

Real-time Specularity Detection Using Unnormalized Wiener Entropy

Qing Tian and James J. Clark

Dept. of Electrical and Computer Engineering & Centre for Intelligent Machines

McGill University

Montreal, Quebec H3A 0E9, Canada

Email: {qtian, clark}@cim.mcgill.ca

Abstract—Specularity is a common phenomenon in the real world and can provide both challenges and opportunities for computer vision algorithms. In this paper, we propose a new method to detect specularities in real time. This method is based on an unnormalized version of the Wiener Entropy which is commonly used in audio spectral analysis. Experiment results demonstrate our proposed methodology’s efficacy in specularity detection on both natural and synthetic images. Its potential in the 3D movie industry and for helping compute stereo correspondence in the presence of specularities is described.

Keywords-specularity detection; unnormalized Wiener entropy; stereo matching; 3D movie;

I. INTRODUCTION

The visual world is enriched by the presence of specular reflections. From the glistening light of waves on lakes to the sparkling highlights that make fresh fruits mouth watering; from the shiny hood reflection of a luxury car to glistening lips in a modern makeup advertisement, specularities are everywhere. When light is incident on a surface, some of it penetrates to the interior. This part of the incident light can be divided into two components: a portion that is absorbed and is not reflected back, and a portion that is scattered within the bulk material and then reflected back to the surface. This reflected component is called the *diffuse* component of the total reflection. In addition to the light that penetrates the surface, some light is reflected immediately from the surface. This reflected portion of light is the *specular* component of the total reflection. In this paper, a *specularity* is loosely defined as a region whose proportion of specular reflection is high enough that the original diffuse spectrum is obviously changed.

Specularity is a common phenomenon in natural images, and constitutes an unignorable double-edged sword. On one hand, specularities contain information that can be used to infer geometrical structures. On the other hand, specularities present difficulties for many computer vision tasks such as shape from shading, stereopsis, motion analysis, object recognition, image segmentation and so on. While the approach proposed in this paper is generally applicable, our work is motivated by problems arising from 3D cinematography. In today’s movie industry, 3D has experienced growing popularity and more and more movies are shot in 3D. However, many computer vision problems

arise in 3D filming [4]. One of these problems is caused by specularities: due to the differentially polarizing effect of the beamsplitter in a mirror stereo camera rig, specular regions are usually captured differently along the two optical paths. This differential image capture results in mismatching image features in corresponding regions, and can produce binocular fusion problems for movie viewers as well as creating difficulties for technical post-processing of the film.

It is therefore of great importance to reliably detect specularities (preferably in real time). This paper describes a new real-time approach for specularity detection, based on an unnormalized form of Wiener Entropy. Prior specularity detection methods are reviewed in Section 2. In Section 3 and 4, the concept of unnormalized Wiener entropy (UW entropy) along with its application in specularity detection is introduced. Section 5 demonstrates some representative results of our proposed specularity detection method. In Section 6 our method’s application in the 3D movie industry is described and Section 7 presents a specularity aware stereo correspondence approach based on UW entropy.

II. PREVIOUS WORK

Many specularity detection algorithms have been developed over the years. Among them, the most popular category is based on color space analysis. Shafer [17] built a dichromatic model in the RGB color space where a collection of color vectors from a homogeneous dichromatic surface roughly lie in a dichromatic plane spanned by the specular color and diffuse color. The specular reflection can be obtained by simply projecting the pixel color onto the specular line. Later it is observed that these color vectors usually cluster in a ‘skewed-T’ shape and improvements are made in Klinker et al. [9]. However, some assumptions have to be made in this method. First of all, the ‘skewed-T’ shape only exists for well-constrained environments where the gradual change in color can be easily seen. For clipped specularities or specularities far away like those on a lake or specularities of the same color as the diffuse underlying component, it is hard to find such a clearly defined ‘skewed-T’ shape. Even if a gradual change exists, often the corner of the ‘skewed-T’ is not sharp enough for an accurate separation. What’s more, in order to fit a parallelogram, the pixel noise should be small and the distribution of pixel

