Assessing the ODD of AV perception sensors with confidence

If you wish to make a reference to this paper, please use:

"Assessing the ODD of AV perception sensors with confidence, Sensor Assurance Framework project, Met Office & National Physical Laboratory, 2024."

Final Version: May 2024

Contents

About the title

The use of the word 'confidence' in the title has dual meaning, relating to both undertaking sensor assurance confidently and, also, the use of the word in a statistical context. In this latter context, confidence is determined through the systematic assessment of uncertainty, which is the phrase we adopt in this paper, however the authors felt that "*Assessing the ODD of AV perception sensors with uncertainty*" was a less desirable title.

Note about formatting

The paper makes use of footnotes so as not to clutter the flow of the arguments made.
The information they contain is important but not essential to understand the key messages.

Audience and assumed knowledge

This paper assumes working familiarity with the concepts of the Operational Design Domain (ODD), including the desire to evaluate this quantitatively for the purposes of safety assurance, and the role of the virtual test environments (VTE) as a verification tool. It makes reference to the Sensor Assurance Framework (SAF) project, a joint undertaking by the Met Office and National Physical Laboratory (NPL) and funded by the Centre for Connected and Autonomous Vehicles (CCAV). The purpose of the SAF is to a develop a usable and reliable framework for characterising perception sensor performance in different weather-related conditions, although the arguments made here apply to all weather-sensitive aspects of AV functionality.

Executive Summary

The safe operation of connected and autonomous vehicles (CAVs) requires that the impact of the weather on their perception sensors is quantitatively understood and represented in the Operational Design Domain (ODD). Since 2020, the Sensor Assurance Framework[1] (SAF) project has been developing methods for the capturing this weather sensitivity, including an extensive set of measurements of sensor performance in a wide range of comprehensively measured weather conditions.

There is still much value to be extracted from the data gathered to date, however our continuing analysis and our engagement with stakeholders clearly points to the following:

- The effect of weather on sensor performance is immensely ***complex***[2], including many non-obvious interactions between multiple variables and factors.

- Even if we were able to build a 'perfect' Virtual Test Environment (VTE) that faithfully transformed *input* weather and scenery parameters into *output* sensor responses[3], the number of inputs required would be significantly greater than are currently described in existing ODD taxonomies in order to fully explain sensor responses to weather.

- The successful integration of the different test domains (calibration lab, real world[4], emulated[5], virtual) requires *all* stakeholders to achieve a shared understanding of their purpose (and limitations) and also, *critically*, a consistent approach to how the uncertainties associated with them are co-managed.

This discussion paper, produced jointly by the Met Office and the National Physical Laboratory, explores some aspects of this problem in order to stimulate discussion and mutual understanding. We would welcome comments, including any points of contention, on all aspects of this paper.

---

[1] The Sensor Assurance Framework (SAF), project is a joint programme of work undertaken by the UK's national meteorological service and metrology institute, the Met Office and National Physical Laboratory. It is funded by the Centre for Connected and Autonomous Vehicles (UK).

[2] ***Complex*** is used specifically in the sense of the *Cynefin* framework. Wikipedia has a good description, as does the Harvard Business Review article here: https://hbr.org/2007/11/a-leaders-framework-for-decision-making. As described in HBR, in the Cynefin approach "leaders who try to impose order in a complex context will fail, but those who set the stage, step back a bit, allow patterns to emerge, and determine which ones are desirable will succeed". Once a pattern has emerged, it is possible to move the problem space into the ***Complicated*** regime.

[3] This qualification is significant – there will never be a perfect VTE

[4] By 'real world' we mean where sensors are exposed to natural weather i.e. outdoors

[5] By 'emulated' we mean creating physical approximations to real weather e.g. rainfall generation from sprinkler systems

Introduction

Since 2020, the Sensor Assurance Framework (SAF) project has been working to deliver a usable and reliable framework to characterise connected and autonomous vehicle (CAV) sensor performance in different weather-related conditions. This has included operating a weather-sensor testbed at Cardington UK for approximately two years. The testbed comprises a comprehensive network of meteorological measurements and a large number of camera, lidar and radar sensors that view well-characterised targets at a range of distances and viewing directions and over a wide range of weather conditions.

This short discussion paper provides some insights regarding the weather-sensor interaction from our testbed data. It is illustrated using examples of the rainfall-radar and rain/fog-lidar interactions, however the concepts generalise across all weather and sensor types and, in fact, all weather-sensitive aspects of AV system performance. We consider this paper as preliminary work ahead of proposing a structure for the weather-sensor test ecosystem and a pragmatic uncertainty framework, which themselves are key steps in generating a reliable and usable means to support authorisation by demonstrating whether CAVs are operating within their ODD.

The paper first recognises *pre-eminence of the ODD*, a core construct for setting standards and aiding collaboration in CAV safety, before raising the critical issue that the perfectly sensible desire to keep the *ODD simple* is a primary driver of uncertainty in ascribing a definitive ODD boundary. This is made more challenging by *the complexity* in the interaction between the sensor, environment and the scenery and actors. Once an acceptable level[6] of ODD simplicity is achieved, the importance of *harnessing uncertainty* is then discussed, before an initial exploration of how the *different kinds of physical testbed,* such as those we are developing within the SAF help to characterise this uncertainty. Finally, the role of *VTEs as a tool to explore the uncertainty budget* is discussed. *Annex A* provides a visualisation of the multiple sources of uncertainty and *Annex B* demonstrates the differing strengths of testbeds.

---

[6] By ‘acceptable’ we mean that it is understood and used consistently by all users of the ODD

The pre-eminence of the ODD

The Operational Design Domain (ODD) lies at the very centre of safety assurance of CAV systems, including perception sensors. The accepted definition of the ODD is:

> *"Operating conditions under which a given driving automation system or feature thereof is specifically designed to function, including, but not limited to, environmental, geographical, and time-of-day restrictions, and/or the requisite presence or absence of certain traffic or roadway characteristics."* ***SAE J3016 (2021)***

ISO 34503[7] attempts to structure the environmental ODD in as few-as-possible ODD parameters for two reasons:

- to keep the ODD description manageable and, to a large extent, human-readable
- because it is imperative that all ODD parameters must be measurable[8] when and where a CAV is in operation

Our early SAF work influenced these standards by supporting the specification of ODD weather parameters that are at the very least *unambiguous in their interpretation*[9]. However there remains a significant challenge - the current set of ODD parameters are insufficient to fully account for the entirety of impact on the performance of a given sensor (or combination of sensors) from environmental factors. An inescapable consequence is that a *compromise must be made* between simplicity and accuracy, implying that:

- we must find an optimal balance between the simplicity of the ODD parameter set and the consequent loss of accuracy in the ability of these parameters to predict sensor performance;
- we must reflect the implications of this compromise in the design of the test ecosystem and the interpretation of its results;
- we must characterise the remaining uncertainty that arises from our choice of ODD parameters.

We also note, there will always be a number of ODD taxonomies and we must be able to understand the implications of the above bullet points *given* a particular choice of taxonomy.

The aim of the remainder of this paper is to illustrate the above with examples from the Cardington testbed and then provide commentary on the practical implications. These are offered to promote discussion; they are *not* intended to be final recommendations. It is felt that an industry- and regulatory-wide discussion and collective understanding will significantly accelerate progress in CAV V&V[10].

---

[7] ISO 34503:2023 Road Vehicles. Test scenarios for automated driving systems. *Specification for operational design domain*

[8] A more appropriate term would be that all ODD parameters can be *estimated* by some traceable means.

