Robert Sim, Ph.D.

408-6385 Hawthorn Lane
Vancouver, BC, V6T 1Z4
Canada

(604) 221-8381
simra@cs.ubc.ca
http://www.cs.ubc.ca/~simra

Current Position

01/2004–present	NSERC/Canadian Space Agency Postdoctoral Fellow, Laboratory for Computational Intelligence, Department of Computer Science, University of British Columbia (UBC), Vancouver. Supervisor: Prof. James Little Design and implementation of autonomous vision-based robotic explorer. Strategies for information-optimal exploration.

Education

1998–2004	Doctor of Philosophy, Computer Science, McGill University. Thesis title: "On visual maps and their automatic construction." Supervisor: Professor Gregory Dudek.
1996–1998	Master of Science, Computer Science, McGill University. Thesis title: "Mobile robot localization from learned landmarks." Supervisor: Professor Gregory Dudek.
1992–1996	Bachelor of Engineering, Computer Engineering, McGill University.

Research Interests

My primary research activities are in vision-based robotics, and the development of practical systems that facilitate robot autonomy. Of particular interest are problems involving models of the visual world, how they can be efficiently represented, and how they can be learned automatically. A key subproblem is that of active learning and how a robot can be 'curious', collecting observations of the world in a manner that is both efficient and informative. My work has produced the world's first fully autonomous, 100% vision-based robotic explorer and mapper– completely laser-free!

Professional Experience

2003-2005	Postdoctoral Fellow (part-time), Artificial Intelligence Lab, University of Toronto Minimal environment descriptions. Vision-based CAD object model augmentation.

Supervisor: Prof. Sven Dickinson

1998–2003 Lead *Robodaemon* developer and project manager, McGill Mobile Robotics Lab

Design, development and maintenance of *Robodaemon*, a network-oriented, multi-threaded mobile robot controller and simulator

Implemented plugin and module architectures for extending platform.

Designed API's for C++ and Java clients.

Supervisor: Prof. Gregory Dudek

2000–2001 Vice President, University and Academic Affairs, McGill Post-Graduate Students' Society (PGSS).

Representing graduate students on a wide variety of University committees, including Senate, Budget Planning, Research Policy and Intellectual Property.

1999–2000 Ph. D. Senator, PGSS.
Representing Ph. D. students at the University Senate.

1998–1999 Team leader, consultant, McGill University Mobile Robotics Team

Provide expertise to team assembling mobile robot system for American Association for Artificial Intelligence (AAAI) Mobile Robot Competition.

Lead consultant on vision system architecture (1999). Design and implement real-time vision system to assist mobile robot navigation, obstacle avoidance, and object recognition (1998).

Implemented colour training architecture, shape recognition system for distinguishing objects, obstacles and floor.

First place finish in category (both years).

Supervisor: Prof. Gregory Dudek

1995–1995 Research Assistant, Centre for Intelligent Machines.

Design and implement an object recognition system by actively controlling a laser-stripe range sensor.

Supervisor: Prof. Martin D. Levine.

Selected Awards

2004–present Natural Sciences and Engineering Research Council (NSERC) Posdoctoral Fellowship.

2004-present Canadian Space Agency Postdoctoral Supplement.

2000–2002	Hydro Quebec McGill Major Fellowship.
2001	Dean's Award for Student Service to Graduate and Postdoctoral Education.
1998–2000	NSERC PGS B Fellowship.
1998–2000	Canadian Space Agency NSERC Supplement.
1998	Canadian Advanced Technology Alliance Prize for academic excellence.
1996–1998	NSERC PGS A Fellowship.
1996	British Association Medal for Academic Achievement.
1992–1996	Canada Scholar.
1992–1996	J.W. McConnell entrance scholarship.

