Weeks 1-2		**Unit 1: Image Processing Fundamentals**
Lecture 1	Images & Color	
	Video lecture	Digital images, RGB image, color spaces, alpha channel MATLAB image i/o, fundamentals
	Live lecture	Computing tools for image processing and computer vision, histograms, human color vision, color spaces, image file formats, CV course plan overview
Lecture 2	Image Filtering	
	Video lecture	Convolution, filtering kernels, box filter, Gaussian filter, MATLAB filtering and visualizations, shift kernel, denoising, sharpening, median filtering, thresholding
	Live lecture	Convolution deep dive, convolution by hand over examples, linearity of convolution and convolution arithmetic, separable kernels, convolution in practice
Lecture 3	Edge Detection	
	Video lecture	Sobel kernel, derivative of Gaussian, image gradients, 2D filtering, Canny edge detection, Laplacian kernel, digital filters, all with MATLAB demonstrations
	Live lecture	Convolution recap, Sobel kernel in action, derivative vs difference, images as 2D functions, x/y gradients, gradient magnitude and phase, derivative of Gaussian theory, linearity of convolution
Lecture 4	Deep Learning	
	Video lecture	Deep learning fundamentals, functions of functions, backpropagation, supervised learning, encoder/decoder architectures, deep features, dataset augmentation case study: alpha blending and compositing in MATLAB
	Live lecture	Deep/supervised learning fundamentals, "machine learning is just statistics", capabilities of modern AI tools, why you shouldn't believe Elon Musk / Sam Altman, why you shouldn't trust ChatGPT outputs, misuse of AI, dangers of AI, discussion of current AI news
Lecture 5	Signals & Images	
	Video lecture	Fourier series, 2D periodic functions, high-frequency/low-frequency images, Fourier transform in MATLAB, frequency domain, magnitude and phase, linearity of FT, box vs Gaussian filters in freq domain, low-pass/high-pass filters
	Live lecture	why it is called "filtering", what do we mean by "signals", filtering in frequency domain deep dive, low-pass/high-pass/band-pass filters, shifted impulse kernel, denoising in frequency domain
Lecture 6	Sampling & Aliasing	
	Video lecture	Digital signals, sampling theory, sampling in frequency domain, sub-sampling an image in MATLAB Guest lecture: Anti-aliasing in CNNs by R. Zhang of Adobe
	Live lecture	Sampling and aliasing in frequency domain theory deep dive, how to resize an image *Assignment 1 detailed overview*
Assignment 1: Image Processing		
	Review image processing fundamentals in MATLAB following the MATLAB tutorials in video lectures, including frequency domain representations, anti-aliasing, and edge detection.	

Weeks 3-4	**Unit 2: Image Registration**	
Lecture 7	Harris Corner Detection	
	Video lecture	Image features, relating two images of the same scene, Harris corner detection theory and derivation, Harris corner detection implementation in MATLAB
	Live lecture	Step by step derivation of Harris corners, what is an energy function, how mathematical derivations work, Taylor series expansion, Eigenvalues and Eigenvectors, the ellipse equation, Harris corners vs Harris measure, relating edge detection with corner detection, convolution revisited
Lecture 8	Feature Invariance, Detection, and Matching	
	Video lecture	Why we match features, feature invariance, 3D scale space, MOPS, SIFT detection and description, feature discriminability,
	Live lecture	Harris corners rotation invariance, scale non-invariance, scale selection, Gaussian and Laplacian pyramids, their relation to low-pass/high-pass filtering, Laplacian of Gaussian, SIFT and band-pass filtering, local gradients and feature description, feature matching
Lecture 9	Transformations & Image Alignment	
	Video lecture	Image alignment, global image warping, linear transformations, homogeneous coordinates, affine transformations, image warping with homographies, least squares
	Live lecture	Linear transformations deep dive, homogeneous coordinates, image warping and interpolation, least squares derivation deep dive, line fitting, linear systems of equations, Ax=b and pseudo-inverse, under-/over-constrained linear systems
Lecture 10	RANSAC	
	Video lecture	What is an outlier, robustness to outliers, RANSAC implementation and behavior
	Live lecture	Least squares and outliers theory, RANSAC overview, feature detection, description and matching overview, the image stitching pipeline, catch-up discussions *Assignment 2 detailed overview*
Assignment 2: Image Stitching		
	Implement a new algorithm (FAST point detector) using their MATLAB background, integrate the points into existing data structures in MATLAB, and implement an image stitching pipeline using their own photographs.	

Weeks 4-6		**Unit 3: Computer Vision Fundamentals**
Lecture 11	Image Segmentation	
	Video lecture	Histograms and thresholding, uncertainty in segmentation, clustering, k-means, texture representations, mean-shift segmentation, superpixels
	Live lecture	High-level/low-level segmentation, segmentation in modern CV literature, unsupervised learning and clustering, k-means deep dive, hyperdimensional/deep features and k-means, spatial coherency and smoothness losses in optimization and deep learning
Lecture 12	Optical Flow	
	Video lecture	Motion in video, optical flow definition, aperture problem, Lukas-Kanade flow, multi-resolution flow estimation, Shi-Tomasi good features to track
	Live lecture	Key assumptions of optical flow, detailed derivation of Kanade-Lukas, how mathematical derivations work revisited, spatial coherence revisited, linear systems revisited, parallels between Harris corners and Kanade-Lukas deep dive, multi-resolution and working around the small motion assumption
Lecture 13	Cameras	
	Video lecture	Pin-hole camera model, simple lens models, Bokeh and aperture, perspective distortion, focal length, simple 2D projection, examples from movie production, global/rolling shutter
	Live lecture	Pin-hole camera, aperture, focal length, and perspective deep dive; real-world cameras and lenses, 2D projection deep dive, focal length and aperture in real-world photography
Lecture 14	Projection	
	Video lecture	3D projection model, orthographic vs perspective projection, how 3D world becomes a 2D photograph, camera parameters, homogeneous coordinates revisited, intrinsic/extrinsic camera parameters, the projection matrix, lens distortions
	Live lecture	3D coordinate systems, world vs camera coordinates, projection mathematical deep dive, focal length vs distance, orthographic and perspective projection in computer graphics, camera parameters deep dive *How to prepare for the mid-term exam and exams in general*
Lecture 15	Stereo	
	Video lecture	Rectified stereo pairs, epipolar geometry, stereo matching, disparity estimation, stereo as energy minimization, depth vs disparity, stereo reconstruction pipeline, structured lighting and laser scanning
	Live lecture	Stereo matching with structural similarity deep dive, smoothness energy revisited, window search and its similarities with feature description and matching, 1D energy minimization deep dive, 3D reconstruction in modern CV literature, relation of CV topics with CG topics