

Technical Note

Setting a Truelight Profile

Richard Kirk

Document ref.	FL-TL-TN-0416-SettingTLProfile
Creation date	14 July 2008
Last modified	30 November 2010
Revision no.	4.0

Summary

Truelight has been designed to simulate images output to film on a display. To be useful, Truelight must match the current practice in the video and film post production industry. Unfortunately there is no single post-production standard, so Truelight must support a range of ad-hoc standards and methods that have been developed by pioneering post production companies in collaboration with manufacturers of film, scanners, recorders, and colour measuring instruments.

A Truelight profile contains commands that set up Truelight parameters. Truelight profiles are used to set the parameters that rarely change within a project, such as the choice of film recorder or print stock. This gives Truelight users a simple way of selecting all the usual settings in a single operation. It also lets Truelight application developers to restrict the user interface to the most commonly used controls. Truelight profiles can be made and edited using the Truelight viewer utility, or any text editor.

This document does not propose a new post-production standard. This is only a guide for new users who want to set up Truelight to match their current process. Where we quote particular values for parameters, these are values that have worked for us.

For technical details on the Truelight library, see the FilmLight *Truelight Commands* Technical Note. For reference on the standard colour spaces, see FilmLight *Standard Colour Spaces* Technical Note.

Contents

1	Introduction.....	3
2	Cineon RGB to Projected Film colour.....	4
	2.1 Recording to Negative.....	5
	2.2 Cineon Exposure Space.....	6
	2.3 Printing.....	9
	2.4 Lamp Calibration.....	12
	2.5 Input Display Calibration.....	12
3	Matching the Colour on the Display.....	13
	3.1 Flare Correction.....	13
	3.2 Display calibration.....	15
	3.3 White Point Correction.....	16
	3.4 Matching displays.....	19
	3.5 Colour Edits.....	20
4	Working with Truelight.....	23
	4.1 Cube Type Options.....	24
	4.2 Film Colour Space Options.....	27
	4.3 Other Cube Generation Options.....	32
5	Example profiles.....	34
	5.1 Previewing film log images in office lighting.....	34
	5.2 Generating HD from film log images.....	35
	5.3 Doing a 'butterfly' test in a theatre.....	37
	5.4 Matching film images on a light box.....	41
	5.5 Matching one display to another.....	41
	5.6 Using an LCD monitor.....	41
	5.7 Grading in HD.....	42
	5.8 Using an RGB cube on a different display.....	44
	5.9 Rendering DCI XYZ' output.....	45
	5.10 Using black and white film stocks.....	46
	5.11 Going beyond the profile.....	46
6	Debugging.....	48
	6.1 Notes on Comparing Images.....	48
	6.2 Check Truelight is doing something.....	49
	6.3 Use the Truelight Viewer.....	50
	6.4 Check the Calibrations are valid.....	51
	6.5 Check the Profile Commands.....	52
	6.6 Check the Profile Edit Settings.....	53
	6.7 Check the Viewer Truelight Controls.....	54
	6.8 Check Other Controls.....	55
	6.9 Check the Aim Gamma Settings.....	57
	6.10 Check the Grey Patch Values.....	58
	6.11 Check for HD video scaling.....	58
	6.12 Check the Display Calibration.....	59
	6.13 Test Values.....	60

1 Introduction

In the typical film output process, a negative image is output on a film recorder. A positive print is made from the negative. The print is then projected onto a screen in a cinema. Truelight calibrates these processes using a strip of calibration images. These digital images should be output the same way the other images are. This way, we can follow the colours as they go from digital RGB values as sent to the recorder, to negative densities, to positive print densities, then to colours as projected in the cinema, and finally to colours as seen by the human eye. In section 2, we retrace this route, describing the relevant parameters as we meet them.

In section 3, we follow the other branch of the image matching process, from digital RGB values, to colours on the display, then to colours as seen by the human eye. This is easy if the display is being viewed under cinema-like conditions. However, this is not always the case. It is not easy to use a workstation in a room as dark as a cinema. Some lighting can serve as a reference white. If the room is not dark, it is often sensible to turn up the display brightness too. We must match the appearance of the projected film to the display image as viewed in the working environment.

You can work directly in film recorder or display RGB units. However, when working with 3D rendering, it is often more convenient to work in some colour space with uniform intervals in RGB brightness. When exchanging image data with other users, it is often handier to use a standard video colour space, rather than a colour space that is fixed to a particular device. We discuss some of the possibilities in section 4.

Truelight is complex. This complexity allows it to deal with the many ad-hoc working practices we find in the field. It allows us to combine the calibrations we have in different ways to simulate other similar working practices. This is what distinguishes Truelight from the competitors. Section 5 contains some example profiles for different applications. If you are making a new profile, you may find one of these examples provides a good starting point, even if it does not exactly match your situation.

Truelight's complexity gives us great scope for making mistakes. The right thing to do is to follow the flow of the data through Truelight, and find out by experiment where the mistake occurred. There are too many controls in Truelight to expect to get to the right solution by guesswork. However, it is hard to work methodically at a commercial site with people looking over your shoulder and pressuring you for results. If you are in this situation, go to section 6.

The Truelight viewer, if you have it, is a good tool for developing and editing profiles. You can enter commands on the **Profile > Edit** panel. You can see a listing of the current commands in the **Info > Current Profile Commands** with any unused commands greyed out. You can test the transform on particular colours using the **Tools > Colour Patch** utility. You will see the effects of the commands on the current image, or error and warning messages on the panel below. Even if you cannot use the image or the display you want, this can give you useful feedback.

2 Cineon RGB to Projected Film colour

A typical Truelight transform is made from:

- A recorder calibration linking the RGB values sent to the film recorder to the corresponding status M densities on the negative.
- A print calibration linking status M densities on the negative to status A densities on the print.
- A lamp calibration linking status M densities on the negative to CIE $L^*a^*b^*$ on the projected image.

A recorder calibration is specific to the combination of the film recorder and the stock it records on.

A print calibration will be specific to a particular laboratory print process using a particular negative and print stock. However, if we can print our grey patch and ask the laboratory to adjust the exposure so the print has a given density, then most of the variations between laboratories disappear. In theory, there should be separate calibration for each combination of negative and print stock. In practice, the effect of changing the negative stock is small if the recorder was correctly calibrated. For most users, the print calibration is linked to the print stock alone.

The lamp calibration will be specific to the combination of print stock and projector lamp. However, most colour films have dyes with very similar spectral properties, and even monochrome stocks and bleach bypass prints seem to work. Most xenon lamps also give very similar results. Most people use the default internal xenon calibration.

There are some special cases of Truelight calibrations:


- You can load the film recorder with print stock. The Truelight print calibration then links the RGB values sent to the film recorder to the corresponding status A densities on the print.
- There is a special class of recorder calibrations that behave as the combinations of a print and a lamp calibration, linking link the RGB values sent to the film recorder directly to visual CIE $L^*a^*b^*$.

These calibrations are kept in the same folder as the ordinary calibrations. They are distinguished by their internal colour space descriptions. They are also rare – you are unlikely to meet them under normal circumstances.

Finally, Truelight can match two displays. This may be used to convert video data for viewing on a HD monitor. The input display data replaces all the recorder, print, and lamp calibration settings.

2.1 Recording to Negative

We can record out a set of RGB values to a film recorder, and then measure the film on a densitometer. In the simplest case, our transform is just a list of input RGB values and the corresponding output densities. Here is a 3D plot of the densities we get for two different film recorders, looking down the neutral axis...



ArriLaser film recorder

Solitaire film recorder

We used the same test values on each recorder. The ArriLaser film recorder was calibrated to give uniform steps in negative density along the neutral axis. It also has very pure RGB laser primaries, so there is little cross coupling between the channels. The Solitaire recorder has not been calibrated to give uniform density steps. It uses CRT primaries rather than lasers, so there is some overlap between the primary spectra, which reduces the saturation of the colours.

The recorder calibration will probably be named after the recorder device, or the recorder device plus the negative film stock if there is a choice. You can select your recorder calibration using the Truelight command...

```
recorder{ recorder calibration filename }
```

2.2 Cineon Exposure Space

'Cineon' is a Kodak trademark for an image processing package, and an image file format. The image file format has a standard way of representing negative densities in 10-bit data that is been widely used in the film industry. For the rest of this section, we shall use the term 'Cineon' for this colour space.

You may think it very strange to define an image in terms of negative densities. There are many ways of printing an image. You can completely change the appearance of the image by changing the printing stock or the exposure. You can get different densities in the colours on a negative by using a different film recorder, as the Cineon space is only calibrated along the neutral axis. Nevertheless, Cineon space has useful properties for calibration, and it is consistent enough for some people to exchange images.

The most popular Cineon file format packs one pixel's 10-bit RGB values into a 4-byte word. In this note we shall describe Cineon values as 10-bit integers from 0 to 1023.

The graph shows the relationship between Cineon value and negative density for an aim gamma of 1.0.

This graph is a straight line. Each Cineon unit corresponds to a density shift of 0.002. Low Cineon values are light on the negative and dark on the print. Cineon RGB images viewed without colour correction on a monitor look like a washed-out version on the print.

The Kodak formula for the absolute density is...

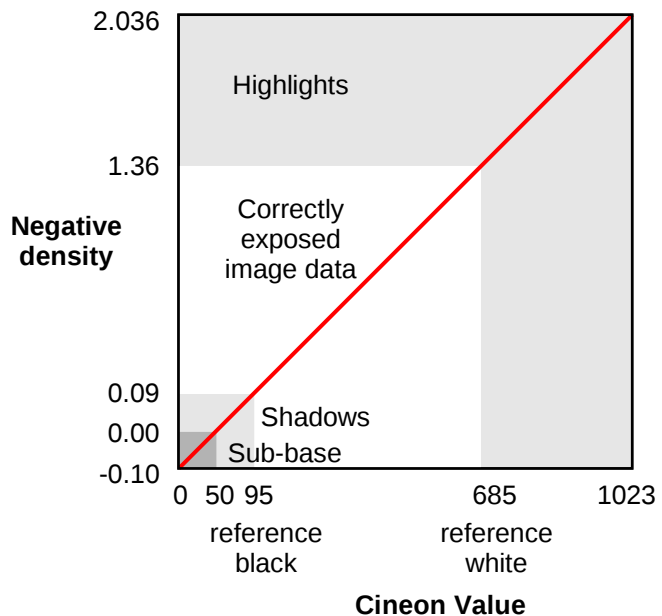
$$D_{absolute} - D_{min} = (value - 95) * 0.002 * aim_gamma$$

Kodak defines absolute RGB 'Dmin' densities for each of their intermediate stocks. In the Kodak documentation this is described as the blackest print black that can be recorded onto that stock. It is approximately the black you get when you take a picture of a 1% black card. They assign the Cineon reference black value 95 to this Dmin density.

The graph shows the film base density as having a base density of 50. If you stick to the Kodak definition, then the actual Cineon base value will depend on the actual choice of film stock. The actual Kodak Cineon value of 5242 base stock is about 42.

In practice, you may not have Dmin values if you are not using Kodak stock, or you may prefer to measure negative densities relative to base. We suggest a nearly equivalent alternative formula...

$$D_{absolute} - D_{base} = (value - 50) * 0.002 * aim_gamma$$



2.2.1 Reference White and Reference Black

Cineon space covers all the useful density range for conventional negatives. A Cineon image usually contains more information than you can see in a single print. You can vary the exposure in a Truelight simulation and see more detail in the highlights or shadows.

There is a subset of the Cineon space, with values between 95 and 685 that is commonly used for *correctly exposed* image data. Kodak documents call 95 the *reference black*, and 685 the *reference white*. Fixed conversions from Cineon to display RGB used to clip values outside this range.

You are not restricted to the *correctly exposed* range. Nevertheless, there are reasons for keeping most of your image within these limits...

The negative exposure curve has its 'toe' beneath the reference black. Down at the toe, accurate interpolation is difficult, and the shape of the tone curve can vary with processing chemistry or the age of the negative film stock.

Most print stocks have a tone curve with the 'shoulder' above the reference white. This print shoulder has all the same problems as the negative toe.

Tube based film recorders limit their output range to the reference white to preserve tube life.

2.2.2 Aim Gammas

If our Cineon step corresponded to a shift in negative density of 0.002 in all three channels, then a set of patches with equal values in all three channels would give a neutral negative. Unfortunately, a neutral negative wedge will produce a print with the darker colours looking distinctly bluish on Vision stock.

Kodak tabulate a set of aim values in their H-387 "Kodak Digital LAD test image" document. If you plot out these values, they lie pretty close to a straight line. They are effectively putting a different scaling on each channel...

