Delft University of Technology

Master's Thesis Biorobotics

Department of Cognitive Robotics

Brain inspired state and input observer for a drone in wind conditions

Author:
Fred Bos

Supervisors:
Prof. Dr. Ir. M. Wisse
Ir. A. Anil Meera
Readers:
Dr. M. Kok
Dr. L. Ferranti

October 16, 2021

MSc Mechanical Engineering,
Faculty of Mechanical, Maritime and Materials Engineering
Delft University of Technology

Abstract

The free energy principle is a recent theory that originates from the neuroscience. It provides a unified framework that combines action perception and learning in the human brain. This research aims to implement the perception aspect of the free energy principle into robotics. This is achieved via the dynamic expectation maximisation (DEM) algorithm. DEM is derived from the free energy principle and provides a novel solution for state and input observation for LTI systems under the influence of coloured noise. This thesis provides an overview of how the DEM observer is derived from the free energy principle. Thereafter, an experimental design is presented in which data is collected to validate the performance of DEM on experimental data. The data is collected from a quadcopter drone flying in wind conditions. A detailed overview of the DEM observer settings is presented and motivated.

This research shows that because of its use of generalized coordinates DEM is able to leverage the coloured nature of the noise for better state estimation. This is demonstrated by the fact that DEM obtains a higher state estimation accuracy than other coloured noise state observers, such as state augmentation and SMIKF. Moreover, in input estimation, DEM is able to obtain similar results as an unknown input observer. Finally, the accuracy vs complexity trade off of DEM is highlighted.

Acknowledgements

I would like to thank my supervisors Prof. Dr. Ir. Martijn Wisse and Ir. Ajith Anil Meera for the assistance and guidance over the course of my graduation. For the past year, during the many online meetings, I could always ask questions about the theory and discuss the progress of the thesis. I would also like to thank Dennis Benders, whose work I could continue thanks to the thorough documentation on his own thesis.

Freddy Bos
October 16, 2021

Contents

List of Figures

Chapter 1: Introduction

Robots are becoming more common in the modern world, automating human labor and assisting in a wide variety of daily human tasks, for example delivery drones. However, in many fields humans still outperform robots. Using more efficient inference mechanisms in the brain, a human can adapt faster to new environments and anticipate better in unexpected conditions. A recent theory in neurosciences aims to provide a common theory that combines action, perception and learning in the human brain. This theory is the free energy principle (FEP) as presented by Karl Friston [1]. The FEP states that biological systems survive by minimizing the information entropy in the brain, in other words, minimize the surprise. The surprise is caused by a combination of the brain's model of the environment, what to expect, and the brain's sensory input, what the brain perceives. Using principles from Bayesian learning, the free energy provides an upper bound on the surprise, meaning that biological systems are constantly minimizing the free energy [2]. By minimizing the free energy humans are capable of efficient action, perception, anticipation and adaptation in their environment, features that would also be beneficial to the development of new robots.
To this end, new research is currently focused on implementing the principles of the FEP in robot control theory. This thesis focuses on one vital aspect of the robot control loop, state and input observation. For smooth and effective operation and control, a robot needs to monitor not only its working environment, but also its internal states. State and input observation is used to determine the robot's internal states while mitigating noise disturbances, and among other things, is used for action selection and fault detection. Dynamic expectation maximisation (DEM) is derived from the FEP and focuses on its perception aspect, providing a new solution to parameter, hyperparameter, state, and input estimation [3]. Moreover, because DEM operates in generalized coordinates, it is capable of leveraging temporal correlations in the noise disturbances for better estimation. Using these features from DEM, a novel state and input observer for linear time invariant (LTI) systems with coloured noise was recently presented and tested in [4]. In this research, the new observer was tested on simulated systems with coloured noise and has shown to outperform the more conventional Kalman filter.
After testing DEM's capabilities on simulated data, this thesis aims to provide an experimental conformation of these results. Moreover, instead of benchmarking the DEM state observer to only the Kalman filter, this research aims to expand the list of benchmarks with state observers that are capable of dealing with coloured noise. In order to validate the performance of DEM on experimental data an experiment is required, in which an LTI system is influenced by coloured noise. In a previous thesis such an experiment was already designed [5]. In the experiment a quadrotor drone was tasked to hover in wind, provided by a blower. The aim was to provide the system with a coloured noise disturbance caused by the wind. The data that was collected from these experiments is used in this thesis to validate the experimental design for DEM and to test DEM's state and input estimation capabilities. Moreover, the goal is to present the results so that they can be published in a paper. Therefore, the main goal of the thesis is:

Use experimental data to validate DEM state and input estimation performance for a system under the influence of coloured noise and present the results in a publishable paper.

To accomplish the main goal, five sub-goals ave been defined:

Validate that the experimental data agrees with the noise and model assumptions in the DEM framework, such as the Laplace assumption and the coloured noise model.

Provide experimental validation of the usefulness of generalized coordinates in state estimation for coloured noise.

Test DEM state and input estimation on experimental data and compare its estimation accuracy against conventional observers for coloured noise such as state augmentation, SMIKF and an unknown input observer.

Demonstrate the influence of priors in the accuracy vs complexity trade of of DEM.

With these goals in mind, the data from the experiments were used to obtain the thesis results. The results have been submitted in the form of a paper [6]. This thesis highlights the underlying theory of DEM state estimation and describes the process of obtaining the results from the data. The thesis is built up in several chapters. First, in chapter 2, the problem of state and input observation in LTI systems is formulated. After that, the works that preceded the paper and the thesis will be summarized in chapter 3. These consist of the previous works on DEM state estimation for LTI systems and a summary of the literature study into different state observers for coloured noise. In chapter 4, the theory and preliminaries from the paper will be revisited and explained more thoroughly. This includes an elaboration on DEM for LTI systems and how this is derived from the FEP. Furthermore, the benchmarks will be briefly explained. After these chapters, the core of the thesis consists of the experimental setup, of which mainly the data processing is important, and the explanation of the different parameters that are used in the observers. In chapter 6 the motivation for the observer settings as chosen in the paper will be motivated. In chapter 7 the results, as given in the paper, will be presented together with some additional results and expanded graphs. Finally, the conclusion of the thesis is presented alongside the recommendations for future research.

Chapter 2: Problem Statement

In this research the main goal is to perform state and input estimation for LTI systems influenced by coloured noise. Therefore, the drone's dynamics are modelled as an LTI system by the state space equations in equation 2.1.

$$\dot{\mathbf{x}} = A\mathbf{x} + B\mathbf{v} + \mathbf{w}, \qquad \mathbf{y} = C\mathbf{x} + \mathbf{z}. \tag{2.1}$$

In the equation, A, B and C are the constant system matrices that represent the drone's dynamics and map the states to the system's output. $\mathbf{x}$ contains the hidden states, $\mathbf{v}$ is the input and $\mathbf{y}$ represents the output. The process and measurement noises are represented by $\mathbf{w}$ and $\mathbf{z}$, respectively. The state space system is represented in a block scheme in figure 2.1 below.

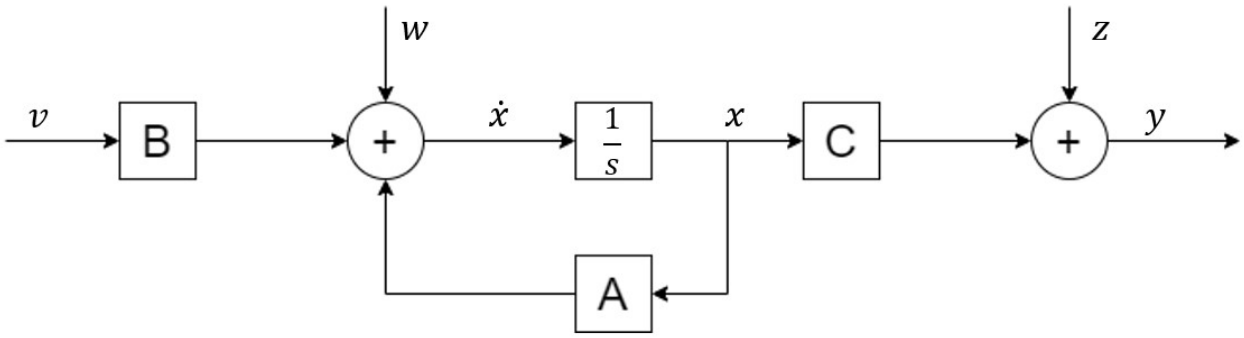

Figure 2.1: A block scheme representing the state space system as used for the drone model.

The process noise has a direct influence on the system's hidden states, and is of the same dimension as $\mathbf{x}$. Process noise is often caused by modelling errors in the system or by an unmodelled external dynamic system acting on the drone. The measurement noise, which is of the same dimension as $\mathbf{y}$, is added to the measurements and is often caused by faults in the measurement equipment and the measurement resolution. In the experiments, the goal is to introduce a large coloured noise disturbance in the form of wind. Therefore, in this research, the noises are assumed as coloured. The goal of DEM is to estimate the hidden states $\mathbf{x}$ and the inputs $\mathbf{v}$ despite the influence of the process and measurement noises, $\mathbf{w}$ and $\mathbf{z}$.

Chapter 3: Related Work

3.1 Free Energy Principle

The FEP is a framework, that finds its origin in neuroscience and proposes a model based inference scheme, explaining action, perception and learning in the brain [2]. It unifies multiple brain theories by introducing the free energy quantity in a mathematical environment, enabling a bridge between neuroscience and robotics [1]. The FEP states that, in order to exist in their environment, biological systems resist their natural tendency to disorder [7]. In neuroscience, this signifies that the brain tries to minimize its information entropy, or surprise. The brain can minimize surprise by interacting with its environment, either through perception and learning, or action. The perception and learning aspects of the FEP are described in the DEM algorithm and allow for a transition from neuroscience to state, parameter and hyperparameter estimation in control systems. The parameter estimation capabilities of DEM were tested in [8], while DEM was transformed to a state and input observer for LTI systems influenced by coloured noise in [3, 4].

3.2 State observers for coloured noise

In literature, noise is often modelled as white noise [9], for which the optimal solution is the Kalman filter [10]. However, when the white noise assumption is violated, the Kalman filter is no longer optimal. To accommodate fair comparison in benchmarking, DEM should be compared to other state observers for coloured noise. Therefore, prior to this thesis, a literature research was performed into the different state observation methods for coloured noise [11]. This section will provide a short overview of the literature review, while in chapter 4, the theory behind the benchmarks will be shortly revisited.

The goal of the literature research was to find the most accurate state observer for LTI systems under the influence of coloured noise, thereby providing an overview of the different solutions for coloured noise state estimation in literature. This resulted in the list of methods as given below:

- **State Augmentation**
  The state augmentation (SA) method, models the process noise as additional states by expending the state space system with the coloured portions of the noise. By eliminating the coloured portion of the noise, the system is transformed to an equivalent higher order system influenced by white noise. Therefore, the Kalman filter can be used for state estimation on the higher order model [12].

- **Measurement Differencing**
  Measurement differencing uses multiple subsequent measurements to eliminate the coloured fraction of the measurement noise in a system, resulting in a system under the influence of white noise. An altered version of the Kalman filter can be used for state estimation on the system with its new pseudo measurements [13].

- **Second Moment Information Kalman Filter**
  In SMIKF, the Kalman filter covariance update equations are expanded to accommodate

for the time correlation in the noise. This transforms the Kalman filter into an observer capable of dealing with coloured noise [14].

- **Adaptive Kalman Filtering**
  In adaptive Kalman filtering the Kalman filter's covariance matrices are constantly updated using the latest measurements. This results in better state estimation because noise modelling errors are mitigated [15].

Based on the literature study, SA and SMIKF are chosen as the benchmarks. The main reason is the fact that, because of the wind in the experimental setup, the process noise is the dominant noise factor. SA and SMIKF were the only two methods that are inherently capable of handling coloured process noise. Measurement differencing only affects coloured measurement noise and adaptive Kalman filtering only mitigates the unmodelled coloured noise by constantly updating the covariance matrices. So in conclusion, SA and SMIKF will be used to benchmark DEM. Finally the Kalman filter will also be added as a benchmark to include a white noise state observer as baseline.

