User Interface Software Projects:
Intro to Kinect

Assoc. Professor Donald J. Patterson
INF 134 Winter 2013

Slides adapted from John MacCormick, Guido Gerig:
http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf
http://www.sci.utah.edu/~gerig/CS6320-S2012/Materials/CS6320-CV-S2012-StructuredLight.pdf

How does the Kinect Gesture Tracking Work

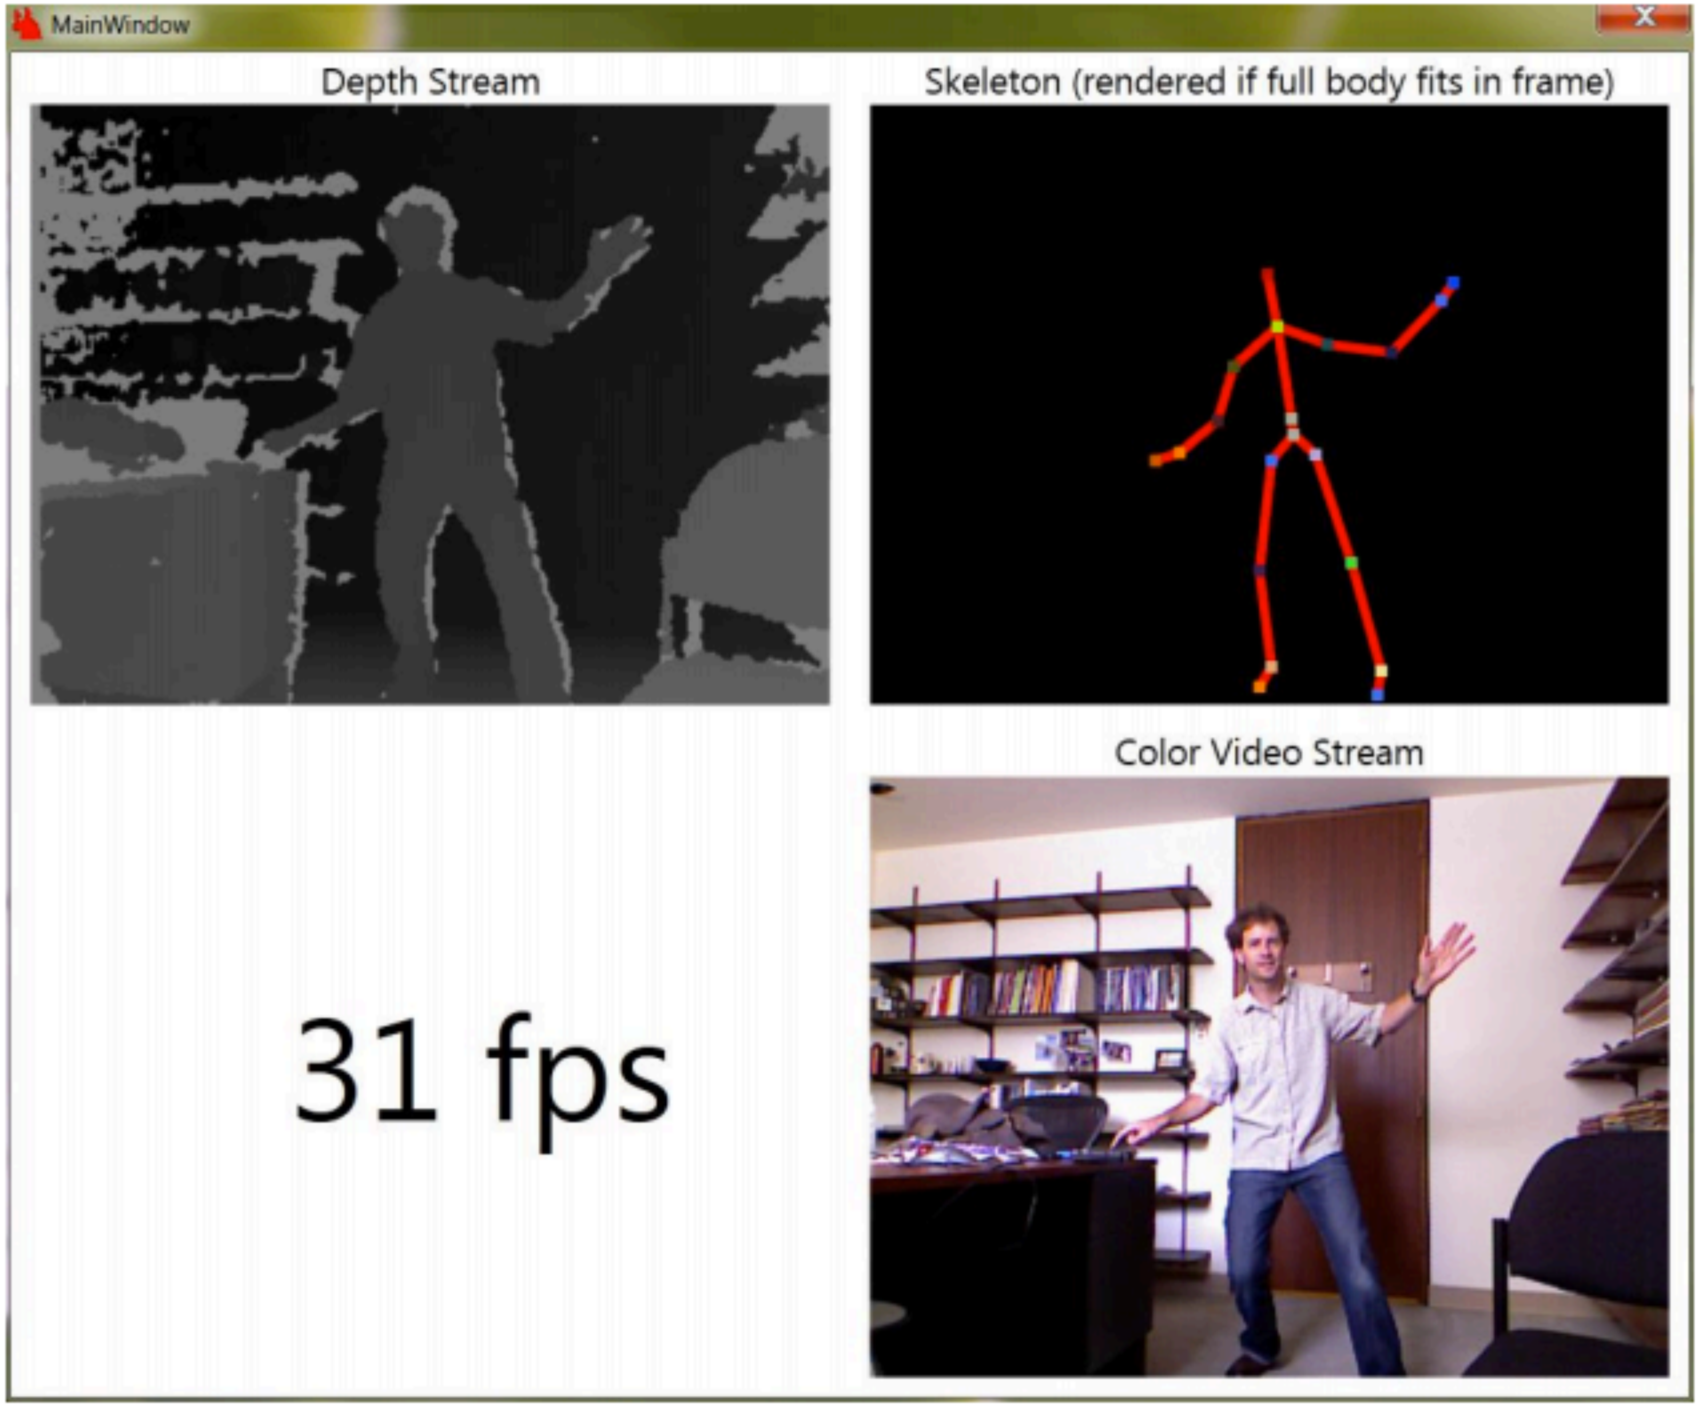

How does the Kinect Gesture Tracking Work

- Depth Video Stream

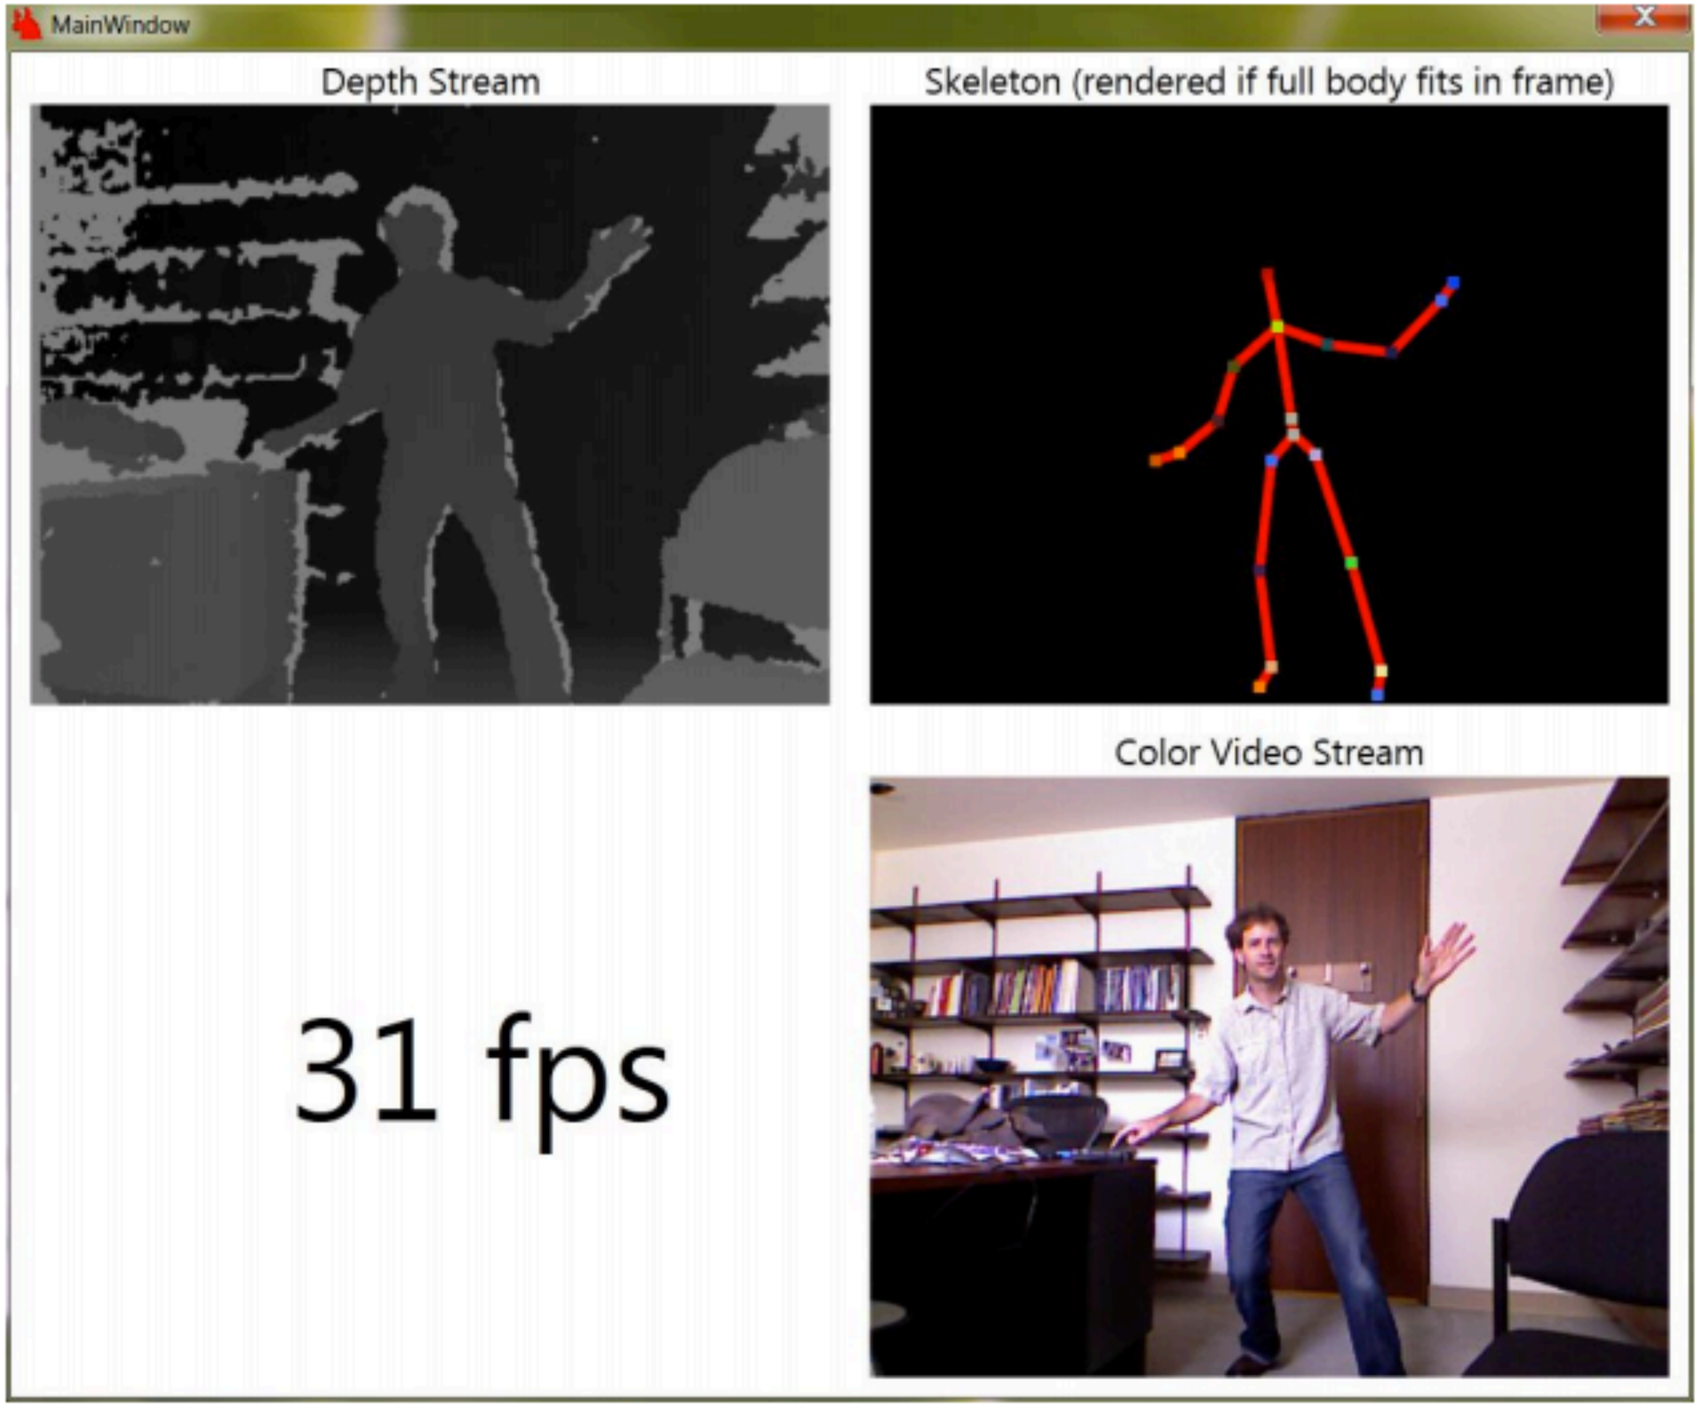

How does the Kinect Gesture Tracking Work

- Depth Video Stream
- Skeleton detection

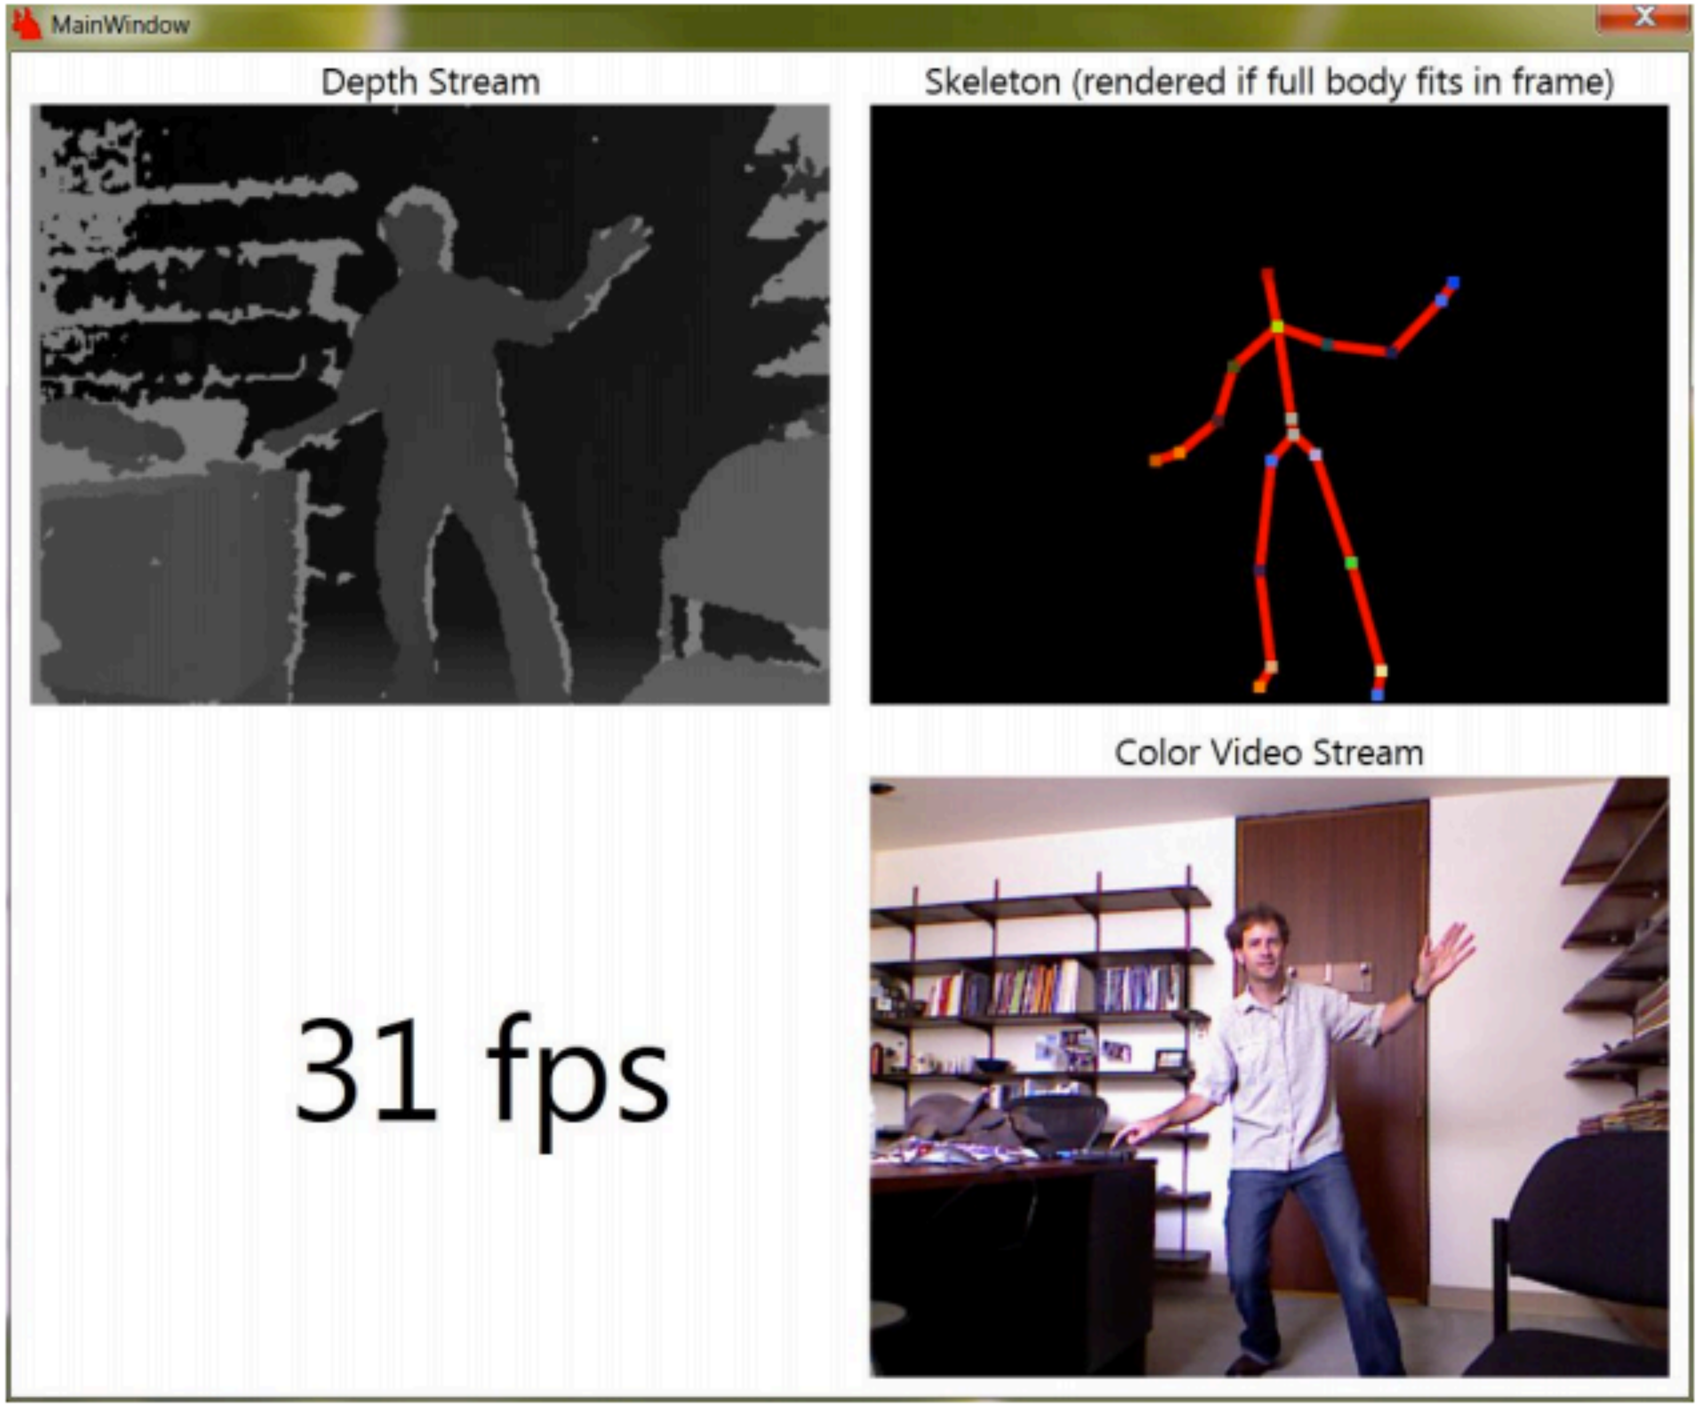

How does the Kinect Gesture Tracking Work

- Depth Video Stream
- Skeleton detection
- Color Video Stream

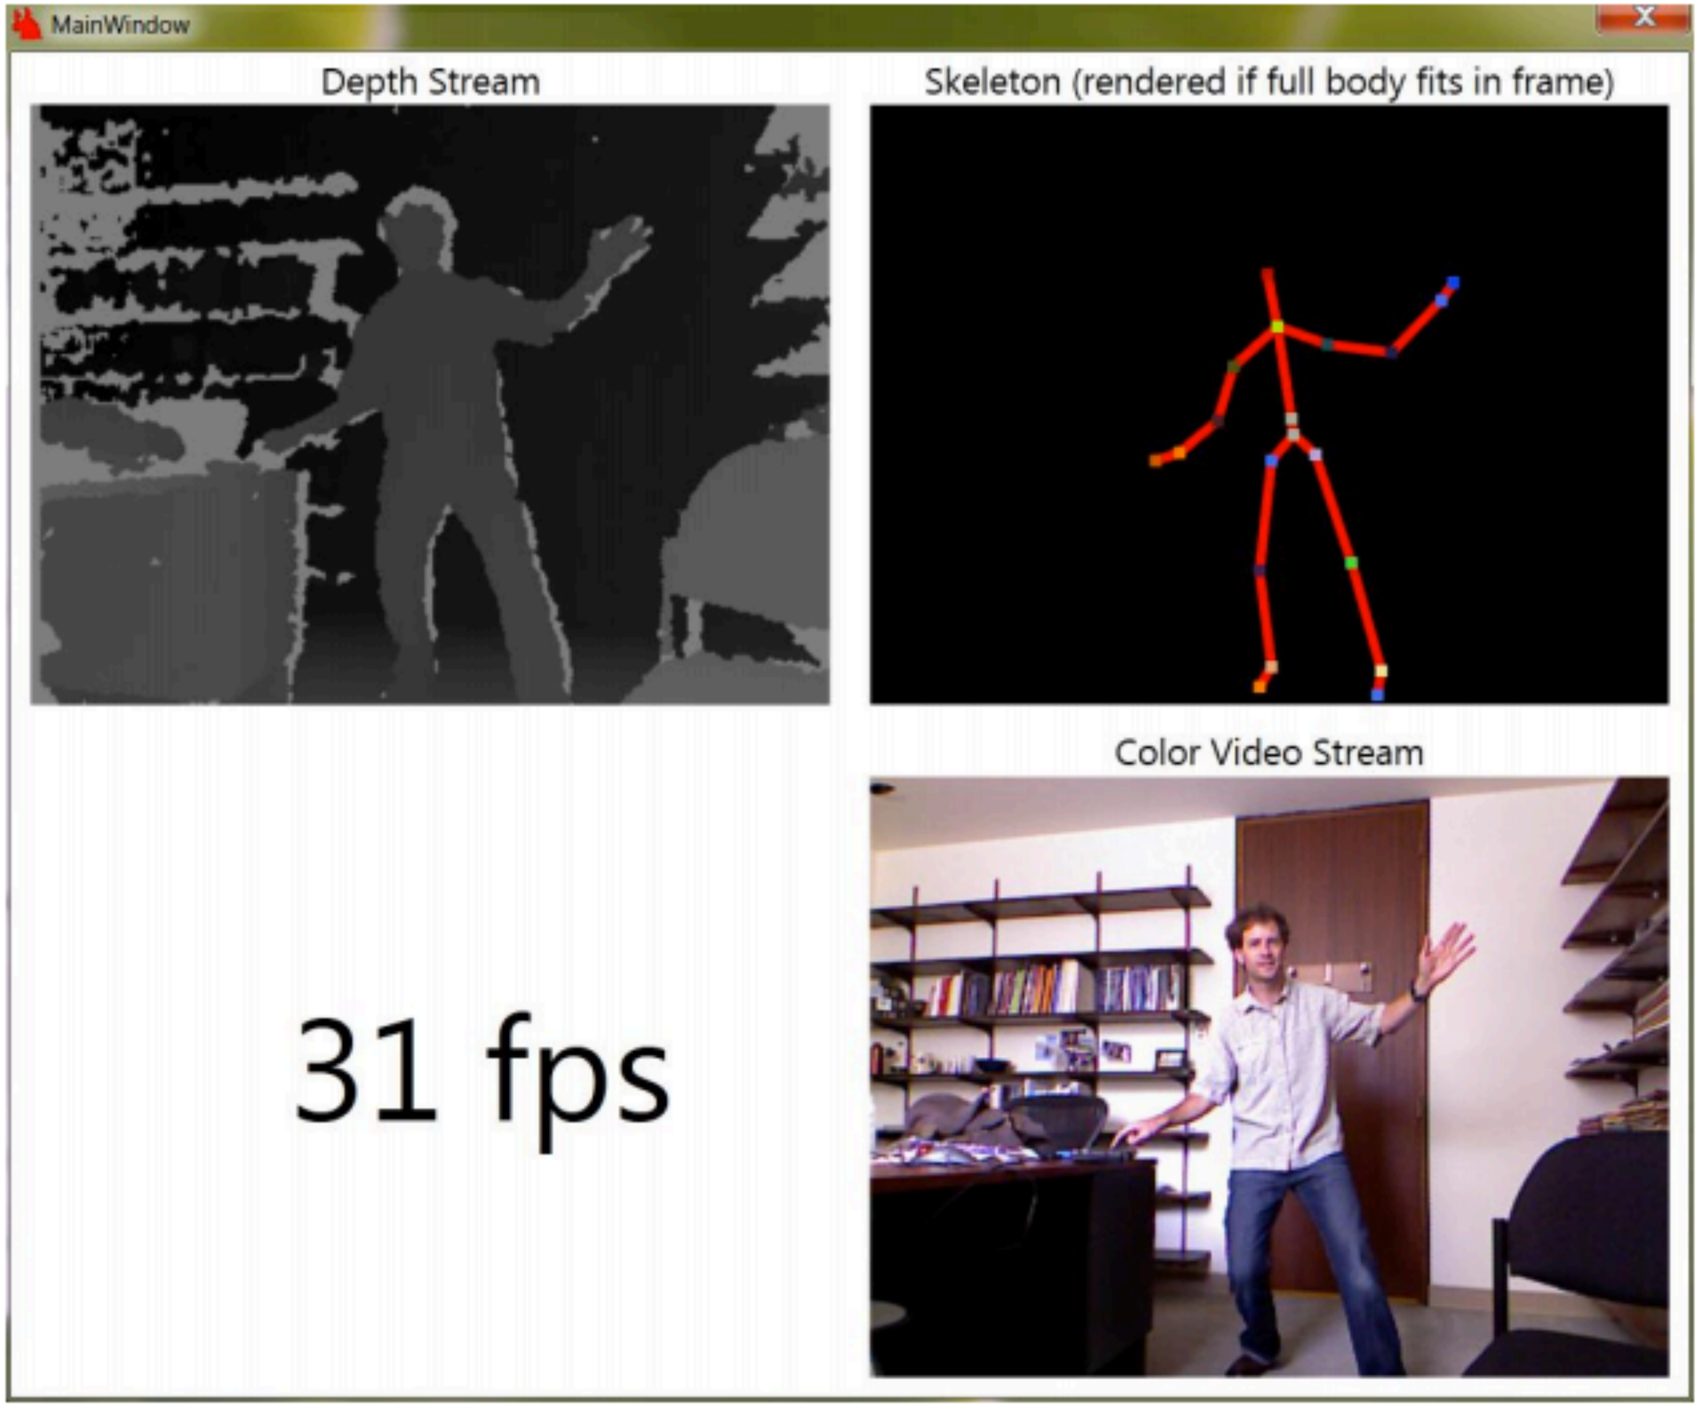

How does the Kinect Gesture Tracking Work

How does the Kinect Gesture Tracking Work

- Several key technologies
  - to compute a depth image
    - Structured light
    - Depth from focus
    - Depth from stereo
  - Machine learning to infer skeleton position

How does the Kinect Gesture Tracking Work

How does the Kinect Gesture Tracking Work

- Structured light
  - The depth map is constructed by analyzing a projected speckle pattern of infrared laser light
    - Microsoft licensed this technology from PrimeSense
  - The depth computation is done by the PrimeSense hardware in the Kinect
  - Details are not public, the following is speculation based on patent applications

How does the Kinect Gesture Tracking Work

How does the Kinect Gesture Tracking Work

- Structured light
  - A Computer Vision concept
  - Based on projecting a known light pattern onto a scene
  - Analyzing how the observed light differs from the known projection
  - Assuming that the differences are due to the topology of the world

How does the Kinect Gesture Tracking Work

How does the Kinect Gesture Tracking Work

- Structured light
  - Technique #1

Structured Light Technique #1

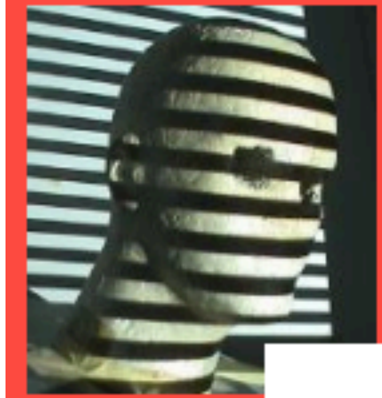

Light Spot Projection 2D

Assume point-wise illumination by laser beam, only 2D

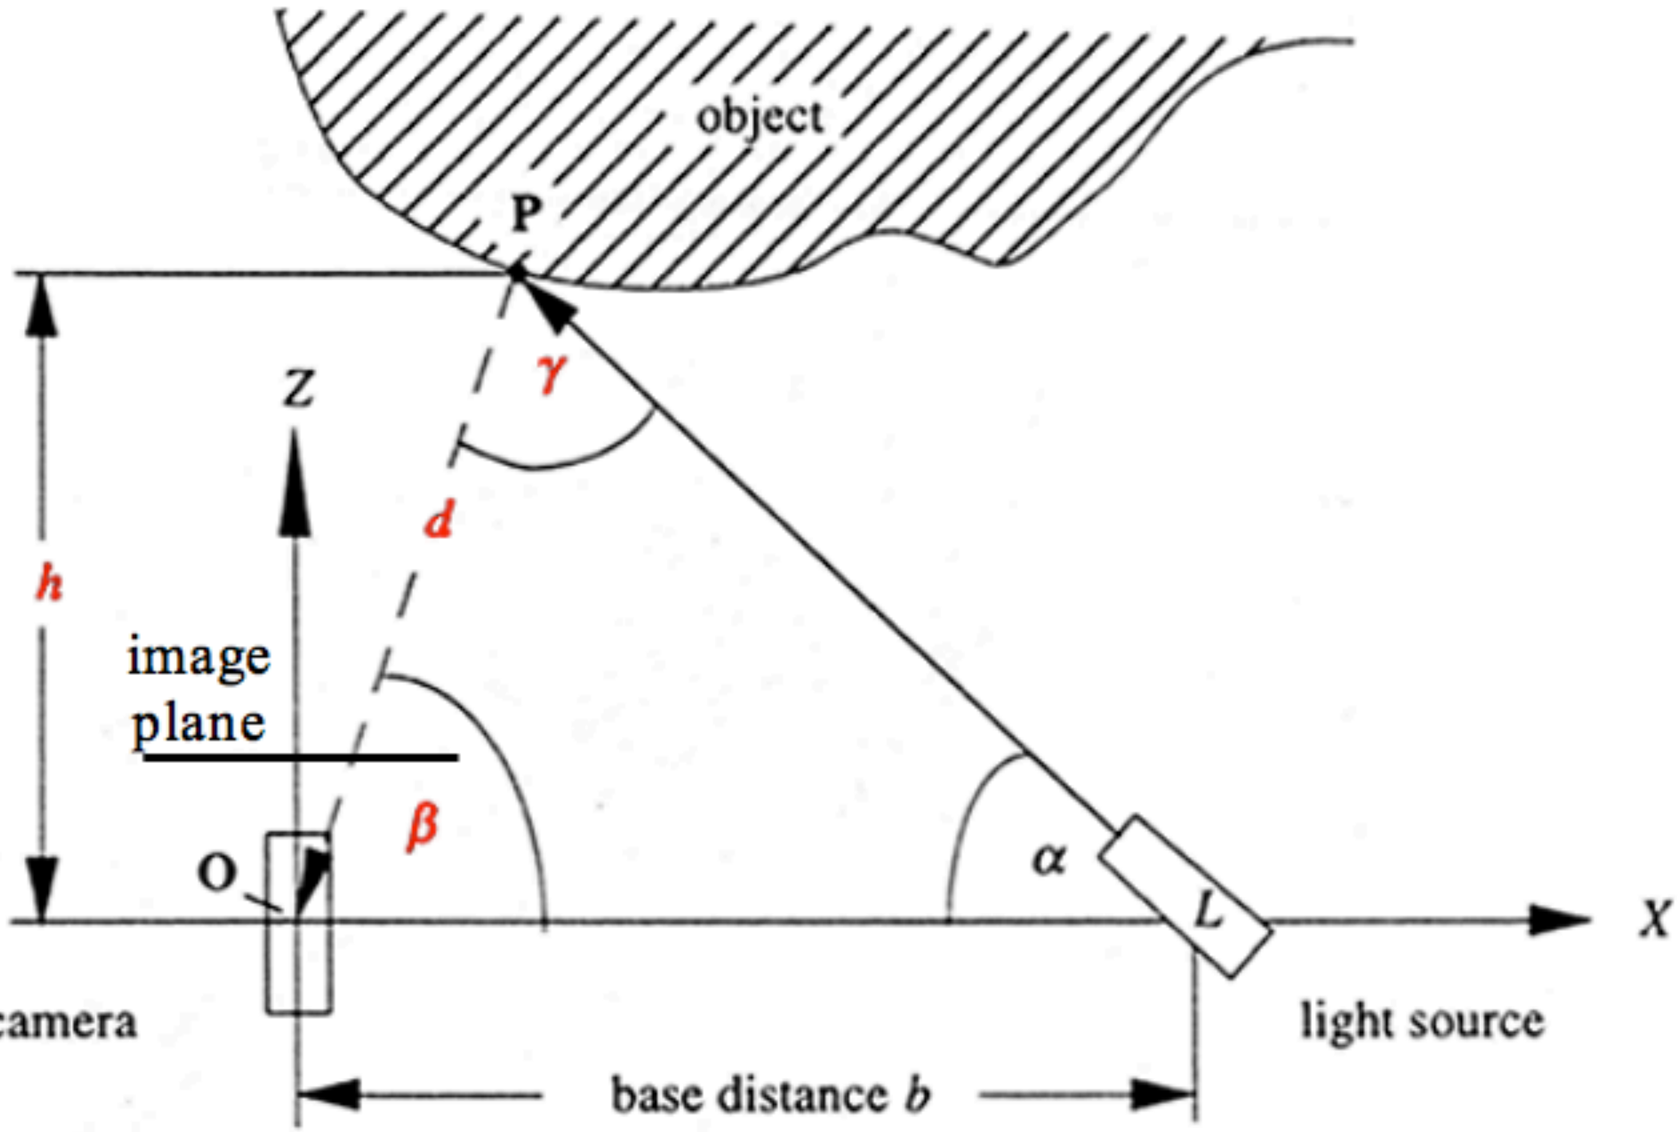

http://www.sci.utah.edu/~gerig/CS6320-S2012/Materials/CS6320-CV-S2012-StructuredLight.pdf