color vectors within a homogeneous material must not be pathological. Last but not least, segmentation into patches of homogeneous diffuse color is required. Similarly, Gershon et al. [7] observed the phenomenon that a highlight typically involves a gradual additive color shift from the object color to the illuminant color. Based on this observation, they developed an algorithm to detect those shifts which are characterized by a “dog leg” behavior in color space between segmented regions. However, besides the previously mentioned drawbacks, there are additional shortcomings. Firstly, neighborhood information among all segmented regions has to be gathered. Also, the transition from the object color to the illuminant color should be smooth, meaning that this algorithm cannot deal with sudden changes in surface curvature. In addition, the algorithm cannot handle more than one light source unless the light sources share the same spectral power distribution. In Bajcsy et al. [1] color pixels measured in RGB are transformed into the S space where color variations are analyzed in terms of brightness, hue, and saturation. Though analysis becomes more intuitive and easier, in theory there is no information gain by just transforming between color spaces. In their method, the scene illumination must be singly colored and estimated using a white reference plate as an active probe, which not only complicates the process but also cannot guarantee an accurate enough estimation.

The previously mentioned methods involve expensive 3D color space analyses. To address the problem of complexity, a relatively fast approach based on a 2D diagram has been developed by Schluns and Koschan [16]. In their method, 3D color descriptors are transformed to a pair of 2D descriptors. Though faster than above approaches, its matte and highlight color estimation procedure is still complex and involves segmentation. Like most methods in this category, it expects objects in the scene to have very saturated colors and exist against a uniform background. Drew [5] showed that the Lambertian part of reflection forms an ellipsoid in the 2D color spaces and the specular part can be considered as outlier. Using a least-median-of-squares method, the ellipsoid can be found to recover the surface orientation due to its linear relationship with the RGB color. Specularities do not fall on this ellipsoid and thus can be detected as a side benefit. However, in this method, several important parameters are set in an ad-hoc manner, an aggregated set of illuminants with spectrally varying light sources is required, and to do regression reliably the non-specular regions should be dominant. Torres et al. [18] introduced a 2D MS diagram describing the relationship between intensity and saturation. They claim that specularities can be located in a well-defined region in this diagram by analyzing the shape of specular clusters. Though this method can work quite well for intensive ‘white’ specularities without any hue, it does not work in rest cases. Furthermore, the region mask selection and some important parameter settings are a little bit ad-hoc and

the segmentation cannot be avoided either. Similarly, in [11], thresholding on intensity and saturation is applied in order to spot the specular pixels. This method confines specular candidates into a particular region of high brightness and unsaturated color in the 2D space of intensity-saturation. However, the fact is that by mixing different specular and diffuse components specularities can actually appear in any color (not necessarily those of high brightness and less saturation). Thus, by using such an axis-aligned bounding region, they let go many genuine specularities. The two ‘not so bright’ specularities in Figure 1, some quite saturated specularities in Figure 5 will all serve as good examples. Different from above methods, Park [15] has presented a new color space transformation, which is similar to the RGB-SUV transformation in [13], where the illuminant direction becomes aligned with one axis. This transformation is advantageous for specular detection because unsaturated specular regions form planar color clusters in the new color space and each of the planar specular clusters corresponds to a specular region. In this way, the complexity has been reduced by one dimension and segmentation is no longer needed. In their work, how to detect saturated pixels that do not lie on above dichromatic planes is also described. However, the illuminant should be evaluated beforehand, which is ill-posed and a difficult task itself.

Color is one of the main features we can extract from a scene and is readily acquired with common digital cameras, perhaps explaining the large number of color-based specular detection methods. That said, many specular detection methods either ignore color (e.g. [3]) or combine it with other cues such as polarization information [19] or multiple views [10] or both [12] to detect specularities at the expense of more complicated hardware devices or more processing time. For the purpose of real-time, in this paper we do not bother to incorporate other physical cues and only adopt the available color information to detect specularities.

III. SPECULARITY AND UW ENTROPY

The effect of a specular reflection in an image over and above the diffuse reflection is twofold: it increases the intensity and it changes the color direction most if not all of the time. The approach to specular detection that we propose in this paper takes into account both of these effects.

