

George Papandreou

Education

- 2003–2009 **Ph.D. in Electrical & Computer Engineering**, *National Technical University of Athens, Greece*.
Ph.D. Thesis: Image Analysis and Computer Vision: Theory and Applications in the Restoration of Ancient Wall Paintings (*Advisor*: Prof. Petros Maragos)
- 1998–2003 **Diploma/M.Eng. in Electrical & Computer Engineering**, *National Technical University of Athens, Greece*.
GPA: 9.54/10 (highest honors), ranked in top 1% of my class at the top engineering school in Greece.
Diploma Thesis: Fast Algorithms for the Evolution of Geodesic Active Contours with Applications in Computer Vision (*Advisor*: Prof. Petros Maragos)

Professional Experience

- 2013– **Research Assistant Professor**, *Toyota Technological Institute at Chicago*.
Research in computer vision and machine learning. Main research directions:
Deep learning for computer vision and speech recognition.
Advancement of the theory and practice of Perturb-and-MAP probabilistic models.
- 2009–2013 **Postdoctoral Research Scholar**, *University of California, Los Angeles*.
Member of the *Center for Cognition, Vision, and Learning (CCVL)*, working with Prof. Alan Yuille. Research in computer vision and machine learning. Main projects:
Developed a novel *Perturb-and-MAP* framework for random sampling and parameter learning in large-scale Markov random fields (NSF, AFOSR, & ONR-MURI, 2010–now).
Building whole-image probabilistic models from patch-level representations to handle in a unified manner visual tasks both at low-level (denoising) and at high-level (classification, recognition) (NSF, 2011–2013).
- 2003–2009 **Graduate Research Assistant**, *National Technical University of Athens, Greece*.
Member of the *Computer Vision, Speech Communication & Signal Processing (CVSP)* group (cvsp.cs.ntua.gr). Participated in national and European research projects in the areas of computer vision and multimodal processing:
THERA: Thesis work on digital restoration of the ancient wall paintings at the pre-historic settlement of Akrotiri, Thera. On site detailed photographic acquisition. Creation of high-resolution mosaics. Research on image inpainting techniques for digitally filling-in missing parts of the wall paintings.
MUSCLE & HIWIRE: Research on facial image analysis and multimodal feature fusion for audiovisual speech processing. Development of a real-time audiovisual speech recognition prototype.
ASPI: Research on speech inversion (recovery of vocal tract geometry) using audiovisual information.
- Summer 2006 **Visiting Researcher**, *Trinity College Dublin, Ireland*.
Project: Image inpainting with complex wavelets. *Mentor*: Prof. Anil Kokaram.

- 2001–2003 **Research Assistant**, *Demokritos Nat. Center for Scientific Research*, Greece.
Member of the Demokritos' Institute of Informatics & Telecommunications, participating in the European research project *WIN*. Worked on design and deployment of wireless WAN network and development of Internet services. Contributed in writing research proposals for follow-up funding. *Mentor*: Dr. Stelios Thomopoulos.

Research Interests

Deep learning, computer vision, multi-modal perception.

Scholarships and Awards

- 2006–2009 **Onassis Public Benefit Foundation**, Graduate studies scholarship.
2007 **IEEE International Workshop on Multimedia Signal Processing (MMSP)**, Student paper contest runner up award.
2004–2007 **Greek State Scholarships Foundation**, Graduate studies scholarship in the area of artificial intelligence and its applications.
1999–2003 **Latsis Public Benefit Foundation**, Undergraduate studies scholarship.
2001 **Technical Chamber of Greece**, Award to top ranking students in the School of E.C.E.
2000–2001 **National Technical University of Athens**, Papakyriakopoulos (awarded twice, 2000 and 2001) and Kritikos (2001) awards for excellence in mathematics.

Research Funding

- 2013 **NVIDIA Corporation**, Academic equipment grant (\$6,000).
Unsupervised Learning of Deep Generative Models for Image Recognition.

