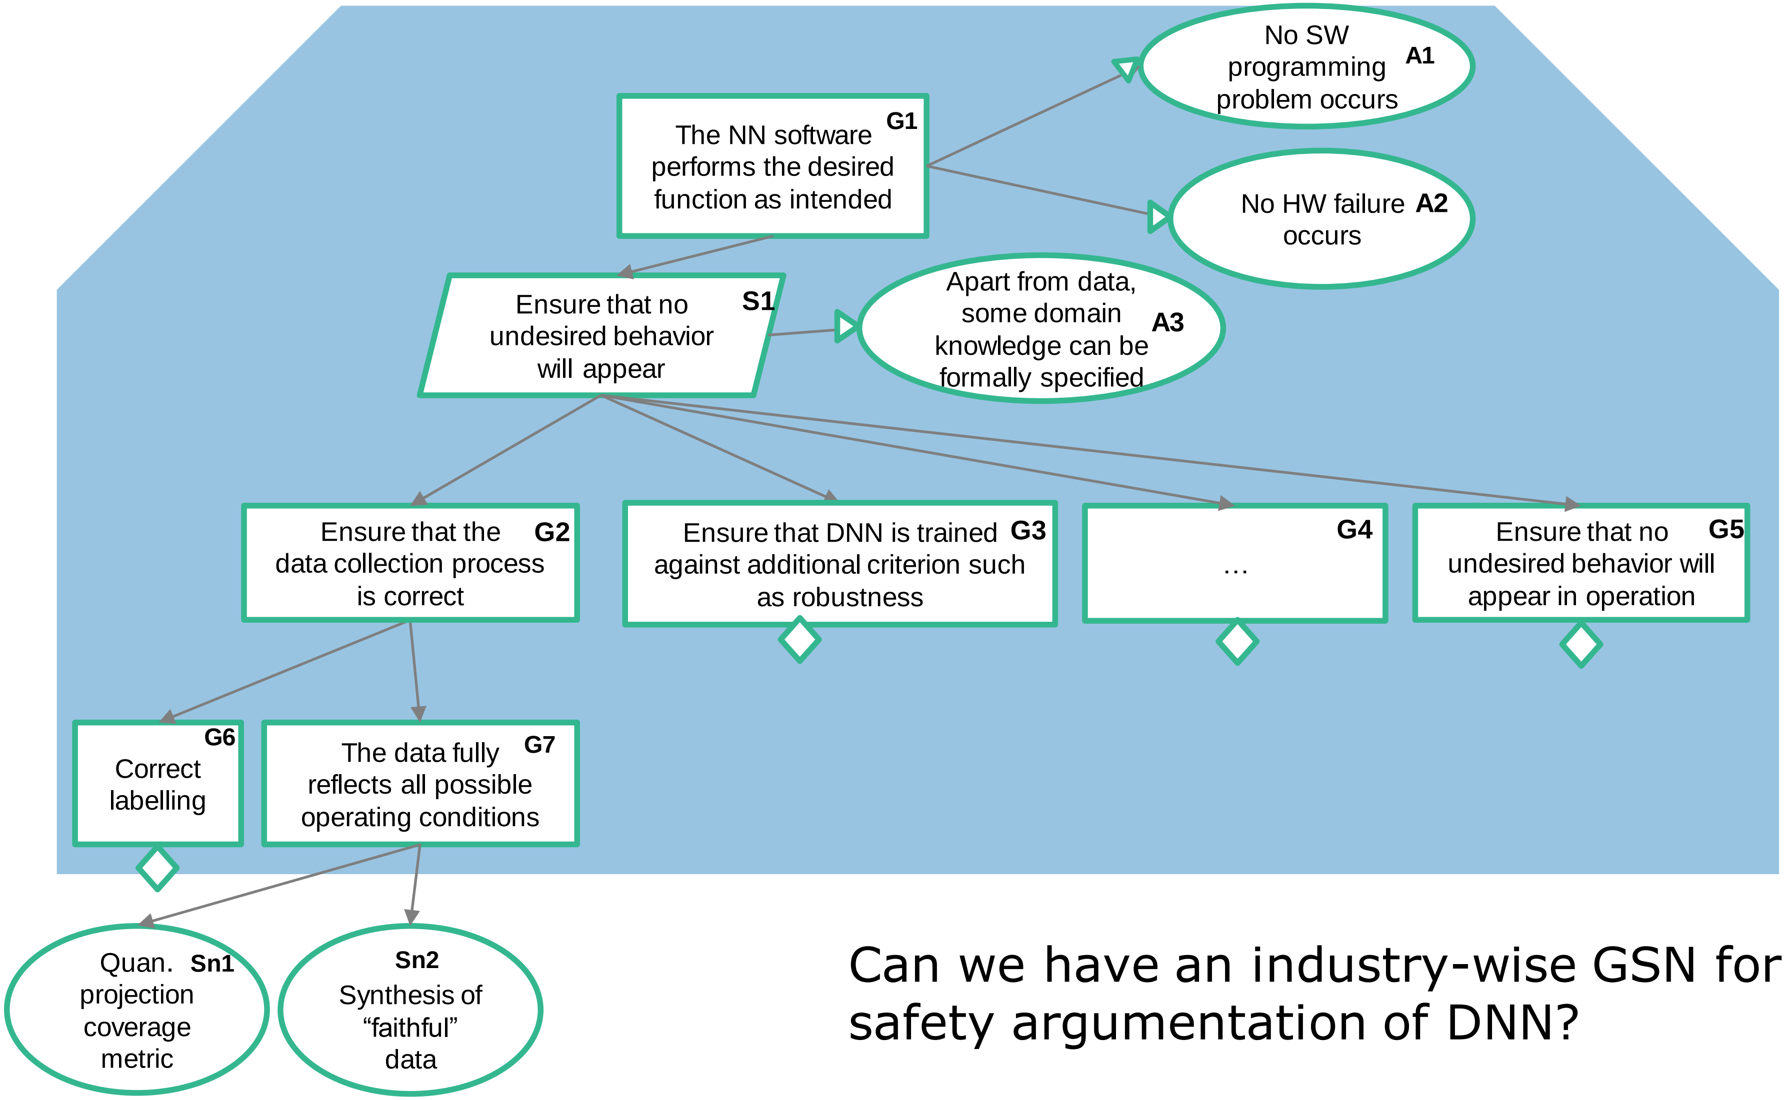

Can we have an industry-wise GSN for safety argumentation of DNN?

Methodology

Data quality

Robust training

Formal Verification

Runtime monitoring

Coverage problem

Combinatorial explosion of scenarios

One possible assignment of "discrete environment operating condition" creates one scenario

- Weather: Sunny/Cloudy/Rainy
- Curve: Straight/Curvy
- Oncoming Car: True/False
- Forward Car: True/False

30 discrete operating conditions => 2^{30} (1 billion) scenarios for testing

- You have definitely more!
- Such a denominator is huge, making most of the "coverage criterion" generate value ≈ 0

Question: Can we have a knob to tune?

"simpler to achieve 100%"

"completeness" more meaningful

Weaker form of “completeness”

The system under analysis takes 3 Boolean inputs x,y,z – a total of 8 input combinations (2^3)

- Each red box is a test case, so we only cover 4/8

But whenever we look at xy hyperplane (via projection), the hyperplane is **fully covered in red**

- Similarly for yz and xz

By fixing number of parameters to be chosen (in this case k=2), we still get a **weaker form of completeness with polynomially bounded test cases**

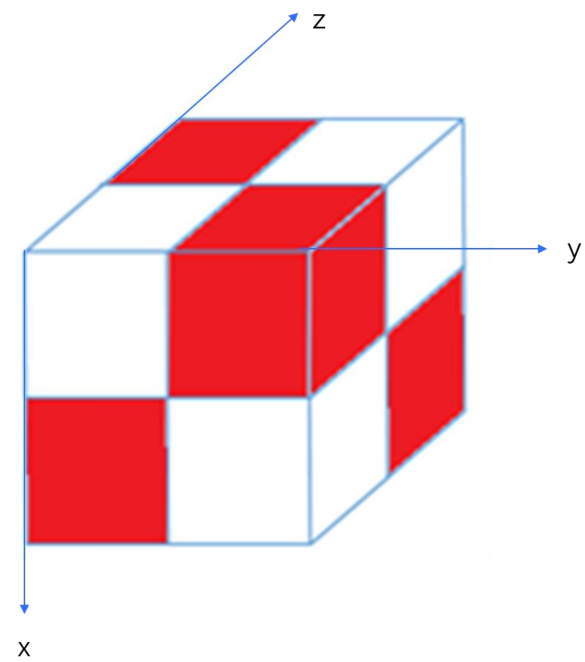

Combinatorial testing and coverage arrays

Operating conditions

- Weather: Sunny/Cloudy/Rainy
- Curve: Straight/Curvy
- Oncoming Car: True/False
- Forward Car: True/False

Combinatorial testing for k-projection: all test cases should cover all possible operating condition tuple

