Computational Photography

Prof. Feng Liu

Spring 2016

http://www.cs.pdx.edu/~fliu/courses/cs510/

05/26/2016

Last Time

- Stereoscopic 3D
  - Human depth perception
  - 3D displays

Today

- Stereoscopic 3D
  - 3D Cinematography
  - Stereoscopic media post-processing

Stereoscopic 3D

History

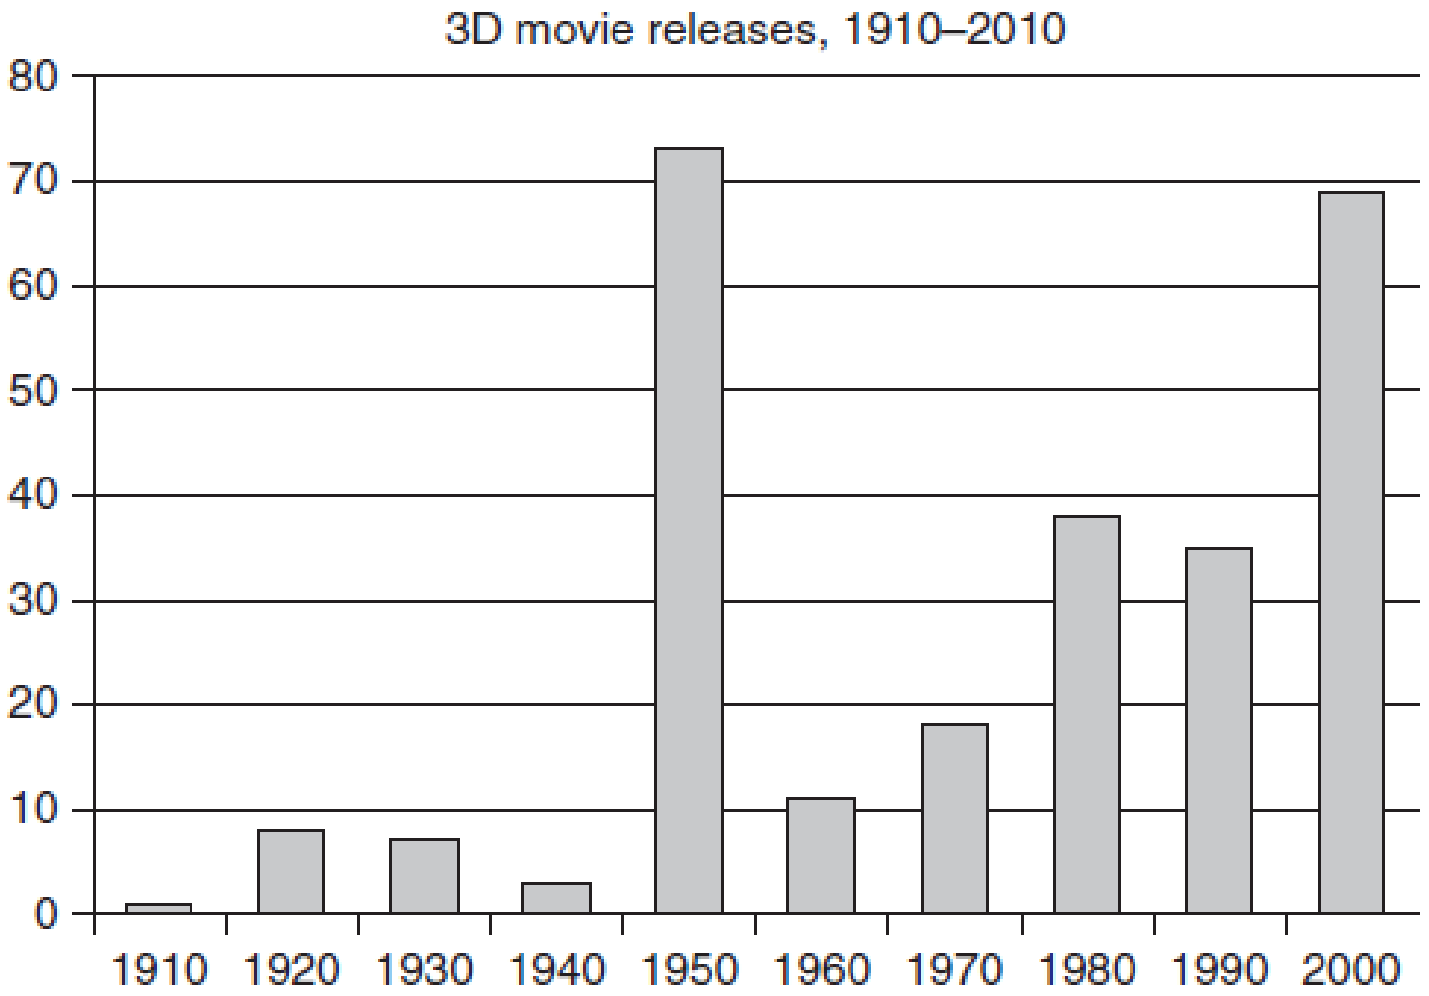

Reprint from *3D Movie Making: Stereoscopic Digital Cinema from Script to Screen*

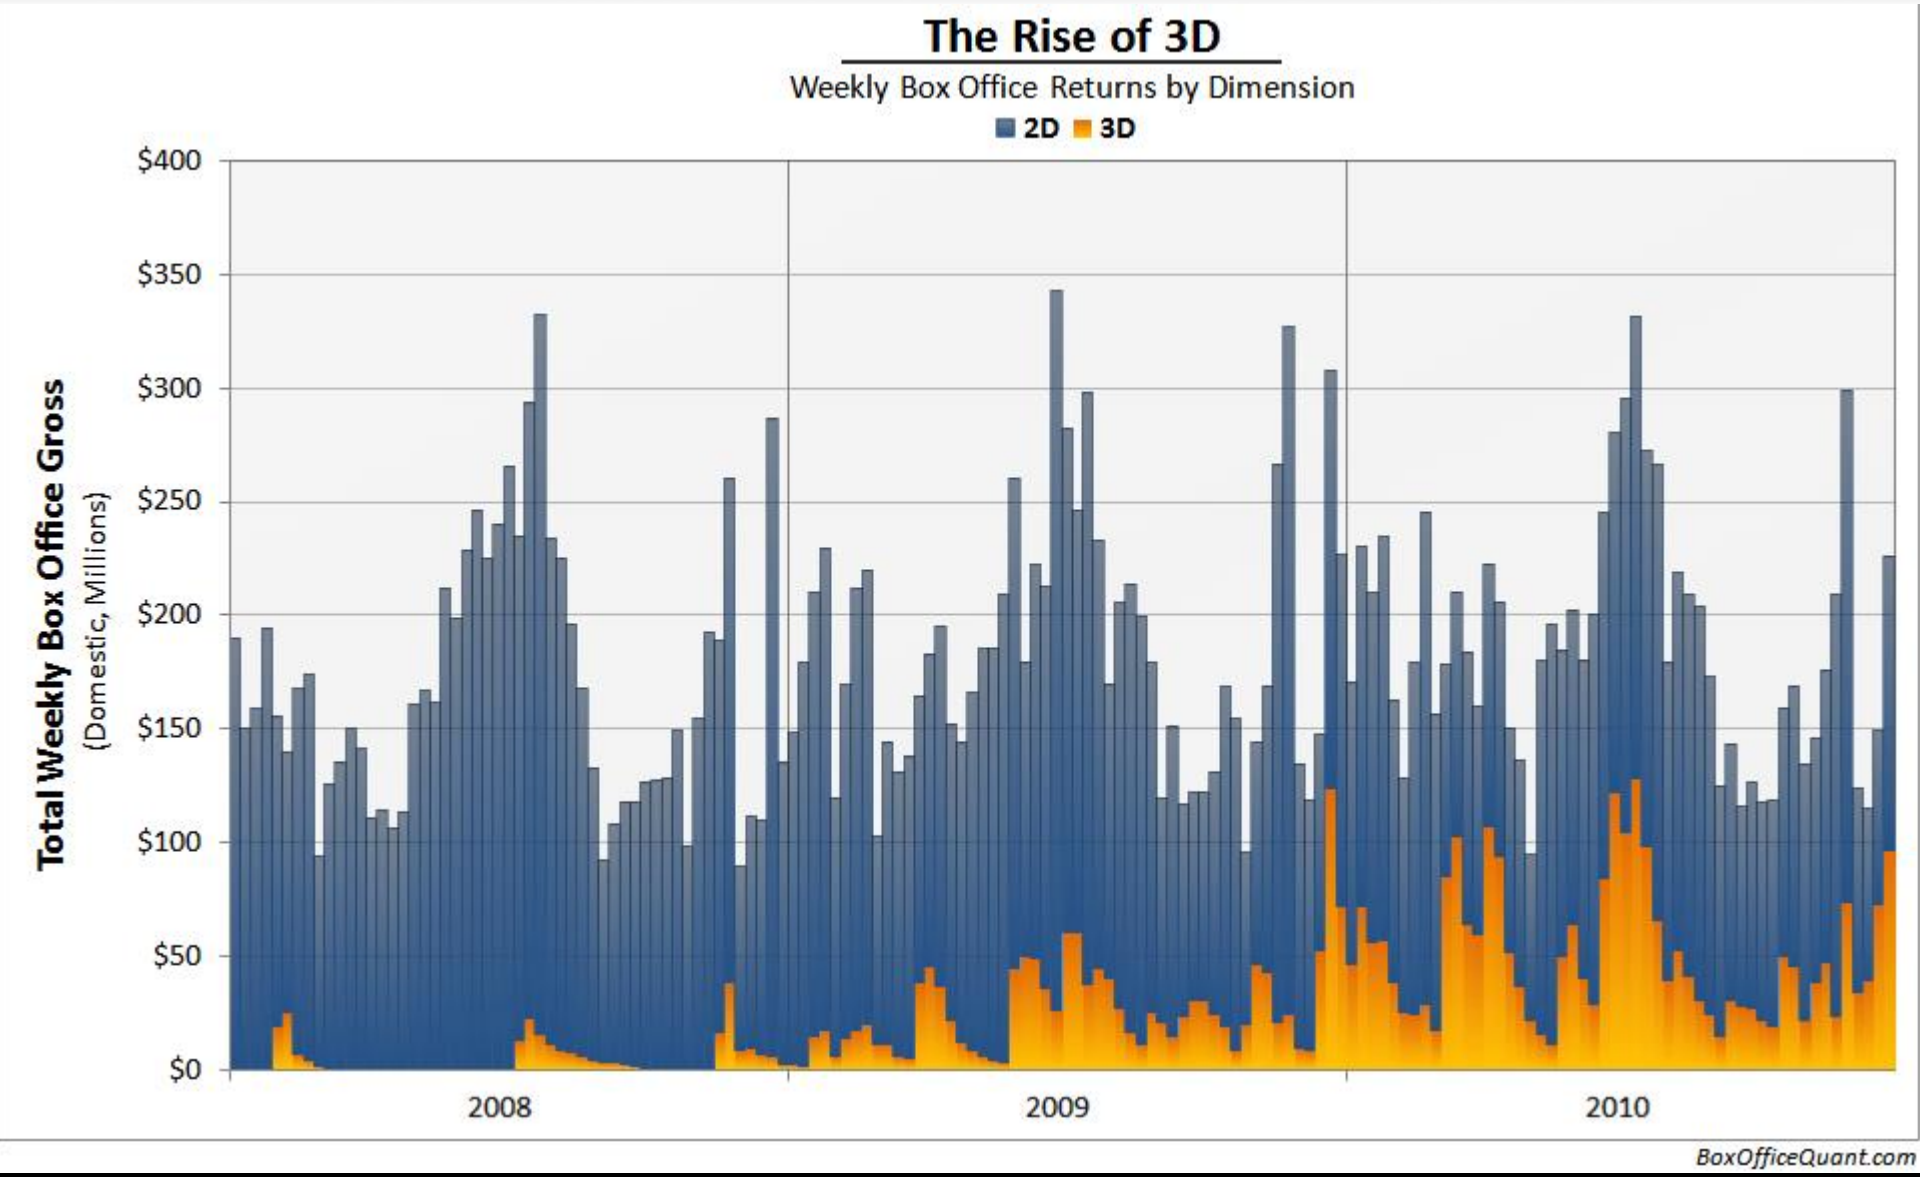
The Rise of 3D
Weekly Box Office Returns by Dimension
2D
3D
Total Weekly Box Office Gross
(Domestic, Millions)
$400
$350
$300
$250
$200
$150
$100
$50
$0
2008
2009
2010
BoxOfficeQuant.com

Ubiquitous Stereoscopic 3D

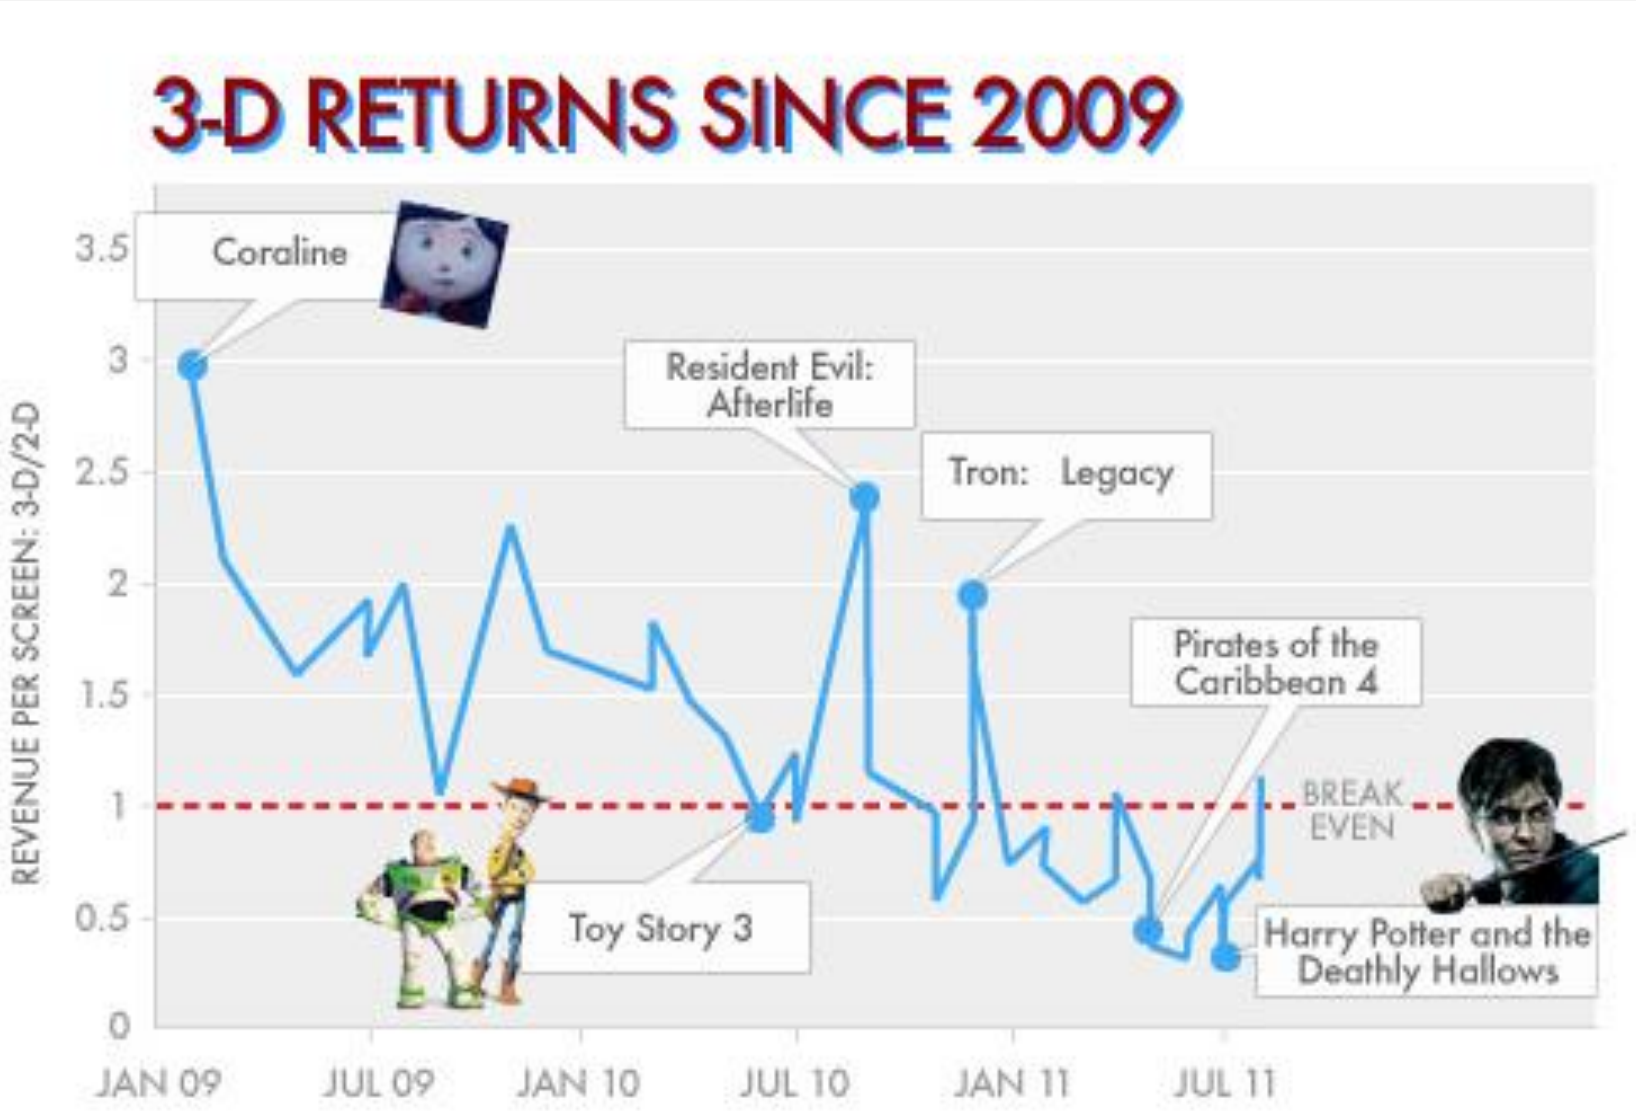

Image source: http://www.slate.com/id/2303814/

Image source: http://www.slate.com/id/2303814/

3D Fatigue

- Blurring vision
- Eyestrain
- Headache

Image source:
http://www.digitalproductionme.com/article-4580-3d--bad-for-you/#.UI-QfGdTDK0

