Towards the Estimation of Non-Uniform Illumination in Real-World Scenes

David Bernecker
with Christian Riess, Sven Pfaller and Elli Angelopoulou

September 27, 2013
Pattern Recognition Lab, University of Erlangen-Nuremberg

FRIEDRICH-ALEXANDER
UNIVERSITÄT
ERLANGEN-NÜRNBERG
TECHNISCHE FAKULTÄT

Illumination Changes Color Appearance

Daylight | Energy saving light | Tungsten light | Yellow light

- With fixed camera settings, one can readily observe that: change in illuminant color → change in object appearance
- These variations make color-based image processing challenging

Color Constancy: Separate Illumination and Material

- Assume: with known illumination, color correction is straightforward
- Let $\mathbf{i} = (i_\mathrm{R} \;\; i_\mathrm{G} \;\; i_\mathrm{B})^\mathrm{T}$ be the illumination color, and $\mathbf{p} = (p_\mathrm{R} \;\; p_\mathrm{G} \;\; p_\mathrm{B})^\mathrm{T}$ the observed pixel color
- Then, the color-corrected pixel is $\mathbf{p}' = \left(\frac{p_\mathrm{R}}{i_\mathrm{R}} \;\; \frac{p_\mathrm{G}}{i_\mathrm{G}} \;\; \frac{p_\mathrm{B}}{i_\mathrm{B}} \right)^\mathrm{T}$
- Hence, we „only“ need an estimate of the illuminant color to normalize an image

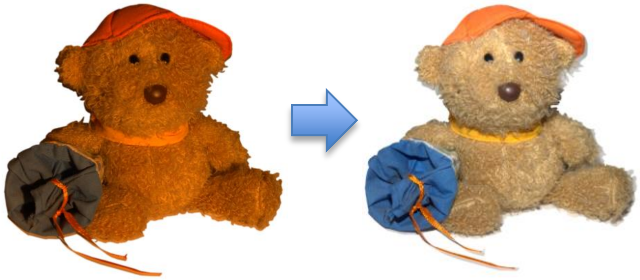

- Many estimators have been proposed, e.g.
  - Gray world
  - Gamut mapping
  - Bayesian color constancy

Example Approach: Gray World / Gray Edge

- Gray world hypothesis: sum of pixel or edge colors is illumination

$$\boldsymbol{l} = k \left(\int \left| \frac{\partial^n \boldsymbol{f}^\sigma(\boldsymbol{x})}{\partial \boldsymbol{x}^n} \right|^p d\boldsymbol{x} \right)^{\frac{1}{p}}$$

Parameters:
- Derivative order n
- p-norm
- Standard deviation of Gaussian smoothing σ

