The Abdus Salam
International Centre for Theoretical Physics

United Nations
Educational, Scientific and
Cultural Organization

IAEA
International Atomic Energy Agency

2017-3

Preparatory School to the Winter College on Optics in Environmental Science

26 - 30 January 2009

Laboratory Exercise - Digital Visible and Infrared Imaging

Brydegaard M.
Lund Institute of Technology
Sweden

Pettersson S.G.
Lund Institute of Technology
Sweden

Laboratory Exercise - Digital Visible and Infrared Imaging

M. Brydegaard and S.G. Pettersson
2008

Contents:

Preparations	**2**
Imaging	**2**
Magnification	3
Resolution	4
f-number	5
Colour fundamentals	**5**
Physiological and evolutionary aspects	5
Origin of colours	6
Additive and subtractive colours	7
Colour spaces	8
Cartesian RGB and CMY	8
Conical HSV, spherical spaces and unit analysis	9
Imagers	**10**
CCD - Charge coupled devices	12
CMOS - Integrated imagers	13
Format and compression	**14**
Irrelevancy and redundancy	15
RAW – Bitmap	16
GIF - Graphical interchange format	16
JPEG - Joint Photographic Expert Group	16
Images in Matlab	**16**
Loading, Presenting	16
Representation and structure	17
Reshaping and restoring	17
Spectral modelling and contrast functions	18
Image manipulation	**21**
Intensity domain methods	21
Contrast stretching and normalization	21
Compression of dynamic range	21
Histogram equalization	22
Spatial domain methods	24
Spatial filtering	24
Frequency domain methods	25
Thermography	**25**
Black body radiation, Planck distribution and emissivity	25
IR detectors	26
IR imaging	27
IR spectra	28
Gas correlation	29
References	**30**
Lab assignments	**31**

Preparations (to be solved before the laboratory exercise):

1. Calculate the MTF for the optical system at the second frequency in figure 3.
2. What is the optimum wavelength for thermography at room temperatures (293 K)?
3. Why is it impossible to describe all light spectra with RGB-values?
4. What is the colour described by the RGB-value (0.2,0.2,0.2)?
5. Get familiar with the Matlab commands mentioned in the Matlab chapter, use them once each and use the *help* on the commands. You can practise on any colour picture of your own choice.
6. What units does a RGB colour space have on the axis? What units does a spherical colour space have on the axis?
7. What does the colour space of a bird look like?
8. 1D histograms are projection of the values of data points along one axis. Why can the data from a colour image not be fully described from three 1D histograms? What does a 2D and a 3D histogram look like?
9. Write a pseudo Matlab code which creates a 2D histogram from two equally long vectors.

Imaging

Imaging means that an optical system handles the light field generated by an object in such a way that a new light field is formed creating an image. The ideal optical system would be a system where the image is equal to the object. However, limitations in all optical systems introduce some errors and differences that are dependent on the system used. The optical system must have a finite resolution, there may occur some magnification, the colour content of the image may be changed or the image will be deformed. The image is also never brighter than the original object if no active optical element such as an image intensifier is used. For a three-dimensional object the image obtained on a flat detector is a projection of the three-dimensional image onto a plane. Object points out of focus will generate spots on the detector with a size that is dependent on the distance between the image point and the image plane; with this in mind also consider that most imaging systems do not have the same focus distance for all wavelengths (chromatic aberration).

We will here only discuss simple image formation with a plane object and a flat image detector such as a CCD camera. Such a case is described by a linear transformation of object points to image points.

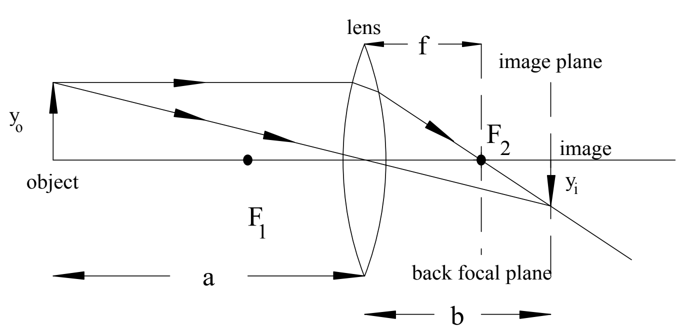

Figure 1. Geometry for simple image formation.

In figure 1 the image formation by a simple lens is illustrated. It is well known from ray optics that the object distance a, the focal length of the lens $f(\lambda)$ and the image distance b are connected by the following relation

$$\frac{1}{a}+\frac{1}{b}=\frac{1}{f(\lambda)}$$

It is important to keep in mind that this relation is a first order truncated solution to the actual problem, and it is valid for "small" incident angles and "thin" lenses.

Magnification
The transversal magnification of an optical system is defined as the ratio of the image size y_i to the object size y_0 or from figure 1:

$$M_T=\frac{y_i}{y_o}=-\frac{b}{a}.$$

The negative sign means that the image is upside-down compared to the object. If a is long compared to b, then b is nearly equal to f which means that the image is formed near the focal plane of the lens.

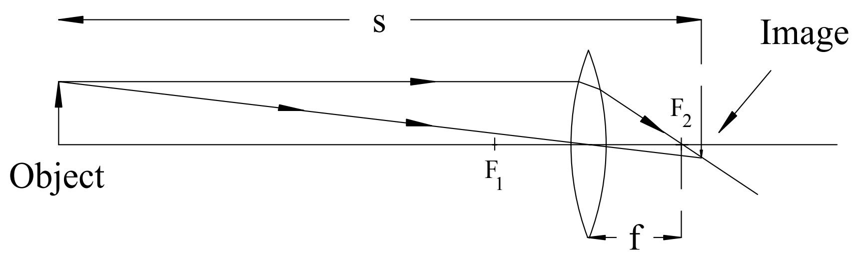

Figure 2. Image formation with long object distance.

Suppose we wish to find the magnification obtained with a camera aimed at a distant object (cf. figure 2). If the centre of the complicated lens system is not easy to find we can measure the distance s between the object and the image plane of the lens (≈ the focal plane). We can then approximate a with $s - f$, b with f and calculate $|M_T|$ as $f/(s - f)$. As an example we put a = 2 m and f = 20 mm. Then $b \approx f$ and $|M_T| \approx 1/100$. If $|M_T|$ = 1/100 (exact) then $b = 101f/100$ so this shows that b is approximately equal to f.

Resolution

Subject to discretization:	Light intensity	Light energy	Space	Time
Domain:	Dynamical -	Spectral -	Spatial -	Temporal -
Discretized by:	Bits	Spectral bands	Pixels / Voxels	Frames
Resolution:	Dynamic -	Spectral -	Spatial -	Temporal -
Res. limited by:	Signal to noise ratio / photons	Channel / illumination bandwidth	Point spread function	Exposure time / flash envelope
Range:	Dynamic -	Spectral -	Field of view	Recording time

Table 1: Comparison terms associated with discretization along various domains.

Due to lens aberrations, diffraction and finite number of pixels the resolution obtained in the image plane of a lens is limited. To measure the resolution one can find the distance between the images of two point sources as they are resolved (Rayleigh's criterion). A better method is to determine the MTF (Modulation Transfer Function) of the system. For different periodic gratings as object functions the MTF determines the visibility of the resulting image functions. The MTF is the ratio of the output modulation to the input modulation where the modulation m is defined as

$$m = \frac{I_{\max} - I_{\min}}{I_{\max} + I_{\min}} \ .$$

In figure 3, three different sinusoidal object functions I(x) (with different spatial frequencies f_x in the x direction) are shown together with the corresponding images obtained in an optical system. This gives three points on the MTF curve (cf. figure 4) for the system.

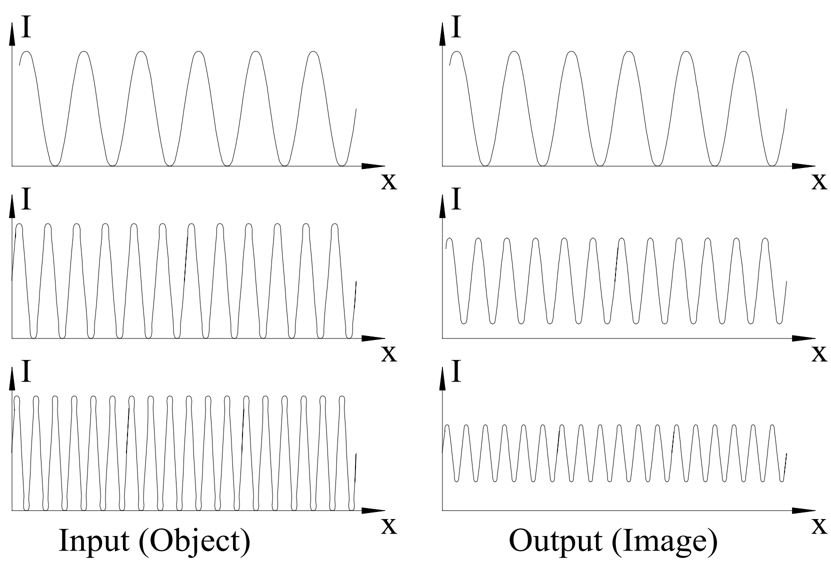

Figure 3. Object functions and image functions in an optical system.

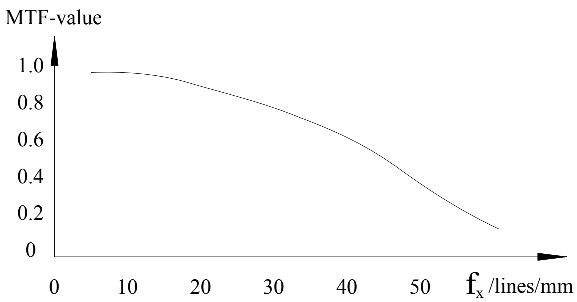

Figure 4. MTF curve for the optical system

From the MTF curve we can estimate the resolution to be about 40 lines/mm (MTF value = 0.5) but the curve gives a more precise description of the optical system than this number.

The *f*-number

The irradiance of the image is determined by the object irradiance and also on the speed of the lens or optical system. This speed is given by the f-number, which is the ratio of the focal length f of the lens divided by the aperture of the lens D. An $f/1.4$ lens is said to be twice as fast as an $f/2$ lens. As the flux density at the image plane varies as $(D/f)^2$ the f-numbers are normally given with a factor $\sqrt{2}$ between consecutive numbers as: 1, 1.4, 2, 2.8, 4, 5.6, 8, 11, 16 etc.

Colour Fundamentals

Physiological and evolutionary aspects

Basically, colours form a concept describing how we humans perceive a given spectrum of visible light between 400 nm and 700 nm. The maximal sensitivity of our vision is centred close to the peak emission of the black body radiation of the sun (which makes a lot of sense because the vision evolved in an environment illuminated by the sun). We lack sensitivity for much shorter wavelength because there is not much natural illumination from the sun passing the atmosphere. Ozone absorbs below 300 nm and air absorbs below 200 nm. However, several insects, birds and reptiles possess ultraviolet sensitivity typically around 360 nm. In the long-wavelengths end of the visible spectrum we lack sensitivity beyond 700 nm. And for much longer wavelengths water begins to absorb. (Our eyes are made of water and they probably evolved in a water environment too). Evolution of snakes, has however come up with pit organs; air filled pin hole camera like devices for imaging with sensitivity between 8-12 μm.

