CS 556: Computer Vision

Lecture 2

Prof. Sinisa Todorovic

sinisa@eecs.oregonstate.edu

OSU Oregon State University

Outline

- Properties of low-level image structure
- From image formation to image interpretation
- Computational paradigms of computer vision

Images are not Collections of Random Pixels

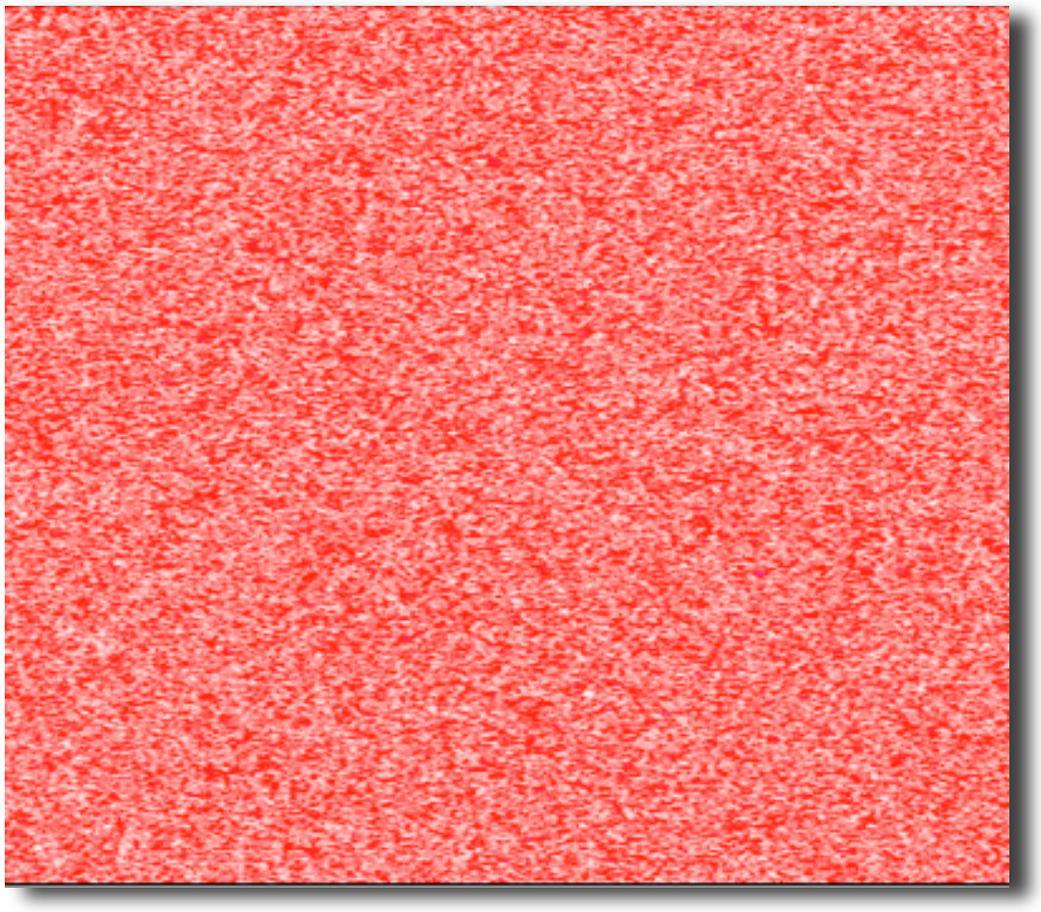

Images are characterized

by structure

Image Structure

- Pixels, 4-adjacency, 8-adjacency, m-adjacency

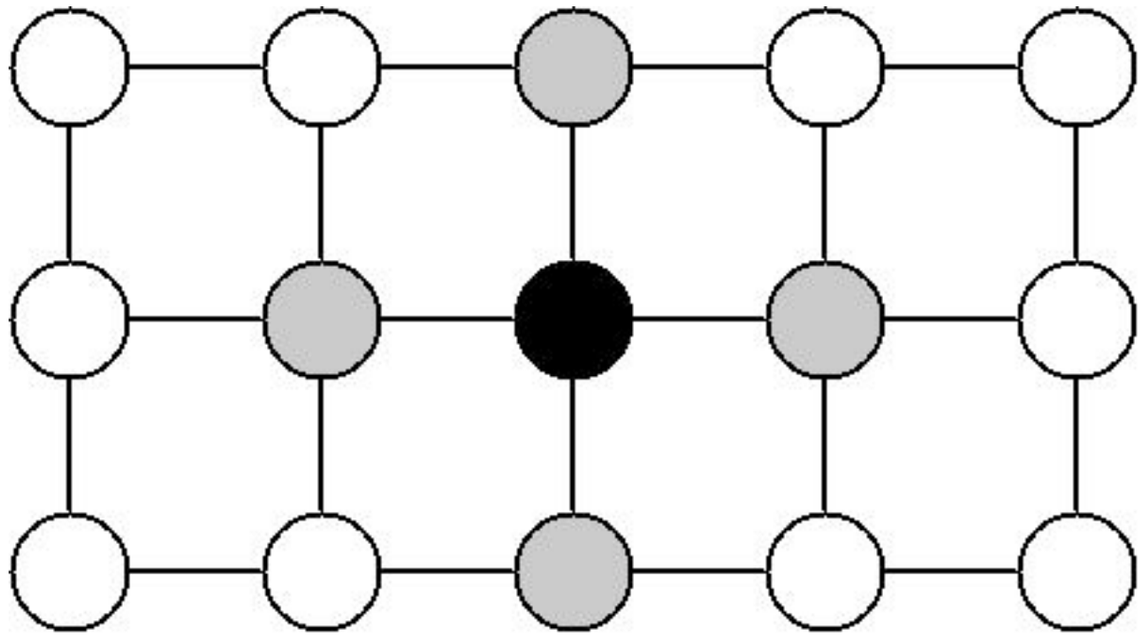

4-adjacency

Image Structure

- Path -- directed, undirected, loop

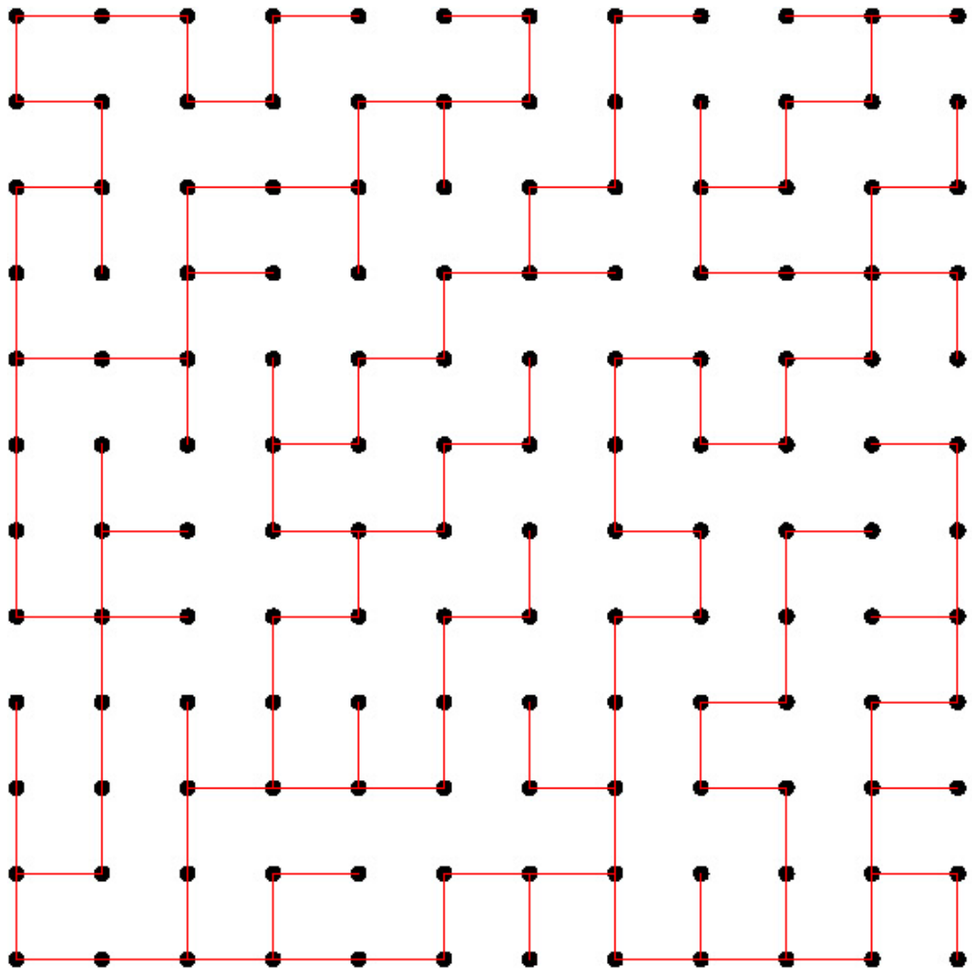