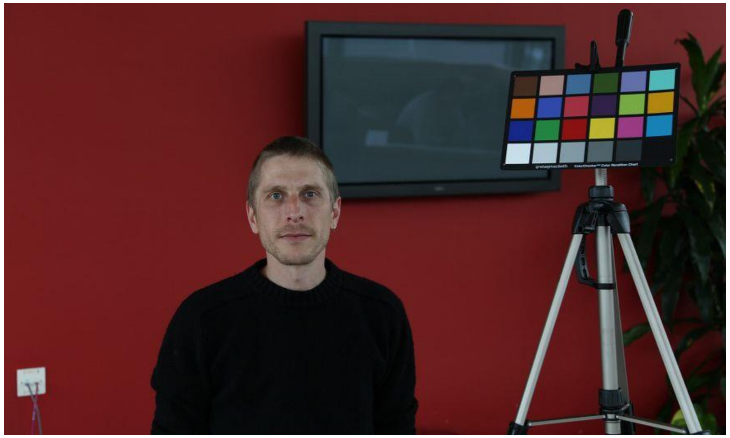

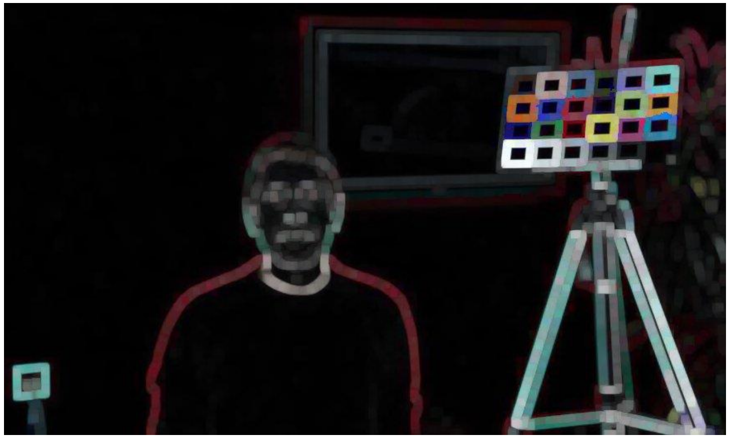

FRIEDRICH-ALEXANDER
UNIVERSITÄT
ERLANGEN-NÜRNBERG
TECHNISCHE FAKULTÄT

Limitation: Assumption of Globally Uniform Illumination

- Classical color constancy algorithms are designed to estimate a single illuminant color for the whole image

- However, many real-world scenes consist of two or more illuminants:

Flash / ambient light

Shadows / direct light

Complex lighting

- **The goal of our research is to estimate local illumination**

Our Prior Work on Local Illuminant Estimation

- Bleier et al. (CPCV 2011):
  scale down spatial support of existing illuminant estimators

Normalization and Pre-processing (e.g., sRGB -> linear RGB) → Superpixel Segmentation → Illuminant Estimation → Post-processing and Re-estimation

Training (↑ Illuminant Estimation)

Input image

Superpixels segmentation

Gray World

What is Wrong with this Approach?

- Using statistical estimators on smaller image patches works, to some extend

- However, there are some serious drawbacks:
  - Segments are treated independently of each other
  - On average, smaller patches lead to larger estimation errors
  - No problem-specific knowledge is used, i.e., we only use an ensemble of „dumb“ global estimators

- When estimating local illuminant colors, we have to solve two joint subproblems:
  - Estimate a **set of illuminant colors** in the scene
  - Estimate the **spatial distribution** of the illuminant colors

Contribution of this Study

- We aim towards decomposing the scene based on its illuminant color distribution

- Assume we have a set of local illuminant estimates, then we seek a **selection of a set of candidate illuminants**
- This selector is heuristically chosen to limit the total number to **only 2 or 3 scene illuminants**
- Each local estimate is then recolored to the closest candidate illuminant

- Findings of this work lead to a recently developed, more complete illuminant color estimator that uses a Conditional Random Field

