

PAPER • OPEN ACCESS

A metrological framework for uncertainty evaluation in machine learning classification models

To cite this article: Samuel Bilson *et al* 2025 *Metrologia* **62** 064001

View the [article online](#) for updates and enhancements.

You may also like

- [Statistical models and computation to evaluate measurement uncertainty](#)
Antonio Possolo
- [Informative Bayesian Type A uncertainty evaluation, especially applicable to a small number of observations](#)
M Cox and K Shirono
- [Probabilistic measurement modelling may overcome the opposition between the Bayesian and the frequentistic views](#)
Giovanni Battista Rossi and Francesco Crenna

A metrological framework for uncertainty evaluation in machine learning classification models

Samuel Bilson^{1,*} , Maurice Cox¹ , Anna Pustogvar^{2,3}  and Andrew Thompson^{1,*} 

¹ Department of Data Science, National Physical Laboratory, Teddington, United Kingdom

² Department of Thermal & Radiometric Metrology, National Physical Laboratory, Teddington, United Kingdom

³ School of Geography, Geology & the Environment, University of Leicester, Leicester, United Kingdom

E-mail: sam.bilson@npl.co.uk and andrew.thompson@npl.co.uk

Received 28 March 2025, revised 3 October 2025

Accepted for publication 5 November 2025

Published 2 December 2025



CrossMark

Abstract

Machine learning (ML) classification models are increasingly being used in a wide range of applications where it is important that predictions are accompanied by uncertainties, including in climate and earth observation, medical diagnosis and bioaerosol monitoring. The output of an ML classification model is a type of categorical variable known as a nominal property in the International Vocabulary of Metrology (VIM). However, concepts related to uncertainty evaluation for nominal properties are not defined in the VIM, nor is such evaluation addressed by the guide to the expression of uncertainty in measurement (GUM). In this paper we propose a metrological conceptual uncertainty evaluation framework for nominal properties. This framework is based on probability mass functions and summary statistics thereof, and it is applicable to ML classification. We also illustrate its use in the context of two applications that exemplify the issues and have significant societal impact, namely, climate and earth observation and medical diagnosis. Our framework would enable an extension of the GUM to uncertainty for nominal properties, which would make both applicable to ML classification models.

Keywords: classification, machine learning, uncertainty evaluation, nominal properties

1. Introduction

1.1. Machine learning (ML) classification from a metrology perspective

The use of ML models offers the potential to bring great benefits to a wide range of societal challenges, for example in

climate and earth observation [1–5], health care [6, 7], and advanced manufacturing [8]. ML models can be used to make sense of large data sets and learn complex relationships within them, and also to automate manual processes that require considerable time or expertise.

Many of these applications involve the measurement of some signal on an instrument and inferring information of value to society from that signal. Traditionally, such problems would have involved developing a physical or empirical model that would relate the measurements to the information of interest. Where the information of interest is a quantity or set of quantities, the empirical model developed would form what is considered in metrology a *measurement model*, providing the relationship between input quantities and a *measurand*, the ‘quantity intended to be measured’ [9].

* Authors to whom any correspondence should be addressed.



Original Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

ML methods, and in particular methods of *supervised learning*, are increasingly being used to establish (define and parametrise) measurement models. In supervised learning, data-driven models are built to infer predictions of the value of some (output) variable given knowledge of the values of other (input) variables. Typically a model *type* is first selected, and then its parameters are optimised. An example of a type of model used in supervised learning is a neural network with a predefined architecture (see section 4.2). The phrase ‘to learn the model’ is commonly used in ML to refer to the optimisation of the model parameters.

The distinctive feature of supervised learning models is that their parameters are learned from training data, rather than purely using physical intuition or laws. Two types of supervised learning are often distinguished: *regression* and *classification*. The distinction is based on the type of variable that is being predicted: a regression model predicts a quantitative variable, while a classification model predicts a categorical variable. In section 2 we define the terms *quantitative* and *categorical* more precisely and provide examples of each.

Both regression and classification models have wide application, but the focus of this paper is on classification models. By a *classification model* we mean a function f , which takes some number of input variables and returns a categorical output variable. Formally speaking, we can denote a *collection* of input variables by X . A collection is like a vector but can include both quantitative and categorical variables. The output categorical variable is denoted by Y , and can take one of K values (often referred to in ML as classes) denoted by $\{c_1, \dots, c_K\}$. Following the convention of the guide to the expression of uncertainty in measurement (GUM) [10], we denote a value taken by any of these variables using lower case. Thus x is an instance of X and y is an instance of Y and these are related through

$$y = f(x), \quad y \in \mathcal{C}_K := \{c_1, \dots, c_K\}. \quad (1.1)$$

In reality, most classification models output a particular type of categorical variable known as a *nominal* variable (sometimes referred to as an unordered categorical variable), or a *nominal property* in the language of the International Vocabulary of Metrology (VIM) [9], in which the categories have no natural ordering. See section 2 for a more detailed discussion of variable types and how they are categorised in the VIM.

More specifically, an *ML classification model* is one that has been learned from labelled training data, by which we mean N samples (or observations) of paired data (x_i, y_i) , $i = 1, \dots, N$, consisting of observations of both the collection of input variables X and the output variable Y . The supervised learning process consists of two stages: a training phase and a prediction phase. In the training phase, the model f is learned from known paired inputs and outputs, which typically involves optimising the parameters of the model. In the prediction phase, the model is then evaluated upon new input variables. ML classification models, which are learned from training data, can be distinguished from *rule-based* classifiers

which use predetermined classification rules which are not dependent upon data.

1.2. Uncertainty evaluation for ML classification

It is always important that an ML classification model is deployed in a trustworthy and quality-assured way. More specifically, ML classification models are currently being used in several applications in which it is crucial that predictions are accompanied by an evaluation of uncertainty. Examples include earth observation such as land cover (LC) classification, mapping of water, ice extent and cloud masking [1–3], environmental monitoring such as bioaerosol classification [4, 5], and medical diagnosis such as the detection of heart arrhythmias from electrocardiogram (ECG) and photoplethysmography (PPG) signals [6, 7]. See section 4 where we illustrate the conceptual contributions of this paper in the context of two case studies: LC classification and atrial fibrillation (AF) detection. The importance of uncertainty evaluation for ML was highlighted in the Strategic Research Agenda developed by the European Metrology Network for Mathematics and Statistics (MATHMET) [11, section 3].

In a metrology context, the outputs of an ML model may also be used to make inferences about other variables, that is, they can be used as input quantities for further measurement models or classification models. It follows that, as well as being important as an end in itself, uncertainty evaluation also enables the outputs of ML models to be used as part of an ongoing traceability chain, where the term *traceability chain* is used to include both traditional metrological traceability chains, where artefacts or reference materials are passed between facilities, and situations where data are passed through multiple processing steps (e.g. [12]). In section 3.3, we provide an example of how the outputs of an ML classification model, along with their accompanying uncertainties, can be included as inputs to other models.

1.2.1. The uncertainty evaluation task for ML classification models.

For classification models, a common way to express the confidence of a prediction is to assign a probability to each of the K possible classes rather than simply returning a single predicted class [9, definition 2.9]. Assigning a probability to each class of a nominal property is equivalent to obtaining the probability mass function (PMF) for the classification. In section 3 we develop the observation that the role of the PMF here is the natural counterpart of the role of the probability density function (PDF) in the measurement of a continuous quantitative variable (continuous quantity). We also show in section 3 how the PMF of a categorical variable can be used as an input into other models, allowing the propagation of uncertainty through multistage measurement models involving nominal properties.

Section 1.1 described how ML classification is a type of supervised learning where a model is learned that relates outputs to inputs, and where an implementation of ML

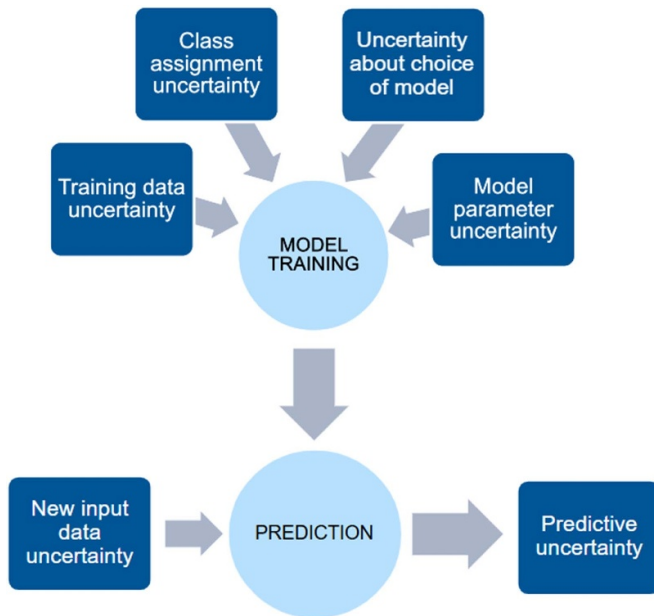


Figure 1. Factors contributing to uncertainty in ML prediction.

classification involves a training phase followed by a prediction phase. In ML classification models, contributions to the overall uncertainty of a prediction are introduced in both the training and prediction phases, which is in contrast to rule-based classifiers in which there is no training phase and where the classification rule is predetermined, for example tolerance thresholds in conformity assessment (see also section 1.2.2).

A crucial assumption in supervised learning is that the model learned on training data is applicable to input data for which predictions are desired. One way to statistically characterise this assumption is to say that the input data (for both training and prediction) are independent and identically distributed (i.i.d.). In practice, this assumption must be validated through appropriate experimentation. For example the i.i.d. assumption can be validated using a suitable hypothesis test. Where the assumption is violated, new input data for which model predictions are required is said to have undergone *distribution shift*; see for example [13, chapter 8]. Uncertainty evaluation methods for supervised learning predictions typically assume the absence of distribution shift, and such methods are only reliable to the extent that this assumption is not violated.

We next outline some of the factors contributing to prediction uncertainty throughout the ML pipeline. These contributory factors are also illustrated in figure 1.

1. **Training data uncertainty.** In the training phase, one contribution to prediction uncertainty is uncertainty in the training data itself. For example, when classifying LC based upon surface reflectance measurements (see our case study in section 4.1), there is uncertainty in the surface reflectance measurements due to imperfect knowledge of aerosol

optical depth (AOD) (input variables), and the LC data is also subject to labelling errors (output data). These training data uncertainties affect the training of the model.

2. **Class assignment uncertainty.** In classification tasks for which ML might be used, there is typically no theoretical justification for the choice of input variables, and therefore no reason to assume that observations of the input variables are sufficient to determine a classification unambiguously. Such tasks are in this sense ill-posed: a given observation for the input variables could occur for data drawn from more than one class. In such cases, there is a limitation, inherent in the modelling task, to determining the class to which a data sample belongs. This limitation is itself dependent upon the choice of input variables. *Feature selection* refers to the common practice in ML of choosing a subset of available input variables, which can help to remove redundant/unnecessary input variables. However, feature selection cannot eliminate class assignment uncertainty if the original selection of variables is insufficient.
3. **Uncertainty about choice of model.** The supervised learning approach relies upon the selection of a certain type of classification model (see section 1.1). The model type is typically selected using empirical testing, and this process introduces uncertainty. Bayesian model averaging is one method that has been proposed to capture uncertainty due to choice of model [14].
4. **Model parameter uncertainty.** ML models need to be optimised using the training data, for example, optimising the values of its parameters. There is, therefore, also uncertainty related to the estimates of the ML model parameters. Some of this uncertainty may be due to the challenges of optimising a model with multiple local optima. The ability to optimise an ML model also depends upon having training data that samples the region of data space that is of interest sufficiently densely. Training data requirements vary depending upon the complexity of the classification task and the chosen ML model.
5. **New input data uncertainty.** In the prediction phase, additional uncertainty is introduced through measurement uncertainties in previously unused (non-training) input data to which the model is being applied.

Uncertainty source (2) can be understood as inherent to the learning task. Uncertainty sources (1) and (5) constitute properties of the measuring system used to obtain the observational data. Uncertainty sources (3) and (4) are both a property of the modelling approach (though modelling approaches are often constrained by the task).

