

Computer Vision

Computer Science Tripos: 16 Lectures by J G Daugman

1. Overview. Goals of computer vision; why they are so difficult.
2. Image sensing, pixel arrays, CCD cameras. Image coding.
3. Biological visual mechanisms, from retina to primary cortex.
4. Mathematical operations for extracting structure from images.
5. Edge detection operators; the Laplacian and its zero-crossings.
6. Scale-space; multi-resolution. Wavelets as visual primitives.
7. Higher brain visual mechanisms; streaming; reciprocal feedback.
8. Texture, colour, stereo, and motion descriptors. Disambiguation.
9. Lambertian and specular surface properties. Reflectance maps.
10. Shape description. Codons; superquadrics and surface geometry.
11. Perceptual psychology and cognition. Vision as model-building.
12. Lessons from neurological trauma and deficits. Visual illusions.
13. Bayesian inference in vision. Classifiers; probabilistic methods.
14. Appearance-based versus model-based vision. Query languages.
15. Vision as a set of inverse problems. Regularisation.
16. Case study: face detection and recognition; face interpretation.



Syllabus for Computer Vision

Aims

The aims of this course are to introduce the principles, models and applications of computer vision, as well as some mechanisms used in biological visual systems that may inspire design of artificial ones. The course will cover: image formation, structure, and coding; edge and feature detection; neural operators for image analysis; texture, colour, stereo, motion; wavelet methods for visual coding and analysis; interpretation of surfaces, solids, and shapes; data fusion; probabilistic classifiers; visual inference and learning; and face recognition.

Lectures

- Goals of computer vision; why they are so difficult. How images are formed, and the ill-posed problem of making 3D inferences from them about objects and their properties.
- Image sensing, pixel arrays, CCD cameras, framegrabbers. Elementary operations on image arrays; coding and information measures.
- Biological visual mechanisms from retina to cortex. Photoreceptor sampling; receptive field profiles; spike trains; channels and pathways. Neural image encoding operators.
- Mathematical operators for extracting image structure. Finite differences and directional derivatives. Filters; convolution; correlation. 2D Fourier domain theorems.
- Edge detection operators; the information revealed by edges. The Laplacian operator and its zero-crossings. Logan's Theorem.
- Scale-space, multi-resolution representations, causality. Wavelets as visual primitives.
- Higher level visual operations in brain cortical areas. Multiple parallel mappings; streaming and divisions of labour; reciprocal feedback through the visual system.
- Texture, colour, stereo, and motion descriptors. Disambiguation and the achievement of invariances.
- Lambertian and specular surfaces. Reflectance maps. Image formation geometry. Discounting the illuminant when inferring 3D structure and surface properties.
- Shape representation. Inferring 3D shape from shading; surface geometry. Boundary descriptors; codons; superquadrics and the "2.5-Dimensional" sketch.
- Perceptual psychology and visual cognition. Vision as model-building and graphics in the brain. Learning to see.
- Lessons from neurological trauma and visual deficits. Visual illusions and what they may imply about how vision works.
- Bayesian inference in vision; knowledge-driven interpretations. Classifiers. Probabilistic methods in vision.
- Object-centred coordinates. Appearance-based *versus* volumetric model-based vision.
- Vision as a set of inverse problems; mathematical methods for solving them: energy minimisation, relaxation, regularisation.
- Approaches to face detection, face recognition, and facial interpretation.

Objectives

At the end of the course students should:

- understand visual processing from both “bottom-up” (data oriented) and “top-down” (goals oriented) perspectives
- be able to decompose visual tasks into sequences of image analysis operations, representations, specific algorithms, and inference principles
- understand the roles of image transformations and their invariances in pattern recognition and classification
- be able to analyse the robustness, brittleness, generalisability, and performance of different approaches in computer vision
- be able to describe key aspects of how biological visual systems encode, analyse, and represent visual information
- be able to think of ways in which biological visual strategies might be implemented in machine vision, despite the enormous differences in hardware
- understand in depth at least one important application domain, such as face recognition, detection, or interpretation

Recommended book

Shapiro, L. & Stockman, G. (2001). *Computer Vision*. Prentice Hall.

Other resources on-line

Annotated Computer Vision Bibliography:

<http://iris.usc.edu/Vision-Notes/bibliography/contents.html>

“The Computer Vision Homepage” (Carnegie Mellon University):

<http://www-2.cs.cmu.edu/~cil/vision.html>

Intel Open Source Computer Vision Library:

<http://www.intel.com/research/mrl/research/opencv/>

1 Overview. Goals of computer vision; why they are so difficult.

Computer vision seeks to generate intelligent and useful descriptions of visual scenes and sequences, and of the objects that populate them, by performing operations on the signals received from video cameras.

Some examples of computer vision applications and goals:

- automatic face recognition, and interpretation of expression
- visual guidance of autonomous vehicles
- automated medical image analysis, interpretation, and diagnosis
- robotic manufacturing: manipulation, grading, and assembly of parts
- OCR: recognition of printed or handwritten characters and words
- agricultural robots: visual grading and harvesting of produce
- smart offices: tracking of persons and objects; understanding gestures
- biometric-based visual identification of persons
- visually endowed robotic helpers
- security monitoring and alerting; detection of anomaly
- intelligent interpretive prostheses for the blind
- tracking of moving objects; collision avoidance; stereoscopic depth
- object-based (model-based) compression of video streams
- general scene understanding

In many respects, computer vision is an “AI-complete” problem: building general-purpose vision machines would entail, or require, solutions to most of the general goals of artificial intelligence. It would require finding ways of building flexible and robust visual representations of the world, maintaining and updating them, and interfacing them with attention, goals and plans.


Like other problems in AI, the challenge of vision can be described in terms of building a *signal-to-symbol converter*. The external world presents itself only as physical signals on sensory surfaces (such as videocamera, retina, microphone...), which explicitly express very little of the information required for intelligent understanding of the environment. These signals must be converted ultimately into symbolic representations whose manipulation allows the machine or organism to interact intelligently with the world.

Although vision seems like such an effortless and immediate faculty for humans and other animals, it has proven exceedingly difficult to automate. Some of the reasons for this include the following:

1. An image is a two-dimensional optical projection, but the world we wish to make sense of visually is three-dimensional. In this respect, vision is “*inverse optics*,” we need to invert the $3D \rightarrow 2D$ projection in order to recover world properties (object properties in space); but the $2D \rightarrow 3D$ inversion of such a projection is, strictly, mathematically impossible.

In another respect, vision is “*inverse graphics*,” graphics begins with a 3D world description (in terms of object and illuminant properties, viewpoint, etc.), and “merely” computes the resulting 2D image, with its occluded surfaces, shading and shadows, gradients, perspective, etc. Vision has to perform exactly the inverse of this process!

A classical and central problem in computer vision is face recognition. Humans perform this task effortlessly, rapidly, reliably, and unconsciously. (We don’t even know quite how we do it; like so many tasks for which our neural resources are so formidable, we have little “cognitive penetrance” or understanding of how we actually perform face recognition.) Consider these three facial images (from Pawan Sinha, MIT, 2002):



Which two pictures show the same person?

Most current computer algorithms select 1 and 2 as the same person, since those images are more similar than 1 and 3.


2. Very few visual tasks can be successfully performed in a purely data-driven way (“bottom-up” image analysis). Consider the next image example: the foxes are well camouflaged by their textured backgrounds; the foxes occlude each other; they appear in several different poses and perspective angles; etc. How can there possibly exist mathematical operators for such



an image that can:

- perform the figure-ground segmentation of the scene (into its objects and background)
- infer the 3D arrangements of objects from their mutual occlusions
- infer surface properties (texture, colour) from the 2D image statistics
- infer volumetric object properties from their 2D image projections
- and do all of this in “real time?” (This matters quite a lot in the natural world “red in tooth and claw,” since survival depends on it.)

Consider now the actual image data of a face, shown as a pixel array with luminance plotted as a function of (X,Y) pixel coordinates. Can you see the face in this image, or even segment the face from its background, let alone recognize the face? In this form, the image reveals both the complexity of the problem and the poverty of the data.



This “counsel of despair” can be given a more formal statement:


Most of the problems we need to solve in vision are *ill-posed*, in Hadamard’s sense that a *well-posed* problem must have the following set of properties:

- its solution exists;
- its solution is unique;
- its solution depends continuously on the data.


Clearly, few of the tasks we need to solve in vision are well-posed problems in Hadamard’s sense. Consider for example the problems of:

- inferring depth properties from an image
- inferring surface properties from image properties
- inferring colours in an illuminant-invariant manner
- inferring structure from motion, shading, texture, shadows, ...

- inferring a 3D shape unambiguously from a 2D line drawing:



- interpreting the mutual occlusions of objects, and stereo disparity
- recognizing a 3D object regardless of its rotations about its three axes in space (e.g. a chair seen from many different angles)
- understanding an object that has never been seen before:



- etc. ...

...but enough counsel of despair. Let us begin with understanding what an image array is.

2 Image sensing, pixel arrays, CCD cameras, image coding.

A CCD video camera contains a dense array of independent sensors, which convert incident photons focused by the lens onto each point into a charge proportional to the light energy there. The local charge is “coupled” (hence CCD) capacitively to allow a voltage ($V=Q/C$) to be read out in a sequence scanning the array. The number of pixels (picture elements) is typically about 442,000 within a grid that is $1/3'' \times 1/3''$ in size, so each pixel sensor is a mere 13 microns across (0.013 mm).

Spatial resolution of the image is thus determined both by the density of elements in the CCD array, and by the properties of the lens which is forming the image. Luminance resolution (the number of distinguishable grey levels) is determined by the number of bits per pixel resolved by the digitizer, and by the inherent signal-to-noise ratio of the CCD array.

Colour information arises (conceptually if not literally) from three separate CCD arrays preceded by different colour filters, or mutually embedded as sub-populations within a single CCD array. It is encoded into the camera output either as a high-frequency “chrominance burst” (to be separately demodulated and decoded); or else signaled on a separate channel (“luma” and “chroma” portions of an S signal); or else provided as three separate RGB colour channels (red, green, blue).

A framegrabber contains a high-speed analogue-to-digital converter (ADC) which discretizes this video signal into a byte stream. Standard pixel array formats are:

- NTSC (North American standard): 640 x 480 pixels, at 30 frames/second (actually there is an interlace of alternate lines scanned out at 60 “fields” per second)
- PAL (European, UK standard): 768 x 576 pixels, at 25 frames/second (actually alternate lines are scanned out at 50 “fields” per second)

Note what a vast quantity of data a video sequence contains! 768×576 pixels/frame \times 25 frames/second = 11 million pixels per second. Each pixel may be resolved to 8 bits in each of the three colour planes, hence 24×11 million = 264 million bits per second! How can we possibly cope with this data flux, let alone understand the objects and events creating such an image stream? ...

2.1 Image formats and sampling theory

Images are represented as rectangular arrays of numbers representing image intensities at particular locations. Each element of such an array is called a pixel, for picture element. A colour image may be represented in three separate such arrays called “colour planes,” containing red, green, and blue components as monochromatic images. An image with an oblique edge might look like:

0	0	0	1	1	0
0	0	1	2	10	0
0	1	2	17	23	5
0	3	36	70	50	10
1	10	50	90	47	12
17	23	80	98	85	30

There are many different image formats used for storing and transmitting images in compressed form, since raw images are large data structures that contain much redundancy (e.g. correlations between nearby pixels) and thus are highly compressible. Different formats are specialized for compressibility, manipulability, or the properties of printers and browsers. Some examples:

- **.jpeg** - ideal for variable compression of continuous colour images, with a “quality factor” (typically 75) that can be specified. Useful range of DCT compression goes from 100:1 (“lossy”) to about 10:1 (almost lossless).
- **.mpeg** - a stream-oriented, compressive encoding scheme used mainly for video (but also multimedia). Individual image frames are **.jpeg** compressed, but an equal amount of redundancy is removed temporally by inter-frame predictive coding and interpolation.
- **.gif** - ideal for sparse binarized images. Only 8-bit colour. Very compressive and favoured for web-browsers and other bandwidth-limited media.
- **.tiff** - a complex, full-colour format (up to 24 bits per pixel with 3 colour planes). Favoured for scanners, quality printing and publishing.
- **.bmp** - a non-compressive bit-mapped format in which individual pixel values can easily be extracted.

In addition there are varieties of colour coordinates used for “colour separation,” such as HSI (Hue, Saturation, Intensity), or RGB (Red, Green, Blue), CMY, etc. But regardless of the sensor properties and coding format used, ultimately the image data must be represented numerically pixel by pixel. Typically this involves the conversion (e.g. by a tool such as **xv**) of the various compressed formats into **.bmp**, with header files for format and dimensions.

The total number of independent pixels in an image array determines the spatial resolution of the image. Independent of this is the grey-scale (or colour) resolution of the image, which is determined by the number of bits of information specified for each pixel. These separate dimensions are illustrated in the following family of images, showing the effects of differing quantization accuracies for spatial and luminance information.

It is typical for a monochromatic (“black & white”) image to have resolution of 8 bits/pixel. This creates 256 different possible intensity values for each pixel, from black (0) to white (255), with all shades of grey in between. A full-colour image may be quantized to this depth in each of the three colour planes, requiring a total of 24 bits per pixel. However, it is common to represent colour more coarsely or even to combine luminance and chrominance information in such a way that their *total* information is only 8 or 12 bits/pixel.

Because quantized image information is thus fundamentally discrete, the operations from calculus which we might want to perform on an image, like differentiation (to find edges) or integration (to perform convolutions or transforms), must be done in their discrete forms. The discrete form of a derivative is a *finite difference*. The discrete form of an integral is a (suitably normalized) *summation*. However, for the sake of conceptual familiarity, it is still commonplace in computer vision to represent such operations using their usual notations from continuous mathematics, with the understanding that the operation itself (as with everything else in Computer Science!) is of course discrete.

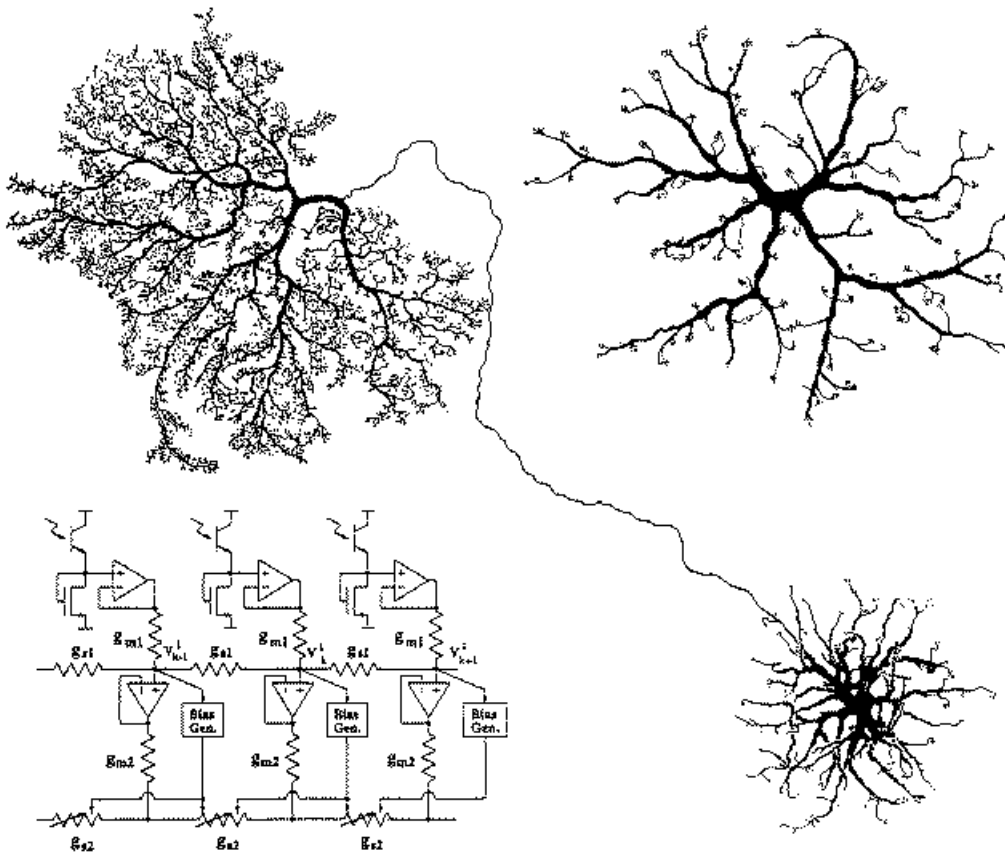
The discreteness of image arrays imposes an upper limit on the amount of information they can contain. One way to describe this is by the total bit count, but this does not relate to the optical properties of image information. A better way is through Nyquist’s Theorem, which tells us that the highest *spatial frequency* component of information contained within the image is equal to one-half the sampling density of the pixel array. (Intuitively, this is because at least two samples are required to signify a single cycle of a sinewave: its peak and its trough.) Thus, a pixel array containing 640 columns can represent spatial frequency components of image structure no higher than 320 cycles/image. For the same reason, if image frames are sampled in time at the rate of 30 per second, then the highest *temporal frequency* component of information contained in the image sequence is 15 Hertz.