Algorithm Overview, Input and Output

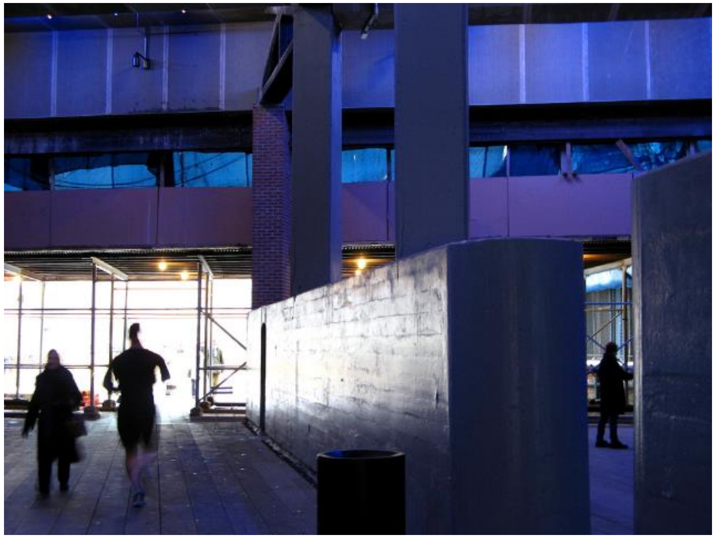

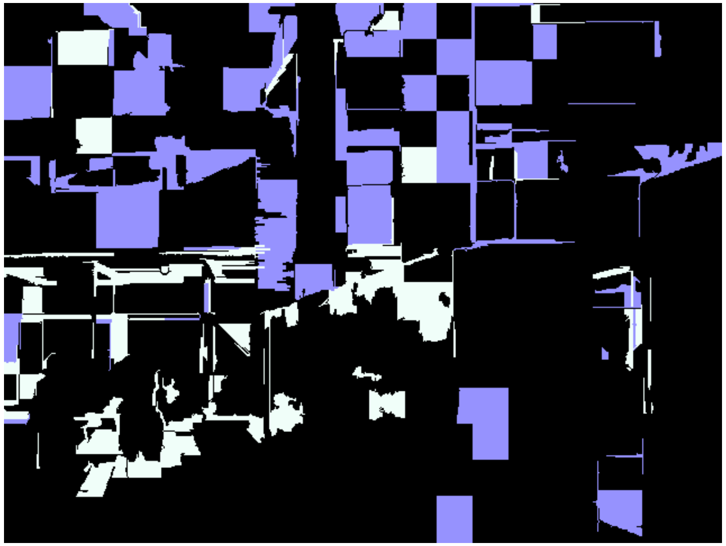

- An input image is
  - Segmented
  - Local illuminant colors are computed (black: little confidence in the estimate)
  - From these colors, few candidates are selected
  - The segments are recolored according to the most similar candidate

Segmentation of the Scene

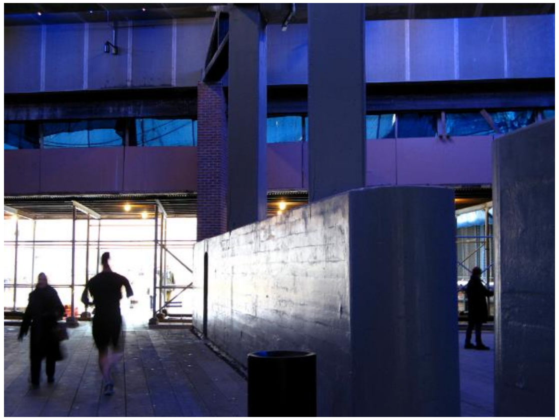
Input image

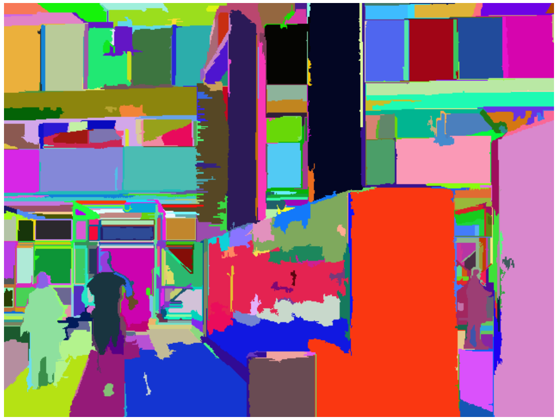
Superpixels


Superpixels w/ grid

- The scene is subdivided in „superpixels“, i.e., areas of similar color (with the graph-based algorithm by Felzenszwalb and Huttenlocher)

- The resulting superpixels are intersected with a grid, to obtain regions of approximately the same size

Local Illuminant Estimation

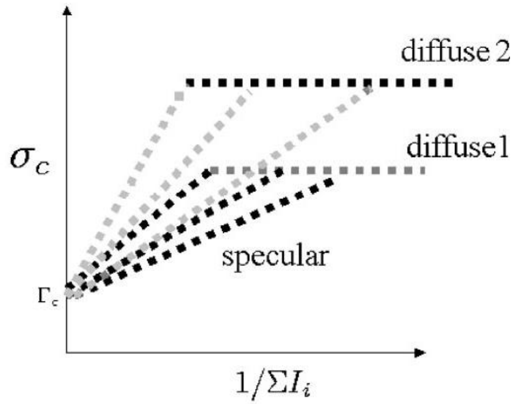

- For each segment, the illuminant color is estimated
- We use a physics-based approach, exploiting the inverse-intensity chromaticity (IIC) space by Tan, Nishino and Ikeuchi (2004):
  - For every color channel, project every pixel to

$$p_c \rightarrow \left(\frac{1}{p_R + p_G + p_B}, \quad \frac{p_c}{p_R + p_G + p_B} \right)$$

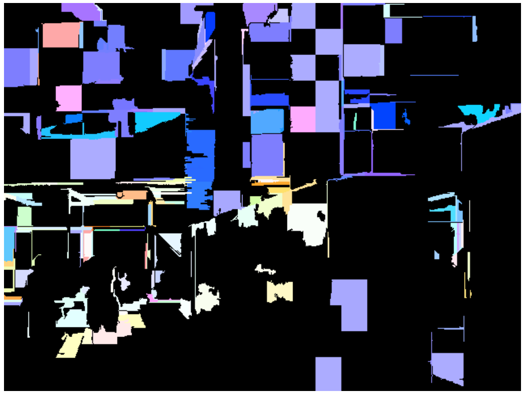

  - Partially specular pixels form a triangular shape.
  - The tip of the triangle intersects the y-axis
  - The position on the y-axis is the illuminant color estimate for color c
  - We added some shape constraints in prior work to suppress outliers

