Laboratory for Reactor Physics
and Systems Behaviour

The Unscented Kalman Filter: Introduction and Application to Point Reactor Kinetics

14/06/2019

Author:
Felix Grimberg, BSc

Supervisors:
Daniel Siefman, MSc
Vincent Lamirand, PhD

Abstract

This report serves two purposes. The first is to introduce the linear Kalman filter and the unscented Kalman filter to readers who are unfamiliar with either or both. To this end, it contains an extensive explanation of these methods, including all equations and interpretations which proved helpful to the author in realizing the second objective. Secondly, this report describes an original application of the unscented Kalman filter to the point reactor kinetics equations. The unscented Kalman filter's capacity to estimate the time-dependent behavior of a thermonuclear reactor is demonstrated, as well as its ability to refine prior estimates of the reactor-specific constants. Limits to the latter ability are expressed.

Contents

List of Tables

List of Figures

1 Introduction

The Kalman filter was first proposed in 1960 and many variants have since been developed [16]. It is a data fusion algorithm with which outputs can be estimated, and their future value predicted, based on noisy inputs and a linear prediction model. It was originally designed as a recursive implementation of the least-squares estimator. One range of applications for the Kalman filter consists of estimating the state[1] of a physical system through a combination of model[2] predictions and noisy observations. Its recursive nature makes it particularly well-suited for real-time applications: Indeed, after obtaining new observations, a new estimate can be produced using only the last calculated estimate, rather than the entire set of earlier observations from which the latter was obtained. While the capacity to filter[3] observations in real time "as they appear" is a principal advantage of the Kalman filter, it may equally be used to smooth[4] a number of observations in one batch [5]. The linear Kalman filter is designed to yield the best linear estimator[5] of the sought-for inputs, provided that the system under consideration is correctly described by a linear model [9]. If the errors follow zero-mean Gaussian distributions with known variances, this estimator is optimal even among nonlinear estimators [4]. More elaborate filters have been developed to suit a wider range of models [7, 15, 3, 1]. One of these newer variants is the unscented Kalman filter, which is designed to estimate the mean and the covariance of any nonlinear model's output [15].

The point reactor kinetics (PRK) equations are a system of coupled ordinary differential equations describing the time-dependent behavior of thermonuclear reactors. It is a crucial part of nuclear safety analysis to correctly assess this behavior, which is determined by time- and energy-dependent properties of materials — most notably, properties of the fuel and the moderator. In the PRK equations, these properties are averaged over all independent variables[6] except time, resulting in a comparatively simple model. However, the resulting averages are reactor-specific because the distribution of the independent variables differs depending on the reactor design. Examples of such reactor-specific average quantities include the delayed neutron fraction and the mean neutron generation time. Their values can be obtained through calculations involving a three-dimensional model of the reactor with measured material and fuel properties [10]. Another complementary approach is to infer them from empirical data.

[1] The term *state* here designates a collection of time-dependent variables related to the physical system of interest. In particular, the value taken by a state variable at any given time must be fully determined by the analysis of the system at said time.

[2] A *model* is a set of equations which describe, often in a simplified way, how a system evolves from one state to another.

[3] In signal processing, a *filter* is a device (physical or numeric) which suppresses unwanted components of a signal. For example, low-pass filters suppress high-frequency oscillations. The Kalman filter reduces the uncertainty by *filtering out* the noise from its input. It does so recursively, i.e. it reuses the result of prior calculations when assimilating new inputs.

[4] A *smoother* estimates each state based on all available observations, whereas a filter estimates the state at time t based only on the observations taken before t.

[5] Here, *best* is understood in terms of minimizing the mean squared error. Note that in many cases, the best linear estimator is outperformed by some nonlinear estimators.

[6] In the context of a mathematical function $y = f(x, t)$, the inputs x and t are called *independent variables*, whereas the output y is a *dependent variable* because its value depends on that of the inputs, x and t.

As shown in the present work, the unscented Kalman filter can be used to estimate (and thus predict) the neutron population in an experimental test reactor more accurately based on experimental data. It can further be used to refine existing estimates of reactor-specific average properties.

The linear Kalman filter is explained in section 2 using a fictitious application case. Section 3 provides a theoretical discussion of the employed models: A simple introduction to the point reactor kinetics equations is followed by a description of the unscented Kalman filter. The design by which the latter is applied to the former is relayed in section 4. Finally, the presented design is tested and results are shown in section 5.

2 The Linear Kalman Filter

The Kalman filter is a data fusion algorithm commonly used to make predictions from noisy data. It can be efficiently implemented as a recursive filter[3] which combines several estimates of a variable, as well as prior knowledge about the uncertainty associated with each of them, to produce a more accurate estimate [16]. The linear Kalman filter is equivalent to the optimal linear least squares estimator. Its mathematical formulation is explained in this section using a fictitious problem. The problem is further used to showcase the performance of the filter.

2.1 Introductory Problem

A simple problem is devised, to which the linear Kalman filter will be applied. The problem consists of a car which moves on a linear trajectory in two dimensions. The relevant properties of the car, i.e. its position and velocity, are parameterized in its *state vector* $\mathbf{x}(t) = [x_1, x_2, \dot{x}_1, \dot{x}_2]^T(t)$. It can be subjected to an acceleration/deceleration profile $\mathbf{a}(t) = [\ddot{x}_1, \ddot{x}_2]^T(t) \parallel \mathbf{d}$ along a constant direction $\mathbf{d}$, and it is assumed to be at rest at the origin at time $t^{(0)}$: $\mathbf{x}(t^{(0)}) = \mathbf{0}$. Furthermore, its position can be observed, as well as a linear combination of its velocity components $\dot{x}_1$ and $\dot{x}_2$. Based on these observations and a postulated model, we are interested in providing estimates of the car's state with the highest possible accuracy at a series of n distinct times $\{t^{(i)}\}_{i=0,\dots,n}$, where $t^{(i)} = i \cdot \Delta t$.

For a sufficiently small time step Δt, the behaviour of the car is assumed to correspond to a linear model:

$$\begin{aligned}
x_1(t^{(i+1)}) &= x_1(t^{(i)}) + \Delta t \cdot \dot{x}_1(t^{(i)}) + \frac{\Delta t^2}{2} \cdot \ddot{x}_1(t^{(i)}) + \textit{random noise} \\
x_2(t^{(i+1)}) &= x_2(t^{(i)}) + \Delta t \cdot \dot{x}_2(t^{(i)}) + \frac{\Delta t^2}{2} \cdot \ddot{x}_2(t^{(i)}) + \textit{random noise} \\
\dot{x}_1(t^{(i+1)}) &= \dot{x}_1(t^{(i)}) + \Delta t \cdot \ddot{x}_1(t^{(i)}) + \textit{random noise} \\
\dot{x}_2(t^{(i)}) &= \dot{x}_2(t^{(i)}) + \Delta t \cdot \ddot{x}_2(t^{(i)}) + \textit{random noise}
\end{aligned} \tag{1}$$

Note the presence of *random noise* terms in eqs. (1): These terms, generally referred to as *process noise*, cause the state vector to be a random variable — hence the need to estimate it. There are two possible sources of process noise: noisy control inputs and inaccuracies in the model. In the problem described above, the *control input* corresponds to the acceleration/deceleration profile $\mathbf{a}(t) = [\ddot{x}_1, \ddot{x}_2]^T(t)$ to which the car is subjected. Assuming that the control input is noisy is then equivalent with acknowledging that the real acceleration/deceleration profile of the car differs slightly from the targeted profile, e.g., due to gusts of wind or a varying slope of the road.

The equations in (1) can be written equivalently in tensor notation, by introducing two matrices: the *state transition matrix*, $\mathbf{F}$, which defines how the state evolves in the absence of a control input, and the *control input matrix*, $\mathbf{B}$, which defines the effect of the control input. The process noise is denoted as a random vector $\mathbf{w}^{(i)}$ whose size matches that of the state vector:

$$\begin{gathered}\mathbf{x}(t^{(i+1)}) = \mathbf{F} \cdot \mathbf{x}(t^{(i)}) + \mathbf{B} \cdot \mathbf{a}(t^{(i)}) + \mathbf{w}^{(i)} \\ \mathbf{F} = \begin{bmatrix} 1 & 0 & \Delta t & 0 \\ 0 & 1 & 0 & \Delta t \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \qquad \mathbf{B} = \begin{bmatrix} \frac{\Delta t^2}{2} & 0 \\ 0 & \frac{\Delta t^2}{2} \\ \Delta t & 0 \\ 0 & \Delta t \end{bmatrix}\end{gathered} \tag{2}$$

Following [4], we will assume for this introductory problem that the process noise $\mathbf{w}^{(i)}$ follows a multivariate zero-mean Gaussian distribution with the *covariance matrix* $\mathbf{\Sigma}_{process}$ defined in eq. (9):

$$\mathbf{w}^{(i)} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_{process}) \tag{3}$$

Similarly, a model can be postulated to describe how observations are generated. As stated in subsection 2, the Kalman filter considered here is designed for linear models. In particular, it requires that each observation be a linear transformation of the state vector. The *observations* are denoted by z_j because they need not correspond exactly to any of the state variables. Indeed, as long as each z_j is fully determined by the state vector $\mathbf{x}$, it contains information that can be used to refine our estimate of the latter[7].

For the hypothetical problem of the moving car, we choose the following model:

$$\begin{aligned} z_1^{(i)} &= x_1(t^{(i)}) + random\ noise \\ z_2^{(i)} &= x_2(t^{(i)}) + random\ noise \\ z_3^{(i)} &= 0.5 \cdot (\dot{x}_1(t^{(i)}) + \dot{x}_2(t^{(i)})) + random\ noise \end{aligned} \tag{4}$$

This model, too, contains *random noise*: In this context, we call it the *observation noise*. It is caused by the limited precision of any used measuring devices: For example, a car's position can be measured by GPS signal, but only to a precision of a few meters. Again, we write

[7] Going one step further, state variables may even be estimated without being observed at all: If they impact other, observed state variables during the state transformation (eq. 2), then the resulting covariance terms (cf. eq. 14) eventually lead to an update of the unobserved variables through eqs. (15) and (16).

eqs. (4) in tensor notation by viewing the observation noise as a random vector $\mathbf{v}^{(i)}$, and by introducing the *observation matrix* $\mathbf{H}$:

$$\mathbf{z}^{(i)} = \left[z_1^{(i)}, z_2^{(i)}, z_3^{(i)}\right]^T = \mathbf{H} \cdot \mathbf{x}(t^{(i)}) + \mathbf{v}^{(i-1)}, \qquad \mathbf{H} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0.5 & 0.5 \end{bmatrix} \tag{5}$$

We further assume that the observation noise $\mathbf{v}^{(i-1)}$ follows a multivariate zero-mean Gaussian distribution with the *covariance matrix* $\mathbf{\Sigma}_{obs}$ defined in eq. (7):

$$\mathbf{v}^{(i-1)} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_{obs}) \tag{6}$$

Two covariance matrices have been introduced in the equations above: $\mathbf{\Sigma}_{process}$ for the process noise, and $\mathbf{\Sigma}_{obs}$ for the observation noise. They must be specified to complete the models (2) and (5). Two scalars, σ_{acc} and σ_{obs}, are introduced to scale the covariance matrices up or down individually.

It is assumed that the coordinates of $\mathbf{z}^{(i)}$ are obtained through separate measurements. Thus, the coordinates of $\mathbf{v}^{(i)}$ are independent from each other, yielding a diagonal covariance matrix:

$$\mathbf{\Sigma}_{obs} = \sigma_{obs}^2 \cdot \mathbf{I}_{3\times 3} \tag{7}$$

The process noise, on the other hand, is assumed to result mainly from a noisy control input[8], i.e. the addition of a Gaussian noise term with variance σ_{acc}^2 to the acceleration inputs $\mathbf{a}(t^{(i)})$. By considering the effect of this hypothetical noise term on the uncertainty of the car's position and velocity after one single time step, we obtain the covariance matrix $\mathbf{\Sigma}_{ideal}$:

$$\mathbf{\Sigma}_{ideal} = \sigma_{acc}^2 \cdot \mathbf{q} \cdot \mathbf{q}^T \qquad \mathbf{q} = \begin{bmatrix} 0.5 \cdot \Delta t^2 \cdot \mathbf{d} \\ \Delta t \cdot \mathbf{d} \end{bmatrix} \tag{8}$$

The orthogonal covariance matrix $\mathbf{\Sigma}_{ideal}$ is not quite satisfactory for the physical example of the moving car, because it would lead to a perfect correlation between any two coordinates of $\mathbf{w}^{(i)}$. An arbitrary departure $\mathbf{\Sigma}_\Delta$ from the ideal case is therefore introduced. Note that the diagonal terms of $\mathbf{\Sigma}_\Delta$ must be larger than its extra-diagonal terms to mitigate the correlation between individual components of $\mathbf{w}^{(i)}$. We then compute $\mathbf{\Sigma}_{process}$ in a multiplicative way, in order to preserve the scaling effect of σ_{acc} (N.B.: the symbol $\circ$ is used here to denote the Hadamard or *element-wise* product):

$$\mathbf{\Sigma}_{process} = \mathbf{\Sigma}_\Delta \circ \mathbf{\Sigma}_{ideal} = \sigma_{acc}^2 \cdot \mathbf{\Sigma}_\Delta \circ (\mathbf{q} \cdot \mathbf{q}^T) \qquad \mathbf{\Sigma}_\Delta = \begin{bmatrix} 1.15 & 0.7 & 0.75 & 0.73 \\ 0.7 & 1.1 & 0.73 & 0.76 \\ 0.75 & 0.73 & 1 & 0.95 \\ 0.73 & 0.76 & 0.95 & 1.015 \end{bmatrix} \tag{9}$$

[8]The acceleration of a car cannot be dictated to absolute precision for many reasons, including unpredictable variations of the friction forces which counteract its movement.

2.2 Application of the Kalman Filter

The Kalman filter described in [4] is applied to the previously described problem. The estimates which it provides are plotted against, among others, the true value and the observations, in figure 1. To this end, the problem and the filter are implemented in Python 3.7.3, using the open-source libraries pykalman, NumPy and Matplotlib.

2.2.1 Mathematical Formulation

The matrices $\mathbf{F}$, $\mathbf{B}$, $\mathbf{H}$, $\mathbf{\Sigma}_{process}$, and $\mathbf{\Sigma}_{obs}$, are stored after setting key parameters: the time step Δt, the number of time steps n to be considered, and the scaling factors σ_{acc} and σ_{obs}. Then, a discrete acceleration profile $\{\mathbf{a}^{(i)}\}_{i=0,\dots,n-1}$ is defined. The products $\{\mathbf{B}\cdot\mathbf{a}^{(i)}\}_{i=0,\dots,n-1}$ are precomputed and saved as *transition offsets* $\{\mathbf{b}^{(i)}\}_{i=0,\dots,n-1}$:

$$\mathbf{b}^{(i)} = \mathbf{B}\cdot\mathbf{a}^{(i)}, \qquad i = 0,\dots,n-1 \tag{10}$$

At each time step, a number of variables are computed:

- $\mathbf{x}(t^{(i)})$[9]: The *true state* of the car at time $t^{(i)}$ is randomly generated according to eq. (2).
- $\mathbf{z}^{(i)}$[9]: A vector of observations is randomly generated using the previously calculated true state $\mathbf{x}(t^{(i)})$, according to eq. (5). Note that the uncertainty of the observations does not vary between time steps.
- $\hat{\mathbf{x}}^{(i)}$: A *purely model-based estimate* is computed, which does not take into account any of the observations. It is computed using the deterministic part of eq. (2):

$$\hat{\mathbf{x}}^{(i)} = \mathbf{F}\cdot\hat{\mathbf{x}}^{(i-1)} + \mathbf{B}\cdot\mathbf{a}^{(i-1)} \tag{11}$$

- $\mathbf{\Sigma}^{(i)}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}$: The *covariance matrix* $\mathbf{\Sigma}^{(i)}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}$ quantifies the uncertainty associated with $\hat{\mathbf{x}}^{(i)}$ as an estimate for $\mathbf{x}(t^{(i)})$. By observing that $\hat{\mathbf{x}}^{(i)}$ is entirely deterministic, we find that the covariance $\mathbf{\Sigma}^{(i)}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}$ of the error $|\hat{\mathbf{x}}^{(i)} - \mathbf{x}(t^{(i)})|$ is (the same as) the covariance of the true value $\mathbf{x}(t^{(i)})$, if the latter is viewed as a random variable at the time $t^{(0)}$. Using the Gaussianity of the noise terms $\{\mathbf{w}^{(i)}\}_{i=0,\dots,n-1}$, we can therefore compute $\mathbf{\Sigma}^{(i)}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}$ iteratively as:

$$\mathbf{\Sigma}^{(i)}_{\hat{\mathbf{x}}\hat{\mathbf{x}}} = \mathbf{F}\cdot\mathbf{\Sigma}^{(i-1)}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}\cdot\mathbf{F}^T + \mathbf{\Sigma}_{process} \tag{12}$$

- $\hat{\mathbf{x}}^{(i|i-1)}$: This is called the *prediction step* of the Kalman filter. At this stage, an estimate of $\mathbf{x}(t^{(i)})$ is computed using the deterministic part of eq. (2) and the best available estimate

[9] The true states $\mathbf{x}(t^{(i)})$, observations $\mathbf{z}^{(i)}$, and the best possible estimates $\hat{\mathbf{x}}^{(i|i)}$ along with their error covariances $\mathbf{\Sigma}^{(i|i)}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}$, are computed at once by using the pykalman module.

of the state $\mathbf{x}(t^{(i-1)})$ at the **previous** time step:

$$\hat{\mathbf{x}}^{(i|i-1)} = \mathbf{F} \cdot \hat{\mathbf{x}}^{(i-1|i-1)} + \mathbf{B} \cdot \mathbf{a}^{(i-1)} \tag{13}$$

- $\mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(i|i-1)}$: The uncertainty associated with $\hat{\mathbf{x}}^{(i|i-1)}$ as an estimate for $\mathbf{x}(t^{(i)})$ is quantified by the *covariance matrix* $\mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(i|i-1)}$ of the error $|\hat{\mathbf{x}}^{(i|i-1)} - \mathbf{x}(t^{(i)})|$, which is calculated in the same way as $\mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(i)}$:
$$\mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(i|i-1)} = \mathbf{F} \cdot \mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(i-1|i-1)} \cdot \mathbf{F}^T + \mathbf{\Sigma}_{process} \tag{14}$$
- $\hat{\mathbf{x}}^{(i|i)}$[9]: During the *update step* of the Kalman filter, $\hat{\mathbf{x}}^{(i|i-1)}$ and $\mathbf{z}^{(i)}$ are combined to provide the best possible estimate of the state $\mathbf{x}(t^{(i)})$ at the **current** time step. $\hat{\mathbf{x}}^{(i|i)}$ is given by eq. (16).
- $\mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(i|i)}$[9]: Analogously, the uncertainty associated with $\hat{\mathbf{x}}^{(i|i)}$ as an estimate for $\mathbf{x}(t^{(i)})$ is quantified by the covariance matrix $\mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(i|i)}$ of the error $|\hat{\mathbf{x}}^{(i|i)} - \mathbf{x}(t^{(i)})|$, which is given by eq. (17).

The state of the car at time $t^{(0)}$ is assumed to be perfectly known, hence $\hat{\mathbf{x}}^{(0)} = \hat{\mathbf{x}}^{(0|0)} = \mathbf{x}(t^{(0)}) = \mathbf{0}$ and $\mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(0)} = \mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(0|0)} = \mathbf{0}_{4\times4}$. The equations for $\hat{\mathbf{x}}^{(i|i)}$ and $\mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(i|i)}$ are derived in [4]:

$$\mathbf{K}^{(i)} = \mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(i|i-1)} \cdot \mathbf{H}^T \cdot \left(\mathbf{H} \cdot \mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(i|i-1)} \cdot \mathbf{H}^T + \mathbf{\Sigma}_{obs}\right)^{-1} \tag{15}$$

$$\hat{\mathbf{x}}^{(i|i)} = \hat{\mathbf{x}}^{(i|i-1)} + \mathbf{K}^{(i)} \cdot \left(\mathbf{z}^{(i)} - \mathbf{H} \cdot \hat{\mathbf{x}}^{(i|i-1)}\right) \tag{16}$$

$$\mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(i|i)} = \mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(i|i-1)} - \mathbf{K}^{(i)} \cdot \mathbf{H} \cdot \mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(i|i-1)} \tag{17}$$

2.2.2 Results

The previously described calculations are carried out for $\Delta t = 1\text{s}$, $n = 16$, and for σ_{acc} and σ_{obs} as stated in figure 1. The resulting $x_1(t^{(i)})$, $z_1^{(i)}$, $\hat{x}_1^{(i)}$, and $\hat{x}_1^{(i|i)}$, are plotted in figure 1 (top)[10].

The calculated standard deviations of the estimates are shown in the center plot of figure 1. The standard deviations associated with $z_1^{(i)}$, $\hat{x}_1^{(i)}$, $\hat{x}_1^{(i|i-1)}$, and $\hat{x}_1^{(i|i)}$, are σ_{obs}, $\sigma^{(i)} = \sqrt{\left(\mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(i)}\right)_{1,1}}$, $\sigma^{(i|i-1)} = \sqrt{\left(\mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(i|i-1)}\right)_{1,1}}$, and $\sigma^{(i|i)} = \sqrt{\left(\mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(i|i)}\right)_{1,1}}$, respectively. In this model problem, we have the luxury of knowing the true state of the car at each of the times $\{t^{(i)}\}_{i=0,\dots,n}$. This allows us to compute and plot the true error of each estimate at each time step in figure 1 (bottom). In each plot, the blue dotted line traces the entire evolution of the prediction, progressing from

[10] The true states $\mathbf{x}(t^{(i)})$ used to produce figures 1, 2, and 3, are randomly sampled from the model (2). This random sampling was performed multiple times in order to select a particularly illustrative set of images. The same set of true states was used for all three figures. Note that there is no random generator involved in section 5.

the updated estimate $\hat{\mathbf{x}}^{(i-1|i-1)}$ to the prediction $\hat{\mathbf{x}}^{(i|i-1)}$ to the updated estimate $\hat{\mathbf{x}}^{(i|i)}$ and so forth.

Figure 1 illustrates how the predictions produced by the Kalman-filter are *generally* the most accurate and how, *overall*, the update step reduces the estimation error as we move from the blue dot $\hat{x}_1^{(i|i-1)}$ back to the green line $\hat{x}_1^{(i|i)}$ in the bottom plot of the figure. Of course, due to the randomness involved both in the process and in the observation, these generalities do not hold *systematically*. In some cases, the estimate produced by the Kalman filter can even be the the least accurate of the plotted estimates for several time steps in a row. It does, however, minimize the expected squared error, as shown in the literature [4]. Finally, we can see by comparing the standard deviations in figure 1 (center), that the uncertainty of the estimate produced by the Kalman filter at any given time step is lower than that of the observations, while the uncertainty of the purely model-based estimate steadily grows between subsequent time steps. Note that it is possible for $\sigma^{(i|i)}$ to outgrow σ_{obs} in particularly "unlucky" cases where the model and the observations both do worse than their respective standard deviation would suggest. However, the uncertainty of the estimate is always reduced in the update step (eq. 17), as seen in the center plot.

The covariance matrix of the estimate $\hat{\mathbf{x}}^{(i|i)}$ is shown for each time step in figure 2. Figure 3 contains the corresponding matrices of correlation coefficients. As expected, the variance generally increases over time. The apparent division of the matrix into four two-by-two submatrices indicates that the variances of the first two coordinates of the state $\left(\hat{x}_1^{(i|i)} \text{ and } \hat{x}_2^{(i|i)}\right)$ are much greater than those of the remaining coordinates $\left(\hat{\dot{x}}_1^{(i|i)} \text{ and } \hat{\dot{x}}_2^{(i|i)}\right)$. Moreover, the correlation of the errors of $\hat{x}_1^{(i|i)}$ and $\hat{x}_2^{(i|i)}$ decreases over time, as the regular observations provide separate, equally accurate and independent estimates for these two variables. Naturally, after several seconds, the error of the velocity estimate $\hat{\dot{x}}_1^{(i|i)}$ is more strongly correlated with that of the distance estimate $\hat{x}_1^{(i|i)}$ than with that of $\hat{x}_2^{(i|i)}$, and vice versa. This is to be expected, considering that a false estimate of one coordinate of the velocity, say $\hat{\dot{x}}_1^{(i|i)}$, leads to the accumulation over time of a growing error on the corresponding coordinate of the distance, $\hat{x}_1^{(i|i)}$, while leaving the other coordinate unaffected.

