



COMP 141: Probabilistic Robotics for Human-Robot Interaction

Instructor: Jivko Sinapov
www.cs.tufts.edu/~jsinapov

Today

- Introduction to Computer Vision

Reading Assignment

- You should be reading Chapter 4 (focus on 4.1 and 4.2)

Project Timeline

Homework 2

Research Papers Presentations

- Each team is required to present one article together
- Presentations should be ~ 20 minutes
- Sign-up sheet to be posted on Canvas

What is an image?

A grayscale image

Index	0	1	2	3	4	5	6	7	8	9
0	1	2	3	4	5	6	7	8	9	10
1	11	12	13	14	15	16	17	18	19	20
2	21	22	23	24	25	26	27	28	29	30
3	31	32	33	34	35	36	37	38	39	40
4	41	42	43	44	45	46	47	48	49	50
5	51	52	53	54	55	56	57	58	59	60
6	61	62	63	64	65	66	67	68	69	70
7	71	72	73	74	75	76	77	78	79	80
8	81	82	83	84	85	86	87	88	89	90
9	91	92	93	94	95	96	97	98	99	100

An RGB Image

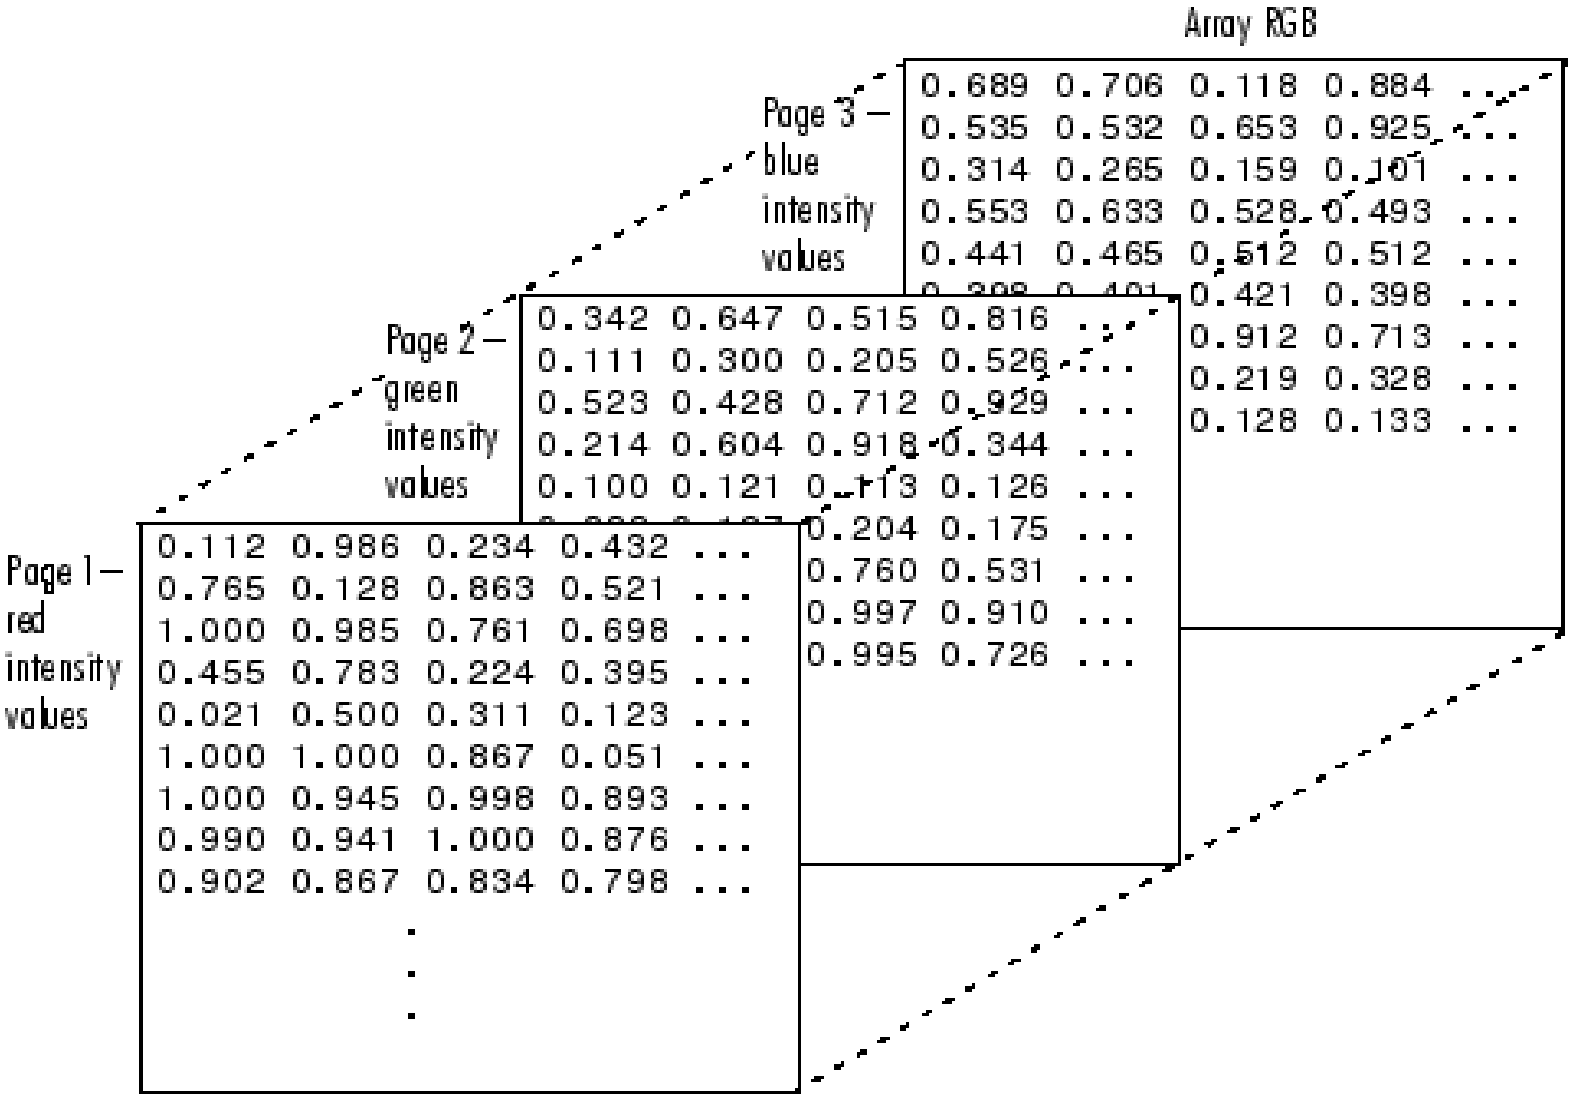
Array RGB
Page 3 — blue intensity values
0.689 0.706 0.118 0.884 ...
0.535 0.532 0.653 0.925 ...
0.314 0.265 0.159 0.101 ...
0.553 0.633 0.528 0.493 ...
0.441 0.465 0.512 0.512 ...
0.421 0.398 ...
0.912 0.713 ...
0.219 0.328 ...
0.128 0.133 ...
Page 2 — green intensity values
0.342 0.647 0.515 0.816 ...
0.111 0.300 0.205 0.526 ...
0.523 0.428 0.712 0.929 ...
0.214 0.604 0.918 0.344 ...
0.100 0.121 0.113 0.126 ...
0.204 0.175 ...
0.760 0.531 ...
0.997 0.910 ...
0.995 0.726 ...
Page 1 — red intensity values
0.112 0.986 0.234 0.432 ...
0.765 0.128 0.863 0.521 ...
1.000 0.985 0.761 0.698 ...
0.455 0.783 0.224 0.395 ...
0.021 0.500 0.311 0.123 ...
1.000 1.000 0.867 0.051 ...
1.000 0.945 0.998 0.893 ...
0.990 0.941 1.000 0.876 ...
0.902 0.867 0.834 0.798 ...

Intensity Levels

- 2
- 32
- 64
- 128
- **256 (8 bits)**
- 512
- …
- 4096 (12 bits)

How did computer vision start?

“In 1966, Marvin Minsky at MIT asked his undergraduate student Gerald Jay Sussman to “spend the summer linking a camera to a computer and getting the computer to describe what it saw”. We now know that the problem is slightly more difficult than that!”

Human vs Computer Vision

What we see

0	3	2	5	4	7	6	9	8
3	0	1	2	3	4	5	6	7
2	1	0	3	2	5	4	7	6
5	2	3	0	1	2	3	4	5
4	3	2	1	0	3	2	5	4
7	4	5	2	3	0	1	2	3
6	5	4	3	2	1	0	3	2
9	6	7	4	5	2	3	0	1
8	7	6	5	4	3	2	1	0

What a computer sees

Image Plane v.s. Image Array

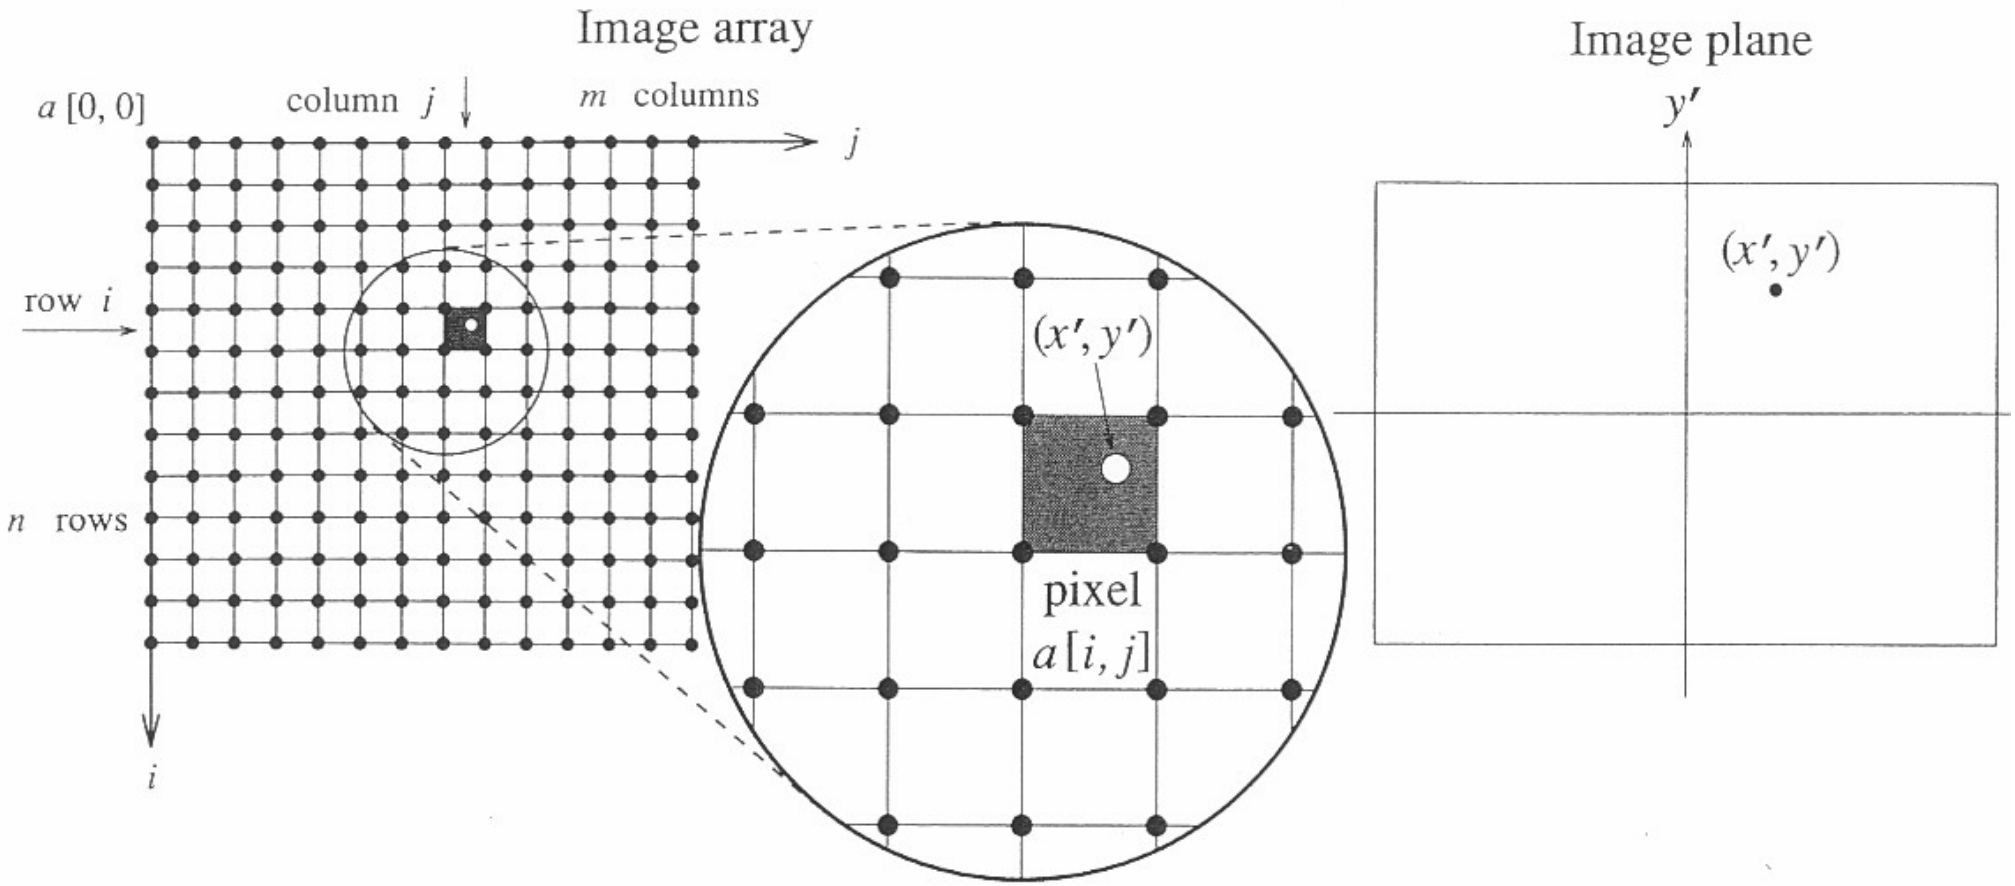

[Jain, Kasturi, and Schunck (1995). Machine Vision, Ch. 1]

Point Operations

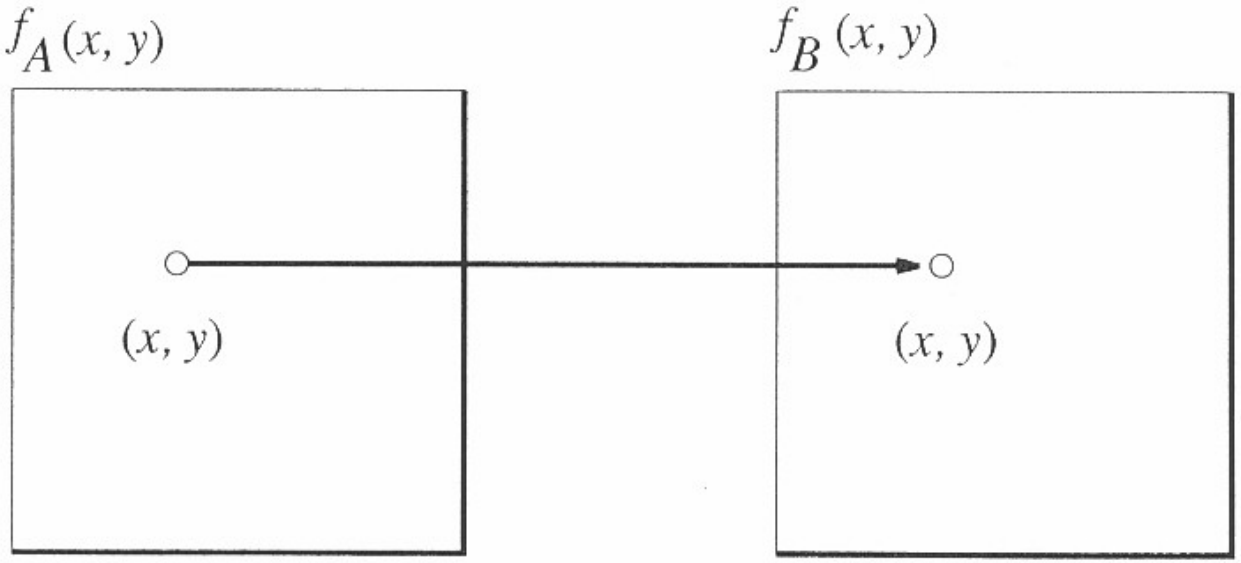

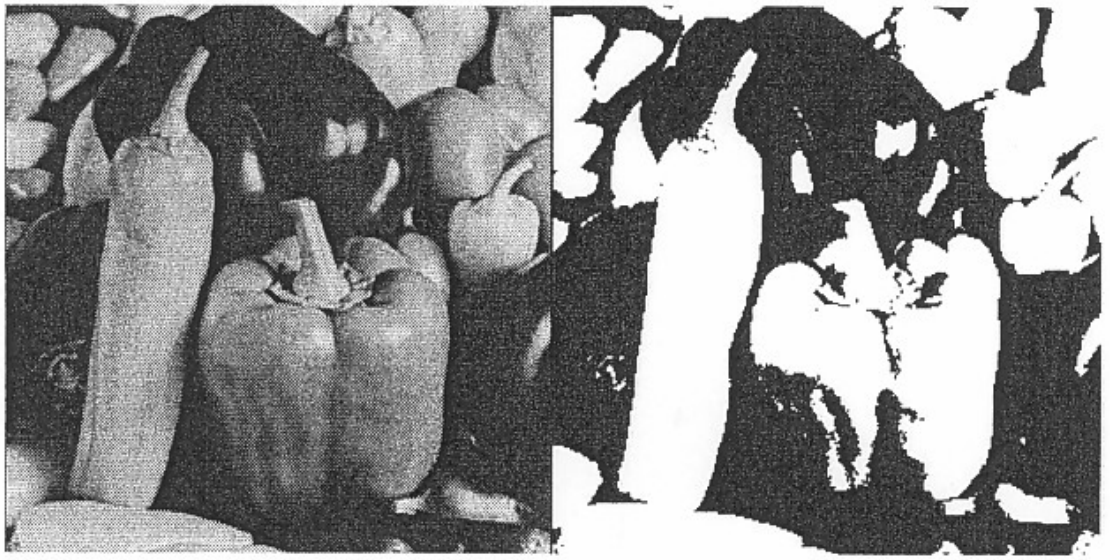

[Jain, Kasturi, and Schunck (1995). Machine Vision, Ch. 1]

Local Operations

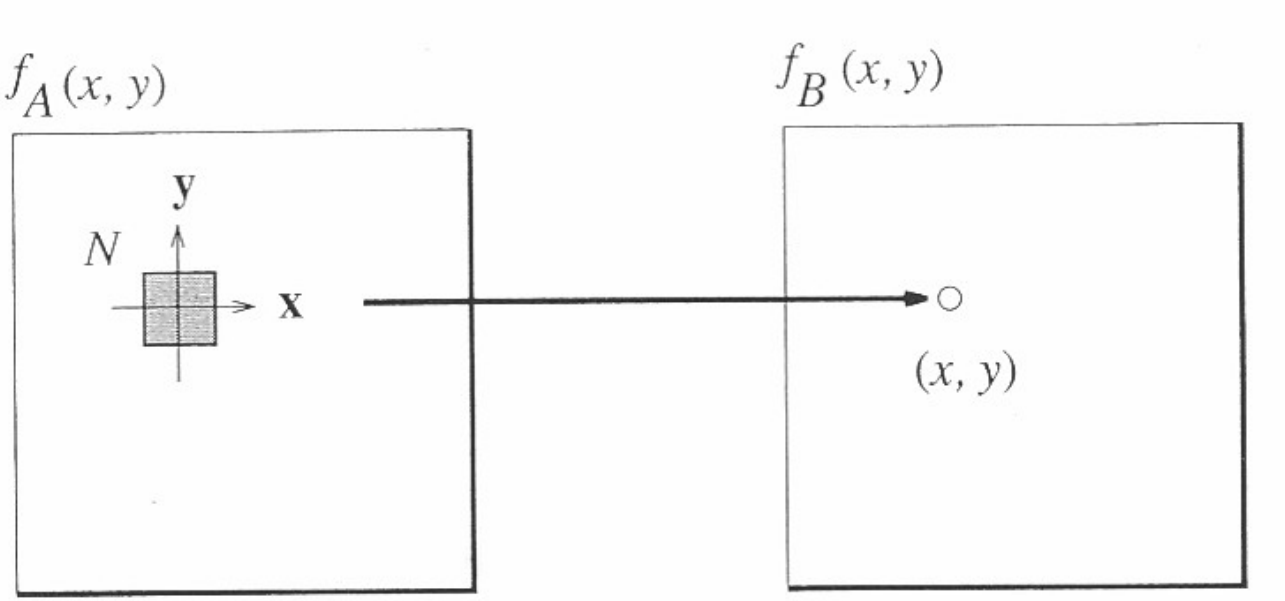

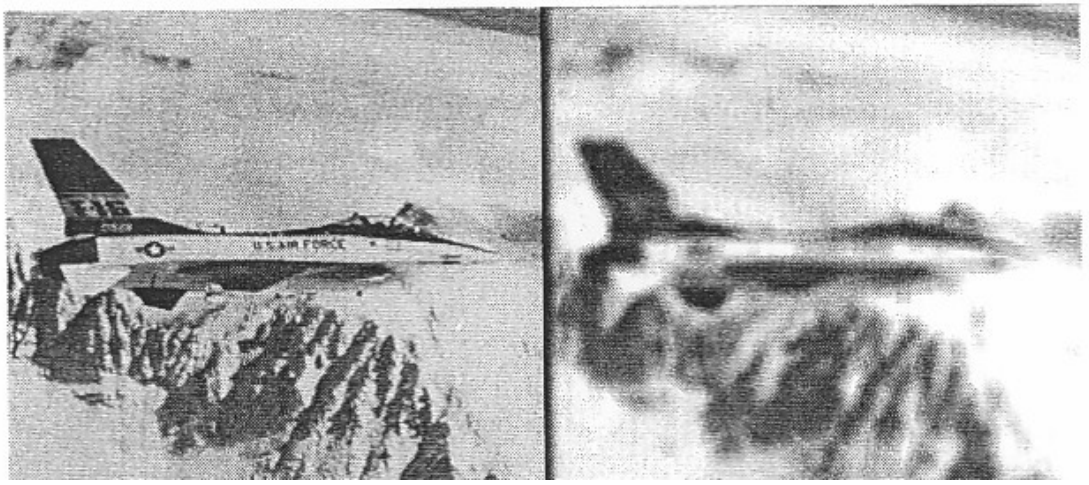

[Jain, Kasturi, and Schunck (1995). Machine Vision, Ch. 1]

