

MSc CGVI

Computer Graphics, Vision and Imaging

The MSc CGVI (computer graphics, vision and imaging) is a one-year full-time advanced Masters course at UCL providing training in computer graphics, virtual reality, machine vision and imaging technology from world-leading experts.

The fields of graphics, vision and imaging are converging and increasingly rely on one another. This unique and timely course allows students to specialize in any of the areas while providing the essential grounding in the others. Graduate prospects are excellent for careers in computer graphics and special effects, medical imaging, robotics, TV and film, gaming, vision engineering, image and video databases, surveillance and many related applications. The course also leads on naturally to PhD study or industrial research and development. Specifically, the course provides

- Mathematical and algorithmic principles that underlie computer graphics, machine vision and imaging science.
- Theoretical understanding and practical experience of the standard techniques in the areas.
- Training with standard software libraries and development tools.
- Training in research and development methodology.
- The ability to evaluate new techniques and assess viability for industrial or commercial exploitation.
- The opportunity to work on an extended cutting edge research project supervised by a leading expert in the area

MSc-CGVI graduates' unique combination of technical skills makes them first choice for the best employers, since similar-looking Master's programs do not give the same breadth of topics and technical detail. Feedback from employers regularly confirms the value of the broad technical knowledge our graduates acquire.

The core areas covered are Computer Graphics, Medical Image Processing, Virtual Environments and Machine Vision:

Core Area 1: Computer Graphics

Computer-generated imagery is becoming more and more ubiquitous and is used in a vast number of applications, including scientific visualization, special effects, virtual news-anchors and computer games. Computer graphics is concerned with methods to digitally synthesize and manipulate visual and geometric content that is used in these applications. It focuses on the computational and mathematical challenges of image generation as opposed to aesthetic issues. There are a variety of subfields covering the different aspects of image synthesis: *modelling* deals with the acquisition, manipulation and representation of shapes and appearance; *animation* is concerned with the movement and deformation of characters; and *rendering* (see figures) deals with the actual image synthesis given a scene description.



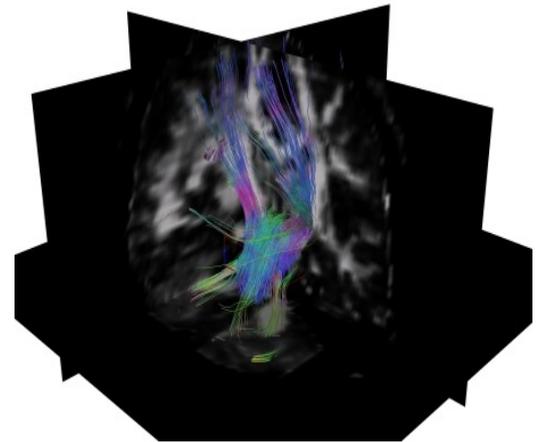
The field of computer graphics has matured over the past decade and it is now possible to render rather realistic looking images. Yet, many open and challenging problems remain that limit the generation of computer-generated imagery. For instance, the computation times of high-quality, realistic images are very long. A single movie frame often takes hours to compute, yet many applications such as flight-simulators and architectural walkthroughs mandate immediate feedback. Simple methods for these applications exist but require a substantial decrease in quality and realism to meet time constraints. It is an open research question how to maintain high quality and high frame rates. Other challenges include efficient authoring of realistic scenes (geometry, appearance, etc.). Currently, most data are authored by hand, which allows for flexibility but is also very tedious. Acquisition of real-world objects is a viable alternative, but most current techniques are cumbersome to use. More research is needed to make acquisition of real objects as easy as using a digital camera. In summary, computer graphics is an academically challenging field due to its high computational demand but at the same time research results can have immediate practical impact.

Core Area 2: Medical Image Processing

Medical imaging aims to produce pictures of the body that reveal information about its normal internal structure and mechanisms, as well as changes that occur in disease (e.g. figure illustrates tracing fibre direction using diffusion MRI). Most medical imaging techniques work by solving an inverse problem that relates properties of the tissue, such as x-ray absorption or water content, inside the body to measurements that an imaging device makes outside the body. For example, a magnetic resonance imaging (MRI) scanner measures the relative strengths of magnetic fields oscillating at different frequencies and infers an image of water density, which provides contrast between different types of tissue. Automated *image processing* methods are increasingly responsible for extracting useful information from such images and highlighting areas that may be abnormal or indicate disease.

For example, one of the most important medical image computing problems is *registration*: to compare two brain scans of the same patient at different times, we need to remove differences in the images that arise from different positioning of the head inside the scanner. Image registration aligns one image to the other so we can subtract them and look for changes.

Another key challenge is to compute reliable *biomarkers*. A biomarker is a feature of the image that shows the presence of a disease or the effect of a treatment. They may be very simple features, such as the image intensity in a certain region of the brain, or require sophisticated computations and comparisons with atlases, such as the degree of atrophy in grey matter of Alzheimer's patients compared to the average brain. Modern medical imaging combines knowledge of biology, medicine, imaging devices and image analysis techniques to identify the most descriptive biomarkers.



Core Area 3: Virtual Environments



A virtual environment is a real-time computer simulation that exploits human ability to understand 3D data. Virtual environments, especially those in computer games or training simulators, thus often look like depictions of real places. However, their form is only limited by our imaginations, and their utility is only limited by their ability to communicate information effectively. In terms of graphics, the current state of the art of virtual environments can be found in the games industry. Although successful as a medium, games provide a very simple user-interface for interacting with 3D data: one must use the joystick or simple motion controls and the environments themselves are usually quite restricted in their behaviour. Compare this with our experience of the everyday world, where behaviour is very complex, and we experience multimodal sensory information from all around us.

