George Konidaris

http://cs.brown.edu/people/gdk
gdk@brown.edu

Academic Employment

- **Brown University**
  Associate Professor, Department of Computer Science — July 2022–present.
  John E. Savage Assistant Professor of Computer Science — July 2019–June 2022.
  Assistant Professor, Department of Computer Science — Sep. 2016–June 2019.

- **Duke University** — Sep. 2014–Sep, 2016.
  Assistant Professor, Departments of Computer Science & Electrical and Computer Engineering.

- **MIT Computer Science and Artificial Intelligence Laboratory** — Jan. 2011–Aug. 2014.
  Postdoctoral research with Professors Leslie Kaelbling and Tomas Lozano-Perez.

Education

- **Doctor of Philosophy**, Computer Science. — Sep. 2004–Dec. 2010.
  Computer Science Department, University of Massachusetts Amherst.
  Dissertation: *Autonomous Robot Skill Acquisition*, advisor: Prof. Andrew G. Barto.

- **Master of Science**, Artificial Intelligence (with distinction). — Sep. 2002–Sep. 2003.
  School of Informatics, University of Edinburgh.
  Dissertation: *Behaviour-Based Reinforcement Learning*, supervisor: Dr. Gillian Hayes.

- **Bachelor of Science with Honours**, Computer Science (with distinction). — Jan. 2001–Dec. 2001.
  School of Computer Science, University of the Witwatersrand.
  Research Report: *Axial Line Placement in Deformed Urban Grids*, supervisor: Prof. Ian Sanders.

- **Bachelor of Science**. — Jan. 1998–Dec. 2000.
  Computer Science and Computational & Applied Mathematics (with distinction).
  University of the Witwatersrand.

Industrial Experience

- **Co-founder, Realtime Robotics**
  Chief Scientific Advisor — Jan. 2024–present.
  Chief Roboticist — March 2016–Dec. 2023.
  Co-founded a startup to commercialize our invention of a specialized robot motion planning processor capable of sub-millisecond planning (see publications 66, 81, and 84).

- **Co-founder, Lelapa AI** — February 2023–present.
  Co-founded a commercial AI lab to develop new AI technologies by and for Africans.

Teaching Experience

- **Brown University**
  - Artificial Intelligence (CSCI 1410) — Spring 2017, Fall 2017–2019, Fall 2021–2022.
  - Reintegrating AI (CSCI 2951X) — Spring 2018, Spring 2020–2021, Fall 2021, Spring 2023–2024.

- **Duke University**
  - Hierarchical Reinforcement Learning and Planning (CPS 590.2) Fall 2014.
  - Introduction to Artificial Intelligence (CPS 270) Spring 2015, 2016.
  - Decision Making for Robotics and Autonomous Systems (CPS 590.1) Fall 2015.

- **Guest Lectures**
  - *Robot Motion Planning on a Chip* December 3rd 2020.
    Cornell University, course ECE 4960 (Fast Robots).
  - *Reinforcement Learning for Robotics* May 7th 2013.
    Massachusetts Institute of Technology, course 6.S064 (Introduction to Machine Learning).
  - *Hierarchical Reinforcement Learning*
    Brown University, course CSCI2951-F (Learning and Sequential Decision Making). Nov. 5th 2012.
  - *Reinforcement Learning* Dec. 1st 2011,
    Massachusetts Institute of Technology, course 6.867 (Machine Learning). Nov. 27th 2012.

- **External Author and Examiner** Jun. 2008–Jul. 2014
  University of London, International Programmes. Developed subject guide and accompanying CD for undergraduate Artificial Intelligence (2910310/CO3310) course, and set and marked projects and final exams. This course is taken by approximately 40 correspondence students every year from around the world, who obtain credit toward a University of London degree.

- **Curriculum Design** Jun. 2011–August 2011
  MIT-Singapore Alliance. Collaborated with the MIT Mechanical Engineering department to design an introductory course (and associated course notes) on numerical programming in Python. A Matlab version is now being used to teach Numerical Computation for Mechanical Engineers (2.086) at MIT.

- **Teaching Assistant** (7 semesters) Sep. 2004–Dec. 2010
  Department of Computer Science, University of Massachusetts Amherst. Graded written and programming assignments, conducted discussion sections and held office consultations for courses ranging from introductory Java programming to graduate-level algorithms.

- **Teaching Assistant** Jan.– Jul. 2002
  School of Computer Science, University of the Witwatersrand, Johannesburg. Co-lectured Basic Computer Organisation (first semester, first year of Computer Science) to approximately 180 students, and tutored Fundamental Algorithmic Concepts (second semester, first year) to a class of approximately 120 students.

- **Tutor** Jan.–Nov. 2001
  School of Computer Science, University of the Witwatersrand, Johannesburg. Tutored and marked tests and assignments for first-year undergraduate courses.

Postgraduate Research Student Supervision

At Brown University:

PhD Dissertations

- Matthew Corsaro, *Learning Task-Specific Grasps.* May 2023
- Ben Abbatematteo, *Exploiting Structure for Efficient Robotic Manipulation* October 2023
- Cameron Allen, *Structured Abstractions for General-Purpose Decision Making* October 2023
- Eric Rosen, *Abstraction for Autonomous Human-Robot Interaction* February 2024
  (Co-advised with Prof. Stefanie Tellex.)
- Akhil Bagaria, *Skill Discovery for Exploration and Planning* August 2024

At Duke University:

PhD Dissertations

- Benjamin Burchfiel, *3D Object Representations for Robot Perception.*   Sep 2019

- Barrett Ames, *Towards Efficient and Robust Robot Planning.*   Oct 2022

At the University of the Witwatersrand:

PhD Dissertations

- Michael Mitchley, *Adaptive Value Function Approximation in Reinforcement Learning.*   Dec 2015
  using Wavelets. (Co-advised with Prof. Ebrahim Momoniat)

- Dean Wookey, *Representation Discovery Using a Fixed Basis in Reinforcement Learning.*   Aug 2016
  (Co-advised with Prof. Clint van Alten)

- Pravesh Ranchod, *Skill Discovery From Multiple Related Demonstrators.*   Dec 2017
  (Co-advised with Prof. Clint van Alten and Dr. Benjamin Rosman)

- Steven James, *Learning Portable Symbolic Representations.*   Sep 2021
  (Co-advised with Prof. Benjamin Rosman.)

MSc Dissertations

- Warwick Masson, *Reinforcement Learning with Parameterized Actions.*   Aug 2016
  (Co-supervised with Dr. Pravesh Ranchod)

- Steven James, *The Effect of Simulation Bias on Action Selection in Monte Carlo Tree Search.*   Aug 2016
  (Co-supervised with Dr. Pravesh Ranchod and Dr. Benjamin Rosman)

- Craig Bester, *Multi-Pass Deep Q-Networks for Reinforcement Learning with*   March 2019
  Parameterised Action Spaces. (Co-advised with Mr. Steven James and Dr. Pravesh Ranchod.)

Invited Talks

Conferences

- *Signal to Symbol (via Skills).* Third Conference on Robot Learning, Osaka, Japan, November 18th 2019.

- *Robot Skill Acquisition.* Third EUCogIII (European Network for the Advancement of Artificial Cognitive Systems, Interaction and Robotics) Members Conference, Palma de Mallorca, April 10th 2013.

Workshops

- *Reintegrating AI: Skills, Symbols, and the Sensorimotor Dilemma.* IROS 2023 Workshop on Task and Motion Planning: from Theory to Practice. October 5th 2023.

- *Reintegrating AI: Skills, Symbols, and the Sensorimotor Dilemma.* ICRA 2023 Workshop on on Effective Representations, Abstractions, and Priors for Robot Learning. May 29th 2023.

- *Symbolic Foundations for Planning and Communication.* ICRA 2023 Workshop on Life-Long Learning with Human Help. May 29th 2023.

- *Reintegrating AI: Skills, Symbols, and the Sensorimotor Dilemma.* CORL 2022 Workshop on Aligning Robot Representations with Humans, December 15th 2022.

- *Learned Abstractions for Intelligent Behavior.* CORL 2022 Workshop on Learning, Perception, and Abstraction for Long-Horizon Planning, December 15th 2022.

- *Reintegrating AI: Skills, Symbols, and the Sensorimotor Dilemma.* IROS 2022 Workshop on Life-long Learning of High-level Cognitive and Reasoning Skills, October 23rd 2022.

