


## Core Courses Syllabi

### CV706 - Neural Networks for Object Recognition and Detection

<b>Title</b>	Neural Networks for Object Recognition and Detection
<b>Code</b>	CV706
<b>Loading</b>	4 Credit-hours
<b>Prerequisites</b>	<ul style="list-style-type: none"> <li>A thorough understanding of the material covered in the following Masters course:             <ol style="list-style-type: none"> <li>CV 703 Visual Object Recognition and Detection</li> </ol> </li> <li>Hands-on experience with Python and Pytorch or equivalent language/ library</li> </ul>
<b>Catalog Description</b>	<p>This course provides focused coverage of special topics on object recognition and detection. The students will develop skills to critique the state-of-the-art works on visual object recognition and detection. Moreover, students will be required to implement papers with the following aims: (1) reproduce results reported in the seminal research papers, and (2) improve the performance of the published works. This course builds upon concepts from Human and Computer Vision (CV701), Visual Object Recognition and Detection (CV702) and assumes familiarity with fundamental concepts in machine learning and optimization.</p>
<b>Goal</b>	<p>This PhD course aims to inculcate a deeper understanding of algorithms for image classification, object detection, instance segmentation, semantic segmentation and panoptic segmentation, so the students can become capable of doing research in these specialized topics and can implement these methods for building systems for real-world scene understanding. A significant goal of this course is to enhance students' teamwork skills by requiring them to participate in group projects and develop their communication and analytical skills by engaging them in reading group activities.</p>
<b>Content</b>	<p>The course covers three modules: <b>(I)</b> Image Classification, <b>(II)</b> Object Detection, and <b>(III)</b> Image Segmentation.</p>
<b>Recommended Textbooks</b>	<p>I. Goodfellow, Y. Bengio, and A. Courville, <i>Deep Learning</i>, MIT Press, 2016.</p>
<b>Recommended References &amp; Supplemental Material</b>	<p>Relevant research papers, tech reports, and surveys for each topic, where needed, are identified in the teaching plan ahead. In addition, the following textbooks will be useful:</p> <ol style="list-style-type: none"> <li>D. A. Forsyth, and J Ponce, <i>Computer vision: A Modern Approach</i>, 2nd edition, Pearson, 2003.</li> <li>R. Szeliski, <i>Computer Vision: Algorithms and Applications</i>, Springer Verlag, 2011.</li> </ol>



Teaching Week	Topics
1	<p><b>Image Classification (Module I)</b></p> <p>This module covers the following topics: large-scale image classification and fine-grained image classification.</p> <p><b>Lecture</b></p> <ul style="list-style-type: none"><li>• Overview of image classification</li><li>• Relevant papers and assigned reading:<ul style="list-style-type: none"><li>- A. Krizhevsky, I. Sutskever, and G.E. Hinton, “Imagenet classification with deep convolutional neural networks”, <i>Advances in neural information processing systems</i>, 2012.</li><li>- K. Simonyan, and A. Zisserman, “Very deep convolutional networks for large-scale image recognition”, <i>International conference on learning representations</i>, 2016.</li><li>- K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition”, <i>IEEE conference on computer vision and pattern recognition</i>, 2016.</li><li>- J. Hu, L. Shen, and G. Sun, “Squeeze-and-excitation networks”, <i>IEEE conference on computer vision and pattern recognition</i>, 2018.</li></ul></li></ul> <p><b>Lab</b></p> <ul style="list-style-type: none"><li>• Discussion on choosing a relevant paper to implement for the project</li><li>• Start Project-1 work</li></ul>
2	<p><b>Image Classification (contd.)</b></p> <p><b>Lecture</b></p> <ul style="list-style-type: none"><li>• Introduction to large-scale image classification</li><li>• Reading group activity on selected papers related to “large-scale image classification”:<ul style="list-style-type: none"><li>- A. Krizhevsky, I. Sutskever, and G.E. Hinton, “Imagenet classification with deep convolutional neural networks”, <i>Advances in neural information processing systems</i>, 2012.</li><li>- K. Simonyan, and A. Zisserman, “Very deep convolutional networks for large-scale image recognition”, <i>International conference on learning representations</i>, 2016.</li></ul></li></ul> <p><b>Lab</b></p> <ul style="list-style-type: none"><li>• Continue Project-1 work</li></ul>
3	<p><b>Image Classification (contd.)</b></p> <p><b>Lecture</b></p> <ul style="list-style-type: none"><li>• Continue reading group activity on selected papers related to “large-scale image classification”:<ul style="list-style-type: none"><li>- K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition”, <i>IEEE conference on computer vision and pattern recognition</i>, 2016.</li><li>- J. Hu, L. Shen, and G. Sun, “Squeeze-and-excitation networks”, <i>IEEE conference on computer vision and pattern recognition</i>, 2018.</li></ul></li></ul> <p><b>Lab</b></p> <ul style="list-style-type: none"><li>• Continue Project-1 work</li></ul>