Research in virtual environments falls in to three themes: *hardware*, *software*, and *perception*. Hardware research investigates increasing the display quality. For visual displays the state of the art are CAVE-like displays. These surround the user with display surfaces, "immersing" them in virtual environment so that they see an egocentric view. The user's head is tracked, so images appear life-sized and in stereo: a chair in a CAVE-like system can be walked around and looks like it could be sat on. Despite their success in many applications, such displays still have many deficiencies: they aren't bright, their dynamic range is low and their end-to-end latency is tens of milliseconds. There are also significant challenges in stimulating other senses such as touch, taste, smell, hearing, balance, proprioception, body awareness, etc.

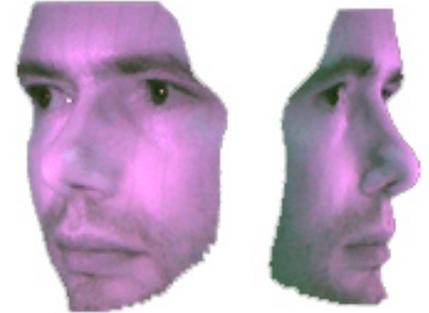
As virtual environments increase in complexity, the related software becomes challenging to manage. Different physical effects (e.g. weather and crowd simulation) are modelled in very different ways but must be integrated in our virtual environments. Perhaps the biggest challenge in this area is that of creating plausible virtual humans that can interact with the user. The final research theme concerns what is necessary for hardware and software systems from a perceptual point of view. Despite the relative primitiveness of our current displays, users have very strong reactions to the virtual environments in certain situations. Users in the UCL CAVE-like system who experience standing on a virtual ledge over a precipice suffer strong stress reactions, even though they know the drop is an illusion; Thus the brain is quite successfully interpreting some set of minimal cues and if we can understand more about how the brain comes to an operational model of its environment, we can better design displays.

Core Area 4: Computer Vision

Computer vision is concerned with developing artificial systems which extract information from image or video data. It is a hybrid subject at the junctions between neuroscience, machine learning, geometry, mathematics and electrical engineering. Automated vision is the opposite process to computer graphics: the aim is to construct a model of the 3d scene from images rather than vice-versa. Unfortunately, the mapping from scenes to images is many to one: there may be several possible configurations of the world that produce exactly the same image. This is one of the reasons vision is challenging.

Vision is also difficult because of the sheer variation in appearance of objects. For example, when we compare images of two dogs (e.g. a poodle and a great dane), the actual pixel values have very little in common. However, humans have no problem recognizing both images as examples of dogs and distinguishing them from images of cats: we know vision is possible as the human visual system provides an existence proof. In fact more than one third of the brain is involved in vision, suggesting that considerable computational resources are required.

Vision tasks include: *reconstruction* in which we build a three dimensional model of the scene from one or more images (right), *camera tracking* where we identify the movement of the camera relative to the scene, *object detection* in which we determine if a certain type of object (e.g. a dog) is in the image and *segmentation* (above right) in which we divide the image up into meaningful regions. Applications of machine vision include robotics, face recognition, content based image and video retrieval, building 3d models from photos and many industrial applications.



MSc CGVI Program Structure

The MSc CGVI runs for one calendar year starting in the last week of September and ending one week into September the following year. It is not possible to take the courses on a part-time basis. The MSc course consists of three elements, all of which must be completed and passed for the student to obtain a degree:

- A core component consisting of essential advanced material for all students. All material in this part is compulsory.
- Options to allow tailoring of the material to individual interests. Students choose four topics from the list below.
- A research project.

Taught courses consist of about 30 hours of lectures and are assessed in the examination period (May / June). In addition, most have associated practical sessions and coursework. Syllabi are appended to the end of this document.

Core Components

- Mathematical Methods, Algorithmics and Implementations (Term 1)
- Image Processing (Term 1)
- Computer Graphics (Term 1)
- Research Methods (Term 2)

Options (select 4)

- Machine Vision (Term 1)
- Geometry of Images (Term 2)
- Virtual Environments (Term 1)
- Advanced Rendering and Animation (Term 2)
- Inverse Problems in Imaging (Term 2)
- Computational Modelling for Biomedical Imaging (Term 2)
- Computational Photography and Capture (Term 2)
- Graphical Models (Term 1)

The project work, which starts after the exams in June and occupies students full-time until early September, is intended to provide an extended opportunity to plan, execute and evaluate a significant piece of work, working closely with an expert in the field. Projects will either be related to a problem of industrial interest or to a topic near the leading edge of research. Examples of past projects are detailed on the next page.

Example Projects

Name: Mahdi MohammadBagher **Year:** 2008-2009

Project title: Screen-Space Percentage-Closer Soft Shadows (SS-PCSS)

Since rendering soft shadows is computationally expensive, Mahdi proposed rendering Percentage-Closer Soft Shadows (PCSS), one of the state of the art techniques in soft shadow rendering, inside a screen-space rendering loop. Edge-aware filtering such as cross-bilateral filtering is required to address the issue of losing the sense of edges in screen space and this was implemented efficiently using a separable kernel. The results are visually comparable to traditional soft shadow algorithms as well as the ground truth while being very fast to compute.



