Optical Engineering

OpticalEngineering.SPIEDigitalLibrary.org

Review of visualization methods for passive polarization imaging

Andrew W. Kruse
Andrey S. Alenin
J. Scott Tyo

SPIE.

Andrew W. Kruse, Andrey S. Alenin, J. Scott Tyo, "Review of visualization methods for passive polarization imaging," *Opt. Eng.* **58**(8), 082414 (2019), doi: 10.1117/1.OE.58.8.082414.

Review of visualization methods for passive polarization imaging

Andrew W. Kruse,* Andrey S. Alenin, and J. Scott Tyo

University of New South Wales Canberra, School of Engineering & IT, Canberra, Australia

Abstract. Collected sources for the different types of visualizations used in the field of polarization imaging are not extensive. Here, we survey and review the different visualization techniques in passive polarimetric imaging. Analysis of the methods is done by applying various concepts from the field of visualization. We provide recommendations for choosing a visualization based on the data structure, spatial frequency, and analysis goals. © 2019 Society of Photo-Optical Instrumentation Engineers (SPIE) [DOI: [10.1117/1.OE.58.8.082414](https://doi.org/10.1117/1.OE.58.8.082414)]

Keywords: polarization; visualization; color; color vision; polarimetric imaging.

Paper 181720SS received Dec. 17, 2018; accepted for publication Feb. 21, 2019; published online Apr. 3, 2019.

1 Introduction

Polarimetric imaging is the measurement of the spatially resolved polarization of light. Like spectral imaging, polarimetry adds additional layers of information to traditional intensity imaging. While the spectral measurements often carry information on material or chemical composition, the polarimetric measurements often carry information on surface properties. In order to explore this information, the polarization measurements can be mapped to the realm of the human visual system in a process called visualization. Since the beginning of polarization imaging,¹ researchers have been trying to figure out ways in which to “see” what is being measured. This is indicative of our preference toward using our incredible built-in pattern-finding machine in order to process information. It seems, however, that the fascination has primarily been focused on what we can see rather than how. Thus, in the literature, the topic of visualization is often treated as an afterthought. While many sources include short reviews of a few polarization visualization strategies,^{2–6} no large-scale review has been attempted. In addition, no review has attempted to bring the wealth of knowledge accumulated in the general field of data visualization to the application of polarization.

One way to view the purpose of data visualization is as the process of asking and answering questions between the data references or independent variables, and characteristics or dependent variables. Andrienko and Andrienko⁷ separate these questions into direct and indirect search queries. A direct query would be “for this reference, what is its characteristics?” while an indirect query would be “where are the references with this characteristic?” The effectiveness of data visualization relies on exploiting the pattern-finding processes of the human visual system in order to answer these queries in the most efficient and accurate manner. The visualization can be used to confirm or communicate the presence of known relations or to conduct an exploratory analysis looking for the presence of unknown relations. In the exploratory stage, it can be beneficial to approach

the data using many iterations of different visualizations in order to find those relations.

Snik et al.³ point out that the polarization variables themselves are not always the end goal but are instead often used to produce information that the user cares about. They also discuss how important the users’ expertise in polarization is to designing and constructing a polarization visualization system. While many of the topics discussed in this review are broadly applicable to these situations, we will assume that the user is interested in visually portraying polarization as data. While this review is focusing on data visualization, there are other paradigms that can be applied to polarization. For example, Ratliff and Tyo⁸ describe two methods for seamlessly blending the polarimetric information into the intensity image as a way to depict how polarimetric imaging can produce greater detail than unpolarized imaging. That is, the final image does not differentiate between information from polarized or unpolarized regions. Instead, polarization is used in a way to add detail to the original image that is easily accessible to those unfamiliar with polarization. They offer two heuristic methods to accomplish this, although much more can be done using this idea. Another paradigm, particularly in the field of biology, is portraying as closely as possible the experience of polarization vision using our polarization-blind human visual system. The well-known Bernard and Wehner paper uses this paradigm to relate polarization vision to color vision.⁹ How and Marshall¹⁰ discuss how objects could be detected by an organism’s polarization visual system using a metric of polarization distance, analogous to color difference. While these paradigms are useful for other purposes, they are out of the scope of this review on data visualization.

2 Univariate Colormaps

2.1 Color

Color is one of the most ubiquitous tools for data visualization, given how much of our visual system is devoted to processing it. The multidimensional nature of color vision allows many different ways to encode information, but not all of them are effective. The effectiveness of any given color visualization is a product of the color model in which it was designed. In short, a colorspace is any multidimensional way in which color can be coded: sRGB, HSV, etc.

*Address all correspondence to Andrew W. Kruse, E-mail: andrew.w.kruse@gmail.com

A color appearance model, on the other hand, encodes the color in terms of how it appears. For this review, we will be using the perceptually uniform color model CAM02-UCS.¹¹ In a perceptually uniform space, Euclidean distances between coordinates are meant to represent perceived color differences. In this section, we use both specific terms using this colorspace, as well as the more general terms. When the terms lightness (J'), colorfulness (M'), and hue (lowercase h) are used, they are referring to the cylindrical coordinates of CAM02-UCS. For this review, the terms luminance, brightness, and achromatic channel are all interchangeable with one another, as are the terms saturation, chromatic magnitude, chroma used interchangeably here. Note occasionally we will use terms specific to hue, saturation, value (HSV) colorspace, such as (uppercase) H , S , V . When these are used, they are specific to their definitions given in HSV and not the general concepts they represent. Additionally, the term isoluminant refers to when the achromatic channel is constant.

The simplest, most accessible, and most common form of polarization visualization is the univariate (single variable) colormap. The technique consists of matching the value of the variable to an ordered color array in order to construct a color image out of a 2-D array of values. Every value is represented by a single color in the color array. This method of colormapping is usually a built-in function to many data processing packages and is customizable by the myriad of available colormaps within the package, from online resources, or by creating a custom color array. With the increased availability of colormaps in general, there is an increase in maps designed both poorly and well in terms of supporting the required analysis tasks. Some maps are designed purely for aesthetics, others for function; however, not all maps designed for function have a basis in an accurate perceptual color model. While there are some good examples^{4,12,13} of the use of effective colormaps used in polarization, the vast majority of colormaps used are either the rainbow colormap,^{14–17} which has been theoretically^{18–20} and experimentally²¹ shown to be ineffective for most analysis tasks, or a colormap inappropriate for the data structure.^{22–24} Both reviews by Zhou and Hansen²⁵ and Bernard et al.²⁶ provide a general review of colormaps in visualization.

2.2 Data Types

The literature on colormaps does not show a common classification method, but the main attributes for choosing a particular colormap rely on data type and task.^{25,27} The most common classification on data types used is nominal, ordinal, interval, and ratio,^{27,28} and other classifications include discrete,²⁹ continuous, monotonic, diverging,³⁰ periodic,³¹ and high and low spatial frequency.²⁸ Variables may be classified by more than one of those listed, as they are not all necessarily mutually exclusive.

Discrete versus continuous: Discrete data types allow for a set number of possible values or categories. Continuous data allow for any value within some range. Continuous data can be displayed as discrete by restricting the number of colors.

Categorical: Also known as nominal, this data type has no underlying order. Colormaps for categorical data should maximize the ability to distinguish between categories without following a perceptual order. This can be realized by

selecting a wide variety of distinct hues. One example of categorical data for polarization is target detection, where regions are nominally labeled as target, background, or other.

Interval: This data type has the quality, where the interval between values is important. If the purpose is to represent the intervals between values as equal, the perceptual color difference should correspond to the difference in values.²⁷ An exception to this rule can occur where the difference in values can be represented by nonlinear perceptual differences in order to improve differentiation.³² Interval data can be discrete or continuous.

Monotonic: While data can be organized in an infinite amount of ways, the most common ways are monotonic, diverging, and periodic. These labels describe fundamental structures within the context of interval or ratio data. Monotonic structure is the most common, where values are ordered from low to high. When this order is meant to be observed, the colormap must also have a monotonically increasing perceptual order. For color, this can be achieved via either the achromatic channel (lightness, brightness, and intensity) or by the magnitude of the chromatic channel (saturation, chroma, and colorfulness). Hue is not an effective channel for ordering because it has no inherent beginning or end. Even though some orderings can be constructed, such as the spectrum or green–yellow–red of a stoplight, there is no purely perceptual reason for these constructions. Unlike saturation or luminance, a given set of hues will be arranged by users very inconsistently.³³

Diverging: This structure has values that fall between two extremes with a critical point separating the two regions. Functionally, colormaps for diverging data structures can be viewed as a concatenation of two monotonic colormaps in opposite directions with a common color at the critical point. The critical point often has some physical or mathematical significance, such as the zero value separating positive and negative values. By using a colormap with a diverging structure, the observer is easily able to tell which side of the scale each value is on. The critical point can also be a way to establish dichotomies such as high or low without necessitating that the critical point has any significance.³⁴ In this sense, even data that is naturally classified as monotonic can be reimagined as diverging toward either the upper end or the lower end. This can be very effective for indicating whether values should be classified as high or low by the observer. The drawbacks to this would be oversimplification by imposing such a dichotomy, neglect of the importance of middle ranged values, and difficulty comparing values by their relative difference. This oversimplification can be useful when a more rigorous analysis is not needed, such as in presentations or in a learning environment.

