

A Standardised Polarisation Visualisation for Images

Alexander Wilkie*
Charles University in Prague
Czech Republic

Andrea Weidlich†
RTT AG
Munich, Germany

Abstract

We discuss issues surrounding the visualisation of the polarisation properties of light stored in the pixels of an image. In order to facilitate comparisons between the work of different researchers, and in order to aid the development and debugging of polarisation-capable rendering systems, we propose a set of four standardised, easily comprehensible visualisations. These cover all aspects of polarisation that are relevant for rendering research, but also for optical design tasks. However, the so-called luminance scaled overlay form of the proposed visualisations is conceivably also useful and instructive to a wider, non-technical audience, and can be used for educational purposes.

CR Categories: I.3.6 [Computer Graphics]: Methodology and Techniques—Standards; I.3.m [Computer Graphics]: Miscellaneous

Keywords: polarisation, visualisation, predictive rendering

1 Introduction

For human observers, the polarisation state of light is an ephemeral quantity. While the human eye does have a very weak capability to detect polarised light [Koe85], this plays practically no role in our overall perception of a visual stimulus. The fact that this capability is almost unknown is evidence of this. However, light is a transverse electromagnetic wave, so its polarisation state is actually a fairly important physical characteristic. This is particularly true from a viewpoint of highly accurate image synthesis, since the polarisation state of a light wave can have significant influence on how it is reflected and refracted by phase boundaries [WUT*04]. Even though we cannot perceive the polarisation status of the resulting light, the intensity of light that is reflected across multiple specular interfaces can, in some configurations, be significantly mispredicted by a renderer that does not take the phenomenon into account.

So far, the topic of the polarisation of light has, to our knowledge, not been explicitly discussed in visualisation literature. It has also only received occasional attention in the rendering community [WK90, HTSG91, TTW94, San97, FGH99, WTP01, DCWP02, WUT*04, WW08]. As a consequence, a lot of the infrastructure we take for granted in normal rendering research is hard to find. There are no widely available rendering toolkits that offer this feature, but also no common image formats which would allow the storage of such information in a rendered image. Nor are there any tools

*e-mail: wilkie@cgg.mff.cuni.cz

†e-mail: andrea.weidlich@rtt.ag

commonly available that visualise the polarisation state of the light stored in an image.

In this paper, we document the polarisation visualisation capabilities of our own rendering research toolkit, which, as far as we know, is currently the only existing, non-trivial rendering engine to support this feature. There have been a few other research systems with such capabilities in the past [WK90, TTW94, San97], but to our knowledge, these systems have either long been abandoned, or were never more than small experimental systems. There are also some commercial systems that support polarisation raytracing, such as [Bre, Opt], but these are mainly focused on optical component, lighting and coating design, and not so much on general image synthesis.

We do not present this functionality of our system in order to claim intellectual precedence for any of the used visualisation methods. All of them have probably, at some point in the past, been used by physicists, biologists, and other researchers that had to deal with the visualisation of polarisation state, usually in non-synthetic images. However, we are not aware of an actual discussion and comparison of these visualisation techniques in computer graphics (or any other) literature, and there are of course also no standards for such visualisations in our field yet.

The goal of this paper is therefore twofold:

1. to briefly discuss the problem of polarisation visualisation and its physics background, and
2. to present one particular set of visualisation conventions for the polarisation state of rendered images

The latter is done in the hope that the conventions we adopted for our visualisations are also used by others who are working in this area. The main benefit of using such a canonical visualisation is course that it becomes much easier to compare results and images. Since highly accurate, predictive image synthesis is a topic of increasing interest in the graphics community, and since this will lead to a renewed interest in polarisation rendering, we assume that work in this direction will become more common in the future. Therefore, adopting such conventions early on would be helpful.

