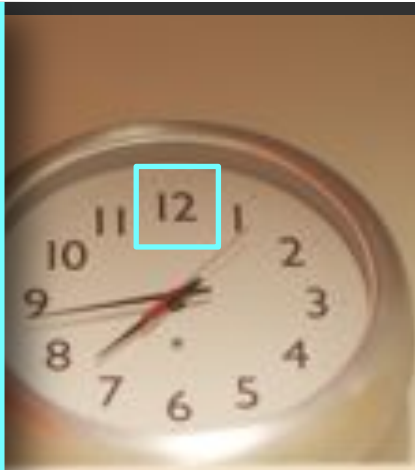
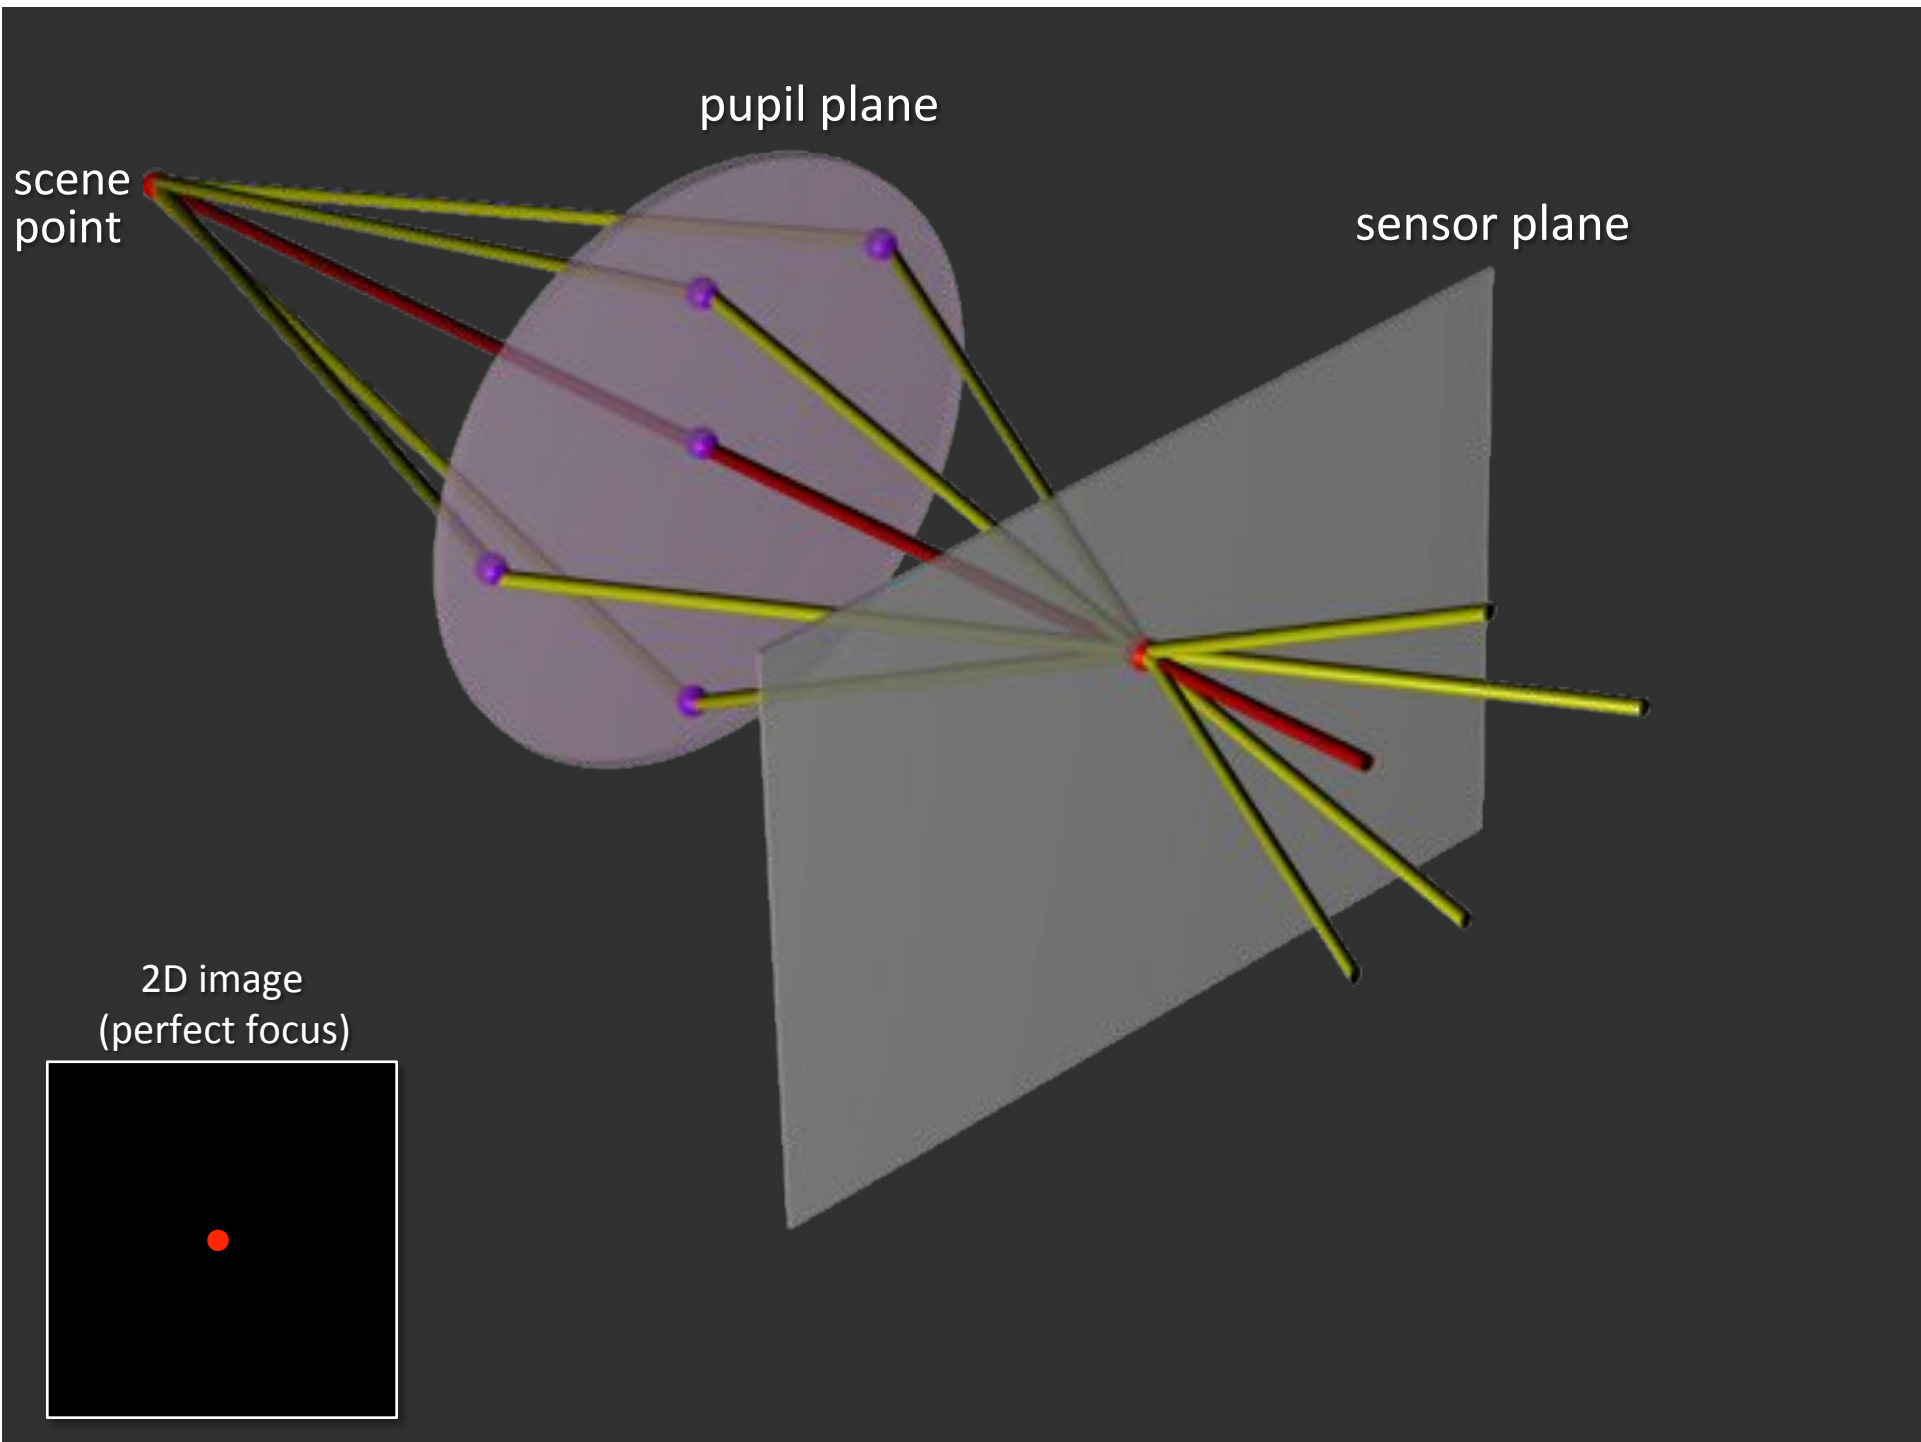


What Does an Aberrated Photo Tell Us About the Lens and the Scene?

