Color:
Inference from color

D.A. Forsyth

University of Illinois at Urbana Champaign

Important fact

- Humans appear to discount the effects of illuminant color
    - This is referred to as color constancy
    - Example:
        - view white plate in yellow incandescent light
            - report “white” (not ”yellow”)
        - view white plate in blue fluorescent light
            - report “white” (not “blue”)
    - It isn’t perfect:
        - traditional example:
            - buy clothing in store light
            - view in sunlight, and regret

Color constancy

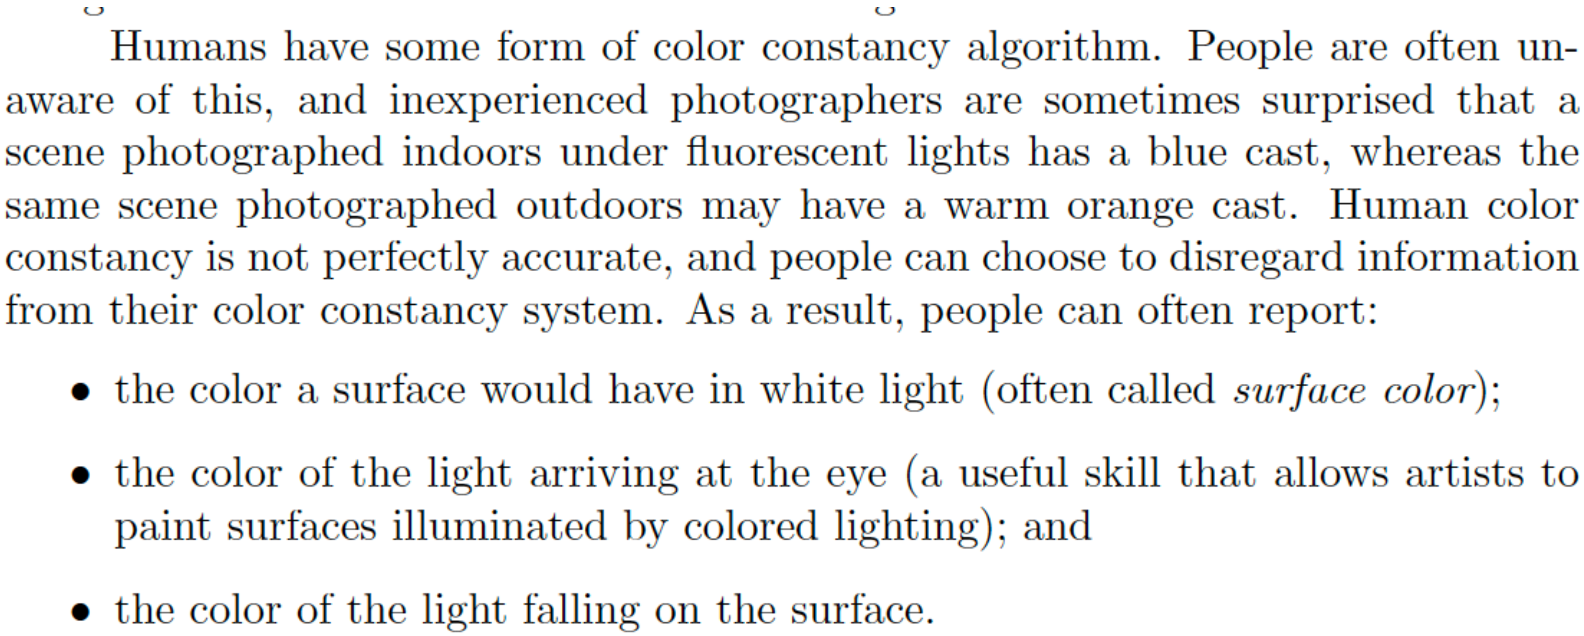

Humans have some form of color constancy algorithm. People are often unaware of this, and inexperienced photographers are sometimes surprised that a scene photographed indoors under fluorescent lights has a blue cast, whereas the same scene photographed outdoors may have a warm orange cast. Human color constancy is not perfectly accurate, and people can choose to disregard information from their color constancy system. As a result, people can often report:

- the color a surface would have in white light (often called *surface color*);
- the color of the light arriving at the eye (a useful skill that allows artists to paint surfaces illuminated by colored lighting); and
- the color of the light falling on the surface.