Name: Fabrizio Pece **Year:** 2008-2009

Project title: High Dynamic Range for Dynamic Scenes

Digital cameras cannot capture the dynamic range of colours (ratio between dark and bright regions) in the real world. High Dynamic Range (HDR) photography overcomes this limitation, but unfortunately it is not suitable for dynamic scenes. Moving objects produce undesirable 'ghosts' in the HDR images. Fabrizio's project developed techniques to detect moving objects in a scene described by a bracketed exposure sequence and to erase the ghosts generated by these movements in the corresponding HDR images.

Name: Thomi Mertzanidou **Year:** 2007 - 2008

Project Title: Image parsing.

Thomi's project involved image parsing which is the process of trying to associate a label with each pixel of the image. In a small region of the image, the visual information is very ambiguous (e.g. the sky and sea look much the same). However, by combining information about context it is possible to reason about what is going on. For instance, we may see a chair next to a table, above the floor and surrounded by wall, but are unlikely to see a chair above a table on top of a window. Thomi's project investigated incorporating context into image labeling.



Name: Frederic Besse **Year:** 2007 - 2008

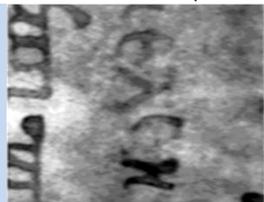
Project Title: Panoramic videos and time - flow manipulation

The goal of Frederic's project was to explore time-flow editing, to generate dynamic panoramic videos. He used time-flow editing to create composite images in which different parts of the image come from different moments in time. For example, in the original video the entire stadium (pictured left) collapsed simultaneously. However, with a modified time front different pixels come from different times in the video so that the right half collapses first. He then combined this idea with image panorama techniques to make panoramic videos.

Name: Laura Panagiotaki **Year:** 2006-2007

Project Title: Automated camera placement

Current video-game camera control techniques are criticised for inadequate capture of game action. Laura's project developed and evaluated a set of real-time automated algorithms for the movement and placement of virtual cameras. These addressed key limitations of existing techniques by drawing on cinematographic principles to drive the autonomous control system. The camera reacts to changes in the game environment in real-time, and allows control parameters to be tailored to maximise dramatic impact and playability.



Name: Saurabh Sethi **Year:** 2006-2007

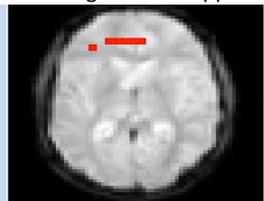
Project Title: Text retrieval from Archimedes palimpsest.

Saurabh's project developed new image-processing techniques for recovering hidden text from multispectral images of an ancient document: the Archimedes palimpsest. As was common in the middle ages, the author took old parchment, scraped off the original writing and wrote his own text on top. Now historians are interested in the text written underneath, which is still discernible in places. This project used state-of-the-art computer vision techniques to enhance the hidden text.

Name: Umar Mohammed **Year:** 2005-2006

Project Title: Generative Models for Face Recognition

Face recognition algorithms have many real-world applications, but current approaches are not sufficiently reliable for widespread market acceptance. Umar's project developed a new approach to face recognition based on recent developments in machine learning. The algorithm calculates the probability that faces have an underlying common cause (they come from the same person). Umar's experiments demonstrate several advantages of this approach over the current state of the art.



Name: Jania Aghajanian **Year:** 2005-2006

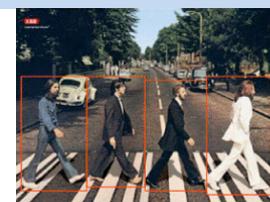
Project Title: Predicting cognitive states from fMRI using support vector machines.

Jania's project wrote an image processing system for a mind reading competition run by the Human Brain Mapping organization in 2006. The competition provided functional MRI scans taken every few seconds from subject watching a movie inside an MRI scanner. Competitors had to write a computer program to determine the contents of the movie from the subjects scans. Jania combined image processing techniques with machine learning to predict cognitive states from features of the fMRI data.

Name: Sun Hyun Lee **Year:** 2005-2006

Project Title: Pedestrian Detection and Segmentation based on human silhouette models in a still image

Sung Hyun's project is on pedestrian detection in static images, which is a challenging task as the shape and appearance of human beings exhibits considerable variability. He built a model that describes the family of shapes that human beings can take and then searched the image for these shapes. He took real world images and fitted this model using an iterative method that built up a color model for each detected person. Finally the pedestrian was segmented using a Markov Random Field.



Our Graduates

The MSc CGVI has been running for more than ten years (formerly named MSc VIVE) and many of our graduates have gone on to have exciting careers. Below are a sample of our graduates to show you possible career paths after the MSc CGVI.

