CS 413 Introduction to Ray and Vector Graphics

Chapter 10: Thin Lens

Instructor: **Joel Castellanos**

e-mail: joel@unm.edu

Web: http://cs.unm.edu/~joel/

Farris Engineering Center: 319

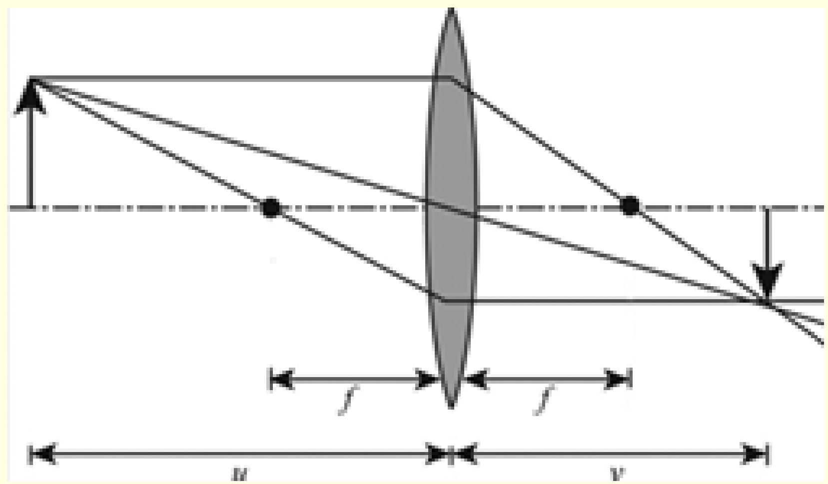

Ray Tracer V: Thin Lens

Read chapter 10 (Depth of Field), then go back and read short chapter 6 (Mapping Samples to a Disk)

Render the same scene as in version 4, and add 3 sliders:

1) Lens diameter
2) Focal plane distance
3) Number of anti-aliasing points per pixel (which some of you already have from past versions)

