Sample Exam Questions for the Course NPGR025 „Introduction to Colour Science“

held at the Faculty of Mathematics and Physics, Charles University Prague,
by Alexander Wilkie, KSVI MFF UK

Version 1.19, January 5th, 2024

Note on how to use this document: this collection of sample exam questions is merely intended to make your personal preparation for the exam easier by giving a rough outline of some of the more popular questions that can be asked by the examiner. There is *no* guarantee of any kind that only these questions will be asked, or that they will be asked in the exact wording that you find on these pages. The provided answers are intentionally brief in most cases, and are only intended to give you a basic idea of what the correct answer is. For a number of the questions, they do *not* constitute a complete answer!

Green text has been updated/added since version 1.17.

Give a short overview over the historical development of colour science

Linear colour models of antiquity (with no concept of the difference between additive and subtractive colour mixing) – first 3D models in mediaeval times – Leonardo da Vinci (4 primaries) – Lambert (first mention of CMY) – Newton (Prism, Rainbow) – first colour circles (Harris) – Goethe as „adversary“ of Newton, „dark colours“, „edge colours“ – Otto Runge (first colour sphere) – Munsell with the first colour ordering system, systematic gathering of scientific results starts with the formation of the CIE

Name the 15 fundamental causes of colour, and describe the main ones in detail

See the slides or the book for the list of effects. It is not crucial that you know all 15 effects, as the differences between some of them are rather subtle by the standards of non-physicists. However, the main effects (such as blackbody radiation, induced emission, charge transfer, diffraction, dispersion etc.) should be known.

Describe the structure of the human eye

See the drawings on the slides; important points are the nature and the location of fovea and blind spot, and the consequences of shape and size of the fovea (i.e. the fact that it is only 1-2 degrees in your field of view, and that the brain compensates by „filling in the rest“, and scanning eye movements).

What is the function of the iris?

The iris has the same function as the aperture diaphragm in a camera; the size of the pupil opening governs the trade-off between depth of field and the amount of light that reaches the retina.

What is the main drawback of a single lens camera set-up? (which is what the human eye amounts to)

Cameras with a single lens are prone to chromatic aberration errors. This is a lesser concern in the case of the human eye, since we only see colours well in the fovea, i.e. a narrow region around the main optical axis. And chromatic aberrations only occur near the edges of the projected image.

Describe structure and function of the retina

A thin layer of tissue which (starting from the inside of the eye) contains the neural circuitry responsible for various post-processing tasks, and only in the outermost layer the actual light sensitive cells. Behind that is a layer of (in humans) weakly reflecting material; cats have a much more strongly reflective material there, which gives some photons a second chance to hit a light sensitive cell. This is only useful in low-light conditions, though, and comes at the expense of some visual acuity, so not all primate eyes have this feature. The outermost layer are the blood vessels that supply oxygen and nutrients to the entire retina.

What types of light sensitive cells are present in the retina?

1. **Cones**, sensitive to colour, with 3 sub-types that are sensitive to long, medium and short wavelengths (so they are collectively abbreviated as LMS cells). They are comparatively fast to adapt and respond, but only active at daylight levels of light. The three types of cones are not referred to as RGB, as they are effectively yellow-green (L), green (M) and blue (S) sensitive.
2. **Rods (R)**, only sensitive to brightness, slow to react and adapt, but active at low light levels

How are they distributed across the retina?

LMS cells are practically only in the fovea, with the exception of short wavelength cells - these are more or less evenly distributed. Rods are fairly evenly distributed, with a *decrease* in the fovea, to make room for the LMS cells.

Why are our opponent colour pairs red/green and yellow/blue, if we do not have yellow and red sensitive cones?

See the slide which shows the schematic of how the retina is wired: basically, our perception of red comes from the subtraction of the M from the L signal, and the Y-B channel from a subtraction between the luminance channel and the blue channel.

What is the *adaptation* of the human eye?

Adaptation is how the eye to changing illumination levels, i.e. which type of light sensitive cell is predominantly active, and how fast the regeneration of rhodopsin is performed. The word is also sometimes used to refer to the focusing state of the eye, but this is incorrect: the focusing state is known as *accommodation*.

What types of image processing are being carried out at the retinal level?

Edge enhancement, averaging of several individual sensor cells, re-coding of the „native" LMS+R signals that emanate from the sensor cells to a signal that (roughly) resembles L*a*b*.

What signals are sent to the brain from the retina?
How is colour encoded on the optic nerve?

Even though we have four types of sensor cell (LMS+R), only three signals are sent back to the visual cortex in the brain (one red-green channel, one blue-yellow, one luminance). The creation of this compound signal is done by the neurons in the intermediary layer of the retina.

Which colour space intrinsically corresponds to the structure of human colour vision?

CIE L*a*b*, since its structure roughly corresponds to the signal that is sent from the eye to the brain.

What are metameric colours? Is this a common phenomenon?

The human eye is not a device that is capable of analysing the spectral composition of light; it only has three wide-band sensors. For these three sensors, an arbitrary number of spectral inputs can cause the same stimulus, i.e. many different input spectra can result in the same LMS cell response. Conversely, stimuli that match for one observer, may differ for another one.

This is an extremely common phenomenon; practically all our colour reproduction technologies just produce metamers of the original colours. Imagine, just for a moment, how complex colour reproduction would be if we would have to match the *spectral composition* of every colour we want to reproduce! In this regard, the fact that our eyes filter information is a huge bonus.

What types of metameric failure are there?

There are two main types of metameric failure:

1. **Illuminant metameric failure**. This is the “classical” form of metameric failure: colours that match under one particular illuminant no longer match under a different one. This is a regularly occurring issue in colour science reproduction, since modern low-energy luminaires all have spiky emission spectra, while natural light sources usually have smooth spectra.
2. **Observer metameric failure**: this form is less well known than the first one, but is still of considerable practical relevance. It occurs when an observer does not perceive two colours as being identical, while another observer does. A classical case where this can be an issue are cameras: for various technical reasons, modern digital cameras do not have response functions that exactly match the human eye. Which can lead to colours that human observers perceive as being similar being shown as different in photos, and vice versa. This is not a common problem, but it can occur - particularly for spiky illumination spectra.

Is the phenomenon of metamerism of any relevance for colour reproduction? Or is this something that is nice to know, but irrelevant for any real-world application?

The answer to this question is highly dependent on the type of colour reproduction technology you are considering:

- For **self-luminous devices** (such as projectors, or monitors) the phenomenon is largely *irrelevant*, since we directly observe whatever colours the device is producing, and environmental factors do not have a chance to influence our perception.
  There is one potentially important exception to this statement, though:
  if one e.g. photographs a monitor with a digital camera, it is potentially possible to run into observer metamerism issues, as digital cameras have sensitivities that are close to, but not identical to, human eyes!
- For **devices which produce output that is not self luminous** (i.e. printers), metamerism is potentially *very relevant* (but note that „potentially“ is an important qualifier here!).
  All that such devices produce is a coloration on some reflective medium (usually paper), which then needs an illumination of some sort to be observed. And all the operator of the printing machine can influence are the colours that are being printed, but not the light sources that are later being used when the document is being viewed. In particular, the operator can only make sure that his output contains the correct metamers for one given target illumination. If a different illumination is being used, all bets are off as to the colorimetric accuracy of the result.
  However, and this is the reason for the „potentially“ qualifier used earlier, this is *not* a serious problem in most cases. Dyes, inks and pigments used in most normal printing processes are, amongst other considerations, specifically selected for their low metameric variance, so most printed output only exhibits this effect very weakly in most viewing scenarios. The notable exception to this are settings where printouts are illuminated by fluorescent or LED illumination,

which tends to have very spiky emission spectra. Under such light sources, problems with colour reproduction accuracy can even be immediately visible even to the untrained eye in some cases.

How many types of colour blindness are there?

Two or three, depending on how you count. Two if you consider the L*a*b* channels that fail due to the condition (R-G and B-Y – this is the common way to count them, since this is what is noticed by the affected person), and three if you count the types of sensor cell that can fail (L, M and S – this way of counting is not common, and restricted to scientific literature about the topic).

