Learning to Control PDEs with Differentiable Physics

Seminar – Deep Learning in Physics

TUM, Summer Semester 2022

Barnabás Börcsök Advisor: Lukas Prantl

Overview

- Differentiable Physics
 - Motivation
 - Forward Physics Simulation
 - Differentiable ("backwards") Physics
 - Examples
- Partial Differential Equations (PDEs)
- Results of the paper
- Other examples and outlook
 - Enhancing simulations with Neural Networks (NNs)
 - Inductive Physical Bias

Overview

- Differentiable Physics
 - Motivation
 - Forward Physics Simulation
 - Differentiable ("backwards") Physics
 - Examples
- Partial Differential Equations (PDEs)
- Results of the paper
- Other examples and outlook
 - Enhancing simulations with Neural Networks (NNs)
 - Inductive Physical Bias



- Modelling
- Simulating ... the real world
- Rendering









- Modelling
- Simulating
- Rendering



- Modelling
- Simulating
- Rendering



- Modelling
- Simulating
- Rendering







- Modelling
- Simulating
- Rendering

 $(x_{i+1}, v_{i+1}) = f(x_i, v_i)$



- Modelling
- Simulating
- Rendering

 $(x_{i+1}, v_{i+1}) = f(x_i, v_i)$



- Modelling
- Simulating
- Rendering



The **observation** should be the same.







(Forward) Physics Simulation

Physical simulations are much more complex ...

 $(x_{i+1}, v_{i+1}) = f(x_i, v_i)$





Our Model

Artistic Simulation of Curly Hair, Hayley et. al

Each strand is made up of springs (Hooke's law) Simulating 579 hair strands (44,552 points)



Source: https://github.com/taichi-dev/difftaichi

Elastic Object + Fluid



Source: https://github.com/taichi-dev/difftaichi



2D Indirect Fluid Control

Source:

Learning to Control PDEs with Differentiable Physics, Holl et. al

 $(x_{i+1}, v_{i+1}) = f(x_i, v_i)$



 $L(x_n) = |x^* - x_n|_2^2$ "Loss function"



 $L(x_n) = |x^* - x_n|_2^2$ "Loss function"

Let's compute $\frac{\partial L}{\partial x_0}$ and $\frac{\partial L}{\partial v_0}$.



 $L(x_n) = |x^* - x_n|_2^2$ "Loss function"

Let's compute $\frac{\partial L}{\partial x_0}$ and $\frac{\partial L}{\partial v_0}$.

Using the chain rule, a.k.a. "backpropagation"



 $L(x_n) = |x^* - x_n|_2^2$ "Loss function"

Let's compute $\frac{\partial L}{\partial x_0}$ and $\frac{\partial L}{\partial v_0}$.

Using the chain rule, a.k.a. "backpropagation"







Gradient descent iteration 0 ...

... and iteration 100

Source: https://github.com/taichi-dev/difftaichi#differentiable-billiard-simulator-python3-billiardspy

Elastic Object + Fluid 450 gradient descent iteration



Source: https://github.com/taichi-dev/difftaichi



2D Indirect Fluid Control

100 training samples, each with 16 time steps

Learns to move >99% of the volume into the bucket

Source:

Learning to Control PDEs with Differentiable Physics, Holl et. al

Supervised vs. Diff. Physics loss



Learning to throw, first 4 test cases

Overview

- Differentiable Physics
 - Motivation
 - Forward Physics Simulation
 - Differentiable ("backwards") Physics
 - Examples
- Partial Differential Equations (PDEs)
- Results of the paper
- Other examples and outlook
 - Enhancing simulations with Neural Networks (NNs)
 - Inductive Physical Bias

Partial Differential Equations (PDEs)

- The most fundamental description of evolving systems
- From quantum mechanics to turbulent flows
- For a physical system u(x, t)

$$\frac{\partial \boldsymbol{u}}{\partial t} = \mathcal{P}\left(\boldsymbol{u}, \frac{\partial \boldsymbol{u}}{\partial \boldsymbol{x}}, \frac{\partial^2 \boldsymbol{u}}{\partial \boldsymbol{x}^2}, ..., \boldsymbol{y}(t)\right)$$
$$\frac{\partial \boldsymbol{u}}{\partial t} = \mathcal{P}\left(\boldsymbol{u}, \frac{\partial \boldsymbol{u}}{\partial \boldsymbol{x}}, \frac{\partial^2 \boldsymbol{u}}{\partial \boldsymbol{x}^2}, ...\right) + \boldsymbol{F}(t)$$

• Agent only has access to o(u) (e.g. density, but not velocity)

Partial Differential Equations (PDEs)

$$\frac{\partial \boldsymbol{u}}{\partial t} = \mathcal{P}\left(\boldsymbol{u}, \frac{\partial \boldsymbol{u}}{\partial \boldsymbol{x}}, \frac{\partial^2 \boldsymbol{u}}{\partial \boldsymbol{x}^2}, \ldots\right) + \boldsymbol{F}(t)$$

• A Solver moves the system forward by a time increment Δt

$$\boldsymbol{u}(t_{i+1}) = \text{Solver}[\boldsymbol{u}(t_i), \boldsymbol{y}(t_i)] = \boldsymbol{u}(t_i) + \Delta t \cdot \mathcal{P}(\boldsymbol{u}(t_i), ..., \boldsymbol{y}(t_i))$$

