

# Todd Hester

---

## CONTACT INFORMATION

116 NE 58th St  
Seattle, WA 98105

Voice: (508) 816-3495  
E-mail: [todd.hester@gmail.com](mailto:todd.hester@gmail.com)  
Web: [www.cs.utexas.edu/~todd/](http://www.cs.utexas.edu/~todd/)

My work is focused on applying cutting edge machine learning to products that will impact many users in a positive way. I have many years of experience in applying ML to real world problems and doing research to advance the state of the art in applied ML. I drive my teams to find the right fit between product needs and machine learning research to have the most impact possible.

## PROFESSIONAL EXPERIENCE

### **DeepMind**, London, UK; Seattle, WA, USA

*Staff Research Scientist*

**Jan 2016 - Present**

- Lead applied research team of 12+ research scientists and engineers
- Lead multiple research efforts and applied research projects
- Set direction for larger 90 person team as part of small leads group
- Apply DeepMind's machine learning research to Google products in areas such as recommender systems and industrial controls
- Perform research to advance the state of the art in applied machine learning
- Make connections between product needs and research advances.

### **Nest Labs**, Palo Alto, CA, USA

*Algorithms Engineer*

**April 2013 to Dec 2015**

- Led a team of machine learning algorithm engineers
- Developed new machine learning algorithms deployed on the Nest thermostat:
  - Developed Auto-Schedule algorithm for Nest Learning Thermostat
  - Developed thermal model and HVAC control algorithm
  - Developed HVAC optimization based on utility rate plans and demand response events.

### **Science Applications International Corporation**, McLean, VA USA

*Consultant*

**April 2013 to March 2014**

- Developed novel reinforcement learning techniques to be applied for specific application.

### **University of Texas at Austin** Austin, TX USA

*Assistant Instructor, Teaching Assistant*

**Sept 2006 - May 2013**

- Instructor for CS378 (Autonomous Intelligent Robotics) and CS344M (Autonomous Multiagent Systems)
- TA for 3 courses; Outstanding TA Award for CS 393R (Autonomous Robotics)

### **Ace Academy**, Austin, TX USA

*Robotics Instructor*

**September 2008 to May 2010**

- Taught LEGO Mindstorms Robotics to elementary and middle school students.
- Taught ALICE programming to elementary school students.

**Spaulding Rehabilitation Hospital Motion Analysis Lab**, Boston, MA USA

*Research Engineer*

**May 2005 to June 2006**

- Used machine learning techniques to predict clinical scores for Parkinson’s Disease and stroke patients from wearable sensor data.

**Motorola**, Marlboro, MA USA

*Software Support Engineer*

**June 2004 to December 2004**

- Analyzed, reproduced, debugged, and solved customer issues with the BSR64000 broadband cable router.

**Sun Microsystems**, Burlington, MA USA

*Circuit Design Engineer*

**June 2003 to December 2003**

- Worked on circuit design and layout for the UltraSparc V.
- Maintained, tested, and characterized custom latch and flop library.

**Air Force Research Laboratory**, Hanscom AFB, MA USA

*Research Engineer*

**October 2001 to March 2003**

- Designed, developed, and programmed a temperature controller device to drive multiple devices in a hyperspectral imager.

EDUCATION

**The University of Texas at Austin**, Austin, TX USA

Ph.D., Department of Computer Science, December 2012, GPA: 4.0

- Dissertation: *TEXPLORE: Temporal Difference Reinforcement Learning for Robots and Time-Constrained Domains*
- Advisor: Professor Peter Stone

**Northeastern University**, Boston, MA USA

B.S., Computer Engineering, April 2005, GPA: 3.9

- *Summa cum Laude*, with Honors

PUBLICATIONS

**Journal Articles**

- [1] T. Hester and P. Stone, “Intrinsically motivated model learning for developing curious robots,” *Artificial Intelligence*, 2015.
- [2] T. Hester and P. Stone, “TEXPLORE: real-time sample-efficient reinforcement learning for robots,” *Machine Learning*, vol. 90, no. 3, pp. 385–429, 2013.
- [3] S. Patel, R. Hughes, T. Hester, J. Stein, M. Akay, J. G. Dy, and P. Bonato, “A novel approach to monitor rehabilitation outcomes in stroke survivors using wearable technology,” *Proceedings of the IEEE*, vol. 98, no. 3, 2010.
- [4] S. Patel, T. Hester, R. Hughes, N. Huggins, A. Flaherty, D. Standaert, J. Growdon, and P. Bonato, “Processing wearable sensor data to optimize deep-brain stimulation,” *IEEE Pervasive Computing*, vol. 7, no. 1, pp. 56–61, 2008.

