


# Introduction to recursive machine learning

## Applications

Juan Pablo Carbajal  
juanpablo.carbajal@ost.ch

Eastern Switzerland University of Applied Sciences OST  
Institute for Energy Technology IET  
Scientific Computing and Engineering Group SCE

January 2023 - Rapperswil



Recap and Q&A

$$\begin{aligned}
 \text{while no measurement: predict} & \begin{cases} \hat{\mathbf{m}}_k = \mathbf{A}_k \mathbf{m}_{k-1} \\ \hat{\mathbf{P}}_k = \mathbf{A}_k \mathbf{P}_{k-1} \mathbf{A}_k^\top + \mathbf{Q}_{k-1} \end{cases} \\
 \text{predict measurement} & \begin{cases} \hat{\mathbf{y}}_k = \mathbf{H}_k \hat{\mathbf{m}}_k \\ \hat{\mathbf{S}}_k = \mathbf{H}_k \hat{\mathbf{P}}_k \mathbf{H}_k^\top + \mathbf{R}_k \end{cases} \\
 \text{on measurement: update} & \begin{cases} \mathbf{K}_k = \hat{\mathbf{P}}_k \mathbf{H}_k^\top \hat{\mathbf{S}}_k^{-1} \\ \mathbf{m}_k = \hat{\mathbf{m}}_k + \mathbf{K}_k (\mathbf{y}_k - \hat{\mathbf{y}}_k) \\ \mathbf{P}_k = \hat{\mathbf{P}}_k - \mathbf{K}_k \hat{\mathbf{S}}_k \mathbf{K}_k^\top \end{cases}
 \end{aligned}$$

Implement these functions in a programming language:

- $\mathbf{m}, \mathbf{P} \leftarrow \text{kf\_predict}(\mathbf{m}, \mathbf{P}, \mathbf{A}, \mathbf{Q})$
- $\mathbf{m}, \mathbf{P}, \mathbf{K} \leftarrow \text{kf\_update}(\mathbf{m}, \mathbf{P}, \mathbf{H}, \mathbf{R}, \mathbf{y})$
- $\mathbf{y}, \mathbf{S} \leftarrow \text{kf\_measure}(\mathbf{m}, \mathbf{P}, \mathbf{H}, \mathbf{R})$

1. Propagation of uncertainty to output using best estimate
2. Kalman gain
3. Update mean by projecting error on Kalman gain
4. Update cov by propagating uncertainty on measurement backward

Do not forget to update your measurement prediction after and update.

Which matrices participate in each function?

How would you determine which states are most sensitive to the residuals  $\mathbf{y}_k - \hat{\mathbf{y}}_k$ ?

These libraries are relatively easy to follow, and have a book for documentation:

- FilterPy
  - <https://filterpy.readthedocs.io/en/latest/>
  - book: [https://nbviewer.org/github/rlabbe/Kalman-and-Bayesian-Filters-in-Python/blob/master/table\\_of\\_contents.ipynb](https://nbviewer.org/github/rlabbe/Kalman-and-Bayesian-Filters-in-Python/blob/master/table_of_contents.ipynb)
- ekfukf
  - MATLAB/Octave <https://github.com/EEA-sensors/ekfukf>,
  - Octave installable <https://github.com/kakila/ekfukf>
  - book [http://users.aalto.fi/~ssarkka/pub/cup\\_book\\_online\\_20131111.pdf](http://users.aalto.fi/~ssarkka/pub/cup_book_online_20131111.pdf) (if you like it, buy it!)

They are not performant (pykalman is abandoned), for prototyping.

Consider Julia: <https://github.com/JuliaGNSS/KalmanFilters.jl>

There are many libraries for Kalman filter, search for the one that suits you.

In my experience, I use simple ones for quickly prototyping, then implement what I need case specific. Algorithms are very simple.

Difficulty in inverting matrices: good algorithms are already available.

Most relevant matrices are symmetric positive definite.