[9] For example, BSI 1883 & ISO34503 insist that it is insufficient merely to use the units of mm/h for rainfall intensity – it is necessary to *also* specify the time and area over which that has been evaluated so that ODD values can be meaningfully compared. There remains work to be done here, but for the purposes of this paper they will be treated as solved.

[10] V&V: Verification and validation

ODD simplicity vs accuracy (or confidence)

We will draw on the example of rainfall attenuation of CAV radar to illustrate the discussion.

Rainfall affects sensor performance through multiple mechanisms. These include path loss of the transmitted and reflected beam due to raindrop absorption and scattering; refraction and/or internal reflection due to water films or droplets on the sensor and target surfaces; multi-path reflections from wet surfaces; and the reflectivity of the rain surrounding the target.

We will consider *only* the path loss in this section as it is sufficient to demonstrate our point and we will assume that we are using an unambiguous definition of rainfall intensity[11]. It should be remembered that each of the mechanisms may respond to different properties of rainfall and will need to be considered at some point.

Rainfall intensity is the rate of accumulation of rainwater falling onto a horizontal surface. It can be calculated by considering the volume of each individual rainfall droplet and its fall speed as a function of droplet size and then summing over all raindrop sizes, described by a drop size distribution (DSD). **Path loss** is caused by the absorption and scattering of the radar beam as it encounters each raindrop, and it is also summed over the complete DSD. The contribution of each rain droplet to radar path loss is a highly non-linear function of droplet size. The result is that the relationship *between* rainfall intensity and path loss is non-unique; it is possible to have many path-loss values for a single rainfall rate. This is illustrated in Figure 1, where the two-way path loss coefficient for a 77 GHz radar is shown for many real DSDs, all measured in the UK.

The figure shows that real DSDs are indeed capable of generating a very wide spread of path loss values for the same headline rainfall rate[12]. The loss-rain rate relationship that results from the commonly adopted Marshall-Palmer "M-P" (1948) DSD[13] is also shown in the figure. The graphs may be interpreted as follows:

- For any given rainfall rate, the M-P path loss is the single value that might be inferred if only a **single** parameter, rainfall rate, is used to describe rain in the ODD, instead of using the full detail of the DSD itself.
- The spread of the other path loss values around the line is indicative of the uncertainty that arises from *not* using additional ODD parameters to capture the sensitivity of path loss to droplet size i.e. ***the potential loss of accuracy caused by ODD simplification***.

---

[11] We assume it is the average intensity of rainfall as measured by a standard rain gauge over one minute.

[12] Obviously, this spread is only significant at the margins of sensor performance, i.e. near the ODD boundary. However, if other pathways such as water films on the radar radome or target are also significant, this spread may be consequential at lower rainfall intensities.

[13] Marshall, J.S., and W.M. Palmer, 1948: The distribution of raindrops with size. *J. Meteor.*, **5**, 165-166

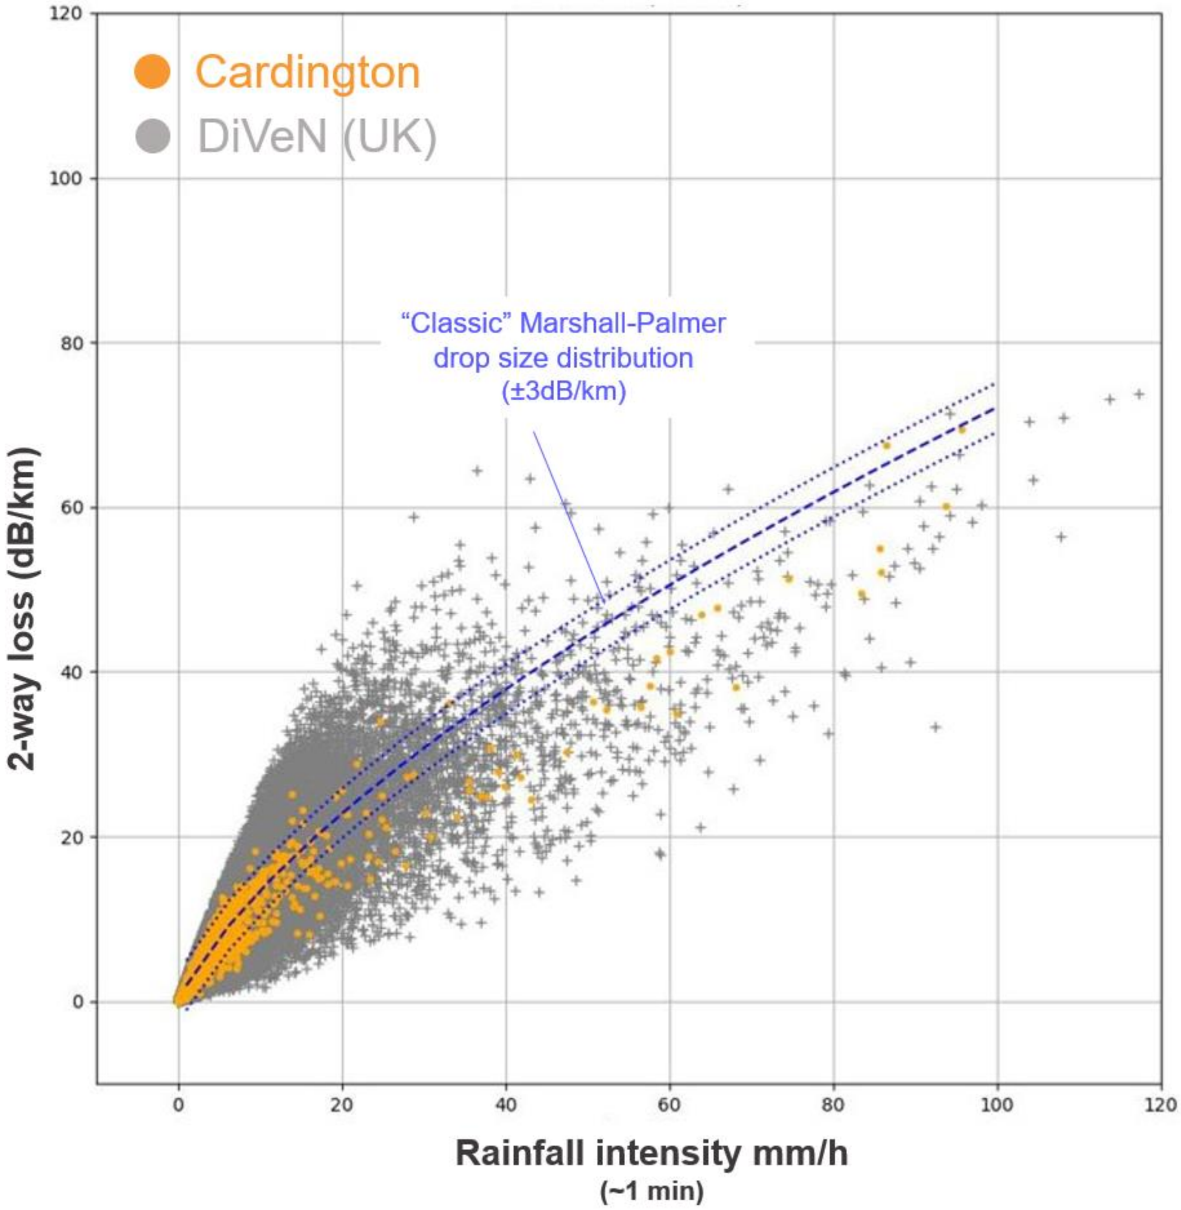

Figure 1: Rainfall intensity and 2-way path loss values calculated[14] for DSDs measured in the UK. Orange points were acquired at our Cardington testbed, the grey points were derived from the DiVeN[15] project, a network of 14 disdrometers deployed across at a range of locations and altitudes across the UK. Also shown on the plot is the relationship that would arise from the commonly assumed Marshall-Palmer DSD, with the dotted 3dB/km line showing where the path loss is a factor of 2 more or less than the central value.