Stereoscopic 3D Camera

Stereo Photo

Left

Right

Red-cyan anaglyph

Stereopsis

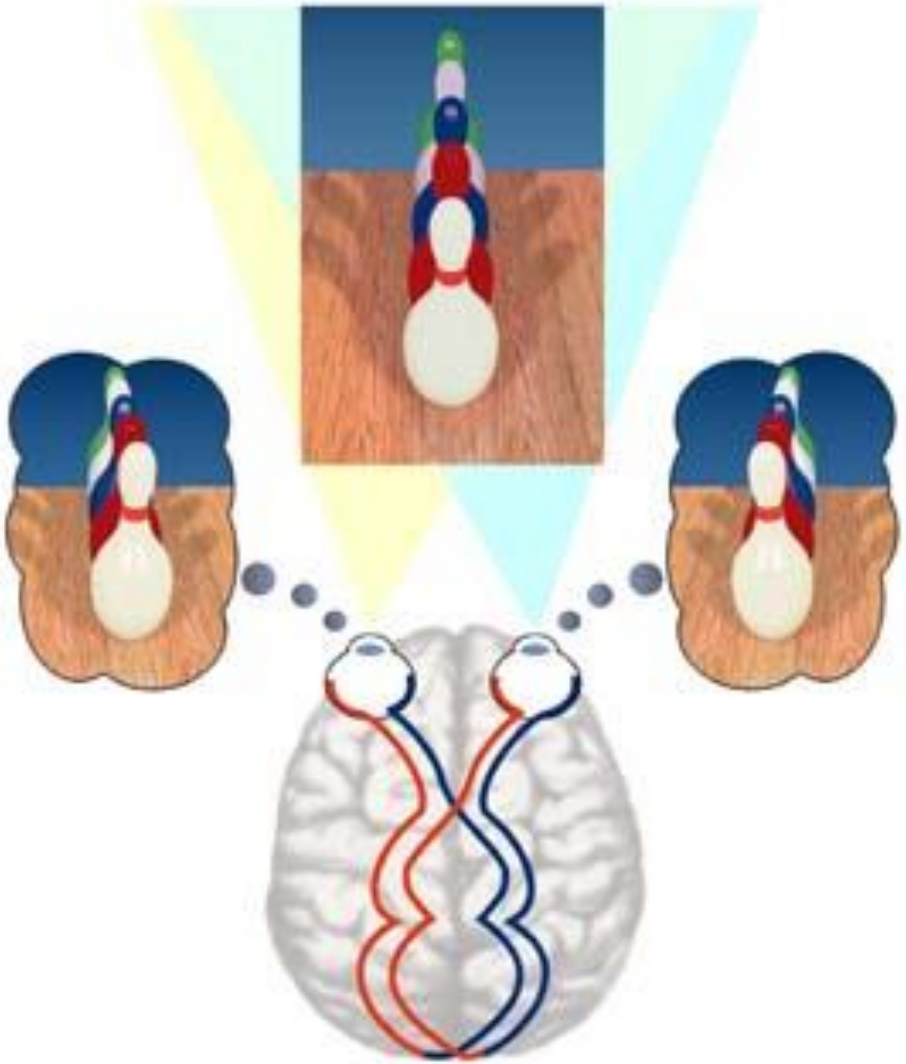

Disparity and Perceived Depth

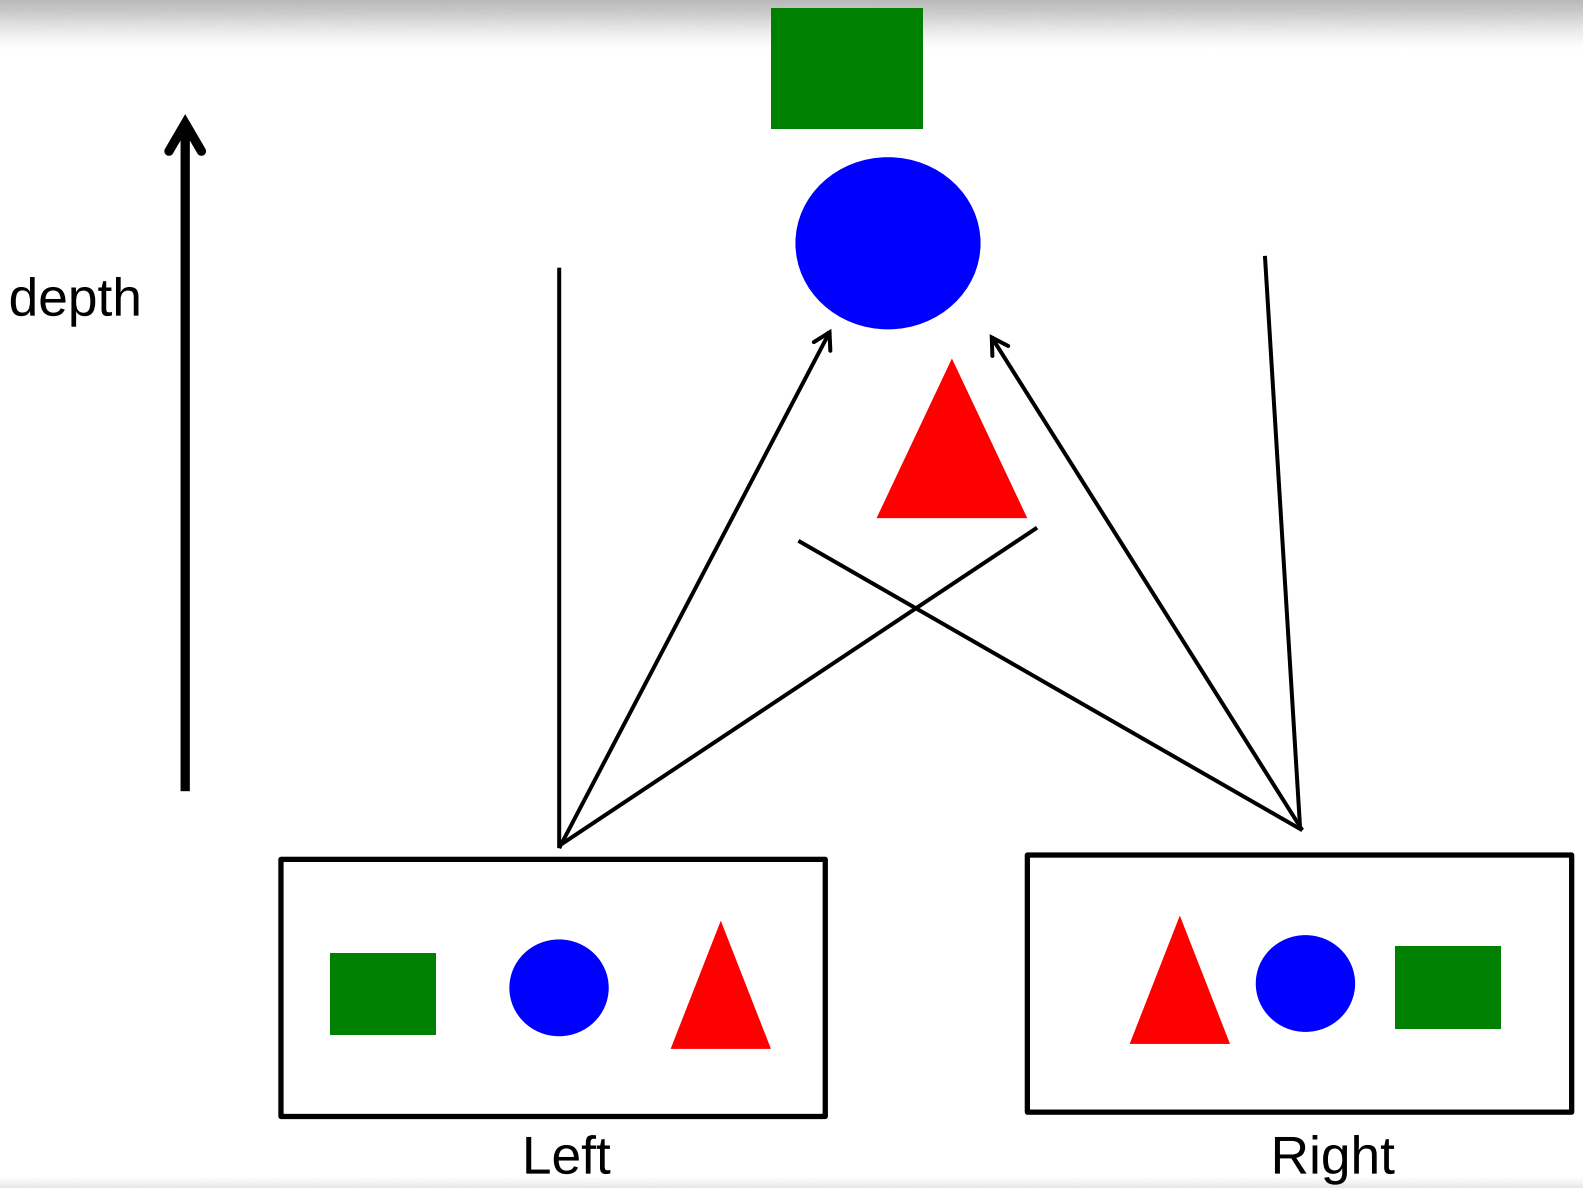

Disparity and Perceived Depth

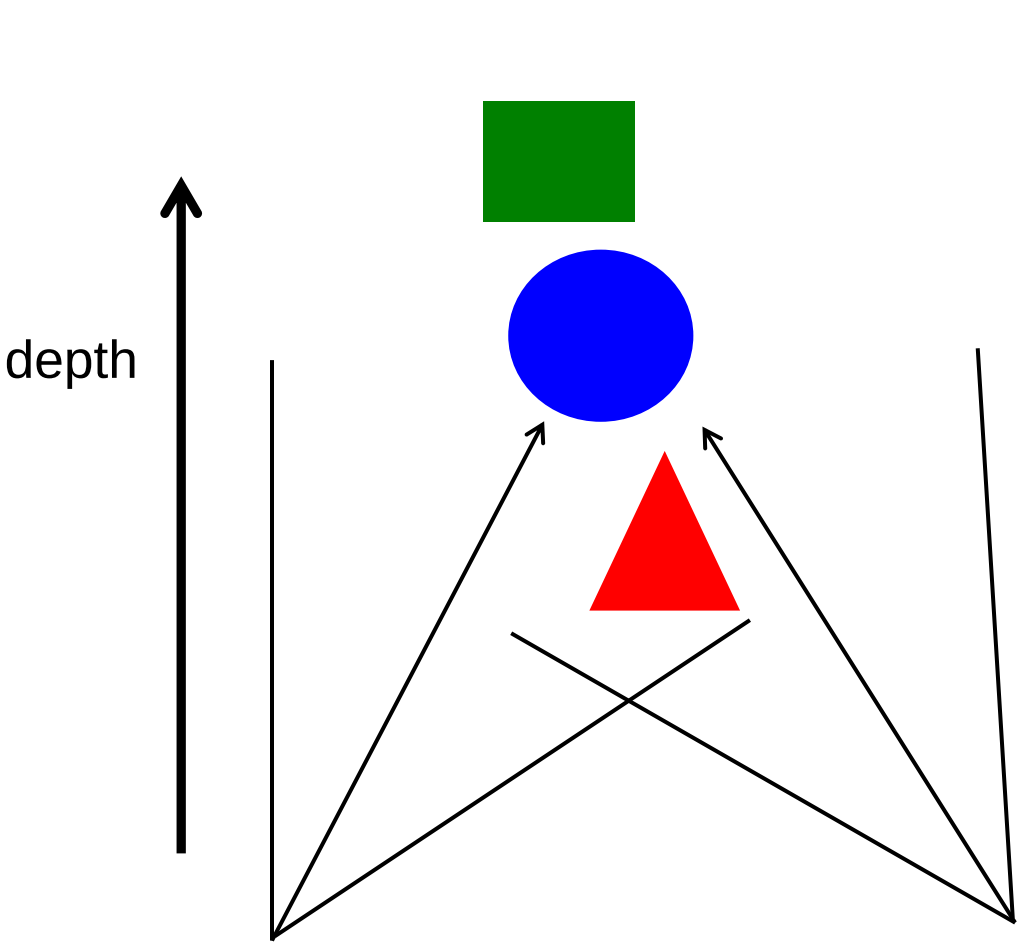

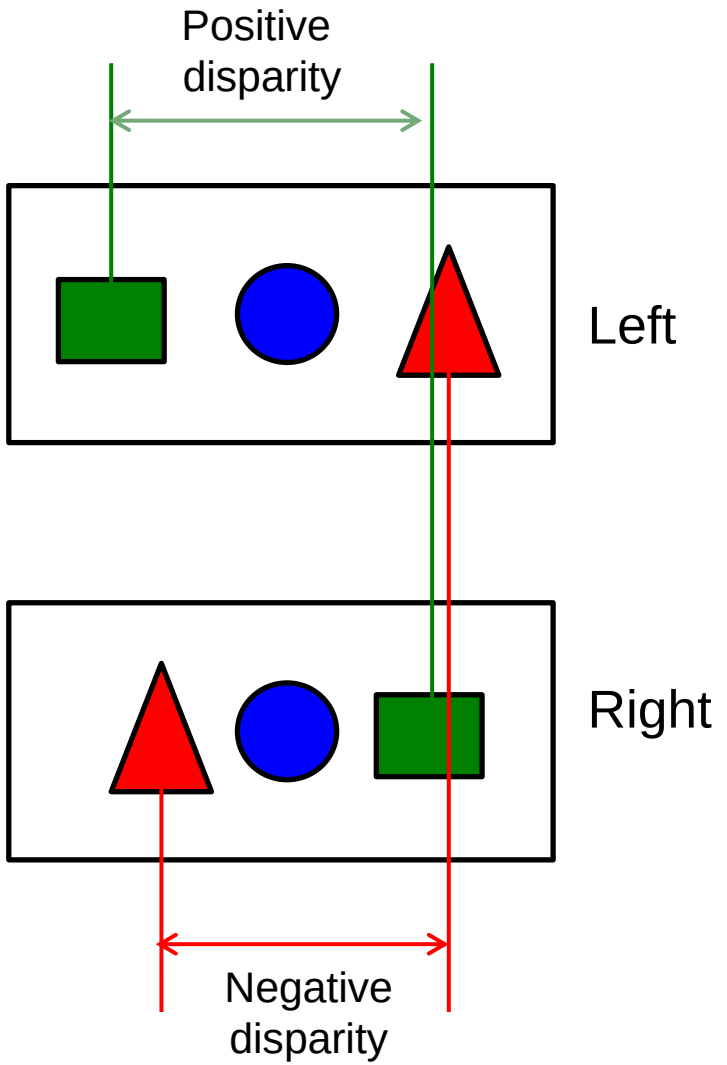

Vergence-accommodation

- There is an area around it where vergence and accommodation agree, which is called zone of comfort. This discrepancy could damage the visual acuity before the age of 8.