The human eye contains two kinds of photoreceptor cells, the rods and cones. Rods and cones contain various visual pigments called chromoproteins (combinations of large proteins called opsins and a cartenoid pigment called retinene) which are activated by specific wavelengths of light. The activation of the visual pigments then stimulates an action potential along the nerve fibres trailing from the ends of the rods and cones to the optic nerve which, in turn, carries the message to the brain to create an image. Roughly 125 millions of rods cells are intermingled nonuniformly on the retina. The ensemble is exceedingly sensitive, performing in light too dim for the cones to respond to, yet it is unable to distinguish colour. In contrast, the ensemble of 6 or 7 million cones provides spectral resolution. It performs in bright light giving detailed, coloured views but is fairly insensitive at low light levels. The cones are of three kinds, thus providing three colour channels, or primary colours, each responding to a specific portion of the visible spectrum (figure 5). Other species than Homo sapiens have different sensitivity curves. Bulls have only two discrete colour channels while bird and reptiles might have four. Some biological vision system might even have a more continuous sensitivity.

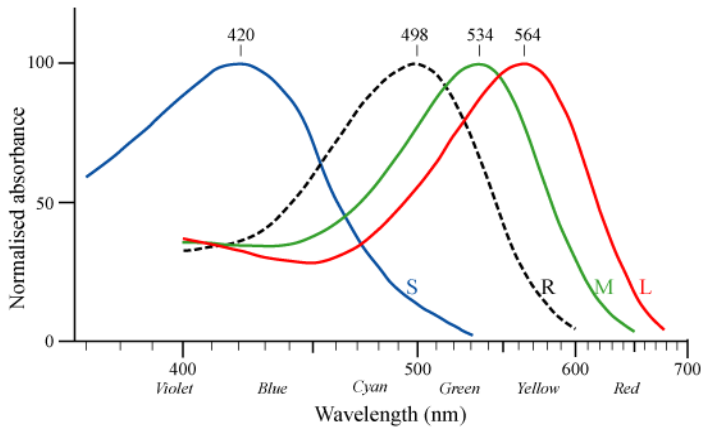

Figure 5. The sensitivity of the different pigment in the eye[7].

The red-sensitive or ρ-cones respond over a range of about 540 to 760 nm, with a peak around 610 nm (orange-yellow). The green-sensitive or γ-cones respond over a range of about 470 to 650 nm, with a peak around 565 nm, considerably overlapping the range of the ρ-cones. The blue-sensitive or β-cones, of which there are fewer than the other types, respond over a range of about 380 to 540 nm, with rather less overlap. There is also a small hump of sensitivity of the ρ-cones, centred on 440 nm, which enables us to distinguish blue from violet.

The contribution ***U*** to a certain cone type, or colour channel, ***k*** is given by:

$$U_{channel(k)} = \int_0^{\infty} I(\lambda) S_k(\lambda) d\lambda$$

where ***I*** is light intensity and $\boldsymbol{S_k}$ is the sensitivity for channel ***k***.

A light spectrum contributing equally much to channel β, γ and ρ ($u_\beta = u_\gamma = u_\rho$) gives the sensation of white light. This implies that a white lamp designed for humans must be expected to give a completely different response to e.g. birds. One can apply the same paradigm of colour channel contribution to any kind of photo electronic detector, including spectrometers. Obviously the last mentioned have much more colour channels, much narrower channels and extend to non visible ranges such as UV or IR. Since ***U*** is proportional to the incident light intensity on the detector, it makes sense to assign an intensity unit to ***U***.

A number of physical phenomena such as reflection, scattering, absorption, fluorescence and interference can change the spectral content of emitted light from an object with respect to the spectral content of the illumination light. This property of changing the light spectrum is considered to be the colour of objects. Sometimes the term is more applicable than other times; what is the colour of a tomato? What is the colour of a compact disc?

The classical idea of the red colour of a tomato does not arise from the reflected light for a simple reason; Specular light reflected at a surface cannot know anything about the underlying tomato volume since it has never been into it. The specular reflex of a tomato has similar colour as the illumination; it is the little white spot on a tomato picture. The light which penetrates into the tomato volume will scatter around multiple times and a fraction of the light will escape the surface again. The probability for escape for blue and green wavelengths is less than for the red wavelengths, which is why the tomato appear red to us. The probability of backscattered light is mainly governed by absorption and scattering probabilities. While structural and surface colour maintain the polarization of the light source, light with has experienced multiple scattering in an object looses its polarization. In this way polarization filters can be used to separate specular and backscattered light, carrying chemical information in terms of influence from the absorption. The object colour property varies obviously from object to object, depending of the phenomena governing light interaction.

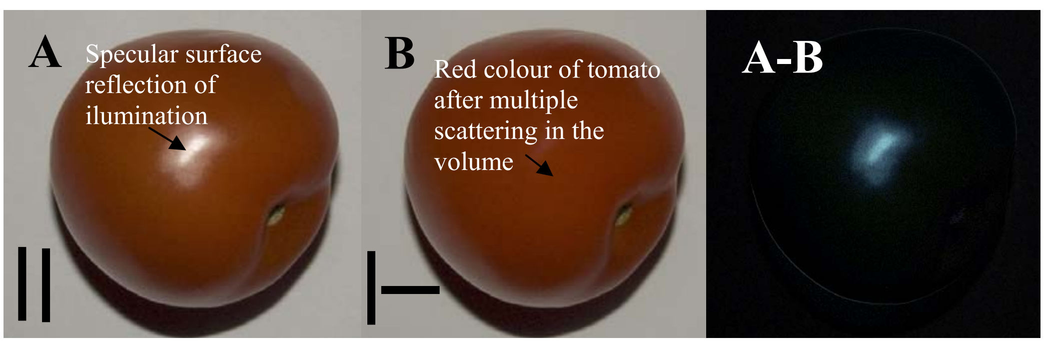

Figure 6: Polarization can reveal origin of change in reflected light spectrum. A picture with parallel polarization filters in front of objective and flash, B picture with perpendicular filter, only allow multiple scattered light to pass. A-B contribution from specular reflexes.

Additive and subtractive colour

It is clear that we apply the diffuse term colour both to the colour of light, and to colour we paint on the wall, and it becomes crucial to distinguish between mixing light emission spectra (additive colour mixing) and light absorption spectra (subtractive colour mixing). When projecting an image on a screen it makes sense to describe it in the primary additive colours, red, green and blue, whereas when printing an image on a paper, it makes more sense to describe the picture in terms of primary subtractive colours, yellow, magenta and cyan.

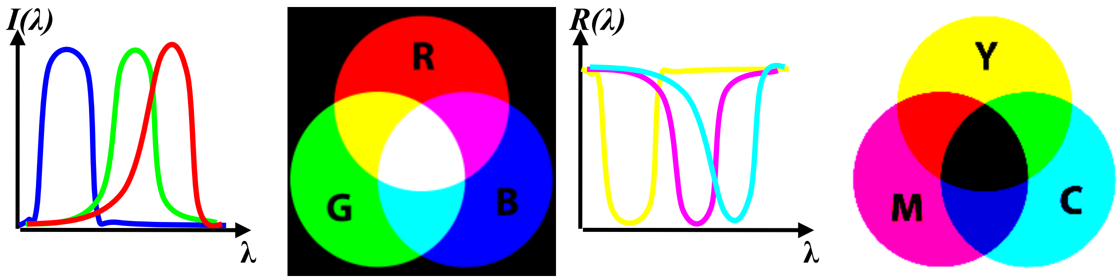

Figure 7: Left: a red, green and blue flash light illuminating a white wall. Right: yellow, magenta and cyan paint on a white sheet of paper illuminated by white light.

Colour spaces

Since we happen to have just three discrete colour channels in our eyes, it is somewhat fruitful to think in 3D colour "spaces". Observations (pixels) of objects in an image can be described spectrally by a position in such colour spaces. We can equally well think of a spectrometer colour space, the difference is just that we may have 4000 colour channels and a 4000D colour space. Similar objects in the image in term of spectral response will be represented by clouds or clusters in such colour spaces.

RGB colour space

The Cartesian red, green and blue (RGB) colour space is arranged in such a way that each axis represents the value from a certain colour channel. The axis are normalized and can have values between 0 and 1 (none and full intensity) for each colour channel, within this cube all possible combinations of the colour channels can be expressed. Images in the RBG colour model consist of three independent image planes, one for each colour. When fed into a RGB monitor, these three images combine on the phosphor screen to produce a composite colour image. Most colour cameras used for acquiring digital images utilise the RGB format.

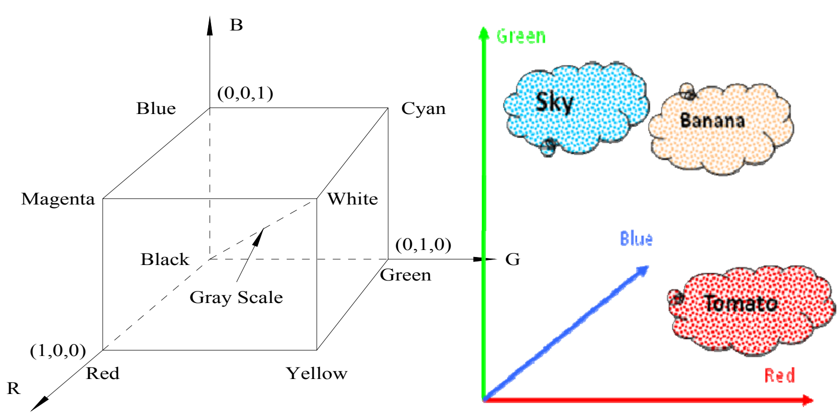

Figure 8: Left: positions in the Cartesian colour space represent different colour nuances. Right: Different object represented by clusters in the RGB colour space.

The CMY colour model
As indicated previously, cyan, magenta and yellow are the secondary colours of light or, alternatively, the primary colours of pigments. For example, when a surface coated with cyan pigment is illuminated with white light, no red light is reflected from the surface. Thus, cyan subtracts red light from reflected white light, which itself is composed of equal amounts of red, green and blue light.

As most printers and copiers (they deposit coloured pigments on paper) require CMY data input, a conversion of the RGB data to CMY data is necessary. This conversion is performed using the simple operation:

$$\begin{bmatrix} C \\ M \\ Y \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} - \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

where, again, the assumption is that all colour values have been normalised to the range [0, 1].

Conical HSV and spherical spaces and unit analysis
As known from basic maths, the same spatial point can be expressed in many different coordinate systems; this is of course also true for colour observations. One common colour space is the cylindrical or conical; Hue, Saturation, Value (HSV) space. Here colours are represented by a hue determining the colour, a saturation determining how grey or how coloured a pixel is, and finally, a brightness value, indicating the light intensity. We understand that nothing changes in the information; only the axes have different interpretation. One particularly interesting colour space is the spherical colour space, since just one axis carries intensity units, while the remaining angles just indicate the relations between the RGB values. In many scenarios illumination intensity is not the same over pictures, by rejecting intensity axis and considering only unit-less values we can reduce data and avoid several problems.