Light Spot Projection 2D

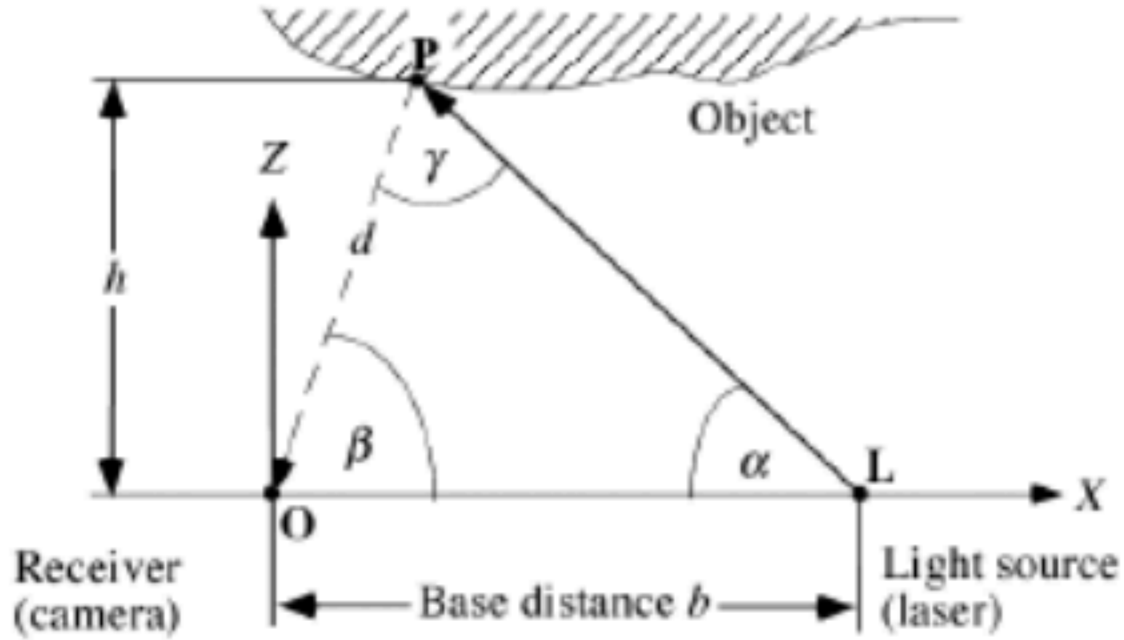

O, **L**, and **P** define a triangle, and we determine the position of P by *triangulation*, using basic formulas about triangles such as the law of sines:

$$\frac{d}{\sin\alpha} = \frac{b}{\sin\gamma}$$

It follows that

$$d = \frac{b\cdot\sin\alpha}{\sin\gamma} = \frac{b\cdot\alpha}{\sin(\pi-\alpha-\beta)} = \frac{b\cdot\alpha}{\sin(\alpha+\beta)}$$

and, finally, $\mathbf{P} = (d\cdot\cos\beta, d\cdot\sin\beta)^T$. Note that β is determined by the position of the projected (illuminated) point **P** in the 1D image.

http://www.sci.utah.edu/~gerig/CS6320-S2012/Materials/CS6320-CV-S2012-StructuredLight.pdf

Structured Light Technique #1

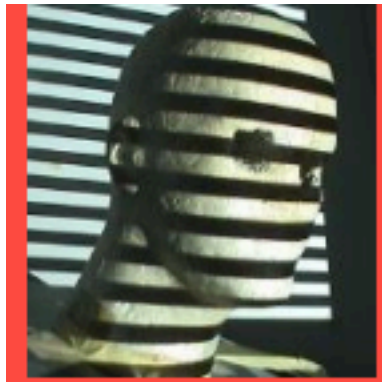

Light Spot Projection 2D

- Coordinates found by triangulation
  - β can be found by projection geometry
  - $d = b*\sin(\alpha)/\sin(\alpha + \beta)$
  - $X_0 = d*\cos(\beta)$
  - $Z_0 = h = d*\sin(\beta)$
- Concept:
  - known b and α
  - β defined by projection geometry
  - Given image coordinate u and focal length f -> calculate β
  - Given b, α, β -> calculate d

http://www.sci.utah.edu/~gerig/CS6320-S2012/Materials/CS6320-CV-S2012-StructuredLight.pdf

How does the Kinect Gesture Tracking Work

How does the Kinect Gesture Tracking Work

- Structured light
  - Technique #2

Light Stripe Scanning – Single Stripe

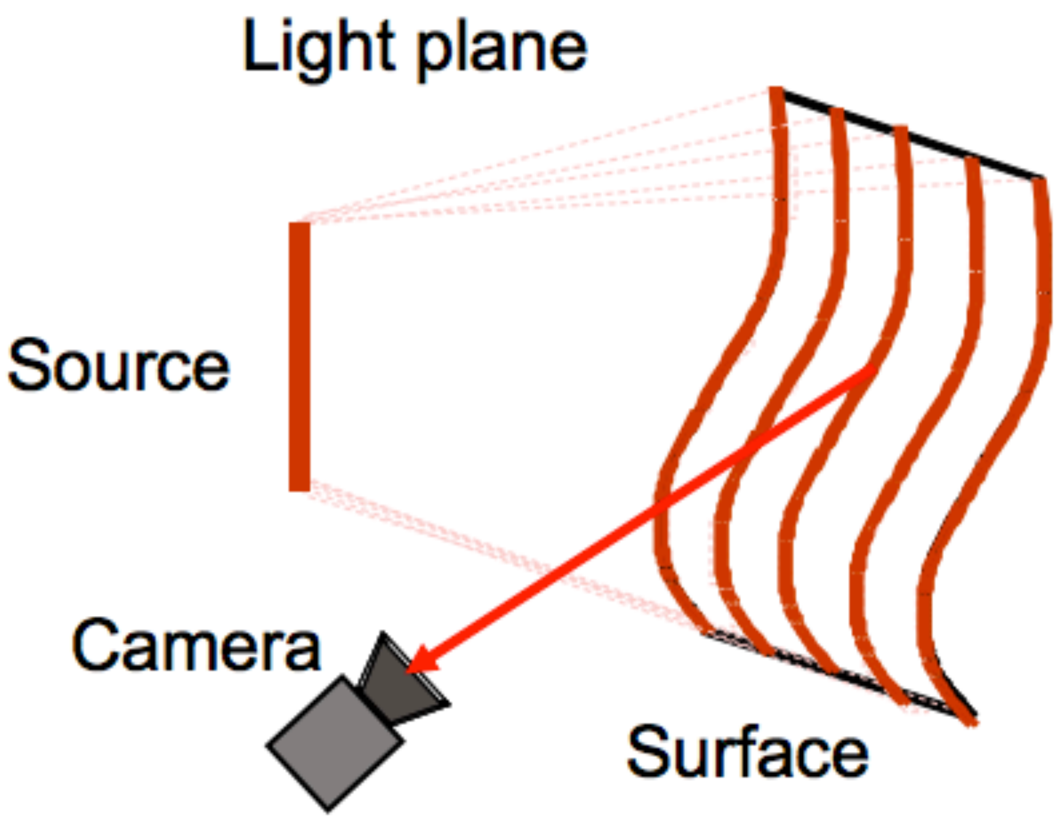

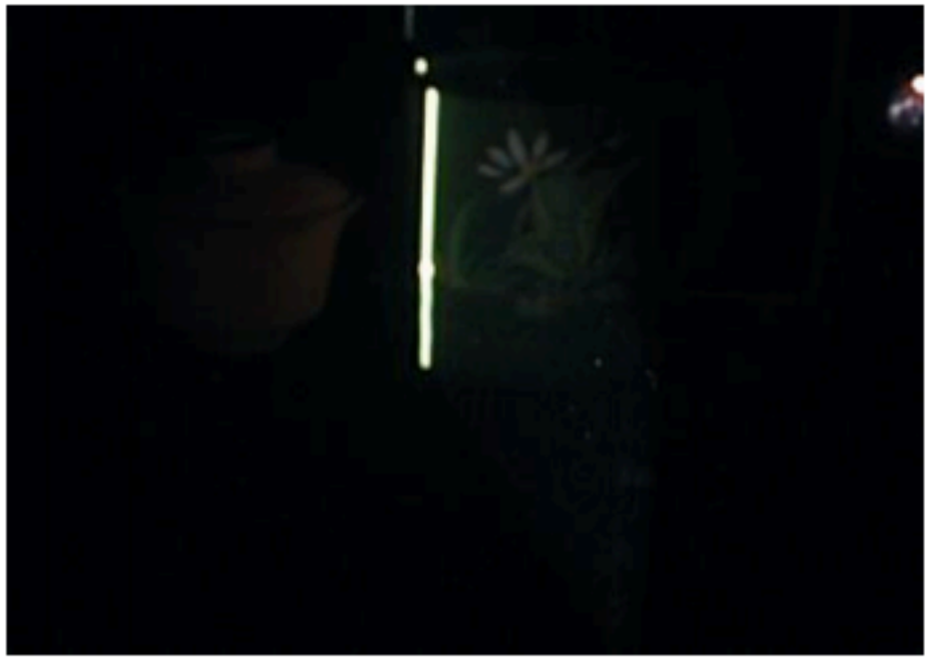

- Optical triangulation
  - Project a single stripe of laser light
  - Scan it across the surface of the object
  - This is a very precise version of structured light scanning
  - Good for high resolution 3D, but needs many images and takes time

Courtesy S. Narasimhan, CMU

http://www.sci.utah.edu/~gerig/CS6320-S2012/Materials/CS6320-CV-S2012-StructuredLight.pdf

Structured Light Technique #2

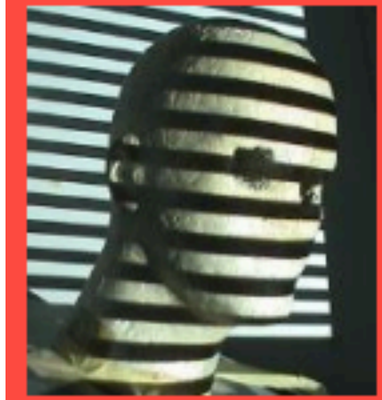

Light Stripe Projection

Structured light is the projection of a light pattern (ray, plane, grid, encoded light, and so forth) under calibrated geometric conditions onto an object whose shape needs to be recovered.