Let us first consider the color direction change due to the specular reflection. We find that in general if the color of the illuminant is not aligned with the color of the diffuse reflecting surface, the effect of the specular component is a tendency to make the spectrum of the mixed light more uniformly distributed, or in other words, less spiky. (here the spectral flatness or uniformity has nothing to do with intensity or magnitude.) The reason is that most light sources in the real world, from the natural sunlight to man-made fluorescent light, have a quite uniform spectral distribution compared to the diffuse component. Even if the light

source's spectrum is not that highly uniformly distributed compared to that of the diffuse component, the resulting mixed spectrum will also get flatter. Take the light from a red led (which has a quite spiky spectrum) for example, when it shines on a different colored surface (for example greenish or bluish surface), the resulting channel values of the specular region also become more uniformly distributed than the underlying diffuse color. Even when the light source has exactly the same spectrum (channel distribution) as the diffuse surface, the resulting ratio over the channels will remain the same, but will not decrease. This may account for the biological fact that the specularity on the same colored object seems relatively less identifiable. In summary, the uniformizing tendency of spectrum is generally valid for most specularities.

In audio spectral analysis, Wiener Entropy is used to measure spectral uniformity or flatness [6], [14]. Suppose we model the spectrum of the imaged light as a distribution of spectral samples (channel values) S_k , which could be the set of N output values from a spectrometer, or the $N = 3$ RGB channels of a digital camera. The Wiener Entropy is then defined as the log of the ratio of the geometric mean of the S_k to the arithmetic mean of the S_k :

$$W = \log \left(\frac{\exp\left(\frac{1}{N} \sum_{k=1}^N \log S_k\right)}{\frac{1}{N} \sum_{k=1}^N S_k} \right) = \log \left(\frac{\sqrt[N]{\prod_{k=1}^N S_k}}{\frac{1}{N} \sum_{k=1}^N S_k} \right) \quad (1)$$

which can be written as:

$$W = \frac{1}{N} \log \left(\frac{\prod_{k=1}^N S_k}{\left(\frac{1}{N} \sum_{k=1}^N S_k\right)^N} \right) \quad (2)$$

The Wiener Entropy can be interpreted as the Kullback-Liebler divergence between $\{S_k\}$ and a uniform distribution. Often, in audio spectral analysis applications, the term inside the logarithm is termed the Spectral Flatness with the logarithm then providing a value expressed in decibels. The Wiener entropy can be seen to capture the (generally true) uniformizing tendency on the spectrum of the specular reflection, in that the value of the Wiener entropy should be higher (approaching zero) for specular regions as compared with nearby non-specular regions (assumed to have similar body reflection colors). However, as previously mentioned, in general the Wiener entropy, or the Spectral Flatness measure, can only capture one part of the effect of the specular component. To handle the intensity enhancing effect, we add a term as follows, proportional to the log of the intensity, as represented by the sum of the spectral components:

$$W' = W + \frac{1}{N} \log \left(\frac{1}{N} \sum_{k=1}^N S_k \right)^N \quad (3)$$

This will nominally cancel out the denominator inside the logarithm in equation (2), resulting in a loss of normalization. However, we write it as above to make clear the

effect of the added term. For computational simplicity, in the remainder of the paper we will omit the logarithm operation and the scale factor of $1/N$ noting that in so doing we obtain a monotonically related measure. This gives us what we call the *unnormalized Wiener entropy* (or UW entropy):

$$H = \exp(N \cdot W') = I^N \exp(N \cdot W) \quad (4)$$

which can be written as:

$$H = I^N \prod_{k=1}^N \left(\frac{S_k}{I} \right) = \prod_{k=1}^N S_k \quad (5)$$

where I is the *intensity*

$$I = \frac{1}{N} \sum_{k=1}^N S_k \quad (6)$$

In our unnormalized Wiener entropy measure, the spectral samples (channel values) are joined together in a multiplicative manner, which not only cares about the intensity change of each individual sample, but also takes into consideration the ratio or proportion among them. For example, in a trichromatic vision system, if the highest sensor response is normalized to one, each of the other two responses becomes the ratio of itself to the highest response. Thus the multiplication of the three reveals the overall ratio among all three sensor values. Dynamically speaking, when one of the sensor values is changed, depending on the proportion of the changed channel in all channels, the same amount of change can produce quite different results. If the change occurs in a channel whose value is of smaller proportion, the multiplication as well as channel ratio is changed more obviously than if the same amount of change is applied to the channel of bigger proportion. The addition of specular reflection to the diffuse component has similar properties: when a red light sheds on an 'already' red object the resulting specularity is not so obvious compared to the case where a different color light source is used. Thus, the proposed unnormalized Wiener entropy measure can also well reflect the different results of various combinations of diffuse and specular components.