- *Reintegrating AI: Skills, Symbols, and the Sensorimotor Dilemma.* RSS 2022 2nd Workshop on Scaling Robot Learning, June 27th 2022.

- *Realtime Robotics: From RSS to Factory Floor.* ICRA 2022 WS on Challenges in Applying Academic Research to Real World Robotics, May 23rd 2022.

- *Signal to Symbol (via Skills).* NeurIPS 2021 Deep Reinforcement Learning Workshop, December 13th 2021.

- *Signal to Symbol (via Skills).* 5th Workshop on Semantic Policy and Action Representations for Autonomous Robots, at IROS, September 27th 2021.

- *Signal to Symbol (via Skills).* IndabaX South Africa, September 17th 2021.

- *Grounded Symbolic Representations for Human-Robot Communication.* ICDL 2021 Workshop on Human-Aligned Reinforcement Learning for Autonomous Agents and Robots, August 27th 2021.

- *Signal to Symbol (via Skills).* RSS 2021 Workshop on Declarative and Neurosymbolic Representations in Robot Learning and Control, July 15th 2021.

- *Learning Plannable Task Models.* RSS 2021 Workshop on Integrating Planning and Learning. July 12th 2021.

- *Signal to Symbol (via Skills).* NeurIPS 2020 Workshop on Biological and Artificial Reinforcement Learning, December 12th 2020.

- *Learning Portable Skills and Symbols.* AAAI 2020 Workshop on Generalization in Planning, New York City, February 7th 2020.

- *Learning Abstract Models for Symbolic High-Level Planning.* Third International Workshop on Intrinsically-Motivated Open-Ended Learning, Rome, Italy, October 5th 2017.

- *Robots, Skills, and Symbols.* IROS 2017 Second Workshop on Machine Learning Methods for High-Level Cognitive Capabilities in Robotics, Vancouver Canada, September 28th 2017.

- *High-Level Symbolic Representations for Planning.* RSS 2017 Workshop on: The What Without the How: Specifying Planning Problems in Robotics, Cambridge MA, July 15th 2017.

- *Combining State and Temporal Abstraction.* ICML 2016 Workshop on Abstraction in RL, New York City, June 23rd 2016.

- *What Are Representations For?* RSS 2016 Workshop on Geometry and Beyond—Representations, Physics, and Scene Understanding for Robotics, Ann Arbor, Michigan, June 19th 2016.

- *Avoiding Learning by Exploiting Structure.* ICRA 2016 Workshop on Nature Versus Nurture in Robotics, May 20th 2016.

- *Robots, Skills, and Symbols.* ICRA 2016 Workshop on Task-driven Perceptual Representations: Sensing, Planning and Control under Resource Constraints, Stockholm, May 16th 2016.

- *Learning Symbolic Representations for Planning.* RSS 2015 Workshop on Learning Reusable Concepts in Robotics, Rome, July 16th 2015.

- *Robots, Skills, and Symbols.* NIPS 2014 Workshop on Autonomously Learning Robots, Montreal, December 12th 2014.

- *Robots, Skills, and Symbols.* IJCAI 2013 Workshop on Machine Learning for Interactive Systems: Bridging the Gap between Perception, Action and Communication, Beijing, August 4th 2013.

- *Robots, Skills, and Symbols.* Schloss Dagstuhl Seminar on Mechanisms of Ongoing Development in Cognitive Robotics, February 13th 2013.

Universities and Research Laboratories

- MIT. *Reintegrating AI: Skills, Symbols, and the Sensorimotor Dilemma*, Center of Brains, Minds, and Machines (Quest Seminar Series), October 18th 2022.

- Navy Center for Applied Research in Artificial Intelligence. *Signal to Symbol (via Skills)*, NCARAI Seminar, December 7th 2020.

- Carnegie Mellon University. *Signal to Symbol (via Skills)*, Robotics Institute Seminar, November 16th 2018.

- MIT. *Robot Motion Planning on a Chip*, Robotics Seminar, November 1st 2016.

- University of Pennsylvania. *Robots, Skills, and Symbols*, April 15th 2016.

- Indiana University. *Robots, Skills, and Symbols*, December 4th 2015.

- MIT. *Robots, Skills, and Symbols*, October 24th 2014.

- Brown University. *Robots, Skills, and Symbols*, October 23rd 2014.

- Texas A&M. *Robots, Skills, and Symbols*, October 6th 2014.

- University of Texas at Austin. *Robots, Skills, and Symbols*, October 3rd 2014.

- University of Michigan. *Robots, Skills, and Symbols*, September 19th 2014.

- Oregon State University. *Robots, Skills, and Symbols*, August 27th 2014.

- Harvard University. *Autonomous Robot Skill Acquisition*, Machine Learning Tea, September 19th 2012.

- University College London, Gatsby Computational Neuroscience Unit. *Autonomous Robot Skill Acquisition*, seminar, December 9th 2011.

- Rutgers. *Autonomous Robot Skill Acquisition*, DCS Colloquium, December 7th 2011.

- MIT. *Skill Acquisition in Continuous Reinforcement Learning Domains*, July 27th 2010.

- Brown University. *Skill Acquisition in Continuous Reinforcement Learning Domains*, July 26th 2010.

- Williams College. *Toward Autonomous Robot Skill Acquisition*, CS Colloquium, November 21st 2008.

- Rutgers. *Toward Autonomous Robot Skill Acquisition*, DCS Colloquium, September 22nd 2008.

- University of Massachusetts Amherst. *Agent Space vs. Problem Space: Knowledge and Skill Transfer in Reinforcement Learning*, Machine Learning and Friends Lunch, December 1st 2005; *Sensorimotor Abstraction Selection for Autonomous Robot Skill Acquisition*. Machine Learning and Friends Lunch, 30th April 2008.

- University of Edinburgh. *Behavior-Based Reinforcement Learning*, December 4th 2003.

- University of the Witwatersrand. *Behavior-Based Reinforcement Learning*, 5th October 2003.

- University of Bath. *Behavior-Based Reinforcement Learning*, BAI Summer Seminar Series, September 23rd 2003.

Academic Service and Memberships

- Action Editor, Journal of Machine Learning Research, November 2015—present.

- Action Editor, Journal of Artificial Intelligence Research, October 2017—September 2020.

- Advisory Board, Deep Learning Indaba, September 2017—present.

- Local Co-Chair (with Michael Frank), The Multi-disciplinary Conference on Reinforcement Learning and Decision Making, 2022.

- Area Co-Chair (with Ifat Levy), The Multi-disciplinary Conference on Reinforcement Learning and Decision Making, 2019.

- Robotics Chair (with Prof. Gregory Dudek, Dr. Brad Knox, and Dr. Nick Hawes), AAAI 2014–2016.

- Co-organizer, co-founder (with Stefanie Tellex, Matt Walter and Brian Scassellati), and steering committee, *Northeast Robotics Colloquium* (NERC), 2012—present.

- Co-organizer (with Lorenzo Riano, Alessandro Saffioti, Nick Hawes, Siddharth Srivastava, and Moritz Tenorth), *IROS 2014 Workshop on AI and Robotics*, September 14th 2014.

- Co-organizer (with Alessandro Saffioti, Nick Hawes, and Moritz Tenorth), *AAAI 2014 Workshop on AI and Robotics*, July 27–28th 2014.

- Co-organizer (with Byron Boots, Nick Hawes, Todd Hester, Tekin Meriçli, Lorenzo Riano, Benjamin Rosman and Peter Stone), *AAAI 2013 Workshop on Intelligent Robotic Systems*, July 14–15th 2013.

- Co-organizer (with Gerhard Neumann, Freek Stulp, and Jan Peters), *RSS Workshop on Hierarchical and Structured Learning for Robotics*, June 28th, 2013.

- Co-organizer (with Byron Boots, Nick Hawes, Todd Hester, Bhaskara Marthi, Lorenzo Riano and Benjamin Rosman), *Designing Intelligent Robots: Reintegrating AI II*, AAAI 2013 Spring Symposium.

- Co-organizer (with Byron Boots, Stephen Hart, Todd Hester, Sarah Osentoski and David Wingate), *Designing Intelligent Robots: Reintegrating AI*, AAAI 2012 Spring Symposium.

- Co-organizer (with Özgür Şimşek), *Abstraction in Reinforcement Learning* ICML/UAI/COLT 2009 Workshop.