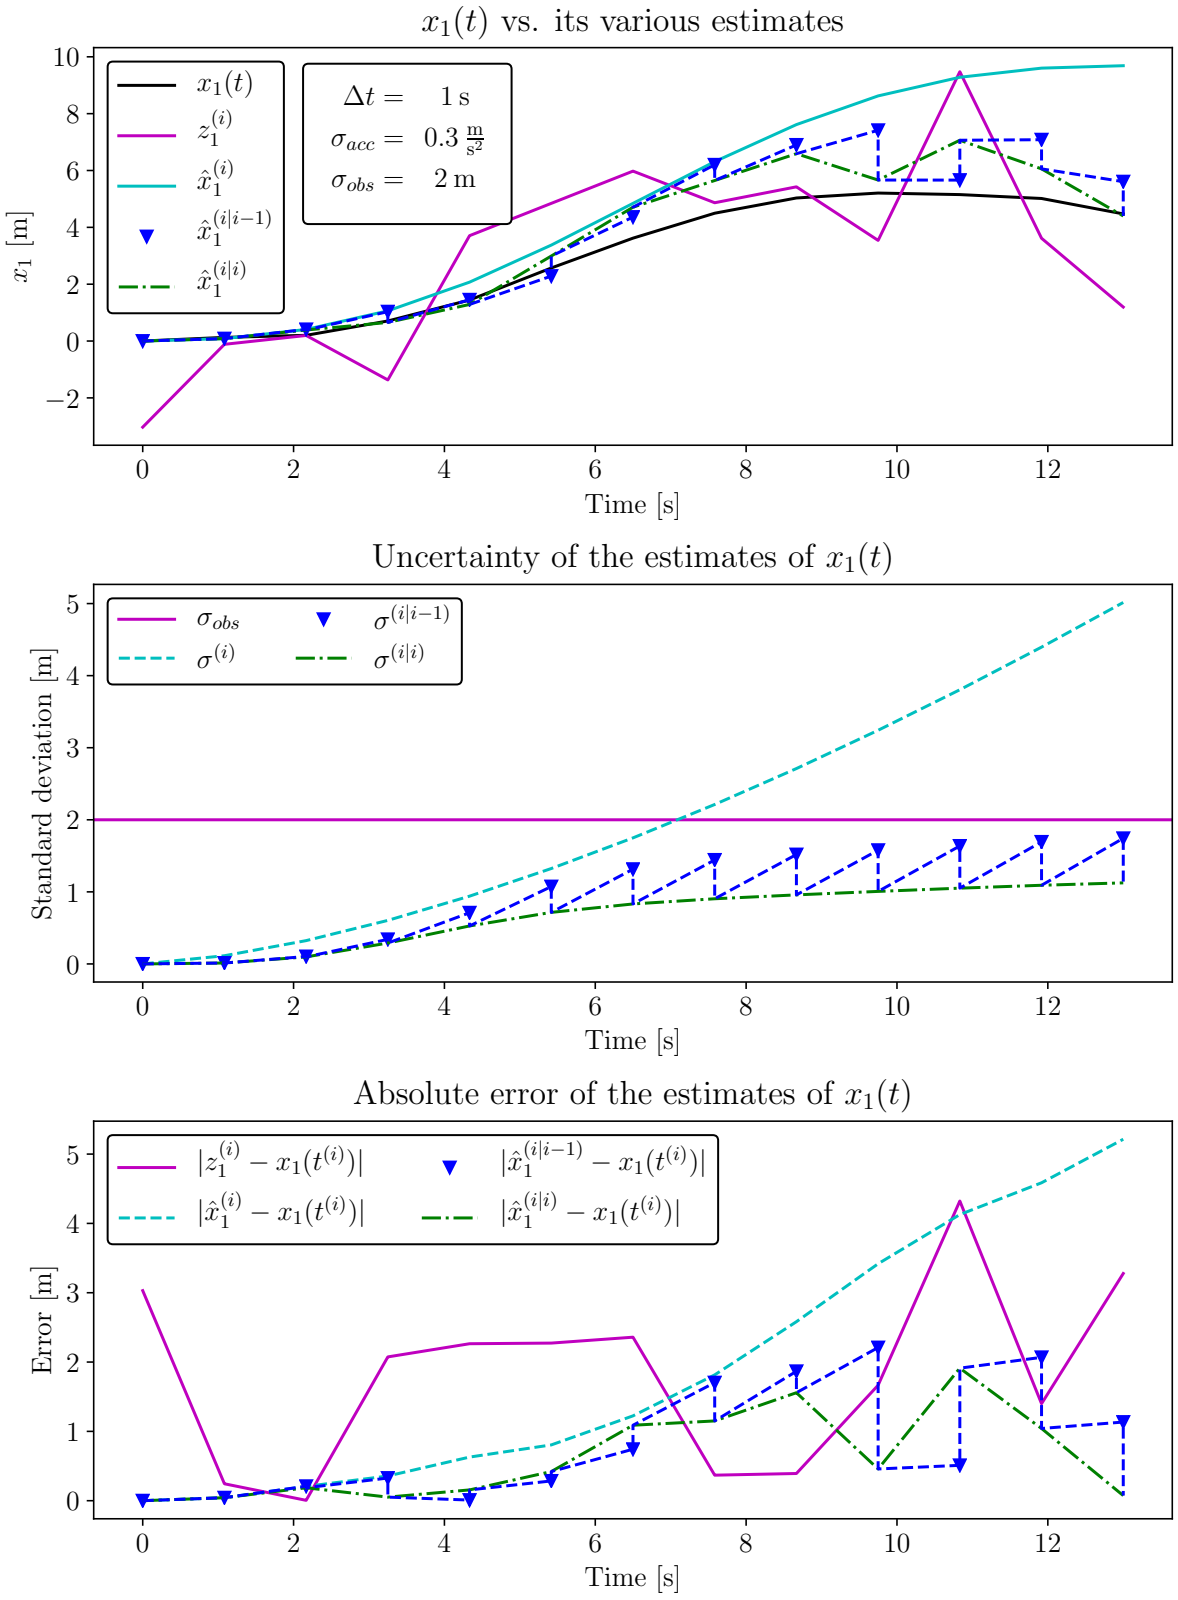

Figure 1: Estimation of the state variable $x_1(t)$ with the Kalman filter. This figure was obtained for $\Delta t = 1\text{s}$, $n = 13$, and for the values of σ_{acc} and σ_{obs} specified in the top plot.

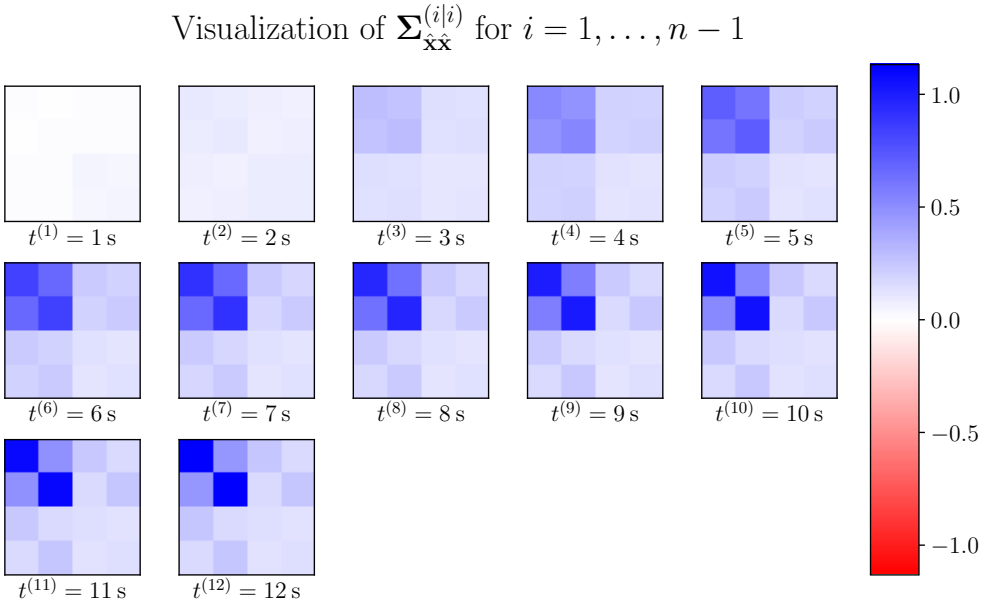

Figure 2: Visualization of the covariance matrix $\mathbf{\Sigma}^{(i|i)}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}$ obtained for $\Delta t = 1\mathrm{s}$, $n = 13$, $\sigma_{acc} = 0.3\,\mathrm{m\,s^{-2}}$ and $\sigma_{obs} = 2\,\mathrm{m}$.

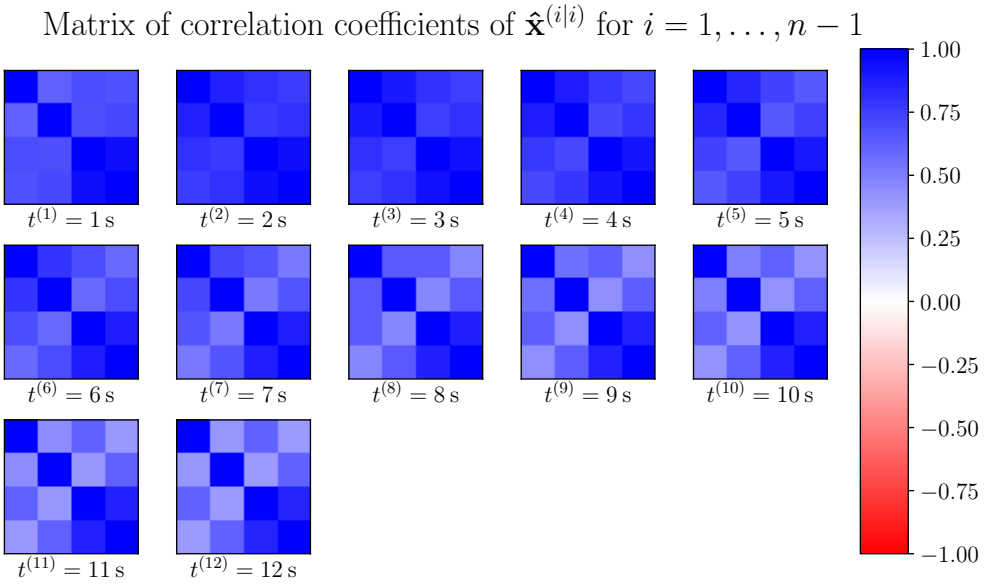

Figure 3: Visualization of the correlation coefficients between coordinates of the estimated state $\hat{\mathbf{x}}^{(i|i)}$. $\Delta t = 1\mathrm{s}$, $n = 13$, $\sigma_{acc} = 0.3\,\mathrm{m\,s^{-2}}$ and $\sigma_{obs} = 2\,\mathrm{m}$.

3 Theoretical Background

This section lays out the theory upon which the present work is based. On the one hand, there are the point reactor kinetics (PRK) equations, which form a simple model to describe the time-dependent behavior of thermonuclear reactors. The equations are first introduced from simplified physical principles in section 3.1.1. Then, typical situations such as reactor start-up and shut-down are discussed in section 3.1.2 based on these equations. On the other hand, this work is also founded on the unscented Kalman filter (UKF), which is contextualized in section 3.2.1 and explained in section 3.2.2.

3.1 Point Reactor Kinetics

Reactor kinetics is the study of power transients[11] in nuclear reactors. Such transients occur during normal start-up and shut-down, as well as in various accidental scenarios. Understanding these events is therefore an important task in reactor safety analysis [8]. The Point Reactor Kinetics equations describe this transient behavior under a set of simplifying assumptions.

The transient behavior of nuclear reactors is mainly driven by the *neutron flux* $\Phi(t, \mathbf{r}, E, \mathbf{\Omega})$, which represents the total distance travelled per unit time by neutrons moving through a point $\mathbf{r}$ in a given direction $\mathbf{\Omega}$ with a given kinetic energy E. In reality, the described quantity is subject to random fluctuations and the neutron flux can be interpreted as a time- and space-dependent instantaneous average [8]. Most importantly, the neutron flux determines the rate of nuclear fission reactions, which in turn dictates the rate at which heat is released in the reactor — in other words: the thermal reactor power.

3.1.1 Interpretation and Mathematical Formulation

In essence, the rate of change of the neutron flux equals the difference between the rate of neutron production on one side, and the rates of neutron absorption[12] and leakage[13] on the other:

$$\frac{d}{dt}\Phi = Production - (Absorption + Leakage) \tag{18}$$

However, all of these rates depend on the neutron flux itself:

$$\begin{aligned} Production &= p(\Phi, \dots) \\ Absorption &= a(\Phi, \dots) \\ Leakage &= l(\Phi, \dots) \end{aligned} \tag{19}$$

[11] The term *transient* here denotes the evolution of a physical quantity which is not constant in time.

[12] The term *neutron absorption* accounts for neutrons that cause the fission of a fuel nucleus, as well as those which are captured by non-fuel nuclei.

[13] Some neutrons cross the boundary of the reactor and are therefore lost to the chain reaction of nuclear fission taking place inside it.

Due to this interdependence, the neutron flux is governed by a differential equation known as the *transport equation* [8].
The rate of neutron production, in turn, is almost entirely (generally more than 99.3 %) accounted for by neutrons created in the fission of fuel nuclei. Crucially, however, this is not the only neutron source: Firstly, external sources can be applied. And secondly, some unstable fission products have a (low) non-zero probability of releasing neutrons after a random delay [11]. These fission products are generally referred to as *neutron precursors*. The important thing to note here is that this phenomenon introduces a delay between the fission of a fuel nucleus and the emission of a neutron from one of the fission products. Neutrons emitted by neutron precursors are therefore referred to as *delayed neutrons*, as opposed to the *prompt neutrons* created in the fission process itself.
The existence of delayed neutrons makes it possible for the neutron population to rise slowly. Indeed, the "time taken for one neutron to continue the chain process by fissioning the next fuel nucleus" is very short: it is generally of the order of 10^{-3} s to 10^{-6} s, depending on reactor type [11]. The average time for one neutron to cause another fission is known as *reproduction time* or *mean neutron generation time*. If, on average[14], each neutron produces more than one prompt neutron (a state known as *prompt supercriticality*), then the neutron flux increases exceedingly fast due to the short reproduction time — fast enough to melt the entire fuel within a fraction of a second. Prompt supercriticality must therefore be avoided at all times. Luckily, the delayed neutrons create a viable window where the following two conditions are satisfied simultaneously: Firstly, and most importantly, the prompt neutrons produced by one fissioning nucleus cause slightly *less* than one other nucleus to fission on average[14] (the reactor is <u>not</u> prompt supercritical). Yet, secondly, when the prompt and delayed neutrons produced by one fissioning nucleus are added together, they cause slightly *more* than one other nucleus to fission on average[14] and the reactor power increases over time. The delayed neutrons, which tip the balance in this situation, are added to the system significantly later, leaving enough time for control mechanisms to intervene and for heat to be transferred away from the fuel.
Due to their critical role, neutron precursors and delayed neutrons must be included in any model attempting to accurately represent the transient behaviour of a reactor that is not prompt supercritical.

The different neutron precursors are empirically divided into a number of *precursor groups* based on their half life time [8]. The rate of decay of each neutron precursor group is proportional to the group's concentration in the fuel. On the other hand, the rate of production of neutron precursors is roughly proportional to the reaction rate of nuclear fission in the fuel and therefore also to the neutron flux:

$$\tfrac{d}{dt}Concentration_\ell = Production_\ell(\Phi, \dots) - Decay_\ell(Concentration_\ell, \dots) \tag{20}$$

This results in a coupled system of differential equations, consisting of the transport equation and the equations which govern the concentration of each precursor group.
Much like the neutron flux $\Phi(t, \mathbf{r}, E, \mathbf{\Omega})$, the dependent variables in this set of equations are functions of time, space, energy and direction. The *point reactor kinetics* (PRK) equations are

[14] The averages marked with this footnote include all neutrons, including those which leave the system.

obtained by assuming isotropic[15] conditions in a homogeneous[16] reactor and then averaging over the energy range through weighted integration, thus leaving time as the only remaining independent[6] variable. They are derived in [8] and stated as follows:

$$\frac{d}{dt}n(t) = \frac{\rho(t) - \beta(t)}{\Lambda(t)}n(t) + \sum_{\ell} \lambda_\ell c_\ell(t) + q(t) \tag{21}$$

$$\frac{d}{dt}c_\ell(t) = \frac{\beta_\ell(t)}{\Lambda(t)}n(t) - \lambda_\ell c_\ell(t) \quad \text{for each precursor group } \ell \tag{22}$$

The following variables and constants are used in eqs. (21) and (22):

- $n(t)$: The *neutron population* actually represents the number density of neutrons per unit volume, but it is proportional both to the total neutron population in the reactor (due to the assumption of homogeneity) and to the thermal reactor power. It is closely related to the neutron flux, Φ.
- $c_\ell(t)$: This variable represents the concentration (number density) of neutron precursors of group ℓ. It is proportional to the total number of neutron precursors of group ℓ due to the assumption of homogeneity.
- $\rho(t)$: The *reactivity* $\rho(t)$ is defined as:

  $$\rho(t) = \frac{k_{eff}(t)-1}{k_{eff}(t)} \tag{23}$$

  With k_{eff} denoting the ratio of neutron production to losses (absorption and leakage) in the neutron balance equation (eq. 18). k_{eff} is known as the *effective multiplication factor*.
  In the point reactor kinetics (PRK) model, $\rho(t)$ is viewed as an external constraint imposed on the reactor. This is an approximation, because internal variables of the reactor, such as the fuel and moderator temperatures, influence the reactivity.
- $\beta_\ell(t)$: The *fractional yield* of each precursor group ℓ represents the proportion of neutrons emitted as delayed neutrons by precursors of group ℓ, relative to the total number of neutrons emitted (delayed and prompt). It varies depending on the fuel composition and the reactor type, among others.
- $\beta(t)$: The *delayed neutron fraction* corresponds to the proportion of neutrons emitted as delayed neutrons (by any precursor), relative to the total number of neutrons

[15] Under *isotropic* conditions, all physical quantities are independent of the direction. In the case of the transport equation, isotropic conditions correspond to the assumption that the number of neutrons traversing any given point, at any given speed, during any given time window and in any given direction, will, in expectation, be equal for all possible directions.

[16] A physical system is called *homogeneous* if, and only if, the physical quantities of the system are independent of the actual location inside the system. The assumption of a homogeneous reactor is a considerable simplification, since it implies that the neutron flux is the same in the moderator, as it is in the fuel.

emitted (delayed and prompt). It is clear from their definitions that the delayed neutron fraction is simply the sum of each precursor group's fractional yield:

$$\beta(t) = \sum_{\ell} \beta_\ell(t) \tag{24}$$

- $\Lambda(t)$ [s] : Conceptually, the *mean neutron generation time*, also called *reproduction time*, is described as the mean time from neutron emission until it causes the fission of another fuel nucleus [11]. It is defined as follows:

$$\Lambda(t) = \frac{n(t)}{RR(t)\cdot\nu} \tag{25}$$

  Where $RR(t)$ denotes the reaction rate of fission and ν is the total number of neutrons emitted per fission. Note that ν accounts for delayed neutrons as well as prompt neutrons. The product $RR(t) \cdot \nu$ hence represents the rate at which prompt neutrons and precursors are produced, without accounting for the delay between the creation of a neutron precursor and the emission of a delayed neutron.

- λ_ℓ [s^{-1}] : Neutron precursors are characterized by the fact that they undergo radioactive decay. Each *decay constant* λ_ℓ characterizes the rate of decay of precursor group ℓ. Conceptually, it is the inverse of the mean lifetime of nuclei in said group.

- $q(t)$ [s^{-1}]: Finally, $q(t)$ represents the neutrons added to the system by an external source.

Eqs. (21, 22) can be written in matrix form. The state vector is labelled $\mathbf{d}(t)$, to avoid confusion with the state vector $\mathbf{x}(t)$ used in the Kalman filter. The resulting equations are:

$$\mathbf{d}(t) = [n(t), c_1(t), \ldots, c_{N_\ell}(t)]^T \tag{26}$$

$$\frac{d}{dt}\mathbf{d}(t) = \mathbf{A}(t)\cdot\mathbf{d}(t) + \mathbf{q}(t)$$

$$\mathbf{A}(t) = \begin{bmatrix} \frac{\rho(t)-\beta(t)}{\Lambda(t)} & \lambda_1 & \cdots & \lambda_{N_\ell} \\ \frac{\beta_1(t)}{\Lambda(t)} & -\lambda_1 & & (0) \\ \vdots & & \ddots & \\ \frac{\beta_{N_\ell}(t)}{\Lambda(t)} & (0) & & -\lambda_{N_\ell} \end{bmatrix} \qquad \mathbf{q}(t) = \begin{bmatrix} q(t) \\ 0 \\ \vdots \\ 0 \end{bmatrix} \tag{27}$$

3.1.2 Typical Power Transients

Figures 4 shows the evolution of the neutron population (normalized with its initial value) over time if the reactivity jumps from 0 to a certain value at time $t = 0$. Before the reactivity insertion, the reactor is at equilibrium state: $\frac{d}{dt}n(t) = \frac{d}{dt}c_\ell(t) = 0$. For the supercritical, but not prompt supercritical, reactivity $\rho = 0.15 \cdot \beta$, the neutron population sees a sudden increase

by about 20 %, followed by a much slower exponential increase. This sudden change is known as the *prompt jump*. It is equally present in the case of a negative reactivity insertion, for which the neutron population suddenly decreases by a fixed amount and then slowly decays over time.

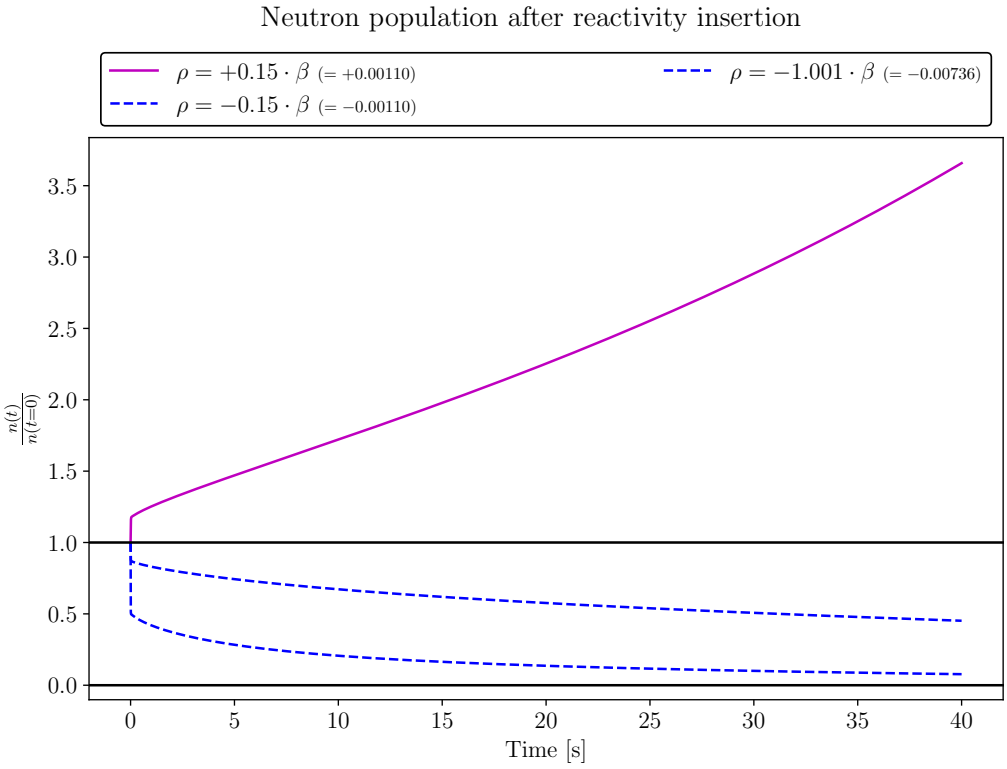

Figure 4: Normal power transients during operation: positive reactivity safely below β for reactor start-up, and negative reactivity for reactor shut-down.