Given k being a constant, the number of test cases needed is **polynomially bounded**, $\binom{n}{k} 2^k$

■ (Sunny, Curvy, Oncoming, No Forward)

□ (Cloudy, Curvy, No Oncoming,Forward)

	Straight	Curvy
Sunny		■
Cloudy		□
Rainy		

	Oncoming (yes)	Oncoming (No)
Sunny	■	
Cloudy		□
Rainy		

	Forward (yes)	Forward (No)
Sunny		■
Cloudy	□	
Rainy		

	Straight	Curvy
Oncoming (yes)		■
Oncoming (no)		□

	Straight	Curvy
Forward (yes)		□
Forward (no)		■

	Oncoming (yes)	Oncoming (No)
Forward (yes)		□
Forward (no)	■	

For autonomous driving, things may be a bit more complicated

- Certain combination of operating conditions (expressed as domain knowledge) may not be feasible, and one should not consider it
  - K-projection coverage + constraint in the domain

- One would like to place different emphasis over different scenarios
  - K-projection coverage + quantitative aspects

- In the paper, we consider these two extensions at once

Result






Data collected for OEM X highway pilot project (during my tenure at fortiss)

- Used in testing
- Used in assume-guarantee verification

Limitation

- There seems to be some further improvements in specification + data collection, e.g.,
  - Disciplined method for data labelling and the effect on uncertainty
    - If you have some error in labelling bounding boxes, it makes no sense to pursue prediction perfection
    - Class imbalance and their mediation
  - Quantifying similarity measure between simulation engine and reality, and to understand their impact