---

[14] Note the only 'measurement' here is the DSD itself; the rainfall intensity and the path loss are calculated by integrating their theoretical size-dependent contributions across all droplet sizes and a spherical assumption for the droplet shape.

[15] Pickering, B. S., Neely III, R. R., and Harrison, D.: The Disdrometer Verification Network (DiVeN): a UK network of laser precipitation instruments, Atmos. Meas. Tech., 12, 5845–5861, https://doi.org/10.5194/amt-12-5845-2019, 2019.

We also note the following:

- Our working hypothesis for the absence of orange (‘Cardington’) points above the M-P line, based on the findings of SAF1[16], is due the absence of orographic (mountain) rain, which cannot occur at Cardington.
- It should be noted that an equivalent plot will apply to the more complex case of lidar path loss and there will also be some degree of correlation *between* lidar and radar responses.

A simplified set of ODD parameters inescapably leads to increased uncertainty associated with the use of those ODD parameters as indicators of the safe operating envelope.

The degree of simplification then determines the role of the components of the test ecosystem, with increased simplification moving the emphasis <u>away</u> from precise measurement and <u>towards</u> characterising the uncertainty space.

---

[16] SAF proof of concept report available at: https://www.npl.co.uk/getattachment/national-challenges/digital/Assuring-Autonomous-systems/Proof-of-concept-report.pdf

Embracing complexity

The previous section made an implicit simplifying assumption; that the dominating source of uncertainty arose from having insufficient ODD parameters to fully explain the atmospheric losses of the CAV sensor (and therefore that the remainder of the sensor-target response could be predicted accurately using the current set of ODD parameters). Our measurements at Cardington indicate that this is *not* the case. There can be a strong dependency of sensor response to very small details in, for example, the viewing geometry of the sensor and target. This may then be further exacerbated by an additional complex interaction with environmental factors such as condensation of water on the target and sensor. This is demonstrated on a notable weather day at Cardington, shown in Figure 2, where an early morning torrential downpour was followed by a period of more moderate rain and finally a period of overnight fog.

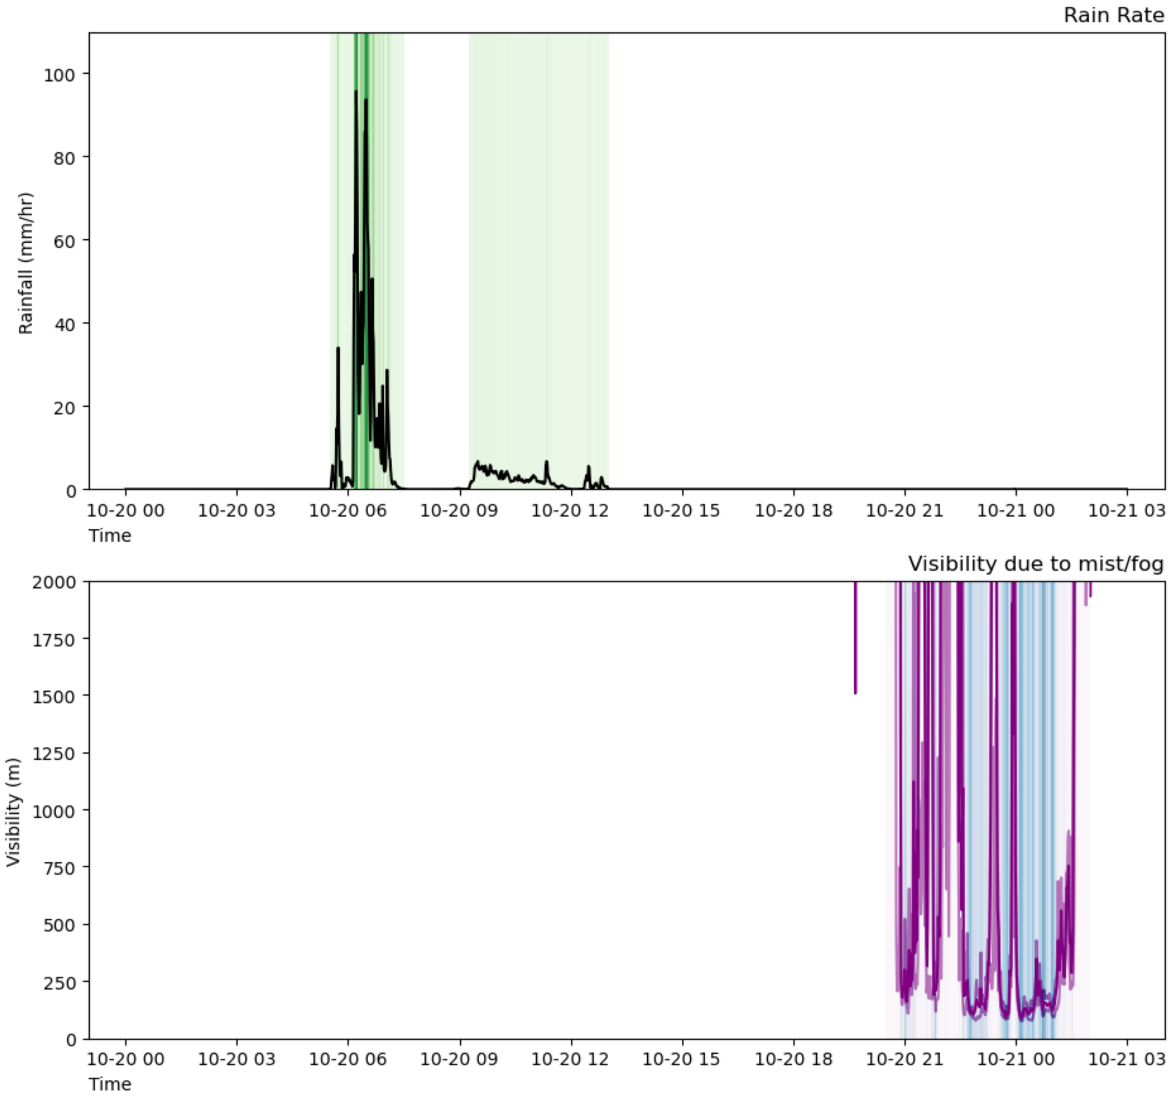

Figure 2: Weather data from Cardington on the 20th October 2022. The upper plot shows the representative 1-minute average ***rainfall intensity****. The green shading in the background of the plot, which darkens with increasing rainfall intensity, provides the weather context for the CAV sensor responses in Figure 3. The lower plot shows the* ***meteorological visibility*** *during the overnight fog event. The paler purple lines give an indication of the variability across the testbed and the blue background shading darkens with increasing fog water content to allow comparison with the sensor response.*