Local Operations

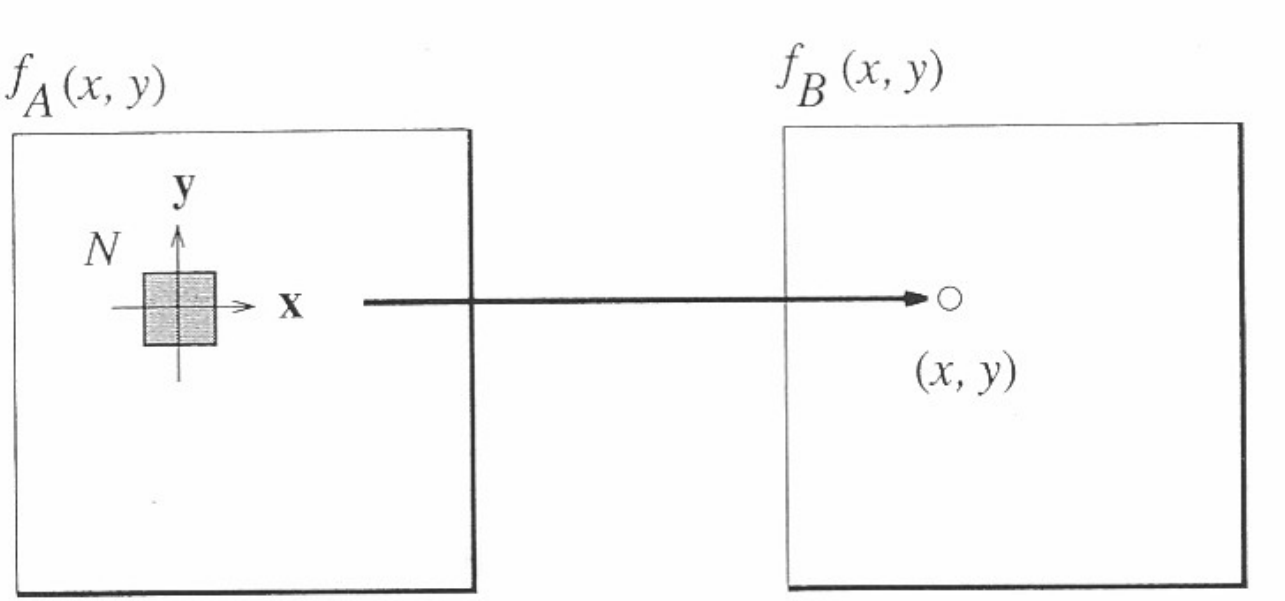

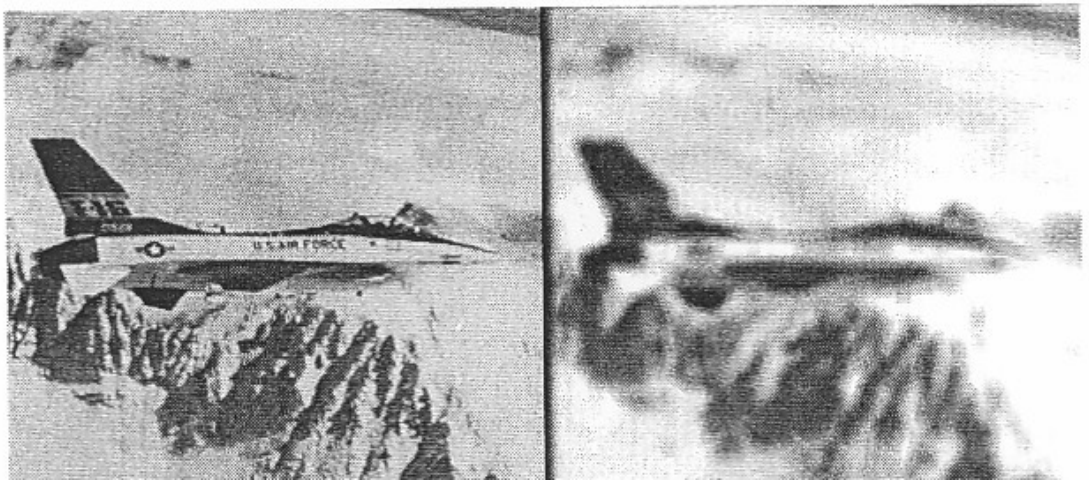

[Jain, Kasturi, and Schunck (1995). Machine Vision, Ch. 1]

Edge Detection

[https://www.mathworks.com/matlabcentral/fileexchange/51124-shannon-edge-detector-for-grayscale-images]

Early Example of Edge Detection by Robots

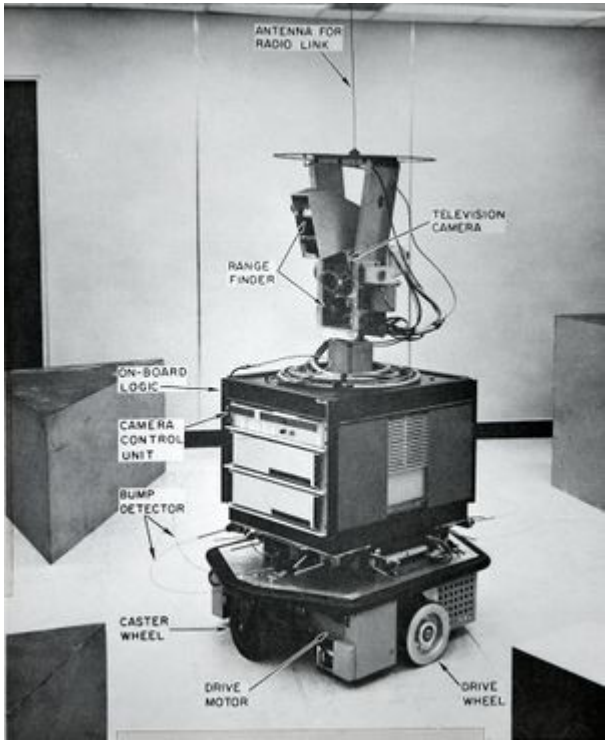

Global Operations

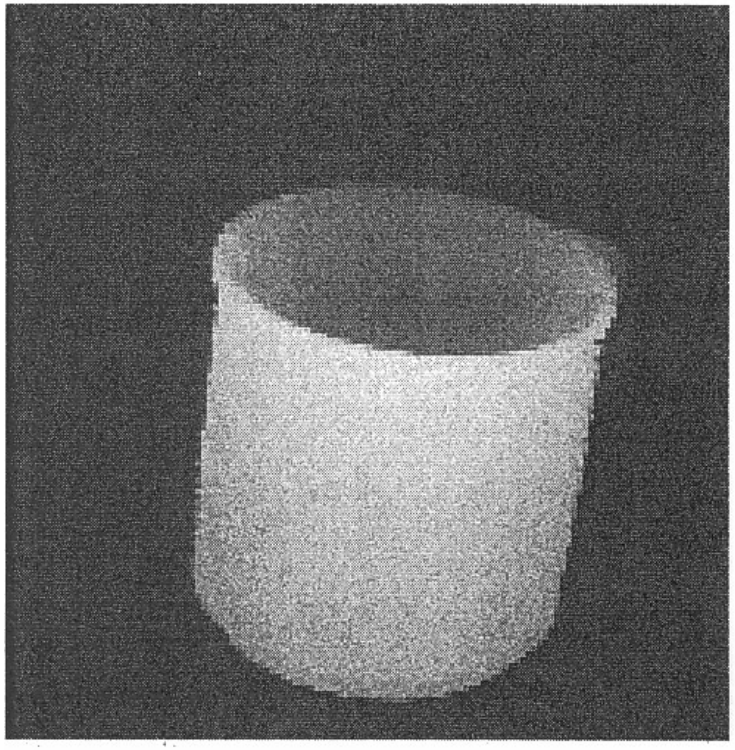

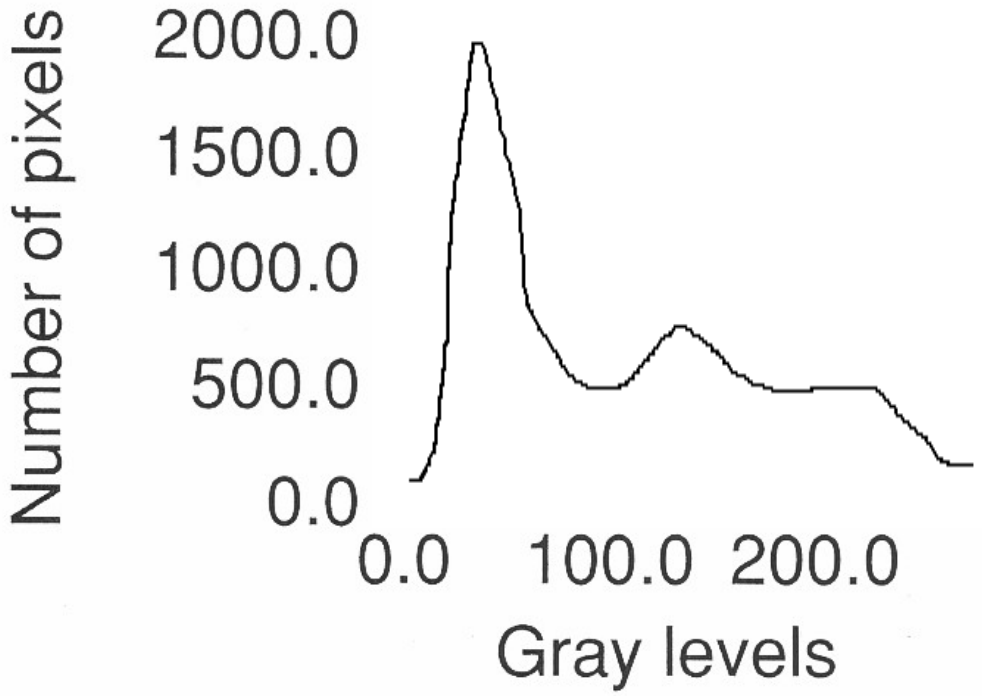

$$P = O_{\text{global}}\{f[i,j]\}$$

[Jain, Kasturi, and Schunck (1995). Machine Vision, Ch. 1]

Thresholding an Image

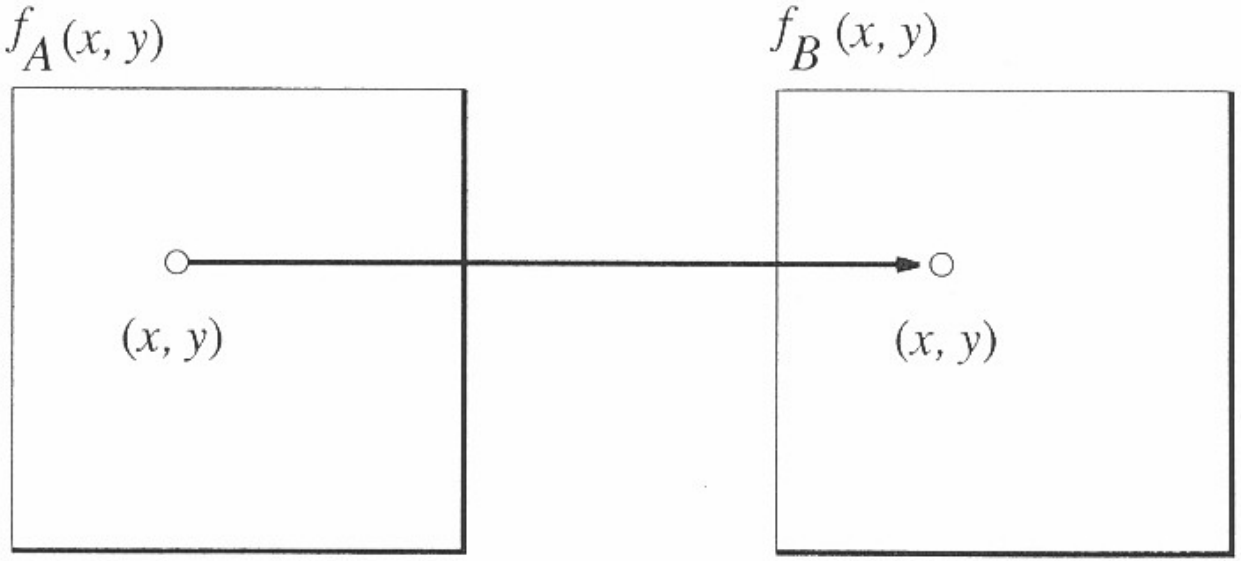

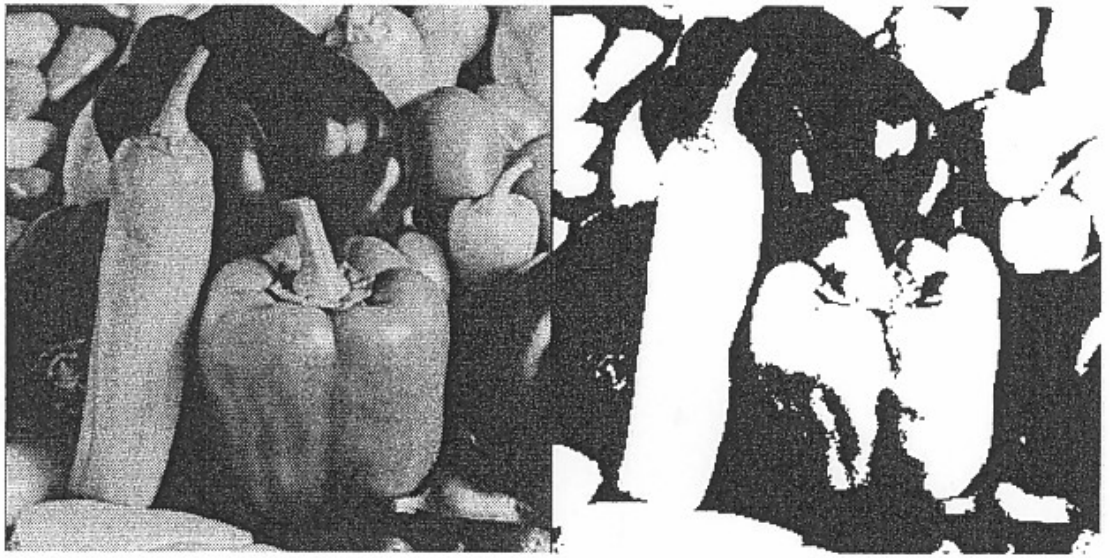

[Jain, Kasturi, and Schunck (1995). Machine Vision, Ch. 1]

Dark Image on a Light Background

$$F_T[i,j] = \begin{cases} 1 & \text{if } F[i,j] \leq T \\ 0 & \text{otherwise.} \end{cases}$$

[Jain, Kasturi, and Schunck (1995). Machine Vision, Ch. 2]

Selecting a range of intensity values

$$F_T[i,j] = \begin{cases} 1 & \text{if } T_1 \leq F[i,j] \leq T_2 \\ 0 & \text{otherwise.} \end{cases}$$

[Jain, Kasturi, and Schunck (1995). Machine Vision, Ch. 2]

Generalized Thresholding

A general thresholding scheme in which the intensity levels for an object may come from several disjoint intervals may be represented as

$$F_T[i,j] = \begin{cases} 1 & \text{if } F[i,j] \in Z \\ 0 & \text{otherwise} \end{cases} \qquad (2.4)$$

[Jain, Kasturi, and Schunck (1995). Machine Vision, Ch. 2]

Thresholding Example (1)

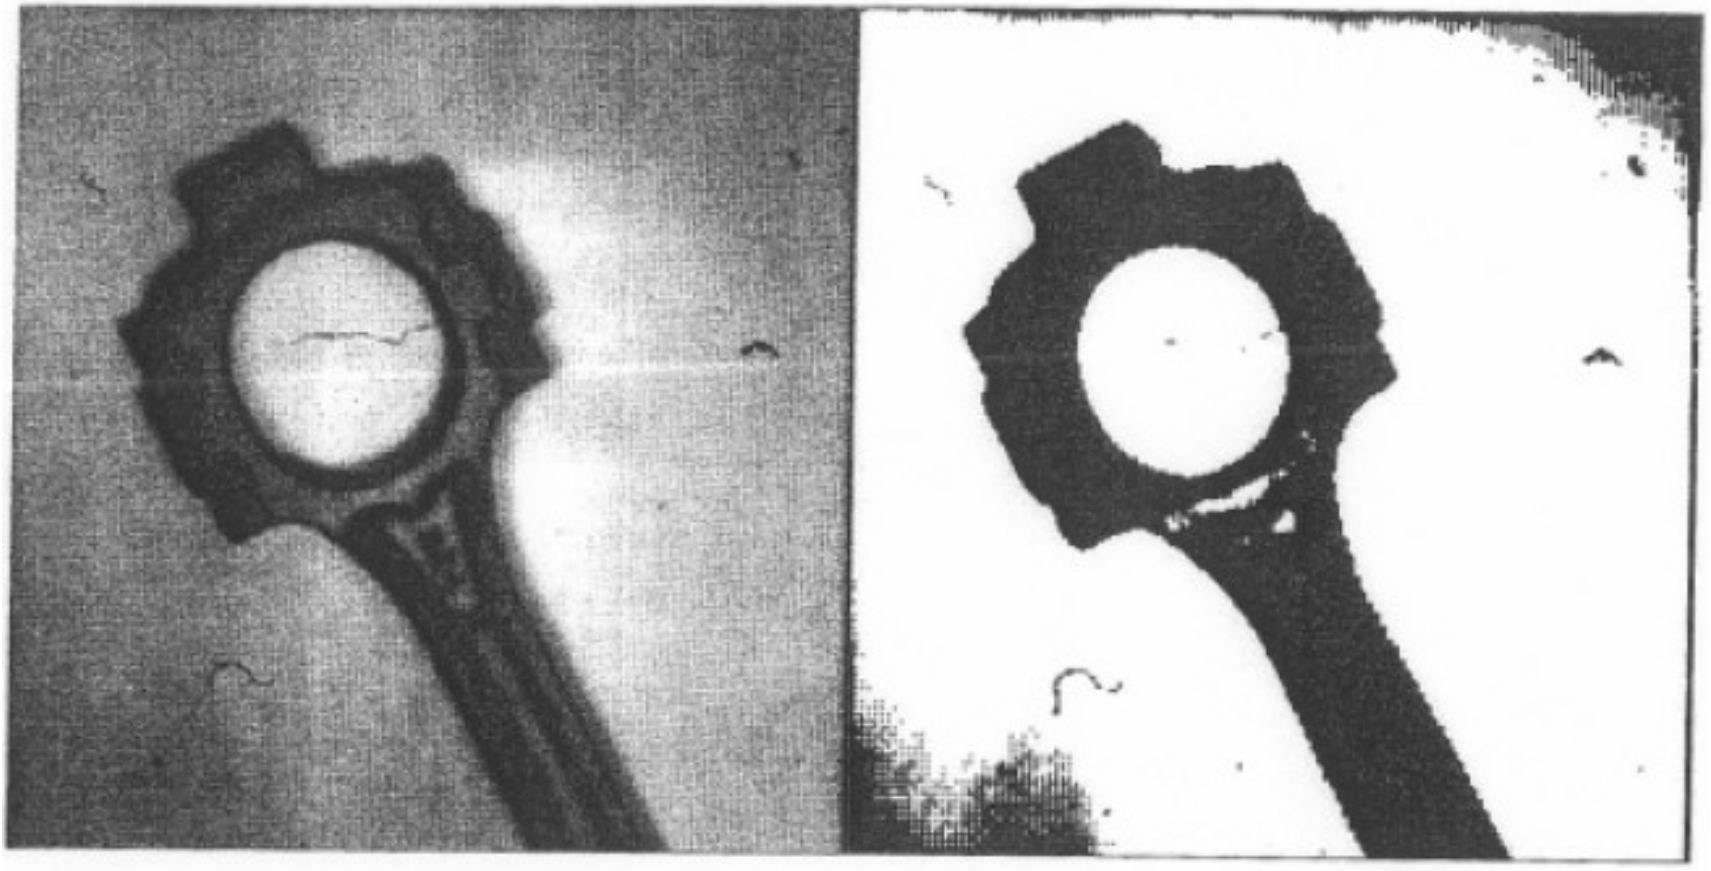

[Jain, Kasturi, and Schunck (1995). Machine Vision, Ch. 2]

Thresholding Example (2)

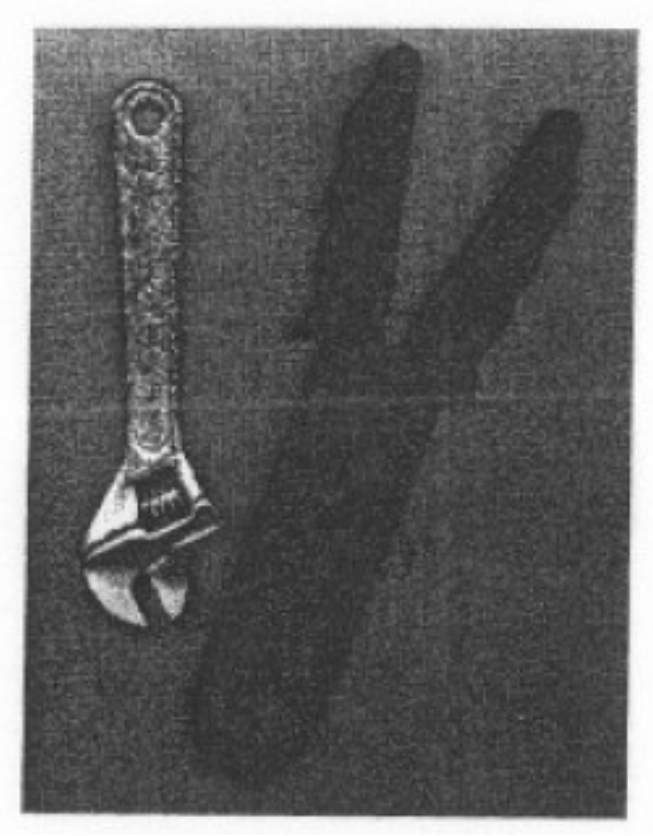

Original grayscale Image

[Jain, Kasturi, and Schunck (1995). Machine Vision, Ch. 2]

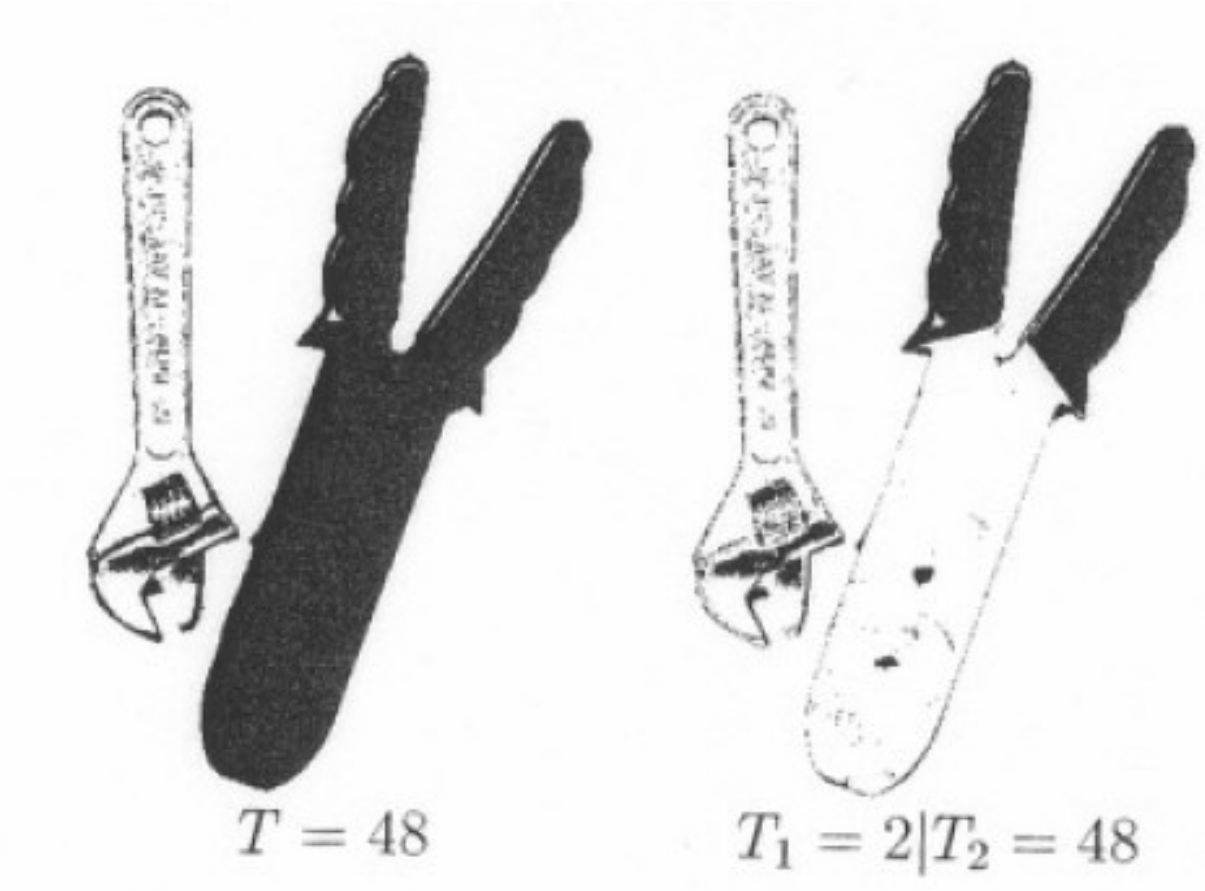

[Jain, Kasturi, and Schunck (1995). Machine Vision, Ch. 2]