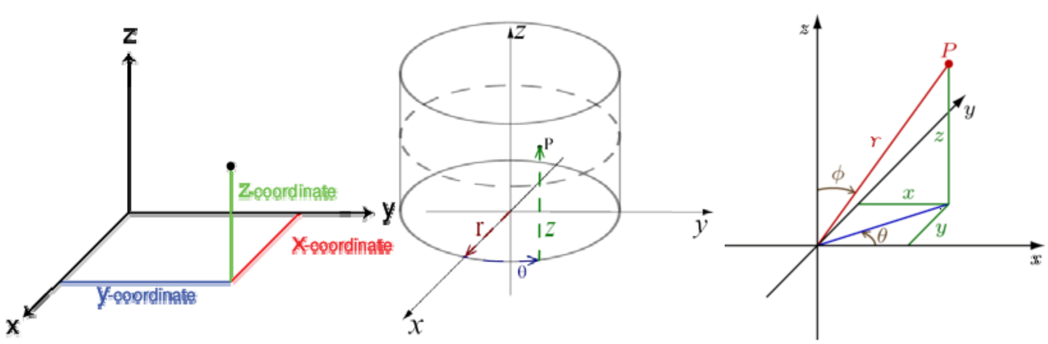

Figure 9: The same colour represented in Cartesian, cylindrical and spherical coordinates.

Since it is not a spectral property of a carrot how much light one shines on it, advantages can be taken by rejecting intensity.

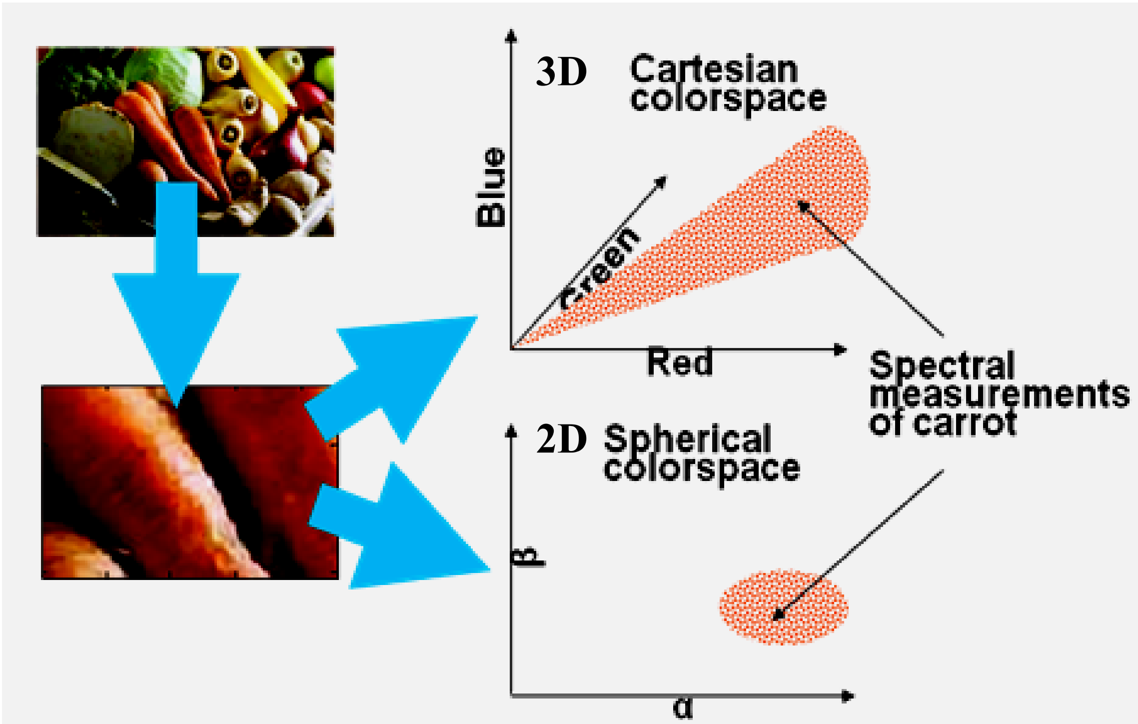

Figure 10: Illustration on how carrot pixels are represented by a cone in the RGB space but only a circle in spherical coordinates when intensity in the r component is rejected.

Imagers

Imaging devices consist of arrays of light detectors; we refer to them as *pixels*. We will discuss two technologies; CCD and CMOS imagers. The performance of light detectors in general is given by the sensitivity, **S**(λ). If the incident light on the detector is **I**(λ), the signal from the light detector is given by:

$$u = \int_0^\infty I(\lambda)S(\lambda)d\lambda$$

where **I** is light intensity and **S** is the sensitivity.

Few systems are perfect and neither are light detectors. Not only light gives rise to signal readout. Noise is a general concern, and the dominant noise sources are referred to as dark current, thermal noise and shot noise, depending on the detector type. Noise can be reduced considerably by lowering the temperature, but cooling systems take up space and energy both in the application and in the budget. Worst of all, cooling the detector yields condensing of water, which distorts both measurement and equipment. Adding the noise **w** the detector now has the following read out:

$$u = \int_0^\infty I(\lambda)S(\lambda)d\lambda + w$$

Furthermore, CCD and CMOS are accumulating imagers where, the signal accumulates during an exposure (alternatives do exists). The signal after an exposure is proportional to the exposure time Δt assuming that the imaged scenario and the illumination was constant over the exposure time.

$$U = \int_{\Delta t} u(t)dt$$

This kind of time integration can alternatively be performed digitally in the signal processing. Such approaches yield slight differences in noise characteristics. In some research where exceptionally fast imaging is required, one can take advantage of photo diode arrays, which provide continuous signal in time.

Accumulating imagers have a maximum limit of accumulated signal. This is generally referred to as *full well capacity*. If one desires to make an image of a still standing rain cloud, one could arrange an array of buckets (pixels) on the ground and derive the shape and density of the cloud by measuring the level in each bucket after a certain rain *exposure time*. If all buckets are put in at one instance of time and collected in another instance of time, it is referred to as a *global shutter*. If buckets are put one by one by a man, and later collected one by one in the same orders as they were put by another man we call it a *rolling shutter* (Meingast, 2005). In the last case the exposure time correspond to the time a single bucket was standing in the rain, not the time for producing the entire image. The two shutter techniques will give the same result as long as the cloud remains still during exposure. The bucket size would be the *full well capacity*. The ratio between bucket-filled area and the total area is referred to as the *fill factor*. If the cloud in some place is particularly dense the rain will fill up the bucket completely at certain places. We refer to this phenomenon as *saturation*. Saturation destroys the relation between the buckets, since we have no way of knowing if the cloud was just slightly denser at this spot or if it was many million of times denser for some reason. Of course we can change the exposure time and take several pictures - one of the million times bigger spike and one of the rests of the cloud, with the saturated bucket. However, this will not be possible on the CCD since charge from one pixel will spread to neighbouring pixels. Think of it as water not only floating over the bucket side and wasting the water, but ever distorting the neighbouring pixels by spilling water into those buckets as a result of the saturation. This phenomenon is referred to as *blooming*. Now, let us imagine that every bucket has centimetre marks to quantify the rain levels. Converting the levels into cm would be called *discrimination*, and the space between the marks would be the *dynamic resolution.* A level change smaller than ½ cm will be neglected. *Dynamic resolution* is mostly specified in bits or dB. E.g., 8 bit resolution means discrimination in 256 levels, between the determined min and max value, corresponding decibels would be $20*\log_{10}(2^8)$=48dB. New imagers might be *logarithmic;* this corresponds to having cone-shaped buckets, where little rain gives big level change, while much rain gives less level change. To finish this small illustrative example, think of dark current and noise as random contribution to the bucket array as the birds flying over and dropping who knows what in the buckets.

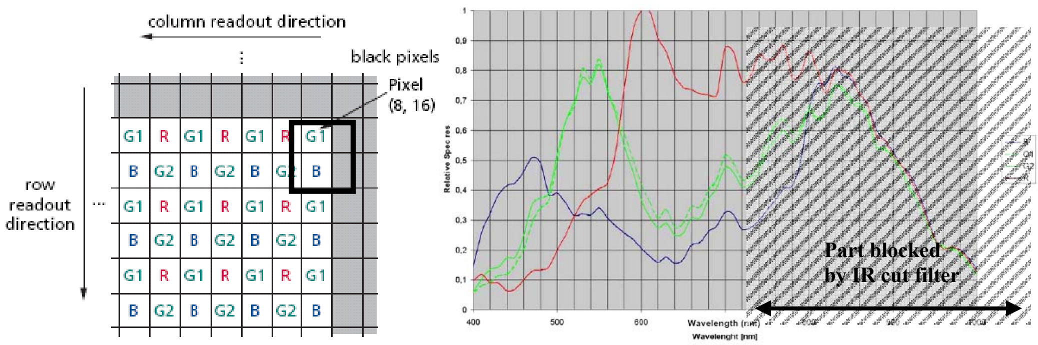

Figure 11 above: Red, Green, Blue Bayesian colour filter. And typical spectral response, compare with Fig. 5 (Source: Micron)

Colour imagers are detector arrays with colour filters superimposed on the chip. In this way only certain bandwidths of light are able to activate certain pixels. One colour pixel is later created from several monochromatic pixels with colour filters by spatial interpolation. As understood, photon economy is poor since light is lost in the filters, and furthermore spatial resolution is lost since it requires several pixels to create one colour pixel. The superimposed colour filters emulate the human eye in respect to wavelength; as such there is not much use of colour imagers in science and spectroscopy. However, the acquired information should be treated as 3D data with two spatial and one spectral dimension, thus we consider the absorbed light intensity $I(x,y,\lambda)$, this makes it easy to generalize for images with higher spectral resolution than 3. The standard red, green and blue colour filter arrangement is referred to as a Bayesian colour filter.

CCD – Charge Coupled Devices

Charge coupled devices (CCD) take advantage of the photoelectric effect. When photons impact on a pixel, they remove electrons and charge accumulates. From the back side of the pixel the opposite charge is applied during exposure so that the accumulated photon charge remains on the pixel. Readout is done only at the borders of the CCD; therefore the charge from each pixel must be moved to the border of the sensor. This is done by sequentially moving the potentials on the backside of the CCD. The result is a displacement of all photon charges along one axis. The row of charges which left the sensing area is later moved to an amplifier and ADC is performed one by one, by the same method that was used to move charges in on the sensor. If the exposure time is too long, charge at certain pixels will eventually overflow and leak to the neighbouring pixels - the phenomenon of blooming. It not only distorts the linear relation for that particular pixel but even distorts neighbouring values. This is a considerable issue when requiring high dynamic resolution in images with extreme contrast, as for example spectroscopy with large peaks of excitation.

