

# DENNIS STRELOW

4970 Cherry Avenue unit 212, San Jose, CA 95118 · (408) 598-1929  
[strelow@gmail.com](mailto:strelow@gmail.com) · [dennis-strelow.com](http://dennis-strelow.com) · [linkedin.com/in/dennis-strelow](https://www.linkedin.com/in/dennis-strelow)

---

Seeking a position in software development and/or computer vision.

In brief: I have more than 20 years of training and experience in software and 3D computer vision, and practical experience in massively parallel computing and optimization. Among other work, I have a Ph.D. in computer science from Carnegie Mellon and I was a Google software engineer for 10 years.

## EMPLOYMENT AND EDUCATION

### SUMMER 2016 – SPRING 2018

#### **DISTINGUISHED FELLOW**, MAGIC LEAP, SUNNYVALE, CA

- Real-time visual-inertial odometry (VIO) on Magic Leap's augmented reality glasses, for estimating the glasses' 6 DOF pose and linear velocity.
- Real-time IMU calibration on the glasses: estimation of gyro and accelerometer biases, time offset between IMU and camera observations, and 3 DOF rotation between the IMU and camera.
- Real-time visual and inertial gravity estimation for rendering content upright on the glasses.
- Extensive evaluation with live quality assurance testers and recordings with Vicon ground truth.

### SPRING 2006 – SPRING 2016

#### **[SENIOR] SOFTWARE ENGINEER**, GOOGLE, MOUNTAIN VIEW, CA

*Google Research embedded with Google Geo, Summer 2012 – Spring 2016*

- Improve Google Maps' 3D city models by adding high-resolution Street View imagery to models textured with lower-resolution aerial images. Tasks included: ultra-wide baseline matching between aerial images and Street View panoramas, tree detection in Street View images, implementing public takedown requests for Street View imagery used in the models.
- Show semantic information from 2D Maps (e.g., business type icons) in 3D in Street View.
- Database and web UI improvements for indoor Street View.
- Show labels for landmarks (e.g., craters) in Google Planets (tour the solar system inside Google Maps).

*Google Research, Summer 2007 – Spring 2012*

- General and nested Wiberg minimization: beat expectation-maximization for minimizing functions of two sets of variables. Demonstrated for L1 minimization, least squares, and maximum likelihood; on bundle adjustment, projective bundle adjustment, and Poisson matrix factorization. Intern Qifan Wang extended L1 Wiberg minimization to max margin optimization for multiple instance learning.
- "Winner-take-all" hashing for image classification and patch matching.
- Group sparse coding: mixed-norm optimization for learning sparse coding dictionaries for local image descriptors, and image classification using the resulting dictionaries.
- Ultra-fast visual features for Image Safe Search classification.
- VOC image classification with chi-squared SVMs.
- Saliency-based autocropping for Google News images.
- Papers in ECCV, ICCV, NIPS, and PAMI. Sole author of a CVPR oral paper (2.5% acceptance rate).
- Massively parallel programming with MapReduce for machine learning tasks.

*Google Image Search, Spring 2006 - Summer 2007*

- "Find Similar Images": developed interactive demo running out of the data center and improved the algorithm's visual features.
- Core Image Search classes: image class and color space conversion.

**SPRING 2005 – SPRING 2006**

**RESEARCH SCIENTIST, HONEYWELL, MINNEAPOLIS, MN**

- Visual odometry for ground vehicles using stereo.
- Particle filter for tracking ground vehicles with a sparse network of vibration sensors.
- Algorithm for automatically transferring detail between datasets of different modalities (e.g., transfer MRI detail to a lower-resolution CT image, transfer satellite image detail to a lower-resolution 3D terrain map).
- Pan-tilt-zoom camera calibration.

**FALL 1998 – FALL 2004**

**PH.D. IN COMPUTER SCIENCE, CARNEGIE MELLON UNIVERSITY, PITTSBURGH, PA**

- Ph.D. thesis focused on making image-based, six degree of freedom motion estimation robust for autonomous vehicles and 3D modeling. Supervised by Professor Sanjiv Singh.
- Tightly integrated visual and inertial odometry, using both Levenberg-Marquardt and an IEKF.
- Robust image feature tracking: epipolar estimation with SIFT followed by 1-D Lucas-Kanade tracking along epipolar lines, and other heuristics.
- Relocalization for long image sequences using SIFT.
- Structure-from-motion and 6 DOF camera-to-mirror calibration for non-central catadioptric cameras.
- Real-time trinocular stereo for highway obstacle detection, on reconfigurable vision computer (with Toyota).
- Visual stair detection for humanoid robots using perceptual grouping (with Honda).
- NASA Graduate Student Researcher Program (GSRP) fellowship, summer 2001 - summer 2004.
- Teaching assistant for data structures and computer vision classes.
- Papers in IJRR, CVPR, and other conferences.

**FALL 1996 – SUMMER 1998**

**SOFTWARE ENGINEER, K2T INC., PITTSBURGH, PA**

- Commercialization of Carnegie Mellon's algorithms for 3D modeling from video.
- Structure-from-motion with the factorization method (including orthographic, weak perspective, paraperspective projection models) and bundle adjustment, Lucas-Kanade feature tracking, camera calibration, absolute orientation, and textured polyhedral modeling.
- End-to-end system for semi-automatic 3D modeling of architecture from still photos, including GUI.

