Learning low-dimensional Lagrangian models of high-dimensional trajectories via Fermat Components Analysis

Paul Vernaza and Daniel D. Lee
GRASP Laboratory
University of Pennsylvania, Philadelphia, PA 19104
{vernaza,ddlee}@seas.upenn.edu

We consider the problem of learning succinct models that adequately describe high-dimensional physical systems from observed trajectory data, and how these models might be used to solve associated high-dimensional learning problems.

One such problem is a high-dimensional regression problem we refer to as *learning interpolation*. Suppose we are given a training set of m sampled trajectories $\{x_i(t_j) \in \mathbb{R}^N \mid i \in \{1 \dots m\}, j \in \{1 \dots T\}\}$, where we assume that the trajectories are generated by some physical system and that N is large. Given just the endpoints $z(t_1)$, $z(t_T)$ of an unknown trajectory $z(t)$, we would like to infer a plausible sequence $\{z_2, z_3, \dots, z_{T-1}\} \in \mathbb{R}^{N \times (T-2)}$. This can be viewed as a regression problem from $\mathbb{R}^{2N}$ to $\mathbb{R}^{N \times (T-2)}$. If we are to solve this problem for $N = 1000$ and $T = 100$, for example, a naive regression approach is clearly hopeless.

We show how this problem can be solved effectively by building a low-dimensional model of the system's Lagrangian dynamics from which we can efficiently extract the answer to such problems as the learning interpolation problem. The existence of such a model is motivated by the presence of *conservation laws* in real physical systems that vastly simplify their analysis. Our method finds a low-dimensional basis in which the given trajectories appear to satisfy as many conservation laws as possible, which vastly simplifies the problem of learning the Lagrangian. We refer to this procedure as Fermat Components Analysis.

Once an estimate of the Lagrangian is obtained, finding the solution to the interpolation problem can be formulated as a high-dimensional minimum-cost path problem. Remarkably, we show that it is possible to efficiently find a path that *globally optimizes* this high-dimensional path planning problem, with a complexity that depends only on the dimensionality of the low-dimensional basis that characterizes the Lagrangian. In fact, we can show that it is possible to directly compute the high-dimensional *value function* thanks to a symmetry of the value function that results from the assumption of low-dimensional structure in the Lagrangian. This itself is an interesting result that has many potential applications to optimal path planning problems in robotics and control.

Our approach differs from previous approaches aimed at similar problems principally in our assumption of a physical origin for the trajectory data. This stands in contrast to methods such as Locally Weighted Projection Regression [1] and Gaussian Process Dynamical Models [2], which do not take advantage of the special structure induced by this physical assumption.

We have applied our techniques to the problem of finding plausible interpolations of human motion capture data from key frames, as shown in Figure 1. The trajectories in this problem are 990-dimensional, corresponding to the three-dimensional positions of 330 markers tracked by the motion

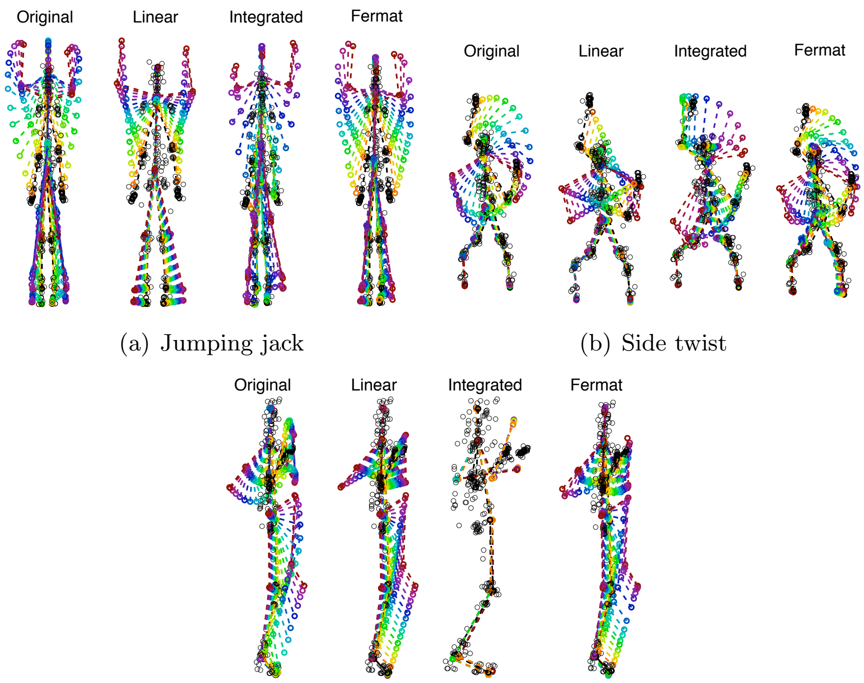

(a) Jumping jack

(b) Side twist

(c) Knee touch

Figure 1: Visualization of experiments in producing novel trajectories from pairs of key frames. Leftmost image in each set is the true, held-out trajectory. Rightmost image is the trajectory reconstructed from the key frames by our method. Other images depict the output of different learning methods.

capture system. Our method exhibits better qualitative and quantitative results than baseline methods, consisting of linear interpolation and kernel-based regression of the system dynamics.

Finally, we give preliminary results showing the application of our method to the classification problem of labeling previously-unseen motion sequences given labeled examples.

References

[1] Sethu Vijayakumar and Stefan Schaal. LWPR : An O(n) algorithm for incremental real time learning in high dimensional space. In *Proc. of Seventeenth International Conference on Machine Learning*, 2000.

[2] Jack Wang, David Fleet, and Aaron Hertzmann. Gaussian process dynamical models. In Y. Weiss, B. Schölkopf, and J. Platt, editors, *Advances in Neural Information Processing Systems 18*, pages 1441–1448. MIT Press, Cambridge, MA, 2006.

Topic: Estimation, prediction, and sequence modeling
Preference: oral