Fundamentals of Digital

Jeremy Daalder

Contents

About The Book Itself

Some History

This book has been written by Jeremy Daalder, the director of Image Science.

The Image Science website https://imagescience.com.au also has hundreds of free articles that extend this material, so please explore it - particularly the knowledge section.

This material was used extensively in the now sadly no longer extant Digital Fine Print course at the Royal Melbourne Institute of Technology (RMIT), and has since served as the basis for many guest lectures at other colleges and studios.

In many cases, rather than repeat content, I have used links to other sections of the book or to our website where that content appears, so just follow the links and read the content, then return to where you were (using the 'Back' button on your browser if you're reading online) - and carry on from there.

How To Read The Book

These notes are ideally read in order from start to finish - as each chapter builds directly on the last.

Often along the way I'll gloss over the finer detail of a particular point in order to maintain clarity. So you may find yourself thinking, at times, that I've over-simplified things, or that I'm not telling the full story. You're right, I'm not - and it's quite deliberate. I've taught these concepts to literally thousands of people now, and I've found that it is important to get the basics in place firmly before piling on the nitty-gritty technical detail.

I very much hope you find these notes useful in your quest to achieve the Digital Fine Print. They represent many hours of work, printing, testing and thinking, and while they're definitely flawed, opinionated, and overly wordy, I do think there is some valuable content in them.

All mistakes are of course my fault, and I apologise in advance for any errors, overly strong opinions, or over-generalisations that may be present. I'm passionate about this stuff, and I'm a direct person, and these notes are very much a reflection of that!

About the Author and the Business - Image Science

Image Science - https://imagescience.com.au - was started by Jeremy Daalder in 2002 due to a dearth of really well-informed digital imaging companies in Melbourne at that time. Our first service was film scanning - as it was impossible to source really excellent film scanning in Melbourne at that time - and to this day we still scan film for photographers working with, for example, National Geographic. Our goal with that first scanning service was to be the best in the country and right up there with the best in the world. Soon, we realised we could do the same for colour management and fine art printing. We are now fortunate to count a huge number of Australia's best and most successful artists, photographers and illustrators as our clients.

Through the years we've been in business, education and promotion of a higher end digital community has been a fundamental part of what we do. Jeremy, the owner of Image Science, has been invited to teach professionally all over the world, including to NATO, IBM, American Express, General Electric Finance, Myer, Nokia, and at RMIT, NMIT and many other educational institutions. He has been a judge at the AIPP/VPPY Print Awards, many many camera clubs, and has consulted with most major labs and organisations including the National Library, National Gallery, and many more.

Image Science has three core principles/goals:

1. **Try and provide services of the highest standard and consistency at approachable prices** - through in depth understanding of the digital print process, and the application of ongoing methodological testing to process improvement.
2. **Provide products we use and have an intimate knowledge of**, and back them with unparalleled, personal support.
3. **Be highly open in sharing knowledge** gained along the road to achieving the best possible prints and digital imaging processes

Our website is one of the more popular imaging websites in the world - with hundreds of free articles on digital imaging, attracting more than 30000 unique visitors a month, so please visit our Knowledge Base - you might find it very useful!

Services we can help you with:

- Fine Art Printing
- Premium Quality Film Scanning
- Colour Management (ICC profile creation)
- Consulting on pretty much all aspects of digital imaging and printing, computing and backup systems
- Information on and consulting about exhibitions, fine art print sales, and the use of the web for those purposes.

Products we can supply you with (along with extensive material on actually using them properly!):

- Eizo and BenQ Colour accurate monitors.
- Colour Management Equipment - Colorimeters, Spectrophotometers etc
- Fine Art Inkjet Materials (Epson pigment inks, and papers from Hahnemuhle, Canson, Ilford, Red River, Epson etc)
- Sundry other things like print packaging products

The Book's Technical Process

This book has taken several forms over the years, but it it currently stored as a single collection of Markdown files (more specifically Pandoc Markdown)

A bash script, using the excellent Pandoc document conversion tool converts this markdown material into static HTML files that are pulled into our Craft based website.

The same markdown is also converted into ePub and PDF ebook formats, and from the ePUB book kindlegen is use to make the mobi format.

That's all a bit of a mouthful eh?

Basically it just means there is one set of source files from which both the website and ebooks are all automatically generated.

While this means there are some slight compromises in the formatting, it is the best way I have found to make the material maintainable and accessible to the largest audience.

The web version is, and always will be, freely available.

However, you can purchase the eBook formats (you get all of ePub, mobi and PDF in one download), if you'd prefer to read the material on your eReader or print it out. All these files are free of nasty DRM, but we do hope you won't share them beyond your immediate circle as of course a LOT of time and effort went in to the making of this book.

Introduction

Overview

This is an introductory course to Digital Imaging. Specifically - Digital Imaging Stuff Related to Really Fantastic Printing. We call this area **Digital Fine Print**. At the heart of it is a process known as **Colour Management**.

The goal of this book is to cover all the fundamentals of getting & making good digital image files so that you can get, or make yourself, great prints from those files.

To cover this, we're going to go right back to the basics of digital files and how computers represent colour. There will be some simplifications and rules of thumb along the way. To avoid getting bogged down in too much detail - which is very easy with colour, as it's an extraordinarily complex and tricky little domain to work in - we're going to focus on the basics of really understanding how a digital file is put together and what are the best decisions to make early on in the process so you have maximum potential for flexibility and quality later when you finally get to the point of putting ink on the page.

Your page may actually be paper, or fabric, or a banner, or glass, or metal etc. - it doesn't really matter what the physical material ultimately is - the principles of good file construction and colour management from end to end are always essentially the same.

The Digital Fine Print

This book is written from the perspective of our domain - that is the perspective of trying to achieve the best possible digital fine art prints.

It may be that your goal is humbler than this - indeed it usually is - but whatever your domain, this book will guide you towards achieving higher quality results. So whether you're an illustrator looking to sell prints of your work on Etsy, or a digital fabric designer looking to achieve an edge by designing your own fabric, or a graphic designer that works on food packaging - this book can help, as the principles of creating high quality digital master files, and moving from them to quality prints - are the same no matter what the domain.

To begin, you need to develop a solid understanding of the fundamentals of digital files and colour and from there you will find that understanding and controlling the downstream print processes becomes much easier - whether you are doing the printing yourself or your working with a print operator.

A course like this should be a standard part of all serious courses in the visual arts, but unfortunately this material is rarely taught.

Our Thoughts on the Digital Fine Print

To kick off, though, here's some thought on what we're all about at Image Science.

What is the Digital Fine Print?

Defining the 'digital' part is easy - it means, of course, a print made using modern digital tools (i.e. the computer), rather than a traditional analogue based process (i.e. chemistry based prints). Hybrid techniques are certainly included, for example prints made using a digital negative which is then printed traditionally.

Defining **the fine print**, however, proves quite a challenge.

Perhaps as a starting point, it is easier to define what a fine print is not.

One thing a Fine Art Photographic print is not is whatever grubby bit of lustre paper you get back from your local photo lab on a Thursday afternoon. Another thing a fine print isn't is whatever your untamed, office-level inkjet printer spits out on a piece of paper designed to last ten minutes on a sunny day before it disappears in a puff of smoke.

An actual definition of a fine print may be impossible as it almost certainly means different things to different people. But there are some basic qualities I think a print must possess to be termed a Fine Art Print:

- Clear thought about the goals/meaning/emotion of an image, from capture to print
- (Which results in…) Fantastic content
- Quality materials
- Excellent printing technique beyond what is normally achievable

If a Fine Print has these qualities the result is a print that:

- The eye is willing to look at for an extended period of time
- People can enjoy having up on their wall, for years
- Can command, and deserves, a substantial price
- Is superbly decorative, or has great depth and meaning, and ideally (and rarely) achieves both of these things

My personal view is:

> A Fine Print is your chance to respectfully request from other people the most valuable thing of all – some of their time.

This is a real privilege, not to be taken lightly, and shouldn't be forgotten. You are asking someone to spend some of their short lives in contemplation of your work and it is your responsibility to create something that is worthy of that time. They don't necessarily have to like it, they don't have to agree with it, but it should be worth the time they spend looking at it. Anything less is an insult. Life is simply too short for bad, sloppy, poorly thought out art.

That's how I separate a Fine Print from an everyday print, and it works for me.

Achieving the Fine Print

For me, I like to consider the print as a whole. A print is much more than just the image itself. Obviously content is of paramount importance, and if you stuff that up, no amount of good printing technique will save you. Ansel Adams perhaps said it best when he said a sharp print of a fuzzy concept is still a fuzzy concept. So a Fine Print starts right at the beginning – with a clear concept, and with image creation technique to realise that concept.

> Almost all print failures occur at this starting point - because the starting image is simply not good enough.

However, this book is not about that stage - improving your image creation abilities is of course a life long journey.

This book is about the secondary stages of taking a great image and making from it the best possible print.

Once you have the image content nailed, it is good craft (and also good business practice) to present that content in the best possible way. This is what can lift your work from merely competent workman-like images to Art, to Beauty, to Feeling - to all those wonderful things a great print can achieve. That is what this book is all about. Almost every print in the world can be improved upon in some way through the application of better and/or more appropriate printing technique and selection of materials – with the goal of creating a better, more beautiful, more meaningful image.

Clear thinking on their concepts, and mastery of their craft, are the two main things that make the great ones great.

I am a photographer by background - and it is one of the strange things about the photographers that, amongst all artists, they are the only ones who seem to consider the medium of delivery largely irrelevant. It's either 'lustre' or 'gloss' and that is as far as most photographer's thinking goes. This is phenomenally short sighted – there is a world of materials out there suitable for photographic reproduction to be explored, and each different medium brings something different to the table. And in a world where photography is increasingly seen as an unskilled trade at worst, or a simple, Uncle-Joe-can-do-it craft at best, we as Photographers must explore anything we can use to make our images more expressive, more meaningful, and more beautiful, or we risk extinction. It's as simple as that.

While artists in the more general sense are often more mindful of the medium of their message than photographers, especially with their originals, they are often fundamentally limited in their ability to make a living by their low volume output. Thus for artists, prints are also very very important as a means of multi-streaming their content. And thus, ultimately, they are doing something very similar to photography - taking an image and from it attempting to produce high quality reproduction prints.

While a few savant types can apparently break wind and achieve artistic greatness, the rest of us have only one realistic chance of fundamentally improving the art that we create. Above and beyond all the technical material that is to follow, there is a greater goal here – to get you thinking about HOW you make art - the craft of your art - so that you can make better art!

Contrary to popular myth, creativity does not well up from the inner psyche and burst forth in a fully completed, realised, and saleable form. It takes excellent craft to be able to create important, valuable, and expressive works of art. Good craft (and years and years of practice) is what separated Jimi Hendrix from that guy who won the air guitar competition. The point isn't whether you like Jimi Hendrix, but there's no denying he was exceptionally good at his craft, and he used that craft to get his message across successfully.

So for me, it works to think of the Fine Print as an extended journey from capture or creation, to output, with each stage of that journey requiring careful consideration.

It starts with an output driven approach to capture – In my case I will specifically light things, and alter my photographic technique, where possible, with an eye to achieving the best final print. The initial capture is just the first step on a much longer journey.

(And if I were an illustrator looking to sell reproduction prints, I would originate my works on clean, smooth relatively bright white paper, to ease the scanning and retouching process that would follow the creation of the original (and probably stick to A3 size as that's the typical limit of what is easily & affordably scannable!).

I then use careful editing, with accurate tools for visualising what I am working on and how the final output process will affect the final image. In determining that output process I give careful consideration of the best and most suitable output materials/forms for that specific image. Finally, I use the best possible output techniques available to produce the final print.

At all stages during the fine print process, I give consideration to both the initial impact AND the longevity of the image – that is, I try and keep my eye on both how my image initially effects the observer, and how the perception of that image changes over time. Obviously sacrificing initial impact is a disadvantage, but it is often worthwhile in achieving a print that stands up to extended and repeated viewing.

Much more on this later when we talk about how the eye works and what it wants, versus how the brain works and what it wants!

On the 'Rules' of Fine Printing

There are no rules. Well, actually there a lot of rules, but one rule rules them all. And that is:

> If it looks good, it looks good.

Tautology, yes, but worth stating.

Everything else I talk about is all designed to help you make beautiful prints, prints that make people really want to look at (and potentially buy) your work. But there are no absolutes in this stuff when it comes finally to what you put on to paper (or even what you use as paper!). It's your call.

> But simply saying 'I like it' isn't good enough.

Unless you take pictures for the sole purpose of putting them in your shoe box and occasionally peeking at them and giving yourself a little pat on your back, then you have to acknowledge the goal is to get other people to like/appreciate/understand your work.

So the other rules are really just rules of thumb that should be followed to make your work look good. Because there are some near universally agreed principles as to what works in a print, and what doesn't. And those prints people nearly universally love and admire pretty much all follow a basic set of rules, almost all of the time. To be perfectly clear, I am talking here about the technical aspects of the print, and NOT the actual image content.

Of course don't take everything I or anyone else says as gospel.

Part of what makes you the artist you are will be the decisions you make on the path to achieving a beautiful, expressive print – and sometimes that means, even necessitates, breaking the rules. But be prepared to argue your case – breaking the rules for an aesthetic purpose is great, doing it just because your technique was poor/you were lazy/you ran out of money - not so great.

Beautiful prints take time, technique, and quality materials.

That's just the way it is.

Now, let's get started.

The Beginning - Digital Imaging Fundamentals

The best place to start is the beginning! We discuss just what the Fine Print is, how we can move towards achieving really beautiful prints, and talk about the rules for making a fine print (and when they can be broken!).

We then really kick things off by moving on to the most critical step of Digital Fine Print production - creating an accurate and suitable print production environment, including notes on your work area, screen and video card, and of course, the essential process of monitor calibration.

Groundwork - Fundamentals of Colour, the Human Eye and Digital Images

Let's begin - first we need to talk about colour itself. Because all great prints are, in the end, just artfully arranged colour on a medium.

Colour doesn't exist without someone to experience it, so let's start *right* at the beginning.

The Incredible Human Eye

... and Its Fickle Nature

The human eye is an awesomely sensitive measuring instrument, capable of seeing the universe in ways your camera can only dream of.

> Printing is a con - just a great big trick.

It's all about using a pathetically limited medium, a single, generally flat, piece of paper, to fool the human eye into thinking that it is seeing something vaguely related to rich and wonderful three dimensional reality.

Our challenge is to use the very limited tools we have available – contrast, colour - that sort of thing, to create a miniature world, capable of evoking emotion/stimulating thought/firing of the visual pleasure synapses.

What the world presents to the eye, and what the eye can take in, are so far beyond anything realisable on paper that it is amazing we can get away with this trick at all.

Lets look at contrast as an example. Imagine yourself on a beautiful tropical beach. Your eye, quite comfortably, will see detail in fluffy white clouds, and in the deepest shadows of the palm trees. I don't know the exact figure, but it is usually quoted as being over 100,000:1 – that means you can comfortably look at a scene and see detail in the brightest object, where the brightest object is at least one hundred thousand times brighter than the dimmest object you can still see the detail of! Your camera can't come close to capturing anything like that brightness range. Prints, even made with the very best printing systems, can't achieve anything like that range of contrast – not even close - indeed printing rarely achieves over 200:1 contrast.

Looking at colour, we've got the same problem. The human eye can see a very broad range of colours indeed. Although, by animal standards, we're pretty feeble. You might think a crow is a black bird but put it under a black (UV) light and you'll see that it isn't – it's actually very colourful (to other birds with eyes sensitive to UV light, anyway). This is, perhaps surprisingly, quite relevant, because it makes us very conscious that when dealing with colour, we only care about what is commonly referred to as the visible part of the electromagnetic spectrum – that is, the only colour that really counts is the colour we humans can see.

From our high school science lessons, we're aware (hopefully) that off one end of the spectrum is Ultra-Violet and off the other is Infra-Red, but we all know that the ROYGBIV rainbow is all that really counts to us (ROYGBIV = Red Orange Yellow Green Blue Indigo Violet). Already this should be giving you the feeling that colour is all about the *observer* – colour doesn't exist unless its seen (if a tree falls in the forest and nobody hears it, are its leaves still green?).

To get back to the topic, the human eye can see a great deal of colour. And of course, just like with contrast, we can't print anything like that amount of colour. We work within a pathetically small subset of colours the eye can see. You'll see precisely how small soon.

What The Eye Wants

The eye, being such a sophisticated sensor, capable of such amazing things, craves these very things. Like any addict, it craves a bigger hit. It's just like when you're listening to music, and over the evening, you turn it up, then up some more, then eventually, around 5 o'clock in the morning, the neighbours are thumping on your walls and you've got yourselves a little problem with the landlord.

The rods and cones in your eye, those things that actually detect light and send the signal to your brain, get saturated with signal. So as you look at an image for an extended period, and Photoshop it from mere snapshot to work of Art, be aware that your eye is going to start lying to you. It's going to tell you that your image needs a little more contrast. Maybe a little more saturation. And then a bit later, hmmm, maybe it needs a bit more punch! I know, I'll boost the contrast some more. And add a little colour.

Pretty soon, your beautiful fine art portrait looks like a Hunter S Thompson novel.

The point here is that the eye needs breaks – when you're working on images, what you see is not the same as what someone looking at your image for the first time sees. It's not the same as what you will see if you go away, have a coffee, and come back to the image. So take regular breaks and give your eyes a chance to rest and replenish.

The Adaptability Issue

The human eye is an automatic adaptor - this is great for us, as it means when we move from different contexts (say indoors under warm interior lights to outdoors under cool blue sky) - we don't experience massive and jarring colour shifts as the colour of light changes - indeed, we basically don't perceive the colour of light changing at all unless its a very significant shift. However, devices like cameras do not automatically adapt in the same way. A camera, without being supplied a whitepoint (or if its auto whitepoint calculation fails) - will perceive a distinct orange cast to indoor light, and a distinct blue cast to outdoor light.

This automatic adaption is very handy for our day to day vision, but it fundamentally makes the human eye a very poor tool for *objective* colour measurement because our eyes instantly and subconsciously adapt to all sort of changes.

Colour - Light, Objects, Sensor

Colour seems obvious, really. The Rainbow - ROYGBIV - right?

Well, that's only one part of the story. Colour is a remarkably fluid thing and pinning it down completely is next to impossible. Long, dull books (with a LOT of unpleasant maths in them) have been written on colour theory. We're just going to cover the basics, so you can see why all this stuff is so hard, and so you can understand the tools you're going to be using.

Colour always involves three things – a light source, an object to reflect that light, and a sensor to receive and interpret that light. Think the sun, the ocean, and your eye. Specifically, the visible wavelengths of light are bounced off the object and stimulate the rods and cones in our retina to send signals to your brain. Obviously, your perception of that colour will change if you vary any one of those three things – that is, if you use the moon instead of the sun, the sandy beach instead of the ocean, or you're a little red/green colourblind.

So pinning down colour can be hard – no object *IS* a particular colour, it just *appears* to be that colour under certain conditions. There are no absolutes in colour – unless you control all three of those variables, you can never absolutely say what the colour of something will be and how it will be perceived by others. This is a Big Problem.

So now we're going to talk about the basics of the three things involved in colour:

- Light
- Objects
- Sensors

Light

Light can be thought of as a wave, and we can see wavelengths that fall between 380 and 700 nm. This is the **visible spectrum**, and what we commonly refer to as, simply, light. From 380 to 700nm, we have named the colours, and it looks kind of like this:

```
(UV) 380 nm - Violet Indigo Blue Green Yellow Orange Red - 700 nm (IR)
```

(While IR (Infra Red) and UV (Ultra Violet) are not generally visible to the eye, they both do have some impact on fine printing, so don't forget about them completely. IR is most relevant at the image capture end as digital camera sensors are very sensitive to IR, hence the IR filter in most modern cameras. UV is most relevant in printing because they show up in fluorescent optical brighteners).

The light we see is typically made up of many wavelengths, that is, our eyes simultaneously see and mix multiple wavelengths bouncing off an object, and form a perception of a particular colour from this mixing process.

White is what we get when light is made up of equal amounts of all the wavelengths of the visible spectrum. But think about a white piece of paper for a moment. The eye is pretty amazing, almost no matter what the light source, whether it's the sun or tungsten lights or glow worms, our eyes will pretty much always look at a piece of paper under any light source and go, in a very simple and basic human way - 'that's basically white'. This is called **colour constancy**.

Our eye is great at ignoring, or coping with, different light sources. It's one of the things that makes colour really hard. Our eyes are very adaptable, but digital systems are not. That's why, for example, we have to define a white point when developing raw images from a camera.

We can see there's no true definition of white, no such thing as 'absolute white'. So we use a term known as **colour temperature** to talk about the actual colour of white light that we mean. Based on the Kelvin System, and it's somewhat counter intuitive - *higher* degrees of Kelvin actually imply a *bluer* whitepoint. To remember this I think of Bunsen burners back in high school chemistry - as you gave the fire more oxygen it burned more powerfully and hotter, and the flame changed colour from yellow to blue. Here are some common white points encountered in digital imaging:

- 9000K – very blue white
- 6500K – so called 'daylight' on a pretty much cloud free summer's day
- 5000K – many fine art papers are down around here, a moderately warm white
- 3000K – traditional tungsten light bulbs emit a quite warm white
- 2000K - candle light - positively red/yellow at this point

Objects

Objects reflect light. Which particular mix of wavelengths they reflect determine the colour of the object. They can of course, only reflect those wavelengths that are falling on them in the first place. (That's why fluorescent tubes can make things look funny, because they simply don't output certain wavelengths of light, so objects can't reflect them, so then those objects look oddly coloured under fluorescent lights).

That's about all you really need to know about objects initially, except it is worth considering one other thing in more detail ...

Fluorescence is a property of objects whereby they change one type of wavelength into another – and it is most noticeable in really, really white things. Like your teeth, if you brush them. You know what teeth look like in a night club with those UV 'black' lights? Brighter than white. Well, the same stuff they use in toothpaste is also used in detergents for clothes, and more importantly, in papers - to make them appear really white. They're called **Optical Brighteners** and some people get pretty upset about them. Probably unnecessarily, but the jury is still out on that to a certain degree. We'll talk about optical brighteners more later.

Fluorescence is often important in colour management and digital printing because it can make things unpredictable, and that's a big problem in managing colour. So we try and avoid it in both lights and papers, if we can.

Sensors (e.g eyes), and RGB vs. CMYK

Sensors – your eye is one, your camera another, and scanners are yet another. In some ways, they're all very similar, in other ways, very different indeed. As discussed, the eye is very flexible in its response, whereas cameras and scanners are very much fixed in theirs.

The eye, your camera, and scanners are all fundamentally Red Green Blue sensors. This is really important – it's the basis of all colour reproduction. Our eyes, and these machines, detect and mix these three colours to see all the rest.

On paper, however, we subtract reflectivity from the paper using the opposites of these colours (Cyan, Magenta, Yellow), to produce all the visible colours.

So, in either colour system, we can use just three primary colours to make all the others, which is pretty amazing. How that works is pretty complicated, but it will do for now just to know that it does, generally, work.

RGB is the language of light.

Red, Green and Blue are known as the **additive primaries** – because we add them together to get all the colours.

CMY is the language of ink.

Cyan Magenta and Yellow are known as the **subtractive primaries**, because we use these colours to subtract reflectivity from white (i.e. paper) to get all the colour.

Some of you are bound to be thinking, right about now (if you're still awake) it is CMY **K** not just CMY. So what is that **K** all about? The K stands for black, and it is added into the ink mix for two main reasons – one, to get better, purer blacks, and two, because it is cheap. And for most printers, cheap works. When we talk about CMYK later, we'll go into more details.

It's worth noting that there's a pretty fundamental difference in the nature of light based systems and ink based systems.

In light systems, you get more saturation by adding more light together, so your saturated colours tend to be at the higher end of the brightness scale.

With ink systems, you make more saturation by adding more ink (and therefore subtracting from the paper's reflectivity) - so you tend to get your more saturated colours at the lower end of the brightness scale.

The two sets of colours (RGB and CMY) are complimentary (or opposites) – here's the digital colour star:

As you can see, The digital colour star shows you the way these things work – if you're trying to get rid of blue, you need to add yellow (or reduce both cyan and magenta). There's no such thing as a 'yellowy blue'!

It's really worth drawing yourself one of these (or copy this and print it out) - and keeping it handy for a while. It really helps when you're trying to make really subtle changes to colour to have a good grip on how the colours all relate with each other.

Some Concepts & Basic Terminology

We're just going to introduce a bit of colour related language now.

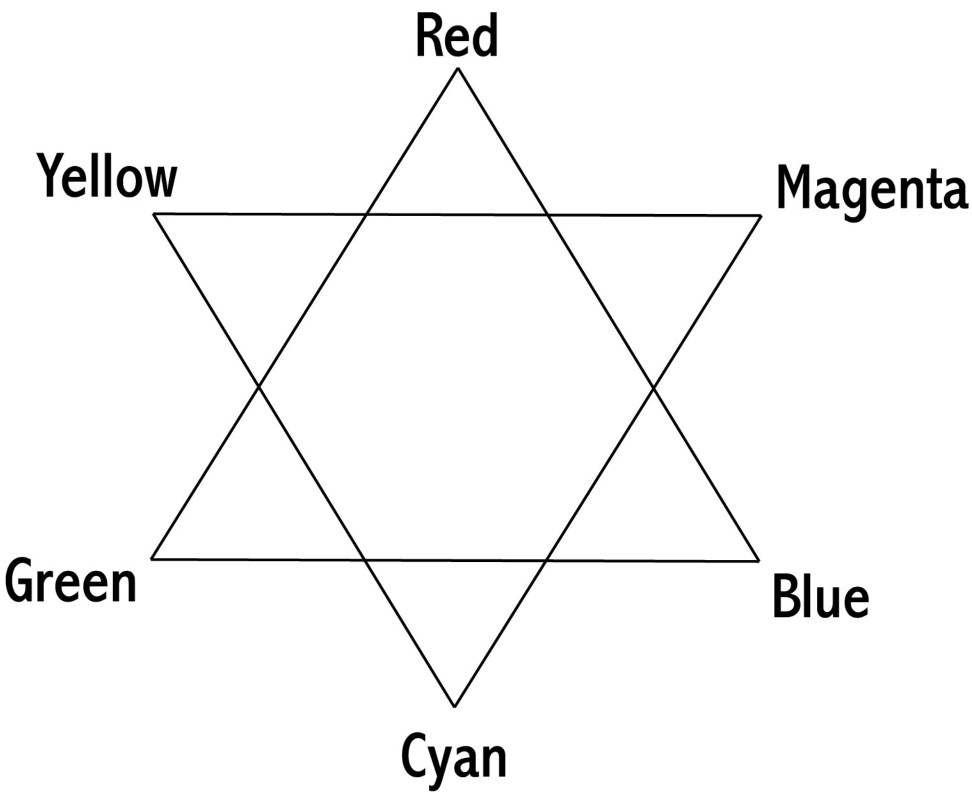

The Digital Colour Star

Standard Lights (D50/D65) and The Standard Observer

The upshot of all the above is you need to start thinking of colour as a fluid thing, dependent very much on three separate parts (Light, Object, Sensor). That is, the properties of the original light source, the object you're looking at, and what you're using to look at it (i.e. the sensor involved).

