


Image Processing & Projective geometry

Arunkumar Byravan


Partial slides borrowed from Jianbo Shi & Steve Seitz

Color spaces

RGB – Red, Green, Blue



HSV – Hue, Saturation, Value



Why HSV?

- HSV separates *luma*, or the image intensity, from *chroma* or the color information
 - Different shades of a color have same hue, different RGBs
- **Advantage:** Robustness to lighting conditions, shadows etc
 - Easy to use for color thresholding!
 - Fast conversion from RGB to HSV (*Python: colorsys*)
- Other relevant color spaces: YCbCr, Lab

HSV example

- GIMP



MPC – Last lecture

- Repeat:
 - Pick a set of controls (Ex: linear velocity, steering angle)
 - Simulate/Rollout using internal model (for time T)
 - Compute error (Ex: distance from desired path, desired fixed point etc.)
 - Choose controls that minimize error
 - Execute control for $T' \ll T$
- Key questions:
 - How to generate rollouts?
 - How to measure error?

MPC - Racecar

- Rollouts – Image templates
 - Pixel tracks of potential paths taken by the car
 - How to generate them?
- Errors – Distance from track in image
 - How to measure error?
 - Other error metrics:
 - Distance to target point
 - Parameterized error (line, spline etc)


How to generate templates?

- Heuristic approach:
 - Generate arcs of varying curvature
 - Associate with a control (linearly interpolate controls & "invent" a mapping)

How to generate samples?

- Geometric approach – use motion model
- Generate rollouts from kinematic model based on controls
 - Linearly interpolate controls, rollout a trajectory for fixed horizon “T”
- “Project” rollouts onto image to generate templates
 - Imagine how each rollout would look like when seen from the camera

How to generate templates?




Rollouts from the kinematic motion model, points in
frame of car, $z = 0$


How do we generate
image templates from
these rollouts?

- Projective geometry!

Camera Extrinsics



Origin frame for rollouts




Red is camera frame, white is rollout origin frame

Camera Extrinsics

- We have 3D points in robot frame (white)
- Transform points to camera frame through extrinsics:

$$\text{Camera frame} \rightarrow \begin{pmatrix} x' \\ y' \\ z' \\ 1 \end{pmatrix} = \begin{bmatrix} \mathbf{R} & \mathbf{t} \\ \mathbf{0} & 1 \end{bmatrix} \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix} \leftarrow \text{Robot frame}$$

- \mathbf{R} = 3x3 rotation matrix
- \mathbf{t} = 3x1 translation vector
- For racecar, extrinsics can be measured (and is constant)




Red is camera frame, white is rollout origin frame


Camera Extrinsics

- Extrinsics allow us to transform 3D points to camera frame of reference
- We need to figure out how to get the image of these points as seen by the camera

Image formation process




Pinhole camera model



- Add a barrier to block off most of the rays
 - This reduces blurring
 - The opening known as the **aperture**
 - How does this transform the image?

Pinhole camera model



- Pinhole model:
 - Captures **pencil of rays** – all rays through a single point
 - The point is called **Center of Projection (COP)**
 - The image is formed on the **Image Plane**
 - **Effective focal length f** is distance from COP to Image Plane

Homemade pinhole camera






Why so
blurry?




<http://www.debevec.org/Pinhole/>

Camera with Lens




Digital camera





a b

FIGURE 2.17 (a) Continuous image projected onto a sensor array. (b) Result of image sampling and quantization.