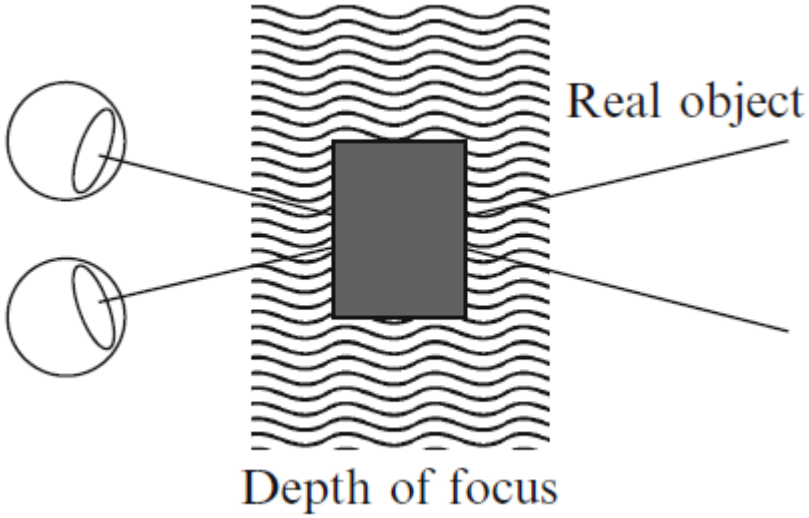

real world

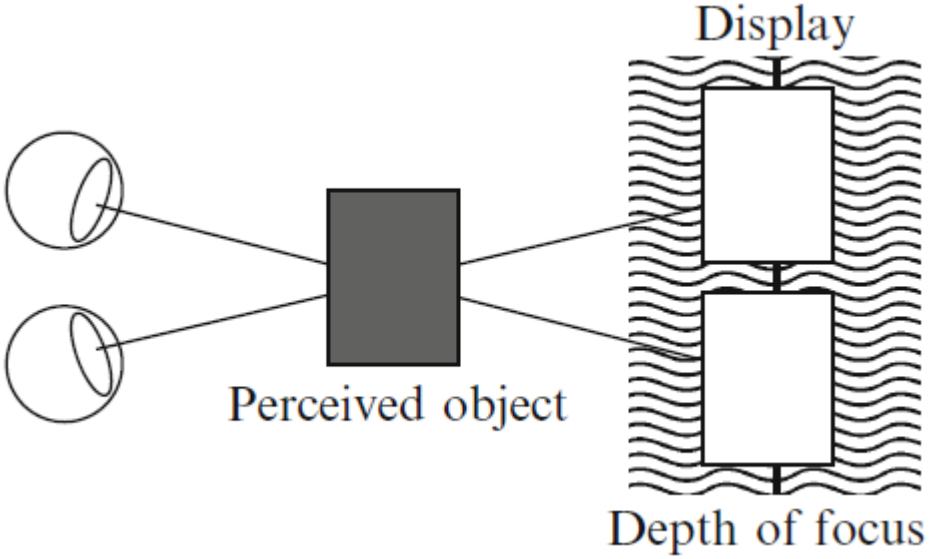

stereoscopic displays

Stereoscopic Comfort Zone

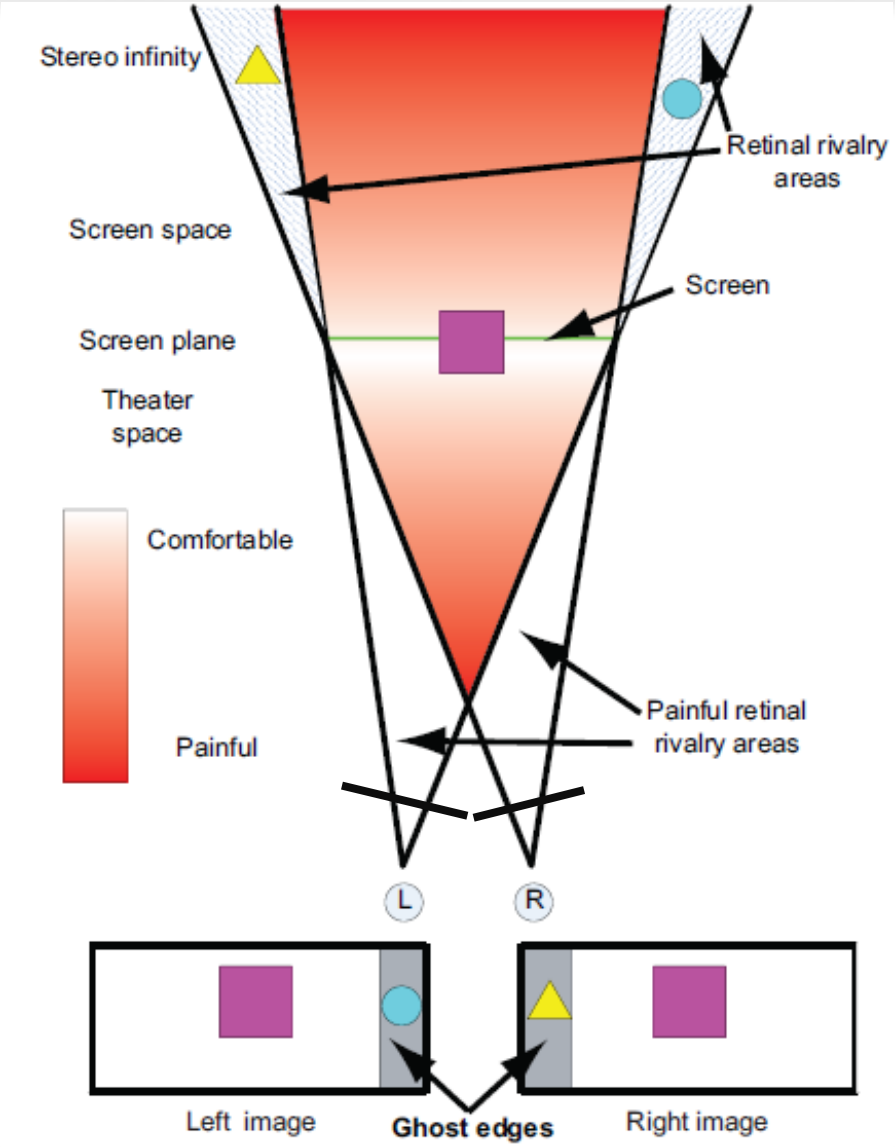

Monocular Object

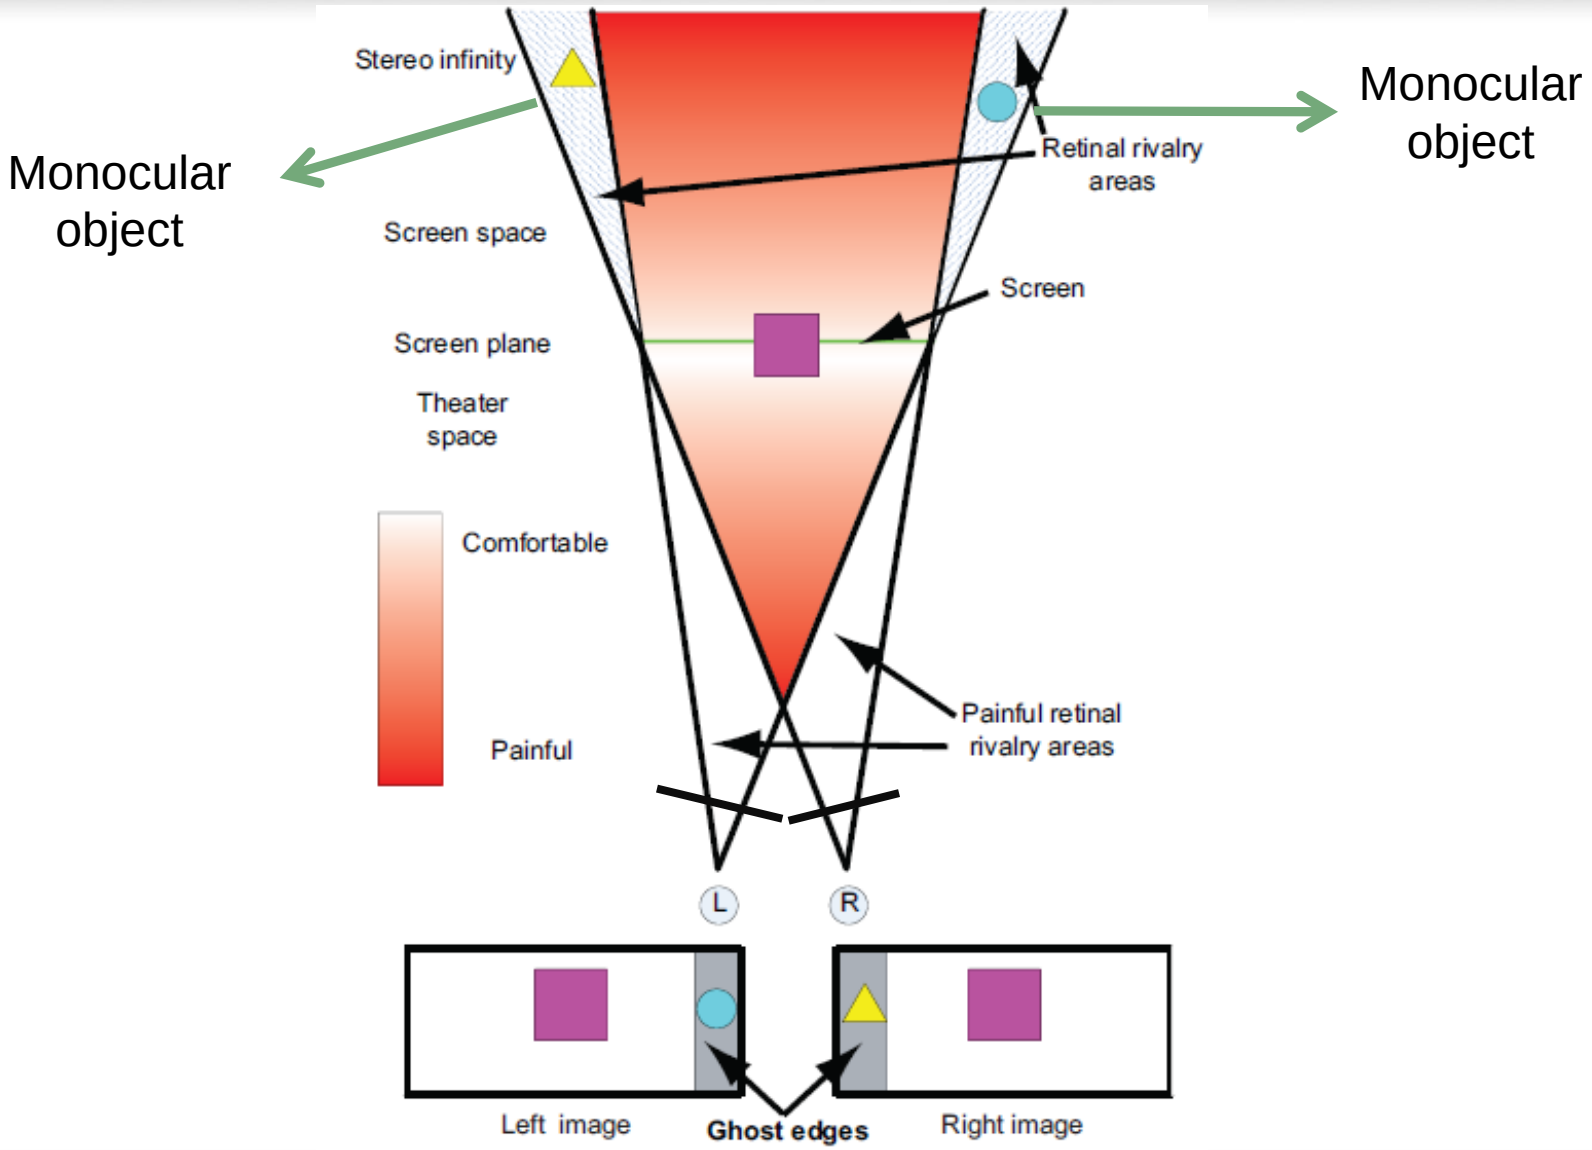

Stereo Window Violation

(a) Ideal 3D perception

(b) Actual 3D perception

When an object with negative disparities is cut by the screen edge, it suffers from the stereo window violation. That is, the object is perceived in front of the screen, but is occluded by the screen edge.

