Update on FG-AI4AD related activities

CITS meeting

9th September 2020

Geneva, Switzerland

Operational Safety

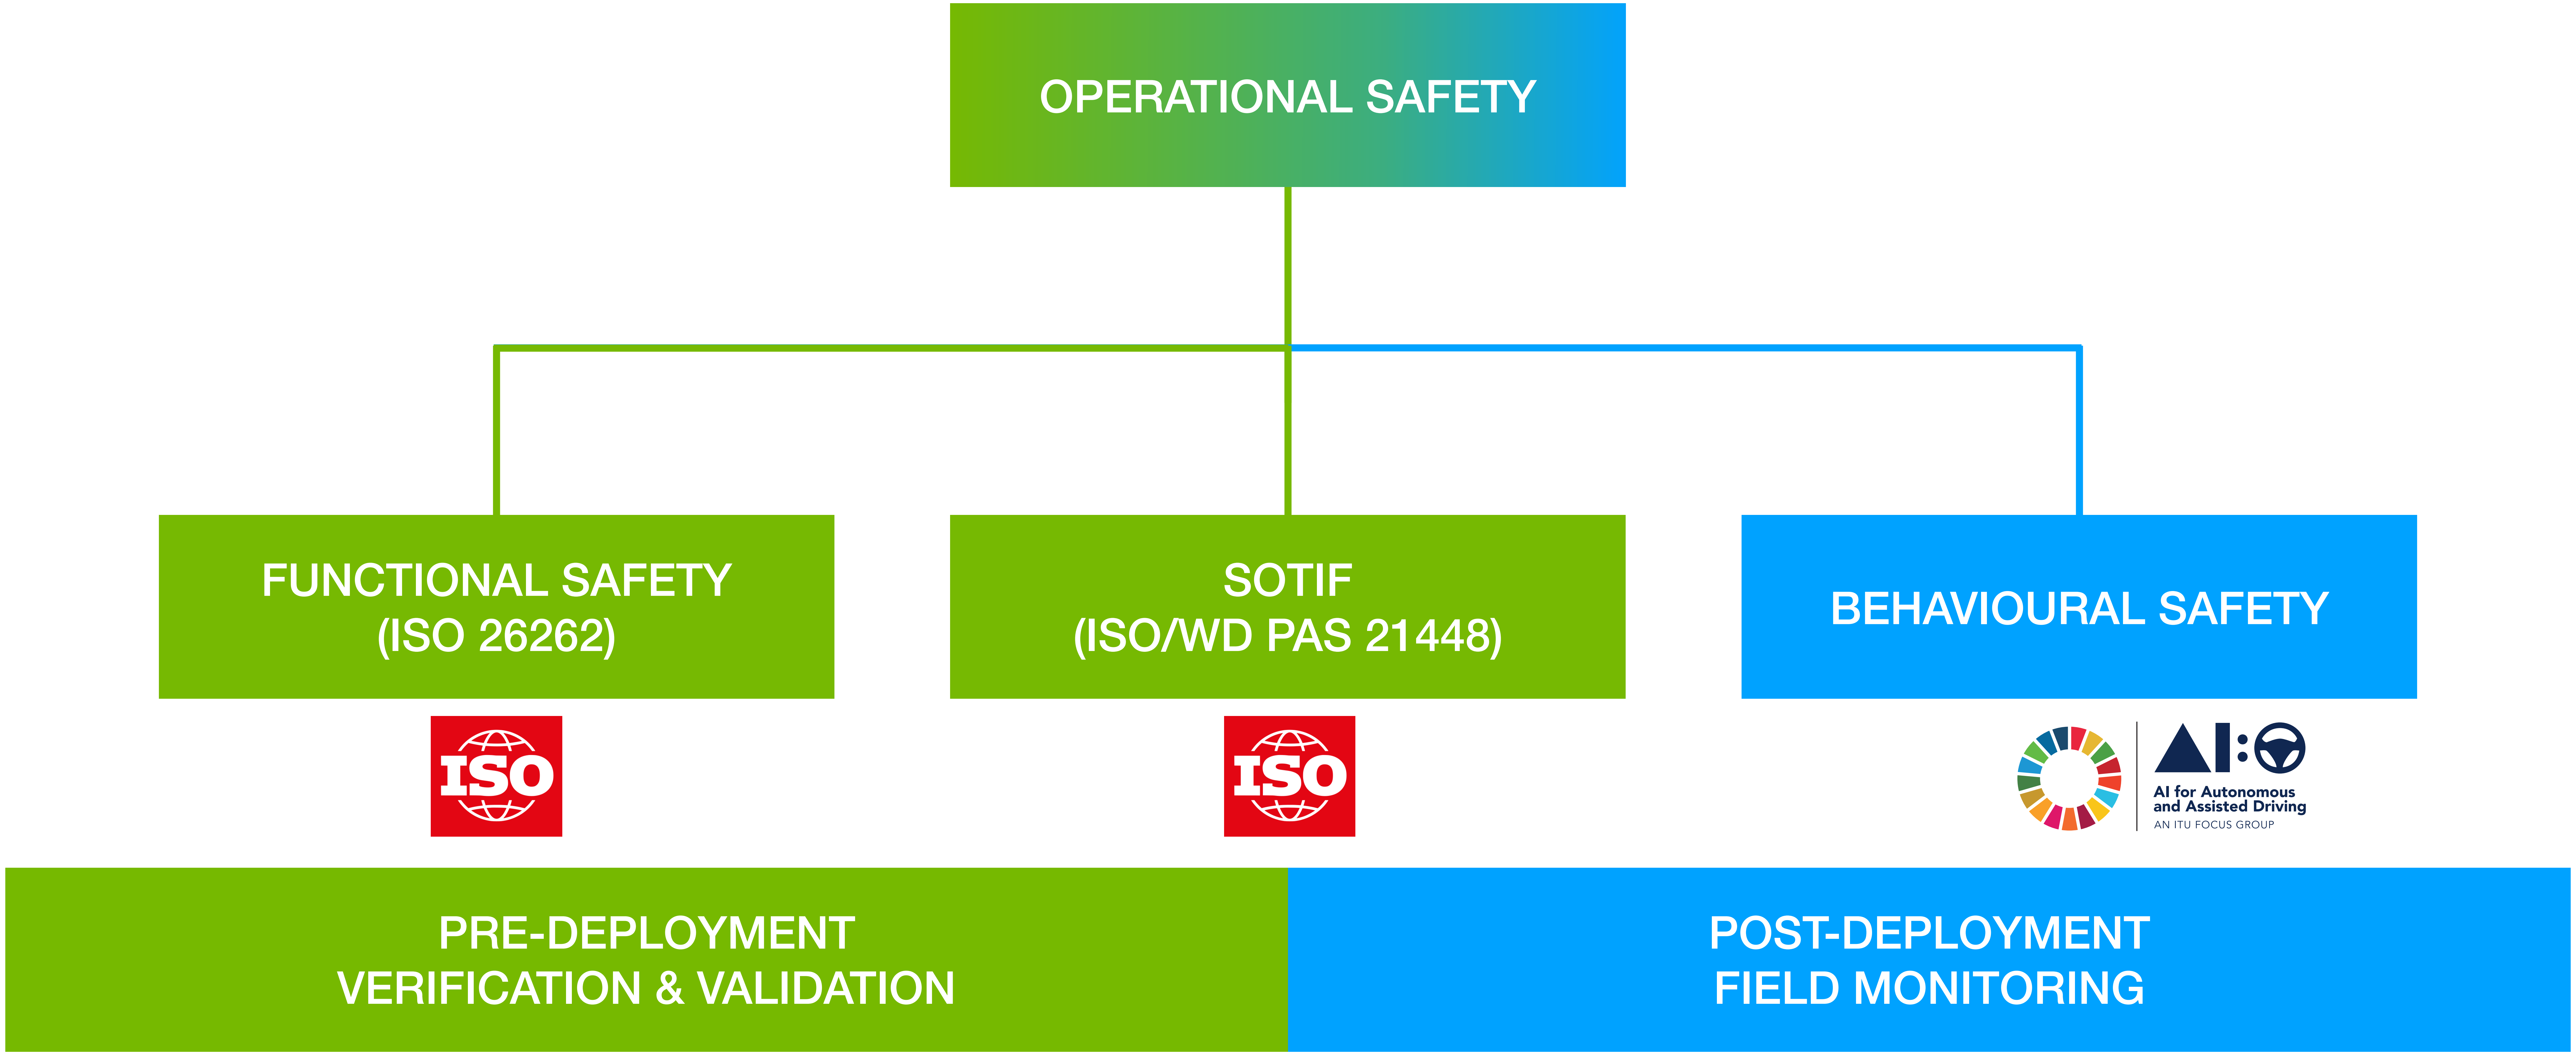

TNO innovation for life https://publications.tno.nl/publication/34626550/AyT8Zc/TNO-2018-streetwise.pdf

Global Regulatory Landscape

UNECE

UNECE WP29

Agreement concerning the Establishing of Global Technical Regulations for Wheeled Vehicles…

