

Robust control strategies for musculoskeletal models using deep reinforcement learning

Łukasz Kidziński & Carmichael Ong



August 7, 2018



Ref: Gillette Children's Specialty Healthcare



Ref: OpenSim (<https://www.youtube.com/watch?v=HLFQM1e-vJw>), Kat Steele



Ref: Emergence of Locomotion Behaviours in Rich Environments, Google DeepMind

How to combine modern control frameworks with knowledge and expertise embeded in musculoskeletal models?

Contents

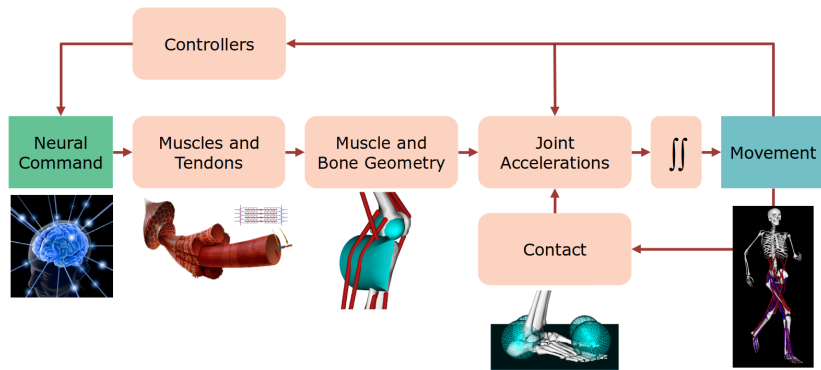
- 1 Motivation
- 2 Reinforcement learning (RL)
 - Problems and examples
 - Framework
 - Basic terminology
 - Training
- 3 Application of RL in biomechanics
 - Learning to Run challenge
 - Tips & tricks
 - AI for prosthetics
- 4 Reinforcement learning in osim-rl
 - Installation
 - Live demo

What can we do with Reinforcement Learning?

1. Model adaptation of kinematics
2. Generate data for statistical models
3. Study:
 - motor control
 - muscle synergies
4. Synthesize physiologically accurate motion

OpenSim

OpenSim is a freely available software that allows you to build, exchange, and analyze musculoskeletal models and dynamic simulations of movement. (<http://opensim.stanford.edu/>)




← “Classical” simulations




Ref: NNAISENSE

Reinforcement learning framework

Reinforcement learning models a **policy** (a decision process) of **an agent** interacting with **an environment** by taking **actions** optimizing **rewards**.



Ref: Wikipedia




Ref: Very good friends and their child

Agent



- Tries to achieve the *objective* (long-term)
- Tries to maximize the *reward* (short-term)
- Interacts with the environment

Reward/Penalty



Clapping



Falling

Observation and action



Observation:

- State of the world
- Sensory input
- State of muscles, joints and bones
- etc.

Action:

- Muscle excitations

Policy:

- Links observations with actions

Mathematical definition of a policy

Policy maps states to actions

Let S be a space of states, A be a space of actions.
Policy is a function $P_\theta : S \rightarrow A$.

Any statistical model can be a policy

Normally, a policy is parametrized with some parameters $\theta \in \mathbb{R}^p$.


We want to find the policy maximizing rewards

Let $R : S \rightarrow \mathbb{R}$ be the reward function. We optimize

$$\arg \max_{\theta} \sum_{t=1}^T R(s_{t,\theta}),$$

where $\{s_{t,\theta}\}_{1 \leq t \leq T}$ is a trajectory dependent on the policy.

The most popular model is a neural network



Stacked linear regressions with nonlinear transformations in outputs.


Training a policy

We want to find the policy maximizing rewards.

Generic optimization approach:

- Start with an empty (random) policy
- Repeat until convergence
 - Use the policy with noise to explore environment
 - Improve the policy given the experience

Toy problem



Pendulum Swing-up

- **Environment:** Pendulum (one joint, one degree of freedom)
- **Objective:** Swing it up
- **Observation:** Angular position & velocity (a vector in \mathbb{R}^2)
- **Action:** Apply torque
- **Reward:** Negative distance from 0 velocity, 0 angle


Exploration of the environment



Ref: Matthew Sheen, Reinforcement Learning - Pendulum Swing-up

<https://www.youtube.com/watch?v=YLAWnYAsai8>

Solve with policy gradient (DDPG)



Let $Q : S \times A \rightarrow \mathbb{R}$ be an approximation of a state-action value.
Let $P_\theta : S \rightarrow A$ be the best policy so far.