Physical Pinhole Camera

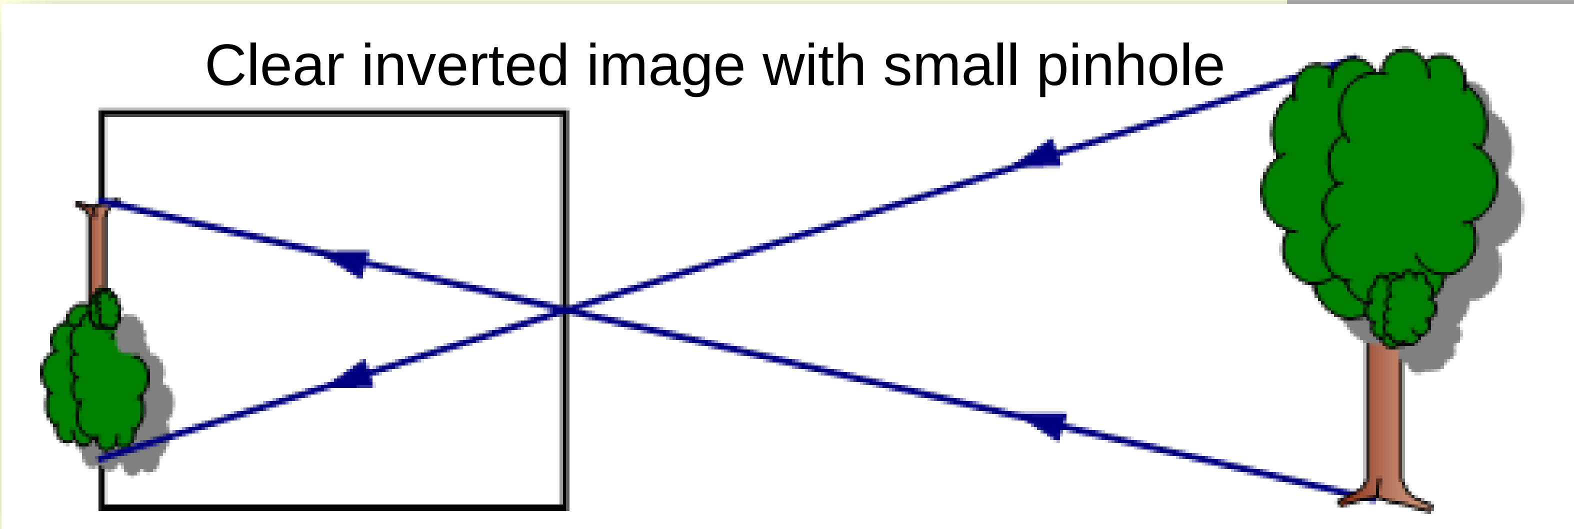

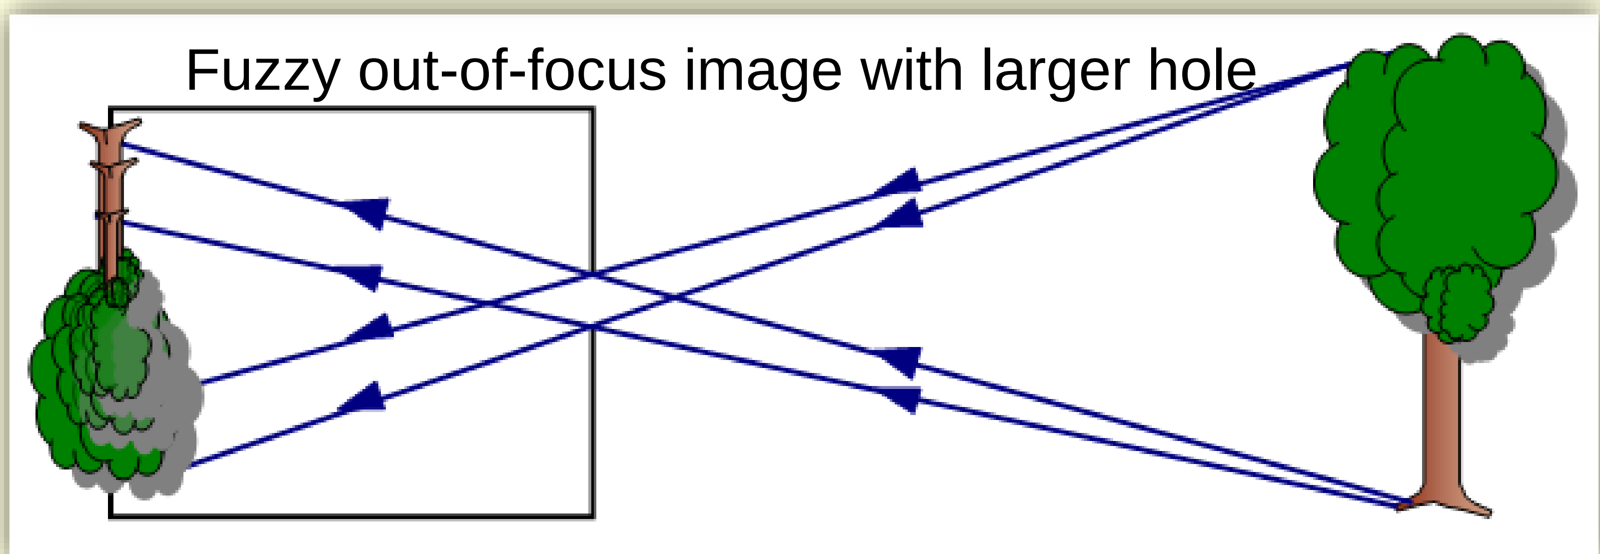

Iris and Pupil of Human Eye

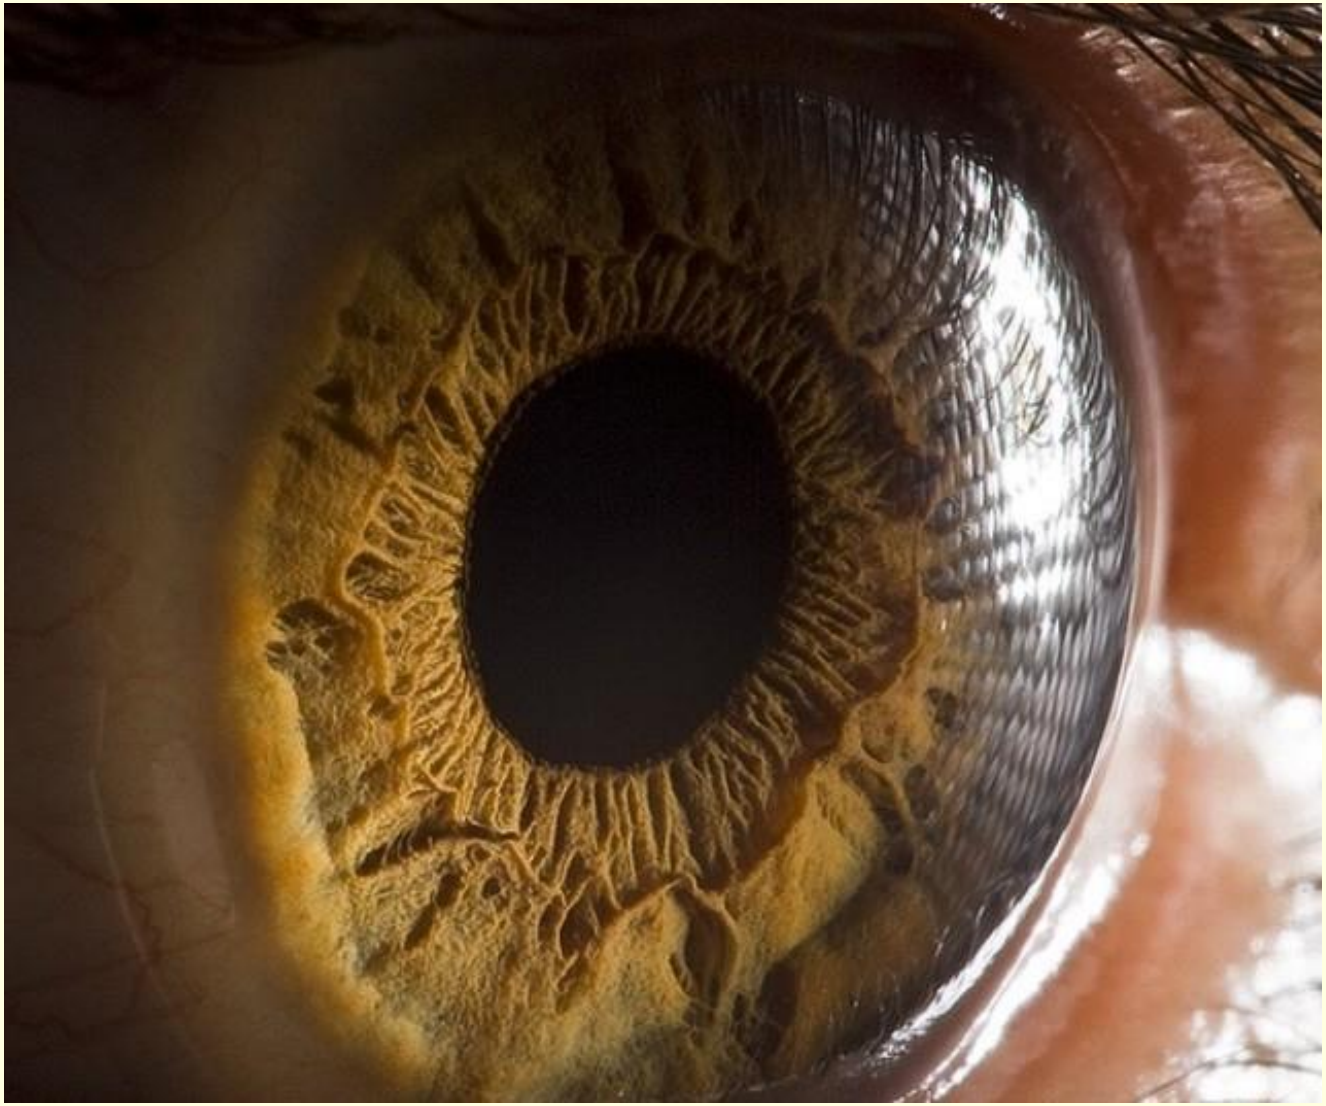

The ***iris*** is a thin, circular structure in the eye, responsible for controlling the diameter of the ***pupil*** and thus the amount of light reaching the retina.

Lens Aperture

Aperture ≠ Shutter

Aperture

Shutter

Why color reflections?