Periodic or cyclic structure has no beginning or end; instead the values indicate a phase, angle, or orientation. When the colormap is listed in an array, the arbitrarily determined endpoints must be the same color to avoid discontinuity.

Spatial frequency: Data structures can also be classified by whether the data contain high spatial frequencies. At low frequencies, there is no issue with relying on chromatic channels alone, whereas the human visual system needs the achromatic channel to discriminate shapes and structures defined by high spatial frequency.³⁵ Thus, colormaps for

such data must include a significant amount of achromatic variation in order to express form.^{27,36}

2.3 Task Types

In addition to the data type, colormaps are most effective when they are designed for particular tasks. Tasks can vary from low level, e.g., reading values, to high level, e.g., qualitative questions, complex pattern recognition, or decision-making. High level tasks may utilize several low-level tasks and are generally more specific for supporting some conclusion. The low level tasks most common for colormaps have been extensively discussed in the literature. They include comparison, identification, localization, and segmentation. Bernard et al.²⁶ provide a large number of metrics in order to quantitatively analyze colormaps for supporting a variety of tasks.

Comparison: This deals with assessing relative differences between data points as well as recognizing the order of the values. Queries that can be answered by the comparison task are “is this value higher or lower than that value?” and “is the value increasing over this region?” For this task to be supported, the colormaps must maintain a linear relationship between perceptual distance in either or both achromatic or the chromatic magnitude channel (colorfulness, saturation, and chroma). Figure 1(c) shows a colormap constructed by linearly increasing both lightness and colorfulness in a single hue. Using this colormap, the

magnitude of the value of the points in the plot can be relatively compared to one another. For example, given two points, if there is a difference between their values, it is immediately obvious which one is greater. The amount of difference will be proportional to the perceptual difference. Figure 1(b) shows a comparison colormap using a diverging data structure assumption. For any value, the user can compare magnitudes from the critical point using the relative colorfulness difference.

Identification and look-up: Identification has observers identify values with a certain characteristic. This task supports description search queries like “where are all areas where the value is the same as this one?” and “which value is the most prominent in this set of data?” Similarly, look-up is a type of identification that allows users to look-up the value of a data point. A look-up task describes the query “what is this value?” or “where are the areas with this chosen value?” Look-up is one of the most prominent tasks because it is a method for visually “reading” the data. In this way, a look-up task would necessitate a key, while identification would not. For color, hue is an effective channel for identification and look-up because, as opposed to the magnitude-based chromatic and achromatic channels, hue is an identity channel without an implicit ordering.³³ It is instinctive to see all green areas as having the same characteristics, and that they have different characteristics to areas that are blue. The same cannot be said for two different gray levels or

Fig. 1 Colormaps for different tasks. (a) Identification and/or look-up task. Each data point is categorized by its hue so that points with similar values can be located. Values are also easily looked up given the key. (b) Combined comparison of extreme and identification of sign. Here, the hue indicates the value of the sign, and the colorfulness indicates how close to either extreme a value is, allowing comparison of the magnitude of extremeness. (c) Comparison task. Colormap increases in both lightness and colorfulness, indicating a clear perceptual progression from low values to high values, allowing comparison of relative distances in value. (d) Localization task for the value 0.5. Here, this value has been given significance so that the values close to 0.5 are highlighted. Data are a random selection of 100 points from a set of 200 by 200 points with random values that have been smoothed with a Gaussian filter to create random clusters.

two saturation levels that differ by the same amount. Additionally, because the hue channel is not inherently ordered, the distinctiveness between hues is not strongly related to the amount of hue difference if the difference is large enough. For example, blue, red, and purple hues are all sufficiently distinct from green that the amount of hue difference does not result in the perceptual ordering of the hues; that is, they are characterized by their hue and not by their proximity to green. This attribute is effective for identification because it perceptually categorizes data and is effective for look-up because its attributes defining characteristics to data can be compared to a key.

Figure 1(a) depicts a colormap supporting identification and look-up. In this, every data point is represented by some hue, which gives a identifying characteristic depending on the data point's value. For any data point, it is easy to identify the other points with a similar value in the midst of all the data points. Additionally, the look-up task can be supported since it is simple to determine the value associated with the point using the key. In contrast, Fig. 1(c), which uses a colormap just for the comparison task, does not support either identification or look-up to a significant degree. Figure 1(b) uses a diverging colormap not only to support comparisons of magnitudes but also to support identifying the sign of the value. Both Figs. 1(a) and 1(b) use hue for identification, however, the identifiers the hue indicate are significantly different.

Localization: This task is targeted to display the location of all values within a proximity to a certain characteristic or value. A query for this task would be "where are all the areas with a value of 0.5?" Note that this differs from look-up in that look-up is more general for all the values displayed, whereas localization only deals with a specific characteristic. Since studies have shown humans have their visual attention directed toward colors with higher intensity or saturation,³⁷ colormaps that support localization have the value or characteristic of importance to be encoded by colors that are particularly more colorful and/or bright. In Fig. 1(d), the localization colormap is constructed such that only values near 0.5 are colorful and are quickly located. A user attempting to locate this value using one of the other colormaps in this figure would take significantly longer and be more prone to error. In Fig. 2, an angle of polarization (AoP) measurement of a daytime scene of the University of Arizona is depicted using a colormap designed to highlight only values close to the AoP of the sky. Thus, areas that are highlighted in red are most likely reflecting the polarization signature of the sky.

Combined tasks: Often, the visualization is desired to be able to handle several tasks simultaneously. Mittelstädt et al.³⁰ provide colormaps for any combination of look-up, comparison, and localization tasks for monotonic and diverging data types. Generally, the more tasks required of the colormap, the less optimal it will be for any individual task, much like how the solution to an optimization problem with many variables would not be optimal for any individual variable.

2.4 AoP Colormap

While the 1-D colormaps for monotonic or diverging types of data are plentiful, few colormaps are available that support periodic data. When AoP data are mapped by a nonperiodic

Fig. 2 AoP of one of the spectral channels taken on ground MSPI at University of Arizona, courtesy of Christine Bradley and Russell Chipman.³⁸ $\lambda = 660$ nm. The chosen colormap is specific for a strict localization task for AoP values that are close in value to the polarization signature from the sky (90 deg). This allows users easily identify regions in which the polarization is likely due to reflection from the sky. Regions in black are below 5% DoLP.

map, the discontinuity at zero is marked by abrupt changes in the pseudocolor image. These abrupt changes in appearance do not correspond to any physical meaning in the data, since the periodic data are recorded in a monotonic array with arbitrary endpoints. Additionally, the perceptual difference between points is exaggerated for angles, where the difference is greater than 90 deg. For example, given a perceptually linear colormap, the difference between angles of 15 deg and 165 deg would have a perceptual difference corresponding to 150 deg, even though they are only 30 deg away from each other in an 180 deg period. Additionally, the difference between 15 deg and 90 deg would only appear to have half of the perceptual difference, even though the actual difference is twice as much as between 15 deg and 165 deg.

Of the periodic maps that do exist, it is not common to find one which takes into account the necessity to support shape discrimination, which is an integral part of spatially resolved AoP data. At high spatial frequencies, human vision is more responsive to luminance changes than color differences, meaning that shape information is more efficiently encoded by luminance than color alone. Often, periodic colormaps are created to be isoluminant with the intention of uniformity in emphasis.^{31,39} While isoluminant periodic maps would be effective for situations where the AoP data have low resolution spatial structure to it, such as the polarization of the sky, it is not appropriate in the situations with high resolution due to the lack of contrast. This is illustrated in Fig. 3, where there are both isoluminant

Fig. 3 AoP in radians of objects on a lab bench visualized with (a) HSV periodic colormap, (b) perceptually uniform isoluminant colormap, (c) perceptually uniform with varying lightness, (d) same as (c) but with variability Δ mapped to transparency with a black background. Data are provided by Samuel Powell and Justin Marshall, University of Queensland.

and luminance-varying periodic colormaps applied to high-frequency AoP data. The most common periodic map is created by cycling through the hues in HSV while saturation and value are at maximum. While the value channel would suggest it to be isoluminant, more accurate color models would predict large variations in luminance.⁴⁰ Additionally, the variations are not primarily a function of the hue difference. Thus, the contrast between any two colors in the colormap does not indicate the difference between the values the colors represent. This means that the impression of the shapes of objects in the image may be defined primarily by the contrast between arbitrary hues rather than by the most significant angular variations. Figure 4 shows how the perceived color difference between consecutive colors in the HSV hues colormap varies depending on the value in the colormap. This figure also depicts the lightness varying nonuniformly across the colormap, with rapid changes at some areas (particularly cyan to blue) and long stretches of isoluminance (yellow to cyan).