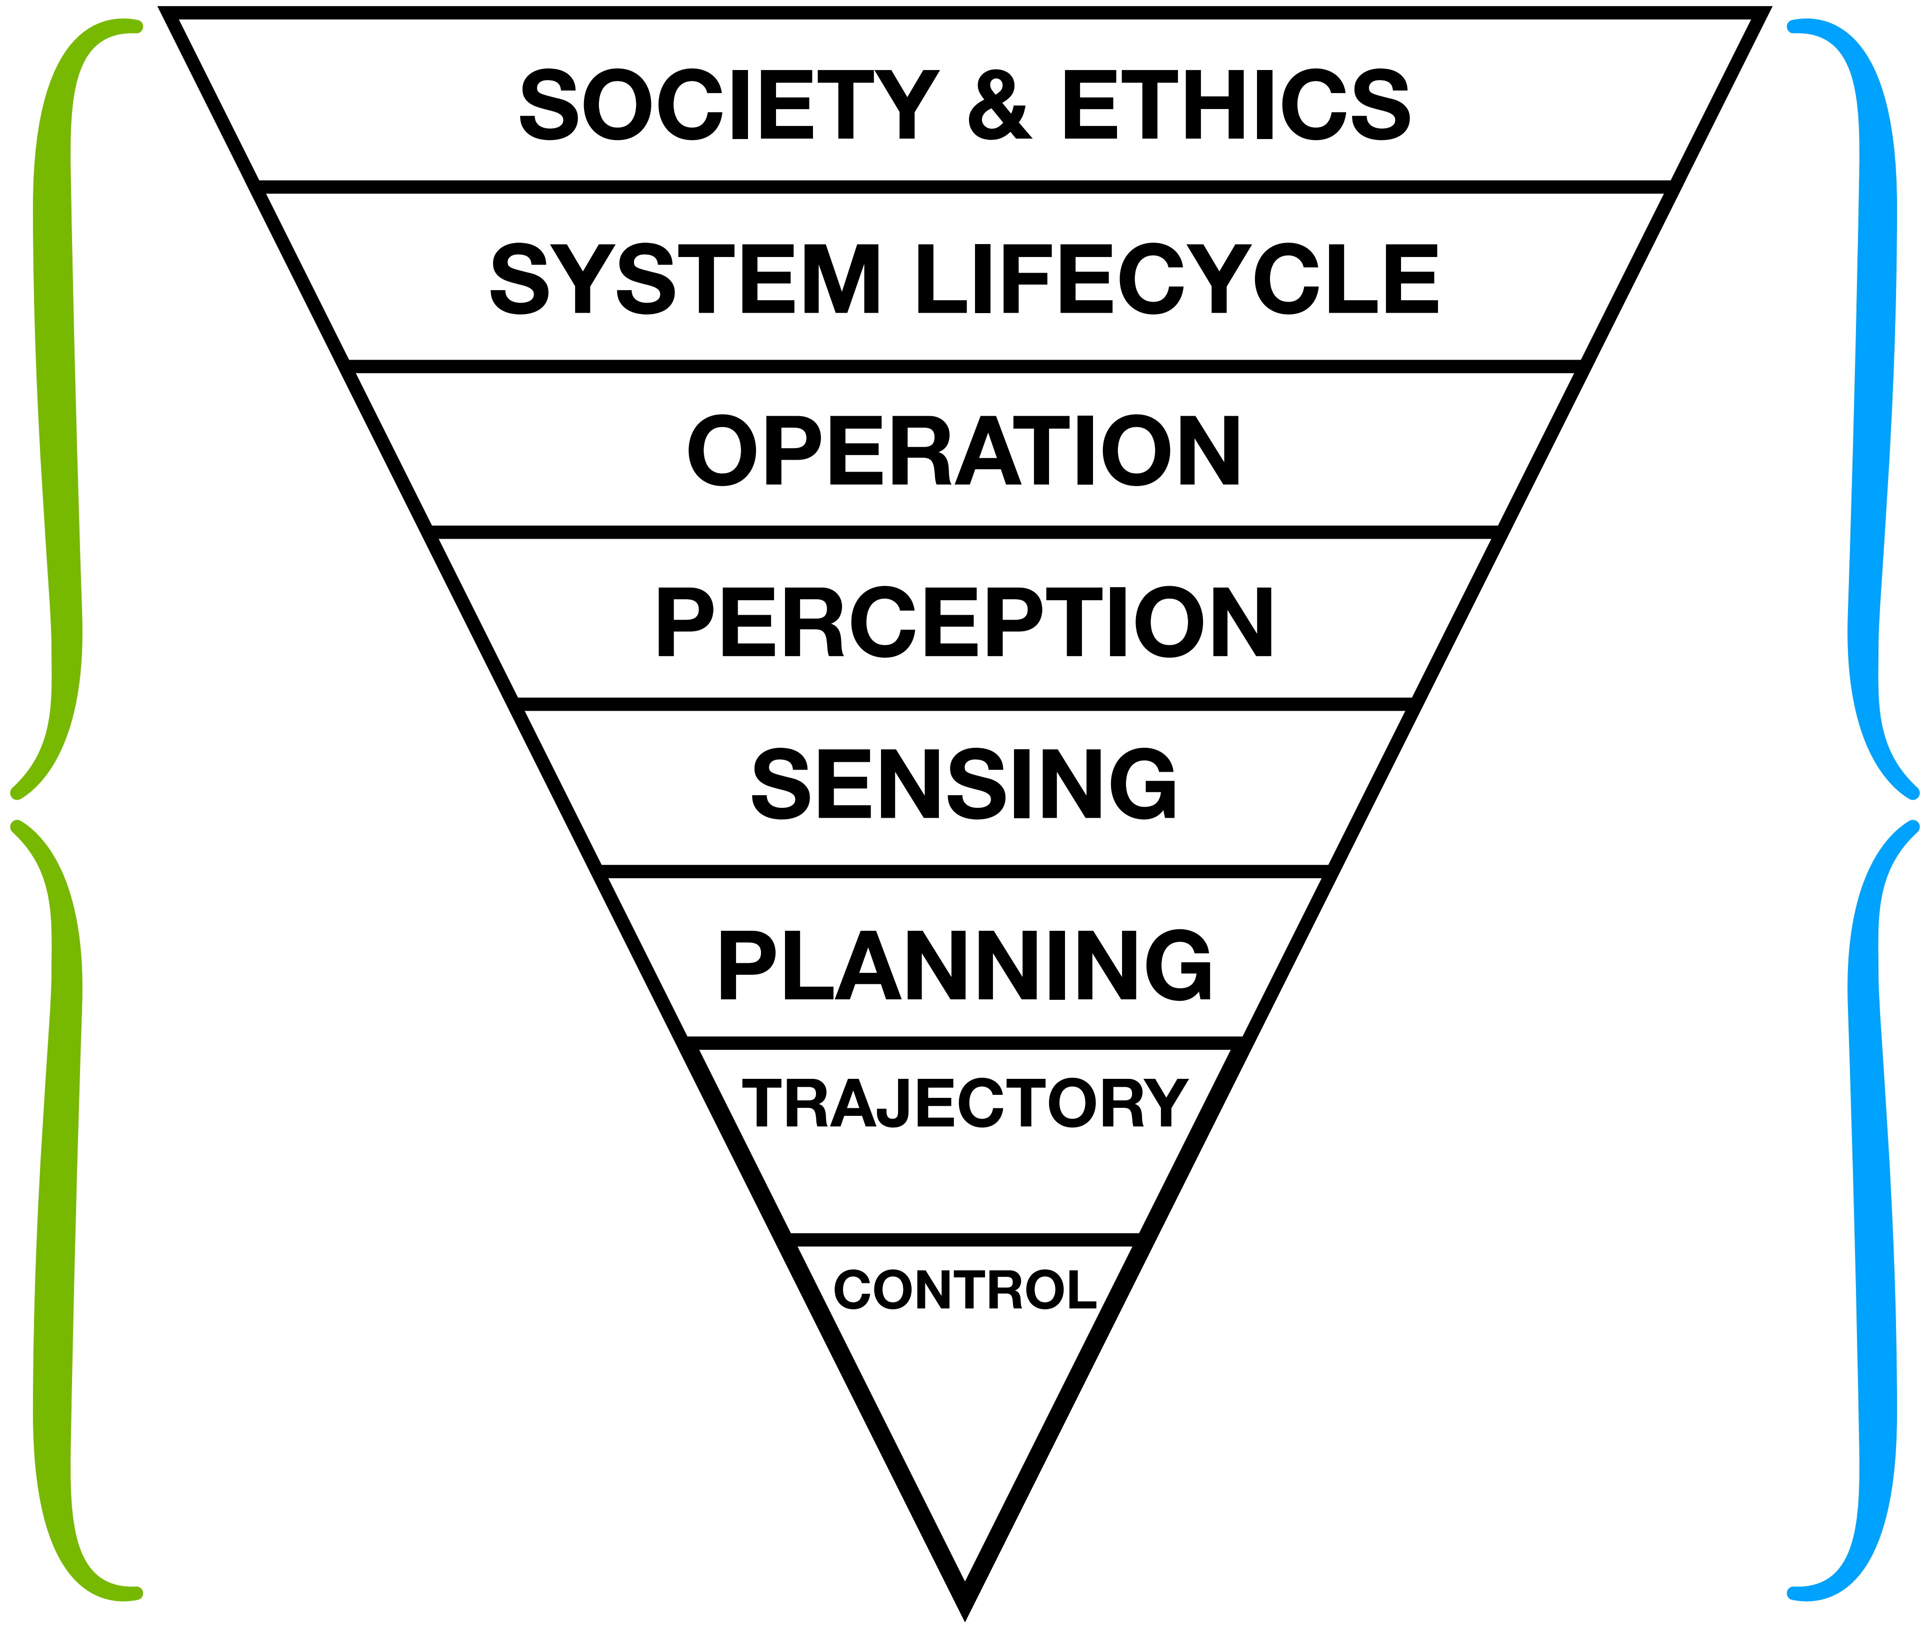

UNECE

UNECE WP1

Convention on Road Traffic

Road User & Driver Behaviour

PRE-DEPLOYMENT VERIFICATION & VALIDATION

POST-DEPLOYMENT FIELD MONITORING

Safety Standard Landscape

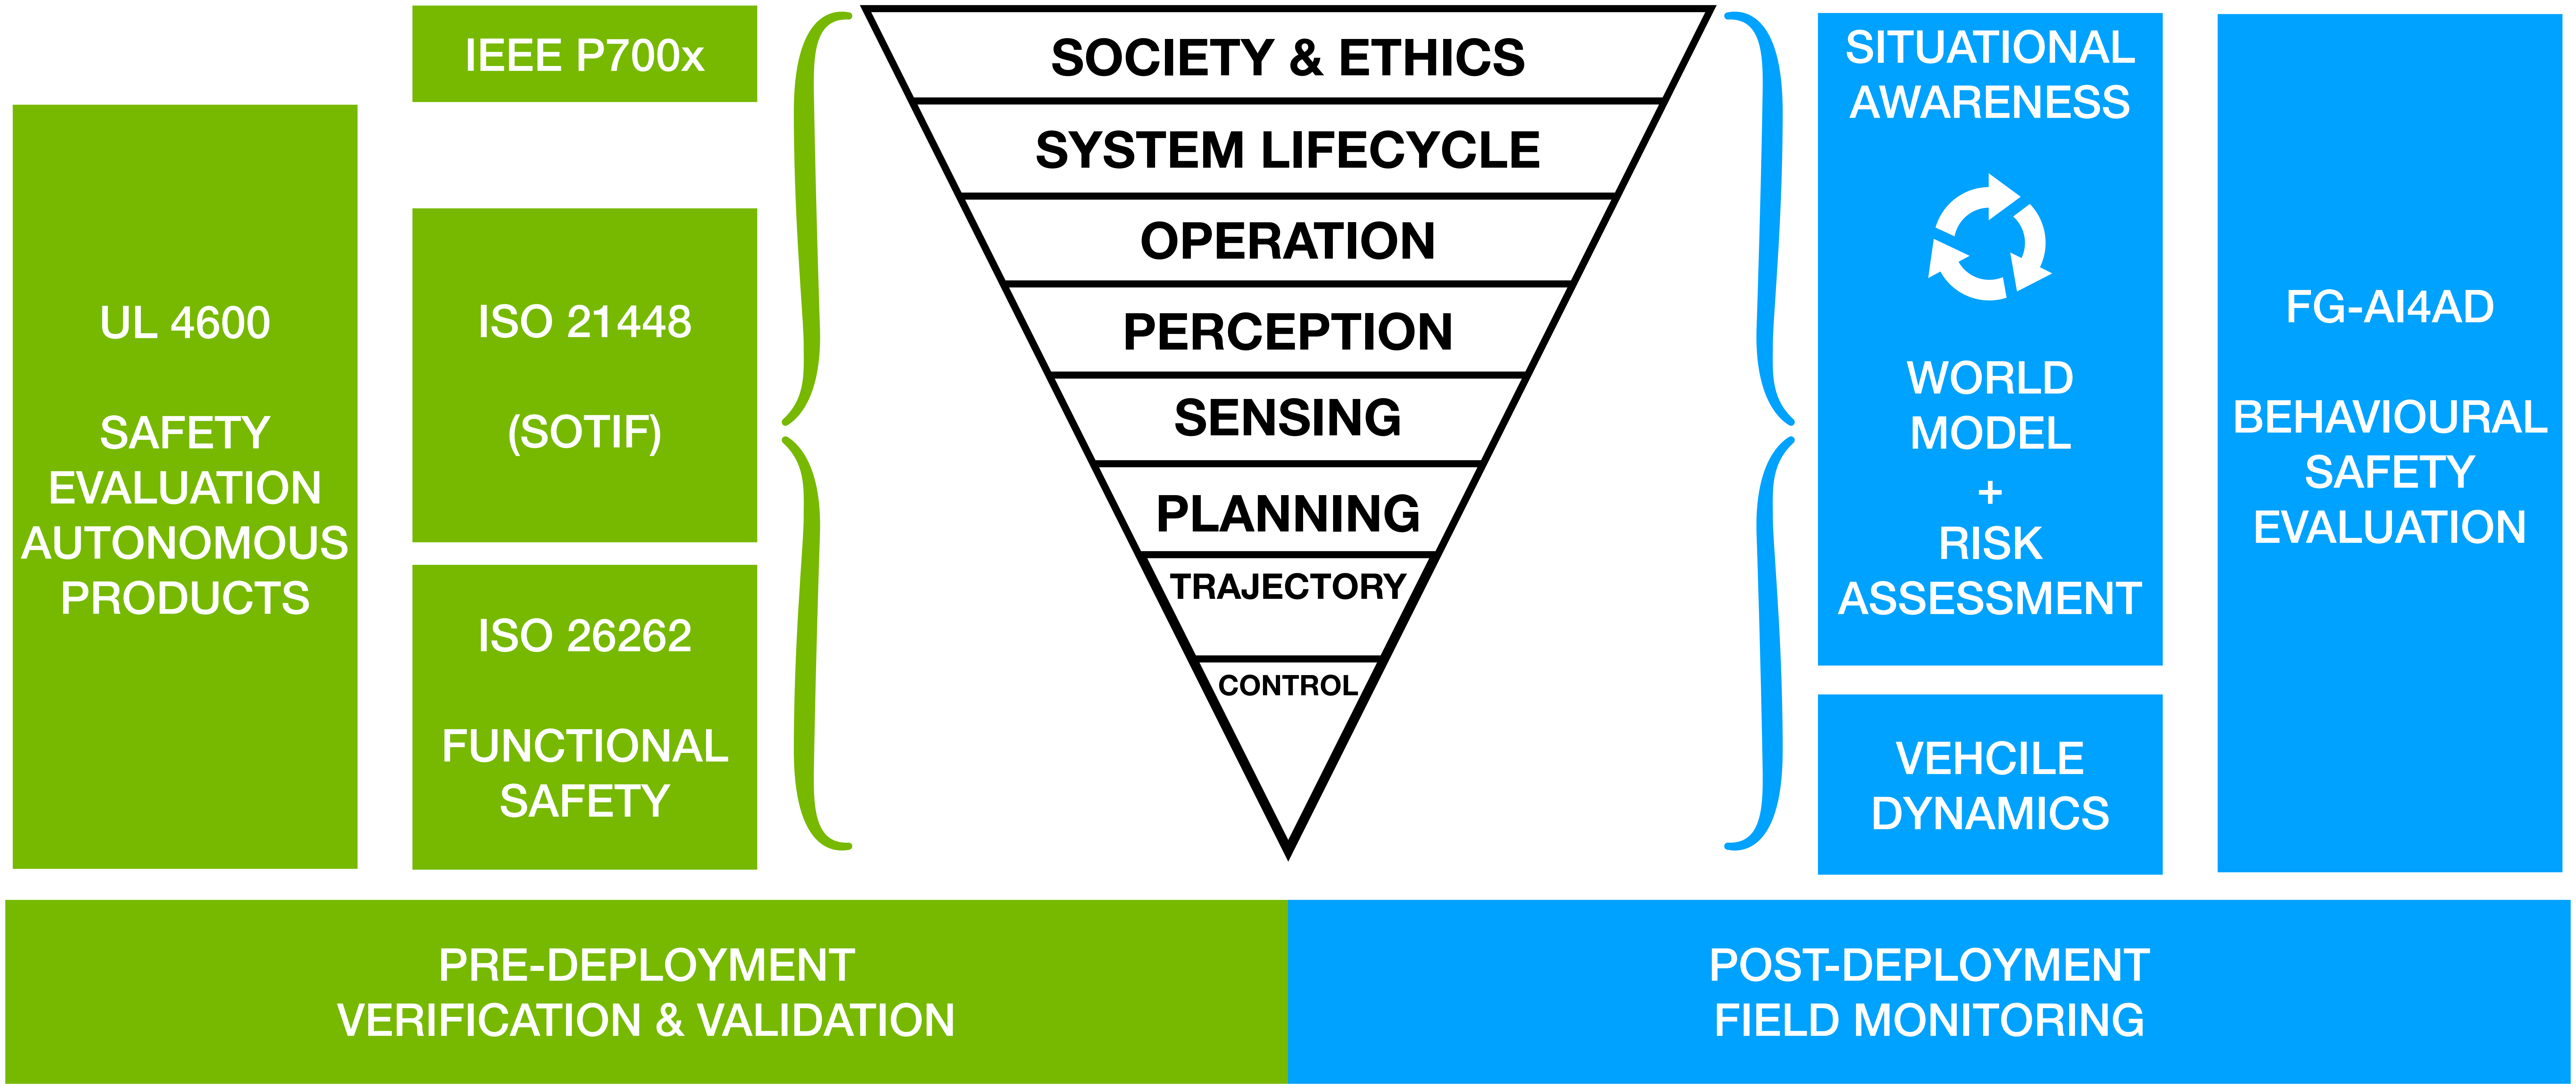

AI for Autonomous and Assisted Driving
AN ITU FOCUS GROUP

Field Monitoring - Leading Measures & Metrics

POST-DEPLOYMENT FIELD MONITORING

LAGGING MEASURES

Observations of safety outcomes or harm

LEADING MEASURES

Reflect performance, activity, and prevention

The Molly Problem

Self-Driving Ethics Revisited

A young girl called Molly is crossing the road alone and is hit by unoccupied self-driving vehicle. There are no eye-witnesses.

– *The Molly Problem*

AI for Autonomous and Assisted Driving
AN ITU FOCUS GROUP

What are the reasonable expectations for what happens next?

– *The Molly Problem: A young girl called Molly is crossing the road alone and is hit by unoccupied self-driving vehicle. There are no eye-witnesses.*

Would you expect the self-driving software;

a) to be aware of the collision?

YES/NO

– *The Molly Problem: A young girl called Molly is crossing the road alone and is hit by unoccupied self-driving vehicle. There are no eye-witnesses.*

Would you expect the self-driving software;

b) to bring the vehicle to a safe stop at the collision site?

YES/NO

– *The Molly Problem: A young girl called Molly is crossing the road alone and is hit by unoccupied self-driving vehicle. There are no eye-witnesses.*

Would you expect the self-driving software;

c) to indicate a hazard to other drivers?

YES/NO

– *The Molly Problem: A young girl called Molly is crossing the road alone and is hit by unoccupied self-driving vehicle. There are no eye-witnesses.*

Would you expect the self-driving software;

d) to alert the emergency services?

YES/NO

– *The Molly Problem: A young girl called Molly is crossing the road alone and is hit by unoccupied self-driving vehicle. There are no eye-witnesses.*

Would you expect the self-driving software;

e) to be able to recall information about the collision required to explain what happened?

YES/NO

– *The Molly Problem: A young girl called Molly is crossing the road alone and is hit by unoccupied self-driving vehicle. There are no eye-witnesses.*

AI for Autonomous and Assisted Driving
AN ITU FOCUS GROUP

Would you expect this recalled information to confirm;

a: the time of the collision?

YES/NO

– *The Molly Problem: A young girl called Molly is crossing the road alone and is hit by unoccupied self-driving vehicle. There are no eye-witnesses.*

Would you expect this recalled information to confirm;

b: the location of the collision?

YES/NO

– *The Molly Problem: A young girl called Molly is crossing the road alone and is hit by unoccupied self-driving vehicle. There are no eye-witnesses.*

Would you expect this recalled information to confirm;

