Data-assimilation meets automatic differentiation for identification of dynamical systems from irregularly-sampled, noisy data

Authors:

- Matthew Levine, Eric and Wendy Schmidt Center & Broad Institute of MIT and Harvard (`levinema@broadinstitute.org`)
- Iñigo Urteaga, Basque Center for Applied Mathematics & Ikerbasque, Basque Foundation for Science (`iurteaga@bcamath.org`)

Abstract:

We present advances on leveraging automatic differentiation —computer-based evaluations, via repeated application of the chain rule, of partial derivatives of software defined functions— for learning and predicting continuous (dynamics), discrete (observation) dynamical systems that underpin real-world messy time-series data.

We study (unknown) stochastic dynamical systems $\dot{x} = f(x,t) + L(x,t)\dot{w}$, where $x \in \mathbb{R}^{d_x}$, $x(0) = x_0 \sim \mathcal{N}(\mu_0, \Sigma_0)$, f is a state and/or time-dependent drift function, L a possibly state and/or time-dependent diffusion coefficient, and $\dot{w}$ the derivative of a d_x-dimensional Brownian motion with covariance Q. Data are observed at *arbitrary times* $\{t_k\}_{k=1}^{K}$ collected via a *noisy measurement* process $y(t) = h\big(x(t)\big) + \eta(t)$, where $h : \mathbb{R}^{d_x} \mapsto \mathbb{R}^{d_y}$ and $\eta(t) \sim \mathcal{N}(0, \Sigma_\eta)$. We denote the collection of all parameters as $\theta = \{f, L, \mu_0, \Sigma_0, Q, h, \Sigma_\eta\}$. Given a sequence of irregularly sampled and noisy observations $Y_K = [y(t_1), \dots, y(t_K)]$, we wish to (i) filter —estimate $p(x(t_K)|Y_K, \theta)$, (ii) smooth —estimate $p(\{x(t)\}_t|Y_K, \theta)$ (iii) predict —estimate $p(x(t > t_K)|Y_K, \theta)$, and (iv) infer parameters —estimate $p(\theta|Y_K)$, for systems with linear and non-linear unknown functions f and h.

We merge machine learning tools (i.e., automatic differentiation) with state-of-the-art data-assimilation to solve all these interconnected Bayesian inference problems [1]. We devise a framework that allows for differentiation through filtering/smoothing algorithms [2] and the SDE solver. By virtue of this novel synergy, we enable usage of modern optimization and inference techniques (e.g., stochastic gradient descent, Hamiltonian Monte Carlo) for learning and parameter inference of continuous-time dynamics. Our work opens up novel research directions on uncertainty quantification and the combination of mechanistic and machine-learning models for improved dynamical system identification from irregularly-sampled, noisy data.

References:

[1] Simo Särkkä and Lennart Svensson (2023). Bayesian Filtering and Smoothing. Second Edition. Cambridge University Press.

[2] Yuming Chen, Daniel Sanz-Alonso, and Rebecca Willett. “Autodifferentiable ensemble Kalman filters.” SIAM Journal on Mathematics of Data Science 4, no. 2 (2022): 801-833.