The evolution of the neutron population in a prompt supercritical reactor is shown on a semilogarithmic scale in figure 5. The reactivity is merely 0.1 % above the prompt supercriticality threshold, and yet he initial neutron population is multiplied billionfold in the first 2 s — and the theoretical reactor power with it. In reality, reactors are designed in such a way that the reactivity does not remain constant during such a power excursion. Instead, it is lowered by *negative temperature feedback* when the fuel and moderator temperatures rise as a consequence of the spike in power. Still, it is worth noting that the same amount of negative reactivity induces a relatively slow decrease of the neutron population and the thermal reactor power, as seen in figure 4. After the full 40 s, while the supercritical reactor is over 100 orders of magnitude above the initial neutron population, the subcritical one is still at 8 % thereof. This is due to the slow decay of the neutron precursors, which continue to produce neutrons and heat for after shut-down. The heat produced by the neutron precursors' decay after shut-down is known as *decay heat*.

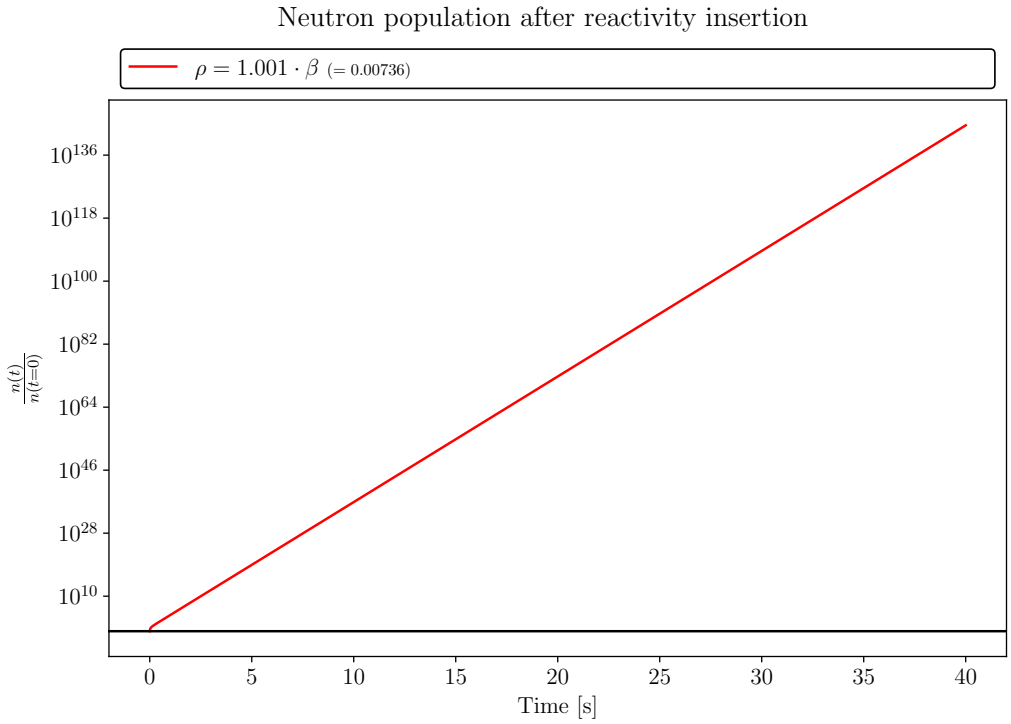

Figure 5: Semilogarithmic plot of a prompt supercritical excursion with constant reactivity.

3.2 The Unscented Kalman Filter

In this work, the point reactor kinetics (PRK) model is combined with experimental observations with the double aim of yielding a more reliable prediction of the time-dependent reactor behavior and more accurate estimates of the reactor properties and the reactivity. The latter are viewed as independent variables for simplicity. This double goal is achieved by applying the unscented Kalman filter (UKF) to the PRK equations and using the resulting model to filter the observations. The UKF is compared with other algorithms in section 3.2.1 and thoroughly explained in section 3.2.2.

3.2.1 Context

The PRK equations (26, 27) form a nonlinear model and thus cannot be written as in the linear eqs. (2, 3). For instance, the reactivity $\rho(t)$, which is viewed as a control input, is multiplied with the neutron population $n(t)$ (a state variable) in eq. (21) to obtain the rate of change of the neutron population. In the notation used for eq. (2), this corresponds to the multiplication of $\mathbf{x}(t^{(i)})$ with $\mathbf{a}(t^{(i)})$. The linear Kalman filter introduced in section 2 is therefore unfit for general application to the PRK equations, although it may be used for state estimation in the simplest case where the matrix $\mathbf{A}(t)$ from eq. (27) does not depend on the state $\mathbf{d}(t)$ and is taken to be correctly known.

A natural approach to overcoming this issue is to apply the Kalman filter to a linearized version

of the nonlinear model. This approach is known as the *extended Kalman filter* (EKF), where the state transition matrix and the observation matrix are computed at each time step as the first term in a Taylor series expansion of the general state transition and observation models, respectively [15]. Thus, the EKF requires Jacobians to be computed for both models at each time step [7]. Two main issues of the EKF are raised in the literature: For one, the linearization of a model introduces an approximation error which may be unacceptably high for some models [15]. However, it has been shown that a first-order approximation using Taylor series expansion is sufficient for the accurate solution of the PRK equations [12]. Secondly, it has been pointed out that certain non-linear models, such as neural network models, bring a rich, multi-model structure to the distribution of the state, which the Gaussian approximation fails to capture appropriately [2]. Nevertheless, the EKF has successfully been used to model reactor kinetics using the PRK equations [7].

Both of the aforementioned issues of the EKF are addressed by Monte-Carlo methods, such as the *particle filter* and the *ensemble Kalman filter* (EnKF) [2]. These filters involve the discrete approximation of the probability density function[17] of the state by means of a large number of weighted sample points ("particles") which are scattered over the state space[18] and propagated through the system at each time step [1, 3]. One of the main drawbacks of the particle filter is its computational complexity, caused by the large number of particles required to provide sufficient accuracy [15]. Moreover, in certain applications, the particles tend to collapse to a point after a number of time steps [1].

The *unscented Kalman filter* (UKF) achieves higher-order accuracy than the EKF with the same order of computational complexity [15]. Like the particle filter, it uses a sampling approach and does not require the explicit computation of derivatives. The main difference lies in the selection of the sample points or particles: The UKF is built around the entirely deterministic *unscented transform* (UT), which selects a minimal set of weighted sample points called *sigma points* [15, 14]. Figure 6 (taken from [15]) illustrates the difference in the prediction step between the particle filter (left, based on general sampling methods), the EKF (center), and the UKF (right, based on the UT).
The sigma points are a special set of particles, chosen to entirely capture the mean and covariance of any distribution with as little particles as possible [15, 14].

By using the UT for particle selection, the UKF requires orders of magnitude less sample points than general sampling methods such as Monte-Carlo estimation, resulting in reduced computational complexity [15]. The drawback is that, contrary to general sampling methods, only the first two moments (mean and covariance) of the state distribution are accurately tracked in the unscented transform. Compared to the EKF, however, the mean and covariance are tracked more accurately, especially so in the presence of strong nonlinearities. Although this advantage seems to be of little importance in the context of point reactor kinetics, it is not

[17] The distribution of an n-dimensional random variable $\mathbf{X}$ is entirely described by its *probability density function* (PDF), commonly denoted as $f(\mathbf{x})$. The PDF $f(\mathbf{x})$ is the derivative of the *cumulative density function* (CDF) $F(\mathbf{x})$, which is itself defined as the probability that a random vector sampled from $\mathbf{X}$ is at most $\mathbf{x}$: $F(\mathbf{x}) = \mathbb{P}[\mathbf{X} \leq \mathbf{x}]$

[18] The *state space* is the set of all values which the state could possibly take.

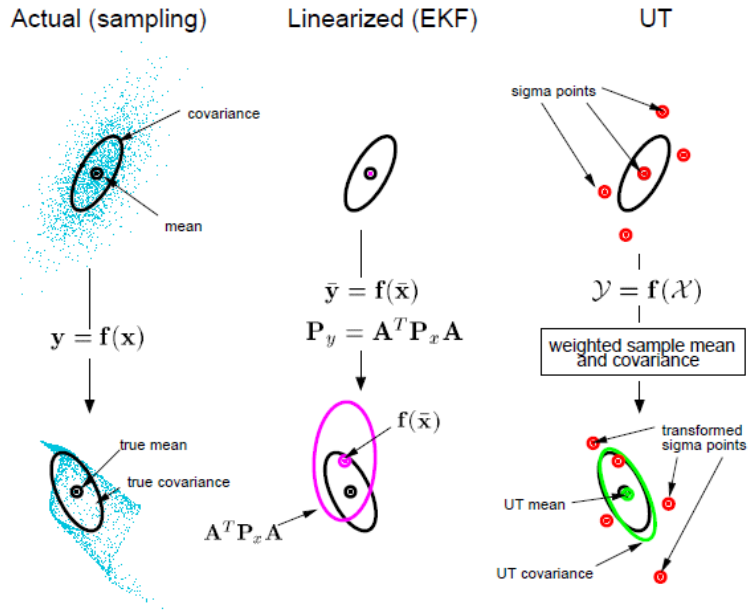

Figure 6: Illustration of the prediction step for three different filters: particle filter (left), EKF (center), and UKF (right). Taken from [15].

mitigated by known major disadvantages and has therefore led to the spread of the UKF and its inclusion in the open-source pykalman library for Python [13, 14, 15, 12].

Both the EKF and the UKF may be used for parameter estimation [7]. This is accomplished by including the parameters of interest in the state along with the system's dependent variables. The state transition function is then specified in such a way, that the parameters remain constant in the absence of noise. With each observation, eq. (16 b) causes their values to be updated depending on their influence on the observed quantities, as well as on their previous (estimated) uncertainty.

3.2.2 Mathematical Formulation

The unscented Kalman filter is designed for a general nonlinear model, consisting of a general state transition function F (eq. 28) and a general observation function H (eq. 29):

$$\mathbf{x}(t^{(i+1)}) = F\left(\mathbf{x}(t^{(i)}), \mathbf{w}^{(i)}, t^{(i)}\right) \tag{28}$$

$$\mathbf{z}^{(i)} = H\left(\mathbf{x}(t^{(i)}), \mathbf{v}^{(i-1)}, t^{(i)}\right) \tag{29}$$

Control inputs may be accounted for by specifying a dependence of the functions F or H on the independent variable $t^{(i)}$.

Like the linear Kalman filter, the UKF is divided into a *prediction step* and an *update step*. Before the prediction step, sigma points $\boldsymbol{\mathcal{X}}^{(i-1|i-1)}$ are computed to represent the distributions

of the process noise $\mathbf{w}^{(i-1)}$ and the observation noise $\mathbf{v}^{(i-1)}$, as well as the estimated distribution of the state $\mathbf{x}\left(t^{(i-1)}\right)$:

$$\mathbf{w}^{(i-1)} \sim \mathcal{N}\left(\mathbf{0}, \mathbf{\Sigma}_{process}\right) \qquad \mathbf{v}^{(i-1)} \sim \mathcal{N}\left(\mathbf{0}, \mathbf{\Sigma}_{obs}\right)$$

$$\mathbf{x}\left(t^{(i-1)}\right) \textit{ follows approximately } \mathcal{N}\left(\hat{\mathbf{x}}^{(i-1|i-1)}, \mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(i-1|i-1)}\right)$$

The sigma points are computed based on the sigma point function $S(\boldsymbol{\mu}, \mathbf{\Sigma})$, which is described in the literature [15, 14] (cf. figure 6). Their associated weights $\boldsymbol{\mathcal{W}}$ are constants[19].

$$\begin{bmatrix} \boldsymbol{\mathcal{X}}_x^{(i-1|i-1)} \\ \boldsymbol{\mathcal{X}}_w^{(i-1)} \\ \boldsymbol{\mathcal{X}}_v^{(i-1)} \end{bmatrix} = S\left(\begin{bmatrix} \hat{\mathbf{x}}^{(i-1|i-1)} \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix}, \begin{bmatrix} \mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(i-1|i-1)} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{\Sigma}_{process} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{\Sigma}_{obs} \end{bmatrix}\right) \tag{30}$$

These sigma points are subsequently used in the prediction step, to approximate the distributions of the state $\mathbf{x}\left(t^{(i)}\right)$ and the observation $\mathbf{z}^{(i)}$. First, the sigma points are transformed by the state transition and observation functions, respectively:

$$\boldsymbol{\mathcal{X}}_x^{(i|i-1)} = F\left(\boldsymbol{\mathcal{X}}_x^{(i-1|i-1)}, \boldsymbol{\mathcal{X}}_w^{(i-1)}, t^{(i)}\right) \tag{31}$$

$$\boldsymbol{\mathcal{Z}}^{(i)} = H\left(\boldsymbol{\mathcal{X}}_x^{(i|i-1)}, \boldsymbol{\mathcal{X}}_v^{(i-1)}, t^{(i)}\right) \tag{32}$$

Subsequently, the mean and the covariance of the estimated distribution of $\mathbf{x}\left(t^{(i)}\right)$ and $\mathbf{z}^{(i)}$ are computed as the empirical mean and variance of the transformed sigma points:

$$\hat{\mathbf{x}}^{(i|i-1)} = \boldsymbol{\mathcal{W}} \cdot \boldsymbol{\mathcal{X}}_x^{(i|i-1)} = \sum_k \mathcal{W}_k \cdot \left(\boldsymbol{\mathcal{X}}_x^{(i|i-1)}\right)_k \tag{33}$$

$$\mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(i|i-1)} = \sum_k \mathcal{W}_k \left[\left(\boldsymbol{\mathcal{X}}_x^{(i|i-1)}\right)_k - \hat{\mathbf{x}}^{(i|i-1)}\right] \left[\left(\boldsymbol{\mathcal{X}}_x^{(i|i-1)}\right)_k - \hat{\mathbf{x}}^{(i|i-1)}\right]^T \tag{34}$$

$$\hat{\mathbf{z}}^{(i|i-1)} = \boldsymbol{\mathcal{W}} \cdot \boldsymbol{\mathcal{Z}}^{(i)} = \sum_k \mathcal{W}_k \cdot \boldsymbol{\mathcal{Z}}_k^{(i)} \tag{35}$$

$$\mathbf{\Sigma}_{\hat{\mathbf{z}}\hat{\mathbf{z}}}^{(i|i-1)} = \sum_k \mathcal{W}_k \left[\boldsymbol{\mathcal{Z}}_k^{(i)} - \hat{\mathbf{z}}^{(i)}\right] \left[\boldsymbol{\mathcal{Z}}_k^{(i)} - \hat{\mathbf{z}}^{(i)}\right]^T \tag{36}$$

The covariance of the estimated distribution of $\mathbf{x}\left(t^{(i)}\right)$ with that of $\mathbf{z}^{(i|i-1)}$ is computed analogously as:

$$\mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{z}}}^{(i|i-1)} = \sum_k \mathcal{W}_k \left[\left(\boldsymbol{\mathcal{X}}_x^{(i|i-1)}\right)_k - \hat{\mathbf{x}}^{(i|i-1)}\right] \left[\boldsymbol{\mathcal{Z}}_k^{(i)} - \hat{\mathbf{z}}^{(i|i-1)}\right]^T \tag{37}$$

The update step is similar to eqs. (15, 16, 17) of the linear Kalman filter. By introducing the mean $\hat{\mathbf{z}}^{(i|i-1)} = \mathbf{H} \cdot \hat{\mathbf{x}}^{(i|i-1)}$ and covariance $\mathbf{\Sigma}_{\hat{\mathbf{z}}\hat{\mathbf{z}}}^{(i|i-1)} = \mathbf{H} \cdot \mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(i|i-1)} \cdot \mathbf{H}^T$ of the estimated distribution

[19] The weights associated with the sigma points depend only on the dimensionality of the state space, which also determines the number of sigma points.

function for the i-th observation $\mathbf{z}^{(i)}$, as well as the covariance $\mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{z}}}^{(i|i-1)} = \mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(i|i-1)} \cdot \mathbf{H}^T$, we can rewrite eqs. (15, 16) as:

$$\mathbf{K}^{(i)} = \mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{z}}}^{(i|i-1)} \cdot \left(\mathbf{\Sigma}_{\hat{\mathbf{z}}\hat{\mathbf{z}}}^{(i)} + \mathbf{\Sigma}_{obs} \right)^{-1} \tag{15 b}$$

$$\hat{\mathbf{x}}^{(i|i)} = \hat{\mathbf{x}}^{(i|i-1)} + \mathbf{K}^{(i)} \cdot (\mathbf{z}^{(i)} - \hat{\mathbf{z}}^{(i)}) \tag{16 b}$$

$$\mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(i|i)} = \mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(i|i-1)} - \mathbf{K}^{(i)} \cdot \mathbf{\Sigma}_{\hat{\mathbf{z}}\hat{\mathbf{x}}}^{(i|i-1)} \tag{17 b}$$

$\hat{\mathbf{x}}^{(i|i)}$ and $\mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(i|i)}$ are indeed computed by eq. (16 b), respectively eq. (17 b), in the UKF. The Kalman gain $\mathbf{K}^{(i)}$ is computed without $\mathbf{\Sigma}_{obs}$, because the observation noise is already taken into account in eqs. (30) and (32):

$$\mathbf{K}^{(i)} = \mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{z}}}^{(i|i-1)} \cdot \left(\mathbf{\Sigma}_{\hat{\mathbf{z}}\hat{\mathbf{z}}}^{(i)} \right)^{-1} \tag{38}$$

4 Application of the Unscented Kalman Filter to the Point Reactor Kinetics Equations

The unscented Kalman filter is applied to the point reactor kinetics (PRK) model with $N_\ell = 6$ precursor groups. The delayed neutron fraction β, the fractional yields β_ℓ and the mean generation time Λ are assumed to be constant with respect to time. Measurements from an experiment conducted on 19/10/2016 inside the experimental zero-power reactor CROCUS at EPFL are used as an example application. In this experiment, the reactor was subjected to a step reactivity insertion, so the reactivity ρ is equally taken to be constant. β, β_ℓ, Λ, λ_ℓ and ρ may safely be viewed as constant in the particular case of CROCUS, mainly because temperature variations and the fuel burn-up remain negligible at all times. They are hence referred to as *independent variables*[6]. The independent variables are estimated along with the (still) dependent variables: the neutron population $n(t)$ and the concentration of neutron precursors of each group, $c_\ell(t)$. The model is implemented in Python 3.7.3, using the open-source libraries pykalman, NumPy, SciPy, and Matplotlib.

4.1 Experimental Set-up

The measurement taken in CROCUS on 19/10/2016 consists of a series of neutron counts, $z^{(i)}$, taken at intervals of 0.1 %. The neutron counts are already recorded while the reactor is brought to critical state and stabilised. Then, one fully inserted control rod is quickly raised by 600mm, which corresponds to a step reactivity insertion of 112 ± 6 pcm[20].

The exact time of the reactivity insertion is not marked in the sequence of recorded counts. The method by which it is estimated is illustrated in Figure 7: In a first step, the initial 1000 observations are smoothed. The time at which the smoothed signal reaches its minimum is

[20] The unit pcm stands for *pour cent-mille*. It is simply a factor of 10^{-5}.

selected as an approximation of the time at which the reactivity insertion took place. The first neutron count after that is denoted as $z^{(0)}$. This approach is not generally applicable, because the minimum value of the smoothed signal could lie anywhere before the reactivity insertion. It is therefore recommended for future experiments, that the time at which the reactivity insertion occurs be measured with sufficient precision to identify the corresponding observation. Nevertheless, the described method of estimation is sensible in this particular case, because the minimum value is located near the change of overall slope.

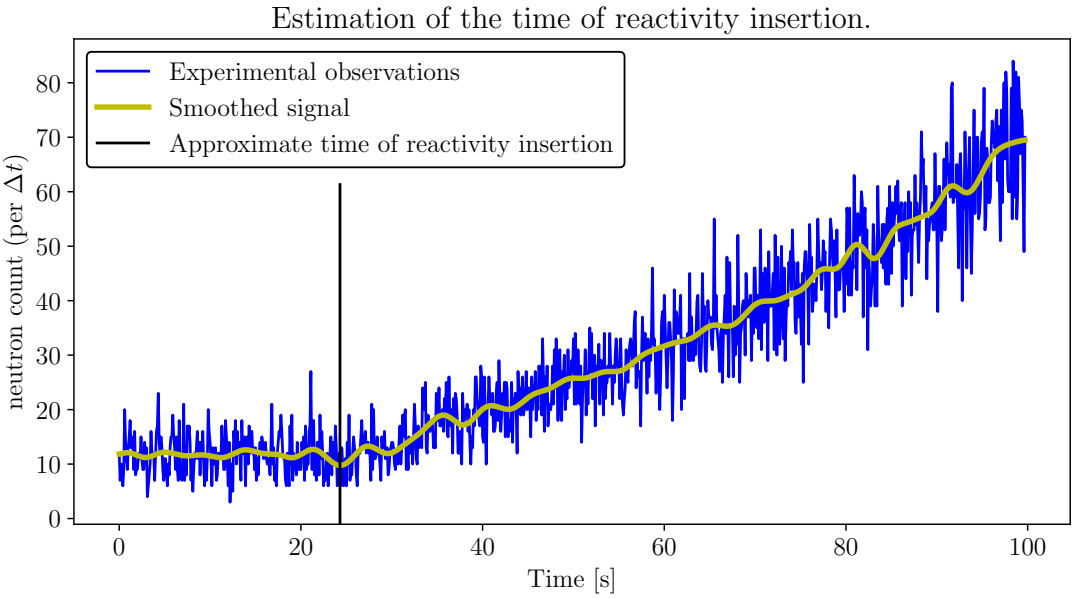

Figure 7: Estimating the time at which the reactivity insertion took place. The first 1000 observations are smoothed and the minimum of the smoothed signal is found.

4.2 State

In eq. (26), the state $\mathbf{d}(t)$ is defined as a vector containing the $N_\ell+1 = 7$ dependent variables of the point reactor kinetics equations. For the purpose of parameter estimation, a $3 \cdot N_\ell+3 = 21$-dimensional state vector $\mathbf{x}(t)$ is defined, containing the independent variables along with the dependent variables:

$$\mathbf{x}(t) = [n(t), c_1(t), \ldots, c_6(t), \rho, \beta_1, \cdots, \beta_6, \lambda_1, \cdots, \lambda_6, \Lambda]^T \tag{39}$$

Note that the delayed neutron fraction β is not included in the state as it is fully determined by the fractional yields $\beta_1, \cdots, \beta_6$ (cf. eq. 24).