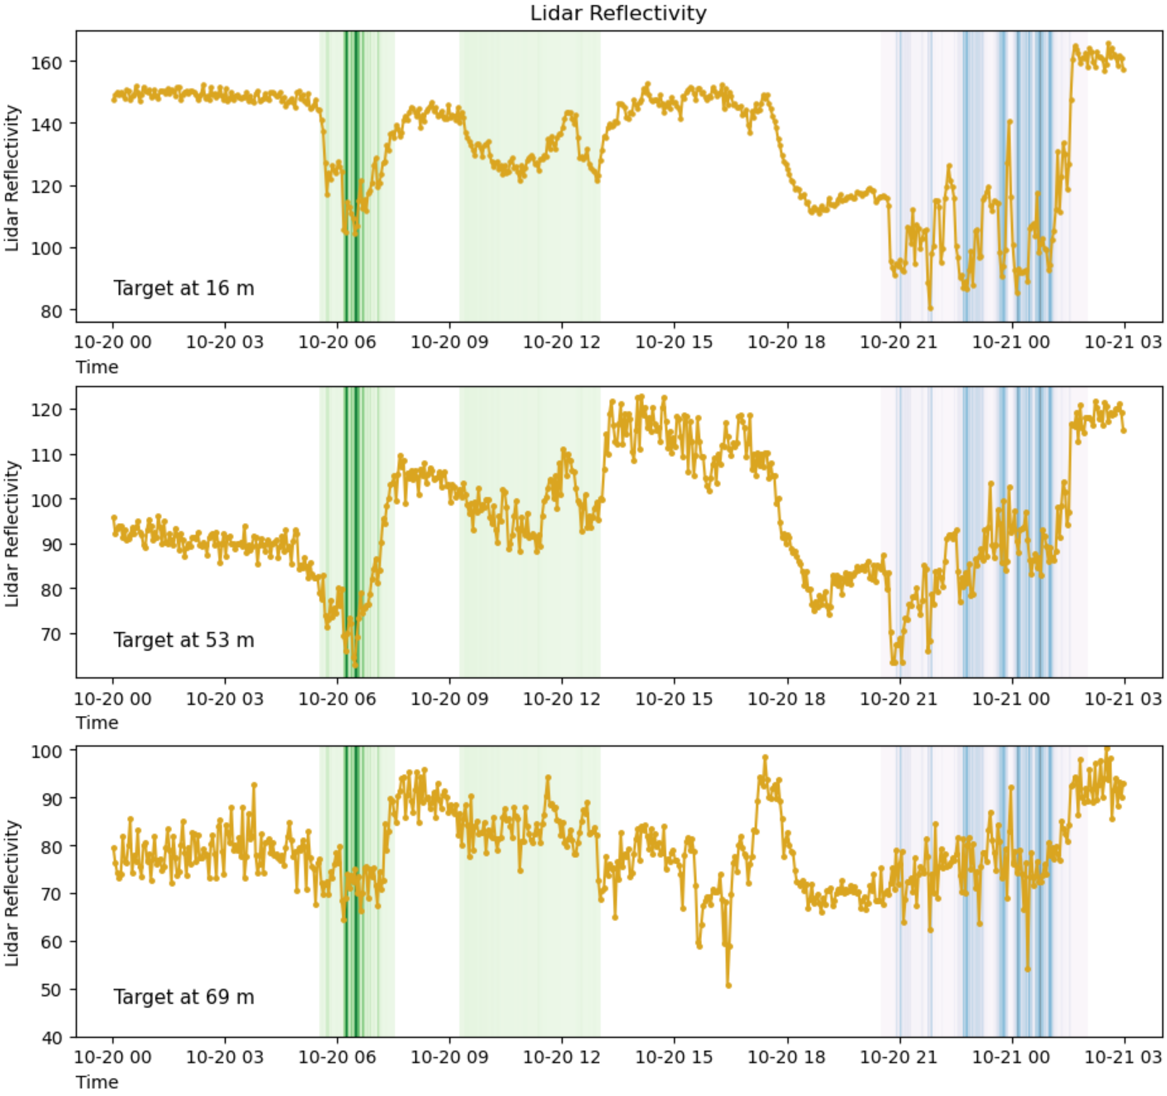

Figure 3: The effective lidar reflectivity (in arbitrary units) from three identical targets at (from top to bottom) 16 m, 53 m and 69 m range with time (expressed as data time e.g. 10-20 03 is 03:00 20th Oct). Note it is the variation that is of interest here rather than the absolute values, which are not directly comparable between the targets. While the meteorology experienced by the targets was to all intents and purposes "the same", the difference between *the traces is striking. The drop in reflectivity at about 17:00 is likely to be due to be due to condensation forming on either or both of the sensor and target surfaces as they cooled during the evening.*

Figure 3 show that it is possible to simultaneously observe positive, null *and* negative impacts on the lidar return from similar targets, when impacted by ostensibly the same weather. This both supports the importance of further detailed field measurements to attempt to understand the degradation mechanisms as comprehensively as possible and should prompt the community to collectively consider our aspirations for explicit VTE modelling of sensor response to the environment.

The real-world sensor-environment-target interaction can be very complex. This may pose practical limitations on our ability to fully model it within a VTE and to uniquely ascribe sensor performance level to even a comprehensively-described environmental ODD. This further suggests an increased role of the test ecosystem towards characterising the uncertainty space.

Harnessing uncertainty

When considering uncertainty, the following statements are true:

- there are many different sources of uncertainty ranging from experimental error to, as previously discussed, the description of the ODD itself and also real-world complexity;
- no single measure of uncertainty suits all use cases;
- larger sources of uncertainty can quickly dominate.

It is our observation that when uncertainty *is* mentioned in discussions relating to ODD transgression, the focus is predominantly on the uncertainty that a ***precisely known*** ODD threshold will be crossed. Much less attention is given to the ***ODD threshold itself being uncertain.*** They are equally important[17] and it is imperative to consider the two alongside each other:

When assessing the degree of confidence that a CAV (or CAV subsystem) is within its ODD, the uncertainty calculation must explicitly include the contribution of both the uncertainty in the ODD parameter and that of the ODD boundary itself.

This is illustrated in Annex A.

The benefits of this wider consideration of uncertainty extend well beyond greater confidence that the vehicle behaviour is understood close to the edge of the ODD boundary. It provides the basis for a framework that will:

- provide greater understanding into the ***role of each test environment*** in V&V;
- facilitate the ***combination of evidence*** from different types of testbed;
- ***optimise the design***[18] of the different elements of the sensor test ecosystem;
- characterise the ***remaining uncertainty***, the consequences of ***which can be assessed in the VTE;***
- ultimately, enable ADSs (*et al.*) to manage implicitly uncertain components of the ODD more intelligently and, thereby, facilitate a larger functional ODD (i.e. allow the CAV to operate safely in a wider range of conditions, and thus, for more of the time)

In summary:

The comprehensive treatment of uncertainty, including that of the ODD boundaries, does not only contribute to the net confidence in the V&V. It can be proactively used to shape the design of the overall test ecosystem.

---

[17] A possible analogy is as follows. Consider driving a car on a long journey that includes many different speed limits on the way. Staying within the legal speed limit involves both the accurate speedometer and clear road signage. During part of the journey, there may be insufficient road signage and therefore the speed limit may be unclear. The dominant source of error in addressing the question “Am I driving legally?” is not the speedometer, it is the knowledge of legal speed limit. Given the choice of investment in a better speedometer or a few more road signs, the choice would be the signage.

[18] …and thereby prevent unnecessary over-investment

The role of different physical testbeds in characterising and/or reducing uncertainty

Minimising the uncertainty in our knowledge of the response of the sensors[19] to weather lies very much at the heart of the SAF project. Table 1 summarises the various dimensions, or components, of the sensor-weather interaction that must be investigated.