Huixuan Tang
Kyros Kutulakos
University of Toronto



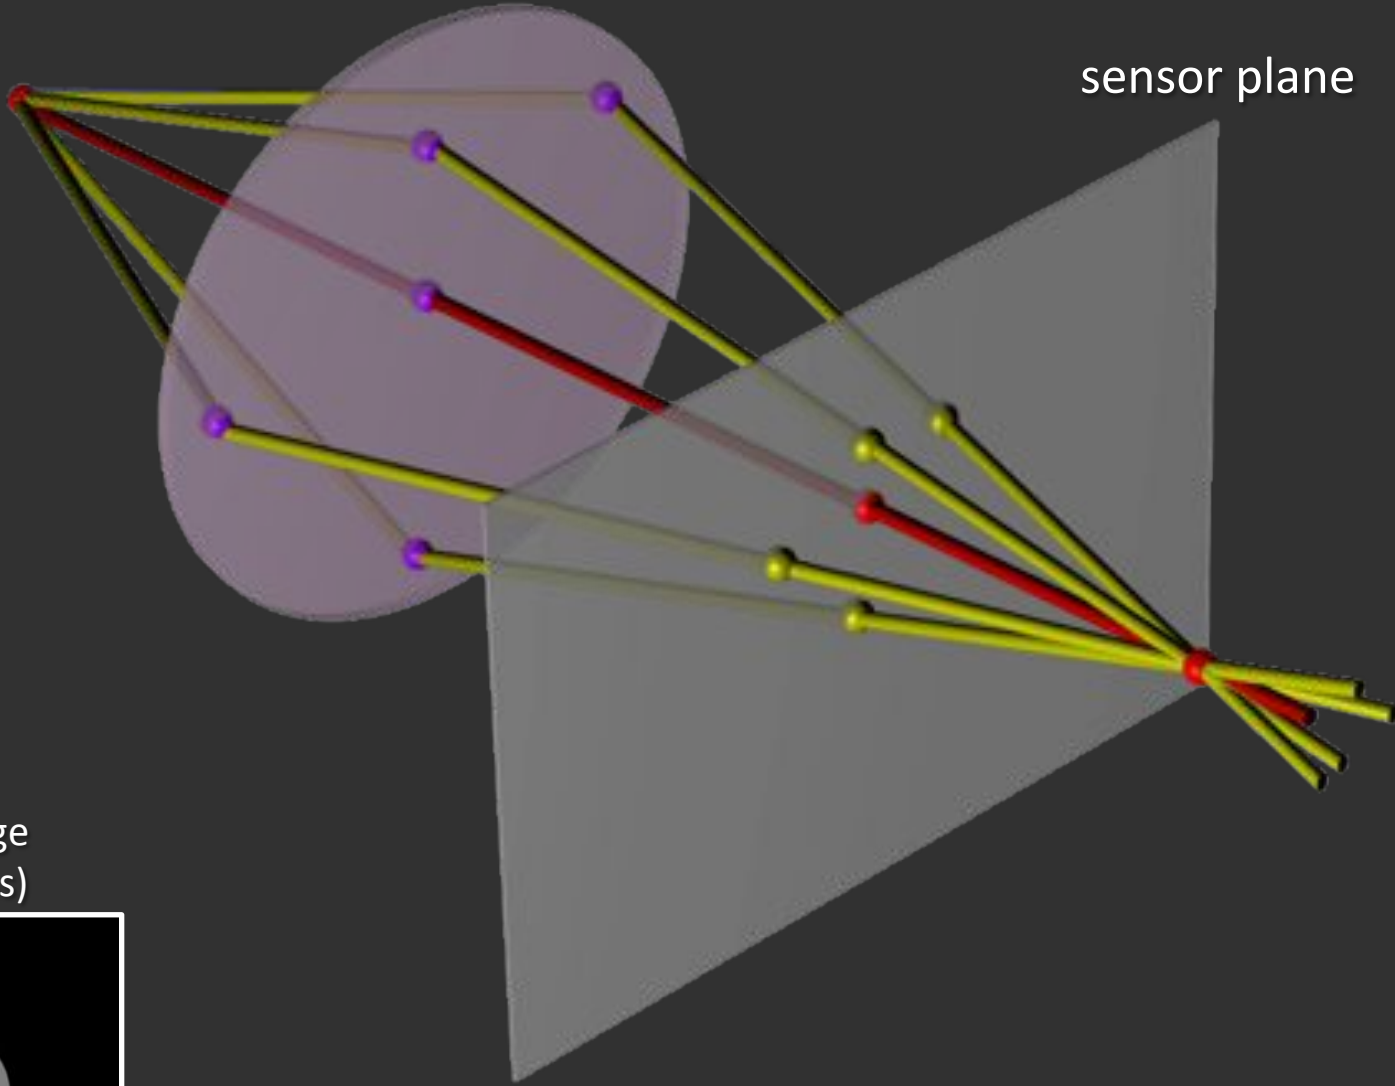



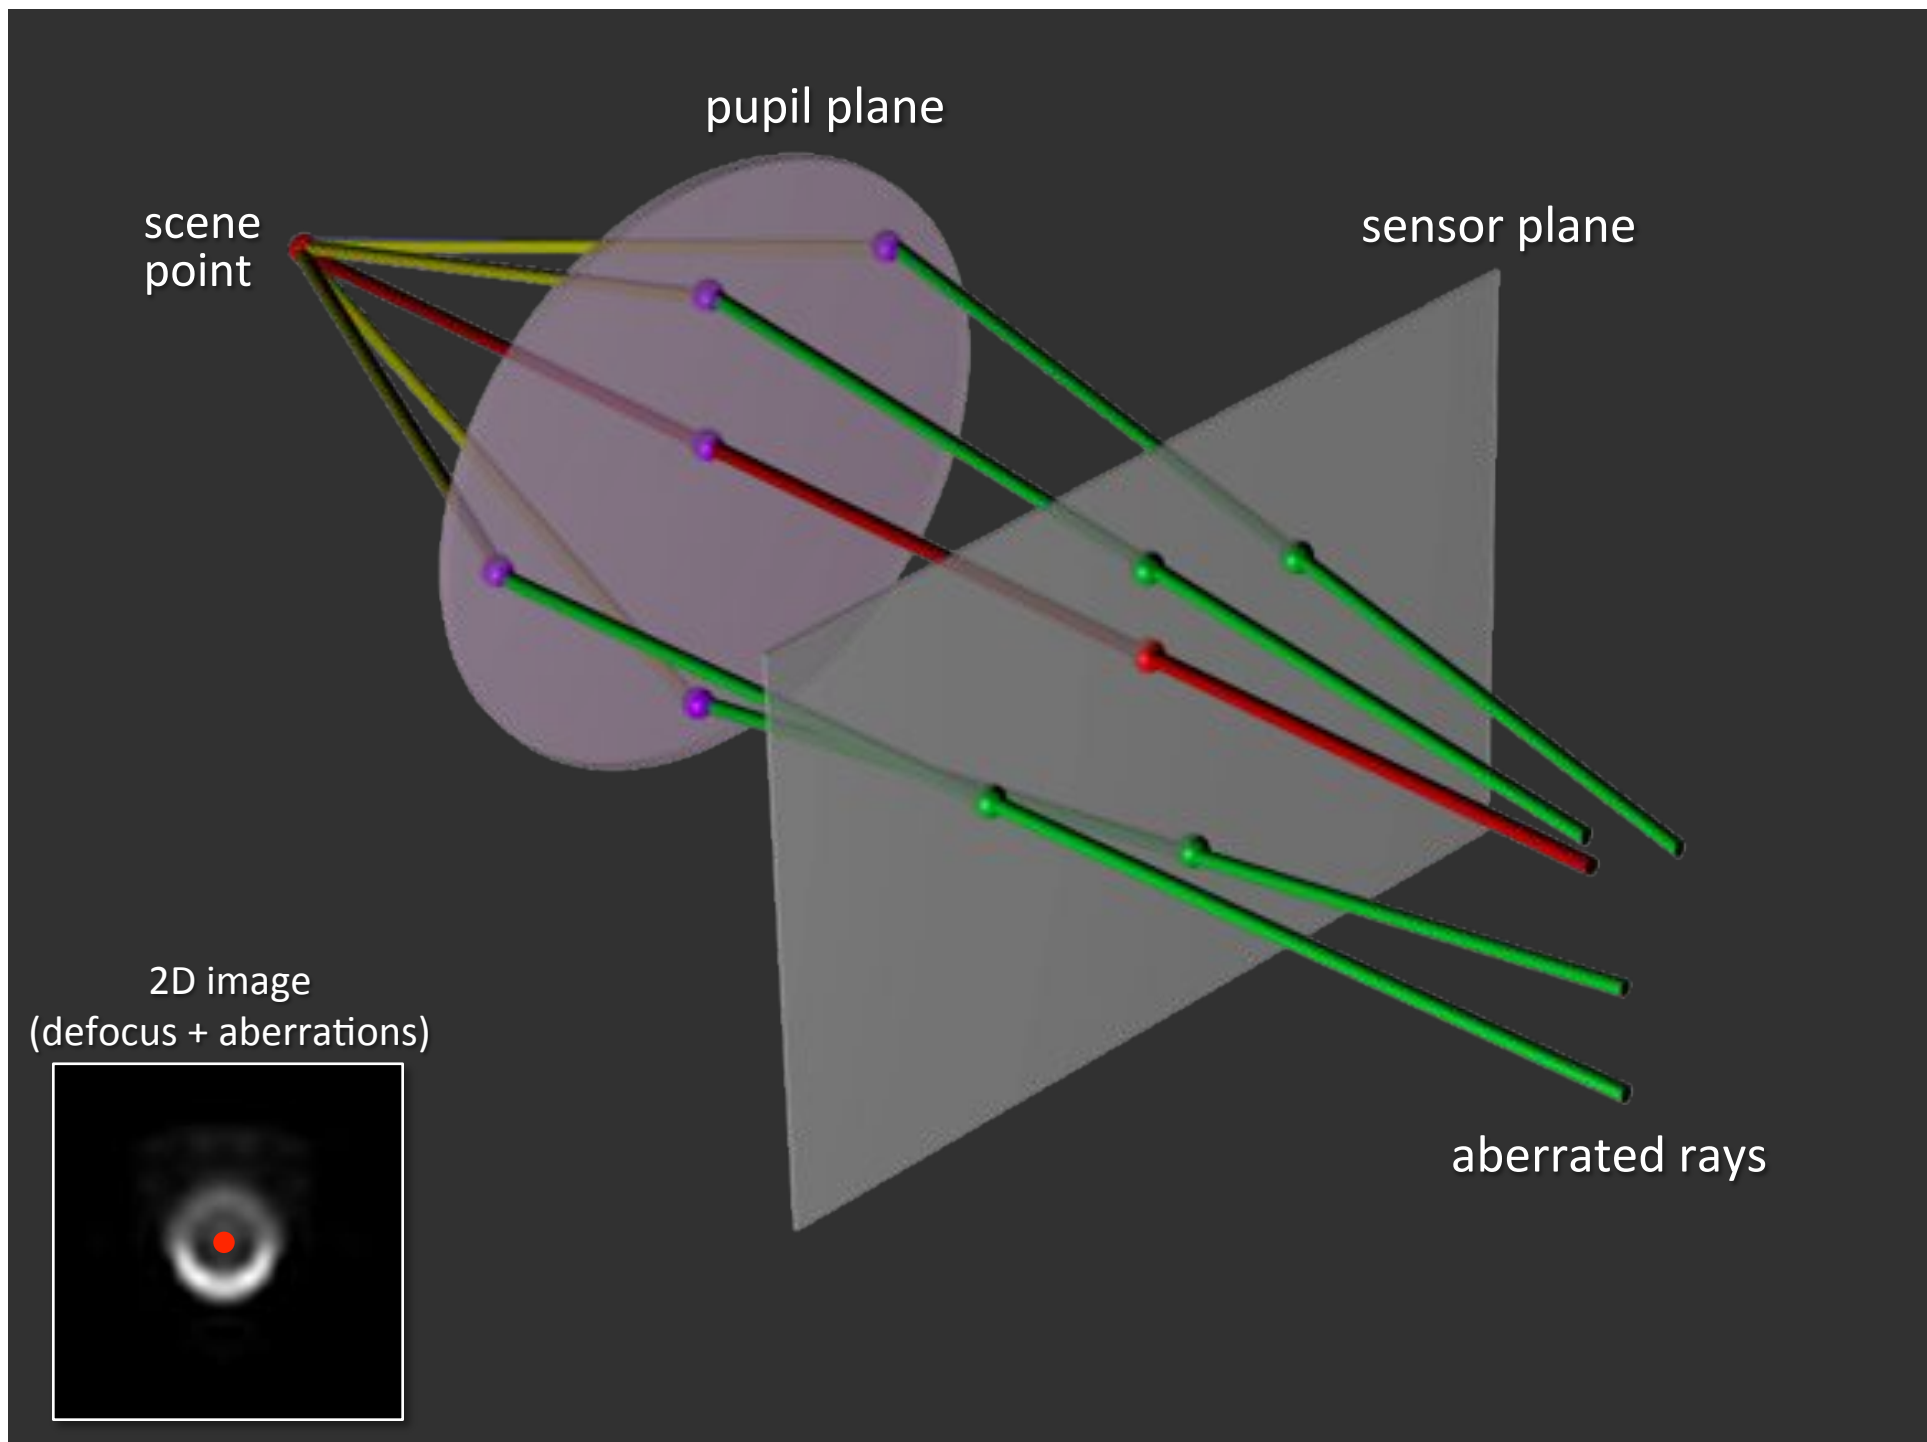
scene point

pupil plane

sensor plane

2D image
(defocus)






scene
point


pupil plane

sensor plane

2D image
(defocus + aberrations)

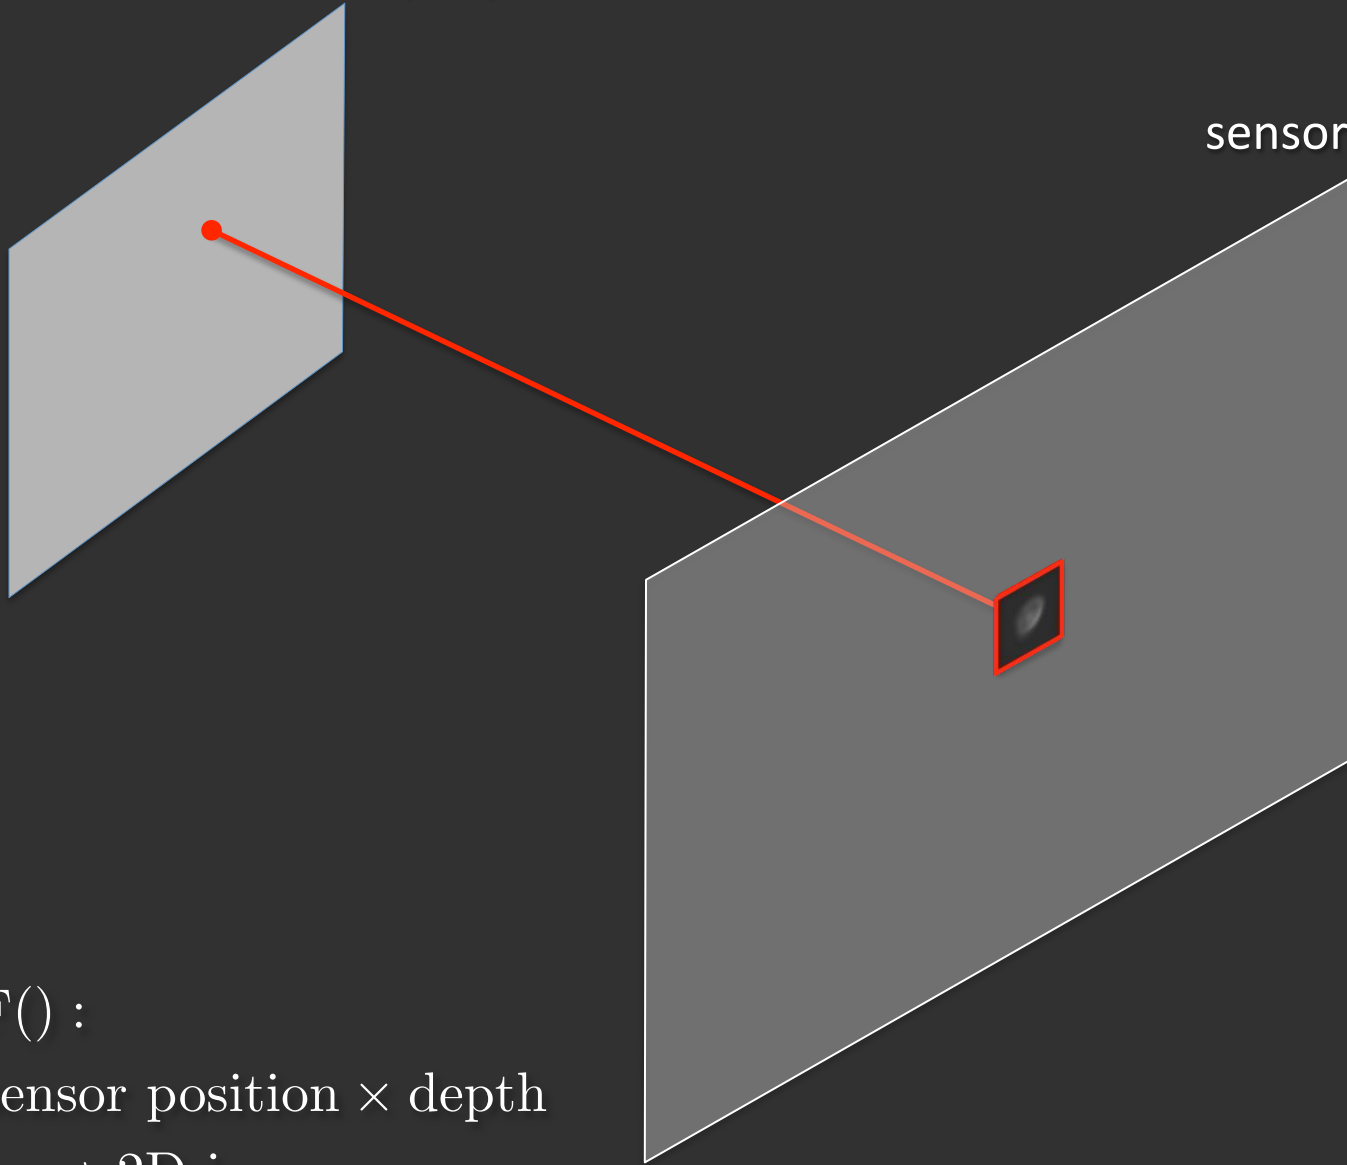


aberrated rays



depth plane D_1

sensor plane




PSF() :

sensor position \times depth

\rightarrow 2D image

depth plane D_1


sensor plane



PSF() :