What forms of colour blindness are there? Why are there only two forms of colour blindness, even though we have three types of receptors in our eyes that can fail?

The two classic forms of colour blindness are actually only selective inabilities to perceive some colours: there is red-green blindness (comparatively common, esp. in males), and yellow-blue blindness (uncommon, similar incidence in both sexes). There is also a third, extremely rare condition that results in total inability to perceive any sort of colour.
A failure of either the L and M sensor cells leads to a symptomatically similar failure of the **R-G** channel. The **B-Y** channel can continue to function properly in both of these cases. Conversely, a failure of the S cones only affects the **B-Y** channel. So both L and M blindness – as they could be called, if they were named for the type of sensor cells that fail – are (almost) similar in their effects on the patient, and therefore subsumed under the term red-green blindness, even if there are actually two different causes for it. Also, note that there is no such thing as „yellow blindness“.

Colour Spaces

What are colour spaces used for? Why are they called colour „spaces“?

For the organisation of the colour stimuli that can be observed by humans.
Because three dimensions are needed for a model that provides a good organisation of these stimuli.

What basic types of colour space do you know?

- **Device Colour Spaces (DCS)**
  These are the colour spaces that are implicitly defined by actual colour output devices – the RGB of a given monitor, or the CMYK space of an offset printing machine.
- **Colour Metric Spaces**, or **Universal Colour Spaces (UCS)**
  Colour spaces that contain *all* colours that can be perceived by human observers belong in this category. Examples are CIE RGB, CIE XYZ, CIE L*a*b*, CIE L*u*v*
- **Colour Ordering Systems (COS)**
  These systems primarily facilitate the selection of colours. Note that a few UCS, such as CIE L*a*b*, can also be used as a COS (although L*a*b* is usually referred to as RAL Design when used in this way). Two sub-classes of COS exist:
    - **Absolute Colour Ordering Systems** An example of such a system would be the Munsell COS – at least in theory, it contains all visible colours, and each Munsell colour value describes a unique, globally valid colour (i.e. each Munsell colour has one, and *only* one, corresponding L*a*b* value). NCS, DIN and RAL Design are other examples.
    - **Relative Colour Ordering Systems** All systems that just re-order the colour values of some „parent“ colour space so that they become more user-friendly belong in this category; HVS and HLS are typical examples. HLS colour coordinates are *not* unambiguous colour values; if one does not know for *which* RGB parent space these coordinates were specified, they are just a relative reference in some RGB space.

What are Colour Collections? How do they relate to Colour Spaces?

Colour Collections are just what the name implies: indexed collections of colours, such as the RAL, BS2660 or FS595 colour lists. Their only relationship to genuine 3D colour spaces is that they are sometimes confused with them, since both collections and 3D colour spaces can be used to select and specify colours. Major paint manufacturers, such as Dulux or Alpina, often also have their own in-house colour collections that are presented in an atlas for customers to choose colours from.

Why is there more than one Universal Colour Space? UCS all cover the entire gamut of human vision, so why bother having so many of them?

All UCS cover the entire gamut of human colour perception, and are therefore of similar size. They do differ – and in some cases differ significantly – with respect to their internal structure, though. Each UCS is tailored to some specific need, be it close relationship to the CIE CME (CIE XYZ), uniformity coupled with the possibility of having a chromaticity diagram (L*u*v*), or uniformity coupled with coordinate axes that are fairly direct correlates of the signals the visual system uses (L*a*b*). It should be noted that the space of human colour perception is inherently non-euclidean due to the various non-linearities involved in the process, so there never will be a „normal“, euclidean 3D UCS space that faithfully captures *all* aspects of human perception.

What is the difference between light and reflection colours / solid colours?

Light colours are self-luminous emission colours, while *reflection colours* (also referred to as *solid colours*) are the result of incident illumination being reflected off some surface. Main difference: light colours are *independent* (in the sense that the colour sensation they evoke in an observer is a stand-alone phenomenon), while reflection colours are, in addition to the colour of the surface itself, always also *dependent on the colour of the illumination*.

What can you say about the purity of light and reflection colours?

Reflection colours that are caused by regular absorption from colourant molecules in practice are never be as pure as saturated light colours; the latter can e.g. be pure spectral colours, which can never be re-created as normal reflection colours. See the images on the slides about bad gap and band pass colours. Reflection colours are therefore only a subset of all possible colours. See the points about structure colour below for exceptions, though.

Are there exceptions to the rule that solid colours are never as pure as light colours?

Structure colours (e.g. interference colours on butterfly wings) can be just as spectrally pure as light colours.

What are structure colours? Describe their properties? Differences to solid colours?

Structure colour is caused by interference effects due to micro-structures on object surfaces. Classical examples are some minerals (opals, labradorite), butterfly wings, beetle carapaces, or peacock feathers. Typical properties of structure colours are:

- They are highly resistant to ageing and bleaching (loss of colour due to UV radiation). As long as the micro-structure of the material remains intact, the colours remain exactly as vivid as they always were. Peacock feathers found in Egyptian tombs were just as colourful as when the bird was alive, 3000 years ago. Compare this to the loss of colour intensity in most fabrics and other organic matter over time.
- They almost always display at least some degree of gonio-chromism (change of hue and intensity with viewing angle, like on the surface of a CD-ROM or DVD disc). By contrast, normal, pigment- or dye-

based colours usually do not change hue a lot with viewing angle: any changes that do occur are near grazing angle, and are usually not due to the pigments.

What are subtractive and additive colour mixing? Corresponding colour spaces?

Additive: mixing of coloured lights, RGB
Subtractive: mixing of colours, e.g. on paper, CMY

What do the letters in „CMYK“ stand for (in particular: what does the „K“ stand for)?

Cyan **M**agenta **Y**ellow **K**ey. In the language of the printing industry, „key“ has long been used as the name of the black separation. **B**, which theoretically could also have been used as an abbreviation of „black“, is not used; instead, it signifies **B**lue.

Colour Matching Experiments (CME)

What is the goal of the CME?

A quantitative metrology of colour, i.e. the ability to describe all visible colours by an unambiguous number (or, as it turns out in practice, a triplet of numbers).

What are – at least compared with other metrology tasks, such as measurements of length and weight – the fundamental issues with this?

One should always keep in mind that this is about a metrology for a subjective sensation of the human mind, which can never be as rigorously measured as purely external things like weight and length.

How does one avoid the problem of not being able to conduct measurements directly on the human optical nerve?

By performing indirect comparison experiments. Discriminability of stimuli is decided at the retinal level, i.e. before the brain comes into play, and yields direct correlates of the sensor functions. Due to the nature of the comparison experiments that are conducted, one can then establish a theory of the absolute behaviour of human vision.

How does one conduct a CME? What colours are shown to the test person?

See the slides for images. Important detail: for the CIE CME, the test persons were only shown spectral test colours. This is not a requirement for a CME to work, though - it just makes it more a lot more efficient, and the experiment more reproducible.

What is the point of having „negative“ light in a CME?

The negative light is needed for those cases where the test light cannot be matched with a normal mixture of the given primaries.

How do you compute the CIE RGB values for an arbitrary input spectrum?

One has to perform the multiplication / integration shown on the slides. The input spectrum has to be multiplied point-wise with each of the three CIE RGB primaries, and the area under the resulting curve has to be computed. These three values are then the RGB coordinates.

What is the result of the CIE CME?

The standardised CIE RGB space – i.e., colloquially speaking, *„the one with the negative component in the red matching curve“*.

What are the units on the axes of the result diagram that you get when performing a CME? What kind of information can you directly see by examining this diagram? What kind of information requires additional calculations?

The *x*-axis stands for the wavelength of the test light (the test persons were only asked to match monochrome lights), while the *y*-axis shows the intensity of the three primaries that were needed for a match. For a given monochrome light, one can directly see the needed RGB intensities that are needed for a match (in the plot, this is shown for blue-green light of 510nm). One can not, however, directly see the RGB coordinates of complex, non-monochrome input spectra.