General probabilistic inference: PyMC3 <https://docs.pymc.io/en/stable/>

Use your code or a library to implement:

- recursive linear regression  $y_k = t_k a + b + \epsilon_k$
- Delayed linear regression  $y_k = w_0 + w_1 y_{k-1} + \dots + w_d y_{k-d} + \epsilon_k$
- Extreme Learning Machines

$$y_k = \sum_{d=1}^D \Phi[\mathbf{y}_{k-d}] \mathbf{w}_d^\top$$

$$\Phi[\mathbf{x}] = [\phi_1(\mathbf{x}) \quad \dots \quad \phi_n(\mathbf{x})], \quad \mathbf{w}_d = [w_{d,1} \quad \dots \quad w_{d,n}]$$

How does it differ from:

$$y_{d,k} = \Phi[y_{k-d}] \mathbf{w}^\top$$

Coffee break: 10 minutes

Participant examples

- Frédéric Bless: Control of a heat pump test bench
- Rafael Graf: Incineration plant model validation. Inspired by <https://www.frontiersin.org/articles/10.3389/fenrg.2020.00049/full>
- Pirmin Weigele: Predictive damping in electromagnetic force compensation
- Matthias Frommelt: End-of-life of Li-ion batteries prediction using real-time data



**Lunch: 1 hour**

## Discretization of ODEs

# What's an ODE?

Ordinary differential equations (ODE) define a function (the unknown) via a relation of its derivatives:

$$\frac{d^2x(t)}{dt^2} + \gamma \frac{dx(t)}{dt} + \omega^2 x(t) = w(t)$$

That's a **2nd order** equation. For low order we can shorten it to:

$$\ddot{x} + \gamma \dot{x} + \omega^2 x = w(t)$$

High order equations can (almost always) be converted to a **system of 1st order** equations:

$$x_1(t) := x(t), \quad x_2(t) := \dot{x}(t)$$
$$\dot{\mathbf{x}} = \begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -\omega^2 & -\gamma \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} w(t) = \mathbf{F}\mathbf{x} + \mathbf{L}w(t)$$

We exchange order for dimension.

In general  $\mathbf{F}$  and  $\mathbf{L}$  can depend on  $x$  and time. If they are constant or depend only on time, the system is **linear** (affine):

$$\dot{\mathbf{x}} = \mathbf{F}(t)\mathbf{x} + \mathbf{L}(t)w_x(t)$$

$$\dot{\mathbf{y}} = \mathbf{F}(t)\mathbf{y} + \mathbf{L}(t)w_y(t)$$

$$\mathbf{z} := \mathbf{x} + a\mathbf{y}$$

$$\begin{aligned}\dot{\mathbf{z}} &= \dot{\mathbf{x}} + a\dot{\mathbf{y}} = \mathbf{F}(t)\mathbf{x} + \mathbf{L}(t)w_x(t) + a\mathbf{F}(t)\mathbf{y} + a\mathbf{L}(t)w_y(t) = \\ &= \mathbf{F}(t)(\mathbf{x} + a\mathbf{y}) + \mathbf{L}(t)(w_x(t) + aw_y(t))\end{aligned}$$

The derivative of the combination follows the same equation with the combined inputs.

# 1st order homogenous ODE

Defines a 1-dimensional system:  $\dot{x} = Fx$ . Integrating both sides from 0 to  $t$ .

$$\int_0^t \dot{x}(\tau) d\tau = x(t) - x(0) \rightarrow x(t) = x(0) + \int_0^t Fx(\tau) d\tau$$

Recursively insert the formal solution in its expression:

$$\begin{aligned}x(t) &= x(0) + \int_0^t F \left[ x(0) + \int_0^{\tau_1} Fx(\tau_2) d\tau_2 \right] d\tau_1 = \\&= x(0) + Fx(0)t + \int_0^t \int_0^{\tau_1} F^2 \left[ x(0) + \int_0^{\tau_2} Fx(\tau_3) d\tau_3 \right] d\tau_2 d\tau_1 = \\&= x(0) + Fx(0)t + \frac{1}{2} F^2 x(0) t^2 + \int_0^t \int_0^{\tau_1} \int_0^{\tau_2} F^3 x(\tau_3) d\tau_3 d\tau_2 d\tau_1 = \dots \\x(t) &= x(0) + Fx(0)t + \frac{1}{2} F^2 x(0) t^2 + \frac{1}{6} F^3 x(0) t^3 + \dots = \\&= \left( 1 + Ft + \frac{F^2 t^2}{2!} + \frac{F^3 t^3}{3!} + \dots \right) x(0) = \exp(Ft)x(0)\end{aligned}$$