CMOS sensor

Bayer grid




Estimate RGB
at 'G' cels from
neighboring
values

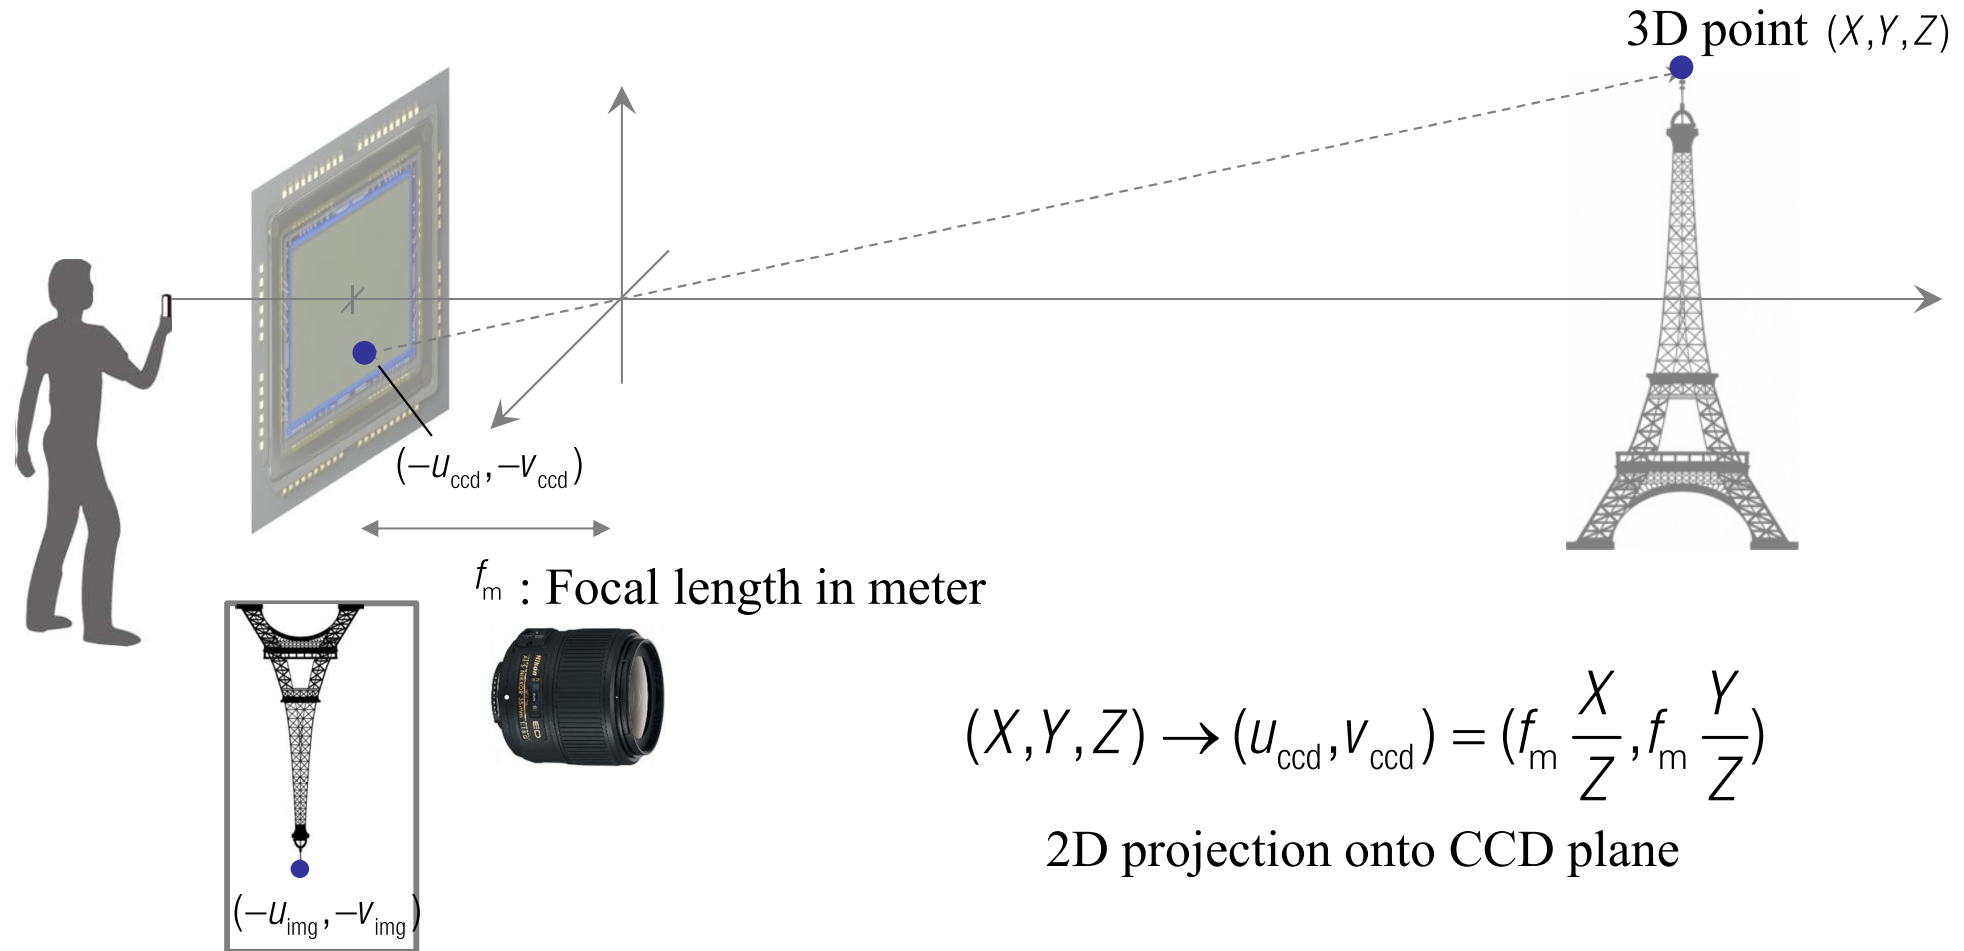
[http://www.cooldictionary.com/
words/Bayer-filter.wikipedia](http://www.cooldictionary.com/words/Bayer-filter.wikipedia)