Teaching Week	Topics
4	<p><b>Image Classification (contd.)</b></p> <p><b>Lecture</b></p> <ul style="list-style-type: none"><li>• Introduction to fine-grained image classification</li><li>• Reading group activity on selected papers related to “fine-grained image classification”:<ul style="list-style-type: none"><li>- T. Y. Lin, A. R. Chowdhury, and S. Maji, “Bilinear cnn models for fine-grained visual recognition”, <i>IEEE international conference on computer vision</i>, 2015.</li><li>- Z. Akata, S. Reed, D. Walther, H. Lee, and B. Schiele, “Evaluation of output embeddings for fine-grained image classification”, <i>IEEE conference on computer vision and pattern recognition</i>, 2015.</li><li>- H. Zheng, J. Fu, T. Mei, and J. Luo, “Learning multi-attention convolutional neural network for fine-grained image recognition”, <i>IEEE international conference on computer vision</i>, 2017.</li></ul></li></ul> <p><b>Lab</b></p> <ul style="list-style-type: none"><li>• Preparation of presentation on Project-1 work</li><li>• Continue Project-1 work</li></ul>
5	<p><b>Image Classification (contd.)</b></p> <p><b>Assessment 1.1</b></p> <ul style="list-style-type: none"><li>• Presentation of the projects by different groups</li></ul> <p><b>Lab</b></p> <ul style="list-style-type: none"><li>• Peer review of project reports</li></ul> <p><b>Assessment 1.2</b></p> <ul style="list-style-type: none"><li>• In-class exam covering Module I – Image Classification</li></ul>
6	<p><b>Object Detection (Module II)</b></p> <ul style="list-style-type: none"><li>• This module covers single-stage and two-stage object detectors for conventional imagery and aerial imagery.</li></ul> <p><b>Lecture</b></p> <ul style="list-style-type: none"><li>• Overview of object detection</li><li>• Relevant papers and assigned reading:<ul style="list-style-type: none"><li>- R. Girshick, “Fast R-CNN”, <i>IEEE international conference on computer vision</i>, 2015.</li><li>- S. Ren, K. He, R. Girshick, and J. Sun, “Faster R-CNN: Towards real-time object detection with region proposal networks”, <i>Advances in neural information processing systems</i>, 2015.</li><li>- J. Dai, Y. Li, K. He, and J. Sun, “R-FCN: Object detection via region-based fully convolutional networks”, <i>Advances in neural information processing systems</i>, 2016.</li><li>- W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C. Y. Fu, and A. C. Berg, “SSD: Single shot multibox detector”, <i>European conference on computer vision</i>, 2016.</li><li>- J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, “You only look once: Unified, real-time object detection”, <i>IEEE conference on computer vision and pattern recognition</i>, 2016.</li></ul></li></ul> <p><b>Lab</b></p> <ul style="list-style-type: none"><li>• Discussion on choosing a relevant paper to implement for the project</li><li>• Start Project-2 work</li></ul>



