Computational Illumination

Frédo Durand

MIT - EECS

Some Slides from Ramesh Raskar (MIT Medialab)

High level idea

- **Control the illumination to**
  - Lighting as a post-process
  - Extract more information

Flash/no-flash

Flash Photography Enhancement via Intrinsic Relighting

Elmar Eisemann

MIT/Artis-INRIA

Frédo Durand

MIT

Introduction

Satisfactory photos in dark environments are challenging!

Introduction

Available light:

+ **original lighting**

- **noise/blurriness**
- **color**

Introduction

Flash:

+ **details**

+ **color**

- **flat/artificial**
- **flash shadows**
- **red eyes**

Introduction

Our approach:

Use no-flash image relight flash image

Introduction

Our approach:

Use no-flash image relight flash image

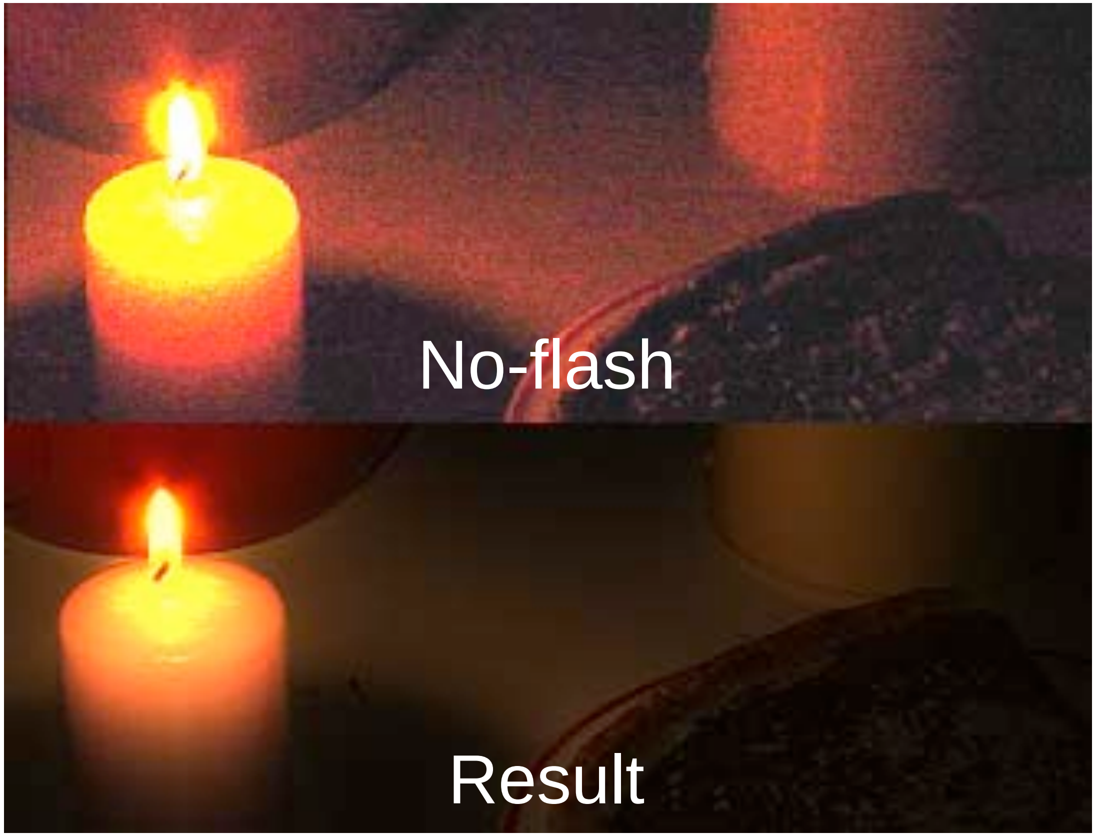

+ original lighting
+ details/sharpness
+ color

Introduction

One approach: Blend the two photos

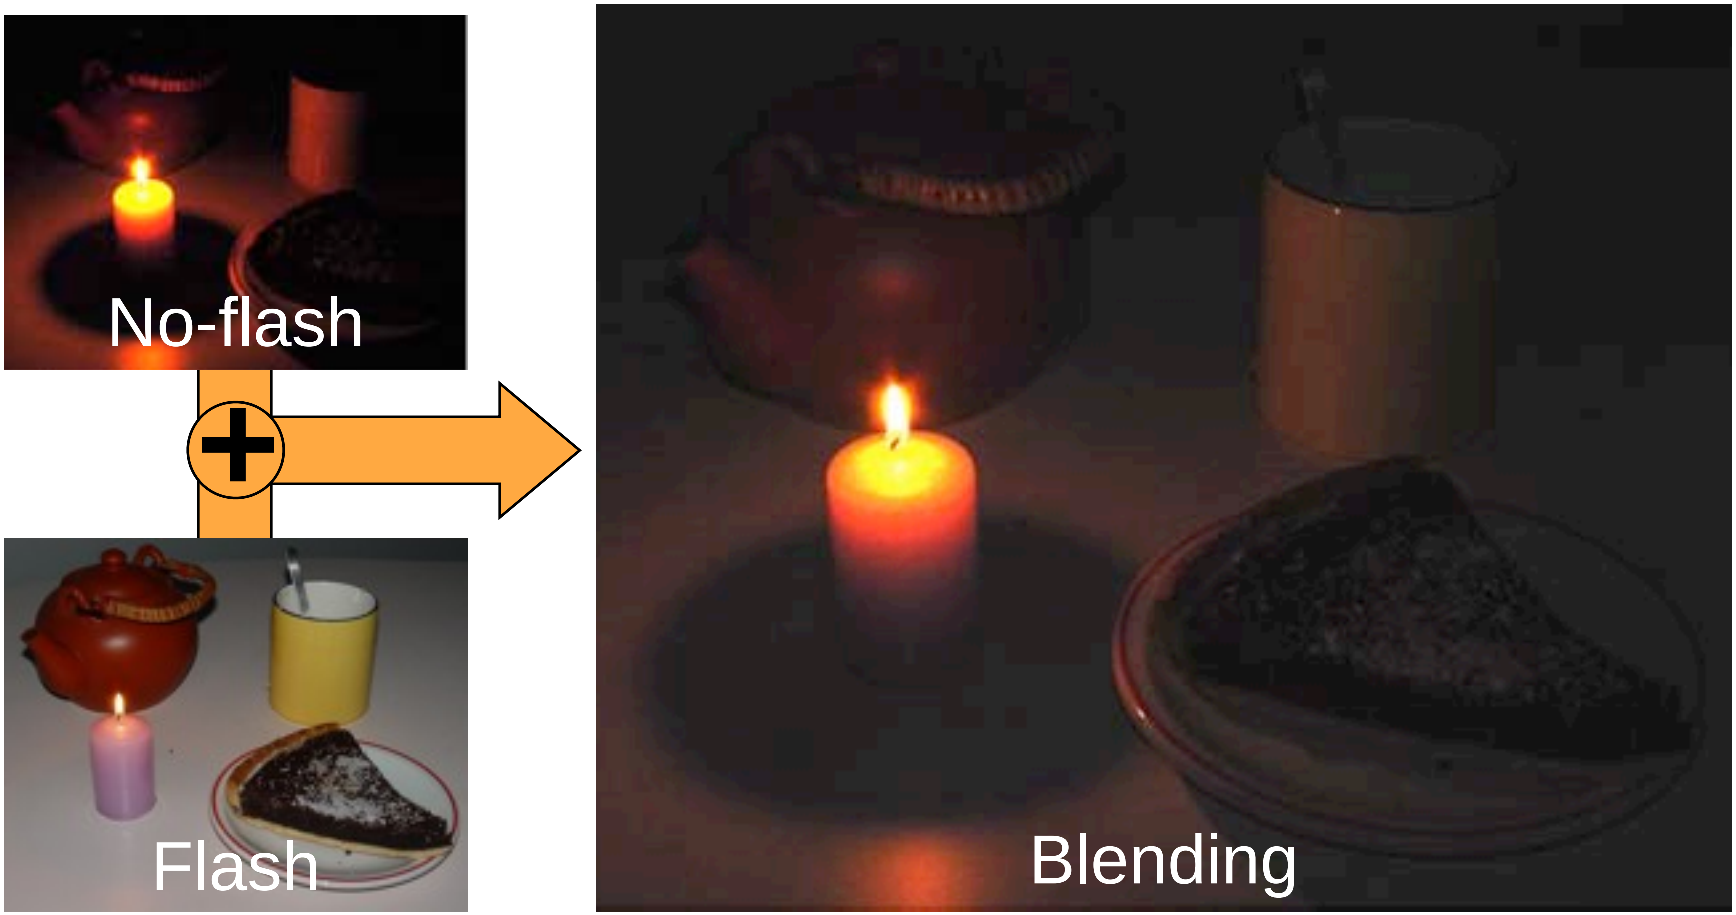

Introduction

One approach: Blend the two photos

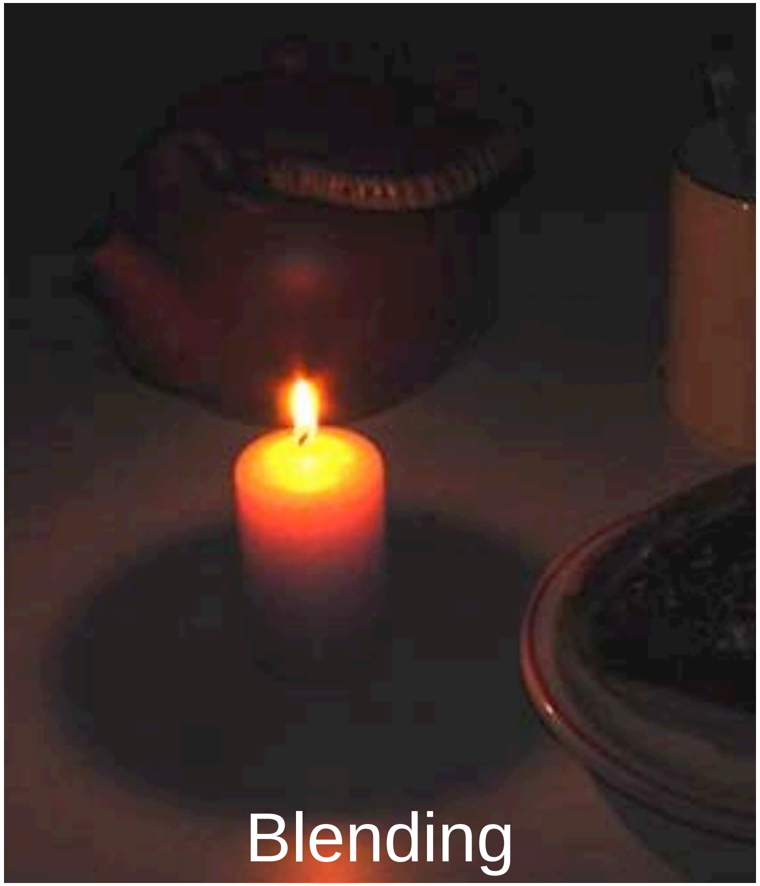

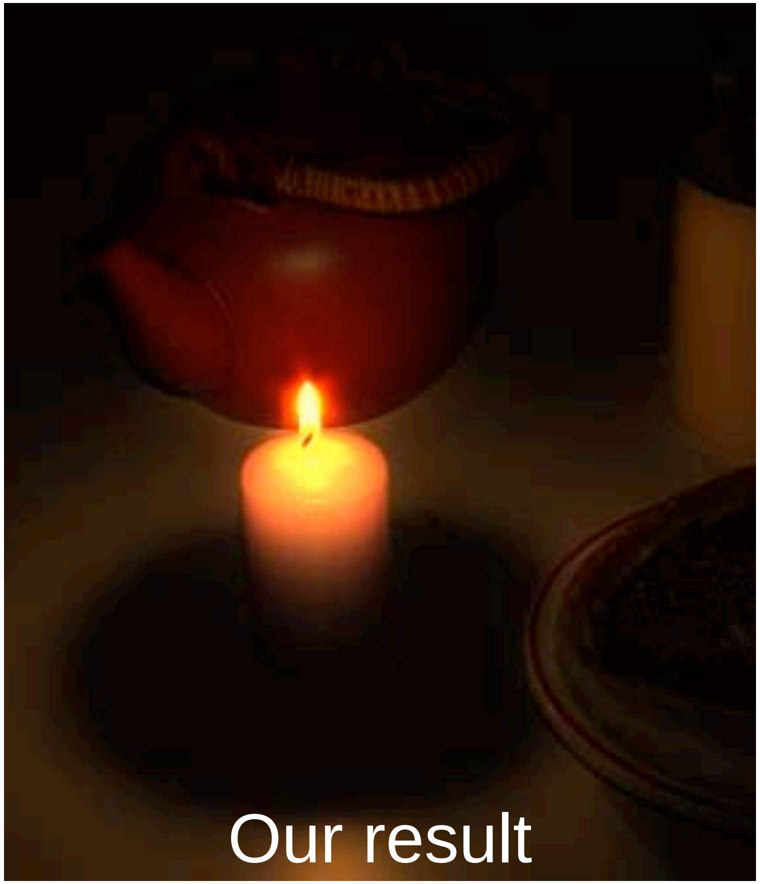

Introduction

One approach: Blend the two photos

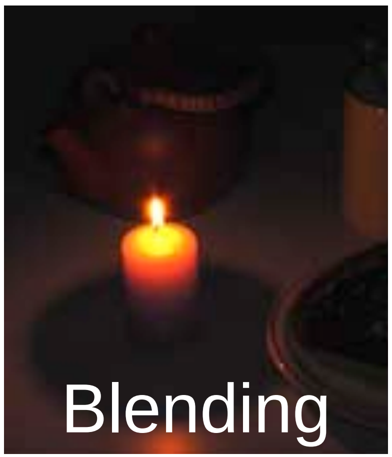

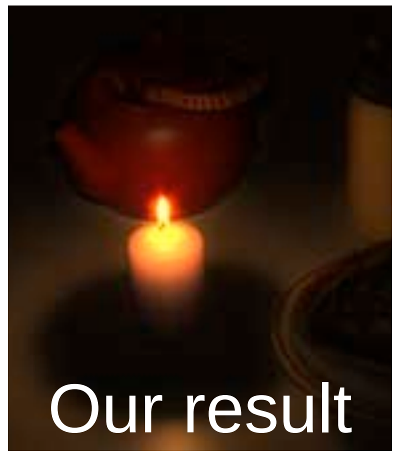

Introduction

One approach: Blend the two photos

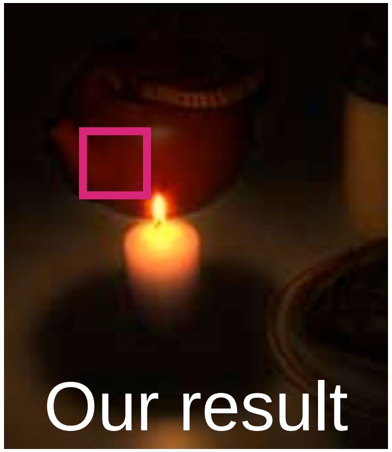

Introduction

One approach: Blend the two photos

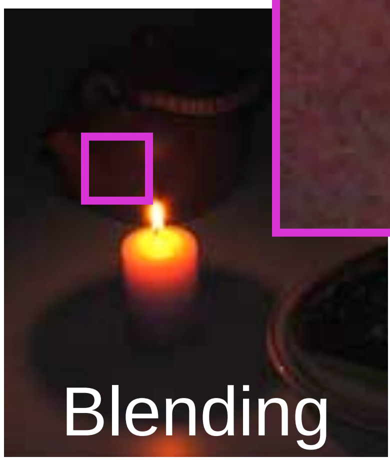

Our Solution:

more details and less noise

Overview

- **Related Work**
- **Our Approach**
- **Results**
- **Conclusion and Future Work**

Overview

- Related Work
- **Our Approach**
- **Results**
- **Conclusion and Future Work**

Related Work

Tone Mapping of High Dynamic Range Images

Decouple detail / large-scale information

- Tumblin et al. [1999]
- Durand et al. [2002]
- Choudhury et al. [2003]

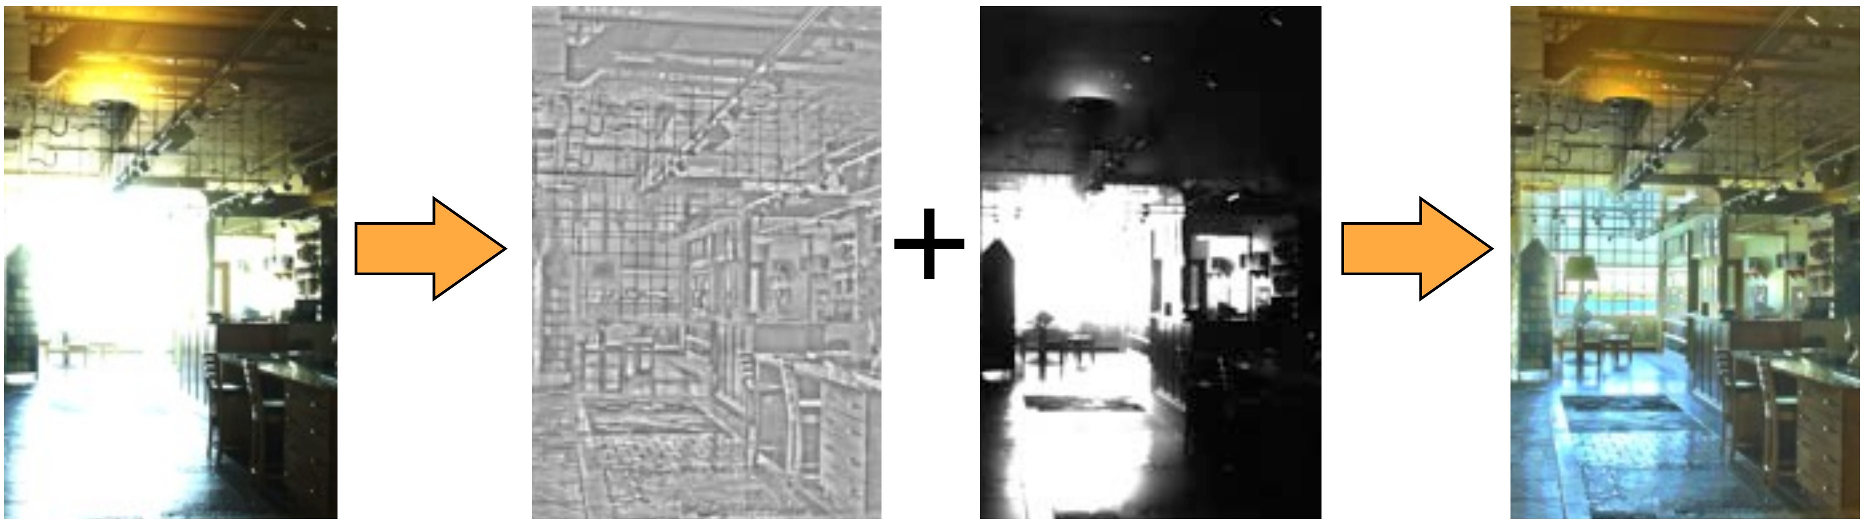

Input Detail Large-Scale Output

Related Work

Petschnigg et al.[2004]:

