Appendix: Inverse articulated-body dynamics from video via variational sequential Monte Carlo

Dan Biderman[1] **Christian A. Naesseth**[1] **Luhuan Wu**[1] **Taiga Abe**[1]

Alice C. Mosberger[1] **Leslie J. Sibener**[1] **Rui Costa**[1]

James Murray[2] **John P. Cunningham**[1]

[1] Columbia University, New York, USA
`{db3236, ca2733, lw2827, ta2507, acm2246`
`ljs2203, rc3031, jpc2181}@columbia.edu`
[2] University of Oregon, Oregon, USA
`jmurray9@uoregon.edu`

1 Sequential Monte Carlo review

The weighted set of particles in SMC are obtained iteratively for $t = 0, \dots, T$. For $t = 0$ we sample particles from a proposal distribution $\mathbf{x}_0^i \sim q(\mathbf{x}_0)$ and compute weights $w_0^i = \frac{f(\mathbf{x}_0^i)g(\mathbf{y}_0|\mathbf{x}_0^i)}{q(\mathbf{x}_0^i)}$. For $t > 0$, we first *resample* "ancestral indices" $a_{t-1}^i \in \{1, .., N\}$, i.e., we choose particles according to their importance weights w_{t-1}^i. Next, we propose new states $\mathbf{x}_t^i \sim q(\mathbf{x}_t^i|\mathbf{x}_{t-1}^{a_{t-1}^i})$, append them to the previous resampled states, and compute updated importance weights $w_t^i = \frac{f(\mathbf{x}_t^i|\mathbf{x}_t^{a_{t-1}^i})g(\mathbf{y}_t|\mathbf{x}_t^i)}{q(\mathbf{x}_t^i|\mathbf{x}_t^{a_{t-1}^i})}$. The proposal distribution is an important design choice. One common choice is to set $q(\mathbf{x}_t|\mathbf{x}_{t-1}) = f(\mathbf{x}_t|\mathbf{x}_{t-1})$, known as the Bootstrap Particle Filter [2].

1.1 Sequential Monte Carlo pseudocode

Sequential Monte Carlo[5, 6, 3]
Require: proposal distribution $q(\mathbf{x}_t|\mathbf{x}_{t-1})$ from which we could sample and that can be evaluated
Require: transition distribution $f(\mathbf{x}_t|\mathbf{x}_{t-1})$ [our rigid-body dynamics] that can be evaluated
Require: emission distribution: $g(\mathbf{y}_t|\mathbf{x}_t)$ [our forward kinematics] that can be evaluated
  for t in $0 : T$ **do**
    if t=0 **then**
      propose initial states $\mathbf{x}_0^i \sim q(\mathbf{x}_0)$
      compute importance weights $w_0^i = \frac{f(\mathbf{x}_0^i)g(\mathbf{y}_0|\mathbf{x}_0^i)}{q(\mathbf{x}_0^i)}$
    else if $t \geq 0$ **then**
      resample indices $a_{t-1}^i \sim \text{Categorical}(w_{t-1}^i / \sum_l w_{t-1}^l)$
      propose states $\mathbf{x}_t^i \sim q(\mathbf{x}_t^i|\mathbf{x}_{t-1}^{a_{t-1}^i})$
      append state $\mathbf{x}_{1:t}^i = \{\mathbf{x}_{1:t-1}^{a_{t-1}^i}, \mathbf{x}_t^i\}$
      compute importance weights $w_t^i = \frac{f(\mathbf{x}_t^i|\mathbf{x}_t^{a_{t-1}^i})g(\mathbf{y}_t|\mathbf{x}_t^i)}{q(\mathbf{x}_t^i|\mathbf{x}_t^{a_{t-1}^i})}$

Workshop on Differentiable Vision, Graphics, and Physics in Machine Learning at NeurIPS 2020.

end if
end for
sample last index $b_T^i \sim \text{Categorical}(w_T^i / \sum_l w_T^l)$
return trajectory $\mathbf{x}_{1:T}^{b_T}$ =0

1.2 Multiple importance-weighted SMC samplers

To improve torque inference, we propose the following algorithm. We run M SMC samplers, and for each sampler in $m = 1, \dots, M$ we compute the expectation for the test function $h(\cdot)$, e.g., posterior mean or variance, over all particles. Then, we compute an outer-level expectation $\widehat{h}_{\text{agg}}$ over all SMC samplers,

$$\widehat{h}_{\text{agg}} = \sum_{i=1}^{M} \frac{\widehat{p}(\boldsymbol{y}_{0:T})[\boldsymbol{u}^m]}{\sum_{n=1}^{M} \widehat{p}(\boldsymbol{y}_{0:T})[\boldsymbol{u}^n]} \widehat{h}(\boldsymbol{u}^m), \quad \widehat{h}(\boldsymbol{u}^m) := \sum_{i=1}^{N} w_T^{m,i} h(\mathbf{x}_{0:T}^{m,i}), \tag{1}$$

where $\boldsymbol{u}^m$ denotes all the random variables generated in SMC sampler m, and $\widehat{p}(\boldsymbol{y}_{0:T})[\boldsymbol{u}^m]$ its log-marginal likelihood estimate. To this end, we have implemented a GPU-supported parallel SMC sampling in PyTorch. Our approach could be seen as a variant of the nested SMC method [4]. Moreover, for each sampler, we run Forward-Filtering Backward-Sampling [1].