A recent solution intended to introduce luminance variation while maintaining perceptual uniformity is to parametrize a circle in CAM02-UCS space that is tilted in the lightness axis.⁴⁰ As can be seen in Fig. 3(c), the shape information carried by the luminance is much greater than in the isoluminant case. The color difference (ΔE) between sequential steps in this colormap is constant. Additionally, the color differences between any two angles are only functions of the difference between them and not the individual angles themselves as in HSV. Unlike the isoluminant colormap, the lightness variation between sequential colors in this colormap contributes to their total color difference, which is just the Euclidean distance between the colors. Since the lightness variation is not constant, the component of the color difference from the chromatic channels is also not constant. This means the hues are not varying linearly with angle. While it is simple to instead restrict the colormap to uniform hues with nonuniform color differences, it is uncertain at the moment which one would perform better

for the variety of visualization tasks if a significant difference exists at all.

While this solution seems to work well in most situations, introducing lightness variation into a univariate periodic colormap creates some issues. Unlike monotonic or diverging data structures which have inherent directions for increasing or decreasing lightness, periodic colormaps cannot introduce lightness variation without simultaneously introducing bias because that breaks the rotational symmetry of the colormap. This trade-off of rotational symmetry in favor of contrast is generally beneficial, as evident in Fig. 3, but it is potentially possible to avoid this trade-off altogether with methods more sophisticated than predetermined one-to-one values-to-colors mappings. To maintain rotational symmetry, the achromatic channel would need to be decoupled from the nominal AoP values. The achromatic channel could then be available for providing contrast and shape discrimination independent of the specific AoP-to-hue mapping. The lightness value of any pixel would be determined such that it promotes the type of shape discrimination the user desires, e.g., corresponding AoP gradients to shading gradients or creating contrast between objects with differing AoP. Currently, no such algorithm has been proposed.

One recurring issue with AoP colormapping is that it has no zero point corresponding to a low amount of signal. That is, the AoP derived from measurements with low signal will be highly sensitive to noise, but without additional processing, the noise will be displayed with the same vibrancy as areas with a consistent signal. This is because unlike other colormaps, periodic colormaps are usually maxed out in terms of colorfulness. The attention drawn to those colorful pixels that are carrying no useful information reduce the amount of attention that could be turned toward the pixels carrying more information. Often, areas below a certain threshold degree of linear polarization (DoLP) are eliminated by setting to black.¹⁰ This threshold, however, may be difficult to set without multiple iterations and does not necessarily conform to the uniformity in the AoP.

Fig. 4 Color metrics for four different colormaps: jet, viridis, and HSV as defined in the Python package matplotlib, as well as the isoluminant periodic colormap made in CAM02-UCS, labeled as UCS. (a) Colormaps shown as a set of 256 sequential colors. (b) Color difference between steps in the colormap, measured using the CAM02-UCS color difference metric. This gives one measure of the rate of perceptual change across the colormap. In this plot, the colormaps jet and HSV are very nonuniform having an irregular color difference, while viridis and the UCS colormaps are uniform with a constant color difference. (c) Lightness as a function of the colormap step. Both jet and viridis are meant for monotonic data, however, only viridis exhibits monotonically increasing lightness. The lightness for the isoluminant colormap UCS has a constant lightness. While isoluminance is not required for a periodic colormap, the lightness profile of HSV does not have a consistent rate of change that would enable shape discrimination evenly.

A solution that does not rely on the DoLP threshold is using the local spatial statistics of the AoP. One option is the variability metric Δ introduced by Tyo et al.⁴¹

$$\Delta = \sqrt{1 - [E_N(\cos \psi)^2 + E_N(\sin \psi)^2]}, \quad (1)$$

where $E_N(x)$ is the expected value of argument x in the pixel neighborhood N and ψ is the AoP. Although originally used for the HSV color fusion method, it is immediately evident that this metric would be useful when looking just at AoP. One way to implement this is to use a threshold in the variability to determine whether pixels should be invalidated before colormapping. Most visualization packages with a colormapping function have an option to set the color for invalid pixels. Most likely, this color would be black in order to differentiate more easily from the other pixels. Another option that would produce a smoother transition between valid and invalid pixels would be to map the variability Δ to transparency, as shown in Fig. 3(d). In this figure, all of the areas with ambiguous AoP are dark, allowing the user to draw all their attention to just the areas with a consistent AoP.

2.5 Degree of Linear Polarization Colormap

DoLP, as a monotonic value ranging between 0 and 1, is most suited to be represented by a monotonic, continuous

colormap. Generally, the tasks performed fall into the comparison and look-up, that is “what is the value of the DoLP here, and how does it compare to the DoLP over there?” Most commonly, the DoLP image has high spatial frequency content, which requires the colormap to include achromatic variation. Given these tasks and the data structure, an appropriate choice would be what is called the spiral colormap.³⁶ These include matplotlib’s “viridis,” as shown in Fig. 5(a)³⁹ and MATLAB’s “parula,”²⁰ in which the color array is ordered with a monotonically increasing achromatic channel and spans roughly a half cycle in hue, so that each end of the array is on either end of one of the color opponent channels, usually the blue-yellow axis. Additionally, these colormaps avoid the simultaneous presence of red and green to be more colorblind friendly.³⁹ Compare the spiral colormap in Fig. 5(a) with the rainbow colormap in Fig. 5(b), particularly the appearance of the gradient in DoLP across the sky. With the spiral colormap, the achromatic channel increasing across this range indicates a clear, intuitive progression in value from left to right. This is not true for the rainbow colormap, where the gradient is primarily indicated by a change in hue, which does not have the same intuitive directional component of the achromatic channel. There is also a significant difference in how the two colormaps steer user attention, which can be judged on where in the colormap the brightest, most vivid colors appear. For the spiral colormap, these colors appear at the maximum, indicating

Fig. 5 DoLP of a beach scene visualized with (a) spiral colormap “viridis” (matplotlib), (b) colormap “jet” (matplotlib), and (c) asymmetric lightness-increasing diverging colormap, d) symmetric diverging colormap. Data are provided by Samuel Powell and Justin Marshall, University of Queensland.

that observers will be visually steered to locations with the highest value. In the rainbow colormap, the brightest color is yellow, which appears in the upper middle range. This indicates that observer attention will be steered not toward the highest values, but to some value that likely does not carry any special significance.

As previously mentioned, monotonic data structures can utilize a diverging colormap in order to impose a simple high-low dichotomy. This is extremely effective for resolving a search query, such as “is this value low or high?” and “where are the areas of high/low DoLP?” The typical diverging colormap would involve a mirrored two-hue strategy, where the color channels are symmetric around the critical point, differing only in hue. Note that it is good practice to avoid colors that may be indistinguishable for colorblind observers (e.g., red-green) when possible. This symmetry is useful for first identifying the range (low or high) and then performing comparisons (higher, lower) within the identified range. Figure 5(d) uses a symmetric diverging colormap on DoLP data to clearly identify whether values should be considered low or high due to the fact that observers only need to conduct a visual search for a single hue corresponding to either extreme. The observer will automatically see the image with a dichotomic paradigm (e.g., red versus blue), which can be used as a tool for the visualization creator to establish a low vs high categorization.

With a symmetric colormap, it is not easy to do comparisons across identifiers, i.e., “what is the relative difference between this high value and this low value,” since the path between the values is not represented by one of the ordered color channels. Instead, the magnitudes from either extreme or from the critical point can be compared across identifiers, such as “is this high value closer to the extreme than this low value?” Such tasks are most useful for symmetric variables such as the polarized Stokes parameters S_1 , S_2 , S_3 , or ellipticity, since relative differences spanning the zero point are less important than the magnitude difference. For monotonic data, however, if the colormap is not symmetric in the

achromatic channel, but instead increases monotonically, the task of comparing values across identifiers becomes easier. This is because, like in the case of spiral colormaps, the ordering is accomplished through the achromatic channel. An example of this type of colormap is given in 5C, where both the low/high identification is present while allowing for comparison through the lightness channel. For a localization task where dealing with searching only for areas of high polarization, the diverging map can be constructed in a way that has the low side be shades of gray while the high side is colorful. With this method, the visual attention is steered only toward the areas of high polarization. This can be valuable if the background has little polarization yet retains some shape information due to the nature of edge-effects producing small fluctuations in the DoLP.

