# Shape Constancy and Shape Recovery: wherein Human and Computer Vision Meet

*Zygmunt Pizlo*

Purdue University, U.S.A.

Acknowledgments: Y. Li, T. Sawada, R.M. Steinman

*Support: NSF, US Department of Energy, AFOSR*

# 3D Shape Perception

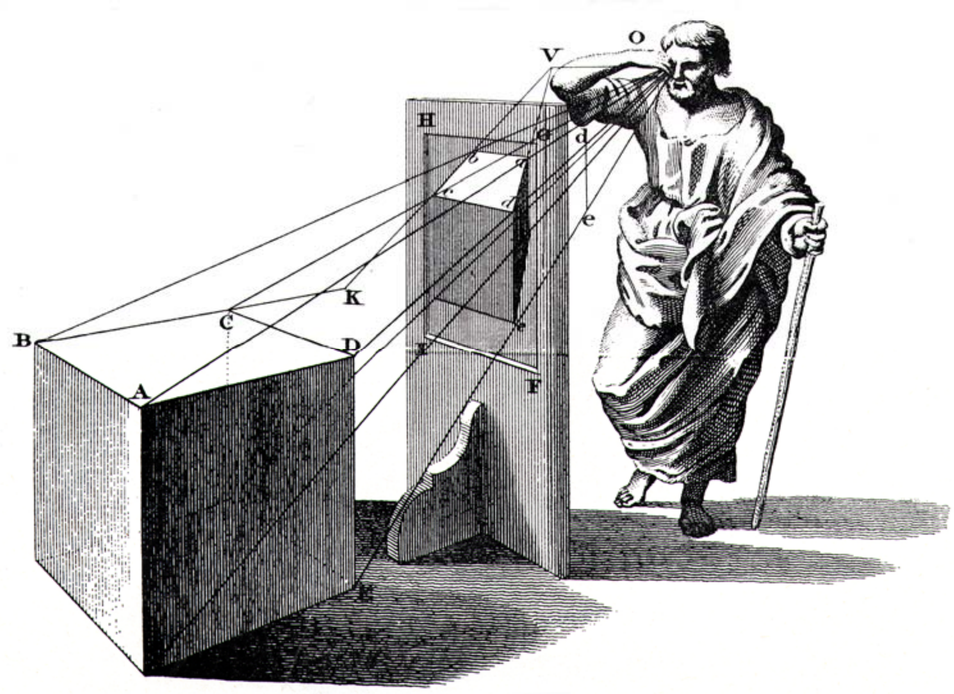


Objects out there are 3D and the percepts are 3D.
But the retinal image is 2D.

How is the 3rd dimension recovered in the percept?

# 3D shape recovery

- Perceptual mechanisms required for 3D shape *recovery* are different from the mechanisms required for *reconstructing* color, size, or surfaces.
- This will not come as a surprise once one realizes that 3D shapes are *special*. They are special because they are *complex* in the sense that the number of parameters needed to describe the shape of a 3D object is (theoretically) infinite (∞).
- This contrasts markedly with the small number of parameters needed to describe color (3), size (1), or surface orientation (2).

- The fact that 3D shape is complex allows one to make use of very *effective* constraints for recovering 3D shape.
- Note that these *shape constraints* must be defined as *shape invariants* because if they are not, they will confound 3D shape with irrelevant aspects of the viewing conditions, e.g., viewing direction or illumination.
- Note that these shape constraints must also be *effective* in the sense that they allow one to tell different 3D shapes apart. They are not *effective* with ellipses or ellipsoids because these shapes are not complex. There are no effective shape invariants for these shapes.

- Now that you know the requirements for effective shape constraints, I will show you how our computational model, making use of such effective constraints, recovers the 3D shape of a sufficiently complex object from one of its 2D images.
- The main shape constraint used by our model is 3D mirror symmetry. Symmetry is combined with 3 additional constraints, representing the likelihood function.