sensor position \times depth

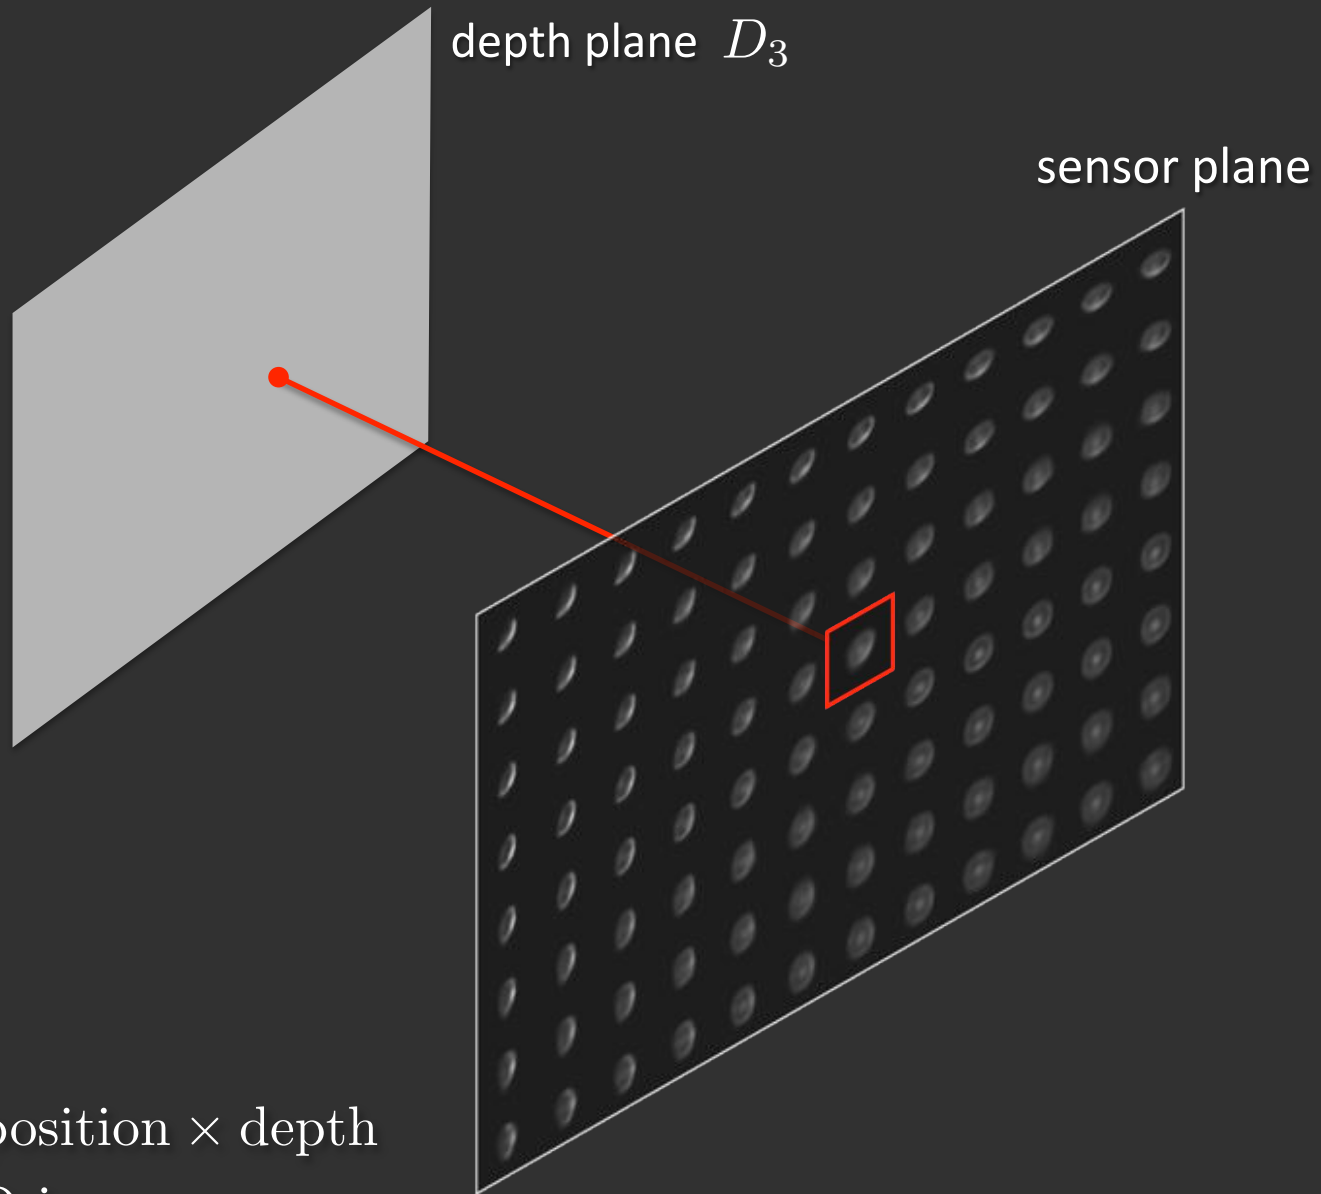
\rightarrow 2D image




PSF() :

sensor position \times depth

\rightarrow 2D image




PSF() :
sensor position \times depth
 \rightarrow 2D image



PSF() :
sensor position \times depth
 \rightarrow 2D image

- depends on depth & lens
- 1 photo \rightarrow 4D scene-dependent slice of PSF
- hard to capture
- hard to invert

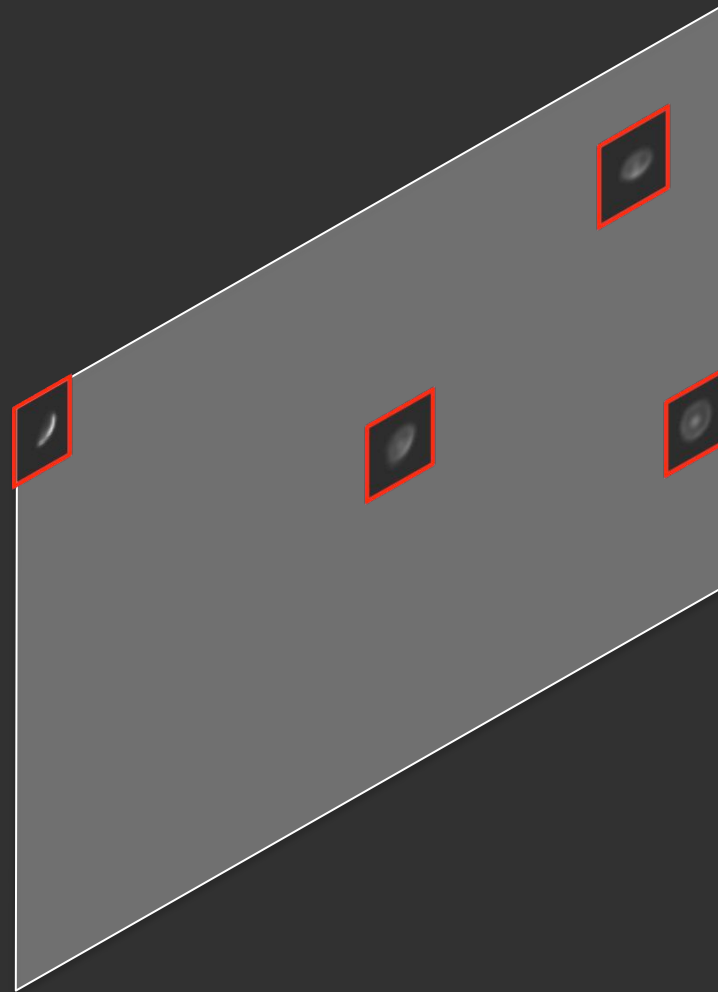


PSF() :
sensor position \times depth
 \rightarrow 2D image

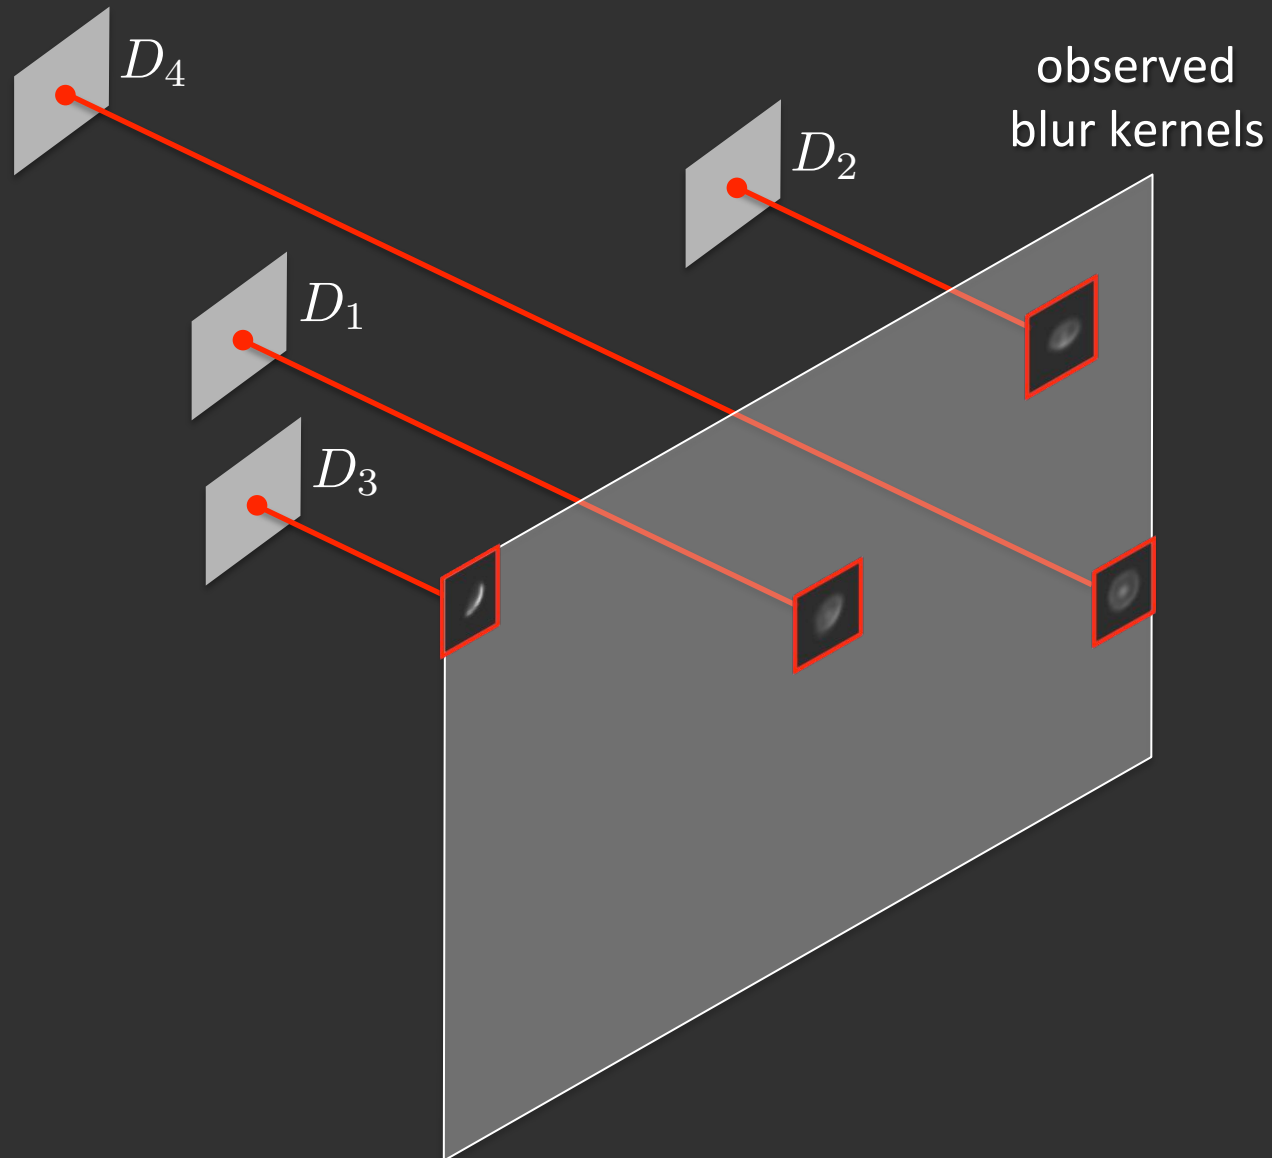
- depends on depth & lens
- 1 photo \rightarrow 4D scene-dependent slice of PSF
- **hard to capture**
- **hard to invert**

given $K \geq 1$ blur kernels from one photo ...


observed
blur kernels




... can we infer depth?




... can we infer the full 5D PSF?



... can we infer the full 5D PSF?



... and can we infer the lens parameters?