Computation of the Segment Confidence

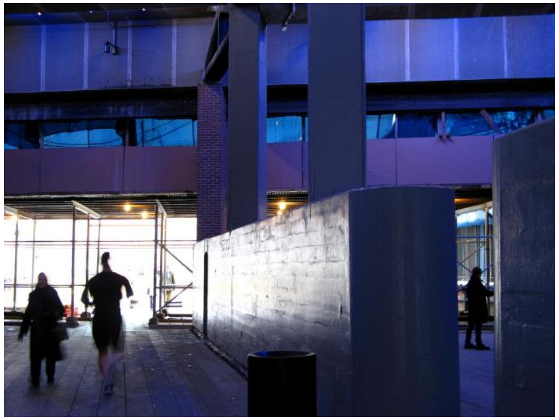

Input image

Confidence map

- IIC space relies on some specular reflection in the segment
- Segments with purely diffuse reflection are likely to produce wrong estimates
- We use the specularity segmentation by Tan and Ikeuchi (2005) to estimate the amount of specular pixels per segments.

Selecting Candidate Illuminants

- Every per-segment illuminant estimate is a potential scene illuminant
- The goal is to select two or three scene illuminants from them

- To do so, we define a distance function between two illuminants $\mathbf{i}_s$ and $\mathbf{q}$ as

$$d(\mathbf{i}_s, \mathbf{q}) = \begin{cases} 1 - 9 \cdot a(\mathbf{i}_s, \mathbf{q}) & \text{if } a(\mathbf{i}_s, \mathbf{q}) \leq 0.1 \\ 0.1 - (a(\mathbf{i}_s, \mathbf{q}) - 0.1)/9 & \text{otherwise} \end{cases}$$

where $a(\mathbf{i}_s, \mathbf{q}) = \cos^{-1}(\mathbf{i}_s \circ \mathbf{q})$ is the angular distance between two RGB-vectors

- $d(\mathbf{i}_s, \mathbf{q})$ is just a piecewise-defined linear function of angular distance

Selecting Candidate Illuminants

- The objective function is then

$$Q_{\text{opt}} = \underset{Q}{\text{argmin}} \sum_{s \in S} \left(1 - \sum_{\mathbf{q} \in Q} d(\mathbf{i}_s, \mathbf{q})\right)^2 c_s$$

  where
  - S is the set of all segments,
  - Q is the set of all illuminant estimates,
  - c_s is a linear weighting factor from the confidences, and
  - $d(\mathbf{i}_s, \mathbf{q})$ is the distance between two illuminants, as defined on the previous slide
- Q_{opt} rewards a selection of illuminant candidates, such that a maximum number of segment estimates is close to exactly one illuminant candidate
- We constrain the cardinality of Q_{opt} to two or three

