Cross dissolve without cross fade: further details

Mark Grundland [*] Rahul Vohra Gareth P. Williams Neil A. Dodgson

Computer Laboratory, University of Cambridge, Cambridge, United Kingdom

Overview

This document contains additional explanatory text supporting our Eurographics 2006 paper:

> "Cross dissolve without cross fade: preserving contrast, color and salience in image compositing", M. Grundland, R. Vohra, G. P. Williams, and N. A. Dodgson, *Computer Graphics Forum*, **25**(3), E. Gröller and L. Szirmay-Kalos (guest editors), 2006

That paper is terse in order to fit within the ten page limit set by Eurographics. This document provides more detailed background to this research and an extended discussion of related work.

Contents

---

[*] Contact: `mark@eyemaginary.com`, `http://www.eyemaginary.com/compositing`

1 Background

From the advent of photography, images have been combined with images, an artistic technique known as photomontage [Ade86]. In the nineteenth century, combination printing processes were developed to composite foreground figures with background landscapes since early cameras could not capture both in sharp focus. Before long, the photographic evidence for the existence of ghosts started to appear. Artistically, photomontage applies contrasting images to juxtapose ideas, overlaying images to express layers of meaning. This use of composite imagery was pioneered during the 1920s by the Dada movement in irreverent collages and, by the 1930s, photomontage was appearing in the popular posters of war propaganda. Today, composite imagery has become so ubiquitous across film and advertising that the audience hardly takes any notice of its presence. However, visual artists are always on the lookout for the means to make the old appear new again.

Before digital imaging, ensuring good tone and smooth transition during cross dissolve required considerable skill, so that mathematical models [Sal68] were developed to assist in camera shutter control. Digital image compositing traces its roots [Smi95] to the 1977 invention by Smith and Catmull of the alpha-channel to encode the transparency of each pixel, enabling an image to represent both the colour and the shape of an object. Porter and Duff [PD84] enumerated the fundamental image blending modes, the operators that determine how semitransparent images can be combined. All based on linear interpolation, these ideas have shaped the basic infrastructure of computer graphics, becoming incorporated into many image storage formats as well as image editing systems. The practical aspects of image compositing are well covered in Brinkmann's textbook [Bri99]. The efficient implementation of image compositing is detailed in Blinn's excellent tutorials [Bli94a, Bli94b]. Our work reformulates the linear interpolation operation used in standard image compositing, showing how better to account for contrast, colour and salience with only a moderate amount of extra computation added to the process. Compared to general purpose approaches to colour image blending, our techniques exhibit improved image quality. Compared to special purpose approaches to colour image blending, our techniques exhibit greater flexibility by being able to simultaneously blend multiple, independent images with variable degrees of transparency. Of course, by being more widely applicable, our techniques may be less adapted to the constraints of specific applications.

There is a lack of general purpose models of image blending which do not display some intrinsic tone or contrast bias. Consider the basic task of rendering a composite with equal contributions from multiple, independent component images. For simplicity, assume the images have good tone balance, exhibiting approximately uniform luminance histograms. There are two basic mathematical models of combining information, averaging and selection, and two basic models physical models of light interaction, absorption and emission. When averaging linearly interpolates between pixel values, the composite tends to lose contrast with the inclusion of each component, eventually turning grey. Conversely, when selection keeps the pixel value with the highest absolute magnitude, with neutral grey as the origin of the tone scale, the composite tends to gain contrast with the inclusion of each component, eventually turning into grainy noise. When images are seen as stacked sheets of a light absorbing material, the composite tends darken with the inclusion of each component, eventually turning black. Conversely, when images are seen as stacked sheets of a light emitting material, the composite tends brighten with the inclusion of each component, eventually turning white. By contrast, none of our novel image blending techniques display a predetermined preference that favours any single colour, while each one has a parameter that places the degree of contrast gain or loss under user control. For example, by default, our contrast preserving technique generates composites that reproduce both the mean tone and contrast of their components.

2 Related work

Many previously proposed alternatives to linear interpolation are not suitable for compositing colour images. The application of realistic models of mixing colour pigments, such as watercolours [CAS*97], requires the specification of paint parameters that are normally unavailable in image compositing. On the other hand, extending binary imaging frameworks, such as boolean operators [Har92], to handle colour images has proven difficult without losing the algebraic properties that make them attractive in the first place. As binary images represent shapes, blending between them is a contour interpolation problem. Applying mathematical morphology, it can be solved by recursively calculating the unknown median contour that lies mid way between two known contours [BMT94]. This method can be extended to greyscale by treating a greyscale image as a height map and separately considering its contours at each elevation [NA99]. While a level-set approach [Whi00] can effectively model greyscale image blending, it does not readily extend to processing colour images. Due to the difficulty of imposing an ordering relation on image colours, applying morphological processing to colour image blending can lead to severe artefacts, so that image smoothing and blending by linear interpolation is required to reduce their visibility [Iwa02].