Channel:	Red	Green	Blue
Cineon interval:	0.966*0.002	1.063*0.002	1.082*0.002
Equivalent aim gamma:	0.966	1.063	1.082

We can put this scaling into Truelight with the command...

```
aimGamma{ 0.966, 1.063, 1.082 }
```

Truelight will then fit the recorder calibration to the correct gradient along the neutral axis for the Kodak settings. The default is not to fit to aim gammas at all. If you want to explicitly turn off the aim gamma fitting, you can do so with the command...

```
aimGamma{ }
```

In theory, there should be a separate set of aim gammas for each combination of negative and print stocks. In practice, you cannot always find published values, particularly for non-Kodak stocks. You can determine the aim gammas yourself, but users rarely take the trouble. There is not a great difference between print stocks, and if you use very non-standard aim gammas, you will have trouble when you exchange Cineon data with other users.

The ArriLaser Carlos program has a control to set the print range for each channel. This is another form of aim control. The Carlos default settings are :

<i>Channel:</i>	<i>Red</i>	<i>Green</i>	<i>Blue</i>
Density range:	1.89	2.046	2.046
Equivalent aim gamma:	0.9237	1.0	1.0

We can put this scaling into Truelight with the command...

```
aimGamma{ 0.9237, 1.0, 1.0 }
```

If you output the Truelight calibration strip with the same aim gammas, then fitting the aim gammas should have little or no effect. Try removing the **aimGamma{}** line, and see whether the images change significantly. If they do, then your gamma settings are not compatible with your other data.

The aim gamma fitting is useful for removing any calibration errors, so it is worth leaving on even if the ranging should not be necessary. You can also change the aim gamma without having to recalibrate. This can be useful if you have to exchange work with a site that uses different aim gammas.

2.3 Printing

You can select your recorder calibration using the Truelight command...

```
print{ print calibration filename }
```

The print calibration will probably be named after the print stock. There may be longer names for print stocks with unusual processing, such as bleach bypass.

There are several parameters that relate to the exposures used at the laboratory when printing...

2.3.1 Printer Point Settings

When a laboratory makes a print from the negative, they have controls on the printer lamp house that vary the red, green, and blue exposure through the negative. The units on these controls are usually referred to as 'points'. There are 12 points to a stop. If you turn up the point setting, more light goes through the negative, and the positive print gets darker. Increasing the printer setting by one point is roughly equivalent to decreasing the corresponding negative density by 0.025.

The absolute printer point setting is not much use outside the film laboratory. Different labs use different settings, and these settings are often changed to track the changing chemistry in the baths. The labs can do this by printing a standard 'grey slate', and adjusting the printer light settings until the slate hits a certain density. Outside the printer lab, it is common to ask for "two points more red" or other relative shifts in printer points, rather than ask for absolute settings.

The Truelight printer point boost settings simulate these printer point corrections. You can set a printer point shift using the Truelight command...


```
printerPoints{ red, green, blue }
```

In practice, this command is used to compensate for the day-to-day laboratory printing errors when comparing the display to an actual print. It is not generally useful or sensible to put a printer point setting in a profile. Instead, we usually set the exposure using a grey slate. This is described in the next section.

2.3.2 Printing a Grey Slate


If you took the same negative to the same laboratory on different days, you will probably be able to see some colour shift between the two prints. The day-to-day variations in laboratory printing can change the appearance of your images. It is very likely that some error, or change in the bath chemistry will alter your image appearance at some time during a large project. A useful safeguard is to print out a grey slate with every batch of recorded images.

The Kodak LAD test image has a grey patch...



In our experience, this image is not suitable for this work. The grey patch is at the edge of the image. Flat grey values will often print and project as different colours in the middle and at the edge of a print. A portrait test image is can give a useful visual check, but there are several similar Kodak images with different colour balances, so using this one can lead to mistakes. We would recommend making your own grey slate image where the grey patch fills the whole field, and attach a separate portrait image if you want that too.

The grey slate image might look something like this...



You can make patches like this with the Truelight viewer. Open an image and a profile, then select the **View > Enable Grey Slate** option. The text is made from the Truelight parameters, and the image size is fitted to the current image. You can save the image with the **Image > Save As** option. This will prompt you to turn off the profile, so the saved image contains the right grey value.

Our example uses the Cineon values and aim densities from the Kodak H-387 document. Some people use different Cineon values. Some people chose to use an aim density that is the same in each channel, as this avoids possible errors from getting the densities in the wrong order. It does not matter which exact values you use, provided you know what they are and you stick to them. If you do not have any reason to choose other values, then the Kodak values are as good as any, and other people in the industry will probably recognize them.

The Truelight library will fit the calibration to a grey slate. You can set the grey slate Cineon values and aim densities using...

```
greyValue{ red, green, blue }  
greyStatusA{ red, green, blue }
```

If you do not specify any values, the default settings are...

```
greyValue{ 445, 445, 445 }  
greyStatusA{ }
```

The empty **greyStatusA{}** command disables the grey point fitting. You can add it to your profile if you want to explicitly disable the grey point fitting. However the Truelight library omits most default settings when saving a profile, so this line may disappear. The default **greyValue{}** setting is one of the rare exceptions to this rule – its default is a real value, so you may find this line gets inserted in your profile automatically.

You can use the **greyStatusA{}** command with the density values measured from your grey patch. Normally we would correct for everyday deviations from aim densities using the **printerPoints{}** command. This can be a useful option for doing a very precise fit provided the same densitometer, or one very similar has been used for all measurements.

When the Truelight library generates a colour transform, it calculates the effective shift in print exposure needed to fit the aim densities to the aim values. If a printer point shift has been specified then this shift is added after the grey slate values have been fitted.

If the Truelight calibration test strip was printed with the correct aim densities, then fitting the grey slate values should have little or no effect. You can test whether this is the case by turning off the grey slate fitting by removing the **greyStatusA{}** command, and seeing whether the image appearance changes significantly.

We strongly recommend using a grey slate with all calibrations and prints. The only sensible reason for turning off the grey slate fitting would be if your images had been output and printed at the same time as the test strip.

2.4 Lamp Calibration

The projector lamp calibration links Status A densities to the projected print CIE L*a*b* with respect to some fixed white point. This fixed white point is typically the projected film base, though we can produce calibrations to the open gate white, or to something else. The calibration is a product of the film, the projector, and the screen material. Most colour films have similar, smooth spectral characteristics. The optics and screen material should also have smooth spectral characteristics. The lamp calibration is almost independent of these factors, and mostly depends on the projector lamp.

You can select your film projector lamp calibration using the Truelight command...

```
lamp{ lamp calibration filename }
```

The lamp calibration will probably be named after the projector lamp type. Most projectors will probably be reasonable matches to the internal xenon arc lamp calibration.

The calibration does not contain the white point itself. To match a display to a projector, set the display white Yu'v' controls to the projector white point CIE Yu'v' values. Typical projector lamp colours can vary with time, and with lamp house focus. The colour can also change slightly when the projector lens is changed. If you are comparing a projection with a display, you may want to let both the projector and display warm up, then fine-tune the display white settings by eye.

In our experience, a generic xenon lamp spectrum can be fitted closely to specific examples of the lamp by just by changing this white point. There will be some errors in the exact spectral detail, but these are smaller than the variations in spectrum with time, as the arc changes due to electrode erosion.

Print stocks can scatter light slightly. This is called the Callier effect. The effect in color film is subtle. As well as the default **Xenon** calibration, the Truelight library has three lamp calibrations with scatter correction: **XenonCP30** (Agfa), **Xenon3513** (Fuji), and **Xenon2583** (Kodak). In our experience, one of these lamp calibrations will correct for all print stocks made by the same company.

2.5 Input Display Calibration

If you define **inputDisplay{}**, then all your other recorder, printer, film, and lamp settings described in the earlier sections are ignored.

The input display calibration is used to match one display to another. A typical use is to display images rendered for a video monitor on a calibrated computer display. The input display transform is the inverse of a list or a formula display calibration taken from the display calibrations directory. You can select an input display calibration using...

```
inputDisplay{ display calibration filename }
```

3 Matching the Colour on the Display

In the last section we described how we set up our profile parameters to get from device RGB to a display colour in CIE L*a*b* coordinates. In this section, we shall describe the remaining settings to get a visual match on our display.

The light levels from a projector and a display can match, but the images may still not look the same. There may be external physical reasons for this, such as stray light landing on the projector screen. There may be internal physical reasons for this - a bright surround may add light that scatters within the eye, even though no light lands on the screen. There are psychophysical effects in the eye-brain system. There are other effects of experience: we are happy to use tungsten lighting temperatures when looking at reflection copy, but a monitor with a tungsten white point will look very orange.

See the later sections of *FilmLight Standard Colour Spaces Technical Note* for more on these psychophysical effects.

3.1 Flare Correction

We can compensate for flare light to some extent. We know if we add an overall wash of stray light to a display, or we add a bright border, or we turn up the ambient lighting, then we are less able to see the shadow details. We can compensate for this by increasing the slope of the tone curve to boost the contrast of these shadow details.

We have three parameters to the model: the intensity, the adaptivity factor, and the flare factor. We have two sets of these parameters: one for the current display, and one for the reference projection.

The projection settings default to...

```
projectionWhite{ 16 }  
projectionFlare{ 0.01 }  
projectionAdapt{ 0.1 }
```

This corresponds to a 16 foot-lambert display with a 0.16 foot-lambert 1% flare signal and a dark surround. These are values for typical cinema viewing conditions. These are rarely set in a profile.

The display settings default to the same values...

```
displayWhite{ 16 }  
displayFlare{ 0.01 }  
displayAdapt{ 0.1 }
```

When the display values match the projector values, there is no flare correction.

The **displayWhite{}** and **displayAdapt{}** commands have only a subtle effect on the displayed image. Only **displayFlare{}** has a marked effect. If there is significant flare, users must set **displayFlare{}** to match their viewing conditions.

3.1.1 Setting the Flare Correction

Usually the operator will set the display flare setting to match their local conditions, and leave the other controls at their default settings. If they are in an open office, they may have to change their lighting conditions as the sun comes out or goes in.

Telecine grading suites often have D65 lighting on the wall behind their monitor. This provides a suitable reference for judging the absolute white point, and also provides a low level of ambient lighting, which makes working easier. This light will not be illuminating the monitor screen directly, but there will be reflections from the walls, which are typically grey rather than black, and there will be light scattered within the eye. The monitor white is brighter: 30 foot-lamberts rather than our cinema 16 foot-lamberts.

If we are going to have a cinema projector in the telecine suite, then we are going to need black walls and no lighting. We may need to make our monitor dimmer to match the projector.

It is difficult to make a flare correction based on instrument measurements. If you want to make an accurate correction, we suggest the following:

Select an image with interesting detail in the deep shadows. A real image seems to work better than an artificial one. Set the monitor brightness to the working value.

View the image with no flare correction and with the lights off. Remember the appearance.

Turn the telecine suite lighting up to your usual working settings. If possible, use some preset on the lighting controls so you know you get back to the same illumination each time.

View the image with the lights up. Vary the display flare correction until the appearance matches your memory of the image without flare correction in a dark room. Check that the same features are visible under both conditions. If necessary, go back to the uncorrected image in the dark to remind yourself how it looked.

Check the effects of the correction on a coloured image. The flare correction will make the mid-tones brighter, and will desaturate pastel colours. The effect should be small unless you have overdone the flare correction.

We have left the display adaption at the dark surround setting. This is probably right as the immediate surround to the screen is probably dark coloured and only weakly illuminated even in the working conditions. You can try using **displayAdapt{0.5}**, the correct setting for a grey surround. If you do, you will probably find that the effect is almost the same as using more flare correction.

Telecine suites are one of the few cases where it makes sense to put flare corrections into a profile.

3.2 Display calibration

The display calibration links CIE colour coordinates to display RGB values.

You can select your display calibration using...

```
display{ display calibration filename }
```

The display calibration will probably be named **monitor** or be named after the display type or reference colour space. The default name is 'monitor'. This name is used by convention for the preferred monitor on any given system for viewing images. This allows us to have a single profile file referencing the local **monitor** display on each workstation that uses it. There is no 'monitor' calibration supplied with the Truelight software – this should be made locally using a Truelight probe and a software application.

There are two types of display calibration. There are fixed white point calibrations that take in CIE L*a*b* where the white point is always 100,0,0; and there are variable white point calibrations that take in CIE XYZ, and the white point can be set using **white_Y{}**, **white_u{}** and **white_v{}** commands. You can also convert your display calibration to a L*a*b* calibration so the device white is mapped onto 100,0,0 without having to find the Yu'v' settings, with...

```
displayLab{on}
```


Most display calibrations have variable white points. Fixed white point display calibrations are used for...

- Displays with limited dynamic range such as LCD displays. To make the best use of the available range, we shall want to map our white point onto the display white.
- Defined colour spaces, such as the ITU rec709 video standard.
- Devices with white point controls such as Sony HD monitors.

3.2.1 Video Calibrations

This plot shows experimental data measured off a Sony HD monitor calibrated in a darkened room using the pluge test image; the Truelight SonyHD fixed calibration; and the tone curve function taken from the ITU rec709 standard.

These measured values are typical of well set up HD monitors. A new HD monitor will give a smooth curve that gets steeper towards the shadows. An older monitor may have a background glow that stops our shadow range going beyond about 14 stops. We have never found a display that comes close to the ITU rec709 standard.



