



# Optimal sequential decision making with probabilistic digital twins




## Theoretical foundations

Christian Agrell<sup>1,2,3</sup>  · Kristina Rognlien Dahl<sup>2,3</sup> · Andreas Hafver<sup>1</sup>

Received: 19 August 2022 / Accepted: 20 February 2023

Published online: 24 March 2023

© The Author(s) 2023 

### Abstract

In this study, we present a formal definition of the probabilistic digital twin (PDT). Digital twins are emerging in many industries, typically consisting of simulation models and data associated with a specific physical system. In order to define *probabilistic* digital twins, we discuss how epistemic uncertainty can be treated using measure theory, by modelling epistemic information via  $\sigma$ -algebras. A gentle introduction to the necessary mathematical theory is provided throughout the paper, together with a number of examples to illustrate the core concepts. We then introduce the problem of optimal sequential decision making. That is, when the outcome of each decision may inform the next. We discuss how this problem may be solved theoretically, and the current limitations that prohibit most practical applications. As a numerically tractable alternative, we propose a generic approximate solution using deep reinforcement learning together with neural networks defined on sets. We illustrate the method on a practical problem, considering optimal information gathering for the estimation of a failure probability.

**Keywords** Probabilistic digital twin · Epistemic uncertainty · Sequential decision making · Partially observable Markov decision process · Deep reinforcement learning

## 1 Introduction

### 1.1 Probabilistic digital twins

The use of digital twins has emerged as one of the major technology trends the last couple of years. In essence, a digital twin (DT) is a digital representation of some physical system, including data from observations of the physical system, which can be used to perform forecasts, evaluate the consequences of potential actions, simulate possible future scenarios, and in general inform decision making without requiring interference with the physical

system. From a theoretical perspective, a digital twin may be regraded to consist of the following two components:

- A set of *assumptions* regarding the physical system (e.g. about the behaviour or relationships among system components and between the system and its environment), often given in the form of a physics-based numerical simulation model.
- A set of *information*, usually in the form of a set of observations, or records of the relevant actions taken within the system.

✉ Kristina Rognlien Dahl, kristina.r.dahl@bi.no; Christian Agrell, christian.agrell@dnv.com; Andreas Hafver, andreas.hafver@dnv.com | <sup>1</sup>Group Research and Development, DNV, Høvik, Norway. <sup>2</sup>Department of Economics, BI Norwegian Business School, Oslo, Norway. <sup>3</sup>Department of Mathematics, University of Oslo, Oslo, Norway.



In some cases, a digital twin may be desired for a system whose attributes and behaviours are not deterministic, but stochastic. For example, the degradation and failure of physical structures or machinery is typically described as stochastic processes. A system's performance may be impacted by weather or financial conditions, which also may be most appropriately modelled as stochastic. Sometimes the functioning of the system itself is stochastic, such as supply chain or production chains involving stochastic variation in demand and performance of various system components.

Even for systems or phenomena that are deterministic in principle, a model will never give a perfect rendering of reality. There will typically be uncertainty about the model's structure and parameters (i.e. epistemic uncertainty), and if consequences of actions can be critical, such uncertainties need to be captured and handled appropriately by the digital twin. In general, the system of interest will have both stochastic elements (aleatory uncertainty) and epistemic uncertainty.

If we want to apply digital twins to inform decisions in systems where the analysis of uncertainty and risk is important, certain properties are required:

1. The digital twin must capture uncertainties: *This could be done by using a probabilistic representation for uncertain system attributes.*
2. It should be possible to update the digital twin as new information becomes available: *This could be from new evidence in the form of data, or underlying assumptions about the system that have changed.*
3. For the digital twin to be informative in decision making, it should be possible to query the model sufficiently fast: *This could mean making use of surrogate models or emulators, which introduces additional uncertainties.*

These properties are paraphrased from [1], which provides a detailed discussion on the use of digital twins for on-line risk assessment.

A digital twin that complies with these properties is referred to as a *probabilistic digital twin* (PDT). Other formulations of PDTs have also been proposed in the recent years. In [2] the PDT is characterized through probabilistic machine learning from a physical system's *observables*, that are associated with properties such as parameters or state variables of a physics-based model. The PDT is first trained as a predictive model before it is exercised in a decision making context. A similar framework is outlined in [3], where the PDT takes the form of a dynamic decision network (an extension of a dynamic Bayesian network with decision nodes). The time evolution of the physical asset and the dynamical updating of the PDT is emphasized with

this approach. Other applications of PDTs can be found in [4], here also starting from a Bayesian network perspective, or [5] which instead deals with uncertainty modelling through polynomial chaos expansion. These methods may differ in their specific mathematical implementation, but the general idea is the same, with the main goal of combining physics-based and probabilistic modelling to support decision making under uncertainty.

In this paper we propose a mathematical framework for defining PDTs, starting from a measure-theoretic perspective. This is not in conflict with the current literature on PDTs, but we take a slightly more generic approach which will give us the vocabulary to properly separate between aleatory and epistemic uncertainty. This will let us analyze the effect of gathering information, which is needed e.g. for optimal experimental design, and which is also important for many safety-critical applications.

With respect to the three properties of PDTs described above, we will build on the Bayesian probabilistic framework which is a natural choice to satisfy Items 1 and 2.

A numerical model of a complex physical system can often be computationally expensive, for instance if it involves the numerical solution of nontrivial partial differential equations. In a probabilistic setting this is prohibitive, as a large number of evaluations (e.g. PDE solves) is needed for tasks involving uncertainty propagation, such as prediction and inference. Applications towards real-time decision making also sets natural restrictions with respect to the runtime of such queries. This is why property 3 is important, and why probabilistic models of complex physical phenomena often involve the use of approximate alternatives, usually obtained by "fitting" a computationally cheap model to the output of a few expensive model runs. These computationally cheap approximations are often referred to as *response surface models*, *surrogate models* or *emulators* in the literature.

Introducing this kind of approximation for computational efficiency also means that we introduce additional epistemic uncertainty into our modelling framework. By *epistemic* uncertainty we mean, in short, any form of uncertainty that can be reduced by gathering more information (to be discussed further later on). In our context, uncertainty may in principle be reduced by running the expensive numerical modes instead of the cheaper approximations.

## 1.2 Sequential decision making

Many interesting sequential decision making problems arise from the property that our knowledge about the system we operate changes as we learn about the outcomes. That is, each decision may affect the epistemic uncertainty which the next decision will be based upon.

For certain applications it is also important that decisions are robust with respect to what we do not know, i.e. with respect to epistemic uncertainty. Although we will not restrict the framework presented in this paper to any specific type of sequential decision making objectives, we will mainly focus on problems related to *optimal information gathering*. That is, where the decisions we consider are related to acquiring information (e.g. by running an experiment) in order to reduce the epistemic uncertainty with respect to some specified objective (e.g. estimating some quantity of interest). This is referred to as *optimal experimental design* in the statistics literature [6] and *active learning* within machine learning [7].

If we have a probabilistic model of some data generating process (for instance given by a PDT), we can simulate posterior distributions (of some quantity of interest) corresponding to what we might observe after a specific action. If we combine this with a utility function, given as a functional of the posterior distribution, we can try to optimize our decisions with respect to this utility. Methods of this sort are called Bayesian optimal designs [8]. Some examples are utility functions based on the Kullback–Leibler divergence between the prior and posterior [9, 10], mutual information [11], predictive uncertainty [6] or the cost of uncertainty [12]. For further details see e.g. the review in [8] and the references therein.

In the sequential setting, it is quite common to search for a myopic strategy, which means that we look one time step into the future and only consider the immediate consequence of a decision. This is of course generally suboptimal, compared to considering all future decisions and the observations they might produce. However, for non-trivial models, finding an exact solution for the multi-step alternative is usually computationally intractable. An alternative that has received increasing attention the recent years is to make use of reinforcement learning to find an approximate solution in the multi-step case. Some examples of this approach by the use of policy gradient reinforcement learning can be found in [13–15]. The examples presented in this paper (Sect. 5) show a different approach using Q-learning.

An example that we will consider in this paper, is the problem of optimal experimental design for structural reliability analysis. This is when information is acquired sequentially in order to evaluate the reliability of a physical operation, something which is a relevant application of a PDT. The quantity of interest here is the probability of structural failure (e.g. structural collapse due to loads caused by wind and waves). Assume that to estimate this probability, a large number of physics-based simulation (e.g. Finite Element Analysis) related to the relevant failure mechanism (e.g. fracture) is needed. Usually, only a small number of such simulations can be performed in practice.

The *experimental design problem* is then to decide which experiments (here simulation runs) to perform in order to build a surrogate model that can be used to estimate a failure probability with sufficient level of confidence. This is a problem that has received considerable attention (see e.g. [16–22]). These methods all make use of a myopic (one-step lookahead) criterion to determine the “optimal” experiment, as a multi-step or full dynamic programming formulation of the optimization problem becomes numerically infeasible. In [16] the authors consider the case where there are different types of experiments to choose from. Here, the myopic (one-step lookahead) assumption can still be justified, but if the different types of experiments are associated with different costs, then it can be difficult to apply in practice (e.g. if a feasible solution requires expensive experiments with delayed reward). We will revisit this example in this paper, but with a method that looks multiple steps ahead.

### 1.3 Contribution of this paper

In this paper we will review the mathematical framework of sequential decision making, and connect this to the definition of a PDT. Traditionally, there are two main solution strategies for solving discrete time sequential decision making problems: Maximum principles, and dynamic programming. We review these two solution methods, and conclude that the PDT framework is well suited for a dynamic programming approach. However, dynamic programming suffers from the curse of dimensionality, i.e. possible sequences of decisions and state realizations grow exponentially with the size of the state space. In fact, when doing numerical backward induction in dynamic programming, the objective function must be computed for each combination of values. This makes the method too computationally demanding to be applicable in practice for problems where the state space is large, see Agrell and Dahl [16] for a discussion of this. Hence, we are typically not able to solve a PDT sequential decision making problem in practice directly via dynamic programming.

As a generic solution to the problem of optimal sequential decision making we instead propose an alternative based on reinforcement learning. This means that when we consider the problem of finding an optimal decision policy, instead of truncating the theoretical optimal solution by e.g. looking only one step ahead, we try to approximate the optimal policy. This approximation can be done by using e.g. a neural network. Here we will frame the sequential decision making setup as a Markov decision process (MDP), in general as a *partially observable* MDP (POMDP), where a *state* is represented by the information available at any given time. This kind of state specification is often referred to as the *information state-space*. As

a generic approach to deep reinforcement learning using PDTs, we propose an approach using neural networks that operate on the information state-space directly.

**Main contributions.** In this paper we will:

- (i) Propose a mathematical framework for modelling epistemic uncertainty based on measure theory.
- (ii) Present a mathematical definition of the probabilistic digital twin (PDT). This is a mathematical framework for modelling physical systems with aleatory and epistemic uncertainty.
- (iii) Introduce the problem of sequential decision making, and illustrate how this problem can be solved (at least in theory) via maximum principle methods or the dynamic programming principle.
- (iv) Discuss the curse of dimensionality for these solution methods, and illustrate how the sequential decision making problem in the PDT can be viewed as a partially observable Markov decision process.
- (v) Explain how reinforcement learning (RL) can be applied to find approximate optimal strategies for sequential decision making in the PDT, and propose a generic approach using a deep sets architecture that enables RL directly on the information state-space. We end with a numerical example to illustrate this approach.

The paper is structured as follows: In Sect. 2 we introduce epistemic uncertainty and suggest modeling this via  $\sigma$ -algebras. In Sect. 3, we present the mathematical framework, as well as a formal definition, of a probabilistic digital twin (PDT), and discuss how such PDTs are used in practice. Then, in Sect. 4, we introduce the problem of stochastic sequential decision making. We discuss the traditional solution approaches, in particular dynamic programming which is theoretically a suitable approach for decision problems that can be modelled using a PDT. However, due to the curse of dimensionality, using the dynamic programming directly is typically not tractable. We therefore turn to reinforcement learning using function approximation as a practical alternative. In Sect. 5, we show how an approximate optimal strategy can be achieved using deep reinforcement learning, and we illustrate the approach with a numerical example. Finally, in Sect. 6 we conclude and sketch some future works in this direction.

## 2 A measure-theoretic treatment of epistemic uncertainty

In this section, we review the concepts of epistemic and aleatory uncertainty, and introduce a measure-theoretic framework for modelling epistemic uncertainty.

### 2.1 Motivation

In uncertainty quantification (UQ), it is common to consider two different kinds of uncertainty: Aleatory (stochastic) and epistemic (knowledge-based) uncertainty. We say that uncertainty is epistemic if we foresee the possibility of reducing it through gathering more or better information. For instance, uncertainty related to a parameter that has a fixed but unknown value is considered epistemic. Aleatory uncertainty, on the other hand, is the uncertainty which cannot (in the modeller's perspective) be affected by gathering information alone. Note that the characterization of aleatory and epistemic uncertainty has to depend on the modelling context. For instance, the result of a coin flip may be viewed as epistemic, if we imagine a physics-based model that could predict the outcome exactly (given all initial conditions, etc.). However, under most circumstances it is most natural to view a coin flip as aleatory, or that it contains both aleatory and epistemic uncertainty (e.g. if the bias of the coin is unknown). Der Kiureghian and Ditlevsen [23] provide a detailed discussion of the differences between aleatory and epistemic uncertainty.