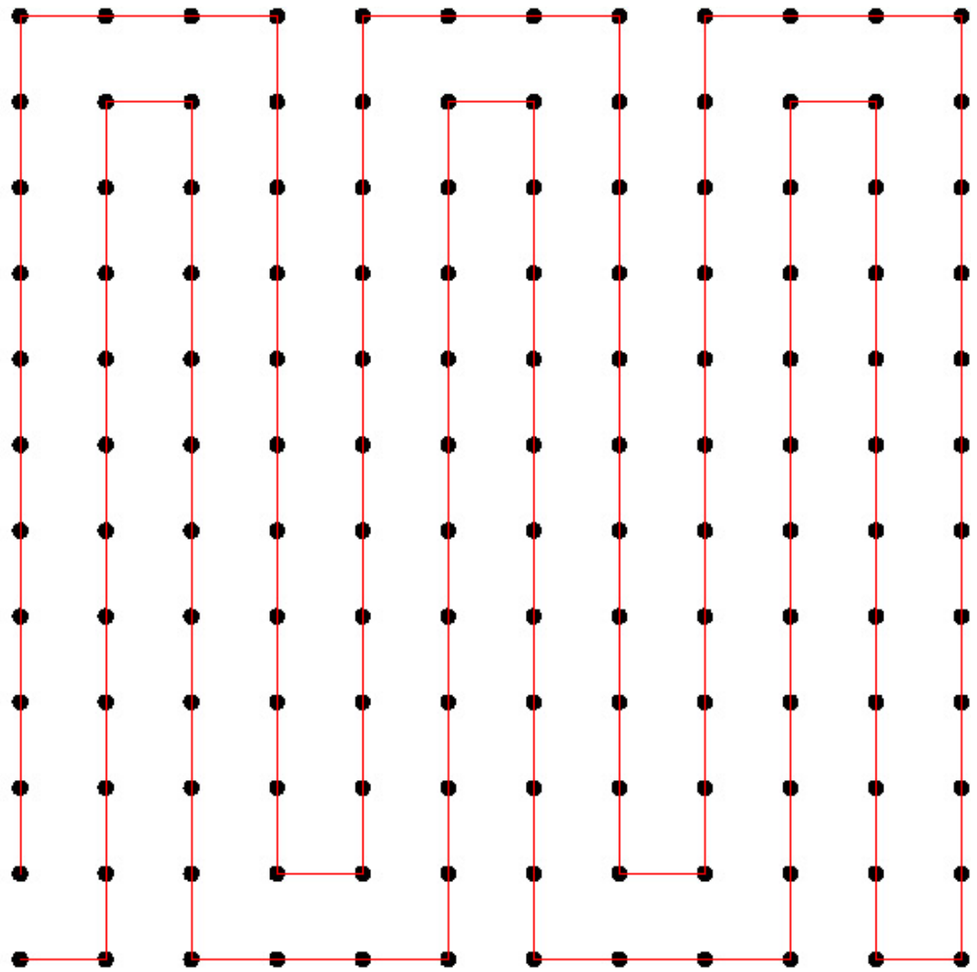

Image Structure

- Edge = Connected pixels with high gradient values

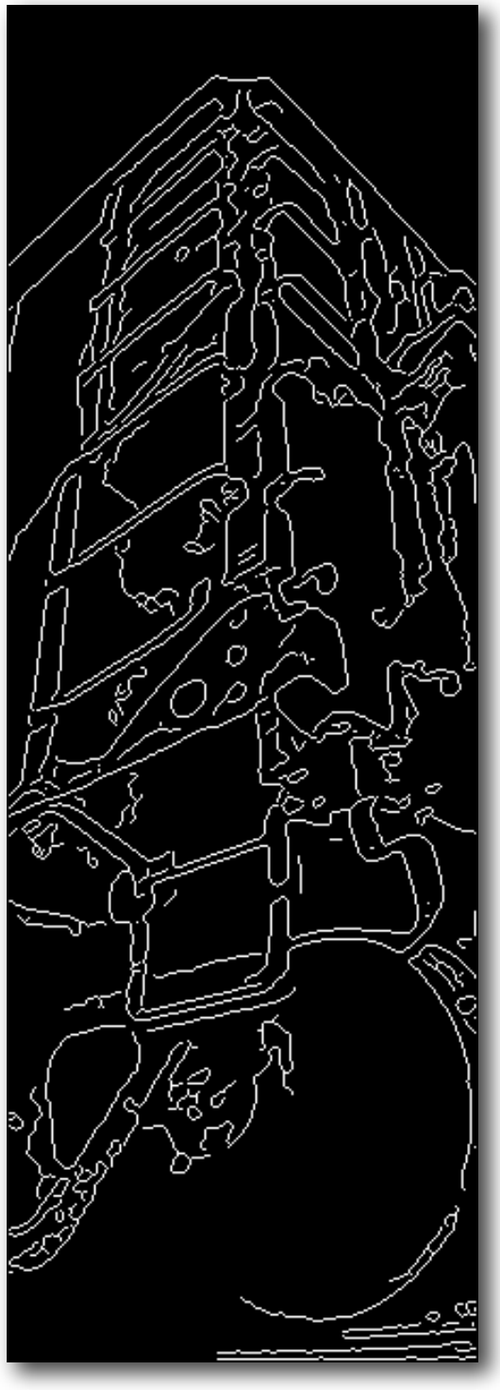

Image Structure

- T-junctions

Image Structure

- T-junctions

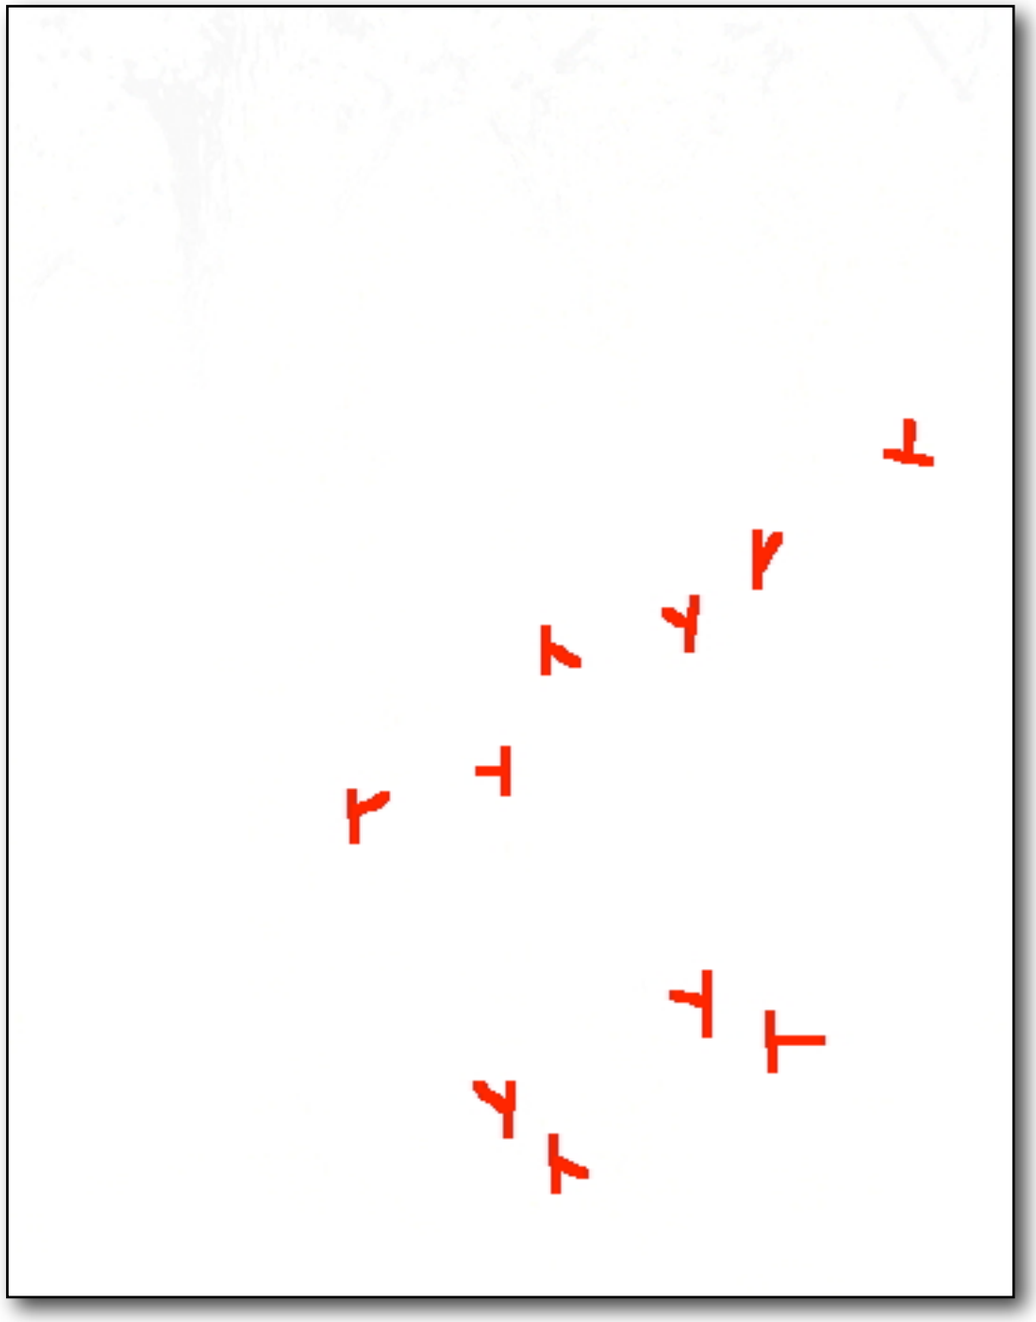

Image Structure

- T-junctions

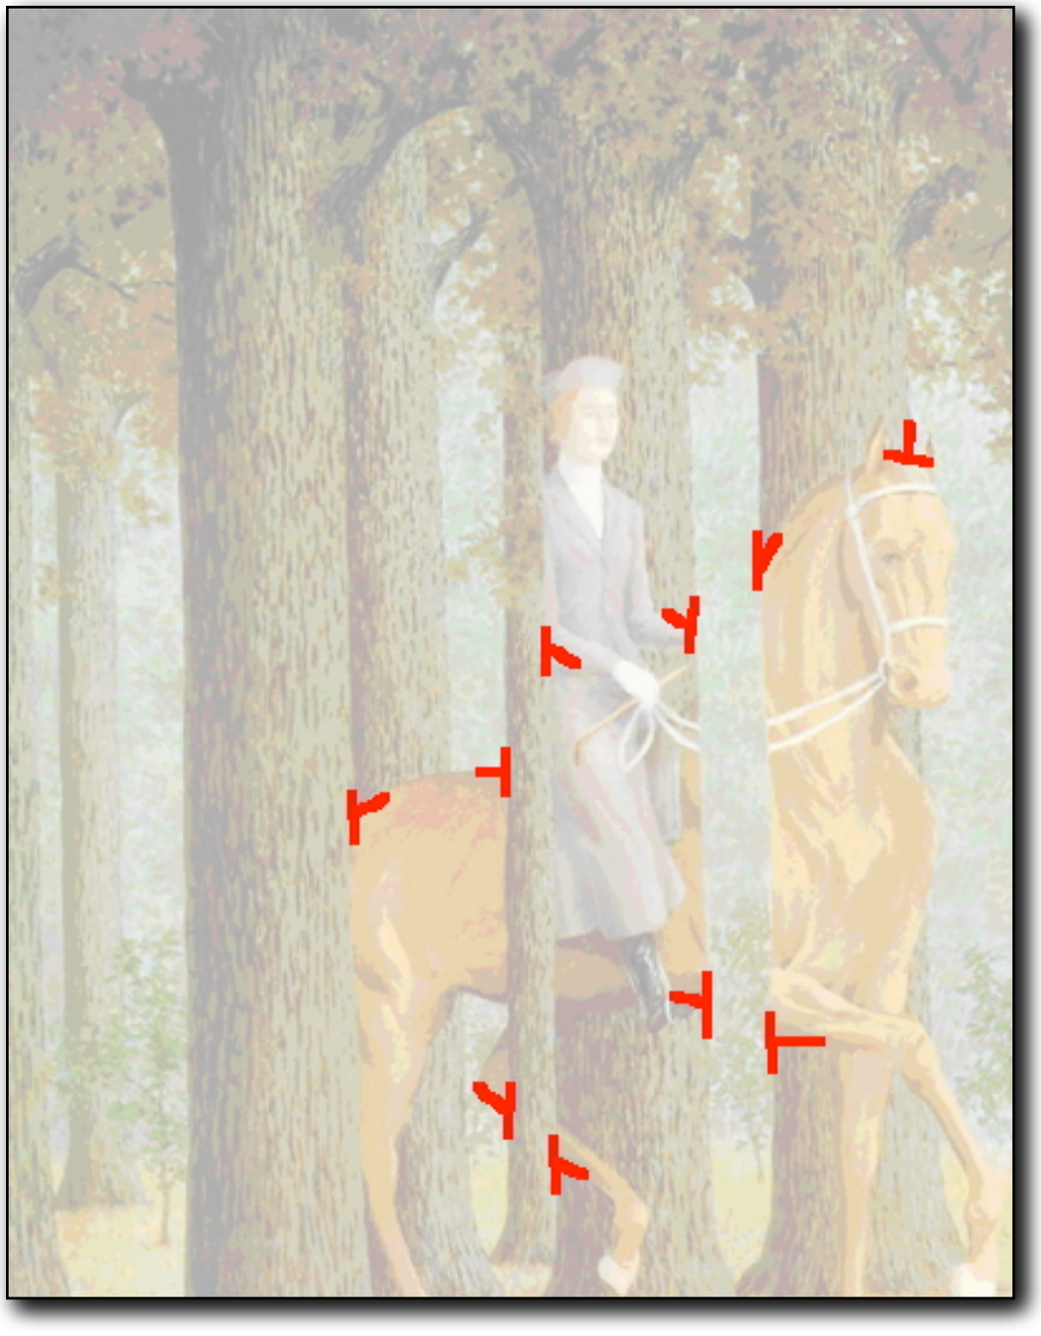

Image Structure

- Interest points = corners, textured patches

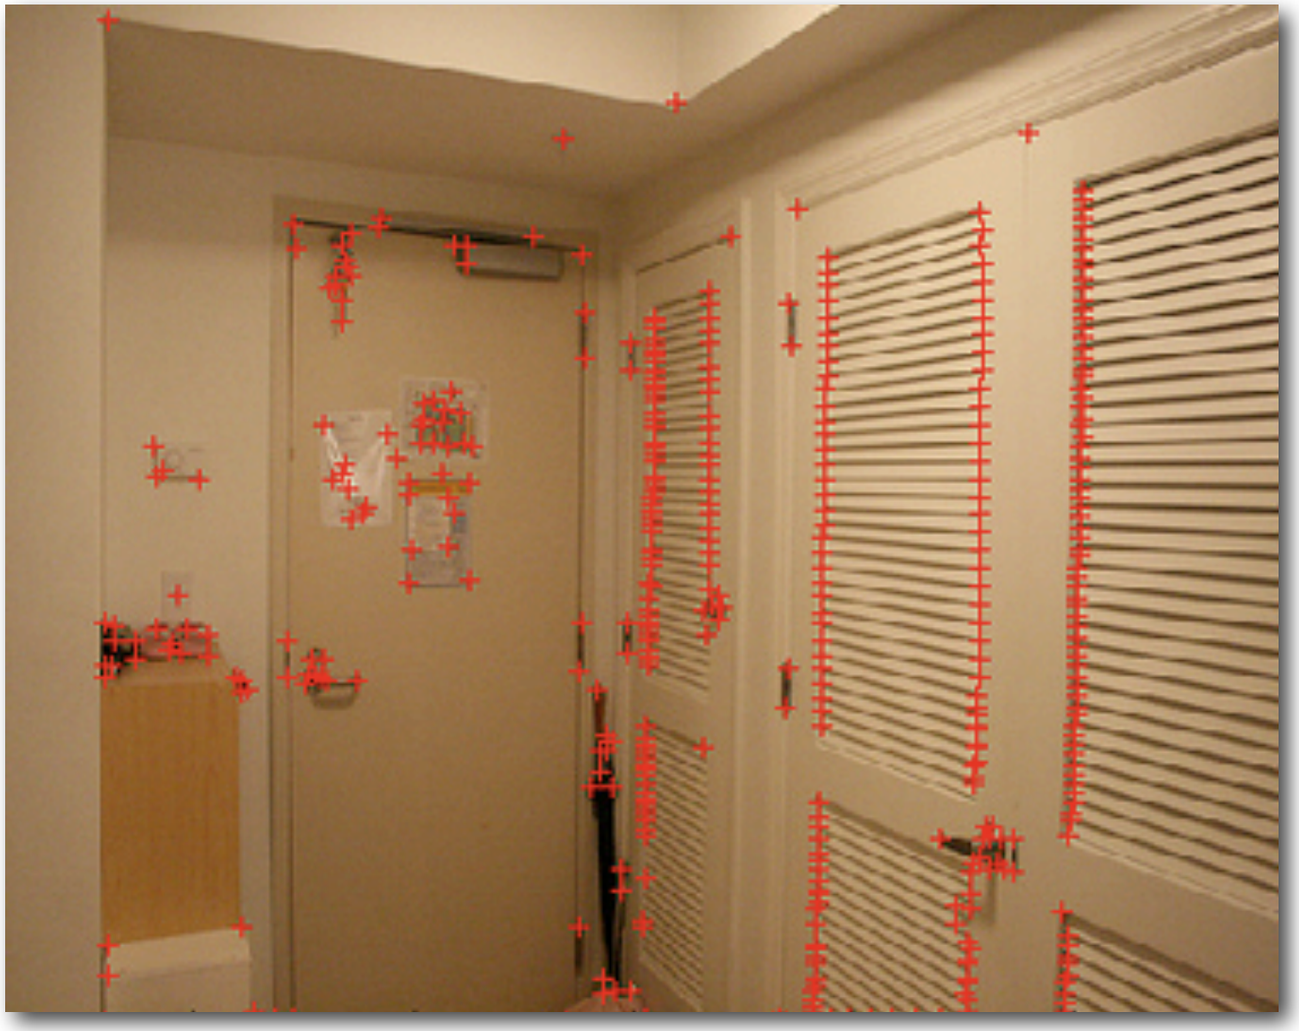

Image Structure

- Specularity = Highlights

Image Structure

- Lambertian surface = isotropic reflectance
- Specular surface = zero reflectance except at an angle

