Does Color Really Help in Dense Stereo Matching?

Michael Bleyer[1] and Sylvie Chambon[2]

[1]Vienna University of Technology, Austria
[2]Laboratoire Central des Ponts et Chaussées, Nantes, France

Dense Stereo Matching

(Left Image)

(Right Image)

Dense Stereo Matching

(Left Image)

(Right Image)

(Disparity Map)

Dense Stereo Matching

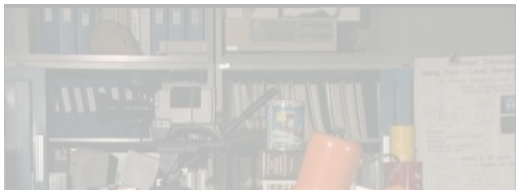

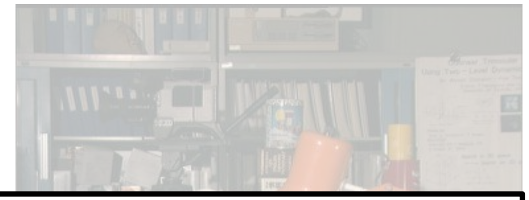

Key question of this paper:

Can we improve matching performance by using color information?

(Disparity Map)

Why Should Color Help?

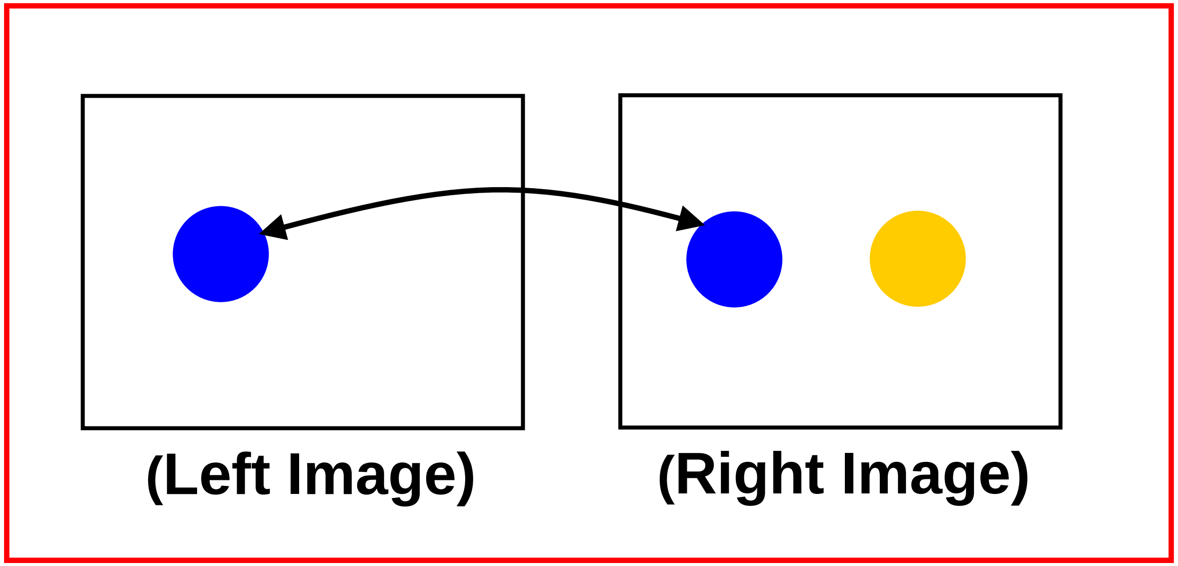

Why Should Color Help?

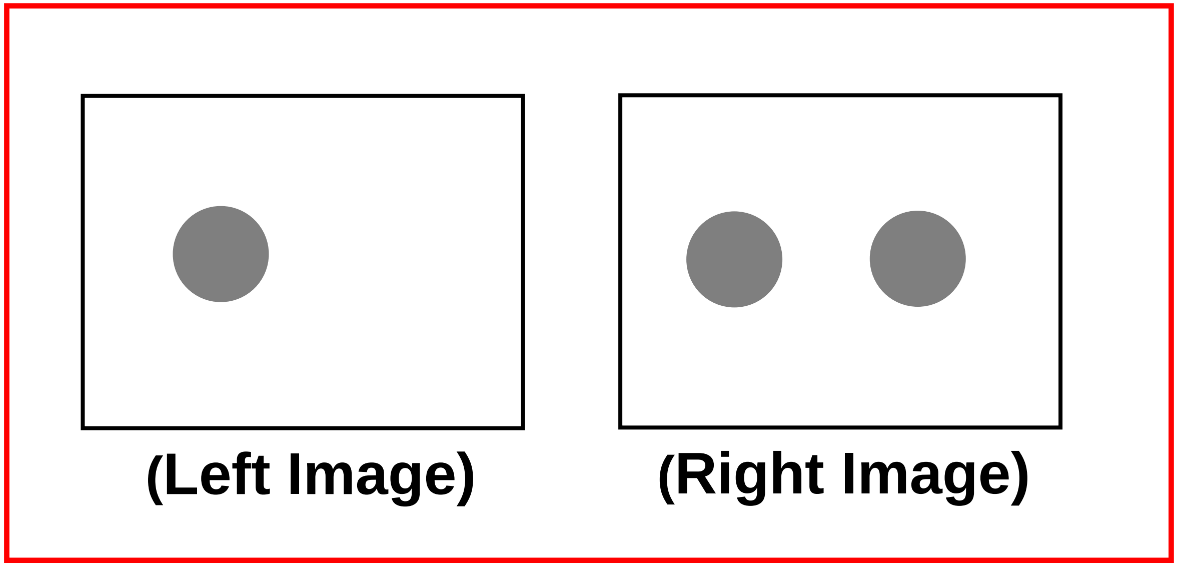

Conversion to grey-scale

Why Should Color Help?

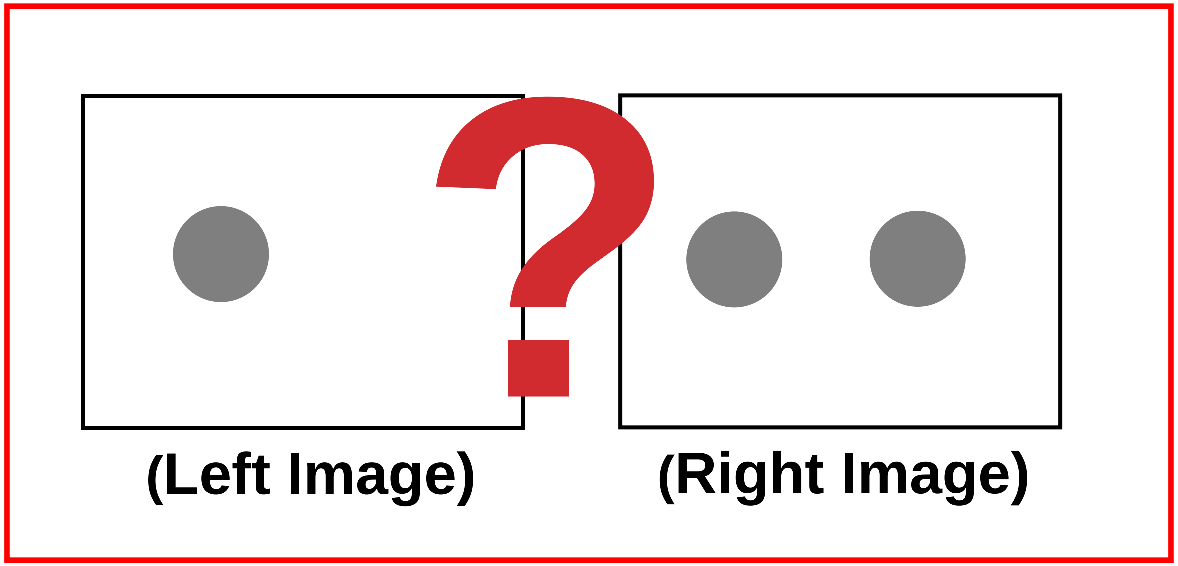

Conversion to grey-scale

Why Should Color Help?