2.6 Diverging Variables

Unlike DoLP and intensity, the polarimetric Stokes parameters S_1 , S_2 , S_3 , and ellipticity, have a diverging data structure centered around a critical value of zero. For the linear Stokes parameters, the zero separates orthogonal polarization states. For the circular variables, the zero separates left-handedness from right. As such, symmetric diverging colormaps are the most appropriate for visualizing these variables. If the data contain high spatial frequency information, then the colormap additionally must implement some achromatic variation. All other types of colormaps will not support the essential task of identifying the sign of the values, particularly those values closer to the critical value. Figure 6 depicts how the choice of colormap for the S_1 parameter makes a large difference in being able to differentiate between no signal and a small, consistent amount of signal. The top row depicts S_1 data of a mantis shrimp as it was originally taken by Sam Powell and Justin Marshall, using both a diverging map and a grayscale map. The second row shows this data with the background normalized S_1 value

Fig. 6 Normalized Stokes parameter S_1 of a mantis shrimp visualized using (a and c) diverging colormap, (b and d) grayscale. (a) and (b) Images are created from the data, and (c) and (d) data for creating images have been altered by adding a value of 0.1 to areas with low S_1 signal. This is a demonstration of how the diverging colormap clearly distinguishes between no signal (white) from little signal (light red), whereas the grayscale colormap does not. Data are provided by Samuel Powell and Justin Marshall, University of Queensland.

artificially increased by a constant value of 0.1. Using the diverging colormap, the original data clearly have a white background, indicating no significant value for S_1 . With the artificially increased data, the diverging colormap has a distinct light red background, which is consistent with the change in the data. In contrast, neither grayscale images have an indication whether or not the background images have a significant S_1 component, and the difference between them is hardly noticeable.

2.7 Discrete and Contour

For any of continuous colormaps used for the polarimetric variables, there is a analogous discrete colormap that can be obtained by sampling the colormap at regular intervals. Generally, discrete colormaps are used to simplify the visualization in order to efficiently show trends, patterns, and symmetries. In polarization literature, discrete colormaps have been used to display angular symmetries in the sky.^{42,43} In these cases, discrete colormaps act effectively as a more visually stimulating alternative to contour lines, as evident in Fig. 7. In this figure, the outlines of the discrete boundaries serve to give a more easily identifiable form, which allow the patterns of DoLP around the singularities to be more apparent. In fact, contour lines have been used alongside discrete maps in order to better outline the boundaries.⁴² As contour lines become more complicated, using elements such as color or texture can significantly improve the differentiation of boundaries.⁴⁴ Some examples of the use of texture for defining boundaries in polarization literature exist, such as using texture to form boundaries between polarized and unpolarized regions,⁵ as well as a boundary between measured and unmeasured areas.⁴⁵

Fig. 7 Polarization of the sky simulated according to equations defined by Berry et al.⁴² with Arago and Babinet singularities. Discrete colormap for expressing contours.

3 RGB

While univariate colormaps are effective for visualizing a single polarimetric variable, it is often the case that polarization attributes that are spatially correlated be represented by a single multivariate visualization. Given that colors are most often encoded using the red, green, and blue primaries of display monitors, it is reasonable to see why the RGB color cube is often intuitively used as the space to map the multidimensionality of polarimetric variables. These primaries have been used to represent the Stokes parameters² or multispectral DoLP.³⁸ The main issue with this technique

is that the RGB colorspace is not a color appearance model; it does not represent how color is perceived.³⁴ Rather, the RGB primaries are more representative of the voltages the monitor needs to supply to the individual subpixels. Thus, the way a computer represents a color is far different than the perceptual color attributes we would ascribe to them. For example, a color encoded as $[1, 0.8, 0.5]$, when displayed on a monitor would be perceived as a pale orange. The descriptive words attributed to the color: “pale” for its saturation and “orange” for its hue are used rather than labeling the color as “fully red, with a good amount of green, and some blue” as we would if we interpreted it similarly to the monitor. In fact, the monitor’s ability to additively mix a few colors to trick the human visual system into perceiving a broad range of colors is precisely why monitors are able to function. Therefore, representing variables with the RGB primaries are likely to produce visualizations in which the values are difficult to interpret. It is not enough simply to match the dimensionality of RGB color cube to the number of channels; the values must be encoded according to the way color is interpreted in order to be effective at communicating information.

Additionally, using the RGB channels as if they were independently accessible and linear can create ineffectiveness in the final result. Each primary is significantly different in luminance, resulting in significantly different visibility and contrast profiles for each channel. Blue, as the darkest primary, will be much harder to distinguish from areas with low signal, even if the signal representing blue is very strong. Additionally, the choice of chromaticity of the RGB channels is meant to maximize the amount of available colors instead of representing equally spaced hues. Thus, the transition from one primary to another is unsurprisingly not a uniform transition in hue. Since the hue of any RGB-encoded color is determined by the two strongest primaries, having a nonuniform hue presents issues. As shown by Kruse et al.,⁴⁰ this nonuniformity results in higher density of green hues compared to any other nominal color. Thus, the viewer of a three-channel RGB-encoded visualization may be more likely to overestimate the magnitude of the channel represented by green.

Since hue is most suited for producing categorical distinctions and saturation for comparing magnitudes, restructuring the visualization to exploit this can result in more effective visualizations. Consider the method of visualizing DoLP taken at three wavelengths via RGB encoding. The presence of color indicates that the polarization is wavelength-dependent. The resulting hue is an indication of some categorical information on the ratio of the stronger two signals. Note that these categories are not evenly distributed: there are more distinct hues between red and green than there are between green and blue. Berlin⁴⁶ posited that the categorical distinctions for different hues are not arbitrary and are consistent across cultures. This hypothesis has since been challenged,⁴⁷ yet if one accepts the set of red, green, blue, and yellow to be sufficiently accepted as primary colors, there is still at least one more categorical color between red and green than between green and blue. This additionally would mean that the wavelengths encoded in another permutation (red, green, and blue assigned to ascending wavelengths) would result in the areas being categorized completely different, violating the visualization design principle of “representation invariance.”⁴⁸

Instead of having the categorizations be implemented in such an unintentional way, the visualization can be constructed such that the categorizations are made more explicit. Using the convention of red, green, blue, yellow, orange, and purple as categorical hues, the different possible states of wavelength ratios can be differentiated by hue. Red, blue, and yellow can represent states where one signal (DoLP₁, DoLP₂, DoLP₃) is stronger than the median. Green, orange, and purple can indicate states where one signal is weaker than the median. The second set of colors represent combinations of the first (DoLP₂ and DoLP₃, DoLP₁ and DoLP₃, DoLP₁ and DoLP₂). By restricting the possible values the hue can take, not only is the number of categories for each state explicitly and evenly set, but also the ability to extract the meaning from the hue is made dramatically easier. If the deviations of one channel are meant to be highlighted, the saturation can be encoded by the difference from the median. Here are three likely visualizations for this scenario: one to indicate where a channel is significantly greater than the others, one to indicate where a channel is significantly lower than the others, and one where both are indicated (with the preference of overlapping categories going to the more extreme case). Here, we will label these “strong,” “weak,” and “combined,” respectively. The luminance channel can represent the average DoLP, since this channel carries much of the shape information. Due to the fact that not all three dimensions can be fully utilized independently to the extent of any display gamut, the mapping methods formulated within the uniform color fusion method [Eqs. (10a)–(10d)] introduced by Kruse et al.⁴⁹ can be used to map this method into the uniform color space CAM02-UCS to produce a categorical color fusion:

$$\begin{aligned}
 J' &= f_{J'}(\text{DoLP}_{\text{mean}}^x), \\
 M'_{\text{strong}} &= f_{M'}[s(\text{DoLP}_{\text{max}} - \text{DoLP}_{\text{med}})], \\
 M'_{\text{weak}} &= f_{M'}[s(\text{DoLP}_{\text{med}} - \text{DoLP}_{\text{min}})], \\
 M'_{\text{comb}} &= \max(M'_{\text{strong}}, M'_{\text{weak}}), \\
 h_{\text{strong}} &= \begin{cases} 0.185(\text{red}) & \text{DoLP}_3 = \text{DoLP}_{\text{max}} \\ 1.723(\text{yellow}) & \text{DoLP}_2 = \text{DoLP}_{\text{max}} \\ -1.775(\text{blue}) & \text{DoLP}_1 = \text{DoLP}_{\text{max}} \end{cases}, \\
 h_{\text{weak}} &= \begin{cases} 2.385(\text{green}) & \text{DoLP}_3 = \text{DoLP}_{\text{min}} \\ 0.976(\text{orange}) & \text{DoLP}_2 = \text{DoLP}_{\text{min}} \\ -1.079(\text{purple}) & \text{DoLP}_1 = \text{DoLP}_{\text{min}} \end{cases}, \\
 h_{\text{comb}} &= \begin{cases} h_{\text{strong}} & M'_{\text{strong}} > M'_{\text{weak}} \\ h_{\text{weak}} & M'_{\text{strong}} \end{cases}, \quad (2)
 \end{aligned}$$

where subscripts max, min, med, and mean are the maximum, minimum, median, and mean DoLP signal for that pixel, with numerical subscripts referring to spectral channel. Function f_x is the mapping function for a color channel x given in Eqs. (10a)–(10d). J' , M' , and h are the CAM02-UCS color channels of lightness, colorfulness, and hue, respectively. The exponent in the lightness channel is there for contrast enhancement. The values for the hues given here are not of particular significance; they are just here to help with converting color spaces. While this differentiation is effective for three-channel scenarios, the “strong”

Fig. 8 Three-channel spectral DoLP visualizations taken on ground MSPI at University of Arizona, courtesy of Christine Bradley and Russell Chipman,³⁸ $\lambda_1, \lambda_2, \lambda_3 = 470, 660, 865$ nm. (a)–(c) Categorical color fusion method based on Eq. (2). Center colorbar indicates difference from the median for the (a) highest DoLP, (b) lowest DoLP, and (c) combined. Hue indicates which wavelength is highest or lowest. (d) DoLP₁, DoLP₂, DoLP₃ as RGB primaries. (e) Diverging color fusion with colorbar for the standard deviation of each pixel between the three wavelengths.

and “weak” cases could apply to more wavelengths, however, the number of distinct hues will max out at around eight (the six already named plus cyan and pink). Thus, the number of combined cases, where both strong and weak signals are indicated, would be restricted to four wavelengths (four hues for strong, four for weak). For the individual cases of indicating strong or weak signal, this visualization method may support up to eight wavelengths. This number is determined by the list of 12 distinct colors for labeling detailed by Ware⁴⁴ excluding black, white, gray, and brown.

In Figs. 8(a)–8(c), the “strong,” “weak,” and “combined” categorical color fusions are demonstrated. Here, the hue unambiguously indicates that one of the wavelengths has a significant difference to the others. The RGB-encoded image in Fig. 8(d) displays many colors, however, it is difficult to determine exactly how the color attributes relate to the DoLP for each wavelength. Compare the color of the sky in this figure. In the RGB image, the blue color would indicate that the DoLP₃ is the strongest. This agrees with the “strong” image [Fig. 8(a)], which shows that this wavelength has a moderately higher DoLP than the median. In the “weak” image [Fig. 8(b)], the sky is indicated strongly with green, indicating that there is a significant lack of DoLP₁. In the combined image, the sky is green because the weakness of DoLP₁ is more extreme than the strength of DoLP₃. In the RGB image, the only way to assess the strength of DoLP₁ is by determining the saturation of the sky, which does not produce a categorical shift like a change in hue. Therefore, the color of the sky is shown as categorically strong in DoLP₃, even though the more defining attribute is its lack of DoLP₁. While the definition of “defining attribute” need not align with how it is used here in all situations, it is used here to provide an example for the process of aligning user-chosen categorizations to visual categorizations.

While those methods are quite effective in highlighting areas in which the polarization is decorrelated between the channels, it can also be useful to highlight areas in which the polarization channels are correlated. If the saturation is instead negatively proportional to the pixel standard deviation, areas in which the channels are similar would be colorful. This can be formulated as a single-hue color fusion:

$$M' = f_{M'}[(1 - \sigma_\lambda)], \quad (3)$$

so that as the standard deviation approaches zero, the colorfulness approaches maximum. In this situation, it would be useful to apply the mapping only to pixels meeting a minimum DoLP threshold in order to desaturate unpolarized areas. Additionally, both correlation and decorrelation can be visualized using a diverging color fusion:

$$M' = f_{M'}(|s - \sigma_\lambda|/s), \quad (4a)$$

$$h = \begin{cases} 0.185(\text{red}) & \sigma_\lambda < s \\ -1.775(\text{blue}) & \sigma_\lambda > s \end{cases}, \quad (4b)$$

where s is some chosen value corresponding that distinguishes correlated from uncorrelated standard deviation values. Figure 8(e) uses the diverging color fusion method to identify whether the DoLP of regions are wavelength-dependent or not. The median standard deviation was around 0.05, which was used as the critical point.

4 Color Fusion

While mapping data into the RGB cube may initially be the most intuitive way to map polarization into color, for a long time researchers have understood that representing the polarization parameters using a more perceptually based color model leads to a more natural way to represent polarization with color. Because the dimensions of linear

polarization are remarkably similar to that of color vision, the technique of directly relating the polarization channels to their colorimetric counterpart is as old as polarization imaging itself. While the technique is often attributed to the work of Bernard and Wehner⁹ on physiological similarities of systems for color vision and polarimetric vision, it was actually Walraven⁵⁰ who introduced mapping the polarimetric imaging channels into color channels for producing images. This was introduced in a single paragraph at the end of his seminal paper, formulated here as follows:

$$\begin{aligned} \text{POL} &\rightarrow \text{saturation} \\ \psi &\rightarrow \text{hue}, \\ I &\rightarrow \text{grayscale}, \\ \text{POL} &= I(1 - \text{DoLP}), \end{aligned} \quad (5)$$

with arrows (instead of equals) indicating the function mapping to color was not well-defined. This method did not precisely define saturation or hue in terms of a color space. Due to the ambiguity of the terms “saturation” and “hue,” it would be difficult to reproduce or analyze this method. Following Walraven’s initial solution, the work of Solomon⁶ gave a more complete description of the mapping of polarization to color. His method was based off of mapping DoLP and AoP into the hue and saturation channels defined by a color space that was introduced by Faugeras.⁵¹ Even though the color channels were now well-defined, Solomon did not account for the irregular shape of perceptual color spaces, which requires some clipping of the gamut in order to evenly map polarimetric channels into color coordinates.⁴⁹ The calculations necessary to convert those color channel into display coordinates were likely too computationally expensive given the technology at that time. Coupled with the relative obscurity of Faugeras’ color space, it is not surprising that this method was never replicated in any subsequent publications. The first iteration of the well-known method of mapping into the HSV color space was introduced by Wolff and Mancini,⁵² mapping the parameters as follows:

$$I = V, \quad (6a)$$

$$\text{DoLP} = S, \quad (6b)$$

$$2 \cdot \psi = H, \quad (6c)$$

where the factor of 2 extends the 180 deg period of ψ into the 360 deg period of H . In a subsequent paper, Wolff et al.⁵³ state that the color is rendered in HSI space, contradicting their earlier use of HSV. While they likely were referring to the same color space, the ambiguous use of word choice can lead to drastically different implementations. Today, the majority of implementations assume the HSV space, however, some applications⁵⁴ are based on the assumption that Wolff et al.⁵³ refer to other color spaces. This illustrates how the process of mapping polarimetric channels into color channels has historically been ambiguous, and the necessity to be more precise when defining terms.

The main advantage of the methods that map to color spaces is their ability to leverage the benefits of human color vision to create an easily understandable yet spatially correlated depiction of polarimetric phenomenon. The HSV

space, unlike HSI or HSL, has peak chroma when S and V are at maximum. The bright, colorful areas that draw the most attention correspond to the strongest polarization and are likely to be very important. In contrast, both HSL and HSI spaces have the most vibrant colors toward the middle range of lightness/intensity, meaning that the areas with high intensity will not be colorful. Even though those areas are drawing the most attention, they do not lose their contextual physical attributes defined by the intensity of the area and its surrounding. Bowers et al.⁵⁵ identify the usefulness of this technique and apply it to thermal polarimetric imaging, with value channel instead representing the thermal levels.

Tyo et al.⁵⁶ introduced a simplified method using two orthogonal hues for polarimetric difference imaging. While this type of imaging does not even capture the full linear polarimetric parameters, this work is often inaccurately cited as source of the general HSV method. The motivation behind this method is to reduce the attention-steering nature of noise once it is colormapped without additional processing.