In this paper, we have two main reasons for distinguishing between epistemic and aleatory uncertainty. First, we would like to make decisions that are robust with respect to epistemic uncertainty. Secondly, we are interested in studying the effect of gathering information. Modelling epistemic uncertainty is a natural way of doing this.

In the UQ literature, aleatory uncertainty is typically modelled via probability theory. However, epistemic uncertainty is represented in many different ways. For instance, Helton et al. [24] consider four different ways of modelling epistemic uncertainty: Interval analysis, possibility theory, evidence theory (Dempster-Shafer theory) and probability theory.

In this paper we take a measure-theoretic approach. This provides a framework that is relatively flexible with respect to the types of assumptions that underlie the epistemic uncertainty. As a motivating example, consider the following typical setup used in statistics:

#### **Example 2.1** (A parametric model)

Let  $\mathbf{X} = (Y, \theta)$  where  $Y$  is a random variable representing some stochastic phenomenon, and assume  $Y$  is modelled using a given probability distribution,  $P(Y|\theta)$ , that depends on a parameter  $\theta$  (e.g.  $Y \sim \mathcal{N}(\mu, \sigma)$  with  $\theta = (\mu, \sigma)$ ). Assume that we do not know the value of  $\theta$ , and we therefore consider  $\theta$  as a (purely) epistemic parameter. For some fixed value of  $\theta$ , the random variable  $Y$  is (purely) aleatory, but in general, as the true value of  $\theta$  is not known,  $Y$  is associated with both epistemic and aleatory uncertainty.

The model  $\mathbf{X}$  in Example 2.1 can be decoupled into an aleatory component  $Y|\theta$  and an epistemic component  $\theta$ . Any property of the aleatory uncertainty in  $\mathbf{X}$  is determined by  $P(Y|\theta)$ , and is therefore a function of  $\theta$ . For instance, the probability  $P(Y \in A|\theta)$  and the expectation  $E[f(Y)|\theta]$ , are both *functions of  $\theta$* . There are different ways in which we can choose to address the epistemic uncertainty in  $\theta$ . We could consider intervals, for instance the minimum and maximum of  $P(Y \in A|\theta)$  over any plausible value of  $\theta$ , or assign probabilities, or some other measure of belief, to the possible values  $\theta$  may take. However, in order for this to be well-defined mathematically, we need to put some requirements on  $A$ ,  $f$  and  $\theta$ . By using probability theory to represent the aleatory uncertainty, we implicitly assume that the set  $A$  and function  $f$  are *measurable*, and we will assume that the same holds for  $\theta$ . We will describe in detail what is meant by measurable in Sect. 2.2 below. Essentially, this is just a necessity for defining properties such as distance, volume or probability in the space where  $\theta$  resides.

In this paper we will rely on probability theory for handling both aleatory and epistemic uncertainty. This means that, along with the measurability requirement on  $\theta$ , we have the familiar setup for Bayesian inference:

**Example 2.2** (A parametric model—Inference and prediction)

If  $\theta$  from Example 2.1 is a random variable with distribution  $P(\theta)$ , then  $\mathbf{X} = (Y, \theta)$  denotes a complete probabilistic model (capturing both aleatory and epistemic uncertainty).  $\mathbf{X}$  is a random variable with distribution

$$P(\mathbf{X}) = P(Y|\theta)P(\theta).$$

Let  $I$  be some piece of information from which Bayesian inference is possible, i.e.  $P(\mathbf{X}|I)$  is well defined. We may then define the updated joint distribution

$$P_{\text{new}}(\mathbf{X}) = P(Y|\theta)P(\theta|I),$$

and the updated marginal (predictive) distribution for  $Y$  becomes

$$P_{\text{new}}(Y) = \int P(Y|\theta)dP(\theta|I).$$

Note that the distribution  $P_{\text{new}}(\mathbf{X})$  in Example 2.2 is obtained by only updating the belief with respect to epistemic uncertainty, and that

$$P_{\text{new}}(\mathbf{X}) \neq P(\mathbf{X}|I) = P(Y|I, \theta)P(\theta|I).$$

For instance, if  $I$  corresponds to an observation of  $Y$ , e.g.  $I = \{Y = y\}$ , then  $P(Y|I) = \delta(y)$ , the Dirac delta at  $y$ , whereas  $P(\theta|I)$  is the updated distribution for  $\theta$  having observed one realization of  $Y$ . In the following, we will

refer to the kind of Bayesian updating in Example 2.2 as *epistemic updating*.

This epistemic updating of the model considered in Examples 2.1 and 2.2 should be fairly intuitive, if

- 1 All epistemic uncertainty is represented by a single parameter  $\theta$ , and
- 2  $\theta$  is a familiar object like a number or a vector in  $\mathbb{R}^n$ .

But what can we say in a more general setting? It is common that epistemic uncertainty comes from lack of knowledge related to *functions*. This is the case with probabilistic emulators and surrogate models. The input to these functions may contain epistemic and/or aleatory uncertainty as well. Can we talk about isolating and modifying the epistemic uncertainty in such a model, without making reference to the specific details of how the model has been created? In the following we will show that with the measure-theoretic framework, we can still make use of a simple formulation like the one in Example 2.2.

## 2.2 The probability space

Let  $\mathbf{X}$  be a random variable containing both aleatory and epistemic uncertainty. In order to describe how  $\mathbf{X}$  can be treated like in Examples 2.1 and 2.2, but for the general setting, we will first recall some of the basic definitions from measure theory and measure-theoretic probability.

To say that  $\mathbf{X}$  is a *random variable*, means that  $\mathbf{X}$  is defined on some *measurable space*  $(\Omega, \mathcal{F})$ . Here,  $\Omega$  is a set, and if  $\mathbf{X}$  takes values in  $\mathbb{R}^n$  (or some other measurable space), then  $\mathbf{X}$  is a so-called *measurable function*,  $\mathbf{X} : \Omega \rightarrow \mathbb{R}^n$  (to be defined precisely later). Any randomness or uncertainty about  $\mathbf{X}$  is just a consequence of uncertainty regarding  $\omega \in \Omega$ . As an example,  $\mathbf{X}$  could relate to some 1-year extreme value, whose uncertainty comes from day to day fluctuations, or some fundamental stochastic phenomenon represented by  $\omega \in \Omega$ . Examples of natural sources of uncertainty are weather or human actions in large scale. Therefore, whether modeling weather, option prices, structural safety at sea or traffic networks, stochastic models should be used.

The probability of the event  $\{\mathbf{X} \in E\}$ , for some subset  $E \subset \mathbb{R}^n$ , is really the probability of  $\{\omega \in \mathbf{X}^{-1}(E)\}$ . Technically, we need to ensure that  $\{\omega \in \mathbf{X}^{-1}(E)\}$  is something that we can compute the probability of, and for this we need  $\mathcal{F}$ .  $\mathcal{F}$  is a collection of subsets of  $\Omega$ , and represents

all possible events (in the “ $\Omega$ -world”). When  $\mathcal{F}$  is a  $\sigma$ -algebra<sup>1</sup> the pair  $(\Omega, \mathcal{F})$  becomes a measurable space.

So, when we define  $\mathbf{X}$  as a random variable taking values in  $\mathbb{R}^n$ , this means that there exists some measurable space  $(\Omega, \mathcal{F})$ , such that any event  $\{\mathbf{X} \in E\}$  in the “ $\mathbb{R}^n$ -world” (which has its own  $\sigma$ -algebra) has a corresponding event  $\{\omega \in \mathbf{X}^{-1}(E)\} \in \mathcal{F}$  in the “ $\Omega$ -world”. It also means that we can define a probability measure on  $(\Omega, \mathcal{F})$  that gives us the probability of each event, but before we introduce any specific probability measure,  $\mathbf{X}$  will just be a measurable function<sup>2</sup>.

- We start with assuming that there exists some measurable space  $(\Omega, \mathcal{F})$  where  $\mathbf{X}$  is a measurable function.

The natural way to make  $\mathbf{X}$  into a random variable is then to introduce some probability measure<sup>3</sup>  $P$  on  $\mathcal{F}$ , giving us the probability space  $(\Omega, \mathcal{F}, P)$ .

- Given a probability measure  $P$  on  $(\Omega, \mathcal{F})$  we obtain the probability space  $(\Omega, \mathcal{F}, P)$  on which  $\mathbf{X}$  is defined as a random variable.

We have considered here, for familiarity, that  $\mathbf{X}$  takes values in  $\mathbb{R}^n$ . When no measure and  $\sigma$ -algebra is stated explicitly, one can assume that  $\mathbb{R}^n$  is endowed with the Lebesgue measure (which underlies the standard notion of length, area and volume, etc.) and the Borel  $\sigma$ -algebra (the smallest  $\sigma$ -algebra containing all open sets). Generally,  $\mathbf{X}$  can take values in any measurable space. For example,  $\mathbf{X}$  can map from  $\Omega$  to a space of functions. This is important in the study of stochastic processes.

### 2.3 The epistemic sub $\sigma$ -algebra $\mathcal{E}$

In the probability space  $(\Omega, \mathcal{F}, P)$ , recall that the  $\sigma$ -algebra  $\mathcal{F}$  contains all possible events. For any random variable  $\mathbf{X}$  defined on  $(\Omega, \mathcal{F}, P)$ , the knowledge that some event has occurred provides information about  $\mathbf{X}$ . This information may relate to  $\mathbf{X}$  in a way that it only affects epistemic uncertainty, only aleatory uncertainty, or both. We are

<sup>1</sup> This means that 1)  $\Omega \in \mathcal{F}$ , 2) if  $S \in \mathcal{F}$  then also the complement  $\Omega \setminus S \in \mathcal{F}$ , and 3) if  $S_1, S_2, \dots$  is a countable set of events then also the union  $S_1 \cup S_2 \cup \dots$  is in  $\mathcal{F}$ . Note that if these properties hold, many other types of events (e.g. countable intersections) will have to be included as a consequence.

<sup>2</sup> By definition, given two measure spaces  $(\Omega, \mathcal{F})$  and  $(\mathbb{X}, \mathcal{X})$ , the function  $\mathbf{X} : \Omega \rightarrow \mathbb{X}$  is measurable if and only if  $\mathbf{X}^{-1}(A) \in \mathcal{F}$  for each  $A \in \mathcal{X}$ .

<sup>3</sup> A function  $P : \mathcal{F} \rightarrow [0, 1]$  such that 1)  $P(\Omega) = 1$  and 2)  $P(\cup E_i) = \sum P(E_i)$  for any countable collection of pairwise disjoint events  $E_i$ .

interested in specifying the events  $e \in \mathcal{F}$  that are associated with epistemic information alone. It is the probability of these events we want to update as new information is obtained. The collection  $\mathcal{E}$  of such sets is itself a  $\sigma$ -algebra, and we say that  $\mathcal{E} \subseteq \mathcal{F}$  is the sub  $\sigma$ -algebra of  $\mathcal{F}$  representing epistemic information.

We illustrate this in the following examples. In Example 2.3, we consider the simplest possible scenario represented by the flip of a biased coin, and in Example 2.4 a familiar scenario from uncertainty quantification involving uncertainty with respect to functions.

#### Example 2.3 (Coin flip)

Define  $\mathbf{X} = (Y, \theta)$  as in Example 2.1, and let  $Y \in \{0, 1\}$  denote the outcome of a coin flip where “heads” is represented by  $Y = 0$  and “tails” by  $Y = 1$ . Assume that  $P(Y = 0) = \theta$  for some fixed but unknown  $\theta \in [0, 1]$ . For simplicity we assume that  $\theta$  can only take two values,  $\theta \in \{\theta_1, \theta_2\}$  (e.g. there are two coins but we do not know which one is being used).

Then,  $\Omega = \{0, 1\} \times \{\theta_1, \theta_2\}$ ,  $\mathcal{F} = 2^\Omega$  and  $\mathcal{E} = \{\emptyset, \Omega, \{(0, \theta_1), (1, \theta_1)\}, \{(0, \theta_2), (1, \theta_2)\}\}$ .

#### Example 2.4 (UQ)

Let  $\mathbf{X} = (\mathbf{x}, \mathbf{y})$  where  $\mathbf{x}$  is an aleatory random variable, and  $\mathbf{y}$  is the result of a fixed but unknown function applied to  $\mathbf{x}$ . We let  $\mathbf{y} = \hat{f}(\mathbf{x})$  where  $\hat{f}$  is a function-valued epistemic random variable.

