Perception-Based Image Editing

Diego Gutierrez
Universidad de Zaragoza
Instituto de Investigación en Ingeniería de Aragón
http://giga.cps.unizar.es/∼diegog/

***Abstract*—Image editing and post-processing techniques have matured over the years, making it difficult (verging on impossible) to assess whether an image has been digitally enhanced or doctored somehow. However, complex manipulations still rely on skilled user input, painstakingly painting over pixels.**

We present our recent and on-going work on advanced image editing techniques, extending current tools by leveraging the limitations of the human visual system and the wealth of information available in a high dynamic range image. Working in perceptual space, psychophysics plays an important role in assessing the validity of the results.

I. Introduction

The human visual system is not a linear light meter. How we interpret images (light) depends on multiple factors, some well-known, some still a mystery. Take for instance the image in Figure 1: we all see two spirals (one green, one blue) on a pink background. If we look closer, we will notice that there are also some orange strips. There does not seem to be a lot more in this image. Well, actually, we have seen *more* than there actually is: in reality, the green and blue colors are exactly the same! A quick Photoshop test will confirm this. So what is going on?

It turns out that our visual system is designed to interpret visual information relying heavily on contrast and other contextual information. In other words, we cannot tell the exact physical magnitude of, say, luminance (an objective magnitude). Instead, we can only judge *brightness* (a subjective measure), that is, we can only tell whether something is lighter or darker than its surroundings. The same concept applies to color: the ”green” spiral in Figure 1 is crossed by orange stripes, whereas for the ”blue” they turn magenta. So our brain computes color based on local information and comes out wrongly with two very different colors when there is only one.

Any image-editing algorithm that works in pixel-value space will miss out on the clear fact that the two spirals are perceived very differently, since the pixel values for both are exactly the same ((0, 255, 150) in RGB space, to be precise). In our group we are exploring algorithms that work in *perceptual space* instead, where there exists a clear distinction between the two spirals. Given that our perception, as we have seen, is not perfect, it makes sense to think that working in perceptual space we can sometimes get away with imperfect simulations. The key is to understand which imperfections will not be noticed by a human observer, and which will be easily spotted and thus must be avoided. In the following we will present this recent line of work, including both published results and

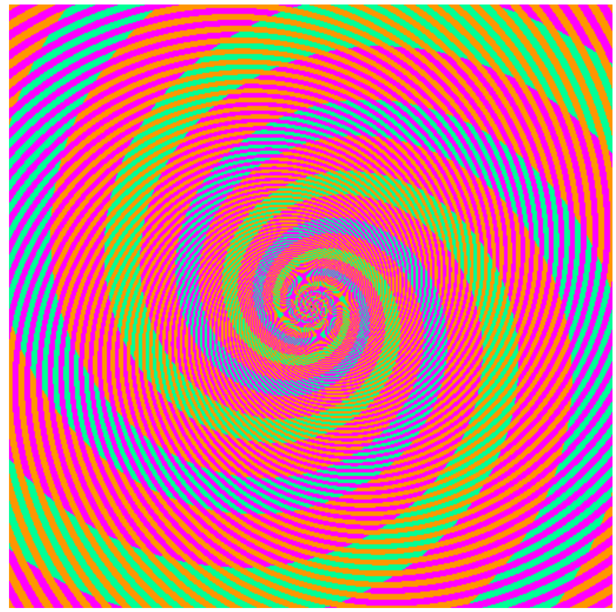

Fig. 1. The perceived green and blue spirals are just a visual effect. In reality, both colors are exactly the same. Image from http://blogs.discovermagazine.com/badastronomy/

on-going research. Our overall goal is to extend the set of tools available to artists to effect high level changes in single images, without the need to painstakingly paint over all pixels.

II. Perception-based image edits

Figure 2 shows a first example of an extreme image editing: procedural caustics are simulated in the image based on statistical information of the input image. First, an approximate depth map is obtained from the input image, which will be used as a rough representation of the objects geometry. Second, the recovered geometry is analyzed to establish likely caustic patterns that such an object may cast. This analysis takes the form of symmetry detection, for which we employ an algorithm that works in frequency space and makes minimal assumptions on its input. Finally, the luminance channel of the image is varied according to the projected caustic patterns [1].