Teaching Week	Topics
7	<p><b>Object Detection (contd.)</b></p> <p><b>Lecture</b></p> <ul style="list-style-type: none"><li>• Introduction to two-stage object detectors</li><li>• Reading group activity on selected papers related to “two-stage object detectors”:<ul style="list-style-type: none"><li>- R. Girshick, “Fast R-CNN”, <i>IEEE international conference on computer vision</i>, 2015.</li><li>- S. Ren, K. He, R. Girshick, and J. Sun, “Faster R-CNN: Towards real-time object detection with region proposal networks”, <i>Advances in neural information processing systems</i>, 2015.</li><li>- J. Dai, Y. Li, K. He, and J. Sun, “R-FCN: Object detection via region-based fully convolutional networks”, <i>Advances in neural information processing systems</i>, 2016.</li></ul></li></ul> <p><b>Lab</b></p> <ul style="list-style-type: none"><li>• Continue Project-2 work</li></ul>
8	<p><b>Object Detection (contd.)</b></p> <p><b>Lecture</b></p> <ul style="list-style-type: none"><li>• Introduction to single-stage object detectors</li><li>• Reading group activity on selected papers related to “single-stage object detectors”:<ul style="list-style-type: none"><li>- W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C. Y. Fu, and A. C. Berg, “SSD: Single shot multibox detector”, <i>European conference on computer vision</i>, 2016.</li><li>- J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, “You only look once: Unified, real-time object detection”, <i>IEEE conference on computer vision and pattern recognition</i>, 2016.</li></ul></li></ul> <p><b>Lab</b></p> <ul style="list-style-type: none"><li>• Continue Project-2 work</li></ul>
9	<p><b>Object Detection (contd.)</b></p> <p><b>Lecture</b></p> <ul style="list-style-type: none"><li>• Applying object detectors on aerial images</li><li>• Reading group activity on selected papers related to “object detection in aerial images”:<ul style="list-style-type: none"><li>- G. S. Xia, X. Bai, J. Ding, Z. Zhu, S. Belongie, J. Luo, M. Datcu, M. Pelillo, and L. Zhang, “DOTA: A large-scale dataset for object detection in aerial images”, <i>IEEE conference on computer vision and pattern recognition</i>, 2018.</li><li>- J. Ding, N. Xue, Y. Long, G. S. Xia, and Q. Lu, “Learning RoI transformer for detecting oriented objects in aerial images”, <i>IEEE conference on computer vision and pattern recognition</i>, 2019.</li></ul></li></ul> <p><b>Lab</b></p> <ul style="list-style-type: none"><li>• Preparation of presentation on Project-2 work</li><li>• Continue Project-2 work</li></ul>



