Fixed-Viewpoint Volumetric Display

Kurt Akeley

SID Bay Area Chapter Meeting

23 March 2010

Image blur affects perceived scale

Original depth of field

The problem with stereo displays

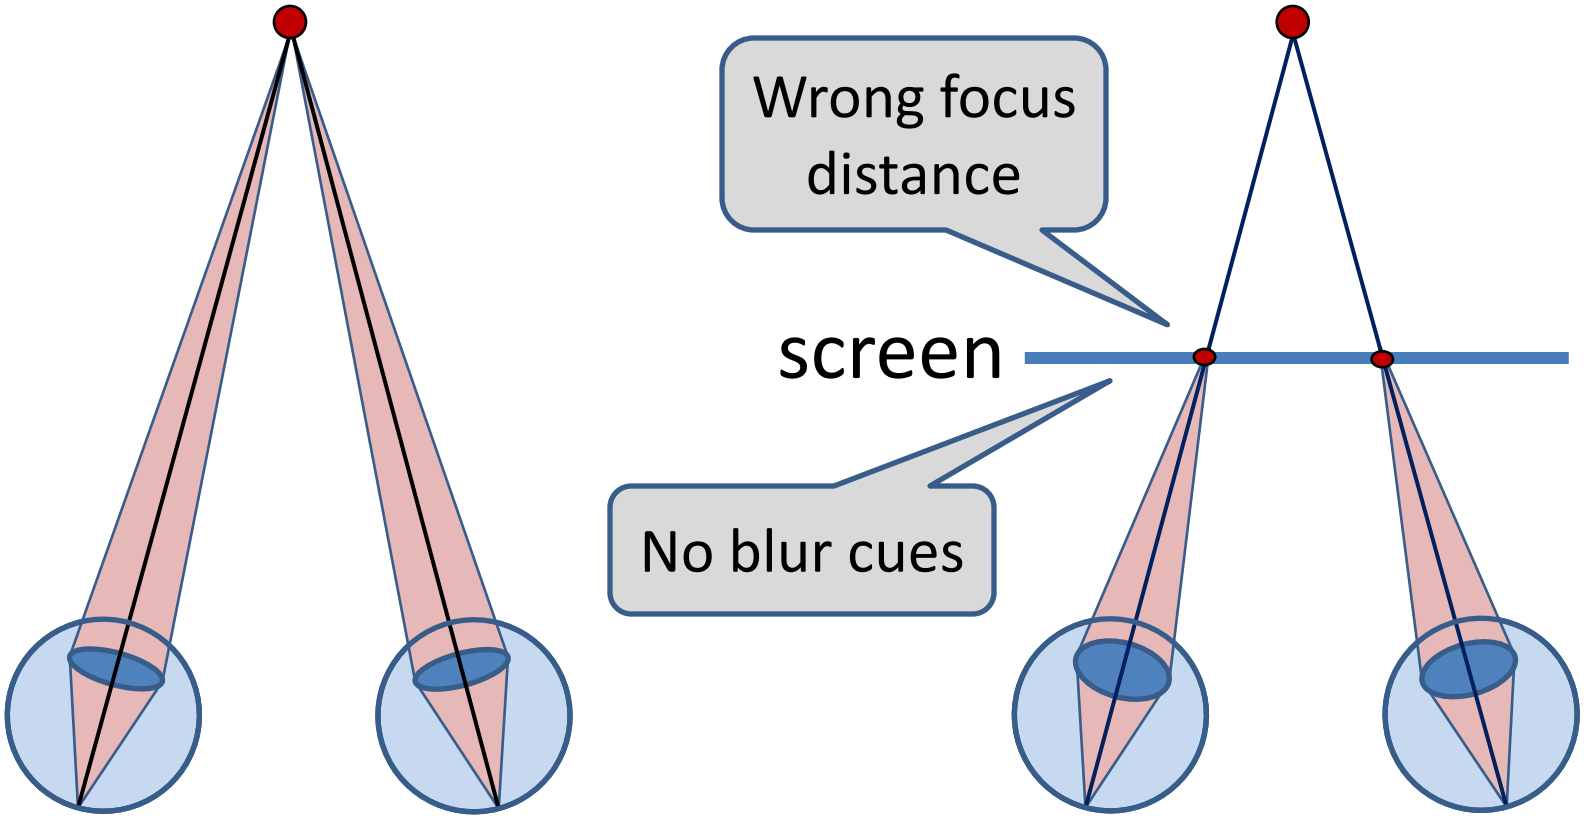

Real world

Stereo display

Our goals

- Build displays that solve this problem:
  - Are geometrically correct
  - Stimulate ~correct focus distance (accommodation)
  - Stimulate ~correct retinal blur
- Use these displays to do science
  - And encourage other researchers to do so too
- Use the science to inform stereo practitioners
  - They already do some clever things
- *Eventually* use the technology to develop better professional and consumer displays

Fatigue due to decoupling

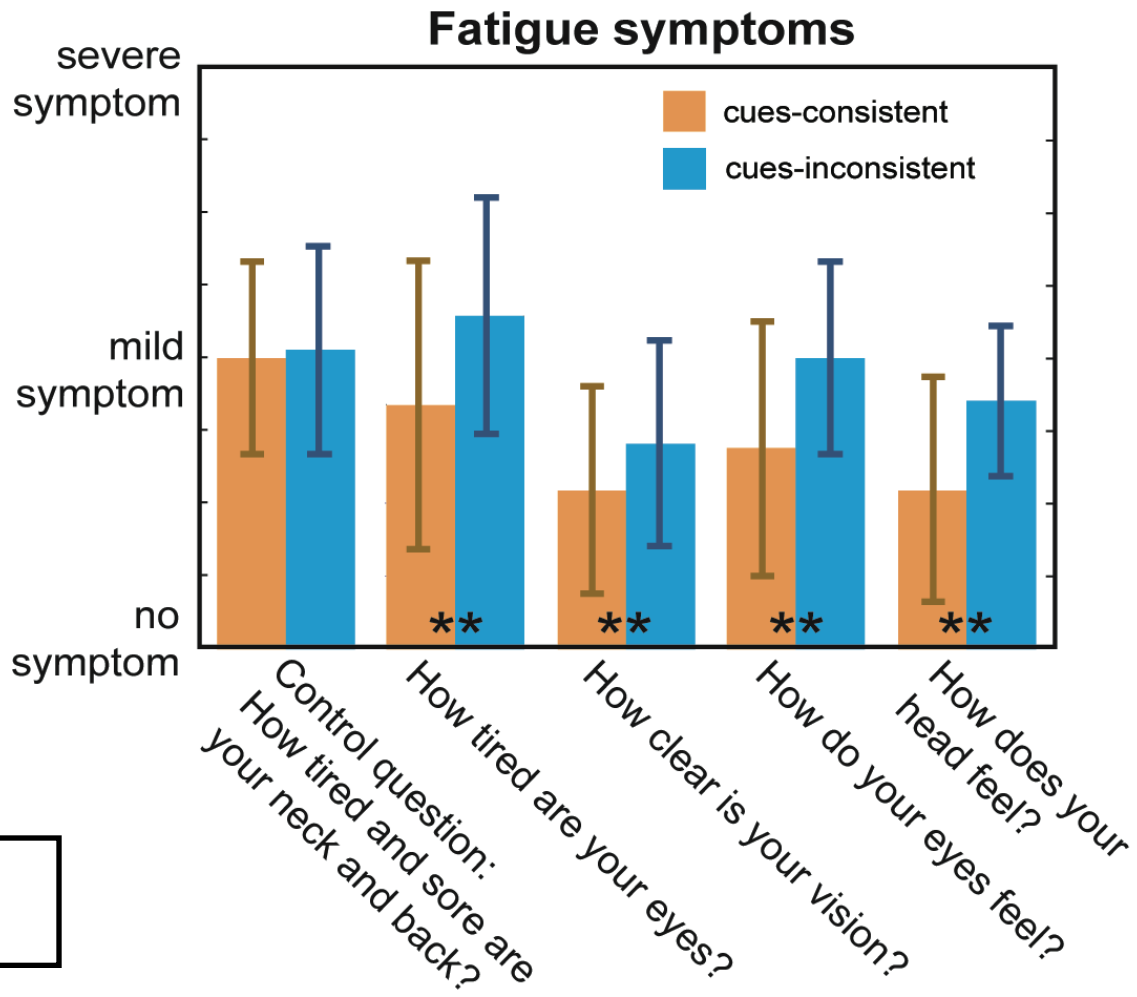

11 subjects
17 sets of trials

David M. Hoffman, Ahna R. Girshick, Kurt Akeley, and Martin S. Banks, Vergence-accommodation conflicts hinder visual performance and cause visual fatigue, in *Journal of Vision*, vol. 8, no. 3, pp. 1-30, March 2008.

Outline

- Volumetric display
- Fixed-viewpoint volumetric display
- *Practical* fixed-viewpoint volumetric display
- Comparison of display approaches

Outline

- Volumetric display
- Fixed-viewpoint volumetric display
- *Practical* fixed-viewpoint volumetric display
- Comparison of display approaches

Volumetric display

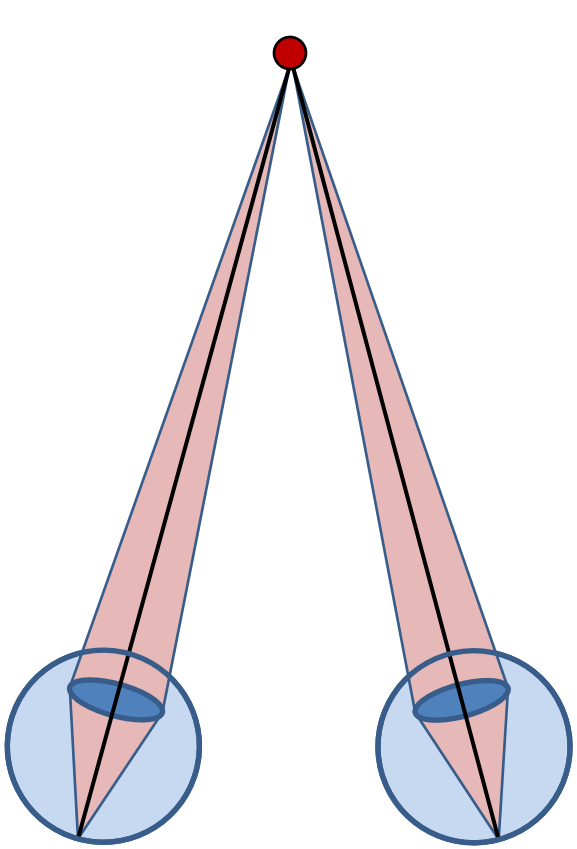

Real world

Additive

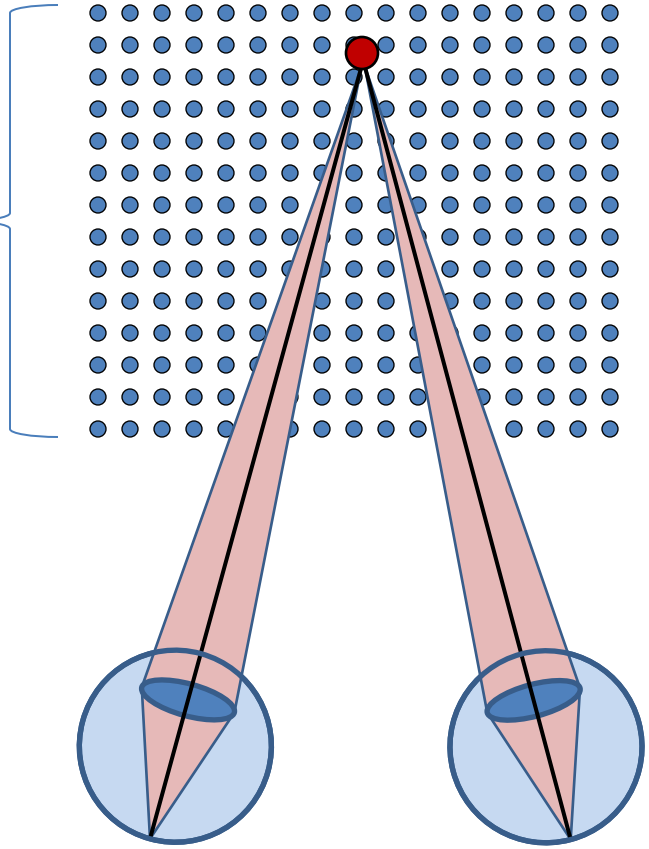

Volumetric display

Actual volumetric displays

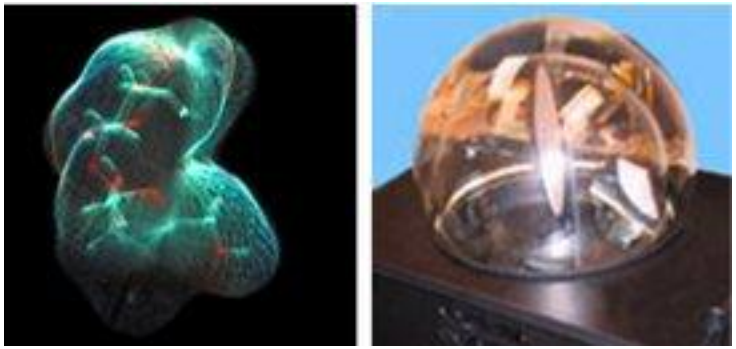

Actuality Systems Perspecta™

Gregg Favalora et al., 100-million-voxel volumetric display, Proceedings of SPIE, vol. 4712, pp. 300-312, 2002.