- many similarities
- part of this year‘s proceedings
- discussion at the end

Overview

- **Related Work**
- Our Approach
- **Results**
- **Conclusion and Future Work**

Our Approach - Main Idea

Our Approach - Main Idea

no-flash

flash

Our Approach - Main Idea

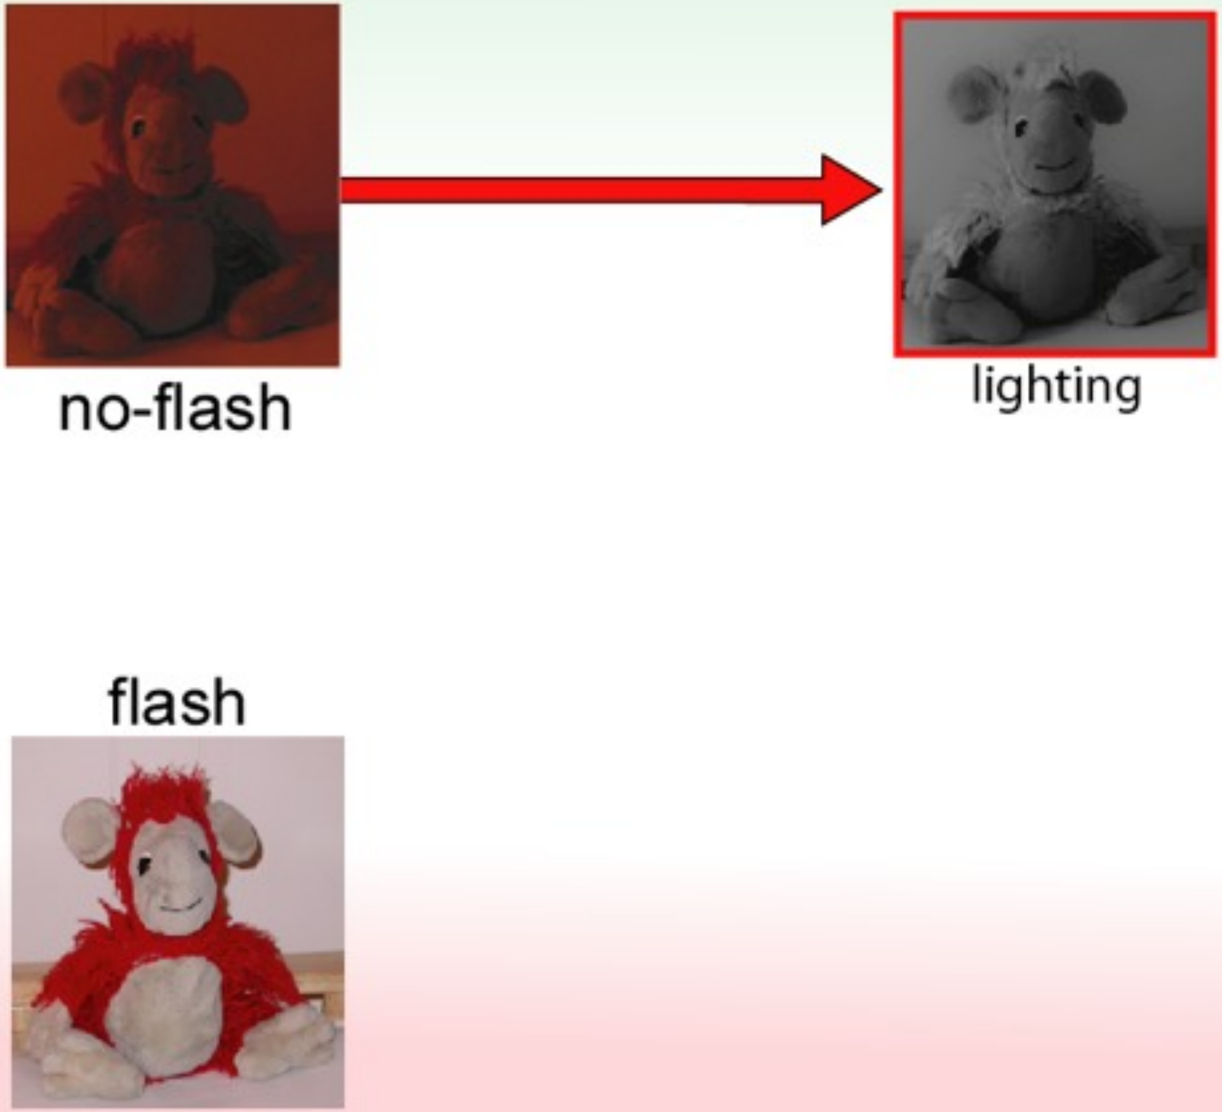

Our Approach - Main Idea

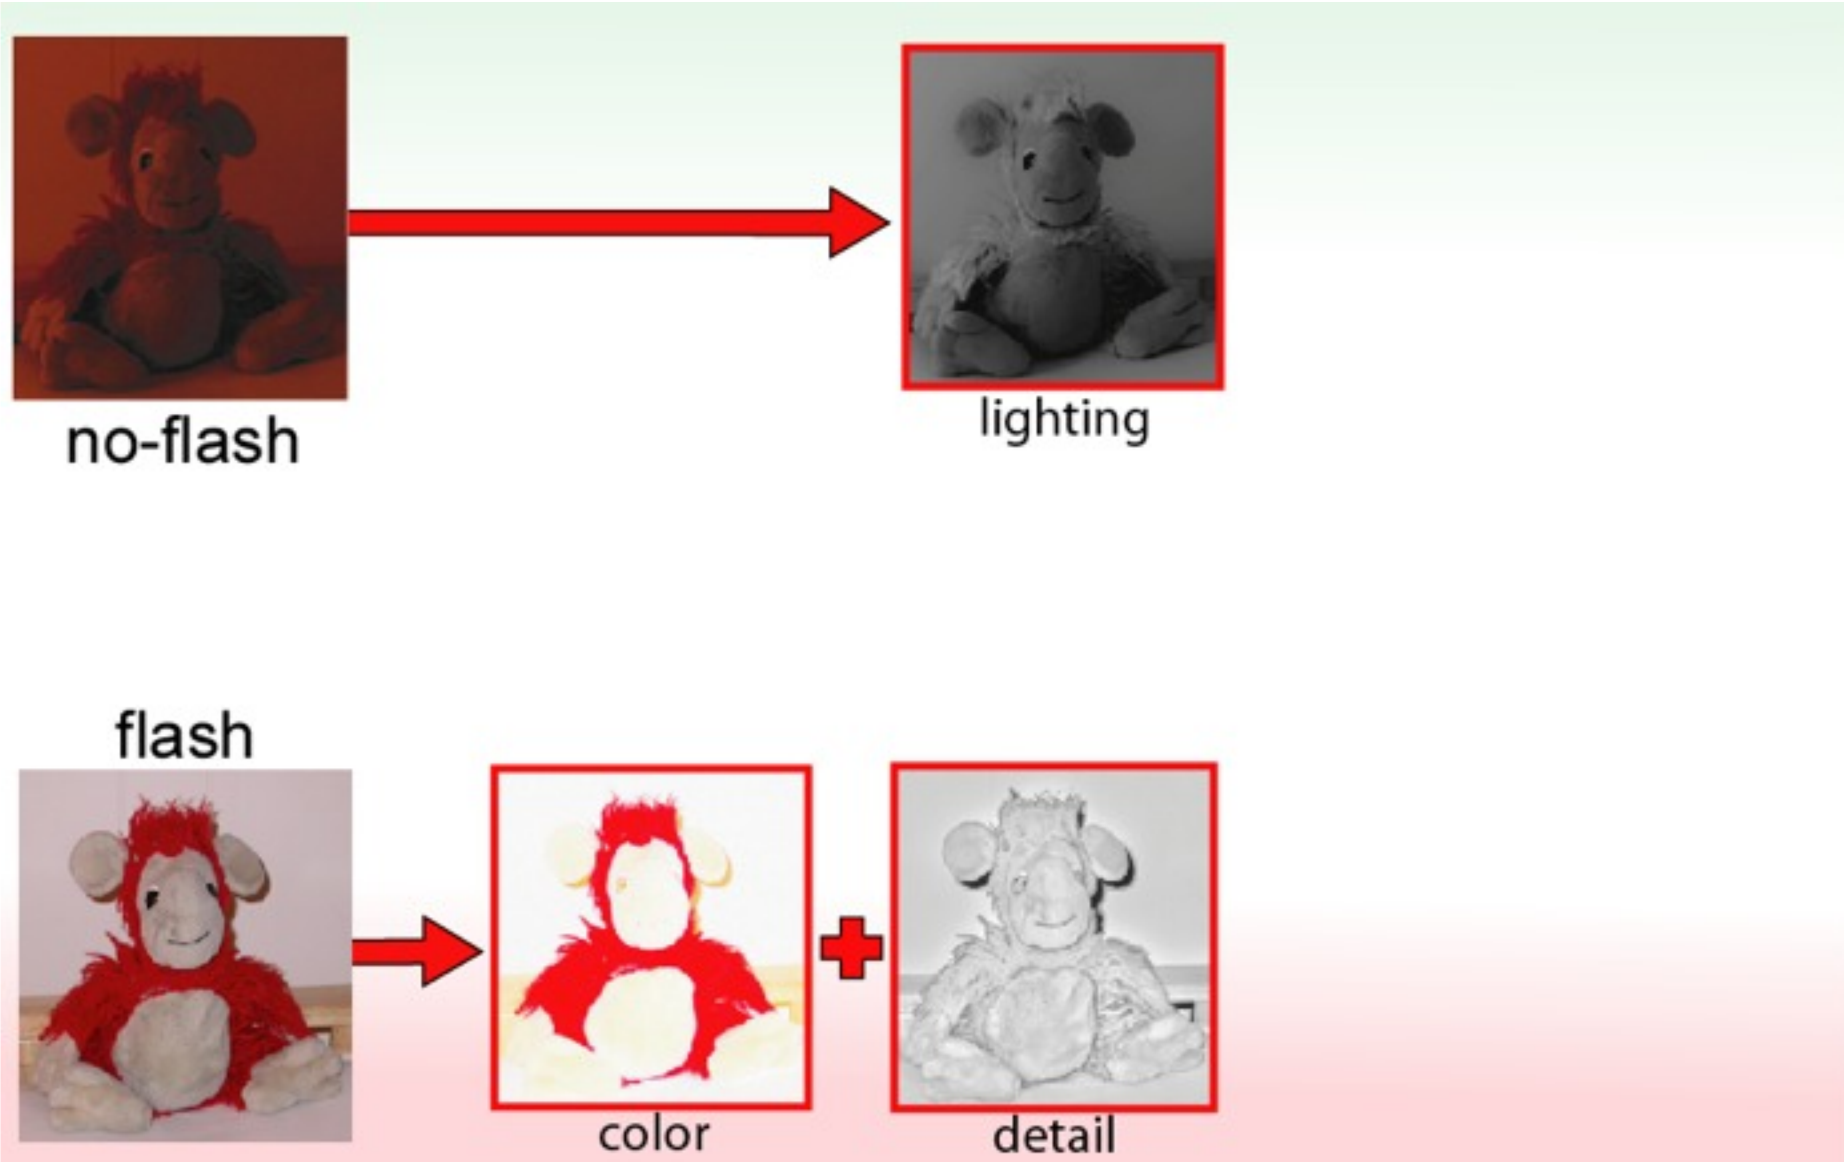

Our Approach - Main Idea

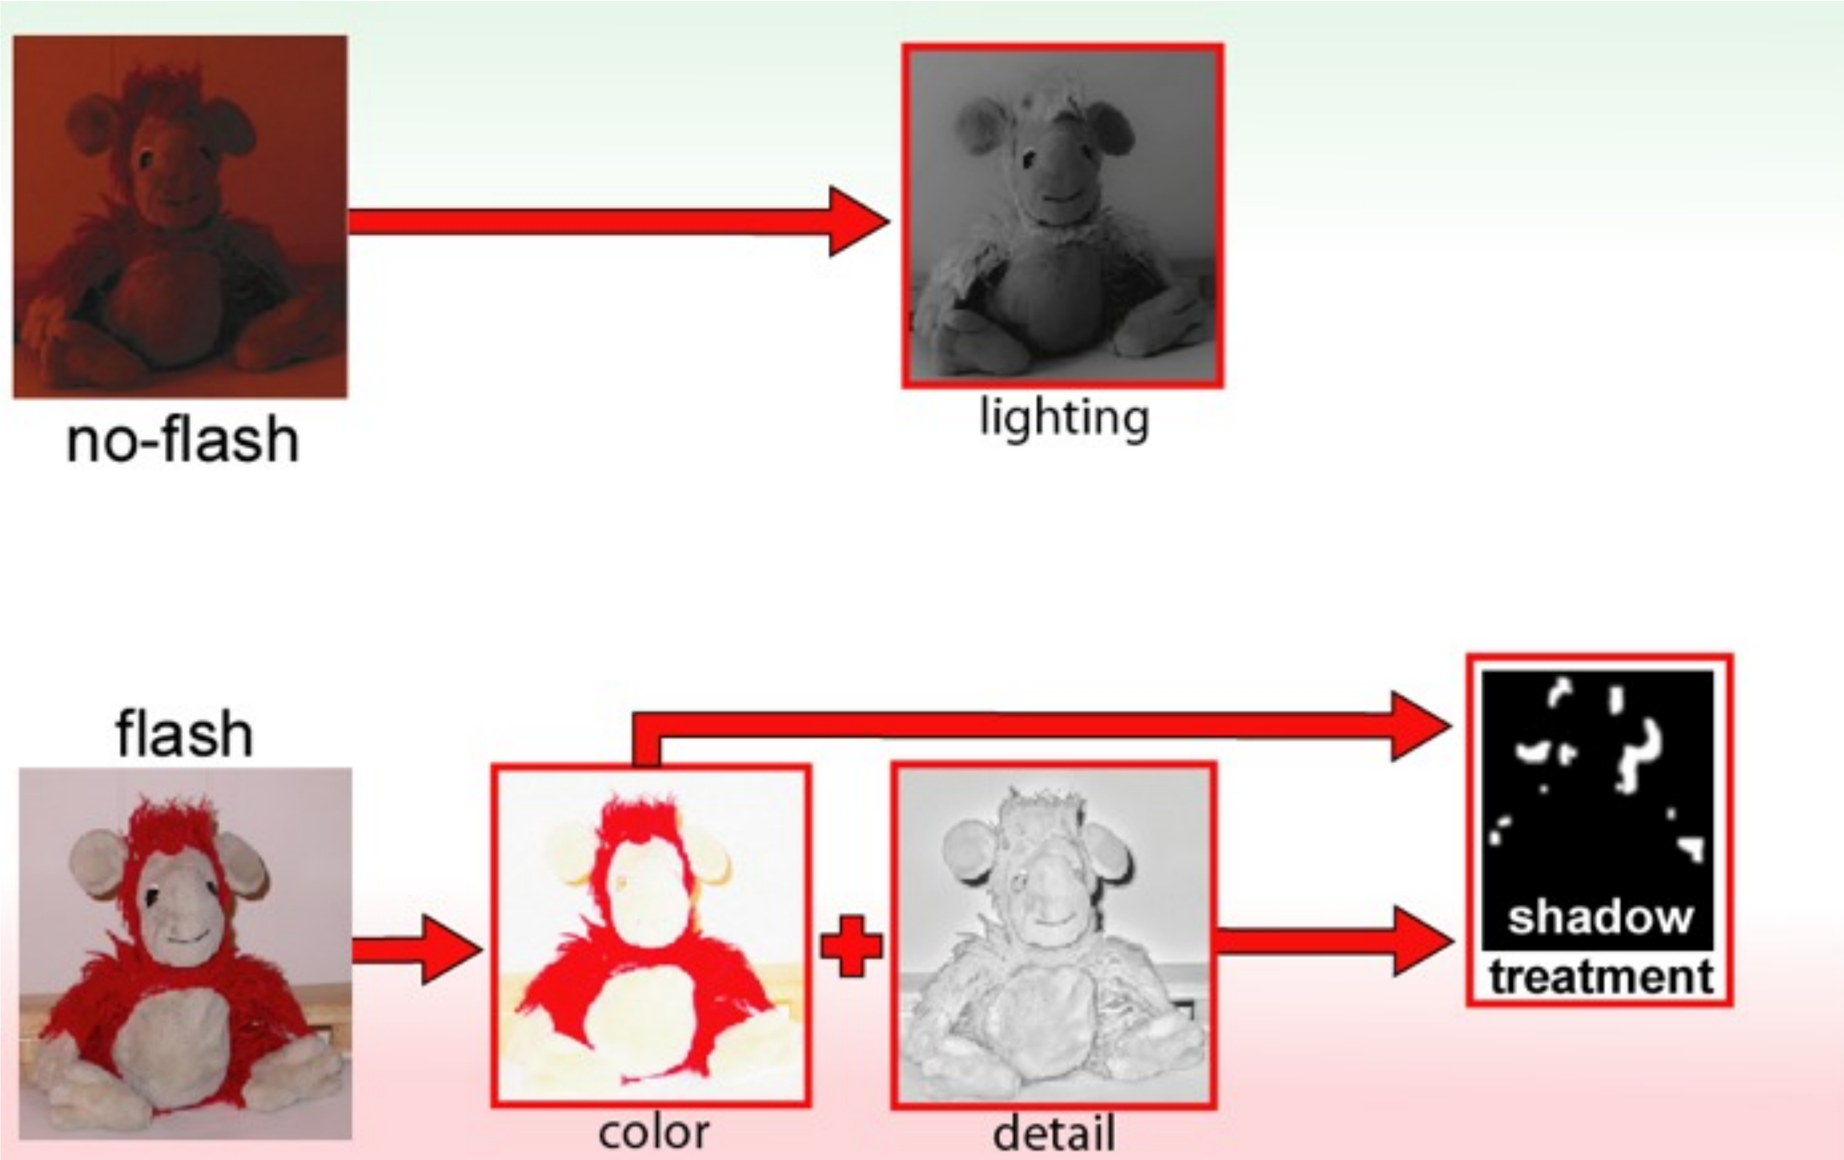

Our Approach - Main Idea

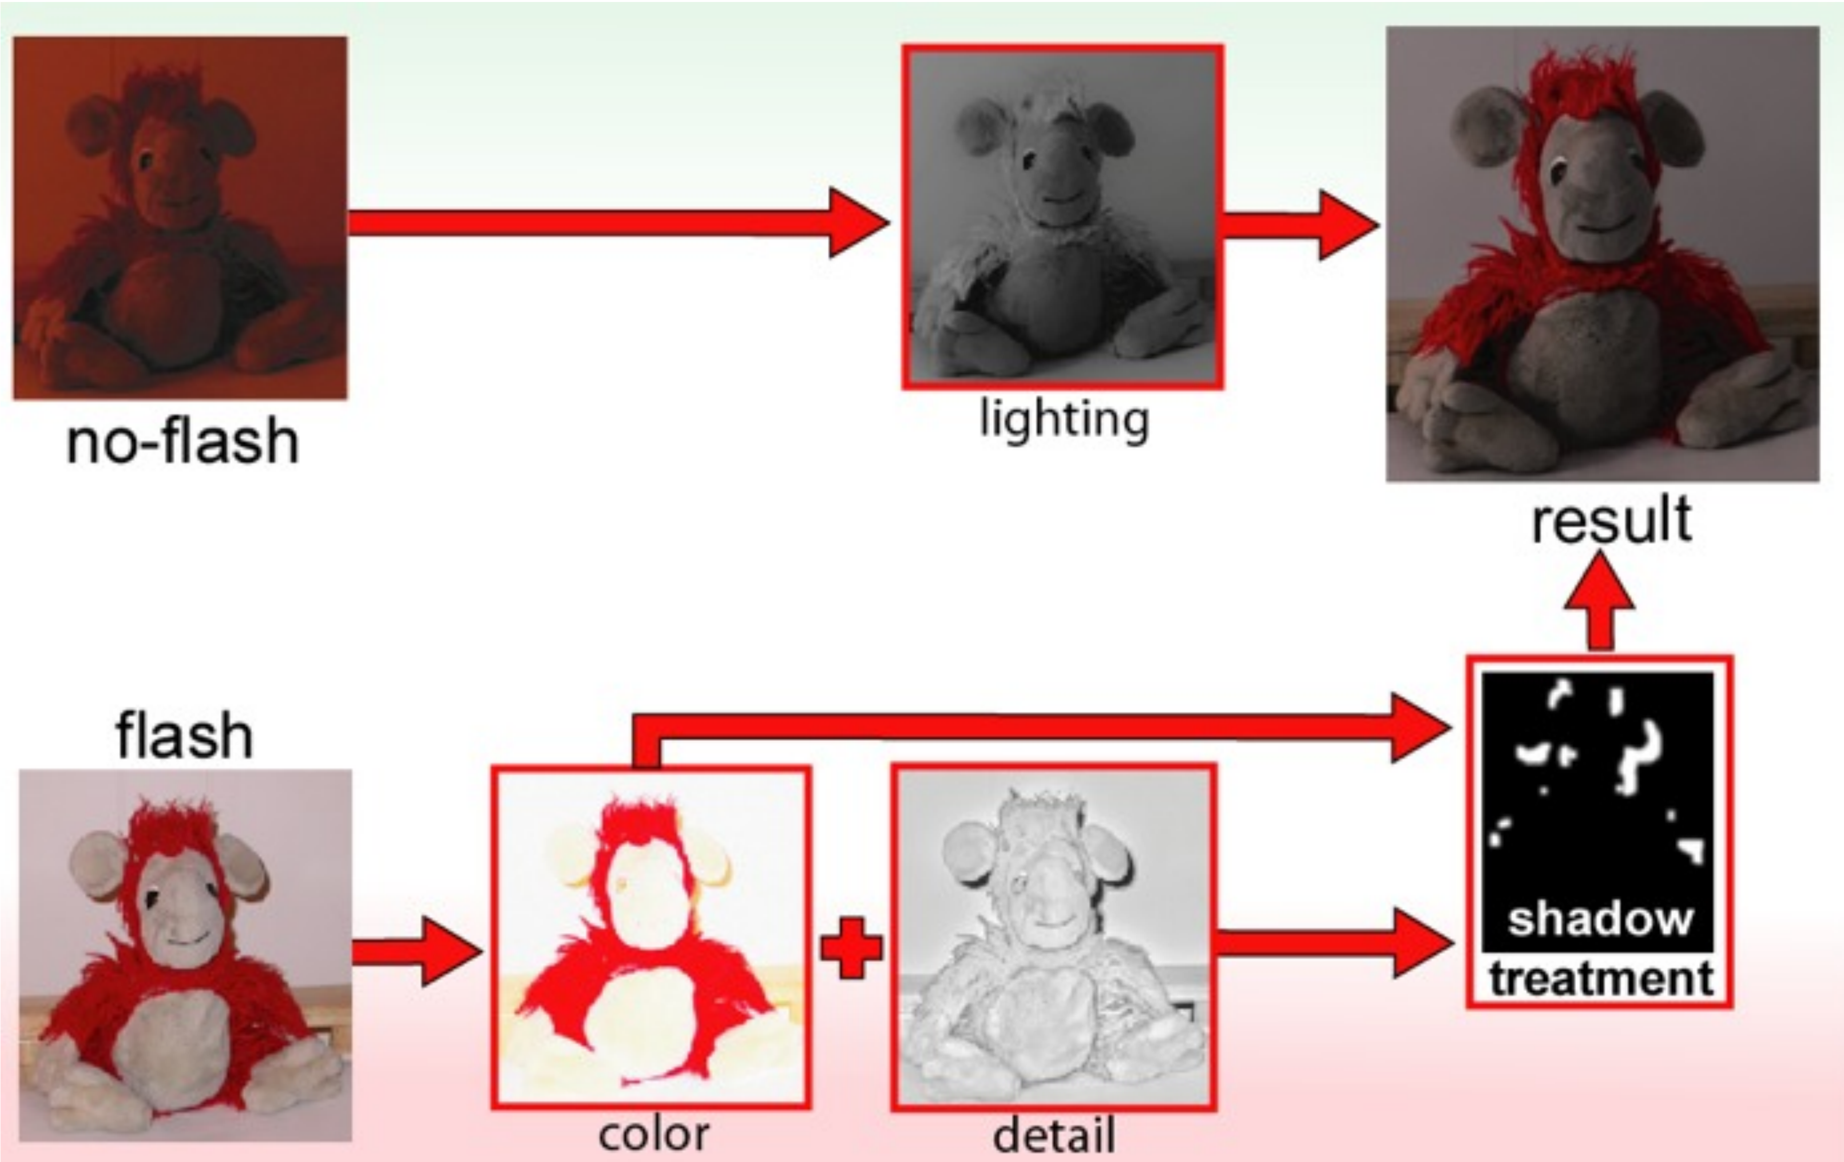

Our Approach

no-flash

flash

Our Approach

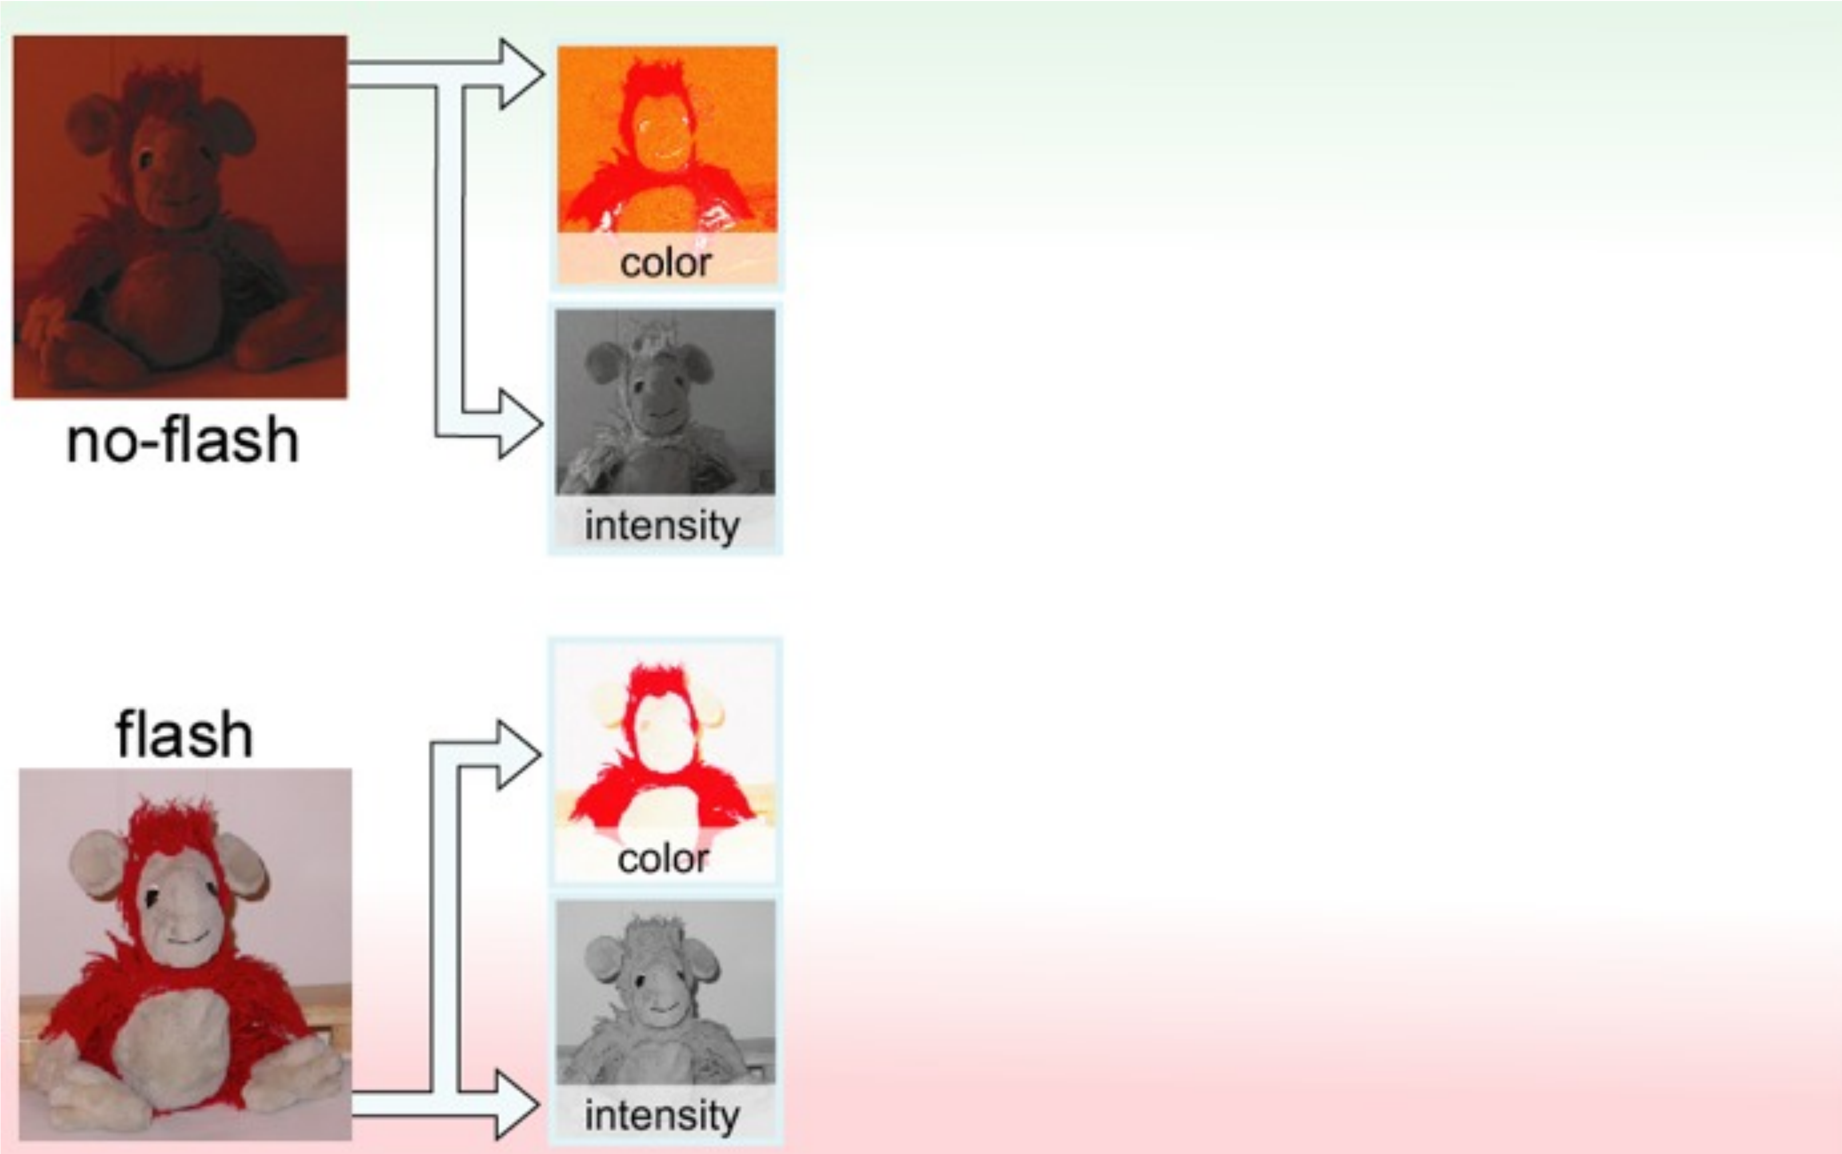

Our Approach

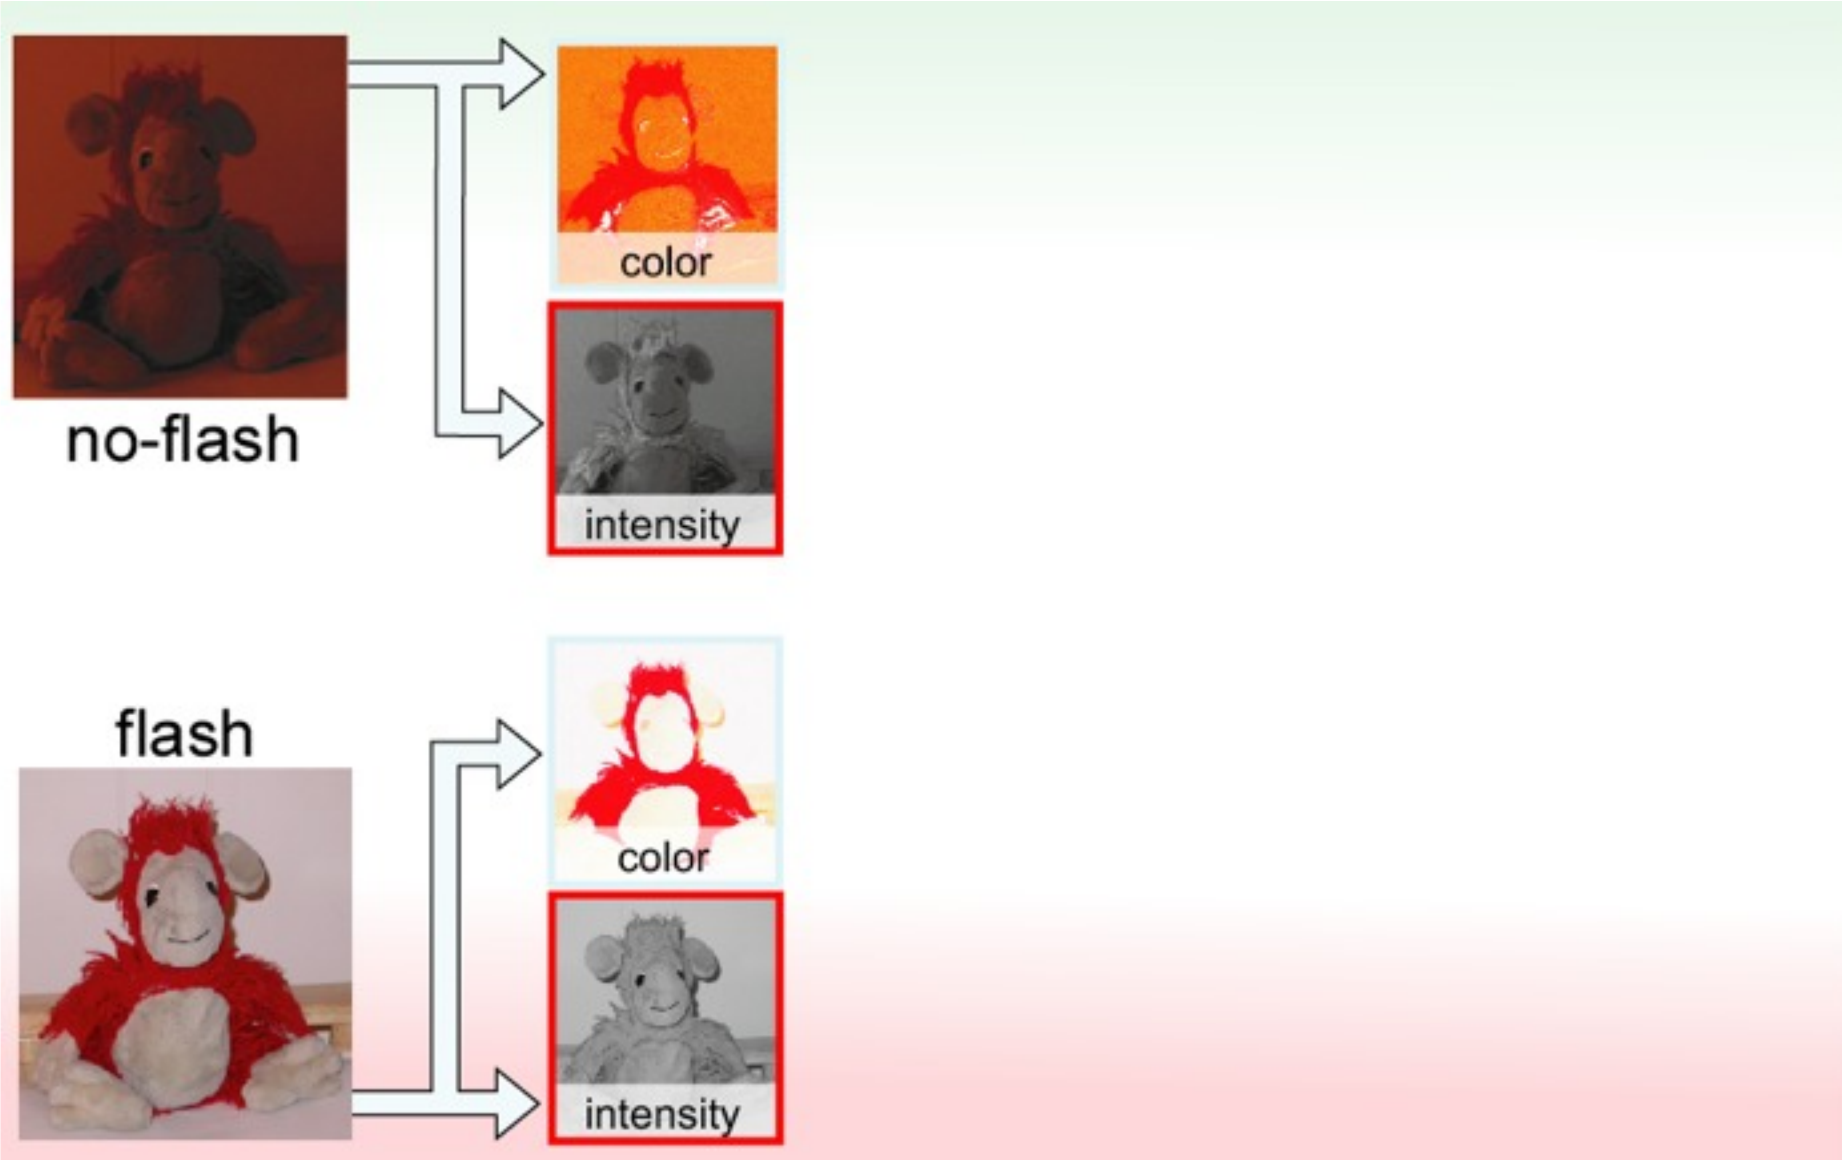

Our Approach

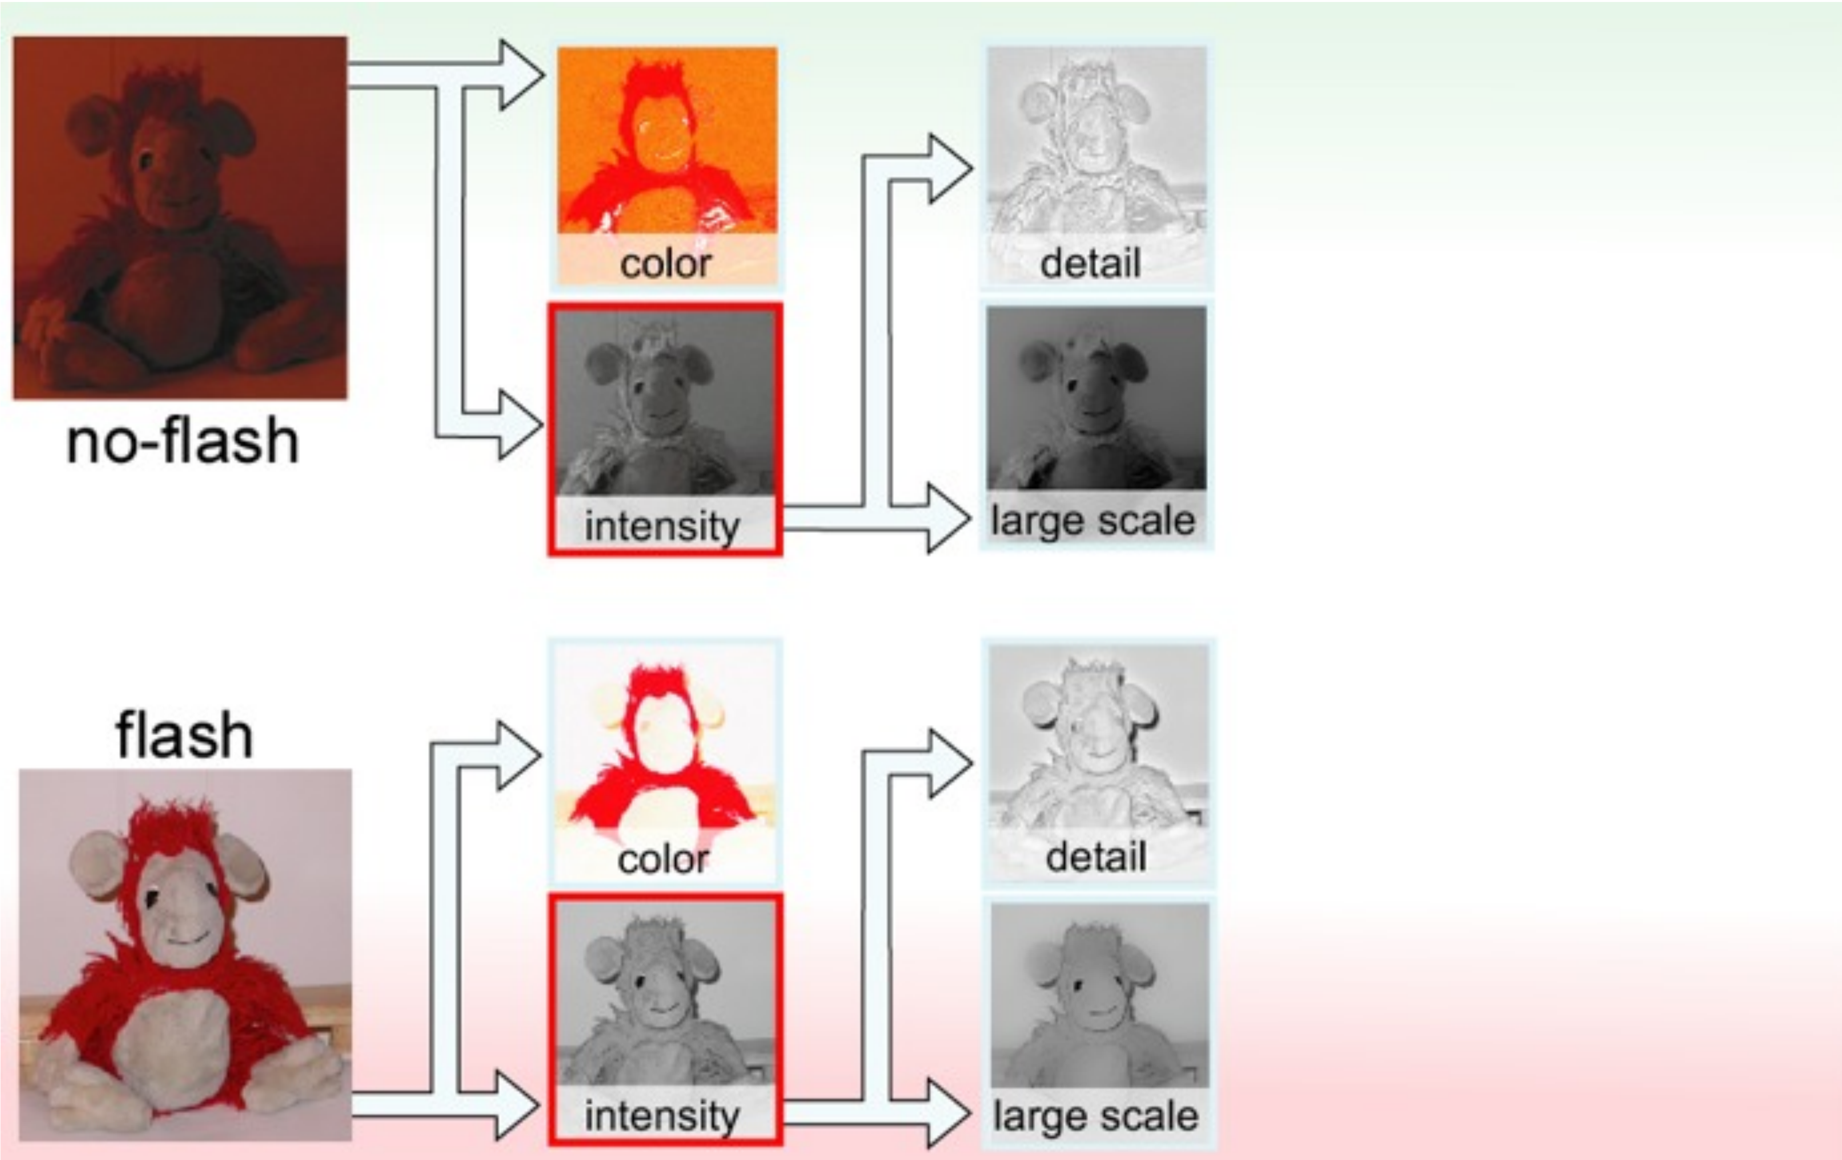

Our Approach

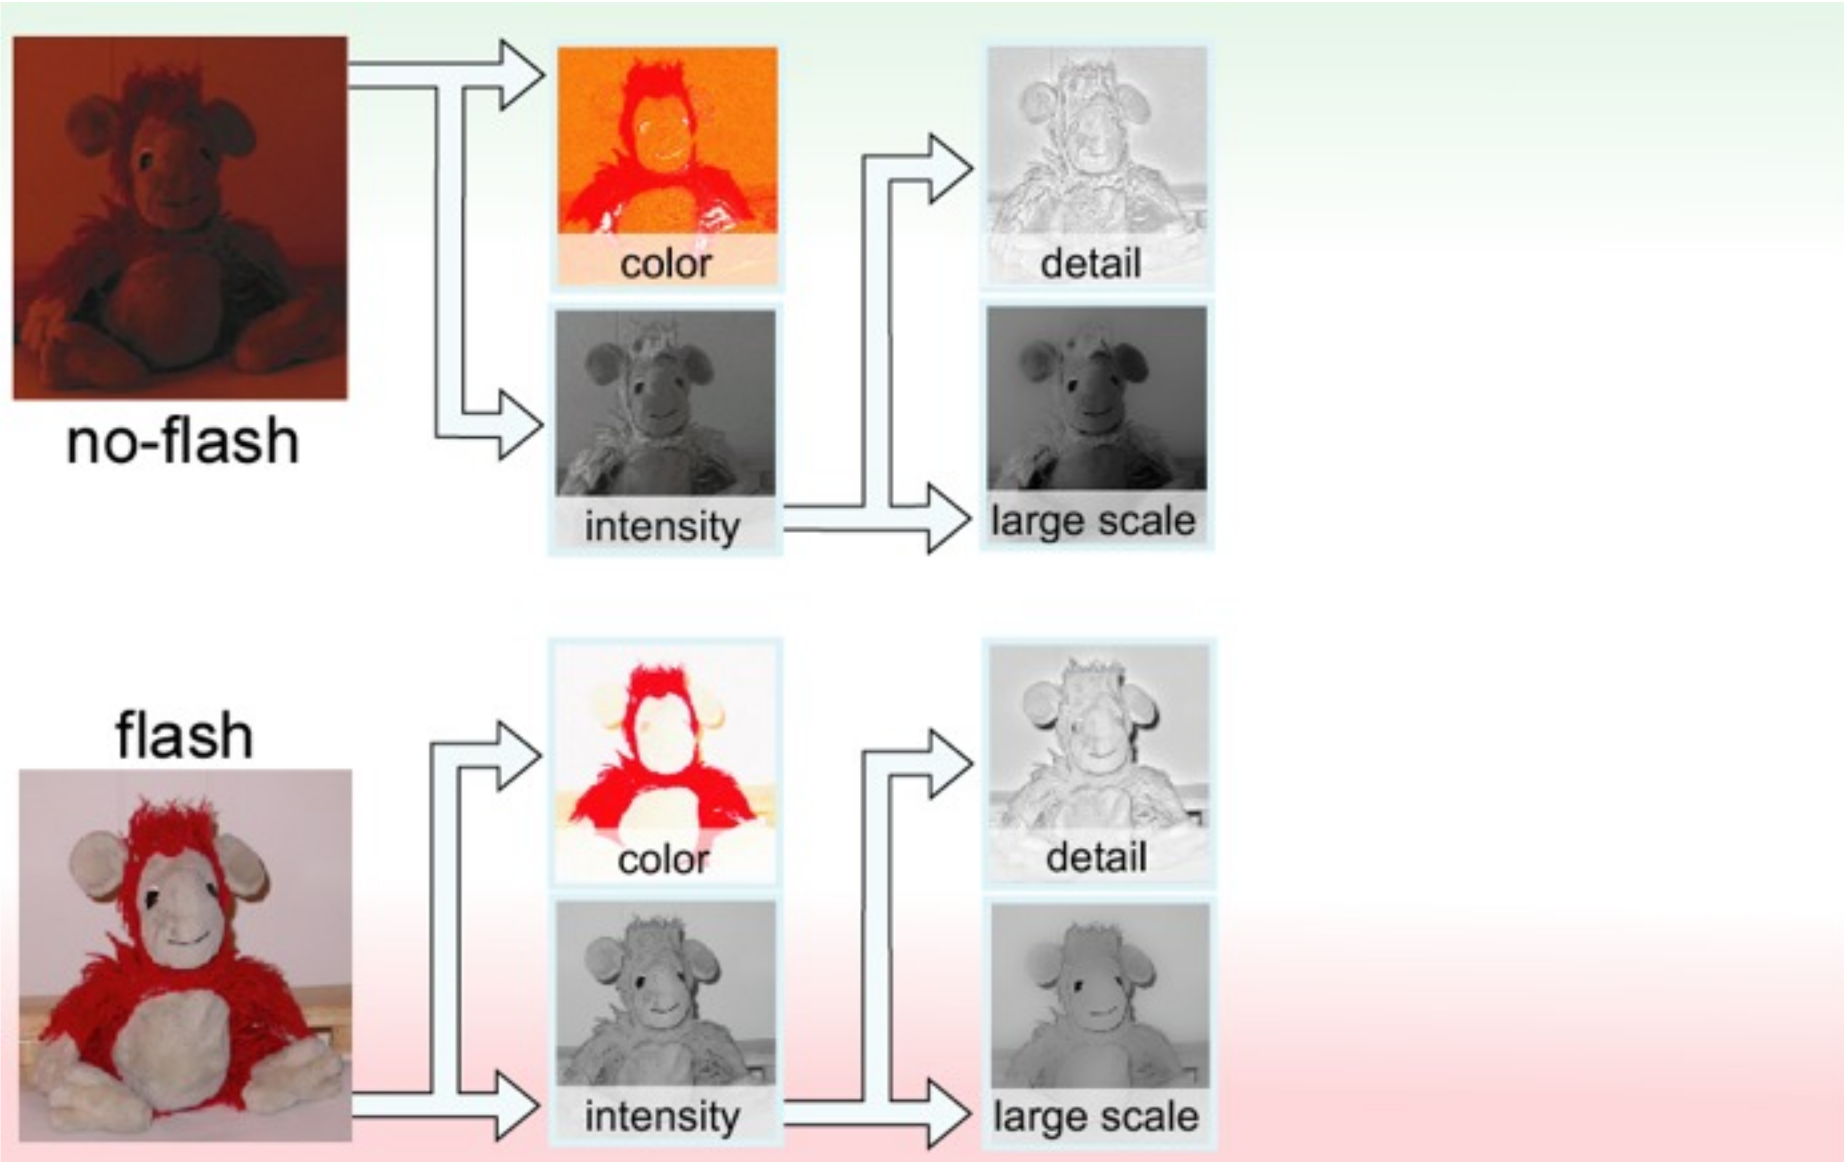

Our Approach

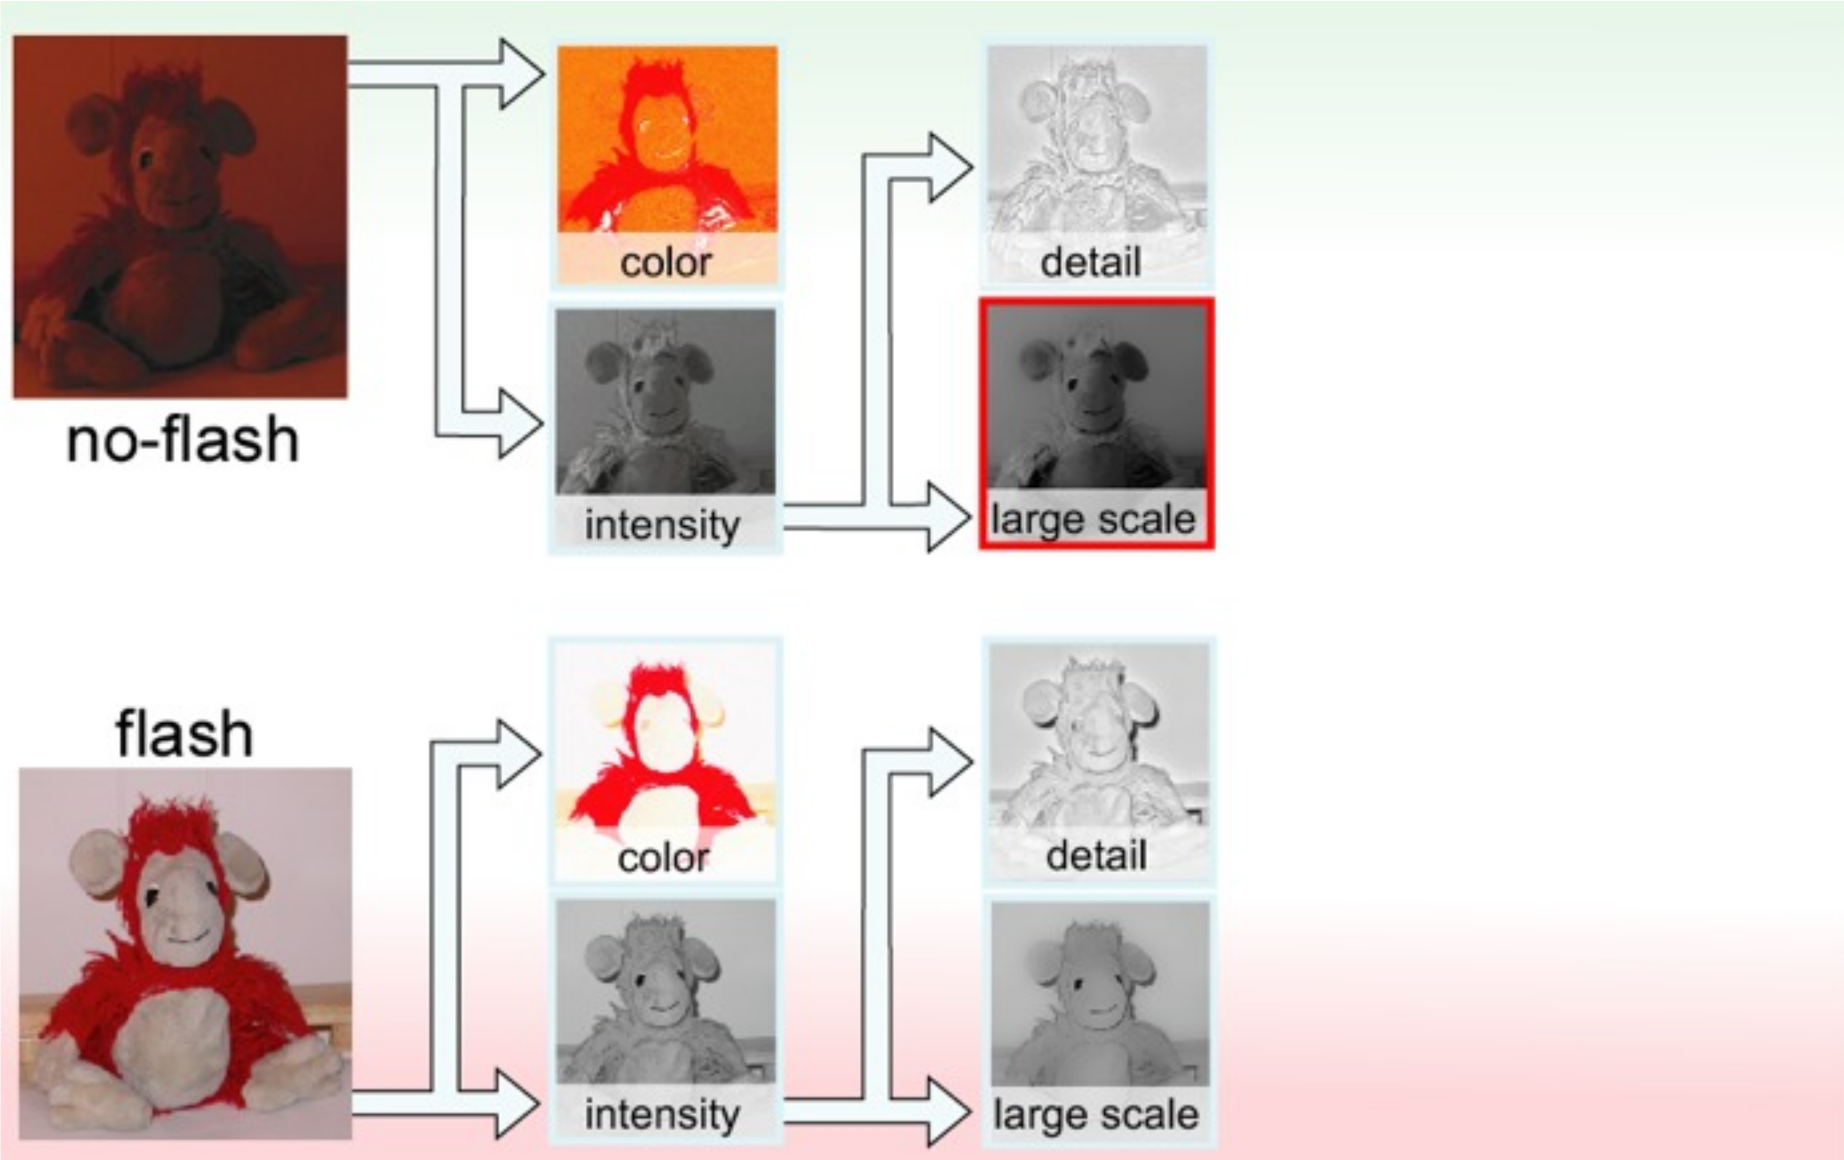

Our Approach

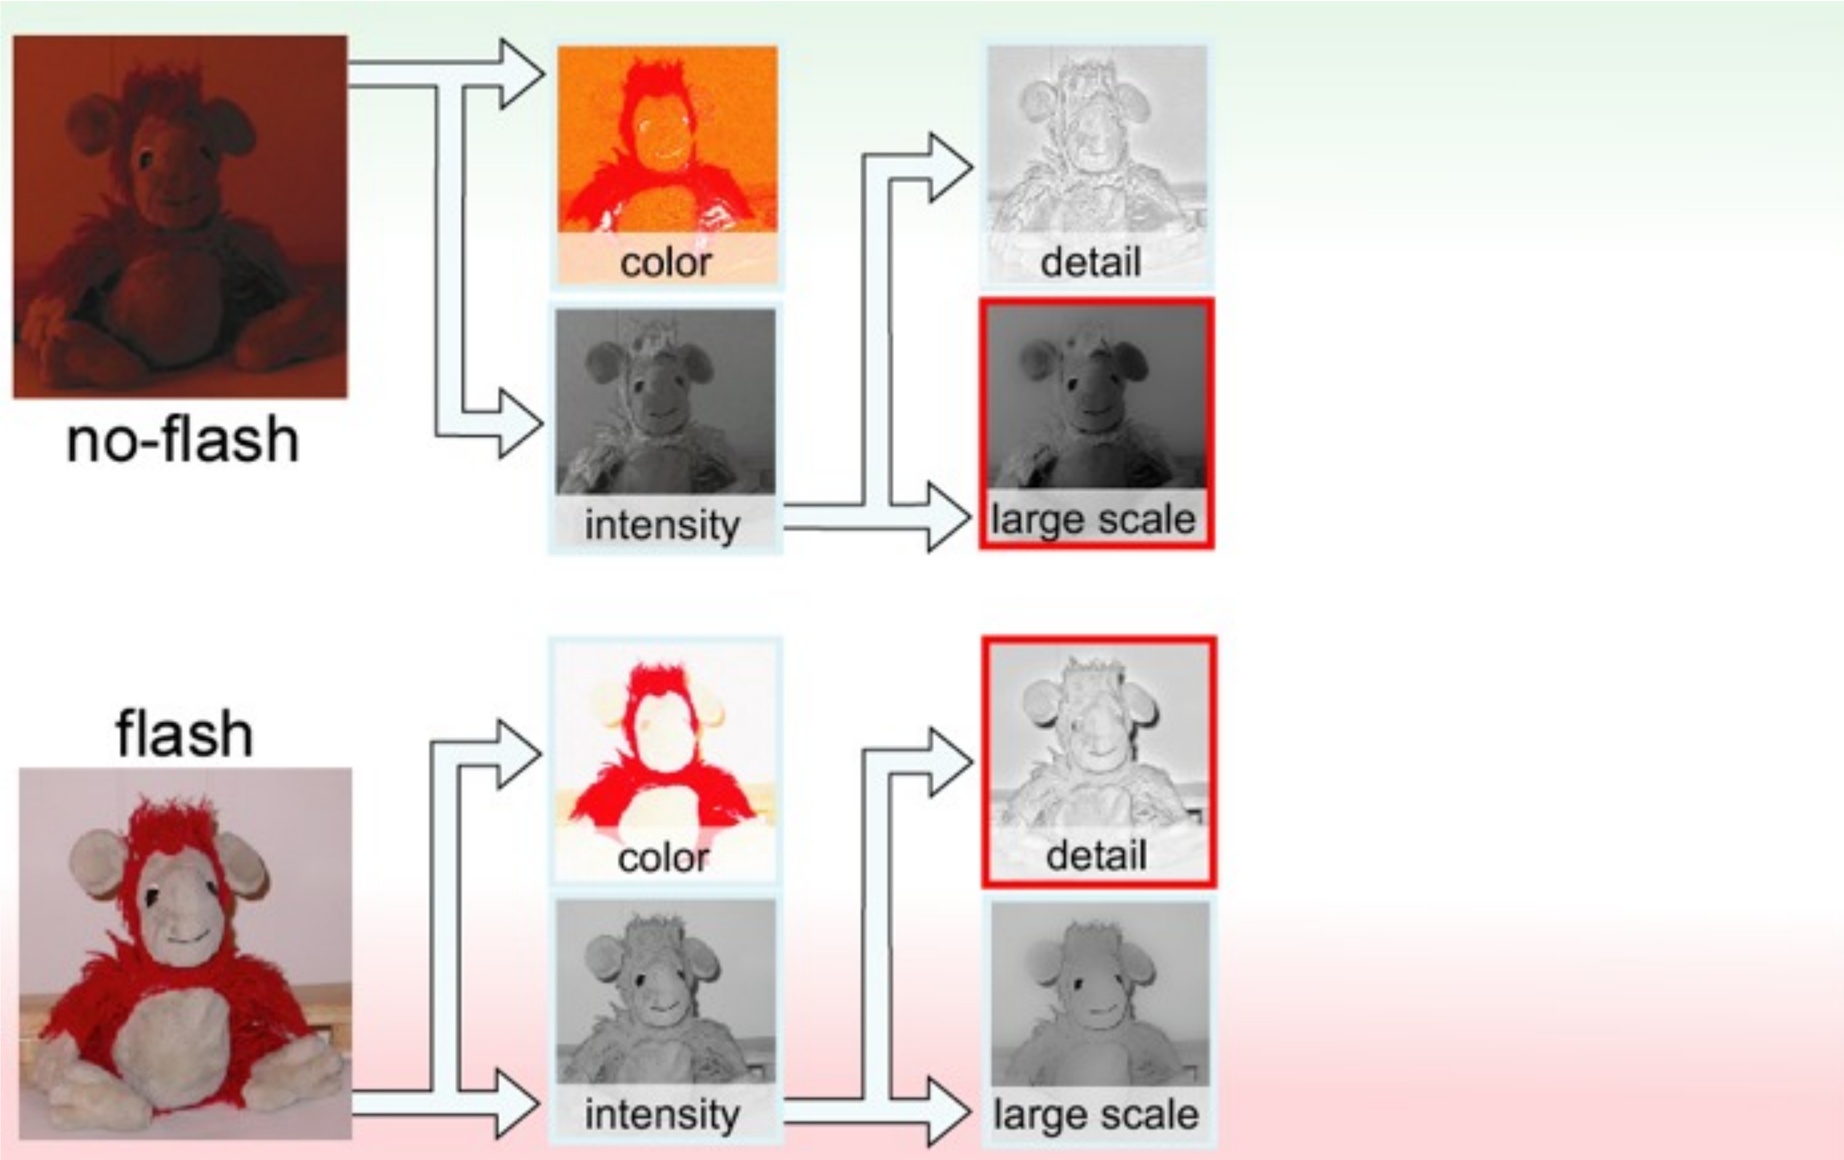

Our Approach

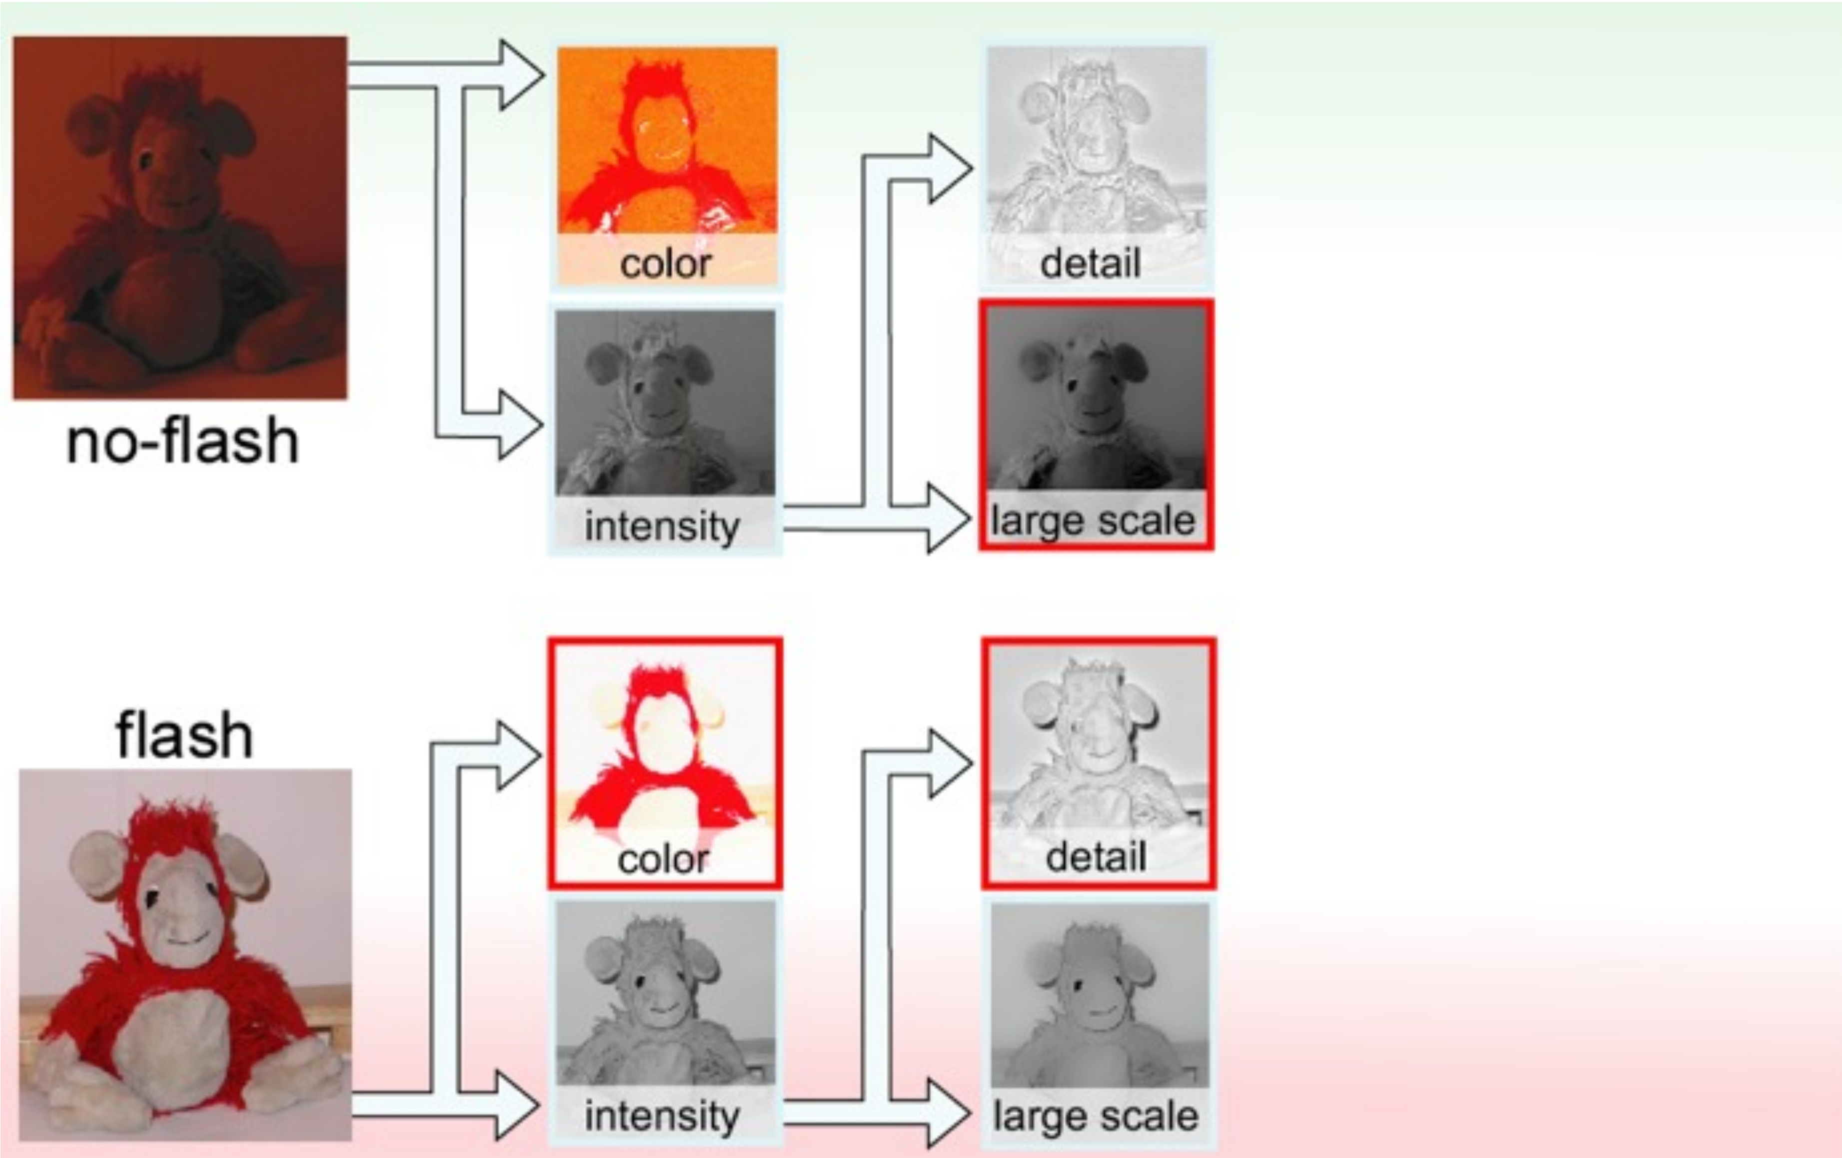

Our Approach

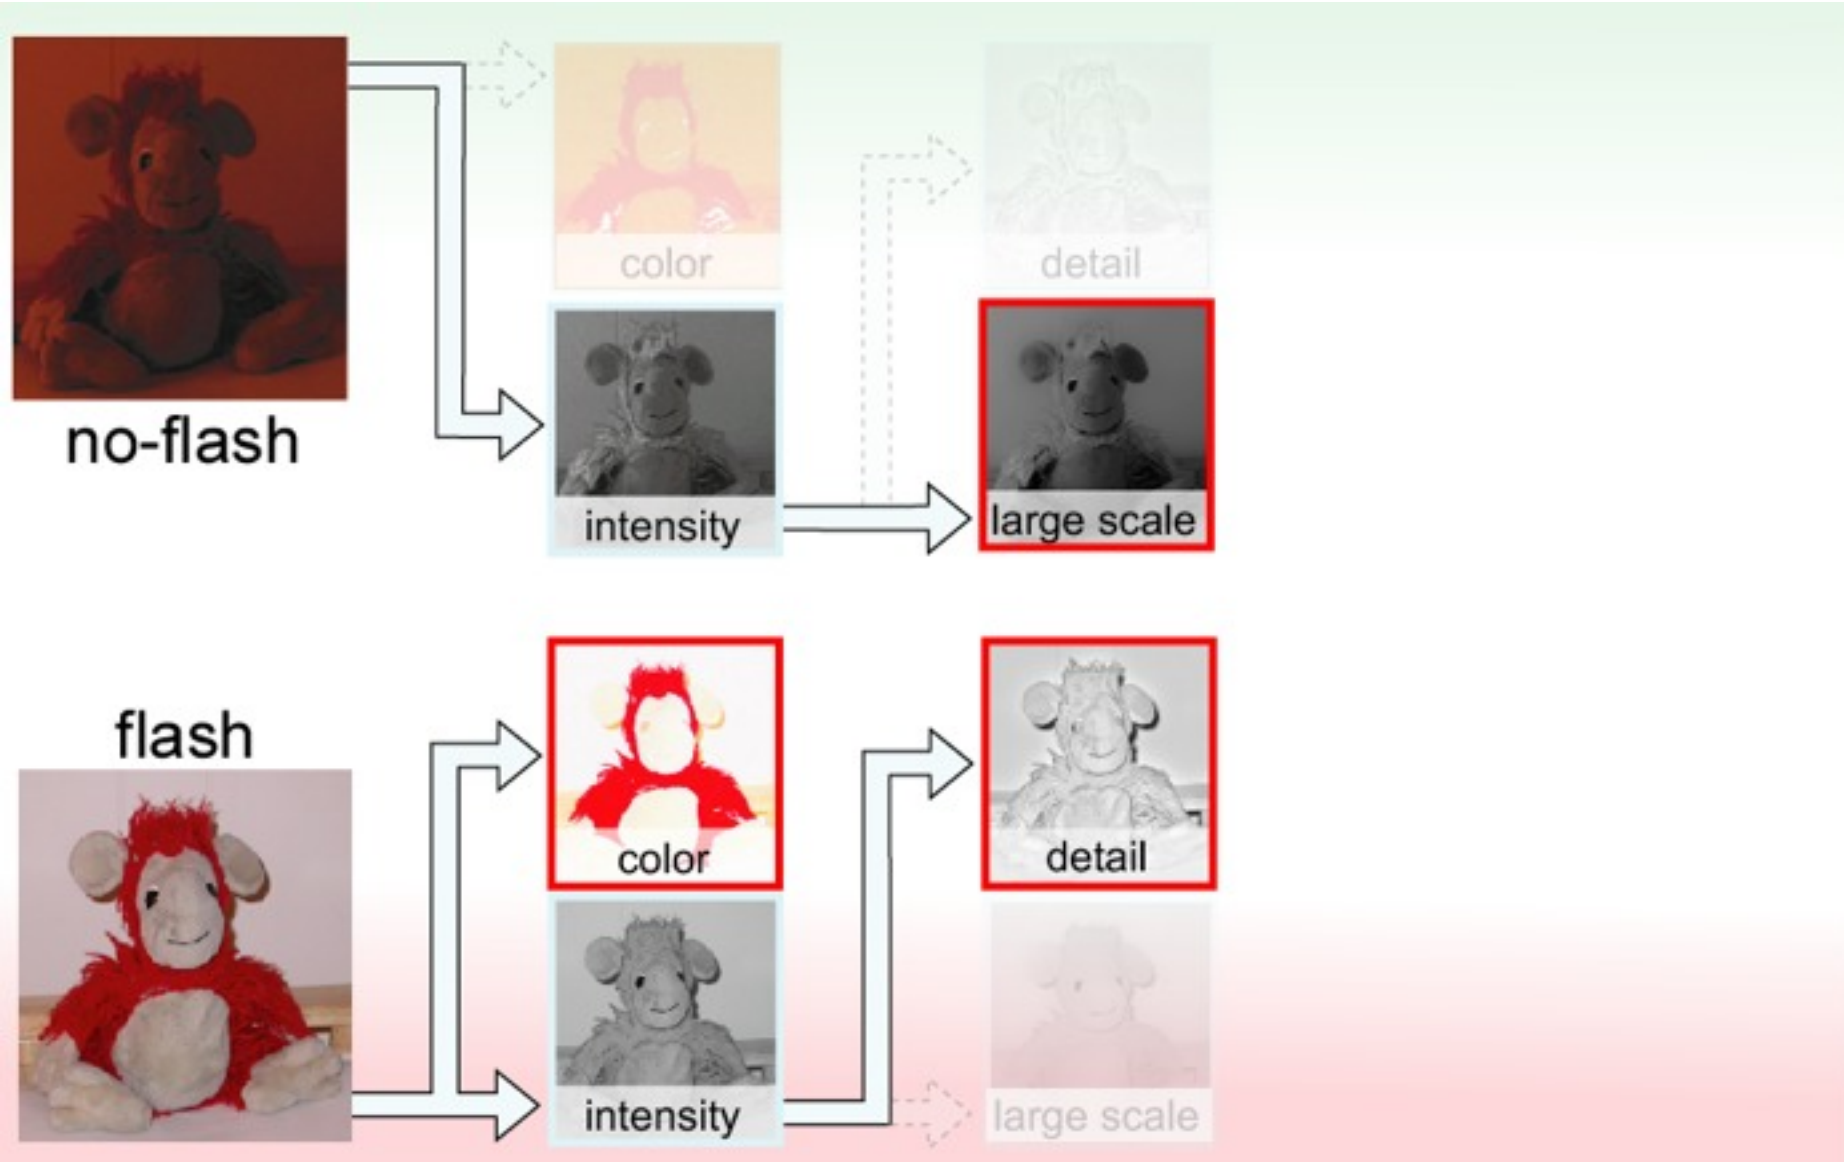

Our Approach

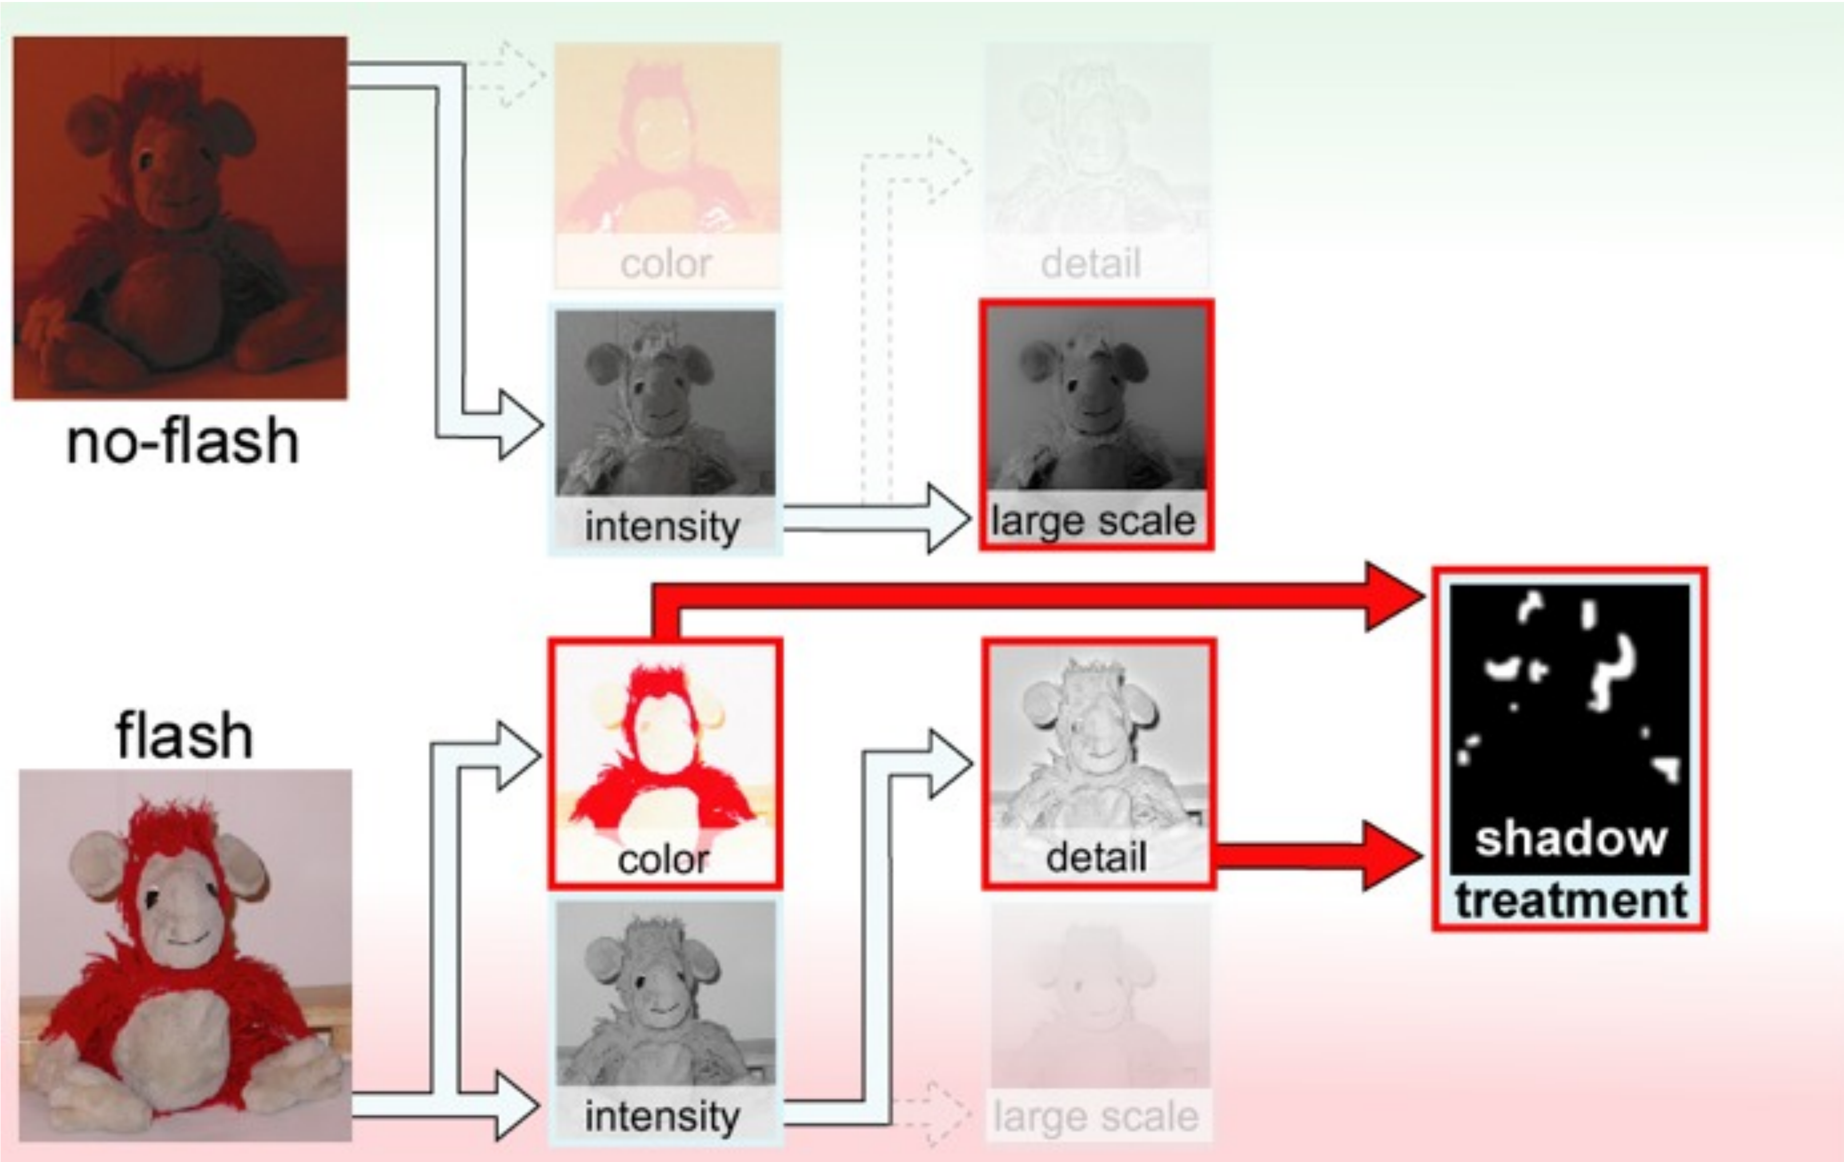

Our Approach

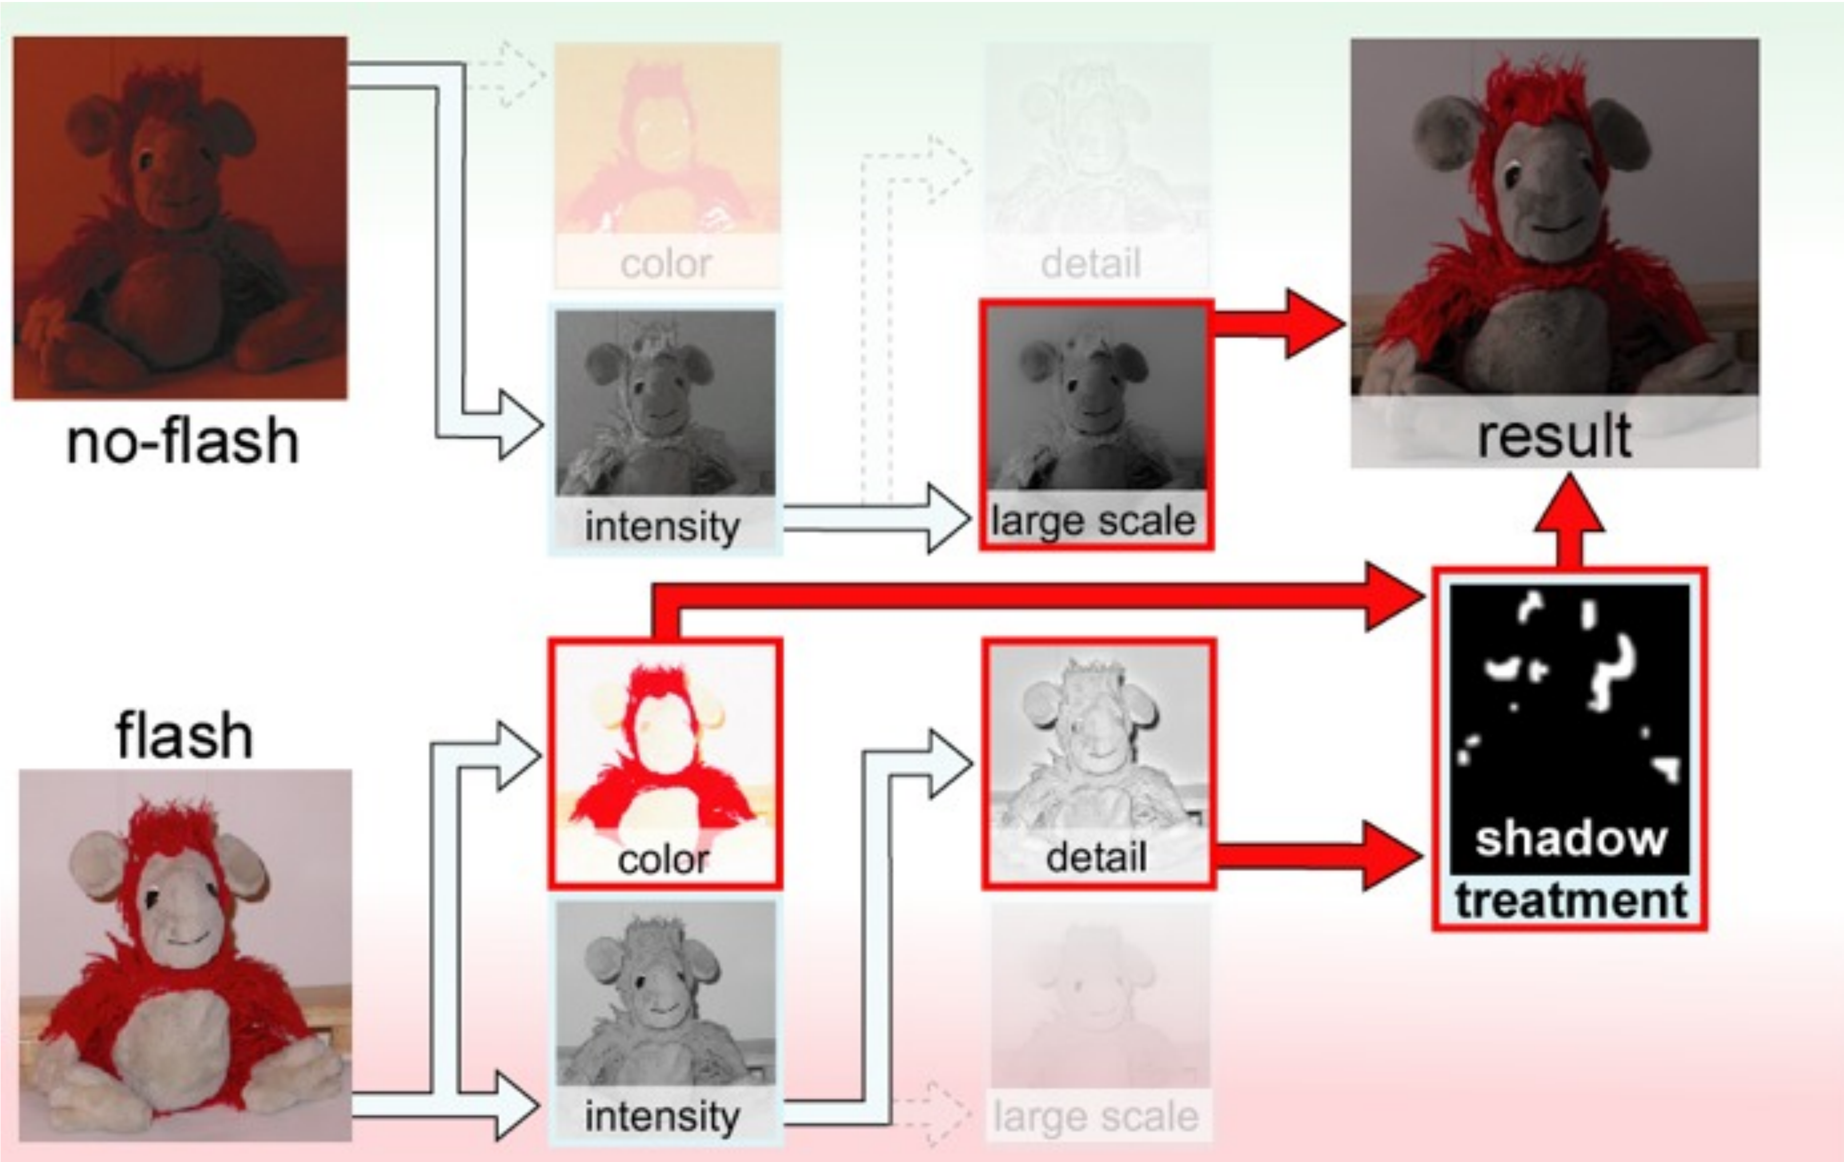

Our Approach

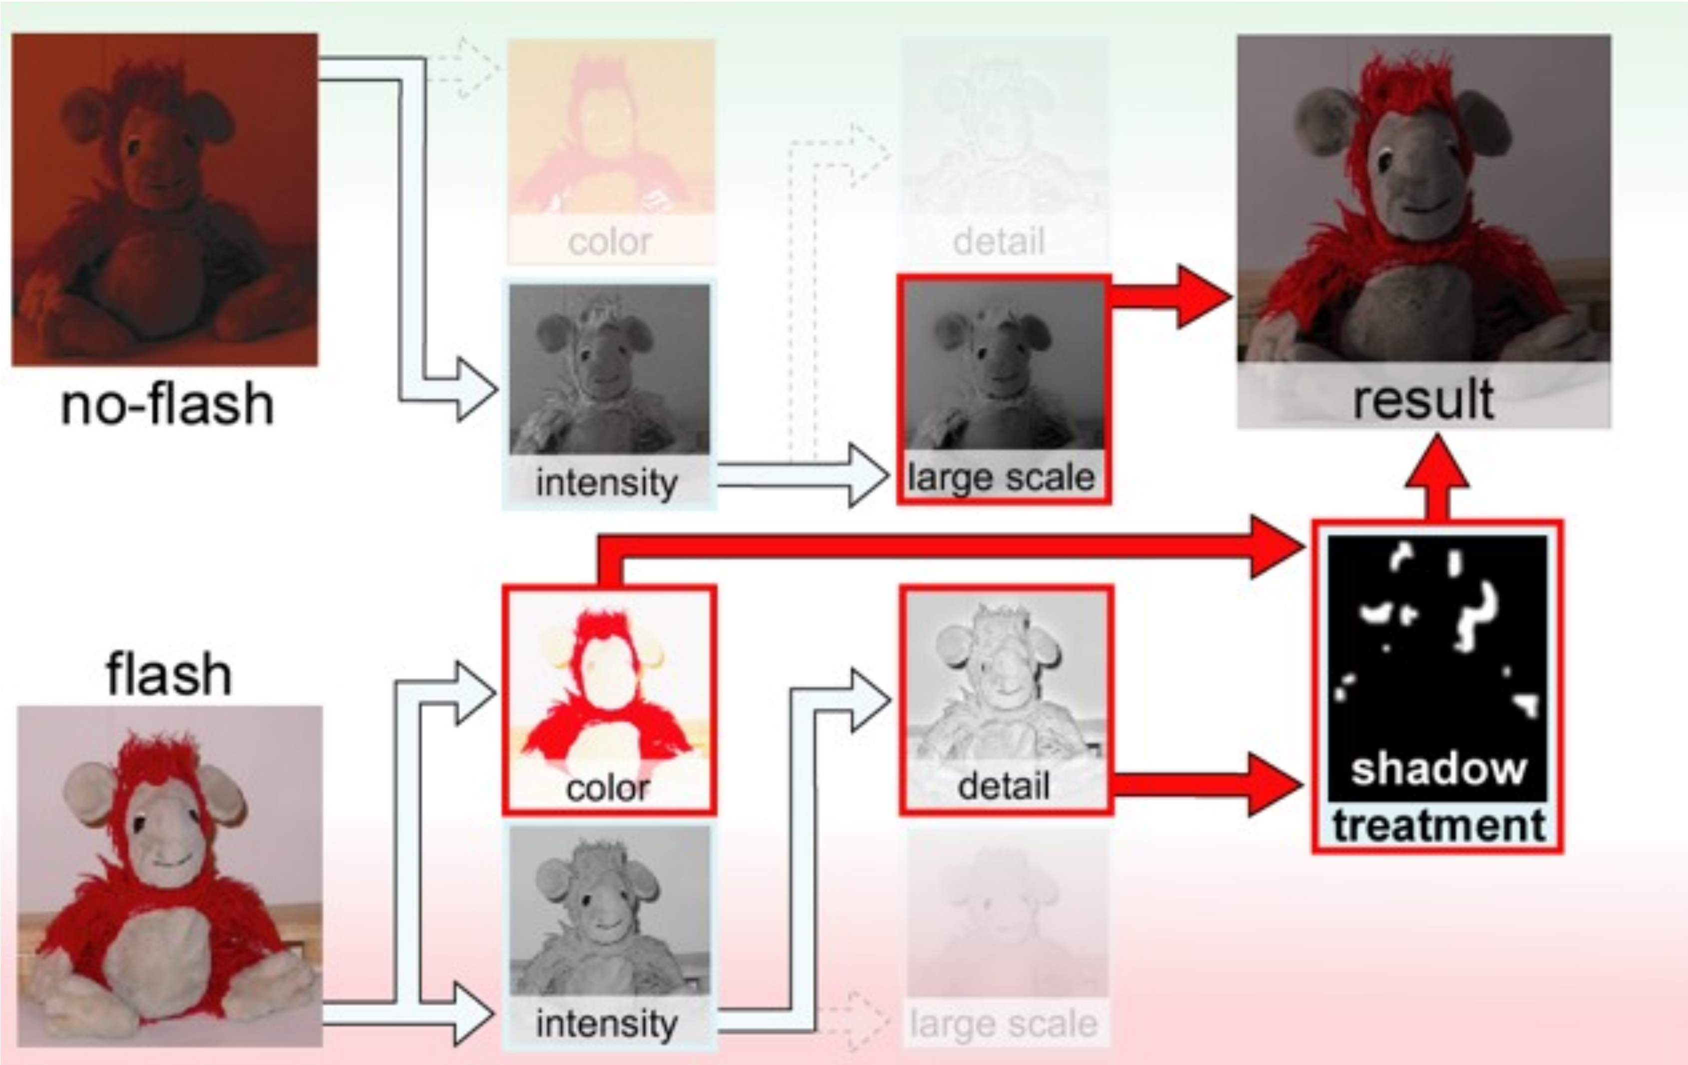

Our Approach

Registration

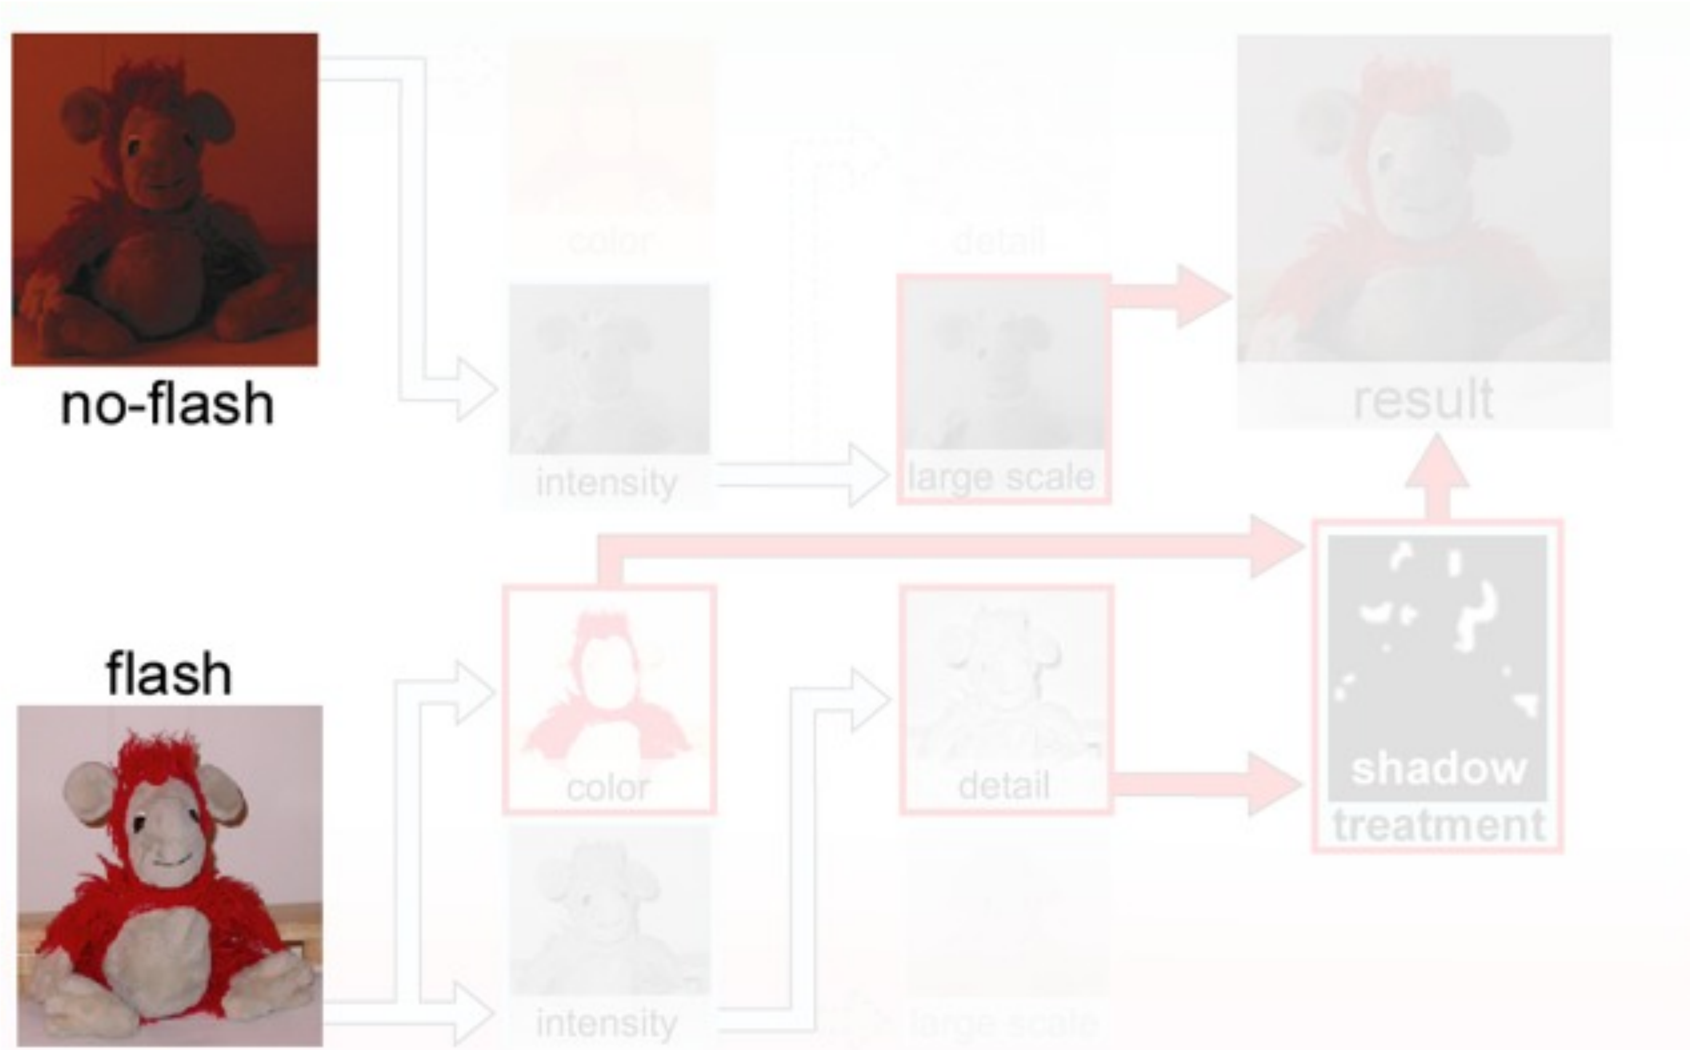

Registration

- **Align to compensate for camera movement**
- **Difficult because lighting changes**
- **Edge detection**

No-flash

Flash

- **See also Ward[2004], Kang[2003]**

Our Approach

Registration

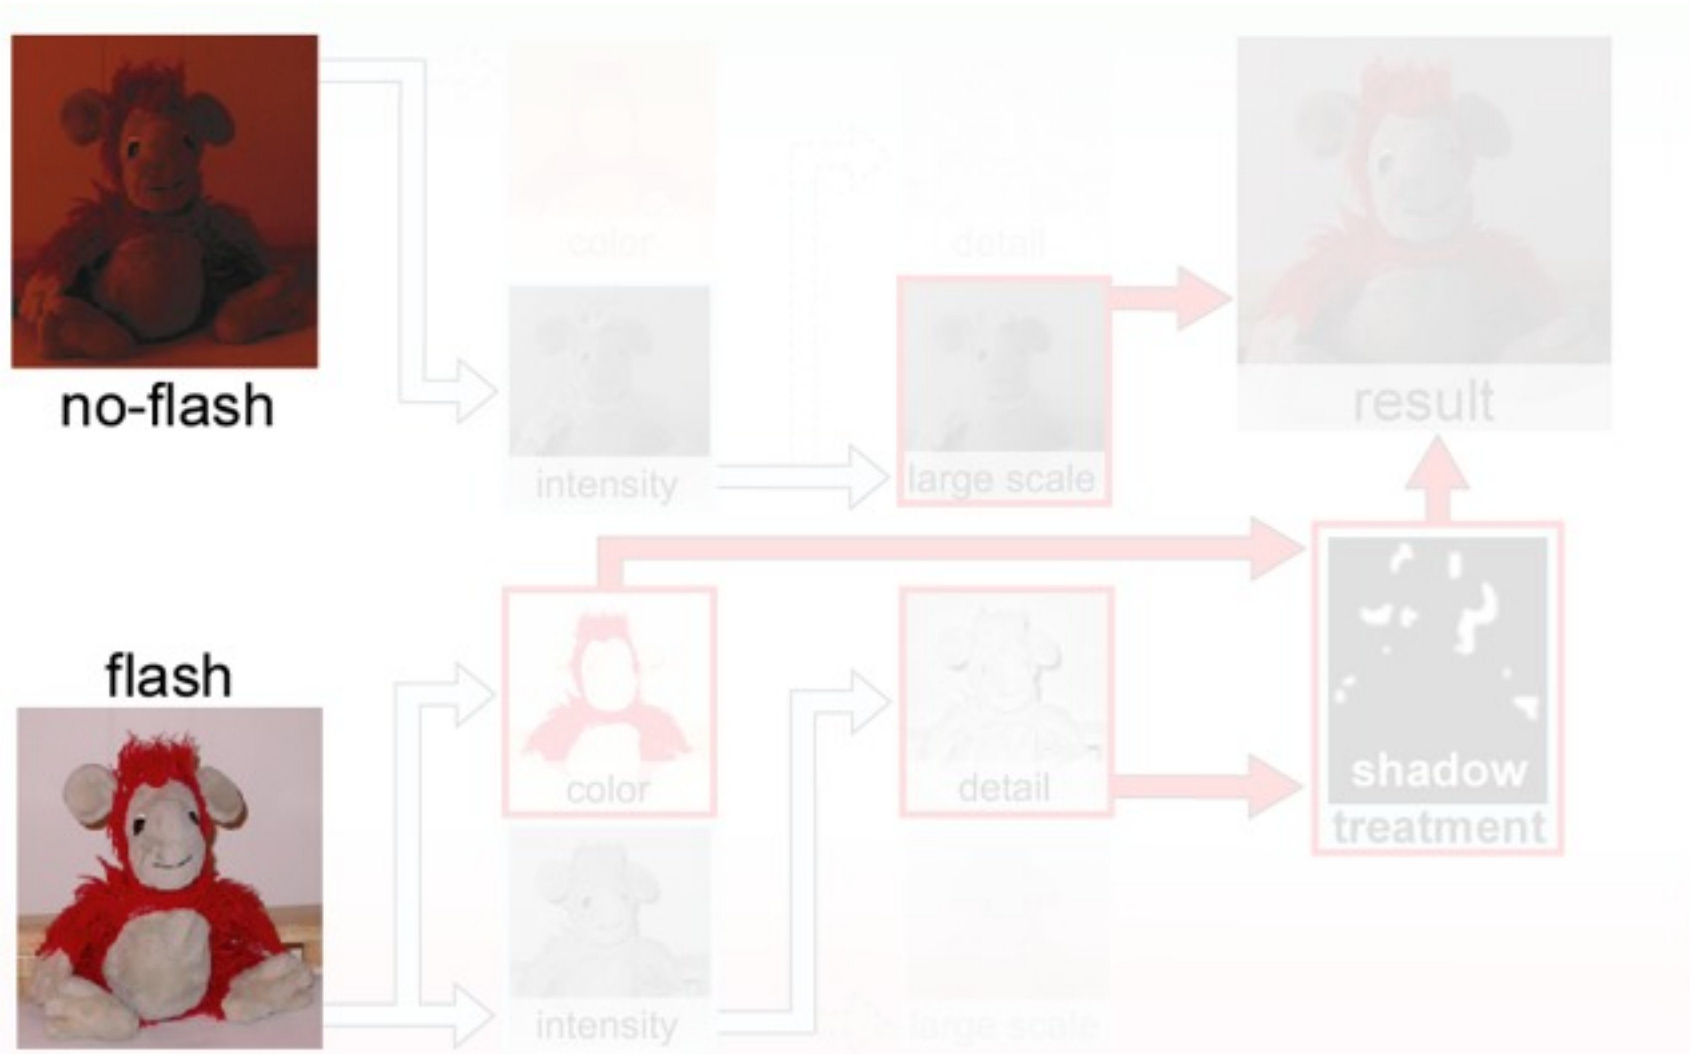

Our Approach

Decomposition

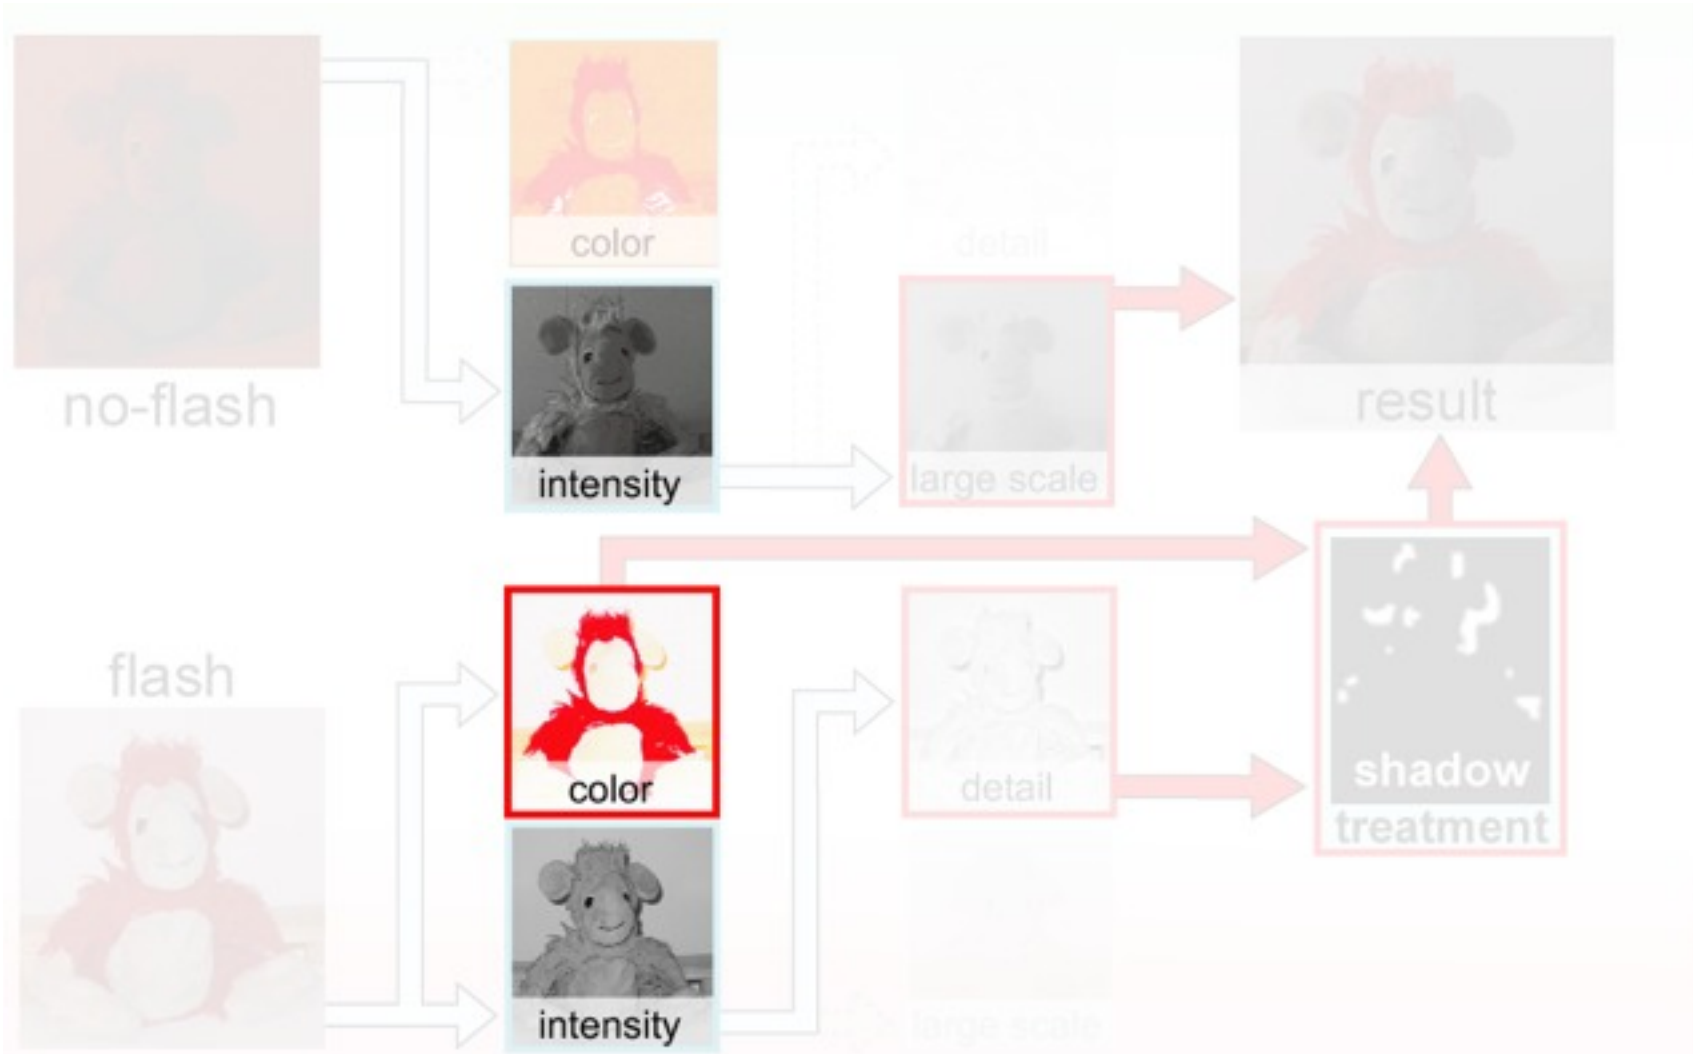

Decomposition

Color / Intensity:

original

Decomposition

Color / Intensity:

original intensity color

Our Approach