# 3D mirror symmetry

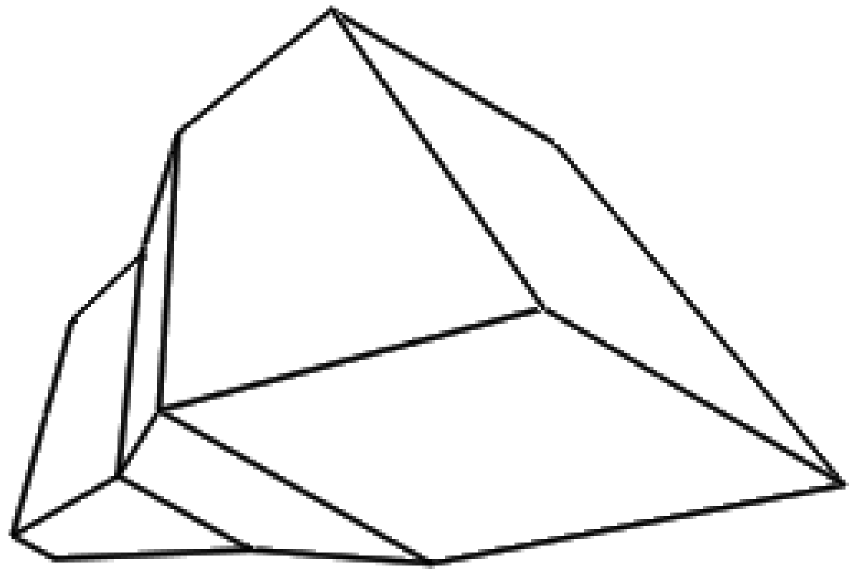

3D mirror symmetry is an *effective* constraint because there are features of the 3D symmetric shape that are preserved (invariant) in any of its 2D images. The invariant feature shown in this figure is the parallelism of the line segments connecting pairs of points symmetric in 3D.

# How 3D mirror symmetry helps 3D shape recovery

3D mirror symmetry greatly reduces the family of possible 3D interpretations. It does this because a single 2D image of a mirror symmetric shape is equivalent to two 2D images of an arbitrary shape. Vetter & Poggio (1994) showed how this works.

# Vetter & Poggio's (1994) Method

- The recovery of a 3D mirror-symmetric shape from a *single* 2D orthographic image is equivalent to the recovery of an arbitrary 3D shape from *two* 2D images.
- The second image (called the "virtual image") is produced from the given 2D image, simply by 2D mirror reflection.
- Recall that *two* 2D images determine a 3D shape "out there" up to only one free parameter (Huang & Lee, 1989, Koenderink & van Doorn, 1991).