Chapter 4: Theory

This chapter will discuss the theory that is used throughout the thesis. First, the different noise models that can be used to model the process and measurement noise will be highlighted. After that, the fundamentals of DEM and how it originates from the FEP will be elaborated. Finally, the principles of the state estimation benchmarks will be explained.

4.1 Noise Modelling

The state observers in this research assume three different types of noise. The Kalman filter assumes random white noise, SA and SMIKF assume auto regressive (AR) noise and DEM assumes Gaussian convoluted noise. The working principles of each of the observers is built upon these noise assumptions and determines the observer design. Therefore, the three noise types will be highlighted in this section.

White Noise

When noise is white noise it means that the noise is a random noise signal. This means that while the noise is zero mean and has a determined variance, the signal at each timestep is completely independent from the signal in any previous timestep. This means there is no temporal correlation in the noise. Throughout this thesis, noise colour will often be determined via its autocorrelation. In the case of white noise the autocorrelation immediately drops to zero. For visualisation, a white noise signal and its autocorrelation are plotted in figure 4.1.

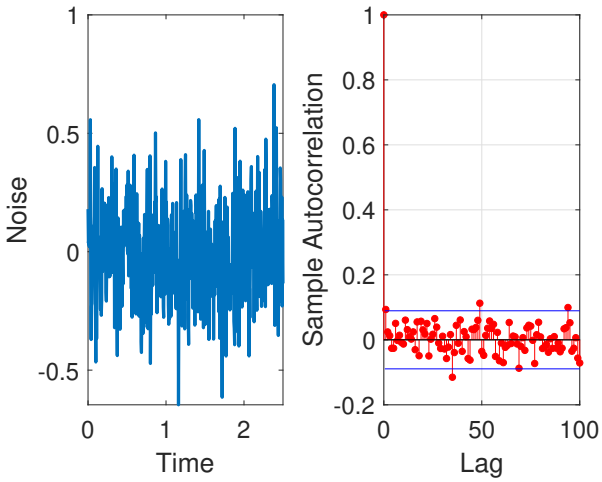

Figure 4.1: A simulated example of white noise together with its autocorrelation function.

It can be clearly observed from the figure that the autocorrelation immediately drops to zero, indicating that the noise is white.

Gaussian Convoluted Noise

In contrast to white noise, Gaussian convoluted noise does contain a degree of temporal correlation between its samples and is therefore coloured. In the rest of this thesis, Gaussian

convoluted noise will be referred to as Gaussian noise for brevity. Gaussian noise is obtained when convoluting white noise, with zero mean and a given variance, with a Gaussian kernel of the form:

$$K(t) = \frac{1}{\sqrt{2\pi}\sigma} exp(-\frac{1}{2}(\frac{t}{\sigma})^2) \tag{4.1}$$

This convolution smooths the white noise and results in a coloured noise signal. The degree of smoothness, or colour, is dictated by the noise's smoothness value σ from equation 4.1. The t in the equation represents the time. A visualisation of a Gaussian noise signal is given in figure 4.2 below.

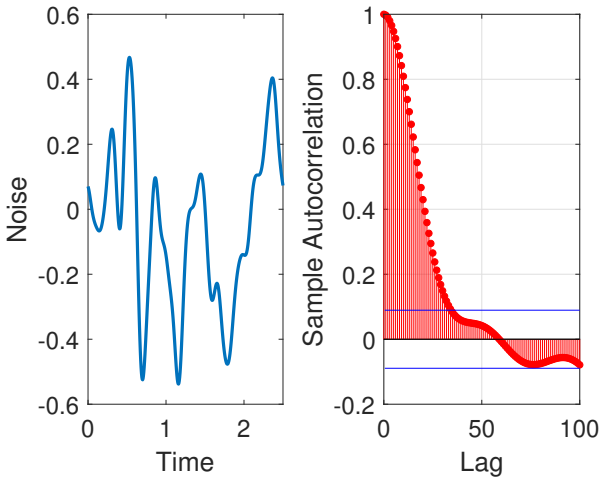

Figure 4.2: A simulated example of Gaussian convoluted noise together with its autocorrelation function.

From the figure it can be clearly observed that in contrast to the white noise in figure 4.1, the noise appears as a more predictable signal. Moreover, in the case of the Gaussian noise, the autocorrelation follows a slowly decaying Gaussian curve, indicating that the noise is coloured. The coloured Gaussian noise is the noise model that is used in DEM.

Auto Regressive Noise

Auto regressive (AR) noise is formed by shaping a white noise through a linear system. In the discrete domain this is governed by equation 4.2.

$$w_{k+1} = \sum_{i=0}^{m-1} \phi_i w_{k-i} + \omega_k \tag{4.2}$$

Here, w_k represents the noise, m the order of the AR model and ϕ the AR coefficients. ω is a white noise signal, that has zero mean and a given variance, that is shaped through the AR system. The degree of colour in the AR noise is determined by the AR coefficients ϕ, that link the current noise sample to previous samples. An AR noise signal and the autocorrelation of the signal have been visualized in the figure below.

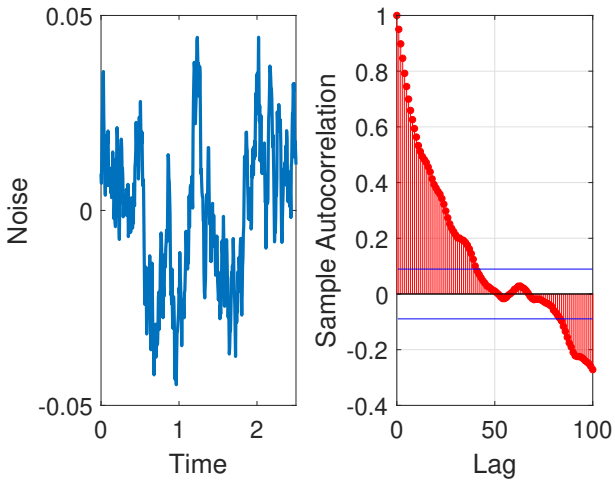

Figure 4.3: A simulated example of Auto regressive noise together with its autocorrelation function.

Similar to the Gaussian noise it can be observed that the signal does not appear as random as the white noise signal. The autocorrelation in the figure is slowly decaying as well, which means that the noise is coloured. In this figure, the AR model is a first order AR model where $\phi = 0.96$. This means that the autocorrelation is decreasing with a factor of 0.96 for each step. Both the SA and SMIKF benchmarks use an AR model to model coloured noise. The AR coefficients from the modelled noise are used in the observer designs of both benchmarks.

4.2 Dynamic Expectation Maximisation

As discussed in section 3.1, DEM describes the perception and learning aspects of the FEP. This section will briefly explain how DEM is derived from the FEP and how this results in a state and input observer for coloured noise, based on the theory from [3, 4].

4.2.1 Variational Free Energy

As previously stated, when considering the perception aspect of the FEP, the brain minimizes the surprise it experiences by minimizing the entropy between its model of the environment and its sensory input. The model of the environment is given as m and the internal states are described by ϑ. ϑ consists of the system's states, inputs, parameters and hyperparameters. This research focuses on just the states and inputs. The sensory data from the environment is depicted with y. In essence, the brain has to determine ϑ based on m and y. In Bayesian inference this is described by the posterior density $p(\vartheta|y,m)$. However, calculating the posterior density is intractable. Therefore, in the brain, the posterior density is approximated by the recognition density $q(\vartheta)$. A measure of similarity between the two probability densities is provided by the Kullback-Leibler divergence D_{KL} as given in equation 4.3.

$$D_{KL}(q(\vartheta)||p(\vartheta|y,m)) = \langle \ln q(\vartheta) \rangle_{q(\vartheta)} - \langle \ln p(\vartheta|y,m) \rangle_{q(\vartheta)} \tag{4.3}$$

Where the angle brackets denote the expected value under a given density. The equation can be simplified to:

$$D_{KL}(q(\vartheta)||p(\vartheta|y,m)) = \langle \ln q(\vartheta) \rangle_{q(\vartheta)} - \langle \ln p(y,\vartheta) \rangle_{q(\vartheta)} + ln(p(y|m)) \tag{4.4}$$

Rearranging the terms leads to an expression for the free energy as given in equation 4.5.

$$\begin{aligned} ln(p(y|m)) &= F + D_{kl}(q(\vartheta)||p(\vartheta|y,m)) \\ F &= G + H = \langle \ln p(y,\vartheta)\rangle_{q(\vartheta)} - \langle \ln q(\vartheta)\rangle_{q(\vartheta)} \end{aligned} \tag{4.5}$$

In this equation, F is the free energy and $ln(p(y|m))$ is the log-evidence of the model. G is the expectation of the internal energy $U = \ln p(y,\vartheta)$ and H is the entropy of the recognition density. Because $ln(p(y|m))$ is independent of $q(\vartheta)$ and the D_{KL} is always positive by definition, minimizing the D_{KL} between the posterior and the recognition density is equivalent to maximizing the free energy F. The problem of matching the recognition density $q(\vartheta)$ to the posterior density $p(\vartheta|y,m))$ has now changed to maximizing the free energy F. Moreover, because only the states and inputs are considered, the term H becomes independent of the parameters ϑ [3]. Therefore, maximizing F is equivalent to maximizing the expectation of the internal energy also called the variational free energy (VFE), which is given as $V(\vartheta) = \langle U(\vartheta,y)\rangle_{q(\vartheta)}$. In DEM, the VFE is maximized by performing a gradient ascent over the VFE for the parameters ϑ, which in this research consists of the state and inputs.

In summary, the problem of obtaining the intractable posterior density is now simplified to an optimisation problem in the form of a gradient ascent over the VFE.

4.2.2 Generative model

In the FEP, the brain's model of the environment is also called the generative model. In the case of the drone, the generative model is the state space model that describes its dynamics. However, in the FEP framework this model is slightly altered by the incorporation of generalized coordinates. The use of generalized coordinates also provides DEM with one of its main strengths, the capability to deal with coloured noise. With generalized coordinates, not only the hidden states but also their derivatives are modelled. This means the dynamics are not only modelled for a single state space system as depicted in chapter 2, but also for the derivatives of the states, inputs, outputs and the noise. The generative model in generalized coordinates is modelled as follows:

$$\begin{aligned} x' &= Ax + Bv + w & y &= Cx + z \\ x'' &= Ax' + Bv' + w' & \dot{y} &= Cx' + z' \\ &... & &... \end{aligned} \tag{4.6}$$

In these equations the $'$ operator denotes the derivative of the variable. The system can be compactly written, by use of the tilde operator, as:

$$\begin{aligned} \dot{\tilde{x}} = D_x\tilde{x} &= \tilde{A}\tilde{x} + \tilde{B}\tilde{v} + \tilde{w} \qquad \tilde{y} = \tilde{C}\tilde{x} + \tilde{z} \\ \text{where } D_x &= \begin{bmatrix} 0 & 1 & & & \\ & 0 & 1 & & \\ & & . & . & \\ & & & 0 & 1 \\ & & & & 0 \end{bmatrix}_{(p+1)\times(p+1)} \otimes I_{n\times n} \end{aligned} \tag{4.7}$$

In this equation, D_x represents the shift matrix, which performs the derivative operation on the generalized state vector $\tilde{x}$. The matrix D_v does the same for the inputs and is of size $r(d+1) \times r(d+1)$. p and d represent the embedding order for the hidden states and the input respectively, whereby the embedding order determines the amount of derivatives per state. The generalized system matrices, $\tilde{A}$, $\tilde{B}$, and $\tilde{C}$ are of the form:

$$\tilde{A} = I_{p+1} \otimes A, \quad \tilde{B} = I_{p+1} \otimes B, \quad \tilde{C} = I_{p+1} \otimes C \tag{4.8}$$

These equations transform the system matrices from the state space system in equation 2.1 to generalized system matrices. In equations 4.7 and 4.8, I denotes the identity matrix and $\otimes$ the Kronecker tensor product. The result of a system in generalized coordinates is that more dynamic information is considered in state and input estimation. This includes the derivative information of the noise. This gives DEM its capability of leveraging coloured noise for better state and input estimation. Highlighting the usefulness of generalized coordinates is also one of the main goals of this research.