Key assumptions

- 3rd order Seidel aberrations
- geometric optics
- rotationally-symmetric lens
- ≥ 1 unknown lens elements
- known optical center

related work


direct aberration
measurement



[Liang et al 94][Ng & Hanrahan 06]

requires physical
access to camera


empirical PSF
modeling



[Joshi et al 08][Kee et al 11]

significant data acquisition,
limited inference power

blind/non-blind
aberration correction




[Shih et al 12][Schuler et al 12]

interaction w/ defocus
not well understood


related work

direct aberration measurement




[Liang et al 94][Ng & Hanrahan 06]

empirical PSF modeling




[Joshi et al 08][Kee et al 11]

blind/non-blind aberration correction



[Shih et al 12][Schuler et al 12]




[Favaro 10]

depth from defocus

assumes aberration-free imaging

related work

direct aberration
measurement




[Liang et al 94][Ng & Hanrahan 06]

requires physical
access to camera

inference from 1
aberrated photo

empirical PSF
modeling




[Joshi et al 08][Kee et al 11]

significant data acquisition,
limited inference power

physics-based model
of 5D PSF formation

blind/non-blind
aberration correction




[Shih et al 12][Schuler et al 12]

interaction w/ defocus
not well understood


unified analysis of
aberration & defocus

three main results



single-point inference

regardless of deblurring quality,
cannot predict lens blur at other
depths/image locations



multi-point inference

reconstruction of depth & 5D PSF
may be possible from just 1 shot
despite ambiguity
in Seidel coefs & K defocus levels




practical modeling/inference of 5D PSFs

instead of modeling the PSF directly,
model the aberrated rays

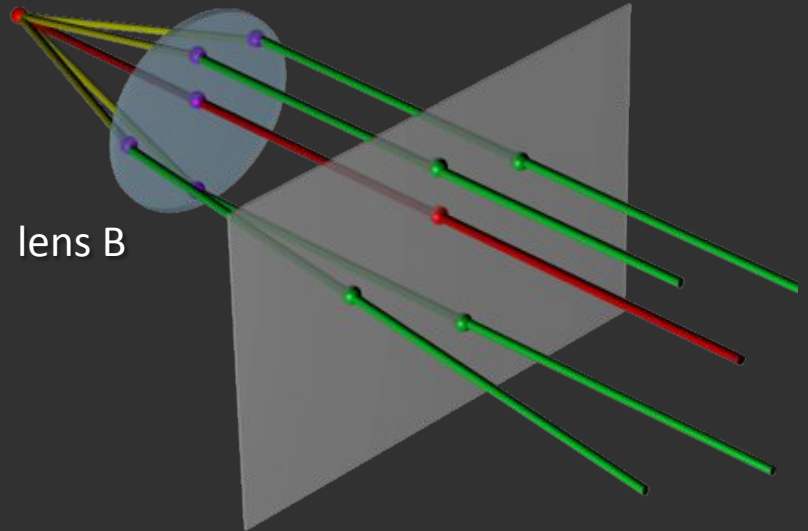
blur-equivalent arrangements

point-lens arrangements that send aberrated rays to identical pixels


point at depth D



point at depth D'




ray-sensor intersections




Blur-consistent arrangements

point-lens arrangements that send aberrated rays to identical pixels

point at depth D



point at depth D'





blur kernel



Blur-consistent arrangements


point-lens arrangements that send aberrated rays to identical pixels




approach: we characterize the space of blur-equivalent arrangements

- i. the PSF integral
- ii. the ray intersection function
- iii. single-source inference
- iv. multi-source inference
- iv. preliminary experiments