D. Turmukhambetov (2010-2011)	Daniyar's project was concerned with intrinsic images, supervised by Gabriel Brostow. He is now pursuing a PhD with Jan Kautz at UCL.
Philippe Leveux (2010-2011)	Philippe did his project on light fields for dynamic objects under the supervision of Jan Kautz. He now works for Vision Semantics Ltd in London.
Sam Birley (2009-2010)	Sam did a project on high-dynamic-range rendering in the context of computer games supervised by Jan Kautz. He now works at Headstrong Games (part of Kuju) in London.
Fiona Liu (2009-2010)	Fiona did a project on tracking faces under the supervision of Anthony Steed. After completing her MSc, she started work for Framestore in London.
Andrew Sweet 2008-2009	Andrew did a project on medical imaging supervised by Professor Daniel Alexander. After completing the M.Sc. he went to work at the INRIA research institute in Sophia Antipolis, France.
Cristina Amati 2008-2009	Cristina's project was supervised by Gabriel Brostow and concerned the automatic animation of hand-drawn images. She now works for London-based computer vision/graphics company, Anthropics.
Frederic Besse 2007-2008	Frederic did a project in computational photography supervised by Dr. Jan Kautz. After the masters he went to work for Double Negative Visual Effects where he is a research and development programmer.
Thomi Mertzaniidou 2007-2008	Thomi did a project on computer vision. After the MSc she decided to stay at UCL study for a Ph.D. in medical image processing, supervised by Dave Hawkes.
Saurabh Sethi 2006-07	Saurabh did a project on computer vision techniques for enhancing ancient document images with Danny Alexander. After completing the MSc, he took an internship with Siemens Corporate Research in Princeton, USA.
Zhexin Yang 2005-2006	Zhexin did a computer graphics project on skin-texture modelling for computer graphics with Celine Loscos. After completing the MSc, she took a job with the computer-games company Real Time Worlds in Dundee, Scotland.
Guang Yang 2005-06	Guang did a project with Dr. Simon Prince on generating random photorealistic faces. After completing the MSc, he moved back to China to take a job with Siemens as a software developer.
Zhuomin Liu 2005-06	Zhuomin did a project with Dr. Celine Loscos on representing avatars in virtual environments. After completing the MSc, she found a job working as a software developer for a digital TV company in the UK.
Sung-Hyun Lee 2005-06	Sung-Hyun did a computer vision project on detecting pedestrians in CCTV footage with Dr. Simon Prince. After the MSc he got a job working in image processing for Samsung in Korea.
Jania Aghajanian 2005-06	Jania did a medical imaging/machine learning project on predicting cognitive states from fMRI with Dr. Daniel Alexander. Since she has done an internship with Google, UK and is now doing a PhD in computer vision at UCL.
Sylvia Xeuni Pan 2004-05	Sylvia did a computer graphics project supervised by Celine Loscos. After the MSc, she started a PhD on designing realistic avatars for virtual environments and, specifically, on realistic facial expressions for virtual humans.
Jamie Wither 2004-05	Jamie did a computer graphics project on photorealistic rendering supervised by Prof. Mel Slater. After the MSc he went to INRIA, Rhones-Alpes in Grenoble, France to do a PhD.
Kate Bergel 2004-05	Kate did a computer graphics project on texture synthesis. After the MSc, Kate worked for the computer-games company Electronic Arts and then moved on to work for a movie-rendering company.
Mitchelle Chen 2004-05	Mitchelle did a graphics project. After the MSc she went to work for Hewlett-Packard in China on a collaborative project with the Chinese ministry of education and various universities to create large-scale digital museums.
Hyun-Jung Kim 2004-05	Hyun-Jung did a computer graphics project supervised by Dr. Celine Loscos. After the MSc he took a job a digital map company in Korea making digital maps and car navigation software.
Shahrum Nedjati-Gilani 2003-04	Shahrum did a medical imaging project on detected multiple-sclerosis lesions in MRI. After the course, he started an EngD, sponsored by Philips Medical Systems, on image reconstruction in diffusion MRI.
Shree Pavar 2003-04	Shree did a project on adding realistic gestures in avatars in virtual environments. After the MSc, he started an EngD in the Centre for Computational Sciences, which is part of the chemistry department at UCL.

Russel Freeman 2003-04	Russel did a computer graphics project supervised by Dr. Anthony Steed. After the MSc, he took a PhD, also supervised by Anthony Steed, on mixed reality.
Ashutosh Chhibbar 2003-04	After completing the MSc VIVE, Ashutosh landed a job with Qinetiq, but decided to reject the offer due to an opportunity to move to California and create a start-up company whose central focus is vision and imaging.
Katrien Jacobs 2002-03	Katrien did a computer graphics project supervised by Dr. Celine Loscos. She completed her PhD, also supervised by Celine Loscos, on radiosity in 2006 and is now a consultant in the commercial sector.
Dan Borthwick 2000-01	After completing the course, Dan went to work for the computer games company Ideaworks3D. He has worked on titles such as: The Sims, Need For Speed Underground and Final Fantasy VII: Dirge of Cerberus.

CGVI Staff



Name: Professor Daniel Alexander

Roles: Personal tutor. **Lecture courses:** Image Processing, Research Methods.

Profile: Professor Alexander is a Reader in Imaging Sciences in and also works within the UCL Center for Medical Image Computing (CMIC). He has worked in the areas of computer vision, image and audio processing and medical imaging for over 10 years and has over 100 peer-reviewed publications. Daniel is best known for his work in neuroimaging and magnetic resonance imaging, in particular in diffusion MRI. He has strong industry links with MRI scanner manufacturers and pharmaceutical companies, such as Philips and GSK. He is associate editor for IEEE Transactions on Medical Imaging.

Name: Prof. Simon Arridge

Roles: Personal tutor. **Lecture courses:** Optimization, Geometry of Images.

Profile: Prof. Arridge is Professor of Image Processing and head of the Vision and Imaging Sciences Group and also a member of the UCL CMIC. He has worked in the areas of inverse problems, image processing and medical imaging for over 20 years. Simon is best known for his work in inverse problems, specifically classical and stochastic solutions to ill-posed problems, with application to optical tomography and other medical imaging modalities. He is a member of the editorial board for the Institute of Physics journal Inverse Problems.



Name: Dr. Gabriel Brostow

Roles: Personal tutor. **Lecture courses:** Image Processing, Computational Photography and Capture

Profile: Dr. Brostow is a Lecturer in the Vision and Imaging Sciences group. His research is about both the analysis and synthesis of motion. This area includes interesting research problems in the areas of: Data-driven animation, applied computer vision, computational photography and video, motion perception and segmentation, action and object recognition, and performance capture. He is an active member of both the vision and graphics research communities, and has strong ties with the main special effects and animation companies.