From equation (4), we can see that the unnormalized Wiener entropy consists of two parts: one is the brightness related map I^N and the other is the product of the proportions of each channel, which we will refer to as the uniformity map, $\exp(N \cdot W)$ (which captures spectral flatness). In our algorithm the two parts of the proposed unnormalized Wiener entropy capture the color-direction-changing and intensity-increasing characteristics of specularities and work 'hand in hand' to successfully detect them. For example, consider Figure 1 where the left hand image is derived from a synthetic scene and contains four specularities.

The brightness map corresponding to the lefthand image is shown in the center image in this figure. In the brightness map we can see that the two most obvious specular




Figure 1. A synthetic image containing multiple specularities (left) with the associated brightness (center) and uniformity (right) maps.

highlights stand out from their neighborhood due to their high brightness. However, the other two specularities are relatively dim compared to the very bright non-specular regions and are drowned in the non-specular neighborhood. By simply and naively searching for bright pixels like [3] does, we cannot effectively identify these specularities. Therefore, brightness alone is not enough for specularity detection because it only focuses on the sum of the channel values and completely ignores their relationship. Fortunately, the uniformity part in the proposed unnormalized Wiener entropy can capture the other very important characteristic, which is specular region generally has a more uniform spectrum (or distribution over channels) and by combining it with the brightness information in a complementary yet compact manner, we can successfully identify the specularities. The uniformity map is shown in the righthand image of Figure 1. In the uniformity map, all four specular regions are of a relatively high value compared to non-specular regions, demonstrating the uniformity map’s ability to complement the brightness information in specularity detection. Figure 2 shows another example, this time on a natural image. The uniformity part helps the brightness part clearly pick out the many specularities in the scene.




Figure 2. A natural image containing multiple specularities (left) with the associated brightness map (center) and UW Entropy map (right).

It is worth noting that as mentioned previously many methods examining both brightness and saturation information have been proposed to detect specularities such as [18], [11]. However, quite different from those methods which separately threshold on brightness and saturation values, in our UW Entropy the brightness related part and the uniformity part have a cooperative and complementary relationship and one cannot survive without the other. For the image in Figure 1, while our UW Entropy can use the high uniformity to help identifying the two not so bright

specularities, the saturation information in naive thresholding methods (similar to our uniformity), though somewhat informative, cannot save the two not so bright specularities from being labelled as non-specular by the brightness test. Besides the case of not-so-bright specularities, later results in the experiment part will also justify our UW Entropy’s superiority in handling specularities that are quite saturated such as the reddish specularity on the teapot, the bluish specularity on the apple, and the yellowish specularities over the sea surface in Figure 5. What’s more, though similar, our uniformity part and saturation are definitely different. Our uniformity part is actually more suitable than saturation for specularity detection due to the previously described advantages of multiplication. Last but not least, our UW Entropy’s simple and compact form is highly desirable for real time applications such as those in the 3D movie industry.

However, we have to admit that though in most situations, as described previously, the uniformity part increases in specular areas as compared with nearby non-specular regions due to the additional specular component, it may fail in one rare case. The case is when the diffuse spectral component is highly uniformly distributed over the channels (gray objects for example) compared to the specular component. Only in this case, the uniformity part decreases. Now it is time for the brightness related part to help and UW entropy, as a whole, will still be successful in specularity detection. In this case relatively speaking the brightness (the number of additional photons loosely speaking) should be so high that the channel ratio is adequately changed to appear specular because there is no channel of obviously smaller proportion where just a tiny amount of additive specular component can make a considerable difference on the channel ratios. As can be seen, even in this case where the uniformity map is lower for specularities, the proposed unnormalized Wiener entropy still works because the brightness map’s significant contribution help to counteract the effect of the uniformity map and makes the unnormalized Wiener entropy higher in specular regions. The reddish specularity on the gray teapot in Figure 5 will serve as a good example of our method’s success in this case. Our method is thus seen to handle most possible cases encountered in practice, and is therefore superior to methods which naively assume a particular kind of light source or specular region color (e.g. nearly white and of low saturation).

IV. SPECULARITY DETECTION BASED ON UW ENTROPY

We can define a specularity as a region where the unnormalized Wiener entropy of the total reflection is much higher than that of the underlying diffuse component. Specifically we say that specular regions are detected when:

$$H(R_f(\lambda)) - H(D(\lambda)) > \tau' \quad (7)$$

where $R_f(\lambda)$ is the spectrum of the total reflection, $D(\lambda)$ is the spectrum of the diffuse component, H indicates the

UW entropy, and τ' is a particular threshold.