Image blending can be performed on a variety of different image representations. As an image can be described through a gradient field and its boundary conditions, image blending can be performed by combining colour gradients instead of colour values [PGB03, ADA*04]. Reconstructing the composite image to fit its prescribed gradients requires solving a sparse linear system of equations with as many unknowns as pixels. Hence, gradient domain approaches are substantially more complicated to implement and calculate than the techniques presented in our paper. Also, as colour fidelity of gradient domain methods depends on the treatment of boundary conditions, they can occasionally exhibit colour artefacts [PGB03] or require additional user input to ensure correct colour handling [ADA*04]. Alternatively, multiresolution image pyramids [BA83] apply filtering and subsampling to decompose an image into successive levels of detail, enabling image blending to be performed at each level independently. In effect, only image features of comparable scales are directly combined with each other. By allowing low frequency image structures to be blended over larger regions than high frequency image details, this approach improves the visual coherence of the composite. Our techniques can take advantage of the popular Gaussian low-pass and Laplacian bandpass image pyramids [BA83] to perform additive colour mixing [KLL96] on RGB colour channels. It would also be possible to extend our work to provide additional interesting visual transitions by blending different image frequencies at different times [SH88].

Image stitching [Sze05] combines image fragments without revealing the seams between them. Image stitching methods are generally more suited for cloning opaque images, while our techniques are designed to support mixing semitransparent images. We use continuous image mattes to specify the opacity of component images while image stitching often starts out with binary image masks to describe the shape of component images. Apart from image and texture cloning [BA83, NOT98, PGB03, KSE*03], image stitching is also used to piece together coherent mosaic images [Mil77, BA83, SHC01, ADA*04, LZP*04, Sze05, ZLP*06] from aligned views of the same scene with varying viewpoint, timing, focus or illumination. To reduce the visibility of the seams between component images, their path can be optimised by using a dynamic programming [Mil77] or a graph-cut algorithm [KSE*03, ADA*04] to partition the image pixels. When compared to other image stitching methods, seam minimisation in the gradient domain appears most effective [LZP*04, ZLP*06]. A wavelet approach to seam optimisation [SHC01] considers the trade-off between the distortion of the components and the coherence of the composite, an intriguing strategy that could be adapted to image mixing. Across a variety of image representations, linear interpolation is often used to blend the image elements near the seams. As a result, several authors [NOT98, KSE*03, LZP*04, Sze05, ZLP*06] report visible blurring and contrast loss along the seams. Our image blending techniques could prove useful in some of these situations. Another solution [NOT98] is to overlay high contrast image features on top of the seam.

Image fusion [ZB99] integrates sensor images into an informative composite. Unlike our techniques, which do not require the component images of a composite to be related, image fusing methods operate on aligned views of the same scene captured by various means. While we allow the user to specify the relative contribution of each component, image fusion often relies entirely on the content of the sensor images to determine their relevance. Apart from the visualisation of remote sensing and multispectral imagery [BK93, PX99], image fusion can also be used to merge photographic images [BK93, ZB99] taken under varying exposure, focus, or illumination conditions. Using a wavelet decomposition [ZB99] or a multiresolution image pyramid [BK93], these techniques derive the composite by choosing at each resolution the component image features with the best contrast. As previously observed [BK93, PX99], in the presence of conflicting high contrasts, selection causes instability and averaging causes cancellation. On possible solution [PX99] is to apply selection to high frequency, foreground details and apply averaging to low frequency, background structures.

Finally, image blending for image mixing is employed by various graphics algorithms. Most notably, image morphing [Wol98], used in image based rendering methods [SD96], creates a smooth visual transition between images by first applying image warping to align them and then applying image blending to interpolate between them. For image morphing to be convincing, user interaction is typically required to specify the visual correspondence between the images. When this correspondence proves insufficient to account for image differences, such as regions of incompatible texture, image blending by linear interpolation is well known to cause contrast and ghosting artefacts [SD96, Whi00, Iwa02]. Our saliency preserving image blending technique could prove useful in reducing these problems. Also it is applicable to rendering pictorial summarisations of video sequences [MB96] or image collections [RKK*05].

References

[ADA*04] AGARWALA A., DONTCHEVA M., AGRAWALA M., DRUCKER S., CURLESS B., SALESIN D., COHEN M.: Interactive digital photomontage. In *Proceedings of SIGGRAPH*, (2004), 294-302.

[Ade86] ADES D.: *Photomontage*, Thames & Hudson, 1986.

[BA83] BURT P. J., ADELSON E. H.: A multiresolution spline with application to image mosaics. *ACM Transactions on Graphics* 2, 4 (1983), 217-236.

[BK93] BURT P. J., KOLCZYNSKI R. J.: Enhanced image capture through fusion. In *Proceedings of the International Conference on Computer Vision*, (1993), 173-182.