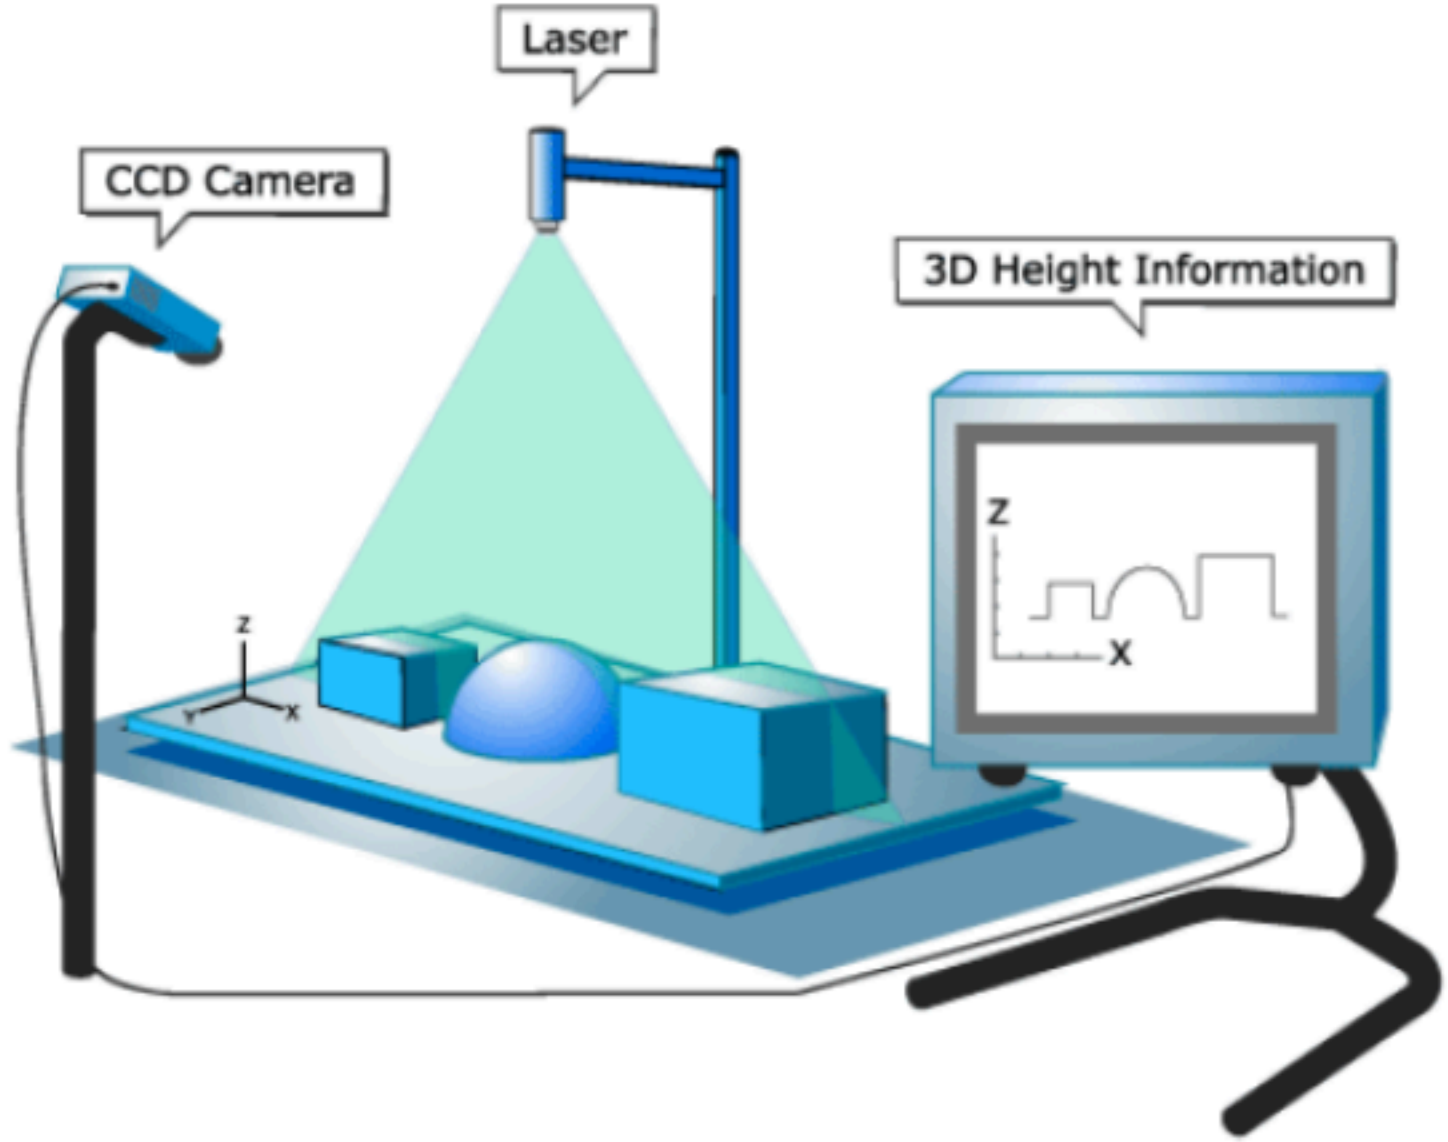

http://www.sci.utah.edu/~gerig/CS6320-S2012/Materials/CS6320-CV-S2012-StructuredLight.pdf

Structured Light Technique #2

Triangulation

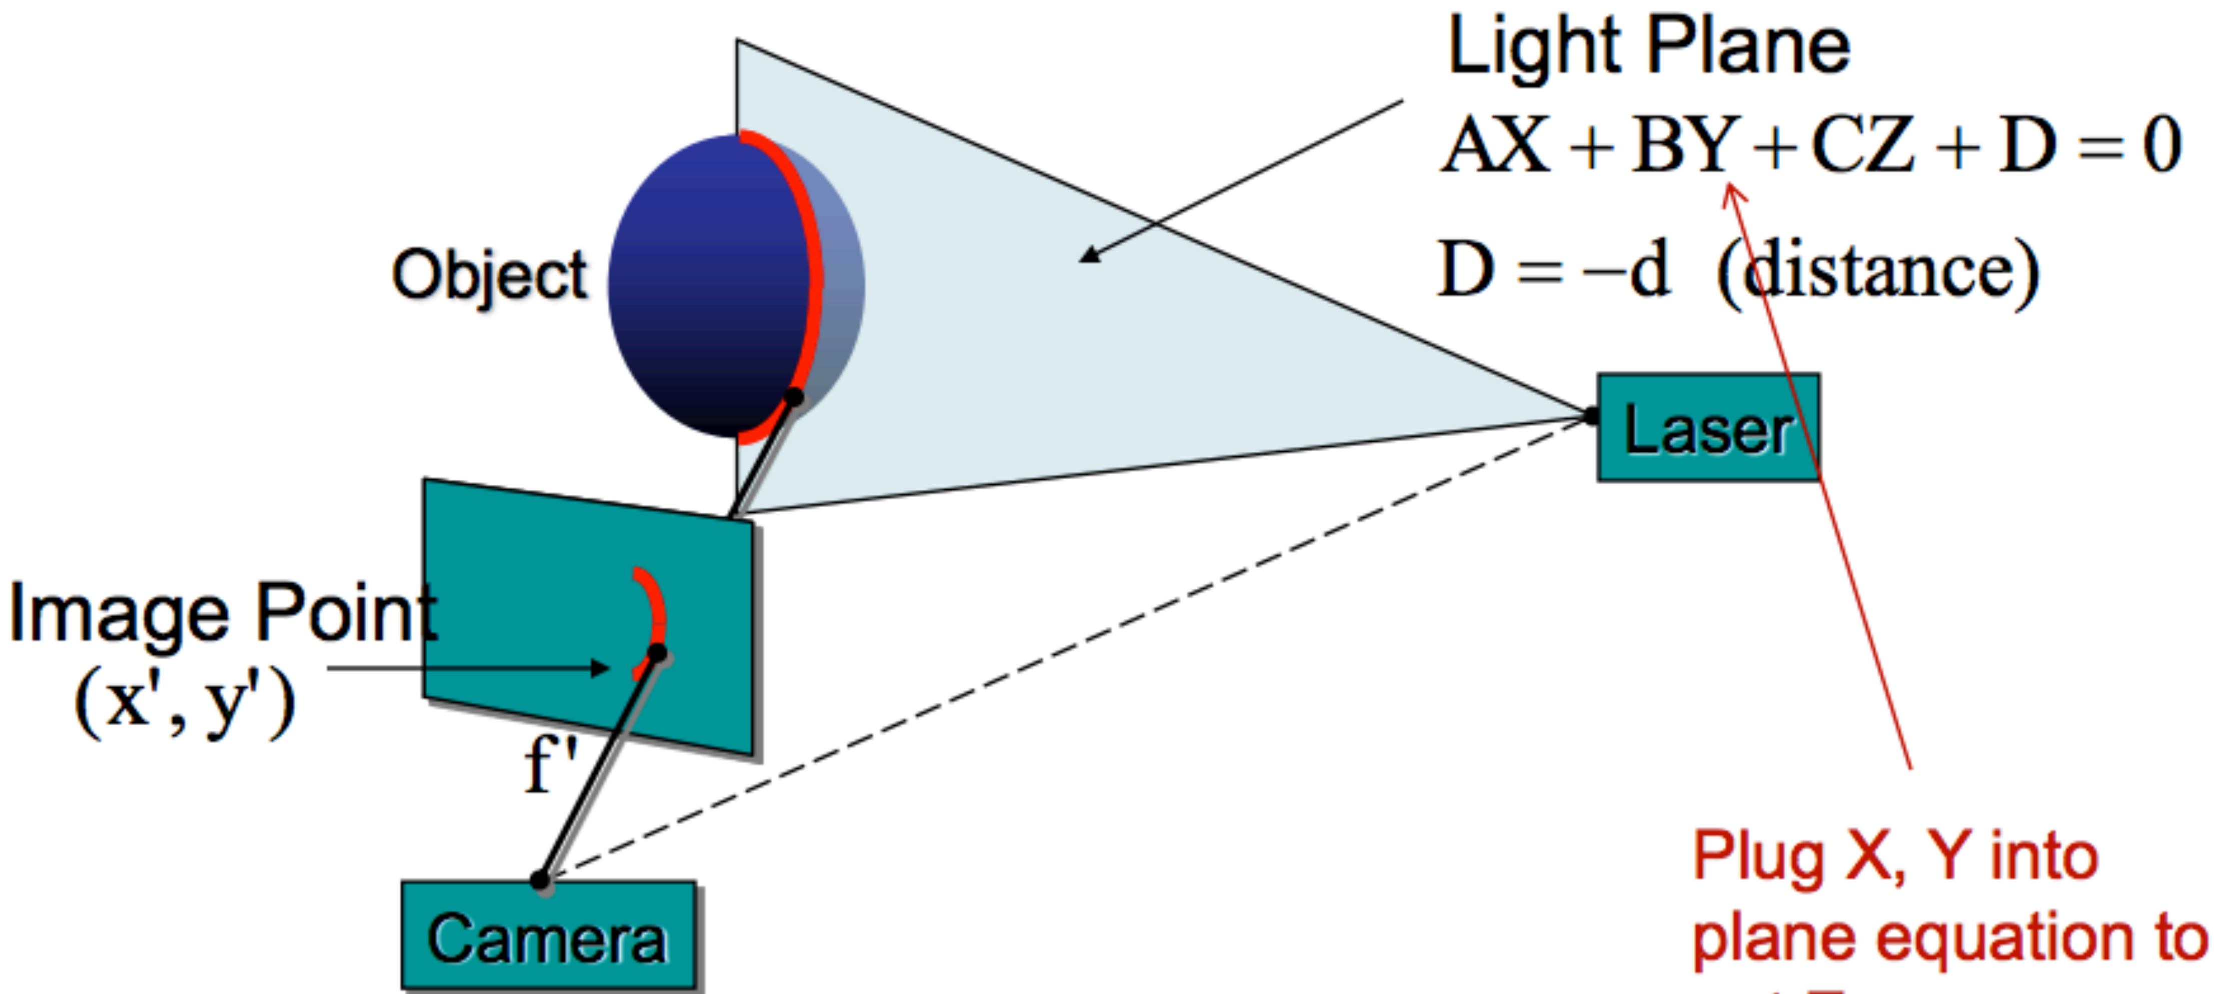

- Depth from ray-plane triangulation:
  - Intersect camera ray with light plane

$$\boxed{X = x'Z / f' \qquad Y = y'Z / f' \qquad Z = \frac{-Df'}{Ax' + By' + Cf'}}$$

Courtesy S. Narasimhan, CMU

http://www.sci.utah.edu/~gerig/CS6320-S2012/Materials/CS6320-CV-S2012-StructuredLight.pdf

Structured Light Technique #2

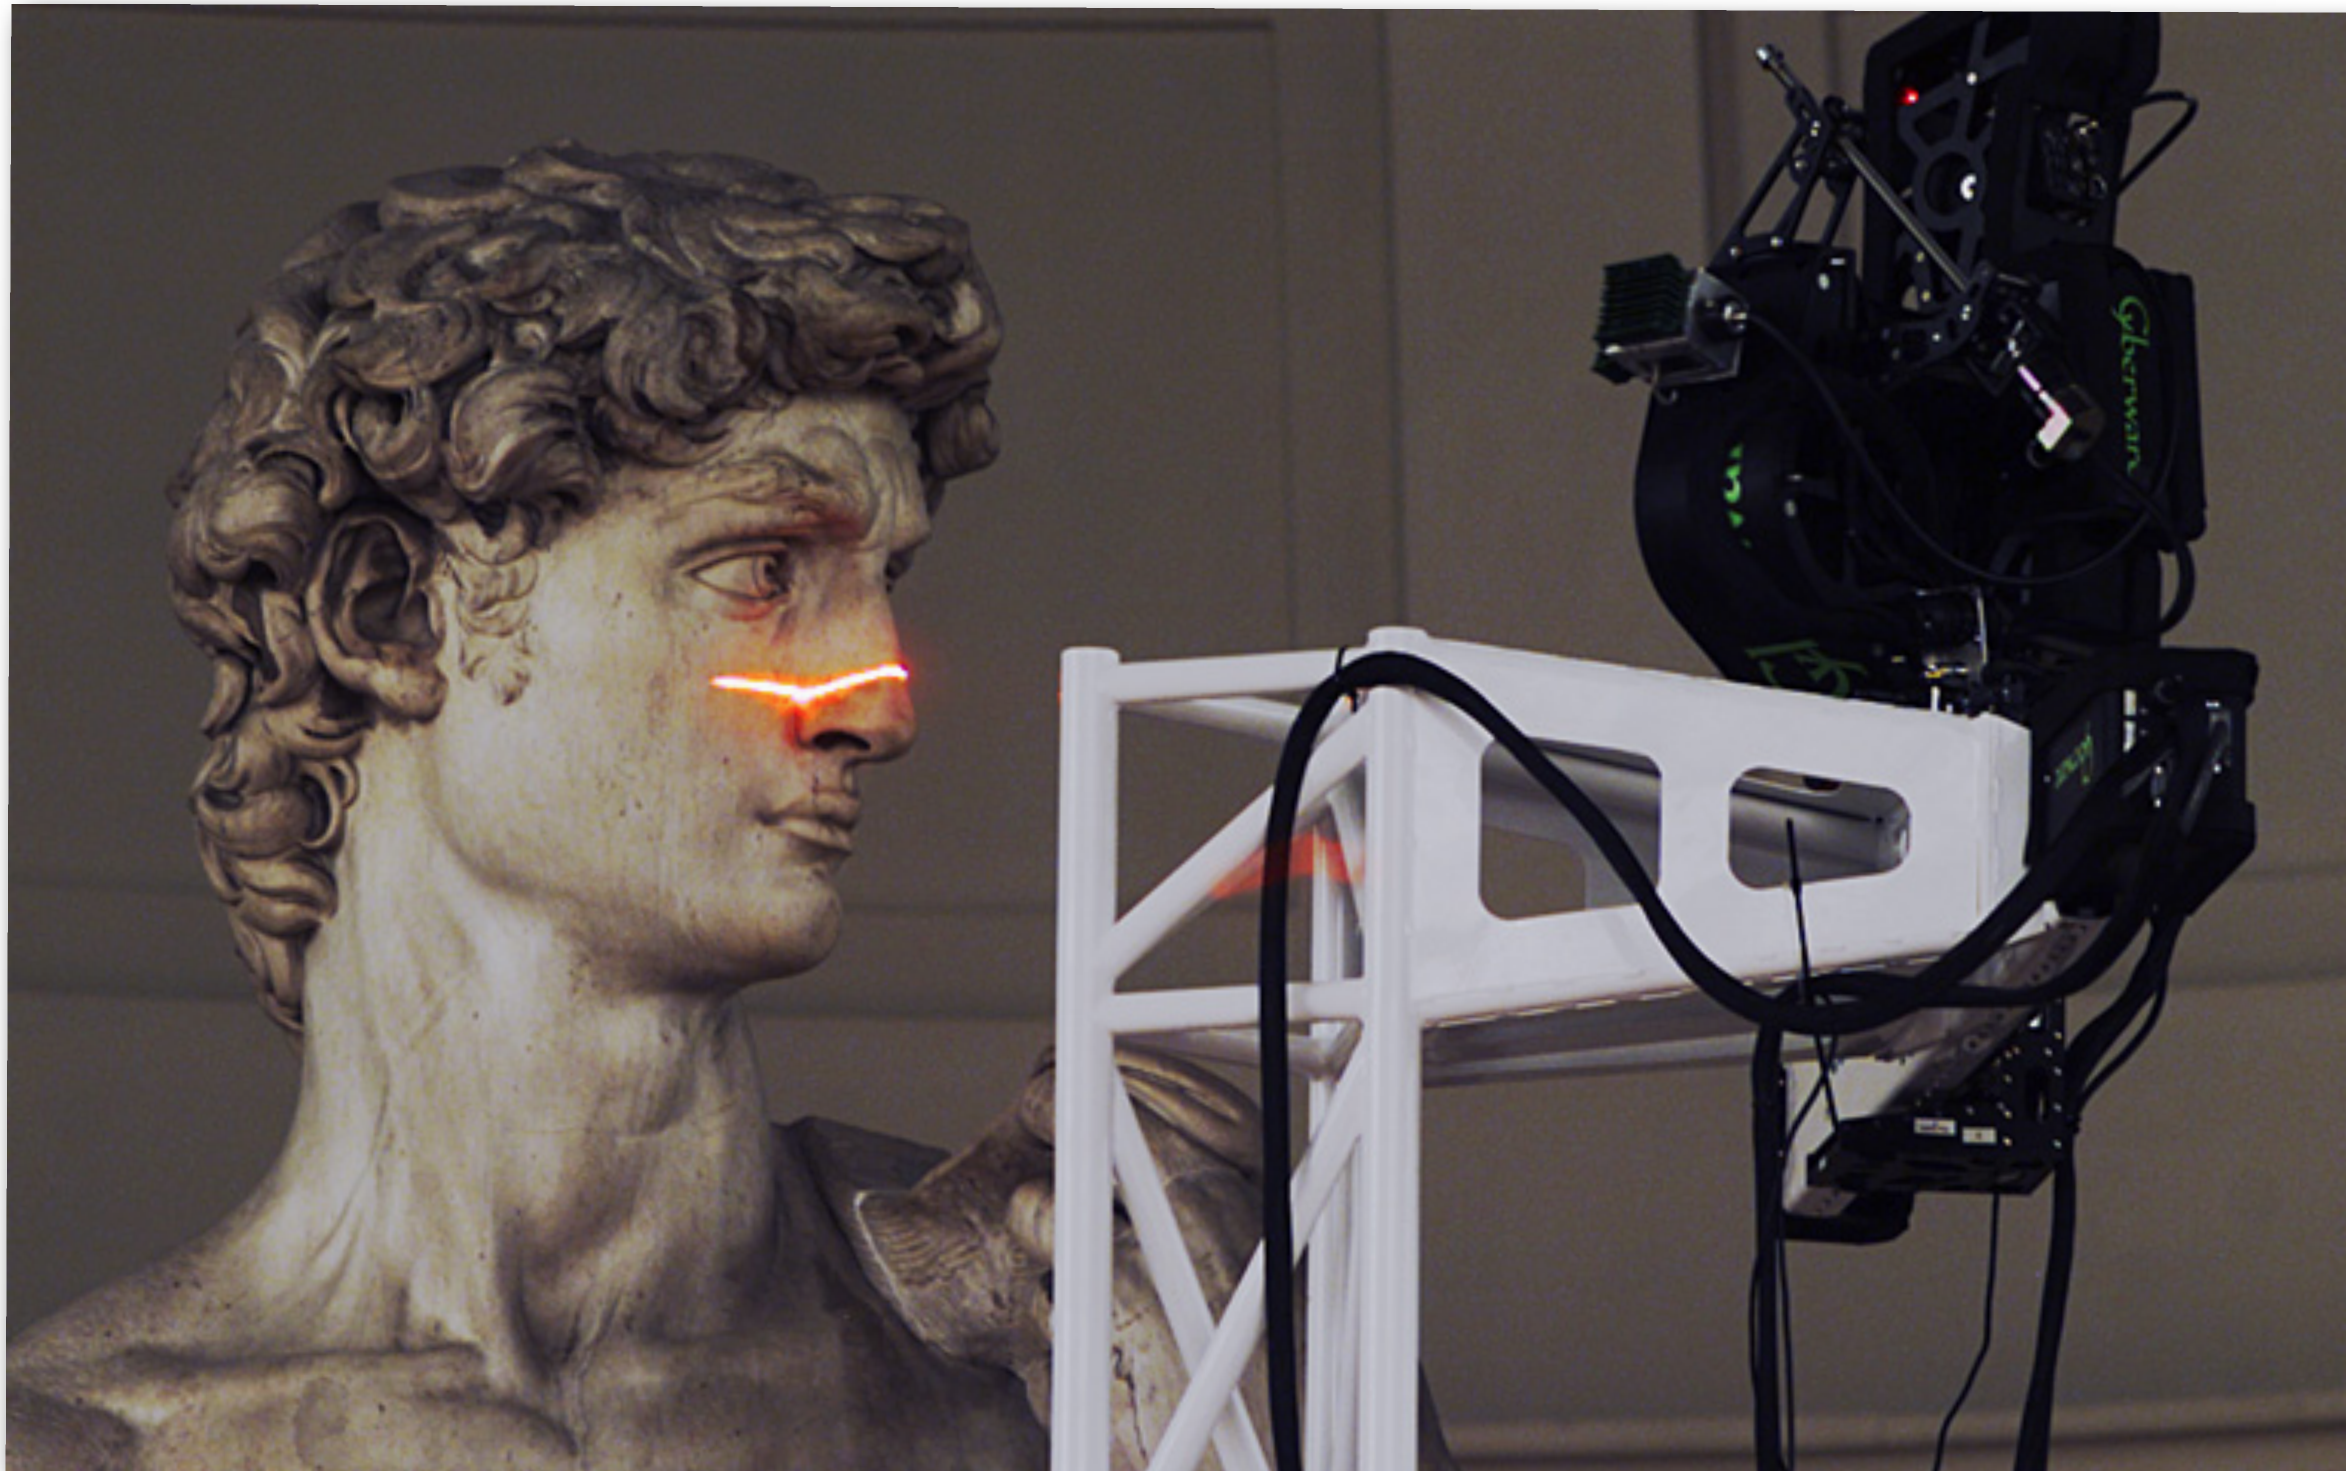

http://graphics.stanford.edu/projects/mich/more-david/scanner-head-and-david-head-s.jpg