Color constancy

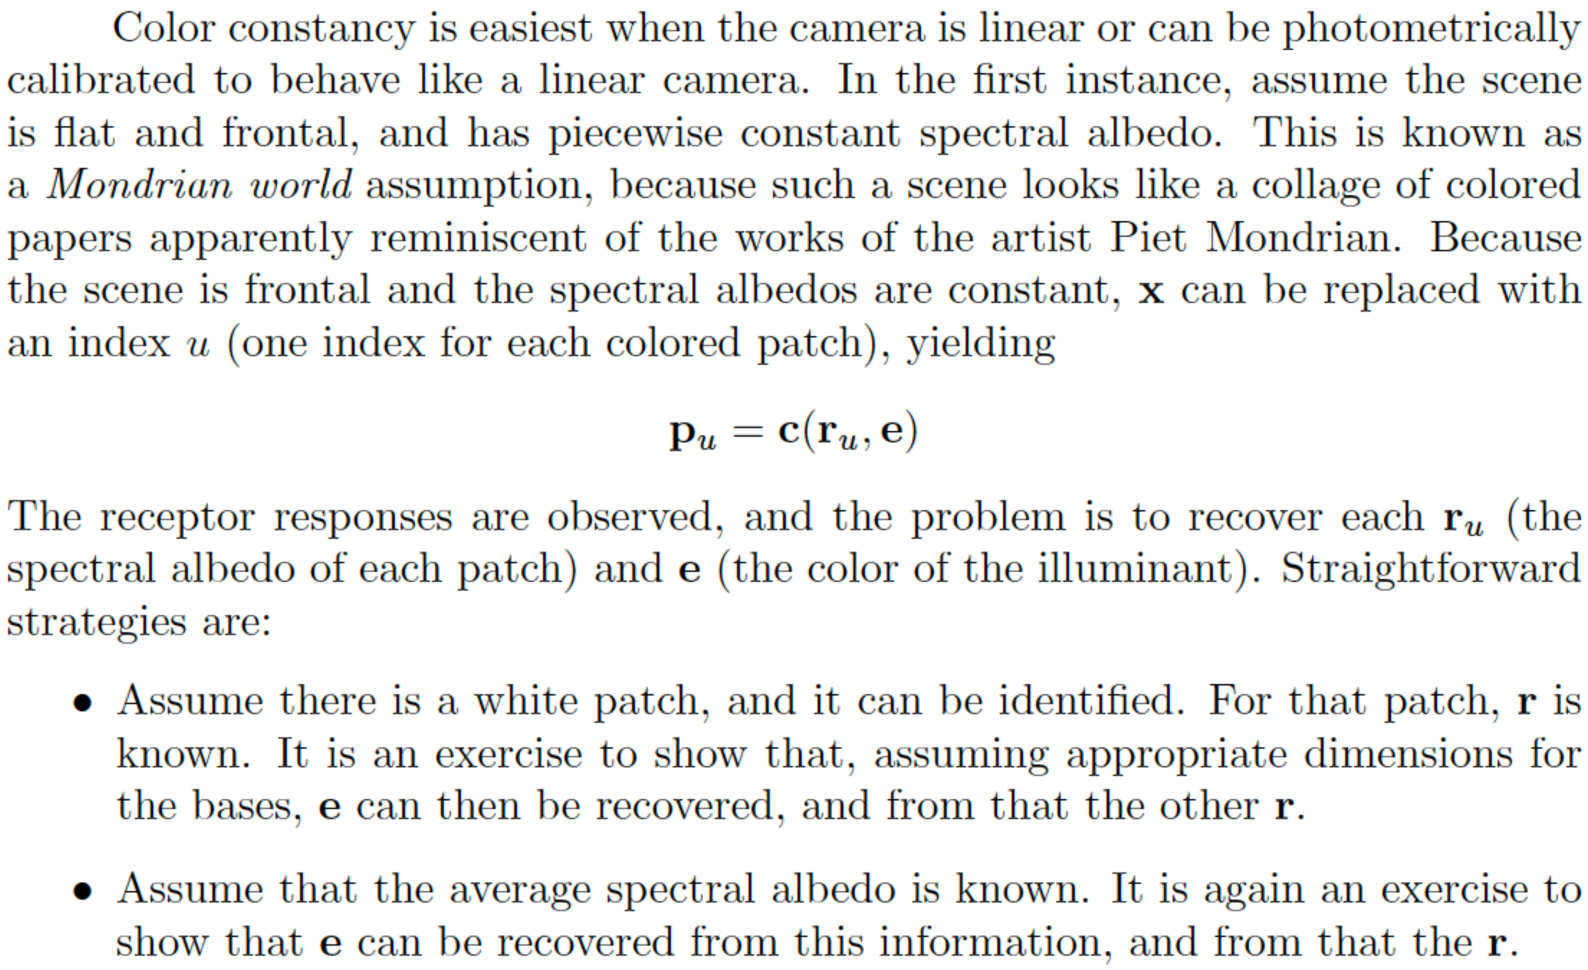

Color constancy is easiest when the camera is linear or can be photometrically calibrated to behave like a linear camera. In the first instance, assume the scene is flat and frontal, and has piecewise constant spectral albedo. This is known as a *Mondrian world* assumption, because such a scene looks like a collage of colored papers apparently reminiscent of the works of the artist Piet Mondrian. Because the scene is frontal and the spectral albedos are constant, $\mathbf{x}$ can be replaced with an index u (one index for each colored patch), yielding

$$\mathbf{p}_u = \mathbf{c}(\mathbf{r}_u, \mathbf{e})$$

The receptor responses are observed, and the problem is to recover each $\mathbf{r}_u$ (the spectral albedo of each patch) and $\mathbf{e}$ (the color of the illuminant). Straightforward strategies are:

- Assume there is a white patch, and it can be identified. For that patch, $\mathbf{r}$ is known. It is an exercise to show that, assuming appropriate dimensions for the bases, $\mathbf{e}$ can then be recovered, and from that the other $\mathbf{r}$.
- Assume that the average spectral albedo is known. It is again an exercise to show that $\mathbf{e}$ can be recovered from this information, and from that the $\mathbf{r}$.