**Books**

- [1] T. Hester, *TEXPLORE: Temporal Difference Reinforcement Learning for Robots*. Studies in Computational Intelligence, Springer-Verlag, 2013.

## Book Chapters

- [1] S. Barrett, K. Genter, Y. He, T. Hester, P. Khandelwal, J. Menashe, and P. Stone, “UT Austin Villa 2012: Standard platform league world champions,” in *RoboCup 2012: Robot Soccer World Cup XVI* (X. Chen, P. Stone, L. E. Sucar, and T. V. der Zant, eds.), Lecture Notes in Artificial Intelligence, Springer Verlag, 2013.
- [2] T. Hester and P. Stone, “Learning and using models,” in *Reinforcement Learning: State of the Art* (M. Wiering and M. van Otterlo, eds.), Berlin, Germany: Springer Verlag, 2011.
- [3] P. Stone, M. Quinlan, and T. Hester, “The essence of soccer: Can robots play too?,” in *Soccer and Philosophy*, vol. 51 of *Popular Culture and Philosophy*, Open Court Publishing Company, 2010.
- [4] S. Kalyanakrishnan, T. Hester, M. Quinlan, Y. Bentor, and P. Stone, “Three humanoid soccer platforms: Comparison and synthesis,” in *RoboCup 2009: Robot Soccer World Cup XIII* (J. Baltes, M. G. Lagoudakis, T. Naruse, and S. S. Ghidary, eds.), pp. 140–152, Springer Verlag, 2010.

## Refereed Conference Papers

- [1] M. Vecerík, O. Sushkov, D. Barker, T. Rothörl, T. Hester, and J. Scholz, “A practical approach to insertion with variable socket position using deep reinforcement learning,” in *IEEE International Conference on Robotics and Automation (ICRA)*, May 2019.
- [2] T. Hester, M. Vecerík, O. Pietquin, M. Lanctot, T. Schaul, B. Piot, A. Sendonaris, G. Dulac-Arnold, I. Osband, J. Agapiou, J. Z. Leibo, and A. Gruslys, “Deep q-learning from demonstrations,” in *Association for the Advancement of Artificial Intelligence (AAAI)*, Feb 2018.
- [3] T. Hester, M. Lopes, and P. Stone, “Learning exploration strategies in model-based reinforcement learning,” in *The Twelfth International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2013.
- [4] T. Hester and P. Stone, “Intrinsically motivated model learning for a developing curious agent,” in *The Eleventh International Conference on Development and Learning (ICDL)*, Nov 2012. Paper of Excellence Award.
- [5] T. Hester, M. Quinlan, and P. Stone, “RTMBA: A real-time model-based reinforcement learning architecture for robot control,” in *IEEE International Conference on Robotics and Automation (ICRA)*, May 2012.
- [6] T. Hester and P. Stone, “Real time targeted exploration in large domains,” in *The Ninth International Conference on Development and Learning (ICDL)*, August 2010.
- [7] T. Hester, M. Quinlan, and P. Stone, “Generalized model learning for reinforcement learning on a humanoid robot,” in *IEEE International Conference on Robotics and Automation (ICRA)*, May 2010.
- [8] T. Hester and P. Stone, “Generalized model learning for reinforcement learning in factored domains,” in *The Eighth International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2009.
- [9] T. Hester and P. Stone, “Negative information and line observations for monte carlo localization,” in *IEEE International Conference on Robotics and Automation (ICRA)*, May 2008.

- [10] N. K. Jong, T. Hester, and P. Stone, “The utility of temporal abstraction in reinforcement learning,” in *The Seventh International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2008.
- [11] P. Boissy, T. Hester, D. M. Sherrill, H. Corriveau, and P. Bonato, “Monitoring mobility assistive device use in post-stroke patients,” in *29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBS)*, 2007.
- [12] S. Patel, T. Hester, R. Hughes, N. Huggins, D. Standaert, A. Flaherty, and P. Bonato, “Using wearable sensors to enhance DBS parameter adjustment for Parkinson’s disease patients through measures of motor response,” in *3rd IEEE/EMBS International Summer School on Medical Devices and Biosensors*, pp. 141–144, 2006.
- [13] T. Hester, D. M. Sherrill, M. Hamel, K. Perreault, P. Boissy, and P. Bonato, “Identification of tasks performed by stroke patients using a mobility assistive device,” in *28th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBS)*, pp. 1501–1504, 2006.
- [14] P. Boissy, T. Hester, D. M. Sherrill, H. Corriveau, and P. Bonato, “Monitoring mobility assistive device use in patients after stroke,” in *16th Congress of the International Society of Electrophysiology and Kinesiology (ISEK)*, 2006.
- [15] T. Hester, R. Hughes, D. M. Sherrill, S. Patel, N. Huggins, A. Flaherty, and D. Standaert, “Adjusting DBS settings to optimize Parkinson’s control therapy,” in *16th Congress of the International Society of Electrophysiology and Kinesiology (ISEK)*, 2006.