Depth of Field

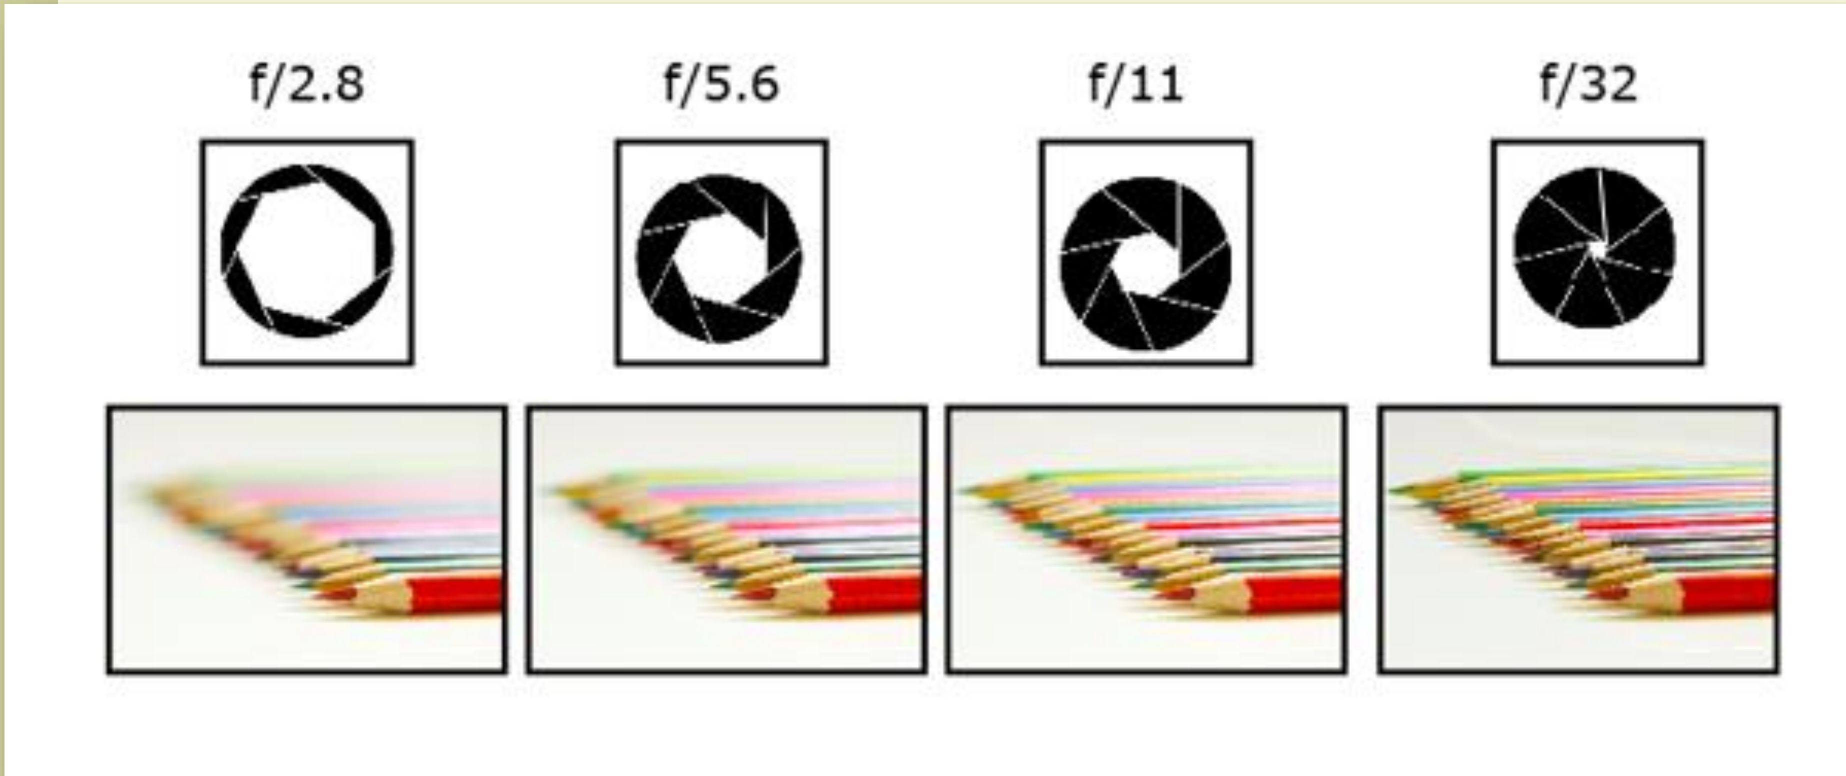

Large aperture lenses are expensive

Canon, Prime 50mm f/1.8 USM Lens:	$125.00
Canon, Prime 50 mm f/1.4 USM Lens:	$399.00
Canon, Prime 50 mm f/1.2 USM Lens:	$1,549.00