If you want to produce images for viewing on a particular HD monitor, then we recommend you calibrate the HD monitor as though it was an ordinary computer CRT. You can do a L*a*b* calibration, or you can do an XYZ calibration and fit the white point using **displayLab{on}**.

If you are exporting HD data for viewing on some unknown HD device, we suggest you use the fixed SonyHD display calibration. HD displays are not all identical. There are often visible differences between similar displays. However, the SonyHD calibration is our best guess of current good practice.

If you are importing HD data from some remote site, you may want to see how close your HD monitor is to SonyHD. Do a Truelight calibration of the monitor. Fit the white point using **displayLab{on}**, then switch between the HD monitor calibration and SonyHD using the **display{}** command. The difference should be hard to see. Or you can compare the calibrations graphically using the **tl-plot** utility.

Most monitors only offer two controls – gain and offset – which is not enough to match the whole tone curve. However, if you have a Truelight probe, then the Truelight calibration utility **Interactive Setup** allows you to adjust your display using a white and a grey patch. If you can get a 500:1 contrast between the white patch and the grey one, the settings should be close to SonyHD. See the FilmLight Technical Note: *Standard Colour Spaces* video section for the theory behind this.

3.3 White Point Correction

If your profile contains the following lines...


```
whiteValue{ 1023,1023,1023 }
white_Y{ 16 }
white_u{ 0.1978 }
white_v{ 0.4683 }
```

...then it will map the Cineon 1023,1023,1023 point onto the CIE L*a*b* 100,0,0 white. If the display calibration has a variable white point, then the L*a*b* 100,0,0 point is mapped onto a CIE XYZ white defined by the **white_Y{}**, **white_u{}** and **white_v{}** parameters. If you do not define **whiteValue{}**, then we skip fitting any Cineon value to 100,0,0. The default Yu'v' settings correspond to a 16 foot-lambert white matched to D65.

3.3.1 Cineon White Point

The **whiteValue{1023}** setting is good if you view Cineon data on a display, and your recorder can cover the entire Cineon range. If your recorder can only go up to the reference white, then you would probably chose the reference white instead.

If you have strong pastel colours in your images, then you may have to choose your white point more carefully...



The diagram shows the colour gamut of a graphics monitor (grey), and projected vision film (colour), fitted so the black and white points coincide. The points are plotted in L*a*b* space, with the and black-white (luminance) axis vertical. The diagram shows the fit you would get with **whiteValue{1023}** and the display white point matched to the projector plus film base white.

Clearly there are colours that we can get on one medium that we cannot on the others. The film red, green, and blue corners of the film gamut are a lot lower (dimmer) than the monitor ones. We could perhaps fit one of these corners by varying the **whiteValue{}** setting, but it is unlikely we could do a good job of fitting all of them.

Suppose you were using an inverse transform to convert images from video (monitor) to film. All the monitor pastel colours except for a few yellows and pinks cannot be rendered on film. If we set **whiteValue{685}**, then the film gamut will be stretched vertically until the video white lay on the film reference white. The white regions in your video image will then be rendered as the film reference white. This may look a bit grey if you view the film images on a light box, but will look fine in a darkened cinema. You will also get all the pastel colours in your video signal, because you can achieve film densities that are lower than your white point.

The Cineon white point fitting is disabled with the command:

```
whiteValue{ }
```

There are sometimes experimental reasons for wanting to turn off the Cineon white point fitting. For most work, it is sensible to set it to a known value.

3.3.2 Display White Point

The commands the **white_Y{}**, **white_u{}** and **white_v{}** set the display white point. These commands are ignored if the display has a fixed white point or we have set **displayLab{on}**.

The default white point is 16 foot-lamberts matched to D65. This does not correspond to any cinema or video standard, so you will probably want to edit these values.

16 foot-lamberts is the US standard figure for an open-gate projector white. The brightness may be given in candelas per square meter. One candela per square meter is 0.2919 foot-lamberts. If our **whiteValue{}** setting matched film base, then our white point will be the white of the same projector loaded with film base. This will be dimmer and more yellow than the open gate standard.

30 foot-lamberts is a typical video white value. If you are viewing your image under typical video conditions, and you do not have a reference projector to compare to, you may choose to fit your white point to your monitor.

There are standards for cinema white chroma, but in our experience many real projectors do not match these well. The projector chroma can also change with bulb life, with lamp house focus, and with time if the arc position jumps. The default D65 chroma fits the video standards, and falls within the range of colour temperatures we have measured in cinemas.

If you want to match a particular projector, you can measure the projector Yu'v' values with some colour measuring instrument. However, unless your projector calibration was based on the same instrument, small differences in calibration can give visible differences. The Yu'v' values are best fitted by eye.

If you have a projector next to your display, we suggest letting both warm up for 30 minutes, then loading a simple image such as the grey slate. Set the grey slate densities for the particular print using the **greyStatusA{}** command, then adjust the white point settings by eye until the projector and display look the same. If your projector cannot hold a single frame, and you do not have a strip of your grey slate, you can use a loop of film base and a flat white image. This will not be as good: Cineon 1023 is not the same as unexposed print, and white is an extreme point in the colour space. You may have to fine-tune the settings with a real image.

If you have a $L^*a^*b^*$ display calibration or **displayLab{on}**, then the **white_Y{}** command will have no effect. There is a second set of commands that works even in $L^*a^*b^*$:

```
brightness{ L*a*b* luminance scale factor }  
shift_a{ L*a*b* chroma a shift factor }  
shift_b{ L*a*b* chroma b shift factor }
```

This scales the perceived brightness or luminance L^* while leaving the chroma a^* and b^* unchanged. The default scale value is 1.0. If you are not trying to match an absolute light level, then brightness can be a useful control that lets the user get the brightest image from their display. If it is a user control, then having a fixed value in the profile may cause problems.

3.4 Matching displays

The **outputDisplay{}** command allows us to simulate one display on another. This command transforms the cube output values. It is applied after all the other operations.

Suppose you have generated a transform from Cineon RGB to video using **display{SonyHD}**. You may want to view the results on your local monitor, even though it is non-standard. If you replace **display{SonyHD}** with **display{monitor}** the gamut alarm will show the gamut on your local monitor, which is perhaps not what you want. If you set **outputDisplay{monitor}** instead, then Truelight generates a cube using the **SonyHD** calibration, and then will warp it to give a matching view on your local monitor with the **SonyHD** white point mapped onto the **monitor** white point. Because the gamut alarm is applied before **outputDisplay{..}**, the gamut alarm will show the gamut of the **SonyHD** display. The cube edits described in the next section are also applied before **outputDisplay{..}**, so these edits would be applied in the **SonyHD** RGB colour space, before converting to the output display.

In the more general case, the **outputDisplay{..}** command will transform from the current output space to the given display. If the display or the current output space is $L^*a^*b^*$ or you have selected **displayLab{on}** then **outputDisplay{..}** will map the white point of the output space gamut onto the white point of the display gamut. If the display and the current output space are both XYZ, then it should make an absolute match.

If you are using **invert{on}** or **cubeType{inputGamutClip}**, then the transform to RGB is taken from the input instead of the output. If you are using **transform{..}** or **importCube{..}** then Truelight has no input or output transforms, but you can define a transform using **display{..}**.

The **outputDisplay{}** command is also used for rendering to standard colour spaces. Suppose you were wanting to render to **DCIxyz**. If you set **display{DCIxyz}** then you would be unable to use the gamut alarm because the CIE XYZ gamut contains all possible colours. If you get the image to display correctly on your reference display using the gamut alarm, the cube edits and all the other tools, you can then add **outputDisplay{DCIxyz}** to your profile, and get the same colours and edits encoded as DCI XYZ'.

If you set **outputDisplay{SonyHD}** then you might see a shift in your white point that you did not get with **DCIxyz**. This is because **SonyHD** is an $L^*a^*b^*$ transform. You can edit **SonyHD** to make an XYZ video calibration with your house white standard. An easier option might be to use **tl-cal** to make an XYZ calibration of your video display.

3.5 Colour Edits

There are psychophysical effects that can cause our image appearance to be modified by factors outside the image. Earlier, we described the flare controls that attempted to correct for the colours surrounding the image. If your image has flare, then the right thing to do is to black out your room properly. The flare controls are only for people who cannot black out their room, and we know these controls affect the colour accuracy.

These edits should be used with similar care. The Truelight system will get your colours as close as it can on measurements alone. We measure over a thousand colour patches, and interpolate the colours in between. This assumes that the other colours lie on some smooth function we fit through the measured points. Over much of the colour space, this approximation works, but the interpolation will not necessarily follow what the film does in all cases. If the Truelight calibration has been done properly, then you should have a good colour match between monitor and the film over most of the colour space. If this is not so, then you should look for the cause before using any edits.

These edits work in display RGB. If you want the edits to continue to be useful after you have recalibrated the display, it is worth taking extra care to balance the white and black points of your display, so the RGB tone curves all match.

3.5.1 Contrast and Saturation

We can justify using subjective colour edits if they are compensating for some psychophysical effect.

There are known effects that can raise or lower the perceived colour saturation level in an image. We often find that people prefer a monitor image with 5% desaturation to the best match by instruments. You can do a 5% desaturation with....

```
saturation{ 0.95 }
```

The saturation control can also be used with the gamut alarm to give some idea of how far out of gamut a troublesome colour is.

Truelight will try to find an exact match for each colour. Some displays such as LCDs have a smaller contrast range than the film they are trying to match. The highlights will be rendered correctly, but the shadow details will be lost. You can get a better visual match by reducing the luminance range, while preserving the chroma, with...

```
contrast{ 0.95 }
```

HD monitors can show the opposite problem: Truelight is matching the light levels in the projection, but differences in flare conditions make the shadow detail more visible. We could add flare, but this makes the shadows look grey or 'milky'. This may look fine on a projected image but can look unnatural on a CRT. Sometimes, it is necessary to artificially extend the range of a video signal to get the right looking trace on a scope. In either case, you can use the contrast control can reduce the luminance range, while keeping the colours the same.

Both of these controls may be used to match a digital proof to a projected print. However, unless both controls are at their default settings, Truelight will not match the in-gamut colours exactly.

3.5.2 Colour Edit Functions

You can apply an arbitrary function to the data in several places within Truelight using...

```
valueFormula { Cineon formula }  
LabFormula   { L*a*b* formula }  
rgbFormula   { display RGB formula }
```

These formula commands give a very general control, but they are not easy to use.

3.5.3 Interactive Colour Edits

Truelight supports colour editing on the cube output RGB values. A typical Truelight cube will have output values in floating-point display RGB, mostly in the range 0..1. These can be used to fine tune the colour transform to remove any remaining colour shifts.

A typical edit command might look like this...

```
editA{ in{123,150,115} out{120,150,110} r{80} e{4} /{255} }
```

This will map the input colour (123,150,115) onto the output colour (120,150,110). This edit will affect colours in the input colour space within a radius of 80 of the input colour.

The optional scale argument `/255` ranges the input and output colours and the range values from 0..255 to 0..1.

The optional eccentricity argument `e{4}` increases the range of the edit in the luminance direction: this is often useful to change all instances of a particular colour over a wide luminance range.

These values are best set using some user interface. The default scale is 1.0. The default radius is 0.5 after scaling. The default eccentricity is 1.0.

You may think, if you put in enough edits, we should be able to get all the colours right, even without Truelight. In practice, this does not work. Each colour edit carries the risk that it will drag some other colour out of calibration. Because colour space is three-dimensional, each colour has a lot of potential neighbours. Unless Truelight can get all the colours to a good starting point, by your 29th secondary colour correction, you may have undone much of the work of the first one. This is why we have restricted ourselves to eight colour edits in total. We encourage people to use these edits only when most of the colours are coming out right, to use as few edits as possible, and to keep these edits smooth and long-ranged.

The Truelight viewer sets these edit controls with a graphical interface that allows the user to pick and edit colours interactively. You can pick the source and destination colours, and the luma and chroma range. The application turns these settings into an edit command.

You can view the extent the edit as a greyscale by adding `gamut{on}` to the edit arguments. Colours affected by the edit will appear light against a dark background.

There are 8 possible edits from `editA` to `editH`. The edits are applied in alphabetical order, stopping at the first edit with `gamut{on}` if there is one.

All these edits can be disabled using the command...

```
editCube{ off }
```

If you edit in RGB, your edits parameters are fixed to your display RGB space. If you recalibrate your display then these edits may no longer look right. All these edits can be done in L*a*b* space with the command...

```
editLab{on}
```

See the *Truelight Commands* technical note for more details.

4 Working with Truelight

Truelight makes a transform from input log Cineon to output display RGB. It uses this transform to make a cube transform to process image data quickly. The cube has 1D input tables to convert values into cube levels and fractions, followed by a 3D table of output values.

This table shows the various options in the order they are used when generating a cube...