The infinite sum inside the parenthesis is the Taylor expansion of  $\exp(Ft)$ . Check it by expanding  $e^x$  around 0, i.e.

$$f(x) = \sum_{n=0}^{\infty} \frac{1}{n!} \left. \frac{d^n f}{dx^n} \right|_{x_0} (x - x_0)^n$$

is the Taylor expansion of  $f(x)$  around  $x_0$

$$= x(0) + Fx(0)t + \int_0^t \int_0^{\tau_1} F^2 x(\tau_2) d\tau_2 d\tau_1 =$$

# 1st order N-dimensional ODE

Homogenous

The homogenous case, defines a N-dimensional system:  $\dot{\mathbf{x}} = \mathbf{F}\mathbf{x}$ .

Following the same procedure as before

$$\mathbf{x}(t) = \left( \mathbf{I} + \mathbf{F}t + \frac{\mathbf{F}^2 t^2}{2!} + \frac{\mathbf{F}^3 t^3}{3!} + \dots \right) \mathbf{x}(0) := \exp(\mathbf{F}t)\mathbf{x}(0)$$

That's the definition of the **matrix exponential**

The matrix exponential is not the exponential of each entry in the matrix (element-wise exponential). Consider

$$\begin{aligned}\mathbf{F} &= \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \\ \mathbf{F}^2 &= \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} = \mathbf{0} \\ \rightarrow \mathbf{F}^n &= \mathbf{0}\end{aligned}$$

Then

$$\exp(\mathbf{F}t) = \mathbf{I} + \mathbf{F}t = \begin{bmatrix} 1 & t \\ 0 & 1 \end{bmatrix} \neq \begin{bmatrix} e^0 & e^t \\ e^0 & e^0 \end{bmatrix} = \begin{bmatrix} 1 & e^t \\ 1 & 1 \end{bmatrix}$$


# 1st order N-dimensional ODE

Inhomogenous

For  $\dot{\mathbf{x}} - \mathbf{F}\mathbf{x} = \mathbf{L}\mathbf{w}(t)$ . We get the solution between  $t_0$  and  $t$ :

$$\mathbf{x}(t) = \exp(\mathbf{F}(t - t_0))\mathbf{x}(0) + \int_{t_0}^t \exp(\mathbf{F}(t - \tau))\mathbf{L}\mathbf{w}(\tau)d\tau$$

In the scalar case this looks like



Multiply both sides by  $\exp(-\mathbf{F}t)$

$$\exp(-\mathbf{F}t)\dot{\mathbf{x}} - \exp(-\mathbf{F}t)\mathbf{F}\mathbf{x} = \exp(-\mathbf{F}t)\mathbf{L}\mathbf{w}(t)$$

You can prove from the definition of the matrix exponential the property

$$\frac{d}{dt} \exp(-\mathbf{F}t) = -\exp(-\mathbf{F}t)\mathbf{F}$$

Just like the scalar case, but here the order of multiplication is important. Observe that:

$$\frac{d}{dt} [\exp(-\mathbf{F}t)\mathbf{x}(t)] = \exp(-\mathbf{F}t)\dot{\mathbf{x}} - \exp(-\mathbf{F}t)\mathbf{F}\mathbf{x} = \exp(-\mathbf{F}t) (\dot{\mathbf{x}} - \mathbf{F}\mathbf{x})$$

Then the equation reduces to

$$\frac{d}{dt} [\exp(-\mathbf{F}t)\mathbf{x}(t)] = \exp(-\mathbf{F}t)\mathbf{L}(t)\mathbf{w}(t)$$

and by integrating both sides from  $t_0$  to  $t$  we get the solution. Note that

$$\int_{t_0}^t \frac{d}{d\tau} [\exp(-\mathbf{F}\tau)\mathbf{x}(\tau)]d\tau = \exp(-\mathbf{F}t)\mathbf{x}(t) - \exp(-\mathbf{F}t_0)\mathbf{x}(t_0)$$

For the general ODE

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t), t)$$

If we move from  $t$  to  $t + \Delta t$  we get the solution

$$\mathbf{x}(t + \Delta t) = \mathbf{x}(t) + \int_t^{t+\Delta t} \mathbf{f}(\mathbf{x}(\tau), \tau) d\tau$$

