

Interactive Shadow Editing from Single Images

Han Gong, Darren Cosker

Department of Computer Science, University of Bath

Abstract. We present a system for interactive shadow editing from single images which includes the manipulations of shape, distribution, sharpness and darkness of shadows according to the features of existing shadows. We first obtain a shadow-free image, shadow boundary and its registered sparse shadow scales using an existing shadow removal method. The modifiable features of the shadow are synthesised from the sparse shadow scales. According to the user-specified shadow-shape and its attributes, our system generates a new shadow matte and composites it into the original image, while also allowing editing of existing shadows. We share our executable for open comparison in community.

1 Introduction

Shadows are ubiquitous in natural images. Advanced shadow editing, which may include the manipulations of different shadow properties, is often desirable for Graphical artists. A shadow generally consists of two parts which are the umbra and penumbra. The umbra is the darkest region of shadow with constant illumination while the penumbra, i.e. shadow boundary, has transitioning illumination between the fully dark and lit area. Shadow effects can be represented as a multiplicative scale field, i.e. a matte. An image \mathcal{I}_c can be represented as a Hadamard product of a shadow scale field \mathcal{S}_c and a shadow-free image \mathcal{I} as follows:

$$\mathcal{I}_c = \mathcal{I} \circ \mathcal{S}_c \quad (1)$$

where c is a channel of RGB colour space. The lit area's scale are 1 and the other areas' scales are between 0 and 1.

In this paper, we propose a system for interactive shadow editing from single images which synthesises the features of existing shadow and preserves the naturalness of newly generated shadows. Our system provides friendly and flexible user-controls for defining the shape, darkness, softness, and colour of existing or new shadows. Furthermore, it does not require users to manually analyse and adjust the shadow properties in an image, as the manipulatable shadow properties are automatically synthesised by our system. Compared with the artificial shadows generated by Computer Graphics rendering, our shadow modification is based on the existing real shadows and our newly modified and added shadows can be highly consistent with the original shadows in a scene. The potential usages of this tool include but are not limited to: (1) moving existing objects and their shadows from one image to another; and (2) artistic modification of shadow.

This paper focuses on a shadow editing framework that utilises the output information from an existing shadow extraction method. Based on a state-of-the-art shadow removal dataset [1], we have both visually and quantitatively tested our system on various shadow scenes, which includes scenes with strong texture background, broken, soft and coloured shadows.

2 Related work

Naive shadow editing using existing image manipulation software, e.g. direct change of brightness or blurring for shadow boundaries, requires a considerable amount of manual adjustment to align the appearances of modified shadow to the original shadow. This alteration requires delicate editing steps and unavoidably results in unnatural artefacts around shadow boundaries. These features of shadow editing are not available in current image manipulation software, e.g. Photoshop and GIMP.

Detection and removal of shadow are the prerequisites for extracting the features of a shadow. Approaches to shadow removal from single images can be categorised as either automatic [2–4] or user-aided [5–10, 1]. The difference between these is whether the shadow detection process is automatic or user-guided. Automatic shadow removal methods generally rely on intrinsic image decomposition [2, 3] or shadow feature learning [4]. User-aided methods generally achieve more accurate and reliable shadow detection results at the practical cost of varying degrees of user input, such as quad map [5], shadow boundary [6, 8], and sample of intensity [9, 10, 1]. As for the removal phase, some of them [2–5, 7] adopt pixel-wise optimisation while others [11, 6, 8–10, 1] rely on the analysis of intensity profiles passing through shadow boundary.

Recent shadow removal work [5, 7, 8] presents some basic examples of shadow editing including complete removal, duplication, distortion and sharpness adjustment of the original shadow. The applications of this work only apply simple image manipulation to the original shadow matte and do not provide a manipulatable model for arbitrary shadow modification. A shadow editing system is first proposed in [11] which requires users to specify some boundary points of shadow. Their shadow edge model is manipulatable and supports controls for sharpness, darkness and shape of shadows. However, users are only allowed to move the specified boundary points of shadow which limits its range of amendable shapes and are not allowed to add new shadow segments. In summary, natural shadow editing for arbitrary shapes and properties is still unexplored.

2.1 Contributions

Given our review of state-of-the-art approaches, we propose the following two contributions:

1) A model for shadow editing from single images We propose a model that analyses the features of an existing shadow in an image and provides parameters for users to edit the extracted shadow and add new shadows. Unlike

previous work (e.g. [11]), our shadow editing model is fast, highly flexible and provides many more controls of shadow properties. The model is also universal and compatible with several shadow removal methods.