Truelight Command	Notes
All commands not listed below	The transform from input RGB to output RGB is generated as a list of input and output values. Any non-numeric values are removed.
invert{}	If this is enabled, the list input and outputs are swapped around.
cubeInFormula{}	This formula transforms the list input values.
cubeLutFormula{}	The cube is set up with linear input tables and evenly spaced cube levels. The cube look-up tables and levels are then transformed using the supplied formula, if there is one. This does not change the transform, but it can redistribute the cube levels to make the interpolation more accurate.
linearizeCube{}	If this is enabled, the cube input tables and levels are warped so the cube output values are at even intervals along the R=G=B diagonal. This can make the interpolation more accurate, so it is enabled by default. You should turn it off if you are using cubeLutFormula{} .
cubeType{ intervals{ fitGamut{on off}	The cube values are generated from the transform. The cubeType{} command can select processing options such as the gamut alarm. When a gamut clipping cubeType{} option is selected, fitGamut{} determines how out of gamut points are clipped. The default fitGamut{on} projects out of gamma points onto the gamut surface in L*a*b* space.
cubeOutFormula{}	The cube values are transformed using the supplied formula. This affects the overall transform.
monotonicCube{}	With this on, the output values in the 3D cube are forced to increase with the corresponding input value. This can remove artefacts from extrapolating noisy data. It is enabled by default, but it may not be appropriate for all colour spaces.
outputDisplay{}	Convert from the current output RGB (typically display{} RGB) to outputDisplay{} RGB.
editA{ } ... editH{ }	The colour edit controls are described in the previous section.
cubeInput{}	The cube input LUTs are transformed using the fixed formulas for linear and other spaces. This happens after the cube values have been generated, so it affects the transform.

The rest of this section will describe these options, and how they can be used.

4.1 Cube Type Options

The **cubeType{}** command lets us make several useful types of cube from the same transform. The **invert{}** command inverts the transform before the cube is generated. There are several useful combinations of these two commands...


4.1.1 Working with Film Data

```
cubeType{ straight }
```

This is the default **cubeType{}** setting, so there is no need to put this line in a profile. It makes a cube that transforms from the input space to the output space. The cube should be similar to the original floating point transform, except that the input values are restricted to the in-gamut range.

```
cubeType{ gamutAlarm }
```

This command makes a cube with a gamut alarm. The gamut alarm clips the values outside the 0..1 gamut, and replaces the in-gamut RGB colours with the corresponding R=G=B neutral. When you transform an image through a gamut alarm cube, the colours that cannot be matched in the display show up against a monochrome background. Colours that are interpolated from in gamut and out-of-gamut points are saturated in proportion.



Gamut alarm on standard Kodak LAD test image

The 3-d gamut plot shows there are colours that we can get on film that we cannot get on our display, and vice-versa.


This looks like a serious problem. In fact, these extreme colours are rarely found in real scenes captured on film. Often, reds and yellows go just out of the CRT gamut, but the overall appearance of the image is hardly affected. However, it is important to identify which colours on our display should not be trusted.

The gamut alarm is a useful tool to turn on and off when you are working. You are unlikely to want it permanently selected in a profile.

The gamut alarm shows the out of gamut colours clearly, but it is much harder to spot highlights and shadows. Truelight therefore supports a gamut map option that shows all out of gamut values as bright colours. You can get a gamut map with...

```
cubeType{ gamutMap }
```

This sets red if values are too low, green if values are too high, yellow if values are too high and too low, and the average grey if all the values are in gamut. In this example, I have turned up the brightness so we can see what the overflow looks like...



Computer graphics 8-bit RGB images typically use the whole range 0-255. Legal 8-bit video RGB is restricted to the range 16-235. Video values outside this legal range have little practical use. Most of the industry seems to favour working with full-ranged computer graphics images, and converting to the video range on output. The video cube types are rarely used these days. They are...

```
cubeType{ videoGamutAlarm }
```

This is similar to the ordinary gamut alarm, but it highlights values that are outside the legal 8-bit video gamut range 16..235. This is the same as...

```
cubeInFormula{ (#*255-16)/219 }
cubeType{ gamutAlarm }
cubeOutFormula{ (#*219+16)/255 }
```

The video gamut alarm is a useful tool to turn on and off when you are working. You should not want it permanently selected in a profile.

```
cubeType{ videoGamutClip }
```

This command generates cubes where all the output values are clipped to the legal 8-bit video gamut range 16..235. If your application has a gamut alarm option, it may overwrite this setting, and not restore it afterwards. If you are using the Truelight hardware, this has an independent legalizer option. This is the better option if you can use it.

4.1.2 Working with Video Data

We can generate an inverse cube with the **invert{}** command...

```
invert{ on }
```

An inverse cube will convert from display RGB to log Cineon. An inverse cube generated with **display{rec709}** or **display{SonyHD}** can be used to transfer images from video to film.

```
invert{ on }  
cubeType{ inputGamutAlarm }
```

The **cubeType{inputGamutAlarm}** command generates a cube where the output colour space is the same as the input space. Colours that would have transformed to out-of-gamut values are unchanged; all other colours are replaced with neutrals. This is useful when combined with the **invert{}** command. Then, when you pass a display RGB image through the cube, the colours that are in the display gamut, but cannot be rendered on film will show up in colour. This is the equivalent of the gamut alarm for people who work in display RGB.

```
invert{ on }  
cubeType{ inputGamutClip }
```


The **cubeType{inputGamutClip}** command also generates a cube where the output colour space is the same as the input space. This cube simulates the effects of transforming to the output space, clipping, then transforming back again. It can be used with the **invert{}** command to simulate the effects of clipping to the film gamut when transforming from display RGB to Cineon.

The **outputDisplay{..}** command converts the current cube output space to fit the given display. If you have an inverse cube, then **outputDisplay{..}** with **cubeType{gamutAlarm}** can give an equivalent to the **cubeType{inputGamutAlarm}**.

4.2 Film Colour Space Options

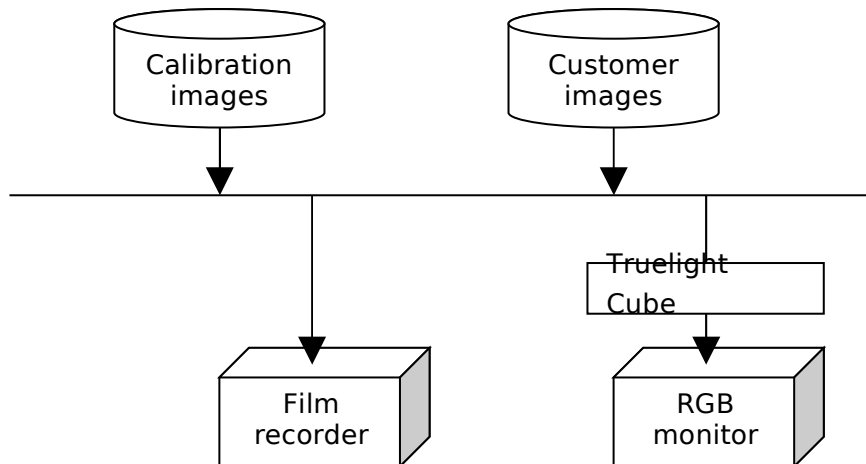
So far, we have been describing film colour as though all the film images were in Cineon log density space. The **printerPoint{}** and **aimGamma{}** settings rely on the RGB values being proportional to negative density. However, some film recorders cannot work in Cineon log. We can still calibrate them using the film colour space options.

The Truelight calibration involves outputting a series of images with colour patches like this...



Truelight Calibration Strip Image

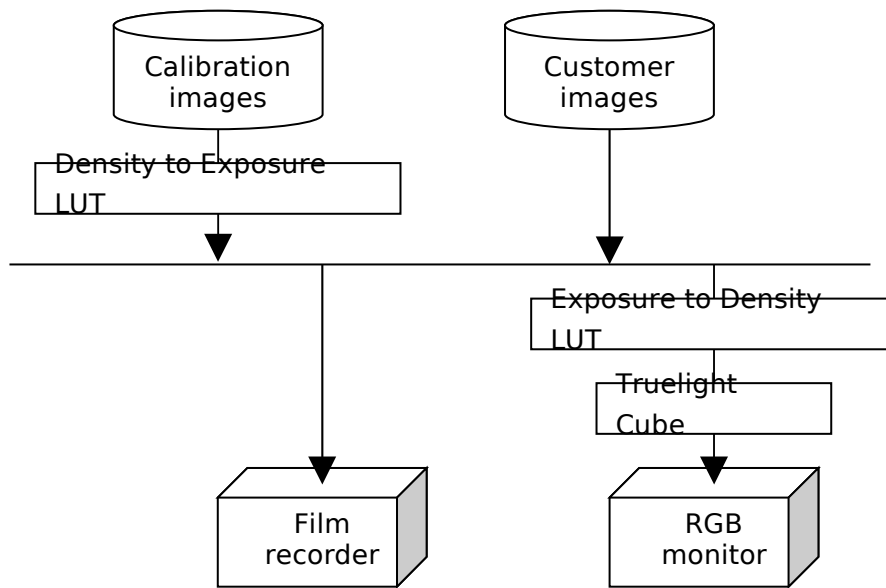
The calibration strip contains colour patches that span the whole RGB colour space with visually even intervals. This strip must be output in the same way that all the other images are output. If everything but the monitor worked in Cineon density space, then our data flow would look like this...



Data flow when working in density space

Suppose we kept the same data flow but switched to working in exposure units. If we use the same calibration images, we would get a set of patches with even exposure intervals. We would have a lot of light patches, and very few dark ones. The Truelight interpolation would also be done in a colour space where equal intervals did not appear as equal contrasts. You would get a calibration, but the results would not look good.

Let us add look-up tables so the Truelight components still work in density, but the customer's images and the film recorder use exposure space. Our data flow would then look like this...



Data flow when working in exposure space

The calibration images would only need converting from density to exposure units once, so their LUT operation would not affect productivity.

The exposure to density LUT must operate on every Truelight viewed image. However, the Truelight cube already has a look-up table on its input. We can merge the look-up table with the other one, and avoid any loss in productivity.

This example shows how we would manage the conversion between density and exposure space. In practice, we can replace our density space with any well-behaved colour space that gives a good set of colours on the Truelight test strip; our recorder calibration will always give us the exact relationship between the RGB values and the negative density. We can also replace our exposure to density transform with any smooth function that can be fitted into 1-D look-up tables.

Because we have considerable freedom in our choice of 'density' and 'exposure' spaces, we support some fixed transforms between the two. The normal Truelight colour space is Cineon density, also known as Cineon log. Kodak also defines fixed formulas for linear and video Cineon colour spaces. For values between 0 and 1...

$$Video > 0.081 \quad Linear = \left(\frac{Video + 0.099}{1.099} \right)^{\frac{1}{0.45}}$$

$$Video < 0.081 \quad Linear = \left(\frac{Video}{4.5} \right)$$

$$Log = \frac{685}{1023} + \frac{300}{1023} \cdot \log_{10} \left(\frac{Linear + 0.0109}{1.0109} \right)$$

Note: Cineon video is not conventional video. It is just a one-dimensional transform of the Cineon negative density space to give a more 'video-like' tone curve on typical print stocks. We do not encourage use of Cineon video space, because it is too easy to confuse with conventional video.

We also would like to discourage people from describing colour spaces as 'linear' without some qualification. This Cineon colour space is in linear units of exposure through the negative. Other popular 'linear' colour spaces include exposure to the negative, light levels in the scene or rendered scene, light levels in the final image, and CRT voltage.

If we have a Cineon log calibration, then we can make a cube that takes in log, linear, or video data using one of the following...

```
cubeInput{ log }
cubeInput{ linear }
cubeInput{ video }
```

The default setting is 'log'. These settings are only appropriate for Cineon-like colour spaces, so **cubeInput{}** is not appropriate if you are using inverse cubes.

There is a more versatile function that transforms the recorder calibration data as well as the cube LUTs using an arbitrary formula. The command checks that the formula does not have any cross-coupling between the channels. It then applies the formula to the transform input and the cube look-up tables. The command is...

```
cubeLutFormula{ formula }
```

The command with the linear-to-log formula would be...

```
cubeLutFormula{ (685/1023) + (300/1023)*log10((# + 0.0099)/1.099) }
```

If you add just this line to your profile, then you will probably see little effect. This is because you have not changed your recorder transform, but just re-arranged where the cube intervals come.

If we are not in Cineon space, then the aim gamma fitting will not be appropriate. If you have an **aimGamma{}** command, you should remove it.