4.3 State Transition Function and Process Noise

As discussed in section 3.2.1, the state transition function must be specified in such a way, that the independent variables remain constant in the absence of noise. Given the state $\mathbf{x}(t^{(i)})$ and a process noise vector $\mathbf{w}^{(i)}$, the state transition function $F\left(\mathbf{x}(t^{(i)}), \mathbf{w}^{(i)}, t^{(i)}\right)$ returns the state $\mathbf{x}(t^{(i+1)})$ as follows:

1. The state is split into a vector of dependent variables, $\mathbf{d}$, and a vector of independent variables, $\mathbf{e}$:

$$\mathbf{d}(t^{(i)}) = \begin{bmatrix} x_1(t^{(i)}) \\ \vdots \\ x_7(t^{(i)}) \end{bmatrix} = \begin{bmatrix} n(t^{(i)}) \\ c_1(t^{(i)}) \\ \vdots \\ c_6(t^{(i)}) \end{bmatrix} \qquad \mathbf{e}(t^{(i)}) = \begin{bmatrix} x_8(t^{(i)}) \\ \vdots \\ x_{21}(t^{(i)}) \end{bmatrix} = \begin{bmatrix} \rho \\ \beta_1 \\ \vdots \\ \beta_6 \\ \lambda_1 \\ \vdots \\ \lambda_6 \\ \Lambda \end{bmatrix} \tag{40}$$

2. The matrix $\mathbf{A}(t^{(i)})$ (cf. eq. 27) is computed using the independent variables $\mathbf{e}(t^{(i)})$.

3. The system in eq. (27) is solved in the interval $\left[t^{(i)}, t^{(i+1)}\right]$ for $\mathbf{A}(t) = const. = \mathbf{A}(t^{(i)})$, setting the initial condition $\mathbf{d}_0 = \mathbf{d}(t^{(i)})$. It is solved using the explicit Runge-Kutta method of order 5, which is described in [6] and implemented in the open-source SciPy package.

4. As explained in the next paragraph, the noise is factored in multiplicatively for the dependent variables only, using the solution of the previous step, $\mathbf{d}^*(t^{(i+1)})$:

$$\mathbf{d}(t^{(i+1)}) = \left[d_1^*(t^{(i+1)}) \cdot \left(1 + w_1^{(i)}\right), \cdots, d_7^*(t^{(i+1)}) \cdot \left(1 + w_7^{(i)}\right)\right]^T \tag{41}$$

5. $\mathbf{d}(t^{(i+1)})$ is concatenated back together with $\mathbf{e}(t^{(i)})$:

$$\mathbf{x}(t^{(i+1)}) = F\left(\mathbf{x}(t^{(i)}), \mathbf{w}^{(i)}, t^{(i)}\right) = \begin{bmatrix} \mathbf{d}(t^{(i+1)}) \\ \mathbf{e}(t^{(i)}) \end{bmatrix} \tag{42}$$

For the UKF as well as the linear Kalman filter, the process noise $\mathbf{w}^{(i)}$ is defined by eq. (3) as a multivariate Gaussian random variable (GRV) with covariance $\mathbf{\Sigma}_{process}$ and mean $\mathbf{0}$. However, contrary to the linear filter, $\mathbf{w}^{(i)}$ can now contribute to the state transition function in a variety of ways. A distinction is made between the independent variables $\mathbf{e}$ and the dependent variables $\mathbf{d}$:
The former, $\mathbf{e}$, may safely viewed as unchanging, at least for the duration of the experiment. This is realized by specifying the uncertainty of their prior estimates in the covariance of the

initial state $\mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(0|0)}$ (cf. section 4.5) and leaving $\mathbf{e}$ entirely unaffected by process noise. This suggests that the process noise needs only be $N_\ell + 1 = 7$-dimensional, however the pykalman package requires that $\mathbf{w}$ and $\mathbf{x}$ have the same number of dimensions. Furthermore, the variance of each coordinate of $\mathbf{w}^{(i)}$ must be strictly positive for $\mathbf{\Sigma}_{process}$ to be positive definite, which is required by the sigma function $S(\boldsymbol{\mu}, \mathbf{\Sigma})$ in eq. (30). For that reason, the last 14 coordinates of $\mathbf{w}^{(i)}$ are simply ignored.
As seen in eq. (41), the dependent variables $\mathbf{d}$ are affected by process noise in a multiplicative way. This is partly motivated by the fact that they increase by several orders of magnitude between the beginning and the end of the measurement. Adding to $\mathbf{d}$ a noise term of constant variance would therefore have a vastly different impact at different times in the simulation. Instead, the process noise is viewed as a relative change caused by random effects (cf. eq. 41). This relative change is specified to have the same variance $\sigma_{process}^2 \ll 1$ for all dependent variables. The covariance of any two distinct coordinates of $\mathbf{w}$ is set to 0 for simplicity. As the coordinates of $\mathbf{w}$ corresponding to the independent variables are ignored, $\mathbf{\Sigma}_{process}$ is simply defined as a multiple of the identity matrix:

$$\mathbf{\Sigma}_{process} = \mathbf{I}_{21\times 21} \cdot \sigma_{process}^2 \tag{43}$$

4.4 Observation Function and Observation Noise

The single observed quantity $z^{(i)}$ is the number of neutrons recorded by the probe located in CROCUS in the interval $\left[t^{(i-1)}, t^{(i)}\right)$. The observation function $H\left(\mathbf{x}(t^{(i)}), v^{(i-1)}, t^{(i)}\right)$, along with the variance σ_{obs}^2 of the (zero-mean Gaussian) observation noise v, should be able to explain the observed neutron counts $z^{(i)}$ given the true states $\mathbf{x}(t^{(i)})$.

Each neutron count $z^{(i)}$ is viewed as a Poisson-distributed discrete random variable whose parameter is proportional to the neutron population:

$$z^{(i)} \sim Poisson(\mu^{(i)}) \qquad \mu^{(i)} = \mathbb{E}\left[z^{(i)}\right] = Var\left(z^{(i)}\right) \propto n(t^{(i)}) \tag{44}$$

The parameter $\mu^{(i)}$ is unknown, as is the scaling factor between $\mu^{(i)}$ and $n(t^{(i)})$. For simplicity, this scaling factor is set to one:

$$\mu^{(i)} = n(t^{(i)}) \text{ by assumption} \tag{45}$$

This assumption is equivalent to setting an initial condition for the system of coupled ordinary differential equations (ODEs) in eqs. (21, 22).

It is possible to obtain a truly Poisson-distributed random variable with parameter $\mu^{(i)} = n(t^{(i)})$ using the *quantile function*[21] of the Poisson distribution and the CDF of the Gaussian

[21] The *quantile function* $F_X^{-1}(p)$, also known as *percent point function*, is defined as the inverse of the cumulative distribution function (CDF) $F_X(x) = \mathbb{P}[X \leq x]$.

distribution:

$$v^{(i)} \sim \mathcal{N}(0,1) \implies F^{-1}_{Poiss\left(\mu^{(i)}\right)}\left(F_{\mathcal{N}(0,1)}\left(v^{(i)}\right) \right) \sim Poisson\left(\mu^{(i)}\right) \tag{46}$$

Nonetheless, as only the mean and the covariance of distributions are accurately tracked in the UKF, it is superfluous to model the third and higher moments of the observations' true distribution. On the contrary, it suffices to specify $H\left(\mathbf{x}(t^{(i)}), v^{(i-1)}, t^{(i)}\right)$ such that it yields a continuous GRV with mean and variance equal to μ, like so:

$$v^{(i)} \sim \mathcal{N}(0,1) \implies \mu^{(i)} + \sqrt{\mu^{(i)}} \cdot v^{(i)} \sim Poisson\left(\mu^{(i)}\right) \tag{47}$$

Accordingly, we specify the observation function and the observation noise as:

$$\sigma^2_{obs} = 1 \tag{48}$$

$$H\left(\mathbf{x}(t^{(i)}), v^{(i-1)}, t^{(i)}\right) = z^{(i)} + v^{(i)} \cdot \sqrt{z^{(i)}} \tag{49}$$

4.5 Initial State

The initial state is known to a certain accuracy. In the context of the UKF, it is viewed as a multivariate GRV, whose mean $\hat{\mathbf{x}}^{(0|0)}$ and covariance $\mathbf{\Sigma}^{(0|0)}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}$ must be specified. For simplicity, only the diagonal terms of the covariance matrix $\mathbf{\Sigma}^{(0|0)}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}$ are specified[22], while the extra-diagonal terms are set to 0. This is equivalent to believing that the errors on any two coordinates of the initial state are uncorrelated.

The means and variances of the reactor properties $\beta_\ell, \lambda_\ell, \Lambda$ are initialized with an estimate provided by the Laboratory for Reactor Physics and Systems Behaviour (LRS) at EPFL. The uncertainties reported on them are calculated by randomly sampling the nuclear data of the Serpent model from the ENDFB/VII.1 covariance matrices [10]. The reactivity ρ, on the other hand, is deduced from the position of the control rod in the experiment by following an empirical rule provided by the LRS. An empirical standard deviation is provided together with this rule. The initial estimates of these independent variables are listed in table 1.
The means of the dependent variables, on the other hand, are initialized with the stationary solution to the PRK equations, using the means listed in table 1 for the independent variables. This does not correspond to the experiment, where the reactor was not in a stationary state at the time of the reactivity insertion. Nevertheless, as the initial state is entirely unknown, it cannot and needs not be specified more accurately. The stationary solution is obtained from eqs. (21, 22) by applying the stationary conditions:

$$\rho = 0 \qquad \frac{d}{dt}\mathbf{x}(t) = \mathbf{0} \tag{50}$$

[22]The k-th diagonal term of the covariance matrix of a multivariate random variable $\mathbf{x}$ is the variance of its k-th coordinate, x_k.

		Initial estimate		
	Unit	Mean (μ)	Standard deviation (σ)	$\left\lvert\frac{\sigma}{\mu}\right\rvert$
ρ	pcm	112	6	5.4 %
β_1	pcm	23.8149	1.211	5.1 %
β_2	pcm	126.101	2.88393	2.3 %
β_3	pcm	122.866	2.75795	2.2 %
β_4	pcm	283.823	4.14416	1.5 %
β_5	pcm	126.319	2.89798	2.3 %
β_6	pcm	52.5235	1.77496	3.4 %
λ_1	s^{-1}	0.013 353 5	$3.066\,38 \times 10^{-6}$	0.0 %
λ_2	s^{-1}	0.032 612 3	$9.880\,13 \times 10^{-6}$	0.0 %
λ_3	s^{-1}	0.121 058	$2.044\,21 \times 10^{-5}$	0.0 %
λ_4	s^{-1}	0.305 665	$1.411\,16 \times 10^{-4}$	0.0 %
λ_5	s^{-1}	0.861 038	$6.096\,98 \times 10^{-4}$	0.0 %
λ_6	s^{-1}	2.892 02	$2.933\,29 \times 10^{-3}$	0.0 %
Λ	s	$4.686\,78 \times 10^{-5}$	$9.682\,23 \times 10^{-8}$	0.0 %

Table 1: Means and standard deviations of the initial estimates of the independent variables.

However, when these conditions are applied, the resulting system of $N_\ell + 1$ equations is linearly dependent[23] and has rank N_ℓ. To find an unique solution, one additional condition must be specified. In conformity with eq. (45), the initial neutron population is set to the value of the first observation (which serves as an estimate of the parameter $\mu^{(0)}$:

$$n(t^{(0)}) = z^{(0)} \approx \mu^{(0)} \tag{51}$$

Finally, as their values at the beginning of the measurement are essentially unknown, the initial standard deviation of each dependent variable is set proportional to its initial mean:

$$\left(\mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(0|0)}\right)_{k,k} = \left(\sigma_{initial} \cdot \hat{x}_k^{(0|0)}\right)^2 \qquad k = 1, \cdots, 7 \tag{52}$$

The factor $\sigma_{initial}$ should be as large as possible, while still ensuring that the sigma points $\boldsymbol{\mathcal{X}}_x^{(0|0)}$ are positive. Indeed, the initial values of the dependent variables are approximated as GRV, which has a positive probability of taking a negative value. As $\sigma_{initial}$ increases, so does the spread of that GRV, and the sigma points grow further apart to account for it. By trial and error, it was determined that some sigma points become negative as soon as $\sigma_{initial}$ exceeds the value of 0.575.

[23] Eq. (21) simply becomes the negative sum of eqs. (22) due to the relation in eq. (24).

5 Results

The measurement described in section 4.1 is filtered using the UKF ($\sigma_{initial} = 0.5$, $\sigma_{process} = 10^{-3}$), to estimate the (scaled-down, cf. eq. 45) neutron population $n(t)$, the reactivity ρ, and the reactor properties $\beta_\ell, \lambda_\ell, \Lambda$. The parameters to be chosen are $\sigma_{process}$ (eq. 43) and $\sigma_{initial}$ (52). For a first execution, $\sigma_{initial}$ is set to 0.5, safely below the limit discussed in section 4.5. The standard deviation $\sigma_{process}$ is set to 10^{-3}. This is equivalent to assuming that the discrete PRK model produces a true relative error in the order of 0.1 % after 0.1 s[24].

5.1 Estimation of the Neutron Population

In figure 8 (top), the observed neutron counts $z^{(i)}$ are plotted against the predicted neutron population $n_{PRK}(t)$, calculated using the PRK model (without Kalman filtering), and the filtered neutron population $n_{filt}(t)$ estimated using the UKF.

It appears in figure 8 (top) that the reactivity insertion was initially overestimated, as the purely model-based estimate $n_{PRK}(t)$ grows slightly faster than the measured neutron count. Although the filtered signal $n_{filt}(t)$ fits the observations much better quantitatively, it is worth noting that its shape does not entirely conform with the PRK theory: In the first 10 seconds, $n_{filt}(t)$ is not monotonically increasing. This is likely an artefact caused by the observed neutron counts and by the uncertain time of the reactivity insertion (compare with the smoothed signal in figure 7).

Figure 8 (center) shows the estimated standard deviation $\sigma_{observation}$ of the observations $z^{(i)}$, the estimated standard deviation σ_{PRK} of the model-based estimates $n_{PRK}(t)$, and the estimated standard deviation of the filtered estimate $n_{filt}(t)$. All three increase exponentially with time due to the exponential increase of the true neutron population $n(t)$ itself. However, the variance of the filtered signal is more than one order of magnitude smaller than that of the observations, while σ_{PRK} is the largest by several orders of magnitude. Note that $\sigma_{observation}$ is not involved in any calculations, but instead serves exclusively for comparison with σ_{filt}. It is calculated as $\sigma^{(i)}_{observation} = \sqrt{z^{(i)}}$ in reference to eq. (44) (cf. eq. 51). $n_{PRK}(t)$ and σ_{PRK} are obtained by applying a simplified version of the UKF where the update step is skipped entirely. In this modified filter, eqs. (53) and (54) are substituted for eqs. (16 b) and (17 b), respectively.

$$\hat{\mathbf{x}}^{(i|i)} = \hat{\mathbf{x}}^{(i|i-1)} \tag{53}$$

$$\mathbf{\Sigma}^{(i|i)}_{\hat{\mathbf{x}}\hat{\mathbf{x}}} = \mathbf{\Sigma}^{(i|i-1)}_{\hat{\mathbf{x}}\hat{\mathbf{x}}} \tag{54}$$

[24] 0.1 s corresponds to the time between subsequent observations.

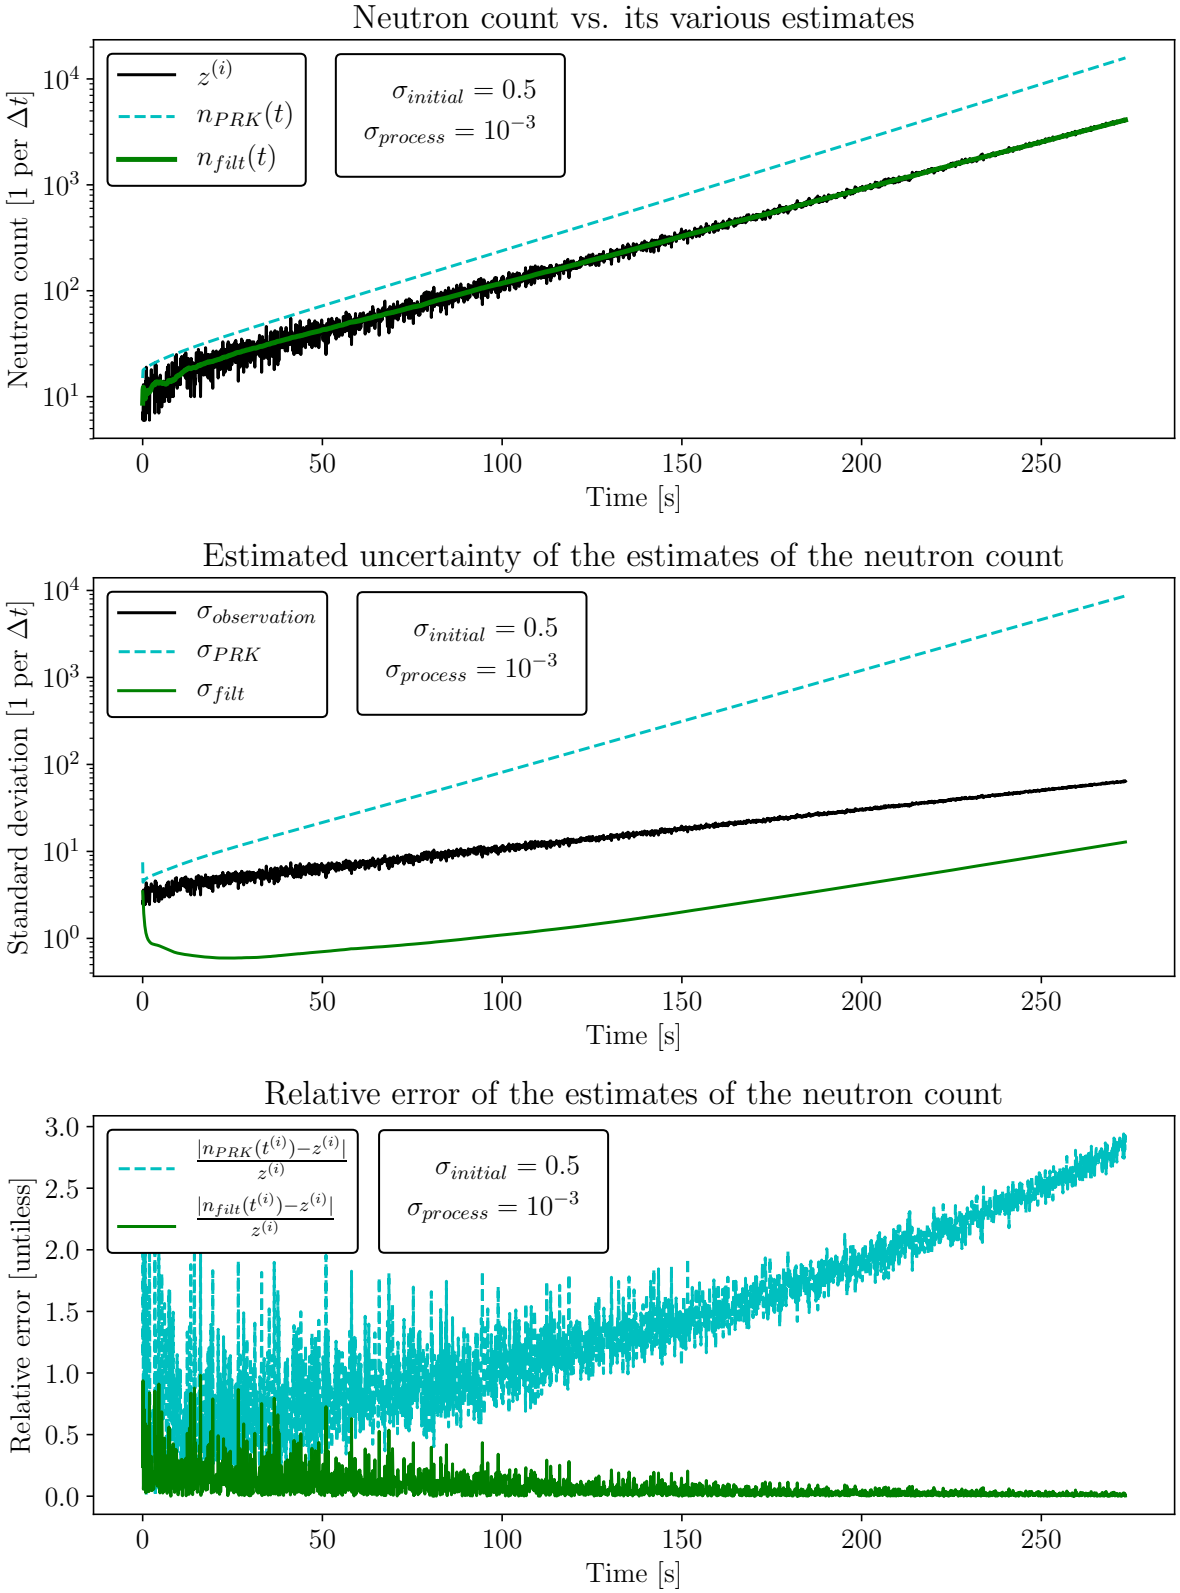

Figure 8: Estimates of the neutron population for $\sigma_{initial} = 0.5$, $\sigma_{process} = 10^{-3}$, their estimated variance, and their relative error with respect to the observations. The time between subsequent observations is 0.1 s.

The relative error of the estimates $n_{PRK}(t)$ and $n_{filt}(t)$ is shown in figure 8 (bottom). It is defined as the absolute value of the relative difference between the estimate and the corresponding observation. Though this definition differs from the true (relative) error in that the observations are themselves different from the true state, it is consistent with the concept of least-squares estimation. Following the exponential increase of the neutron population $n(t)$, the relative errors also follow an exponential progress. The relative error of $n_{filt}(t)$, whose standard deviation is consistently close to one tenth of $\sigma_{observation}$ according to figure 8 (center), decays exponentially. The relative error of $n_{PRK}(t)$, on the contrary, grows exponentially.

5.2 Estimation of the Independent Variables

Figure 8 is obtained by estimating the independent variables and the dependent variables simultaneously. The sought-for result is a more accurate estimate of the independent variables, which can be used to predict power transients more accurately. In this section, the estimation of the independent variables is presented and the possibility of overfitting is discussed.

5.2.1 Estimation with the UKF

The final estimates of the independent variables for $\sigma_{initial} = 0.5$ and $\sigma_{process} = 10^{-3}$ are listed in table 2. The evolution of the estimates, along with their estimated variances, is further shown in figure 9. Each plot shows the initial mean (straight black line) and the initial 68% confidence interval[25] (dashed black line) of one variable, along with the evolution of its estimate (green line). The shaded region indicates the estimated 68% confidence interval[25] of the estimate, based on the covariance matrix $\mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{\mathbf{x}}}^{(i|i)}$. Note that there is no shaded region for β_{filt}: This is because β_{filt} is not directly estimated. The plotted values are calculated from eq. (24). Finally, the numerical values of the estimates at the final time step are shown in the top left plot.

It is immediately apparent in figure 9 that some estimates are more accurate after the simulation than initially, while others are entirely unchanged. By looking at table 2, we confirm that the estimate of the reactivity ρ is substantially improved (its standard deviation is divided by four), while the final estimates of the decay constants λ_ℓ and of the mean neutron generation time Λ equal their initial estimates almost exactly.
Generally speaking, three factors influence the extent to which a specific state variable x_j is updated in a given step i of the UKF (eqs. 16 b, 38, 37): How much the observation $\mathbf{z}^{(i)}$ differs from the expected observation $\hat{\mathbf{z}}^{(i)}$, how large the estimated variance of $x_j^{(i|i-1)}$ was before the update, and how strongly x_j really impacts the observed quantities $\mathbf{z}$. In the present case, the observed quantity is the one-dimensional neutron count $z^{(i)}$, which makes the Kalman gain $\mathbf{K}^{(i)}$ a column vector. Following eq. (16 b), the difference $\hat{z}^{(i)} - z^{(i)}$ does not affect how much the state variables are updated relative to each other, but rather scales the entire update. How much each state variable is updated in relation to the others is determined by the Kalman gain

[25] When a Gaussian random variable with mean μ and variance σ^2 is realized, 68.27 % is the probability that it falls at most σ away from μ (σ is the standard deviation): $X \sim \mathcal{N}\left(\mu, \sigma^2\right) \implies P\left(|X - \mu| \leq \sigma\right) = 0.6827$

		Final estimate		Relative difference		
	[Unit]	Mean (μ)	Standard deviation (σ)	$\frac{\Delta\mu}{\mu_{initial}}$	$\frac{\Delta\sigma}{\sigma_{initial}}$	$\left\lvert\frac{\Delta\mu}{\sigma_{initial}}\right\rvert$
ρ	pcm	102.7	1.4	-8.3 %	-75.6 %	157.5 %
β_1	pcm	24.1063	1.202	+1.2 %	-0.8 %	24.1 %
β_2	pcm	127.183	2.83239	+0.9 %	-1.8 %	37.5 %
β_3	pcm	123.337	2.75030	+0.4 %	-0.3 %	17.1 %
β_4	pcm	284.465	4.13787	+0.2 %	-0.2 %	15.5 %
β_5	pcm	126.446	2.89707	+0.1 %	-0.03 %	4.4 %
β_6	pcm	52.5493	1.77486	+0.05 %	-0.0 %	1.5 %
λ_1	s^{-1}	0.013 353 5	$3.066\,37 \times 10^{-6}$	+0.0 %	-0.0 %	0.01 %
λ_2	s^{-1}	0.032 612 3	$9.880\,10 \times 10^{-6}$	+0.0 %	-0.0 %	0.4 %
λ_3	s^{-1}	0.121 058	$2.044\,21 \times 10^{-5}$	+0.0 %	-0.0 %	0.4 %
λ_4	s^{-1}	0.305 665	$1.411\,16 \times 10^{-4}$	+0.0 %	-0.0 %	0.6 %
λ_5	s^{-1}	0.861 038	$6.096\,97 \times 10^{-4}$	-0.0 %	-0.0 %	0.02 %
λ_6	s^{-1}	2.892 02	$2.933\,29 \times 10^{-3}$	-0.0 %	-0.0 %	0.09 %
Λ	s	$4.686\,53 \times 10^{-5}$	$9.681\,68 \times 10^{-8}$	-0.0 %	-0.0 %	2.5 %

Table 2: Means and standard deviations of the final estimates of the independent variables, after filtering the observations with the UKF. The difference between the final values and the initial values, relative to the initial values, is listed in the "Relative difference" column.