Image Structure

- Region = Connected set of pixels
- Region boundary, inner and outer contour

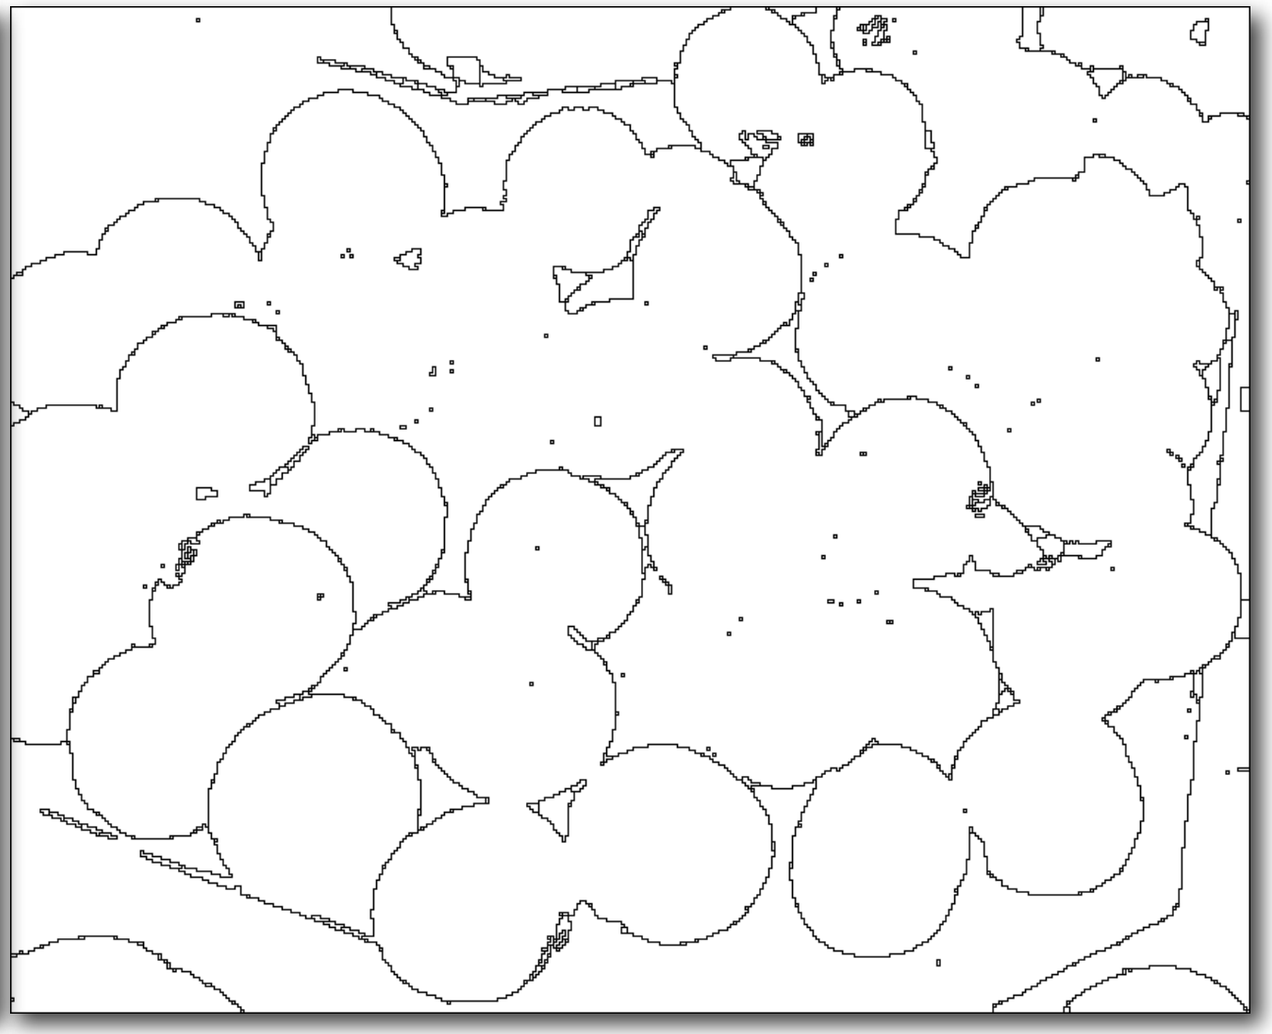

Image Structure

- Texture

Image Structure

- Dynamic texture

Image Structure

- Foreground - Background

From Image Formation to Image Interpretation

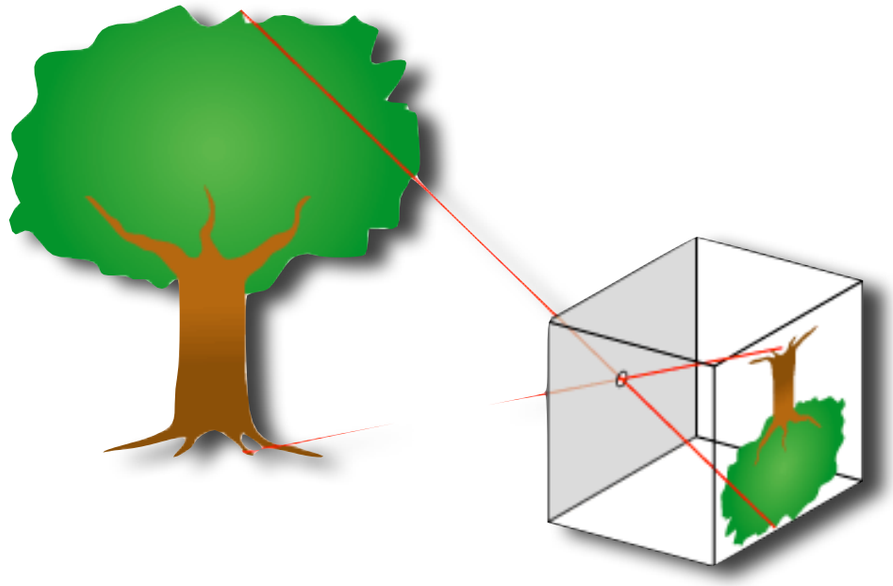

From Image Formation to Image Interpretation

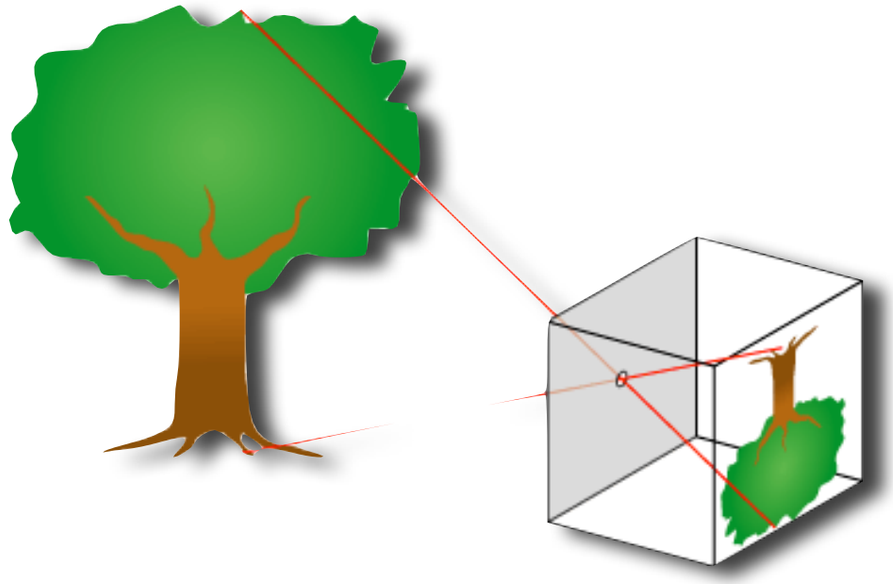

Image interpretation is usually defined in terms of

- Objects, scenes, activities, events
- Spatiotemporal relations between objects, scenes, ...

From Image Formation to Image Interpretation

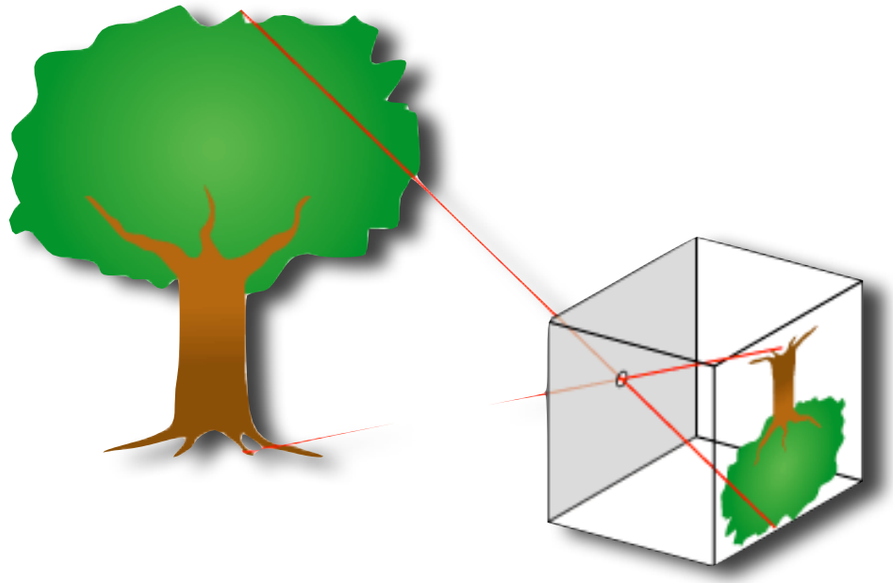

Image properties depend on:

- Imaging conditions
- Visual properties of the 3D world

From Image Formation to Image Interpretation

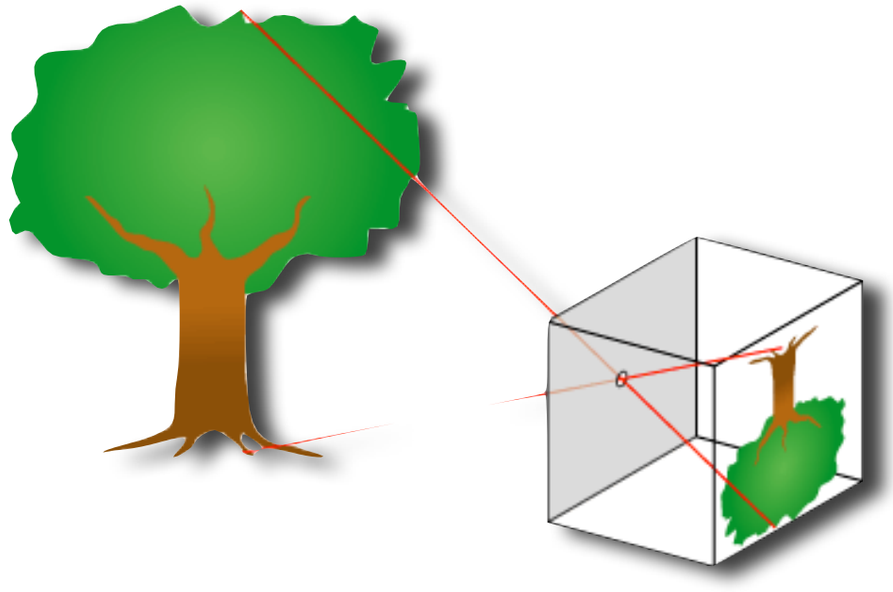

- Image acquisition parameters:
- Camera distance, viewpoint, motion
- Camera intrinsic parameters (e.g., lens)
- Illumination or Brightness
- Occlusion and Clutter

Are there general constraints

the real world poses

on image properties?

Properties of Real World Objects

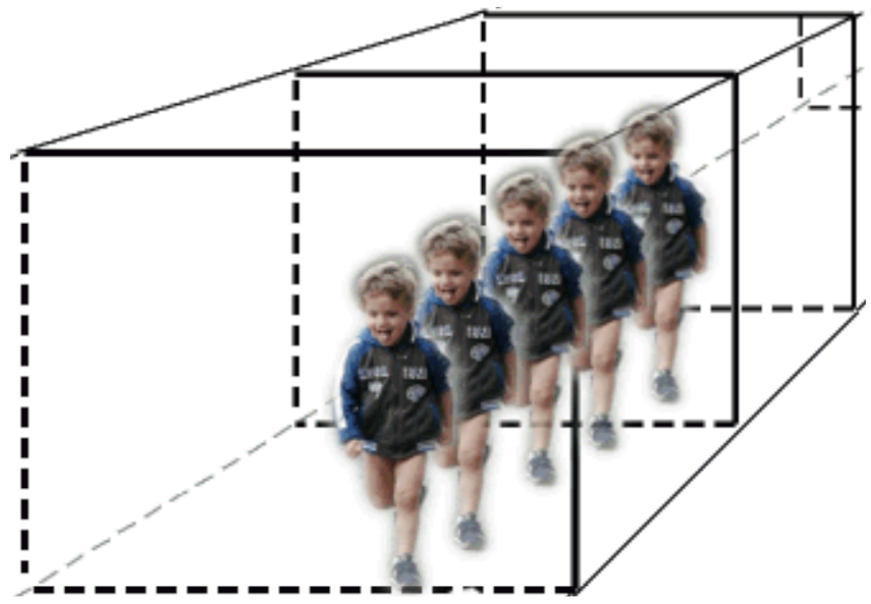

Sources: N. Ahuja, Pritch et al 07

Properties of Real World Objects

- Occupy finite 3D volume

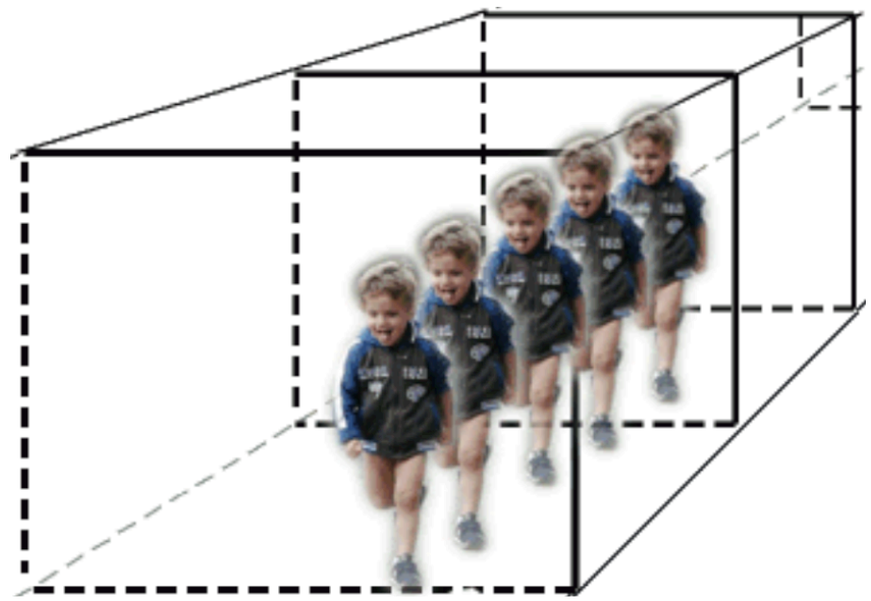

Sources: N. Ahuja, Pritch et al 07

Properties of Real World Objects

- Occupy finite 3D volume
- Cohesive -- Contiguous & Continuous

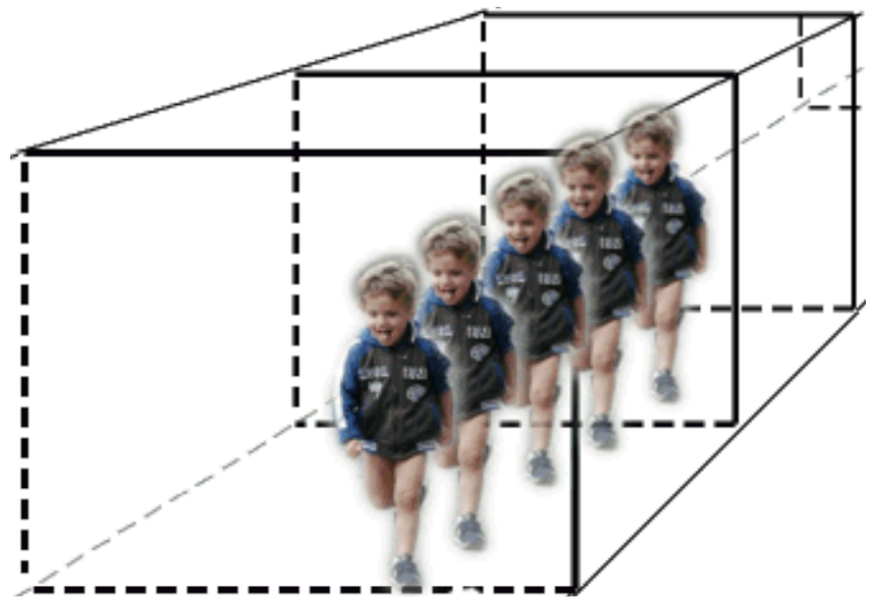

Sources: N. Ahuja, Pritch et al 07

Properties of Real World Objects

- Occupy finite 3D volume
- Cohesive -- Contiguous & Continuous
- Have locally smooth shape

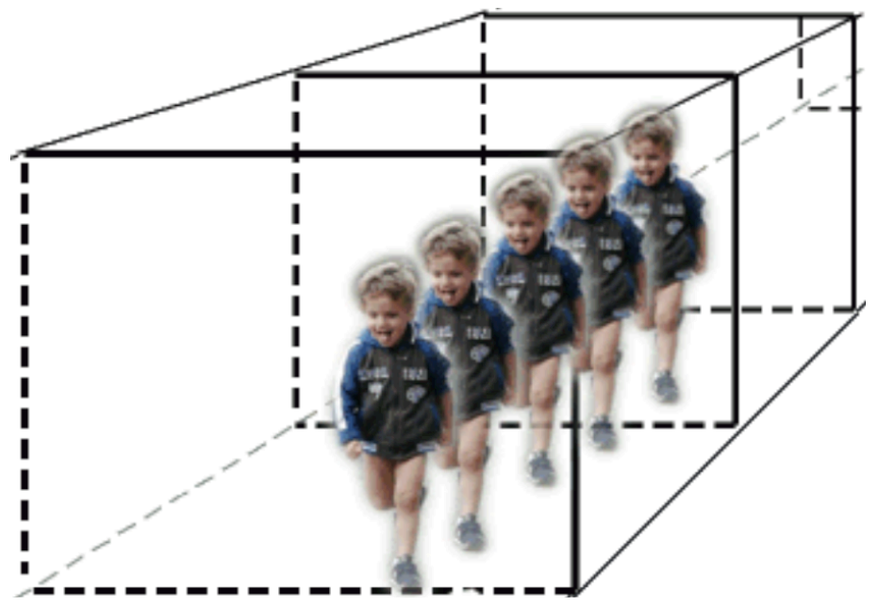

Sources: N. Ahuja, Pritch et al 07

Properties of Real World Objects

- Occupy finite 3D volume
- Cohesive -- Contiguous & Continuous
- Have locally smooth shape
- Made locally of the same material

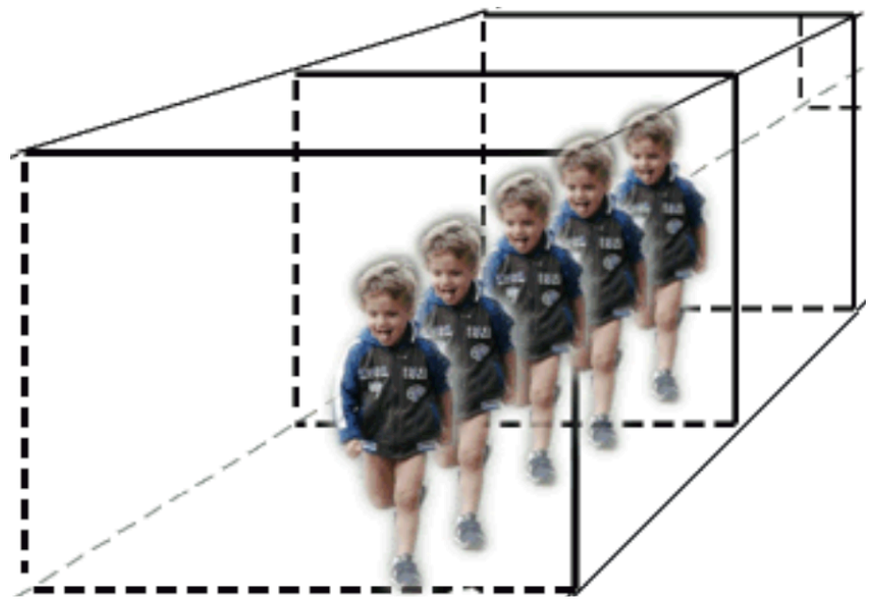

Sources: N. Ahuja, Pritch et al 07

Properties of Real World Objects

- Occupy finite 3D volume
- Cohesive -- Contiguous & Continuous
- Have locally smooth shape
- Made locally of the same material
- Opaque or partially transparent