Many modern ML classifiers were originally developed with the aim of identifying the (single) correct class. However, the need to go beyond deterministic classifiers and evaluate uncertainty has been recognised in the ML community; see for example [15]. Many, though not all, ML classifiers are probabilistic, in the sense that they return class probabilities [16]. Methods differ according to which sources of uncertainty are

captured and whether a Bayesian or frequentist approach is adopted.

In section 4 we present two case studies based on societal applications that rely on a metrological analysis of probabilistic ML classifiers.

In our first case study (section 4.1), a Bayesian approach is taken using a quadratic discriminant analysis (QDA) model [17]. This method explicitly models four of the five sources of uncertainty described above and illustrated in figure 1. The method falls into the general category of *generative modelling* in which the input variables in each class are assumed to follow some statistical distribution, which in the case of QDA is a multivariate Gaussian distribution. Other examples of generative classifiers include linear discriminant analysis and Naive Bayes; see [16] for more details.

Many modern ML classifiers are instead examples of *discriminative modelling* in which classification is achieved without making statistical assumptions about the data. In our second case study (section 4.2), a discriminative modelling approach based on neural network models known as Monte Carlo (MC) dropout is used [18, 19]. This method captures uncertainty arising from both the data and the model, but it does not explicitly model the various sources of uncertainty in figure 1. The method does, however, have the advantage of being applicable to high-complexity neural networks, which are known to be especially able to capture complex non-linear relationships. Other examples of discriminative classifiers include logistic regression, support vector machines and tree-based methods such as random forests; see [20] for more details.

1.2.2. Related metrological frameworks. Evaluation of the uncertainty of quantitative outputs from measurement models is addressed by the GUM suite of documents [10, 21–23] and the related VIM vocabulary [9]. Central to the GUM is the concept of the propagation of measurement uncertainty through measurement models whose parameters have been determined. This scenario would correspond to a trained ML model in which estimates of the parameters are used. The GUM can also treat the parameters that define the measurement model as input quantities in their own right, propagating uncertainty (and where there are multiple such parameters, their covariance) as with other input quantities.

The GUM framework is explained for models obtained through regression of parametric models (empirical measurement models) in GUM-6 [24]. Uncertainty associated with model parameter estimates is addressed in [24, clause 7]. Uncertainty concerning the choice of model is discussed briefly in [24, clause 8.5]. The inclusion of effects due to measurement uncertainties in either the input or output variables is addressed in [24, clauses 9 and 10].

The frameworks developed in the GUM suite of documents can thus, in principle, be applied to any ML regression model. However, the GUM documents mentioned above do not address classification models since the concept of

measurement uncertainty evaluation has been developed in the metrology community (and in the GUM and VIM documents) for quantitative variables, or *quantities* in the language of the VIM and the GUM. The output of a classification model, on the other hand, is not a quantitative variable, but rather an unordered categorical variable, or a nominal property.

A certain kind of classification model is addressed in the fourth document of the GUM suite dedicated to conformity assessment [25]. In conformity assessment, the result of a measurement, which represents a quantity (e.g. mass concentration of mercury in a sample of industrial water), is converted into a nominal variable, which has two categories: *conforming* and *non-conforming*. The conversion is implemented by identifying whether the result of a measurement falls into an interval of permissible values set by one or two tolerance limits. Such conversion essentially represents a binary classification task. Binary classification performed by applying a threshold to a quantity (similar to conformity testing) has many applications. For example, in earth observation, operational cloud masking of satellite imagery has been historically performed by applying empirically derived thresholds to top-of-atmosphere reflectance measurements, e.g. see [26] (figure 2). However, it is important to note that with the recent progress of ML learning, and specifically deep learning, more efforts have been made in developing more advanced cloud masking algorithms (e.g. [27]).

Similarly to a probabilistic ML classifier, in conformity assessment the assigned category is provided with a probability value, the *conformance probability*. However, in conformity assessment the classifier is rule-based, there is no training phase, and therefore uncertainty related to a training phase does not need to be taken into account. For these reasons, the framework in [25] cannot be used for a comprehensive uncertainty evaluation of ML classifiers, which rely on training.

In a standard conformity assessment a threshold is fixed (e.g. a certain critical mass concentration of mercury based on which a water sample is considered to be conforming or non-conforming). In contrast, in a binary classification performed alike conformity assessment (e.g. cloud masking described above), a threshold is often empirically derived (e.g. based on expert knowledge of top-of-atmosphere reflectance of a cloud vs non-cloud). In the latter case, there would be an additional uncertainty source associated with the choice of the threshold, which will affect the classification results, and this uncertainty source is not considered in a standard conformity assessment [25].

1.3. Summary of contributions

- We propose a metrological conceptual uncertainty evaluation framework for nominal properties which is applicable to ML classification. Our framework is inspired by existing vocabularies and frameworks from the fields of qualitative chemical analysis [28] and clinical laboratory sciences [29]. It is based upon the determination of the PMF of the nominal variable returned by the ML classifier, and the expression

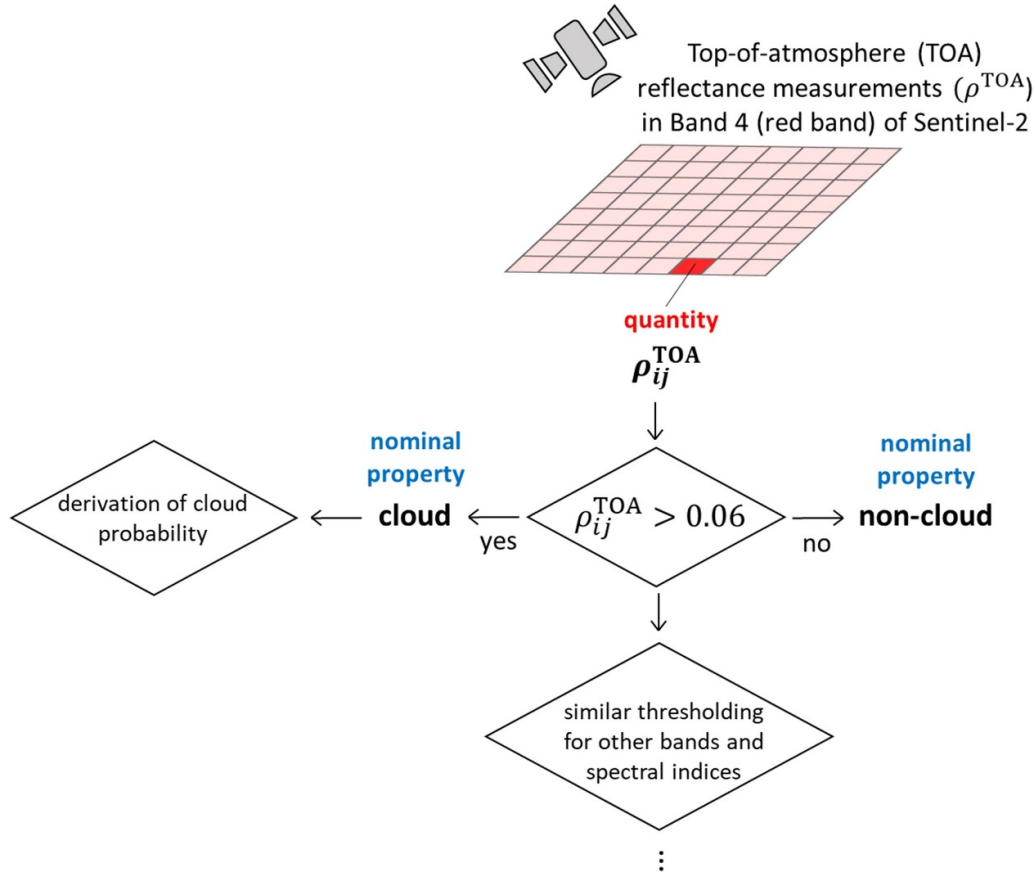


Figure 2. Cloud masking of Sentinel-2 satellite imagery.

of uncertainty in terms of summary statistics derived from it. We review some possible summary statistics, discussing their benefits and drawbacks. The framework represents a natural analogue to the determination of the PDF for a quantitative variable and the derivation from it of summary statistics such as standard uncertainty. We also show how conditional probability can be used to enable uncertainty evaluation for measurement models that have quantitative and categorical variables as inputs.

- We illustrate our conceptual framework by applying it in the context of two case studies where a metrological approach adds value. The first is the use of ML for classification of LC based upon surface reflectance measurements made by a spaceborne optical sensor. The second is the use of ML for detection of AF based upon PPG signals used in wearable devices.

1.4. Outline of contents

In section 2, we give background on nominal properties, and we describe how such properties are addressed in metrology. In section 3, we present our framework for uncertainty evaluation of nominal properties based on discrete probability distributions and summary statistics. We also address the propagation of uncertainty through multistage measurement

models. We present our two ML case studies illustrating the framework in section 4, before concluding in section 5.

2. Nominal variables

2.1. Nominal versus other variables

Classification tasks deal with nominal variables. To explain what nominal variables are, we will use Stevens’ theory of scales of measurement [30]. For a deeper review of Stevens’ theory and its relevance to metrology we refer the reader to [31]. Stevens’ theory distinguishes four types of scales of measurement: *nominal*, *ordinal*, *interval* and *ratio* (see table 1). A nominal scale has multiple categories (or classes) with no natural ordering, for example a biological species or a blood group. Although these categories can be represented by numbers, the numbers do not have any quantitative meaning. Properties measured on a nominal scale can only be compared based on their equivalence, i.e. same or different.

An ordinal scale also has multiple categories, but with a natural ordering, for example wind force on the Beaufort scale or stages of cancer. Properties measured on an ordinal scale can be compared based on their equality and their order, e.g. higher/lower, stronger/weaker. Distances between categories, however, are not necessarily meaningful. For example, the

Table 1. Stevens' scales of measurement.

Scale	Examples	Central tendency statistic	Dispersion statistic
Nominal	Biological species, blood group	Mode	See Section 3
Ordinal	Wind force on the Beaufort scale, stage of cancer	Median, mode	Interquartile range (IQR)
Interval	Longitude and latitude, year	Mean, median, mode	Standard deviation, IQR
Ratio	Length, mass	Mean, median, mode	Standard deviation, IQR

difference in the severity of cancer stages I and II is not the same as the difference in the severity of cancer stages III and IV [32].

Unlike for an ordinal scale, for an interval scale the differences are meaningful. Examples of properties measured on an interval scale include longitude and latitude, year, etc. A distinctive feature of this scale is that it does not have a meaningful ('true') zero, which is chosen arbitrarily.

Properties measured on a ratio scale do have a 'true' zero and represent the most encountered properties in physics, for example length, mass, and electrical power.

One of the important aspects of Stevens' theory is that different types of scales have different statistics representing central tendency and dispersion. From a statistical point of view, a nominal scale represents the lowest (the most primitive) scale, while a ratio scale represents the highest scale. Statistics that apply to the lower scales also apply to the higher scales, but not the other way around [33]. The central tendency of properties measured on a nominal scale can only be expressed by the mode (although strictly speaking, the mode is representative of the most frequent class, not the central class). The central tendency of properties measured on an ordinal scale can be expressed by the mode and the median. The notions of central tendency and dispersion are important for understanding how to express uncertainties. In section 3, we discuss expressions of uncertainty for nominal variables.

The nomenclature for nominal variables varies across scientific communities. In the mathematical sciences, nominal and ordinal variables are often merged into a single group, *categorical variables* [20, 33], with nominal variables sometimes referred to as *unordered categorical variables* and ordinal variables referred to as *ordered categorical variables* [20]. In addition, nominal variables are sometimes called *qualitative variables*, and are distinguished from *quantitative variables* (variables measured on interval and ratio scales), with some ambiguity around the classification of ordinal variables [33]. However, in some international standards, the term *qualitative variable* (or *qualitative property*) is used as an umbrella term for both nominal and ordinal variables [32].

In this paper, we use the term *categorical variable* as an umbrella term for nominal and ordinal variables. We do not use the term *qualitative variable* to avoid confusion because, as mentioned above, there is disagreement about whether ordinal variables are included [32, 33]. We also use interchangeably the term *variable* (more common in the language of data

science) and the term *property* (more common in the language of metrology).

