

Fitting a Kalman Smoother to Data

Shane Barratt Stephen Boyd

March 7, 2020

Abstract

This paper considers the problem of fitting the parameters in a Kalman smoother to data. We formulate the Kalman smoothing problem with missing measurements as a constrained least squares problem and provide an efficient method to solve it based on sparse linear algebra. We then introduce the Kalman smoother tuning problem, which seeks to find parameters that achieve low prediction error on held out measurements. We derive a Kalman smoother auto-tuning algorithm, which is based on the proximal gradient method, that finds good, if not the best, parameters for a given dataset. Central to our method is the computation of the gradient of the prediction error on the held out measurements with respect to the parameters of the Kalman smoother; we describe how to compute this at little to no additional cost. We demonstrate the method on population migration within the United States as well as data collected from a smartphone’s IMU+GPS system while driving.

1 Introduction

Kalman smoothers are used widely to estimate the state of a linear dynamical system from noisy measurements [Swe58, Kal60]. In the traditional formulation, the dynamics and output matrices are considered fixed attributes of the system; the covariance matrices of the process and sensor noise are tuned by the designer, within some limits, to obtain good performance in simulation or on the actual system. For example, it is common to use noise levels in the Kalman smoother well in excess of the actual noise to obtain practical robustness [Gel74, §8].

In this paper we take a machine learning approach to the problem of tuning a Kalman smoother. We start with the observation that (by our definition) only the output is observed. This implies that the only way we can verify that a Kalman smoother is working well is to compare the outputs we predict with those that actually occur, on new or unseen test data, *i.e.*, data that was not used by the Kalman smoother. In machine learning terms, we would consider this output prediction error to be our error, with the goal of minimizing it. We consider the noise covariance matrices, as well as the system matrices, as *parameters* that can be varied to obtain different estimators, in this case, different Kalman smoothers. These

are varied, within limits, to obtain good test performance; this final Kalman smoother can then be checked on entirely new data.

To do this we formulate the Kalman smoothing problem, with missing observations, as a simple least squares problem, with a coefficient matrix that depends on the parameters, *i.e.*, the system and noise covariance matrices. We show how to efficiently compute the derivative of the test error with respect to the parameters, and use a simple proximal gradient method to update them to improve the test error. This method yields a Kalman smoother *auto-tuning* method. It uses one or more observed output sequences, and the usual prior knowledge in determining the starting system matrices as well as a description of the set over which we are allowed to vary them.

The same formulation works for tuning robust Kalman smoothers, where the process and sensor noises are assumed to have a non-Gaussian distribution, typically with fatter tails. In this case the least squares formulation of the Kalman smoother becomes a convex optimization problem, and the effect of the parameters is even less obvious, and therefore harder to tune manually. Our auto-tuning method extends immediately to such problems.

In summary, the contributions of this paper are:

- We describe a Kalman smoother auto-tuning method that requires only a dataset of measurements, which may have missing entries.
- We describe an efficient method for computing the gradient of the prediction error with respect to the Kalman smoother parameters, that incurs little to no additional computational cost on top of already smoothing.
- We provide an open-source implementation of the aforementioned ideas and illustrate the method via numerical examples that use real data.

Related work. The Kalman filter was independently invented by Swerling [Swe58] and Kalman [Kal60] around 1960, and one of its original applications was for space aircraft tracking in the Apollo navigation system [MS85]. The Kalman filter assumes *a priori* knowledge of the system matrices and noise statistics. Indeed, in his ground-breaking paper, Kalman remarked on the difficulty of identifying such parameters:

In real life, however, the situation is usually reversed. One is given the covariance matrix [of the state] and the problem is to get [the dynamics] and the statistical properties of [the disturbance]. This is a subtle and presently largely unsolved problem in experimentation and data reduction.

Despite its wide use and success, practitioners employing the Kalman smoother still have to resort to manually tuning its parameters. As a result, many have proposed methods for automatically tuning the parameters in Kalman smoothers. One of the first methods proposed was to jointly learn the parameters and state/output sequence using expectation-maximization [SS82]. More recent approaches employ different optimization algorithms, including the simplex algorithm [Pow02], coordinate descent [ACM⁺05], genetic algorithms

[OS00], nonlinear programming using finite differencing for the gradient [AE12], Bayesian optimization [CHJA18], and reinforcement learning [GES07].