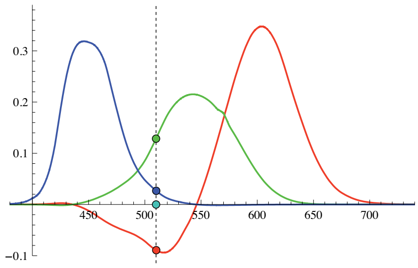

What is the difference between the 1932 and 1964 CIE colour spaces? Do the 1964 versions replace the 1932 ones? Which of the two is relevant for Computer Graphics?

The difference is in the viewing geometry of the test samples that the matching is done for. The original 1932 experiments were done for a viewing angle of 2°, which roughly corresponds to the size of the fovea, i.e. the area on the retina that contains most of the colour-sensitive photoreceptors. The 1964 experiments used a 10° field of view, and provided matching curves for so-called *wide-field viewing conditions*. Therefore, the 1964 version of the matching curves does not replace the 1932 one; they merely describe a different kind of viewing environment.
Since Computer Graphics usually deals with computations that attempt to determine the colour of individual pixels (which subtend a tiny viewing angle), the 1932 primaries are almost exclusively used.

What are the differences between CIE RGB and other RGB colour spaces, eg sRGB?

There are no fundamental differences, apart from the fact that different primaries are used. The only noteworthy differences to commonly used RGB spaces are that CIE RGB has monochrome primaries, and that it has no gamma. Also, it was probably the first RGB space that was specified in a rigorous fashion, and it is the basis for (resp. the ancestor of) CIE XYZ.

What is the purpose of the YIQ/YUV/YCbCr/YPbPr colour spaces?

These all originate from TV broadcasting, and all use some form of luminance/colour separation that was a) originally used to allow backwards compatibility with black and white TV sets, and b) designed to optimise resource usage with respect to the bandwidth needed for transmission (much lower resolution in the colour channels, compared to the luminance channel).

Could one build a monitor that directly displays CIE RGB?

In principle, **yes** - and there is a chance that such devices will actually appear in the foreseeable future. Projectors that use red, green and blue lasers as primaries are now a technological possibility (and indeed, in some areas like cinemas, probably the future), so such devices even might become somewhat common. Also, modern OLED displays have near monochrome primaries, so they already come close to the sort of large gamut a true CIE RGB display would have.

However, it is important to note that even a perfect CIE RGB display that uses 100% monochrome primaries which exactly match the frequencies used in the original 1932 experiments **would not actually be able to reproduce all colours a human observer can perceive**. All a CIE RGB monitor would have is a fairly large gamut: but it would **not** cover all of CIE XYZ. The only reason the colour matching experiment could match all colours with the CIE RGB primaries was because in that particular setting, one could “cheat” by adding the primaries to the stimulus, if no match on the mixture side was possible via positive mixture. But this trick is not possible when displaying an image: all one can do there is additive mixture of positive RGB values.

(This question is mainly intended as a nit-picking counterpart to the question „Could one build a monitor that directly displays CIE XYZ?“, to which the answer is of course „no“ - unless one starts a huge circus and builds a spectral device with at least 10 (or more) evenly spaced monochrome primaries, which then together could reproduce almost arbitrary pixel spectra… i.e. something that is not really feasible anytime soon)

Bonus question: why does only the red matching curve have a negative section?

If the CME is done with non-monochromatic primaries (which is perfectly possible), all three curves have negative sections. However, the CIE RGB primaries are a) monochrome, and b) at frequencies that - either by design or coincidence - minimise the negative areas for two of the curves. That only one of the curves has negative values is *not* an intrinsic feature of a CME!

CIE XYZ

What issues remained with CIE RGB that one wanted to address in future developments? What colour space was developed to satisfy these requirements?

The counter-intuitive negative area in the red channel; one wanted a colour space that contained all visible colours, and that only had positive coordinates for them.

End result: CIE XYZ. The transform between CIE RGB and CIE XYZ was found by trial and error, according to the set goals of „all positive coordinates“, and „G should transform to V(λ)“.

What kind of mathematical transformation is used to convert values between CIE RGB and CIE XYZ, and between arbitrary RGB spaces?

In both cases, a 3x3 matrix is used - a linear transform, in other words.

What are the main properties of CIE XYZ? What were the design goals of CIE XYZ?

- all components should be positive everywhere
- the Y component should have the shape of V(λ)
- which turns X and Z into pure colour channels

The modified matching curves have exactly the same semantics as the CIE RGB curves: for a monochrome stimulus, one can directly read the XYZ values from the diagram. For complex stimuli, an integration is necessary.

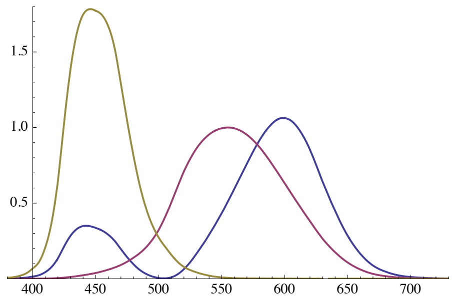

How does one compute the CIE XYZ coordinates for a given input spectrum?

One has to perform the multiplication / integration shown on the slides. The input spectrum has to be multiplied point-wise with each of the three CIE XYZ primaries, and the area under the resulting curve has to be computed. These three values are then the XYZ coordinates.

Is the CIE XYZ colour space suited for light and reflectance calculations?

No, mainly because – unlike RGB – it is not closed under multiplication. However, if it is used for this purpose, the results are actually fairly plausible (as long as they remain in the gamut of XYZ colours).

What are the V(λ) / V´(λ) functions?

A function that describes the general brightness sensitivity of the human eye for varying wavelengths. This function has a maximum in the green area of the spectrum for daylight adapted eyes – the light blue **V(λ)** curve. For eyes that are adapted to low light conditions, the sensitivity changes to the pink **V´(λ)** curve, with a maximum in the blue area of the visible range.

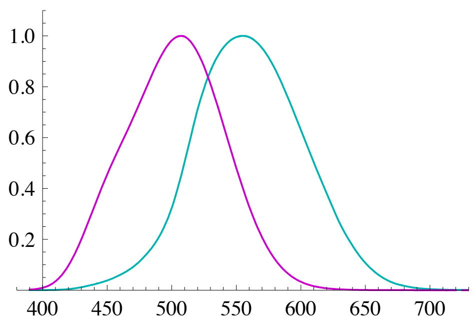

What is the x,y chromaticity diagram (CD)? Please sketch it, and describe it!

Colloquially speaking, you can „remove“ the luminance channel from CIE XYZ colour coordinates, so that only the pure colour information encoded in the colour values remains. In a diagram of the resulting 2D coordinate system only colour hues are visible, and no luminance information.

The diagram has the shape of a shoe sole / tongue / surf sail / camel back / whatever (see slides), lies in the positive quadrant of the plane, and is organised so that pure spectral colours lie on the outside, and that saturation decreases toward the neutral point in the middle of the diagram. One key property of this diagram is that for two given colours, all other colours that can be created through additive mixture of these two lie on a straight line between them. Also, for three given colours, all mixable colours lie inside the triangle defined by them.

Does the CD really contain all visible colour hues?

In principle, yes – but there are some issues with this. In particular, there are certain hues which seem to be missing – with brown being the most prominent example. Technically speaking, brown is just a dark, de-saturated orange, but it is perceived as being a separate hue by most people. Due to the missing luminance information, this and other similar „special“ hues are not present in the CD.

Why is there a straight line at the bottom of the CD?

This line is a connection between the shortest and the longest wavelength that humans can perceive, and it is usually referred to as the *magenta line*. All intermediate colours that lie between the opposite ends of the spectrum (i.e. shades of purple and pink that result from mixtures of red and violet) are located above this line.
The statement that *„no real colours lie on the magenta line“* (this is stated in some older computer graphics textbooks) is true, but - at least in the opinion of the lecturer of this course - also potentially confusing. Since the extreme ends of the visible spectrum (i.e. the two sharp corners of the CD) are indeed the colours where human vision, as it were, *ends*, any colours mixed from these „end colours“ are themselves not really perceptible either. However, even colours infinitesimally close to the magenta line are already (barely) perceptible: only the colours on the actual magenta line itself are not valid colours anymore. So mentioning this specific feature of the CD is of course mathematically rigorous, but somewhat pointless.