Application of Structured Light Scanning

Application of Structured Light Scanning

http://www.flickr.com/photos/nathaninsandiego/6165296066/sizes/z/in/photostream/

Application of Structured Light Scanning

http://graphics.stanford.edu/projects/mich/head-of-david/head-of-david.html

Application of Structured Light Scanning

Application of Structured Light Scanning

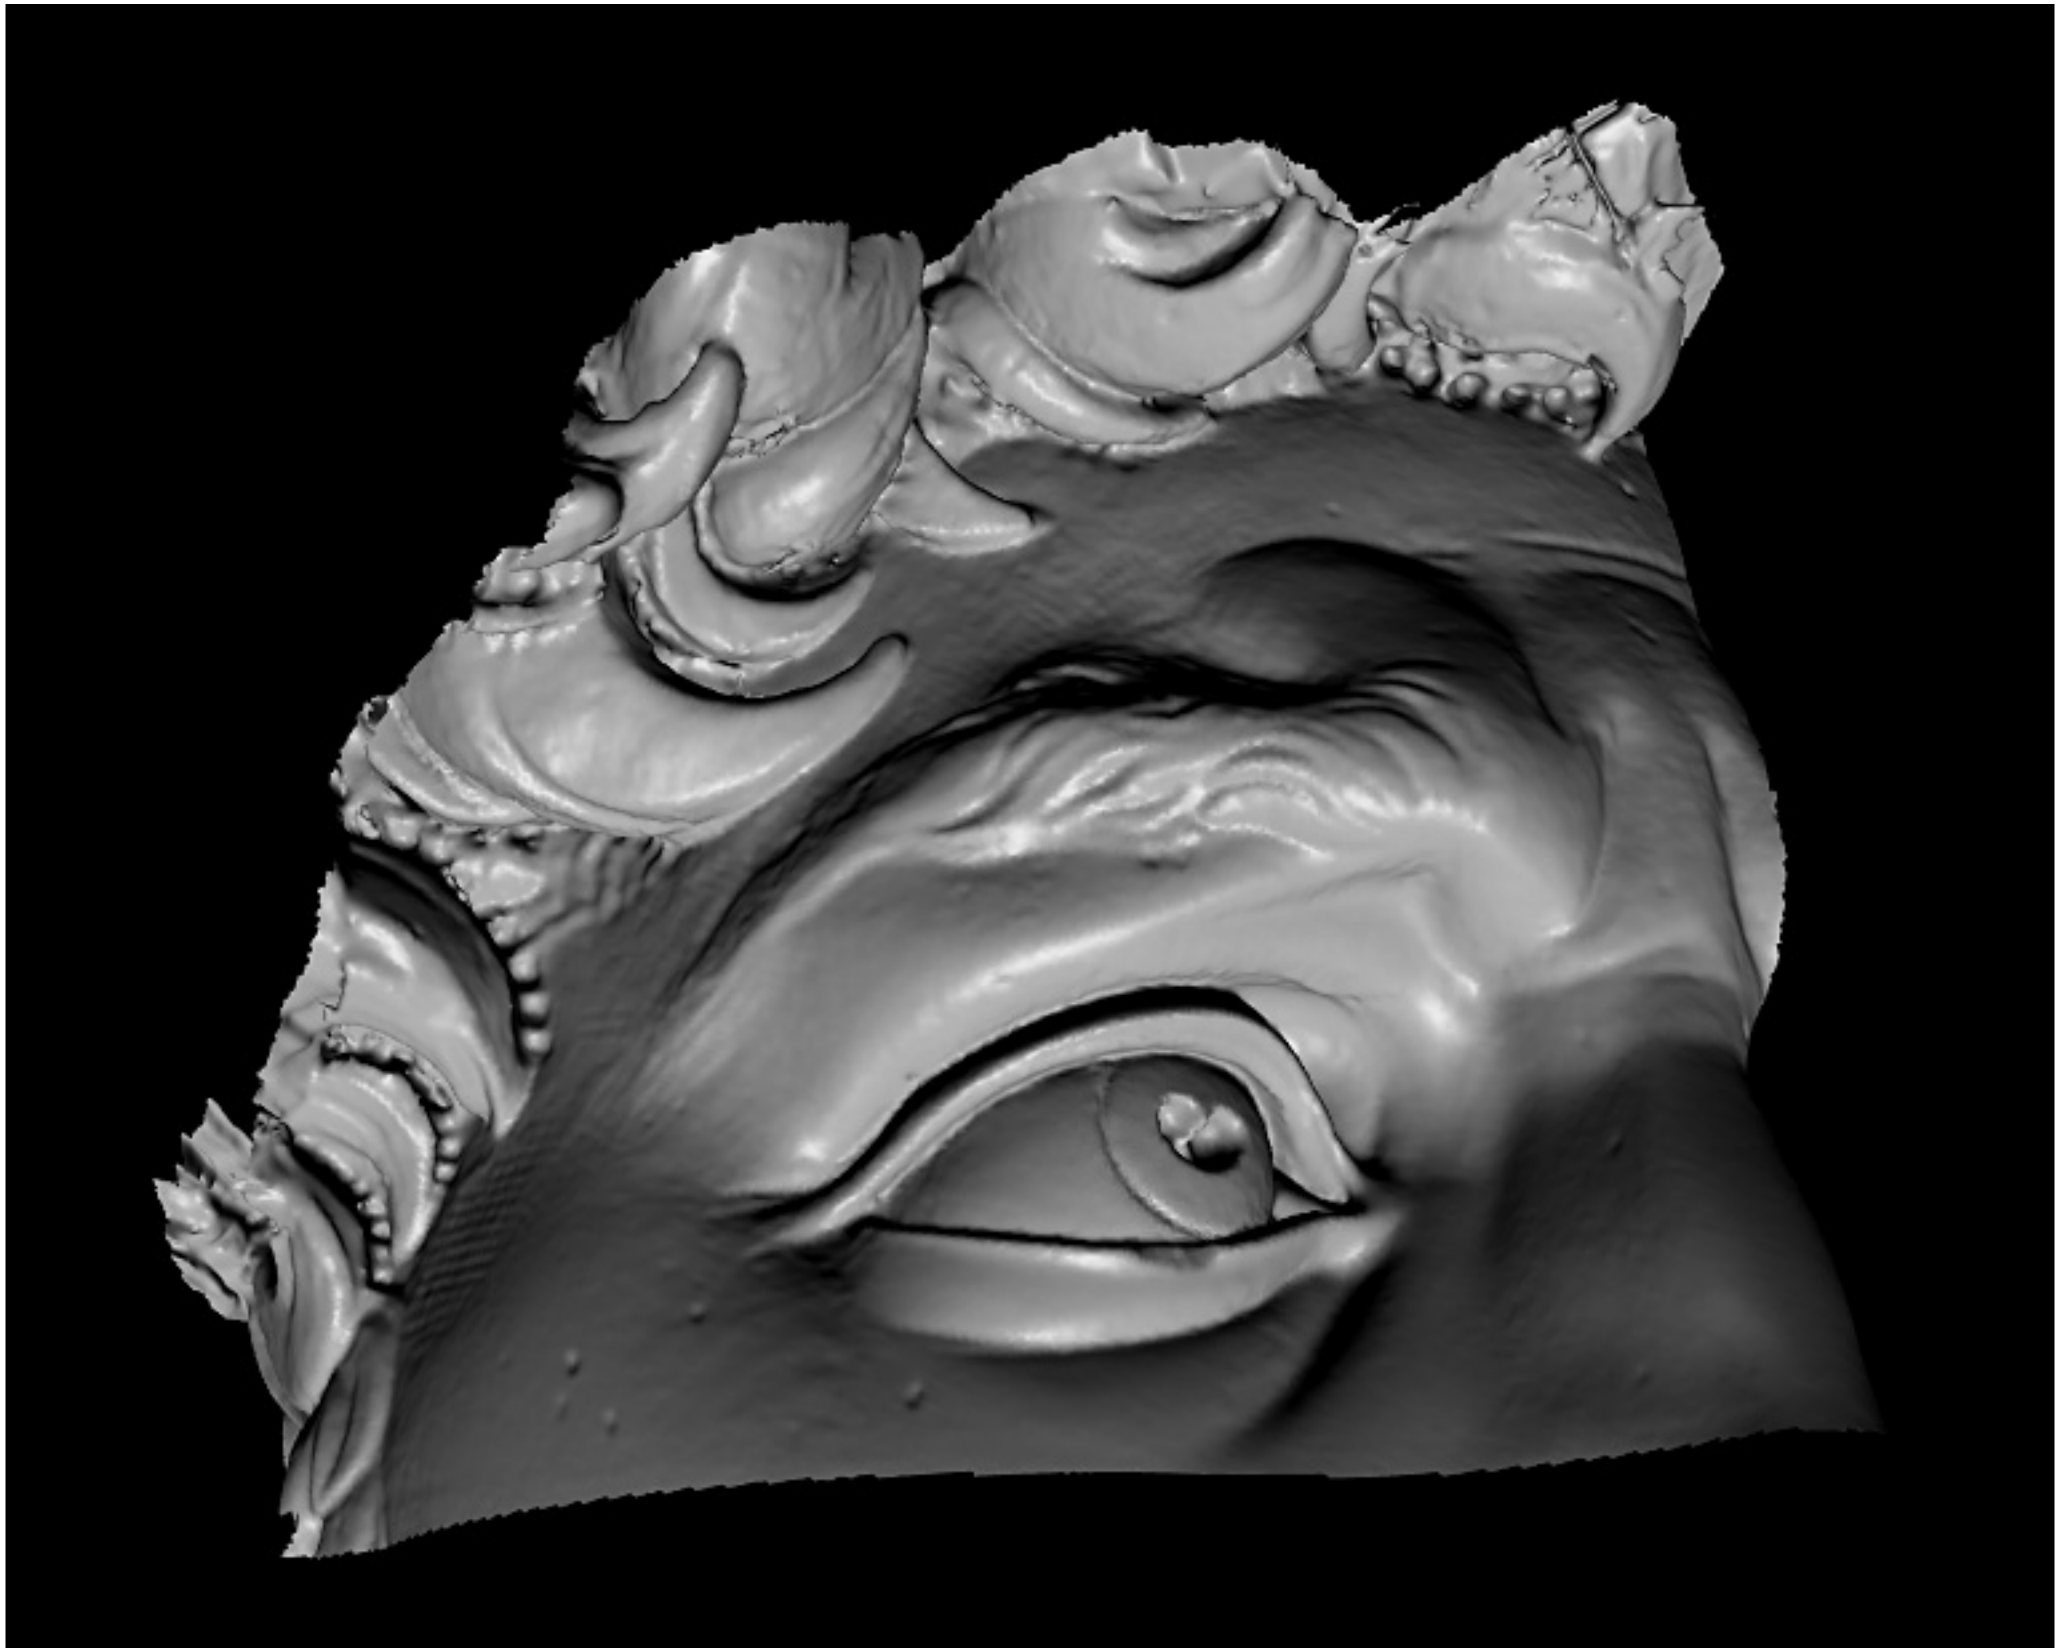

Application of Structured Light Scanning

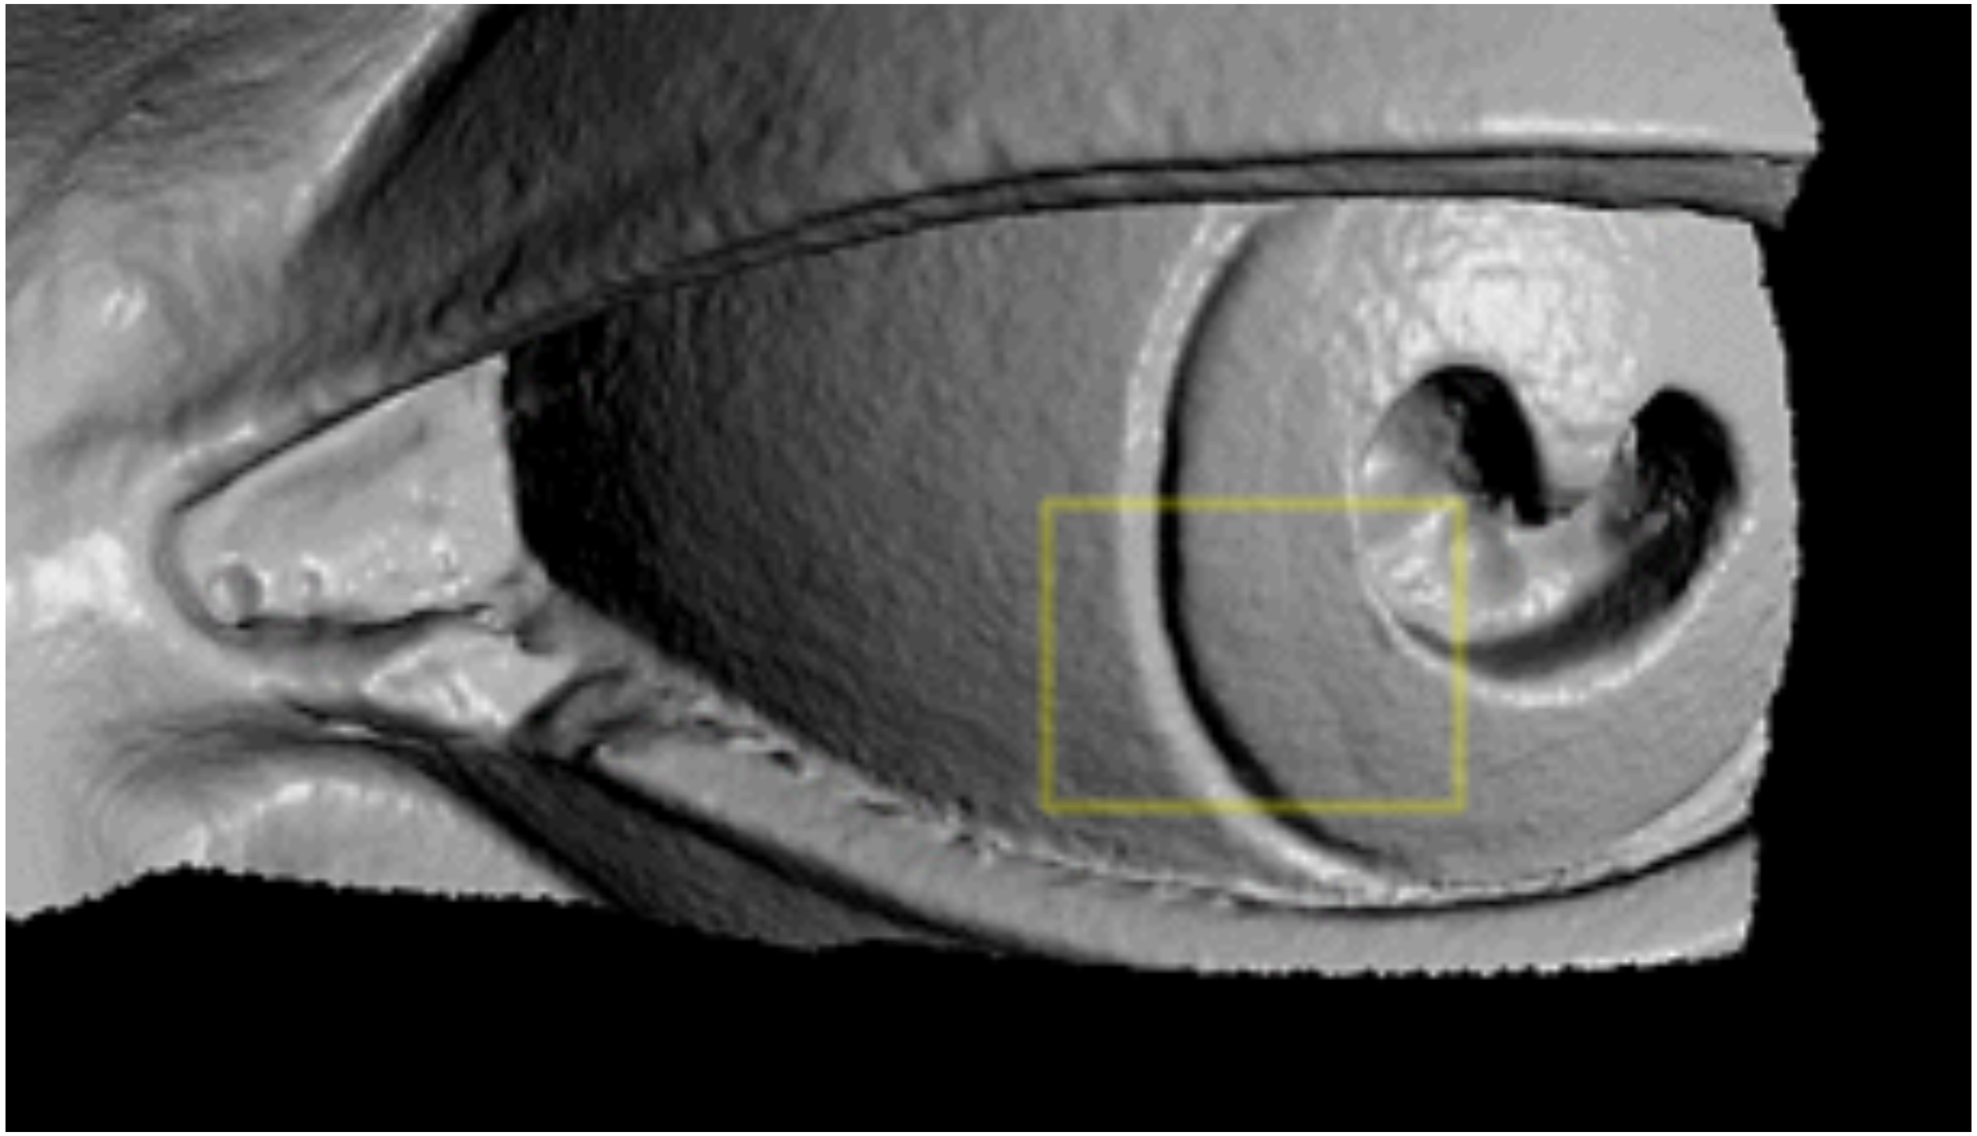

Application of Structured Light Scanning

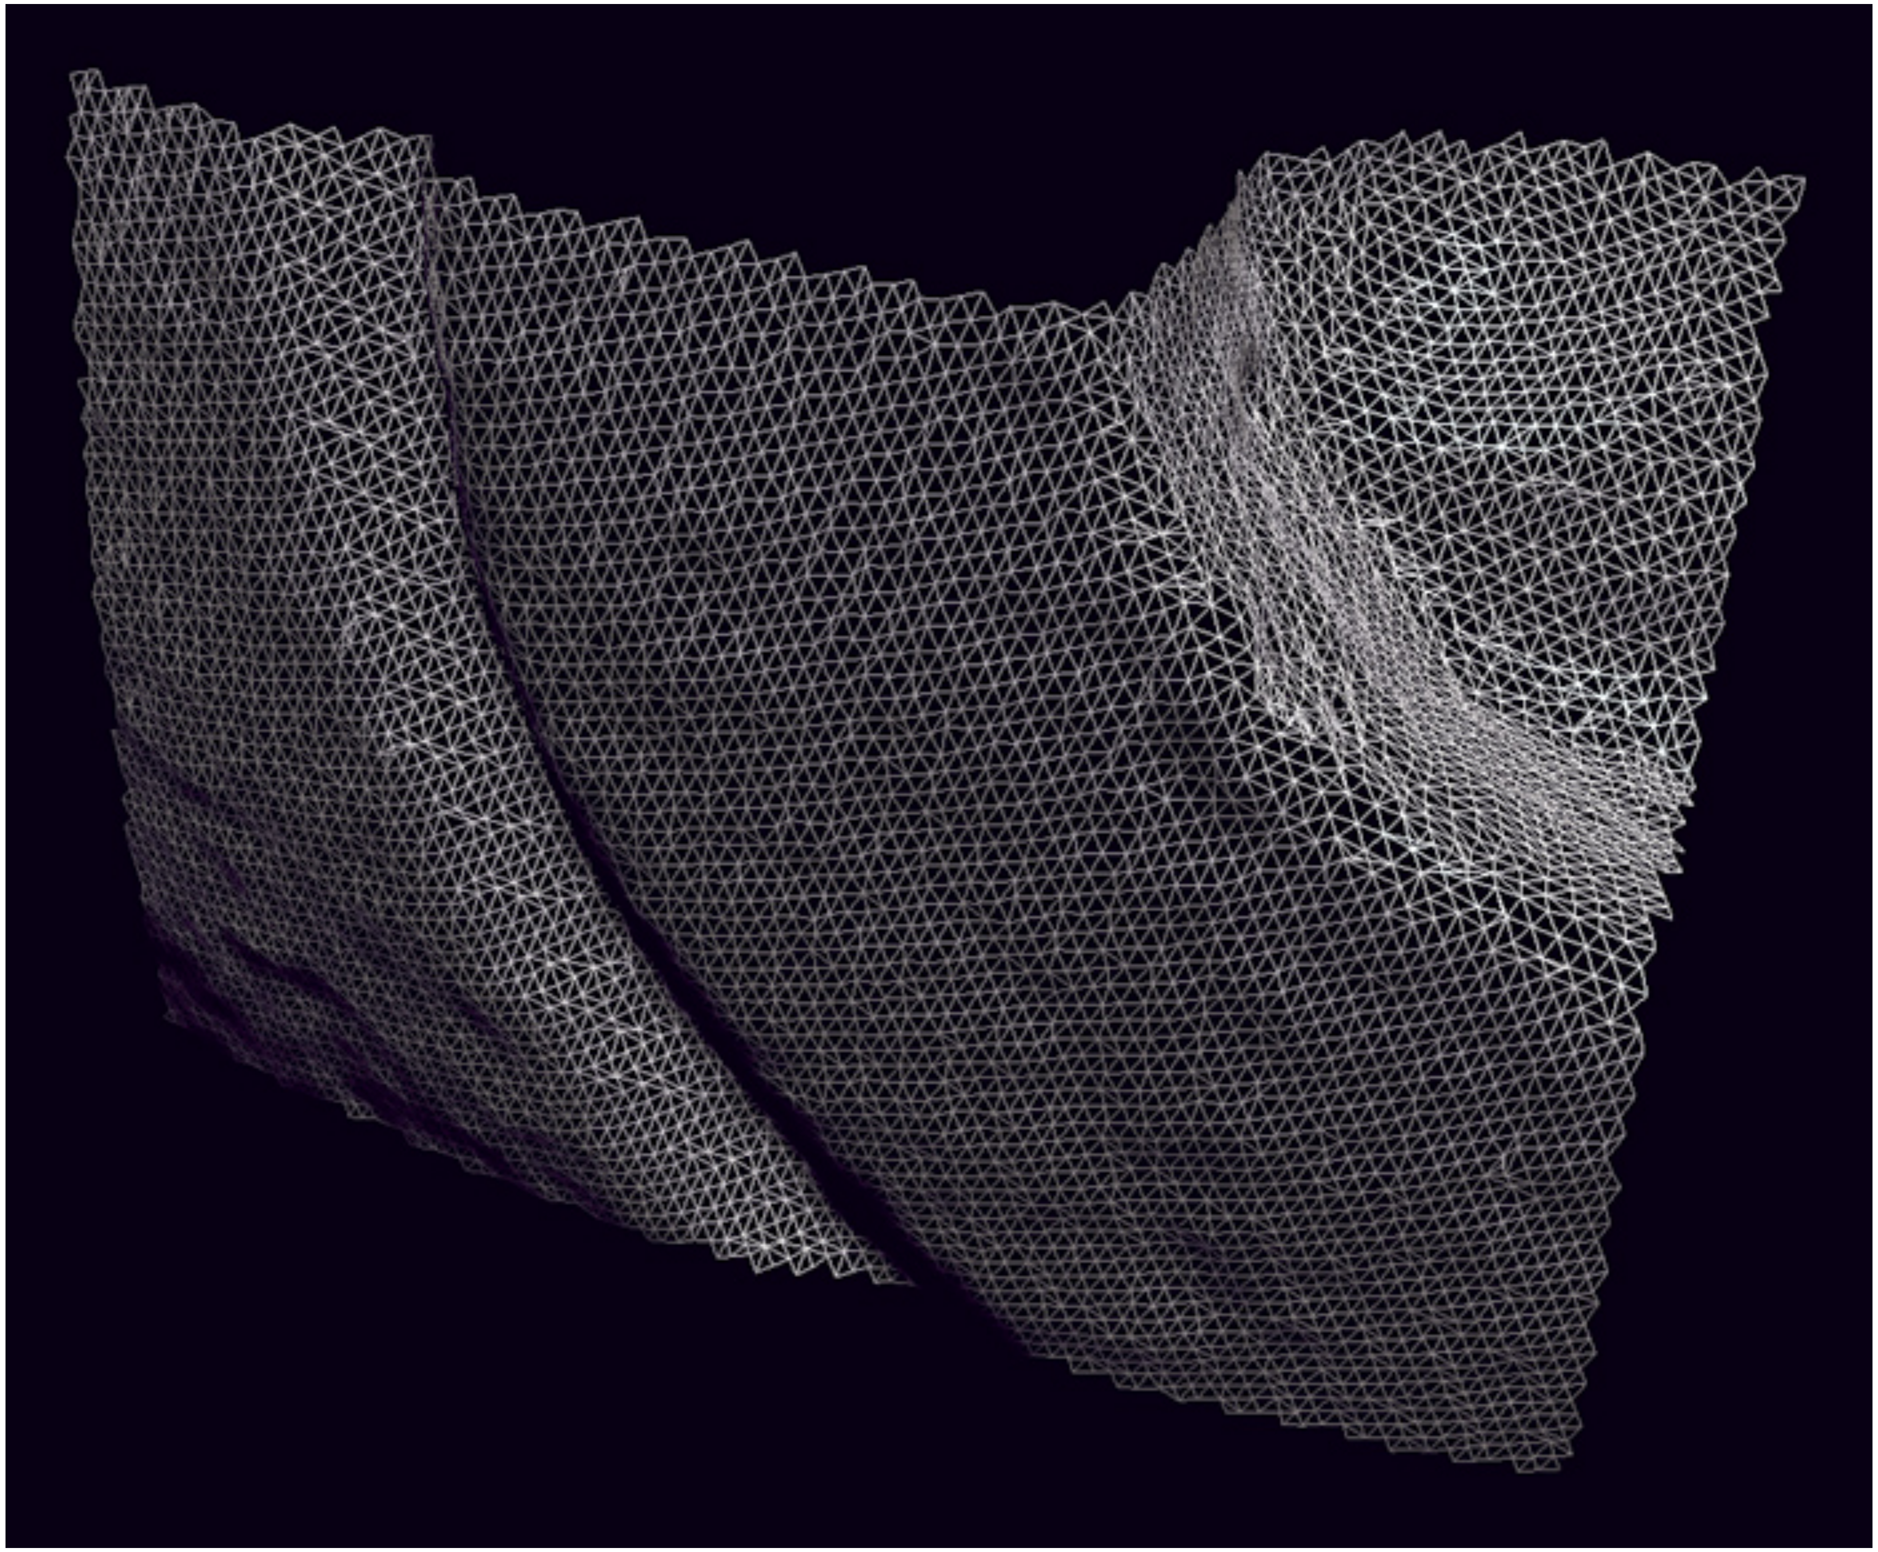

Digital Michelangelo Project:
http://graphics.stanford.edu/projects/mich/

How does the Kinect Gesture Tracking Work

How does the Kinect Gesture Tracking Work

- Structured light
    - Technique #3

Real Time by Color Coding

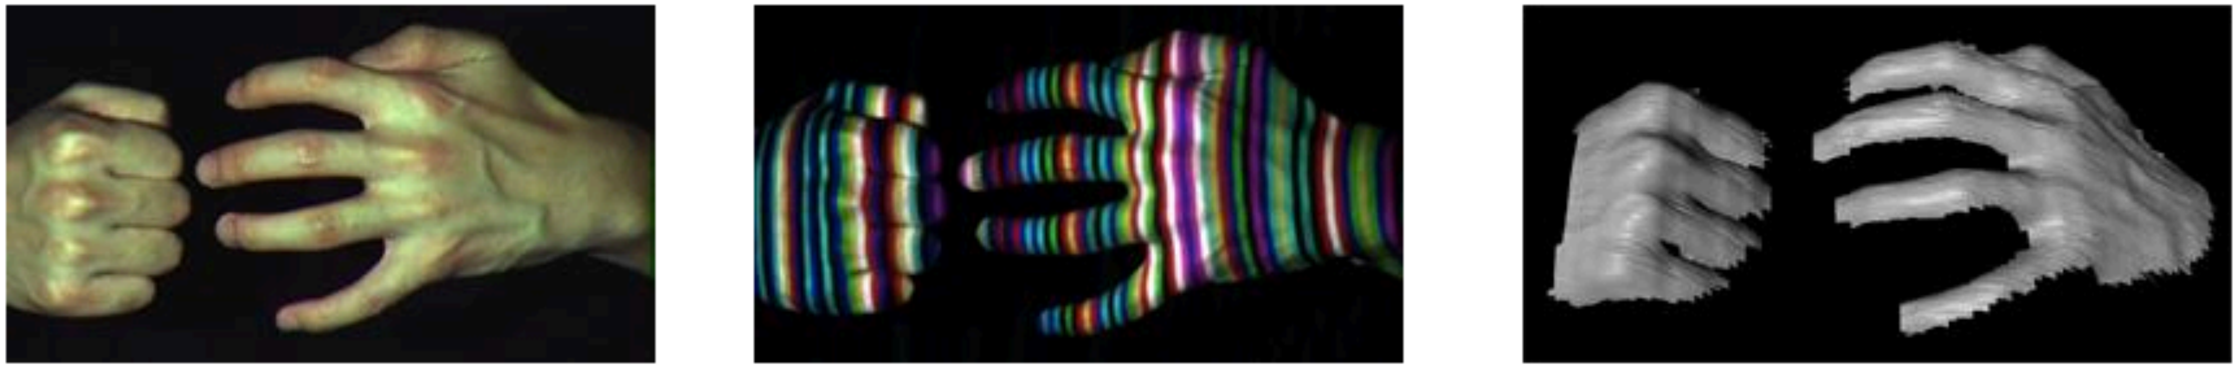

Works despite complex appearances

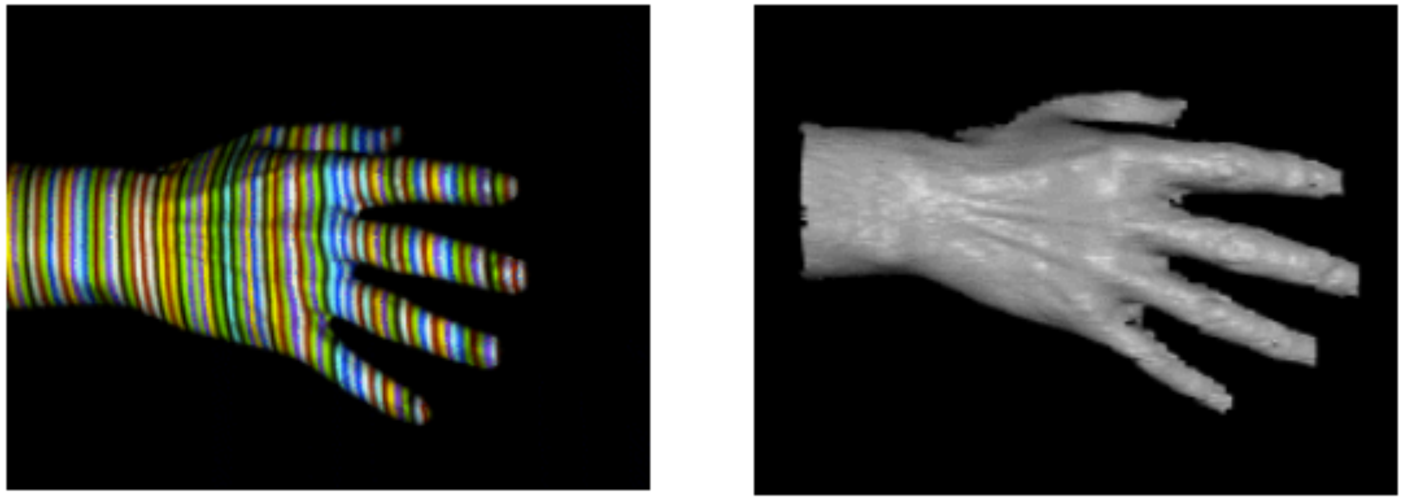

Works in real-time and on dynamic scenes