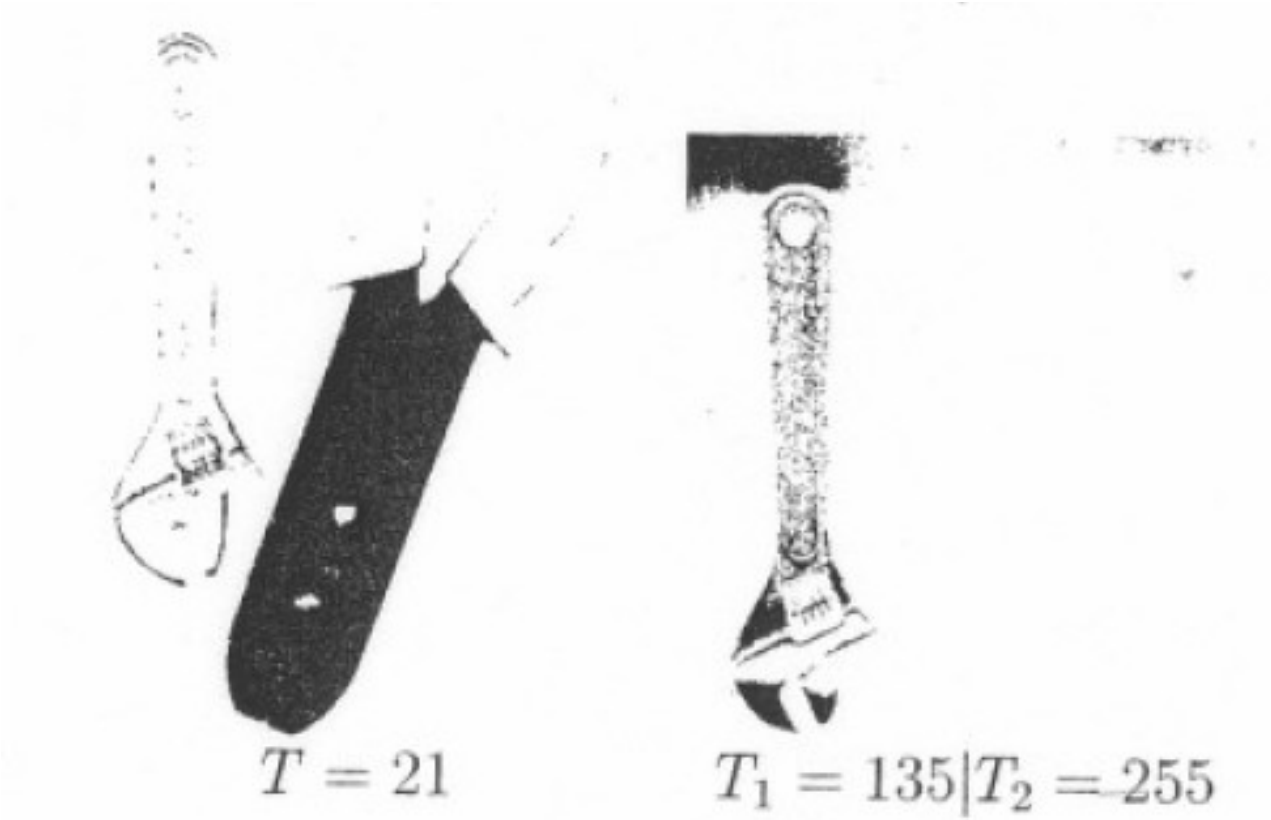

[Jain, Kasturi, and Schunck (1995). Machine Vision, Ch. 2]

Area of a Binary Image

$$A = \sum_{i=1}^{n} \sum_{j=1}^{m} B[i, j].$$

[Jain, Kasturi, and Schunck (1995). Machine Vision, Ch. 2]

This figure now becomes important

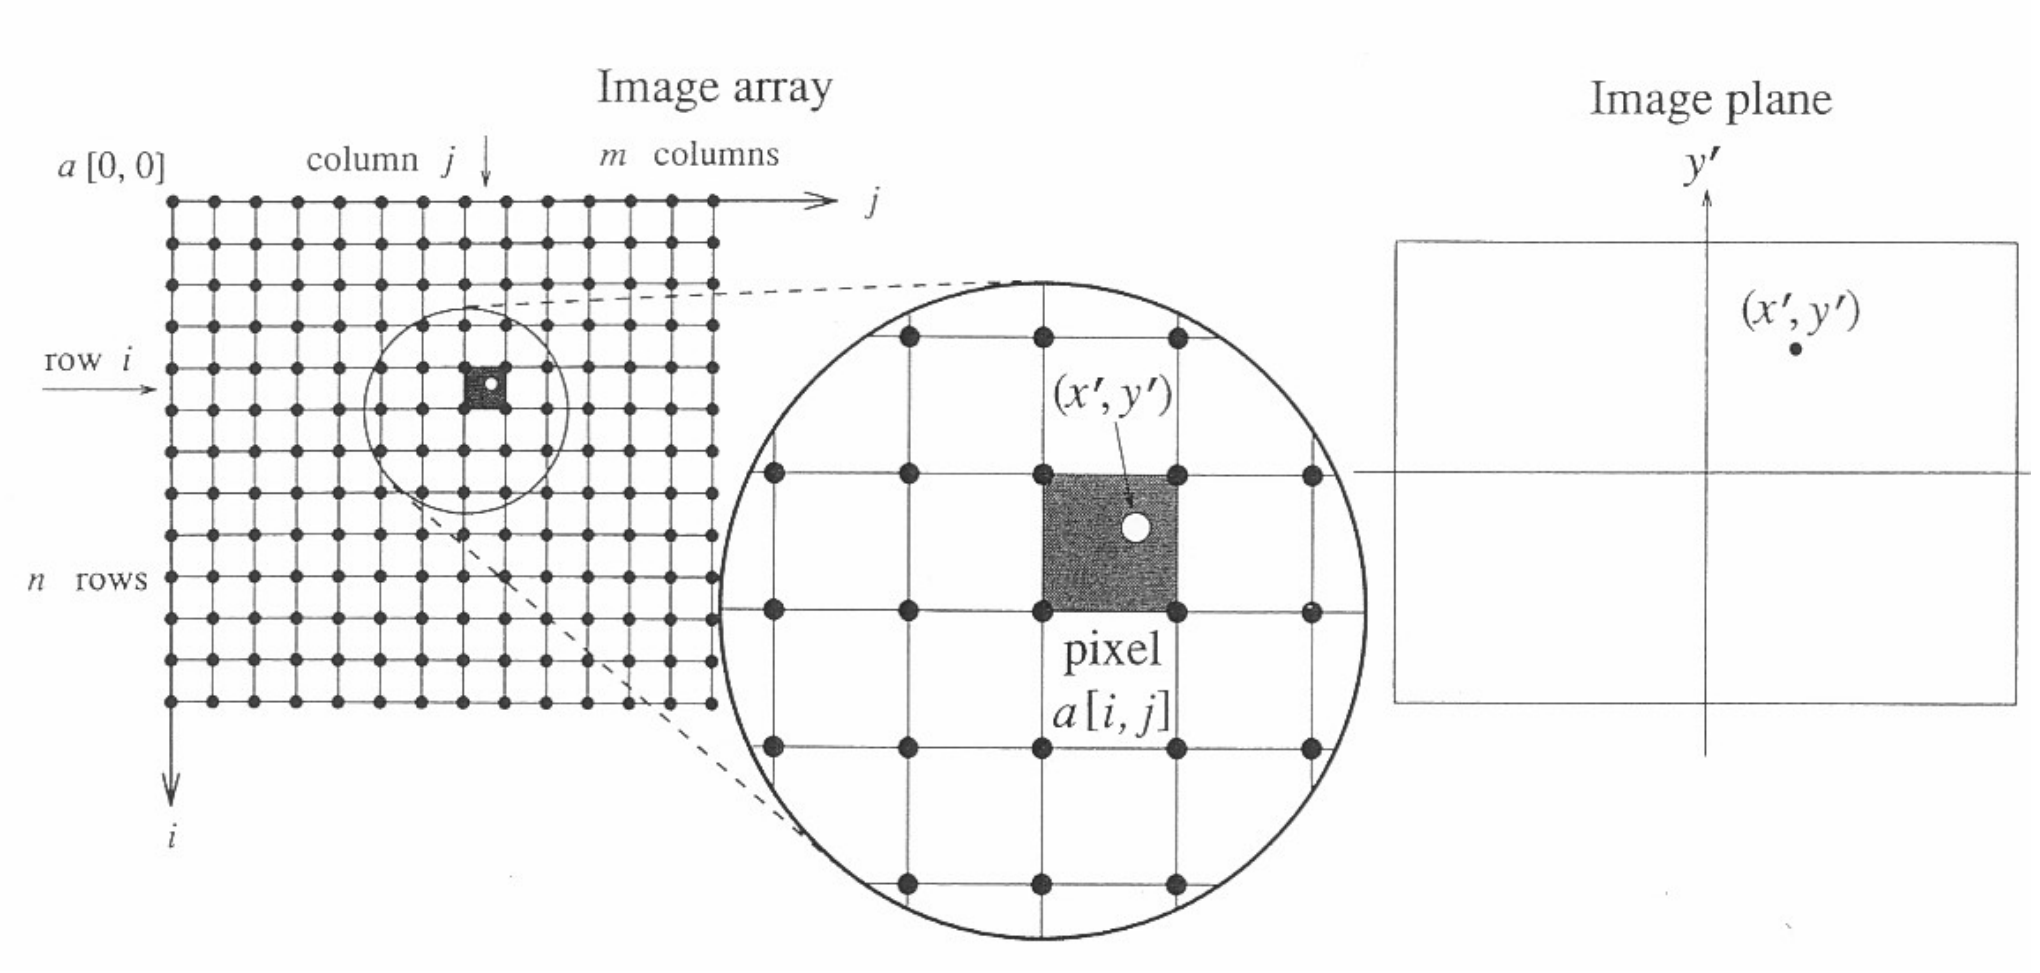

[Jain, Kasturi, and Schunck (1995). Machine Vision, Ch. 1]

Calculating the Position of an Object

$$\bar{x} \sum_{i=1}^{n} \sum_{j=1}^{m} B[i,j] = \sum_{i=1}^{n} \sum_{j=1}^{m} jB[i,j]$$

$$\bar{y} \sum_{i=1}^{n} \sum_{j=1}^{m} B[i,j] = \sum_{i=1}^{n} \sum_{j=1}^{m} iB[i,j]$$

[Jain, Kasturi, and Schunck (1995). Machine Vision, Ch. 2]

The center is given by

$$\bar{x} = \frac{\sum_{i=1}^{n}\sum_{j=1}^{m} jB[i,j]}{A}$$

$$\bar{y} = \frac{\sum_{i=1}^{n}\sum_{j=1}^{m} iB[i,j]}{A}.$$

[Jain, Kasturi, and Schunck (1995). Machine Vision, Ch. 2]

Horizontal and Vertical Projections

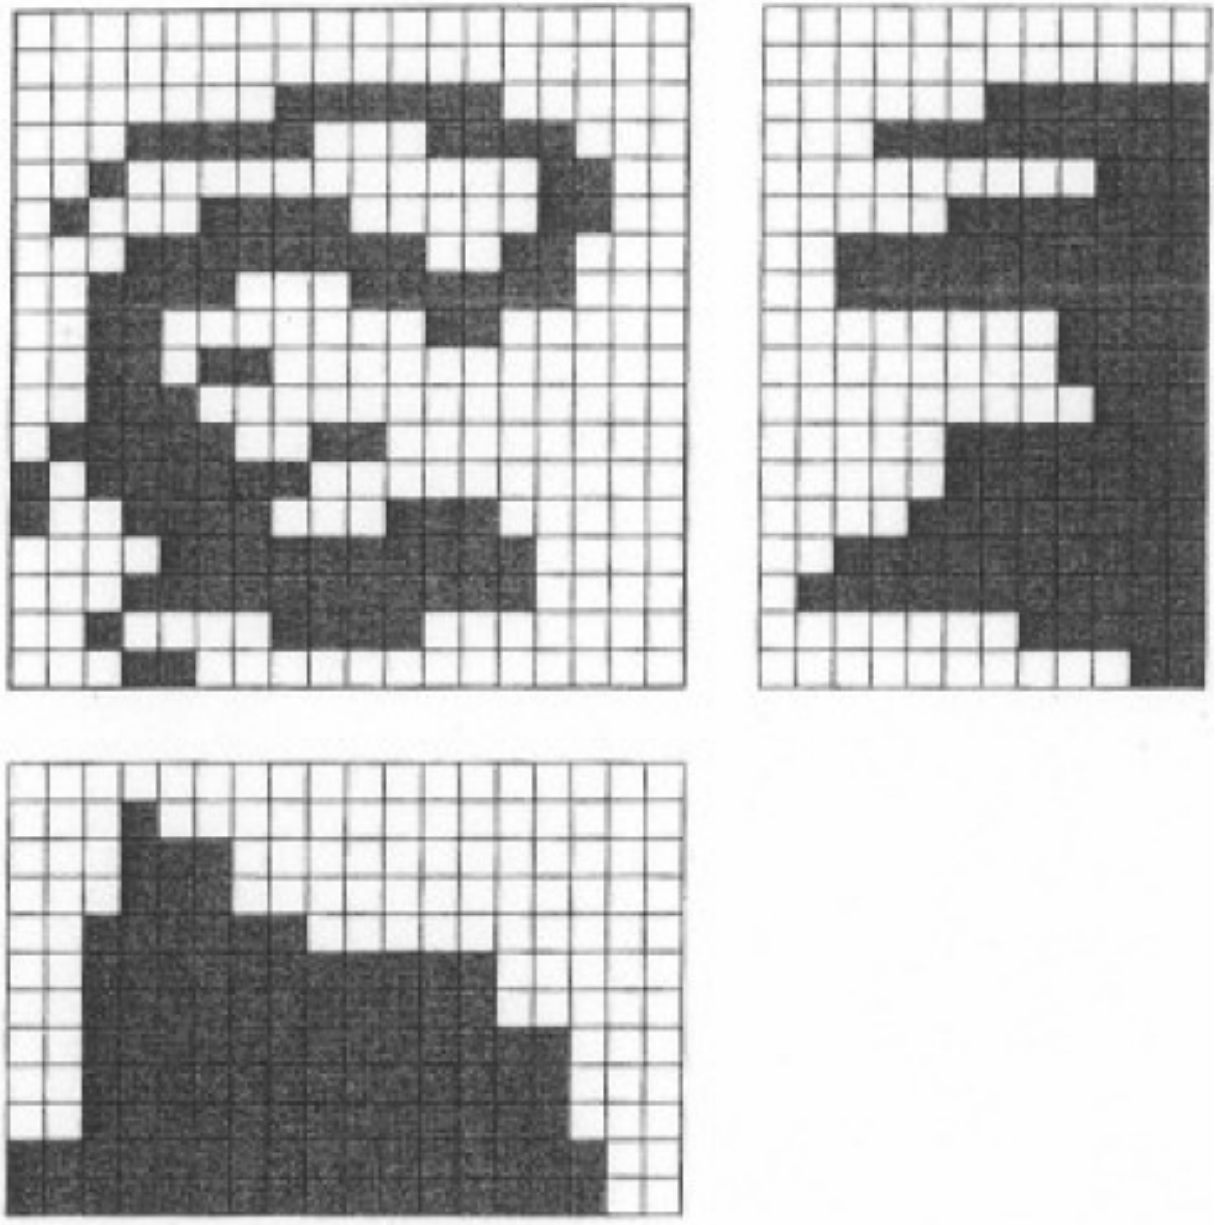

[Jain, Kasturi, and Schunck (1995). Machine Vision, Ch. 2]

Projection Formulas

$$H[i] = \sum_{j=1}^{m} B[i,j]$$

$$V[j] = \sum_{i=1}^{n} B[i,j].$$

[Jain, Kasturi, and Schunck (1995). Machine Vision, Ch. 2]

Diagonal Projection

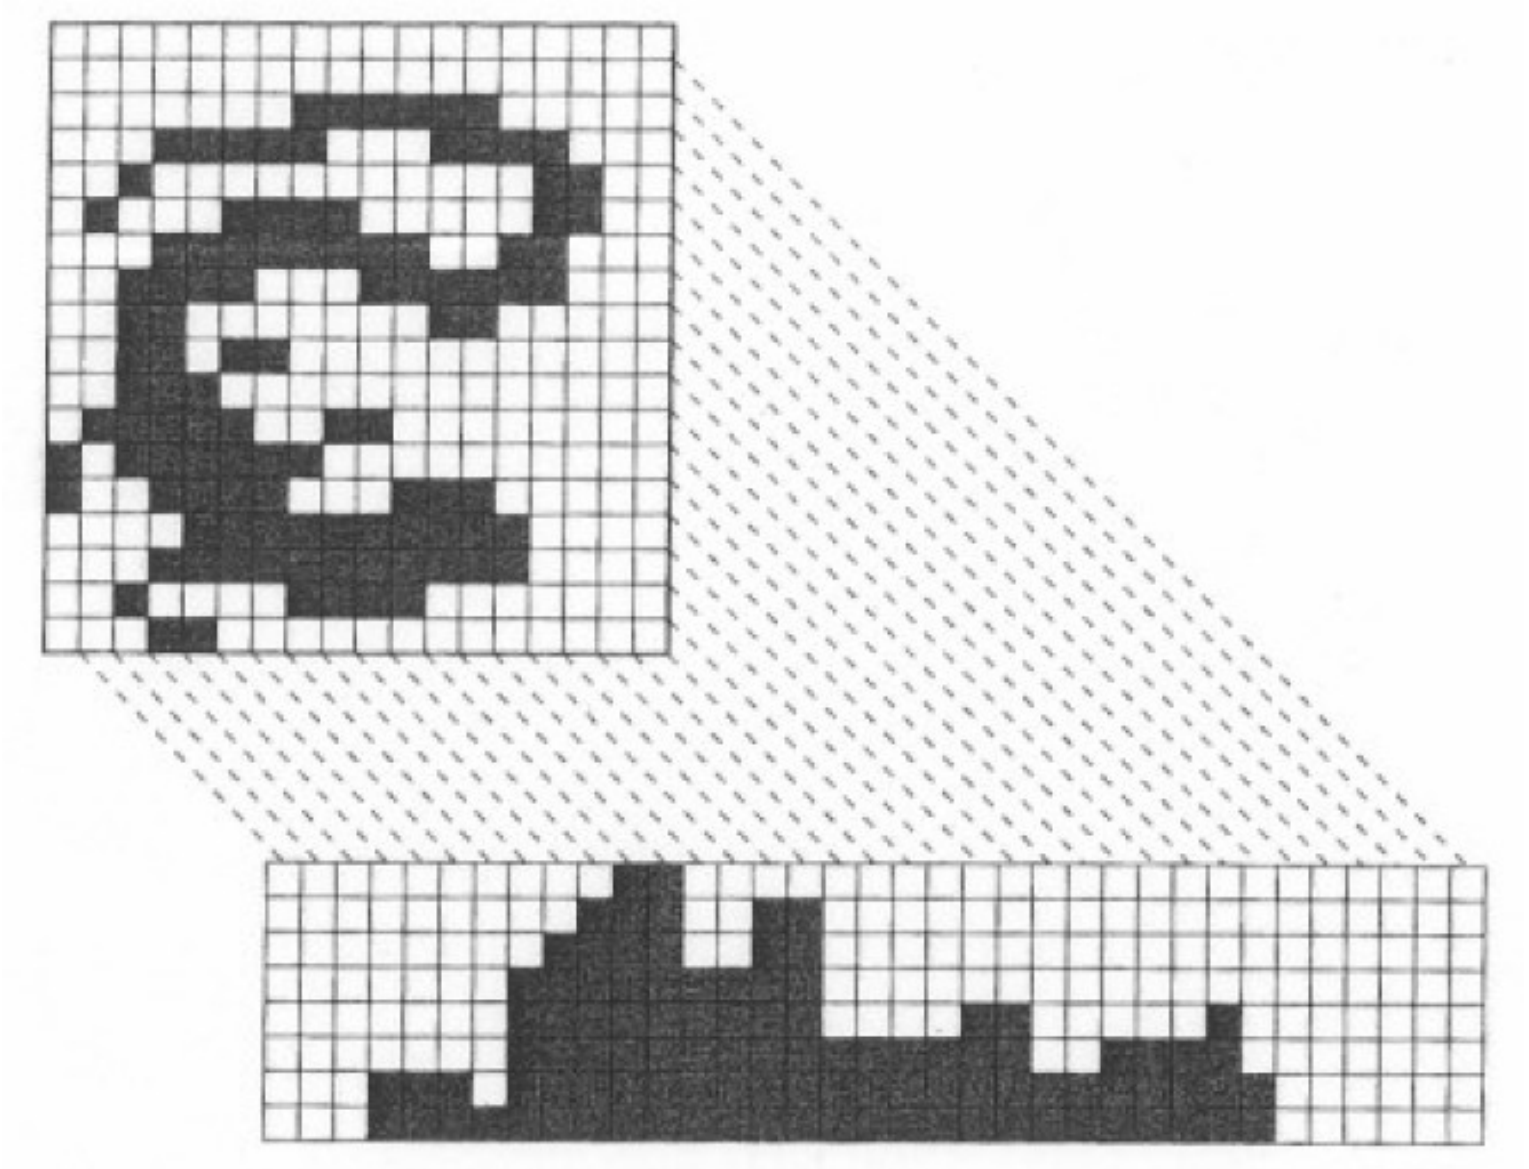

[Jain, Kasturi, and Schunck (1995). Machine Vision, Ch. 2]

The area and the position can be computed from the H and V projections

$$A = \sum_{j=1}^{m} V[j] = \sum_{i=1}^{n} H[i]$$

$$\bar{y} = \frac{\sum_{i=1}^{n} iH[i]}{A}$$

$$\bar{x} = \frac{\sum_{j=1}^{m} jV[j]}{A}.$$

[Jain, Kasturi, and Schunck (1995). Machine Vision, Ch. 2]

Neighbors and Connectivity

4-Connected

4-neighbors $[i+1, j]$, $[i-1, j]$, $[i, j+1]$, $[i, j-1]$

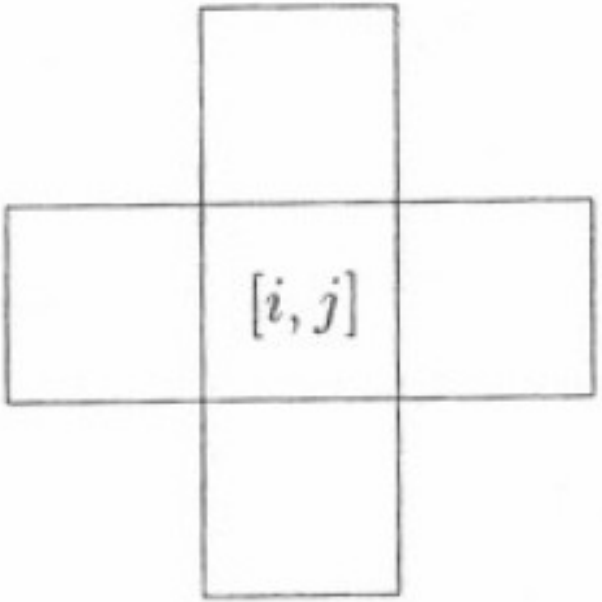

[Jain, Kasturi, and Schunck (1995). Machine Vision, Ch. 2]

8-connected

8-neighbors $[i+1, j+1]$, $[i+1, j-1]$, $[i-1, j+1]$, $[i-1, j-1]$ plus all of the 4-neighbors

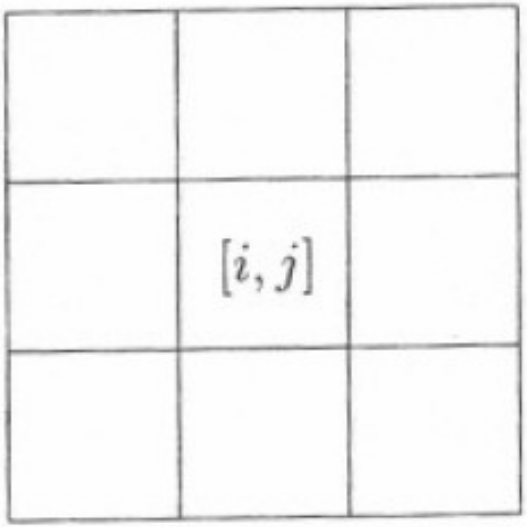

[Jain, Kasturi, and Schunck (1995). Machine Vision, Ch. 2]