4.2.3 Noise modeling

To accommodate the use of generalized coordinates the derivatives of the noise have to be well defined. As discussed in section 4.2.3, the noise in the FEP is modelled as Gaussian noise. By modelling the noise as Gaussian noise, it is possible to obtain the covariance for all noise derivatives. This is accomplished through the use of the temporal precision matrix $S(\sigma^2)$ as given in equation 4.9, where precision is defined as the inverse of the noise variance.

$$S(\sigma^2) = \begin{bmatrix} 1 & 0 & -\frac{1}{2\sigma^2} & .. \\ 0 & \frac{1}{2\sigma^2} & 0 & .. \\ -\frac{1}{2\sigma^2} & 0 & \frac{3}{4\sigma^4} & .. \\ .. & .. & .. & .. \end{bmatrix}^{-1}_{(p+1)\times(p+1)} \tag{4.9}$$

in this matrix the value σ is the aforementioned smoothness value that determines the degree of colour in the noise. Using S, the generalized precision matrices for the noise can be defined as:

$$\tilde{\Pi}_w = S \otimes \Pi_w, \quad \tilde{\Pi}_z = S \otimes \Pi_z, \quad \tilde{\Pi}_v = S \otimes \Pi_v \tag{4.10}$$

4.2.4 Generalized output

The generalized output $\tilde{y}$ from equation 4.7 is often unobtainable in practical applications since it is often impossible to measure all derivatives of a system's output. This obstacle is overcome by using multiple discrete measurements $\hat{y}$ to calculate the generalized outputs, using a Taylor expansion as described in equation 4.11.

$$\hat{y} = \begin{bmatrix} \cdots \\ y(t-dt) \\ y(t) \\ y(t+dt) \\ \cdots \end{bmatrix}_{m(p+1)} = (E \otimes I_m)\tilde{y} \tag{4.11}$$

Where E is calculated as:

$$E_{ij} = \frac{[(i - ceil(\frac{p+1}{2}))dt]^{j-1}}{(j-1)!} \tag{4.12}$$

Here, i and j are integers from 1 to $p+1$ and *ceil* denotes that the value between brackets is rounded up to the first integer. In DEM, the generalized output is calculated by reversing the Taylor-expansion, resulting in the following pre-processing step:

$$\tilde{y} = (E^{-1} \otimes I_m)\hat{y} \tag{4.13}$$

The same principle is used to embed the higher orders of the generalized input $\tilde{v}$. So by using equation 4.13 it is possible to obtain the higher order derivatives of the inputs and outputs, at the cost of a time-delay.

4.2.5 DEM for LTI systems

As discussed in section 4.2.1 DEM performs a gradient ascent over the VFE. Since this research focuses on state and input estimation, the parameters ϑ consist of the states and inputs, denoted by $X = [\tilde{x}\,\tilde{v}]^T$. For X, the gradient ascent takes the following form [4]:

$$\dot{X} = kV(t)_X + DX \tag{4.14}$$

In this equation, $V(t)_X$ is the derivative of $V(t)$ with respect to X and D is the shift matrix containing both D_x and D_v. The VFE, $V(t)$, is obtained by using the Laplace assumption for the internal energy $U(y, \vartheta)$.

Laplace Approximation

The Laplace approximation is used to provide a way to calculate the internal energy and thereby provides an expression for the VFE. As discussed in section 4.2.1 the internal energy is given as $U = \ln p(\tilde{y}, \vartheta)$. This can be expanded to: [8]

$$\ln p(\tilde{y}, \vartheta) = \ln p(\tilde{y}, \tilde{x}, \tilde{v}) = \ln p(y|\tilde{x}, \tilde{v}) + \ln p(\tilde{x}|\tilde{v}) + \ln p(\tilde{v}) \tag{4.15}$$

The Laplace approximation is used to approximate these probability density functions as Gaussian distributions. These are given in the three equations below:

$$p(\tilde{y}|\tilde{x}, \tilde{v}) = \frac{1}{\sqrt{(2\pi)^{n_y(p+1)}|\tilde{\Sigma}_z|}} e^{-\frac{1}{2}(\tilde{y}-\tilde{C}\tilde{x})^T\tilde{\Pi}_z(\tilde{y}-\tilde{C}\tilde{x})} \tag{4.16}$$

$$p(\tilde{x}|\tilde{v}) = \frac{1}{\sqrt{(2\pi)^{n_x(p+1)}|\tilde{\Sigma}_w|}} e^{-\frac{1}{2}(D_x\tilde{x}-\tilde{A}\tilde{x}-\tilde{B}\tilde{v})^T\tilde{\Pi}_w(D_x\tilde{x}-\tilde{A}\tilde{x}-\tilde{B}\tilde{v})}, \tag{4.17}$$

$$p(\tilde{v}) = \frac{1}{\sqrt{(2\pi)^{n_v(p+1)}|\tilde{\Sigma}_v|}} e^{-\frac{1}{2}(\tilde{v}-\eta^{\tilde{v}})^T\Pi_{\tilde{v}}(\tilde{v}-\eta^{\tilde{v}})}, \tag{4.18}$$

In these equations n_y, n_x and n_v are the dimensions of the output, state and input vectors respectively. $\tilde{\Sigma}$ denotes the generalized covariance matrix for each of the noises respectively. Here $\tilde{\Sigma}$ is the inverse of the precision matrix given in section 4.2.3. Finally, $\eta^{\tilde{v}}$ is introduced as the prior on the inputs. Because the internal energy consists of the natural logarithms of each of these equations, the exponential will disappear from the equations. Moreover, because the VFE is differentiated in equation 4.14, the constant terms in front of the exponential drop out. Therefore, the Laplace approximation transforms the VFE into a comprehensive definition when considering DEM state and input estimation for LTI systems, given in the equation below:

$$V(t) = -\frac{1}{2}\tilde{\epsilon}^T\tilde{\Pi}\tilde{\epsilon} \tag{4.19}$$

Where $\tilde{\epsilon}$ is the error between the sensory input and the generative model, given as:

$$\tilde{\epsilon} = \begin{bmatrix} \tilde{y} - \tilde{C}\tilde{x} \\ \tilde{v} - \eta^{\tilde{v}} \\ D_x\tilde{x} - \tilde{A}\tilde{x} - \tilde{B}\tilde{v} \end{bmatrix} \tag{4.20}$$

The derivative of V with regard to X is:

$$V(t)_X = -\tilde{\epsilon}_X^T\tilde{\Pi}\tilde{\epsilon} \tag{4.21}$$

Where $\tilde{\epsilon}_X$ is given as:

$$\tilde{\epsilon}_X = \begin{bmatrix} -\tilde{C} & O \\ O & I \\ D_x - \tilde{A} & -\tilde{B} \end{bmatrix} \tag{4.22}$$

By substituting these equations into equation 4.14, the linear state and input observer that results is: [4]

$$\dot{X} = \begin{bmatrix} \dot{\tilde{x}} \\ \dot{\tilde{v}} \end{bmatrix} = \begin{bmatrix} A_1 \\ A_2 \end{bmatrix} \begin{bmatrix} \tilde{x} \\ \tilde{v} \end{bmatrix} + \begin{bmatrix} B_1 \\ B_2 \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{y}} \\ -\eta^{\tilde{v}} \end{bmatrix} \; and \; Y = X, \tag{4.23}$$

where Y is the output of the observer and,

$$A_1 = \begin{bmatrix} D_x & O \end{bmatrix} - k_x \begin{bmatrix} \tilde{C}^T \tilde{\Pi}_z \tilde{C} + D_a^T \tilde{\Pi}_w D_a & -D_a^T \tilde{\Pi}_w \tilde{B} \end{bmatrix}$$

$$A_2 = \begin{bmatrix} O & D_v \end{bmatrix} - k_v \begin{bmatrix} -\tilde{B}^T \tilde{\Pi}_w D_a & \tilde{B}^T \tilde{\Pi}_w \tilde{B} + \tilde{\Pi}_v \end{bmatrix}$$

$$With, D_a = D_x - \tilde{A}$$

$$B_1 = k_x \begin{bmatrix} \tilde{C}^T \tilde{\Pi}_z & O \end{bmatrix}, \; B_2 = k_v \begin{bmatrix} O & -\tilde{\Pi}_v \end{bmatrix}$$

This is the DEM state and input observer for LTI systems under the influence of coloured noise. The observer was already tested on simulated systems [4], where it was compared to the Kalman filter. Throughout the rest of this thesis, this observer will be tested on experimental data and compared to other benchmarks for coloured noise.

4.3 Observer Benchmarks

This section will briefly highlight the theory behind the three benchmark state observers that are used to test DEM.

Kalman Filter

One of the most widely used state observers is the Kalman filter [10]. The Kalman filter uses the LTI system from equation 2.1 in combination with a white noise model to obtain an estimate for the hidden states. The covariance matrices of the process and measurement noise are denoted by Q and R respectively. The observer determines the states by a weighted average between a prior prediction of the estimate $\hat{x}^-$ and the systems output y. The weighted average is dictated by the Kalman gain K which is calculated from the prior covariance P^-. The Kalman gain is obtained by minimizing the covariance of the squared error of the posteriori value. The result of minimizing the square of the error for the Kalman gain, is that the Kalman filter provides the optimal solution [16].

The result is a state observer consisting of two steps, a predict step and an update step. The algorithm can be initiated after providing the priors for Q, R, $\hat{x}_0$ and P_0. The equations are

presented in the formulas below [17]:

$$\begin{aligned}
&\text{Predict step:}\\
&\quad \hat{x}_{k+1}^{-} = A\hat{x}_k + Bu_k\\
&\quad P_{k+1}^{-} = AP_kA^T + Q\\
&\text{Update step:}\\
&\quad K_k = P_k^{-}C^T(CP_k^{-}C^T + R)^{-1}\\
&\quad \hat{x}_k = \hat{x}_k^{-} + K_k(y_k - C\hat{x}_k^{-})\\
&\quad P_k = (I - K_kC)P_k^{-}
\end{aligned} \tag{4.24}$$

This is the conventional Kalman filter for state observation. As discussed, this filter is optimal if the white noise assumption is correct. However, most practical applications do not satisfy the white noise model assumption.

State Augmentation

The Kalman filter is often adapted to coloured process noise with the state augmentation (SA) method [16, 12, 14]. In this method, the coloured noise is modelled by an AR system as described in section 4.2.3. If the process noise is modeled as in equation 4.2 the AR coefficients can be incorporated into the system matrix A, thereby transforming the system to an equivalent system with a larger state vector. The state vector is now expanded to also track the coloured portion of the noise, modelled by the AR coefficients. The state vector now becomes $x_a = \begin{bmatrix} x & w \end{bmatrix}^T$. The system is then expended to:

$$\begin{bmatrix} x_{k+1} \\ w_{k+1} \end{bmatrix} = \begin{bmatrix} A & I \\ 0 & \phi \end{bmatrix} \begin{bmatrix} x_k \\ w_k \end{bmatrix} + \begin{bmatrix} 0 \\ \omega_k \end{bmatrix} \tag{4.25}$$

And the measurement is given as:

$$y_k = \begin{bmatrix} C & 0 \end{bmatrix} \begin{bmatrix} x_k \\ w_k \end{bmatrix} + z_k \tag{4.26}$$

This means that, in essence, the system has been transformed into an equivalent higher order LTI system under the influence of Gaussian white noise ω, of the form:

$$\begin{aligned}
x_{k+1}^a &= A^a x_k^a + \omega_k^a\\
y_k &= C^a x_k^a\\
&\text{with}\\
x^a = \begin{bmatrix} x \\ w \end{bmatrix} A^a = \begin{bmatrix} A & I \\ 0 & \phi \end{bmatrix} &\omega^a = \begin{bmatrix} 0 \\ \omega \end{bmatrix} C^a = \begin{bmatrix} C & 0 \end{bmatrix}
\end{aligned} \tag{4.27}$$

The process noise covariance matrix then becomes:

$$Q^a = \begin{bmatrix} O & O \\ O & Q \end{bmatrix} \tag{4.28}$$

Where Q is the original covariance matrix of the process noise. Because the new system is only subjected to white noise, it is now possible to use the conventional Kalman filter. When the AR model is of higher order, the same principle is used to add more states to the system. When augmenting the system with the measurement noise, the result is that the measurement noise covariance R becomes zero. This results in ill-conditioned Kalman filter equations. However, in this research the experimental setup ensures the measurement noise is of much smaller magnitude than the process noise, mitigating this problem.

