

BACKWARD AND INVERSE OPPORTUNITIES FOR PROBNUM IN MACHINE LEARNING

Philipp Hennig

Heilbronn Workshop – 27 March 2022

EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN




FACULTY OF SCIENCE
DEPARTMENT OF COMPUTER SCIENCE
CHAIR FOR THE METHODS OF MACHINE LEARNING

Many thanks to our supporter and sponsors:



CyberValley

DFG



The Men in Black

This talk contains work by many people, including many in the room today



Filip Tronarp




Nico Krämer



Nathanael Bosch



Jonathan Schmidt




Marvin Pförtner

A very 2021 inference task

Mixed Information Sources




Schmidt, Krämer, Hennig, 2021, NeurIPS 2021



Is this

- ▶ a machine learning task? (regress on $I(t)$)
doesn't work without mechanistic knowledge



Is this

- ▶ a machine learning task? (regress on $I(t)$)
doesn't work without mechanistic knowledge
- ▶ a simulation problem? (solve SIRD ODE)
we don't know β , though!
- ▶ an *inverse* problem (estimate β)
we really care about $I(t)$, though!


$$\frac{d}{dt} \begin{bmatrix} S(t) \\ I(t) \\ R(t) \\ V(t) \\ D(t) \end{bmatrix} = \begin{bmatrix} -\beta(t)S(t)I(t)/P - v(t) \\ \beta(t)S(t)I(t)/P - \gamma I(t) - \eta I(t) \\ \gamma I(t) \\ v(t) \\ \eta I(t) \end{bmatrix}$$

$$x'(t) = f(x(t), u(t)), \quad p(\mathbf{y} | x) = \prod_n^N \mathcal{N}(y_n; H(x(t_n)), \Sigma_n), \quad p(x) = \mathcal{GP}(m_x, k_x), \quad p(u) = \mathcal{GP}(m_u, k_u)$$

An *inverse problem* seems to be

- ▶ another word for an *inference* problem (inferring latent quantities from observations) (wikipedia).
- ▶ about inferring the object x in $y = D(x)$, where D is a known operator (here: the ODE integral) from data y .

In both cases, it seems the tough part, arguably, is the ill-posedness. But the data y is already probabilistic, too!



- Define some loss $L(u)$, e.g.

$$L(u) := \sum_i -\log p(y_i | \hat{x}(u)) = \sum_i (y_i - H(\hat{x}(u)))^2 + \text{const..}$$


- Compute the gradient $\nabla_u L(u^{(i)})$ with automatic differentiation. Examples:
 - `numppro` tutorial, using `jax's dopri5`
 - `Turing.jl` tutorial, using `diffeq.jl` solvers

Note how the user is discouraged from even thinking about the ODE solver.

The Probabilistic View on ODEs

[Schober, Duvenaud & P.H., 2014. Schober & P.H., 2016. Kersting & P.H., 2016, Tronarp, Kerstin, P.H., 2019, Bosch, Tronarp, P.H., 2021, ...]

$$x'(t) = f(x(t), t), \quad x(t_0) = x_0$$



`scipy.integrate.solve_ivp(f, t_span, x_0)`

\Rightarrow

`probnum.diffeq.probsolve_ivp(f, t_span, x_0)`

$$x'(t) = f(x(t), t), \quad x(t_0) = x_0$$

`scipy.integrate.solve_ivp(f, t_span, x_0)` \Rightarrow `probnum.diffeq.probsolve_ivp(f, t_span, x_0)`

$$x'(t) = f(x(t), t), \quad x(t_0) = x_0$$

`scipy.integrate.solve_ivp(f, t_span, x_0)` \Rightarrow `probnum.diffeq.probsolve_ivp(f, t_span, x_0)`

$$x'(t) = f(x(t), t), \quad x(t_0) = x_0$$


`scipy.integrate.solve_ivp(f, t_span, x_0)` \Rightarrow `probnum.diffeq.probsolve_ivp(f, t_span, x_0)`