Examples of Paths

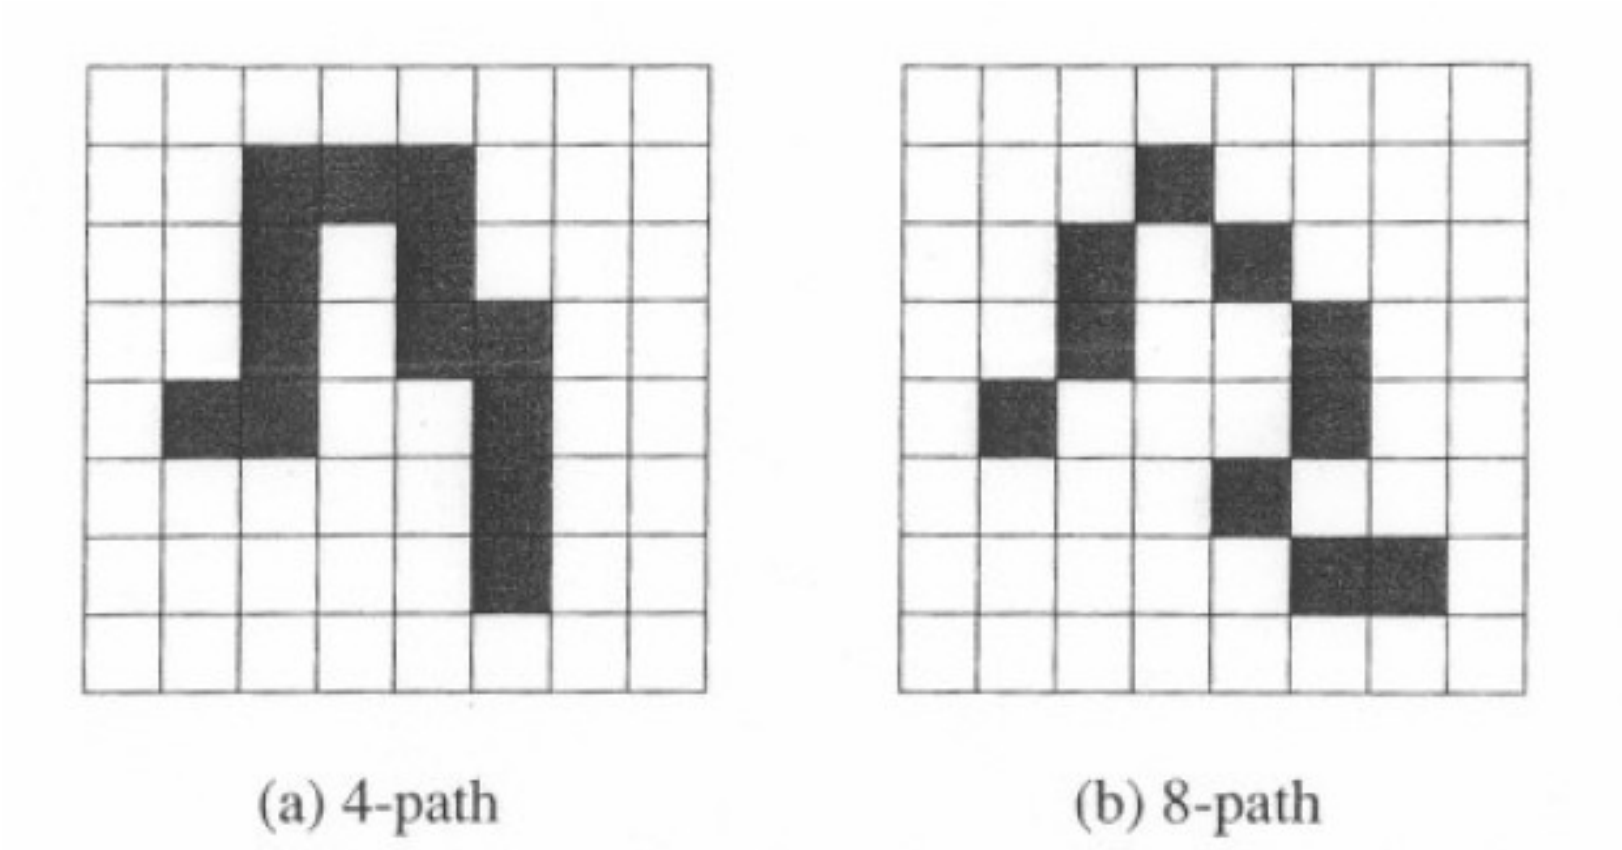

(a) 4-path

(b) 8-path

[Jain, Kasturi, and Schunck (1995). Machine Vision, Ch. 2]

Boundary, Interior, and Background

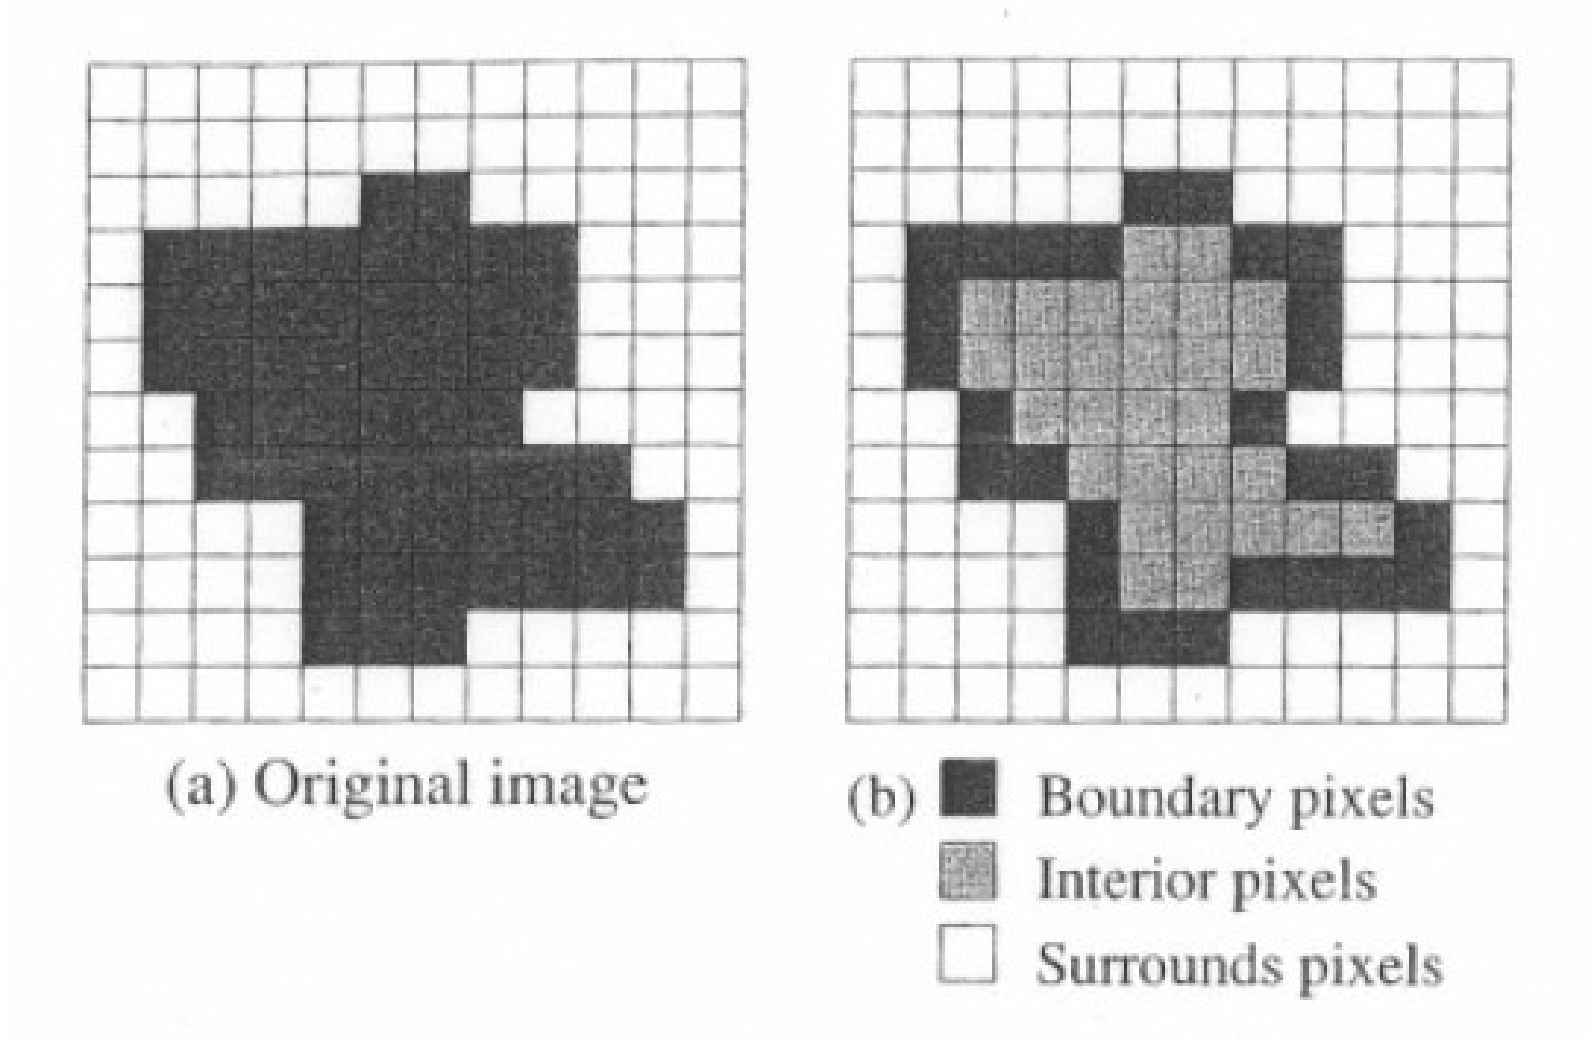

[Jain, Kasturi, and Schunck (1995). Machine Vision, Ch. 2]

An Image (a) and Its Connected Components (b)

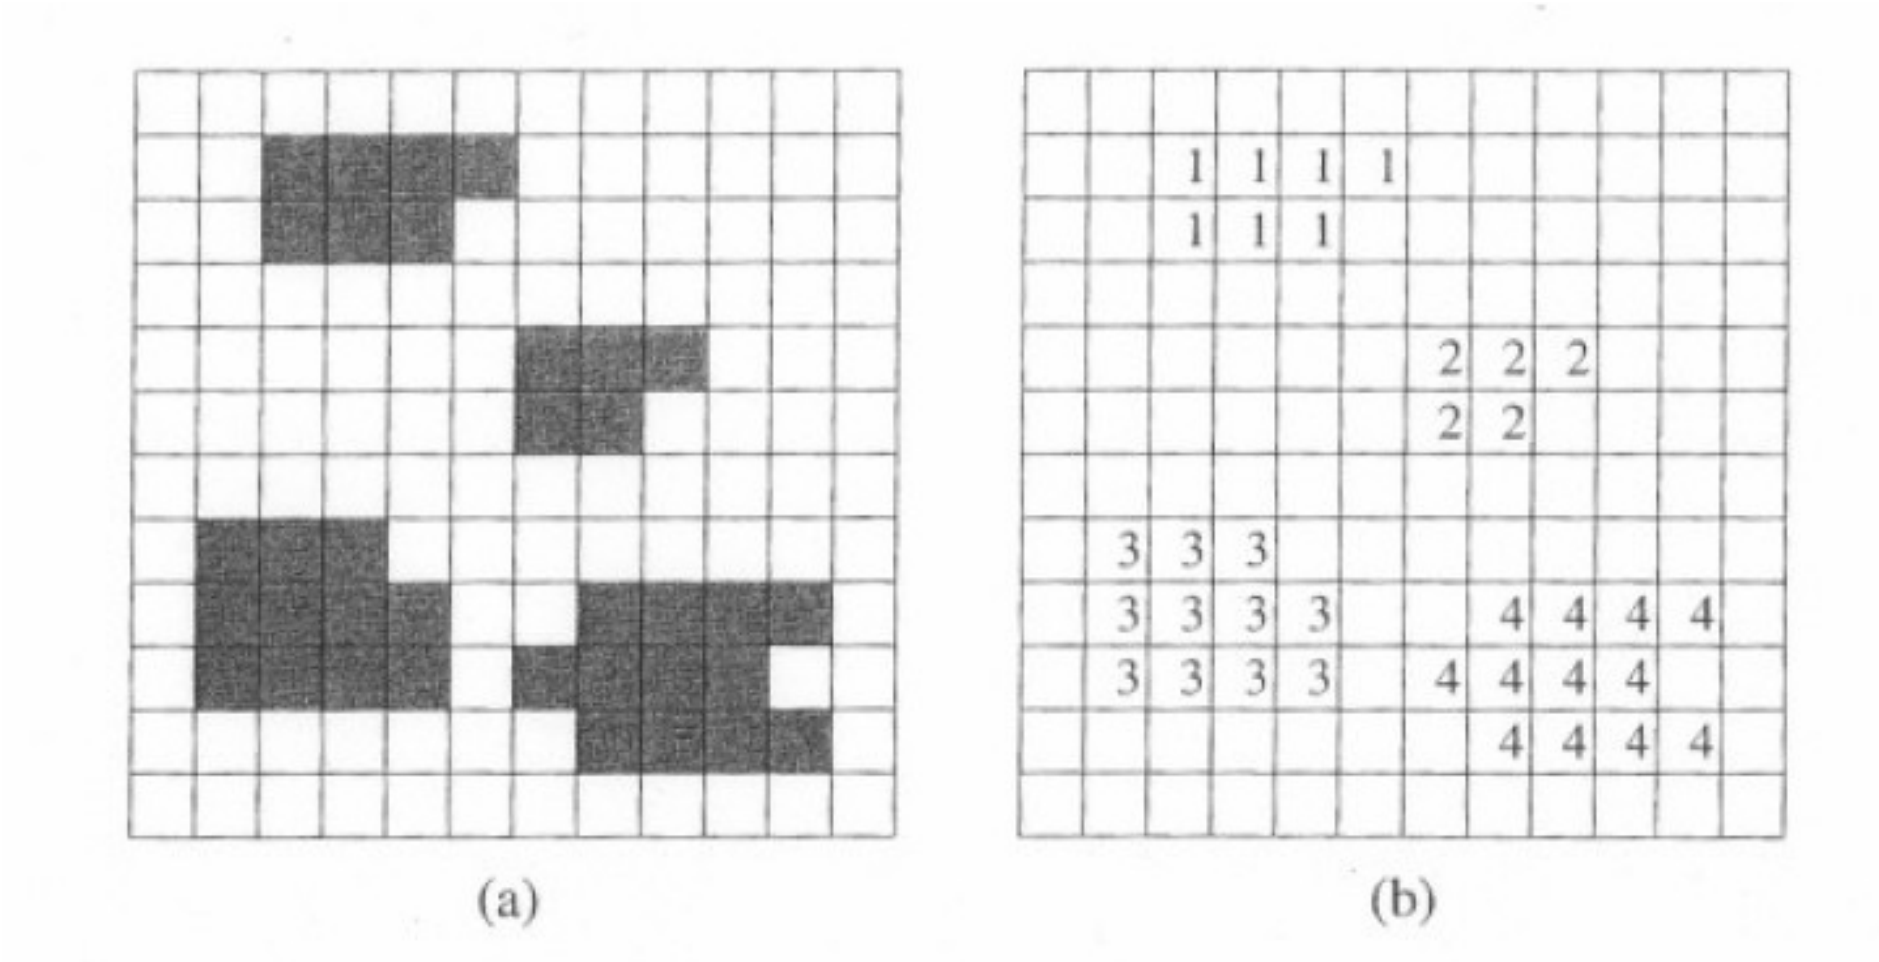

[Jain, Kasturi, and Schunck (1995). Machine Vision, Ch. 2]

Color Perception

The RGB Color Space

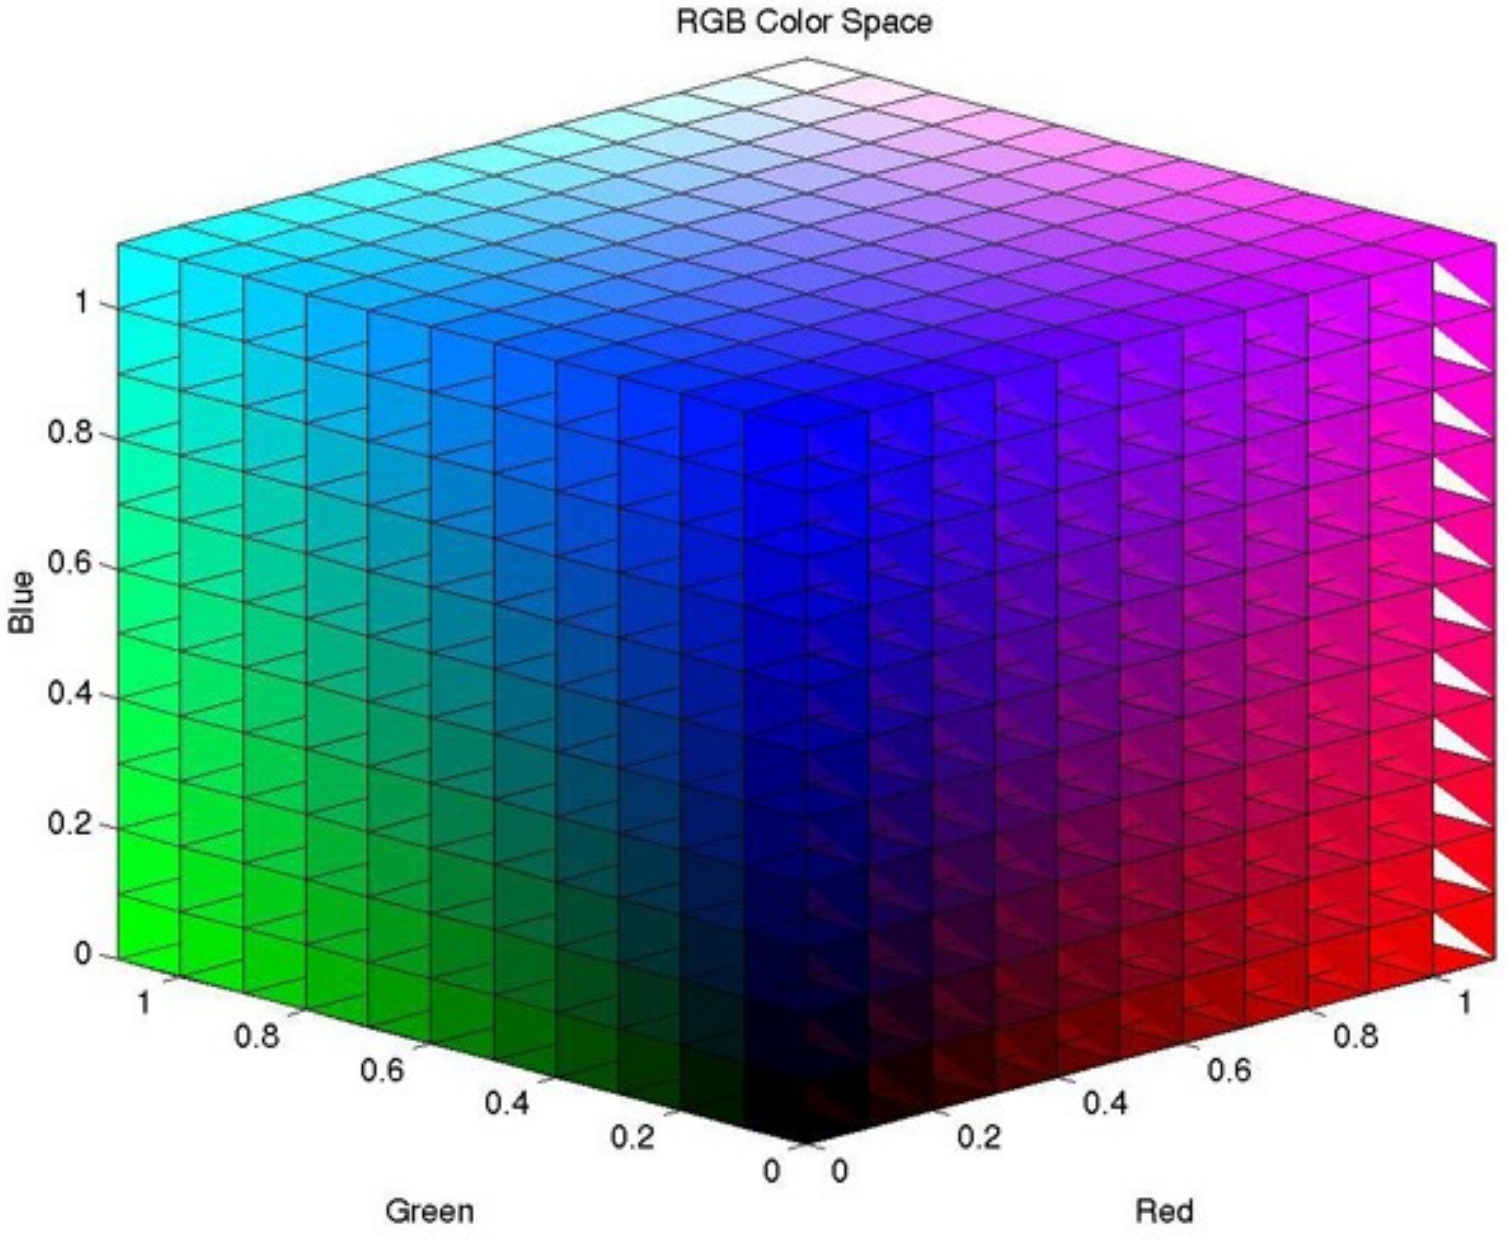

[http://www.arcsoft.com/images/topics/darkroom/what-is-color-space-RGB.jpg]

The RGB Color Space

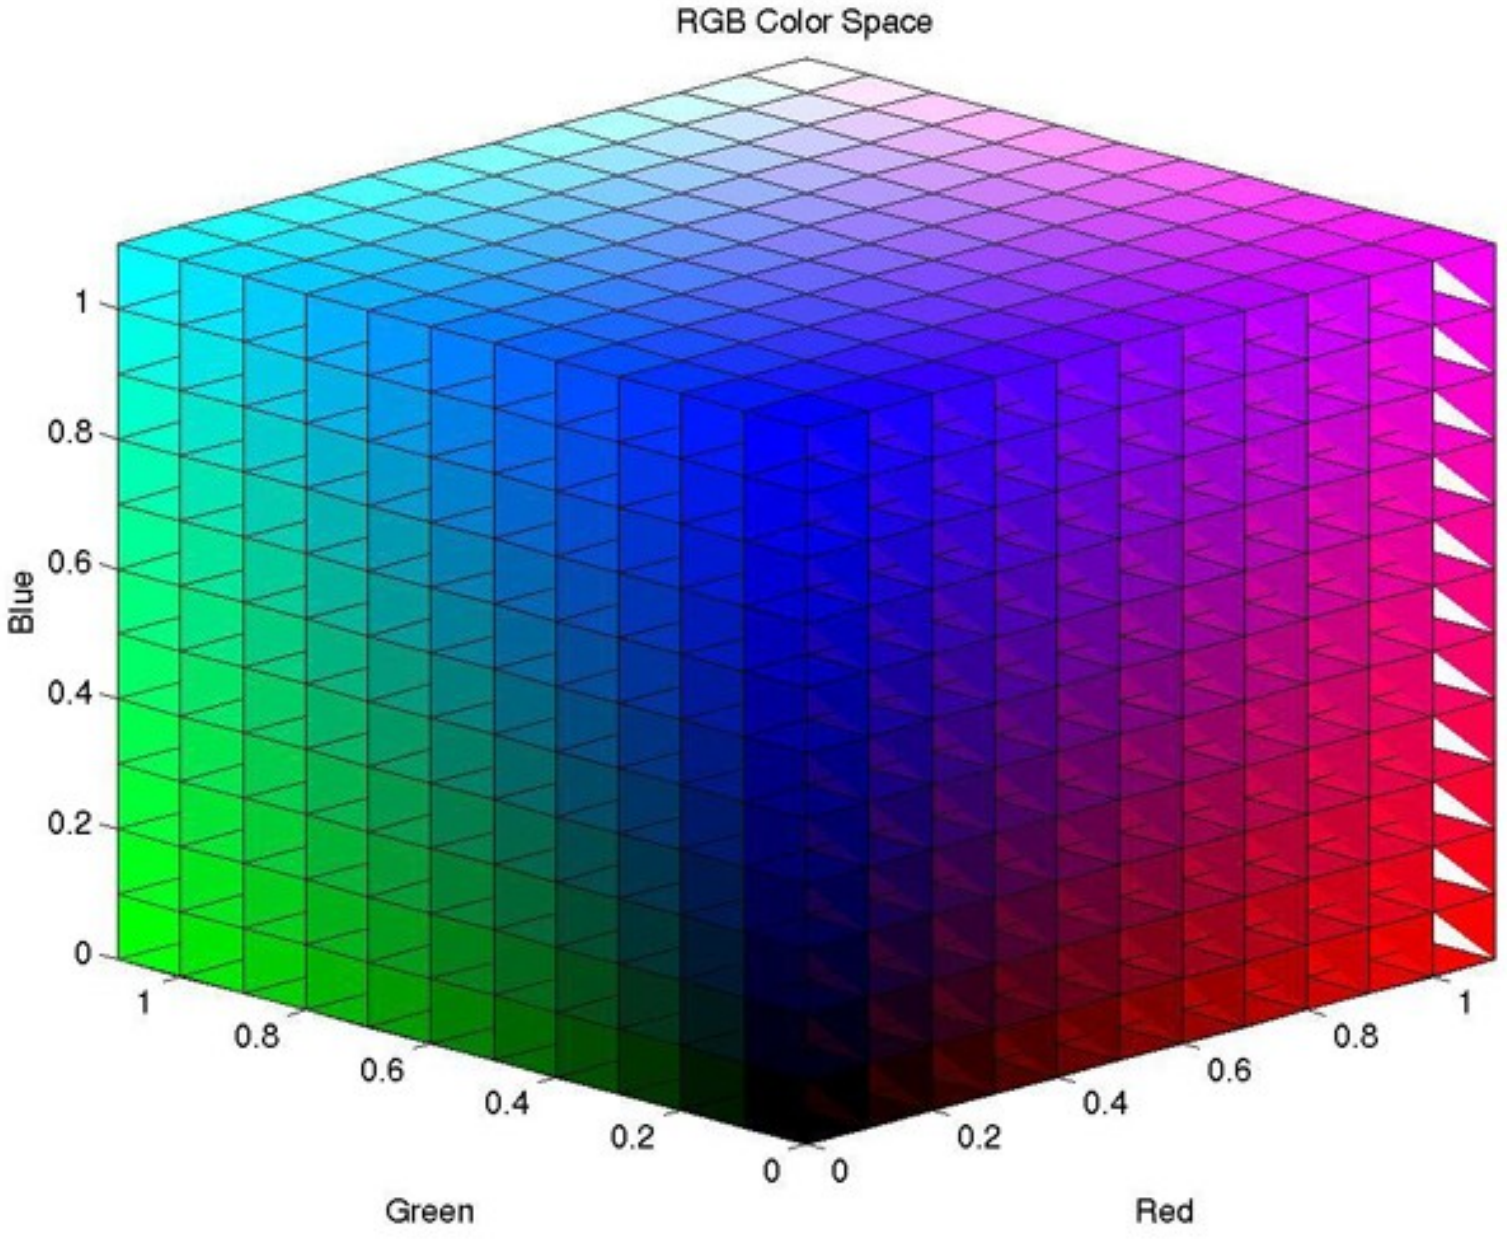

[http://www.arcsoft.com/images/topics/darkroom/what-is-color-space-RGB.jpg]