2 Related Work and Background

While for a large number of purposes it is sufficient to describe light as an electromagnetic wave that travels linearly through space as a discrete ray (or a set of such rays), such a wavetrain in fact also oscillates in a plane perpendicular to its propagation. The exact description of this phenomenon requires more than just the notion of radiant intensity, which the conventional representation of light provides, and which is the basis of practically all representations of light in computer graphics.

A good introduction to this phenomenon is given by Shumaker [Shu77], and we refer the reader to this, and other related literature [BW64, Koe85] for an in-depth discussion. The only exception is information on the Stokes Vector formalism itself, which we briefly recap in section 2.2.

2.1 Polarisation Rendering

As far as the implementation of polarisation support in a renderer is concerned, several publications have addressed the question which mathematical formalism, and, consequently, data structures are most suitable for this purpose [WK90, San97, FGH99, WTP01, DCWP02]. Wolff and Kurlander [WK90], who were the first to implement a polarisation renderer, and Tannenbaum et al. [TTW94] both used the Coherency Matrix formalism to describe the transversal oscillation pattern of light. Sankararayanan [San97], Freniere et al. [FGH99] and Wilkie et al. [WTP01] all proposed using the Stokes Vector formalism instead. The main goal of these efforts was to find an appropriate way to describe, and perform rendering-related calculations with, polarised light. Both earlier groups of authors settled for the notation suggested by general physics literature, while the Stokes Vector formalism used by later investigations proved to be easier to handle and implement in a rendering framework.

The two formalisms are interchangeable, can be converted into each other [Alm92], and mainly differ in how easily they can be integrated into a graphics workflow. In this regard, the advantage of Stokes Vectors is that their components are direct correlates of observables; this was the main reason this formalism was (and still is) favored by the optical community. This better comprehensibility also makes it easier to use this formalism in a renderer. In this paper, we specify how one extracts a polarisation visualisation from an image that contains Stokes Vector data. Since the two formalisms are interchangeable, these definitions are of course also directly applicable to Coherency Matrices.




Figure 1: a) Three-dimensional view of the propagation of a perfectly linearly polarised light wave which is rotated by an angle of α from the x -axis. b) Two-dimensional view of the plane of oscillation for this wave. c) The Stokes Vector formalism: four numbers are used to describe the polarisation state of a wave. The first component S_0 encodes the radiant intensity, components S_1 and S_2 encode linear polarisation in two different reference frames rotated by 45° , and S_3 is used to describe the circular polarisation component.

2.2 The Stokes Vector Formalism

For a given wavelength, four numbers are sufficient to describe the polarisation state of a given transversal wave; the entire group of parameters is usually referred to as *Stokes Vector*, while an individual parameter is referred to as *Stokes Component*. This description only uses real-valued terms to describe all polarisation states of optical radiation. It also uses a noncomplex description of ray weights, or attenuation factors, in the form of Müller matrices [Shu77].

Actually, only three real-valued parameters would be required to describe a general polarisation ellipse. The slightly redundant, but very convenient, four value formalism with parameters $\{S_0, S_1, S_2, S_3\}$ has originated, and proven itself, in the optical measurements community. For graphics purposes, it has the key advantage that the first component S_0 of this 4-vector is the unpolarised radiant intensity of the light wave in question, i.e. the same quantity that a nonpolarising renderer uses. Components S_1 and S_2 describe the preference of the wave towards linear polarisation at zero and 45 degrees, respectively, while the fourth, S_3 , encodes preference for right-circular polarisation. While the first component is obviously always positive, the values for the three latter parameters are bounded by $[-S_0, S_0]$; e.g. for a radiant intensity of $S_0 = 2$, a four-tuple of $\{2, 0, 0, -2\}$ would indicate light which is totally left circularly polarised. Figure 1 gives a graphical representation of the individual Stokes Components. Note that the parameters S_1 to S_3 are also under the constraint $S_0 \geq \sqrt{S_1^2 + S_2^2 + S_3^2}$.

