Virtual ACS 2021 Surgeons and Engineers: A Dialogue on Surgical Simulation Meeting

Research

Novel Application of Reinforcement Learning to Automate Surgical Subtasks Rendered in a Virtual Soft-Body Simulation

Alexandra Tan Bourdillon, BS; Animesh Garg, PhD; Hanjay Wang, MD; Joseph Woo, MD; Marco Pavone, PhD; and Jack H. Boyd, MD.

Yale School of Medicine, New Haven, CT; University of Toronto, Toronto, ON, Canada; Stanford School of Medicine, Stanford, CA

Introduction: The revolutions in artificial intelligence hold tremendous capacity to augment human achievements in surgery, but robust integration of deep learning algorithms with high-fidelity surgical simulation remains a challenge. We present, to our knowledge, the first application of reinforcement learning for automating surgical dissection in a graphical simulation.

Methods: Using the Unity3D game engine, we integrated the “Machine Learning-Agents” Unity package with a graphical environment rendering a planar, deformable tissue modeled by Nvidia’s FleX particle simulator (uFlex package). Our network comprised 2 hidden layers and 128 hidden units. Proximal Policy Optimization rewarded tissue collision along the desired path, simulating dissection. TensorFlow analytics informed hyperparameter tuning; constant and proportional rewards were evaluated. Task automation consisted of two stages: one-dimensional motion (forward and backward) followed by two-dimensional (adding lateral movement).

Results: In the first stage with one-dimensional motion, a downward trend in loss values was achieved. Episode length significantly decreased over 30,000 steps ($p < 0.001$, $R^2 = 0.38$) and better performance was obtained using constant rewards compared to proportional rewards ($p<0.05$, $R^2 = 0.32$). In stage two allowing for two dimensions of freedom, cumulative reward progressed slightly over 50,000 iterations ($p = 0.17$) and the rate of accomplishing the simulated task approached 50% completion. Entropy of the agent’s actions increased and then decreased over the training period ($p < 0.001$, $R^2 = 0.67$).

Conclusions: This work marks a promising methodology to develop automated programs for performing surgical subtasks, but further work is needed to enhance simulation complexity and computational performance. Ultimately, advancements in computer vision will allow for the translation of virtual algorithms in the physical world.

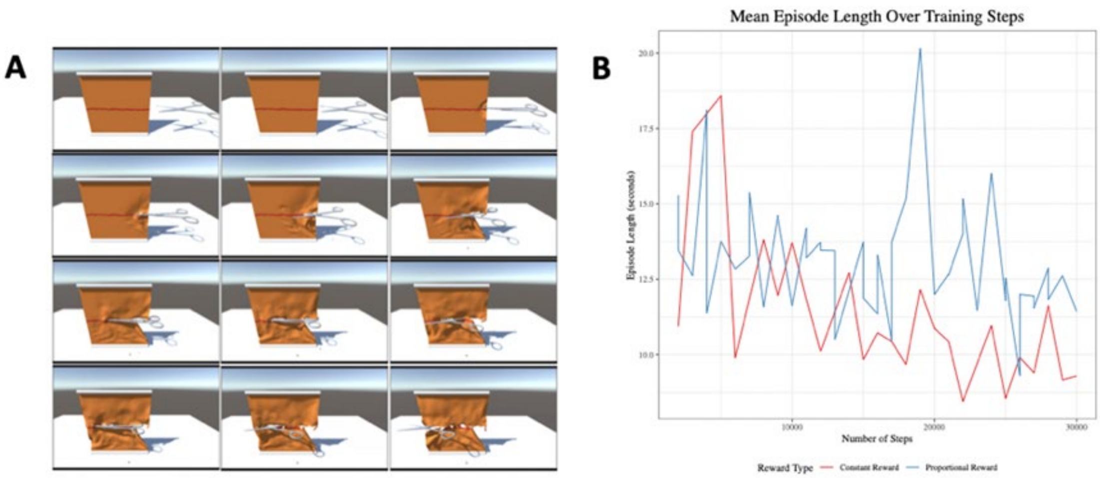

Figure: Agent scissors autonomously cut deformable tissue in a virtual environment (A). Training with constant rewards rather than proportional rewards achieved better performance, as seen by greater decreases in episode length (B).