Technical Skills

Languages	C++, C, Perl, Matlab.
Core skills	Scientific computing, Pthreads and concurrency, TCP/IP Sockets, real-time systems, distributed systems.
Graphics and UI	OpenGL, Gtk+, Gtkmm, MFC.
Markup/WWW	LaTeX, AJAX, RSS, HTML, CSS, custom languages.
Other skills	LAMP architecture, CGI programming, parsers.

Publications

Published Journal Articles

[1] P. L. Sala, R. Sim, A. Shokoufandeh, and S. J. Dickinson, "Landmark selection for vision-based navigation," *IEEE Transactions on Robotics*, 2005. To appear.

[2] R. Sim and G. Dudek, "Learning generative models of scene features," *International Journal of Computer Vision*, vol. 60, pp. 45–61, October 2004.

[3] R. Sim and G. Dudek, "Learning environmental features for pose estimation," *Image and Vision Computing, Elsevier Press*, vol. 19, no. 11, pp. 733–739, 2001.

Journal Articles in Submission

[4] R. Sim, P. Elinas, and J. J. Little, "A study of the Rao-Blackwellised particle filter for efficient and accurate vision-based SLAM," *International Journal of Computer Vision/International Journal of Robotics Research*, 2005. Special Issue on Vision in Robotics. Submitted Nov. 2005.

[5] I. Rekleitis, R. Sim, G. Dudek, and E. Milios, "Collaborative exploration for visual map construction," *IEEE Transactions on Robotics*, 2004. Submitted October, 2004.

Refereed Conferences

[6] R. Sim, P. Elinas, M. Griffin, A. Shyr, and J. J. Little, "Design and analysis of a framework for real-time vision-based SLAM using Rao-Blackwellised particle filters," in *Proceedings of the 3rd Canadian Conference on Computer and Robotic Vision (CRV)*, (Québec City, QC), CIPPRS, IEEE Press, June 2006. To appear.

[7] P. Elinas, R. Sim, and J. J. Little, "σSLAM: Stereo vision SLAM using the Rao-Blackwellised particle filter and a novel mixture proposal distribution," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, (Orlando, FL), IEEE, IEEE Press, May 2006. To appear.

[8] R. Sim, "Stabilizing information-driven exploration for bearings-only SLAM using range gating," in *Proceedings of Intelligent Robots and Systems (IROS)*, (Edmonton, AB), pp. 2745–2750, IEEE/RSJ, IEEE Press, 2005.

[9] R. Sim, M. Griffin, A. Shyr, and J. J. Little, "Scalable real-time vision-based SLAM for planetary rovers," in *IEEE IROS Workshop on Robot Vision for Space Applications*, (Edmonton, AB), pp. 16–21, IEEE, IEEE Press, August 2005.

[10] R. Sim, P. Elinas, M. Griffin, and J. J. Little, "Vision-based SLAM using the Rao-Blackwellised particle filter," in *Proceedings of the IJCAI Workshop on Reasoning with Uncertainty in Robotics (RUR)*, (Edinburgh, Scotland), pp. 9–16, 2005.

[11] R. Sim, "Stable exploration for bearings-only SLAM," in *Proceedings of the IEEE International Conference on Robotics and Automation*, (Barcelona, Spain), pp. 2422–2427, IEEE, IEEE Press, April 2005.

[12] R. Sim and N. Roy, "Global A-optimal robot exploration in SLAM," in *Proceedings of the IEEE International Conference on Robotics and Automation*, (Barcelona, Spain), pp. 673–678, IEEE, IEEE Press, April 2005.

[13] R. Sim and G. Dudek, "Learning generative models of invariant features," in *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 3, (Sendai, Japan), pp. 3481–3488, IEEE/RSJ, IEEE Press, September 2004.

[14] P. L. Sala, R. Sim, A. Shokoufandeh, and S. J. Dickinson, "Landmark selection for vision-based navigation," in *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 3, (Sendai, Japan), pp. 3131–3138, IEEE/RSJ, IEEE Press, September 2004.