Our approach is inspired by previous research on automatically tuning hyper-parameters in least squares [BB19]. Our paper departs from prior work on tuning Kalman filters in several ways. Since our Kalman smoother can deal with missing measurements, we can hold out measurements and use those to evaluate the smoother. Also, our method makes explicit use of the gradient of the loss with respect to the parameters, leading to a more efficient optimization algorithm than black box (or zeroth order) techniques, such as genetic algorithms and finite differencing.

2 Kalman smoother

System model. We consider a linear system with dynamics

$$x_{t+1} = Ax_t + w_t, \quad t = 1, \dots, T-1, \quad (1)$$

and output or sensor measurements

$$y_t = Cx_t + v_t, \quad t = 1, \dots, T. \quad (2)$$

Here $x_t \in \mathbf{R}^n$ is the state, $w_t \in \mathbf{R}^n$ is the process noise, $y_t \in \mathbf{R}^p$ is the output or sensor measurement, and $v_t \in \mathbf{R}^p$ is the sensor noise, at time t . The matrix $A \in \mathbf{R}^{n \times n}$ is the state dynamics matrix and $C \in \mathbf{R}^{p \times n}$ is the output matrix. We make the standard statistical assumptions that w_1, \dots, w_{T-1} are IID $\mathcal{N}(0, W)$ and v_1, \dots, v_T are IID $\mathcal{N}(0, V)$, where the symmetric positive definite matrices W and V are the process and sensor noise covariance matrices, respectively.

Missing measurements. We assume throughout that only the sequence y_t is observed. Indeed, we will assume that not all of the measurements are available to us. To model this, we modify the output equation (2) so that $y_t \in (\mathbf{R} \cup \{?\})^p$, where $?$ denotes a missing value. We have

$$(y_t)_i = (Cx_t + v_t)_i, \quad (t, i) \in \mathcal{K}, \quad (3)$$

where $\mathcal{K} \subseteq \{1, \dots, T\} \times \{1, \dots, p\}$ is the set of (scalar) outputs that are available. For $(t, i) \notin \mathcal{K}$, we take $(y_t)_i = ?$. We refer to entries of y_t that are real as *known measurements* and entries of y_t that have the value $?$ as *missing measurements*.

Smoothing. The goal in smoothing is to reconstruct or approximate the missing measurements given the known measurements. Since the outputs and states are jointly Gaussian, the maximum likelihood and conditional mean estimates of the missing output values are the same, and can be found as the solution of the constrained least squares problem

$$\begin{aligned} & \text{minimize} && \sum_{t=1}^{T-1} \|W^{-1/2}(\hat{x}_{t+1} - A\hat{x}_t)\|_2^2 + \\ & && \sum_{t=1}^T \|V^{-1/2}(\hat{y}_t - C\hat{x}_t)\|_2^2, \\ & \text{subject to} && (\hat{y}_t)_i = (y_t)_i, \quad (t, i) \in \mathcal{K}, \end{aligned} \quad (4)$$

with variables $\hat{x}_1, \dots, \hat{x}_T$ and $\hat{y}_1, \dots, \hat{y}_T$.

We note that problem (4) has a simple and widely used recursive solution for \hat{x}_t when $\mathcal{K} = (1, \dots, T) \times (1, \dots, p)$, and also when $T \rightarrow \infty$. This recursive solution is often referred to as the Kalman filter [Kal60].

Solving the Kalman smoothing problem. There are many ways to solve the Kalman smoothing problem (4). One method is to eliminate the equality constraint [BV04, §4.2.4] and solve the resulting unconstrained least squares problem, which has a banded coefficient matrix. This method has time and space complexity of order $T(n+p)^2$. We give some details on another method that has roughly the same complexity, but is simpler since it does not require eliminating the equality constraints.