So the question is how to discretize the integral in the right-hand side. Choosing a quadrature defines an integration method. Euler method is obtained by the approximation


$$\int_t^{t+\Delta t} \mathbf{f}(\mathbf{x}(\tau), \tau) d\tau \approx \mathbf{f}(\mathbf{x}(t), t) \Delta t$$



# Discretization: example

Euler method 1D

A comparison of discretization methods. We will see some of the others later on. Clearly Euler is not the best.




source: *s\_euler\_exp.m*

# Discretization: example

Euler method 2D

$$\dot{\mathbf{x}} = \begin{bmatrix} 0 & 1 \\ -\omega^2 & -\gamma \end{bmatrix} \mathbf{x} \quad \dot{\mathbf{x}} = \begin{bmatrix} x_2 \\ -\omega^2 \sin(x_1) - \gamma x_2 \end{bmatrix}$$



source: *s\_euler\_exp.m*

# Discretization

## Linear time invariant (LTI) systems

When  $\mathbf{w}(t)$  is stochastic this is an instance of the Euler-Murayama method.

The system  $\dot{\mathbf{x}} - \mathbf{F}\mathbf{x} = \mathbf{L}\mathbf{w}(t)$  has an exact solution between  $t$  and  $t + \Delta t$


$$\mathbf{x}(t) = \exp(\mathbf{F}\Delta t)\mathbf{x}(0) + \exp(\mathbf{F}\Delta t) \int_t^{t+\Delta t} \exp(\mathbf{F}(t - \tau))\mathbf{L}\mathbf{w}(\tau)d\tau$$

So the approximation is only for the last integral. For example:

$$\hat{\mathbf{x}}(t_{k+1}) = \exp(\mathbf{F}\Delta t) [\hat{\mathbf{x}}(t_k) + \mathbf{L}\mathbf{w}(t_k)\Delta t]$$

# Discretization

Linear time invariant (LTI) systems




source: *s\_exp\_step\_input.m*

# Discretization

Linear time invariant (LTI) systems

This also works when  $w(t)$  is a random process.



source: *s\_exp\_step\_input.m*

When  $w(t)$  is stochastic this is an instance of the Euler-Murayama method. The stochastic case is subtle, e.g. the discretized noise has a variance that depends on the time step  $\Delta t$ . For details, refer to Simo Särkkä and Arno Solin (2019). Applied Stochastic Differential Equations.

There are stochastic methods that generalize the deterministic ones, e.g. Stochastic Runge-Kutta. For second order system check the Stochastic Verlet algorithm.

The data model is  $y_k = p(t_k, n) + \epsilon_k$ , where

$$p(t, n) := \sum_{s=0}^n a_s t^s$$

is a polynomial of degree  $n$ . Consider the system

$$\begin{aligned} \dot{x}_1 &= x_2 & \dot{x}_2 &= x_3 & \dot{x}_3 &= 0 \\ \begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \dot{x}_3 \end{bmatrix} &= \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \end{aligned}$$

if we start at  $\mathbf{x}(0) = [a_0 \quad a_1 \quad 2a_2]$  we get

$$x_3(t) = 2a_2 \quad x_2(t) = a_1 + 2a_2 t \quad x_1(t) = a_0 + a_1 t + a_2 t^2$$

we recover  $p(t, 2)$  on  $x_1(t)$ !

Can be generalized to any finite degree:

$$\dot{\mathbf{x}} = \mathbf{F}\mathbf{x}, \quad F_{ij} = \delta_{i(j+1)}, \quad 1 \leq i, j \leq n + 1$$
$$\mathbf{x}(0)^\top = [a_0 \quad a_1 \quad 2a_2 \quad \dots \quad n!a_n]$$

has  $p(t, n)$  in  $x_1(t)$ .

Then filter

$$\mathbf{A} = \exp(\mathbf{F}\Delta t), \quad \mathbf{Q} \neq \mathbf{0}$$
$$\mathbf{H} = [1 \quad 0 \quad \dots \quad 0], \quad \mathbf{R} \neq \mathbf{0}$$

implements polynomial regression with drift model in the coefficients.

See `s_polyreg.py`

Coffee break: 10 minutes



**Parameter estimation**

## What to estimate?

The linear Kalman filter is defined with the parameters

$\mathbf{A}$ ,  $\mathbf{H}$ ,  $\mathbf{Q}$ ,  $\mathbf{R}$ ,  $\mathbf{P}_0$ ,  $\mathbf{m}_0$ .