Methodology

Data quality

Robust training

Formal Verification

Runtime monitoring

Provably Robust Training

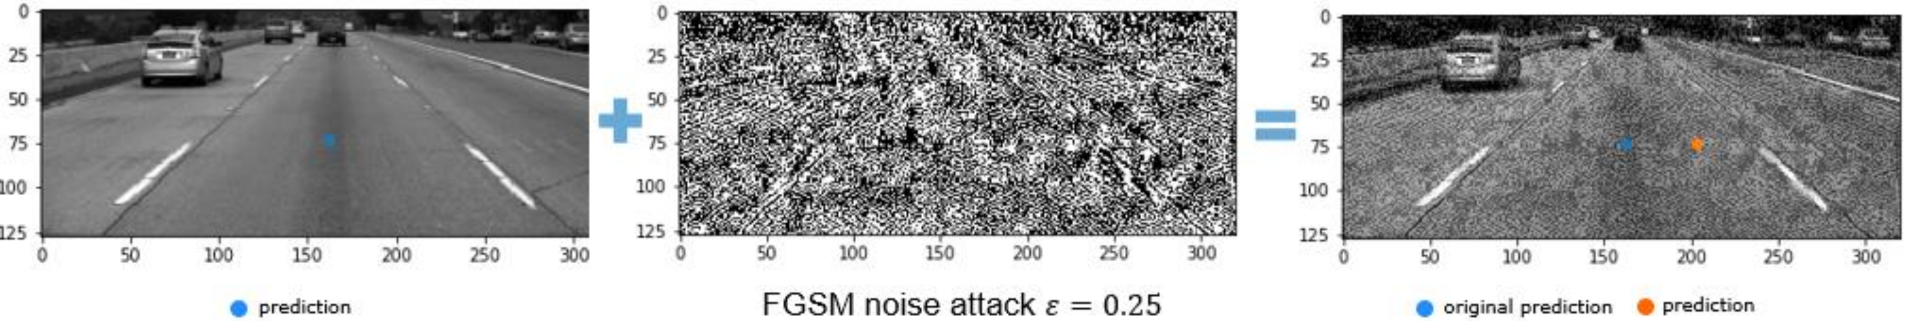

Standard training techniques are subject to noise and adv. attacks

Provably robust DNN training technique can resist attacks

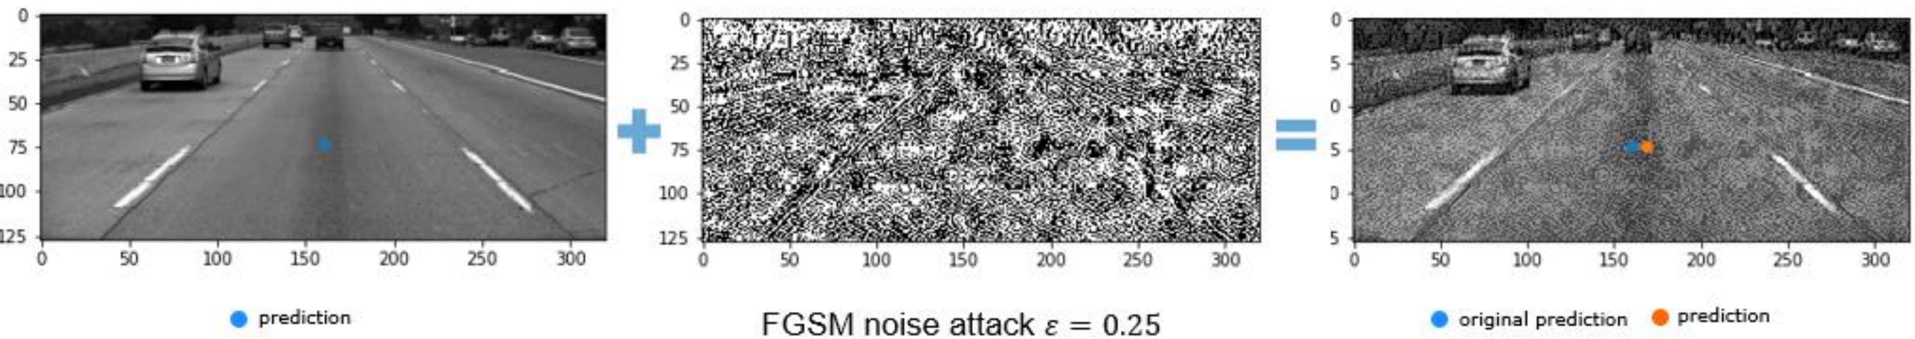

Behind provably robust training

Provably Robust DNN Training

= Standard DNN training

\+ New neuron layers with symbolic bound propagation techniques

(to estimate the worst-case effect due to perturbation)