Projection from 3D to 2D




3D point projection (Metric space)



3D point projection (Metric space)



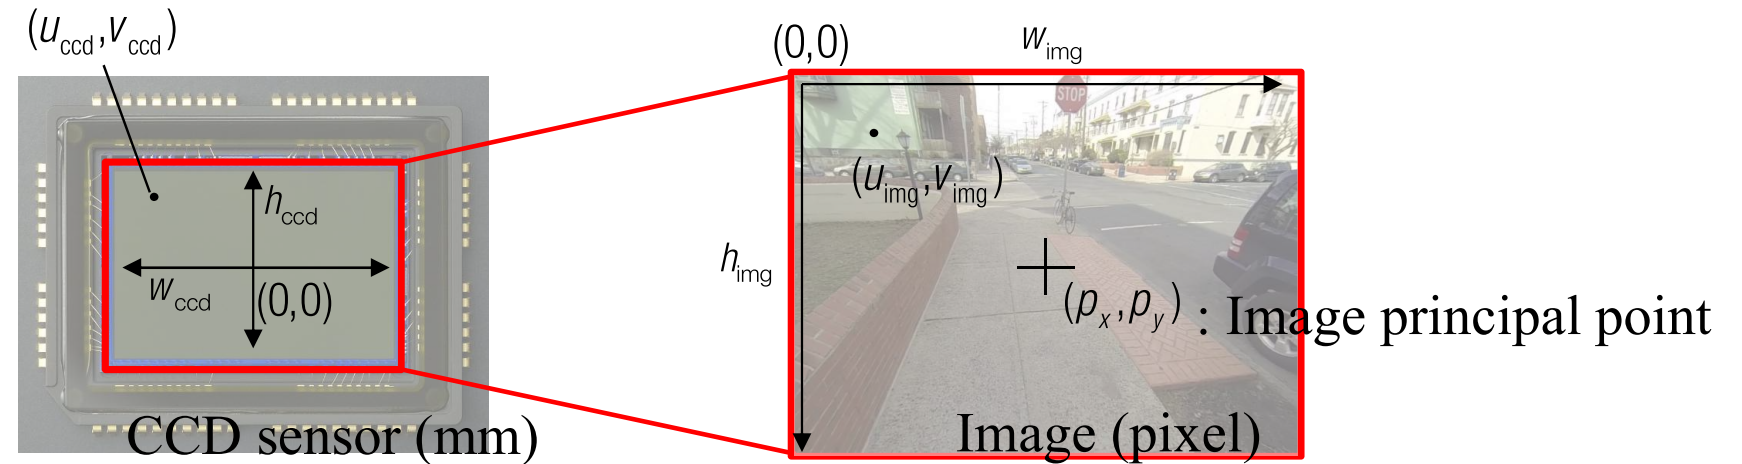
3D point projection (Metric space)



$$(X, Y, Z) \rightarrow (u_{\text{ccd}}, v_{\text{ccd}}) = \left(f_m \frac{X}{Z}, f_m \frac{Y}{Z} \right)$$