**FALL 1994 – SUMMER 1996**

**M.S. IN COMPUTER SCIENCE, UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN**

- M.S. thesis with the biological imaging group at the National Center for Supercomputing Applications (NCSA).
- Thesis developed a new algorithm for constructing polyhedral surface models from volumetric data (e.g., cortex models from MRI data) and the visualization and manipulation of these models in NCSA's CAVE virtual reality environment.
- Classes in computer vision, image processing, artificial intelligence, numerical analysis, programming languages, algorithms, computer architecture, and operating systems.
- Teaching assistant for intro CS for engineers and artificial intelligence classes. Excellent Teaching Assistant award for spring 1995 and fall 1995.

**SUMMER 1995**

**GRADUATE RESEARCH ASSISTANT, LOS ALAMOS NATIONAL LABORATORY, LOS ALAMOS, NM**

- Medical image analysis: active-contour-like algorithm for finding rib cage volumes in computed tomography (CT) images.
- Document image analysis: word segmentation using clustering.

#### SUMMER 1994

#### RESEARCH ASSISTANT, NATIONAL SOLAR OBSERVATORY, SUNSPOT, NM

- Algorithms for tracking erratic sunspot motion in solar images, using blob detection and similarity.
- Algorithms for measuring flare intensity in solar images.

#### FALL 1990 – SPRING 1994

#### B.S. IN COMPUTER SCIENCE AND MATHEMATICS, UNIVERSITY OF WISCONSIN-MADISON

- Advanced computer science classes: artificial intelligence, compilers, computer architecture, and theoretical computer science.
- Advanced math classes: numerical analysis, numerical linear algebra, advanced calculus, differential equations, abstract algebra, and symbolic logic.
- Graduated with Distinction (GPA in top of graduating class).

#### SUMMER 1993

#### SUMMER STUDENT, UNIVERSITY OF NEBRASKA-LINCOLN

- Investigated AVHRR satellite image restoration, showing that optimal restoration using small convolution kernels should include a model of the display.
- Image processing and computer vision classes.
- NSF Research Experience for Undergraduates (REU) program.

## PAPERS

Dennis Strelow, Qifan Wang, Luo Si, and Anders Eriksson. General, nested, and constrained Wiberg minimization. *IEEE Pattern Analysis and Machine Intelligence (PAMI)*, Vol. 38, No. 9, September 2016, pp. 1803-1815.

Dennis Strelow. General and nested Wiberg minimization: L2 and maximum likelihood. *European Conference on Computer Vision (ECCV 2012)*, Florence, Italy, October 2012.

Dennis Strelow. General and nested Wiberg minimization. **Oral presentation (2.5% acceptance rate)**, *Computer Vision and Pattern Recognition (CVPR 2012)*, Providence, RI, June 2012.

Jay Yagnik, Dennis Strelow, David A. Ross, and Rwei-sung Lin. The power of comparative reasoning. *International Conference on Computer Vision (ICCV 2011)*, Barcelona, Spain, November 2011.

Samy Bengio, Fernando Pereira, Yoram Singer, and Dennis Strelow. Group sparse coding. *Neural Information Processing Systems (NIPS 2009)*, Vancouver, December 2009.

Dennis Strelow and Sanjiv Singh. Long-term motion estimation. *International Symposium on Experimental Robotics (ISER 2006)*, Rio de Janeiro, July 6-10 2006.

Dennis Strelow and Sanjiv Singh. Motion estimation from image and inertial measurements. *The International Journal of Robotics Research*, Vol. 23, No. 12, December 2004, pp. 1157-1195.

Dennis Strelow and Sanjiv Singh. Online motion estimation from image and inertial measurements. *Workshop on the Integration of Vision and Inertial Sensors (INERVIS 2003)*, Coimbra, Portugal, June 2003.

Dennis Strelow and Sanjiv Singh. Reckless motion estimation from omnidirectional image and inertial measurements. *IEEE Workshop on Omnidirectional Vision and Camera Networks (OMNIVIS 2003)*, Madison, Wisconsin, June 2003.

Dennis Strelow and Sanjiv Singh. Optimal motion estimation from visual and inertial measurements. *IEEE Workshop on Applications of Computer Vision (WACV 2002)*, Orlando, Florida, December 2002.

Henele Adams, Sanjiv Singh and Dennis Strelow. An empirical comparison of methods for image-based motion estimation. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2002)*, Lausanne, Switzerland, October 2002.

Sanjiv Singh, George Kantor and Dennis Strelow. Recent results in extensions to simultaneous localization and mapping. *International Symposium on Experimental Robotics (ISER 2002)*, Sant'Angelo d'Ischia, Italy, July 2002.

Dennis Strelow, Jeffrey Mishler, David Koes and Sanjiv Singh. Precise omnidirectional camera calibration. *IEEE Computer Vision and Pattern Recognition (CVPR 2001)*, Kauai, Hawaii, December 2001.