3 Visualising the Polarisation of Light

There are two basic options for visualising the polarisation state of light captured in an image:

3.1 Glyph-based

One can overlay **glyphs**, which are a concept with a long tradition in visualisation [War02, LKH09], over the image in question. One possibility for this are arrows or lines to describe the plane and magnitude of vibration of light in a given region of an image. This has the advantage of yielding very intuitive results, and the disadvantage that only approximative visualisations are possible. For the purposes of predictive rendering research, this approach is less useful, since for such work, one is often interested in the state of each pixel, e.g. for debugging, calibration and evaluation purposes. Glyphs are, however, a technique favoured by many in the natural sciences, such as biologists who are just interested in demonstrating how entire image regions appear in the field of view of creatures that are capable of discerning polarisation [Weh01]. The polarisation state of each individual pixel is not relevant for such purposes.

3.2 False Colour

Alternatively, one can use **false colour** to individually visualise the polarisation state of each pixel; this is what we refer to as **pixel accurate polarisation visualisation**. This yields images that are not as easy to read as those where glyphs are used. However, if the visualisation colours are properly chosen, and if a reference is provided, they provide unambiguous information to the viewer, and can be used to accurately analyse the results generated by a polarisation-aware renderer. In various forms, they are of course also popular in the natural sciences; examples of this can be seen in [BGH04, KMSH08]. As mentioned earlier, there are two main tasks that rendering engineers might need such visualisations for:

1. Visualisation of the polarisation in an image for some given external task, e.g. in order to assess, via a false colour image, the likelihood that polarising sunglasses can effectively reduce glare in a given viewing situation. Or to quantitatively show the degree to which a physically plausible BRDF model emulates the polarisation characteristics of real surfaces.

- Employing these visualisations as one of the tools used for verification and debugging of a polarisation renderer. Relevant information that can be gained from such visualisation is whether polarisation occurs in those areas of the image that it ought to, and if it is of the correct magnitude and form. Competent assessment of this requires some experience by the person reading the visualisations, but this is true for many other types of visualisations as well.

It is worth noting that in the second scenario, a pixel-accurate visualisation is the only way to verify that the polarisation information is appropriately consistent across image regions.

Recently, a false colour polarisation visualisation scheme has been proposed [NHHG08] which directly encodes the entire polarisation state of a pixel into CIE $L^*u^*v^*$ space. The results obtained with it are aesthetically pleasing, but the main drawback is that such integrated visualisations can be rather hard to interpret and use for the application areas mentioned earlier in this section. These very often require inspection of individual aspects of the polarisation state, which a single integrated visualisation can not readily provide. The analogy between the (I, Q, U) components of the Stokes formalism and the channels of CIE $L^*u^*v^*$ is an intriguing concept, but the polarisation phenomena one can easily correlate with scene features, such as the presence of linear polarisation in specular reflections, and that are of interest during e.g. the debugging of a polarisation renderer or the assessment of a polarisation-capable BRDF model, seem to warrant a more detailed approach.

3.3 Combined Visualisations

Both methods can of course be combined, as can for instance be seen in [HGPW98]. Since we focus on the predictive rendering aspects of this problem, and are interested in pixel-accurate techniques, glyph-based visualisation is of lesser interest to us, and we also do not deal with combined techniques.

4 False Colour Visualisations of Polarisation

For each pixel of an image that contains polarisation information in Stokes Vector form, we store, in addition to the radiant intensity, three additional signed quantities that encode the polarisation state for each spectral sample. Which means that if we restrict ourselves to the generation of two-dimensional output (i.e. false colour images), we do not have enough independent channels to simultaneously visualise all the information contained in such an image. Appropriate data reduction, and separation of the available information into several well-defined plots, therefore becomes the order of the day.