Decomposition

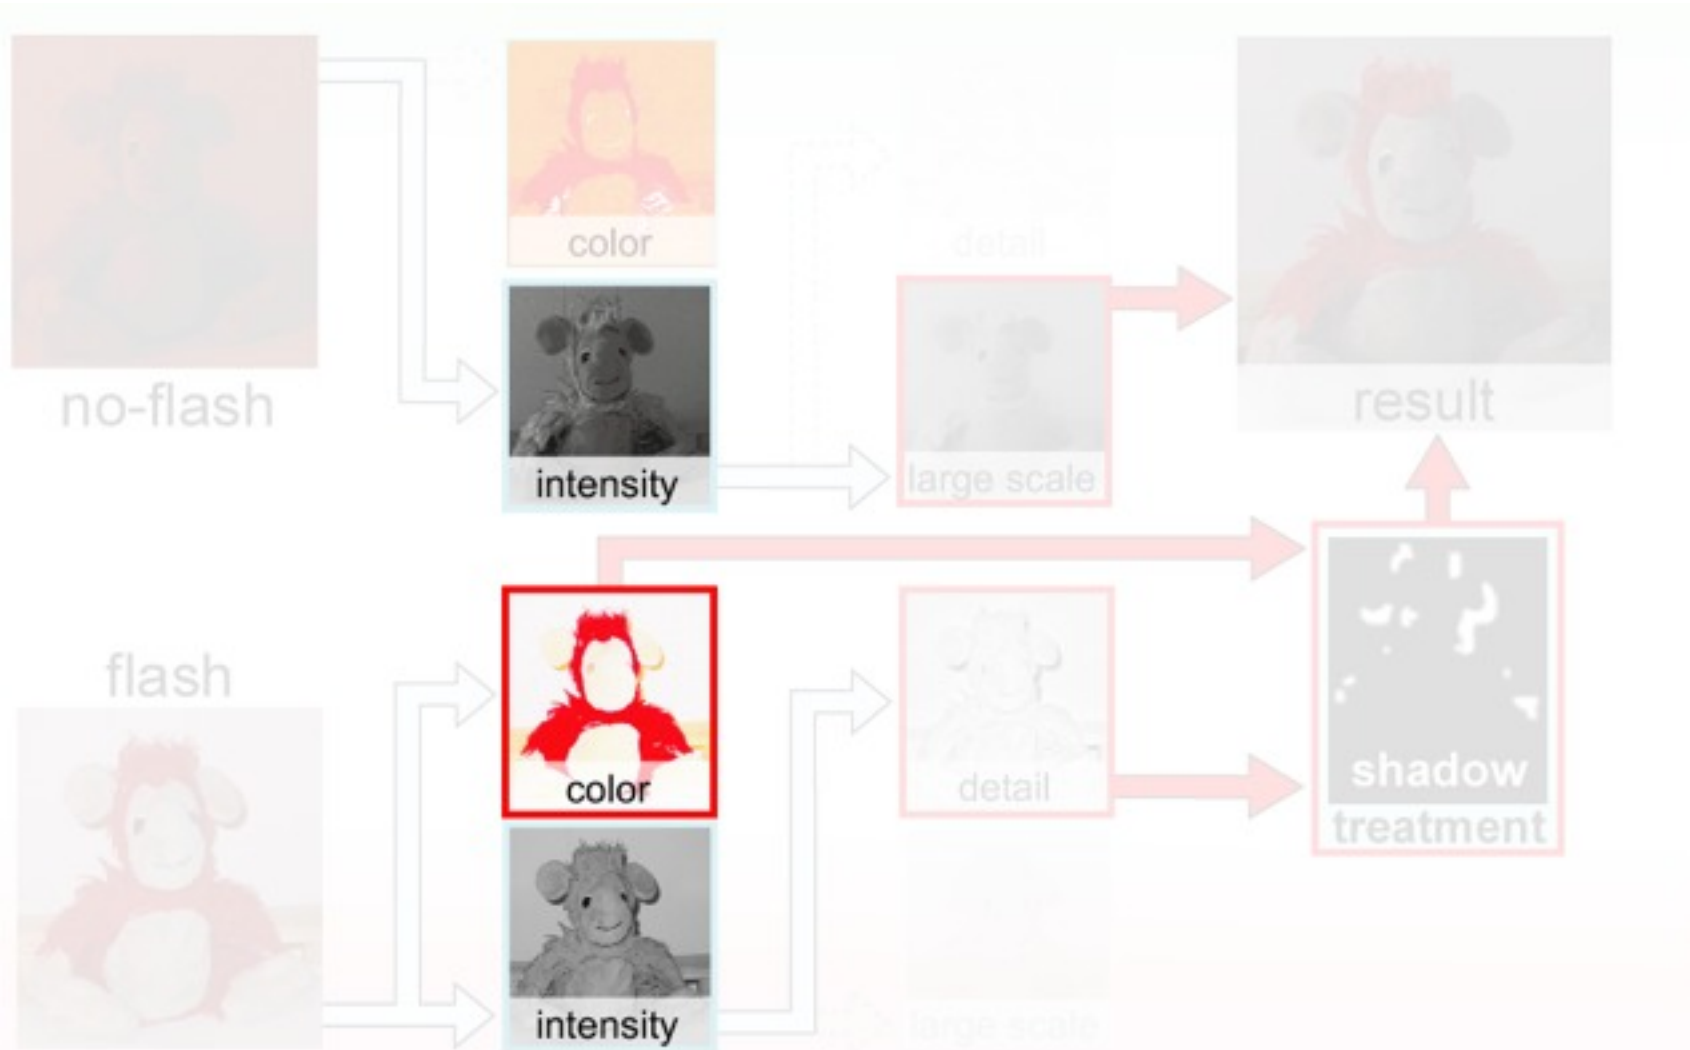

Our Approach

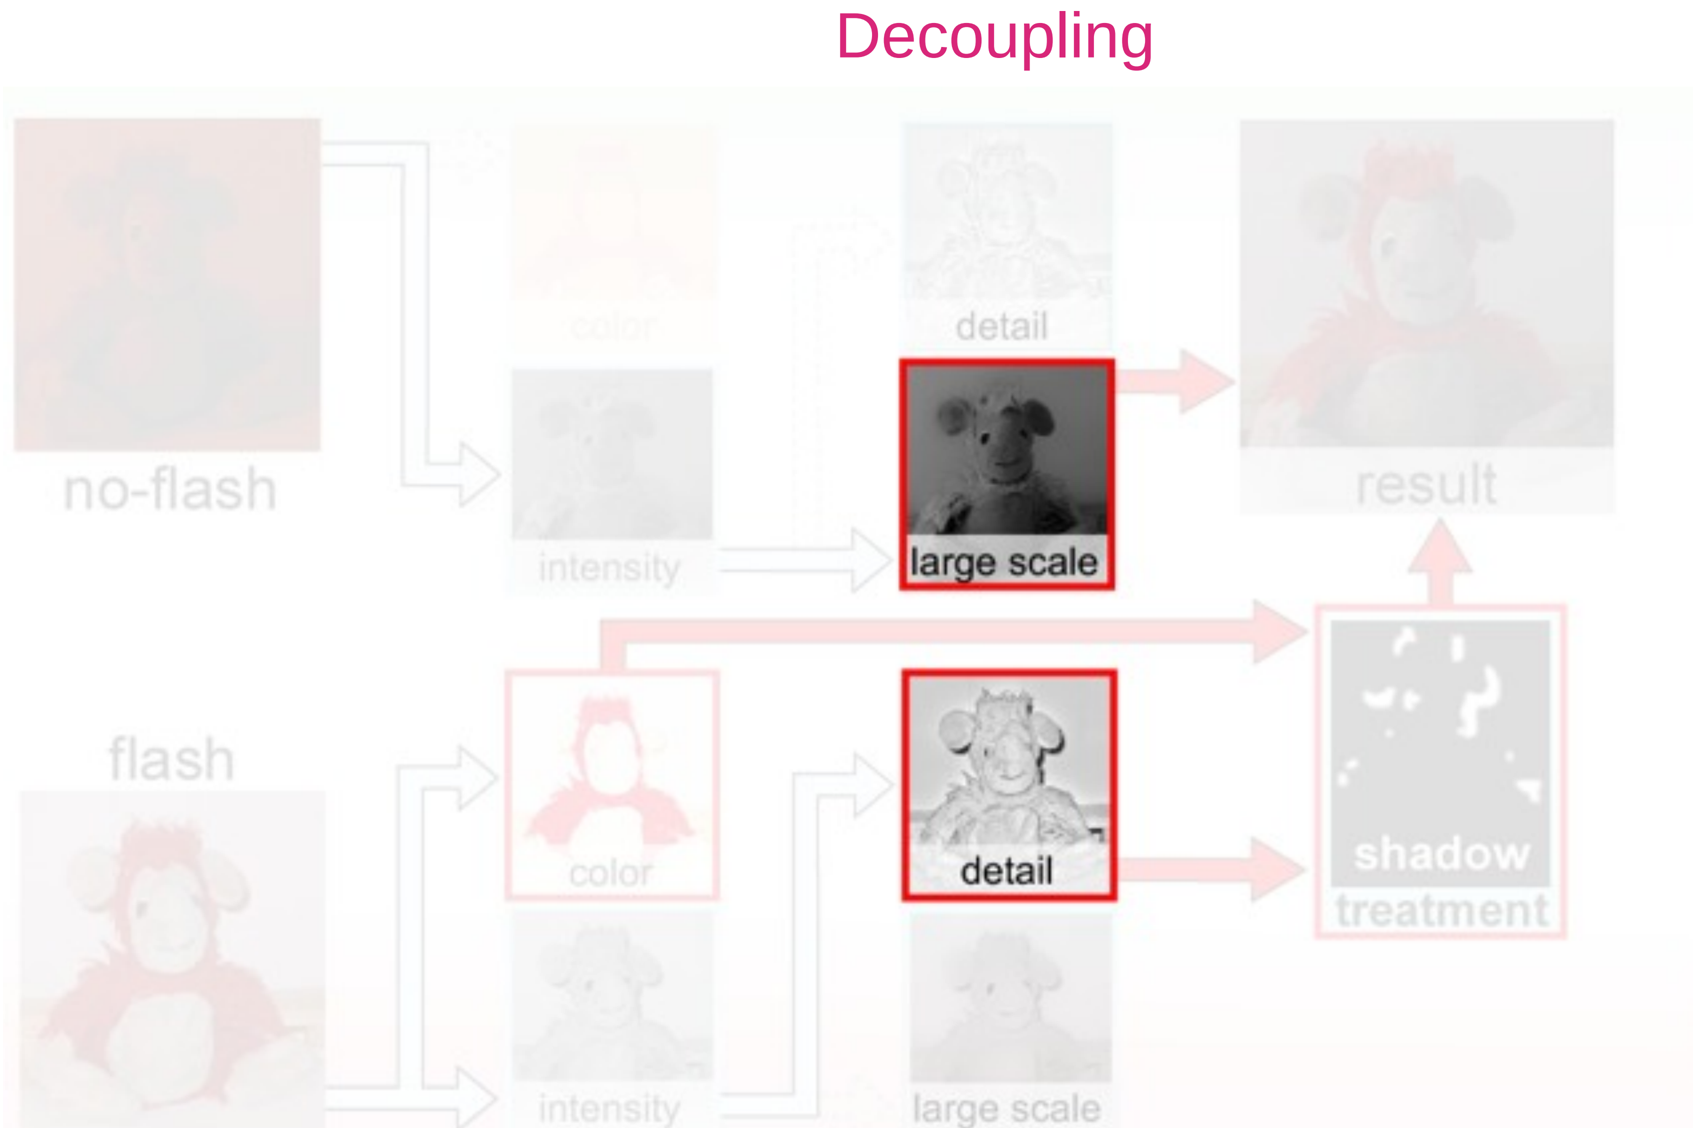

Decoupling

- **Lighting : Large-scale Variation**
- **Texture : Small-scale Variation**

Large-scale

Small-scale

Decoupling

- **Lighting : Large-scale Variation**
- **Texture : Small-scale Variation**
- Lighting : Large-scale Variation
- Texture : Small-scale Variation

Large-scale

Small-scale

Large-scale Layer

- **Gaussian filter**

$$I'(p) := \sum_q \overbrace{I(q)}^{image} \overbrace{f(p-q)}^{spatial}$$

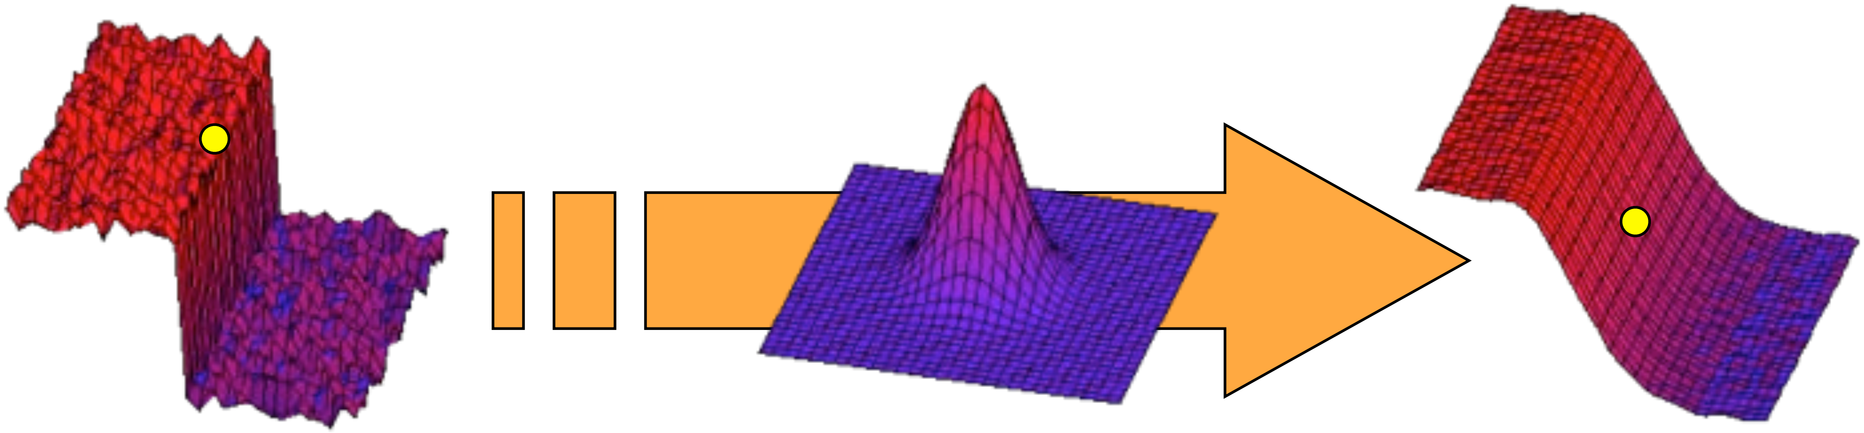

Large-scale Layer

- **Bilateral filter**
  - Smith and Brady 97
  - Tomasi & Manducci 98

$$I'(p) := \sum_q \overbrace{I(q)}^{image} \boxed{\overbrace{f(p-q)}^{spatial}} \overbrace{g(I(p)-I(q))}^{influence}$$

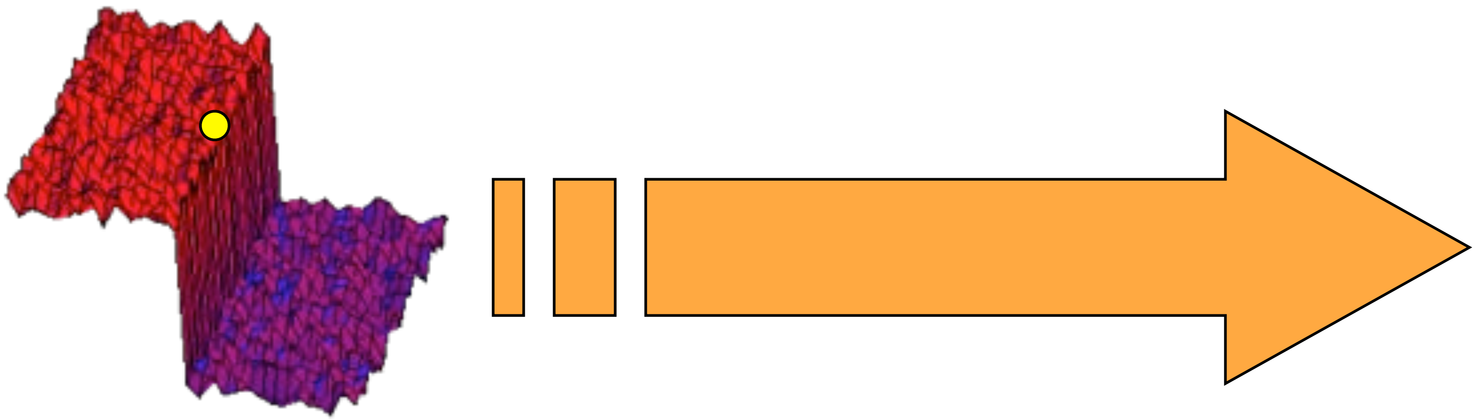

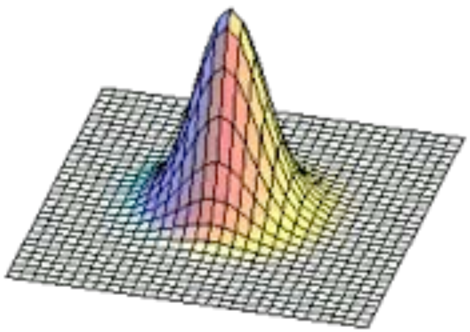

Large-scale Layer

- **Bilateral filter**
  - Smith and Brady 97
  - Tomasi & Manducci 98

$$I'(p) := \sum_q \overbrace{I(q)}^{image}\ \overbrace{f(p-q)}^{spatial}\ \boxed{\overbrace{g(I(p)-I(q))}^{influence}}$$

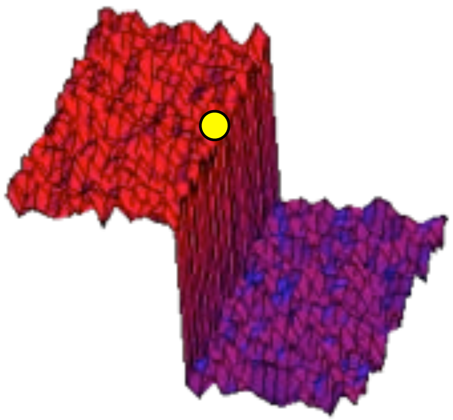

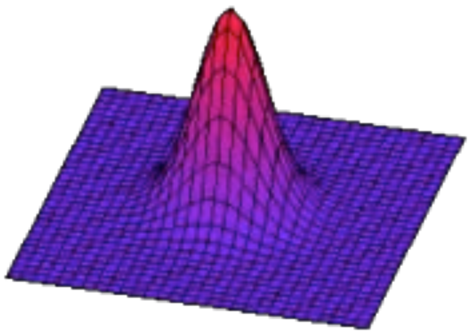

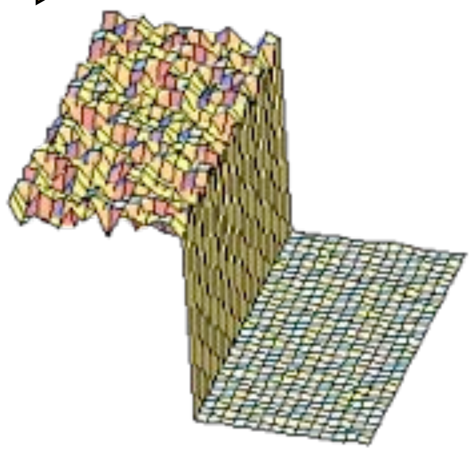

Large-scale Layer

- **Bilateral filter**
  - Smith and Brady 97
  - Tomasi & Manducci 98

$$I'(p) := \sum_q \overbrace{I(q)}^{image} \boxed{\overbrace{f(p-q)}^{spatial} \overbrace{g(I(p)-I(q))}^{influence}}$$

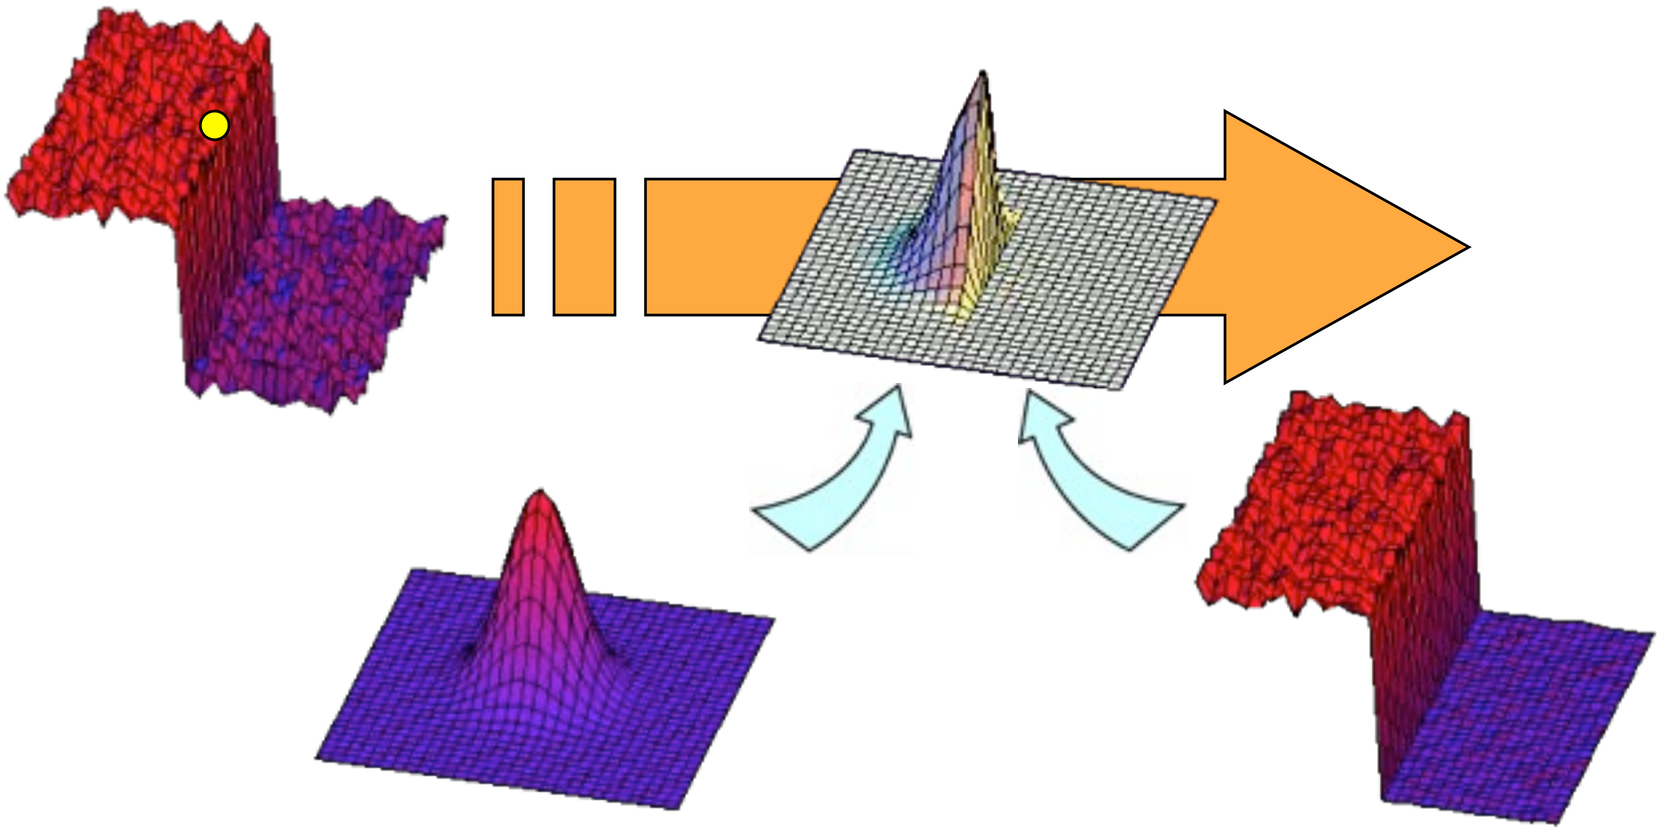

Large-scale Layer

- **Bilateral filter**
  - Smith and Brady 97
  - Tomasi & Manducci 98

$$I'(p) := \sum_q \overbrace{I(q)}^{image} \overbrace{f(p-q)}^{spatial} \overbrace{g(I(p)-I(q))}^{influence}$$