2) Easy user interaction design for shadow editing An easy interface is proposed for users to freely and quickly define the shape of shadow and control various properties of shadow.

To summarise, we believe our contributions are important in this area due to our significant improvements in extracting and utilising controllable properties from existing shadow and the ease in which a user may interact with a shadow scene.

3 Interactive Shadow Editing Model

In this section, we overview our algorithm first in brief then expend on technical details for each step. Our algorithm consists of three steps (see Fig. 2):

1) Pre-processing (§3.1) A shadow-free image, shadow boundary information, and shadow scale attenuation profiles are first obtained. This information is converted into a field of shadow softness and a synthesised intensity attenuation profile.

2) Synthesis of shadow matte (§3.2) Based on the pre-processed information, a shadow model with various tunable parameters is generated. A modified shadow matte is synthesised using a distance transform.

3) Composition of shadow (§3.3) The modified shadow is seamlessly composited into the original shadow image.

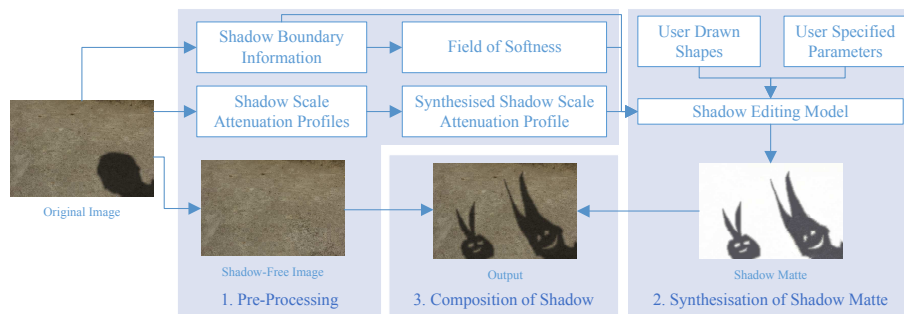


Fig. 1. Our shadow editing pipeline. The 3 main blocks in this chart correspond to the 3 main steps of our shadow editing system.

3.1 Pre-Processing

Pre-processing provides a shadow-free image, shadow boundaries, a field of shadow softness and a synthesised shadow scale attenuation profile required by our shadow editing model.

Shadow extraction Shadow extraction is the first step for shadow editing which provides necessary information for realistic shadow re-creation. To extract shadow from a single image, a state-of-the-art shadow removal method [1] is applied. This shadow removal method gives a shadow-free image, a shadow mask, and many shadow scale attenuation profiles perpendicular to the shadow boundary. Some other shadow removal methods [11, 6, 8–10] rely on intensity profile analysis are also compatible to our system.

Synthesisation of shadow scale attenuation Shadow creation is simplified by synthesising intensity attenuation. All extracted shadow scale profiles are re-scaled to a unique length (length of the longest profile by default). A synthesised shadow scale attenuation profile is then generated by profile-wise averaging all re-scaled shadow scale profiles. The averaging process cancels texture noise. To accelerate the computation for variable penumbra widths, a look-up table like profile is pre-computed. The synthesised shadow scale attenuation profile is evenly re-sampled to a large number (1000 in our implementation) of data sites using a piecewise cubic Hermite polynomial [12]. Shadow scales on wide penumbra can thus be queried by finding the closest scale from the large number of data sites.

Generation of softness field When the shape of shadow is changed from the original shadow image, it is unknown how soft the penumbra of the newly created parts should be. This is solved by generating a field of shadow softness from known penumbra widths of original shadow boundary points. The problem is equalised to in-painting an image with known pixel values (shadow scale profile lengths) at original shadow boundary points and unknown values for the remaining of pixels. A spring-metaphor based in-painting method [13] is adopted to smoothly interpolate and extrapolate the unknown values.

3.2 Synthesisation of shadow matte

Our goal is to generate a shadow matte according to user specified shadow properties. Given a shadow mask, whether the original or not, its crisp shadow boundaries can be located using a method for boundary tracing from a binary image [14]. The resulting boundaries should exclude the boundaries of the image border. A typical solution to generate a shadow matte is to re-generate the sparse scales of sampling lines perpendicular to shadow boundaries and in-paint for the other unknown shadow scales in image like the algorithm to form a dense shadow

scale field from sparse scales described in [1]. However, this can be computationally costly for interactive performance as the in-painting process is comparatively slow and every modification of shadow revokes the in-painting process. Instead, a linear time Euclidean distance transform [15] is used to generate the shadow matte quickly. The procedure for the generation of shadow mattes is described in Algorithm 1 where χ is a function computes Euclidean distance transform [15]