2D projection onto CCD plane

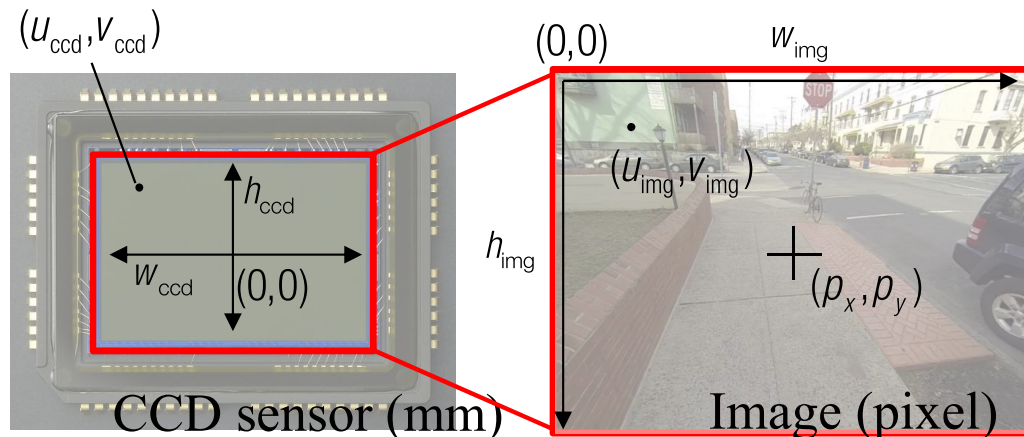

3D point projection (Pixel space)



$$\frac{u_{\text{ccd}}}{w_{\text{ccd}}} = \frac{u_{\text{img}} - p_x}{w_{\text{img}}} \quad \frac{v_{\text{ccd}}}{h_{\text{ccd}}} = \frac{v_{\text{img}} - p_y}{h_{\text{img}}}$$

$$u_{\text{img}} = u_{\text{ccd}} \frac{w_{\text{img}}}{w_{\text{ccd}}} + p_x \quad v_{\text{img}} = v_{\text{ccd}} \frac{h_{\text{img}}}{h_{\text{ccd}}} + p_y$$

3D point projection (Pixel space)



$$(X, Y, Z) \rightarrow (u_{\text{ccd}}, v_{\text{ccd}}) = \left(f_m \frac{X}{Z}, f_m \frac{Y}{Z} \right)$$

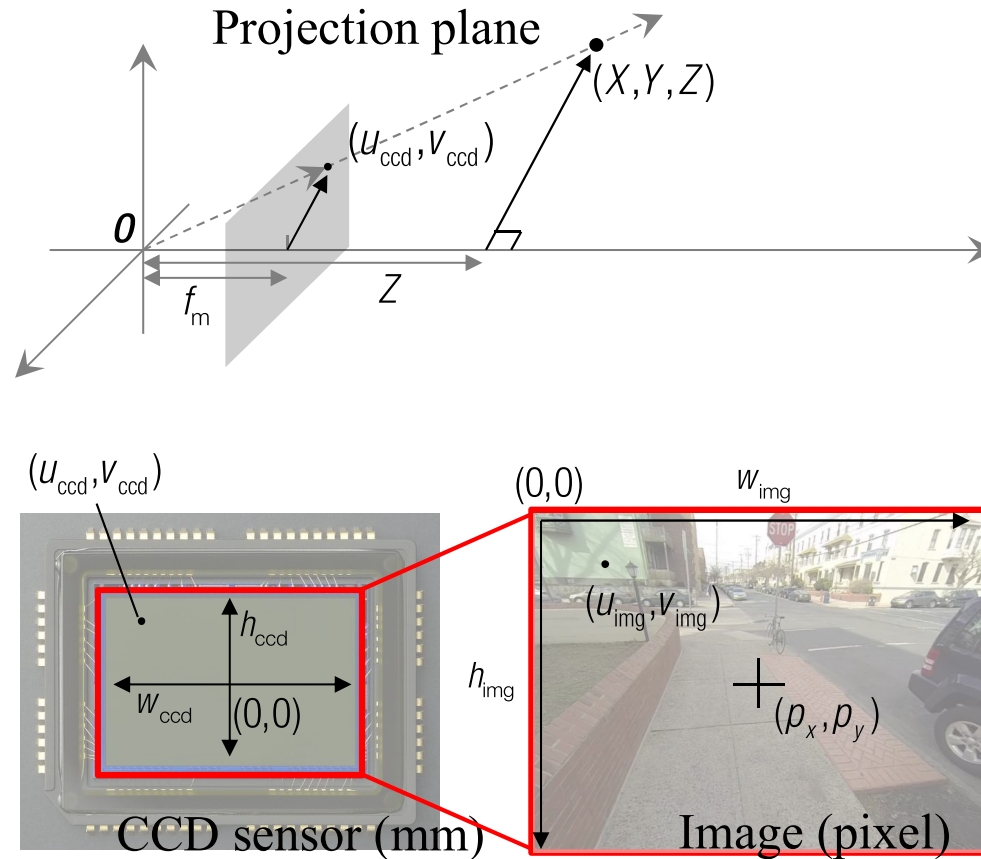
$$u_{\text{img}} = u_{\text{ccd}} \frac{w_{\text{img}}}{w_{\text{ccd}}} + p_x = f_m \frac{w_{\text{img}}}{w_{\text{ccd}}} \frac{X}{Z} + p_x$$

