A challenging control task for automated driving is lane merging. Previous research utilized rule-based or optimal control approaches to solve the merging problem[2], but there is an increasing interest in using learning-based methods to train an automated merging policy[3], partly due to the super-human intelligence displayed by deep reinforcement learning (DRL) in playing board games[4]. DRL algorithms use deep (multi-layer) neural nets as the policy representation and, once trained, do not require the computing power of model-predictive control[5]. In our research, we examined decentralized decision making and control for on-ramp merging using DRL. The goal was to merge autonomously from an on-ramp to a highway while avoiding stops and collisions. The on-ramp automated vehicle obtained the states of other vehicles with its own sensors; there was no centralized control from roadside units. We used the states of the automated (merging) and highway vehicles as the input to the DRL training framework, and the acceleration command for the automated vehicle as the output. After training convergence, the learned policy enabled the automated vehicle to decelerate to merge behind (or accelerate to merge ahead) of a highway vehicle, much like human driving. We tested the learned policy with 16,975 merging episodes and observed 0 stops and only 1 collision, which is better performance for highway merging than that achieved by humans[6].

Integration of Machine Learning with Dynamics and Control: From Autonomous Cars to Biomechatronics

DR. JOHN McPHEE, FCSME, FCAE, FASME, FEIC

Professor J. McPhee is the Canada Research Chair in System Dynamics at the University of Waterloo, which he joined in 1992. Prior to that, he held fellowships at Queen's University, and the Université de Liège, Belgium.

He pioneered the use of linear graph theory and symbolic computing to create real-time models and model-based controllers for multi-domain dynamic systems, with applications ranging from autonomous vehicles to rehabilitation robots and sports engineering. His research algorithms are a core component of the widely-used MapleSim modelling software, and his work appears in over 130 journal publications.

Prof. McPhee is the past Chair of the International Association for Multibody System Dynamics, a co-founder of two international journals and three technical committees, a member of the Golf Digest Technical Panel, and an Associate Editor for five journals. He is a Fellow of the Canadian Society of Mechanical Engineers, the Canadian Academy of Engineering, the American Society of Mechanical Engineers, and the Engineering Institute of Canada. He has received eight Best Paper Awards, the prestigious NSERC Synergy Award from the Governor General of Canada, and the PEO Engineering Excellence Medal in 2019.

IMAGES COURESTY OF THE RESEARCHER / UNIVERSITY OF WATERLOO

RESEARCHERS AT THE UNIVERSITY OF WATERLOO are strategically combining artificial intelligence (AI) methods with dynamics and model-based control approaches, to solve real-world problems that are too difficult or computationally intensive to tackle with either AI or conventional methods on their own. Our goal is to combine the best features of machine learning and model-based control theories, in contrast to computer scientists that are seeking purely AI solutions to dynamics and control problems. This article presents our recent work towards this goal, with an emphasis on two very active research applications: Autonomous Cars and Biomechatronics.

Autonomous Cars

To engage the CSME reader, we start with our autonomous car, the Autonomoose (Figure 1). From rolling dyno and track testing at Waterloo, we developed a physics-based model of the vehicle and tires. However, the details of the powertrain control were proprietary and unknown to us. Data collection from the on-board CAN bus allowed us to train a double-layer perceptron neural network representation of the powertrain, which was easily integrated with the vehicle and tire models using symbolic computing. By combining the physics-based and AI-based models into a single hybrid representation, we could create a model-predictive controller (MPC) for longitudinal speed that was both accurate and computationally efficient; further speedup was obtained by differentiating the symbolic model to obtain exact expressions for the gradients needed by the optimizer[1]. The MPC was deployed on the Autonomoose as a ROS (Robot Operating System) node, and preliminary track testing has demonstrated the improved performance of this hybrid controller in comparison to previous PI-based approaches.

Biomechatronics

The strong performance of reinforcement learning in the previous project motivated us to use DRL in our end-effector-based stroke rehabilitation robot, for which we have developed a model-predictive controller. For effective rehabilitation, the weights of the controller should be tuned for each individual patient. Manually tuning for different patients is time-consuming and will not lead to the best performance. Inspired by[7], we are designing an adaptive controller with a DRL-based automatic tuner. Our goal is to develop a subject-specific controller, using measurements of the position and force at the robot end-effector, which will adapt to different patient movement and strength characteristics.