- Journal Reviewing:

  ○ Journal of Machine Learning Research (2006, 2007, 2008, 2010, 2012, 2013, 2015).
  ○ Journal of Artificial Intelligence Research (2006, 2007, 2011, 2012, 2013, 2014, 2015, 2016, 2020, 2021).
  ○ International Journal of Robotics Research (2013, 2016, 2017, 2018).
  ○ Autonomous Robots (2015).
  ○ Artificial Intelligence (2012, 2013, 2014, 2015).
  ○ IEEE Transactions on Robotics (2005, 2013, 2014, 2015, 2017).
  ○ Machine Learning (2009, 2011).
  ○ Robotics and Autonomous Systems (2010, 2011).
  ○ IEEE Transactions on Autonomous Mental Development (2010).
  ○ Computational Intelligence (2011, 2012, 2013).
  ○ IEEE Transactions on Neural Networks (2009, 2010).
  ○ South African Computer Journal (2011, 2012).

- Conference Reviewing:

  ○ Neural Information Processing Systems (2011, 2012, 2013, 2014, 2015, 2020).
  ○ International Conference on Machine Learning (2009, 2012, 2013, 2014, 2015, 2017, 2018, 2019).
  ○ Robotics: Science and Systems (2013, 2014, 2015, 2017, 2018, 2019, 2020).
  ○ International Joint Conference on Artificial Intelligence (2009, 2011, 2013, 2015, 2016).
  ○ AAAI Conference on Artificial Intelligence (2010, 2013, 2016, 2017).
  ○ IEEE International Conference on Robotics and Automation (2010, 2011, 2012, 2013, 2015, 2016).
  ○ IEEE/RSJ International Conference on Intelligent Robots and Systems (2012, 2013, 2014, 2015, 2016).
  ○ International Conference on Automated Planning and Scheduling (2015).
  ○ ACM/IEEE International Conference on Human-Robot Interaction (2010).
  ○ IEEE Conference on Development and Learning (2010, 2011, 2012).
  ○ North East Student Colloquium on Artificial Intelligence (2006, 2007, 2008).
  ○ SAICSIT Annual Research Conference (2009, 2010, 2011, 2012, 2013).

- Senior Program Committee Member:

- International Joint Conference on Artificial Intelligence (2013, 2015, 2017, 2018).
  - AAAI Conference on Artificial Intelligence (2017, 2018).
  - International Conference on Robot Learning (2017, 2018)
  - IEEE/RSJ International Conference on Intelligent Robots and Systems (2017)
  - IEEE International Conference on Robotics and Automation (2017, 2018)
  - International Conference on Learning Representations (2019)
- Membership of Professional Societies:
  - Association for the Advancement of Artificial Intelligence (since 2006).
  - International Society for Adaptive Behavior (2004–2007).
  - South African Institute for Computer Scientists and Information Technologists (since 2001).

Grants and Funding

- *REPRISM: Flexible Embodied Problem-Solving by Manipulating the Representational Prism.* Office of Naval Research MURI. Grant number N000142412603, award amount: $7,462,073.

- *Designing the Sensorimotor Substrate.* Boston Dynamics AI Institute. Award amount: $1,500,000.

- *Establishing a Unifying Artificial Intelligence Model for Generally-Capable Robots.* Office of Naval Research (with Co-PI Stefanie Tellex). Grant number N00014-22-1-2592, award amount: $3,359,988.

- *Humanoid Robotics Platform for Investigating the Sensorimotor Basis of Compositional Reasoning.* Air Force Office of Scientific Research (DURIP Program) (with Co-PI Stefanie Tellex). Grant number FA9550-21-1-0308, award amount: $391,761.

- *Skill Learning for Goal-Directed Behavior by Humanoid Robots in Complex Unstructured Environments.* Office of Naval Research (with PI Stefanie Tellex). Grant number N00014-21-1-2584, award amount: $1,913,500.

- *Decision-Theoretic Sequential Decision Making for Observer-Aware, Goal-Directed Behavior in Swarms.* Office of Naval Research. Grant number N00014-21-1-2200, award amount: $405,000.

- *Learning Task-Specific Representations for Broadly Capable Reinforcement Learning Agents.* National Science Foundation (with Co-PI Professor Michael Littman). Grant number 1955361, award amount: $1,199,684.

- *CAREER: Learning Symbolic Representations for Robot Manipulation.* National Science Foundation. Grant number 1844960, award amount: $549,988.

- *Online, Incremental Skill and Representation Acquisition for Lifelong Learning.* DARPA Seedling. Grant number W911NF1820268, award amount: $360,890.

- *Hidden Parameter Markov Decision Processes: Exploiting Structure in Families of Tasks.* National Science Foundation (with PI Professor Finale Doshi-Velez). Grant number 1717569, award amount: $450,000 (Brown University share $208,000).

- *Constructing Abstract Hierarchies for Robust, Real-Time Control.* AFOSR Young Investigator Program. Grant number FA9550-17-1-0124, award amount: $356,861.

- *Bridging the Gap Between Low-Level Robot Control and Flexible High-Level Task Planning.* DARPA Young Faculty Award. Grant number D15AP00104, award amount: $492,826 + $64,127 *Director's Fellowship* (additional third year).

- *Low-Power, Real-Time Motion Planning for Complex Robots in Unstructured Environments.* Robotics Fast Track Grant (with Co-PI Professor Daniel Sorin), award amount: $99,242.28.

- *CRCNS: Representational Foundations of Adaptive Behavior in Natural and Artificial Agents.* National Institutes of Health (with PI Professor Matthew Botvinick and Co-PI Professor Samuel Gershman). Grant number 1R01MH109177-01, award amount: $1,183,417.

- *Robotics Activities at Association for the Advancement of Artificial Intelligence (AAAI) 2016.* National Science Foundation. Award number 1600043, award amount: $17,500.00.

Honors and Awards

- Amazon Research Award, for *Learning Composable Manipulation Skills*. July 2022.

- IROS 2021 Best RoboCup Paper Award, for *Multi-Resolution POMDP Planning for Multi-Object Search in 3D*. Kaiyu Zheng, Yoonchang Sung, George Konidaris, and Stefanie Tellex, September 2021.

- The Karen T. Romer Prize for Undergraduate Advising and Mentoring, Brown University, 2021.

- The IJCAI-JAIR Best Paper Prize, for *From Skills to Symbols: Learning Symbolic Representations for Abstract High-Level Planning*. George Konidaris, Leslie Kaelbling, and Tomas Lozano-Perez, 2020.

- NSF CAREER Award, 2019.

- Best Paper nomination at the 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems, for *Bounded-Error LQR-Trees*. Barrett Ames and George Konidaris. November 2019.

- Richard B. Salomon Faculty Research Award, Brown University, 2018.

- DARPA Director's Fellowship, 2017.

- AFOSR Young Investigator Program Award (class of 2017).

- DARPA Young Faculty Award (class of 2015).

- Best Paper nomination at Robotics: Science and Systems, for *Policy Search for Multi-Robot Coordination under Uncertainty*. Christopher Amato, George Konidaris, Ariel Anders, Gabriel Cruz, Jonathan P. How and Leslie P. Kaelbling. July 2015.

- MIT Intelligence Initiative (I^2) Postdoctoral Fellowship, April 2013.

- Best Student Video, AAAI 2011 Video Competition, for *Autonomous Robot Skill Acquisition*, Scott Kuindersma and George Konidaris. August 2011.

- The Rank Xerox Prize for the best Artificial Intelligence MSc dissertation, University of Edinburgh, 2003.

- Commonwealth Scholarship (ref. ZACS-2002-344), Association of Commonwealth Universities, for study at the University of Edinburgh, 2002–2003.

- The Liberty Life Gold Medal for outstanding performance in Computer Science Honours, University of the Witwatersrand, 2001.

- The Altech Systems Prize for the best Computer Science Honours Research Report, University of the Witwatersrand, 2001.

- The Colin James Young Award for the best project in any area of the Mathematical Sciences, University of the Witwatersrand, 2001.

- The Computer Science Alumni Medal for the Best Student Tutor, University of the Witwatersrand, 2001.

- Wits Interactive Computing Group (ICG): President (2001), Head of Programming (1999 and 2000), Programming Champion (2000 and 2001).