Focal length in pixel

$$v_{\text{img}} = v_{\text{ccd}} \frac{h_{\text{img}}}{h_{\text{ccd}}} + p_y = f_m \frac{h_{\text{img}}}{h_{\text{ccd}}} \frac{Y}{Z} + p_y$$

Focal length in pixel

3D point projection (Pixel space)



$$(X, Y, Z) \rightarrow (u_{ccd}, v_{ccd}) = \left(f_m \frac{X}{Z}, f_m \frac{Y}{Z} \right)$$


$$u_{img} = u_{ccd} \frac{w_{img}}{w_{ccd}} + p_x = \underbrace{f_m \frac{w_{img}}{f_x w_{ccd}}}_{\text{Focal length in pixel}} \frac{X}{Z} + p_x$$

Focal length in pixel

$$v_{img} = v_{ccd} \frac{h_{img}}{h_{ccd}} + p_y = \underbrace{f_m \frac{h_{img}}{f_y h_{ccd}}}_{\text{Focal length in pixel}} \frac{Y}{Z} + p_y$$


Focal length in pixel

3D point projection (Pixel space)




$$(X, Y, Z) \rightarrow (u_{img}, v_{img}) = \left(f_m \frac{w_{img}}{w_{ccd}} \frac{X}{Z}, f_m \frac{h_{img}}{h_{ccd}} \frac{Y}{Z} \right)$$