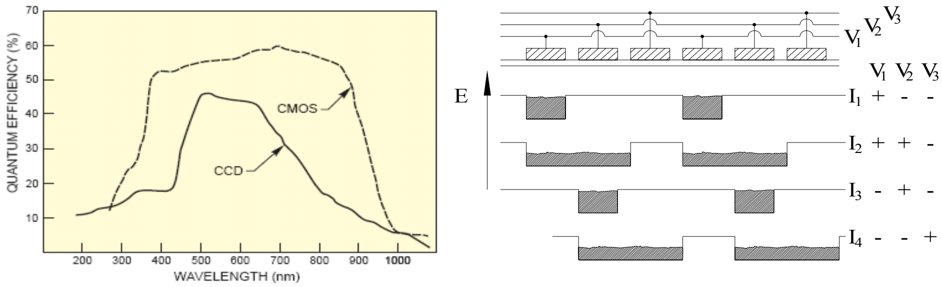

Figure 12. Left: Comparison of quantum efficiencies for different technologies. (Source: Oriel) Right: Schematics of movement of charge in time (Source: www.astro.virginia.edu)

CCDs are comparatively expensive, not at least CCDs for astronomical and scientific use. Sensor chips easily cost 20,000 $ especially for improved sensibility sensors. Also CCDs are analog and require considerable additional circuitry for amplification, filtering, discrimination and interface. Several efforts can be made to increase or extent the spectral response of CCDs. Scientific CCD are often “back thinned” and light impacts on the chip from behind instead. This option allows new chip geometries where light has to travel less inside the silicon before it generates the photon charge. Because of bulk absorption in silicon, especially in the blue region, back-thinned CCDs increases quantum efficiency significally. However, it decreases the red response. Coatings for improving transmission and trapping photons inside the sensitive region can also be applied. Generally, high wavelength response extension is limited by the photon lack of energy to get the electron over the band gap, while blue or UV response extension is limited by the penetration depth for the light in the sensor chip. The last mentioned limitation can be overcome by down-conversion; this technique refers to a phosphor coating transforming UV efficiently into longer wavelengths where better spectral response prevails.

CMOS Imager - Complementary Metal–Oxide–Semiconductor Imager

Complementary Metal–Oxide–Semiconductor Imagers (CMOS imagers) represent a young technology, which goes back to the beginning of the millennium. Markets show rapid development of CMOS imagers, and they are expected to outcompete CCDs in the coming decades. Also they have enabled various new application such as inexpensive webcams and imagers in cell phones. In comparison to the CCD, where the charge is accumulated on the pixel itself, in the CMOS imager the pixel consists of a photodiode. The photons can be absorbed in the depletion layer of the photodiode and generate a reverse current which accumulates on a nearby capacitor until read out. For this reason, every single pixel will even contain a small circuit for amplification, accumulation, reset and read out. Since such circuits also takes space on the chip the photosensitive ratio or *fill factor* is generally lower than for CCDs. Also the pixels are typically L shaped, which can generate certain artefacts in the images under special conditions. (Yaddid-Pecht, 2005)

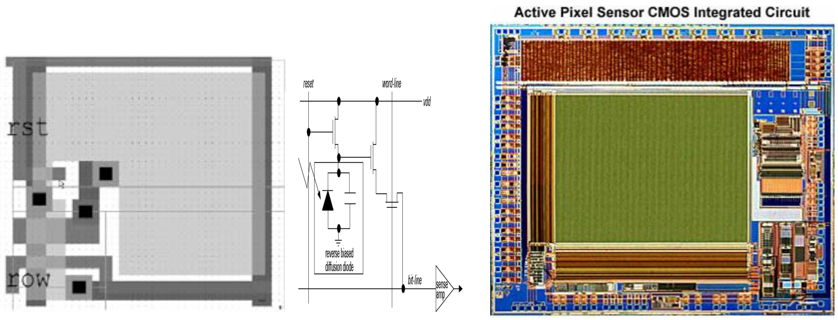

Figure 14: Left: Typical L-shaped CMOS pixel revealing pixel circuitry in lower left corner. (Source: Yadid-Pecht, Senior Member, IEEE) Middle: Example of a CMOS pixel circuit. (Source: www.emeraldinsight.com)Right: CMOS image sensor layout. The additional amplification, signal processing, memory; interface is placed on the same chip. (Source: www.microscopy.fsu.edu)

Inexpensive CMOS imagers typically come with *rolling shutters*. This mean that rows of pixels are not exposed to the image in the same time space. First a reset row will roll down over the sensor and after a readout-row will roll over the sensor. The time between a reset and readout pulse passes a certain row is the exposure time. Thus, with a rolling shutter the image frames are not orthogonal to the time axis. Image artefacts occur when the scenery is moving or the illumination changes. For this purpose most webcams have the option to compensate for indoor/outdoor illumination and 50/60 Hz power frequency. More expensive CMOS imager version have global shutters as CCDs, however. Uniformity is slightly worse for CMOS than for CCDs; this is due to wafer quality. Electronic noise is worse for CMOS, due to the on-chip circuitry.

CMOS advantages over CCD are several. CMOS is a standard chip fabrication procedure, which makes CMOS imagers cheap in comparison to CCDs, especially considering the whole system, since electronics can be incorporated on the chip. Megapixel sensors can be bought for a couple of dollars. Furthermore additional circuitry can be placed on the same chip; this includes analog signal processing, ADC (analog digital converter), frame memory, image processors, compression algorithms, computer interface and communication protocols. This means that basically nothing more than a cable is required to connect a CMOS imager to a computer, and the sensor output will be delivered digitalized. Processing can take place at pixel levels, for instance a logarithmic function can be applied to avoid saturation, and much higher dynamical resolution can be achieved by CMOS imagers by converting pixels to floating-point-like values individually (*photonfocus.com*, 2007). Since CMOS imagers accumulate the photon charge on a isolated capacitor, there will be no blooming and overflowing of charge from one pixel to another. Quantum efficiency is surprisingly better for CMOS even thought the fill factor is smaller; this is due to pixel-based amplification. Power consumption is less for CMOS which makes them attractive to mobile applications. CMOS also provide much faster readout, and frame rates can be in several hundreds of Hz.

Image file formats

There are two main types of graphic files; vector and bitmap. Vector graphics describe shapes mathematically in terms of faces (triangular base elements from which all other shapes can be constructed), whereas bitmap files describe shapes as a pattern of dots or pixels. We will here mainly discuss bitmap files. A bitmap image file is a file where information about an image is stored so that computer programs can display an image on the screen similar to the original image taken by a video camera, scanned or recorded on a digital camera. The result is dependent on the program used, the graphical card in the computer and the type of monitor used. An image file normally consists of a header, where information about the image is given and an array of ASCII-numbers describing the image; modern digital still cameras even store information on zoom, exposure, focus etc. The image can either be given in raw format or in a compressed form. The main operation for compression in general it that a substantial part of the raw data is often irrelevant for the purpose which the data will be used for, and also data might be redundant. The relevancy of pictures is judged in the JPG compression, and is determined by a picture quality factor and the human perception of images. For this reason it often becomes hard to enhance or process JPG compressed images, since the small details have been removed in the compression. Redundancy can be compared to a Swedish phone book. The surname Svensson is rather common and is not written out for every person, but only the first Svensson, in this way we save pages in the book, but we can still derive the complete name without losses, GIF compression operates in that way.

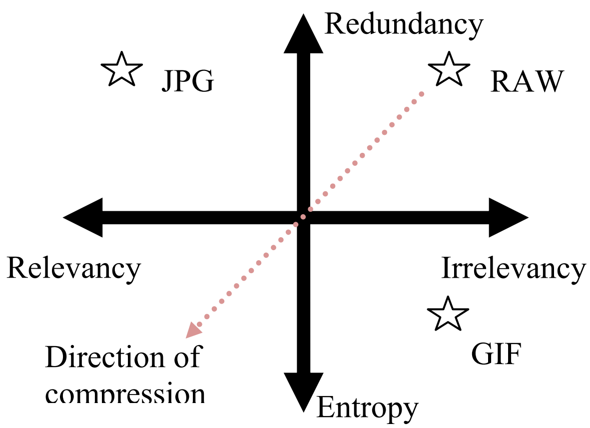

Figure 15: Compression methods of data.

If compression is used it can either be lossless or lossy. While an uncompressed image encodes the value of each pixel in an image, lossless compression looks for areas containing pixels of the same value and encodes the area. As graphic images can represent many more colours than the human eye can perceive or printers can print, it is possible to use a compression scheme that loses some of the information with little or no degradation of the uncompressed image when it is retrieved. With compression the image file size can be reduced down to 1/20th of the original file size. Making image files smaller is important when the file is to be transmitted across networks or for archiving libraries of images.

One important factor affecting the image quality is the dynamic range measured in bit depth. The bit depth of an image makes up the colour values assigned to each pixel in the image. In Table 2 the naming convention is easily understood and based on simple mathematics.

Table 2

Bit depth	**Binary description**	**Colours available**
1-bit	$2^{1^{1}}$	black and white
4-bit	2^4	16 colours
8-bit	2^8	256 colours
16-bit	2^{16}	65 536 colours
24-bit	2^{24}	16 777 216 colours

Commercial colour cameras typically discretize the signal from each colour channel assigning 8 bits. Dynamical range above 8 bit becomes of increased importance when further analysis needs to be done to pictures, which is why many industrial cameras digitalize with 10 or 12 bit.

RAW - Raw bitmap format

A raw image file contains minimally processed data from the image sensor of a digital camera or image scanner. Raw files are so named because they are not yet processed and ready to be used with a bitmap graphics editor or printed. Basically raw files contain equally many matrixes as colour channels containing light intensities at each pixel.

GIF – Graphics Interchange Format

GIF is an extensible file format for the lossless, portable, well-compressed storage of raster images. This format can replace many common uses of TIFF (a rather complicated file format). Indexed-colour, grayscale, and truecolour images are supported, plus an optional alpha channel for transparency (allowing underlying bitmaps to be mixed gradually with the current bitmap). Sample depths range from 1 to 16 bits per component (up to 48bit images for RGB, or 64bit for RGBA).

GIF compression operates along the redundancy axis, and to not remove any information from the image, it simply reduces the names for the most abundant colours. Imagine a picture of a big pink elephant. Starting from upper left corner we could start describing the pixels: pink, pink, pink, pink, etc. But if we instead just refer to “pink” as “p”, the same information is now given by: p, p, p, p, etc. thereby greatly reducing the representation. GIF is powerful in compressing drawings with many spatial details and few spectral details.

JPEG - Joint Photographic Experts Group

JPEG has the smallest file size of all the graphic formats and therefore is a space saver as well as downloads faster. JPEG's also support huge compression rates and often compresses to only 1/10 to 1/20 size. JPEG utilizes "lossy" compression which changes the original image by removing the spatial details, which it judges to be irrelevant for the observer, during the compression process. To ensure minimal loss of image quality only save an image as a JPEG when you are completely finished editing it. The amount of degradation that occurs in an image when converting it to JPEG increases significantly with the amount of compression that you use. Most programs provide a choice between amount of compression and therefore of image quality. The JPEG format supports 24-bit full colour images and is great to use for photographs for this reason. JPEG should be used for photographs or naturalistic artwork and not for illustrations, cartoons, lettering or images with sharp edges as they tend to blur in JPEG. JPEG does not support transparency but progressive JPEGs act like interlaced images.