If  $\mathbf{x}$  is defined on a probability space  $(\Omega_{\mathbf{x}}, \mathcal{F}_{\mathbf{x}}, P_{\mathbf{x}})$  and  $\hat{f}$  is a stochastic process defined on  $(\Omega_f, \mathcal{F}_f, P_f)$ , then  $(\Omega, \mathcal{F}, P)$  can be defined as the product of the two spaces and  $\mathcal{E}$  as the projection  $\mathcal{E} = \{\Omega_{\mathbf{x}} \times A \mid A \in \mathcal{F}_f\}$ .

In the following, we assume that the epistemic sub  $\sigma$ -algebra  $\mathcal{E}$  has been identified.

Given a random variable  $\mathbf{X}$ , we say that  $\mathbf{X}$  is  $\mathcal{E}$ -measurable if  $\mathbf{X}$  is measurable as a function defined on  $(\Omega, \mathcal{E})$ . We say that  $\mathbf{X}$  is independent of  $\mathcal{E}$ , if the conditional probability  $P(\mathbf{X}|e)$  is equal to  $P(\mathbf{X})$  for any event  $e \in \mathcal{E}$ . With our definition of  $\mathcal{E}$ , we then have for any random variable  $\mathbf{X}$  on  $(\Omega, \mathcal{F}, P)$  that

- $\mathbf{X}$  is purely epistemic if and only if  $\mathbf{X}$  is  $\mathcal{E}$ -measurable,
- $\mathbf{X}$  is purely aleatory if and only if  $\mathbf{X}$  is independent of  $\mathcal{E}$ .

### 2.4 Epistemic conditioning

Let  $\mathbf{X}$  be a random variable on  $(\Omega, \mathcal{F}, P)$  that may contain both epistemic and aleatory uncertainty, and assume that the epistemic sub  $\sigma$ -algebra  $\mathcal{E}$  is given. By epistemic conditioning, we want to update the epistemic part of the uncertainty in  $\mathbf{X}$  using some set of information  $I$ . In Example 2.3 this means updating the probabilities  $P(\theta = \theta_1)$  and

$P(\theta = \theta_2)$ , and in Example 2.4 this means updating  $P_f$ . In order to achieve this in the general setting, we first need a way to decouple epistemic and aleatory uncertainty. This can actually be made fairly intuitive, if we rely on the following assumption:

**Assumption 2.5** There exists a random variable  $\theta : \Omega \rightarrow \Theta$  that generates<sup>4</sup>  $\mathcal{E}$ .

If this generator  $\theta$  exists, then for any fixed value  $\theta \in \Theta$ , we have that  $\mathbf{X}|\theta$  is independent of  $\mathcal{E}$ . Hence  $\mathbf{X}|\theta$  is purely aleatory and  $\theta$  is purely epistemic.

We will call  $\theta$  the *epistemic generator*, and we can interpret  $\theta$  as a signal that reveals all epistemic information when known. That is, if  $\theta$  could be observed, then knowing the value of  $\theta$  would remove all epistemic uncertainty from our model. As it turns out, under fairly mild conditions one can always assume existence of this generator. One sufficient condition is that  $(\Omega, \mathcal{F}, P)$  is a *standard* probability space, and then the statement holds up to sets of measure zero. This is a technical requirement to avoid pathological cases, and does not provide any new intuition that we see immediately useful, so we postpone further explanation to Appendix A.

**Example 2.6** (Coin flip—epistemic generator)

In the coin flip example, the variable  $\theta \in \{\theta_1, \theta_2\}$  which generates  $\mathcal{E}$  is already specified.

**Example 2.7** (UQ – epistemic generator)

In this example, when  $(\Omega, \mathcal{F}, P)$  is the product of an aleatory space  $(\Omega_{\mathbf{x}}, \mathcal{F}_{\mathbf{x}}, P_{\mathbf{x}})$  and an epistemic space  $(\Omega_f, \mathcal{F}_f, P_f)$ , we could let  $\theta : \Omega = \Omega_{\mathbf{x}} \times \Omega_f \rightarrow \Omega_f$  be the projection  $\theta(\omega_{\mathbf{x}}, \omega_f) = \omega_f$ .

Alternatively, given only the space  $(\Omega, \mathcal{F}, P)$  where both  $\mathbf{x}$  and  $\hat{f}$  are defined, assume that  $\hat{f}$  is a Gaussian process (or some other stochastic process for which the Karhunen-Loève theorem holds). Then there exists a sequence of deterministic functions  $\phi_1, \phi_2, \dots$  and an infinite-dimensional variable  $\theta = (\theta_1, \theta_2, \dots)$  such that  $\hat{f}(\mathbf{x}) = \sum_{i=1}^{\infty} \theta_i \phi_i(\mathbf{x})$ , and we can let  $\mathcal{E}$  be generated by  $\theta$ .

The decoupling of epistemic and aleatory uncertainty is then obtained by considering the joint variable  $(\mathbf{X}, \theta)$  instead of  $\mathbf{X}$  alone, because

$$P(\mathbf{X}, \theta) = P(\mathbf{X}|\theta)P(\theta). \quad (1)$$

From (1) we see how the probability measure  $P$  becomes the product of the epistemic probability  $P(\theta)$  and the aleatory probability  $P(\mathbf{X}|\theta)$  when applied to  $(\mathbf{X}, \theta)$ .

Given new information,  $I$ , we will update our beliefs about  $\theta$ ,  $P(\theta) \rightarrow P(\theta|I)$ , and we define the epistemic conditioning as follows:

$$P_{\text{new}}(\mathbf{X}, \theta) = P(\mathbf{X}|\theta)P(\theta|I).$$

## 2.5 Two types of assumptions

Consider the probability space  $(\Omega, \mathcal{F}, P)$ , with epistemic sub  $\sigma$ -algebra  $\mathcal{E}$ . Here  $\mathcal{E}$  represents *epistemic information*, which is the information associated with *assumptions*. In other words, an epistemic event  $e \in \mathcal{E}$  represents an assumption. In fact, given a class of assumptions, the following Remark 2.8, shows why  $\sigma$ -algebras are appropriate structures.

**Remark 2.8** Let  $\mathcal{E}$  be a collection of assumptions. If  $e \in \mathcal{E}$ , this means that it is possible to assume that  $e$  is true. If it is also possible to assume that that  $e$  is false, then  $\bar{e} \in \mathcal{E}$  as well. It may then also be natural to require that  $e_1, e_2 \in \mathcal{E} \Rightarrow e_1 \cap e_2 \in \mathcal{E}$ , and so on. These are the defining properties of a  $\sigma$ -algebra.

For any random variable  $\mathbf{X}$  defined on  $(\Omega, \mathcal{F}, P)$ , when  $\mathcal{E}$  is a sub  $\sigma$ -algebra of  $\mathcal{F}$ ,  $\mathbf{X}|e$  for  $e \in \mathcal{E}$  is well defined, and represents the random variable under the assumption  $e$ . In particular, given any fixed *epistemic event*  $e \in \mathcal{E}$  we have a corresponding *aleatory distribution*  $P(\mathbf{X}|e)$  over  $\mathbf{X}$ , and the conditional  $P(\mathbf{X}|\mathcal{E})$  is the random measure corresponding to  $P(\mathbf{X}|e)$  when  $e$  is a random epistemic event in  $\mathcal{E}$ . Here, the global probability measure  $P$  when applied to  $e$ ,  $P(e)$ , is the belief that  $e$  is true. In Sect. 2.4 we discussed updating the part of  $P$  associated with epistemic uncertainty. We also introduced the epistemic generator  $\theta$  in order to associate the event  $e$  with an outcome  $\theta(e)$ , and make use of  $P(\mathbf{X}|\theta)$  in place of  $P(\mathbf{X}|\mathcal{E})$ . This provides a more intuitive interpretation of the assumptions that are *measurable*, i.e. those whose belief we may specify through  $P$ .

Of course, the measure  $P$  is also based on assumptions. For instance, if we in Example 2.1 assume that  $Y$  follows a *normal* distribution. One could in principle specify a (measurable) space of probability distributions, from which the normal distribution is one example. Otherwise, we view the normality assumption as a *structural* assumption related to the probabilistic model for  $\mathbf{X}$ , i.e. the measure  $P$ . These kinds of assumptions cannot be treated the same way as assumptions related to measurable events. For instance, the consequence of the complement assumption “ $Y$  does not follow a normal distribution” is not well defined.

<sup>4</sup> There exists some measurable space  $(\Theta, \mathcal{T})$  and a  $\mathcal{F}$ -measurable function  $\theta : \Omega \rightarrow \Theta$  such that  $\mathcal{E} = \sigma(\theta)$ , that is the smallest  $\sigma$ -algebra containing all of the sets  $\theta^{-1}(T)$  for  $T \in \mathcal{T}$ .

In order to avoid any confusion, we split the assumptions into two types:

- 1 The measurable assumptions represented by the  $\sigma$ -algebra  $\mathcal{E}$ , and
- 2 the set  $M$  of structural assumptions underlying the probability measure  $P$ .

This motivates the following definition.

**Definition 2.9** (Structural assumptions) We let  $M$  denote the set of structural assumptions that defines a probability measure on  $(\Omega, \mathcal{F})$ , which we may write  $P_M(\cdot)$  or  $P(\cdot | M)$ .

We may also refer to  $M$  as the *non-measurable* assumptions, to emphasize that  $M$  contains all the assumptions not covered by  $\mathcal{E}$ . When there is no risk of confusion we will also suppress the dependency on  $M$  and just write  $P(\cdot)$ . Stating the set  $M$  explicitly is typically only relevant for scenarios where we consider changes being made to the actual system that is being modelled, or for evaluating different candidate models, e.g. through the marginal likelihood  $P(I|M)$ . In practice one would also state  $M$  so that decision makers can determine their level of trust in the probabilistic model, and the appropriate level of caution when applying the model.

As we will see in the upcoming section, making changes to  $M$  and making changes to how  $P_M$  acts on events in  $\mathcal{E}$  are the two main ways in which we update a *probabilistic digital twin*.

### 3 The Probabilistic Digital Twin

The object that we will call *probabilistic digital twin*, PDT for short, is a probabilistic model of a physical system. It is essentially a (possibly degenerate) probability distribution of a vector  $\mathbf{X}$ , representing the relevant attributes of the system, but where we in addition require the specification of epistemic uncertainty (assumptions) and how this uncertainty may be updated given new information.

Before presenting the formal definition of a probabilistic digital twin, we start with an example showing why the identification of epistemic uncertainty is important.

#### 3.1 Why distinguish between aleatory and epistemic uncertainty?

The decoupling of epistemic and aleatory uncertainty (as described in Sect. 2.4) is central in the PDT framework. There are two good reasons for doing this:

- 1 We want to make decisions that are *robust with respect to epistemic uncertainty*.
- 2 We want to study the effect of gathering information.

Item 1 relates to the observation that decision theoretic approaches based on expectation may not be robust. That is, if we marginalize out the epistemic uncertainty (and considering only  $E_\theta[P(\mathbf{X}|\theta)] = \int P(\mathbf{X}|\theta)dP_\theta$ ). We give two examples of this below, see Examples 3.1 and 3.2. Item 2 means that by considering the effect of information on epistemic uncertainty, we can evaluate the value of gathering information. This is discussed in further detail in Sect. 4.6.

In the context of PDTs, it is essential that the sources of epistemic uncertainty are understood. First of all, because we need to understand how it may be reduced (as its definition is a reducible uncertainty), but also if we want to understand the criticality with respect to assumptions made with imperfect or lack of knowledge. Within risk analysis of engineering systems, which is a relevant application area for PDTs, this is important. This is discussed from a general risk perspective in [25, 26] and specifically for PDTs in [1].

#### Example 3.1 (Coin flip – robust decisions)

Continuing from the coin flip example (see Example 2.3), we let  $\theta_1 = 0.5$  and  $\theta_2 = 0.99$ . Assume that you are given the option to guess the outcome of  $\mathbf{X}$ . If you guess correct, you collect a reward of  $R = 10^6$ , otherwise you have to pay  $L = 10^6$ . A priori your belief about the bias of the coin is that  $P(\theta = 0.5) = P(\theta = 0.99) = 0.5$ . If you consider betting on  $\mathbf{X} = 0$ , then the expected return, obtained by marginalizing over  $\theta$ , becomes  $P(\theta = 0.5)(0.5R - 0.5L) + P(\theta = 0.99)(0.99R - 0.01L) = 490.000$

This is a scenario where decisions supported by taking the expectation with respect to epistemic uncertainty is not robust, as we believe that  $\theta = 0.5$  and  $\theta = 0.99$  are equally likely, and if  $\theta = 0.5$  we will lose  $10^6 50\%$  of the time by betting on  $\mathbf{X} = 0$ .

#### Example 3.2 (UQ – robust decisions)

This example is a continuation of Examples 2.4 and 2.7.

In structural reliability analysis, we are dealing with an unknown function  $g$  with the property that the event  $\{y = g(\mathbf{x}) < 0\}$  corresponds to failure. When  $g$  is represented by a random function  $\hat{g}$  with epistemic uncertainty, the failure probability is also uncertain. Or in other words, if  $\hat{g}$  is epistemic then  $\hat{g}$  is a function of the generator  $\theta$ . Hence, the failure probability is a function of  $\theta$ . We want to make use of a conservative estimate of the failure probability, i.e. use a conservative value of  $\theta$ .  $P(\theta)$  tells us how conservative a given value of  $\theta$  is.