$\mathbf{K}^{(i)}$, which is a scalar multiple of the covariance $\mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{z}}^{(i|i-1)}$ since the covariance $(\sigma_{\hat{z}\hat{z}}^2)^{(i)}$ is scalar (eq. 38). Finally, by eq. (37), each coordinate of $\mathbf{\Sigma}_{\hat{\mathbf{x}}\hat{z}}^{(i|i-1)}$ depends on how much this coordinate varies in the k different sigma points $\left(\boldsymbol{\mathcal{X}}_x^{(i|i-1)}\right)_k$, as well as how far the resulting $\mathcal{Z}_k^{(i)}$ strays from $\hat{z}^{(i)}$ when $\left(\boldsymbol{\mathcal{X}}_x^{(i|i-1)}\right)_k$ is indeed different from $\hat{\mathbf{x}}^{(i|i-1)}$. The former (i.e., how far the sigma points are spread along the j-th dimension) is a direct consequence of the estimated variance of the corresponding coordinate $x_j^{(i|i-1)}$ of the estimate of the state prior to the current step. The latter, on the other hand, is a combined result of that same estimated variance and of how strongly the state variable x_j really does impact the neutron population $n(t^{(i)})$.

This leaves two plausible reasons for the major difference between the estimated independent variables: A difference in the standard deviations of the initial estimates, or a difference in each variable's impact on the neutron population. Clearly, the former is the case for the decay constants λ_ℓ and the mean neutron generation time Λ: Table 1 shows that the standard deviation of each of these variables' initial estimate was three or more orders of magnitude smaller than the estimate itself, whereas the same ratio is about 5 % for ρ or β_1.

However, the estimates of ρ and β_1 are not updated to the same degree either, despite their similar relative uncertainties. To a slightly lesser extent, the same can be said when comparing ρ to any fractional yield β_ℓ. This suggests that the fractional yields impact the neutron population less than the reactivity. To investigate whether this is the case, the result of a 20 % change in each of these independent variables is plotted in figure 10. Figure 10 was obtained like the figures in section 3.1.2, using the PRK equations with the initial estimates of all independent

variables and without Kalman filtering.

Figure 10 confirms that, after upwards of 100 s, a change in reactivity has a greater effect on the neutron population than the same relative change in any single fractional yield. Furthermore, the figure explains the differences between the fractional yields of the various precursor groups. Note that the order of decreasing $\frac{\Delta\sigma}{\sigma_{initial}}$ and $\left|\frac{\Delta\mu}{\sigma_{initial}}\right|$ is: group 2 > group 1 > groups 3 & 4 > group 5 > group 6 (cf. table 2). When the precursor groups are ordered by decreasing impact on the neutron population, group 1 comes slightly after, rather than slightly before, groups 3 & 4. This discrepancy is readily explained by the particularly high initial $\left|\frac{\sigma}{\mu}\right|$ ratio of group 1 (cf. table 1). Hence, jointly, figure 10 and table 1 explain the different relative magnitude of the update for the individual precursor groups and the reactivity. The estimated uncertainty of the final estimate for the reactivity remains surprisingly low, particularly when its relative change is compared to that of the estimated uncertainties of the fractional yields.

Independent variable estimation with UKF ($\sigma_{initial} = 0.5$, $\sigma_{process} = 10^{-3}$)

Final estimates of the indep. variables

All values	$\beta_1 = 24.106$	$\lambda_1 = 1335.3\,\mathrm{s}^{-1}$
$\times\mathbf{10^{-5}}$	$\beta_2 = 127.18$	$\lambda_2 = 3261.2\,\mathrm{s}^{-1}$
	$\beta_3 = 123.33$	$\lambda_3 = 12105\,\mathrm{s}^{-1}$
$\rho = 102.73$	$\beta_4 = 284.46$	$\lambda_4 = 30566\,\mathrm{s}^{-1}$
$\beta = 738.08$	$\beta_5 = 126.44$	$\lambda_5 = 86103\,\mathrm{s}^{-1}$
$\Lambda = 4.6865\,\mathrm{s}$	$\beta_6 = 52.549$	$\lambda_6 = 289202\,\mathrm{s}^{-1}$

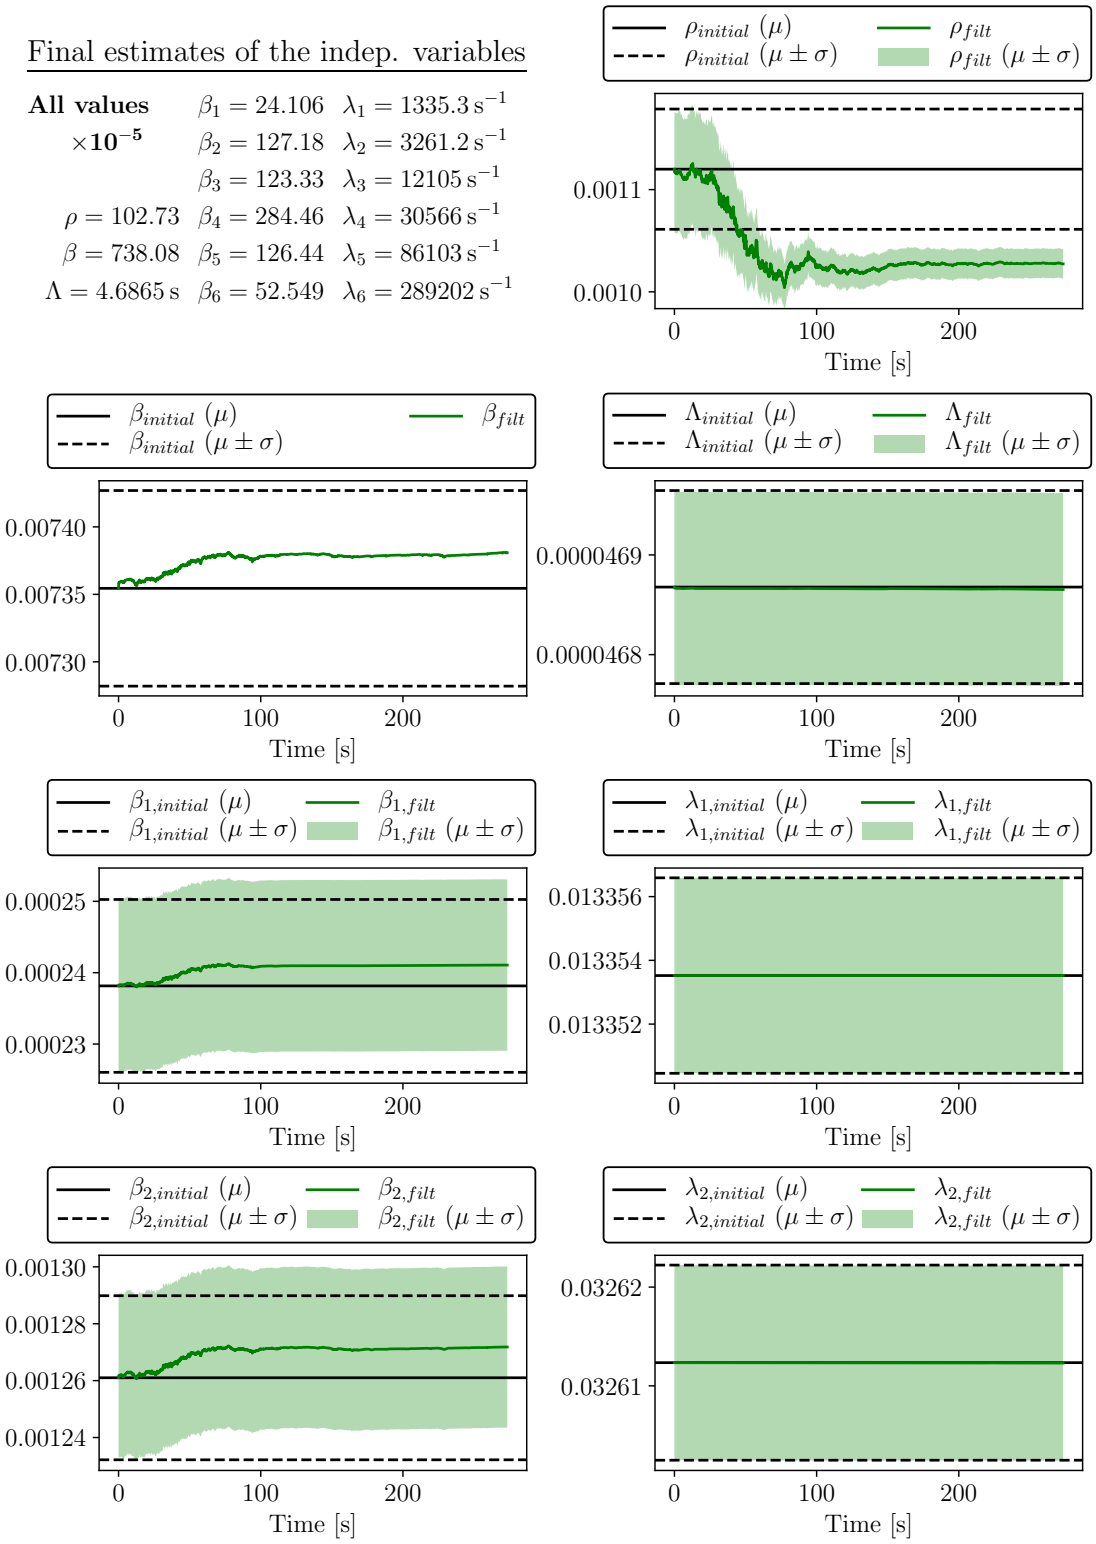

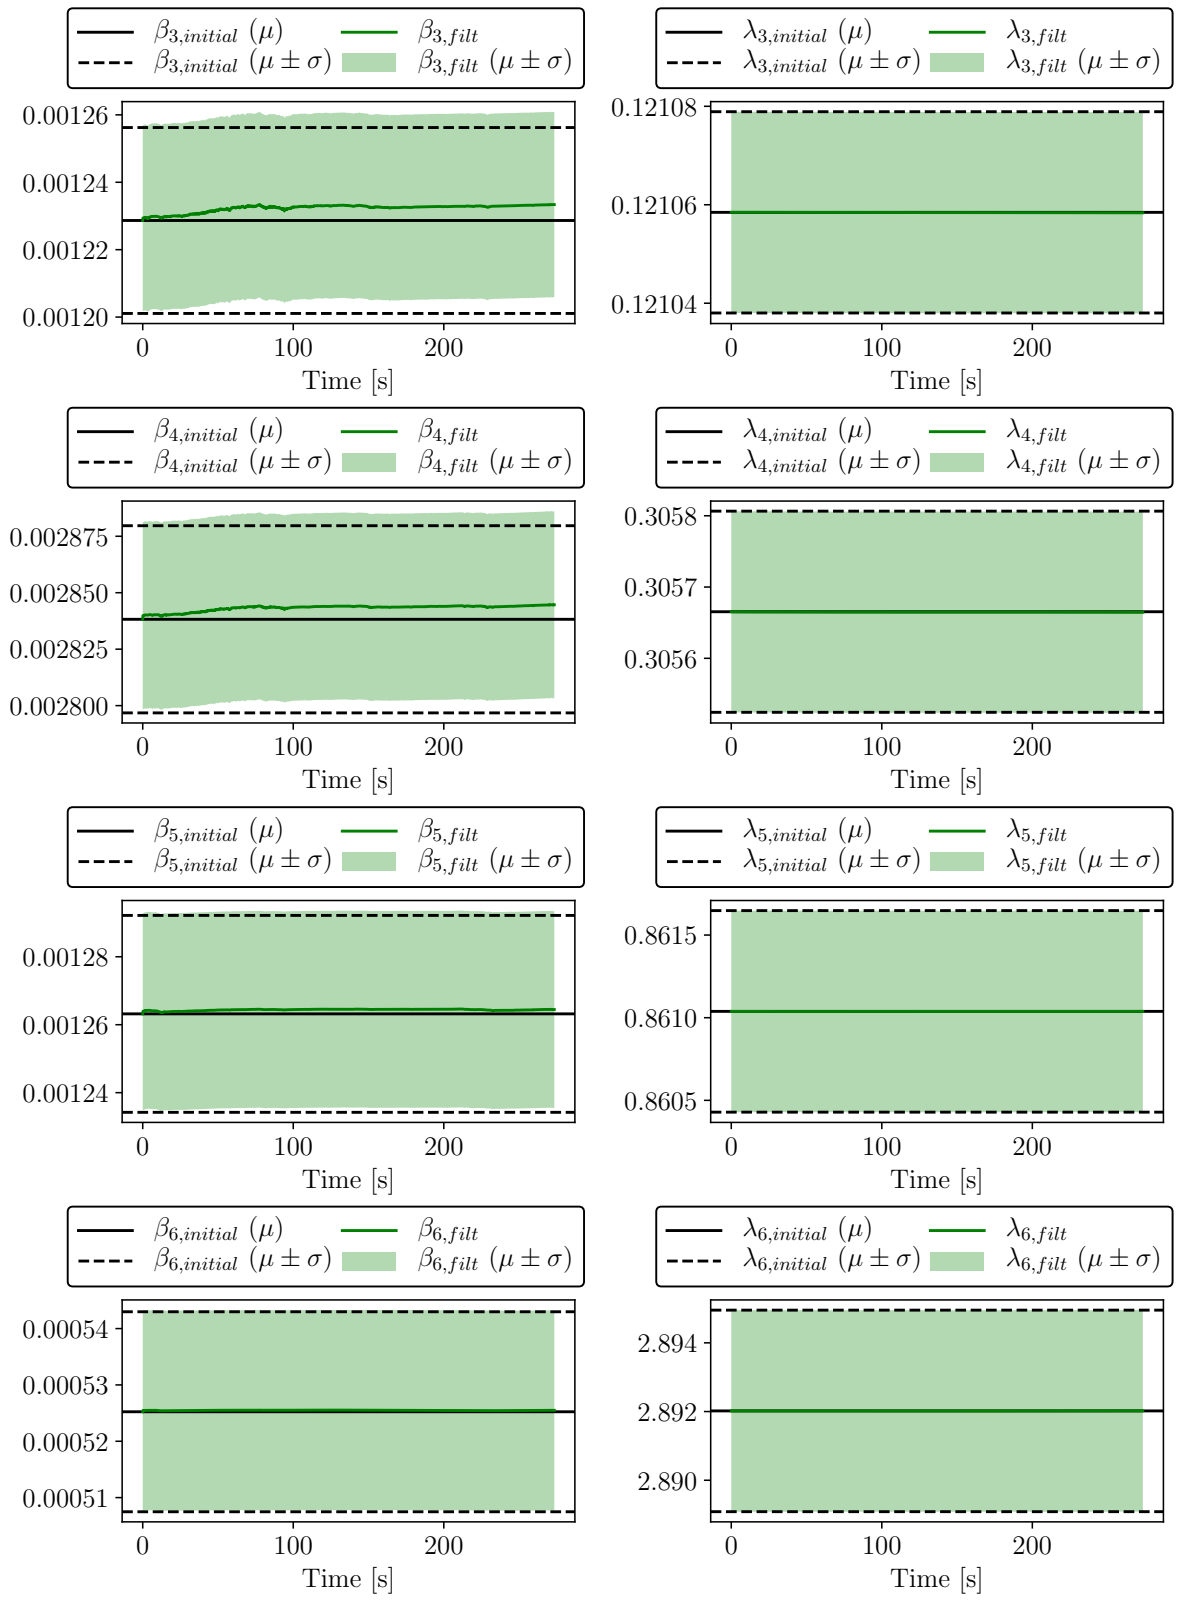

Figure 9: Estimates of the independent variables for $\sigma_{initial} = 0.5$, $\sigma_{process} = 10^{-3}$, and their estimated 68% confidence intervals. The time between subsequent observations is 0.1 s.

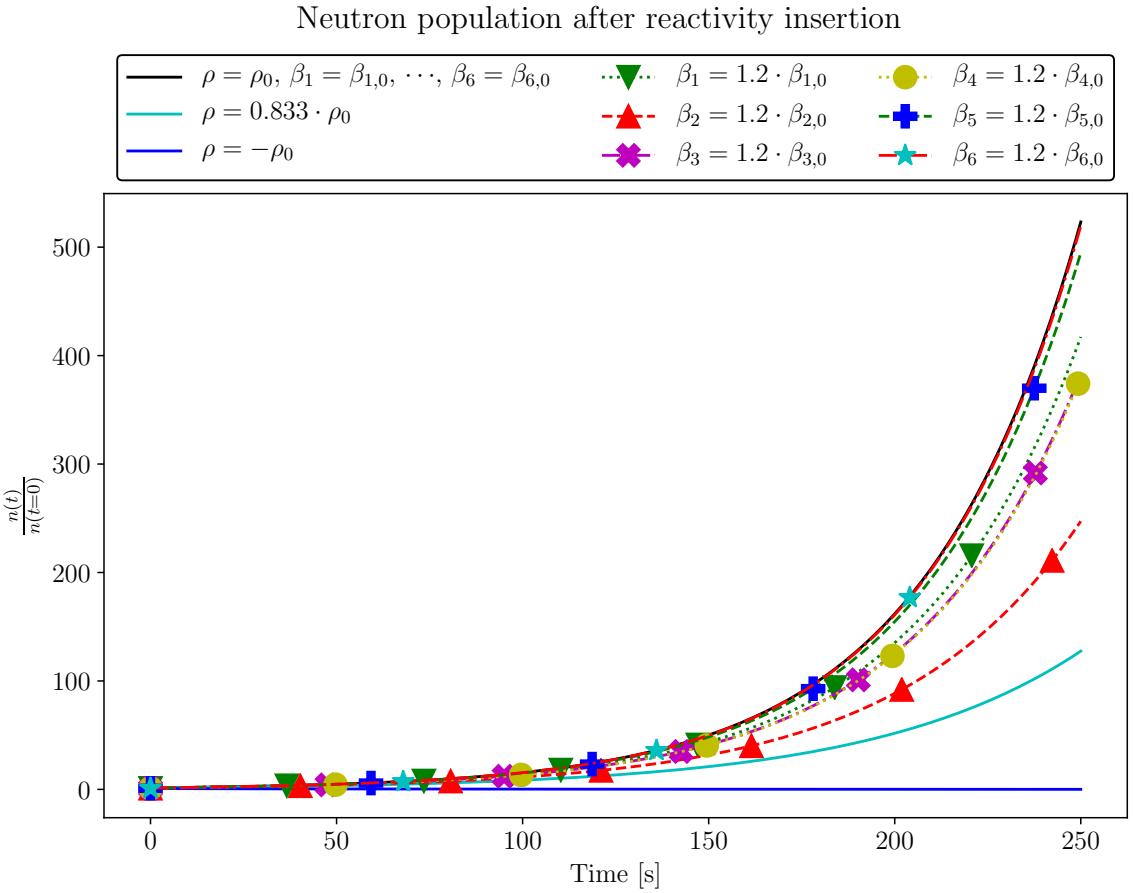

Figure 10: Effect of different independent variables on the neutron population.

5.2.2 Discussion of Overfitting

When 15 independent variables are estimated based on 1-dimensional observations, overfitting may sometimes occur. Figure 11 shows that this possibility is real in the case of the PRK model with the present experimental set-up. Indeed, when considering the entire duration of the positive reactivity insertion (stretching over 250 s), a given signal may be equally well explained by contradictory combinations of independent variables. This is showcased by the two lines with triangle (respectively cross) markers. One is obtained by *increasing* β_2, while the other is obtained by *de*creasing it and simultaneously decreasing ρ to overcompensate the decrease in β_2.

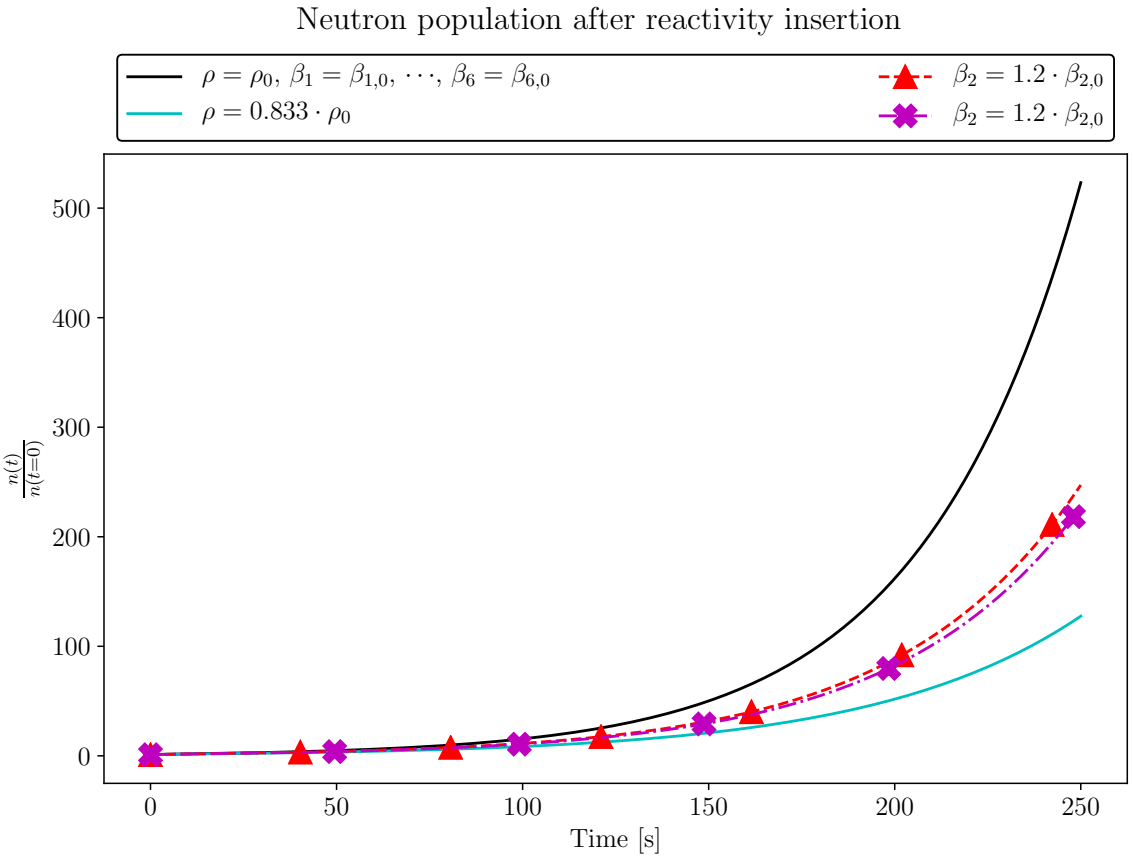

Figure 11: Possibility of overfitting the PRK model.