Images in Matlab

Matlab with *Image Processing* TOOLBOX supports four basic types of images: indexed images, intensity images, binary images and RGB images. Images are loaded either by the command *I=imread('picture.bmp');* or simply by dragging the image into the workspace window of Matlab. Images can be reduced in size with *imresize*, this increases the processing speed. For any given Matlab commando help is displayed by e.g.: *help imread*

The basic data type in Matlab is the rectangular matrix, an ordered set of real elements. The intensity in each pixel is generally represented by unsigned 8 bit integer (*uint8*), meaning that it can be assigned any integer value between 0 and 255. This means that the image will take up one byte for every pixel for every colour channel in the memory. However it generates several problems to operate with unsigned integers, examples of those are: 200+200=255, 1 / 2 = 1 and 3-7 = 0. For this reason it is convenient to convert to double representation temporally. Use *I=double(I)/255;*

A single pixel is selected from an image matrix by using normal matrix subscripting. For example *I(2,15,2)* returns the value of the pixel at row 2 and column 15 in the green colour channel of the image I. The spectrum of the same single pixel is accessed by *I(2,15,:)*.

Images can be presented at the screen using either: *image, imshow, imagesc*.

Geometrical information from images can be acquired by *[x y]=getpts;* Taking into account that the imaging system might have different resolution along each axis.

Intensity of a image can be given by taking the mean value along third dimension (spectral dimension): *Intensity=mean(I,3);*. Normalization of mono coloured images can be performed by *Inorm=(I-min(min(I)))/(max(max(I))-min(min(I)))*; or a more relaxed version: *Inorm=(I-mean2(I))/(k*std2(I))+0.5;* Saturation of a picture between Imin and Imax is performed by *I=min(max(I,Imin),Imax);*

Spatial and spectral resolution of a picture is accessed by *res=size(I)* where two first element are spatial and third on is spectral.

A sub-image containing certain object in the picture can be produced using *Isub=imcrop(I);*

Especially interesting in the field of multi-spectral imaging becomes the command *reshape*, with which we can temporally discard and restore spatial dimensions. This is useful since spectral processing is based on the spectra of the pixel and not on their position x,y. Therefore it is useful to consider a pixel at position x,y rather as just an n^{th} observation. As one can understand we will now apply exactly the same methods to a pixel as we would have done by any other spectroscopic point measurement. The main difference is that, by imaging we get a large number of point measurements in a very short time. By remembering the original spatial resolution, we can reconstruct the image once done with the spectral processing. We can discard spatial information like this *Obs=reshape(I,[res(1)*res(2) res(3)]);* Assuming that above resolution command was executed. We reconstruct the picture by: *I=reshape(Obs,[res(1) res(2) res(3)]);* Every colour channel image is now a column vector, as such, we have observations along the rows and spectral information along the columns.

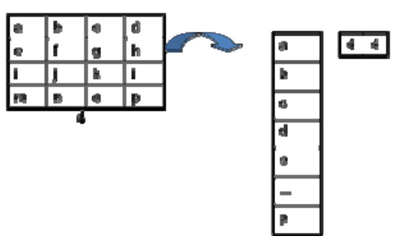

Figure 16: When processing images along the spectral dimension it can be fruitful to temporally discard spatial information.

Obs is a matrix with spectral observation along the rows, and colour channels along the columns. The RGB pixels can be observed in their colour space with *plot3(Obs(:,1), Obs(:,2), Obs(:,3),'.').*

1D histograms are projection of observation clusters in the colour space on a single axis. E.g. a histogram of intensities in the green colour channel can be generated using *hist(Obs(:,2))*. In greyscale images we get a complete overview of light intensities by a 1D histogram. Such histograms can be plotted with *imhist*. For more than one colour channels, we cannot describe the intensities fully by several 1D histograms, instead we will need a ND histogram. This can easily understood, considering that the 1D histograms are projection of a ND cloud. In a CT scanner we can not reconstruct the interior of the body using only three projections. Assume for instance a volume of 10x10x10=1000 voxels. Three projections along x,y,z axis will yield three pictures of 10x10 = 300 pixels. We can not solve 1000 unknown data point only using 300 variables.

Transformation to spherical colour space: *[alpha beta r]=cart2sph(Obs(:,1), Obs(:,2), Obs(:,3))* and visa versa with *sph2cart*. For higher dimensionality in the colour space we can define spherical coordinates: *r = sqrt(sum(Obs.^2,2)); angles=atan(Obs(:,1:end-1)./ Obs(:,2:end))*

Contrast function can be found by expanding the spectral information in on, e.g., polynomial base functions and finding the coefficients. Polynomial expansion of spectral coefficient can be done by *.^n* and training of contrast function can be performed by the backslash operator \. Providing the right result for the function in vector Y, by *contrastfunction=expandedObs\Y.*

If unit-less description of observations are used, the model will be insensitive to variations in intensity. *contrastfunction=expandedAngles\Y.*

A multi-dimensional sub space of base functions can be thought of as ND-cube with the different observations expanded in e.g. powers along each axis of the cube. A position in the cube will indicate certain inter combination of expanded observations. This is similar to the relative simple 2D cosine transform.

$$
\begin{array}{c c}
pix_{1\ldots N} & \begin{array}{cccc} \beta^0 & \beta^1 & \ldots & \beta^m \end{array} \\
\begin{array}{c} \alpha^0 \\ \alpha^1 \\ \ldots \\ \alpha^m \end{array} &
\begin{bmatrix} 1 & \beta & \ldots & \beta^m \\ \alpha & \beta\alpha & & \\ \ldots & & \ldots & \\ \alpha^m & & & \beta^m\alpha^m \end{bmatrix}
\end{array}
$$

$$
\xrightarrow{\mathrm{Re}shape}
$$

$$
\begin{array}{c c}
 & \begin{array}{cccccc} \beta^0\alpha^0 & \beta^1\alpha^0 & \alpha^1\beta^0 & \beta^1\alpha^1 & \ldots & \beta^m\alpha^m \end{array} \\
\begin{array}{c} pix_1 \\ pix_2 \\ \ldots \\ pix_N \end{array} &
\begin{bmatrix} 1 & 0 & 0.8 & 0 & \ldots & 0 \\ 1 & 0.9 & 0.1 & 0.09 & \ldots & 0.0081 \\ \ldots & \ldots & \ldots & \ldots & \ldots & \ldots \\ 1 & 0.1 & 0.1 & 0.01 & \ldots & 0.00001 \end{bmatrix} = \Phi
\end{array}
$$

Figure 17: One RGB pixel can be described by one intensity component and two angular components here α and β. The shape of a cloud of observation points expressed in α and β, can be approximated by a ND polynomial.

Spatial domain filtering can be performed by *medfilt2* and *filter2*. While spatial frequency domain methods can be done by transforming the image with *fft2* and *ifft2*. One might further find use of the command *fftshift*.

Example displaying successive Matlab commandos applied in spectral modeling:

```
I = imread('mytomatopicture.bmp');  % Load picture
I = double(I)/255;                   % Converts to double representation

%%%%%%%%%%%%%%%%%%%%%%%%%%%
% Construction of function:
%%%%%%%%%%%%%%%%%%%%%%%%%%%
tomato = imcrop(I);             % Promt used to select tomato pixels
notTomato = imcrop(I);          % Promt used to select not tomato pixels

resTomato = size(tomato);       % Determines resolution of selection
resNotTomato = size(notTomato);

% Discard spatial information

tomato = reshape(tomato,[resTomato(1)* resTomato(2) resTomato(3)]);
notTomato    =    reshape(notTomato,[resNotTomato(1)*    resNotTomato(2)
resNotTomato(3)]);

obs = [tomato ; notTomato]; % Places all observation in the same matrix
          % pixel number as rows and spectral information along columns

% Construct matching answer, ones for tomato and zeros else
corectAnswer = [ones(length(tomato),1); zeros(length(notTomato),1)];

% Casts to spherical coordinates
[a b r] = cart2sph(obs(:,1), obs(:,2), obs(:,3));

% Expand angular components of observations
expandedObs  =  [a.^0.*b.^0  a.^0.*b.^1  a.^0.*b.^2  a.^1.*b.^0  a.^1.*b.^1
a.^1.*b.^2 a.^2.*b.^0 a.^2.*b.^1 a.^2.*b.^2];

% Solve contrast function coefficients
contrastFunction=expandedObs\corectAnswer;

%%%%%%%%%%%%%%%%%%%%%%%%%
% Evaluation of function:
%%%%%%%%%%%%%%%%%%%%%%%%%
res = size(I);                  % Determines resolution of entire image
I = reshape(I,[res(1)* res(2) res(3)]); % Discard spatial information
[a b r] = cart2sph(I(:,1), I(:,2), I(:,3)); % Transforms to spherical

% Expands angular components
expandedObs  =  [a.^0.*b.^0  a.^0.*b.^1  a.^0.*b.^2  a.^1.*b.^0  a.^1.*b.^1
a.^1.*b.^2 a.^2.*b.^0 a.^2.*b.^1 a.^2.*b.^2];

% Processes image, using                contrast function
modelAnswer = expandedObs*contrastFunction;

% Restores spatial informatino of the now monochromatic picture
modelAnswer = reshape(modelAnswer,[res(1) res(2) 1]);

% Saturates result underzero and over one
modelAnswer = min(max(modelAnswer,0),1);
imagesc(modelAnswer)                           % Displays result
```

Image Enhancement

The principal objective of enhancement techniques is to process an image so that the result is more suitable than the original image for a specific application. This means that image enhancement is problem-oriented. Thus, for example, a method that is quite useful for enhancing x-ray images may not be the best approach for enhancing pictures of Mars transmitted by a space probe.

We will discuss three different categories of image enhancement: intensity domain methods, spatial domain methods and frequency domain methods. The spatial domain methods can be either based on point processing, which modifies the grey level of a pixel independently of the nature of its neighbours, or mask processing. Masks, defined as small sub-images, are used in local processing to modify each pixel in the image to be enhanced. The frequency domain methods are based on the Fourier transform.

Intensity Domain Methods

Some processing methods are based only on the intensity of single pixels. We will denote the intensity of pixels before and after processing by r and s, respectively. The enhancement is described by the transformation function $s = T(r)$.

Contrast stretching

Low-contrast images can result from poor illumination, lack of dynamic range in the imaging sensor, or even wrong setting of a lens aperture during image acquisition. The idea behind contrast stretching is to increase the dynamic range of the grey levels in the image being processed. Figure 18 shows a typical transformation used for contrast stretching where L is the number of graylevels. The locations of points (r_1, s_1) and (r_2, s_2) control the shape of the transformation function. For instance, if $r_1 = s_1$ and $r_2 = s_2$, the transformation is a linear function that produces no changes in grey level. If $r_1 = r_2$, $s_1 = 0$ and $s_2 = L - 1$, the transformation becomes a *tresholding function* that creates a binary image. Intermediate values of (r_1, s_1) and (r_2, s_2) produce various degrees of spread in the grey levels of the output image, thus affecting its contrast.