Teaching Experience

- 2014 **BASIS-14**, *Full-day Tutorial in conjunction with CVPR*.
Teaching (with I. Kokkinos, A. Bronstein, M. Bronstein) the tutorial BASIS-14 (BASes for Images and Surfaces). Linear and non-linear image and surface analysis methods as basis-based expansions.
- 2013–now **Machine Learning Seminar**, *Toyota Technological Institute at Chicago*.
Co-organizer (with R. Tomioka) of the TTI-C Machine Learning reading group.
- 2010–2013 **Group Seminar**, *UCLA*.
Organized Yuille's group weekly seminar. Led discussion on several of the presented papers.
- 2009–2013 **Student Mentoring**, *UCLA*.
Helped in supervising the research of PhD, MSc, and visiting students in Yuille's group.
- 2006–2009 **Graduate Teaching Assistant**, *National Technical University of Athens*, Computer Vision class (School of E.C.E., 8th semester), *Instructor*: Prof. Petros Maragos.
Served as lab and teaching assistant. Contributed to student homework grading and development of new lab exercises and teaching material. Assisted diploma thesis students working on computer vision projects at the CVSP group.
- 2007–2008 **Guest Class Lectures**, *National Technical University of Athens*, Computer Vision class (School of E.C.E., 8th semester), *Instructor*: Prof. Petros Maragos.
Gave invited class lectures on optical flow computation and image modeling in scale-spaces.

Invited Talks and Demonstrations

- 2014 **SIAM Conference on Imaging Science**, *Perturb-and-MAP Random Fields: Reducing Random Sampling to Optimization with Applications in Computer Vision*, Mini-Symposium on Challenges in Inverse Problems in Imaging.
- 2013 **ASA/IMS Spring Research Conference (SRC 2013) on Statistics in Industry and Technology**, *Perturb-and-MAP Random Fields: The Interplay between Random Sampling and Optimization, with Applications in Computer Vision*, Session on Statistics in Cognition and Vision.
- 2013 **Toyota Technological Institute, Chicago**, *Probabilistic Machine Learning in Computer Vision and Multimodal Perception*, Host: Prof. R. Urtasun.
- 2012 **University of California, Irvine**, *Random Sampling and Optimization in Probabilistic Modeling for Computer Vision*, Host: Prof. M. Welling.
- 2012 **NIPS 2012**, *Perturb-and-MAP Random Fields*, NIPS 2012 Workshop: Perturbations, Optimization, and Statistics.
- 2012 **Imperial College London, U.K.**, *Probabilistic Machine Learning in Computer Vision and Multimodal Perception*, Host: Dept. of Computing.
- 2012 **Univ. of Cambridge, U.K.**, *Random Sampling and Optimization in Probabilistic Modelling for Computer Vision*, Host: Dept. of Engineering.
- 2011 **EPFL, Lausanne, Switzerland**, *Perturb-and-MAP Random Fields*, Host: Prof. M. Unser.
- 2011 **UAB/Computer Vision Center, Barcelona, Spain**, *Perturb-and-MAP Random Fields*, Host: Prof. M. Vanrell.
- 2010 **Akrotiri Excavation, Thera, Greece**, *High resolution photo capture and digital restoration of missing parts in the wall paintings of Thera*, Princeton Univ. and Akrotiri Excavation 2010 Summer School – Reassembling and Studying the Thera Frescoes.
- 2009 **PENED Workshop**, *Digital Restoration of Missing Parts in the Wall Paintings of Thera*, PENED project Workshop: Digital Cultural Heritage Technologies with Applications at the Pre-Historic Settlement of Akrotiri-Thera.
- 2009 **Demokritos Nat. Center for Scientific Research, Athens, Greece**, *Multiresolution image models with application to image segmentation and digital restoration of missing parts in ancient wall paintings from Akrotiri-Thera*, Host: Dr. G. Potamianos.
- 2009 **Akrotiri Excavation, Thera, Greece**, *Capturing high resolution photos and automatically filling-in gaps in Thera wall paintings*, Princeton Univ. and Akrotiri Excavation 2009 Summer School – Reassembling and Studying the Thera Frescoes.
- 2009 **Institute for Language and Speech Processing, Athens, Greece**, *Audiovisual speech analysis*, Host: Dr. A. Vataki.
- 2008 **University of California, Los Angeles**, *Multi-resolution techniques for efficient image analysis and modeling*, UCLA Image Processing Research Group, Host: Prof. L. Vese.
- 2007-2008 **Real-time audiovisual speech recognition demonstrator**, Presented at the demo session of three IEEE conferences (MMSP-07, ICASSP-08, CVPR-08).
- 2006 **Trinity College Dublin**, *Feature uncertainty in multimodal fusion and learning*, SIGMEDIA group, Host: Prof. A. Kokaram.
- 2005 **MUSCLE Workshop**, *Audiovisual speech recognition*, MUSCLE Network of Excellence Workshop, Paris.