# Vetter & Poggio's (1994) Method

***Real image***

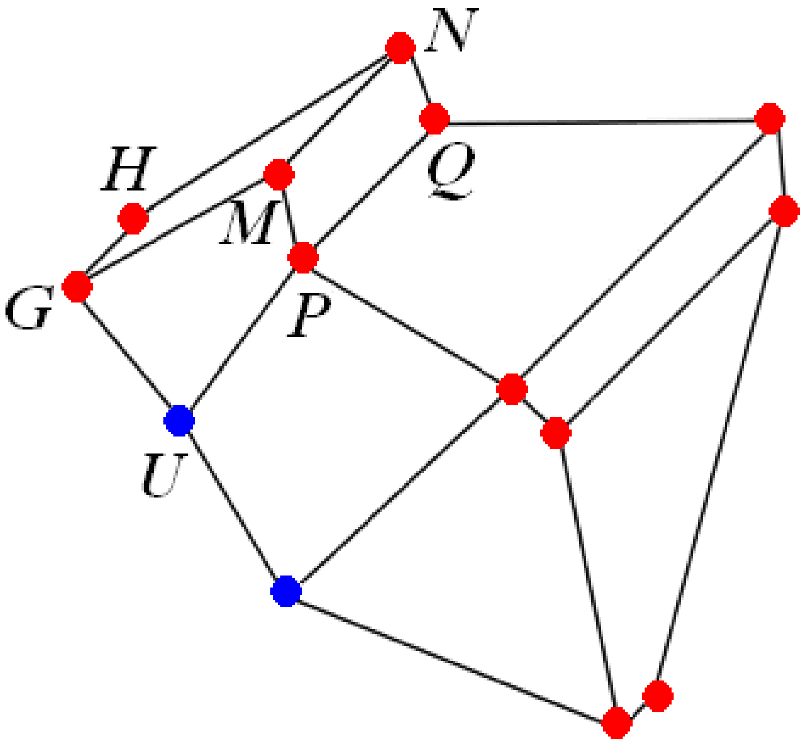

# Vetter & Poggio's (1994) Method

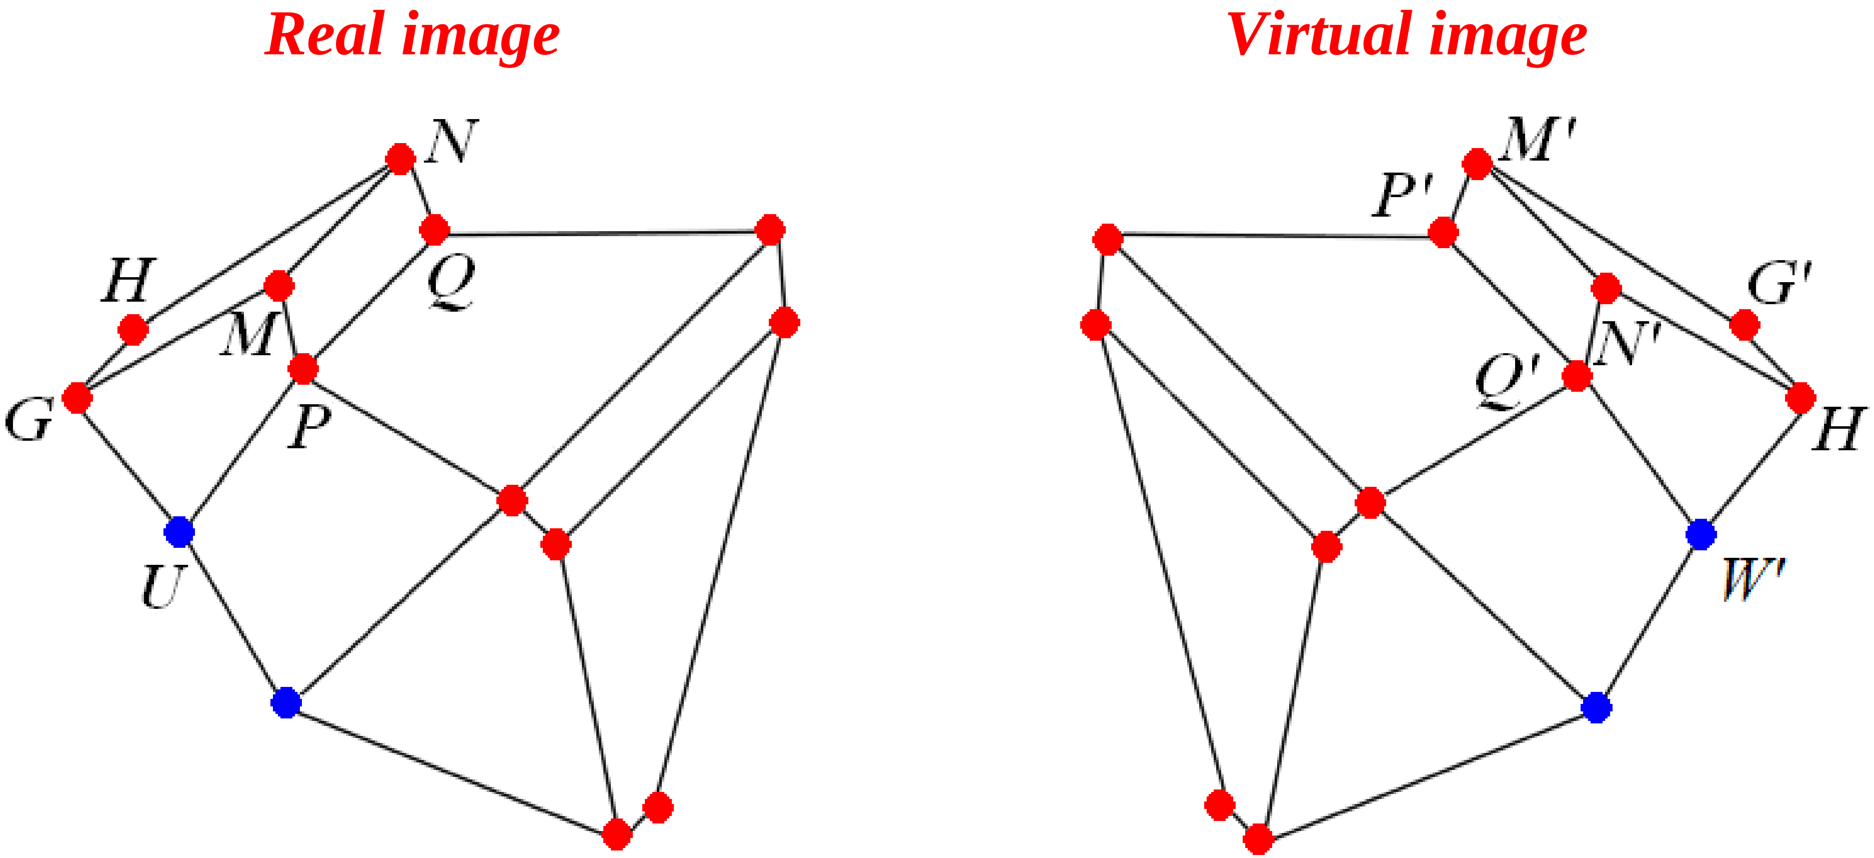


3D reflection of a 3D symmetric shape can be undone by 3D rigid rotation. Two 2D images determine the unknown 3D symmetric shape with only one degree of freedom.