Figure 3 shows another example of a complex image edit, which would require painstakingly painting over pixels by a skilled user. The image on the left is the original picture; on the right, the effect of thick fog has been simulated [3]. The algorithm leverages the findings by Narasimhan and Nayar [4], who model the effects of different kinds of atmospheric

Fig. 2. Example of light transport editing. Top left, original image. Top right, transparent mill following the approach in [2]. Notice the absence of caustics. Bottom: final result, with caustics added with our algorithm. Images from [1].

Fig. 3. Another example of complex light transport editing achieved through image editing techniques. Left: original image. Right: simulated fog in screen-space. Images from [3].

haze and fog by measuring their characteristic point-spread function. In our work, the user simply draws a mask separating foreground and background objects and sets some intuitive fog parameters: its corresponding point-spread function, plus color desaturation, are automatically applied based on the relative distance of the objects in the image.

In [5], we propose a novel skin shader which translates the simulation of subsurface scattering from texture space to a screen-space diffusion approximation, thus converting the expensive subsurface scattering computations into a fast image-based post-process. It naturally scales well while maintaining a perceptually plausible result. This technique allows to ensure real-time performance even when several characters may appear on-screen at the same time, whereas previous state-of-the-art algorithms [6] would perform more poorly as the number of translucent objects increase. Figure 4 shows some results.

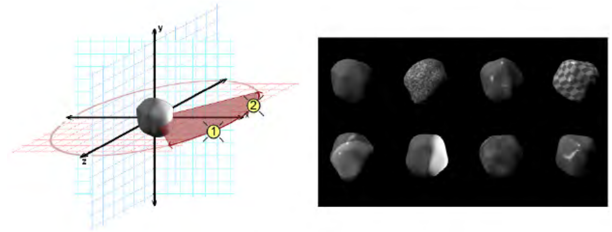

Fig. 6. Measuring the accuracy of human vision with respect to light detection. Left: scheme of two diverging lights applied to the same object. Right: one of the images shown to the participants in the user test. Images from [10].

III. Psychophysics and User Tests

When working in perceptual space, psychophysics and user tests become a crucial way to validate the results [7]. In the procedural caustics work [1], for instance, psychophysics were run to show how the results were perceptually on par with photon-mapped caustics, but without the need for any 3D geometry, whereas the skin shading results [5] were perceptually validated against a state-of-the-art algorithm [6]. In the case of the simulation of participating media [3], the user tests were designed to prove that the method could provide realistic looking results in a fraction of the time it would take (on average) using an image-editing program like Photoshop. An added advantage is that no skilled user input is required, since most of the algorithm is automatic.

Other times, psychophysics are used to gather information prior to actually designing the algorithms themselves. For instance, in the context or reverse tone mapping of low dynamic range *overexposed* images, psychophysical experiments showed that a simple operator based on gamma expansion avoids the errors introduced by other, more complex methods. With this gained insight, we proposed a method to automatically set a suitable gamma value for each image, based on the image key and empirical data [8], and use it as a simple, global reverse tone mapping operator. Figure 5 shows an analysis of the error introduced by existing operators, compared to our simple gamma expansion.

Finally, in [10] we rely on psychophysics to quantify a well-known aspect of human perception: its inability to detect light directions accurately in an image. Since that is actually an ill-posed problem for which no precise solution can be inferred, the goal is to understand the limits of our human visual system in order to design light detection algorithms within perceptual limits: as long as the error of the algorithm is less than the accuracy of our perception, the results, although physically inaccurate, will be *perceived* as correct.

IV. Conclusions

We have shown several results of our perception-based, image editing techniques, a line of work currently being developed at the Universidad de Zaragoza, with the collaboration of several other academic institutions and private corporations such as Adobe Inc. Even though it is still in

Fig. 4. Real-time, screen-space skin rendering. Image from [5].