### 3.2 The attributes $\mathbf{X}$

To define a PDT, we start by considering a vector  $\mathbf{X}$  consisting of the *attributes* of some system. This means that  $\mathbf{X}$  is a representation of the physical object or asset that we are interested in. In general,  $\mathbf{X}$  describes the physical *system*. In addition,  $\mathbf{X}$  must contain attributes related to any type of information that we want to make use of. For instance, if the information consists of observations, the relevant observable quantities, as well as attributes related to measurement errors or noise, may be included in  $\mathbf{X}$ . In general, we will think of a *model* of a system as a set of assumptions that describes how the components of  $\mathbf{X}$  are related or behave. The canonical example here is where some physical quantity is inferred from observations including errors and noise, in which case a model of the physical quantity (physical system) is connected with a model of the data generating process (observational system). We are interested in modelling dependencies with associated uncertainty related to the components of  $\mathbf{X}$ , and treat  $\mathbf{X}$  as a random variable.

The attributes  $\mathbf{X}$  characterise the state of the system and the processes that the PDT represents.  $\mathbf{X}$  may for instance include:

- *System parameters* representing quantities that have a fixed, but possibly uncertain, value. For instance, these parameters may be related to the system configuration.
- *System variables* that may vary in time, and whose value may be uncertain.
- *System events* i.e. the occurrence of defined state transitions.

In risk analysis, one is often concerned with risk measures given as quantified properties of  $\mathbf{X}$ , usually in terms of expectations. For instance, if  $\mathbf{X}$  contains some extreme value (e.g. the 100-year wave) or some specified event of failure (using a binary variable), the expectations of these may be compared against risk acceptance criteria to determine compliance.

### 3.3 The PDT definition

Based on the concepts introduced so far, we define the PDT as follows:

**Definition 3.3** (Probabilistic Digital Twin) A Probabilistic Digital Twin (PDT) is a triplet  $(\mathbf{X}, A, I)$ , where  $\mathbf{X}$  is a vector of attributes of a system,  $A$  contains the *assumptions* needed to specify a probabilistic model, and  $I$  contains *information* regarding actions and observations:

- $A = ((\Omega, \mathcal{F}), \mathcal{E}, M)$ , where  $(\Omega, \mathcal{F})$  is a measure space where  $\mathbf{X}$  is measurable, and  $\mathcal{E}$  is the sub  $\sigma$ -algebra representing epistemic information.  $M$  contains the structural assumptions that define a probability measure  $P_M$  on  $(\Omega, \mathcal{F})$ .
- $I$  is a set consisting of events of the form  $(d, o)$ , where  $d$  encodes a description of the conditions under which the observation  $o$  was made, and where the likelihood  $P(o|\mathbf{X}, d)$  is well defined. For brevity, we will write this likelihood as  $P(I|\mathbf{X})$  when  $I$  contains multiple events of this sort.

When  $M$  is understood, and there is no risk on confusion, we will drop stating the dependency on  $M$  explicitly and just refer to the probability space  $(\Omega, \mathcal{F}, P)$ .

It is important to note that consistency between  $I$  and  $P(\mathbf{X})$  is required. That is, when using the probabilistic model for  $\mathbf{X}$ , it should be possible to simulate the type of observations given by  $I$ . In this case the likelihood  $P(I|\mathbf{X})$  is well defined, and the epistemic updating of  $\mathbf{X}$  can be obtained from Bayes' theorem.

Finally, we note that with this setup the information  $I$  may contain observations made under different conditions than what is currently specified through  $M$ . The information  $I$  is generally defined as a set of events, given as pairs  $(d, o)$ , where  $d$  encodes the relevant action leading to observing  $o$ , as well as a description of the conditions under which  $o$  was observed. Here  $d$  may relate to modifications of the structural assumptions  $M$ , for instance if the causal relationships that describe the model of  $\mathbf{X}$  under observation of  $o$  is not the same as what is currently represented by  $M$ . This is the scenario when we perform controlled experiments. Alternatively,  $(d, o)$  may represent a passive observation, e.g.  $d = \text{'measurement taken from sensor 1 at time 01:02:03'}$  and  $o = 1.7$  mm. We illustrate this in the following example.

#### Example 3.4 (Parametric regression)

Let  $(x_1, x_2)$  denote two physical quantities where  $x_2$  depends on  $x_1$ , and let  $(y, \varepsilon)$  represent an observable quantity where  $y$  corresponds to observing  $x_2$  together with additive noise  $\varepsilon$ . Set  $\mathbf{X} = (x_1, x_2, y, \varepsilon)$ .

We define a model  $M$  corresponding to  $x_1 \sim p_{x_1}(x_1|\theta_1)$ ,  $x_2 = f(x_1, \theta_2)$ ,  $y = x_2 + \varepsilon$  and  $\varepsilon \sim p_\varepsilon$ , where  $p_{x_1}$  is a probability density depending on the parameter  $\theta_1$  and  $f(\cdot, \theta_2)$  is a deterministic function depending on the parameter  $\theta_2$ .  $\theta_1$  and  $\theta_2$  are epistemic parameters for which we define a joint density  $p_\theta$ .

Assume that  $I = \{(d^{(1)}, o^{(1)}), \dots, (d^{(n)}, o^{(n)})\}$  is a set of controlled experiments, where  $d^{(i)} = (\text{set } x_1 = x_1^{(i)})$  and  $o^{(i)}$  is a corresponding observation of  $y | (x_1 = x_1^{(i)}, \varepsilon = \varepsilon^{(i)})$  for a selected set of inputs  $x_1^{(1)}, \dots, x_1^{(n)}$  and unknown i.i.d.




Fig. 1 A standard regression model as a PDT

$\epsilon^{(i)} \sim p_\epsilon$ . In this scenario, regression is performed by updating the distribution  $p_\theta$  to agree with the observations:

$$p_\theta(\theta|I) = p_\theta(\theta_1|\theta_2)p_\theta(\theta_2|I) = \frac{1}{Z} \left( \prod_{i=1}^n p_\epsilon(o^{(i)} - f(x_1^{(i)}, \theta_2)) \right) p_\theta(\theta), \tag{2}$$

where  $Z$  is a constant ensuring that the updated density integrates to one.

If instead  $I$  corresponds to direct observations,  $d^{(i)} = (\text{observe } y^{(i)})$ ,  $o^{(i)} = y^{(i)}$ , then  $p_\theta(\theta|I)$  corresponds to using  $x_1$  instead of  $x_1^{(i)}$  and multiplying with  $p_{x_1}(x_1|\theta_1)$  in (2).

Note that the scenario with controlled experiments in Example 3.4 corresponds to a different model than the one in Fig. 1. This is a familiar scenario in the study of causal inference, where actively setting the value of  $x_1$  is the do-operator (see Pearl [27]) which breaks the link between  $x_1$  and  $x_2$ .

### 3.4 Corroded pipeline example

To give a concrete example of a system where the PDT framework is relevant, we consider the following model from Agrell and Dahl [16]. This is based on a probabilistic structural reliability model for assessment of corroded pipelines, which is the basis for the recommended practice DNV GL RP-F101 [28]. It is a model of a physical failure mechanism called pipeline burst, which may occur when the pipeline’s ability to withstand the high internal pressure has been reduced as a consequence of corrosion. We will describe just a general overview of this model, and refer to (Example 4 in Agrell and Dahl [16]) for specific details. Later, in Sect. 5.6, we will revisit this example and make use of reinforcement learning to search for an optimal way of updating the PDT.

A graphical representation of the model is shown in Fig. 2. The steel thickness  $t$  and the diameter  $D$  represent the **pipeline** geometry, and the material is represented by the ultimate tensile strength  $s$ . The pipeline has a **defect**, in the form of a rectangle with depth  $d$  and length  $l$ . For some pipeline  $(D, t, s)$  with defect  $(d, l)$ , the pipeline’s pressure resistance **capacity** (the maximum differential pressure the pipeline can withstand before bursting) can be determined. By running a Finite




Fig. 2 Graphical representation of the corroded pipeline structural reliability model. The shaded nodes  $d$ ,  $p_{FE}$  and  $\mu_m$  have associated epistemic uncertainty

Element simulation of the physical mode of failure, we obtain the theoretical capacity  $p_{FE}$ . We then model the true capacity of the pipeline as a function of  $p_{FE}$  and the **model discrepancy**  $X_m$  (representing the difference between true and theoretical capacity). We assume that  $X_m$  is independent of the type of pipeline and defect, and also that  $\sigma_m$  is fixed, where it is only the mean  $\mu_m$  that is inferred from experiments. The limit state equation representing the transition from a safe to failed state is expressed as  $g = p_c - p_d$ , where  $p_d$  is the pressure **load**. The probability of failure is defined as  $P(g \leq 0)$ .

If we let  $\mathbf{X}$  be the random vector containing all of the nodes in Fig. 2, then  $\mathbf{X}$  represents a probabilistic model of the physical system. In this example, we want to model some of the uncertainty related to the defect size, the model uncertainty, and the capacity as epistemic. We assume that the defect depth  $d$  has a fixed but unknown value, that can be inferred through observations that include noise. Similarly, the model uncertainty  $X_m$  can be determined from experiments. Uncertainty with respect to  $p_{FE}$  comes from the fact that evaluating the true value of  $p_{FE}|(D, t, s, d, l)$  involves a time-consuming numerical computation. Hence,  $p_{FE}$  can only be known for a finite, and relatively small set of input combinations. We can let  $\hat{p}_{FE}$  denote a stochastic process that models our uncertainty about  $p_{FE}$ . To construct a PDT from  $\mathbf{X}$  we will let  $\hat{p}_{FE}$  take the place of  $p_{FE}$ , and specify that  $d, \mu_m$  and  $\hat{p}_{FE}$  are epistemic, i.e.  $\mathcal{E} = \sigma(d, \mu_m, \hat{p}_{FE})$ .

If we want a way to update the epistemic uncertainty based on observations, we also need to specify the relevant data generating process. In this example, we assume that there are three different ways of collecting data:

- 1 **Defect measurement:** We assume that noise perturbed observations of the relative depth,  $d/t + \varepsilon$ , can be made.
- 2 **Computer experiment:** Evaluate  $p_{FE}$  at some selected input  $(D, t, s, d, l)$ .
- 3 **Lab experiment:** Obtain one observation of  $X_m$ .


As the defect measurements require the specification of an additional random variable, we have to include  $\varepsilon$  or  $(d/t)_{\text{obs}} = d/t + \varepsilon$  in  $\mathbf{X}$  as part of the complete probabilistic model. This would then define a PDT where epistemic updating is possible.

The physical system that the PDT represents in this example is rarely viewed in isolation. For instance, the random variables representing the pipeline geometry and material are the result of uncertainty or variations in how the pipeline has been manufactured, installed and operated. And the size of the defect is the result of a chemical process, where scientific models are available. It could therefore be natural to view the PDT from this example as a component of a bigger PDT, where probabilistic models of the manufacturing, operating conditions and corrosion process, etc. are connected. This form of modularity is often emphasized in the discussion of digital twins, and likewise for the kind of Bayesian network type of models as considered in this example.

## 4 Sequential decision making

We now consider how the PDT framework may be adopted in real-world applications. As with any statistical model of this form, the relevant type of applications are related to prediction and inference. Since the PDT is supposed to provide a one-to-one correspondence (including uncertainty) with a real physical system, we are interested in using the PDT to understand the consequences of actions that we have the option to make. In particular, we will consider the discrete sequential decision making scenario, where we get the opportunity to make a decision, receive information related to the consequences of this decision, and use this to inform the next decision, and so on.

In this kind of scenario, we want to use the PDT to determine an *action* or *policy* for how to act optimally (with respect to some case-specific criterion). By a *policy* here we mean the instructions for how to select among multiple actions given the information available at each discrete time step. We describe this in more detail in Sect. 4.3 where we discuss how the PDT is used for *planning*. When we make use of the PDT in this way, we consider the PDT as a “mental model” of the real physical system, which an agent uses to evaluate the potential consequences of actions. The agent then decides on some action to make, observes the



**Fig. 3** A PDT as a mental model of an agent taking actions in the real world. As new experience is gained, the PDT may be updated by changing the structural assumptions  $M$  that defined the probability measure  $P$ , or updating belief with respect to epistemic events through conditioning on the new set of information  $I$ . The changes in structural assumptions and epistemic information are represented by  $\Delta M$  and  $\Delta I$  respectively. As part of the planning process, the PDT may simulate possible scenarios as indicated by the inner circle

outcome, and updates its beliefs about the true system, as illustrated in Fig. 3.

Whether the agent applies a policy or just an action (the first in the policy) before collecting information and updating the probabilistic model depends on the type of application at hand. In general it is better to update the model as often as possible, preferably between each action, but the actual computational time needed to perform this updating might make it impossible to achieve in practice.