The RGB Color Space

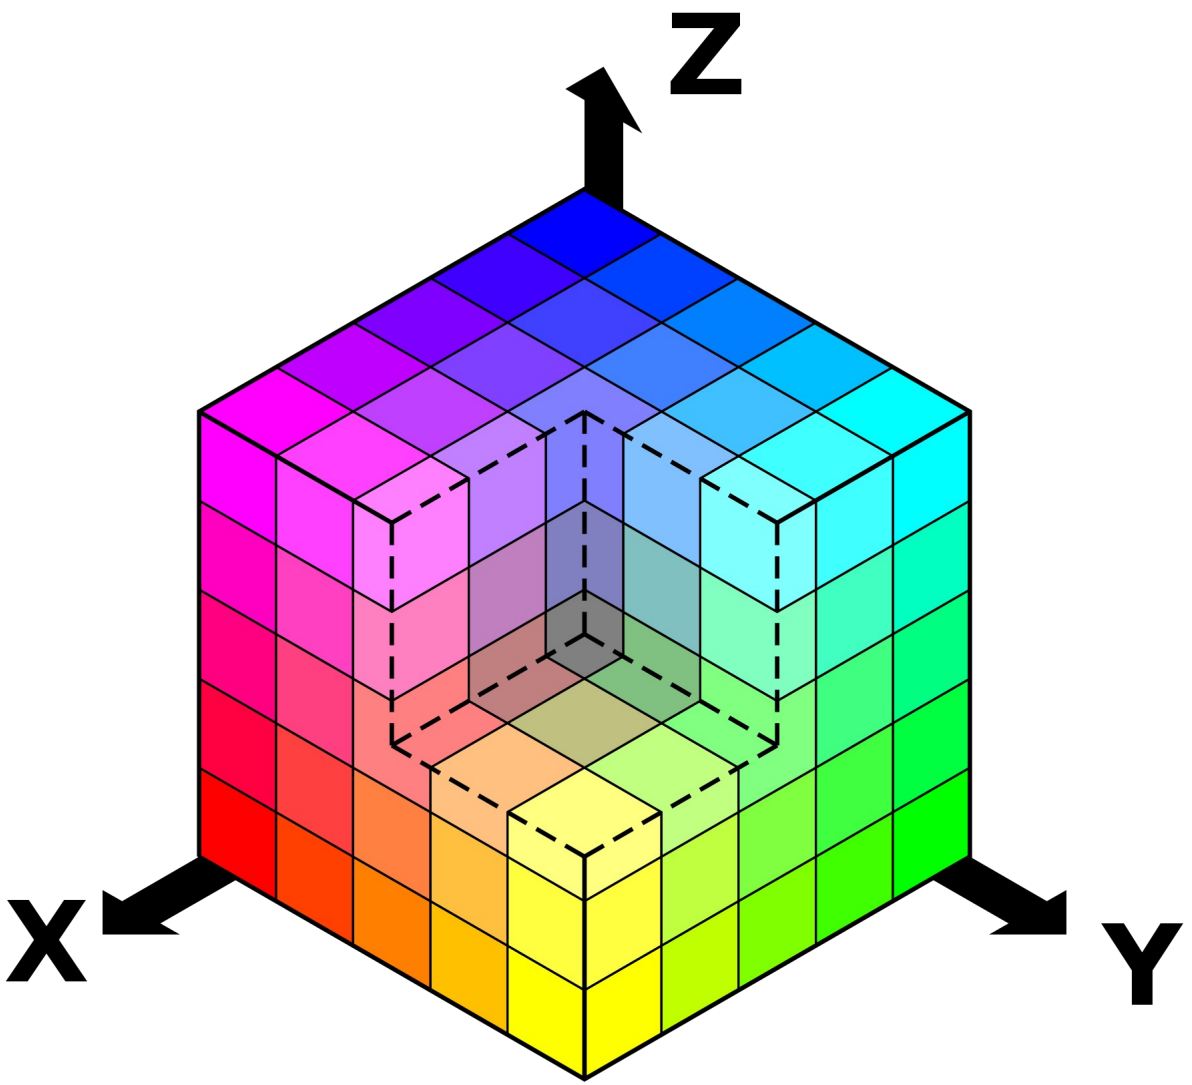

https://upload.wikimedia.org/wikipedia/commons/thumb/1/11/RGBCube_b.svg/2000px-RGBCube_b.svg.png

3D Scatter Plot for a patch of skin

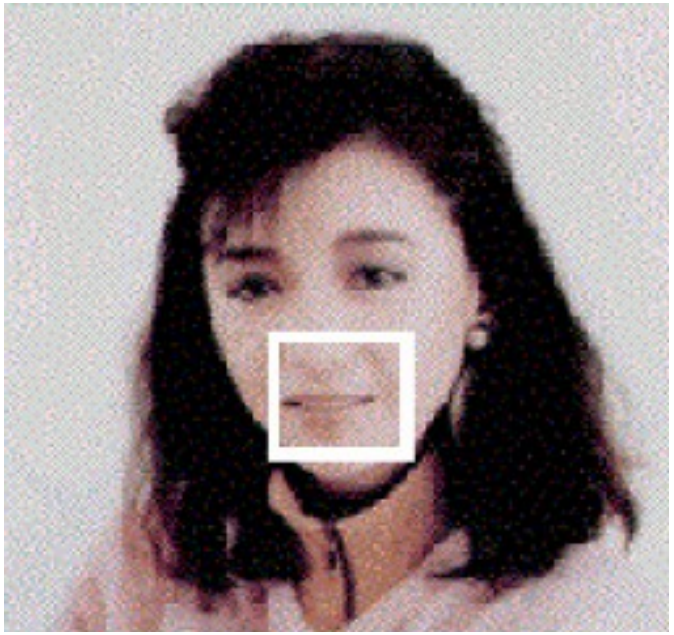

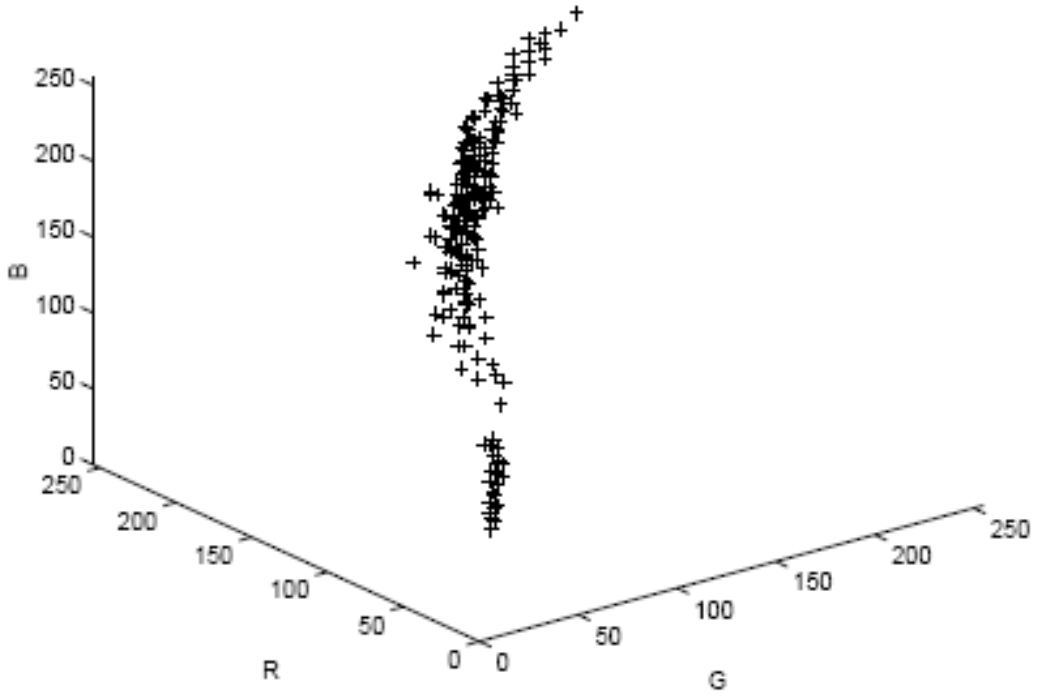

3-D scatterplot

The HSV Color Space

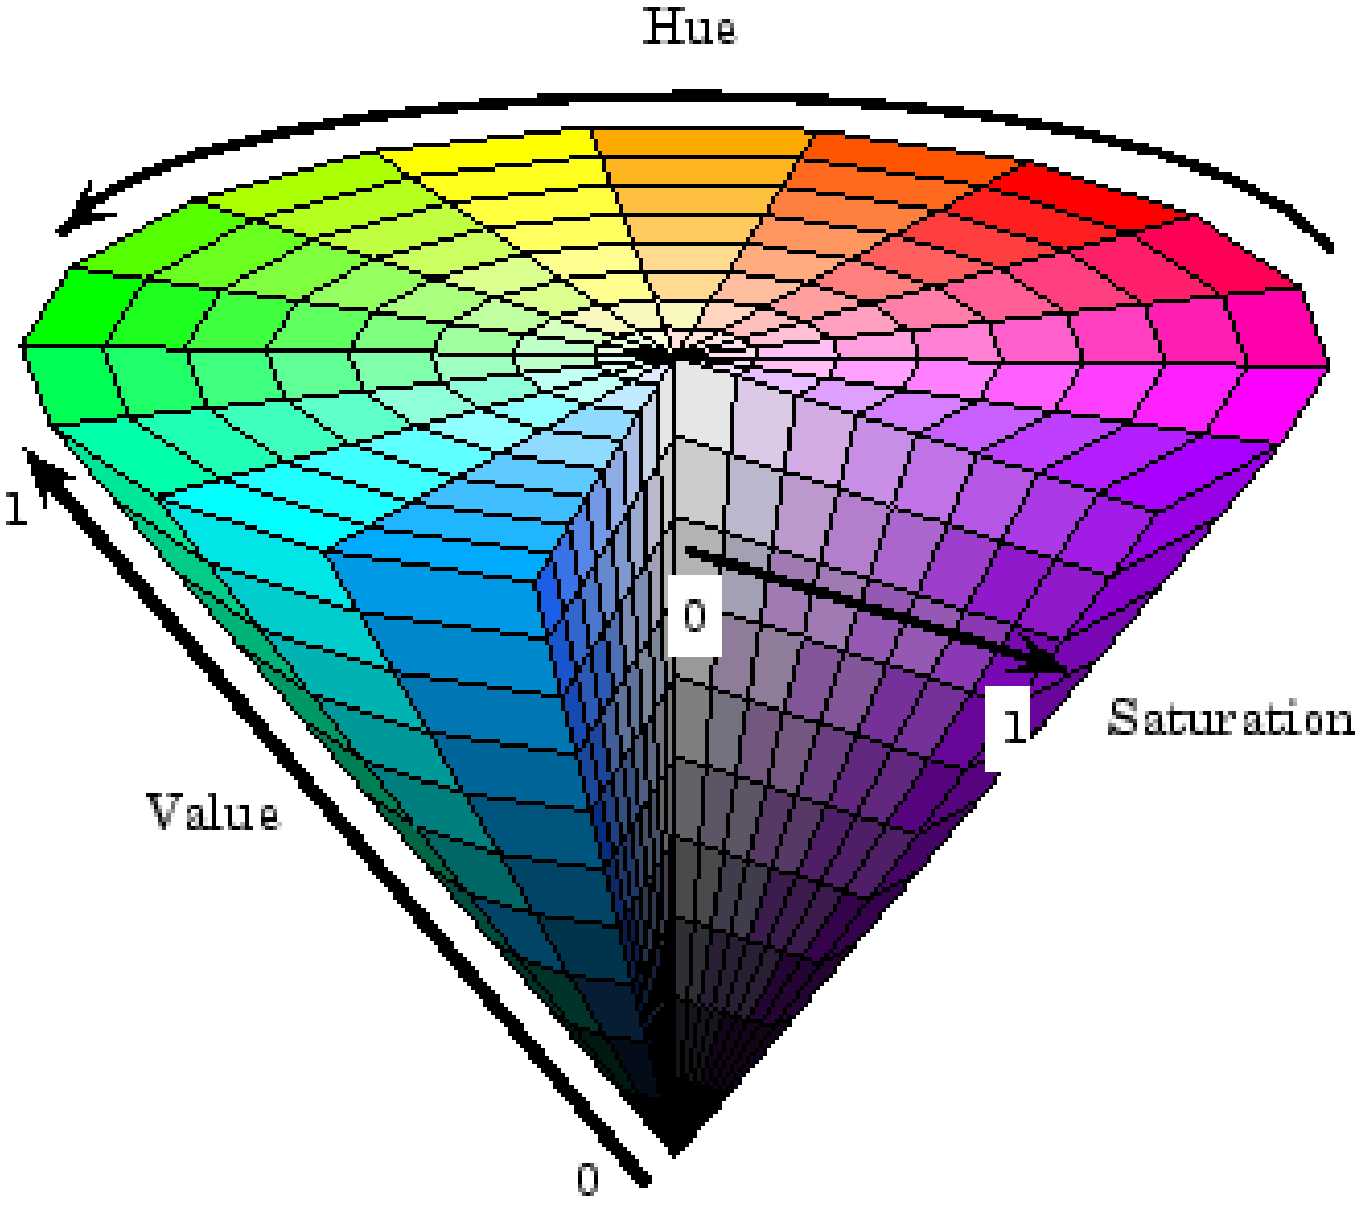

Color-based Tracking

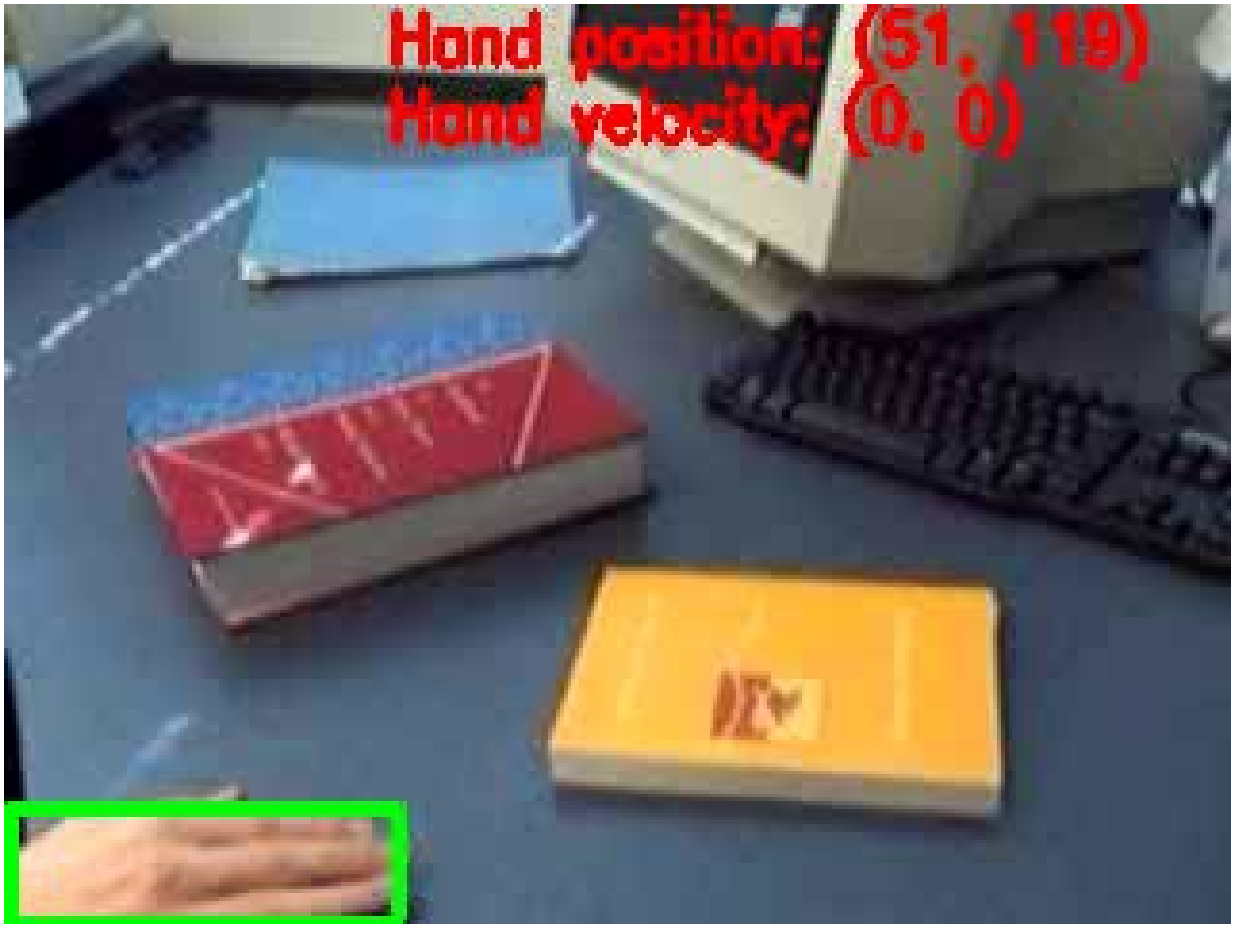

Color-based Tracking

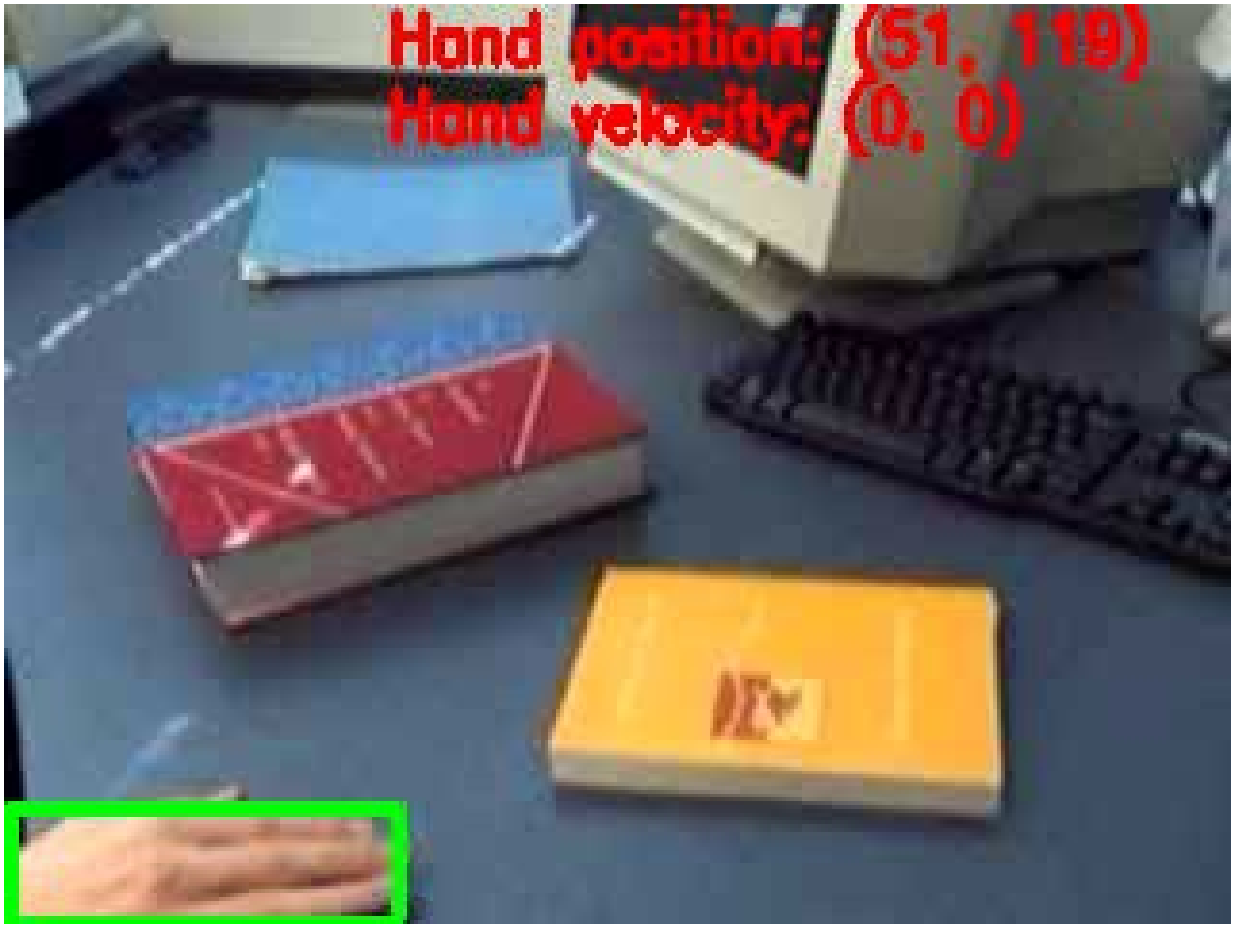

How should we determine the min and max thresholds for each color channel?