Algorithm 1: Generation of Shadow Matte

```

input : point set of shadow boundary  $\mathbf{B}$ , field of shadow softness  $F$ ,
        synthesised shadow scale attenuation profile  $A$ , shadow mask  $N$ 
output: shadow matte  $S$ 
1 Initialise  $M$  as a zero matrix in the size of original image;
2 Initialise  $S$  as a matrix of ones in the size of original image;
3  $M(\mathbf{B}) \leftarrow 1$ ; /* mark shadow boundary points */
4  $(D, L) \leftarrow \chi(M)$ ; /* perform a Euclidean distance transform */
5 foreach point  $p \in \mathbf{B}$  do
6    $E \leftarrow \{x | L(x) = \phi(p)\}$ ; /* find pixels closet to current point */
   /* divide the point set  $E$  into lit and shadow parts */
7    $E_l \leftarrow \{x | x \in E \cap N(x)\}$ ; /* lit part */
8    $E_s \leftarrow \{x | x \in E \cap x \notin E_l\}$ ; /* shadow part */
   /* perform distance conversion */
9    $R \leftarrow D(E)$ ;
10   $R(E_l) = \max(0.5 + D(E_l)/F(p), 1)$ ;
11   $R(E_s) = \min(0.5 - D(E_s)/F(p), 0)$ ;
12   $r \leftarrow 1000$ ; /* number of cached intensity attenuation data sites */
13  foreach RGB colour channel  $c \in \{1, 2, 3\}$  do convert distance to shadow
   scale
14  |  $S(E, c) \leftarrow A(\max(\kappa(rR), 1), c)$ ;
15  end
16 end

```

(its first output D refers to a matrix of Euclidean distance and its second output L refers to a label index matrix of closest boundary point), ϕ is a function which returns the index of a pixel in image, κ is a function that truncates each value of a set of data.

Algorithm 1 already handles variable shadow shape. To provide more controls for softness, darkness, and colour of shadow, more parameters are added to enable tuning:

Softness of Shadow A parameter of global softness adjustment is provided. This is achieved by simply multiplying the field of softness F in Algorithm 1 with a scaling factor v_s as follows:

$$F_n = \max(Fv_s, 2) \quad (2)$$

where F_n is the modified field of softness and an additional saturation operation is applied to ensure the width of penumbra is at least 2 pixel wide.

Darkness of Shadow A parameter for darkness adjustment is provided. This is achieved by modifying the synthesised shadow scale attenuation profile A in Algorithm 1 using a scaling factor V_d which controls the variation of shadow scale attenuation. The profile A is modified as follows:

$$\begin{cases} m(c) = \mu(A(1..r, c)) \\ A_n(1..r, c) = \max(v_d(A(1..r, c) - m(c)) + m(c), 0) \end{cases} \quad (3)$$

where A_n is the updated profile, c is the RGB channel index, r is the number of pre-computed data sites in Algorithm 1, μ is a function which computes the mean of a set of values. An additional saturation operation is applied to ensure the shadow scale is non-negative.

Colour of Shadow Three parameters of colour adjustment are provided. Each of them controls the intensity strength of its RGB colour channel. The 3 parameters are represented as a 1-by-3 vector V_k . The adjustment is done by further modifying Eq. 3 as follows:

$$\begin{cases} m(c) = \mu(A(1..r, c)) \\ o_k(c) = A(1, c) \\ d_k(c) = V_k(c) - o_k(c) \\ q_k(c) = (1 - V_k(c))/(1 - o_k(c)) \\ A_n(1..r, c) = \max(q_k(c)v_d(A(1..r, c) - m(c)) + m(c) + d_k(c)/2, 0) \end{cases} \quad (4)$$

where o_k is an intensity vector representing the original colour of the shadow, d_k and q_k are the error and ratio between original and adjusted colours of shadow respectively which change the variation and mean of shadow scale attenuation profile A .

After the adjustment, A_n may contain shadow scale values greater than 1, it is thus normalised by dividing all data sites by its lit end value as follows:

$$A_n(1..r, c) = A_n(1..r, c)/A_n(r, c) \quad (5)$$

3.3 Composition of Shadow

According the Eq. 1, the edited shadow can be composited using the Hadamard product of the edited shadow matte and the shadow-free image. This approach of composition ensures that the original properties of the surface background, e.g. texture and reflectance, are preserved.

4 User Interaction

In this section, we describe our user interaction and its related algorithm. Fig. 2 shows the prototype of our user interface. This interface can be divided into a drawing section and a configuration section.