Color reduces the ambiguity problem

(Left Image) (Right Image)

Conversion to grey-scale

Contradicting Statements in the Literature

- Several previous color evaluation studies:
  - In the context of local methods:
    - [Chambon, IJRA05], [Mühlmann, IJCV02], …
  - In the context of global methods:
    - [Bleyer, ISPRS08]
  - In the context of radiometric insensitive match measures (local and global methods)
    - [Hirschmüller, PAMI09]

Color helps

Color helps

Color does not help

Contradicting Statements in the Literature

- Several previous color evaluation studies:
  - In
  - In
    - [Bleyer, ISPRS08]
  - In the context of radiometric insensitive match measures (local and global methods)
    - [Hirschmüller, PAMI09]

helps

helps

Color does not help

Goal of our benchmark:
Find out who is right

Remainder of this Talk

- Benchmark design:
  - Competing energy functions
    - Dissimilarity functions
    - Color spaces
  - Performance metrics
- Benchmark results:
  - Standard dissimilarity functions
  - Radiometric insensitive functions

Energy Function

- Quality of disparity map *D* measured by a standard energy function:

$$E(D) = \sum_{p \in I} m(p, d_p) + \sum_{<p,q> \in N} s(d_p, d_q)$$

Data Term:

- **Measures pixel dissimilarities**
- **Evaluated in this paper**

Smoothness Term:

- **Measures amount of spatial smoothness**

Energy Function

- Quality of disparity map *D* measured by a standard energy function:

$$E(D) = \sum_{p \in I} m(p, d_p) + \sum_{<p,q> \in N} \boxed{s(d_p, d_q)}$$

- Modified Potts Model:

$$s(d_p, d_q) = \begin{cases} 0 & \text{if } d_p = d_q \\ P_1 & \text{if } |d_p - d_q| = 1 \\ P_2 & \text{if } |d_p - d_q| > 1 \end{cases}$$

- $P_1 := P_2 / 3$
- P_2 is tuned individually for each energy function

Energy Function

- Quality of disparity map *D* measured by a standard energy function:

$$E(D) = \sum_{p \in I} \boxed{m(p, d_p)} + \sum_{<p,q> \in N} s(d_p, d_q)$$

- Various competing implementations:
  - 4 dissimilarity functions
  - 9 color spaces
- Leads to 4 * 9 = 36 competing energy functions

Dissimilarity Functions

- Absolute difference of colors
  - Cannot handle radiometric distortions
  - Pixel-based
- Mutual Information (MI)
  - Handles radiometric distortions
  - Pixel-based
- Zero mean Normalized Cross-Correlation (NCC)
  - Handles radiometric distortions
  - Window-based
- Census
  - Handles radiometric distortions
  - Window-based

Dissimilarity Functions

- Absolute difference of colors
  - Cannot handle radiometric distortions

We incorporate color by
- Computing the dissimilarity function individually for each color channel
- Summing up the values over the 3 channels

  - Window-based
- Census
  - Handles radiometric distortions
  - Window-based

Color Spaces

- Primary systems:
  - ***RGB, XYZ;***
- Luminance-chrominance systems:
  - ***LUV, LAB, AC_1C_2, YC_1C_2;***
- Statistical independent component systems:
  - ***$I_1I_2I_3$, $H_1H_2H_3$;***
- Use of intensity values only:
  - ***Grey;***

Name	Definition
Grey	$I = 0.299R + 0.587G + 0.114B$
XYZ	$\begin{pmatrix} X \\ Y \\ Z \end{pmatrix} = \begin{pmatrix} 0.607 & 0.174 & 0.200 \\ 0.299 & 0.587 & 0.114 \\ 0.000 & 0.066 & 1.116 \end{pmatrix} \begin{pmatrix} R \\ G \\ B \end{pmatrix}$
LUV	$L = \begin{cases} 116\,(Y/Y_w)^{\frac{1}{3}} - 16 & \text{if } Y/Y_w > 0.01 \\ 903.3\,Y/Y_w & \text{otherwise} \end{cases}$ $U = 13L(u' - u'_w)$ with $u' = \frac{4X}{X+15Y+3Z}$ $V = 13L(v' - v'_w)$ with $v' = \frac{9Y}{X+15Y+3Z}$ X_w, Y_w, Z_w: white reference components
LAB	$A = 500(f(X/X_w) - f(Y/Y_w))$ $B = 200(f(Y/Y_w) - f(Z/Z_w))$ $f(x) = \begin{cases} x^{1/3} & \text{if } x > 0.008856 \\ 7.787x + \frac{16}{116} & \text{otherwise} \end{cases}$
AC_1C_2	$\begin{pmatrix} A \\ C_1 \\ C_2 \end{pmatrix} = \begin{pmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{\sqrt{3}}{2} & \frac{-\sqrt{3}}{2} & 0 \\ \frac{-1}{2} & \frac{-1}{2} & 1 \end{pmatrix} \begin{pmatrix} R \\ G \\ B \end{pmatrix}$
YC_1C_2	$\begin{pmatrix} Y \\ C_1 \\ C_2 \end{pmatrix} = \begin{pmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ 1 & \frac{-1}{2} & \frac{-1}{2} \\ 0 & \frac{-\sqrt{3}}{2} & \frac{\sqrt{3}}{2} \end{pmatrix} \begin{pmatrix} R \\ G \\ B \end{pmatrix}$
$I_1I_2I_3$	$\begin{pmatrix} I_1 \\ I_2 \\ I_3 \end{pmatrix} = \begin{pmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{2} & 0 & \frac{-1}{2} \\ \frac{-1}{4} & \frac{-1}{4} & \frac{1}{2} \end{pmatrix} \begin{pmatrix} R \\ G \\ B \end{pmatrix}$
$H_1H_2H_3$	$\begin{pmatrix} H_1 \\ H_2 \\ H_3 \end{pmatrix} = \begin{pmatrix} 1 & 1 & 0 \\ 1 & -1 & 0 \\ \frac{-1}{2} & 0 & \frac{-1}{2} \end{pmatrix} \begin{pmatrix} R \\ G \\ B \end{pmatrix}$

(Conversion from *RGB* to other color spaces)

Energy Minimization

- Accomplished via Dynamic Programming (DP)
- DP is performed on simple tree structures [Bleyer,VISAPP09]:
  - Trees contain horizontal and vertical smoothness edges (scanline streaking problem)
- Simple method for occlusion handling