The rehabilitation robot measurement data can be augmented with the motion of the user. Currently, this information is extracted from markers or wearable technologies that are cumbersome and have high setup time. This is the motivation for a related project that is using deep convolutional neural networks (CNNs)

FIG. 2: POSE ACTIVATIONS BEFORE AND AFTER TRAINING THE POSE ESTIMATION LAYERS FOR HUMAN ACTION RECOGNITION.

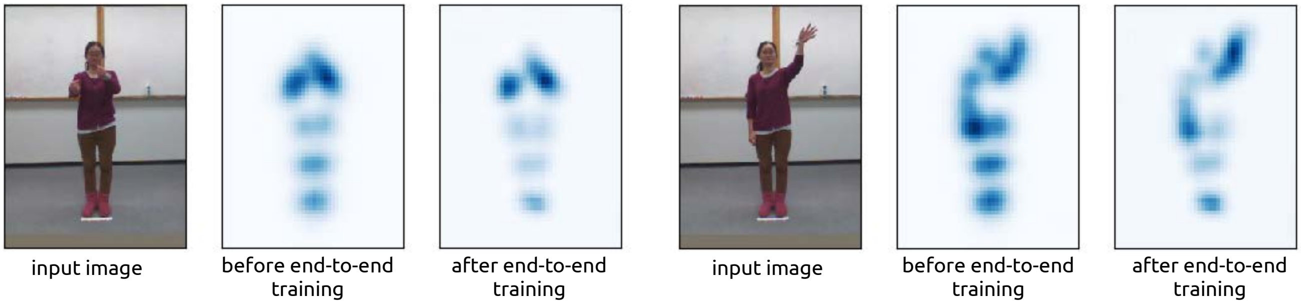

to simultaneously recognize actions and perform markerless tracking of human motion[8]. In this project, a custom CNN architecture was designed to circumvent some of the challenges associated with training the complex neural network architectures commonly used for action recognition by exploiting the spatial information in pre-trained human pose estimation layers. More specifically, the generated pose data is re-projected using a stack of 3D convolutions in a seamless network architecture that integrates human pose with action recognition. The advantage of this framework is that action predictions are based on pose information rather than appearance, which is advantageous in single-environment applications where the variation in human movement is critical. On Multimodal Human Action Dataset (UTD-MHAD)[9], a 27-class multi-modal action recognition dataset, the proposed method – requiring only an RGB camera – outperformed several methods using richer data streams from depth cameras and inertial sensors. Interestingly, it was found that the spatial activations produced by the pose layers changed after training the model for action recognition. Figure 2 demonstrates that the activations for the joints that contributed most to the action, such as the wrists and ankles, were greatly accentuated after end-to-end training. Conversely, the non-moving joints such as the knees and hips were attenuated. This interesting phenomenon can be leveraged to gain new insights into discovering the key defining characteristics for specific human actions.

Vision-based artificial intelligence has also advanced our exoskeleton rehabilitation robotics research. While many robotic vision applications have focused on autonomous navigation, recent research has considered energy savings from environment sensing. For instance, autonomous and connected vehicles use future environment information to achieve energy-efficient driving through optimal power management. Environment recognition can likewise improve the energy-efficient control of lower-limb biomechatronic devices (e.g., exoskeletons and prostheses) that assist seniors and rehabilitation patients with walking[10]. The implementation of such energy-efficient controllers, however, remains fundamentally contingent upon their ability to accurately predict oncoming environments. Therefore, our research has focused on developing an accurate environment recognition system with robotic vision and deep learning (Figure 3)[11]. Using a wearable RGB camera, approximately 2 million images were collected while walking through unknown outdoor and indoor environments; the labelled images were uploaded to IEEE DataPort. A 10-layer deep CNN was developed and trained using five-fold cross-validation to automatically recognize three different oncoming environments: level-ground, incline staircases, and decline staircases. The environment recognition system achieved 95% overall image classification accuracy. Extending these preliminary findings, our next-generation system focuses on 1) improving image classification accuracy by using larger and more diverse training datasets, and 2) minimizing onboard computational and memory storage requirements using efficient deep CNNs designed for mobile and embedded vision applications. By developing more accurate and efficient prediction of future walking environments, we can achieve more energy-efficient control of robotic lower-limb exoskeletons and prostheses.