[15] C. Georgiades, A. German, A. Hogue, H. Liu, C. Prahacs, A. Ripsman, R. Sim, L. A. Torres, P. Zhang, M. Buehler, G. Dudek, M. Jenkin, and E. Milios, "AQUA: an aquatic walking robot," in *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 3, (Sendai, Japan), pp. 3525–3531, IEEE/RSJ, IEEE Press, September 2004.

[16] R. Sim and G. Dudek, "Self-organizing visual maps," in *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, (San Jose, CA), pp. 470–475, AAAI Press, July 2004.

[17] R. Sim, G. Dudek, and N. Roy, "Online control policy optimization for minimizing map uncertainty during exploration," in *Proceedings of the IEEE International Conference on Robotics and Automation*, vol. 2, (New Orleans, LA), pp. 1758–1763, IEEE Press, April 2004.

[18] R. Sim and G. Dudek, "Effective exploration strategies for the construction of visual maps," in *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, vol. 3, (Las Vegas, NV), pp. 3224–3231, IEEE Press, October 2003.

[19] R. Sim and G. Dudek, "Examining exploratory trajectories for minimizing map uncertainty," in *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI) Workshop on Reasoning with Uncertainty in Robotics (RUR)*, (Acapulco, Mexico), pp. 69–76, Morgan Kaufmann, August 2003.

[20] R. Sim and G. Dudek, "Comparing image-based localization methods," in *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI)*, (Acapulco, Mexico), pp. 1560–1562, Morgan Kaufmann, August 2003.

[21] G. Dudek and R. Sim, "Robodaemon - a device independent, network-oriented, modular mobile robot controller," in *Proceedings of the IEEE International Conference on Robotics and Automation*, vol. 3, (Taipei, Taiwan), pp. 3434–3440, IEEE Press, May 2003.

[22] R. Sim and G. Dudek, "Learning generative models of scene features," in *IEEE Conference on Computer Vision and Pattern Recognition*, (Hawaii), pp. 406–412, IEEE Press, December 2001.

[23] I. Rekleitis, R. Sim, G. Dudek, and E. Milios, "Collaborative exploration for map construction," in *2001 IEEE International Symposium on Computational Intelligence in Robotics and Automation*, (Banff, AB), pp. 296–301, July 2001.

[24] I. Rekleitis, R. Sim, G. Dudek, and E. Milios, "Collaborative exploration for the construction of visual maps," in *2001 IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, (Hawaii), pp. 1269–1274, October 2001.

[25] R. Sim and G. Dudek, "Learning visual landmarks for pose estimation," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, (Detroit, MI), pp. 1972–1978, IEEE Press, May 1999.

[26] R. Sim and G. Dudek, "Learning environmental features for pose estimation," in *Proceedings of the 2nd IEEE Workshop on Perception for Mobile Agents*, (Ft. Collins, CO), pp. 7–14, IEEE Press, June 1999.

[27] R. Sim and G. Dudek, "Learning and evaluating visual features for pose estimation," in *Proceedings of the Seventh IEEE International Conference on Computer Vision (ICCV)*, (Kerkyra, Greece), pp. 1217–1222, IEEE Press, Sept 1999.

[28] R. Sim and G. Dudek, "Mobile robot localization from learned landmarks," in *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, (Victoria, Canada), pp. 1060–1065, IEEE Press, October 1998.

Invited Articles

[29] R. Sim and G. Dudek, "Visual landmarks for pose estimation," *Canadian Artificial Intelligence*, pp. 13–17, Spring 1999.

Abstracts

[30] R. Sim and G. Dudek, "On visual maps and their automatic construction," in *Precarn/IRIS Proceedings*, p. 1, June 2003.

[31] R. Sim and G. Dudek, "Learning landmarks for robot localization," in *Proceedings of the National Conference on Artificial Intelligence SIGART/AAAI Doctoral Consortium*, (Austin, TX), pp. 1110–1111, SIGART/AAAI, AAAI Press, July 2000.