Given there are so many variables in colour, it's hard to talk sensibly about colour without defining a few things. So we have to define absolutes, or reference points, so that we can get to grips with all this.

Standard Light Sources

The most important of these definitions are the **reference light sources**. There is no sensible way to talk about the colour of something under *any* possible light source. So first we define some basic, sensible light sources. You may have heard of these before – in particular you may have come across D50 and D65.

- **D50** is a light source with a whitepoint of 5000 kelvin.
- **D65** is a light source with a whitepoint of 6500 kelvin.

The whitepoint is the most important thing, but of course they both have a bit more to them than that - notably they each have defined spectral output curves which define exactly how much of each given wavelength of light they emit.

Once you have a precisely defined light source, you have nailed one of the three variables to the wall.

In theory anyway. In practice, you will almost never come across true D50 lighting. Fortunately, there are a lot of 5000K balanced lights out there that come acceptably close. You can buy lights that come pretty close to the D50 standard and this is an excellent thing to do if you're going to be evaluating colour and prints with any seriousness.

The Standard Observer

The next thing we nail to the wall is the range of colour the human eye can see – the only colour that is really relevant. This has been measured for us, and is called the 2° Standard Observer. This results in a particular colour space called CIELAB, which is a description of colour as the human eye sees it – roughly. More on this later.

Objects

We've now standardised two of the three variables – the light source and the observer. All that is left now are the objects, i.e. prints.

Well, the objects we create are ink on paper, and how skilfully we put it there is what determines how good a printer we are. Of course, unless we fancy grinding our own pigments (great video here!), we have to use the inks we can easily buy to produce our prints, so the trick is to really understand the behaviour of those inks and the printers using them, to create prints which look good - under standard light sources to standard observers!

In the real world, the reality is that we don't normally have precisely the standard light source OR the standard observer available to us, so in reality we'll be working within acceptable tolerances, but it all works pretty well in the end. But you have to hang your hat somewhere, and for colour, the standards above are the commonly used hat racks.

Gamut

Gamut is just a word that means the total range of colours of something. The gamut of the eye are all the colours an eye can see. The gamut of a printer is all the colours that printer can reproduce (which is determined by the inks the printer is using, the paper/coating the inks are being put on, and the driver technology that is actually determining how that ink is being put on the paper).

Metamerism

A true definition of metamerism is quite complex, but metamerism in practice is the phenomenon whereby something you print changes colour unpredictably under different light sources.

This is most relevant to black and white printing, and is still a very common problem. Basically, say you can achieve nice, neutral greys out of an inkjet printer, or so you think. Take your nice neutral great print outside and it suddenly changes colour, taking on a unpleasant cyan or magenta cast. This is metamerism in practice, and in many cases, it is an unsolvable problem without pretty fundamentally changing your inks.

Optical Brighteners

Optical Brightening Agents (often referred to as OBs, OBAs, or Fluorescent Whitening Agents - FWAs) are chemicals used in many contexts - for example, they are put into detergents to achieve brighter whites and stronger colours. Toothpaste is another common place they are found.

OBs convert light through fluorescence, right at the far blue edge of the visible spectrum (i.e. UV light) into visible light, thus making whites appear cooler and brighter, and colours more saturated. This is done for marketing reasons (i.e. look at our lovely bright whites) and also to allow the manufacturer some tolerance in the paper making process (e.g. if the paper base of a particular batch is a little more yellow than spec., OBs can be used to cheaply bring the paper to a consistent state).

The jury is still out for many people on whether OBs are of real practical significance when considering the longevity of a print. Some people believe OBs will simply fade away over time, leaving the paper no warmer than it would have been anyway, and no warmer than papers originally sold without them. Others believe the chemicals are likely to have an effect on the long term survival of a paper (either causing greater degradation over time, or leaving the paper more yellow than it would have been were the OBs never present), and thus should be avoided at all costs. Some people just don't like warmer papers and so are willing to accept OBs no matter what the cost!

The reality is that many prints made by great printers of days past, people like Ansel Adams, were made on papers with OBs in them. Undoubtedly they have a different visual appearance today to prints made on papers that never had OBs in them, but it turns out that many of these prints are still the preferred prints today! The character they have taken on (through the OBs fading) is actually still appreciated! Thus, as usual, the answer is not simple.

However, in terms of colour management, OBs are a definite problem. Firstly, they tend to promote metamerism, making colour casts on black and white prints under different light sources much worse, and they also mean that what you print now will not be what is on the page in a few years (as OBs fade much more quickly than paper coatings without them)…thus all your good work in colour management is negated after just a few years if you use papers with OBs in them. It may be the fading is very minor, so this may not be a real practical problem, of course, but some very bright white papers have a LOT of OBs in them (e.g. a lot of the Epson papers) so these should definitely be avoided.

If in doubt, and particularly in the gallery market, it is probably best to avoid them, or at least papers that have lots of them in them.

File Types

Let's talk about the basics of digital image files. Understanding these is at the core of getting started with producing high quality files that give us plenty of flexibility later in the process.

Continuous Tone vs. Spot Colour

Just some terminology to describe the two types of imagery:

1. **Continuous tone** images include a range of colour densities, as opposed to line art which

includes only solid colours (e.g. black text in a book). Photos are the most obvious example.

2. **Spot Colour / Line Art** images are just those having only a few chosen 'spot' (or solid) colours in them.

File Formats & Compression

Digital image files can have many formats. The most fundamental decision with file formats is whether to use a lossy compressed or uncompressed/lossless file. The difference is the same as with MP3s and CDs - one file has been subjected to a mathematical technique that dramatically reduces file size whilst only typically marginally reducing the reproduction quality.

In digital imaging, we should always try to use uncompressed file formats for our work. Modern computers are cheap and capable of manipulating huge files with minimal delay. So always use file formats that maintain the full integrity of your file. Typically, this means using either your applications native format (.PSD, for example) - or a standardised uncompressed file format like TIFF. If in doubt, a TIFF is pretty much always a safe bet. The observant among you will now say 'but wait - when I save a TIFF it asks me how I want to compress the image. Indeed it does, but it's using a lossless (non damaging) form of compression, so you can essentially ignore this.

The only time to use lossy compressed files (e.g. JPG) - is when we are sending them somewhere - out to the internet, via email etc. Even then, modern file sending services often make it easy to send the large, original files, so resist using lossy compressed files wherever you can.

Vector vs Bitmap

There are fundamentally two types of digital image files:

1. **Bitmap Images** - These describe the image in question using dots (=pixels) (like a digital camera produces) - usually these are generated by cameras or scanners, or made using software like Photoshop. They are often also called 'rendered' images.

2. **Vector Images** - These describe the image in terms of lines, shapes, and splines (which may be packaged up as fonts etc.). These usually made in Desktop Publishing apps like InDesign or Illustration packages like Illustrator.

Choosing vector or bitmap is a fundamental decision.

Vector files are best, really - if you can use them they have many advantages - much smaller file sizes, and you can scale them to *any* size without experience any quality loss.

But of course, we can't always use vectors. Cameras and scanners always produce bitmaps, and if we're working with that sort of imagery, at the least those elements will be represented as bitmaps. Basically, anything with a great deal of fine detail is best represented by a high resolution bitmap. (More on resolution in a second).

Formats like PDF and .INDD are actually *container* formats that can package up elements - some of which may be vector, some bitmap - in many cases mixing both types within the one package.

The idea is to use the most appropriate type wherever you can and avoid mixing where you can. That is - avoid dropping fonts over a low resolution photo for example - when you scale this later, your font (now rendered as pixels) - will go all blocky.

> Ultimately all vector images are 'rendered' into bitmap images when finally printed - in the end, **all images end up being dots.**

Control over precisely how this rendering is done can be important and it's why, for example, at Image Science we only accept fully rendered files - no vectors allowed! This way we can avoid all common rendering issues like missing fonts or variable transparency behaviours in different app versions - the rendering has already been done by the image maker, who is of course the best person to control that process.

But rendering in many typical workflows happens just right before print, in the RIP (Raster Image Processor, which is either a dedicated RIP program or a 'software RIP' as found in Photoshop, Illustrator etc). This common process mistake is a classic cause of many print issues.

Pixels and Resolution

Knowing how a digital image actually works is really useful, because it helps you avoid common mistakes.

Knowing the basic of pixels, bits, resolution, and colour models, help you solve real world problems like – how big can I print a file and get away with it? How should I sharpen a print? How do I move between CMYK and RGB most effectively?

So we're just going to wiz through the basics of digital images, ending up in the strange land of colour spaces, where we're going to spend some real time.

Pixels - Greyscale

A pixel is a single dot of a single colour.

A bitmap digital image is made up of a grid of pixels.

To start simply, in a grey scale file, the value of that pixel indicates how bright it is, from:

`0 - Black`

to

`255 - White`

People think of these things as a theoretical black and white, but it's better to start thinking of printing right from the start. Thus you might want to start thinking of these values as indicating D-Max (maximum density), or maximum ink black, and D-Min or Paper White (the colour of your printed medium with absolutely no ink on it).

Dynamic Range and Clipping

So, this is the tonal range we actually have to work with – we obviously can't print anything darker than maximum ink black (determined by the machine, media and inkset we use), and we can't print anything lighter than paper white (determined by the type of media we choose).

But we are free to roam within those boundaries, completely as we see fit. Indeed we need to roam within these boundaries, and right up to the fence - or we will get dull, lifeless prints lacking the full tonal range. But we really need to be careful stepping out of these boundaries. Because over the fence bad things live.

While I've asked you to start thinking of the boundaries as ink black and paper white, it's important to remember that this isn't necessarily the case – it's only the case if our output medium is a print. If your final output medium is a monitor, then 0 and 255 are simply the blackest black and brightest white the screen can reproduce. The distance from black to white in any system is referred to as the **contrast range** or **dynamic range** of the system – i.e. the range of brightness levels a system can reproduce.

If we don't get close to these boundaries, we won't be taking advantage of the previously mentioned pathetically small amount of contrast (dynamic range) available to us, and that is bad because our prints will be flat and muddy looking, but if we step outside of those boundaries (called clipping), we're going to be in very big trouble because we're going to run out of detail and start committing some of the cardinal sins of printing - burned out highlights or shadows without detail.

When we're manipulating tonality in our images, be it with levels, curves, or any other tool, it's very important that we make sure we are using the complete dynamic range we can (in general, occasionally you may deliberately want to keep your images within a narrow contrast range for effect), and the Threshold View in Photoshop is a really handy tool for this. You've probably seen this before, but just in case you haven't, here's the scientific way to determine the black and white point in your photograph, thus maximising your images use of the available dynamic range.

The image that follows is, at the moment, flat and lifeless, and the histogram for the image shows us why – the tonality is compressed and we're not taking full advantage of the available dynamic range available to us. This can be seen because there are no pixels at either end of the tonal scale, that is, no pixels have values below about level 17 and above level 237. This is a typical example of a 'raw' scan you would get back from a high quality scanning lab (like Image Science!) – it is deliberately scanned in an unclipped form, so that maximum detail is obtained from the sensor.

A typical lab scan will try and make the image 'more appealing' on screen, because it will be done by a scanning operator with very little training – this is what you will get back from a typical lab:

You can see from the histogram that the image has been clipped – and the resultant image is way too high in contrast – printable detail in both the shadows and the highlights has been lost (e.g. the hair is now a black blob on the girl's head).

So, starting with the first image which has all the detail from the original capture, rather than the second which is irreparably damaged, we now use the threshold view in levels to scientifically determine the black and white points.

Open levels. Now, while holding down the ALT key (option on a Mac), drag the black slider to the right – the image will go completely white at first, and then as you drag the slider to the right you will see some pixels change colour. Those pixels that are changing colour are clipping in that channel (i.e. they will be set to (0,something, something). When the pixels go completely white,

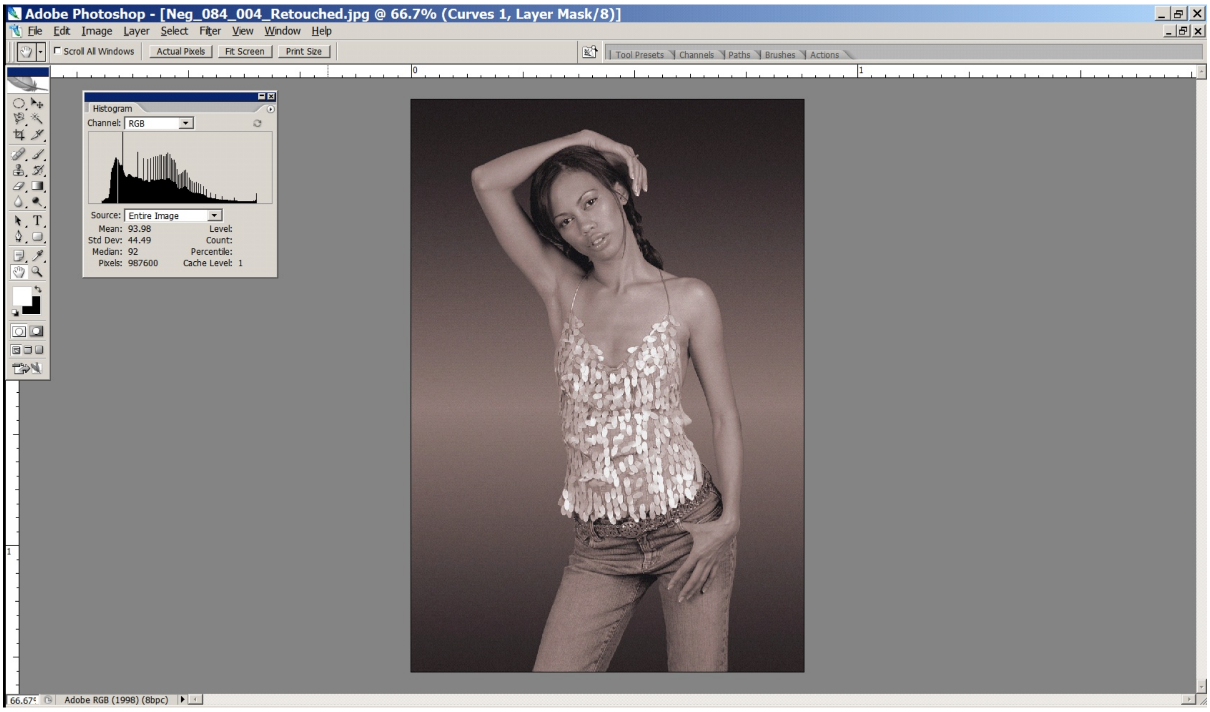

Good - Raw Scan

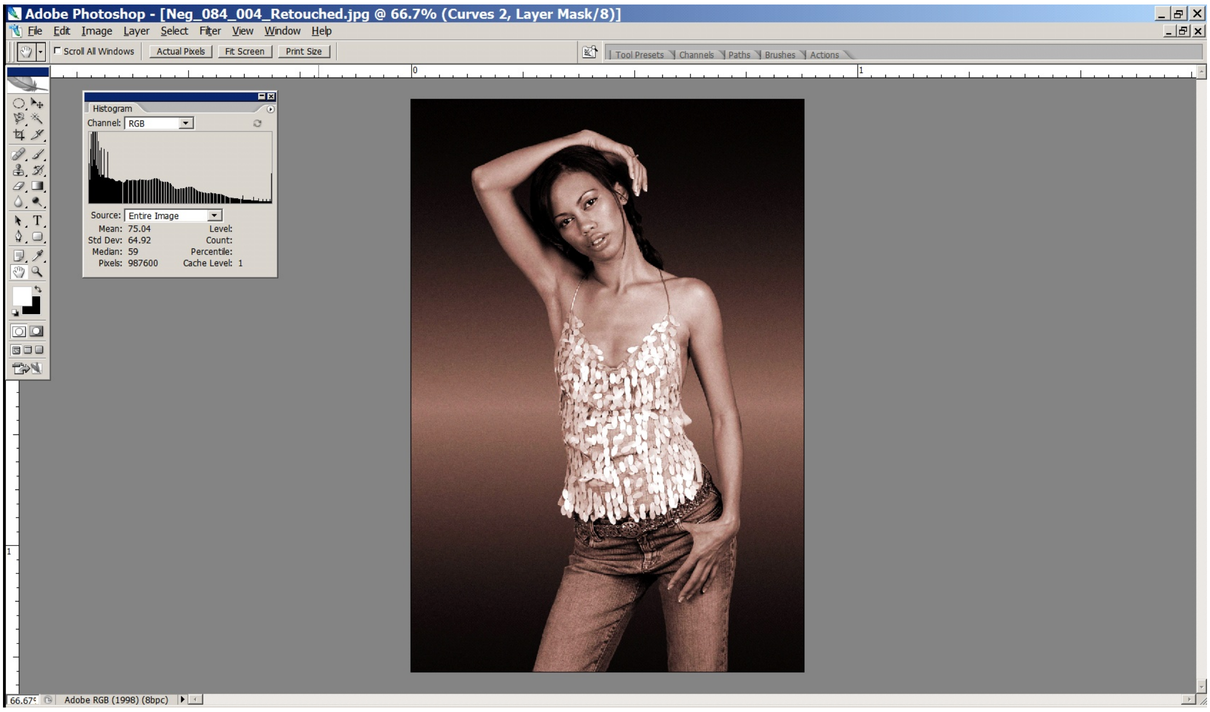

Bad - Clipped Scan

they are clipping in all three channels (ie are set permanently to

```
0,0,0
```

– maximum black).

Exactly where you want to set the black point is up to you – if you are trying to achieve maximum black in areas (e.g. an image with a totally black background) then this process can be used to set those pixels to absolute maximum black, so that when printed there will be no chance of weak, muddy blacks. However for most images, we generally only want the very darkest of pixels to print as max black, and to have detail in the image wherever possible, so we tend to set the black point to the value just on where the first pixels change colour:

Here I am allowing the very darkest areas of her to go pure black.

Next, we do the same for the white point – hold down ALT (option) while dragging the white point slider to see where the threshold view lies – remember, any pixels that turn white will be permanently set to 255 (ie paper white)

Here, I am deliberately allowing a few of the specular highlights off the girl's top to be set to pure white. In images without specular highlights, I would set the white point to the level just before ANY pixels turn white. These highlights are deliberately specular (i.e. have no detail) and small – otherwise we would want their value to remain below 255 (i.e. paper white).

The result of this process is an image with appropriate contrast - making full use of the dynamic range available to us, but still with detail across the entire file (the histogram proves the details are there!). The histogram runs all the way from 0 to 255. The little lump at the far right of the histogram are the specular highlights we have set to 255.

Rule Number One

So, we've looked at how to stretch the tonality of an image across the full tonal range to take advantage of what little contrast range we've got available to us. The reason we use the threshold view is because it explicitly tells us which pixels we're going to clip. Clipping is when we cause values in our file to rise above or fall below the threshold – that is, they are (permanently) set to 0 or 255. 0, as we know, is maximum black, and it isn't necessarily a big problem, but it does mean we've lost any printable detail in those pixels. 255 is paper white, and this means, in reality, that we'll be sending a signal to the printer to not lay down any ink at all for those pixels. This is bad, and it leads us directly to the first rule of printing…

Here's the First Rule (and remember, every rule has an exception, and these rules are more like guidelines, but this first one is pretty much set in stone!):

Rule Number One

> Never, ever, print paper white in your print.

It's ugly, it stands out like a sore thumb and anyone who tells you different is blind or wrong. Printing paper white in your print area is the number one thing that will make your prints look dodgy and