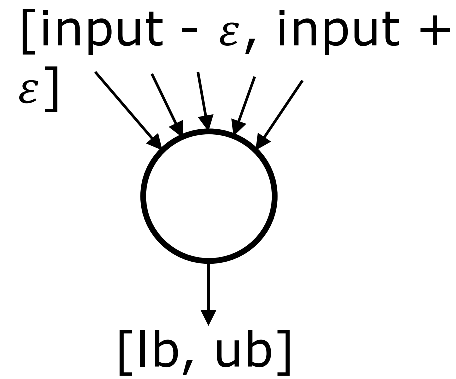

\+ New robust loss function

(to understand if the worst-case effect is contained inside the allowed tolerance)

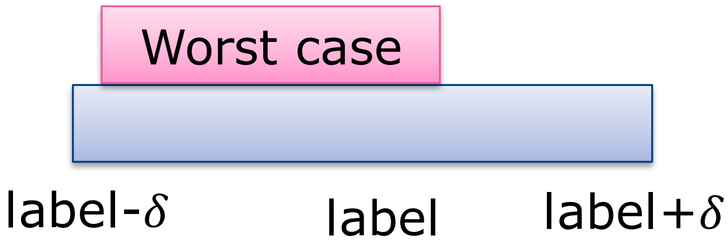

Limitation

- Going beyond bit-level perturbation into feature-level perturbation

- The robust-accuracy tradeoff

- Even with zero loss (in the training dataset), the created provable guarantees will still be lost if you are not careful in post-processing algorithms (such as non-max suppression; see SafeComp'20)

Methodology

Data quality

Robust training

Formal Verification

Runtime monitoring

Runtime monitoring

Decision supported by prior similarities in training?

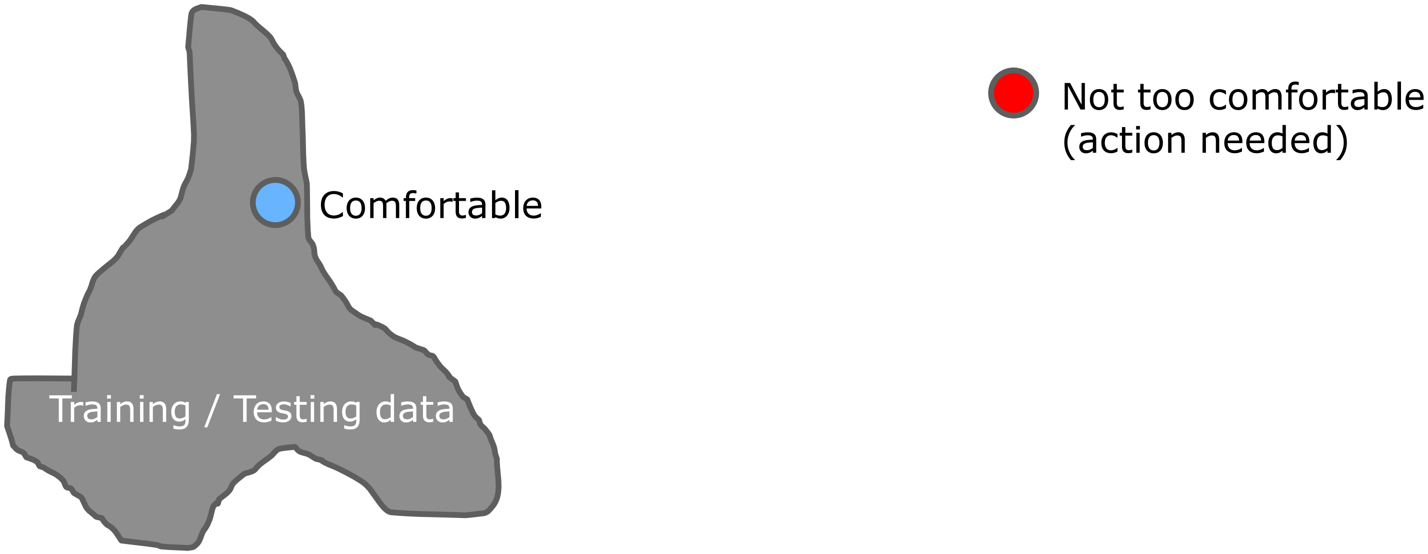

We might need a bag of techniques

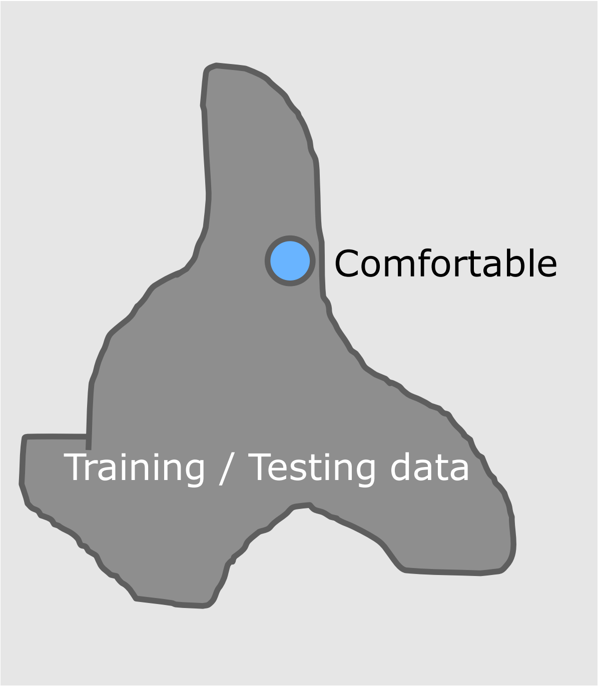

abstraction

Not too comfortable
(action needed)