 Differentiable Solvers can efficiently compute the derivatives with respect to any of the inputs:

 $\partial \boldsymbol{u}(t_{i+1})/\partial \boldsymbol{u}(t_i)$ and $\partial \boldsymbol{u}(t_{i+1})/\partial \boldsymbol{y}(t_i)$



Goal

- i. Find optimal trajectory Observation Predictor (OP) agent
- ii. Move the system along itCorrector Network agent:Control Force Estimator (CFE)



(+ minimize the force applied)

Results

- 1. Burger's Equation
- 2. Incompressible Fluid Flow
 - 1. Natural fluid flow reconstruction
 - 2. Shape transition
- 3. Incompressible Fluid with Indirect Control

Results

- 1. Burger's Equation
- 2. Incompressible Fluid Flow
 - 1. Natural fluid flow reconstruction
 - 2. Shape transition
- 3. Incompressible Fluid with Indirect Control



Overview

- Differentiable Physics
 - Motivation
 - Forward Physics Simulation
 - Differentiable ("backwards") Physics
 - Examples
- Partial Differential Equations (PDEs)
- Results of the paper
- Other examples and outlook
 - Enhancing simulations with Neural Networks (NNs)
 - Inductive Physical Bias

• Numbers in, Numbers out



$$VGG - 16\left(\begin{bmatrix} (0.8, 0.6, 0.15) & \cdots & (0.87 & 0.4, 0.2) \\ \vdots & \ddots & \vdots \\ (0.3, 0.5, 0.4) & \cdots & (0.3, 0.7, 0.2) \end{bmatrix}\right) = [.02, \dots, .001, 0.4221, \dots]$$

- Pros:
 - Fast
- Con
 - "black box", hard to explain
 - Unreliable outside known domain
 - Training

- Pros:
 - Fast
- Con
 - "black box", hard to explain
 - Unreliable outside known domain
 - Training

Idea: Let's combine the "fast" parts and the "reliable" parts

- Pros:
 - Fast
- Con
 - "black box", hard to explain
 - Unreliable outside known domain
 - Training

Idea: Let's combine the "fast" parts and the "reliable" parts

Inductive Physical Bias

Inductive Physical Bias

- Assumptions used while learning
- Underlying physical system

Inductive Physical Bias

- Assumptions used while learning
- Underlying physical system

$$\frac{\partial \boldsymbol{u}}{\partial t} = \mathcal{P}\left(\boldsymbol{u}, \frac{\partial \boldsymbol{u}}{\partial \boldsymbol{x}}, \frac{\partial^2 \boldsymbol{u}}{\partial \boldsymbol{x}^2}, \ldots\right) + \boldsymbol{F}(t)$$

Deep Learning for Cloud Rendering

• Approximating the Radiative Transfer Function

$$L(\mathbf{x},\omega) = \int_0^b T(\mathbf{x},\mathbf{x}_u)\mu_s(\mathbf{x}_u) \int_{S^2} p(\omega \cdot \widehat{\omega})L(\mathbf{x}_u,\widehat{\omega}) \,\mathrm{d}\widehat{\omega} \,\mathrm{d}u + T(\mathbf{x},\mathbf{x}_b)L(\mathbf{x}_b,\omega),$$

• By learning the indirect incoming radiance function

$$L_i(\mathbf{x},\omega) = \int_{S^2} p(\omega \cdot \widehat{\omega}) (L(\mathbf{x},\widehat{\omega}) - L_d(\mathbf{x},\widehat{\omega})) \,\mathrm{d}\widehat{\omega}.$$

Deep Learning for Cloud Rendering – Results



24x faster convergence than path tracing (PT) on average!

Main takeaway

- Don't throw away existing knowledge!
- Leverage already existing methods
- Deep Learning is not magic, and shouldn't be treated as such

Overview

- Differentiable Physics
 - Motivation
 - Forward Physics Simulation
 - Differentiable ("backwards") Physics
 - Examples
- Partial Differential Equations (PDEs)
- Results of the paper
- Other examples and outlook
 - Enhancing simulations with Neural Networks (NNs)
 - Inductive Physical Bias

Sources

Papers:

- Learning to Control PDEs with Differentiable Physics [Holl et al. 2020] (+ the supplemental materials) <u>https://ge.in.tum.de/publications/2020-iclr-holl/</u>
- DiffTaichi: Differentiable Programming for Physical Simulation
 <u>https://arxiv.org/abs/1910.00935</u>
- Deep Scattering: Rendering Atmospheric Clouds with Radiance-Predicting Neural Networks, by Kallweit et al. [2017] <u>https://la.disneyresearch.com/publication/deep-scattering/</u>

Talks:

- Differentiable Physics (for Deep Learning), Overview Talk by Nils Thuerey [<u>https://youtu.be/BwuRTpTR2Rg</u>, 40m48s]
- Differentiable Physical Simulation and AI @ NeurIPS 2020 DiffCVGP workshop [https://youtu.be/i2072iMe9ug, 29m59s]
 - Differentiable physics and the Taichi programming language
 - Code repo with examples: <u>https://github.com/taichi-dev/difftaichi</u>