Dimension	Description
ODD parameters	Optimising the use of conventional meteorological parameters by ensuring they are adequately described and understood by all users (i.e. correct interpretation of “rain rate”, “visibility”, etc)
ODD limitations	Understanding how the adoption of conventional environmental ODD parameters only may limit the extent to which they are good *predictors* of sensor response (which is critical in VTEs)
All weather types	Coverage of all relevant weather such as snow, hail, hazes, different rain types, dust, solar geometry… even space weather
Baseline weather-sensor response	Highest quality measurement of sensor response to weather using *reference* systems and targets
Degradation pathways	e.g. Exploring the relative importance of the action of weather on the sensor/target *surfaces* compared to the response to the *propagation* environment
Behaviours (e.g. vibration/motion)	Characterising variability of sensor response due to non-meteorological factors (and which may combine with weather effects in a non-linear manner)
Scenery (e.g. road)	Characterising the variability of sensor response due to surface & real-world target types (and which may themselves vary with the weather)

Table 1: The 'principal components' of the weather-sensor interaction. This is not considered to be comprehensive at this point.

In an ideal world, unconstrained by cost and time, all the components would be addressed by a single real-world testbed. However, there are many practical reasons why this will not be the case. For example, fully understanding all *degradation pathways* (which might require much manual intervention, such as sensor or target cleaning), while capturing *all weather types* (which might involve working at remote locations) and simultaneously understanding *ODD limitations* (which might require more detailed weather measurements using specialist, high-maintenance instrumentation) may not be realisable at the same time. The task may have to be split across multiple testbeds, each with more limited scope (but sharing some dimensions in common), which will necessarily reduce the amount of information available regarding the correlations between these dimensions.

---

[19] And, as importantly, characterising the remaining uncertainty

Understanding the magnitude of the uncertainties associated with these components can help us to prioritise the role and therefore required capabilities of each testbed in the ecosystem. It is also an essential exercise so that we can understand how a given set of test measurements from a testbed contributes to increasing confidence in quantifying the ODD for the sensor.

Our approach to distributing the measurement goals across the SAF testbeds is briefly described in Annex B. It is important to note our testbeds are *research* capabilities and therefore do not operate under the very tight operational timescales that will be required to test, for example, a newly-developed sensor where speed to market is a factor. It is very likely that the industry ecosystem will involve a number of reference testbeds that employ well-characterised reference sensors that are able to span a wide range of weather conditions over long periods of time in combination with, for example, weather emulation testbeds carrying similar reference sensors that can be linked back to these reference testbeds.

The multi-dimensionality of the weather-sensor response means that it is very unlikely that a single testbed will be able to adequately characterise a sensor's response to the relevant ODD parameter.

We expect that measurements of sensor measurements will have to be combined with information gathered at other testbeds using reference sensors.

VTEs as a key tool to quantify the impact of uncertainties

The previous sections argue that, regardless of the quality of the testbeds, there will remain inescapable but consequential uncertainties around all elements of sensor assurance as it relates to weather. The existence of these uncertainties implies that there is likely to be a point where there is no additional return from investing further in the complexity and capability of the physical testbeds.

The impact of the resulting residual uncertainty on the V&V process must be quantified and this is where simulation offers great potential.

We propose the following statements of guidance, which will be further explored in the SAF project.

VTE approaches must be able to generate quantitative assessments of confidence arising from uncertainties in the input parameters, including the uncertainty in the performance of sensors and sensor combinations close to the ODD edge.

and

When characterising uncertainty associated with data from physical sensor-weather testbeds, the resulting descriptions must be suitable to generate an ensemble of input scenarios into a VTE such that it is possible to produce meaningful statistics to inform V&V decisions.

Concluding remarks

This paper discusses the challenge of quantifying the environmental ODD for CAV sensors. It explores the complexity of the weather-sensor interaction and the implications that this has on the confidence with which we can determine the ODD boundary. A major theme running through the document is the need to characterize and then exploit uncertainty information in order to both optimise the design of the test ecosystem and then use the results it produces confidently.

In the section "The pre-eminence of the ODD", the point is that made that ODD parameters must be measurable (or at least can be estimated) when the CAV is in operation[20]. Confidence that a CAV sensor is safely within its ODD therefore depends on the certainty of the ***separation*** between the current value of the relevant weather parameter(s) and the ODD boundary (see Annex A). This is just as much a function of our ability to monitor current weather conditions as it is to determine the ODD value in the test ecosystem – the largest error source dominates. This must be a key consideration in the test ecosystem design and how the environmental ODD is used when the CAV is deployed.

Finally, we would like to emphasise that this paper is to promote discussion and challenge in order to further mutual understanding. We welcome your comments.

Acknowledgments

This discussion paper is an output from the joint Met Office and NPL Sensor Assurance Framework (SAF) project. In addition to the listed contributing specialists, we are indebted to the wider project teams of both organizations and to the stakeholders who engage with the project across UK government, regulatory bodies, standards organisations, academia and industry.

The project is funded and supported by the Centre for Connected and Autonomous Vehicles (CCAV). We are also grateful to the support of Innovate UK in sponsoring the initial design phase of the Cardington testbed.

---

[20] The 'live' values of the ODD parameters when the CAV is on the road is now widely referred to as the Current Operational Domain (COD)

Annex A: The importance of the uncertainty budget

The following figures give a visual representation of how uncertainty manifests itself when considering the environmental ODD boundary for a CAV sensor. The figure captions present the narrative.

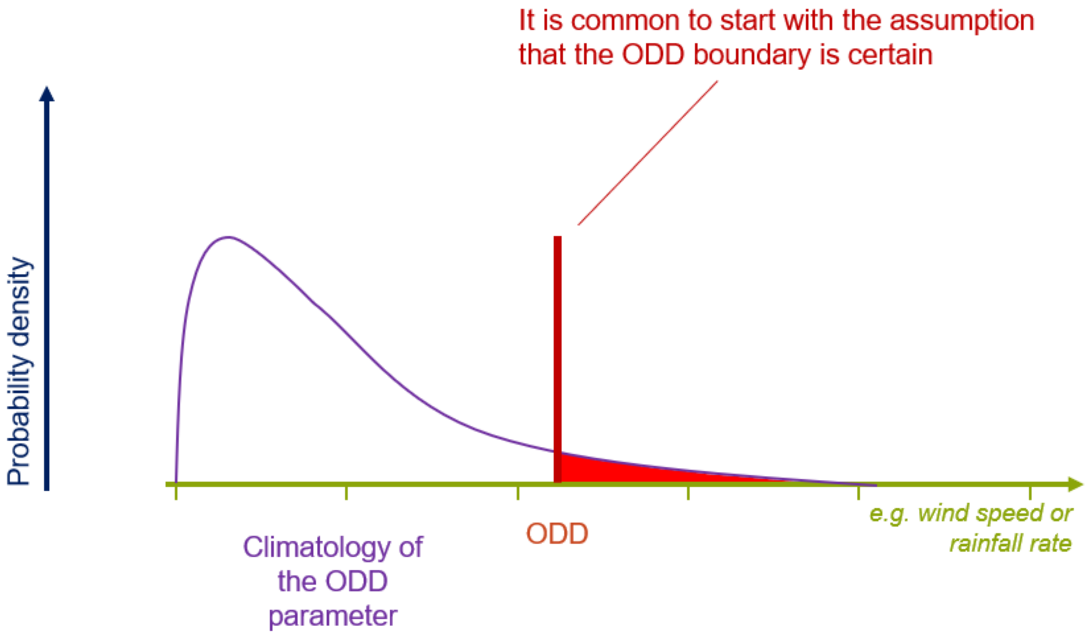

Figure 4: For a given location and a given environmental ODD parameter (e.g. rainfall rate or wind speed), there will be a statistical distribution (or climatology) of parameter values. This is often referred to as the Target Operational Domain (TOD) and is represented here a simplified probability density function. If the ODD boundary is taken as a certain value, the probability that the ODD boundary is exceeded at any given time is indicated by the red-shaded area.