4.1 Information Reduction and Analysis

It is first worth noting that for the purposes of polarisation visualisation, it is very often convenient to limit oneself to analysing the average properties of each pixel, instead of polarisation information found in the entire pixel spectrum. The polarisation state of the light in a pixel can vary significantly with wavelength even for neutrally coloured light, especially if the light in question originates from a specular reflection. Consequently, colour changes – such as the removal of a white highlight on a layered, varnished surface, to reveal the coloured layer underneath [WW07] – can be induced by the use of polarisation filters. However, as the visualisation proposed by [NHHG08] shows, the inclusion of such




Figure 2: *The test image used for the proposed polarisation visualisations. In the upper image, the scene is shown with Lambertian surfaces, and slightly zoomed out, to clarify the overall set-up. Also, the two green objects were added to the scene to give a reference for the slight tilt of the rear mirror surfaces. The lower image shows the actual test image with specular surfaces. The dark reflective objects are made of black glass. Conveniently for our purposes, this material exhibits linear polarisation in its reflections, but no refractions – these would just be distracting here. The metallic reflector objects are made of gold, which, like most conductors, is capable of giving rise to elliptical polarisation in reflected light. The lighting is uniform D65, and the floor is a neutral grey surface. The arrangement of the reflector shapes was chosen so that they exhibit a large number of inter-reflections that one can see. Some of the inter-reflections appear almost black, but are of course just very dark.*

wavelength dependencies in a visualisation does little to clarify the overall presence, strength and nature of polarisation effects in an image. Derived properties such as the average polarisation can be more useful than the plots of [NHHG08] for certain tasks we are interested in, like assessment of the impact a polarisation filter will have on an image, analysis of the performance of a polarisation-capable BRDF model, or even just for debugging purposes. In the latter case, one of the main requirements is the ability to easily inspect an entire image in order to assess whether the expected types of polarisation are present on the surfaces they should be on, and if any polarisation that is present is of the right type and orientation.

If we average the polarisation properties across channels, this means that the data we wish to visualise for each pixel has just four primary channels, three of which are signed quantities: the unsigned radiant intensity (Stokes Component 0), plus three signed Stokes Components 1 to 3 that encode the actual polarisation state. Apart from the first (which basically is a monochrome version of the original image) and the fourth channels, direct visualisation of




Figure 3: Visualisation 1 – the degree of polarisation (DOP). The more intense the red colour, the more strongly the pixels are polarised. As described in the text, the difference between (b) "direct overlay" and (c) "luminance scaled overlay" is that in case (c), the visualisation colour is additionally scaled with the luminance of the target pixel. For most purposes this type of plot might be used for, the two overlay visualisations are superior to the "raw data" image. Overlay visualisation (b) has advantages in those situations where one wants to clearly see all strongly polarised pixels, irrespective of their luminance; this is useful during debugging (for which even type (a) can have its uses), or to generally inspect the polarising effect of a specular object. Type (c) corresponds to the potential visual prominence of polarised pixels – see figure 8 for a relevant example. Figure 9 shows representative Fresnel reflectance plots for dielectrics and metals, gives a brief background discussion of the DOP distributions seen in this visualisation, and the differences between metals and dielectrics in this regard.

the Stokes Components does not yield particularly useful insights, though. In a false colour visualisation, three unsigned quantities could conceivably be encoded in the R, G and B channels of the result image – or, as in the visualisation proposed by [NHHG08], the three channels of CIE $L^*u^*v^*$ space. However, the Stokes Components 1 to 3 are signed quantities, which means that it is more appropriate to use opponent colours (i.e. red-green and blue-yellow) for each channel. In any given output image, there are only two such channels available, which means that at most two Stokes Components can be directly visualised in a single image.

4.2 Information to be Visualised

Instead of the "raw" Stokes Components, one rather wants to display several derived properties, which are directly useful in understanding the polarisation properties of the light in an image. These are:

1. the overall degree of polarisation
2. the type of polarisation: pure linear or circular polarisation, or the general case of elliptical polarisation
3. the oscillation plane of the linear polarisation component
4. the chirality of any circular polarisation

Also, for debugging purposes, it is handy to have a fifth, additional mode in a image polarisation visualisation tool - one which flags pixels with physically implausible polarisation information. This is primarily useful to check whether a polarisation-aware renderer is operating correctly, though, and is not a visualisation of the polarisation state as such. In section 5, we discuss each of these four (plus one) visualisations, and show examples.