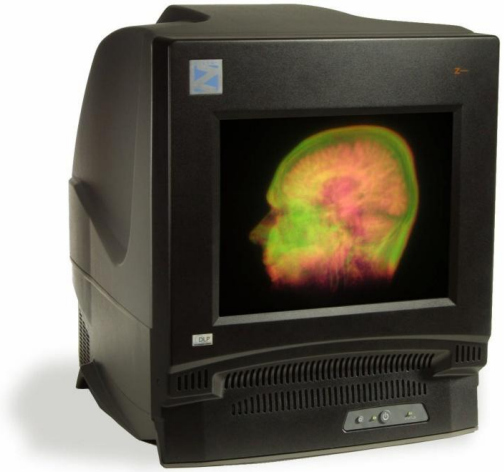

Vista3D DepthCube™

Alan Sullivan et al., A solid-state multi-planar volumetric display, SID 58.3, 2003

My definition of volumetric

- The light field is the sum of a 3-D volume of diffuse pixels (voxels), or it appears to be.
  - sum: voxels are non-occluding
  - 3-D volume: light comes from the correct distance
  - diffuse: light from a voxel is radiated approximately equally in all directions
  - appears: clever optics may be used to create the effect

Rotating parts ≠ volumetric

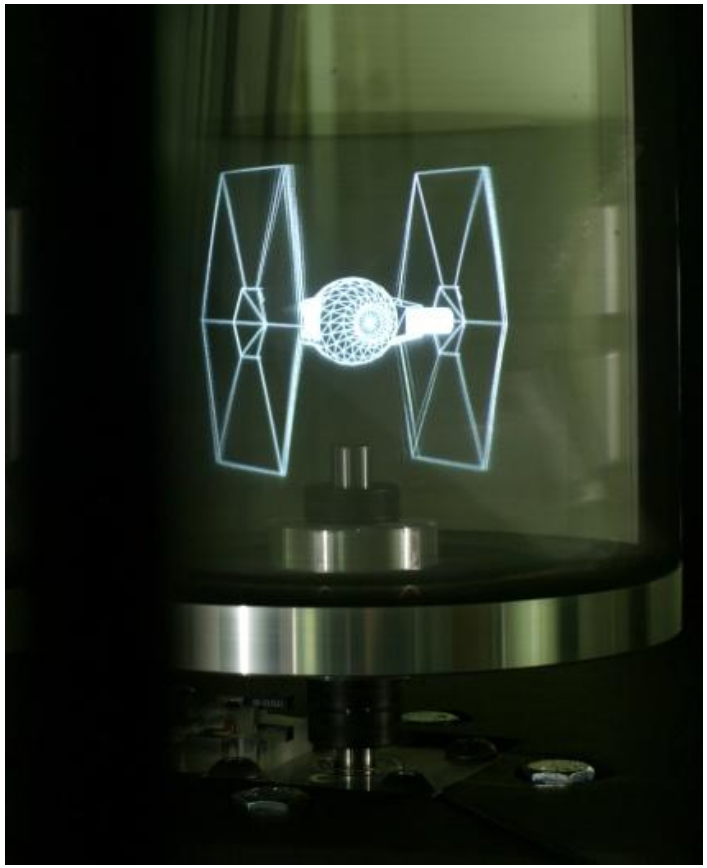

USC/Fakespace Labs/Sony

Jones et al., Rendering for an interactive 360 degree light field display, *SIGGRAPH '07: ACM SIGGRAPH 2007 papers*, 2007.

This is a *light-field* display,
Not a *volumetric* display!

Volumetric display: desirable features

- Auto-multiscopic
  - Auto-stereoscopic
  - Multiple viewers
  - Motion parallax
- Focus cues are nearly correct
- A true 3-D display ☺

Auto-multiscopic

Volumetric display: major failings

- No view-dependent shading
  - No reflections or specularity
  - No hidden-surface elimination
- Too many voxels!
  - O(n^3) for unconstrained viewing
  - 100-million voxel …
- Not immersive
  - Viewers look at it
- Poor quality *volume rendering*
  - Sum along lines of sight is unusual

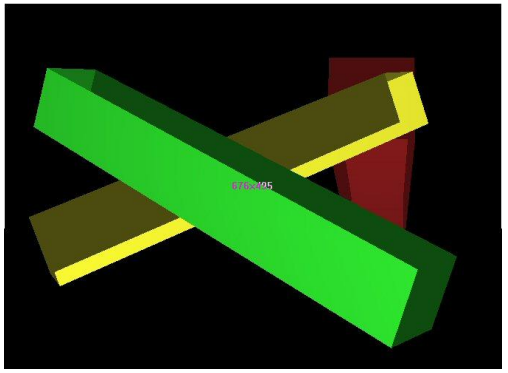

Can we fix the voxels?

- Occluding voxels
  - Correct hiding, but not view-dependent shading
  - I know of no working examples
- Directional (non-diffuse) voxels
  - This corrects visibility and view-dependent shading
  - But it cannot be implemented (O($\boldsymbol{n^5}$) is hopeless)
  - And its five degrees of freedom are overkill
    - Any light field is specified by four degrees, O($\boldsymbol{n^4}$)

Outline

- Volumetric display
- Fixed-viewpoint volumetric display
- *Practical* fixed-viewpoint volumetric display
- Comparison of display approaches

Idea: constrain the viewing position

- Weak constraint (range of viewing positions)
  - Addresses the voxel-count problem
    - Less resolution is required in depth
    - Reduces O(n^3) to O(n^2)

DepthCube™

Idea: constrain the viewing position

- Weak constraint (range of viewing positions)
  - Addresses the voxel-count problem
    - Less resolution is required in depth
    - Reduces O(n^3) to O(n^2)

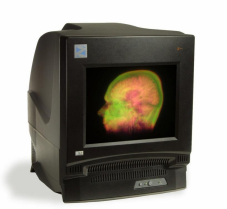

DepthCube™

- Strong constraint (rigidly fixed viewing position)
  - *S*olves the shading problem
    - Can render for a single viewpoint
    - There are minor errors—we'll come back to them later …
  - Solves the immersion problem
    - Head-mount the display
    - Use optical cleverness to implement long focus distances

Diopters

$$\text{diopters} = \frac{1}{\text{meters}}$$

My corrected range (1 D)

My accommodation range (1 D)

Young adult's accommodation range (8 D)

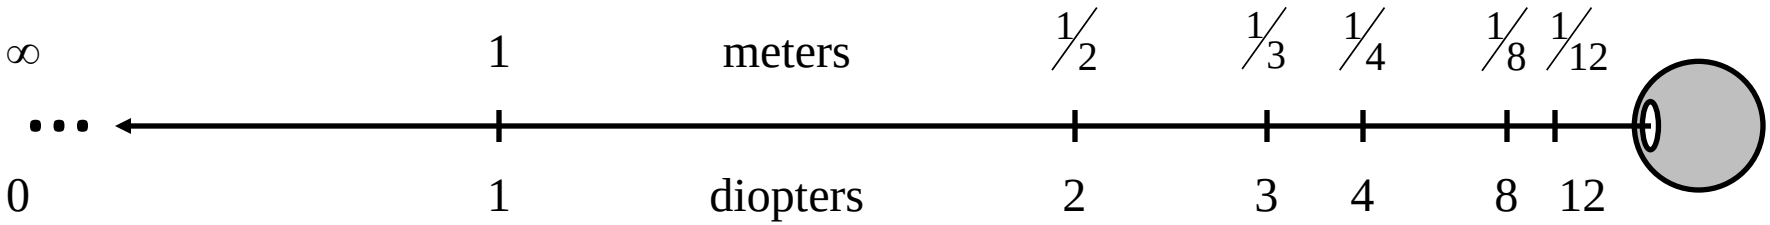

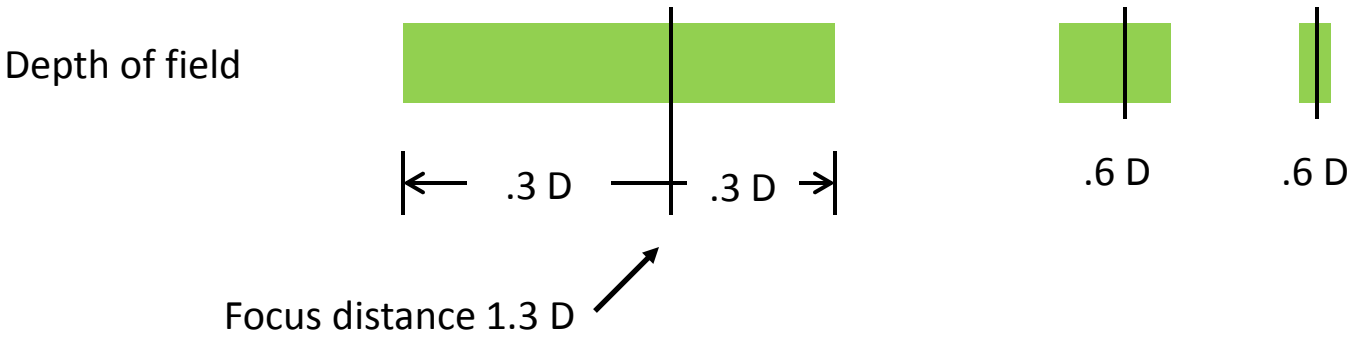

How much depth resolution?