Let $N = T(n+p)$ and define the vector $z \in \mathbf{R}^N$ as $z = (\hat{x}_1, \dots, \hat{x}_T, \hat{y}_1, \dots, \hat{y}_T)$. Using the variable z , we can express the estimation problem (4) compactly as the constrained least squares problem [BV18, §17.3]

$$\begin{aligned} & \text{minimize} && \frac{1}{2} \|Dz\|_2^2 \\ & \text{subject to} && Bz = c, \end{aligned} \tag{5}$$

where $B \in \mathbf{R}^{|\mathcal{K}| \times N}$ is a selector matrix and $c \in \mathbf{R}^{|\mathcal{K}|}$ contains the corresponding entries of y_t . Concretely, if we denote the elements in \mathcal{K} as $\mathcal{K}_1, \dots, \mathcal{K}_{|\mathcal{K}|}$, then if $\mathcal{K}_j = (t, i)$, the j th row of B is $e_{Tn+tp+i}$ and the j th entry of c is $(y_t)_i$. The matrix $D \in \mathbf{R}^{N-n \times N}$ is given by

$$D = \begin{bmatrix} D_{11} & 0 \\ D_{21} & D_{22} \end{bmatrix},$$

where

$$\begin{aligned} D_{11} &= \begin{bmatrix} -W^{-1/2}A & W^{-1/2} & 0 & 0 \\ 0 & \ddots & \ddots & 0 \\ 0 & 0 & -W^{-1/2}A & W^{-1/2} \end{bmatrix} \\ D_{21} &= \begin{bmatrix} -V^{-1/2}C & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & -V^{-1/2}C \end{bmatrix} \\ D_{22} &= \begin{bmatrix} V^{-1/2} & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & V^{-1/2} \end{bmatrix}. \end{aligned}$$

The matrices D and B are evidently very sparse, since each have a density of approximately $1/N$.

The optimality conditions for (5) can be expressed as

$$\begin{bmatrix} 0 & D^T & B^T \\ D & -I & 0 \\ B & 0 & 0 \end{bmatrix} \begin{bmatrix} z \\ v \\ \eta \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ c \end{bmatrix},$$

where $\eta \in \mathbf{R}^{|\mathcal{K}|}$ is the dual variable for the equality constraint and $v = Dz$. The KKT matrix, denoted by

$$M = \begin{bmatrix} 0 & D^T & B^T \\ D & -I & 0 \\ B & 0 & 0 \end{bmatrix},$$

is also very sparse, since B and D are sparse.

We assume for the remainder of the paper that M is full rank (if it is not, we can add a small amount of regularization to make it invertible). Therefore we can solve the KKT system using any method for solving a sparse system of linear equations, *e.g.*, a sparse LU factorization. The complexity of a sparse LU factorization mainly depends on the fill-in of the factorization, making it hard to analyze, but we have observed that the time complexity is roughly linear in T .

Judging a Kalman smoother. Suppose we have gathered a sequence of outputs denoted $y_1, \dots, y_T \in (\mathbf{R} \cup \{?\})^p$. We can judge how well a Kalman smoother is working on this sequence of observations by obscuring a fraction of the known outputs and comparing the outputs predicted by the Kalman smoother to those that actually occurred.

The first step in judging a Kalman smoother is to mask some fraction (*e.g.*, 20%) of the non-missing entries in the observations, denoted by the set $\mathcal{M} \subseteq (1, \dots, T) \times (1, \dots, p)$, resulting in a *masked trajectory* $\tilde{y}_1, \dots, \tilde{y}_T$. That is, we let $(\tilde{y}_t)_i = ?$ for $(t, i) \in \mathcal{M}$ and $(\tilde{y}_t)_i = (y_t)_i$ for $(t, i) \notin \mathcal{M}$.

We then solve the smoothing problem (4) with $y_t = \tilde{y}_t$ and known set $\mathcal{K} \setminus \mathcal{M}$, resulting in a predicted output trajectory $\hat{y}_1, \dots, \hat{y}_T$.

In order to judge the Kalman smoother, we calculate the *prediction error*, which is the squared difference between the predicted output trajectory and the actual trajectory in the entries that we masked,

$$L = \sum_{(t,i) \in \mathcal{M}} ((\hat{y}_t)_i - (y_t)_i)^2. \quad (6)$$

The goal in the sequel will be to adjust the parameters to minimize this error.

3 Kalman smoother auto-tuning

In this section we describe how to automatically tune the parameters in a Kalman smoother (that is, the dynamic matrices and covariance matrices) to minimize the prediction error on the held-out measurements (6). Once the parameters have been tuned, the Kalman smoother can be tested on another (unseen) output sequence.