In cases where accurate prior estimates are unavailable for more than one independent variable, the estimation should be performed multiple times using different measurements to inspect the consistency of the results. In the present case, fortunately, the presence of accurate prior estimates for all independent variables makes gross overfitting easy to detect (see figure 13 for an example). Moreover, such overfitting is prevented by the Kalman filter thanks to the existence of accurate prior estimates. Indeed, we observe in figure 9 and table 2 that all independent variables are updated in a consistent way: The fractional yields are estimated higher, while the reactivity is estimated lower. Each of these changes by itself results in a slower rise of the

neutron population. All other independent variables are virtually unchanged.

5.3 Influence of the Parameters $\sigma_{initial}$ and $\sigma_{process}$

In the model presented in section 4, the choice of $\sigma_{initial}$ and $\sigma_{process}$ is free (partly free for the former). The influence of these freely chosen parameters is tested in this section.

5.3.1 $\sigma_{initial}$: Uncertainty of the Dependent Variables' Initial Estimates

The filter is tested with $\sigma_{initial} = 10$, $\sigma_{initial} = 3$ and $\sigma_{initial} = 0.05$, while $\sigma_{process} = 10^{-3}$ is held constant.

$\sigma_{initial} = 10$: As stated in section 4.5, for a value of $\sigma_{initial} > 0.575$, some sigma points of the first prediction step become negative, which results in nonphysical model behavior in eq. (31). In fact, for $\sigma_{initial} = 10$, this translates into extreme fluctuations as seen in figures 12 and 13. With $\sigma_{initial} = 10$, the model crashes because one of the estimated covariance matrices is not positive definite. Figures 12 and 13 were obtained by manually replacing the negative eigenvalue of the affected covariance matrix with a small positive number. This problem was not encountered in the production of any of the other figures.

$\sigma_{initial} = 3$: This value is still above the safe threshold of 0.575. Nevertheless, the performance of the filter is overall the same as with $\sigma_{initial} = 0.5$. Both the neutron population in figure 14 and the independent variables in figure 15 exhibit mild fluctuations in the first 20 s, respectively the first 100 s. However, the means and estimated standard deviations for the estimates of all independent variables are nearly the same for both parameter settings. For example, the estimated standard deviation of the final estimate for the reactivity merely rises from 1.43 pcm to 1.45 pcm.

$\sigma_{initial} = 0.05$: A particularly low value of $\sigma_{initial}$ means that the initial values of the dependent variables are fairly well-known, which is precisely not the case. As seen in figures 16 and 17, this discrepancy leads to a rapid and substantial adaptation of the independent variables. The true (and underestimated) initial error of the dependent variables is essentially spread evenly across all variables.

These findings show that the value of $\sigma_{initial}$ is crucial to the success of the proposed filtering. If it is specified too far above the empirical limit of 0.575, then spurious effects take over entirely. On the other hand, it must adequately reflect the amount of trust accorded to the initial estimates of the dependent variables.

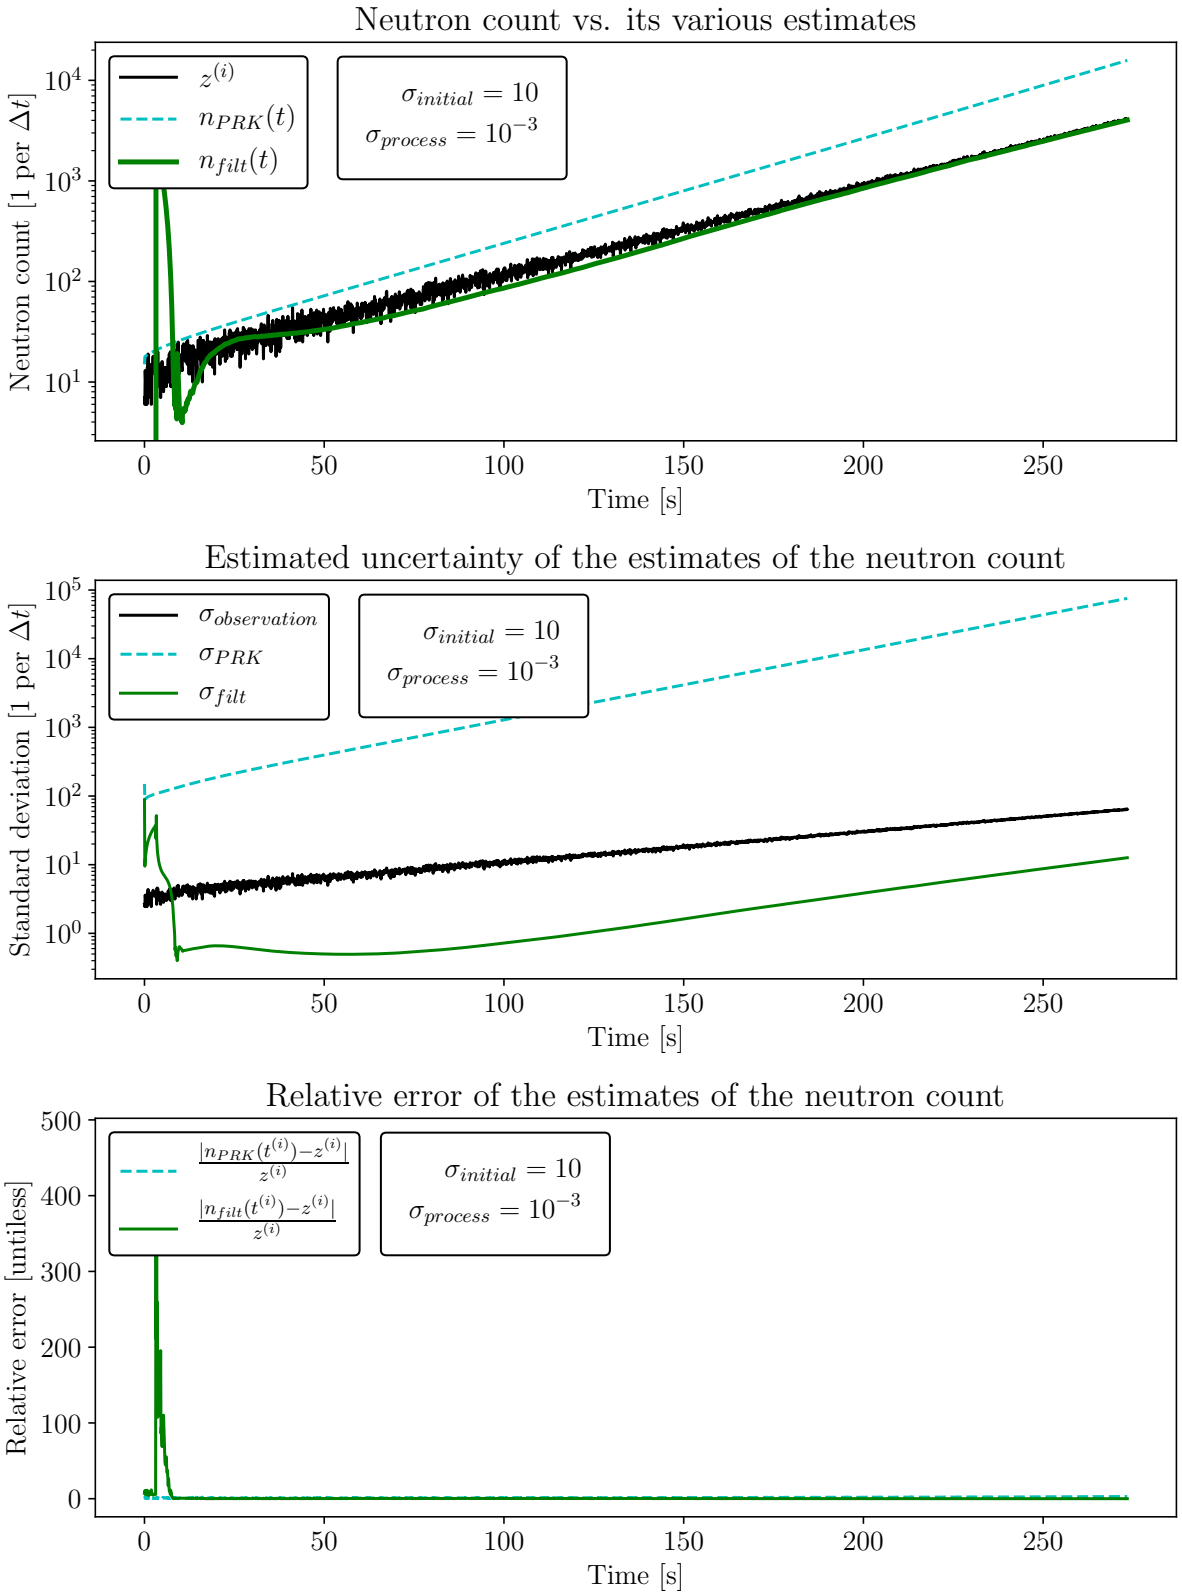

Figure 12: Estimates of the neutron population for $\sigma_{initial} = 10$ and $\sigma_{process} = 10^{-3}$, their estimated variance, and their relative error with respect to the observations. The time between subsequent observations is 0.1 s. This figure was obtained by "fixing" an estimated covariance matrix that was not positive definite.

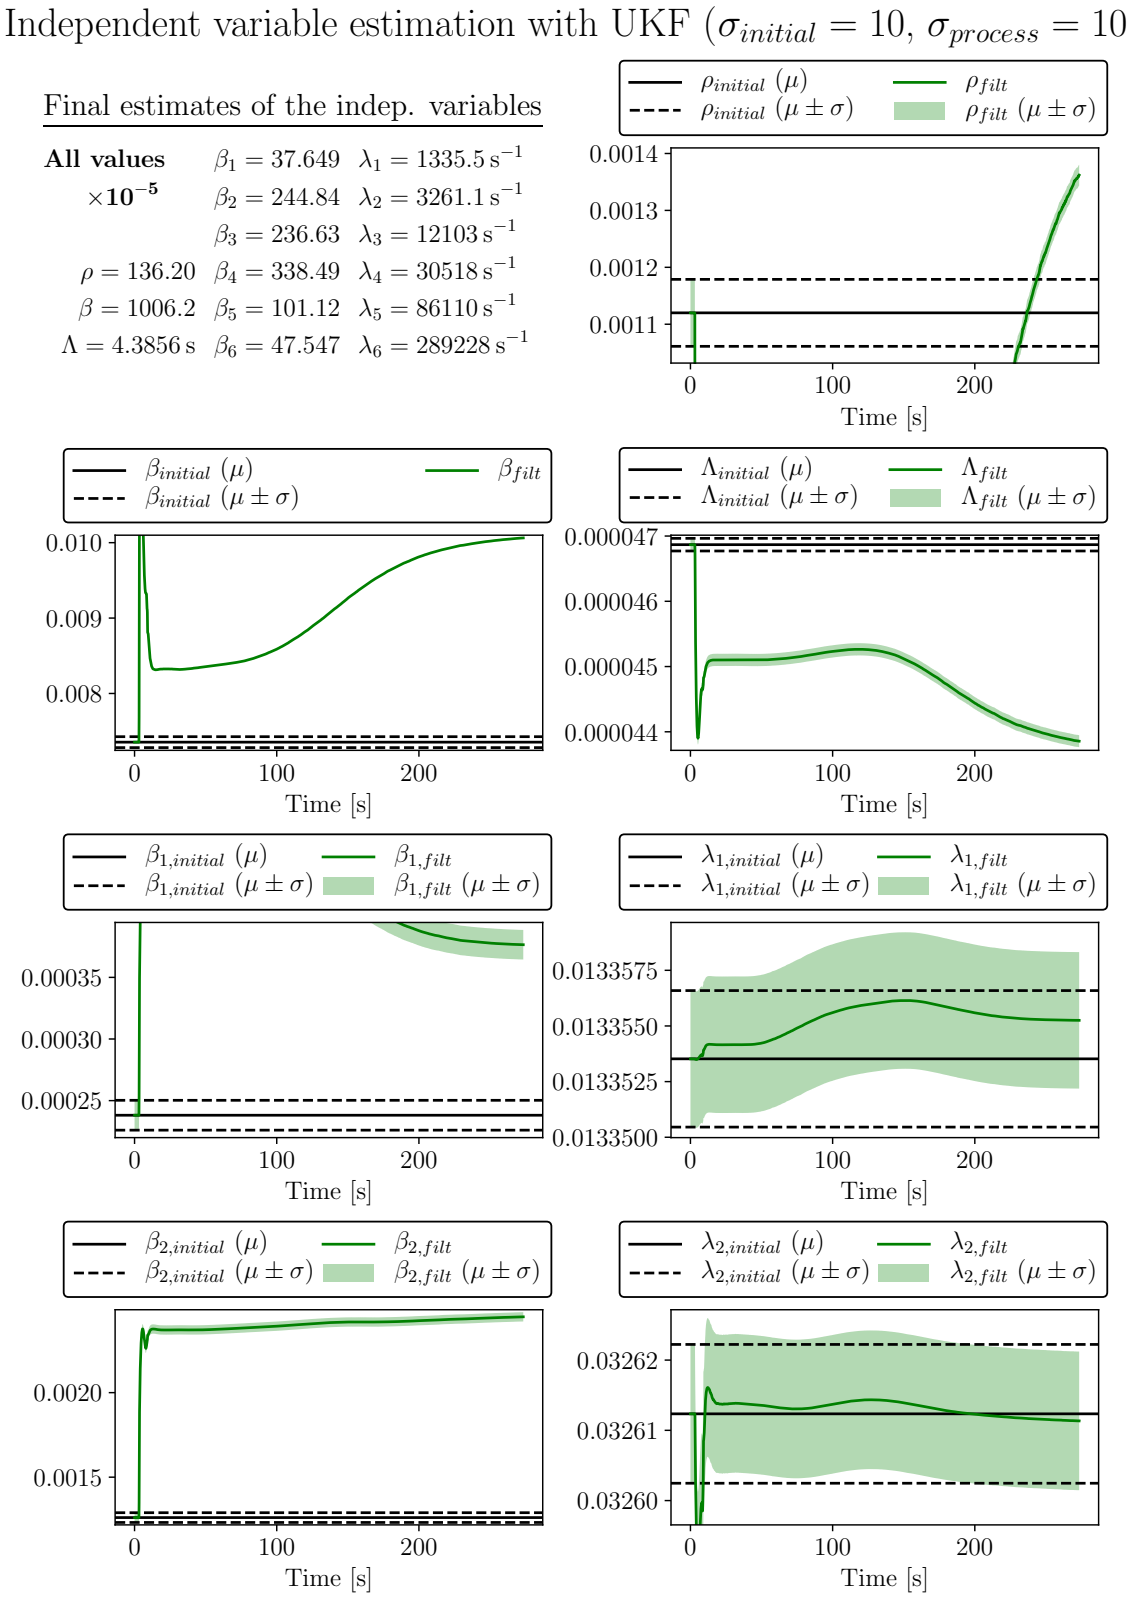

Figure 13: Estimates of the most relevant independent variables for $\sigma_{initial} = 10$ and $\sigma_{process} = 10^{-3}$, and their estimated 68% confidence intervals. The time between subsequent observations is 0.1 s. This figure was obtained by "fixing" an estimated covariance matrix that was not positive definite.

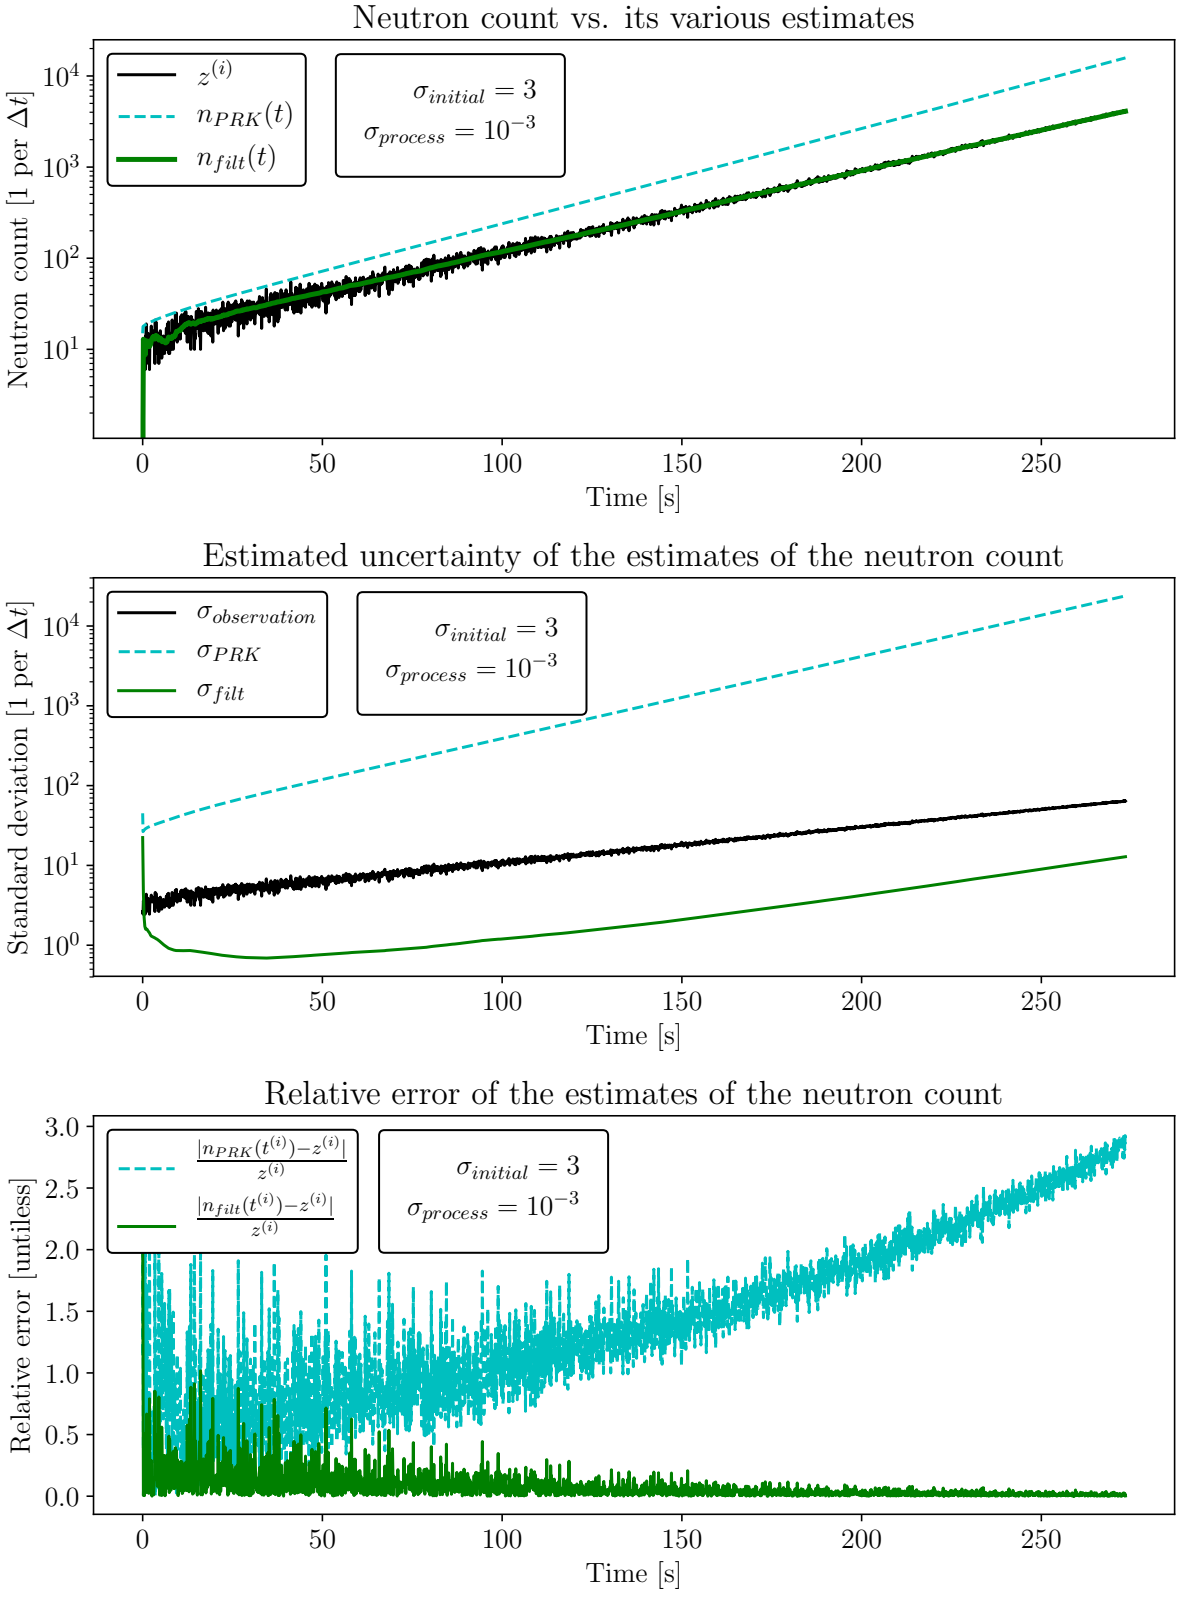

Figure 14: Estimates of the neutron population for $\sigma_{initial} = 3$ and $\sigma_{process} = 10^{-3}$, their estimated variance, and their relative error with respect to the observations. The time between subsequent observations is 0.1 s.