[32] R. Sim and G. Dudek, "Position estimation from learned landmarks," in *Precarn/IRIS Proceedings*, p. 1, June 1998.

[33] R. Sim and G. Dudek, "Navigating by the stars: A landmark-based method for mobile robot localization," in *Precarn/IRIS Proceedings*, p. 1, June 1997.

[34] M. Bolduc, E. Bourque, G. Dudek, N. Roy, and R. Sim, "Autonomous exploration: An integrated systems approach," in *Proceedings of the AAAI National Conference on Artificial Intelligence*, pp. 779–780, AAAI Press, July 1997.

Technical Reports

[35] P. L. Sala, R. Sim, A. Shokoufandeh, and S. J. Dickinson, "Technical report: Optimal landmark acquisition for vision-based navigation," tech. rep., Dept. Computer Science, University of Toronto, 2004. http://www.cs.toronto.edu/∼psala/Papers/MORDP-TR.pdf.

[36] R. Sim, S. Polifroni, and G. Dudek, "Comparing attention operators for learning landmarks," Tech. Rep. MRL-03-02, Mobile Robotics Lab, McGill University, Montreal, QC, October 2001.

[37] R. Sim and G. Dudek, "Self-organizing visual maps," Tech. Rep. MRL-03-01, Mobile Robotics Lab, McGill University, Montreal, QC, October 2003.

[38] R. Sim, "Bayesian exploration for mobile robots," Tech. Rep. CIM-03-02, Centre for Intelligent Machines, McGill University, Montreal, QC, June 2000.

[39] R. Sim, "Mobile robot localization from learned landmarks," Tech. Rep. CIM-98-03, Centre for Intelligent Machines, McGill University, Montreal, QC, 1998.

Statement of Research Interests

My main research objectives are to develop autonomous vision-based robotic systems that can learn to work in cooperation with their human counterparts. The potential applications for such systems range from household and workplace assistants (including the performance of household chores, and providing medical assistance), to operation in remote and dangerous environments (such as deep sea and deep space exploration).

My research spans three main areas, elaborated further below, aimed at the ultimate goal of producing intelligent autonomous systems. My work on *vision-based mapping* enables a robot to automatically construct a visual representation of its environment. My subsequent work on *autonomous exploration* is aimed at simultaneously optimizing the rate at which a robot explores its environment and the quality of the resulting map. My ongoing work takes these two core problems (mapping and exploration), and aims to develop *scalable vision-based representations* for very large environments, with the aim of facilitating the general navigational and search capabilities enjoyed by humans. In the following sections I will elaborate on these ideas further.

Vision-based Mapping

My research focuses on vision-based sensing. Cameras are inexpensive passive sensors that provide a rich stream of information about a robot's surroundings. By contrast, conventional laser range-finding sensors, such as those commonly found on mobile robots, are expensive, energy-emitting and information-impoverished. My graduate work concentrated on the question of how the visual world can be represented, and how such representations can facilitate robot localization and navigation. The central contribution of this work was to develop the *visual map* framework, an architecture for learning models of environmental features as a robot moves through the world. The resulting models are useful for solving inference problems, such as, "What is the probability I am in the kitchen", or "How far am I from the stove?". A key feature of this work is that it represents a turn-key solution to robotic environment learning that does not depend on extensive calibration of the robot's sensors or supervision as the robot explores the world.

The core idea of the visual map framework is to use a model of visual saliency to extract local features from images, and to subsequently learn a generative model of each feature as a function of the robot's pose in the world. These models enable the robot to predict future observations of the features from arbitrary poses. In addition, we compute an independent noise model for each feature, producing a probability distribution over observations conditioned on the robot's pose. These models can be subsequently applied in a Bayesian framework for solving localization and path planning tasks.