The latest amendment to the HSV method came from Tyo et al.,⁴¹ due to the inability to see polarization in dark regions. This is because the channel S has its perceptual magnitude (measured in colorfulness or chroma), proportional to V . That is, dark regions will always be not colorful, even when they are fully saturated. Since regions with low irradiance often have higher polarization,⁵⁷ this characteristic of the HSV model can make significant portions of the polarization information inaccessible to observers. The method they proposed to prevent this was to have DoLP map into V when the intensity is lower than DoLP:

$$H = 2 \cdot \psi, \quad (7a)$$

$$S = \text{DoLP}, \quad (7b)$$

$$V = \max(I, \text{DoLP}). \quad (7c)$$

While this method works well to eliminate the sources that can cause some polarization to be missed, it inadvertently increases the amount of false polarization. In order to avoid this, Tyo et al.⁴¹ employed a method to determine the appropriate mapping method at the pixel level based on the variability Δ of the AoP, given above in Eq. (1). When the variability is high, the standard mapping was used for that pixel. When the variability is low, the amplified mapping was used:

$$V = f(\Delta, t) = \begin{cases} \max(I, \text{DoLP}) & \Delta \leq t \\ I & \Delta > t \end{cases} \quad (8)$$

An optimal threshold t may be dependent on several factors that affect the statistics of the angular distribution like instrument, spectrum, amount of illumination, etc. Additionally, there must be some decision made to the level of acceptable variability in the angular distribution to define what is considered false polarization. In practice, the threshold was applied to create a binary map indicating the mapping method. This allows for morphological processes, such as opening and closing to fill in gaps before mapping into color.

While this method greatly reduces the amount of false polarization from dark regions, it does not affect the other cases of false polarization. Kruse et al.⁴⁹ argue that simply reducing the saturation in areas with high variability

Fig. 9 IHIP data, $^{58} \lambda = 4.17 \mu\text{m}$. (a) and (b) [Eqs. (9a) and (9)] Tyo et al.⁴¹ HSV color fusion method, max degree of polarization (DoP) set to 0.4. (c) and (d) [Eqs. (10a)–(10d)] Kruse et al.⁴⁹ uniform color fusion method, max DoP set to 0.4. (b) and (d) represent the data in (a) and (c) rotated by 20 deg. Note how, when using the HSV method, the AoP gradient over the hood of the truck in (a) is not maintained after rotation in (b), where the sharp color difference between cyan and blue (evident in Fig. 4) perceptually split the hood into two sections. This is not the case in (c) and (d), which use CAM02-UCS due to its perceptual uniformity.

accounts for eliminating much more of the sources of false polarization. This instead uses an assumption that areas with high angular variability should not be displayed as real sources of polarization.

To implement this, the DoLP can be multiplied by the previously defined binary map so that all areas above the variability threshold are reduced to zero. The mapping functions would then be as follows:

$$S = D \cdot \text{DoLP}, \quad (9a)$$

$$V = \max(I, S), \quad (9b)$$

where D is the binary threshold map. Due to the amplification of dark regions on a single pixel basis, these techniques can create excessive pixelated contrast. To the viewer, this can be jarring, distracting, unaesthetic, and can convey a diminished sense of quality. To reduce this, the binary map D can be smoothed prior to using Eqs. (9a) and (9b), for example, by using a Gaussian filter.

While the cylindrical coordinate system of HSV is useful as an approximation of human color vision for color-picking, there are many issues caused by using it as the model for visualization.²⁹ The main two issues of HSV are nonuniformity and channel mixing.⁴⁹ Nonuniformity describes how perceptual channels get distorted, meaning that differences in the polarization values are not proportional to the perceptual differences. Channel mixing describes how the individual channels of HSV map to multiple perceptual channels simultaneously, meaning that extracted polarization information from a perceptual channel is dependent on the other polarimetric variables. Both of these aspects are apparent in Fig. 4, with the nonuniformity evident from the highly varying color difference, and the channel-mixing with the varying lightness since value V is held constant. In the color fusion

technique, these aspects are made visible comparing identical data where in one set, the AoP is rotated by some angle. In Fig. 9, the rotation between left and right columns creates a significantly different appearance on the hood of the truck. In Fig. 9(a), the hue indicates a smooth gradient across the hood, whereas in Fig. 9(b), the hue and lightness change so rapidly that the hood appears to be two distinct sections. In comparison, the color fusion designed to be uniform and without channel-mixing shows no such phenomenon in Figs. 9(c) and 9(d).

The method Solomon introduced would not have had these issues because it was based on a uniform color space. Combining ideas from Solomon and Tyo et al., and with a more contemporary color space, Kruse et al.⁴⁹ introduced mapping methods to map the channels of AoP, DoLP, and intensity into the color channels of hue h , colorfulness M' , and lightness J' , defined in the uniform color space CAM02-UCS. The move from using saturation to colorfulness is supported by two main arguments. First, saturation does not have a uniform correlation in CAM02-UCS. There is a nonuniform correlation in the nonuniform color model CAM02, however, the uniform transformation does not have a defined term for saturation. More importantly, the ability to perform comparison tasks using a saturation channel is highly dependent on lightness. In contrast, the comparison task would likely not be affected by changes in lightness, due to the fact that the magnitude of color difference along the colorfulness axis would not change at different lightness levels. The mapping equations are as follows:

$$h = f_h(\psi) = 2 \cdot \psi, \quad (10a)$$

$$M' = f_{M'}(\text{DoLP}) = \max(c) \cdot \text{DoLP}, \quad (10b)$$

$$J' = f_{J'}(I) = I \cdot (J'_1 - J'_0) + J'_0, \quad (10c)$$

Fig. 10 IHIP data,⁵⁸ $\lambda = 4.17 \mu\text{m}$ [Eq. (11)]. Diverging color fusion method, max DoP set to 0.3.

$$J'_0, J'_1 = \{J' : c(J') = M'\}, \quad (10d)$$

which describe the technique of placing polarization channels into more accurate color channels. The curve c represents the maximum M' for any J' that allows any hue with that colorfulness and lightness to map into a given RGB color space. Therefore, c is the boundary curve that when rotated along the J' axis defines the subset of the display gamut which is rotationally symmetric. This rotational symmetry allows for two of the parameters to be independent or unmixed. Due to the limitations of human color vision, at least one of the channels J' or M' must be dependent on the other in order to have a reasonably large range in both. For this reason, I is linearly mapped to the J' range that can be achieved with the current M' . An analogous visualization for circular polarization can be constructed using diverging color fusion⁵⁸ by having the hue indicate the handedness:

$$h = \begin{cases} 0.185(\text{red}) & S_3 > 0 \\ -1.775(\text{blue}) & S_3 < 0 \end{cases} \quad (11)$$

An example of this is shown in Fig. 10.

Because AoP often describes categorical information about an object, e.g., its physical orientation or whether it is reflecting a polarized light source, it is very convenient that hue is both categorical and periodic. DoLP, on the other hand, is generally not as definitively categorical on the level of small deviations. That is, it is often the case where the particular value of an object's AoP is more descriptive of the source of the polarizing effect than how much it is polarized. It would be more likely that two areas with the same AoP have a shared source even if they differ in DoLP. Thus, having the visualization primed for assisting in identification tasks for AoP and comparison tasks for DoLP is often the most effective combination. Still, there are an infinite number of tasks and task combinations using this framework that could be implemented to construct different visualizations.⁵⁹ In addition to the traditional color fusion, the other types discussed here include diverging, categorical, and single-hue. It is important to keep in mind the importance of supplying some sort of key when using these types of color fusion techniques. Without some sort of indication of how the viewer is meant to interpret the color values, the visualization mainly functions as a source of aesthetics.

