Reflectance Normalization in Illumination-Based Image Manipulation Detection

Dr.-Ing. Christian Riess, Sven Pfaller, Elli Angelopoulou, Ph.D.
September 7th, 2015

Pattern Recognition Lab (CS 5)
University of Erlangen-Nuremberg, Germany

FRIEDRICH-ALEXANDER UNIVERSITÄT ERLANGEN-NÜRNBERG
TECHNISCHE FAKULTÄT

Illumination Direction as a Forensic Cue

- Idea by Johnson and Farid 2005 and 2007 [1]:
  - From intensities and surface normals one can solve for the position of the light source, (under some assumptions)
  - At occluding contours, surface normals are 2-D vectors in image plane
  - Collect intensities and 2-D normals along an object contour
  - Calculate a 2-D projection of the incident light direction
  - If direction is inconsistent across objects, the image is manipulated.

Surface normals within objects are 3-D (and hard to estimate from a 2-D image):

Surface normals at occluding contours lie in the 2-D image plane. Estimation is simpler: just estimate the contour normal.

[1] M. K. Johnson and H. Farid: Exposing Digital Forgeries in Complex Lighting Environments, IEEE Transactions on Information Forensics and Security, vol. 2, no. 3, Sep. 2007, pp. 450-461.

Example Images (Johnson/Farid 2007)

Pictures from: M. K. Johnson and H. Farid: Exposing Digital Forgeries in Complex Lighting Environments, IEEE Transactions on Information Forensics and Security, vol. 2, no. 3, Sep. 2007, pp. 450-461.

Limitations of the Johnson/Farid Approach

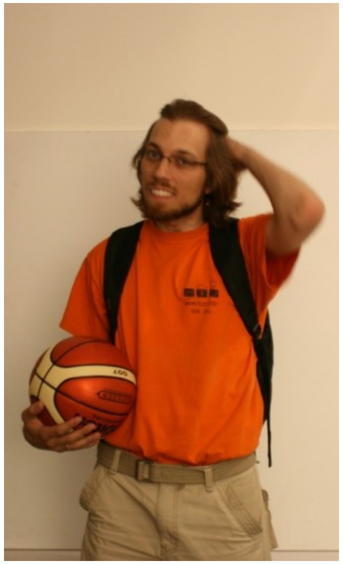

- Lighting direction is determined by solving a system of linear equations, with the assumptions of
  1) Lambertian reflectance on contours
  2) Infinitely distant light source
  3) Wide angular coverage by the contours (about 130°)
  4) Identical color (and no texture) on contours
  5) No self-shadowed (concave) contours

- In practice, constraint 3) oftentimes conflicts with constraints 4) and 5)

In this work, we relax constraint 4) by normalizing the color of the contour

Ideal solution: Intrinsic Image Decomposition

- Model: $\text{Observation} = \text{Reflectance} \cdot \text{Shading}$

- Example (from the MIT database [2])

Picture Reflectance Shading

Shading is free from color differences and texture…
->
This is the best possible input for estimating the lighting direction!

[2] R. Grosse, M. K. Johnson, E. Adelson and W. T. Freeman: A Ground-Truth Dataset and Baseline Evaluations for Intrinsic Image Algorithms, IEEE International Conference on Computer Vision, 2009, pp. 2335-2342.

Prior Work on Intrinsic Image Decomposition

- We implemented/experimented with two recent methods:
  - Shen et al. [3]: Wavelet decomposition, followed by a factorization with a sparsity constraint on the reflectance
  - Gehler et al. [4]: Energy minimization based on the assumptions that shading varies smoothly and scene consists of only few colors
- Both approaches achieve very good result on the controlled MIT laboratory data set
- Sadly, we failed to transfer these results to our real-world data

[3] L. Shen, P. Tan, and S. Lin: Intrinsic Image Decomposition using a Local and Global Sparse Representation of Reflectance, IEEE Computer Vision and Pattern Recognition, 2011, pp. 697-704.
[4] P. V. Gehler, C. Rother, M. Kiefel, L. Zhang, and B. Schölkopf:: Recovering Intrinsic Images with a Global Sparsity Prior on Reflectance, Advances in Neural Image Processing, 2011, pp. 765-773.

Intrinsic Contours instead of Intrinsic Images

- For our task, we actually only require **„intrinsic contours"**
- In other words, 1-D processing suffices:
  - Sort the normals by angle,
  - Cluster the pixels by color,
  - Multiply cluster brightness by a factor, such that normals from different clusters with the same angle exhibit the same brightness

Intrinsic Contours instead of Intrinsic Images

- We use k-means for clustering

- Let $w(\nu_1, \nu_2) = \exp\left(\frac{\cos^{-1}(\nu_1^{\mathrm{T}}\nu_2)}{\sigma}\right)$

  a Gaussian similarity between two normals ν_1 and ν_2.