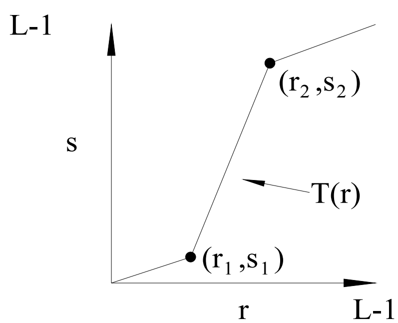

Figure 18. A transformation function suitable for contrast stretching

Compression of dynamic range
Sometimes the dynamic range of the processed image far exceeds the capability of the display device. In this case only the brightest parts of the image are visible on the display screen. This problem is encountered when the Fourier spectrum of an image is to be displayed. A suitable transformation that will compress the dynamic range is the following intensity transformation:

$$s = c\log(1 + r)$$

where c is a scaling constant and the logarithm function performs the desired compression. If the Fourier spectrum has a dynamic range of 0 - $2.5 \cdot 10^6$, then $\log(1+|r|)$ ranges from 0 to 6.4. With a scaling factor c of 255/6.4, a good result is obtained on an 8-bit system.

Grey-level slicing
Highlighting a specific range of grey levels in an image often is desired. One approach is to display a high value for all grey levels in the range of interest and low values for all other grey level. This transformation, shown in figure 19a, produces a binary image. Another approach, based on the transformation shown in figure 19b, brightens the desired range of grey levels but preserves the background and grey-level tonalities in the image.

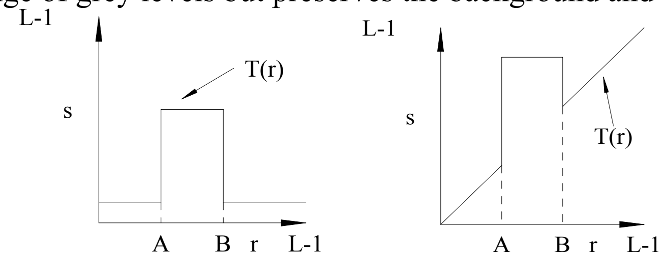

Figure 19. Intensity-level slicing: left a transformation function that highlights a range [A,B] of intensities while diminishing all other to a constant, low level, right a transformation that highlights a range [A,B] of intensities but preserves all others

Histogram equalisation
The histogram of a digital image with grey levels in the range [0, L-1] is a discrete function $p(r_k) = n_k/n$, where r_k is the kth grey level, n_k is the number of pixels in the image with that grey level, n is the total number of pixels in the image, and $k = 0, 1, 2, ...L - 1$. For a dark image this histogram has its high values for low grey levels while for a bright image high values of the histogram are concentrated towards the bright end of the grey scale range. If the histogram is concentrated in a small range of grey levels the image is of low contrast and similarly a histogram with significant spread corresponds to an image with high contrast.

For an image with low contrast it is possible to increase the dynamic range by contrast stretching as discussed above, but it may be better approach to use histogram equalisation. This means that the grey level values in a way are equally distributed within its total range. To illustrate histogram equalisation we consider figure 20a which shows a 512 x 512, 8-bit image (of a welding) that is dark and has very poor dynamic range.

Besides the image is shown the histogram which show a concentration to low grey level values.

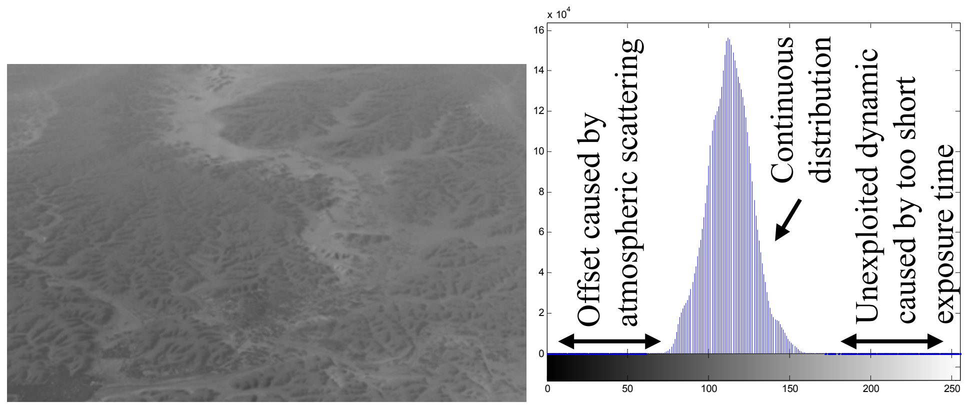

(a) (b)

Figure 20 (a) Original image of Saharan city from an airplane (b) its histogram

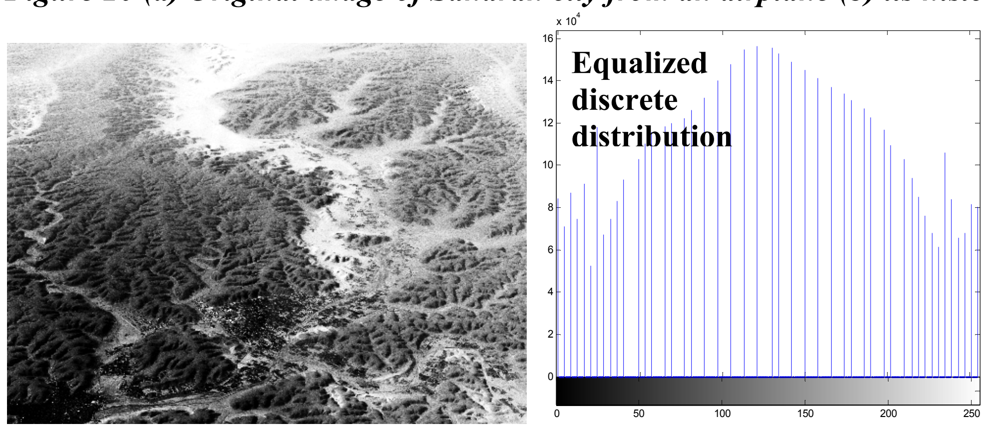

Figure 21 (a) Image subjected to histogram equalisation and (b) its histogram.

With histogram equalisation using a transformation given by the relation

$$s_k = T(r_k) = \sum_{j=0}^{k} \frac{n_j}{n} = \sum_{j=0}^{k} p_r\left(r_j\right)$$

where s_k are the new grey level values, a uniform histogram is obtained. The histogram is now distributed within the total grey level range. In figure 21a is shown the image enhancement resulting from this transformation and in figure 21b its histogram. Note that in the figures the horizontal range is [0 - 255] and the vertical is [0 - 30K]. To obtain a histogram as defined above the number of pixels (n_k) should be divided by $(512)^2$ (= n) and similarly the grey level values should be divided by 255.

In above histogram we have counted the number of measurements / pixels with a single parameter, namely the grey level or light intensity, which fall into a certain interval bin. Naturally if we observe more parameters, it is free for us to arrange the bins in a 2D or ND matrix in order to analyse the simultaneous values of different parameters, e.g. the simultaneous counts of pixels who have a red value between 220 and 230, while the green value is between 110 and 120.

Spatial domain methods

Spatial filtering

The transfer function and the impulse or point spread function of a linear system are inverse Fourier transformations of each other. A low pass filter is a filter that removes the high frequency components of the Fourier domain. The low frequency components are unaffected. The high frequency components are responsible for the edges and other sharp features in the image. The effect of a lowpass filter is thus image blurring. Similarly a high pass filter reduces the low frequency components and leaves the high frequency components untouched. The net effect of such a filter is a sharpening of edges and other sharp details. A third type of spatial filter is the bandpass filter which acts on the frequencies between the low and high parts of the Fourier domain.

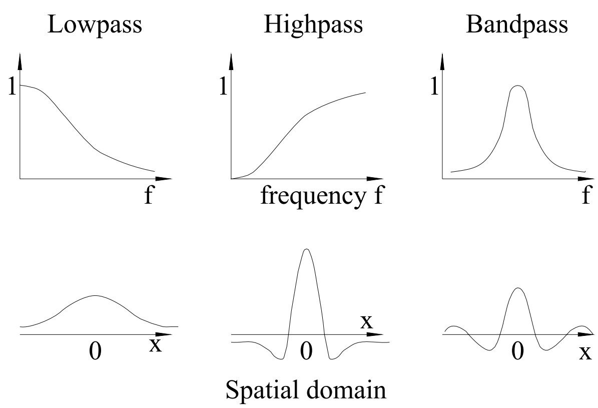

Figure 22. Top: cross sections of basic shapes for circularly symmetric frequency domain filters. Bottom: cross sections of corresponding spatial domain filters

Spatial filtering can be performed either in the spatial domain or in the Fourier domain. In the Fourier domain it is quite clear how such a filter should be designed (see below). To design a filter in the spatial domain we will first study figure 22. This figure shows cross sections of circularly symmetric lowpass, highpass, and bandpass filters in the frequency domain and their corresponding filters in the spatial domain. Using these figures it is possible to design a suitable spatial filter. A mask is moved over the image (this results in a convolution) and the grey level of the image point in the centre of the mask is replaced by a new level found from the coefficients of the mask. As an example the mask in figure 18 will give the grey level R given by the following product:

w_1	w_2	w_3
w_4	w_5	w_6
w_7	w_8	w_9

$$R = w_1 z_1 + w_2 z_2 + \ldots\ldots + w_9 z_9$$

Figure 23. A 3 x 3 mask with arbitrary coefficients (weights)

A simple low pass filter would be a filter with all coefficients having the value 1. To compensate for an increase in the total grey level of the image a normalisation factor of 1/9 is used. In figure 24 is shown a 3 x 3 and a 5 x 5 low pass filter. Such a filter is useful for smoothing the image and reducing high frequency noise.

$\frac{1}{9}$ x

1	1	1
1	1	1
1	1	1

$\frac{1}{25}$ x

1	1	1	1	1
1	1	1	1	1
1	1	1	1	1
1	1	1	1	1
1	1	1	1	1

$\frac{1}{9}$ x

-1	-1	-1
-1	8	-1
-1	-1	-1

Figure 24a. Spatial low pass filters of various sizes. Figure 24b. A 3 x 3 high pass filter

To obtain a highpass filter we see from figure 17 that the coefficients should be positive near its centre and negative in the outer periphery. In figure 24b is shown the basic 3 x 3 highpass filter. More details of the properties of these and other spatial filters are given in ref. 1.

Frequency Domain Methods

Filtering in the frequency domain is in principle straightforward. First the Fourier transform of the image is calculated then the transform is multiplied with the appropriate filter transfer function. When the inverse transform of the result is taken the filtered image is obtained. For an ideal filter with a sharp cut-off frequency there is the problem of ringing. This problem is possible to avoid if a filter is used that have no sharp discontinuity.

Thermography

Blackbody radiation

There is no fundamental difference between radiation in the different bands of the electromagnetic spectrum, that is, they are all governed by the same laws and the only difference are those due to differences in wavelength.