- Human DOF is approximately 1/3 D
- Human accommodation range is 8 D
- Suggests resolution of 24 in depth
- Required spatial resolution is in the thousands
  - E.g. 5400 for 60 pix/deg and 90 deg field of view
- Hence O($\boldsymbol{n^2}$)

Display schematic

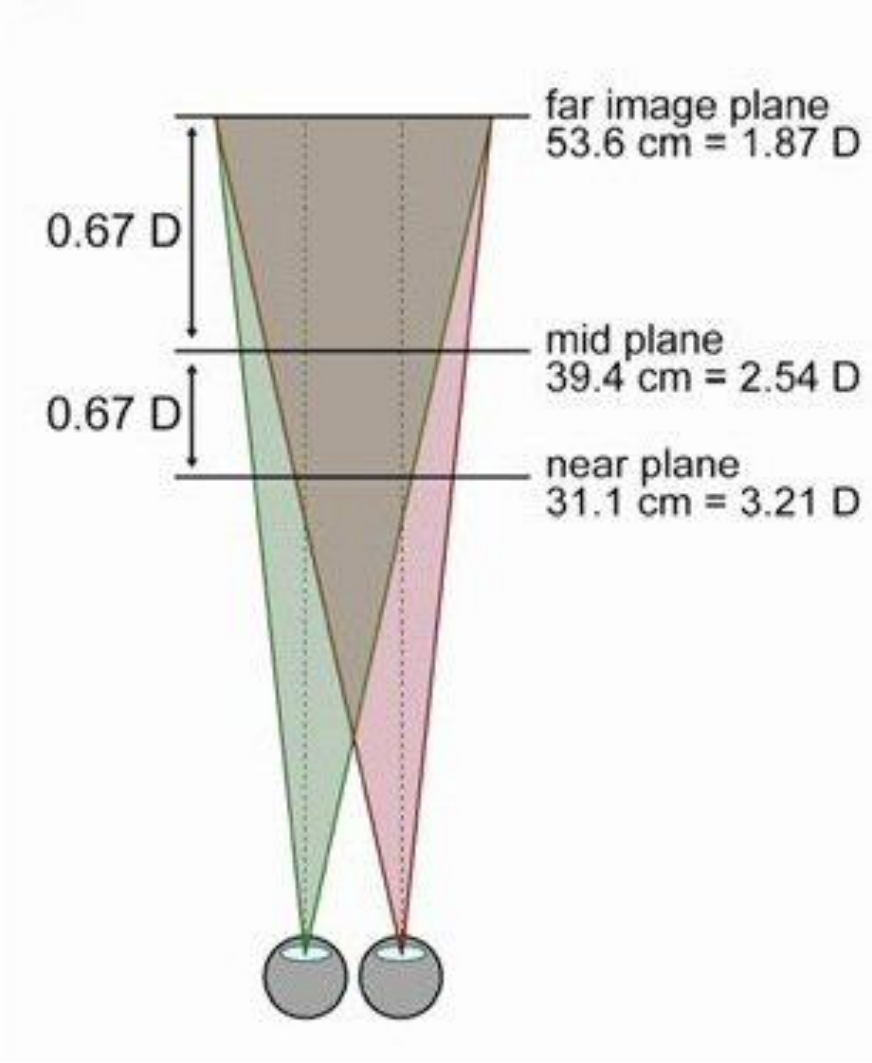

Simple (laboratory) fixed-viewpoint volumetric display

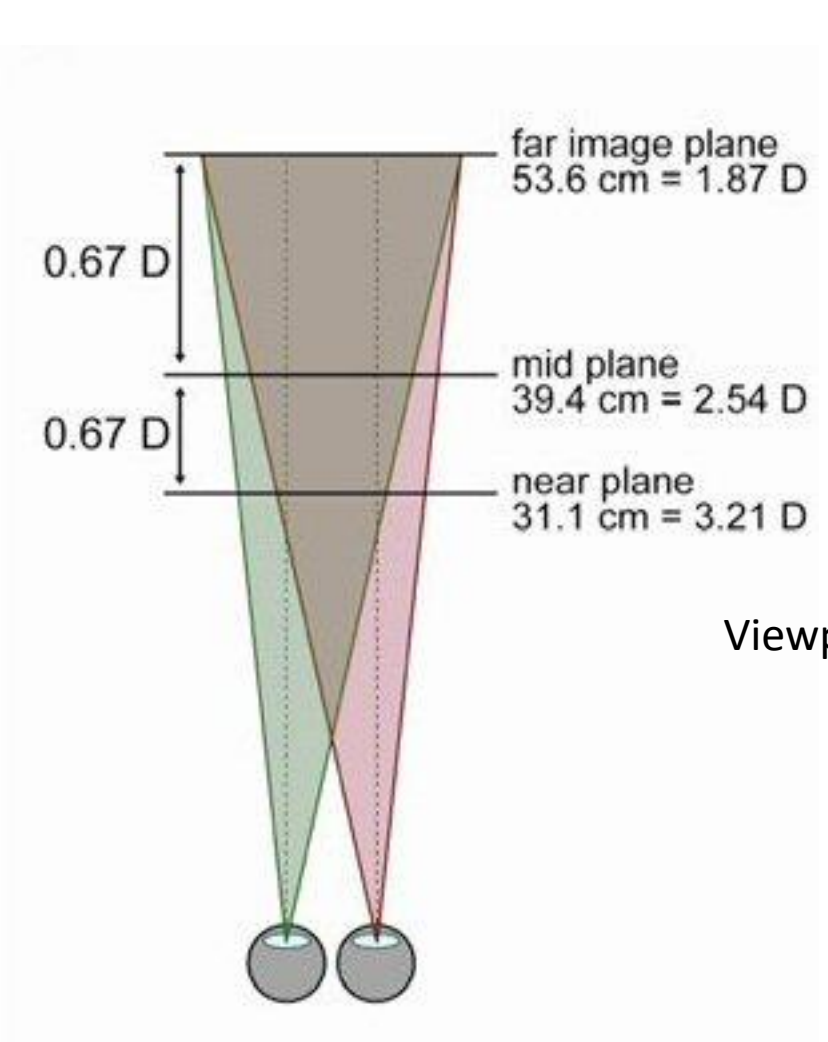

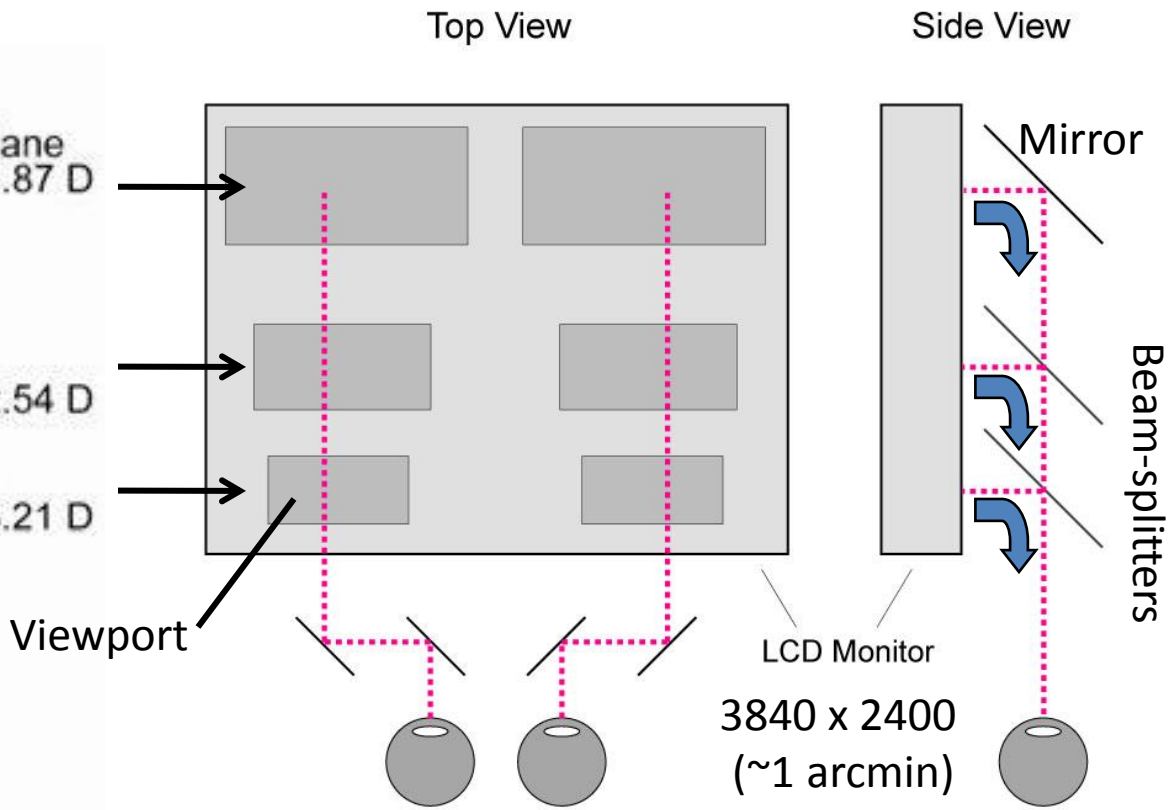

Simple (laboratory) fixed-viewpoint volumetric display

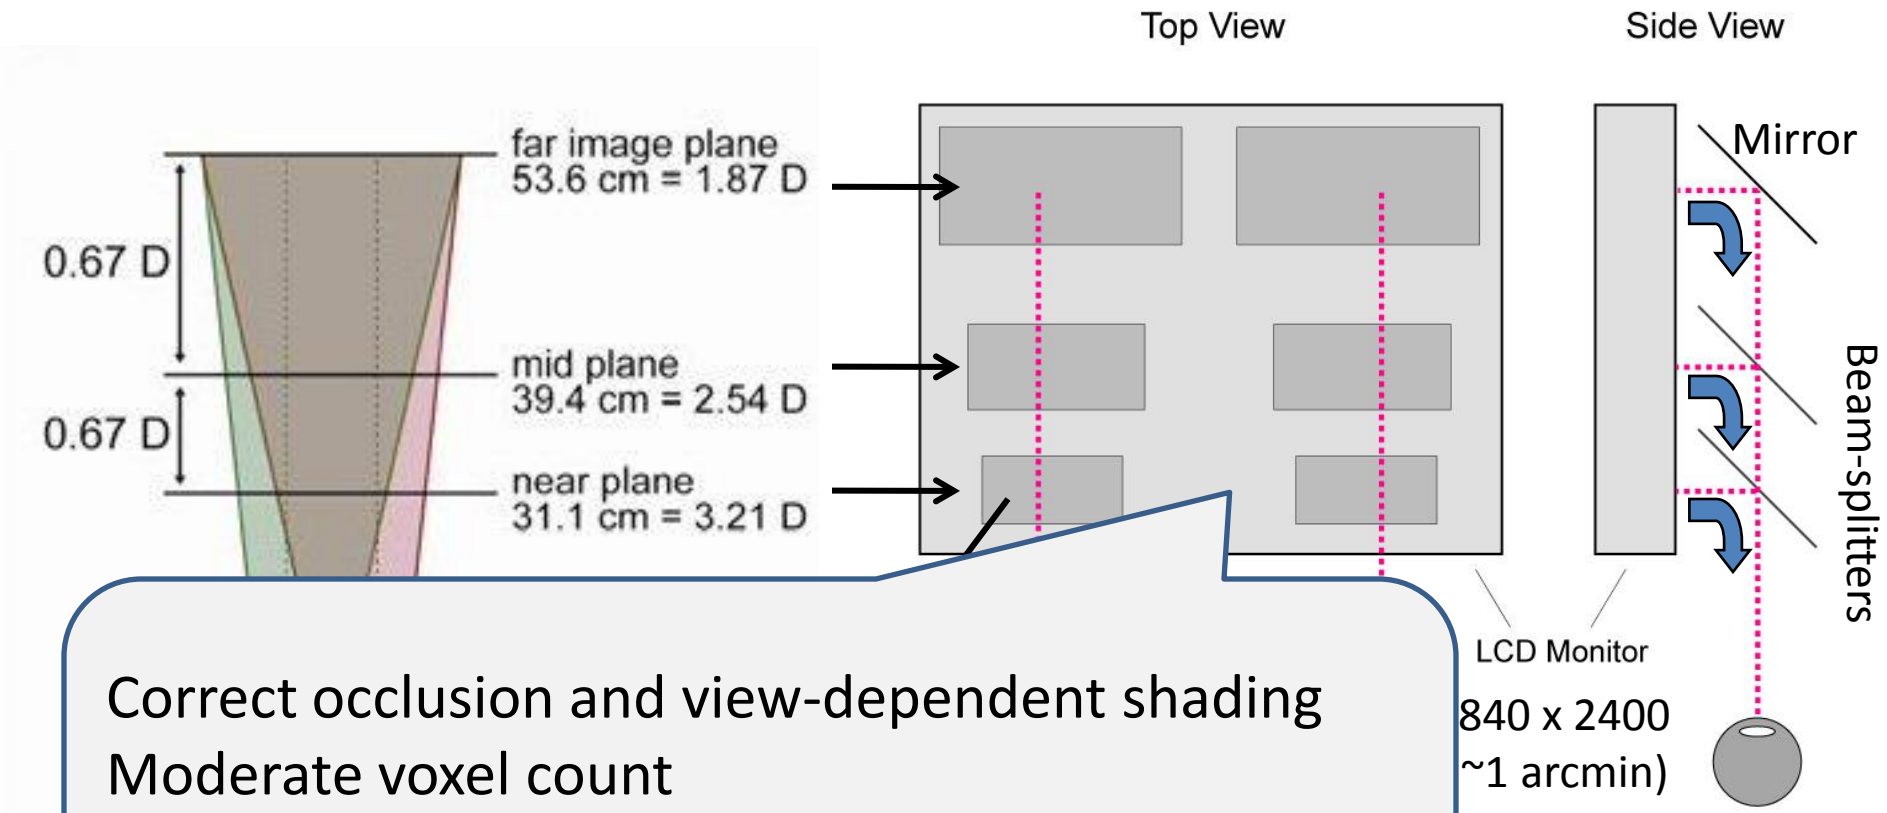

Correct occlusion and view-dependent shading
Moderate voxel count
Limited depth range and FOV (OK for science use)
Not immersive (OK for science use)