- Normalization factors $\mathbf{r}$ are the solution of

$$\mathbf{W}\mathbf{r} = \mathbf{0}$$

where $W_{j,c_{\{1,2\}}} = w(\nu_1, \nu_2)p(\nu_{\{1,2\}})$ or 0,

$p(\nu_{\{1,2\}})$ is the pixel intensity,
$c_{\{1,2\}}$ the cluster index,
and the first entry of $\mathbf{r}$ is set to 1.

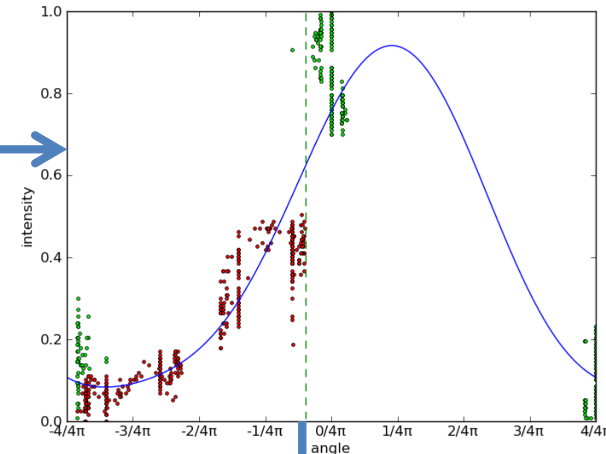

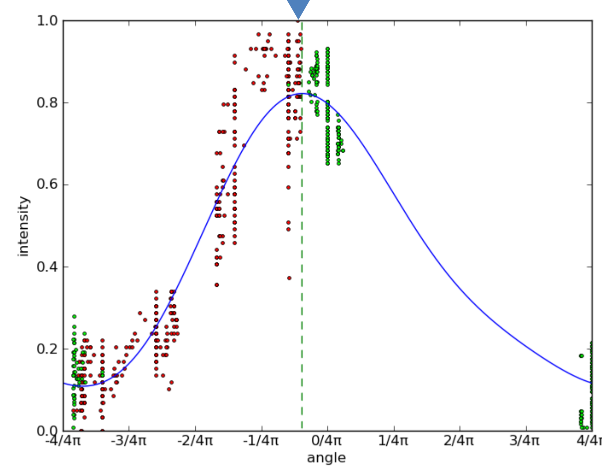

Capturing Ground-Truth Data

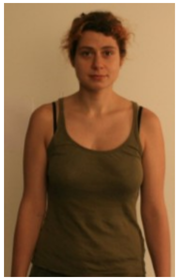

- We captured 10 individuals in a room with controlled lighting
- Direct illumination was incident at 0°, 45°, 90° (close distance, not „infinitely“ far away
- Indirect illumination was either „on“ or „off“
- Contours are manually annotated

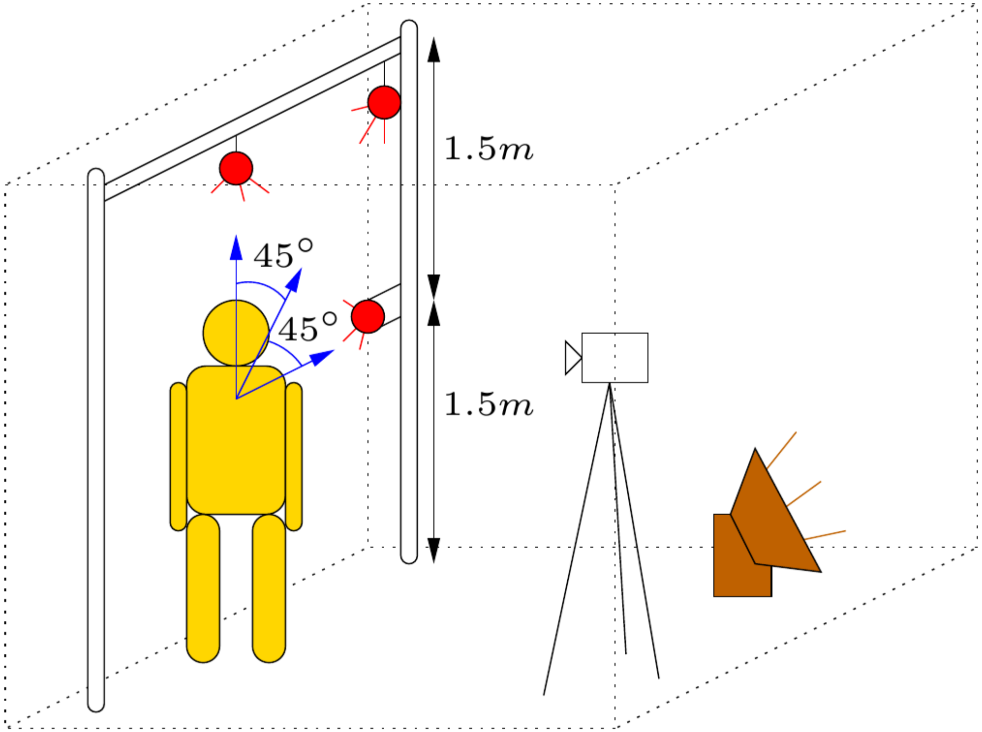

Quantitative Evaluation

- We compared the proposed normalization to the Johnson/Farid method („Original“) and to the preprocessing by Gehler et al.
- The proposed method gently incorporates multi-colored contours.