2.2. Nominal variables in metrology

The authoritative metrological documents used to focus solely on quantities (quantitative properties). This focus changed about a decade ago, when definitions of nominal properties and ordinal quantities were introduced in the 3rd edition of the VIM, the VIM3 [9, definitions 1.30 and 1.26]. In the preparation of the 4th edition of the VIM (the VIM4), the Joint Committee for Guides in Metrology (JCGM) is holding a discussion on including an entire chapter on nominal properties, motivated by the growing interest in a metrologically rigorous handling of these properties.

Unlike the VIM, the GUM suite of documents is entirely focused on quantities, with the only exception being the fourth document of the GUM suite dedicated to conformity assessment [25], previously mentioned in section 1.

Although the VIM and the GUM currently provide negligible coverage of nominal variables, these variables have been addressed in a number of standards and recommendations produced by international bodies working in various areas of science and industry, where nominal variables are frequently used. For example, the International Federation of Clinical Chemistry and Laboratory Medicine (IFCC) and International Union of Pure and Applied Chemistry (IUPAC) developed 'Vocabulary on nominal property, examination, and related concepts for clinical laboratory sciences (IFCC-IUPAC Recommendations 2017)' [29]. Also, nominal variables and approaches to the evaluation of uncertainties associated with such variables were discussed in 'Assessment of performance and uncertainty in qualitative chemical analysis (EURACHEM/CITAC Guide)' [28]. Recommendations on handling nominal variables specifically related to reference materials, which have nominal properties, were recently systemised in the International Standard ISO 33406:2024 'Approaches for the production of reference materials with qualitative properties' [32]. Discussion on nominal variables and their uncertainty evaluation can also be found in publications by national metrology institutes (e.g. [34]), as well as in publications by the members of JCGM Working Group 2 [35]. It is important to note that although many of the above references address evaluation and expression of uncertainties associated with nominal variables, no agreement has yet been reached on this topic.

3. Uncertainty evaluation for nominal properties

In this section, we outline approaches for uncertainty evaluation of nominal properties and further propagation through measurement models. In section 3.1, we define PDFs and PMFs, illustrating that PMFs are the natural probability distribution for nominal properties. In section 3.2, we describe various statistics that could be used to express uncertainty of nominal variables, highlighting their strengths and weaknesses and the extent of their departure from accepted notions of uncertainty in the metrology community. In section 3.3, we show, with use of the PMF, how uncertainties of nominal properties can be propagated through measurement models having categorical inputs.

3.1. Probability distributions

As stated in the GUM [10] and other metrology-related publications [34], the uncertainty associated with any variable is fully expressed by its probability distribution, which is the distribution around a measured value (often taken as the mean of the distribution) in which the ‘true’ value is expected to lie. This distribution can either be empirical (i.e. based upon the frequency of occurrence of the variable), or have an assumed form (such as Gaussian). For continuous quantitative variables \mathbf{x} that can take values in the set $\mathcal{X} \subset \mathbb{R}^D$, where D is the dimension of the continuous input space, the probability distribution is defined in terms of a PDF [36][37, Definition 2.26], which is a map or function

$$p : \mathcal{X} \rightarrow \mathbb{R}^+, \quad \mathbf{x} \mapsto p(\mathbf{x}), \quad (3.1)$$

where

$$\int_{\mathcal{X}} p(\mathbf{x}) \, d\mathbf{x} = 1. \quad (3.2)$$

Condition (3.2) ensures that expression (3.1) is properly normalised and can act as a PDF.

Similarly, a nominal variable can take values in the set $\mathcal{C}_K = \{c_1, \dots, c_K\}$, where K is the number of possible classes (or categories). Thus, the probability distribution is defined in terms of the PMF [16][37, definition 2.24]:

$$p : \mathcal{C}_K \rightarrow [0, 1], \quad c_k \mapsto p_k := p(c_k), \quad (3.3)$$

where

$$\sum_{k=1}^K p_k = 1. \quad (3.4)$$

Condition (3.4) ensures the expression (3.3) obeys the correct probability sum rule and can act as a PMF.

The PMF of a nominal property can also be fully described in terms of its *probability vector*

$$\mathbf{p} = [p_1, \dots, p_K] \in [0, 1]^K, \quad (3.5)$$

where each element of the array can take values on the interval $[0, 1]$, and condition (3.4) also holds.

Given the analogy with PDF for quantitative variables, the PMF is the natural candidate for expressing uncertainty in nominal properties and has been motivated in metrology-related publications such as [35, 38], as well as the recent international standard [32] introduced in section 2.2.

3.2. Expressions of uncertainty

Expressions of uncertainty for quantitative variables such as the standard deviation are well known in the metrology community, and are used in the GUM [10]. The standard deviation is one example of a summary statistic that expresses the uncertainty of a quantitative variable in terms of a single positive real number. More formally, given a quantitative variable \mathbf{X} with PDF $p(\mathbf{x})$, a statistic u for uncertainty is evaluated using a function F such that

$$u = F(p(\mathbf{x})), \quad u \in [0, \infty). \quad (3.6)$$

The definition of uncertainty according to the VIM [9, 2.26] is, a ‘non-negative parameter characterising the dispersion of the quantity values being attributed to a measurand, based on the information used’. A statistic of dispersion is typically assumed to involve some notion of distance, and it has been questioned whether such a notion exists for nominal properties [35]. The PMF of a nominal property can be interpreted in terms of expected distance with respect to a simple metric. Using the Hamming distance on a single categorical variable [16, chapter 25]

$$s_{jk} = \begin{cases} 0 & c_j = c_k, \\ 1 & \text{otherwise,} \end{cases} \quad j, k \in \{1, \dots, K\}, \quad (3.7)$$

the expected distance to class c_k of a nominal property following the distribution given by \mathbf{p} is given by

$$\sum_{j=1}^K p_j s_{jk} = \sum_{j \neq k} p_j = 1 - p_k. \quad (3.8)$$

Various so-called *statistics of variation* [39] have been proposed for nominal properties, and we present an overview of some of these in section 3.2.3. Given our observation in the previous paragraph that all summary statistics of a PMF have a distance interpretation, we will refer to them throughout as statistics of dispersion.

The VIM also defines a notion of *standard measurement uncertainty* for quantities, namely the standard deviation, which captures dispersion about a central tendency, namely the mean [9, (2.30)]. For nominal properties, the only statistic of central tendency is the mode. Similarly to the notion of distance in (3.7), we can also define a meaningful notion of *distance from the modal class* (or classes in the multimodal case) for nominal properties (see appendix A). Equipped with this notion of distance, some of the statistics we present in section 3.2.3 can indeed be viewed as statistics of dispersion

about the modal class(es). Writing $\hat{p} = \max(\mathbf{p})$ for the probability assigned to the modal class(es), also known as the *modal probability*, those statistics that can be expressed as a function of \hat{p} alone can be interpreted as dispersion about central tendency.

That said, some of the statistics we present in section 3.2.3 are not in general statistics of dispersion from a central tendency. Such statistics do not, therefore, represent a counterpart to the notion of standard measurement uncertainty for quantities.

Such a distinction is only relevant for multi-class nominal variables ($K > 2$). For binary nominal variables ($K = 2$), all statistics of dispersion (from the mode or otherwise) can be expressed as a function of \hat{p} alone, and so capture dispersion about central tendency. For each of the statistics of dispersion defined in section 3.2.3, we state their explicit reduction to a function of \hat{p} alone in the binary case.

3.2.1. Uncertainty and confidence. A distinction is made in the literature e.g. [35, 40], between the *confidence* in the prediction of an ML classification model, and the predictive *uncertainty*. Given a probability vector $\mathbf{p} = [p_1, \dots, p_K]$ of an ML classification model $y = f(x)$, $y \in \mathcal{C}_K$, the classification of the output is typically given by the mode of the PMF, defined as the *modal class* \hat{c} . The confidence in the prediction can then be expressed as modal probability \hat{p} , which is the probability that y is the modal class. Note that this notion of confidence is distinct from a *confidence interval* which is one way of expressing uncertainty for quantitative variables in the GUM and the VIM.

Using the modal probability as a form of uncertainty expression has advantages in the sense that it is directly associated with the model prediction, and as such aligns with the language for expressing uncertainty of quantitative variables [10]. For example, the output of a classification model $\hat{y} = f(x)$ can be written as the pair (\hat{c}, \hat{p}) , where \hat{p} is the probability of the prediction $\hat{y} = \hat{c}$. This language can be extended [35] to include multiple classes, i.e. the ML classification model predicts $\hat{y} = \{c_1, c_2\}$ with probability $p = p_1 + p_2$. However, confidence is not the same as uncertainty. A probability equal to 1 would express complete certainty in the output of a classification model, thus the uncertainty should be 0. It is possible, though, to define a statistic of dispersion about a central tendency which is a linear transformation of confidence (see section 3.2.3). Thus, even though uncertainty and confidence are not the same, they are intimately related.

3.2.2. Properties of an uncertainty statistic. In general, a statistic F expressing the uncertainty associated with a nominal variable y with PMF \mathbf{p} , should satisfy the following properties (adapted from [39, 41]):

1. F gives a non-negative real scalar value: $F(\mathbf{p}) \in [0, \infty)$.

2. F gives 0 if and only if there is no uncertainty. That is, $F(\mathbf{p}^1) = 0$ if and only if the elements of \mathbf{p}^1 take the form

$$p_k^1 := \begin{cases} 1 & c_k = \hat{c}, \\ 0 & \text{otherwise,} \end{cases} \quad k = 1, \dots, K. \quad (3.9)$$

3. F is maximum when there is maximal uncertainty. This is when all probabilities are equal, or equivalently when all classes could be the mode:

$$F(\mathbf{p}) \leq F(\mathbf{p}_{\text{unif}}), \quad (3.10)$$

where $(p_{\text{unif}})_k = 1/K$, $k = 1, \dots, K$ is the categorical uniform distribution.

A normalised statistic F_{norm} can then be defined for nominal properties:

$$F_{\text{norm}}(\mathbf{p}) = \frac{F(\mathbf{p})}{F(\mathbf{p}_{\text{unif}})} \in [0, 1], \quad (3.11)$$

where $F_{\text{norm}}(\mathbf{p}_{\text{unif}}) = 1$ represents maximal uncertainty.

In summary, a (normalised) statistic F expressing the uncertainty associated with the output probabilities of an ML classification model is minimum (zero) when the model predicts a single class with 100% probability, and is maximum (one) when the model predicts equal probabilities to all classes.

In the following analysis, we will consider only normalised statistics that satisfy the above properties, for a fair comparison between expressions of uncertainty (see section 3.2.4). As such, we do not consider statistics that quantify different aspects of uncertainty, such as the expected difference of information to a reference class [42], which has been gaining popularity in LC classification (see case study in section 4.1). This statistic fails on property 3, since the statistic is minimum for a uniform distribution, whereas it should be maximum if it is to represent uncertainty.

3.2.3. Statistics of dispersion. In this section, we give an overview of some of the statistics of dispersion that have been proposed.

Wilcoxon’s variation ratio (WVR)

WVR [41], also known as the mean deviation from the mode, is a natural extension of the concept of confidence (see section 3.2.1) in a classification model prediction to an expression of uncertainty, defined by

$$u_{\text{WVR}}(\mathbf{p}) = 1 - \frac{K\hat{p} - 1}{K - 1}. \quad (3.12)$$

For unimodal distributions only, the WVR can be expressed in terms of the expected distance from the mode. Thus the WVR is analogous to the standard deviation for continuous variables. See appendix A for a derivation of WVR in terms of expected distance from the mode.

For binary classification ($K = 2$), expression (3.12) reduces to

$$u_{\text{WVR}}(\mathbf{p}) = 2(1 - \hat{p}). \quad (3.13)$$

The WVR expresses the remaining uncertainty after taking into account the modal probability \hat{p} of the classification result. However, it does not take into account the entire PMF, and in particular variation in probabilities of the remaining classes. This issue is only relevant in the case of non-binary classification problems, such as the one described in section 4.1.