c: the speed of the vehicle at the point of collision?

YES/NO

– *The Molly Problem: A young girl called Molly is crossing the road alone and is hit by unoccupied self-driving vehicle. There are no eye-witnesses.*

Would you expect this recalled information to confirm;

d: if Molly was detected by the software?

YES/NO

– *The Molly Problem: A young girl called Molly is crossing the road alone and is hit by unoccupied self-driving vehicle. There are no eye-witnesses.*

Would you expect this recalled information to confirm;

e: when Molly was detected by the software?

YES/NO

– *The Molly Problem: A young girl called Molly is crossing the road alone and is hit by unoccupied self-driving vehicle. There are no eye-witnesses.*

Would you expect this recalled information to confirm;

f: when the risk of collision was identified?

YES/NO

– *The Molly Problem: A young girl called Molly is crossing the road alone and is hit by unoccupied self-driving vehicle. There are no eye-witnesses.*

Would you expect this recalled information to confirm;

g: when mitigating action was taken to avoid the collision?

YES/NO

– *The Molly Problem: A young girl called Molly is crossing the road alone and is hit by unoccupied self-driving vehicle. There are no eye-witnesses.*

Would you expect this recalled information to confirm;

h: what mitigating action was taken to avoid the collision?

YES/NO

– *The Molly Problem: A young girl called Molly is crossing the road alone and is hit by unoccupied self-driving vehicle. There are no eye-witnesses.*

Would you expect this recalled information to confirm;

i: whether the mitigating action was executed successfully?

YES/NO

– *The Molly Problem: A young girl called Molly is crossing the road alone and is hit by unoccupied self-driving vehicle. There are no eye-witnesses.*

Question #1

Does this align with the three FG-AI4AD behavioral proofs?

– *The Molly Problem: A young girl called Molly is crossing the road alone and is hit by unoccupied self-driving vehicle. There are no eye-witnesses.*

Expected behavioural proofs for AI Software on our roads

Prove AI Software never engages in **careless, dangerous** or **reckless driving behaviour.**

In accordance to Article 7 of the Geneva Convention on Road Traffic ***"not to endanger"***

Prove AI Software **meets, or exceeds**, the performance of a **competent** and **careful** human driver

In accordance with Article 10 of the Geneva Convention on Road Traffic "**reasonable and prudent**" driving

Prove AI Software remains **aware**, **willing** and **able** to **avoid collisions** at all times

In accordance to Article 7 of the Geneva Convention on Road Traffic "**shall avoid all behaviour that might cause damage to persons, or public or private property**."

Question #2

Does this align with NTSB collision investigation requirements?

