

Yannis Kalantidis

Education

- 2009–2014 **Ph.D. in Computer Science**, *School of Electrical and Computer Engineering*, National Technical University of Athens, Greece.
Ph.D. Thesis: Clustering and nearest neighbor methods for visual search
- 2002–2009 **Diploma/M.Eng. in Electrical and Computer Engineering**, *School of Electrical and Computer Engineering*, National Technical University of Athens, Athens, Greece.
Diploma Thesis: Content- and geometry-based image retrieval

Working experience

- Feb 2017 – Present **Research Scientist**, *Facebook Research*, Menlo Park, CA,
Research and development on video representations, temporal segmentation, web-scale learning, visual relationship detection, multi-modal learning.
- Jan 2015 – Dec 2016 **Postdoctoral Researcher & Research Scientist (2016)**, *Yahoo Research*, San Francisco, CA,
(2 years) Research and development on web-scale visual search, distributed backend for approximate nearest neighbor search, image & video representation, mutlimodal classification, native ad selection, personal media search. Collaborated with Stanford on the Visual Genome project, <https://visualgenome.org/>, [IJCV, 2017].
- 2009–2014 **Image Video and Multimedia Laboratory, School of Electrical and Computer Engineering, National Technical University of Athens**, *Research and Development in FP7 European Projects*,
(5 years) Notably:
2009: Imagination (under ICCS/NTUA)
2009-2011: WEKNOWIT (under CERTH/ITI) .

Scientific Publication, Patent & Citation Records

- Publications Papers published in CVPR, ICCV, ECCV, IJCV, CVIU, ACM MM, ICMR, WSDM, CHI and other.
- Citations Google Scholar citations: **1118 citations, h-index: 13, i10-index: 15** (July 2018)
- Patents 2 US Patents, 4 US Patent Applications & 2 Defensive Publications.

Teaching experience

- Fall 2018 **Mini-Course on Computer Vision**, *UCSF Department of Radiation Oncology*, University of California, San Francisco, Invited by Prof Gilmer Valdes. In planning stage, details and dates to be determined.
- Fall 2016 **Guest lecturer at CS131 - Computer Vision: Foundations and Applications**, *Computer Science Department*, Stanford University, Course taught by Prof. Fei-Fei Li and Juan Carlos Niebles.
I was invited to give three lectures on basic computer vision principles (motion, gestalt and segmentation).
- 2009–2014 **Image and Video Analysis, Lab assistant/lecturer**, *School of Electrical and Computer Engineering*, National Technical University of Athens, Greece.
- 2010–2014 **Programming Techniques, Teaching Assistant**, *School of Electrical and Computer Engineering*, National Technical University of Athens, Greece.

2009–2010 **Introduction to Programming, Teaching Assistant**, *School of Electrical and Computer Engineering*, National Technical University of Athens, Greece.

Internships

May–July 2010 **Research Intern at Yahoo Research Barcelona**, *supervisor: Roelof Van Zwol, collaborators: Lluís Garcia Pueyo, Michele Trevisiol*, Barcelona, Spain.

May–August 2012 **Research Intern at Yahoo Research Santa Clara**, *supervisor: Lyndon Kennedy, collaborators: Li-Jia Li, David A. Shamma*, San Francisco & Santa Clara, USA.

May–August 2014 **Research Intern at Yahoo Labs San Francisco**, *supervisor: David A. Shamma, collaborators: Lyndon Kennedy*, San Francisco, USA.

Awards & Scholarships

2017 **Outstanding Reviewer Award**, *International Conference on Computer Vision*, (ICCV).

2016 **Best Paper Award**, *NIPS Large-Scale Computer Vision Systems Workshop*.

2015 **Outstanding Reviewer Award**, *International Conference on Computer Vision*, (ICCV).

2013 **Thomaidion Award for Scientific Publication**, *Getting the Look: Clothing Recognition and Segmentation for Automatic Clothing Suggestions in Everyday Photos*, National Technical University of Athens.

2012 **Thomaidion Award for Scientific Publication**, *VIRaL: Visual Image Retrieval and Localization*, National Technical University of Athens.

2011 **Thomaidion Award for Scientific Publication**, *Retrieving landmark and non-landmark images from community photo collections*, National Technical University of Athens.