$$x'(t) = f(x(t), t), \quad x(t_0) = x_0$$

`scipy.integrate.solve_ivp(f, t_span, x_0)` \Rightarrow `probnum.diffeq.probsolve_ivp(f, t_span, x_0)`

$$x'(t) = f(x(t), t), \quad x(t_0) = x_0$$

`scipy.integrate.solve_ivp(f, t_span, x_0)` \Rightarrow `probnum.diffeq.probsolve_ivp(f, t_span, x_0)`




$$z_0 | X(t_0) \sim \delta(x(t_0) - x_0) \quad z_m | X(t_m) \sim \delta(X^{(1)}(t_m) - f(X^{(0)}(t_m)))$$

- ▶ Use a **tractable** (linear Gaussian) *stochastic* differential equation as a prior for the **intractable** solution of the *nonlinear ordinary* differential equation

$$dX(t) = FX(t) dt + LdW(t) \quad \text{with} \quad X^{(i)}(t) = \frac{d^i}{dt^i} x(t), i = 1, \dots, \nu$$

- ▶ Consider *information operators* Z_i to link evaluations of the vector field f to x
- ▶ run the *extended Kalman filter (EKF)* to propagate uncertainty through f .



1 **procedure** EXTENDEDFILTER($m_{t-1}, P_{t-1}, A, Q, H, R, y$)

2 $m_t^- = Am_{t-1}$
3 $P_t^- = AP_{t-1}A^\top + Q$
4 $r = y - Hm_t^-$

5 $S = HP_t^-H^\top + R$
6 $K = P_t^-H^\top S^{-1}$

7 $m_t = m_t^- + Kr$

8 $P_t = (I - KH)P_t^-$

9 **return** $(m_t, P_t), (m_t^-, P_t^-)$

10 **end procedure**

// predictive mean with $A = \int \exp(F(\Delta t))$

// predictive covariance, with $Q = \int_0^{\Delta t} e^{F\tau} L L^\top e^{F\tau} d\tau$

// residual, with $y = 0, H = \partial h_{10} / \partial X|_{x=m_t^-}$

// innovation covariance

// gain

// updated mean


// updated covariance

Returning to our “Inverse Problem”



The real world is not described by an ODE, but regression alone doesn't help either

Schmidt, Krämer, Hennig, 2021, NeurIPS 2021




$$\frac{d}{dt} \begin{bmatrix} S(t) \\ I(t) \\ R(t) \\ V(t) \\ D(t) \end{bmatrix} = \begin{bmatrix} -\beta(t)S(t)I(t)/P - v(t) \\ \beta(t)S(t)I(t)/P - \gamma I(t) - \eta I(t) \\ \gamma I(t) \\ v(t) \\ \eta I(t) \end{bmatrix}$$

Not forward/inverse, but *mixed information*



blurring the boundaries of the black box

[Tronarp, Kersting, Särkkä, Hennig, 2019; Schmidt, Krämer, Hennig, , NeurIPS 2021]



to solve ODE $\frac{d}{dt}x(t) = f(x(t), t)$, model with SDE $dX(t) = F_X X(t) dt + L_X dW_X(t)$ and observation model (information operator)


$$Z_m | X(t_m) \sim \delta(E_X^{(1)} X(t_m) - f(E_X^{(0)} X(t_m)))$$

Not forward/inverse, but *mixed information*



blurring the boundaries of the black box

[Tronarp, Kersting, Särkkä, Hennig, 2019; Schmidt, Krämer, Hennig, , NeurIPS 2021]



natively (within same “forward” solve) combine with physical observations of the trajectory