– *The Molly Problem: A young girl called Molly is crossing the road alone and is hit by unoccupied self-driving vehicle. There are no eye-witnesses.*

NTSB: HWY18MH010, Tempe, Arizona - Uber ATG

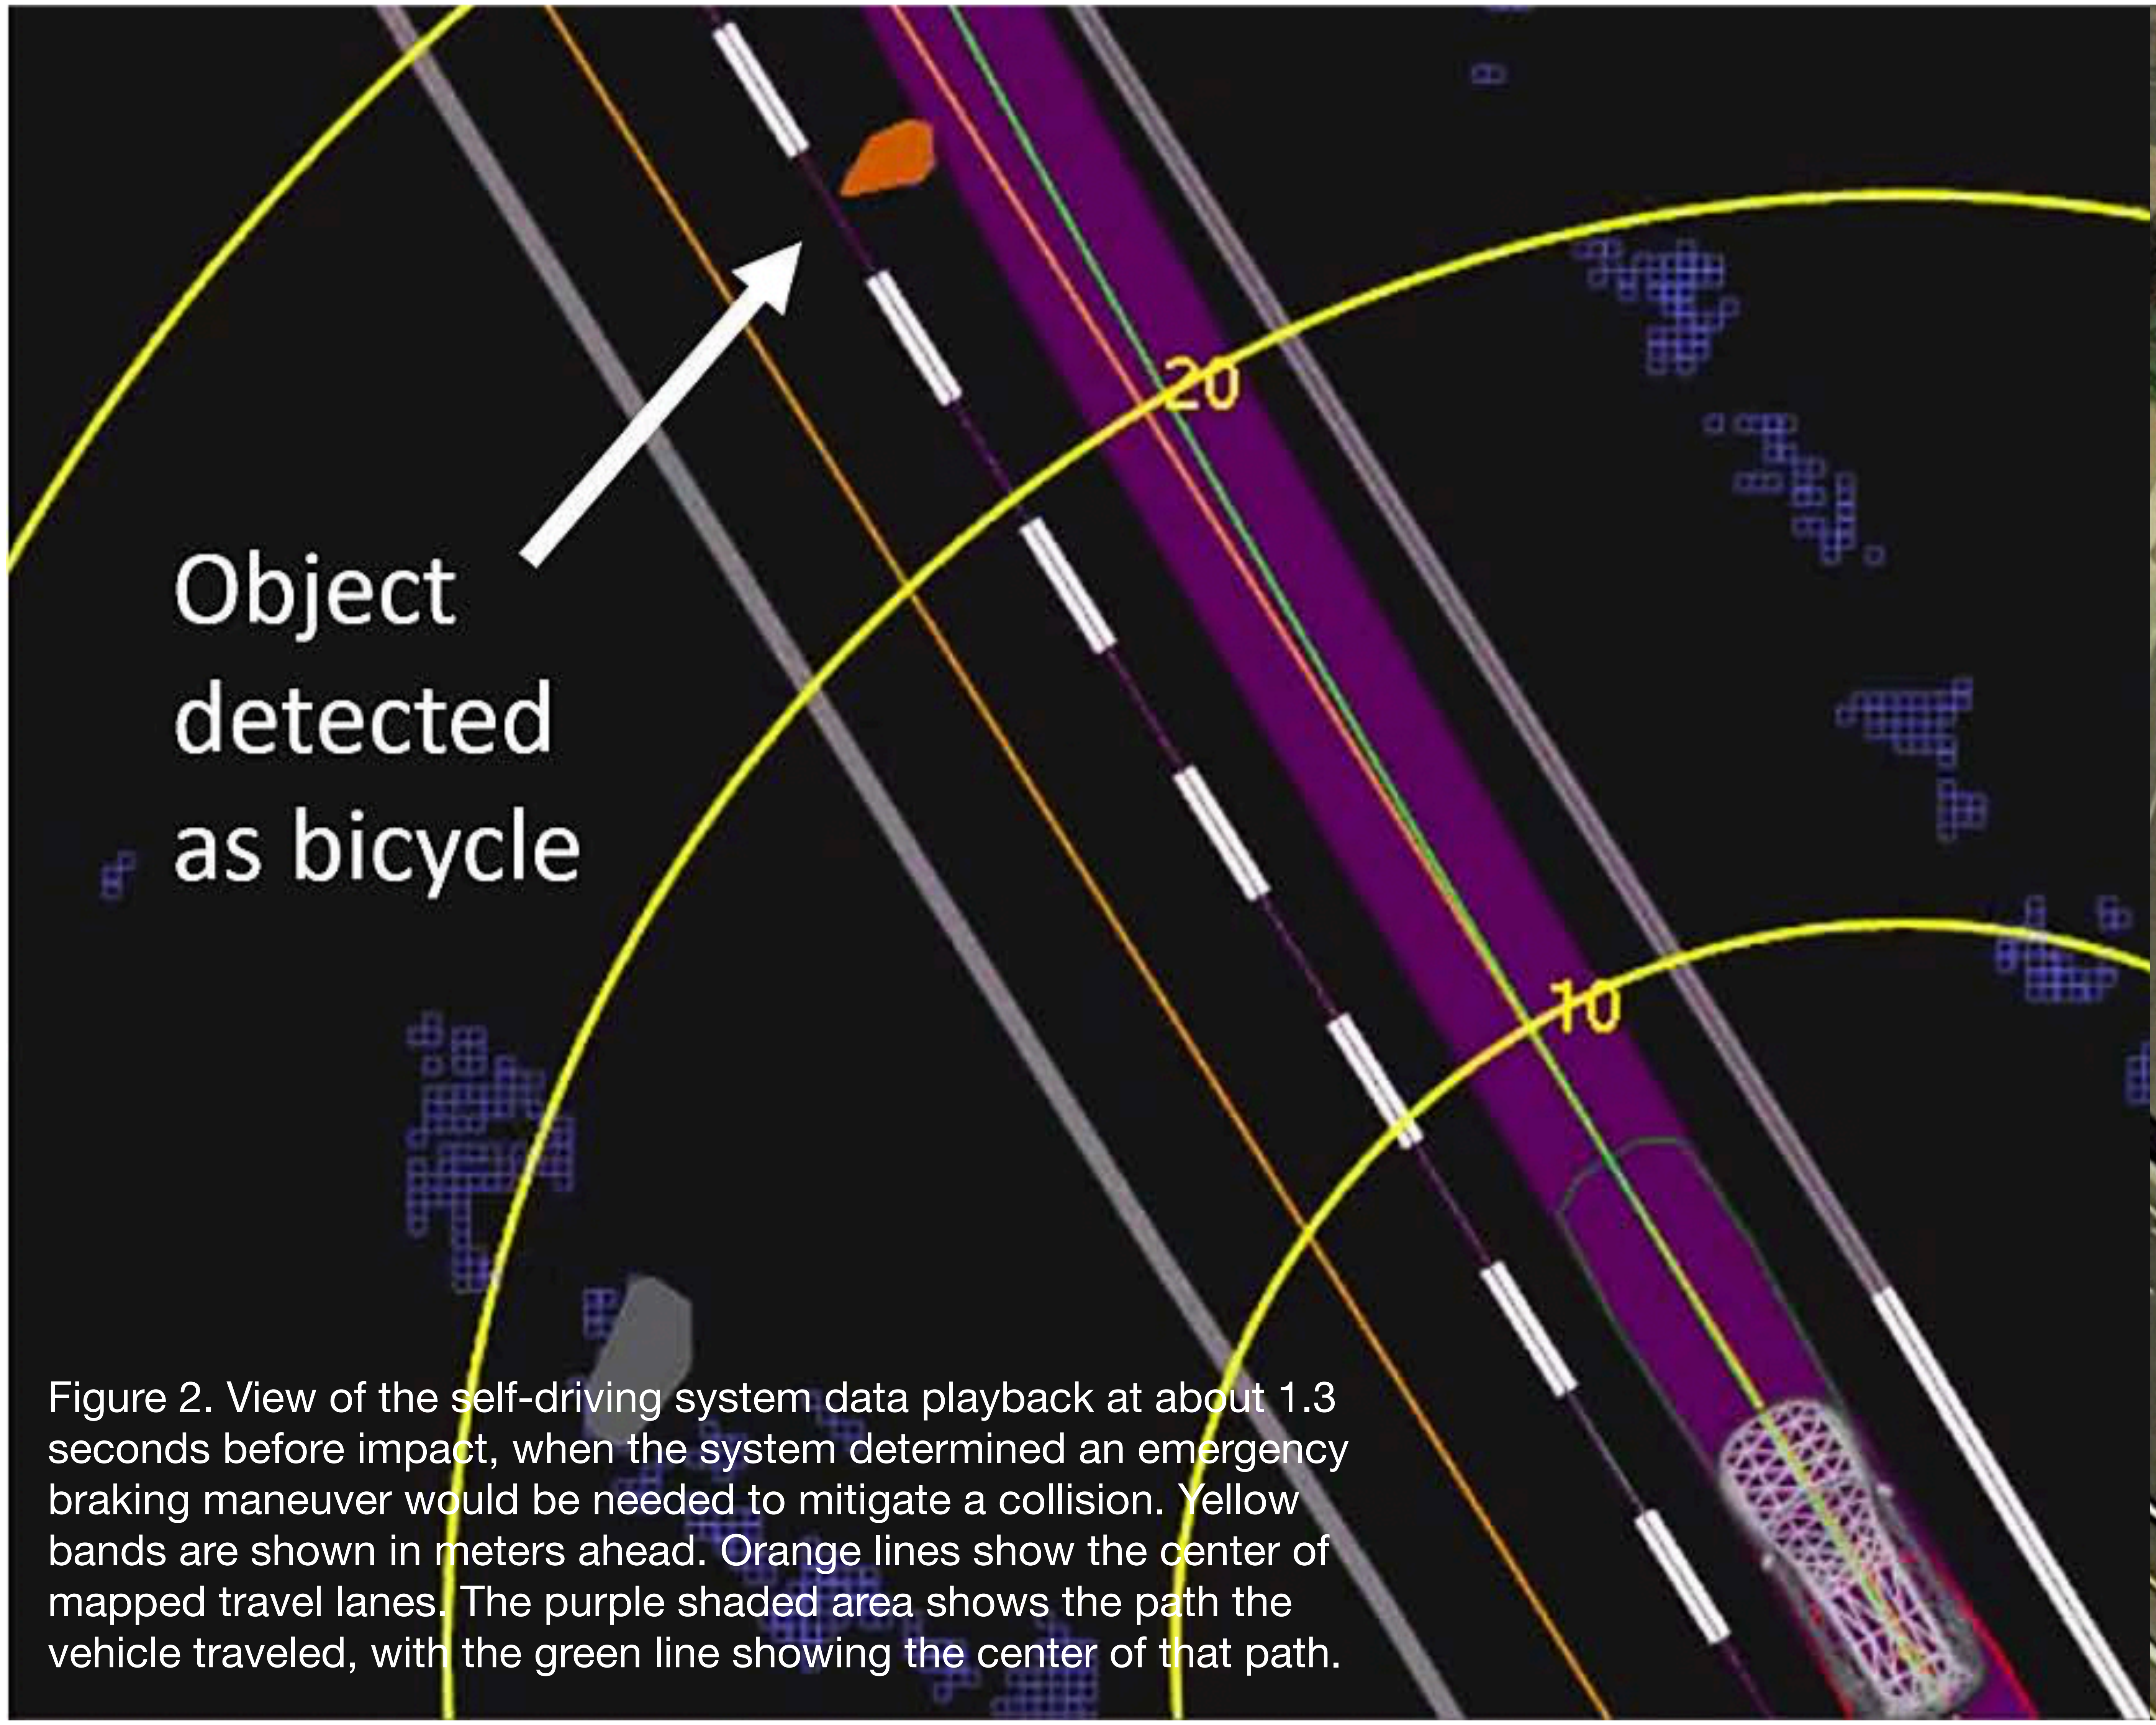

Figure 2. View of the self-driving system data playback at about 1.3 seconds before impact, when the system determined an emergency braking maneuver would be needed to mitigate a collision. Yellow bands are shown in meters ahead. Orange lines show the center of mapped travel lanes. The purple shaded area shows the path the vehicle traveled, with the green line showing the center of that path.