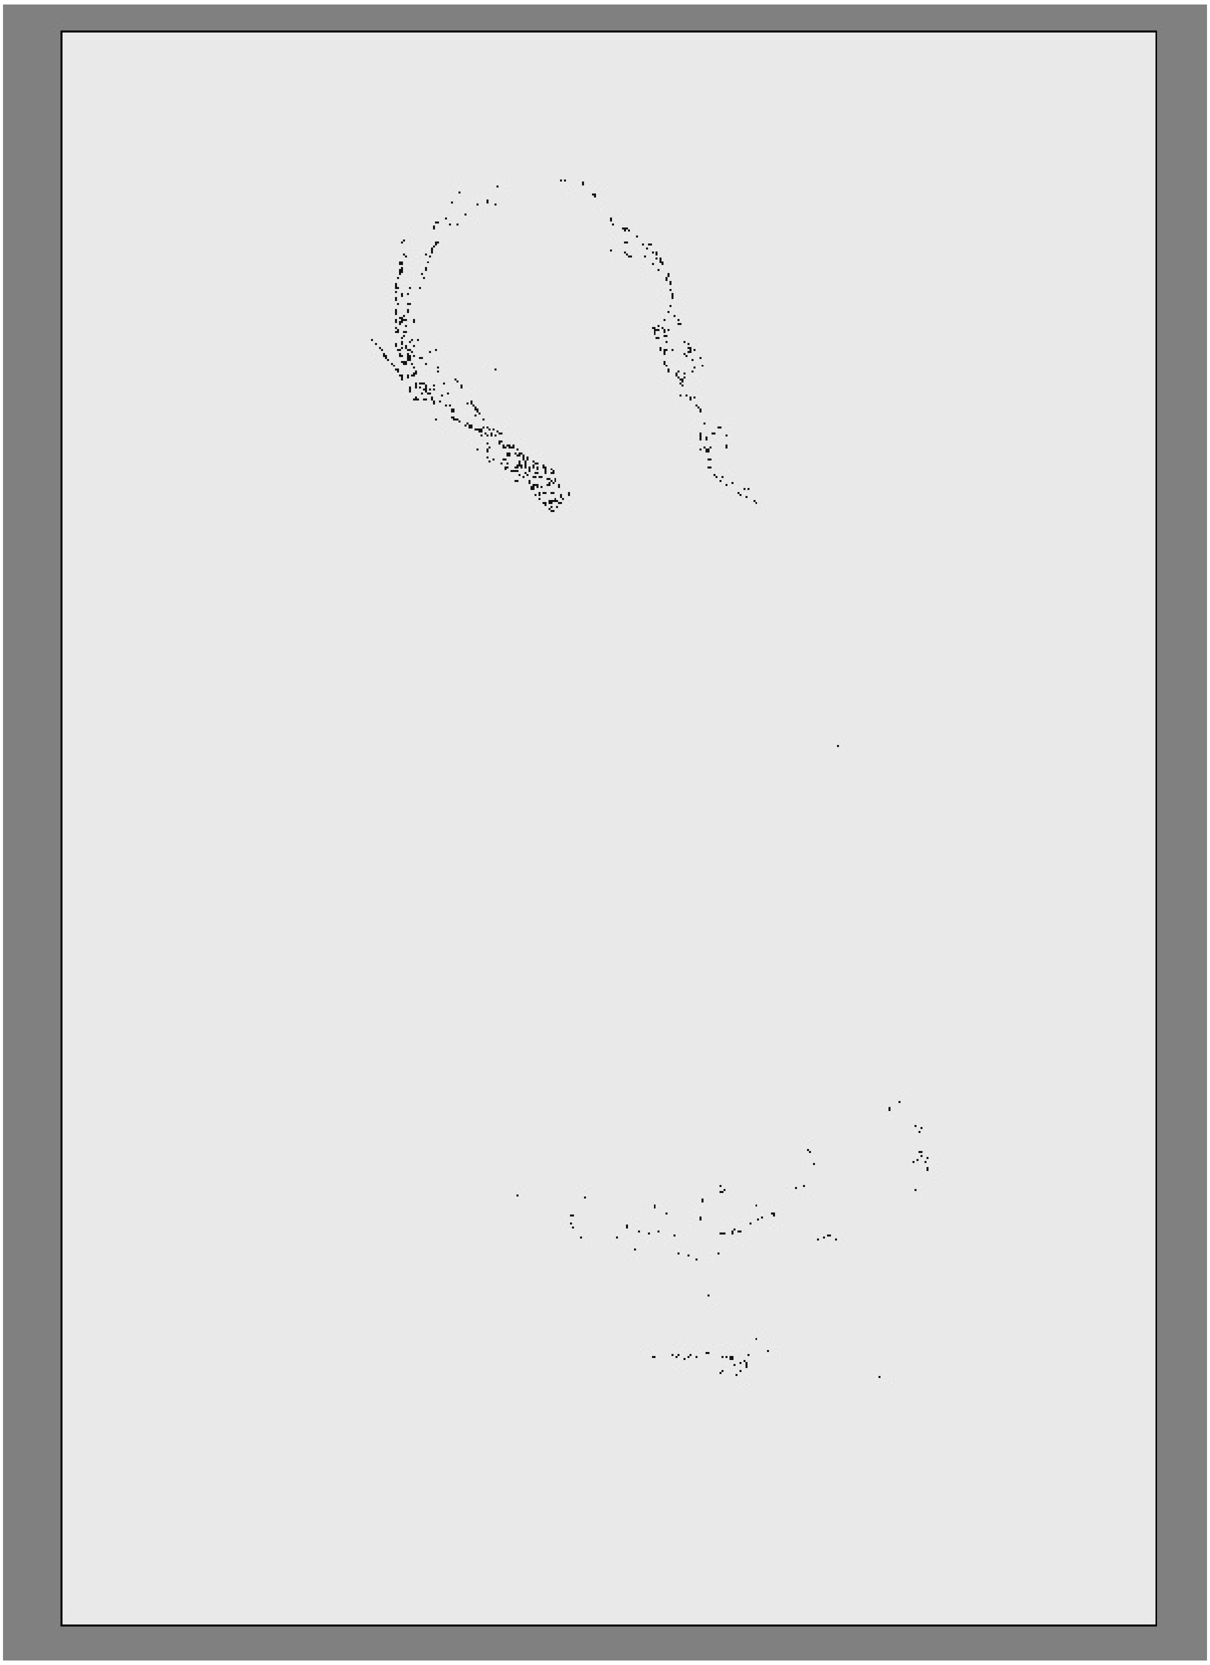

Setting the Black Point

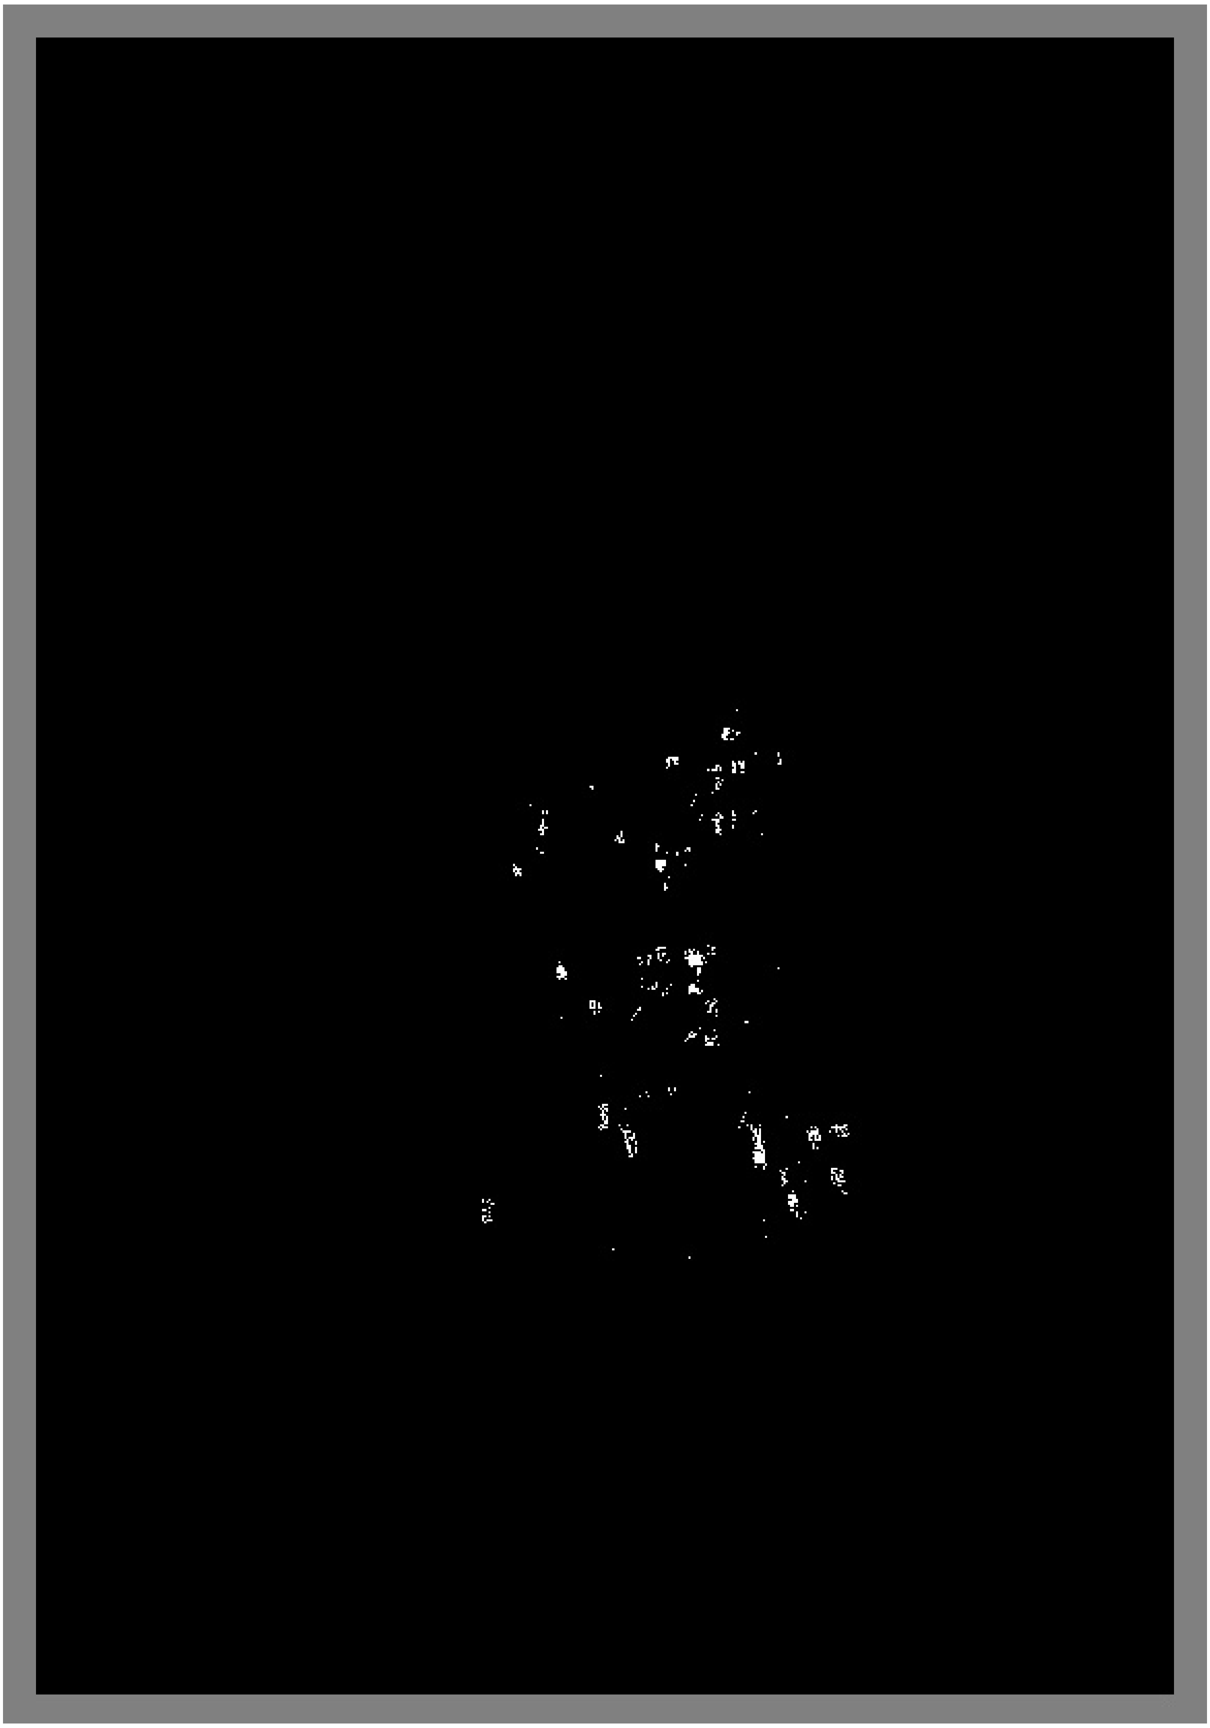

Setting the White Point

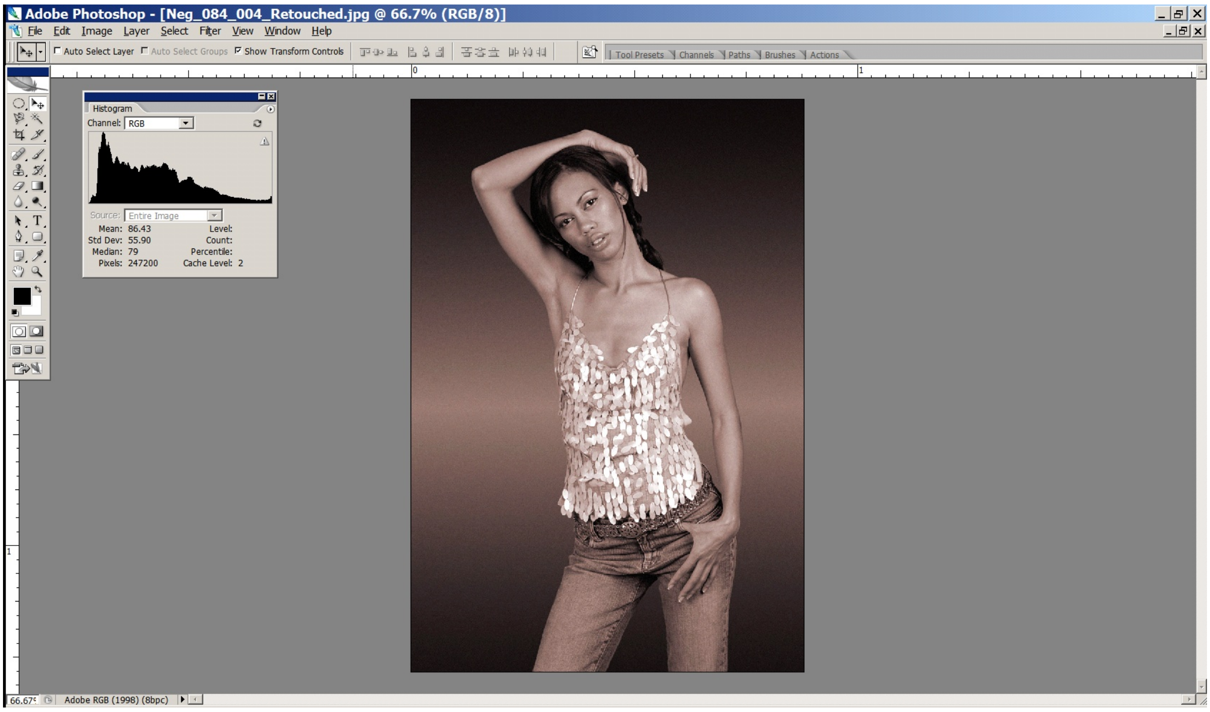

The Properly Prepared Image

amateur. Printing paper white in your print joins your image with the world outside your print. It defeats the fundamental goal of a fine print – that of capturing and holding the eyes attention.

Exceptions to Rule Number One

Specular highlights (tiny reflections, like those on the edges of metal), are ok to print as paper white. Even then, only if they are *very* small (like 0.5mm or less) and there aren't thousands of them. Also, when you're doing cut outs of objects (e.g. in a brochure) - it's ok to pop those on pure paper white, of course. Rule number one is about highlights within the boundaries of your image.

Addendum to Rule Number One – About Blacks

Printing full ink black, on the other hand, is often very necessary. Without it your blacks will look weak and grey. But printing too much of it is fatiguing (and dull) to the eye. Unless you're doing a commercial type shot on a completely black background, be wary of large, detail-less areas in a print. Black paper doesn't have much of a story to tell.

Resolution

We've talked about individual pixels, but a single pixel on its own is no good to us. We need millions of them all squished together to make a decent image.

A complete grey-scale digital image is thus made up of a grid of pixels:

The resolution of an image is how many pixels it has – X by Y pixels.

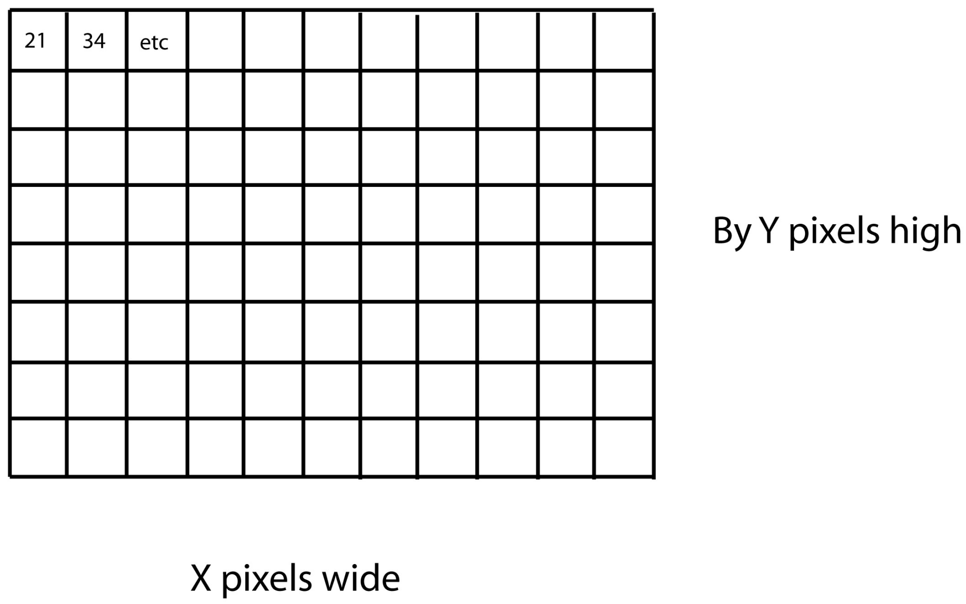

Pixel Grid

Some examples: a typical digital camera file has 3500 by 2200 pixels. A typical 35mm film scan has 5400 by 3600 pixels. In Photoshop, I can create an image with 12000 by 12000 pixels if I like. It will take a lot of memory, but it will also mean I have lots of resolution to work with later. In general, more resolution is always better - it gives us more options with print size later.

PPI vs DPI & Screens/Dot Patterns

PPI - Pixels Per Inch

DPI - Dots Per Inch

Unfortunately, in the world of imaging, PPI and DPI are used almost interchangeably. They are two quite different things and the confusion of the two leads to more confusion than almost any other single thing in digital imaging from the file production end of things.

The first things to understand is that in digital imaging, the only thing that really counts about a file is how many pixels are in a file. Terms like megapixels, DPI and file sizes in megabytes only confuse the issue. In the end, all digital images are simply X pixels by Y pixels big (by Z bits of colour data but we can ignore that for now).

This is the only absolute measure of the quantity of information in a file (nb it has nothing to do with the quality of information in a file!)

In almost all cases, unless you are talking about physically how your printer is laying ink down on the paper, you are actually dealing with PPI - pixels per inch.

Dots Per Inch is an old printing term and has almost no place in modern digital imaging at the creation end.

DPI is a measure of how many tiny, tiny droplets of ink a printer is laying down in its dither pattern to form one inch of a print. Most Epson inkjets, for example, operate in Photo mode at 1440 DPI (can be as low as 720 DPI and as high as 5760 DPI). You tell the printer which mode to print in in the driver - this is usually camouflaged by the use of modes 'Photo', 'Best Photo', 'Photo RPM' or 'Fine', 'SuperFine'. In the bigger printers you can usually choose the DPI directly as, say, 2880 DPI. Incidentally, almost all printers operate best for general photographic usage at around 1440 DPI. The higher DPI modes like 'Photo RPM' are useful if you're printing really high key (i.e. all light toned) shots, but next to useless for general printing. In fact, they are worse than useless, they are positively damaging - way too much ink is laid down on the paper, resulting in seriously impaired shadow detail. And of course the more ink you use, the more ink you pay for (an Epson R800 running in 'Best Photo' mode is less than half as expensive to run as one left in 'Photo RPM' mode for instance, and in 99.9% of photos, you'll get worse quality from the 'Photo RPM" mode!).

Pixels Per Inch is the key term. It is a description of the logical number of pixels from your original image (X pixels by Y pixels, remember) that will be used to tell the printer to print one inch on paper. Assuming a sharp original photograph with good technique (see resolution discussion below), the higher the PPI, the better the quality print you can achieve - this is testable as true even well beyond most claims of 360 PPI being the most you need ... 600 PPI images can easily be seen to be much sharper again if this data is available at good quality from the original file).

PPI is a logical term - changing the PPI of a particular file does not in any way affect the file itself - it is simply a decision about how many pixels of the available pixels you will use to print an inch on page. You can choose any number you like - from 1 to infinity. The de facto standard for high quality, photographic printed images is 300 PPI - that is, for each inch of the printed image, there must be 300 source pixels to use.

This is why the 'resample' check box, in the Image Size dialogue, is the single most important (and dangerous) control in Photoshop! When you resample an image, you are actually changing the number of pixels in your image (i.e. changing the value of X and Y) - adding some or throwing them away. You should only do this if you are making an explicit and informed decision to do so, because no single other thing will affect the quality of information available to you from your file as this!

An example will make all this clear:

Say we scan a 35mm tranny at 4000PPI - this will result in a file that has 5400 by 3600 pixels.

We now want to make a print of 12 by 8 inches. This means the PPI we have available from our file for a print of this size is:

```
5400/12 = 450 PPI (or 3600/8 = 450).
```

We are choosing to use 450 pixels to represent one inch of our print.

The printer driver will now translate those 450 logical pixels into 1440 physical dots per inch (DPI) and produce a very high quality print for us.

If we wanted to know the maximum print size we can achieve at good quality, and we know from experience that given a sharp original outputted on an Epson inkjet, 240 PPI is sufficient, we can calculate the maximum print size by taking the total number of pixels available to us (5400 on the long edge) and dividing it by the PPI required to give us, in inches, the size of the print:

```
5400/240 = 22.5 inches
```

```
3600/240 = 15 inches
```

So, our final print will be 22.5 by 15 inches, with 240 pixels used to represent each inch. And of course, the printer will actually use its 1440 dots per inch to actually print that image to paper.

Hopefully that is clear - it can be a bit confusing at first. Just in case, you can read another version of the same stuff here to get a different perspective on it.

Pixels - Colour

In reality a colour image is really three grids exactly like the one above, overlaid, one for each channel – the red channel, the green channel and the blue channel.

Each pixel in a colour file is actually a composite of three colour components representing the saturation of light for that colour – a red value, a green value, and a blue value. Mixed together these three colour components give us the pixels final colour. Here are some examples:

```
(0,0,0)          - No red, No green, No blue
                 = BLACK/Ink Black

(128,128,128)    - Midtone RGB  = 50% GREY
(255, 0, 0)      - Most saturated red
(0, 255, 0)      - Most saturated green
(0, 0, 255)      - Most saturated blue

(255,255,255)    - Most saturated R,G,B
                 = WHITE/Paper White
```

The range we work between is still black (0,0,0) and white (255,255,255) but at the end of each tonal scale for each colour is not the brightest pixel of that colour, but the most saturated. So with respect to each colour we're working between zero saturation and maximum saturation. If we clip a particular colour channel, just as we talked about clipping the grey channel above, we are clipping the channel to the maximum saturated colour. This is bad because it's going to lead to lumps of undifferentiated colour, which will really stand out in a fine print.

Bits and Bit Depths

Why are the numbers we use always between 0 and 255 and not something nicer, like 0 and 100.

Actually, they're not. Well, they always are IF we're talking about 8 bit images, then the maximum numbers are indeed 0 and 255. However, if we're talking about 16 bit images, the real numbers are 0 and 65536! But it's pretty hard to draw a graph on screen that is 65000 odd columns wide, so Photoshop always pretends the numbers are between 0 and 255.

Why 0 and 255 though? Well, this has to do with binary numbers. Basically, if you use 8 bits, or a maximum of 8 ones/zeros, the biggest number you can represent is 255.

```
000000000 = 0
```

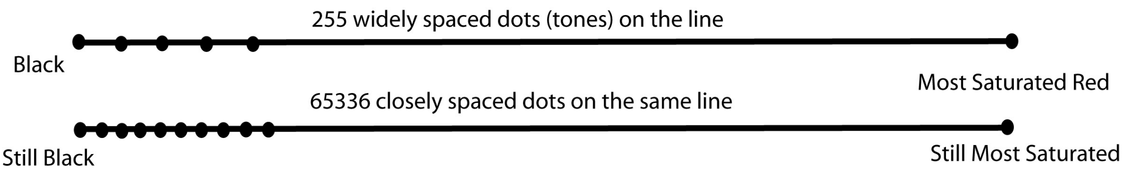

Line is the same length, there are just more dots (ie tones) on it!

8 bit and 16 bit

```
111111111 = 255
```

And if we use 16 bits:

```
0000000000000000 = 0
1111111111111111 = 65536
```

Colour, in digital form, is like points on a line. The starting point on that line is black, or the absence of that colour. The ending point on that line is the most saturated version of that colour (let's use red).

Note – and this is really important – whether you are using 8 or 16 bit images, the end points of those lines do not change.

0 is still black and 255 (or 65536 or whatever) is still most saturated red. There are just more gradations between the end points.

This means 16 bit does not mean more gamut is available to us (as everyone normally claims).

It does mean that there are more points on the line – that is, more tones in-between black and our most saturated red. So it means we have more smoothness of tonality between black and colour, not that we have extended our gamut. People get this wrong all the time.

With 8 bit, by mixing three colours (R,G,B), each between 0 and 255, together, we can get any one of 16.7 million possible colours. With 16 bit, by mixing our three values between 0 and 65536 together, there are more than 281 trillion tones possible. But the range of those tones (the gamut) is the same regardless of bit depth.

> 16 bit is always better – always use it if you can.

Here's *why* 16 bit is better - because of the very nature of Photoshop (it's really just a big maths engine), the operations you perform on colour numbers invariably result in a loss of tonality. Here's a simple example with three pixels having three different brightness values:

```
Pixel one:    9
Pixel two:    8
Pixel three:  7
```

We want to halve the brightness of these three pixels:

```
9/2 = 4.5
```

Hmmm, that's no good, this is a computer and dumb, so we aren't allowed to have floating point (i.e. decimal) numbers like 4.5 - we can only store a value of 4 or 5. Ok, so let's make a decision to round all pixels down.

```
Pixel one:    9/2 = 4
Pixel two:    8/2 = 4
Pixel Three:  7/2 = 3.5.  Round down as per the above = 3.
```

So as a result of just one operation we get:

```
Pixel one:     4
Pixel two:     4
Pixel three:   3
```

Hmmm, suddenly instead of having a nice gentle tonal drop, we've compressed the tonalities of pixel one and two together such that they are now the same value. Stretch this example out over a file with millions of pixels, and maybe do several hundred operations over a long editing session, and you've got a big problem pretty quickly – all your tonality is bunching up around the same values.

The real, tangible result of this in your print is **colour banding** – that ugly, hyper-digital look where smooth tonality transitions are replaced by crude bands of tones.

With an 8 bit image, we have at most 255 levels for each colour. This means it doesn't take too many operations before we run out of levels. With 16 bit images, there are 65000 odd levels, so considerably less chance the pixels will land on the same values, and therefore considerably less chance of banding.

The thing with levels is, we only really go down in the number we're actually using. Sure, 255 (or 65536) are available in theory, but odds are when we capture/scan, we won't use all of them – because they don't exist in the original, or our sensor is incapable of actually differentiating all those levels. Then, with a few basic operations in PS like levels or curves, we end up using even less (due to maths stuff above, like rounding errors). Finally, when we get to printing, we translate the few colour levels we are using into the printer's colour levels. So in the end, we're probably not using that many at all! Luckily, it doesn't actually take too many to create convincingly smooth tonality, but this is much easier to achieve if you start off with high bit depth scans/captures and then work in 16 bit in PS (and 32 bit editing is on the horizon).

Even if you don't follow this, **the take home point is to generate, capture and work in 16 bit whenever possible**.

Colour Models

So far we've kind of skipped over some detail - we've discussed pixels, each having a number representing a colour - and we've glossed over alternatives and discussed the RGB model for colour. But there are other colour models. RGB is the easiest to understand conceptually, and the most logical to work in, so that's where we started. But their are other Colour Models.

Models (RGB, CMYK, LAB etc)

Computers represent colours with numbers. In different Colour Models the numbers mean different things. Here are the three most important colour models:

> **RGB** - three numbers, each representing an amount of Red, Green or Blue light, mixed to make the final tone
>
> **CMYK** - four numbers per colour, each representing an amount of ink laid on a page, mixed to make the final tone
>
> **LAB** - a mathematical system to define specific points (colours) within the gamut visible by the human eye.

We'll look at each of these in more detail later.

RGB vs CMYK in a fight to the death!

Why did we start off not even acknowledging really CMYK?

After all, pretty much all printing is based on CMYK right? Well, yes, most printing is. Although that is changing with lots of alternative inksets and methods for printing, the fact remains the vast bulk of printing at its core is based on CMYK ink.

Traditionally, therefore, the print domain has been very CMYK focussed in its language and technique.

There is a HUGE issue with this, though. CMYK is the language of ink - and it's a messy, dirty, ever changing language. The fundamental problem with CMYK is that it describes a concrete physical thing - CMYK numbers are in fact percentages of physical inks going onto a physical page. The problem with this is that we rarely know precisely what the RIGHT percentages of ink on the page should be. We don't know the behaviour of the press, or the media. When originating our images, we might not even know IF We'll end up using a press, or whether our media will be a low quality newsprint or a high quality coated stock.

> So when creating our images, it's totally inappropriate to be specifying them as ink levels on a page at this point. We simply don't know enough to do this sensibly at this stage.

Also - CMYK behaves weirdly because it represents the weird, chemical world of inks - they simply don't always behave as you'd logically expect. Almost never, for example, does mixing an even amount of C, M and Y result in a neutral grey ink. And if it did - change the paper and your neutral will no longer be neutral to the new whitepoint, so your specified ink values are wrong! Another classic issue is you would logically think you'd get your best black using 100% of C,M and Y. But in fact you usually don't - for a variety of reasons. You might experience *hooking* - where the tones actually get lighter as you add more ink! Or you may encounter physical issues like *pooling* - where too much ink goes on to the page and you begin to get, basically, a puddle problem.

The tools you use are all inverted as well - logically and at a base human level, pulling something UP should make something brighter, stronger or more saturated. But in the CMYK world, up is down, black is white, and the tools simply don't make as much sense as in the RGB world.

RGB is logical and well behaved. Use it.

It separates your initial colour representation from the physical - from the mess of actual ink on a page. It is designed to match the way your eye and brain work with colour. It's easier, and creating with it properly gives a much higher quality result in the end if you know what you are doing - as it maintains your flexibility for longer, and means you will ultimately make just one conversion from a well behaved system (RGB) to your final ink-on-page ink colours. If you start off working in ink, you usually end up needing to convert from one weird ink system to another - and this always results in a loss of colour integrity. Work in RGB, create your RGB master files, and then from there move in the most controlled manner you can to print - and your RGB master files can often end up being very useful in multiple domains. If you tie yourself to physical ink too early in your process, you simply reduce your options later.

People who have been in the printing industry for a while are often VERY resistant to RGB in general. This is really based on their history (often taught at TAFE right on the press by the buckets of ink) - and not on rational thought or logic or any real in depth understanding of digital processes. Because big printers cost big money, there are a lot of wealthy old people without a lot of digital knowledge running those big printers. These printers are being rapidly undercut by smaller, smarter operators with a much more modern digital approach, thankfully.