Information-driven Exploration

In the course of my graduate studies, a key question that emerged was the problem of map construction in the face of uncertainty. As a robot moves through the world, its actions have uncertain outcomes, and errors in the robot's position estimate can become unbounded. While many researchers in the robotics community have considered this problem in the context of simultaneous localization and mapping (SLAM), my main interest in the problem is from the perspective of exploration. In its most general form, how does a robot go about actively learning about the world in such a way that its inferences are accurate and robust? Expressed more colloquially, the goal of this work is to imbue a robot with a healthy sense of curiosity.

My Ph.D. research examined a variety of policies for exploring an unknown environment and building a visual map, with the aim of optimizing the accuracy and utility of the map. In my post-doctoral work I have considered this problem from the perspective of *information-optimal* exploration in the context of solving the SLAM problem. One common approach to solving this problem is to apply an Extended Kalman Filter (EKF), which assumes a joint Gaussian distribution over map features and robot positions, and makes linear approximations to the underlying non-linear motion and observation processes. My work on this problem has demonstrated that these simplifying assumptions cause difficulties for information-optimal exploration approaches that can lead to incorrect maps. I have demonstrated new exploratory policies that address these problems while ensuring a high rate of information gain. I have also been actively collaborating with Nicholas Roy at M.I.T. on the *global exploration* or *non-myopic planning* problem. That is, how can a robot plan for exploration in a way that optimizes the rate at which it learns *over the long term*? This approach is in contrast to most solutions which make myopic or greedy decisions about how to explore the world. We have proposed a planning approach that uses dynamic programming to compute a good approximation to the globally optimal exploratory plan.

Large-scale environment representations

In my time as a postdoctoral fellow at UBC and Toronto, the core questions I have been addressing are those of exploring and representing very large visual domains. As the extent of a robot's environment grows, the number of visual features that must be stored becomes very large. Furthermore, core cognition problems, such as recognizing when one has arrived at a previously visited location by an alternate route, become more difficult to solve. At UBC, working in cooperation with James Little and his students, I have developed a scalable architecture for visual map learning that enables a robot equipped with a camera to explore a large environment and construct a robust map in real-time. Our approach employs Monte-Carlo state estimation methods based on Rao-Blackwellised particle filters. Among the technological achievements of this architecture is the ability to build a map using a camera and no direct knowledge of the robot's (or human operator's) actions.

My related work at the University of Toronto, in collaboration with Pablo Sala, Sven Dickenson, and Ali Skohoufandeh, considered the problem of selecting an optimal subset of visual features that enables a robot to localize accurately without storing excessive redundancy in the map. This problem is also related to minimum description length (MDL) encoding. Our work proved the intractability of finding the minimal map and evaluated a variety of methods for approximating optimal solutions. My latest work with Sven Dickinson, sponsored in part by MacDonald, Dettwiler and Associates (MDA), has been to examine the problem of model augmentation, using real imagery to improve impoverished CAD models, with a specific focus on developing robust visual models of satellites for on-orbit servicing.

Future Work

Most current environment representations in robotics, with a few exceptions, explicitly encode the entire domain in which a given robot is situated. However, it is reasonable to assume that humans do not explicitly encode their entire visual world, and yet we are capable of quickly learning a representation and navigating within previously unseen environments. How, then, might a robot establish the same capability, acquiring the key features of very large worlds without explicitly representing the entire structure? For example, how can a robot learn generalizations, such as the concepts that consecutive room numbers in a building are usually grouped locally in space, and often alternate on either side of a hall? In my opinion, this representational gap must be overcome to achieve systems that are truly adaptive and autonomous, and will become the key problem for robotics research over the next decade. My goal for the near future is to examine how techniques in machine learning could be applied to enable a robot to infer a map of its world with only limited exploration. A solution to this problem would be a stepping stone towards making other useful inferences, such as how to safely navigate without an explicit obstacle representation, or likely locations to search for survivors in a rescue operation.