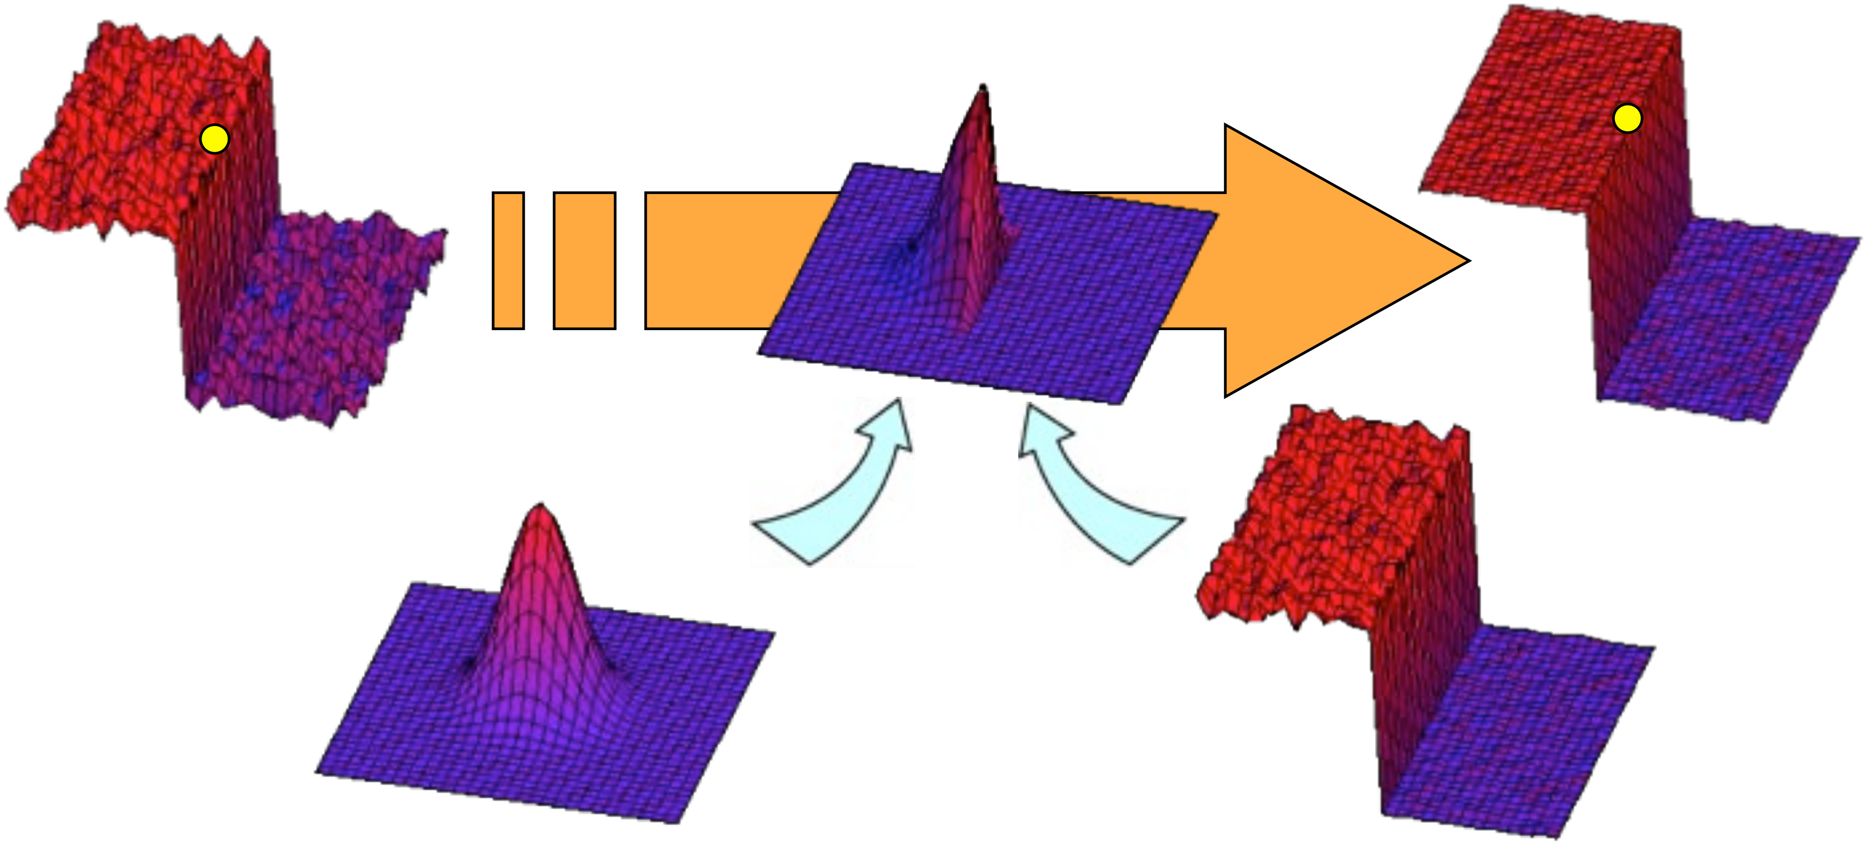

Large-scale Layer

- **Bilateral filter**

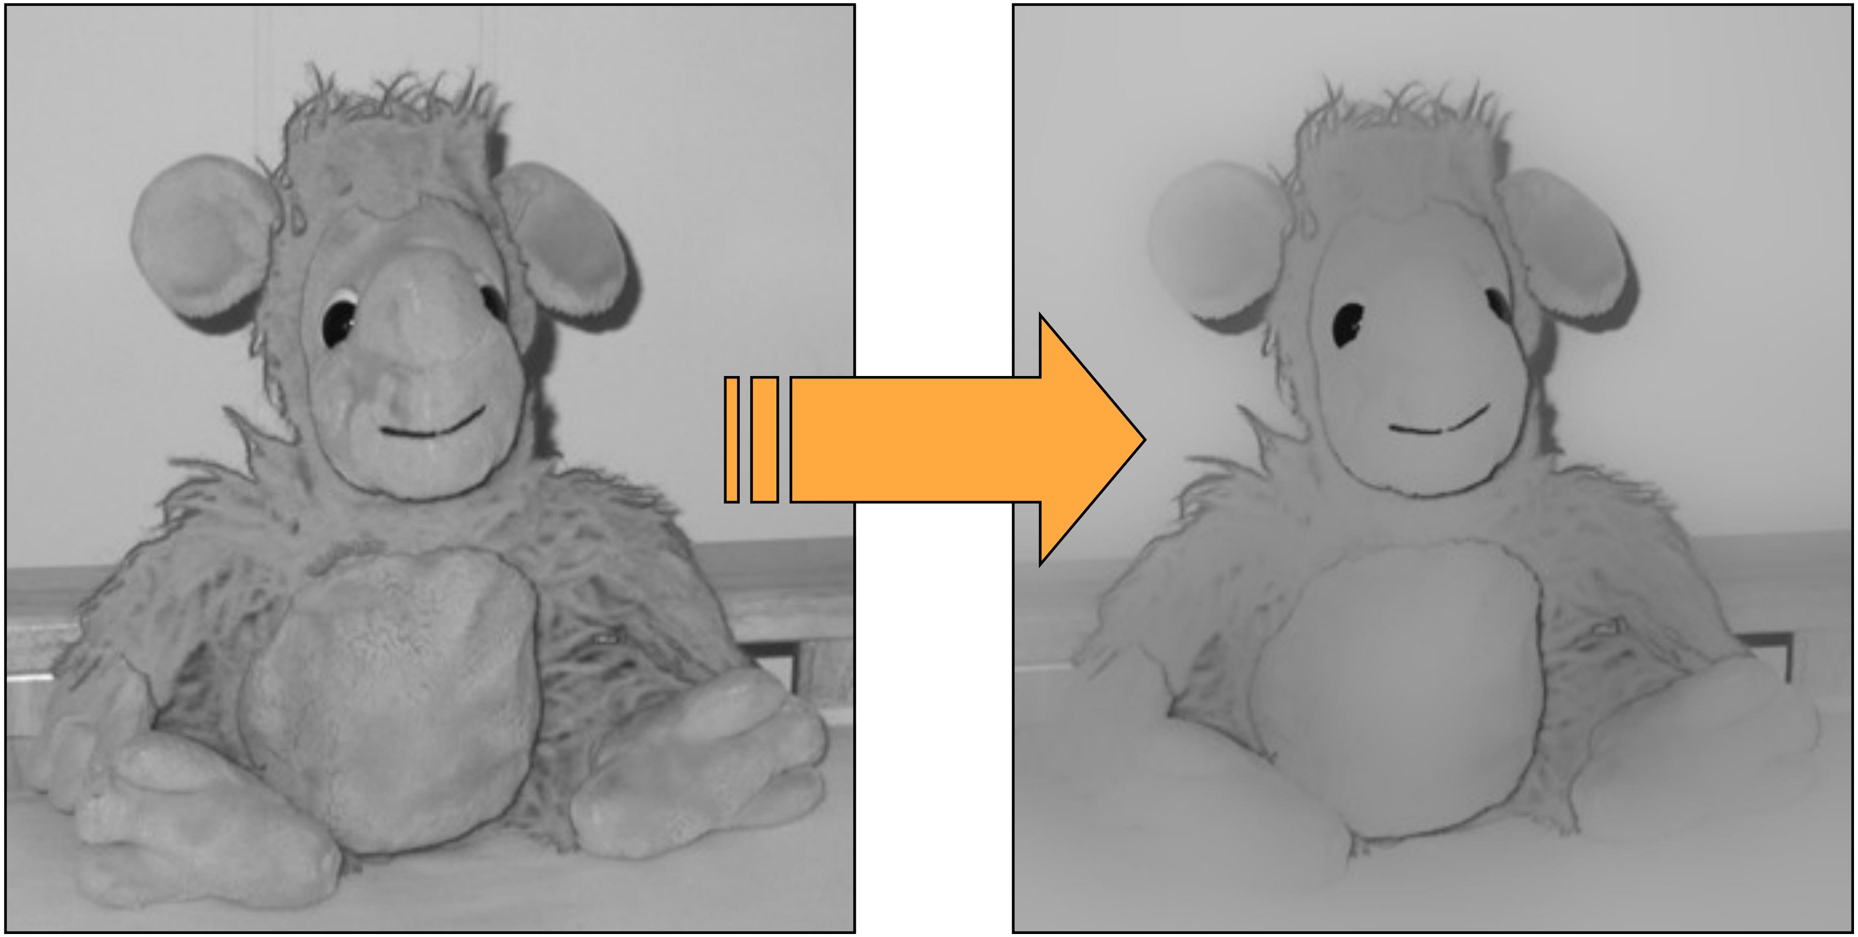

Large-scale Layer

- **Cross Bilateral filter**

  - Better smoothing for no-flash
  - Value penalization based on flash image

  - ➢ edge stopping from flash image

Detail Layer

Intensity

/

Large-scale

=

Detail

Recombination: Large scale * Detail = Intensity

Recombination

*

=

Large-scale
No-flash

Detail
Flash

Intensity
Result

Recombination: Large scale * Detail = Intensity

Recombination

shadows

Intensity Result

*

Color Flash

≈

Result

Recombination: Intensity * Color = Original

Our Approach

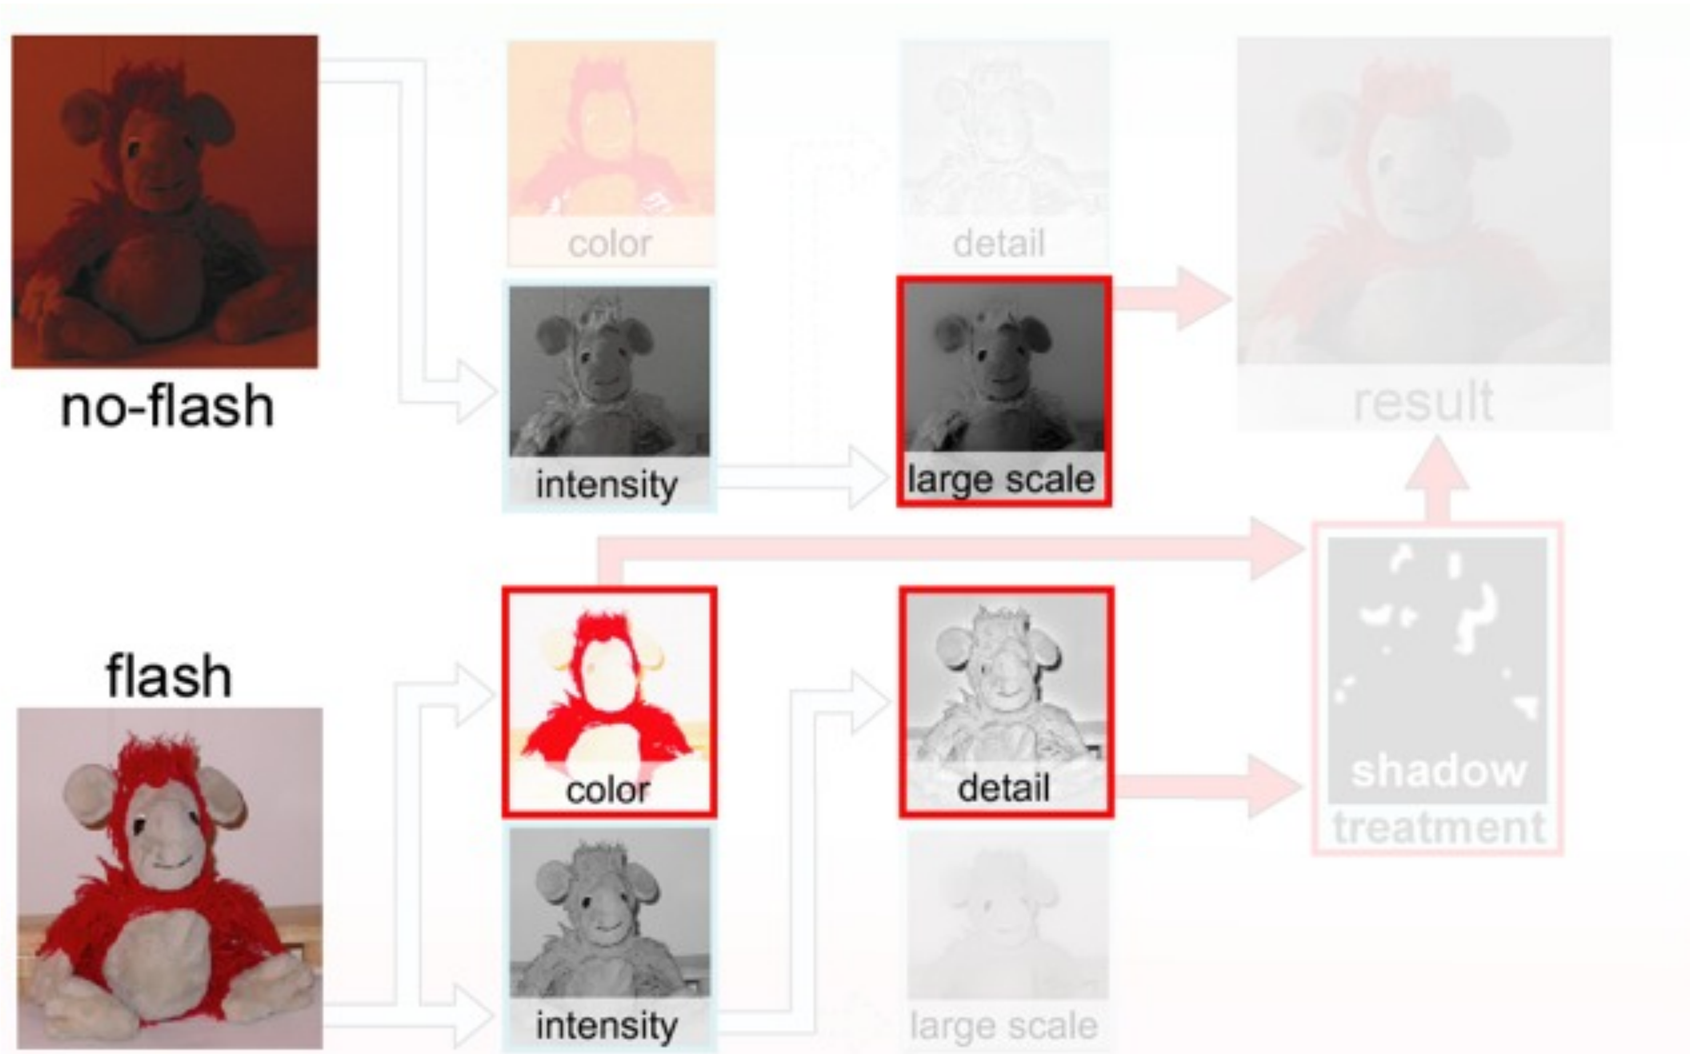

Our Approach

Shadow Detection/
Treatment

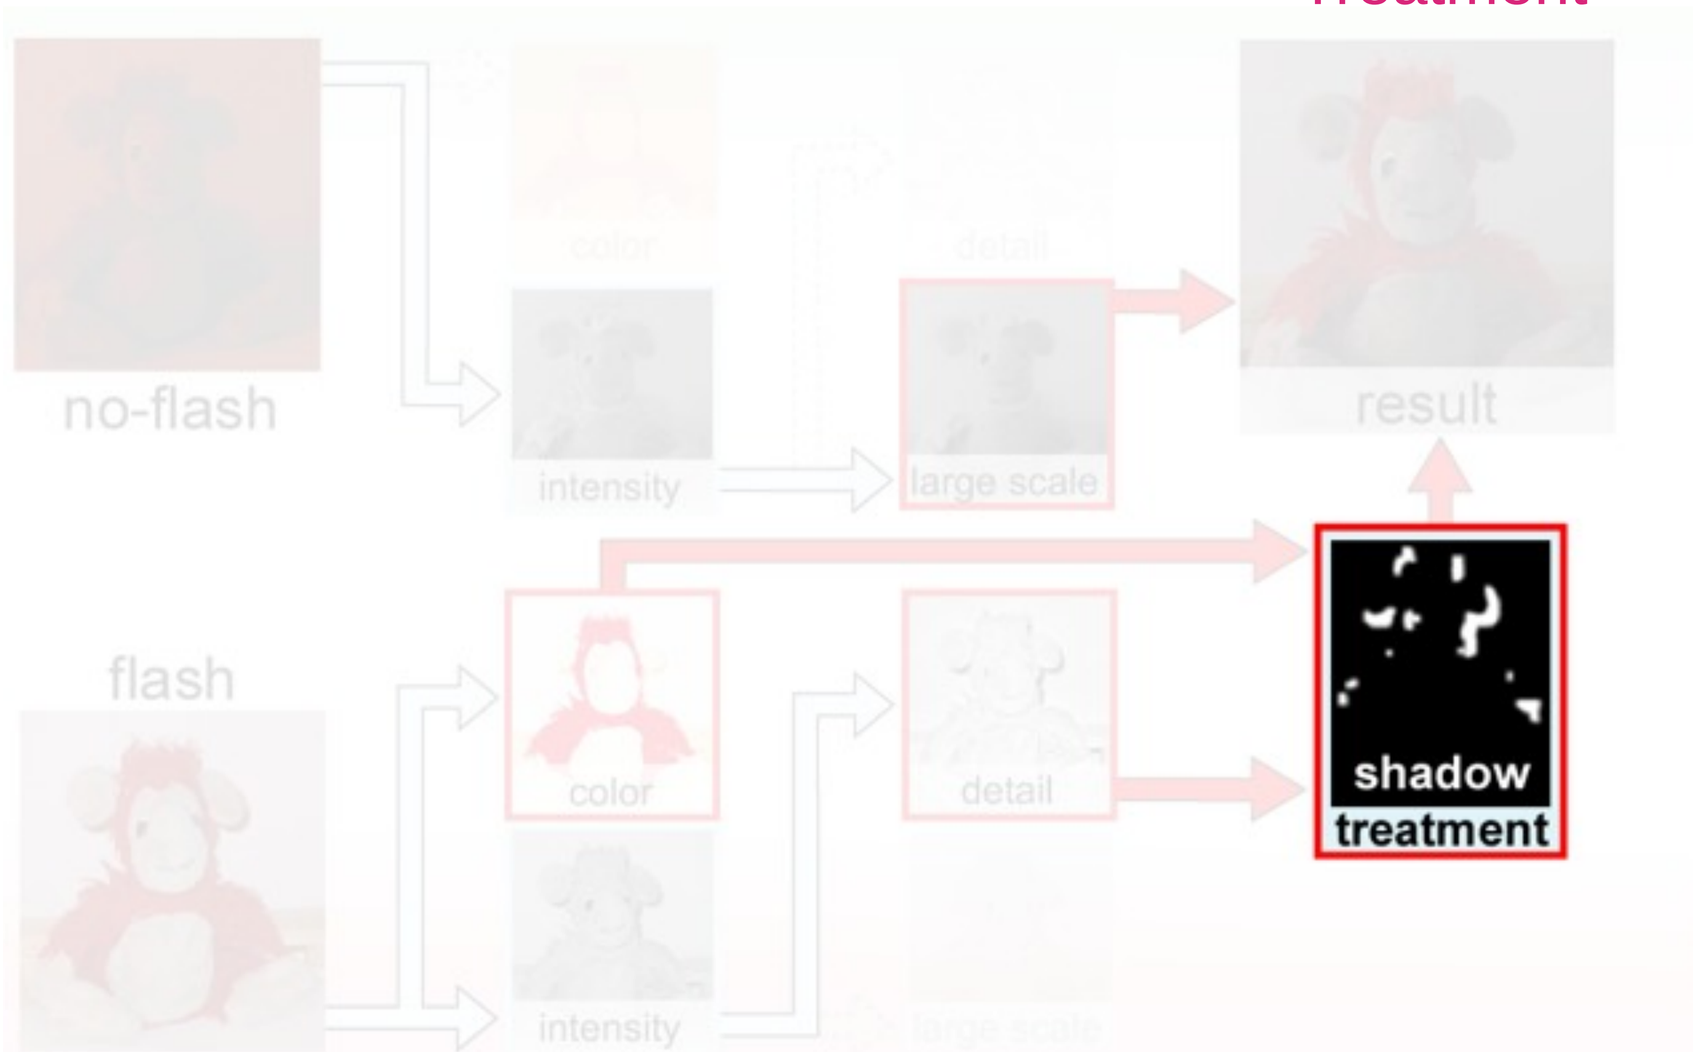

Problem

Shadow Correction – Why?

Shadow Correction – Why?

Several artifacts:

Shadow Correction – Why?

Several artifacts:
- at shadow boundary

Shadow Correction – Why?

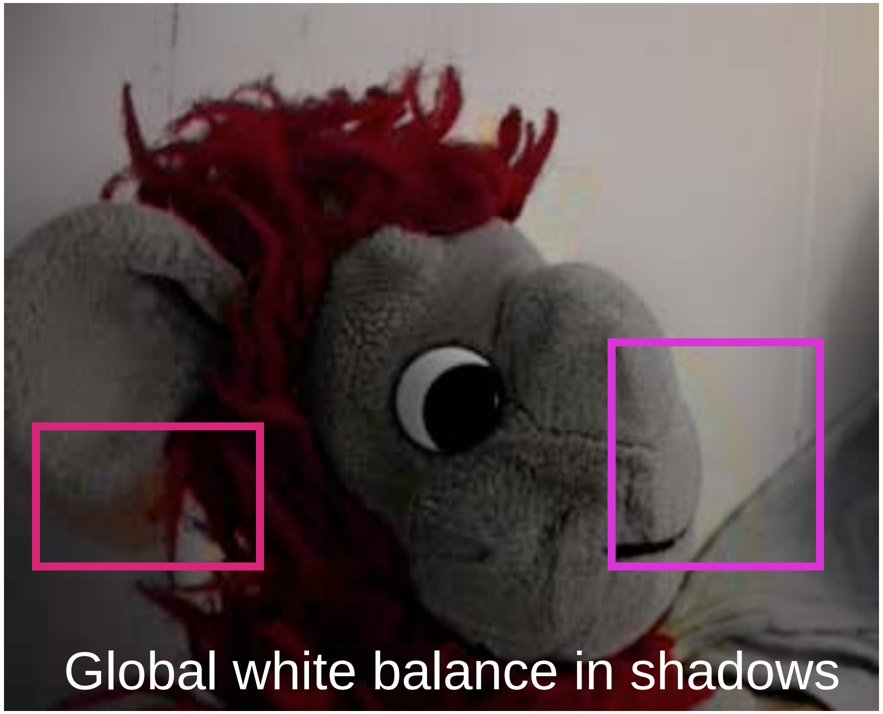

Several artifacts:

- at shadow boundary
- inside shadows (color bleeding)

Shadow Correction – Why?

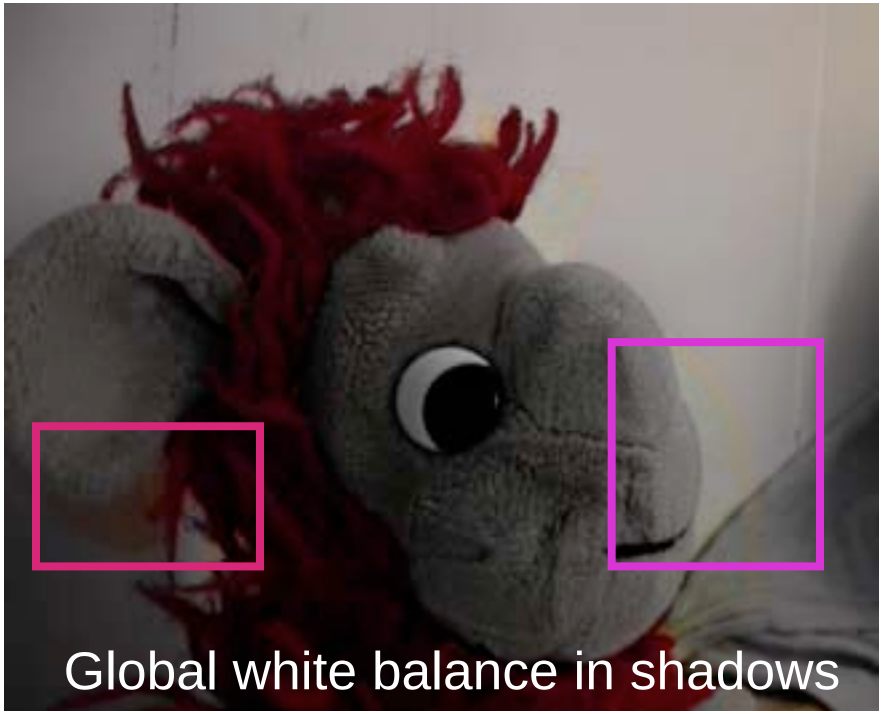

Several artifacts:

- at shadow boundary
- inside shadows (color bleeding)

➢ shadows need to be corrected !

Shadow Correction – Why?

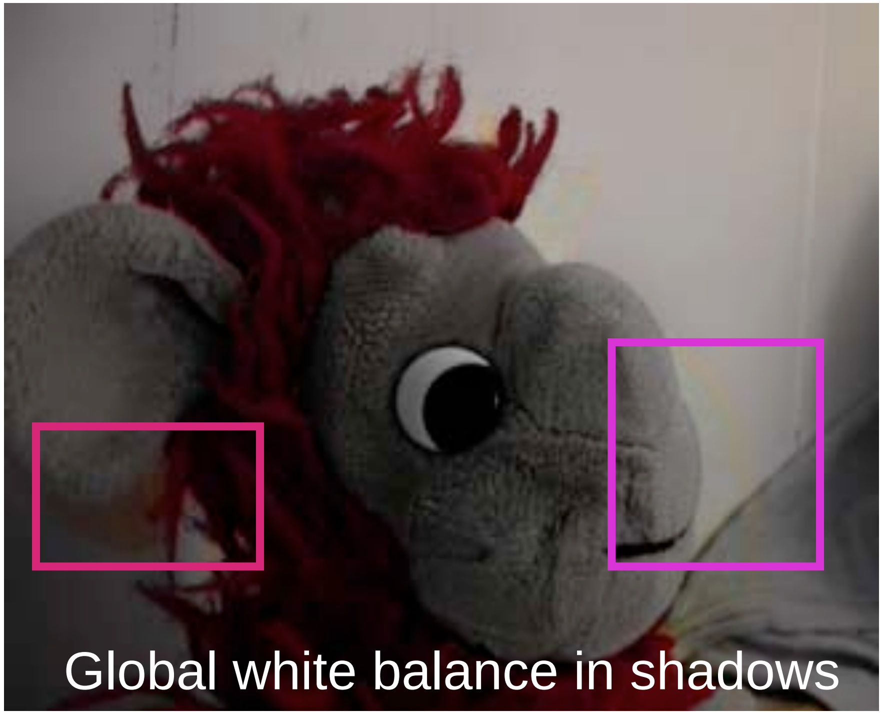

Several artifacts:

- at shadow boundary
- inside shadows (color bleeding)

➢ shadows need to be corrected !

... and detected

Shadow Detection

Shadow Detection

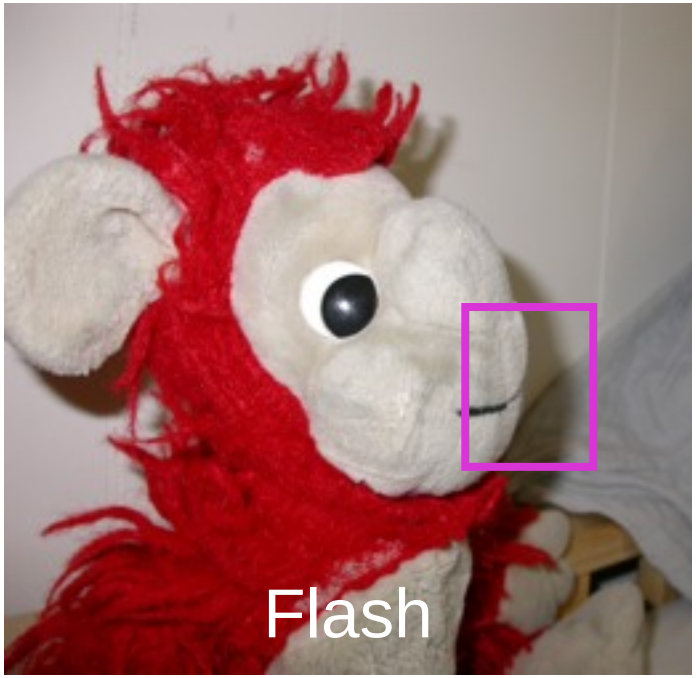

Shadow Detection

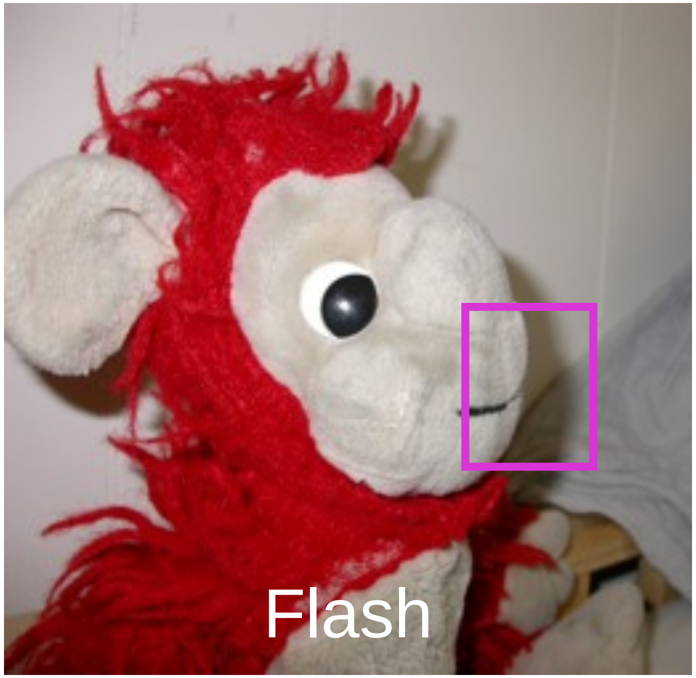

- Umbra

Shadow Detection

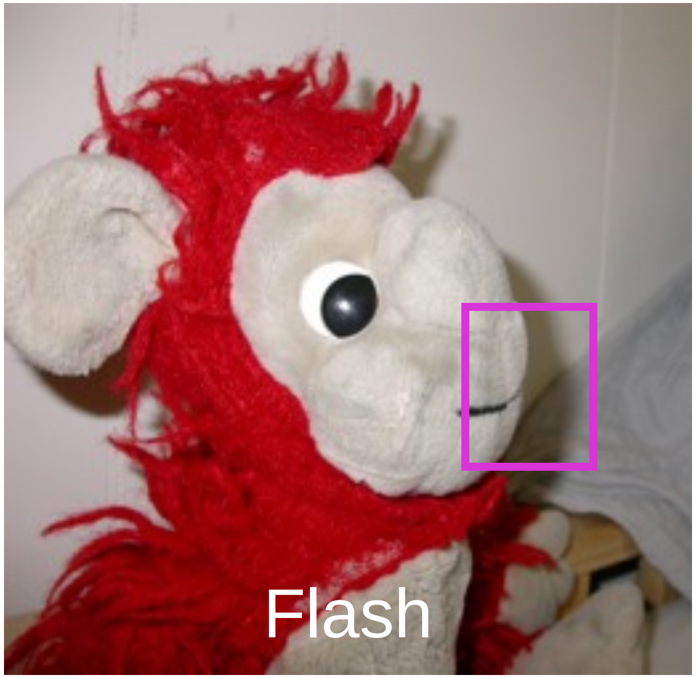

- Umbra
- Penumbra

Shadow Detection

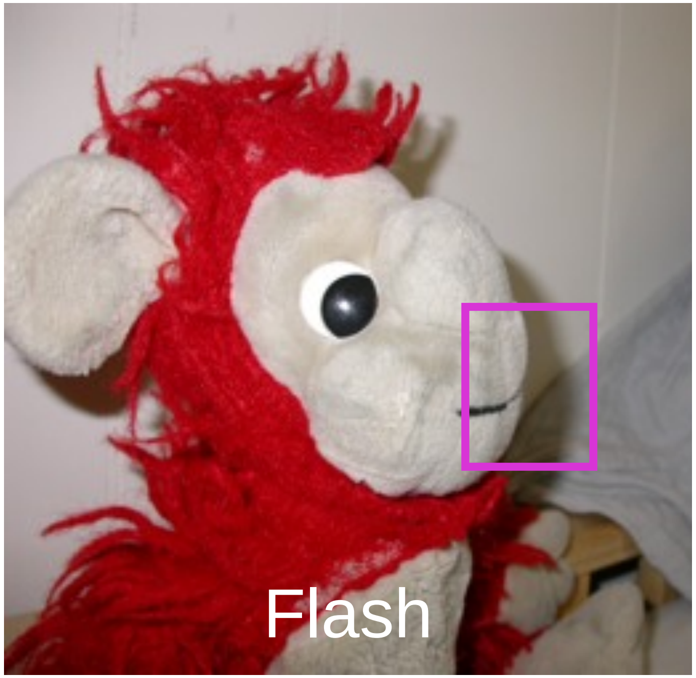

- Umbra
- Penumbra

➢ Detection in two steps

Shadow Detection

Shadow Detection

Umbra detection

Shadow Detection

Umbra detection

- **uniform, scattered light from flash**

Shadow Detection

Umbra detection

- **uniform, scattered light from flash**

➢ **Difference of the two photos ΔI reveals these regions**

Shadow Detection

Umbra detection

- **uniform, scattered light from flash**

➢ **Difference of the two photos ΔI reveals these regions**

Flash

−

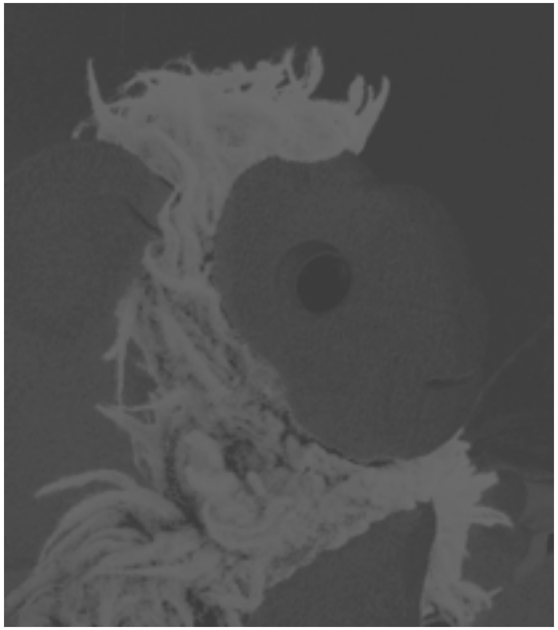

No-flash

=

ΔI

Shadow Detection

Umbra detection

- **Difference ΔI = light added by the flash**

- **Goal: Find a threshold for ΔI**

ΔI

Shadow Detection

Umbra detection

- **Umbra Detection Threshold (details in paper)**

ΔI

Shadow Detection

Umbra detection

- **Umbra Detection Threshold (details in paper)**

ΔI

Shadow Detection

Umbra detection

- **Umbra Detection Threshold (details in paper)**

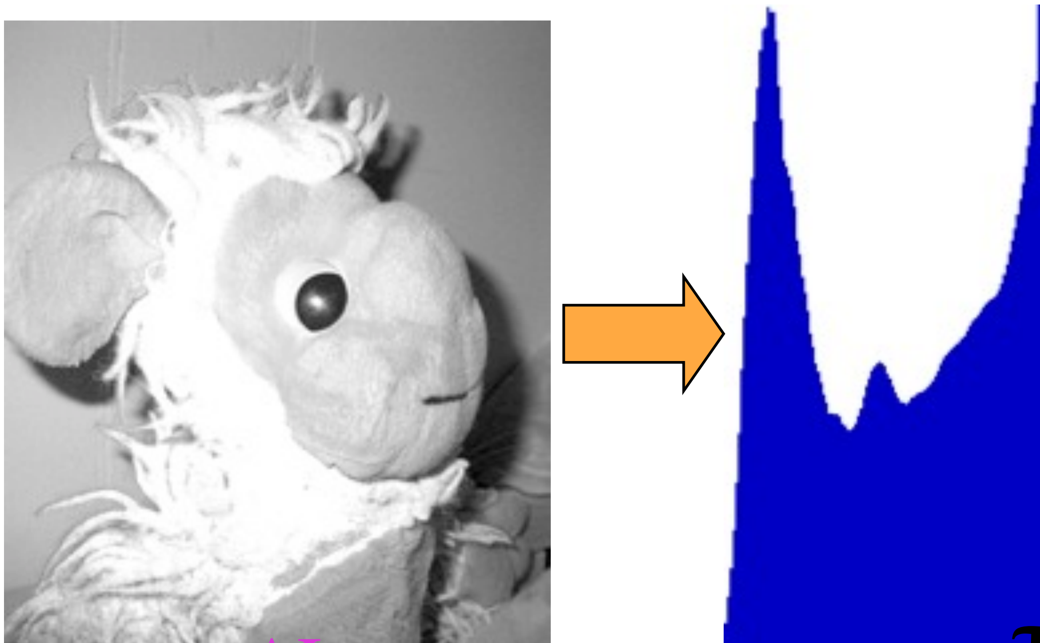

ΔI τ_{cut}

Shadow Detection

Umbra detection

- **Umbra Detection Threshold (details in paper)**

ΔI

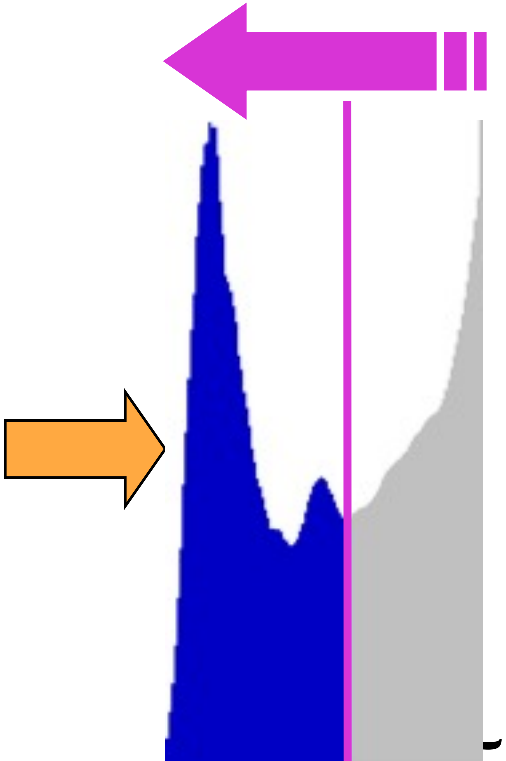

τ_{cut}

Shadow Detection

Umbra detection

- ➢ **Umbra Detection Threshold**
  (details in paper)

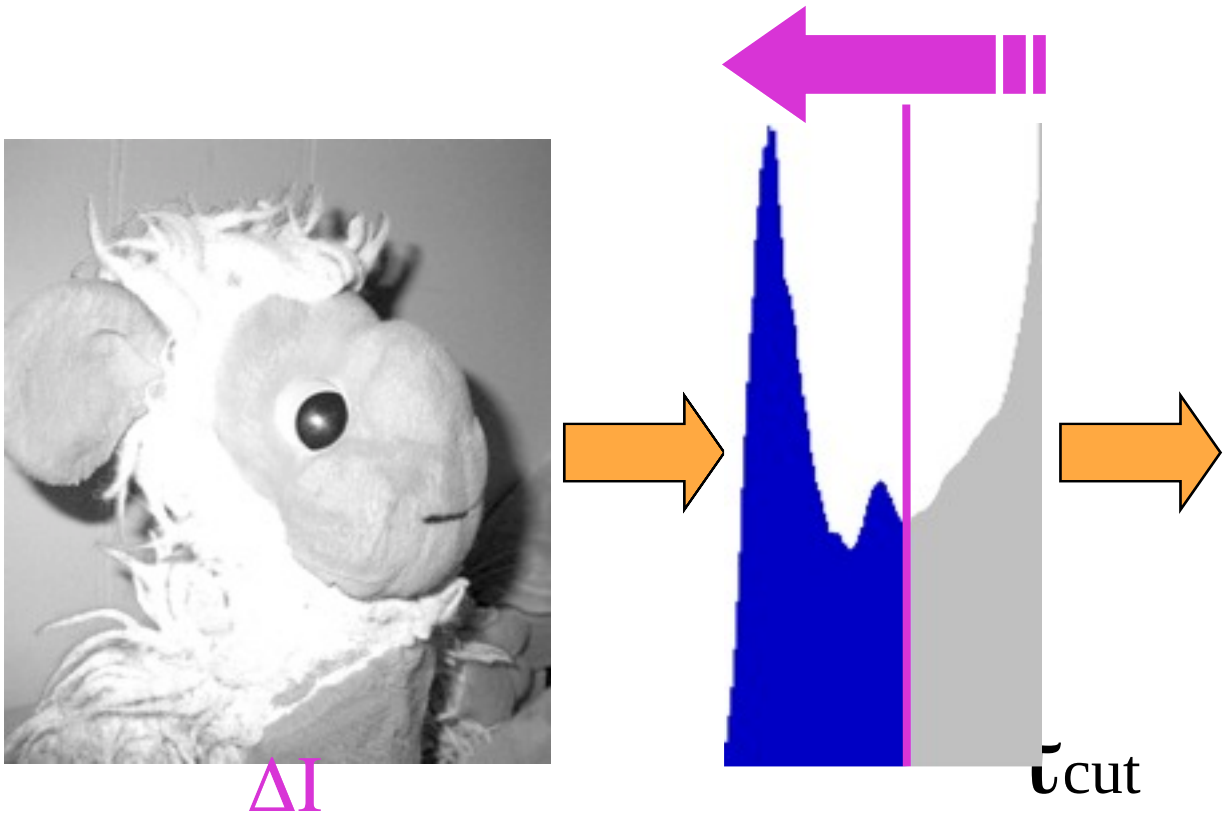

Shadow Detection

Umbra detection

- **Umbra Detection Threshold (details in paper)**

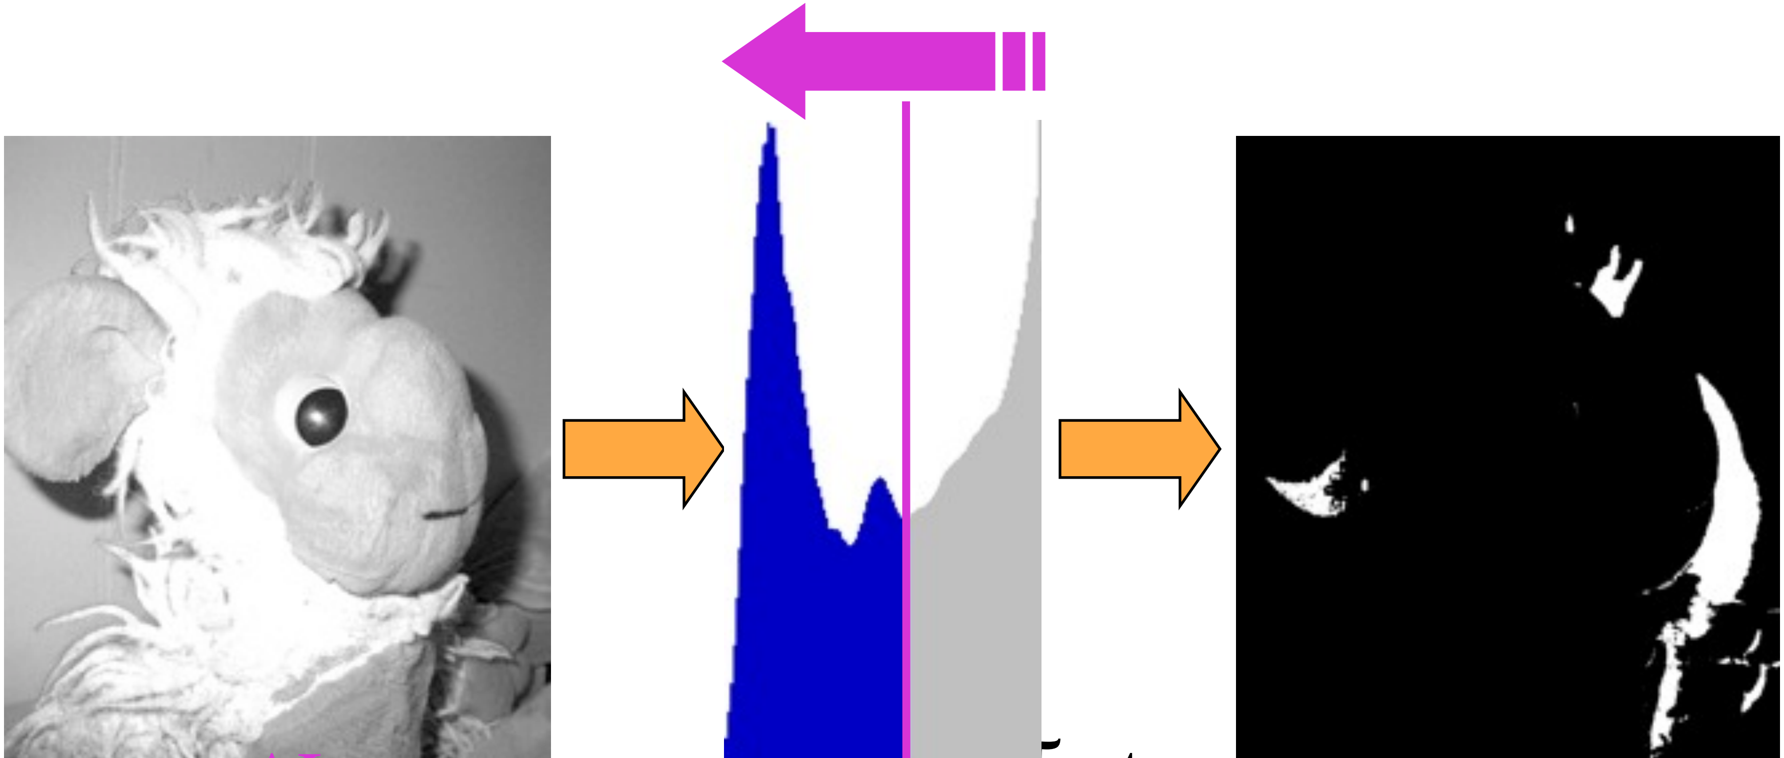

ΔI τ_{cut}

Shadow Detection

Umbra detection

- **Umbra Detection Threshold (details in paper)**

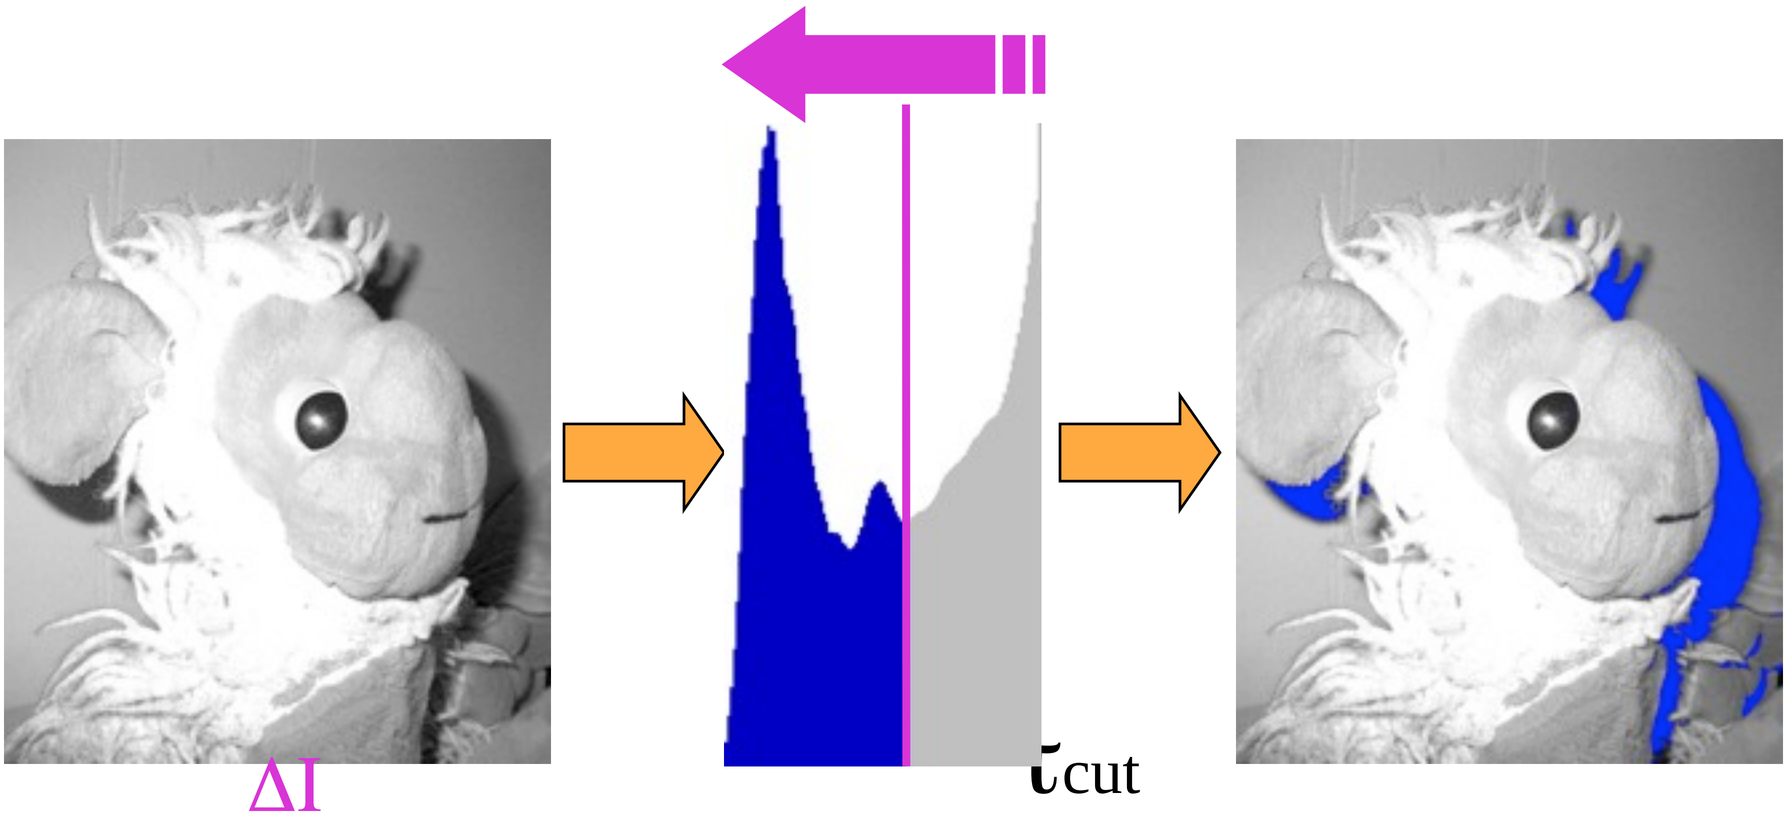

Shadow Detection

Umbra detection

- **Umbra Detection Threshold (details in paper)**

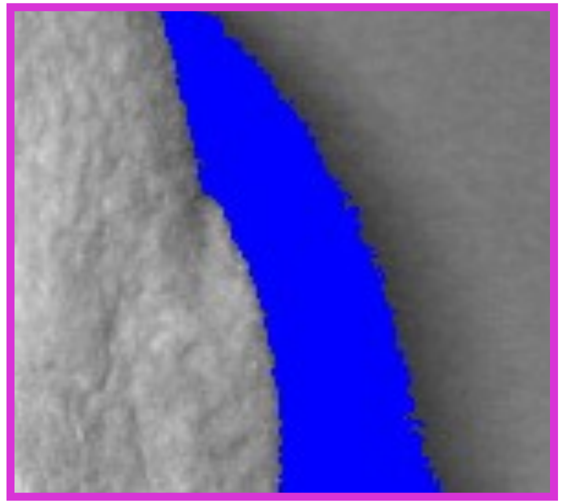

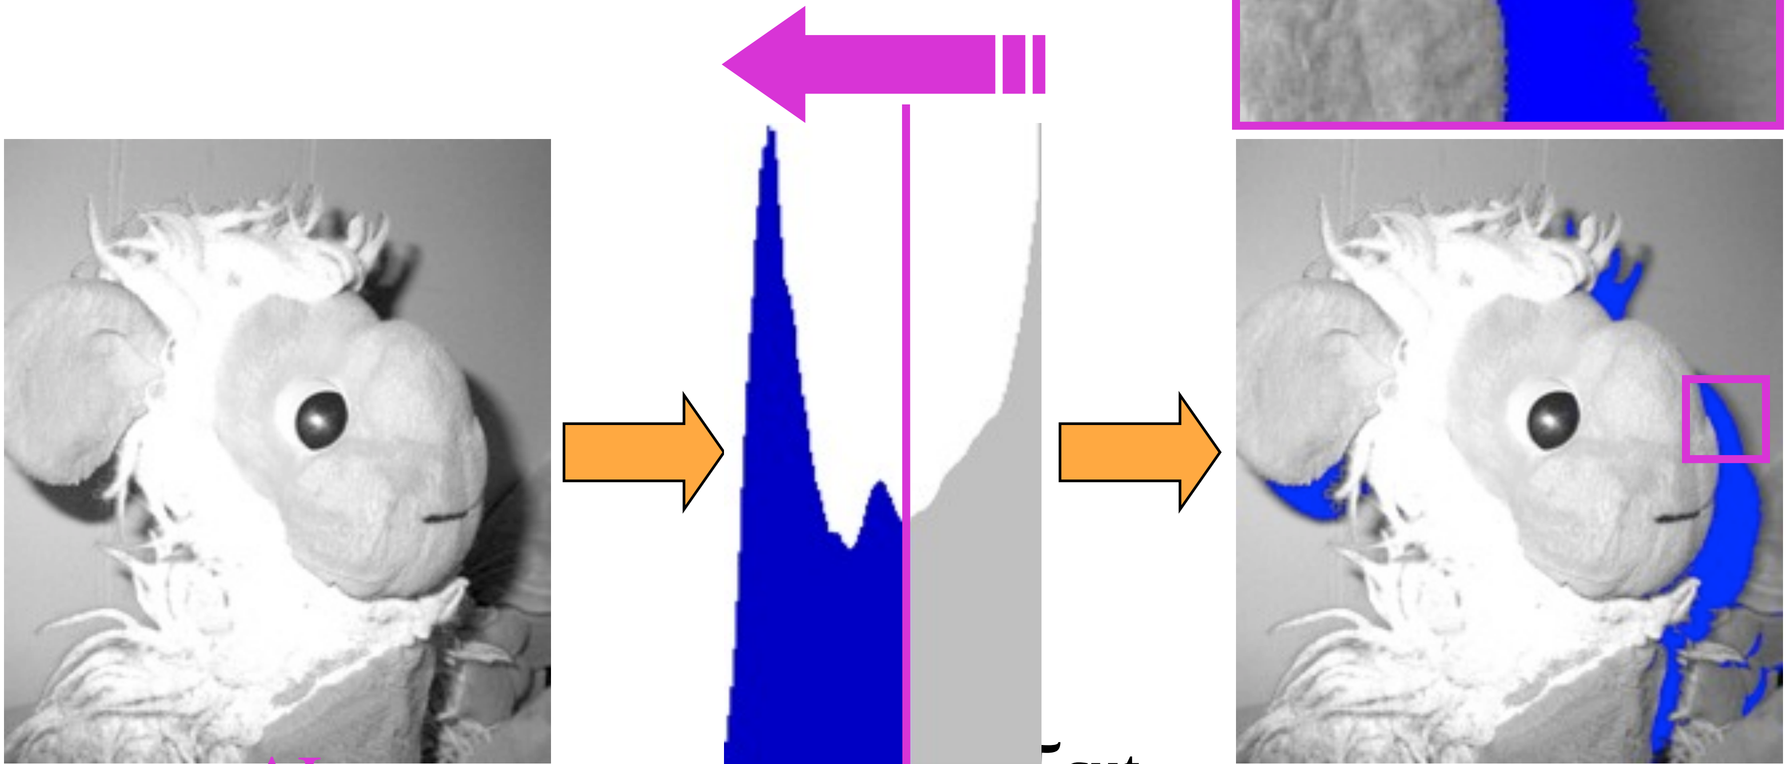

Shadow Detection

Umbra detection

- Umbra Detection Threshold
  (details in paper)