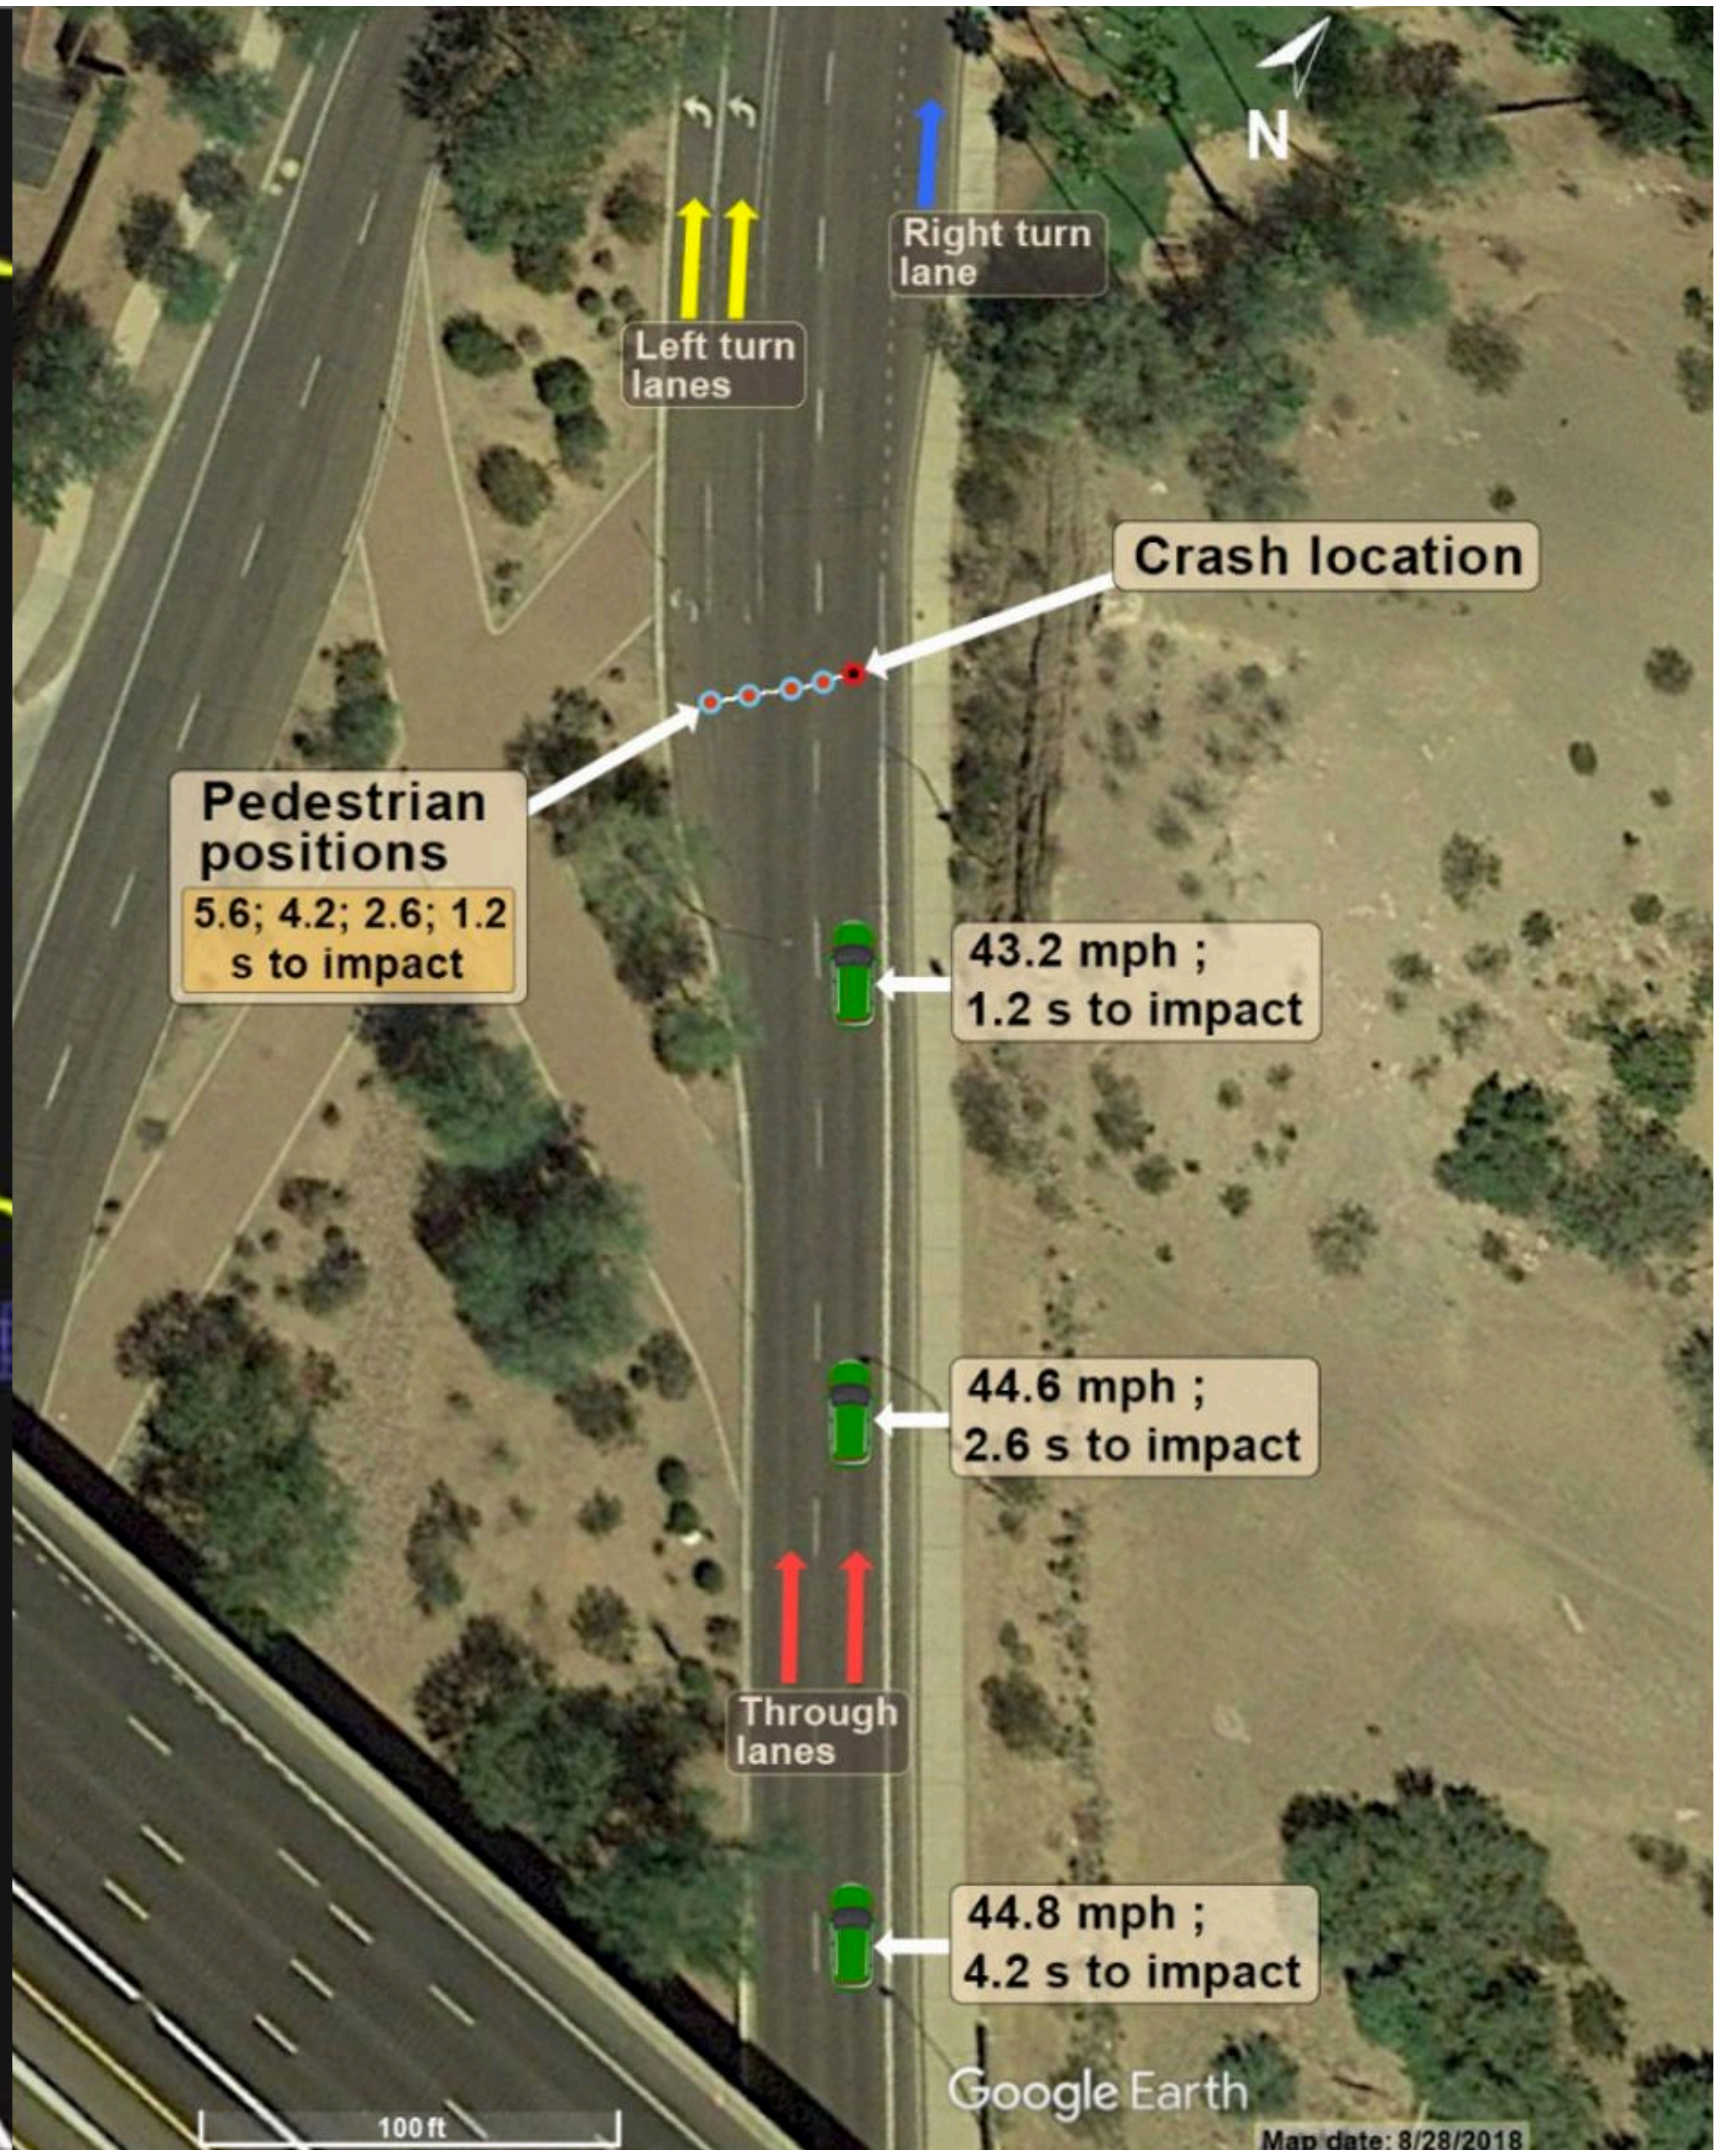

https://www.ntsb.gov/investigations/pages/hwy18fh010.aspx

AI for Autonomous and Assisted Driving
AN ITU FOCUS GROUP

NTSB: HWY18MH010, Tempe, Arizona - Uber ATG

Time (s) relative to impact	Speed (mph)	Classification and Path Prediction[a]	Other Events / Details[b]
-9.9	35		Vehicle begins to accelerate from 35 mph due to an increased speed limit
-5.8	44		Vehicle reaches the speed of 44 mph
-5.6	**44**	**Classification: Vehicle - by radar** **Path prediction: None; not on the path of the SUV**	**Radar makes the first detection of the pedestrian and estimates its speed.**
-5.2	45	Classification: **Other** - by lidar Path prediction: Static; not on the path of the SUV	Lidar detects an unknown object; this is the first detection of that object by lidar, the tracking history is unavailable, and its velocity cannot be determined. ADS predicts the object's path as static.
-4.2	**45**	**Classification: Vehicle - by lidar** **Path prediction: Static; not on the path of the SUV**	**Lidar classifies a detected object as a vehicle; this is a changed classification of an object and without a tracking history. ADS predicts the object's path as static.**
-3.9		Classification: Vehicle - by lidar Path prediction: The left through lane (adjacent to the SUV); not on the path of the SUV	Lidar retains the classification "vehicle", and based on the tracking history and the assigned goal, ADS predicts the object's path as traveling in the left through lane.
3.8 > - 2.7	45	Classification: alternated several times between **vehicle and other** - by lidar Path prediction: alternated between static and left lane; neither were considered on the path of the SUV	The object's classification alternates several times between vehicle and an unknown. At each change, the object's tracking history is unavailable, and ADS predicts the object's path as static. When the detected object's classification remained the same, ADS predicts the path as traveling in the left through lane.
-2.6	**45**	**Classification: Bicycle - by lidar** **Path prediction: Static; not on the path of the SUV**	**Lidar classifies a detected object as a bicycle; this is a changed classification of the object, and without a tracking history. ADS predicts the bicycle's path as static.**
-2.5	45	Classification: **Bicycle** - by lidar Path prediction: The left through lane (adjacent to the SUV); not on the path of the SUV	Lidar retains the classification "bicycle" and based on the tracking history and the assigned goal, ADS predicts the bicycle's path as traveling in the left through lane.