Universal variation ratio (UVR)

The UVR [43] is an extension of the WVR that explicitly takes account of multimodal distributions, defined as

$$u_{\text{UVR}}(\mathbf{p}) = \frac{K^2}{K^2 - 1} \left(1 - \frac{\hat{p}}{m} \right), \quad (3.14)$$

where m is the number of modes. This extension allows expression of the UVR in terms of the expected distance from the mode in the multimodal case (see appendix A for a derivation), which cannot be achieved for the WVR.

By explicitly including the number of modes, the UVR is discontinuous as \hat{p} varies. This becomes clear in binary classification, where $m = 1$ in all cases except for a uniform distribution when $m = 2$. Therefore, in binary classification expression (3.14) reduces to

$$u_{\text{UVR}}(\mathbf{p}) = \begin{cases} \frac{4}{3}(1 - \hat{p}) & \hat{p} > 0.5, \\ 1 & \hat{p} = 0.5. \end{cases} \quad (3.15)$$

It can be shown [43] that $u_{\text{UVR}}(\mathbf{p}) \leq u_{\text{WVR}}(\mathbf{p})$ for all unimodal distributions.

Standard deviation from the mode (SDM)

Both the WVR and UVR take account of the modal probability \hat{p} , but not the distribution among the remaining classes, which might be important to consider. Motivated by finding a statistic based on the deviation from the mode, which takes account of all values in the PMF, Kvå lseth [44] defined the SDM for nominal variables. In normalised form this standard deviation is

$$u_{\text{SDM}}(\mathbf{p}) = 1 - \sqrt{\frac{\sum_{k=1}^K (\hat{p} - p_k)^2}{K - 1}}. \quad (3.16)$$

For binary classification, expression (3.16) reduces to (see appendix B):

$$u_{\text{SDM}}(\mathbf{p}) = 2(1 - \hat{p}), \quad (3.17)$$

and so in the binary case WVR and SDM are equivalent.

The SDM has the property that $u_{\text{SDM}}(\mathbf{p}) \leq u_{\text{WVR}}(\mathbf{p})$ for all unimodal distributions.

One advantage of using the SDM is that for a large number n of samples used to estimate the output PMF, such as the

output trees of random forests, the SDM is normally distributed in the unimodal case [44] $\hat{u}_{\text{SDM}} \sim \mathcal{N}(u_{\text{SDM}}(\mathbf{p}), \sigma_{\text{SDM}}^2(\mathbf{p}))$ with

$$\sigma_{\text{SDM}}^2(\mathbf{p}) = \frac{\hat{p}(1 - K\hat{p})^2 + \sum_{k=1}^K p_k(\hat{p} - p_k)^2}{n(K - 1)^2(1 - u_{\text{SDM}}(\mathbf{p}))^2} - \frac{(1 - u_{\text{SDM}}(\mathbf{p}))^2}{n}. \quad (3.18)$$

This result allows the use of confidence intervals for the statistic.

Entropy

The information (or Shannon) entropy (not to be confused with thermodynamic entropy) is an expression of information content, originally motivated to describe the theory of signal communication [45].

The normalised version (or relative probability entropy) is defined for nominal properties as

$$H(\mathbf{p}) = \mathbb{E}[-\log_K \mathbf{p}] = -\sum_{k=1}^K p_k \log_K p_k. \quad (3.19)$$

In the binary classification case, expression (3.19) reduces to the original Shannon entropy in bits

$$H(\mathbf{p}) = -\hat{p} \log_2(\hat{p}) - (1 - \hat{p}) \log_2(1 - \hat{p}). \quad (3.20)$$

Information entropy is commonly used in the existing literature as a statistic for expressing the uncertainty in nominal variables [41, 46, 47], as well as the recent international standard [32]. It is also the most commonly used statistic in the ML community for summarising uncertainty in predictions of classification models [15, 19, 40]. The principle of maximum entropy (uncertainty) is also used in Bayesian inference (see section 4.1.3) for selection of prior probabilities [48].

Mari [35] suggested an extension to information entropy given (in normalised form) by

$$H^*(\mathbf{p}) = \frac{K^{H(\mathbf{p})} - 1}{K - 1} = \frac{\prod_{k=1}^K p_k^{-p_k} - 1}{K - 1}. \quad (3.21)$$

Such a transformation can be performed to any expression of uncertainty $F(\mathbf{p})$ for nominal properties to produce a new statistic $F^*(\mathbf{p}) = (K^{F(\mathbf{p})} - 1)/(K - 1)$, which also satisfies the properties in section 3.2.2. This transformation has the property $F^*(\mathbf{p}) \leq F(\mathbf{p}) \forall \mathbf{p}$, and so will lower the variation in the values of the new statistic (see tables 4–6). This allows the statistic to be more robust against small perturbations of the PMF.

One alternative variant on the information entropy is called the α -quadratic entropy [49], defined as

$$H_\alpha(\mathbf{p}) = \frac{K^{2\alpha-1}}{(K-1)^\alpha} \sum_{k=1}^K p_k^\alpha (1 - p_k)^\alpha, \quad \alpha \in (0, 1]. \quad (3.22)$$

If α is close to zero, this statistic is not very sensitive to changes in \mathbf{p} , but for α close to one, this statistic has higher

sensitivity [50], and has been used in evaluating uncertainties in LC classification models [42, 50].

Both of the statistics given in (3.21) and (3.22) reduce to statistics of dispersion about the mode in the binary case, but we omit the details for the sake of brevity.

Index of qualitative variation (IQV)

The IQV [41], also known as the quadratic score [51], is another common expression of uncertainty that is widely accepted in the social sciences [43]. It is defined in normalised form as

$$u_{\text{IQV}}(\mathbf{p}) = \frac{K}{K-1} \left(1 - \sum_{k=1}^K p_k^2 \right), \quad (3.23)$$

which for binary classification reduces to

$$u_{\text{IQV}}(\mathbf{p}) = 4\hat{p}(1-\hat{p}). \quad (3.24)$$

The IQV is a special case of the α -quadratic entropy in expression (3.22) with $\alpha = 1$.

Coefficient of nominal variation (CNV)

A variant of the IQV proposed in [52] is the CNV defined as

$$u_{\text{CNV}}(\mathbf{p}) = 1 - \sqrt{1 - u_{\text{IQV}}(\mathbf{p})} = 1 - \sqrt{\frac{K \sum_{k=1}^K p_k^2 - 1}{K-1}}. \quad (3.25)$$

This transformation ensures that $u_{\text{CNV}}(\mathbf{p}) \leq u_{\text{IQV}}(\mathbf{p})$ for all distributions, with equality for \mathbf{p}^1 and \mathbf{p}_{unif} . In terms of the normed distance $d(\mathbf{p}, \mathbf{q}) = \|\mathbf{p} - \mathbf{q}\|_2$ between two PMFs \mathbf{p} and \mathbf{q} , the CNV can be rewritten as

$$u_{\text{CNV}}(\mathbf{p}) = 1 - \frac{d(\mathbf{p}, \mathbf{p}_{\text{unif}})}{d(\mathbf{p}^1, \mathbf{p}_{\text{unif}})}, \quad (3.26)$$

giving the CNV a natural interpretation in terms of the distance of the PMF from a uniform distribution. For binary classification, expression (3.25) reduces to (see appendix B):

$$u_{\text{CNV}}(\mathbf{p}) = 2(1-\hat{p}), \quad (3.27)$$

which is equivalent to the WVR and the SDM.

3.2.4. Comparing statistics of uncertainty. The statistics described in section 3.2 have various advantages and disadvantages. In this section we compare their properties and performance on a few example predictive probability distributions. One particular example in binary classification is the PMF \mathbf{p}_A shown in figure 3. Since \mathbf{p}_A is halfway between total certainty and maximal uncertainty, i.e. $\hat{p}_A = 0.75$, it was suggested in [39] that the statistic of uncertainty should also be halfway between the extreme values, i.e. 0.5. We also consider the multi-class PMFs \mathbf{p}_B and \mathbf{p}_C , which have the same model prediction \hat{y} , with the same confidence (modal probability) of 0.5. The example PMFs are visualised in figure 3 for comparison.

Table 2 compares the various expressions of uncertainty for nominal properties. The first point to note is the result that $F(\mathbf{p}_B) \geq F(\mathbf{p}_C)$ for all F considered, even though at first glance it might appear more obvious which is the best prediction of the classification model in \mathbf{p}_B . This result is due to the fact that the statistics are expressing the certainty that the mode can either be one of two possibilities in \mathbf{p}_C (i.e. $(p_C)_1 + (p_C)_2 = 0.96$), as the distribution is peaked about these two possible classes. Whereas for \mathbf{p}_B , the probability distribution is spread over all classes, leading to more uncertainty in the output class.

If we expect the uncertainty of \mathbf{p}_A to be 0.5, then we see that WVR, SDM and CNV satisfy this condition. In fact, WVR, SDM and CNV have the same expressions for binary classification (see appendix B). These three statistics also have the same values for \mathbf{p}_B because expressions (3.16) and (3.25) simplify to (3.12) when all the non-modal classes have the same probability (see appendix B). Thus, we only see a difference in values for \mathbf{p}_C , with WVR and CNV giving the highest and lowest values respectively.


If we wish the statistic to use explicitly the modal probability (or confidence), as well as information about the entire distribution, then SDM is the only option. However, the SDM shows little difference in uncertainty between \mathbf{p}_B and \mathbf{p}_C , where a larger difference from studying figure 3 might be expected. In this case, CNV might be the more appropriate choice, but at the cost of no longer explicitly using information about the mode.

3.3. Uncertainty propagation with nominal properties


In section 3.2, we considered various ways in which the uncertainty in nominal properties can be expressed through a summary statistic. However, to propagate such uncertainty through a chain of measurement models, it is more appropriate to consider the full PMF. For quantitative properties, it is more difficult to propagate the full PDF analytically due to intractable integrals, and it is commonly the case that the full PDF is unknown and modelled using simpler forms such as assuming a Gaussian distribution. In such cases, summary statistics, such as the standard deviation, are propagated using the law of propagation of uncertainty [21]. Alternatively, where the PDF is well known, sampling techniques such as MC can propagate the full PDF [22].

In the case of a PDF having an analytical form, such as a Gaussian or rectangular distribution, the full PDF can be directly inferred from summary statistics, such as the standard deviation. On the other hand, the summary statistics for nominal properties described in section 3.2 cannot be used to infer the PMF except in simplistic examples such as binary classification. It is therefore recommended to use the full PMF in uncertainty propagation through multistage measurement models that involve nominal properties.

Since, for nominal properties, the PMF can be defined explicitly, it can be used to obtain analytical results for propagation through measurement models. Consider an ML classification model f with outputs $y \in \mathcal{C}_K$, which have a



(a) Example binary classification PMF: $\mathbf{p}_A = [0.75, 0.25]$.



(b) Example multi-class classification PMFs: $\mathbf{p}_B = [0.5, 0.1, 0.1, 0.1, 0.1, 0.1]$ and $\mathbf{p}_C = [0.5, 0.46, 0.01, 0.01, 0.01, 0.01]$.

Figure 3. Example predictive PMFs as output from ML classifiers.

Table 2. A comparison of expressions of uncertainty for nominal properties.

Statistic	Uses	Uses	\mathbf{p}_A	\mathbf{p}_B	\mathbf{p}_C
	confidence	full PMF			
WVR (3.12)	Y	N	0.50	0.60	0.60
UVR (3.14)	Y	N	0.33	0.51	0.51
SDM (3.16)	Y	Y	0.50	0.60	0.56
H (3.19)	N	Y	0.81	0.84	0.50
H^* (3.21)	N	Y	0.75	0.69	0.29
IQV (3.23)	N	Y	0.75	0.84	0.65
CNV (3.25)	N	Y	0.50	0.60	0.40

predictive distribution $\mathbf{p} = [p_1, \dots, p_K]$. Assume the outputs y become inputs to the next measurement model g , which has both nominal and quantitative input variables, in the form of a D -dimensional vector \mathbf{x} , and outputs a quantitative variable $z \in \mathbb{R}$:

$$z = g(\mathbf{x}, y), \quad \mathbf{x} \in \mathcal{X} \subset \mathbb{R}^D, y \in \mathcal{C}_K, z \in \mathbb{R}. \quad (3.28)$$

Since there are K possible values for y , we can describe g as a model conditioned on a nominal variable, and therefore express g in terms of K models $z = g(\mathbf{x}, c_k) = g_k(\mathbf{x})$, where $\mathbf{x} \in \mathcal{X}_k \subset \mathbb{R}^D$ for $k = 1, \dots, K$. This expression can be thought of as modelling a system which has K different regimes, defined by the classes c_k .

