Andrew Kae

CONTACT INFORMATION

Andrew Kae
New York, NY USA

E-mail: andrew.kae@gmail.com
Site: `http://www.andrewkae.com`

RESEARCH INTERESTS

Computer vision, machine learning, deep learning, graphical models.

I have worked on tasks such as domain adaptation, visual search, and semantic segmentation for images/videos. My ideal role is to use machine learning to create practical models for tasks in computer vision and other interesting domains.

EDUCATION

University of Massachusetts Amherst, Amherst, MA

Ph.D. in Computer Science, May 2014

- Advisor: Erik Learned-Miller
- GPA : 3.66

Cornell University, Ithaca, NY

M.Eng in Computer Science, December 2004

- GPA : 3.82

B.A. in Computer Science, May 2003

- GPA : 3.31

WORK EXPERIENCE

Oath, New York, NY

Research Scientist in Video Intelligence Group **August 2017 - Present**

My work focuses on building prototypes for video understanding tasks and publishing new research. In particular, I have worked on:

- Performing domain adaptation between images and videos to build a video classifier.
- Prototyping a new ad format using image synthesis techniques from computer vision.

Curalate, New York, NY

Research Engineer **November 2014 - August 2017**

My work involves deploying visual search applications by incorporating state-of-the-art work in machine learning and computer vision.

- Deployed a system for visual fashion search. Given an image we first perform object detection to find the clothing items in the image and find the closest matches (by visual similarity) from a product catalog.
- Built a prototype to perform A/B testing of different page layouts.

Yahoo! Labs, Santa Clara, CA

Intern in Display Advertising Group **Summer 2010**

I helped develop an algorithm for improved serving of relevant display relevant ads (KDD workshop 2011, Patent).

- Developed a system to classify display ads into a taxonomy by extracting OCR features from display ads and using them in conjunction with additional features.
- Used Hadoop to process millions of ads.

Lockheed Martin - Systems Integration, Owego, NY

Embedded Software Engineer in Simulations Group **Spring 2005 - Spring 2006**

I worked in the Simulations group which developed and supported software systems to simulate flight hardware.

Academic Research Experience

University of Massachusetts Amherst, Amherst, MA

Research Assistant in Computer Vision Lab **September 2007 - May 2014**

My work focuses on the semantic labeling of faces, which is the task of assigning category labels (such as hair or skin) to pixels in a face image. In particular, I have worked on:

- Incorporating a restricted Boltzmann machine (RBM) to model global label shape within a discriminative framework, for the semantic labeling of real-world face images (CVPR 2013).
- Extending previous work to model both temporal and shape dependencies using a conditional restricted Boltzmann machine (CRBM) for the semantic labeling of faces in YouTube videos (CVPR 2014).

In the past I have also worked on building character recognition systems. My work includes:

- Learning a clean-word list (a subset of words that are believed to be correctly recognized from an initial translation) from a noisy document with high confidence and using this list to improve character recognition within the document (CVPR 2010, JMLR 2012).
- Learning document-specific character models by exploiting language statistics (ICDAR 2009, IJDAR 2011).

Teaching Experience

University of Massachusetts Amherst, Amherst, MA

Teaching Assistant **Spring 2012 - December 2013**

I was a TA for the following courses: Introduction to Programming, Machine Learning, Graphical Models, and Software Engineering. For these courses, I taught discussion sections, held office hours, and graded homework assignments and exams.

Cornell University, Ithaca, NY

Teaching Assistant **Fall 2004**

I was a TA for the introductory Artificial Intelligence course. I graded homework assignments and exams, and wrote up one of the homework assignments.

Publications (Refereed)

Andrew Kae, Benjamin Marlin, and Erik Learned-Miller. *The Shape-Time Random Field for Semantic Video Labeling.* Computer Vision and Pattern Recognition (CVPR), 2014.

Andrew Kae*, Kihyuk Sohn*, Honglak Lee, and Erik Learned-Miller. *Augmenting CRFs with Boltzmann Machine Shape Priors for Image Labeling.* Computer Vision and Pattern Recognition (CVPR), 2013.
*The first and second authors made equal contributions and should be considered co-first authors.

Gary B. Huang, Andrew Kae, Carl Doersch, and Erik Learned-Miller. *Bounding the Probability of Error for High Precision Optical Character Recognition.* Journal of Machine Learning Research (JMLR), 2012.

Andrew Kae, Kin Kan, Vijay K Narayanan, Dragomir Yankov. *Categorization of Display Ads using Image and Landing Page Features.* The Third Workshop on Large-scale Data Mining: Theory and Applications'11 (LDMTA'11), in conjunction with SIGKDD2011.

Andrew Kae, David A. Smith, Erik Learned-Miller. *Learning on the Fly: A font-free approach towards multilingual OCR.* IJDAR, 2011.

Andrew Kae, Gary Huang, Erik Learned-Miller, Carl Doersch. *Improving State-of-the-Art OCR through High-Precision Document-Specific Modeling.* Computer Vision and Pattern Recognition (CVPR), 2010.

Andrew Kae, Erik Learned-Miller. 2009. *Learning on the Fly: Font free approaches to difficult OCR*

problems. International Conference on Document Analysis and Recognition (ICDAR), 2009.

Publications (Unrefereed)

Andrew Kae, Gary Huang, Erik Learned-Miller. 2009. *Bounding the Probability of Error for High Precision Recognition*. Technical Report UM-CS-2009-031, School of Computer Science, University of Massachusetts, Amherst, 2009.

Conference Presentations

(**Spotlight**) Andrew Kae*, Kihyuk Sohn*, Honglak Lee, and Erik Learned-Miller. *Augmenting CRFs with Boltzmann Machine Shape Priors for Image Labeling*. Computer Vision and Pattern Recognition (CVPR), 2013.
*The first and second authors made equal contributions and should be considered co-first authors.

(**Spotlight**) Andrew Kae, Gary Huang, Erik Learned-Miller, Carl Doersch. *Improving State-of-the-Art OCR through High-Precision Document-Specific Modeling*. Computer Vision and Pattern Recognition(CVPR), 2010.

(**Oral**) Andrew Kae, Erik Learned-Miller. 2009. *Learning on the fly: Font free approaches to difficult OCR problems*. International Conference on Document Analysis and Recognition (ICDAR), 2009.

Patents

- Andrew Kae, Kin Fai Kan, Vijay K. Narayanan. *Automatic classification of display ads using ad images and landing pages*. Publication number US8732014 B2.

Awards

- NSF East Asia and Pacific Summer Institute (EAPSI) Fellowship 2011.
- Computer Vision and Patten Recognition (CVPR) Doctoral Consortium 2014.

Links

- `https://techcrunch.com/2014/05/04/whichbeers-wants-to-help-you-find-the-right-beer/`
- `https://www.meetup.com/deeplearn/events/238950120/`

Services

Reviewer for
- IEEE Conference on Computer Vision and Pattern Recognition (CVPR).
- International Conference on Computer Vision (ICCV).

Graduate Coursework

- *Computer Science* : Graphical Models, Machine Learning, Database Design and Implementation, Advanced Algorithms, Mining Text and Images with Grid Computation, Information Retrieval.
- *Statistics* : Math Statistics I, II.

Skills

- *Languages*: Python, Scala, Java
- *Frameworks*: AWS, OpenCV, Numpy, PyTorch, TensorFlow, Matlab

References

- Erik Learned-Miller (elm@cs.umass.edu)
  School of Computer Science, University of Massachusetts Amherst

- Benjamin Marlin (marlin@cs.umass.edu)
  School of Computer Science, University of Massachusetts Amherst