Publications

Journal Articles

1. R. Rodriguez-Sanchez and G.D. Konidaris. Learning Abstract World Models for Value-Preserving Planning with Options. *Reinforcement Learning Journal* 1(1), August 2024.

2. A. Pacheck, S. James, G.D. Konidaris, and H. Kress-Gazit. Automatic Encoding and Repair of Reactive High-Level Tasks with Learned Abstract Representations. *International Journal of Robotics Research* 42(4-5), pages 263–288, April 2023.

3. M. M. Baker, A. New, M. Aguilar-Simon, Z. Al-Halah, S.M.R. Arnold, E. Ben-Iwhiwhu, A.P. Brna, E. Brooks, R.C. Brown, Z. Daniels, A. Daram, F. Delattre, R. Dellana, E. Eaton, H. Fu, K. Grauman, J. Hostetler, S. Iqbal, C. Kent, N. Ketz, S. Kolouri, G.D. Konidaris, D. Kudithipudi, E. Learned-Miller, S. Lee, M.L. Littman, S. Madireddy, J.A. Mendez, E.Q. Nguyen, C.D. Piatko, P.K. Pilly, A. Raghavan, A. Rahman, S.K. Ramakrishnan, N. Ratzlaff, A. Soltoggio, P. Stone, I. Sur, Z. Tang, S. Tiwari, K. Vedder, F. Wang, Z. Xu, A. Yanguas-Gil, H. Yedidsion, S. Yu, G.K. Vallabha. A Domain-Agnostic Approach for Characterization of Lifelong Learning Systems. *Neural Networks* volume 160, pages 274-296, March 2023.

4. B. Ames, J. Morgan, and G.D. Konidaris. IKFlow: Generating Diverse Inverse Kinematics Solutions. *IEEE Robotics and Automation Letters* 7(3), pages 7177–7184, July 2022.

5. O. Kroemer, S. Niekum, and G.D. Konidaris. A Review of Robot Learning for Manipulation: Challenges, Representations, and Algorithms. *Journal of Machine Learning Research* 22(30), pages 1–82, January 2021.

6. B. Keller, M. Draelos, K. Zhou, R. Qian, A.N. Kuo, G.D. Konidaris, K. Hauser, and J.A. Izatt. Optical Coherence Tomography-Guided Robotic Ophthalmic Microsurgery via Reinforcement Learning from Demonstration. *IEEE Transactions on Robotics* 36(4), 1207–1218, August 2020.

7. G.D. Konidaris. On The Necessity of Abstraction. *Current Opinion in Behavioral Sciences* 29 (Special Issue on Artificial Intelligence), pages 1–7, October 2019.

8. E. Rosen, D. Whitney, E. Phillips, G. Chen, J. Tompkin, G.D. Konidaris, and S. Tellex. Communicating And Controlling Robot Arm Motion Intent Through Mixed Reality Head-Mounted Displays. *The International Journal of Robotics Research* 38(12-13), October 2019.

9. C. Amato, G.D. Konidaris, L.P. Kaelbling, and J.P. How. Modeling and Planning with Macro-Actions in Decentralized POMDPs. *Journal of Artificial Intelligence Research* 64, pages 817–859, March 2019.

10. Y. Zhou, B. Burchfiel, and G.D. Konidaris. Representing, Learning, and Controlling Complex Object Interactions. *Autonomous Robots* 42(7), pages 1355–1367, April 2018.

11. E.L. Nelson, N.E. Berthier, and G.D. Konidaris. Handedness and Reach-to-Place Kinematics in Adults: Left-Handers Are Not Reversed Right-Handers. *Journal of Motor Behavior* 50:4, pages 381–391, July 2018.

12. G.D. Konidaris, L.P. Kaelbling, and T. Lozano-Perez. From Skills to Symbols: Learning Symbolic Representations for Abstract High-Level Planning. *Journal of Artificial Intelligence Research* 61, pages 215–289, January 2018.

13. C. Amato, G.D. Konidaris, A. Anders, G. Cruz, J. How, and L.P. Kaelbling. Policy Search for Multi-Robot Coordination under Uncertainty. *The International Journal of Robotics Research* 35(14), pages 1760–1778, December 2016.

14. D. Wookey and G.D. Konidaris. Regularized Feature Selection in Reinforcement Learning. *Machine Learning* 100(2), pages 655–676, September 2015.

15. S. Niekum, S. Osentoski, G.D. Konidaris, S. Chitta, B. Marthi, and Andrew G. Barto. Learning Grounded Finite-State Representations from Unstructured Demonstrations. *The International Journal of Robotics Research*, 34(2), pages 131–157, February 2015.

16. E.L. Nelson, G.D. Konidaris, and N.E. Berthier. Hand preference status and reach kinematics in infants. *Infant Behavior and Development*, 37(4), 615–623, November 2014.

17. G.D. Konidaris, I. Scheidwasser and A.G. Barto. Transfer in Reinforcement Learning using Common Features. *Journal of Machine Learning Research* 13:1333–1371, May 2012.

18. E.L. Nelson, G.D. Konidaris, N.E. Berthier, M.C. Braun, M.S.F.X. Novak, S.J. Suomi and M.A. Novak. Kinematics of reaching and implications for handedness in rhesus monkey infants. *Developmental Psychobiology* 54(4), pages 460–467, May 2012.

19. G.D. Konidaris, S.R. Kuindersma, R.A. Grupen and A.G. Barto, Robot Learning from Demonstration by Constructing Skill Trees. *The International Journal of Robotics Research* 31(3), pages 360–375, March 2012.

20. G.D. Konidaris and G.M. Hayes. An Architecture for Behavior-Based Reinforcement Learning. *Adaptive Behavior* 13(1), pages 5–32, March 2005.

Highly Refereed Conference Papers

21. X. Liu and A. Shah and G.D. Konidaris and S. Tellex and D. Paulius. Grounding Spatio-temporal Navigation Commands Using Large Language and Vision Models. Accepted, *Proceedings of the 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems*, October 2024.

22. R. Zhang, H. Fu, Y. Miao, and G.D. Konidaris. Model-based Reinforcement Learning for Parameterized Action Spaces. In *Proceedings of the Forty-first International Conference on Machine Learning*, July 2024.

23. H. Fu, P. Sharma, E. Stengel-Eskin, G.D. Konidaris, N. Le Roux, M-A Côté, and X. Yuan. Language-Guided Skill Learning with Temporal Variational Inference. In *Proceedings of the Forty-first International Conference on Machine Learning*, July 2024.

24. M. Merlin, S. Parr, N. Parikh, S. Orozco, V. Gupta, E. Rosen, and G.D. Konidaris. Robot Task Planning Under Local Observability. In *Proceedings of the 2024 IEEE Conference on Robotics and Automation*, May 2024.

25. B. Abbatematteo, E. Rosen, S. Thompson, M.T. Akbulut, S. Rammohan, and G.D. Konidaris. *Composable Interaction Primitives: A Structured Policy Class for Efficiently Learning Sustained-Contact Manipulation Skills.* In Proceedings of the 2024 IEEE Conference on Robotics and Automation, May 2024.

26. X. Liu, A.J. Shah, E. Rosen, M. Jia, G.D. Konidaris, and S. Tellex. LTL-Transfer: Skill Transfer for Temporal Task Specification. In *Proceedings of the 2024 IEEE Conference on Robotics and Automation*, May 2024.

27. A. Bagaria, B.M. Abbatematteo, O. Gottesman, M. Corsaro, S. Rammohan, and G.D. Konidaris, Effectively Learning Initiation Sets in Hierarchical Reinforcement Learning. In *Neural Information Processing Systems 37*, December 2023.

28. E. Rosen, S. James, S. Orozco, V. Gupta, M. Merlin, S. Tellex, and G.D. Konidaris. Synthesizing Navigation Abstractions for Planning with Portable Manipulation Skills. In *Proceedings of the 7th Annual Conference on Robot Learning*, November 2023.

29. R. Ma, L. Lam, B.A. Spiegel, A. Ganeshan, R. Patel, B. Abbatematteo, D.P. Paulius, S. Tellex, and G.D. Konidaris. Skill Generalization With Verbs. In *Proceedings of the 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems*, October 2023.

30. S. Shaw, D. Jha, A. Raghunathan, R.I. Corcodel, D. Romeres, G.D. Konidaris, and D. Nikovski. Constrained Dynamic Movement Primitives for Collision Avoidance in Novel Environments. In *Proceedings of the 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems*, October 2023.