Name: Dr. Lewis Griffin

Roles: Personal tutor. **Lecture courses:** Geometry of Images.

Profile: Dr. Griffin is Senior Lecturer in the Vision and Imaging Sciences group and is also a member of the UCL CMIC. His area of expertise is Computational Vision with particular emphasis on Spatial Vision and Colour Vision. Lewis is well known for his recent work on using the human vision system to inspire computer vision algorithms and systems. He is a member of the editorial board of the Elsevier journal Computers in Medicine and Biology.



Name: Dr. Simon Julier

Roles: Personal tutor. **Lecture courses:** Mathematical Methods, Algorithms and Implementation, Virtual Environments.

Profile: Dr. Julier is Senior Lecturer in the Virtual Environments and Computer Graphics group. His key interests are in all aspects of nonlinear estimation, data fusion, and augmented reality. Simon is best known for his work on Kalman filtering and the SLAM (Simultaneous Localization and Mapping) problem in robotics and computer vision. He has twice been cochair of the IEEE Virtual Reality conference.

Name: Dr. Jan Kautz

Roles: Personal tutor, Course Director. **Lecture courses:** Computer Graphics, Advanced Rendering and Animation.

Profile: Dr. Kautz is a Reader in the Virtual Environments and Computer Graphics group. His research focuses on realistic computer graphics, with particular interest in real-time rendering, illumination computation, realistic materials, and image- and video-based rendering. Jan publishes regularly at the top international computer graphics conferences and received the Eurographics Young Researcher Award 2007. He is a regular program committee member for major computer graphics conferences and has links with the BBC graphics and visualization department.



Name: Dr. Niloy Mitra

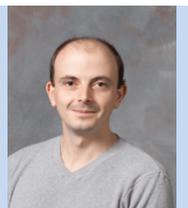
Roles: Personal tutor. **Lecture courses:** Image Processing.

Profile: Dr. Mitra is Senior Lecturer in the Virtual Environments and Computer Graphics group. His background is in geometry processing. His research interests include geometric modelling, shape analysis, scan alignment and symmetry and structure detection. Niloy regularly publishes in, and reviews for all the top computer graphics journals and conferences. In 2008, he received a Microsoft Outstanding Young Faculty Fellowship. He currently has active research projects in collaboration with other leading universities.

Name: Dr. Anthony Steed

Roles: Personal tutor. **Lecture courses:** Virtual Environments.

Profile: Prof. Steed is Professor in Virtual Environments and head of the Virtual Environments and Computer Graphics group. His specialist research areas are real-time interactive virtual environments, with particular interest in mixed-reality systems, large-scale models and collaboration between immersive facilities. Anthony is a member of the editorial board for the journal Presence and is a regular program committee member and cochair of the top international conferences on virtual reality and human-computer interaction. He has strong links with computer-games companies, such as EA, and other companies interested in virtual reality, computer graphics and visualization, such as the BBC.



Name: Dr. Tim Weyrich

Roles: Personal tutor. **Lecture courses:** Computational Photography and Capture

Profile: Dr. Weyrich is Senior Lecturer in the Virtual Environments and Computer Graphics group. His research interests are appearance acquisition, 3D reconstruction, point-based graphics, and interactive applications in art and cultural heritage preservation. His work covers a wide range of topics, including human skin appearance capture, graphics chip design, automated generation of bas-relief sculpture and the assembly of fractured Minoan frescoes. He is an active member of both the vision and graphics research communities, and his ongoing research concentrates on development and analysis of interactive applications in real-world scenarios, collaborating, for instance, with archaeologists and with skin biologists.

Selected Syllabi

Name of Course: **Mathematical Methods Algorithms and Implementations**

Aims: To understand analytical and numerical methods for image processing, graphics and image reconstruction. To provide a rigorous mathematical approach: in particular to define standard notations for consistent usage in other modules. To present relevant theories and results. To develop algorithmic approach from mathematical formulation through to hardware implications.

Syllabus:

- *Linear Algebra via Geometry:* Vectors; matrices; eigenvalues; kernel spaces; singular value decomposition; co-ordinate systems; orthogonalisation; lines; planes; rotation and translation
- *Probability and Estimation:* Forward probability; common probability distributions; Monte Carlo sampling; moments; inverse probability; Bayes Theorem; random variables; maximum likelihood estimation
- *Calculus :* Ordinary differential equations (complementary functions and particular integrals); partial differential equations (separation of variables)
- *Fourier Transforms:* Calculating Fourier series and transforms; interpreting Fourier series; Fast Fourier Transforms
- *Basic Algorithms:* Dynamic programming; sorting; tree searches
- *Programming in Matlab*

Name of Course: **Image Processing**

Aims: The first half of this course introduces the digital image, describes the main characteristics of monochrome images, how they are represented and how they differ from graphics objects. It covers algorithms for image manipulation, characterisation, segmentation and feature extraction. The second half proceeds to a more formal treatment of image filtering and the role and implications of Fourier space, and more advanced characterisation and feature detection techniques such as edge and corner detection, together with multiresolution methods, treatment of colour images and template matching.