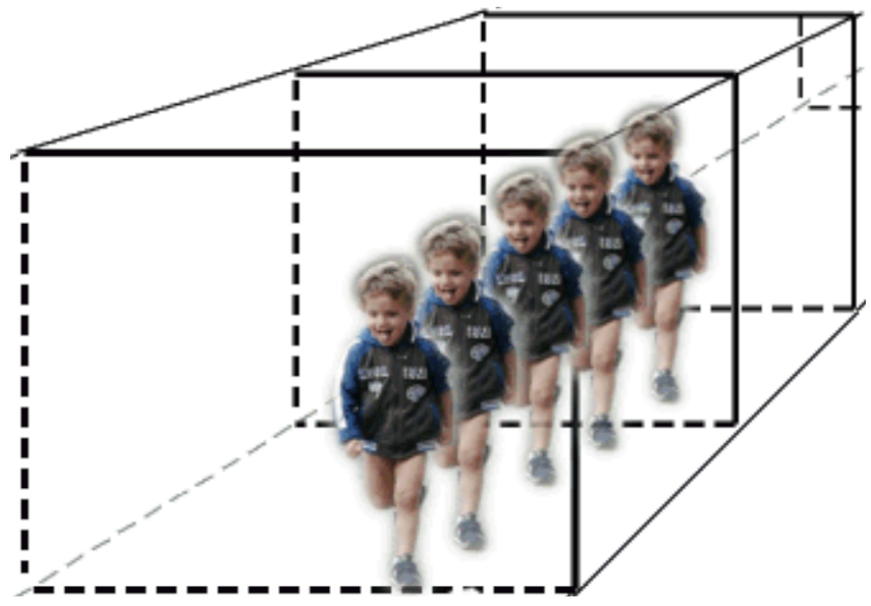

Sources: N. Ahuja, Pritch et al 07

Properties of Real World Objects

- Occupy finite 3D volume
- Cohesive -- Contiguous & Continuous
- Have locally smooth shape
- Made locally of the same material
- Opaque or partially transparent

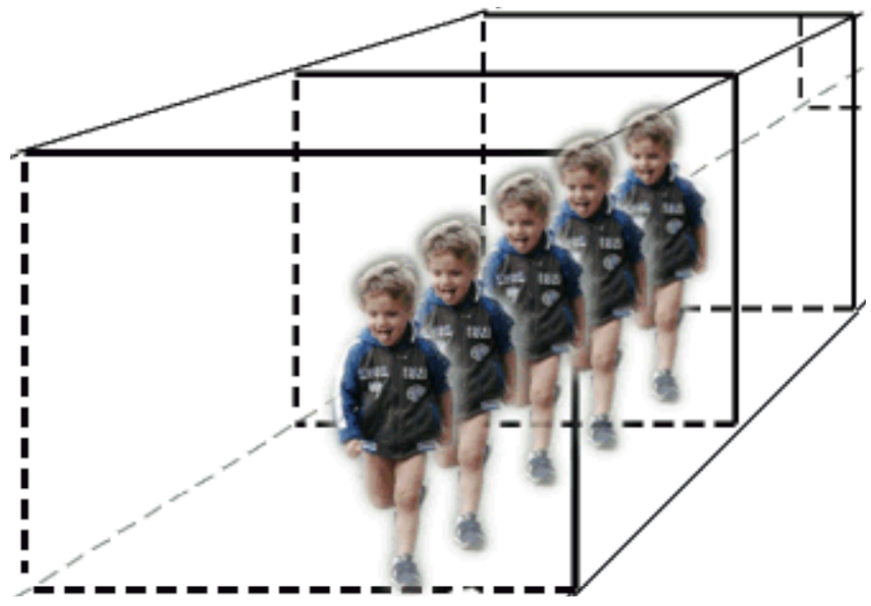

- Have certain photometric properties (e.g., reflectance, specularity)

Sources: N. Ahuja, Pritch et al 07

Properties of Real World Objects

- Occupy finite 3D volume
- Cohesive -- Contiguous & Continuous
- Have locally smooth shape
- Made locally of the same material
- Opaque or partially transparent

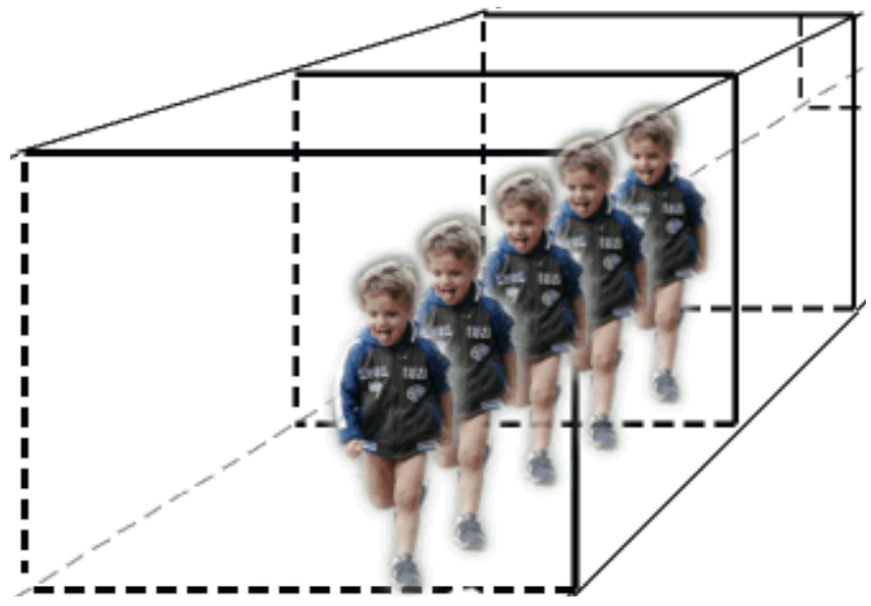

- Have certain photometric properties (e.g., reflectance, specularity)
- Occupy distinct locations

Sources: N. Ahuja, Pritch et al 07

Properties of Real World Objects

- Occupy finite 3D volume
- Cohesive -- Contiguous & Continuous
- Have locally smooth shape
- Made locally of the same material
- Opaque or partially transparent

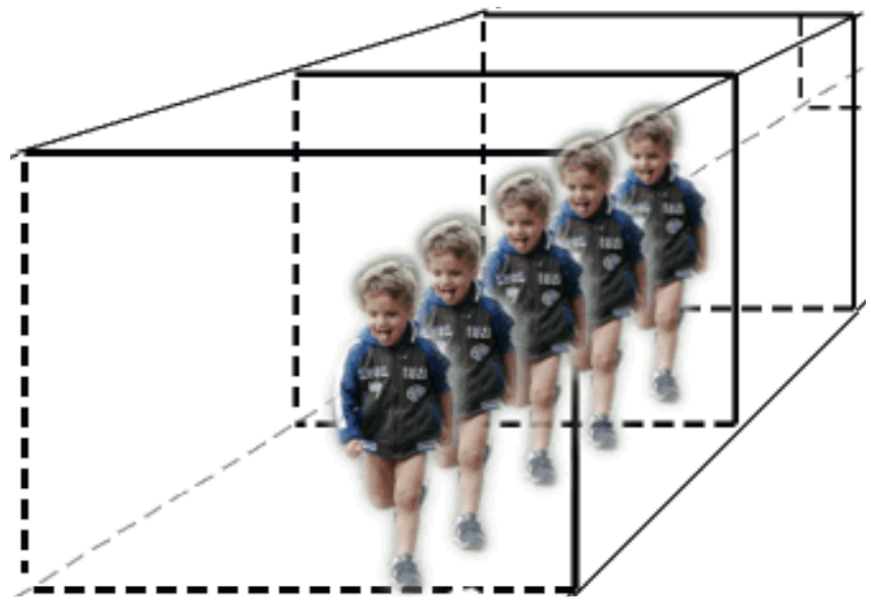

- Have certain photometric properties (e.g., reflectance, specularity)
- Occupy distinct locations
- Form characteristic spatial configurations with other 3D objects

Sources: N. Ahuja, Pritch et al 07

Properties of Real World Objects

- Occupy finite 3D volume
- Cohesive -- Contiguous & Continuous
- Have locally smooth shape
- Made locally of the same material
- Opaque or partially transparent

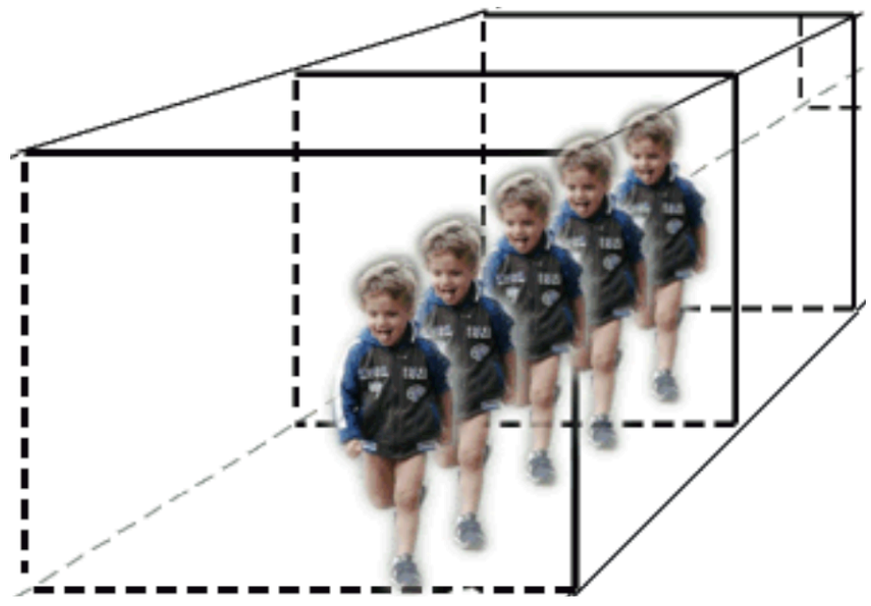

- Have certain photometric properties (e.g., reflectance, specularity)
- Occupy distinct locations
- Form characteristic spatial configurations with other 3D objects
- Form characteristic temporal configurations with other 3D objects

Sources: N. Ahuja, Pritch et al 07

Properties of Real World Objects

- Occupy finite 3D volume
- Cohesive -- Contiguous & Continuous
- Have locally smooth shape
- Made locally of the same material
- Opaque or partially transparent

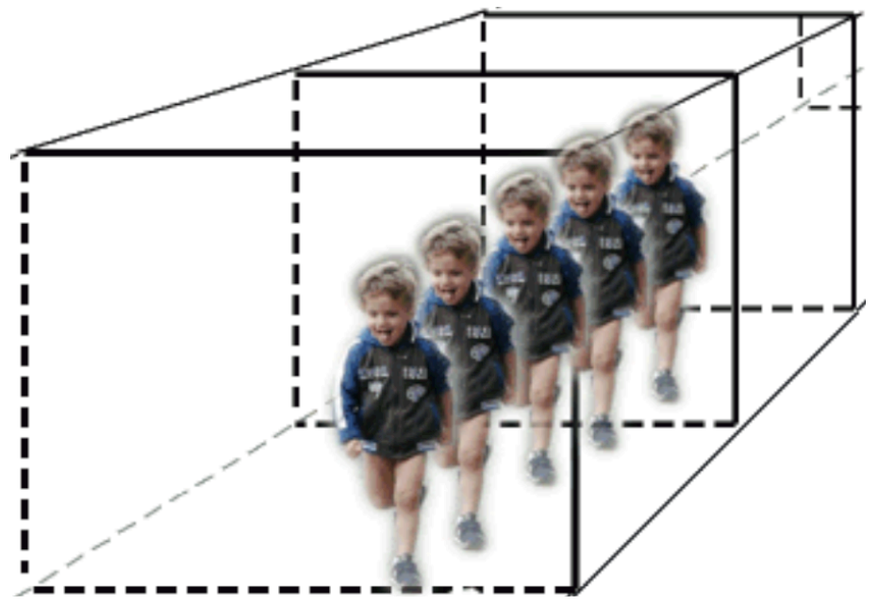

- Have certain photometric properties (e.g., reflectance, specularity)
- Occupy distinct locations
- Form characteristic spatial configurations with other 3D objects
- Form characteristic temporal configurations with other 3D objects
- Consist of parts that are valid 3D objects in their own right

Sources: N. Ahuja, Pritch et al 07

What low-level image properties capture the general constraints of real world?

Image Formation + Real World Properties → 2D Objects

Finite 3D volume + Cohesive → 2D objects = Regions

- Region boundaries coincide with 2D object boundaries
- Characteristic size and shape of 2D objects

Source: N. Ahuja

Image Formation + Real World Properties → 2D Objects

Finite 3D volume + Cohesive → 2D objects = Regions

- Region boundaries coincide with 2D object boundaries
- Characteristic size and shape of 2D objects

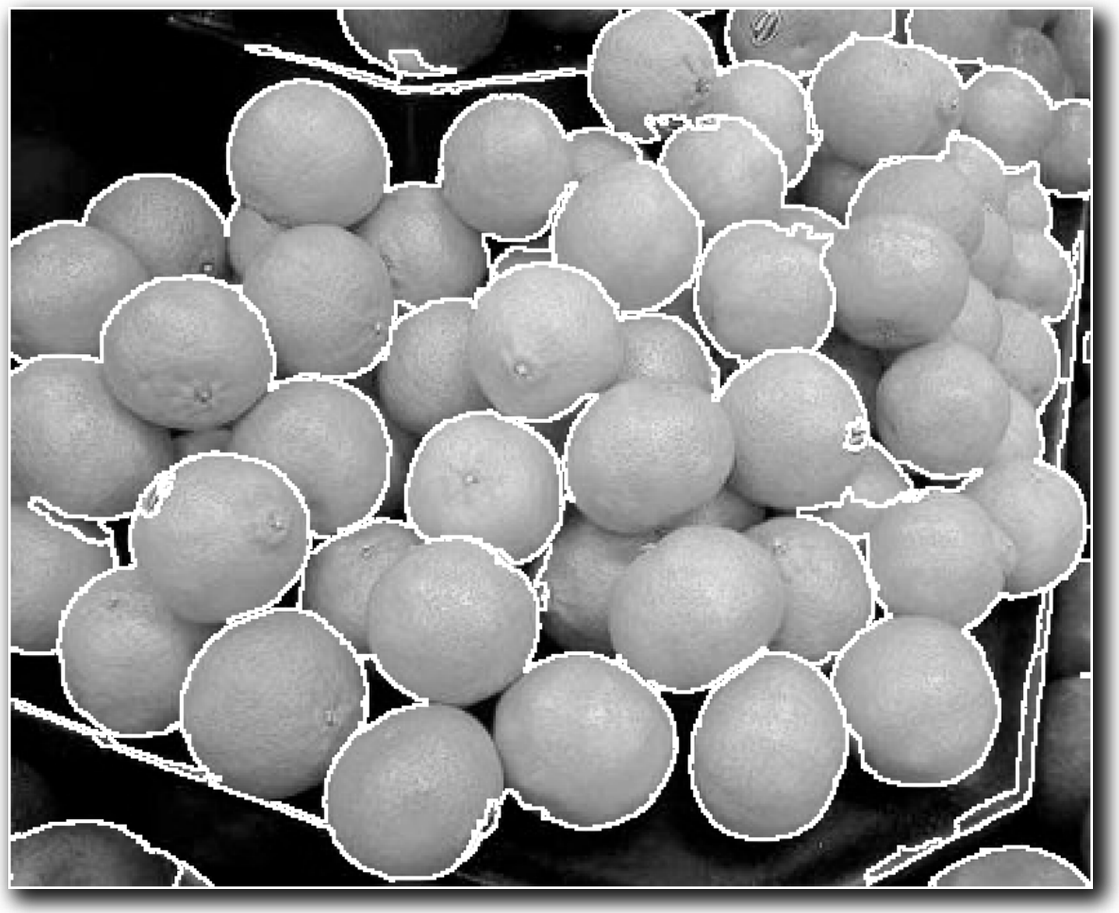

Source: N. Ahuja

Image Formation + Real World Properties → 2D Objects

- Opaqueness + Photometry → Color, brightness of 2D objects
- 3D local smoothness → Locally smooth shape, color of 2D objects
- 3D locations → 2D spatial layout, partial occlusion
- Compositionality → Smaller regions embedded within larger ones