L*a*b and L*u*v

Which problematic property the CIE XYZ colour space was documented by MacAdam?

CIE XYZ is not perceptually uniform, i.e. Euclidian distance within the colour space does not correlate well with how different two colours are perceived by a human observer - at least not across the entire colour space. This was documented through experiments that measured how far (in XYZ coordinates) one could move from a given stimulus before a difference became apparent to a test person: the resulting ellipses are very different in size for different regions of the (x,y) chromaticity diagram (see slides). Please note that this non-linearity is not just a two-dimensional issue: the MacAdam ellipses are usually only visualised in the (x,y) chromaticity diagram, in two dimensions. But in CIE XYZ, the luminance axis is also not perceptually uniform, so all three axes need to be distorted to achieve some semblance of uniformity.

Why is perceptual uniformity important?
What kind of application requires this property?

Because without it, one cannot reasonably quantify differences between colours through the euclidean distance between them. This is of crucial importance for a number of applications, such as quality control, where a uniform threshold of error tolerance is required, regardless of which hue is being measured.

Which colour spaces were defined to solve this problem?

L*a*b* / L*u*v*

What is the difference between the Y channel of CIE XYZ, and the L channel of CIE L*a*b*/L*u*v*? Both are monochrome „brightness“ channels - why do they use different letters?

The meaning of the letters **Y** and **L** for the brightness channels of colour models is unfortunately not completely consistent for all models. Usually, **Y** denotes a channel that is unlimited in terms of maximal brightness (such as with CIE XYZ), while **L** (for **L**uminance) is used for relative measures w/r to some white point (as with CIE L*a*b*, where L is in the range 0..100). This semantic difference is not strictly adhered to in all cases, though. The one property that is reasonably consistent is that **L** is only used for channels that are more or less perceptually uniform.

Why were *two* perceptually uniform colour spaces defined in 1976?

Because each of them has some desirable property that the other does not have, and that cannot be combined in a single model due to the inherent non-linearity of human perception.

What are the advantages and disadvantages of CIE L*a*b*/L*u*v*?

L*a*b* has the advantage that the co-ordinate axes are physiologically meaningful, since they correspond to the channels of human colour vision.

L*u*v* has the advantage that one can still draw a chromaticity diagram for it; this was a desired feature for use in the TV industry at the time.

How do you compute the perceived distance between colours in CIE L*a*b*/L*u*v*?

As Euclidean distance, e.g. in the case of L*a*b* $\Delta E = \sqrt{\Delta L^2 + \Delta a^2 + \Delta b^2}$.

What is a Chromatic Adaptation Transform (CAT)?

A - fundamentally inaccurate, but still quite useful - method to predict the appearance of solid colours under varying illuminants *from their colour values alone*. CATs are 3x3 matrices that are solely derived based on the source and destination white points (i.e. illumination colours).

Which CATs do you know?

Von Kries transform (older, mostly deprecated), and Bradford transformation (most popular transform at the moment, used e.g. by Photoshop). CIECAM02 also has its own CAT, that is structurally similar to the other two (also a 3x3 matrix, just different coefficients).

What is the *white balance* operation in a digital camera?

White balance computations is the term used to describe the process of removing a colour cast that was caused by the illumination that was present when an image was taken.

A good example would be the behaviour of our visual system in a room that is only illuminated by candlelight; such an illumination is actually very orange in colour. However, while being immersed in the scene, one would not perceive the objects in the room as being particularly orange, but rather as still exhibiting their "real" colours. The colour cast of the ambient illumination would certainly be noticeable, and would be perceived as "warm" in this case, but one would not be under the impression of being in an environment that is largely reduced to varying hues of orange. Our perception adapts to its surroundings, and compensates the dominant colour cast to some degree.

A digital camera records images in its fixed device colour space, regardless of the environment conditions. If no correction is applied, images taken in the candle-lit room would be perceived as being much too orange, if viewed on screen under normal daylight conditions later. The scene *was* very orange (just like the camera recorded it), but was not *perceived* that way by an observer that was immersed in it. Therefore a correction that is intended to mimic the adaptation process of the human eye is applied, with the goal that the photographs then represent a reasonable approximation of how the scene was perceived.

What does a digital camera use CATs for?

To perform the white balance computations discussed in the previous question. A Bradford transform (or von Kries, or CIECAM02) is the technical implementation of such a colour shift.

What is a *Colour Appearance Model* (CAM)?

A mathematical model that allows a prediction of how colours will be perceived under a specified set of extended environmental conditions. This includes white adaptation, but also the surround and background luminance of a given viewing situation.

What is typically the output of a CAM? How are these output values used?

A CAM usually yields 6-7 *correlates of human perception*, like brightness, lightness, chroma, colourfulness, and so on. Note that these are not directly displayable quantities, and that the 6-7 properties are somewhat redundant. The space of these values is what is referred to as „CAM space“ in the flowchart in the next Q&A pair. Normally, these values are only used as inputs to *another* CAM that is run „in reverse“.

What is the relationship between ICC profiles, CATs, and CAMs?

In their default usage scenario, ICC profiles are only used for colour transforms between devices that are subject to roughly the same viewing conditions (such as white adaptation, or surround). Theoretically, the mechanisms in an ICC workflow could be used to perform transformations between dissimilar viewing conditions, and this is actually sometimes done. However, it is best to explicitly use CATs and CAMs for transforms between different viewing situations. As can be seen in the following diagram, ICC profiles play a role even in the usage scenarios that perform CATs or CAMs.

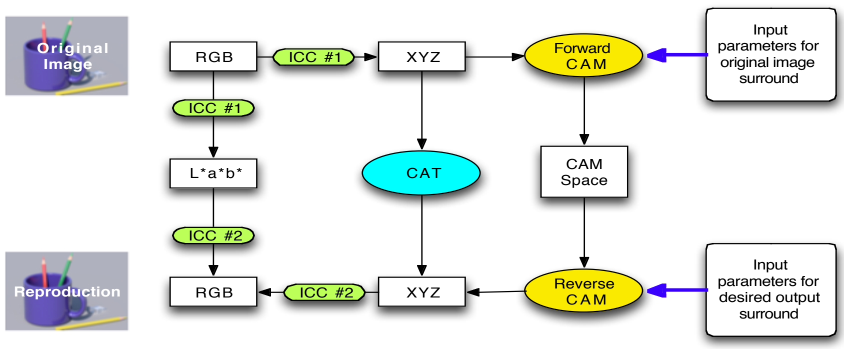

What is the main difference to CATs?

CATs are solely driven by the difference between source and destination white points, and can be seen as just one aspect of a CAM (albeit an important one).

CAMs go beyond CATs in that they include other factors which affect the viewing situation, such as shape and size of the viewed colour stimulus, the surround and background luminance. This enables a much more sophisticated prediction of colour appearance under varying conditions.

What is the most common CAM in use today?

CIE CAM 2002, which is usually simply referred to as CIECAM02.

What is the main limitation of conventional CAMs like CIECAM02?

CIECAM02 deals with a single stimulus, so it is by definition not capable of handling complex appearance prediction for images (e.g. predict the various appearance phenomena one can experience when viewing complex images, i.e. optical illusions and such). Techniques like iCAM06 or the recent work of Kim et. al. *(„Modeling human color perception under extended luminance levels"*, SIGGRAPH 2009) go beyond plain CAMs, and are, at least up to a point, still research work in progress.

What is an *Image Difference Predictor*?

IDPs are models that attempt to reliably quantify how different a (human) observer will perceive two images to be. A simple metric of this kind would be e.g. the average delta CIELab difference between the images, or the RMS difference between the colour coordinates. But while simple metrics do yield results (in the sense that they are numerically stable, and are always able to compute some number), their practical usefulness is rather limited. Human perception is far too complex for such simple approaches to yield robust correlates of human difference perception. Modern approaches, like e.g. the techniques introduced by Martin Čadík et al., are based on training data obtained from human observers.