Qualitative Results

Qualitative Results

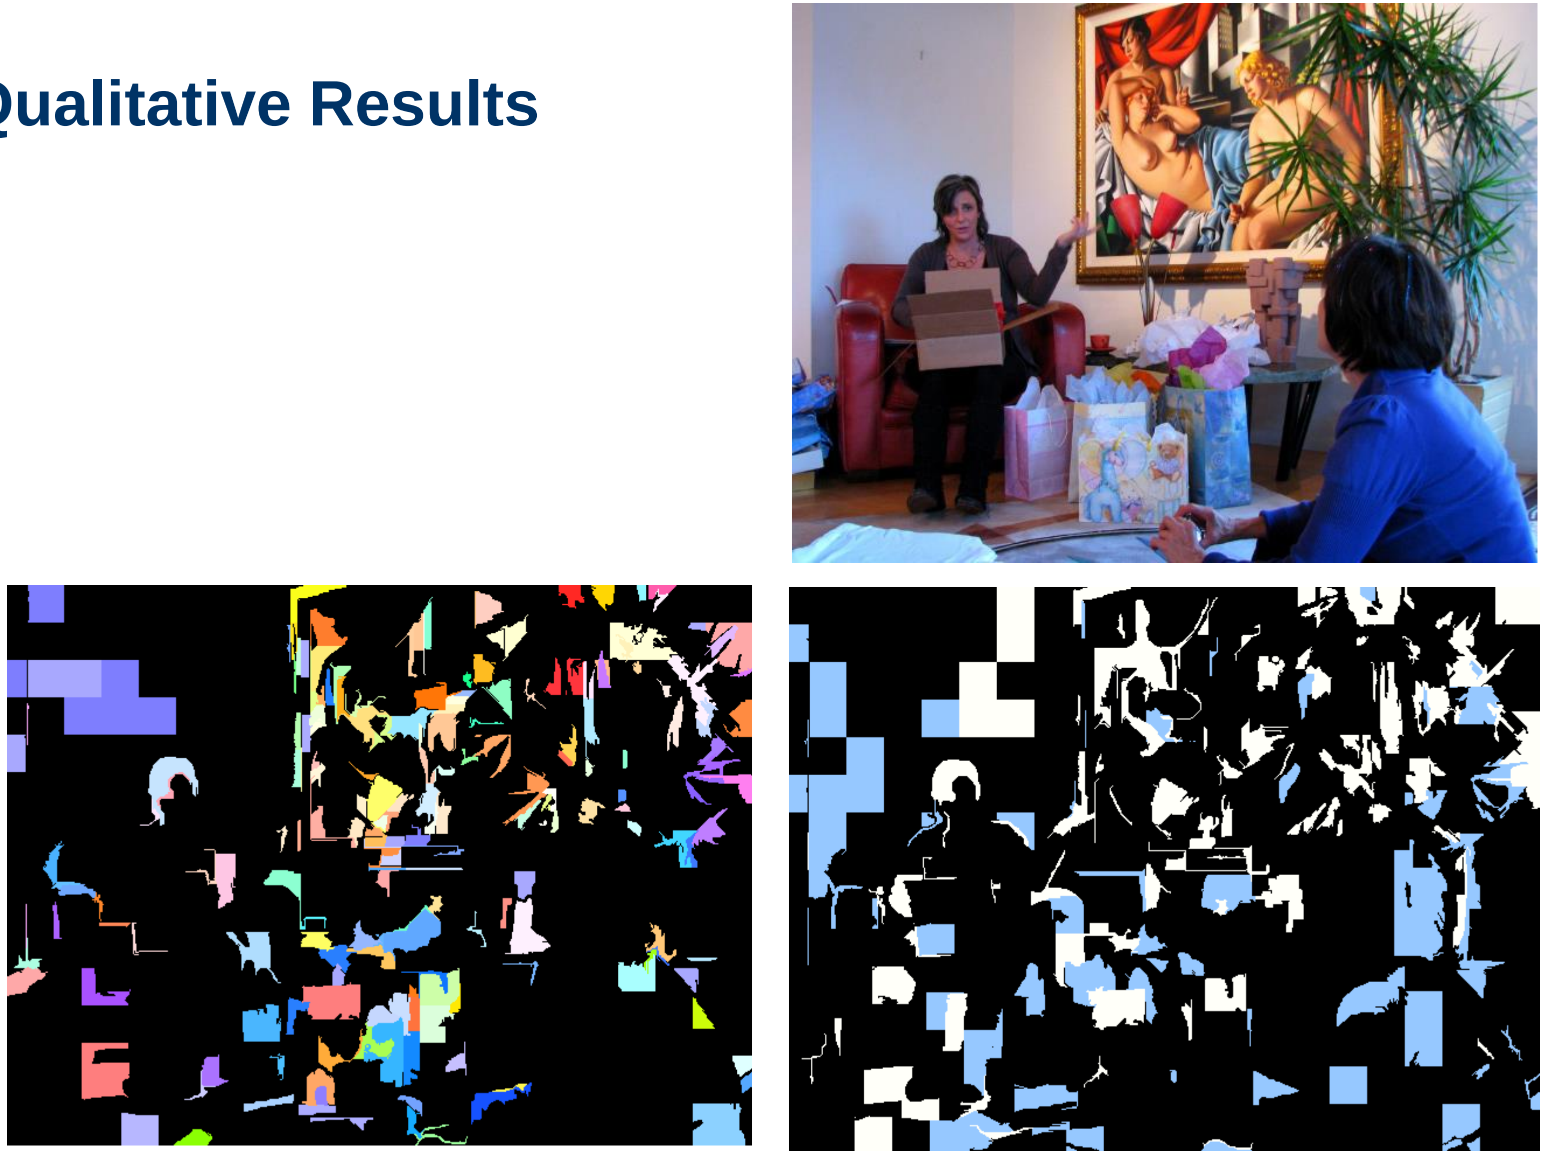

Qualitative Results

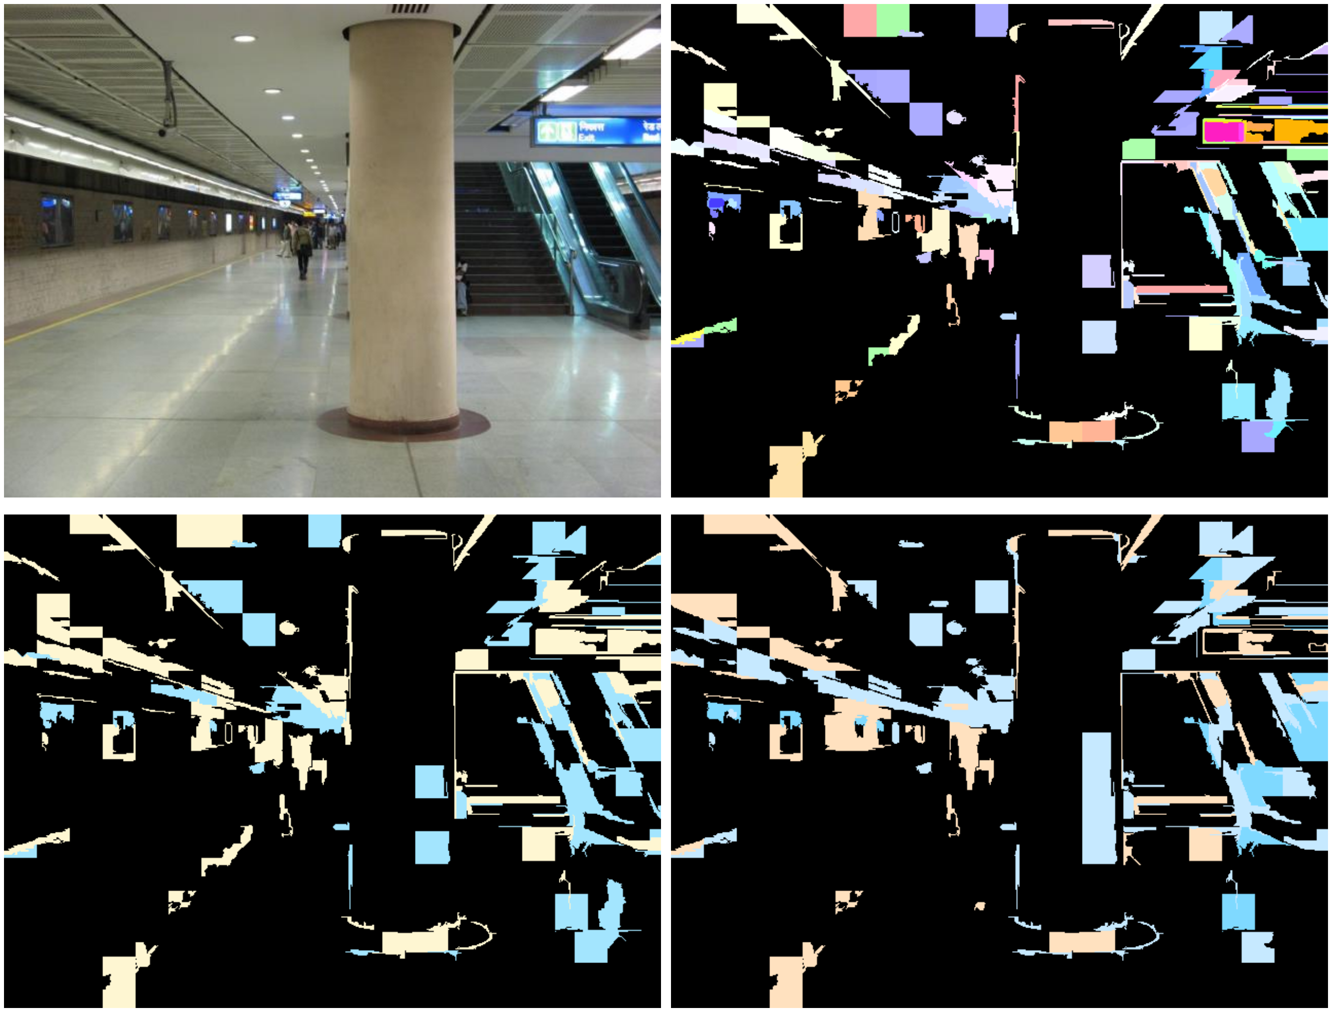

Discussion

- Selection of two or three illuminants is an open issue – although two might be enough in many scenarios