Algorithm 3.1 *Kalman smoother auto-tuning.*

given initial hyper-parameter vector $\theta^1 \in \Theta$, initial step size t^1 , number of iterations n_{iter} , tolerance ϵ .

for $k = 1, \dots, n_{\text{iter}}$

1. *Filter the output sequence.* Let $\hat{y}_1, \dots, \hat{y}_T$ be the solution to (4).
2. *Compute the gradient of the prediction error.* $g^k = \nabla_{\theta} L(\theta)$.
3. *Compute the gradient step.* $\theta^{k+1/2} = \theta^k - t^k g^k$.
4. *Compute the proximal operator.* $\theta^{\text{tent}} = \mathbf{prox}_{t^k r}(\theta^{k+1/2})$.
5. **if** $F(\theta^{\text{tent}}) \leq F(\theta^k)$:
 - Increase step size and accept update.* $t^{k+1} = (1.5)t^k$; $\theta^{k+1} = \theta^{\text{tent}}$.
 - Stopping criterion.* **quit** if $\|(\theta^k - \theta^{k+1})/t^k + (g^{k+1} - g^k)\|_2 \leq \epsilon$.
6. **else** *Decrease step size and reject update.* $t^{k+1} = (0.5)t^k$; $\theta^{k+1} = \theta^k$.

end for

Kalman smoother parameters. A Kalman smoother has four parameters, which we denote by

$$\theta = (A, W^{-1/2}, C, V^{-1/2}) \in \mathbf{R}^{n \times n} \times \mathbf{R}^{n \times n} \times \mathbf{R}^{p \times n} \times \mathbf{R}^{p \times p}.$$

Evidently, this parametrization of the Kalman smoother is not unique. For example, if $T \in \mathbf{R}^{n \times n}$ is invertible, then $\tilde{x}_t = T x_t$, $\tilde{A} = T A T^{-1}$, $\tilde{W} = T^{-1} W T^{-T}$, $\tilde{C} = C T^{-1}$, and $\tilde{V} = V$ gives another representation of (4). As another example, scaling W and V by $\alpha > 0$ gives an equivalent representation of (4).

Auto-tuning problem. The prediction error L in (6) is a function of the parameters, and from here onwards we denote that function by $L(\theta)$. To tune the Kalman smoother, we propose solving the optimization problem

$$\text{minimize } F(\theta) = L(\theta) + r(\theta), \tag{7}$$

with variable θ (the parameters of the Kalman smoother), where $r : \Theta \rightarrow \mathbf{R}$ is a regularization function. Here Θ denotes the set of allowable parameters and can, for example, include constraints on what parameters we are allowed to change. (The function r evaluates to $+\infty$ for $\theta \notin \Theta$, thus constraining θ to be in Θ .)

The objective function $F : \Theta \rightarrow \mathbf{R}$ is composed of two parts: the prediction error and the regularization function. The first term encourages the Kalman smoother to have the same outputs as those observed, and the second encourages the parameters to be simpler or closer to an initial guess.

Regularization functions. There are many possibilities for the regularization function r ; here we describe a few. Suppose we have some initial guess for A , denoted A_{nom} . We could then penalize deviations of A from A_{nom} by letting, *e.g.*,

$$r(\theta) = \|A - A_{\text{nom}}\|_F^2.$$

As another example, suppose we suspected that C was low rank; then we could use

$$r(\theta) = \|C\|_*,$$

where $\|C\|_*$ is the nuclear norm of C , *i.e.*, the sum of the singular values of C . This regularizer encourages C to be low rank. Of course, any combination of these regularization functions is possible.

Allowable sets. There are also many possibilities for Θ , the allowable set of parameters. One option is to only allow certain entries of A to vary by letting the set of allowable A matrices be

$$\{A \mid A_{ij} = (A_{\text{nom}})_{ij}, (i, j) \in \Omega\}$$

for some set Ω . If we wanted to keep A fixed, we could let $\Theta = \{A_{\text{nom}}\}$. Another sensible option is to let A vary within a box by letting the set of allowable A matrices be

$$\{A \mid \|A - A_{\text{nom}}\|_{\infty} \leq \rho\},$$

for some nominal guess A_{nom} and maximum allowable change $\rho > 0$.