Syllabus:

- *Introduction to the digital image:* Why digital images? The digital camera. Data types and 2d representation of digital images. Characteristics of grey-level digital images. Discrete sampling model. Quantisation. Noise processes. Image attributes.
- *Segmentation:* Thresholding and thresholding algorithms. Performance evaluation and ROC analysis. Connected components labelling. Region growing and region adjacency graph (RAG). Split and merge algorithms. Image Transformations. Grey level transformations. Histogram equalization.
- *Geometric transformations:* Affine transformations. Polynomial warps.
- *Morphological operations:* Erode and dilate as max and min operators on binary images. Open, close, thinning and other transforms. Medial axis transform. Introduction to grey-level morphology
- *Feature Characterisation:* Calculation of region properties. Moment features. Boundary coding. Fourier descriptors: Line descriptors from boundary coding and from moments.
- *Image filtering:* Linear and non-linear filtering operations. Image convolutions. Separable convolutions. Sub-sampling and interpolation as convolution operations.
- *Edge detection:* Alternative approaches. Edge enhancement by differentiation. Effect of noise, edge detection and Canny implementation. Edge detector performance evaluation.
- *Corner detection:* Image structure tensor. Relationship to image auto-correlation. Characterisation and Harris corner detector. Sub-pixel accuracy and performance evaluation
- *Colour images:* Representations of colour in digital images. Colour metrics. Pixel-wise (point) operations. Colour invariants and Finlayson colour constancy algorithm
- *Template matching:* Similarity and dissimilarity matching metrics. L2 metric and relationship to cross-correlation. Image search and multi-resolution algorithms. 2D object detection, recognition, location

Name of Course: Computer Graphics

Aims: To introduce the fundamental concepts of 3D computer graphics and give the students all the knowledge needed for creating an image of a virtual world from first principles.

Syllabus:

- *Introduction:* The painter's method.
- *Creating an image using ray tracing:* Ray casting using a simple camera. Local illumination. Global illumination with recursive ray tracing
- *Specifying a general camera:* World / image coordinates. Creation of an arbitrary camera. Ray tracing with an arbitrary camera.
- *Constructing a scene:* Definition of polyhedral. Scene hierarchy. Transformations of objects / rays. Other modelling techniques.
- *Acceleration Techniques:* Bounding volumes. Space subdivision.
- *From ray tracing to projecting polygons:* Graphics pipeline. Transforming the polygons to image space. Sutherland-Hodgman clipping. Weiler-Atherton clipping
- *Polygon rasterization/Visible surface determination:* Scan conversion. Z-buffer. Interpolated shading. Texture mapping. OpenGL. Back face culling.
- *Shadows:* Shadow volumes. Shadow buffer.
- *The nature of light:* Transport theory, Radiance, luminance, radiosity. The radiance equation.
- *Radiosity method:* Classical radiosity. Substructuring. Progressive refinement.
- *Parametric surfaces:* Bezier Curves. B-Splines Curves.

Name of Course: Machine Vision

Aims: The course addresses algorithms for automated computer vision. It focuses on building mathematical models of images and objects and using these to perform inference. Students will learn how to use these models to automatically find, segment and track objects in scenes, perform face recognition and build three-dimensional models from images.

Syllabus:

- *Two-dimensional visual geometry:* 2d transformation family. The homography. Estimating 2d transformations. Image panoramas.
- *Three dimensional image geometry:* The projective camera. Camera calibration. Recovering pose to a plane.
- *More than one camera:* The fundamental and essential matrices. Sparse stereo methods. Rectification. Building 3d models. Shape from silhouette.
- *Vision at a single pixel:* background subtraction and color segmentation problems. Parametric, non-parametric and semi-parametric techniques. Fitting models with hidden variables.
- *Connecting pixels:* Dynamic programming for stereo vision. Markov random fields. MCMC methods. Graph cuts.
- *Texture:* Texture synthesis, super-resolution and denoising, image inpainting. The epitome of an image.
- *Dense Object Recognition:* Modelling covariances of pixel regions. Factor analysis and principle components analysis.
- *Sparse Object Recognition:* Bag of words, latent dirichlet allocation, probabilistic latent semantic analysis
- *Face Recognition:* Probabilistic approaches to identity recognition. Face recognition in disparate viewing conditions.
- *Shape Analysis:* Point distribution models, active shape models, active appearance models.
- *Tracking:* The Kalman filter, the Condensation algorithm.

Name of Course: Virtual Environments

Aims: The purpose of this course is to introduce students to the main concepts and practical issues in constructing and understanding Virtual Environments, and how people respond to a VE experience. Given the background of the course teachers, the focus on the technical side will be more on the visual aspects of VEs. A central theme of the course will also be that the understanding of VEs can be best understood through the concepts of presence and shared presence.

- Syllabus:**
- *Introduction:* Virtual Environment Technology. Requirements. Applications.
 - *Interaction:* 3D Interaction tasks. Tracking. Input devices. System affordances. 3D Widgets
 - *Presence:* Immersion and presence. Meaning and utility of presence. Measuring presence.
 - *Displays:* 3D and Stereo Viewing: HMDs, CAVEs and desks. Graphics Architectures.
 - *Programming Virtual Environments:* Programming models. Simulation and animation . Programming for distribution ves.
 - *Devices:* Haptic devices. Sound simulation. Augmented reality.

Name of Course: Geometry of Vision

Aims: To introduce the generalisation of image processing to n-Dimensional data : volume data, scale space, time-series and vectorial data. To understand the principles of image processing in n-dimensions, time-series analysis and scale space, and to understand the relations between geometric objects and sampled images.