SMIKF

Apart from SA, the literature research also provided the second moment information Kalman filter (SMIKF) as a viable option for state estimation for systems influenced by coloured noise. For this research, only the variant of SMIKF that models coloured process noise is implemented. The main principle behind SMIKF is that the temporal correlation of the noise is included in the Kalman filter's covariance calculation. With SMIKF, this correlation is modelled by a first order AR noise model in which the colour is governed by the AR parameter ϕ. Using this definition of the process noise, the predict step for the covariance evolves into equation 4.29 below.

$$\begin{aligned} P_k^- &= E[\tilde{x}_k \tilde{x}_k^T] \\ &= E[(A(x_{k-1} - \hat{x}_{k-1}^-) + w_k)(...^T)] \\ &= AE[(x_{k-1} - \hat{x}_{k-1}^-)(x_{k-1} - \hat{x}_{k-1}^-)^T]A^T + AE[\hat{x}_{k-1} w_k] + E[w_k \hat{x}_{k-1}^T]A^T + E[w_k w_k^T] \\ P_k^- &= AP_{k-1}A^T + AP_{\hat{x}_{k-1}w} + P_{w\hat{x}_{k-1}}A^T + P_w \end{aligned} \tag{4.29}$$

The covariances between the noise and the hidden states are defined in the following equations:

$$\begin{aligned} P_{\hat{x}_{k-1}w} &= (I - K_{k-1}C)P_{w_{k-1}w_k} \\ P_{w\hat{x}_{k-1}} &= P_{\hat{x}w_{k-1}}^T \end{aligned} \tag{4.30}$$

Where K_{k-1} is the Kalman gain of the previous timestep and $P_{w_{k-1}w_k}$ is the covariance between the noise in the current timestep and the noise in the previous timestep given by equation 4.31 below.

$$P_{w_{k-1}w_k} = \phi P_{w_{k-1}} \tag{4.31}$$

$P_{w_{k-1}}$ is the variance of the process noise at the previous time-step. Finally P_{w_k} is given as the noise variance at the current timestep.

$$P_w = \phi^2 P_{w_{k-1}} + Q \tag{4.32}$$

These equations redefine the predict step of the Kalman filter for a model with coloured process noise. Since only the coloured process noise is considered for this research, the update equations of the Kalman filter remain unchanged.

4.4 Overview of the Theory

In this chapter, the preliminaries and the theory of this research have been elaborated. First, the three different noise models that are used by DEM and its benchmarks were discussed, which consisted of white noise, Gaussian noise and AR noise. Second, a short overview of how the DEM state and input observer is derived from the FEP was given. By making use of generalized coordinates and the Laplace approximation, the result is a gradient ascent over the variational free energy for the states and inputs and a state and input observer for LTI systems under the influence of coloured noise. In section 4.3, the observers that will be used to benchmark DEM for state estimation, the Kalman filter, SA and SMIKF, were briefly highlighted. One final benchmark that is missing is the unknown input observer (UIO) that is used for input estimation. This benchmark was implemented in [4] and originated from [18]. In this research this implementation is used to provide insight into the input estimation accuracy of DEM.

Chapter 5: Experimental Design

This chapter presents an overview of the experiments in this research. First, the setup of the experiments and measurements is explained. Second, the selection of the drone's dynamic model is motivated. Finally, the data selection and preparation are elaborated.

5.1 Experimental Setup

The goal of the experiment was to obtain data from a quadrotor drone, which could be analyzed and used for state and input estimation. The drone that was used for these experiments is a Parrot AR.drone 2.0, shown in figure 5.1a. The data was obtained from the experiments as performed in [5]. This section will provide a short overview of the setup of these experiments.

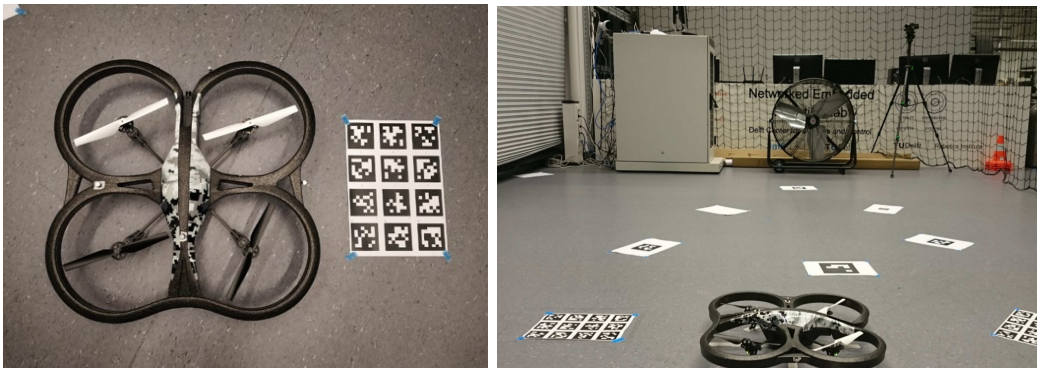

(a) A picture of the AR drone 2.0 and a space saving alignment.

(b) The lab setup with the drone and the blower.

Figure 5.1: Pictures of the drone and the lab setup.

In the experiments, the drone was tasked to keep a constant position while hovering in the air. While keeping a constant position and orientation, the drone's movements were measured both by the motion capture OptiTrack system and by its own internal IMU. The OptiTrack system was capable of recording accurate position and orientation measurements, while the IMU recorded angular velocity and acceleration. A blower was placed in the y direction of the drone in order to provide the wind influence to the system. A picture of the lab setup is given in figure 5.1b. With this setup nine experiments were conducted. Five experiments were conducted in which the drone was hovering in wind conditions, supplied by the blower, and four experiments were conducted while there was no wind. The data from these experiments, both the OptiTrack data and the data internally recorded in the drone, was saved in ROS bag-files.

5.2 Model Selection

The goal of the drone's model selection is to provide a simple dynamic model which accommodates a large coloured noise influence. Therefore, an LTI white box model was selected. The linearized white box model results in unmodelled dynamics of the nonlinear drone system which, along with the wind influence, will be assumed as part of the coloured noise. Since the roll angle around the x-axis is perpendicular to the wind direction, this angle experiences the largest wind

disturbance. Therefore, the system will be limited to the roll angle and roll rate, ϕ and $\dot{\phi}$. A linearized state-space model for the drone's roll angle and rate is given in equation 5.1.

$$\begin{bmatrix} \dot{\phi} \\ \ddot{\phi} \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \phi \\ \dot{\phi} \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 & 0 \\ \frac{c_{B\phi}}{I_{xx}} & -\frac{c_{B\phi}}{I_{xx}} & -\frac{c_{B\phi}}{I_{xx}} & \frac{c_{B\phi}}{I_{xx}} \end{bmatrix} \begin{bmatrix} pwm_1 \\ pwm_2 \\ pwm_3 \\ pwm_4 \end{bmatrix}, \tag{5.1}$$
$$y = \begin{bmatrix} 1 & 0 \end{bmatrix} \begin{bmatrix} \phi \\ \dot{\phi} \end{bmatrix}$$

In this equation, the *pwm* values are the motor Pulse Width Modulation values that are applied to the drone as input. I_{xx} is the drone's moment of inertia around the x-axis and $c_{B\phi}$ is a thrust coefficient that models the relation between the *pwm* values and the thrust, generated by the drone's rotors. In this experiment, the value for $c_{B\phi}/I_{xx}$ is 0.3748.
In order to perform the noise analysis for the experiments, the process noise has to be isolated. In this research this is possible because of the small magnitude of the measurement noise compared to the process noise. The following formula is used to determine the process noise that acts on the drone:

$$w_k = x_{k+1} - A_d x_k - B_d v_k \tag{5.2}$$

In this equation, w_k is the process noise vector, A_d and B_d are the discretized system matrices and x_k and v_k are the discrete measurements of the states and inputs.

5.3 Data Selection and preparation

The data is extracted from the ROS bags in Matlab. The data consists of several topics that were recorded from ROS. The OptiTrack data contains the OptiTrack positions and orientations, the NavData contains the linear velocities and the *pwm* motor values for the system's input. Finally, the IMU data contains the angular velocities. In Matlab, the data is formatted into data files, that can be used in the state and input observers, in the following steps.

- Because the OptiTrack system starts recording earlier than the drone, the first step is to link the correct OptiTrack start time to the start time of the drone.
- The second step is to select the correct time frames for each of the wind modes. This is accomplished by visually inspecting the data and cutting out the correct intervals, based on a combination of orientation, rotation, velocity and wind speed.
- The OptiTrack is recorded in 120 Hz while the Ardrone data is recorded in 200 Hz. Therefore, the 200 Hz data is interpolated so that both datasets have the same timestep of 0.0083.
- After this step, the data is converted to the correct units, degrees to radians.
- The measurement data now contains the following quantities from the data.

OptiTrack Position	x	y	z
NavData Linear Velocity	$\dot{x}$	$\dot{y}$	$\dot{z}$
OptiTrack Orientation	ϕ	θ	ψ
IMU Angular Velocity	$\dot{\phi}$	$\dot{\theta}$	$\dot{\psi}$
NavData Motor PWM values	v		

- The data for all 9 experiments was then cut to the length of the shortest experiment, so all experiments would be equal in length. Therefore, all experiments were cut to 10 seconds, or 1200 samples.
- In order to accommodate for more accurate input estimation, the inputs are normalized by dividing them by a normalizing constant and multiplying the B-matrix with said constant. Here the constant is given as:

$$C_v = \frac{1}{max(v) - min(v)} \tag{5.3}$$

  Moreover, since the input values contained a non-zero mean, because the rotors have to counter gravity, the mean was subtracted from the inputs. This did not effect the current system, since only the roll angles are examined and not the z position or velocity.
- Since the model is restricted to only the roll angle, ϕ and $\dot{\phi}$ are isolated from the data.
- The model, as discussed in section 5.2, is added to the data files.

The result is nine data files of 1200 samples, consisting of the roll angle and rate, ϕ and $\dot{\phi}$, the input v the timestep $dt = 0.0083$, and the drone's model in the form of the A, B and C matrices. In the new data files v and B have been normalized. These data files are subsequently used to tune the observers, analyze the noise, and perform state and input estimation. The ninth experiment is used to tune the observers and the other eight, four with wind and four withoutwind, are used to obtain the results.

Chapter 6: Observer Tuning

This chapter motivates the different settings of the DEM state and input observer. Moreover, the settings of the other benchmarks are briefly discussed. The DEM setup can be schematically represented by figure 6.1.

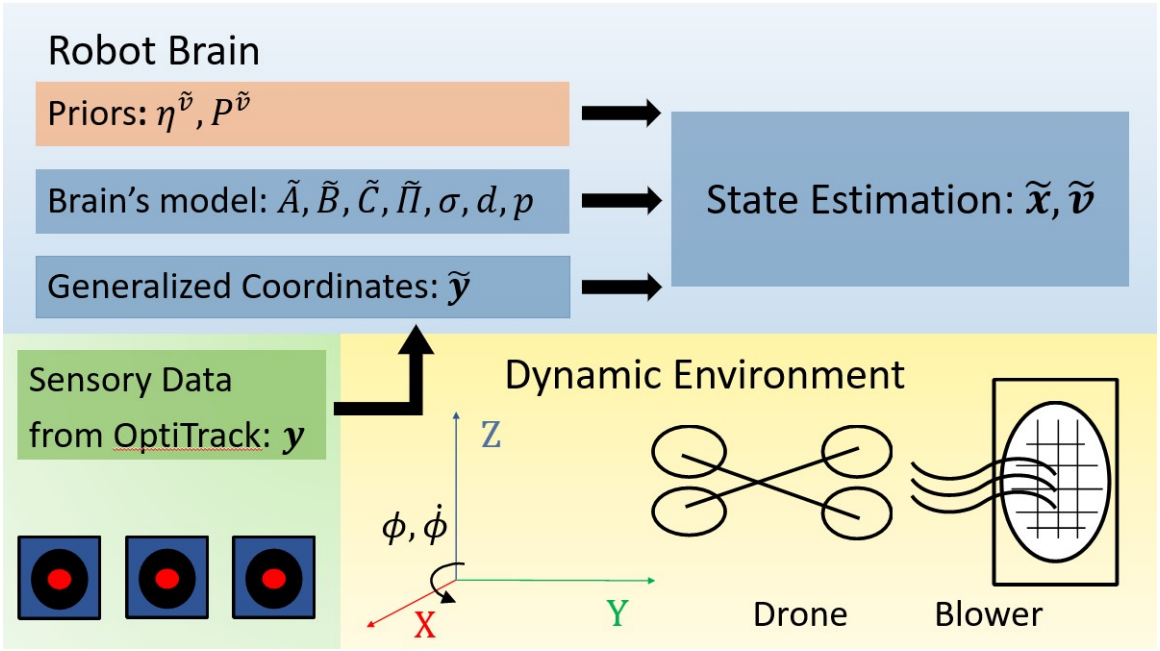

Figure 6.1: A schematic representation of DEM state and input estimation for a drone flying in wind conditions, including all quantities and signals.