Conclusion

We are currently witnessing an explosion of research into the application of artificial intelligence to engineering problems. In the Motion Research Group at the University of Waterloo, we are combining AI tools with dynamic modeling and control theories. This hybrid approach has increased the efficiency, accuracy, and real-time practicality of our engineering solutions, and fostered exciting new projects in the fields of Autonomous Cars and Biomechatronics. We are optimistic that the rapidly-evolving progress in deep learning will further advance the frontiers of dynamics and control of future engineering systems. *– A. Hashemi, Y. Lin, W. McNally, B. Laschowski, B. Hosking, A. Wong, and J. McPhee*

FIG. 3: SCHEMATIC OF THE ENVIRONMENT RECOGNITION SYSTEM INCLUDING DEEP CONVOLUTIONAL NEURAL NETWORK.

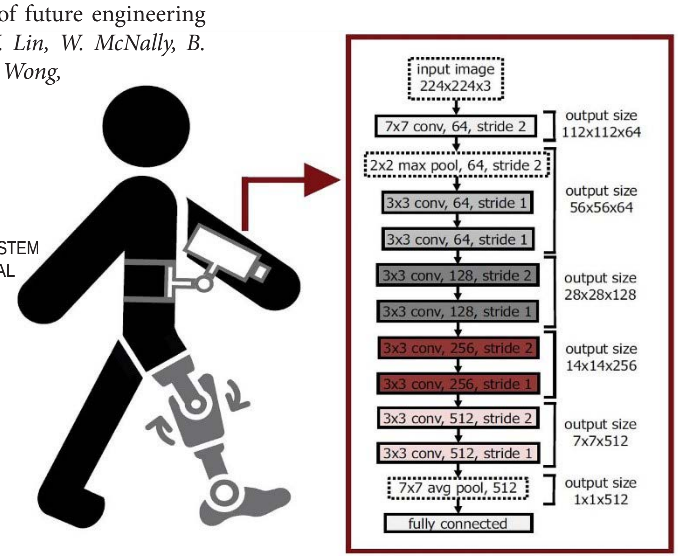

References

[1]Hosking B, McPhee J. (2018). Powertrain Modeling and Model Predictive Longitudinal Dynamics Control for Hybrid Electric Vehicles. SAE Technical Paper 2018-01-0996, Detroit, United States.

[2]J. Rios-Torres and A. A. Malikopoulos, A survey on the coordination of connected and automated vehicles at intersections and merging at highway on-ramps, *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 5, pp. 1066–1077, 2016.

[3]Y. Hu, A. Nakhaei, M. Tomizuka, and K. Fujimura, Interaction-aware decision making with adaptive strategies under merging scenarios, arXiv preprint arXiv:1904.06025, 2019.

[4]D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, et al., Mastering the game of go with deep neural networks and tree search, *Nature*, vol. 529, no. 7587, p. 484, 2016.

[5]T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, Continuous control with deep reinforcement learning, arXiv preprint arXiv:1509.02971, 2015.

[6]Y. Lin, J. McPhee, and N. L. Azad. Decision-Making and Control for Freeway On-Ramp Merging Using Deep Reinforcement Learning, arXiv preprint arXiv:1909.12967, 2019.

[7]M. Sedighizadeh, and A. Rezazadeh. Adaptive PID Controller Based on Reinforcement Learning for Wind Turbine Control. In *Proceedings of world academy of science, engineering and technology* (Vol. 27, pp. 257-262), 2008.

[8]W. McNally, A. Wong and J. McPhee, STAR-Net: Action Recognition using Spatio-Temporal Activation Reprojection, *16th Conference on Computer and Robot Vision* (CRV), Kingston, ON, Canada, 2019, pp. 49-56. doi: 10.1109/CRV.2019.00015

[9]C. Chen, R. Jafari and N. Kehtarnavaz, UTD-MHAD: A multimodal dataset for human action recognition utilizing a depth camera and a wearable inertial sensor, *IEEE International Conference on Image Processing* (ICIP), Quebec City, QC, 2015, pp. 168-172.

[10]Laschowski B, McPhee J, and Andrysek J. (2019). Lower-Limb Prostheses and Exoskeletons with Energy Regeneration: Mechatronic Design and Optimization Review, *ASME Journal of Mechanisms and Robotics*, 11(4), pp. 040801-040801-8. DOI: 10.1115/1.4043460.

[11]Laschowski B, McNally W, McPhee J, and Wong A. (2019). Preliminary Design of an Environment Recognition System for Controlling Robotic Lower-Limb Prostheses and Exoskeletons, *IEEE International Conference on Rehabilitation Robotics*, pp. 868-873. DOI: 10.1109/ICORR.2019.8779540.