Professional Activities

Journal Paper Reviewing

IEEE Transactions on Pattern Analysis and Machine Intelligence (2006–now), IEEE Transactions on Image Processing (2005–now), International Journal of Computer Vision (2006), IEEE Transactions on Multimedia (2010–now), Pattern Recognition (2008–now), Signal Processing (2010–now), Image and Vision Computing (2011–now), IEEE Signal Processing Letters (2009–now), IEEE Transactions on Systems, Man, and Cybernetics – Part B (2011–now), SIAM Journal on Scientific Computing (2010–now), EURASIP Journal on Advances in Signal Processing (2010–now), EURASIP Journal on Image and Video Processing (2009–now).

Conference Technical Program Committee Member

Int. Conf. on Comp. Vision (ICCV) (2013), IEEE Int. Conf. on Comp. Vision and Pat. Rec. (CVPR) (2012, 2013, 2014), Int. Conf. on Neural Information Processing Systems (NIPS) (2012, 2013), European Conf. on Comp. Vision (ECCV) (2012, 2014), Int. Conf. on Machine Learning (ICML) (2012, 2013), Int. Conf. on Pat. Rec. (ICPR) (2010, 2012, 2014), ICCV Workshop on Graphical Models for Scene Understanding (2013), IEEE CVPR Workshop on Structured Prediction (2013), IEEE CVPR Workshop on Perceptual Organization in Computer Vision (2012), ECCV Workshop on Higher-Order Models and Global Constraints in Computer Vision (2012), ACM Conf. on Multimedia (2010), Eurographics (2010), NIPS Workshop on Deep Learning and Unsupervised Feature Learning (2010).

Conference Organization

NIPS Workshop on Perturbations, Optimization, and Statistics (2013) [co-organized with Tamir Hazan, Alexander Rakhlin, and Daniel Tarlow], NIPS Workshop on Perturbations, Optimization, and Statistics (2012) [co-organized with Tamir Hazan and Daniel Tarlow].

Tutorial Organization

Full day tutorial on “Bases for Images and Surfaces” in conjunction with CVPR-2014 [co-organized with Iasonas Kokkinos, Alex Bronstein, Michael Bronstein].

Professional Memberships

Senior Member of the Institute of Electrical and Electronics Engineers (IEEE) (2003–now). Member of Association for Computing Machinery (ACM) (2003–now), Society for Industrial and Applied Mathematics (SIAM) (2008–now), Technical Chamber of Greece (2004–now).

Open-source Research Software

- 2008 **GAC++**, *A C++ toolbox for geometric active contours and other related PDE-based computer vision models*, (GPL license).
- 2008 **AAMtools**, *A MATLAB toolbox for building active appearance models and fitting them to still and moving images*, (GPL license).

Computer Skills

Programming, Proficient in C/C++, Matlab, CUDA GPU programming. Working knowledge of OpenGL, Java.

Other, System administration (Linux, Windows), document processing in LaTeX.