The quantities that have to be supplied to the DEM observer, prior to state and input estimation are $\eta^{\tilde{v}}$, $P^{\tilde{v}}$, $\tilde{A}$, $\tilde{B}$, $\tilde{C}$, σ, d and p. The generalized system matrices are calculated using the white box model matrices as given in section 5.2, and equation 4.8. The rest of the settings will be explained in the sections below.

6.1 Input Priors

The input prior $\eta^{\tilde{v}}$ determines the prior beliefs on the inputs. For state estimation it is assumed that the inputs are known so the priors are set to the measured inputs. For input estimation, only the first input is estimated. This means that the priors for the second, third and forth input remain unchanged. For input estimation, the prior for the first input is set to an incorrect value of either 0.5 or 1.
The precision on the input priors $P^{\tilde{v}}$ determines the weight DEM gives to the priors. So when the precision is very high, DEM will focus on following the prior instead of exploration. Intuitively, a higher precision means more certainty on the prior. For state estimation, the correct priors are supplied to DEM, so the precision is set to a high value of e^8 for all inputs. When performing input estimation on the first input, its precision is lowered to 1. The prior precisions on the input are converted to generalized precision matrices using the temporal correlation matrix from equation 4.9.

6.2 Smoothness Value

One requirement for DEM state and input estimation is that the noise smoothness value σ has to be known a priori and supplied to the DEM observer. Therefore, the ninth experiment was used to fit a smoothness value to the noise that is used in the calculation of the temporal correlation matrix $S(\sigma^2)$ throughout the experiments. Two different methods are used to obtain the smoothness value for the experiments. In the first method, the autocorrelation of the process noises w_1 and w_2 is calculated and compared to a range of autocorrelations for generated Gaussian convoluted noise. Based on the SSE between the two autocorrelations, the best smoothness value is selected for the noise. The second method was to run the DEM state estimation algorithm for a varying smoothness value and obtain the best smoothness value based on state estimation SSE.

Fitting by autocorrelation

First, the autocorrelation of the process noise was compared to the autocorrelations of a range of simulated Gaussian noise signals. The SSE's between the autocorrelations are reported in figure 6.2 below.

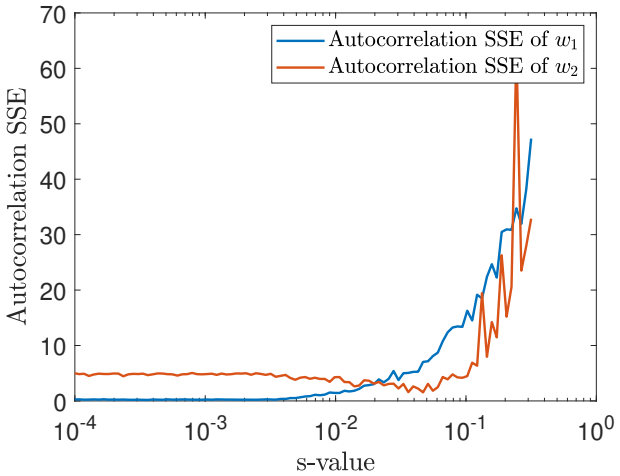

Figure 6.2: The autocorrelation SSE between the process noise and the simulated Gaussian noise with varying σ.

From the figure it can be observed that the SSE of w_1 converges to zero for lower s-values. w_2 appears to have an optimal value around 0.05. In order to provide a more decisive result a genetic algorithm was used to optimize the SSE between autocorrelations for the value σ. This resulted in a value of 0.0041 for w_1 and a value of 0.0394 for w_2. The autocorrelation functions of the noises, generated with these smoothness values, are presented in figure 6.3 alongside the autocorrelation functions of the process noise.

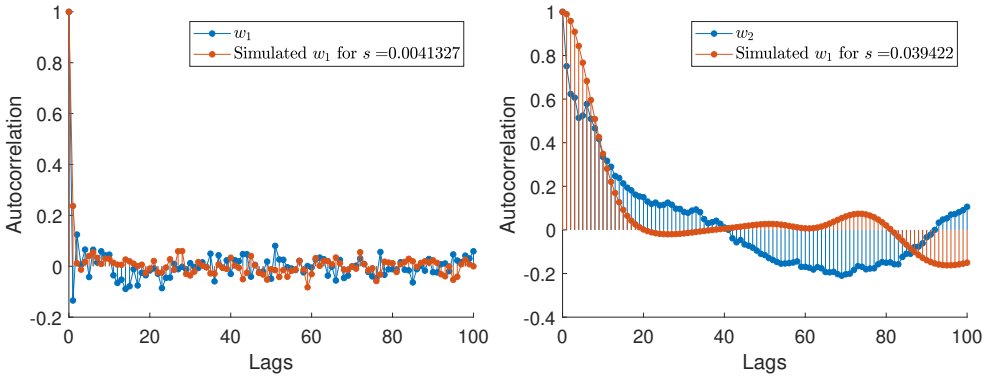

(a) Autocorrelation functions of the process noise w_1 and the best fit.

(b) Autocorrelation functions of the process noise w_2 and the best fit.

Figure 6.3: The results of the optimisation based on autocorrelation SSE fits.

Optimizing for DEM SSE

Next to the fit based on autocorrelation, the smoothness value was also tested based on the DEM state estimation performance. DEM state estimation was performed for a varying s-value and the resulting SSE's have been reported in figure 6.4. The Kalman filter SSE is added for comparison.

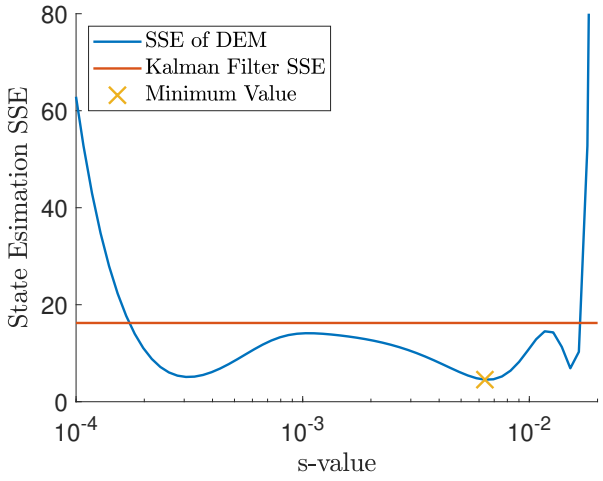

Figure 6.4: The state estimation SSE for DEM and the Kalman filter with varying s-value and a minimum at 0.006.

From the figure it can be observed that DEM outperforms the Kalman filter for a range of s-values up to 0.01. The optimal state estimation performance was recorded at 0.006, which is also close to the optimal value for w_1 from the previous section.
Combining the results of the two sections, the value of σ is set to 0.006 for the remainder of the experiments. The main reason is the fact that this value provided the best state estimation SSE. Moreover, the value can not be increased above 0.1 as figure 6.4 shows DEM becomes unstable after that value.

6.3 Noise Precisions

This section discusses how the noise precisions are determined.

Measurement Noise

In order to determine the standard deviation of the measurement noise, the OptiTrack measurement of ϕ is examined while the drone is stationary. The signal is visualized in figure 6.5. The standard deviation of the signal indicated between the dotted lines is 8.12×10^{-9}.

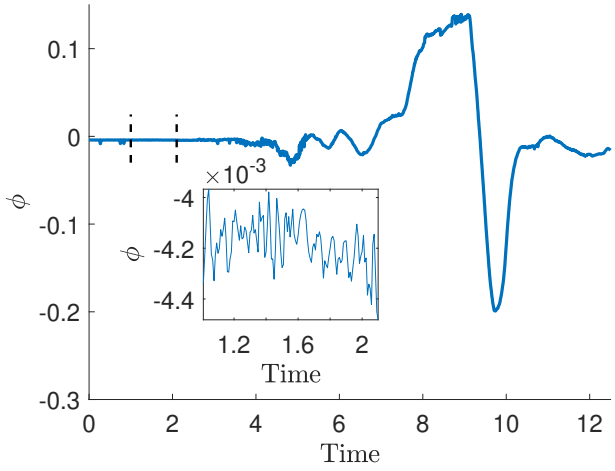

Figure 6.5: The standard deviation of the measurement noise is determined from static measurements from OptiTrack.

Process Noise

As discussed in section 5.2, equation 5.2 is used to isolate the noise from the signals. Prior to state estimation the noise w is used to calculate the noise precisions for each experiment. The calculated noise precisions $P_{w_{calc1}}$ and $P_{w_{calc2}}$ are supplied to the observers for each experiment. This allows for the observers to operate at optimal settings and allows for fair comparison between DEM and the benchmarks, since the benchmarks are supplied the fitted AR models of the noise for each experiment.

6.4 Embedding Order

The embedding orders of the states and inputs, p and d, are set to 6 and 2 respectively. The main reason for this choice is to remain consistent with the previous literature [3, 4, 5].

6.5 Input Estimation Tuning

For input estimation the observer settings are slightly different than the settings for state estimation. The two main differences are the C matrix of the LTI system and the process noise precision Π_w.
Because DEM is benchmarked against UIO, the system has to comply to the observer requirements of UIO. With the state space system as presented in equation 5.1 these observability requirements are not met [19]. The requirement that has to be met is: $rank(CE) = rank(E)$, where E is the part of B that models the unknown inputs. When estimating only the first input E becomes $[0; \frac{c_{B\phi}}{I_{xx}}]$, so:

$$rank(\begin{bmatrix}1 & 0\end{bmatrix}\begin{bmatrix}0 \\ \frac{c_{B\phi}}{I_{xx}}\end{bmatrix}) = 0 \neq rank(E)$$

Therefore, for input estimation benchmarking the C matrix is changed to the identity matrix. This makes the system fully observable and satisfies the observability requirement of UIO.
After changing the C matrix, the second alteration to the observer settings is the process noise precision matrix Π_w. In order to allow for exploration in input estimation, Π_w is changed to a fixed value of e^3.

6.6 Benchmark Tuning

For the benchmarks SA and SMIKF the AR constants of the noise are required. These are obtained using the Matlab function for AR fitting prior to state estimation. To reflect the $p = 6$ setting of DEM, SA is supplied with a sixth order AR noise model. SMIKF can only be implemented for a first order AR model.