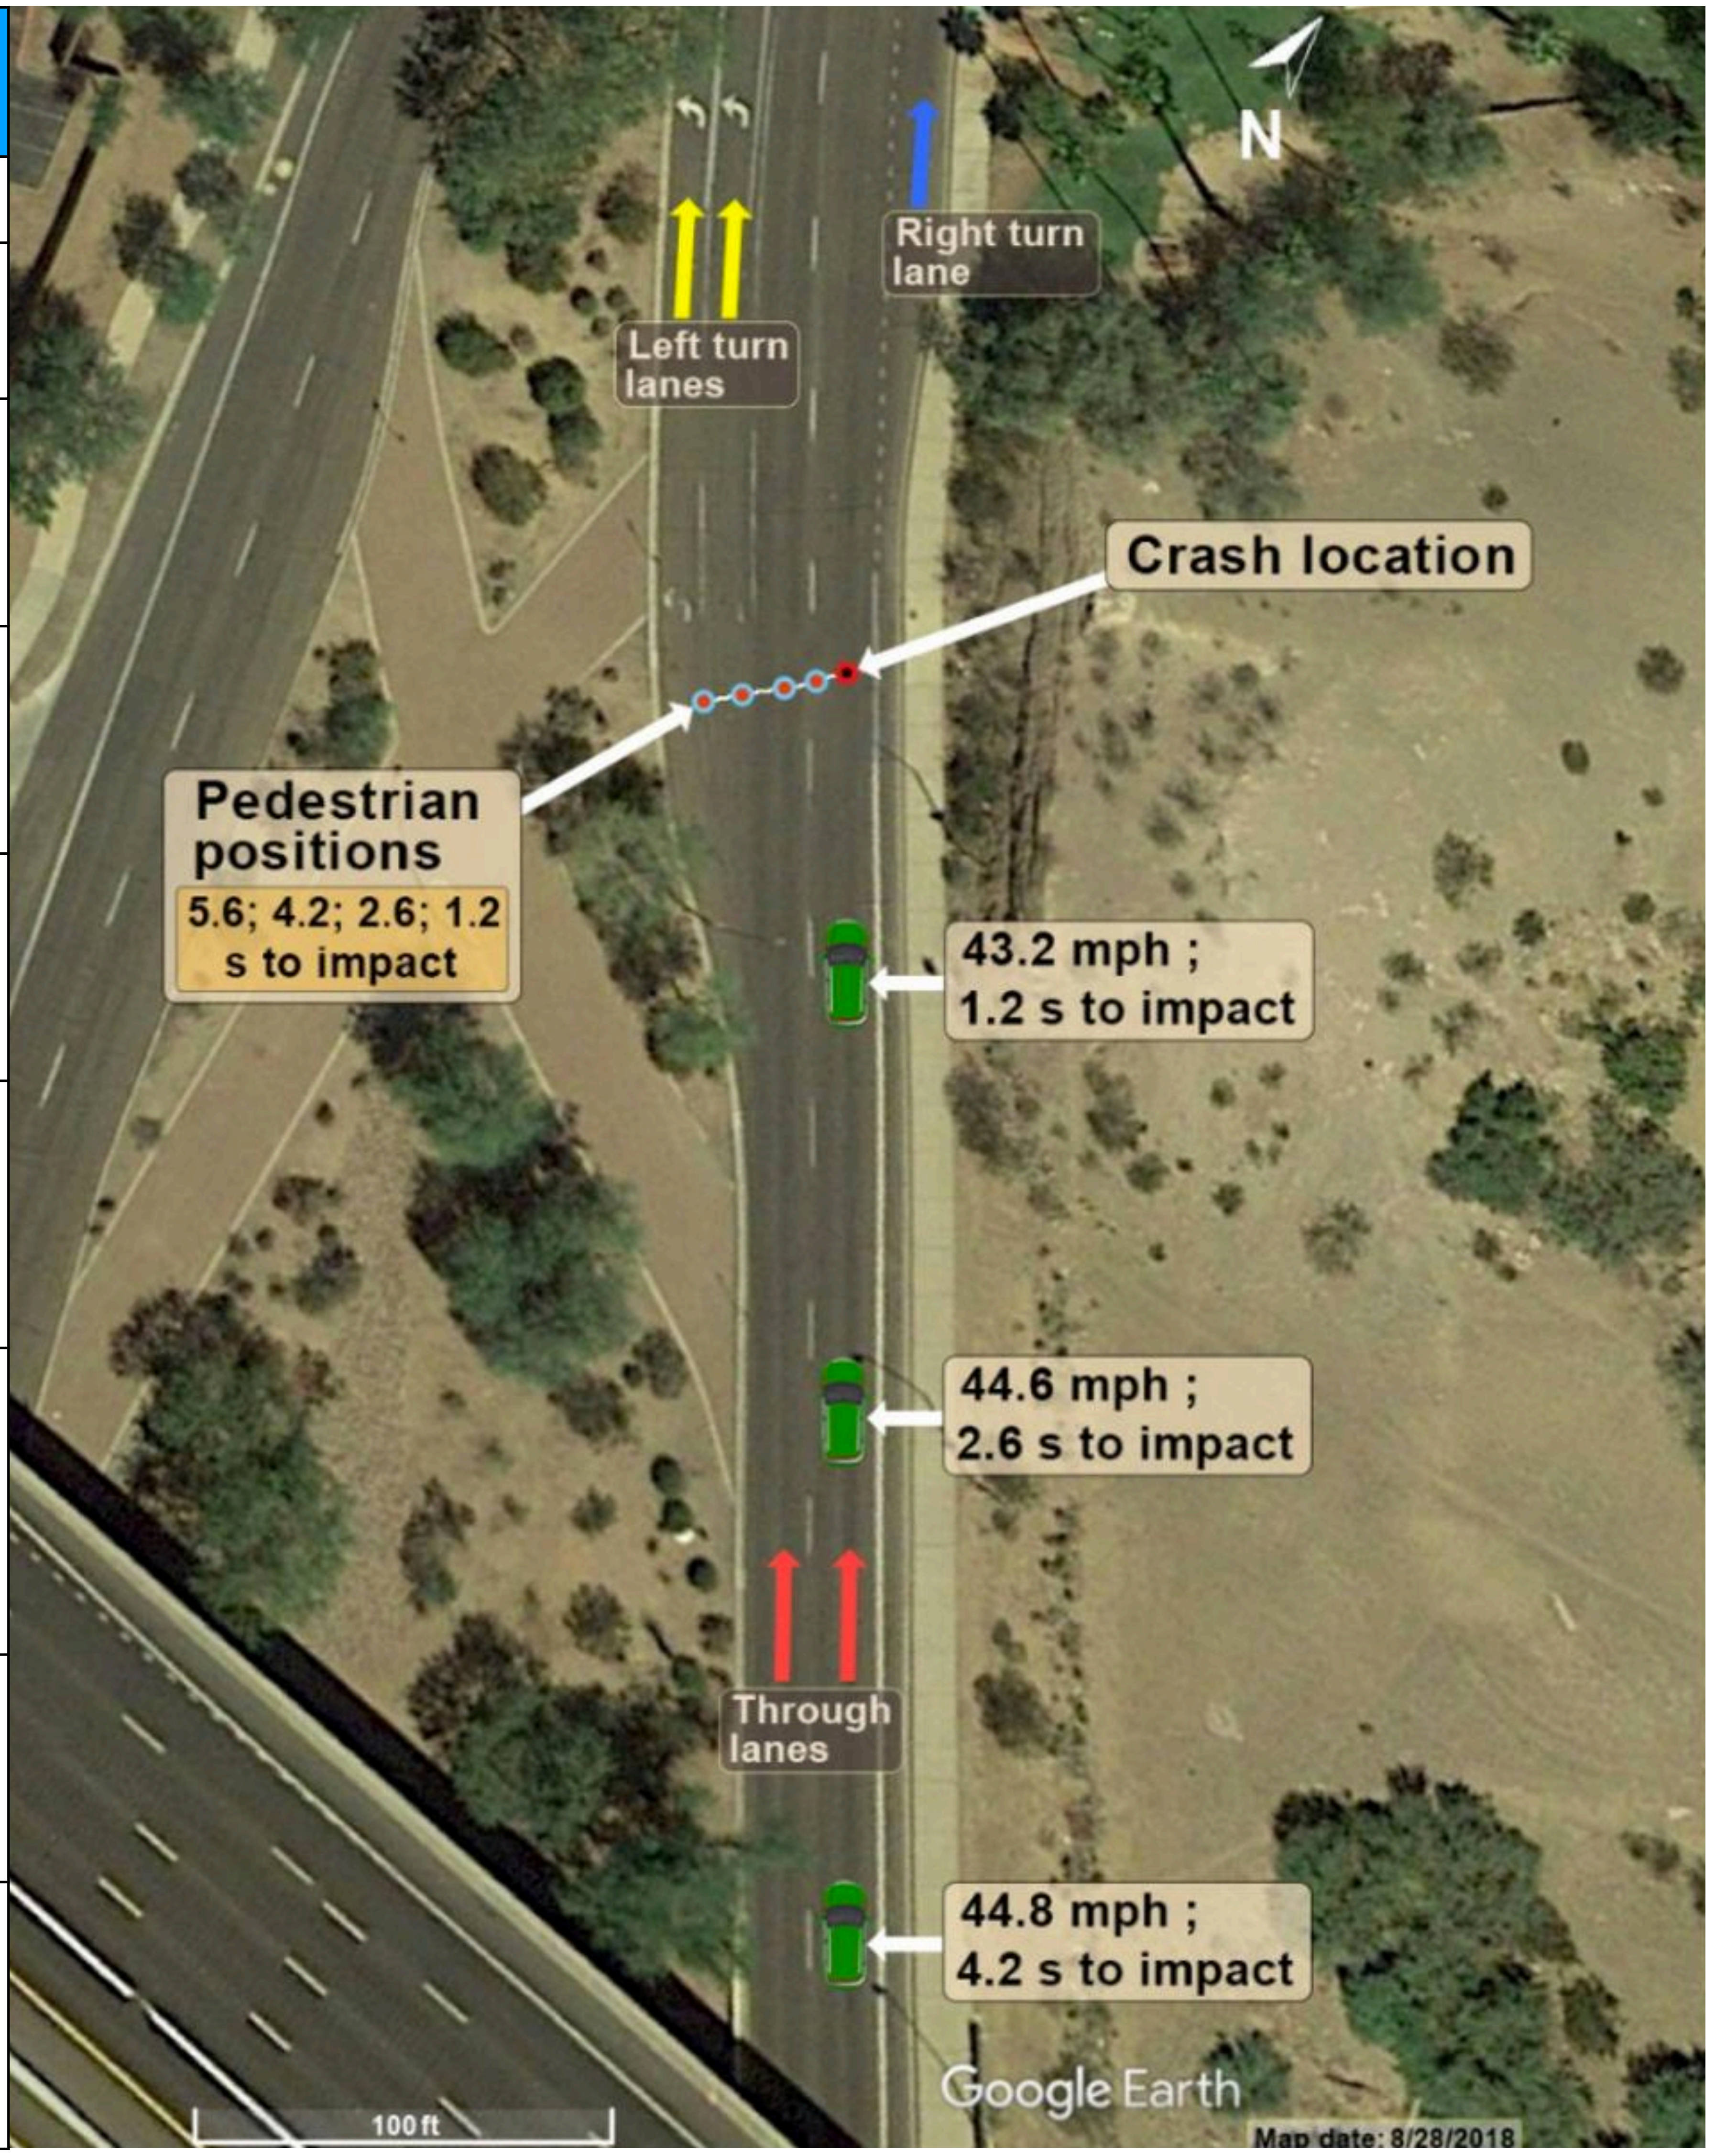

https://www.ntsb.gov/investigations/pages/hwy18fh010.aspx

NTSB: HWY18MH010, Tempe, Arizona - Uber ATG

Time (s) relative to impact	Speed (mph)	Classification and Path Prediction[a]	Other Events / Details[b]
-1.5	*44*[c]	<u>Classification</u>: ***Unknown*** - by lidar <u>Path prediction</u>: *Static*; partially on the path of the SUV	Lidar detects an unknown object; since this is a changed classification, and an unknown object, it lacks tracking history and is not assigned a goal. ADS predicts the object's path as static. Although the detected object is partially in the SUV's lane of travel, the ADS generates a motion plan around the object (maneuver to the right of the object); this motion plan remains valid—avoiding the object—for the next two data points.
-1.2	**43**	**<u>Classification</u>: *Bicycle* - by lidar** **<u>Path prediction</u>: The travel lane of the SUV; fully on the path of the SUV**	**Lidar detects a bicycle; although this is a changed classification and without a tracking history, it was assigned a goal. ADS predicts the bicycle to be on the path of the SUV.** **The ADS motion plan—generated 300 msec earlier—for steering around the bicycle was no longer possible; as such, this situation becomes hazardous.** **- Action suppression begins**
-0.2	40	<u>Classification</u>: ***Bicycle*** - by lidar <u>Path prediction</u>: The travel lane of the SUV; fully on the path of the SUV	Action suppression ends 1 second after it begins. The situation remains hazardous; as such, ADS initiates a plan for vehicle slowdown. *An auditory alert was presented to indicate that the controlled slowdown was initiating* [d]
-0.02	39		Vehicle operator takes control of the steering wheel, disengaging the ADS.
IMPACT			
0.7	37		Vehicle operator brakes

[a] Only changes in object classification and path prediction are reported in the table.The last reported values persist until a new one is reported.

[b] The process of predicting a path of a detected object is complex and relies on the examination of numerous factors, beyond the details described in this column

c The vehicle started decelerating due to the approaching intersection, where the pre-planned route includes a right turn at Curry Road. The deceleration plan was generated 3.6 seconds before impact.

[d] *While the system generated a plan for the vehicle slowdown, due to a slight communication delay, the data is unclear on whether the implementation of the slowdown plan started before the operator took control prior to impact.*