Stereo Window Violation

Left

Right

More Visual Fatigue Sources

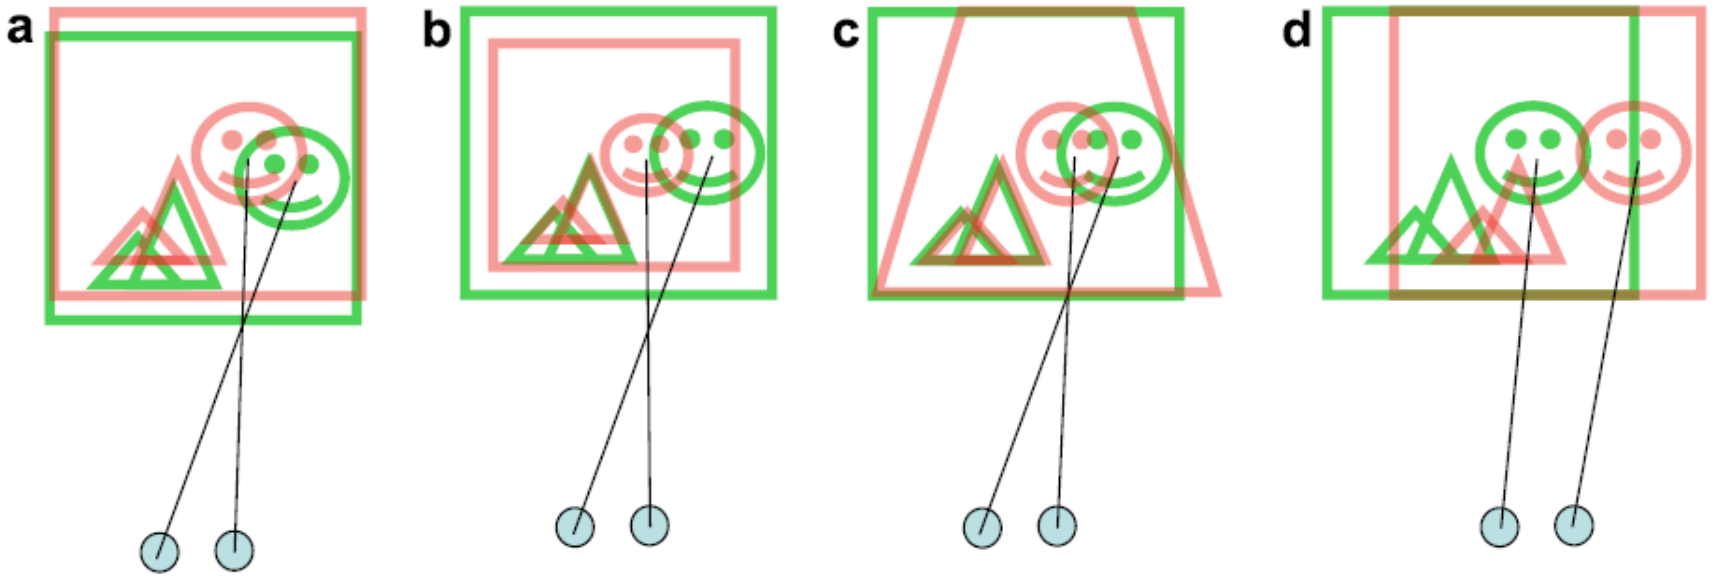

Slide credit: Y.Y Chuang

Enable Warping on Stereoscopic Images

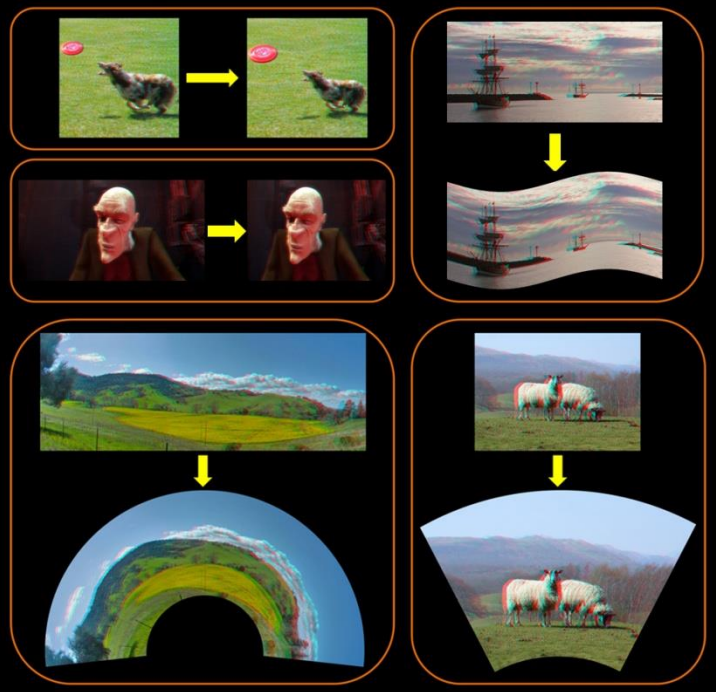

Yuzhen Niu
Wu-Chi Feng
Feng Liu

Portland State University

Portland State UNIVERSITY

SIGGRAPH ASIA 2012

Image Warping

Naïve Stereo Warping

Left

Right

Naïve Stereo Warping

Left

Right

90° Rotation

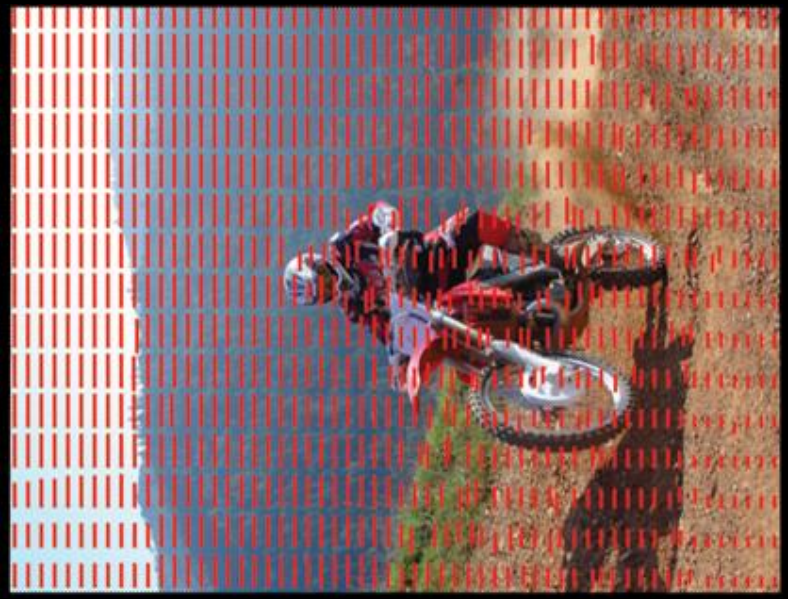

Left

Right

More Stereo Warping

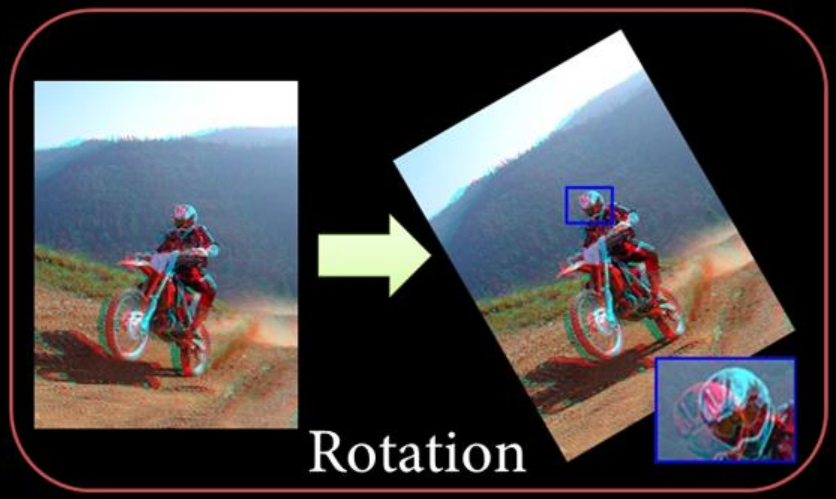

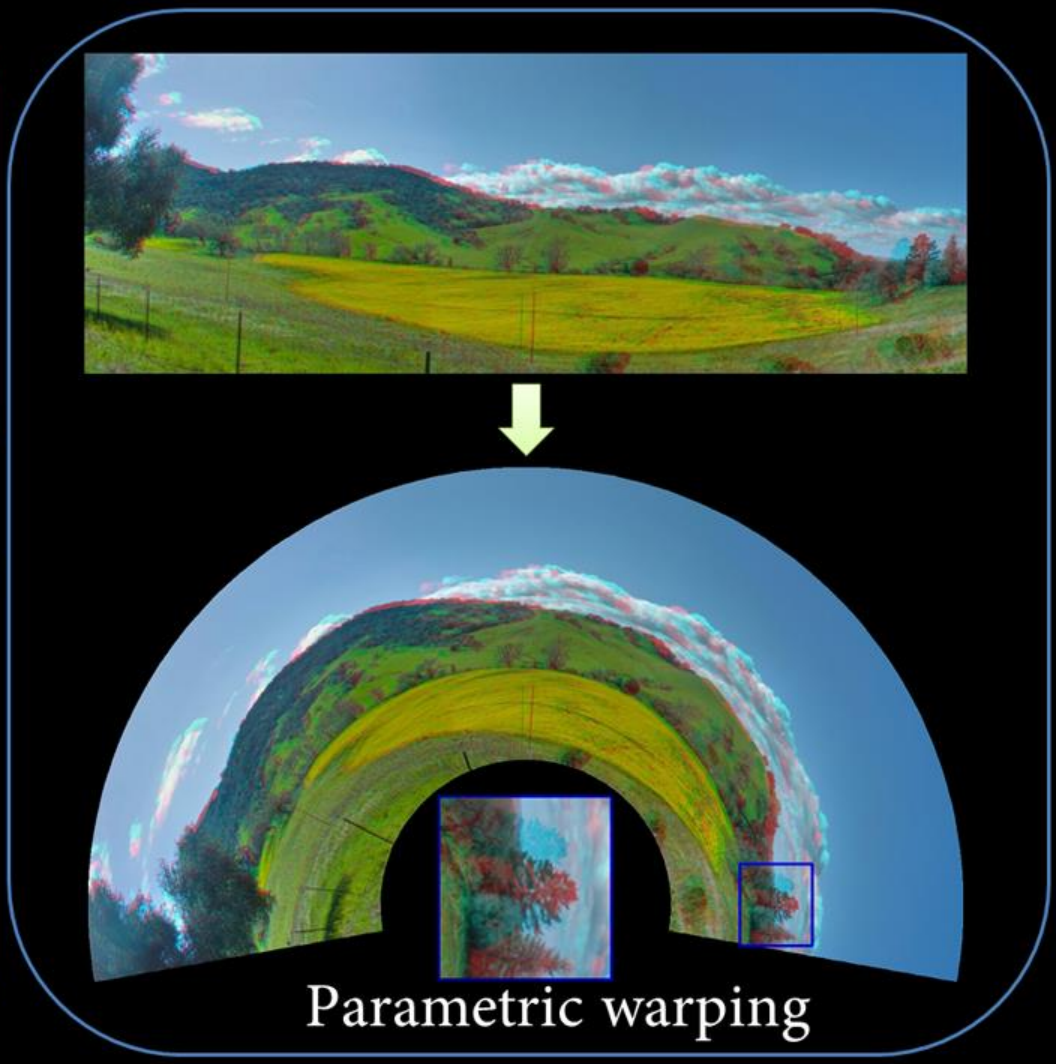

Goals for Stereo Warping

- Warp the left and right view consistently
- Avoid introducing vertical disparities
- Maintain good horizontal disparities

Naïve Stereo Warping

- Warp the left and right view consistently
- ~~Avoid introducing vertical disparities~~
- ~~Maintain good horizontal disparities~~

Our Solution

- Warp the left image using the user-specified warping
- Estimate the target disparity map for the warping result
- Warp the right image guided by the target disparity map

Overview

Left input

Right input

Overview

Left result

Right input

Overview

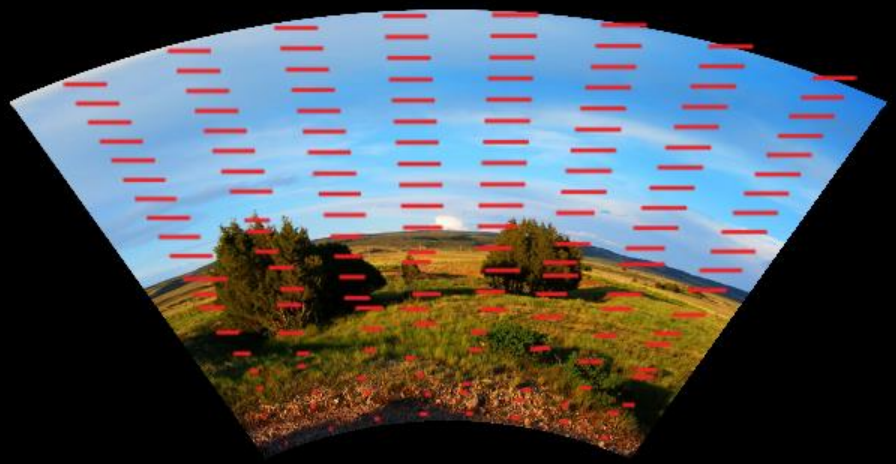

Left result
w/target disparity map

Right input

Overview

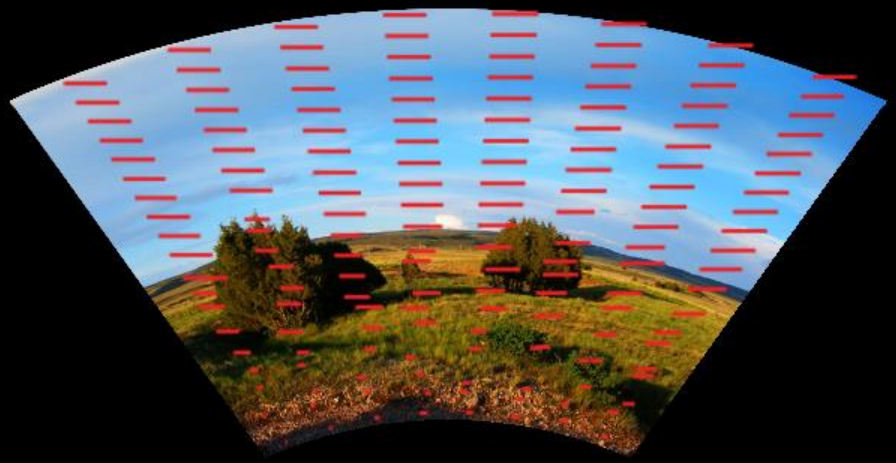

Left result

Right result

Overview

Input

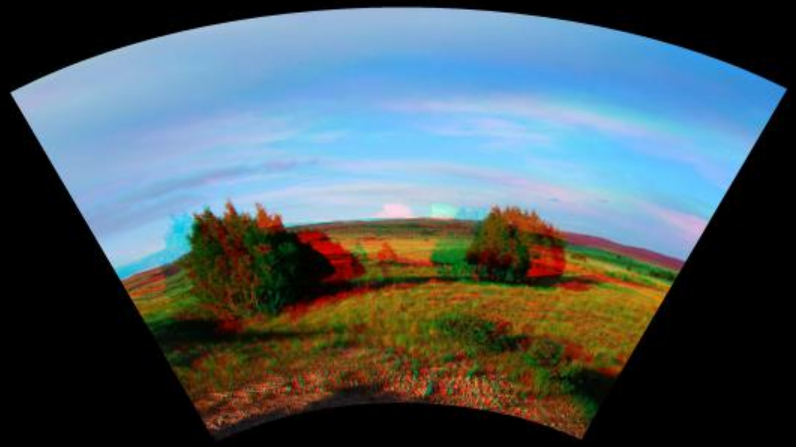

Warping result

Our Solution

1. Warp the left image using the user-specified warping
2. Estimate the target disparity map for the warping result
3. Warp the right image guided by the target disparity map

Our Solution

1. Warp the left image using the user-specified warping
2. Estimate the target disparity map for the warping result
   - Estimate the original disparity map
     - SIFT-based feature matching [Lowe 2004] and
     - Optical flow [Sun et al. 2010]
   - Compute the target disparity map
3. Warp the right image guided by the target disparity map