5 Markers

Often, the color and intensity of the image are desired to be displayed as they are, rendering colormapping strategies useless. As such, since the original image is meant to be unaltered, there is a popular collection of methods for superimposed markers that depict polarization states at their spatial positions. The markers can be lines,^{43,60} ellipses,^{4,61} moving dots,⁵ or arbitrary symbols.⁵ The markers can represent polarization states at specific spatial intervals,⁴³ or placed in a randomized fashion,⁶⁰ or depict the average states of local pixel neighborhoods.⁴

For linear markers, encoding the DoLP into the length of the line or arrow may result in biases in the tasks. A psychophysical study by Craven⁶² exhibited consistently that lines that are oriented more vertically are estimated to appear longer than lines oriented horizontally. Additional evidence for this phenomena is demonstrated more recently by Zhu and Ma.⁶³ In order to reduce the impact of this bias, it would be beneficial to also encode DoLP into additional channels, such as color, using the colormapping methods discussed above. Using multiple channels to encode a single variable may seem excessive, but the use of redundant visual cues is well established in the visualization literature.⁶⁴

The debate over the optimal placing of indicator lines for vector field visualizations is still ongoing. Ware,⁴⁴ using the Field et al.⁶⁵ theory of contour perception, argues that placing vector lines in a grid-like fashion creates false contours and does not strongly stimulate the neurons responsible for detecting the actual contours. Instead, he argues that end-to-end lines, where the end of each line points toward the end of another line, are the most effective for establishing a perceived contour. However, this method inherently adds complexity to the visualization algorithm by requiring a method to position them in such an alignment. A middle-path approach is the "jittered" or randomized line segment location. Without the grid pattern, the false contours are less likely to be perceived, and end-to-end positioning is more likely to occur. Given that measurements are virtually always taken on a grid, randomization can be implemented by allowing the position of each line segment to vary within the local area of the grid. Such a technique is exemplified perfectly by Berry⁶⁰ in a visual comparison of revisualizing the measured AoP of the daylight sky with randomly shifted line segments that was originally published using a strict grid. The contours for the randomized lines are much more apparent than for the lines on a grid, as apparent in Fig. 11. Note that the decision between grid and jittered is not important if contour or trend detection are not important tasks.

On the other hand, one user study for vector flow visualizations carried out by Laidlaw et al.⁶⁶ did not find much statistically significant difference between the flow visualizations using a grid versus randomized. However, the tasks involved were not representative of polarization vision applications. Additionally, the types of flow visualizations were not similar to the vector visualizations used in polarization imaging. Given that this is the only user study that is somewhat related, it should not be dismissed entirely.

The method for determining the type and orientation of the marker based on a local pixel average is most evidently useful in situations where spatial features are significantly larger than the individual data points and those individual

Fig. 11 Polarization of the sky simulated according to equations defined by Berry et al.,⁴² with Arago and Babinet singularities. Vector maps for AoP using (a) a grid, (b) a randomized pattern, (c) with DoLP encoded into length, and (d) with DoLP encoded into length and color.

data points have significant variance, such that a randomly selected data point is not adequately representative of its neighbors. This rules out simulated data⁶¹ as well as sparse data.^{60,67} The averaging method has been used effectively for biological polarization imaging due to high amount of noise.⁴

6 Miscellaneous Methods

While the most commonly used visualizations have been covered, there are a handful of visualizations introduced by Yemelyanov et al.⁵ that have not been practically implemented outside of the initial publication. The first method involves the visual system's ability to cluster objects based on coherent motion.⁶⁸ While this is not a widely used visualization technique, there is considerable psychophysical evidence⁶⁹ that this is a underused channel for data visualization.⁴⁴ The method overlays randomly seeded dots, where the motion of each dot is either set to random or coherent based on a threshold in the polarization difference image. The orientation is set to either horizontal or vertical based on whether the polarization difference is positive or negative. While the horizontal/vertical alignments are artifacts of polarimetric difference imaging, it is easy to see how the movement can be altered to be extendable to linear, circular, and elliptical polarization, using the polarization ellipse as the movement pattern. The ability to separate random movement from coherent movement may be underutilized in imposing a binary movement pattern of either fully random or fully coherent. It may be underestimating the visual system's ability to detect smaller differences in motion, but the extent to which it is underestimated may only be determined by conducting user studies. Instead of the binary movement patterns, the coherence of the movement pattern can be proportional to the degree of polarization. Perhaps the greatest

benefit that is unique to this method is the ability to preserve the unpolarized image, like other types of markers, while depicting a spatially resolved polarization information, like methods using color. Size, frequency, transparency, and amplitude, among others, are some parameters of the moving dots that have not yet been explored.

Yemelyanov et al.⁵ introduced two methods for temporally modulating the unpolarized and polarized images. In one method, the visualization transforms from the unpolarized image to the color fusion image and back over a set period of frames. In the second, the intensity of polarized regions modulated, where positive and negative polarization differences are set a half cycle apart. Like the moving dots, this type of method has not undergone user studies that might indicate its effectiveness. These methods, while largely unknown, have significant potential and warrant further investigation.

7 Conclusion

Figure 12 illustrates some of the decisions that must be made in order to choose an appropriate visualization method. This includes whether or not the data is multivariate or univariate, what the data type is, and what tasks are being chosen. While this diagram is not exhaustive, performing the exercise of going through each decision is a good way to get in the mindset of visualization design. The diagram is meant to be only a simple approximation to the questions one should ask when deciding on a particular visualization. This is not meant to be a set of rules to follow but rather a method for producing an effective visualization that suits the needs of the person using it. In fact, the future for polarization visualization is much more open-ended, with a seemingly limitless number of avenues in which to design.

Acknowledgments

1. G. Garlick, G. Steigmann, and W. Lamb, "Differential optical polarization detectors," US Patent 3,992,571 (1976).
2. J. R. Schott, "Fundamentals of polarimetric remote sensing," in *Tutorial Texts in Optical Engineering*, Vol. **TT81**, SPIE, Bellingham, Washington (2009).
3. F. Snik et al., "An overview of polarimetric sensing techniques and technology with applications to different research fields," *Proc. SPIE* **9099**, 90990B (2014).