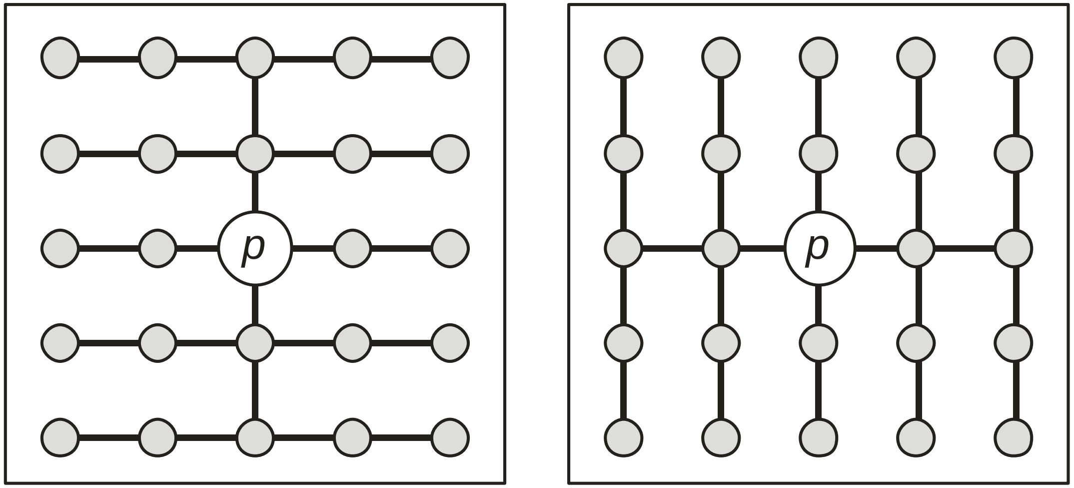

Error Metrics

- 3 error metrics to evaluate the quality of a disparity map:
  1. Percentage of wrong pixels in **all** regions
  2. Percentage of wrong pixels in **radiometric distorted** regions
  3. Percentage of wrong pixels in **radiometric clean** regions
- Definitions:
  - **Wrong pixel** means absolute disparity error > 1 (analogously to Middlebury)
  - **Radiometric distorted** means a pixel has different intensity/color in left and right images
  - When I say **average error** I mean the average error computed over all 30 test pairs

Error Metrics

- 3 error metrics to evaluate the quality of a disparity map:
  1. P
  2. P … **d** regions
  3. P … gions

How can we extract radiometric distorted regions?

- Definitions:
  - **Wrong pixel** means absolute disparity error > 1 (analogously to Middlebury)
  - **Radiometric distorted** means a pixel has different intensity/color in left and right images
  - When I say **average error** I mean the average error computed over all 30 test pairs

Extraction of Radiometric Affected Regions

- Compute absolute intensity difference between corresponding pixels (Correct correspondences known from ground truth disparity data)

Moebius left image (new Middlebury set)

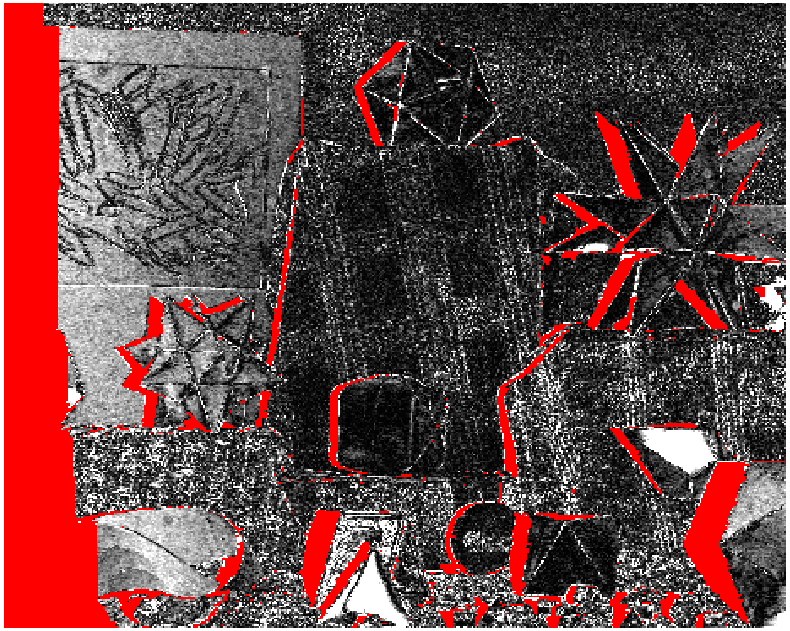

Ground truth data costs (bright pixels have high pixel dissimilarity)

Extraction of Radiometric Affected Regions

- Compute absolute intensity difference between corresponding pixels (Correct correspondences known from ground truth disparity data)

Moebius left image (new Middlebury set)

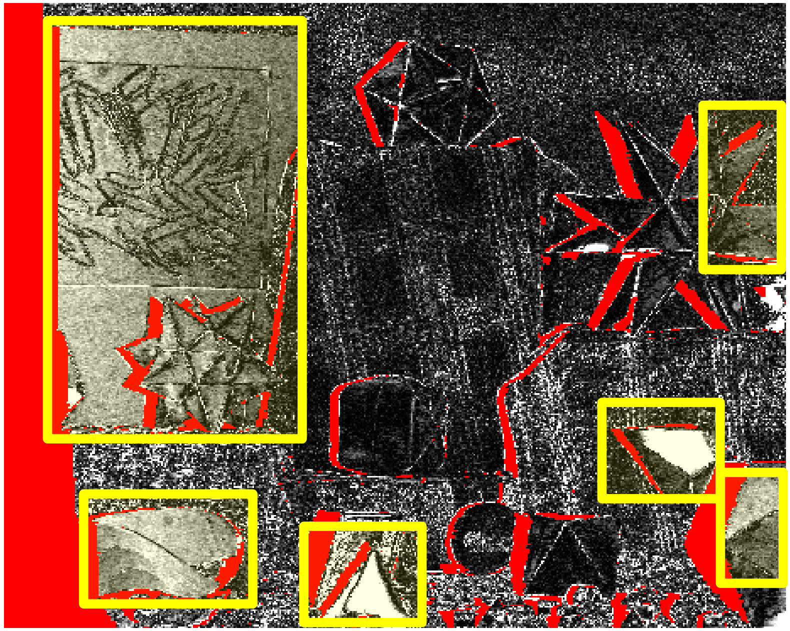

Ground truth data costs (bright pixels have high pixel dissimilarity)

Extraction of Radiometric Affected Regions

- Compute absolute intensity difference between corresponding pixels (Correct correspondences known from ground truth disparity data)

