

Yannis Kalantidis

Education

- 2009–2014 **Ph.D. in Computer Science**, *School of Electrical and Computer Engineering*, National Technical University of Athens, Greece.
2002–2009 **Diploma/M.Eng. in Electrical and Computer Engineering**, National Technical University of Athens, Athens, Greece.

Working Experience

- Feb 2017 – Present **Research Scientist**, *Facebook AI*, Menlo Park, CA,
Research and development on video understanding, large-scale learning, and vision & language..
Jan 2015 – Dec 2016 **Research Scientist**, *Yahoo Research*, San Francisco, CA,
Research and development on web-scale visual search, classification from noisy data, video representation, adversarial & unsupervised learning. Collaborated with Stanford on the Visual Genome project [IJCV, 2016] (<https://visualgenome.org/>).

Scientific Publication, Patent & Citation Records

- Publications Papers published in CVPR, NeurIPS, ICCV, ECCV, TPAMI, IJCV, AAAI, CVIU, WSDM, CHI and other.
Citations Google Scholar citations: **1777 citations**, **h-index: 15**, **i10-index: 20** (July 2019)
Patents 2 US Patents, 4 US Patent Applications & 2 Defensive Publications.

Selected Recent Publications

- Y. Chen, H. Fan, B. Xu, Z. Yan, **Y. Kalantidis**, M. Rohrbach, Y. Shuicheng, J. Feng. *Drop an Octave: Reducing Spatial Redundancy in Convolutional Neural Networks with Octave Convolution*. **ICCV 2019**.
Y. Chen, M. Rohrbach, Z. Yan, Y. Shuicheng, J. Feng, **Y. Kalantidis**. *Graph-Based Global Reasoning Networks*. **CVPR, 2019**.
L. Zhou, **Y. Kalantidis**, X. Chen, J. Corso, M. Rohrbach. *Grounded Video Description*. **CVPR (oral), 2019**.
B. Xiong, **Y. Kalantidis**, D. Ghadiyaram, K. Grauman. *Less is More: Learning Highlight Detection from Video Duration*. **CVPR 2019**.
Z. Shou, Z. Yan, **Y. Kalantidis**, L. Sevilla-Lara, M. Rohrbach, X. Lin, S.-F. Chang. *DMC-Net: Generating Discriminative Motion Cues for Fast Compressed Video Action Recognition*. **CVPR 2019**.
J. Liang, L. Jiang, L. Cao, **Y. Kalantidis**, L.J. Li and A. Hauptmann. *Focal Visual-Text Attention for Memex Question Answering*. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, **2019**.
J. Zhang, **Y. Kalantidis**, M. Rohrbach, M. Paluri, M. Elhoseiny. *Large-Scale Visual Relationship Understanding*. **AAAI, 2019**.
Y. Chen, **Y. Kalantidis**, J. Li, Y. Shuicheng, J. Feng. *A²-Nets: Double Attention Networks*. **NeurIPS, 2018**.
Y. Chen, **Y. Kalantidis**, J. Li, Y. Shuicheng, J. Feng. *Multi-Fiber Networks for Video Recognition*. **ECCV, 2018**.
R. Krishna, ..., **Y. Kalantidis et al.** *Visual Genome: Connecting Language and Vision Using Crowdsourced Annotations*. **IJCV 2017**.
L. Jiang, **Y. Kalantidis**, L. Cao, J. Tang and A. Hauptmann. *Delving Deep into Personal Photo & Video Search*. **WSDM, 2017**.
S. Chancellor, **Y. Kalantidis**, J.A. Pater, M. De Choudhury, D.A. Shamma *Multimodal Classification of Moderated Online Pro-Eating Disorder Content*, **CHI, 2017**.
P. Garrigues, S. Farfadi, H. Izadnia, K. Boakye and **Y. Kalantidis**. *Tag Prediction at Flickr: a View from the Darkroom*. **LSCVS NeurIPS Workshop (best paper award), 2016**.
Y. Kalantidis, C. Mellina and S. Osindero. *Cross-dimensional Weighting for Aggregated Deep Convolutional Features*. **ECCVw, 2016**.
Y. Avrithis, **Y. Kalantidis**, E. Anagnostopoulos and I. Z. Emiris. *Web-scale image clustering revisited*. **ICCV (oral), 2015**.
Y. Kalantidis and Y. Avrithis. *Locally Optimized Product Quantization for Approximate Nearest Neighbor Search*. **CVPR, 2014**.

Research Interests

- Image and Video representation learning [ECCV 2016, ECCV 2018, NeurIPS 2018, CVPR 2019, ICCV 2019]
Personal media search and multimodal question answering [WSDM 2017, TPAMI 2019]
Vision and language [IJCV 2017, AAAI 2019, CVPR 2019]
Video description [CVPR 2019] and summarization [CVPR 2019]
Web-scale clustering [ECCV 2012, ICCV 2015] and learning [NeurIPSw 2016]
Computer Vision for global challenges and social impact [CVPR 2019 Workshop – Lead Organizer]

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>).

Programming skills

- Programming Python, C/C++, familiar with Spark, Caffe2, Tensorflow, pyTorch.