ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Automatic and accurate shadow detection using near-infrared information

http://ivrg.epfl.ch

C. Fredembach, D. Rüfenacht and S. Süsstrunk

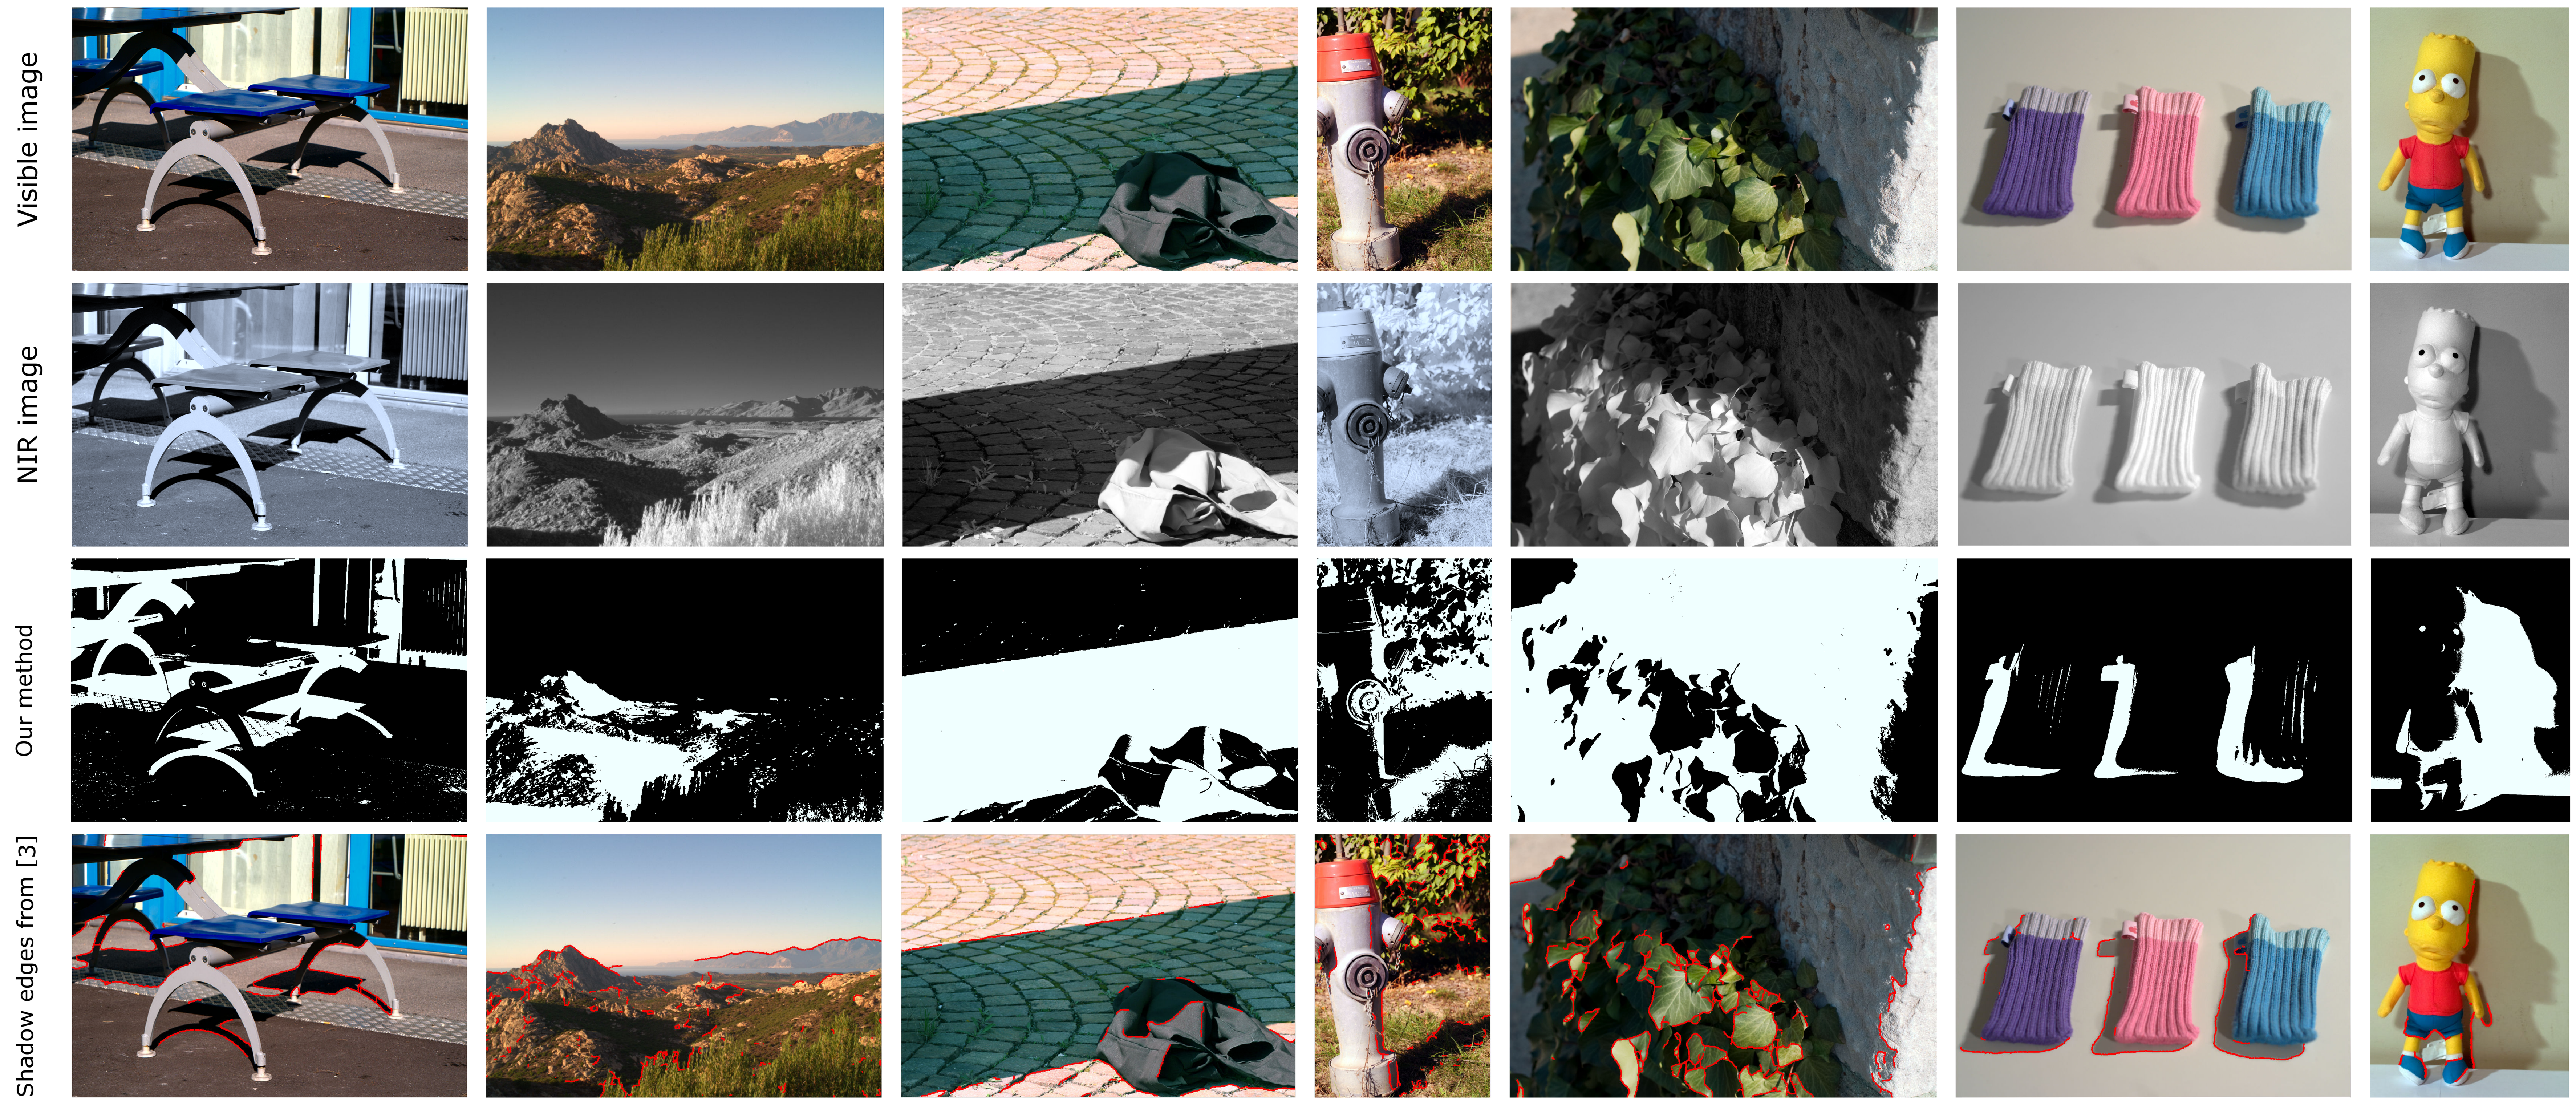

Figure 1: Our method automatically and accurately detects shadows by using both visible and near-infrared (NIR) images. We first compute an initial shadow candidate map by only considering the dark areas of both scene representations. The results are refined by incorporating color to NIR ratios, which contain valuable information because of the very distinct spectra