Source: N. Ahuja

Image Formation + Real World Properties → 2D Objects

- Opaqueness + Photometry → Color, brightness of 2D objects
- 3D local smoothness → Locally smooth shape, color of 2D objects
- 3D locations → 2D spatial layout, partial occlusion
- Compositionality → Smaller regions embedded within larger ones

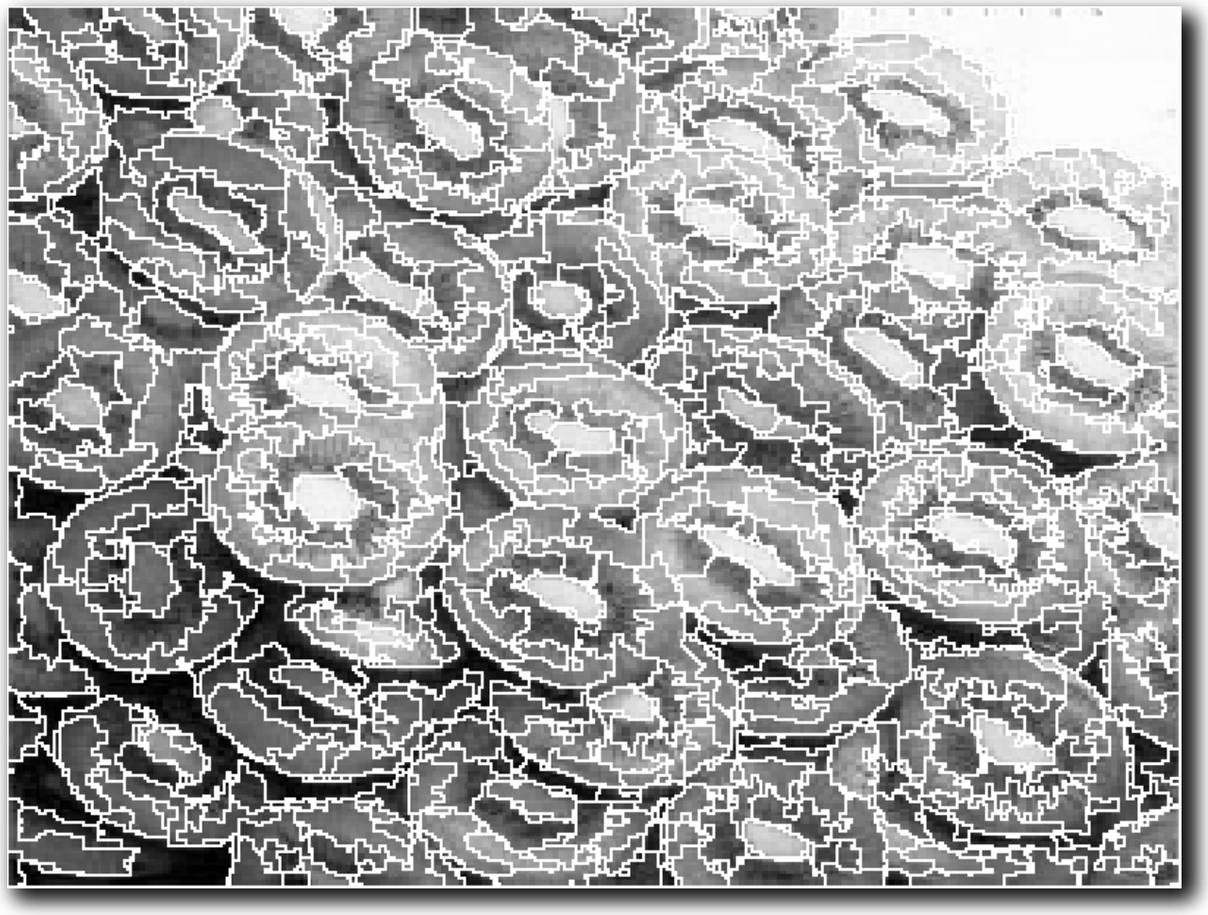

Source: N. Ahuja

From Low-Level Image Properties to Interpretation

From Low-Level Image Properties to Interpretation

Interpretation = Traversal Across Knowledge Levels

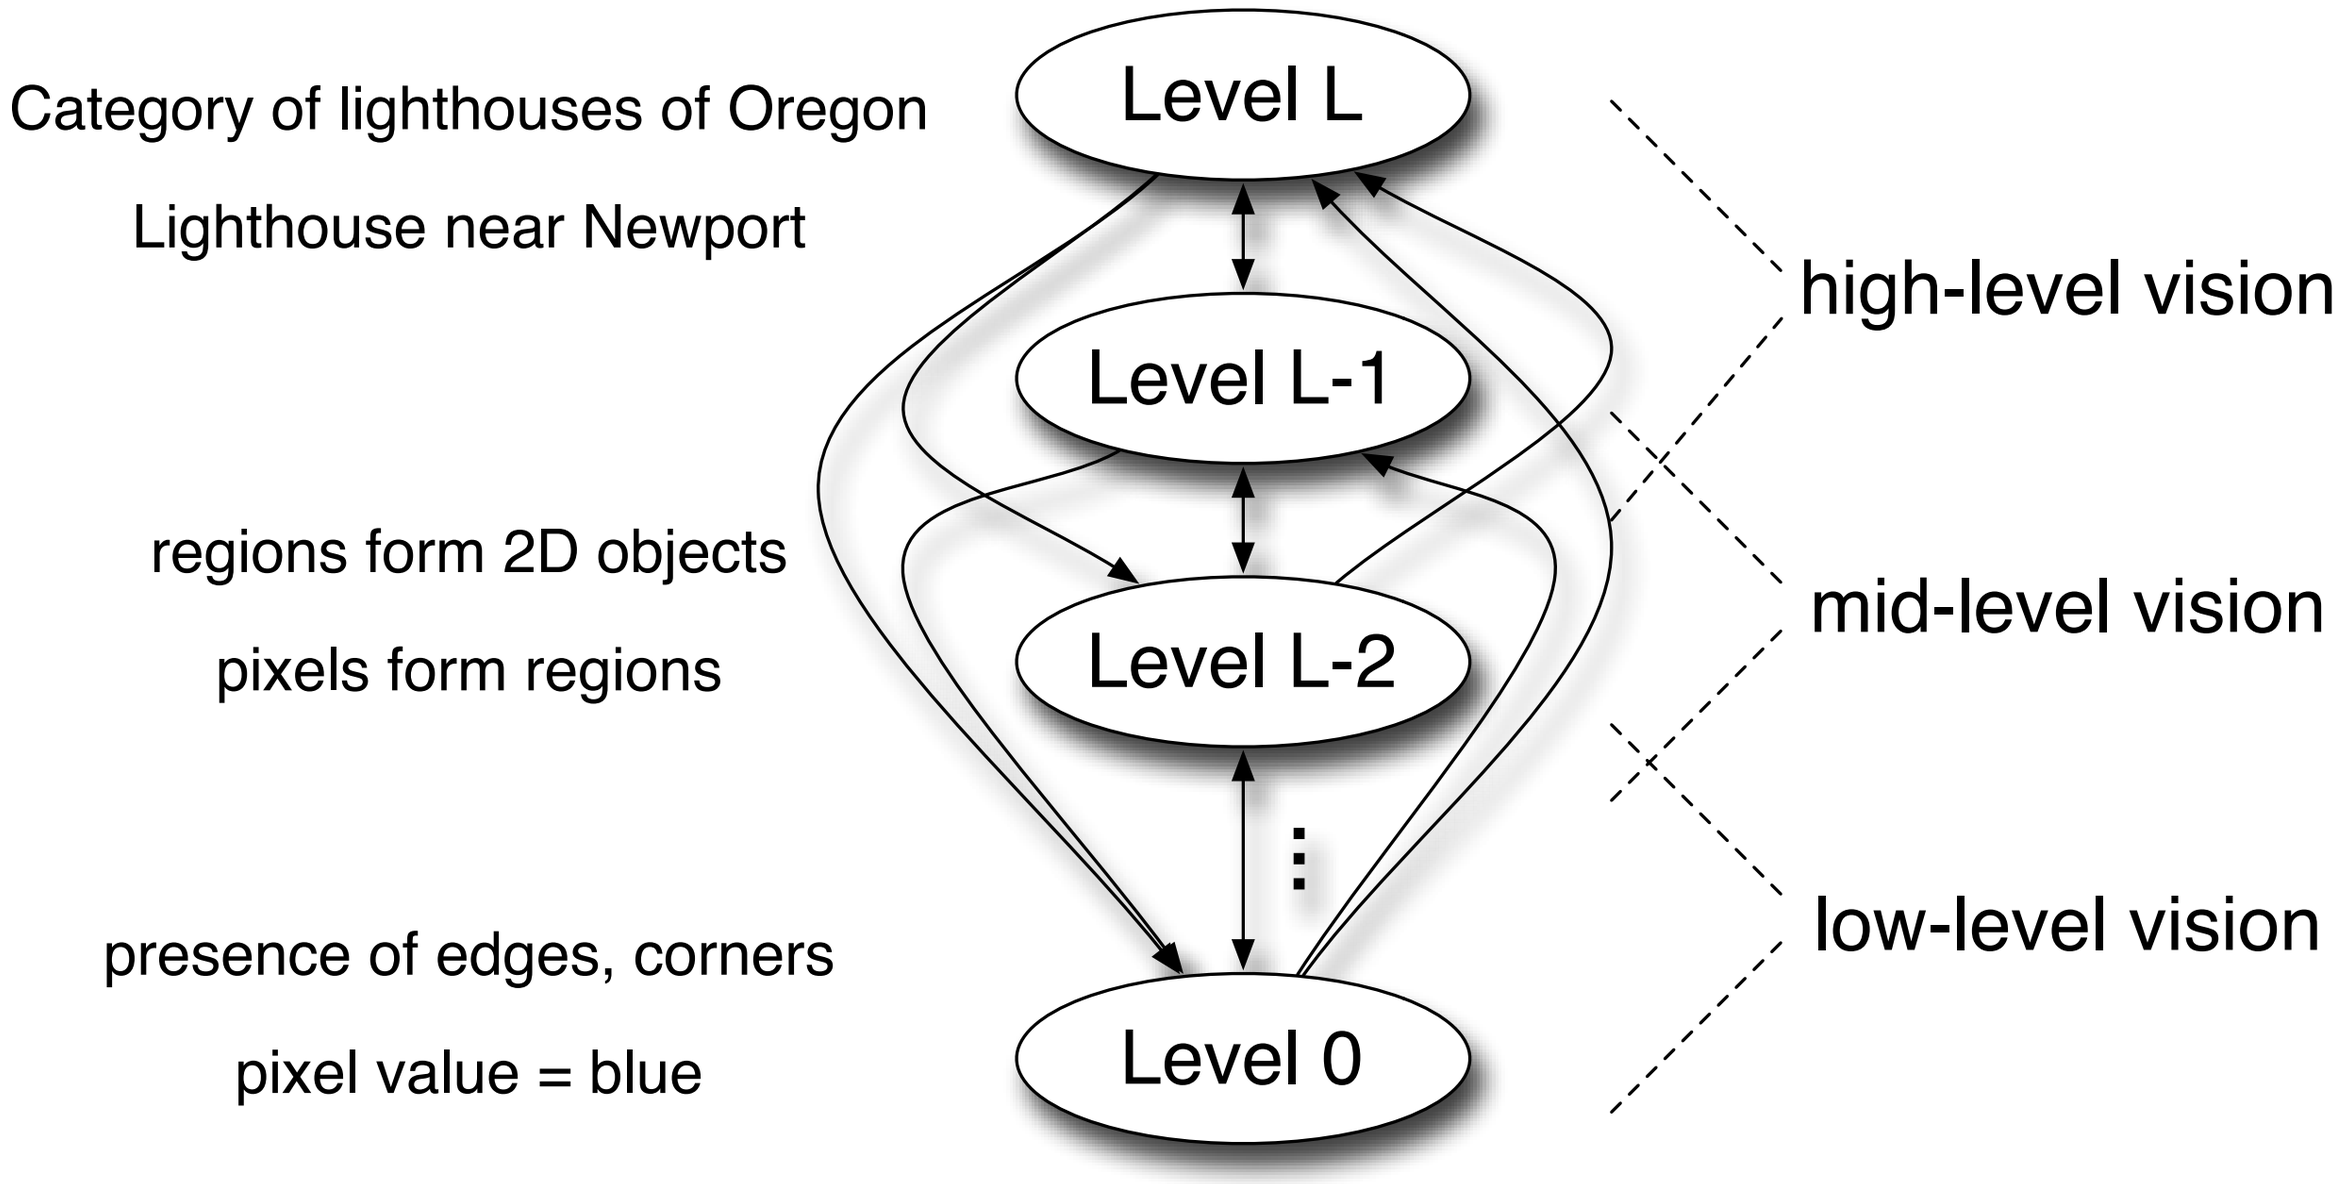

- Bottom-up traversal -- e.g., perceptual grouping
- Top-down traversal -- e.g., context

Source: N. Ahuja

Bottom-up Reasoning Important, But...

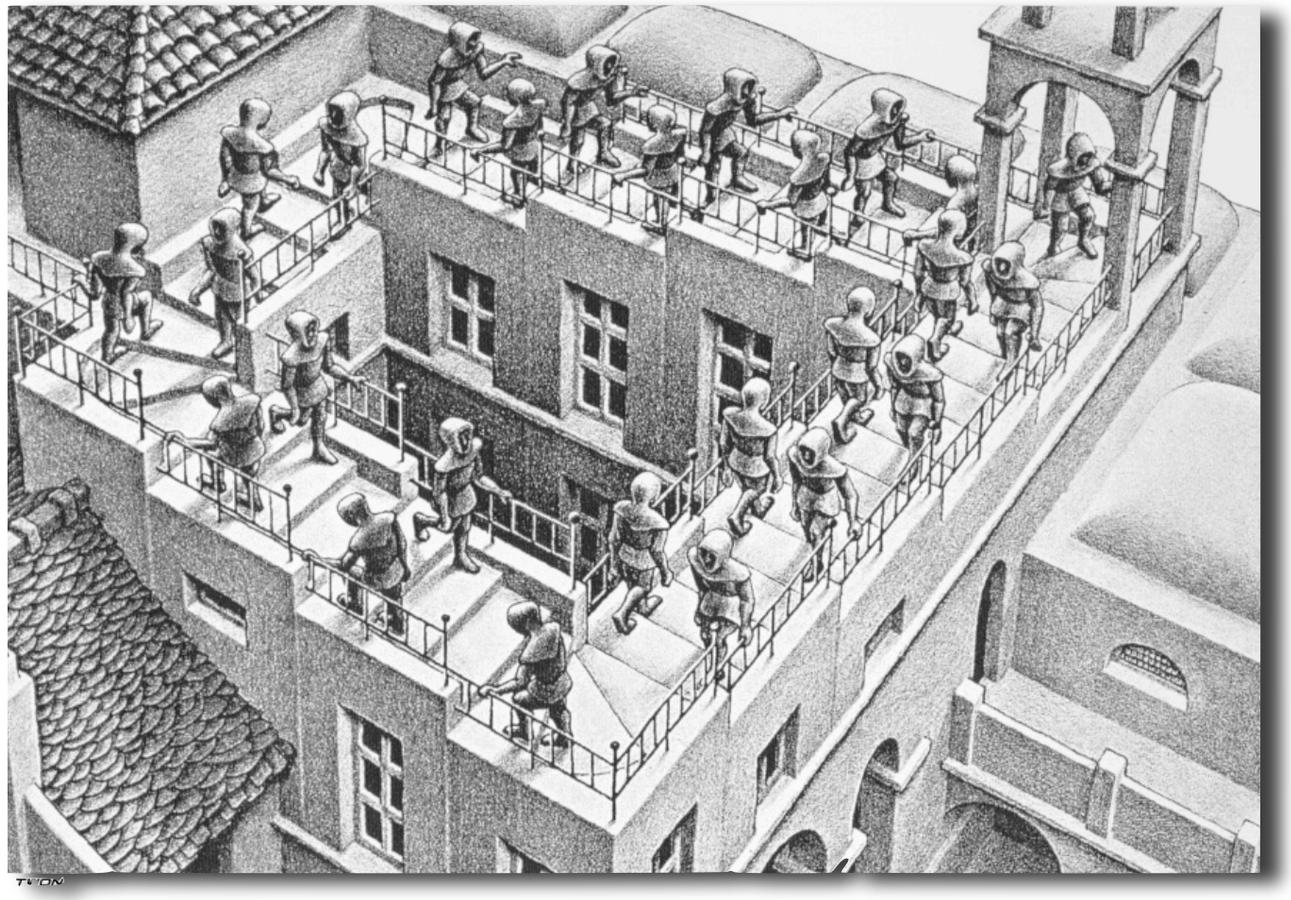

- A sense made from analyzing low-level constraints, imposed by real world, last only until cognitive scrutiny takes over
- Good syntax does not imply good semantics

Source: N. Ahuja

How to conduct image interpretation?