Color constancy: gamuts as a cue

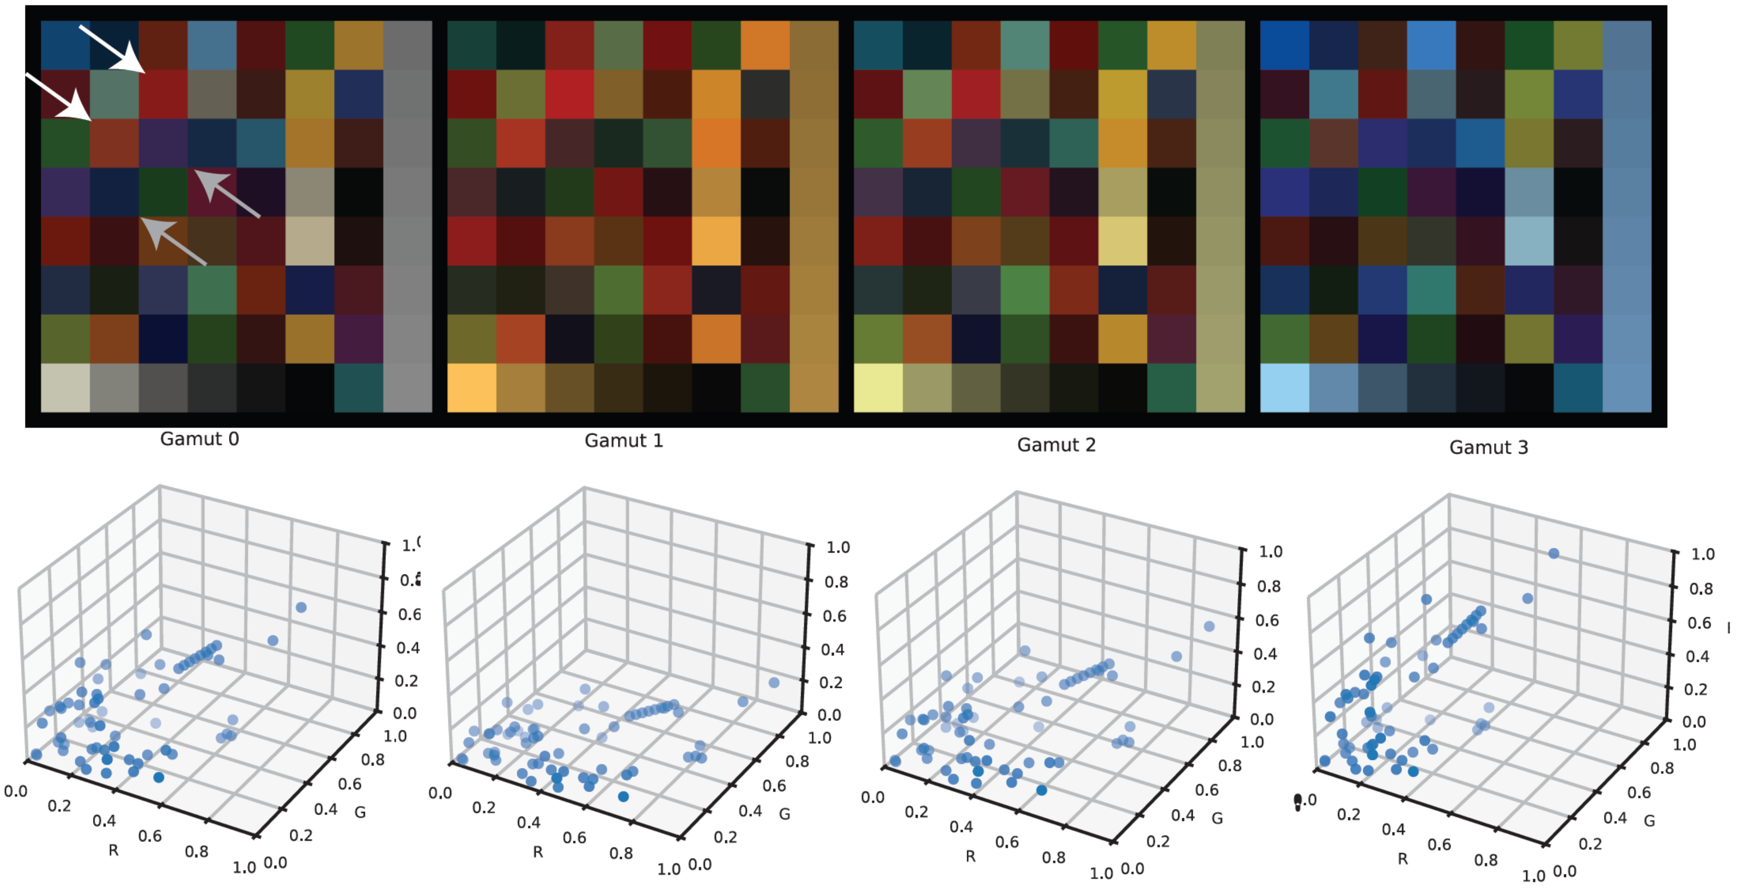

Color constancy has morphed

- White balancing
  - Make the image look as though the light is white
- Intrinsic images
  - Report the albedo at each point
    - actually, not the only intrinsic: depth, normal, etc.
    - but different – you can’t get real training data easily
- Relight the scene
  - Make the image look as though the light had changed