4.3 Stand-alone vs. Overlay Polarisation Visualisations

For each of the four forms of polarisation visualisation introduced in the next sections, it is worth noting that the output can be either viewed on its own (usually with unpolarised pixels set to black), or

overlaid on a black and white version of the original image. The overlay can be created in two ways: in a form that we refer to as *direct overlay*, and one which we call *luminance scaled overlay*.

In both cases, the visualisation colour that is composed onto the black and white image has an alpha component, which is determined by the degree of polarisation of the pixel. This ensures that colour is only applied to polarised pixels, and that unpolarised pixels remain achromatic.

The first type, *direct overlay*, applies this visualisation colour irrespective of the pixel brightness – as can be seen in the examples, this yields information for all areas of the image that contain polarised pixels.

As its name implies, the second type, *luminance scaled overlay*, additionally scales the overlay colour with the luminance of the target pixel. The advantage of this is that the polarisation information is introduced into the image in a very natural way, and that there is an implicit suppression of visually insignificant polarisation states – dark areas of the image might have high chroma values if they are strongly polarised, but this will not be strongly visible.

However, depending on what exactly one is doing, the latter point is also a potential problem – so the type of overlay has always to be chosen according to the needs of the person doing the visualisation. Also, because there is no reliable way to distinguish between the two types of overlay just by looking at a visualisation, it is vital that one always explicitly specify which of the two types of overlay is used, if such a plot is produced.

5 The Four Proposed Canonical Polarisation Visualisations

To demonstrate the proposed visualisations, we intentionally chose a very simple test scene, shown in figure 2, that maximises the occurrence of polarised light in some parts of the image. The image contains much more linearly than circularly polarised light, but this is fairly typical for most scenes. Also, please note that while it is not incorrect to refer to the visualisations described in sections 5.2

and 5.4 as dealing with circular polarisation, the polarisation caused by the golden object in the test image should correctly be referred to as *elliptical* – part linear, part circular. However, we often just use the expression "circular polarisation" in the text in an attempt to avoid unnecessary confusion.

In the spirit of trying to define a standard, we would suggest that not only the semantics of these plots be used, but also the exact same colour schemes that we use here. Not because the particular choices made by us are, in any sense, better than conceivable alternatives, but because some sort of standard is necessary – and, all other things being equal, these particular colours seem to serve this purpose well.

5.1 Visualisation 1 – Degree of Polarisation (DOP)

The overall degree of polarisation DOP of a given electromagnetic wave is determined by the formula [Shu77]

$$DOP = \frac{\sqrt{S_1^2 + S_2^2 + S_3^2}}{S_0}$$

Colouring pixels according to their DOP gives an immediate insight into which parts of the image contain polarised radiation, but gives no indication of its type. We recommend that red be used as the indicator colour, and figure 3 shows an example of such a plot.



Figure 5: Colour coding of the linear polarisation false colour scheme proposed for use in visualisation 3 – the oscillation plane visualisation seen in figure 6. Linearly polarised light with the indicated oscillation plane would e.g. be encoded as yellow. In figure 6, the fringe of the sphere, which in its entire circumference contains linearly polarised light of all orientations, is a good example of what this looks like when used in an image.