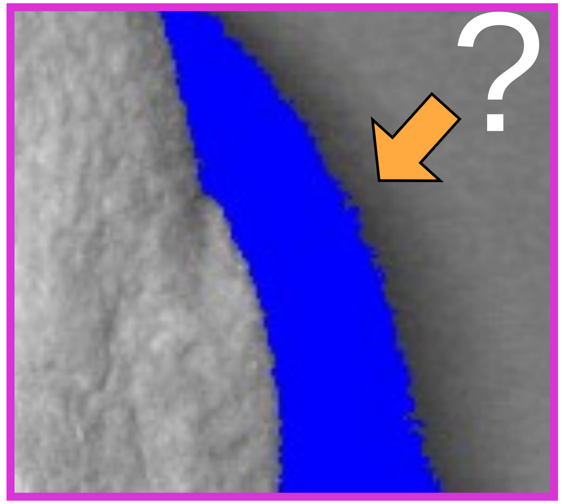

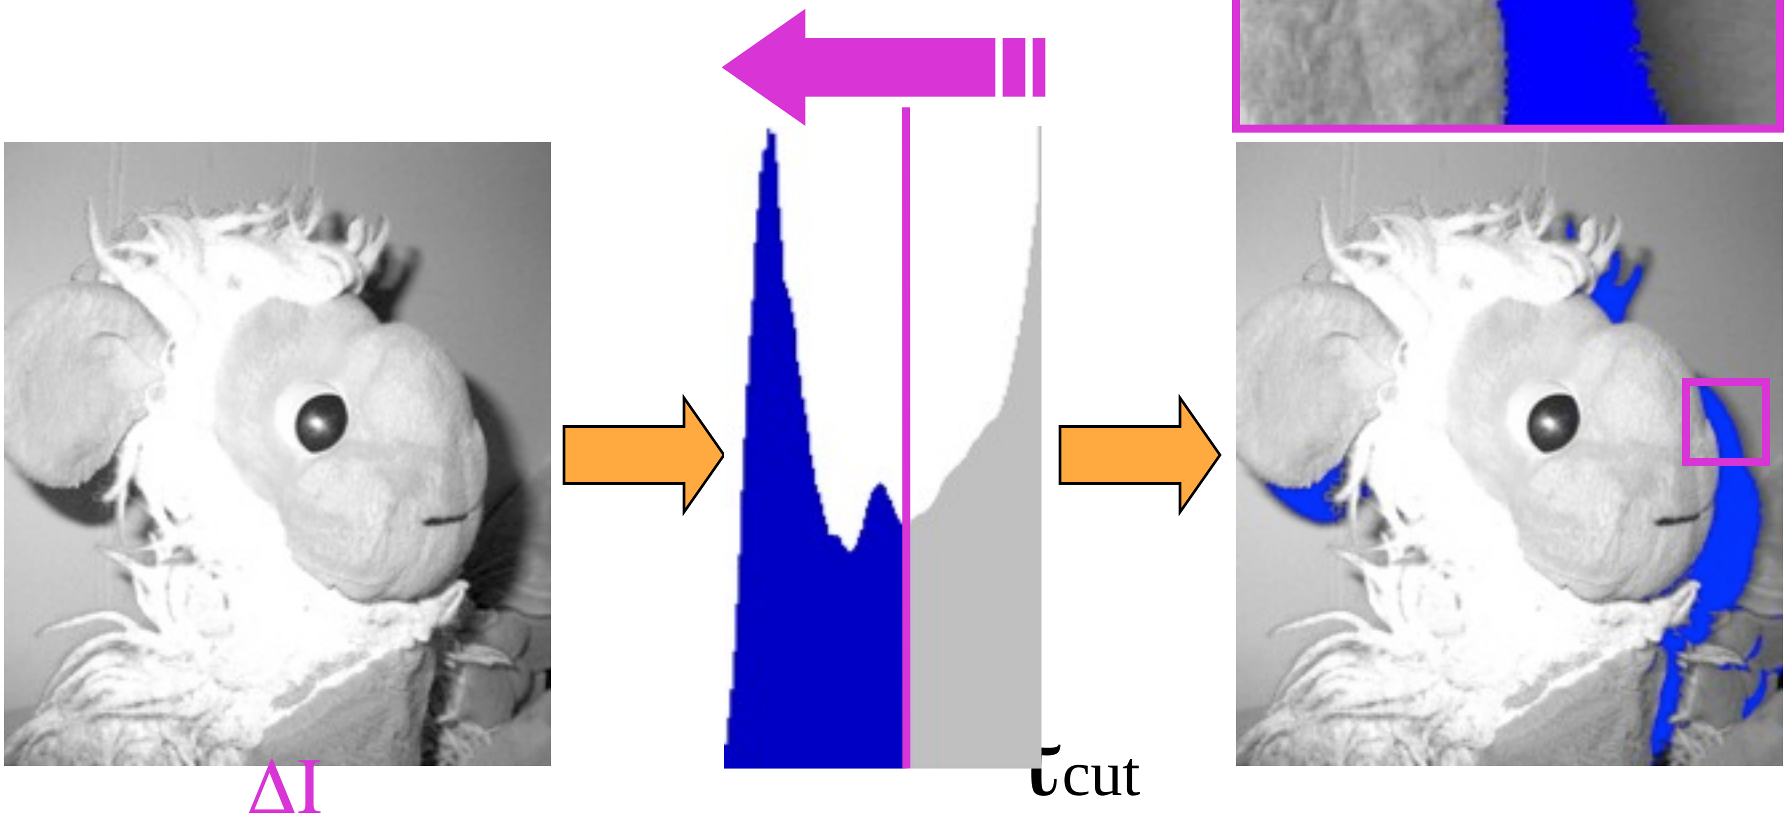

Shadow Detection

Shadow Detection

No-flash

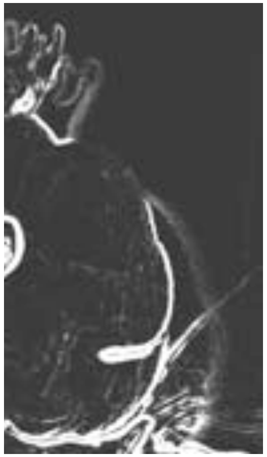

flash

Shadow Detection

Penumbra detection

- ➢ strong gradient at boundary
- no strong gradient in no-flash image
- connected to umbra

No-flash

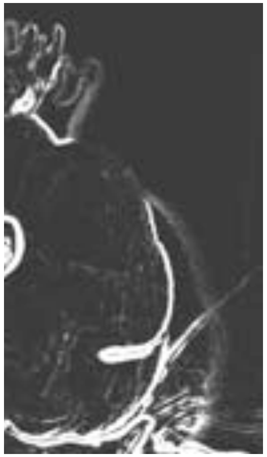

flash

Shadow Detection

Penumbra detection

- strong gradient at boundary
- no strong gradient in no-flash image
- connected to umbra

No-flash

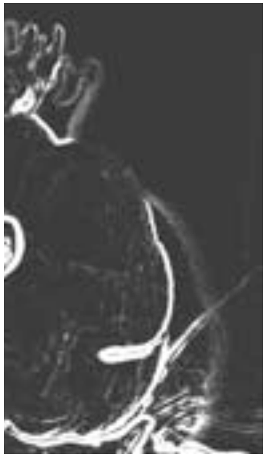

flash

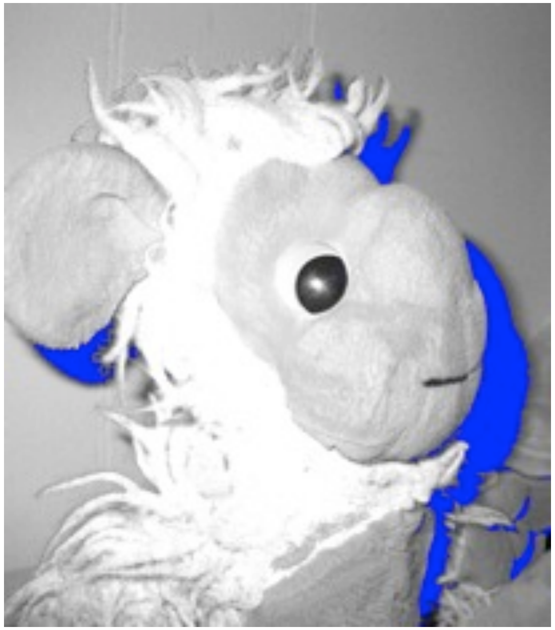

Umbra

Shadow Detection

Penumbra detection

- strong gradient at boundary
- no strong gradient in no-flash image
- connected to umbra