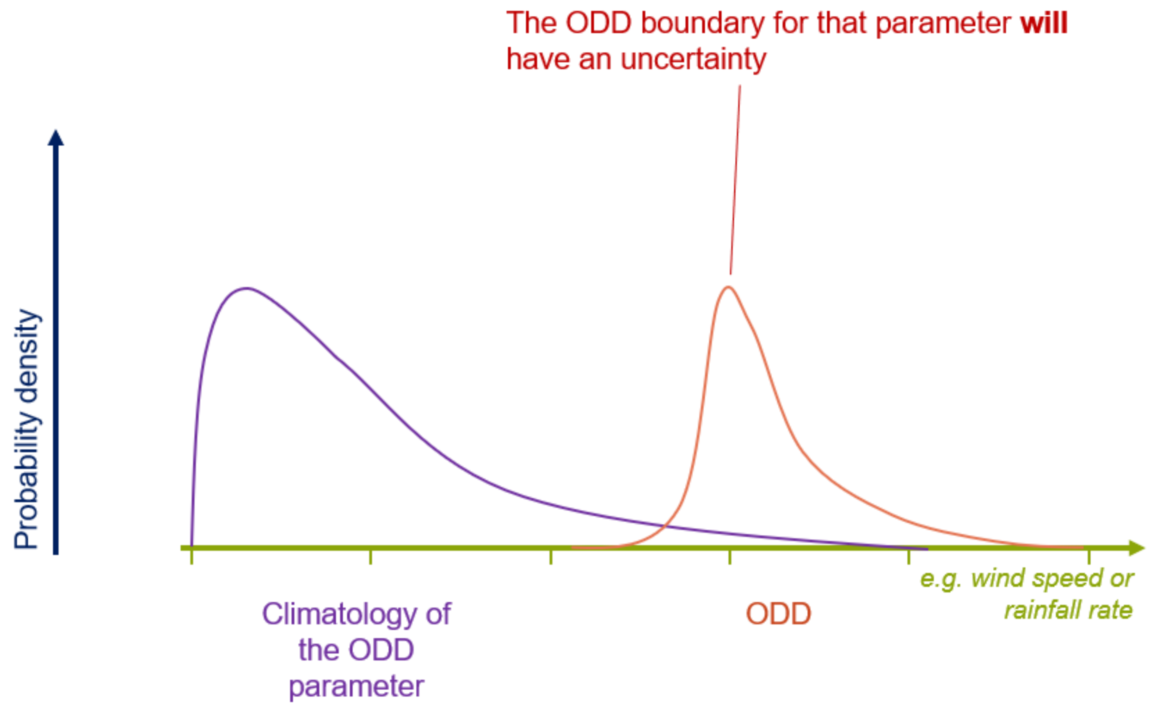

The ODD boundary for that parameter **will** have an uncertainty

Figure 5: For the reasons stated in the paper, the ODD boundary is also uncertain and this uncertainty is represented as a range of possible values, again in the form of a probability density. There is an overlap between the curves. Note we have not attempted to ensure the areas under our curves are consistent!

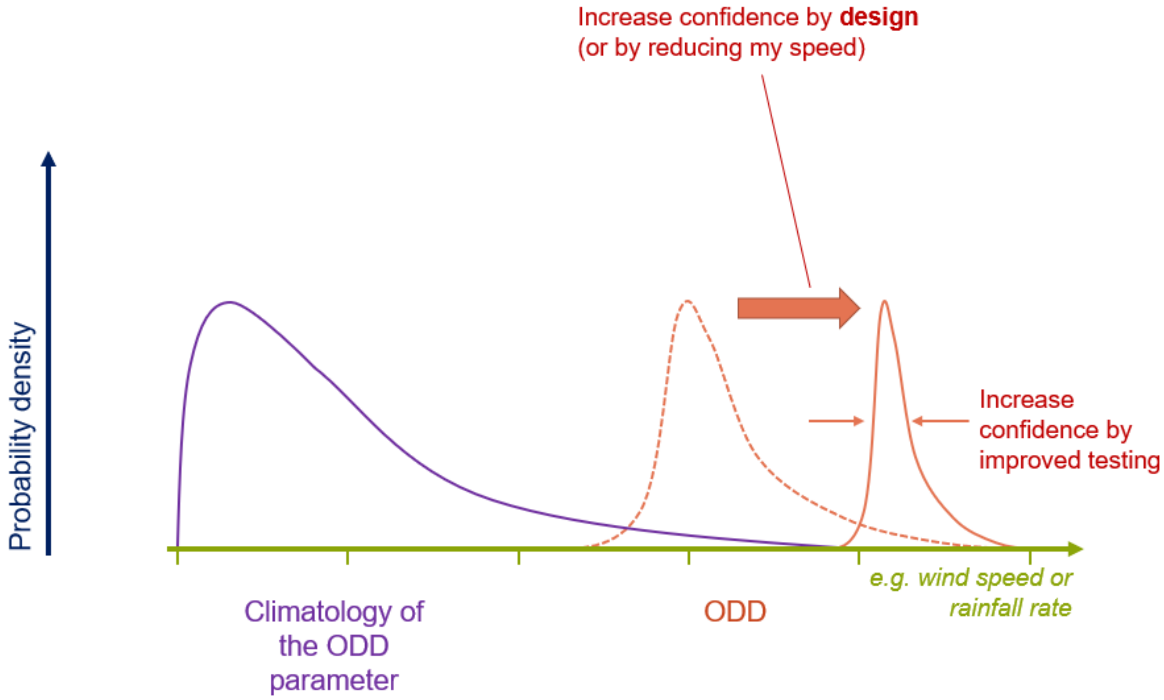

Increase confidence by **design** (or by reducing my speed)

Increase confidence by improved testing

Figure 6: When the uncertainty in the ODD boundary itself is accepted, the means to manage this become clearer. Greater confidence in remaining within the nominal operating window for the sensor can be achieved by improved design of the sensor ***or*** *by addressing the main sources of uncertainty in the characterisation of the ODD boundary.*