Demo

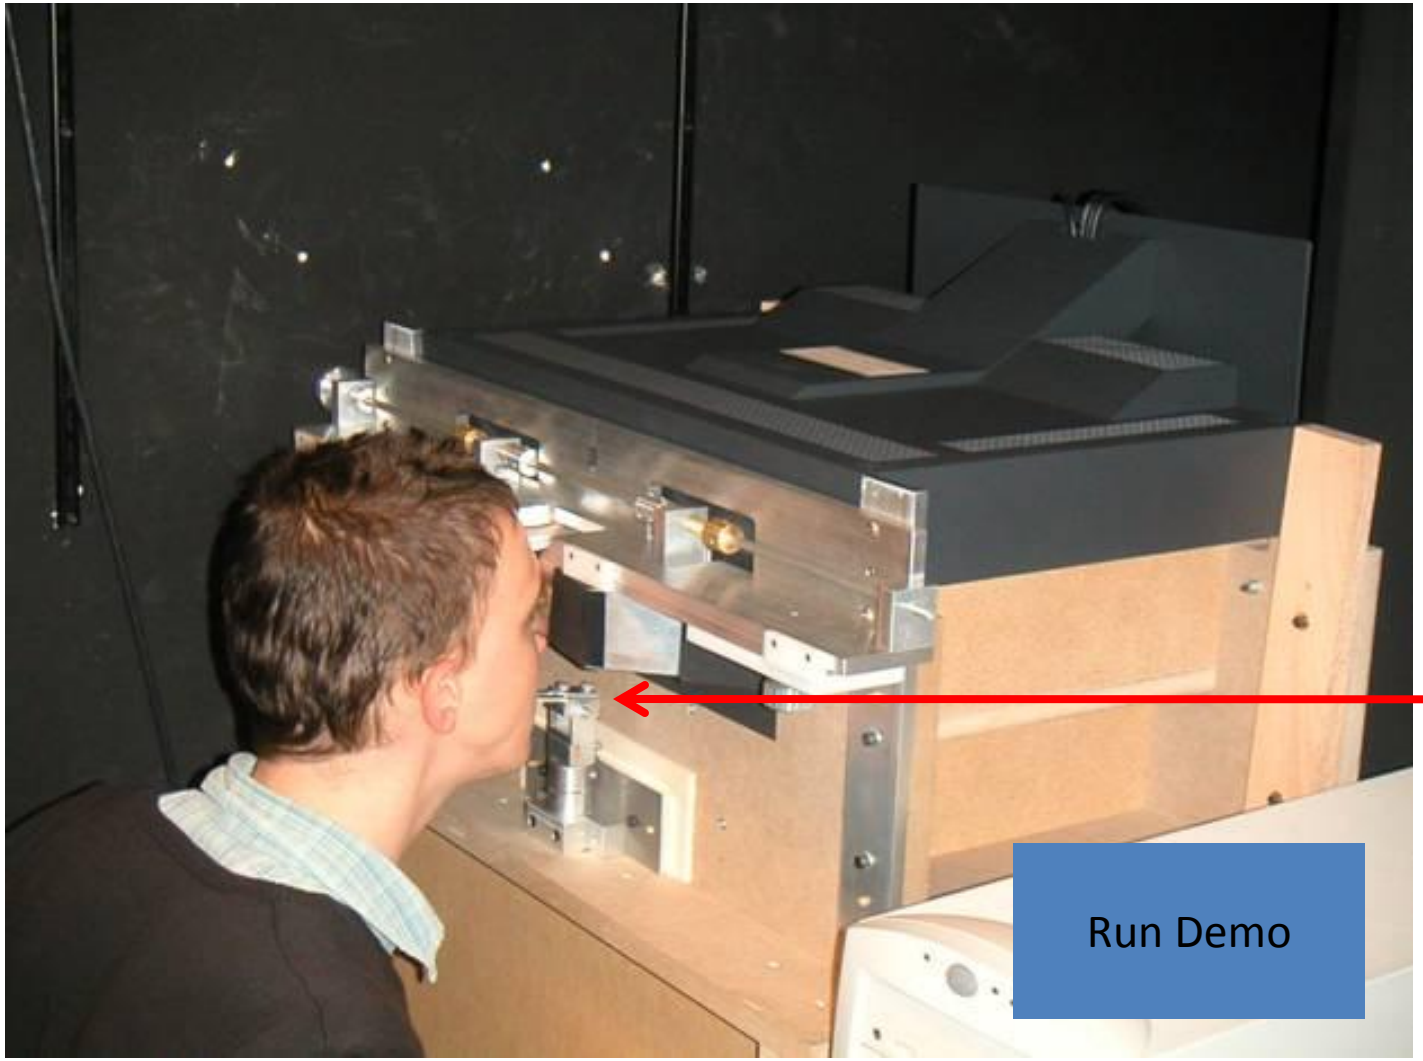

Depth filtering

(aka line-of-sight filtering)

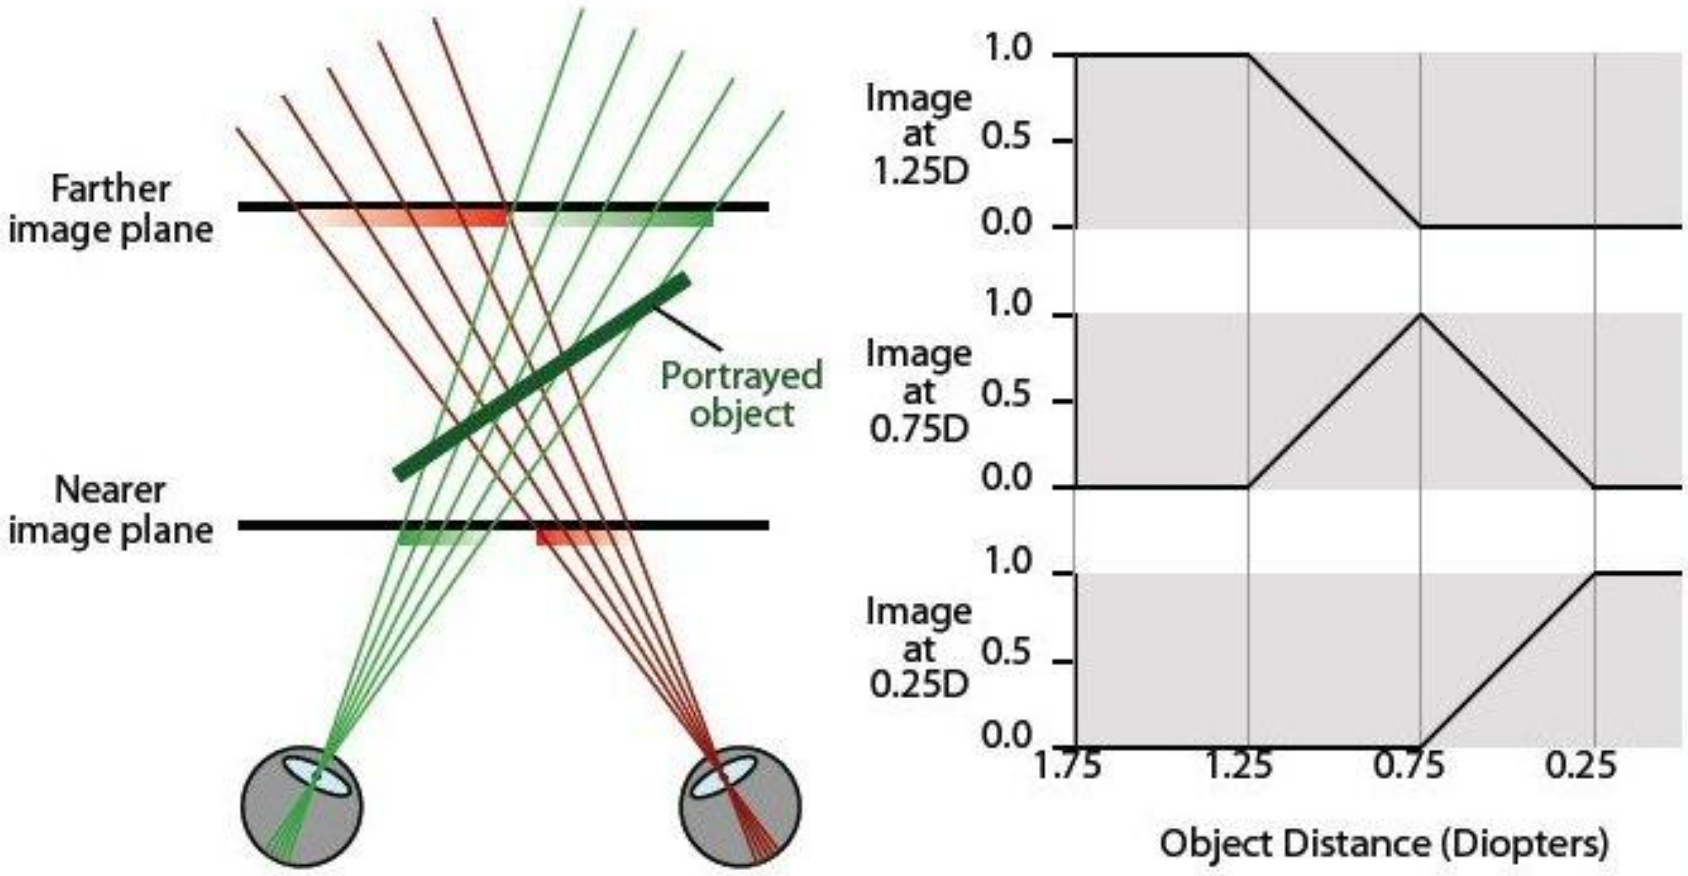

Kevin J. MacKenzie and Simon J. Watt, A display with multiple focal planes can stimulate continuous variations in accommodation, *Vision Sciences Society Annual Meeting*, 2009.

