by approximately 12%. This type of automated optimization is more time-efficient than a conventional manual tuning approach. Since policy search was conducted over the parameters already present in Ford's production start-stop controller, the updated parametrization obtained by policy search can be directly and immediately incorporated into the start-stop controller calibration process. Prior to deployment, the new parametrization would have to go through a verification process, which would typically include an objective verification of performance targets (such are engine restart time under a wide range of conditions including stopping on an incline) and a subjective verification of customer acceptability (possibly through a customer clinic).

This work also shows that the use of multiple policies could further improve the performance of the start-stop controller. Multiple policies could be implemented via manual driver selection (e.g. high sensitivity versus low sensitivity) or through automatic policy selection. Our analysis found that even with just two selectable policies, the better performing policy can change from stopping event to stopping event. We could thereby gain a further 8% improvement if we were able to accurately identify which could be applied to each stop. Our experiments showed that the right policy to apply at each stop could not be determined based on the driver's behavior *before* the stop; however, they suggested that the driver's behavior *during* the stop might be more useful, thus indicating directions for improving the policy used for start-stop control. We conclude that a feed-forward prediction based on the driver behavior immediately leading up to the stopping event would be best suited to this problem. The key challenge in implementing an adaptive start-stop controller will be finding a reliable feed-forward indicator.

The future development of a method for selecting one from a set of possible policies for use online presents a more complicated path to deployment than using a single optimized policy. The two (or more) policies are simply different sets of parameters that can be swapped out of Ford's existing parametrized start-stop controller, but determining the thresholds for switching between policies will need to be calibrated. For instance, a specialized policy might only be selected over a base policy if the likelihood that it will improve performance is greater than 95%. With this approach, an online adaptive start-stop controller can be deployed with minimal risk.

Acknowledgments

This research was funded by a grant from the Ford-MIT Alliance.

References

Bishop, J.; Nedungadi, A.; Ostrowski, G.; Surampudi, B.; Armiroli, P.; and Taspinar, E. 2007. An engine start/stop system for improved fuel economy. SAE Technical Paper 2007-01-1777.

Bishop, C. 2006. *Pattern Recognition and Machine Learning*. New York, NY: Springer.

Bodik, P.; Armbrust, M.; Canini, K.; Fox, A.; Jordan, M.; and Patterson, D. 2008. A case for adaptive datacenters to conserve energy and improve reliability. Technical Report UCB/EECS-2008-127, University of California at Berkeley.

Csepinszky, A. 2011. Operational results and conclusions of the fot execution phase of eurofot european large scale field operational test. In *18th ITS World Congress*.

Dalamagkidis, K., and Kolokotsa, D. 2008. Reinforcement learning for building environmental control. In Weber, C.; Elshaw, M.; and Mayer, N., eds., *Reinforcement Learning, Theory and Applications*. I-Tech Education and Publishing. 283–294.

Ford Motor Company. 2013. Ford motor company 2013 annual report. http://corporate.ford.com/doc/ar2012-2012%20Annual%20Report.pdf.

Furushou, M.; Nishizawa, K.; Iwasaki, T.; and Tahara, M. 2012. Stop-start system with compact motor generator and newly developed direct injection gasoline engine. SAE Technical Paper 2012-01-0410.

Kohl, N., and Stone, P. 2004. Machine learning for fast quadrapedal locomotion. In *Proceedings of the Nineteenth National Conference on Artificial Intelligence*, 611–616.

Kolter, J.; Jackowski, Z.; and Tedrake, R. 2012. Design, analysis, and learning control of a fully actuated micro wind turbine. In *Proceedings of the American Control Conference*.

Murphey, Y.; Masrur, M.; Chen, Z.; and Zhang, B. 2006. Model-based fault diagnosis in electric drives using machine learning. *IEEE/ASME Transactions on Mechatronics* 11:290–303.

Navigant Research. 2013. Stop-start vehicles. http://www.navigantresearch.com/research/stop-start-vehicles.

Park, J.; Chen, Z.; Kiliaris, L.; Kuang, M.; Masrur, M.; Phillips, A.; and Murphey, Y. 2009. Intelligent vehicle power control based on machine learning of optimal control parameters and prediction of road type and traffic congestion. *IEEE Transactions on Vehicular Technology* 58(9):4741 – 4756.

Prokhorov, D., ed. 2008. *Computational Intelligence in Automative Applications*, volume 132 of *Studies in Computational Intelligence*. Springer.

Rosenfeld, A.; Bareket, Z.; Goldman, C.; Kraus, S.; LeBlanc, D.; and Tsimoni, O. 2012. Learning driver's behavior to improve the acceptance of adaptive cruise control. In *Proceedings of the Twenty-Fourth Innovative Applications of Artificial Intelligence Conference*, 2317–2322.

Sutton, R., and Barto, A. 1998. *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press.

Sutton, R.; McAllester, D.; Singh, S.; and Mansour, Y. 2000. Policy gradient methods for reinforcement learning with function approximation. In *Advances in Neural Information Processing Systems 12*, 1057–1063.

Theocharous, G.; Mannor, S.; Shah, N.; Gandhi, P.; Kveton, B.; Siddiqi, S.; and Yu, C.-H. 2006. Machine learning for adaptive power management. *Intel Technology Journal* 10:299–312.

Vong, C.-M.; Wong, P.-K.; and Li, Y.-P. 2006. Prediction of automative engine power and torque using least squares support vector machines and Bayesian inference. *Engineering Applications of Artificial Intelligence* 19:277–287.

Wang, X.; McGee, R.; and Kuang, M. 2013. Vehicle system control for start-stop powertrains with automatic transmissions. SAE Technical Paper 2013-01-0347.

Williams, R. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning* 8:229–256.