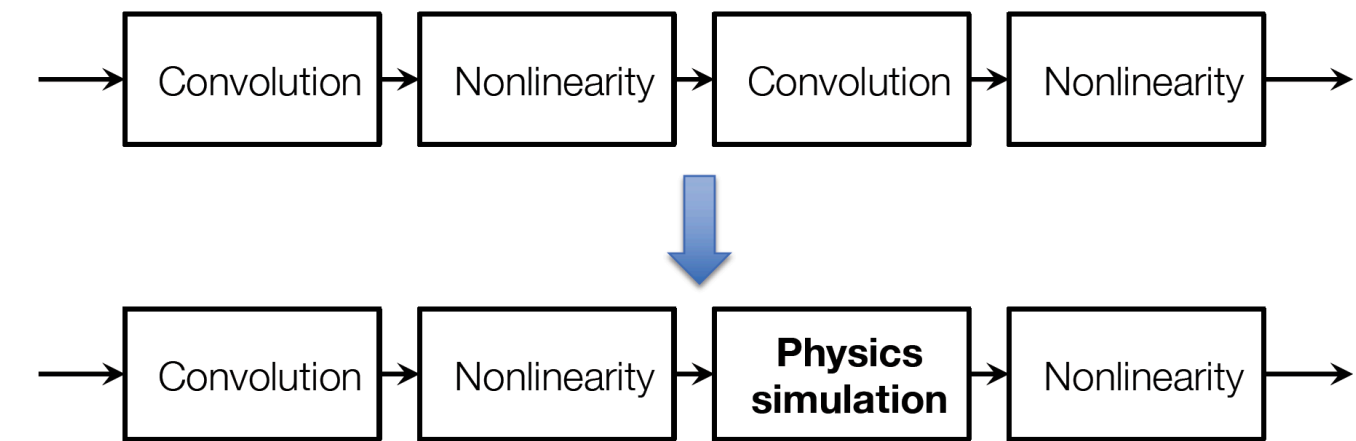


## Introduction & Motivation

- Simulation-based models have strong physical knowledge and learn efficiently, but are inflexible; learned models are flexible, but require extensive (re)training and predictive precision decays fast
- We develop a differentiable physics engine, which has both precise knowledge of physics and can be embedded in an end-to-end learning system



- Previous similar work have either used numerical differentiation methods or relied solely on auto-differentiation. We propose finding the analytic gradients by differentiating the dynamics equations
- This system contributes to a recent trend of incorporating components with structured modules into end-to-end learning systems such as deep networks.

## Dynamics LCP

- Rigid body dynamics can be framed as a linear complementarity problem (LCP)
- Newtonian dynamics represented in terms of velocity, using a discrete time step
- Constraints are added to enforce rigid body dynamics

$$\begin{aligned} \mathcal{M}\dot{v} &= f^{(c)} + f \quad \Rightarrow \quad \mathcal{M}(v_{t+dt} - v_t) = dt f_t^{(c)} + dt f_t \\ \mathcal{J}_e \lambda_e &= 0 \quad \text{Equality constraints} \\ (\lambda_c, \mathcal{J}_c v + c) &\in \mathcal{C} \quad \text{Contact constraints} \\ (\lambda_f, \mathcal{J}_f v + E\gamma) &\in \mathcal{C} \\ (\mu \lambda_c - E^T \lambda_f, \gamma) &\in \mathcal{C} \end{aligned} \quad \left. \vphantom{\begin{aligned} \mathcal{J}_e \lambda_e &= 0 \\ (\lambda_c, \mathcal{J}_c v + c) &\in \mathcal{C} \\ (\lambda_f, \mathcal{J}_f v + E\gamma) &\in \mathcal{C} \\ (\mu \lambda_c - E^T \lambda_f, \gamma) &\in \mathcal{C} \end{aligned}} \right\} \text{Friction constraints}$$

where  $\mathcal{C}(a, b) = \{a \geq 0, b \geq 0, a^T b = 0\}$

- Equality constraints define joints, contact constraints prevent penetrations and friction constraints define frictional forces
- Contact and friction constraints are inequalities and also have a complementarity term ( $a^T b = 0$ ), which characterizes the LCP formulation
- Solvable via primal-dual interior point method