The conversion from RGB to CMYK is really the printer's actual job - it's their job to understand their printer, to receive well and sensibly specified digital colours and translate those into the best actual colours on the page. If the printer CAN tell you a precisely accurate, specific CMYK language to deliver your files in, then you can control the conversion yourself (as the last step, from your RGB master file directly in to the best matching ink representations in CMYK) - and there can be advantages to this - but it does mean more work for you of course, and ultimately the person putting the ink on the page is in the best position to determine the best ink values - so wherever possible, work in RGB and work with printers who understand this is the way modern digital imaging works.

Colour Management Theory

So far we have covered the very basics - that colour is really made up of three things (light, objects, and sensors, resulting in a perception of colour). We've introduced basic terminology (standard light, standard observer, gamut, metamerism, and optical brighteners).

We've discussed the fundamentals of digital files - that we should be using Vector or Bitmap as appropriate, if bitmap then high resolution bitmap, and we should be storing our images as high quality, uncompressed files.

We've also covered the basics of pixels, dynamic range, clipping, and bit depths.

Finally, we discussed colour models and in particular RGB and CMYK.

So, we've really begun to discuss colour as numbers in digital files. But we all know that different screens and devices produce wildly varying colour. And that's because those colour numbers aren't being treated with respect. So now we're going to talk about how we can work properly with colour numbers to actually get the colours we want.

The Colour Chain - in Theory

The image below shows the ideal colour chain in theory – as the image moves from input to output, colour is maintained consistently across devices. It's a nice idea, but as you've probably experienced, it doesn't normally work this way.

The Colour Chain - in Practise

The way it normally works, in practice, is more like the image below. Colour varies considerably across devices.

The Problem

So far we've talked about how colours in the digital imaging world are really triplets of red, green and blue values. When these three additive primaries are mixed together, we get a colour – in 8 bit, any one of about 16.7 million possible colours, in 16 bit any one of about 281 trillion colours. (It's worth noting that the human eye can't actually distinguish anything like this amount of different colours, it's more about avoiding banding than anything else - as explained in chapter one).

The problem is that all those numbers mean different things to each different device. The thing we've glossed over until now, and it's a very big thing, is the real meaning of those numbers. At the moment they're just numbers really – nothing more.

This whole section is about giving meaning to those numbers.

Just what *does* 255 red mean? We know it's 'our most saturated red', which is a part of the puzzle. But what is our most saturated red? Is it the monitor's most saturated red? The printer's best red? Or the camera's red? What if one day we change monitors – does that suddenly mean the colours in our photos are actually going to change?

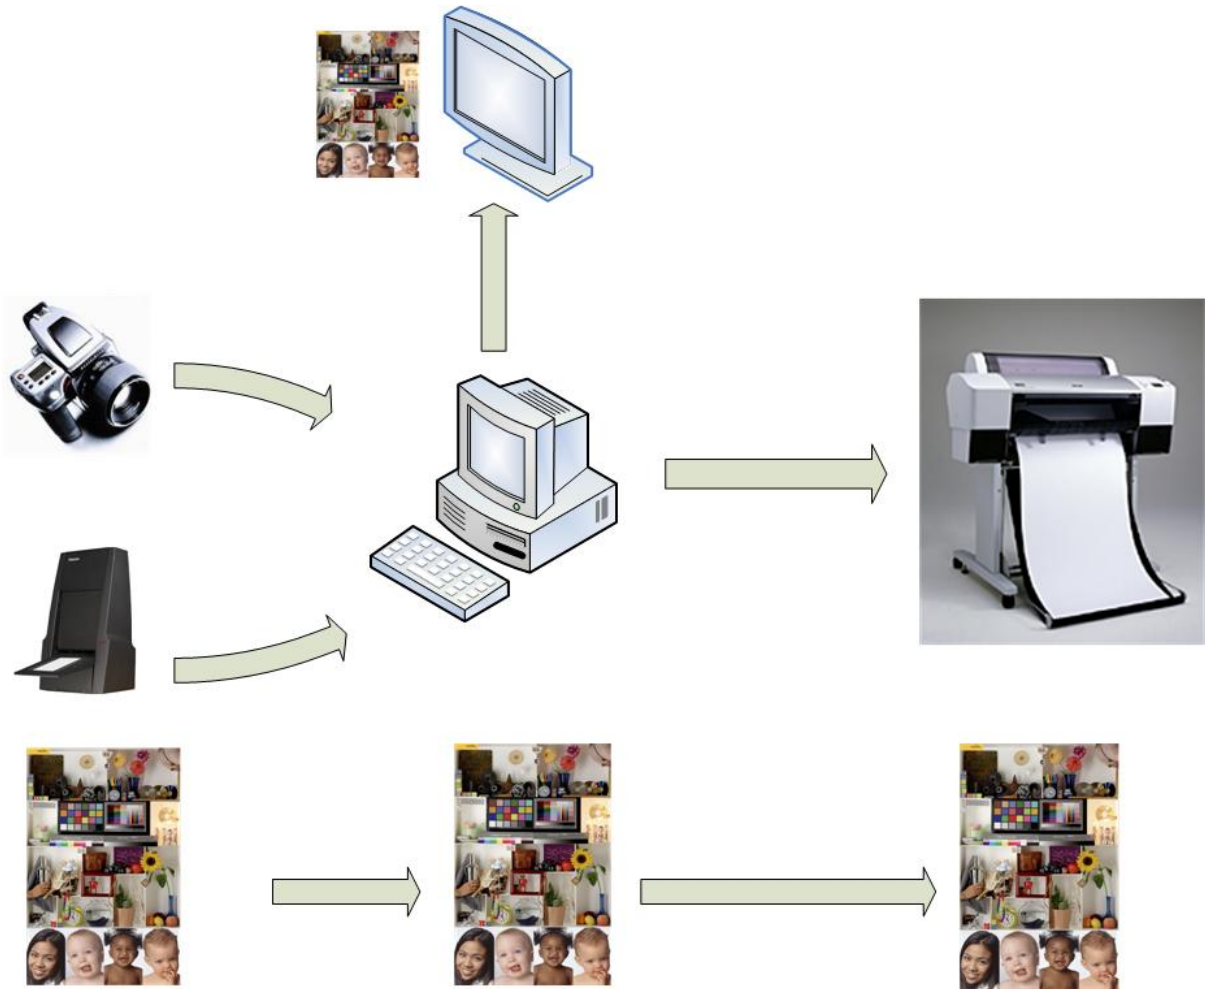

The Colour Chain - In Theory

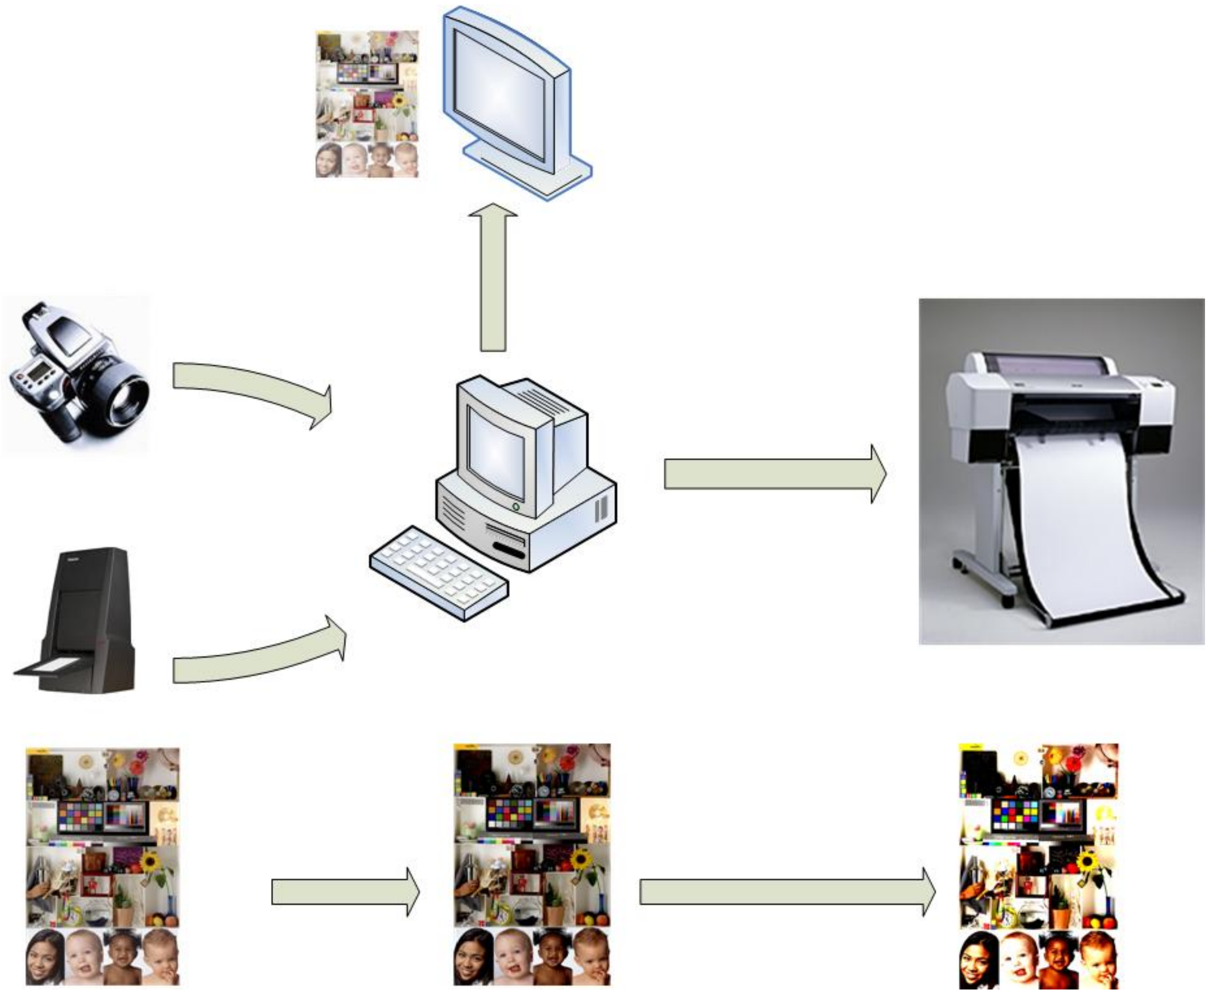

The Colour Chain - In Practise

It should be pretty obvious that it's no good using our devices as an absolute reference for what a particular colour number means – because they're all going to be different. Pretty dramatically different. And yet this is precisely what most people do every day - they sit in front of any old monitor tweaking colour and crossing their fingers - endlessly wasting time, in a wild Hail Mary hope their monitor will bear some resemblance to their final printed output. It's madness, really!

Not only are the different types/brands of devices different in their colours, but each device itself is almost certain to change its colour output over time as the device ages, or ink formulations change, or whatever. Even devices of the same brand, same model number and bought on the same day, will in practise have different colour outputs – due to manufacturing tolerances.

So what do we do? Well, just like the situation with colour above, where we defined standard light sources and standard observers, we're going to create some standards for what colours are. These are absolute standards for colour that by design have nothing to do with physical devices. Once we've got some definite idea of what we're talking about when we talk about specific colours, perhaps we've got a slim chance of convincing our devices to produce those colours.

A Quick Definition

We've said colour is a very fluid thing, highly dependent on the light source, the object properties and the observer. We've created a few standards for light sources and observers (the reference light sources D50 and D65, and the 2° Standard Observer). Only when given a specific light source and specific observer, can we talk about colours in absolute terms.

Absolute colour is a colour, under a specific light source observed by a specific observer. I'm sick of typing that so from now on we're going to refer to it as **absolute colour**. (This is my definition, not a universally recognised term – it just saves me some typing!).

The Solution - Colour Management

Colour management is all about giving meaning to colour numbers. It's about bringing colour consistency to colour across all the devices we might use – from the input end of things (cameras, scanners etc) through our image editing environment, to the final output medium (prints, screens etc).

It's not really about making a perfect screen to print match, as many people think. In fact, given the gamuts of printers and screens are pretty fundamentally different, it is pretty much impossible to achieve a perfect screen to print match across all images (so called 'what you see is what you get' nirvana). However, you can come surprisingly close, and colour management tools offer the best methods currently available for dealing with the problems as they can, and do, arise. It's really What You See and What You Get is What You Mean - (WYSAWYGIWYM?)

Colour Management in the Wild

Many people are able to find perfectly adequate working methods that suit their needs at various levels of implementation of colour management workflows – ranging from no colour management at all (eyeballing and guesswork) to the ideal chain – colour management at all stages of the process.

This course isn't about finding methods that are adequate, or quite good, it's about finding methods that are superbly accurate, methods that allow us to achieve excellent, predictable and repeatable results. Methods allowing us to print as well as anyone in the world can print – that's the goal.

The need for good colour management is unarguable. If someone tells you they don't need it, or it doesn't work – and someone inevitably will along the way - it just means they are ignorant, incapable of seeing the difference, or simply not working to a high level of quality. Avoid where possible!

It's simply the case that many people have not really developed an ability to SEE a fine print, let alone produce one. They have simply not trained their vision to appreciate the more subtle differences and the level of complexity that good prints have. It takes years, indeed a good chunk of a lifetime, to really develop the ability to see, and to properly analyse, fine prints, and each of us are at different levels.

The internet is full of colour management mythology and misunderstanding. Not surprising really, because it is a complex field and a relatively young one. It's also difficult to decide how much to learn about the field – how much of the knowledge is really necessary for practical work. But be wary of people armed with a little knowledge – it's a dangerous thing. They'll swear blind that something works a particular way, when it is obvious to anyone who has actually tested it that it doesn't.

In this course, the aim is to teach you to be a little bit smarter than the average bear. We want to go beyond the very basics, to get you to a real level of understanding - all with the goal of making a real, tangible difference to the quality of your work. This next section is quite heavy on theory, but pretty soon we'll be putting this all to work for us, and if you follow all the theory through, you'll be in a much better place when it comes to understanding the practise later on.

So, let's move on to Colour Models and Colour Spaces. These are the core concepts in colour management and learning to understand them, and use them properly, will put you ahead, in colour management terms, of at least 90% of the other working image makers in the world in terms of what you can achieve in printing.

> Good colour management is really what puts the Fine in Fine Printing.

Colour Models Revisited

We've previously touched briefly on colour models. We said there are many ways for computers to represent colour, and RGB is one such colour model – the one we are going to spend 90% of our time talking about.

CMYK is another important colour model we're going to come across in practice, but there is a third colour model that forms the theoretical basis for all colour management and is worth discussing – LAB, the colour model that best represents the human eye's perception.

LAB - the Universal Language

The LAB colour model is an attempt to model colour as the human eye sees it. **LAB represents the gamut of colour the human eye can see.** LAB is therefore, by definition, the boundary for colour that we care about.

LAB is also known as the (profile) **connection space**, because it is in fact the space that all the colour management systems do their maths in, including Photoshop's Colour Engine (ACE). It's sometimes called the universal colour space, or the translation colour space, and when you see the colour space graphs, like the one above, they are actually plots of colour spaces within the LAB colour space.

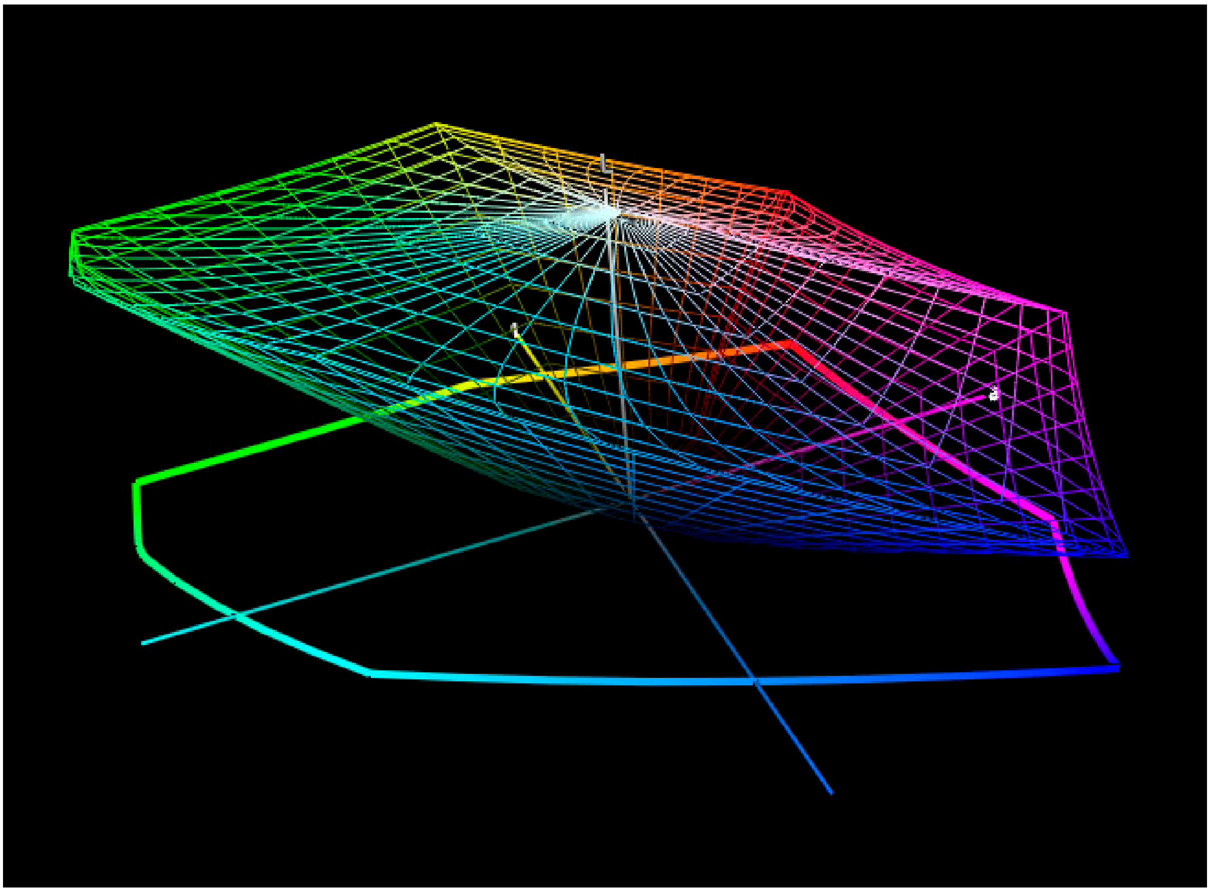

Colour Space

LAB is defined by three components – L for lightness, and A and B for chroma, or colour, components. In theory, every colour we can see has an LAB value. In practice, like most mathematical models, LAB isn't perfect, but it works pretty well and we can let the colour scientists worry about the remaining imperfections.

LAB is a colour model using three primaries (like RGB) except those primaries are a little harder to understand.

> **L*** (L star) represents the lightness of a colour
>
> 0 = black, 100 = white
>
> **A*** represents how red or green a colours is
>
> Negative indicates green, positive indicates red
>
> **B*** represents how blue or yellow a colour is
>
> Negative indicates blue, Positive indicates yellow

Any colour, represented in any other colour model, has in fact a corresponding LAB value hidden in the background. When people are getting really technical about colour, and trying to pin down absolute colours, they typically use LAB values. Spectrophotometers usually give an L.A.B. reading for a particular colour (and usually corresponding spectral data for that colour as well).

We'll be referring back to LAB at various times, you don't need to understand too much about it, the key point is that LAB values are the mathematical language for absolute colours we've been talking about.

Just keep in the back of your mind that all colours can be translated to LAB values if necessary. LAB values are like English – pretty much everyone speaks it. (Really LAB values are more like what Esperanto was supposed to be – the universal language of colour).

LAB - Why not LAB?

If we have a universal language that defines all the colours the human eye can see, why don't we just make all our devices use LAB, and why don't we then just work on our files in LAB mode rather than RGB mode in Photoshop? Why not leave all this RGB and CMYK stuff behind?

It *is* in fact quite possible to work in LAB, and some people regularly do, for a variety of reasons. However for general purpose usage, LAB isn't a great place to be because it makes colour significantly harder to understand than the RGB model.

Using RGB (and the digital colour wheel) is conceptually simple and matches how our brain thinks about colour. Through isolating colours and clearly identifying the primaries and their opposites, RGB is much easier than trying to manipulate colours in the odd language of A and B.

Also, RGB colour models generally have a simple to understand neutral axis – that is as our numbers range:

```
0,0,0 through 128,128,128 to 255, 255, 255
```

Using the arrow on the status bar, choose 'Document Profile' to show the profile for any image

we have:

```
black, neutral shades of grey from dark to light, white.
```

So RGB (the language of light) is a very convenient and conceptually simple place in which to work.

So, let's go back to the RGB colour model, but keep in the back of your mind that all colours can be translated to LAB values if necessary.

Colour Spaces - Dictionaries and Boundaries for Colour

Colour spaces are dictionaries for colour. All those RGB colour numbers in your file are given real, precise meaning**only by the colour space attached to that image** (i.e. colour spaces define those colour numbers in absolute colour terms).

If your file does not have a dictionary (i.e. colour space) attached, then the numbers in your file don't really mean *anything*, and you are completely at the mercy of whatever device is displaying/printing your image at the time.

Photoshop and colour spaces.

For any given file, Photoshop will tell you what colour space it is in:

An image with a colour space attached is said to be **tagged**. An image without a colour space attached is said to **untagged**.

Untagged images are dangerous to work with, because their colour numbers have no real meaning, and results will be unpredictable. From now on, **you should do your best to work always and only with tagged files**.

We'll talk about how to do this in a second.

But first...

What Colour Spaces are actually doing...

Colour spaces actually translate RGB triplets to absolute LAB values. This is how they give meaning to colour numbers. All 16.7 million different colours, in 8 bit, and 281 trillion colours, in 16 bit, get mapped to specific points in the LAB colour model, representing some absolute colour within the total range of colours the human eye can see.

As an example, take:

```
(128, 0, 0) - a medium saturated red.
```

In Untagged RGB (i.e. if you are NOT working in a defined colour space)

```
(128, 0, 0) = unknown LAB value
```

...which results in a monitor's best guess at this red, or printer's best guess at this red. These won't match. There's no real meaning for this set of numbers.

In a defined colour space (eg. AdobeRGB):

```
(128, 0, 0) = an absolute red
```

....defined by the colour space to mean precisely:

```
LAB (31, 55, 46)
```

This will be displayed and printed as the same absolute red by both your monitor and printer (if everything is working, calibrated, and both devices are physically capable of producing that red).

Colour Spaces Define Boundaries

By defining all colour numbers in a colour system, colours spaces therefore also **define the boundaries** of colours representable within that space.

They are abstract mathematical models that define what colour can be held within them. If a colour cannot be held within the particular model (i.e. colour space), then for all intents and purposes, that colour does not, and cannot, exist in that colour space.

As an analogy, imagine we have a bottomless container that can hold only fruits. We can put an infinite amount of anything that is a fruit into that container – apples, oranges, bananas, it doesn't matter because they're all fruits. But if we try and put a bicycle into the container, we can't do it. A bicycle simply can't exist in our bottomless fruit container because we can't put a bicycle into the container, and we can't use what's inside the container to make one either. Bicycles therefore cannot exist in this container.

In other words, we know all pixels have three RGB values between 0 and 255 which define their colour. Colour spaces define what those numbers actually mean – i.e. how dark is 0, how bright

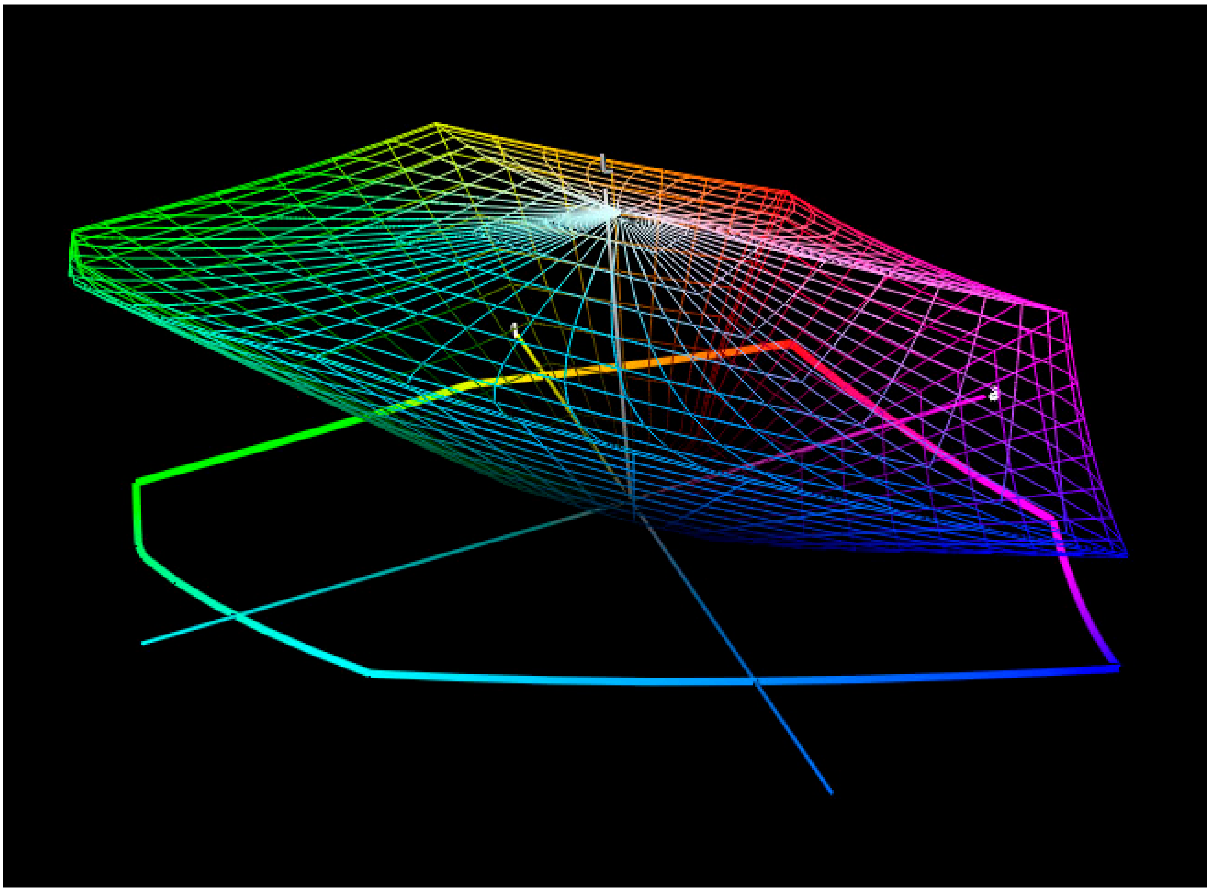

3D graph of a Colour Space

and saturated is 255, and the exact tone (in absolute colour terms) of everything in-between. But there is no such thing as -1, or 256. Colours outside of the boundaries simply cannot exist within the model because the end points are defined and there is no way to store values beyond those endpoints.