31. I. Idrees, T. Yun, N. Sharma, N. Gopalan, S. Tellex, and G.D. Konidaris. Improved Inference of Human Intent by Combining Plan Recognition and Language Feedback. In *Proceedings of the 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems*, October 2023.

32. R. Rodriguez-Sanchez, B.A. Spiegel, J. Wang, R. Patel, G.D. Konidaris, and S. Tellex. RLang: A Declarative Language for Describing Partial World Knowledge to Reinforcement Learning Agents. In *Proceedings of the Fortieth International Conference on Machine Learning*, July 2023.

33. S. Lobel, A. Bagaria, and G.D. Konidaris. Flipping Coins to Estimate Pseudocounts for Exploration in Reinforcement Learning. In *Proceedings of the Fortieth International Conference on Machine Learning*, July 2023.

34. H. Fu, S. Yu, S. Tiwari, M.L. Littman, and G.D. Konidaris. Meta-Learning Parameterized Skills. In *Proceedings of the Fortieth International Conference on Machine Learning*, July 2023.

35. H. Fu, J. Yao, O. Gottesman, F. Doshi-Velez, and G.D. Konidaris. Performance Bounds for Model and Policy Transfer in Hidden-parameter MDPs. In *Proceedings of the Eleventh International Conference on Learning Representations*, May 2023.

36. O. Gottesman, K. Asadi, C. Allen, S. Lobel, G.D. Konidaris, and M.L. Littman. Coarse-Grained Smoothness for Reinforcement Learning in Metric Spaces. In *Proceedings of the 26th International Conference on Artificial Intelligence and Statistics*, April 2023.

37. S. Lobel, S. Rammohan, B. He, S. Yu, and G.D. Konidaris. Q-Functionals for Value-Based Continuous Control. In *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence*, February 2023.

38. H. Fu, S. Yu, M.L. Littman, and G.D. Konidaris. Model-based Lifelong Reinforcement Learning with Bayesian Exploration. In *Advances in Neural Information Processing Systems 35*, December 2022.

39. S. Tiwari and G.D. Konidaris. Effects of Data Geometry in Early Deep Learning. In *Advances in Neural Information Processing Systems 35*, December 2022.

40. C. Lovering, J.Z. Forde, G.D. Konidaris, E. Pavlick, and M.L. Littman. Evaluation beyond Task Performance: Analyzing Concepts in AlphaZero in Hex. In *Advances in Neural Information Processing Systems 35*, December 2022.

41. H. Abdul-Rashid, M. Freeman, B, Abbatematteo, G.D. Konidaris, and D Ritchie. Learning to Infer Kinematic Hierarchies for Novel Object Instances. In *Proceedings of the 2022 International Conference on Robotics and Automation*, May 2022.

42. S. Shaw, B. Abbatematteo, and G.D. Konidaris. RMPs for Safe Impedance Control in Contact-Rich Manipulation. In *Proceedings of the 2022 International Conference on Robotics and Automation*, May 2022.

43. E. Hsiung, H. Mehta, J. Chu, X. Liu, R. Patel, S. Tellex, and G.D. Konidaris. *Generalizing to New Domains by Mapping Natural Language to Lifted LTL*. In Proceedings of the 2022 International Conference on Robotics and Automation, May 2022.

44. M. Berg, S. Tellex, and G.D. Konidaris. Using Language to Generate State Abstractions for Long-Range Planning in Outdoor Environments. In *Proceedings of the 2022 International Conference on Robotics and Automation*, May 2022.

45. K. Zheng, R. Chitnis, Y. Sung, G.D. Konidaris, and S. Tellex. Towards Optimal Correlational Object Search. In *Proceedings of the 2022 International Conference on Robotics and Automation*, May 2022.

46. S. James, B. Rosman, and G.D. Konidaris. Autonomous Learning of Object-Centric Abstractions for High-Level Planning. In *Proceedings of the The Tenth International Conference on Learning Representations*, April 2022.

47. S. Lobel, A. Bagaria, C. Allen, O. Gottesman, G.D. Konidaris. Optimistic Initialization for Exploration in Continuous Control. In *Proceedings of the Thirty-Sixth AAAI Conference on Artificial Intelligence*, February 2022.

48. C. Allen, N. Parikh, O. Gottesman, and G.D. Konidaris. Learning Markov State Abstractions for Deep Reinforcement Learning. In *Advances in Neural Information Processing Systems 34*, December 2021.

49. B. Abbatematteo, E. Rosen, S. Tellex, and G.D. Konidaris. Bootstrapping Motor Skill Learning with Motion Planning. In *Proceedings of the 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems*, September 2021.

50. M. Corsaro, S. Tellex, and G.D. Konidaris. Learning to Detect Multi-Modal Grasps for Dexterous Grasping in Dense Clutter. In *Proceedings of the 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems*, September 2021.

51. K. Zheng, Y. Sung, G.D. Konidaris, and S. Tellex. Multi-Resolution POMDP Planning for Multi-Object Search in 3D. In *Proceedings of the 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems*, September 2021.

52. A. Bagaria, J. Senthil, M. Slivinski, and G.D. Konidaris. Robustly Learning Composable Options in Deep Reinforcement Learning. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence*, August 2021.

53. C. Allen, M. Katz, T. Klinger, G.D. Konidaris, M. Riemer, and G. Tesauro. Efficient Black-Box Planning Using Macro-Actions with Focused Effects. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence*, August 2021.

54. A. Bagaria, J. Senthil, and G.D. Konidaris. Skill Discovery for Exploration and Planning using Deep Skill Graphs. In *Proceedings of the Thirty-eighth International Conference on Machine Learning*, pages 521–531, July 2021.

55. B. Tang, M. Corsaro, G.D. Konidaris, S. Nikolaidis, and S. Tellex. Learning Collaborative Pushing and Grasping Policies in Dense Clutter. In *Proceedings of the 2021 International Conference on Robotics and Automation*, May 2021.

56. K. Asadi, N. Parikh, R. Parr, G.D. Konidaris, and M.L. Littman. Deep Radial-Basis Value Functions for Continuous Control. In *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence*, February 2021.

57. J. Roy and G.D. Konidaris. Visual Transfer For Reinforcement Learning Via Wasserstein Domain Confusion. In *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence*, February 2021.

58. E. Rosen, N. Kumar, N. Gopalan, D. Ullman, G.D. Konidaris, and S. Tellex. Building Plannable Representations with Mixed Reality. In *Proceedings of the 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 11146–11153, October 2020.

59. S. Murray, G.D. Konidaris, and D.J. Sorin. Roadmap Subsampling for Changing Environments. In *Proceedings of the 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5664–5670, October 2020.

60. S. James, B. Rosman, and G.D. Konidaris. Learning Portable Representations for High-Level Planning. In *Proceedings of the Thirty-Seventh International Conference on Machine Learning*, pages 4682–4691, July 2020

61. N. Gopalan, E. Rosen, G.D. Konidaris, and S. Tellex. Simultaneously Learning Transferable Symbols and Language Groundings from Perceptual Data for Instruction Following. In *Robotics: Science and Systems XVI*, July 2020.

62. Y. Jinnai, J. Park, M.C. Machado, and G.D. Konidaris. Exploration in Reinforcement Learning with Deep Covering Options. In *Proceedings of the Eighth International Conference on Learning Representations*, April 2020.

63. A. Bagaria and G.D. Konidaris. Option Discovery using Deep Skill Chaining. In *Proceedings of the Eighth International Conference on Learning Representations*, April 2020.

64. B. Abbatematteo, S. Tellex, and G.D. Konidaris. Learning to Generalize Kinematic Models to Novel Objects. In *Proceedings of the Third Conference on Robot Learning*, November 2019.

65. V. Cohen, B. Burchfiel, T. Nguyen, N. Gopalan, G.D. Konidaris, and S. Tellex. Grounding Language Attributes to Objects Using Bayesian Eigenobjects. In *Proceedings of the 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1187–1194, November 2019.

66. B. Ames and G.D. Konidaris. Bounded Error LQR-Trees. In *Proceedings of the 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 144–150, November 2019.

67. A. Pacheck, G.D. Konidaris, and H. Kress-Gazit. Automatic Encoding and Repair of Reactive High-Level Tasks with Learned Abstract Representations. In *Robotics Research: the 18th Annual Symposium*, October 2019.