We can still fit aim gammas with **cubeLutFormula{}** provided our film recorder calibration can be transformed to Cineon units. Suppose our film recorder has been calibrated to cover the density above film base range 0.0 to 1.2 in units of transmission. For aim gammas to work, the recorder calibration will have to be in Cineon units. We can convert the calibration values from Cineon to transmission with **valueFormula{}**, then get the cube intervals in the right places with **cubeLutFormula{}**...

```
valueFormula{ (1 - 0.1^(#*2.046 - 0.1))/(1 - 0.1^1.2)
cubeLutFormula{ log10(1 - #*(1 - 0.1^1.2))/-1.2 }
```

We can use **cubeLutFormula{}** with inverse cubes. The formula is applied after the transform has been inverted, so it will be applied to the display RGB. If you are working in linear light level display RGB, then you might use...

```
invert{}
cubeLutFormula{ #^0.4 }
```

The **cubeLutFormula{# ^0.4}** command makes a cube with gamma 0.4 input LUTs. This formula stretches out the shadows and compress the highlights in your linear display RGB, so the cube interpolation effectively happens in gamma 2.5 space. This gives the same transform, but the more perceptually uniform distribution of cube intervals gives smoother-looking results.

```
Invert{}
cubeLutFormula{ (#*255-16)/219 }
```

If you have video input, this generates a cube with the first interval at 16 and the last at 235 when working with 8-bit data. This is useful for legalized video data: it puts the first and last cube intervals at the gamut limits, and avoids wasting LUT range on the illegal values.


```
cubeLutFormula{ max(1/255, min(251/255, #)) }
```

This can be useful if you are driving a DVI device via a HD video link. RGB data is coded as YCrCb before sending down a HD link, and is decoded to RGB at the other end. HD video RGB values are usually confined to the 8-bit range 16-235. We can use RGB values outside this range, but values beyond 1-251 can give overflow or underflow errors when converting to or from YCrCb.

```
linearizeCube{ off }
```

The cube generation has an option to linearize the cube. The input LUTs are warped to give cube intervals with uniform intervals along the diagonal in the output space. In general this gives smoother transforms, so it is enabled by default. However, the effects of following this with an arbitrary **cubeLutFormula{}** setting are not fully explored.

You can combine `cubeLutFormula{}`, `linearizeCube{}`, and `cubeInput{}`. If you combine them, then this is the order in which they operate on the LUTs. Combining them is not harmful, but the results may be difficult to predict. If you are using `cubeLutFormula{}` or `cubeInput{}`, you might prefer to explicitly turn the cube linearization off.



Cube points plotted in 3D looking down the blue axis with `linearizeCube{off}` and `linearizecube{on}`

The diagram shows the position of the cube points in the input space. With linearization off, Truelight uses uniform intervals in the input space. With linearization on, we get large intervals at the extremes, where large changes in negative density only give us a small change in the print, and smaller intervals in the centre where we get most of the image data. It may look surprising, but the results look right.

However, there are exceptions when we are looking at raw Cineon data. Suppose you want to see whether your display could show all the shadow detail in your Cineon image. In the Truelight viewer, you could simulate clipping to the display gamut in Cineon space with...

```
cubeType{inputGamutClip}
fitGamut{off}
linearizeCube{ off }
```

By default, Truelight clips colours to gamut in L*a*b* space, preserving the hue were possible. The `fitGamut{off}` option turns off this L*a*b* gamut clip, and the output RGB values are just clipped.

If you turn off the profile (F6), the image should look the same provided the colours are all in gamut. However, the cube is accurate near the Cineon limits only if we turn off the cube linearization.

4.3 Other Cube Generation Options

There are some commands that do not affect the transform itself, but modify how the cube does the transform. Most profiles use the default settings, but there may be special cases when these parameters are useful.

```
fitGamut { on }
```

This is the default setting. If the interpolated output values are not in gamut, the library routines search for the closest in-gamut point in L*a*b* with matching hue, and - where possible - matching luminance. This should give better results because the interpolated value will lie closer to the points it is being interpolated from. You can revert to the older, simpler approach with...

```
fitGamut { off }
```

You can make cubes with between 2 and 40 points per axis with...

```
intervals { number of cube intervals }
```

The default value of 16 seems to work for most conditions. You can try increasing the number of levels for the badly behaved colour spaces such as the gamut alarm. More intervals can give better results and the interpolation is just as fast, but the cubes take much longer to calculate. Larger cubes may not be compatible with Truelight hardware.

```
intervals { 0 }
```

This is a special case: it gives a transform with input look-up tables but no cube. This is faster than the regular cube transform, but it is incapable of doing any 3-D correction. This can be used to generate look-up tables for conventional display hardware. It has mostly been used to show off the advantages of the full 3-D transform.

```
monotonicCube{ off }
```

The cube generation includes a monotonic filter. In general, we would expect each output value to increase if we increased the corresponding input value. However, small experimental errors and/or interpolation overshoot can give values that go against the trend. The monotonic filter corrects these values. Usually, this is a good thing to do, so the default setting is 'on'. However, you can turn off this filter for output spaces that are intentionally non-monotonic.

You can make cubes for uploading to the Truelight hardware with...

```
cubeHardware{ 10-bit }  
cubeFilename{ file name }
```

The '10-bit' argument selects the 10-bit Truelight hardware build version. If you have Truelight hardware, then the software that generates the data will probably insert lines like these. These commands should not normally appear in a profile. The option to save to a named file may be disabled for some versions of Truelight.

```
title{ title string }
```

The title does not affect the transform. It is used to describe and identify the profile. The first 16 characters of this string are echoed on the Truelight hardware front panel and the HD signal (in upper case only) when the cube is selected or loaded.

```
exportType{ type }  
exportOptions{ string }
```

If you have the option to save cubes in other formats then the Truelight software will probably insert lines like these for you. These commands should not normally appear in a profile.

5 Example profiles

Here are some simple examples of profiles for typical Truelight applications.

5.1 Previewing film log images in office lighting

If your images are in film log space in an office environment, then your profile might look like this...

```
title{ simulating print on monitor }
recorder{myRecorderCalibration}
print{myVisionCalibration}
display{monitor}
displayFlare{0.055}
brightness{1.25}
greyValue{445, 445, 445}
greyStatusA{1.09, 1.06, 1.03}
aimGamma{0.9237, 1.0, 1.0}
whiteValue{1023, 1023, 1023}
```

Here, we use the standard grey patch values, and the default aim gammas for an ArriLaser using the Carlos program. We shall keep these values for the other examples. We are using the default xenon lamp calibration, and the default D65 16 ft-lambert white point. The **whiteValue{}** command maps the Cineon white limit onto this white point. You can probably leave out this line if we are not being fussy about the exact white point, as most calibrations will naturally produce a white point close to this.

This is probably a bit dim for office use, so the brightness has been turned up a bit. If you are doing this, check that the highlight values can be displayed, and you are not reading the limits of the display RGB.

The office lighting will be adding flare to the image, so this user has turned up the flare correction. This may need adjusting as the light levels vary. You can do this if you have an image with known shadow detail.

To view the image with the gamut alarm on, add the line...

```
cubeType{gamutAlarm}
```

This will highlight all the colours that you can get on film but you cannot match on your monitor.

A Shake node user with a Truelight licence might use these settings. If they did not have a separate Truelight licence, they would have to use the internal recorder and print calibrations. These will give a good match for standard printing conditions.

5.2 Generating HD from film log images

If you want to make an HD version of your log images, your profile might look like this...

```
title{ simulating print on monitor }
recorder{myRecorderCalibration}
print{myVisionCalibration}
display{SonyHD}
greyValue{445, 445, 445}
greyStatusA{1.09, 1.06, 1.03}
aimGamma{0.9237, 1.0, 1.0}
whiteValue{685, 685, 685}
videoBlack{0.04, 0.04, 0.04}
```

This is a bit like a film preview profile, except that we are using the generic Sony HD calibration instead of our monitor calibration. This is an L*a*b* calibration, with the (100,0,0) white point on the display white (1,1,1). The **whiteValue{}** command maps the Cineon reference white onto this white point.

The Cineon reference white is an arbitrary choice for the white point, but there are several reasons to pick a point below Cineon maximum white:

- The Cineon RGB space has more range than you can reliably print. Most sensible Cineon data is confined to the range 90-685. The HD RGB space is all viewable colours. It then makes sense to map the sensible Cineon exposure sub-range onto the entire HD range.
- Film pastel colours can look grey. We can cure this if we depress the white point from the clear film base to a slight grey. This grey cast will be visible when you look at a print on a light box, but it will be invisible when you project the film in a cinema. This makes the film pastel colours look cleaner because they can use exposures that are greater than the white point.
- The print film tone curve goes very flat at the highlights and the shadows. This means we need large changes in the negative density to make a small change in the highlight density. A small variation in the processing chemistry can change the highlight density as much as a large shift in negative density. We do not have good control over the lightest portion of the print tone curve, so we should avoid it if we can.

10-bit Video RGB is conventionally ranged from 64-940. Values outside this range are not necessarily harmful, but they will cause the video to fail QA. Most video hardware, including the Truelight hardware, allows us to scale computer graphics data ranged 0-1023 to the conventional 64-940 video ranges as part of the transform to YCrCb for 422 encoding. If you are using 444 RGB video, then there is no transform to YCrCb. You may have to put in the ranging explicitly by adding the line...

```
cubeOutFormula{ (max(0,min(1,#))*(940 - 64) + 64)/1023 }
```

This profile does not have a flare correction. Films are normally viewed in the dark with dark surrounds, while HD editing suites normally have grey walls and a D65 backlight. You may decide that some flare correction is appropriate.

This profile does not shift the white point. The film white point is mapped onto the video white. A film projector loaded with clear film base is generally yellower than the video D65 white. However, many people are happy to work with different white points for film and video because:

- The HD and film images are being displayed in different rooms.
- This difference appears as an overall colour cast. The colour contrasts within the image will look the same.
- The projected film will be viewed in a darkened image. Under these conditions, the eye can readily adapt to different white points.
- HD editing suites normally have grey walls and a D65 backlight. People will expect to see a D65 white point under these conditions.

You can shift the HD white point with the Truelight L*a*b* controls, but this will make the HD signal look odd on the scope as the white levels will be different for each channel.

If you want to shift the white point, use the **brightness{}**, **shift_a{}** and **shift_b{}** commands, as the Yu'v' controls will not work with the L*a*b* SonyHD calibration. Alternatively, you could use a measured XYZ calibration for your particular monitor, and the white point Yu'v' controls. Either of these will give you an image with the right white point, but the HD signal will now look odd on the scope. You could change your HD monitor controls to match its white to the film white, but that would not work if the signal were viewed on a different monitor. In general, we cannot get both the white point and the scope trace looking right at the same time. If you are not viewing the film and the HD at the same time, it may be easier to live with two different white points.

The RGB image we generate should look sensible on the monitor, but the corresponding HD signal will not necessarily look sensible on a scope. The film black is not a true black, so our black signal may have non-zero RGB values. As conventional video is expected to go down to zero in the shadows, we may have to distort the tone curves slightly to get the scope signal to look right. The **videoBlack{}** command forces the Cineon value ranged 0-1 of (0.04, 0.04, 0.04) to the video black (0,0,0). This choice of Cineon value is a bit arbitrary, but something between 0.0 and 0.09 ought to work.

The **videoBlack{}** correction affects the extreme shadows. It is not supposed to have a significant effect on the image appearance. If this command changes the image appearance significantly, then there is something wrong with your data or your choice of black point. For example if you added a flare correction, the **videoBlack{}** command may undo it. Get the black point as close as you can to the video black using the other controls, and then try applying **videoBlack{}** again.


There is also a **videoWhite{}** command. It maps the given Cineon colour onto the video (1,1,1) white in a similar way.

5.3 Doing a 'butterfly' test in a theatre

In a 'butterfly' test, a film and a digital image are projected side by side on the same screen. Each image has one side blanked off, and one image is flipped over horizontally, so it looks as if the image has been reflected about the midline. This test should be done on a selection of images, including a set of grey patches, a digital 'Macbeth' chart, a scene with faces and flesh tones, a scene with sky tones, a scene with deep shadow detail, and so forth.

A butterfly test is not part of the production process. It is just an extreme test of Truelight's ability to match film. To do a 'butterfly' test we recommend outputting the film images you are going to project at the same time as the Truelight calibration strip and a grey slate image. This avoids most colour errors due to variations in the recording or film processing.


It is not easy to set up a butterfly test. This is usually done by masking off the regions of each projected image using pieces of card against the projection booth window. If you can get the two images the same size and the same brightness, then your images may look like this...



Sometimes the two images will appear to be a different colour. The flipped image will look unfamiliar, so you may be seeing some of the colours afresh. Small differences in sharpness or grain structure can also have unexpected effects.

There is another phenomenon that can make the colours look different. The projection lenses will be corrected to give the minimum chromatic aberration on the plane of the projected image. The pieces of card will be in a very different plane, so the masked edge may appear to be coloured.

In this image, we have simulated a coloured fringe close to the edge where one image dissolves into the other. This colour change does not extend much beyond the region of the dissolve, but the whole of the left-hand image may look blue, and the right-hand one look yellow.



Here is the same image, but with a black line down the centre. We have masked out the coloured fringe, and the effect has disappeared. If you can do so, we recommend hanging a strip of black material between your two images. If you cannot do this, try blocking out the join between the images using a pencil or a finger.



Your butterfly test profile might look like this...

```
title{ Butterfly test }
recorder{myRecorderCalibration}
print{myVisionCalibration}
lamp{myProjectorCalibration}
display{myDigitalProjectorCalibration}
white_Y{13.5}
white_u{0.1973}
white_v{0.4985}
greyValue{445, 445, 445}
greyStatusA{1.01, 1.03, 0.98}
aimGamma{0.9237, 1.0, 1.0}
whiteValue{1023,1023,1023}
```

In this case, we have set **whiteValue{}** to match the match the film white to the display white point. We have measured the grey slate values on an Xrite TR310 densitometer and put these values in the **greyStatusA{}** command, so this setting should match the particular print we are using exactly, without a printer point correction. The flare should be exactly the same on both sides of the screen, so no flare correction should be necessary.