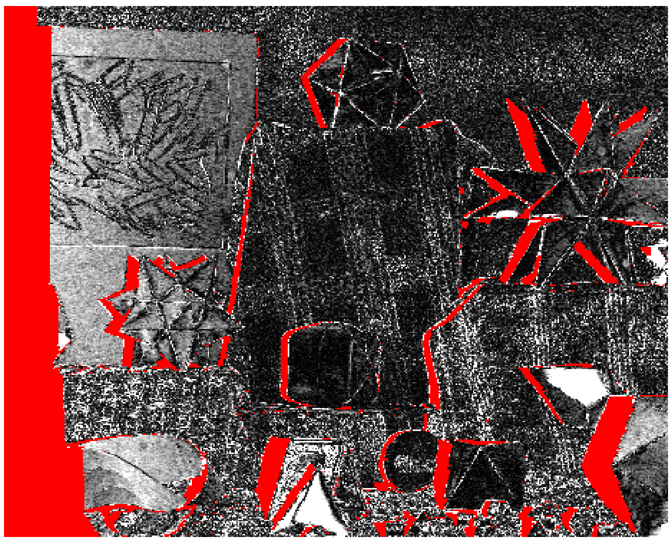

Median Filter

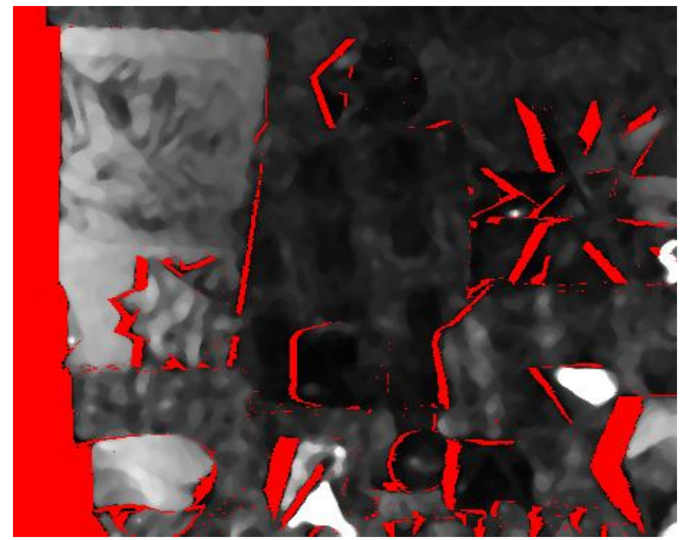

Ground truth data costs

Radiometric distorted regions

Test Set

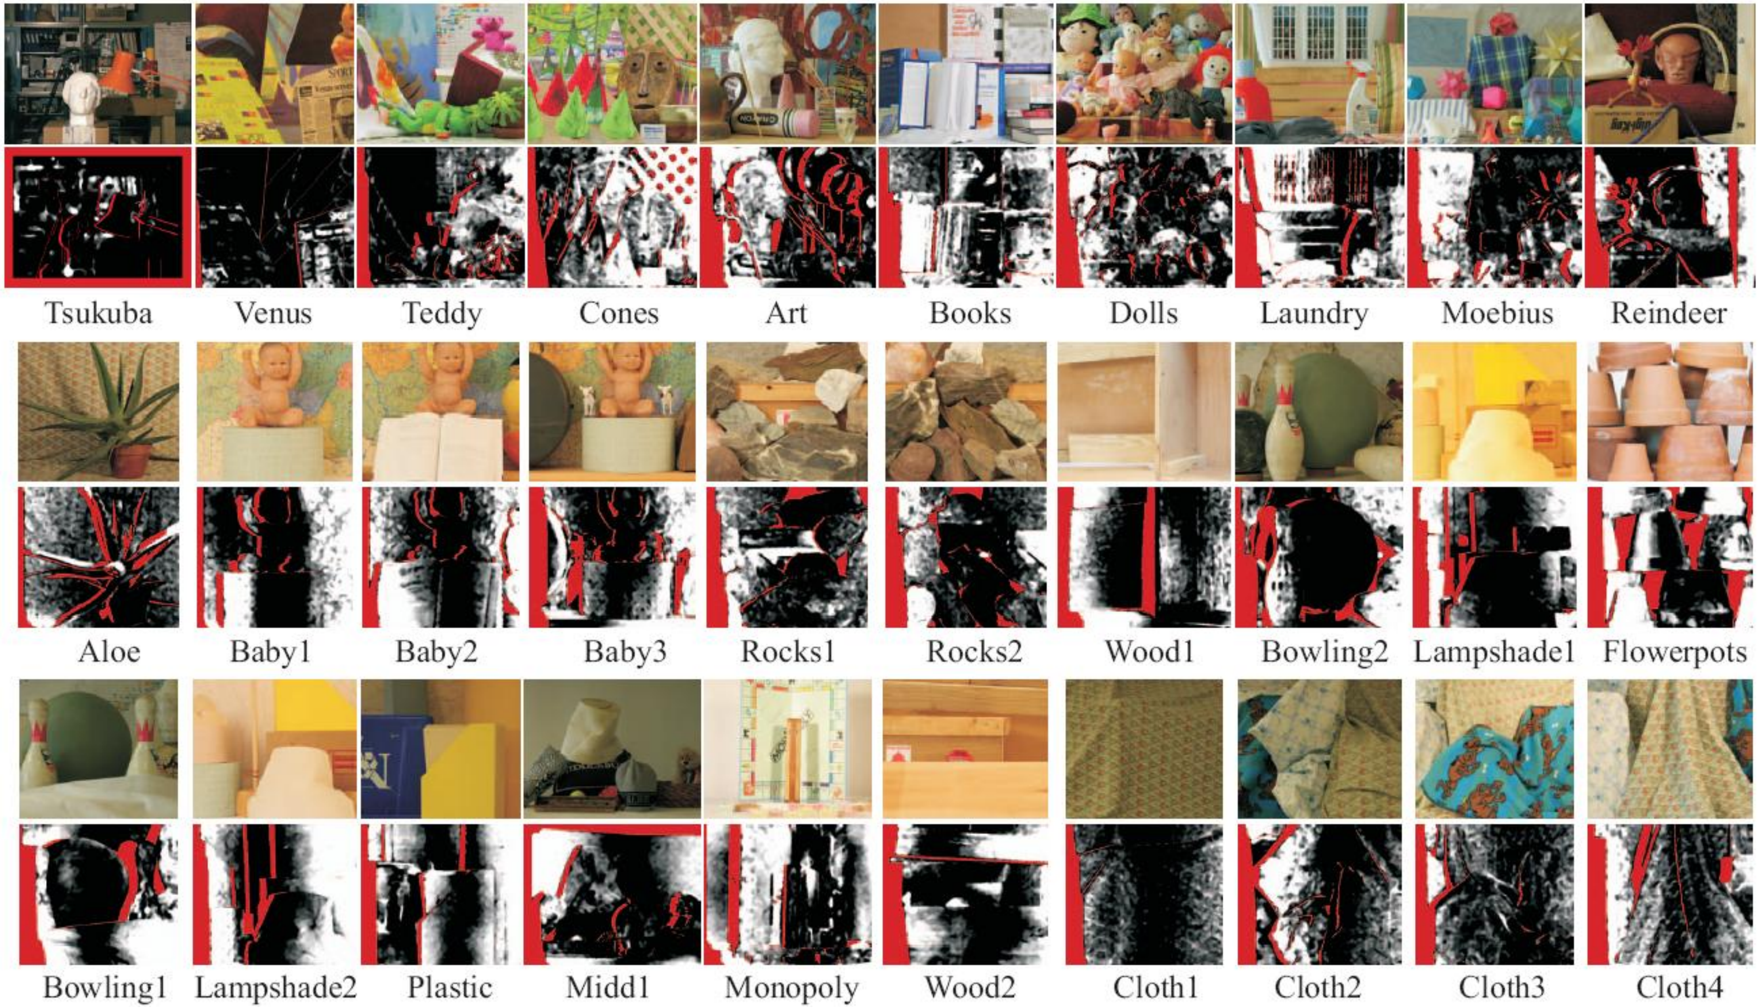

30 test pairs (Middlebury set) with corresponding images showing radiometic distortions