David Marr's Paradigm -- Feedforward Approach

original image

primal sketch

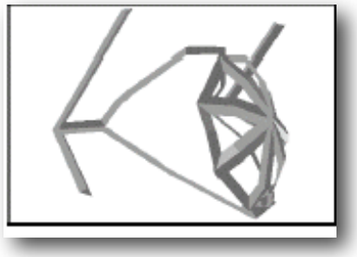

2.5D sketch

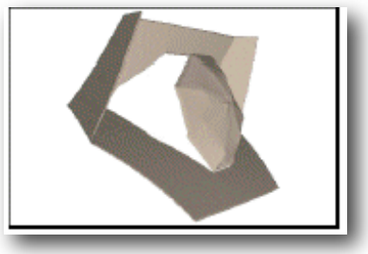

3D model

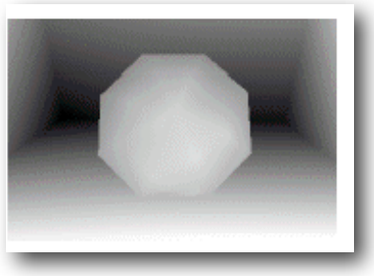

- Pros:
  - Simple
- Cons:
  - Progression often not clear
  - Early mistakes unrecoverable
  - No synergy between the stages

Source: D. Hoiem

David Marr's Paradigm

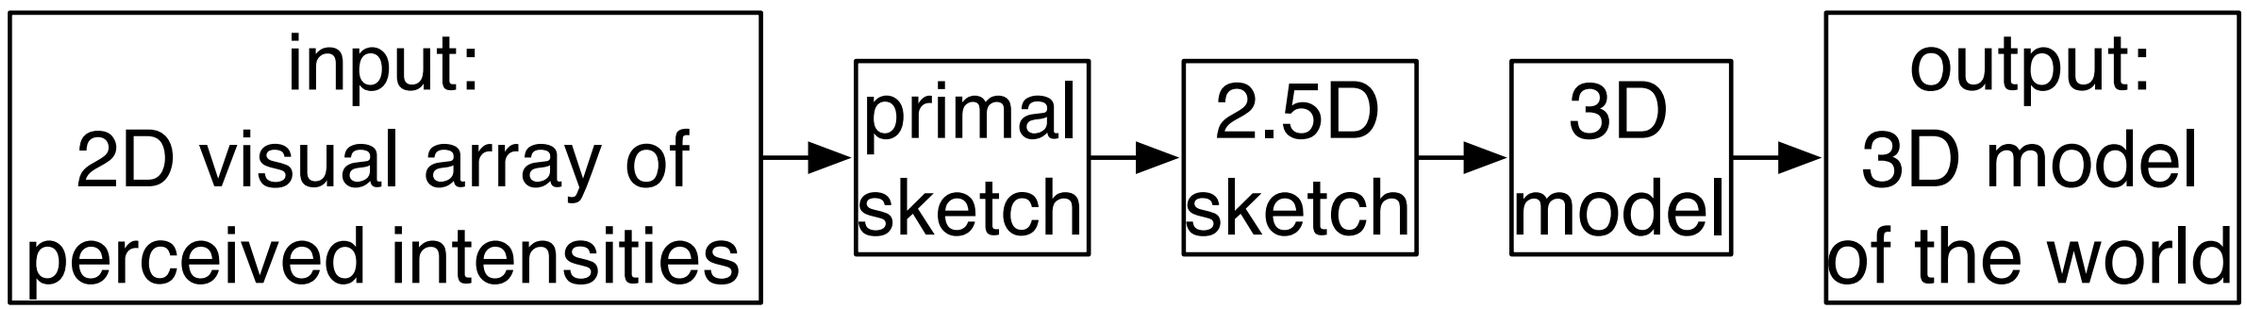

- Human vision = Stagewise processing

- From an image, we can reason what is where in the world
  ⇒ We must have representations of visual information

- Recognition matches 3D models from a catalogue to the extracted 3D representation in output

David Marr's Paradigm

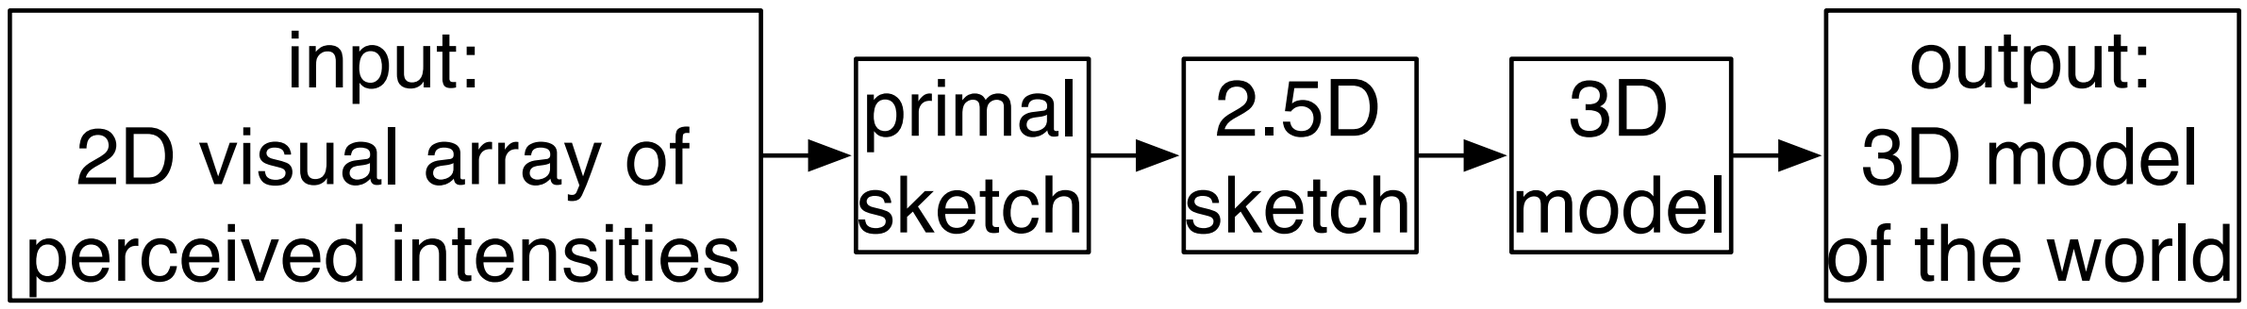

- Perceived intensities depend on:

  geometry reflectance

  illumination viewpoint

- Perceived intensities ↔ Structure of the visible surfaces

- Primal sketch = Feature extraction (edges, corners, regions)

- Each feature is a 5-tuple:

  type position contrast

  orientation scale

David Marr's Paradigm

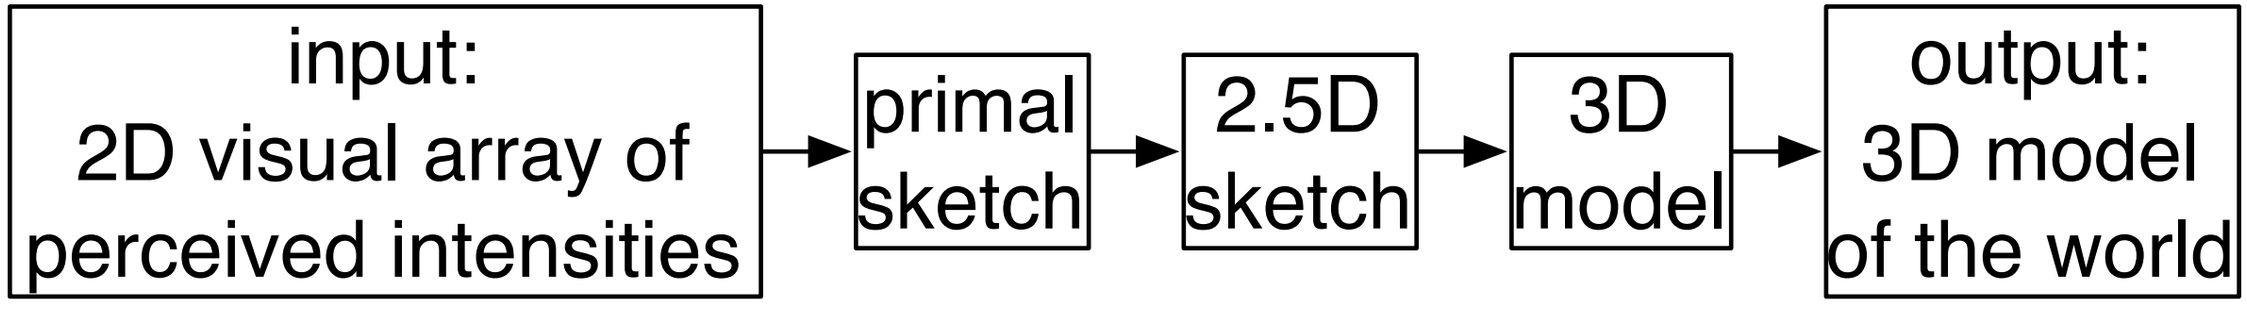

- 2.5D sketch =
  - Textures
  - Depth map
  - Surface orientation
  - Discontinuities

- Viewer centered coordinate system

David Marr's Paradigm

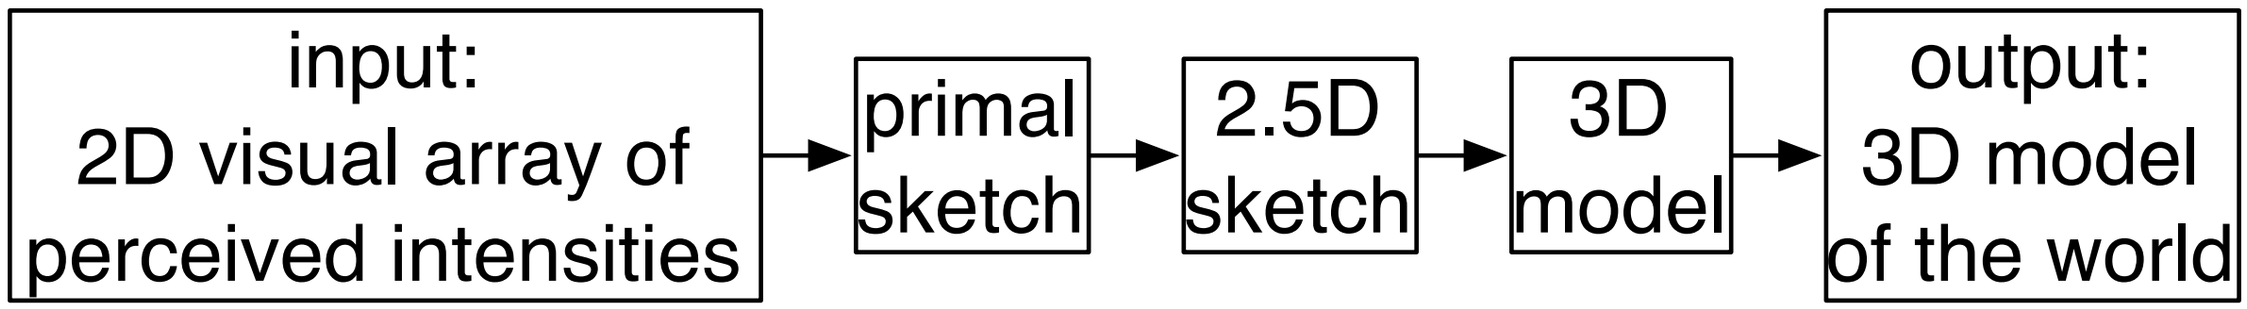

- 3D model = Modular and hierarchical representation of 3D world
- Representation in terms of volumetric and surface primitives
- Hierarchy of 3D models wrt specificity of information they carry