- Need very few images (one or two).
- But needs a more complex correspondence algorithm

Zhang et al, 3DPVT 2002

http://www.sci.utah.edu/~gerig/CS6320-S2012/Materials/CS6320-CV-S2012-StructuredLight.pdf

How does the Kinect Gesture Tracking Work

How does the Kinect Gesture Tracking Work

- Structured light
  - Do it yourself

How does the Kinect Gesture Tracking Work

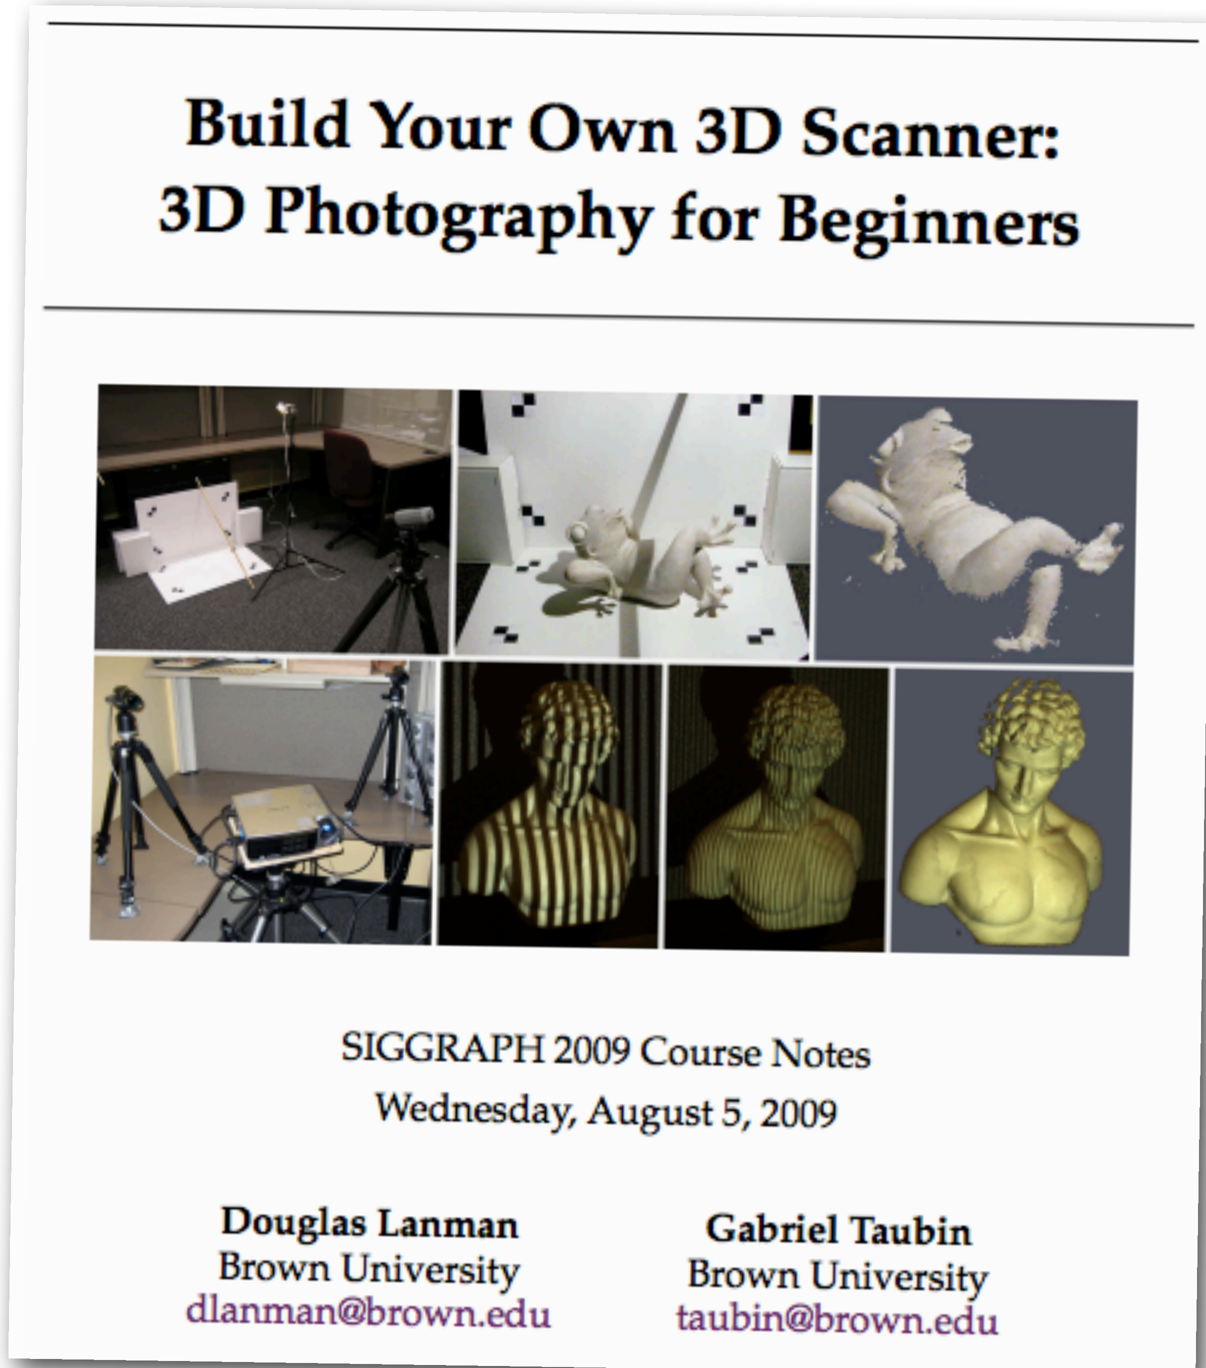

- Course notes: http://mesh.brown.edu/byo3d/notes/byo3D.pdf
- Slides: http://mesh.brown.edu/byo3d/slides.html
- Source code: http://mesh.brown.edu/byo3d/source.html

How does the Kinect Gesture Tracking Work

Dual Photography

Pradeep Sen* Billy Chen* Gaurav Garg* Stephen R. Marschner†
Mark Horowitz* Marc Levoy* Hendrik P.A. Lensch*

*Stanford University †Cornell University

SIGGRAPH2005

How does the Kinect Gesture Tracking Work

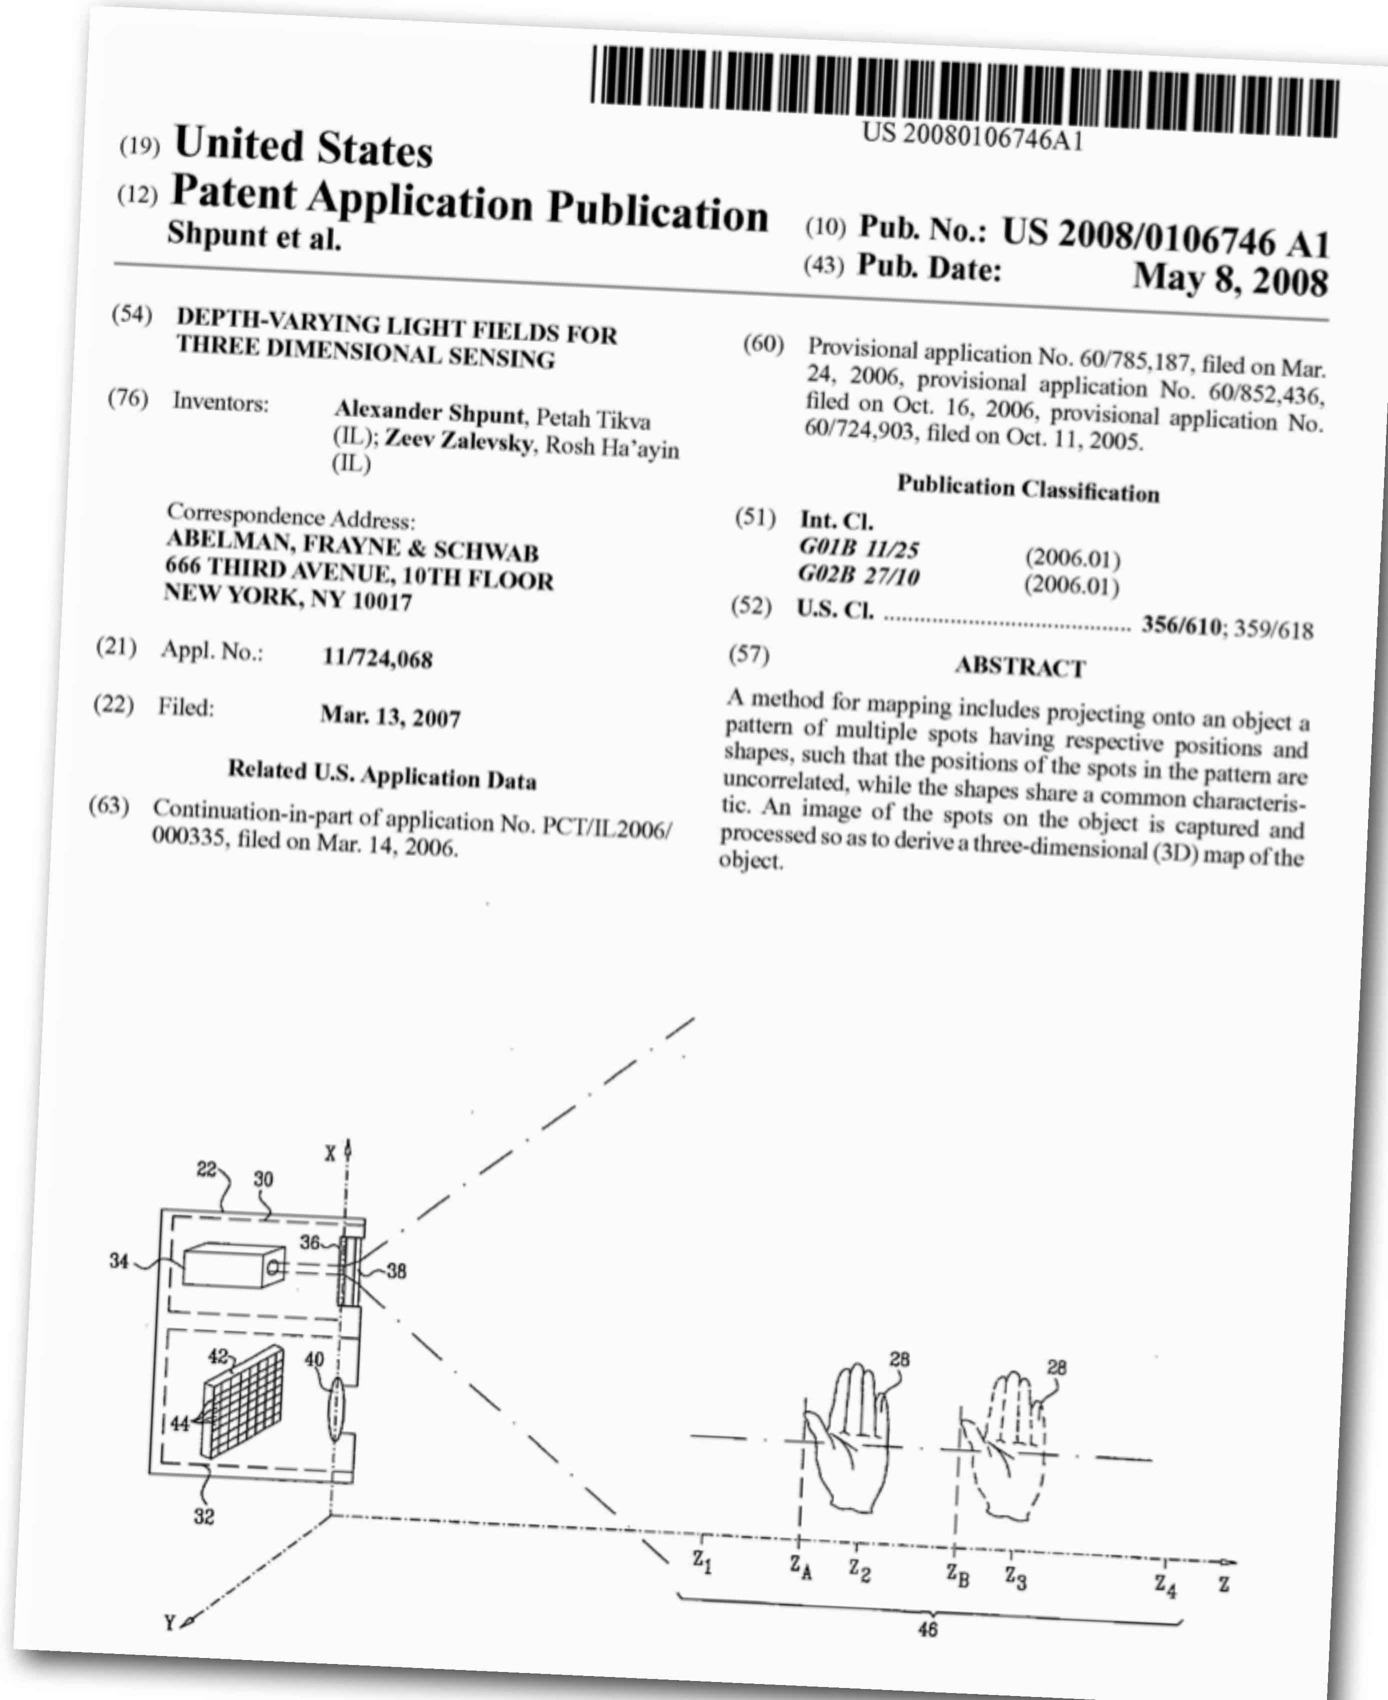

US 20080106746A1

(19) **United States**
(12) **Patent Application Publication**
Shpunt et al.

