Manolis Savva

msavva@sfu.ca · 650-485-1626 · https://msavva.github.io · 715 Sheraton Dr, Sunnyvale, CA 94087

Education

Stanford University
Ph.D. in Computer Science (Thesis: Body-centric Understanding of 3D Environments)
Conferred September 2016

Stanford University
MS in Computer Science
Conferred December 2012

Cornell University
B.A. in Physics and Computer Science
Conferred December 2009

Refereed Publications

Text2Shape: Generating Shapes from Natural Language by Learning Joint Embeddings
Kevin Chen, Christopher B. Choy, Manolis Savva, Angel X. Chang, Thomas Funkhouser, Silvio Savarese
ACCV 2018

Functionality Representations and Applications for Shape Analysis
Ruizhen Hu, Manolis Savva, Oliver van Kaick
Eurographics STAR, Computer Graphics Forum 2018

Im2Pano3D: Extrapolating 360 Structure and Semantics Beyond the Field of View
Shuran Song, Andy Zeng, Angel X. Chang, Manolis Savva, Silvio Savarese, Thomas Funkhouser
Proceedings of CVPR 2018

Matterport3D: Learning from RGB-D Data in Indoor Environments
A. Chang, A. Dai, T. Funkhouser, M. Halber, M. Nießner, M. Savva, S. Song, A. Zeng, Y. Zhang
Proceedings of 3DV 2017

Cross-modal Attribute Transfer for Rescaling 3D Models
Lin Shao, Angel X. Chang, Hao Su, Manolis Savva, Leonidas Guibas
Proceedings of 3DV 2017

ScanNet: Richly-annotated 3D Reconstructions of Indoor Scenes
Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, Matthias Nießner
Proceedings of CVPR 2017

Physically-Based Rendering for Indoor Scene Understanding Using Convolutional Neural Networks
Yinda Zhang, Shuran Song, Ersin Yumer, Manolis Savva, Joon-Young Lee, Hailin Jin, Thomas Funkhouser
Proceedings of CVPR 2017

Semantic Scene Completion from a Single Depth Image
Shuran Song, Fisher Yu, Andy Zeng, Angel X. Chang, Manolis Savva, Thomas Funkhouser
Proceedings of CVPR 2017

PiGraphs: Learning Interaction Snapshots from Observations
Manolis Savva, Angel X. Chang, Pat Hanrahan, Matthew Fisher, Matthias Nießner
Proceedings of ACM SIGGRAPH 2016

Activity-centric Scene Synthesis for Functional 3D Scene Modeling
Matthew Fisher, Manolis Savva, Yangyan Li, Pat Hanrahan, and Matthias Nießner
Proceedings of ACM SIGGRAPH Asia 2015

Text to 3D Scene Generation with Rich Lexical Grounding
Angel X. Chang, Will Monroe, Manolis Savva, Christopher Potts, and Christopher D. Manning
Proceedings of ACL 2015

SceneGrok: Inferring Action Maps in 3D environments
Manolis Savva, Angel X. Chang, Pat Hanrahan, Matthew Fisher, and Matthias Nießner
Proceedings of ACM SIGGRAPH Asia 2014

Learning Spatial Knowledge for Text to 3D Scene Generation
Angel X. Chang, Manolis Savva, and Christopher D. Manning
Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing

TransPhoner: Automated Mnemonic Keyword Generation
Manolis Savva, Angel X. Chang, Christopher D. Manning, and Pat Hanrahan
Proceedings of CHI 2014

Example-based Synthesis of 3D Object Arrangements
Matthew Fisher and Daniel Ritchie and Manolis Savva and Thomas Funkhouser, and Pat Hanrahan
Proceedings of ACM SIGGRAPH Asia 2012

GraphPrism: Compact Visualization of Network Structure
Sanjay Kairam, Diana MacLean, Manolis Savva, and Jeffrey Heer
Advanced Visual Interfaces 2012

ReVision: Automated Classification, Analysis and Redesign of Chart Images
Manolis Savva, Nicholas Kong, Arti Chhajta, Fei-Fei Li, Maneesh Agrawala, and Jeffrey Heer
Proceedings of ACM UIST 2011, Notable Paper Award

Characterizing Structural Relationships in Scenes Using Graph Kernels
Matthew Fisher, Manolis Savva, and Pat Hanrahan
Proceedings of ACM SIGGRAPH 2011

Back-action-evading Measurements of Nanomechanical Motion
Jared Hertzberg, Tristan Rocheleau, Tchefor Ndukum, Manolis Savva, Aashish Clerk, and Keith Schwab
Nature Physics vol. 6, no. 3, pp. 213–217, 2009

Technical Reports and Preprints

On Evaluation of Embodied Navigation Agents
Peter Anderson, Angel X. Chang, Devendra Singh Chaplot, Alexey Dosovitskiy, Saurabh Gupta, Vladlen Koltun, Jana Kosecka, Jitendra Malik, Roozbeh Mottaghi, Manolis Savva, Amir R Zamir
arXiv:1807.06757 [cs.AI], Jul 2018