## Differentiable physics

- Optimality conditions for LCP can be written compactly

$$\mathcal{M}x + A^T y + G^T z + q = 0$$

$$Ax = 0$$

$$Gx + Fz + s = m$$

$$s \geq 0, z \geq 0, s^T z \geq 0.$$

where:

$$\begin{aligned} x &:= -v_{t+dt} & q &:= -\mathcal{M}v_t - dt f_t & s &:= \begin{bmatrix} a \\ \sigma \\ \zeta \end{bmatrix} & F &:= \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & E \\ \mu & -E^T & 0 \end{bmatrix} \\ y &:= \lambda_c & A &:= \mathcal{J}_c & & & & \\ z &:= \begin{bmatrix} \lambda_c \\ \lambda_f \\ \gamma \end{bmatrix} & G &:= \begin{bmatrix} \mathcal{J}_c & 0 \\ \mathcal{J}_f & 0 \\ 0 & 0 \end{bmatrix} & m &:= \begin{bmatrix} c \\ 0 \\ 0 \end{bmatrix} \end{aligned}$$

- Using matrix differential calculus, we now take the differentials of the system above

$$d\mathcal{M}x + \mathcal{M}dx + dA^T y + A^T dy + dG^T z + G^T dz + dq = 0$$

$$dAx + Adx = 0$$

$$dz \circ (Gx + Fz - m) + z \circ (dGx + Gdx + dFz + Fdz - dm) = 0$$

or in matrix form:

$$\begin{bmatrix} \mathcal{M} & G^T & A^T \\ D(z^*)G & D(Gx^* + Fz^* - m) + F & 0 \\ A & 0 & 0 \end{bmatrix} \begin{bmatrix} dx \\ dz \\ dy \end{bmatrix} = \begin{bmatrix} -d\mathcal{M}x^* - dA^T y^* - dG^T z^* - dq \\ -D(z^*)dGx^* - D(z^*)dFz^* + D(z^*)dm \\ -dAx^* \end{bmatrix}$$

- This system is linear in the unknowns ( $dx, dy, dz$ ), simple to solve for desired differentials

- By defining

$$\begin{bmatrix} dx \\ dz \\ dy \end{bmatrix} := \begin{bmatrix} \mathcal{M} & G^T & A^T \\ D(z^*)G & D(Gx^* + Fz^* - m) + F & 0 \\ A & 0 & 0 \end{bmatrix}^{-T} \begin{bmatrix} \left(\frac{\partial \ell}{\partial x^*}\right)^T \\ 0 \\ 0 \end{bmatrix}$$

we can derive the gradients

$$\frac{\partial \ell}{\partial q} = -d_x$$

$$\frac{\partial \ell}{\partial \mathcal{M}} = -\frac{1}{2}(d_x x^T + x d_x^T)$$

$$\frac{\partial \ell}{\partial m} = D(z^*)d_z$$

$$\frac{\partial \ell}{\partial G} = -D(z^*)(d_z x^T + z d_x^T)$$

$$\frac{\partial \ell}{\partial A} = -d_y x^T - y d_x^T$$


$$\frac{\partial \ell}{\partial F} = -D(z^*)d_z z^T$$

- Importantly, since we have already solved the LCP in the forward pass, we can compute the backward pass with just one additional solve based upon the LU-factorization of the LCP matrix