Now iteratively:

- Update approximation of Q from history of trials
- Optimize $J(\theta) = Q(s, P_\theta(s))$
- Compute the gradient of J
- Update θ

Contents

- 1 Motivation
- 2 Reinforcement learning (RL)
 - Problems and examples
 - Framework
 - Basic terminology
 - Training
- 3 Application of RL in biomechanics
 - Learning to Run challenge
 - Tips & tricks
 - AI for prosthetics
- 4 Reinforcement learning in osim-rl
 - Installation
 - Live demo


NIPS 2017: Learning to Run challenge

Task:

- **Environment:** A model with 18 muscles and 9 DOF
- **Objective:** Go as far as possible in 10 seconds
- **Observation:** State of the model
- **Action:** Muscle excitations
- **Reward:** Speed

Some details:

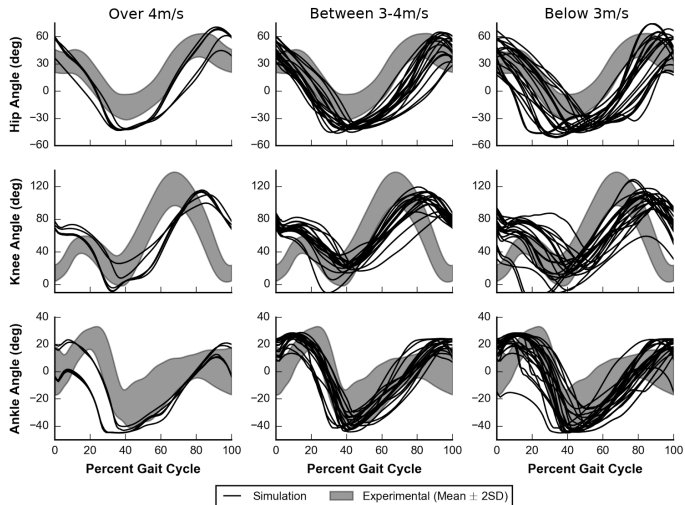
- Obstacles in the ground
- 2D model (the model doesn't fall sideways)
- Don't use experimental data





Ref: NNAISENSE

Experimental vs synthesized data



Speeding up exploration

- Frameskip: instead of sending signals every 1/100 of a second, keep the same control for, for example, 5 frames.
- Binary actions: excitations 0 or 1 instead of values in the interval $[0,1]$.
- Reward shaping: modifying the reward for training in such a way that it still makes the model train faster for the actual initial reward.

Ref: Kidziński, Łukasz, et al. "Learning to Run challenge solutions: Adapting reinforcement learning methods for neuromusculoskeletal environments." arXiv preprint arXiv:1804.00361 (2018).



NIPS 2017

NIPS 2017: Learning to Run

Reinforcement learning environments with musculoskeletal models



By [Stanford Neuromuscular Biomechanics Laboratory](#)

Completed

78340

Views

578

Participants

2154

Submissions




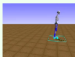
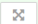

224

FOLLOW

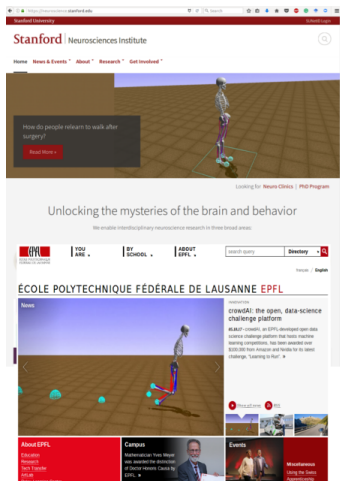
[Overview](#) [Leaderboard](#) [Discussion](#) [Dataset](#) [Winners](#)

Round 1

Round 2

△ #	Participant	Media	Mean Reward per Simulation	Entries	Last Submission (UTC)	
01.	 USTC-IMCL		44.618	47	Mon, 30 Oct 2017 03:47	
02.	 Megvii-...		43.968	57	Sat, 4 Nov 2017 10:16	