- The optimization criterion is not perfectly well-behaved for single-illuminant scenes:
  encouraging each segment to belong to exactly one candidate illuminant encourages a minimum distance of the candidate illuminants

- However, qualitative results on several real-world scenes look promising

Outlook: Ground Truth

- Meanwhile, we developed a way to obtain quantitative ground truth from a series of linear images under two illuminants:

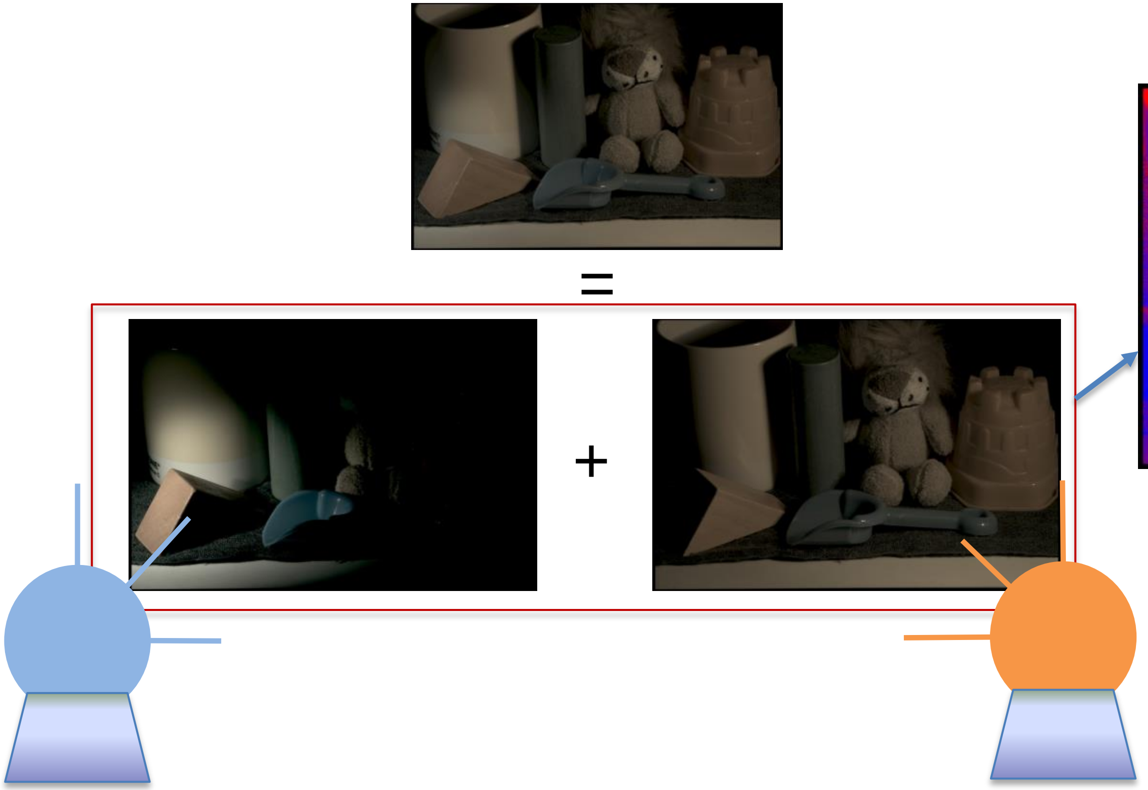

Per-illuminant influence
in the two-illuminant image

Outlook: Conditional Random Field

- We pursued the ideas of this talk in a Conditional Random Field framework (accepted in IEEE TIP two days ago)

- Here, an energy functional is minimized to automatically select the proper number of illuminants

Summary

- Color constancy aims to separate object colors from illumination colors (e.g., for object recognition, white balancing...)
- Color constancy is underconstrained: more unknowns than observations
- Under non-uniform illumination, one has to estimate the colors of the illuminants and their spatial distribution
- We present a straightforward optimization approach to solve both tasks simultaneously
- In further work, we extended this functional to a more principled energy functional

Thank you for your attention!

FRIEDRICH-ALEXANDER
UNIVERSITÄT
ERLANGEN-NÜRNBERG

TECHNISCHE FAKULTÄT