Homogeneous coordinates



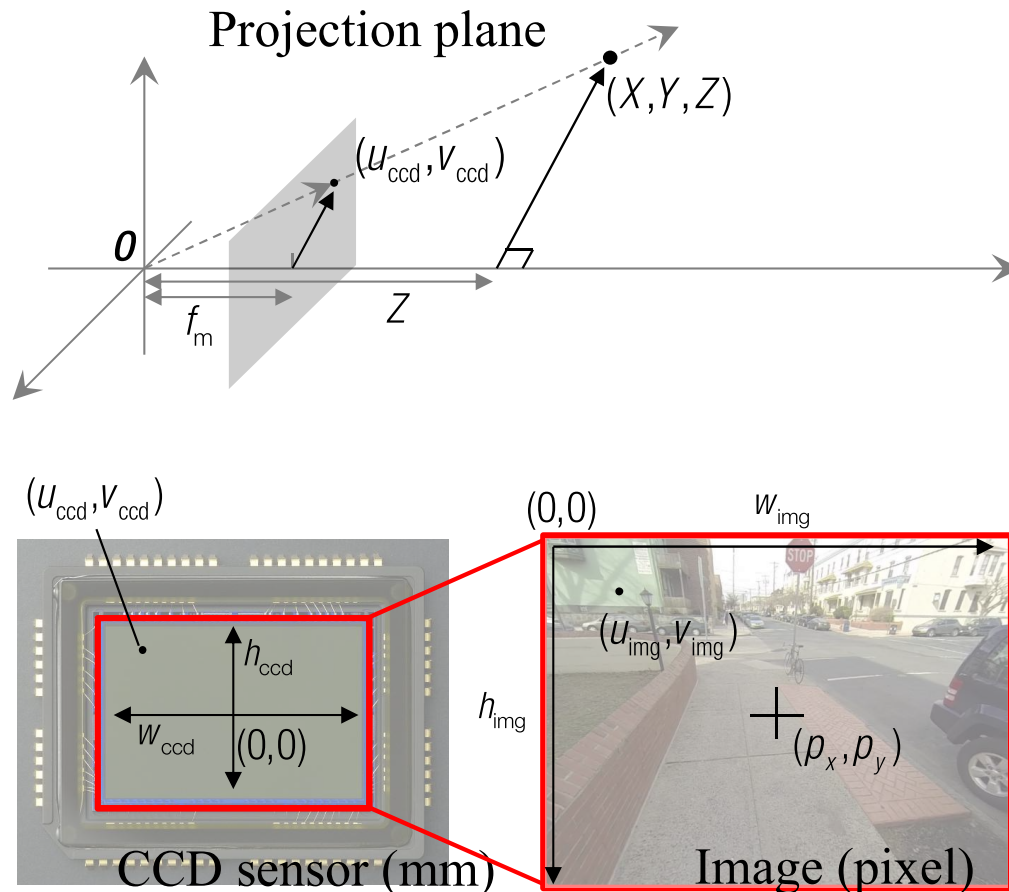
$(x, y) \rightarrow (x, y, 1)$: A point in Euclidean space (\mathbb{R}^2) can be represented by
a homogeneous representation in Projective space (\mathcal{P}^2) (3 numbers).
 $= f(x, y, 1)$
 $= \lambda(x, y, 1)$

Homogeneous coordinates



$\lambda(x, y, 1)$ $= (X, Y, Z)$: 3D point lies in the 3D ray passing 2D image point.
Homogeneous coordinate

3D point projection (Pixel space)




$$(X, Y, Z) \rightarrow (u_{\text{ccd}}, v_{\text{ccd}}) = \left(f_m \frac{X}{Z}, f_m \frac{Y}{Z} \right)$$

$$u_{\text{img}} = f_x \frac{X}{Z} + p_x \quad v_{\text{img}} = f_y \frac{Y}{Z} + p_y$$


$$\lambda \begin{bmatrix} u_{\text{img}} \\ v_{\text{img}} \\ 1 \end{bmatrix} = \begin{bmatrix} f_x & p_x \\ & f_y & p_y \\ & & 1 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}$$

Homogeneous representation

Camera intrinsics



$$\lambda \begin{bmatrix} u_{\text{img}} \\ v_{\text{img}} \\ 1 \end{bmatrix} = \underbrace{\begin{bmatrix} f_x & p_x \\ t_y & p_y \\ & 1 \end{bmatrix} \mathbf{K}}_{\text{Camera intrinsic parameter}} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}$$



Camera intrinsic parameter
: metric space to pixel space

Putting it all together – Generating templates

- Procedure:
 - Generate rollouts based on kinematic car model (robot frame)
 - Transform points to camera frame based on **camera extrinsics**
 - Project points to pixel space using **camera intrinsics**

Diagram illustrating the transformation from the Robot frame to the Camera frame using extrinsics, and then projecting to pixel space using intrinsics.

Extrinsics: The transformation from the Robot frame to the Camera frame is defined by the extrinsic matrix $\begin{bmatrix} R & t \\ \mathbf{0} & 1 \end{bmatrix}$. The equation is:

$$\begin{pmatrix} x' \\ y' \\ z' \\ 1 \end{pmatrix} = \begin{bmatrix} R & t \\ \mathbf{0} & 1 \end{bmatrix} \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix}$$

Labels: Camera frame (pointing to the left vector), Robot frame (pointing to the right vector).