MINOS: Multimodal Indoor Simulator for Navigation in Complex Environments
Manolis Savva, Angel X. Chang, Alexey Dosovitskiy, Thomas Funkhouser, Vladlen Koltun
arXiv:1712.03931 [cs.LG], Dec 2017

Learning Where to Look: Data-Driven Viewpoint Set Selection for 3D Scenes
Kyle Genova, Manolis Savva, Angel X. Chang, Thomas Funkhouser
arXiv:1702.04405 [cs.CV], Apr 2017

SceneSuggest: Context-driven 3D Scene Design
Manolis Savva, Angel X. Chang, Maneesh Agrawala
arXiv:1703.00061 [cs.GR], Feb 2017)

SceneSeer: 3D Scene Design with Natural Language
Angel X. Chang, Mihail Eric, Manolis Savva, Christopher D. Manning
arXiv:1703.00050 [cs.GR], Feb 2017

ShapeNet: An Information-Rich 3D Model Repository
Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu
arXiv:1512.03012 [cs.GR], Dec 2015

Workshop Papers, Tutorials, and Extended Abstracts

Linking WordNet to 3D Shapes
Angel X. Chang, Rishi Mago, Pranav Krishna, Manolis Savva, Christiane Fellbaum
Proceedings of Global WordNet Conference 2018

Directions in Shape Analysis towards Functionality
Ruizhen Hu, Oliver van Kaick, Youyi Zheng, Manolis Savva
SIGGRAPH Asia Course, 2016

SHREC'17 Track: Large-Scale 3D Shape Retrieval from ShapeNet Core55
M. Savva, F. Yu, H. Su, A. Kanezaki, T. Furuya, R. Ohbuchi, Z. Zhou, R. Yu, S. Bai, X. Bai, M. Aono, A. Tatsuma S. Thermos, A. Axenopoulos, G. Th. Papadopoulos, P. Daras, X. Deng, Z. Lian, B. Li, H. Johan, Y. Lu, S. Mk
Eurographics Workshop on 3D Object Retrieval 2017

SHREC'16 Track: Large-Scale 3D Shape Retrieval from ShapeNet Core55
M. Savva, F. Yu, H. Su, M. Aono, B. Chen, D. Cohen-Or, W. Deng, H. Su, S. Bai, X. Bai, N. Fish, J. Han E. Kalogerakis, E. G. Learned-Miller, Y. Li, M. Liao, S. Maji, A. Tatsuma, Y. Wang, N. Zhang, Z. Zhou

Eurographics Workshop on 3D Object Retrieval 2016

Semantically-Enriched 3D Models for Common-sense Knowledge
Manolis Savva, Angel X. Chang, and Pat Hanrahan
CVPR 2015 Vision meets Cognition Workshop

On Being the Right Scale: Sizing Large Collections of 3D Models
Manolis Savva, Angel X. Chang, Gilbert Bernstein, Christopher D. Manning, Pat Hanrahan
SIGGRAPH Asia 2014 Workshop on Indoor Scene Understanding: Where Graphics meets Vision

Learning Affordance Maps by Observing Interactions
Manolis Savva, Angel X. Chang, Matthew Fisher, Matthias Nießner, and Pat Hanrahan
CVPR 2014 Workshop on Functionality, Physics, Intentionality and Causality

Interactive Learning of Spatial Knowledge for Text to 3D Scene Generation
Angel X. Chang, Manolis Savva, and Christopher D. Manning
Proceedings of the ACL 2014 Workshop on Interactive Language Learning, Visualization, and Interfaces

Semantic Parsing for Text to 3D Scene Generation
Angel X. Chang, Manolis Savva, and Christopher D. Manning
Proceedings of the ACL 2014 Workshop on Semantic Parsing

Invited Talks

Computer Science Department, University College London — London, UK
Understanding 3D Environments through Embodiment — September 2018

Beijing Film Academy Roundtable — Vancouver, BC
3D Content Creation for Learning through Simulation — August 2018

Embodied Agents and Environments Workshop, FAIR, Facebook — Menlo Park, CA
MINOS: Multimodal Indoor Simulator — February 2018

Visual Models Tech Talk, Google Brain, Google — Mountain View, CA
Generation of 3D Environments through Embodied Analysis and Synthesis — February 2018

Intel/NSF Visual and Experiential Computing Retreat, Intel Labs — Santa Clara, CA
From Virtual to Real and Back Again — December 2017

Visual Computing Center, Shenzhen University — Shenzhen, China
Towards Holistic 3D Scene Understanding — July 2017

Perceptual Computing Group, Intel — Jerusalem, Israel
Towards Holistic 3D Scene Understanding — February 2017

Adobe Research — San Jose, CA
PiGraphs for Text to Interaction Snapshot Generation — May 2015

Vicarious — Union City, CA
Common-sense Knowledge for Virtual Environments — October 2015

Computer Science Department, University of California, Berkeley — Berkeley, CA
Semantic Understanding of Objects, Actions, and Environments — September 2014

Employment

Visiting Researcher — Menlo Park, CA
Facebook AI Research — July 2018 –
Research on 3D simulation platforms for embodied AI agents.