Depth filtering avoids discontinuities

- Must filter to avoid visible changes
  - Obvious to people familiar with computer graphics
  - Easy to demonstrate (demo after presentation)
  - Easy to show in theory
    - Viewer can't focus on both depths simultaneously
    - Different blurs sum to a visible discontinuity

No depth filter → visible discontinuity

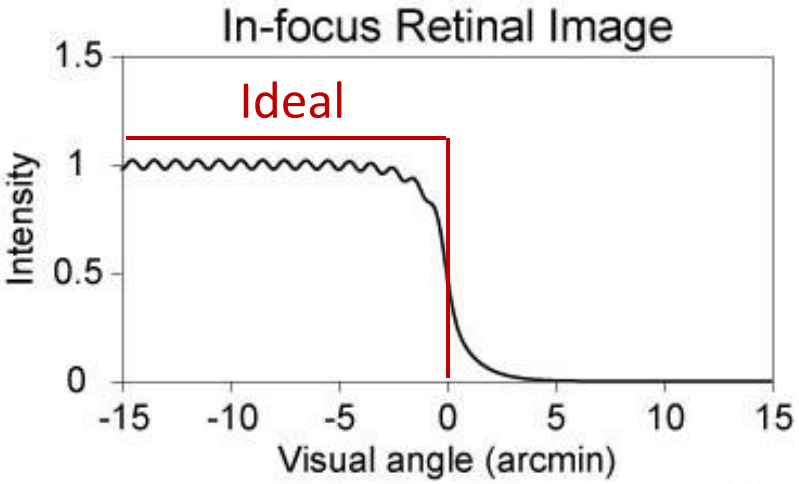

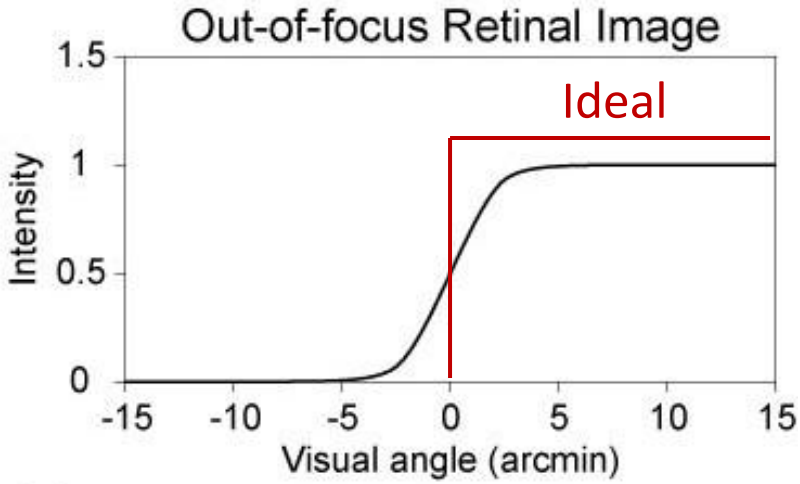

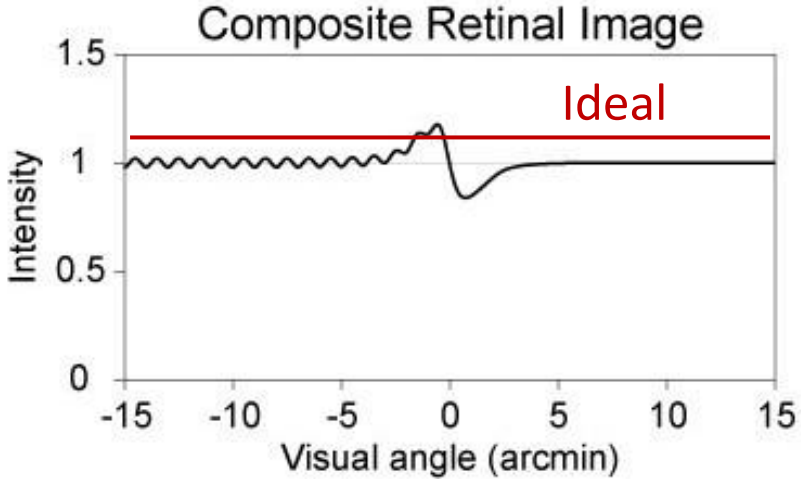

Depth filtering guides accommodation

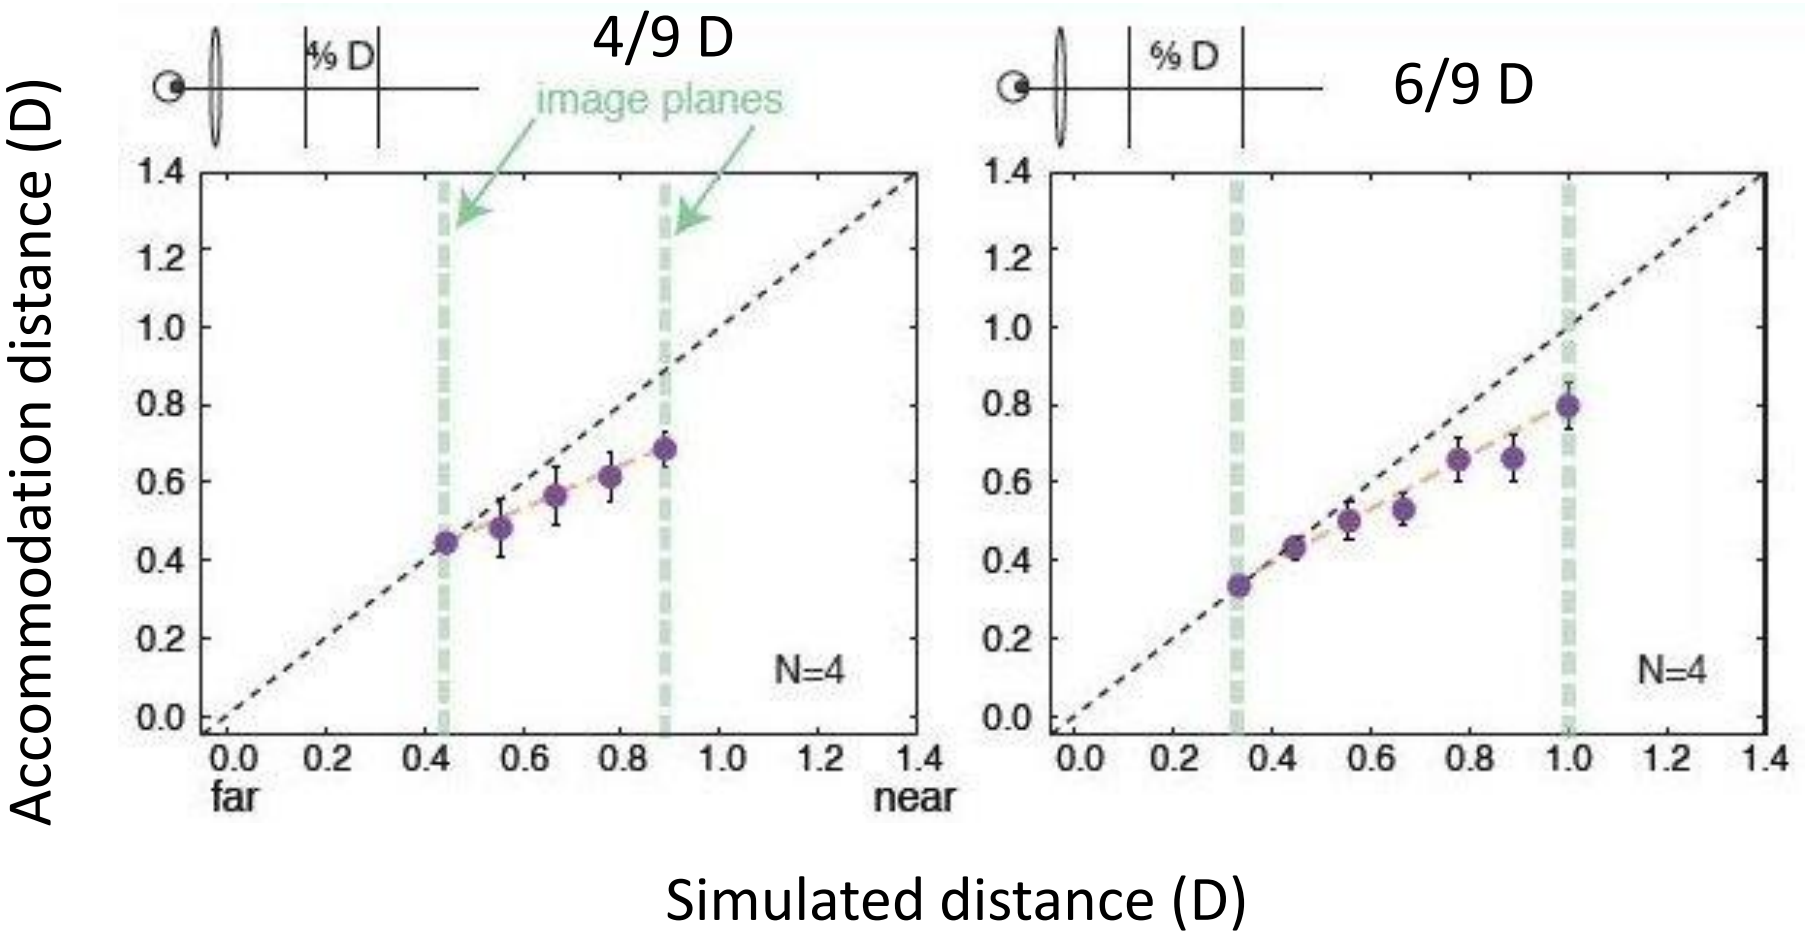

Kevin J. MacKenzie and Simon J. Watt, A display with multiple focal planes can stimulate continuous variations in accommodation, *Vision Sciences Society Annual Meeting*, 2009.