Teaching Week	Topics
10	<p><b>Object Detection (contd.)</b></p> <p><b>Assessment 2.1</b></p> <ul style="list-style-type: none"><li>• Presentation of the projects by different groups</li></ul> <p><b>Lab</b></p> <ul style="list-style-type: none"><li>• Peer review of project reports</li></ul> <p><b>Assessment 2.2</b></p> <ul style="list-style-type: none"><li>• In-class exam covering Module II – Object Detection</li></ul>
11	<p><b>Image Segmentation (Module III)</b></p> <p>This module covers:</p> <ol style="list-style-type: none"><li>1. Instance segmentation,</li><li>2. Semantic segmentation, and</li><li>3. Panoptic segmentation.</li></ol> <p><b>Lecture</b></p> <ul style="list-style-type: none"><li>• Overview of image segmentation</li><li>• Relevant papers:<ul style="list-style-type: none"><li>- K. He, G. Gkioxari, P. Dollár, R. and Girshick, “Mask R-CNN”, <i>IEEE international conference on computer vision</i>, 2017.</li><li>- S. Liu, L. Qi, H. Qin, J. Shi, and J. Jia, “Path aggregation network for instance segmentation”, <i>IEEE conference on computer vision and pattern recognition</i>, 2018.</li><li>- J. Long, E. Shelhamer, and T. Darrell, “Fully convolutional networks for semantic segmentation”, <i>IEEE conference on computer vision and pattern recognition</i>, 2015.</li><li>- L. C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, A.L., “Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs”, <i>IEEE transactions on pattern analysis and machine intelligence</i>, 2016.</li><li>- H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, “Pyramid scene parsing network”, <i>IEEE conference on computer vision and pattern recognition</i>, 2017.</li><li>- A. Kirillov, K. He, R. Girshick, and P. Dollár, “Panoptic segmentation”, <i>IEEE conference on computer vision and pattern recognition</i>, 2019.</li><li>- A. Kirillov, R. Girshick, K. He, and P. Dollár, “Panoptic feature pyramid networks”, <i>IEEE conference on computer vision and pattern recognition</i>, 2019.</li></ul></li></ul> <p><b>Lab</b></p> <ul style="list-style-type: none"><li>• Discussion on choosing a relevant paper to implement for the project</li><li>• Start Project-3 work</li></ul>
12	<p><b>Image Segmentation</b></p> <p><b>Lecture</b></p> <ul style="list-style-type: none"><li>• Instance segmentation</li><li>• Reading group activity on selected papers related to “instance segmentation”:<ul style="list-style-type: none"><li>- K. He, G. Gkioxari, P. Dollár, R. and Girshick, “Mask R-CNN”, <i>IEEE international conference on computer vision</i>, 2017.</li><li>- S. Liu, L. Qi, H. Qin, J. Shi, and J. Jia, “Path aggregation network for instance segmentation”, <i>IEEE conference on computer vision and pattern recognition</i>, 2018.</li></ul></li></ul> <p><b>Lab</b></p> <ul style="list-style-type: none"><li>• Continue Project-3 work</li></ul>



Teaching Week	Topics
13	<p><b>Image Segmentation</b></p> <p><b>Lecture</b></p> <ul style="list-style-type: none"><li>• Semantic segmentation</li><li>• Reading group activity on selected papers related to “semantic segmentation”:<ul style="list-style-type: none"><li>- J. Long, E. Shelhamer, and T. Darrell, “Fully convolutional networks for semantic segmentation”, <i>IEEE conference on computer vision and pattern recognition</i>, 2015.</li><li>- L. C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, A.L., “Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs”, <i>IEEE transactions on pattern analysis and machine intelligence</i>, 2016.</li><li>- H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, “Pyramid scene parsing network”, <i>IEEE conference on computer vision and pattern recognition</i>, 2017.</li></ul></li></ul> <p><b>Lab</b></p> <ul style="list-style-type: none"><li>• Continue Project-3 work</li></ul>
14	<p><b>Image Segmentation</b></p> <p><b>Lecture</b></p> <ul style="list-style-type: none"><li>• Panoptic segmentation</li><li>• Reading group activity on selected papers related to “panoptic segmentation”<ul style="list-style-type: none"><li>- A. Kirillov, K. He, R. Girshick, and P. Dollár, “Panoptic segmentation”, <i>IEEE conference on computer vision and pattern recognition</i>, 2019.</li><li>- A. Kirillov, R. Girshick, K. He, and P. Dollár, “Panoptic feature pyramid networks”, <i>IEEE conference on computer vision and pattern recognition</i>, 2019.</li></ul></li></ul> <p><b>Lab</b></p> <ul style="list-style-type: none"><li>• Preparation of presentation on Project-3 work</li><li>• Continue Project-3 work</li></ul>
15	<p><b>Image Segmentation</b></p> <p><b>Assessment 3.1</b></p> <ul style="list-style-type: none"><li>• Presentation of the projects by different groups</li></ul> <p><b>Lab</b></p> <ul style="list-style-type: none"><li>• Peer review of project reports</li></ul> <p><b>Assessment 3.2</b></p> <ul style="list-style-type: none"><li>• In-class exam covering Module III – Image Segmentation</li></ul>