Solution method. The auto-tuning problem (7) is in general nonconvex, even if Θ and r are convex, so it is very difficult to solve exactly. Therefore, we must resort to a local or heuristic optimization method to (approximately) solve it. There are many methods that we could use to (approximately) solve the auto-tuning problem (see, *e.g.*, [DR56, LM79, Sho85, Nes13b]). In this paper we employ one of the simplest, the proximal gradient method [Mar70, Nes13a], since F is differentiable in θ (see Sec. 3).

The proximal gradient method is described by the iteration

$$\theta^{k+1} = \mathbf{prox}_{t^k r}(\theta^k - t^k \nabla_{\theta} L(\theta)),$$

where k is the iteration number, $t^k > 0$ is a step size, and the proximal operator of the function tr is defined as

$$\mathbf{prox}_{tr}(\nu) = \underset{\theta \in \Theta}{\operatorname{argmin}} (tr(\theta) + (1/2)\|\theta - \nu\|_2^2).$$

When Θ is a convex set and r is convex, evaluating the proximal operator requires solving a (small) convex optimization problem. Also, the proximal operator often has a (simple) closed-form expression [PB14]. We note that r need not be differentiable.

We employ the proximal gradient method with the adaptive step size scheme and stopping condition described in [BB19]. The full algorithm for Kalman smoother auto-tuning is summarized in Algorithm 3.1.

Computing the gradient. Evidently, the proximal gradient method requires computing the gradient of the prediction error with respect to the parameters, denoted $\nabla_{\theta} L(\theta)$. The sensitivity analysis of Kalman smoothing has previously been considered in the forward direction, *i.e.*, how changes in the parameters affect the output [Gel74, §7].

To do this, we first form the gradient of L with respect to $\hat{y}_1, \dots, \hat{y}_T$, given by

$$\nabla_{(\hat{y}_t)_i} L = \begin{cases} 2((\hat{y}_t)_i - (y_t)_i) & (t, i) \in \mathcal{M}, \\ 0 & \text{otherwise.} \end{cases}$$

Next we form the gradient of L with respect to the solution to (5), which is given by

$$g = \nabla_{(z, \eta, v)} L = (0, \nabla_{\hat{y}_1} L, \dots, \nabla_{\hat{y}_T} L, 0, 0).$$

Next we solve the linear system

$$M \begin{bmatrix} q_1 \\ q_2 \\ q_3 \end{bmatrix} = -g,$$

which only requires a backsolve if we have already factorized M . Since the KKT system M is invertible, the prediction error is indeed differentiable.

The next step is to form the gradient of L with respect to the coefficient matrix D , which is given by

$$G = \nabla_D L = D(q_1 z^T + z q_1^T).$$

Since θ only affects D at certain entries, we only need to compute G at those entries. That is, we compute G at the entries

$$G = \begin{bmatrix} G_{11} & 0 \\ G_{21} & G_{22} \end{bmatrix},$$

where

$$\begin{aligned} G_{11} &= \begin{bmatrix} G_{11} & G_{12} & 0 & 0 \\ 0 & \ddots & \ddots & 0 \\ 0 & 0 & G_{(T-1)(T-1)} & G_{(T-1)T} \\ G_{T1} & 0 & 0 & \end{bmatrix}, \\ G_{21} &= \begin{bmatrix} G_{T1} & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & G_{(2T-1)T} \end{bmatrix}, \\ G_{22} &= \begin{bmatrix} G_{(T+1)T} & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & G_{(2T-1)(2T-1)} \end{bmatrix}, \end{aligned}$$

which we can efficiently do since

$$G_{ij} = (Dq_1)_i z_j + (Dz)_i (q_1)_j.$$

The final step is to form the gradients with respect to the parameters, which are given by

$$\begin{aligned} \nabla_A L &= -(W^{-1/2})^T \sum_{t=1}^{T-1} G_{tt}, \\ \nabla_{W^{-1/2}} L &= \sum_{t=1}^{T-1} G_{t(t+1)} - G_{tt} A^T, \\ \nabla_C L &= -(V^{-1/2})^T \sum_{t=1}^T G_{(T-1+t)t}, \\ \nabla_{V^{-1/2}} L &= \sum_{t=1}^T G_{(T+t)(T+1+t)} - G_{(T-1+t)t} C^T. \end{aligned}$$

The complexity of computing the gradient is roughly the same complexity as solving the original problem, since it requires the solution of another linear system. However, the time required to compute the gradient is often lower since we cache the factorization of M .