What is the connection between iCAMs and Image Difference Predictors?

The problems faced by these two types of model are, at least up to a point, quite related: however, the techniques used by the two communities are different. IDPs rely on machine learning, and one can expect that such models will become excellent at the task of determining whether a human observer would see a difference between two images. However, even though machine learning approaches for image *creation* have made enormous progress in the past years (think Deepl, midjourney and such), there is no standard model resp. metric for this task yet that could be routinely used by everyone.

Colour Ordering Systems (COS)

What are *Colour Ordering Systems* used for?

They are used to choose colours according to certain criteria. Often, the ability to easily choose several matching colours at once in some systematic fashion (e.g. by picking three hues that are 120 degrees apart in hue space) is also an important feature that a COS offers.

What is a *colour atlas*, and what is something like that used for?

A ring binder (common), or an actual book (rare) with physical samples of the colours that are defined by a particular COS. Often, the samples can be removed from the pages of the atlas (they are usually small cards, or colour chips), and used for in-situ comparisons, and design evaluations. Some systems also offer loose samples in boxes. The benefit of such an atlas - which is normally used by designers - is that it provides a source of real colour samples that can be viewed alone, placed next to each other, or evaluated in context (i.e. compared to existing objects that one wants to match the colour of - they are actually a crude colour measurement tool as well). For some types of design work, such a hands-on tool is far more convenient and intuitive to use than just having an on-screen colour picker.

Give some examples of COS

Munsell, NCS, RAL Design, DIN, Coloroid.

Describe the Munsell System

Defined around 1900 by Albert Munsell as a perceptually uniform colour space. Based on experiments with colour chip comparisons. Irregular shape of the colour solid as shown on the slides. 5 (!) basic colours which - as a slight oddity, since no other colour space has this as a genuine primary - include purple. Colour coordinates are Chroma, Hue and Value. A colour atlas is sold for this system.

Describe the NCS System

See the slides, developed in Sweden by the local standards authority, primarily for interior design purposes. A colour atlas is available.

Describe the RAL Classic and RAL Design systems

RAL Classic – a.k.a. „the four digit colour codes found in European DIY stores“ – is a simple enumeration of colours that have been registered with RAL (***R**eichs**a**usschuss für **L**ieferbedingungen* – literally *„Imperial Committee for Delivery Specifications“*, which sounds incredibly anachronistic to current ears) in Germany during the past 100+ years. RAL was founded in the 1920s to serve as an entity that provides standardisation of all sorts of aspects that are relevant to an efficient commercial environment, with colours being one of them. The RAL Classic codes are loosely grouped by their major number, with e.g. numbers of the form 1xxx being yellow hues, 3xxx red, and 5xxx blue. However, this was just done for convenience‘s sake, and is not really systematic.

RAL Design is a front-end for CIE L*a*b*, or rather, CIE LCH. A colour atlas for this system is available.

Describe the FS 595 system

Like RAL Classic, this is a pure colour collection with no inherent ordering, and is primarily used for government contracting purposes in the USA. Colour numbers are given as five digit numbers that are preceded by the letters „FS“. For some colours, a common name may be available. Example: FS 36375 Ghost Grey, a colour used in the camouflage of some U.S. warplanes. The first digit indicates glossiness!

Describe the Pantone system

Again a pure enumeration of colours, but with an added feature: the system also contains standardised instructions on how to create the colour in question from a given, small number of basic colorants. Main use: the printing, plastic and branding industries. Example: „Starbucks Green“ is Pantone 3425 C.

Describe the Coloroid system

A mathematically defined colour system that was derived based on experiments with observers, who, amongst other things, were asked to rank colour sequences according to aesthetic criteria. One of the goals was to develop a colour system that enables the user to automatically find aesthetically pleasing colour sequences. The system has the interesting property that its aesthetic choices seem to reflect the aesthetics of the time during which the experiments were made – the 1970s.

Measurement Devices

Can you directly convert radiometric quantities to photometric ones, and vice versa?

No, you always need the original input spectrum to compute the other type of unit. Knowing the numerical value of a radiometric unit does, in general, not help you in determining the corresponding photometric value. The exception to this rule is monochrome radiation (such as lasers).

What is the difference between radiometry and photometry?

Both concern themselves with the measurement of electromagnetic radiation. Radiometry does this over a freely defined frequency range (that can, but does not have to, coincide with the visible range), while photometry is inherently limited to the visible range. In addition to this range limitation, one always multiplies the input spectrum with V(λ) before computing a photometric unit. For all radiometric units, there is a photometric counterpart that has been weighted with V(λ).

A practical example is best to clarify this difference: for a given spectral distribution (the blue curve), a radiometric measure of incident energy corresponds to the area under the curve:

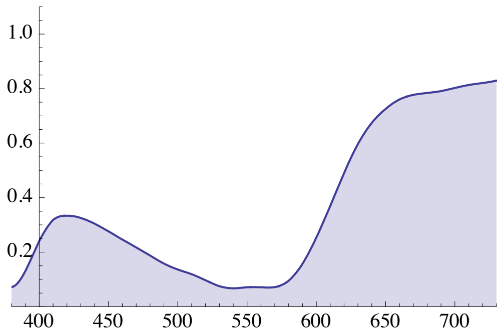

The corresponding photometric quantity is computed by first performing a point-wise multiplication of the original spectrum with V(λ), and only *then* computing the area under the resulting curve:

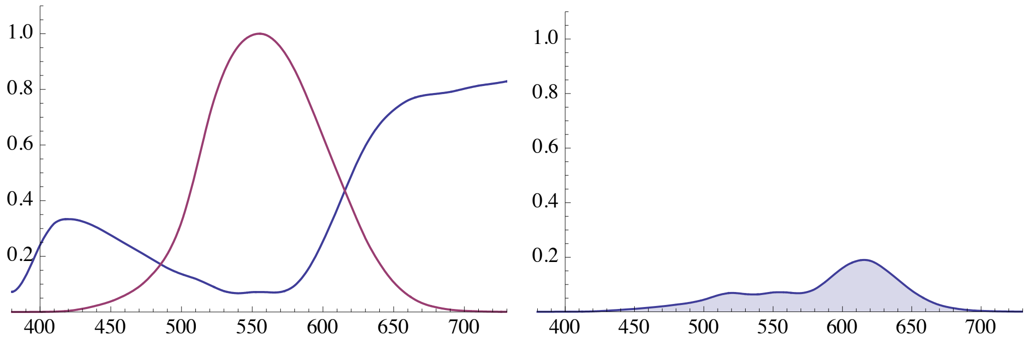

This yields a measurement that is a direct correlate of the human perception of brightness. In this particular example, one can clearly see that in terms of integral measure, the photometric quantity is much smaller than the radiometric one, since the input spectrum has comparatively little energy in the green range of the spectrum, to which human eyes are most sensitive.

Who needs photometric quantities?
Why do we not use radiometric ones all the time, and be done with it?

Typical users of photometric quantities are illumination engineers, or interior designers, who have to meet certain mandated, ergonomically viable illumination levels for their designs.

For such applications, the key advantage of photometric units is that they directly correspond to how bright a person will perceive a given light to be – *independent* of its colour, and spectral composition. Such a prediction is in no way possible with the corresponding radiometric units.

Their achromatic nature relegates photometric units to secondary importance within *colour* science in the narrow sense of the word. However, they are so frequently encountered in practice in colour science settings (i.e. as specifications of brightness during monitor calibration routines) that they have to be covered in this course.

What are reference light sources (RLS)?
What types of RLS do you know?

Such light sources of exactly known emissivity are used to calibrate measurement devices.