Cross section through a Thin Lens

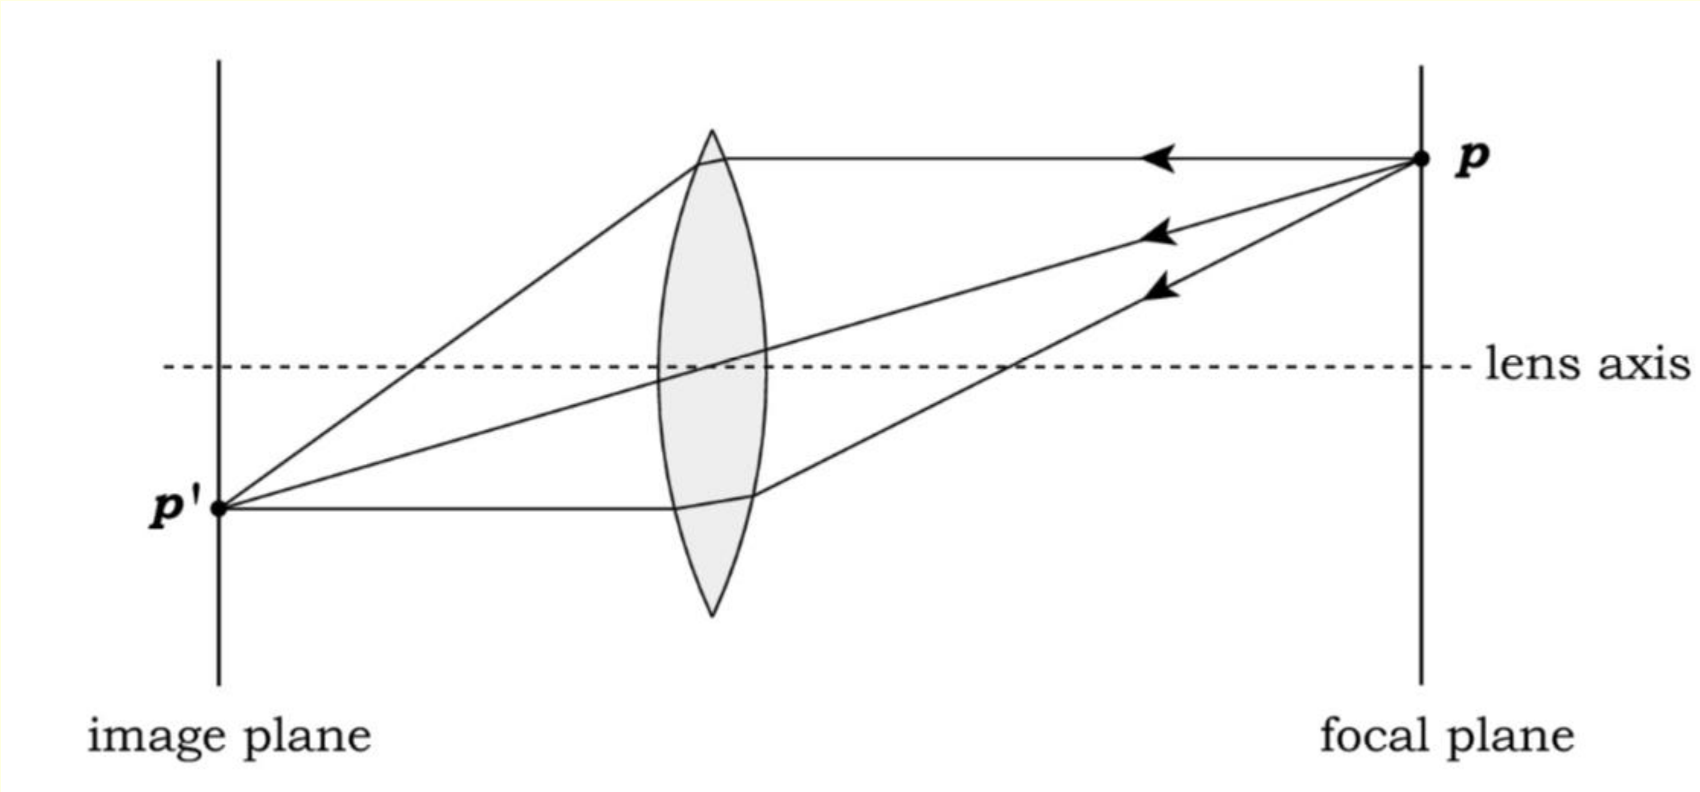

Every point, p, on the focal plane has a corresponding image p' on the image plane.