# Vetter & Poggio's (1994) Method

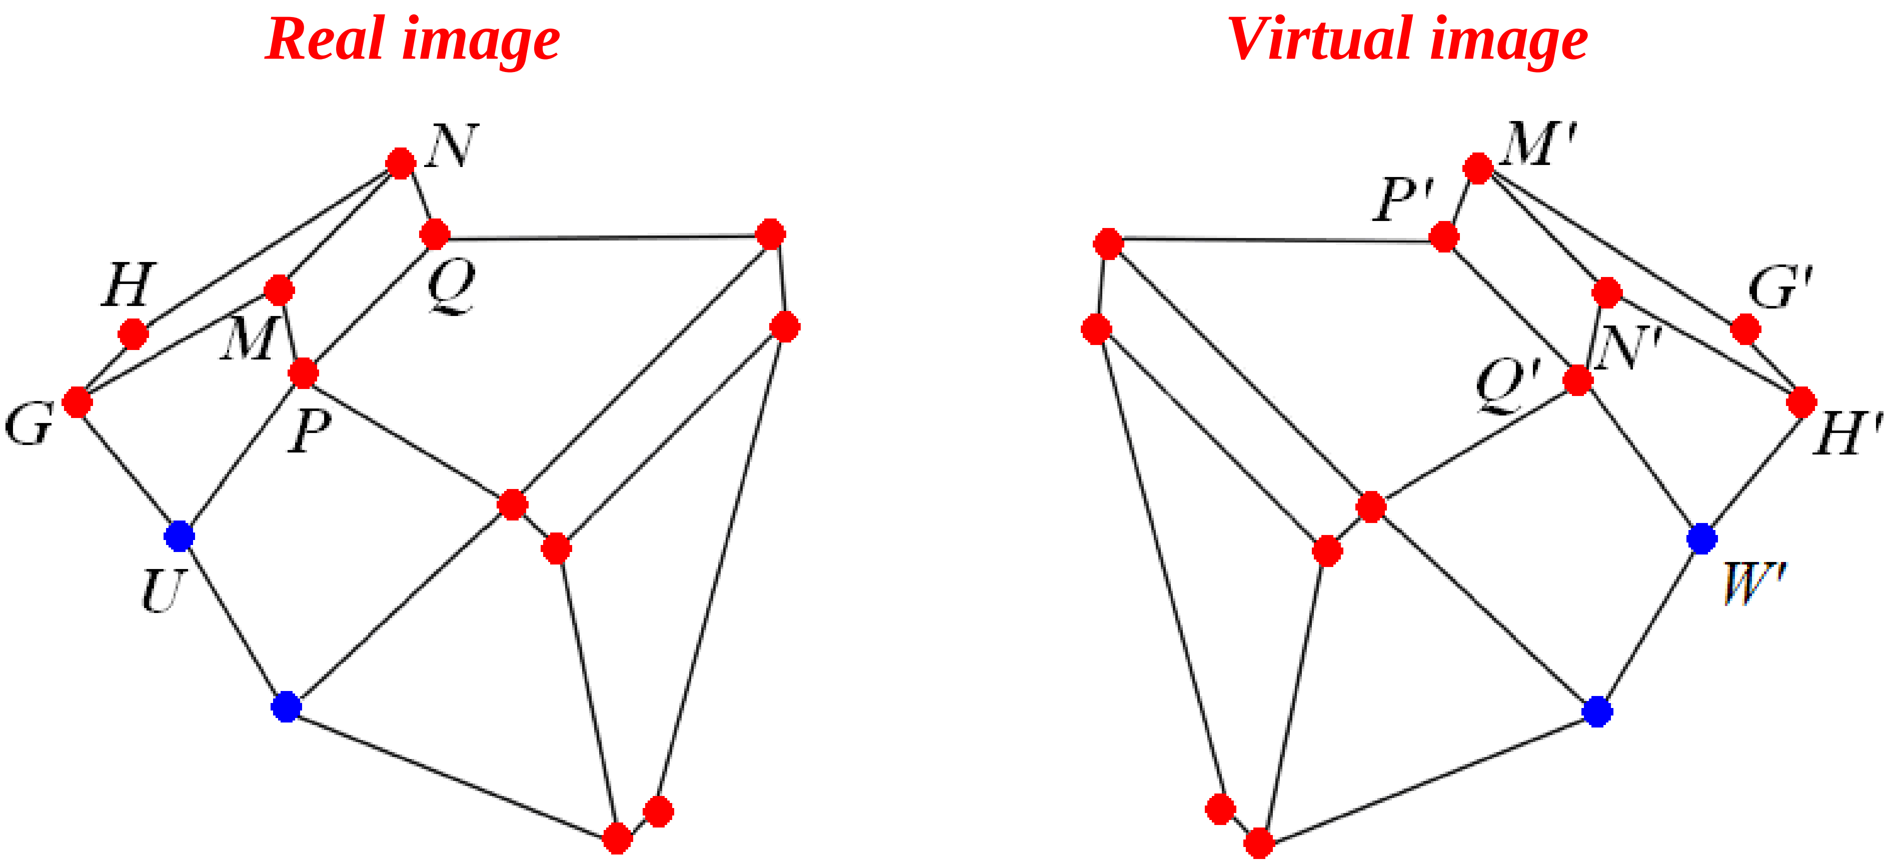


The one-parameter family of possible 3D symmetric interpretations will be shown on the next slide.

# An illustration of a one-parameter family of 3D shapes recovered from the 2D image on the left

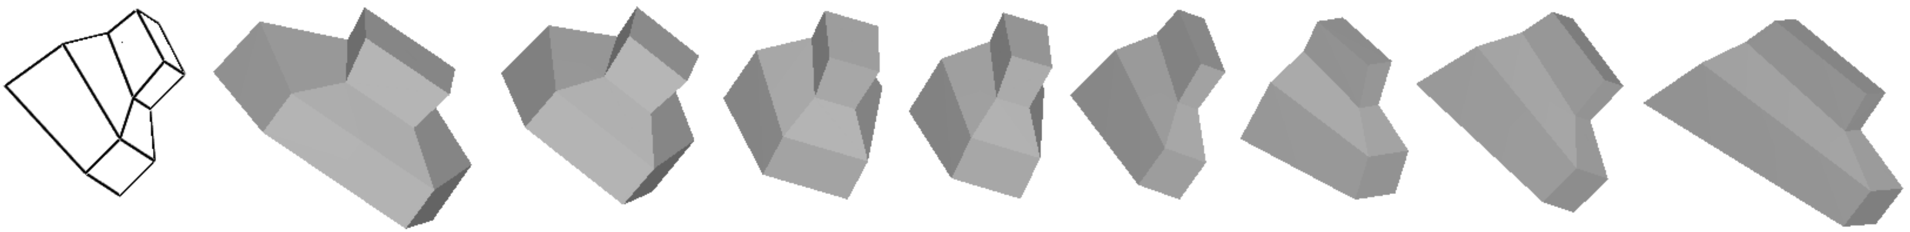

The shaded shapes are possible 3D symmetric interpretations of the line drawing. The 3D shapes differ with respect to an aspect ratio

Now that you have seen the power of 3D mirror symmetry, I will describe the additional constraints used in our 3D recovery model.