No-flash

flash

Umbra

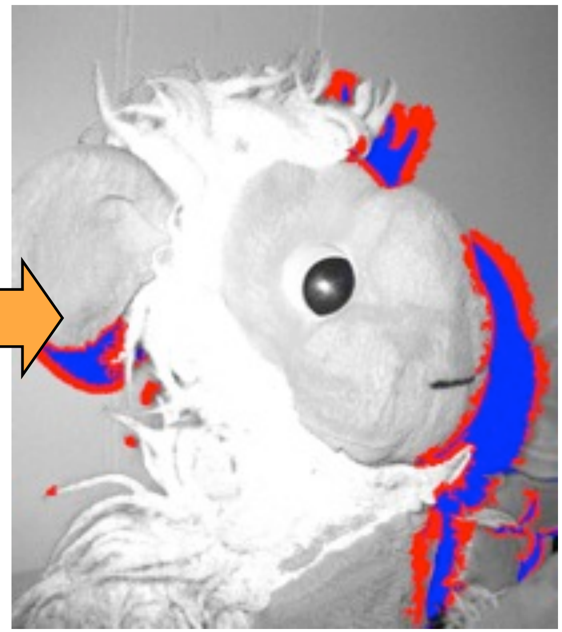

Penumbra

Shadow Correction

- **Need a robust correction**
- **Correct color and detail**

Binary shadow mask

Shadow Color Correction

Flash color

Shadow Color Correction

Flash color

Shadow Color Correction

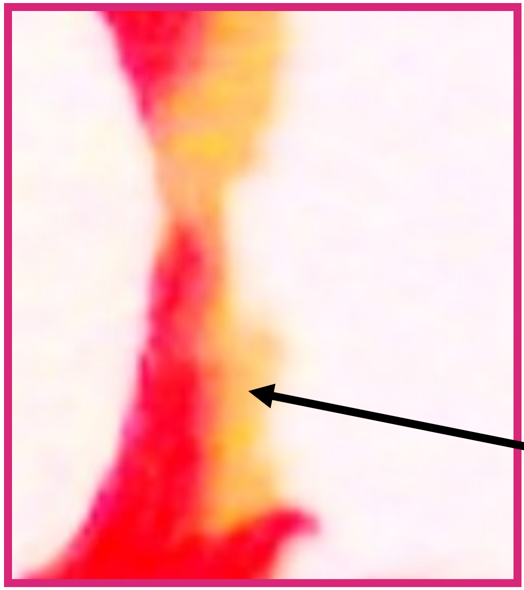

Wrong color

Shadow Color Correction

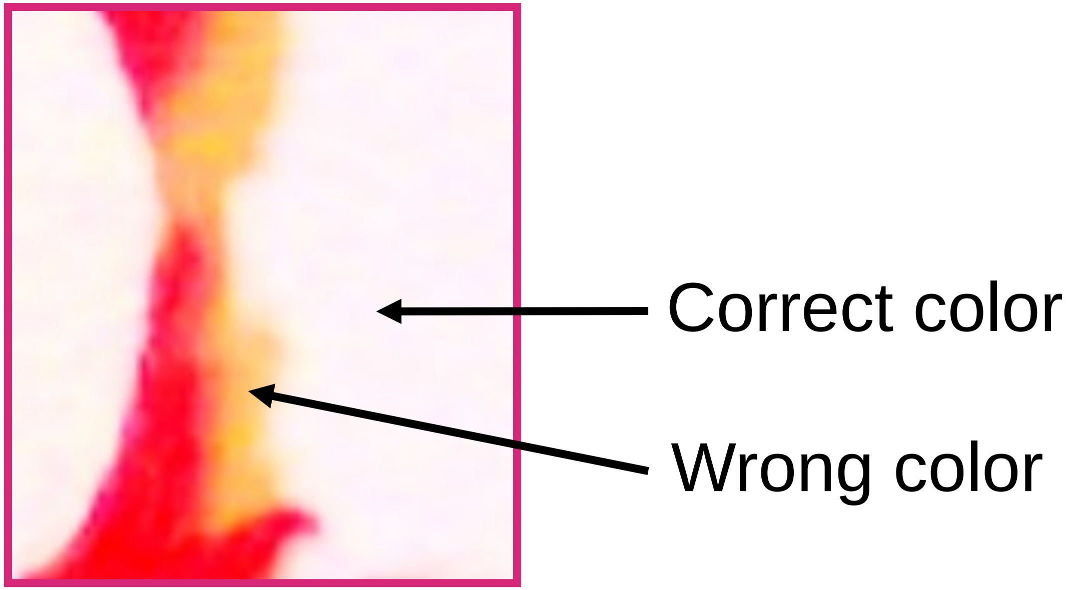

Shadow Color Correction

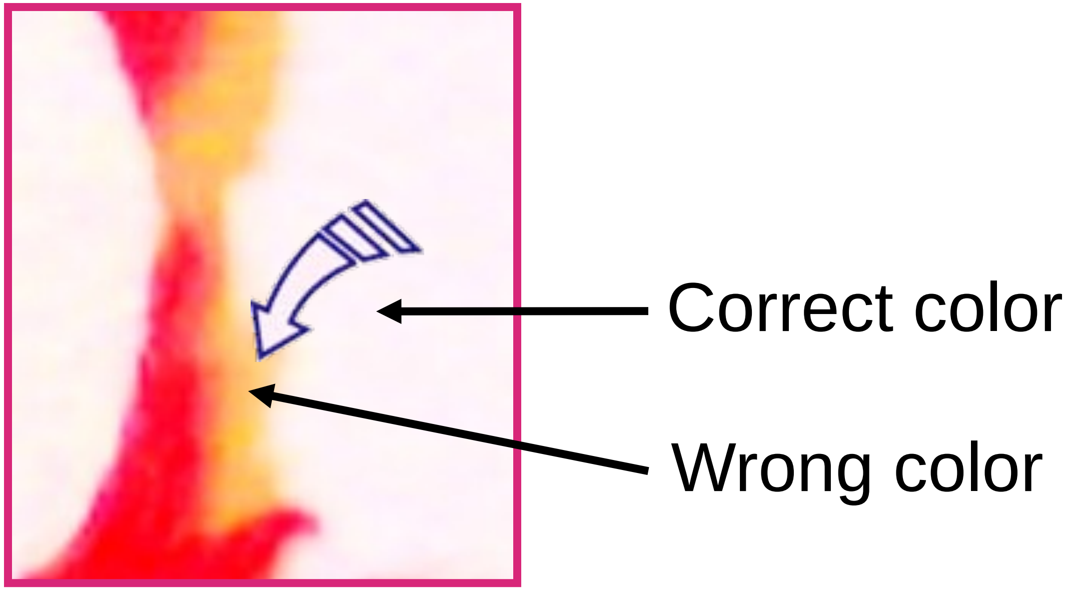

Fill in shadow from similar surrounding

Shadow Color Correction

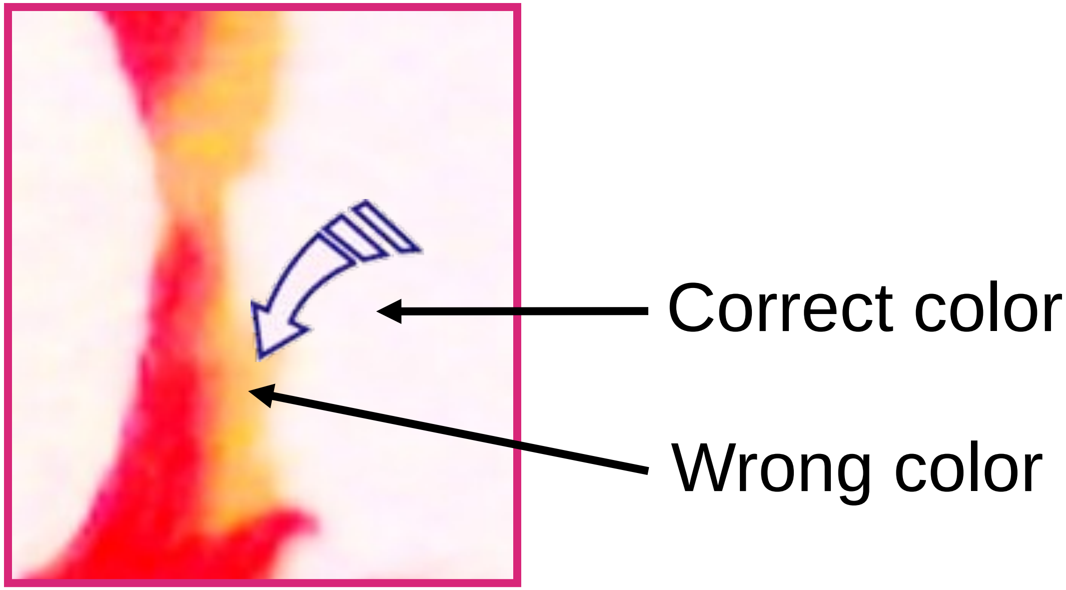

Fill in shadow from similar surrounding

Shadow Color Correction

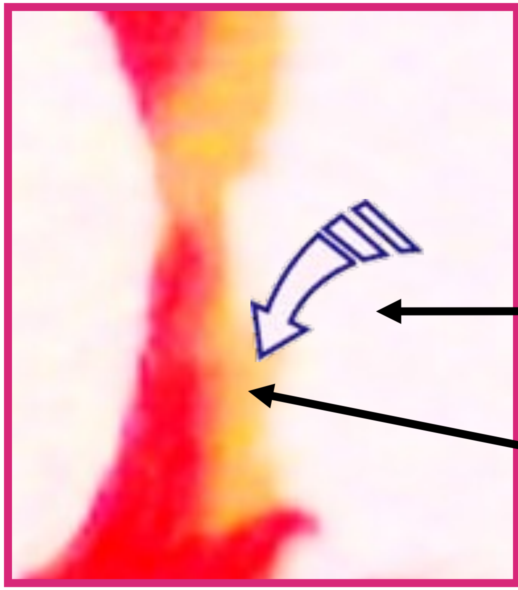

Correct color

Wrong color

Fill in shadow from similar surrounding

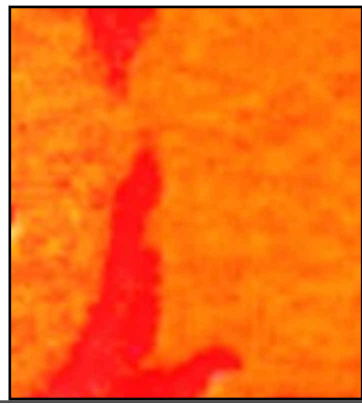

No-flash colors

Shadow Color Correction

Shadow Color Correction

Shadow Color Correction

No-flash

Flash

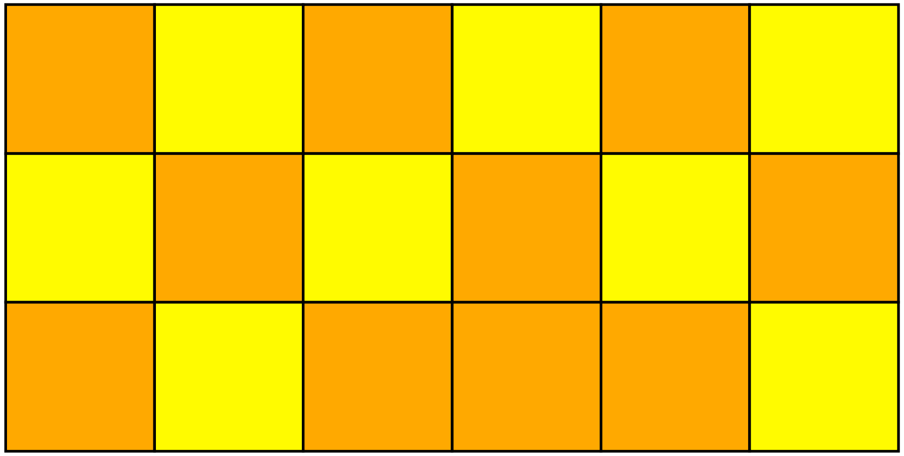

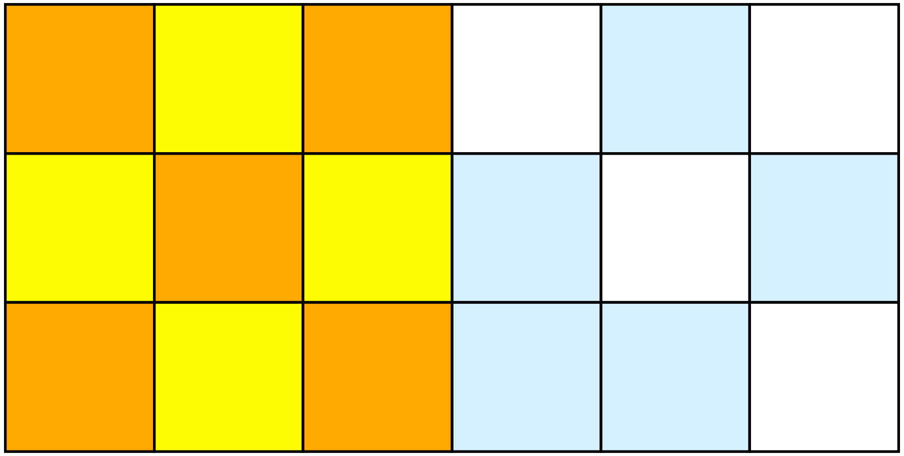

Shadow Color Correction

No-flash

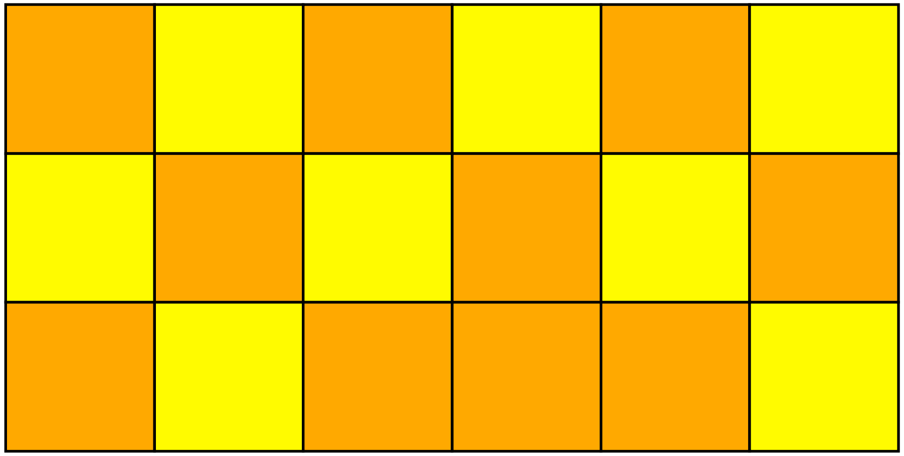

Flash

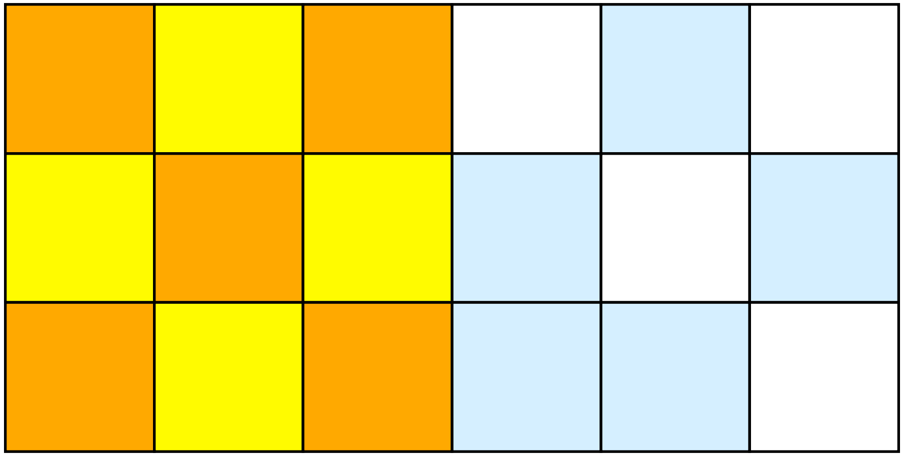

Shadow Color Correction

No-flash

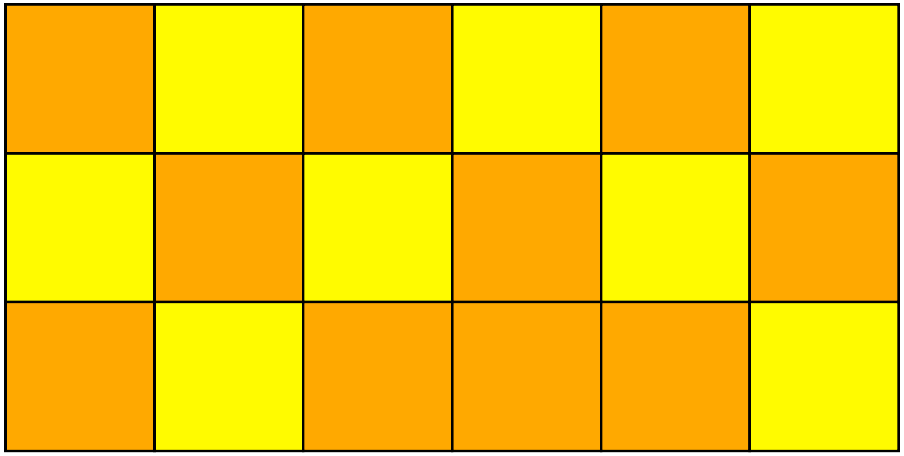

Flash

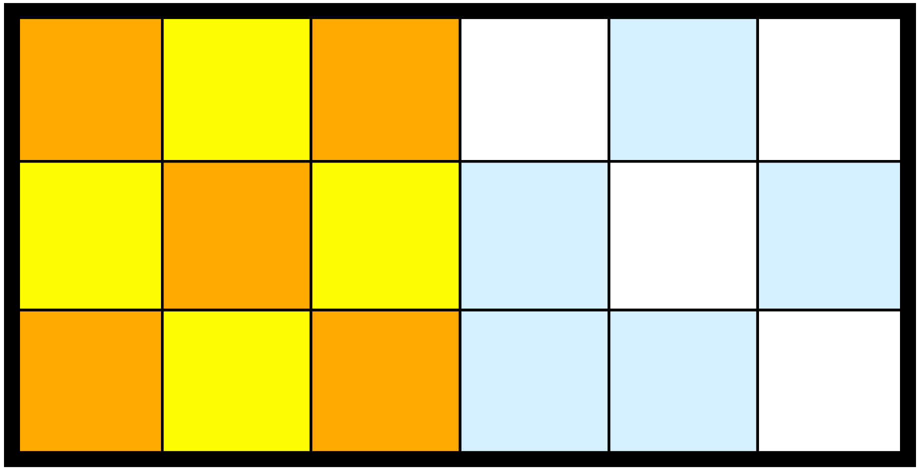

Shadow Color Correction

No-flash

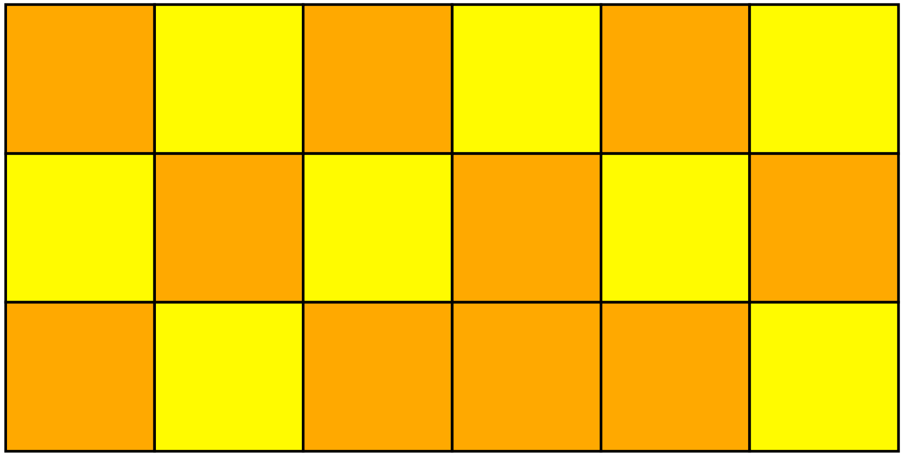

Flash

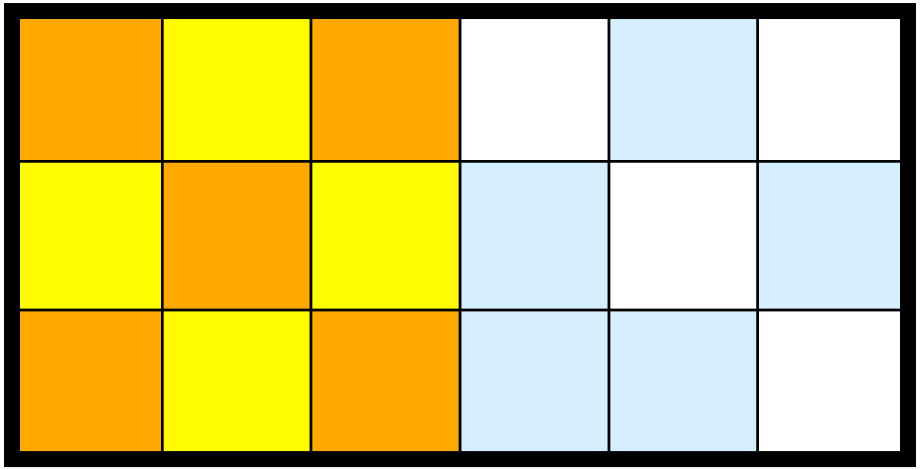

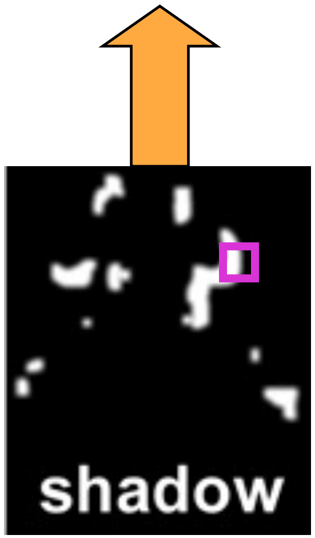

Shadow Color Correction

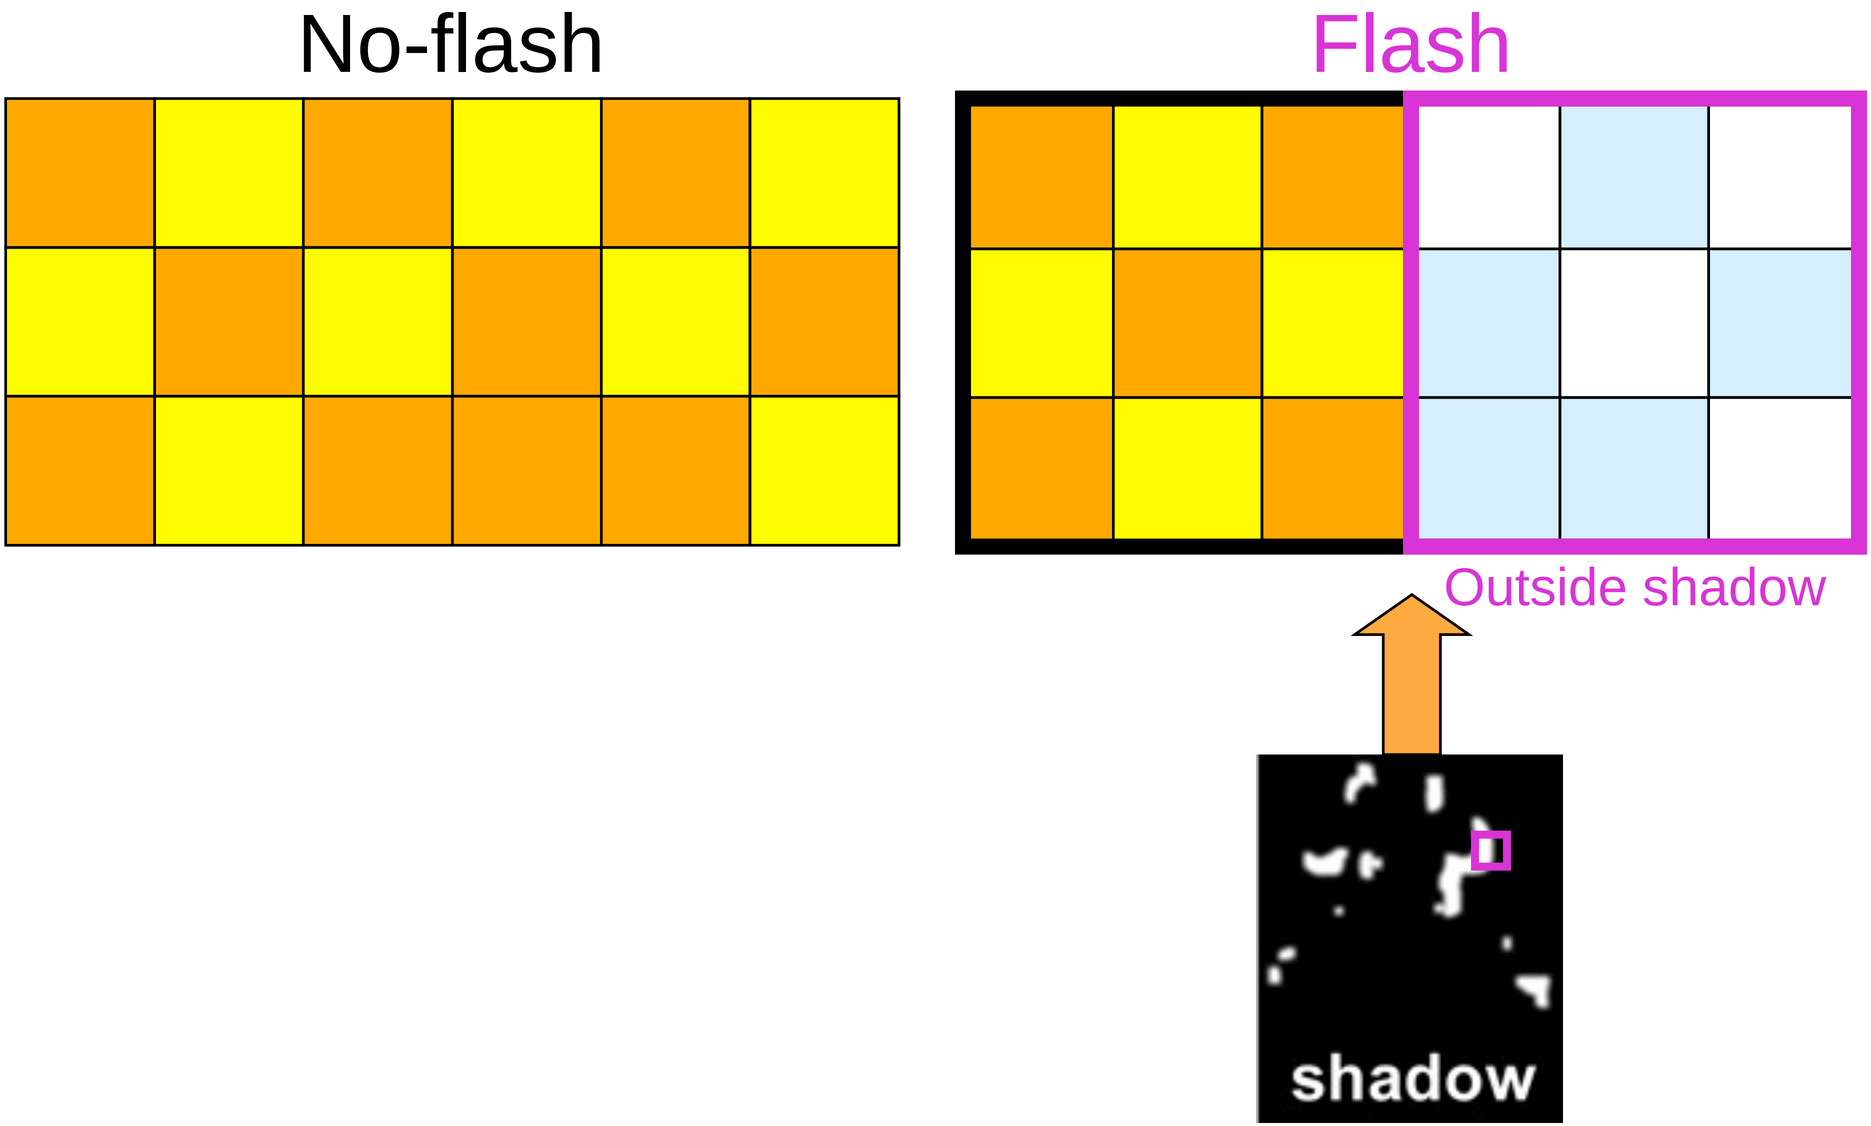

Shadow Color Correction

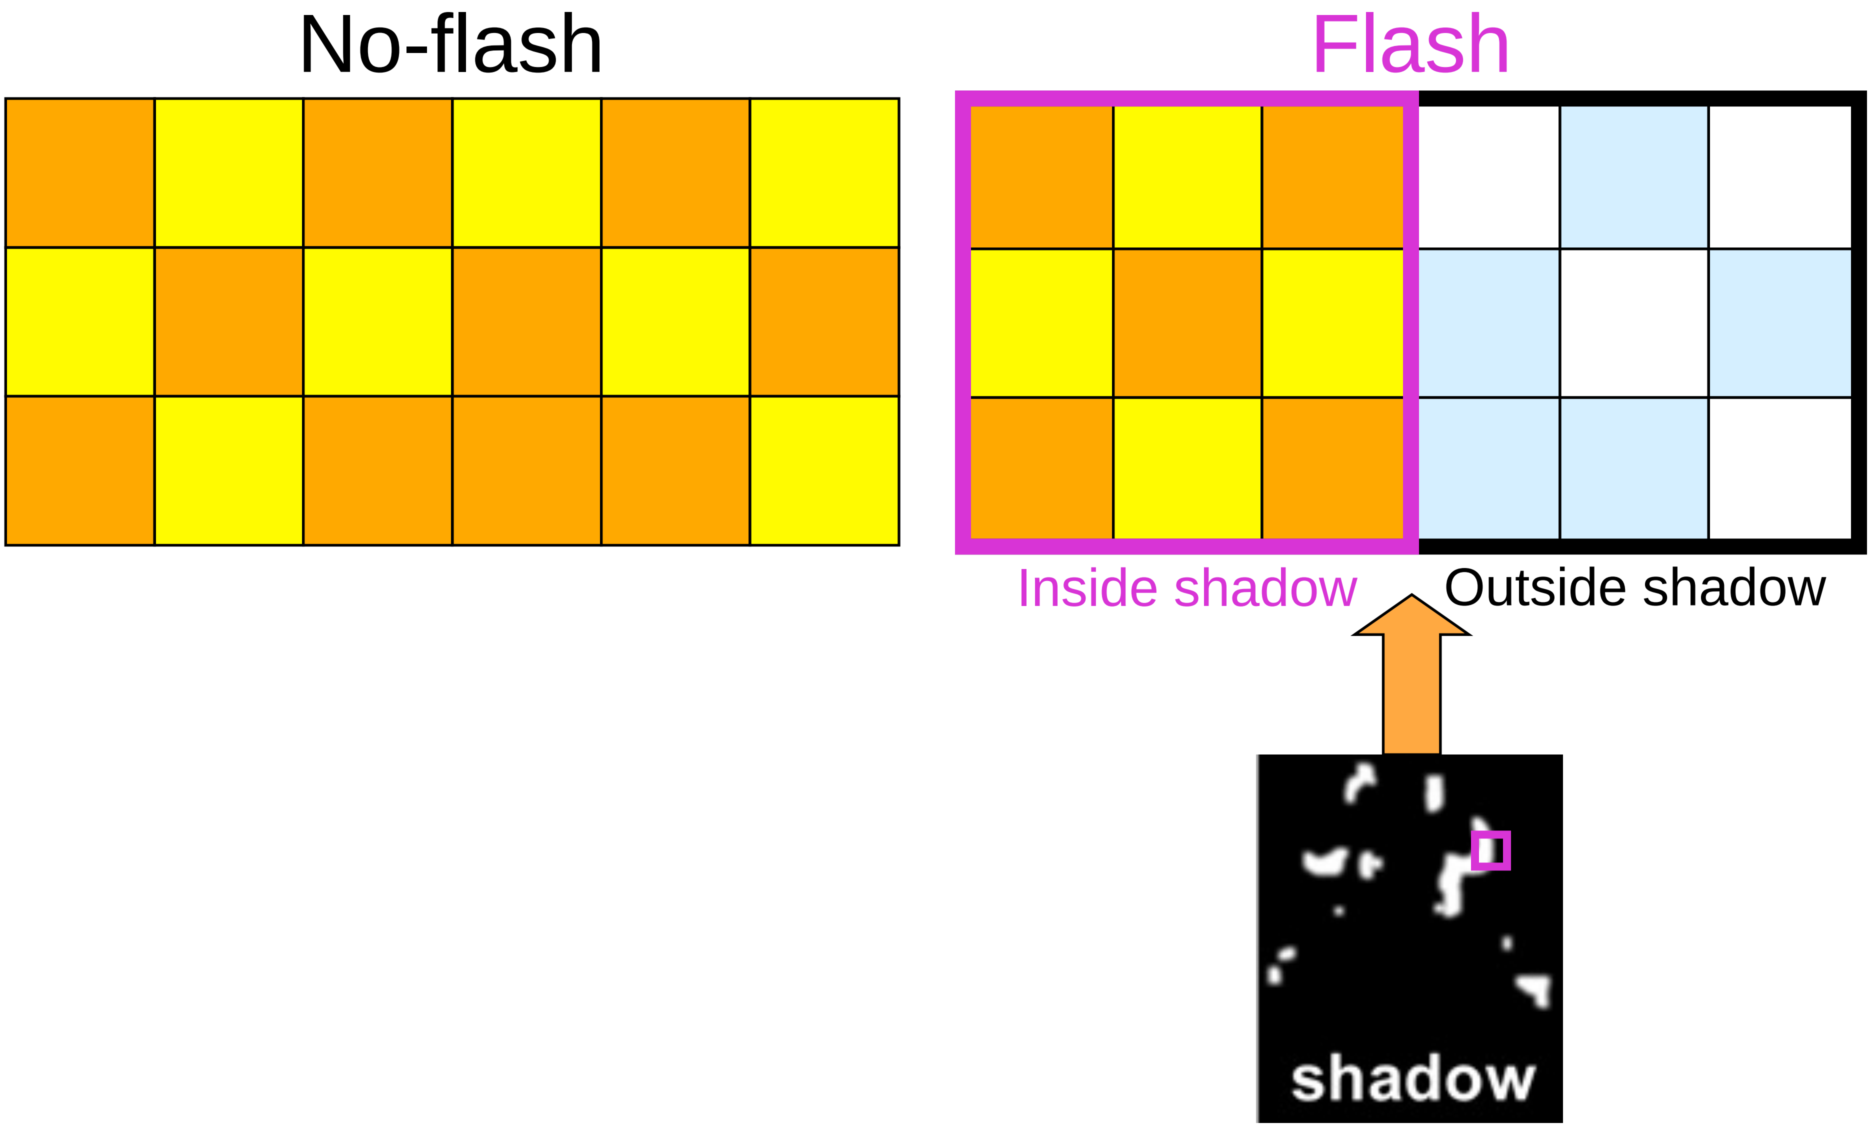

Shadow Color Correction

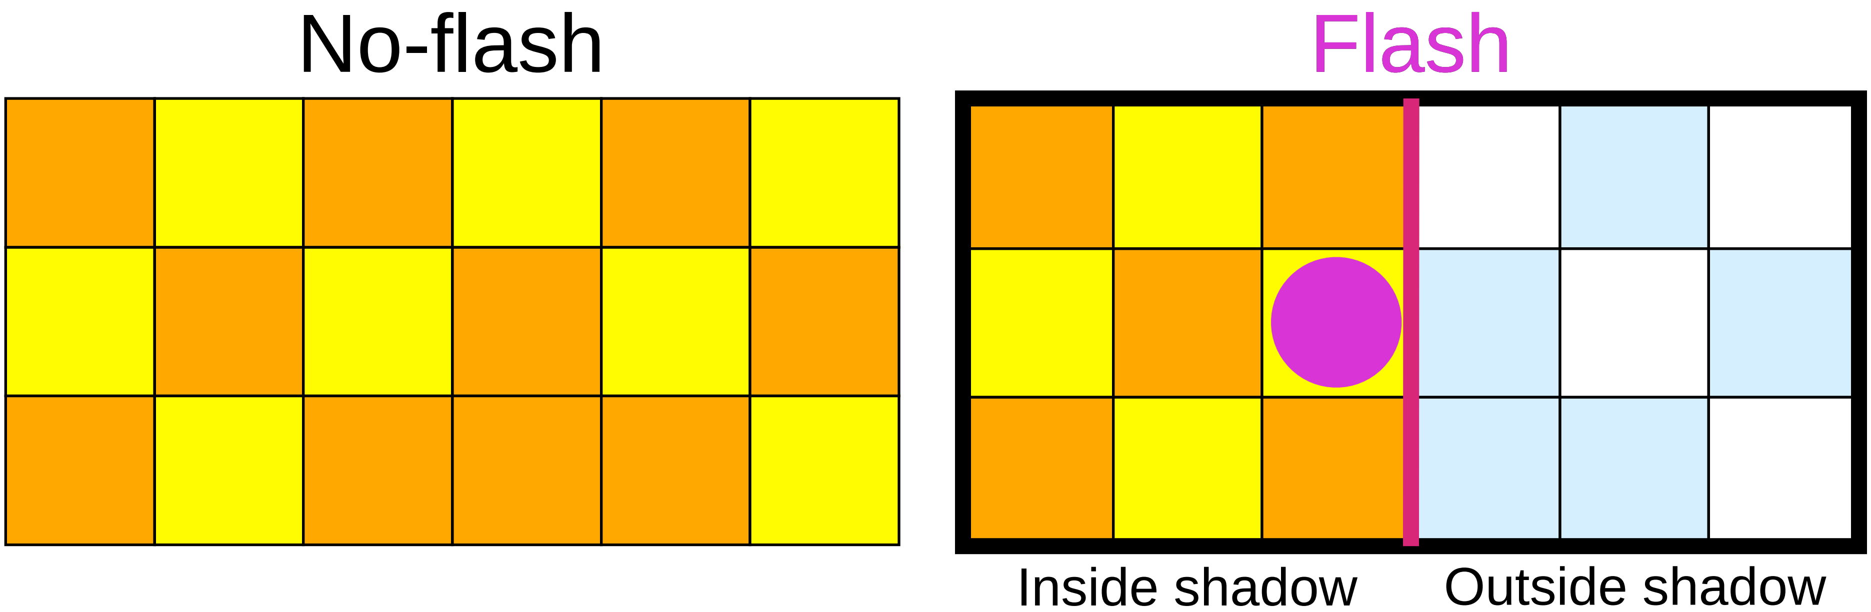

Select pixel in shadow

Shadow Color Correction

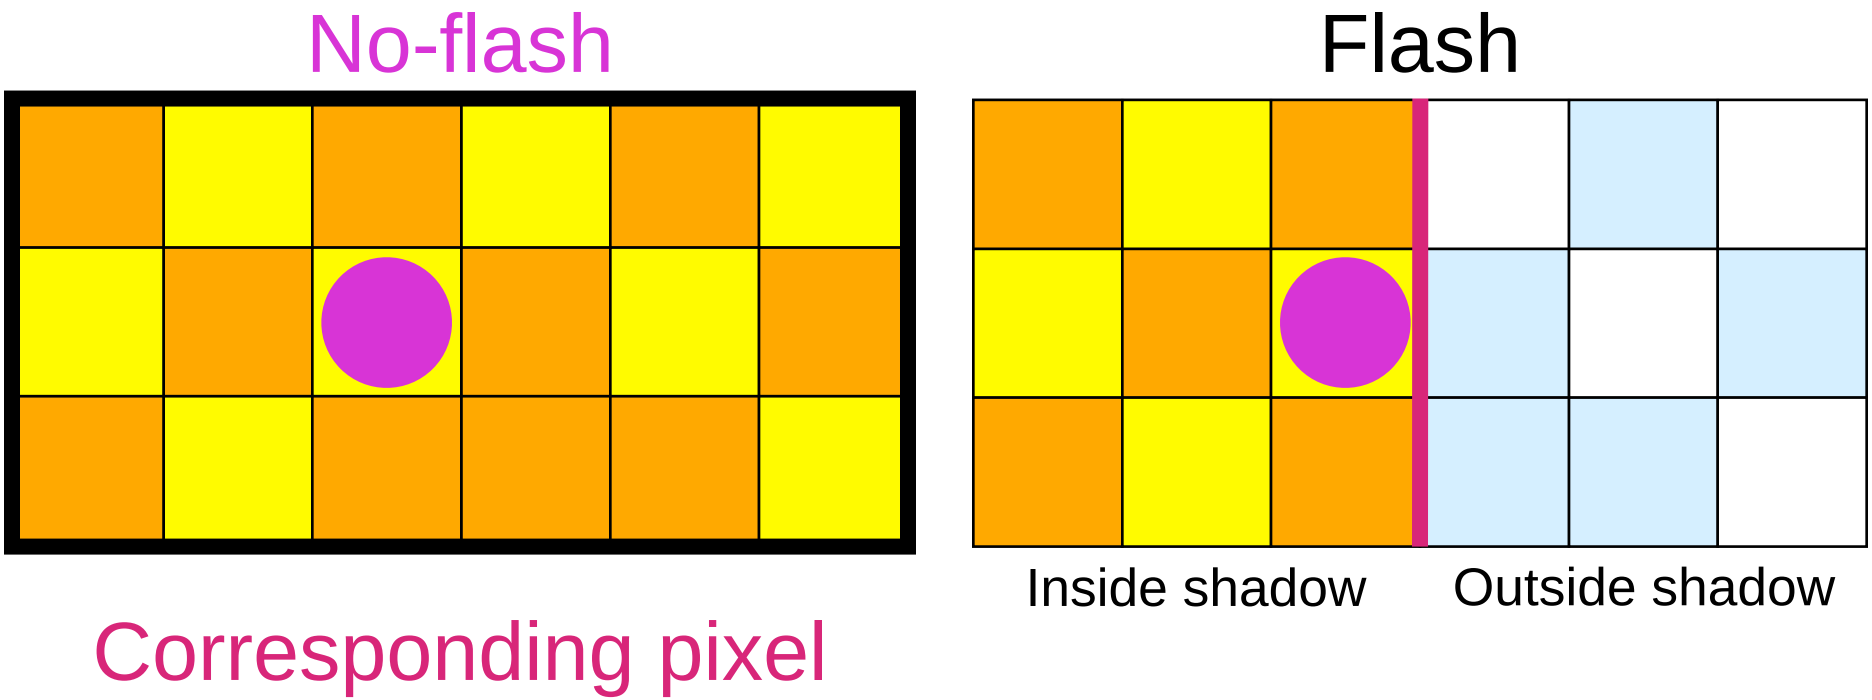

Shadow Color Correction

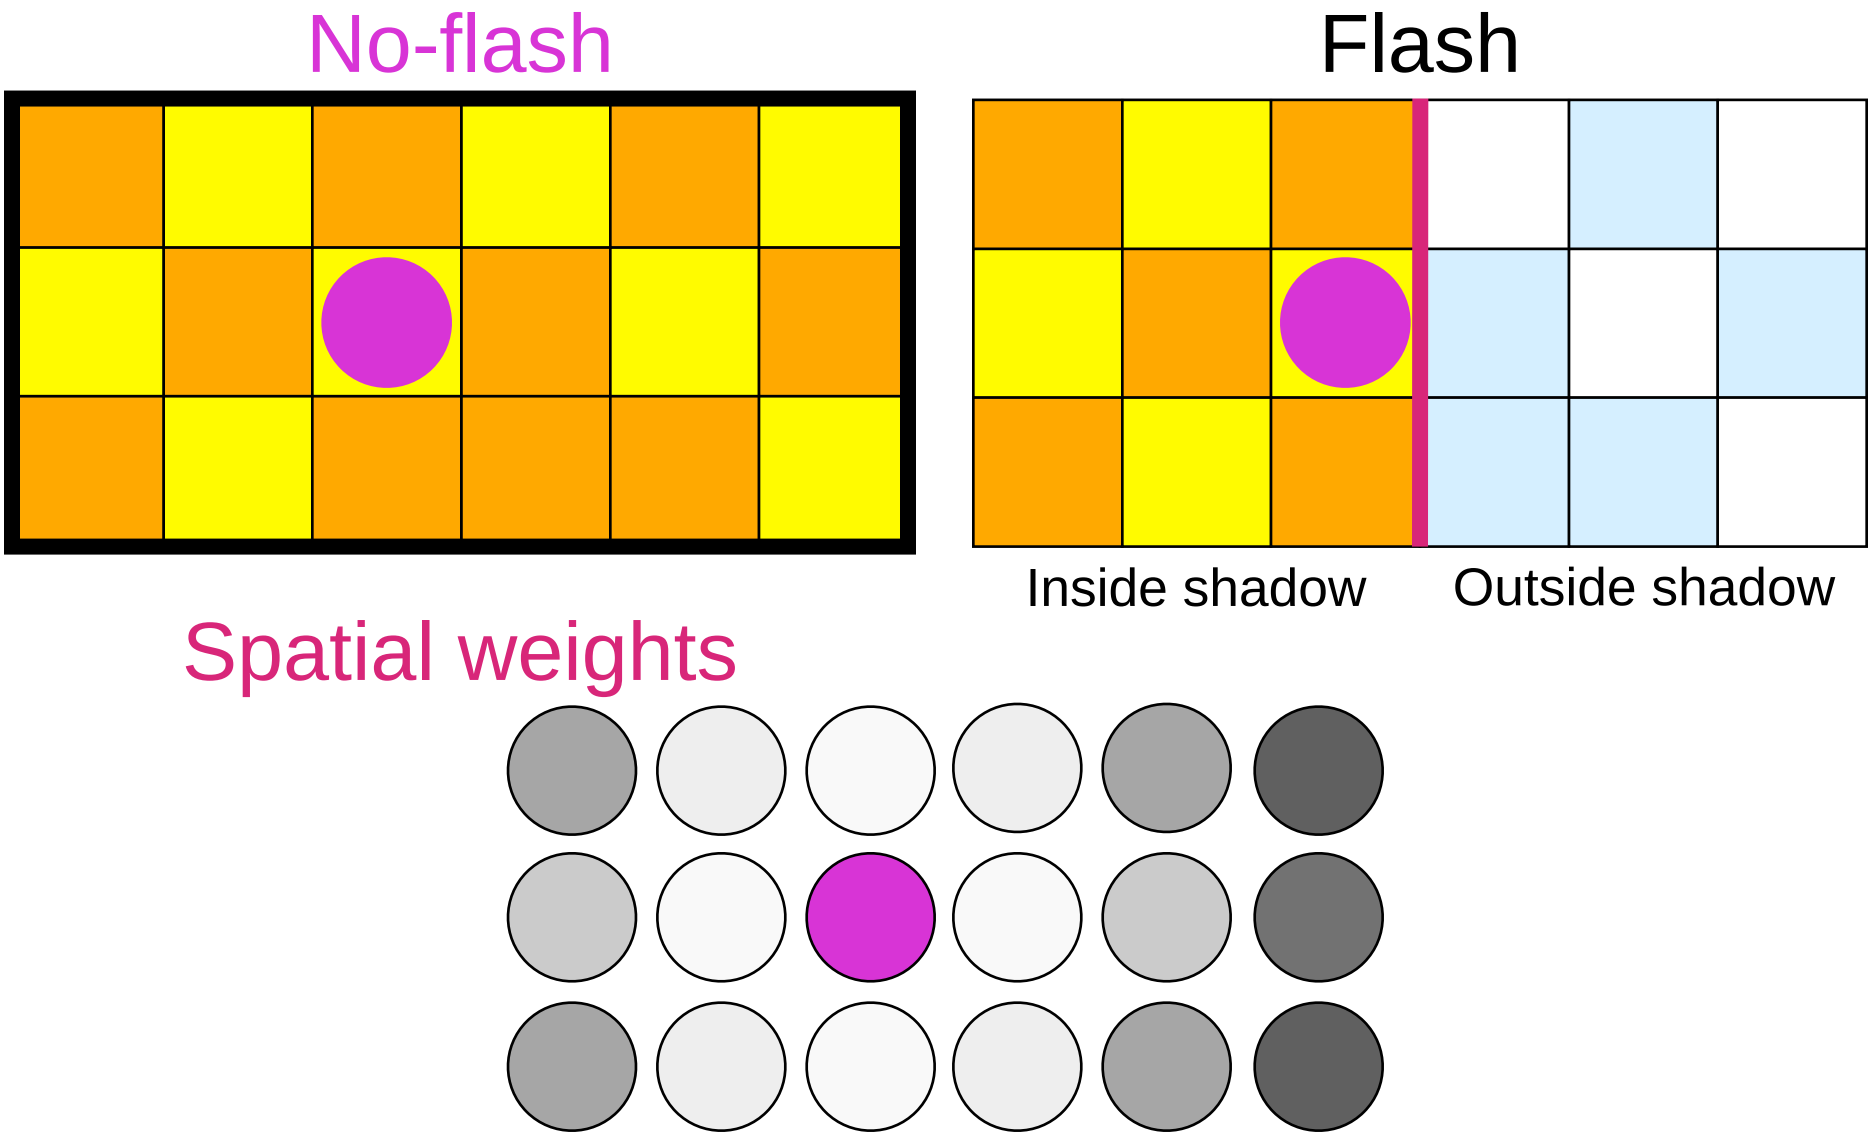

Shadow Color Correction

No-flash

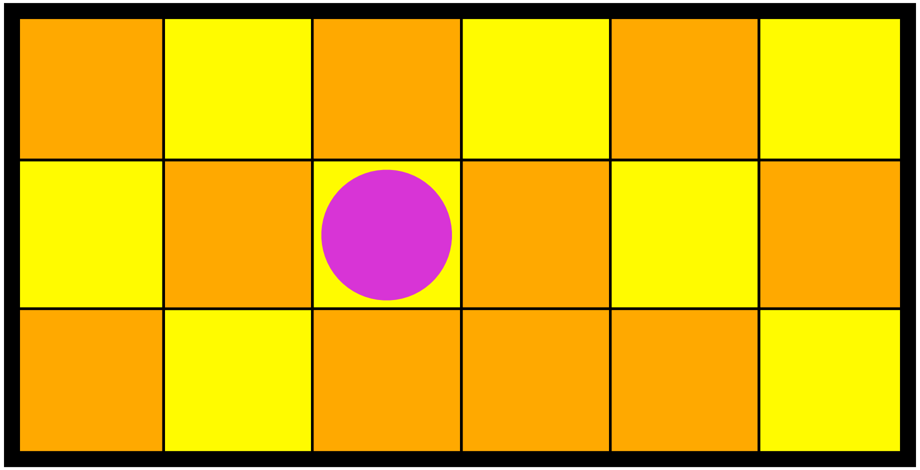

Flash

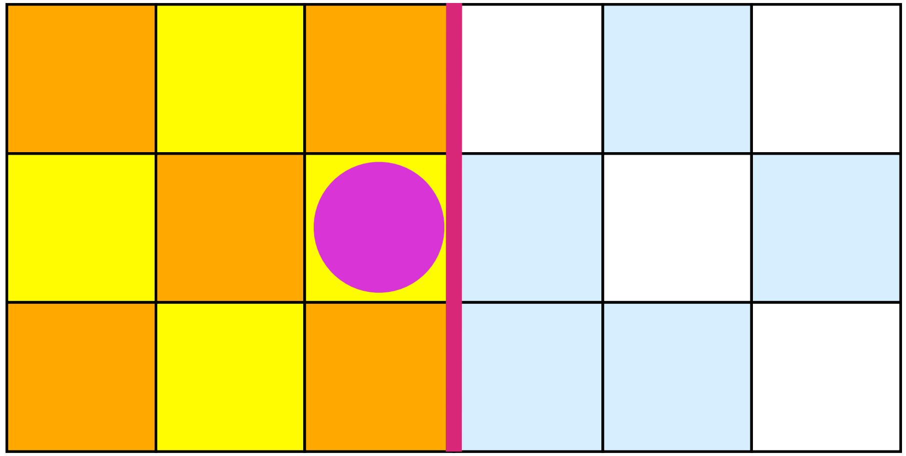

Inside shadow Outside shadow

Spatial and Color weights

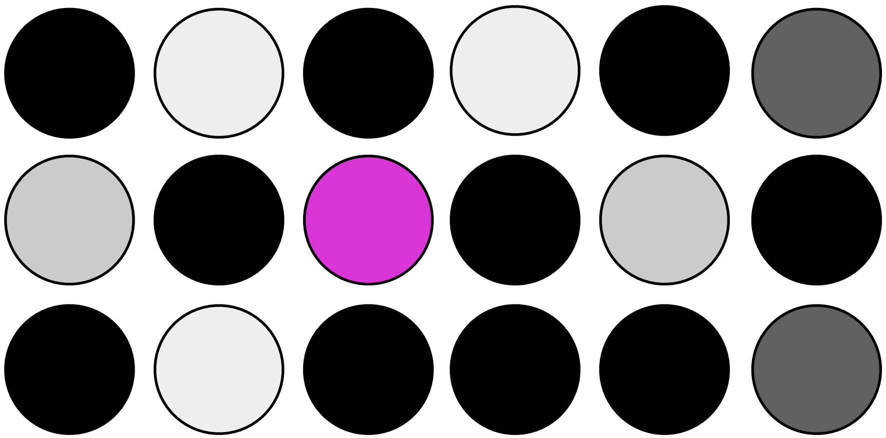

Shadow Color Correction

No-flash

Flash

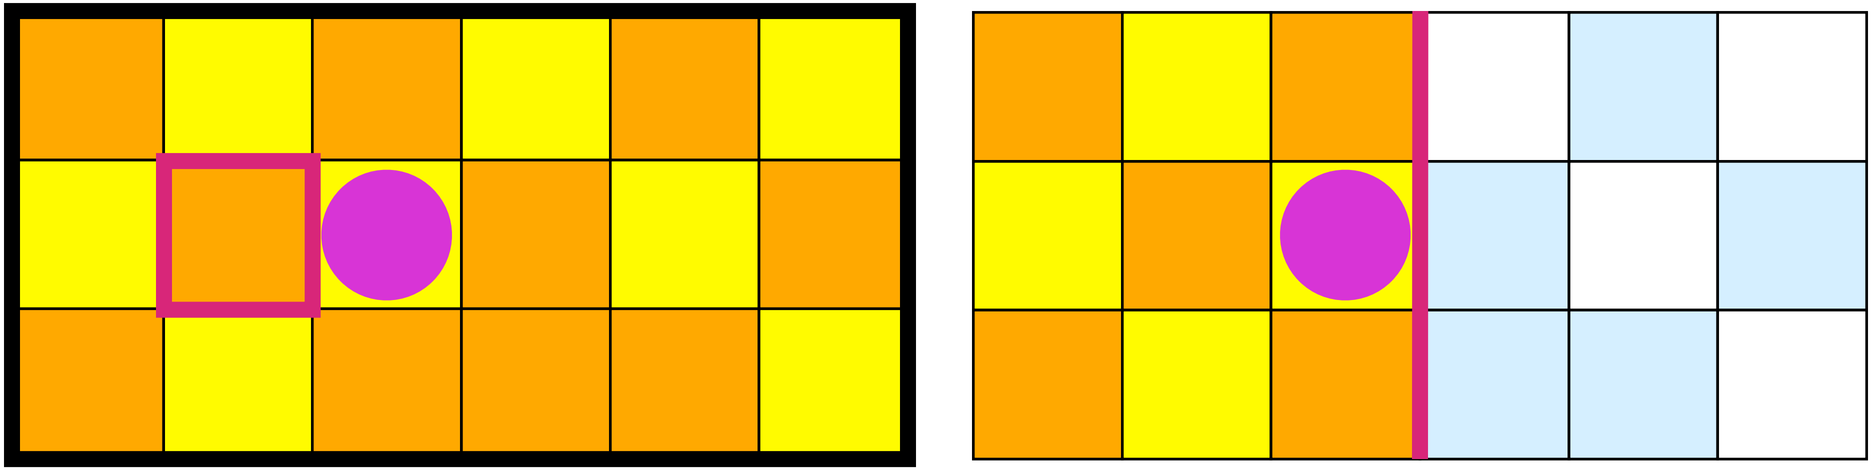

Inside shadow

Outside shadow

Spatial and Color weights

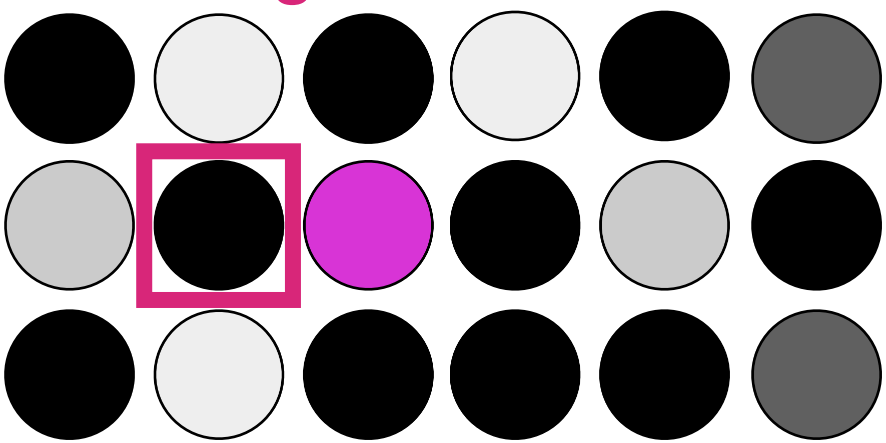

Shadow Color Correction

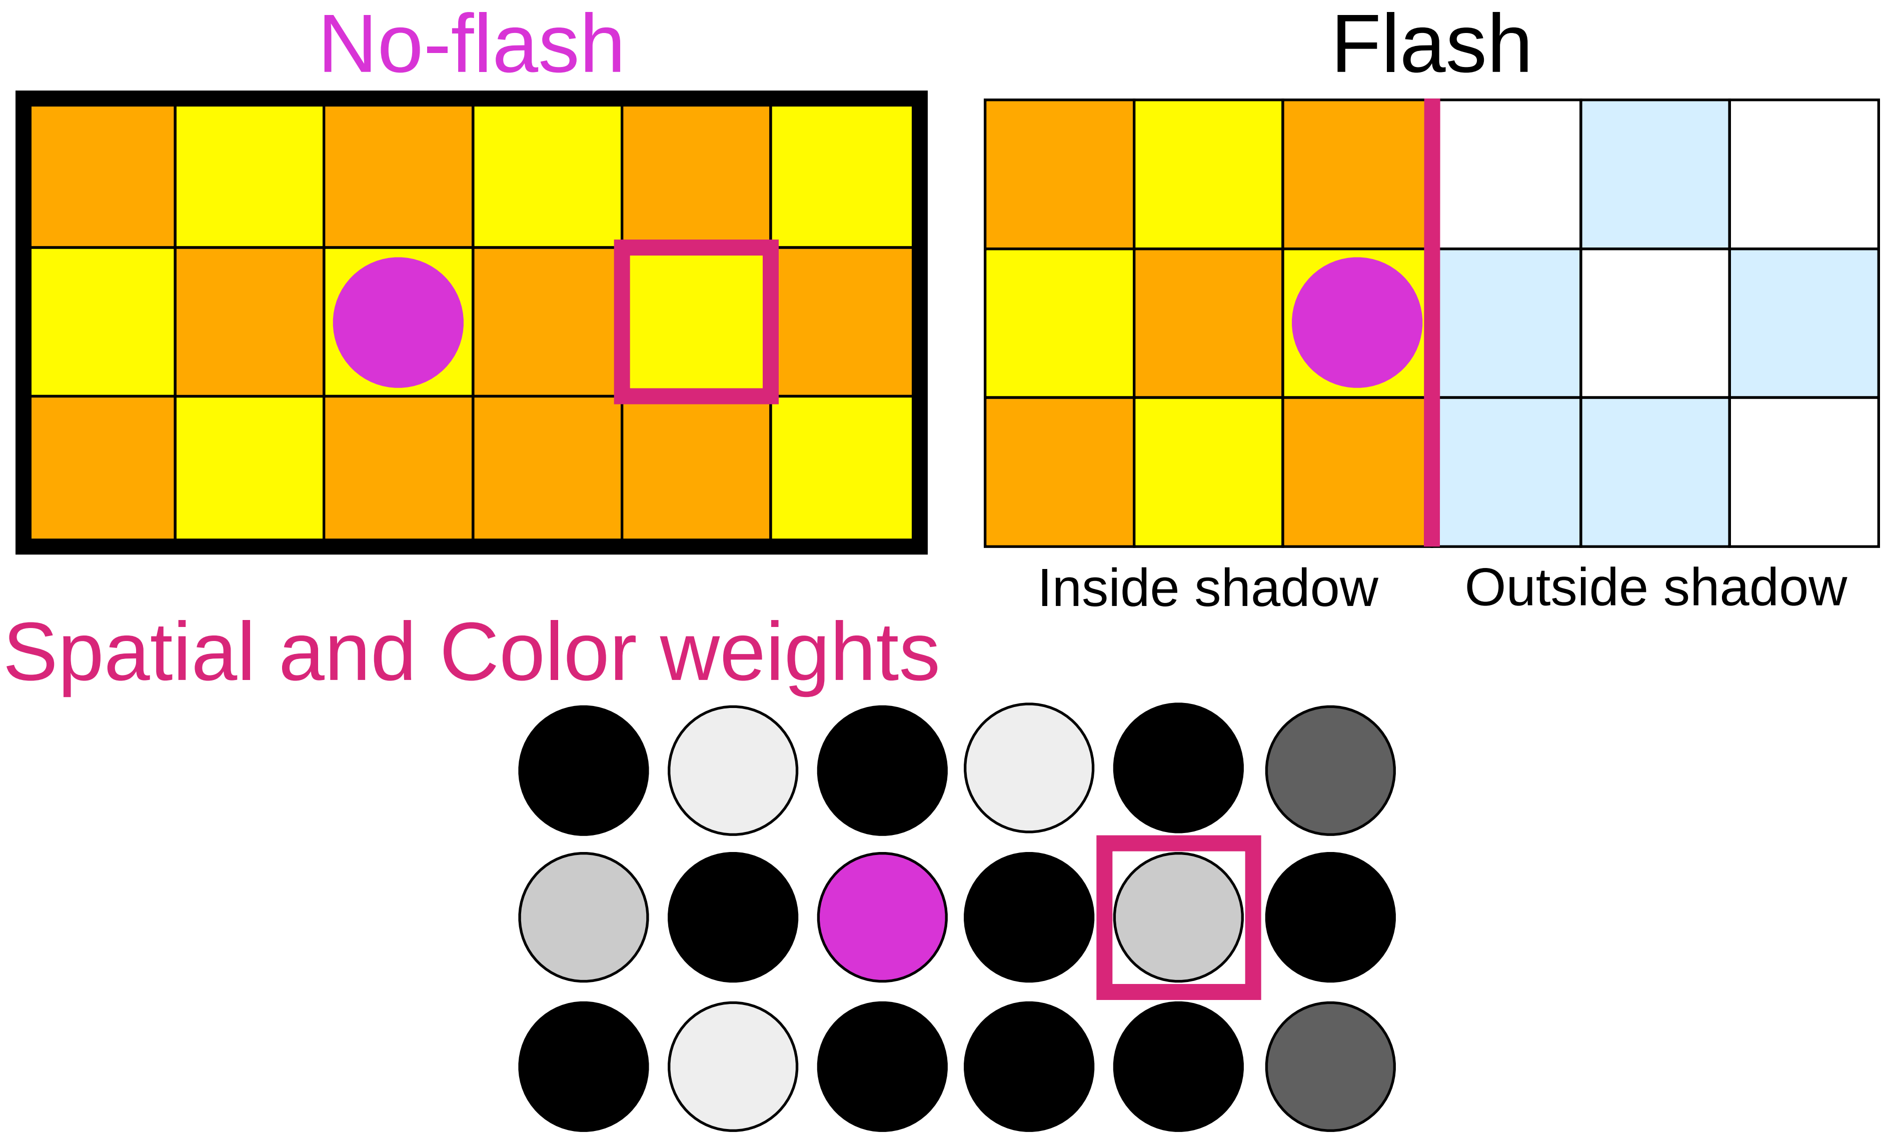

Shadow Color Correction

No-flash

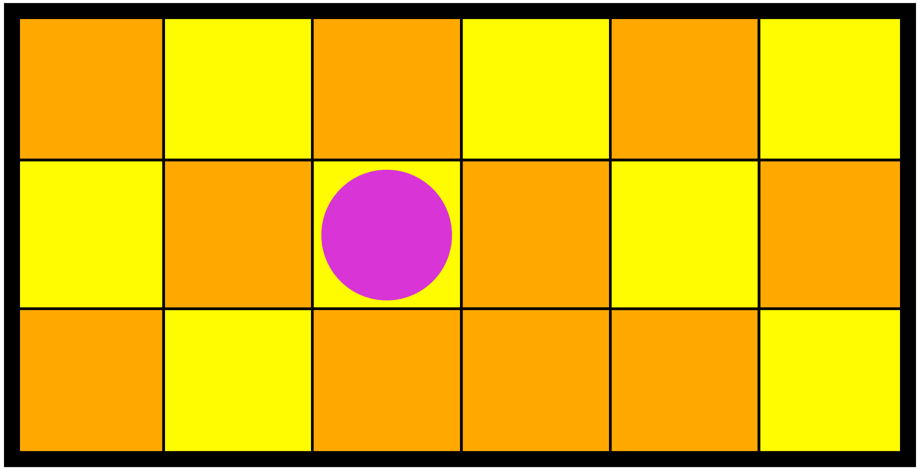

Flash

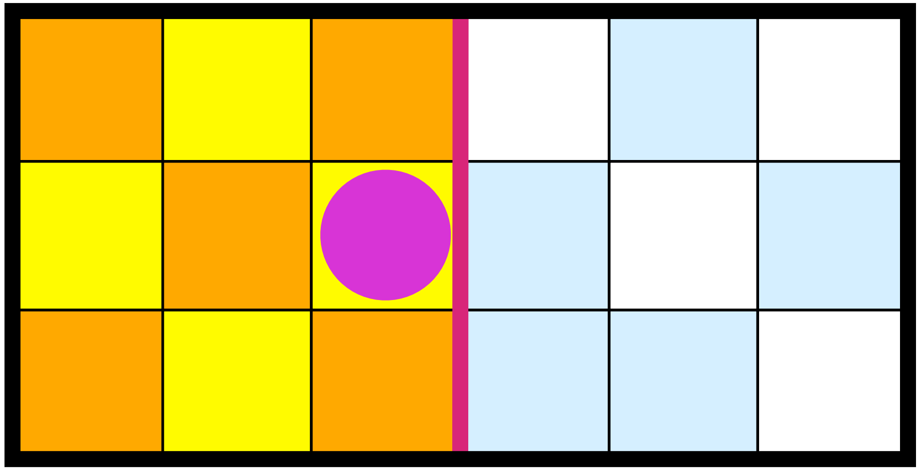

Inside shadow   Outside shadow

Use shadow mask

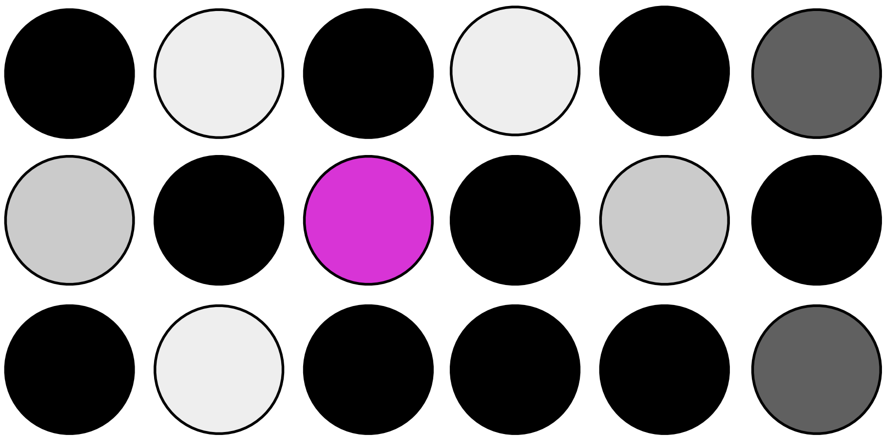

Shadow Color Correction

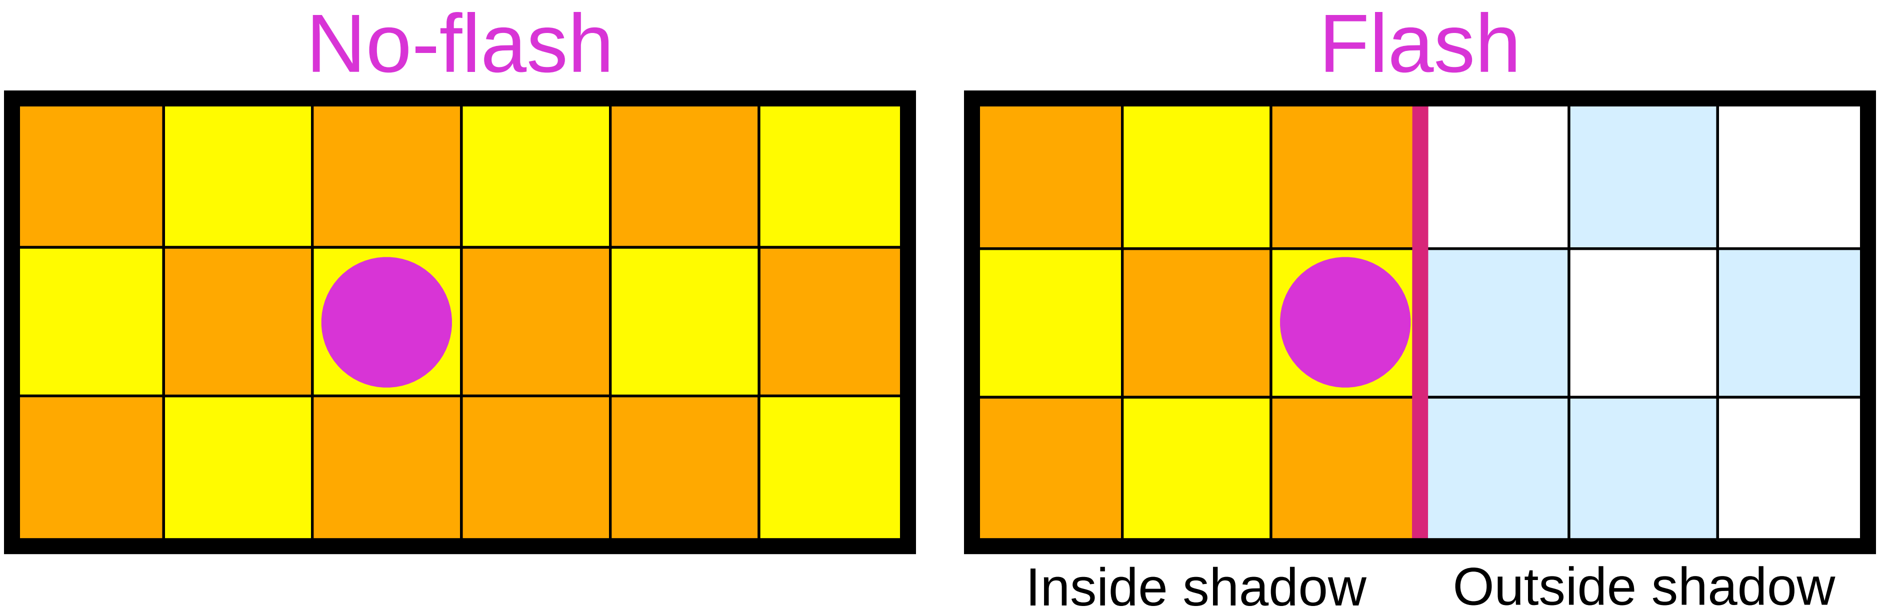

Use shadow mask

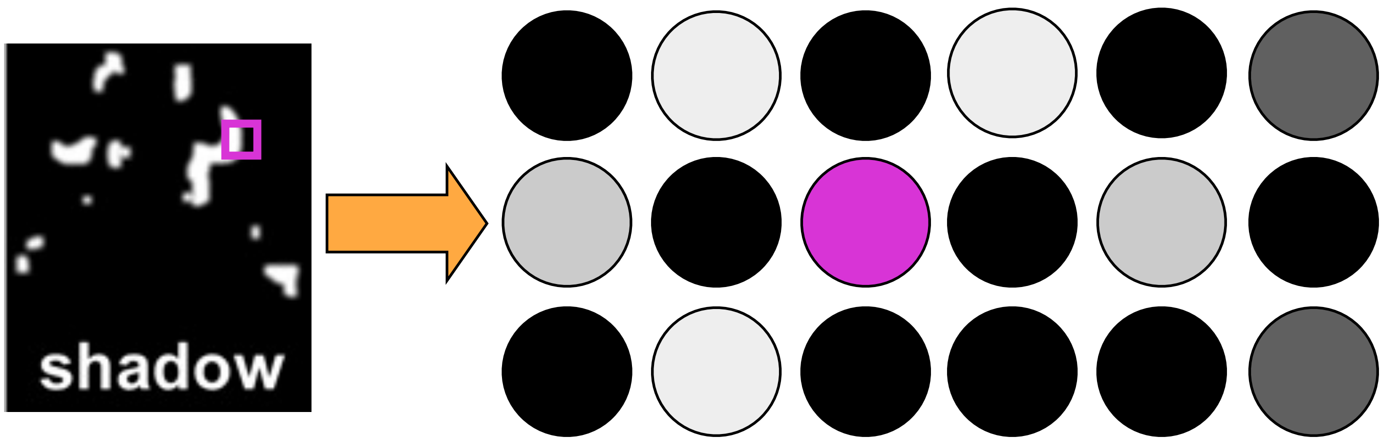

Shadow Color Correction

No-flash

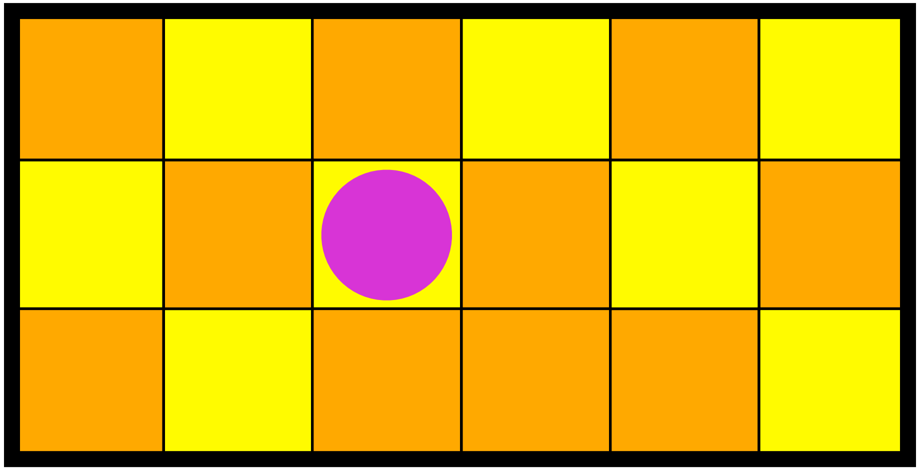

Flash

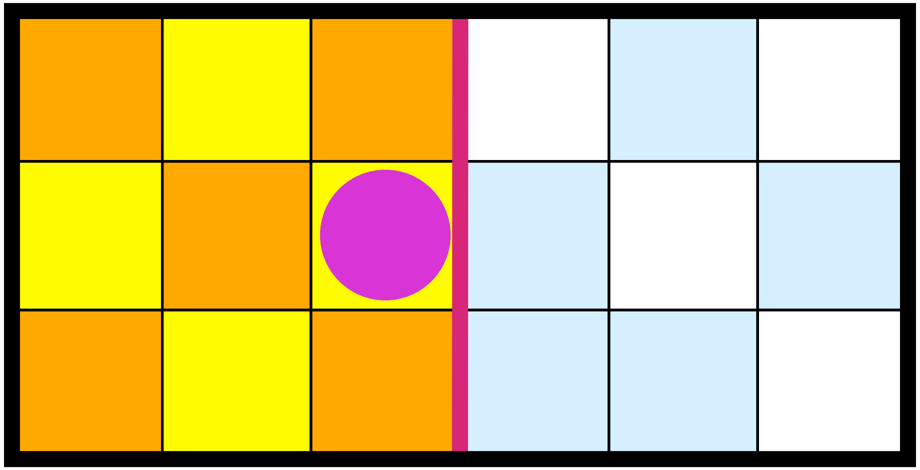

Inside shadow Outside shadow