Circle of Confusion

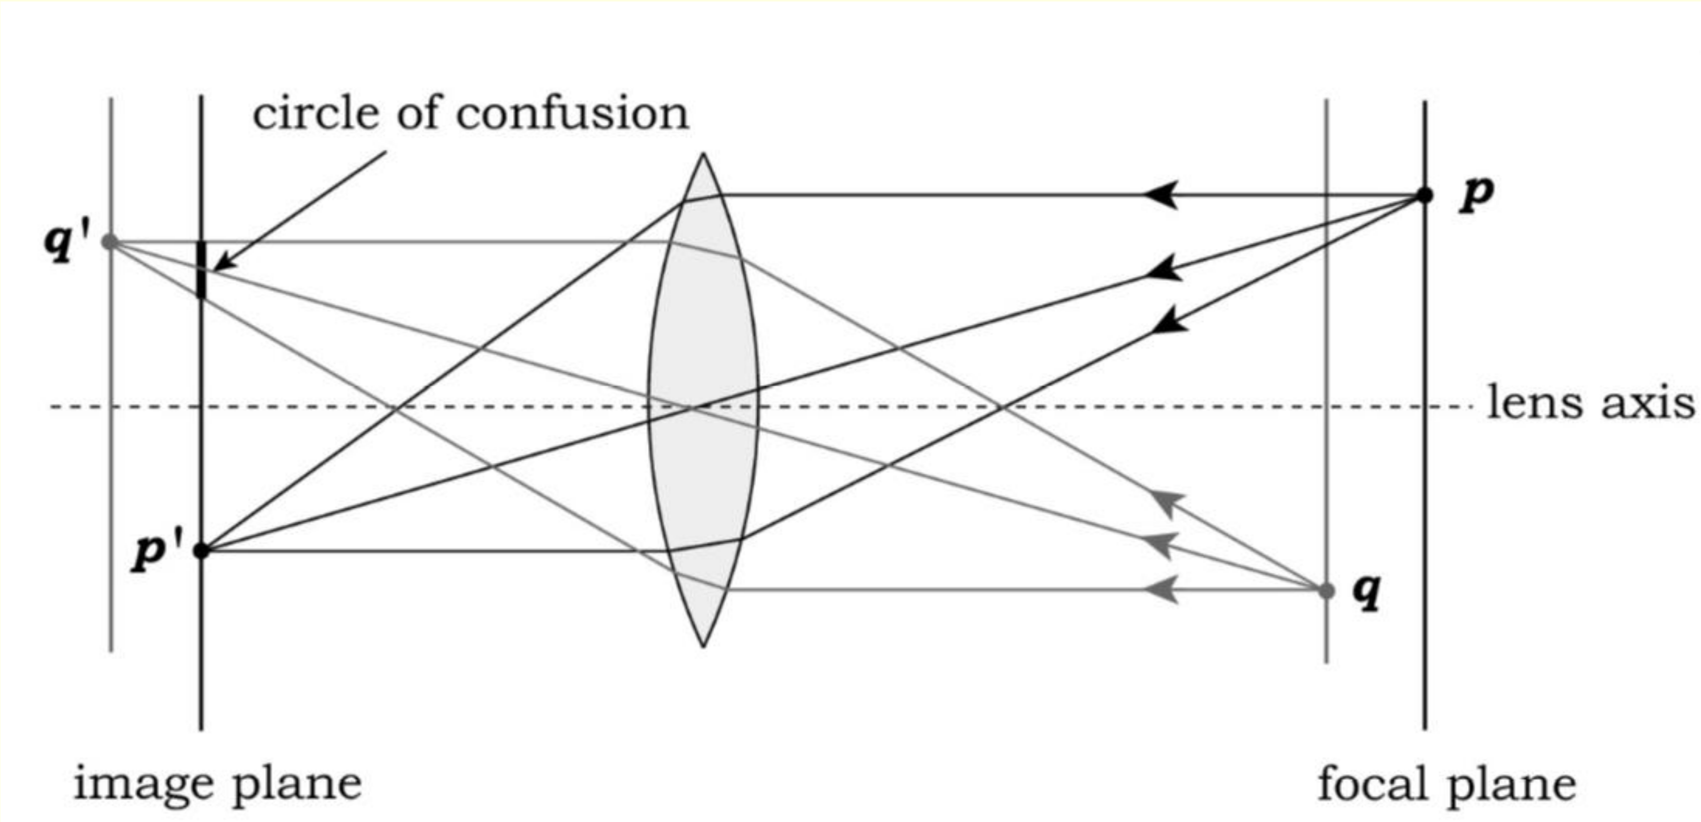

Rays starting a point q go through the image plane at different locations, with the result that q will appear out of focus.