6.7 Overview of the Observer Settings

In the previous sections it was discussed how all observer settings were selected. a more comprehensible overview of the settings is given below:

State Estimation
For state estimation the priors on the input are given as:

$$\eta^{\tilde{v}} = v, \qquad P^{\tilde{v}} = \begin{bmatrix} e^8 & 0 & 0 & 0 \\ 0 & e^8 & 0 & 0 \\ 0 & 0 & e^8 & 0 \\ 0 & 0 & 0 & e^8 \end{bmatrix}, \qquad \tilde{\Pi}_v = S(\sigma^2) \otimes P^{\tilde{v}}$$

The system matrices are:

$$\tilde{A} = I_{p+1} \otimes \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, \qquad \tilde{B} = I_{p+1} \otimes \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0.3748 & -0.3748 & -0.3748 & 0.3748 \end{bmatrix}, \qquad \tilde{C} = I_{p+1} \otimes \begin{bmatrix} 1 & 0 \end{bmatrix}$$

The smoothness value and the embedding orders are selected as given below:

$$\sigma = 0.006, \qquad d = 2, \qquad p = 6$$

Finally, the noise precisions are given as:

$$\tilde{\Pi}_w = S(\sigma^2) \otimes \begin{bmatrix} P_{w_{calc1}} & 0 \\ 0 & P_{w_{calc2}} \end{bmatrix}, \quad \tilde{\Pi}_z = S(\sigma^2) \otimes 8.12 \times 10^{-9}$$

Input Estimation
The quantities that are changes for input estimation compared to state estimation are given as follows:

$$\eta_1^{\tilde{v}} = 1, \qquad P^{\tilde{v}} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & e^8 & 0 & 0 \\ 0 & 0 & e^8 & 0 \\ 0 & 0 & 0 & e^8 \end{bmatrix}, \qquad \tilde{\Pi}_v = S(\sigma^2) \otimes P^{\tilde{v}}$$

$$\tilde{C} = I_{p+1} \otimes \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad \tilde{\Pi}_w = S(\sigma^2) \otimes \begin{bmatrix} e^3 & 0 \\ 0 & e^3 \end{bmatrix}$$

The data as collected in chapter 5 together with the settings in this chapter are used to test DEM in state and input estimation. The results are presented in the next chapter.

Chapter 7: Results

After preparing the data and tuning the observers, this chapter provides an overview of the results that were obtained from the experiments. The first eight experiments, four with wind and four without wind, were considered in the analysis. The aim is to satisfy the goals as stated in chapter 1. The results are divided into three sections, noise analysis, state estimation, and input estimation. In the noise analysis, it will be examined if the noise assumptions of the FEP and DEM are satisfied. In the state estimation section, DEM will be compared to the benchmarks SA and SMIKF and the usefulness of generalized coordinates will be highlighted. In the input estimation section, DEM will be compared to UIO based on the estimation of the first input of the system. Finally, the accuracy vs complexity trade off of DEM is presented.

7.1 Noise Analysis

For the noise analysis section, 400 samples were examined for each experiment. The noise was isolated as discussed in section 5.2. The standard deviation of the measurement noise was already determined in section 6.3. However, the standard deviations of the states and process noises are calculated for each experiment. For both the experiments with and without wind the averages are reported in table 7.1.

	ϕ	$\dot{\phi}$	w_ϕ	$w_{\dot{\phi}}$
Without wind	0.00615	0.0444	0.000427	0.0258
With wind	0.0505	0.280	0.000934	0.0612

Table 7.1: The standard deviations of the states, ϕ and $\dot{\phi}$, and the process noises, w_ϕ and $w_{\dot{\phi}}$, for the experiments with and without wind.

From the table it can be observed that introducing wind to the system does indeed cause a disturbance to the states and the noise. Moreover, since the measurement noise was determined to be 8.12×10^{-9}, the process noise is much larger in magnitude than the measurement noise. Therefore, equation 5.2 holds.

Noise Colour

One of the core features of this research is testing DEM as a state observer for coloured noise. Therefore, the autocorrelations of the noise were examined to determine whether the noise is indeed coloured. The autocorrelation functions of the noise for both the experiments with and without wind are presented in figure 7.1.

In the figures, it can be observed that none of the process noise autocorrelations immediately drops to zero, as would be expected for a white noise autocorrelation function. Moreover, the process noise on the unobserved state $\dot{\phi}$ is clearly more coloured than the process noise on ϕ. This is mainly caused by the fact that ϕ is directly observed. Finally, the noise has a more Gaussian character for the experiments in wind conditions, while the noise in the experiments for hovering without wind follows a sinusoidal pattern. This indicates that introducing wind to the system results in coloured Gaussian process noise, as assumed in the FEP.

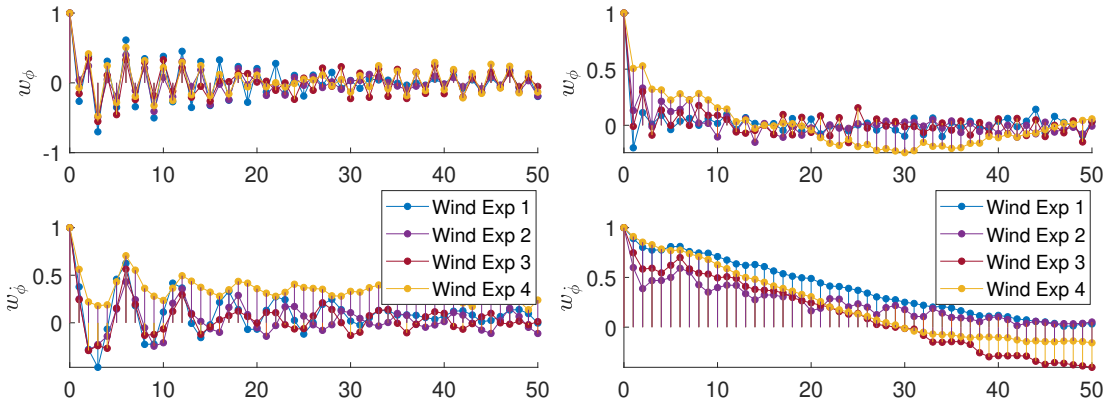

(a) The autocorrelation for all experiments without wind.

(b) The autocorrelation for all experiments with wind.

Figure 7.1: The autocorrelation functions of the process noise.

Validity of the Laplace Assumption

As stated in section 4.2, the Laplace assumption provides a vital link in the bridge from neuroscience and the FEP to DEM for LTI systems. Therefore, an important step in benchmarking DEM on experimental data is to verify that the noises follow a Gaussian distribution. In figure 7.2 the histograms of the process noises w_ϕ and $w_{\dot{\phi}}$ are presented for the experiments without and the experiments with wind. A Gaussian fit is added for reference.

As can be observed from the figures, the noise follows a Gaussian distribution for both w_ϕ and $w_{\dot{\phi}}$ for all experiments. This accommodates use of the Laplace assumption in the current experimental setup.

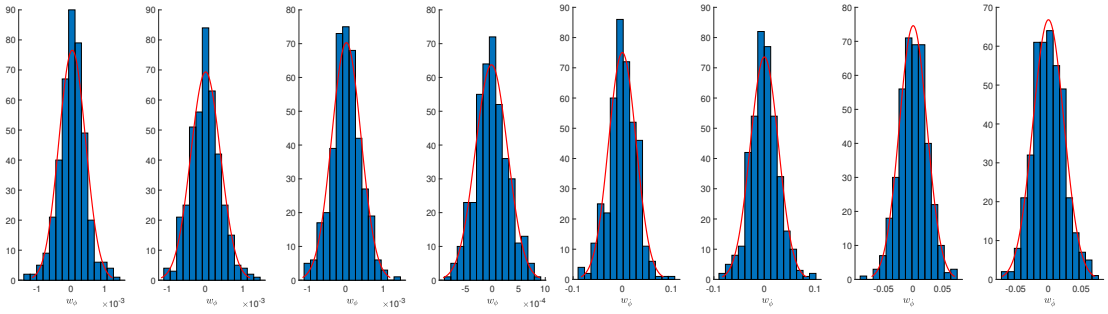

(a) Histograms of the process noise w_ϕ with a Gaussian fit for no wind conditions.

(b) Histograms of the process noise $w_{\dot{\phi}}$ with a Gaussian fit for no wind conditions.

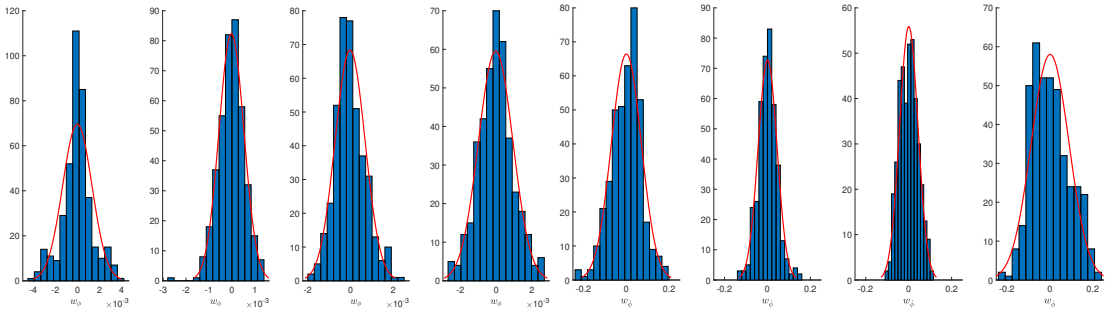

(c) Histograms of the process noise w_ϕ with a Gaussian fit for wind conditions.

(d) Histograms of the process noise $w_{\dot{\phi}}$ with a Gaussian fit for wind conditions.

Figure 7.2: The process noise from the experiments for with and without wind conditions.

7.2 State Estimation and Benchmarking

In this section, the capabilities of DEM's state estimation will be examined and compared to the benchmarks, KF, SA and SMIKF. Moreover, the impact of generalized coordinates on state estimation accuracy will be examined. The observers are tuned as described in section 6.7. The experiments are divided into 5 segments of 2 seconds to accommodate for outlier detection, which can occur in some experiments when one of the benchmarks becomes unstable, thereby distorting the results.

State Estimates

DEM is benchmarked against the Kalman filter, SA and SMIKF based of the estimation accuracy of the hidden state $\dot{\phi}$. The accuracy is measured in terms of the sum of squared errors (SSE) between the estimate of $\dot{\phi}$ and the measured state. An example of the state estimations is depicted in figure 7.3a, where the real state, the DEM estimate, and the benchmark estimates are plotted together. It can be observed that, although all state observers follow the trend, the DEM and SA estimates represent the real state the closest. The mean SSE in state estimation of all experiments and all segments, for the experiments with and without wind, is given in figure 7.3b.

From figure 7.3b it can be observed that, although SA outperforms DEM for the experiments without wind, DEM outperforms all other benchmarks for the experiments with wind.

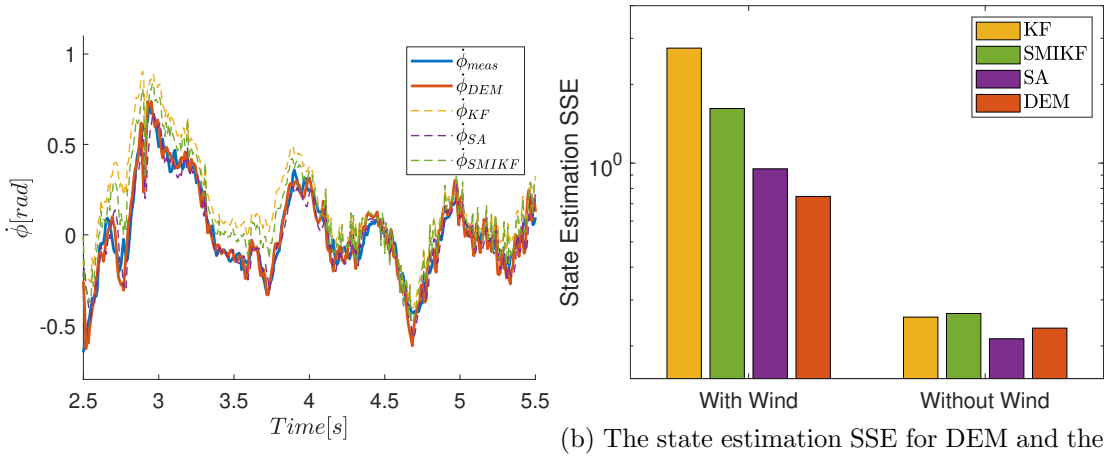

(a) DEM state estimation in wind conditions, including the benchmarks.

(b) The state estimation SSE for DEM and the benchmarks for both wind and no wind conditions.

Figure 7.3: The state estimation results of DEM and the benchmarks.

Free Energy Curve

As explained in section 4.2, DEM performs a gradient over the VFE. To demonstrate this aspect of DEM, the state $\dot{\phi}$ was arbitrarily varied for each timestep. After embedding the states, using equation 4.13, the free energy was calculated for each state, for each timestep. The result is a free energy curve that moves through time. Independent of calculating the free energy curves, DEM state estimation was performed and the free energy of the estimated states was calculated. The two results have been combined in figure 7.4.