blur formation model




blur kernel




$$= \int_{x^2 + y^2 \leq R}$$

ray-sensor intersections



vignetting function



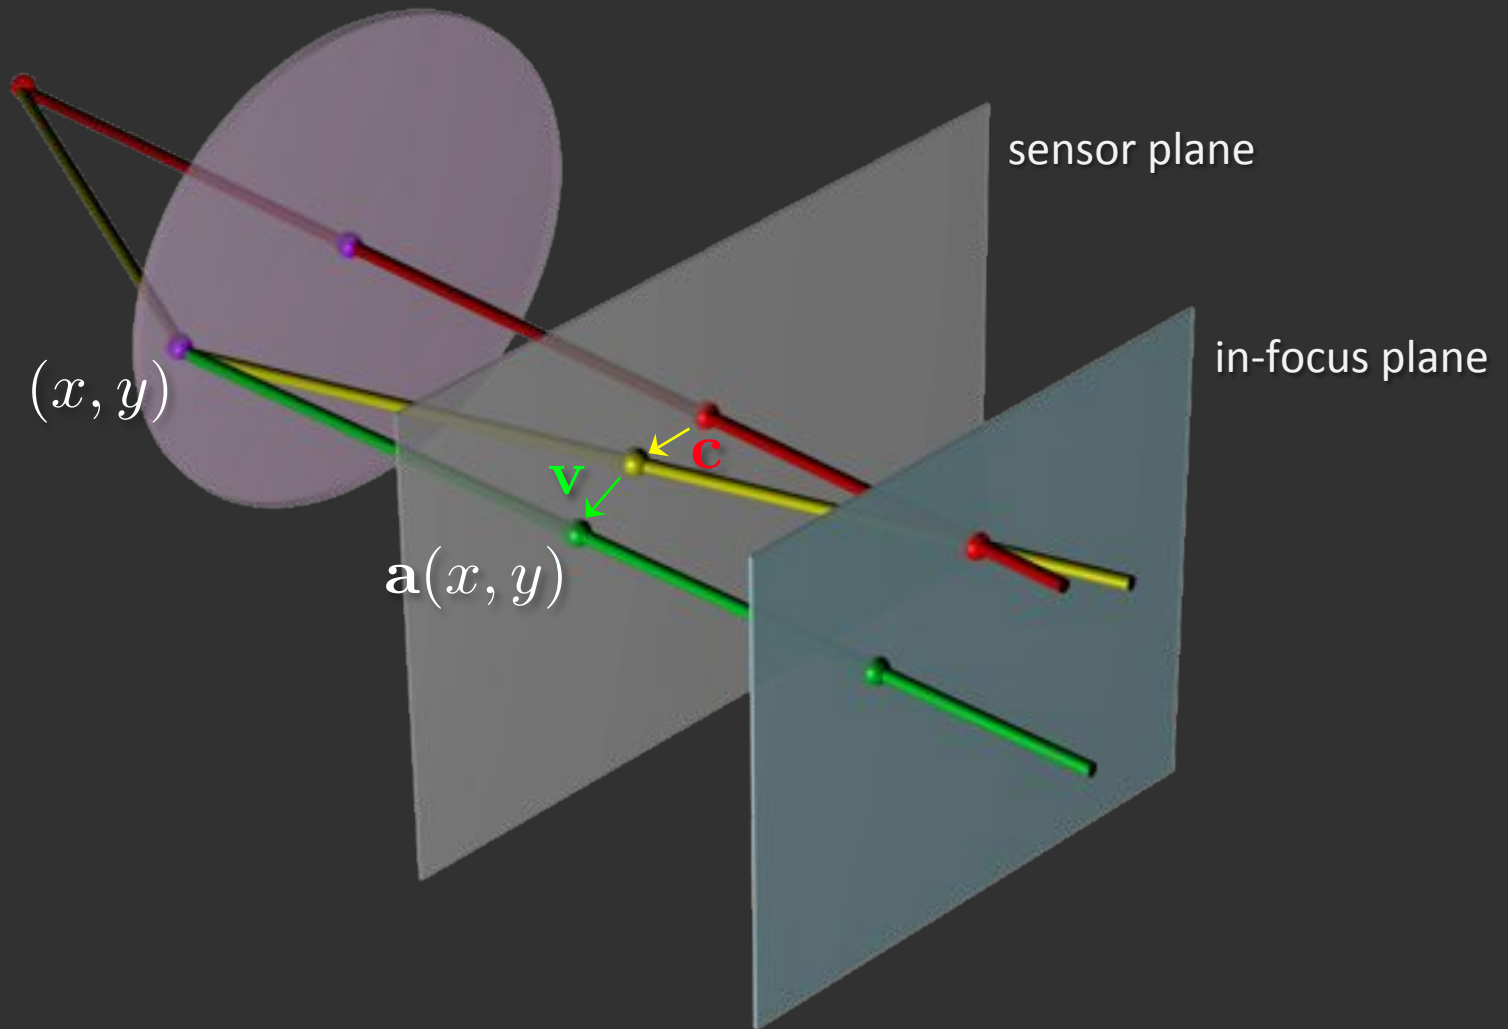
$dx dy$

complex function of lens params & defocus

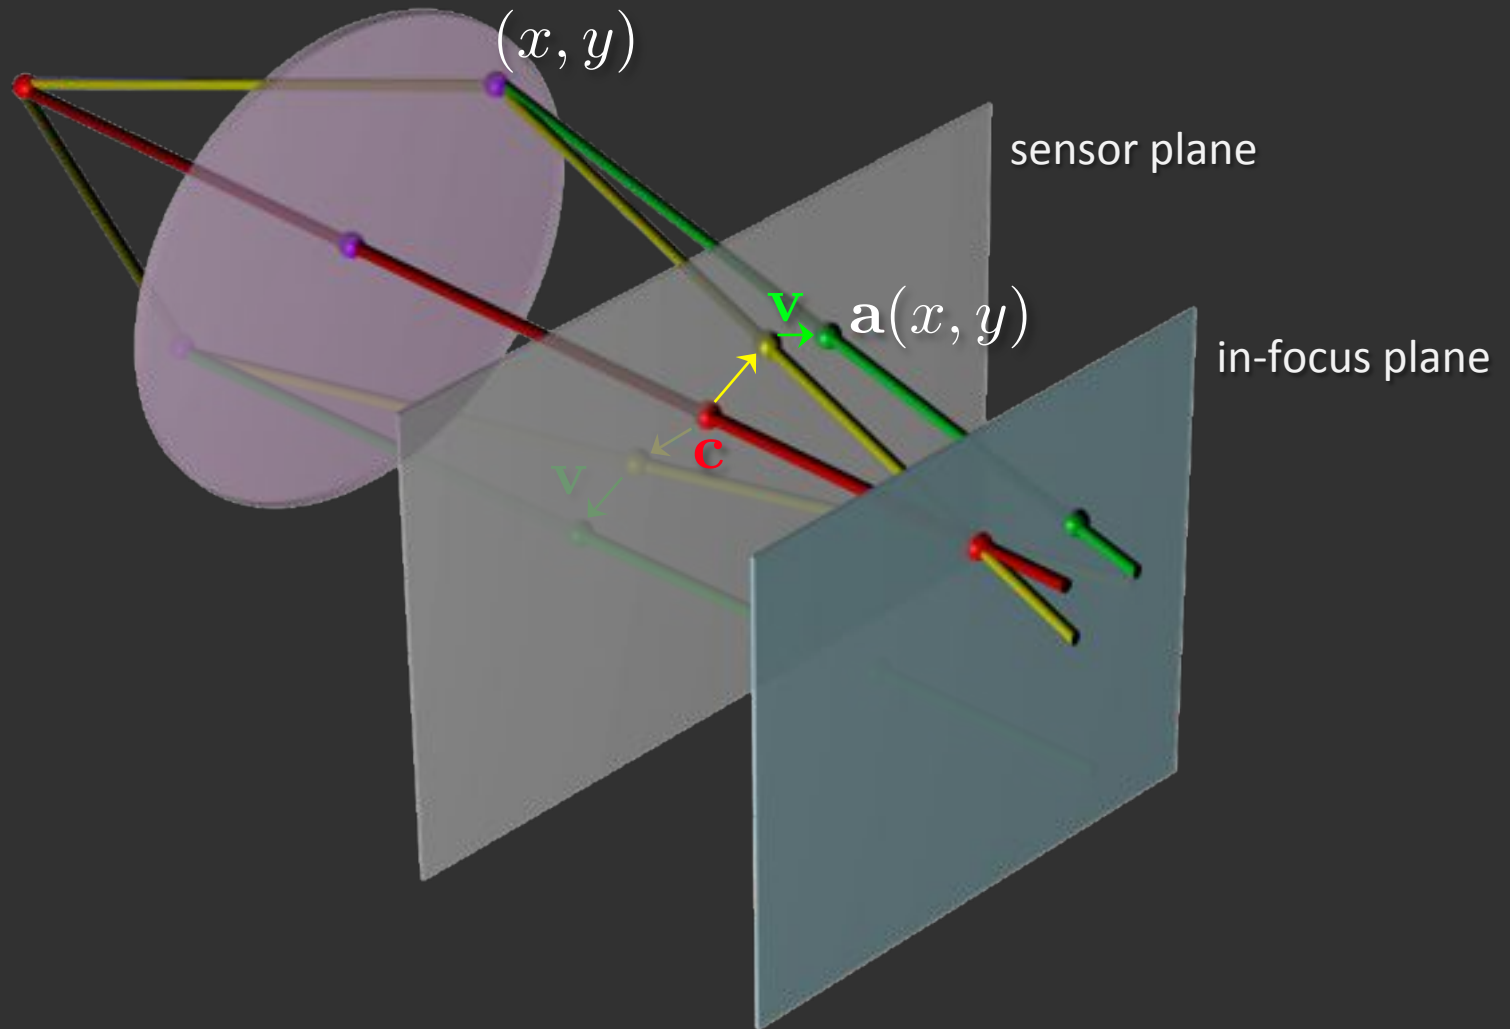
simple 3rd order polynomial

intersection of discs

ii. the ray intersection function
(for 3rd order Seidel aberrations)




$$\begin{array}{ccccccc}
 \text{ray intersection} & = & \text{perspective} & + & \text{defocus} & + & \text{aberration} \\
 \mathbf{a}(x, y) & & \mathbf{c} & & (x, y) \cdot D & & \mathbf{v}
 \end{array}$$



$$\text{ray intersection } \mathbf{a}(x, y) = \text{perspective } \mathbf{c} + \text{defocus } (x, y) \cdot D + \text{aberration } \mathbf{v}(x, y, \mathbf{c}, D)$$

the Seidel displacement field




$$\mathbf{v}(x, y, \mathbf{c}, D)$$

- 3rd order polynomial of pupil coordinates, perspective projection & defocus
- fully determined by 5 lens-specific aberration parameters


$$\text{ray intersection } \mathbf{a}(x, y) = \text{perspective } \mathbf{c} + \text{defocus } (x, y) \cdot D + \text{aberration } \mathbf{v}(x, y, \mathbf{c}, D)$$

the Seidel displacement field



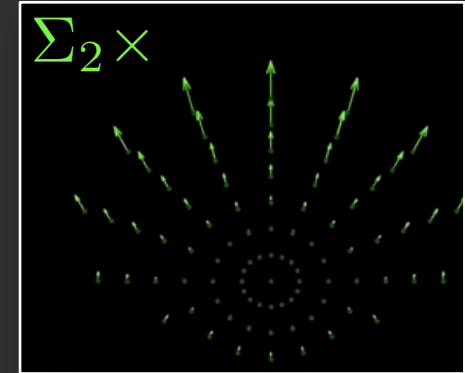
$\mathbf{v}(x, y, \mathbf{c}, D)$

=




spherical

+




coma



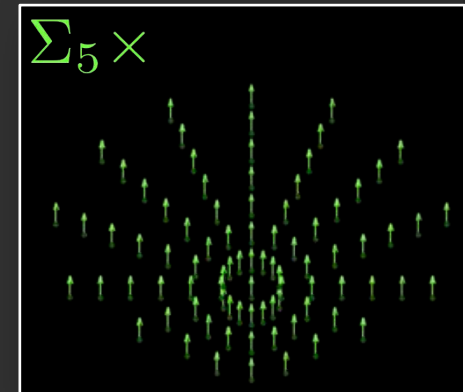
astigmatism

+



field curvature

+



field distortion

$$\text{ray intersection } \mathbf{a}(x, y) = \text{perspective } \mathbf{c} + \text{defocus } (x, y) \cdot D + \text{aberration } \mathbf{v}(x, y, \mathbf{c}, D)$$

the ray intersection polynomial

$$\left(\begin{array}{c} \left[\begin{array}{c} 0 \\ 0 \\ 0 \\ D \\ |c| \end{array} \right] + \left[\begin{array}{ccccc} 1 & D & D^2 & D^2 & D^3 \\ 0 & |c| & D|c| & 2D|c| & 3D|c|^2 \\ 0 & 0 & |c|^2 & 0 & 2D|c|^2 \\ 0 & 0 & 0 & |c|^2 & D|c|^2 \\ 0 & 0 & 0 & 0 & |c|^3 \end{array} \right] \left[\begin{array}{c} \Sigma_1 \\ \Sigma_2 \\ \Sigma_3 \\ \Sigma_4 \\ \Sigma_5 \end{array} \right] \end{array} \right)^T \left[\begin{array}{c} [x^2 + y^2](x, y) \\ (3x^2 + y^2, 2xy) \\ (x, 0) \\ (x, y) \\ (1, 0) \end{array} \right]$$