2005 **Erasmus Scholarship**, *Exchange Studies at Lund Univesity*, Lund, Sweden.

Development projects

VIRaL Principal developer in the Visual Image Retrieval and Localization tool (On-line demo: <http://viral.image.ntua.gr>).

LOPQ Principal developer in the open source code for Locally Optimized Product Quantization (Available on Yahoo's github: <https://github.com/yahoo/lopq>).

Crow Principal developer in the open source code for Cross-dimensional feature Weighting for convolutional features (Available on Yahoo's github: <https://github.com/yahoo/crow>).

Patents & Defensive Publications

US Patent **Visual Clothing Retrieval**, *Lyndon Kennedy, Li-Jia Li, Ioannis Kalantidis*, US Patent 9,460,518, Pub. Date: 4 October 2016.

US Patent **Locally optimized feature space encoding of digital data and retrieval using such encoding**, *Lyndon Kennedy, Ioannis Kalantidis, David Shamma*, US Patent 9,659,214, Pub. Date: 2017/5/23.

US Patent Application **Method and system for selecting supplemental content using visual appearance**, *Ioannis Kalantidis, Lyndon Kennedy, Ayman Farahat, Ricardo Baeza-Yates, Simon Osindero and David A. Shamma*, US Patent App. 15/097,658, 2017/10/19.

US Patent Application **Content sharing recommendations**, *Bart Thomee, Andrew Stadlen, Austin Shoemaker, Ioannis Kalantidis, Dan Tasse and David A. Shamma*, US Patent App. 15/134,633, 2017/10/26.

US Patent Application **Computerized system and method for modifying a media file by automatically applying security features to select portions of media file content**, *Bart Thomee, Ioannis Kalantidis, Lyndon Kennedy*, US Patent App. 15/190,315, 2017/12/28.

- US Patent Application **Selective content insertion into areas of media objects**, *Bart ThomÁlle, Ioannis Kalantidis, Clayton Ellis Mellina, Simon Osindero*, US Patent App. 15/204,450, 2018/1/11.
- Defensive Publication **Scalable Triangulation-based Logo Recognition**, *Yannis Kalantidis, Michele Trevisiol, Lluís Garcia Pueyo and Roelof van Zwol*, Defensive Publication, Docket No.: ID11-7243, Pub. Date: April 29 2011.
- Defensive Publication **Cross-dimensional Weighting for Aggregated Deep Convolutional Features**, *Yannis Kalantidis, Clayton Mellina, Simon Osindero, David A. Shamma and Gerry Pesavento*, Defensive Publication, Docket No.: ID15-11464, Pub. Date: TBD.

Invited Talks

- University of San Francisco**, *Pushing the Frontiers of Video Understanding*, Seminar Series in Data Science, April 20th 2018, San Francisco, CA.
- Universita Politecnica Catalunya**, *Computer Vision at Facebook research*, Deep Learning Open Lecture, November 20th 2017, Barcelona, Spain.
- Columbia University**, Digital Video and Multimedia Laboratory, invited by Prof. Shih-Fu Chang, August 19th 2014, New York, NY.
- Stanford University**, Image, Video and Multimedia Systems Group, invited by Prof. Bernd Girod, August 1st 2014, Stanford, CA.
- FXPAL**, *Nearest Neighbor Search and Clustering for Large Scale Visual Search*, Fuji-Xerox Palo Alto Laboratory, July 29th 2014, Palo Alto, CA.
- Stanford University**, Stanford Vision Lab, invited by Prof. Fei-Fei Li, July 28th 2014, Stanford, CA.

Reference/Recommendation Letters

The following scientists would be willing to provide recommendation letters

Please contact them in the following order.

Yannis Avrithis, *Senior Researcher, Inria Rennes*, (PhD Supervisor).

Simon Osindero, *Senior Researcher, Deepmind*, (Co-author).

David A. Shamma, *Senior Researcher, FXPAL*, (Postdoc Supervisor).

Ricardo Baeza-Yates, *Professor, Universitat Pompeu Fabra, Universidad de Chile & CTO at NTENT*, (Director at Yahoo Research).