The spectrum of the light coming from specular regions do not only depend on the specular component but also on the underlying diffuse component. However, based only on the sensor responses of the total reflection such as a normal RGB photo, in theory it is almost impossible to separate the diffuse and specular component perfectly. To get around this problem, in this paper we take an alternative perspective with a simplifying assumption. From equation (7), we have

$$\begin{aligned} H(R_f(\lambda)) - H(D(\lambda)) &> \tau' \\ H(R_f(\lambda)) &> \tau' + H(D(\lambda)) \\ H(R_f(\lambda)) &> \tau \end{aligned} \quad (8)$$

where $\tau = \tau' + H(D(\lambda))$. Equation (8) suggests that instead of trying to directly measure the difference between the UW entropy of the total reflection and that of the diffuse component, we can consider just the total reflection as a whole. However, to make the threshold $\tau = \tau' + H(D(\lambda))$ universal, the UW Entropy of the diffuse component $H(D(\lambda))$ has to be constant over the region under consideration or at least having little variation. This is quite possible within one object. However, for the whole scene, different colored objects and backgrounds usually have very different colors and the UW Energies of them are most likely not the same. The good news is that even though in a scene the background's and the objects' UW entropies are not the same, most of the time they are quite similar compared to the much higher UW Entropies of the specular regions because in the real world most objects often have a color that has a small value in at least one of the color channels such as a red apple, green tree, yellow flower, brown desk, and so on. It follows that for an image a single universal threshold τ can be applied to the unnormalized Wiener entropy map to separate specular and non-specular regions without pre-segmentation. However, the determination of the threshold based on the image becomes an important process.




Figure 3. Local Surface Model

To proceed with the development of a threshold determi-

nation method, we make a number of modeling assumptions. The first is to locally consider the cross-section of a surface as a segment of a circular arc (a flat surface can be seen as a special arc whose radius is large). This is shown in Figure 3 where S is the light source position, O is the center of the circular arc AB , P_1 is the intersection point of the surface with the line SO , P_2 is another point on the surface, and the viewer V is assumed to be at S .

We also adopt the Blinn-Phong specular reflection model [2] and the Lambertian diffuse reflection model. Under these assumptions the intensity of channel i is:

$$\begin{aligned} S_i &= \alpha_s \int_{\lambda_1}^{\lambda_2} I_{nc}(\lambda) \rho_i(\lambda) d\lambda \cdot \cos^n \theta + \\ &\quad \alpha_{di} \int_{\lambda_1}^{\lambda_2} I_{nc}(\lambda) \rho_i(\lambda) d\lambda \cdot \cos \theta \\ &= \int_{\lambda_1}^{\lambda_2} I_{nc}(\lambda) \rho_i(\lambda) d\lambda \cdot (\alpha_s \cdot \cos^n \theta + \alpha_{di} \cdot \cos \theta) \end{aligned} \quad (9)$$

where $I_{nc}(\lambda)$ is the spectrum of the incident light, α_s is the specular albedo (often assumed to be channel independent in the literature so that the specular light has the same color with the light source), α_{di} is the diffuse albedo of channel i , $\rho_i(\lambda)$ is the spectral sensitivity of the i^{th} channel, and $\lambda \in (\lambda_1, \lambda_2)$ is the sensing range. The unnormalized Wiener entropy can thus be expressed as follows:

$$\begin{aligned} H_R &= \prod_{i=1}^N S_i = \left(\prod_{i=1}^N \int_{\lambda_1}^{\lambda_2} I_{nc}(\lambda) \rho_i(\lambda) d\lambda \right) \cdot \\ &\quad (\alpha_s \cdot \cos^n(\theta) + \alpha_d \cdot \cos \theta)^N \end{aligned} \quad (10)$$

where α_d is a value that lies somewhere between the maximum and the minimum of the channels' diffuse albedos α_{di} . The existence of such a value can be easily proved mathematically. The first part (the multiplication of integrals) does not change with the lighting and viewing direction. The last part, on the other hand, does depend on θ . This geometry related factor ($h(\theta) = (\alpha_s \cdot \cos^n \theta + \alpha_d \cdot \cos \theta)^N$) can be decomposed into a diffuse part ($f(\theta) = (\alpha_d \cdot \cos \theta)^N$) and a specular related part ($g(\theta) = h(\theta) - f(\theta)$). The variation of these components with θ are illustrated in Figure 4. It is worth noting that in addition to helping select the threshold, Figure 4 also explains our methodology's success in specularity detection from another perspective.