68. S. Kim, K. Asadi, M.L. Littman, and G.D. Konidaris. DeepMellow: Removing the Need for a Target Network in Deep Q-Learning. In *Proceedings of the Twenty Eighth International Joint Conference on Artificial Intelligence*, pages 2733–2739, August 2019.

69. S. Murray, W. Floyd-Jones, G.D. Konidaris, and D.J. Sorin. A Programmable Architecture for Robot Motion Planning Acceleration. In *Proceedings of the 30th IEEE International Conference on Application-specific Systems, Architectures and Processors*, July 2019.

70. Y. Jinnai, J. Park, D. Abel, and G.D. Konidaris. Discovering Options for Exploration by Minimizing Cover Time. In *Proceedings of the 36th International Conference on Machine Learning*, pages 3130–3139, June 2019.

71. Y. Jinnai, D. Abel, D. Hershkowitz, M.L. Littman, and G.D. Konidaris. Finding Options that Minimize Planning Time. In *Proceedings of the 36th International Conference on Machine Learning*, pages 3120–3129, June 2019.

72. N. DeMarinis, S. Tellex, V. Kemerlis, G.D. Konidaris, and R. Fonseca. Scanning the Internet for ROS: A View of Security in Robotics Research. In *Proceedings of the 2019 International Conference on Robotics and Automation*, May 2019.

73. S.Y. Gadre, E. Rosen, G. Chien, E. Phillips, S. Tellex, and G.D. Konidaris. End-User Robot Programming Using Mixed Reality. In *Proceedings of the 2019 International Conference on Robotics and Automation*, May 2019.

74. A. Levy, G.D. Konidaris, R. Platt, and K. Saenko. Learning Multi-Level Hierarchies with Hindsight. In *Proceedings of the Eighth International Conference on Learning Representations*, May 2019.

75. B. Burchfiel and G.D. Konidaris. Hybrid Bayesian Eigenobjects: Combining Linear Subspace and Deep Network Methods for 3D Robot Vision. In *Proceedings of the 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 6843–6850, October 2018.

76. B. Ames, A. Thackston, and G.D. Konidaris. Learning Symbolic Representations for Planning with Parameterized Skills. In *Proceedings of the 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 526–533, October 2018.

77. D. Abel, Y. Jinnai, Y. Guo, G.D. Konidaris, and M.L. Littman. Policy and Value Transfer for Lifelong Reinforcement Learning. In *Proceedings of the Thirty-fourth International Conference on Machine Learning*, pages 20–29, July 2018.

78. G. Andersen and G.D. Konidaris. Active Exploration for Learning Symbolic Representations. In *Neural Information Processing Systems 30*, pages 5009–5019, December 2017.

79. T. Killian, S. Daulton, G.D. Konidaris, and F. Doshi-Velez. Robust and Efficient Transfer Learning with Hidden Parameter Markov Decision Processes. In *Neural Information Processing Systems 30*, pages 6250–6261, December 2017.

80. D. Whitney, E. Rosen, E. Phillips, G.D. Konidaris, and S. Tellex. Comparing Robot Grasping Teleoperation Across Desktop and Virtual Reality with ROS Reality. In *Robotics Research: the 17th Annual Symposium*, December 2017.

81. E. Rosen, D. Whitney, E. Phillips, G. Chen, J. Tompkin, G.D. Konidaris, and S. Tellex. Communicating Robot Arm Motion Intent Through Mixed Reality Head-mounted Displays. In *Robotics Research: the 17th Annual Symposium*, December 2017.

82. B. Burchfiel and G.D. Konidaris. Bayesian Eigenobjects: A Unified Framework for 3D Robot Perception. In *Robotics: Science and Systems XIII*, July 2017.

83. S. James, G.D. Konidaris, and B. Rosman. An Analysis of Monte Carlo Tree Search. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, pages 3576–3582, February 2017.

84. S. Murray, W. Floyd-Jones, Y. Qi, G.D. Konidaris and D. Sorin. The Microarchitecture of a Real-Time Robot Motion Planning Accelerator. In *Proceedings of the The 49th Annual IEEE/ACM International Symposium on Microarchitecture*, October 2016.

85. G.D. Konidaris. Constructing Abstraction Hierarchies Using a Skill-Symbol Loop. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*, pages 1648–1654, July 2016.

86. F. Doshi-Velez and G.D. Konidaris. Hidden Parameter Markov Decision Processes: A Semiparametric Regression Approach for Discovering Latent Task Parametrizations. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*, pages 1432–1440, July 2016.

87. S. Murray, W. Floyd-Jones, Y. Qi, D. Sorin and G.D. Konidaris. Robot Motion Planning on a Chip. In *Robotics: Science and Systems XII*, June 2016.

88. Y. Zhou and G.D. Konidaris. Representing and Learning Complex Object Interactions. In *Robotics: Science and Systems XII*, June 2016.

89. W. Masson, P. Ranchod, and G.D. Konidaris. Reinforcement Learning with Parameterized Actions. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pages 1934–1940, February 2016.

90. P. S. Thomas, S. Niekum, G. Theocharous, and G. D. Konidaris. Policy Evaluation using the Ω-Return. In *Advances in Neural Information Processing Systems 29*, pages 334–342, December 2015.

91. P. Ranchod, B. Rosman, G.D. Konidaris. Nonparametric Bayesian Reward Segmentation for Skill Discovery Using Inverse Reinforcement Learning. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 471–477, September 2015.

92. G.D. Konidaris, L.P. Kaelbling, and T. Lozano-Perez. Symbol Acquisition for Probabilistic High-Level Planning. In *Proceedings of the Twenty Fourth International Joint Conference on Artificial Intelligence*, pages 3619–3627, July 2015.

93. C. Amato, G.D. Konidaris, A. Anders, G. Cruz, J.P. How, and L.P. Kaelbling. Policy Search for Multi-Robot Coordination under Uncertainty. In *Robotics: Science and Systems XI*, July 2015.

94. C. Amato, G.D. Konidaris, G. Cruz, C. Maynor, J.P. How, and L.P. Kaelbling. Planning for Decentralized Control of Multiple Robots Under Uncertainty. In *Proceedings of the 2015 IEEE International Conference on Robotics and Automation*, pages 1241-1248, May 2015.

95. G.D. Konidaris, L. Kaelbling and T. Lozano-Perez. Constructing Symbolic Representations for High-Level Planning. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, pages 1932–1940, July 2014.

96. N. Hollingsworth, J. Meyer, R. McGee, J. Doering, G.D. Konidaris and L. Kaelbling. Optimizing a Start-Stop Controller using Policy Search. In *Twenty-Sixth Annual Conference on Innovative Applications of Artificial Intelligence*, pages 2984–2989, July 2014.

97. B.C. da Silva, G.D. Konidaris, and A.G. Barto. Active Learning of Parameterized Skills. In *Proceedings of the Thirty First International Conference on Machine Learning*, pages 1737-1745, June 2014.

98. B.C. da Silva, G. Baldassarre, G.D. Konidaris, and A.G. Barto. Learning Parameterized Motor Skills on a Humanoid Robot. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 5239–5244, May 2014.

99. C. Amato, G.D. Konidaris and L.P. Kaelbling. Planning with Macro-Actions in Decentralized POMDPs. In *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems*, pages 1273–1280, May 2014.

100. G. Goretkin, A. Perez, R. Platt Jr., G.D. Konidaris. Optimal Sampling-Based Planning for Linear-Quadratic Kinodynamic System. *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 2429–2436, May 2013.

101. S. Niekum, S. Osentoski, G.D. Konidaris and A.G. Barto. Learning and Generalization of Complex Tasks from Unstructured Demonstrations. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5239–5246, October 2012.

102. B.C da Silva, G.D. Konidaris and A.G. Barto. Learning Parameterized Skills. In *Proceedings of the Twenty-Ninth International Conference in Machine Learning*, pages 1679–1686, June 2012.

103. A. Perez, R. Platt, G.D. Konidaris, L. Kaelbling and T. Lozano-Perez. LQR-RRT*: Optimal Sampling-Based Motion Planning with Automatically Derived Extension Heuristics. In *Proceedings of 2012 IEEE International Conference on Robotics and Automation*, pages 2537–2542, May 2012.