It helps to see this visually, so look at the image below. You can see a 3 dimensional shape - think of it as a funny shaped box - with black at the bottom and white at the top. Running up the middle, like a spine, are the greys – from black to white - the least saturated colours. Around the edges are the most saturated colours that this colour space (AdobeRGB 1998) can contain (the colours represented by the 255 numbers in your file if there are any). Colours more saturated (or lighter, or darker) than the colours around the edges of this colour space cannot be represented by this colour space.

To put it all another way – LAB defines all the colours we can see. RGB colour spaces define a smaller subset of those colours, and translate the numbers in RGB form in a particular file, to absolute values in LAB.

The Two Types of Colour Space

There are in fact two types of colour spaces.

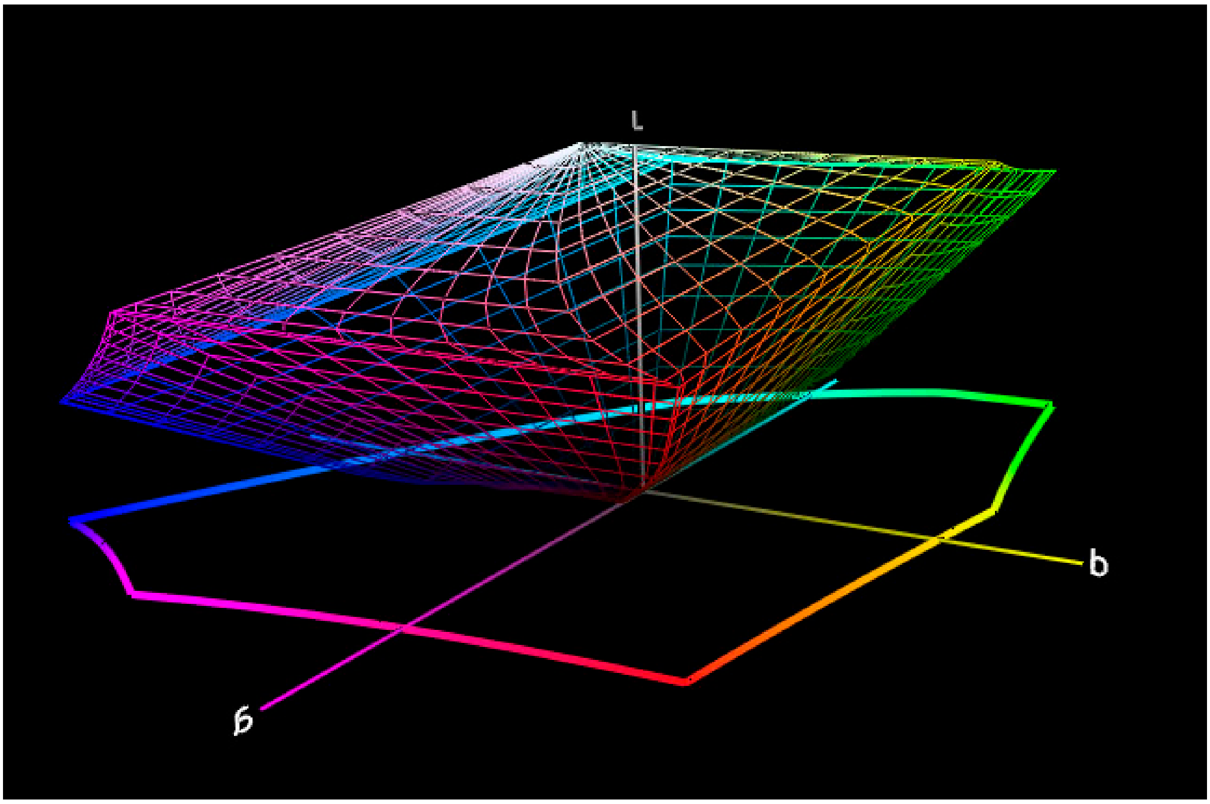

...And from the other side (note the L, A and B axes)

Abstract (= Device Independent) Colour Spaces

The first are abstract, theoretical colour spaces. There are many popular pre-defined colour spaces which we'll discuss later, and it is in fact quite easy to create one's own colour space. These types of spaces are called device independent colour spaces. They are often also referred to as working spaces, abstract spaces, theoretical spaces...they're all the same thing. **The key point is they (deliberately) have no particular relationship to a specific device.**

Device Dependent Colour Spaces

The second type of colour spaces are just the opposite – they are based on measurements of a specific device's behaviour, and thus are inexorably linked to that device – they are said to be device dependent colour spaces. More commonly, they are referred to as profiles, and they are the key to really getting colour management to work for you in practice. More on these in a moment.

Anatomy of Colour Spaces

(AKA Profiles, ICC Profiles, Tags)

As we saw above, colour spaces are definitions of actual colours – a way of tying all those RGB colour numbers to specific colours in LAB, or the gamut of the human eye.

We saw pictures of colour spaces, and have discussed them as boundaries for colour. It's time to break them down a little further to see what makes them tick.

All colour spaces have the same general anatomy (i.e. properties that make them a colour space). They all have three basic characteristics:

- Colourants
- Black and White Points
- Tone Reproduction Characteristics

Colourants

The first and biggest of these are the colourants – also referred to as the primaries. These are definitions of the three primary tones of that colour space – the most saturated versions of that particular colour.

In RGB colour spaces, we have Red, Green and Blue colourants.

In CMY colour spaces, we have Cyan, Magenta and Yellow colourants. (In reality, because subtractive primaries, and specifically inks, are pretty complicated, there are sometimes other colorants defined, we can ignore this for now)

The exact colour (and density, or ability to absorb light) of these colorants (e.g. the phosphors in CRT monitors, or the inks in printers etc) are what, fundamentally, determine the gamut of a colour space. In theoretical colour spaces, the colours are chosen, not measured.

Black Point and White Point

Other than the primaries, the other two main properties of colour spaces are the black point and the white point.

The black point is all about density – how dark is the black (i.e. how much light does the black absorb). Density is something we haven't really discussed yet. There are lots of way to measure density, or at least lots of ways to express the density value of something. Every colour has a density – not only is it a colour, but you can measure, in brightness terms, how much light the colour is reflecting (irrespective of colour). I prefer to think about dynamic range more than density because it makes more sense to me – I can say that, for example, a printer has a dynamic range of 200:1, which means it's black is 200 times less reflective than its white. That's easier to understand that saying the printer can achieve a density of 2.1, or whatever. But a lot of old school black-and-whiters really like talking about densities achievable on papers (this is largely because they own relatively cheap densitometers rather than spectrophotometers as far as I can tell). These days, it is more fashionable to talk about black points in LAB values, rather than densities.

The white point is more about colour – what is the colour of the white. Remember how we talked about a white piece of paper, under different light sources, always appearing white? This is because your eye constantly and instantly performs a process known as white point adjustment. It's really the same thing as when an eye traverses a print – the first point your eyes goes to, involuntarily and unequivocally, is the brightest point in that print. Your eye then uses that as the reference point for all other colours in the print. Essentially, your eye white balances to the brightest highlight in a print. So it's a pretty fundamental thing.

While each paper has its own white point (under a standard light source, to a standard observer – remember, colour is never absolute!), the fact is that what we consider to be white is highly flexible, and the true white point of a paper will, in practice, always be just as much about the light shining on the paper as it is about the paper itself.

White point should not to be confused with the colour temperature of white. With colour temperature, we're talking about the colour of light. With white point we're talking about the colour and brightness of the brightest white a device/medium can achieve (with respect to a specific light source and observer).

Tone Reproduction Characteristics

So far we've talked about all the things which define the edges of a colour space – the colourants (most saturated colours achievable) and the black and white points (darkest and brightest points achievable).

We haven't talked about what goes on between black and white, or between black and the most saturated colours.

The final part of the anatomy of a colour space is a set of curves (sometimes tables) that describe the behaviour of colours in-between black (0) and 255 (most saturated). That is how they travel from one end to the other - is it a liner progression or are there bumps in the road?

The Wacky World of Device Colour Spaces

Ok, now we're going to put it all together and move away from the theoretical into the practical.

Device dependent colour spaces are measurements of a device's total colour gamut – as defined by the three things above. From black to white, through all the colours, they are a measurement of what a device can do, be it a monitor, a printer, or anything else.

They're also called profiles, concrete colour spaces, and device specific colour spaces. They're all the same thing. Most often you will hear the name profiles used.

We can do this same process with ANY device that outputs colours. Screen, printer, plotter, the basic approach is the same.

Profiles

Device profiles are really translation tables between the numbers of colours in a file and the absolute colours we're trying to produce on that device.

Perceptually accurate profiles are the key to achieving good colour. They allow you to get the best possible results out of your devices. As we move from input (capture/creation) to output (prints), accurate device profiles are the only things that will keep colour consistent across multiple devices. Working without profiles is really working blind and it is staggering, in an industry so concerned and dependent on colour, that so many people are still not using accurate profiles in their work. There are only a handful of labs in this country that offer really accurate profiles for their printing services. Without them, it is damn near pot luck what colour you will get back from a lab.

Why Device Profiles are not enough...

Remember, measured profiles are device dependent colour spaces. And colour spaces define the boundaries of the colours we can reproduce by defining all those number from 0 to 255 as LAB values.

Working in any particular colour space (profile) absolutely defines the limit of the colours we can work in.

So which profile do we work in? Do we work in our monitor's colour space so that we can accurately see all the colours we're working on – no, because there are all sorts of colours are monitor can't accurately reproduce that we might want to print. So do we work in our printer's colour space? No, we might get a printer later on with a much better colour gamut, and we don't want the file to be forever limited by our current printer's colour space.

What we need is a colour space that is independent of any particular device, and preferably a good compromise for ALL the devices we are going to use – from input all the way through to output. Bring on the device independent colour space.

Device Independent Colour Spaces

AKA Theoretical Colour Spaces or Working Spaces.

You've almost come certainly come across a few of these already – sRGB, AdobeRGB 1998 and ProPhoto are probably the three most common. There are actually many of these available, and it is also pretty easy to make your own as well.

But why do we have device independent colour spaces?

Well, we have them specifically because it's no good being pushed around by our devices, and the devices we actually use will change over time anyway. To separate colour from specific devices, we use working spaces.

Device independent colour spaces are designed with specific purposes in mind. As an example, the popular working space AdobeRGB was designed to be a space large enough to contain all of the typical gamuts of monitors, and both CMYK and inkjet printers. We'll talk more about specific popular working spaces later.

We use a sensible working space - with its nice, theoretical behaviour, specifically to separate colour from the wanton vagaries of our colour producing devices. They allow us to specify colour precisely and theoretically without concern for specific devices.

The New Colour Chain in Practise

So here's the colour management workflow as we actually use it in practice:

As you can see, colour is managed at all stages.

This is achieved by using colour spaces (profiles) to define the colours at each stage from input to output.

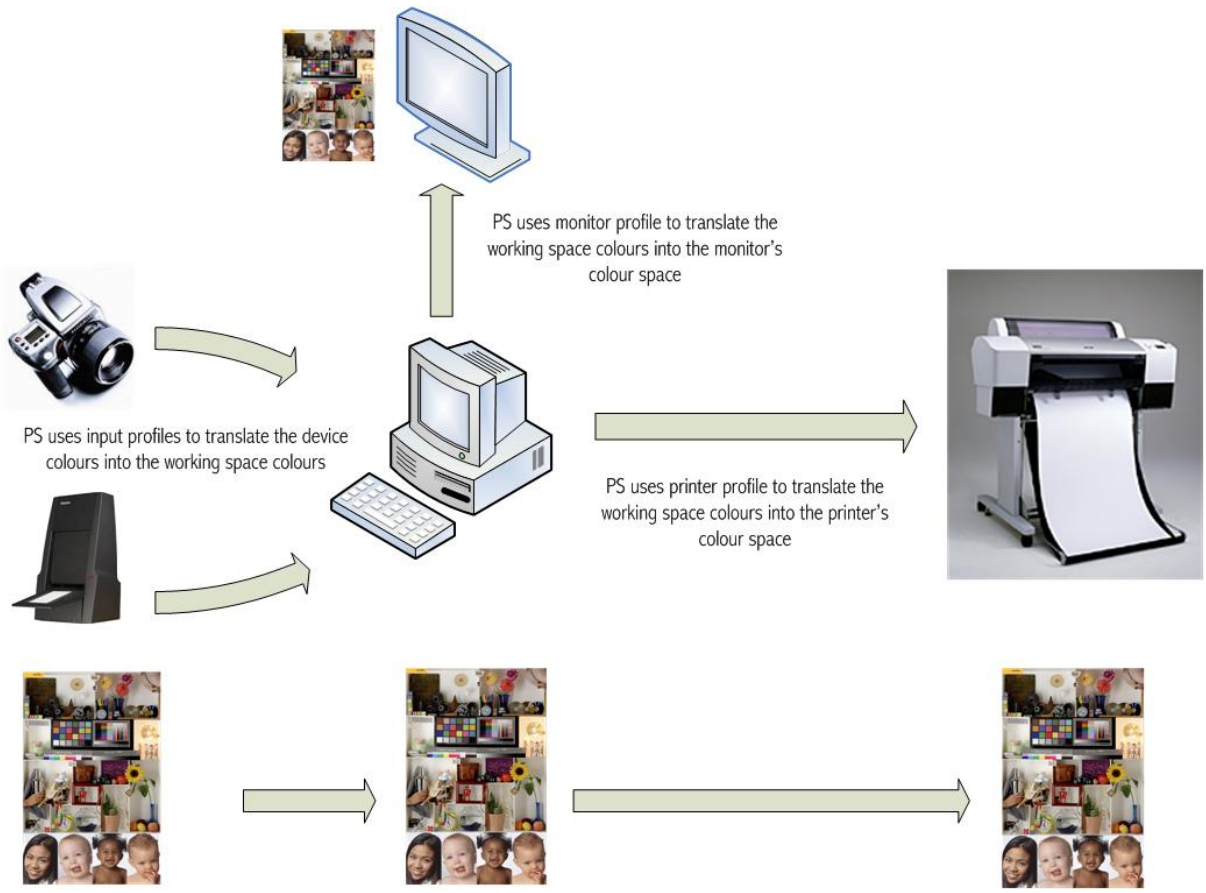

Colour Managed Workflow

At the input end, colour is translated accurately using scanner of camera profiles (in truth camera profiles are very hard to build), into the abstract working space. As we edit the image in Photoshop, we stay in this colour space, with Photoshop using our monitor profile to display the colours as accurately as possible on our screen. As we move from the working space to the output space, output profiles are used to translate from the working space colours to the output space colours for the best possible match.

It All Works, Right?

So that's it, right? Everything is hunky dory now, because we've got systems to translate all the different colour numbers into one universal colour language and get all our devices to show the same colour?

Well, that's not quite it, of course. But we're getting closer and that's all the real theory most people need to understand.

The next stage is to implement this colour workflow in the real world, and to start getting serious about device profiles.

Once we've done that, we can get some pretty spectacular results from all of this stuff, and solve many of the problems when we inevitably run into them.

So stay tuned, because in the next section the rubber really hits the road!

Colour Management Practise - Part One

In the last chapter we went through the basic theories of managing colour in a digital imaging workflow.

We discussed what we want, versus what we normally get – and discussed why this problem occurs. It's all about the numbers in the digital files and the fact that they lack any real meaning if the files are untagged – each device will just reproduce those numbers as they see fit, and because each device is using different colorants to produce its colour, the colours all come out differently on different devices.

So we then introduced the concept of colour spaces – dictionaries for colour that translate from the numbers in a file into absolute real colours – (where absolute colours are specific colours under a specific light source - and all have value in the LAB colour space which defines the gamut of the human eye).

So, at this stage we know that all pixels are just RGB number triplets like `(128,128,128)`. And we know that colour spaces give real meaning to those numbers (i.e. define a translation between the numbers and absolute LAB values).

We also know that there are two types of colour spaces – those that are device dependent (based on measuring the physical properties of a device) and those that are deliberately device independent (theoretical colour spaces that abstract colour away from individual devices). We use the device independent colour spaces specifically because we don't want our colours to depend on the mood of our inkjet on a particular day – we want to talk about what colours we want in absolute terms, and then get those colours out of our devices.

We went into some detail of the anatomy of a colour space – that each colour space has 3 defining characteristics – namely the colorants (primaries) used to mix all the other colours, the black/white points, and the tone characteristic curves.

We considered colour spaces as boundaries for colour – they define what all the numbers from 0 to 255 mean for each colour, and there is no -1 and 256. So by defining the end points we are defining the maximum achievable gamut in that colour space. If we're talking about a device profile – that is measured – i.e. we measure the most saturated colours, the darkest black, the brightest white, and all the colours in-between to come up with the profile. In an abstract colour space, we choose what the end points are. We'll get into more detail about this today, and really convert all this theory into working practice.

Making it work in practice

This chapter is all about making colour management work in practice. Even if all the theory makes your head spin and you can't quite get your head around it, the actual USE of colour management is actually pretty simple. You'll probably find the theory becomes more understandable as you consider how it works in practice, and we'll go over several examples to make sure it all makes sense.

Some Real World Examples - Device Gamuts vs. Working Spaces

It's worth looking briefly at some real world measurements now, to give you a stronger idea of what we're talking about.

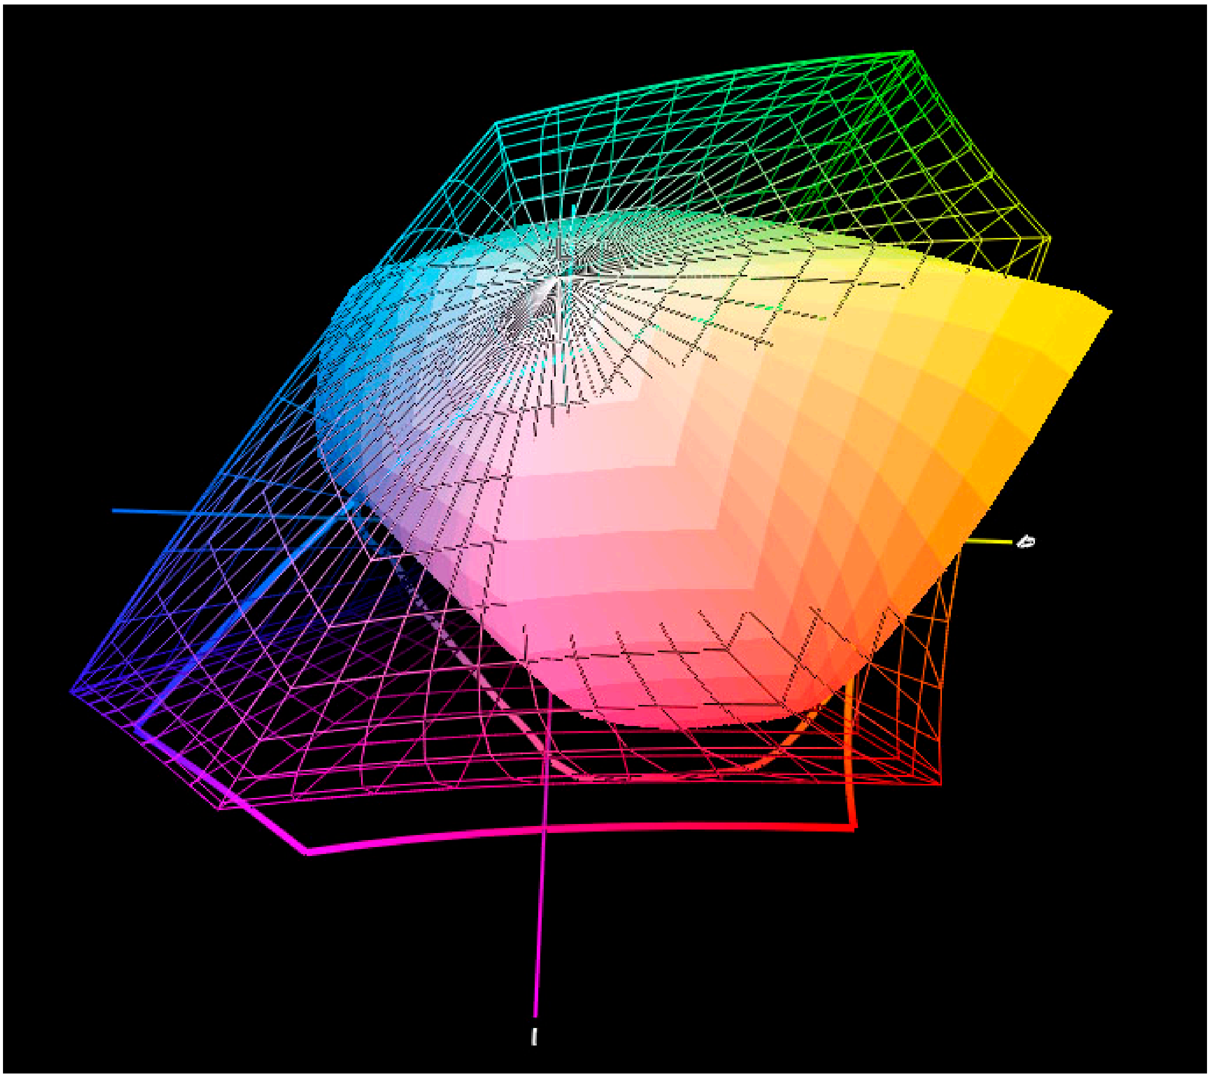

Gamut of Screen vs. Print

The image below is a plot (seen from above, looking down on the whitepoint) of the gamuts of a high quality Eizo LCD monitor (the wire-frame graph) versus the measured gamut of an Epson 7800 inkjet printer using Hahnemuehle Photo Rag paper.

You can see how the monitor's gamut extend much deeper into the red, greens and blues than the printers – not surprising because a monitor is inherently an RGB based device (its colorants are red, green and blue). The printer's gamut extends more heavily into the cyans, magentas, and yellows – again not that surprising since the colorants of the device are CMY.

The graph shows clearly one of the fundamental problems of colour management – what we can see on our screens is vastly different to what we can print. There's a whole range of saturated yellows our printer can achieve that we can't see on screen, and vast numbers of blues, greens and reds that we can see on our screen that we can't print.

Notice that these plots only really show you gamut volume – i.e. These plots only really show you the edges of the gamut, the most saturated colours that exist for the profile. These plots don't really tell you anything about what is going on *inside* of a profile.

Here's a plot of the gamut of a modern-ish digital camera (Canon EOS 1DS MK2) versus the Epson

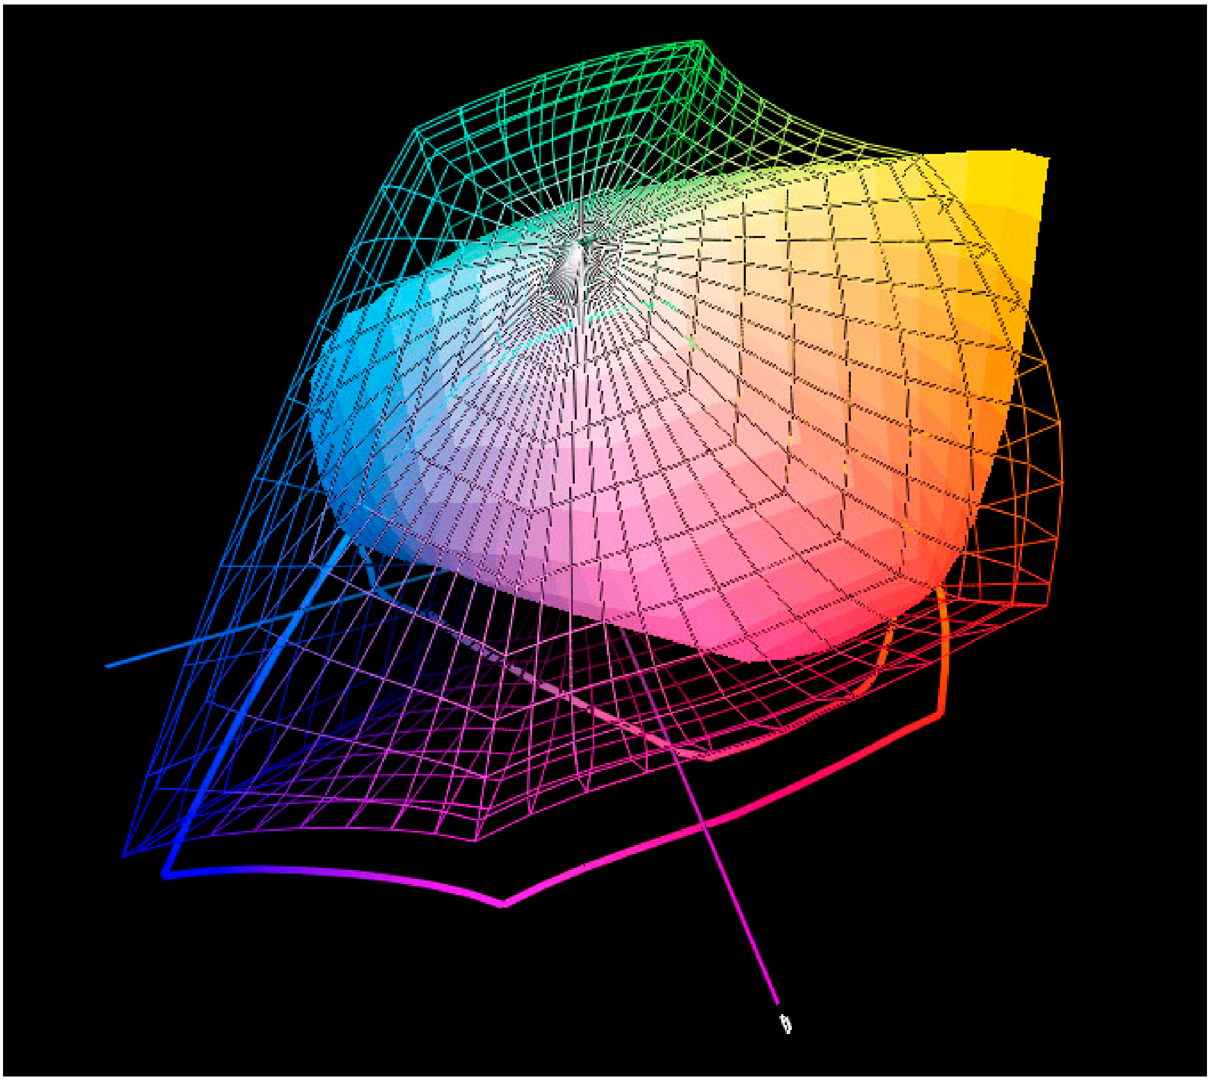

Gamut of Camera vs. Print

7800 – as you can see the gamut of the camera far exceeds what is printable on paper.