### 4.1 Mathematical framework of sequential decision making

In this section, we briefly recap the mathematical framework of stochastic, sequential decision making in discrete time. We first recall the general framework, and in the following Sect. 4.2, we show how this relates to our definition of a PDT.

Let  $t = 0, 1, 2, \dots, N - 1$  and consider a discrete time system where the state of the system,  $\{x_t\}_{t \geq 0}$  is given by

$$x_{t+1} = f_t(x_t, u_t, w_t), \quad t = 0, 1, 2, \dots, N - 1. \quad (3)$$

Here,  $x_t$  is the *state of the system* at time  $t$ ,  $u_t$  is a *control* and  $w_t$  is a *noise*, or random parameter at time  $t$ . Note that the control,  $u_t$ , is a decision which can be made by an *agent*

(the controller) at time  $t$ . This control is to be chosen from a set of admissible controls  $\mathcal{A}_t$  (possibly, but not necessarily depending on time). Also,  $f_t, t = 0, 1, 2, \dots, N - 1$  are functions mapping from the space of state variables (state space), controls and noise into the set of possible states of  $\{x_t\}_{t \geq 0}$ . The precise structure of the state space, set of admissible controls and the random parameter space depends on the particular problem under consideration. Note that due to the randomness in  $w_t, t = 0, 1, 2, \dots, N - 1$ , the system state  $x_t$  and control  $u_t, t = 1, 2, \dots, N - 1$  also become random variables.

We remark that because of this, the state equation is sometimes written in the following form,

$$x_{t+1}(\omega) = f_t(x_t(\omega), u_t(\omega), w_t) \tag{4}$$

where  $\omega \in \Omega$  is a scenario in a scenario space  $\Omega$  (representing the randomness). Sometimes, the randomness is suppressed for notational convenience, so the state equation becomes  $x_{t+1} = f_t(x_t, u_t), t = 0, 1, 2, \dots, N - 1$ .

Note that in the state Eq. (3) (alternatively, Eq. (4)),  $x_{t+1}$  only depends on the previous time step, i.e.  $x_t, u_t, w_t$ . This is the *Markov property* (as long as we assume that the distribution of  $w_t$  does not depend on past values of  $w_s, s = 0, 1, \dots, t - 1$ , but only on  $x_t, u_t$ ). That is, the next system state only depends on the previous one. Since this Markovian framework is what will be used throughout this paper as we move on to reinforcement learning for a probabilistic digital twin, we focus on this. However, we remark that there is a large theory of sequential decision making which is not based on this Markovianity. This theory is based around maximum principles instead of dynamic programming.

The aim of the agent is to minimize a cost function under the state constraint (3) (or alternatively, (4)). We assume that this cost function is of the following, additive form,

$$E \left[ g(x_N) + \sum_{t=0}^{N-1} h_t(x_t, u_t, w_t) \right]$$

where the expectation is taken with respect to an a priori given probability measure. That is, we sum over all *instantaneous rewards*  $h_t(x_t, u_t, w_t), t = 0, 1, \dots, N - 1$  which depend on the state of the system, the control and the randomness and add a *terminal reward*  $g(x_N)$  which only depends on the system state at the *terminal time*  $t = N$ . This function is called the *objective function*.

Hence, the stochastic sequential decision making problem of the agent is to choose admissible controls  $u_t, t = 0, 1, 2, \dots, N - 1$  in order to optimize

$$\begin{aligned} \min_{u_t \in \mathcal{A}_t, t \geq 0} E \left[ g(x_N) + \sum_{t=0}^{N-1} h_t(x_t, u_t, w_t) \right] \\ \text{such that } x_{t+1} = f_t(x_t, u_t, w_t), \\ t = 0, 1, 2, \dots, N - 1. \end{aligned} \tag{5}$$

Typically, we assume that the agent has *full information* in the sense that they can choose the control at time  $t$  based on (fully) observing the state process up until this time, but that they are not able to use any more information than this (future information, such as inside information).

This problem formulation is very similar to that of continuous time stochastic optimal control problem.

**Remark 4.1** (A note on continuous time) This framework is parallel to that of stochastic optimal control in continuous time. The main differences in the framework in the continuous time case are that the state equation is typically a stochastic differential equation, and the sum is replaced by an integral in the objective function. For a detailed introduction to continuous time control, see e.g. Øksendal [29].

Other versions of sequential decision making problems include inside information optimal control, partial information optimal control, infinite time horizon optimal control and control with various delay and memory effects. One can also consider problems where further constraints, either on the control or the state, are added to problem (5).

In Bertsekas [30], the sequential decision making problem (5) is studied via the *dynamic programming algorithm*. This algorithm is based on the *Bellman optimality principle*, which says that an optimal policy chosen at some initial time, must be optimal when the problem is re-solved at a later stage given the state resulting from the initial choice.

## 4.2 Sequential decision making in the PDT

Now, we show how the sequential decision making framework from the previous section can be used to solve sequential decision making problems in the PDT.

We may apply this sequential decision making framework to our PDT by letting

$$x_t := \mathbf{X}_t.$$

That is, the state process for the PDT sequential decision making problem is the random vector of attributes  $\mathbf{X}_t$ . Note that in Definition 3.3, there is no time-dependency in the attributes  $\mathbf{X}$ . However, since we are interested in considering sequential decision making in the PDT, we need to assume that there is some sort of development over time (or some indexed set, e.g. information) of the PDT.

Hence, the stochastic sequential decision making problem of the PDT-agent is to choose admissible controls  $u_t$ ,  $t = 0, 1, 2, \dots, N - 1$  in order to optimize

$$\min_{u_t \in \mathcal{A}, t \geq 0} E \left[ g(\mathbf{X}_N) + \sum_{t=0}^{N-1} h_t(\mathbf{X}_t, u_t, w_t) \right] \quad (6)$$

such that  $\mathbf{X}_{t+1} = f_t(\mathbf{X}_t, u_t, w_t)$ ,  
 $t = 0, 1, 2, \dots, N - 1$ .

Here, the set of admissible controls,  $\{u_t\}_{t \geq 0} \in \mathcal{A}$ , are problem specific. So are the functions  $h_t, g_t$  and  $f_t$  for  $t \geq 0$ . Given a particular problem, these functions are defined based on the goal of the PDT-agent as well as the updating of the PDT given new input.

### 4.3 Planning in the PDT

In this section, we discuss how the PDT can be used for *planning*. That is, how we use the PDT to identify an optimal policy, without acting in the real world, but by instead simulating what will happen in the real world given that the agent chooses specific actions (or controls, as they are called in the sequential decision making literature, see Sect. 4.1). We use the PDT as a tool to find a plan (policy), or a single action (first action of policy), to perform in the real world.

In order to solve our sequential decision making problem in the PDT, we have chosen to use a reinforcement learning formulation. This essentially corresponds to choosing the dynamic programming method for solving the optimal control problem. Because we will use a DPP (Dynamic Programming Principle) approach, we need all the assumptions that come with this: A Markovian framework, or the possibility of transforming the problem to something Markovian. We need the Bellman equation to hold in order to avoid issues with time-inconsistency. In order to ensure this, we for example need to use exponential discounting and not have e.g. conditional expectation of state process in a non-linear way in the objective function. Finally, our planning problem cannot have state constraints.

Starting with an initial PDT as a digital representation of a physical system given our current knowledge, we assume that there are two ways to update the PDT:

1. Changing or updating the structural assumptions  $M$ , and hence the probability measure  $P_M$ .
2. Updating the information  $I$ .

The structural assumptions  $M$  are related to the probabilistic model for  $\mathbf{X}$ . Recall from Sect. 2.4, that these assumptions define the probability measure  $P_M$ . Often,

this probability measure is taken as given in stochastic modeling. However, in practice, probability measures are not given to us, but *decided* by analysts based on previous knowledge. Hence, the structural assumptions  $M$  may be updated because of new knowledge, external to the model, or for other reasons the analysts view as important.

Updating the information is our main concern in this paper, since this is related to the agent making costly decisions in order to gather more information. An update of the information also means (potentially) reducing the epistemic uncertainty in the PDT. Optimal information gathering in the PDT will be discussed in detail in the following Sect. 4.6.

### 4.4 MDP, POMDP and their relation to DPP

In this section, we briefly recall the definitions of Markov decision processes, partially observable Markov decision processes and explain how these relate to the sequential decision making framework of Sect. 4.1.

*Markov decision processes* (MDPs) are discrete-time stochastic control processes of a specific form. An MDP is a tuple

$$(S, A, P_a, R_a),$$

where  $S$  is a set of states (the state space) and  $A$  is a set of actions (action space). Also,

$$P_a(s, s') = P_a(s_{t+1} = s' | a_t = a, s_t = s)$$

is the probability of going from state  $s$  at time  $t$  to state  $s'$  at time  $t + 1$  if we do action  $a$  at time  $t$ . Finally,  $R_a(s, s')$  is the instantaneous reward of transitioning from state  $s$  at time  $t$  to state  $s'$  at time  $t + 1$  by doing action  $a$  (at time  $t$ ).

An MDP satisfies the *Markov property*, so given that the process is in state  $s$  and will be doing  $a$  at time  $t$ , the next state  $s_{t+1}$  is conditionally independent of all other previous states and actions.

#### Remark 4.2 (MDP and DPP)

Note that this definition of an MDP is essentially the same as our DPP framework of Sect. 4.1. In the MDP notation, we say *actions*, while in the control notation, it is common to use the word *control*. In Sect. 4.1, we talked about *instantaneous cost functions*, but here we talk about *instantaneous rewards*. Since minimization and maximization problems are equivalent (since  $\inf\{z\} = -\sup\{-z\}$ ), so are these two concepts. Furthermore, the definition of the transition probabilities  $P_a$  in the MDP framework corresponds to the Markov assumption of the DPP method. In both frameworks, we talk about the system states, though in the DPP framework we model this directly via Eq. (3).

A generalization of MDPs are *partially observable Markov decision processes* (POMDPs). While an MDP is a 4-tuple, a POMDP is a 6-tuple,

$$(S, A, P_a, R_a, \bar{\Omega}, O).$$

Here (like before),  $S$  is the state space,  $A$  is the action space,  $P_a$  gives the conditional transition probabilities between the different states in  $S$  and  $R_a$  gives the instantaneous rewards of the transitions for a particular action  $a$ .

In addition, we have  $\bar{\Omega}$ , which is a *set of observations*. In contrast to the MDP framework, with POMDPs, the agent no longer observes the state  $s$  directly, but only an observation  $o \in \bar{\Omega}$ . Furthermore, the agent knows  $O$  which is a *set of conditional observation probabilities*. That is,

$$O(o|s', a)$$

is the probability of observing  $o \in \bar{\Omega}$  given that we do action  $a$  from state  $s'$ .

The objective of the agent in the POMDP sequential decision problem is to choose a policy, that is actions at each time, in order to optimize

$$\max_{(a_t) \in A} E \left[ \sum_{t=0}^T \lambda^t r_t \right] \tag{7}$$

where  $r_t$  is the reward earned at time  $t$  (depending on  $s_t, a_t$  and  $s_{t+1}$ ), and  $\lambda \in [0, 1]$  is a number called the discount factor. The discount factor can be used to introduce a preference for immediate rewards as opposed to more distant rewards, which may be relevant for the problem at hand, or used just for numerical efficiency. Hence, the agent aims to maximize their expected discounted reward over all future times. Note that it is also possible to consider problem (7) over an infinite time horizon or with a separate terminal reward function as well. This is similar to the DPP sequential decision making framework of Sect. 4.1.

In order to solve a POMDP, it is necessary to include memory of past actions and observations. Actually, the inclusion of partial observations means that the problem is no longer Markovian. However, there is a way to Markovianize the POMDP by transforming the POMDP into a *belief-state MDP*. In this case, the agent summarizes all information about the past in a belief vector  $b(t)$ , which is updated as time passes. See [31], Chapter 12.2.3 for details.

### 4.5 MDP (and POMDP) in the PDT framework

In this section, we show how the probabilistic digital twin can be incorporated in a reinforcement learning framework, in order to solve sequential decision problems in the PDT.

In Sect. 4.2, we showed how we can use the mathematical framework of sequential decision making to solve optimal control problems for a PDT-agent. Also, in Sect. 4.4, we saw (in Remark 4.2) that the MDP (or POMDP in general) framework essentially corresponds to that of the DPP. In theory, we could use the sequential decision making framework and the DPP to solve optimal control problems in the PDT. However, due to the curse of dimensionality, this will typically not be practically tractable. In order to resolve this, we cast the PDT sequential decision making problem into a reinforcement learning, in particular a MDP, framework. This will enable us to solve the PDT optimal control problem via deep reinforcement learning, in which there are suitable tools to overcome the curse of dimensionality.

To define a decision making process in the PDT as a MDP, we need to determine our state space, action space, (Markovian) transition probabilities and a reward function.