What Is A Good Disparity Map

√ No vertical disparities

What Is A Good Disparity Map

√ No vertical disparities

? Preserve original (horizontal) disparities

Uniform Scaling

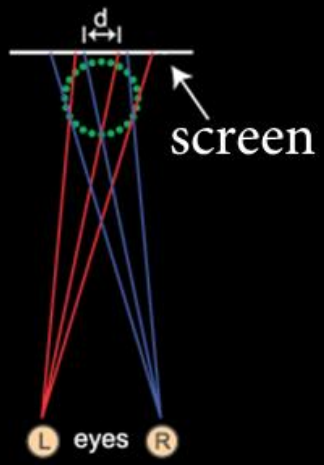

Input

Uniform Scaling

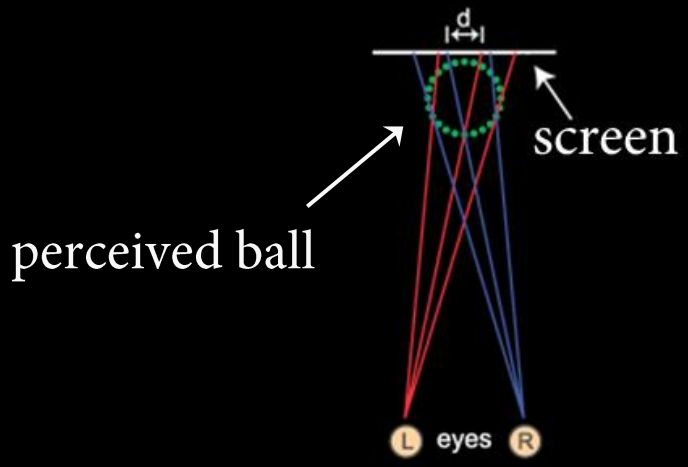

Input

Uniform Scaling

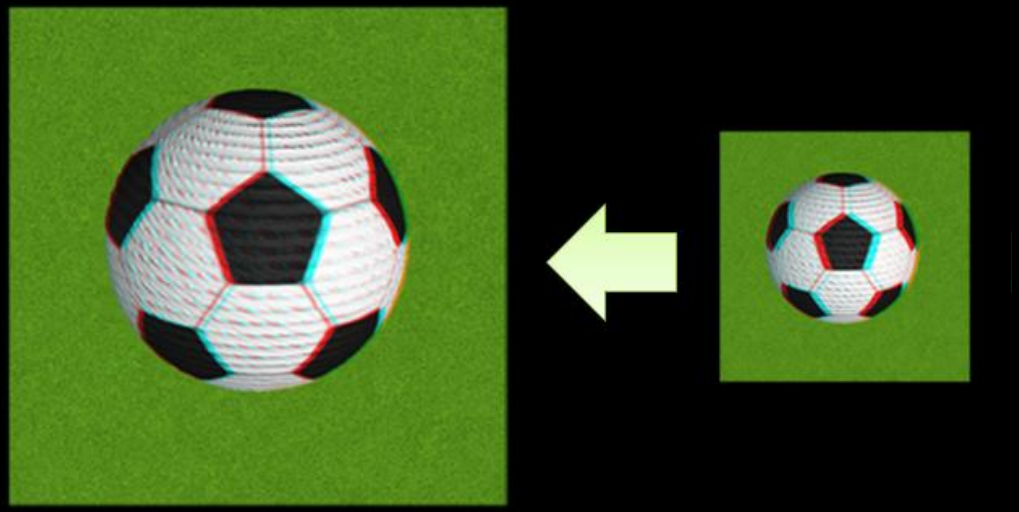

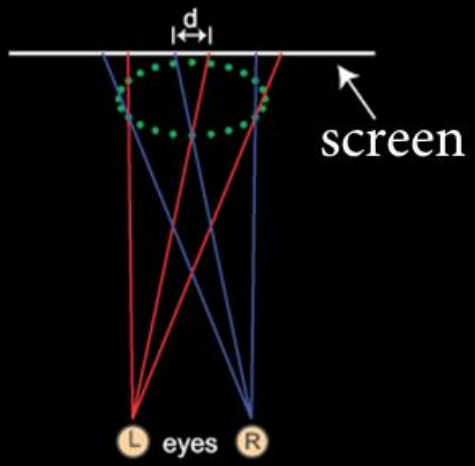

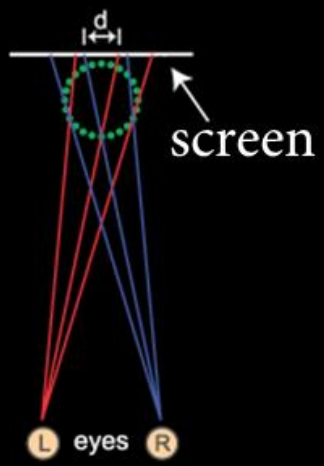

Flattened object

Input

Our result

Uniform Scaling

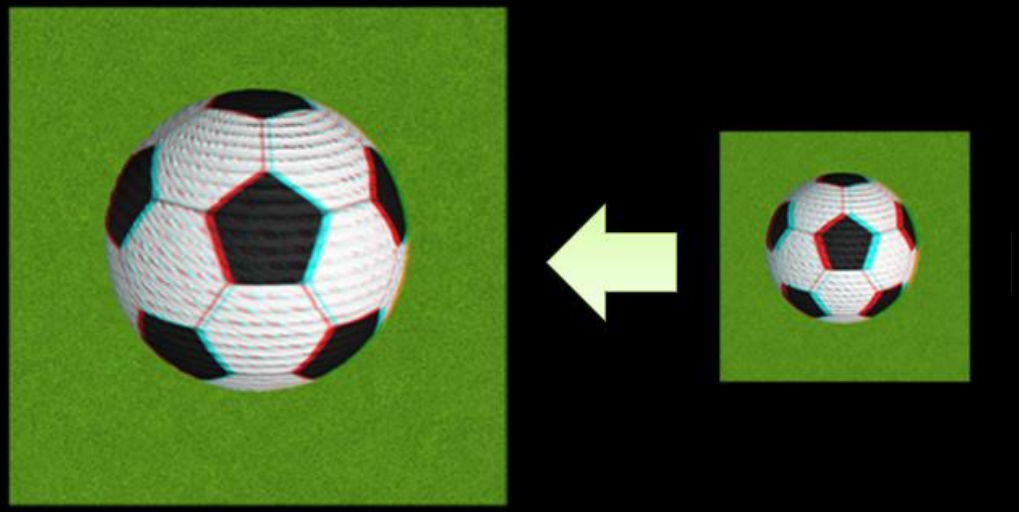

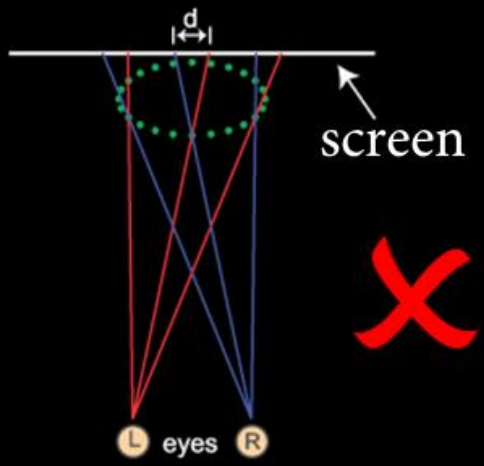

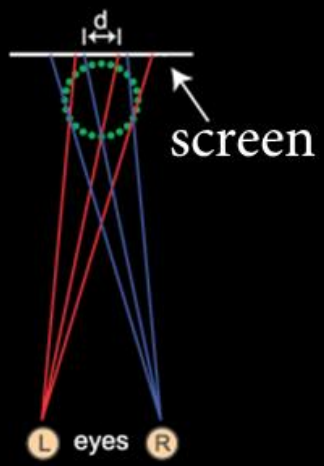

Flattened object

Input

Our result

Uniform Scaling

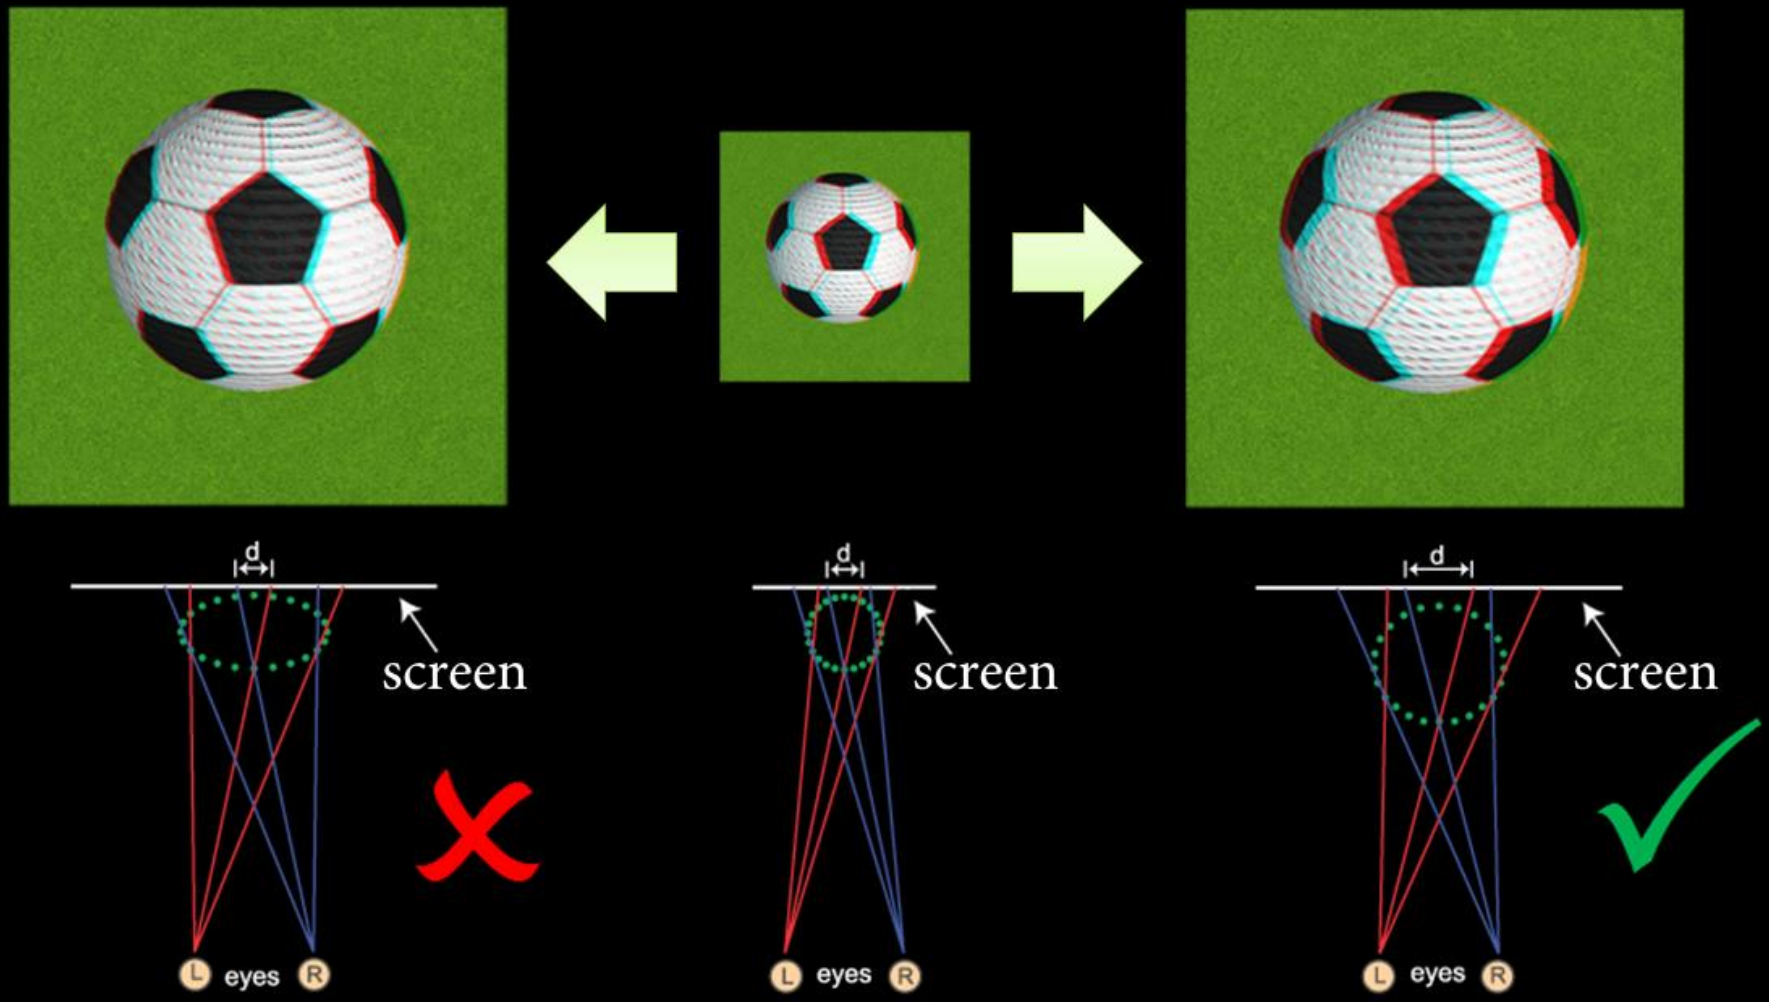

Flattened object

Input

Our result

What Is A Good Disparity Map

√ No vertical disparities

√ Preserve perceived 3D shape

Target Disparity Map Estimation

Goal: preserve the perceived 3D shape

Idea: match the perceived depth range change to the local warp in 2D image space

Perceived Depth

- Depend on
  - Raw image disparity
  - Viewing distance
  - Screen size

Perceived Depth

- Depend on
  - Raw image disparity
  - Viewing distance ?
  - Screen size ?

Target Disparity Map Estimation

Goal: preserve the perceived 3D shape

Idea: match the perceived depth range change to the local warp in 2D image space

Target Disparity Map Estimation

Goal: preserve the perceived 3D shape

Idea: match the disparity range change to the local warp in 2D image space

Target Disparity Map Estimation

Goal: preserve the perceived 3D shape

Idea: match the disparity range change to the local warp in 2D image space

$$\sum_{d_i} \sum_{d_j \in N(d_i)} \left(\left(\hat{d}_i - \hat{d}_j \right) - s_i \left(d_i - d_j \right) \right)^2$$

$$s.t. \quad \hat{d}_{\min} = s d_{\min}$$

$\hat{d}$: target disparity

d : input disparity

s_i : local image scaling factor

Keep the objects with small disparities in the comfort zone after warping

Local Image Scaling Factor

- Find the best fitting similarity transform to the local image warping

- Use the similarity transformation scaling factor

Perceived Depth vs Disparity

- Perceived depth does not linearly depends on disparity
- But, when objects have small disparities, perceived depth nearly linearly depends on disparity
  - objects are closer to the screen
  - small screen sizes

Our Solution

- Warp the left image using the user-specified warping
- Estimate the target disparity map for the warping result
- Warp the right image guided by the target disparity map

Disparity-guided Image Warping

- Content-preserving warping
  - As-rigid-as-possible shape manipulation [Igarashi et al. 2005]
  - Feature-aware texturing [Gal and Cohen-Or 2006]
  - Content-preserving warps for 3d video stabilization [Liu et al. 2009]
  - Nonlinear disparity mapping for stereoscopic 3d [Lang et al. 2010]

Content-preserving Warping

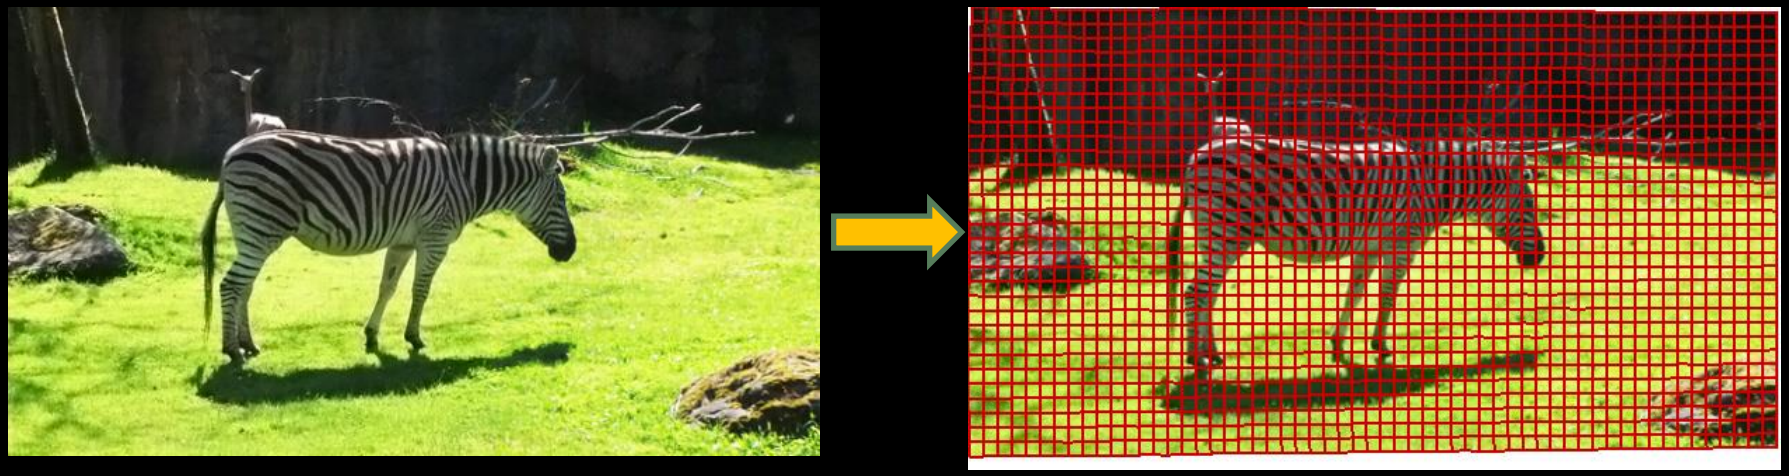

- Build a grid mesh from input image
- Warp input image by least-squares minimization
  - Data term: move features to target positions
  - Smoothness term: avoid visual distortion
    - Similarity transformation constraints
  - Solved by a linear solver