(10) **Pub. No.: US 2008/0106746 A1**
(43) **Pub. Date: May 8, 2008**

(54) **DEPTH-VARYING LIGHT FIELDS FOR THREE DIMENSIONAL SENSING**

(76) Inventors: **Alexander Shpunt**, Petah Tikva (IL); **Zeev Zalevsky**, Rosh Ha'ayin (IL)

Correspondence Address:
ABELMAN, FRAYNE & SCHWAB
666 THIRD AVENUE, 10TH FLOOR
NEW YORK, NY 10017

(21) Appl. No.: **11/724,068**

(22) Filed: **Mar. 13, 2007**

Related U.S. Application Data

(63) Continuation-in-part of application No. PCT/IL2006/000335, filed on Mar. 14, 2006.

(60) Provisional application No. 60/785,187, filed on Mar. 24, 2006, provisional application No. 60/852,436, filed on Oct. 16, 2006, provisional application No. 60/724,903, filed on Oct. 11, 2005.

Publication Classification

(51) **Int. Cl.**
G01B 11/25 (2006.01)
G02B 27/10 (2006.01)

(52) **U.S. Cl.** **356/610**; 359/618

(57) **ABSTRACT**

A method for mapping includes projecting onto an object a pattern of multiple spots having respective positions and shapes, such that the positions of the spots in the pattern are uncorrelated, while the shapes share a common characteristic. An image of the spots on the object is captured and processed so as to derive a three-dimensional (3D) map of the object.

How does the Kinect Gesture Tracking Work

- Structured light
  - Kinect Style

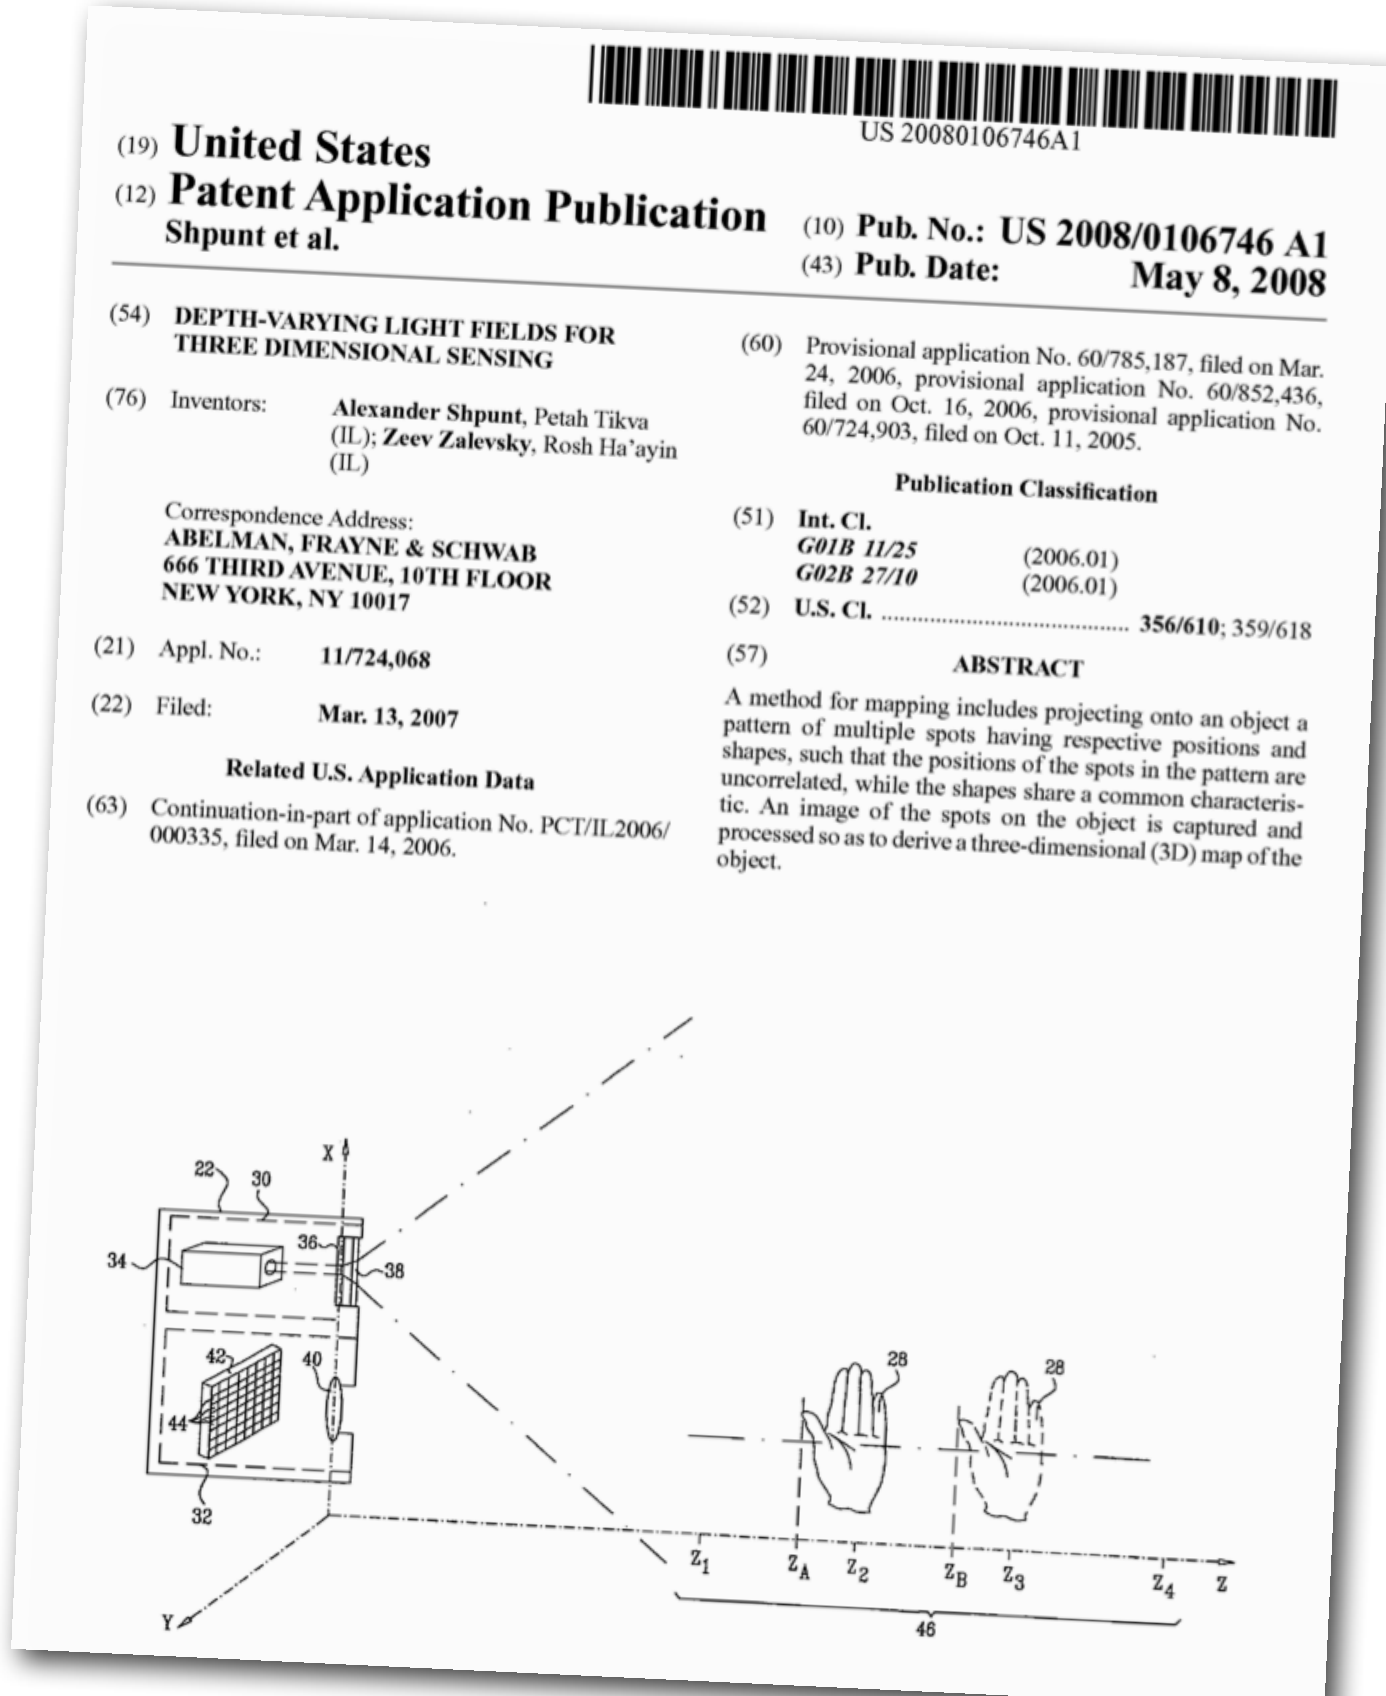

US 20080106746A1

(19) **United States**
(12) **Patent Application Publication**
Shpunt et al.

(10) **Pub. No.: US 2008/0106746 A1**
(43) **Pub. Date: May 8, 2008**

(54) **DEPTH-VARYING LIGHT FIELDS FOR THREE DIMENSIONAL SENSING**

(76) Inventors: **Alexander Shpunt**, Petah Tikva (IL); **Zeev Zalevsky**, Rosh Ha'ayin (IL)

Correspondence Address:
ABELMAN, FRAYNE & SCHWAB
666 THIRD AVENUE, 10TH FLOOR
NEW YORK, NY 10017

(21) Appl. No.: **11/724,068**

(22) Filed: **Mar. 13, 2007**

Related U.S. Application Data

(63) Continuation-in-part of application No. PCT/IL2006/000335, filed on Mar. 14, 2006.

(60) Provisional application No. 60/785,187, filed on Mar. 24, 2006, provisional application No. 60/852,436, filed on Oct. 16, 2006, provisional application No. 60/724,903, filed on Oct. 11, 2005.

Publication Classification

(51) **Int. Cl.**
G01B 11/25 (2006.01)
G02B 27/10 (2006.01)

(52) **U.S. Cl.** **356/610**; 359/618

(57) **ABSTRACT**

A method for mapping includes projecting onto an object a pattern of multiple spots having respective positions and shapes, such that the positions of the spots in the pattern are uncorrelated, while the shapes share a common characteristic. An image of the spots on the object is captured and processed so as to derive a three-dimensional (3D) map of the object.