Thin Lens Simulation Layout

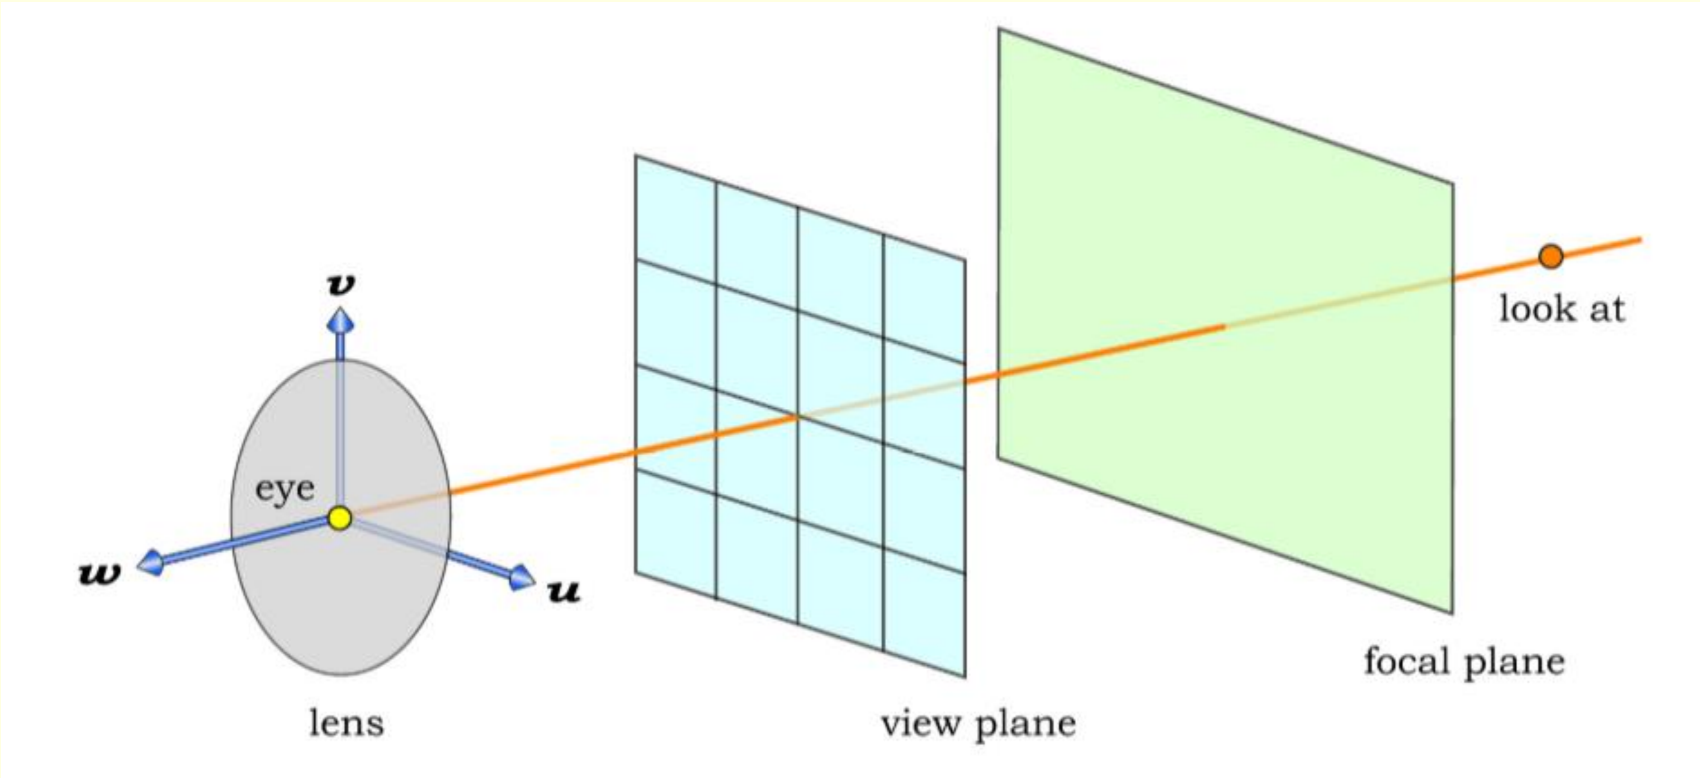

A thin lens has the view plane on the opposite side of the lens as the image plane, but our simulation will give the same result with this layout