Stefanos Kollias, *Professor, National Technical University of Athens*, (PhD Supervisor).

Petros Maragos, *Professor, National Technical University of Athens*, (PhD Committee).

Ioannis Emiris, *Professor, University of Athens*, (Co-author).

Giorgos Toliás, *Researcher, CMP, University of Prague*, (Co-author).

Lyndon Kennedy, *Senior Researcher, FXPAL*, (Internship Supervisor).

Lluís Garcia-Pueyo, *Senior Software Engineer, Facebook*, (Co-author).

Li-Jia Li, *Director of Research, Google AI*, (Co-author).

Roelof Van Zwol, *Director of Innovation, Netflix*, (Internship Supervisor).

Research Interests

Efficient spatio-temporal representations [MF-Net, ECCV 2018 Submission]

Learning Temporal Segmentation [SD-Net, ECCV 2018 Submission]

Large-scale Visual Relationship Understanding [Arxiv 2018]

Approximate nearest neighbor search [LOPQ, CVPR 2014]

Web-scale clustering [AGM, ECCV 2012] [IQ-Means, ICCV 2015]
Aggregated Deep Convolutional Features [Crow, ArXiv 2016]
Deep embeddings for Personal Media Search [WSDM 2017]
Distributed hashing and visual recommendations [LOH, ArXiv 2016]
Visually congruent selection and placement of native advertisements [ICPR 2016]
Geometry indexing [FMH, ACM Multimedia 2010] [FMS, CVIU 2014]
Clothing segmentation, recognition and classification [CsCR, ICMR 2013]
Geo and multiple-view clustering, location recognition [Scene Maps, ACM Multimedia 2010, MTAP 2011]
One-shot visual feature selection [SymCity, ACM Multimedia 2012]
Logo recognition [MSDT, ICMR 2011]
Question answering for personal media search [MemexQA, Arxiv 2017]
Multimodal Classification of Moderated Online Pro-Eating Disorder Content [CHI 2017]

Programming skills

Programming Python, C/C++, Matlab, PIG.
Familiar with Java, HIVE.
Other Caffe, Tensorflow, Hadoop, Spark, OpenCV, \LaTeX , Pgfplots

Professional Memberships

2009–today **Institute of Electrical and Electronics Engineers**, [IEEE].
2009–today **Technical Chamber of Greece**, [GTC].

Reviewer in journals

2010–today **Multimedia Tools and Applications**, [MTAP], Springer.
2012–today **Signal Processing: Image Communication**, [SPIC], Elsevier.
2014–today **Transactions on Circuits and Systems for Video Technology**, [CSVT], IEEE.
2014–today **Transactions on Multimedia**, [TMM], IEEE.
2015–today **Transactions on Image Processing**, [TIP], IEEE.
2015–today **Transactions on Intelligent Systems and Technology**, [TIST], ACM.
2015–today **Signal Processing Letters**, [SPL], IEEE.
2015–today **Transactions on Knowledge and Data Engineering**, [TKDE], IEEE.
2015–today **The VLDB Journal**, [VLDBJ], Springer.
2016–today **Transactions on Pattern Analysis and Machine Intelligence**, [PAMI], IEEE.
2016–today **Computer Vision and Image Understanding**, [CVIU], Elsevier.

Committee member in conferences & workshops

2016–today **IEEE Conference on Computer Vision and Pattern Recognition**, *Technical Program Committee Member*, [CVPR].
2016–today **IEEE European Conference on Computer Vision**, *Technical Program Committee Member*, [ECCV].
2015–today **IEEE International Conference on Computer Vision**, *Technical Program Committee Member*, [ICCV], (**Outstanding Reviewer Award** in 2015 and 2017).
2013–today **ACM Multimedia Conference**, *Technical Program Committee Member*, [ACM-MM].

- 2011 **IEEE Workshop on Content-based Multimedia Indexing**, *Technical Program Committee Member*, [CBMI].
- 2009 **IEEE Workshop on Content-based Multimedia Indexing**, *Local Chair Member*, [CBMI].