From Figure 4 we can see that the specular related part (the function in orange $g(\theta)$) falls off more quickly than the diffuse part (the function in red $f(\theta)$). When the geometry related factor $h(\theta)$ shown in blue (and also the UW entropy) decreases down to one half its maximum value (indicated by the blue spot in Figure 4), the specular related part $g(\theta)$ has already been weakened greatly while the diffuse part $f(\theta)$ has not changed very much. Thus, 0.5 is a reasonable threshold on the UW Entropy. If the UW Entropy decreases but to a value above 0.5, the specular related component




Figure 4. Threshold Selection ($n = 3$, $\alpha_s = 0.2$, $\alpha_d = 0.8$, $n = 60$ (a typical value according to [2])). The blue curve represents the geometric factor which varies depending on the viewing direction. The blue spot indicates the point at which the geometric factor is one half its maximum value. The red curve indicates the diffuse component of the geometric factor and the orange curve indicates the specular component.

has not been attenuated enough compared to the diffuse component, meaning that the image is still in the specular region. On the other hand, if the UW entropy decreases to a value below 0.5, then, according to Figure 4, the specular part begins to change very slowly while the diffuse part begins to decrease dramatically, which strongly indicates that the specular ‘boundary’ has already been reached. Also, by examining the histogram of UW Entropy distribution over many images with specularities, generally speaking, specular and non-specular pixels are gathered in two different clusters and a large gap between the non-specular pixels with low UW entropies and the specular pixels with high UW entropies can be found. Intuitively, this explains why a threshold of 0.5 successfully divides the specular and non-specular regions. As a matter of fact, in our experiments a threshold of 0.5 works well for all images tested and thus can be considered as a general threshold. It is also worth noting that because of this wide gap, the threshold can be chosen from a rather large range and 0.5 is just a convenient but in no way unique threshold. Also due to this large gap, the threshold selection method is quite robust to noise.

V. EXPERIMENTAL RESULTS

In our experiment, we have analyzed many images with specularities. Some of them are from peer-reviewed papers, some are collected over the Internet, while others are taken or simulated by ourselves. Typical examples of specular detection using the proposed UW Entropy based method are shown in Figure 5 for a mixed set of synthetic and natural images. In each of these examples, a general threshold value of 0.5 was used, relative to the maximum UW entropy value over the entire scene. As can be seen from those examples, our unnormalized Wiener entropy based method works quite well on both synthetic and natural scenes under illuminants with different colors and saturations. In addition,

the proposed method requires no extra information on the illuminant, object material, reflectance maps, or the imaging process. Given only a color image, the method can locate the specularities with very little processing.

It must be noted that the computational complexity of the proposed approach is very low. As can be seen from equation (5), the computation of the unnormalized Wiener entropy only requires the multiplication of the sensor channel values (e.g. 2 multiplications in the case of an RGB camera). The implementation of the thresholding operation is likewise of very low complexity. This means that the specular detection algorithm is extremely fast and easily implemented in simple hardware for real-time applications. However, since any specular color can be generated in a diffuse manner in theory, like all other specular detection methods based on color alone, our approach may detect some false positives sometimes. In this case, other indispensable physical cues must be adopted to remove false positives, which is beyond the scope of this paper.

VI. APPLICATION IN THE 3D MOVIE INDUSTRY

Modern cinema has undergone a revolution with the advent of high-quality 3D film-making. Currently 3D movies are created from stereo imagery captured with binocular camera rigs (in addition to synthetic imagery). The binocular imagery after post-production editing is presented to viewers typically using circularly polarized video projectors. Specularities are a very important issue in 3D cinema, for two main reasons. The first is that specular reflections generally appear somewhat different in the left and right eye images, due to the change in viewing direction between the eyes. It is important that this (usually slight) difference be captured and rendered accurately to provide a natural appearing 3D display. The second reason has to do with the way in which 3D imagery is captured for high-quality 3D cinema. Typically, in order to use high-quality fast camera lenses with large apertures while still having an inter-camera baseline comparable to the human one, images are captured with *mirror rigs*. A mirror rig employs a beam splitter consisting of a half-silvered mirror which passes half of the light directly through the mirror to one camera and the other half is reflected to another camera. One of the drawbacks of the mirror rig setup is that light passing along the two optical paths undergo a differential polarization [4]. The result is that when linearly polarized light, such as provided by specular reflection, enters the beam splitter, the reflected light may be attenuated relative to the light in the direct path. This means that in viewing a scene with specularity, one camera may capture the specularities while the other does not. A viewer looking at the resulting stereo imagery may have difficulty fusing the specular region due to the excessive mismatch. The use of quarter wave retarders can reduce this effect, but at the cost of reducing light levels and possibly affecting color balance.