These images are really here just to remind you that all this theory has a very practical context – profiles clearly define the range of colours a device can produce. And abstract spaces clearly define the total gamut of colours you can work with in your files. And clearly, these two gamuts are not the same!

More on Wacky Device Profiles

A device profile is like a map of a device's behaviour across its entire gamut. Your ability to take advantage of that gamut is totally determined by the accuracy of that map. But the map reflects the true behaviour of the device as measured, and it may not be a pretty thing. It certainly isn't a great place, in general, to work with colour.

The device profile may well reflect all sorts of things you don't expect/desire – one very obvious and common issue is that devices rarely have linear, neutral grey scales. That is, if you're working in the device's colour space, you can't rely on all equal value RGB triplets (`10,10,10` and `20,20,20` etc) being a neutral grey, like you can in a well behaved working space. Device profiles also reflect the actual gamma of that device, which may differ considerably from the standard 2.2 used to reflect the gamma of a print (i.e. tone curve from shadow to highlight). So actually working within a device profile may (will) mean you get unexpected results with both colour and density.

We'll talk more about this when we talk about Early Binding Vs Late Binding – the key point to take on board is that device profiles reflect the real world chemistry of inks, or the real world physics of phosphors, so they're not necessarily nice, neat things, and they don't always behave like we might logically expect.

Putting it all together

In practice, colour management works like this:

RGB values are created by a camera or scanner. These colours are converted by the scanner profile into the closest matching colours available in our working space. Then, just before we send these colour numbers to the printer, the printer profile translates those numbers into the best matching colours the printer can produce.

To break that down, we're going to look at how it all works, from capture to output, for just one single pixel in a file.

The Journey of a Single Pixel

The easiest way to really work out what is going on is to imagine the life of a single pixel all the way from capture to output – that is, how is the colour on just one pixel managed (kept consistent) all the way from initial capture all the way to final output.

If we can manage the colour of one pixel all the way through the chain, then all we have to do is follow exactly the same formula for all the pixels in our file and we'll have successfully managed colour from capture to output.

Doing it backwards

We're going to start at the print and work our way back to the original colour in the original scene. This is because its always easier to work towards something if you know exactly what it is you are working towards – this is true whether we're talking more generally about pre-visualisation of your final image, or just the specifics of the colour of one tiny dot on a page.

Here's our goal – we want to produce the absolute colour as defined in LAB as:

```
(20,-15,25) - a leafy green colour.
```

We want to produce a single dot of this absolute colour on a nice piece of fibre based Hahnemuehle Photo Rag paper – that is, we want to lay down some mixture of ink, that when viewed under a standardised light source (D50), appears as that absolute colour.

So how do we get there? How do we get the printer to produce this colour?

Well, first we have to know if the colour is in our printer's gamut at all – that is, is there anyway to tell the printer to produce this colour at all? If the colour is not achievable on the printer at all, then we've failed – the best thing we can do is lay down the closest thing to this colour, and hope it is close enough to be convincing. If the colour IS in the printer's gamut, then we just want to send the printer the right numbers to get it to produce this colour.

It's the printer profile that tells us whether or not a particular colour is in a printer's gamut (remember the printer's profile is a device dependent colour space created by measuring the printer's total gamut). The printer profile tells us what RGB values we need to send to the printer to get the printer to actually print that colour. In this case, the printer profile tells us that to get this particular shade of leafy green on to paper, we have to send the printer the following numbers:

`RGB (71,100,14)`

....the colour numbers the printer actually needs to receive to trigger it to produce leafy green

`(= LAB(20,-15,25))`

So, to achieve this colour on the printer, we have to send the printer a single RGB pixel with those numbers.

We've just taken one step back from our goal, and worked out how to achieve that goal.

How do we get to `RGB(71,100,14)`?

We have to send our printer `RGB (71,100,14)` to achieve the leafy green that we want, but how do we get to this number?

Well, we're photographers and we create images by capturing light with a sensor. Whether that sensor is a digital camera or film, it doesn't really matter. Either way, the sensor in question generates RGB colour numbers to represent the colours of the real world. Problem is, all the sensors generate different colour numbers. Could we just tell our camera to always record colours as our printer sees those colours? Well, we could, but we'd be forever tied to that particular printer – and we know that printers can't reproduce a lot of colours of the real world. We really want to take full advantage of the complete abilities of our camera (which are already pretty feeble) and not limit our camera's ability to record the world based on our printer. Especially since there is almost certainly a printer in the near future that will be able to produce a wider gamut.

Hopefully you can see we need an independent step in the middle. That is, take the best possible colour we can record at the input end, and translate that into the best possible colour we can achieve at the output end.

In this case, imagine we are taking a photograph of a plant. This plant, when we shoot it, contains the same leafy green we want to reproduce – the absolute colour `LAB(20,-15,25)`. We'll shoot this image on film and scan it, to keep things simple for now.

We know the scanner produces RGB values, not LAB values. In this case, the scanner produces the values of `(30,70,25)` for the same shade of green. If we sent this directly to the printer, we wouldn't get the same green – we'd get the printer's version of `(30,70,25)` which is a completely different green. So we have to translate from the scanner's green, to the printer's green. Which is pretty easy – we just convert (30,70,25) from the scanner into the printer's equivalent green (71,100,14) and we're done.

In reality, it is Photoshop that sits in the middle and does these conversions. To do these conversions, it looks at the RGB <-> LAB tables in each profile, and it basically works like this:

```
Scanner (30,70,25)  = LAB     (20,-15,25)
LAB     (20,-15,25) = Printer (71,100,14)
```

So you can see how profiles are being used to translate between different device dependent RGB values by converting those numbers into universal, absolute colours in LAB.

A printer profile is really just a big table of these match-ups – to print such and such a LAB colour, the printer requires such and such RGB.

The scanner profile is just the same – a big table that says, for each colour, when the scanner generates such and such RGB, it really means such and such LAB.

So we go from RGB -> LAB -> RGB.

This works, and if you have complete tables representing the gamuts of the scanner and the printer, it is sufficient if all we want to do is send files directly from the scanner to the printer – the colours will be consistent across these devices and the job of colour management is done.

Journey of a Single Pixel, Visualised

We've just discussed the life of a single pixel from capture by scanner through to final print on fibre based paper. And we've succeeded in translating colour between devices by using profiles (our colour dictionaries). This is a great leap forward, but not yet nearly enough. All this isn't much good to us if we can't actually see and edit our files along the way.

So far, we have three different values for the colour we want to achieve –

```
The LAB absolute colour:
LAB(20,-15,25)

The RGB colour our scanner generates for this absolute colour:
RGB (30,70,25)

The RGB colour our printer needs to be given to print this colour:
RGB (71,100,14)
```

Your monitor is also an RGB device, and it too has its own colour response. So it too needs to be sent an RGB colour that causes it to display its best match for the same absolute LAB colour. In this case it is your monitor's profile that supplies the RGB values that best match the LAB colour we want to achieve, and the results is:

```
RGB (49, 24, 5)
```

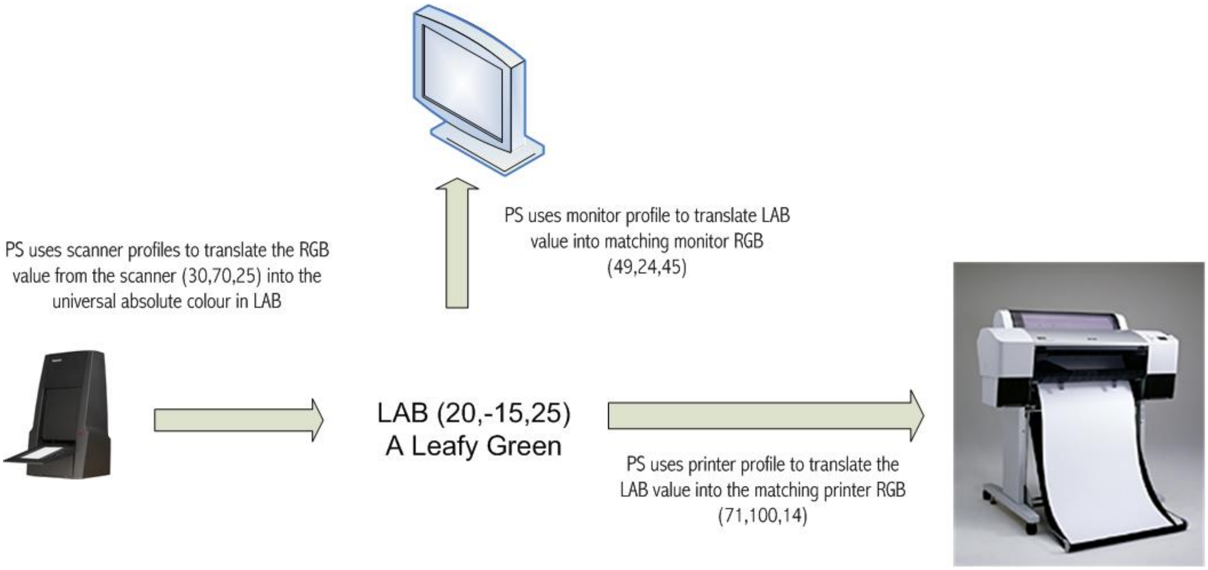

Journey of a Singple Pixel

So, for just this one pixel, we have a pretty complicated situation, which looks like this:

The result of all this?

One single pixel of a leafy green colour that is consistent from capture, through visualisation, to print. Multiply by several million and you have yourself a fully colour managed workflow! At this point we can capture the colour, see the colour on our screens, and print the colour. But what if we want to actually edit that colour?

Fine printing is very much about managing tonal relationships between pixels in your files – and for that we need Photoshop.

More on LAB

From all this, we can see LAB is pretty important. It sits in the middle of all the devices and acts as the universal language that translates from RGB numbers into consistent, absolute colours. Photoshop actually does all of its maths in LAB. As PS converts colours from one device to another, it uses the tables in profiles to convert LAB values to RGB. This is why LAB is called the Profile Connection Space.

Why don't we always use LAB then?

If we have a universal language that defines all the colours the human eye can see, why don't we just make all our devices use LAB directly, and why don't we just work on our files in Photoshop in LAB mode rather than RGB mode in Photoshop?

It is in fact quite possible to work in LAB, and many people regularly do, for a variety of reasons. However for general purpose usage, LAB isn't a great place to be because it makes colour harder to understand than the RGB model.

Using RGB (and the digital colour wheel) is conceptually simple – by isolating the colours and clearly identifying the primaries and their opposites – this is much easier than trying to manipulate colours in the language of A and B. Also, RGB working colour spaces generally have a neutral axis – that is from 0,0,0 through 128,128,128 to 255,255,255 we have from black, through shades of neutral grey, to white. So RGB (the language of light) is a very convenient and conceptually simple place in which to work.

LAB isn't so friendly and obvious. You can try it for yourself – convert one of your files into LAB and attempt to edit the colours. You will find strange things happen. It's not really important why this is the case, just that it IS the case.

So working in LAB is no good to us.

The Working Space

Instead of working directly in LAB, which doesn't make much conceptual sense, we (generally) choose to work in RGB, because RGB is the language of light, the language of the human eye, and the language of most of our devices, so it's the easiest one to understand and use.

This means we need yet another translation to occur, from the LAB colours in the middle of all this, to the working RGB space. We'll use a popular working space called AdobeRGB as an example, and later on we'll discuss popular working spaces in general.

In this example, the AdobeRGB values for our leafy green are `(44,56,18)`

From Scanner to working space it goes like this:

The scanner colour:

```
(30, 70 , 25)
```

is generated by the scanner when it sees the absolute colour

```
LAB(20,-15,25))
```

which is converted to the equivalent working space colour

```
AdobeRGB(44,56,18)
```

From working space to screen it goes like this:

```
AdobeRGB(44,56,18)
```

is converted to

```
LAB(20,-15,25)
```

which is converted to the monitor's RGB:

`(49,24,45)`

From working space to printer it goes like this:

`AdobeRGB(44,56,18)`

is converted to

`LAB(20,-15,25)`

which is converted to the printer's RGB

`(71,100,14)`

...and thus all the devices, *finally*, are producing the same colour because profiles are converting the actual absolute colour to the RGB numbers that correspond to that actual colour, and we've also got our file into an easy to understand abstract working space for editing.

Colour Management, At Last!

In summary: **Ignore LAB!**

While LAB is the colour model used in the background, and it is very useful to understand what it is, for the most part, you can ignore it. You'll probably only ever work with RGB files and think in RGB values for a good long while until your understanding of this deepens over time.

In the end, in practice, colour management looks like this:

It All Works, Right?!

The end of the journey

Using accurate device profiles, we have translated the colours between devices and abstract working spaces, all the while keeping the colour consistent – even though all our devices use different numbers to represent the same colours, we get the same colour on each device – which is what colour management is all about.

So there are two major components to a colour managed workflow:

- The device profiles that accurately describe the behaviour and gamut of devices
- The abstract working spaces we use while in Photoshop.

Good colour management is all about getting/making good device profiles and making informed decisions about working spaces. So that is what we'll talk about next.

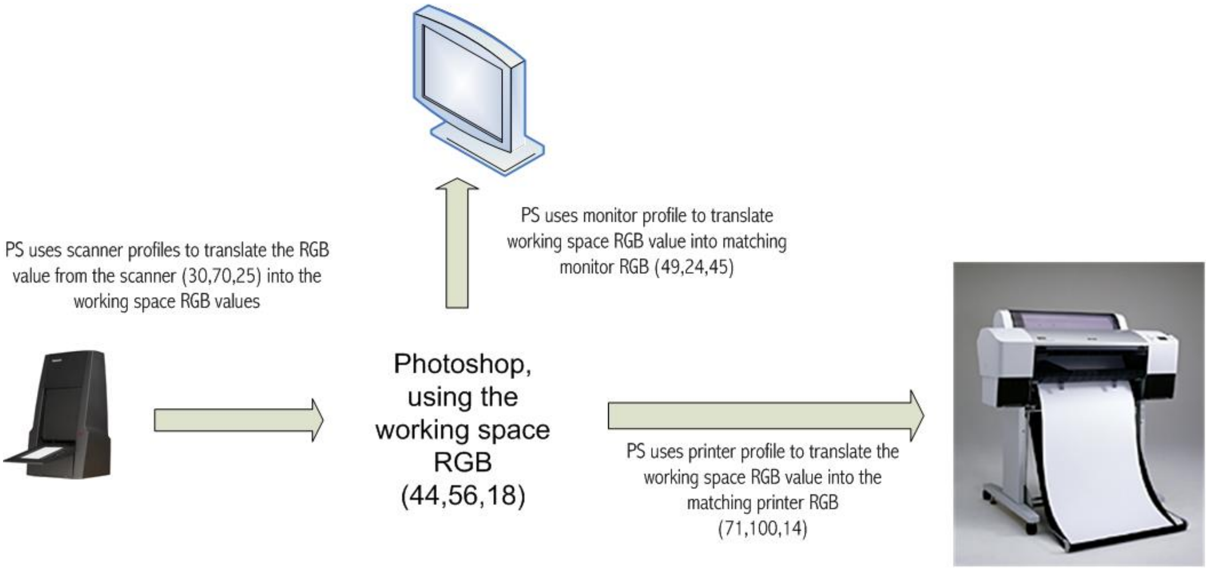

One Pixel, Fully Colour Managed

One More Thing!

The above example shows us what happens when everything works nicely – but of course in practice there are always some glitches in the system. Not only are there glitches in the system, but there are some inherent limits in the system – and we've already seen one major reason why (the gamuts of devices do not match up nearly as much as we might like).

The first thing to achieve is to get colour management working perfectly in simple situations – like reproducing one, in gamut, colour that all the devices in the chain can comfortably deal with. Once we've achieved that, we can look at solving the problems when we inevitably run into them. And in reality, even with the flaws in the system that remain, we're well on our way to good results.

Colour Management Practise - Part Two

In the last section we began to move from the theory of colour management into the practice of colour management.

We looked at the real measured gamuts of a variety of devices, and we discussed the life of a single colour managed pixel all the way from capture, through a working space, to display on screen and to print.

We then had a look at actually getting good device profiles – the key to colour management in practice We looked at capture, screen and output profiles, and concluded that there are a bare minimum set of requirements in colour management:

- An accurately calibrated screen using a hardware based measurement device
- Accurate output profiles describing the printer's behaviour
- Careful selection of the working space (more about this in this section!)

In This Episode...

We're going to finish off putting it all into practice and get you to the point where you can soft proof and then make a proper colour managed print. Along the way we will set up Photoshop correctly and discuss RGB working spaces in greater detail.

We're also going to discuss the limits of the system, because until now we've really been looking at simple situations – where there is a colour that is in gamut for ALL of our devices from capture to print. In reality, we know the world is vast and rich in colour, colour far beyond what we can capture and print, so we need to look at how we deal with those pesky colours that are out of gamut.

Journey II: Journey of a Single Out of Gamut Pixel

In the first Journey of a Pixel we followed the life of one leafy tree green colour from capture to print, as carried through our camera, scanner, screen and printer. The leafy green colours happened to be an in gamut colour for all of our devices – all of them had the physical means to reproduce that colour. We now need to look at what happens to colours that are out of gamut – that is, where some (or all) of our devices simply cannot capture or reproduce that colour.

If we can't capture the colour in the first place...

Then we have no hope of reproducing it further down the line – hence the need for high quality camera sensors and lenses (Even more importantly, high quality film, and film scans, if we're still shooting film). So the best way to solve this problem is to throw money at it....if you want to be a quality photographer, you need to use quality tools.

The thing is, cameras will always see and record something, even if it isn't an accurate colour. We went over this when discussing input profiles, so I won't repeat it. Suffice it to say, we're going to get some sort of colour captured, and we want to reproduce this further down the chain as best as is possible. But we can't improve the whole process if we don't improve the front end. It's just that improving the front end is so expensive (i.e. buying newer, better devices) and so dependent on what the sensor makers come up with, there is little we can actually do about it.

Let's assume...

We've captured the colour (or the closest thing to it the sensor can produce) – now we want to carry that colour all the way along our devices, through visualisation, to print.

Let's say we've got a really deep blue. This colour is in our camera's gamut, but not the screen's gamut and not the printer's gamut. How can we see that this is the case, and how can colour management help us solve the problem?

Capture

We use a digital camera and a raw converter to get the file into Photoshop. We'll use ProPhoto for our colour space this time - we'll cover that in a bit more detail soon, but let's assume we somehow arrive at a file in ProPhoto such that our pixel is:

```
ProPhoto RGB (10,10,80)
(= LAB (3,15,-68))  (an extremely deep and saturated blue)
```

This colour is not within our monitor's gamut – that is, our monitor has no way of displaying this colour. So Photoshop has no choice but to convert it to the nearest matching colour our screen can display.

```
ProPhoto RGB (10,10,80)
    = LAB (3,15,-68)
    = nearest matching LAB (9, 29, -57)
= Monitor RGB (0,0,110)
```

As you can see, the closest colour our monitor can use is 110 blue, and that is quite a different absolute colour from the absolute colour we're trying to achieve. But we can only display our best approximation of the blue on screen and hope that is sufficiently useful to be worth something.

So we can't actually see this colour on screen. What happens when we try and print it?

Inkjets are not wonderfully good at very deep blues. This is a fact of life – it's simply not very easy to mix a deep blue using CMY colourants. So the first place colour management might help us is to say – maybe an inkjet printer is not the right machine for achieving the best print of this image – maybe an RGB light based printer, like a lamda, might be a better tool for the job. Let's see what happens.

The file is sent from the working space to the inkjet printer.

(We're going to assume the Relative Colorimetric intent for now, more on this in just a second)

```
ProPhoto RGB (10,10,80)
    = LAB (3,15,-68)
    = becomes closest LAB (19,5,-46)
= Printer RGB (2, 93, 199)
```

So the final colour numbers we're sending to the printer are WAY different to the original ProPhoto RGB numbers when trying to achieve the best possible match for this colour – i.e. `(10,10,80)`

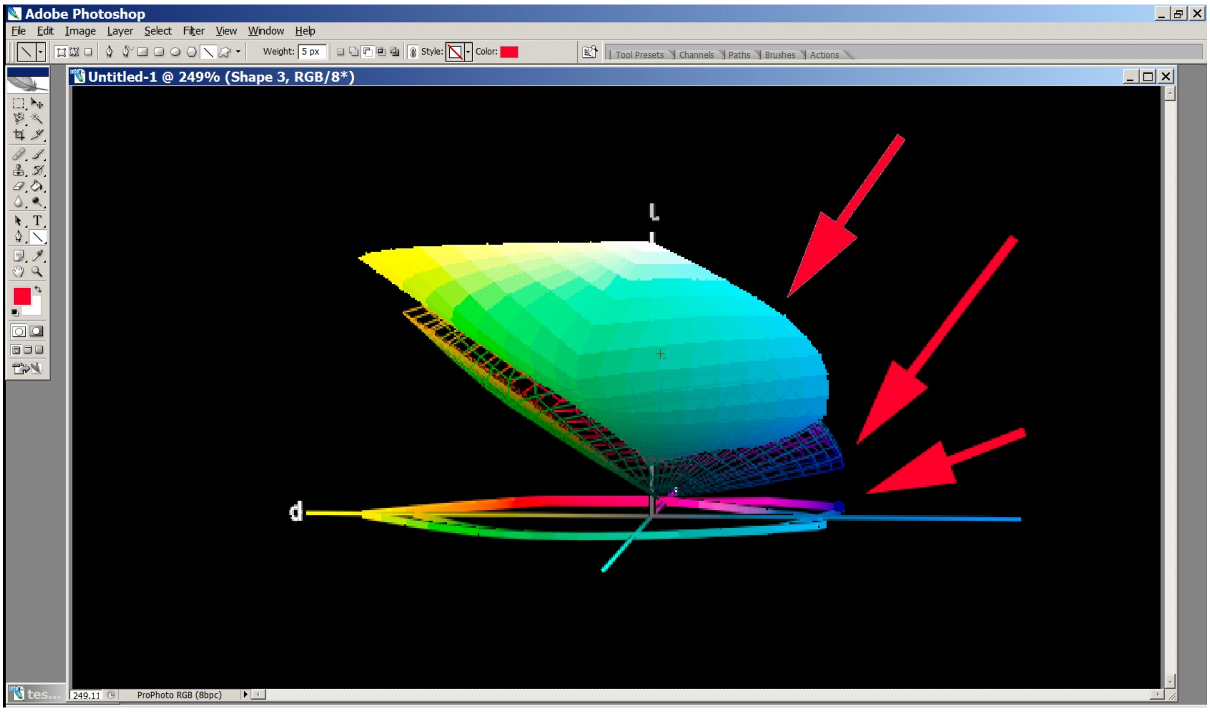

Out Of Gamut Colour

becomes (7,108,202)! And this is just the printer's best approximation of that colour, which is a long way off the actual colour (which we can see by comparing the LAB values).

The file is sent from the working space to the Lamda printer.

ProPhoto RGB (10,10,80)
= LAB (3,15,-68) = becomes closest LAB (17, 7, -49) = Printer RGB (1,31,132)

The RGB numbers are still way different, but a little bit less different. The LAB numbers are also closer together. So the match isn't perfect, but better, i.e. the colour is out of this printer's gamut as well, but not by as much.

A visual representation looks like this:

The bottom arrow points to the blue dot that is the actual colour plotted in the LAB space (i.e. plotted as a point within the gamut of the human eye). The wire-frame graph (middle arrow) represents the gamut of a Lamda printer, and you can see how the gamut of that printer extends more deeply into the deep blues than that of the inkjet printer (the solid gamut plot pointed to by the top arrow). This, in a nutshell, explains why inkjet printer manufacturers are all busy adding blue inks into their printers – specifically (and obviously) to extend the blue gamut!

So, for this particular out of gamut colour, we're going to get closest to its reproduction by sending the file to the Lamda printer. Of course, we might not have a Lamda, but in both cases, use of the profile will give us the best matching blue that is within the printers gamut.

From one pixel to a glorious image...?

This is just one pixel – obviously, if we have a lot of these pixels that are out of gamut, our final print is going to be a long way from matching reality. If we simply take all our colours that are not in the printer's gamut and map them all to the edge of the printer's gamut, then we're going to get lots of blobs of undifferentiated colour. There are in fact 4 different ways the computer automatically deals with out of gamut colours, and these are known as **rendering intents.**

Rendering Intents

Rendering intents are algorithms (recipes) for the automatic mapping of colours from one profile to another.

Rendering intents deal with the out of gamut colours – those colours that do exist in our starting profile (usually a working space) but don't exist in the gamut of our output device (monitor or printer).

There are two basic ways you can transfer colours that are our of gamut from one space to the other – you can map them all to the edge of the new spaces (i.e. gamut clipping the colours to their most saturated equivalents in the destination space) or you can re-map ALL of the image's colours, trying to preserve the relationships between colours (gamut compression).

Basically, as you try and fit 12 clowns into a mini, and you run out of room at 8, you can either chop pieces off the clowns, or you can squish the clowns into the mini by getting them all to sit closer together!

Each rendering intent produces a new image from the source image, with the new image in the output colour space.

There are no right and wrongs with these intents – you can use any, as it pleases you - the decision is purely aesthetic.

Relative Colorimetric

(The Gamut Clipping Intent)

This transfers all in gamut colours to the destination image exactly. All out of gamut colours are clipped to the closest reproducible hue. This is the chop off approach, and tends to work well with images as the process preserves as many of the original colours as exactly as possible. However, if large areas of your image are out of gamut, the result is large areas all having a single hue (the closest, most saturated colour to the original). This means the result can be large areas of undifferentiated tonality. So in practice, this works well with images where only a small amount of the colours fall outside of the destination gamut.

(This rendering intent is the one Photoshop uses to display images on your monitor).