- Syllabus:**
- *Basic Image Operations:* Fourier Transforms. Convolution and Differentiation in Fourier Domain. Recursive Filters. Marching Square/ Cubes. Level Set Methods.
 - *Introduction to Differential Geometry:* Images as functions. Taylor Series expansion and the Koenderick jet. Properties of the local Hessian. Definition of extrema and saddle points. Ridges in n-dimensions. Image invariants up to fourth order.
 - *Curvature:* Contour curvature. Image curvature. 3D curvature and the Weingarten mapping. Gaussian and mean curvatures.
 - *Linear Scale Space:* Introduction and background. Formal properties. Gaussian kernels and their derivatives.
 - *Non-linear Scale Space:* Motivation. Edge-effected diffusion (Perona-Malik). Classification of Alvarez and Morel. Euclidian and Affine shortening flow. Numerical methods for computing scale spaces.
 - *Feature Space:* Introduction. Definitions of feature space. Clustering.
 - *Statistical Methods:* Linear and non-linear discriminant functions. Supervised learning Unsupervised learning .
 - *Bayesian and Information Theoretic Approaches:* Bayesian Image Restoration
 - *Markov Random Fields*
 - *Definitions of Entropy and Mutual information*
 - *Deconvolution with image priors* (statistical and structural)

Name of Course: Advanced modelling, graphics and visualization

Aims: To provide an understanding of advanced modelling, animation and rendering techniques. The students will know how to create complex models, how to animate a camera, rigid objects and characters, and finally how to make images look realistic by physically simulating the interaction of light with the objects in the scene.

- Syllabus:**
- *Modelling:* Curves and surfaces. Progressive meshes. Texture generation.
 - *Animation:* Rigid body, keyframe interpolation and motion capture. Forward and inverse kinematics. Dynamics, spring and mass, particle simulation, flock simulation.
 - *Rendering:* Physics of light. Photon tracing. Radiometrics. BRDF. Radiance equation. Deterministic techniques. Advanced radiosity and ray tracing. Montecarlo techniques. Distributed ray tracing. Path tracing. Particle tracing. Hardware rendering. Image-based rendering. Inverse illumination and digital image composition. Managing complex scenes.

Name of Course: Inverse Problems in Imaging

Aims: To introduce the concepts of optimisation, and appropriate mathematical and numerical tools. Applications in image processing and image reconstruction.

- Syllabus:**
- *Introduction:* Example problems. Data Fitting Concepts. Variational and Iterative Concepts.
 - *Mathematical Tools:* Linear Algebra. Solving Systems of Linear Equations. Over and Under Determined Problems. Eigen-Analysis and SVD. Preconditioning. Variational Methods: Calculus of Variation. Multivariate Derivatives. Frechet and Gateaux Derivatives. Regularisation. Tikhonov and Generalised Tikhonov. Non-Quadratic Regularisation. Non-Convex Regularisation.
 - *Numerical Tools:* Non-Gradient Methods. Simplex Method. Powell's Method. Descent Methods. Steepest Descent. Conjugate Gradients. Line Search. Newton Methods. Gauss Newton and Full Newton. TrustRegion and Globalisation. Quasi-Newton. Inexact Newton.
 - *Unconstrained Optimisation:* Least-Squares Problems. Linear Least Squares. Non-linear Least Squares. Non-Quadratic Problems. Poisson Likelihood. Kullback-Leibler Divergence.
 - *Regularisation:* Parameter Selection. Discrepancy Principles. The L-Curve Method. Generalised Cross-Validation. Constrained Optimisation Equality Constraints. Lagrangian Penalties. Inequality Constraints. Positivity Constraints. Upper and Lower Bounds. Active Sets. Primal-Dual Methods. Primal-Dual Interior Point Methods.
 - *Bayesian Approach:* Bayesian Priors and Penalty Functions. Maximum Likelihood and Maximum A Posteriori. Best Linear Unbiased Estimation. Expectation-Minimisation. Posterior Sampling. Confidence-Limits. Monte Carlo Markov Chain.
 - *Applications:* Image Deblurring. Deconvolution. Anisotropic Denoising. Linear Image Reconstruction. Tomographic Reconstruction. Reconstruction from Incomplete Data.
 - *Non-Linear Parameter Estimation:* General Concepts. Direct and Adjoint Differentiation.
 - *Other Approaches:* Simulated Annealing. Genetic Algorithms.

Name of Course: Computational Photography and Capture

Aims: The module is designed to be self-contained, introducing the theoretical and practical aspects of modern photography and capture algorithms to students with only limited mathematical background. The two primary aims are i) to introduce universal models of colour, computer-controlled cameras, lighting and shape capture, and ii) to motivate students to choose among the topics presented for either continuing study (for those considering MSc's and PhD's) or future careers in the fields of advanced imaging.

- Syllabus:**
- *Introduction:* Cameras, sensors and colour. Computational photography.
 - *Basic operations on photographs:* Blending and compositing. Background subtraction and matting. Warping, morphing, mosaics and panoramas. High-dynamic range imaging, tone mapping. Hybrid Images. Flash Photography. Stylised rendering using multi-flash.
 - *Image and Video Synthesis:* Image inpainting. Texture synthesis. TIP, Video Textures. Temporal sequence re-rendering. Speech animation. Controlled video sprites. Video-based rendering. Motion magnification. Non-photorealistic rendering and animation. Colourisation and colour Transfer.
 - *Advanced Image Editing:* Intrinsic images. Vectorising raster images. Poisson image editing. Seam carving. De-blurring, De-hazing. Coded Aperture imaging.
 - *Image-based Rendering:* Image-based modelling and photo editing. View-dependence. Light-dependence. The plenoptic function. Factored time-lapse video. Computational time-lapse Video. Video synopsis and indexing.
 - *Appearance Capture and Modelling:* Structured light acquisition (stripe, edge and phase-shift encoding). ShadowCuts. Photometric stereo. Dual photography. Separation of global and local reflectance. Image-based BRDF and BSSRDF measurements.