- The action space  $A$ : These are the possible actions within the PDT. These may depend on the problem at hand. In the next Sect. 4.6, we will discuss optimal information gathering, where the agent can choose between different types of experiments, at different costs, in order to gain more information. In this case, the action space is the set of possible decisions that the agent can choose from in order to attain more information.
- The state space  $S$ : We define a state as a PDT (or equivalently a *version* of a PDT that evolves in discrete time  $t = 0, 1, \dots$ ). A PDT represents our belief about the current *physical* state of a system, and it is defined by some initial assumptions together with the information acquired through time. In practice, if the structural assumptions are not changed, we may let the information available at the current time represent a state. This means that our MDP will consist of *belief-states*, represented by information, from which inference about the true physical state can be made. This is a standard way of creating a MDP from a POMDP, so we can view the *PDT state-space* as a space of *beliefs* about some underlying partially observable physical state. Starting from a PDT, we define the state space as all updated PDTs we can reach by taking actions in the action space  $A$ .
- The transition probabilities  $P_a$ : Based on our chosen definition of the state space, the transition probabilities are the probabilities of going from one level of information to another, given the action chosen by the agent. For example, if the agent chooses to make decision (action)  $d$ , what is the probability of going from the current level of information to another (equal or better) level. This is given by epistemic conditioning of the PDT with respect to the given information set

$I = \{(d, o)\}$ , which is based on the decision  $d$  and the new observation  $o$ . When it comes to updates of the structural assumptions  $M$ , we consider this as deterministic transitions.

- The reward  $R_a$ : The reward function, or equivalently, cost function, will depend on the specific problem at hand. To each action  $a \in A$ , we assume that we have an associated reward  $R_a$ . In the numerical examples in Sect. 5, we give specific examples of how these rewards can be defined.

As mentioned in Sect. 4.3, there are two ways to update the PDT: Updating the structural assumptions  $M$  and updating the information  $I$ . If we update the PDT by (only) adding to the information set  $I$ , we always have the Markov property.

If we also update  $M$ , then the preservation of the Markov property is not given. In this case, using a maximum principle deep learning algorithm instead of the DPP based deep RL is a possibility, see [32].

**Remark 4.3** Note that in the case where we have a very simple PDT with only discrete variables and only a few actions, then the RL approach is not necessary. In this case, the DPP method as done in traditional optimal control works well, and we can apply a planning algorithm to the PDT in order to derive an optimal policy. However, in general, the state-action space of the PDT will be too large for this. Hence, traditional planning algorithms, and even regular RL may not be feasible due to the curse of dimensionality. In this paper, we will consider deep reinforcement learning as an approach to deal with this. We discuss this further in Sect. 5.

Note that what determines an optimal action or policy will of course depend on what *objective* the outcomes are measured against. That is, what do we want to achieve in the real world? There are many different objectives we could consider. In the following we present one generic objective related to optimal information gathering, where the PDT framework is suitable.

## 4.6 Optimal information gathering

A generic, but relevant, objective in optimal sequential decision making is simply to “improve itself”. That is, to reduce epistemic uncertainty with respect to some quantity of interest. Another option, is to consider maximizing the Kullback–Leibler divergence with respect to epistemic uncertainty as a general objective. This would mean that we aim to collect the information that “will surprise us the most”. See for instance [8] for a review of some common alternatives.

By definition, a PDT contains an observational model related to the data generating process (the epistemic conditioning relies on this). This means that we can simulate the effect of gathering information, and we can study how to do this optimally. In order to define what we mean by an optimal strategy for gathering information, we then have to specify the following:

- *Objective*: What we need the information for. For example, what kind of decision do we intend to support using the PDT? Is it something we want to estimate? What is the required accuracy needed? For instance, we might want to reduce epistemic uncertainty with respect to some quantity, e.g. a risk metric such as a failure probability, expected extreme values, etc.
- *Cost*: The cost related to the relevant information-gathering activities.

Then, from the PDT together with a specified *objective* and *cost*, one alternative is to define the optimal strategy as the strategy that minimizes the (discounted) expected cost needed to achieve the objective (or equivalently achieves the objective while maximizing reward).

**Example 4.4** (Coin flip – information gathering) Continuing from Example 3.1, imagine that before making your final bet, you can flip the coin as many times as you like in order to learn about  $\theta$ . Each of these test flips will cost 10.000. You also get the opportunity to replace the coin with a new one, at the cost of 100.000.

An interesting problem is now how to select an optimal strategy for when to *test*, *bet* or *replace* in this game. And will such a strategy be robust? What if there is a limit on the total number of actions than can be performed? In Sect. 5.5 we illustrate how reinforcement learning can be applied to study this problem, where the coin represents a component with reliability  $\theta$ , that we may *test*, *use* or *replace*.

## 5 Deep Reinforcement Learning with PDTs

In this section we give an example of how reinforcement learning can be used for planning, i.e. finding an optimal action or policy, with a PDT. The reinforcement learning paradigm is especially relevant for problems where the state and/or action space is large, or dynamical models where specific transition probabilities are not easily attainable but where efficient sampling is still feasible. In probabilistic modelling of complex physical phenomena, we often find ourselves in this kind of setting.

### 5.1 Reinforcement Learning (RL)

Reinforcement learning, in short, aims to optimize sequential decision problems through sampling from a MDP (Sutton and Barto [33]). We think of this as an agent taking actions within an environment, following some policy  $\pi(a|s)$ , which gives the probability of taking action  $a$  if the agent is currently at state  $s$ . Generally,  $\pi(a|s)$  represents a (possibly degenerate) probability distribution over actions  $a \in A$  for each  $s \in S$ . The agent’s objective is to maximize the amount of reward it receives over time, and a policy  $\pi$  that achieves this is called an *optimal policy*.

Given a policy  $\pi$  we can define the *value* of a state  $s \in S$  as

$$v_\pi(s) = E \left[ \sum_{t=0}^T \lambda^t r_t \mid s_0 = s \right] \tag{8}$$

where  $r_t$  is the reward earned at time  $t$  (depending on  $s_t, a_t$  and  $s_{t+1}$ ), given that the agent follows policy  $\pi$  starting from  $s_0 = s$ . That is, for  $P_a$  and  $R_a$  given by the MDP,  $a_t \sim \pi(a_t|s_t), s_{t+1} \sim P_{a_t}(s_t, s_{t+1})$  and  $r_t \sim R_{a_t}(s_t, s_{t+1})$ . Here we make use of a discount factor  $\lambda \in [0, 1]$  in the definition of cumulative reward. If we want to consider  $T = \infty$  (continuing tasks) instead of  $T < \infty$  (episodic task), then  $\lambda < 1$  is generally necessary.

The *optimal value function* is defined as the one that maximizes (8) over all policies  $\pi$ . The optimal action at each state  $s \in S$  then corresponds to acting greedily with respect to this value function, i.e. selecting the action  $a_t$  that in expectation maximises the value of  $s_{t+1}$ . Likewise, it is common to define the *action-value* function  $q_\pi(s, a)$ , which corresponds to the expected cumulative return of first taking action  $a$  in state  $s$  and following  $\pi$  thereafter. RL generally involves some form of Monte Carlo simulation, where a large number of *episodes* are sampled from the MDP, with the goal of estimating or approximating the optimal value of states, state-action pairs, or an optimal policy directly.

Theoretically this is essentially equivalent to the DPP framework, but with RL we are mostly concerned with problems where optimal solutions cannot be found easily and some form of approximation is needed. By the use of flexible approximation methods combined with adaptive sampling strategies, RL makes it possible to deal with large and complex state and action spaces.

### 5.2 Function approximation

One way of using function approximation in RL is to define a parametric function  $\hat{v}(s|\mathbf{w}) \approx v_\pi(s)$ , given by a set of weights  $\mathbf{w} \in \mathbb{R}^d$ , and try to learn the value function of an optimal policy by finding an appropriate value

for  $\mathbf{w}$ . Alternatively, we could approximate the value of a state-action pair,  $\hat{q}(s, a|\mathbf{w}) \approx q_\pi(s, a)$ , or a policy  $\hat{\pi}(a|s, \mathbf{w}) \approx \pi(a|s)$ . The general goal is then to optimize  $\mathbf{w}$ , using data generated by sampling from the MDP; the RL literature contains many different algorithms designed for this purpose. In the case where a neural network is used for function approximation, it is often referred to as *deep reinforcement learning*. One alternative, which we will make use of in an example later on, is the deep Q-learning approach (DQN) as introduced by van Hasselt et al. [34], which represents the value of a set of  $m$  actions at a state  $s$  using a multi-layered neural network

$$\hat{q}(s|\mathbf{w}) : S \rightarrow \mathbb{R}^m. \tag{9}$$

Note here that  $\hat{q}(s|\mathbf{w})$  is a function defined on the state space  $S$ . In general, any approximation of the value functions  $v$  or  $q$ , or the policy  $\pi$  are defined on  $S$  or  $S \times A$ . A question that then arises, is how can we define parametric functions on the state space  $S$  when we are dealing with PDTs? We can assume that we have control over the set of admissible actions  $A$ , in the sense that this is something we define, and creating parametric functions defined on  $A$  should not be a problem. But as discussed in Sect. 4.5,  $S$  will consist of *belief-states*.

### 5.3 Defining the state space

We are interested in an MDP where the transition probabilities  $P_a(s, s')$  correspond to updating a PDT as a consequence of action  $a$ . In that sense,  $s$  and  $s'$  are PDTs. Given a well-defined set of admissible actions, the state space  $S$  is then the set of all PDTs that can be obtained starting from some initial state  $s_0$ , within some defined horizon.

Recall that going from  $s$  to  $s'$  then means keeping track of any changes made to the structural assumptions  $M$  and the information  $I$ , as illustrated in Fig. 3. From now on, we will for simplicity assume that updating the PDT only involves epistemic conditioning with respect to the information  $I$ . This is a rather generic situation. Also, finding a way to represent changes in  $M$  will have to be handled for the specific use case under consideration. Assuming some initial PDT  $s_0$  is given, any state  $s_t$  at a later time  $t$  is then uniquely defined by the set of information  $I_t$  available at time  $t$ . Representing states by information in this way is something that is often done to transform a POMDP to a MDP. That is, although the true state  $s_t$  at time  $t$  is unknown in a POMDP, the information  $I_t$ , and consequently our *belief* about  $s_t$ , is always known at time  $t$ . Inspired by the POMDP terminology, we may therefore view a PDT as a belief-state, which seems natural as the PDT is essentially a way to encode our beliefs about some real physical system.



Hence, we will proceed with defining the state space  $S$  as the *information state-space*, which is the set of all sets of information  $I$ . Although this is a very generic approach, we will show that there is a way of defining a flexible parametric class of functions on  $S$ . But we must emphasize that if there are other ways of compressing the information  $I$ , for instance due to conjugacy in the epistemic updating, then these are probably much more efficient. Example 5.1 below shows exactly what we mean by this.

**Example 5.1** (Coin flip – information state-space) In the coin flip example (Example 2.3), all of our belief with respect to epistemic uncertainty is represented by the number  $\psi = P(\theta = \theta_1)$ . Given some observation  $Y = y \in \{0, 1\}$ , the epistemic conditioning corresponds to

$$\psi \rightarrow \frac{\beta_1(y)\psi}{\beta_1(y)\psi + \beta_2(y)(1 - \psi)}$$

where, for  $j = 1, 2$ ,  $\beta_j(y) = \theta_j$  if  $y = 0$  and  $\beta_j(y) = 1 - \theta_j$  if  $y = 1$ .

In this example, the information state-space consists of all sets of the form  $I_t = \{y_1, \dots, y_t\}$  where each  $y_i$  is binary. However, if the goal is to let  $I_t$  be the representation of a PDT, we could just as well use  $\psi_t$ , i.e. define  $S = [0, 1]$  as the state space. Alternatively, the number of heads and tails (0 s and 1 s) provides the same information, so we could also make use of  $S = \{0, \dots, N\} \times \{0, \dots, N\}$  where  $N$  is an upper limit on the total number of flips we consider.

#### 5.4 Deep learning on the information state-space

Let  $S$  be a set of sets  $I \subset \mathbb{R}^d$ . We will assume that each set  $I \in S$  consists of a finite number of elements  $y \in \mathbb{R}^d$ , but we do not require that all sets  $I$  have the same size. We are interested in functions defined on  $S$ .

A function  $f$  is permutation invariant if  $f(\{\mathbf{y}_1, \dots, \mathbf{y}_N\}) = f(\{\mathbf{y}_{\kappa(1)}, \dots, \mathbf{y}_{\kappa(N)}\})$  for any permutation  $\kappa$ . It can be shown that under fairly mild assumptions, such functions have the following decomposition

$$f(I) = \rho \left( \sum_{\mathbf{y} \in I} \phi(\mathbf{y}) \right). \quad (10)$$

These sum decompositions were studied by Zaheer et al. [35] and later by Wagstaff et al. [36], which showed that if  $|I| \leq p$  for all  $I \in S$ , then any continuous function  $f : S \rightarrow \mathbb{R}$  can be written as (10) for some suitable functions  $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^p$  and  $\rho : \mathbb{R}^p \rightarrow \mathbb{R}$ . The motivation in [35, 36] was to enable supervised learning of permutation invariant and set-valued functions, by replacing  $\rho$  and  $\phi$  with flexible function approximators, such as Gaussian

processes or neural networks. Other forms of decomposition, by replacing the *summation* in (10) with something else that can be learned, have also been considered by Soelch et al. [37]. For reinforcement learning, we will make use of the form (10) to represent functions defined on the information states space  $S$ , such as  $\hat{v}(s|\mathbf{w})$ ,  $\hat{q}(s, a|\mathbf{w})$ , or  $\hat{\pi}(a|s, \mathbf{w})$ , using a neural network with parameter  $\mathbf{w}$ . In the remaining part of this paper we present two examples showing how this works in practice.