How does the Kinect Gesture Tracking Work

FIG. 3A

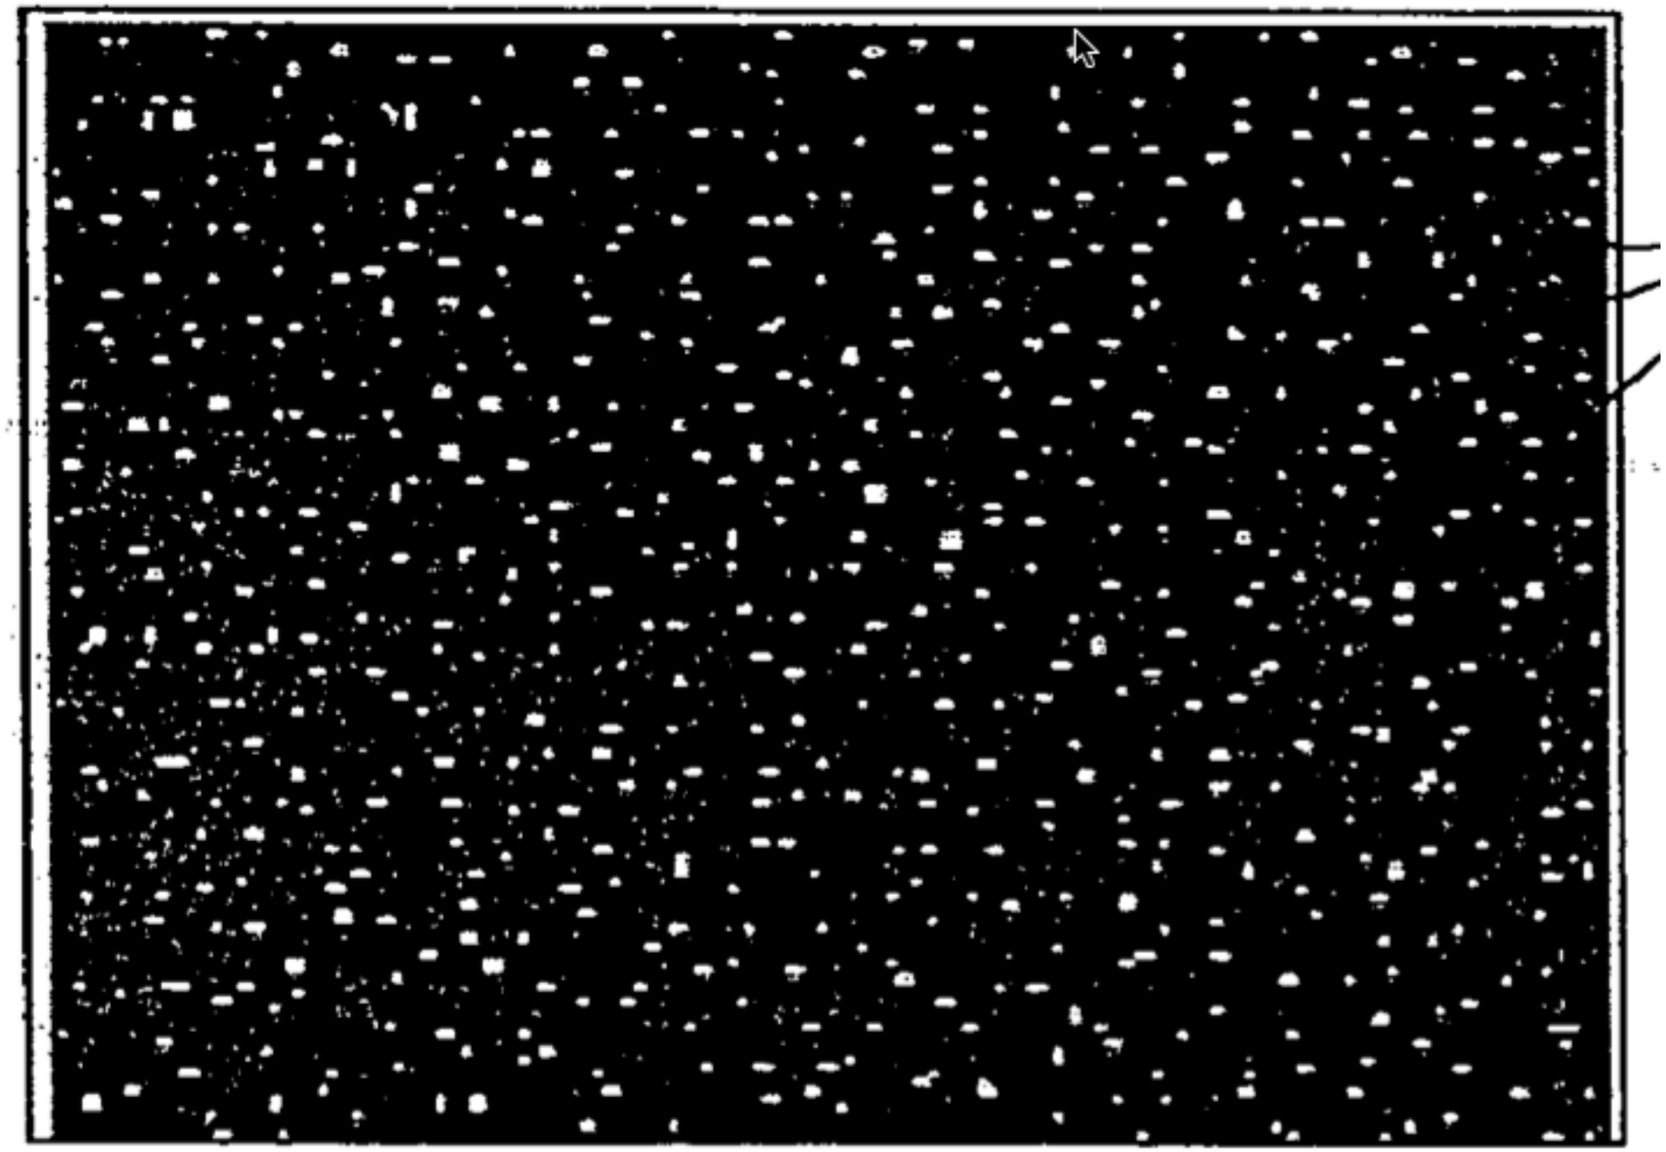

How does the Kinect Gesture Tracking Work

FIG. 3A

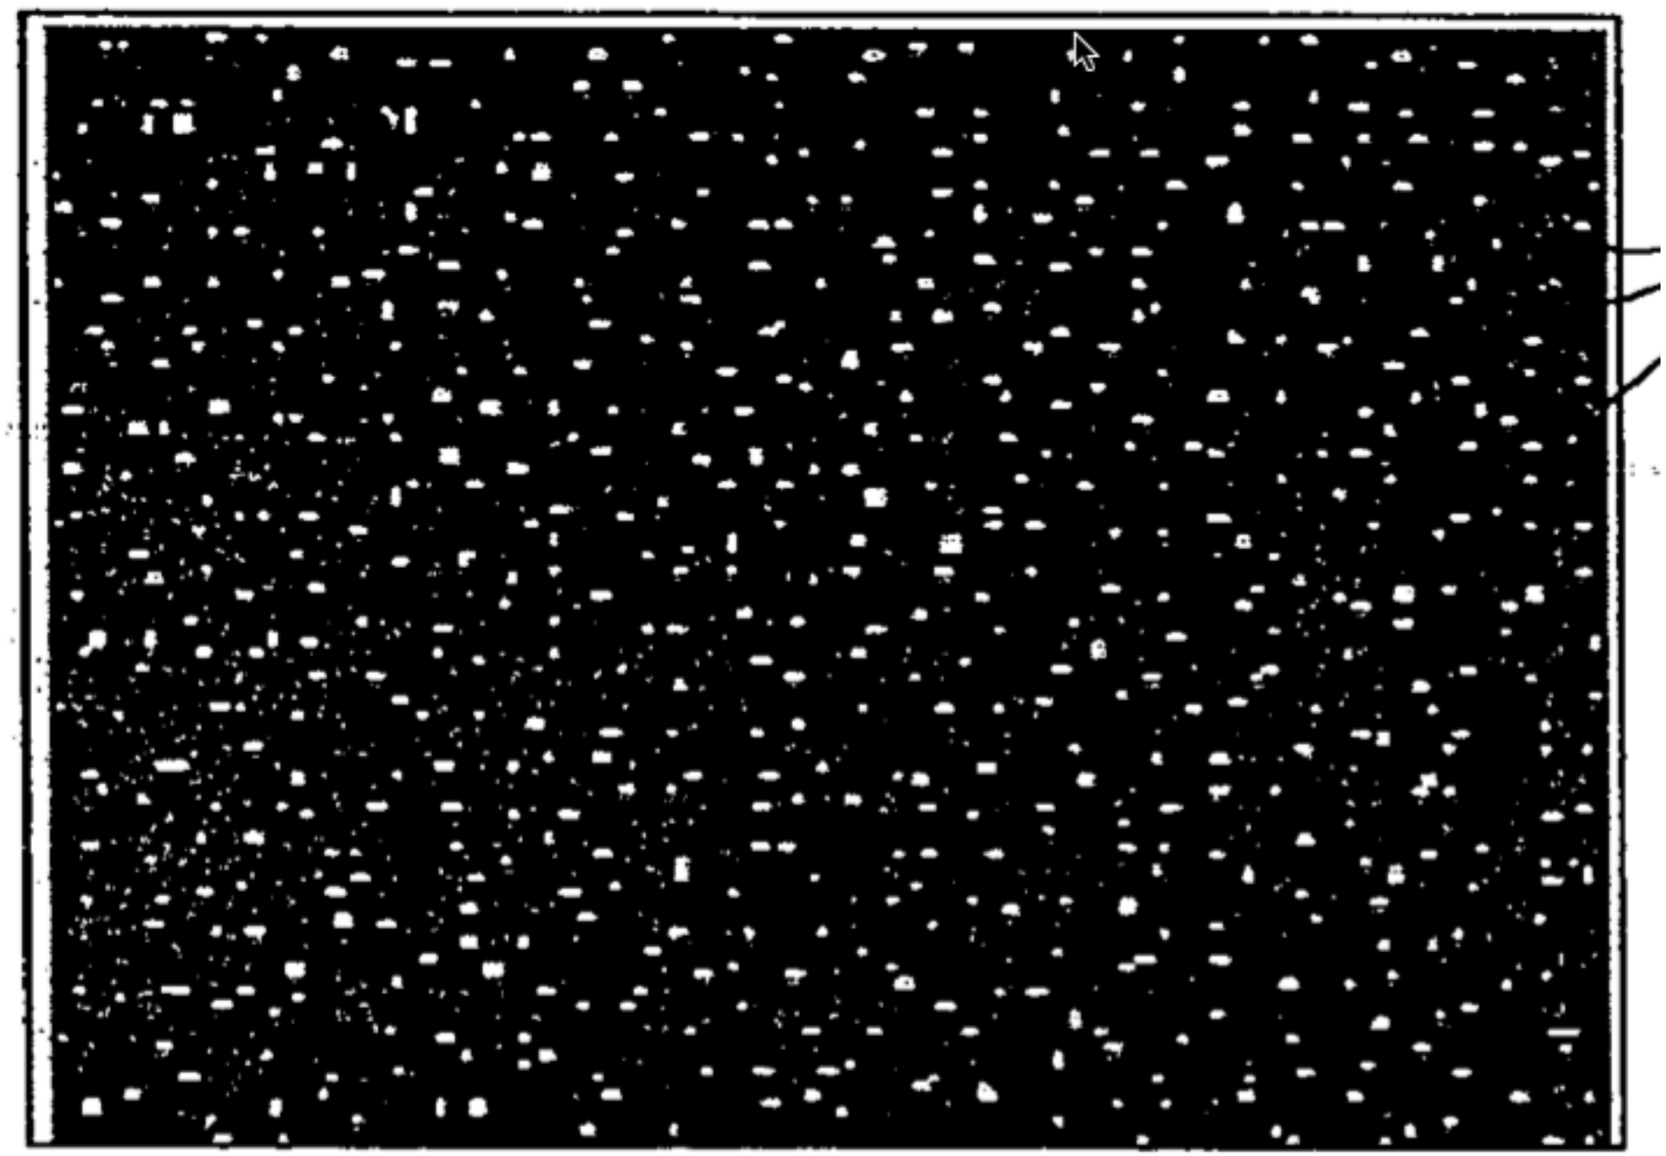

Patent: US 2008/0106746 A1

How does the Kinect Gesture Tracking Work

- Structured light

FIG. 3A

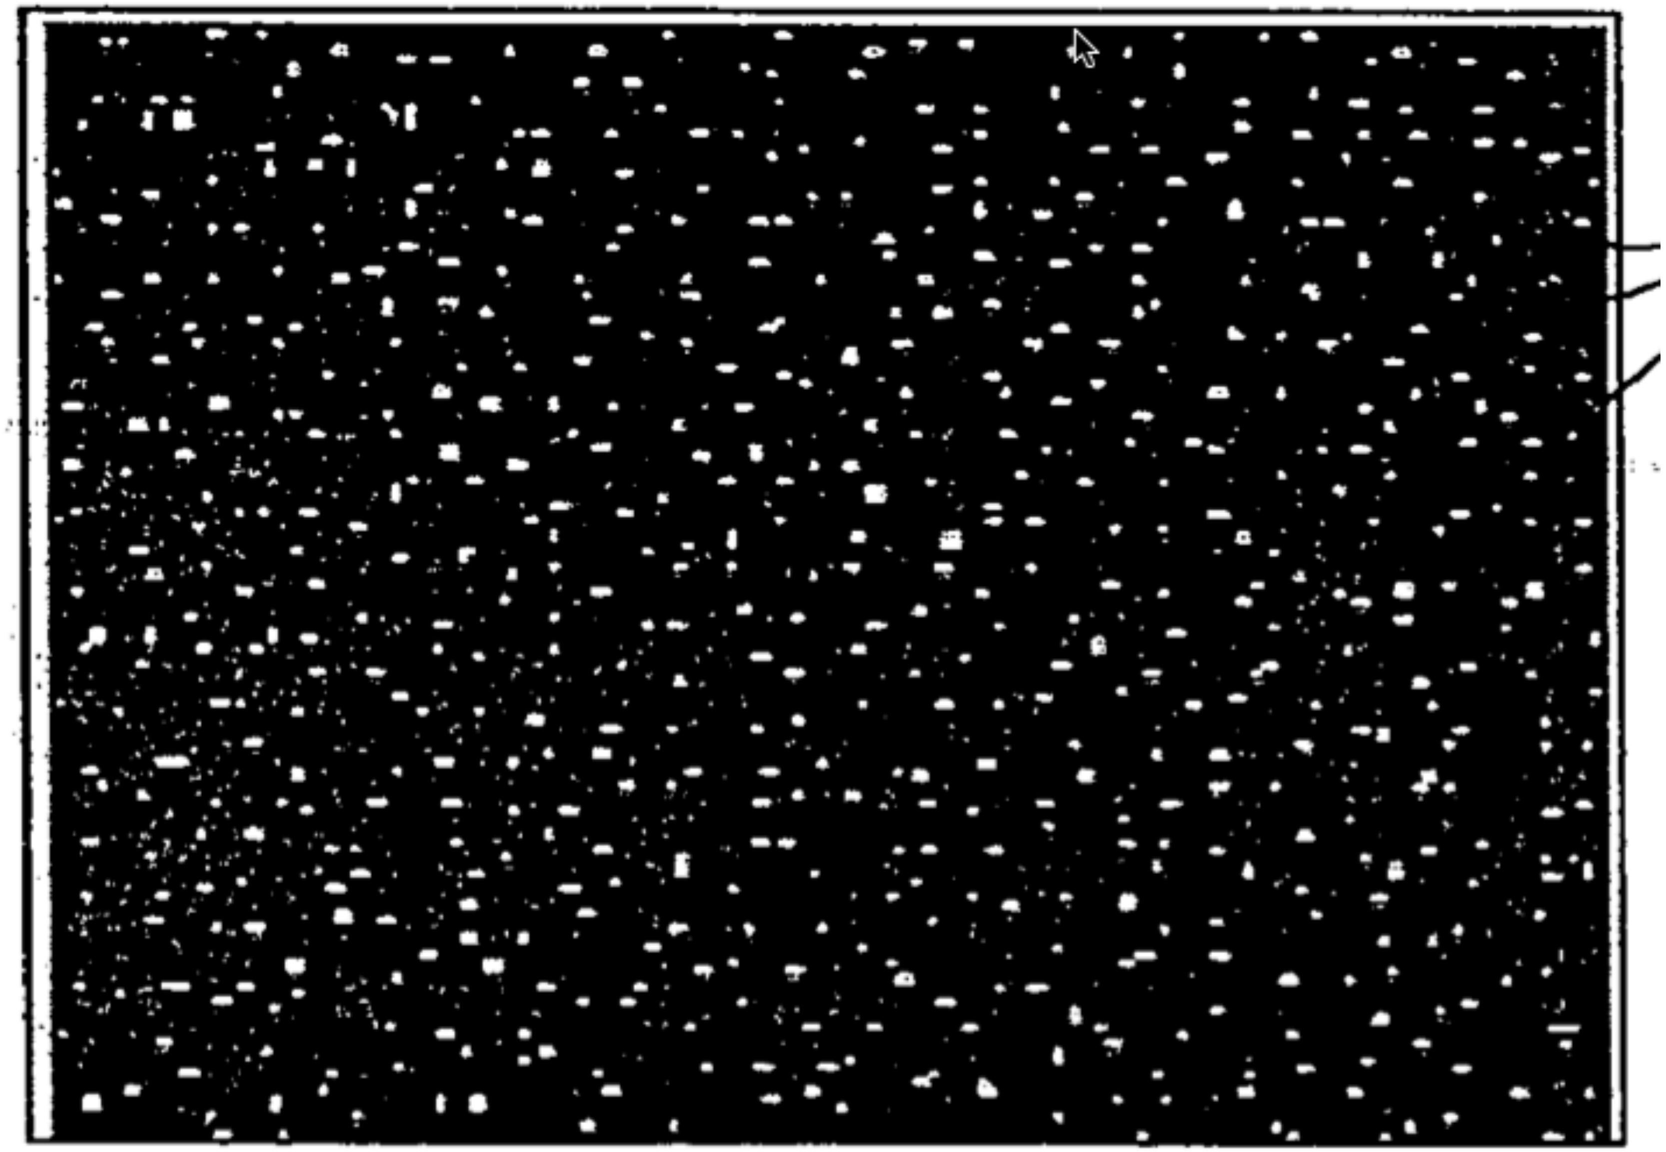

Patent: US 2008/0106746 A1

How does the Kinect Gesture Tracking Work

- Structured light
  - The kinect uses random spots created by an infrared laser

FIG. 3A

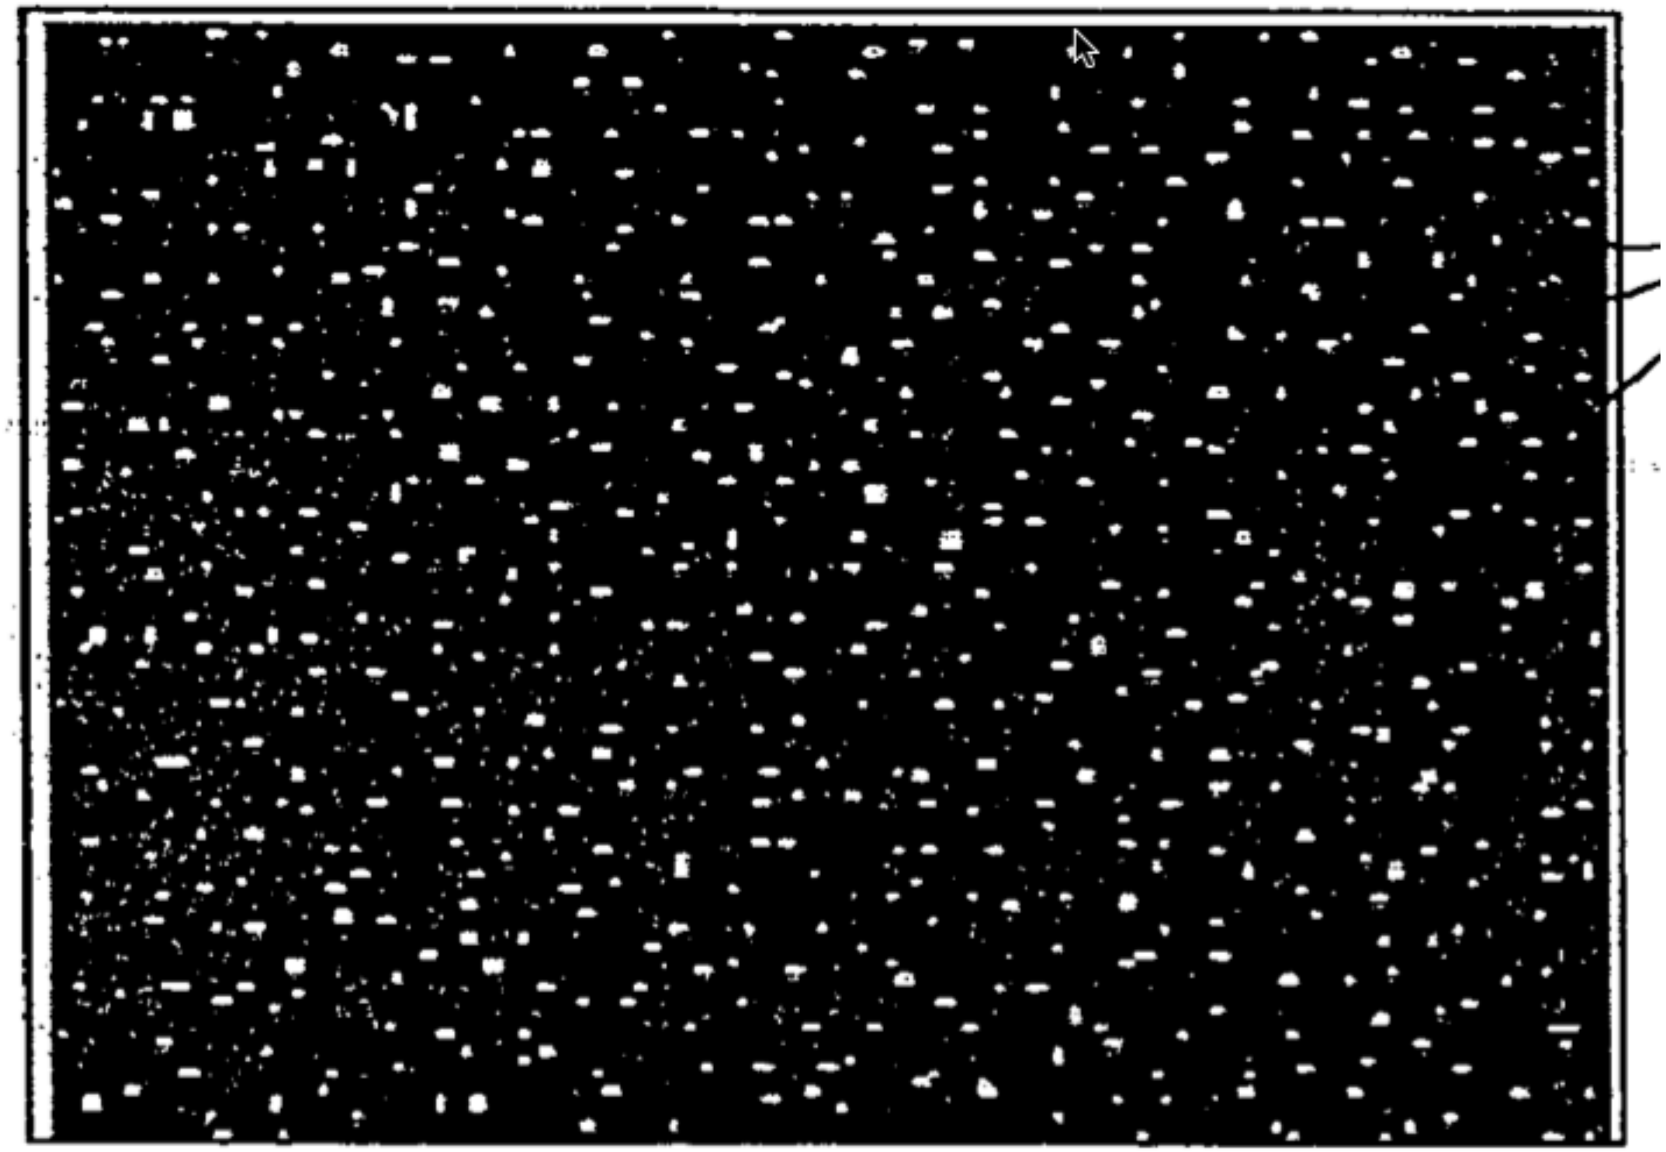

Patent: US 2008/0106746 A1

How does the Kinect Gesture Tracking Work

FIG. 3A

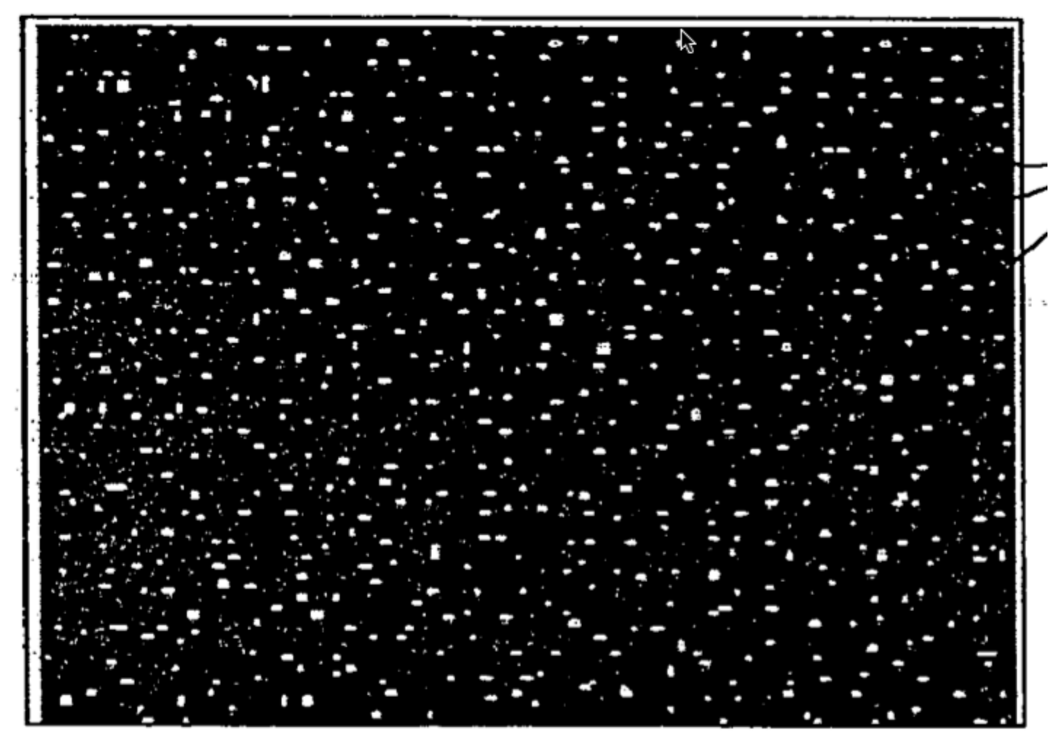

How does the Kinect Gesture Tracking Work

FIG. 3A

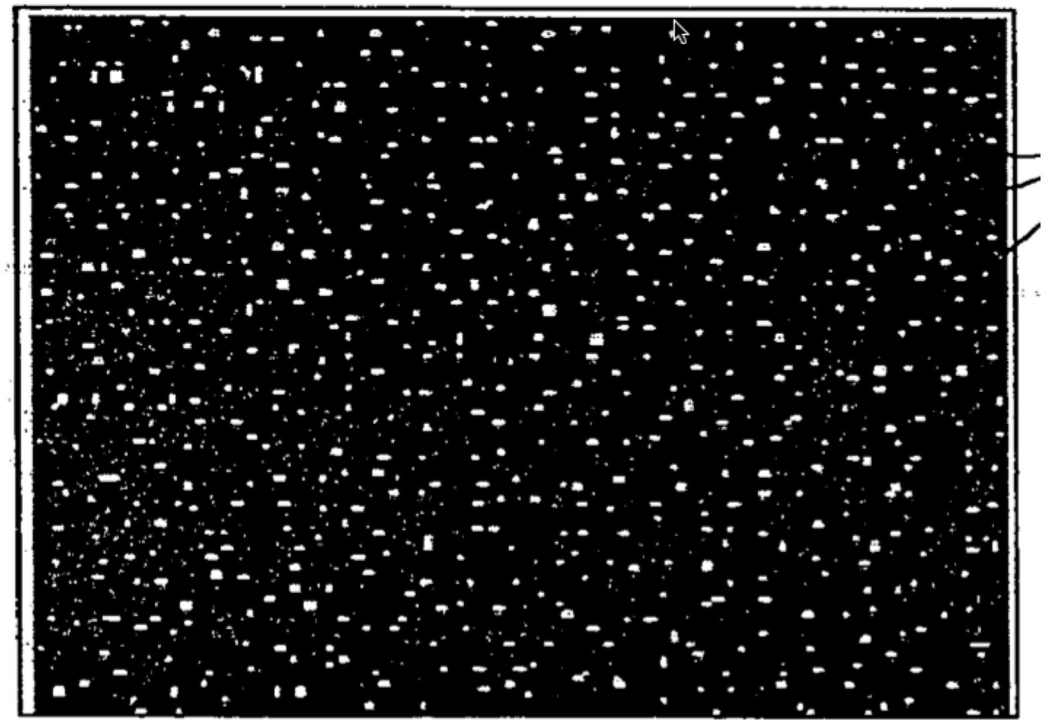

Patent: US 2008/0106746 A1

How does the Kinect Gesture Tracking Work

- Structured light

FIG. 3A

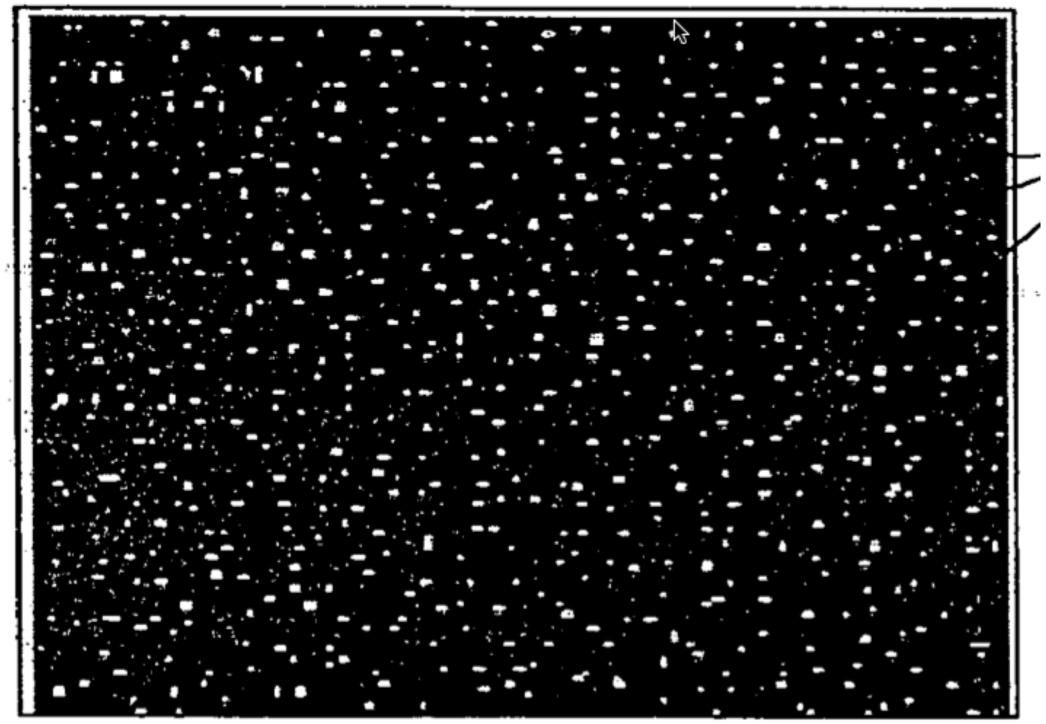

Patent: US 2008/0106746 A1

How does the Kinect Gesture Tracking Work

- Structured light
  - The kinect uses random spots created by an infrared laser

FIG. 3A

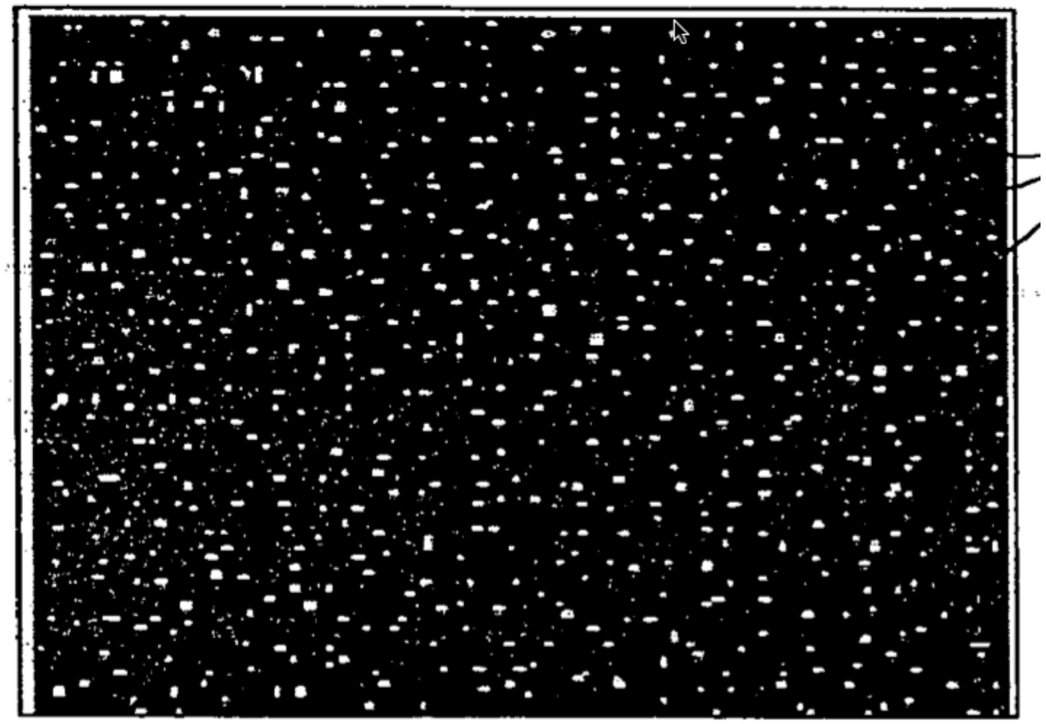

Patent: US 2008/0106746 A1

How does the Kinect Gesture Tracking Work

- Structured light
  - The kinect uses random spots created by an infrared laser
  - The RGB camera is not used for depth computation