Languages

English (Fluent), **German** (Good), **Spanish** (Basic), **Greek** (Native)

Scientific Publications

Refereed Journal Articles

- [1] G. Papandreou, A. Katsamanis, V. Pitsikalis, and P. Maragos. Adaptive multimodal fusion by uncertainty compensation with application to audiovisual speech recognition. *IEEE Transactions on Audio, Speech and Language Processing*, 17(3):423–435, March 2009.
- [2] A. Katsamanis, G. Papandreou, and P. Maragos. Face active appearance modeling and speech acoustic information to recover articulation. *IEEE Transactions on Audio, Speech and Language Processing*, 17(3):411–422, March 2009.
- [3] S. Lefkimmiatis, P. Maragos, and G. Papandreou. Bayesian inference on multiscale models for Poisson intensity estimation: Applications to photon-limited image denoising. *IEEE Transactions on Image Processing*, 18(8):1724–1741, August 2009.
- [4] G. Papandreou and P. Maragos. Multigrid geometric active contour models. *IEEE Transactions on Image Processing*, 16(1):229–240, January 2007.

Refereed Conference Proceedings

Note: ICCV, CVPR, and NIPS are top conferences in computer vision and pattern recognition. Only around 25% of the submissions get accepted (less than 4% as oral presentations).

- [1] G. Papandreou. Deep Epitomic Convolutional Neural Networks. In *arXiv:1406.2732*.
- [2] G. Papandreou, L. Chen, and A. Yuille. Modeling the Appearance of Image Patches with a Generic Dictionary of Mini-Epitomes. In *Proc. IEEE Int. Conf. on Comp. Vision and Pat. Rec. (CVPR)*, Columbus, OH, June 2014.
- [3] L. Chen, G. Papandreou, and A. Yuille. Learning a Dictionary of Shape Epitomes with Application to Semantic Labeling. In *Proc. IEEE Int. Conf. on Computer Vision (ICCV)*, Sydney, Australia, Dec. 2013.
- [4] G. Papandreou and A. Yuille. Perturb-and-MAP random fields: Using discrete optimization to learn and sample from energy models. In *Proc. IEEE Int. Conf. on Computer Vision (ICCV)*, Barcelona, Spain, Nov. 2011, pages 193–200. **Oral Presentation.**
- [5] G. Papandreou and A. Yuille. Efficient Variational Inference in Large-Scale Bayesian Compressed Sensing. In *Proc. IEEE Workshop on Information Theory in Computer Vision and Pattern Recognition (in conjunction with ICCV-11)*, Barcelona, Spain, Nov. 2011, pages 1332–1339. **Oral Presentation.**
- [6] G. Papandreou and A. Yuille. Gaussian sampling by local perturbations. In *Proc. Int. Conf. on Neural Information Processing Systems (NIPS)*, Vancouver, B.C., Canada, Dec. 2010.