Results

Absolute Differences as Dissimilarity Function

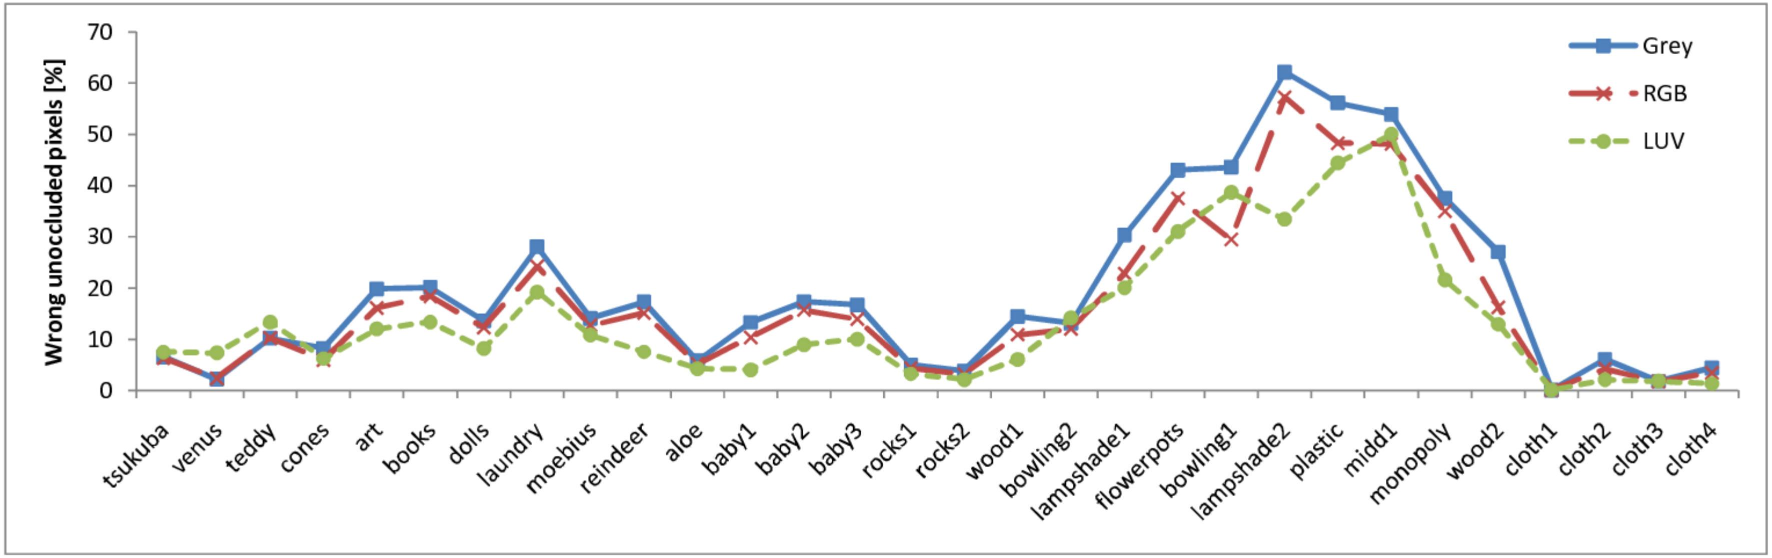

- Grey-scale matching nearly always performs worst.
- *LUV* performs better than *RGB*
- Identical to the results of [Bleyer,ISPR08]

Absolute Differences as Dissimilarity Function

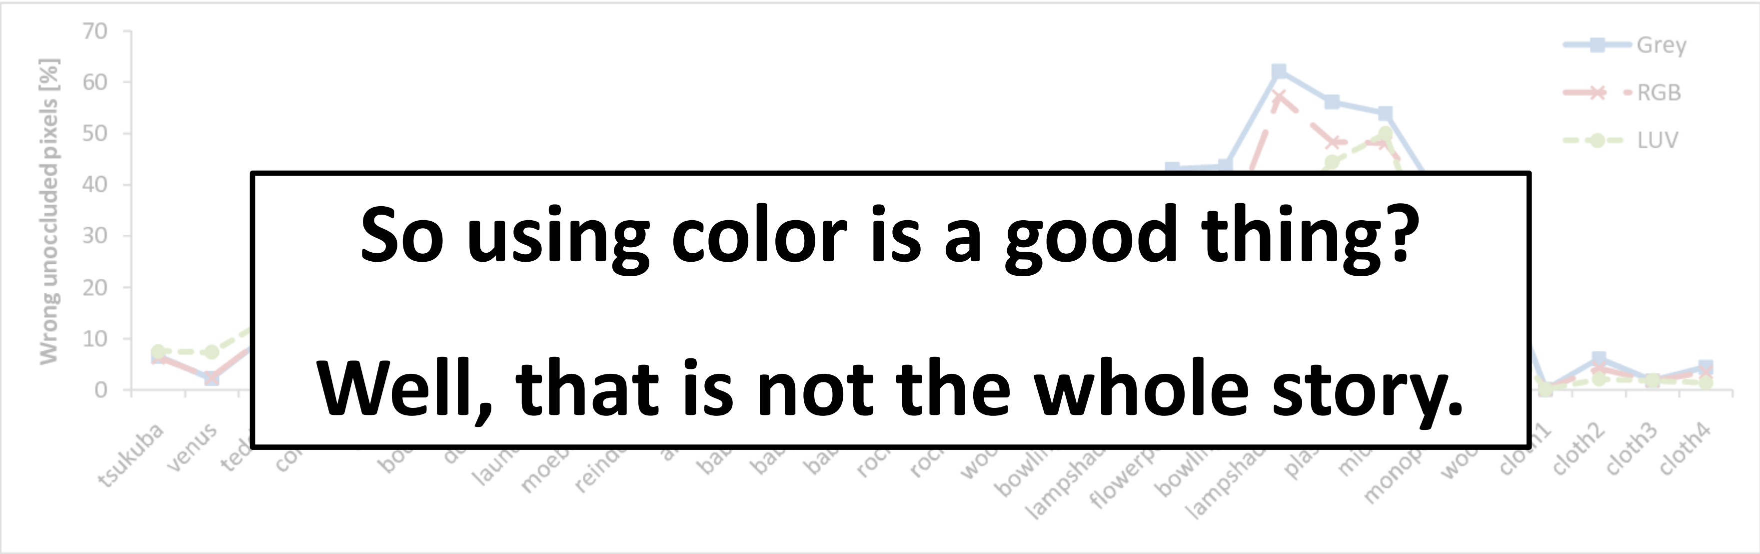

So using color is a good thing?

Well, that is not the whole story.

- Grey-scale matching nearly always performs worst.
- *LUV* performs better than *RGB*
- Identical to the results of [Bleyer,ISPR08]

Where Does Color Help?