Color Histograms

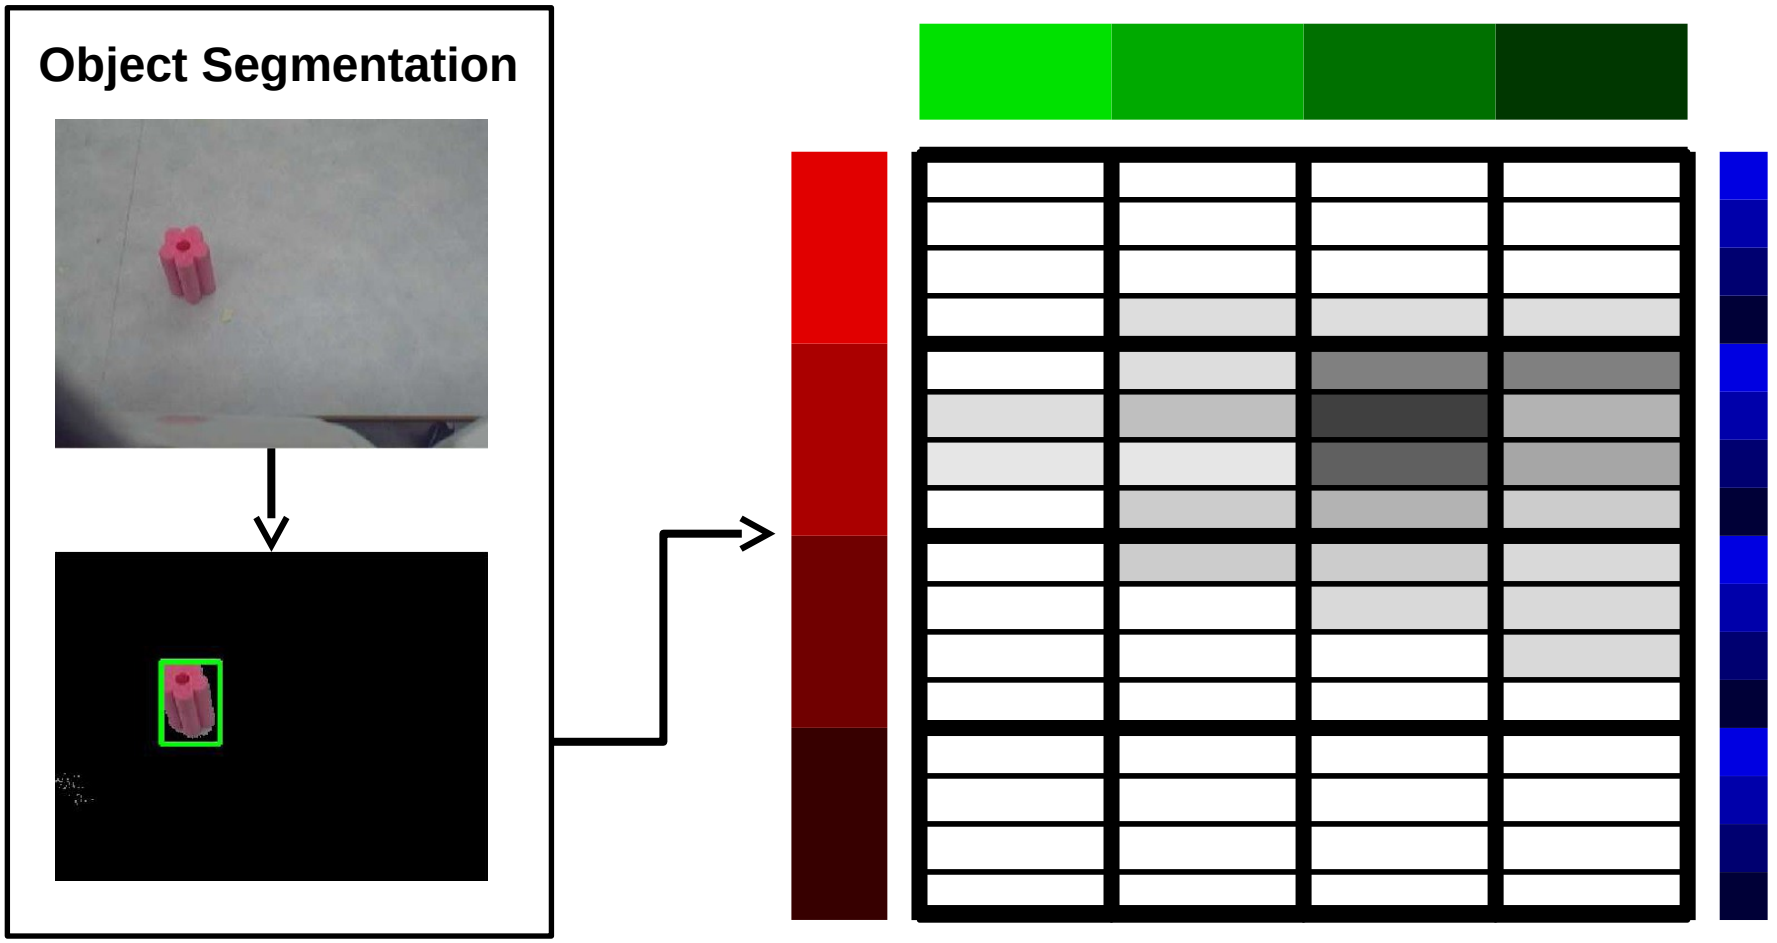

Motion

What is this?

What is this?

What action is being performed?

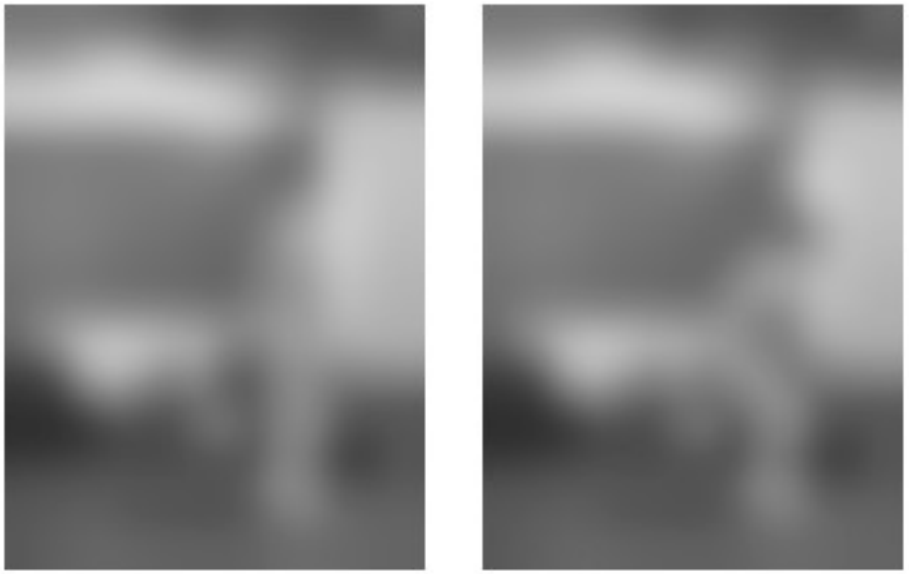

Frame 10

20

Frame 30

40

Motion Energy Image (MEI)

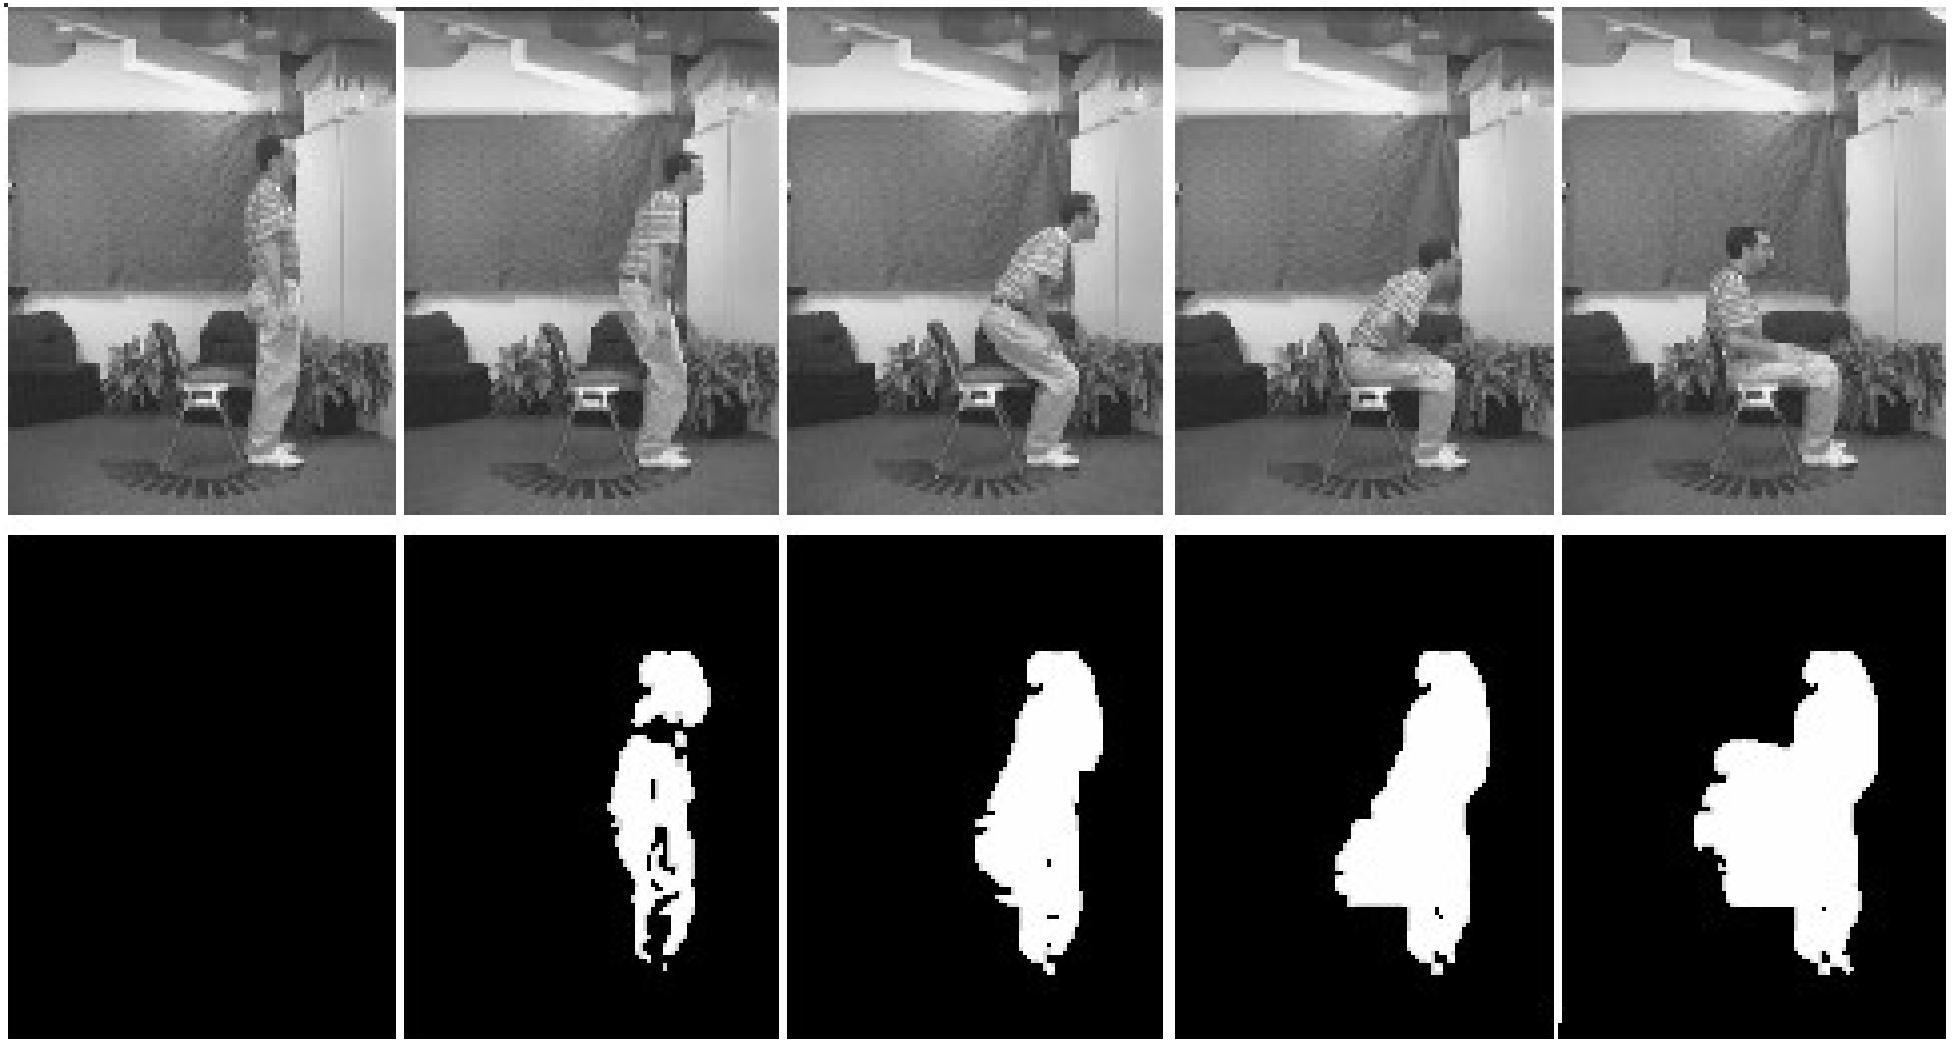

[http://www.cse.ohio-state.edu/~jwdavis/CVL/Research/MHI/mhi.html]

Average MEI for various viewing angles

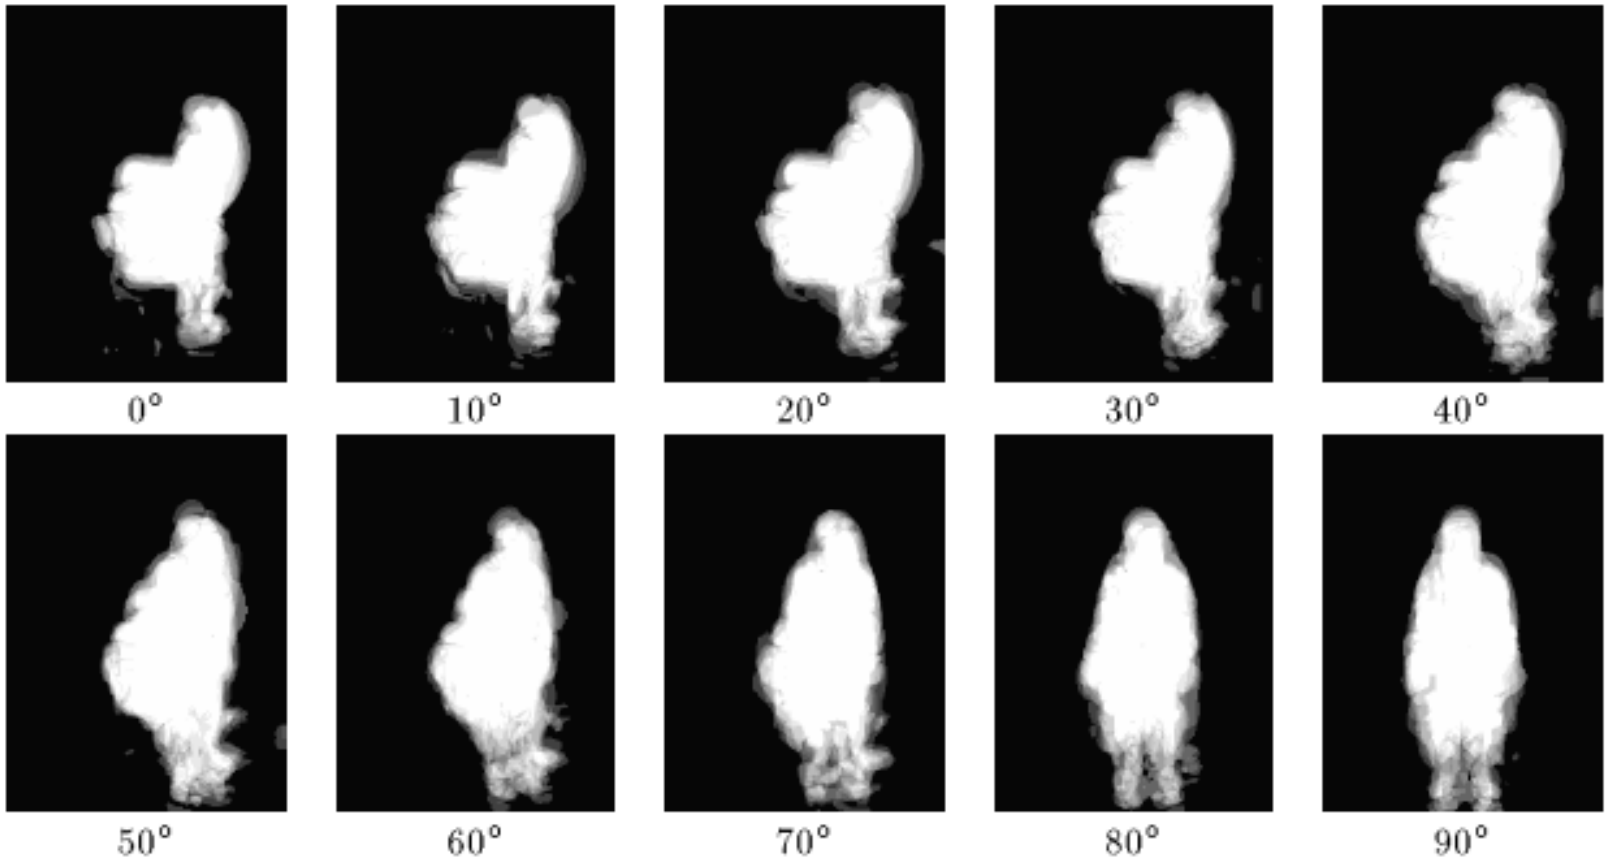

Motion History Image (MHI)

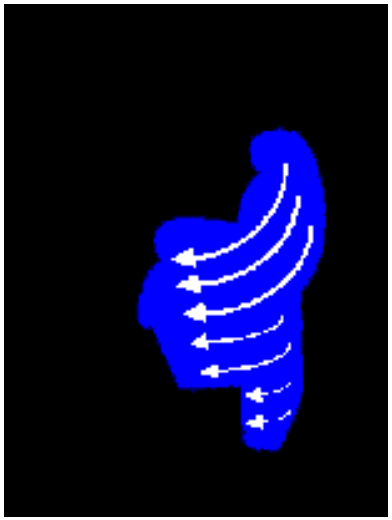

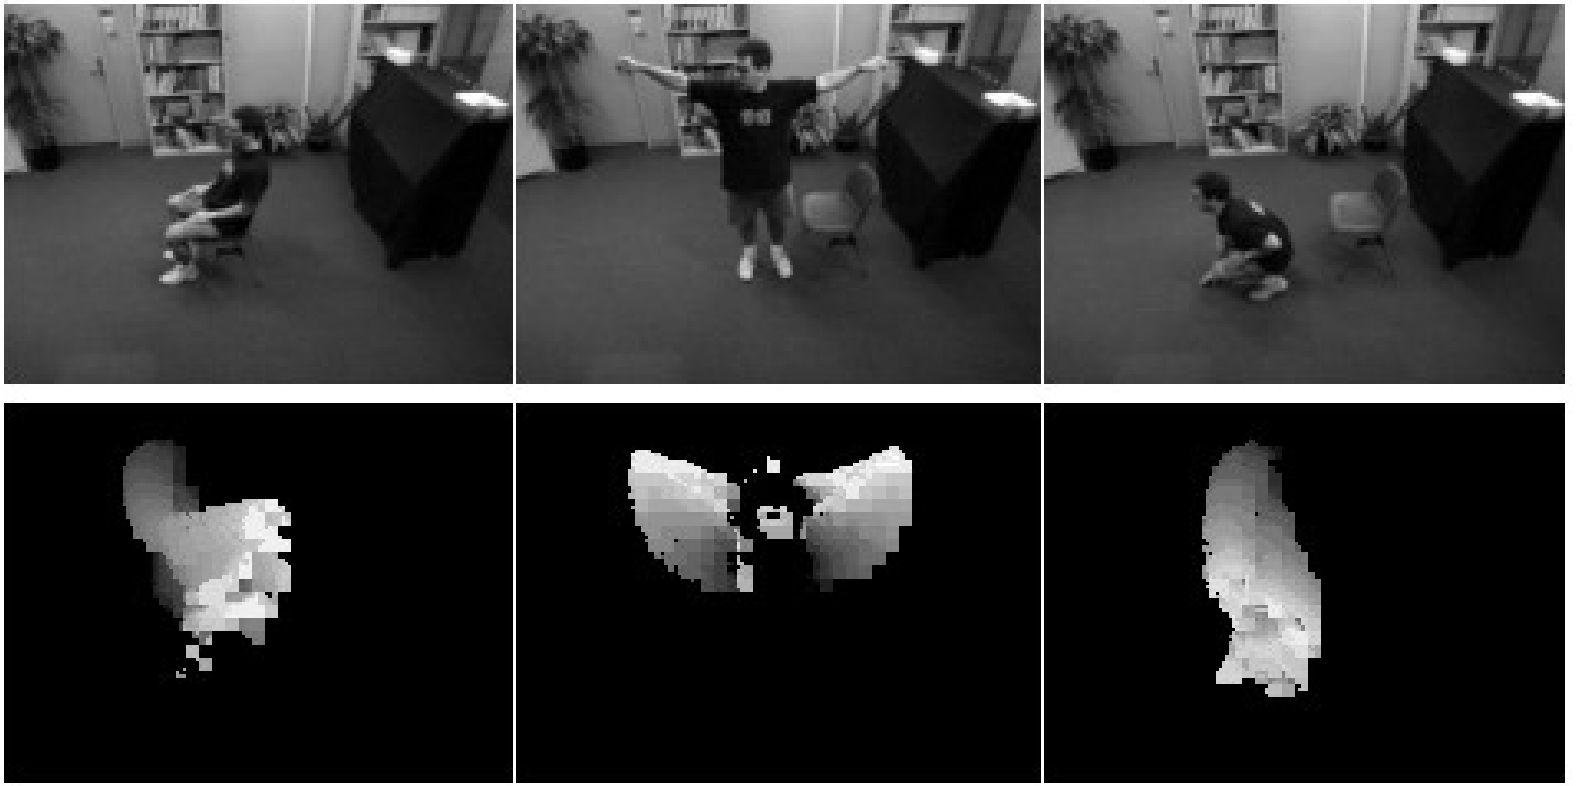

[http://www.cse.ohio-state.edu/~jwdavis/CVL/Research/MHI/mhi.html]

Definitions

- Image Sequence $I(x, y, t)$
- Binary Images
  indicating regions of motion $D(x, y, t)$
- Binary Motion Energy Image $E_{\tau}(x, y, t)$

Motion Energy

$$E_\tau(x, y, t) = \bigcup_{i=0}^{\tau-1} D(x, y, t - i)$$

Motion History

$$H_\tau(x,y,t) = \begin{cases} \tau & \text{if } D(x,y,t) = 1 \\ \max\ (0, H_\tau(x,y,t-1) - 1) & \\ \text{otherwise} & \end{cases}$$

The result: more recently moving pixels appear brighter

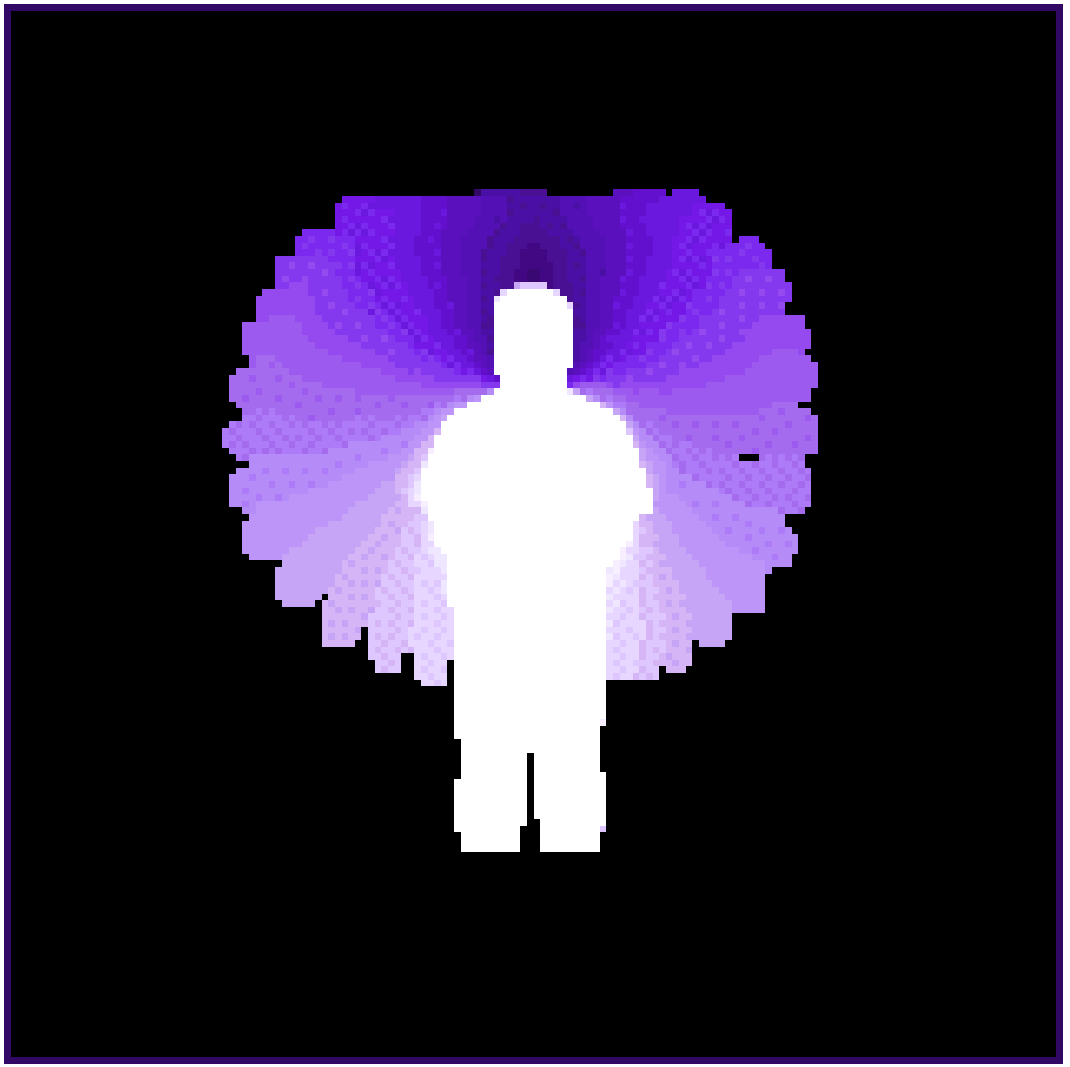

[http://www.cse.ohio-state.edu/~jwdavis/CVL/Research/MHI/mhi.html]