3 Biological visual mechanisms, from retina to primary cortex.

A strategy that has long inspired researchers in Computer Vision, whether they work on low-level problems (such as sensor design, image coding, and feature extraction), or high-level problems (such as pattern recognition, inference, and visual learning), is:

Neurobiological Visual Principles \Rightarrow Machine Vision

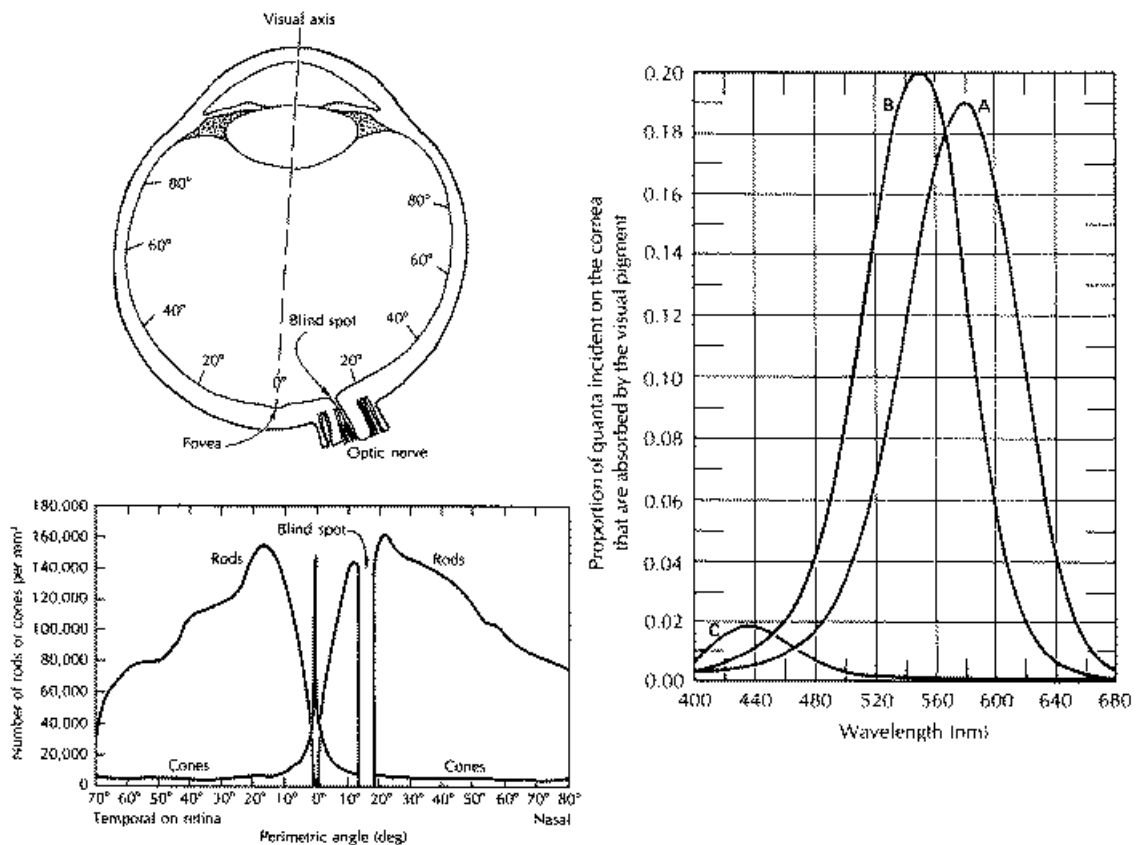
The structure of biological nervous tissue and the nature of events that occur in it are utterly different from those found in computing hardware. Yet since the only general-purpose visual systems that exist today are the biological ones, let us learn what we can from “wetware.”



Neurons are sluggish but richly interconnected devices having both analogue and discrete aspects. Fundamentally they consist of an enclosing membrane that can separate electrical charge (hence there is generally a voltage difference between inside and out). The membrane is a lipid bilayer that has a capacitance of about $10,000 \mu\text{Farads}/\text{cm}^2$, and it also has pores that are differentially selective to different ions (mainly Na^+ , K^+ , and Cl^-). These ion species enter or leave a neuron through protein pores studding its lipid membrane, acting as conductances (hence as resistors). The resistors for Na^+ and K^+ have the further crucial property that their resistance is not constant, but voltage-dependent. Hence as more positive ions (Na^+) flow into the neuron, the voltage becomes more positive on the inside, and this further reduces the membrane’s resistance to Na^+ , allowing still more to enter. This catastrophic breakdown in resistance to Na^+ constitutes a nerve impulse. Within about a msec a slower but opposite effect involving K^+ takes over, eventually restoring the original voltage. Following a short refractory period of about 2 msec during which ions are actively pumped back


in opposite directions to reach their original electro-osmotic equilibrium concentrations, the neuron is ready for action again. Meanwhile, the impulse thus generated propagates down the axon, at a speed of about 100 m/sec. This signalling pulse can be described as discrete, but the antecedent summations of current flows into the neuron (from various influences by other neurons) which caused the catastrophic impulse are fundamentally analogue events.

Overall, the human brain contains about 100 billion neurons (10^{11}). On average each neuron may have connections with about 1,000 to 10,000 others, and so the total number of synapses (= junctions between neurons) in the brain is a staggering 10^{15} . Yet balanced against this massive connectivity, is the surprising sluggishness of neurons: as indicated above, the time course of nerve impulse generation prevents “clocking” of nerve pulses any faster than about 300 Hz. Neural activity is fundamentally asynchronous: there is no master clock on whose edges the events occur. A further contrast with computing systems is that it is rarely possible to distinguish between processing and communications, as we do in computing. In the brain, there are just impulses implementing both, by exchange of signals amongst neurons. It is not so much a hierarchical architecture as a parallel one, with reciprocal connections amongst different areas. About $2/3^{rd}s$ of the brain receives visual input; we are quite fundamentally visual creatures. There are some 30 known different visual areas, of which the primary visual cortex in the occipital lobe at the back of the brain has been the most extensively studied.




The mammalian eye is formed from a collapsed ventricle of the brain. The retina is about 1 mm thick and contains about 120 million light-sensitive photoreceptors, of which only 6 million are cones (in 3 wavelength-selective classes nominally red, blue, and green) and the vast remainder are rods which do not discriminate in wavelength. The visible spectrum of light consists of wavelengths in the range of 400nm - 700nm. Rods are specialised for much lower light intensities than cones; they subserve our “night vision” (hence the absence of perceived colour at night), and they pool together their responses (hence their much poorer spatial resolution). Cones exist primarily near the fovea, in about the central 20° (see diagram), where their responses remain individual and thus they detect with high spatial resolution. But cone

light sensitivity is much less than rods, functioning only at higher light levels, and so we really have a dual system with two barely overlapping dynamic ranges. The total dynamic range of human vision (range of light intensities that can be processed) is a staggering 10^{11} to 1. At the low end, we can reliably “see” individual photons (i.e. reliably have a visual sensation when at most a few photons reach the retina in a burst).




Rods and cones are distributed across the retina in jointly embedded hexagonal lattices but with varying relative densities, depending on eccentricity (distance from the fovea, measured in degrees). The hexagonal lattices are imperfect (incoherent rather than crystalline), which is believed to help prevent aliasing of high resolution information.



The retina is a multi-layered structure, containing 3 nuclear layers (of neurons) plus 2 plexiform layers (for interconnections amongst the neurons). Paradoxically, the photoreceptors are at the back, so light must first travel through all of the rest of the retina before being absorbed

by the pigments in the rods and cones. There are basically two directions of signal flows in the retina: longitudinal (photoreceptors → bipolar cells → ganglion cells); and lateral (via horizontal cells in the outer plexiform layer, and amacrine cells in the inner plexiform layer).



In only a very crude sense can one describe the retina as an “image capture” device like a camera, having analogue input phototransducers that convert photons into voltage changes, and discrete output devices that send pulses down the optic nerve. This simple view is quickly discarded by recognising that there are 120 million “input channels” (the photoreceptors, similar in a sense to pixels), but only 1 million “output channels” (the axons of the ganglion cells which constitute the optic nerve). Clearly the retina is already doing a lot of processing of the image, and it sends its coded results to the brain: not merely a raw converted image array. The retina *is a part* of the brain.

The nature of retinal signal processing might be summarised as:

- image sampling by photoreceptor transducers, with pooling of signals from rods
- spatial centre-surround comparisons implemented by bipolar cells (direct central input from photoreceptors, minus surround inhibition via horizontal cells)
- temporal differentiation by amacrine cells, subserving motion sensitivity
- separate coding of “sustained” versus “transient” image information by different classes of ganglion cells (large receptive fields \Leftrightarrow transient; small fields \Leftrightarrow sustained)
- initial colour separation by “opponent processing” channels (yellow vs blue; red vs green) coupled sometimes with spatial opponency (on-centre, off-surround)
- generation of nerve impulse spikes in a parallel temporal modulation code on the 1 million fibres of the optic nerve from each eye (= 2nd Cranial Nerve)

There is both convergence (fan-in) and divergence (fan-out) of signals through the retina:





Fig. 5. Actual convergence and divergence through successive arrays along the cone bipolar pathway to the ON-beta ganglion cell in the area centralis. In this circuit connections between successive arrays are much narrower than their potential connections. Text describes structure-function relationships of this circuit.



3.1 Receptive field structure in the retina


The spatial structuring of excitatory and inhibitory influences amongst neurons in the retina gives them their properties as image operators. Similarly for the temporal structure of their interactions. In both space and time, retinal neurons can thus be described as filters; and to the extent that they act as linear devices (having the properties of proportionality and superposition of responses to components of stimuli), their behaviour can be fully understood (and even predicted for arbitrary images) through Fourier analysis and the other tools of linear systems analysis. An important aspect of retinal receptive fields – as distinct from those found in most neurons of the visual cortex – is that their spatial structure is isotropic, or circularly symmetric, rather than oriented.

- Photoreceptors respond to light by hyperpolarising (the voltage across the cell membrane becomes more negative inside, for vertebrates; the opposite is true for invertebrates). Their “receptive field” is just their own cross-section for absorbing light, a small disk about $3\ \mu$ in diameter on the human retina, about a minute of visual arc.
- Horizontal cells pool together the responses from large numbers of photoreceptors within a local area. With these “surround” signals, they inhibit bipolar cells (hence the name).
- Bipolar cells are the first to have a “centre-surround” receptive field structure: their response to light in a central disk is opposite from their response to light in the local surrounding area. Field boundaries are circular and roughly concentric (i.e. annular).
- Amacrine cells are “on-off” in temporal, as opposed to spatial, terms.
- Ganglion cells combine these spatial and temporal response properties and thus serve as integro-differential image operators with specific scales and time constants. Moreover they convert their responses to impulses in a spike frequency code, traveling down their axons which are the fibres of the optic nerve to the thalamus and thence on to the primary visual cortex in the brain.

3.2 Visual cortical architecture and receptive field structure


The optic nerve from each eye splits into two halves at the optic chiasm, each portion continuing on to only one of the two cerebral hemispheres of the brain. The optic nerve portion containing signals from the nasal half of each retina crosses over to project only to the contralateral (opposite) brain hemisphere; whereas the optic nerve portion bearing signals from the temporal half of each eye projects only to the ipsilateral (same side) brain hemisphere. Since the optical image on each retina is inverted, this means that the left-half of the visual world (relative to the point of gaze fixation) is directly “seen” only by the right brain; and the right-half of the visual world only by the left brain. It is almost interesting to ask why we don’t see some kind of seam going down the middle... (Ultimately the two brain hemispheres share all of their information via a massive connecting bundle of 500 million commissural fibres called the corpus callosum.)


The optic nerve projections to each visual cortex pass first to a 6-layered structure called the lateral geniculate nucleus (LGN), in a polysensory organ of the midbrain called the thalamus. It is an intriguing fact that this so-called “relay station” actually receives 3 times more descending (efferent) fibres projecting back down from the cortex, as it does ascending (afferent) fibres from the eyes. Could it be that this confluence compares cortical feedback representing hypotheses about the visual scene, with the incoming retinal data in a kind of predictive coding or hypothesis testing operation? Several scientists have proposed that “vision is graphics” (i.e. what we see is really our own internally generated 3D graphics, modelled to fit the 2D retinal data, with the model testing and updating occurring here in the thalamus).



Schematic of the central visual pathway in the human. (from Popper and Eccles, 1977)

The right-eye and left-eye innervations from each LGN to the primary visual cortex in the occipital lobe of that hemisphere are inter-woven into “slabs,” or columns, in which neurons receive input primarily from just one of the eyes. These ocular dominance columns have a cycle of about 1 mm and resemble fingerprints, as seen in the following figures. Clearly each hemisphere is trying to integrate together the signals from the two eyes in a way suitable for stereoscopic vision, by computing the relative retinal disparities of corresponding points in the two images. The disparities reflect the relative positions of the points in depth, as we will study later with stereoscopic visual algorithms.



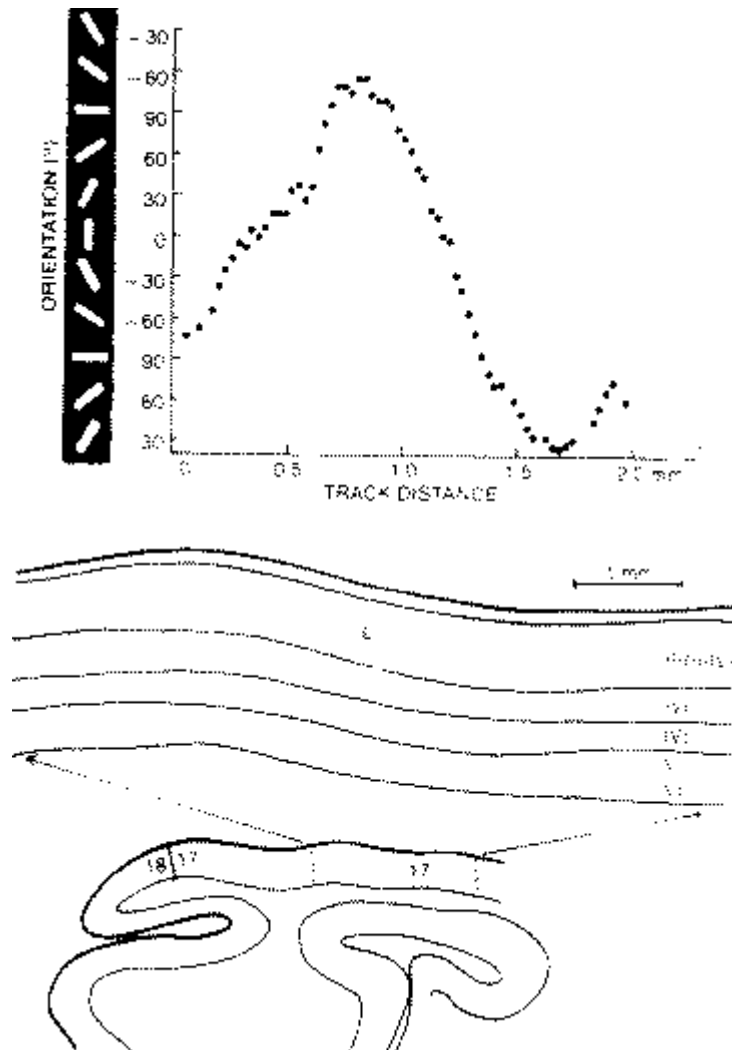


RELATION BETWEEN OCULAR DOMINANCE AND ORIENTATION COLUMNS. Scheme in which the ocular dominance and orientation columns run at right angles to each other. An example of a complex cell is shown in an upper layer, receiving its inputs from two simple cells that lie in two neighboring ocular dominance columns, but share the same orientation column. (From Hubel and Wiesel, 1972)




ANATOMICAL CONFIRMATION of ocular-dominance columns came from various staining methods and from axonal-transport autoradiographs such as those shown in color on page 85. This composite autoradiograph visualizing the pattern over an area some 10 millimeters wide was made by cutting out and pasting together the regions representing layer IV in a number of parallel sections; the one in bottom illustration on page 85 and others at different depths.