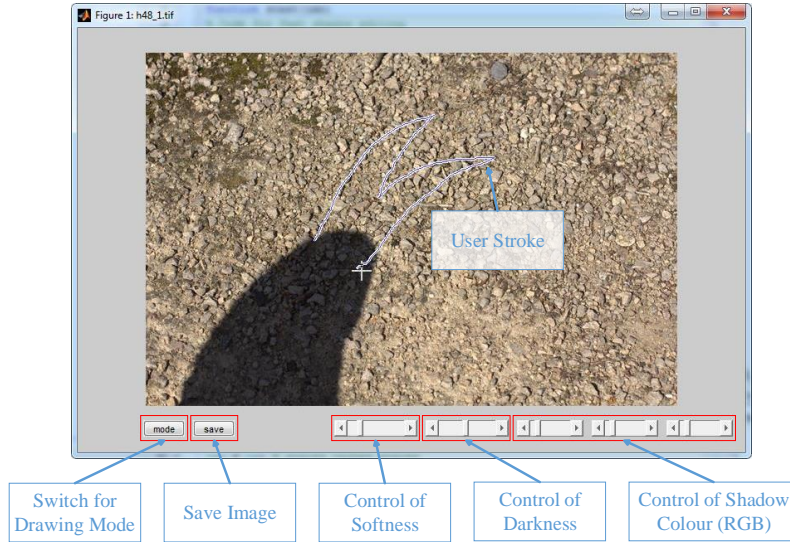


Fig. 2. Graphic User Interface.

4.1 Drawing Section for Shape Modification

In previous work [11], users are required to specify multiple sparse control points defining shadow boundary. This user interaction requires multiple delicate mouse clicks to ensure accuracy of shadow modification. After the initial specification, users can drag the previously defined boundary points to change the shape of shadow. However, its modifiable shapes are limited by the initial boundary points which also disallows users to add new segments of shadow. The interaction can be cumbersome for complex shape editing as users have to move multiple boundary points. To alleviate this burden as well as to increase the freedom of editing, we introduce an interaction to modify the shape of shadow by simply adding and subtracting user drawn shapes.

At start-up, our system detects, removes shadow and loads the original shadow according to the original shadow shape. Users then draw the boundary of new shape for addition or subtraction. The system automatically connects the two ends of a boundary curve using a straight line when the boundary is not closed. The drawing actions have two types:

Shadow Addition Shadow addition is equivalent to adding or sometimes merging a shape to the original shadow shape. This is done by applying a logical *or* image operation as follows:

$$N_n = N \vee N_a \quad (6)$$

where N refers to the shadow mask in Algorithm 1, N_a is the additional shape drawn by user, N_n is the updated shadow mask.

Shadow Subtraction Similar to Eq. 6, shadow subtraction is done by applying two logical image operations as follows:

$$N_n = N \wedge \neg N_a \quad (7)$$

As our shadow matte generation in Algorithm 3 already handles arbitrary shadow masks, users can draw any complicated shapes rather than changing a single existing shape in previous work [11].

4.2 Configuration Section for Other Controls

The other non-drawing controls are placed in the configuration section. The "mode" toggle button switches drawing mode between addition and subtraction of shape. The parameters for tuning shadow editing model are controlled by scrollbars.

5 Evaluation

In this section, we demonstrate our shadow editing system with examples of different types of shadow and show the quantitative evaluation result of our shadow reconstruction based on a state-of-the-art shadow removal dataset [1]. Our executables and data are made available to the community.

5.1 Demonstration of Shadow Editing

The shadow editing steps for various modification and shadow scenes are show in Fig. 3. Please also see our supplementary material for the video demonstration of these examples. Our algorithm is implemented in MATLAB script (non-MEX) and it gives interactive performance on a 2.4G hz machine. Fig. 3 also provides a qualitative evaluation of the shadow removal and editing algorithms. The real shadows are modified seamlessly in the images of different scenes. This qualitative demonstration reflects the compatibility of our shadow editing tool. The shadow editing process only changes the illumination and the background features, e.g. texture and reflectance, are preserved.

5.2 Quantitative Evaluation of Shadow Reconstruction

Although our system can produce new shadows perceptually similar to the original shadow, there may still be minor errors introduced in shadow composition phase. To evaluate any error and judge the quantitative performance of our system, we reconstruct the original shadow image using our shadow editing model with the extract shadow information and compute the Root Mean Square Error (RMSE) between the generated image and the original image. The evaluation is based on a state-of-the-art shadow removal dataset [1] containing 214 categorised shadow images in variable scenes. As shown in Table 1, our system produces negligible errors of shadow reconstruction and has similar accuracies for all 4 categories. The accuracy of reconstruction is often not crucial for shadow editing but can be important for photo forging.




