Abstract

Dynamic Latent Spaces with Statistical Finite Elements
Andrew Wang, Christ's College
1 June 2022

Time-dependent physical phenomena in nature are often modelled by partial differential equations (PDEs), which are widely used in science and engineering, for example, for structural mechanics, tidal waves, heat transfer or chemical diffusion. It has become clear now that the best models of physical phenomena are constrained not only by these laws of physics, but also by observed data. However, it is still challenging to assimilate real-world data with underlying physics in a coherent, statistical way and with a quantifiable uncertainty, so that model estimates can be interpreted with risk, because of the variability and complexity of practical data.

A new data-driven approach that encapsulates inherent uncertainties in the physical world is pivotal in developing digital twins, which are making an increasing impact in a growing number of science and engineering fields, for example in structural health monitoring, climate modelling, precision medicine and agricultural monitoring. Digital twins are useful in estimating the true state of the physical twin, and forecasting their future. For this purpose, all of these digital twins applications are characterised by the need to assimilate large quantities of sensor data into a virtual, "digital twin" physics model, whether that model be of sea surface temperature, the heart, or a structural beam [21].

The statistical finite element method (statFEM) has been developed in [6] in order to tackle this problem of data assimilation. The method enhances the finite element method in solving spatiotemporal PDEs by providing a framework to assimilate noisy data in a Bayesian manner using stochastic filtering. However, with the widespread use of more advanced sensing technologies, it has become vital to achieve the same style of data assimilation for complex datasets, such as video and audio, where the mapping between data and underlying variables is non-linear, spatially dependent in 2D, and harder to estimate. Our work focuses on developing a statistical framework around statFEM for complex datasets, leveraging recent advances in machine learning and computer vision techniques. We approach this problem using representation learning; as in statFEM, the physics model is contained in a dynamic low-dimensional latent

space, but we learn the mapping to and from higher-dimensional data using deep unsupervised learning.

This project describes the use of the Variational Autoencoder (VAE) [16], a form of deep generative model, which is capable of being trained in order to probabilistically model the data likelihood and a meaningful low-dimensional latent space, using deep convolutional neural networks to model the mapping. We train and test the VAE neural networks on computer vision datasets to show that it can quickly learn a map from data to smooth latent spaces, where the coordinates of inferred points encode some information characteristic of the dataset, for example, the shape of a handwritten digit or the spatial position of an object in the image.

Building on this, we show that the Kalman VAE (KVAE) [9] allows us to model time dependency in the dataset in order to model linear dynamics. We show that by factorising the log-likelihood following [9], we can separate the problem of inferring a dynamic latent space into two parts: the standard time-independent VAE latent space inference to model a complex data-generating process, and a Kalman filter to infer sequences in the state-space in the context of a dynamical model. We note that there has been much interest in developing autoencoder-based models for dynamic latent space modelling; however, the specific use of the Kalman filter, which is a form of sequential Bayesian inference, will be of use to us below. Training on videos generated from a dynamical model, we show that the KVAE can infer a latent space similar to the generating physics, and also generate realistic observations forwards in time from the latent space.

However, for our purposes, and for digital twins applications, we do not actually want to learn the dynamical model, as was the focus in [9]. Instead, we wish to fix the dynamics to a dynamical model derived from a physical PDE. Since the dynamical model in the KVAE is ultimately a linear state-space transition model, we show that in the case of a linear PDE, we can use statFEM to solve the PDE and directly fix the parameters of the transition model, where the statFEM stochastic filter becomes the Kalman filter in the KVAE. We name our resulting *physics-informed* model the statFEM-KVAE for this reason. We train and test statFEM-KVAE using videos generated from simple linear PDEs, such as of waves following the wave equation. We show that the trained model can extract not only a meaningful, but the correct latent space, that represents a Bayesian combination of the underlying PDE variables (e.g. wave position) and the data (videos of waves moving). Our model provides a statistical framework for the motivating aims outlined above, for example for digital twins applications: it can continually estimate and calibrate a physical model given data, predict and interpolate in time and space, and return all predictions with an associated uncertainty.