The expected value of the output z can be expressed in terms of the expected values of the functions $g_k(\mathbf{x})$:

$$\begin{aligned} \mathbb{E}[z] &= \sum_{k=1}^K p_k \int_{\mathcal{X}_k} d\mathbf{x} q_k(\mathbf{x}) g(\mathbf{x}, c_k) \\ &= \sum_{k=1}^K p_k \int_{\mathcal{X}_k} d\mathbf{x} q_k(\mathbf{x}) g_k(\mathbf{x}) = \sum_{k=1}^K p_k \mu_k, \end{aligned} \quad (3.29)$$

where $q_k(\mathbf{x})$ are the PDFs of the input variable $\mathbf{x} \in \mathcal{X}_k$ and $\mu_k = \mathbb{E}_{\mathbf{x} \sim q_k(\mathbf{x})}[g_k(\mathbf{x})]$. The means μ_k can be evaluated in the standard way for quantitative input variables.

We can also compute the variance of the output z in terms of summary statistics of the functions $g_k(\mathbf{x})$:

$$\begin{aligned} \text{Var}[z] &= \mathbb{E}[z^2] - \mathbb{E}[z]^2 \\ &= \sum_{k=1}^K p_k \mathbb{E}[g_k(\mathbf{x})^2] - \sum_{j,k=1}^K p_j p_k \mu_j \mu_k \\ &= \sum_{k=1}^K p_k (\sigma_k^2 + \mu_k^2) - \sum_{j,k=1}^K p_j p_k \mu_j \mu_k \end{aligned} \quad (3.30)$$

where $\sigma_k^2 = \text{Var}_{\mathbf{x} \sim q_k(\mathbf{x})}[g_k(\mathbf{x})]$ can be evaluated using the methods described in [21] for computing expressions of uncertainty in quantitative variables.

For this simple example, the moments of the measurand given nominal property inputs can be determined analytically. More generally, MC techniques can be applied by sampling directly from the PMF to determine the output PDF $p(z)$. Proceeding in this way is relevant to measurement applications that depend upon a prior classification, such as in determining particle concentration from counting measurements [53].

4. Case studies

In this section we present two applications of uncertainty evaluation using ML classification models: LC classification (section 4.1), and AF detection using PPG signal data (section 4.2). In both cases we describe the background on the challenge, methods for obtaining class probabilities, and exploration of the contributions of different sources of uncertainty.

4.1. LC classification

4.1.1. Background. LC maps produced by classification of satellite imagery provide essential information on the Earth's biophysical cover. This information serves as an important input into a variety of applications, e.g.

- Hydrological and climate modelling [54, 55], where LC maps are used for setting up land surface conditions;
- National greenhouse gas accounting [56], where LC maps along with additional data are used for estimation of emissions associated with anthropogenic land use;
- Biodiversity monitoring [57], where LC maps provide information on habitats and their fragmentation, etc.

Evaluation of uncertainties associated with LC maps is important both in its own right for the interpretation of such maps, and for the propagation of these uncertainties further—to the downstream applications of LC maps mentioned above. As stated in section 3.3, propagation of uncertainties of a nominal variable requires its full PMF. This PMF should include contributions of all uncertainty sources mentioned in section 1.2 and illustrated in figure 1.

LC mapping is a multi-class classification problem, which currently is widely approached by ML. The most popular ML classification models used for LC mapping are tree-based models (such as random forests) [58–61], and convolutional neural networks (CNNs) [62, 63]. There are many LC classification schemes. Some classification schemes include only a few, very broad, classes, e.g. tree cover, grassland, cropland, etc [64], whereas other classification schemes provide finer details, e.g. they distinguish broad-leaved and coniferous forest [65]. Some classification schemes have a nested structure where broad classes are further split into narrower classes [60]. It is important to note that in the present case study we approach LC classification in similar fashion to how it is commonly approached in the literature—as a classification problem with a univariate output, i.e. each pixel considered

separately. An important extension would be to perform classification with a multivariate output. In this case, classification would produce a collection of output variables (a collection of LC classes) corresponding to multiple pixels as opposed to each single pixel. Such an extension would make sense given how strongly correlated EO data typically is in the spatial domain.

4.1.2. Data. As previously mentioned in section 1.1, a supervised ML classification model is trained on a data set consisting of pairs of several input variables and a categorical output. In our LC classification study, the input variables represented surface reflectance measurements performed in ten spectral bands of Sentinel-2 satellite over a single pixel. The output variable represented LC class in that pixel. Information on LC classes (i.e. LC labels) is typically obtained either manually (by visual interpretation of high/very high resolution satellite imagery or by ground surveys), automatically (e.g. based on pre-existing LC maps), or hybrid-like (following a combination of the two methods) [66]. In our study, LC labels were obtained by an automatic collection procedure following a modified approach of UK Centre for Ecology and Hydrology (UKCEH), used for the collection of training data for the production of UK annual LC maps [59]. This method relies on the assumption that LC change is gradual. It uses UKCEH LC maps produced for three previous years to find pixels with stable LC classes (i.e. unchanged LC classes throughout these years). Further filtering was performed to identify such unchanged pixels where the RF classifier used by the UKCEH to produce their LC maps yielded a probability greater than 95%. In other words, we only retained pixels, where the LC label was assigned with the highest confidence. For the present case study, LC labels were collected over a 20 km × 20 km area of interest (AOI) in Scotland. The collection process yielded 189 142 unchanged pixels (i.e. 4.7% of the AOI) corresponding to forest, cropland, grassland and settlements classes.

The input variables (i.e. surface reflectance measurements) for the 189 142 unchanged pixels were retrieved from Sentinel-2 images collected over the AOI on the 1st of June, 2020 and 2021. Although space-borne surface reflectance measurements have many uncertainty sources associated with them, in the present study we choose to focus on one uncertainty source. This uncertainty source is related to the imperfect knowledge of AOD used in the atmospheric correction of the satellite imagery. It represents just one (although ordinarily dominant) contribution to the uncertainty sources (1) and (5) described in section 1.2 and illustrated in figure 1. The AOD uncertainties were introduced into the atmospheric correction process by MC sampling (we used 25 MC samples). The process yielded 25 realisations of surface reflectance measurements for each of the 189 142 unchanged pixels. The relatively low number of MC samples used in this study reflects its preliminary nature. The current implementation serves as a proof-of-concept, and we plan to increase both the number of MC samples and the range of uncertainty sources in future

iterations of the framework. See [17] for more details on the data used in this study.

4.1.3. Propagation of uncertainties within Bayesian classification models. As illustrated in figure 1, ML classification models come with multiple sources of uncertainty. Most of these sources can be described in terms of probability distributions for particular parametric classification models (such as Naive Bayes, logistic regression, discriminant analysis [16] and neural networks [67]) that naturally lend themselves to a Bayesian treatment. In this section we summarise a framework that can be followed to incorporate these sources of uncertainty for Bayesian generative ML classifiers, which explicitly model the input data uncertainties. For specific details of the framework and its implementation on LC data, with reference to various uncertainty sources, see [17].

Let $\mathcal{D} = \{(\mathbf{x}_i, y_i) | i = 1, \dots, N\}$ be the training data consisting of N pairs of input and output observations, where $\mathbf{x}_i \in \mathcal{X} \subset \mathbb{R}^p$ and $y_i \in \mathcal{C}_K$. Then we can model the sources of uncertainty described in figure 1 in the following way:

- The training data uncertainty (1) from a measurement process can be described in terms of a probability distribution $p(\mathbf{x})$.
- The class assignment uncertainty of the ML classification task (2) derives from the joint empirical probability distribution $p(\mathbf{x}, y)$. This can be separated into two contributions via Bayes' rule

$$p(\mathbf{x}, y = c_k) = p(\mathbf{x} | y = c_k) p(y = c_k), \quad k = 1, \dots, K, \quad (4.1)$$

$p(\mathbf{x} | y = c_k)$ is the *class conditional* probability distribution describing the distribution of the input data for each class. $p(y = c_k)$ is the probability distribution describing the prior class distribution (or *class imbalance*) in the training data.

- ML models also come with model hyperparameters, which describe the choice of model (3), such as model architecture or optimisation method. Various hyperparameter optimisation tuning methods [68] can be used to determine the optimal model choice. Bayesian optimisation is one popular approach [69], and for such an approach, quantification of the uncertainties in choice of model has been attempted [70].
- A parametric classification model can be described by a probability distribution $p(y_i | \mathbf{x}_i, \theta)$, where θ are the parameters of the model. The uncertainty associated with such a model is given by the probability distribution of the model parameters. The likelihood function \mathcal{L} of the model is expressed in terms of the model distribution on the training data, and any additional training data uncertainty

$$\mathcal{L}(\theta | \mathcal{D}) = p(\mathcal{D} | \theta) = \prod_{i=1}^N p(\mathbf{x}_i, y_i | \theta) = \prod_{i=1}^N p(y_i | \mathbf{x}_i, \theta) p(\mathbf{x}_i). \quad (4.2)$$

In a Bayesian setting, Bayes' rule is used to determine the posterior distribution in the model parameters (4) (otherwise

known as *Bayesian inference*)

$$p(\theta | \mathcal{D}) = \frac{p(\mathcal{D} | \theta) p(\theta)}{p(\mathcal{D})}. \quad (4.3)$$

The Bayesian approach introduces an additional uncertainty $p(\theta)$, which is the prior distribution on the model parameters. This can be determined from expert knowledge, or one can use conjugate priors [16] if known for analytic determination of the posterior. $p(\mathcal{D})$ is known as the *evidence*, and is typically intractable for more complex classification models such as CNNs. In such cases one must resort to sampling techniques such as Markov Chain MC [71], approximation techniques such as variational inference [16] or MC dropout [19] to estimate the posterior distribution.

- Just as with training input data, new input data also comes with measurement uncertainty (5). This can be described in terms of a probability distribution $p(\mathbf{x}_{\text{new}})$.
- The probability distribution associated with the new input data uncertainty on an input \mathbf{x}_{new} (known as the *posterior predictive*) can be evaluated as the expectation of the ML model probability of class k with respect to the uncertainty in the model parameters $p(\theta | \mathcal{D})$:

$$p(y_{\text{pred}} = c_k | \mathbf{x}_{\text{new}}, \mathcal{D}) = \mathbb{E}_{\theta | \mathcal{D}} [P(y_{\text{pred}} = c_k | \mathbf{x}_{\text{new}}, \theta)]. \quad (4.4)$$

The posterior predictive distribution of the ML classification model can then be written in the form of a PMF with probability vector $\mathbf{p} = [p_1, \dots, p_K]^\top$, where $p_k := p(y_{\text{pred}} = c_k | \mathbf{x}_{\text{new}}, \mathcal{D})$. This PMF can then be used to propagate uncertainties in further measurement models where the output of the ML classification model is used as an input, either analytically (as shown in section 3.3), or through MC sampling of the PMF.

4.1.4. Numerical examples. We apply the Bayesian framework described in section 4.1.3 to the datasets described in section 4.1.2, using the Bayesian QDA (BQDA) model presented in [17]. Specifically we compute the posterior predictive probabilities on test sets (90% of available pixels) from years 2020 and 2021, with training set taken from the remaining 10% of pixels from year 2020.

The confusion matrices (comparing class predictions and 'true' class labels) from taking the maximum probability as the predicted class are given in figure 4. We include the true/false positive and true/false negative rates to see how the model performs on each class, as well as computing the overall classification loss, which is the misclassification rate, or fraction of class predictions not from the true labelled class.

For evaluation of the performance of the output class probabilities from the classification models, we use two expected values of proper scoring rules [72], namely the multi-class cross-entropy loss (XE) and the Brier score (BS). Proper scoring rules are a principled way to assess the quality of the predictive probabilities produced by ML classification models [73].