Visiting Research Collaborator — Princeton, NJ
Princeton University — Sep 2017 – Sep 2018
Mentored by Prof. Tom Funkhouser while on sabbatical at Google and Stanford. Research in 3D scene understanding and human-centric analysis of 3D environments.

Research Engineer Contractor — Mountain View, CA
AutoRoboto LLC, on-site at Google — Sep 2017 – July 2018
Mentored by Prof. Tom Funkhouser while on sabbatical at Google and Stanford. Research in 3D scene understanding and human-centric analysis of 3D environments.

Postdoctoral Research Associate — Princeton, NJ
Princeton University — Aug 2016 – Aug 2017
Mentored by Prof. Tom Funkhouser. Research in 3D scene understanding and human-centric analysis of 3D environments.

Research Intern Tokyo, Japan
Square Enix Co., Ltd. Fall 2013
Mentored by Remi Driancourt. Investigated geometric analysis methods for 3D model part segmentation and recombination in order to enable automated synthesis of object variations. Outcome was a prototype system and research talk to the advanced technologies division of Square Enix.

Research Assistant Ithaca, NY
Program of Computer Graphics, Cornell University May 2009 – May 2010
Mentored by Prof. Steve Marschner. Acquired material appearance as Bidirectional Texture Functions using gonioreflectometer experimental setup; investigated data compression and real-time rendering of captured data

Research Assistant Ithaca, NY
Laboratory of Atomic and Solid State Physics, Cornell University Fall 2007
Mentored by Prof. Keith Schwab. Designed, prototyped and implemented microwave cavity electromagnetic filter and cryogenic probes for achieving near absolute zero cooling of nano-mechanical resonator circuits.

Teaching and Mentoring

Teaching Fellow Stanford, CA
Introduction to Computer Graphics and Imaging (Stanford CS 148) Summer 2015
Instructor for course. Created lecture material, taught lectures, designed and graded assignments and exams

Research Mentor Stanford, CA
Stanford RA and CURIS Programs 2013 – 2016
Mentored two masters students in their research assistanships and four undergraduate students as part of the Stanford CS Undergraduate Research Internship program. Students contributed significantly to active research projects

Course Assistant Stanford, CA
Introduction to Computer Graphics and Imaging (Stanford CS 148) Summer 2011
Advised students in office hours, designed and graded exams and programming assignments, gave guest lectures

Teaching Assistant Ithaca, NY
Introduction to Scientific Computing (Cornell CS 3220) Spring 2010
Advised students in office hours, designed and graded exams and programming assignments

Teaching Assistant Ithaca, NY
Introduction to Computer Graphics (Cornell CS 4620) Fall 2009
Advised students in office hours, graded exams and programming assignments

Service

Program Committee: CVM 2019; GI 2019; Pacific Graphics 2018; SIGGRAPH Asia 2017 Briefs and Posters; SMI 2018

Reviewer: CoRL 2017; CVPR 2018; CHI 2014,2017; EuroGraphics 2018; NAACL-HLT 2018 SpLU; SIGGRAPH 2013,2015-18; SIGGRAPH Asia 2012,2016-18; SMI 2018, TVCG 2013,2018, UIST 2014,2016-17

Workshop Organizer: ECCV 2018 workshop on Visual Learning and Embodied Agents in Simulation Environments; ICCV 2017 workshop on Learning to See from 3D Data; Eurographics 3DOR 2016/2017 SHREC Track — Large-scale 3D Shape Retrieval from ShapeNet Core55

Skills

Languages: Modern Greek and Bulgarian (native tongues), English (fluent), Japanese (advanced proficiency), Mandarin (intermediate proficiency), German (basic proficiency)

Martial Arts: practitioner of Aikido, served as president of Cornell Aikido Club for 4 years

Honors and Awards

SGP 2018 Dataset Award for ShapeNet

Stanford Graduate Fellowship 2012-2015

ACM UIST Notable Paper Award (for ReVision paper, at UIST 2011)

Stanford School of Engineering Akiko Yamazaki and Jerry Yang Engineering Fellowship 2010-2011

CASP–Fulbright undergraduate scholarship (US-Cyprus exchange student scholarship, 2005-2009)

Undergraduate Teaching Assistant Excellence Award (Computer Science Department, Cornell University, 2009)

Distinguished Leadership Award for Aikido Club presidentship (Cornell Student Activities Office, 2009)

Robert J. Smith award for most promising student of Japanese (Asian Studies Department, Cornell University, 2006)

Highest international score award for GCE A–Level Physics Advanced Extension (administered by Edexcel, 2003)

Graduating class valedictorian (2003 class of American Academy Larnaca, Cyprus)

References

Pat Hanrahan
Canon USA Professor of Computer Science
hanrahan@cs.stanford.edu

Thomas Funkhouser
Professor of Computer Science
funk@cs.princeton.edu

Leonidas J. Guibas
Paul Pigott Professor of Computer Science and Electrical Engineering
guibas@cs.stanford.edu

Niloy Mitra
Professor
n.mitra@cs.ucl.ac.uk

Silvio Savarese
Associate Professor
ssilvio@stanford.edu