Single-colored contour

	Median [°]	Mean [°]	Within 22.5°
Original	10.7	13.6	25/30
Gehler et al.	**9.1**	**12.5**	**26/30**
Proposed	10.9	14.1	24/30

Multi-colored contour

	Median [°]	Mean [°]	Within 22.5°
Original	40.2	56.5	10/30
Gehler et al.	33.0	50.7	13/30
Proposed	**12.6**	**13.0**	**26/30**

Open Issues / Failure Cases

- SNR in dark pixels is significantly worse
- Noise is accordingly upscaled, rendering the estimate less reliable

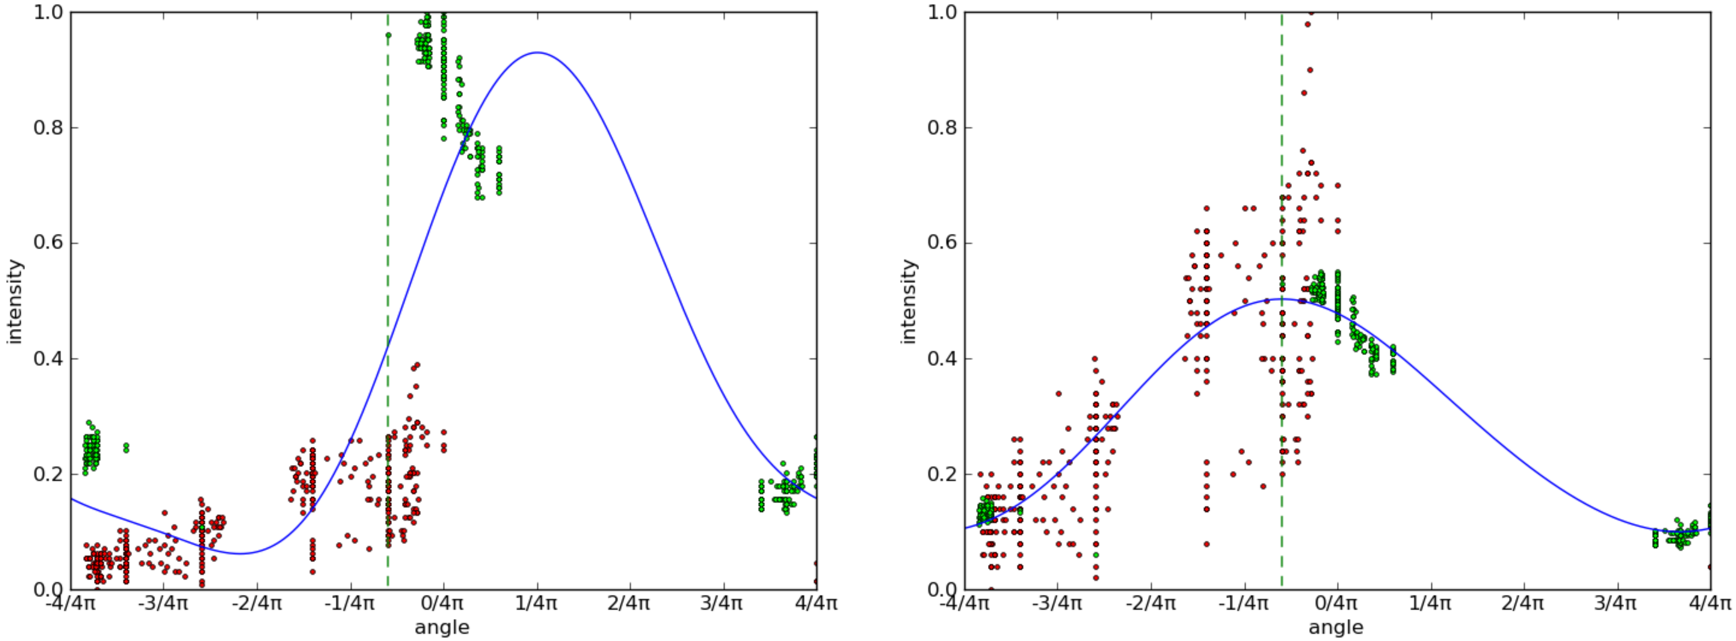

Summary

- 2-D estimation of the lighting direction is pretty, but suffers from many restrictive constraints

- In this work, we examined intrinsic image decomposition to relax the method‘s dependency on single-colored surfaces

- It turned out that a direct brightness normalization along the contours works best
- For evaluation, we created and annotated benchmark data with 10 individuals illuminated under 3 angles
- In future work, we will investigate the robustness of the algorithm more closely

Thank you for your attention.