Orthogonal to the ocular dominance columns in the cortical architecture, there runs a finer scale sequence of orientation columns. Neurons in each such column respond only to image structures that have a certain preferred orientation (such as bars or edges). The columns form a regular sequence of systematically changing preferred orientations. This is one of the most crystalline properties seen in visual cortical architecture:



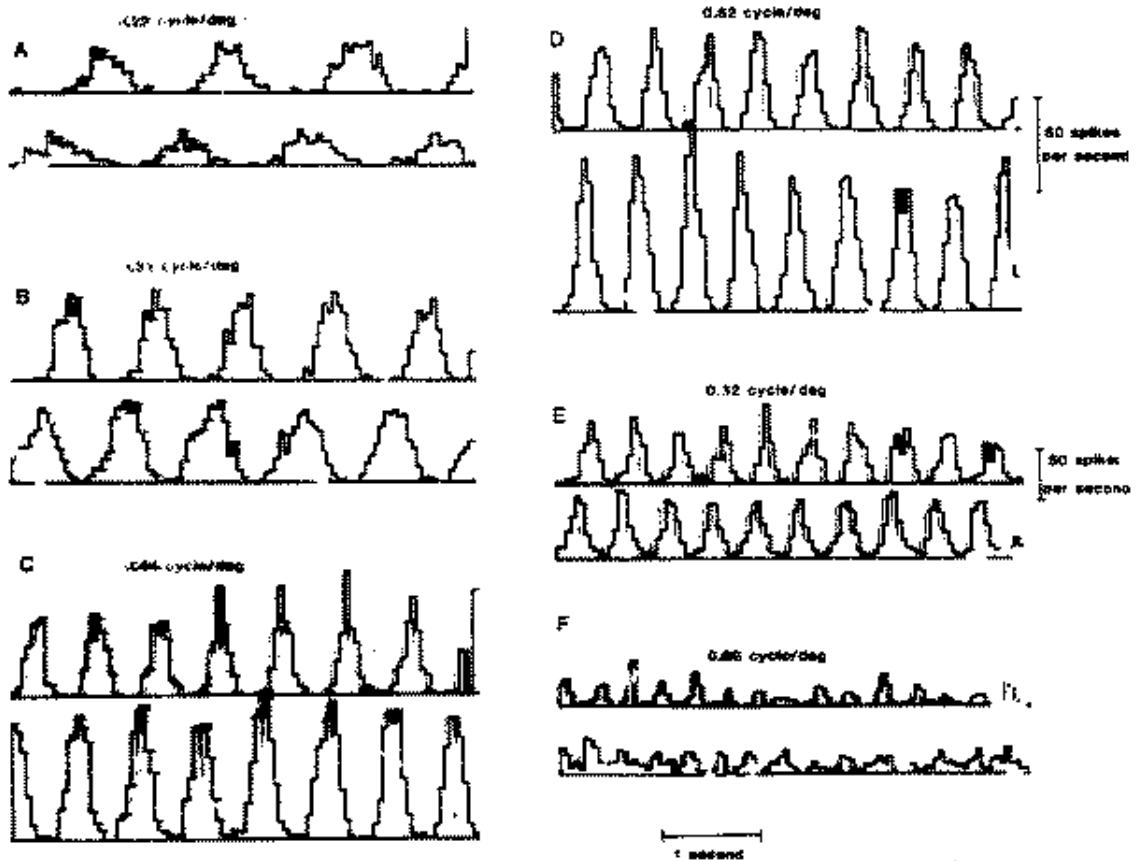
When individual neurons in the visual cortex are probed with microelectrodes during light stimulation of the retina, their functional properties are revealed by demarcating the region of visual space over which they respond (as indicated by a change in their firing rate). Areas where they are excited by light are indicated by + marks; areas where light inhibits them are indicated by - marks. Their plotted receptive fields then seem to reveal 5 main spatial “degrees of freedom:”

1. Position of their receptive field in visual space, both horizontally...
2. ...and vertically;
3. Size of their receptive field;
4. Orientation of the boundaries between excitatory and inhibitory regions;
5. Phase, or symmetry of the receptive field (bipartite or tripartite types).




The phase variable is particularly revealing. By recording from adjacent pairs of neurons simultaneously, with a kind of “double-barrelled” micro-electrode, it was found that neurons having the same receptive field location, the same field size and the same orientation preference, actually had a quadrature phase relationship. Adjacent neurons would form pairs whose modulated receptive field structure showed a 90° spatial phase offset. Several examples of such quadrature pairs of cortical visual neurons are shown in the following spike histograms recorded in response to a drifting sinusoidal luminance grating.


Finally, by plotting the actual amount by which a neuron is excited or inhibited by light, as a function of the coordinates of the stimulus within its receptive field, we obtain a 2D function called its receptive field profile. These turn out, for about 97% of the neurons, to be very closely described as 2D Gabor wavelets (or phasors). Some examples of empirically measured profiles are shown in the top row of the lower figure; the ideal theoretical form of such a wavelet (which we will define later) is shown in the middle row; and the difference between these two functions in the bottom row; the differences are nearly nil and statistically insignificant. So, it appears that the visual cortex of the brain evolved a knowledge of the valuable properties of such wavelets for purposes of image coding and analysis!



2D Receptive Field Profiles



Fitted 2D Gabor Phasors



Residuals



4 Mathematical operations for extracting structure from images.

Most image processing and feature encoding operations are related at least indirectly to the theory of Fourier Analysis.

Even if the operations never actually require computing a Fourier transform, their underlying principles and concepts (such as scale; edge or motion energy; filtering; directional derivative; textural signature; statistical structure; etc.) should be understood at least partially in “spectral” (i.e. Fourier) terms.

In addition to this explanatory role, Fourier analysis can sometimes be used to construct useful visual representations that are invariant under translation (change in position), rotation, and dilation (change in size). This is therefore the representation underlying certain pattern classification and recognition applications, such as optical character recognition (OCR).

Finally, even many operations in pattern recognition that might not seem related in any way to Fourier analysis, such as computing correlations, convolutions, derivatives, differential equations, and diffusions, are much more easily implemented in the Fourier domain. (Powerful algorithms like the FFT make it easy to go back and forth rapidly between the image and Fourier domains).

For all of these reasons, we will review some principles and techniques of Fourier analysis with a view to understanding some of the basic operations in computer vision. Applications include edge detection operators, analysis of motion, texture descriptors, and wavelet-based feature detectors.

Consider an image as a greyscale luminance distribution over the (x, y) plane: a real-valued (indeed, a positive-valued) two-dimensional function $f(x, y)$.

Any image can be represented by a linear combination of basis functions:

$$f(x, y) = \sum_k a_k \Psi_k(x, y) \quad (1)$$

where many possible choices are available for the expansion basis functions $\Psi_k(x, y)$. In the case of Fourier expansions in two dimensions, the basis functions are the bivariate complex exponentials:

$$\Psi_k(x, y) = \exp(i(\mu_k x + \nu_k y)) \quad (2)$$

where the complex constant $i = \sqrt{-1}$. A complex exponential contains both a real part and an imaginary part, both of which are simple (real-valued) harmonic functions:

$$\exp(i\theta) = \cos(\theta) + i \sin(\theta) \quad (3)$$

which you can easily confirm by the power-series that define the transcendental functions such as \exp , \cos , and \sin :

$$\exp(\theta) = 1 + \frac{\theta}{1!} + \frac{\theta^2}{2!} + \frac{\theta^3}{3!} + \cdots + \frac{\theta^n}{n!} + \cdots, \quad (4)$$

$$\cos(\theta) = 1 - \frac{\theta^2}{2!} + \frac{\theta^4}{4!} - \frac{\theta^6}{6!} + \cdots, \quad (5)$$

$$\sin(\theta) = \theta - \frac{\theta^3}{3!} + \frac{\theta^5}{5!} - \frac{\theta^7}{7!} + \cdots, \quad (6)$$

(It has been said that the most remarkable and far-reaching relationship in all of mathematics is the simple yet counterintuitive “Euler Relation” implied by Eq (3) above: $e^{i\pi} + 1 = 0$, which also contains the five most important mathematical constants, and symbolizes the subject of harmonic analysis.)

Fourier Analysis computes the coefficients a_k that yield an expansion of the image $f(x, y)$ in terms of complex exponentials:

$$f(x, y) = \sum_k a_k \exp(i(\mu_k x + \nu_k y)) \quad (7)$$

where the parameters μ_k and ν_k define the coordinates of the 2D Fourier domain. These (μ_k, ν_k) coordinates are called vector spatial frequencies, and the array of them must span the (μ, ν) Fourier plane in a uniform cartesian lattice.

It is often useful to think of the (μ, ν) Fourier plane as resolved into polar coordinates, where $\omega = \sqrt{\mu^2 + \nu^2}$ is (scalar) spatial frequency and $\phi = \tan^{-1}(\nu/\mu)$ is (scalar) orientation.

Each Fourier coefficient a_k is computed as the orthonormal projection of the entire image $f(x, y)$ onto the conjugate Fourier component $\exp(-i(\mu_k x + \nu_k y))$ associated with that coefficient:

$$a_k = \int_X \int_Y \exp(-i(\mu_k x + \nu_k y)) f(x, y) dx dy \quad (8)$$

Note that these computed Fourier coefficients a_k are complex-valued. To get a complete representation in the 2D Fourier domain for an image with $n \times n$ pixels, the number of (μ_k, ν_k) vector frequency components whose associated coefficients a_k must be computed is also $n \times n$.

4.1 Useful Theorems of 2D Fourier Analysis

Many important steps in computer vision such as feature extraction and invariant pattern recognition depend at least partly on a small set of Fourier

theorems. We will review some main ones here, together with their direct consequences for practical computer vision applications. In every case, the input image is denoted $f(x, y)$, and its 2D Fourier transform (given by the set of computed coefficients a_k spanning the Fourier plane) is denoted by $F(\mu, \nu)$.

Shift Theorem: Shifting the original pattern in (x, y) by some 2D displacement (α, β) merely multiplies its 2DFT by $\exp(-i(\alpha\mu + \beta\nu))$. Thus the 2DFT of the shifted pattern $f(x - \alpha, y - \beta)$ is: $F(\mu, \nu) \exp(-i(\alpha\mu + \beta\nu))$.

Practical Application: The power spectrum of any isolated pattern is thus translation-invariant: it does not depend on where the pattern is located within the image, and so you don't have to find it first. The power spectrum is defined as the product of the pattern's 2DFT, $F(\mu, \nu)$, times its complex conjugate, $F^*(\mu, \nu)$, which just requires that the sign $(-)$ of the imaginary part of $F(\mu, \nu)$ gets reversed. You can easily see that the power spectrum of the shifted pattern $f(x - \alpha, y - \beta)$, namely:

$$\exp(-i(\alpha\mu + \beta\nu))F(\mu, \nu) \exp(i(\alpha\mu + \beta\nu))F^*(\mu, \nu)$$

is equal to the power spectrum of the original unshifted pattern, namely: $F(\mu, \nu)F^*(\mu, \nu)$. Thus the power spectrum is translation-invariant.

Similarity Theorem: If the size of the original pattern $f(x, y)$ changes (shrinks/expands), say by a factor α in the x -direction, and by a factor β in the y -direction, becoming $f(\alpha x, \beta y)$, then the 2DFT of the pattern, $F(\mu, \nu)$, also changes (expands/shrinks) by the reciprocal of those factors and with similarly scaled amplitude. It becomes: $\frac{1}{|\alpha\beta|}F(\frac{\mu}{\alpha}, \frac{\nu}{\beta})$.

Rotation Theorem: If the original pattern $f(x, y)$ rotates through some angle θ , becoming $f(x \cos(\theta) + y \sin(\theta), -x \sin(\theta) + y \cos(\theta))$, then its 2DFT $F(\mu, \nu)$ also just rotates through the same angle. It becomes: $F(\mu \cos(\theta) + \nu \sin(\theta), -\mu \sin(\theta) + \nu \cos(\theta))$.

Practical Application: Size- and orientation-invariant pattern representations can be constructed by these relationships. Specifically, if the Fourier domain (μ, ν) is now mapped into log-polar coordinates (r, θ) where $r = \log(\sqrt{\mu^2 + \nu^2})$ and $\theta = \tan^{-1}(\nu/\mu)$, then any dilation (size change) in the original pattern becomes simply a translation along the r -coordinate; and any rotation of the original pattern becomes simply a translation along the orthogonal θ -coordinate in this log-polar Fourier domain. But we saw earlier that translations are made immaterial by taking a power spectrum, and so these effects of dilation and rotation of the pattern are eliminated in such a representation.

Combined with the translation-invariant property of the power spectrum, we now see how it becomes possible to represent patterns in a manner that is independent of their position in the image, their orientation, and their size (i.e. the Poincaré group of transformations) These principles are routinely exploited in machine optical character recognition; in military recognition of aircraft profiles; and in “optical computing” generally.

Convolution Theorem: Let function $f(x, y)$ have 2DFT $F(\mu, \nu)$, and let function $g(x, y)$ have 2DFT $G(\mu, \nu)$. The convolution of $f(x, y)$ with $g(x, y)$, which is denoted $f * g$, combines these two functions to generate a third function $h(x, y)$, whose value at location (x, y) is equal to the integral of the product of functions f and g after one is flipped and undergoes a relative shift by amount (x, y) :

$$h(x, y) = \int_{\alpha} \int_{\beta} f(\alpha, \beta) g(x - \alpha, y - \beta) d\alpha d\beta \quad (9)$$

Thus, convolution is a way of combining two functions, in a sense using each one to blur the other, making all possible relative shifts between the two functions when computing the integral of their product to obtain the output as a 2D function of these amounts of shift.

Convolution is extremely important in vision because it is the basis for filtering. It is also the essential neural operation in the brain’s visual cortex, where each neuron’s receptive field profile is convolved with the retinal image. In the above integral definition, if the minus (–) signs were simply replaced with (+) signs, the new expression would be the correlation integral.

The Convolution Theorem states that convolving two functions $f(x, y)$ and $g(x, y)$ together in the image domain, simply multiplies their two 2DFT’s together in the 2D Fourier domain:

$$H(\mu, \nu) = F(\mu, \nu) G(\mu, \nu) \quad (10)$$

where $H(\mu, \nu)$ is the 2DFT of the desired result $h(x, y)$.

This is extremely useful as it is much easier just to multiply two functions $F(\mu, \nu)$ and $G(\mu, \nu)$ together, to obtain $H(\mu, \nu)$, than to have to convolve $f(x, y)$ and $g(x, y)$ together (if the kernel is larger than tiny, say larger than about 5 x 5) to obtain $h(x, y)$. Of course, exploiting the Convolution Theorem means going into the 2D Fourier Domain and computing the 2DFT’s of $f(x, y)$ and $g(x, y)$, and then performing yet another

(inverse) FFT in order to recover $h(x, y)$ from the resulting $H(\mu, \nu)$. But with available powerful and fast 2D-FFT algorithms, this is very efficient.

Practical Application: Filtering. The starting-point of all feature extraction and image understanding operations is the filtering of an image $f(x, y)$ with some set of filters $g_k(x, y)$. Filtering is a linear operation implemented by the convolution of an image $f(x, y)$ with filter kernel(s) $g_k(x, y)$. The resulting output “image” $h_k(x, y)$ then normally undergoes non-linear operations of various kinds for image segmentation, motion detection, texture analysis, pattern recognition, and object classification.

The 2D discrete convolution of an image array with a 2D filter kernel can be represented algebraically in the following form, where the earlier continuous integrals have now been replaced by discrete summations:

$$\text{result}(i, j) = \sum_m \sum_n \text{kernel}(m, n) \cdot \text{image}(i - m, j - n)$$

Simple C program for performing image convolutions

In the following simple example, the array **image** is being convolved with the (typically much smaller) array **kernel**, in order to generate a new image array **result** as the output of the convolution. (Problems with array boundaries have been ignored here for simplicity.) Discrete convolution such as illustrated here is the key operation for all image processing and front-end stages of computer vision.

```
int  i, j, m, n, sum, image[iend][jend],
    kernel[mend][nend], result [iend][jend];

for (i = mend; i < iend; i++) {
    for (j = nend; j < jend; j++) {
        sum = 0;
        for ( m = 0; m < mend; m++) {
            for ( n = 0; n < nend; n++ ) {
                sum += kernel[m][n] * image[i-m][j-n];
            }
        }
        result[i][j] = sum/(mend*nend);
    }
}
```

If we chose to implement the convolution in the Fourier domain because the kernel array was large, then of the four nested **for loops** in the C code

above, the inner two loops would be entirely eliminated. Instead, the only operation inside the outer two **for loops** would be a multiplication:

$$\text{Result}[i][j] = \text{Kernel}[i][j] * \text{Image}[i][j];$$

but the program would have to be preceded by FFTs of `kernel[i][j]` (trivial) and of `image[i][j]`, and followed by an FFT of `Result[i][j]`. Since the complexity of a 2D FFT is on the order of $n^2 \log_2(n)$ where n^2 is the number of pixels, plus n^2 multiplications in the nested two loops, the total complexity of the Fourier approach is $n^2(2 \log_2(n) + 1)$. In contrast, the number of multiplications in the explicit convolution above (not including all the array-addressing) is `iend*jend*mend*nend` (note that `iend*jend = n^2`). Hence you can calculate that the trade-off point occurs when the convolution kernel size `mend*nend` is about $\approx 2(\log_2(n) + 1)$: a very small convolution kernel indeed, roughly 5 x 5 for a 512 x 512 image. For convolutions larger than this tiny one, the Fourier approach is faster.

Differentiation Theorem: Computing the derivatives of an image $f(x, y)$ is equivalent to multiplying its 2DFT, $F(\mu, \nu)$, by the corresponding frequency coordinate raised to a power equal to the order of differentiation:

$$\left(\frac{d}{dx}\right)^m \left(\frac{d}{dy}\right)^n f(x, y) \xrightarrow{2DFT} (i\mu)^m (i\nu)^n F(\mu, \nu) \quad (11)$$

A particularly useful implication of this theorem is that isotropic differentiation, which treats all directions equally (for which the lowest possible order of differentiation is 2nd-order, known as the Laplacian operator ∇^2) is equivalent simply to multiplying the 2DFT of the image by a paraboloid:

$$\nabla^2 f(x, y) \equiv \left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}\right) f(x, y) \xrightarrow{2DFT} -(\mu^2 + \nu^2) F(\mu, \nu) \quad (12)$$

Practical Application: Multi-Resolution Edge Detection.