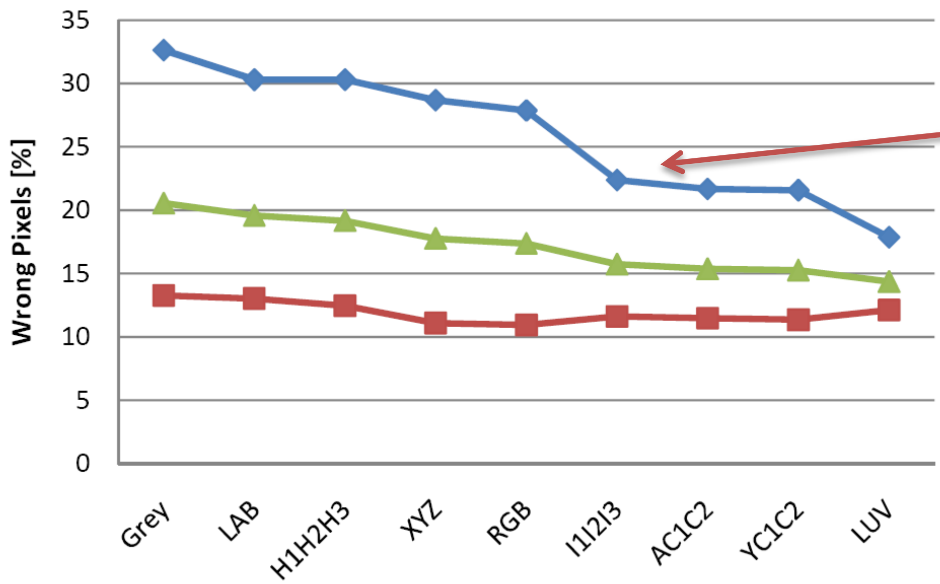

Avg. error
radiometric distorted
regions

- Color effectively improves performance in **radiometric distorted** regions.

Where Does Color Help?

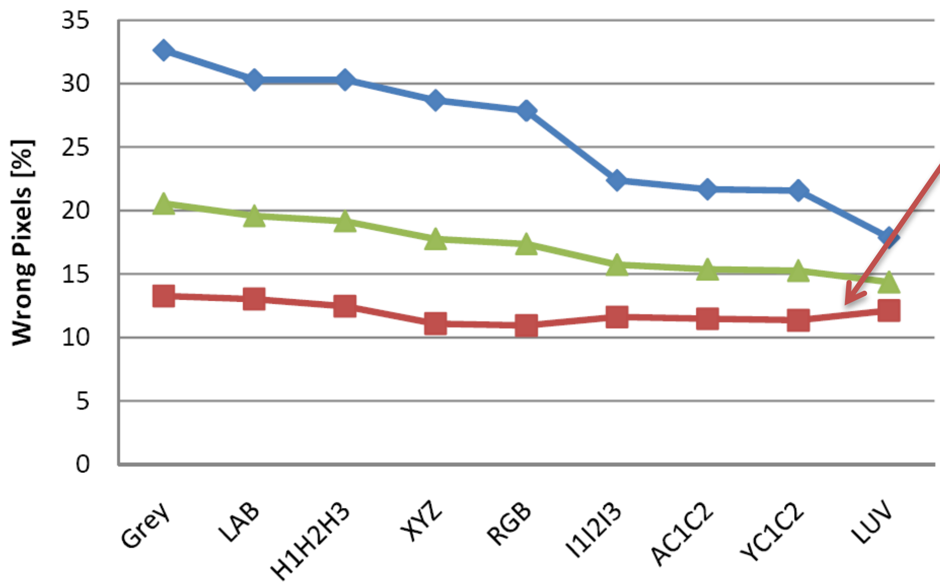

Avg. error
radiometric clean regions

- Color effectively improves performance in **radiometric distorted** regions.
- Color makes little difference in **radiometric clean** regions.

Where Does Color Help?

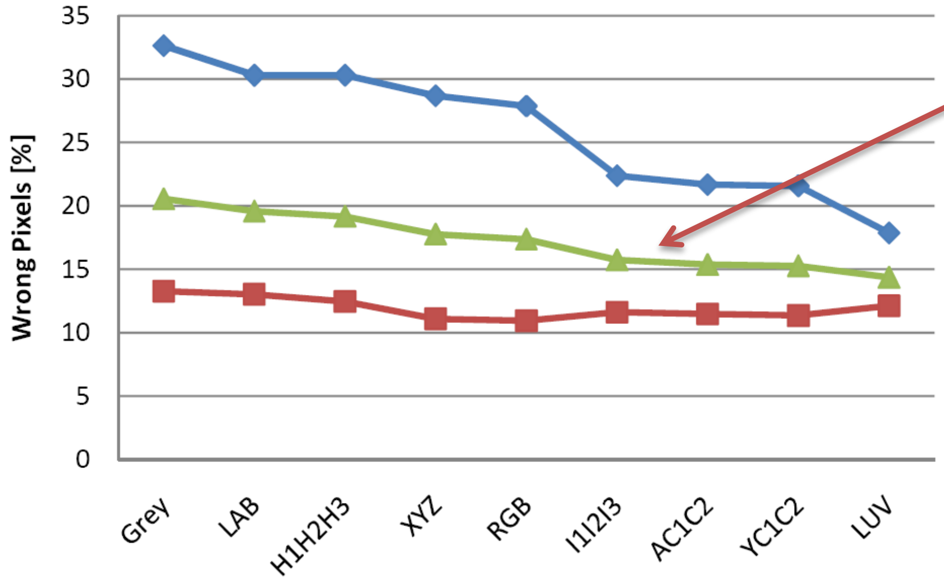

Avg. error
all regions

- Color effectively improves performance in **radiometric distorted** regions.
- Color makes little difference in **radiometric clean** regions.
- The **overall** improvement is largely due to considerably improved performance in **radiometric distorted** regions.

Where Does Color Help?

**If color mostly helps in radiometric distorted regions,
why not directly using a radiometric insensitive match measure?**

- Color … rted regions.
- Color makes little difference in **radiometric clean** regions.
- The **overall** improvement is largely due to considerably improved performance in **radiometric distorted** regions.

Why Not Directly Using a Radiometric Insensitive Measure?

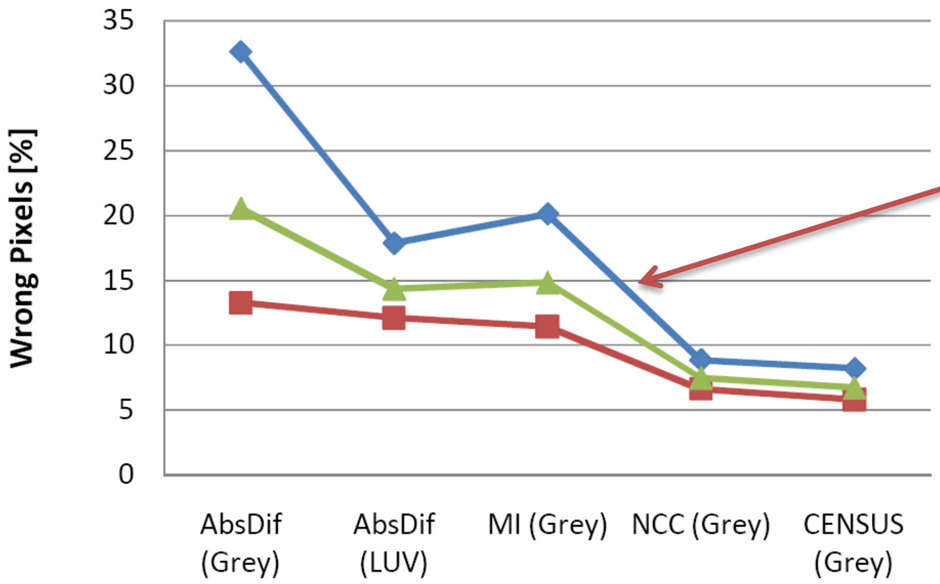

Avg. error
radiometric distorted
regions

- ZNCC and Census considerably improve performance in **radiometric distorted** regions.
  - They are more effective than color in this respect

Why Not Directly Using a Radiometric Insensitive Measure?