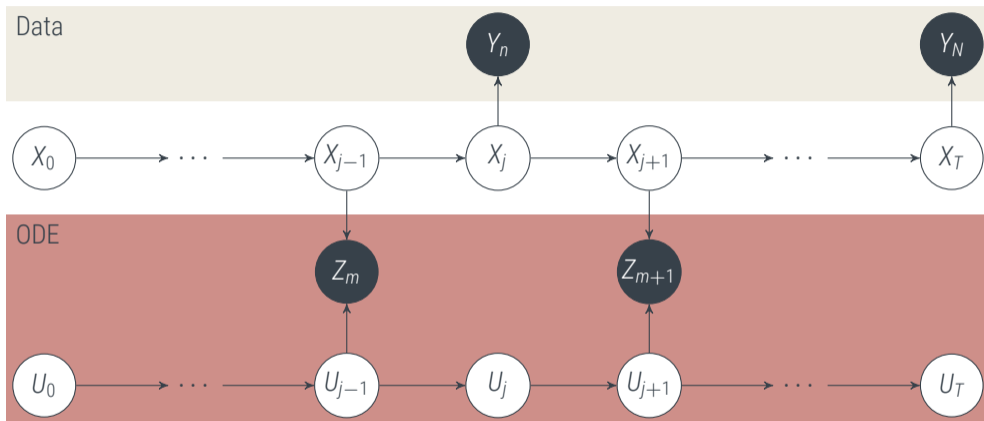
$$Y_n | X(t_n) \sim \mathcal{N}(HE_X^0 X(t_n), R)$$

Not forward/inverse, but *mixed information*



blurring the boundaries of the black box

[Tronarp, Kersting, Särkkä, Hennig, 2019; Schmidt, Krämer, Hennig, , NeurIPS 2021]



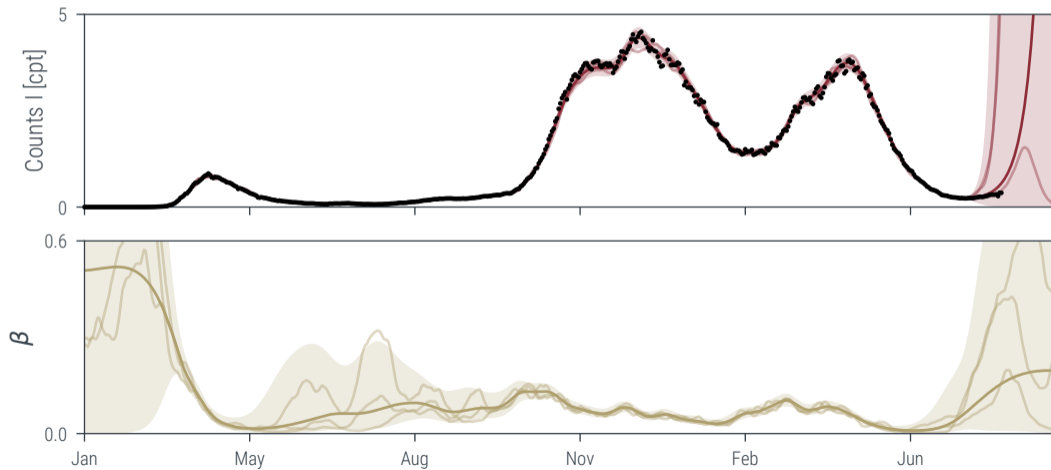
propagate uncertainty about ODE (e.g. from a latent force U) through the extended Kalman filter to solve $\frac{d}{dt}x(t) = f(x(t), u(t), t)$ with $dU(t) = F_U U(t) dt + L_U dW_U(t)$.