5 Edge detection operators; the information revealed by edges.

Computer vision applications invariably begin with *edge detection*, be the edges straight, curvilinear, or closed boundary contours. There are several reasons why edges are important, and why detecting the edges in a scene can be regarded as an elementary form of constructing a *signal-to-symbol converter*:


- Edges demarcate the boundaries of objects, or of material properties.
- Objects have parts, and these are typically joined with edges.

- The three-dimensional distribution of objects in a scene usually generates occlusions of some objects by other objects, and these form occlusion edges which reveal the geometry of the scene.
- Edges can be generated in more abstract domains than luminance. For example, if some image property such as colour, or a textural signature, or stereoscopic depth, suddenly changes, it forms a highly informative “edge” in that domain.
- Velocity fields, containing information about the trajectories of objects, can be organized and understood by the movements of edges. (The motions of objects in space generates velocity discontinuities at their edges.)
- The central problem of stereoscopic 3D depth vision is the “correspondence problem:” matching up corresponding regions of two images from spatially displaced cameras. Aligning edges is a very effective way to solve the correspondence problem. The same principle applies to measuring velocities (for image frames displaced in time, rather than displaced in space) by tracking edges to align corresponding regions and infer velocity (ratio of object displacement to temporal interval).

In summary, DISCONTINUITIES = INFORMATION.



Original profile



Mask = [-1 1]



Mask = [1 -2 1]

An intuitive way to find edges is to compute the derivative of a (1D) signal, as this will be large where the luminance is changing rapidly. Since image arrays

are discrete, we must use the *finite difference* representation of a derivative, and this is implemented by a *convolution*: If our (1D) luminance array is $L[n]$ (sequence of pixels, index n), then the first-order finite difference operator $(h[0], h[1]) = (-1, 1)$ when convolved with $L[n]$ would generate an output which is large in amplitude only where $L[n]$ has edges (see previous figure).

However, note an important disadvantage of this approach: “rightward edges” (say, from dark to bright) generate the opposite sign from “leftward edges” (say, from bright to dark). We would prefer to generate the same detection signal regardless of the *polarity* of the edge.

A solution is to convolve the discrete luminance data $L[n]$ instead with the *second finite difference operator*, defined as $(h[-1], h[0], h[1]) = (1, -2, 1)$ and look for the *zero-crossings* of this operator. These correspond to peaks or troughs of the first finite difference operator that we considered above, and thus they reveal the edges, regardless of their polarity. Similarly for $(-1, 2, -1)$.

In the two-dimensional case, we have the choice of using *directional derivative* operators, or *non-directional* ones. An example of a directional operator is one which integrates (sums) pixels in one direction, but differentiates (differences) them in the perpendicular direction. Clearly, such an operator will detect edges only in a specific orientation; - namely the orientation along which the integration was done. A example of such a directional edge detector is the following 3 x 3 array:

-1	2	-1
-1	2	-1
-1	2	-1

In comparison, an *isotropic* operator such as the Laplacian (sum of second derivatives in two perpendicular orientations) has no preferred direction; that is the meaning of isotropy. It will detect edges in all orientations. The next picture illustrates such an effect. A discrete approximation to the Laplacian operator ∇^2 in just a 3 x 3 array is:

-1	-2	-1
-2	12	-2
-1	-2	-1

Notice how each of these simple 3 x 3 operators sums to zero when all of their elements are combined together. These types of operators (of which there are obviously numerous other examples, differing in array sizes as well as element composition) are called *filters*, because of their spectral consequences

for favouring some spatial frequency bands and orientations at the expense of others. Their zero-sum property means that they are insensitive to the overall brightness value of a scene, as we would desire: they have “no DC term.” (Their Fourier transform is equal to zero at the origin.) They also may, or may not, have a certain preferred, or characteristic *direction*; a certain phase or *symmetry* (even or odd); and a certain *scale*, defined by the spacing between changes of sign in the elements in (larger) arrays.




Figure: Illustration of edge-detection by convolution with an isotropic Laplacian operator, and marking the zero-crossings of the result of the convolution.

Edges in images are defined at different scales: some transitions in brightness are gradual, others very crisp. More importantly, at different scales of analysis, different edge structure emerges.

Example: an image of a leopard that has been low-pass filtered (or analyzed at a coarse scale) has edge outlines corresponding to the overall form of its body. At a somewhat finer scale of analysis, image structure may be dominated by the contours of its “spots.” At a still finer scale, the relevant edge structure arises from the texture of its fur.

In summary, non-redundant structure exists in images at different scales of analysis (or if you prefer, in different frequency bands).

The basic recipe for extracting edge information from images is to use a multi-scale family of image filters (convolution kernels). A wide variety of these are in standard use, differing in terms such as:

- isotropic (circularly symmetric), or anisotropic (directional)
- self-similar (dilates of each other), or not self-similar
- separable (expressible as product of two 1D functions), or not
- degree of conjoint uncertainty in the information resolved
- size of support (number of “taps,” or pixels, in the kernel)
- preferred non-linear outputs (zero-crossings; phasor moduli; energy)
- theoretical foundations (e.g. Logan’s Theorem)

5.1 The Laplacian $\nabla^2 G_\sigma(x, y) * I(x, y)$ and its zero-crossings. Logan’s Theorem.

One highly influential idea due to Marr (1981), that is frequently exploited for edge detection in machine vision systems, is to convolve the image with a multi-scale family of isotropic (non-directional) blurred 2nd-derivative filters, and to retain only their output zero-crossings. These correspond well to the edges in the image, at each chosen scale.

One primary motivation for doing this comes from Logan’s Theorem (1977) concerning the “richness” of Laplacian zero-crossings for band-limited signals. What Logan proved (albeit only in the 1D case) is that subject to two constraints, the zero-crossings alone suffice to represent the signal completely (i.e. it could be perfectly recovered from just its zeros, up to a scale factor).

This is a truly remarkable result. Consider the fact that a signal is continuous and dense, but in any finite interval it will have only a finite (countable) number of zero-crossings (e.g., 7). How can those 7 points completely determine what the signal does everywhere else within this finite interval??

The two constraints are:

1. The signal must be strictly bandlimited to one octave, or less. This means that its highest frequency component must be no more than twice its lowest frequency component.
(This constraint is much more powerful than it may appear.)
2. The signal must have no complex zeros in common with its Hilbert Transform. This effectively excludes purely amplitude-modulated signals. For example, a pure sinewave whose amplitude is merely modulated will have exactly the same zero-crossings as the unmodulated sinusoid, so their zero-crossings would not distinguish between them. Thus AM signals cannot be represented by zero-crossings.

The $\nabla^2 G_\sigma(x, y)$ filter kernel that is convolved with the image serves to bandpass-filter it. In the 2D Fourier domain, as we have seen, the spectral consequence of the Laplacian operator $\nabla^2 \equiv \left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} \right)$ is to multiply the image spectrum by a paraboloid: $(\mu^2 + \nu^2)$. Clearly this emphasizes the high frequencies at the expense of the low frequencies, and eliminates the DC component entirely (hence the output is centered around a mean of zero).

Blurring the Laplacian by a Gaussian $G_\sigma(x, y)$ of scale σ , simply limits the high-frequency components. The 2DFT of a Gaussian is also a Gaussian, with reciprocal dimension (by the Similarity Theorem discussed earlier). The scale parameter σ determines where the high-frequency cut-off occurs.

The resulting bandwidth of a $\nabla^2 G_\sigma(x, y)$ filter is about 1.3 octaves, regardless of what value for scale parameter σ is used. Note that this doesn't *quite* satisfy the first constraint of Logan's Theorem.

Note also that by commutativity of linear operators, the order in which these steps are applied to the image $I(x, y)$ doesn't matter. First computing the Laplacian of the image, and then blurring the result with the Gaussian, is equivalent to first convolving the image with the Gaussian and then computing the Laplacian of the result:

$$\nabla^2 [G_\sigma(x, y) * I(x, y)] = G_\sigma(x, y) * \nabla^2 I(x, y) \quad (13)$$

Moreover, both of these sequences are equivalent to just convolving the image with a single filter kernel, namely the Laplacian of a Gaussian: $[\nabla^2 G_\sigma(x, y)] * I(x, y)$. Clearly this is the preferred implementation, since it just involves a single convolution.

Some open theoretical issues in this approach are:

1. It is not clear how to generalize the constraint of one-octave bandlimiting to the case of 2D signals (images). E.g. should their 2DFT be confined to an annulus in the Fourier plane, whose outer radius is twice its inner radius?; or to four squares in the four quadrants of the Fourier plane that satisfy the one-octave constraint on each frequency axis? The first method doesn't work, and clearly the second filter is no longer isotropic!
2. Whereas the zeros of a 1D signal (soundwave) are denumerable [countable], those of a 2D signal (image) are not. Rather, they form "snakes" that are continuous contours in the plane.
3. As a practical matter, the $\nabla^2 G_\sigma(x, y) * I(x, y)$ approach to edge extraction

tends to be very noise-sensitive. Many spurious edge contours appear that shouldn't be there. This defect inspired the development of more sophisticated non-linear edge detectors, such as Canny's, which estimates the local image signal-to-noise ratio (SNR) to adaptively optimize its local bandwidth. This, however, is very computationally expensive.

4. Finally, strong claims were originally made that $\nabla^2 G_\sigma(x, y) * I(x, y)$ edge-detecting filters describe how human vision works. In particular, the receptive field profiles of retinal ganglion cells were said to have this form. However, counterexamples reveal several visual tasks that humans are able to perform, effortlessly and pre-attentively, which we could not perform if our visual systems functioned in this way.

6 Scale-space; multi-resolution. Wavelets as visual primitives.

Images contain information at multiple scales of analysis, so detecting visual features (such as edges) must be done across a range of different scales.

- An interesting property of edges as defined by the zero-crossings of multi-scale operators whose scale is determined by convolution with a Gaussian, is that as the Gaussian is made coarser (larger), new edges (new zero-crossings) can never appear. They can only merge and thus become fewer in number. This property is called causality. It is also sometimes called 'monotonicity,' or 'the evolution property,' or 'nice scaling behaviour.'
- One reason why causality is important is that it ensures that features detected at a coarse scale of analysis were not spuriously created by the blurring process (convolution with a low-pass filter) which is the normal way to create a multi-scale image pyramid using a hierarchy of increasing kernel sizes. One would like to know that image features detected at a certain scale are "grounded" in image detail at the finest resolution.
- For purposes of edge detection at multiple scales, a plot showing the evolution of zero-crossings in the image after convolution with a linear operator, as a function of the scale of the operator which sets the scale (i.e. the width of the Gaussian), is called scale-space.
- Scale-space has a dimensionality that is one greater than the dimensionality of the signal. Thus a 1D waveform projects into a 2D scale-space. An image projects into a 3D scale space, with its zero-crossings (edges) forming surfaces that evolve as the scale of the Gaussian changes. The scale of the Gaussian, usually denoted by σ , creates the added dimension.
- A mapping of the edges in an image (its zero-crossings after such filtering operations, evolving with operator scale) is called a scale-space fingerprint.

Several theorems exist called “fingerprint theorems” showing that the Gaussian blurring operator uniquely possesses the property of causality. In this respect, it is a preferred edge detector when combined with a bandpass or differentiating kernel such as the Laplacian.

- However, other non-linear operators have advantageous properties, such as reduced noise-sensitivity and greater application for extracting features that are more complicated (and more useful) than mere edges.

6.1 2D Gabor “Logons;” Quadrature Pair Wavelets

The family of filters which uniquely achieve the lowest possible conjoint uncertainty (i.e. minimal dispersion, or variance) in both the space domain and the Fourier domain are the complex exponentials multiplied by Gaussians. These are sometimes known as Gabor wavelets, or “logons.” In one dimension:

$$f(x) = \exp(-i\mu_0(x - x_0)) \exp(-(x - x_0)^2/\alpha^2)$$

This is a Gaussian localized at position x_0 , complex modulated at frequency μ_0 , and with size or spread constant α . It is noteworthy that such wavelets have Fourier Transforms $F(\mu)$ with exactly the same functional form, but with their parameters merely interchanged or inverted:

$$F(\mu) = \exp(-ix_0(\mu - \mu_0)) \exp(-(\mu - \mu_0)^2\alpha^2)$$

Note that for the case of a wavelet $f(x)$ centered on the origin so $x_0 = 0$, its Fourier Transform $F(\mu)$ is simply a Gaussian centered on the modulation frequency $\mu = \mu_0$, and whose width is $1/\alpha$, the reciprocal of the wavelet’s space constant. This shows that it acts as a bandpass filter, passing only those frequencies that are within about $\pm \frac{1}{\alpha}$ of the wavelet’s modulation frequency μ_0 .

Dennis Gabor (1946) named these wavelets “logons” from the Greek word for information, or order: *logōs*. Because of the optimality of such wavelets under the Uncertainty Principle, Gabor proposed using them as an expansion basis to represent signals. In particular, he wanted them to be used in broadcast telecommunications for encoding continuous-time information. He called them the “elementary functions” for a signal. Unfortunately, because such functions are mutually non-orthogonal, it is very difficult to obtain the actual coefficients to be used with the elementary functions in order to expand a given signal in this basis. (Gabor himself could not solve this problem, although he went on to invent holography and to win the Nobel Prize in Physics in 1974.)

When a family of such Gabor functions are parameterized to be self-similar, i.e. they are dilates and translates of each other so that they all have a common

template (“mother” and “daughter”), then they constitute a (non-orthogonal) *wavelet basis*. Today it is known that infinite classes of wavelets exist which can be used as the expansion basis for signals. Because of the self-similarity property, this amounts to representing or analyzing a signal at different scales. This general field of investigation is called *multi-resolution analysis*, and we have already encountered its importance for extracting edge features.

6.2 Generalization of Wavelet Logons to 2D for Image Analysis




Figure 1: The real part of a 2D Gabor wavelet, and its 2D Fourier transform.

An effective method for extracting, representing, and analyzing image structure is the computation of the 2D Gabor wavelet coefficients for the image. This family of 2D filters were originally proposed as a framework for understanding the orientation-selective and spatial-frequency-selective receptive field properties of neurons in the brain’s visual cortex, as well as being useful operators for practical image analysis problems. These 2D filters are conjointly optimal in extracting the maximum possible information both about the orientation and modulation of image structure (“what”), simultaneously with information about 2D position (“where”). The 2D Gabor filter family uniquely achieves the theoretical lower bound on joint uncertainty over these four variables in the Uncertainty Principle when it is suitably generalized.

These properties are particularly useful for texture analysis because of the 2D spectral specificity of texture as well as its variation with 2D spatial position. These wavelets are also used for motion detection, stereoscopic vision, and many sorts of visual pattern recognition such as face recognition. A large and growing literature now exists on the efficient use of this non-orthogonal expansion basis and its applications.

Two-dimensional Gabor wavelets have the functional form:

$$f(x, y) = e^{-[(x-x_0)^2/\alpha^2 + (y-y_0)^2/\beta^2]} e^{-i[u_0(x-x_0) + v_0(y-y_0)]}$$

where (x_0, y_0) specify position in the image, (α, β) specify effective width and length, and (u_0, v_0) specify modulation, which has spatial frequency $\omega_0 = \sqrt{u_0^2 + v_0^2}$ and direction $\theta_0 = \arctan(v_0/u_0)$. (A further degree-of-freedom not included above is the relative orientation of the elliptic Gaussian envelope, which creates cross-terms in xy .) The 2D Fourier transform $F(u, v)$ of a 2D Gabor wavelet has exactly the same functional form, with parameters just interchanged or inverted:

$$F(u, v) = e^{-[(u-u_0)^2\alpha^2 + (v-v_0)^2\beta^2]} e^{-i[x_0(u-u_0) + y_0(v-v_0)]}$$

The real part of one member of the 2D Gabor filter family, centered at the origin $(x_0, y_0) = (0, 0)$ and with unity aspect ratio $\beta/\alpha = 1$ is shown in the Figure, together with its 2D Fourier transform $F(u, v)$.

By appropriately parameterizing them for dilation, rotation, and translation, 2D Gabor wavelets can form a complete self-similar (but non-orthogonal) expansion basis for images. If we take $\Psi(x, y)$ to be some chosen generic 2D Gabor wavelet, then we can generate from this one member a complete self-similar family of 2D wavelets through the generating function

$$\Psi_{mpq\theta}(x, y) = 2^{-2m}\Psi(x', y')$$

where the substituted variables (x', y') incorporate dilations in size by 2^{-m} , translations in position (p, q) , and rotations through orientation θ :

$$x' = 2^{-m}[x \cos(\theta) + y \sin(\theta)] - p$$

$$y' = 2^{-m}[-x \sin(\theta) + y \cos(\theta)] - q$$

It is noteworthy that as consequences of the similarity theorem, shift theorem, and modulation theorem of 2D Fourier analysis, together with the rotation isomorphism of the 2D Fourier transform, all of these effects of the generating function applied to a 2D Gabor mother wavelet $\Psi(x, y) = f(x, y)$ have corresponding identical or reciprocal effects on its 2D Fourier transform $F(u, v)$. These properties of self-similarity can be exploited when constructing efficient, compact, multi-scale codes for image structure.