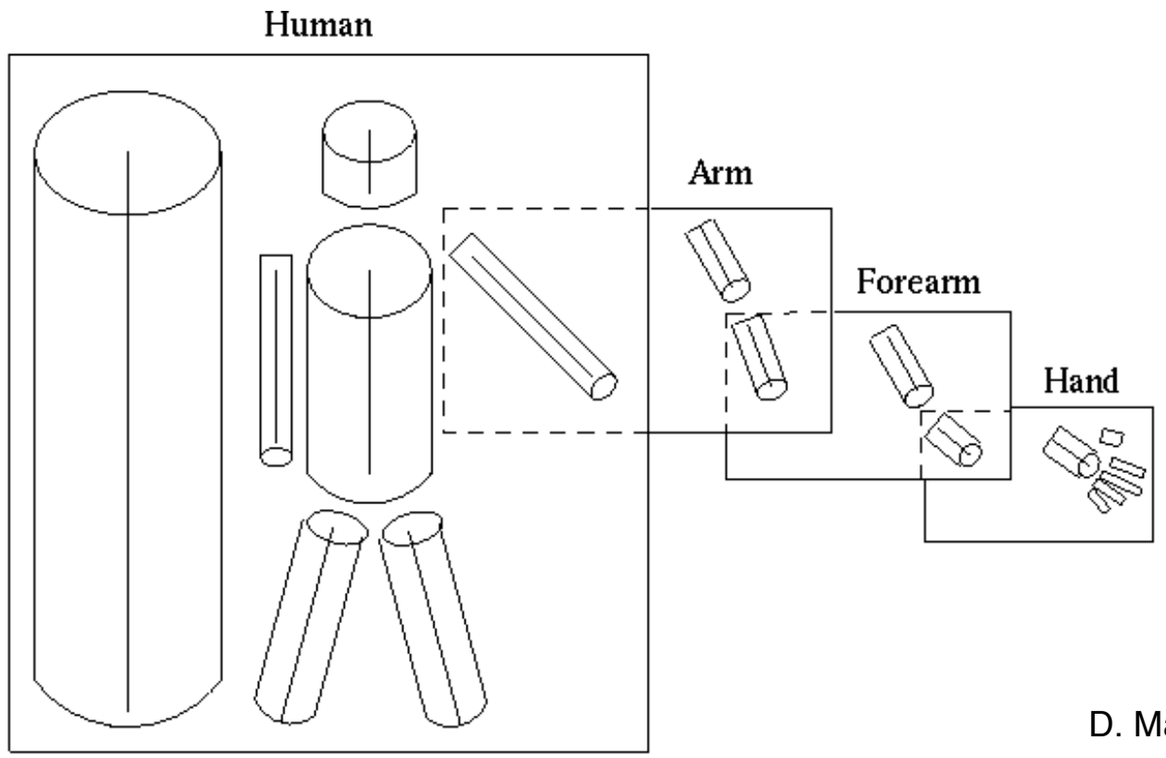

D. Marr. Vision. W. H. Freeman and Co., 1982

David Marr's Paradigm

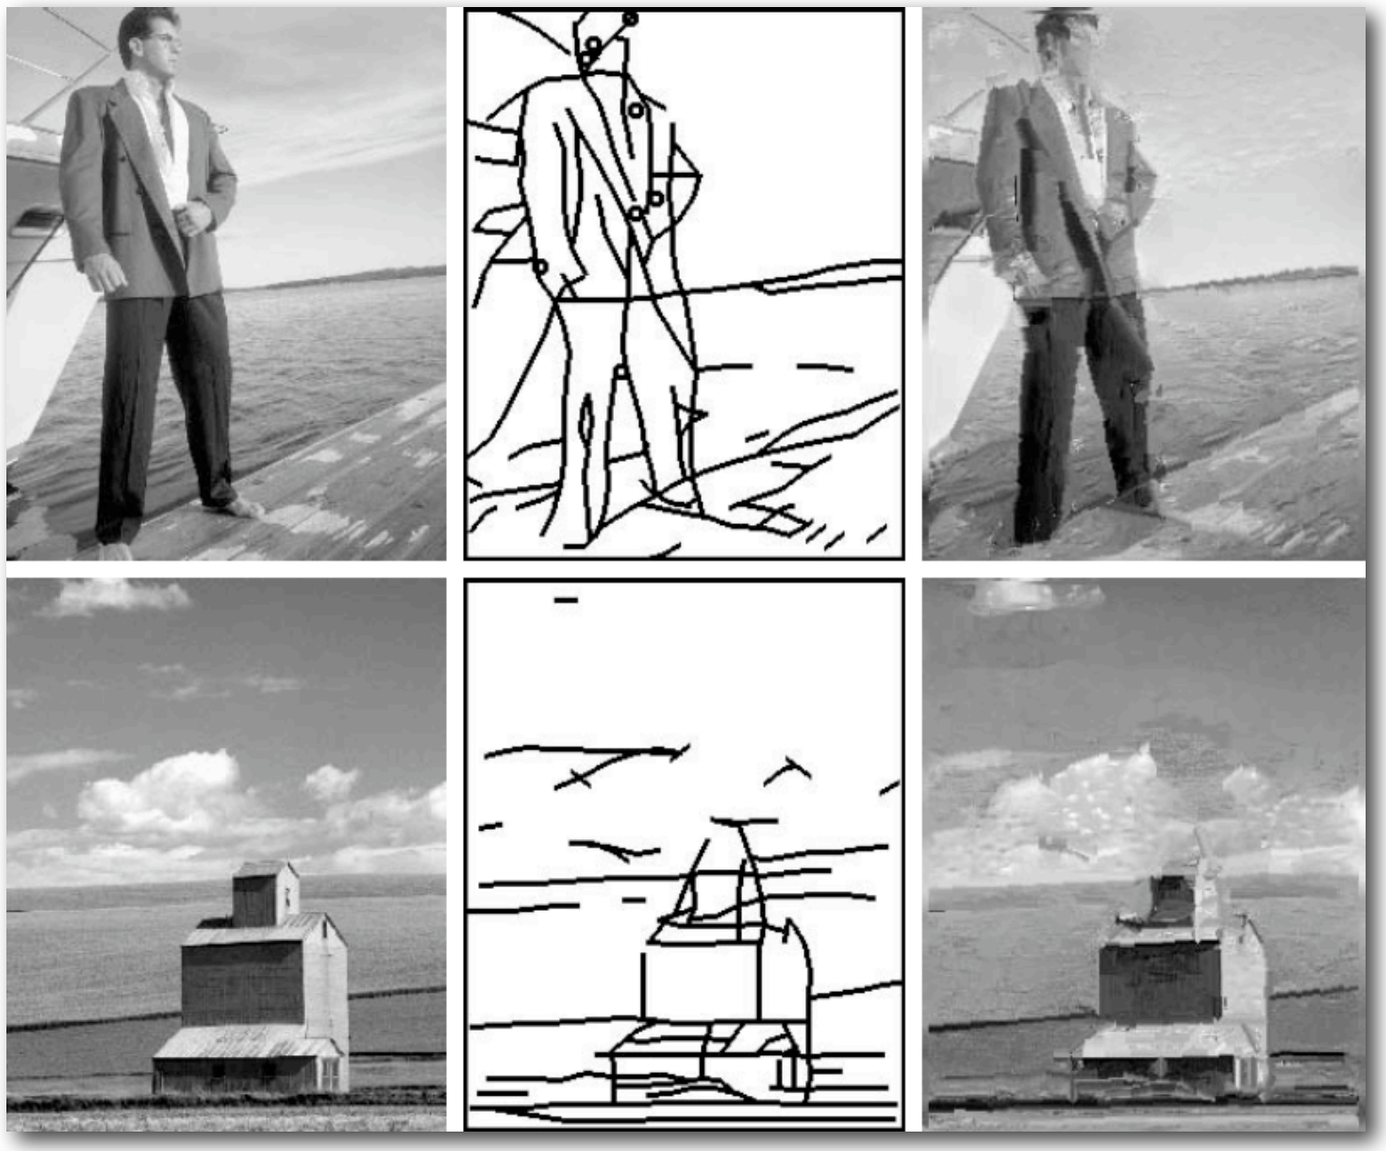

original image

Source: S.C. Zhu et al.

David Marr's Paradigm

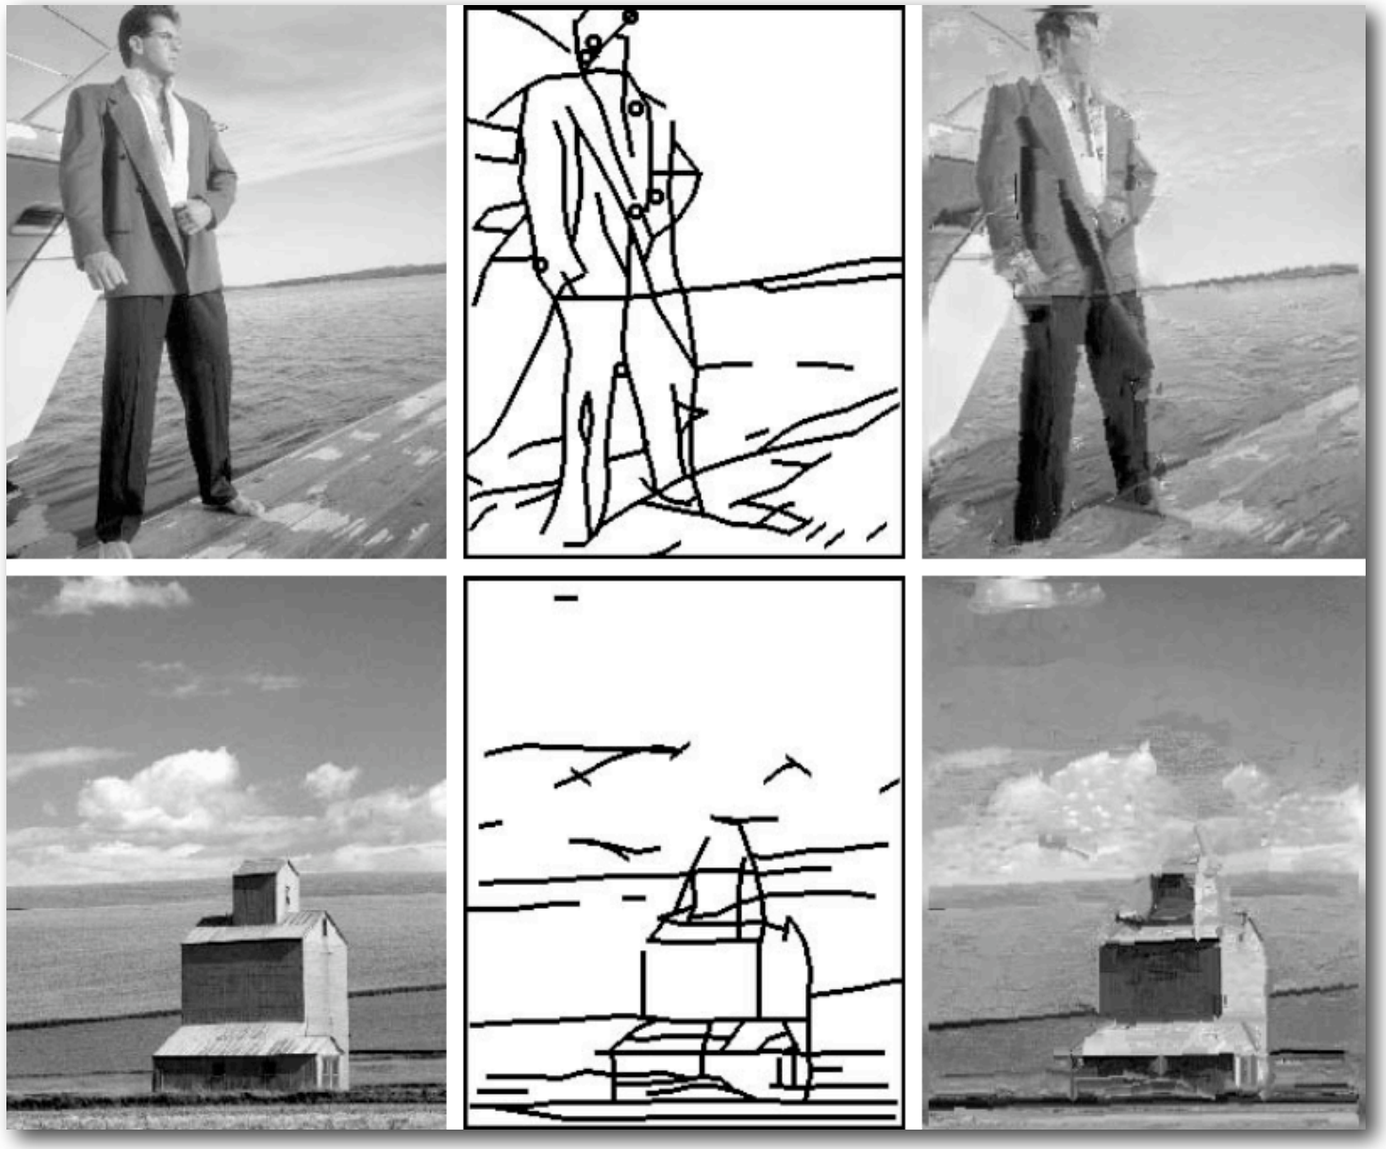

original image | primal sketch | 2.5D sketch

Source: S.C. Zhu et al.

Recognition by Components

Biederman: Recognition-by-Components

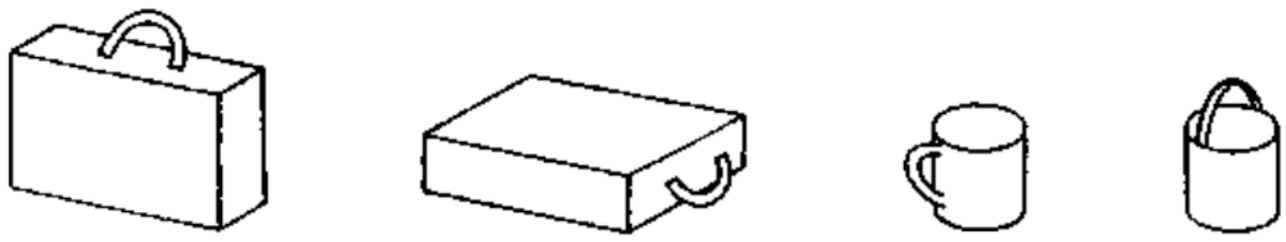

Source: Biederman 87

Biederman: Recognition-by-Components

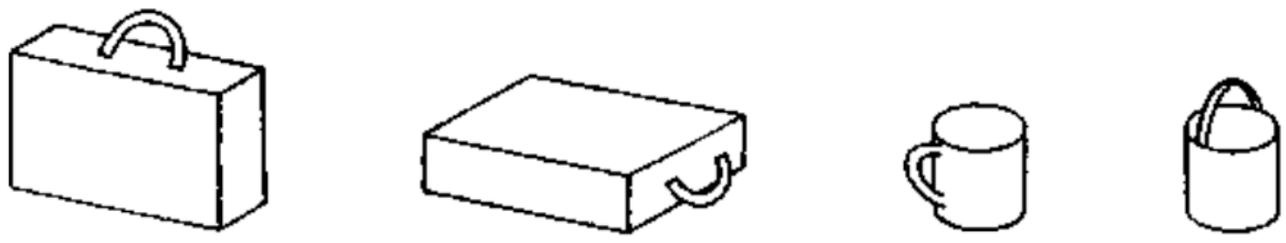

- Objects consist of generalized-cone components -- geons

Source: Biederman 87

Biederman: Recognition-by-Components

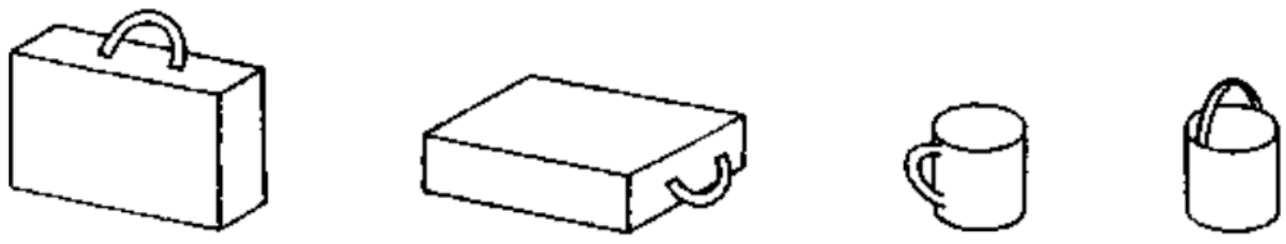

- Objects consist of generalized-cone components -- geons
- There is a finite (small) number of distinct geons, shared by objects

Source: Biederman 87

Biederman: Recognition-by-Components

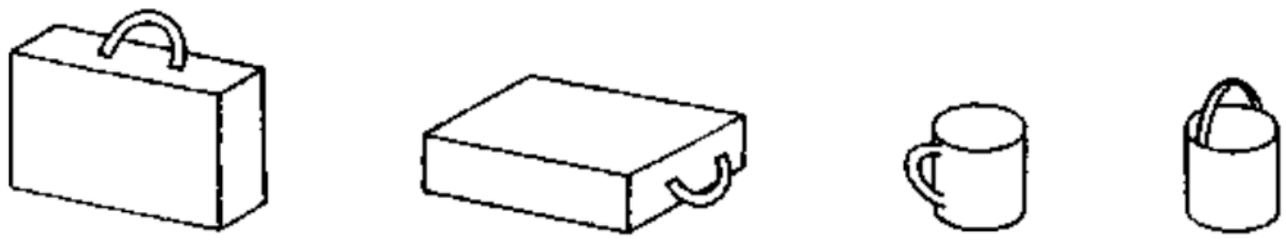

- Objects consist of generalized-cone components -- geons
- There is a finite (small) number of distinct geons, shared by objects
- Different arrangements of the same geons produce different objects

Source: Biederman 87

Biederman: Recognition-by-Components

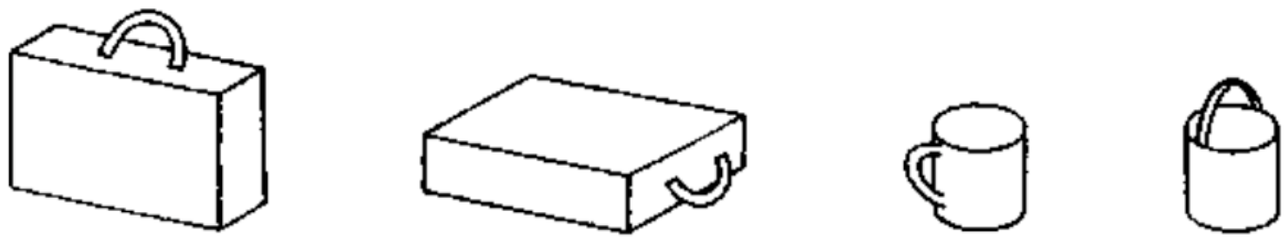

- Objects consist of generalized-cone components -- geons
- There is a finite (small) number of distinct geons, shared by objects
- Different arrangements of the same geons produce different objects
- Geons can be robustly detected from properties of edges: collinearity, curvature, symmetry, parallelism, and co-termination

Source: Biederman 87

Biederman: Recognition-by-Components

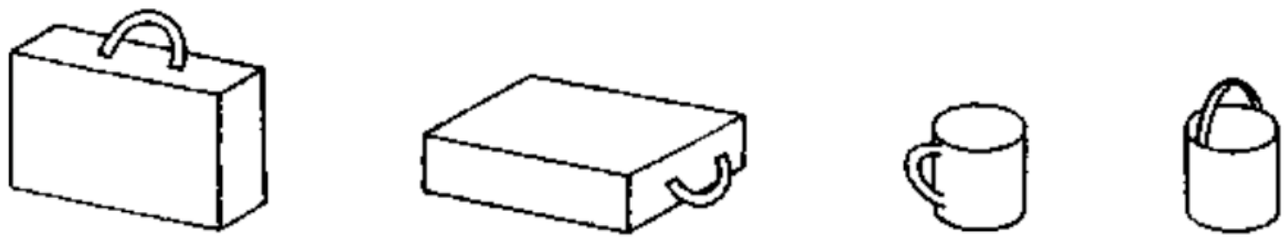

- Objects consist of generalized-cone components -- geons
- There is a finite (small) number of distinct geons, shared by objects
- Different arrangements of the same geons produce different objects
- Geons can be robustly detected from properties of edges: collinearity, curvature, symmetry, parallelism, and co-termination
- Objects can be readily recognized from detected geons even if occluded, rotated in depth, deformed, etc.