Use shadow mask

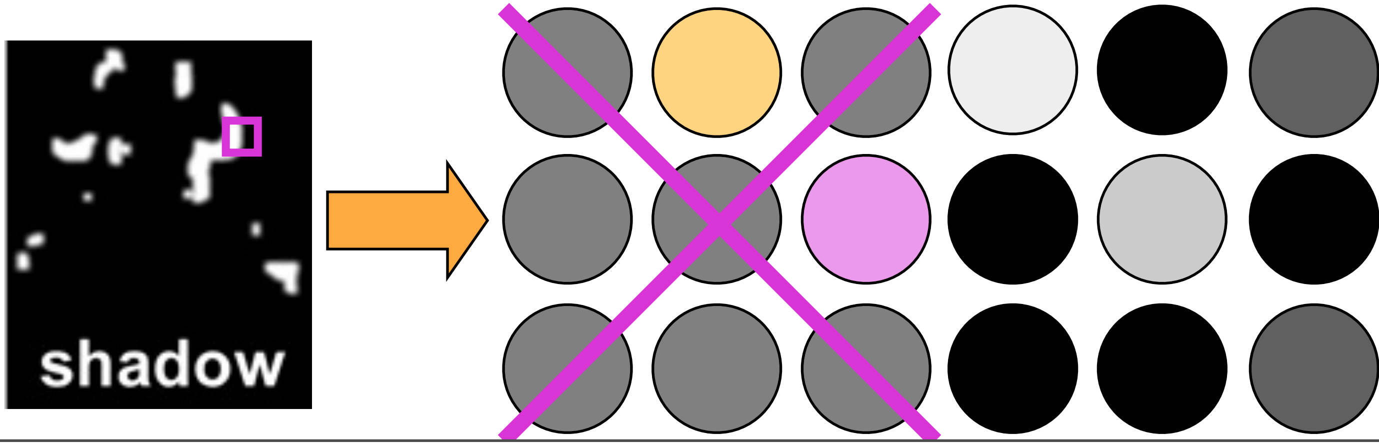

Shadow Color Correction

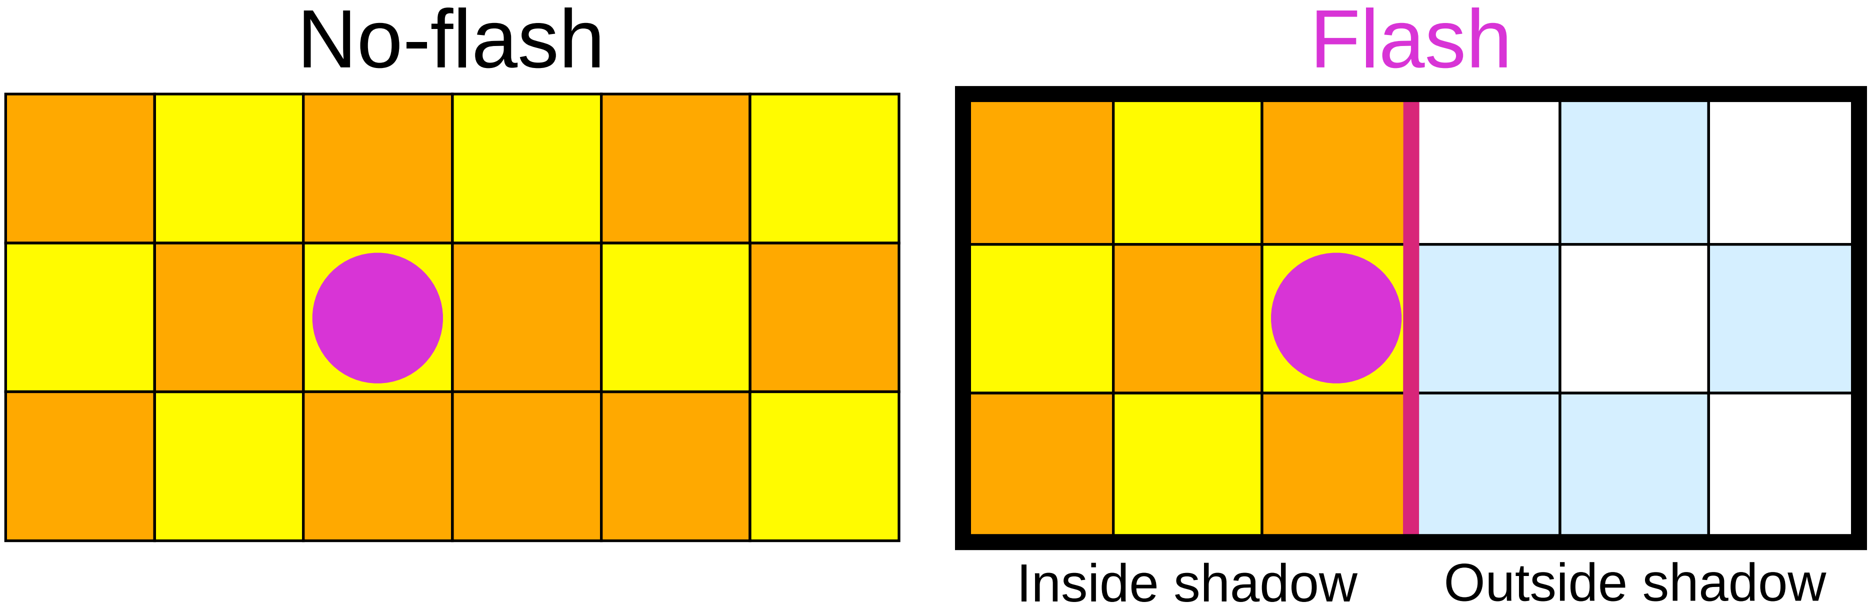

Use shadow mask

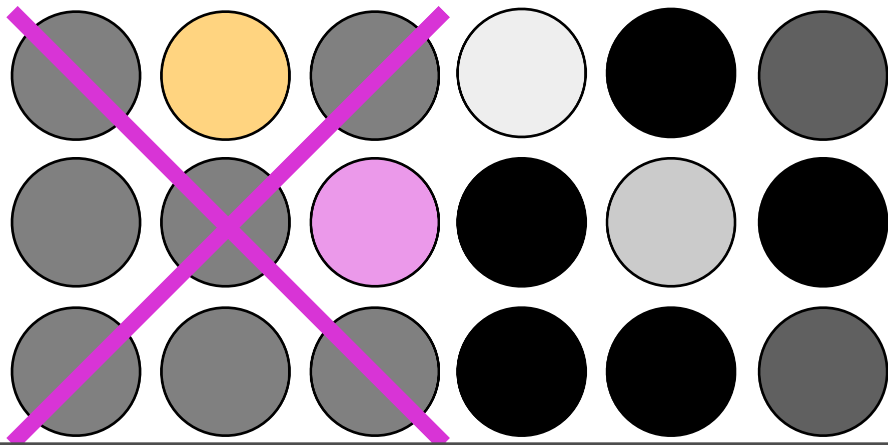

Shadow Color Correction

No-flash

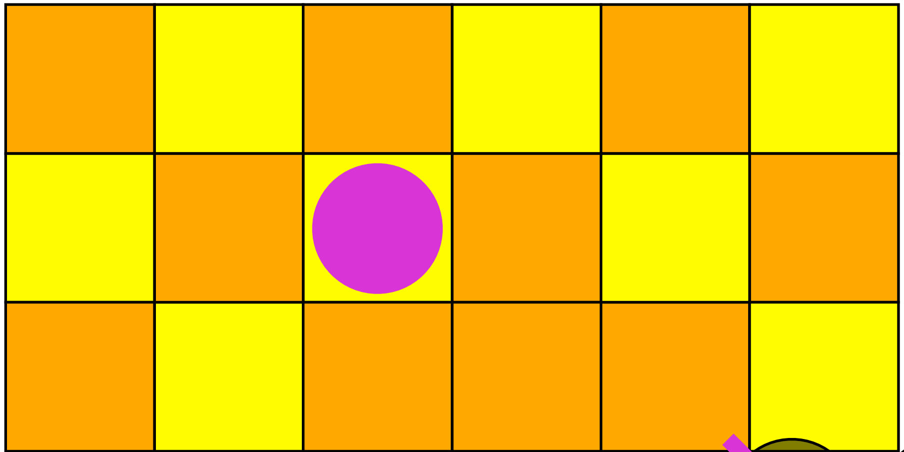

Flash

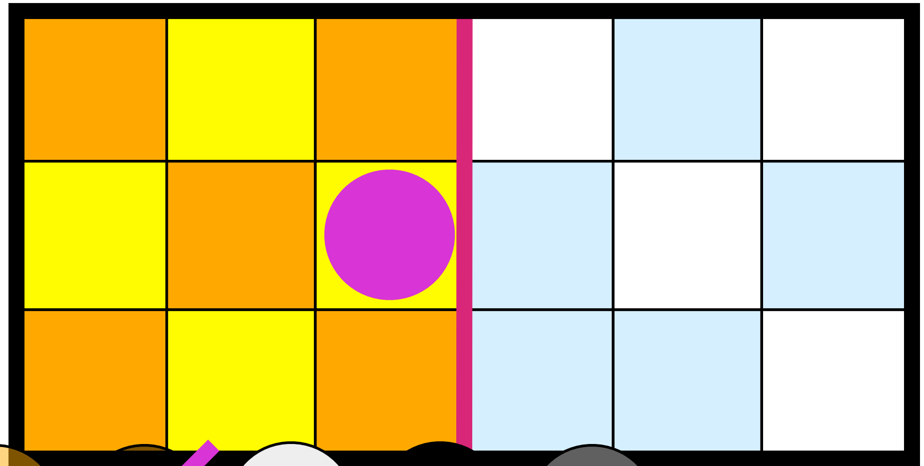

Use shadow mask

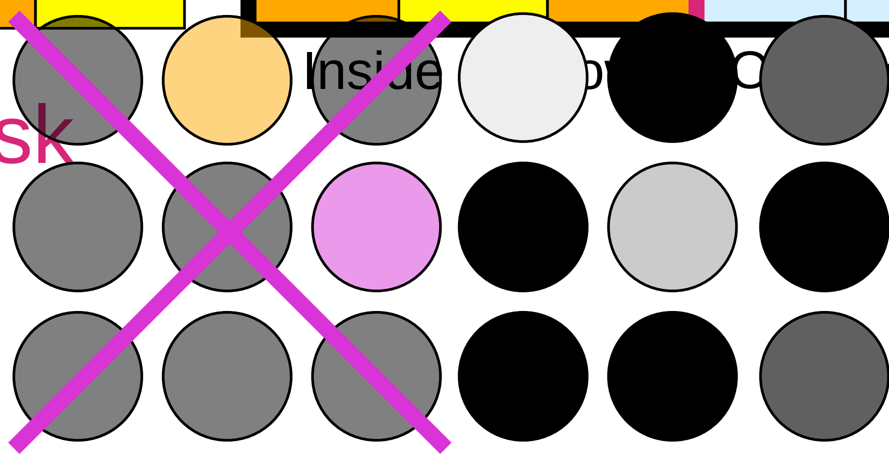

Shadow Color Correction

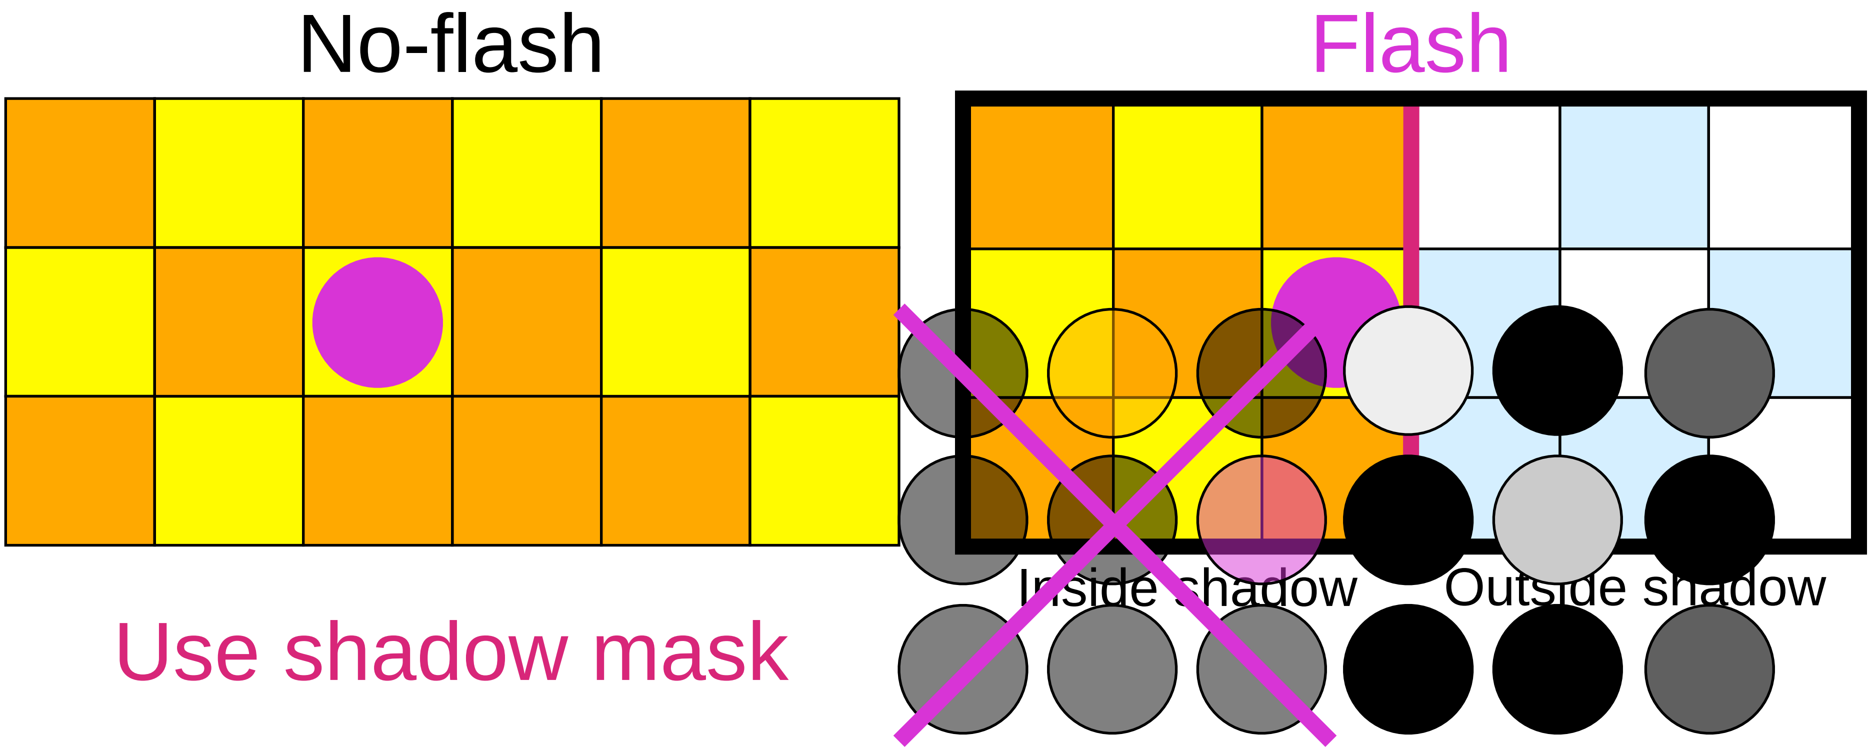

Use shadow mask

Shadow Color Correction

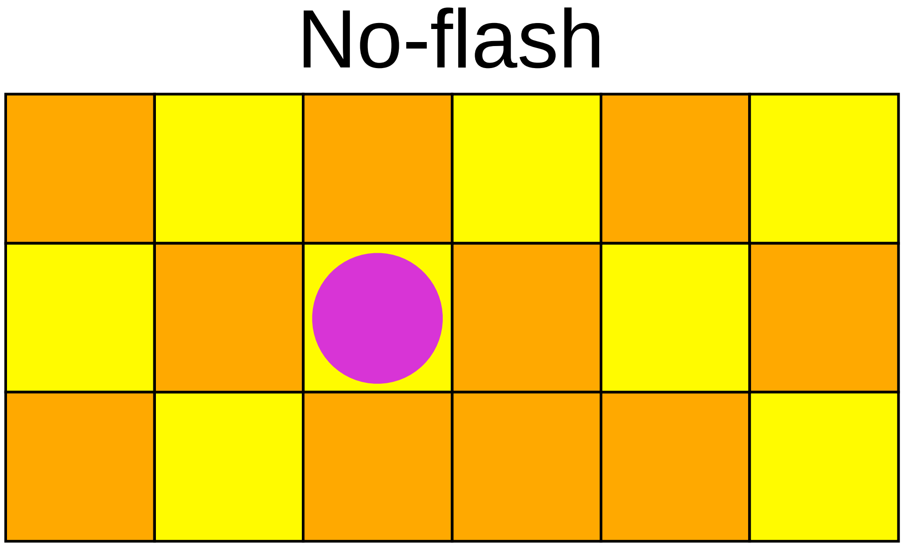

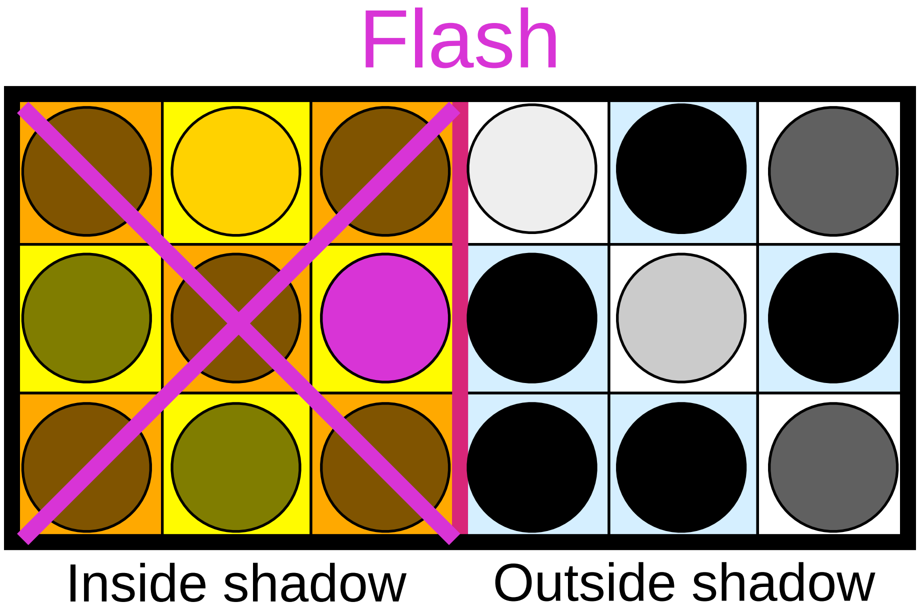

Use shadow mask

Shadow Color Correction

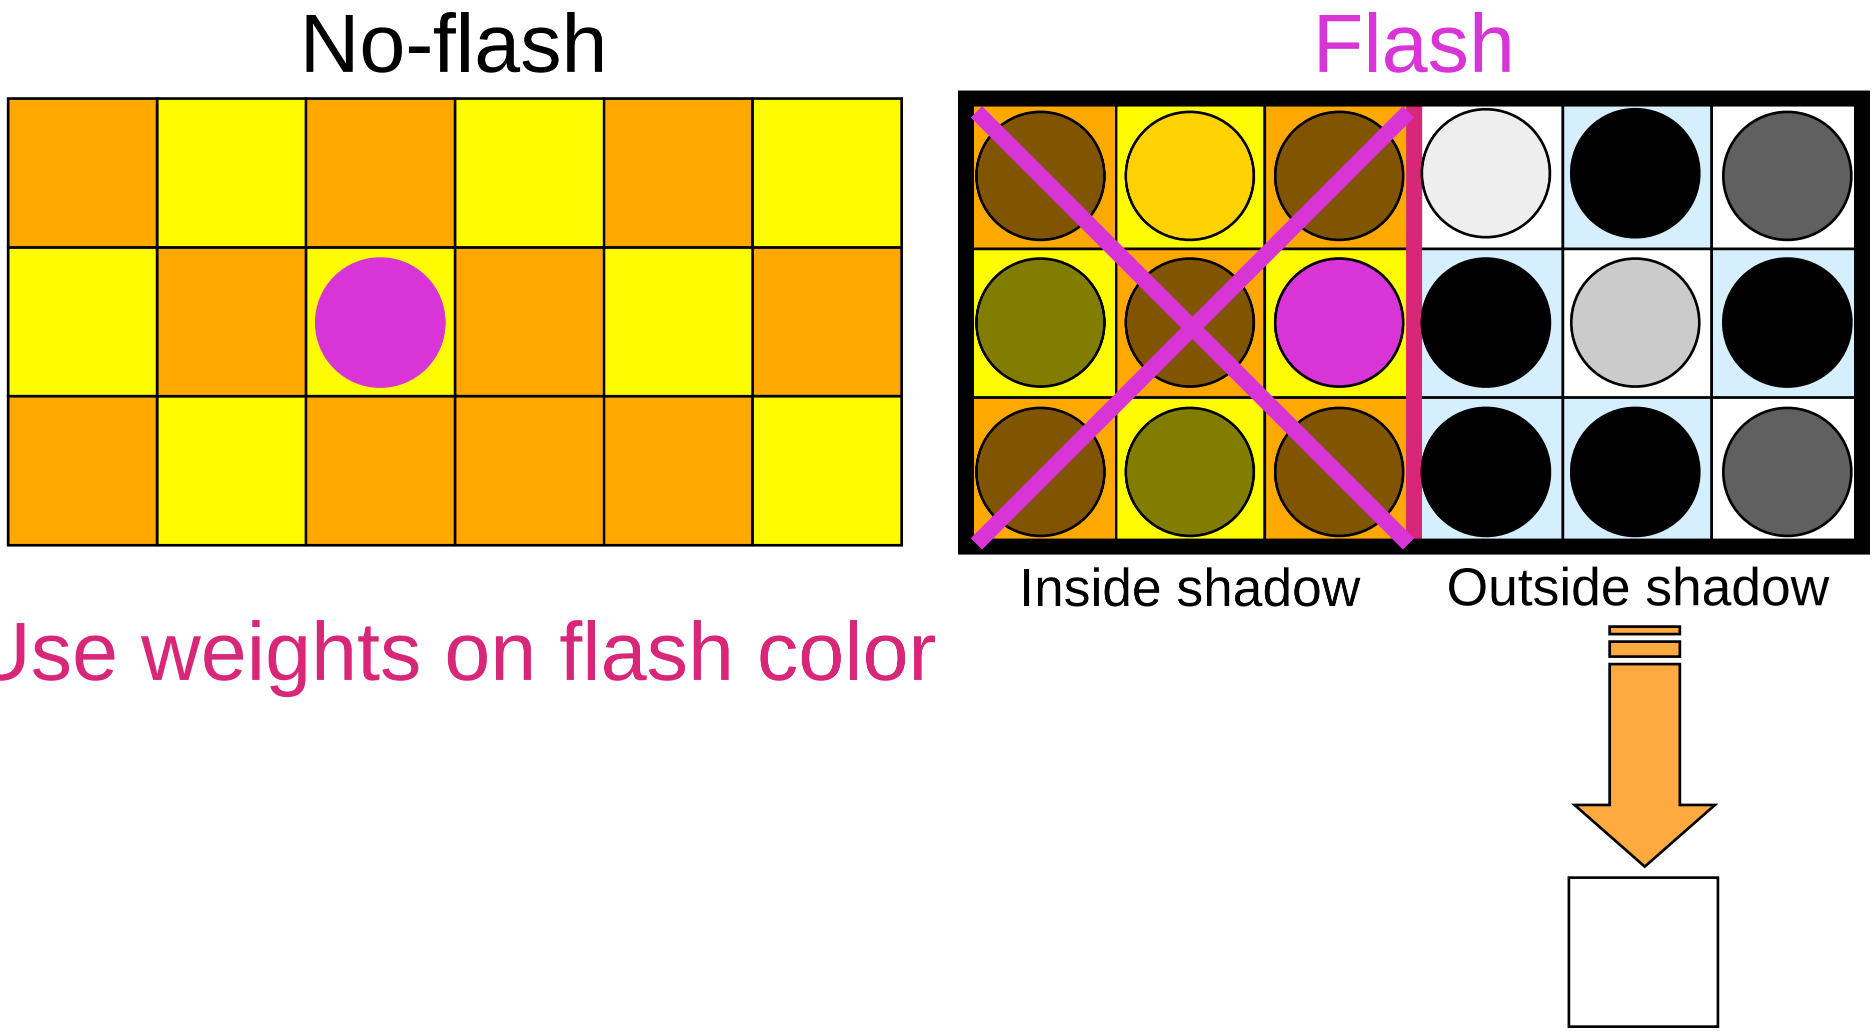

Use weights on flash color

Shadow Color Correction

No-flash

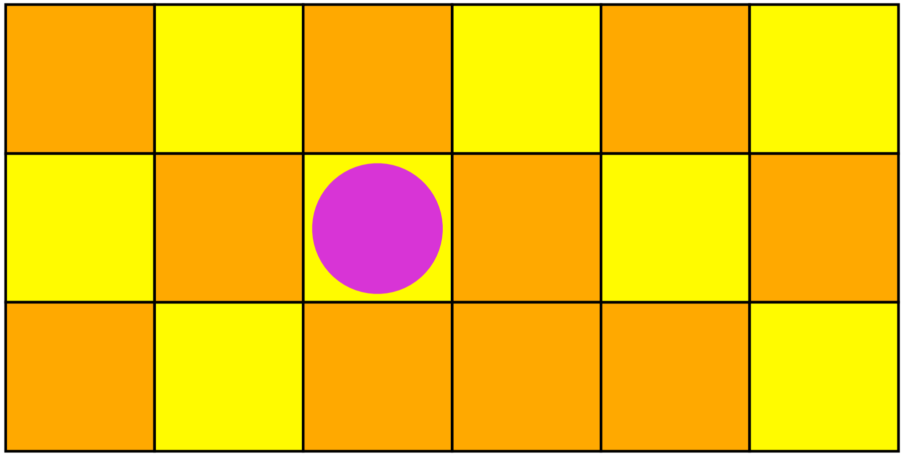

Flash

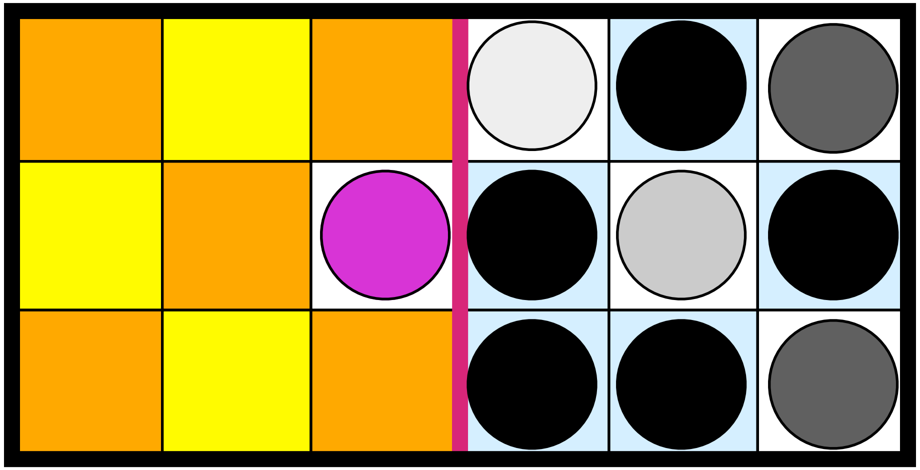

Inside shadow

Outside shadow

Replace shadow pixel

Shadow Color Correction

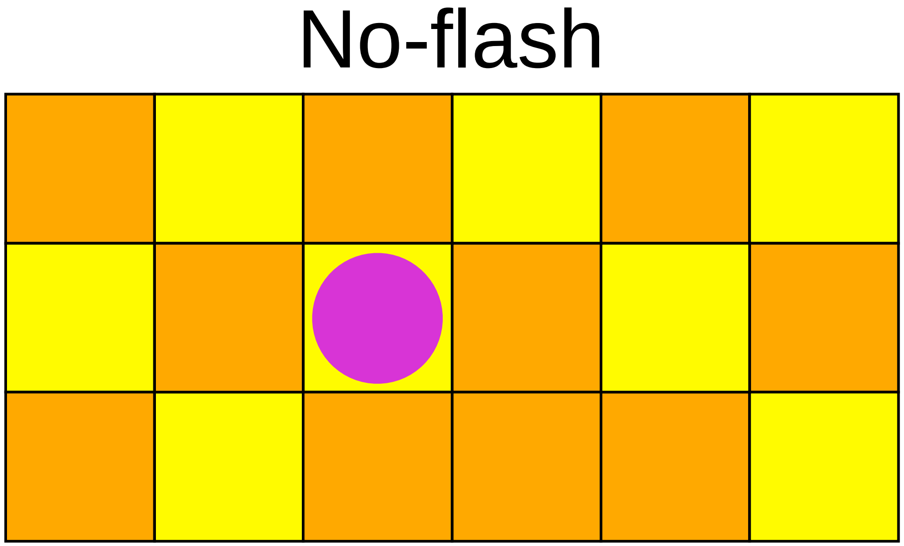

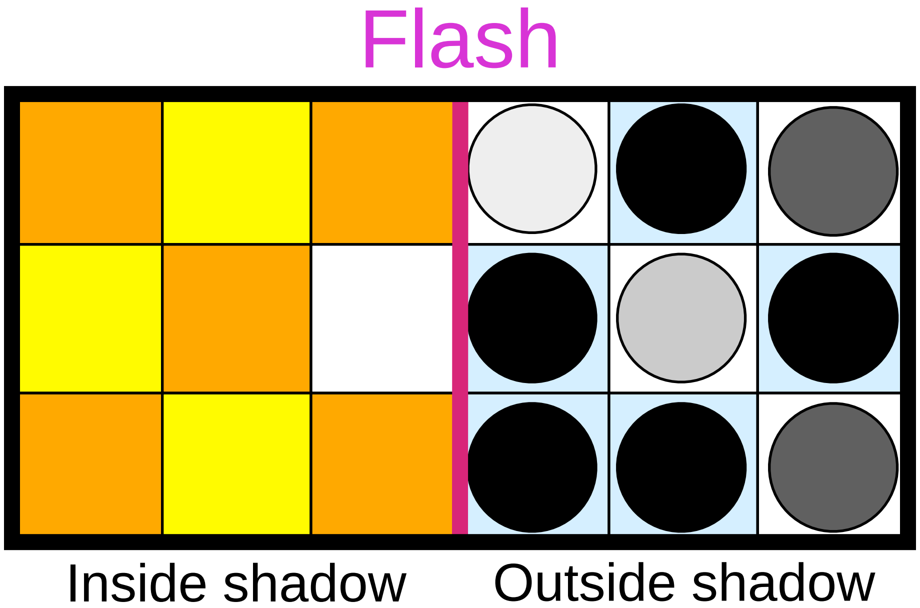

Proceed for all shadow pixels

Shadow Color Correction

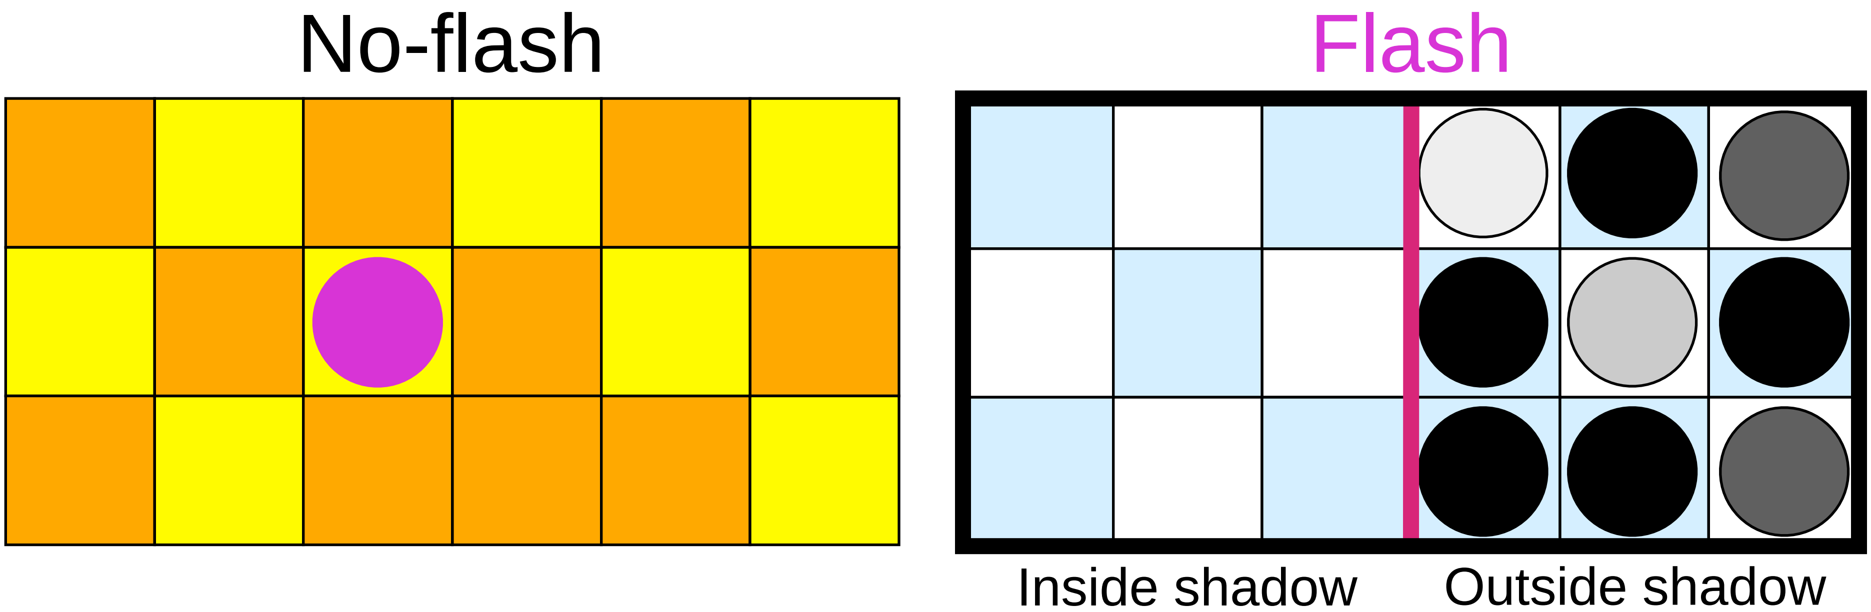

Proceed for all shadow pixels

Our Approach

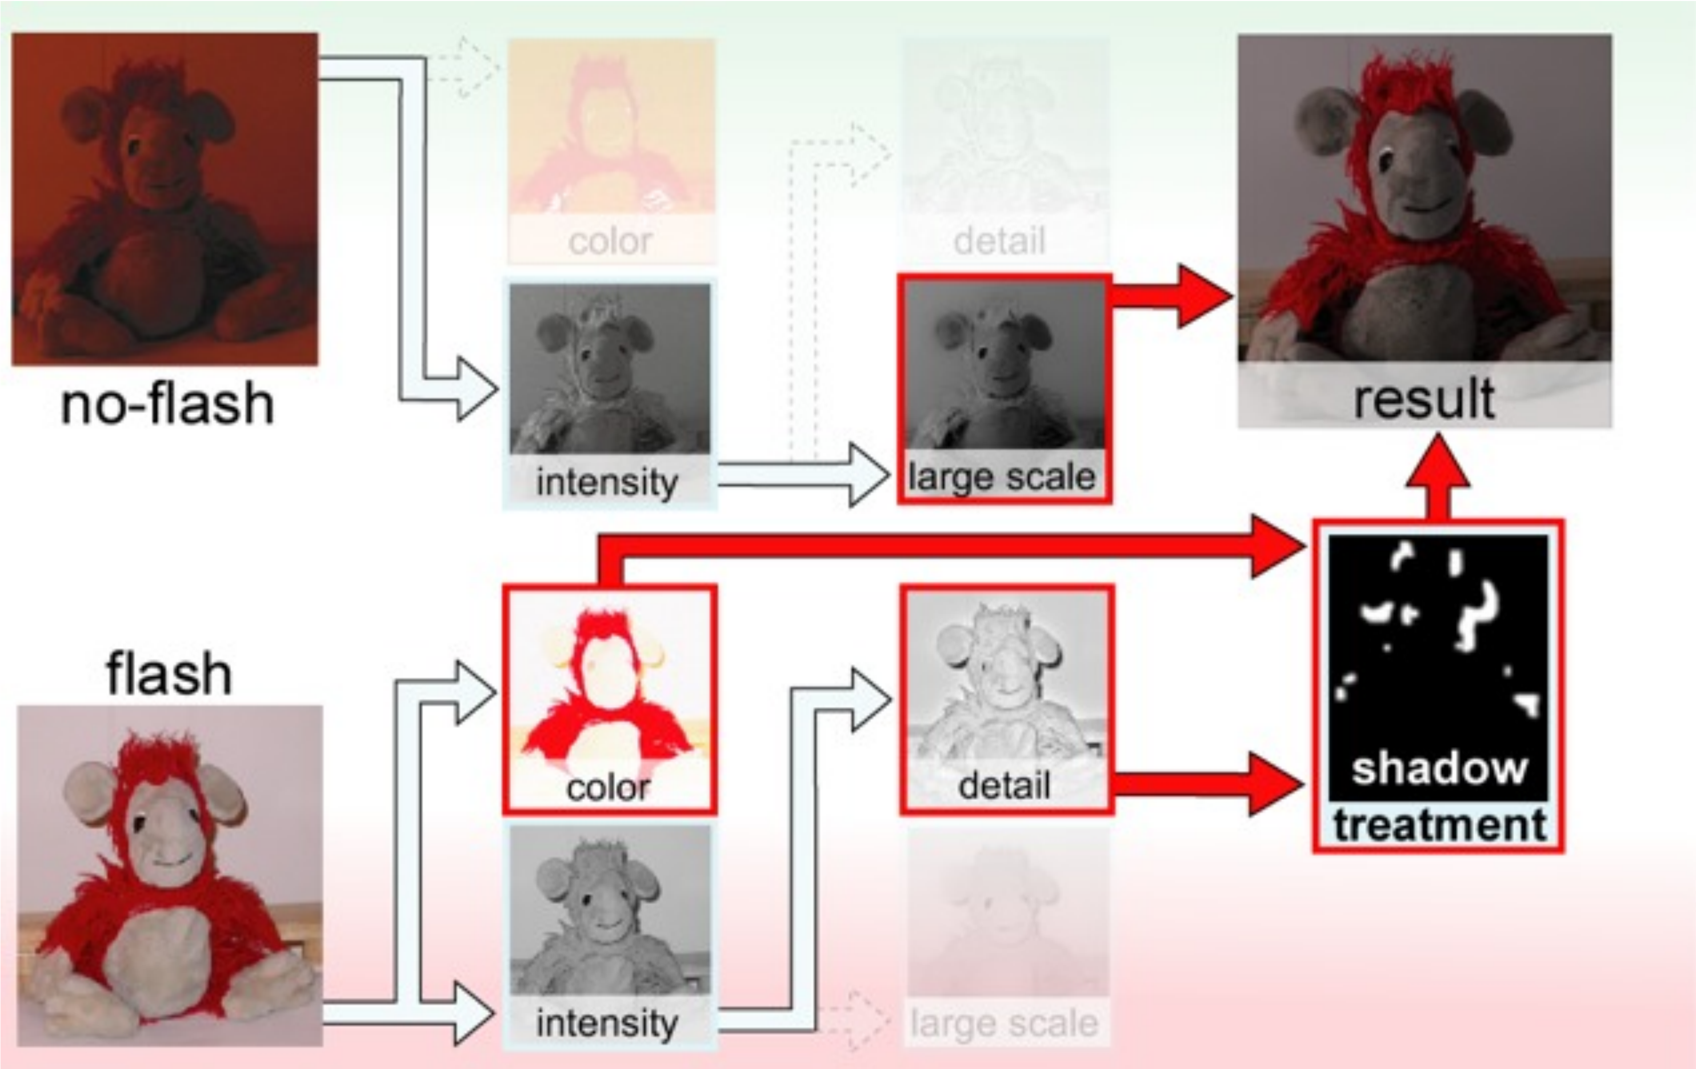

Our Approach

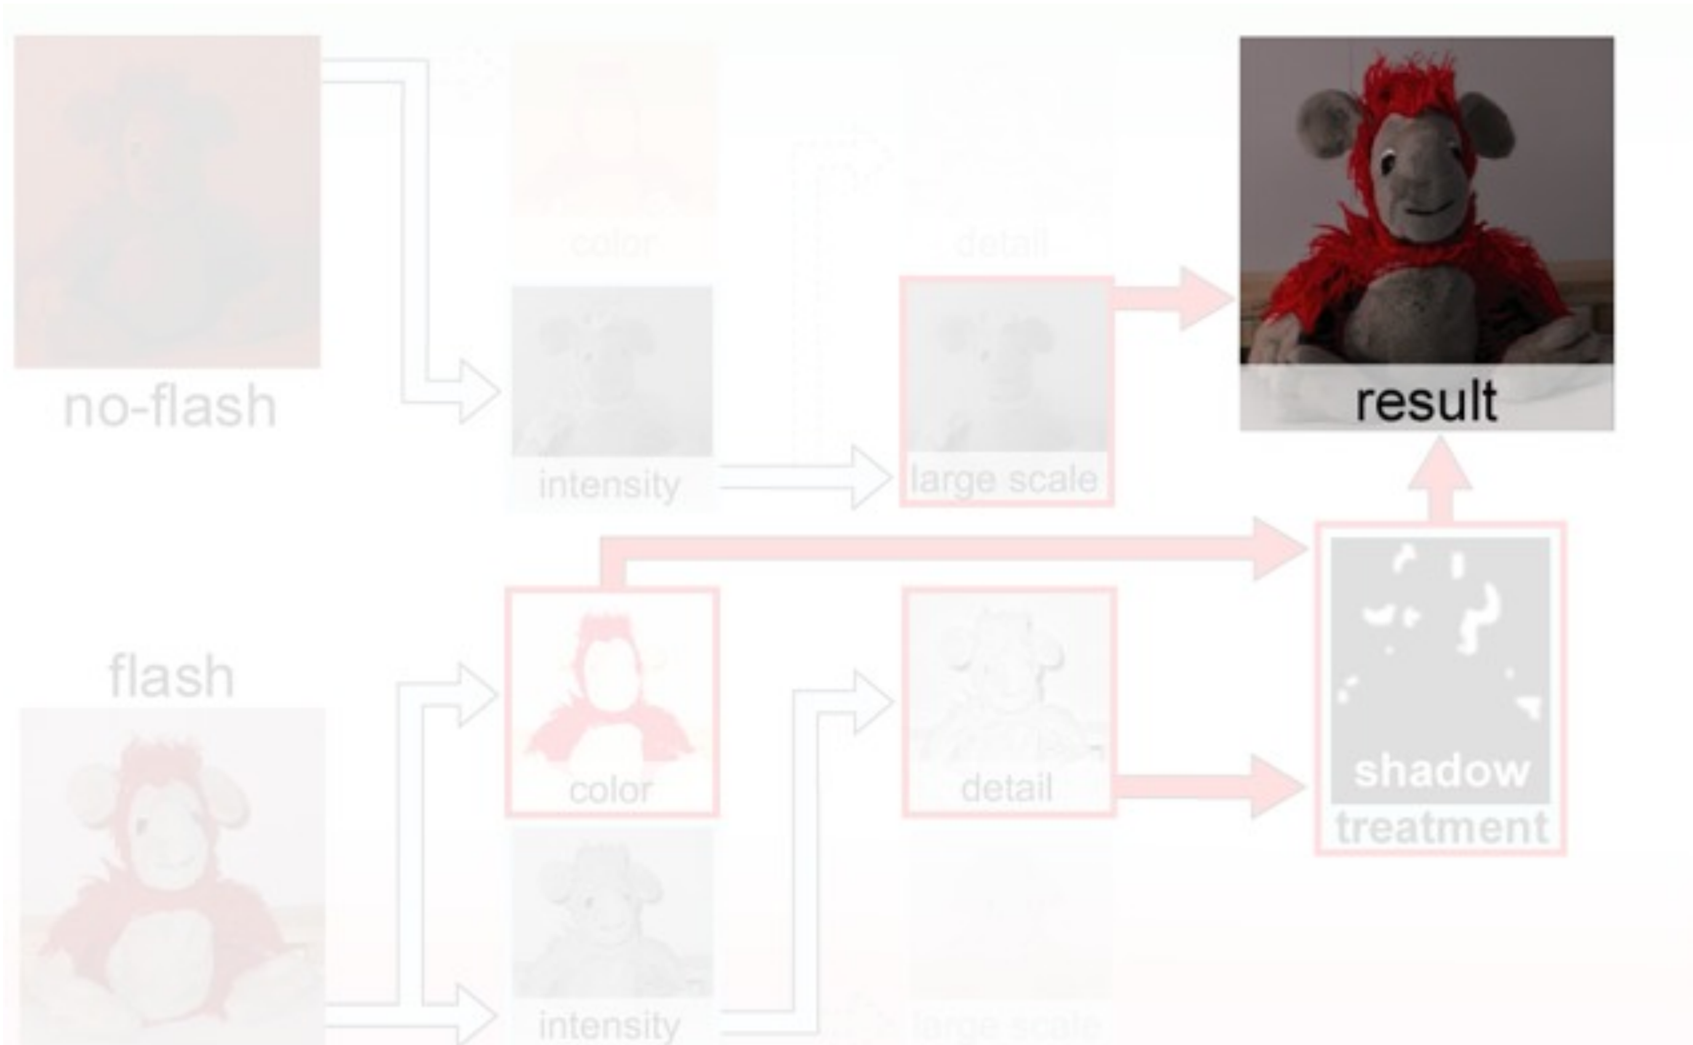

Overview

- Related Work
- Our Approach
- Results
- **Conclusion and Future Work**

Results

No-flash

Flash

Results

No-flash

Flash

Results

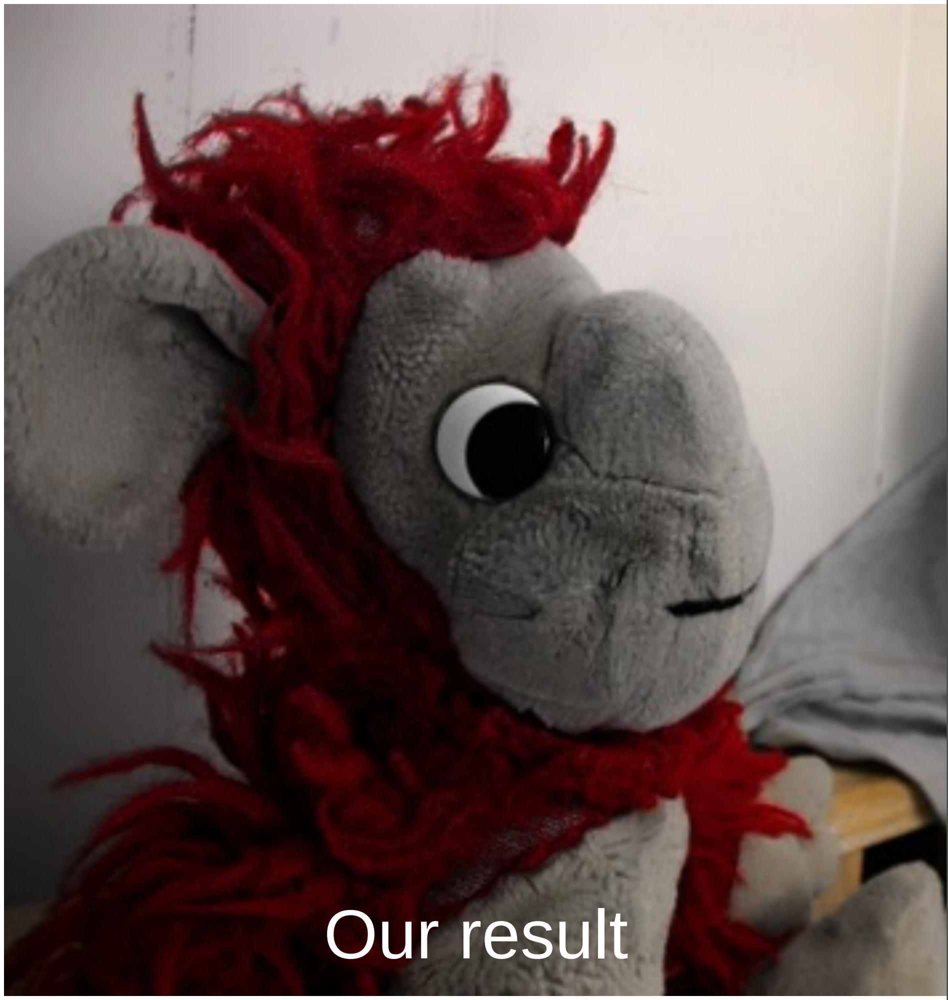

Results

No-flash

Flash

Results

No-flash

Flash

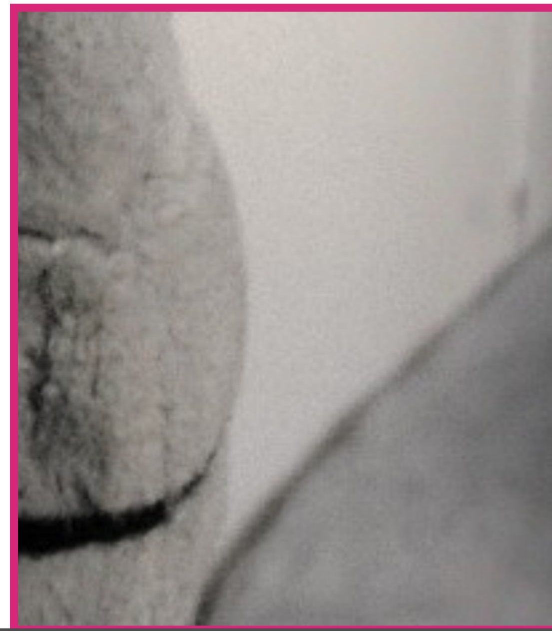

Results

No-flash

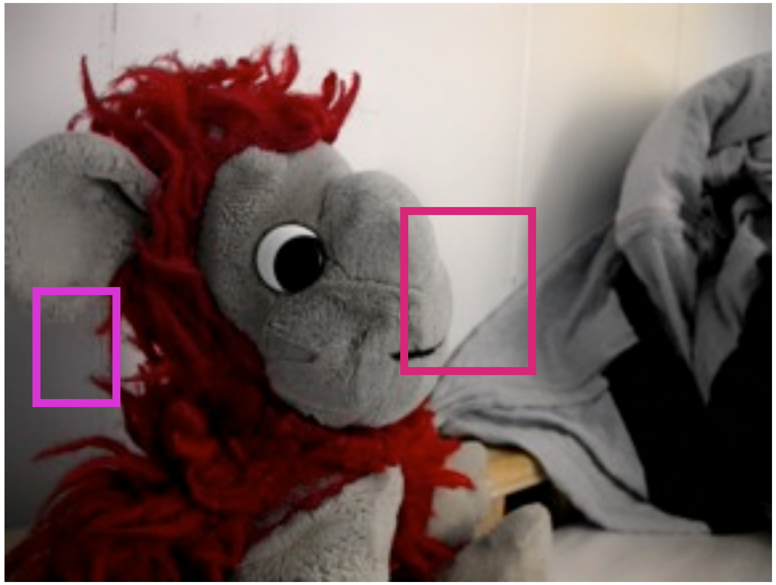

Flash

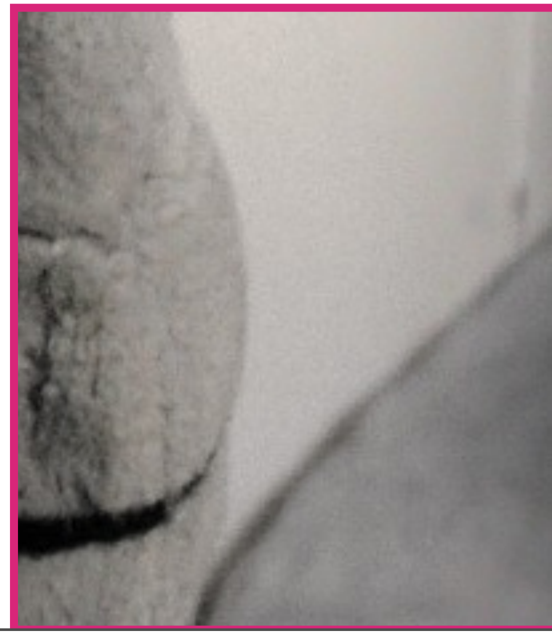

Results

No-flash

Flash

Results

No-flash

Flash

Results

Results

Results

Results

Results

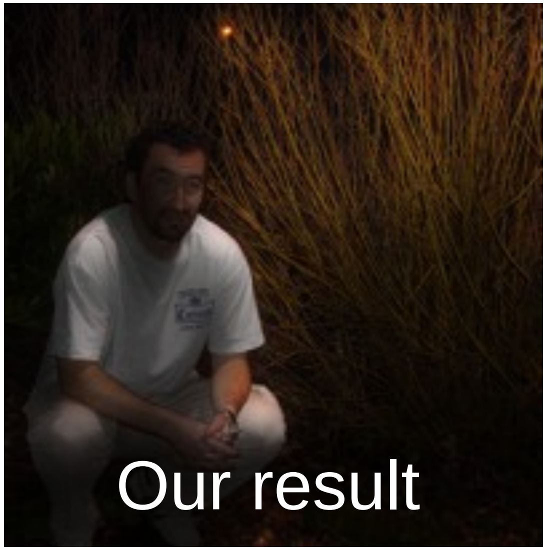

Results

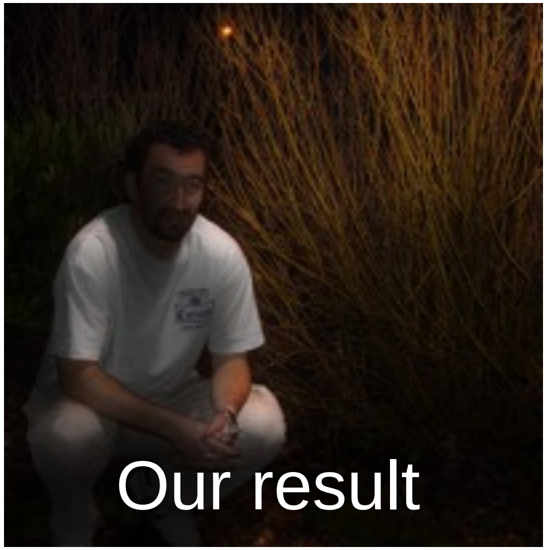
Our result

Deduce “distance“ to camera

Exploit $1/r^2$ flash intensity falloff

Results

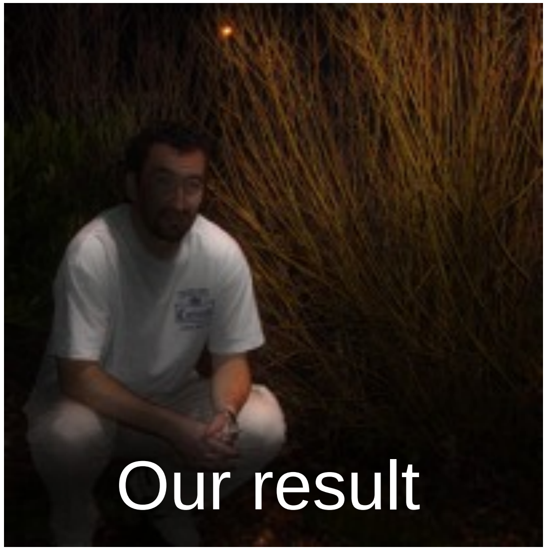

Deduce “distance“ to camera

Exploit $1/r^2$ flash intensity falloff

Results

(Inverse) White balance to original illumination

Results

(Inverse) White balance to original illumination

Results

(Inverse) White balance to original illumination

Results

(Inverse) White balance to original illumination

Results

No-flash

Flash

Results

No-flash

Flash

Overview

- Related Work
- Our Approach
- Results
- Conclusion and Future Work

Conclusion

Improving photography in dim environments

- **Capture original lighting**
- **Add sharpness/details**

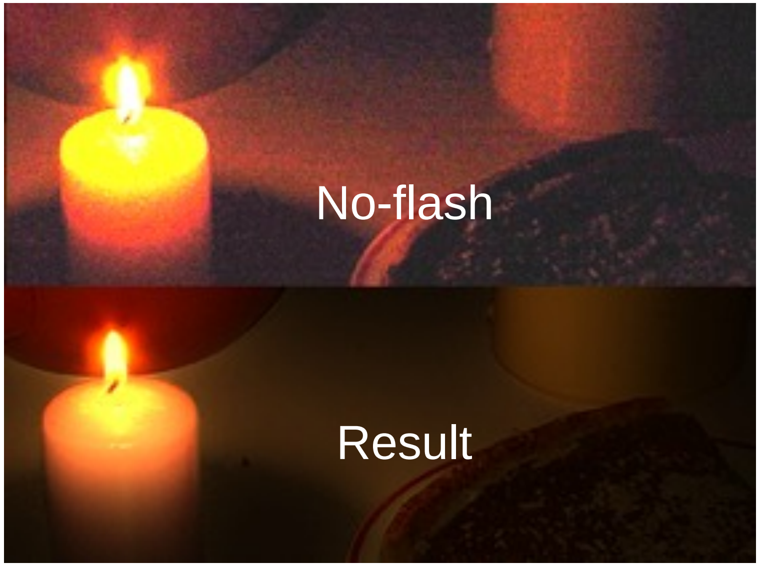

- Correct flash shadows

- Pseudo distance
  (emphasize foreground)

- white balancing

- Cross Bilateral Filter

NPR multiflash

Non-photorealistic Camera: Depth Edge Detection and Stylized Rendering using Multi-Flash Imaging