The completeness of 2D Gabor wavelets as an expansion basis for any image can be illustrated by reconstruction of a facial image, in stages. Note how efficiently the facial features, such as the eyes and mouth, are represented using only a handful of the wavelets. Later we will see how this can be exploited both for automatic feature localization, and for face recognition.



Reconstruction of Lena: 25, 100, 500, and 10,000 Two-Dimensional Gabor Wavelets

6.3 Unification of Domains


Until now we have viewed “the image domain” and “the Fourier domain” as very different domains of visual representation. But now we can see that the “Gabor domain” of representation actually embraces and unifies both of these other two domains. How?

In the wavelet equations above, the scale constant α (and β in the 2D case) actually builds a continuous bridge between the two domains. If the scale constant is set very large, then the Gaussian term becomes just 1 and so the expansion basis reduces to the familiar Fourier basis. If instead the scale constant is made very small, then the Gaussian term shrinks to a discrete delta function (1 only at the location $x = x_0$, and 0 elsewhere), so the expansion basis implements pure space-domain sampling: a pixel-by-pixel image domain

representation. This allows us to build a continuous deformation between the two domains when representing, analyzing, and recognizing image structure, merely by changing a single scaling parameter in this remarkable, unifying, expansion basis.

A “philosophical” comment about 2D Gabor wavelets.


Aristotle defined vision as “knowing what is where.” We have noted the optimality (conjoint uncertainty minimization) property of 2D Gabor wavelets in the two domains for extracting structural (“what”) and positional (“where”) information. Thus if we share Aristotle’s goal for vision, then we cannot do better than to base computer vision representations upon these wavelets. Perhaps this is why mammalian visual systems appear to have evolved their use; the receptive field profiles of isolated neurons in the brain’s visual cortex, as determined by the spatial distribution of excitatory and inhibitory inputs to each so-called “simple cell,” can be well-described as quadrature-paired 2D Gabor wavelets. At the present time, this is basically the standard model for how the brain’s visual cortex represents the information in the retinal image.



6.4 Detection of Facial Features by Quadrature Gabor Wavelet Energy

One illustration of a practical application of such image operators is in the automatic localization of facial features. Interestingly, most facial features themselves can be described by only a handful of wavelets, since such features are after all just localized undulations having certain positions, orientations, spatial frequencies, and phases. By taking the modulus (sum of the squares of the real and imaginary parts) of a facial image after convolving it with complex-valued 2D Gabor wavelets, key facial features (eyes and mouth) are readily detected; we may call this a Quadrature Demodulator Neural Neural. This capability is illustrated in the next two Figures.

Quadrature Demodulator Network



Neural Network for Image Analysis. The above neurobiologically-inspired network performs image demodulation using 2D Gabor wavelets, in order to find salient features in the image that have a characteristic orientation and scale or frequency composition. The operation of the biphasic receptive fields (representing even- and odd-symmetric visual cortical neurons) is described by:

$$g(x, y) = \int_{\alpha} \int_{\beta} e^{-((x-\alpha)^2 + (y-\beta)^2)/\sigma^2} \cos(\omega(x - \alpha)) I(\alpha, \beta) d\alpha d\beta$$

$$h(x, y) = \int_{\alpha} \int_{\beta} e^{-((x-\alpha)^2 + (y-\beta)^2)/\sigma^2} \sin(\omega(x - \alpha)) I(\alpha, \beta) d\alpha d\beta$$

and the demodulated output at the top of the network resembles that of the brain's "complex cells" which combine inputs from the quadrature simple cells as:

$$A^2(x, y) = g^2(x, y) + h^2(x, y)$$




Illustration of Facial Feature Detection by Quadrature Filter Energy. Left panel: original image. Right panel (clockwise from top left): the real part after 2D Gabor wavelet convolution; the imaginary part; the modulus energy; and this energy superimposed on the original (faint) image, illustrating successful feature localization.

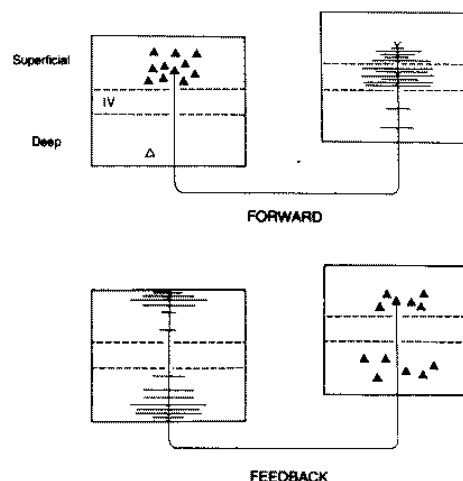
These operations can be performed by the **Quadrature Demodulator Network** shown in the previous Figure.

7 Higher brain visual mechanisms; streaming; reciprocal feedback

Besides the primary visual cortex in the occipital lobe, there are at least 30 further visual areas distributed across the parietal and temporal cortices of the brain. Many of these are specialised for particular kinds of visual processing, including colour (V4), motion (MT), stereo (Area 18), and facial and other form processing areas. There is a pronounced functional streaming, or division of labour, for form, colour, and motion processing; some neuroscientists have proposed a fundamental division into “two visual systems” along lines such as magno and parvo (fast/slow) or even conscious and unconscious vision.





The existence of so many distinct visual areas in the brain almost begs the question of “how the visual world gets put back together again.” An intriguing aspect of this architecture is the pattern of reciprocating feedback connections.



In general there are pairwise reciprocating connections between visual areas from the deep layers in one area to superficial layers in another area, whose deep layers in turn project back to the superficial layers of the first. Just as noted earlier with the massive feedback projections from primary visual cortex back down to the LGN (where it meets afferent signals ascending from the eyes), these reciprocating projection pathways are perhaps suggestive of a kind of “hypothesis generation and testing” iterative strategy for understanding the visual environment and the objects that populate it.

The fovea tends to be represented in all visual areas, and the mapping from the retina is retinotopic (meaning that adjacent points in the retinal image usually project to adjacent points in a given cortical map); but typically there is a highly pronounced geometrical distortion. In part this reflects a great over-representation of the fovea, which is called cortical magnification factor. In the foveal projection to primary visual cortex, about 6mm of neural tissue is devoted to 1 degree of visual angle, whereas in the periphery, 1mm of neural tissue handles about 6 degrees. It has been proposed that the geometrical distortion in visual mapping actually serves a specific mathematical role, that of achieving pattern representations that are invariant to rotation and dilation because of log-polar projection. Crudely speaking, this converts a polar grid (whose concentric circles have geometrically-increasing radii) into a cartesian grid with a nearly uniform lattice. Thus changes in object distance (hence image size) become just translations along one axis, while rotations become just translations along the other axis, thereby facilitating pattern recognition.





8 Texture, colour, stereo, and motion descriptors. Disambiguation.

Many seemingly disparate tasks in computer vision actually share a common formal structure: to convert ill-posed, insoluble problems of inference from raw data, into well-posed problems in which we can compute object properties disambiguated from the image-formation processes which confound the raw luminance data itself.


One obvious aspect of this issue is the fact that images are 2D projections of 3D data which could, in principle, arise equally well from many different constellations of worlds and objects. A more subtle aspect is the fact that the information received as an image is the compound product of several factors that are difficult to disambiguate: (1) the nature, geometry, and wavelength composition of the illuminant(s); (2) properties of the objects imaged, such as: spectral reflectances; surface shape; surface albedo; surface texture; geometry, motion, and rotation angle; and (3) properties of the camera (or viewer), such as (i) geometry and viewing angle; (ii) spectral sensitivity; (iii) prior knowledge, assumptions, and expectations. The aim of this lecture is to study how these many factors can be disambiguated and even exploited, in order to try to make objective inferences about object and world properties from these ambiguous and confounded image properties.

8.1 Texture information.

Most surfaces are covered with texture of one sort or another. Texture can serve not only as a helpful identifying feature, but more importantly as a cue to surface shape because of the foreshortening it undergoes as it follows the shape of the object if one can assume that it has some uniform statistics along the surface itself. The following patterns illustrate the inference of surface slant and of 3D surface shape from texture cues when they are combined with the assumption of texture uniformity on the surface itself:



Texture is also a useful cue to image segmentation by parsing the image into local regions which are relatively homogeneous in their textural properties. Here are some illustrations:




How can one measure something as ill-defined as a “textural signature?” What is texture, anyway?

As implied by the root of the word, which links it to textiles, texture is defined by the existence of certain statistical *correlations* across the image. These can be almost anything, from quasi-periodic undulations as one might see in water ripples or in woven fabrics, to repetitive but highly punctate features. Many natural scenes, such as woodlands, grasslands, mountain ranges and other terrains, have such properties which give them a distinctive identifying visual signature. The unifying notion in all of these examples is *quasi-periodicity*, or repetitiveness, of some features.

The detection of quasi-periodicity is best done by Fourier methods. There are deep and multi-faceted links between many topics in statistics (such as time-series analysis, correlation, moments) and Fourier analysis. These links arise from the fact that the eigenfunctions of the Fourier transform, complex exponentials (sinusoids in quadrature), are of course periodic but also have a specific scale (frequency) and direction (wavefront). Thus they excel in detecting the existence of a correlation distance and direction, and in estimating the relative “power” represented in various components of quasi-periodic correlated structures.

Unfortunately, these eigenfunctions are globally defined, but we wish to use local regional information as a basis for texture-based image segmentation. Hence the ideal solution is to “window” the sinusoids so that they analyze the image characteristics only within a local region and thus extract the spectral statistics as a function that varies with location. The optimal set of windowing functions are bivariate Gaussians, since their joint spatial/spectral localization

is greater than that of any other function. The product of complex exponentials times bivariate Gaussians are called *2D Gabor wavelets*, and these form a complete basis for image analysis and representation. The pictures below illustrate successful segmentation of collages of textured natural scenes, as well as of textured artificial objects, using such 2D wavelets for local spectral analysis to infer and measure their textural discriminators.




8.2 Colour information.

Colour is a nearly ubiquitous property of surfaces. Just like texture, it can serve both in object identification and in scene segmentation. But the fundamental difficulty in using the wavelength composition of images to infer the colour properties (“spectral reflectances”) of objects, is the fact that the wavelengths received depend as much upon the *illuminant* as upon the spectral reflectances of the surface that is scattering back the light. When a yellow banana is illuminated in bluish light, the image that it forms obviously has a very different wavelength composition than when it is illuminated in reddish light. The central mystery of human colour perception is the fact that the banana still appears yellow (“colour constancy”). In computer vision, how can we possibly achieve this same vital capability of inferring an inherent underlying *object property* from a confounded (i.e., illuminant-wavelength dependent) set of *image properties*?

To give the problem a slightly more formal presentation:

- Let $I(\lambda)$ represent the wavelength composition of the illuminant (i.e. the amount of energy it contains as a function of wavelength λ , across the visible spectrum from about 400 nanometers to 700 nm).
- Let $O(\lambda)$ represent the inherent spectral reflectance of the object at a particular point: the fraction of incident light that is scattered back from its surface there, as a function of the incident light's wavelength λ .
- Let $R(\lambda)$ represent the actual wavelength mixture received by the camera at the corresponding point in the image of the scene.

Clearly, $R(\lambda) = I(\lambda)O(\lambda)$. The problem is that we wish to infer the “object colour” (its spectral reflectance as a function of wavelength, $O(\lambda)$), but we only know $R(\lambda)$, the actual wavelength mixture received by our sensor. So unless we can measure $I(\lambda)$ directly, how could this problem of inferring $O(\lambda)$ from $R(\lambda)$ possibly be solved?



One simple idea that has been proposed is to try actually to measure $I(\lambda)$ directly, by searching for highly specular (shiny, metallic, glassy) regions in an image where the reflected light might be a fairly faithful copy of $I(\lambda)$. This might be a glint from someone's glasses or from a shiny doorknob. Then at all other points in the image we need only to divide the $R(\lambda)$ we receive there by our other specular “measurement” of $I(\lambda)$, and we can then compute the desired $O(\lambda)$ across the image.

Clearly, this method has several weaknesses: (1) there may be no specular surfaces in the image; (2) those that there are may themselves affect somewhat the wavelength composition that they reflect (e.g. metals which have a brassy colour); and (3) the method is neither robust nor stable, since global inferences about scene interpretation depend critically upon uncertain measurements at (what may be just) a single tiny point in the image.

A more stable and interesting approach was developed by Dr E Land, founder of Polaroid, and is called the Retinex because he regarded it as modelled after

biological visual systems (RETINa + cortEX). Land's critical observation was that (contrary to almost universal popular belief), the colour perceived in an area of a scene is *not* determined by the wavelength composition of light received from that area (!). A simple experiment proves this: illuminate a scene, such as a bowl of fruit containing (say) a yellow banana, a red tomato and a green pepper, with three different narrowband light sources, each of which contains a different wavelength (say red, green, or blue) and with adjustable intensities. (No other light sources are present.)

The first observation is that even under drastic changes in the intensities of each of the three illuminators, the objects maintain exactly their normal colours. Obviously the wavelength mixture reaching the eye from each object is drastically changing, in proportion to the illuminators, but there are no changes in perceived colours. The phenomenon does not depend upon knowing the natural colours for objects identifiable by (say) their shape; a collage of patches of coloured paper cut into random shapes, forming a *mondrian*, produces exactly the same effect.


The second observation is that even when the wavelength composition of light reflected from each object is exactly the same (i.e. the three light sources are adjusted separately for each object to ensure that the light reflected in the three wavebands as measured by a spectral photometer is exactly the same for each of the objects individually), they *still* retain their natural colours. The banana still looks yellow, the tomato still looks red, and the pepper still looks green, even when each one is sending *identical* wavelength “messages” to your eyes. This is rather miraculous.

The Retinex algorithm attempts to account for this remarkable biological phenomenon, and to provide a means to achieve similar *colour constancy* in computer vision systems so that they may “discount the illuminant” and infer the spectral reflectance properties of objects, independent of the composition of their illumination. Only a cursory description of Retinex will be given here.

The key idea is that *the colours of objects or areas in a scene are determined by their surrounding spatial context*. A complex sequence of ratios computed across all the boundaries of objects (or areas) enables the illuminant to be algebraically discounted in the sense shown in the previous Figure, so that object spectral reflectances $O(\lambda)$ which is what we perceive as their colour, can be inferred from the available retinal measurements $R(\lambda)$ without explicitly knowing $I(\lambda)$.

8.3 Stereo information

Important information about depth can be obtained from the use of two (or more) cameras, in the same way that humans achieve stereoscopic depth vision by virtue of having two eyes. Objects in front or behind of the point in space at which the two optical axes intersect (as determined by the angle between them, which is controlled by camera movements or eye movements), will project into different relative parts of the two images. This is called *stereoscopic disparity*.




This “error signal” becomes greater in proportion to the distance of the object in front or behind the point of fixation, and so it can be calibrated to obtain a depth cue. It also becomes greater with increased spacing between the two eyes or cameras, since that is the “base of triangulation.” (That is why the German Army in WWI introduced V-shaped binocular “trench periscopes” to increase stereoscopic visual acuity, for breaking camouflage by increasing the *effective* spacing between the viewer’s two eyes to almost a meter.)

The essence of making use of such stereoscopic disparity cues is the need to solve the Correspondence Problem. In order to infer that the cylinder is in a different position relative to the background objects in the two frames shown, it is first necessary to detect the correspondence of the background objects in the two frames, or at least of their edges. This puts the two frames “into registration,” so that the disparity of the foreground object can be detected.

Unfortunately, current algorithms for solving the Correspondence Problem tend to require very large searches for matching features under a large number of possible permutations. It is difficult to know which set of features in the two frames to select for comparison in evaluating the degree of alignment, when trying to find that relative registration which generates maximum correlation between the two background scenes.

One helpful approach here is to use a “multi-scale image pyramid,” which steers the search in a coarse-to-fine fashion to maximize its efficiency. In initially sparsely sampled (coarsely blurred and under-sampled) images, the permutation-matching space of possible corresponding points is greatly attenuated compared with full-resolution images.

After an adequate alignment match is found for low-resolution (blurred) copies of the image pair, the process repeats on somewhat higher resolution (less blurred) copies of the image pair but over a search space that has been greatly curtailed by having first found the coarse-scale solution. Such “pyramid” processes usually increment in one-octave steps (factors of two in improved resolution), from coarse to fine, spanning a total of perhaps four or five levels before the final solution is determined to within single-pixel precision.



Once the Correspondence Problem has thereby been solved, the inference of depth from object disparity in the two image frames is then just a matter of triangulation and “look-up” from a calibration table which includes information about the spacing between the two cameras (or eyes) and their focal lengths. (See the above simplifying diagram, for the case that the two cameras’ optical axes are parallel and hence converged at infinity.) Specifically, if the two cameras have focal length f and the optical centres of their lenses (remember the trench periscopes!) are separated by a distance b , and the disparity in the projections of some object point onto the two images (in opposite directions relative to their optical axis) is α in one image and β in the other image, then the distance d to the object in front of the two lenses is simply:

$$d = fb / (\alpha + \beta)$$

8.4 Motion information

Only a few vision applications actually involve just static image frames. That is basically vision “off-line;” – but the essence of an effective visual capability must be for real-time use in a dynamic environment. This requires the ability to detect and measure motion, and to be able thereby to draw inferences quickly (such as time-to-collision).