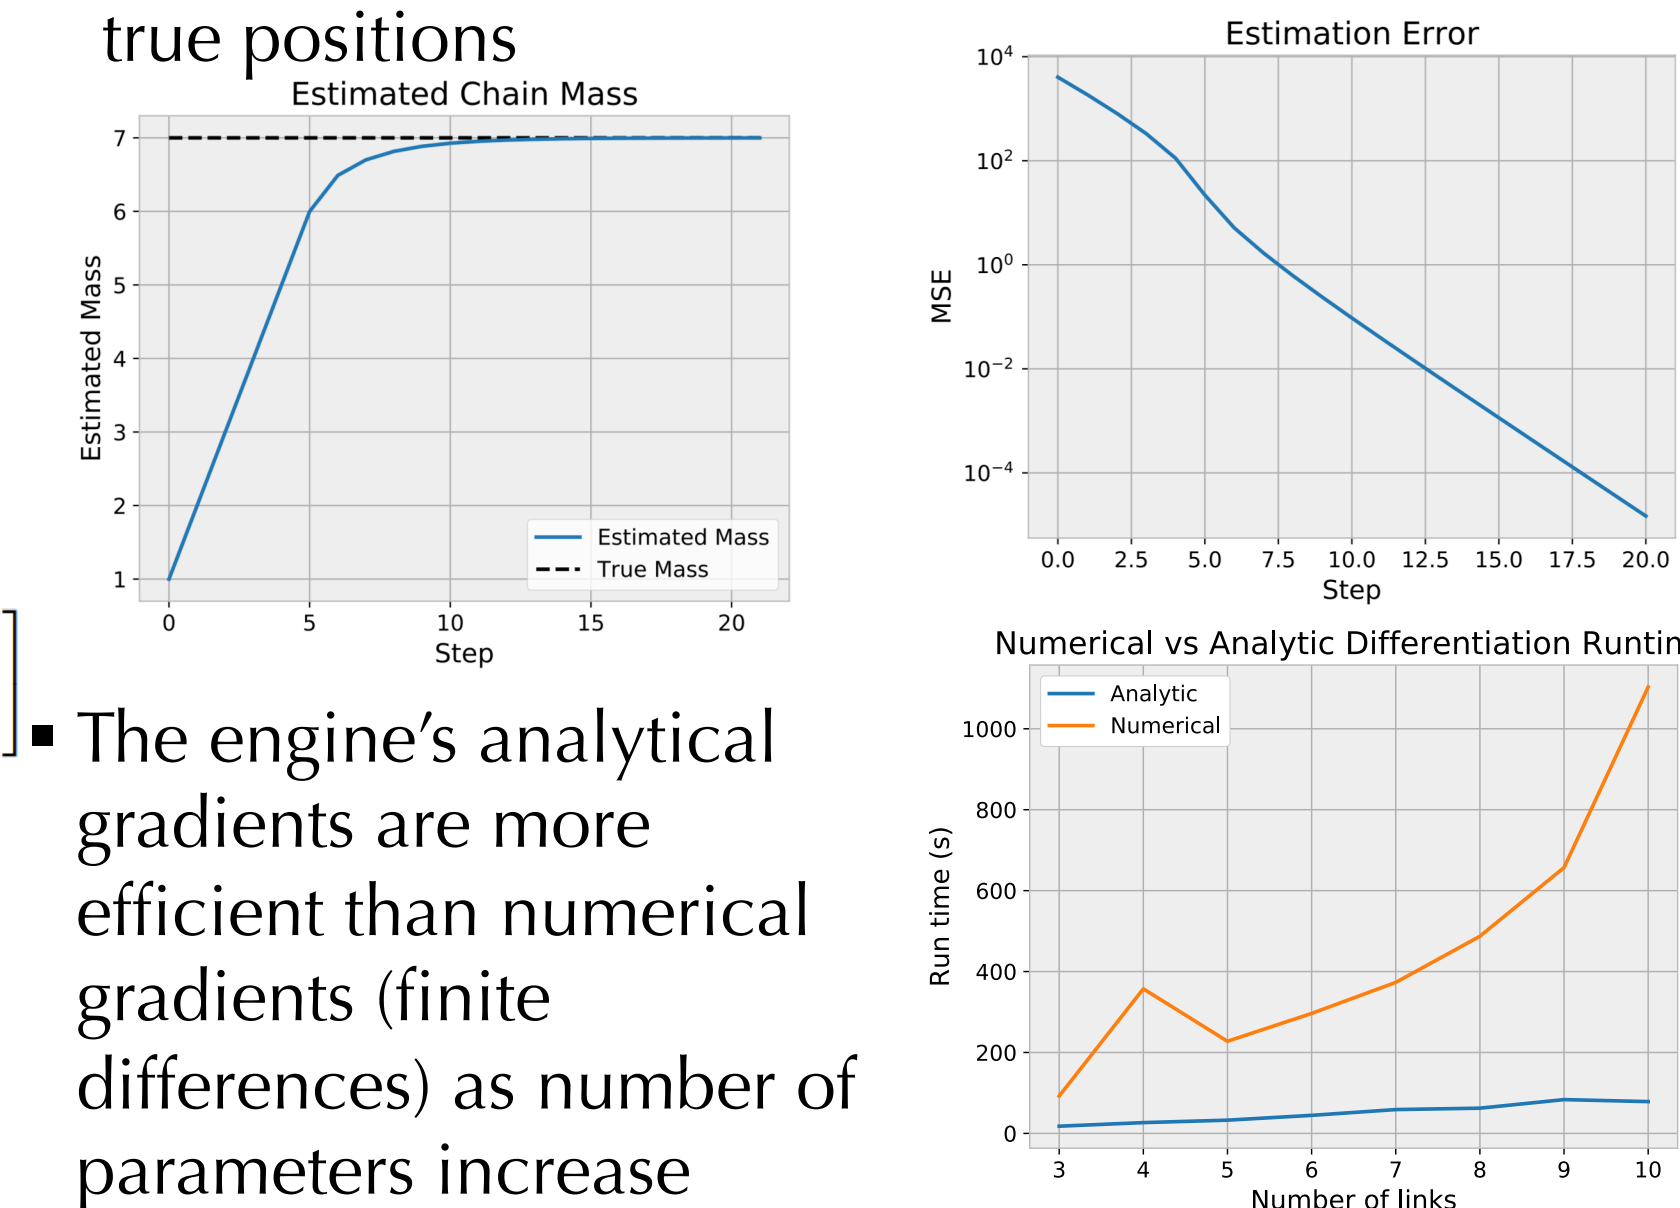
- We can effectively differentiate through the simulation at no additional cost to just running the simulation itself**