Content-preserving Warping

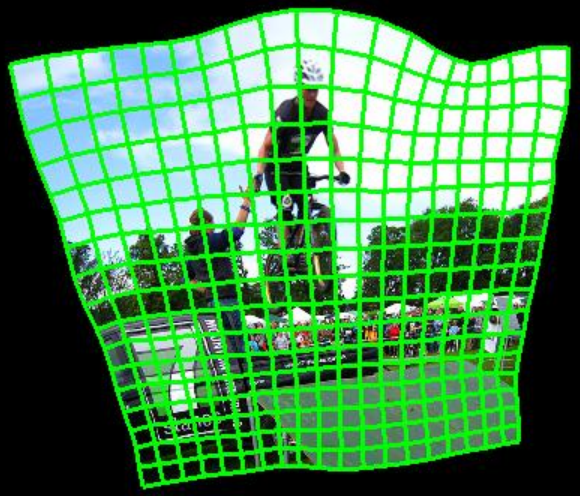

Left result

Right result

Our Observation

Left

Right

Left

Right

Left

Right

Our Observation

Left

Right

Left

Right

Left

Right

Similar

Pre-warping

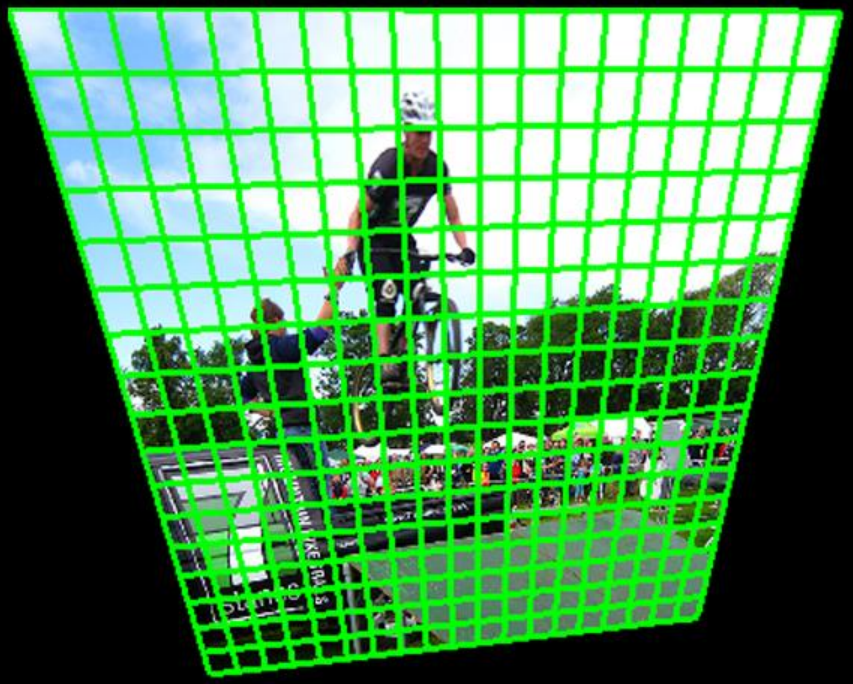

Left result

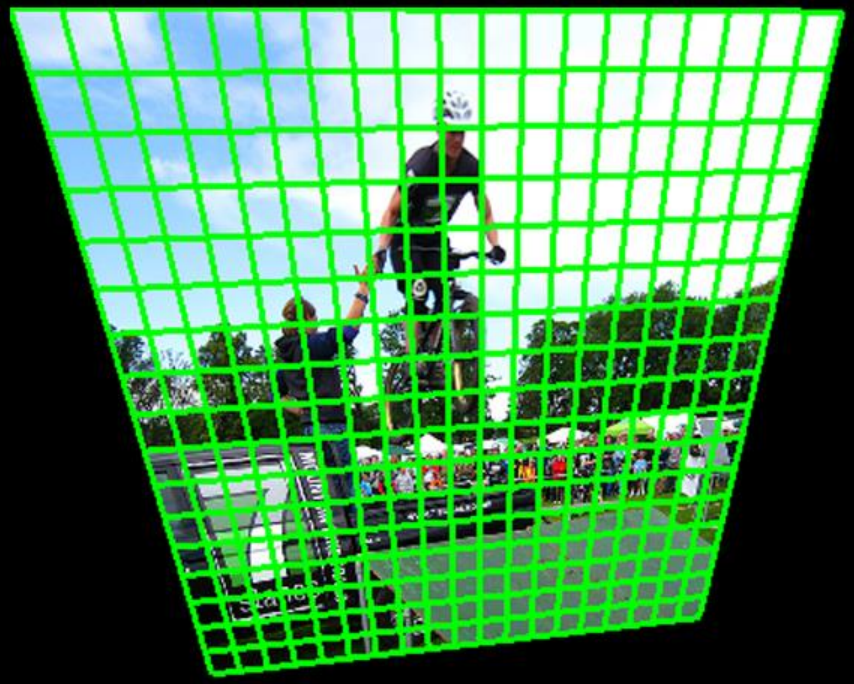

Right pre-warping result

Pre-warping

Left result

Right pre-warping result

Content-preserving Warping

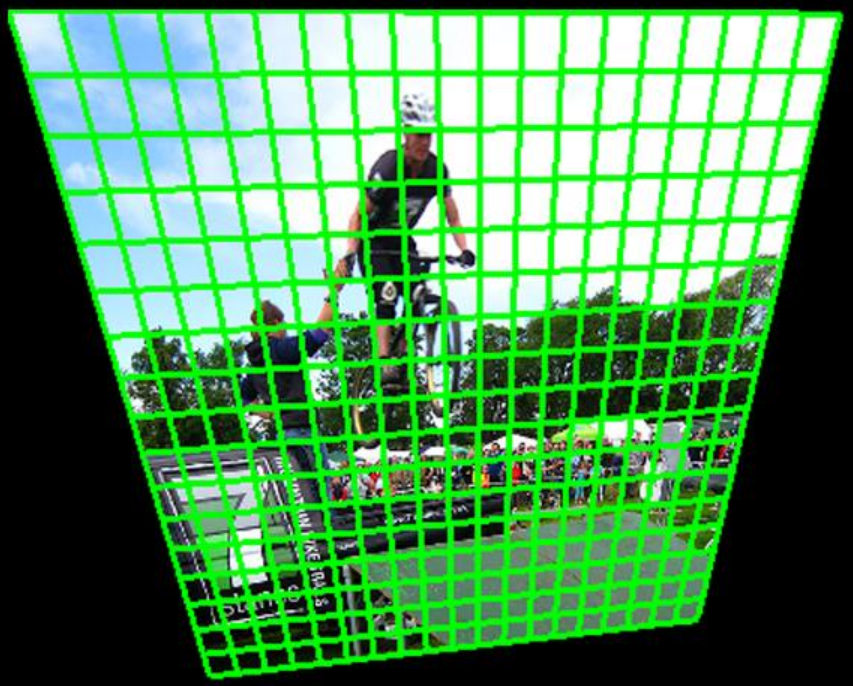

Left result

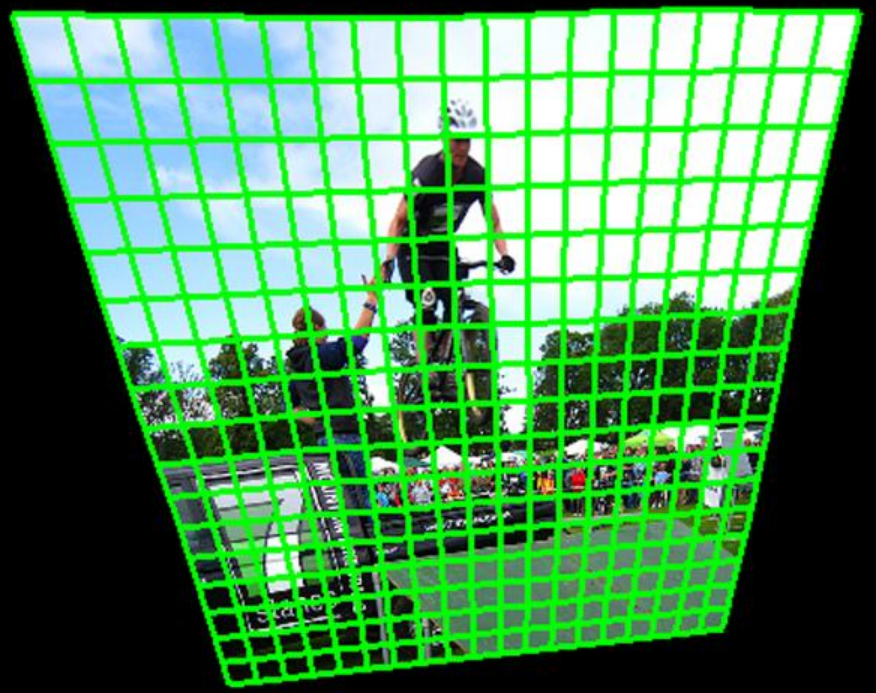

Right result

Content-preserving Warping

Left result

Right result

Input

Warping result

Result: rotation

- No change to disparity map
  - Rotating a rectified stereo camera will not change the disparity map

Result: rotation

Input

Our result

Result: rotation

Input

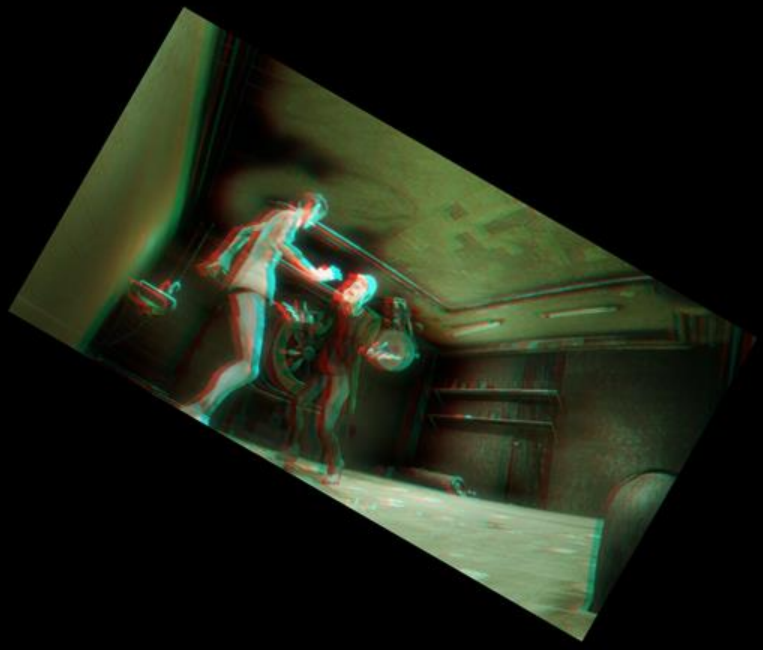

Our result

Result: rotation

Input

Our result

Result: rotation

Input

Our result

Result: rotation

Input

Our result

Result: rotation

Input

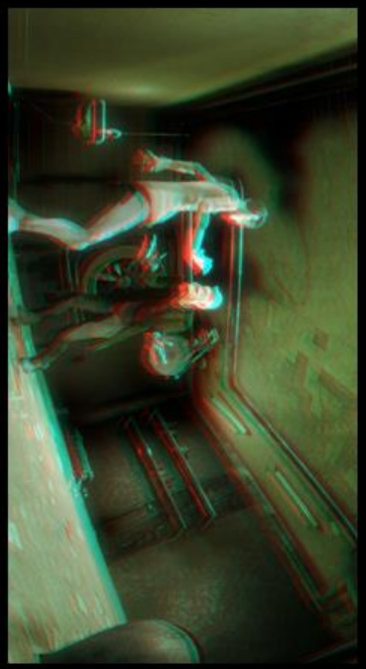

Our result

Result: similarity transformation

Input

Our result

Uniformly scaling the disparity values

Result: affine transformation

Input

Our result

Uniformly scaling the disparity values

Result: perspective transformation

Input

Our result

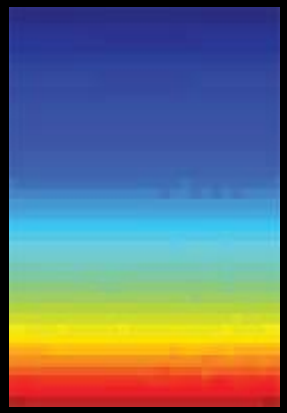

Scaling map

More parametric warping

Input

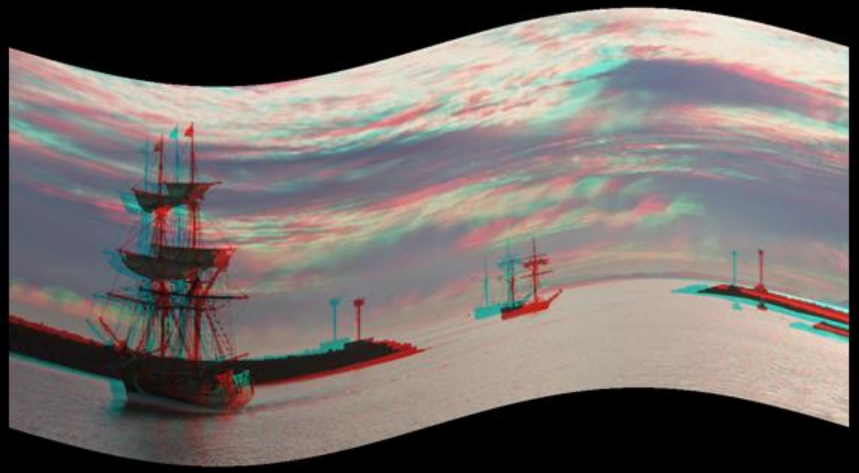

Our result

More parametric warping

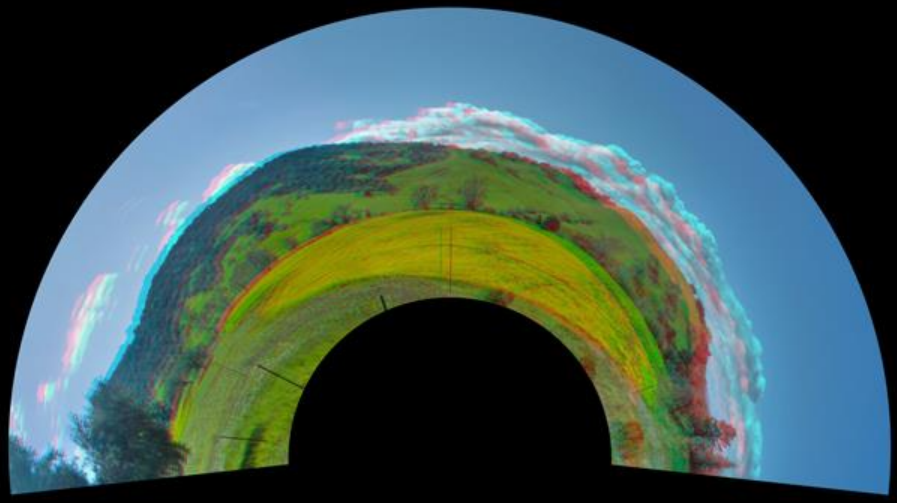

Input

Our result

Image Retargeting

Input

Our result

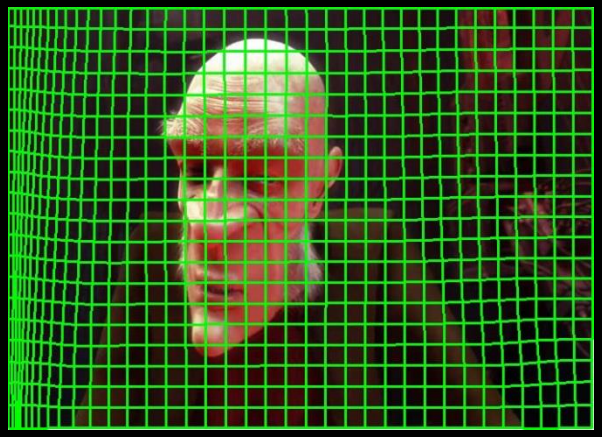

Left

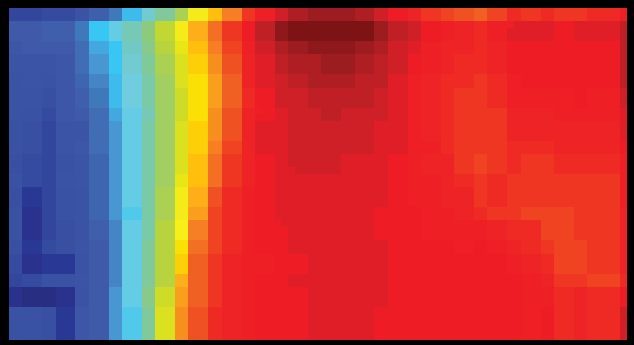

Scaling map

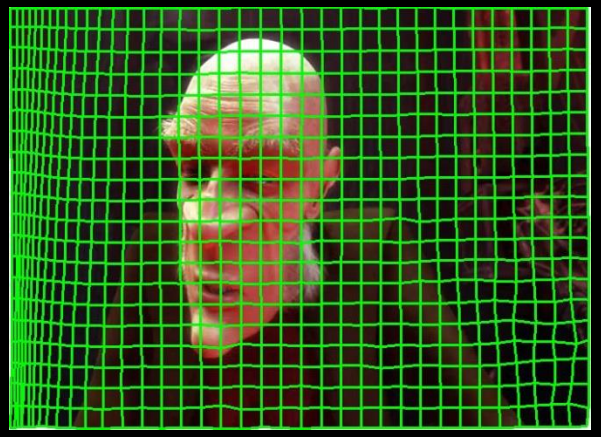

Right

Object Resizing

Input

Our result

Content-preserving Warping

Input

Our result

Keystone Correction for Stereoscopic Cinematography

Feng Liu, Yuzhen Niu, and Hailin Jin. Keystone Correction for Stereoscopic Cinematography. IEEE 3-D Cinematography 2012, Providence, RI, June 2012