There are two ways to build such light sources in practice:

- **Reference Lamps**
  Incandescent bulbs are subject to large variations between individual bulbs, and they also alter their emission characteristics significantly over their lifespan due to ageing. However, if manufactured under carful quality control, and if their ageing process is properly monitored, such lamps can be used as lesser sources of device calibration.
- **Phase Boundaries**
  During a phase transition – e.g. when liquid water still has some pieces of ice floating in it – the temperature of a substance will remain at the melting point, until all of the substance has reached the next phase state (i.e. until all the ice has finally molten). This kind of effect can be used to construct a device that emits blackbody radiation at a very precisely known temperature, and therefore very precisely known spectral composition. As an added bonus, this radiation also remains stable over some time, if the experiment is properly set up. To be useful for emission calibrations in the visible range, though, molten metal has to be used, which severely limits the applicability of this method to any setting except reference laboratories.

What is the purpose of a *viewing booth*?

Viewing booths are simple devices that allow the viewing of objects under controlled lighting conditions. Usually, such booths are just large boxes with neutral grey inside colouring, and a selection of representative lightsources (D65, A, fluorescent) mounted inside.

What are reflectance standards (RS)?
What types of RS do you know?

These are reflective objects used to calibrate measurement devices. There are two forms:

- **Transfer Standards**
  These objects are measured to very demanding tolerances, and used for inter-device calibration purposes. They have to be very resistant against ageing and environmental influences.
- **Device Standards**
  These are objects that are associated with a given device, and used to repeatedly calibrate the device (the properties of the reference object are often programmed into the firmware of the device, and the object has the same serial number as the main device).

What types of measurement device do you know?
Give a brief description of each

- **Luminance Meters**
  Such devices just measure perceived brightness in Candela/m^2, and do not take any colour information into account (mainly used in photography to compute exposure times and such). The sensor is usually a single photocell or photodiode.
- **Colour Measurement Devices**, a.k.a. **Colorimeter**
  Such devices directly measure colour values. They have three photocells that have been modified to roughly have sensitivities that correspond to CIE X, Y and Z (the modification is usually done by gluing coloured filters over the photocells). It is worth noting that it is *not* impossible to have a sensor that has a sensitivity that directly corresponds to CIE XYZ – after all, our eyes also have that characteristic. It is just impossible to build a *display* that uses XYZ as primaries.
  Some of these devices have internal light sources (usually a „classical“ standard light, such as D65), and can be used to measure reflectance colours; the returned colour values are then of course only valid for the light source that is built into the device.
- **Spectro-Radiometer** (for the visible range only, a.k.a. **Spectro-Photometer** - see next question for details on this)
  Such devices measure the spectral composition of incident radiation. Versions with and without internal light sources exist; only the former is capable of measuring reflectance colours.

What is the difference between a Spectro-Radiometer and a Spectro-Photometer?

There is no difference, at least not as far as the basic optical technology in the device is concerned. The generic term for a device that is capable of measuring electromagnetic radiation in a frequency-resolved manner is *spectro-radiometer*. This expression can be used for all such devices, even for those that are built to resolve radiation in areas of the spectrum that are far from the visible range. Devices that are designed only for the visible range, in particular those used by the printing and photographic industries, are sometimes referred to as *spectro-photometers*, though, in analogy to other devices used there (e.g. photometers). This nomenclature is of debatable merit, since all other devices with the *photo*- prefix usually measure photometric quantities (e.g. lumens), and not radiance values like spectro-photometers are still bound to do.

Can you use a colorimeter to calibrate a monitor?

Yes, because a monitor is self luminous, and one does not have to know the spectral composition of the light it emits to judge its appearance.

Why do graphic artists still use spectrophotometers most of the time?

Many graphic artists need a spectrophotometer to calibrate their printer. Once they have such a device (which can, apart from calibrating printers, of course also be used for monitors), it would be pointless to buy an extra colour measurement device just to calibrate their screens with.

When is the purchase of a plain colour measurement device for calibration purposes justified?

If it is certain that only self-luminous output devices will be calibrated (say, if the device is purchased solely to occasionally calibrate a home cinema set-up), a colorimeter is sufficient.

Why does one need a spectral measurement device to profile a printer? When is a colorimeter sufficient for this task?

Spectral reflectance data is needed to accurately predict what a given printout will look like under an illumination that is different from the one installed in the measurement device.

If (and only if) one is solely interested in calibrating a printer to produce output that is valid for the illuminant that is in-built into a particular colorimeter, then such a device is also sufficient for the calibration of a printer. In practice, this will not be a very common usage scenario, though, especially since most colorimeters on the market do not contain a lightsource, and are intended for emission measurements only.

Printing Technology

What types of printing technology do you know?

Laser, ink jet, offset and many others – but these three are the most common ones.

What was *Pantone Hexachrome*?

A (now deprecated) standardised print technology that used pure green and orange in addition to the usual CMYK primaries; the abbreviation is CMYKGO. This could only be used on offset machines, but (unsurprisingly) offered a much larger colour gamut than plain CMYK. There are several such systems (Hexachrome was a product of Pantone Inc.), and the collective term for them is „high fidelity printing“.

Describe how a laser printer works

See the slides. From a colour perspective, the most important aspect is that the toner particles are not pure colorant - there is a significant amount of carrier substances involved, and the colorants are often pigments, not dyes. As a consequence of this, the achievable colour purity is often below that of inkjet and offset printing processes, and more in the league of enamel paints. However, in recent years, the colour quality of laser printers has improved considerably, and is now (almost) on par with ink-jets.

What are the advantages and disadvantages of laser printers?

Advantages: speed, average to low cost per page, stability of output

Disadvantages: cost of device, sometimes still average to low colour quality, colour laser printers are quite complicated (4 small laser printers in one device, one for each channel of CMYK)

Rate the colour output quality of laser printers?

Medium to low, at least compared to dye and ink based systems - but this is slowly improving.

Describe how an ink-jet printer works

See slides.

What are the advantages and disadvantages of ink-jets?

Advantages: low device cost, good colour quality, can print on overhead transparencies (this is a *very* obsolete advantage, though)

Disadvantages: low speed, high costs per page, heads can dry out (which usually means an expensive replacement is needed), results not very stable (at least with cheap ink - there are also archival inks that are guaranteed to last decades without fading)

Rate the colour output quality of inkjets?

Normally very good to excellent, especially on special paper.

Describe the principle of offset printing

See the slides, important facts are that this is a) a technology that is exclusively used in the domain of professional printing, b) the only technology that is routinely used to print with more than CMYK as primaries (CMYKGO, spot colours), and c) a technology that is orders of magnitude more powerful than the others in terms of output capacity (pages per minute).

What are the advantages and disadvantages of offset printing machines?

Advantages: very low cost per page (provided you need a huge number of copies to begin with), very high speed (once the print job has been set up), easily expandable

Disadvantages: only for professional printers, high setup costs, plates have to be produced

Rate the colour output quality of offset printing machines?

Very good (normal CMYK) to excellent (Hexachrome).

ICC Profiles

What is an ICC profile used for? How is this goal achieved in practice?

To completely characterise the properties of an input or output device with respect to colours.

By analysing the difference between what *should* happen during an I/O process that involves colour (printing, scanning, display on a monitor), measuring the actual result, and storing the difference in the profile. This difference data can then be used to compensate the behaviour of the device accordingly. Note that for input devices, objects with known colours (colour charts etc.) have to be available.

What is the *Profile Connection Space* (PCS)?

The colour space that is used as the centre of the star topology used in ICC transforms. ICC profiles either offer transformations from the PCS to a device colour space (for output devices), or from a device colour space to the PCS (for input devices). Normally, CIE L*a*b* is used as default PCS for ICC profiles.

What problems can occur in practice when working with ICC profiles?

For most applications, the standard star topology is sufficient. There is at least one area where it can cause problems, though, and that is if pre-separated CMYK data, and/or files that contain spot colours are present in a workflow that requires that this additional information be retained. If such files are handled using the standard star topology, any information about their special properties (such as black composition, or spot colour separations) is lost during the transform via the PCS. L*a*b* just stores colour values, and has no concept of spot colours or black composition.
The solution for this problem is to avoid the star topology in these rare cases, and to use *device link* profiles for these transforms.

What is a *Device Link ICC Profile*?