Summer schools attended

- GSS 2013 **Graduate Summer School: Computer Vision**, *UCLA, Los Angeles, California*, 22 July - 9 August 2013, Full UCLA Scholarship.
- SAGA 2011 **SAGA (ShApes Geometry and Algebra) Fall School**, *Vilnius, Lithuania*, 27-30 September 2011, Full SAGA Scholarship.
- VS3 2010 **Vision and Sports Summer School 2010**, *ETH, Zurich, Switzerland*, 16-20 August 2010.
- ESSIR 2009 **7th European Summer School in Information Retrieval**, *Padova, Italy*, 31st August - 4 September 2009.

Languages

Greek Native
English Fluent
German Basic

Certificate of Proficiency in English
Zertificat Deutsch als Fremdsprache, Test DaF (grade 3)

List of Publications

Selected publications

Y. Chen, **Y. Kalantidis**, J. Li, Y. Shuicheng, J. Feng. Double Attention Networks. In *Neural Information Processing (NIPS) (Submission under review)*, 2018.

Y. Chen, **Y. Kalantidis**, J. Li, Y. Shuicheng, J. Feng. Multi-Fiber Networks. In *European Conference on Computer Vision (ECCV)*, 2018.

J. Zhang, **Y. Kalantidis**, M. Rohrbach, M. Paluri A. Elgammal, M. Elhoseiny. Large-Scale Visual Relationship Understanding. In *arxiv:1804.10660*, 2018.

R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, **Y. Kalantidis**, L.-J. Li, D.A. Shamma, M. Bernstein and L. Fei-Fei. Visual Genome: Connecting Language and Vision Using Crowdsourced Dense Image Annotations. In *International Journal of Computer Vision (IJCV)*, 2017.

S. Chancellor, **Y. Kalantidis**, J.A. Pater, M. De Choudhury, D.A. Shamma Multimodal Classification of Moderated Online Pro-Eating Disorder Content In *Conference on Human Factors in Computing Systems (CHI)*, 2017.

L. Jiang, **Y. Kalantidis**, L. Cao, S. Farfadi, J. Tang and A. Hauptmann. Delving Deep into Personal Photo and Video Search. In *Web Search and Data Mining (WSDM)*, 2017.

L. Jiang, L. Cao, **Y. Kalantidis**, S. Farfadi and A. Hauptmann. MemexQA: Visual Memex Question Answering. In *arXiv:1708.01336*, 2017.

P. Garrigues, S. Farfadi, H. Izadinia, K. Boakye and **Y. Kalantidis**. Tag Prediction at Flickr: a View from the Darkroom. In *Large Scale Computer Vision Systems Workshop, NIPS, 2016, Best Paper Award Winner*.

Y. Kalantidis, C. Mellina and S. Osindero. Cross-dimensional Weighting for Aggregated Deep Convolutional Features. In *Web-scale Vision and Social Media (VSM) Workshop, ECCV, 2016*.

Y. Kalantidis, A. Farahat, L. Kennedy, R. Baeza-Yates and D.A. Shamma. Visual Congruent Ads for Image Search. In *International Conference on Pattern Recognition (ICPR)*, 2016.

Y. Kalantidis, L. Kennedy, H. Nguyen, C. Mellina and D.A. Shamma. LOH and behold: Web-scale visual search, recommendation and clustering using Locally Optimized Hashing. In *Web-scale Vision and Social Media (VSM) Workshop, ECCV, 2016*.

Y. Avrithis, **Y. Kalantidis**, E. Anagnostopoulos and I. Z. Emiris. Web-scale image clustering revisited. In *International Conference on Computer Vision (oral) (ICCV)*, Santiago, Chile, December 2015.

Y. Kalantidis and Y. Avrithis. Locally Optimized Product Quantization for Approximate Nearest Neighbor Search. In *Computer Vision and Pattern Recognition (CVPR)*, Columbus, OH, June 2014.

G. Toliás, **Y. Kalantidis**, and Y. Avrithis. Towards large-scale geometry indexing by feature selection. *Computer Vision and Image Understanding (CVIU)*, March 2014.

Y. Kalantidis, L. Kennedy and L.-J. Li. Getting the Look: Clothing Recognition and Segmentation for Automatic Clothing Suggestions in Everyday Photos. In *International Conference on Multimedia Retrieval (Oral paper) (ICMR)*, Dallas, TX, April 2013.

