Timothée Cour
San Francisco, CA
Visa Status: US Green Card Holder
timothee.cour at gmail.com
http://www.timotheecour.com

Topics of interest

Machine Learning: deep learning, weakly supervised learning. **Computer Vision:** large scale object detection, recognition, segmentation, image retrieval. **Software Engineering:** Full stack.

Education

University of Pennsylvania Philadelphia, PA, USA
*Ph.D. in Computer and Information Science (**Advisor**: Ben Taskar)* *2009*
"Weakly supervised learning from multiple modalities: exploiting video, audio& text for video understanding"
Committee: Kostas Daniilidis, Fernando Pereira, Camillo J. Taylor, Andrew Zisserman

*M.S. in Computer and Information Science (**Advisor**: Jianbo Shi)* *2005*

Ecole Polytechnique (the **best rated** French engineering school) Paris, France
B.S. Applied Mathematics and Computer Science *2003*

Research Fellowship from University of Pennsylvania and from the Ecole Polytechnique Foundation
Carnot Foundation Fellowship, 2003 (**3 awardees** from Ecole Polytechnique each year)

Work Experience

Startup *may 2016-current*

Google Inc. Mountain View, CA, USA
Senior Software Engineer, Google Maps. *nov 2012-may 2016*
In charge of Street View Privacy (until Q1'15), a complex and critical infrastructure. Major redesign, including a universal license plate detector (for blurring) using deep nets and special hardware, running in production worldwide (**petabytes** of image data). This **enabled Street View to launch in APAC and 30 other** previously problematic countries, while reducing CPU cost and engineering time by $> 10x$.
Tech lead for Road Detection (since Q2'14). Initiated and drove a project (5 engineers) that automatically detects and synthesizes road networks from aerial imagery and other data sources using deep learning, **running worldwide**, helping us identify new and missing roads in Maps.
Designed and implemented internal tools for Geo Imagery: a generic visualizer to interactively query and visualize images and other data in *any* standard google format, and an interactive map with server side plugins for all geolocated data exported to Maps.

NEC Labs Cupertino, CA, USA
Research Scientist, Media Analytics department *2010-nov 2012*
Real time pedestrian, face, and hand detection, large scale image retrieval with state-of-the-art results on 4 public datasets, gesture recognition from webcam or kinect sensor, 3D reconstruction.
Member of the **winning team** at the *ImageNet Challenge*, 2010 (1K categories, 1M images).
Demos at the NEC Open House in Tokyo in 2010 (image retrieval) and 2011 (gesture based mouse control).

INRIA / Ecole Normale Supérieure - Willow project Paris, France
Postdoctoral researcher with Jean Ponce and Francis Bach. *2009-2010*
Learning from weakly annotated images, multiple instance boosting, video alignment using plot summaries.

GRASP Lab, University of Pennsylvania Philadelphia, PA, USA
Graduate Research Fellow with Ben Taskar and Jianbo Shi *2003-2009*
Weakly-supervised learning, object detection& segmentation, graph matching, approximate inference.

Microsoft Research Redmond, WA, USA

Research Internship with Paul Viola and Michael Shilman *summer 2005*
Recognition of handwritten mathematical expressions using efficient two-dimensional parsing.

Supervision and mentoring

James Liu (*Google*, 2016), YuHui Chen (*Google*, 2015), Kevin Lai (*University of Washington*, 2012), Olivier Duchenne (*ENS*, 2011), Jiangping Wang (*UIUC*, 2011), Ugo Jardonnet (*ENS*, 2010), Remi Cuingnet (*Ecole Polytechnique*, 2006), Pierre Fournier (*Ecole Polytechnique*, 2005), Florence Benezit (*Ecole Polytechnique*, 2004), Nicolas Gogin (*Ecole Polytechnique*, 2004).

Teaching

University of Pennsylvania: *Teaching Assistant: Computer Vision (CSE 399), Mathematical Foundations of Computer Science (CSE 260), Automata, Computability, and Complexity (CSE 262).* *2004 - 2006*
Ecole Polytechnique: *Tutoring in Maths and Physics for undergraduate students.* *2000-2003*

Invited talks

Various talks at Google (2013-2016).
Talking Pictures: Temporal Grouping and Dialog-Supervised Person Recognition. *Oxford Brookes/KTH/VGG/Willow* Workshop, 2010.

Weakly Supervised Learning for Video Understanding and Object Recognition. *NEC Labs*, Cupertino, 2010; *Intel Labs*, Seattle, 2010; *Google*, Mountain View, 2010, *MSR-INRIA* Workshop on Computer Vision and Machine Learning, 2010.

Talking Pictures: Temporal Grouping and Dialog-Supervised Person Recognition in Video. *CVPR* 2009 Workshop on Visual and Contextual Learning from Annotated Images and Videos.

Movie/Script: Alignment and Parsing of Video and Text Transcription. *Google*, Mountain View, 2008.

Object segmentation and recognition in image datasets and movies. *ENS*, Paris, 2007.