4 Experiments

In this section, we describe our implementation of Kalman smoother auto-tuning, as well as the results of some numerical experiments that illustrate the method. All experiments were performed on a single core of an unloaded Intel i7-8770K CPU.

Implementation. We have implemented the Kalman smoother auto-tuning method described in this paper as an open-source Python package, available at

https://github.com/cvxgrp/auto_ks.

Our CPU-based implementation has methods for performing Kalman smoothing with missing measurements and for tuning the matrices in the Kalman smoother (Algorithm 3.1). Our only dependencies are `scipy` [JOP⁺01], which we use for sparse linear algebra, and `numpy` [WCV11], which we use for dense linear algebra.

Performance. We ran our Kalman smoothing function on random problems with $n = p = 10$. figure 1 shows the execution time, averaged over ten runs, of solving the smoothing problem (denoted as *forward* in the figure), as well as computing the derivative with respect to the parameters (denoted as *backward* in the figure). As expected, the time required to compute the solution and its derivative is roughly linear in the length of the sequence T . Empirically, we found that the time required to compute the derivative is roughly half of the time required to compute the solution. We remark that our method is very efficient and effortlessly scales to extremely large problem sizes.

4.1 Human migration example

Suppose we have n states, where the i th state has a population x_i . At some cadence, say yearly, a fraction of people in each state decide to move to another state. We take noisy measurements of the population in some of the states and wish to infer the population in every state, including those we have not even measured.

Dynamics. The vector $x_t \in \mathbf{R}^n$ denotes the population in each state at year t . The dynamics are described by

$$x_{t+1} = Ax_t + w_t,$$

where $A \in \mathbf{R}_+^{n \times n}$. Here A_{ij} denotes the fraction of the population in state j that move to state i each year. The process noise w_t includes births, deaths, and immigration from or to outside places.

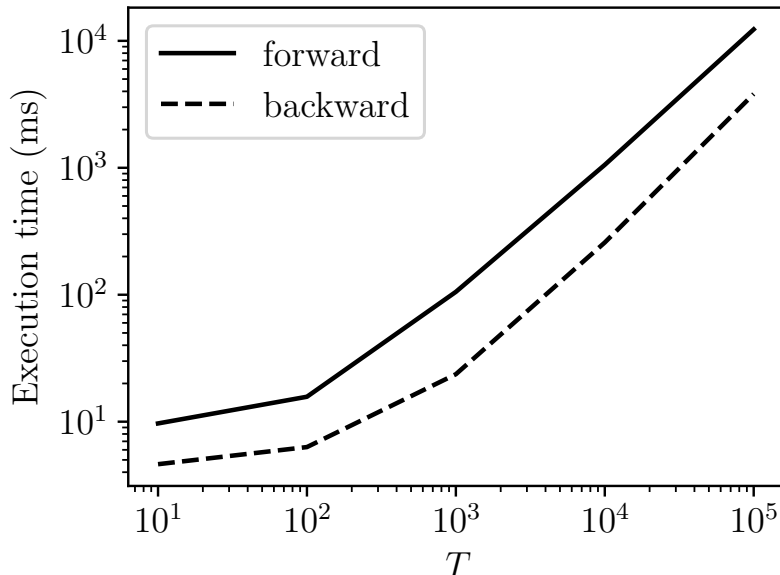


Figure 1: Method timings for a random problem with $n = p = 10$.

Outputs. Each year, we take noisy measurements of the populations in some of the states. The outputs are described by

$$y_t = x_t + v_t.$$

We use the set of known measurements \mathcal{K} to denote the measurements we actually have access to.

United States population data. We gathered yearly population data (in millions of people) for the $n = 48$ states in the continental U.S from the U.S. Census Bureau [cen19]. The data includes all years from 1900 to 2018.

Experiment details. Our goal is to learn the dynamics matrix, dynamics covariance, and output covariance via Kalman smoother auto-tuning. To this end, we use the regularization function $r(\theta) = 0$ and allowable set

$$\Theta = \{(A, \mathbf{diag}(w), C_0, \mathbf{diag}(v)) \mid A \in \mathbf{R}_+^{n \times n}, w, v \in \mathbf{R}_+^n\}.$$

We initialize the parameters as

$$A_0 = I, \quad W_0^{-1/2} = (30)I, \quad C_0 = I, \quad V_0^{-1/2} = (10)I$$

For each year, we pick 30 out of the 48 states at random to be measured. In each year, of those measured, we pick 12 at random to be missing and 5 at random to be part of the test set. We ran the method for 50 iterations with $t_0 = 1 \times 10^{-4}$.