UCL Computer Science Department

University College London (UCL) is a multi-faculty higher education institution based and a constituent college of the University of London. It was founded in 1826 and was the first British university to admit students regardless of race, class, religion and gender. Today it has more than 25000 staff and students. UCL's main campus is in Bloomsbury in central London, close to Covent Garden, the British Museum and Oxford street.

UCL is a member of the G5 group of super-elite British universities. In 2009 it was ranked the 4th (Times Higher Education) and 22nd (Shanghai Jiao Tong) best university in the world and 2nd (Times Higher Education) and 3rd (Shanghai Jiao Tong) in Europe.

The computer science department at UCL has more than 60 academic staff and several hundred post-doctoral staff and research students. Other than vision and graphics, the department has particular strengths in networks, software engineering, machine learning and human computer interaction. Members of the department are also involved in the Centre for Medical Image Computing (a collaboration between medical physicists and computer scientists) and the Centre for Statistics and Machine Learning (a collaboration between statisticians and computer scientists). We particularly encourage CGVI students to choose projects that combine the core vision/graphics element with ideas from these other areas.

The department also houses a CAVE-like virtual reality facility. A CAVE is a room in which the user is presented with high-resolution stereo-pair images projected in real-time on 3 walls and the floor. When viewed through lightweight shutterglasses, the left/right stereo images are presented separately to the left and right eyes respectively, producing the illusion of 3D objects appearing both within and beyond the walls of the CAVE. The images are presented with reference to the users viewpoint, which is continuously updated via a head-tracking unit; thus even as the user moves around in a CAVE the environment displayed will always be perspective-correct.



The department has its own dedicated support team and operates its own extensive and regularly updated computer system. This system is a distributed workstation environment comprising of more than 500 workstations and servers connected through a series of 100/1000 Mbps, Ethernet-based, local area networks. Of these approximately 120 are currently available to undergraduate and graduate taught students. All machines run SUN Solaris, Linux and Windows XP.

CGVI Entrance Requirements

The MSc is designed for people who have a first or upper second class honours BSc (or equivalent) in a numerate discipline such as Computer Science, Mathematics, Engineering or the Physical Sciences. Industrial experience may compensate for lesser degrees or lack of technical qualification. Candidates do not require previous knowledge of graphics, vision or imaging, although some familiarity can be an advantage. Successful candidates will already have good programming skills in a language such as C, C++, Java and/or Matlab. Applicants should have good competence and familiarity with mathematics including vector and matrix algebra, calculus, geometry and probability and statistics. Familiarity with more advanced mathematics such as linear algebra, differential equations and Fourier transforms is an advantage. However, some revision material is available and early parts of the course cover and revise all these areas.

All students are required to demonstrate proficiency in English (UCL Level *Good*). For those whose first language is not English a number of qualifications are acceptable. A score of 100 in TOEFL together with a 24/30 in the reading and writing subtests and 23/30 in the listening and speaking subtests. A score of 7.0 in the IELTS examination with a minimum score of 6.0 in each of the four sub-tests is acceptable as well. These figures are absolute MINIMUM requirements on which we do not compromise, and we should make it clear that some students who have only these minimal qualifications have found it very hard to pass the course. In particular, there is one component in the research methods module which requires an extended written report that cannot be passed if the written English is not sufficiently clear. Moreover, the MSc project thesis is 20% of the final marks for the MSc, and is graded largely on clarity. A complete list of all acceptable English language qualifications is available from the Admissions and General Enquiries Office. Overseas students can seek further advice here: <http://www.ucl.ac.uk/prospective-students/international-students/>

Admissions

Applications may be submitted at any time, but the course begins at the end of September each year. We recommend early application to secure a place for the next academic year. You can apply online or download an application form from:

<http://www.ucl.ac.uk/admission/graduate-study/application-admission/>

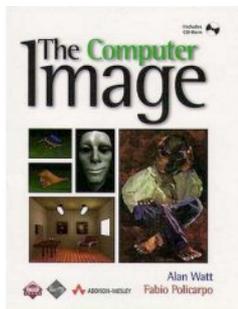
For further information about graduate study at UCL and the application process, see the general UCL information for prospective students:

<http://www.ucl.ac.uk/prospective-students/graduate-study/>

or contact the UCL registry on +44 (0) 207 380 3000. In particular, information about fees can be found at:

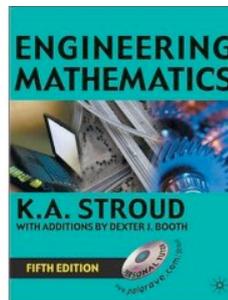
<http://www.ucl.ac.uk/prospective-students/graduate-study/fees-costs>

Recommended Reading



The Computer Image by Alan H. Watt and Fabio Policarpo

This is a nice general book on image processing, machine vision and computer graphics. It's somewhat outdated now but still provides a readable introduction to the course content.



Engineering Mathematics 5th ed. by K.A. Stroud

This book will help you prepare for the mathematical content of the course.

Contact

If you have any further questions about the course, please contact either the MSc CGVI administrator, Tracy Williams via advancedmsc-admissions@cs.ucl.ac.uk or the course director, Jan Kautz via j.kautz@cs.ucl.ac.uk.