Abstraction-based monitor: if “not too comfortable”, then it is truly problematic

We might need a bag of techniques

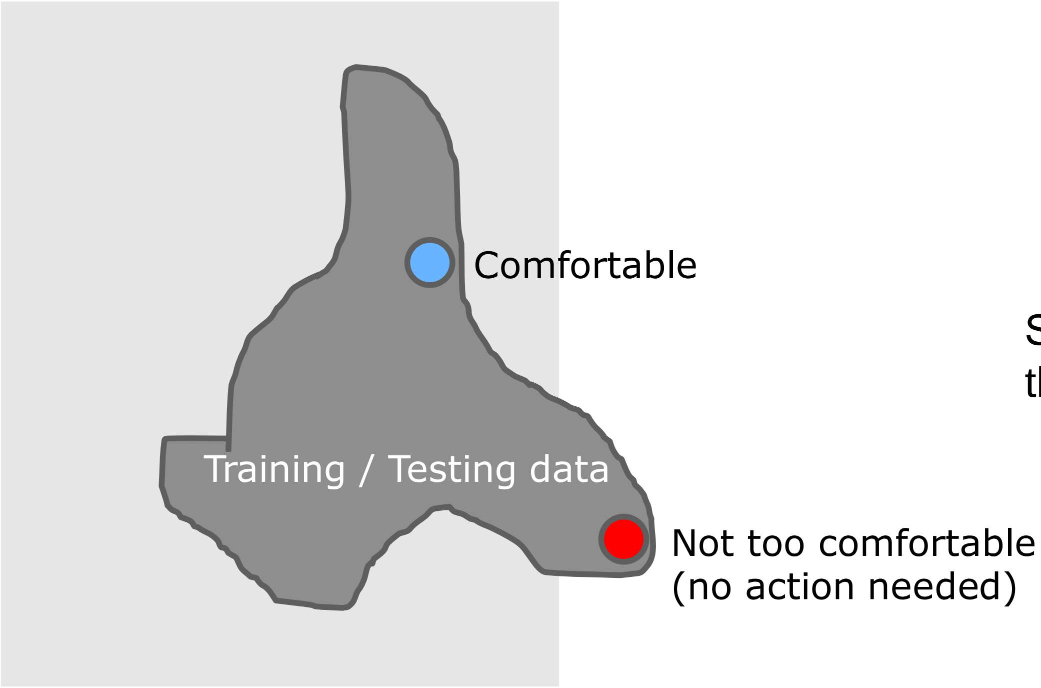

Monitor 1

Not too comfortable
(action needed)

Standard monitor: if “not too comfortable”, then it may be something OK

We might need a bag of techniques

Arsenal

- Abstraction based on neuron activation patterns (value bounds, activation sequences)
- Drop-out and majority vote
- Noise and majority vote
- Autoencoder with reconstruction loss
- …

Limitation

- Things need to be scalable on 3D object detection

Methodology

Data quality

Robust training

Formal Verification

Runtime monitoring

Formal verification of neural networks

Summary of approaches (numerous papers in two years)

