

Learning Lagrangian Dynamics

Clancy Rowley

Joint work with
Mohammadreza Ahmadi and Ufuk Topcu (UT Austin)



Overview

- ▶ Many mechanical systems are described by Lagrangian dynamics
- ▶ Properties of the Lagrangian (e.g., symmetries) lead to physical properties of the system (e.g., conservation laws)

Overview

- ▶ Many mechanical systems are described by Lagrangian dynamics
- ▶ Properties of the Lagrangian (e.g., symmetries) lead to physical properties of the system (e.g., conservation laws)

Idea

- ▶ Instead of learning the dynamics directly, learn the Lagrangian
- ▶ Known symmetries (a type of side information) can be incorporated, and the models obtained will therefore respect conservation laws (e.g., energy, momentum, angular momentum)

Lagrangian dynamics

- ▶ Consider a $2n$ dimensional state space with variables $q = (q_1, \dots, q_n)$ (“positions”) and $\dot{q} = (\dot{q}_1, \dots, \dot{q}_n)$ (“velocities”).
- ▶ Define a real-valued function $L(q, \dot{q})$.
- ▶ The corresponding dynamics are given by the Euler-Lagrange equations:

$$\frac{d}{dt} \left(\frac{\partial L}{\partial \dot{q}_i} \right) - \frac{\partial L}{\partial q_i} = 0, \quad i = 1, \dots, n.$$

Lagrangian dynamics

- ▶ Consider a $2n$ dimensional state space with variables $q = (q_1, \dots, q_n)$ (“positions”) and $\dot{q} = (\dot{q}_1, \dots, \dot{q}_n)$ (“velocities”).
- ▶ Define a real-valued function $L(q, \dot{q})$.
- ▶ The corresponding dynamics are given by the Euler-Lagrange equations:

$$\frac{d}{dt} \left(\frac{\partial L}{\partial \dot{q}_i} \right) - \frac{\partial L}{\partial q_i} = 0, \quad i = 1, \dots, n.$$

- ▶ For mechanical systems, the Lagrangian is kinetic energy minus potential energy:

$$L(q, \dot{q}) = \frac{1}{2} \dot{q}^T M \dot{q} - V(q)$$

Lagrangian dynamics

- ▶ Consider a $2n$ dimensional state space with variables $q = (q_1, \dots, q_n)$ (“positions”) and $\dot{q} = (\dot{q}_1, \dots, \dot{q}_n)$ (“velocities”).
- ▶ Define a real-valued function $L(q, \dot{q})$.
- ▶ The corresponding dynamics are given by the Euler-Lagrange equations:

$$\frac{d}{dt} \left(\frac{\partial L}{\partial \dot{q}_i} \right) - \frac{\partial L}{\partial q_i} = 0, \quad i = 1, \dots, n.$$

- ▶ For mechanical systems, the Lagrangian is kinetic energy minus potential energy:

$$L(q, \dot{q}) = \frac{1}{2} \dot{q}^T M \dot{q} - V(q)$$

- ▶ The corresponding equations are equivalent to Newton's equations:

$$\frac{d}{dt} M \dot{q} = -\nabla V(q),$$

where the force is $-\nabla V$.

Learning Lagrangians versus learning dynamics

Suppose the dynamics of a system are given by

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

where $x = (q, \dot{q})$ is the state.

- ▶ One approach to learning is to use data (e.g., samples of $x(t_i)$) to try to find the function $f : \mathbb{R}^{2n} \rightarrow \mathbb{R}^{2n}$.

Learning Lagrangians versus learning dynamics

Suppose the dynamics of a system are given by

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

where $x = (q, \dot{q})$ is the state.

- ▶ One approach to learning is to use data (e.g., samples of $x(t_i)$) to try to find the function $f : \mathbb{R}^{2n} \rightarrow \mathbb{R}^{2n}$.
- ▶ If the dynamics come from a Lagrangian $L(q, \dot{q})$, then a different approach is to use data to try to find the Lagrangian $L : \mathbb{R}^{2n} \rightarrow \mathbb{R}$.

Learning Lagrangians versus learning dynamics

Suppose the dynamics of a system are given by

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

where $x = (q, \dot{q})$ is the state.

- ▶ One approach to learning is to use data (e.g., samples of $x(t_i)$) to try to find the function $f : \mathbb{R}^{2n} \rightarrow \mathbb{R}^{2n}$.
- ▶ If the dynamics come from a Lagrangian $L(q, \dot{q})$, then a different approach is to use data to try to find the Lagrangian $L : \mathbb{R}^{2n} \rightarrow \mathbb{R}$.
- ▶ Here we take the latter approach.

Why?

Why would we want to do this?

Symmetries

- ▶ Symmetries are a type of side information. For instance, the Lagrangian may be independent of a particular variable q_j , or invariant to the action of some group (e.g., rotations).
- ▶ For instance, suppose $L(q, \dot{q})$ is independent of position q_j (for some j).
- ▶ Then the Euler-Lagrange equation for that variable is

$$\frac{d}{dt} \frac{\partial L}{\partial \dot{q}_j} = 0,$$

so the “generalized momentum” $\partial L / \partial \dot{q}_j$ is conserved.