4. Y. L. Gagnon and N. J. Marshall, "Intuitive representation of photopolarimetric data using the polarization ellipse," *J. Exp. Biol.* **219**(16), 2430–2434 (2016).
5. K. M. Yemelyanov et al., "Bio-inspired display of polarization information using selected visual cues," *Proc. SPIE* **5158**, 71–84 (2003).
6. J. E. Solomon, "Polarization imaging," *Appl. Opt.* **20**(9), 1537–1544 (1981).
7. N. Andrienko and G. Andrienko, *Exploratory Analysis of Spatial and Temporal Data: A Systematic Approach*, Springer Science & Business Media, New York (2006).
8. B. M. Ratliff and J. S. Tyo, "Moving towards more intuitive display strategies for polarimetric image data," *Proc. SPIE* **10407**, 104070E (2017).
9. G. D. Bernard and R. Wehner, "Functional similarities between polarization vision and color vision," *Vision Res.* **17**(9), 1019–1028 (1977).
10. M. J. How and N. J. Marshall, "Polarization distance: a framework for modelling object detection by polarization vision systems," *Proc. R. Soc. B: Biol. Sci.* **281**(1776), 20131632–20131632 (2014).
11. M. R. Luo, G. Cui, and C. Li, "Uniform colour spaces based on CIECAM02 colour appearance model," *Color Res. Appl.* **31**(4), 320–330 (2006).
12. L. M. Eshelman et al., "All-sky polarization measurements of the total solar eclipse on 21 August 2017," *Proc. SPIE* **10655**, 106550L (2018).
13. K. P. Gorton, A. J. Yuffa, and G. W. Videen, "Enhanced facial recognition for thermal imagery using polarimetric imaging," *Opt. Lett.* **39**(13), 3857–3859 (2014).
14. W. Zhang et al., "Angle of sky light polarization derived from digital images of the sky under various conditions," *Appl. Opt.* **56**(3), 587–595 (2017).
15. J. Craven-Jones et al., "A polarization system for persistent chemical detection," *Proc. SPIE* **9613**, 96130B (2015).
16. G. M. Calabrese et al., "Polarization signaling in swordtails alters female mate preference," *Proc. Natl. Acad. Sci. U. S. A.* **111**(37), 13397–13402 (2014).
17. D. A. Lavigne et al., "Evaluation of active and passive polarimetric electro-optic imagery for civilian and military targets discrimination," *Proc. SPIE* **6972**, 69720X (2008).
18. A. Light and P. J. Bartlein, "The end of the rainbow? Color schemes for improved data graphics," *Eos, Trans. Am. Geophys. Union* **85**(40), 385–391 (2004).
19. D. Borland and R. M. T. Li, "Rainbow color map (still) considered harmful," *IEEE Comput. Graphics Appl.* **27**(2), 14–17 (2007).
20. S. Eddins, "Rainbow color map critiques: an overview and annotated bibliography," *MathWorks Technical Articles and Newsletters*, Vol. **25**, 92238v00 (2014).
21. M. Borkin et al., "Evaluation of artery visualizations for heart disease diagnosis," *IEEE Trans. Visual. Comput. Graphics* **17**(12), 2479–2488 (2011).
22. L. Patrick, M. Graham, and P. Stanley, "Lens array stokes imaging polarimeter," *Meas. Sci. Technol.* **22**(6), 065901 (2011).
23. W. Zhang et al., "Robust sky light polarization detection with an S-wave plate in a light field camera," *Appl. Opt.* **55**(13), 3518–3525 (2016).
24. E. Vexberg and M. A. Golub, "Optical imaging system which is sensitive to a degree of light polarization and coherence," *Proc. SPIE* **8165**, 816515 (2011).
25. L. Zhou and C. D. Hansen, "A survey of colormaps in visualization," *IEEE Trans. Visual. Comput. Graphics* **22**(8), 2051–2069 (2016).
26. J. Bernard et al., "A survey and task-based quality assessment of static 2D colormaps," *Proc. SPIE* **9397**, 93970M (2015).
27. S. Silva, J. Madeira, and B. S. Santos, "There is more to color scales than meets the eye: a review on the use of color in visualization," in *Proc. 11th Int. Conf. Inf. Visual.*, pp. 943–950 (2007).
28. L. D. Bergman, B. E. Rogowitz, and L. A. Treinish, "A rule-based tool for assisting colormap selection," in *Proc. Visual.*, p. 118 (1995).
29. C. A. Brewer, "Color use guidelines for data representation," in *Conf. Sel. Stat. Graphics, Am. Stat. Assoc.*, pp. 55–60 (1999).
30. S. Mittelstädt et al., "Colorcat: guided design of colormaps for combined analysis tasks," in *Proc. Eurographics Conf. Visual.*, Vol. **2**, p. 7 (2015).
31. K. M. Thyng et al., "True colors of oceanography: guidelines for effective and accurate colormap selection," *Oceanography* **29**(3), 9–13 (2016).
32. C. Tominski, G. Fuchs, and H. Schumann, "Task-driven color coding," in *Proc. IEEE's 12th Int. Conf. Inf. Visual.*, pp. 373–380 (2008).
33. T. Munzner, *Map Color and Other Channels*, Book Section 10, AK Peters Visualization Series, A K Peters/CRC Press, New York, pp. 218–241 (2014).
34. K. Moreland, *Diverging Color Maps for Scientific Visualization*, Springer, Berlin Heidelberg, pp. 92–103 (2009).
35. K. T. Mullen and W. H. A. Beaudot, "Comparison of color and luminance vision on a global shape discrimination task," *Vision Res.* **42**(5), 565–575 (2002).
36. C. Ware, "Color sequences for univariate maps: theory, experiments and principles," *IEEE Comput. Graphics Appl.* **8**(5), 41–49 (1988).
37. N. Camgöz, C. Yener, and D. Güvenç, "Effects of hue, saturation, and brightness: part 2: attention," *Color Res. Appl.* **29**(1), 20–28 (2004).
38. C. L. Bradley et al., "Spectral invariance hypothesis study of polarized reflectance with ground-based multiangle spectropolarimetric imager (ground MSPi)," *Proc. SPIE* **9613**, 96130U (2015).
39. N. Smith and S. V. D. Walt, "A better default colormap for matplotlib," *SciPy2015* (2015).
40. A. W. Kruse et al., "Polarization-color mapping strategies: catching up with color theory," *Proc. SPIE* **10407**, 104070G (2017).
41. J. S. Tyo, B. M. Ratliff, and A. S. Alenin, "Adapting the HSV polarization-color mapping for regions with low irradiance and high polarization," *Opt. Lett.* **41**(20), 4759–4762 (2016).
42. M. V. Berry, M. R. Dennis, and J. R. L. Lee, "Polarization singularities in the clear sky," *New J. Phys.* **6**(1), 162–162 (2004).
43. G. Horvath et al., "Polarization portrait of the arago point: video-polarimetric imaging of the neutral points of skylight polarization," *Naturwissenschaften* **85**(7), 333–339 (1998).
44. C. Ware, *Information Visualization: Perception for Design*, Elsevier, Amsterdam (2012).
45. G. Horvath et al., "Ground-based full-sky imaging polarimetry of rapidly changing skies and its use for polarimetric cloud detection," *Appl. Opt.* **41**(3), 543–559 (2002).
46. B. Berlin, *Basic Color Terms: Their Universality and Evolution*, University of California Press, Berkeley, pp. 164–171 (1969).
47. B. A. C. Saunders and J. van Brakel, "Colour: an exosomatic organ?" *Behav. Brain Sci.* **20**(2), 212–220 (1997).
48. G. Kindlmann and C. Scheidegger, "An algebraic process for visualization design," *IEEE Trans. Visual. Comput. Graphics* **20**(12), 2181–2190 (2014).
49. A. W. Kruse et al., "Perceptually uniform color space for visualizing trivariate linear polarization imaging data," *Opt. Lett.* **43**(11), 2426–2429 (2018).
50. R. Walraven, "Polarization imagery," *Proc. SPIE* **112**(1), 164–167 (1977).
51. O. D. Faugeras, "Digital color image processing and psychophysics within the framework of a human visual model," PhD Thesis, University of Utah (1976).
52. L. B. Wolff and T. A. Mancini, "Liquid crystal polarization camera," in *Proc. IEEE Workshop Appl. Comput. Vision*, pp. 120–127 (1992).
53. L. B. Wolff et al., "Liquid crystal polarization camera," *IEEE Trans. Rob. Autom.* **13**(2), 195–203 (1997).
54. Y. Zhao et al., "Object separation by polarimetric and spectral imagery fusion," *Comput. Vision Image Understanding* **113**(8), 855–866 (2009).
55. D. L. Bowers et al., "Evaluation and display of polarimetric image data using long-wave cooled microgrid focal plane arrays," *Proc. SPIE* **6240**, 62400F (2006).
56. J. Tyo, E. Pugh, and N. Engheta, "Colorimetric representations for use with polarization-difference imaging of objects in scattering media," *J. Opt. Soc. Am. A* **15**(2), 367–374 (1998).
57. M. J. Duggin, "Imaging polarimetry in scene element discrimination," *Proc. SPIE* **3754**, 108–117 (1999).
58. J. Craven-Jones et al., "Infrared hyperspectral imaging polarimeter using birefringent prisms," *Appl. Opt.* **50**(8), 1170–1185 (2011).
59. A. W. Kruse et al., "Overview of visualization strategies for polarimetric imaging data," *Proc. SPIE* **10655**, 106550S (2018).
60. M. V. Berry, "Nature's optics and our understanding of light," *Contemp. Phys.* **56**(1), 2–16 (2015).
61. O. Korotkova et al., "Coherence and polarization properties of far fields generated by quasi-homogeneous planar electromagnetic sources," *J. Opt. Soc. Am. A* **22**(11), 2547–2556 (2005).
62. B. J. Craven, "Orientation dependence of human line-length judgements matches statistical structure in real-world scenes," *Proc. R. Soc. London. Ser. B: Biol. Sci.* **253**(1336), 101–106 (1993).
63. J. E. Zhu and W. J. Ma, "Orientation-dependent biases in length judgments of isolated stimuli," *J. Vision* **17**(2), 20 (2017).
64. P. L. Rheingans, "Task-based color scale design," *Proc. SPIE* **3905**, 35–43 (2000).
65. D. J. Field, A. Hayes, and R. F. Hess, "Contour integration by the human visual system: Evidence for a local 'association field'," *Vision Res.* **33**(2), 173–193 (1993).
66. D. H. Laidlaw et al., "Comparing 2D vector field visualization methods: a user study," *IEEE Trans. Visual. Comput. Graphics* **11**(1), 59–70 (2005).
67. L. Woltjer, "The polarization and intensity distribution in the crab nebula derived from plates taken with the 200-inch telescope by Dr. W. Baade," *Bull. Astron. Inst. Neth.* **13**(478), 301–311 (1957).
68. K. M. Yemelyanov et al., "Display of polarization information by coherently moving dots," *Opt. Express* **11**(13), 1577–1584 (2003).
69. S. Limoges, C. Ware, and W. Knight, "Displaying correlations using position, motion, point size or point colour," in *Proc. Graphics Interface*, **770**, 262–265 (1989).

Andrew W. Kruse received his BS degree in optical engineering from the University of Rochester in 2016 and is currently working towards a PhD at the University of New South Wales in Canberra, Australia, in the Advanced Sensing Laboratory (ASL) research group. His

research interests include visual analysis, data visualizations, and polarization imaging.

Andrey S. Alenin is a lecturer at the School of Engineering and IT at UNSW Canberra, Australia. He received his PhD in optical sciences from University of Arizona in 2015. His polarimetry work focuses on efficient reconstruction techniques within channeled and/or partial polarimeters.

J. Scott Tyo is a professor of electrical engineering and the head of the School of Engineering and IT at UNSW Canberra. He has been a faculty member at the US Naval Postgraduate School, the University of New Mexico, and the College of Optical Sciences at the University of Arizona. His research group studies advanced optical sensing, focusing in polarimetry. He was the 2014 recipient of the SPIE GG Stokes award for his work on modulated polarimeters.