Source: Biederman 87

Biederman: Recognition-by-Components

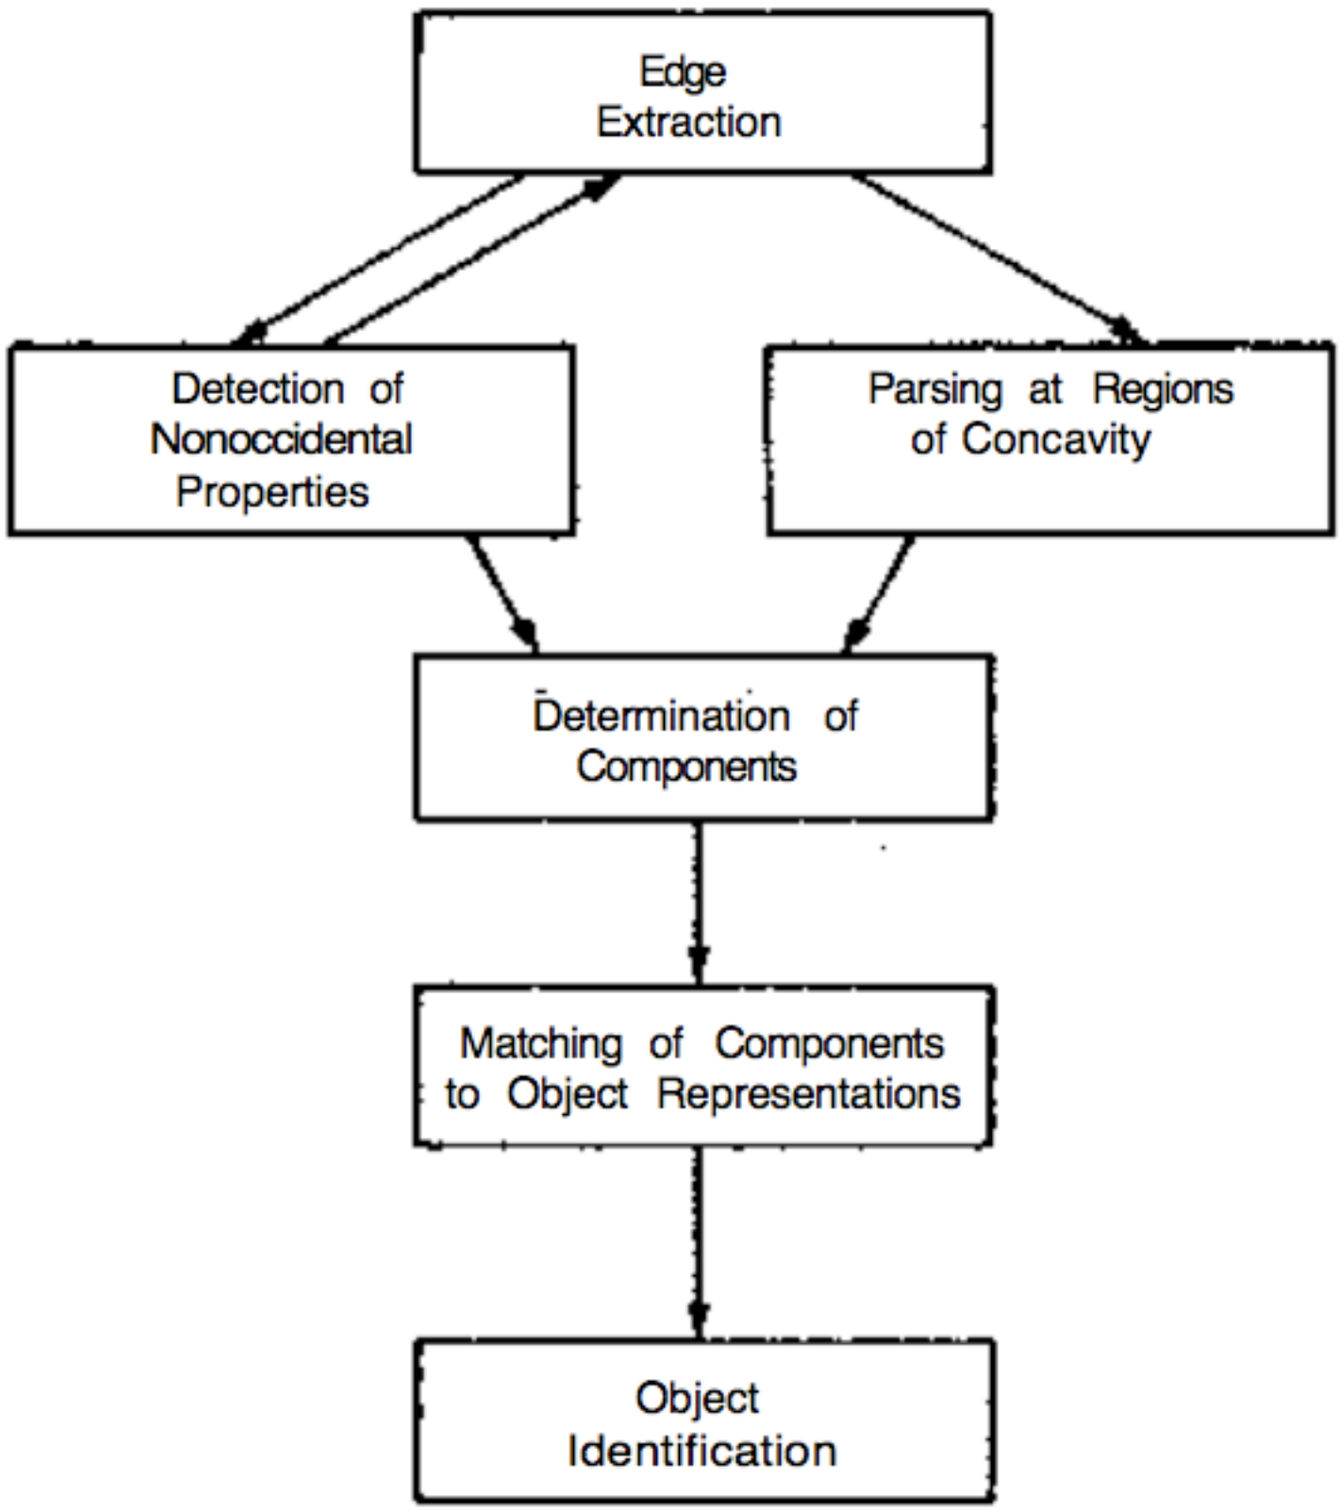

Source: Biederman 87

Intrinsic Images

Barrow & Tenenbaum -- Intrinsic Images

Input Image

Scene Analysis Processes

Surface Orientation
Object/Viewpoint
Occlusion/Depth

Intrinsic Images

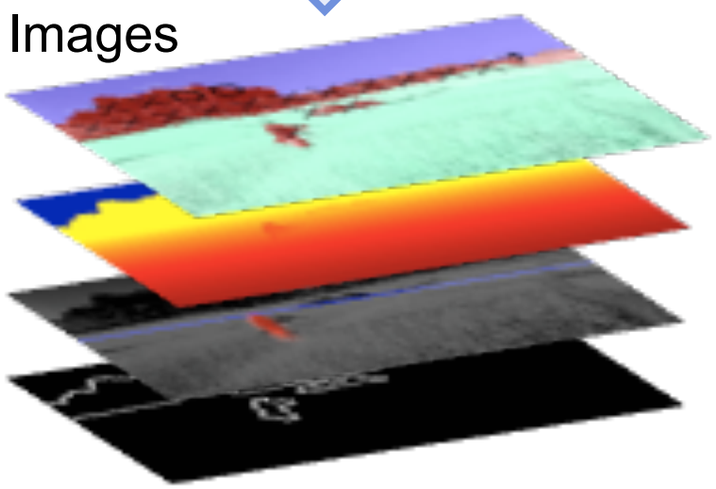

- Intrinsic images = Confidence maps
- Pros:
  - Synergistic contextual interactions
  - Easily extended
  - Robust
  - Consistency enforced via soft constraints
- Cons:
  - Complex
  - How to fuse flawed and disparate intrinsic images into coherent interpretation?

Source: D. Hoiem

Barrow & Tenenbaum -- Intrinsic Images

Surfaces

Surface Maps

Depth, Boundaries

Occlusions

Support

Horizon, Object Maps

Objects and Viewpoint

Boundaries

Horizon, Object Maps

Viewpoint/Size
Reasoning

Source: D. Hoiem

Output of Individual Intrinsic Images

input

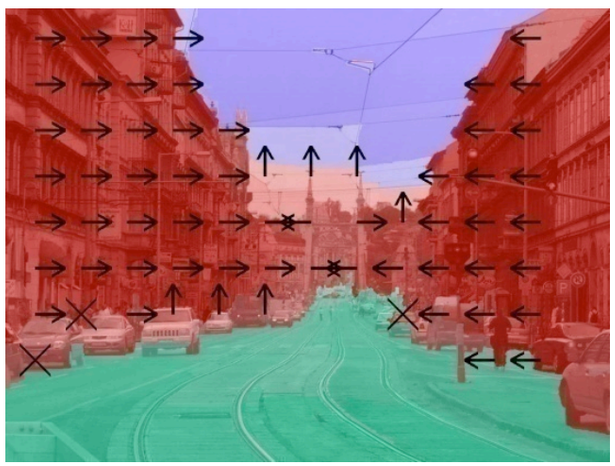

surface map

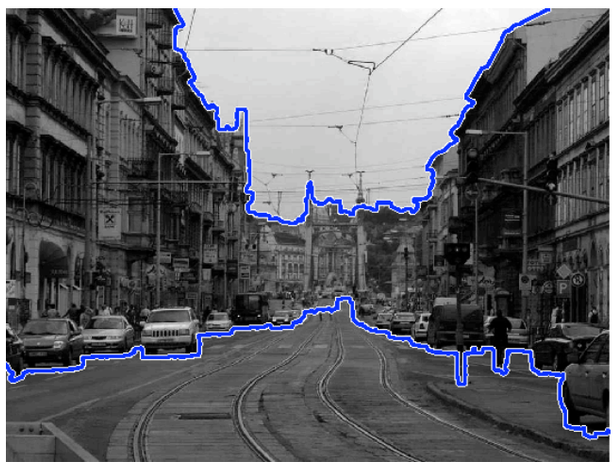

occlusion boundaries

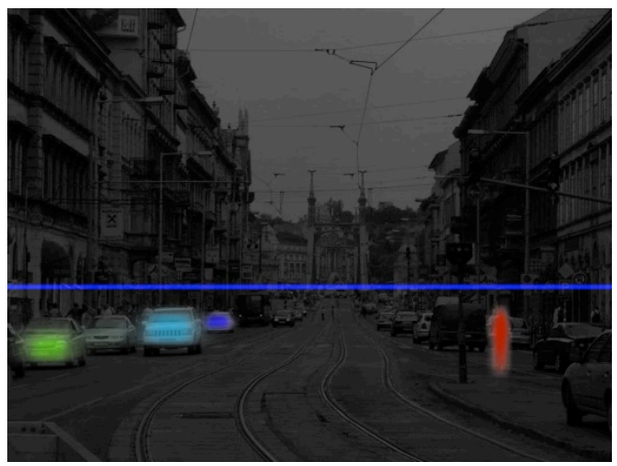

object/horizon detection

Source: D. Hoiem

Results by Fusing all Intrinsic Images

input

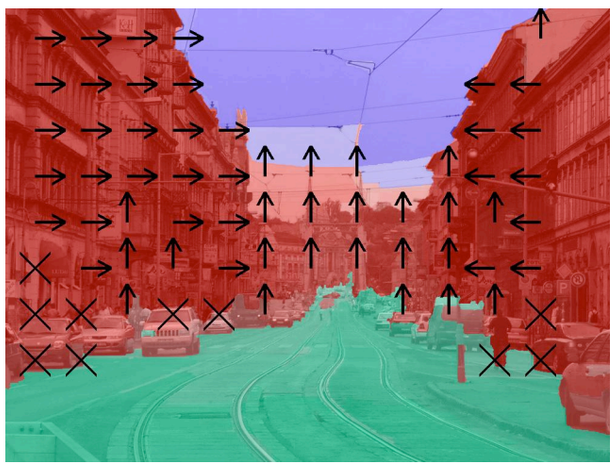

surface map

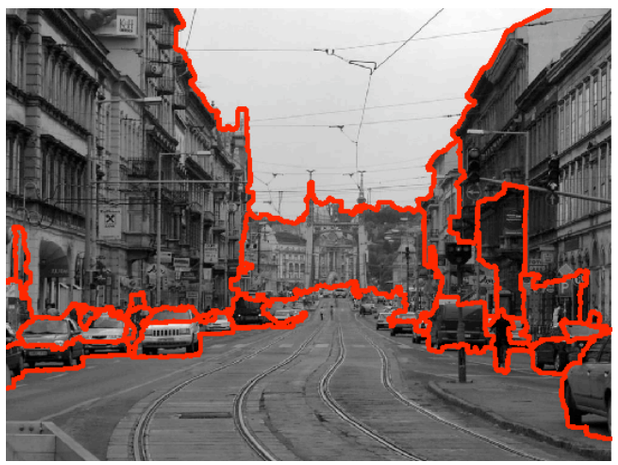

occlusion boundaries

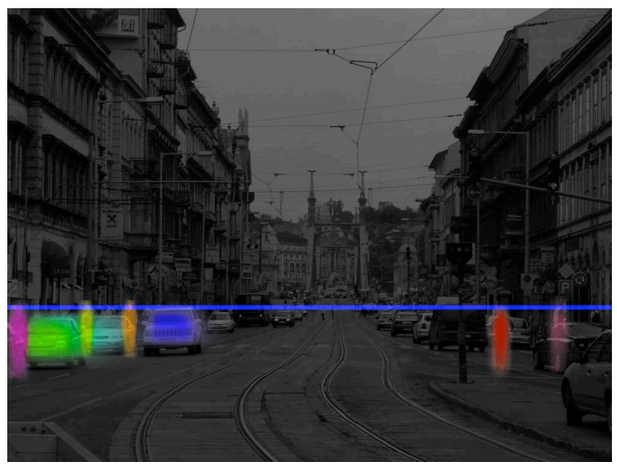

object/horizon detection

Source: D. Hoiem

Next Class

- Image features
  - Color
  - Edges
  - Interest points
- Homework 1