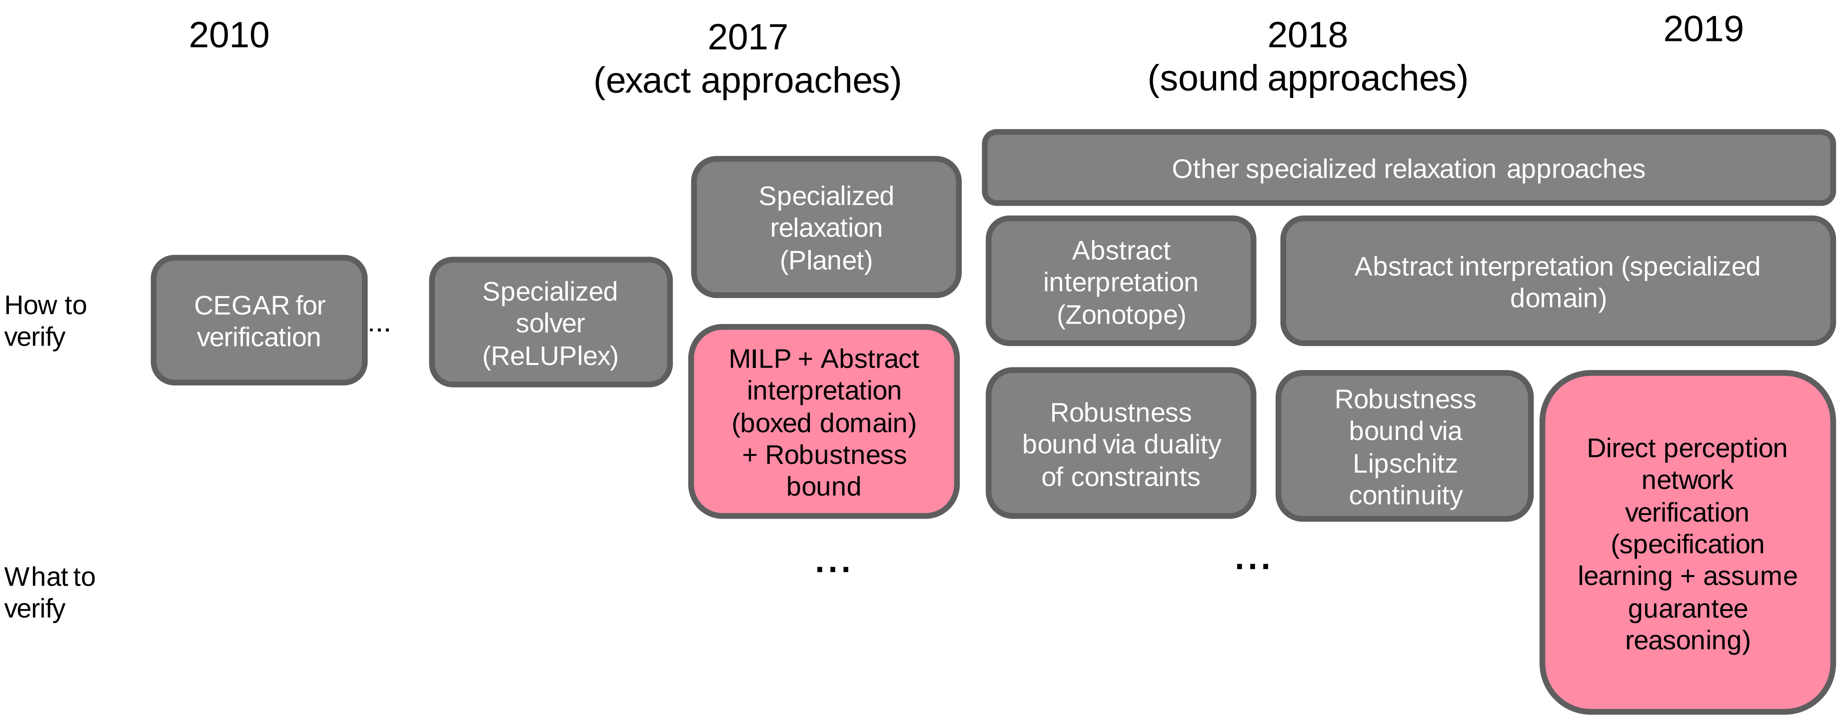

The ultimate challenge – Image from autonomous driving

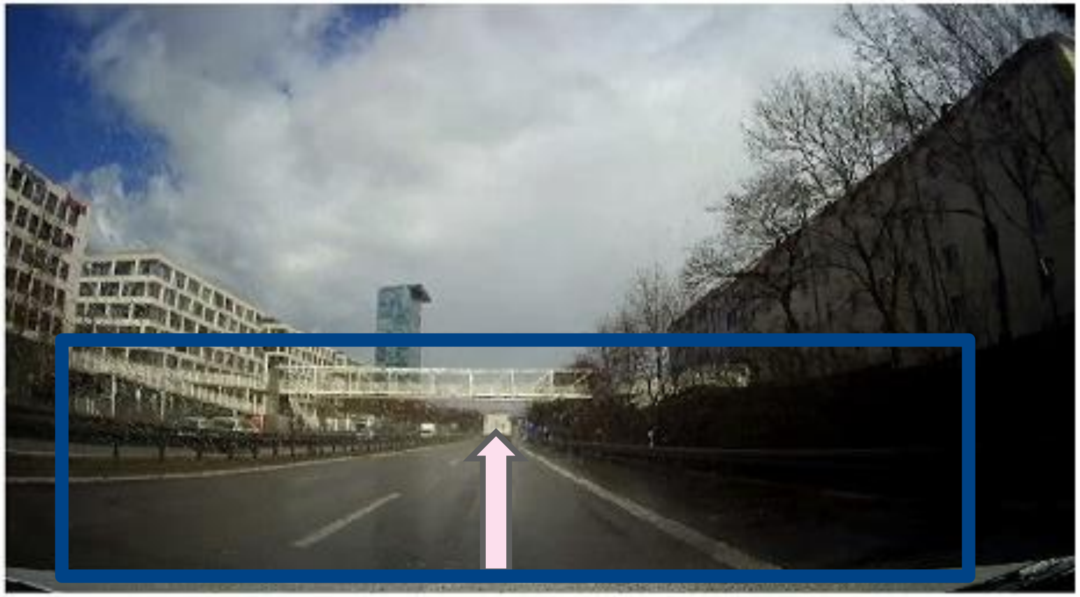

For illustration only (not output from real network)

- Large input space
  - Lane detection: 400x150= 60k pixels (RGB)
  - MNIST: 28x28=784 pixels (greyscale)
- Information rich (beyond characters)

Verification in practice

E.g., we want to prove that "if **the road bends to the left**", the neural network path planner never output to steer to the right"

We need to handle

1. **Specification problem**
   - What kind of input characterizes "**the road bends to the left**"

     They need to specified as **constraints over input variables**
   - What kind of input characterize the ODD?
     - If you just use $[-1, 1]^N$ (i.e., unconstrained), where N is the number of pixels, you very likely will get a counter example
2. **Scalability problem**

   Static analysis won't give you the precision you need; exact methods via constraint solving can't scale that well