White balancing

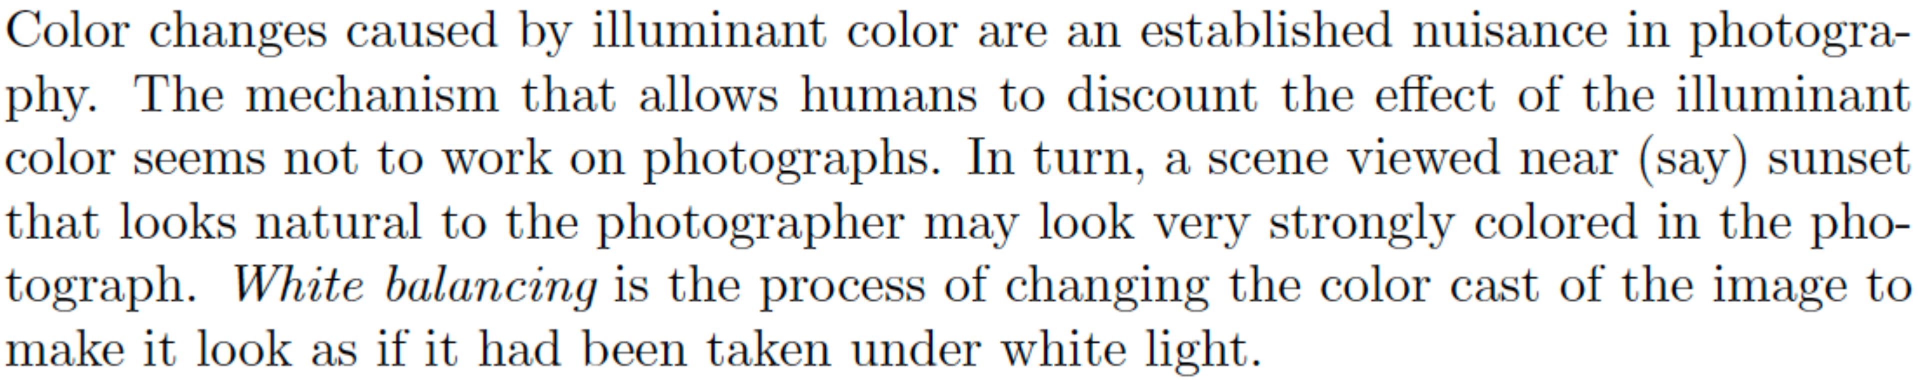

Color changes caused by illuminant color are an established nuisance in photography. The mechanism that allows humans to discount the effect of the illuminant color seems not to work on photographs. In turn, a scene viewed near (say) sunset that looks natural to the photographer may look very strongly colored in the photograph. *White balancing* is the process of changing the color cast of the image to make it look as if it had been taken under white light.

White balancing

Instead, assume there are one or more references in the image. In practical photography, the reference is usually a surface that the photographer knows to be neutral and light in color (perhaps a wall painted white indoors, or the body of a seagull). Write $\mathbf{p}_{i,\mathbf{t}}$ for the color of the i'th reference surface under the target light and $\mathbf{p}_{i,\mathbf{e}}$ for the color of that surface under the illuminant of the photograph. Then the image can be corrected by finding

$$\underset{\mathcal{M}}{\operatorname{argmin}} \sum_i \left(\mathbf{p}_{i,\mathbf{t}} - \mathcal{M}\mathbf{p}_{i,\mathbf{e}}\right)^T \left(\mathbf{p}_{i,\mathbf{t}} - \mathcal{M}\mathbf{p}_{i,\mathbf{e}}\right)$$

which is a simple least-squares problem (exercises). The material of Section 3.2.3 implies that $\mathcal{M}$ is a general 3×3 matrix (exercises). It is quite usual to assume that $\mathcal{M}$ is a diagonal matrix instead, and compute one scale for each channel. In the simplest case, one simply scales by the pixel values for a white surface (exercises).