104. G.D. Konidaris, S. Niekum and P.S. Thomas. TD_γ: Re-evaluating Complex Backups in Temporal Difference Learning. *Advances in Neural Information Processing Systems 24*, pages 2402–2410, December 2011.

105. G.D. Konidaris, S.R. Kuindersma, R.A. Grupen and A.G. Barto. Autonomous Skill Acquisition on a Mobile Manipulator. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*, pages 1468–1473, August 2011.

106. G.D. Konidaris, S. Osentoski and P.S. Thomas. Value Function Approximation in Reinforcement Learning using the Fourier Basis. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*, pages 380–385, August 2011.

107. G.D. Konidaris, S.R. Kuindersma, A.G. Barto and R.A. Grupen. Constructing Skill Trees for Reinforcement Learning Agents from Demonstration Trajectories. In *Advances in Neural Information Processing Systems 23*. pages 1162–1170, December 2010.

108. G.D. Konidaris and A.G. Barto. Skill Discovery in Continuous Reinforcement Learning Domains using Skill Chaining. In *Advances in Neural Information Processing Systems 22*, pages 1015–1023. December 2009.

109. G.D. Konidaris and A.G. Barto. Efficient Skill Learning Using Abstraction Selection. In *Proceedings of the Twenty First International Joint Conference on Artificial Intelligence*, pages 1107–1112, July 2009.

110. G.D. Konidaris and A.G. Barto. Sensorimotor Abstraction Selection for Efficient, Autonomous Robot Skill Acquisition. In *Proceedings of the 7th IEEE International Conference on Development and Learning*, pages 151–156, August 2008.

111. L. Georgopoulos, G.M. Hayes and G.D. Konidaris. A Forward Model of Optic Flow for Detecting External Forces. In *Proceedings of the IEEE/RSJ 2007 International Conference on Intelligent Robots and Systems*, pages 913–918, October 2007.

112. G.D. Konidaris and A.G. Barto. Building Portable Options: Skill Transfer in Reinforcement Learning. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence*, pages 895–900, January 2007.

113. G.D. Konidaris and A.G. Barto. An Adaptive Robot Motivational System. In *From Animals to Animats 9: Proceedings of the 9th International Conference on the Simulation of Adaptive Behavior*, pages 346–356, September 2006.

114. G.D. Konidaris and A.G. Barto. Autonomous Shaping: Knowledge Transfer in Reinforcement Learning. In *Proceedings of the Twenty Third International Conference on Machine Learning*, pages 489–496, June 2006.

115. S. Rauchas, B. Rosman, G.D. Konidaris and I.D. Sanders. Language Performance at High School and Success in First Year Computer Science. In *Proceedings of the SIGCSE 2006 Technical Symposium on Computer Science Education*, pages 398–402, March 2006.

116. F.J. Stewart, T. Taylor and G.D. Konidaris. METAMorph: Experimenting with Genetic Regulatory Networks for Artificial Development. In *Proceedings of the VIIIth European Conference on Artificial Life*, pages 108–117, September 2005.

117. G.D. Konidaris and G.M. Hayes. Estimating Future Reward in Reinforcement Learning Animats using Associative Learning. In *From Animals to Animats 8: Proceedings of the 8th International Conference on the Simulation of Adaptive Behavior*, pages 297–304, July 2004.

118. G.D. Konidaris, T. Taylor and J.C.T Hallam. HydroGen: Automatically Generating Self-Assembly Code for Hydron Units. In *Proceedings of the Seventh International Symposium on Distributed Autonomous Robotic Systems*, pages 33–42, June 2004.

Chapters in Books

119. A.G. Barto, G.D. Konidaris, and C.M. Vigorito. Behavioral Hierarchy: Exploration and Representation. In *Computational and Robotic Models of the Hierarchical Organization of Behavior*, Baldassarre, Gianluca; Mirolli, Marco (Eds.), pages 13–46, Springer, Berlin, October 2013.

Lightly Refereed Conference Papers, Workshops, Symposia and Posters

120. B. Quartey, A. Shah, and G.D. Konidaris. Exploiting Contextual Structure to Generate Useful Auxiliary Tasks. In the *NeurIPS 2023 Workshop on Generalization in Planning*, December 2023.

121. R. Rodriguez-Sanchez, and G. Konidaris. Learning Abstract World Models for Value-preserving Planning with Options. In the *NeurIPS 2023 Workshop on Generalization in Planning*, December 2023.

122. M. Fishman, N. Kumar, C. Allen, N. Danas, M. Littman, S. Tellex, and G. Konidaris. Task Scoping: Generating Task-Specific Simplifications of Open-Scope Planning Problems. In the *IJCAI Workshop on Bridging the Gap Between AI Planning and Reinforcement Learning*, August 2023.

123. B.A. Spiegel and G.D. Konidaris. Guided Policy Search for Parameterized Skills using Adverbs. In The *ICML 2023 Workshop on Interactive Learning with Implicit Human Feedback*, July 2023.

124. B. Abbatematteo, C. Robbins, K. Sherry, J. Mokkamakkul, E. Rosen, S. Thompson, M. Stein, and G.D. Konidaris. Sensorized Manipulation Challenge Kit for Benchmarking Robotic Manipulation. In the *RSS 2023 Workshop on Experiment-oriented Locomotion and Manipulation Research*, July 2023.

125. E. Rosen, B. Abbatematteo, S. Thompson, T. Akbulut, and G.D. Konidaris. On the Role of Structure in Manipulation Skill Learning. In the *CoRL 2022 Workshop on Learning, Perception, and Abstraction for Long-Horizon Planning*, December 2022.

126. J.Z. Forde, C. Lovering, G.D. Konidaris, E. Pavlick, and M.L. Littman. Where, When & Which Concepts Does AlphaZero Learn? Lessons from the Game of Hex. In *Proceedings of the AAAI-22 Workshop on Reinforcement Learning in Games*, February 2022.

127. H. Fu, S. Yu, M.L. Littman, and G.D. Konidaris. Bayesian Exploration for Lifelong Reinforcement Learning. In *The NeurIPS 2021 Deep Reinforcement Learning Workshop*, December 2021.

128. C. Allen, N. Parikh, and G.D. Konidaris. Learning Markov State Abstractions for Deep Reinforcement Learning. In *The NeurIPS 2020 Deep Reinforcement Learning Workshop*, December 2020.

129. N. Parikh, Z. Horvitz, N. Srinivasan, A. Shah, and G.D. Konidaris. Graph Embedding Priors for Multi-task Deep Reinforcement Learning. In *The 4th Knowledge Representation and Reasoning Meets Machine Learning Workshop at NeurIPS 2020*, December 2020.

130. A. Bagaria, J. Crowley, J.W.N. Lim, G.D. Konidaris. Skill Discovery for Exploration and Planning using Deep Skill Graphs. In *the 4th Lifelong Learning Workshop at ICML 2020*, July 2020.

131. R. Rodriguez-Sanchez, R. Patel, and G.D. Konidaris. On the Relationship Between Structure in Natural Language and Models of Sequential Decision Processes. In *the First Workshop on Language in Reinforcement Learning at ICML 2020*, July 2020.

132. M. Merlin, N. Parikh, E. Rosen, and G.D. Konidaris. Locally Observable Markov Decision Processes. In the *ICRA 2020 Workshop on Perception, Action, Learning: From Metric-Semantic Scene Understanding to High-level Task Execution*, June 2020.

133. A. Bagaria and G.D. Konidaris. Option Discovery using Deep Skill Chaining. In the *NeurIPS 2019 Workshop on Deep Reinforcement Learning*, December 2019.

134. S. Kim and G.D. Konidaris. Adaptive Temperature Tuning for Mellowmax in Deep Reinforcement Learning. In the *NeurIPS 2019 Workshop on Deep Reinforcement Learning*, December 2019.

135. S. Kim, K. Asadi, M.L. Littman, and G.D. Konidaris. Removing the Target Network from Deep Q-Networks with the Mellowmax Operator (Extended Abstract). In *Proceedings of the 2019 International Conference on Autonomous Agents and Multiagent Systems*, May 2019.

136. Y. Jinnai, D. Abel, J.W. Park, D.E. Hershkowitz, M.L. Littman, and G.D. Konidaris Skill Discovery with Well-Defined Objectives. In *Proceedings of the ICLR Workshop on Structure and Priors in Reinforcement Learning*, May 2019.