Guidance degrades for large separations

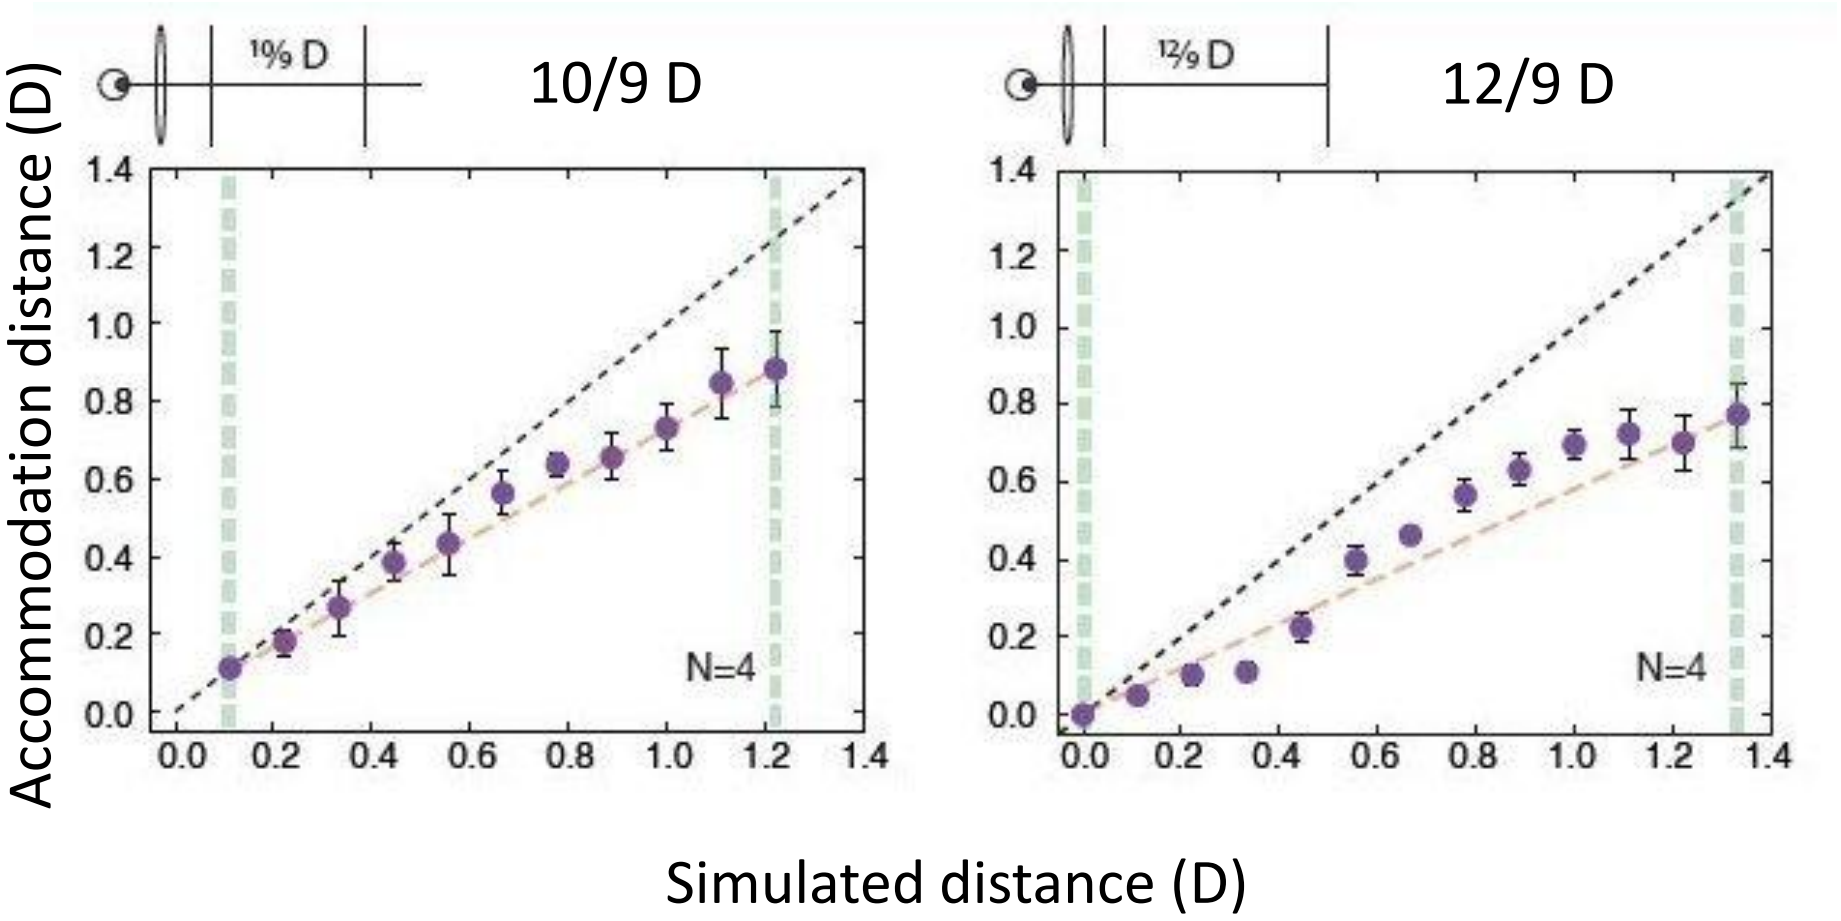

Kevin J. MacKenzie and Simon J. Watt, A display with multiple focal planes can stimulate continuous variations in accommodation, *Vision Sciences Society Annual Meeting*, 2009.

Fixed-viewpoint volumetric pros

- All cues correct (to an engineering tolerance)
- Moderate voxel count (5-10 image planes)
- Possibility of immersion (we'll see how soon)
- Transparency (defocus laser scan can't do this)
- No eye tracking required (position is fixed)
  - Optical center is ~5 mm ahead of rotational center
  - Calibrate to rotational center, not optical center
    - Correct alignment in fixation direction
    - Increasing error in periphery, but this isn't noticed

Fixed-viewpoint volumetric cons

- Head-mounting
  - Great for immersion, but awkward and unpopular
  - Especially poor for tele-communication
  - Impractical without further refinement …
- Volumetric → no light is occluded
  - Visible errors at silhouettes
    - depth-blending argument explains this
  - True (4-D) light-field required to correct this
    - Integral imaging or holography
    - O($\boldsymbol{n^4}$) voxels, so significant engineering challenge

Outline

- Volumetric display
- Fixed-viewpoint volumetric display
- *Practical* fixed-viewpoint volumetric display
- Comparison of display approaches

Practical FVV displays

- Prototype approach, no optics
  - Limited maximum simulated distance
  - Bulky (head mounted to display)
- Add a lens between the eye and the voxels
  - Extends simulated distance to infinity
  - Greatly reduces bulk
  - Can operate with a single display surface
  - Proposed by Rolland et al. in 1999

J. P. Rolland, M. W. Krueger, and A. A. Goon, Dynamic focusing in head-mounted displays, in *Proceedings of SPIE*, vol. 3639, pages 463-470, 1999.

Fixed-power lens

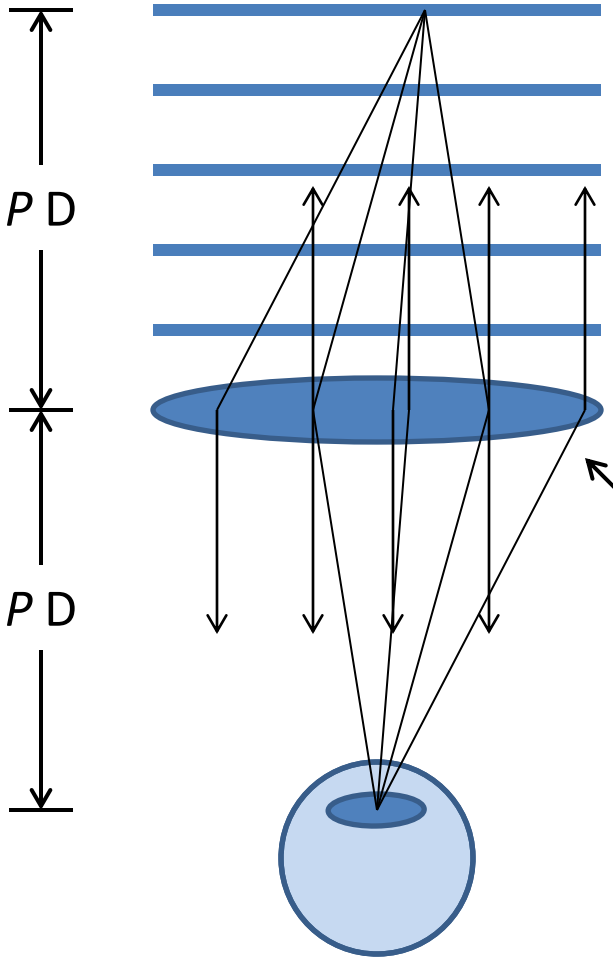

0 D (infinity) focal distance

0.2*P* D focal distance

0.4*P* D focal distance

0.6*P* D focal distance

0.8*P* D focal distance

P D focal distance

Lens power is *P* D

Badal optics

Determines the metric depth of the display

Transparent, emissive image planes (e.g., OLED)

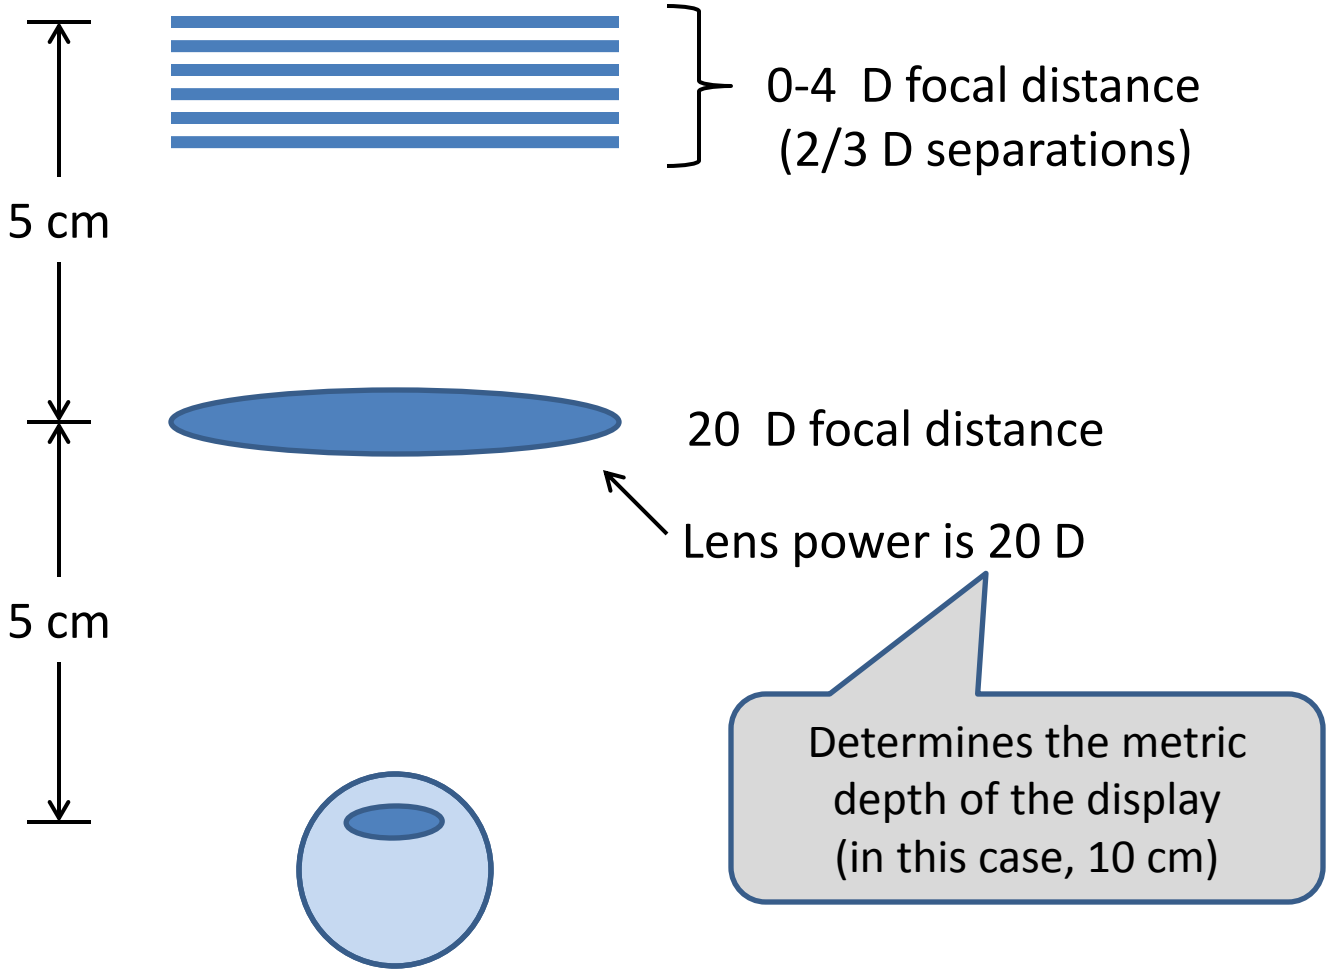

Bangor system (mirrors again)