[Bli94a] BLINN J. F.: Compositing 1: Theory. *IEEE Computer Graphics & Applications 14*, 5 (1994), 83-87.

[Bli94b] BLINN J. F.: Composting 2: Practice. *IEEE Computer Graphics & Applications 14*, 6 (1994), 78-82.

[BMT94] BARRETT W., MORTENSEN E., TAYLOR D.: An image space algorithm for morphological contour interpolation. In *Proceedings of Graphics Interface*, (1994), 16-24.

[Bri99] BRINKMANN R.: *The Art and Science of Digital Compositing*, Morgan Kaufmann, 1999.

[CAS*97] CURTIS C. J., ANDERSON S. E., SEIMS J. E., FLEISCHER K. W., SALESIN D. H.: Computer-generated watercolor. In *Proceedings of SIGGRAPH* (1997), 421-430.

[Har92] HARRINGTON S. J.: Generalizations of boolean imaging operations to the continuous-tone domain. *Journal of Electronic Imaging 1*, 3 (1992), 266-276.

[Iwa02] IWANOWSKI M.: Image morphing based on morphological interpolation combined with linear filtering. *Journal of WSCG 10*, 1 (2002), 233-240.

[KLL96] KONIK H., LOZANO V., LAGET B.: Color pyramids for image processing. *Journal of Imaging Science & Technology 40*, 6 (1996), 535-542.

[KSE*03] KWATRA V., SCHODL A., ESSA I., TURK G., BOBICK A.: Graphcut textures: Image and video synthesis using graph cuts. In *Proceedings of SIGGRAPH*, (2003), 277-286.

[LZP*04] LEVIN A., ZOMET A., PELEG S., WEISS Y.: Seamless image stitching in the gradient domain. In *Proceedings of the European Conference on Computer Vision*, Lecture Notes in Computer Science, (2004), vol. 3024, part 4, 377-389.

[MB96] MASSEY M., BENDER W.: Salient stills: Process and practice. *IBM Systems Journal 35*, 3-4 (1996), 557-573.

[Mil77] MILGRAM D. L.: Adaptive techniques for photomosaicking. *IEEE Transactions on Computers 26*, 11 (1977), 1175-1180.

[NA99] NOVINS K., ARVO J.: The morphological cross-dissolve. In *Proceedings of SIGGRAPH: Conference Abstracts and Applications*, (1999), 257. Technical Sketch.

[NOT98] NAKAMURA K., OHKI M., TOTSUKA T.: Image blending by feature overwrite. In *Proceedings of the International Conference on Image Processing*, (1998), vol. 1, 226-230.

[PD84] PORTER T., DUFF T.: Compositing digital images. In *Proceedings of SIGGRAPH* (1984), 253-259.

[PGB03] PEREZ P., GANGNET M., BLAKE A.: Poisson image editing. In *Proceedings of SIGGRAPH* (2003), 313-318.

[PX99] PETROVIC V., XYDEAS C.: Computationally efficient pixel-level image fusion. In *Proceedings of Eurofusion*, (1999), 177-184.

[RKK*05] ROTHER C., KUMAR S., KOLMOGOROV V., BLAKE A.: Digital tapestry. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*, (2005), vol. 1, 589-596.

[Sal68] SALT B. G. D.: Mathematics in aid of animation. In *The Technique of Film Animation*, 2nd ed., (1968), 333-345.

[SD96] SEITZ S. M., DYER C. R.: View morphing. In *Proceedings of SIGGRAPH*, (1996), 21-30.

[SH88] STEIN C. S., HITCHNER L. E.: The multiresolution dissolve (image processing). *SMPTE Journal 97*, 12 (1988), 977-984.

[SHC01] Su M.-S., Hwang W.-L., Cheng K.-Y.: Variational calculus approach to multiresolution image mosaic. In *Proceedings of the International Conference on Image Processing*, (2001), vol. 2, 245-248.

[Smi95] Smith A. R.: Alpha and the history of digital compositing. *Microsoft Tech Memo 7*, 1995.

[Sze05] Szeliski R.: Image alignment and stitching. In *Handbook of mathematical models in computer vision*, (2005), 275-294.

[Whi00] Whitaker R. T.: A level-set approach to image blending. *IEEE Transactions of Image Processing 9*, 11 (2000), 1849-1861.

[Wol98] Wolberg G.: Image morphing: A survey. *Visual Computer 14*, 8-9 (1998), 360-372.

[ZB99] Zhong Z., Blum R. S.: A categorization of multiscale-decomposition-based image fusion schemes with a performance study for a digital camera application. *Proceedings of the IEEE 87*, 8 (1999), 1315-1326.

[ZLP*06] Zomet A., Levin A., Peleg S., Weiss Y.: Seamless image stitching by minimizing false edges. In *IEEE Transactions on Image Processing 15*, 4 (2006), 969–977.