#### Refereed Workshop Papers

- [1] G. Dulac-Arnold, D. J. Mankowitz, and T. Hester, “Challenges of real-world reinforcement learning,” in *ICML Workshop on Reinforcement Learning for Real Life (RLRL)*, 2019.
- [2] T. Hester and P. Stone, “Intrinsically motivated model learning for a developing curious agent,” in *AAMAS Adaptive Learning Agents (ALA) Workshop*, June 2012.
- [3] T. Hester and P. Stone, “TEXPLORE: Real-time sample-efficient reinforcement learning for robots,” in *Proceedings of the AAAI Spring Symposium on Designing Intelligent Robots: Re-integrating AI*, March 2012.
- [4] S. Barrett, K. Genter, T. Hester, M. Quinlan, and P. Stone, “Controlled kicking under uncertainty,” in *Proceedings of the 5th Humanoids Workshop on Humanoid Soccer Robots*, Dec 2010.
- [5] T. Hester and P. Stone, “An empirical comparison of abstraction in models of Markov Decision Processes,” in *Proceedings of the ICML/UAI/COLT Workshop on Abstraction in Reinforcement Learning*, June 2009.
- [6] T. Hester, D. M. Sherrill, P. Bonato, M. Hamel, K. Perreault, and P. Boissy, “Using wearable sensors to analyze the quality of use of mobility assistive devices,” in *Third Annual Workshop on Wearable and Implantable Body Sensor Networks (BSN)*, pp. 127–130, April 2006.
- [7] T. Hester, R. Hughes, D. M. Sherrill, B. Knorr, M. Akay, J. Stein, and P. Bonato, “Using wearable sensors to measure motor abilities following stroke,” in *Third Annual Workshop on Wearable and Implantable Body Sensor Networks (BSN)*, pp. 5–8, April 2006.

- [8] S. Patel, D. Sherrill, R. Hughes, T. Hester, T. Lie-Nemeth, P. Bonato, D. Standaert, and N. Huggins, “Analysis of the severity of dyskinesia in patients with Parkinson’s disease via wearable sensors,” in *Third Annual Workshop on Wearable and Implantable Body Sensor Networks (BSN)*, pp. 123–126, April 2006.

### Patents

- [1] T. Hester, A. J. Minich, and G. A. Heitz III, “Enhanced automated environmental control system scheduling using a preference function,” Apr. 3 2018. US Patent 9,933,177.
- [2] T. A. Hester, E. J. Fisher, and P. Khandelwal, “Predictively controlling an environmental control system,” Jan. 16 2018. US Patent 9,869,484.
- [3] T. A. Hester, A. J. Minich, and G. A. Heitz III, “Enhanced automated control scheduling,” Sept. 26 2017. US Patent 9,772,116.
- [4] S. Y. Shafi, T. Hester, J. Ben-Meshulam, and S. R. Dey, “Identification of similar users,” Sept. 5 2017. US Patent 9,756,478.
- [5] M. R. Malhotra, S. Le Guen, J. A. Boyd, J. T. Lee, and T. Hester, “Learned overrides for home security,” Dec. 13 2016. US Patent 9,520,049.

### Pending Patents

- [1] M. R. Malhotra, S. Le Guen, J. A. Boyd, J. T. Lee, and T. Hester, “Operating a security system,” Mar. 5 2019. US Patent App. 10/223,896.
- [2] R. A. Evans, J. Gao, M. C. Ryan, G. Dulac-Arnold, J. K. Scholz, and T. A. Hester, “Optimizing data center controls using neural networks,” July 19 2018. US Patent App. 15/410,547.
- [3] J. Crimins, S. Ruffner, A. Minich, T. Hester, and A. Sahl, “Thermostat algorithms and architecture for efficient operation at low temperatures,” Apr. 12 2018. US Patent App. 15/286,564.
- [4] W. Greene, S. McGaraghan, J. Crimins, S. Ruffner, A. Minich, T. Hester, A. Sahl, and P. Subramani, “Architecture for thermostat control during time-of-use intervals,” Dec. 21 2017. US Patent App. 15/187,562.
- [5] J. Crimins, S. Ruffner, A. Minich, T. Hester, A. Sahl, and P. Subramani, “Architecture for thermostat control during peak intervals,” Dec. 21 2017. US Patent App. 15/187,565.
- [6] I. Karp, L. Stesin, C. Pi-Sunyer, M. A. McBride, A. Dubman, J. Lyons, S. W. Kortz, G. J. Hu, A. Surya, A. Thelen, *et al.*, “Methods and apparatus for using smart environment devices via application program interfaces,” July 6 2017. US Patent App. 15/380,767.
- [7] P. Verhoeven and T. Hester, “Coordinating energy use of disparately-controlled devices in the smart home based on near-term predicted hvac control trajectories,” Apr. 13 2017. US Patent App. 14/881,807.
- [8] P. Verhoeven and T. Hester, “Persistent home thermal comfort model reusable across multiple sensor and device configurations in a smart home,” Feb. 23 2017. US Patent App. 14/832,675.