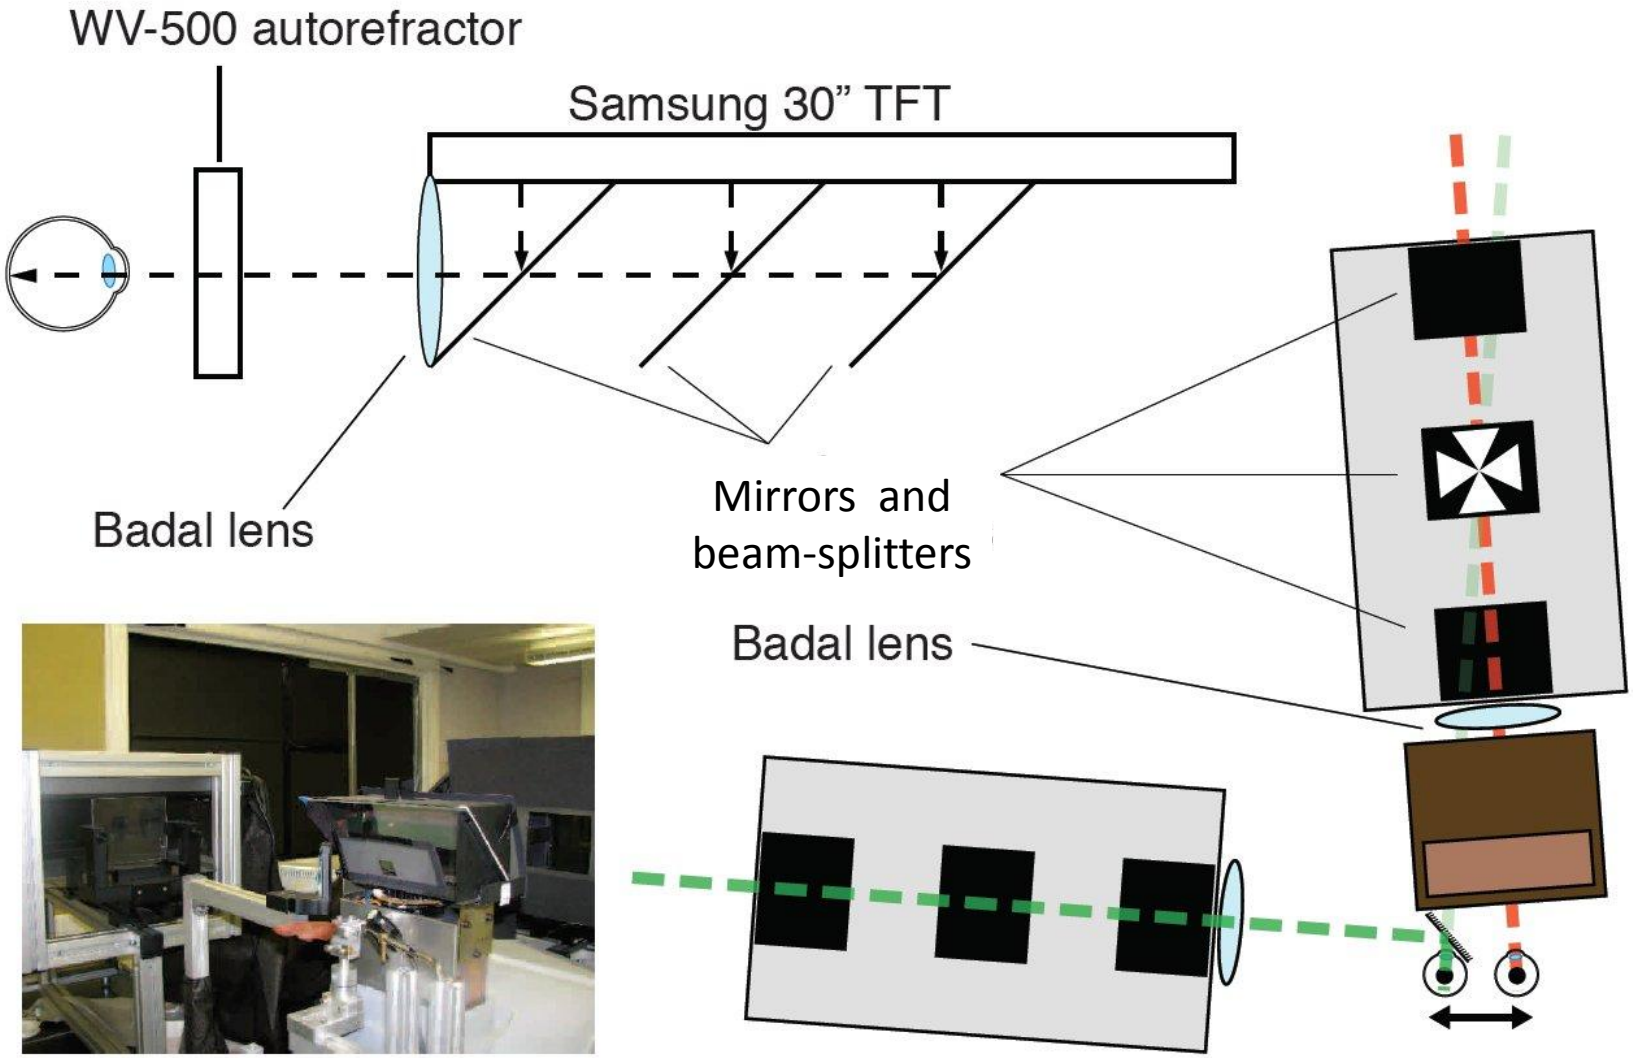

Kevin J. MacKenzie and Simon J. Watt, A display with multiple focal planes can stimulate continuous variations in accommodation, *Vision Sciences Society Annual Meeting*, 2009.