DENSO
Crafting the Core

Learning input specifications for formal verification

- Constraints over input variables → constraints over new output variables

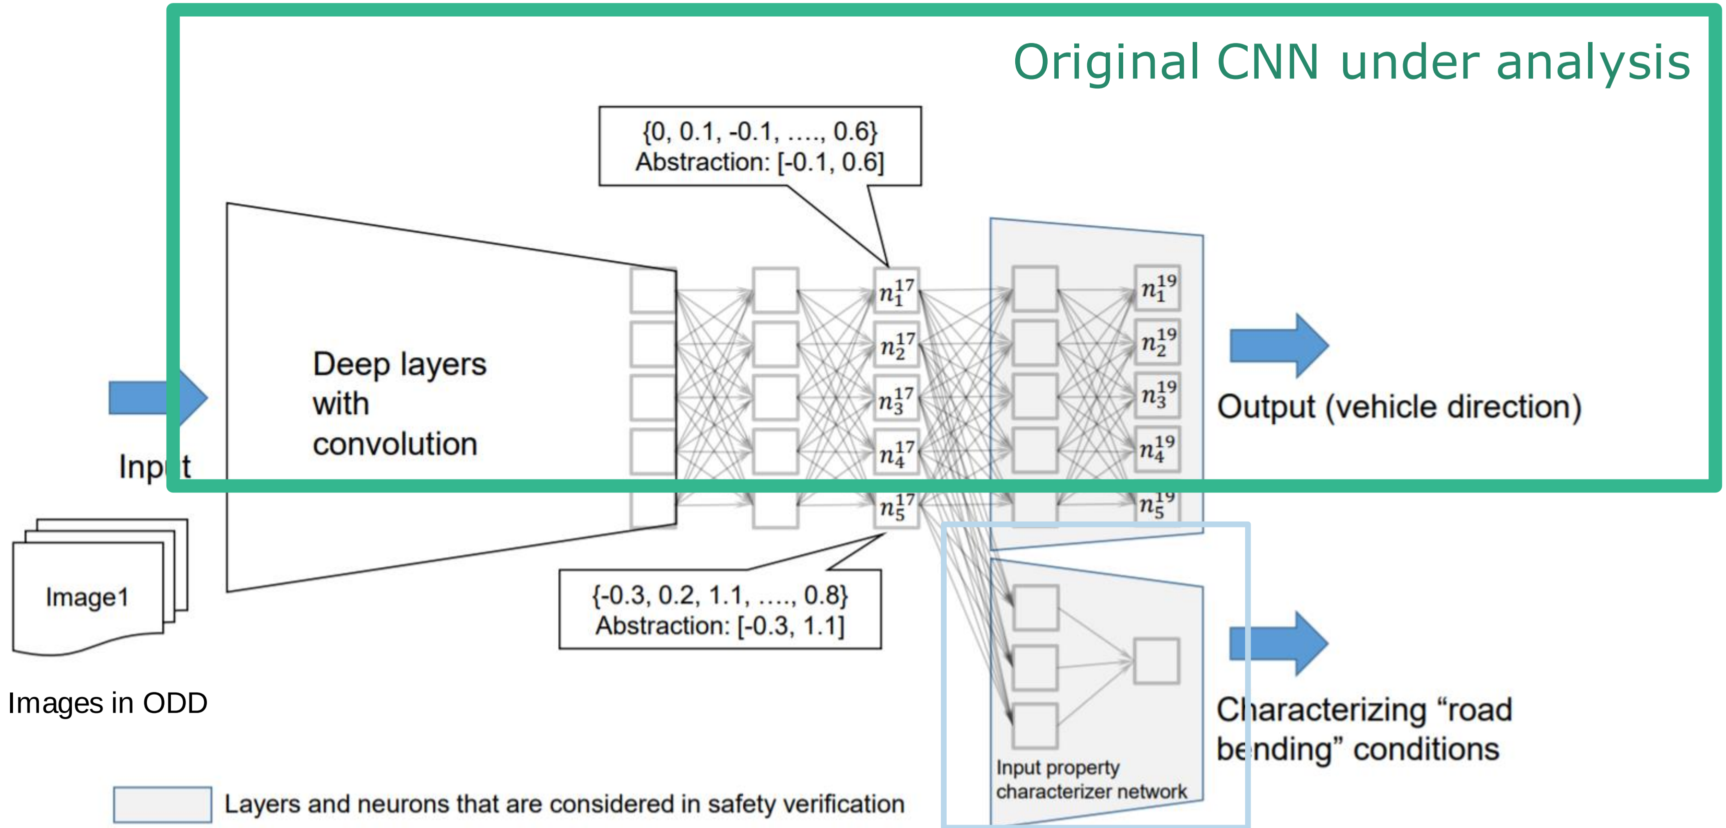

Created for verification purposes, not used in production

ODD and scalable verification

- Characterizing ODD now turned into the boxed abstraction
  - The boxed abstraction, acting as an assumption, needs to be monitored in runtime (assume-guarantee reasoning)