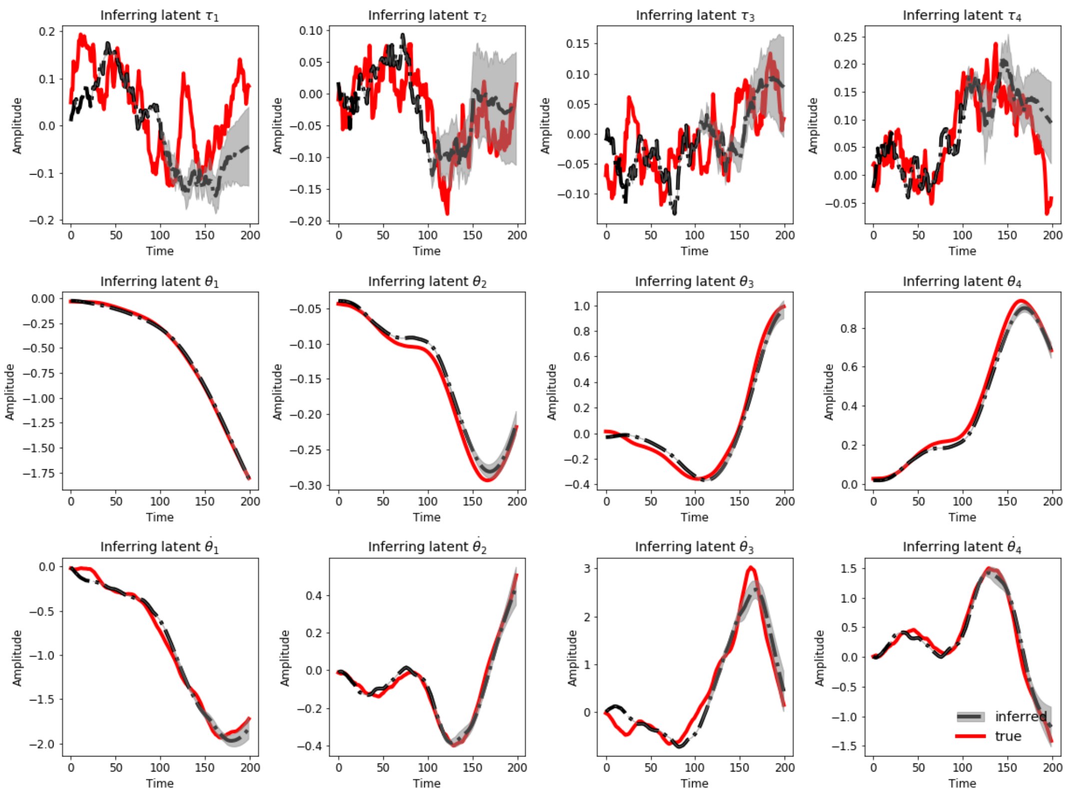

Figure 1: Successful state inference for the 3D arm model. Conventions as in Figure 2 in the main text except that each joint has two axes of rotation.

References

[1] A. Doucet and A. M. Johansen. A tutorial on particle filtering and smoothing: Fifteen years later.

[2] N. J. Gordon, D. J. Salmond, and A. F. Smith. Novel approach to nonlinear/non-gaussian bayesian state estimation. In *IEE proceedings F (radar and signal processing)*, volume 140, pages 107–113. IET, 1993.

[3] G. Kitagawa. Monte carlo filter and smoother for non-gaussian nonlinear state space models. *Journal of computational and graphical statistics*, 5(1):1–25, 1996.

[4] C. Naesseth, F. Lindsten, and T. Schön. Nested sequential Monte Carlo methods. volume 37 of *Proceedings of Machine Learning Research*, pages 1292–1301, Lille, France, 07–09 Jul 2015. PMLR.

[5] C. Naesseth, S. Linderman, R. Ranganath, and D. Blei. Variational sequential monte carlo. In *International Conference on Artificial Intelligence and Statistics*, pages 968–977. PMLR, 2018.

[6] L. Stewart and P. McCarty Jr. Use of bayesian belief networks to fuse continuous and discrete information for target recognition, tracking, and situation assessment. In *Signal Processing, Sensor Fusion, and Target Recognition*, volume 1699, pages 177–185. International Society for Optics and Photonics, 1992.