MHI pyramid

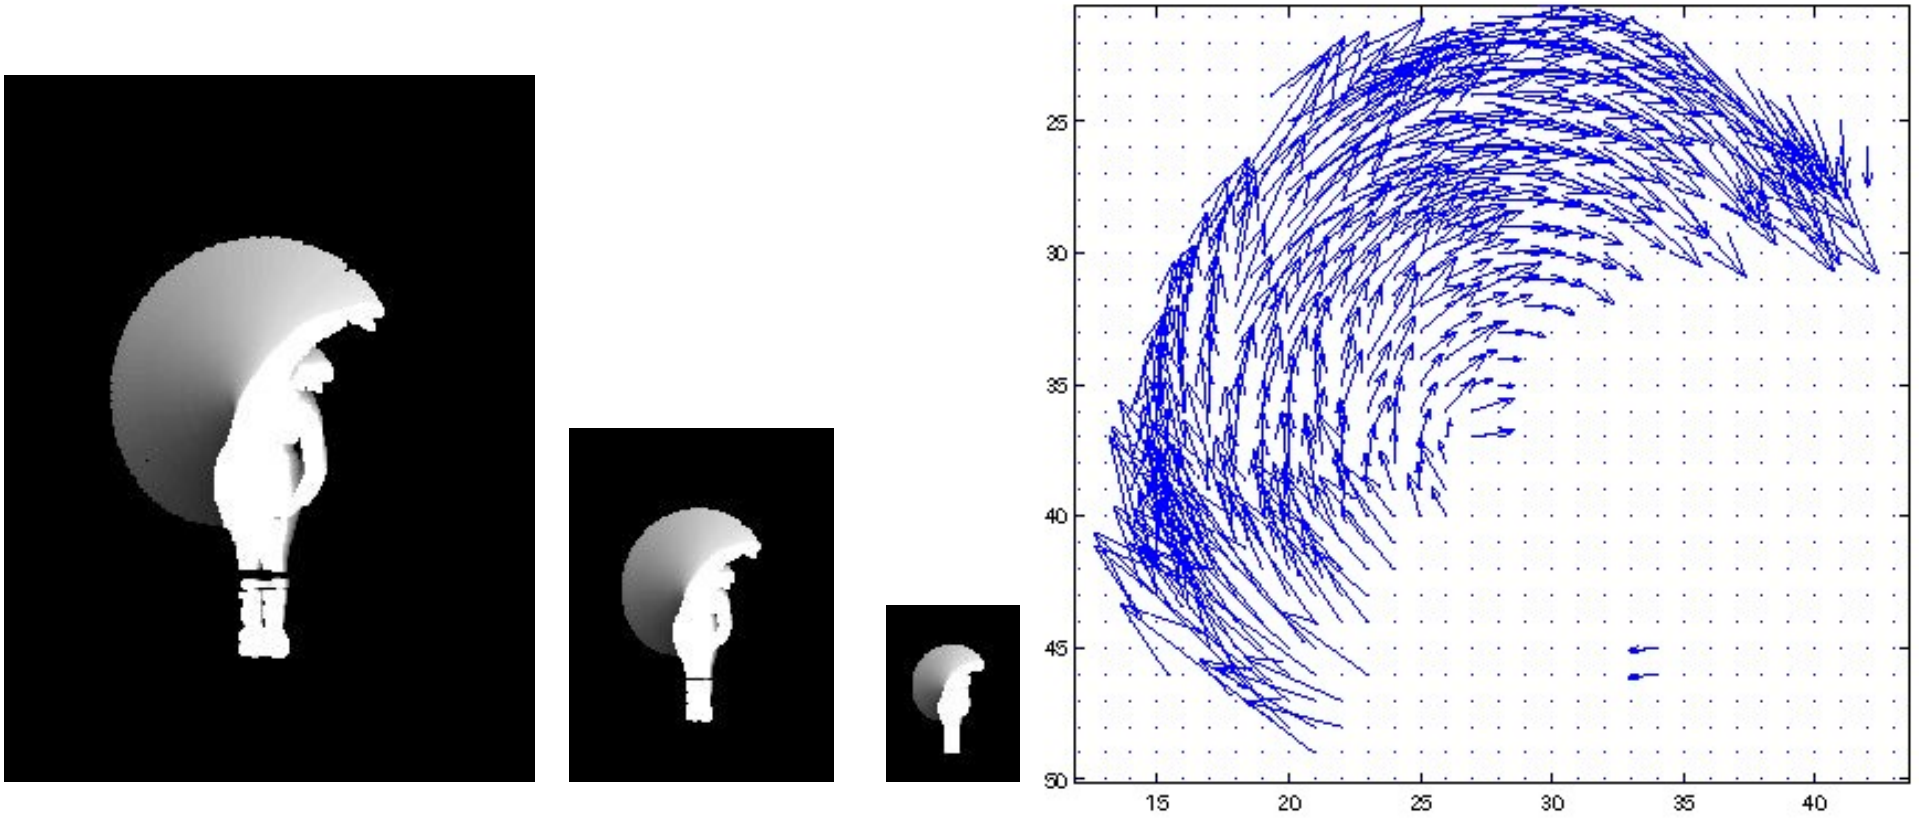

[http://www.cse.ohio-state.edu/~jwdavis/CVL/Research/MHI/mhi.html]

Motion templates for finishing LEFT-ARM-RAISE and FAN-UP-ARMS.

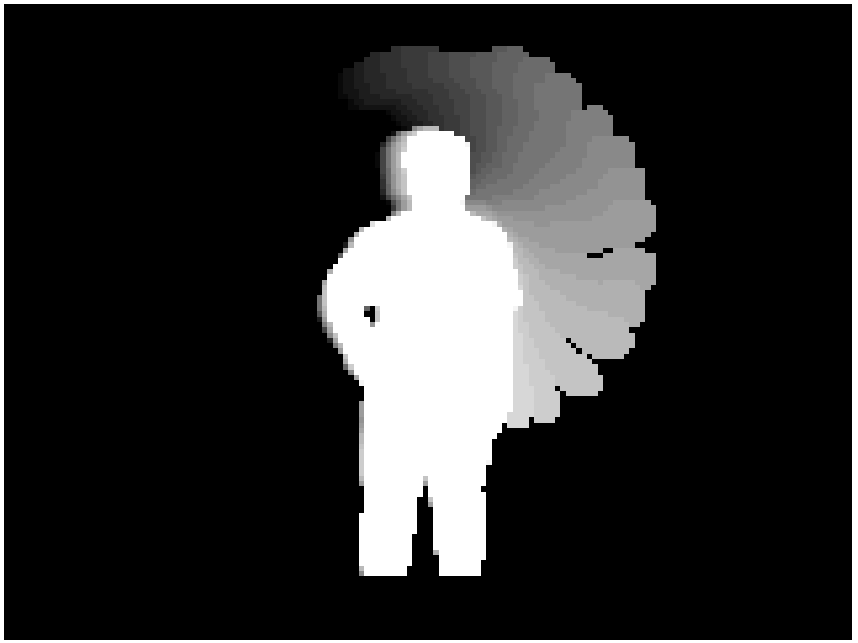

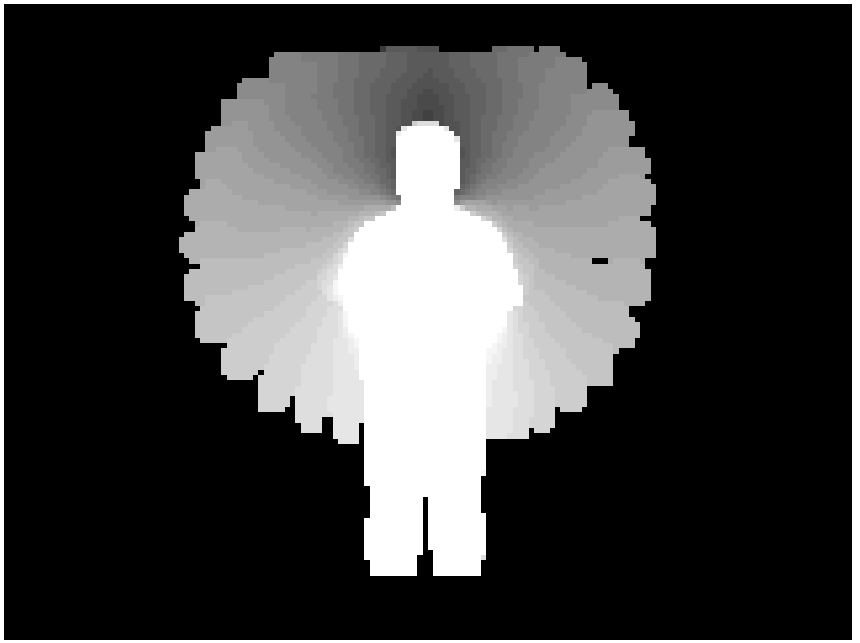

[http://www.cse.ohio-state.edu/~jwdavis/CVL/Research/VirtualAerobics/aerobics.html]

sit-down

sit-down MHI

arms-wave

arms-wave MHI

crouch-down

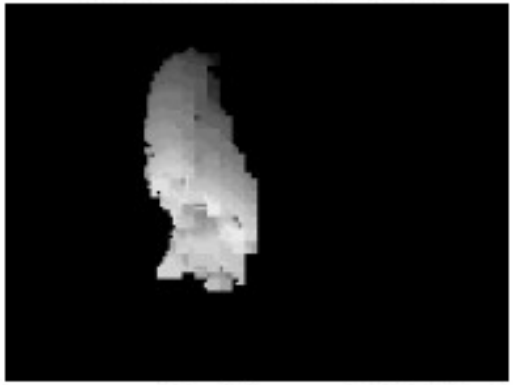

crouch-down MHI

Aerobics Dataset

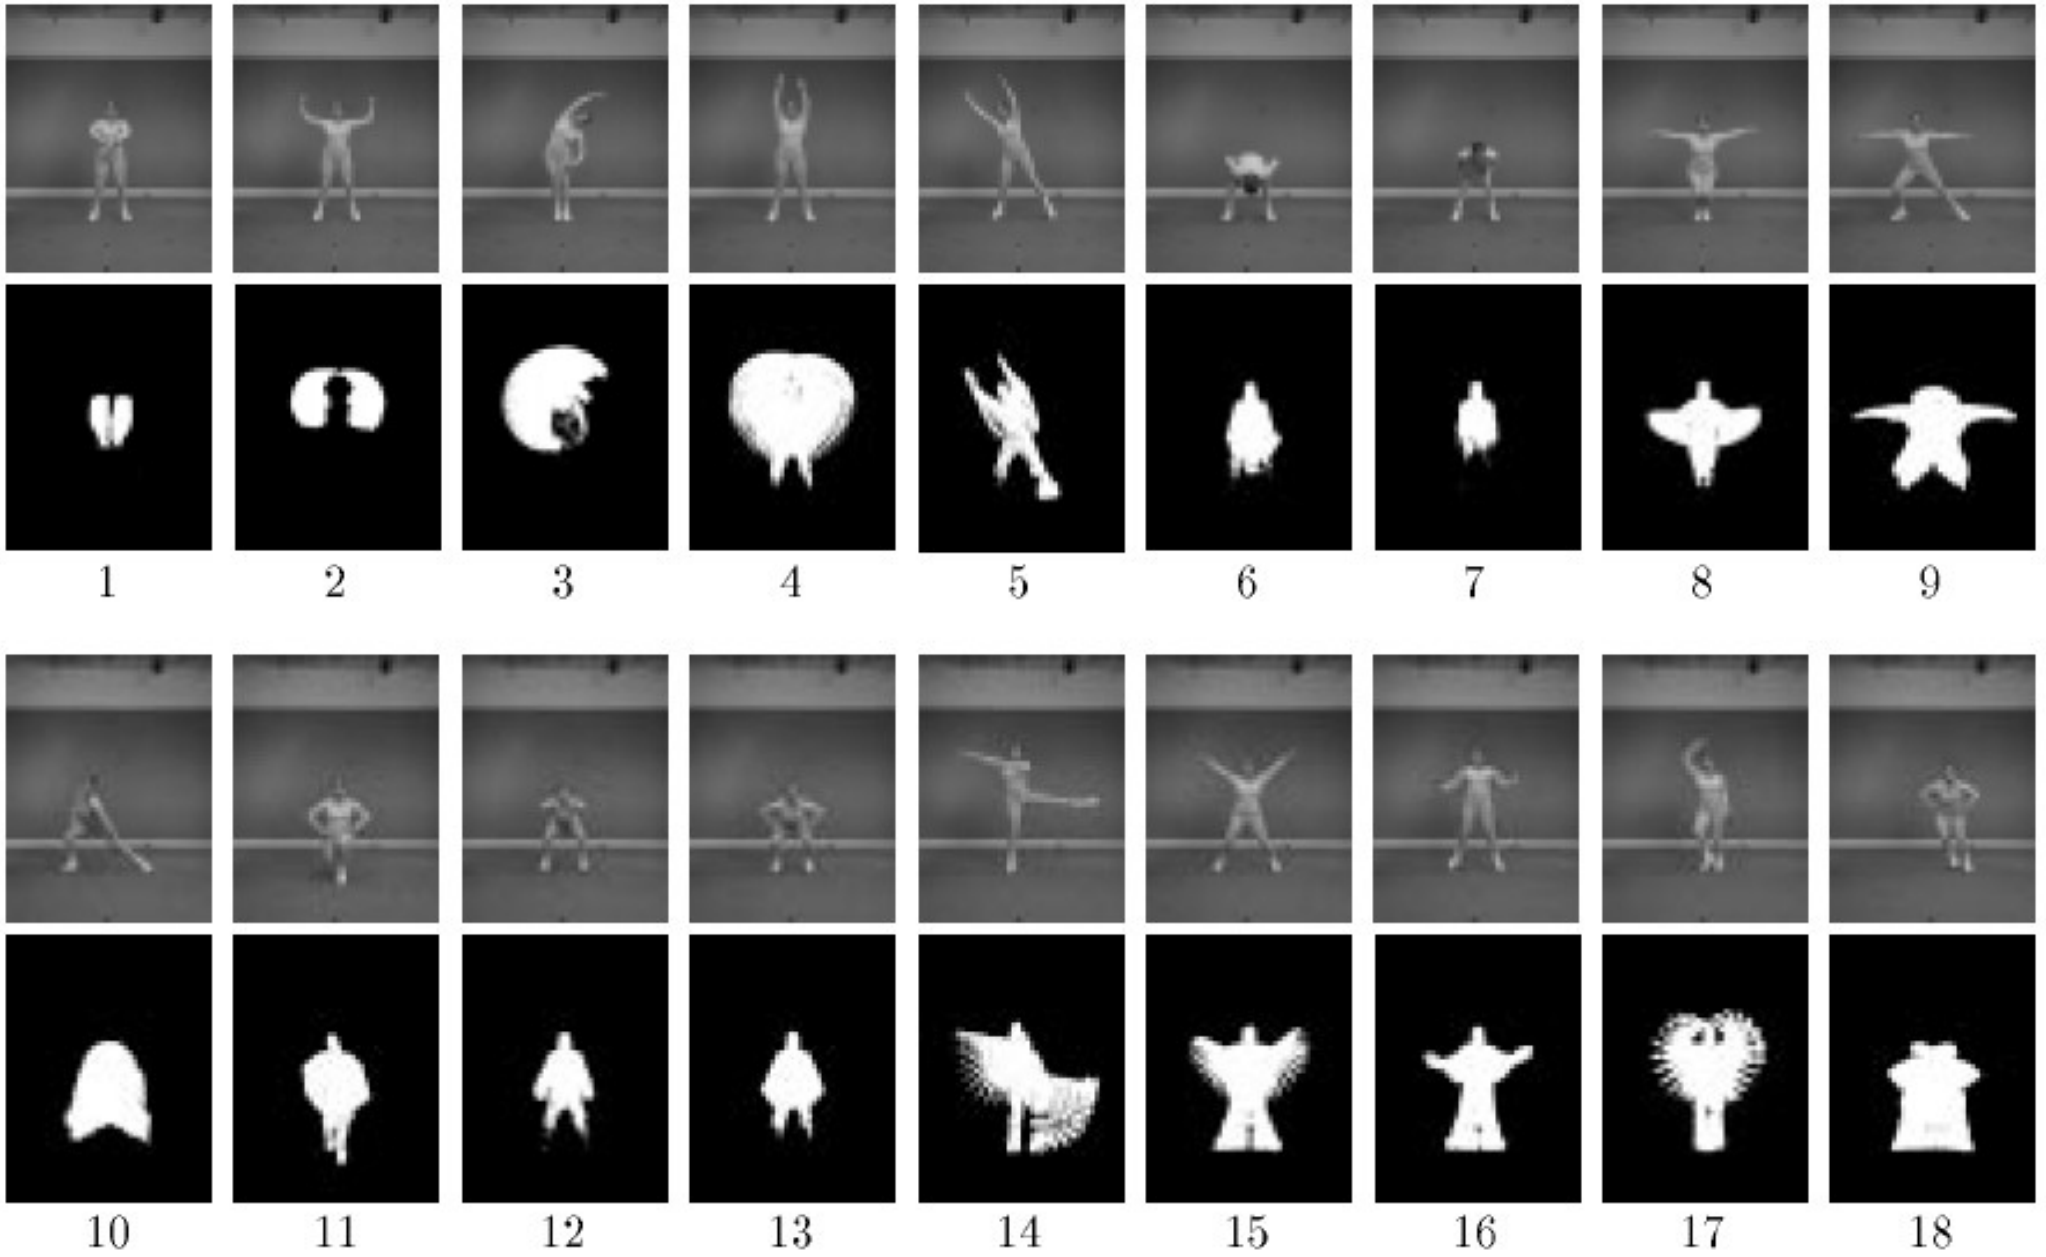

Video

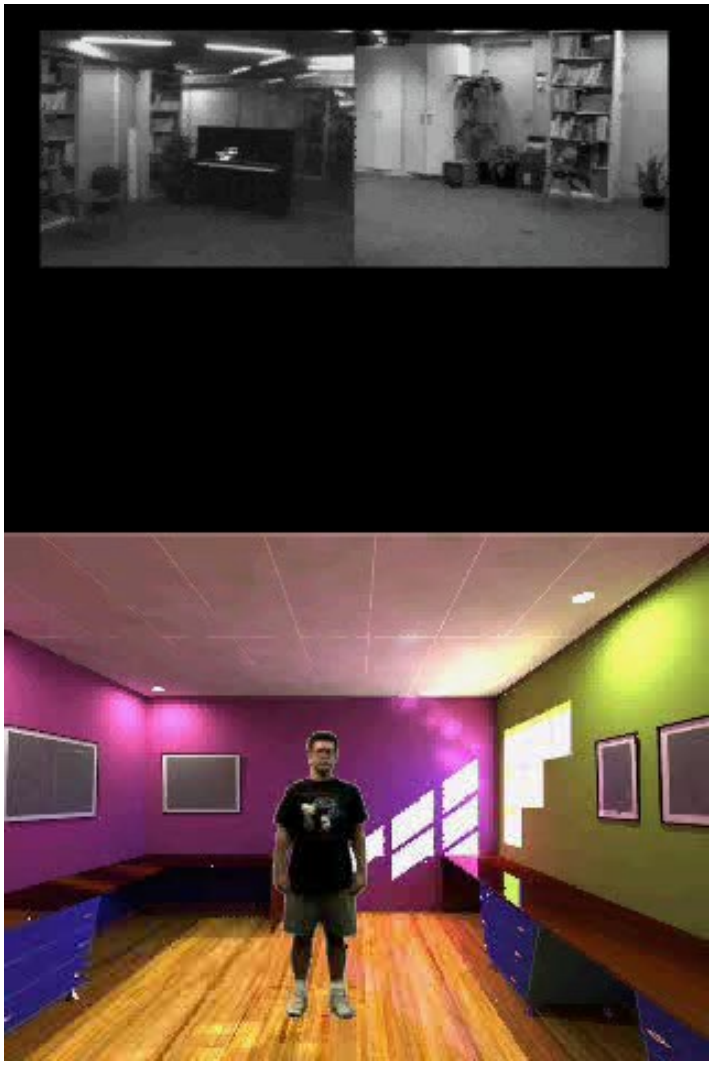

A. Bobick, S. Intille, J. Davis, F. Baird, C. Pinhanez, L. Campbell, Y. Ivanov, A. Schutte, and A. Wilson (1999)

``The Kidsroom: A Perceptually-Based Interactive and Immersive Story Environment"

Presence: Teleoperators and Virtual Environments, Vol. 8, No. 4, 1999, pp. 367-391.