Assume we have an ML classification model and a test dataset with N observations (i.e. pixels in the case of LC). For




Figure 4. Confusion matrices using BQDA. Final column corresponds to the true/false positive rates for each class. Final row corresponds to true/false negative rates.

each observation $i = 1, \dots, N$, the test dataset has a class label $k = 1, \dots, K$. This can be represented by the binary matrix y_{ik} , where $y_{ik} = 1$ if test observation i has class label k , and zero otherwise. The output predictive probabilities of the model for each test observation i and class label k is represented by the probability matrix p_{ik} . The XE and BS for each test observation i are then defined as

$$XE_i = - \sum_{k=1}^K y_{ik} \log(p_{ik}), \quad BS_i = \frac{1}{K} \sum_{k=1}^K (y_{ik} - p_{ik})^2. \quad (4.5)$$

The XE is similar to the information entropy defined in section 3.2.3 except that instead of only using the probability of being the most likely class from the model output, the output probabilities are used to calculate agreement to an independent test dataset, so that ‘predictive performance’ is evaluated. The XE is commonly used in neural networks classifiers as the objective function to minimise during training. The BS is the mean squared difference between the test class label and the predicted probability. It considers all output probabilities, whereas XE only considers the probability associated with the test class label.

To evaluate overall performance, we take expected values over the XE and BS defined in expression (4.5) for each test pixel i , as well as normalising the values such that 0 is perfect prediction, and 1 represents a predictive performance no better than using the prior class distribution of the test dataset as a trivial classifier (i.e. the fraction of LC class abundance within the AOI). The expected XE (EXE) and BS (EBS) are defined as

$$EXE = \frac{1}{N} \sum_{i=1}^N \frac{\sum_{k=1}^K y_{ik} \log(p_{ik})}{\sum_{k=1}^K q_k \log(q_k)}, \quad EBS = \frac{1}{N} \sum_{i=1}^N \frac{\sum_{k=1}^K (y_{ik} - p_{ik})^2}{\sum_{k=1}^K q_k (1 - q_k)}, \quad (4.6)$$

where q_k is PMF of classes k from the test dataset prior to ML classification.

Table 3. Performance metrics for LC classification using BQDA model.

Year	Loss	EXE	EBS
2020	0.012	0.079	0.031
2021	0.057	0.567	0.151

Table 4. Summary of uncertainty statistics for LC classification. Test year: 2020.

Statistic	Median/ 10^{-8}	Mean/ 10^{-3}	IQR/ 10^{-5}	s.d./ 10^{-2}
WVR	4.63	8.10	2.72	5.00
UVR	3.71	6.48	2.18	4.00
SDM	4.63	7.81	2.72	4.72
H	88.8	12.9	24.6	5.84
H^*	41.0	7.35	11.3	3.63
IQV	9.26	12.4	5.45	6.73
CNV	4.63	7.04	2.72	4.02

4.1.5. Results. The resulting metric values for assessment of the performance of the classifiers are given in table 3. As is common with ML classification tasks, we see a degradation in performance when the model is tested on data outside the training domain (in this case a different year). We also observe from figure 4 that the class ‘Settlements’ have the highest false positive rate. This is in large part due to the low abundance of settlement pixels within the AOI, and so the ML classifier will tend to give other LC classes with larger abundances a higher predictive probability.

In addition to the classification results, we also compute the expressions of uncertainty for nominal properties described in section 3.2 using the posterior PMFs obtained within the Bayesian framework of [17]. Tables 4 and 5 show a summary of the behaviour in the chosen uncertainty statistics, for test years 2020 and 2021 respectively.

In both tables, we see that the median and mean values are very small. This is due to the high accuracy or confidence of the LC classifiers, leading to low uncertainty. In both tables, we also see a pattern that the median is orders of magnitude

Table 5. Summary of uncertainty statistics for LC classification. Test year: 2021.

Statistic	Median/ 10^{-5}	Mean/ 10^{-2}	IQR/ 10^{-3}	s.d./ 10^{-2}
WVR	7.44	3.03	3.16	9.94
UVR	5.95	2.42	2.53	7.95
SDM	7.44	2.91	3.16	9.34
H	51.9	4.33	13.2	11.1
H^*	24.0	2.56	6.15	7.17
IQV	14.9	4.48	6.31	13.1
CNV	7.44	2.59	3.16	7.97

smaller than the mean, and the IQR is orders of magnitude smaller than the s.d. This is because the distributions of the uncertainty statistics are highly skewed to the left, as again most of the classifier predictions were highly confident, leading values of the uncertainty statistics to be close to zero.

Note the large difference in scales between tables 4 and 5. This is due to the large classification performance degradation observed in the classification metrics from years 2020 to 2021 (see figure 4), causing the resulting predictive PMFs to be more uncertain. However, the uncertainty statistics still show high certainty in comparison to a uniform distribution. We also see a clear pattern for both years, with the information entropy H having the largest median and IQR, and UVR having the smallest. Since the median and IQR represent the behaviour of the statistics for small values, the entropy is the most sensitive to variations in confident PMFs out of all the uncertainty statistics considered. UVR is the least sensitive in this case. Alternatively, the IQV has the largest s.d., and largest mean for year 2021, indicating that this statistic is more sensitive to variations in uncertain PMFs, i.e. PMFs that are close to uniform. The transformed entropy H^* is the least sensitive, or most robust against small changes in uncertain PMFs.

We observe that WVR, SDM, and CNV have equivalent median and IQR in tables 4 and 5. This is because the median and IQR are describing PMFs with a large \hat{p} in comparison with the other classes. This reduces these expressions for the uncertainty statistics of (3.16)(3.25) to approximately (3.12), essentially acting as a binary classification PMF where all other classes are treated as one. Appendix B gives a derivation of this equivalence.

4.2. AF detection

4.2.1. Background. AF is a heart condition that causes an irregular and often fast heart rate. The global prevalence of AF was estimated to be 59 million individuals in 2019. Patients with AF are at increased risk of stroke, heart failure and dementia, and it is considered to be a 21st century cardiovascular disease epidemic [74].

Digital health solutions are often used for the timely detection of AF, and diagnosis of AF is typically based upon an ECG. Wearable or smartphone-based PPG technology has more recently emerged as having clear potential to make continuous monitoring of AF more widely available. PPG measurements use a light source and a detector on the surface of

the skin to measure changes in the optical properties of the skin's microvascular bed caused by relative changes in blood volume [75].

Detection of AF from a time series segment of PPG measurements can be posed as a binary classification problem with two categories: *positive* (AF is present) and *negative* (AF is not present). Promising results for this classification task have already been reported. For example, a 122-patient study in [76] reported a specificity (Type I error) of 98.9% and a sensitivity (Type II error) of 100% for *Fibrichck*'s PPG-based AF detection smartphone app [77].

4.2.2. Data. We make use of a labelled PPG dataset called *DeepBeat* introduced in [7]. A custom subset of the dataset was selected to ensure a balanced label distribution and no overlap of patients between the training and test datasets. A total of 134 patients were selected, giving a derived dataset consisting of over 100 000 signal segments, with over 50 000 instances of both AF and non-AF. Each segment has a duration of 25 s and a sampling rate of 32 Hz. See [7] for further details on the primary dataset, and see [75] for further details on the derived dataset.

4.2.3. Probabilistic ML model. There are generally two approaches to the AF classification task [78]: (i) a rule-based classifier which uses features relating to pulse wave shape or inter-beat intervals; or (ii) an ML-based classifier which uses a training set of PPG data labelled by a clinical expert to optimise a generic empirical classification model (for example a neural network). This case study focuses on the latter approach.

Uncertainty evaluation is paramount in any classification model which informs medical diagnosis, where the consequences of misdiagnosis can be significant. Confidence in AF classification is affected by signal quality, and an evaluation of uncertainty is needed for clinicians to establish whether a classification can be used to reliably inform diagnosis. This case study therefore considers the use of a probabilistic ML classifier which, for each observation, outputs probabilities of being in each of the two classes.

Recall that the dataset described in [7] consists of signal segments from 134 patients. In [75], data from 110 of the patients was used to train the model, which allows for data from the remaining 24 patients to be used as examples of the use of the trained model for probabilistic AF classification.

We give a brief overview of the ML model and its optimisation, and refer the interested reader to [75] for further details. We use a CNN model which is a close variant of *xresnet1d50* reported in [79], which is itself a 1D variant of *xresnet* [80]. A CNN model consists of a sequence of linear transformations (*convolutions*) interspersed with simple non-linear operations, and it outputs prediction probabilities for each of the classes. Such a model is parametrised by a very large number of weights and biases defining each of the convolutions, and training the model involves optimising these parameters with respect to some loss function.




Figure 5. Confusion matrix for AF classifier on test dataset.

The probabilities returned by a CNN capture contributions to the prediction uncertainty arising from some of the sources of uncertainty mentioned in section 1, namely the class assignment uncertainty (2), the training data uncertainty (1) and the input test data uncertainty (5). These sources of uncertainty are those which are properties of the classification task itself or the measurement system used to obtain the observations. In the ML community, such sources of uncertainty are often referred to as *aleatoric*. Given observations, this is the contribution to the prediction uncertainty that cannot be reduced by choosing a better model or obtaining additional training samples. Aleatoric uncertainty is often distinguished in ML from *epistemic* uncertainty, which is the uncertainty concerning the choice (3) and optimisation (4) of the model [15].

In order to capture in addition contributions to the prediction uncertainty due to model parameter uncertainty (4), an approach known as *MC Dropout* [18, 19] is used. The approach involves enforcing convolution weights to be zero at random with some fixed probability when both training and evaluating the model. A probability distribution is finally obtained which combines the aleatoric and epistemic contributions to the prediction uncertainty [19].

We note two sources of uncertainty which are inadequately captured by the MC dropout approach. Firstly, as is usually the case for uncertainty evaluation of ML models, uncertainty concerning the choice of model is ignored. It follows therefore that reliable uncertainty evaluation assumes the selection of an adequate type of model. Secondly, only uncertainty of the *output* data is explicitly modelled in this approach, and uncertainty concerning *input* data is not explicitly modelled. This is in contrast to the BQDA approach in section 4.1. Approaches which explicitly model uncertainty of the inputs to deep learning models have been considered [81], but the approach is computationally expensive and consequently not practically feasible for the large-scale AF detection task considered here. Full details of the involved steps and their linkage, and how

Table 6. Uncertainty statistics for AF classifier.

Statistic	Median/ 10^{-2}	Mean/ 10^{-1}	IQR/ 10^{-1}	s.d./ 10^{-1}
WVR	6.98	2.83	5.87	3.41
UVR	4.66	1.88	3.92	2.28
SDM	6.98	2.83	5.87	3.41
H	21.8	4.01	8.71	4.10
H^*	16.3	3.76	8.30	4.02
IQV	13.5	3.69	8.29	4.04
CNV	6.98	2.83	5.87	3.41

aleatoric and epistemic uncertainty are captured, can be found in [75].

4.2.4. Numerical examples. We analyse the output predictive probabilities of the AF detection classifier using the same procedure as described in section 4.1.4 for LC classification. The confusion matrix and summary statistics are given in figure 5 and table 6 respectively. The classification performance metrics are: classification loss = 0.209, EXE = 0.874, and EBS = 0.622. Since the loss and other classification metrics are larger than in the case of LC classification, the uncertainty statistics in table 6 are also larger than those in tables 4 and 5. The WVR, SDM and CNV give equivalent values as this is a binary classification problem, and their expressions are equivalent in this case (see appendix B). However, we observe a similar pattern to the uncertainty statistics computed for the LC case study. Namely that the entropy is the most sensitive to variations in the predictive probability distributions, with UVR being the least sensitive.

4.3. Summary of results

The application of our ML classification uncertainty framework across different models and datasets allowed evaluation of the predictive PMFs, whilst taking into account the various

sources of uncertainty described in figure 1. All uncertainty statistics of nominal properties described in section 3.2.3 were evaluated based on the outputted predictive PMFs, with results given in tables 4–6.

The results show consistent behaviour in the predictive PMFs, in that more accurate ML classification models produced uncertainty statistics with values close to zero. The results also showed that the information entropy was the most sensitive overall to variations in the PMF out of all uncertainty statistics considered, with UVR being the most robust. However, if one is more concerned with assessing highly uncertain PMFs, the IQV is the most sensitive. We also demonstrated that three uncertainty statistics (WVR, SDM, and CNV) are equivalent in the case of binary classification, and approximately equivalent for multi-class ML classification models with highly confident predictive PMFs.