Figure 5. Results of specular region detection on a mixed set of synthetic and natural images using the unnormalized Wiener entropy based method. The examples are presented in nine groups of 3 images: the leftmost image of each group is the original image, the middle image of the group is a false-color depiction of the unnormalized Wiener entropy, and the rightmost image of each group shows the detected specular regions after thresholding with 0.5 of the maximum entropy values.

It is important to be able to detect stereo mismatches due to specularities during the filming (in real time), so as to minimize expensive re-filming or post-production processing. The specularity detection method described in this paper is computationally inexpensive and so can be applied in real-time for monitoring by a stereographer. An example of this is shown in Figure 6 where (a) and (b) show a pair of images of the same scene, taken with perpendicular polarizing angles trying to maximize differences across views. The result of our specular detection is also shown in (c) and (d) and demonstrates the success of our method in finding the specular regions. This example also shows the effect of the polarizing filter in removing some of the specularities, which could potentially cause difficulties in binocular fusion for viewers of the stereo imagery.



Figure 6. Left column: A simulated pair obtained from a mirror rig camera setup (source:[4]). Right column: The output of our specularity detection method. Note the differences in specularity between the two images caused by the polarizing effect of the beam-splitter in the mirror rig.

The 3D cinematographer would want to know when stereo mismatches are present. To handle this, we have implemented a Fourier-domain shape matching algorithm trying to find matches between candidate specular regions in the stereo pair. If mismatches are found, the cinematographer has a number of options: he can re-take the imagery after modifying the set to remove the troublesome specularity; a synthetic specularity can be added in to the polarized channel to compensate for the attenuation; or the specularity in the non-polarized channel can be removed by doing inpainting of the specular region. The implementation of these steps is outside of the scope of this paper, however.

VII. APPLICATION TO COMPUTING CORRESPONDENCES IN THE PRESENCE OF SPECULARITY

Beyond the particular application to 3D cinematography, our specularity detection method also has a more general application to stereo vision research in aiding the computation of stereo correspondence in the presence of specularities. Many robust stereo algorithms have been developed to deal with specularities. For example, one class of approaches use mutual information (MI) as the similarity measure (e.g. [8]), which requires no prior information about the corresponding relationship and allows correspondences to be found in spite of some mismatches in brightness. However, improvements can be made if specularities are explicitly modelled, in our case through use of the unnormalized Wiener entropy. Inspired by Kim et al. [8], we presented a new specularity-aware visual correspondence method by embedding our UW entropy into the energy function. Different from the data term in [8] which actually does not incorporate any information about specularity, our specularity-aware mutual information data term is as follows.

$$Data_p^{MI}(d_p) = -\frac{k}{H(p)} \cdot \log(P_d(I_1(p), I_2(p + d_p))) \quad (11)$$




Figure 7. (a) (b) A stereo pair with specularities, which is part of the Tsukuba benchmark stereo pair (to focus on the regions with specularities and also to achieve a high resolution, other non-specular parts of the Tsukuba pair is excluded) (c) The disparity map produced by Kim’s method [8]. (d) The disparity map by our method. It is worth noting that for the two methods all shared parameters are set exactly the same.

where p indicates a pixel, $H(p)$ is our UW entropy at p , k is a constant, I_1 and I_2 are the two intensity images in a stereo pair, and P_d is the joint probability based on the current labeling d (representing disparity). Besides the introduction of $H(p)$, the joint probability distribution P_d is also constructed differently from [8] in our approach: instead of counting every correspondence equally, we count specular correspondences less according to their UW Entropies. Using the Potts model as the penalty term and the same optimization technique – MAP-MRF with graph cuts, we applied our specularity-aware MI-based method to the stereo pair shown in Figure 7. The results of our method and the method in Kim et al. [8] are compared. From the results, we can see that the specular regions have successfully fooled the method in [8] making it assign wrong disparity values to specular regions (indicated by red circles). On the other hand, our method is smarter and can calculate a correct disparity value even in the presence of specularities because specular information has been incorporated into the energy function and the graph cuts optimization method can make more intelligent decisions accordingly, which demonstrates our proposed methods efficacy and superiority.