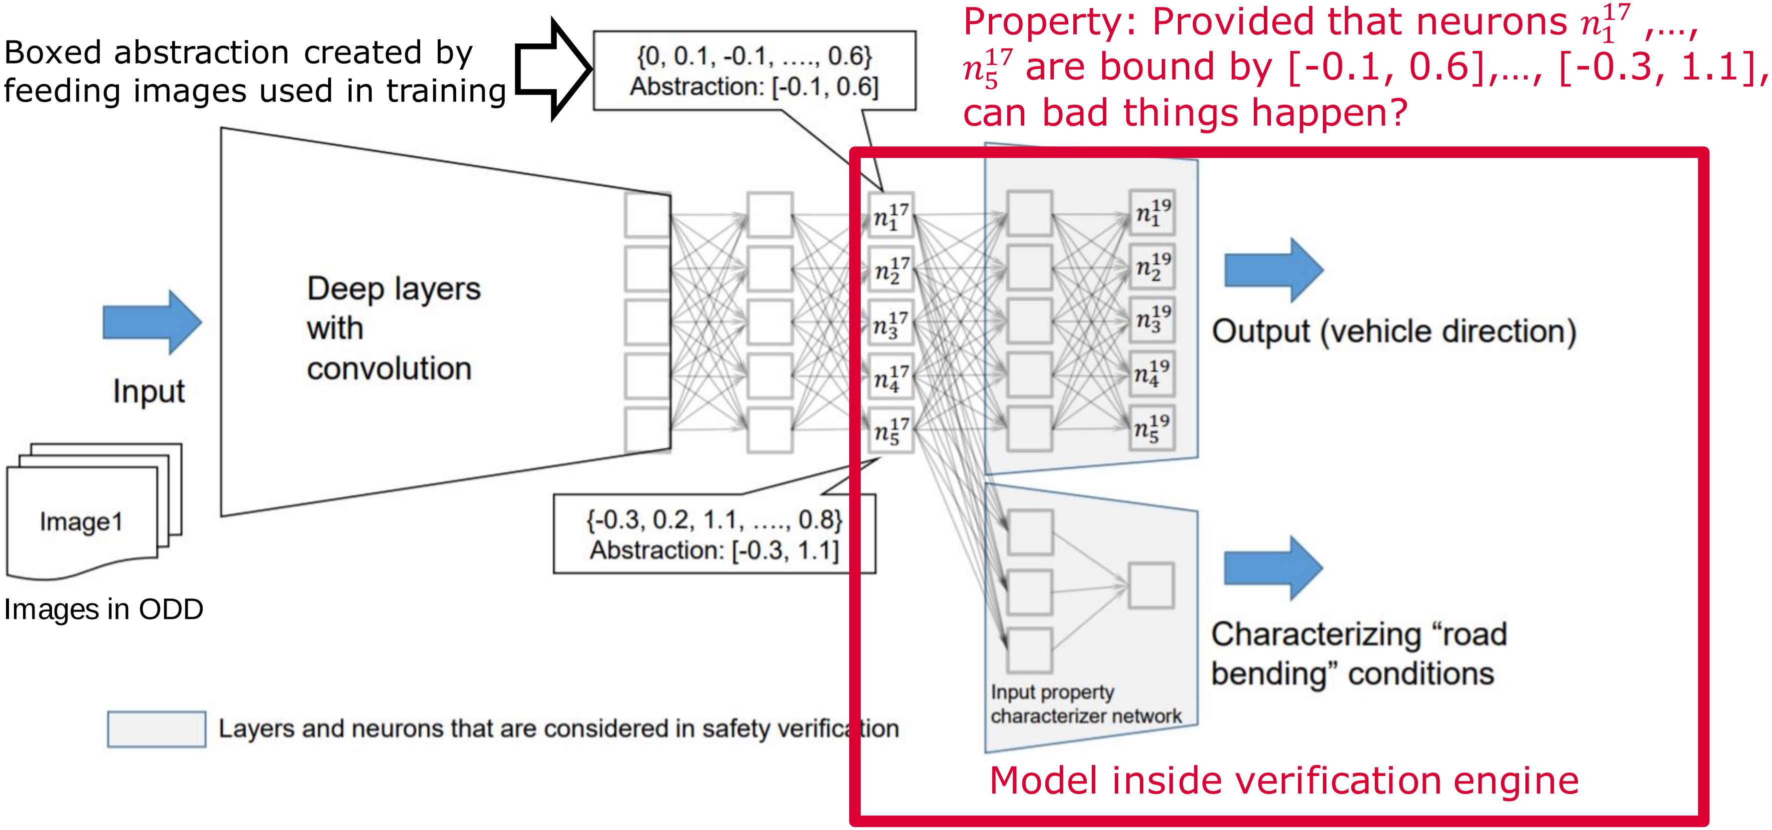

Result and Limitations

In this work (together with an OEM), we were able to prove that *extremely bad things won't happen*

- E.g., if the road is bending hugely to the left, the decision won't suggest to go hugely to the right.

Limitations

- We couldn't prove that "bad things won't happen"
- Maybe formal verification is just a topic not applicable on perception
    - Pushing scalability may be an academic interest, but not for industry

Agenda

- Background
- DNN safety in automated driving
- **Concluding remarks**

Concluding remarks

- Safety of automated driving is now the decisive factor

- We need a disciplined approach for engineering DNN to be used in autonomous driving

- Possible to borrow techniques from other fields (EDA, Control, SE, FM) to bring benefits

DENSO

Crafting the Core