perspective & defocus
aberrations
pupil coords

$$\langle \mathbf{c}, D, R, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5 \rangle$$

specific pupil
to point radius


aberration
parameters

$$\text{ray intersection } \mathbf{a}(x, y) = \text{perspective } \mathbf{c} + \text{defocus } (x, y) \cdot D + \text{aberration } \mathbf{v}(x, y, \mathbf{c}, D)$$

iii. single-source inference

limits of local inference


what does 1 blur kernel in 1 photo tell us about the lens & the scene?




limits of local inference

what does 1 blur kernel in 1 photo tell us about the lens & the scene?

$$\langle \mathbf{c}, D, R, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5 \rangle$$




?



pupil radius ambiguity

Lemma: there is a blur-compatible solution for every $R' > 0$

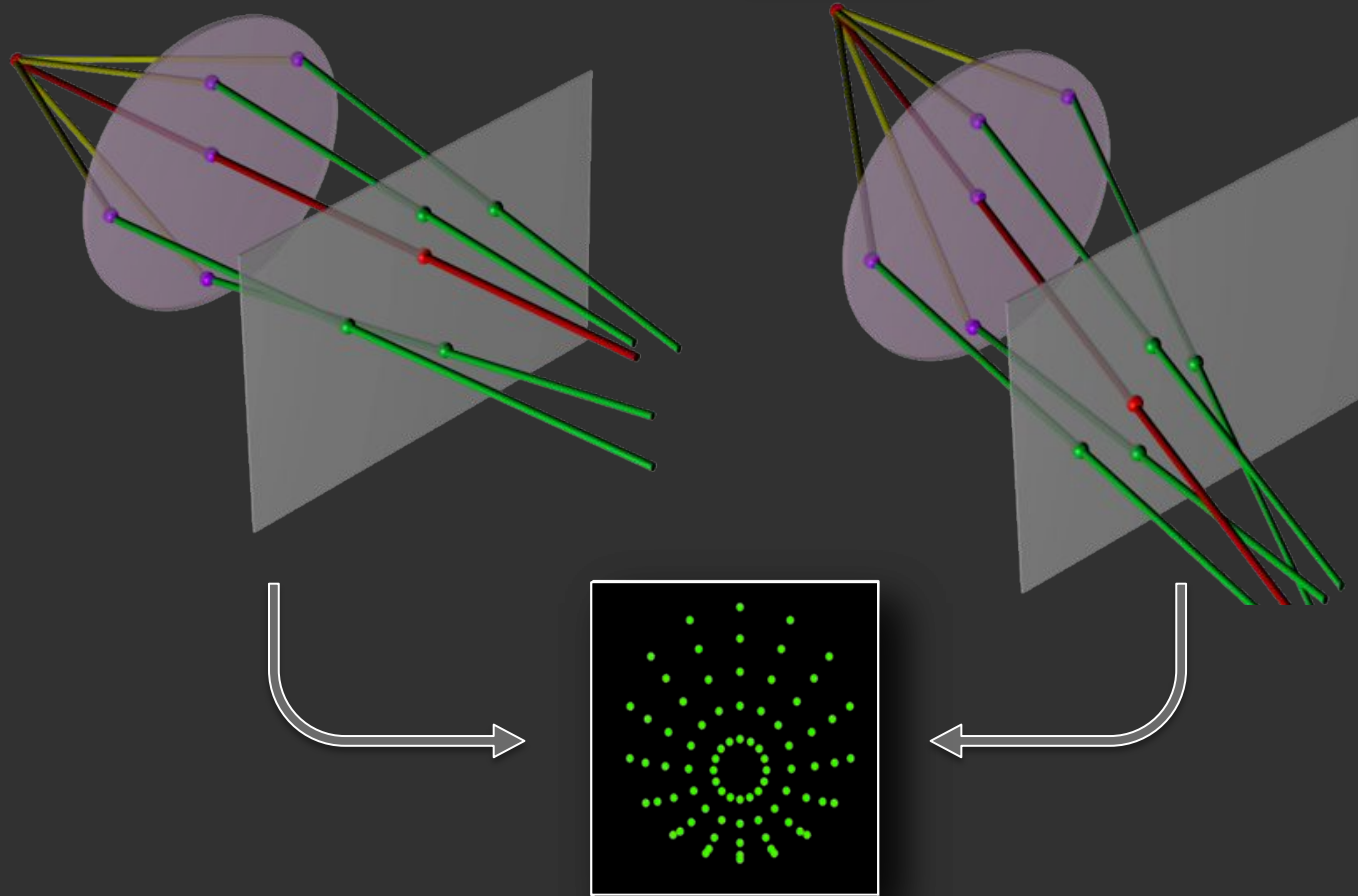
$$\langle \mathbf{c}, D, R, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5 \rangle \quad \langle \mathbf{c}, D', \boxed{R'}, \Sigma'_1, \Sigma'_2, \Sigma'_3, \Sigma'_4, \Sigma_5 \rangle$$



depth & projection ambiguity

Lemma: there is a blur-compatible solution for every $D', c' > 0$


$$\langle \mathbf{c}, D, R, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5 \rangle \quad \langle \mathbf{c}', D', R, \Sigma'_1, \Sigma'_2, \Sigma'_3, \Sigma'_4, \Sigma'_5 \rangle$$



depth & projection ambiguity

Lemma: there is a blur-compatible solution for every $D', c' > 0$

$$\langle \mathbf{c}, D, R, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5 \rangle \quad \langle \mathbf{c}', D', R, \Sigma'_1, \Sigma'_2, \Sigma'_3, \Sigma'_4, \Sigma'_5 \rangle$$



$$\begin{bmatrix} 0 \\ 0 \\ 0 \\ D \\ |c| \end{bmatrix} + \begin{bmatrix} 1 & D & D^2 & D^2 & D^3 \\ 0 & |c| & D|c| & 2D|c| & 3D|c|^2 \\ 0 & 0 & |c|^2 & 0 & 2D|c|^2 \\ 0 & 0 & 0 & |c|^2 & D|c|^2 \\ 0 & 0 & 0 & 0 & |c|^3 \end{bmatrix} \begin{bmatrix} \Sigma_1 \\ \Sigma_2 \\ \Sigma_3 \\ \Sigma_4 \\ \Sigma_5 \end{bmatrix}$$