VIII. CONCLUSION

In this paper, a simple yet powerful methodology detecting specular regions in images is introduced, based on the use of an unnormalized form of Wiener entropy. The approach has a low computational complexity, permitting its use in real-time applications such as monitoring and quality control of stereo imagery in 3D cinematography. We demonstrate that the method works well on both natural and synthetic imagery, and show how UW Entropy can be used to improve the performance of modern correspondence

computation algorithms in the presence of specularities.

REFERENCES

- [1] R. Bajcsy, S. W. Lee, and A. Leonardis. Color image segmentation and color constancy. in *Proceedings of SPIE Perceiving, Measuring and Using Color*, 1250:245–255, 1990.
- [2] J. F. Blinn. Models of light reflection for computer synthesized pictures. *SIGGRAPH Computer Graphics*, 11(2):192–198, 1977.
- [3] G. Brelstaff and A. Blake. Detecting specular reflections using lambertian constraints. *Computer Vision., Second International Conference on*, 26:297–302, 1988.
- [4] I. T. Doser, F. Hassenpflug, and R. Zwing. The glasses are half empty: Turn your sight to autostereoscopic 3-d. *Consumer Electronics Magazine, IEEE*, 1(1):49–60, jan. 2012.
- [5] M. S. Drew. Robust specularity detection from a single multi-illuminant color image. *CVGIP Image Understand*, 59:320–327, 1994.
- [6] S. Dubnov. Generalization of spectral flatness measure for non-gaussian linear processes. *IEEE Signal Processing Letters*, 11(8):698–701, 2004.
- [7] R. Gershon, A. D. Jepson, and J. K. Tsotsos. The use of color in highlight identification. in *Proceedings of the 10th International Joint Conference on Artificial Intelligence*, pages 752–754, 1987.
- [8] J. Kim, V. Kolmogorov, and R. Zabih. Visual correspondence using energy minimization and mutual information. in *Proceedings of the 9th IEEE International Conference on Computer Vision*, 2:1033–1040, oct. 2003.
- [9] G. J. Klunker, S. A. Shafer, and T. Kanade. The measurement of highlights in color images. *International Journal of Computer Vision*, 2:7–32, 1988.
- [10] S. W. Lee and R. Bajcsy. Detection of specularity using colour and multiple views. *Image and Vision Computing*, 10(10):643–653, 1992.
- [11] T. M. Lehmann and C. Palm. Color line search for illuminant estimation in real-world scenes. *J Opt Soc Am A Opt Image Sci Vis.*, 18(11):2679–2691, 2001.
- [12] S. Lin and S. W. Lee. Detection of specularity using stereo in color and polarization space. in *Proceedings of the 13th International Conference on Pattern Recognition*, 1:263–267 vol.1, 1996.
- [13] S. P. Mallick, T. E. Zickler, D. J. Kriegman, and P. N. Belhumeur. Beyond lambert: reconstructing specular surfaces using color. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2:619–626, 2005.
- [14] P. Niyogi and M. M. Sondhi. Detecting stop consonants in continuous speech. *The Journal of the Acoustical Society of America*, 111(2):1063–1076, 2002.
- [15] J. B. Park. Detection of specular highlights in color images using a new color space transformation. *IEEE International Conference on Robotics and Biomimetics*, pages 737–741, 2004.
- [16] K. Schlüns and A. Koschan. Global and local highlight analysis in color images. in *Proceedings of 1st International Conference on Color in Graphics and Image Processing*, 2000.
- [17] S. A. Shafer. Using color to separate reflection components. *Color Research and Application*, 10:210–218, 1985.
- [18] F. Torres, J. Angulo, and F. Ortiz. Automatic detection of specular reflectance in colour images using the ms diagram. *Computer Analysis of Images and Patterns*, 2756:132–139, 2003.
- [19] L. B. Wolff. Using polarization to separate reflection components. in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 363–369, 1989.