Such special ICC profiles allow the direct conversion between two device colour spaces **without** the detour via the PCS. These direct transforms can be tailored to retain black composition, spot colours and similar properties, and to just transform the individual colours between the two colour spaces.

A classical usage scenario for such a device link profile would be the adaptation of a pre-separated PDF/X-1a that has been prepared for Euroscale Coated CMYK to, say, U.S. Sheet-fed Coated CMYK. This is a rather subtle change - slightly different printing inks, but the same paper type. A device link profile can achieve this conversion without altering the black separation that is already present in the file.

What kinds of information does an ICC profile contain?

Usually point sample measurements that allow a compensation of any colour deviations from the desired device performance. Monitor profiles normally only contain the RGB matrix and gamma curves, and several other sub-types of ICC profile with different content also exist (colour lists, colour spaces, etc.)

What are *Rendering Intents* (RI)?

Rendering Intents is the collective name the printing industry uses for various common, pre-defined gamut mapping strategies. Note that “pre-defined” is not an exact term in this context; see the section on PDF-X for a further explanation of this.

What are the main types of Rendering Intent?

The two main types of RI are:

- **Photographic**
  All colours in an image are moved towards de-saturated hues until they all fit into the target gamut. Advantage: the relationships between colours are maintained. Disadvantage: saturation is lost, the image appears to be too grey afterwards in some cases.
- **Colorimetric**
  Here, only those colours that are outside the target gamut are moved, and those that are already inside are left alone. Advantage: saturation is maintained. Disadvantage: information is lost in areas of high chroma.

In addition to these, there are several vendor-specific RIs; since these are not standardised, they are of lesser importance. One example is

- **Saturation**
  Somewhat similar to the Photographic RI, except that the colours may be further altered to maximise the colourfulness of the result.

Are the colour transforms that are at the core of ICC colour correction, and the rendering intents, precisely defined by the ICC?

No, and this causes problems in the printing industry. In particular, PDF/X would not have been necessary if rendering intents were predictably similar across vendors and platforms.

What is the difference between *absolute* and *relative* colorimetric rendering intent?

Absolute CRI really does what one expects - it does not alter in-gamut colours in any way. This has the (perhaps not immediately obvious) side effect that all actively printed colours are reproduced correctly, but that the paper white might not be the correct white in relation to the printed colours (this would not be an issue if printers applied their own white to the medium).

Relative CRI compensates for this effect by altering all in-gamut colours so that they have the right distance to paper white. This ensures a correct relation between the printed and unprinted areas of hardcopy output.

What is *depth compensation*?

Depth compensation is an additional correction that can be applied to Relative CRI. Its goal is to additionally transform the black point of the two colour spaces so that they match. This mainly increases the appearance of photographs, and similar images.

How does one create an ICC profile for a device?

Through measurements, and analysis of this data in special programs such as ProfileMaker, or the profile creation tools of LCMS.

What is a *Matrix Shaper* ICC profile?

An ICC profile for additive RGB output devices that only contains the RGB transformation matrix, and the gamma curves for that device. This is much more efficient than using tabulated measurements, and can be directly used to alter the behaviour of (i.e. calibrate) device drivers.

Give examples of ICC profiles that are not classical device calibration tools?

ICC profiles can be used as generic colour transforms for images: so transformations like “convert to greyscale”, “blue tone” or “sepia tone” can be modelled as ICC profiles. One can even reduce all input colours to a discrete number of outputs, as in the case of the obsolete but technically interesting “web safe colours” profile that converts all input colours to the 256 colours that were deemed “web safe” in old web browsers running on dithered 8 bit colour screens.

Is the presence of an ICC profile for a given projector or screen a guarantee for correct colour output? If no, which technical issues can arise?

For arbitrary hardware, the answer is unfortunately **no**, although this is no longer an issue for most modern hardware.

The main causes of problems that can occur are *dithering on the hardware level* (which causes strong non-linearities, esp. in older, cheaper laptop displays), and *projectors with brightness boost technologies* such as some forms of DLP (which again cause non-linearities that can not easily be compensated).

Is colour management on the screen device driver level sufficient for all scenarios?

No, because for obvious performance reasons a screen device driver has to use a colorimetric rendering intent, which might not be appropriate for all applications.

Also, sophisticated colour manipulation software such as photoshop will always have its own internal colour engine *in addition* to the compensation at device driver level; *both* these components are needed for a proper, colour safe working environment. In particular, if no calibrated screen device driver were present, Photoshop would not know what colours the user is shown, and its internal colour engine could not do its job.

What is the *Image Sum Density* of a print job?

ISD is the sum of the maximal intensities of all the channels in a print job. CMYK fully applied has an ISD of 400%, and Hexachrome 600%. Neither figure is possible with real inks, and real paper, though. Newsprint, which uses the cheapest paper available, can have an ISD as low as 230% (i.e. not even three of the 4 channels in CMYK can be fully utilised), while high-quality coated paper can reach up to 380%.

Due to their different printing technology which utilises pigment particles that are not as prone to dot spread, colour laser printers always use a full ISD of 400%.

What is a *Colour Separation*?

Generally speaking, the *colour separation* of an image is the splitting of that image into the primary colours that are being used by the intended output device. Most images are directly stored in RGB colour space, and only CMYK printing uses a different number of channels. In practice, the CMYK version of a given image, and in particular, the CMYK version tailored for a specific printing machine, are what is usually referred to when speaking of the „separation of an image“. Note that there can be many different CMYK separations for a given RGB image; see „black composition“ for details.

What is a *Spot Colour*?

For a print job, this is a particular solid colour that will be printed with a separate ink *in addition* to the normal CMYK primaries. This is done for specialised print jobs that involve large areas in one specific colour, such as e.g. corporate stationery that is adorned with the company logo in a fixed Pantone hue, or specialised print jobs with e.g. gold or silver inks. Due to the effort involved (in particular, the fact that only offset printing is capable of easily adding primaries), this kind of multi-colour separation is not done outside the professional printing world.

What is *black composition*? Give examples of common black composition strategies? Which types of black composition does a given ICC profile support?

Black composition is the expression used for the strategy that is applied to CMYK colour separation in a print job. CMYK is an ambiguous colour model; colours can e.g. be either combined by using as much K as possible, or as little K as possible (with variations in between). For varying paper types, both strategies offer advantages. Examples are:

- **UCR** - Under Colour Removal
  Complete replacement of CMY with K in all neutral areas of the image only.
- **GCR** – Grey Component Removal
  Replace as much of CMY with black as possible in all areas of the image. Low ISD, low colour brilliance, typical for newsprint.
- **Skeleton Black**
  K is only used in those areas where it is essential; good colour rendition, high ISD.

A single ICC profile only contains one particular black composition, which has to be chosen when the profile is created.

Describe the influence of paper type and printing technology on black composition?

The less the paper is capable of absorbing ink, the higher its quality. Less absorbed ink means lower dot spread for a given amount of ink, and allows for higher overall ISD, which leads to greater contrast ratios. Higher ISD means that more genuine black and CMY can be printed on top of each other. Low ISD (as in newsprint) means that very little CMY can be printed alongside K.

What is the purpose of PDF/X-1a and PDF/X-3? What are the differences?

Both formats attempt to define a subset of PDF that has the same purpose: to improve the consistency and predictability of colour output produced by printers, and in particular professional print shops.