#### 5.5 The “coin flip” example

Throughout this paper we have presented a series of small examples involving a biased coin, represented by  $X = (Y, \theta)$ . In Example 4.4 we ended by introducing a game where the player has to select whether to *bet on*, *test* or *replace* the coin. As a simple illustration we will show how reinforcement learning can be applied in this setting.

But now, we will imagine that the coin  $Y$  represents a *component* in some physical system, where  $Y = 0$  corresponds to the component functioning and  $Y = 1$  represents failure. The probability  $P(Y = 1) = 1 - \theta$  is then the component's *failure probability*, and we say that  $\theta$  is the *reliability*.

For simplicity we assume that  $\theta \in \{0.5, 0.99\}$ , and that our initial belief is  $P(\theta = 0.5) = 0.5$ . That is, when we buy a new component, there is a 50% chance of getting a “bad” component (that fails 50% of the time), and consequently a 50% probability of getting a “good” component (that fails 1% of the time).

We consider a project going over  $N = 10$  days. Each day we will decide between one of the following 4 actions:

1. Test the component (flip the coin once). Cost  $r = -10.000$ .
2. Replace the component (buy a new coin). Cost  $r = -100.000$ .
3. Use the component (bet on the outcome). Obtain a reward of  $r = 10^6$  if the component works ( $Y = 0$ ) and a cost of  $r = -10^6$  if the component fails ( $Y = 1$ ).
4. Terminate the project (set  $t = N$ ),  $r = 0$ .

We will find a deterministic policy  $\pi : S \rightarrow A$  that maps from the information state-space to one of the four actions. The information state-space  $S$  is here represented by the number of days left of the project,  $n = N - t$ , and the set  $I_t$  of observations of the component that is currently in use at time  $t$ . If we let  $S_Y$  contain all sets of the form  $I = \{Y_1, \dots, Y_t\}$ , for  $Y_t \in \{0, 1\}$  and  $t < N$ , then

$$S = S_Y \times \{1, \dots, N\} \quad (11)$$

represents the information state-space. In this example we use the deep Q-learning (DQN) approach described by van Hasselt et al. [34]. Here we define a neural network

$$\hat{q}(s|\mathbf{w}) = \hat{q}(l, n|\mathbf{w}) : S \rightarrow \mathbb{R}^4,$$

that represents the action-value of each of the four actions. The optimal policy is then obtained by at each state  $s$  selecting the action corresponding to the maximal component of  $\hat{q}$ . We implement this by writing  $\hat{q}(l, n|\mathbf{w}) = h(f(l|\mathbf{w}_f), n|\mathbf{w}_h)$ , where  $h(f, n|\mathbf{w}_h) : \mathbb{R}^4 \times \mathbb{R} \rightarrow \mathbb{R}^4$ , and the function  $f(l|\mathbf{w}_f)$  taking sets as input is defined as in (10) using two functions  $\phi(\cdot|\mathbf{w}_\phi) : \mathbb{R} \rightarrow \mathbb{R}^4$  and  $\rho(\cdot|\mathbf{w}_\rho) : \mathbb{R}^4 \rightarrow \mathbb{R}^4$ . We let  $h(\cdot|\mathbf{w}_h)$ ,  $\phi(\cdot|\mathbf{w}_\phi)$  and  $\rho(\cdot|\mathbf{w}_\rho)$  be multi-layer perceptrons so that  $\hat{q}(\cdot|\mathbf{w})$  becomes a feedforward neural network with a combined parameter vector  $\mathbf{w} = (\mathbf{w}_h, \mathbf{w}_\phi, \mathbf{w}_\rho)$  that is optimized using the Q-learning algorithm<sup>5</sup>.

We start by finding a policy that optimizes the cumulative reward over the 10 days (without discounting). To test the model, we run a large number of episodes and store the total reward after each episode. As it turns out, this policy prefers to “gamble” that the component works rather than performing tests. In the case where the starting component is reliable (which happens 50% of the time), a high reward can be obtained by selecting action 3 at every opportunity. The general “idea” with this policy, is that if action 3 results in failure, the following action is to replace the component (action 2), unless there are few days left of the project in which case action 4 is selected. We call this the “unconstrained” policy.

Although the unconstrained policy gives the largest expected reward, there is an approximately 50% chance that it will produce a failure, i.e. that action 3 is selected with  $Y = 1$  as the resulting outcome. One way to reduce this failure probability, is to introduce the constraint that action 3 (using the component) is not allowed unless we have a certain level of confidence in that the component is reliable. We introduce this type of constraint by requiring that  $P(\theta = 0.99) > 0.9$  (a constraint on epistemic uncertainty). The optimal policy under this constraint will start with running experiments (action 1), before deciding whether to replace (action 2), use the component (action

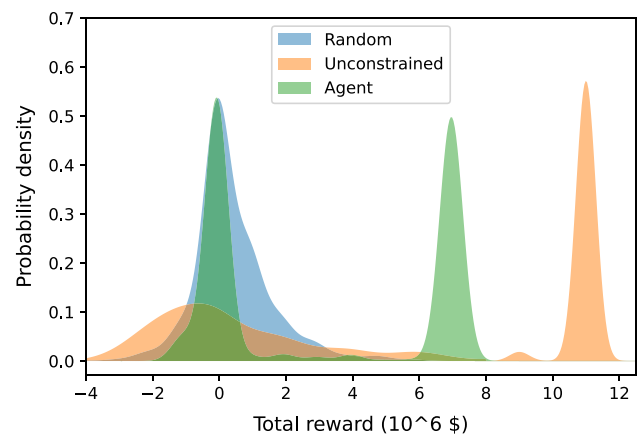


Fig. 4 Total reward after 1000 episodes for a random policy, the unconstrained policy, and the agent which is subjected to the constraint that action 3 is not allowed unless  $P(\theta = 0.99) > 0.9$

3), or terminate the project (action 4). Figure 4 shows the distributions of the cumulative reward over 1000 simulated episodes, for the constrained and unconstrained policies obtained by RL, together with a completely random policy for comparison.

In this example, the information state-space could also be defined in a simpler way, as explained in Example 5.1. As a result the reinforcement learning task will be simplified. Using the different state-space representations, we obtained the same results shown in Fig. 4. Finally, we should note that in the case where defining the state space as in (11) is necessary, the constraint  $P(\theta = 0.99) > 0.9$  is not practical. That is, if we could estimate this probability efficiently, then we also have access to the compressed information state-space. One alternative could then be to instead consider the uncertain failure probability  $p_f(\theta) = P(Y = 1|\theta)$ , and set a limit on e.g.  $E[p_f] + 2 \cdot \text{Std}(p_f)$ . This is the approach taken in the following example concerning failure probability estimation.

### 5.6 Corroded pipeline example


Here we revisit the corroded pipeline example from Agrell and Dahl [16] which we introduced in Sect. 3.4. In this example, we specify the epistemic uncertainty with respect to model discrepancy, the size of a defect, and the capacity  $p_{FE}$  coming from a Finite Element simulation. If we let  $\theta$  be the epistemic generator, we can write the failure probability conditioned on epistemic information as  $p_f(\theta) = P(g \leq 0|\theta)$ . In [16] the following objective was considered: Determine with confidence whether  $p_f(\theta) < 10^{-3}$ . That is, when we consider  $p_f$  as a purely epistemic random variable, we want to either confirm that the failure probability is less than the target  $10^{-3}$  (in which case we can continue operations as normal),

<sup>5</sup> For our experiments we used a PyTorch (v 1.13.0) implementation of the Q-learning algorithm with experience replay (100 episode batch size) and a target network (10 episode frequency) as described in Algorithm 1 in [38]. For the neural network  $\hat{q}(s|\mathbf{w})$ , we used 3 hidden layers of size 150 for  $\phi$ , 2 hidden layers of size 150 for  $\rho$ , and 2 hidden layers of size 50 for  $h$ , all with ReLU activations. During training, random actions were taken with a probability  $\epsilon$  decaying from 0.5 with steps 0.001, the discount factor  $\gamma$  was set to 0.999 and the Adam optimizer was used with a learning rate of 0.001 for gradient descent.

or to detect with confidence that the target is exceeded (and we have to intervene). We say the objective is achieved if we obtain either  $E[p_f] + 2 \cdot \text{Std}(p_f) < 10^{-3}$  or  $E[p_f] - 2 \cdot \text{Std}(p_f) > 10^{-3}$  (where  $E[p_f]$  and  $\text{Std}(p_f)$  can be efficiently approximated using the method developed in [16]). There are three ways in which we can reduce epistemic uncertainty:

- 1 **Defect measurement:** Noise perturbed measurement that reduces uncertainty in the defect size  $d$ .
- 2 **Computer experiment:** Evaluate  $p_{FE}$  at some selected input  $(D, t, s, d, l)$ , to reduce uncertainty in the surrogate  $\hat{p}_{FE}$  used to approximate  $p_{FE}$ .
- 3 **Lab experiment:** Obtain one observation of  $X_{m^r}$ , which reduces uncertainty in  $\mu_{m^r}$ .

The set of information corresponding to defect measurements is  $I_{\text{Measure}} \subset \mathbb{R}$  as each measurement is a real valued number. Similarly,  $I_{\text{Lab}} \subset \mathbb{R}$  as well, and  $I_{FE} \subset \mathbb{R}^6$  when we consider a vector  $\mathbf{y} \in \mathbb{R}^6$  as an experiment  $[D, t, s, d, l, p_{FE}]$ . In this example, we exploit conjugacy in the representation of  $I_{\text{Measure}}$  and  $I_{\text{Lab}}$  as discussed in Example 5.1. All relevant information from  $I_{\text{Measure}}$  is represented by the posterior distribution of the quantity that is measured, which is normally distributed defined by its mean and standard deviation. The same is true for  $I_{\text{Lab}}$ . Hence, the information represented by  $I_{\text{Measure}}$  and  $I_{\text{Lab}}$  can be represented by a single vector in  $\mathbb{R}^4$  (see [16] for details). We therefore define the information state-space as  $S = S_{FE} \times \mathbb{R}^4$ , where  $S_{FE}$  consists of finite subsets of  $\mathbb{R}^6$ . We use RL to determine which of the three types of experiments to perform, and define the action space  $A = \{\text{Measurement, FE, Lab}\}$ . Note that when we decide to run a computer experiment, we also have to specify the input  $(D, t, s, d, l)$ . This is a separate decision making problem regarding design of experiments. For this we make use of the myopic (one-step lookahead) method developed in [16], although one could in principle use RL for this as well. This kind of decision making, where one first decides between different types of tasks to perform, and then proceeds to find a way to perform the selected task optimally, is often referred to as hierarchical RL in the reinforcement learning literature. Actually, [16] considers a myopic alternative for also selecting between the different types of experiments, and it was observed that this might be challenging in practice if there are large differences in cost between the experiments. This was the motivation for studying the current example, where we now define the reward (cost)  $r$  as a direct consequence of  $a \in A$  as follows:  $r = -10$  for  $a = \text{Measurement}$ ,  $r = -1$  for  $a = \text{Lab}$  and  $r = -0.1$  for  $a = \text{FE}$ .



**Fig. 5** Total cost (negative reward) after 100 successful episodes. For the random and benchmark policy, the success rate was around 60% (to achieve the objective within 40 experiments in total), whereas 94% was successful for the RL agent

In this example we also use the DQN approach of van Hasselt et al. [34], where we define a neural network

$$\hat{q}(s, \mathbf{w}) : S = S_{FE} \times \mathbb{R}^4 \rightarrow \mathbb{R}^3,$$

that gives, for each state  $s$ , the (near optimal) value of each of the three actions. The neural network is constructed in the same way as in the “coin flip” example in Sect. 5.5, and trained with the same Q-learning approach (Algorithm 1 [38]).

The objective in this RL example is to estimate a failure probability using as little resources as possible. If an agent achieves the criterion on epistemic uncertainty reduction, that the expected failure probability plus/minus two standard deviations is either above or below the target value, we say that the agent has succeeded and we report the sum of the cost of all performed experiments. We also set a maximum limit of 40 experiments. I.e. after 40 tries the agent has failed. To compare the policy obtained by RL, we consider the random policy that selects between the three actions uniformly at random. We also consider a more “human like” benchmark policy, that corresponds to first running 10 computer experiments, followed by one lab experiment then one defect measurement, then 10 new computer experiments, and so on.