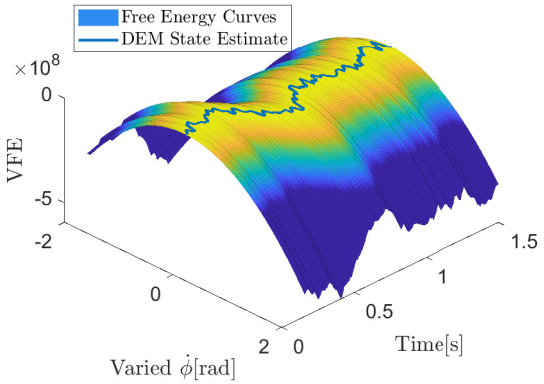

Figure 7.4: Visualisation of the VFE gradient ascent of DEM in state estimation.

The figure shows that, over time, the DEM state estimate lies on top of the free energy manifold. This visually demonstrates the VFE gradient ascent of DEM.

Role of Generalized Coordinates

One of the main strengths of DEM, its capability of dealing with coloured noise, comes from its use of generalized coordinates. To emphasize this characteristic of DEM, state estimation is performed for varying orders of generalized coordinates p. The performance is measured based on state estimation SSE. The results are given in figure 7.5. For this figure, only the experiments under wind conditions have been examined. The datapoints consist of the mean and standard deviations of the SSE for all experiments with wind.

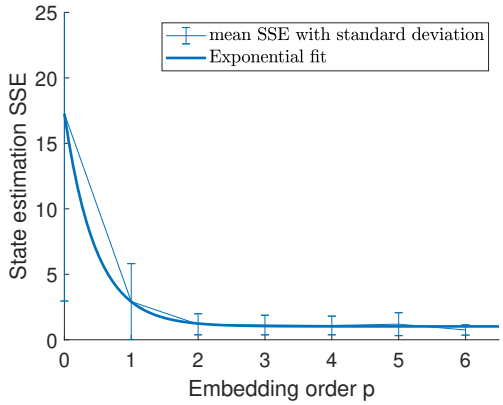

Figure 7.5: SSE drops exponentially with p.

In the figure, an exponential trend can be observed, validating the importance of generalized coordinates in coloured noise handling on experimental data.

7.3 Input Estimation

Besides state estimation, DEM is also capable of input estimation. To highlight the capabilities of DEM input estimation, DEM will be compared to an unknown input observer (UIO) [18]. For the input estimation section, the observer settings have been changed as described in section 6.7. Moreover, in order to provide more insight into the input estimation behaviour of DEM, a simulation with randomized systems has been conducted. Finally, the unique ability of DEM to balance between accuracy and complexity is highlighted.

DEM vs UIO

For input estimation, only the first of the input *pwm* values is estimated. An example of input estimation for DEM and UIO is given in figure 7.6a. It can be observed that both DEM and UIO follow the trend, but do not completely coincide with the measured input. An interesting observation is that they do coincide with each other. The input estimation SSE for all experiments for both DEM and UIO are displayed in the bar graph in figure 7.6b.

From figure 7.6 it can be observed that DEM and UIO obtained similar performance for the experiments in wind conditions, demonstrating the input estimation capabilities of DEM. However, it is also evident that UIO slightly outperforms DEM.

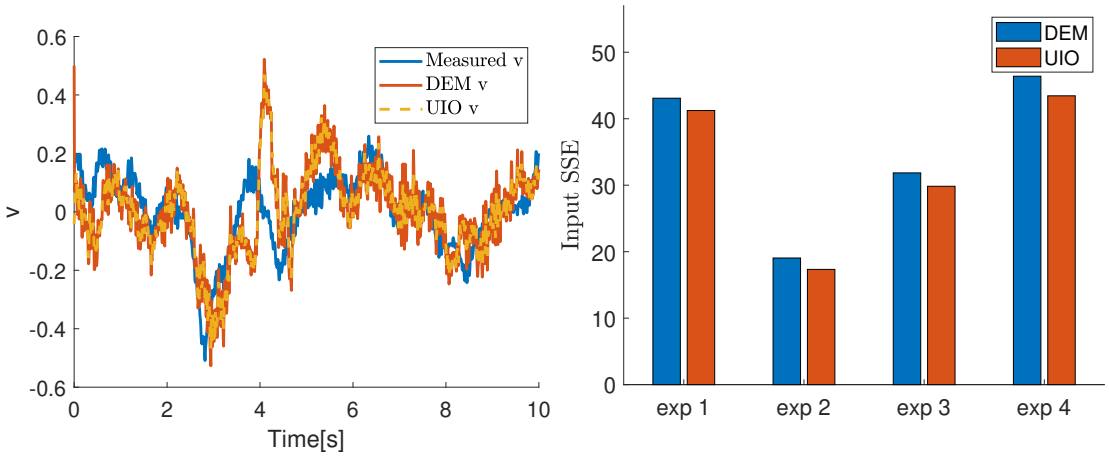

(a) A visualisation of DEM input estimation together with the UIO benchmark.

(b) The input estimation SSE for DEM and UIO for all experiments with wind conditions

Figure 7.6: The results of DEM input estimation compared to UIO.

Input Estimation Simulation

The fact that in the previous section DEM is outperformed by UIO is in contrast with the simulation results of [4]. Therefore, DEM and UIO have been compared on randomized systems with similar noise settings to the experiments with wind. In the simulation, 20 systems were randomly selected with the same dimensions as the state space system from the drone as used in this section, 2 states, 4 inputs and 2 outputs, with $C = [1\ 0; 0\ 1]$. The noise was generated by convoluting a white noise signal with a Gaussian kernel as described in section 4.2.3. The process noise w was simulated with a smoothness value of $\sigma = 0.1$ and a standard deviation of $\sigma_w = exp(-4)$. The measurement noise z was simulated with a smoothness value of $\sigma = 0.01$ and a standard deviation of $\sigma_z = exp(-16)$. For these settings the first input was estimated for 20 systems and the SSE was calculated for both DEM and UIO. The results are shown in figure 7.7a.

From figure 7.7a it can be observed that DEM and UIO perform very similar, with UIO outperforming DEM in some cases, as encountered in the previous section. However, the same simulation was performed while changing the smoothness value of the measurement noise to $\sigma = 0.1$. This provided the system with a more coloured measurement noise and allowed for DEM to stably operate with a higher smoothness value, which is now also set to 0.1. The resulting SSEs are presented in figure 7.7b. From this figure it can be observed that DEM now outperforms UIO in 85% of the simulated systems.

For further investigation the estimated inputs from the simulation were also plotted. Each figure contains the input estimation of eight simulated systems, which are shown in 7.8 and 7.9 below.

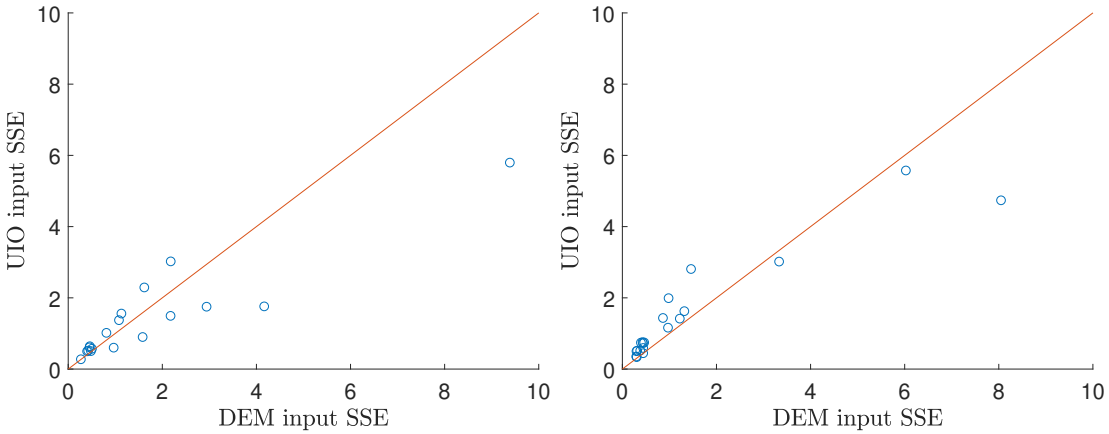

(a) DEM and UIO performance on simulated systems with a measurement noise smoothness value of 0.01.

(b) DEM and UIO performance on simulated systems with a measurement noise smoothness value of 0.1.

Figure 7.7: The results of DEM and UIO input estimation on simulated systems with different measurement noise smoothness values.

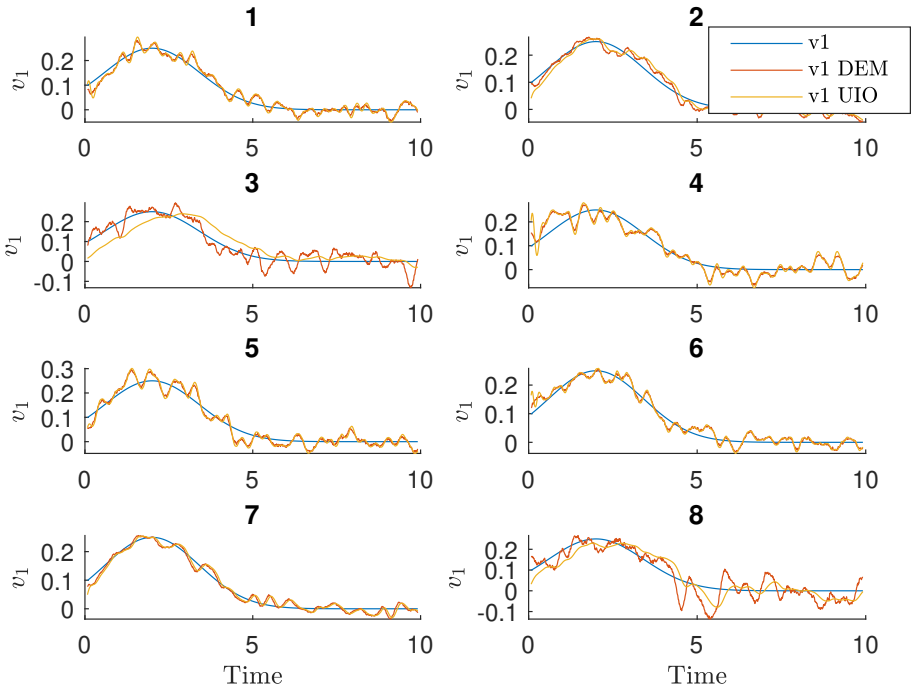

Figure 7.8: The input estimation of DEM and UIO for eight random systems with generated measurement noise with $\sigma = 0.01$.

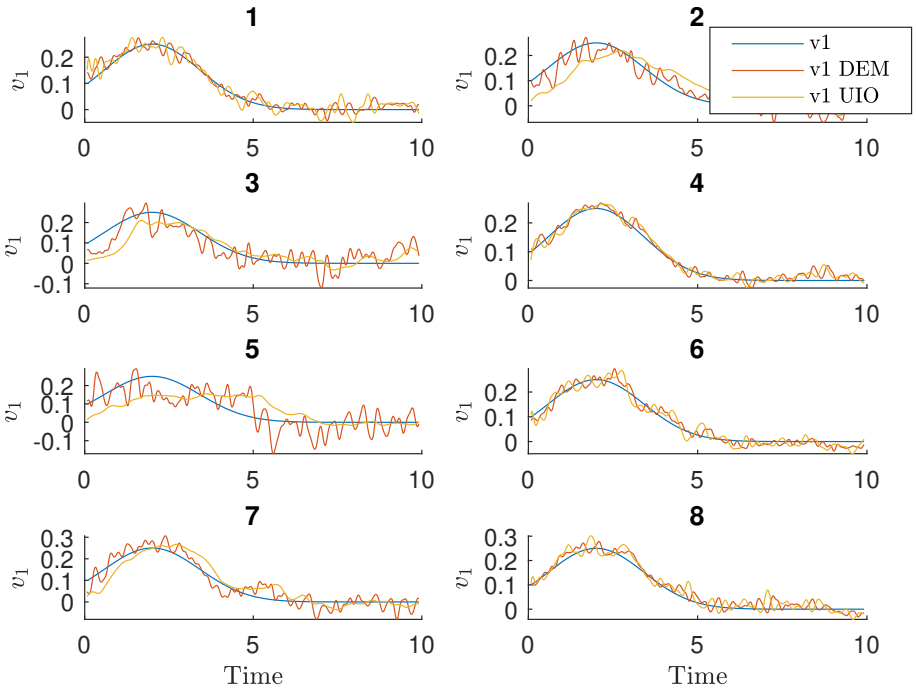

Figure 7.9: The input estimation of DEM and UIO for eight random systems with generated measurement noise with $\sigma = 0.1$.