In a formal sense, the problem of computing motion information from an image sequence is very similar to that of computing stereo information.

- For stereo vision, we need to solve the Correspondence Problem for two images simultaneous in time but acquired with a spatial displacement.
- For motion vision, we need to solve the Correspondence Problem for two images coincident in space but acquired with a temporal displacement.
- The object’s spatial “disparity” that can be measured in the two image frames once their backgrounds have been aligned, can be calibrated to reveal motion information when compared with the time interval, or depth information when compared with the binocular spatial interval.

Among the challenging requirements of motion detection and inference are:

1. Need to infer 3D object trajectories from 2D image motion information.
2. Need to make *local* measurements of velocity, which may differ in different image regions in complex scenes with many moving objects. Thus, a velocity vector field needs to be assigned over an image.
3. Need to disambiguate object motion from contour motion, so that we can measure the velocity of an object regardless of its *form*.
4. Need to measure velocities regardless of the size of the viewing aperture in space and in time (the spatial and temporal integration windows).
5. It may be necessary to assign more than one velocity vector to any given local image region (as occurs in “motion transparency”)
6. We may need to detect a *coherent* overall motion pattern across many small objects or regions separated from each other in space.


The major classes of models and approaches to motion detection are largely inspired by detailed neurobiological studies of motion processing both in the invertebrate eye and in mammalian retina and cortex. Diverse mathematical frameworks have been proposed, but the main classes of models are:

INTENSITY GRADIENT MODELS .

Assume that the local time-derivative in image intensities at a point, across many image frames, is related to the local spatial gradient in image intensities because of object velocity \vec{v} :

$$-\frac{\partial I(x, y, t)}{\partial t} = \vec{v} \cdot \vec{\nabla} I(x, y, t)$$

Then the ratio of the local image time-derivative to the spatial gradient is an estimate of the local image velocity (in the direction of the gradient).



DYNAMIC ZERO-CROSSING MODELS .


Measure image velocity by first finding the edges and contours of objects (using the zero-crossings of a blurred Laplacian operator!), and then take the time-derivative of the Laplacian-Gaussian-convolved image:

$$-\frac{\partial}{\partial t} [\nabla^2 G_\sigma(x, y) * I(x, y, t)]$$

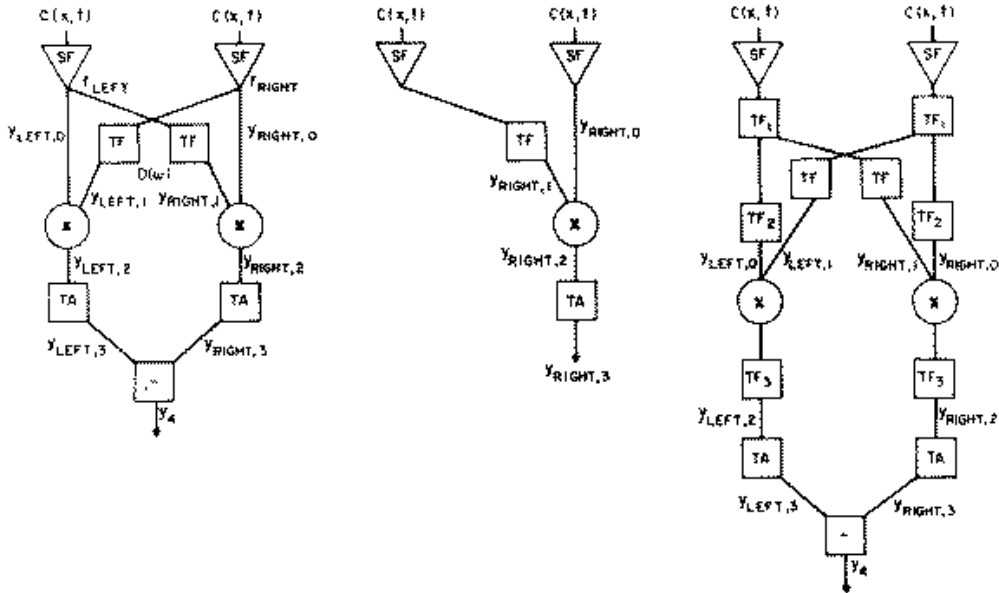
in the vicinity of a Laplacian zero-crossing. The amplitude of the result is an estimate of speed, and the sign of this quantity determines the direction of motion relative to the normal to the contour.

SPATIO-TEMPORAL CORRELATION MODELS .

Image motion is detected by observing a *correlation* of the local image signal $I(x, y, t)$ across an interval of space and and after an interval of time τ . Finding the pair of these intervals which maximizes the correlation between $I(x, y, t)$ and $I(x - v_x\tau, y - v_y\tau, t - \tau)$ determines the two components of image velocity v_x and v_y which we desire to know.



Detailed studies of fly neural mechanisms (above) for motion detection and visual tracking led to elaborated correlation-based motion models:



SPATIO-TEMPORAL SPECTRAL MODELS .

It is possible to detect and measure image motion purely by *Fourier* means. This approach exploits the fact that motion creates a covariance in the spatial and temporal *spectra* of the time-varying image $I(x, y, t)$, whose three-dimensional (spatio-temporal) Fourier transform is defined:

$$F(\omega_x, \omega_y, \omega_t) = \int_X \int_Y \int_T I(x, y, t) e^{-i(\omega_x x + \omega_y y + \omega_t t)} dx dy dt$$

In other words, rigid image motion has a 3D spectral consequence: the local 3D spatio-temporal spectrum, rather than filling up 3-space $(\omega_x, \omega_y, \omega_t)$, collapses onto a 2D inclined plane which includes the origin. Motion detection then occurs just by filtering the image sequence in space and in time, and observing that tuned spatio-temporal filters whose center frequencies are **co-planar** in this 3-space are activated together. This is a consequence of the **SPECTRAL CO-PLANARITY THEOREM**:

Theorem: Translational image motion of velocity \vec{v} has a 3D spatio-temporal Fourier spectrum that is non-zero only on an inclined plane through the origin of frequency-space. Spherical coordinates of the unit normal to this spectral plane correspond to the speed and direction of motion.

Let $I(x, y, t)$ be a continuous image in space and time.

Let $F(\omega_x, \omega_y, \omega_t)$ be its 3D spatio-temporal Fourier transform:

$$F(\omega_x, \omega_y, \omega_t) = \int_X \int_Y \int_T I(x, y, t) e^{-i(\omega_x x + \omega_y y + \omega_t t)} dx dy dt.$$

Let $\vec{v} = (v_x, v_y)$ be the local image velocity.

Uniform motion \vec{v} implies that for all time shifts t_o ,

$$I(x, y, t) = I(x - v_x t_o, y - v_y t_o, t - t_o).$$

Taking the 3D spatio-temporal Fourier transform of both sides, and applying the shift theorem, gives

$$F(\omega_x, \omega_y, \omega_t) = e^{-i(\omega_x v_x t_o + \omega_y v_y t_o + \omega_t t_o)} F(\omega_x, \omega_y, \omega_t).$$

The above equation can only be true if $F(\omega_x, \omega_y, \omega_t) = 0$ everywhere the exponential term doesn't equal 1.

This means $F(\omega_x, \omega_y, \omega_t)$ is non-zero only on the 3D spectral plane

$$\boxed{\omega_x v_x + \omega_y v_y + \omega_t = 0} \quad \text{Q.E.D.}$$

The spherical coordinates $(\theta, \phi, 1)$

$$\phi = \tan^{-1} \left(\omega_t / \sqrt{\omega_x^2 + \omega_y^2} \right)$$

$$\theta = \tan^{-1} (\omega_y / \omega_x)$$


of the inclined spectral plane's unit normal are determined by \vec{v} and correspond to the *speed* (ϕ) and *direction* (θ) of motion:

$$\phi = \sqrt{v_x^2 + v_y^2}$$

$$\theta = \tan^{-1} (v_y / v_x)$$

9 Lambertian and specular surfaces. Reflectance maps.

How can we infer information about the surface reflectance properties of objects from raw measurements of image brightness? This is a more recondite matter than it might first appear, because of the many complex factors which determine how (and where) objects scatter light.



The definitions of the angles i , e , and g .

Some definitions of surface type and properties:

- Surface albedo refers to the fraction of the illuminant that is re-emitted from the surface in all directions, in total. Thus, albedo corresponds more-or-less to “greyness.”

The amount of light reflected is the product of two factors: the albedo of the surface, times a geometric factor that depends on angle.

- A Lambertian surface is “pure matte.” It reflects light equally well in all directions.

Examples of Lambertian surfaces include snow, non-glossy paper, ping-pong balls, magnesium oxide, projection screens,...

A Lambertian surface looks equally bright from all directions; the amount of light reflected depends only on the angle of incidence.

If you looked inside a “Lambertian bottle” with an arbitrarily complex shape, illuminated with a point source of light from any angle, you could never infer the interior shape! It would have uniform brightness everywhere regardless of its actual shape.

- A specular surface is locally mirror-like. It obeys Snell’s law (i.e. the angle of incidence of light is equal to the angle of reflection from the surface), and does not scatter light. Most metallic surfaces are specular.
- The reflectance map is a function $\phi(i, e, g)$ which relates intensities in the image to surface orientations of objects. It specifies the fraction of incident light reflected per unit surface area, per unit solid angle, in the direction of the camera; thus it has units of flux/steradian. It is a function of three variables (see previous Figure): i is the angle of the illuminant, relative to the surface normal N ; e is the angle of a ray of light re-emitted from the surface; and g is the angle between the emitted ray and the illuminant.

There are many types of reflectance functions, each of which is characteristic of certain surfaces and imaging environments. For a Lambertian surface, the reflectance function $\phi(i, e, g) = \cos(i)$. It looks equally bright viewed from all directions; the amount of reflected light depends only on angle of illumination.

For surfaces such as the dusty surface of the moon, the reflectance function $\phi(i, e, g)$ depends only upon the ratio of the cosines of the angles of incidence and emission: $\cos(i)/\cos(e)$, but not upon their relative angle g nor upon the surface normal N . In case you ever wondered, this is why the moon looks like a penny rather than a sphere. Even though the moon is illuminated by a point source (the sun), it does not fade in brightness towards its limbs (as N varies). Surfaces with this property are called *lunar* surfaces.

For a specular surface, the reflectance function $\phi(i, e, g)$ is especially simple: $\phi(i, e, g) = 1$ when $i = e$ and both are coplanar with the surface normal N , so $g = i + e$ (Snell’s law for a pure mirror); and $\phi(i, e, g) = 0$ otherwise.


Typically, surfaces have both specular and matte properties. For example, facial skin may vary from lunar (powdered) to specular (oily). The main purpose of powdering one’s face is to specify s and n in this expression:

$$\phi(i, e, g) = \frac{s(n+1)(2\cos(i)\cos(e) - \cos(g))^n}{2} + (1-s)\cos(i) \quad (14)$$

The first term is the specular component, and the second term is the matte component. s is the fraction of light emitted specularly, and n represents the

sharpness of the specular peak. For glossy paint, typically the exponent n may be about 20. (Obviously as n grows very large, the exponentiated trig function approaches a delta function representing Snell's law: a very sharp power function of angle.)

Typically there is not just one point source of illumination, but rather a multitude of sources (such as the extended light source provided by a bright overcast sky). In a cluttered scene, much of the light received by objects has been reflected from other objects (and coloured by them...) One needs almost to think of light not in terms of ray-tracing but in terms of thermodynamics: a “gas” of photons in equilibrium inside a room.



Clearly, the only way to infer the nature and geometry of surface properties from image properties, given all of these complications in the way that surfaces reflect and scatter light, is to build in certain assumptions about the nature of the surfaces from other kinds of evidence. This requires us to consider the general problem of inference and integration of evidence.

10 Shape description. Codons, superquadrics, surface geometry.

Just as illustrated earlier by the examples of inferring surface and object properties from texture, colour, stereo, and motion information, the shading and brightness variation within an image is another important cue to surface shape.

As with all of these problems, computing “shape-from-shading” requires the disambiguation of many confounding factors. These arise from the

1. geometry of the illuminant (e.g. is the light a point source or extended? If a point source, where is it?) Are there several light sources? How will these affect the shading and shadowing information?

2. reflectance properties of the surface. What kind of surface is it – e.g. Lambertian, or specular, or a combination of both?
3. geometry of the surface (its underlying shape). Are shadows cast?
4. rotations of the surface relative to perspective angle and illuminant.
5. variations in material and surface reflectance properties across space (e.g. variation from Lambertian to specular where skin becomes more oily).
6. variations in surface albedo (“greyness”)

The inference of a surface shape (a relief map, or an object-centred description of a surface) from shading information is an inherently ill-posed problem because the data necessary for the computation is simply not known. One has to introduce ancillary assumptions about the surface material composition, its albedo and specularity parameters, the illumination of the scene and its geometry, before such inferences become possible. It is almost as though the assumptions are more important than the available image data. The computational nature of the inference task then becomes one of *constraint satisfaction*, and solving such a problem is often formulated as an optimisation or *relaxation problem*. Often there are rivalrous alternative solutions. In human vision these can be triggered and alternated (e.g. converting a crater into an apparent mound, reversing the inferred surface shape) simply by changing a cue about the direction of illumination.




10.1 How should shape be represented? Boundary descriptors; codons.

Closed boundary contours can be represented completely by their curvature map $\theta(s)$ (the reciprocal of the local radius of curvature $r(s)$ as a function of position s along the contour). This is the *Fundamental Theorem of Curves*. Local radius of curvature $r(s)$ is defined as the limiting radius of the circle that best


“fits” the contour at position s , in the limit as the arc length Δs shrinks to 0, and the local curvature of the contour there is:

$$\theta(s) = \lim_{\Delta s \rightarrow 0} \frac{1}{r(s)}$$



Closed boundary contours can be expanded with basis functions (such as “Fourier descriptors” of the radius of curvature) from their curvature map, in order to generate a shape description that is invariant to translation, rotation, and dilation. By cataloging a list of all possible combinations of changes in sign of the curvature map relative to the zeroes of curvature, it is possible to generate a restricted “grammar” for the shapes of closed contours. A lexicon of all possible shapes having a certain number of zeroes-of-curvature generates a list of “codons,” from which shapes can be classified and recognized. Interestingly, Logan’s Theorem (about the richness of zero-crossings for capturing bandlimited 1D signals completely) arises again in this context: the curvature map of a closed contour is a bandlimited signal, and it can be described by its zero-crossings; such a description amounts to a shape classification. This is one of several approaches proposing an elementary grammar for shape, and it can be generalised to surfaces.

Codon Triples (5)



The curvature map $\theta(s)$ together with a “starting point” tangent $t(s_o)$ specifies a shape fully. Some nice properties of curvature-map descriptions are:

1. The description is position-independent (i.e., object-centred).
2. The description is orientation-independent (rotating the shape in the plane does not affect its curvature map).
3. The description represents mirror-symmetric shapes simply by a change in sign:


$$\theta(s) \rightarrow \theta(-s)$$

4. Scaling property: Changing the size of a shape simply scales $\theta(s)$ by the same factor. The zero-crossings are unaffected.

$$\theta(s) \rightarrow K\theta(s)$$

10.2 The “2.5-dimensional” sketch

A scheme which David Marr proposed for bridging the gap between 2D image (appearance-based) descriptions and 3D model-based descriptions is called the “2.5-dimensional sketch.” Surface normals are computed and assigned to each point in the image domain, which indicate 3D shape information. Looking at such “pin-cushion” diagrams does effectively convey three-dimensional shape.



10.3 3D Object-centred coordinates. Superquadrics.

Represent solids by the unions and intersections of generalized superquadric objects, defined by equations of the form:

$$Ax^\alpha + By^\beta + Cz^\gamma = R$$

Examples include “generalized cylinders” and cubes (large exponents); prolate spheroids (footballs) and oblate spheroids (tomatoes), when $\alpha = \beta = \gamma = 2$ and when only two of (A, B, C) are equal to each other.


These simple, parametric descriptions of solids, when augmented by Boolean relations for conjoining them, allows one to generate object-centered, “volumetric” descriptions of the objects in a scene (instead of an image-based description) by just giving a short list of 3D parameters and relations, rather like the codon descriptors for closed 2D shapes.

10.4 Deformable parametric models

A powerful approach for representing complex shapes in simple and compact (if only approximate) terms, is the use of *deformable parametric models*. These are especially useful for time-varying objects, such as a human face generating some expression that evolves in time, or a hand gesture. The idea is to find some model (such as the superquadrics) and a parameter set that are fitted to describe the object. They thereby constitute a compact code that can be used for detection and recognition of the object, although fitting such parameters to a 3D object is an “inverse problem” of high computational complexity. If these parameters are then made to evolve in time, as $(A(t), \alpha(t), \beta(t), \dots)$ above, one can encode (compress, generate) an image sequence such as a “talking head” or avatar. This topic unites both computer vision and graphics. The new international MPEG-7 standard for motion image encoding specifies provision for both facial animation and body animation.