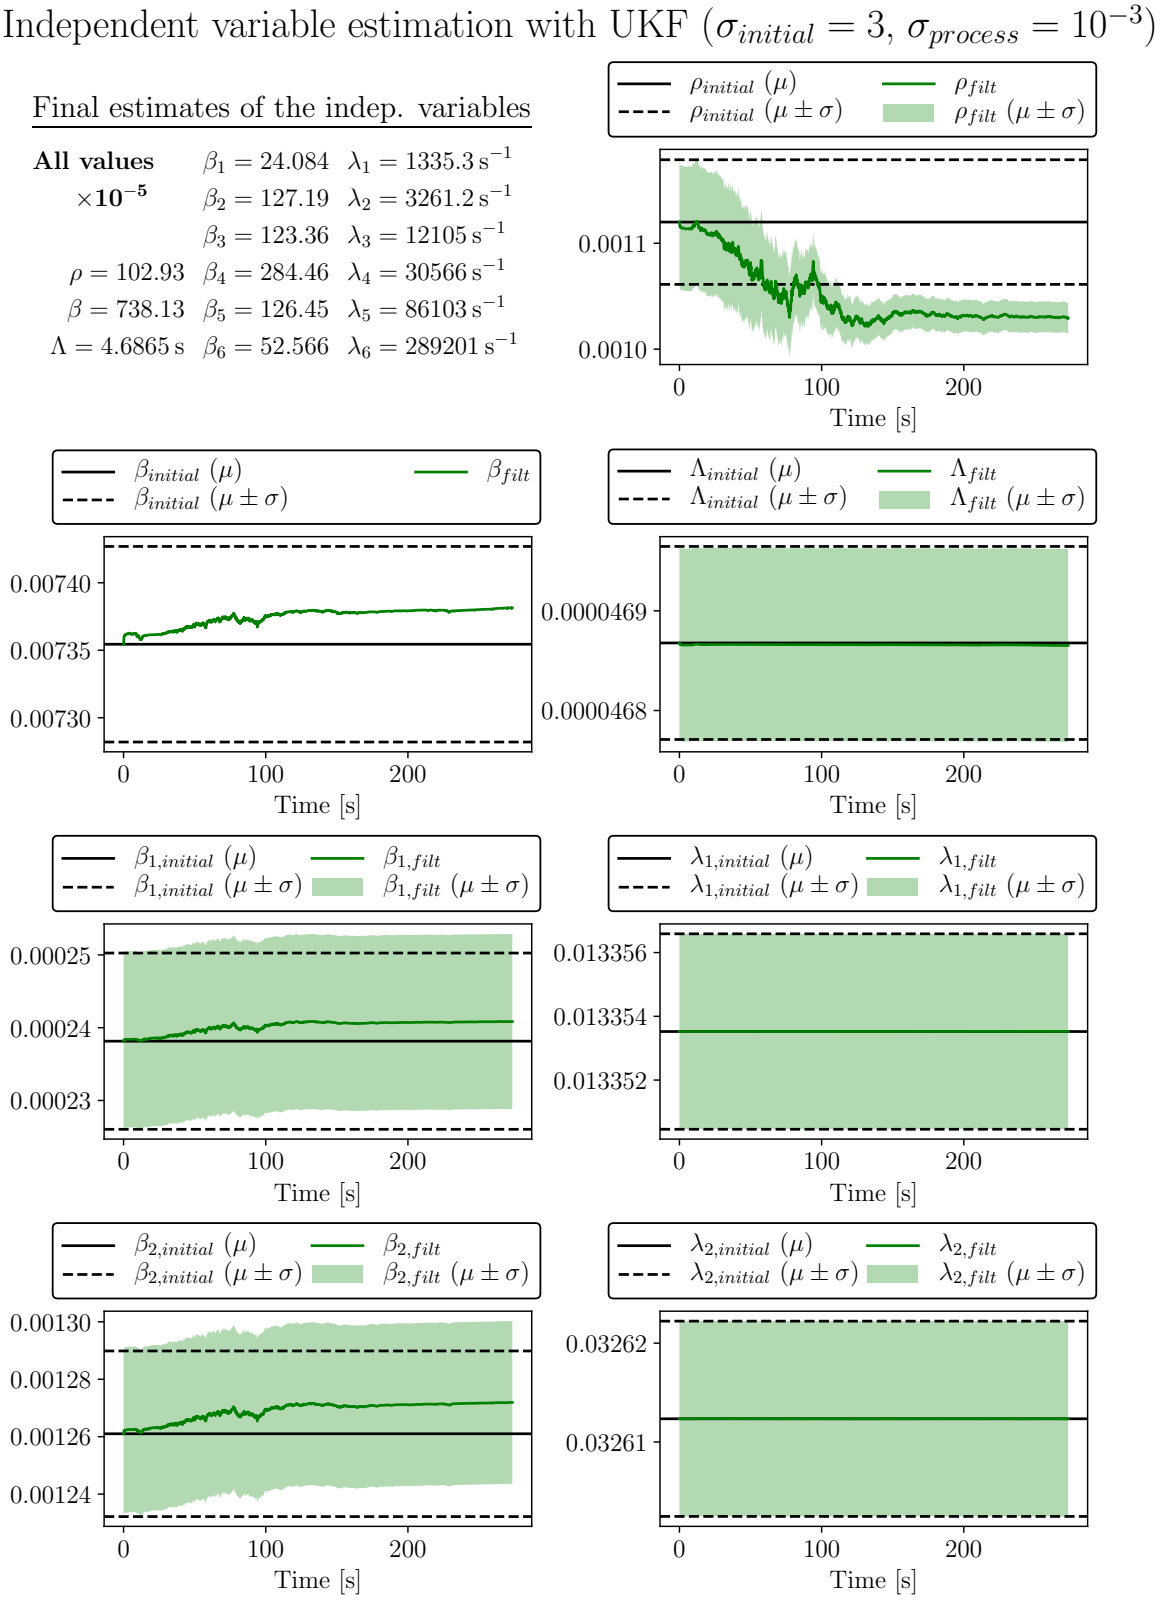

Figure 15: Estimates of the most relevant independent variables for $\sigma_{initial} = 3$ and $\sigma_{process} = 10^{-3}$, and their estimated 68% confidence intervals. The time between subsequent observations is 0.1 s.

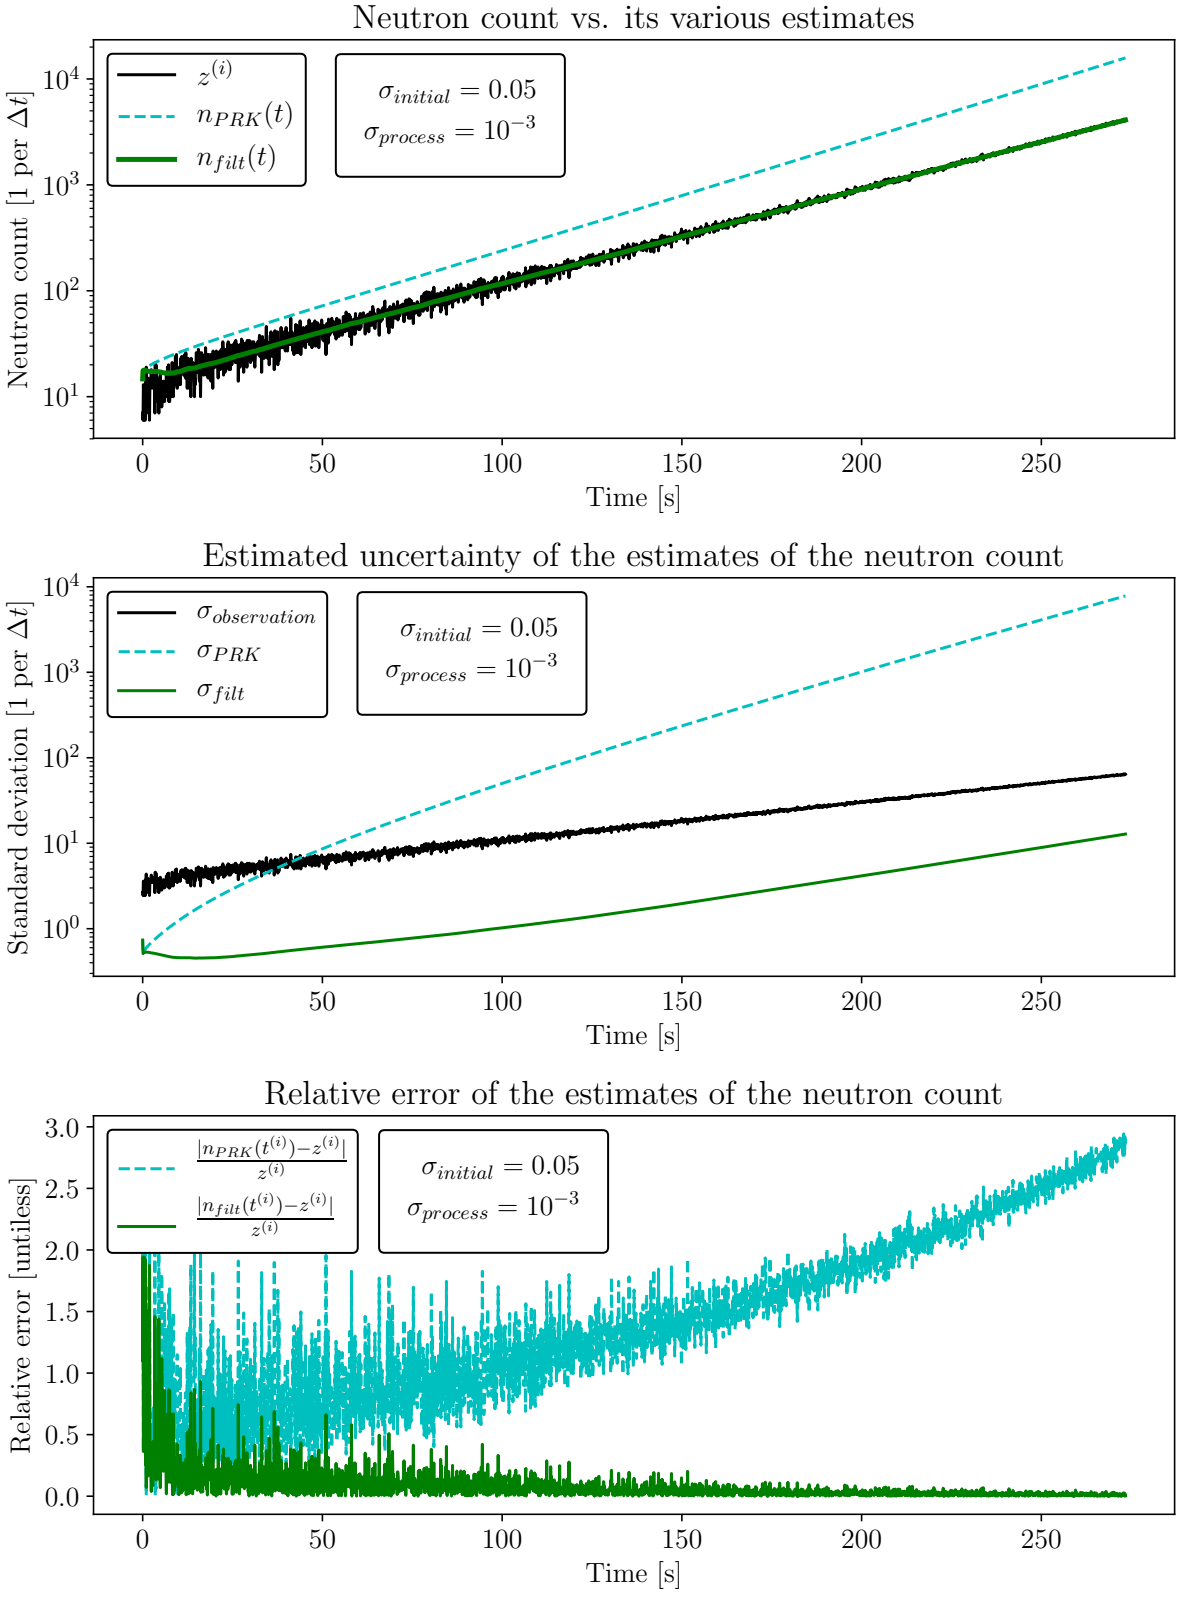

Figure 16: Estimates of the neutron population for $\sigma_{initial} = 0.05$ and $\sigma_{process} = 10^{-3}$, their estimated variance, and their relative error with respect to the observations. The time between subsequent observations is 0.1 s.

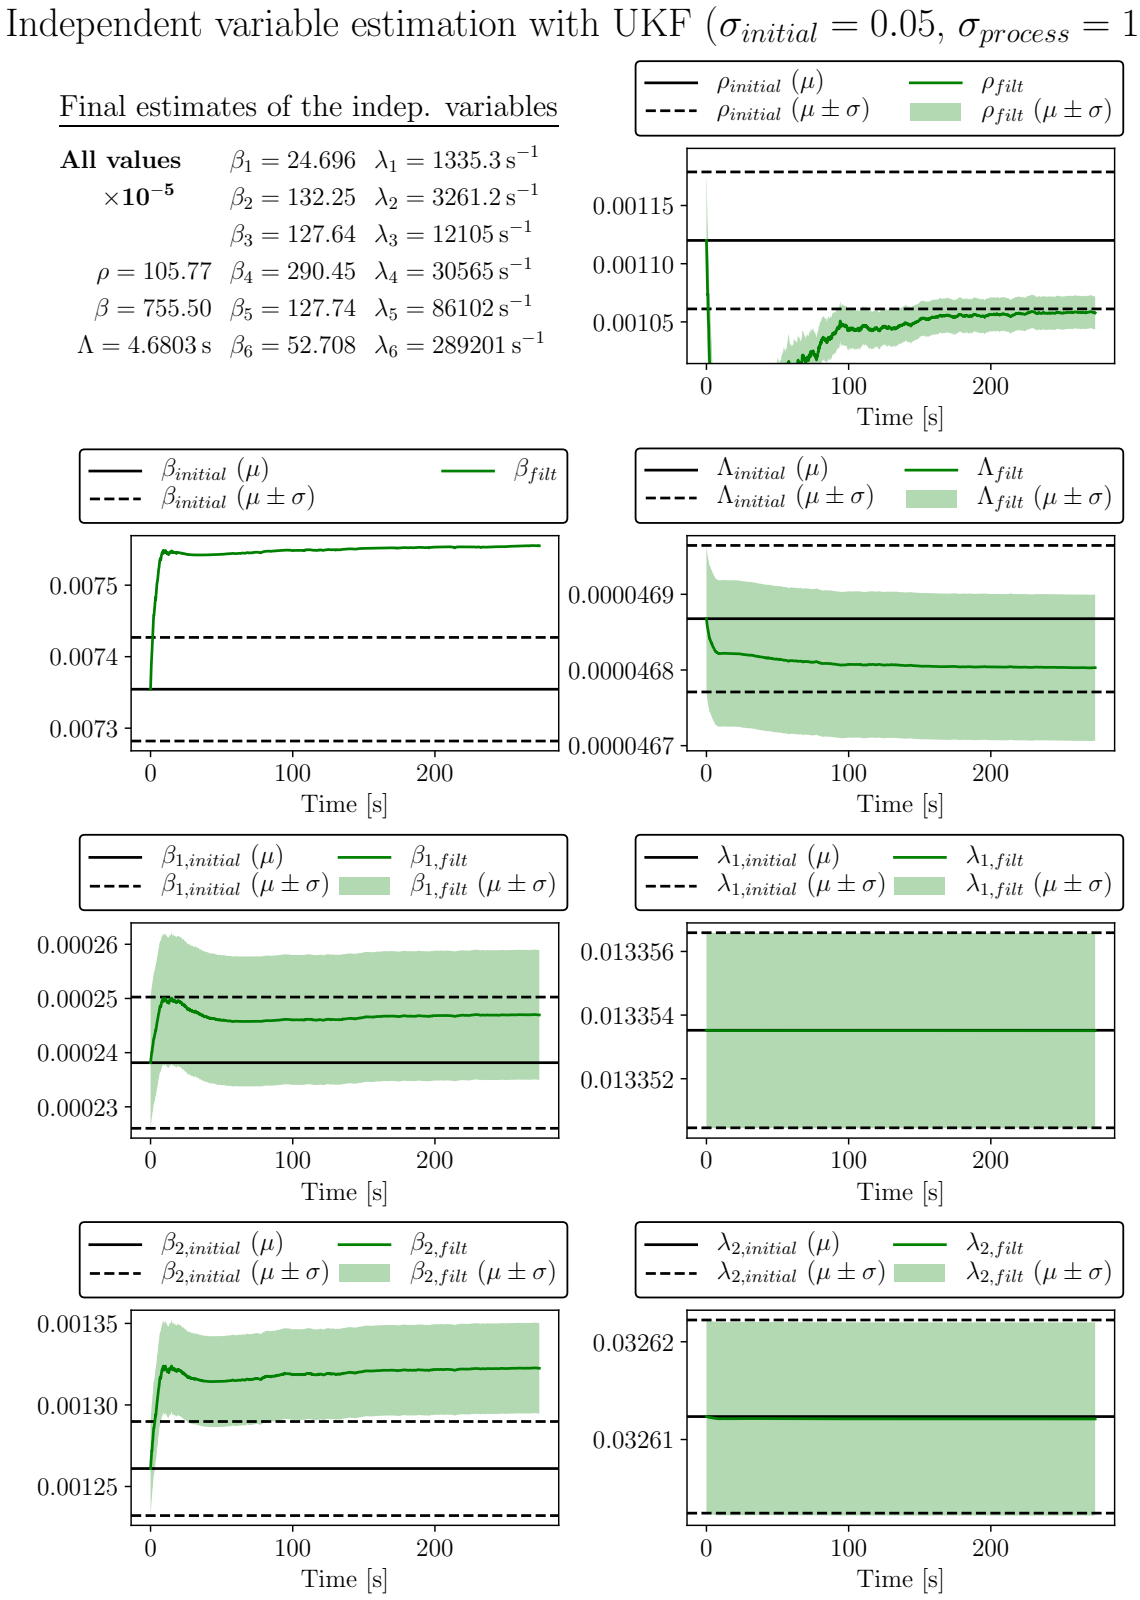

Figure 17: Estimates of the most relevant independent variables for $\sigma_{initial} = 0.05$ and $\sigma_{process} = 10^{-3}$, and their estimated 68% confidence intervals. The time between subsequent observations is 0.1 s.

5.3.2 $\sigma_{process}$: Additional Uncertainty after each Step

The filter is tested with $\sigma_{process} = 10^{-5}$ and $\sigma_{process} = 0.05$, while $\sigma_{initial} = 0.5$ is held constant.

$\sigma_{process} = 10^{-5}$: As shown in figure 18 changing $\sigma_{process}$ from 10^{-3} to 10^{-5} has no discernible effect on the estimate neutron population, however it decreases the estimated uncertainty of the estimate itself. For the most part, this is simply the belief that the model is particularly accurate, producing a relative error in the order of only 10^{-5} in the time between two consecutive observations (cf. section 4.3).

This same belief slightly reduces the estimated uncertainty of the estimated independent variables. This effect is not marked enough to be perceptible in figure 19. For instance, the estimated uncertainty of the reactivity drops from 1.43 pcm to 1.38 pcm.

$\sigma_{process} = 0.05$: Figure 20 shows the estimate of the neutron population with this setting. It fluctuates notably for the entire duration of the experiment, and in fact the estimated uncertainty is nearly equal to that of the observations themselves. By setting $\sigma_{process}$ so high, we are allowing the model to explain the noise in the measurements. This results in an insufficient smoothing of the observed signal and a failure to produce an improved estimate.

The estimation of part of the independent variables with $\sigma_{initial} = 0.5$ and $\sigma_{process} = 0.05$ is shown in figure 21. The means are nearly identical to those obtained with $\sigma_{process} = 10^{-3}$. Only the estimated standard deviation of the reactivity is greater in figure 21. Because more of the signal is explained by process noise, its influence on the estimation of the independent variables is reduced.

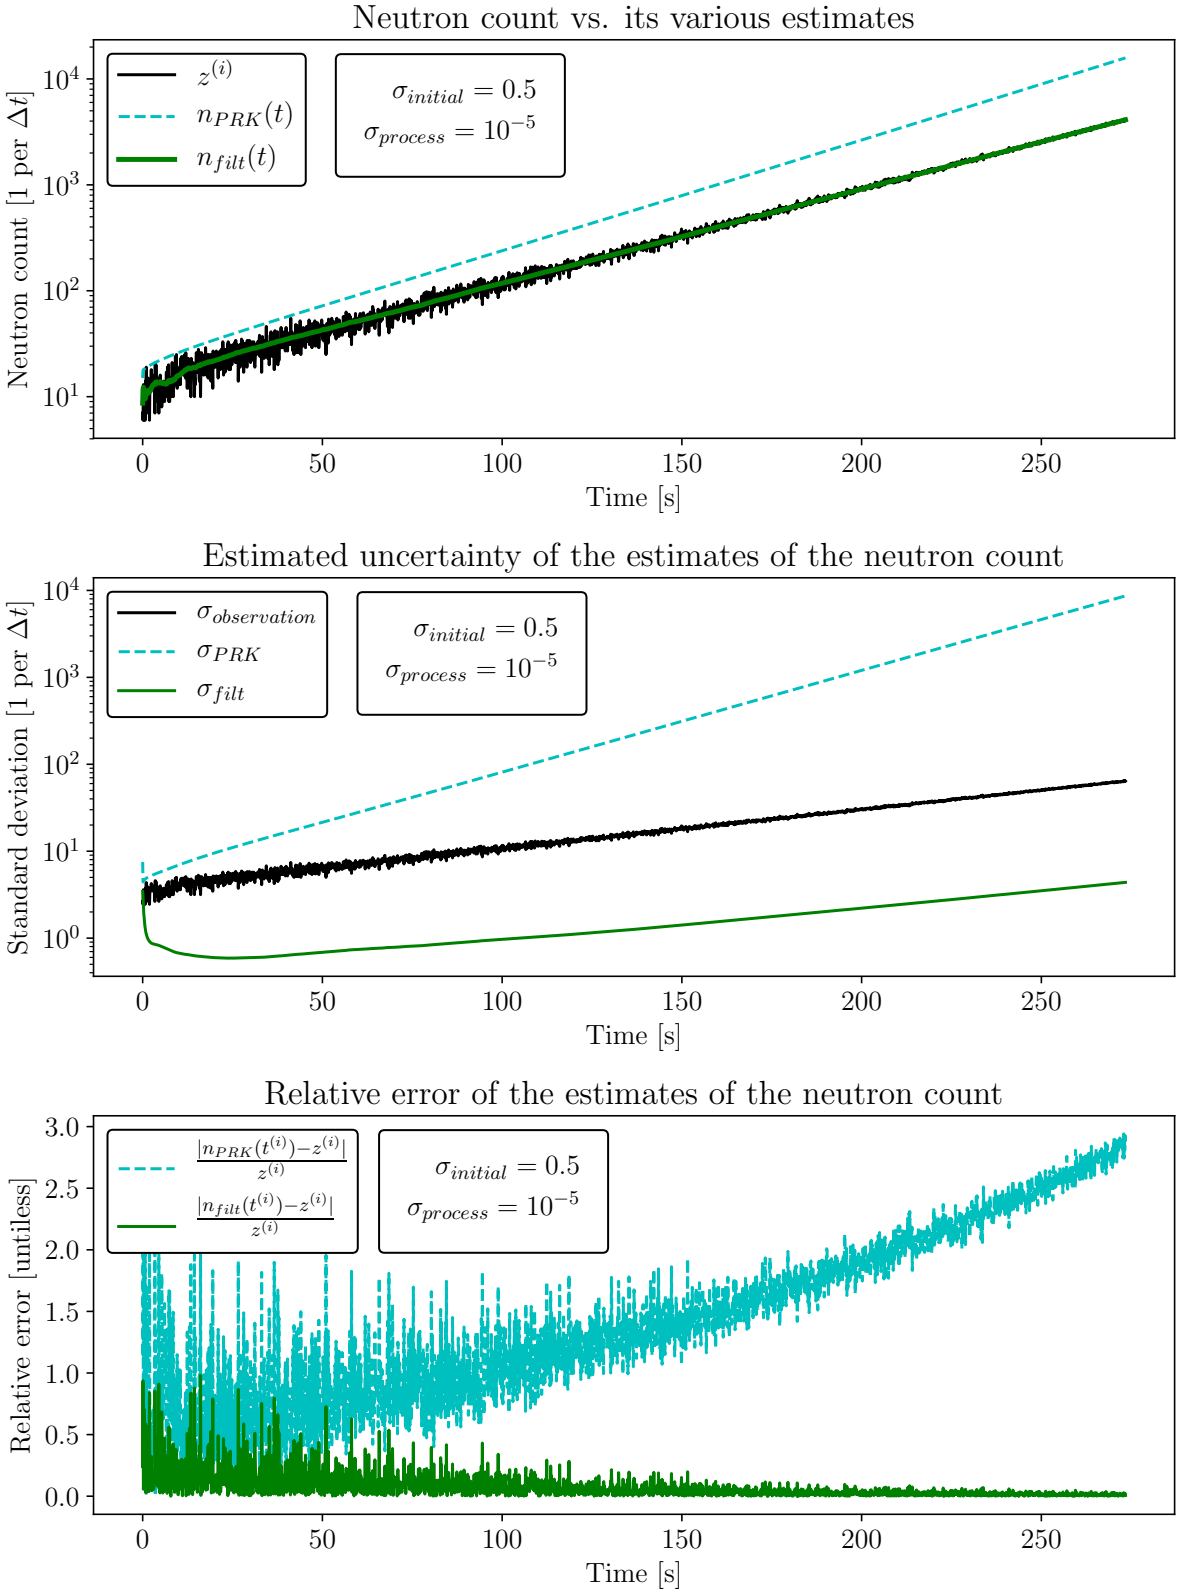

Figure 18: Estimates of the neutron population for $\sigma_{initial} = 0.5$ and $\sigma_{process} = 10^{-5}$, their estimated variance, and their relative error with respect to the observations. The time between subsequent observations is 0.1 s.