5. Conclusion

In this paper, we have presented a metrological framework for evaluating uncertainty of nominal properties, motivated by the increasing use of ML classification models in a wide range of applications. We have identified and listed a number of sources of uncertainty within ML classification models which contribute to the output predictive uncertainty of the nominal property or variable, namely training data uncertainty, class assignment uncertainty, uncertainty about choice of model, model parameter uncertainty, and input data uncertainty, only some of which are currently addressed in the GUM suite of documents.

The work described in this paper has focused primarily on predictive uncertainty of ML classification models, with a specific focus on probabilistic ML classification models, i.e. models that return class probabilities for the output nominal variable. This focus included the evaluation of uncertainty for such nominal variables, which is not currently considered within the VIM and GUM. We have demonstrated that the PMF provides a complete characterisation of the uncertainty associated with nominal properties, and thus the predictive uncertainty in probabilistic ML classifiers. We have also motivated the use of the PMF of a nominal variable, as opposed to a summary statistic, in propagating such uncertainty through a multistage measurement model, representing either a metrological traceability chain or a processing chain for information products. We then provided example approaches using analytic or sampling methods of how to perform such propagation using the PMF.

We have reviewed various statistics that have been proposed within the literature associated with expressing the predictive uncertainty of probabilistic ML classifiers. We provided a comparison between these statistics in terms of their fundamental properties, and applied them in the context of two case studies that rely upon a metrological analysis, namely LC classification, and AF detection. We found that for both case studies, the information entropy, which is the most common statistic for uncertainty within the ML community, is

in general, the most sensitive to variations in the predictive probability distributions. We have also shown that three of the uncertainty statistics (WVR, SDM and CNV) are equivalent for binary classification, and approximately equivalent in the multi-class case for confident ML classification model outputs. We have also illustrated how sources of uncertainty are addressed within the two case studies using different methodologies dependent on the classification modelling procedure.

The framework we have outlined builds upon and enhances previous work in uncertainty evaluation for nominal properties. The framework enables an extension of the GUM to uncertainty evaluation for nominal properties, and provides methods for handling the sources of uncertainty within ML classification, underpinning their use within metrological applications.

Acknowledgment

The authors would like to thank Emma Woolliams (NPL) and Ian Smith (NPL) for providing helpful reviews and feedback. We would also like to thank Richard Brown, Xavier Loizeau, Samuel Hunt, and Bernardo Mota (all NPL) for their contributions and insights during the workshops held at NPL on 27th March and 27th November 2024.

Funding

This work was funded by the UK Government's Department for Science, Innovation and Technology (DSIT) through the UK's National Measurement System (NMS) programmes. Anna Pustogvar's work was also supported by a College of Science and Engineering (CSE) Scholarship from the University of Leicester.

Data sources

The original data used for LC classification study can be found at [82–85]. More details on how this data was processed for the LC classification study are provided in [17]. The original data for the case study on AF detection (*DeepBeat*) can be found at [86] and is described in [7]. The custom data split used was described in [75].

Appendix A. Derivation of uncertainty statistics WVR and UVR in terms of expected distance from mode

We can express the WVR and UVR defined in section 3.2.3 in terms of a simple distance metric from the mode(s), defined as

$$d_k = \begin{cases} 0 & p_k = \hat{p}, \\ 1 & \text{otherwise,} \end{cases} \quad k = 1, \dots, K. \quad (\text{A.1})$$

WVR

Using the fact that $\hat{p} = 1 - \mathbf{d} \cdot \mathbf{p}$ for unimodal distributions, the expression for WVR (3.12) becomes

$$\begin{aligned} u_{\text{WVR}}(\mathbf{p}) &= 1 - \frac{K\hat{p} - 1}{K - 1} \\ &= 1 - \frac{K(1 - \mathbf{d} \cdot \mathbf{p}) - 1}{K - 1} \\ &= \frac{K}{K - 1} \mathbf{d} \cdot \mathbf{p}. \end{aligned} \quad (\text{A.2})$$

Since the probability vector \mathbf{p} represents a PMF, we see that $\mathbf{d} \cdot \mathbf{p} = \mathbb{E}_{\mathbf{p}}[\mathbf{d}]$. This is the expected distance from the mode. Thus, we can express the normalised WVR in terms of the expected distance

$$u_{\text{WVR}}(\mathbf{p}) = \frac{K}{K - 1} \mathbb{E}_{\mathbf{p}}[\mathbf{d}]. \quad (\text{A.3})$$

We see a direct analogy with the variance for continuous variables in terms of the expected squared distance from the mean: $\text{Var}[X] = \mathbb{E}_p[d(X)^2]$, where $d(X) = X - \mathbb{E}_p[X]$.

UVR

For multimodal distributions, $m\hat{p} = 1 - \mathbb{E}_{\mathbf{p}}[\mathbf{d}]$, where m is the number of modes. In this case, the UVR in expression (3.14) can be written as

$$\begin{aligned} u_{\text{UVR}}(\mathbf{p}) &= \frac{K^2}{K^2 - 1} \left(1 - \frac{\hat{p}}{m}\right) \\ &= \frac{K^2}{K^2 - 1} \left(1 - \frac{1 - \mathbb{E}_{\mathbf{p}}[\mathbf{d}]}{m^2}\right) \\ &= \frac{K^2}{(K^2 - 1)m^2} \mathbb{E}_{\mathbf{p}}[\mathbf{d}] + \frac{K^2}{K^2 - 1} \frac{m^2 - 1}{m^2}. \end{aligned} \quad (\text{A.4})$$

Thus the UVR can be expressed as a linear transformation of the expected distance from the mode even when the PMF is multi-modal, whereas expression (A.3) is only valid for unimodal distributions.

Appendix B. Relationship between uncertainty statistics WVR, SDM, and CNV

Although expressions for the uncertainty statistics WVR (3.12), SDM (3.16) and CNV (3.25) give different values in general, there is a relationship between them for certain forms of PMF \mathbf{p} .

Consider the unimodal case where the modal probability is \hat{p} , and all other class probabilities are equal. This PMF can be written as

$$q_k = \begin{cases} \hat{p} & c_k = \hat{c}, \\ \frac{1-\hat{p}}{K-1} & \text{otherwise,} \end{cases} \quad k = 1, \dots, K. \quad (\text{B.1})$$

All binary class ($K = 2$) PMFs can be written in this form. For multi-class classification models where outputs are confident in one class, the predictive output PMFs approximately follow the form of (B.1).

We will now show that WVR, SDM and CNV uncertainty statistics are equivalent in such a scenario. Substituting (B.1) into (3.16) leads to





$$\begin{aligned} u_{\text{SDM}}(\mathbf{q}) &= 1 - \sqrt{\frac{(K - 1) \left(\hat{p} - \frac{1-\hat{p}}{K-1}\right)^2}{K - 1}} \\ &= 1 - \hat{p} + \frac{1-\hat{p}}{K-1} \\ &= 1 - \frac{K\hat{p} - 1}{K - 1} = u_{\text{WVR}}(\mathbf{q}). \end{aligned}$$

Now substituting (B.1) into (3.25) yields

$$\begin{aligned} u_{\text{CNV}}(\mathbf{q}) &= 1 - \sqrt{\frac{K \left(\hat{p}^2 + (K - 1) \left(\frac{1-\hat{p}}{K-1}\right)^2\right) - 1}{K - 1}} \\ &= 1 - \frac{1}{K - 1} \sqrt{(K - 1)(K\hat{p}^2 - 1) + K(1 - \hat{p})^2} \\ &= 1 - \frac{1}{K - 1} \sqrt{K^2\hat{p}^2 - 2K\hat{p} + 1} \\ &= 1 - \frac{1}{K - 1} \sqrt{(K\hat{p} - 1)^2} = u_{\text{WVR}}(\mathbf{q}). \end{aligned}$$

This also explains the equivalence in the binary class uncertainty expressions (3.13), (3.17), and (3.27).

ORCID iDs

Samuel Bilson  0009-0005-1705-2661
 Maurice Cox  0000-0002-6342-7840
 Anna Pustogvar  0009-0005-1729-8792
 Andrew Thompson  0009-0006-5572-0453

References

- [1] ESA Climate change initiative extension (CCI+) phase 1 – new essential climate variables (NEW ECVS) – high resolution land cover ECV (HR_LandCover_cci) – algorithm theoretical basis document (ATBD) (available at: https://climate.esa.int/media/documents/CCI_HRLC_Ph1-D2.2_ATBD_v4.0.pdf)
- [2] ESA Lakes_CCI+ - Phase 2. D2.2. Algorithm theoretical basis document (ATBD) (available at: https://climate.esa.int/media/documents/Lakes_CCI_D2-2_ATBD.pdf)
- [3] Skakun S *et al* 2022 Cloud Mask Intercomparison eXercise (CMIX): an evaluation of cloud masking algorithms for Landsat 8 and Sentinel-2 *Remote Sens. Environ.* **274** 112990
- [4] Buters J *et al* 2024 Automatic detection of airborne pollen: an overview *Aerobiologia* **40** 13–37
- [5] Erb S, Graf E, Zeder Y, Lionetti S, Berne A, Clot B, Lieberherr G, Tummon F, Wullschlegler P and Crouzy B 2024 Real-time pollen identification using holographic imaging and fluorescence measurements *Atmos. Meas. Tech.* **17** 441–51
- [6] Venton J, Harris P M, Sundar A, Smith N A and Aston P J 2021 Robustness of convolutional neural networks to physiological electrocardiogram noise *Phil. Trans. R. Soc. A* **379** 20200262
- [7] Torres-Soto J and Ashley E A 2020 Multi-task deep learning for cardiac rhythm detection in wearable devices *npj Digit. Med.* **3** 116