11 Perceptual psychology and cognition. Vision as model-building.

Opportunities to learn from biological visual systems, for the design of artificial ones, are not limited to low-level neural mechanisms. Insights from *perceptual and cognitive psychology* are also relevant. Consider the illusory contours:




Through Section 8 we investigated various mechanisms for image analysis, such as filters for spatial forms and edge detectors, and special mechanisms


for handling texture, colour, stereo, and motion information. But how does all of this get “put back together again” into a unified visual percept?

This was a motivating question for a school of research in visual perception called *Gestalt Psychology*, associated mainly with Koffka, Köhler, and Wertheimer in Germany in the 1930s. “The whole is greater than the sum of its parts” might almost be a slogan for the Gestaltists, who sought to understand how meaningful wholes (*Gestalten*) are constructed in a way that seems (introspectively) to precede the analysis of parts and properties. Thus was born the study of perceptual organisation. The Gestaltists enumerated a series of principles collectively called the *Law of Prägnanz* that seemed to underlie our perceptual organisation of form and patterns, based on sub-laws of grouping by proximity, similarity, “good continuation,” and closure (the filling-in of missing parts). More recent developments include the idea of a *process grammar* which models objects and shapes in terms of their morphogenesis (the likely sequence of steps in their evolution from simpler forms). A common theme is that vision is inference, going well beyond the given.


An aspect of perceiving scenes as meaningful wholes, as opposed to an atomistic, literal, or elemental description in terms of individual features, is the grouping of features into a 3D model. A classic illustration of this idea of “vision as model-building” is the Necker cube: a set of 12 planar line segments that are always seen as a 3D solid (a cube); yet having two conflicting (bistable) visual interpretations:



Such bistable percepts are examples of perceptual rivalry: two or more alternative ways to interpret the same visual stimulus. Several more examples are given on the next page: Ruben’s vase/faces; girl/witch; man/rat. The key notion is that percepts are hypotheses: they are top-down interpretations that depend greatly on contexts, expectations, and other extraneous factors that go beyond the actual stimulus. Used in an analogy for scientific revolutions, Thomas Kuhn equated scientific paradigms with percepts, and he cited the familiar example of the duck-rabbit: “What were seen as ducks before the revolution, are rabbits afterwards.”




AMBIGUOUS FIGURES



In the examples on the next page, illusory contours are perceived in places where no actual contours exist. In the upper examples the illusory contours even seem to demarcate a region that is “brighter” than its surrounds; and the illusory contours can even take curvilinear trajectories. But defined by what??

In the lower pair of examples, the circle and the square appear significantly deformed in shape by their context. Are such aspects of your visual system

“bugs,” or “features?” Should such inaccuracies in representing patterns also be designed into machine vision systems, intentionally or epiphenomenally?



12 Lessons from neurological trauma and deficits. Visual illusions.

A further source of some insights about how human vision works at higher levels comes from traumas to brain visual areas, either from injury, stroke,

or carcinoma. From neurological studies that began in the 1920's, largely in response to massive numbers of brain injured soldiers from WWI, Karl Lashley formulated two putative principles of brain recovery from injury:

1. The “Law of Mass Action” (all brain areas participate in all tasks)
2. The “Law of Equipotentiality” (any brain area can do any function)




But counter-examples emerged from specific deficits (especially for learning, memory, visual pattern recognition, and language) that remained after specific traumas. Many cognitive-neurological deficits have been documented; these are interesting to study from a computational neuroscience perspective.

Besides numerous aphasias (loss or impaired use of language, including gestural, prosodic, semantic, or syntactic disorders, often highly specific...), there are many agnosias that can result from specific brain injury. An agnosia is a *failure of recognition*. Examples:




- Facial Prosopagnosia: lost ability to recognise faces; yet vision is normal. Only faces cease to be represented or processed as a special class of object. (Reference: Oliver Sacks' *The Man Who Mistook His Wife for a Hat*)
- Achromatopsia: cortical loss of colour vision; but apparently “normal” achromatic (black-and-white) vision
- Astereognosia: loss of ability to perceive three-dimensionality
- Simultanagnosia: inability to perceive simultaneously more than one thing at a time (e.g. multiple elements in a display)
- Neglect and hemi-inattention syndromes: one side of any object is always neglected. Such patients dress themselves only on (say) their right side, and always bump into things with their left side; and will draw a clock face with all the numbers 1 - 12 in the right half only.
- Xanthopsia: perception that all objects are covered with gold paint (!!)

What kind of a “computer” is the brain, that it can display these types of faults when traumatised? What do these phenomena reveal about the nature of the brain, its architecture, its data structures, and its “algorithms?”

Similar questions can be raised about illusions within “normal” vision. We will examine specific illusions of geometry, size, brightness, and motion. As always, the question for us to ask in connection with Computer Vision is, what do these illusions reveal about visual mechanisms? Should we try to design our algorithms in such a way that they too would “suffer” from such illusions, at least as an epiphenomenon [side-effect] of the desired functionality?




The pattern on the left above has a set of long oblique lines which are in fact parallel, but they appear very non-parallel. Why does the presence of the other short lines so severely distort our judgment of orientation and parallelism? In the visual cortex, there are both competitive and cooperative neural processes, operating over both orientation and proximity; presumably these create the illusion seen. In the examples on the right above, the net effect is a continuous bending of the (in fact straight and parallel) lines, bowing them together in one case and apart in the other.



Competition appears also in the domain of size, as seen in the example above on the left: the central disk in fact is the same size in both cases. The right

is an illustration of the Müller-Lyer illusion: the vertical segment is in fact the same length in both cases, contrary to appearances. But there it is the presence of the other oblique lines that somehow cause the illusion.



Finally, the presence of extraneous cues plays a large role in many inferences. The two boxes on the left above are clearly of different sizes; yet when forced to see them as displaced in depth, we judge them to be in reality the same size. [Illusions of brightness and motion can be found on the course website.]


13 Bayesian inference in vision. Classifiers; probabilistic methods.

It is virtually impossible to perform most computer vision tasks in a purely “bottom-up” fashion. Consider the following images, and how impoverished are the data which must support the task of object recognition!



An important “AI” perspective on vision is that vision is knowledge-driven. In this view, all of the front-end image processing is merely a distraction, if not an irrelevancy. What is really needed for vision is not a lot of theorems involving the 2D Fourier transform of the Laplacian of a Gaussian filter, but rather a good interface to an expert system that stores and indexes knowledge about such things as Dalmatian hounds and the general way that dogs behave when following a scent...

This section reviews the basic ideas behind Bayesian inference, which is a method fundamental to probability theory, statistics, and machine learning. Its purpose is to provide a means for integrating *prior* information (such as general knowledge about the sorts of things that populate the world, their properties and relationships, the metaphysics of objects, etc...) with empirical information gathered from incoming image data. This principle is expressed in the form of a basic rule for relating conditional probabilities in which the “antecedent” and “consequent” are interchanged. The value of this method for computer vision is that it provides a framework for continually updating one’s theory of what one is looking at, by integrating continuously incoming evidence with the best available inference or interpretation so far.



13.1 Decisions under uncertainty.

Most real-world tasks (whose solution requires intelligence) involve degrees of uncertainty. Decision-making under uncertainty is especially characteristic in computer vision. The sources of uncertainty may include:

- the nature of the data or signals available
- the inherent problem of classifying or recognizing them
- the unpredictability of the future
- the fact that objects and events have probabilities
- the uncertainty of causation

- the fact that associative knowledge is only probabilistic
- the inherent incompleteness or imperfection of processing
- possible undecidability of a problem, given all available data
- the “ill-posed” nature of many tasks
- inherent trade-offs such as speed versus accuracy

But despite these realities, decisions are required. The framework to adopt is that, in a sense, *the world consists of probabilities*, and that visual processing really amounts to *computing probabilities and assigning them*.

Examples of decisions-under-uncertainty in vision:

- Medical diagnosis; radiology: Is this a tumor? Does the cost of a possible False Alarm (taking a biopsy, frightening the patient unnecessarily) exceed the cost of possibly missing an early diagnosis? What should you do if the odds are 99% that it is just a benign cyst; but if it is a tumour, missing it now could be fatal?
- Military decision-making: a plane is seen approaching your aircraft carrier very low on the horizon and at high speed. Is it friend or foe? How should the costs of the two possible types of error (shooting down one of your own planes, vs allowing the whole aircraft carrier to be sunk) be balanced against their relative probabilities, when making your decision?

Finally, how can decision strategies be updated by the integration of evidence that arrives gradually over time? How should a-priori knowledge about the probabilities of events in the world be combined with available incoming data?


Statistical decision theory is the study of how to optimize certain measures of performance, given the available data and the decision environment as specified by costs/benefits, a-priori knowledge, speed and confidence requirements.

The Bayesian View

A highly influential formalism for integrating prior knowledge about the world (beliefs being expressed in terms of probabilities) with new incoming data (e.g. an image sequence), or of achieving fusion amongst different and possibly incommensurable forms of data, is that of Bayesian inference.

Bayesianism interprets probability as “degree-of-belief,” rather than as “frequency of occurrence,” and argues for weighing all evidence with all possible (imaginable) interpretations and their associated (estimated) probabilities.

Baye’s rule, named after the 17th-century cleric, Thomas Bayes, is a formalism for combining *prior knowledge or beliefs* with empirical observations. It is at once a theory of explanation, a procedure for the integration of evidence, and a protocol for decision-making. Some aspects of Bayesian interpretation in vision are evident in the way we read the following texts, in which the same letter stimulus is read in completely different ways depending on local context:



We begin with an informal statement of Bayes’ rule for drawing inferences from data. If H represents an hypothesis about the “state of the world” (e.g. the object in an image) and D represents the available image data, then the explanatory conditional probabilities $p(H|D)$ and $p(D|H)$ are related to each other and to their unconditional likelihoods $p(H)$ and $p(D)$ as follows:

$$p(H|D) = \frac{p(D|H)p(H)}{p(D)} \quad (15)$$

For example, a human agricultural expert, or an artificial expert system, has knowledge of the form $p(D|H)$: Given a plant (or a hypothetical disease state) H , there is a corresponding conditional probability $p(D|H)$ of observing certain image data D . However, typically the goal of computer vision and pattern recognition is to calculate just the *inverse* of that conditional probability: given image data D , what is the probability $p(H|D)$ that the hypothesis (of plant or disease state H) is true?

Bayes’ rule (Eq. 15) specifies the formal procedure for calculating such inferences $p(H|D)$, given the observations, the unconditional probabilities, and

the prior expert agricultural knowledge $p(D|H)$. It thereby offers a clean and simple interface between a knowledge base and visual data. A key feature of Bayes' Theorem is that it provides a mechanism for repeatedly updating our assessment of a visual hypothesis as more data arrives incrementally. We can apply the rule recursively, using the latest *posterior* as the new *prior* for interpreting the next set of data. In AI, this feature is important because it allows the systematic and real-time construction of interpretations that can be updated continuously as more data arrive in a time series, such as a flow of images or spoken sounds that we wish to understand.

13.2 Statistical Decision Theory

The Bayesian view focuses on the use of *priors*, which allow vision to be steered heavily by one's *a priori* knowledge about the world and the things which populate it. For example, probabilistic priors can express the notion that some events, objects, or interpretations are vastly more probable than others; that matter cannot just disappear, but does routinely become occluded; that objects rarely change their surface colour; that uniform texturing on a complex surface shape is a more likely interpretation than highly non-uniform texturing on a simple or planar shape; that a rigid rotation in three dimensions is a "better explanation" for deforming boundaries (if consistent with same) than actual boundary deformations in the object itself; and so forth. Being able to integrate formally such learned or even "metaphysical" assumptions about the world is one way in which Bayesian inference facilitates a "top-down" or AI-oriented, expert-system-oriented, approach to vision.

However, in many vision tasks, there may be no useful (or strong) priors. We may need to solve pattern recognition problems purely on the basis of some vector of acquired features from a given object or image; the task is to decide whether or not this feature vector is consistent with membership in a particular class or object category. In this sense, the problem of object identification amounts to a "same / different" decision between the presenting feature vector and one (or more) characteristic class feature vectors, even if we don't have any useful priors about the relative likelihoods of the possible object classes or interpretations.

The degree of match between two feature vectors must be computed and formally evaluated to make a decision of "same" or "different." Almost always, there is some similarity between "different" patterns, and some dissimilarity between "same" patterns. This creates a decision environment with four possible outcomes:

1. Hit: Actually same; decision “same”.
2. Miss: Actually same; decision “different”.
3. False Alarm: Actually different; decision “same”.
4. Correct Reject: Actually different; decision “different”.

We would like to maximize the probability of outcomes 1 and 4, because these are correct decisions. We would like to minimize the probability of outcomes 2 and 3, because these are incorrect decisions (“Type II” and “Type I” errors).




Figure 2: Two-choice Decision Environment of Statistical Decision Theory

We can adjust our decision threshold (become more liberal or more conservative) to reflect the costs and benefits of the four possible outcomes. But adjusting the decision threshold has coupled effects on the four outcomes:

- Increasing the “Hit” rate will also increase the “False Alarm” rate.
- Decreasing the “Miss” rate will also decrease the “Correct Reject” rate.

How can we understand these relationships in a theoretical formalism? How can we optimize the decision-making process?

During WWII, a theoretical framework was developed for understanding such decision environments in the context of radar. It was developed at the University of Michigan and became known as Signal Detection Theory, or also as Statistical Decision Theory. The signals received about the “state of the world” are modelled as arising from two noisy probability distributions. In the pattern recognition context, these two correspond to the class-object relationships of “same” versus “different.” (In the schematic diagram above, the terms “authentic” and “imposter” were used.)

When a decision of “same” or “different” is made, based upon the observed similarity and some acceptability threshold, the probabilities of the four possible outcomes can be computed as the four areas lying under these two distributions to either side of the decision criterion. These four probabilities correspond to the shaded areas in the diagram. The computed error probabilities can be directly translated into a *confidence level* that we can assign to any decision that is made in this formalism. The result of being “liberal” or “conservative” in our decision-making is revealed in the ROC curve, for Receiver Operating Characteristic (the name is derived from radar analysis).




Figure 3: ROC Curve for Two-Choice Decision Environments (e.g. making Yes-No decisions for object classification and pattern recognition). ROC curves reveal trade-offs between error rates.

Each point on the ROC curve represents a particular decision strategy. It plots the relationship between the resulting Hit Rate and False Alarm Rate.

Finally, regardless of where our decision threshold is placed, the fundamental decidability of the decision task (or the detectability of the signal detection task) is measured by the quantity “d-prime” (d'). It is defined as the difference between the means of the two distributions, scaled by the square-root of their average variance (a conjoint standard deviation):

$$d' = \frac{|\mu_2 - \mu_1|}{\sqrt{\frac{1}{2}(\sigma_2^2 + \sigma_1^2)}}$$

where the two distributions are characterized by means μ_1 and μ_2 and standard deviations σ_1 and σ_2 . An improvement in d' can result either from pushing the two distributions further apart, or from making one or both of them narrower. In the ROC curve, d' corresponds to how “bowed” the curve is. The bigger d' is, the better; a pattern recognition problem with high decidability will have a large d' , so the curve approaches the upper-left corner. Any value higher than about 3 is great. The Figure below illustrates $d' = 11.36$ for iris recognition.

These considerations illustrate what might be called the “Primary Law of Pattern Recognition”:

The key factor is the relation between within-class variability and between-class variability. Pattern recognition can be performed reliably only when the between-class variability is larger than the within-class variability.





Figure 4: A powerful decision environment with $d' = 11.36$, illustrating how well-separated same/different distributions lead to highly decidable classification and pattern recognition.

13.3 Bayesian pattern classifiers

Consider a two-class pattern classification problem, such as optical character recognition on the space of just two letters, a and b . We compute some set of features x from the image data (for now we don't care what those features are), and we wish to build a Bayesian classifier that will assign a given pattern to one of two classes, C_1 or C_2 , corresponding to the two letter instances.



Whatever are the extracted features x (maybe as simple as the height/width ratio), after collecting these measurements from a large number of samples of letters a and b , we can plot a histogram of how these measurements are distributed for each of the two classes. In general, these histograms will overlap, as illustrated above right. A particular sample of the value of x might come from either class $C_1 \equiv a$ or $C_2 \equiv b$; but the further to the left it is, clearly the more likely it is to have come from class C_1 , other things being equal.

What do we mean by “other things being equal?” Suppose that instances of class C_2 are 100 times more frequent (more probable) than class C_1 . Would we then still say that, given a slightly smallish sampled value x as indicated above, the letter class is more likely to have been C_1 than C_2 ?

No. Now we need to become Bayesians and take into account baseline rates. Define the prior probabilities $P(C_1)$ and $P(C_2)$ as their relative proportions (summing to 1). If we had to guess which character had appeared without our even seeing it, we would always just guess the one with the higher prior probability. Thus since in fact an ‘ a ’ is about 4 times more frequent than a ‘ b ’ in English, and these are the only two cases in this two-class inference problem, we would set $P(a) = 0.8$ and $P(b) = 0.2$.