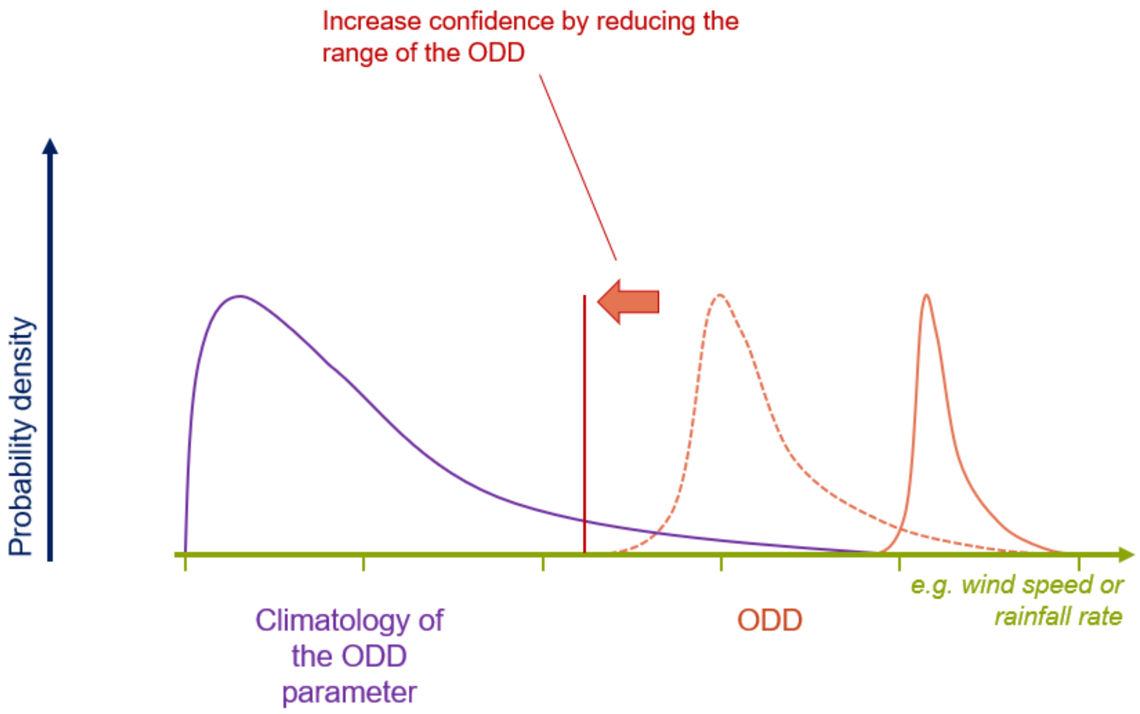

*Figure 7: The uncertainty may also be managed by reducing the ODD threshold to a negligible risk level. This still requires that mitigation is required (e.g. performing a minimum risk manoeuvre, MRM) if the ODD is then breached. This in turn requires that the vehicle is able to recognise that this has happened, either by knowing the current operating domain values (**explicit** COD awareness) or by recognising its sensors are compromised (**implicit** awareness).*

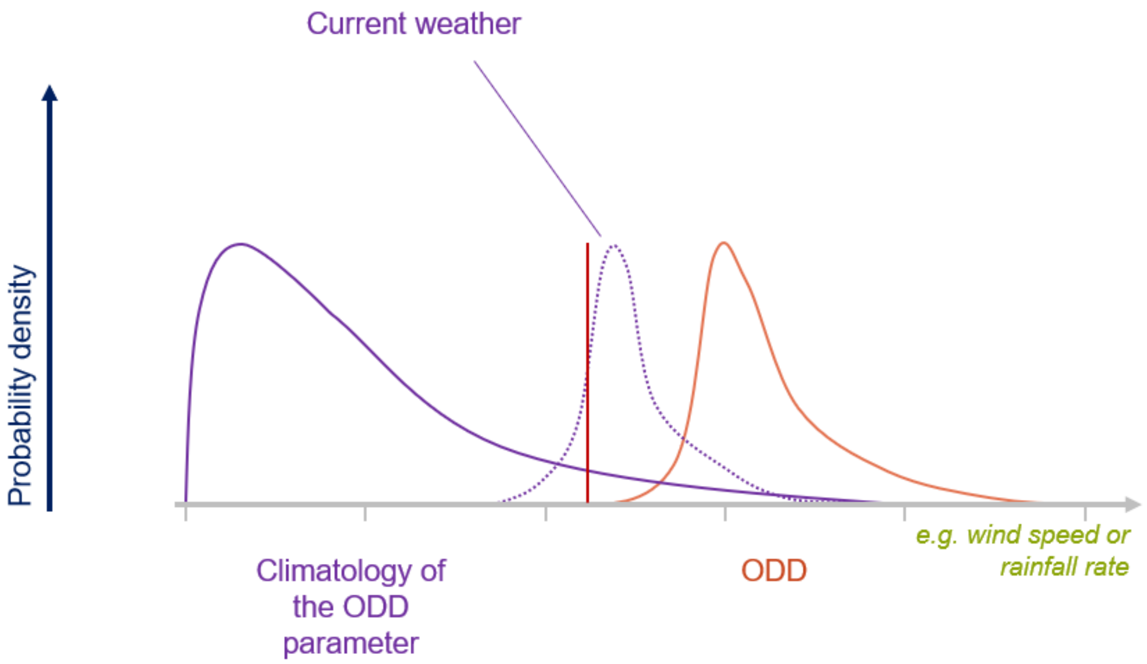

Figure 8: Explicit COD awareness (ODD measurability) also carries uncertainty because there is a limit to which we can ever know the current operational environment

Annex B: SAF testbeds

This Annex illustrates the role of different testbeds using existing and planned testbeds within the SAF project.

The three different testbeds are:

1. The Cardington testbed
2. A new testbed which is under construction at NPL HQ, Teddington
3. A potential mobile testbed designed to experience a wider range of climatic conditions

The plots do not include “emulated” weather facilities nor on-the-road trials, which we also believe to be very important in the test ecosystem. Please note that the plots, presented as spidergrams, should be considered to be illustrative only. The relative importance of any one dimension in the confidence of the V&V process will ultimately be determined by their contribution to the reduction of uncertainty.

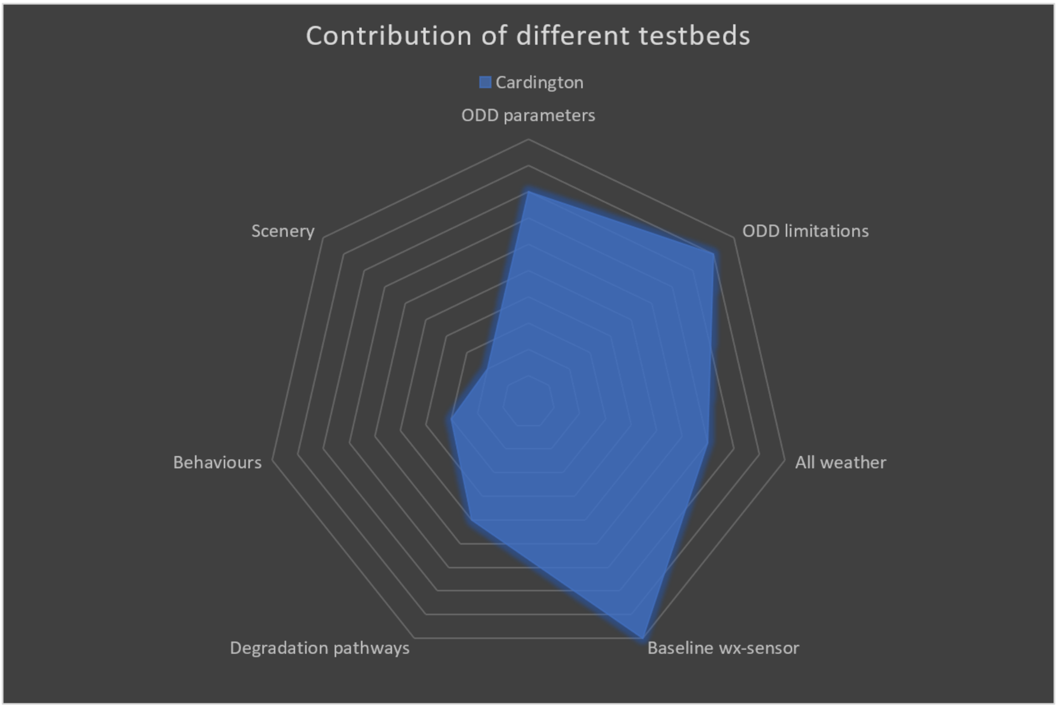

Figure 9: The Cardington testbed and its coverage of the dimensions of testing listed in Table 1. Particular strengths of the testbed are understanding the response of ***reference*** *sensors to the weather and the exploration limitations caused by ODD simplification, the latter being possible because of the large number of meteorological sensors dedicated to measuring droplet size, precipitation type, spatial variability and the illumination environment. With only a grass surface at the testbed, however, we cannot explore the effect of road surface wetting. Also, the rural nature of the site makes the manually intensive exploration of degradation mechanisms more challenging. Cardington’s location on flat, low-lying ground means that we cannot explore the impact of orographic mountain rain and are unlikely to see significant snow.*