We set the Yu'v' values to match the overall brightness and colour cast of the two sides of the images. To begin with, we can roughly match the two using a spot meter and a white patch. The best results are obtained when we try to get the best overall match for all the images by eye.

It is also possible to match a scan of a print to a projection. This is an experimental technique, but it can show how well Truelight can work when we take out the variations in the lab. We measure a print test strip in a scanner, and make a transform with scanner values as input, and Status A values as output. This will work as a recorder transform for a recorder loaded with print stock. We can then display images scanned off print stock. If these images were output at the same time as the test strip, we avoid any process variations. The print images also have all the artefacts of the recording and print process: dark edges, film grain and so on.

When matching by eye, remember:

- Film projectors are less bright at the edges. The digital projector may vary too. Trust the middle of the image over the edges.
- Film printing is rarely uniform - it can be darker and bluer at the edges.
- The line where the images meet can give strange perceptual effects. Try and mask it out.
- Film can bleach in minutes if a single frame is left up. This will make your film image lighter and less yellow. If you can stop your projector on a single frame, always have a reference copy of the frame to check for fading. If you cannot stop your projector, you will have to use a loop of film, so this will probably not be a problem.
- You may have light leaking from one projector onto the other half of the image. Check for stray light that can make the images look different. Cast a shadow onto the screen. Are there any secondary shadows, suggesting light is coming from an unexpected direction?
- The projector lamp colour may vary with position on the screen. It can also vary with time as the lamp electrodes wear, and the arc jumps to a new position. This is a very hard thing to spot, as the arc can be stable, then switch position very rapidly. It can also switch back again just as quickly. There is not a lot you can do about this apart from being aware that it happens.
- Your impression of colour will be affected by the image texture and the sharpness. Make oph images the same size if you can. If there are large differences in sharpness between the two halves, you can try defocusing the sharper projector.
- The white point luminance can drift with edits. It is easy to notice colour shifts and ignore the luminance difference. If you see different colour shifts in the highlights and mid tones, then check the luminance, preferably using a spot meter if you have one.

If you still cannot get a colour to match, remember to check the colour is in gamut on your digital projector.

It can take a long time to get a good match, but it should be possible without using any other colour corrections other than the ones in the example profile. If you are having trouble matching the whites, this may be due to artefacts in how we see colours, rather than the quality of the match in XYZ values See *FilmLight Standard Colour Spaces* for more details.

If you get a good match by this method, then you can use the profile for viewing other images. However, before you do so, remember to reset the **greyStatusA{}** command to the usual aims, as your values are only correct for this particular print.

5.4 Matching film images on a light box

Here is an example of a profile to match film on a 5600K light box.

```
title{matching monitor to 5600K tablet}
recorder{myRecorderCalibration}
print{myVisionCalibration}
lamp{my5600tablet}
display{monitor}
white_Y{26.5}
white_u{0.2197}
white_v{0.4963}
greyValue{445, 445, 445}
greyStatusA{1.09, 1.06, 1.03}
aimGamma{0.9237, 1.0, 1.0}
whiteValue{1023,1023,1023}
printerPoints{0,1.25,0}
```

The Yu'v' values were matched by eye to the light box white. We can do this the same way we set the white point in the 'butterfly' test, though it is hard to get a good match as the images are very different sizes. The grey patch status A is reset to the aim values. In this example, the print looked too pink and too dark, so the user added a printer light shift.

5.5 Matching one display to another

You can display HD on your monitor like this...

```
title{matching monitor to HD}
inputDisplay{SonyHD}
display{monitor}
displayLab{on}
```

You can set your Yu'v' values by hand if you do not want to match the HD white to your device white. You can use the gamut alarm to show the HD colours that you cannot get on your monitor.

5.6 Using an LCD monitor

LCD monitors do not have enough contrast for grading films. Their images may look fine under office lighting but a black screen will look grey in a dark room. We can get the best results from an LCD monitor like this...

```
title{LCD monitor}
contrast{0.95}
display{monitor}
whiteValue{850,850,850}
displayLab{on}
```

We have compressed the contrast range so we should see the shadow detail instead of crushing it. We are fitting the film white to the display to use as much of the display range as possible. We set the white value to give a bit of headroom over the reference white. You may wish to vary the actual values to suit your display and source images.

5.7 Grading in HD

If your images are in HD space, then you can convert them to log images for output using an inverse transform profile...

```
title{ HD to log }
recorder{myRecorderCalibration}
print{myVisionCalibration}
display{SonyHD}
greyValue{445, 445, 445}
greyStatusA{1.09, 1.06, 1.03}
aimGamma{0.9237, 1.0, 1.0}
whiteValue{685, 685, 685}
invert{on}
```

This is almost the inverse of the log to HD conversion. Our input HD is likely to cover the whole range 0-1. We map this range onto the safe Cineon range, with our neutrals stopping at the Cineon reference white. This should map most of the HD colour space onto sensible Cineon colours.

Here, the film white point is mapped onto the Cineon white point, though the projected white is usually a different colour to the HD white. There is no single best **whiteValue{..}** setting. The 685 value in the example gives us some headroom to produce saturated pastels using film exposures beyond the white point. This may be needed for animation or advertising work with saturated colours: most other material will work with a higher value. Many users are happy to work with different white points for film and HD.

To simulate the effects of clipping to the film gamut, add the line...

```
cubeType{inputGamutClip}
```

This cube takes in HD, and outputs HD. If the whole image is in the film gamut, the image should look the same if the profile is bypassed. If some colours are out of gamut, these colours are replaced with the nearest in-gamut colours with the same hue.

To highlight the out of gamut colours, add the line...

```
cubeType{inputGamutAlarm}
```

This will output an HD signal, but all the in-gamut colours are now greys, so the out of gamut colours stand out.

Note: **inputGamutAlarm** and **inputGamutClip** always get you back to the original input colour space.

If the colour is in gamut, then the input will equal the output. Truelight cannot combine our simulation of the film gamut with a simulation of a change in white point. If we want to do this, we have to convert our data from HD to log, and then convert it back into the new display colour space.

There are good reasons for this limitation. If we change our display white point, then we have three gamuts: our initial display gamut, our film gamut, and our final display gamut with the shifted white point. We also need separate Yu'v' controls for the initial display white point, and the final display white point.

Here is a summary of these transforms. The colours are for illustration only, and should not be considered accurate...

<i>Notes</i>	<i>Image view</i>
<p>Here is the original HD data as viewed on an HD monitor.</p> <p>The video data is ranged 0-1, so the white patch will look white in this document. The video white is normally matched to D65.</p>	
<p>HD data converted to log, then viewed on an HD monitor. The log data is also ranged 0-1, so the white looks the same. This has no real meaning, as log images should not be viewed as HD. The grey scale is not quite neutral, as neutral patches do not have equal RGB values in log space.</p>	
<p>Here is the log data, viewed using Truelight on an HD monitor.</p> <p>Truelight is matching the cinema white in hue and luminance. The cinema white will typically be half the brightness of the HD monitor, and more yellow.</p>	
<p>The log data, viewed using Truelight on an HD monitor with matching brightness.</p> <p>In a dark cinema, your eye will adjust to the differences in brightness. This looks similar to our original image, but the cream coloured cast should match the projector white in the cinema.</p>	
<p>The original HD data, viewed using the input gamut clip.</p> <p>This converts the HD to log, clips to gamut, then converts back again. The pink patch above the red gets slightly lighter because of the gamut clip. The white is in gamut, so it stays the same. We do not simulate the cinema white point.</p>	
<p>The original HD data, viewed using the input gamut alarm.</p> <p>This is similar to the gamut clip, except that all in-gamut colours are made neutral. This shows that we cannot match the HD colour on the pink, red, and yellow patches on film. We do not simulate the cinema white point.</p>	

You can view HD data with the correct white point on your monitor using **inputDisplay{}**. You can convert the images from HD to log as outlined at the beginning of this section, then view the log images via Truelight.. Another option is to use the **outputDisplay{}** command. Here is our example profile, modified to show the effects of clipping to gamut on a local monitor...

```
title{HD gamut clip}
recorder{myRecorderCalibration}
print{myVisionCalibration}
display{SonyHD}
cubeType{inputGamutClip}
outputDisplay{monitor}
greyValue{445, 445, 445}
greyStatusA{1.09, 1.06, 1.03}
aimGamma{0.9237, 1.0, 1.0}
whiteValue{685, 685, 685}
invert{on}
```

Here, the **SonyHD** white point is mapped onto the **monitor** white point. This does not guarantee that all the colours on **SonyHD** can be displayed on **monitor**, but the two displays have similar RGB primaries it is likely that most of them will. The monitor white point will not necessarily match your HD monitor, but the relative colours should be accurate. Even if your local monitor is a Sony BMV monitor, there are variations between individual HD monitors that you may want to correct for. However, the RGB values in the image for your display will no longer match the generic HD values: if you want to measure the HD RGB values in the viewer, then you will need to turn off the **outputDisplay{}** command.

5.8 Using an RGB cube on a different display

Suppose you have a cube that transforms from film RGB to **sRGB**. You want to view film material transformed by the cube, but using a non-standard display. Your profile might look like this...

```
title{sRGB cube}
importCube{suppliedCube}
display{sRGB}
outputDisplay{monitor}
```

The cube is just a 3D transform - it contains no record of what the destination colour space is. However, you are sure about the output, then **display{}** and **outputDisplay{}** can be used as shown to convert the cube output.

The **importCube{}** command means our cube has been read in, so all of the usual L*a*b* controls and cube generation options ought not to work. In this special case, we re-use these parameters on the cube data. The gamut alarm and gamut map options will have an effect, and the brightness and shift controls will change the white point.

5.9 Rendering DCI XYZ' output

Some colour tools such as the gamut alarm do not work in DCI XYZ' space. Other tools, such as the cube RGB edits do not work well in spaces other than RGB. It is often easier and safer to work in the display RGB, and convert the output space at the last moment. This approach is also a good thing to do for other standard colour spaces such as sRGB, even though the RGB tools may work for them.

Your DCI XYZ' profile might look like this...

```
title{DCIxyz output graded on Projector}
display{Projector}
outputDisplay{DCIxyz}
greyValue{445}
whiteValue{1023}
greyStatusA{1.09,1.06,1.03}
displayLab{on}
editA{in{0.473,0.454,0.219}out{0.473,0.346,0.219}r{0.26}e{3.6}}
```

The **outputDisplay{}** command transforms the cube RGB values from **display{}** device RGB to **outputDisplay{}** device RGB after all the other commands have been processed. If we put **display{DCIxyz}** instead of using **outputDisplay{DCIxyz}** then our images may not look the same. The **editA{}** command would be operating in DCI XYZ', so you would have to transform the **in{}** and **out{}** colours to DCI XYZ'. There may be other more subtle effects if the **display{}** and **outputDisplay{}** gamuts or white points are different.

In this example, the **displayLab{on}** command forces the **display{}** transform to $L^*a^*b^*$. The **outputDisplay{}** command then maps the white corner of the **display{}** device gamut onto the white corner of the **outputDisplay{}** device gamut. The DCI gamut white point is $X=Y=Z=52.37 \text{ cd/m}^2$, so this may not be the white point you expect.

If we remove the **displayLab{on}** command, and our **Projector** calibration was in XYZ, then **outputDisplay{}** will try for an absolute XYZ match. A DCI projector should be able to achieve a white of 48 cd/m^2 . If your display white is dimmer than this, then your highlights may clip.

There is no easy way to shift the output white point without affecting the cube edits. You should adjust your display settings to give a good working white point before you set up your cube edits. This is a good thing to do anyway, as it makes best use of your display gamut.

5.10 Using black and white film stocks

Truelight handles black and white film stocks just like any other stock. However, because black and white stocks always print as neutral, this can put some restrictions on what Truelight can sensibly do.

Clearly, **invert{on}** is not going to give predictable results: we would not expect to take a black and white display image, and reconstruct the original colour. Less obviously, **greyStatusA{}** may give unexpected results because Truelight cannot find a set of status M values that give the target densities. If you need to fit aim densities, we suggest using the printer point controls. If you are using the Truelight viewer, you can gang all three sliders together. If you set **testValue{}** to the **greyValue{}** settings then the densities get printed out to the log file or the terminal session. If you vary the printer point settings, you can fit the target status A values without the problems you may get from using **greyStatusA{}**.

If you are using black and white print, then you may want to consider using the **Callier{}** command. Black and white film scatters much more light than colour film stock, so the projected image will have considerably more density than the densitometer will measure. Something like **Callier{1.2}** may give sensible results.