5.2 Visualisation 2 – Type of Polarisation (TOP)

The next distinction one can make is whether any polarisation that is present is predominantly linear, or circular. Such a plot does not take the orientation or chirality of any polarised light into account, and just gives an overview of the prevalent polarisation type. The relative degree of linear polarisation $rDOP_L$ of a pixel is given by

$$rDOP_L = \frac{\sqrt{S_1^2 + S_2^2}}{\sqrt{S_1^2 + S_2^2 + S_3^2}}$$

and, complementary to this, the relative degree of circular polarisation $rDOP_C$ is defined as

$$rDOP_C = \frac{S_3}{\sqrt{S_1^2 + S_2^2 + S_3^2}}$$

We propose to use these relative measures, instead of the absolute measures more commonly used in the optics community, because these values seem to give a clearer indication of the polarisation type if only the distinction between the two is of interest. These two factors are only used to blend between the respective colours for linear and circular polarisation, C_L and C_C :

$$C_{TOP} = rDOP_L \cdot C_L + rDOP_C \cdot C_C$$

These resulting colour C_{TOP} is then scaled with the DOP , which gives the TOP plot an overall appearance that is quite similar to the DOP plot shown in section 5.1 – except that the plot colour changes with polarisation type. We recommend that cyan be used for linearly polarised light, and yellow for circularly polarised pixels. Figure 4 shows such a plot.

5.3 Visualisation 3 – Oscillation Plane of the Linear Polarisation Component

The magnitude and direction of any linear polarisation that is present is determined by the intensity and the ratio of the first two Stokes components. If the components S_1 and S_2 are visualised simultaneously using red-green and blue-yellow opponent colour channels, a false-colour encoding like the one shown in figure 5 results.

In particular, for channel S_1 we use the red-green channel, with negative values encoded as red. For S_2 , we use blue-yellow, with negative values shown as blue. As with all our proposed plots, the selected colour is then scaled by the $rDOP_L$, and for overlay plots, the alpha channel of the overlay is also given by this factor.

The "colour wheel" that is created by this simple technique is of course not perceptually uniform. A considerable number of more sophisticated techniques to generate colour wheels exist; [BBS09] gives some pointers on this. But in our opinion, it is debatable whether the additional effort, and the potential technical issues (such as gamut problems, which are not an issue for the simple approach) incurred by the creation of a perceptually uniform colour wheel stand in relation to the visual benefits one might obtain. The comprehensibility of the oscillation plane plot seems to be reasonable, even with our easy to create, non-uniform colour wheel.

5.4 Visualisation 4 – Chirality of the Circular Polarisation

For this plot, one directly visualises the third component of the Stokes vector, and we propose to use the blue-yellow opponent colour pair for this. In such an image, light with a right circular orientation is signified by blue, and left circular by yellow. Figure 7 shows this for our test scene, with the $rDOP_C$ used as the scaling factor for the colour, and the alpha factor.

5.5 Visualisation 5 – Stokes Vector Sanity Check Visualisation

The fifth visualisation offered by our polarisation visualisation application is just a debugging tool, and flags all those pixels for




Figure 4: Visualisation 2 – the dominant type of polarisation (TOP). We use cyan for linearly polarised light, and yellow for circularly polarised pixels. Here, a visualisation of raw data can be superior to the overlay approach if, as it is very often the case in real scenes, only weak circular polarisation, and therefore only weak yellow colouration, is present. If an overlay plot is still desired, overlay type (b) can give a clearer indication of the property one wants to visualise in such cases, and should be preferred over type (c).




Figure 6: Visualisation 3 – the oscillation plane of the linear polarisation component. The false colour scheme shown in figure 5 is used to encode the oscillation direction. Again, the choice of stand-alone or overlay plot seems to depend on the application: the stand-alone plot is appropriate for debugging purposes, since it shows the oscillation direction even in dark areas of the image. An example is the perceptually unimportant (i.e. very dark) reflection on the wall that is purple in this plot; being able to discern this is potentially quite important for debugging purposes. Both types of overlay give a nice visualisation of the behaviour of linearly polarised light on the fringe of the dielectric sphere on the left. The coloured area on the sphere seems to be larger in the stand-alone image because no alpha compositing is done, and weakly polarised pixels retain their colours better. Also, compare (c) with the effect of the polarisation filter shown in figure 8.

which the assertion $DOP \leq 1$ does not hold. This is a hard assertion in the sense that the only reasons for non-conforming pixels to appear in an image are errors in the simulation and/or the input data. We just mention this here because according to our experience, one would not want to write a polarisation visualisation tool without this sort of capability.