Stereoscopic Comfort Zone

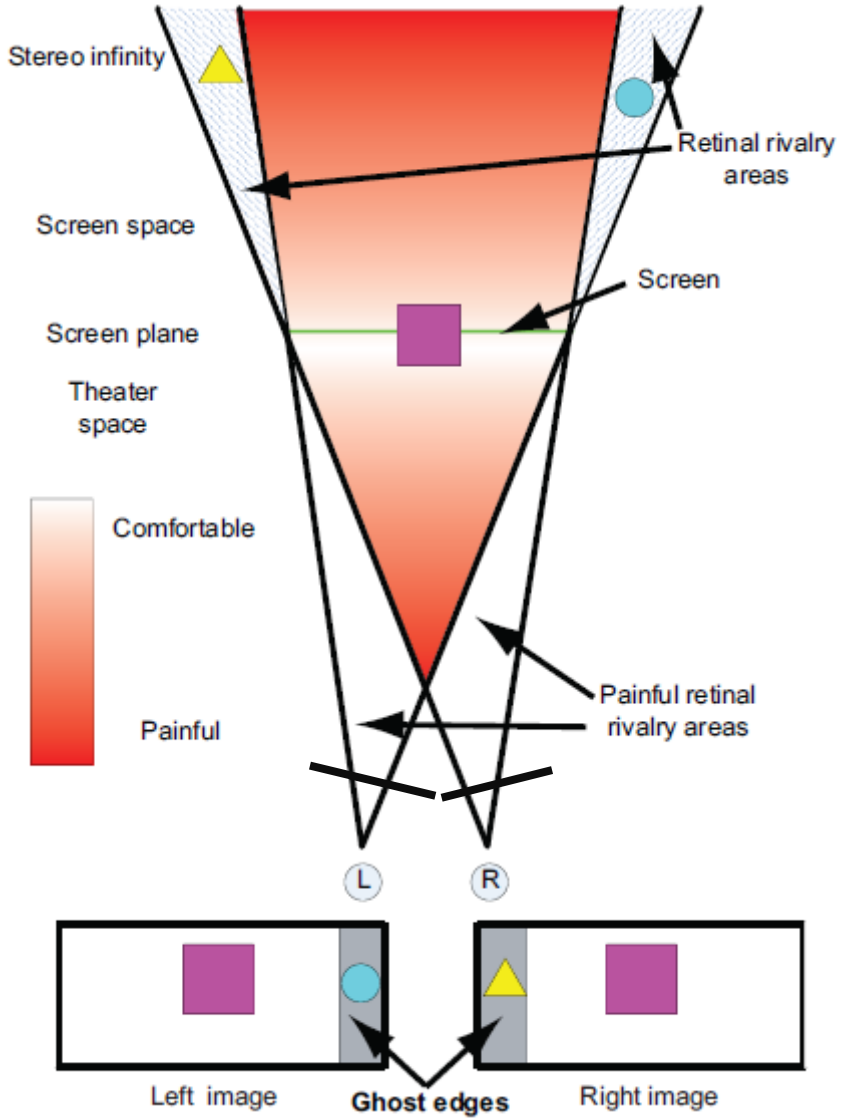

Stereoscopic Camera Model

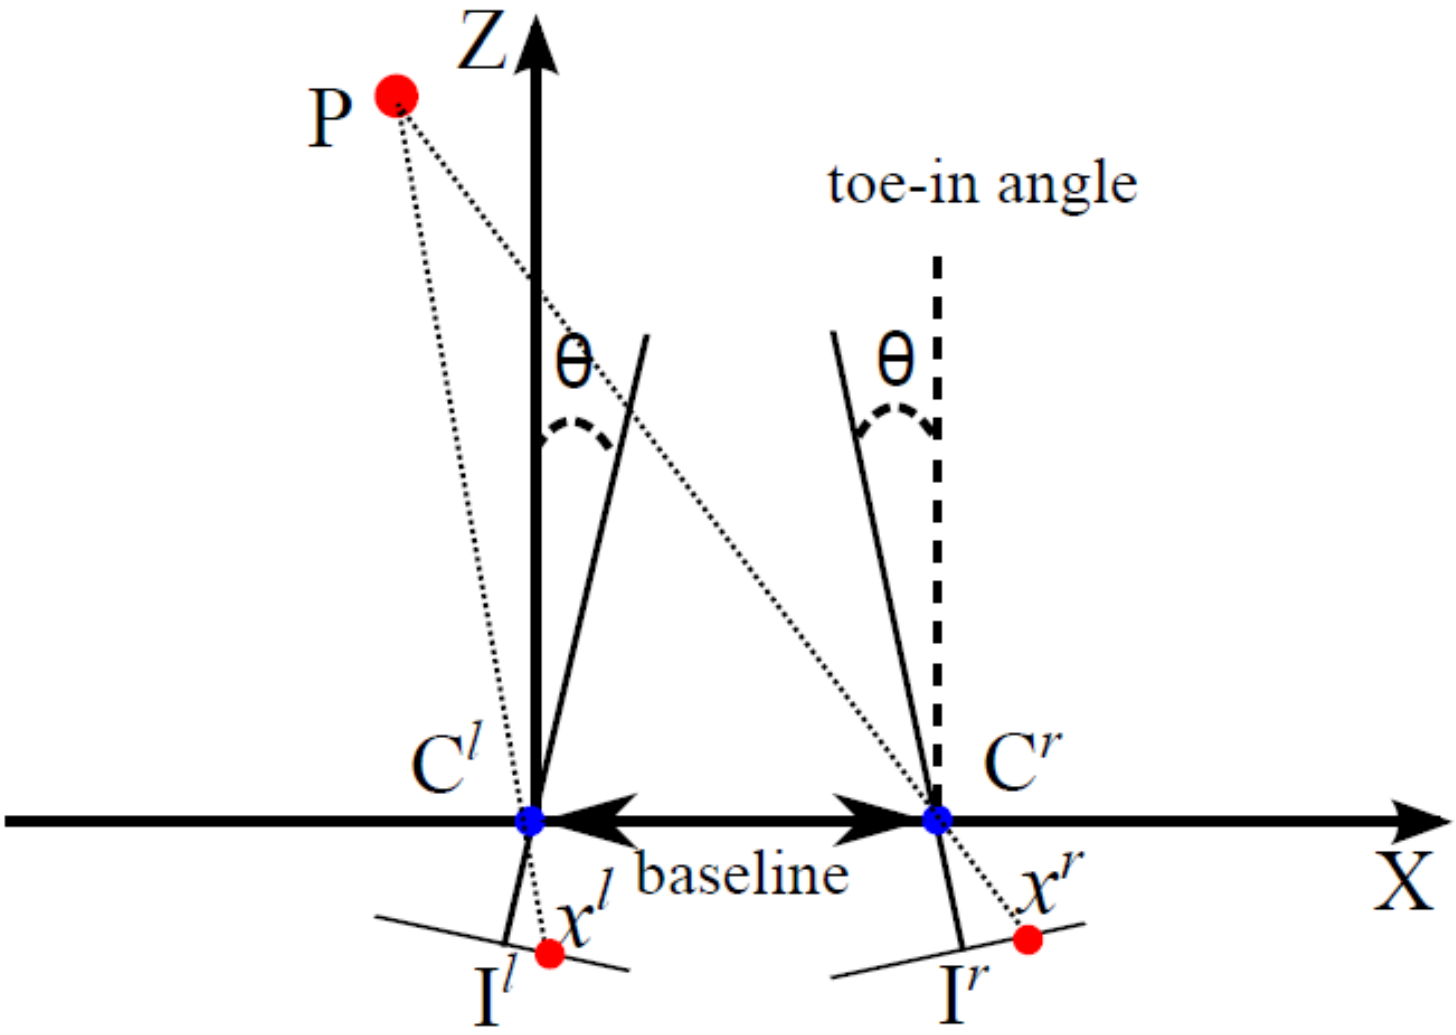

Keystone distortion

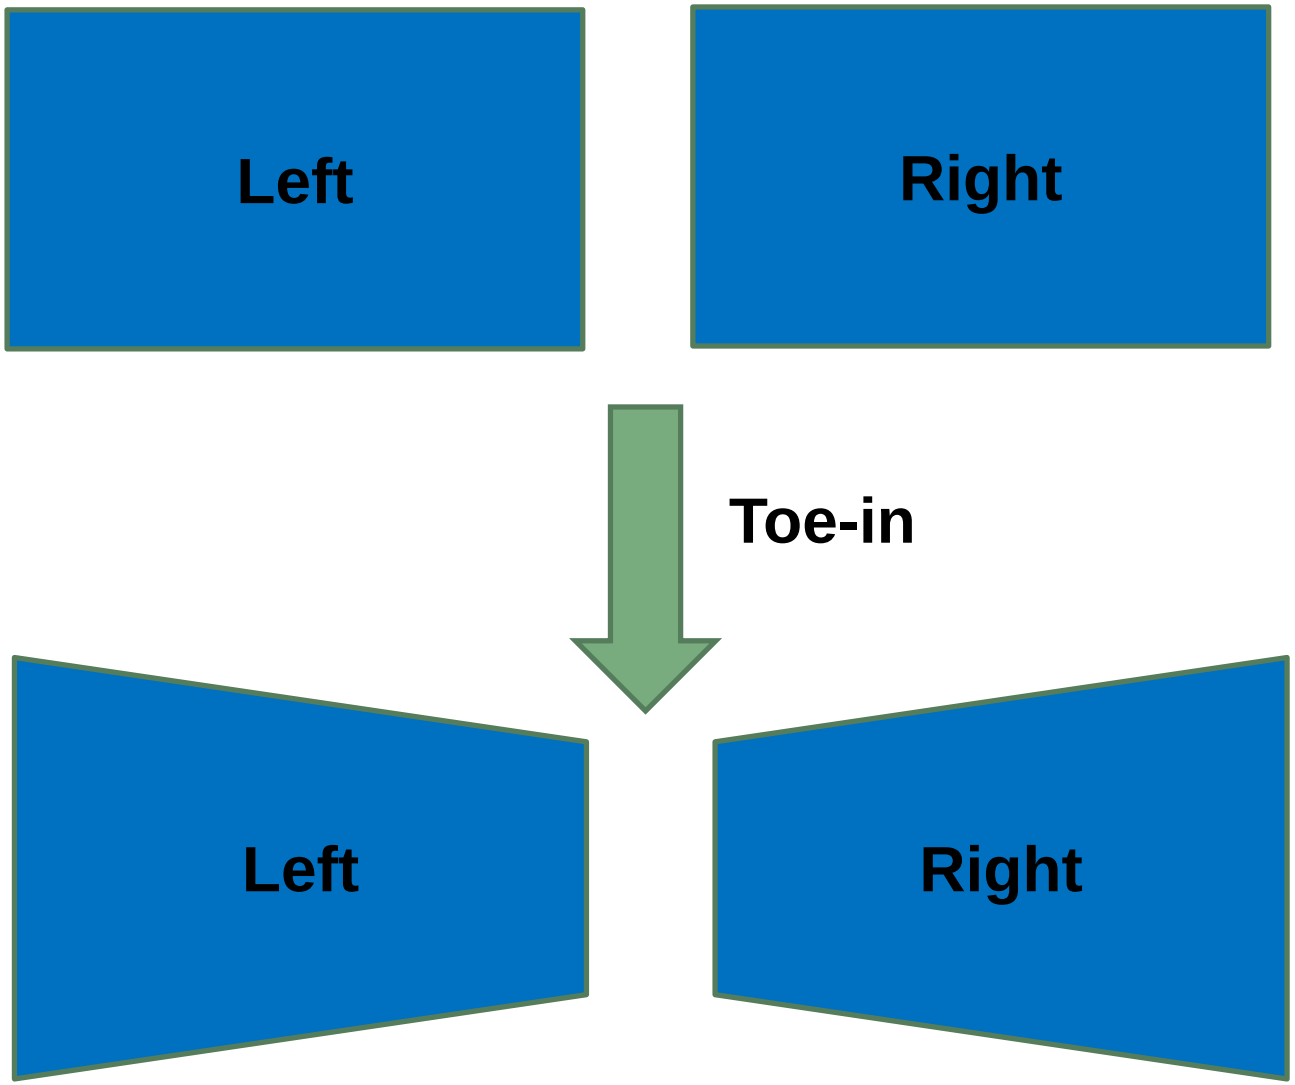

Keystone distortion

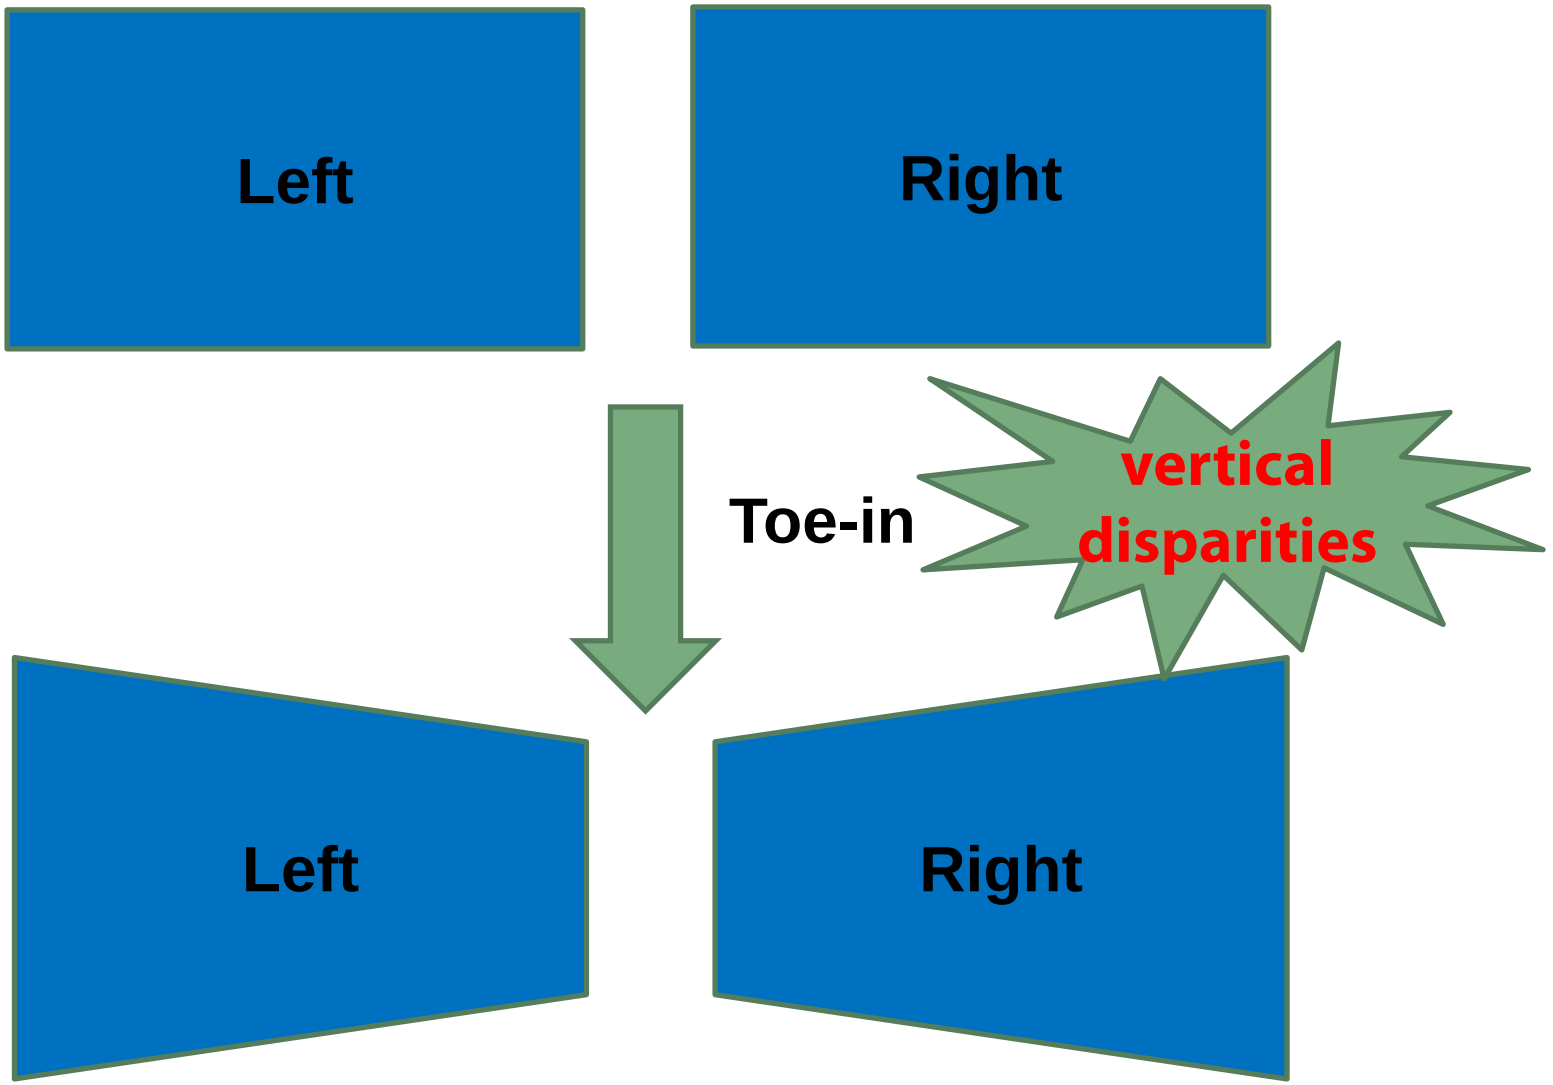

Keystone in projectors

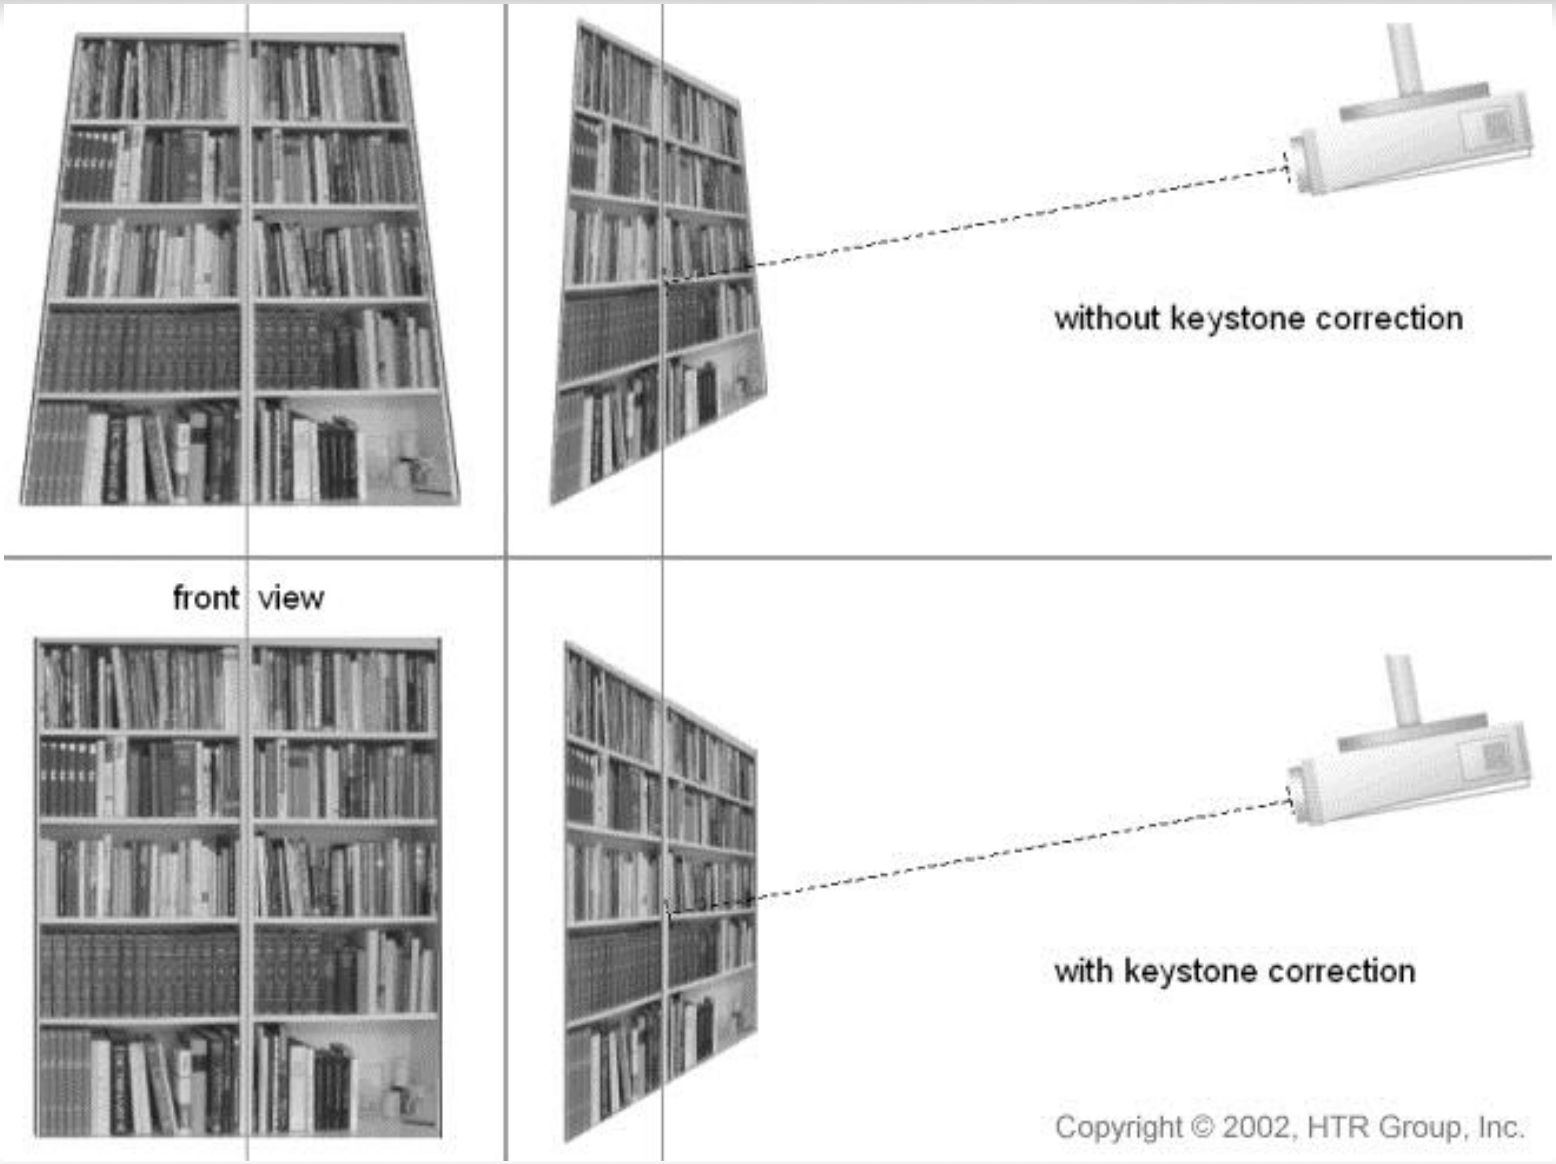

Images from http://htrgroup.com/

Keystone correction for projectors

- Basics: 3D rotation can be modeled by a homography
- Keystone correction [Raskar and Beardsley 01, Li et al. 04, etc]
  - Estimate 3D rotation or homography
  - optical keystone correction by modifying the lens system
  - or digital keystone correction by image warping