Perceptual

(The Gamut Compression Intent)

Aims to preserve relationship between colours within an image. All colours, including in gamut colours, are remapped into the destination space by remapping the end point (the most saturated colour in the source image) to the end point in the new image (the most saturated matching hue in the destination space), and all other colours within gamut by an amount that aims to maintain their visual relationships.

Perceptual tends to work best when you have an image with a lot of out of gamut colours – instead of all being mapped to one hue, the resulting image will have the tonal relationships in tact, that is, there will be some tonal differentiation.

Absolute Colorimetric Intent

(The Proofing Intent)

This rendering intent behaves like Relative above, but the whitepoint is handled differently. It's called the proofing intent because when you make a print with this intent, ink is used to simulate the whitepoint of the source colour space. So one can render absolute colorimetrically from one printer profile to another as a means of simulating the output of the first printer on the second, media colour included. There's a lot more to this, and doing it well, of course.

Saturation Intent

(The Punch Intent)

This rendering intent throws colour accuracy to the wind and simply tries to match the 'punch' or saturation level of the image when moving into the destination space.

Setting Up For Colour Management

(Photoshop Example)

The screenshot below shows how you should set up your colour settings in Photoshop for general purpose imaging use (unless you have a good reason not to!).

The most important thing is the RGB working space – set this to AdobeRGB for now (or read the working spaces section below and choose another sensible working space). Under Colour Management Policies, all should be set to 'Preserve Embedded Profiles' and all 'ask when opening' boxes should be ticked.

Setting Photoshop up in this manner means you will start making explicit choices about how Photoshop will handle colours in files.

Dealing with Colour Management Issues

(Photoshop Example)

By setting your Colour Settings to the above, you may initially find yourselves getting all sorts of colour warning messages when you now open files in Photoshop. While perhaps annoying, in the long term this is a good thing as it will force you to make informed decisions about what is happening to your colour rather than letting Photoshop choose for you (and as you begin to use colour management more consistently, the messages will stop appearing in time).

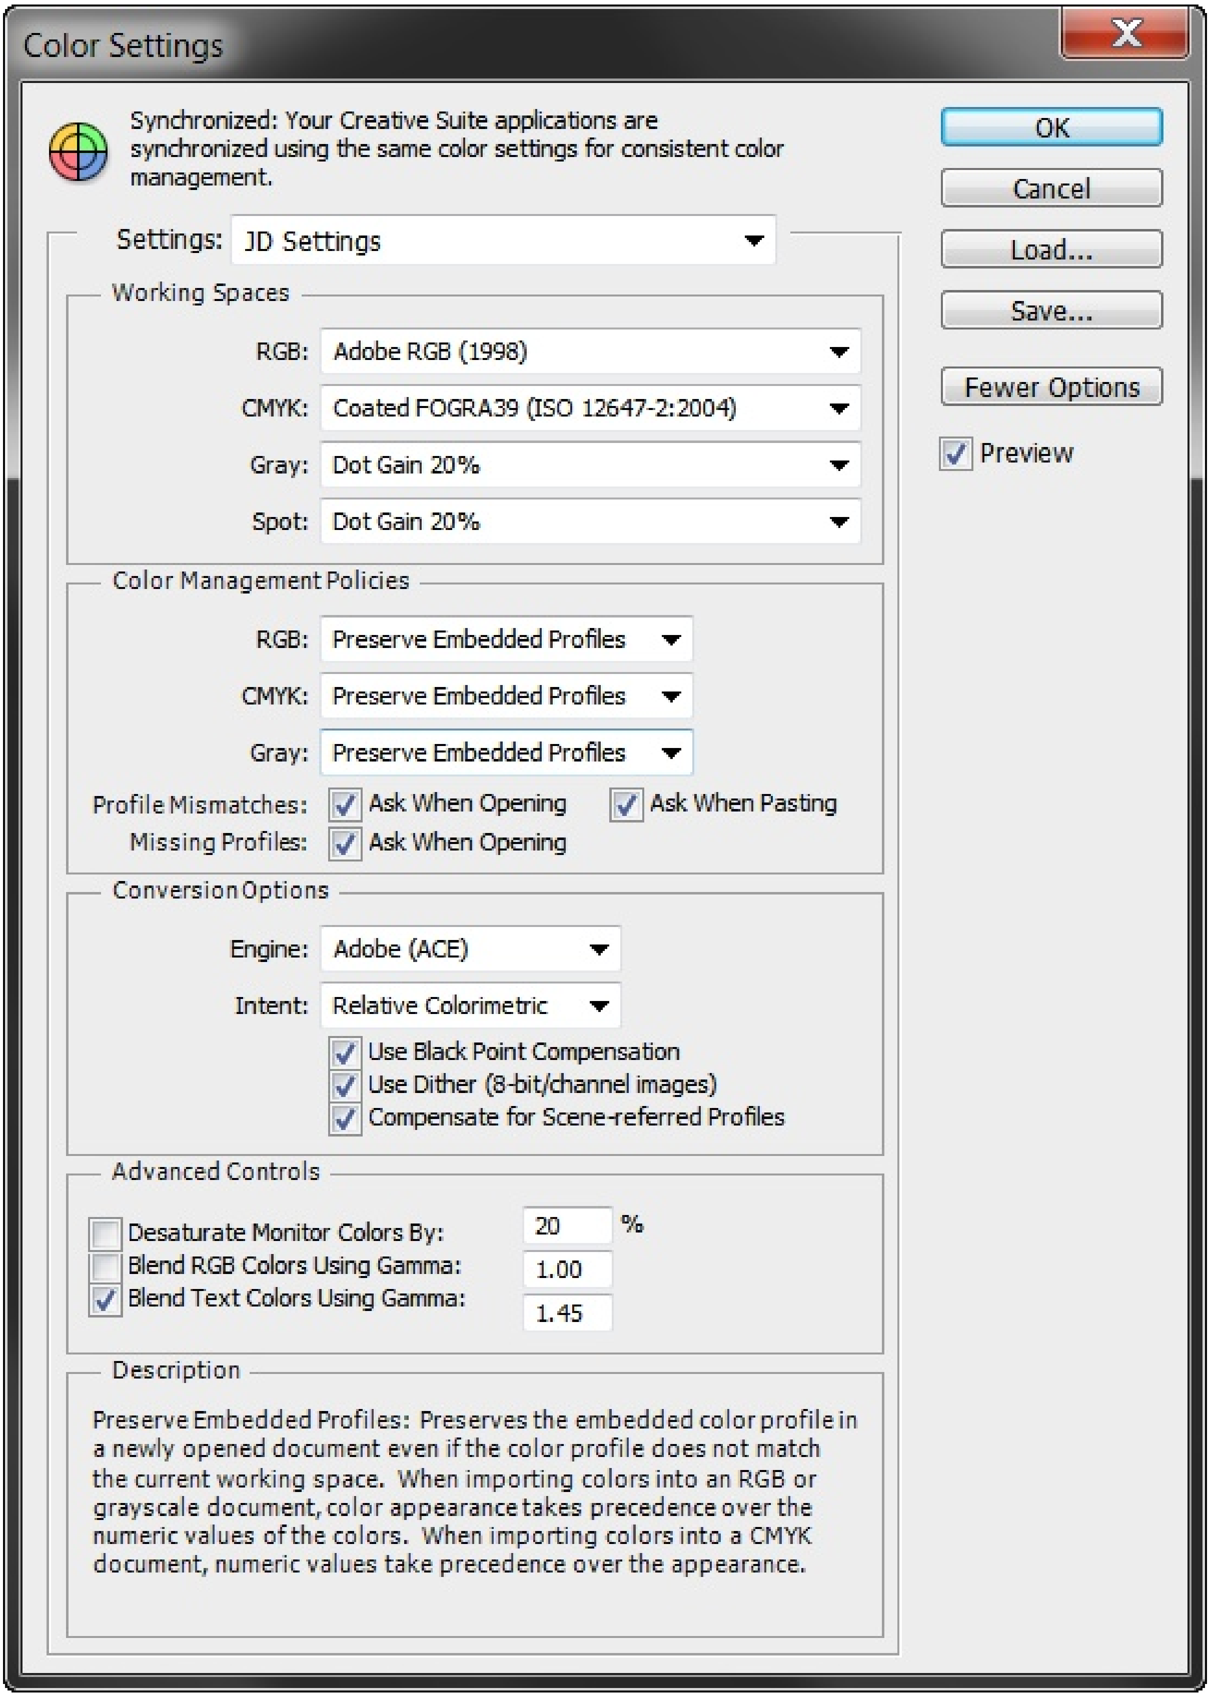

Photoshop Colour Settings (CS6)

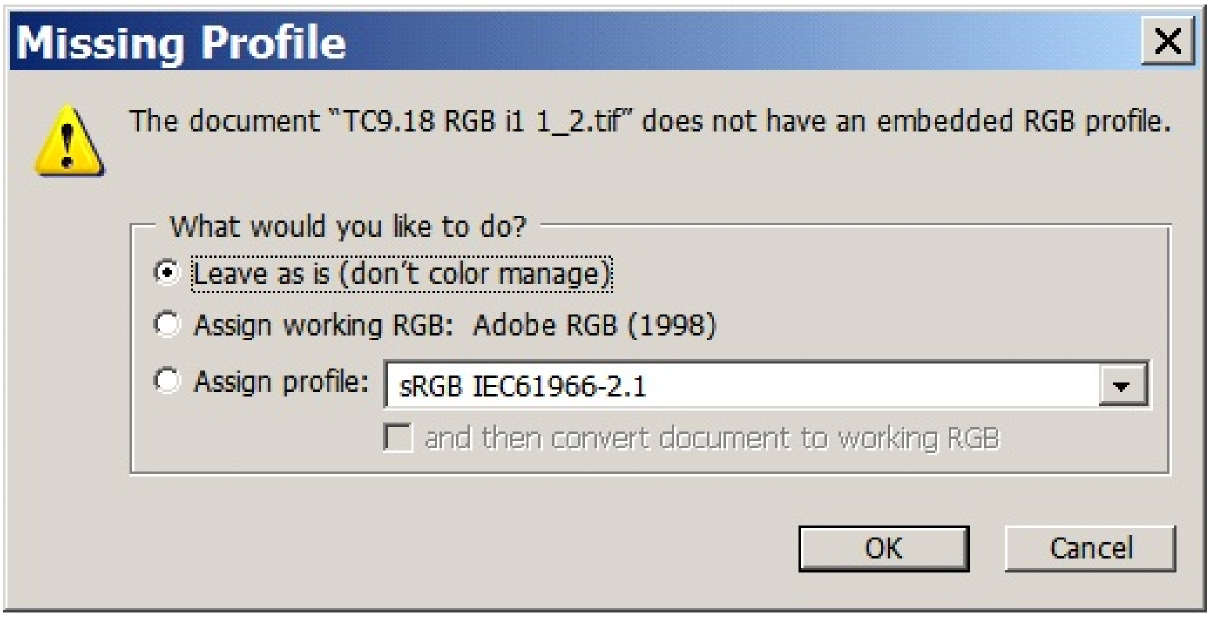

No Embedded Profile Warning

Danger Will Robinson! Untagged Files!

Untagged files are like the ravings of a lunatic. Sure, some of the words are familiar and we now what they mean but overall there's no meaning to the words. There's no message – it's just a series of meaningless words.

In colour terms, sure we know that (0,0,0) is black etc, but none of the other numbers have any real meaning in absolute colour terms. We can only guess. And in managing colour, guessing is always bad.

If you open an untagged file, you will see the following screen:

You have several options:

1. **Do Not Colour Manage**: Photoshop will send the colour numbers directly, unaltered, to your output devices. In general, you only ever use this option when printing out profiling targets for printer profile.
2. **Assign Working RGB**: Photoshop will tag the file with your working space and open the file. This is like saying I'm an English speaker, and I assume this file is in English. If the file really is in English, great. If it's in German – uh oh! You are making an assumption about the colour numbers in a file and telling Photoshop to work with that assumption. If you get that assumption wrong, you have altered every single colour in your file
3. **Assign Profile**: of your choice. Like option 2, you are making an assumption about the numbers in the files, and tagging the file with that assumption. If you get that assumption wrong, you're in trouble. If you get it right, all is well.

So what option should you choose? Well, it depends on how you got the file in the first place. If YOU generated this file, then you probably made a mistake further back in setting up your software (eg. scanner driver or raw converter). See below for more info, but you should definitely be generating input files that are tagged with a profile.

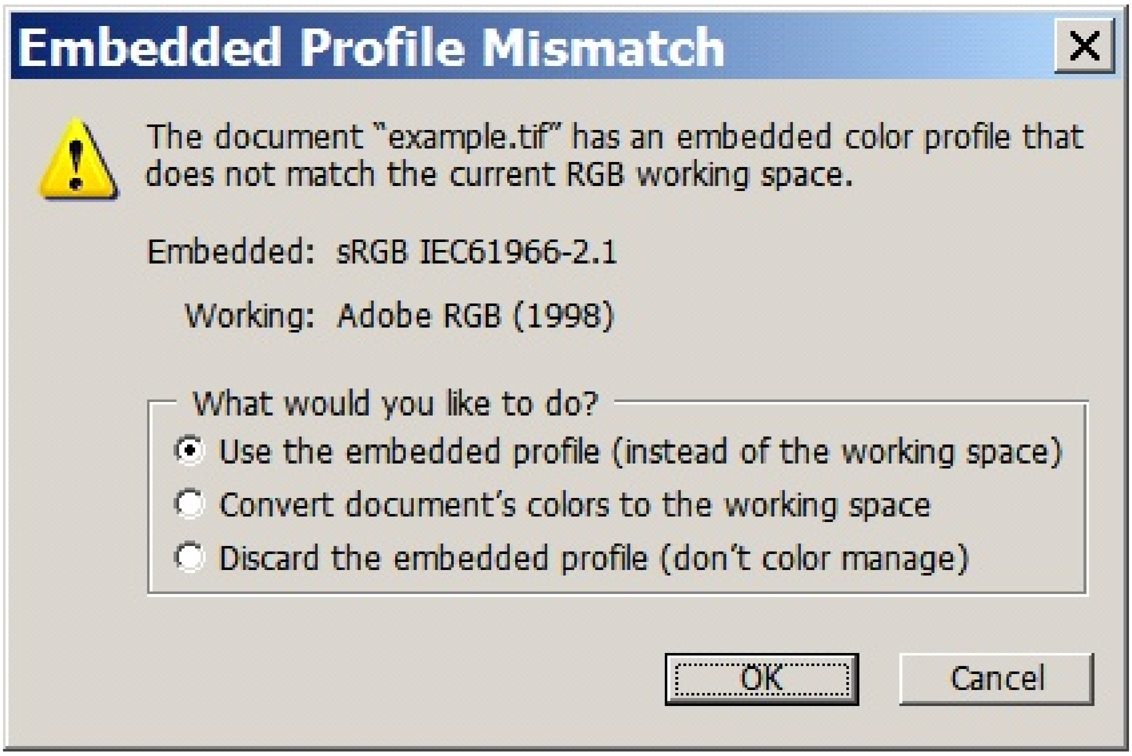

Profile Mis-Match Warning

If someone else generated the file, and it isn't something special like a profiling target, then they made the mistake and it's a clear sign that the person/lab you are working with does not understand digital colour management. They're giving you colour but no means by which to understand that colour. It makes it next to impossible for you to get accurate results, so run (don't walk) to a different service.

Still, you've got the file, so what do you do? I typically tell Photoshop to open the file without colour management first. I then get to see what the file looks like when the colour numbers are sent straight to my screen. I then use the Edit->Assign Profile command to assign a variety of common profiles to the file and make a visual comparison of the results. Often, it will be obvious which 'colour language' a file is in as I choose the profile – the image's colours will fall into place and look right. I can of course choose anything that is pleasing, as it's entirely up to me. I always start with sRGB and then AdobeRGB as these are most likely to be the spaces the file was created in, even if it wasn't tagged. All sorts of cheap devices create sRGB files and then don't tag them, so it's a good place to start.

Danger Will Robinson! This file is not in my working space!

When you open a file that is tagged by a colour space that is NOT your default working space, Photoshop will warn you and you will have the opportunity to make a choice from three options.

1. **Use the embedded profile** – Photoshop uses the embedded profile to translate the colour numbers and you work exactly as if you always work in the embedded profile. In general, you should use this option if you are working on a file that arrives in a working space other than the one you typically work in (eg., you use AdobeRGB and the file arrives in sRGB or ProPhoto). The colours already have meaning and if they have been clipped or whatever, that has already

happened. There is little point translating the file to your normal working space unless you can think of a specific reason to do so.

2. **Convert the document's colours to the working space** – the colour numbers are translated from whatever space they are in to the nearest equivalent numbers in the working space. If your file arrives with a device profile attached, this is the best option to use. E.g. I open a file that is tagged with 'Imacon Flextight 949 scanner profile' – I know I don't want to edit in the scanner's colour space, so I convert it into my working space at this stage.
3. **Discard the embedded profile** – you'll pretty much never want to use this, as it is going backwards really, you're throwing away the dictionary that defines the meaning of the colour numbers.

Understanding Convert Vs. Assign

At any time, not just when opening files, you can convert between profiles, or simply assign profiles to an image. It's important to understand the difference between the two so you know which to use when necessary.

As usual it's all about the numbers in the files.

Assign

(or 'leave the numbers alone, but change the dictionary')

Assigning a profile leaves the RGB numbers exactly the same, but changes the dictionary for those numbers (assign). So, the end result will be a file with all the same RGB numbers in it, but with the definitions of those numbers changed into different absolute colours.

For example:

```
A pixel in an sRGB file(100, 100, 100)
Translated by the sRGB profile to LAB (40,0,0)
```

We then `Edit->Assign Profile->AdobeRGG`

```
The pixel is still RGB (100, 100, 100)
Now translates to LAB (50,0,0) - a different absolute colour
```

So the colour *numbers* have not changed, but the colour *meanings* have.

Convert

(or 'change all the numbers to the equivalent new ones')

Converting to a profile creates an entirely new image, taking all the RGB colour numbers from the original image and changing them to the matching RGB numbers in the new working space (using whatever rendering intent you choose to control what happens to out of gamut colours).

For example: A pixel in an sRGB file (100, 100, 100) Translated by the sRGB profile to LAB (40,0,0)

We then `Edit->Convert Profile->AdobeRGB`

```
The pixel is still RGB (90, 90, 90)
Now translates to LAB (40,0,0) - the same absolute colour
```

So the colour *numbers* have all changed, but the colour *meanings* have not!

Working Spaces - RGB

The colour management world is full of debate as to which colour spaces are best to use, and when. Zealotry is common, and this leads to a lot of conflicting information.

There are several ways to look at the issue. You could, for example, take the attitude that you always want to keep as much colour as possible in your original file, even if you can't see or print that colour - yet. Maybe soon there will be a device created with a significantly wider gamut. Or you could take the attitude that you're never going to achieve that colour (or be able to see it in your monitor) so why concern yourself with it? Both arguments have their merit of course.

You could also look at it from the practical perspective – all professional imaging agencies in Australia currently specify AdobeRGB as their preferred colour space (as defined in the file recommendations for the AIPP, ACMP, 3DAP etc). So maybe we should always use AdobeRGB.

The debate about which working space to use is long, fierce, and very often pointless. It's very easy to use whichever you like as necessary, and if you understand the meaning of your choice it is easy to make the right choice for each situation as you come across it.

We'll first talk about some of the most popular working space, and then look at when you might want to use them.

sRGB

sRGB was created to represent the typical gamut of a cheap PC monitor (so the legend goes). As a general purpose working space it is simply too small for high quality continuous tone work as it clips so many tones we find in, e.g. day to day photography - colours that we can also print. You certainly DON'T WANT TO CAPTURE ALL YOUR IMAGES in sRGB as right at capture you are throwing away LOTS of tones we see, can capture, can visualise on screen, and can print. So make sure you set your camera to AdobeRGB or bigger (if not shooting RAW). If you are using RAW, you can then of course choose your colour space later in the RAW converter.

When you DO want to use sRGB – the web/email/guessing at untagged files

You DO want to use sRGB for all images you are uploading for display web. This is because all browsers are ignorant of colour management, and just send the colour numbers straight to screen. And sRGB numbers are most likely (of all popular working spaces) to be closest to the numbers monitors actually need to produce the correct colours. There is no way you can make your images appear accurate on everyone else's screen – this is a fact of life as it is currently and is unlikely to change for a good long time (until monitor's are accurately self calibrating by default).

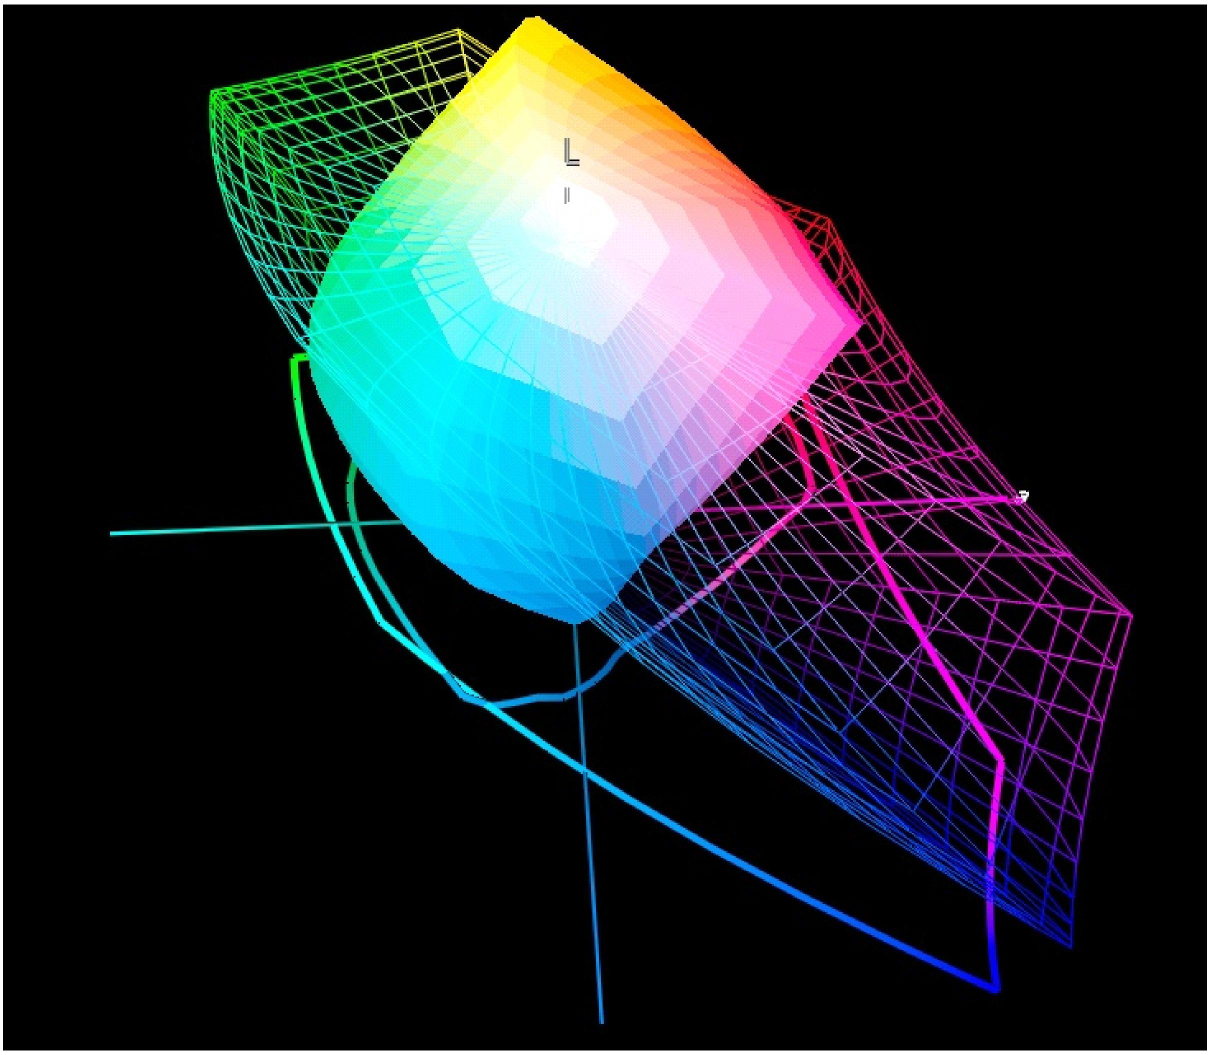

sRGB is really small

You MAY want to use sRGB to send files via email to people that you're pretty sure are not using colour management, as likely most of their applications etc. will assume images are in sRGB, and default settings are typically sRGB.

If you are sending images (e.g. some portrait images) out to a client by email for approval etc, then its probably best to convert them to sRGB first (and of course warn them that the colours will not be a perfect match since their system is likely to be a) crap and b) uncalibrated).

If you receive an untagged file (yikes!) then you can take an educated punt that most of the time sRGB will be the closest thing to a dictionary for those numbers – try assigning sRGB and see where that leaves you.

Above sRGB is pictured against the gamut of a relatively low gamut inkjet printer (the Epson 7800 on a matte paper). You can see that the printer's gamut in blues and greens and yellows extends well beyond the colours that are in sRGB's gamut, which should tell you straight away that sRGB's gamut is generally too limiting for high quality photographic output for colour images.

Adobe RGB 1998

Created by Adobe as a colour space large enough to cope with most general situations and to contain the gamuts of most common output devices (including inkjets, true RGB printers and offset presses), AdobeRGB is a bit of a standard in this country and in the imaging world in general.

Unless you have a good reason not to use this working space (see Raw Conversion later), it should probably be your default choice.

ProPhoto RGB

ProPhoto is a massive colour space, and easily contains the gamut of all current input and output devices. It is so large, it should never be used in 8 bit mode (with only 256 levels, and such a large gamut, the gaps between tones are very large and it is very easy to introduce visible banding).

For a full discussion on the pros and cons of ProPhoto, see the article ProPhoto or ConPhoto?

In general, and for many images, ProPhoto is perhaps a little too large for most uses. But modern devices are increasing in gamut all the time (particularly capture devices) so the argument that AdobeRGB is a little too small is getting stronger all the time. If you're shooting with a high quality modern digital SLR, you may well want to use ProPhoto – but there is a definite way to tell if so (using your raw converter - see the raw conversion section below) - and as a general purpose space, it is needlessly large.

Working Spaces - CMYK

What CMYK space should you use? Well, ideally you won't use one at all!