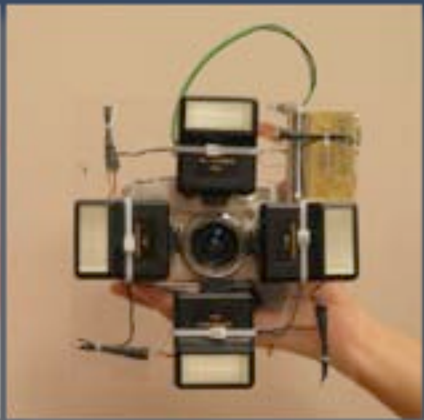

Ramesh Raskar, Karhan Tan, Rogerio Feris, Jingyi Yu, Matthew Turk

Mitsubishi Electric Research Labs (MERL), Cambridge, MA

Depth Edge Camera

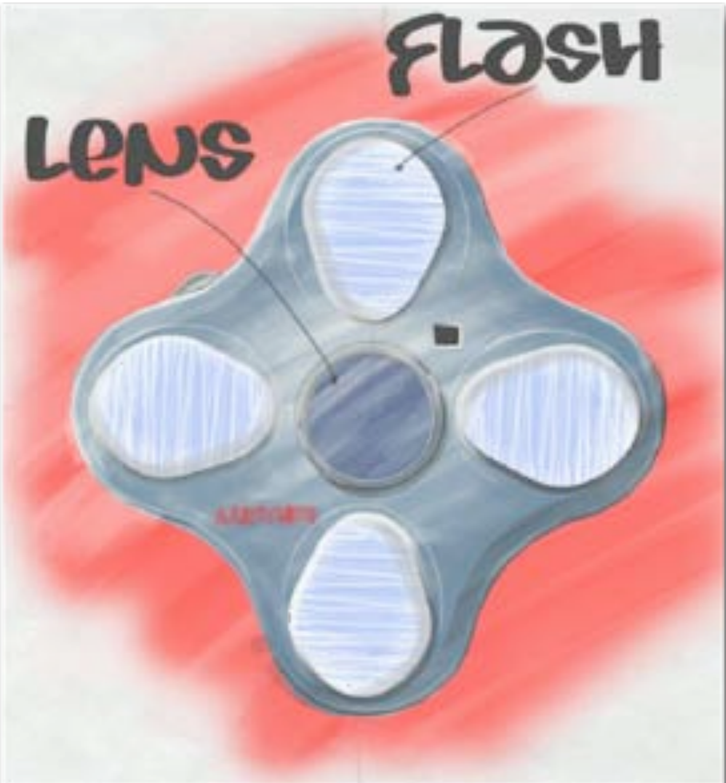

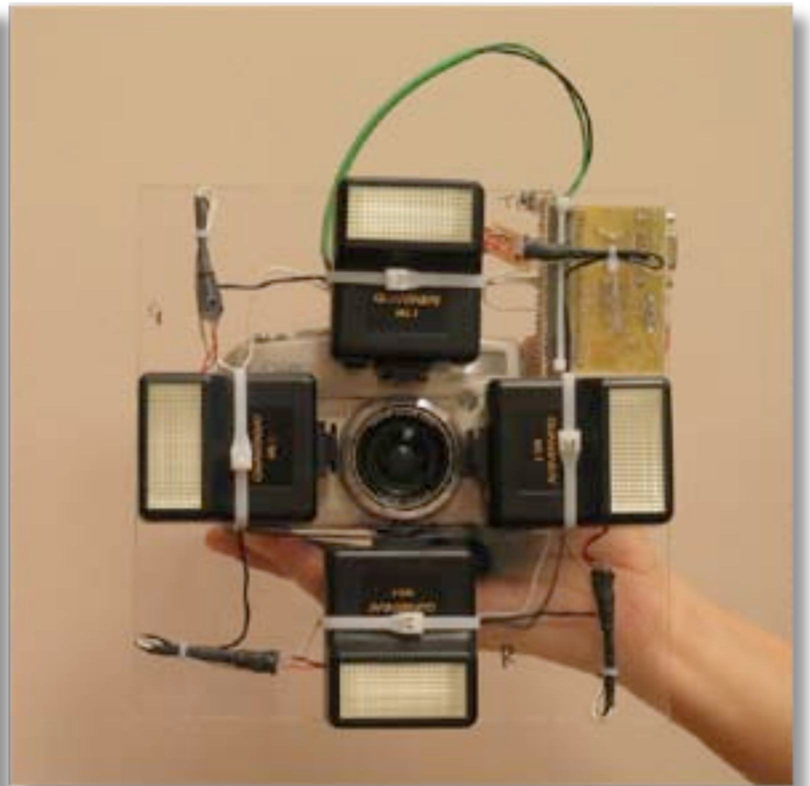

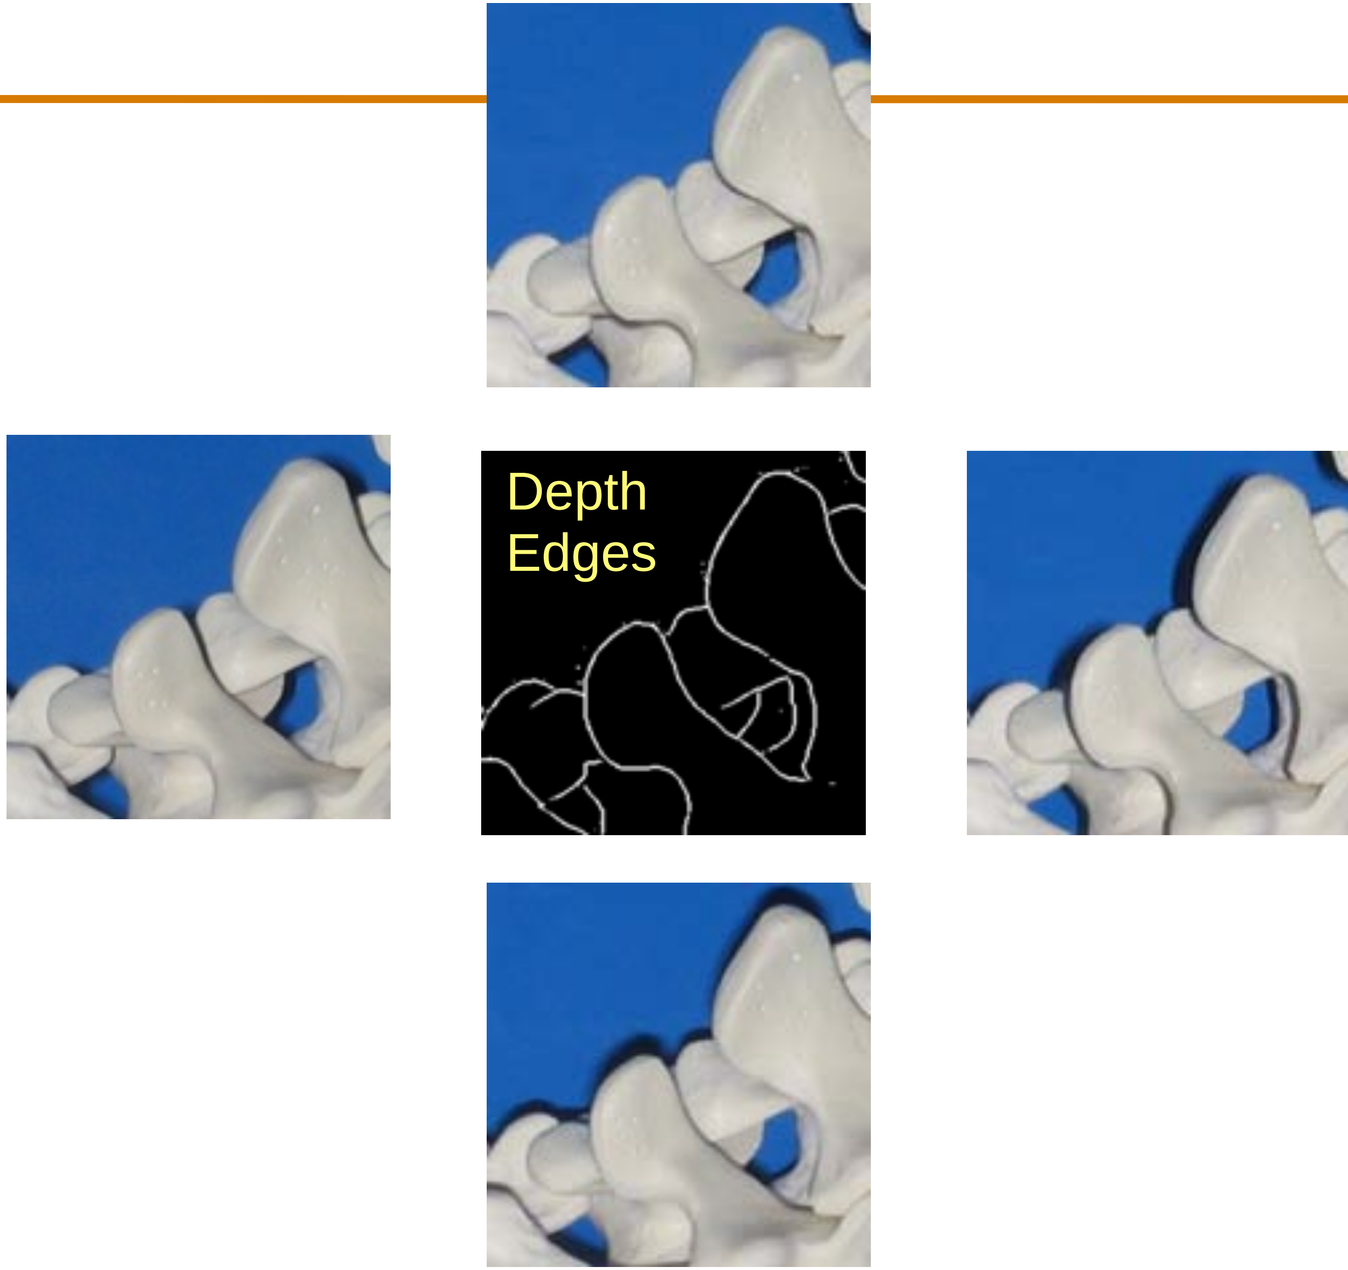
Depth
Edges

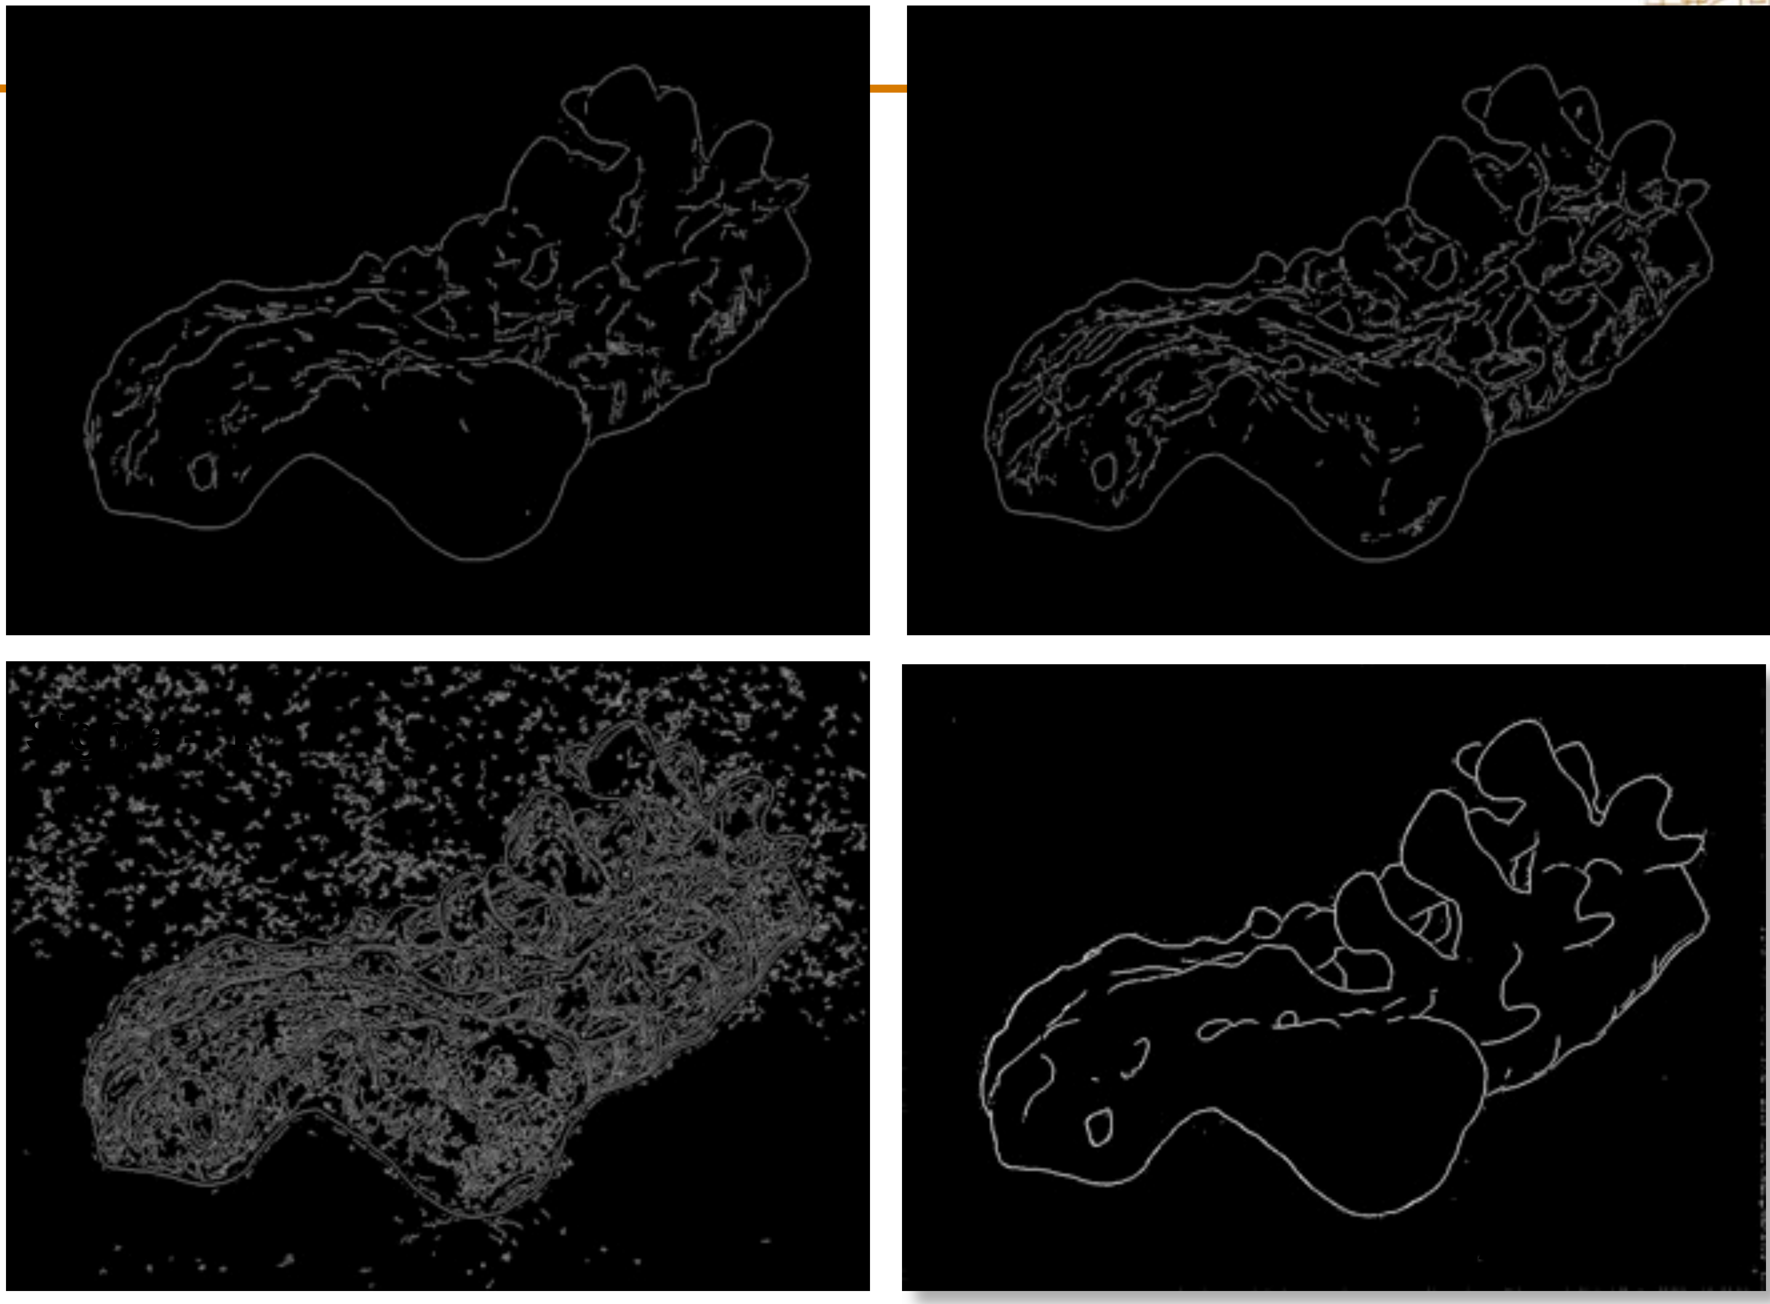


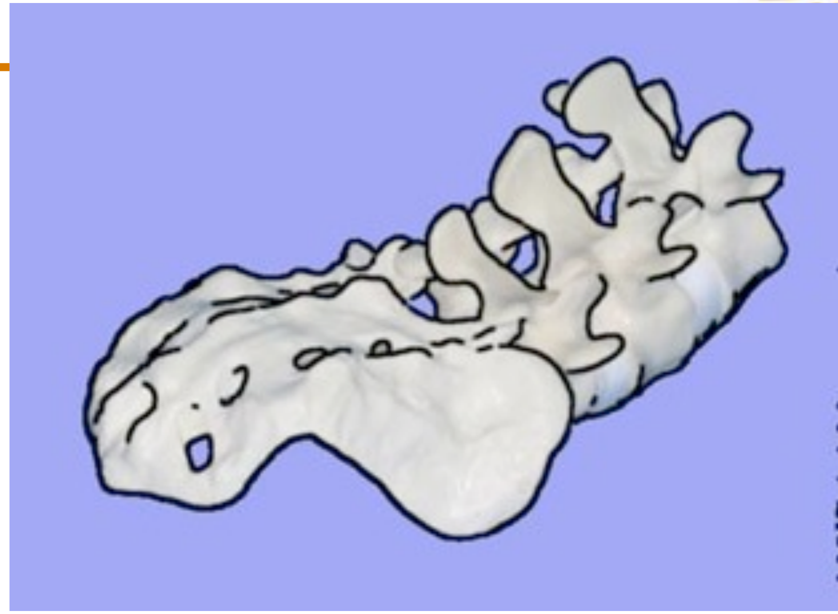


Photo

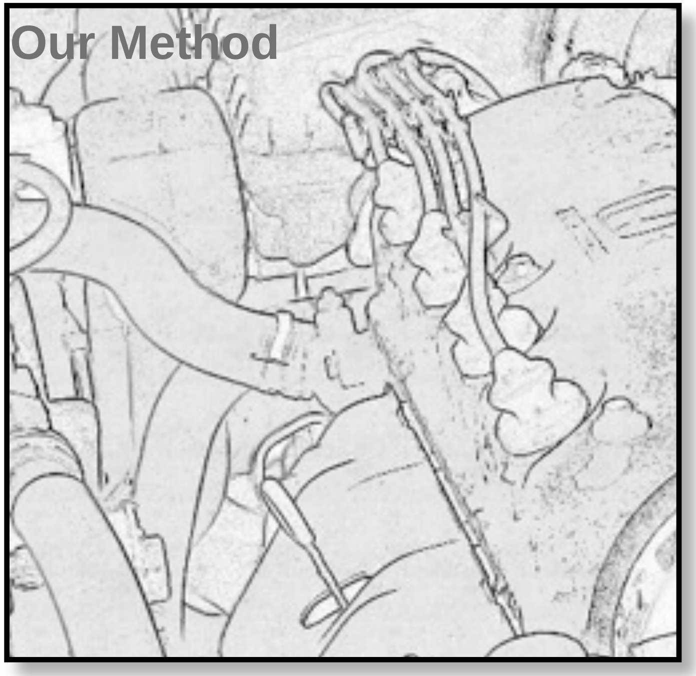
Our Method


Photo
Result
Our Method
Canny Intensity
Edge Detection

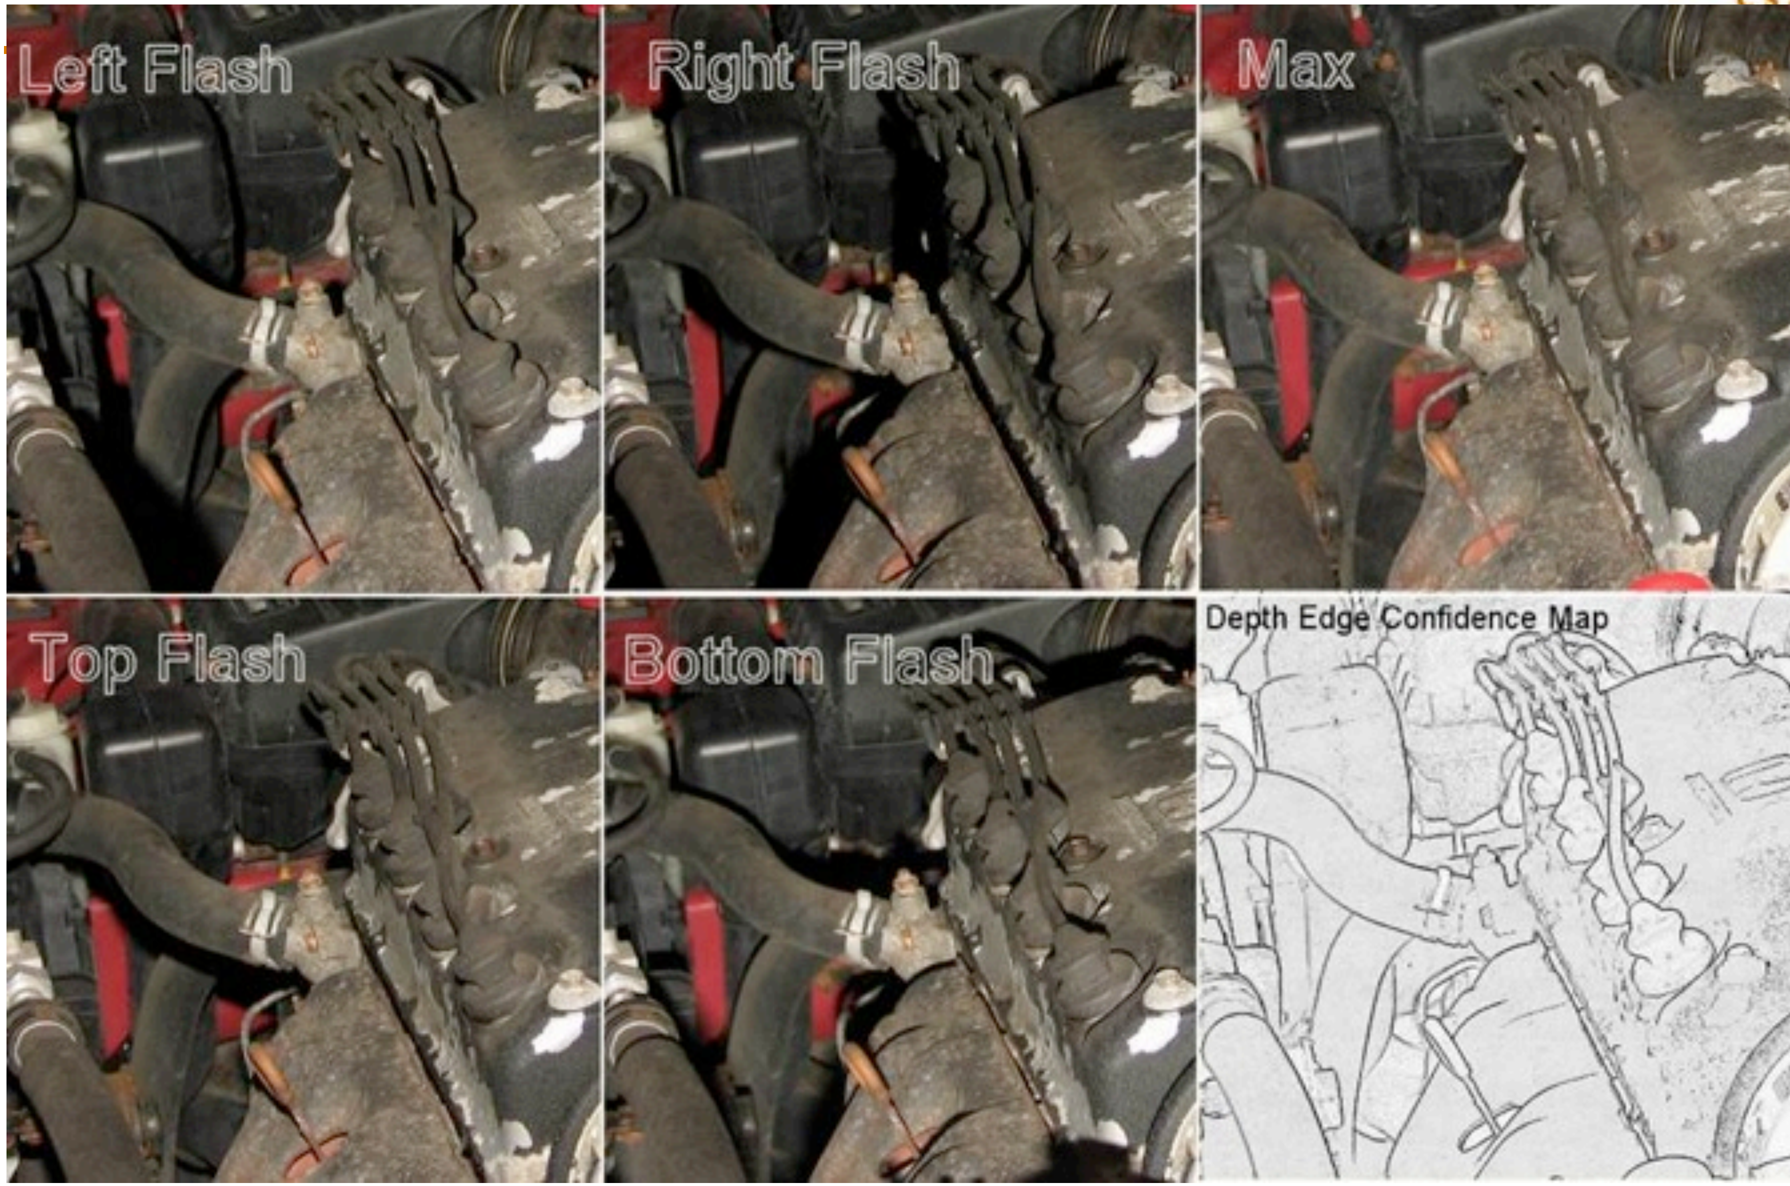
Left Flash
Right Flash
Max
Top Flash
Bottom Flash
Depth Edge Confidence Map

Shadows

Clutter

Many Colors

Highlight Shape Edges

Mark moving parts

Basic colors

Flash matting

Flash Matting

Jian Sun[1] Yin Li[1] Sing Bing Kang[2] Heung-Yeung Shum[1]

[1]Microsoft Research Asia, Beijing [2]Microsoft Research, Redmond, WA

Figure 1: Flash matting results for flower scene with complex foreground and background color distributions. From left to right: flash image, no-flash image, recovered alpha matte, and flower basket with new background.

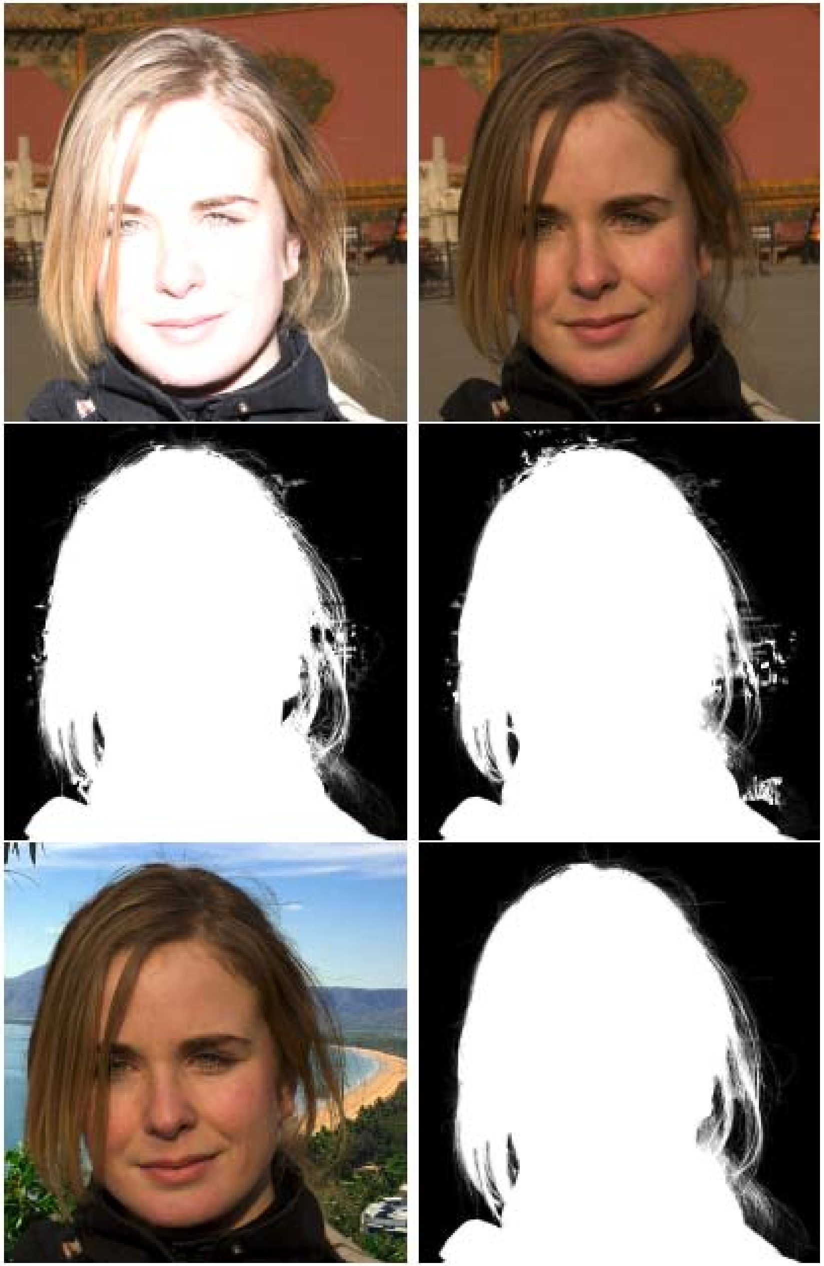

Figure 4: Continuous shot. Top: flash/no-flash pair. Middle: Bayesian matting results on the flash and no-flash images (separate estimation). Bottom: Joint Bayesian flash matting result.

Relighting

A Lighting Reproduction Approach to Live-Action Compositing

SIGGRAPH 2002 Conference Proceedings

Paul Debevec Andreas Wenger Chris Tchou Andrew Gardner Jamie Waese Tim Hawkins

USC Institute for Creative Technologies

Performance Relighting and Reflectance Transformation with Time-Multiplexed Illumination

SIGGRAPH 2005 Papers Proceedings

Andreas Wenger Andrew Gardner Chris Tchou
Jonas Unger Tim Hawkins Paul Debevec
University of Southern California Institute for Creative Technologies

Figure 5: Successive images of an actor lit by the 180-pattern sequence in the span of one twelfth of a second. Tracking and matte frames seen in the two left columns occur ten times each within the lighting basis with an effective rate of 120Hz in the 2160fps sequence.

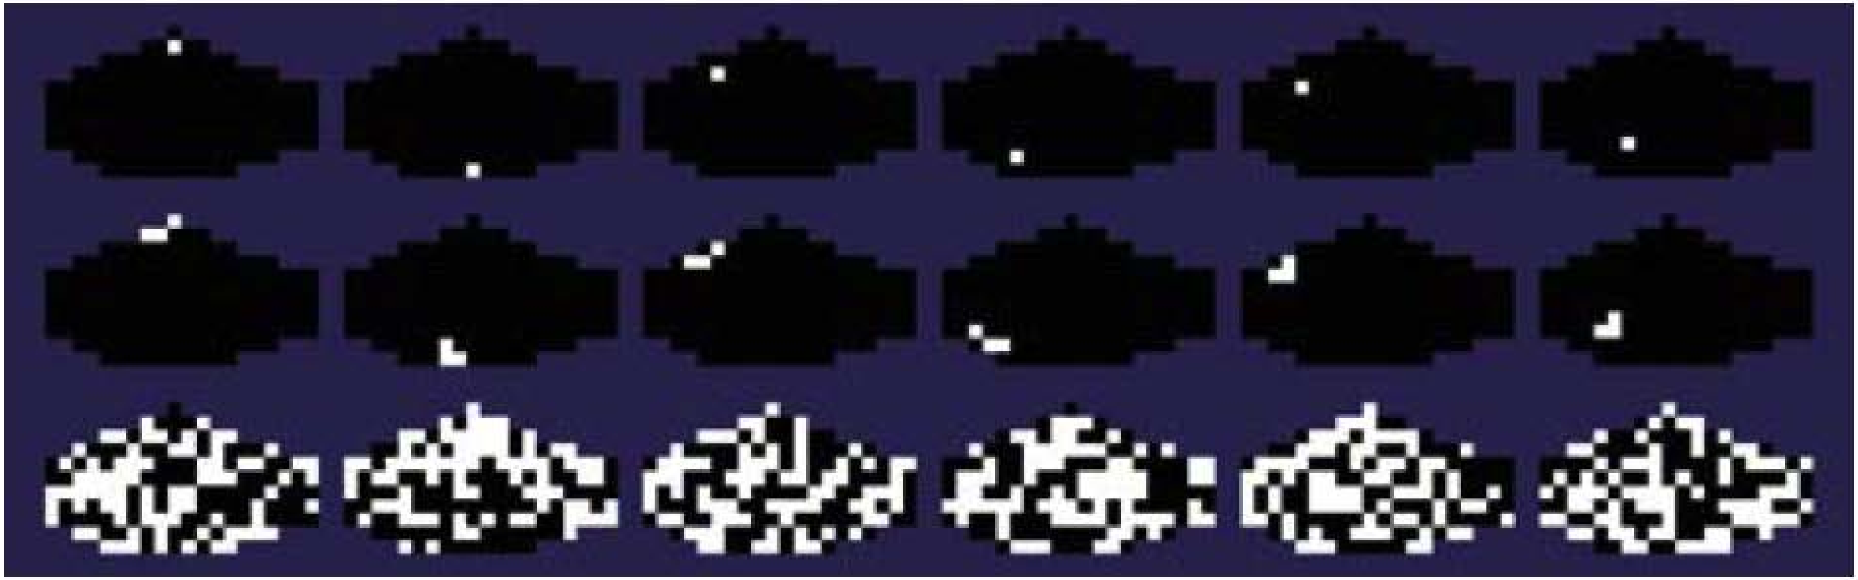

Figure 4: Six elements of each of the three lighting bases used in this work. Top Row: Single lights. Middle Row: Triangles. Bottom Row: Hadamard Patterns

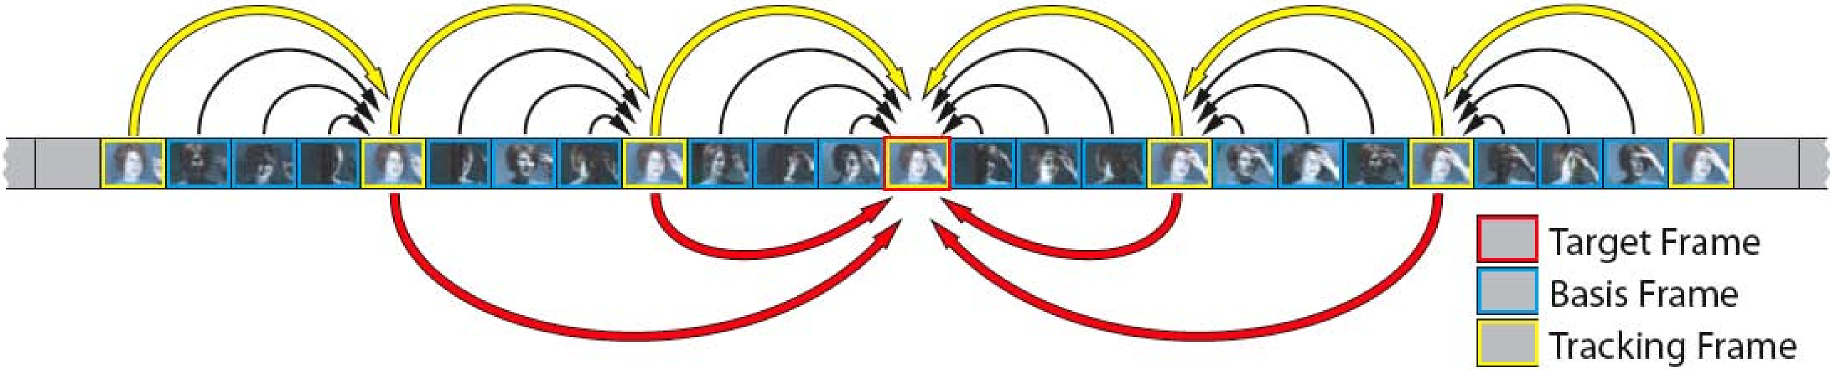

Figure 7: **Motion Compensation Process.** To compensate for the subject motion, basis frames (blue) are warped to match the target output frame (red). Optical flow is calculated between adjacent tracking frames (yellow), and linearly interpolated to warp each basis frame to the nearest tracking frame (black arrows) in the direction of the target output frame. Then, long-range warps (red arrows) are applied to bring each basis frame into alignment with the output frame.

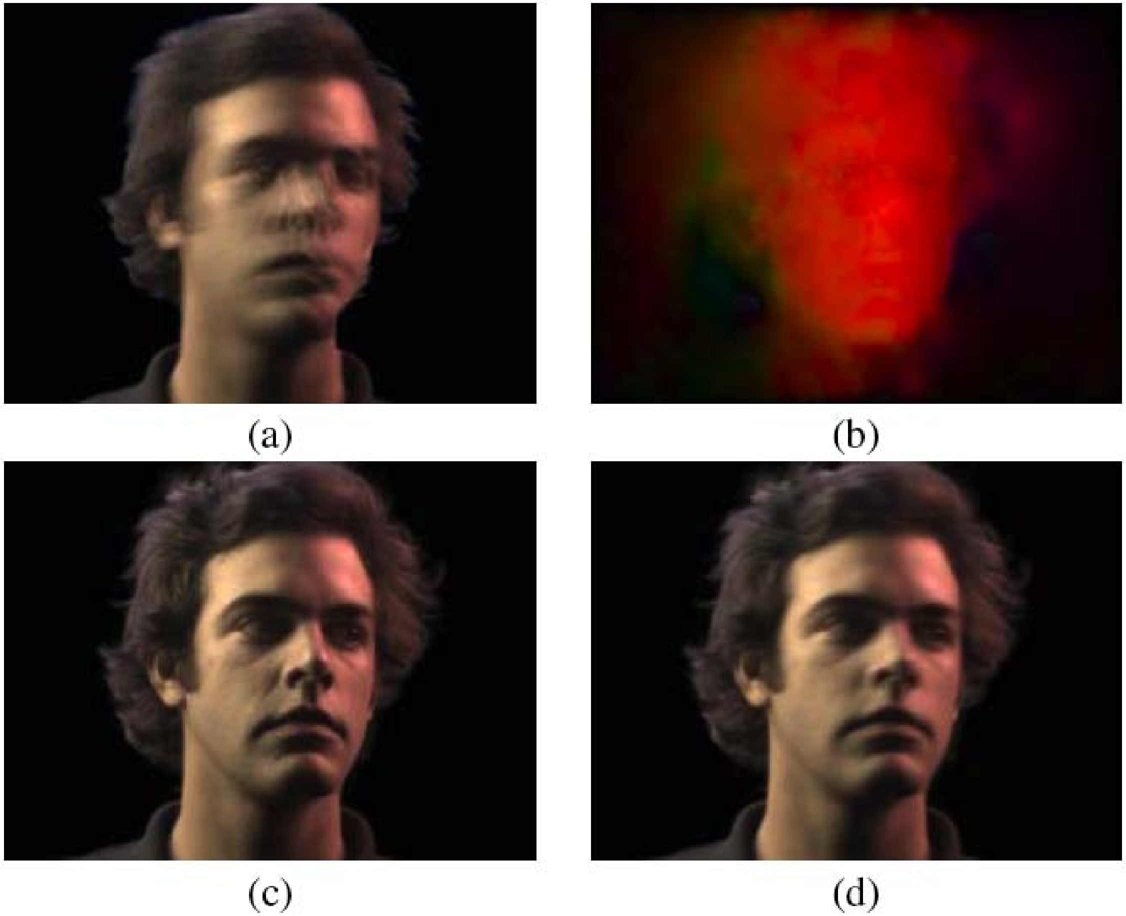

Figure 8: (a) Relit frame without motion compensation, showing smearing. (b) Color-coded optical flow field computed between neighboring tracking frames (c) Stabilized frame where motion compensation has been applied to the basis. (d) Relit frame with synthesized 180-degree shutter motion blur based on the flow field.

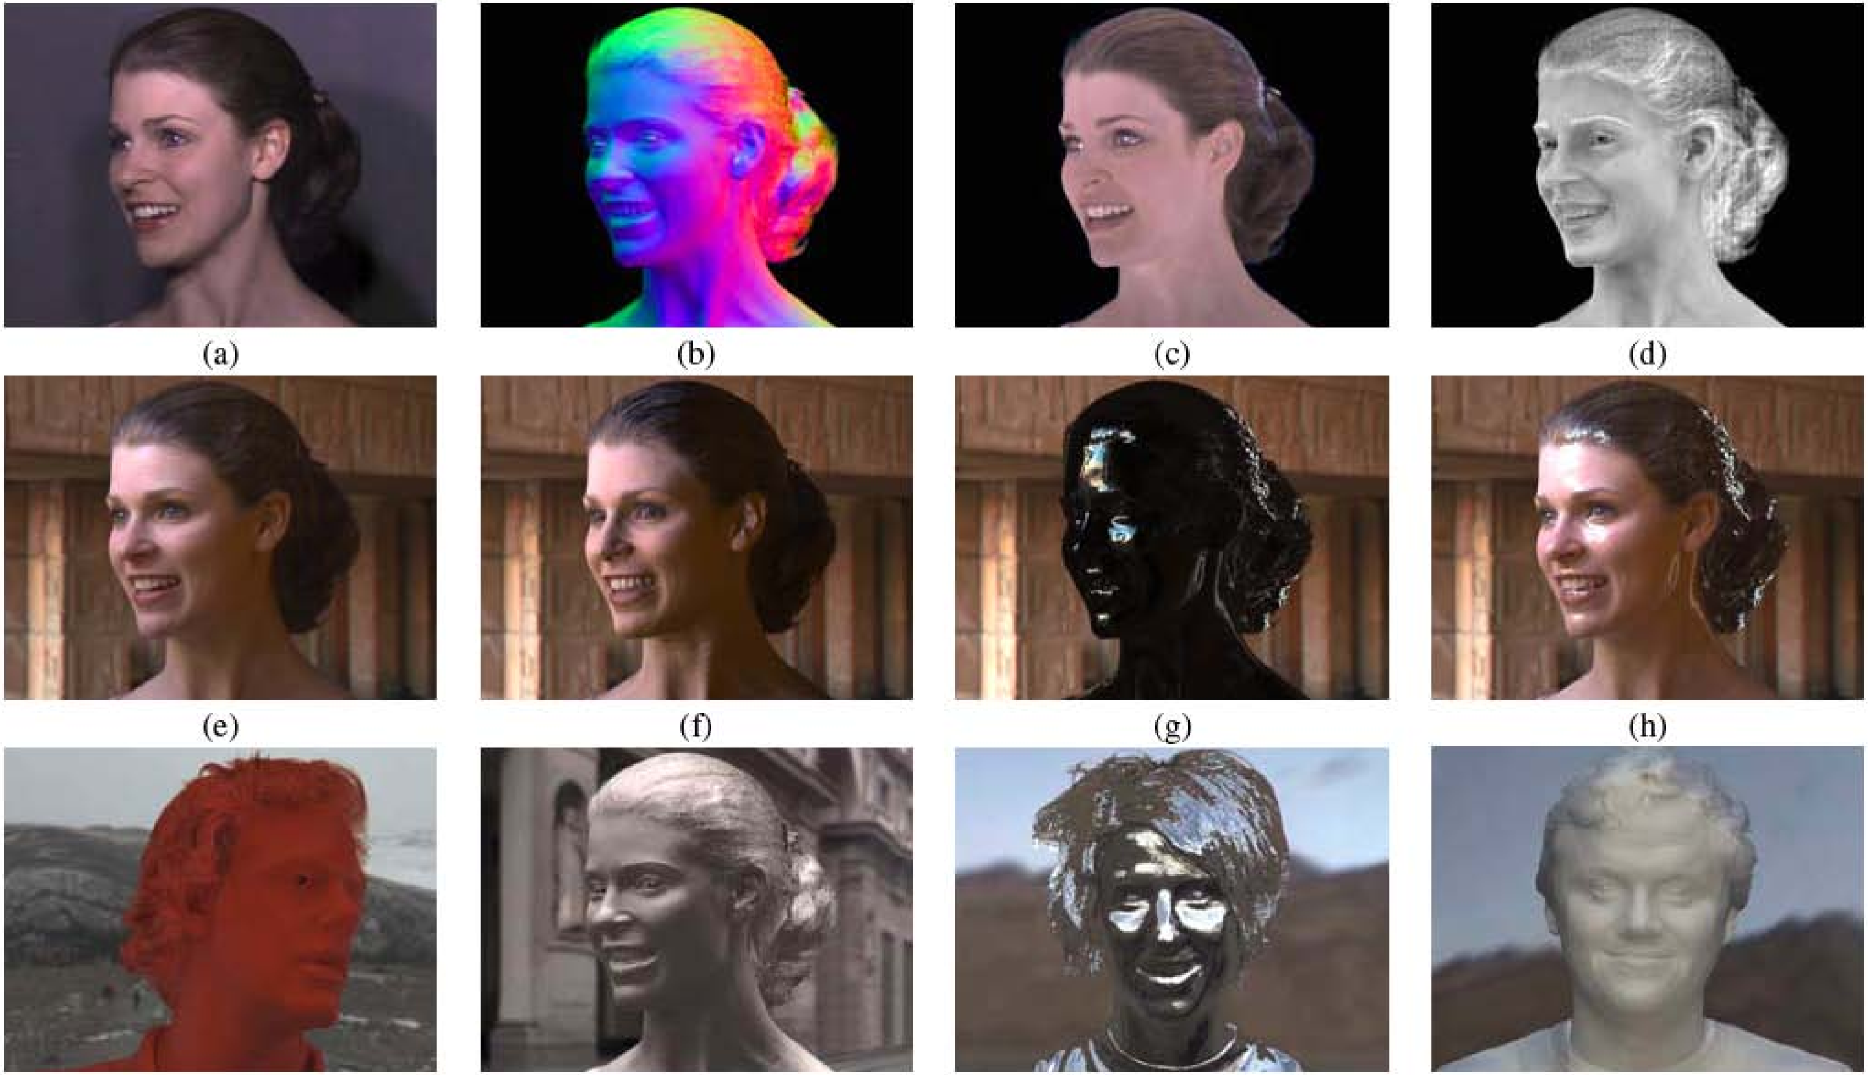

(a) (b) (c) (d)

(e) (f) (g) (h)

(i) (j) (k) (l)

Figure 9: **Reflectance Estimation and Transformation** (a) An original photographed image from the performance with the triangle basis. (b) Estimated surface normals from the full lighting basis (c) Estimated diffuse albedo. (d) Estimated ambient occlusion. (e) Original reflectance re-illuminated by an environment. (f) Diffuse albedo with normals and occlusion illuminated by the environment. (g) Specular reflection in the environment. (h) Specular enhancement. (i) Stylized plastic reflectance with occlusion. (j) Diffuse reflectance with occlusion. (k) Metallic specular reflectance. (l) Diffuse reflectance without occlusion, yielding a translucent appearance.

Dual photography

To appear in the ACM SIGGRAPH 2005 conference proceedings

Dual Photography

Pradeep Sen* Billy Chen* Gaurav Garg* Stephen R. Marschner†
Mark Horowitz* Marc Levoy* Hendrik P. A. Lensch*

*Stanford University †Cornell University

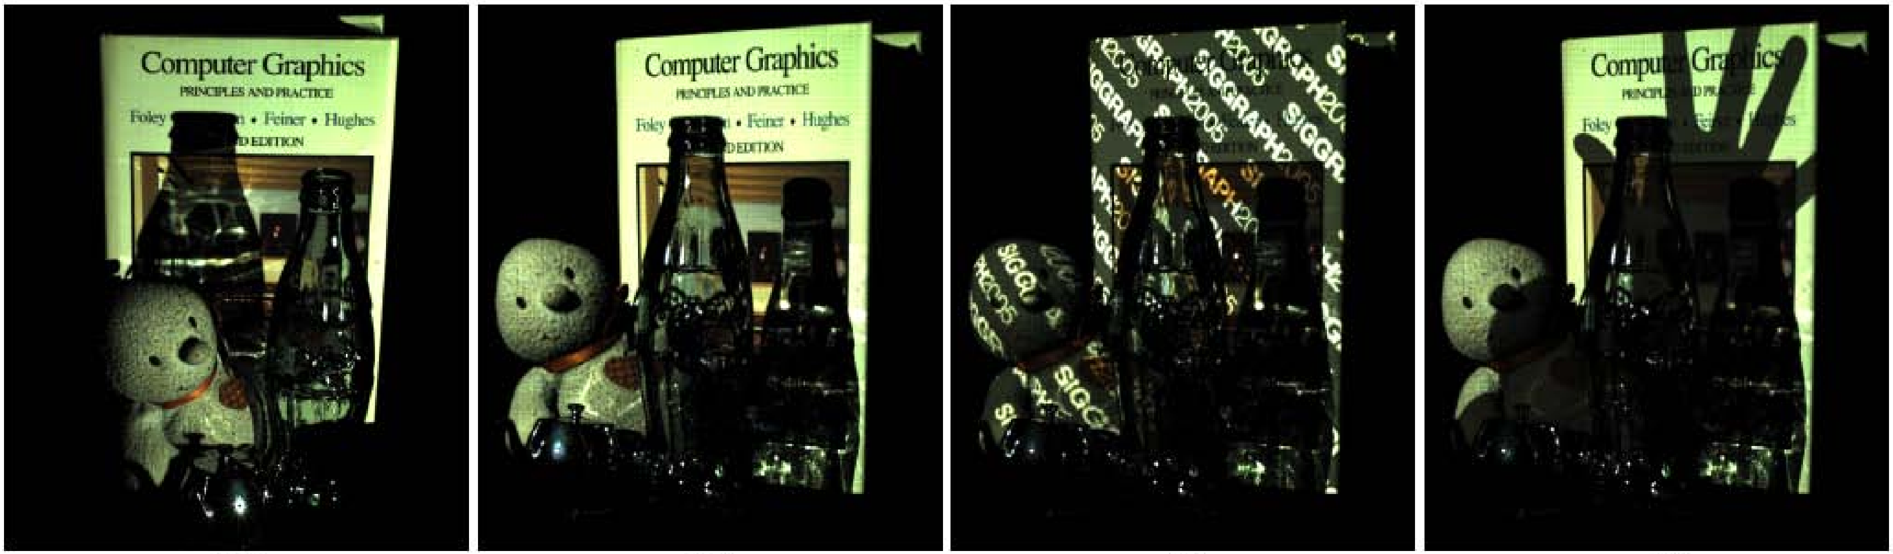

(a) (b) (c) (d)

Figure 1: (a) Conventional photograph of a scene, illuminated by a projector with all its pixels turned on. (b) After measuring the light transport between the projector and the camera using structured illumination, our technique is able to synthesize a photorealistic image from the point of view of the projector. This image has the resolution of the projector and is illuminated by a light source at the position of the camera. The technique can capture subtle illumination effects such as caustics and self-shadowing. Note, for example, how the glass bottle in the primal image (a) appears as the caustic in the dual image (b) and vice-versa. Because we have determined the complete light transport between the projector and camera, it is easy to relight the dual image using a synthetic light source (c) or a light modified by a matte captured later by the same camera (d).

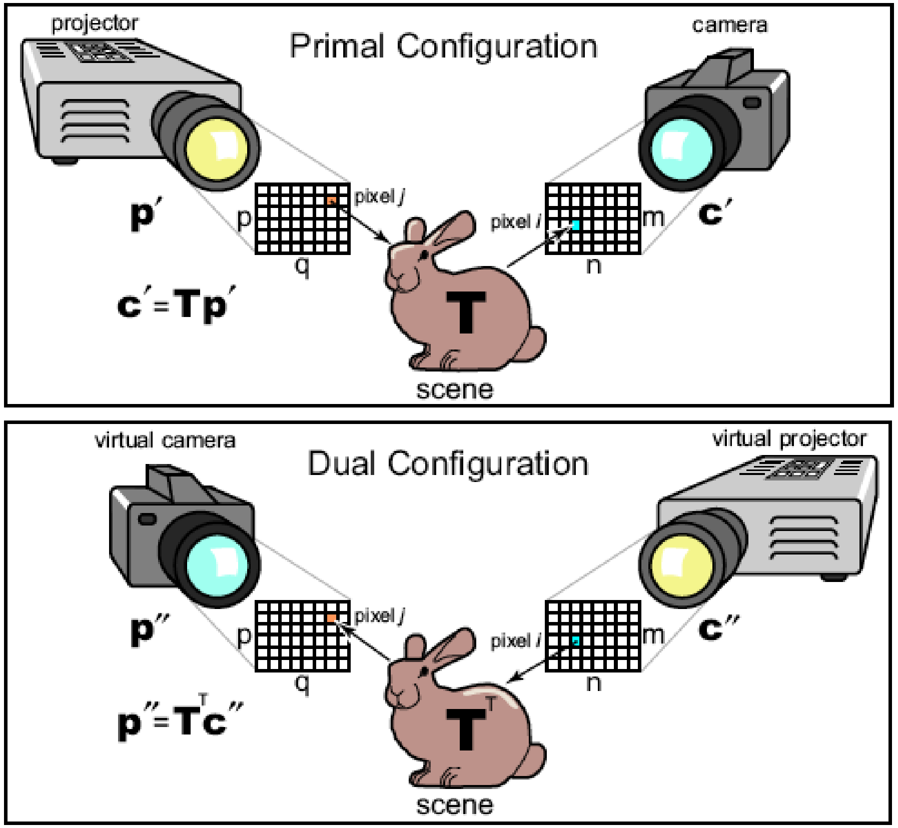

Figure 2: The principle of dual photography. The top diagram shows our primal configuration, with light being emitted by a real projector and captured by a real camera. Matrix $\mathbf{T}$ describes the light transport between the projector and the camera (element $\mathbf{T}_{ij}$ is the transport coefficient from projector pixel j to camera pixel i). The bottom diagram shows the dual configuration, with the positions of the projector and camera reversed. Suppose $\mathbf{T}''$ is the transport matrix in this dual configuration, so that $\mathbf{T}''_{ji}$ is the transport between pixel i of the virtual projector and pixel j of the virtual camera. As shown in Appendix A, Helmholtz reciprocity specifies that the pixel-to-pixel transport is equal in both directions, i. e. $\mathbf{T}''_{ji} = \mathbf{T}_{ij}$, which means $\mathbf{T}'' = \mathbf{T}^T$. As explained in the text, given $\mathbf{T}$, we can use $\mathbf{T}^T$ to synthesize the images that would be acquired in the dual configuration.

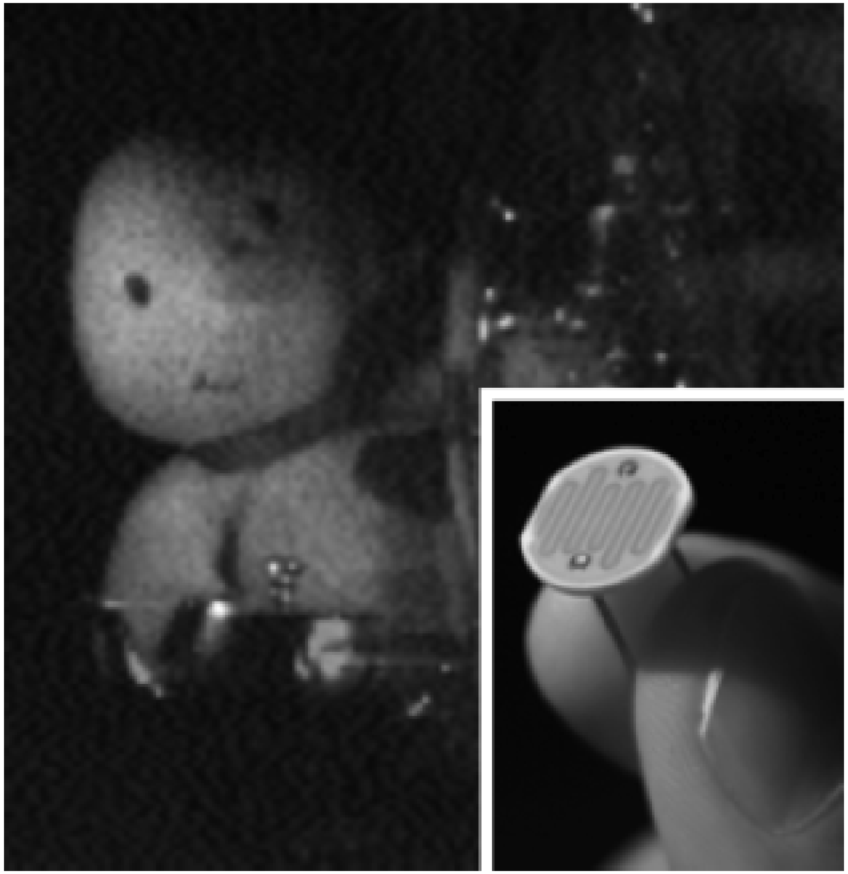

Figure 4: Photography without an imaging sensor. This image was generated using a projector and two photo-resistors like the one shown in the inset. This is a dual image of the scene in Figure 1 and is a view of the scene from the projector's location as illuminated by point light sources at the locations of the two photo-resistors.

Figure 8: Sample scenes. The acquired primal image is on the left, the synthesized dual on the right. Note for example the detail on the pillar in the dual image of the bottom row which is barely visible in the primal due to foreshortening.

(a) (b) (c)

Figure 15: Using a mirror array to emulate multiple virtual projectors. (a) A camera was imaged onto the mirror array in order to emulate multiple virtual light positions. A block pattern scan consisting of 144 high dynamic range images was performed to acquire the scene's transport matrix. The region of interest within the projector's field of view was 864×604 pixels, the final resolution of the dual image. Each camera in the mirror array had an approximate resolution of 800×600 pixels, which is the resulting spatial resolution of our virtual lights. (b) The scene is illuminated by 12 point light sources to create soft shadows. (c) An animated character is embedded in the scene and casts shadows onto the scene.

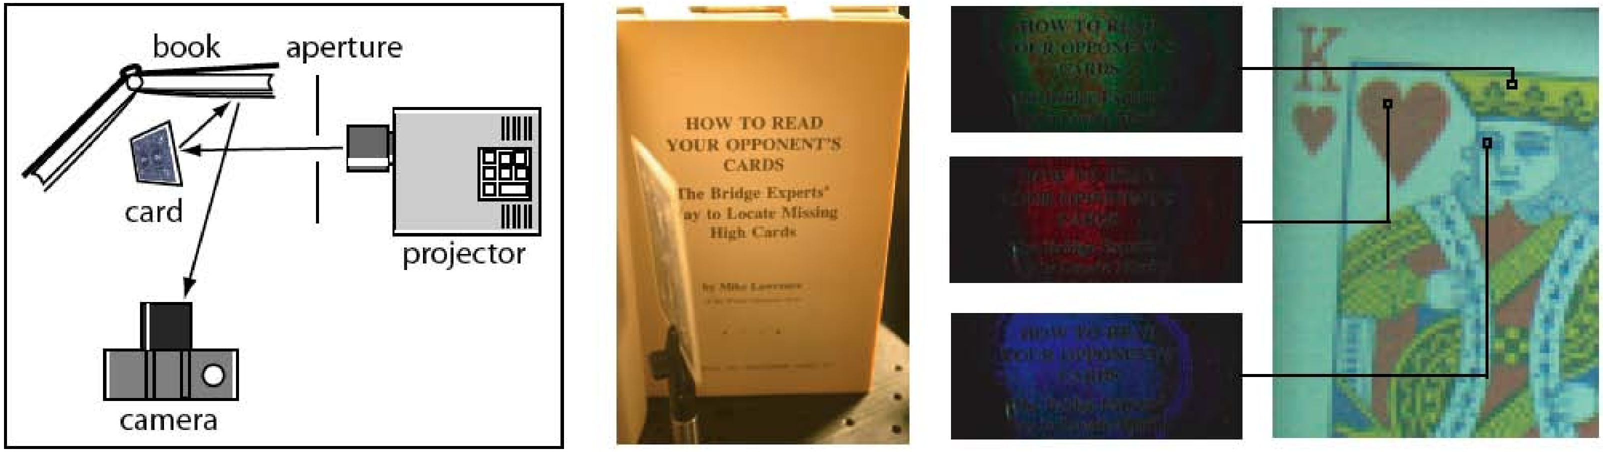

(a) (b) (c) (d)

Figure 16: Dual photography with indirect light transport. (a) A projector illuminates the front of a playing card while the camera sees only the back of the card and the diffuse page of the book. An aperture in front of the projector limits the illumination only onto the card. The card was adjusted so that its specular lobe from the projector did not land on the book. Thus, the only light that reached the camera underwent a diffuse bounce at the card and another at the book. (b) Complete camera view under room lighting. The back of the card and the page of the book are visible. It seems impossible to determine the identity of the card from this point of view simply by varying the incident illumination. To acquire the transport matrix, a 3×3 white pixel was scanned by the projector and 5742 images were acquired to produce a dual image of resolution 66×87. (c) Sample images acquired when the projector scanned the indicated points on the card. The dark level has been subtracted and the images gamma-corrected to amplify the contrast. We see that the diffuse reflection changes depending on the color of the card at the point of illumination. After acquiring the T matrix in this manner, we can reconstruct the floodlit dual image (d). It shows the playing card from the perspective of the projector being indirectly lit by the camera. No contrast enhancement has been applied. Note that the resulting image has been automatically antialiased over the area of each projector pixel.

Separation of global/direct

Fast Separation of Direct and Global Components of a Scene using High Frequency Illumination

Shree K. Nayar[*] Columbia University
Gurunandan Krishnan[†] Columbia University
Michael D. Grossberg[‡] City University of New York
Ramesh Raskar[§] MERL

(a) Scene (b) Direct Component (c) Global Component

Figure 1: (a) A scene lit by a single source of light. The scene includes a wide variety of physical phenomena that produce complex global illumination effects. We present several methods for separating the (b) direct and (c) global illumination components of the scene using high frequency illumination. In this example, the components were estimated by shifting a single checkerboard pattern 25 times to overcome the optical and resolution limits of the source (projector) and sensor (camera). The direct and global images have been brightness scaled by a factor of 1.25. In theory, the separation can be done using just 2 images. When the separation results are only needed at a resolution that is lower than those of the source and sensor, the separation can be done with a single image.

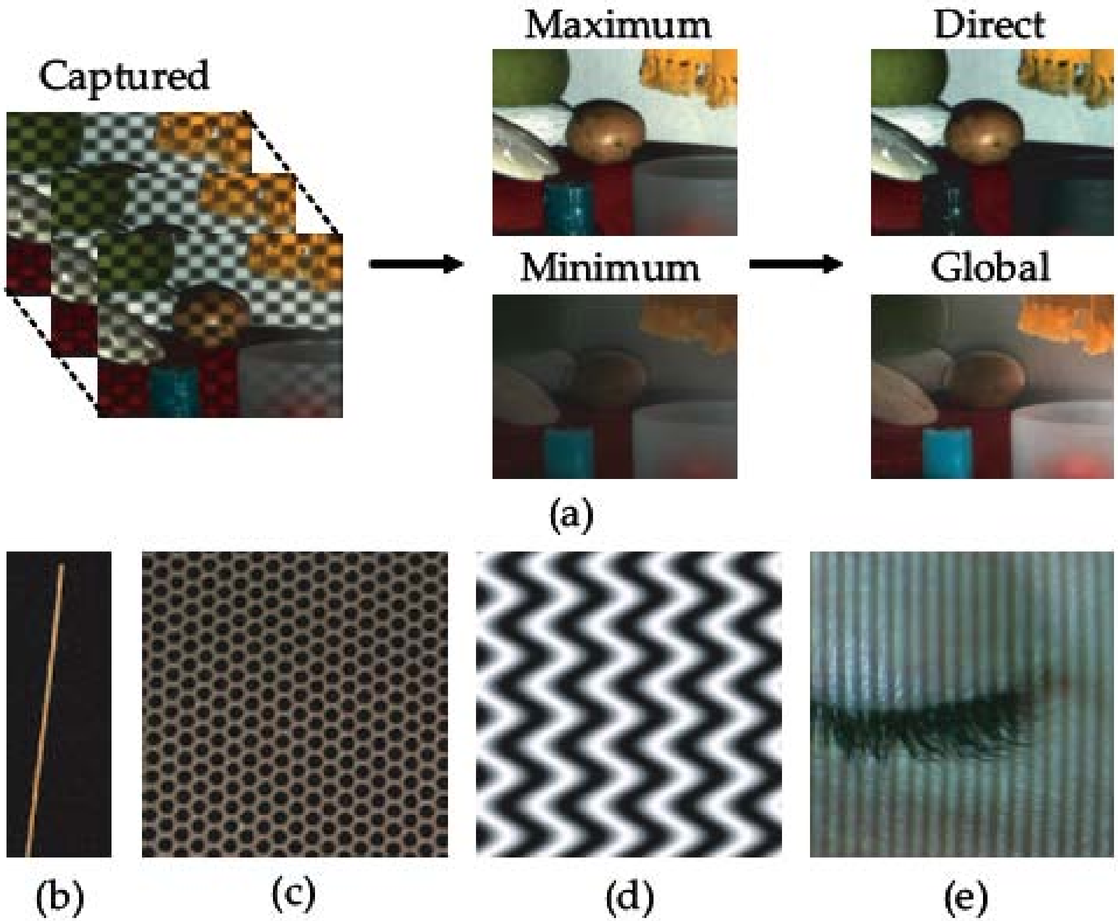

Figure 5: (a) The steps involved in the computation of direct and global images using a set of shifted checkerboard illumination patterns. (b) The line occluder (stick) used to scan scenes lit by a simple source such as the sun. (c) The mesh occluder used to expedite the scanning process. (d) Three shifted versions of this sinusoid-based illumination pattern are sufficient to perform the separation. (e) A magnified part of the face image in Figure 7(d) reveals the stripe pattern used to do the separation from a single image.

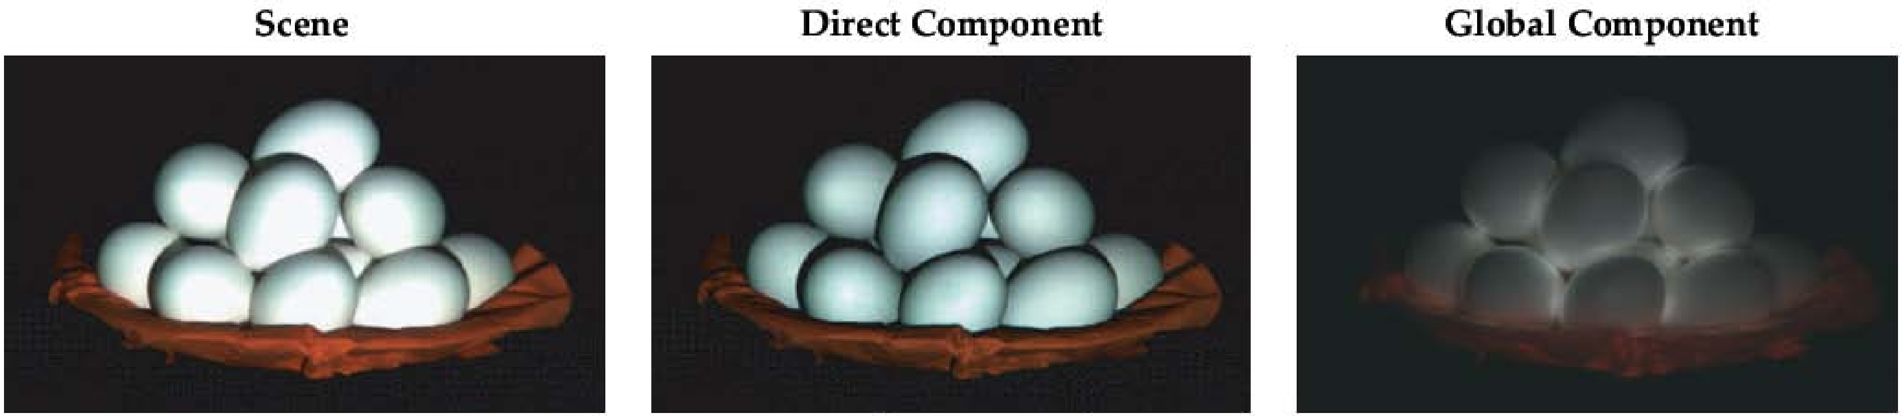

(a) Eggs: Diffuse Interreflections

(b) Wooden Blocks: Diffuse and Specular Interreflections

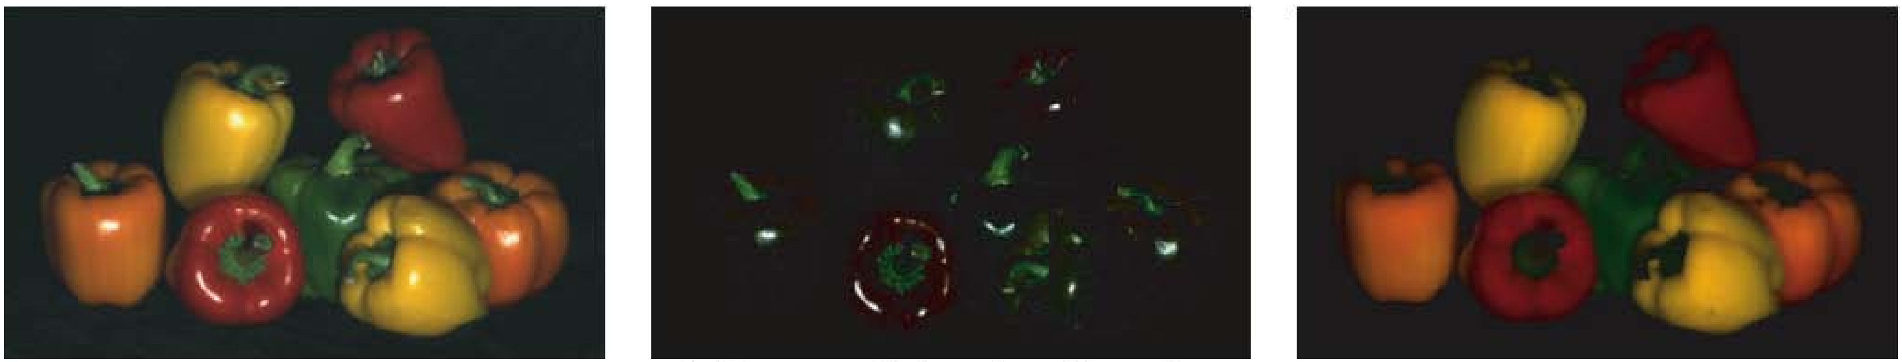

(c) Peppers: Subsurface Scattering

(d) Grapes and Cheese: Subsurface Scattering

(e) Kitchen Sink with Milky Water: Volumetric Scattering

(f) Novel Images

References

- **http://www.debevec.org/Publications/DEBEVEC-IEEE-Computer-200608-high.pdf**