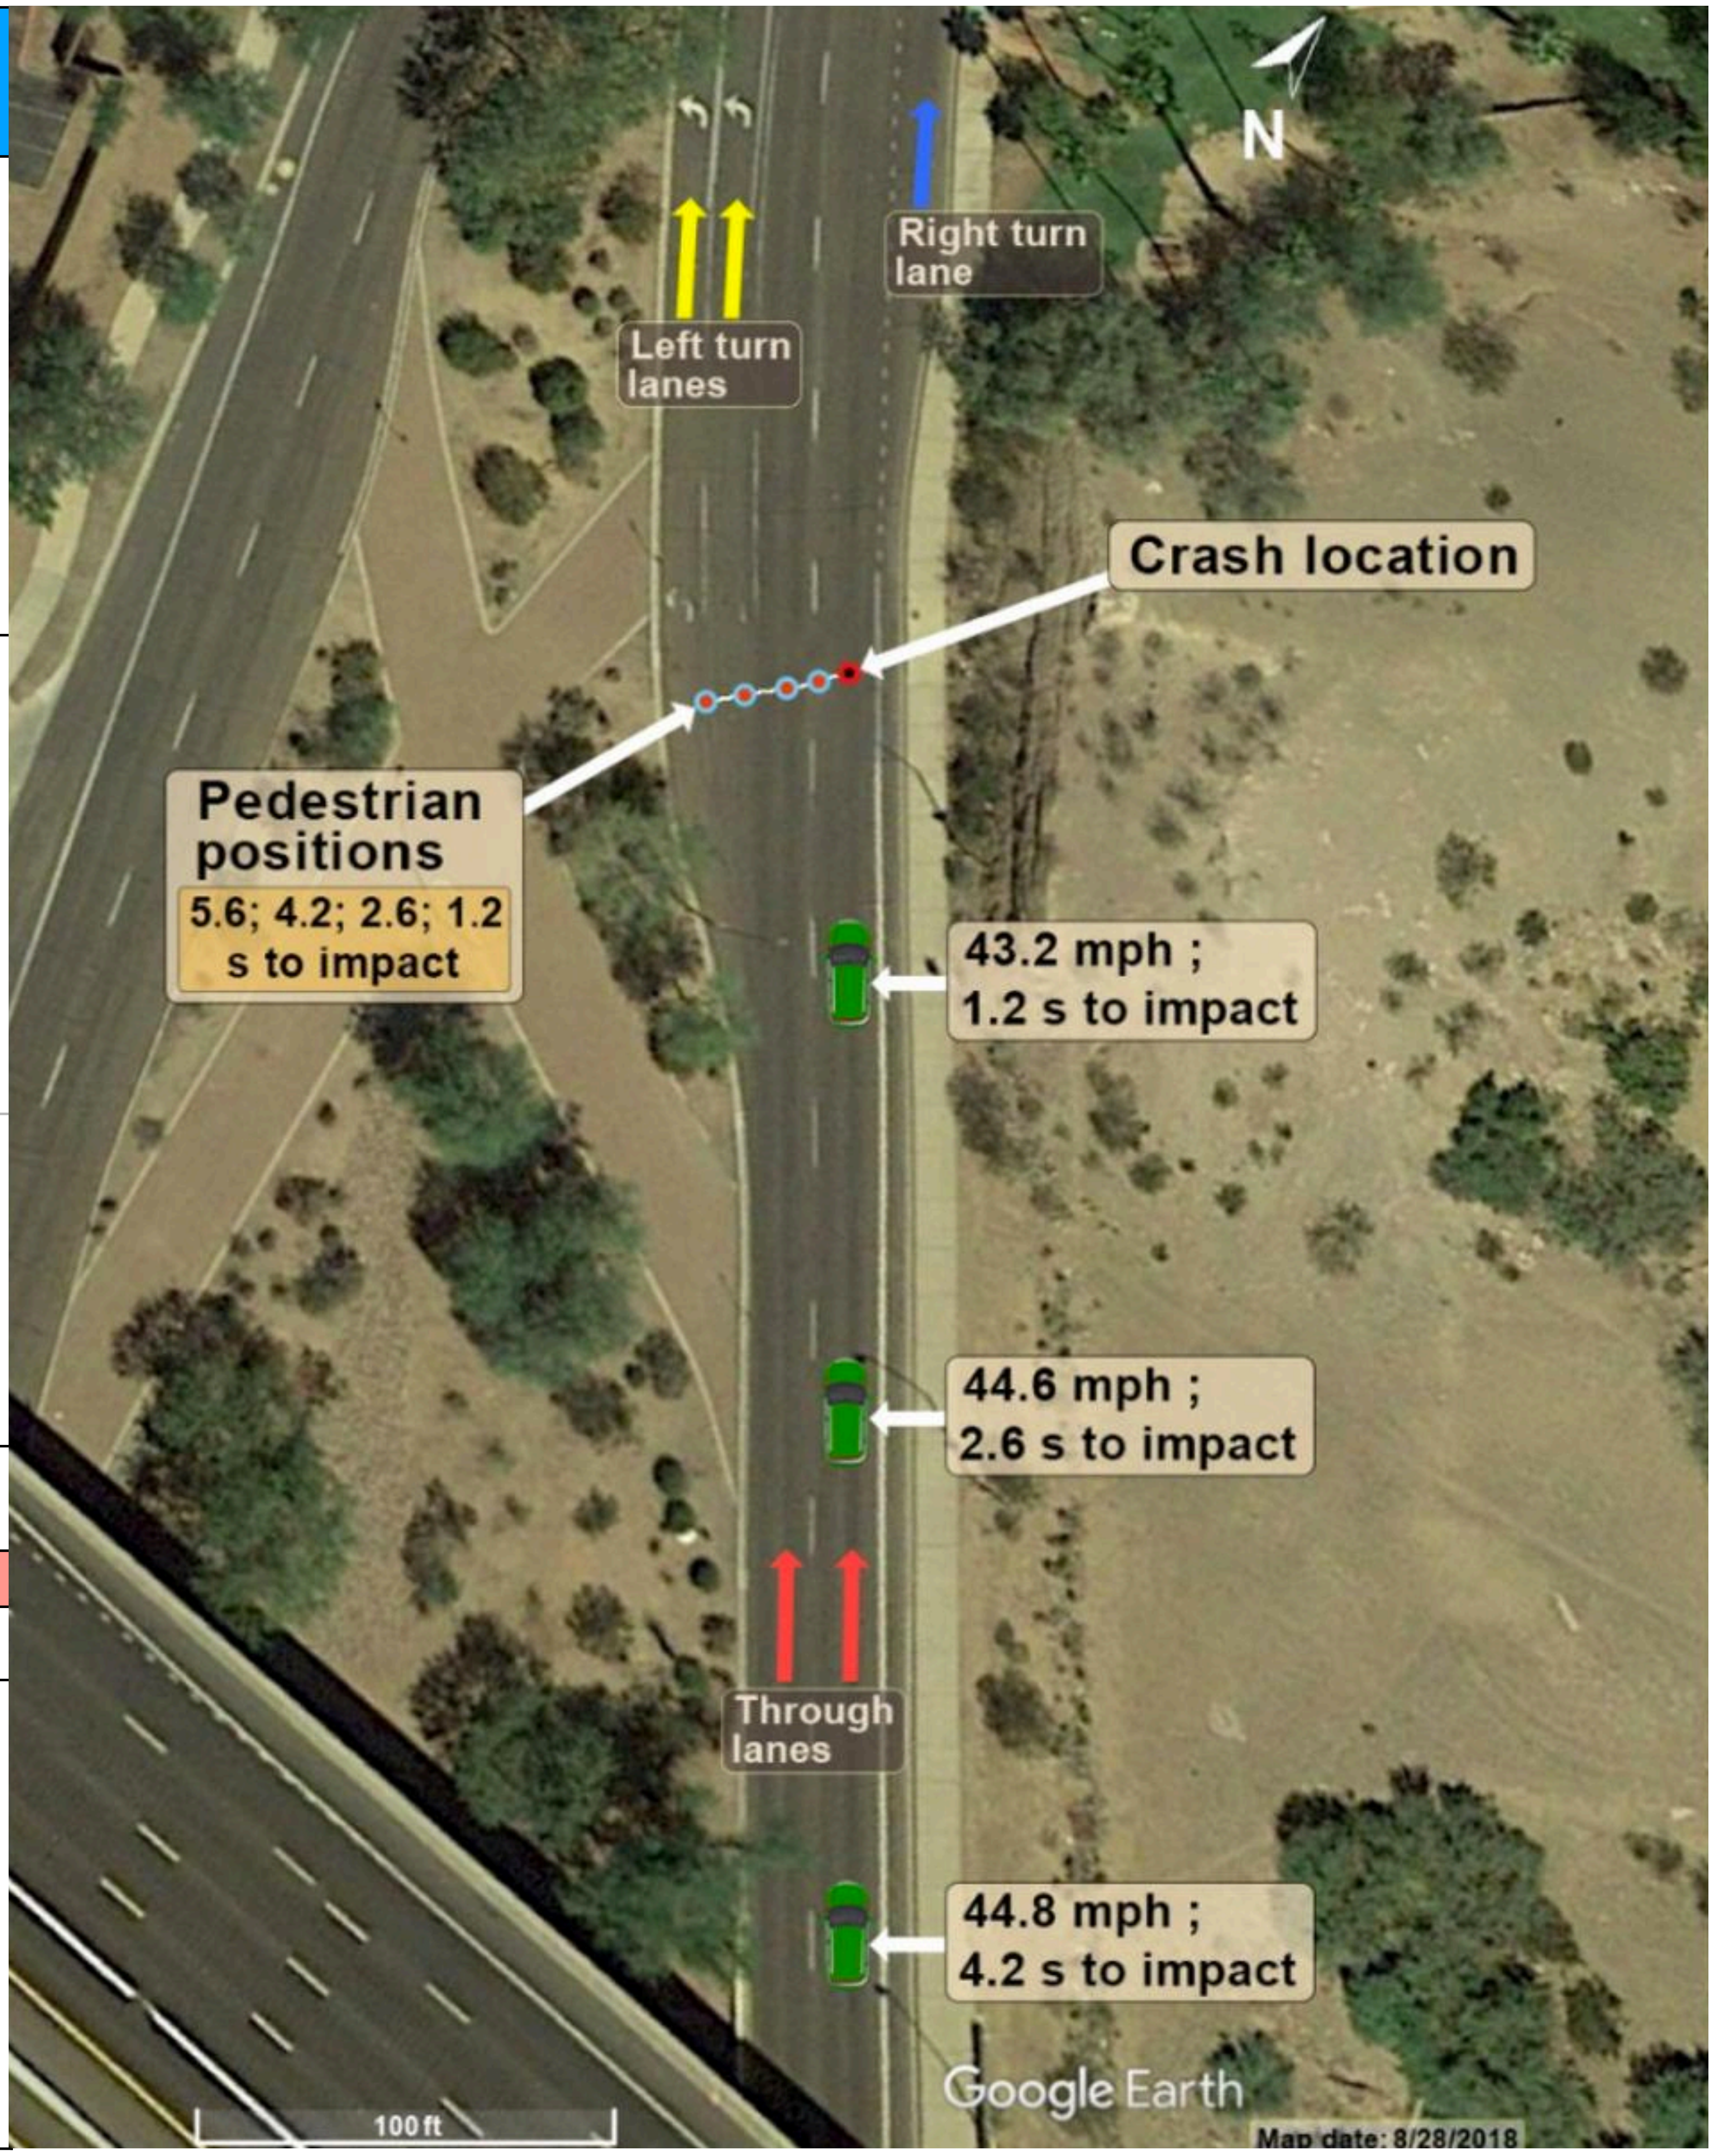

https://www.ntsb.gov/investigations/pages/hwy18fh010.aspx

Question #3

Does this align with proposed data sources for continual evaluation?

– *The Molly Problem: A young girl called Molly is crossing the road alone and is hit by unoccupied self-driving vehicle. There are no eye-witnesses.*

ITU-T Focus Group on AI for Autonomous & Assisted Driving

Data Sources for Behavioral Evaluation

Situation

Does the AD understand the situation?

Extracted from the local world model. Where is the vehicle and where are all the other static/dynamic objects?

Risk

Does the AD understand the level risk?

Prediction of risk presented by the situation. Levels of uncertainty in the models used to make the prediction.

Action

Does the AD execute the correct action?

Control inputs to the vehicle and resultant dynamics.

Outcome

Is the real-world risk in the acceptable?

Using real-time continual monitoring of three input data sources.

Question #4

Does this align with needs of developers, insurers & regulators?

– *The Molly Problem: A young girl called Molly is crossing the road alone and is hit by unoccupied self-driving vehicle. There are no eye-witnesses.*

Leading Measures: Behavioural Safety Evaluation

AI for Autonomous and Assisted Driving
AN ITU FOCUS GROUP

AD BEHAVIOURAL SAFETY EVALUATION

The continuous, on-road, data-driven, science-based, transparent safety evaluation of autonomous and assisted driving software for developers, insurers and regulators.

DEVELOPERS

Applicable for scenario based testing in simulation and proving grounds.

Applicable for on-road real-world verification and validation.

Applicable to post-deployment field monitoring and edge-case scenario capture.

Maximum protection for proprietary technology in perception, planning and control

INSURERS

Common foundation with telematics usage based insurance based on vehicle dynamics.

Places resultant vehicle dynamics in the context of the road traffic situation.