- [7] S. Lefkimmiatis, G. Papandreou, and P. Maragos. Poisson-Haar transform: A nonlinear multiscale representation for photon-limited image denoising. In *Proc. IEEE Int. Conf. on Image Processing (ICIP)*, Cairo, Egypt, Nov. 2009, pages 3853–3856.
- [8] S. Lefkimmiatis, G. Papandreou, and P. Maragos. Photon-limited image denoising by inference on multiscale models. In *Proc. IEEE Int. Conf. on Image Processing (ICIP)*, San Diego, CA, Oct. 2008, pages 2332–2335.
- [9] G. Papandreou and P. Maragos. Adaptive and constrained algorithms for inverse compositional active appearance model fitting. In *Proc. IEEE Int. Conf. on Comp. Vision and Pat. Rec. (CVPR)*, Anchorage, AK, June 2008.
- [10] A. Katsamanis, G. Ananthakrishnan, G. Papandreou, P. Maragos, and O. Engwall. Audiovisual speech inversion by switching dynamical modeling governed by a hidden Markov process. In *Proc. 16th European Signal Processing Conf. (EUSIPCO)*, Lausanne, Switzerland, Aug. 2008. **Oral Presentation.**
- [11] G. Papandreou, P. Maragos, and A. Kokaram. Image inpainting with a wavelet domain hidden Markov tree model. In *Proc. IEEE Int. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*, Las Vegas, NV, Apr. 2008, pages 773–776. **Oral Presentation.**
- [12] A. Katsamanis, G. Papandreou, and P. Maragos. Audiovisual-to-articulatory speech inversion using active appearance models for the face and hidden Markov models for the dynamics. In *Proc. IEEE Int. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*, Las Vegas, NV, Apr. 2008, pages 2237–2240.
- [13] G. Papandreou, A. Katsamanis, V. Pitsikalis, and P. Maragos. Multimodal fusion and learning with uncertain features applied to audiovisual speech recognition. In *Proc. IEEE Workshop on Multimedia Signal Processing (MMSp)*, Chania, Greece, Oct. 2007, pages 264–267. **Oral Presentation.**
- [14] A. Katsamanis, G. Papandreou, and P. Maragos. Audiovisual-to-articulatory inversion using hidden Markov models. In *Proc. IEEE Workshop on Multimedia Signal Processing (MMSp)*, Chania, Greece, Oct. 2007, pages 457–460.
- [15] V. Pitsikalis, A. Katsamanis, G. Papandreou, and P. Maragos. Adaptive multimodal fusion by uncertainty compensation. In *Proc. Int. Conf. on Spoken Language Processing (ICSLP)*, Pittsburgh, PA, Sep. 2006, pages 2458–2461. **Oral Presentation.**
- [16] A. Katsamanis, G. Papandreou, V. Pitsikalis, and P. Maragos. Multimodal fusion by adaptive compensation for feature uncertainty with application to audiovisual speech recognition. In *Proc. 14th European Signal Processing Conf. (EUSIPCO)*, Florence, Italy, Sept. 2006. **Oral Presentation.**
- [17] D. Dimitriadis, N. Katsamanis, P. Maragos, G. Papandreou, and V. Pitsikalis. Towards automatic speech recognition in adverse environments. In *Proc. Hell. Europ. Conf. on Comp. Math. and Appl. (HERCMA)*, Athens, Greece, Sept. 2005. **Oral Presentation.**
- [18] G. Papandreou and P. Maragos. A cross-validated statistical approach to scale selection for image denoising by nonlinear diffusion. In *Proc. IEEE Int. Conf. on Comp. Vision and Pat. Rec. (CVPR)*, San Diego, CA, June 2005, volume I, pages 625–630.
- [19] G. Papandreou and P. Maragos. Image denoising in nonlinear scale-spaces: Automatic scale selection via cross-validation. In *Proc. IEEE Int. Conf. on Image Processing (ICIP)*, Genova, Italy, Sept. 2005, volume I, pages 481–484.
- [20] G. Papandreou and P. Maragos. A fast multigrid implicit algorithm for the evolution of geodesic active contours. In *Proc. IEEE Int. Conf. on Comp. Vision and Pat. Rec. (CVPR)*, Washington DC, June 2004, volume II, pages 689–694.

Book Chapters

- [1] G. Papandreou and A. Yuille. Perturb-and-MAP Random Fields: Reducing Random Sampling to Optimization, with Applications in Computer Vision. In S. Nowozin, P. Gehler, J. Jancsary, and C. Lampert, editors, *Advanced Structured Prediction*. MIT Press, 2014 (in press).
- [2] G. Papandreou, A. Katsamanis, V. Pitsikalis, and P. Maragos. Adaptive multimodal fusion by uncertainty compensation with application to audiovisual speech recognition. In P. Maragos, A. Potamianos, and

P. Gros, editors, *Multimodal Processing and Interaction: Audio, Video, Text*, chapter 4, pages 111–126. Springer-Verlag, New York, 2008.

- [3] P. Maragos, P. Gros, A. Katsamanis, and G. Papandreou. Cross-modal integration for performance improving in multimedia: A review. In P. Maragos, A. Potamianos, and P. Gros, editors, *Multimodal Processing and Interaction: Audio, Video, Text*, chapter 1, pages 3–48. Springer-Verlag, New York, 2008.