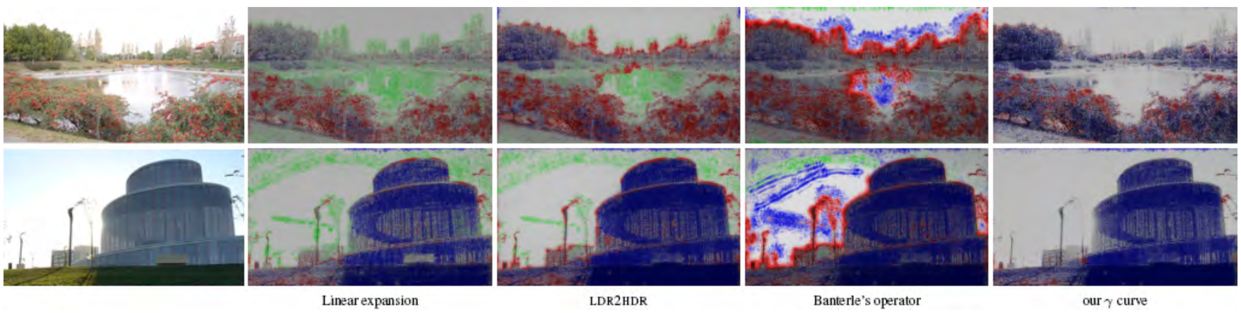

Fig. 5. Comparing the results of several reverse tone mapping operators with the image quality metric from Aydin et al. [9]. The left column shows the original, overexposed images. For the rest of the columns, green, blue and red identify loss of visible contrast, amplification of invisible contrast and contrast reversal respectively after range expansion. Our proposed γ expansion does not lose any contrast, while minimizing gradient reversals. More importantly, it reveals more detail in the most significant areas of the images (trees, grass, bushes and buildings in the images shown). Images from [8].

its infancy, we have already achieved both promising and stimulating results. Other on-going works include acquisition of subsurface scattering properties from single images, or development of non-photorealistic rendering techniques. This is particularly interesting, since we have found that the human visual system is more forgiving in a non-photorealistic context, and thus larger errors go unnoticed.

Acknowledgments

Many thanks to all my colleagues and collaborators, specially to my students: most of the results covered in this paper come from their hard work and dedication. This research has been partially funded by the Spanish Ministry of Science and Technology (TIN2007-63025) and the Gobierno de Aragón (project OTRI 2009/0411), as well as by a generous gift from Adobe Inc.

References

[1] D. Gutierrez, F. J. Seron, J. Lopez-Moreno, M. P. Sanchez, J. Fandos, and E. Reinhard, "Depicting procedural caustics in single images," *ACM Trans. Graph.*, vol. 27, no. 5, pp. 1–9, 2008.

[2] E. A. Khan, E. Reinhard, R. W. Fleming, and H. H. Bülthoff, "Image-based material editing," *ACM Trans. Graph.*, vol. 25, no. 3, pp. 654–663, 2006.

[3] J. Lopez-Moreno, A. Cabanes, and D. Gutierrez, "Image-based participating media," in *Congreso Español de Informatica Grafica, CEIG*, 2008, pp. 179–188.

[4] S. Narasimhan and S. Nayar, "Shedding light on the weather," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2003, pp. 665–672.

[5] J. Jimenez, V. Sundstedt, and D. Gutierrez, "Screen-space perceptual rendering of human skin," *ACM Transactions on Applied Perception*, 2009.

[6] E. d'Eon, D. Luebke, and E. Enderton, "Efficient rendering of human skin," in *Proceedings of Eurographics Symposium on Rendering*, 2007.

[7] J. A. Ferwerda, "Psychophysics 101: how to run perception experiments in computer graphics," in *SIGGRAPH '08: ACM SIGGRAPH 2008 classes*. New York, NY, USA: ACM, 2008, pp. 1–60.

[8] B. Masia, S. Agustin, R. Fleming, O. Sorkine, and D. Gutierrez, "Evaluation of reverse tone mapping through varying exposure conditions," *ACM Trans. Graph.*, vol. 28, no. 5, pp. 1–8, 2009.

[9] T. O. Aydin, R. Mantiuk, K. Myszkowski, and H.-P. Seidel, "Dynamic range independent image quality assessment," *ACM Trans. Graph*, vol. 27, no. 3, p. 69, 2008. [Online]. Available: http://doi.acm.org/10.1145/1360612.1360668

[10] J. Lopez-Moreno, F. Sangorrin, P. Latorre, and D. Gutierrez, "Where are the lights? measuring the accuracy of human vision," in *Congreso Español de Informatica Grafica, CEIG*, 2009, pp. 145–152.