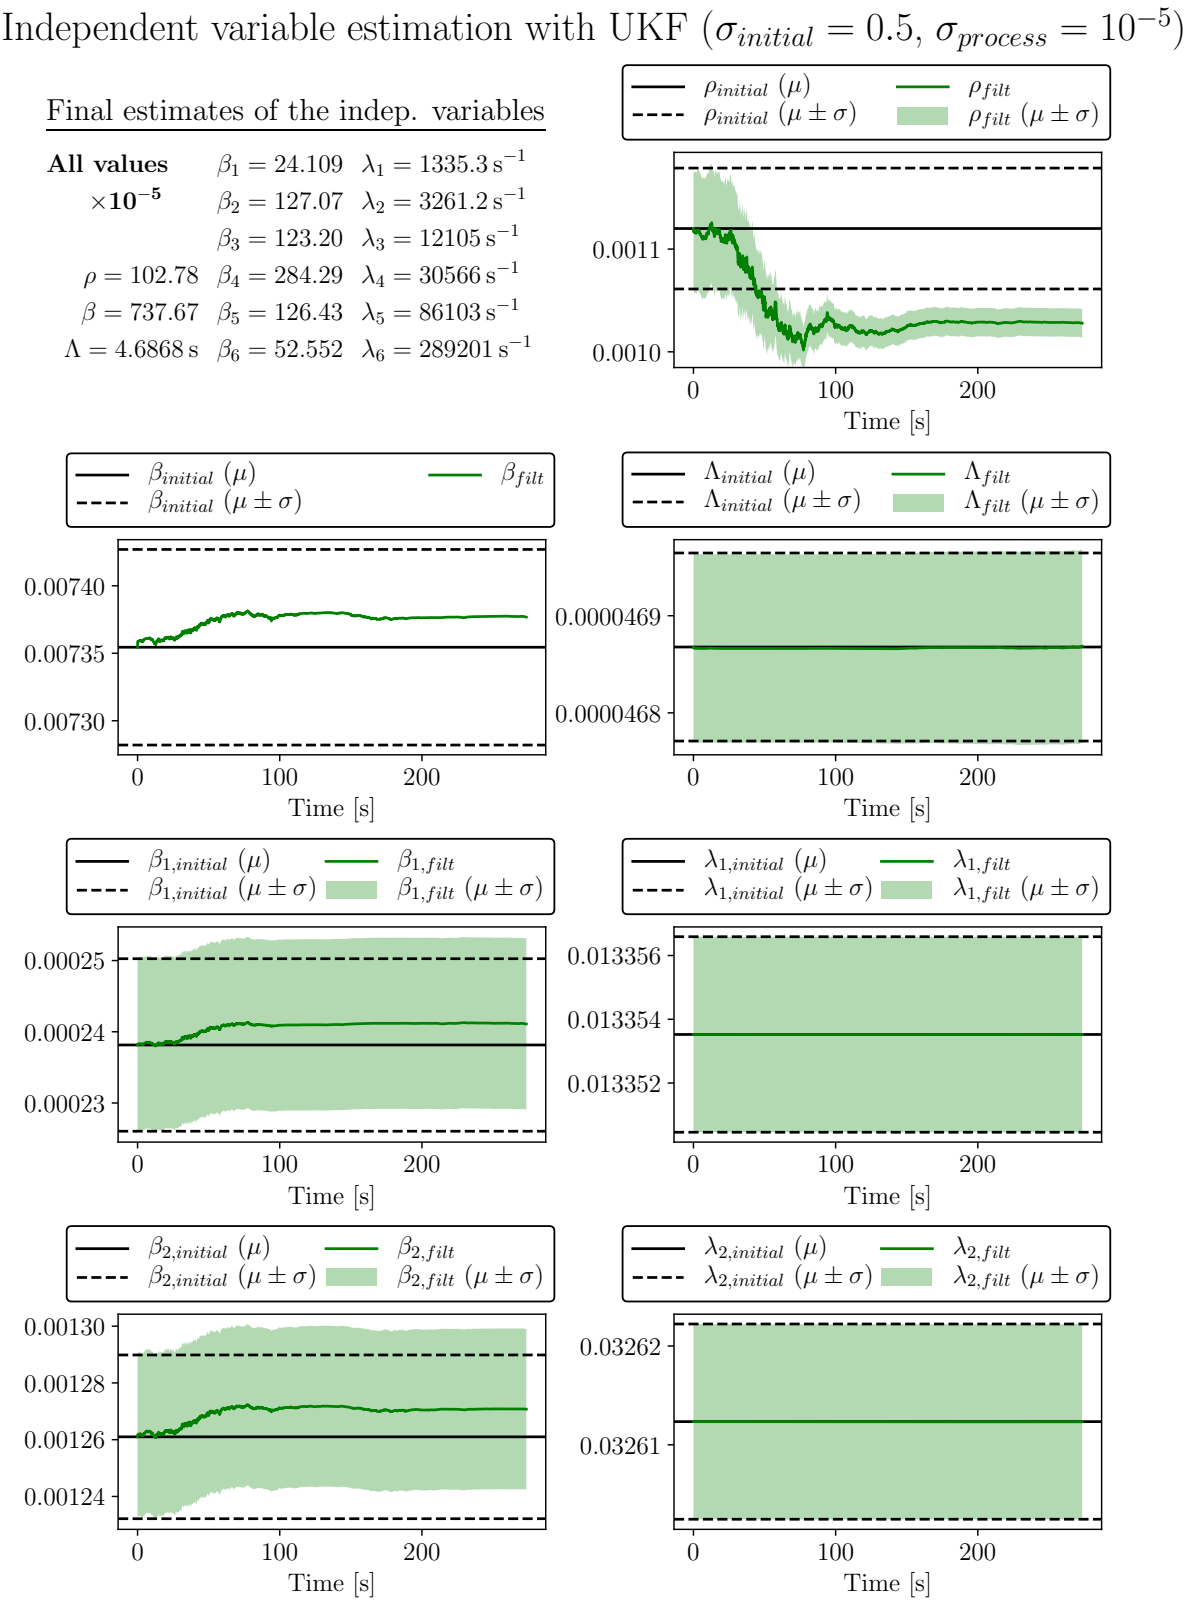

Figure 19: Estimates of the most relevant independent variables for $\sigma_{initial} = 0.5$ and $\sigma_{process} = 10^{-5}$, and their estimated 68% confidence intervals. The time between subsequent observations is 0.1 s.

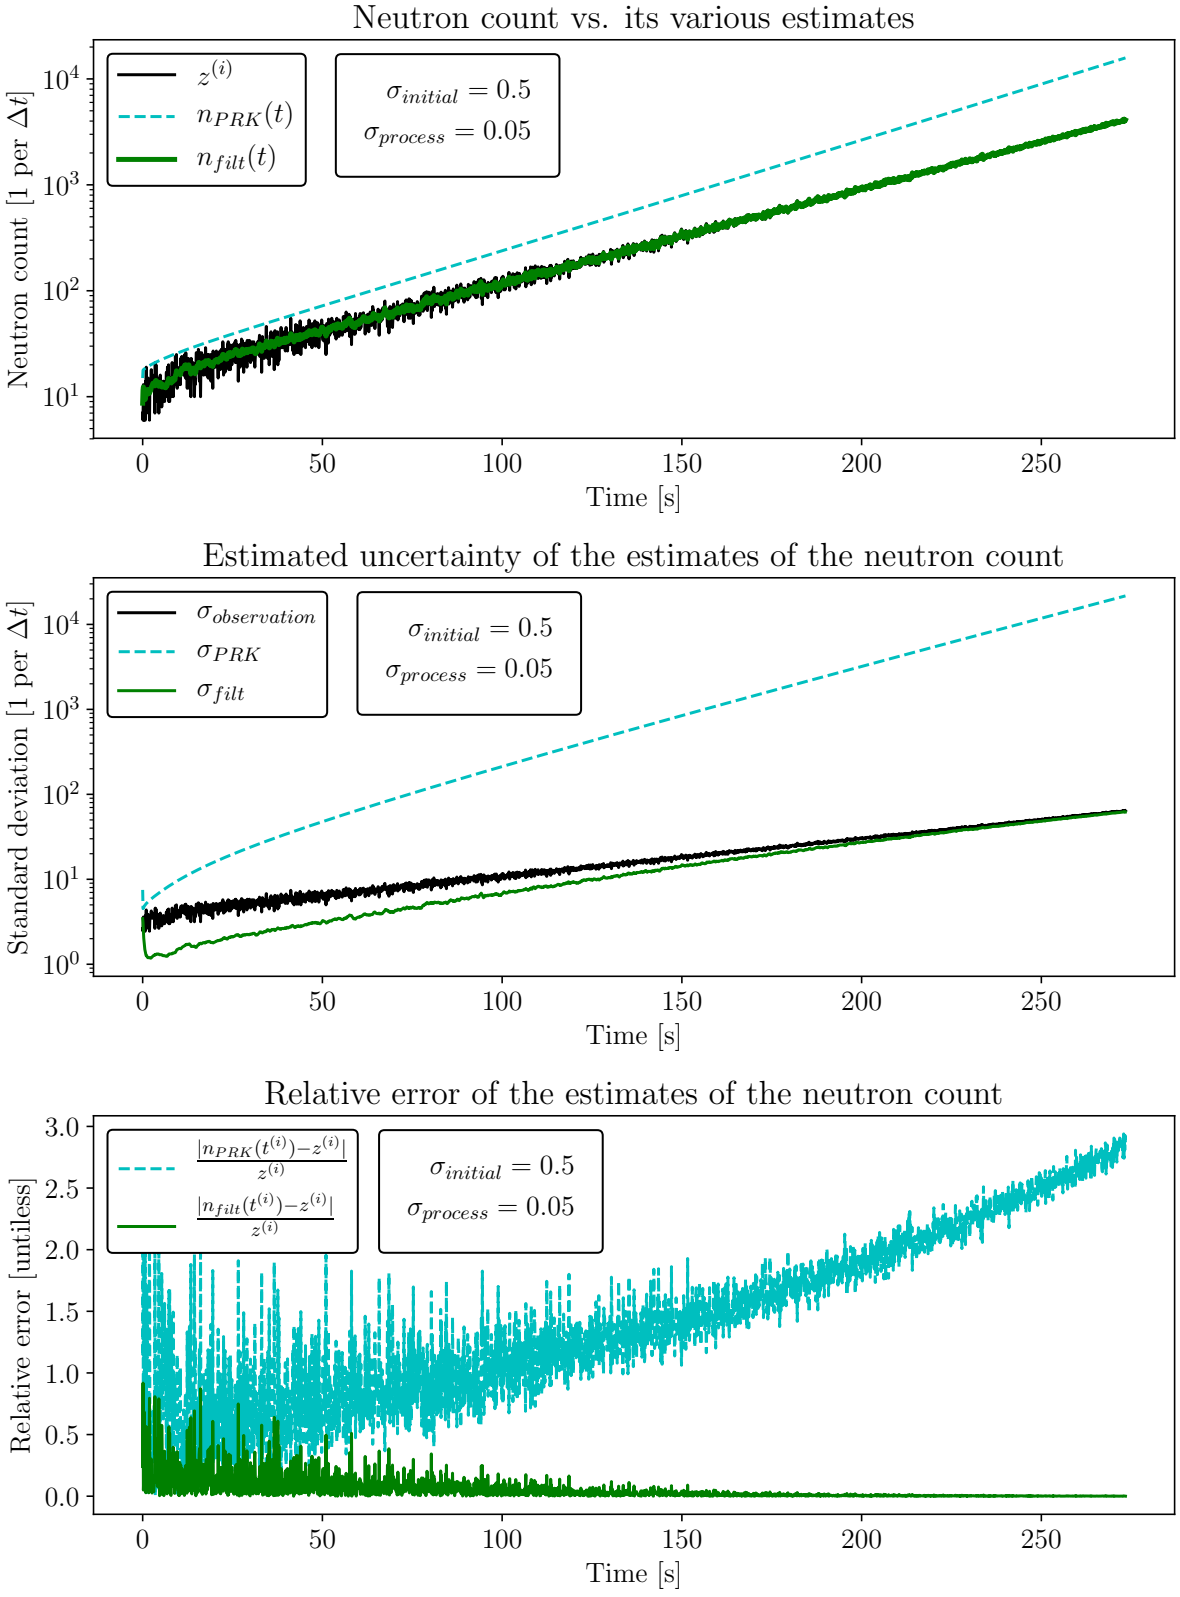

Figure 20: Estimates of the neutron population for $\sigma_{initial} = 0.5$, $\sigma_{process} = 0.05$, their estimated variance, and their relative error with respect to the observations. The time between subsequent observations is 0.1 s.

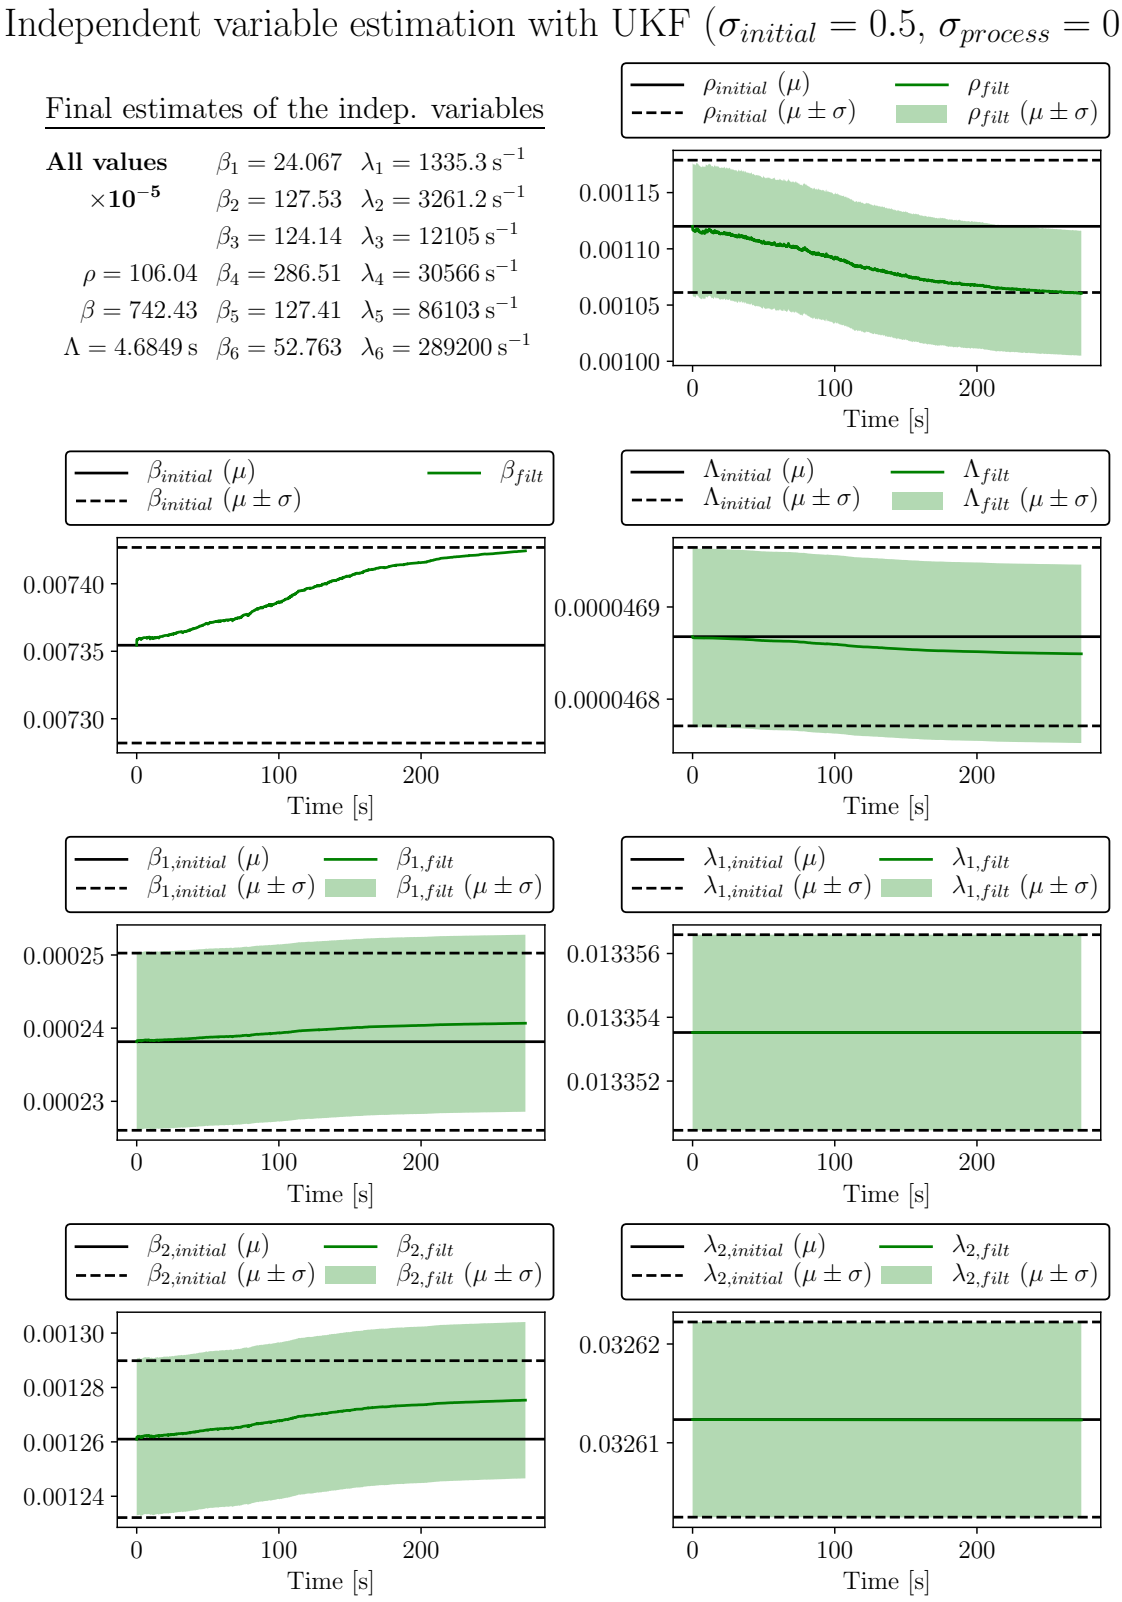

Figure 21: Estimates of the most relevant independent variables for $\sigma_{initial} = 0.5$, $\sigma_{process} = 0.05$, and their estimated 68% confidence intervals. The time between subsequent observations is 0.1 s.

6 Conclusion and Outlook

Existing estimates of the reactor-specific properties involved in the point reactor kinetics (PRK) equations were refined using the unscented Kalman filter (UKF). Firstly, the linear Kalman filter was introduced and applied to a fictitious problem. The theoretical background of the PRK equations and the UKF was then laid out before our application of the UKF to the six-group PRK model was presented. The proposed method was finally applied to a sequence of neutron counts acquired during previous experiments on the experimental thermonuclear reactor CROCUS at the LRS.

A significantly more accurate estimate of the reactivity in CROCUS during the experiment was obtained. The estimates of the — constant, hence more generally useful — reactor properties (the fractional yields β_ℓ and decay constants λ_ℓ of each neutron precursor group ℓ, and the mean neutron generation time Λ) were only marginally improved. This is due in part to the higher accuracy of their previously available estimates (especially λ_ℓ and Λ). For some, an additional inhibitor was their comparatively lower impact on the neutron population (especially β_5 and β_6).

It was shown that the UKF removes the risk of overfitting if the two free parameters of the model, $\sigma_{initial}$ and $\sigma_{process}$, are chosen sensibly. It was further demonstrated that the sensible range for these parameters is reasonably wide. However, trivial implementation errors may potentially set one or more initial values outside of their sensible range. It is therefore recommended to visualize the initial estimates (means *and* variances), as was done in figure 9.

The presented method should be applied to additional experimental data to validate the refined estimates produced in this work. To that end, the time at which the reactivity insertion occurs could be tracked accurately in future experiments. It would further be beneficial to measure the experimental reactivity insertion more precisely, in order to obtain more accurate estimates of the fractional yields. With regards to the estimation of the decay constants and the mean generation time, it appears that the UKF is out-competed by the method by which the initial estimates have been generated. Finally, the UKF could be applied to a more complex reactor kinetics model. For instance, the PRK equations could be augmented by temperature feedback coefficients of reactivity or substituted with a spatio-temporal reactor model.

References

[1] James Carpenter, Peter Clifford, and Paul Fearnhead. "Improved particle filter for nonlinear problems". In: *IEE Proceedings-Radar, Sonar and Navigation* 146.1 (1999), pp. 2–7.

[2] JFG De Freitas et al. "Sequential Monte Carlo methods for optimisation of neural network models". In: *Cambridge University Engineering Department, Cambridge, England, Technical Report TR-328* (1998).

[3] Geir Evensen. "The Ensemble Kalman Filter: theoretical formulation and practical implementation". In: *Ocean Dynamics; Heidelberg* 53.4 (Nov. 2003), pp. 343+i. ISSN: 16167341. DOI: http://dx.doi.org/10.1007/s10236-003-0036-9. URL: https://search.proquest.com/docview/218843560/abstract/3B21258AF8AD4760PQ/1 (visited on 06/10/2019).

[4] Ramsey Faragher. "Understanding the Basis of the Kalman Filter Via a Simple and Intuitive Derivation [Lecture Notes]". In: *IEEE Signal Processing Magazine* 29.5 (Sept. 2012), pp. 128–132. ISSN: 1053-5888. DOI: 10.1109/MSP.2012.2203621. URL: http://ieeexplore.ieee.org/document/6279585/ (visited on 05/28/2019).

[5] R. Frühwirth. "Application of Kalman filtering to track and vertex fitting". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 262.2 (Dec. 1987), pp. 444–450. ISSN: 01689002. DOI: 10.1016/0168-9002(87)90887-4. URL: https://linkinghub.elsevier.com/retrieve/pii/0168900287908874 (visited on 06/04/2019).

[6] E Hairer, SP NORSETT, and G Wanner. *Solving Ordinary, Differential Equations I, Nonstiff problems/E. Hairer, SP Norsett, G. Wanner, with 135 Figures, Vol.: 1*. Tech. rep. 2Ed. Springer-Verlag, 2000, 2000.

[7] Ehsan Hatami, Hassan Salarieh, and Naser Vosoughi. "Design of a fault tolerated intelligent control system for a nuclear reactor power control: Using extended Kalman filter". In: *Journal of Process Control* 24.7 (July 2014), pp. 1076–1084. ISSN: 09591524. DOI: 10.1016/j.jprocont.2014.04.012. URL: https://linkinghub.elsevier.com/retrieve/pii/S0959152414001139 (visited on 05/29/2019).

[8] Alain Hebert, ed. *Applied reactor physics*. Montrfffdfffdal: Presses internationales Polytechnique, 2009. 382 pp. ISBN: 978-2-553-01436-9.

[9] Lindsay Kleeman. "Understanding and Applying Kalman Filtering". In: (), p. 37.

[10] Jaakko Leppänen et al. "The Serpent Monte Carlo code: Status, development and applications in 2013". In: *Annals of Nuclear Energy* 82 (2015), pp. 142–150. ISSN: 0306-4549. DOI: https://doi.org/10.1016/j.anucene.2014.08.024. URL: http://www.sciencedirect.com/science/article/pii/S0306454914004095.

[11] Jeffery Lewins. *Nuclear reactor kinetics and control*. Elsevier, 2013.

[12] David McMahon and Adam Pierson. "A Taylor series solution of the reactor point kinetics equations". In: (), p. 13.

[13] *pykalman — pykalman 0.9.2 documentation*. URL: https://pykalman.github.io/ (visited on 06/07/2019).

[14] Simo Sarkka. *Bayesian Filtering and Smoothing*. Cambridge: Cambridge University Press, 2013. ISBN: 978-1-139-34420-3. DOI: 10.1017/CBO9781139344203. URL: http://ebooks.cambridge.org/ref/id/CBO9781139344203 (visited on 06/07/2019).

[15] Eric A Wan and Rudolph Van Der Merwe. "The unscented Kalman filter for nonlinear estimation". In: *Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium (Cat. No. 00EX373)*. Ieee. 2000, pp. 153–158.

[16] Greg Welch. "An Introduction to the Kalman Filter". In: (1997), p. 16.