## Parameter learning

- A ball of known mass hits chain. Object positions are observed. Task is inferring the chain mass



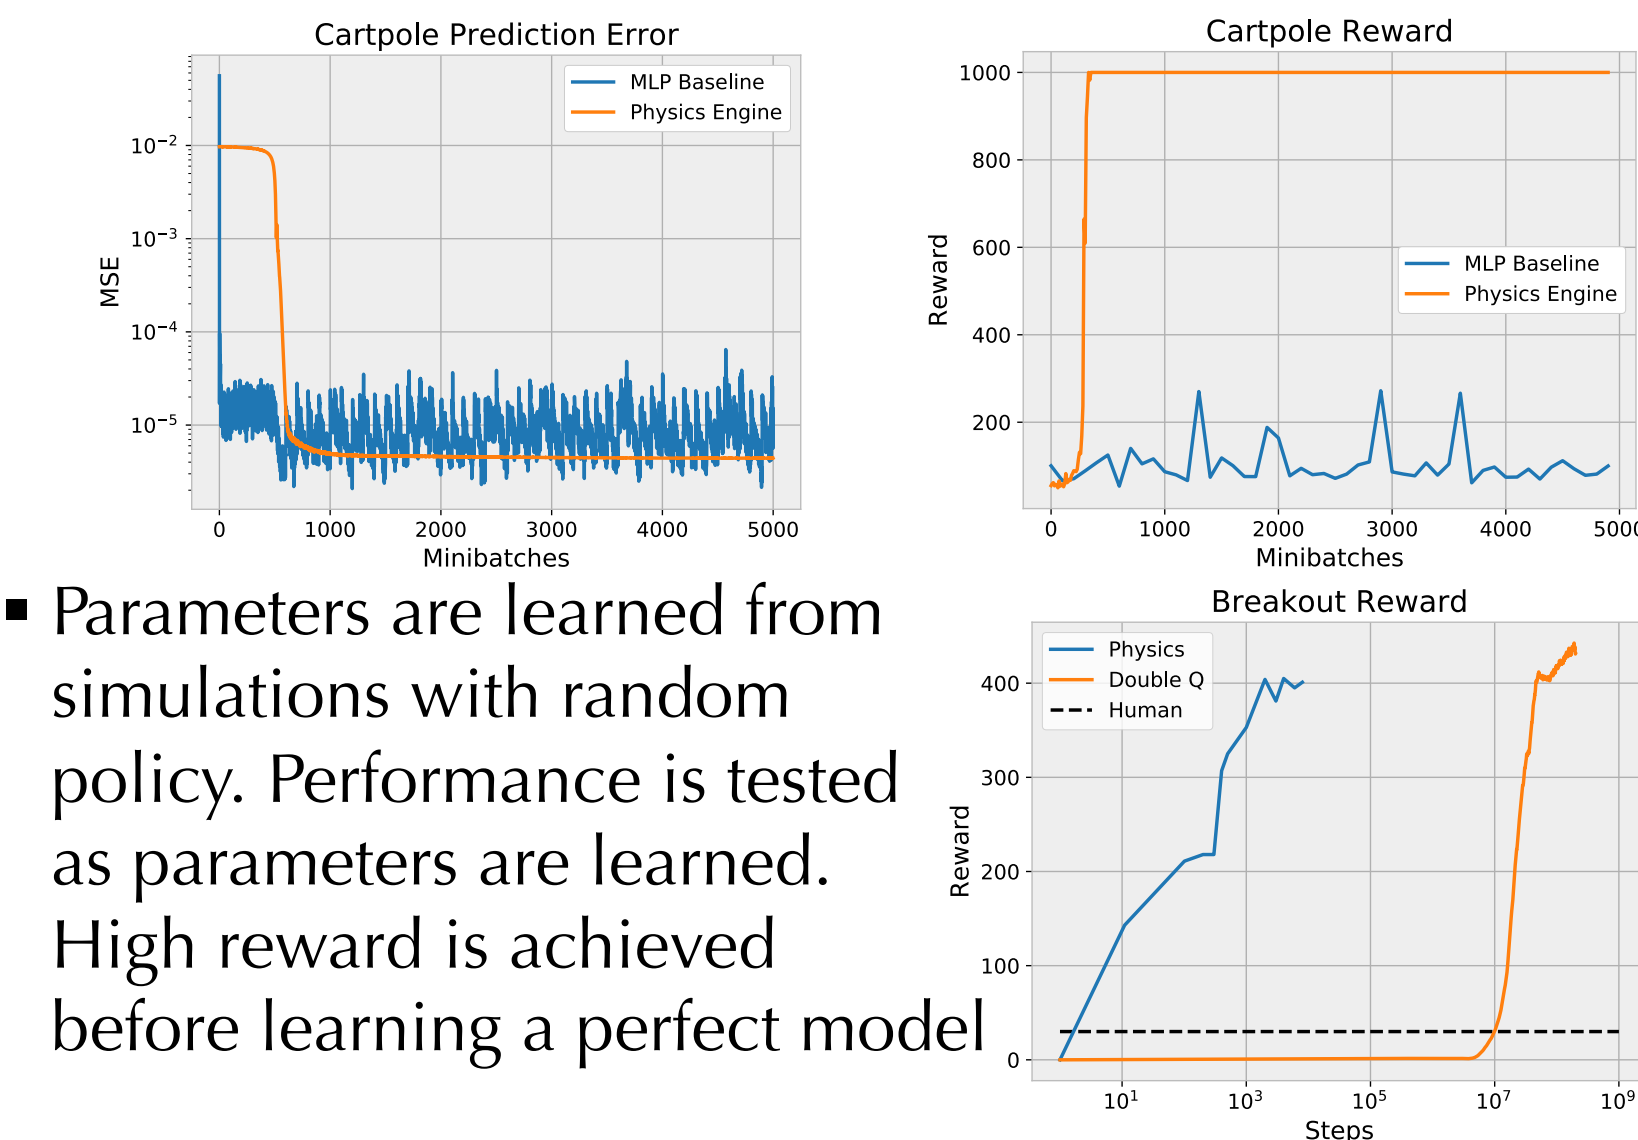
- Simulations are iteratively unrolled starting with an arbitrary mass. Estimated chain mass is minimized using gradient of MSE between the simulated and true positions