- [9] P. Verhoeven and T. Hester, “Persistent thermal model and method of using same for automatically determining the presence of an additional thermal source other than the hvac system being controlled,” Feb. 23 2017. US Patent App. 14/832,702.
- [10] P. P. Reddy, M. Malhotra, E. J. Fisher, T. Hester, M. A. McBride, and Y. Matsuoka, “Intelligent configuration of a smart environment based on arrival time,” Dec. 24 2015. US Patent App. 14/531,805.

### Technical Reports

- [1] D. J. Mankowitz, N. Levine, R. Jeong, A. Abdolmaleki, J. T. Springenberg, T. A. Mann, T. Hester, and M. A. Riedmiller, “Robust reinforcement learning for continuous control with model misspecification,” *CoRR*, vol. abs/1906.07516, 2019.
- [2] T. Pohlen, B. Piot, T. Hester, M. G. Azar, D. Horgan, D. Budden, G. Barth-Maron, H. van Hasselt, J. Quan, M. Vecerík, M. Hessel, R. Munos, and O. Pietquin, “Observe and look further: Achieving consistent performance on atari,” *CoRR*, vol. abs/1805.11593, 2018.
- [3] G. Dalal, K. Dvijotham, M. Vecerík, T. Hester, C. Paduraru, and Y. Tassa, “Safe exploration in continuous action spaces,” *CoRR*, vol. abs/1801.08757, 2018.
- [4] M. Vecerík, T. Hester, J. Scholz, F. Wang, O. Pietquin, B. Piot, N. Heess, T. Rothörl, T. Lampe, and M. A. Riedmiller, “Leveraging demonstrations for deep reinforcement learning on robotics problems with sparse rewards,” *CoRR*, vol. abs/1707.08817, 2017.
- [5] T. A. Mann, H. Penedones, S. Mannor, and T. Hester, “Adaptive lambda least-squares temporal difference learning,” *CoRR*, vol. abs/1612.09465, 2016.
- [6] Nest Labs, “Thermal model and hvac control white paper,” Nov 2015.
- [7] Nest Labs, “Enhanced auto-schedule,” Nov 2014.
- [8] T. Hester, *TEXPLORE: Temporal Difference Reinforcement Learning for Robots and Time-Constrained Domains*. PhD thesis, Department of Computer Science, University of Texas at Austin, Austin, TX, December 2012.
- [9] S. Barrett, K. Genter, T. Hester, P. Khandelwal, M. Quinlan, P. Stone, and M. Sridharan, “Austin Villa 2011: Sharing is caring: Better awareness through information sharing,” Tech. Rep. UT-AI-TR-12-01, The University of Texas at Austin, Department of Computer Sciences, AI Laboratory, January 2012.
- [10] S. Barrett, K. Genter, M. Hausknecht, T. Hester, P. Khandelwal, J. Lee, M. Quinlan, A. Tian, P. Stone, and M. Sridharan, “Austin Villa 2010 standard platform team report,” Tech. Rep. UT-AI-TR-11-01, The University of Texas at Austin, Department of Computer Sciences, AI Laboratory, January 2011.
- [11] T. Hester, M. Quinlan, P. Stone, and M. Sridharan, “UT Austin Villa 2009: Naos across Texas,” Tech. Rep. UT-AI-TR-09-08, The University of Texas at Austin, Department of Computer Sciences, AI Laboratory, December 2009.
- [12] T. Hester, M. Quinlan, and P. Stone, “UT Austin Villa 2008: Standing on Two Legs,” Tech. Rep. UT-AI-TR-08-8, The University of Texas at Austin, Department of Computer Sciences, AI Laboratory, November 2008.