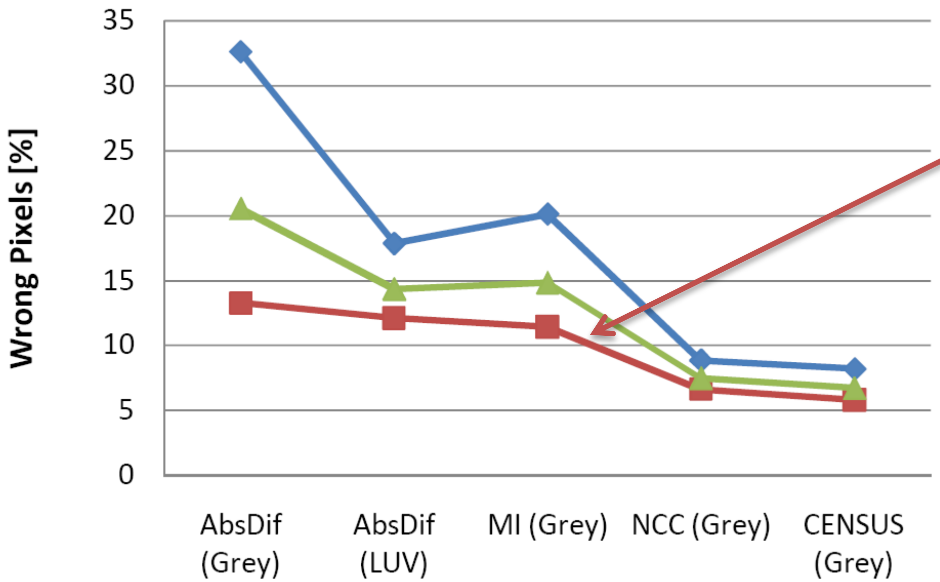

Avg. error
radiometric clean regions

- ZNCC and Census considerably improve performance in **radiometric distorted** regions.
  - They are more effective than color in this respect
- They even improve performance in **radiometric clean** regions.

Why Not Directly Using a Radiometric Insensitive Measure?

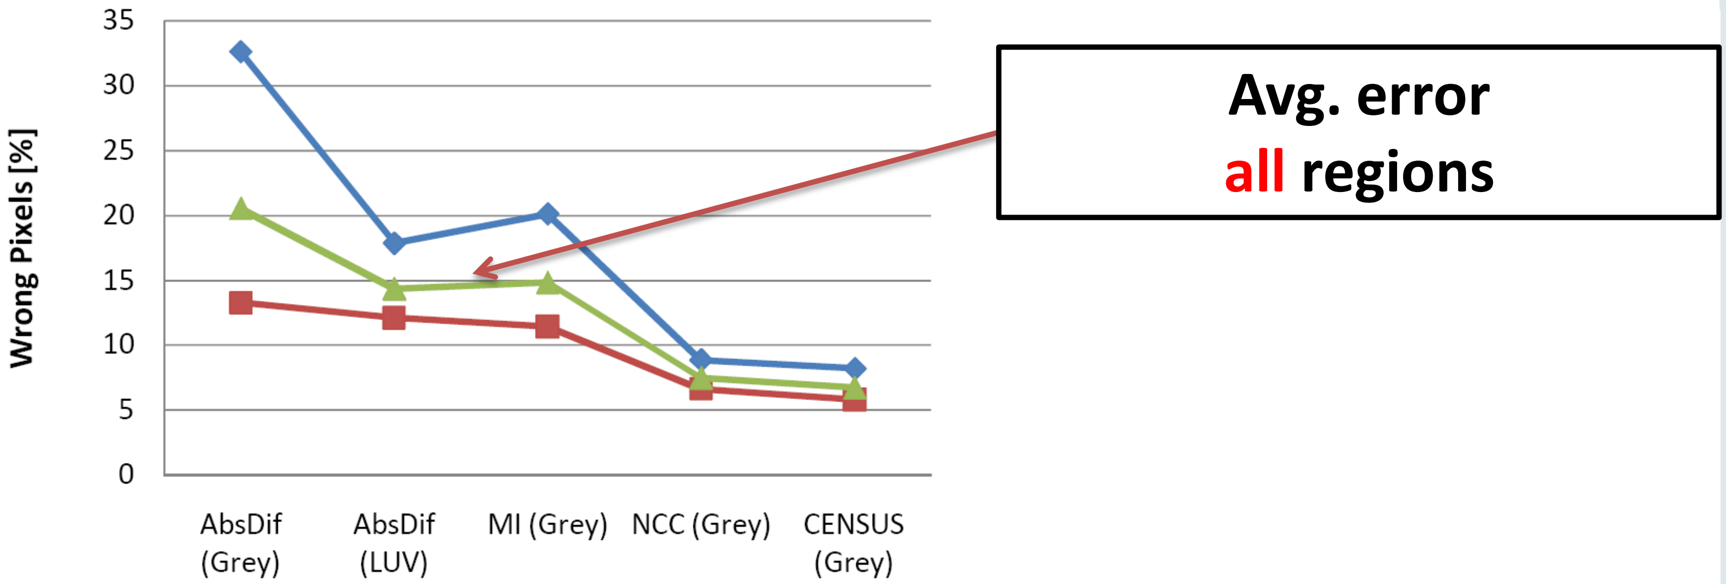

- ZNCC and Census considerably improve performance in **radiometric distorted** regions.
  - They are more effective than color in this respect
- They even improve performance in **radiometric clean** regions.
- ZNCC and Census improve the **overall** performance considerably
  - The overall error drops from **20.5%** [Absdif (Grey)] to **6.7%** [Census (Grey)] **(!)**

Why Not Directly Using a Radiometric Insensitive Measure?

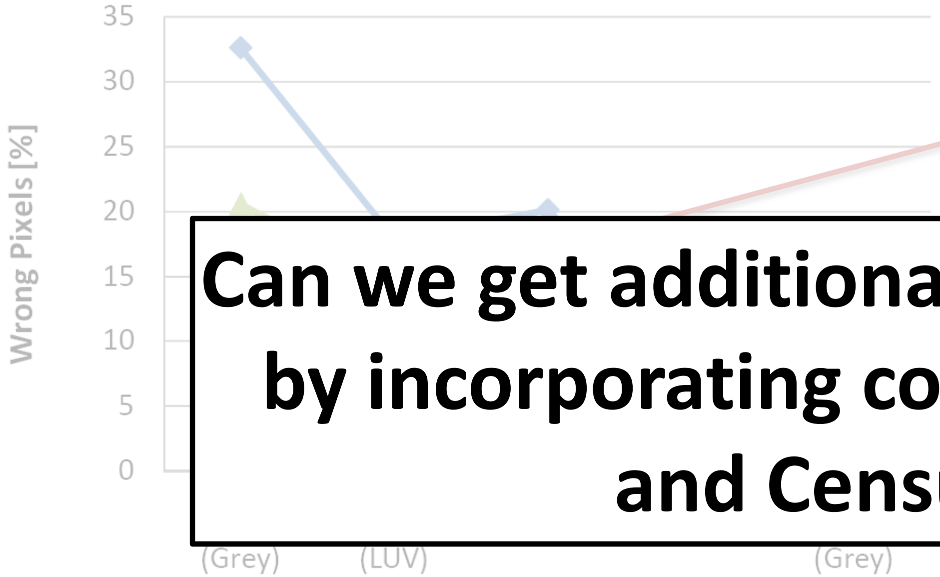

Avg. error
all regions

Can we get additional improvement by incorporating color into ZNCC and Census?