Y. Avrithis and **Y. Kalantidis**. Approximate gaussian mixtures for large scale vocabularies. In *European Conference on Computer Vision (ECCV)*, Florence, Italy, October 2012.

G. Toliás, **Y. Kalantidis**, and Y. Avrithis. Symcity: Feature selection by symmetry for large scale image retrieval. In *ACM Multimedia (Oral paper) (MM 2012)*, Nara, Japan, October 2012. ACM.

Y. Avrithis, **Y. Kalantidis**, G. Toliás, and E. Spyrou. Retrieving landmark and non-landmark images from community photo collections. In *ACM Multimedia (Oral paper) (ACM MM 2010)*, Firenze, Italy, October 2010.

Y. Avrithis, G. Toliás, and **Y. Kalantidis**. Feature map hashing: Sub-linear indexing of appearance and global geometry. In *ACM Multimedia (Oral paper) (ACM MM 2010)*, Firenze, Italy, October 2010.

Y. Kalantidis, L.G. Pueyo, M. Trevisiol, R. van Zwol, and Y. Avrithis. Scalable triangulation-based logo recognition. In *International Conference on Multimedia Retrieval*, Trento, Italy, April 2011.

Y. Kalantidis, G. Toliás, Y. Avrithis, M. Phinikettos, E. Spyrou, P. Mylonas, and S. Kollias. Viral: Visual image retrieval and localization. *Multimedia Tools and Applications (MTAP)*, 2011.

Other publications

L Jiang, LL Cao, Y Kalantidis, S Farfade, AG Hauptmann. Visual Memory QA: Your Personal Photo and Video Search Agent In *AAAI*, demo presentation, 2017.

Y Kalantidis, C Xu, C Mellina Do Silhouettes Dream? *Creativity & Cognition*, Art Track, 2017.

I. Kollia, Y. Kalantidis, K. Rapantzikos, and A. Stafylopatis. Improving semantic search in digital libraries using multimedia analysis. In *Journal of Multimedia*, vol. 7, no 2, pp. 193-204, 2012.

I. Kollia, Y. Kalantidis, K. Rapantzikos, and S. Kollias. Visual analysis for semantic search in digital libraries. In *Constantinides International Workshop on Signal Processing (CIWSP 2013)*, London, 2013.

E. Spyrou, Y. Kalantidis, G. Toliás, Ph. Mylonas, and S. Kollias. Intelligent content retrieval using a visual vocabulary and geometric constraints. In *FUZZ-IEEE 2010, in conjunction with The 2010 IEEE World Congress on Computational Intelligence (WCCI 2010)*, July 2010.

S. Papadopoulos, C. Zigkolis, G. Toliás, Y. Kalantidis, P. Mylonas, Y. Kompatsiaris, and A. Vakali. Image clustering through community detection on hybrid image similarity graphs. In *International Conference on Image Processing (ICIP 2010)*, Hong Kong, September 2010.

E. Spyrou, Y. Kalantidis, and Ph. Mylonas. Exploiting a region-based visual vocabulary towards efficient concept retrieval. In *in Proceedings of Recognising and tracking events on the Web and in real life, in conjunction with SETN 2010 (EVENTS 2010)*, May 2010.

Y. Kalantidis, E. Spyrou, Ph. Mylonas, and S. Kollias. Using a region and visual word approach towards semantic image retrieval. In *Workshop on Semantic Media Adaptation and Personalization (SMAP 2010)*, Limassol, Cyprus, December 2010.

Y. Kalantidis, G. Toliás, E. Spyrou, Ph. Mylonas, Y. Avrithis, and S. Kollias. Visual image retrieval and localization. In *Greek Electrical and Computer Engineering Student Conference*, Thessaloniki, Greece, April 2009.

Y. Kalantidis, G. Toliás, E. Spyrou, Ph. Mylonas, and Y. Avrithis. Visual image retrieval and localization. In *International Workshop on Content-Based Multimedia Indexing (CBMI 2009)*, Chania, Greece, June 2009.

Y. Kalantidis. Content- and geometry-based image search. *Diploma Thesis (National Technical University of Athens, Athens, Greece, February 2009.*