



**Department of Electronic Engineering
The Chinese University of Hong Kong**

SEMINAR

Towards Machines that Perceive and Communicate

By

Dr. Kevin Murphy

Research Scientist, Google in Mountain View

California, USA

Date : 27th March, 2018 (Tuesday)

Time : 10:30am

Place : Room 121, Ho Sin Hang Engineering Building, CUHK

Abstract:

I will summarize some recent work related to visual scene understanding and "grounded" language understanding. In particular, I will various recent results from my group at Google including

- Our DeepLab system for semantic segmentation (PAMI'17, 2nd place COCO-Stuff 2017) [1].
- Our object detection system (CVPR '17, 1st place COCO-Detection 2016) [2].
- Our person detection/pose estimation system [3] (CVPR'17, 4th place COCO-Keypoints 2017)
- Visually grounded referring expressions (CVPR'16) [4].
- Discriminative image captioning (CVPR'17) [5].
- Optimizing semantic metrics for image captioning using RL (ICCV'17, 2nd place on COCO-captions 2017) [6]
- Generative models of images and text (ICLR'18). [7]

I will explain how each of these pieces can be combined to develop systems that can better understand images and words.

[1] <https://arxiv.org/abs/1606.00915>

[2] <https://arxiv.org/abs/1611.10012>

[3] <https://arxiv.org/abs/1701.01779>

[4] <https://arxiv.org/abs/1511.02283>

[5] <https://arxiv.org/abs/1701.02870>

[6] <https://arxiv.org/abs/1612.00370>

[7] <https://arxiv.org/abs/1705.10762>

Biography:

Kevin Murphy is a research scientist at Google in Mountain View, California, where he works on AI, machine learning, and computer vision. Before joining Google in 2011, he was an associate professor (with tenure) of computer science and statistics at the University of British Columbia in Vancouver, Canada. Before starting at UBC in 2004, he was a postdoc at MIT. Kevin got his BA from U. Cambridge, his MEng from U. Pennsylvania, and his PhD from UC Berkeley. He has published over 80 papers in refereed conferences and journals, as well as an 1100-page textbook called "Machine Learning: a Probabilistic Perspective" (MIT Press, 2012), which was awarded the 2013 DeGroot Prize for best book in the field of Statistical Science. Kevin was also the (co) Editor-in-Chief of JMLR (the Journal of Machine Learning Research) 2014-2017.

***** All are welcome to attend *****

For inquires, please contact Prof. X G Wang, Tel. No. 3943 8283