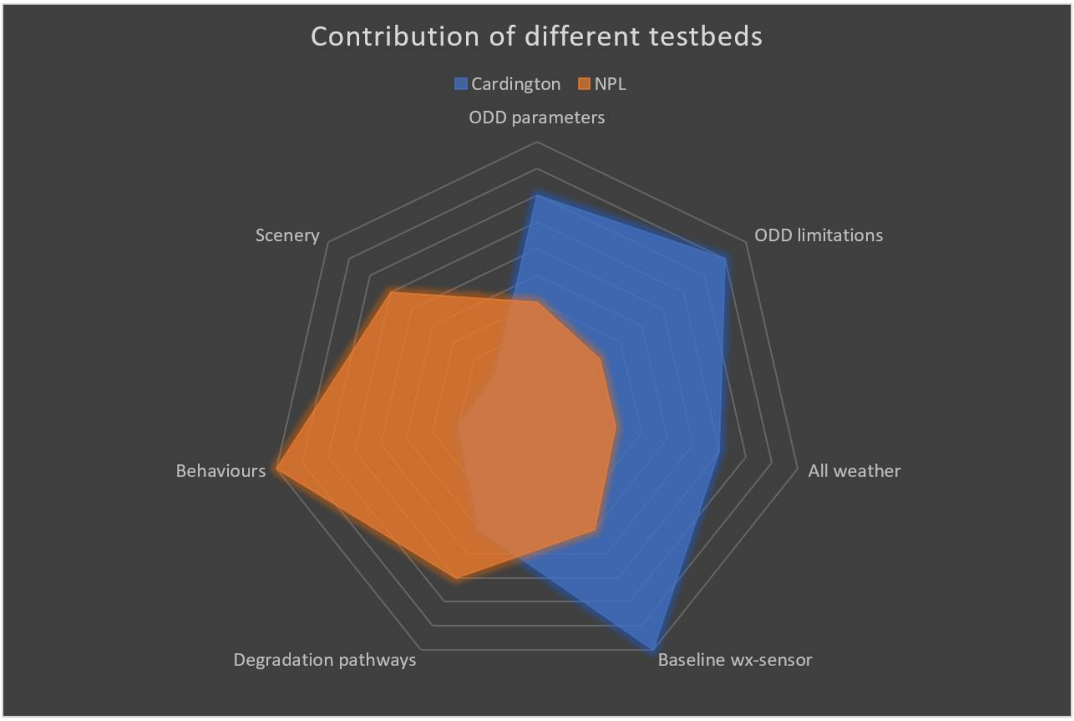

Figure 10: The new SAF testbed at NPL (in orange) has access to a road surface and its accessibility to nearby lab space will enable testing such effects as moving targets and vibration (behaviours) and degradation pathways. However, it is less likely to experience e.g. fog due to its location.

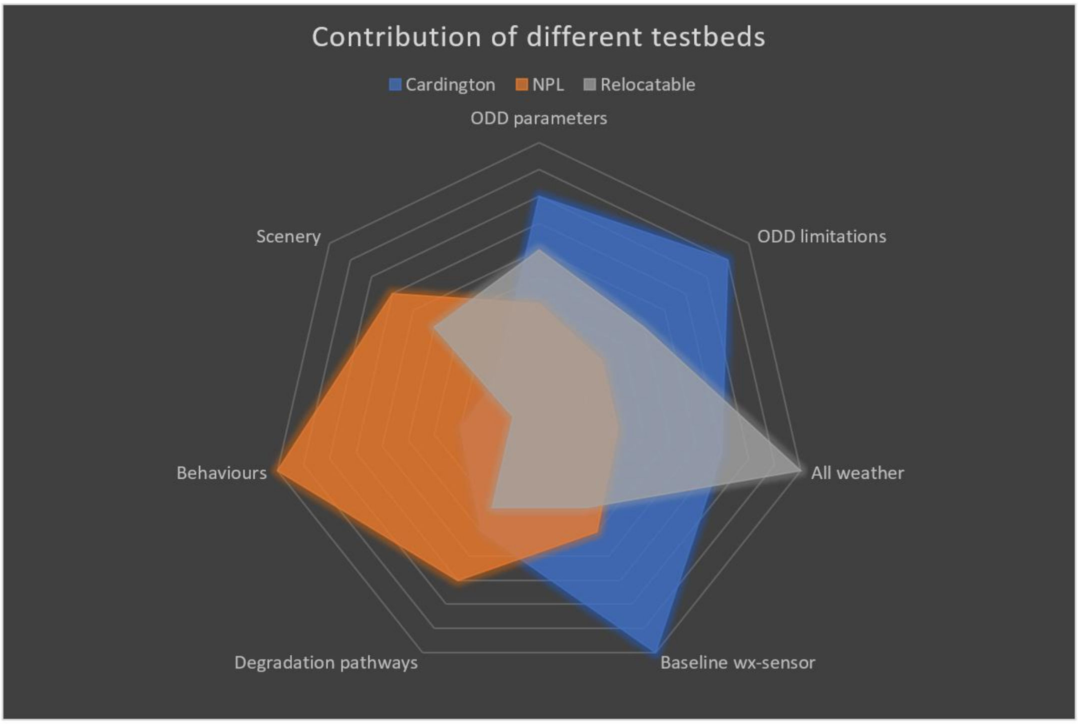

Figure 11: A simplified relocatable facility may provide access to a greater variety of weather conditions. The ability to deploy may require both a simplified specification of both meteorological and CAV sensor capabilities.

Contributors and correspondence

Contact for correspondence	• Dr Dave Jones, Principal Scientific Consultant, Observations R&D, Met Office dave.jones@metoffice.gov.uk
Contributing specialists	• Dr Dave Jones, Principal Scientific Consultant, Observations R&D, Met Office • Dr Richard Dudley, Science Area Lead and Principal Scientist, Electromagnetic technologies Group, National Physical Laboratory • Dr Nawal Husnoo, Senior Scientist, Met Office • Dr Henry Odbert, Science Manager, Met Office • Matthew Fry, Scientist, Met Office • Grace Richardson, Industrial Placement, Met Office • Louise Wright, Head of Science for Data Science, National Physical Laboratory • Andre Burgess, MBA, Security and Resilience Group, Partnerships, National Physical Laboratory • Michael Sczepanski, Business Analyst - Product Futures, Met Office • Dr Stephan Havemann, Senior Scientist, Met Office • Dr Jeremy Price, Scientific Manager, Met Office • Dr Imran Mohamed, Higher Scientist, Electromagnetic Technologies Group, National Physical Laboratory • Dr Maurizio Bevilacqua, Senior Engineer, Data Science Dept, National Physical Laboratory • Joel Rapley, Scientist, Electromagnetic, National Physical Laboratory • Dr Fengping Li, Higher Scientist, Electromagnetic Technologies Group, National Physical Laboratory • Hsun Yang, Higher Scientist, Electromagnetic Technologies Group, National Physical Laboratory.