<https://www.crowdai.org/challenges/nips-2017-learning-to-run>



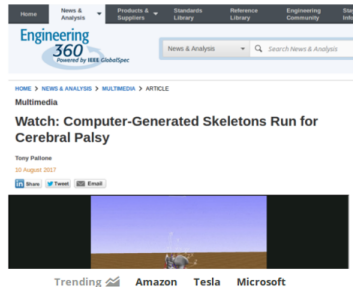
TE News Startups Mobile Gadgets Enterprise Social Europe



NVIDIA Sponsors “Learning to Run” AI Competition at NIPS 2017

September 8, 2017

Participants in the Neural Information Processing Systems (NIPS) conference “Learning to Run” competition are vying for the chance to win an NVIDIA DGX Station, the fastest personal supercomputer for researchers and data scientists.



chatbots
machine learning

Dueling AIs compete in learning to walk, secretly manipulating images and more at NIPS


NIPS 2018: AI for prosthetics

Rules:

- 19 muscles and 14 DOF
- Match requested velocity vector
- 3D model
- Experimental data is allowed

Rewards:

- Solving a real medical problem!
- Travel grants to Stanford & EPFL
- \$250 Google Cloud Credits for start
- 4 NVIDIA GPUS to win (~\$3,000 each)



Contents

- 1 Motivation
- 2 Reinforcement learning (RL)
 - Problems and examples
 - Framework
 - Basic terminology
 - Training
- 3 Application of RL in biomechanics
 - Learning to Run challenge
 - Tips & tricks
 - AI for prosthetics
- 4 Reinforcement learning in osim-rl
 - Installation
 - Live demo

Linking OpenSim with RL (osim-rl)



osim-rl

HOME DOCS NEWS HELP TEAM

V1.5 GITHUB

Reinforcement learning with musculoskeletal models in OpenSim

NIPS 2018: AI for prosthetics

Design artificial intelligence to control human body and predict performance of a prosthetic leg. Participate in the NIPS 2018 challenge to win prizes and fame.

[Learn more about the challenge →](#)

OpenSim RL

Use our musculoskeletal reinforcement learning environment for other projects in computer science, neuroscience, biomechanics, etc.

[Learn more about osim-rl →](#)

Stanford University Berkeley EPA MITRE NVIDIA Amazon Web Services TechCrunch Google Cloud Platform

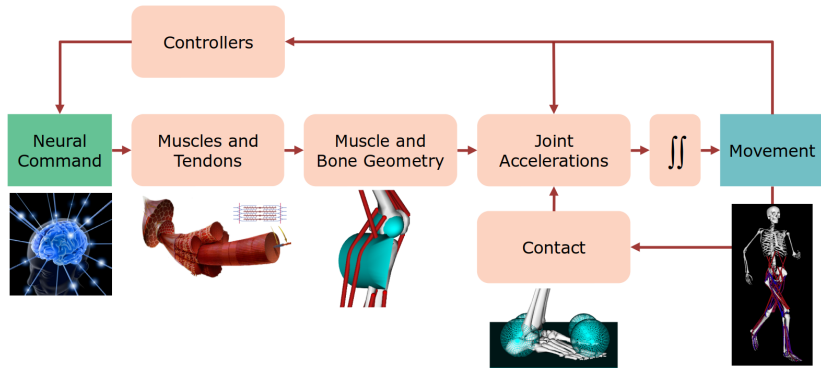
osim-rl is an **interface** for building **robust controllers** for your **OpenSim models** using open **deep reinforcement learning algorithms**.

<http://osim-rl.stanford.edu/>


```

> # Assume that the brain variable is a policy function
> env = ProstheticsEnv(visualize=True)
> observation = env.reset()
> for i in range(200):
>     observation = env.step(brain.action(observation))

```



Arm model



<https://github.com/stanfordnmb1/osim-rl/blob/master/examples/legacy/train.arm.ipynb>

Summary

- Reinforcement learning (RL) allows you to build a robust controller by specifying a high-level objective
- `osim-rl` package allows you to use RL frameworks to control your OpenSim models