The final results from simulating 100 episodes with each of the three policies are shown in Fig. 5.

## 6 Concluding remarks

To conclude our discussion, we recall that in this paper, we have:

- Given a measure-theoretic discussion of epistemic uncertainty and formally defined epistemic conditioning.
- Provided a mathematical definition of a probabilistic digital twin (PDT).
- Connected PDTs with sequential decision making problems, and discussed several solution approaches (maximum principle, dynamic programming, MDP and POMDP).
- Argued that using (deep) reinforcement learning to solve sequential decision making problems in the PDT is a good choice for practical applications today.
- For the specific use-case of optimal information gathering, we proposed a generic solution using deep reinforcement learning on the information state-space.

Further research in this direction includes looking at alternative solution methods and reinforcement learning algorithms in order to handle different PDT frameworks. A possible idea is to use a maximum principle approach instead of a DPP approach (as is done in reinforcement learning). By using one of the MP based algorithms in [32], we may avoid the Markovianity requirement, possible time-inconsistency issues (see e.g. Rudloff et al. [39]) and can also allow for state constraints. For instance, this is of interest when the objective of the sequential decision making problem in the PDT is to minimize a risk measure such as CVaR or VaR, see Cheridito and Stadje [40], and Artzner et al. [41] respectively. Both of these risk measures are known to cause time-inconsistency in the Bellman equation, and hence, the DPP (and also reinforcement learning) cannot be applied in a straightforward manner. Another relevant topic is whether formal methods for model checking can be applied for the reinforcement learning approach considered in this paper. This is particularly important for extending beyond cost optimization towards optimizing decisions that may also impact the level of safety of the physical system. The application of PDTs for safe reinforcement learning (see e.g. [42]) is therefore a relevant topic. This is work in progress.

**Acknowledgements** The main ideas presented in this paper are the culmination of input from recent research activities related to risk, uncertainty, machine learning and probabilistic modelling in DNV and at the University of Oslo. In particular we want to thank Simen Eldevik, Frank Børre Pedersen and Carla Ferreira for valuable input to this paper.

**Author Contributions** All authors contributed equally.

**Funding** This work has been supported by grant 276282 from the Research Council of Norway (RCN) and DNV Group Research and Development (Christian Agrell and Andreas Hafver). The work has also been supported by grant 29989 from the Research Council of Norway as part of the SCROLLER project (Kristina Rognlien Dahl).

**Availability of data and materials** The only relevant data in this paper is synthetic data generated from the probabilistic models as described in the examples.

## Declarations

**Conflict of interest** This work has been supported by grant 276282 from the Research Council of Norway (RCN) and DNV Group Research and Development (Christian Agrell and Andreas Hafver). Agrell and Hafver are employed by Group Research and Development, DNV, Norway. The work has also been supported by grant 29989 from the Research Council of Norway as part of the SCROLLER project (Kristina Rognlien Dahl). Rognlien Dahl was previously employed at Department of Mathematics, University of Oslo, Norway. She is currently employed at the Department of Economics, BI Norwegian Business School, Norway.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## Appendices

### A Existence of the epistemic generator

The purpose of this section, is to explain why Assumption 2.5 (of the existence of a generating random variable for the epistemic  $\sigma$ -algebra) holds under some very mild assumptions.

In order to do this, we consider a *standard probability space*. Roughly, this is a probability space consisting of an interval and/or a countable (or finite) number of atoms. Formally, a probability space is standard if it is isomorphic (up to  $P$ -null sets) with an interval equipped with the Lebesgue measure, a countable (or finite) set of atoms, or a disjoint union of both of these types of sets.

The following proposition says that in a standard probability space, any sub  $\sigma$ -algebra is generated by a random variable up to  $P$ -null sets. For a proof, see e.g. Greinecker [43].

**Proposition A.1** *Let  $(\Omega, \mathcal{F}, P)$  be a standard probability space and  $\mathcal{E} \subseteq \mathcal{F}$  a sub  $\sigma$ -algebra.*

Then there exists an  $\mathcal{F}$ -measurable random variable  $\theta$  such that

$$\mathcal{E} = \sigma(\theta) \text{ mod } 0.$$

Hence, as long as our probability space is standard (which is a mild assumption), we can assume that our sub  $\sigma$ -algebra of epistemic information,  $\mathcal{E}$ , is generated (up to  $P$ -null sets) by a random variable  $\theta$  without loss of generality. Note that for the purpose of this paper, the mod 0 (i.e. up to  $P$ -null sets) is not a problem. Since we are only considering conditional expectations (or in particular, expectations), the  $P$ -null sets disappear.

Actually, this generating random variable,  $\theta$ , can always be modified to another random variable,  $\hat{\theta}$ , which is  $\mathcal{E}$ -measurable (purely epistemic) by augmenting the  $P$ -null sets. This means that  $\theta$  and  $\hat{\theta}$  are the same with respect to conditional expectations.

Furthermore, if  $\mathbf{X}$  is a random variable on this standard probability space,  $\mathbf{X}|\theta$ , is purely aleatory, i.e. independent of  $\mathcal{E}$ . This follows because  $\mathbf{X}|\theta$  is independent of  $\sigma(\theta)$  and independence holds  $P$ -almost surely, so mod 0 does not affect this.

## References

- Hafver A, Eldevik S, Pedersen FB (2018) Probabilistic digital twins . [Online position paper by DNV GL Group Technology and Research; posted 28-August-2018; <https://ai-and-safety.dnvgl.com/probabilistic-twin/>]
- Ghanem R, Soize C, Mehrez L, Aitharaju V (2020) Probabilistic learning and updating of a digital twin for composite material systems. *Int J Numer Meth Eng*. <https://doi.org/10.1002/nme.6430>
- Kapteyn M, Pretorius J, Willcox K (2021) A probabilistic graphical model foundation for enabling predictive digital twins at scale. *Nat Computat Sci* 1:337–347
- Ayello F, Liu G, Yang Y, Cui N (2021) Probabilistic digital twins for transmission pipelines. *Proc. NACE 2021*
- Milton MA (2021) Real-time probabilistic solvers for digital twins of power electronic systems. PhD thesis, College of Engineering and Computing - University of South Carolina
- Chaloner K, Verdinelli I (1995) Bayesian experimental design: a review. *Stat Sci* 10(3):273–304
- Settles B (2009) Active learning literature survey. *Computer Sciences Technical Report 1648*, University of Wisconsin–Madison
- Ryan EG, Drovandi CC, McGree JM, Pettitt AN (2016) A review of modern computational algorithms for Bayesian optimal design. *Int Stat Rev* 84(1):128–154
- Lindley DV (1956) On a measure of the information provided by an experiment. *Ann Math Stat* 27(4):986–1005
- Huan X, Marzouk YM (2013) Simulation-based optimal Bayesian experimental design for nonlinear systems. *J Comput Phys* 232(1):288–317
- Kim W, Pitt M, Lu Z-L, Steyvers M, Myung J (2014) A hierarchical adaptive approach to optimal experimental design. *Neural Comput* 26:1–28
- Boluki S, Qian X, Dougherty ER (2019) Experimental design via generalized mean objective cost of uncertainty. *IEEE Access* 7:2223–2230
- Ashenafi Y, Pandita P, Ghosh S (2021) Reinforcement learning based sequential batch-sampling for Bayesian optimal experimental design. *arXiv preprint arXiv:2112.10944*
- Blau T, Bonilla E, Dezfouli A, Chadés I (2022) Optimizing sequential experimental design with deep reinforcement learning. *arXiv preprint arXiv:2202.00821*
- Shen W, Huan X (2022) Bayesian sequential optimal experimental design for nonlinear models using policy gradient reinforcement learning. *arXiv preprint arXiv:2110.15335*
- Agrell C, Dahl KR (2021) Sequential Bayesian optimal experimental design for structural reliability analysis. *Statist Comput* 31(3):27. <https://doi.org/10.1007/s11222-021-10000-2>
- Bect J, Ginsbourger D, Li L, Picheny V, Vazquez E (2012) Sequential design of computer experiments for the estimation of a probability of failure. *Stat Comput* 22(3):773–793. <https://doi.org/10.1007/s11222-011-9241-4>
- Echard B, Gayton N, Lemaire M (2011) AK-MCS: an active learning reliability method combining Kriging and Monte Carlo Simulation. *Struct Saf* 33(2):145–154. <https://doi.org/10.1016/j.strusafe.2011.01.002>
- Bichon B, Eldred M, Swiler L, Mahadevan S, McFarland J (2008) Efficient global reliability analysis for nonlinear implicit performance functions. *AIAA J* 46:2459–2468. <https://doi.org/10.2514/1.34321>
- Sun Z, Wang J, Li R, Tong C (2017) LIF: a new Kriging based learning function and its application to structural reliability analysis. *Reliabil Eng Syst Saf* 157:152–165. <https://doi.org/10.1016/j.res.2016.09.003>
- Jian W, Zhili S, Qiang Y, Rui L (2017) Two accuracy measures of the Kriging model for structural reliability analysis. *Reliabil Eng Syst Saf* 167:494–505. <https://doi.org/10.1016/j.res.2017.06.028>
- Perrin G (2016) Active learning surrogate models for the conception of systems with multiple failure modes. *Reliabil Eng Syst Saf* 149:130–136. <https://doi.org/10.1016/j.res.2015.12.017>
- Kiureghian AD, Ditlevsen O (2009) Aleatory or epistemic? Does it matter? *Struct Saf* 31(2):105–112 (**Risk Acceptance and Risk Communication**)
- Helton J, Johnson J, Oberkampf W, Sallaberry C (2010) Representation of analysis results involving aleatory and epistemic uncertainty. *Int J Gen Syst* 39:605–646
- Aven T (2014) Risk, surprises and black swans: fundamental ideas and concepts in risk assessment and risk management. *Routledge*, New York
- Aven T (2013) Practical implications of the new risk perspectives. *Reliabil Eng Syst Saf* 115:136–145
- Pearl J (1995) Causal diagrams for empirical research. *Biometrika* 82(4):669–688
- GI DNV (2017) Recommended practice: corroded pipelines DNVGL-RP-F101. DNV GL, Høvik, Norway
- Øksendal B (2003) Stochastic differential equations. *Springer*, Heidelberg, pp 65–84
- Bertsekas DP (1995) Dynamic programming and optimal control, vol 1. Athena Scientific Belmont, MA
- Wiering M, Van Otterlo M (2012) Reinforcement learning. *Adapt Learn Optim* 12(3)
- Li Q, Chen L, Tai C et al (2017) Maximum principle based algorithms for deep learning. *arXiv preprint arXiv:1710.09513*
- Sutton RS, Barto AG (2018) Reinforcement learning: an introduction, 2nd edn. The MIT Press, London
- van Hasselt H, Guez A, Silver D (2016) Deep reinforcement learning with double Q-learning. In: *Proceedings of the AAAI conference on artificial intelligence*, vol 30

35. Zaheer M, Kottur S, Ravanbakhsh S, Poczos B, Salakhutdinov RR, Smola AJ (2017) Deep sets. In: Guyon I, Luxburg UV, Bengio S, Wallach H, Fergus R, Vishwanathan S, Garnett R (eds) *Advances in neural information processing systems*. Curran Associates Inc, USA
36. Wagstaff E, Fuchs F, Engelcke M, Posner I, Osborne MA (2019) On the limitations of representing functions on sets. In: Chaudhuri K, Salakhutdinov R (eds) *Proceedings of the 36th international conference on machine learning*. *Proceedings of Machine Learning Research*, vol 97, pp 6487–6494 (**PMLR**)
37. Soelch M, Akhundov A, van der Smagt P, Bayer J (2019) On deep set learning and the choice of aggregations. In: Tetko IV, Kurková V, Karpov P, Theis FJ (eds) *Artificial Neural Networks and Machine Learning - ICANN 2019*, vol 11727, pp 444–457
38. Mnih V, Kavukcuoglu K, Silver D, Rusu AA, Veness J, Bellemare MG, Graves A, Riedmiller M, Fidjeland AK, Ostrovski G, Petersen S, Beattie C, Sadik A, Antonoglou I, King H, Kumaran D, Wierstra D, Legg S, Hassabis D (2015) Human-level control through deep reinforcement learning. *Nature* 518(7540):529–533
39. Rudloff B, Street A, Valladão DM (2014) Time consistency and risk averse dynamic decision models: definition, interpretation and practical consequences. *Eur J Oper Res* 234(3):743–750
40. Cheridito P, Stadje M (2009) Time-inconsistency of var and time-consistent alternatives. *Financ Res Lett* 6(1):40–46
41. Artzner P, Delbaen F, Eber J-M, Heath D, Ku H (2007) Coherent multiperiod risk adjusted values and Bellman's principle. *Ann Oper Res* 152(1):5–22
42. García J, Fernández F (2015) A comprehensive survey on safe reinforcement learning. *J Mach Learn Res* 16:1437–1480
43. Greinecker, M (2014) Every countably generated sigma-algebra is generated by a real random variable. *Mathematics Stack Exchange* . <https://math.stackexchange.com/q/816337>

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.