Need to be measured for racecar (Approx values in /tf)

Intrinsics: The projection from the normalized camera coordinates to pixel space is defined by the intrinsic matrix K . The equation is:


$$\begin{pmatrix} u \\ v \\ 1 \end{pmatrix} = K \begin{pmatrix} \frac{x'}{z'} \\ \frac{y'}{z'} \\ 1 \end{pmatrix}$$

Intrinsics fixed for a camera (for racecar: /camera/color/camera_info)

Rollout templates



Rollouts from kinematic car model



Projected rollouts

Measuring error (for MPC)


- Template matching using convolution
 - Find template that best matches masked track, choose it



- Issues?


Measuring error (Set-point error)

- Choose set point in image (similar to PID)
 - Find template that gets you closest to set point, choose it




Measuring error (Set-point + Direction error)

- Choose set point in image along with heading (based on track)
 - Find template that gets you closest to set point while oriented correctly
 - Keep track of heading in templates



Measuring error (3D error)

- Instead of generating pixelized templates, project masked track (or set point) back to 3D
- How?
 - Each pixel corresponds to ray in 3D
 - We know that all pixels on track lie on ground plane (known)
 - Solve for ray-plane intersection
- Advantage: Reason in 3D!






Pixel space		Metric space
$\lambda \begin{bmatrix} u_{\text{img}} \\ v_{\text{img}} \\ 1 \end{bmatrix} = \begin{bmatrix} f_x & p_x \\ & p_y \\ & 1 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}$	\mathbf{K}	
$\lambda \mathbf{K}^{-1} \begin{bmatrix} u_{\text{img}} \\ v_{\text{img}} \\ 1 \end{bmatrix} = \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}$		$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix}$
		3D ray

The 3D point must lie in the 3D ray passing through the origin and 2D image point.



Measuring error (fancy error metric)

- Fit a line or curve to pixel/3D track points & your rollouts
 - Compare the errors in parametric space (line / curve co-efficients)