of commonly shadow creating light sources in the NIR band. The simple computations combined with a good heuristic allow us to obtain high quality shadow maps that can be computed in real-time. First row: color images. Second row: NIR images. Third row: Our shadow masks. Fourth row: Shadow edges from [3].

Proposed Framework [1]

We present a method to automatically detect shadows in a fast and accurate manner by employing the inherent sensitivity of digital camera sensors to the near-infrared (NIR) part of the spectrum. Shadows are generally found in the dark parts of an image, be it color or NIR. We consider these dark pixels as shadow candidates. By observing that commonly encountered light sources have very distinct spectra in the NIR, we propose that the ratios of the color channels (red, green and blue) to the NIR give valuable information about impinging illumination, which we employ to assess the shadow candidate pixels.

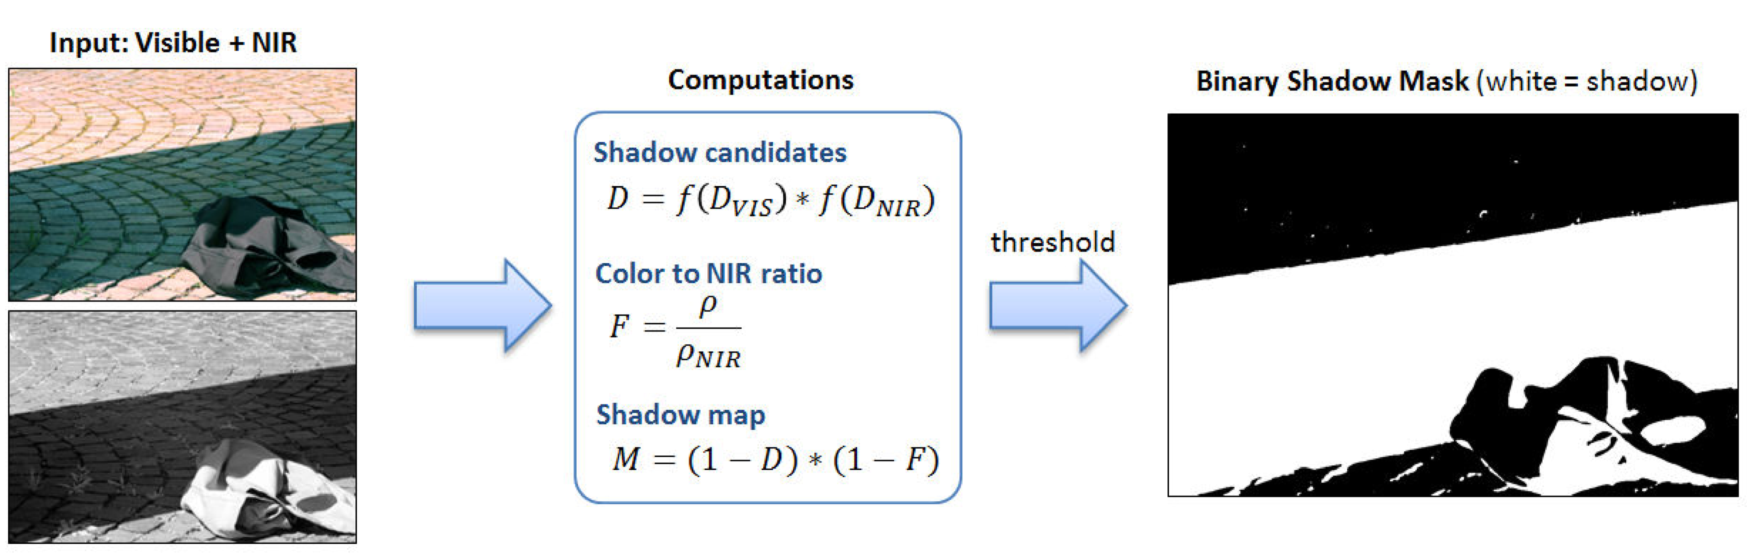

Figure 2: *Our proposed framework.*

Shadow Candidates Selection

Spectral studies of man-made surfaces and colorants show that they are in general much darker in the visible part of the spectrum than in the NIR.