From these figures it can be observed that for a smoothness value of 0.01, six out of eight of the input estimates of DEM and UIO converge to the exact same estimated input. For a smoothness value of 0.1, this is only the case for one of the systems. This indicates that when DEM assumes noise that is less coloured, it behaves the same as UIO. When DEM considers more coloured noise, it can leverage the noise colour to increase its input estimation performance. In the current experimental setup, DEM was slightly outperformed by UIO because the smoothness value is chosen as a small value of 0.006 and can not be increased too much as indicated in section 6.2.

Accuracy v/s complexity

One of DEM's abilities is to moderate between accuracy and complexity in estimation. In input estimation the accuracy vs complexity trade off is dictated by the input priors $\eta^{\tilde{v}}$ and $P^{\tilde{v}}$. For a high complexity DEM focuses on the priors and tries to follow them as closely as possible. This is accomplished with a high prior precision. A high accuracy indicates that the DEM estimate is very close to the real input. This is accomplished by either supplying DEM with the correct prior or lowering the precision to allow DEM to explore. The trade off is visualized in figure 7.10a, where an incorrect prior is supplied to DEM and different input estimation results are obtained for different values of P_v.

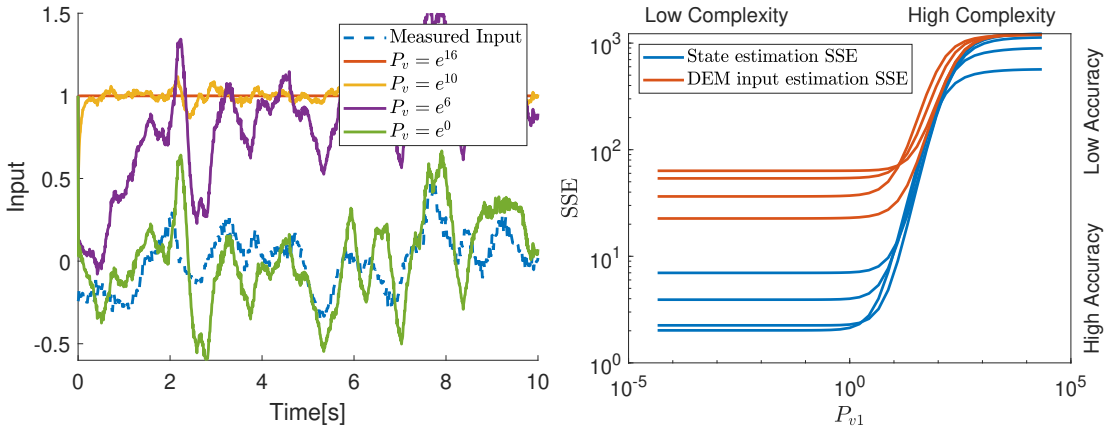

(a) DEM and UIO performance on simulated systems with a measurement noise smoothness value of 0.01.

(b) The SSE of state and input estimation for a varied precision, P_v, on one input with a wrong prior.

Figure 7.10: The results of DEM and UIO input estimation on simulated systems with different measurement noise smoothness values.

It can be observed from figure 7.10a that for a high precision DEM has a high complexity and focuses on the prior. For a lower precision, DEM starts to move towards the measured input, thereby increasing its accuracy. The accuracy vs complexity trade off is also demonstrated in figure 7.10b. In this figure, the states and inputs are estimated for a varying precision P_v for all experiments in wind conditions. It can be observed that for a high precision DEM has a high complexity and a low accuracy. When the precision on the priors in lowered, DEM moves to a low complexity and a high accuracy, indicated by the state and input estimation SSE.

Chapter 8: Conclusion and Recommendations

This section will restate the main goals of the thesis, provide a quick overview of the method, and briefly present the results. Finally, two suggestions for future work will be proposed.

8.1 Conclusion

Originating from neuroscience, the free energy principle proposes a theory that explains action, perception, and learning in the human brain. The FEP states that the brain is constantly minimizing its entropy, or "surprise". The surprise is caused by a combination of the generative model in the brain and the sensory input from the environment. The free energy quantity functions as an upper bound on this surprise. Therefore, in the FEP framework, the brain is constantly minimizing the free energy. This provides a theory that explains the efficiency of the human brain, while providing a mathematical description of the process. In recent studies, this math has been converted to control theory in robotics, with the aim of developing more efficient robots.
The perception aspect of the FEP is described in the dynamic expectation maximisation algorithm. Operating within the FEP framework, DEM is capable of parameter, hyperparameter, state and input estimation by performing a gradient ascent over the free energy. This thesis focuses on the state and input estimation capabilities of DEM. Because of DEM's use of generalized coordinates, DEM is able to leverage the coloured nature of noise for more accurate state and input estimation. Here generalized coordinates signify that the higher order derivatives of the system's states are also considered in estimation. Moreover, by making use of the Laplace assumption and the Gaussian nature of the noise, DEM is converted to a state and input observer for LTI systems. The result is a novel state and input observer for LTI system under the influence of coloured noise.
In recent studies, DEM has shown to outperform the Kalman filter in simulation and on experimental data. This research aimed to extend on these studies and use the experimental data to validate the performance of DEM in robotics. Furthermore, this thesis aimed to expand the list of state estimation benchmarks to state observers for coloured noise, to accommodate for fair comparison. Next to state observation, DEM input estimation was tested on experimental data. On top of the thesis, the goal was to combine the research of the previous studies and the thesis into a publishable paper. Therefore the main goal of the thesis was summarized was:

Use experimental data to validate DEM state and input estimation performance for a system under the influence of coloured noise and present the results in a publishable paper.

The main goal was supported by the following sub-goals:

Validate that the experimental data agrees with the noise and model assumptions in the DEM framework, such as the Laplace assumption and the coloured noise model.

Provide experimental validation of the usefulness of generalized coordinates in state estimation for coloured noise.

Test DEM state and input estimation on experimental data and compare its estimation accuracy

against conventional observers for coloured noise such as state augmentation, SMIKF and an unknown input observer.

Demonstrate the influence of priors in the accuracy vs complexity trade off of DEM.

The experiment consists of a quadcopter drone flying in wind conditions. Using a blower, the wind conditions can be regulated. The drones movements were monitored using a 3D motion capture system. A total of nine experiments were carried out, four without wind and five with wind. The data from the experiments was prepared in Matlab, resulting in 9 experiments of 10 seconds. The ninth experiment was used to tune the DEM observer a well as the other benchmarks. The other eight experiments were used to obtain the results.

The results show that the noise is indeed coloured and follows a Gaussian distribution. This means the use of the Laplace assumption in the derivation of the DEM observer for LTI systems is valid. Furhtermore, this indicates that the experiment complies with the noise assumptions of DEM. The state estimation results demonstrate that, in the wind conditions from the experiments, DEM outperforms the other state observers for coloured noise. Moreover, it is observed that the use of generalized coordinates increases the accuracy of DEM state estimation. This indicates that DEM uses generalized coordinates to leverage the coloured nature of the noise for better state estimation and proves the usefulness of generalized coordinates in robotics. In input estimation, DEM was able to obtain similar performance compared to an unknown input observer. Finally, the DEM accuracy vs complexity trade off was demonstrated. With these results the main- and sub-goals of the thesis have been satisfied. The result is an experimental setup that validates the performance of DEM, as a state observer for LTI systems influenced by coloured noise, and the use generalized coordinates in the field of robotics.

8.2 Recommendations

One of the main problems that arose when testing DEM on experimental data, was the fact that, unlike with simulated data, the smoothness value of the noise is unknown and needs to be calculated and supplied to the observer prior to estimation. Moreover, the smoothness value for all three portions of the noise, w_ϕ, $w_{\dot{\phi}}$ and z, are different. In its current form, DEM uses one value for all types of noise. Because this value is not optimal for all noises, this deteriorates DEM's performance. For practical applications, it would be beneficial to develop a method that enables different smoothness values for different noise levels.

The second consideration is the time delay that is introduced into the state observer by the reverse Taylor expansion. Because the outputs have to be embedded using both past and future measurements a delay is introduced into the algorithm. This is a disadvantage when implementing DEM into an online control scheme for robots. Therefore, a different embedding approach is needed to avoid a delay in online state estimation and control.

Bibliography

[1] K. Friston, "The free-energy principle: a unified brain theory?," Nature reviews neuroscience, vol. 11, no. 2, pp. 127–138, 2010.

[2] K. Friston, J. Kilner, and L. Harrison, "A free energy principle for the brain," Journal of Physiology-Paris, vol. 100, no. 1-3, pp. 70–87, 2006.

[3] K. J. Friston, N. Trujillo-Barreto, and J. Daunizeau, "Dem: a variational treatment of dynamic systems," Neuroimage, vol. 41, no. 3, pp. 849–885, 2008.

[4] A. A. Meera and M. Wisse, "Free energy principle based state and input observer design for linear systems with colored noise," in 2020 American Control Conference (ACC), pp. 5052–5058, IEEE, 2020.

[5] D. Benders, "Ar.drone 2.0 state estimation using dynamic expectation maximization," Master's thesis, 2020. Available at `http://resolver.tudelft.nl/uuid:156157c6-d7f0-4dc1-a55a-b2e4ed66f1c2`.

[6] F. Bos, A. A. Meera, D. Benders, and M. Wisse, "Free energy principle for state and input estimation of a quadcopter flying in wind," 2021.

[7] R. ASHBY, "Principles of the self-organizing dynamic system," Journal of General Psychology, 1947.

[8] A. Anil Meera and M. Wisse, "Dynamic expectation maximization algorithm for estimation of linear systems with colored noise," Entropy, vol. 23, no. 10, p. 1306, 2021.

[9] A. C. Jalfon and Y. Halevi, "On partially augmented observers for systems with coloured noises," International Journal of Control, vol. 57, no. 2, pp. 335–349, 1993.

[10] B. P. Gibbs, Advanced Kalman filtering, least-squares and modeling: a practical handbook. John Wiley & Sons, 2011.

[11] F. Bos, "State estimation for a drone under theinfluence of coloured noise," TU Delft literature review, 2021.

[12] A. Bryson and D. Johansen, "Linear filtering for time-varying systems using measurements containing colored noise," IEEE Transactions on Automatic Control, vol. 10, no. 1, pp. 4–10, 1965.

[13] A. Bryson Jr and L. Henrikson, "Estimation using sampled data containing sequentially correlated noise.," Journal of Spacecraft and Rockets, vol. 5, no. 6, pp. 662–665, 1968.

[14] Z. Zhou, J. Wu, Y. Li, C. Fu, and H. Fourati, "Critical issues on kalman filter with colored and correlated system noises," Asian Journal of Control, vol. 19, no. 6, pp. 1905–1919, 2017.

[15] R. Moghe, R. Zanetti, and M. R. Akella, "Adaptive kalman filter for detectable linear time-invariant systems," Journal of Guidance, Control, and Dynamics, vol. 42, no. 10, pp. 2197–2205, 2019.

[16] D. Simon, Optimal state estimation: Kalman, H infinity, and nonlinear approaches. John Wiley & Sons, 2006.

[17] R. E. Kalman, "A new approach to linear filtering and prediction problems," Transactions of the ASME–Journal of Basic Engineering, vol. 82, no. Series D, pp. 35–45, 1960.

[18] B. A. Charandabi and H. J. Marquez, "Observer design for discrete-time linear systems with unknown disturbances," in 2012 IEEE 51st IEEE Conference on Decision and Control (CDC), pp. 2563–2568, IEEE, 2012.

[19] J. Chen, R. J. Patton, and H.-Y. Zhang, "Design of unknown input observers and robust fault detection filters," International Journal of control, vol. 63, no. 1, pp. 85–105, 1996.