Primary Rays Form the Image

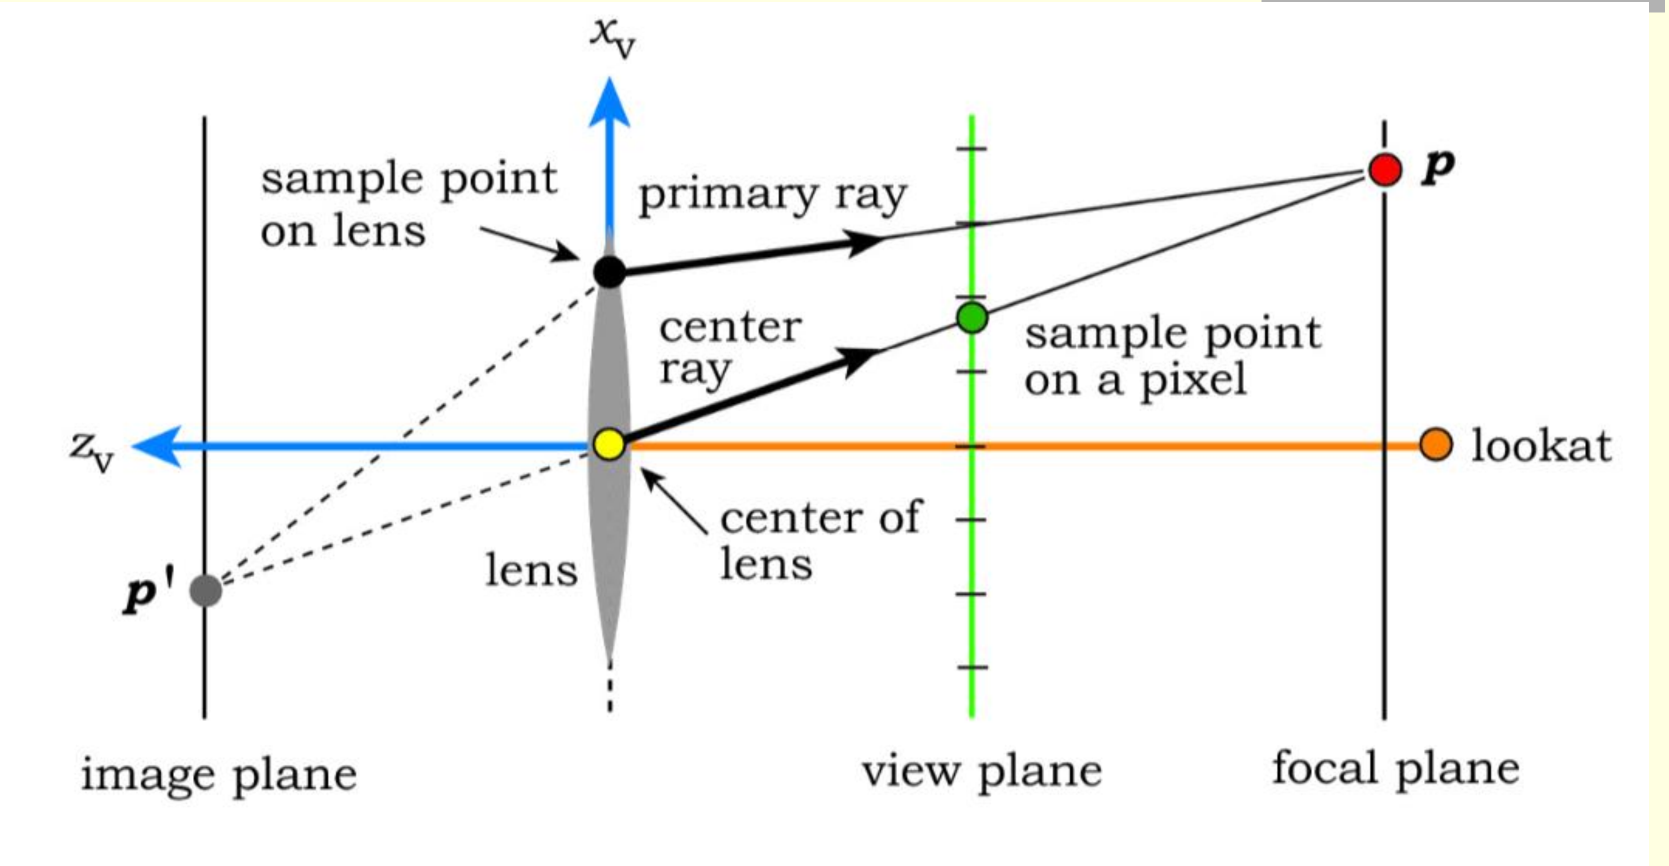

Two rays projected onto the (viewX, viewY) plane.
The center ray starts at the center of the lens.
The primary ray starts at the sample point. Both hit the focal plane at p.

Depth of Field Simulation (1 of 2)

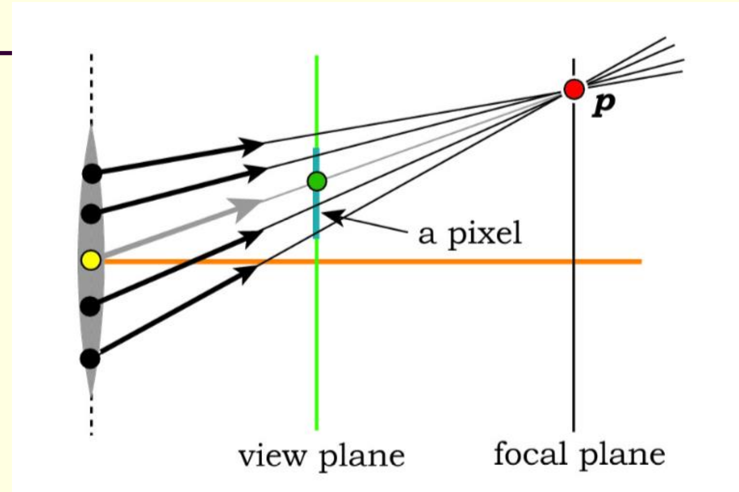

To simulate depth of field:

1) Compute the point ***p*** where the center ray hits the focal plane.
2) Use p and the sample point on the lens to compute the direction of the primary ray so that this ray also goes through ***p***.
3) Ray trace the primary ray into the scene.

Depth of Field Simulation (2 of 2)

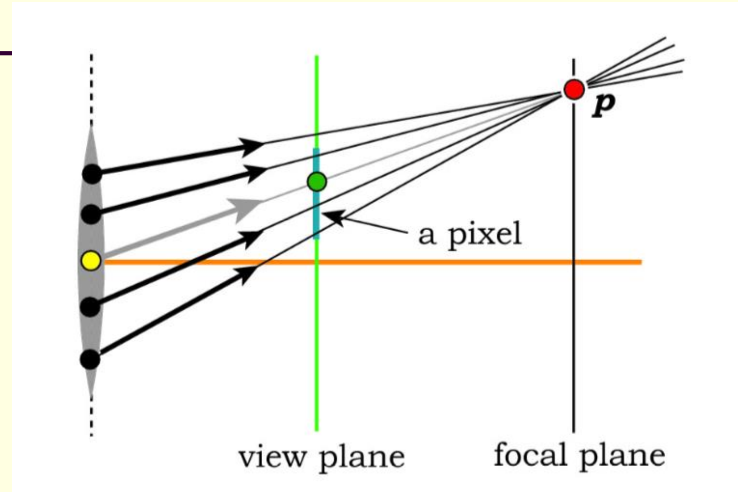

Notes:

1) The center ray doesn't contribute any color to the pixel; it is only used to find $\boldsymbol{p}$.
2) If we use a single center ray for each pixel, all primary rays for a given pixel would hit the focal plane at the same point. This would result in the focal plane being perfectly in focus but without antialiasing.

Rendering a pixel: Moving Average

With depth-of-field, the pixel you start with is not always the pixel the color contributes to.

1) Maintain 4 2D arrays of pixel: red, green, blue and count.

2) At the start of each frame, count[col,row] =0 for all col and row. The 3 color arrays can have any values.

3) Each time you calculate a new red, green, and blue value $(\acute{r}, \acute{g}, \acute{b})$ for pixel[col,row]:

$$r[col,row] = \frac{\acute{r} + count[col,row] \cdot r[col,row]}{count[col,row] + 1}$$

$$count[col,row] = count[col,row] + 1$$