

Date	Heading	Main Paper	Recommended Readings/Papers
01/24/22	LeNet	Gradient-Based Learning Applied to Document Recognition	All about convolutions : A guide to convolution arithmetic for deep learning
	Deep model problems	Understanding the Difficulty of training Deep feedforward Neural Networks	Activation Functions: Comparison of trends in Practice and Research for Deep Learning
01/31/22	AlexNet	ImageNet Classification with Deep Convolutional Neural Networks	Visualizing and Understanding Convolutional Networks
	VGGNet	Very Deep Convolutional Networks for Large-Scale Visual Recognition	Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification
02/07/22	Deep-supervision	Deeply supervised nets	Dropout: A Simple Way to Prevent Neural Networks from Overfitting
	Net-in-Net	Network in Network	
	Knowledge Distillation	Distilling the Knowledge in a Neural Network	
02/14/22	Inception Versions	GoogLeNet: Going Deeper with Convolutions	Inception v4, inception-ResNet and the impact of residual Connections on Learning
		Inception v3: Rethinking the inception architecture for computer vision	Inception v2: Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift,
02/21/22	ResNet	Deep Residual Learning for Image Recognition	
	Identity mapping	Identity Mapping in Deep Residual Networks	
02/28/22	Wide ResNet	Wide Residual Networks	Deep Networks with Stochastic Depth
	ResNeXt	Aggregated Residual Transformations for Deep Neural Networks	
03/07/22	DenseNet	Densely Connected Convolutional Networks	Memory efficient implementation of DenseNets
	U-Net	U-Net: Convolutional Networks for Biomedical Image Segmentation	Inception-ResNet, inception-Densenet
03/14/22	Spring Break		
03/21/22	Transformer	Attention is all you need	
	Vision Transformer	An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale	https://www.youtube.com/watch?v=3B6q4xnuFUE&list=PLcp6ZnH4WYlbSPL-zrFNOyqUiYV6HN6sK&index=3&ab_channel=AIBites
03/28/22	DeiT	Training data-efficient image transformers & distillation through attention	
	Swin Transformer	Swin Transformer: Hierarchical Vision Transformer using Shifted Windows	https://www.youtube.com/watch?v=tFYxJZBAE8&ab_channel=AIBites https://sieunpark77.medium.com/swin-transformers-the-most-powerful-tool-in-computer-vision-659f78744871
04/04/22	Locality in ViT	CvT: Introducing Convolutions to Vision Transformers	Xception: Deep Learning with Depthwise Separable Convolutions
		Local ViT: Bringing Locality to vision transformers	Demystifying local vision transformers: sparse connectivity, weight sharing, and dynamic weight
04/11/22	SAN	Exploring self-attention for image recognition	
	No attention	Do you even need Attention? A stack of feed-forward layers does surprisingly well on ImageNet	
		When Shift Operation Meets Vision Transformer: An Extremely Simple Alternative to Attention Mechanism (partial shift operation)	All you need is a few shifts: designing efficient convolutional neural networks for image classification Active shift: Constructing fast network through deconstruction of convolutions Shift: A Zero FLOP, Zero Parameter Alternative to Spatial Convolutions (grouped shift)
04/18/22	All MLP models for Vision	MLP-Mixer: An all-MLP Architecture for Vision	Feedforward networks for image classification with data-efficient training
		cycleMLP: A MLP like architecture for dense prediction	S2MLP: Spatial-shift MLP architecture for vision
04/25/22	ConvNext	A Battle of Network Structures: An Empirical Study of CNN, Transformer, and MLP	
		A ConvNet for the 2020s	