Dennis Strelow, Jeff Mishler, Sanjiv Singh and Herman Herman. Extending shape-from-motion to noncentral omnidirectional cameras. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2001)*, Maui, Hawaii, October 2001.

Dan Foygel and Dennis Strelow. Reducing web latency with hierarchical cache-based prefetching. *IEEE Workshop on Scalable Web Services (SWS 2000)*, Toronto, Canada, June 2000.

Nathaniel Daw, Seth Goldstein and Dennis Strelow. Embedded compilation for multimedia applications (abstract). *IEEE Symposium on Field-Programmable Custom Computing Machines (FCCM 2000)*, Napa Valley, California, May 2000.

Dennis Strelow, Warren Gardner, Regis Hoffman, Jeff Mishler and Fred Persi. A shape and motion engine for parameterized models. *Proceedings of the 1998 DARPA Image Understanding Workshop*, Monterey, California, November 1998.

Dennis Strelow, Clinton S. Potter and Paul C. Lauterbur. A virtual environment for the visualization of functional MRI data (abstract). *Current progress in functional brain mapping: science and applications*, T. Yuasa, J.W. Prichard and S. Ogawa, editors, Nishumira/Smith-Gordon, November 1998.

Dennis W. Strelow, Clinton S. Potter and Paul C. Lauterbur. The construction and visualization of surfaces from MRI data. NCSA Technical Report 036, August 1996.

E. Reiger, D.F. Neidig, D.W. Engfer and D. Strelow. The role of high-energy protons and electrons in powering the solar white light flare emissions. *Solar Physics*, 167:307-320, 1996.

S. Reichenbach, D. Koehler and D. Strelow. Restoration and reconstruction of AVHRR images. *IEEE Transactions on Geoscience and Remote Sensing*, 33(4):997-1007, May 1995

## **SELECTED PATENTS**

Dennis Strelow, Craig Robinson, and Samuel Felix de Sousa, Jr. Aligning panoramic imagery and aerial imagery.

Dennis Strelow and Jay Yagnik. General and nested Wiberg minimization.

Luca Bertelli, Dennis Strelow, and Sally Goldman. Foreground object detection from multiple images.

Dennis W. Strelow. Method and apparatus for propagating high resolution detail between multimodal data sets.

## **SELECTED TALKS**

### **OCTOBER 2015**

Ultra-wide baseline matching. Technical University of Vienna.

### **DECEMBER 2011 – JUNE 2012**

General and nested Wiberg minimization. Carnegie Mellon, University of Wisconsin-Madison, Willow Garage, and others.

### **OCTOBER 2003 – JANUARY 2006**

Motion estimation from image and inertial measurements. Google, Northrup Grumman, NASA Ames, and others.

### **DECEMBER 2001**

Extending shape-from-motion to noncentral omnidirectional cameras. University of Wisconsin-Madison.

## **TEACHING**

### **SPRING 2000**

Teaching assistant, Computer Vision, Carnegie Mellon University, Computer Science Department.

### **FALL 1999**

Teaching assistant, Data Structures, Carnegie Mellon University, Computer Science Department.

### **FALL 1995**

Teaching assistant, Artificial Intelligence, University of Illinois at Urbana-Champaign, Department of Computer Science.

### **FALL 1994 – SPRING 1995**

Teaching assistant, Computer Science for Engineering Students, University of Illinois at Urbana-Champaign, Department of Computer Science.

## **OTHER**

### **FALL 2015**

Ph.D. committee member, Samuel Felix de Sousa, Jr., Technical University of Vienna.

## SPRING 2008

Ph.D. committee member, Craig Yoshioka, Scripps Institute.

## TECHNICAL SKILLS

- Operating systems: more than 20 years Linux/Unix development.
- Languages: more than 20 years C++ development. Limited experience with many other languages.
- Tools: revision systems including Perforce and git, debuggers (e.g., gdb), performance profilers (e.g., gprof), heap checkers (e.g., valgrind).
- Third-party libraries: LAPACK and image libraries (e.g., jpeglib). Limited experience with OpenCV and Ceres.
- Computer vision: more than 20 years experience in 3D vision, including structure-from-motion, visual odometry, image and inertial fusion, feature extraction and tracking, wide baseline matching, depth from stereo and stereo calibration, omnidirectional cameras, absolute orientation, and camera calibration. Image classification. Limited experience with Markov random fields (Gibbs sampling), image segmentation, and image restoration.
- Numerical methods and optimization: more than 20 years experience, including matrix factorization, Levenberg-Marquardt, maximum likelihood estimation, RANSAC, iterated extended Kalman filters, and L1 minimization via successive linear programming. Very limited experience with branch and bound.
- Parallel computing: Google's MapReduce (> 5 years) and Flume (2 years). Limited experience with RPC, multithreading, and OpenMP.
- Machine learning: limited experience with distance learning and image classification using SVMs, AdaBoost, and sparse coding. Good understanding of probability.
- LaTeX.

Interested in learning: web programming, mobile programming, GPU programming, deep learning; and more visual odometry, machine learning and optimization.