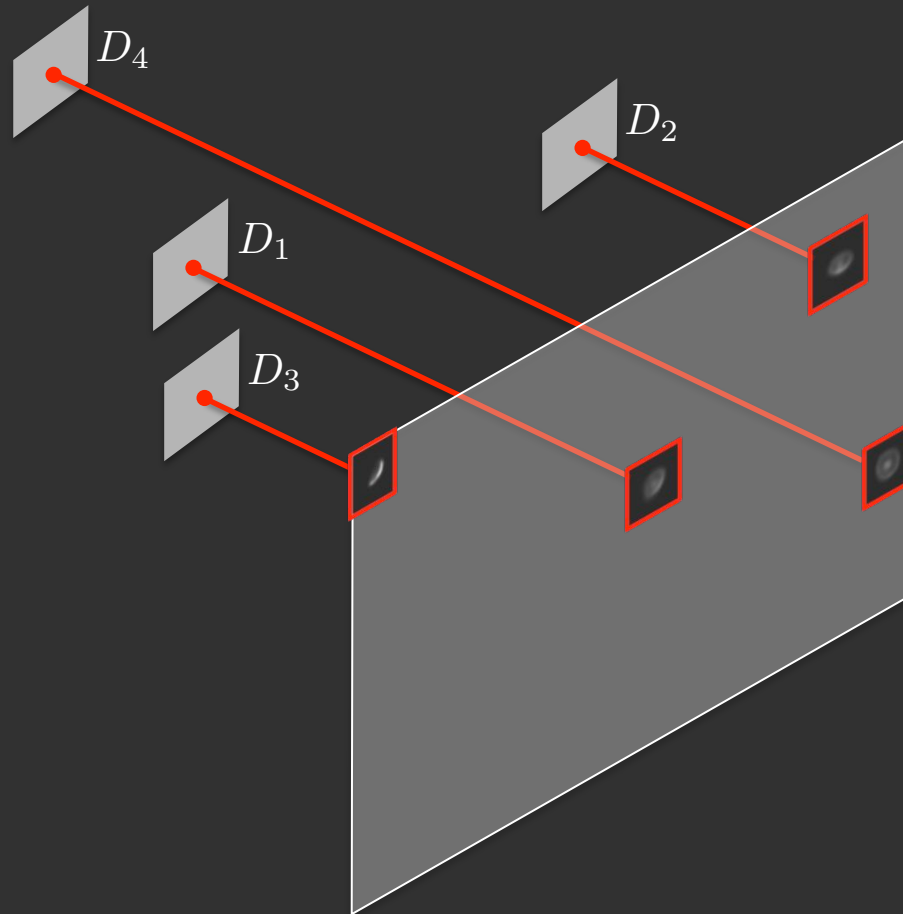
$=$

$$\begin{bmatrix} 0 \\ 0 \\ 0 \\ D' \\ |c'| \end{bmatrix} + \begin{bmatrix} 1 & D' & D'^2 & D'^2 & D'^3 \\ 0 & |c'| & D'|c'| & 2D'|c'| & 3D'|c'|^2 \\ 0 & 0 & |c'|^2 & 0 & 2D'|c'|^2 \\ 0 & 0 & 0 & |c'|^2 & D'|c'|^2 \\ 0 & 0 & 0 & 0 & |c'|^3 \end{bmatrix} \begin{bmatrix} \Sigma'_1 \\ \Sigma'_2 \\ \Sigma'_3 \\ \Sigma'_4 \\ \Sigma'_5 \end{bmatrix}$$

iii. multi-point inference

multi-local inference

what do K blur kernels in 1 photo tell us about the lens & the scene?




multi-local inference

what do K blur kernels in 1 photo tell us about the lens & the scene?


$$\langle \mathbf{c}_1, D_1, R, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5 \rangle$$
$$\langle \mathbf{c}_2, D_2, R, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5 \rangle$$

5 lens unknowns

2K scene unknowns



?





power of multi-local inference

Lemma: space of blur-consistent arrangements is discrete for $K \geq 3$


$$\langle \mathbf{c}_1, D_1, R, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5 \rangle$$

$$\langle \mathbf{c}_2, D_2, R, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5 \rangle$$



v. preliminary experiments

single-point & multi-point inference



ground-truth acquisition of 5D PSF

1. focus at depth D
2. for $D' \in [D \pm 20\text{dof}, D + 20\text{dof}]$
take narrow & wide-aperture shots
estimate local blur kernels [Joshi 08]
3. repeat


inference

1. choose the shot for a pair (D, D')
2. choose K blur kernels from that shot
3. estimate $D' \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5$
4. predict blur kernels for all depths

evaluation

compare predicted &
ground-truth blur kernels

single-point & multi-point inference




ground-truth acquisition of 5D PSF

1. focus at depth D
2. for $D' \in [D \pm 20\text{dof}, D + 20\text{dof}]$
take narrow & wide-aperture shots
estimate local blur kernels [Joshi 08]
3. repeat

41 target positions




single-point & multi-point inference




ground-truth acquisition of 5D PSF

1. focus at depth D
2. for $D' \in [D \pm 20\text{dof}, D + 20\text{dof}]$
take narrow & wide-aperture shots
estimate local blur kernels [Joshi 08]
3. repeat

12x8 spatial position samples in $\frac{1}{4}$ field of view




single-point & multi-point inference




ground-truth acquisition of 5D PSF

1. focus at depth D
2. for $D' \in [D \pm 20\text{dof}, D + 20\text{dof}]$
take narrow & wide-aperture shots
estimate local blur kernels [Joshi 08]
3. repeat

12x8x41 blur kernels



single-point & multi-point inference



ground-truth acquisition of 5D PSF

1. focus at depth D
2. for $D' \in [D \pm 20\text{dof}, D + 20\text{dof}]$
take narrow & wide-aperture shots
estimate local blur kernels [Joshi 08]
3. repeat

inference ($K=1$)


1. choose the shot for a pair (D, D')
2. choose K blur kernels from that shot
3. estimate $\Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4$
4. predict blur kernels for all depths

optimization

$$\max_{\Sigma_1 \dots \Sigma_4} \text{corr}(\text{PSF}_{c, D', \Sigma_1 \dots \Sigma_5}, \text{PSF}_{data})$$


fitted depth = -20 DoF

predicted depth = -20 DoF




fitted depth = -20 DoF

predicted depth = -20 DoF




fitted depth = -20 DoF

predicted depth = 0 DoF




fitted depth = -20 DoF

predicted depth = 20 DoF



multi-point inference




ground-truth acquisition of 5D PSF

1. focus at depth D
2. for $D' \in [D \pm 20\text{dof}, D + 20\text{dof}]$
take narrow & wide-aperture shots
estimate local blur kernels [Joshi 08]
3. repeat

inference ($K > 1$)


1. choose the shot for a pair (D, D')
2. choose K blur kernels from that shot
3. estimate $D' \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4, \Sigma_5$
4. predict blur kernels for all (D, D')