A blackbody is defined as an object which absorbs all radiation that impinges on it at any wavelength. Kirchhoff's law states that a body capable of absorbing all radiation at any wavelength is equally capable of emitting radiation. A simple blackbody consists of a box with an aperture in one of the sides. Any radiation which enters the hole is scattered inside the box and absorbed. By providing such an isothermal cavity with a heater it will generate blackbody radiation and the cavity can be used as a thermal reference in an IR-camera. Max Planck was able to describe the spectral distribution of the radiation from a blackbody in 1900:

$$I_{e\lambda} = \frac{2\pi h c^2}{\lambda^5}\left[\frac{1}{e^{(hc/\lambda kT)} - 1}\right].$$

Here k is Boltzmann's constant and h Planck's constant. Planck's law first derived experimentally then verified with the help of quantized energy gives us a family of curves when plotted, figure 25.

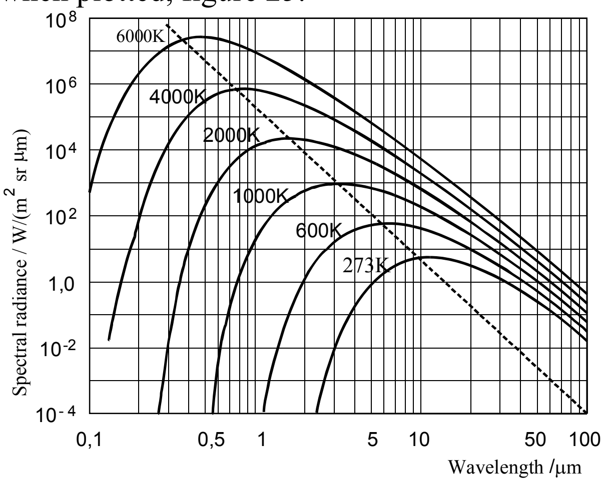

Figure 25. Planck's law plotted at different temperatures.

By derivating Planck's law with respect to λ, and finding the maximum, we have:

$$\lambda_{\max} = \frac{2898 \cdot 10^{-6}}{T} \quad ,$$

which is Wien's displacement law that expresses the common observation that colours vary from red to yellow as the temperature increases. The sun (approx. 6000 K) emits yellow light, peaking at about 0.5 µm in the middle of the visible light spectrum. At room temperature or 300 K, the peak of radiant emittance is at 9.7 µm, in the far infrared.

By integrating Planck's law we obtain the total radiant emittance of a blackbody, the Stefan-Boltzmann law,

$$I_b = \sigma T^4 ,$$

where σ is the Stefan-Boltzmann constant which is equal to 5.7×10^{-8} $Wm^{-2}K^{-4}$. Using the law to calculate the power radiated by the human body, at a temperature of 300 K and an external surface area of $2m^2$, we obtain 1 kW. Fortunately we receive almost equally much from the radiation emitted from our surroundings at room temperature. Further the emitted power have to by multiplied the emissivity, ε, which vary between 0 for a perfect mirror, and 1 for a perfect black body. Most object have an value in between. In the grey body assumption ε is assumed to be independent of λ, also ε is most assumed to be independent of T. Both assumptions are not always entirely valid though. It is noteworthy that a mirror is equally bad at absorption radiation as emitting it self. I day life application of this are thermo flasks, where vacuum prevent conduction heat and a mirror coating prevents radiated heat transfers.

IR detectors

The detection of infrared radiation can be performed in two ways. Infrared radiation is heat transfer or we can picture infrared radiation as photons. A microbolometer detectors consists of a matrix of small bolometers sensitive to heat and each of the bolometers react with a change in resistance. A temperature change causes thus a change in the measured resistance of the individual "pixels" of the matrix. A microbolometer based infrared camera will be demonstrated during the laboratory work.

The other way of detecting infrared light is the use of a semiconductor that captures the incoming photon in a depletion region with a sufficiently small bandgap.

Semiconductor quantum detectors are the most commonly used detectors today. The images that you will process during the laboratory work was captured with a Mercury Cadmium Telluride (MCT) “Signal-Processing-In-The-Element” (SPRITE) detector array cooled with a miniature Stirling cooler to 80 K. In order to prevent the camera from seeing its own heat, the camera cavity, where the imager is situated, is usually placed in a so called cold shield.

IR imaging

As discussed above it is characteristic for the 10μm IR region that objects at room temperature having their peak emission for these wavelengths. In general when changing the wavelength all phenomena changes in significance. Well known examples are the medical X-ray region where tissue becomes transparent, and the microwave region for mobile communication where even concrete walls become transparent. Similarly, the air which is transparent in the visible region, becomes opaque below 300 nm and also in the IR between 5 μm and 8 μm. Thus, when imaging in IR, assumptions can generally not be done from the visible region. Properties of transmission, reflection, absorption and scattering might differ greatly.

In many applications, for instance remote measuring of T, conduction, friction and evaporation heat etc. One interesting application is finding and visualising gas plumes in real-time. Images of gas concentration distributions are needed in several environmental studies. Global monitoring of greenhouse gases from spacecraft-based systems is one example. An industrial application is hydrocarbon emission monitoring from petrochemical installations and leakage detection along pipelines. Indoor monitoring of the working environment and surveillance of hazardous gases are other fields where imaging is useful. Nitrous oxide leaking from anaesthetic masks in operating theatres has been studied with an IR camera using absorption of the gas throughout the spectral region of the camera. The gas can be confined by turbulent flow around the breathing zone of the anaesthetist and ventilation aspects are studied with this method. An active method using a carbon dioxide laser that scans the detector field of view is one way of detecting gas leaks in industrial environments. Another method uses infrared emitters and a retroreflector-screen to visualise gas absorption.

Gas correlation spectrometry is a particularly simple and powerful technique utilising an absorption spectrum in the measuring system that is compared with the spectrum of incoming light and a signal proportional to the number of absorbing molecules of a specific kind is generated. Gas correlation spectrometry has previously been employed in passive point monitoring and in laser radar.

In this laboratory work we will use a new method that makes it possible to visualise gases and estimate concentration in a two-dimensional image based on infrared absorption and gas-correlation techniques. We will first recapitulate the point monitoring principle, see figure 26.

The incoming light is sent through a cell containing the gas to be studied at such high concentration that little or no light can pass at the absorption wavelength, and through a variable filter. With no external gas present the variable filter is adjusted so that equal intensities are obtained in the spectral region cut out by a band-pass filter. Now, if an external pollution plume is present the signal in the gas cell arm is not affected (except a minor broadening of the absorption lines) since no more than full absorption can be obtained. The incoming light in the other arm is reduced due to absorption lines in the plume. The introduced imbalance in the two arms is a measure of the gas concentration in the plume.

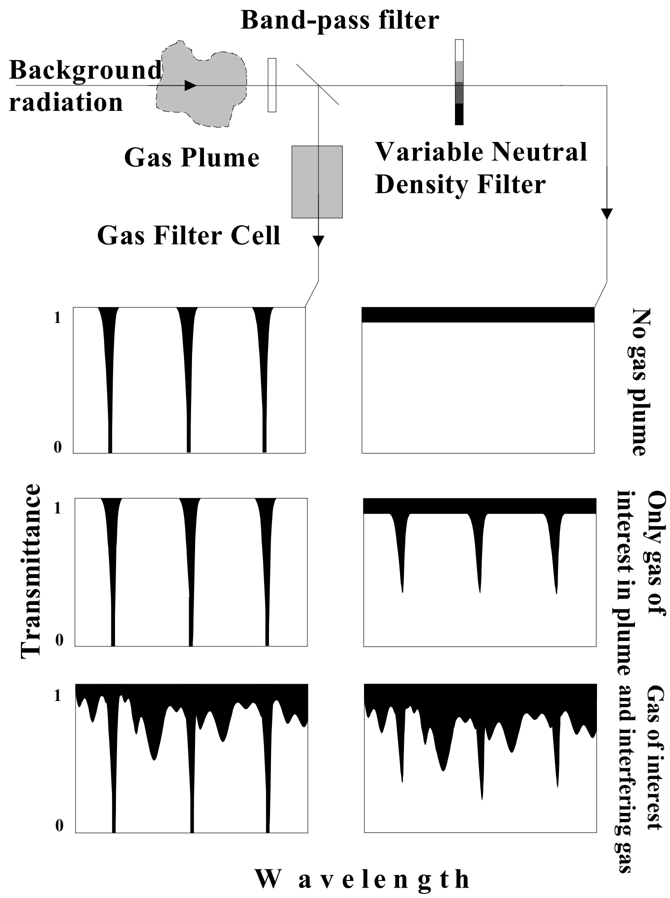

Figure 26 . The gas-correlation principle.

IR spectra

Several interesting and environmentally important gases have vibration-rotation absorption in the infrared region, such as hydrocarbons, sulphur compounds and nitrous compounds. Figure 27 and 28 shows examples of typical spectral features in the fundamental infrared region (2-20 µm).

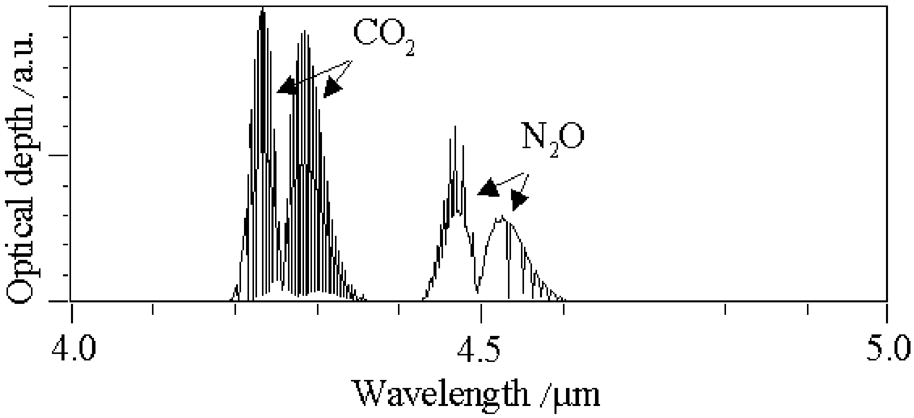

Figure 27. Spectral features of CO_2 and N_2O.

As seen in 27 it is possible to use a band pass filter to specifically select one of the gases. But in the case shown in figure 28 we need a more sophisticated method namely the gas-correlation method to select the gas of interest.