Symmetries

- ▶ Symmetries are a type of side information. For instance, the Lagrangian may be independent of a particular variable q_j , or invariant to the action of some group (e.g., rotations).
- ▶ For instance, suppose $L(q, \dot{q})$ is independent of position q_j (for some j).
- ▶ Then the Euler-Lagrange equation for that variable is

$$\frac{d}{dt} \frac{\partial L}{\partial \dot{q}_j} = 0,$$

so the “generalized momentum” $\partial L / \partial \dot{q}_j$ is conserved.

- ▶ Emmy Noether showed that a similar result holds for symmetries described by arbitrary Lie groups.

Symmetries

- ▶ Symmetries are a type of side information. For instance, the Lagrangian may be independent of a particular variable q_j , or invariant to the action of some group (e.g., rotations).
- ▶ For instance, suppose $L(q, \dot{q})$ is independent of position q_j (for some j).
- ▶ Then the Euler-Lagrange equation for that variable is

$$\frac{d}{dt} \frac{\partial L}{\partial \dot{q}_j} = 0,$$

so the “generalized momentum” $\partial L / \partial \dot{q}_j$ is conserved.

- ▶ Emmy Noether showed that a similar result holds for symmetries described by arbitrary Lie groups.

Idea

- ▶ If there are known symmetries, incorporate these as “side information”.
- ▶ The resulting models will then automatically satisfy the corresponding conservation laws.


Problem setup for learning Lagrangians

We will assume the dynamics satisfy Euler-Lagrange equations of the form

$$\frac{d}{dt} \left(\frac{\partial L}{\partial \dot{q}} \right) - \frac{\partial L}{\partial q} = B(q)u.$$


- ▶ $u(t)$ is a control input, which we get to choose
- ▶ $B(q)$ is known
- ▶ The function $L(q, \dot{q})$ is **unknown**, and we will try to determine it from data.

Learning Lagrangians from data



- ▶ The goal is to learn $L(q, \dot{q})$ given data $\{t_i, u(t_i), q(t_i), \dot{q}(t_i), \ddot{q}(t_i)\}_{i=1}^N$.

Learning Lagrangians from data




- ▶ The goal is to learn $L(q, \dot{q})$ given data $\{t_i, u(t_i), q(t_i), \dot{q}(t_i), \ddot{q}(t_i)\}_{i=1}^N$.
- ▶ Parameterize the Lagrangian:

$$L(q, \dot{q}) = \sum_{i=1}^d \alpha_i \varphi_i(q, \dot{q}),$$

for given basis functions φ_i (e.g., monomials)

Learning Lagrangians from data



- ▶ The goal is to learn $L(q, \dot{q})$ given data $\{t_i, u(t_i), q(t_i), \dot{q}(t_i), \ddot{q}(t_i)\}_{i=1}^N$.
- ▶ Parameterize the Lagrangian:

$$L(q, \dot{q}) = \sum_{i=1}^d \alpha_i \varphi_i(q, \dot{q}),$$

for given basis functions φ_i (e.g., monomials)

- ▶ Find constants α_i that best fit the data.
- ▶ One can then predict the future evolution from E-L equations.

Two warnings

1. Different Lagrangians can lead to the same dynamics
2. If $L(q, \dot{q}) = 0$, the Euler-Lagrange equations are trivially satisfied for any data.


Two warnings

1. Different Lagrangians can lead to the same dynamics
2. If $L(q, \dot{q}) = 0$, the Euler-Lagrange equations are trivially satisfied for any data.



We take steps to find a particular (non-trivial) Lagrangian whose corresponding dynamics agree with the data.

For details, see paper (M. Ahmadi, U. Topcu, C. Rowley, ACC 2018).


Example: forced Duffing equation





Results: forced Duffing equation





Example: Acrobot





Results: Acrobot



Results: Acrobot



Results: Acrobot



Summary and future work

- ▶ Instead of learning dynamics directly ($\dot{x} = f(x)$), we learn a Lagrangian $L(q, \dot{q})$.
- ▶ This provides a way to incorporate physics into learning.
- ▶ For instance, known **symmetries** (a type of side information) lead to **conservation laws**.
- ▶ The method works for some simple examples of mechanical systems, for which it is known how the control input enters.

Summary and future work

- ▶ Instead of learning dynamics directly ($\dot{x} = f(x)$), we learn a Lagrangian $L(q, \dot{q})$.
- ▶ This provides a way to incorporate physics into learning.
- ▶ For instance, known **symmetries** (a type of side information) lead to **conservation laws**.
- ▶ The method works for some simple examples of mechanical systems, for which it is known how the control input enters.

- ▶ Ufuk and his team have a way of quantifying uncertainty, and using barrier certificates for safety verification and safe controller synthesis thanks to sum-of-squares programming.

Summary and future work

- ▶ Instead of learning dynamics directly ($\dot{x} = f(x)$), we learn a Lagrangian $L(q, \dot{q})$.
- ▶ This provides a way to incorporate physics into learning.
- ▶ For instance, known **symmetries** (a type of side information) lead to **conservation laws**.
- ▶ The method works for some simple examples of mechanical systems, for which it is known how the control input enters.

- ▶ Ufuk and his team have a way of quantifying uncertainty, and using barrier certificates for safety verification and safe controller synthesis thanks to sum-of-squares programming.
- ▶ In the future, we want to try these methods on more complex examples, with known symmetries, for which it is not known how the control enters.