Video Deconstruction: Revealing narrative structure through image and text alignment. *NIPS* 2007 Workshop on the grammar of vision.

Publications and Patents

In Computer Vision and Machine Learning, conferences such as CVPR, ICCV, ECCV, NIPS, AISTATS are highly refereed. **Google scholar**: `http://scholar.google.com/citations?user=pkFzb9QAAAAJ`

S. Zhang, M. Yang, **T. Cour**, K. Yu, DN. Metaxas Query specific rank fusion for image retrieval. **PAMI** *2015.*

X. Wang, M. Yang, **T. Cour**, S. Zhu, K. Yu Contextual weighting and efficient re-ranking for vocabulary tree based image retrieval. **US Patent 8,892,542** *2014.*

S. Zhang, M. Yang, **T. Cour**, K. Yu Query Specific Fusion for Image Retrieval. **ECCV** *2012.*

X. Wang, M. Yang, **T. Cour**, S. Zhu, K. Yu, and T. X. Han. Contextual Weighting for Vocabulary Tree based Image Retrieval. **ICCV** *2011.*

T. Cour, B. Sapp, B. Taskar. Learning from Partial Labels. **JMLR** *2011.*

Y. Lin, F. Lv, S. Zhu, M. Yang, **T. Cour**, K. Yu, L. Cao, T. Huang. Large-scale image classification: fast feature extraction and SVM training. **CVPR** *2011.*

T. Cour, B. Sapp, A. Nagle, B. Taskar. Talking Pictures: Temporal Grouping and Dialog-Supervised Person Recognition. **CVPR** *2010.*

T. Cour Weakly Supervised Learning from Multiple Modalities: Exploiting Video, Audio and Text for Video Understanding. **PhD Thesis, University of Pennsylvania** *2009.*

T. Cour, B. Sapp, C. Jordan, B. Taskar. Learning from Ambiguously Labeled Images. **CVPR** *2009.*

T. Cour, C. Jordan, E. Miltsakaki, B. Taskar. Movie/Script: Alignment and Parsing of Video and Text Transcription. **ECCV** *2008.*

T. Cour, J. Shi. Recognizing objects by piecing together the Segmentation Puzzle. **CVPR** *2007.*

T. Cour, J. Shi. Solving Markov Random Fields with Spectral Relaxation. **AISTATS** *2007.*

T. Cour, P. Srinivasan, J. Shi. Balanced Graph Matching. **NIPS** *2006.*

T. Cour, F. Benezit, J. Shi. Spectral Segmentation with Multiscale Graph Decomposition. **CVPR** *2005.* **(500+citations)**

T. Cour, N. Gogin, J. Shi. Learning spectral graph segmentation. **AISTATS** *2005.*

Technical reports or Workshop papers:
T. Cour Convex Relaxations for Markov Random Field MAP estimation. **University of Pennsylvania, Philadelphia** *2008.*

T. Cour, B. Taskar. Video Deconstruction: Revealing narrative structure through image and text alignment. **Workshop on the Grammar of Vision, NIPS** *2007.*

Professional Activities

Member of IEEE (2006 - present).
Best Reviewer award, CVPR 2011.

Program committee / reviewer (conferences): ICCV 2007, 2009; ECCV 2010; CVPR 2009, 2010, 2011, 2012; NIPS 2009, 2010, 2011; ICML 2010, 2012; AAAI 2007; NESCAI 2007; IJCAI, 2009.

Reviewer (journals): PAMI 2007, 2008, 2008, 2009, 2010; JMLR 2010, 2011; CVIU 2007, 2008, 2010; IEEE Transactions on Image Processing 2007; IEEE Transactions on Information Theory 2007.

Open-source software

github: `https://github.com/timotheecour`
[old] Open source software in Matlab, C++, available at `http://www.timotheecour.com`.
Convex Learning from Partial Labels Toolbox: Implementation of the partially labeled multiclass classification algorithm introduced in *Learning from Partial Labels, JMLR 2011.* This also includes the *Annotated Faces on TV Dataset:* faces extracted, aligned and annotated from 8 TV shows.

Graph matching toolbox: State of the art and scalable subgraph matching algorithm. Uses spectral relaxation with affine constraints and bistochastic normalization. 5000+ page views since 2010.

Normalized Cuts Segmentation Code: Image segmentation and data clustering using Normalized Cuts. 15000+ page views since 2010.

Multiscale Normalized Cuts Image segmentation toolbox: Linear time image segmentation based on multiscale graph decomposition. Extensively tested and downloaded. 13000+ page views since 2010.

Technical skills

Programming languages: C++, C, D, Python, Javascript, Go, Matlab, Java, C#.

Libraries: Opencv, Opengl, TensorFlow, Openni, Libfreenect, phobos, stl, boost.

Tools: Docker, SWIG, LaTeX, gdb, llvm, hadoop, git.

Operating systems: Linux, Mac, Windows.

References

Available upon request.