Provides direct insight into the perceived and predicted risk within the AD Software.

Provides key metrics for fully independent assessment of risk after the event

REGULATORS

International harmonisation for behavioural safety of deployed autonomous and assisted driving software for WP.1 & WP.29

Data driven, science-based and transparent

Common data protocol for publication of leading metrics from AD Software

Any country, any Operational Design Domain, any technology architecture

Question #5

Does FG-AI4AD extend existing UNECE EDR/DSSAD provisions?

– *The Molly Problem: A young girl called Molly is crossing the road alone and is hit by unoccupied self-driving vehicle. There are no eye-witnesses.*

YES.

EDR will not trigger on pedestrian impact
DSSAD records only the entity responsible for the DDT

FG-AI4AD provides continual monitoring and threshold based recording for near-miss as well as collision events.

– *The Molly Problem: A young girl called Molly is crossing the road alone and is hit by unoccupied self-driving vehicle. There are no eye-witnesses.*

Question #6

Does FG-AI4AD extend SAE J3197 provisions for the Automated Driving System Data Logger?

– *The Molly Problem: A young girl called Molly is crossing the road alone and is hit by unoccupied self-driving vehicle. There are no eye-witnesses.*

YES.

SAE J3197 records only upon a collision events

FG-AI4AD provides continual monitoring and threshold based recording for near-miss as well as collision events

– *The Molly Problem: A young girl called Molly is crossing the road alone and is hit by unoccupied self-driving vehicle. There are no eye-witnesses.*

AI for Autonomous and Assisted Driving
AN ITU FOCUS GROUP

THANK YOU. STAY SAFE. STAY HEALTHY.

Chair ITU FG-AI4AD Bryn Balcombe: bryn@ada.ngo

General mailing list: fgai4ad@lists.itu.int

Dedicated secretariat email: tsbfgai4ad@itu.int

Dedicated webpage: www.itu.int/en/ITU-T/focusgroups/ai4ad