Table 1: Entries in the learned A matrix.

From	To	Fraction
Texas	California	0.56%
Florida	California	0.42%
California	Texas	0.30%
Texas	Florida	0.29%
Pennsylvania	California	0.28%

Results. The prediction error decreased from 0.0097 to 0.0058. The test error decreased from 0.0041 to 0.0030. The algorithm took 31 seconds to run. Besides the purely numerical results, there are interesting interpretations of the resulting parameters. For example, we can interpret the off-diagonal entries in the A matrix as the fraction of the population in one state that migrates to another state over the course of one calendar year. The biggest such entries are displayed in table 1.

4.2 Vehicle smoothing example

In vehicle smoothing, we have noisy measurements of the position, velocity, and acceleration of a vehicle over time, and wish to infer the true position, velocity, and acceleration at each time step.

Dynamics. The state $x_t = (p_t, v_t, a_t)$ is composed of the position $p_t \in \mathbf{R}^3$, the velocity $v_t \in \mathbf{R}^3$, and the acceleration $a_t \in \mathbf{R}^3$. The dynamics are described by

$$x_{t+1} = \begin{bmatrix} I & hI & 0 \\ 0 & I & hI \\ 0 & 0 & I \end{bmatrix} x_t + w_t.$$

Outputs. Any output vector and linear output matrix is possible. In this specific example we use $y_t = (\hat{p}_t, \hat{a}_t, (\hat{v}_t)_1, (\hat{v}_t)_2)$, so $p = 8$. The output is described by

$$y_t = \begin{bmatrix} I & 0 & 0 & 0 \\ 0 & 0 & 0 & I \\ 0 & I_2 & 0 & 0 \end{bmatrix} x_t + v_t.$$

Data. We used the Sensor Play data recorder iOS application [Bro14] to record the acceleration, attitude, latitude, longitude, heading, speed, and altitude of an iPhone mounted on a passenger vehicle. We converted the latitude and longitude into local North-East-Up coordinates, used the heading to convert the speed into a velocity in local coordinates, and used the attitude to orient the acceleration to local coordinates. We recorded data for a total of 330 seconds with a sampling frequency of 100 Hz, resulting in $T = 33000$ measurements.

Experiment details. Our goal is to learn the state and observation covariance matrices, via Kalman smoother auto-tuning. To this end, we penalize the off-diagonal entries of $W^{-1/2}$ and $V^{-1/2}$ by using the regularization function

$$r(\theta) = \alpha \sum_{i \neq j} (W_{ij}^{-1/2})^2 + \alpha \sum_{i \neq j} (V_{ij}^{-1/2})^2,$$

where α is a hyper-parameter (we use $\alpha = 1 \times 10^{-4}$) and

$$\Theta = \{(A_0, W^{-1/2}, C_0, V^{-1/2}) \mid W \in \mathbf{S}_+^9, V \in \mathbf{S}_+^8\}.$$

We initialize the covariances as $W_0^{-1/2} = I$, $V_0^{-1/2} = (0.01)I$, and initialize A_0 and C_0 as given in Sec. 4.2 and Sec. 4.2 respectively. We consider all measurement indices where the GPS or velocity change as known (since GPS is only useful when it changes) and all acceleration indices as known. We use 20% of the known position measurements as the missing measurements and another 20% as the test measurements. We ran the method for 25 iterations with $t_0 = 1 \times 10^{-2}$.

Results. The prediction error decreased from 13.23 to 2.97. The test error decreased from 16.57 to 1.37. The algorithm took 135 seconds to run. The diagonals of the final state and output covariance matrices were

$$\begin{aligned} W &= \mathbf{diag}(5.7, 14.9, 1.1, 0.8, 1.3, 1.0, 1.0, 1.0, 1.3), \\ V &= \mathbf{diag}(0.2, 0.4, 9.6, 2.7, 2.3, 0.2, 1.4, 2.4). \end{aligned}$$

(Note that these matrices can be scaled and the smoothing result is the same, so only relative magnitude matters.) We observe that there is more state noise in north and east dimensions than up, which makes sense. Also, there is less state noise in velocity than in position. We also observe that there is much higher measurement noise for z direction in GPS, which is true with GPS. In figure 2 we show the position estimates before and after tuning. Visually, we see significant improvement from tuning.