- ZNCC and Census considerably improve performance in **radiometric distorted** regions.
  - They are more effective than color in this respect
- They even improve performance in **radiometric clean** regions.
- ZNCC and Census improve the **overall** performance considerably
  - The overall error drops from **20.5%** [Absdif (Grey)] to **6.7%** [Census (Grey)] **(!)**

Using Color with Radiometric Insensitive Measures

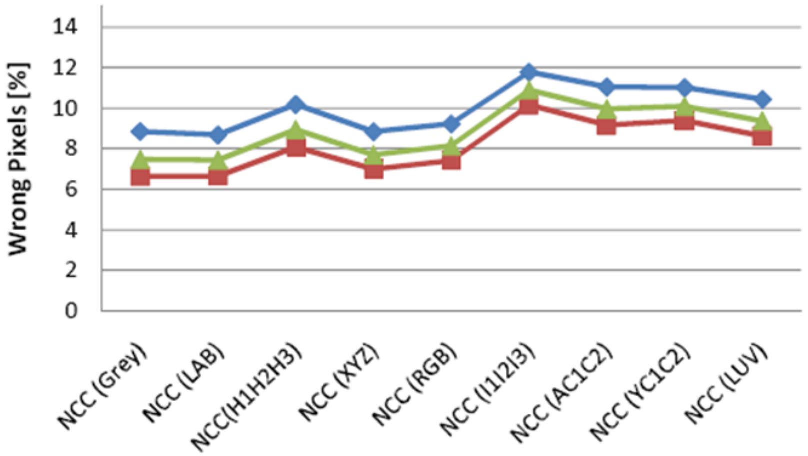

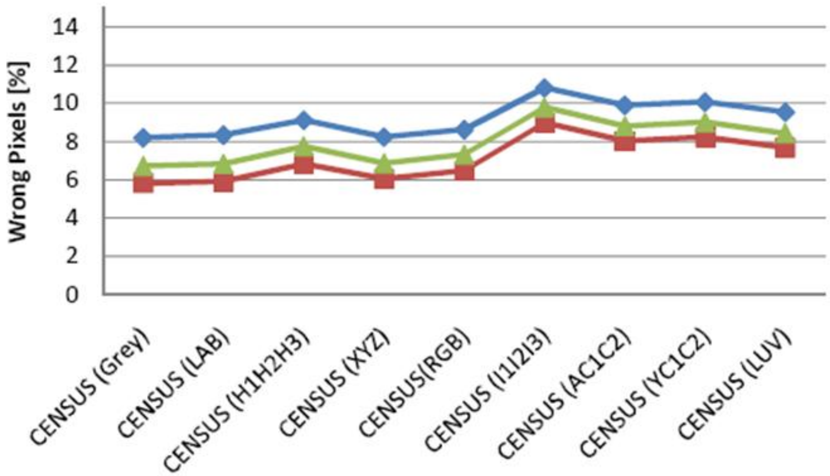

- Seems to be a bad idea:
  - Color even worsens results
- Why is it?
  - Increased robustness of color in radiometric regions is not important anymore
    - NCC and CENSUS do a better job
  - You practically do not lose texture when deleting color
  - Intensity is probably more robustly captured by nowadays cameras (less noise in the intensity channel)

Using Color with Radiometric Insensitive Measures

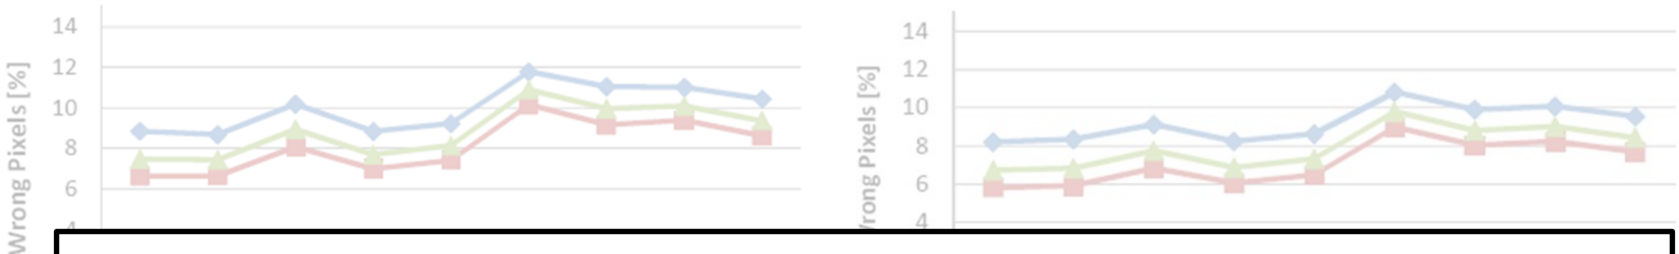

We suggest:

- Use radiometric insensitive match measures

- Do not use color

anymore

– NCC and CENSUS do a better job

- You practically do not lose texture when deleting color
- Intensity is probably more robustly captured by nowadays cameras (less noise in the intensity channel)

Conclusions

- Major argument for color is the increased robustness to radiometric distortions.
- This benefit is low considering that ZNCC and Census do a considerably better job in distorted regions.
- Do not use color in conjunction with ZNCC or Census.

Future work

- Check whether results generalize to other image sets
  - Currently only Middlebury used
- Test other optimization algorithms
  - We have done some preliminary tests with α-expansion and local optimization that confirm our results (not in the paper)
- Use other energy models
  - Second-order smoothness
  - More dissimilarity measures
- Incorporate segmentation-based aggregation schemes
  - Another step on the way to an “optimal” stereo data term

References

- [Bleyer,ISPRS08] M. Bleyer, S. Chambon, U. Poppe, M. Gelautz, Evaluation of Different Methods for Using Colour Information in Global Stereo Matching, Int. Archives of the ISPRS, 2008.
- [Chambon,IJRA05] S. Chambon, A. Crouzil, Colour Correlation-Based Matching, Int. Journal for Robotics and Automation, vol. 20, no. 2, 2005.
- [Mühlmann,IJCV02] K. Mühlmann, D. Maier, J. Hesser, R. Männer. Calculating Dense Disparity Maps from Color Stereo Images, an Efficient Implementation, IJCV, vol. 47, no. 1, 2002.
- [Hirschmueller,PAMI09] H. Hirschmueller, D. Scharstein, Evaluation of Stereo Matching Costs on Images with Radiometric Differences, PAMI, vol. 31, no. 9, 2009.