White balancing

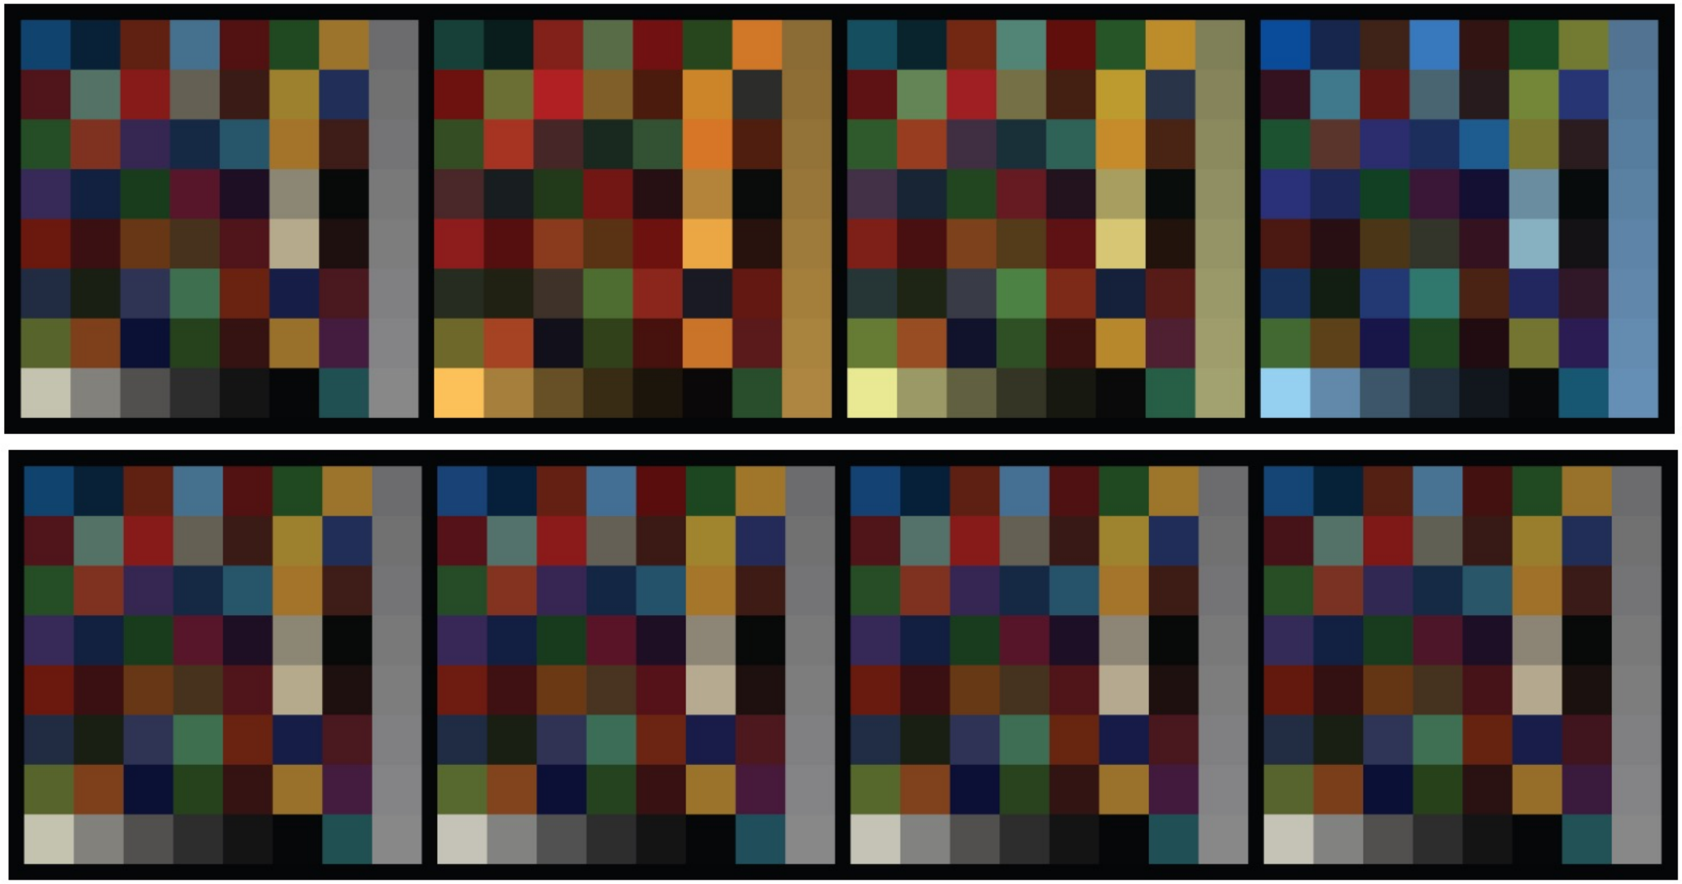

dw: 8.6 dl: 6.1 ls: 5.1 dw: 5.5 dl: 3.1 ls: 2.7 dw: 14.8 dl: 13.1 ls: 8.3

Intrinsic images

- Report color albedo from image
  - simplest:
    - lightness algorithm in R, G, and B
  - harder:
    - obtain dataset
      - render pairs of (true illuminated image, albedo)
      - simulated pairs
        - true image= (piecewise constant albedo)* (slowly varying shading)
    - fire up U-Net
      - BUT sim to real gap
  - even harder:
    - use conditional diffusion model trained rendered data

UNet with simulated data

Forsyth and Rock, 2022

Conditional diffusion model

Image

Mean albedo estimate

Kocsis et al 2024

Relight the scene

- Simplest version
    - remove shadows
        - this can be useful, but is physically weird
- Harder
    - Insert objects
- Even harder
    - Change illuminant in some way
        - and produce physically correct answer

Shadow removal outdoors

- Assume light sources are black body
  - at different temperatures
- You can then infer:
  - light source temperature
  - whether a boundary is a shadow boundary
- Reconstruct by:
  - Differentiate,
  - knock out shadow boundaries,
  - integrate