- The engine's analytical gradients are more efficient than numerical gradients (finite differences) as number of parameters increase

## Control


- Since the physics engine is differentiable, we use it in conjunction with iLQR for control in the *Cartpole* and the Atari *Breakout* tasks



- Parameters are learned from simulations with random policy. Performance is tested as parameters are learned. High reward is achieved before learning a perfect model

## Visual prediction

- After observing 3 frames of a billiard ball-like scene, predict positions 10 frames into the future



- Autoencoder architecture. *Encoder* maps frames into physical predictions. *Engine* steps physics into the future. *Decoder* draws image from physics state
- Training is performed with only partially labeled data. For unlabeled examples, prediction uses the estimated parameters (notice the hats):

$$\hat{\phi}_t = \text{encoder}(x), \quad \hat{\phi}_{t+dt} = \text{physics}(\hat{\phi}_t), \quad \hat{y} = \text{decoder}(\hat{\phi}_{t+dt})$$

- When labels are available, prediction uses the true labels (notice the lack of hats):

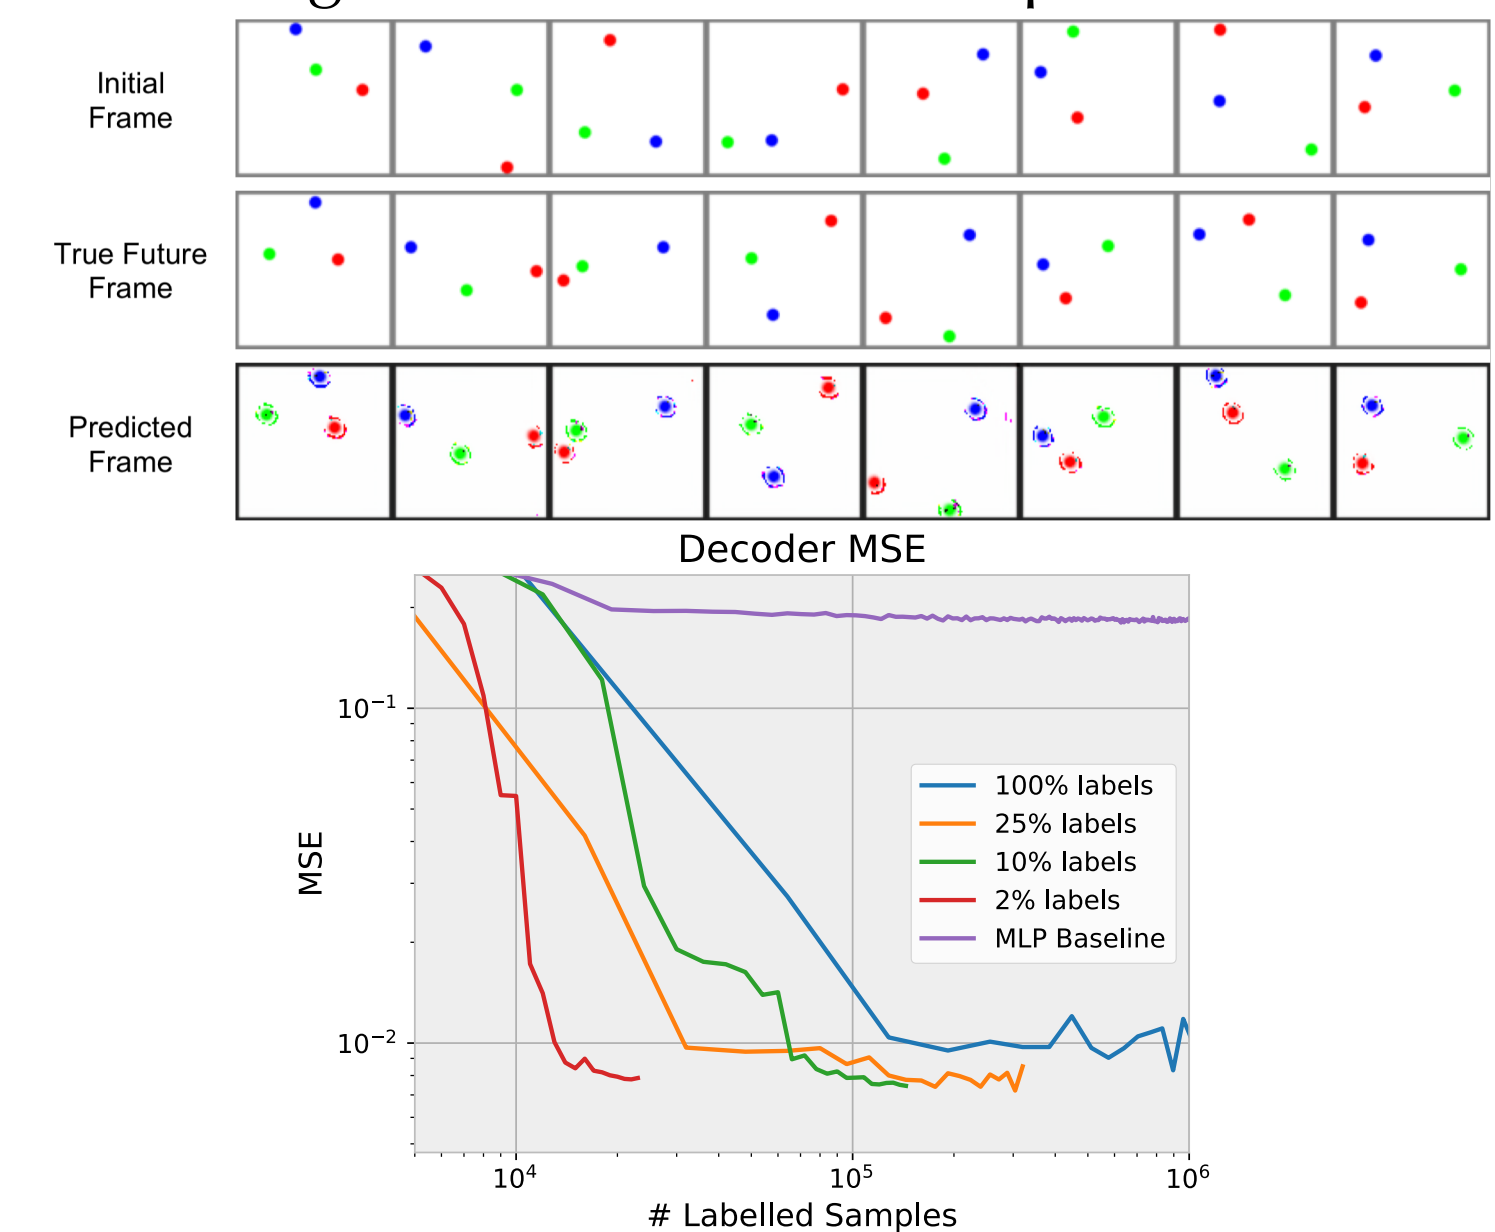
$$\phi_t = \text{encoder}(x), \quad \phi_{t+dt} = \text{physics}(\phi_t), \quad \hat{y} = \text{decoder}(\phi_{t+dt})$$

- Loss is composed of three terms when labels are available, or only the decoder loss when unlabeled

$$\mathcal{L} = \mathcal{L}_{enc} + \mathcal{L}_{phys} + \mathcal{L}_{dec},$$

$$\mathcal{L}_{enc} = \ell(\hat{\phi}_t, \phi_t), \quad \mathcal{L}_{phys} = \ell(\hat{\phi}_{t+dt}, \phi_{t+dt}), \quad \mathcal{L}_{dec} = \ell(\hat{y}, y)$$

- Physics structure embedded into the model allows for learning with few labeled samples



## External resources

- Code available at: <https://github.com/locuslab/lcp-physics>