15-464/15-664 Reference List for April 15, 2019

Following on the optimization discussion from last Wednesday, we started with this paper, where the main idea is to use inverse optimization – i.e., to formulate an optimization problem for which observed motion is optimal.

Wampler, Kevin, Zoran Popović, and Jovan Popović. "Generalizing locomotion style to new animals with inverse optimal regression." *ACM Transactions on Graphics (TOG)* 33.4 (2014): 49.
http://grail.cs.washington.edu/projects/inverse_locomotion/

We then jumped back to 2010 to pick up these two papers which try to create character controllers to recreate a given motion (e.g., rolling, cartwheels) through simulation. These papers use a sampling based approach with the goal of obtaining a single simulation that well matches the input motion data.

Liu, Libin, KangKang Yin, Michiel van de Panne, Tianjia Shao, and Weiwei Xu. "Sampling-based contact-rich motion control." *ACM Transactions on Graphics (TOG)* 29, no. 4 (2010): 128.
https://www.cs.ubc.ca/~van/papers/2010-TOG-sampControl/index.html

Liu, Libin, KangKang Yin, and Baining Guo. "Improving Sampling-based Motion Control." In *Computer Graphics Forum*, vol. 34, no. 2, pp. 415-423. 2015. http://libliu.com/Samcon2/Samcon2.html

For many applications (e.g., games), it is not sufficient to create a single, isolated simulated motion. The following papers uses linear regression to create policies that are sufficiently robust to allow simulated motions to be connected into control graphs.

Liu, Libin, Michiel Van De Panne, and KangKang Yin. "Guided learning of control graphs for physics-based characters." *ACM Transactions on Graphics (TOG)* 35, no. 3 (2016): 29.
https://www.cs.ubc.ca/~van/papers/2016-TOG-controlGraphs/index.html

Motion does not have to come from motion capture data, and this paper discusses extracting reference motion from video and using it to create policies for character animation / simulation.

Peng, Xue Bin, Angjoo Kanazawa, Jitendra Malik, Pieter Abbeel, and Sergey Levine. "Sfv: Reinforcement learning of physical skills from videos." In *SIGGRAPH Asia 2018 Technical Papers*, p. 178. ACM, 2018.
https://xbpeng.github.io/projects/SFV/index.html

Stepping back to 2014, we had a look at this paper, which does not use a reference motion, but instead utilizes sampling in real time to choose character's actions in a game environment. There was a follow on paper, which we did not see in class, which is also listed here.

Hämäläinen, Perttu, et al. "Online motion synthesis using sequential monte carlo." *ACM Transactions on Graphics (TOG)* 33.4 (2014): 51. https://mediatech.aalto.fi/publications/graphics/OnlineSMC/

Hämäläinen, Perttu, Joose Rajamäki, and C. Karen Liu. "Online control of simulated humanoids using particle belief propagation." *ACM Transactions on Graphics (TOG)* 34, no. 4 (2015): 81. https://mediatech.aalto.fi/publications/graphics/C-PBP/

Finally, we looked at a very recent paper on learning locomotion without a reference motion as guidance and also without state as we saw in last week's class. This paper represents the current state of the art in learning locomotion without these priors.

Yu, Wenhao, Greg Turk, and C. Karen Liu. "Learning symmetric and low-energy locomotion." *ACM Transactions on Graphics (TOG)* 37, no. 4 (2018): 144. https://wenhaoyu.weebly.com/publications.html

We took a brief look at reinforcement learning and deep reinforcement learning during this class. If you are interested in learning more, you may want to take a look at this paper.

Arulkumaran, Kai, Marc Peter Deisenroth, Miles Brundage, and Anil Anthony Bharath. "A brief survey of deep reinforcement learning." *arXiv preprint arXiv:1708.05866* (2017). https://arxiv.org/abs/1708.05866