137. S. Kim, K. Asadi, M.L. Littman, and G.D. Konidaris. Removing the Target Network from Deep Q-Networks with Mellowmax Operator (Extended Abstract). In *Proceedings of the 2019 International Conference on Autonomous Agents and Multiagent Systems*, May 2019.

138. S. James, B. Rosman, and G.D. Konidaris. Learning to Plan with Portable Symbols. In the *ICML/IJCAI/AAMAS 2018 Workshop on Planning and Learning*, July 2018.

139. T. Killian, G.D. Konidaris, and F. Doshi-Velez. Transfer Learning Across Patient Variations with Hidden Parameter Markov Decision Processes. *The NIPS 2016 Workshop on Machine Learning for Health*, December 2016.

140. S. James, B. Rosman and G.D. Konidaris. An Investigation into the Effectiveness of Heavy Rollouts in UCT. *The IJCAI 2016 Workshop on General Intelligence in Game-Playing Agents*, July 2016.

141. B. Burchfiel and G.D. Konidaris. Generalized 3D Object Representation using Bayesian Eigenobjects. *The RSS 2016 Workshop on Geometry and Beyond: Representations, Physics, and Scene Understanding for Robotics*, June 2016.

142. C. Amato, S. Omidshafiei, A Agha-mohammadi, G.D. Konidaris, J.P. How, and L.P. Kaelbling. Probabilistic Planning for Multi-Robot Systems. In *Proceedings of the RSS-2015 Workshop on Principles of Multi-Robot Systems*, July 2015.

143. G.D. Konidaris and F. Doshi-Velez. Hidden Parameter Markov Decision Processes: An Emerging Paradigm for Modeling Families of Related Tasks. In *Proceedings of the AAAI 2014 Fall Symposium on Knowledge, Skill, and Behavior Transfer in Autonomous Robots*, November 2014.

144. C. Amato, G.D. Konidaris, J.P. How and L.P. Kaelbling. Decentralized Decision-Making Under Uncertainty for Multi-Robot Teams. In *Proceedings of the 2014 IROS Workshop on The Future of Multiple-Robot Research and Its Multiple Identities*, September 2014.

145. C. Amato, G.D. Konidaris, G. Cruz, C.A. Maynor, J.P. How and L.P. Kaelbling. Planning for Decentralized Control of Multiple Robots Under Uncertainty. In *Proceedings of the 2014 ICAPS Workshop on Planning and Robotics*, June 2014.

146. C. Trewick, P. Ranchod and G.D. Konidaris. Preferential Targeting of HIV Infected Hubs in a Scale-free Sexual Network. In the *Annual Conference of the Computational Social Science Society of the Americas*, August 2013.

147. G.D. Konidaris. Robots, Skills, and Symbols (Extended Abstract). In *Proceedings of the 2013 Workshop on Machine Learning for Interactive Systems*, August 2013.

148. G.D. Konidaris, L.P. Kaelbling and T. Lozano-Perez. Symbol Acquisition for Task-Level Planning. In the *AAAI 2013 Workshop on Learning Rich Representations from Low-Level Sensors*, July 2013.

149. G.D. Konidaris, S. Kuindersma, S. Niekum, R.A. Grupen and A.G. Barto. Robot Learning: Some Recent Examples. In *Proceedings of the Sixteenth Yale Workshop on Adaptive and Learning Systems*, pages 71–76, Center for Systems Science, Dunham Laboratory, Department of Electrical Engineering, Yale University, New Haven CT, June 2013.

150. F. Doshi-Velez and G.D. Konidaris. Transfer Learning by Discovering Latent Task Parametrizations. In the *NIPS 2012 Workshop on Bayesian Nonparametric Models For Reliable Planning And Decision-Making Under Uncertainty*, December 2012.

151. G.D. Konidaris, S.R. Kuindersma, R.A. Grupen and A.G. Barto. Acquiring Transferrable Mobile Manipulation Skills. In the *RSS 2011 Workshop on Mobile Manipulation: Learning to Manipulate*, June 2011.

152. G.D. Konidaris, S.R. Kuindersma, R.A. Grupen and A.G. Barto. CST: Acquiring Skill Trees by Demonstration. In the *ICML 2011 Workshop on New Developments in Imitation Learning*, July 2011.

153. S. Kuindersma, G. Konidaris, R. Grupen, A. Barto. Learning from a Single Demonstration: Motion Planning with Skill Segmentation (poster abstract). In the *NIPS Workshop on Learning and Planning in Batch Time Series Data*, Vancouver, December 2010.

154. G.D. Konidaris and A.G. Barto. Towards the Autonomous Acquisition of Robot Skill Hierarchies (poster abstract). In the *Robotics: Science and Systems Workshop on Bridging the Gap Between High-Level Discrete Representations and Low-Level Continuous Behaviors*, Seattle, June 2009.

155. G.D. Konidaris and S. Osentoski. Value Function Approximation using the Fourier Basis (extended abstract). In the *Multidisciplinary Symposium on Reinforcement Learning*, Montreal, Canada, June 2009.

156. G.D. Konidaris and A.G. Barto. Skill Chaining: Skill Discovery in Continuous Domains (extended abstract). In the *Multidisciplinary Symposium on Reinforcement Learning*, Montreal, Canada, June 2009.

157. G.D. Konidaris. *Autonomous Robot Skill Acquisition (thesis summary)* Doctoral Symposium, 23rd National Conference on Artificial Intelligence (AAAI 2008), July 2008.

158. E.L. Nelson, G.D. Konidaris and N.E. Berthier. *Using Real-Time Motion Capture to Measure Handedness in Infants.* Poster presentation at the XVIth Biennial International Conference on Infant Studies, Vancouver, Canada, March 2008.

159. G.D. Konidaris. A Framework for Transfer in Reinforcement Learning. In the *ICML-06 Workshop on Structural Knowledge Transfer for Machine Learning*, Pittsburgh PA, June 2006.

160. A. Stout, G.D Konidaris and A.G. Barto. Intrinsically Motivated Reinforcement Learning: A Promising Framework for Developmental Robot Learning. In *Proceedings of the 2005 AAAI Spring Symposium on Developmental Robotics*, March 2005.

161. G.D. Konidaris and G.M. Hayes. Anticipatory Learning for Focusing Search in Reinforcement Learning Agents. In the *Second Workshop on Anticipatory Behavior in Adaptive Learning Systems*, Los Angeles CA, July 2004.

162. G.D. Konidaris, D.A. Shell and N. Oren. Evolving Neural Networks to Play the Capture Game. In *Proceedings of the SAICSIT 2002 Postgraduate Symposium*, September 2002.

163. J. Adler, G.D. Christelis, J.A. Deneys, G.D. Konidaris, G. Lewis, A.G. Lipson, R.L. Phillips, D.K. Scott-Dawkins, D.A. Shell, B.V. Strydom, W.M. Trakman and L.D. Van Gool. Finding Adjacencies in Non-Overlapping Polygons. Electronic Paper, *Proceedings of the 2001 SAICSIT Conference*, September 2001.

Dissertations and Technical Reports

164. C.J. Bester, S.D. James, and G.D. Konidaris. *Multi-Pass Q-Networks for Deep Reinforcement Learning with Parameterised Action Spaces.* arXiv:1905.04388, May 2019.

165. C. Allen, K. Asadi, M. Roderick, A-R. Mohamed, G.D. Konidaris, and M.L. Littman. *Mean Actor Critic.* arXiv:1709.00503 [stat.ML], September 2017.

166. D.H. Zewdie and G.D. Konidaris. *Representation Discovery for Kernel-Based Reinforcement Learning.* Technical Report MIT-CSAIL-TR-2015-032, MIT Computer Science and Artificial Intelligence Laboratory, November 2015.

167. G.D. Konidaris. *Autonomous Robot Skill Acquisition.* PhD Thesis, Department of Computer Science, University of Massachusetts, May 2011.

168. G.D. Konidaris and S. Osentoski. *Value Function Approximation in Reinforcement Learning using the Fourier Basis.* Technical Report UM-CS-2008-19, Department of Computer Science, University of Massachusetts Amherst, June 2008.

169. G.D. Konidaris. *Behaviour-Based Reinforcement Learning.* Master's Thesis, School of Informatics, University of Edinburgh, 2003.

170. G.D. Konidaris. *Axial Line Placement in Deformed Urban Grids.* Honours Dissertation, School of Computer Science, University of the Witwatersrand, 2001.