Shadow removal

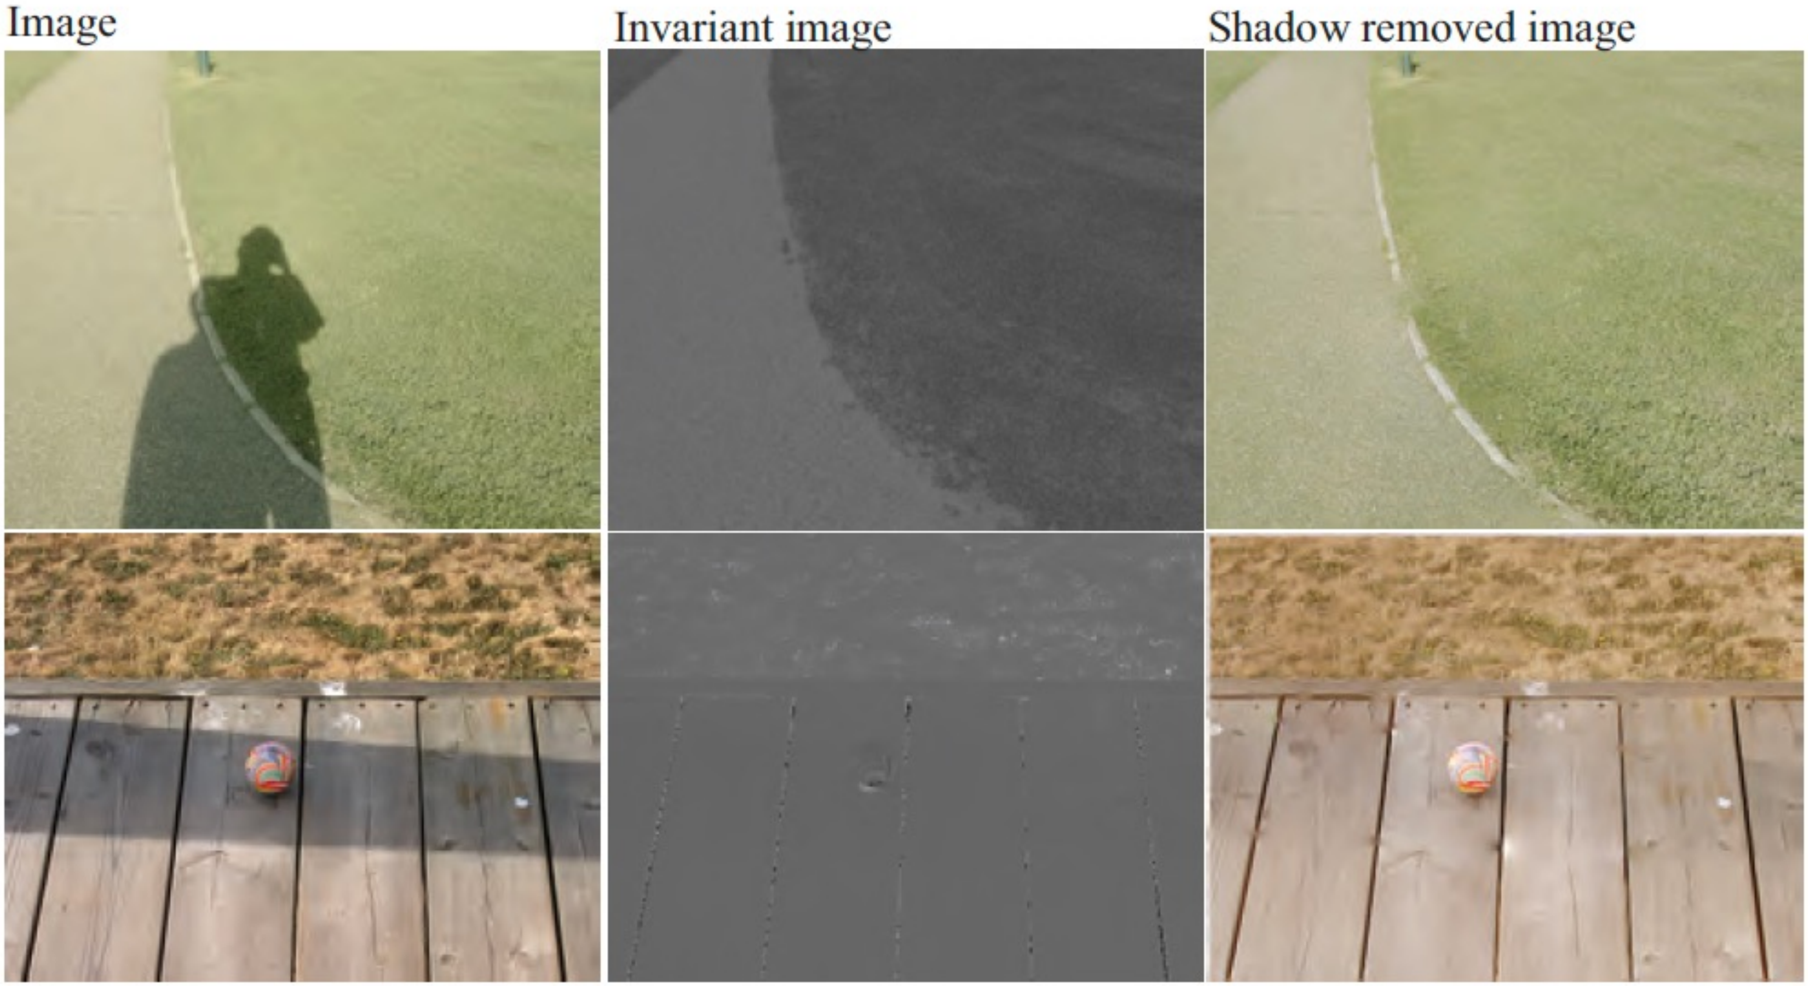

Inserting Computer Graphics

Input image

Estimate geometry

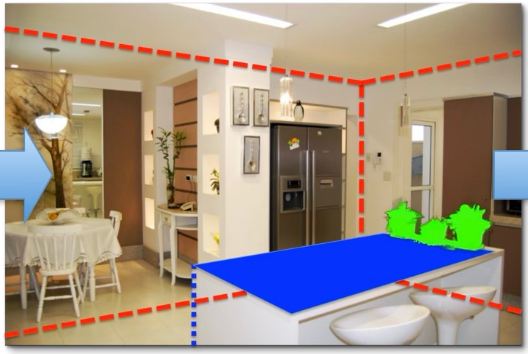

Using above, with manual help

Estimate materials

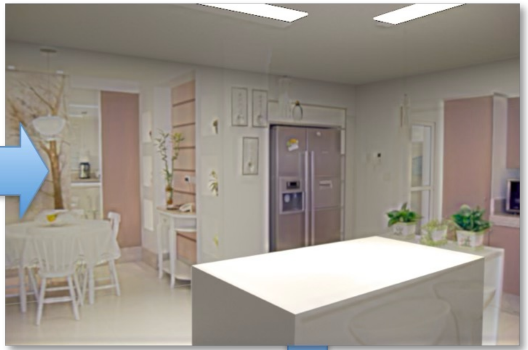

Standard methods (Land 71)

Estimate lighting

Markup, for the moment
Secret sauce: Consistency
Secret sauce: Shafts

Compose & render

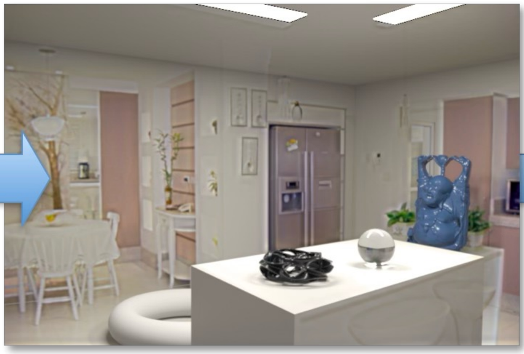

Secret sauce: Physical renderer

Karsch ea 11

Final composite

Compositing by standard
method Debevec 98

Karsch ea 11

Karsch ea 11

Changing the light

- Idea:
  - can recover:
    - albedo
    - normal
    - depth
    - etc
  - get that, render with new light
- This seems not to work all that well
  - likely problem:
    - errors in inferences interact in a bad way

Turn off a visible lamp

Li et al, 2022

Changing the light

- Alternatives:
  - various forms of conditional image generation
  - (rather loosely) a regression problem

Changing the light

Bhattad et al 24

Changing the light

Xing et al 25