BY KATHARINE MILLER

THE "LEARNING TO RUN" CHALLENGE:

Engaging Data Scientists in Biomechanics

Most of us take walking or running for granted. But injury or neurological disease can cause these basic skills to deteriorate to such a degree that they need to be re-learned. To better understand how the brain accomplishes such learning, Stanford University's Mobilize Center challenged the research community to a competition: develop a controller (essentially a brain) that will allow a physiologically-based human model to navigate through a complex obstacle course as quickly as possible. At the December 2017 Neural Information Processing Systems (NIPS) conference, winners of the "NIPS 2017: Learning to Run" challenge will be announced.

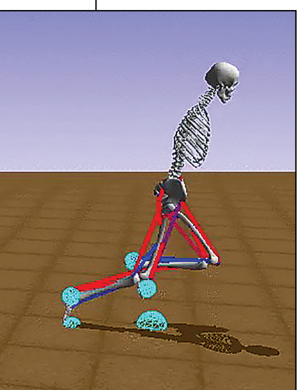

The nature of the challenge has drawn many participants who are more familiar with data science concepts—such as reinforcement learning—than they are with muscles and bones, says **Łukasz Kidziński, PhD**, the Mobilize Center postdoctoral student who co-organized the competition in collaboration with researchers at the University of California, Berkeley, and École polytechnique fédérale de Lausanne in Switzerland. "Because they come at the problem from a different field, they use other types of skills to find new ways to solve challenging biomechanical problems," he says.

For the competition, Kidziński says, "We gave people a model with muscles and bones and interactions and constraints of the human body, and they had to build a brain for this body." This simplified lower body model, which consists of bones and 18 key muscles, was designed in OpenSim, the biomechanical simulation platform developed by Mobilize Center researchers. "Other experiments had shown that these are enough muscles to synthesize human-like gait," Kidziński says. He and his colleagues also set some reasonable constraints on the model, such as limits on the forces that muscles can exert, and a requirement that the model can only go forward. In addition, they built a virtual obstacle course of spheres inserted into the ground. "Something you can trip on," Kidziński says. The purpose was to make the models more generalizable and to see if the model can adapt just as a human would. Those who join the competition create their controllers in OpenSim, which Kidziński and his colleagues modified to include reinforcement learning in a changeable environment.

More than 400 people have joined the competition, which received over 1,500 submissions by early October—three weeks ahead of the deadline for the competition's first round. For November, the top models from the first round have been given a new challenge—a change in muscle strength, for example, or the frequency of obstacles. The victor's model will be the one that goes farthest in a set amount of time.

While there are some cool prizes, most participants are in it for the learning and community recognition. The competition's leaderboard shows a thumbnail video of the models in action. In early October, the top contestant was **Jackie Tseng**, a PhD candidate at Tunghai University in China. She says the most interesting part of the competition has been training her agents and watching them progress from falling down continually, to taking first steps, to crossing over an obstacle, and then to running—and earning a high score. Another competitor, **Anton Pechenko**, a research and development engineer at Yandex, in Moscow, agrees this is the fun part: "There are three things you can watch forever," he says, "Fire, water, and your agent performing actions solving the problem."

Tseng looks forward to the day when computer vision might be added to reinforcement learning systems such as this one, to allow agents to understand and adapt to more complex environments. For now, she says, "Though the half-humanoid model trained in the 'Learning to Run' challenge is still relatively simple, it is a quite important beginning in AI for biomechanics." ☐

DETAILS

Results of the NIPS 2017: "Learning to Run" challenge will be posted at https://www.crowdai.org/challenges/nips-2017-learning-to-run

Acknowledgements:

Coordinators: Carmichael Ong, Mohanty Sharada, Jason Fries, Jennifer Hicks, Joy Ku

Advisors: Sergey Levine, Marcel Salathé, Scott Delp

The Mobilize Center for Mobility Data Integration and Insight is an NIH Big Data to Knowledge (BD2K) Center of Excellence at Stanford University.