For each class separately, we can calculate how likely any particular feature sample value x would be. This is specified by each of the two histograms above, as probability distributions. This gives us $P(x|C_1)$ and $P(x|C_2)$.


Finally, we need to know the unconditional probability $P(x)$ of any measurement value x . We can calculate this by the probability “sum rule:”

$$P(x) = \sum_{k=1}^2 P(x|C_k)P(C_k)$$

Now we have everything we need to apply Bayes’ Rule to calculate the likelihood of either class membership, given some observation x , factoring in the prior probabilities $P(C_k)$, the unconditional probability $P(x)$ of the observed data, and the likelihood of the data given either of the classes, $P(x|C_k)$. The likelihood of class C_k given the data x , is the posterior probability $P(C_k|x)$:

$$P(C_k|x) = \frac{P(x|C_k)P(C_k)}{P(x)} \quad (16)$$

Thus Bayes’ Rule gives us a principled, formal way to perform pattern classifications on the basis of the available data and our knowledge of class baseline rates, and how likely the data would be for each of the classes. We may now plot the likelihoods of each of the classes, as a function of the data x :



We minimise the probability of misclassification if we assign each new input x to the class with the highest posterior probability. Assign x to class C_k if:


$$P(C_k|x) > P(C_j|x) \quad \forall j \neq k$$

Since the denominator in Bayes’ Rule (Eq. 16) is independent of C_k , we can rewrite this *minimum misclassification criterion* simply as:

$$P(x|C_k)P(C_k) > P(x|C_j)P(C_j) \quad \forall j \neq k$$

If we now plot the quantities in this inequality relation as a function of x , we

can see that the minimum misclassification criterion amounts to imposing a decision boundary where the two curves cross each other (arrow):



Because the costs of the two different types of errors are not always equal, as illustrated earlier in the medical example of the biopsy, we may not necessarily want to place our decision criterion at the point where the two curves cross, even though that would minimise the total error. If the decision boundary that we choose is as indicated by the vertical line above, then the total error is equal to the total shaded area. Let R_1 and R_2 be the regions of x on either side of our decision boundary. Then the total probability of error is:

$$\begin{aligned} P(\text{error}) &= P(x \in R_2, C_1) + P(x \in R_1, C_2) \\ &= P(x \in R_2|C_1)P(C_1) + P(x \in R_1|C_2)P(C_2) \\ &= \int_{R_2} P(x|C_1)P(C_1)dx + \int_{R_1} P(x|C_2)P(C_2)dx \end{aligned}$$

Thus the total shaded area is the total probability of error, and obviously we would minimise this (if that were our goal) by putting the decision boundary at the arrow where the two curves cross.

13.4 Discriminant functions and decision boundaries

If some set of functions $y_k(x)$ of the data x are constructed, one function for each class C_k , such that classification decisions are made by assigning an observation x to class C_k if

$$y_k(x) > y_j(x) \quad \forall j \neq k,$$


those functions $y_k(x)$ are called discriminant functions. The decision boundaries between data regions R_j and R_k are defined by those loci in the (normally multi-dimensional) data x at which $y_k(x) = y_j(x)$. A natural choice for discriminant functions would be the posterior probabilities:

$$y_k(x) = P(C_k|x)$$

Equivalently since the denominator $P(x)$ in Bayes' Rule is independent of k , we could choose

$$y_k(x) = P(x|C_k)P(C_k)$$

or any monotonic function of this, since the decision boundaries would remain the same. The following figure illustrates how in even just the case of two-dimensional data, the decision boundaries separating four Gaussian densities (corresponding to four classes) can be rather complex.



15 Vision as a set of inverse problems. Regularization.

Many goals of vision can be described as “inverse problems,” because inferences from data are needed to derive models of the world properties that projected those image data. Such inverse problems are formally ill-posed. A powerful way to overcome the inherently ill-posed nature of many visual tasks is to re-formulate, or regularise them as *optimisation problems*, which can be approached essentially as problems of calculus because the solution to the task corresponds to finding an extremum. The method involves specifying some *functional* (typically an integral operator) consisting of two terms:

1. a “data” term, in which image parameters are represented;
2. a “smoothness” (or regularisation) term, embodying world constraints.

These two terms are combined by some weighting parameter λ specifying their relative importance, i.e. the degree to which either the data dominates, or the background assumptions dominate, the solution.

16 Face detection, recognition, and interpretation



Louis Léopold Boilly, Reunion de Têtes Diverses

The goal of detecting faces and recognizing their identity has long been one of the “Holy Grail” problems in computer vision. It is a hard problem for all of the reasons we have encountered that generally make computer vision hard:

- Faces are surfaces on 3D objects (heads). Therefore the images they project depend on the perspective angle between object and camera, the rotation of the object around its own axes, and the illuminant.
- Facial surfaces have relief, and so parts (e.g. noses) can occlude other parts. Hair can also create random occlusions and shadows.
- Surface relief causes shading and shadows to depend upon the angle of the illuminant, and whether it is an extended or a point source.
- Faces have variable specularities (dry skin may be Lambertian, oily or sweaty skin may be specular). As always, this confounds the interpretation of the reflectance map.
- Parts of faces can move around relative to other parts (eye movements; lip movements; eyebrows and winks).
- Humans put things on their faces (e.g. glasses, cosmetics, cigarettes) and change their facial hair (moustaches, eyebrows). They also use their faces as organs of expression, and so the surface isn’t even rigid. (Ideally one would like not only to be able to detect and recognize faces, but also to interpret and classify their expressions.)

16.1 Issues in detecting, recognising, and interpreting faces

As usual, this domain of computer vision raises questions such as:

1. What is the best representation to use for faces?
2. Must this be treated as a 3D (object-based) or 2D (image-based) problem?
3. How can invariances to size (hence distance), location, pose, and angle of view be achieved? (A face should acquire the same representation under such transformations, for matching purposes.)
4. What are the *generic* (i.e. universal) properties of all faces that we can rely upon, in order to reliably detect the presence of a face?
5. What are the *particular* features that we can rely upon to recognize the identity of any given face?
6. What is the best way to handle “integration of evidence,” and incomplete information, and to make decisions under uncertainty?
7. How can we handle the transformations that can occur in a given person’s face, either through natural, or unnatural means?



The above two images show the same person. Likewise the four below, despite apparent changes in gender, race, and *Gattungswesen* (species-being).




16.2 The Fundamental Problem.

The central issue in pattern recognition is the relation between within-class variability and between-class variability. These are determined by the degrees of freedom spanned by the pattern classes. Ideally the within-class variability should be small and the between-class variability large, so that the classes are well separated. In the case of encoding faces for identity, one would like different faces to generate face codes that are as different from each other as possible, while different images of the same face should ideally generate similar codes across conditions. Several recent investigations of how well this goal is achieved have studied the invariances in face coding schemes under changes in illumination, perspective angle or pose, and expression. Their results have tended to show that there is greater variability in the code for a given face across these three types of changes, than there is among the codes for different faces when these three factors are kept constant. Since reports documenting performance of particular face recognition algorithms have often been based upon trials in which these factors (pose, illumination, and expression) were held artificially constant, the performance statistics in real-world settings have been very disappointing by contrast, with error rates approaching 50%.




The array of images above shows how dramatic are the effects of even only a change in illumination *direction*. Facial expression remains exactly the same. Going across the columns from left to right, the illumination changes from frontal to side; and going down the rows, it changes in elevation. If you compare the 3 images in the last column on the right, it seems almost inconceivable that any means could be found to represent these as images of the same person.

Earlier (on page 4) we saw how dramatically a change in pose angle affects image appearance, even though the expression and illumination remained the same. Appearance-based algorithms for face recognition still tend to judge different faces in the same pose as more similar than identical faces in different poses. Finally now, the images below show how much a given person's face (in each row) can change when she is using it socially as an organ of expression.



For comparison now, when we examine images of *different* faces seen under fixed illumination and with neutral expressions, their (between-class) variability seems tiny compared to the same-person (within-class) variabilities above associated with changes either in illumination or in expression:



When there is variability across two or more dimensions (let us say both face identity and facial expression, as in the images below), then discriminability might benefit from variability *within* a class of the other dimension, but not from variability *between* classes of the other dimension.



For example, facial expressions are more reliably distinguished if there is large variation among the different expressions generated by a given face, but small variation in how a given expression is generated amongst different faces. The consequences of within-class and between-class variability, for single dimensions and across them, are noted in the following table:

Task	Within-Class Variability	Between-Class Variability
Face detection (classes: face / non-face)	bad	good
Face identification (classes: same/different faces)	bad	good
Facial expression interpretation (classes: same/different faces)	good	bad
Facial expression interpretation (classes: same/different expressions)	bad	good

Many of these forms of variation in facial appearances were captured in the painting by Boilly, *Reunion de Têtes Diverses*. In characterizing the within-class variability and the between-class variability of faces, it is clear that (at least over time), the variability of any given face can easily outstrip the variability among contemporary faces. No one would deny that young babies look far more similar to each other, than each does to the adult that it grows into.


Even when all other factors such as pose angle, expression, illumination, and age are held constant, we can distinguish those aspects of facial variation that are genetically inherited (“genotypic features”), from those that primarily reflect development, aging, or environment (“phenotypic features”). Persons who are genetically identical would share all their genotypic features, such as gender, blood group, race, and DNA sequence, whereas phenotypic features can be shared among different individuals only by chance, according to their associated probability distributions.

One source of evidence about the genotypic/phenotypic ratio of facial variation arises from identical (monozygotic) twins. Obviously any pair of twins are always matched in age. Each twin’s appearance changes over time in the normal dramatic way, yet the pair usually remain strikingly similar to each other in appearance at any age. Nobody would deny that identical twins look vastly more similar to each other than unrelated persons do. Since such twins are genetically identical, their similarity in appearance serves to calibrate the extent of *genetic penetrance* for facial structure.

A further, but secondary, indicator of the genetic penetrance of facial appearance is provided by persons who share only 50% rather than 100% of their genes. These include fraternal twins, full siblings, double cousins, and a given parent and offspring. Occasionally the latter pairings have virtually indistinguishable appearance at a similar age, such as Robert F. Kennedy and his son Michael Kennedy in adulthood.

Interestingly, a major part of the computational load of the brain is concerned with “social computation,” a large part of which involves identifying and interpreting faces. It is generally accepted among ethologists and neuroscientists that the main evolutionary pressures that led to the large brains of primates were not “engineering oriented” pressures such as learning to use tools, but rather the demands of sexual competition. Included in those task demands are: seduction; betrayal; assessing power hierarchies and your own place within them [“who is the alpha male here?”]; manipulation of others’ intentions and desires; and interpreting those within others [i.e. the “other minds” problem].

Neurons in some 30 distinct visual areas in primate brains have been found which are selective to faces. Many of these neurons are further tuned not only for the familiarity of a face, but also for its expression, e.g. friendly versus threatening. The following sketch of the rhesus monkey brain identifies numerous areas in the temporal cortex in which face sensitive neurons have been found. Some of them respond mainly on the basis of whether the eyes in a face are averted away from the monkey, or looking straight at it, which is a key threat gesture among many primates. (Try making too much eye contact with the “alpha male” gorilla in a zoo sometime, and you will see how angry he gets with you.) Even we humans are exquisitely sensitive to the amount of eye contact, especially between the genders. Either too much, or too little, is quickly noted (at least unconsciously) and “interpreted.”



16.3 Face detection

Paradoxically, face detection is a harder problem than face recognition, and the performance rates of algorithms are poorer. (This seems paradoxical since detection must precede recognition; but recognition performance is measured only with images already containing faces.) Approaches to face detection still use generic templates, spanning multiple scales (for faces of different distances, hence sizes) and poses. One of the most powerful features aiding face detection is the rather special hue composition of human skin, which is matched by few other types of surfaces. (Racial differences correspond only to variations in saturation, due to differential melanin density, but not to differences in hue.) If skin tone is ignored, most current algorithms for face detection will find lots of faces when “aimed” merely at a nice Persian rug...


16.4 Two-dimensional (appearance-based) approaches to face recognition

Early attempts at face recognition were based on simple “template-matching” on a pixel-by-pixel basis. Although primitive, these methods have re-surfaced now as performing about as well as anything else. Of course, they do require some mechanism for size-invariance to compensate for distance; the brute force approach would require storing virtually an infinite number of templates for different sizes, and so accurate size normalisation is needed instead. Variation in pose angle is normally handled today by storing not just one template for each face, but rather a (3×3) array or a (4×4) array for the different pose angles. Some older 2D approaches to face recognition also sought to enumerate a long vector of facial features, such as measurements of the distances between many points on the face. These include all mutual distances between distinct points such as the pupils, the boundary of the chin, the cantus (corners) of the eyes, corners of the lips, etc. However, these methods were generally just not very robust, despite the combinatorial explosion of possible measurements and metrics.

One of the most prominent approaches to face recognition as a two-dimensional, appearance-based problem is the *Eigenfaces* approach (Kirby and Sirovich; Turk and Pentland). This involves performing a complete Karhunen-Loeve Transform of a large database of faces (e.g. 8,000) to extract the principal components, i.e. the 20 or so main 2D factors along which different faces differ from each other. These may be expressed as the eigenvalues on the eigenvectors (eigenfaces), which form a new abstract kind of feature vector on which to base recognition. Performance is often in the 90% to 95% range. However, a limitation in this method is that many of the principal components simply extract variations due to shading caused by variations in the angle of illumination! Other high-order principal components are extracting variations in the outline of the face due to small errors in size (distance) normalisation. Thus the method is essentially a 2D representation of faces, and lacks invariances to illumination or pose angle, or any real size invariance.

The Eigenfaces approach exemplifies a typical strategy in computer vision, which is *projection to a low-dimensional subspace*. The critical variation in a pattern recognition problem is captured in some low dimensional set of basis vectors, such as the 20 most important “eigenfaces” to emerge from Principal Components Analysis (PCA) of a dataset that is regarded as representative of the problem domain. Those are then treated as basis vectors, in terms of which any face is represented by some linear combination. The weights, or coefficients, that specify each such linear combination are the eigenvalues; in effect they indicate the “relative presence” of each of the eigenfaces within


any given presenting face. They are computed simply by taking the inner product of the presenting face image with each of the eigenfaces. Because the eigenfaces which emerge from the PCA are, by construction, orthogonal, this is a relatively rapid computation. The projection coefficients obtained serve also as expansion coefficients, since this specified linear combination of roughly 20 eigenfaces will superimpose into a very close approximation to the face in question. Thus a face is effectively represented by a small set of numbers: the eigenvalues. Such a “face code” is extremely compact, and databases can be searched very rapidly since the description of each face is a simple feature vector of only 20 numbers. The following picture illustrates 15 eigenfaces computed from PCA as basis vectors, and their linear combination to superimpose into the face in the top left.



16.5 Wavelet approaches to face recognition

Recently much interest has developed in *wavelet* representations of faces. This idea can be applied in either a 2D or a 3D fashion, either to represent image structure or to represent the surface of the face as a 3D model. Because wavelets are localized, they can track changes in facial expression in a local way. This approach essentially treats a face as a kind of *texture*, made up of

various undulations in various positions, sizes, and orientations but without incorporating explicit models for the individual parts of faces. Remarkably, the major facial features such as eyes, lips, and noses can be extremely well represented by just a handful of 2D Gabor wavelets, as can the entire face, as was noted earlier.



To allow for the deformations associated with changes in pose angle or changes in expression, these wavelet mappings (called “Gabor jets”) are often placed on a deformable graph which tolerates distortions relative to fiducial points. Matching can be performed on such distorted graphs in a way that compensates for a limited range of deformations. Performance of this approach is comparable to that of the Eigenfaces. The computed feature vectors can be local Gabor wavelet amplitude or phase information.

Phase-Quadrant Demodulation Code





Figure 5: An encoding approach for faces (and other patterns) by phase-quadrant quantization of complex-valued 2D Gabor wavelet projections of image structure. Bits in the “face code” are set by the quadrant in which the phasor lies for each aspect of facial structure.

16.6 Three-dimensional approaches to face recognition

Today all focus is on modelling faces truly as three-dimensional objects, and performing recognition in those terms, in order to achieve invariance both to pose angle and illumination geometry. Of course, this requires solving the massive ill-posed problems of inferring shape from shading, interpreting albedo vs. variations in Lambertian and specular surface properties, structure from motion, etc. On page 4 we examined how difficult this problem is, and how remarkable it is that we humans seem to be so competent at it. The synthesis of vision as model-building and graphics, to perform face recognition in object-based terms, rather than appearance-based terms, is the future of this field.

16.7 Interpreting facial expressions

Finally, much current effort is devoted to recognising facial expressions, not only as static aspects of pictures but as dynamic sequences, unfolding in time. Motion energy models are used to extract motion signatures from specific parts of faces, and to classify these as expressions. This task also entails vision-as-inverse-graphics to construct models based upon knowledge of the human facial musculature and behaviour. The future of human-machine interaction will rely upon this interpretive aspect of “affective computing.”






Figure 6: Facial expression interpretation as an “inverse problem,” using generative models. Knowledge of the facial musculature (7 main agonist / antagonist muscle pairs), together with intuitive ancestral knowledge, helps to interpret facial expressions by modelling the motor command sequences that would generate them.