# Additional constraints

- Maximal planarity of contours
- Minimum surface area (arg max $V^0/S^3$)
- Maximal 3D compactness (arg max $V^2/S^3$)

Note that these three additional constraints represent the *likelihood* function: they minimize the sensitivity of the 2D image to the changes in the 3D viewing direction.

# Additional constraints

- Maximal planarity of contours
- Minimum surface area (arg max $V^0/S^3$)
- Maximal 3D compactness (arg max $V^2/S^3$)

They implement the *likelihood* for 1D features (contour), 2D features (surface) and 3D features (volume).

Now that you have some intuition about how the model works, let's see a demo of what it can do, and how well it does it…

http://web.ics.purdue.edu/~li135/Demo.html

So far you have seen that the model recovers synthetic 3D objects very well, even when figure-ground organization is noisy.

Next, I will show how well it does with real 2D images of real 3D objects…

http://psyclops.psych.purdue.edu/~tsawada/demo/

You have seen that the model can recover 3D shapes quite well when the shapes are approximately symmetric. But some living symmetric shapes are only rarely geometrically symmetric as they go about their daily business. Humans, for example, are constantly changing the articulation of their arms, legs, heads and torsos. Can our model handle such asymmetric perturbations of a human's 3D shape?

The answer is “yes.” A small elaboration of our symmetry constraint did the trick…

Note, perturbations of symmetry only apply to the configuration of the body’s parts, not to the parts themselves. Today, we will show how the configuration of the body is recovered.