Stereo keystone correction

- Projector keystone correction cannot work
  - Revert the toe-in operation
  - Change the desirable (horizontal) disparity distribution
- Stereo keystone correction requires
  - Eliminate vertical disparities
  - Preserve horizontal disparities

Content-preserving warping

- Non-uniformly move image content to target positions
- Avoid noticeable distortion
- Applications:
  - Video stabilization [Liu et al. ‘09]
  - Disparity editing [Lang et al. ‘10]

Correction by content-preserving warping

- Use a spatially-varying warping method
  - Non-uniformly move image content to remove vertical disparities and preserve horizontal disparities
  - Avoid noticeable image distortion

Stereo keystone correction

- Feature correspondence estimation
- Target feature position estimation
- Image transformation via content-preserving warping

Feature correspondence estimation

Input: left image with disparity and right image

- Detect SIFT features from the left and right image
- Establish feature correspondence [Lowe ‘04]
- Remove outliers using the epipolar geometry constraint [Hartley and Zisserman ‘00]

Target feature position estimation

- Keep the input horizontal coordinates to
  - preserve horizontal disparities
- Average the left and right vertical coordinates for each feature pair to
  - remove vertical disparities

Content-preserving warping

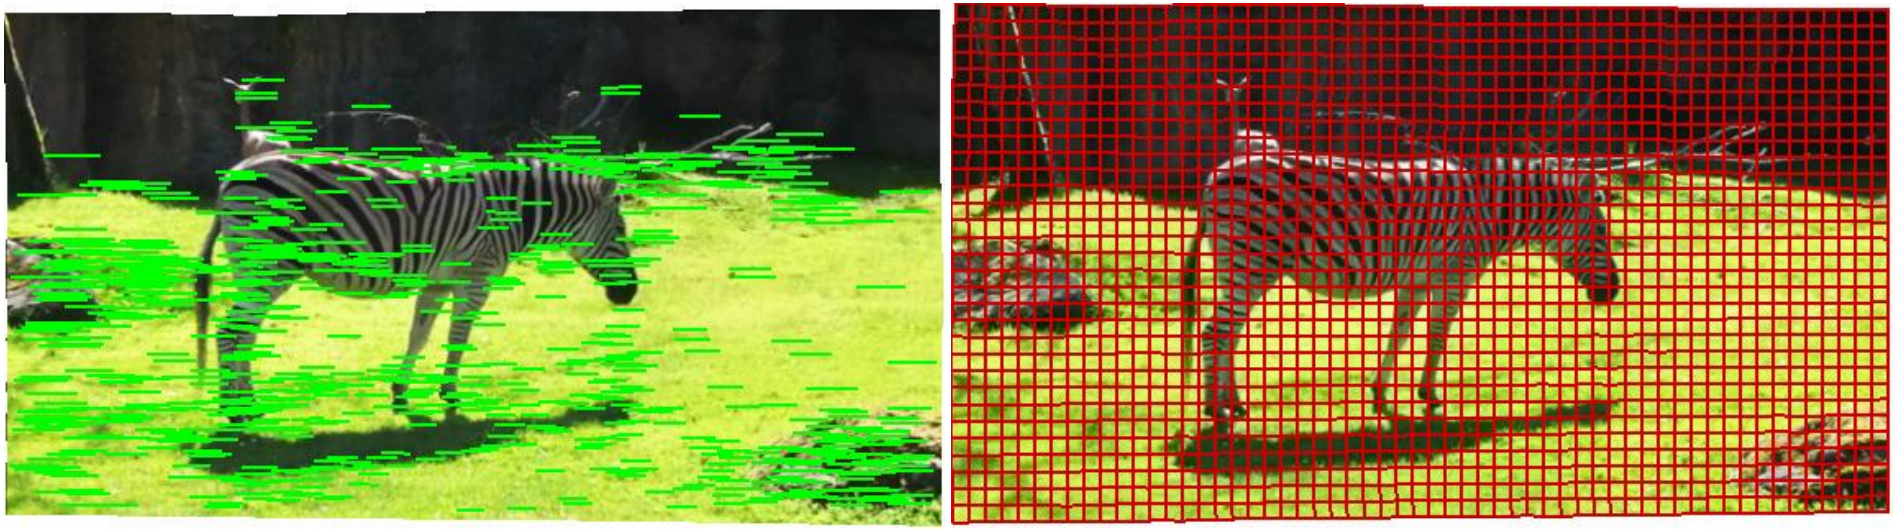

Keystone correction result: left with disparity and right with mesh

Warping algorithm

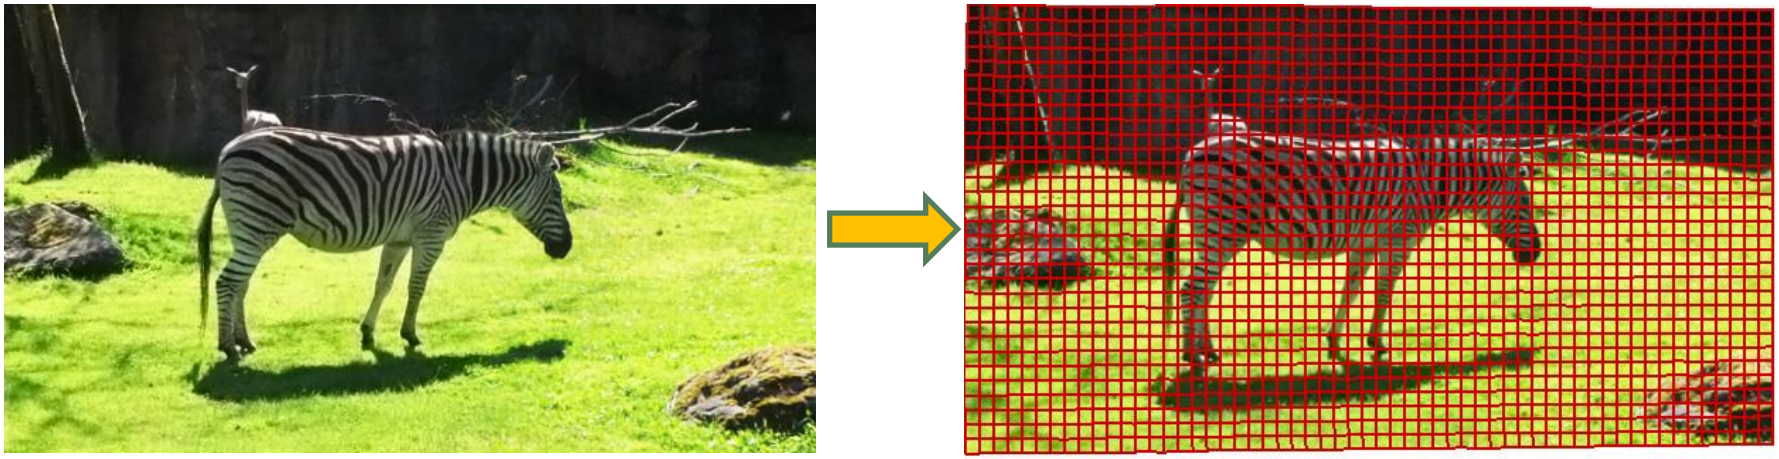

- Build a grid mesh from input image
- Warp input image by least-squares minimization
  - Data term: move features to target positions
  - Smoothness term: avoid visual distortion

Smoothness term: minimize visual distortion

Local similarity transformation constraint

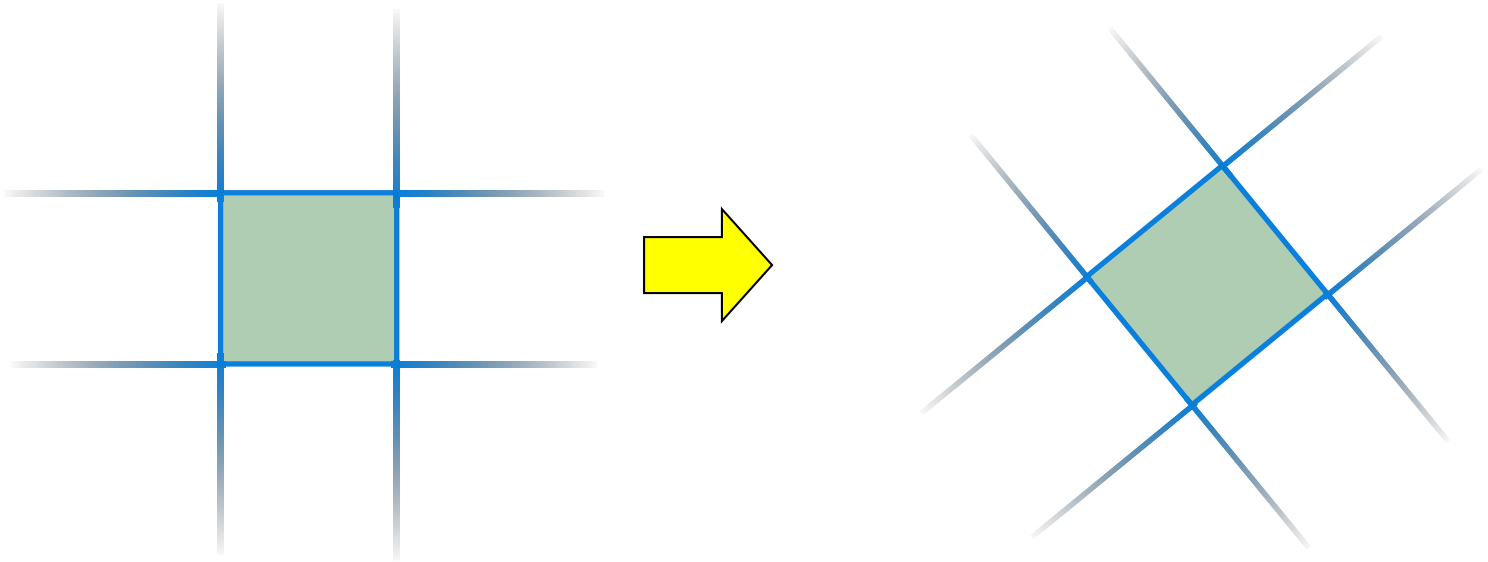

Warping algorithm

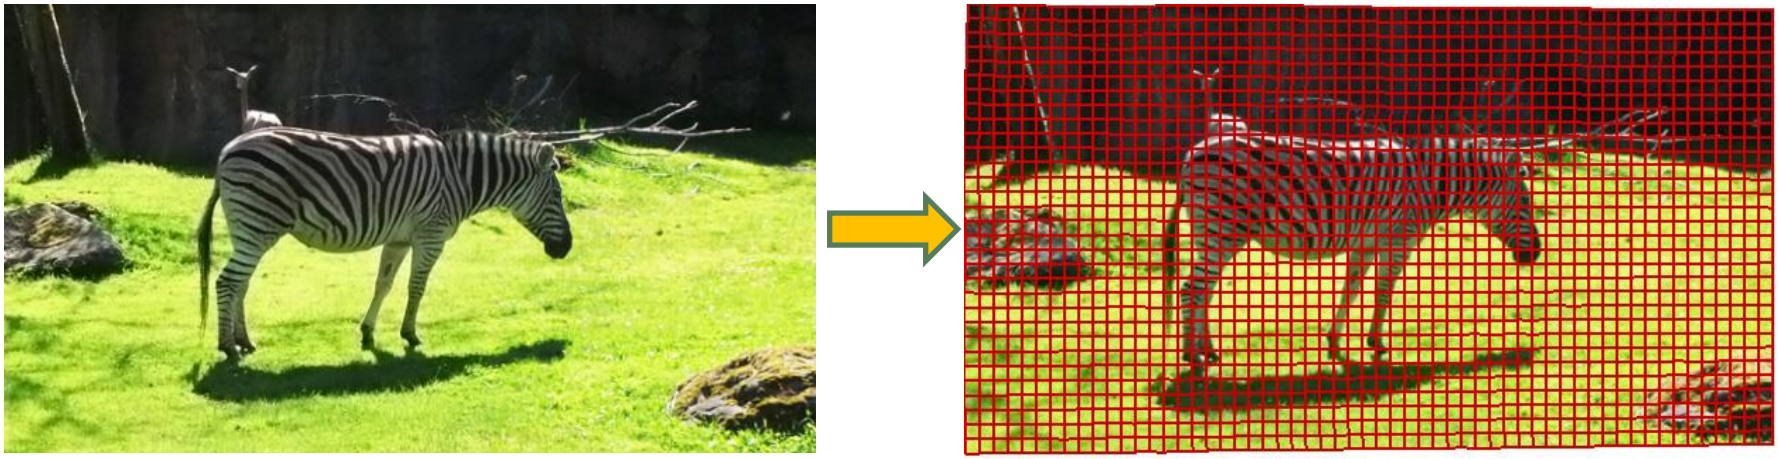

- Build a grid mesh from input image
- Warp input image by least-squares minimization
  - Data term: move features to target positions
  - Smoothness term: avoid visual distortion
  - Solved by a linear solver

Camera-centric disparity editing

- Estimate the relative camera pose between the left and right camera and a sparse set of 3D points
  - 6-point algorithm [Stewenius et al. '05]
- Adjust the baseline and toe-in angle
  - Compute output feature positions
- Content-preserving warping

Disparity adjustment

Input

Disparity adjustment



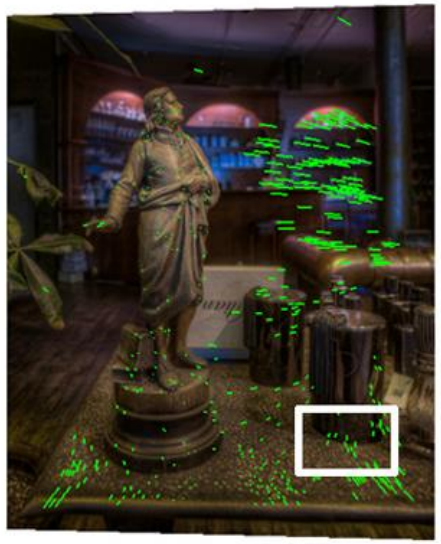
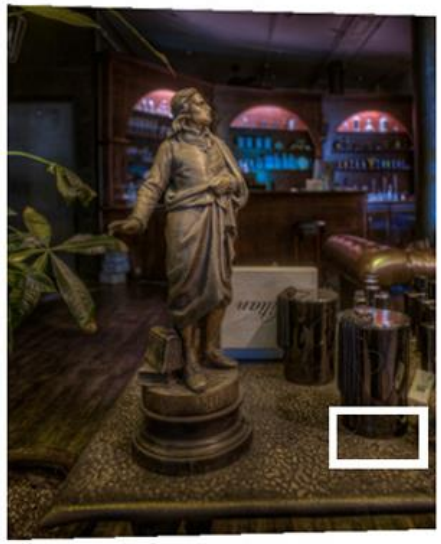

Input

Disparity adjustment

Input

Vertical disparity from 3D rotation

Disparity adjustment

Input

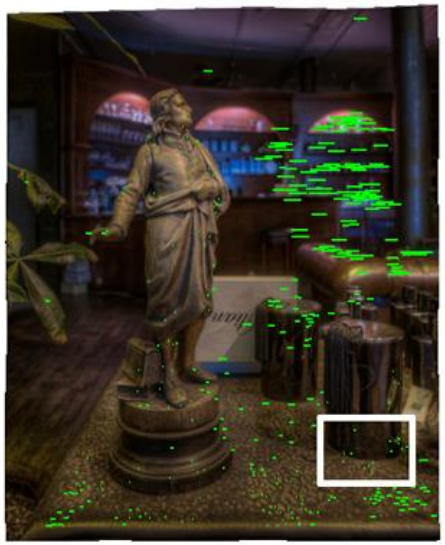

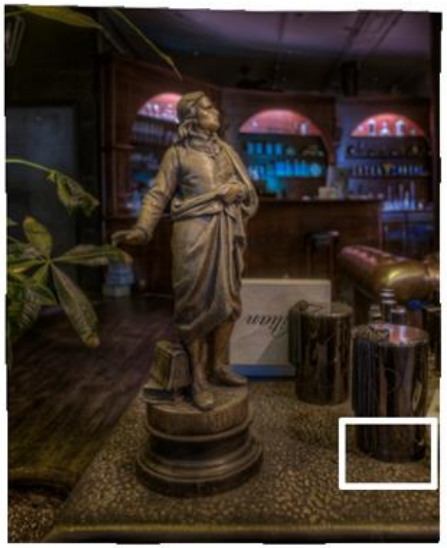

Our result

Disparity adjustment

Examples

Input anaglyph and disparity

Examples: Move the train near the screen

Toe-in result

Examples: Move the train near the screen

Output anaglyph and disparity

Examples: Move the walker near the screen

Examples

Input

Output 1 and 2

Video example


Input sequence


Output sequence

Input

Result

Next Time

- Final Project Presnetation