6 Conclusion

We have discussed the polarisation properties of visible light, and presented one particular set of conventions to visualise them. Adherence to these conventions would make scientific collaboration about this one aspect of predictive image synthesis easier, and might also benefit other disciplines, if the polarisation visualisation conventions developed for rendering were to be commonly accepted.

In the future, we will investigate the unambiguous definition of additional plots that might be useful for areas outside graphics engineering, e.g. for wavelength-dependent differences in the degree of polarisation.

7 Acknowledgements

We would like to thank the anonymous SCCG conference reviewers for their constructive input and criticisms, and the additional literature references they provided.

References

- ALMEIDA J. S.: Radiative transfer for polarized light - equivalence between stokes parameters and coherency matrix formalisms. *Solar Physics* 137 (Jan 1992), 1–14.
- BRATKOVA M., BOULOS S., SHIRLEY P.: oRGB: A practical opponent color space for computer graphics. *IEEE Computer Graphics and Applications* 29, 1 (2009), 42–55.
- BERNÁTH B., GÁL J., HORVÁTH G.: Why is it worth flying at dusk for aquatic insects? polarotactic water detection is easiest at low solar elevations. *The Journal of Experimental Biology* 207 (2004), 755–765.




Figure 7: Visualisation 4 – chirality of the polarisation. The few areas in the test image which exhibit circular polarisation are those in which linearly polarised light (in this case, polarised by reflection of one of the glass objects) is reflected at least twice in a metal surface. The sparsity of this plot is fairly typical; except under special circumstances (e.g. in a Fresnel rhomb), reflection-induced circular polarisation is a rare phenomenon. Visualisations of it are nevertheless important to assess the plausibility of a given polarisation rendering. As in the case of figure 4, raw data visualisations, and overlay type (b) can have distinct advantages if weak chirality data is to be visualised.

BREAUULT RESEARCH: Advanced Systems Analysis Program (ASAP[®]) product pages. <http://www.breault.com/>.

BORN M., WOLF E.: *Principles of Optics*. The Macmillan Company, 1964.

DEVLIN K., CHALMERS A., WILKIE A., PURGATHOFER W.: Tone reproduction and physically based spectral rendering. In *STAR Proceedings of Eurographics 2002* (Saarbruecken, Germany, Sept. 2002), Fellner D., Scopigno R., (Eds.), Eurographics Association.

FRENIERE E. R., GREGORY G. G., HASSLER R. A.: Polarization models for monte carlo ray tracing. In *Optical Design and Analysis Software, Proceedings of SPIE, Volume 3780* (Denver, 1999).

HORVÁTH G., GÁL J., POMOZI I., WEHNER R.: Polarization portrait of the arago point: Video-polarimetric imaging of the neutral points of skylight polarization. *Naturwissenschaften* 85 (1998), 333–339.

HE X. D., TORRANCE K. E., SILLION F. X., GREENBERG D. P.: A Comprehensive Physical Model for Light Reflection. In *Computer Graphics (ACM SIGGRAPH '91 Proceedings)* (July 1991), vol. 25, pp. 175–186.

KRISKA G., MALIK P., SZIVÁK I., HORVÁTH G.: Glass buildings on river banks as "polarized light traps" for mass-swarming polarotactic caddis flies. *Naturwissenschaften* 95 (2008), 461–467.

KOENNEN G.: *Polarized Light in Nature*. Cambridge University Press, 1985.

LIE A. E., KEHRER J., HAUSER H.: Critical design and realization aspects of glyph-based 3d data visualization. In *Proceedings of the Spring Conference on Computer Graphics (SCCG 2009)* (April 2009), pp. 27–34.