The Kid’s Room

[Bobick et al. 1996]

The Blue Monster

[http://vismod.media.mit.edu/vismod/demos/kidsroom/kidsroom.html]

The Technology

[http://vismod.media.mit.edu/vismod/demos/kidsroom/kidsroom.html]

Motion History Templates

Making a ‘Y’

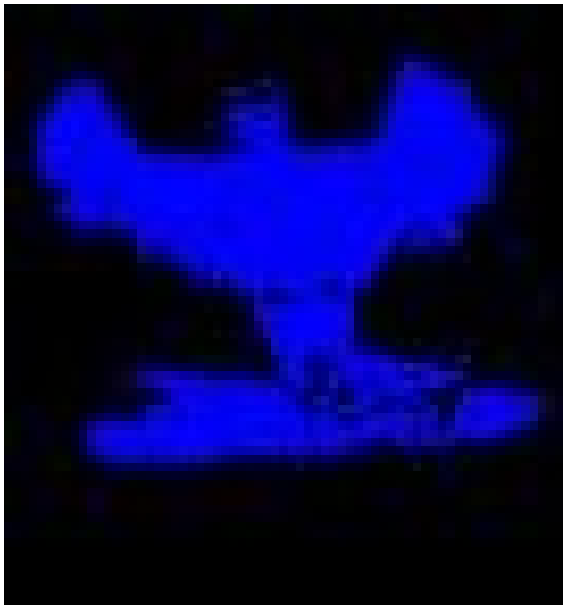

Flapping

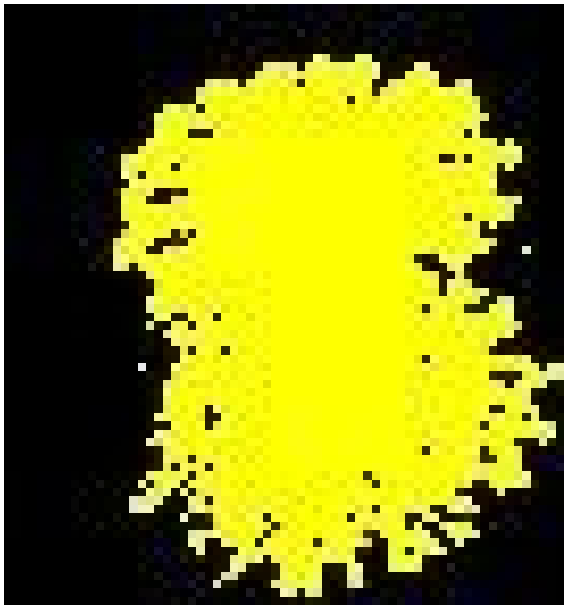

Spinning

[http://vismod.media.mit.edu/vismod/demos/kidsroom/kidsroom.html]

Detecting the Bed

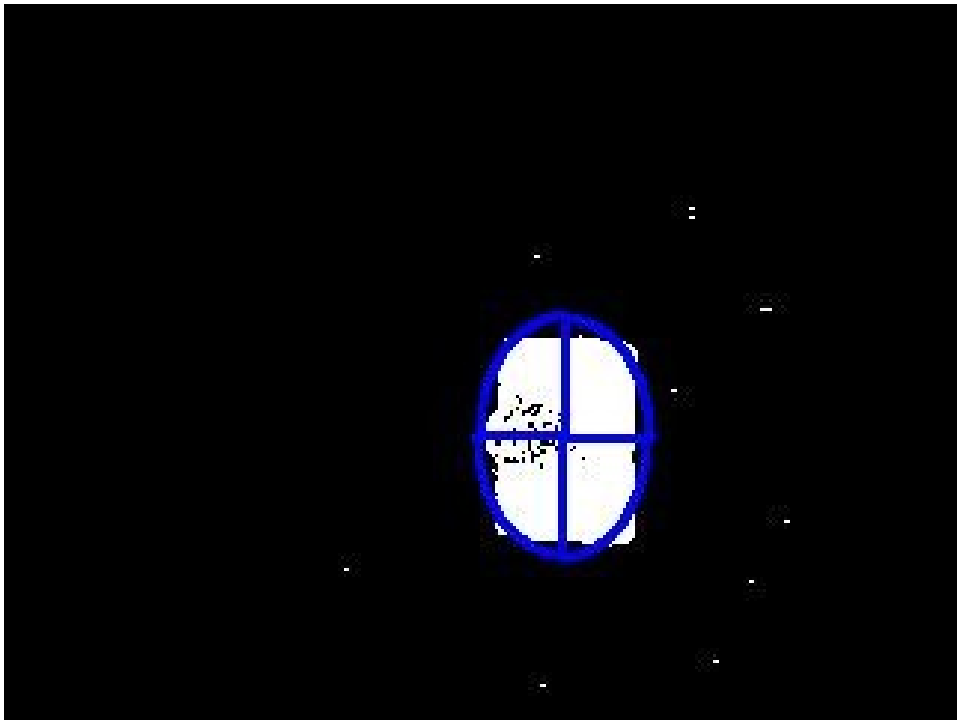

[http://vismod.media.mit.edu/vismod/demos/kidsroom/kidsroom.html]

Man Overboard Detector

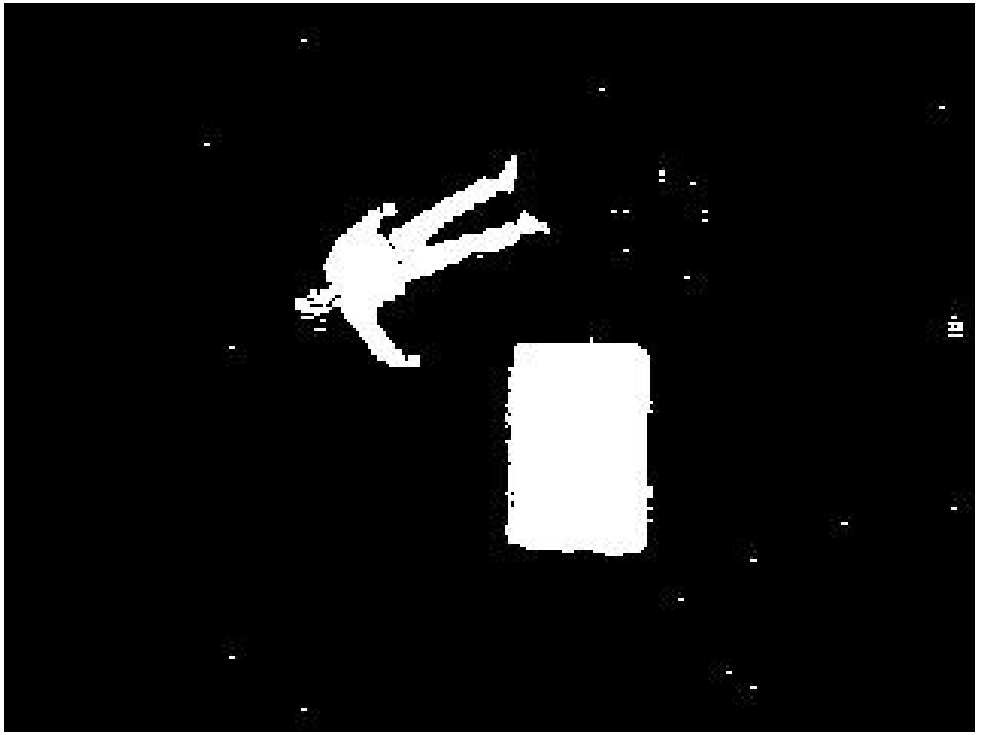

[http://vismod.media.mit.edu/vismod/demos/kidsroom/kidsroom.html]

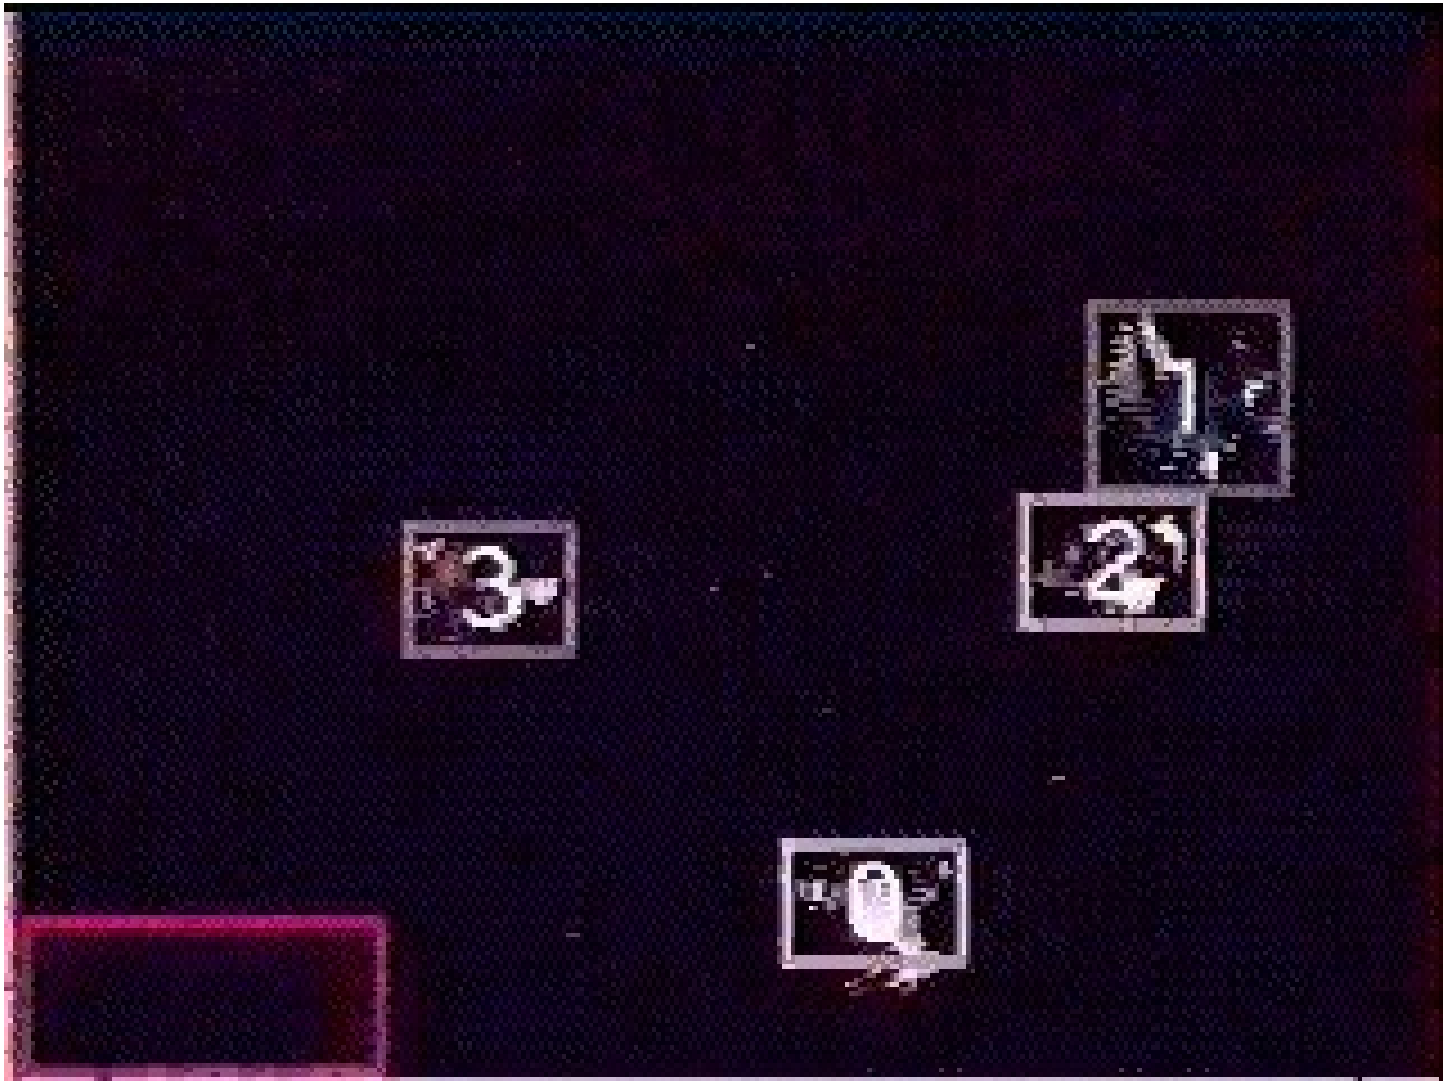
3
2

OpenCV Book and Code

- “Learning OpenCV”

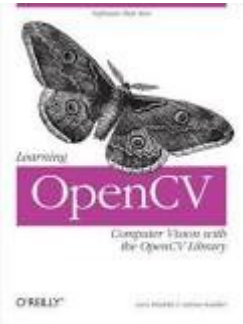

- Code from book is on github:

  https://github.com/Itseez/opencv_extra/tree/master/learning_opencv_v2

OpenCV Book and Code

- “Learning OpenCV”

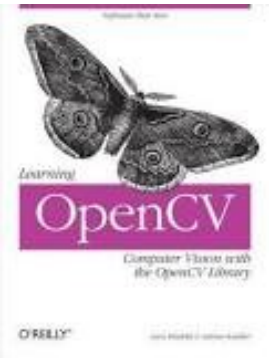

- Code from book is on github:

  https://github.com/Itseez/opencv_extra/tree/master/learning_opencv_v2

OpenCV Tutorials

- Connected Components:
  - http://nghiaho.com/?p=1102
  - https://davidlavy.wordpress.com/opencv/connected-components-in-opencv/

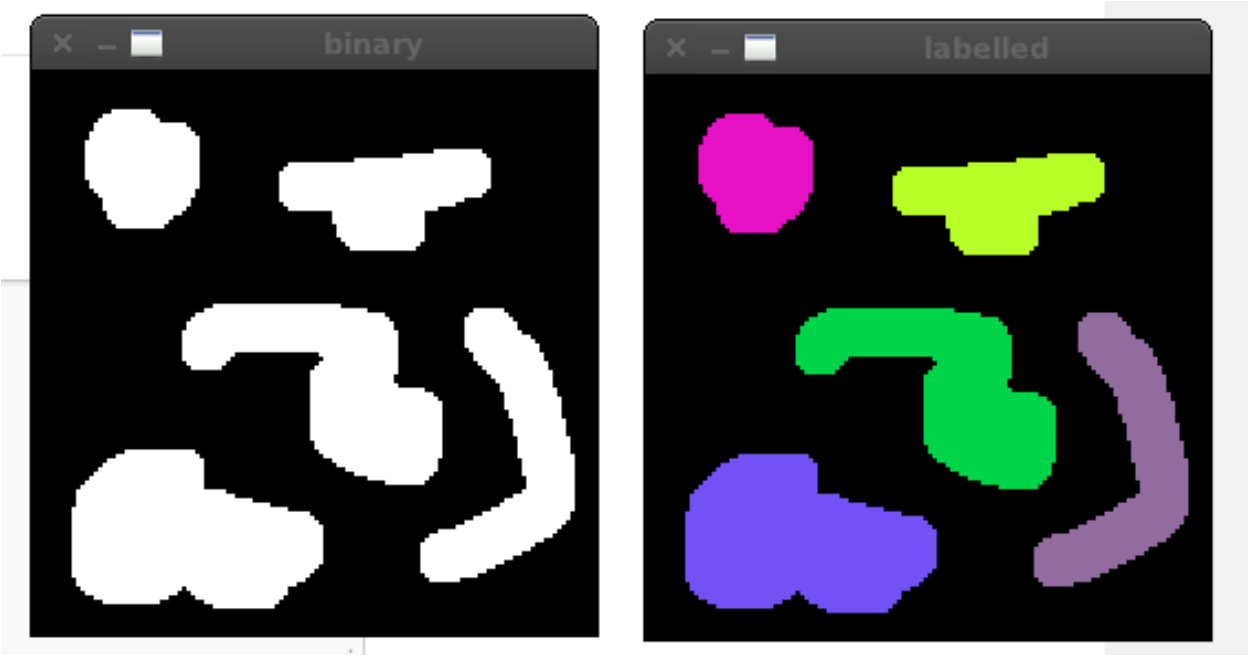

OpenCV Tutorials

- Circle Detection:
  - http://docs.opencv.org/3.1.0/d4/d70/tutorial_hough_circle.html#gsc.tab=0

OpenCV Tutorials

- Face Detection:
  - http://stackoverflow.com/questions/20757147/detect-faces-in-image
  - https://github.com/Itseez/opencv_extra/blob/master/learning_opencv_v2/ch13_ex13_4.cpp

OpenCV Tutorials

- Blog full of OpenCV examples:
  - http://opencvexamples.blogspot.com/

Resources

- OpenCV in ROS:
  - http://wiki.ros.org/vision_opencv
  - http://wiki.ros.org/cv_bridge/Tutorials
  - http://docs.opencv.org/2.4/doc/tutorials/tutorials.html

Grabbing image data with ROS

- Example ROS node that subscribes to an image topic and does image processing:

  http://www.cs.tufts.edu/comp/50AIR/code/comp50_computer_vision.zip

Next time...interest points and registration

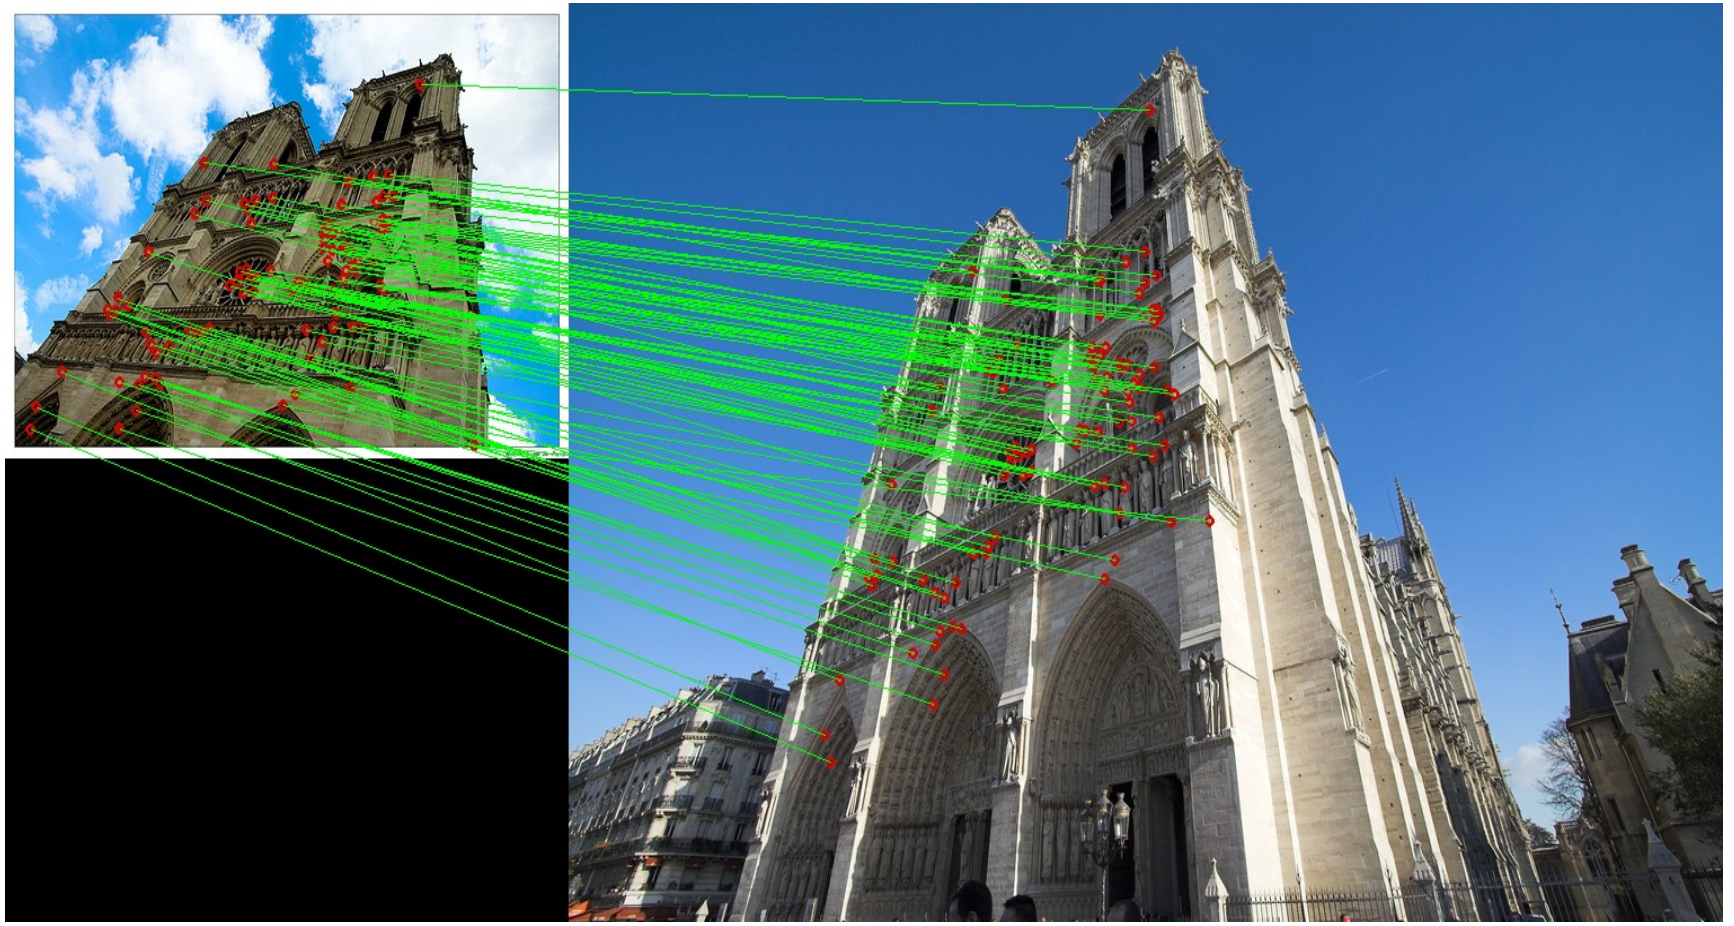

Later in the course...3D Vision

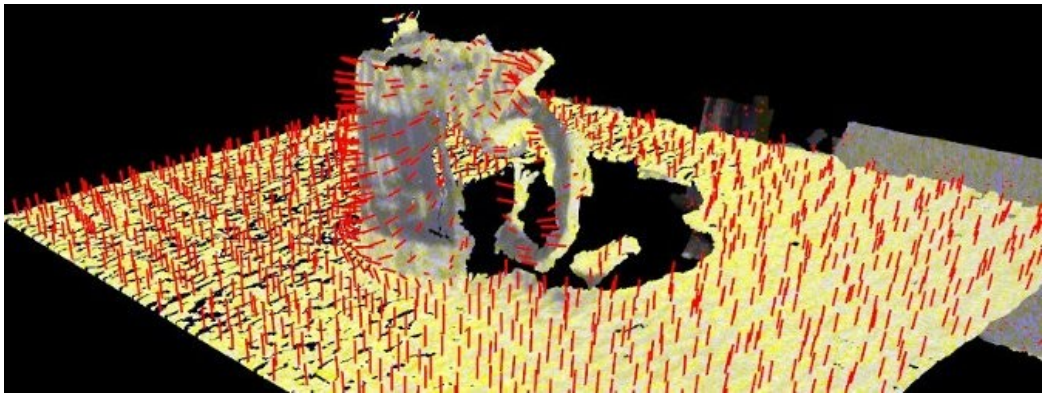

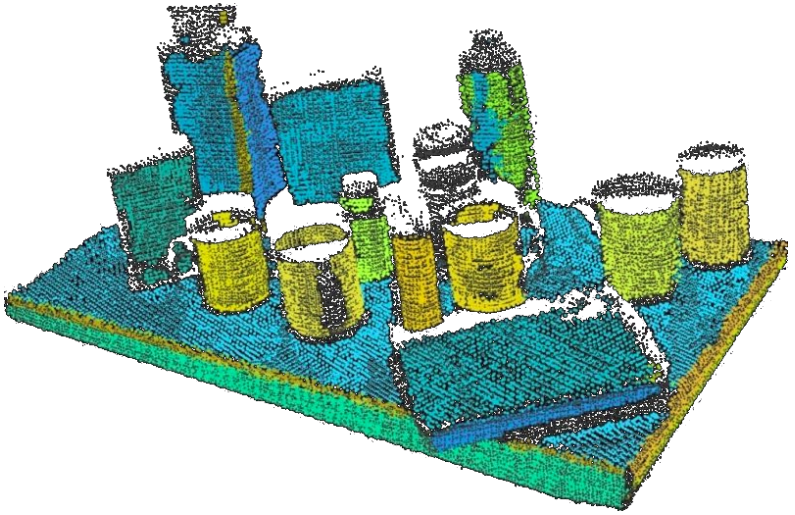

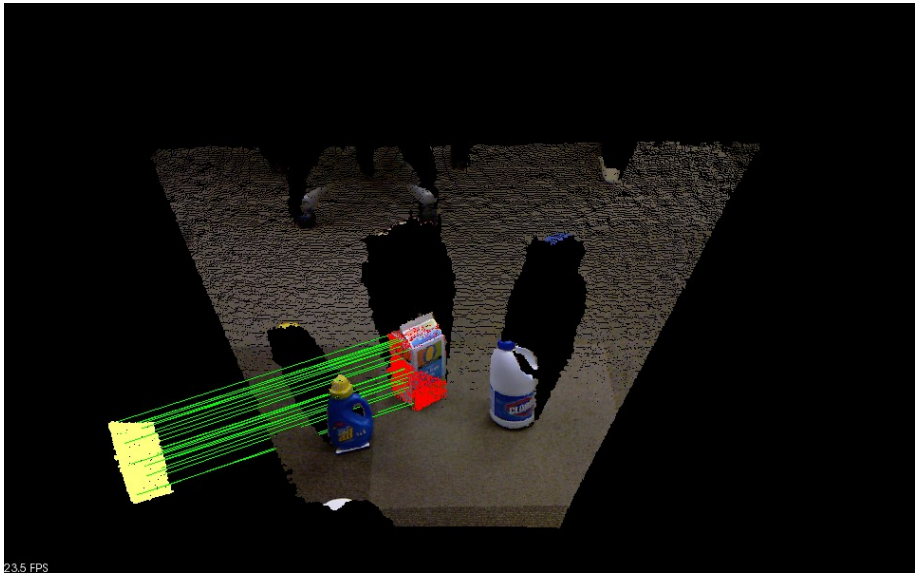

Project Activity

- Get with your group
- Sketch an outline of your proposal and assign individual responsibilities

THE END