$$\int_{VIS} S(\lambda) Q_{R,G,B}(\lambda) d\lambda < \int_{NIR} S(\lambda) Q_{NIR}(\lambda) d\lambda$$

This permits to disambiguate a number of otherwise problematic dark objects/surfaces. The dark maps used to identify shadow pixel candidates are computed as follows. Let ρ_k be the normalized camera sensor response, for $k \in \{R, G, B, NIR\}$. We define two temporary dark maps as follows:

$$D_{VIS} = 1 - \frac{\rho_R + \rho_G + \rho_B}{3} \,;\, D_{NIR} = 1 - \rho_{NIR}$$

To compress the dark (and white) pixels, we apply a sigmoid function f to both D_{VIS} and D_{NIR}. The shadow candidate map is then computed as the product of the two dark maps:

$$D = f(D_{VIS}) f(D_{NIR})$$

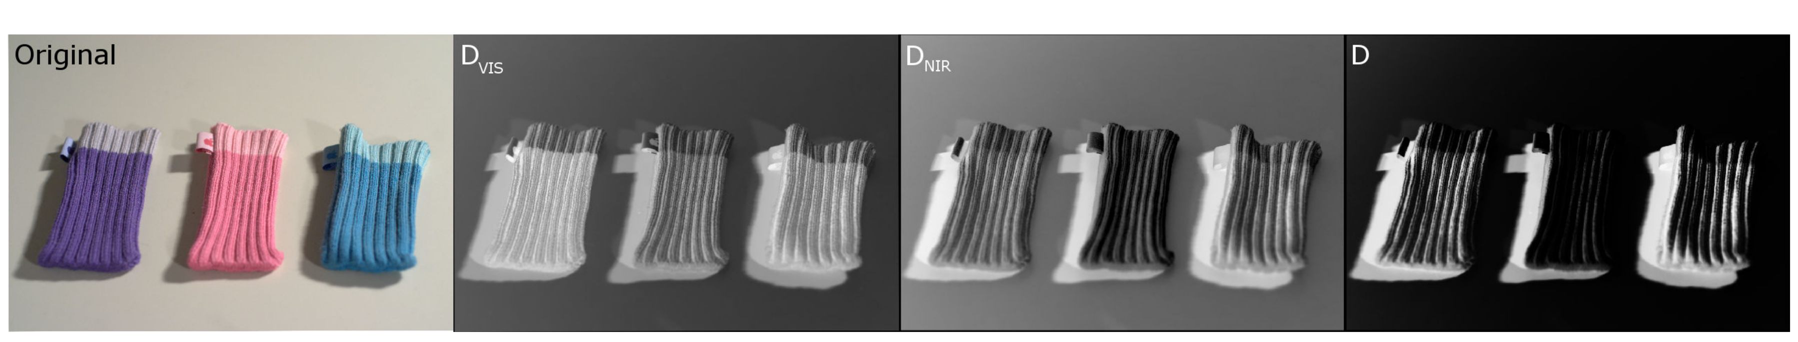

Figure 3: *While the presence of dark objects confounds D_{VIS}, D is quite accurate.*

Color to NIR ratios

The key reason of using color to NIR ratios is illustrated in Figure 4. We see for example that skylight emits very little in the NIR, while daylight emits approximately as much in the NIR as in the visible band. We therefore compute color to NIR ratios, which will have a significant impact on shadow detection. We compute a ratio image F as:

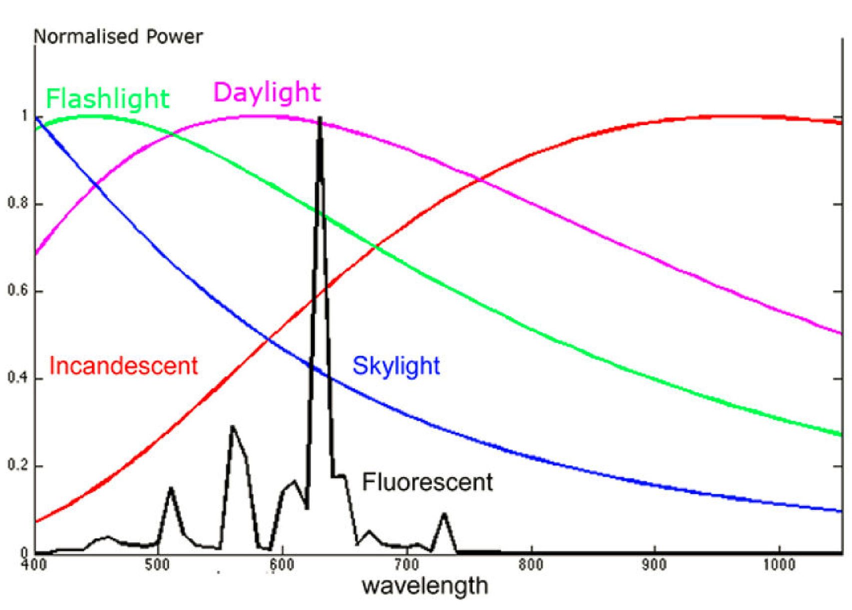

Figure 4: *Measured spectra of typical shadow-creating illuminants.*

$$F = \frac{1}{2} \min(\max(\frac{\rho_k}{\rho_{NIR}}), 2), k \in \{R, G, B\}$$

To obtain the final shadow map, we take into account both the shadow candidates from D, and the ratio image F. Since both D and F have comparable values, we simply compute $M = (1-D)(1-F)$.

To obtain the final shadow mask, we need to binarize M. According to Lischinski *et al.* [2], we compute the histogram of M and calculate the location of its first valley. Let us denote this location as θ. The binary shadow value of each pixel x is then given by:

$$M_{bin}(x) = \begin{cases} 1 & \text{if } M(x) \le \theta \\ 0 & \text{otherwise} \end{cases}$$

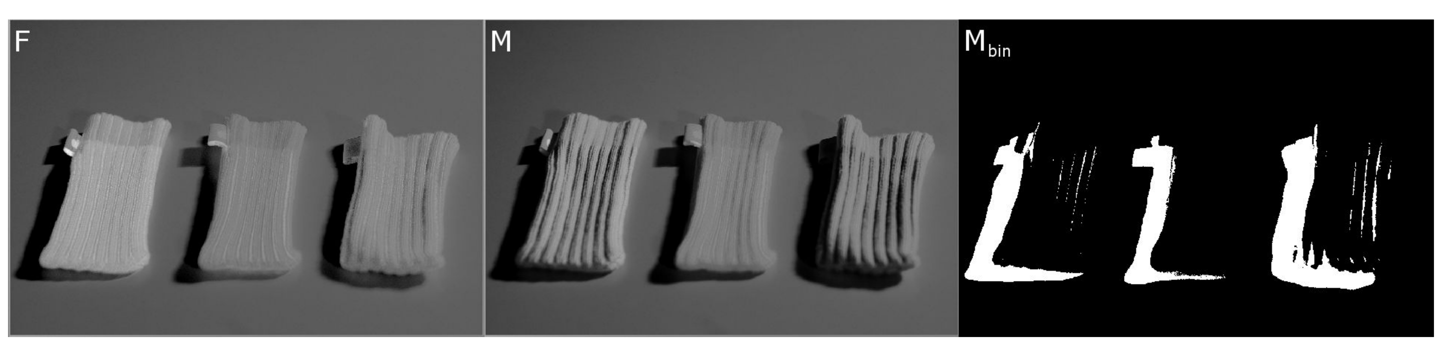

Figure 5: *Color to NIR ratio map F, shadow map M, and binary shadow mask M_{bin}.*

Results

The results are shown in Figure 1. For comparison, we have included the results from the shadow detection algorithm of Lalonde *et al.* [3], which finds shadow edges (red lines). One can appreciate the accuracy of the computed shadow masks by our method. The entire shadow detection procedure consists in a pixel-wise division and multiplication, followed by a histogram-based thresholding, which makes it very fast.

References

[1] C. Fredembach and S. Süsstrunk. *Automatic and accurate shadow detection from (potentially) a single image using near-infrared information.* EPFL Tech Report 165527, 2010.
[2] Y. Shor and D. Lischinski. *The shadow meets the mask: Pyramid-based shadow removal.* Eurographics, 2008.
[3] J-F. Lalonde, A. Efros and S. Narasimhan. *Detecting ground shadows in outdoor consumer photographs.* European Conference on Computer Vision, 2010.