Modern films that appear to be in black and white may actually use colour stocks. Modern colour stocks have continued to develop, whereas monochrome stocks have not. Labs are more familiar with handling colour stocks. Strange as it may seem, you may get better black and white results on colour film these days than true black and white, though the images will be less permanent.

5.11 Going beyond the profile

Suppose you have a project that was printing on Fuji 3513 film, and in the middle of the project it was decided to use 3521 instead. The Truelight profile only supports one set of film processes, so we cannot convert from one film output to another within the profile. This shall be our example of using the **tl-utils** command line utility.

If you run **tl-utils** without arguments, it will give you a list of the possible commands. There are specific help messages on for each command.

Our calibrations can be found in the Truelight root directory. The Truelight internal transforms are not stored as files. In this example we shall use internal transforms. We can get these as files using **tl-utils...**

```
> tl-utils Internal recorder ArriInterneg arri.lis
> tl-utils Internal print Fuji3513 3513.lis
> tl-utils Internal print Fuji3521 3521.lis
```

We can merge the recorder and print transforms into a single transform...

```
> tl-utils Merge arri.lis 3513.lis arri+3513.lis
> tl-utils Merge arri.lis 3521.lis arri+3521.lis
```

These transforms take in log RGB, and output status A density. We want to get a transform that takes in our original log values for 3513 and puts out the log values that would give us a matching print on 3521.

We can do this as follows...

```
> tl-utils Invert arri+3521.lis temp.lis
> tl-utils Merge arri+3513 temp.lis 3513_to_3521.lis
```

What happens when we pass our 445 grey through this transform?

```
> tl-utils Test 445/1023 3513_to_3521.lis
445/1023 -> 428.737,481.294,416.323
```

The second line is the reply from **tl-utils**. The 445 value is divided by 1023 because 445 is a 10-bit value.

We may want our 445 value to stay the same. We can do this by shifting the transform output using a formula...

```
> tl-utils MergeFormula 3513_to_3521.lis \
"#0*1023+(445-428.737),#1*1023+(445-481.294),#2*1023+(445-416.323)" \
/usr/fl/truelight/transforms/3513_to_3521
```

The command is entered as one long command line. In UNIX, if you end a line with a backslash as we have done here, then the command shell knows to wait for more input on the next line. The quotes around the formula stop the command shell from interpreting it.

What happens when we pass our 445 grey through our new transform?

```
> tl-utils Test 445/1023 /usr/fl/truelight/transforms/3513_to_3521
445/1023 -> 445.037,444.994,445.023
```

Our output value is not exactly 445, but it is pretty close.

We have put the new transform into the Truelight **transforms** directory. We can apply this transform to image data by making a profile containing **transform{3513_to_3521}**.

This is not necessarily the best way of doing this job. Instead of keeping the 445 grey the same, we might chose a white value that actually occurs in the list. In this case, if we knew the number of the entry, we could use the **tl-utils Zero** option on our two separate half transforms instead of using the complicated formula. The choice of white point is also important - a white of 840 might work well, but the white of 1020 would be close to the edge of the film gamut where all the points bunch up, so the inverse transform might be unstable. If you are doing something like this for yourself, you will need to experiment.

This is just a brief introduction to what is possible. We have not explained here why some values are multiplied or divided by 1023 while others are not, or how the formulas work. There are more details on the inner workings of Truelight in the *Truelight Commands* technical note. The transforms are all human-readable files, so you can look at the data after each step and check at each stage that you are getting what you expected.

6 Debugging

Truelight is complex. This complexity allows it to deal with the many ad-hoc working practices we find in the field. It allows us to combine the calibrations we have in different ways to simulate other similar working practices. This is what distinguishes Truelight from the competitors. However, it does give us great scope for making mistakes.

The right thing to do is to follow the flow of the data through Truelight, and find out by experiment where the mistake has occurred. There are too many controls in Truelight to expect to get to the right solution by guesswork. However, it is hard to work methodically with people looking over your shoulder and pressuring you for results. If you are in this situation, read on...

6.1 Notes on Comparing Images

Truelight should match your display to your film image. In the end Truelight will pass or fail on the quality of the visual match. However, comparing two images is not always as trivial as it might seem. There are many factors that affect the appearance of an image to us that cannot easily be measured with instruments.

Two copies of the same image on the same monitor can look surprisingly different if you look hard enough. When comparing two images on the same display, try and get the images in the same part of the screen, and flip between them. If this cannot be done, then try and get them at the same scale and note the differences: then swap the positions and note the differences again.

When trying to compare different displays...

- Make the room as dark as you can. If possible, surround both images with black.
- Match the brightness of the two images.
- Match the size of the images, or if you cannot do that, find a position where the images appear the same size.
- Keep the CRT screen perpendicular to your line of sight. If the screen is at an angle of more than 30 degrees, you will be looking through more glass, and your image will appear bluer.
- If you have small patches of flat colour on your images, it may help to block off the rest of the displays when comparing them.
- Use instruments such as spot meters or spectrometers if you can.
- Remember that a film will bleach if left in a projector too long. Kodak recommends limiting colour slide projection times to 15 seconds. If you have a loop of 100 frames of film, this probably equates to about 20 minutes display if the film keeps moving.

6.2 Check Truelight is doing something

Try turning Truelight on and off in your application. Try changing the profile. Do you see any effect? If the controls seem unresponsive, then it may be that Truelight is not refreshing the display, so what you see no longer matches the control settings.

Check for error messages. The Truelight library will return an error if it is unable to find some parameter or file that it needs to build a colour conversion. A well-written application should let you see these error messages. The Truelight viewer displays error messages on the status bar underneath the image.

If Truelight does not seem to be doing anything, look for:

- Missing monitor calibration. The **monitor** calibration does not come with the install but has to be generated.
- Over-zealous file protection on the Truelight files or subdirectories. These are probably in **/usr/fl/truelight**.
- Over-zealous file protection on the serial ports when driving Truelight hardware or probes.
- Expired full Truelight licence. This will stop loading of recorder, print and lamp calibrations. You can check the Truelight licensing with **tl_licence**.
- Incorrect Truelight root directory. Try copying a profile to another name. Restart the application. Can the application see the profile under the new name? If not, you are probably looking at the wrong version of the root directory.

6.3 Use the Truelight Viewer

Your application may use custom hardware, special software, or do the colour conversion in the display board as a texture map. The fault may lie in the application, the hardware, or in the Truelight library.

The Truelight viewer is a static image viewer with the ability to generate cubes. It can also drive the Truelight hardware via ethernet. It defaults to doing the rendering in software: this is slower but should give the same results on all machines.

You can call up the Truelight utility with the command...

```
truelight
```

If you are not using the Truelight hardware, start the Truelight application, select an image using the **Image > Open** menu, then select a Truelight profile using the **Profile > Open** menu. Put up the same image and select the same profile in your other application. The two images should look the same.

If you are using the Truelight hardware, you can check the stored cube looks the same as one generated by the Truelight viewer. Connect the hardware to your computer via ethernet. Display an image via the Truelight hardware. Use **Tools > Truelight Hardware** to bring up the viewer hardware dialogue. Select the **Ethernet** option, and enter the IP address or device name you're your sysadmin). The viewer should connect to the hardware. The hardware has a number of 'slots' that are used to store cubes. The default slot should be the one for the current display. Make sure this is so: if you write to any other slot you risk overwriting valuable data. Press **Go** to read the cube to this slot. The cube will take a second or so to download. If the cubes are different, the display should change.

Your application should match the Truelight viewer. If there is a difference, try the following...

- Disable any hardware acceleration in your application.
- Check the application is using the same Truelight libraries
- Check the application is using the same Truelight root directory

If the image in the Truelight viewer matches your utility but still does not look right, then you may have a Truelight problem. We suggest you continue your investigation using the Truelight viewer. This has many useful diagnostic features that are probably missing in your application, and it makes it easier for us to help over the telephone or e-mail.

6.4 Check the Calibrations are valid

Using the Truelight viewer, open an image. If you are using the Truelight hardware, the viewer image will be on the wrong display for your profile, but this should not matter for these tests. Then open the profile. If the profile and the calibrations can all be read, then the image appearance should change. If there is an error, then the image appearance will stay the same, and you should get an error message on the line beneath the image.

The **monitor** calibration is the most likely one to be missing, as this is not supplied as standard.

If there is a problem with a profile, open the **Profile > Edit** dialogue. You can then edit the current settings. This will not change the profile unless you use the **Profile > Save** option. There are browser options for the recorder, print, lamp, and display calibrations. Try picking a calibration that exists, and see whether the image is re-rendered.

If your image is rendering correctly, you might try substituting some of the internal calibrations to see if they make a significant difference:

- We would not expect to see a significant difference between a customer ArriLaser calibration and the internal ArriLaser calibrations if you set **aimGamma{}**. If you are not using an ArriLaser recorder, then you will not have an internal transform to test against, but you could test against an earlier transform if you have one.
- If there is an internal calibration for your print stock, then we would not expect to see a change if you switch to it, provided you are fitting to a grey patch using the **greyValue{}** command. If you are not fitting to a grey patch, then the result will depend on the match the lab was working to at the time the calibration was made.
- If you are using a custom calibration for a xenon arc projector, try switching to the internal xenon lamp calibration. So far, we have not seen a large difference between one xenon lamp and another. The only time we have seen a significant difference was when there was an error in the custom lamp calibration.
- There are no internal display calibrations. Try switching to the supplied SonyHD calibration. If image artefacts disappear when you switch to SonyHD then your display calibration may be faulty. The SonyHD calibration should not give the right colours on your monitor, but it will usually give a good looking image.

6.5 Check the Profile Commands

The **Info > Current Profile Commands** dialogue lists the current commands as they would be saved if you used the **Profile > Save** option. These commands should be the same as your original profile, but...

- The commands may be in a different order.
- The commands may be formatted differently. Most whitespace is ignored, and the order of the arguments is not significant.
- Comments are skipped. Anything enclosed in curly brackets not following a command name is assumed to be a comment. You can easily disable commands without deleting them altogether by enclosing them in brackets using a text editor.
- Commands with default values may be omitted.
- A few commands such as **greyValue{}** are inserted even if the values are defaulted in the profile.
- If we change a command – **cubeFormula{}**, for example, was re-named **cubeLutFormula{}** - we may read the old form of the command, but we will display and save it in the new form.
- The **profile{}** command can be used to include one profile in another. This command is executed as the profile is read in, so it does not appear in the listing.

If one of the profile commands cannot be read, you should get an error message at the bottom of the main Truelight application window.

If Truelight does not recognize a command, check the command is spelled correctly. Check the **Info > Truelight Library Info** dialogue. This should list all the commands that the current Truelight library recognizes. The first line should give the current library version and build number. The build number will be 'unknown' if you are not using a standard release library.

If a command used to be recognized, but is not recognized now, check your Truelight licences. Some commands require additional licence settings to work, and these may have expired. Check the documentation – you may be using an old form of a command that is no longer supported.

If you do not get an error message, but many of the commands seem to be missing, look in your profile for an extra closing curly bracket in the last parsed command. This will cause the profile reader to prematurely assume we have reached the end of the profile information.

6.6 Check the Profile Edit Settings

If the image looks strange when viewed on the correct display, or does not match the reference, open the **Profile > Edit** dialogue, and check the following...

- The image may look dim. This may be correct: the default brightness is 16 foot-lamberts, which matches typical cinema conditions, but is probably a lot lower than the normal setting of the monitor. This may make the image look dim if you are viewing it next to other windows on the screen, or if the room you are viewing it in is not dark. Open the **Tools > Colour Patch** window and set all the sliders to 1023. The figures on the right are the display RGB values, ranged 0-255. Try increasing the white luminance setting until one of these values becomes 255. This will be the brightest image you can get with the same white point without clipping. If the luminance control has no effect, you may have a L*a*b* display calibration.
- The image may seem to have a colour cast when compared to your reference projector. This may be correct: the default white point u' and v' will match D65, and not your projector. Change these controls until you get a visual match on an image. If possible, check the u' and v' values against instrument readings. We would not expect them to be exactly right, but we would hope the values would be within 0.01.
- The default saturation setting is 1.0. We have noticed when we get a match between a CRT and a projection using instruments, viewers think the CRT image is too saturated. We can compensate for this effect by setting the saturation to 0.95. However, the saturation setting can be affected by other settings. If you find you are wanting to set the saturation outside the range 0.95-1.0, then something else needs fixing.
- The default contrast setting is 1.0. Values less than 1.0 can be used when displaying film images on CRTs and other displays with significantly less dynamic range than film. Values above 1.0 are occasionally used when displaying film images on HD and other monitors with a greater range when we want to keep the blacks to a certain value. If in doubt, set the value to 1.0. This should allow you to match all the colours except for the extreme shadows.
- The **title{}** setting has no effect on the colour calibration, but it does give us a useful string to identify the transform. This string is copied into the hardware cube header, so changing the title string should force a re-render.
- The **Truelight Command** prompt at the bottom of the **Profile Edit** dialogue allows us to enter any Truelight command as a text string, followed by the return key. This allows us to enter Truelight commands that are not supported by the regular user interface. If you want to remove a command in the **Current Profile Commands** listing including those supported in the UI, type the command name here, followed by an empty “{}” argument string.