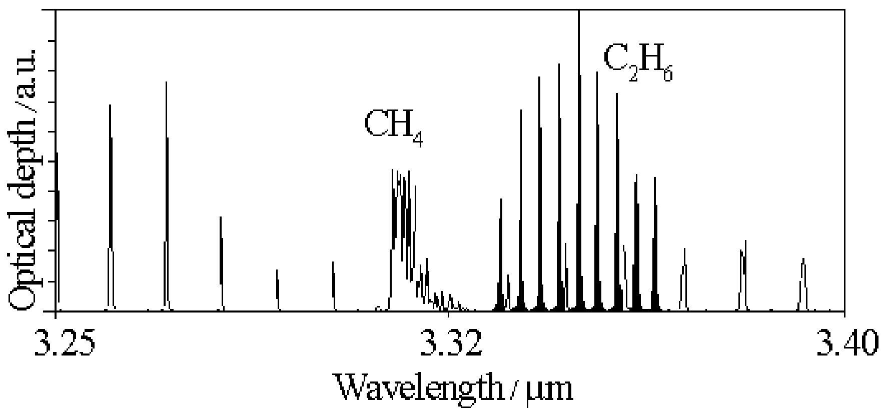

Figure 28. Interfering gases, CH_4 and C_2H_6

Gas correlation

In preparation for the laboratory work two images of the flowing gases was formed at the same time on an IR sensitive camera with a compact Cassegrain split-mirror telescope. A bandpass filter was used to isolate a small spectral region containing absorption features of the gas that was studied. In front of one part of the telescope a short CaF_2 cell filled with a high concentration of the gas was mounted. The other part of the telescope gave a direct image, see figure 24. In this laboratory work the images will be correlated in a computer to eliminate differences in background illumination as well as the interference by other gases. The correlation or image processing consists of mainly two steps, where the first step is to choose overlapping regions of interest in the gas-filtered image and the direct image. The direct image is then divided by, or subtracted from, the gas-filtered image. In order to divide the images without loosing intensity dynamics we multiply the gas-filtered image with a factor. If the images are noisy (salt and pepper noise or shot noise) the best correlation technique is subtraction. The noise in the result image will be increased when dividing noisy images, due to the problem of dividing a large number by a small number, but added when the images are subtracted from each other. The result will be a grey-scale or false colour-coded picture only showing the distribution of the specific gas. Flow and concentration in two dimensions can in this way be presented in near real-time. All gas-correlated images will be presented as an infrared motion picture.

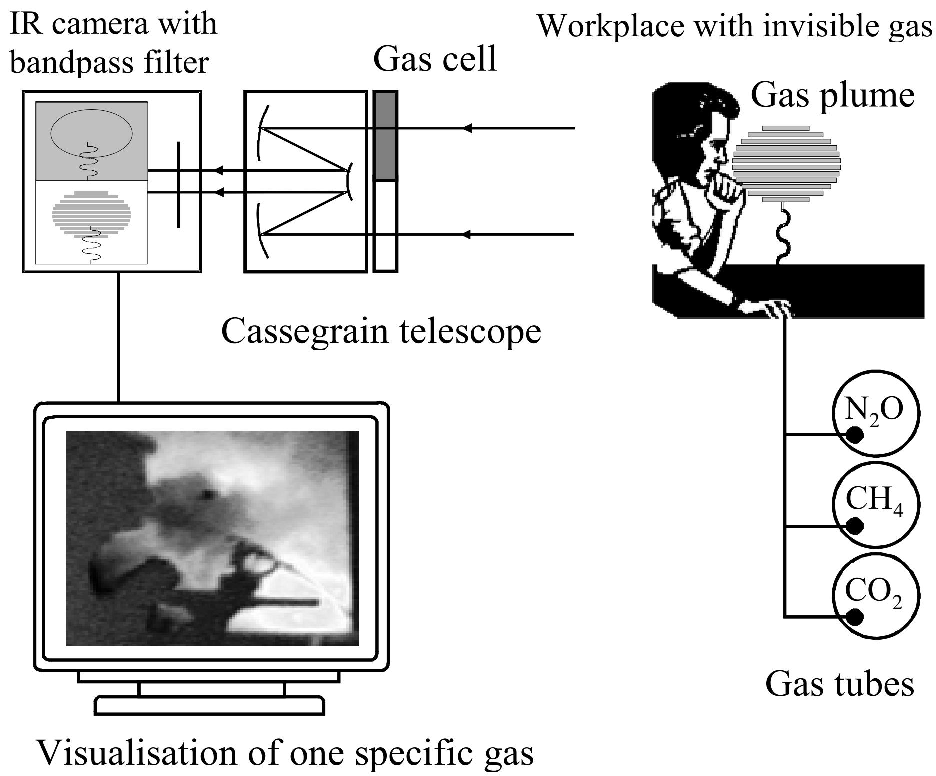

Figure 29. Infrared set-up for gas-correlation spectrometry.

References:

1. R. C. Gonzales, R. E. Wood, Digital Image Processing, Addison Wesley 1993
2. C. Thompson, L. Shure, Image Processing Tooolbox, Math Works Inc. 1995
3. G. Saxby, Practical Holography, Prentice Hall 1994
4. M. Brydegaard, Master thesis: New approaches for applied spectroscopy. LRAP-382. Atomic physics division, Lund 2007.
5. O. Yaddid-Pecht. CMOS Imagers: From Phototransduction to Image Processing. Springer-Verlag, Heidelberg 2005.
6. M. Meingast. Geometric models of rolling-shutter cameras. University of Berkeley, California 2005.
7. Bowmaker J.K. and Dartnall H.J.A., "Visual pigments of rods and cones in a human retina." J. Physiol. 298: pp501-511

Before starting: *Transfer the files in the common folder for the lab to your own folder.*

Pixel size, magnification and resolution

The CCD camera is equipped with a lens with variable aperture and automatic gain control (AGC: see the back of the camera) and an electronic level control (ELC). It is suitable to use an aperture of about 5.6 - 8, as this is where the camera objective gives a rather sharp image.

With a CCD camera an image is recorded of an object with a given size (paper with scale). We use a frame grabber Grabbex that we can handle from within the MatLab toolbox Image Acquisition Toolbox. With a suitable m-file *takeImage.m* we can take a snapshot of the video image from the black and white camera. The frame grabber digitizes the analogue video signal from the camera. *Start the m-file and study the effect of turning on and off the AGC.* Use the commands *imcrop, size* or *getpts*. How many pixels on the screen correspond to 1 cm in the object plane both vertically and horizontally? How would you propose to determine a distance in the image? (Hint: use the commands: *getpts, diff, sum sqrt*). The resolution of the system is determined by studying images of a test object with different distances between lines. Compare the vertical and the horizontal resolution. Place the camera at different distances from the object.

Grey scale images, image format and image information

Take a new image of something in the lab with the program takeImage.m and select a part of the recorded image (small size matrix = faster calculation), use *imresize* or *imcrop* for this purpose. The intensity matrix should have values between 0 and 1. Note! The values in the intensity matrix should be of type *double*. The image can now be presented as perspective plots with functions such as *surfl* or *mesh*. Try these. Try the function *flipud, fliplr* and *'* on the matrix and see what happens. Print the image. What kind of mathematical operation is performed with the subtraction of two identical images with one image displaced one pixel? Displacing can be done by indexing e.g. *I1=I(2:end,:);* and *I2=I(1:end-1,:);* or alternative by translation of the camera.

Use the test chart with horizontal and vertical lines again and determine the MTF-values for the whole image system. Use a distance from the object of less than 1 m. Take an image with *takeImage.m*. Select a subimage with *imcrop* and take the *sum* along one spatial dimension. Simple column sums or row sums can now be generated and these sums can be presented as diagrams. Quantify the resolution of the camera. The magnification (in this case <1) is found from measurements and calculations using simple image geometry. Use lines/mm in the image plane of the objective as units on the x-axis. Print the diagram. Discuss the result together and with the supervisor.

Image enhancement

Study the two images "IMG1.bmp" and "IMG2.bmp". Try to enhance the images with histogram equalization and contrast stretching, use *imhist*, *histeq* and the mathematical operations * / + -.

Colour images

Analysis*:* Study the screen at various regions with a magnifier. One colour image is recorded with the digital camera (Olympus) with 3x8-bit depth. The object consist some coloured objects. You can make use of a pair of polarisers to avoid specular reflexes, confirm this by taking two photos with respectively parallel and perpendicular filters. The image is transferred to the computer using USB cable. Identify the dimensions and the resolution of the data. Mark out several objects in the image and compare their mean spectrum (use: *imcrop, reshape and mean*). Do the spectra make sense? Create histograms of the RGB colour channels; do you make use of the entire dynamic range? How would you propose to stretch the intensities? Divide the RGB picture with its own intensity and look at the result, what information remains? Intensity can be calculated by taking the sum along the spectral dimension using *sum*. (Hint, you will have to divide each colour channel individually or alternatively use *repmat*). Crop a region and discard the spatial information with *reshape*. Plot observations from one or several objects in a RGB space using *plot3(R,G,B,'.')*. Transform to spherical colour space using *cart2sph*. Present histograms again, can you identify the objects? Plot the angular components using *plot(alpha, beta,'.')*. Use the program *angHist2* on the two angular components; observe the histogram with *mesh* or *imagesc*. What does a particular object look like in the histogram? How would you make a function that finds that particular object and nothing else? Expand the angular components in powers from 0 to 2 and take all intercombinations. Use *th=[alpha.^0 alpha.^1 alpha.^2 beta.^1 beta.^2 alpha.*beta … alpha.^2*beta.^2];* Construct a right answer vector *Y=[ones(n,1) ; zeros(m,1)];* and solve the contrast function with backslash. Calculate *Yestimates=th*contrastfunction;* Present a histogram for the result and determine a threshold. Process the image with the same contrast function. Reconstruct the image *Yestimated* for the image and present it as an image. Present *image(Yestimated>treshold)*

Synthesis**:** A coloured object is recorded (with 8-bit depth) with red, green and blue colour filters in front of the black and white CCD camera. The three images are labelled different in the Matlab workspace, create a color picture by indexing, e.g. *I(:,:,1)=Red;*. Filtering the detected light works equally well as filtering the illumination light. Move to the microscope setup. Use the Matlab command *lightoff* and *lighton(n)* where n is 1..13. This will change the LED illuminating the sample. You can get the corresponding wavelength with the command *l=getlambda*. Use the Windows program AMcap to locate a suitable object in the sample. Close AMcap. Use the command *M=multiaq;* to take 13 pictures of the same scenario with different wavelength. Compose and present at least three false colour pictures. Transform to ND "spherical" coordinates and present a false colour picture of the angular components. Which disadvantages does the imaging system suffer from?

IR - Imaging.

Familiarize yourself with the infrared camera, and adjust zoom and focus. Study the temperature of various objects in the room. Study heat conduction, friction heat and evaporation heat. Study the transmission, reflection, absorption and scattering properties of water, glass, plastic glass, whiteboard, black and white plastic bag. Why is the IR lens not made from glass? When using a IR camera in a sea rescue mission, what will appear of a person in the sea? Why can an image not be formed behind a black and white plastic bag in the visible, and why would you possibly succeed in the IR? Can you measure the temperature of the whiteboard and a white plastic bag?

Laboratory report:

The report should primarily show what you have learned during the lab exercise. Include figures from the lab to explain the topic. Below are some checkpoints for the report.

- Discuss quantization and dimensions.
- What are the similarities in the different domains?
- Discuss MTF and the Rayleigh and Nyquist sampling criterion in respect to space and time.
- Discuss orthogonally of dimensions in different modes of acquisition.
- What is the advantage of filling the dynamic range during the acquisition rather than stretching the contrast afterwards? Compare with similarities in other domains.
- Why is the stretched histogram discrete rather than continuous?
- Give reasons for the fact that natural intensity variance is large while variance of relative intensities across the spectral domain are small for a given object.
- How can we discard dimensions, process images and restore the dimensions again?
- What are the similarity of a multivariate histogram and a multivariate model trying to separate one cluster of data from the rest?
- When filtering light in illumination instead of detection what considerations and assumptions must be done?
- Discuss chromatic aberration and it relation to multispectral imaging.
- Discuss IR imaging, possible error sources, and what physical phenomena to take into account.