- [8] Bilson S, Thompson A, Tucker D and Pearce J 2024 A machine learning approach to automation and uncertainty evaluation for self-validating thermocouples *AIP Conf. Proc.* **3230** 090011
- [9] BIPM, IEC, IFCC 2012 International vocabulary of metrology – Basic and general concepts and associated terms (VIM) *Joint Committee for Guides in Metrology, JCGM 200:2012* (<https://doi.org/10.59161/JCGM200-2012>)
- [10] BIPM, IEC, IFCC 2023 Guide to the expression of uncertainty in measurement—Part 1: Introduction *Joint Committee for Guides in Metrology, JCGM GUM-1* (<https://doi.org/10.59161/JCGMGUM-1-2023>)
- [11] E M N for mathematics and statistics 2024 strategic research agenda 2023 - 2033 (available at: www.euramet.org/european-metrology-networks/mathmet/strategy/strategic-research-agenda)
- [12] Woolliams E et al 2025 A metrological framework for addressing uncertainty in satellite and *in situ* earth environmental observations submitted
- [13] Huyen C 2022 *Designing Machine Learning Systems* (O'Reilly Media, Inc)
- [14] Hoeting J A, Madigan D, Raftery A E and Volinsky C T 1999 Bayesian model averaging: a tutorial (with comments by M. Clyde, David Draper and El George and a rejoinder by the authors *Stat. Sci.* **14** 382–417
- [15] Hullermeier E and Waegeman W 2021 Aleatoric and epistemic uncertainty in machine learning: an introduction to concepts and methods *Mach. Learn.* **110** 457–506
- [16] Murphy K P 2012 *Machine Learning: a Probabilistic Perspective* (MIT Press)
- [17] Bilson S and Pustogvar A 2025 Uncertainty-aware Bayesian machine learning modelling of land cover classification (arXiv:2503.21510)
- [18] Gal Y and Ghahramani Z 2016 Dropout as a Bayesian approximation: representing model uncertainty in deep learning *Int. Conf. on Machine Learning* pp 1050–9
- [19] Kendall A and Gal Y 2017 What uncertainties do we need in Bayesian deep learning for computer vision? *NIPS'17: Proc. of the 31st Int Conf. on Neural Information Processing Systems* vol 30
- [20] Hastie T, Tibshirani R and Friedman J 2008 *The Elements of Statistical Learning. Data Mining, Inference and Prediction* 2nd edn (Springer)
- [21] BIPM, IEC, IFCC 2008 Evaluation of measurement data—guide to the expression of uncertainty in measurement *Joint Committee for Guides in Metrology, JCGM* vol 100
- [22] BIPM, IEC, IFCC 2008 Evaluation of measurement data—Supplement 1 to the 'Guide to the expression of uncertainty in measurement'—propagation of distributions using a Monte Carlo method *Joint Committee for Guides in Metrology, JCGM* vol 101
- [23] BIPM, IEC, IFCC 2011 Evaluation of measurement data—supplement 2 to the 'Guide to the expression of uncertainty in measurement'—extension to any number of output quantities *Joint Committee for Guides in Metrology, JCGM* vol 102
- [24] BIPM, IEC, IFCC 2020 Guide to the expression of uncertainty in measurement — Part 6: developing and using measurement models *Joint Committee for Guides in Metrology, JCGM GUM-6:2020* (<https://doi.org/10.59161/JCGMGUM-6-2020>)
- [25] BIPM, IEC, IFCC 2012 Evaluation of measurement data — the role of measurement uncertainty in conformity assessment *Joint Committee for Guides in Metrology, JCGM* vol 106
- [26] ESA *Sentinel-2 MPC – Level 2A Products Algorithm Theoretical Basis Document S2-PDGS-MPC-ATBD-L2A*
- [27] López-Puigdollers D, Mateo-García G and Gómez-Chova L 2021 Benchmarking deep learning models for cloud detection in landsat-8 and sentinel-2 images *Remote Sens.* **13** 992
- [28] Bettencourt da Silva R and Ellison S L R 2021 Eurachem/CITAC Guide: assessment of performance and uncertainty in qualitative chemical analysis (available at: www.eurachem.org/index.php/publications/guides/performance-and-uncertainty-in-qualitative-analysis)
- [29] Nordin G, Dybkaer R, Forsum U and Fuentes-Arderiu X 2018 Vocabulary on nominal property, examination and related concepts for clinical laboratory sciences (IFCC-IUPAC Recommendations 2017) *Pure Appl. Chem.* **90** 913–35
- [30] Stevens S S 1946 On the theory of scales of measurement *Science* **103** 677–80
- [31] White R 2011 The meaning of measurement in metrology *Accreditation Qual. Assur.* **16** 31–41
- [32] ISO 33406:2024 2024 *Approaches for the Production of Reference Materials with Qualitative Properties* (International Organisation for Standardization)
- [33] Agresti A 2002 *Categorical Data Analysis (Wiley Series in Probability and Statistics)* 2nd edn (Wiley)
- [34] Possolo A 2015 Simple guide for evaluating and expressing the uncertainty of NIST measurement results National Institute of Standards and Technology (<https://doi.org/10.6028/NIST.TN.1900>) *Technical Report* NIST TN 1900
- [35] Mari L, Narduzzi C, Nordin G and Trapmann S 2020 Foundations of uncertainty in evaluation of nominal properties *Measurement* **152** 107397
- [36] Evans M J and Rosenthal J S 2004 *Probability and Statistics: The Science of Uncertainty* (Macmillan)
- [37] 3534-1:2006 *Statistics — Vocabulary and symbols Part 1: General statistical terms and terms used in probability* 2006
- [38] Possolo A, Hibbert D B, Stohner J, Bodnar O and Meija J 2024 A brief guide to measurement uncertainty (IUPAC Technical Report) *Pure Appl. Chem.* **96** 113–34
- [39] Kvålseth T O 2011 Variation for categorical variables *International Encyclopedia of Statistical Science* ed M Lovric (Springer) pp 1642–5
- [40] Shirmohammadi S, Amiri M H and Osman H A 2024 Measurement methodology: uncertainty as a predictor of classification accuracy in machine learning-assisted measurements *IEEE Instrum. Meas. Mag.* **27** 37–45
- [41] Wilcox A R 1973 Indices of qualitative variation and political measurement *West. Polit. Q.* **26** 325–43
- [42] Bogaert P, Waldner F and Defourny P 2017 An information-based criterion to measure pixel-level thematic uncertainty in land cover classifications *Stoch. Environ. Res. Risk Assess.* **31** 2297–312
- [43] Rubia J 2024 Measures of variability for qualitative variables using the R software *Open J. Stat.* **14** 259–93
- [44] Kvålseth T O 1988 Measuring variation for nominal data *Bull. Psychonomic Soc.* **26** 433–6
- [45] Shannon C E 1948 A mathematical theory of communication *Bell Syst. Tech. J.* **27** 379–423
- [46] Mari L 2017 Toward a harmonized treatment of nominal properties in metrology *Metrologia* **54** 784–95
- [47] Possolo A 2014 Statistical models and computation to evaluate measurement uncertainty *Metrologia* **51** 228–36
- [48] Jaynes E T 1968 Prior Probabilities *IEEE Trans. Syst. Sci. Cybern.* **4** 227–41
- [49] Pal N and Bezdek J 1994 Measuring fuzzy uncertainty *IEEE Trans. Fuzzy Syst.* **2** 107–18
- [50] Low F, Michel U, Dech S and Conrad C 2013 Impact of feature selection on the accuracy and spatial uncertainty of per-field crop classification using support vector machines *ISPRS J. Photogramm. Remote Sens.* **85** 102–19
- [51] Glasziou P and Hilden J 1989 Test selection measures *Med. Decis. Making* **9** 133–41
- [52] Kvålseth T O 1995 Coefficients of variation for nominal and ordinal categorical data *Perceptual Motor Skills* **80** 843–7

- [53] Sipkens T A, Corbin J C, Grauer S J and Smallwood G J 2023 Tutorial: guide to error propagation for particle counting measurements *J. Aerosol Sci.* **167** 106091
- [54] Bohn T J and Vivoni E R 2019 MOD-LSP, MODIS-based parameters for hydrologic modeling of North American land cover change *Sci. Data* **6** 2
- [55] Wang L, Arora V K, Bartlett P, Chan E and Curasi S R 2023 Mapping of ESA's climate change initiative land cover data to plant functional types for use in the CLASSIC land model *Biogeosciences* **20** 2265–82
- [56] IPCC 2006 Guidelines for National Greenhouse Gas Inventories, The Intergovernmental Panel on Climate Change (IPCC) *Technical Report*
- [57] Skidmore A K et al 2021 Priority list of biodiversity metrics to observe from space *Nat. Ecol. Evol.* **5** 896–906
- [58] Kilcoyne A, Clement M, Moore C, Phillipps G P, Keane R, Woodget A, Potter S, Stefaniak A and Trippier B 2022 Living England: satellite based habitat classification. Technical User Guide, Natural England *Technical Report* NERR108
- [59] Marston C, Rowland C, O'Neil A and Morton R 2022 Land cover map 2021 (10m classified pixels, GB)
- [60] Friedl M A et al 2022 Medium spatial resolution mapping of global land cover and land cover change across multiple decades from landsat *Front. Remote Sens.* **3** 2022
- [61] Venter Z S and Sydenham M A K 2021 Continental-scale land cover mapping at 10 m resolution over Europe (ELC10) *Remote Sens.* **13** 2301
- [62] Brown C F et al 2022 Dynamic world, near real-time global 10 m land use land cover mapping *Sci. Data* **9** 4
- [63] Karra K, Kontgis C, Statman-Weil Z, Mazzariello J C, Mathis M and Brumby S P 2021 Global land use/land cover with Sentinel 2 and deep learning *2021 IEEE Int. Geoscience and Remote Sensing Symp. IGARSS (Brussels, Belgium)* (IEEE) pp 4704–7
- [64] Zanaga D et al 2022 ESA WorldCover 10 m 2021 v200
- [65] EEA 2020 Corine land cover 2018 (raster 100 m) (<https://doi.org/10.2909/960998c1-1870-4e82-8051-6485205ebbac>)
- [66] Moraes D, Campagnolo M L and Caetano M 2024 Training data in satellite image classification for land cover mapping: a review *Eur. J. Remote Sens.* **57** 414
- [67] Jospin L V, Laga H, Boussaid F, Buntine W and Bennamoun M 2022 Hands-on Bayesian neural networks-a tutorial for deep learning users *IEEE Comput. Intell. Mag.* **17** 29–48
- [68] Agrawal T 2021 *Hyperparameter Optimization in Machine Learning* 1st edn (Apress Berkeley) (<https://doi.org/10.1007/978-1-4842-6579-6>)
- [69] Snoek J, Larochelle H and Adams R P 2012 Practical Bayesian optimization of machine learning algorithms *Advances in Neural Information Processing Systems* vol 25, ed F Pereira, C Burges, L Bottou and K Weinberger (Curran Associates, Inc) (available at: <https://dl.acm.org/doi/10.5555/2999325.2999464>)
- [70] Tuo R, Wang W, Camps-Valls G, Ruiz F J R and 2022 Uncertainty quantification for Bayesian optimization *Proc. of The 25th Int. Conf. on Artificial Intelligence and Statistics (Proc. of Machine Learning Research)* vol 151, ed I Valera
- [71] Brooks S, Gelman A, Jones G and Meng X-L 2011 *Handbook of Markov Chain Monte Carlo* (CRC Press)
- [72] Gneiting T and Raftery A E 2007 Strictly proper scoring rules, prediction and estimation *J. Am. Stat. Assoc.* **102** 359–78
- [73] Ferrer L 2023 Analysis and comparison of classification metrics (arXiv:2209.05355)
- [74] Linz D, Gawalko M, Betz K, Hendriks J M, Lip G Y, Vinter N, Guo Y and Johnsen S 2024 Atrial fibrillation: epidemiology, screening and digital health *Lancet* **37** 100786
- [75] Bench C, Strodthoff N, Moulaeifard M, Aston P and Thompson A 2024 Towards trustworthy atrial fibrillation classification from wearables data: quantifying model uncertainty *Computing in Cardiology (https://doi.org/10.22489/CinC.2024.068)*
- [76] Wouters F et al 2025 Comparative evaluation of consumer wearable devices for atrial fibrillation detection: validation study *JMIR Formative Res.* **9** e65139
- [77] Fibricek: advanced heart rhythm monitoring for detection and management of atrial fibrillation (available at: www.fibricek.com) (Accessed 31 December 2024)
- [78] Charlton P H, Kyriacou P A, Mant J, Marozas V, Chowienzyk P and Alastruey J 2022 Wearable photoplethysmography for cardiovascular monitoring *Proc. IEEE* vol 110 pp 355–81
- [79] Strodthoff N, Wagner P, Schaeffter T and Samek W 2020 Deep learning for ECG analysis: benchmarks and insights from PTB-XL *IEEE J. Biomed. Health Inf.* **25** 1519–28
- [80] He T, Zhang Z, Zhang H, Zhang Z, Xie J and Li M 2019 Bag of tricks for image classification with convolutional neural networks *Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition* pp 558–67
- [81] Martin J and Elster C 2023 Aleatoric uncertainty for errors-in-variables models in deep regression *Neural Process. Lett.* **55** 4799–818
- [82] Morton R, Marston C, O'Neil A and Rowland C NERC Environmental Information Data Centre 2020 Land cover map 2017 (20m classified pixels, GB) (<https://doi.org/10.5285/f6f86b1a-af6d-4ed8-85af-21ee97ec5333>)
- [83] Morton R, Marston C, O'Neil A and Rowland C NERC Environmental Information Data Centre 2020 Land cover map 2018 (20m classified pixels, GB) (<https://doi.org/10.5285/b3dfc4c7-c9bd-4a02-bed8-46b2a41be04a>)
- [84] Morton R, Marston C, O'Neil A and Rowland C NERC Environmental Information Data Centre 2020 Land cover map 2019 (20m classified pixels, GB) (<https://doi.org/10.5285/643eb5a9-9707-4fbb-ae76-e8e53271d1a0>)
- [85] Morton R, Marston C, O'Neil A and Rowland C NERC EDS Environmental Information Data Centre 2021 Land cover map 2020 (10m classified pixels, GB) (<https://doi.org/10.5285/35c7d0e5-1121-4381-9940-75f7673c98f7>)
- [86] Torres-Soto J and Ashley E A Deepbeat (available at: <https://github.com/AshleyLab/deepbeat>)