The configuration of this body's parts is geometrically symmetric…

The configuration of this body's parts is geometrically symmetric…

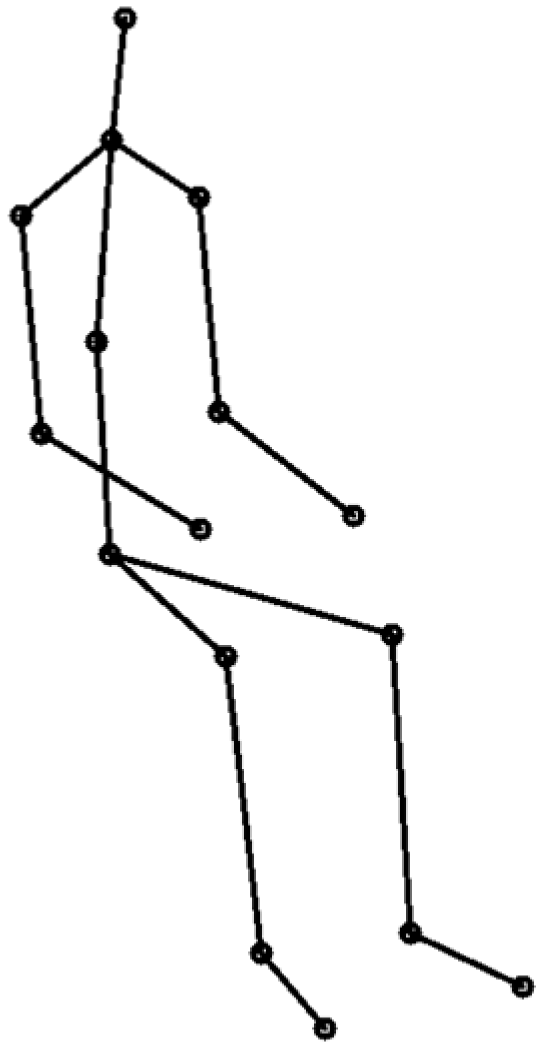

The configuration of this body’s parts is ***not*** geometrically symmetric…

The configuration of this body's parts is ***not*** geometrically symmetric…

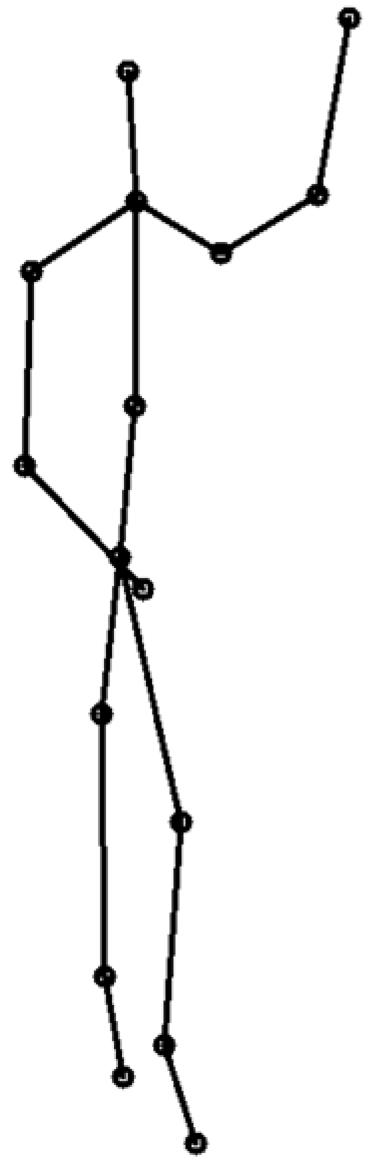

**An asymmetric real body in real action**

**An asymmetric real body in real action**

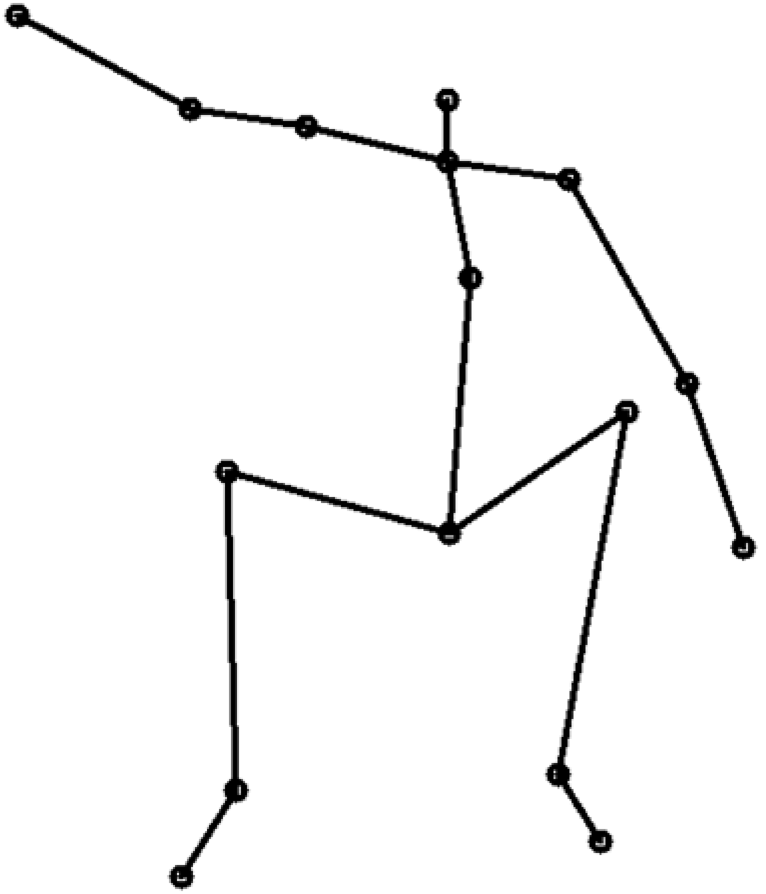

# Summary

- Our model recovers the 3D shape of almost any symmetric or approximately symmetric object.
- No depth cues or familiarity is required.
- Surfaces come *after*, *not before* a 3D shape is recovered. Surfaces can be wrapped around shapes and texture can be added.
- Our model is robust, it tolerates substantial noise in the 2D image.
- It can recover the entire 3D shape, including its back, invisible part.

- We used 4 constraints (symmetry, planarity, minimum surface area and maximum 3D compactness) to recover all of the 3D shapes you have seen.
- Symmetry is special – it represents our *a priori* knowledge about the natural environment – almost all natural objects in natural environments are symmetrical.
- The other 3 constraints can be thought of as representing the probabilistic nature of the viewing conditions, specifically, their *likelihood*, i.e., the recovered 3D shape should be such that the 2D image that was used for the recovery was not degenerate.
- **THE BOTTOM LINE**: only two constraints, *symmetry* and *likelihood* are both necessary and sufficient to recover the 3D shape of an object from one of its 2D images.

# Thank you!