But, if you do - the best choice is probably going to be FOGRA39. This is a modern ISO standard the most printing organisations are trying to achieve - that is, they are running their presses in such a way as the behaviour of those presses is brought into line with this standard. Thus if you must specify actually ink-on-page values, then this is probably the best CMYK space to choose - unless of course they can direct you to another print standard.

We would still strongly advocate initially working in RGB and only converting to FOGRA39 at the end of your workflow, just before printing. Once you have converted to FOGRA39, re-edit your image to maximise the output quality under these conditions and save this file as a specific print ready file separate to your master file.

Raw Conversion

More Notes on Choosing a Working Space

Choosing a working space isn't that hard, it depends on the task at hand – there's no need to be deeply committed to any one colour space. In principle, we should choose colour spaces based on the image itself – that is, what colour space will best hold the capture image data as we move from our capture medium (digital or film scan) to working space.

The **ideal colour space** for editing any particular image is the smallest colour space which does not clip any tones, i.e. is just large enough to contain the most saturated tones captured in a scene, but is no larger. Going right back to the beginning, you will recall that all images of a given bit depth have the same total number of tones – 8 bit images have 1-255 for each colour, 16 bit images have 1-65536 for each colour. What those colours actually mean is defined by the colour space (by defining the colourants, we move the end points, or how saturated the colours are, but we don't actually increase the total number of colours available to us – we just change the size of the gaps between those colours). Thus the best space to use is one that is just bigger than needed to hold our image's colour – that way we will not be wasting any room in our model of colour on colours not present in the image, and will therefore obtain the smallest gaps between tones – and thus the best smoothness in our images.

So, as a general rule, we want to use the smallest working colour space possible, only moving to a bigger space if the actual image colours as captured will be clipped by that colour space.

In practice, choosing a colour space usually comes down to how you are generating/receiving your images.

If you are scanning your own film

If you are scanning your own images, you should set up your scanning software to use the colour space you think is the best fit for film scans. Most people scan into AdobeRGB, but you may want to investigate using Ektaspace by Joseph Holmes (designed to closely match the gamut of E6 film) or ProPhoto (if you want to be absolutely certain you do not clip any colours during the scan). In general, any of these three will provide good results. Your biggest issue is likely to be the fact that cheaper desktop scanners are very poor at tone placement and differentiation, resulting in loss of local contrast (which leads to flat, lifeless prints),

If you are using a premium scanning service

If you are using a premium scanning service, like Image Science's service, you will probably find they have made a default decision about colour spaces (we use AdobeRGB as it is the standard with all major industry bodies). However you can usually request the service scan into any colour space of your choice, or produce scans which are not yet converted but tagged with the scanner's colour space, which leaves you free to do the conversion yourself.

Far more important will be the ability of the scanner operator to extract maximum quality from their machine, and their use of a superior quality machine in the first place. Better scanners aren't just sharper or better at seeing into shadows, they are also much better at the job of tonal separation. This is one of the reasons why scans off a cheap flatbed scanner result in flat, lifeless prints – they tend to compress and bunch up tonality, so local contrast is lost.

If you are shooting digitally

If you shoot RAW, you can choose your colour space in the RAW converter, which is ideal. You can even re-process the same image into multiple colour spaces for multiple uses – e.g. you could process one version of the file into sRGB for sending to an art director via email for approval, and another into AdobeRGB or ProPhoto RGB as your high quality master file to make prints from.

If you shoot jpg/tiff rather than RAW, you must specify which colour space to use in camera, and I would seriously advise you choose AdobeRGB or bigger or you will be severely limiting your gamut right at the time of capture.

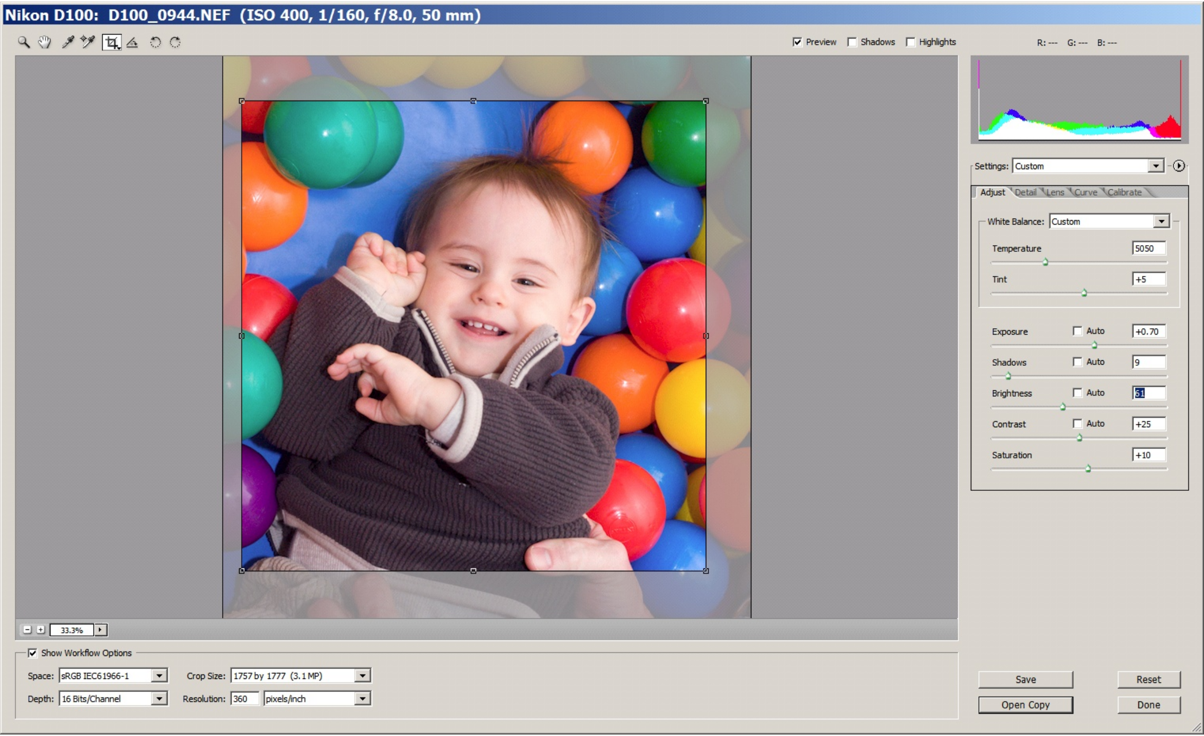

Choosing Colour Space with RAW images

In the section on sRGB above, I said you should use sRGB whenever you are dealing with the web (e.g. putting up your website), or emailing to clients (best to assume they will use a dumb image viewer that send the raw sRGB numbers straight to the monitor), and as the first possibility when attempting to assign meaning to untagged image files. Those are pretty much the only times you want to be using sRGB, but it does have its place, e.g. if you're shooting *only* for putting images on the web you might as well use sRGB and save yourself some time, but this is generally not a common situation.

When shooting RAW, you can definitively choose the best colour space for the job at hand, so let's move on to the next section...

Choosing Working Spaces When Shooting RAW

Your RAW converter can definitively tell you which is the best colour space to use for each specific image.

RAW is an amazing development in photography and anyone using a digital camera and not shooting RAW is missing a huge opportunity to improve the quality of their work by exercising greater control at a more appropriate time in the image making process. When we shoot RAW, the camera records ALL of the sensor's information and we can make decisions, out our leisure, as to what to do with that information once we're away from the frenetic action of the shooting situation.

One of the key things we can do is analyse the colour we have recorded and how that colour will behave as we convert from the camera's colour space into the working space – the RAW converter gives us an extremely powerful tool to help us with this process – the histogram.

Looking at the image above we can see the histogram for this image if we convert the image from the camera's colour space into sRGB. The white histogram indicates there is some exposure clipping (some pixels are set to `0,0,0` maximum black – the deep shadows under the balls) – this is of no

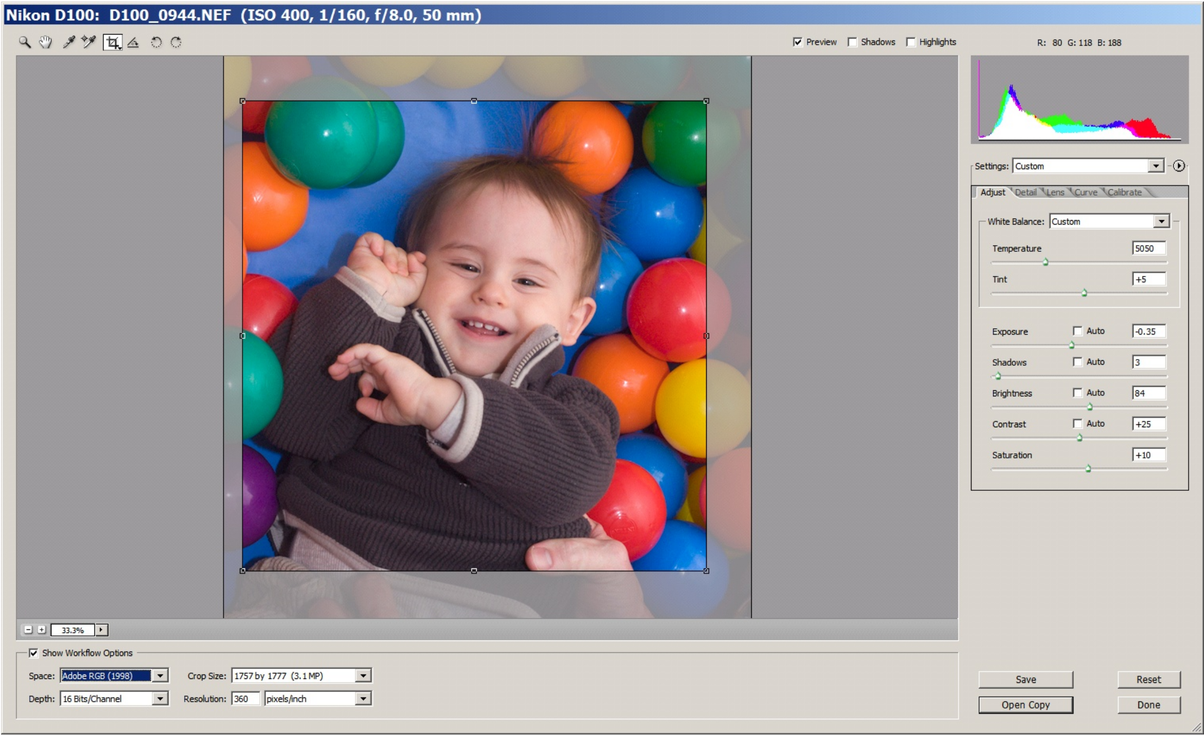

Choosing Colour Space with RAW images

great concern if we're not looking for detail in the deep shadows.

In general, clipping at the shadow end is far less of a worry than clipping at the highlight end (when it comes to overall luminosity) and clipping to black is far less of a worry than saturation clipping (when it comes to colour), so be more concerned about problems on the right side of the histogram than the left. Ultimately, this is because we want some things (deep shadows) in our image to be true black, but we don't really want to break rule number 1 and allow paper white into the print. We also don't want to have large areas of undifferentiated, highly saturated colour in our print.

So, the big red spike on the right hand side of the histogram IS a problem. It indicates that we'll be clipping this images tonality in a big way if we convert it into sRGB. The red balls are simply too saturated to be mapped into sRGB and the result will be large amounts of pixels all having 255 as their red value.

In practical terms, if we convert this image into sRGB, we will forever have an ugly band of undifferentiated tonality on the bright red ball to the right of the child's head. This is about as bad as image problems get – bright red attracts the eye like a moth to a flame, and the first thing the eye will see as it moves to this red is easily visible banding.

So the raw converter is very explicitly telling us that we will be clipping tonality if we convert into this colour space – and (without even yet considering the printer's gamut and the problems it may cause) the net result will be an image converted into the working space with irreparable colour damage.

Looking at the same image being processed into the wider gamut colour space Adobe RGB 1998, we can see that the red clipping has now disappeared completely (we've also got rid of the exposure clipping by decreasing the shadows value). The net result is, even on the most saturated balls, we have now have tonal separation. Where before we had large amount of pixels clipped to 255 red, these are now at values more like 230 – 240 in AdobeRGB. We have better preserved our tonal separation.

In a nutshell, the best colour space to use is the smallest one capable of holding all of the capture image colours. It is image dependent, and you have the tools to evaluate this. Use them!

Using Output Profiles

We have a comprehensive practical guide to using output profiles on our website. Please read those notes and then return here. It covers installation, testing, evaluation and the all critical soft proofing process in depth. (Soft proofing is simulating your printed output on screen before making your actual print and relies very heavily on accurately calibrated equipment!).

A few further points on Advanced Soft Proofing:

There is no white for your eye to white balance against except for the simulated paper white, and no black reference either, other than Photoshop's simulation of ink black. While this proof is technically the most accurate – spectacularly so when the print evaluation environment is correct - many people simply find the screen rendering does not give them a true feeling of how their image will truly look and feel in practice. The elevated black point with matte papers in particular cause many people concern – will my print really look that flat? The answer is yes – under very bright exhibition style lighting, fine art matte photo papers will look very flat, the blacks very weak. But in normal display conditions this is not necessarily the case (a typical home or office is lit to about 500 lux, exhibition lighting is more like 2000 lux typically, so much much brighter!).

While the proof is technically accurate, it may not actually be that indicative of the reality of viewing the print in practice – which brings us full circle, and reminds us once again that colour/dynamic range, is all about the light source and the observer. What's 'real' in some abstract, mathematical sense, may not have any real bearing on what is perceived.

So remember, while the tools are accurate and extremely useful, they are only part of the story, and will not solve every problem you run into!

Issues & Alternatives

Colour Management is by no means perfect, and won't solve all problems. Here are some remaining issues/problem areas.

Early vs Late Binding

There is much debate over Early vs. Late Binding (with RGB vs. CMYK being part of this issue!).

Early binding means working in your output conditions as early as possible. This can be effective is you really know precisely what those conditions are going to be - for example, you're aiming for newsprint, so there's not much point preparing a gorgeous master file for fine art reproduction - you'd perhaps be better off preparing your image with newsprint (and all its limitations) in mind right from the beginning. If you are 100% sure of your output, this can definitely speed your workflow - but it's very limiting and can mean, should your output ever change, you have to repeat all your work. In general, first creating a good master file, and only then targeting print - known as Late Binding - is the more flexible approach and usually ends up saving you time in the long run.

Certainly - if you are an artist wanting to multi-purpose your output - it is MUCH more sensible to

Late Bind. That is, work is a sensible RGB working space and save that as your master file - which is what we've been discussing all this time!

Gamut Mapping is theoretical and NOT based on actual image tones!

Colour management works through the translation of colours in one space (ultimately LAB values) being translated into another space (also LAB values).

One important thing to notice about this process is that it is operating at the gamut level, NOT the actual image level. That is, one entire colour space is mapped into another entire colour space - with no account to which colours within that space are actually being used by a particular image! So this automatic mapping may not result in an optimal translation of colours.

This is a fundamental flaw in the ICC profile system, and can and does cause practical issues at times. But mostly this gamut level mapping works really quite well with continuous tone imagery.

Spot Colours & Optimisation

The corollary to the above is that sometimes ICC profiles are sub-optimal for spot colour reproduction.

That is - for any one particular tone, this automatic gamut to gamut mapping may not produce a wonderfully good match at the output end. And when it comes to perceiving very slight differences in single tones, the human eye is almost unmatched - it is uncannily and frustratingly good at it!

This makes spot colour work, perhaps counter-intuitively, harder than continuous tone work. Small colour error gets 'absorbed' into continuous tone images and accepted by the eye as overall having a very good match. But any small difference on spot tones is immediately obvious.

There are at least four methods - often used together - that can help:

1. Buy a spot ink and print with it - this is obviously VERY effective but also very often not achievable!
2. Optimise your output profiles for particular tones. Modern profiling software allows you to measure an input tone (or group of tones) and say 'this is the most important tone to me - please prioritise reproduction of this tone - and when the software builds the profile (specifically the Perceptual intent) - it will optimise the reproduction of those tones as the possible expense of others. This is a handy technique, although the generated profiles are really only useful for those particular circumstances, so one should always make a general profile before doing this.
3. Automated ring-arounds - a very fast and effective technique is to generate a ring around image - that is, measure the LAB values of the tone you want, then generate an image with a bunch of nearby tones
4. RIPs often have their own variant methods on the above with similar approaches and these can often be really quick and helpful in solving spot tone issues.

Colour Naming Systems (Pantones etc.)

One alternative, old-school approach to all of this is to use ‘named colours’. That is Pantone colours or similar.

And Pantones are great - IF you’re actually going to use a Pantone ink mix at the end to do your printing. In this context, they make perfect sense. Specifying your colours in Pantone Plus Solid (formerly PMS) colours is a quick and easy way to get everyone on the same page.

However - in reality most Pantones are in practise reproduced using standard CMYK inks (aka Process colours). In which case, really all you are getting from Pantone is an expensive (and confusing) naming system - all the process colour libraries are specified as LAB values in practise - when you pick them in Illustrator or whatever, all it does is look up the Pantone <> LAB charts and go from there. It’s no different from any other swatch.

The fact remains, though, they are an easy language for colours and if you’re communicating remotely, having a named colour system can be helpful. But be aware that getting Pantones out of a printer is just like anything else - to do it right, you need good colour management and process control.

Equipment and Environment

Wow! So that was a long and thorough journey through the theory and practise of colour management.

By now, you should have a sound working knowledge of the fundamentals of digital imaging - and how to get files from one end of your workflow to the other with accuracy and control.

Now - the fun really starts! Actually applying the knowledge to your work and seeing improvements.

It's worth finishing with a discussion on setting up your working environment properly - as by now you realise that perception of colour is a complex thing and of course the environment in which you're seeing and working on colour can make a substantial difference.

Also, you simply can't produce great craft with crappy, unreliable tools, so at the end we point you in the direction of specific equipment recommendations. For a relatively modest investment you can set up a truly world class digital darkroom (aka Lightroom!).

Introduction

The state of your working environment - essentially made up by the colour of your walls, and the colour and brightness of your lighting, has a big effect on your ability (or lack of ability) to control colour accurately. It's important to minimise colour distractions in your field of vision, so that you can really concentrate just on the image you are working on and give your eyes the opportunity to accurately perceive colour.

Even wearing a very brightly coloured shirt can cause minor issues as this will reflect in your screen and alter how you see the screen colours!

As with all things, you will need to find your own balance between your wish to achieve accuracy and the lengths and expense you're will to go to achieve this accuracy. For most people a happy balance it not too hard, or too expensive, to find. Below are some tips on both ideal scenarios and scenarios that are more realistic for most of us!

Your Work Room

Ideally your environment should be as neutral as possible. The ideal is a room painted neutral grey (a neutral grey with a defined spectral output across the visible spectrum being better yet).

Of course not everyone wants to paint their work room neutral grey. It's worth being aware, though, that strongly coloured walls can have a fairly big effect on how you will perceive your screen and print colours. For example, if you have warm yellow walls, it is likely this will cause you to create images with a bias towards cool as you will, subconsciously, compensate for this overly warm light.

To overcome this, rather than paint your walls, you could alternatively consider using equipment like a monitor hood and a print viewing box to control the light more directly in the relevant areas of your work room.

If you do want to paint your walls, here are some popular paint companies in Australia and their formulas for neutral grey - please note these figures are just a good starting point and you should test small quantities yourself (preferably using a spectrophotometer) before committing to any substantial order:

Dulux

Use the following formulas for one litre (and choose a low gloss level paint):

```
For Light Grey (CIE Lab approximately 80)
Base: Vivid White
Tinter Formula (in shots, 1/64)
    M (Black) 9
    G (Red Oxide) 2
    EE (Yellow Oxide) 2

For Mid Grey (CIE Lab approximately 65)
Base: Vivid White
Tinter Formula (in shots, 1/64)
    M (Black) 48
    G (Red Oxide) 7
    EE (Yellow Oxide) 12
```

Taubmans

Use the following formulas for one litre (and choose a low gloss level paint):

```
For Light Grey (CIE Lab approximately 80)
Base: White
Tinter Formula (in shots, 1/64)
    B (Black) 7
    F (Red Oxide) 0.5
    C (Yellow Oxide) 1.5

For Mid Grey (CIE Lab approximately 65)
Base: Light
Tinter Formula (in shots, 1/64)
    B (Black) 27
    F (Red Oxide) 1
    C (Yellow Oxide) 5
```

Wattyl

Use the following formulas for one litre (and choose a low gloss level paint):

```
For Light Grey (CIE Lab approximately 80)
Base: WHT
Tinter Formula (in shots, 1/48)
    B (Black) 7.5
    F (Red Oxide) 0.5
    C (Yellow Oxide) 1
```

```
For Mid Grey (CIE Lab approximately 65)
Base: LTB
Tinter Formula (in shots, 1/64)
    B (Black) 19
    F (Red Oxide) 0.75
    C (Yellow Oxide) 1.75
```

Lighting

Perhaps the easiest and most important change you can make is to change the colour and brightness of the lighting in your work area. You can either change all the lighting in your work area using your existing fittings, or buy a print viewing box, or indeed build your own print viewing box or area.

Most lighting in most homes is a long way off the ideal light for digital imaging work, but there are often inexpensive and very effective solutions available that can drastically improve the situation. You may well be quite stunned at how good well made prints can look under really good lighting!

First, the colour of your lights

First, remove any strongly coloured lighting – that means absolutely no normal fluros (the daylight balanced tri-phosphor types are better but still not wonderful). If you have tungsten lights, try swapping them for 'cool white' versions balanced to 5000K, and the same goes for 'down lights'. The idea is to make sure the lights are not adversely affecting your perception of your screens colours by outputting light of fundamentally the wrong colour (i.e. far too warm or too cool).

5000 Kelvin is chosen as the ideal colour temperature for print viewing as it is a happy average between indoor (typically 3000 to 4000K) and outdoor lighting - (can be anything from maybe 4000 to 9000K depending on the weather and the time of day).

5000 Kelvin is representative of a mix of indoor and outdoor light, or the light typically found in most homes and offices during a large part of the day - the sorts of places where people will actually be looking at your prints! At night, when daylight is no longer part of the light mix, the light will typically be warmer than 5000K indoors, but the eye is quite tolerant of shifts to warmth, and this is not typically a problem.

If you want to go a step further, then look specifically for lights that have a high CRI value - CRI stands for colour rendering index. The idea is to get a light source that is as close as possible to the D50 lighting standard (5000 Kelvin with a defined spectral output curve). A CRI of 100 is perfect. High quality lights can achieve a CRI of 98 or 99.

If your work or print display area has down-lights (i.e. halogen lights) It is generally recognised that Solux make the best bulbs (of any type) that are currently available. These are the bulbs that serious museums use - places like the Guggenheim and the Van Gogh museum. They can be purchased from the Image Science colour accurate lighting page. Generally the 4700K black back bulbs are regarded as having the best match to D50, but there is quite a lot of information on their website that you may find useful. The CRI of these bulbs is over 99.

If your work area has fluorescent lighting, then we sell GTI colour accurate tubes that are very good (again see our colour accurate lighting page. Whilst still flourescent and thus not *quite* SoLux quality, they're very often used in commercial proofing environments and a very good option.

I am not currently aware of any lights for standard fittings that are particularly good. This does not necessarily mean they don't exist, just that I am not aware of them. If you have information on any this, I'd be very happy to hear about it.

Second, the brightness of your lights (around your monitor and for viewing prints)

Secondly, the level of lighting should be considered. Colour management systems for screens are based around your screen being the brightest element in your field of vision. This means the ambient lighting level in your room (and specifically around your monitor) should be below that of your screen (generally around 30 to 50% below). This is one big advantage of LCDs – they easily go much brighter than CRTs and therefore you can work in a more comfortably lit room.

However, when viewing prints, there is no one standard for brightness. Generally, colour management systems are built around scenarios where prints are viewed under strong, bright lighting (D50 at 2000 lux). Print viewing boxes typically have dimmer controls so that they can simulate different lighting levels. However, lighting levels within buildings (and of course outside!) can vary dramatically. A typical room in Australia during the daytime is about 500 lux in brightness, but many with large windows will be significantly brighter than this (as much as 2000 lux), while at night under artificial light the level may be as low as 200 to 300 lux.

How do you decide? Well, it depends on the context you're aiming for - 500 lux is probably a nice average for most scenarios, but exhibitions and galleries (and photographic competition judgings!) are generally significantly above this level. Obviously you can't control the final environment your prints will end up in (unless it is your own home), so the best approach is to print for a reasonable average and accept that this is fundamentally out of your control. A very effective print viewing area for not too much money...

Print viewing boxes are very nice, and very professional looking, but certainly expensive (you won't get much change on $1000 for a small, basic unit with many being a LOT more than this). You can, however, build your own print viewing box/area for a lot less money quite easily, or invest in a simple but effective Grafilite system for $100 to $200. See buying guide below.

If you choose to build a print viewing box, then build a box with white or neutral grey walls (from something like MDF, and using the paint formulas above, perhaps). Then install in this box some 12 volt, 35 watt halogen fixtures (available from any lighting store or even Ikea!). Add Solux bulbs and you're done. Total expense should come in under $300.

Or, for something a little more attractive and convenient than a box, you could perhaps paint a feature wall in your home in light grey (formula above) and install a track based lighting system, again with 12 volt and either 35 or 50 watt halogen fixtures. Add Solux bulbs, and a gallery hanging rail system, and you're got yourself a gallery quality print display area (better in fact than 99% of galleries in this country!).

Remember, all electricity work involving 240 volt power must be performed by a licensed electrician!

A note on your computer's desktop...

One of the cheapest and most effective things you can do to improve your colour perception is to get rid of background images, gradients etc. on your computer's desktop. Just set your desktop to black or a neutral grey (I wouldn't recommend white, it is simply too tiring on the eyes).

Looking at anything colourful around your Photoshop window, or immediately before you start using Photoshop, will have a big effect on your ability to see colour accurately.

Equipment Choices - Monitors & Printers etc.

Because equipment changes so frequently, we recommend you refer to the following articles on our website for up to date recommendations on actual equipment selection:

- Monitors for Imaging Work - general tenchnology & what to look for guide
- Monitor Recommendations - specific recommendations of current models
- Monitor Calibrators
- System Calibrators (i.e. these do monitors AND printers and possibly cameras and scanners too)
- Inkjet Printers
- Pantone Guides
- Colour Accurate Lighting
- Computers for Imaging Work