Dynamic (variable-power) lens

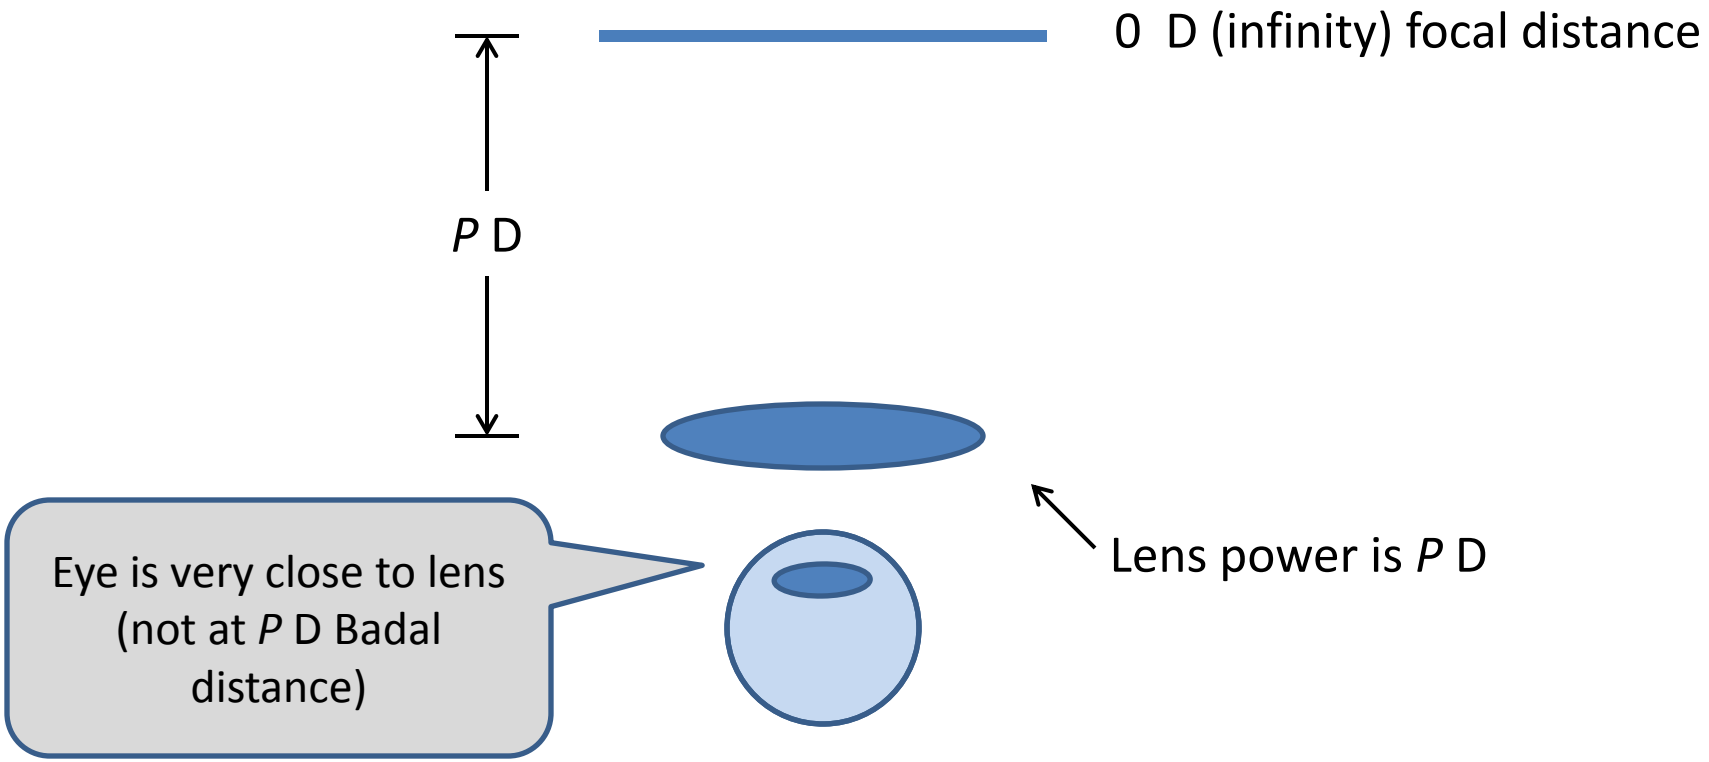

Dynamic (variable-power) lens

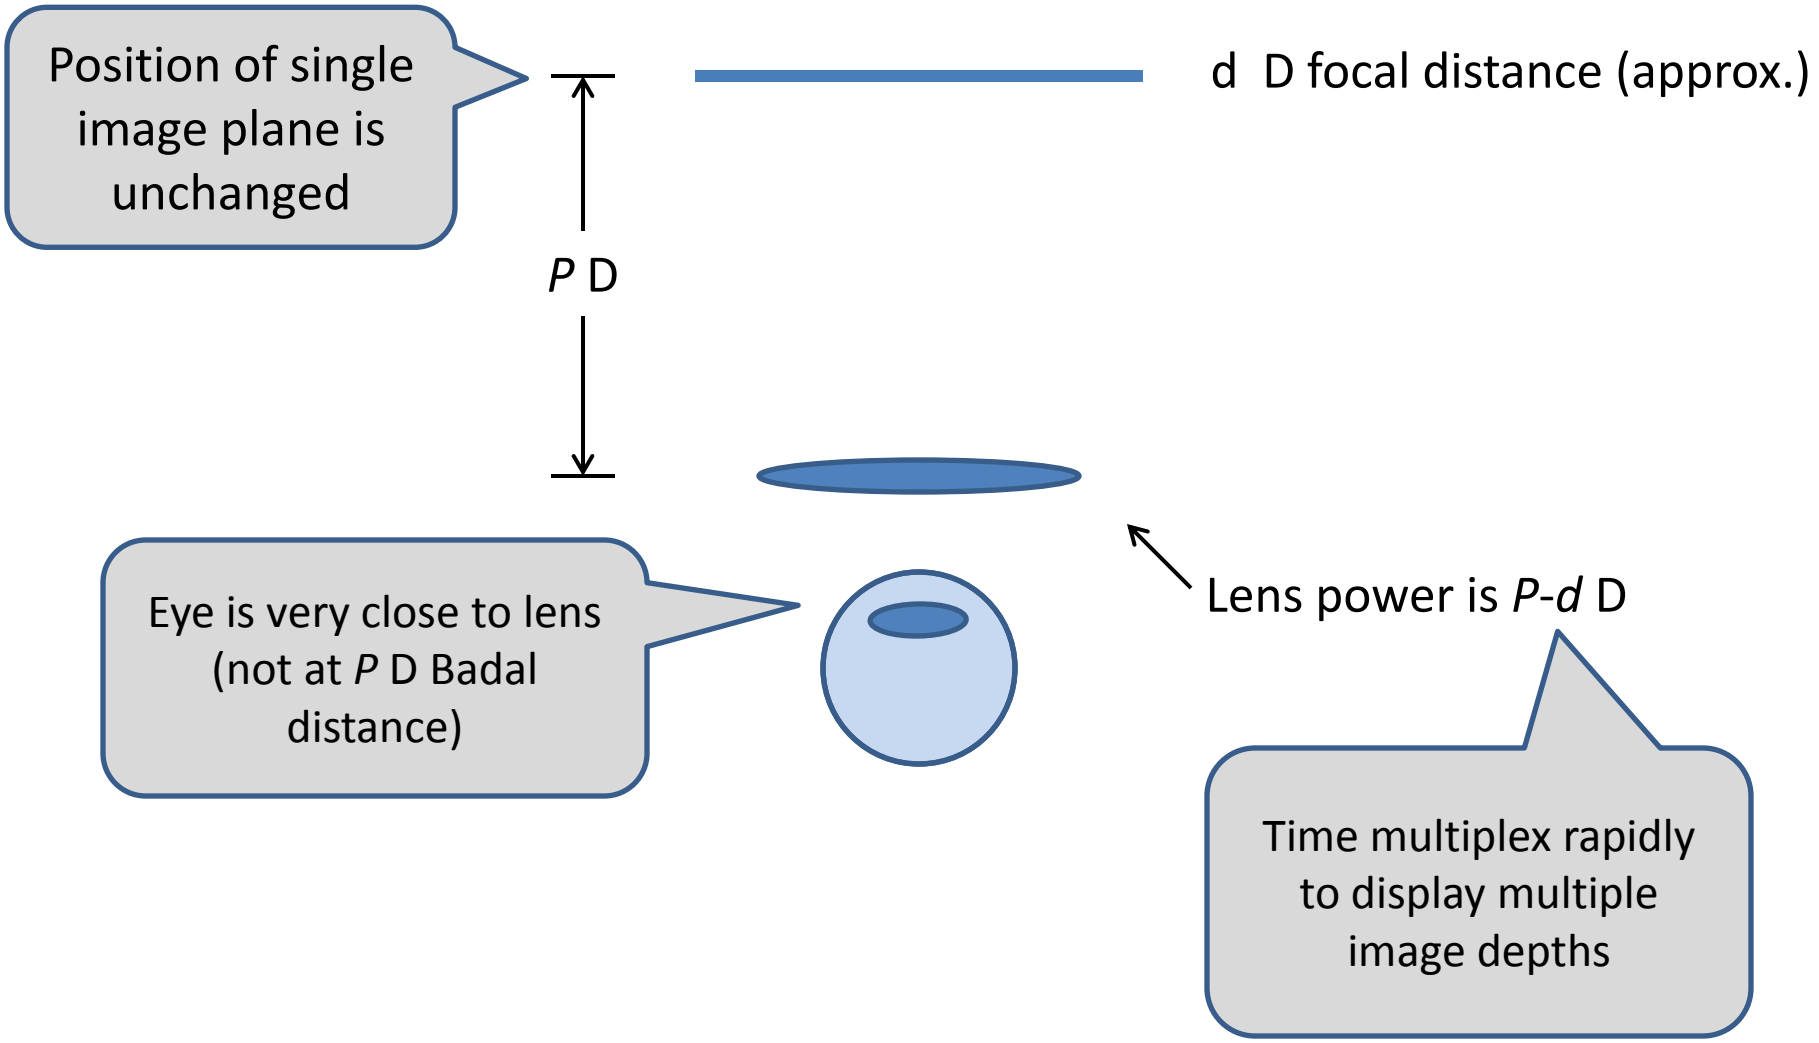

New Berkeley system

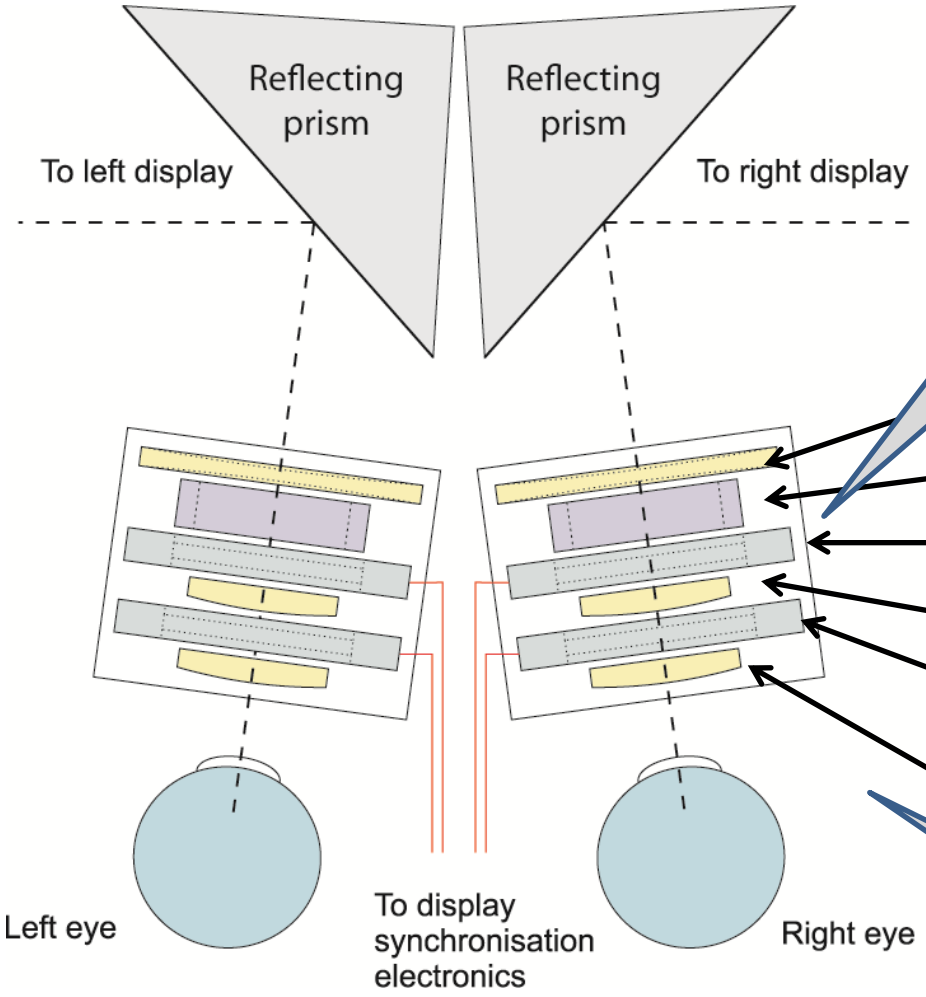

Lenses can be switched at up to 1000 Hz

Static glass lens

Polarizing filter

FLC 1

Birefringent lens 1

FLC 2

Birefringent lens 2

Lens stack delivers four focal distances at 45 Hz

David Hoffman, SPIE slide deck, 2009

New Berkeley system

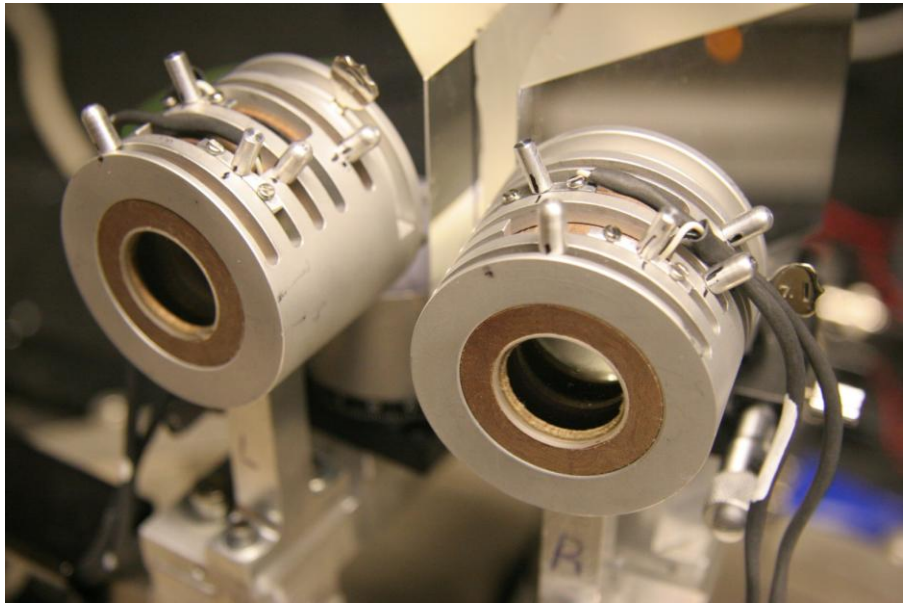

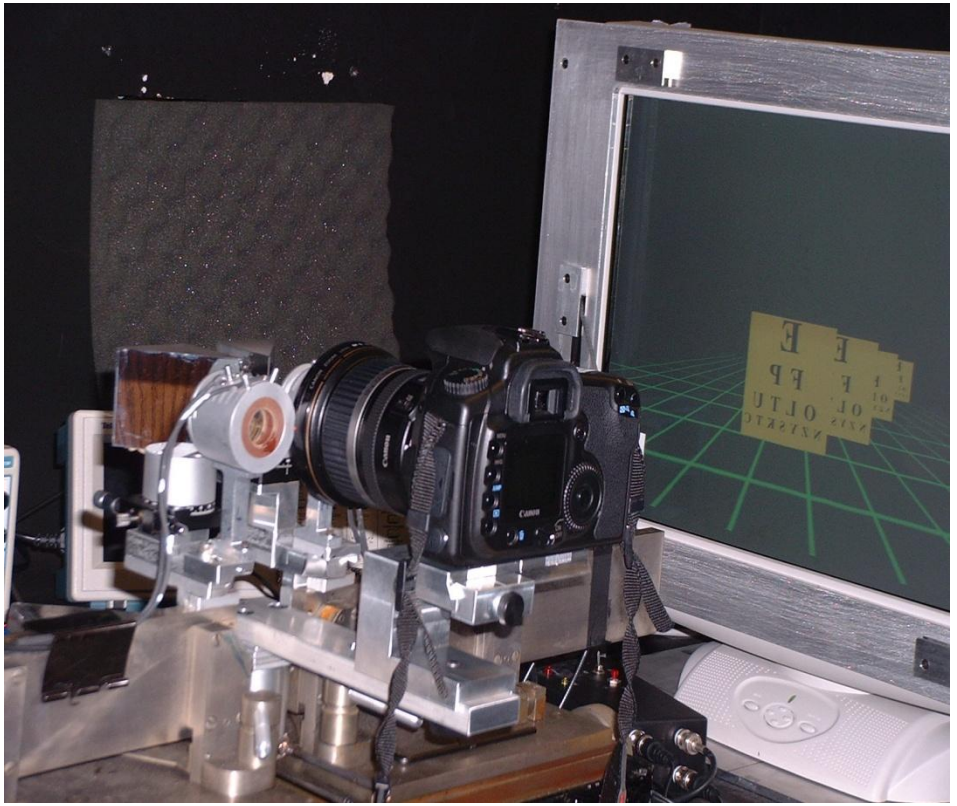

Outline

- Volumetric display
- Fixed-viewpoint volumetric display
- *Practical* fixed-viewpoint volumetric display
- Comparison of display approaches

Display comparison

Display type	Focus cues	Voxel $O(n^x)$	Hidden-surface elimination	Correct silhouettes
Stereo (e.g., theater)	X	2		
Auto-multiscopic volumetric		3	X	X
Fixed-viewpoint volumetric		2		X
Defocused-laser retinal scan		2		X
Track accommodation and render blur		2		
Integral imaging		4		
Holographic		4		

Should “3D” mean that all cues are correct, including stereopsis, head-motion parallax, **and** accommodation and retinal blur?

Fixed-viewpoint volumetric summary

- Nice qualities:
  - Able to create high-quality imagery with focus cues
  - Moderate voxel count
  - Does not require eye tracking
  - Tolerable shortcomings (head-mounting, silhouettes)
- May become a practical solution
  - Useful for scientists now
  - Perhaps for professionals and consumers in the future

Collaborators

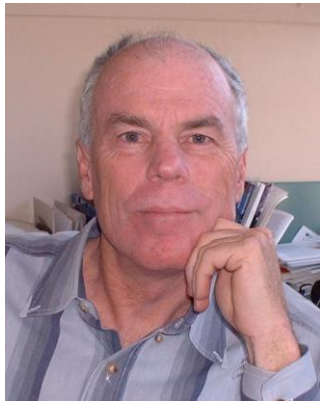

Marty Banks
UC Berkeley

Gordon Love
Durham Univ.

Simon Watt
Bangor Univ., Wales

Ahna Girshick
NYU

David Hoffman
UC Berkeley

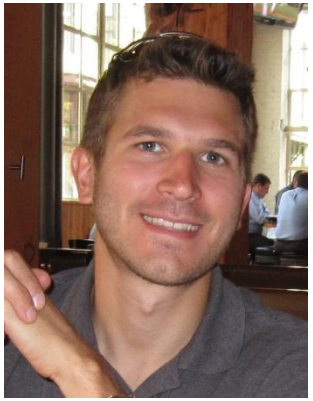

Robin Held
UC Berkeley

Microsoft Research, Silicon Valley, 2010

Demos up front

End