Generic PDF files follow a philosophy that is very common in colour graphics programming: each object in a PDF file can have its own ICC profile attached to it, which is a sensible thing to do from an information preservation viewpoint – no information is destroyed until the file gets printed, and converted to a target colour space. There is one scenario where this is problematical, though: if a designer wants his own proofing system to deliver a reliable proof of what the print job will look like once it runs off an offset press. This is practically impossible to do if he sends his print job to the print shop as normal PDF, because the lack of standardisation in the area of rendering intents means that it is virtually certain that the print shop will use slightly different gamut mapping algorithms than he does on his own proofing system. His proof will in all probability *not* match the output of the print shop. The solution that PDF/X-1a offers is to reduce the contents of the entire PDF file to the colour space of the output printing press, i.e. the contents of the PDF are just plain, pre-separated CMYK. This might seem to be a cosmetic change, but it has one critical advantage: it enables the proofing system of the designer to deliver a reliable proof, since all gamut mapping is now done ahead of time by the designer himself, when the PDF/X-1a is created. All the print shop has to do is to feed the PDF/X-1a to the offset press without any further modifications; since there is no need for colour transforms any more. Since the colour space of offset machines is known to a high degree of precision, nothing can happen at the print shop that alters the colours of the output in unforeseeable ways. PDF/X-3 attempts to achieve the same goal, but without reducing all objects in the PDF to a single target colour space. Instead, provisions are made so that any subsequent colour transforms become predictable.

What is iccMAX / ICC v.5?

These are the next generation of ICC functionality, which will substantially expand the usage cases for ICC profiles (see slides). A reference implementation exists, but as of 2019, no commercial products support iccMax in its full extent. Due to the very large scope of iccMAX (up to a point, it tries to do predictive rendering for print and display output), it is unlikely that it will be adopted in coherent form in the near future: the standard is far too complex for this.

Colour-safe workflows with Photoshop et al.

What basic types of colour workflow do you know? Which of the two is more modern? What issues does the more modern workflow have?

Based on CMYK, and based on L*a*b*. The latter is conceptually more modern, but the intricacies of CMYK workflows are such that L*a*b* will not entirely replace them in the foreseeable future (see also the question about device link profiles).

What is the *working space* of a graphical application? Why is this something that can be altered, and why is this even mentioned in the user interface?

The *working space* of an image is the colour space in which its pixels are assumed to lie. This is worth mentioning because this working space does not have to be identical with the *screen space*, i.e. the device colour space of the screen that one is using to work with the image. In most modern environments, the two will definitely not be the same, although this was an implicit assumption in practically all older image manipulation software.

What is the operation *Assign Profile* used for?

An ICC profile is assigned to an existing image, which means that all the colours in the image are from then on interpreted as being colour values in the colour space defined by the profile (usually a particular RGB or CMYK space). No transformation of colour values takes place - this operation just (re-)defines the *meaning* of the colour values stored in the image.

What is the operation *Convert To Profile* used for?

Here, the user specifies an ICC profile to *convert* the contents of an image to (the image already has to have a profile associated with it). If the new gamut is smaller than the original, gamut mapping is necessary, and a rendering intent has to be selected. An example would be a Wide Gamut RGB image that is converted to sRGB (a much smaller RGB space) for output purposes; this almost certainly requires gamut mapping, and information will be destroyed.

What types of image file formats do you know?

There are two fundamentally different types of image file formats:

1. **Relative images (Type A on the slides).** Brightness bounded by an intrinsic maximum value, very often only 8 bit per channel accuracy, integer channel values, rarely 16 integer bits per channel. Example formats: JPEG, TIFF, GIF, PNG.
   Such images are nowadays often referred to as LDR images (Low Dynamic Range).
2. **Scene-referred images (Type B on the slides).** Contain pixel values in absolute brightness units, practically always encoded in floating point.
   Example formats: Log Luv TIFF, OpenEXR
   These images are also (and more commonly) known as HDR (High Dynamic Range) images.

As of 2024, there is now actually a third category that is a hybrid between the two: **XDR image formats** such as HEIC and JPEG XL. There, a fourth “brightness boost” channel is added to a conventional RGB format, so that an extended luminance range can be displayed on modern OLED displays capable of XDR output. Due to the fact that the RGB channels in such images are encoded as normal relative “Type A” images, they display in a backward-compatible way on older hardware: but on modern displays they look far more appealing due to a substantially increased contrast range. However, they are not true “Type B” images in this mode either, as the “brightness boost” encoded in the fourth channel is monochrome: and the pixel values are usually not assumed to be in absolute brightness units, either. Sometimes, such images are referred to as RGB+E (with E standing for “exponent”, as the boost channel is usually in some exponential encoding).

What are RAW images?

In digital photography, RAW images usually contain the raw output of the CCD that is in the camera. Advantage: no information is lost, and post-processing has more data to work with. Disadvantage: larger file size, and special software like Adobe Photoshop Lightroom is needed to convert this data to a displayable LDR format, such as JPEG, or an XDR image format such as AVIF or HEIC. Slightly confusingly, RAW images are sometimes also referred to as being XDR images themselves: this is technically true, as they exhibit larger dynamic range than true LDR formats (2-3 f-stops), but considerably less than „true“ HDR images (which can cover up to 30 f-stops).

What are XDR images?

XDR images (eXtended Dynamic Range) are either of two types of images that go beyond LDR, but fall short of “real” scene referred HDR. These are the “brightness boost” RGB+E formats such as AVIF, HEIC and JPEG XL on the one hand, and camera RAW files on the other. Both types of images are referred to as being “XDR”: and arguably, camera RAW is more of a true XDR format than the RGB+E ones, as the extended range is stored separately for all three RGB channels there - instead of just a monochrome “brightness boost” as with RGB+E.

The main thing RGB+E has a hard time reproducing are very bright but saturated colours, like the red of a setting sun. These either look reasonably good (read: bright red) in brightness boosted +E mode on an OLED screen: but then, the plain RGB value looks odd as a dim saturated red on normal screens, as the solar disc in a sunset photo is normally “tone mapped” to a neutral or yellow-ish colour, to at least give the impression of it being bright on an LDR display.

Conversely, the classic “off white” solar disc cannot be boosted to a bright red in +E mode anymore, as the colour information has been lost. One would need a true HDR format, such as OpenEXR, to store such information. The main reason OpenEXR is not used for consumer camera work is that unlike RGB+E, it does not have an intrinsic backward compatible mode for LDR displays.

What are HDR images?

Scene-referred images that can contain brightness values that go way beyond those that normal display devices can reproduce. A tone mapping operator of some sort is needed to reduce them to a displayable brightness range. “True” HDR would require the RGB values to be scene-referred, i.e. to be in absolute brightness units. OpenEXR is the main HDR format in use nowadays, although mainstream software and operating systems do not properly support it on most platforms yet (that is, most operating systems can display them in some basic form: but their HDR aspect is often not properly shown).

What types of tone mapping operators do you know?

In order from least sophisticated to most sophisticated:

1. **Global operators**: a single scaling factor (or exponential coefficient) is derived for the entire image (either on an ad hoc basis, or via some simple heuristic), and applied globally to all pixels. Advantages: speed & simplicity. Disadvantage: cannot cope well with some scenarios, such as outdoor photos with bright cloudy skies over dark foregrounds.
2. **Local operators**: similar to local operators insofar as some sort of non-perceptual mapping heuristic is being used, but on a locally varying basis. Advantage: can cope with scenes that global methods cannot handle. Disadvantage: can still fail for some scenarios, needs even more ad hoc parameters (such as kernel size and segmentation thresholds) to work correctly.
3. **Perceptual operators**: attempt to use forward and reverse models of human perception to accurately predict which modified/compressed/tone mapped version of an image, when shown on a particular output device, will elicit a viewer response that is as close as possible to seeing the original scene in reality. Advantage: best results. Disadvantage: still an active research area, principle well understood, but not robust enough for practical use yet.

What is the relationship between tone mapping and gamut mapping?

1. **Tone mapping** only reduces the luminance range of an image to displayable levels.
2. **Gamut mapping** reduces the colour gamut of an image so that it fits into the display gamut of a chosen output device.
3. **Tone reproduction** is sometimes used to describe the whole process of tone and gamut mapping.

3D Print

What issues exist when trying to reproduce colour textures on polyjet 3D printers? What can be done to compensate?

All current polyjet materials are partially translucent: they have to be, to allow curing of the resin via UV light. However, this leads to considerable lateral light scattering in the top layers of the material, which in turn leads to loss of texture contrast. Via simulation of light transport in the polyjet materials, an iterative optimisation can be performed, which arranges the coloured material voxels in the layers immediately beneath the surface in such a way that lateral scattering is minimised (see slides).