Over the long term, my goal is to establish a framework such that robots can learn to operate in novel, dynamic, and highly uncertain environments in a way that is natural (that is, human-like), robust, and useful to their human partners. While the core problems of representation, learning and planning will continue to require attention, the open problem of generalization will emerge as the dominant concern. It is truly exciting to think of the potential successes and applications of these research directions, and I expect that the coming years will yield many new and innovative solutions for enabling robotic autonomy.

Statement of Teaching Interests

Teaching is the foundation of a research-intensive university as it is only through the dissemination of knowledge that scientific progress can be achieved. For this reason, I have made teaching a priority throughout my academic career. In this statement I will describe my teaching experience and elucidate my approach.

As a postdoctoral fellow at the University of British Columbia, I requested the opportunity to teach undergraduate classes. To date, I have taught two core undergraduate courses in computer science.

- CPSC 219, Software Development Laboratory. This second-year summer course focused on programming tools, and lectures were complemented with a significant laboratory-based component. The challenge as a lecturer was to convey a big-picture view of various Unix tools (make, tsch, gcc, perl, java), while providing enough detail to help the students address their immediate goals in the lab. I supplemented the diversity of the material, and the limited contact time in lectures with frequent use of on-line fora, such as WebCT.

- CPSC 315, Introduction to Operating Systems. This third year course had a deeper lecture component, giving me the opportunity to address the core ideas in operating systems design, as well as discuss the historical context of modern operating systems. This course also had a strong emphasis on challenging lab work, and I viewed it as the one chance that many undergraduates in computer science and engineering get to roll up their sleeves and learn the core essentials of modern systems, such as socket programming, resource allocation and concurrency.

The most important task for any teacher is to inspire a desire to learn. In my lectures, I strive to deliver the core content with a focus on the big picture. I present real-world situations, particularly those that are related to my research interests, to exemplify, illustrate or provide context for a topic. I am also a strong believer that students "learn-by-doing", and provide them with assignments and course work that focuses on experimentation and implementation. Finally, communications are an essential aspect of my approach to teaching; through interaction in-class, soliciting feedback through mid-term surveys, or the use of on-line course tools, such as WebCT, to facilitate discussion, I work to maintain a receptive, helpful and responsive presence.

Prior to my arrival at UBC, I had the opportunity teach at the high school level. For four years I gave courses in Calculus and Linear Algebra to advanced Grade 11 students at MIND High School in Montreal. My responsibilities at the time included all aspects of defining the curriculum, teaching, evaluation and grading. I took enormous pleasure from opportunities to share my practical experiences in these subjects. Our annual trips to the Mobile Robotics Lab at McGill were a chance to see many aspects of the material at work.

At MIND I had the opportunity to develop my skills in the classroom, to develop a challenging curriculum, and to develop effective methods of evaluation. Much more significant to me, though, is the opportunity I've had to encourage students to pursue careers in science and engineering, and to inspire a deeper affinity for science and mathematics in those students that went on to pursue non-scientific careers.

My teaching interests in undergraduate computer science and engineering span the breadth of my research interests, as well as my undergraduate experiences. I am comfortable teaching most of your core undergraduate curriculum, with particular emphasis on data structures and programming languages, computer architecture, operating systems, and graphics. At the graduate level, my teaching interests include computer vision, robotics, and machine learning. I am confident that I would be an asset to your program.

Finally, my belief is that graduate supervision is an important part of any professor's teaching dossier. At UBC I have acted as a supervisor for summer undergraduate research assistants, and at McGill I assisted in the supervision of senior undergraduate research projects. I look forward to the opportunity to supervise graduate students and engage them in the process of establishing their own careers in research. My graduate and post-graduate experience has been thoroughly positive and I have developed a good sense of best-practices for graduate supervision.

I am looking forward to the opportunity to teach on a regular basis. I am confident that my passion for learning, for the subject material and for the act of teaching itself will be an essential contribution to your department.