fitted depth = -20 DoF




fitted depth = -20 DoF

predicted depth = -20 DoF




fitted depth = -20 DoF

predicted depth = 0 DoF

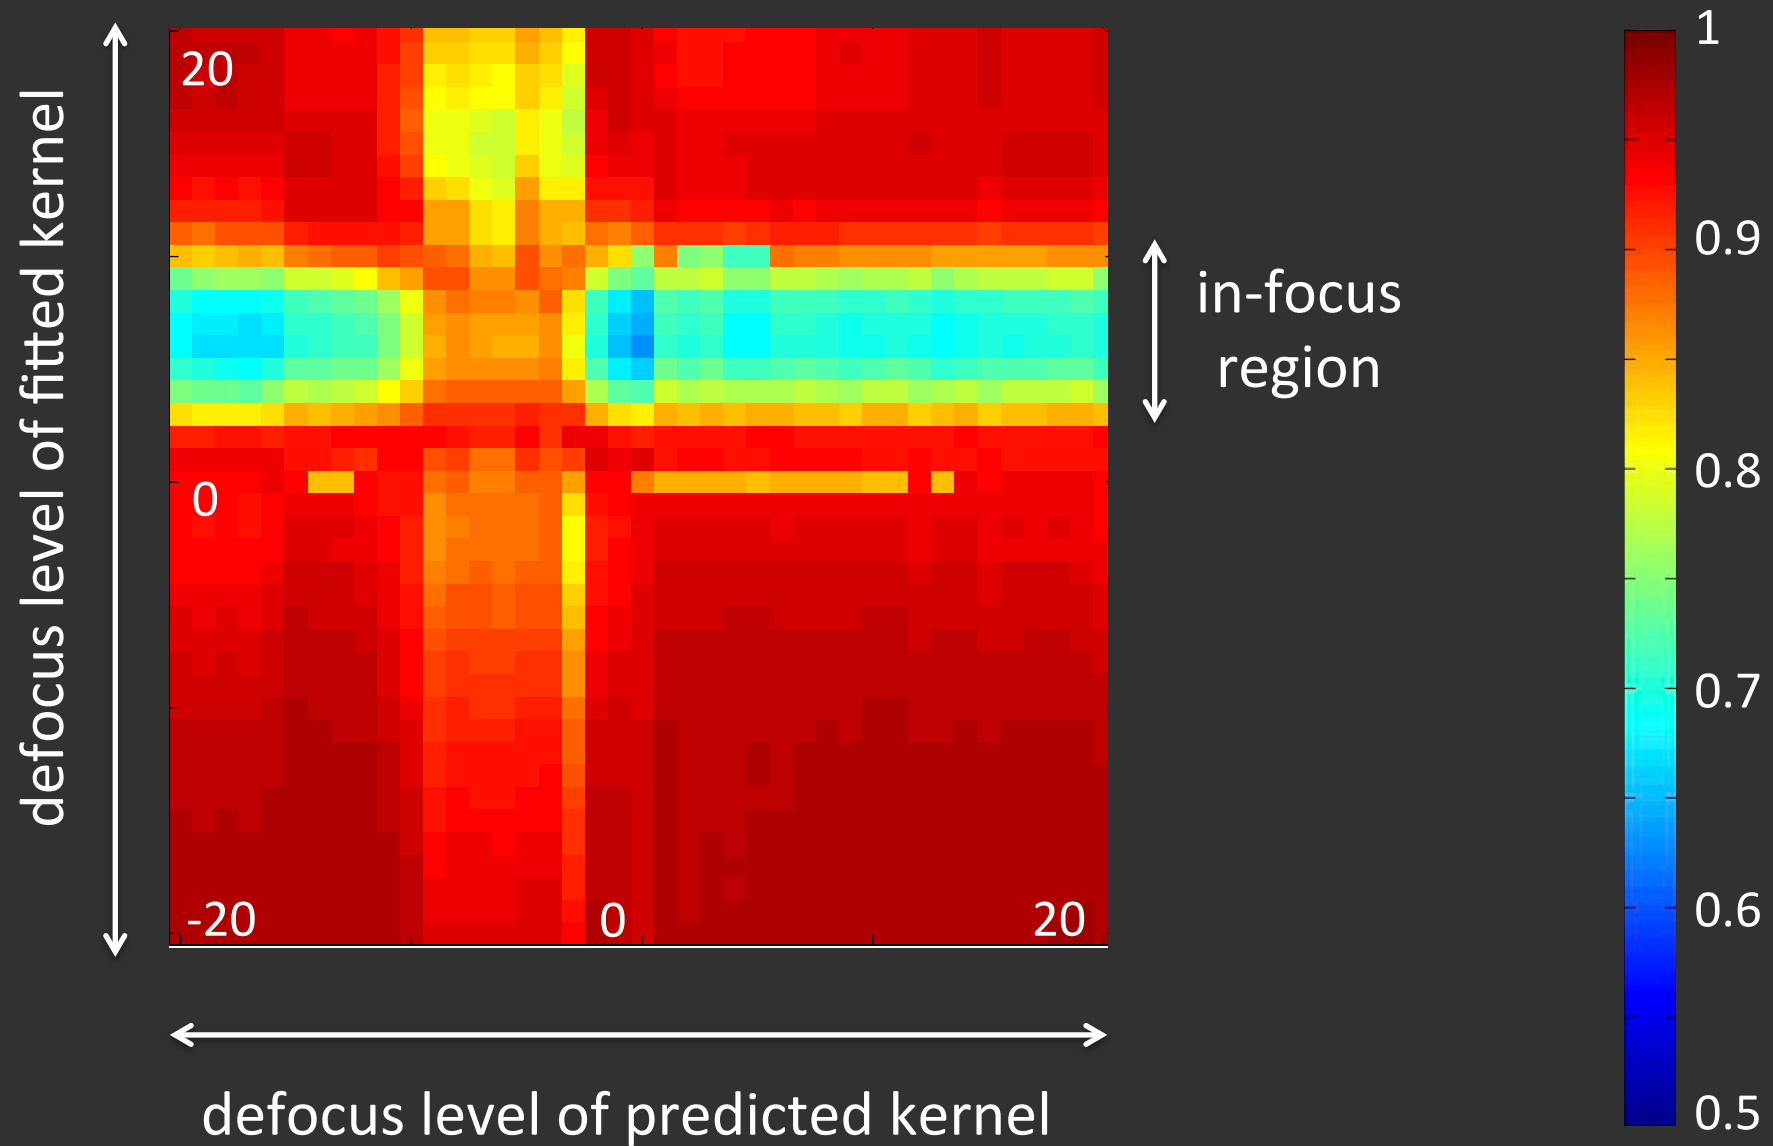


fitted depth = -20 DoF

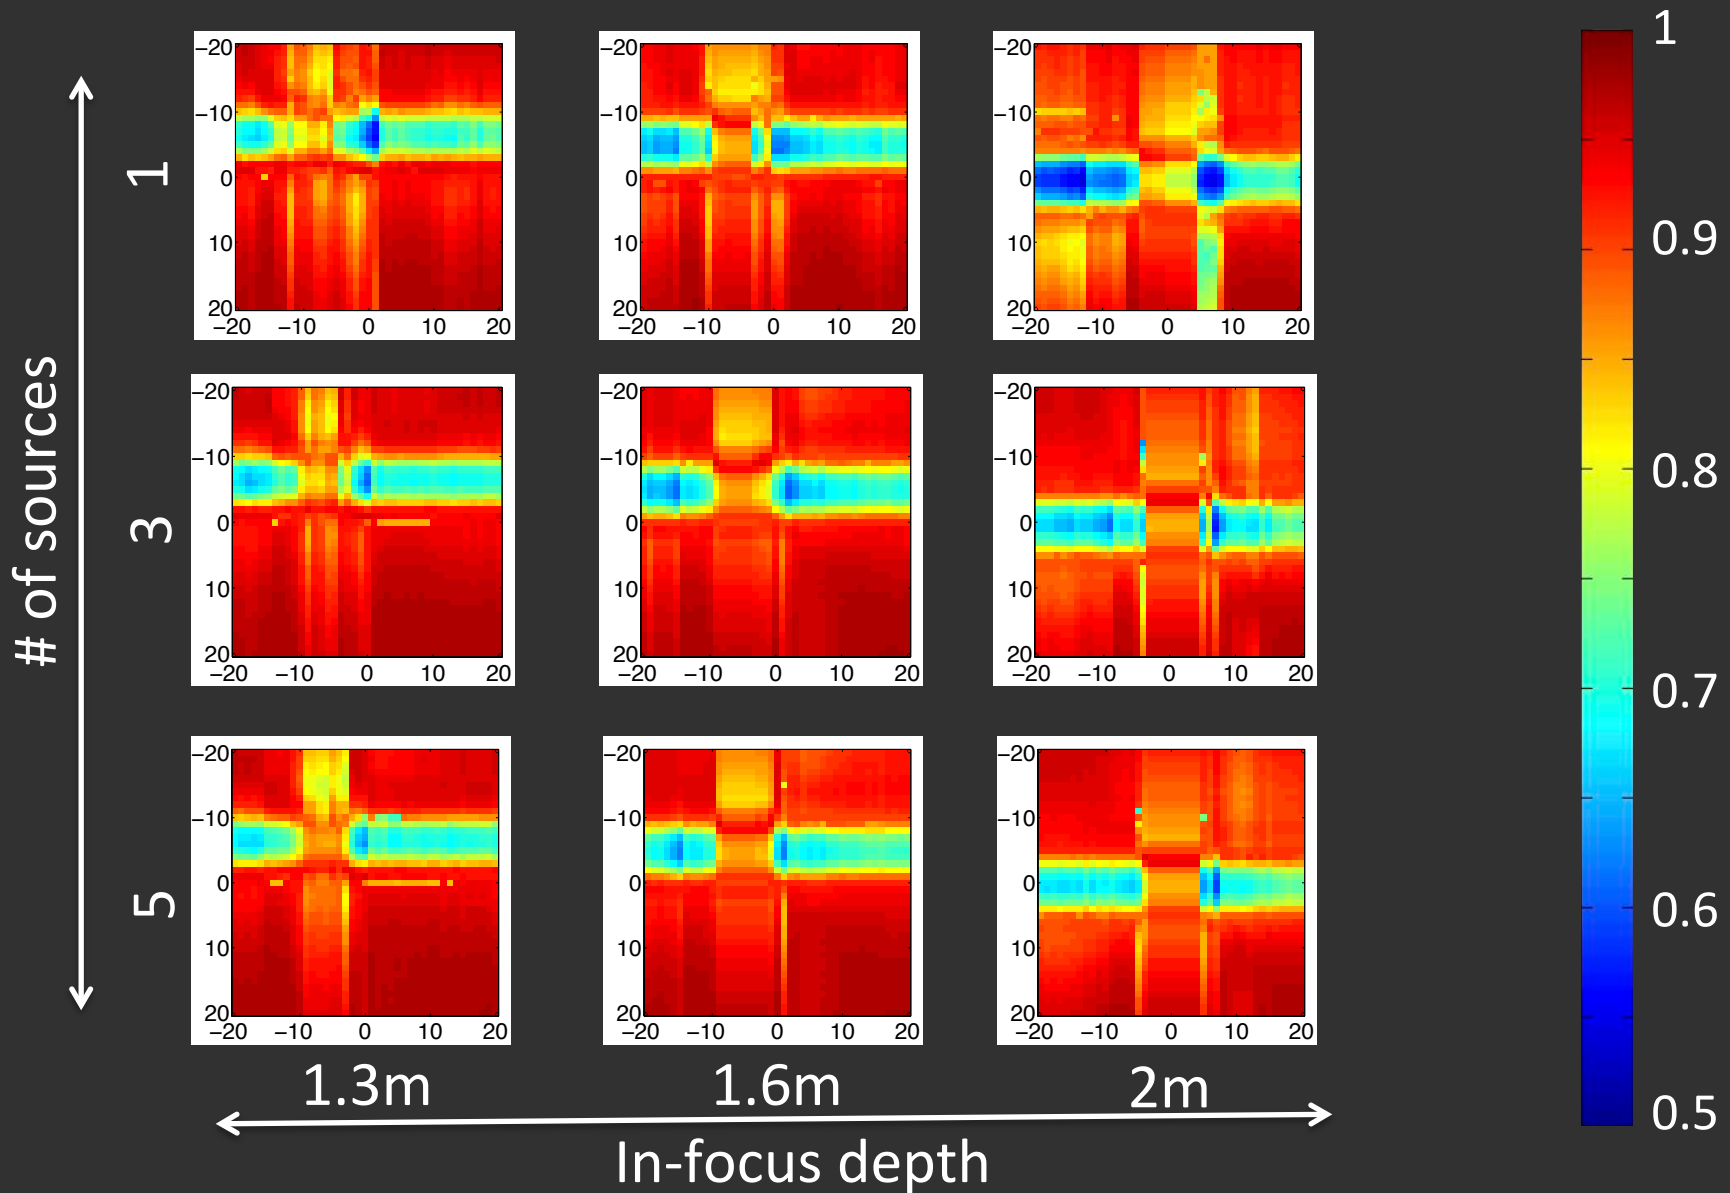
predicted depth = 20 DoF



quantitative evaluation (K=5)



Quantitative results for 3 different focus settings



concluding remarks

aberration & defocus blur give a great deal of info about the lens & scene

one-shot recovery of depth map & 5D lens PSF may be possible

modeling ray-sensor intersections easier & more general than modeling blur kernels directly

open problems

- blind estimation of blur kernels, depth & 5D PSF
- modeling mechanical vignetting
- accounting for changes in focus setting & zoom