Figure 2: Smoothed position estimates before and after tuning.

Acknowledgments

S. Barratt is supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-1656518.

References

- [ACM⁺05] P. Abbeel, A. Coates, M. Montemerlo, A. Ng, and S. Thrun. Discriminative training of Kalman filters. In *RSS*, volume 2, page 1, 2005.
- [AE12] Dylan Asmar and Greg Eslinger. Nonlinear programming approach to filter tuning, 2012.
- [BB19] S. Barratt and S. Boyd. Least squares auto-tuning. *arXiv preprint arXiv:1904.05460*, 2019.
- [Bro14] Philip Broder. Sensor Play - data recorder, iOS App, 2014.
- [BV04] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.
- [BV18] S. Boyd and L. Vandenberghe. *Introduction to applied linear algebra: vectors, matrices, and least squares*. Cambridge University Press, 2018.
- [cen19] U.S. census bureau, resident population by state. <https://fred.stlouisfed.org/release/tables?rid=118&eid=259194>, 2019.
- [CHJA18] Z. Chen, C. Heckman, S. Julier, and N. Ahmed. Weak in the NEES?: Auto-tuning Kalman filters with bayesian optimization. In *Intl. Conf. Information Fusion*, pages 1072–1079. IEEE, 2018.
- [DR56] J. Douglas and H. Rachford. On the numerical solution of heat conduction problems in two and three space variables. *Transactions of the American Mathematical Society*, 82(2):421–439, 1956.
- [Gel74] A. Gelb. *Applied optimal estimation*. MIT press, 1974.
- [GES07] C. Goodall and N. El-Sheimy. Intelligent tuning of a Kalman filter using low-cost mems inertial sensors. In *Proc. Intl. Symposium Mobile Mapping Technology*, pages 1–8, 2007.
- [JOP⁺01] E. Jones, T. Oliphant, P. Peterson, et al. SciPy: Open source scientific tools for Python, 2001. [Online; accessed 9/25/2019].
- [Kal60] R. Kalman. A new approach to linear filtering and prediction problems. *Journal Basic Engineering*, 82(1):35–45, 1960.

- [LM79] P. Lions and B. Mercier. Splitting algorithms for the sum of two nonlinear operators. *Numerical Analysis*, 16(6):964–979, 1979.
- [Mar70] B. Martinet. Brève communication. régularisation d’inéquations variationnelles par approximations successives. *Mathematical Modelling and Numerical Analysis*, 4(R3):154–158, 1970.
- [MS85] L. McGee and S. Schmidt. Discovery of the Kalman filter as a practical tool for aerospace and industry. 1985.
- [Nes13a] Y. Nesterov. Gradient methods for minimizing composite functions. *Mathematical Programming*, 140(1):125–161, 2013.
- [Nes13b] Y. Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2013.
- [OS00] Yaakov Oshman and Ilan Shaviv. Optimal tuning of a Kalman filter using genetic algorithms. In *AIAA GNC Conference*, page 4558, 2000.
- [PB14] N. Parikh and S. Boyd. Proximal algorithms. *Foundations and Trends® in Optimization*, 1(3):127–239, 2014.
- [Pow02] Thomas Powell. Automated tuning of an extended Kalman filter using the downhill simplex algorithm. *Journal Guidance, Control, and Dynamics*, 25(5):901–908, 2002.
- [Sho85] N. Shor. *Minimization methods for non-differentiable functions*, volume 3. Springer Science & Business Media, 1985.
- [SS82] Robert Shumway and David Stoffer. An approach to time series smoothing and forecasting using the EM algorithm. *Journal Time Series Analysis*, 3(4):253–264, 1982.
- [Swe58] P. Swerling. *A proposed stagewise differential correction procedure for satellite tracking and prediciton*. Rand Corporation, 1958.
- [WCV11] S. Walt, S. Colbert, and G. Varoquaux. The numpy array: a structure for efficient numerical computation. *Computing in Science & Engineering*, 13(2):22–30, 2011.