Fig. 3. Demonstration of shadow editing in different scenes: The first two rows are examples of easy scenes. The remaining rows are examples for scenes with strong texture background, broken shadow, coloured shadow, and soft shadow respectively.

	Texture				Soft				Broken				Colour				Other
	1	2	3	M	1	2	3	M	1	2	3	M	1	2	3	M	
E	1(0)	1(1)	3(1)	2(1)	1(0)	1(1)	3(1)	2(1)	1(0)	2(1)	3(1)	2(1)	1(1)	1(0)	2(1)	2(1)	2(1)
E*	0(0)	0(0)	0(0)	0(0)	0(0)	0(0)	0(0)	0(0)	0(0)	0(0)	0(0)	0(0)	0(0)	0(0)	1(1)	0(0)	0(0)

Table 1. Error of shadow reconstruction according to 4 attributes. The intensities of image are normalised and the errors shown in this table are in percent (multiplied by 100). The "E" and "E*" indicate the error score where all pixels in the image are used, and just shadow area pixels respectively. For each score of each attribute, the images with other predominant attributes (degree = 3) are not used. Hence, test cases have a strong single bias towards one of the attributes. "Other" refers to a set of shadow cases showing no markedly predominant attributes (degree = 1). "M" refers to the average score for each category. Standard deviations are shown in brackets.

6 Conclusion

We have presented a shadow editing system with a user-friendly interface. It enables users to freely and quickly modify various properties of existing shadows in images which include shape, darkness, softness, and colour. We have demonstrated and quantitatively evaluated our shadow editing system. Future work includes: 1) the evaluation of the interface with actual users; 2) making use of estimated light sources and geometry from images to provide intelligent suggestions for users to create plausible shadows.

Acknowledgements This work was funded by the China Scholarship Council and the University of Bath.

References

1. Gong, H., Cosker, D.: Interactive shadow removal and ground truth for variable scene categories. In: British Machine Vision Conference (BMVC). (2014)
2. Finlayson, G.D., Drew, M.S., Lu, C.: Entropy minimization for shadow removal. *International Journal of Computer Vision* **85** (2009) 35–57
3. Yang, Q., Tan, K.H., Ahuja, N.: Shadow removal using bilateral filtering. *Image Processing, IEEE Transactions on* **21** (2012) 4361–4368
4. Guo, R., Dai, Q., Hoiem, D.: Paired regions for shadow detection and removal. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **PP** (2012) 1–1
5. Wu, T.P., Tang, C.K., Brown, M.S., Shum, H.Y.: Natural shadow matting. *ACM Transactions on Graphics (TOG)* **26** (2007) 8
6. Liu, F., Gleicher, M.: Texture-consistent shadow removal. In: Proc. of European Conf. on Computer Vision (ECCV). (2008) 437–450
7. Shor, Y., Lischinski, D.: The shadow meets the mask: Pyramid-based shadow removal. *Comput. Graph. Forum* **27** (2008) 577–586
8. Su, Y.F., Chen, H.H.: A three-stage approach to shadow field estimation from partial boundary information. *IEEE Trans. on Image Processing* **19** (2010) 2749–2760
9. Arbel, E., Hel-Or, H.: Shadow removal using intensity surfaces and texture anchor points. *IEEE Trans. Pattern Analysis and Machine Intelligence* **33** (2011) 1202–1216
10. Gong, H., Cosker, D., Li, C., Brown, M.: User-aided single image shadow removal. In: Proc. IEEE Int. Conf. Multimedia and Expo (ICME). Volume 1. (2013) 1
11. Mohan, A., Tumblin, J., Choudhury, P.: Editing soft shadows in a digital photograph. *IEEE Computer Graphics and Applications* **27** (2007) 23–31
12. Fritsch, F.N., Carlson, R.E.: Monotone piecewise cubic interpolation. *SIAM Journal on Numerical Analysis* **17** (1980) 238–246
13. Bertalmio, M., Sapiro, G., Caselles, V., Ballester, V.: Image inpainting. In: Proceedings of the 27th annual conference on Computer graphics and interactive techniques. SIGGRAPH '00 (2000) 417–424
14. Schalkoff, R.J.: Digital image processing and computer vision. Volume 286. Wiley New York (1989)
15. Maurer Jr, C.R., Qi, R., Raghavan, V.: A linear time algorithm for computing exact euclidean distance transforms of binary images in arbitrary dimensions. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **25** (2003) 265–270