No more black box ODE solvers

Example: Covid modelling



Schmidt, Krämer, Hennig, 2021, NeurIPS 2021





Description	Equation	Information operator
First-order ODE	$\dot{y}(t) = f(y(t), t)$	$z(t, Y) := Y^{(1)} - f(Y^{(0)}, t)$
Second-order ODE	$\ddot{y}(t) = f(\dot{y}(t), y(t), t)$	$z(t, Y) := Y^{(2)} - f(Y^{(1)}, Y^{(0)}, t)$
Mass matrix DAE	$M\dot{y}(t) = f(y(t), t)$	$z(t, Y) := MY^{(1)} - f(Y^{(0)}, t)$
Invariances	$g(y(t), \dot{y}(t)) = 0$	$z(t, Y) := g(Y^{(0)}, Y^{(1)})$
Chain rule	$\ddot{y}(t) = J_f(y(t)) \cdot \dot{y}(t)$	$z(t, Y) := Y^{(2)} - J_f(Y^{(0)}) \cdot Y^{(1)}$

Infer the parameters θ of IVP ξ_θ measured with Gaussian noise at solution x


$$\frac{d}{dt}\xi_\theta(t) = f_\theta(\xi_\theta(t)), \quad \phi_\theta(0) = x_0, \quad p(\mathbf{y} | x) = \prod_i \mathcal{N}(y_i; H^\top x, R_\theta)$$

- ▶ We'd like to compute the marginal

$$p(\mathbf{y} | \theta) = \int p(\mathbf{y} | x) \delta(x - \xi_\theta) dx$$

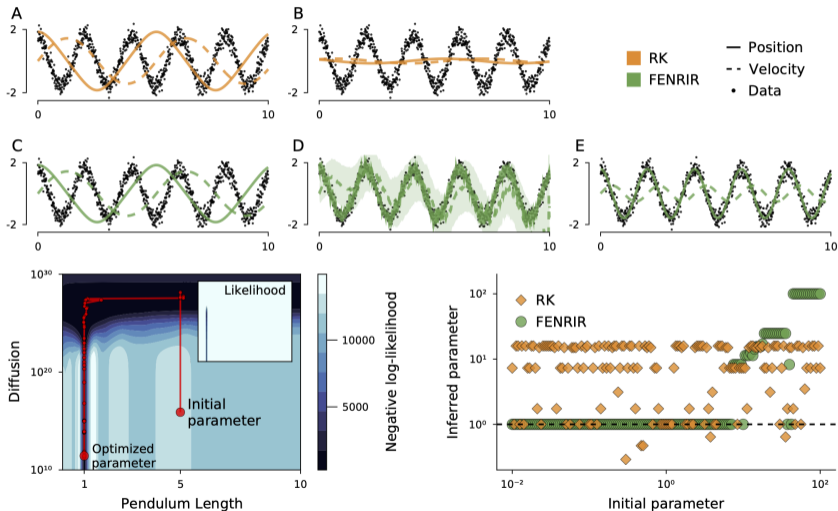
- ▶ Approximate δ with a Gaussian

$$\hat{p}(\mathbf{y} | \theta) = \int p(\mathbf{y} | x) \hat{\delta}_N(x - \xi_\theta) dx$$



Prior Hyperparameters as Regularizers


Tronarp, Bosch, Hennig. *Fenrir: Physics-Enhanced Regression for Initial Value Problems*





Summary

- ▶ *Propagation of Uncertainty* is great, but should not mislead us to keep the rigid structure of classical code
- ▶ instead, sometimes, *information* (the opposite of uncertainty) shouldn't be propagated, but *combined* efficiently
- ▶ because Probnum methods can deal with imprecise quantities natively, changing the order of the computation does not pose a conceptual problem for them. (That doesn't mean changing the order is always a good idea. But it's also not necessarily a bad idea).
- ▶ doing so can break the (artificial) separation between forward and inverse problems.

Re-casting computation as inference allows genuinely new, valuable functionality.


<http://mml.inf.uni-tuebingen.de>


Probabilistic Numerics – Computation as Machine Learning. P. Hennig, H. Kersting, M.A. Osborne, CUP, 2022


<https://www.youtube.com/c/TuebingenML>


[@PhilippHennig5](https://twitter.com/PhilippHennig5)

High-Dimensional ODEs/PDEs



Factorization assumptions allow scaling to millions of dimensions

Krämer, Bosch, Schmidt, PH, arXiv 2110.11812

