

Yannis Kalantidis

12 Rue Auguste Ravier
38100, Grenoble, France
✉ ykalant@image.ntua.gr
📄 www.skamalas.com

Working Experience

- Mar 2020 – present **Senior Research Scientist**, *NAVER LABS Europe*, Grenoble, France, Research and development on resource-constrained representation learning, vision & language, multi-modal learning.
- Feb 2017 – Sept 2019 **Research Scientist**, *Facebook AI*, Menlo Park, CA, Research and development on representation learning, video understanding and vision & language.
- Jan 2015 – Dec 2016 **Research Scientist**, *Yahoo Research*, San Francisco, CA, Research and development on web-scale visual search and clustering, classification from noisy data, deep neural network representations.

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.

Scientific Publication, Patent & Citation Records

- Citations Google Scholar citations (profile link): **6006 citations**, **h-index: 24**, **i10-index: 31**. Since 2017: **5394 citations** (Jan 2022)
- Patents 7 US Patents, 6 US Patent Applications.

Selected Recent Publications

- P. Weinzaepfel, T. Lucas, D. Larlus, **Y. Kalantidis**. *Learning Super-Features for Image Retrieval*. **ICLR, 2022**.
- Y. Kalantidis**, C. Lassance, J. Almazan, D. Larlus. *TLDR: Twin Learning for Dimensionality Reduction*. **arXiv, 2021**.
- M.B. Sariyildiz, **Y. Kalantidis**, D. Larlus, K. Alahari. *Concept Generalization in Visual Representation Learning*. **ICCV, 2021**.
- S. Chun, S.J. Oh, R.S. de Rezende, **Y. Kalantidis**, D. Larlus. *Probabilistic Embeddings for Cross-Modal Retrieval*. **CVPR, 2021**.
- Y. Kalantidis**, M.B. Sariyildiz, N. Pion, P. Weinzaepfel, D. Larlus. *Hard Negative Mixing for Contrastive Learning*. **NeurIPS, 2020**.
- B. Kang, S. Xie, M. Rohrbach, Z. Yan, A. Gordo, J. Feng, **Y. Kalantidis**. *Decoupling Representation and Classifier for Long-Tailed Recognition*. **ICLR, 2020**.
- C.Y. Ma, **Y. Kalantidis**, G. AlRegib, P. Vajda, M. Rohrbach, Z. Kira. *Learning to Generate Grounded Image Captions without Localization Supervision*. **ECCV 2020**.
- 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**, K. Grauman. *Less is More: Learning Highlight Detection from Video Duration*. **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**.
- 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**.
- P. Garrigues, S. Farfade, H. Izadinia, K. Boakye and **Y. Kalantidis**. *Tag Prediction at Flickr: a View from the Darkroom*. **NeurIPS LSCVS Workshop (best paper award), 2016**.
- Y. Kalantidis**, C. Mellina and S. Osindero. *Cross-dimensional Weighting for Aggregated Deep Convolutional Features*. **ECCVw, 2016**.
- Y. Kalantidis** and Y. Avrithis. *Locally Optimized Product Quantization for Approximate Nearest Neighbor Search*. **CVPR, 2014**.

Research Interests

- Summary **Visual representation learning in constrained environments // learning adaptive multi-modal systems**. Representation learning with limited annotations, compute or data; learning reusable and adaptive models for solving multiple tasks; learning transferable representations; modeling and understanding high-dimensional spaces, transfer learning and generalization.
- Self-supervised learning and generalization**, [NeurIPS 2020, ICCV 2021, Arxiv 2021].
- Deep architectures for images and videos**, [ECCV 2016, ECCV 2018, NeurIPS 2018, CVPR 2019a, ICCV 2019, ICLR 2020, 3DV 2021, ICLR 2022].
- Vision and language, multi-modal search**, [IJCV 2017, WSDM 2017, AAAI 2019, CVPR 2019b, TPAMI 2019, CVPR 2021].
- Video understanding**, [CVPR 2019b, CVPR 2019c, CVPR 2019d, ECCV 2020].
- Web-scale nearest neighbor search, clustering and classification**, [ECCV 2012, CVPR 2014, ICCV 2015, NeurIPSw 2016].

Open-source Development Projects

- Yahoo Research Principal developer of the Locally Optimized Product Quantization (yahoo/lopq) and CroW (yahoo/crow) github repos.
- Facebook AI Contributor to a number of open-source Facebook AI github repos, including: facebookresearch/GloRe, facebookresearch/OctConv.
- NAVER LABS Principal developer of the TLDR github repo (naver/tldr).