NEUMANN L., HEGEDUS R., HORVÁTH G., GARCIA R.: Applications of high precision imaging polarimetry. In *Eurographics Symposium on Computational Aesthetics in Graphics, Visualization and Imaging* (June 2008), pp. 89–98.

OPTICAL RESEARCH ASSOCIATES: LightTools product pages. <http://www.opticalres.com/>.

SANKARARAYANAN S.: *Modelling Polarized Light for Computer Graphics*. PhD thesis, Iowa State University, Iowa City, Iowa, 1997.

SHUMAKER J. B.: Distribution of optical radiation with respect to polarization. In *Self-Study Manual on Optical Radiation Measurements, Part 1: Concepts*, Nicodemus F. E., (Ed.). Optical Physics Division, Institute for Basic Standards, National Bureau of Standards, Washington, D.C., June 1977.

TANNENBAUM D. C., TANNENBAUM P., WOZNY M. J.: Polarization and birefringency considerations in rendering. In *Proceedings of SIGGRAPH '94* (July 1994), Glassner A., (Ed.), Computer Graphics Proceedings, Annual Conference Series, ACM SIGGRAPH, ACM Press, pp. 221–222.

WARD M. O.: A taxonomy of glyph placement strategies for multi-dimensional data visualization. *Information Visualization* 1, 3–4 (2002), 194–210.

WEHNER R.: Polarization vision - a uniform sensory capacity? *The Journal of Experimental Biology* 204 (2001), 2589–2596.

WOLFF L. B., KURLANDER D.: Ray tracing with polarization parameters. *IEEE Computer Graphics & Applications* 10, 6 (Nov. 1990), 44–55.

WILKIE A., TOBLER R. F., PURGATHOFER W.: Combined rendering of polarization and fluorescence effects. In *Proceedings of the 12th Eurographics Workshop on Rendering Techniques, London, UK, June 25-27, 2001* (2001), Gortler S. J., Myszkowski K., (Eds.), Springer, pp. 197–204.

WILKIE A., ULBRICHT C., TOBLER R. F., ZOTTI G., PURGATHOFER W.: An analytical model for skylight polarisation. In *Eurographics Symposium on Rendering* (Norrköping, Sweden, 2004), Keller A., Jensen H. W., (Eds.), Eurographics Association, pp. 387–397.

WEIDLICH A., WILKIE A.: Arbitrarily layered micro-facet surfaces. In *GRAPHITE* (2007), Rohl A., (Ed.), ACM, pp. 171–178.

WEIDLICH A., WILKIE A.: Realistic rendering of birefringency in uniaxial crystals. *ACM Trans. Graph* 27, 1 (2008).




Figure 8: The test image with a horizontal (top) and vertical (bottom) linear polariser applied before tone mapping; this sort of simulation is one application for images that contain polarisation information. Note that the main effects of the polariser (i.e. the suppression and enhancement of certain specular highlights on the dielectrics) correspond much better to the coloration in the luminance scaled overlay visualisations, e.g. figures 3c and 6c, than to the direct overlays of the properties in question. But also note the subtle discoloration of the golden sphere – the reason for this are wavelength-dependent polarisation effects, which are not covered by the proposed plots.




Figure 9: Plots of the Fresnel reflectance for typical materials – one metal, and one dielectric. For both materials, the solid lines are the vertically and horizontally polarised reflectance components, which are usually averaged in normal graphics use (dashed line). The greater the absolute difference between them, the greater the polarisation that is caused by the reflection. One can see why metals only cause weak reflective polarisation – the absolute difference is lower for them, due to their higher residual reflectance. The DOP visualisation shown in figure 3, and in particular the DOP of the reflection pattern on the dielectric sphere, are an instructive real-life example of the effects of this well-known reflectance behaviour. That the DOP increases as one goes from the center towards the fringe of the sphere, but decreases again once grazing angles are reached, corresponds directly to the increasing, and then again decreasing split between the two reflectance components.