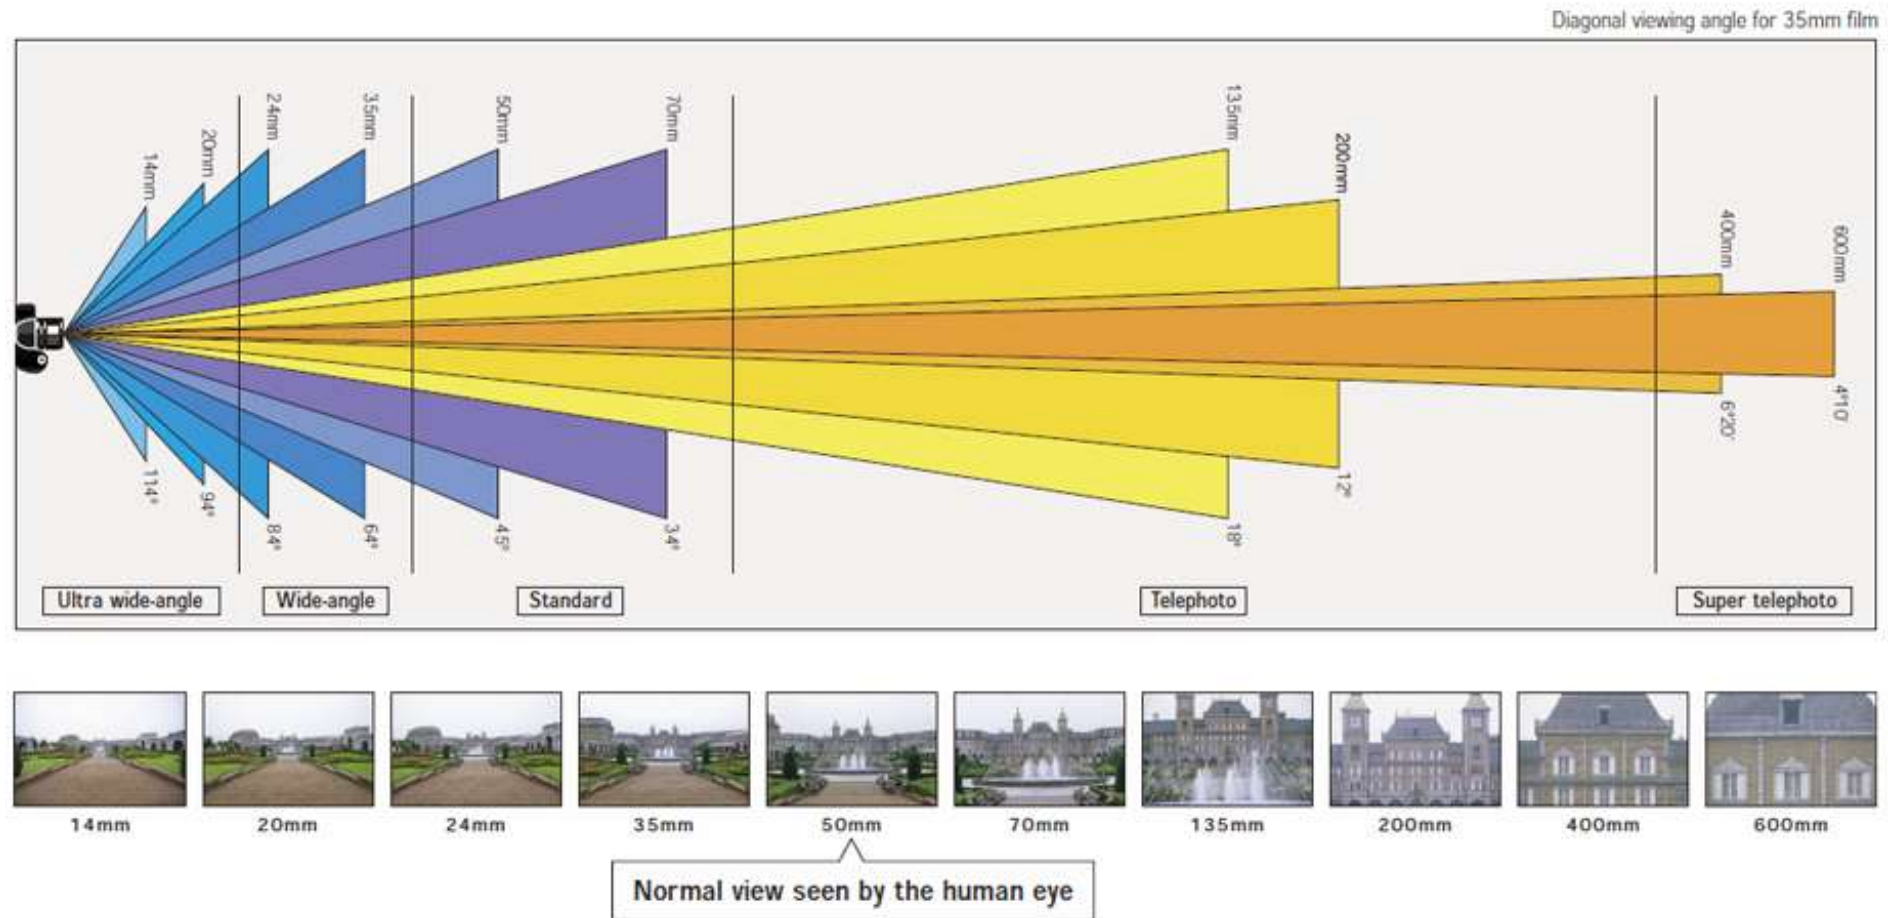
Focal length




Focal length



Focal length




Focal length




Dolly zoom

- <https://www.youtube.com/watch?v=NB4bikrNzMk>


Perspective cues



Perspective cues



Perspective cues




Lens distortion (Fisheye lens)




Multiple models to capture distortion, commonly used is Plumb Bob model

Slide from Jianbo Shi


Lens distortion



Barrel distortion




Pincushion distortion



Moustache distortion

Modeled as a function that changes pixel (u,v) after intrinsics +
extrinsics based projection

Camera Calibration



Images courtesy Jean-Yves Bouguet, Intel Corp.

- Compute camera intrinsic parameters & distortion
- Compute extrinsics between multiple views/different cameras
- **Key idea:** Use a known object of fixed size & match it across multiple scenes -> provides enough constraints to solve for camera parameters
- Good code available online!
 - Intel's OpenCV library: <http://www.intel.com/research/mrl/research/opencv/>
 - Matlab version by Jean-Yves Bouguet: http://www.vision.caltech.edu/bouguetj/calib_doc/index.html
 - Zhengyou Zhang's web site: <http://research.microsoft.com/~zhang/Calib/>