- $\mathbf{A}$ ,  $\mathbf{H}$  reflect knowledge of the physical system (unless data-driven).
- Bounds for the diagonal of  $\mathbf{R}$  can be estimated from the data.
- $\mathbf{Q}$  in general is hidden, unless derived from physical SDE.
- $\mathbf{P}_0$ ,  $\mathbf{m}_0$  can be critical for the performance of filters with long-term memories (eigenvalues of  $\mathbf{A}$  closer to or bigger than 1).

Idea: start by freezing degrees of freedom (à la stochastic gradient decent).

To estimate  $\mathbf{R}$  use any smoother on the data and compute the variance of the residuals between smoothed data and raw data. There is almost always a strong expectation on this parameter.

Many physics based models will have limited memory. An exception would be a fluid model where waves can carry information for long distances and induce long-term temporal correlations. In general systems with echo (transport, waves) can produce these effects.

Bayesian filter for state space models with parameter vector  $\theta$

$$\theta \sim p(\theta), \quad \mathbf{x}_0 \sim p(\mathbf{x}_0|\theta)$$

$$\mathbf{x}_k \sim p(\mathbf{x}_k|\mathbf{x}_{k-1}, \theta)$$

$$\mathbf{y}_k \sim p(\mathbf{y}_k|\mathbf{x}_k, \theta)$$

The idea is to compute the posterior distribution of states and parameters given the measurements

$$p(\mathbf{x}_{0:T}, \theta | \mathbf{y}_{1:T}) \propto p(\mathbf{y}_{1:T} | \mathbf{x}_{0:T}, \theta) p(\mathbf{x}_{0:T} | \theta) p(\theta)$$

where

$$p(\mathbf{y}_{1:T} | \mathbf{x}_{0:T}, \theta) = \prod_{k=1}^T p(\mathbf{y}_k | \mathbf{x}_k, \theta)$$

$$p(\mathbf{x}_{0:T} | \theta) = p(\mathbf{x}_0 | \theta) \prod_{k=1}^T p(\mathbf{x}_k | \mathbf{x}_{k-1}, \theta)$$

We focus on the parameters, so we "average" (marginalize) out the states.

Herein follow chapter 12 of the book Särkkä, S. (2013). Bayesian Filtering and Smoothing.

These are the same relations as before, we just made the parameters explicit.

$p(\theta)$  is the prior over the parameters.

Bayesian filter for state space models with parameter vector  $\theta$

$$\theta \sim p(\theta), \quad \mathbf{x}_0 \sim p(\mathbf{x}_0|\theta)$$

$$\mathbf{x}_k \sim p(\mathbf{x}_k|\mathbf{x}_{k-1}, \theta)$$

$$\mathbf{y}_k \sim p(\mathbf{y}_k|\mathbf{x}_k, \theta)$$

The posterior distribution of parameters given the measurements

$$p(\theta|\mathbf{y}_{1:T}) \propto p(\mathbf{y}_{1:T}|\theta)p(\theta)$$

where (check the "averaged" states)

$$p(\mathbf{y}_{1:T}|\theta) = \prod_{k=1}^T p(\mathbf{y}_k|\mathbf{y}_{1:k-1}, \theta)$$

$$p(\mathbf{y}_k|\mathbf{y}_{1:k-1}, \theta) = \int \overbrace{p(\mathbf{y}_k|\mathbf{x}_k, \theta) p(\mathbf{x}_k|\mathbf{y}_{1:k-1}, \theta)}^{\text{update}} d\mathbf{x}_k$$

measurement                      prediction

Herein follow chapter 12 of the book Särkkä, S. (2013). Bayesian Filtering and Smoothing.

For convenience of the notation we used

$$p(\mathbf{y}_1|\mathbf{y}_{1:0}, \theta) = p(\mathbf{y}_1|\theta)$$

To get the recursion for the posterior of the parameters, the key is to write a recursion for likelihood of the data. We do that next.