6.7 Check the Viewer Truelight Controls

Most of the profile controls are likely to remain fixed for the lifetime of a project, so the average Truelight user should not need to touch them. The **View > Truelight Controls** dialogue has a small set of controls including Printer Points, Brightness, and Flare correction that the average user may need.

- The default printer point settings are 0.0 in each channel. You can set the printer lights by eye if you have correctly set the white point of the lamp. If you are not sure, then the best option is to leave the controls at the default settings.
- The default brightness is 1.0. The default a^* and b^* shift are 0.0. These controls have similar effects to the white $Y_u'v'$ controls, but they operate in $L^*a^*b^*$ space, rather than XYZ, so they will work on both $L^*a^*b^*$ and XYZ display calibrations.
- The display flare correction default value is 0.01. The reference projection flare correction value is also 0.01. If both the flare correction values are the same, then no flare correction is needed. If you are viewing the display in a well blacked-out room, no flare correction should be necessary.
- Any **cubeType{}** options should be turned off.
- You can set all of these controls to their default settings using the **Reset Truelight Controls** button at the bottom of the dialogue.

It is unusual to save these settings in a profile. If you try to do so, the Truelight utility will pop up a warning. However, there are circumstances where this may be the right thing to do.

6.8 Check Other Controls

You may find Truelight controls we have not described yet. Check for these in the **Info > Current Profile Commands** window for a current listing. Many of these are not supported in the UI, but can be accessed via the **Profile > Edit > Truelight Command** text entry.

importCube{} transform{} formula{}	These replace the normal Truelight transform with the supplied cube, transform or formula, overriding all film and display parameters.
inputDisplay{}	Used to match one display to another. Overrides all film parameters.
projectionWhite{} projectionFlare{} projectionAdapt{}	Reference flare correction values. These should be left at the default values unless there are exceptional projection conditions.
whiteValue{}	If this is set, it fits the display white to the supplied Cineon RGB. If you know your recorder cannot get up to 1023, then put in a lower value. The reference white 685 is a common value. Test the effect of fitting the white point by turning it off.
displayWhite{} displayAdapt{}	The other two parameters in the display flare model. The flare correction is only weakly dependent on the white luminance level, and the adaption parameter should be left at the default value for a black surround unless there are exceptional viewing conditions.
videoBlack{} videoWhite{}	Used to get video output to look right on a scope. If they are set properly, they should have little effect on the image. Remove these commands while debugging.
brightness{} shift_a{} shift_b{}	These parameters shift the L*a*b* whiteprint. They do the same sort of thing as the Yu'v' controls but it works with both XYZ and L*a*b* display calibrations. You can use both sets of controls in a profile.
fitGamut{}	Default setting is on . This is usually best.
output{}	This can be used to redirect the Truelight terminal output including the test value listing to a file, the standard outputs, or to null. Unusual.
invert{}	This inverts the Truelight transform, transforming display values to Cineon values. Some of the cubeType{} options are only appropriate for forward transforms, some are only appropriate for inverse transforms. Other cube settings may not work as expected with inverse transforms.
profile{}	It is possible to call one profile from another. You might use this to make a profile that was the same as another one but with the gamut alarm turned on. Unusual.
enableMask{}	Reduces the Truelight licence privileges. Used for testing only. Delete.
valueFormula{} LabFormula{} rgbFormula{}	Formulae applied to the Cineon, L*a*b*, and display RGB values as the Truelight transform is generated. Historically used to correct problems by hand. Video ranging formulae should be moved from valueFormula{} and rgbFormula{} to cubeInFormula{} and cubeOutFormula{} . Try commenting out any other formulae you do not understand – they may be a fix that is no longer needed.



videoBlack{} videoWhite{}	Used to make video data look right on a scope. Delete these commands and get the visual match right first. If the white and black points are close enough, you should then be able to put these commands back without significantly changing the image appearance.
cubeType{}	Used in viewer for gamut alarms and gamut clips.
cubeHardware{}	Option to generate cubes for the Truelight hardware. Remove when testing.
monotonicCube{} linearizeCube{} intervals{} lutLength{}	Cube generation options. Set linearizeCube{off} if you are using cubeLutFormula{} unless you are sure you want to do both.
cubeLutFormula{} cubeInFormula{} cubeOutFormula{}	Used to reposition the cube intervals and re-range the cube values after the Truelight transform has been made. Set linearizeCube{off} if you are using cubeLutFormula{} unless you are sure you want to do both.
cubeInput{}	This performs a fixed transform on the log cube look-up tables so we can accept Kodak linear or Kodak Video input. This was written as a tool for Shake users. It is unlikely to be used in a profile.
edit[A-H]{} editCube{}	The edit[A-H] commands are used to correct colours using a graphic user interface. When trying to solve problems, turn off all such edits with editCube{off} .
cubeFile{} listFile{}	Save the Truelight cube to a file. Should work for hardware cubes, but may need special license permissions for software cubes or lists. Remove when testing.
dumpProfile{} testValue{}	Used for debugging. Remove when done.

Truelight will probably always have other commands that escape the documentation. Some will be under development. Others will be for special customer sites. Check the **Info > Truelight Library Info** for details. If the command exists and you are licensed to use it, then it ought to appear in the listing. Try removing it. It may be an old command or a setting that is only used under special circumstances.

6.9 Check the Aim Gamma Settings

Truelight uses the aim gamma settings to fit the recorder calibration to the correct Cineon aim gamma. Normally, the aim gamma of a recorder should not change. You can turn off and on the **aimGamma{}** settings from the **Profile > Edit** menu command tool. If the recorder was calibrated when the test strip was produced, then the recorder calibration should already fit the aim gammas, so turning off the gamma fitting should have no effect.

It is important to check that the test strip was produced with the same recorder settings. Check that:

- Your images are the same size as the test strip images.
- The output fill factor is the same as the test strip.
- The image file format is the same as the test strip images. Some recorder driver software assumes .cin and .dpx files contain Cineon data, but .rgb and .tif files contain video RGB.
- Some tube recorders need multiple exposures to get the density up in the red channel. The number of exposures should be the same.
- The film stock and batch number is the same as the test strip.
- The film processing is the same as the test strip.

If you can use an image that was output at the same time as the test strip, then most of these will automatically be true. If you get really stuck, then making a new calibration strip with images may be worthwhile. However, to start with, you will have to work with the images you have.

If turning off the gamma fitting makes a significant change, then you probably have the wrong aim gamma settings. In the short term, you can turn the gamma fitting off.

In the longer term, you should try and determine the aim gamma. Unfortunately, many people who use film recorders do not understand the aim gamma setting. It is sometimes helpful to ask how they routinely calibrate the film recorder. You will probably find they output a set of grey patches, and then measure them on a densitometer, and the software generates a look-up table that serves as the monitor calibration. If the exposure values are linearly related to the resulting negative density, then you can calculate the aim gammas. If not, leave the gamma fitting turned off.

The **tl-plot** utility has a gamma slider. This can be used to fit gamma values to recorder profiles.

6.10 Check the Grey Patch Values

You should record out a standard grey patch with each job. This gives the processing laboratory a standard patch for their print density aims, and it provides a useful diagnostic if something goes wrong. If anyone does not routinely output a test patch, we urge them to do so. If you have a profile and an image selected, you can bring up a grey slate the same size as your images and with the right text with **View > Enable Grey Slate**. You can then save this image for output by turning off the profile with **View > Disable Profile**, then saving with **Image > Save As**.

The grey patch Cineon values are set with the **greyValue{}** command.

The grey patch Status A aim densities are set with the **greyStatusA{}** command.

By default, Truelight puts in a shift in negative density to force the Cineon values to give the corresponding Status A densities.

You can turn off and on the grey patch fitting by changing the **greyStatusA{}** setting on the **Profile > Edit** menu command tool. If you have the right values set, and the print was close to the aim densities, then turning the grey patch fitting off and on should have little effect. This is a useful check that you have set the right values. It is important to get these settings right.

If you have a set of images with a grey patch, you have measured the grey patch in a densitometer, and the grey patch is not on the aim densities, the normal thing would be to add a printer point offset to correct for this. However, you can also set the printer points to zero, and enter the exact Status A densities off the grey patch using the **greyStatusA{}** command. This is a good way to set the exact patch densities before setting the display white Yu'v'.

6.11 Check for HD video scaling

Computer graphics 8-bit RGB images typically use the whole range 0-255. Legal 8-bit video RGB is restricted to the range 16-235. Video values outside this legal range have little practical use. Most people using 4:2:2 video work with 0-255 RGB images, and convert on output. People who have 4:4:4 video equipments, on the other hand, often use it without legalization. If you are using hardware that converts between computer graphics and video formats, check whether it is scaling the data for you.

If you are converting computer graphic images to video, set up your HD monitor using a PLUGE test image. Then make a digital PLUGE test image with stripes of 9, 16, and 23 grey. Display this image on your HD monitor. If there is no ranging, then this should show the same contrasts as your original PLUGE test. If the RGB values have been re-ranged, then the contrast between the 9 and the 16 will be visible.

You can also test for scaling by comparing the video and computer graphic RGB values directly if you have the appropriate equipment to sample the video values. The Truelight hardware has a scaling option on input and output.

6.12 Check the Display Calibration

The Truelight calibration utility is used to calibrate displays. If you think your display calibration is in error, try doing it again. How this is done is described in another note, but here are some useful tips:

- The monitor calibration seems to work best with a display contrast of about 2000:1. Too little range will give problems, but too much range can give problems too.
- If you have spent a long time setting up the display using the first half of the probe menu, go back and reconnect to the probe. This will repeat the black point measurement. A small drift in the black point may give a significant shift in the lowlight measurements.
- If you are worried about the calibration, try repeating it without touching the display controls, and saving it with a different name. Go into the Truelight utility, and try switching between the two display calibrations. You should not be able to tell the difference between the two calibrations.
- If you have trouble calibrating the RGB tone curves, then try calibrating just the tone curves using the by eye calibration. If you use your previous calibration as the template for this one when you save, you can keep your previous set of RGB primaries. With care, you can get good results even with poor displays.
- If you are calibrating a display via the Truelight HD video hardware, test the range of your data. You can sample the RGB values in the hardware with the **Tools > Truelight Hardware** cursor control in the Truelight utility. Put up a set of grey patches with known values, then check the values using the cursor. If you get the same values, then the HD conversions between RGB and YCrCb are not shifting the black and white point. If your values are shifting, there are only two likely shifts – either the video 16-235 is being mapped onto the range 0-255 in the hardware, or the shift is happening the other way.
- If your display may not have independent primaries, try doing a full 3-D calibration. This will measure the colours over the whole of the colour space, taking about 30 minutes. The resulting calibration is a list rather than a formula. You can try switching between the 3D calibration and the original one in the Truelight viewer. If the difference is small, use the 1D calibration – it should be smoother.
- If you have set up a 500:1 contrast between grey and white using the **Interactive Setup** option, it should be close to SonyHD. Try comparing your monitor calibration and SonyHD using **tl-plot** or the viewer.

6.13 Test Values

You can get the display RGB values for a given set of Cineon values using the **Tools > Colour Patch** utility. This can be a handy way to explore which colours are coming out correctly, and which ones are not. However, this only gives you the values before and after the Truelight transform. If the output values are clipping, then you may be able to spot this from the output values.

If you cannot see what is going wrong using the colour patch utility, you can get more details of the internal Truelight values for a particular colour printed out on the terminal using the **testValue{}** command. Run the Truelight viewer from a command line in a terminal window, not from the icon. With **testValue{445}**, you should get a listing on the terminal window something like this...

```
test value: rgb = 445.0 445.0 445.0 (0-1023)
test value: neg = 0.69 0.79 0.83 (statusM)
test value: pos = 1.10 1.05 1.01 (statusA)
test value: Lab = 43.2 0.3 -6.2 (film)
test value: Lab = 41.5 0.0 -5.3 (display)
test value: XYZ = 1.92 2.03 2.60 (ft-L)
test value: Yuv' = 2.03 0.1917 0.4548 (ft-L)
test value: RGB = 75.4 87.8 103.3 (0-255)
```

Here, we can see the Kodak LAD grey 445,445,445 gives a Status A density that is close to the Kodak status A aims of 1.09,1.06,1.03. It also gives us the L*a*b* values before and after fitting the white point and doing the flare correction, the display XYZ values (also shown as Yu'v') and the display RGB.

You can turn off this listing by calling **testValue{}** without an argument.

We can measure the Status M and Status A values off film with a densitometer. Try turning off the grey patch and gamma fitting, and anything else that might affect the internal values. You should be able to get close to the measured values. This should serve as a check for the recorder and print calibrations.

You can read the monitor XYZ or Yu'v' if you have the right instruments. You can use these to check the display calibration. Do not expect a close match unless you are using the same instrument that did the original calibration.

The two sets of white values will differ if you are fitting the white point using **LabFormula{}**, or compensating for flare. Try turning off anything that might affect the Lab values. You should get one set of Lab values when everything is turned off.