We know how to compute the posterior distribution of parameters given the measurements at a given step  $k$

$$p(\boldsymbol{\theta}|\mathbf{y}_{1:k}) \propto p(\mathbf{y}_{1:k}|\boldsymbol{\theta})p(\boldsymbol{\theta})$$

We define log-likelihood  $\phi_k(\boldsymbol{\theta})$ :

$$\begin{aligned}\exp(\phi_k(\boldsymbol{\theta})) &:= p(\mathbf{y}_{1:k}|\boldsymbol{\theta}) = \prod_{n=1}^k p(\mathbf{y}_n|\mathbf{y}_{1:n-1}, \boldsymbol{\theta}) \\ &\quad \underbrace{p(\mathbf{y}_{1:k-1}|\boldsymbol{\theta})}_{\exp(\phi_{k-1}(\boldsymbol{\theta}))} \\ &= p(\mathbf{y}_k|\mathbf{y}_{k-1}, \boldsymbol{\theta}) \prod_{n=1}^{k-1} p(\mathbf{y}_n|\mathbf{y}_{1:n-1}, \boldsymbol{\theta}) \\ &= p(\mathbf{y}_k|\mathbf{y}_{k-1}, \boldsymbol{\theta}) \exp(\phi_{k-1}(\boldsymbol{\theta}))\end{aligned}$$

Taking logs

$$\phi_k(\boldsymbol{\theta}) = \phi_{k-1}(\boldsymbol{\theta}) + \log [p(\mathbf{y}_k|\mathbf{y}_{k-1}, \boldsymbol{\theta})], \quad \phi_0(\boldsymbol{\theta}) = \log [p(\boldsymbol{\theta})]$$

Herein follow chapter 12 of the book Särkkä, S. (2013). Bayesian Filtering and Smoothing.

If the prior  $p(\boldsymbol{\theta})$  is uniform, this is the same as maximum likelihood (ML). If the distribution reflect our prior knowledge about the parameters, then it is maximum a posteriori probability (MAP).


$\phi_k(\boldsymbol{\theta})$  is the log-likelihood, sometimes the negative is defined.

Maximizing the log-likelihood gives optimal estimates. This should remind you of Viterbi, if you know it.

With the filters we have been studying, the log-likelihood has an explicit formula.

$$\mathbf{v}_k(\boldsymbol{\theta}) := \mathbf{y}_k - \mathbf{H}(\boldsymbol{\theta})\mathbf{m}_k(\boldsymbol{\theta})$$

$$\phi_k(\boldsymbol{\theta}) = \phi_{k-1}(\boldsymbol{\theta}) - \frac{1}{2} \log [2\pi \det \mathbf{S}_k(\boldsymbol{\theta})] - \frac{1}{2} \mathbf{v}_k^\top(\boldsymbol{\theta}) \mathbf{S}_k^{-1}(\boldsymbol{\theta}) \mathbf{v}_k(\boldsymbol{\theta})$$



To compute the likelihood we need to run over the whole data for each new value of  $\boldsymbol{\theta}$ . Each run over the whole data is done at fixed parameter value.

# Maximum likelihood

Example: oscillator

$$\mathbf{x}_k = \exp\left(\begin{bmatrix} 0 & 1 \\ -\omega^2 & 0 \end{bmatrix} \Delta t\right) \mathbf{x}_{k-1} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} w_{k-1}, \quad w_{k-1} \sim \mathcal{N}(0, q)$$
$$y_k = \begin{bmatrix} 1 & 0 \end{bmatrix} \mathbf{x}_k + r_k, \quad r_k \sim \mathcal{N}(0, \sigma_y^2)$$

`s_estimate_oscillator.py` does

- Line search for  $\omega$
- Optimization for  $\omega, q, \sigma_y, \mathbf{m}_0$

Gradient based methods will not be sensitive to the data if started far away from a local maximum: likelihood tends to be sharp and wiggly, vanishing gradients. Alternative: informative priors, e.g. tight bounds on the parameters.

If you want to optimize many parameters of the model, check the Expectation Maximization algorithm.

**Wrap-up**



- Linear Kalman filter with Gaussian distributions ✓
- Non-linear
  - Local linearization (Extended Kalman filter)
  - Gaussian approximation (Unscented Kalman filter, General Gaussian filters: GHKF, CKF)
  - Particle filters, ...
  - Brute force (naive MC, simulation of trajectories)
- Other noise models (check generalizations of KF, Särkkä's book references)
- Recursive or Batch?
  - Data: fixed size or incoming?
  - Estimation parameters: recorded data (batch) or online?