FIG. 3A

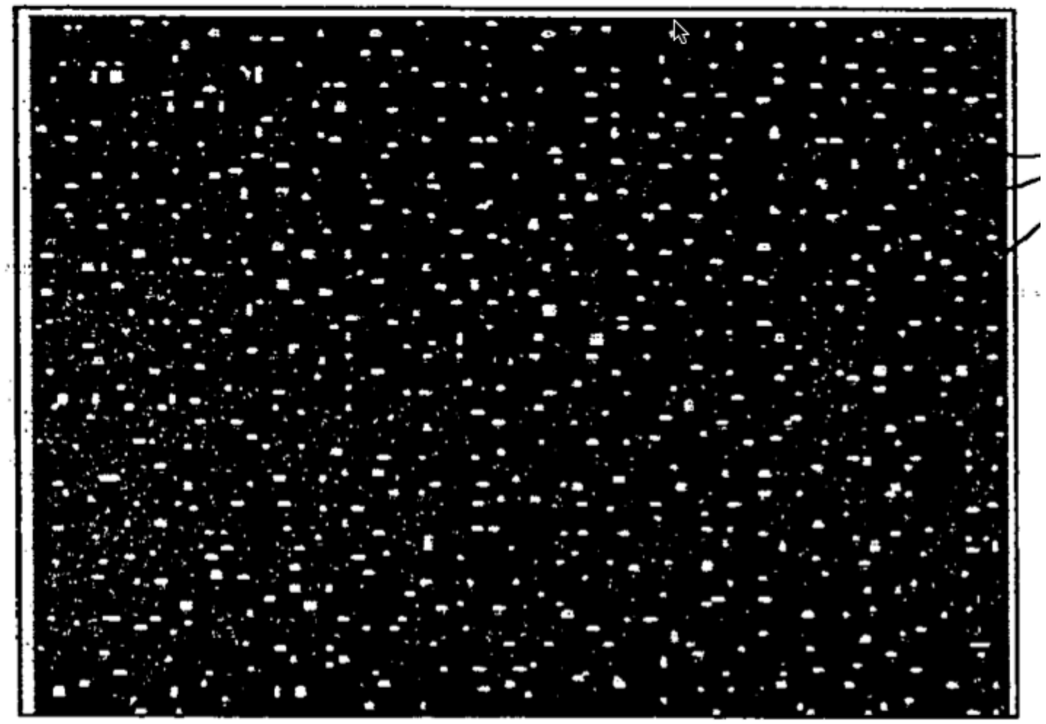

Patent: US 2008/0106746 A1

How does the Kinect Gesture Tracking Work

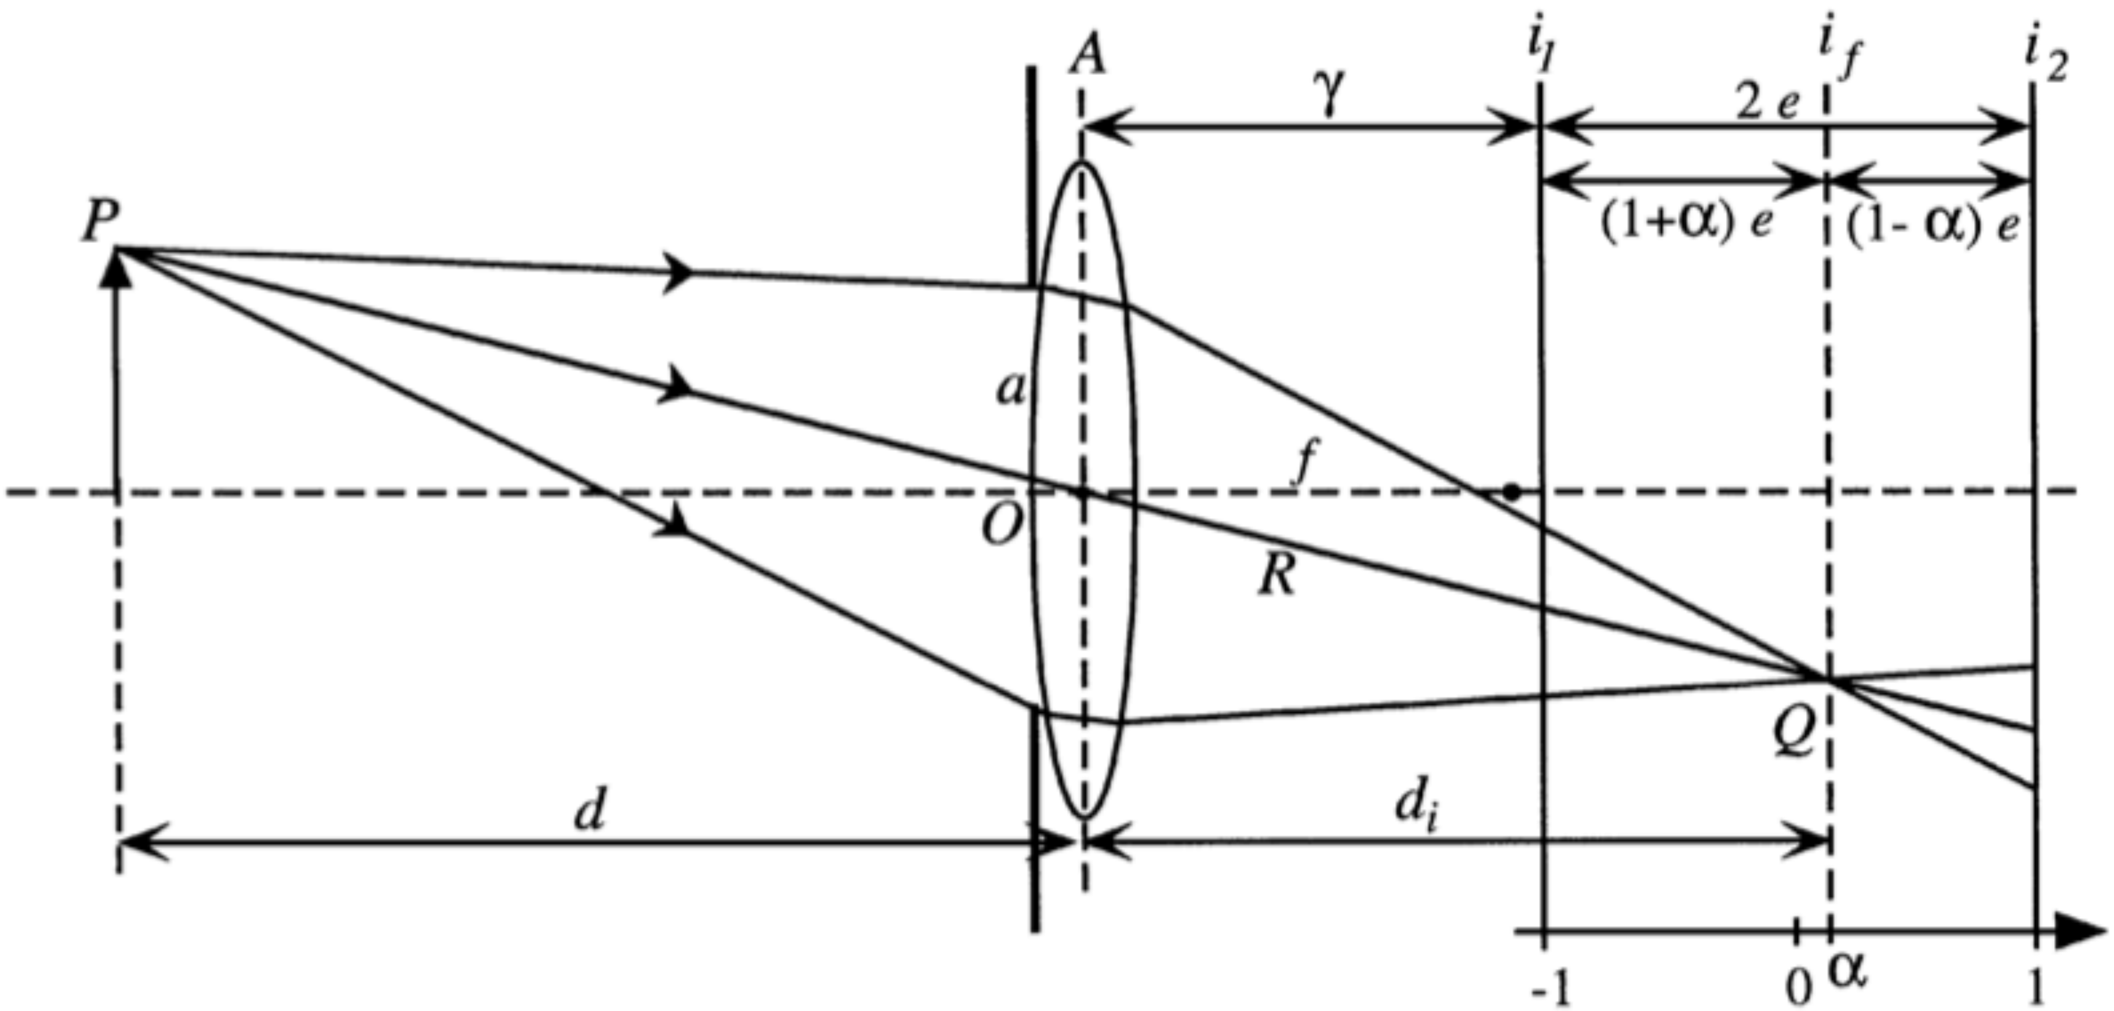

How does the Kinect Gesture Tracking Work

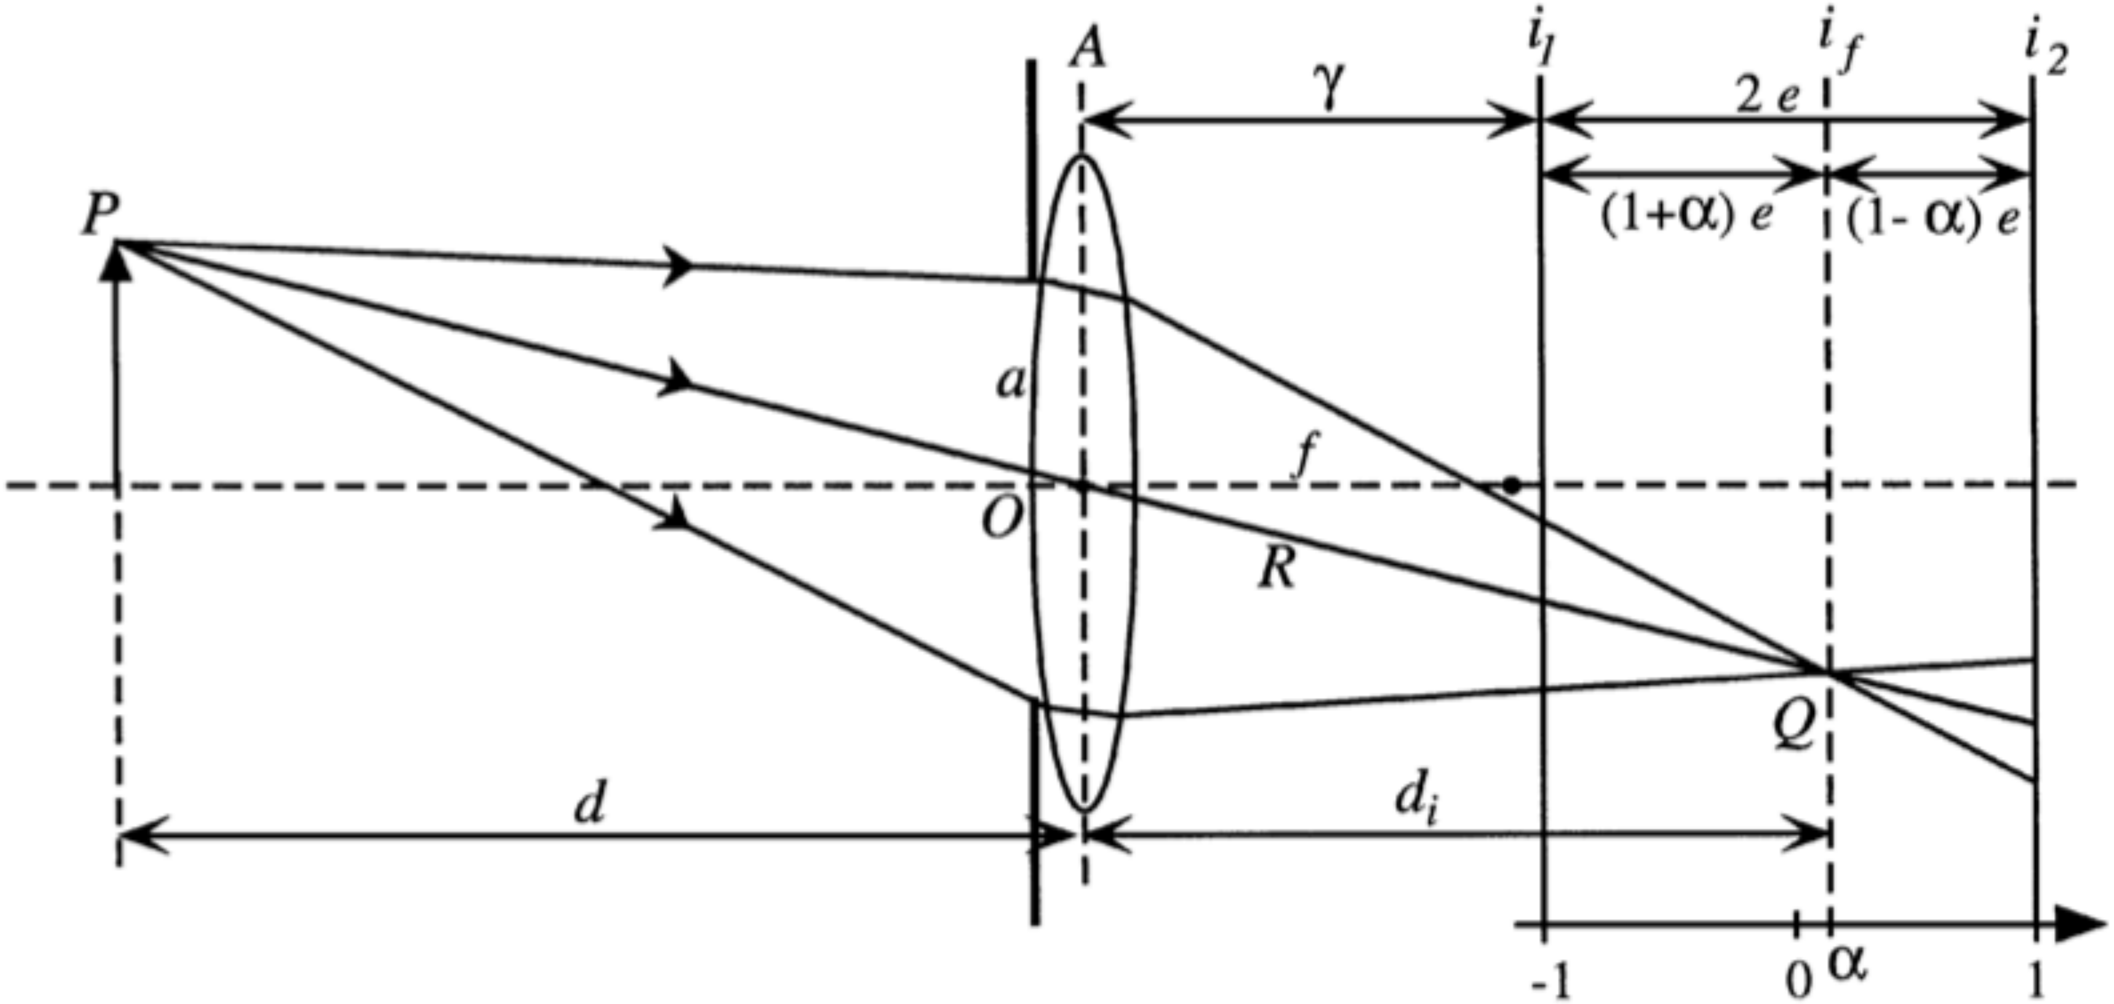

Rational Filters for Passive Depth from Defocus by Watanabe,Nayar

How does the Kinect Gesture Tracking Work

- Depth from focus

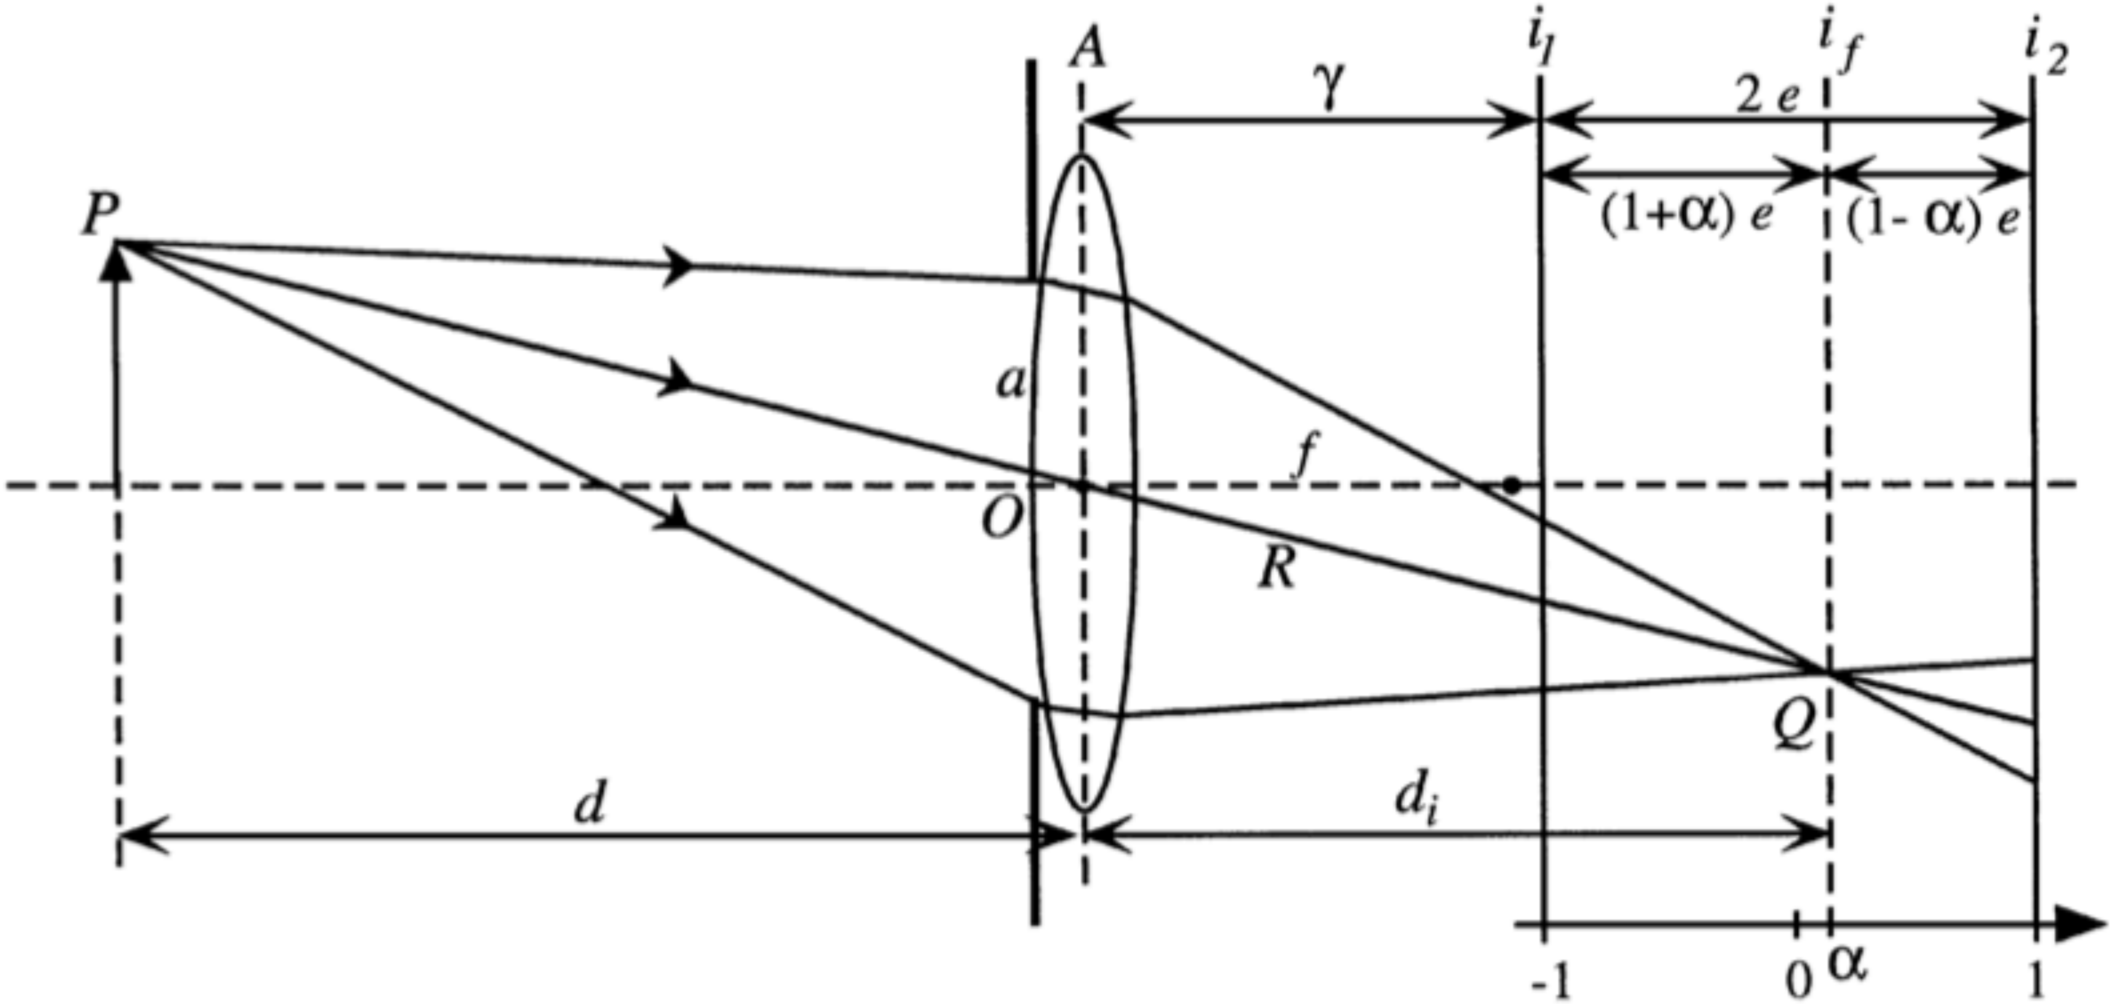

Rational Filters for Passive Depth from Defocus by Watanabe,Nayar

How does the Kinect Gesture Tracking Work

- Depth from focus
  - Blurry things are further away

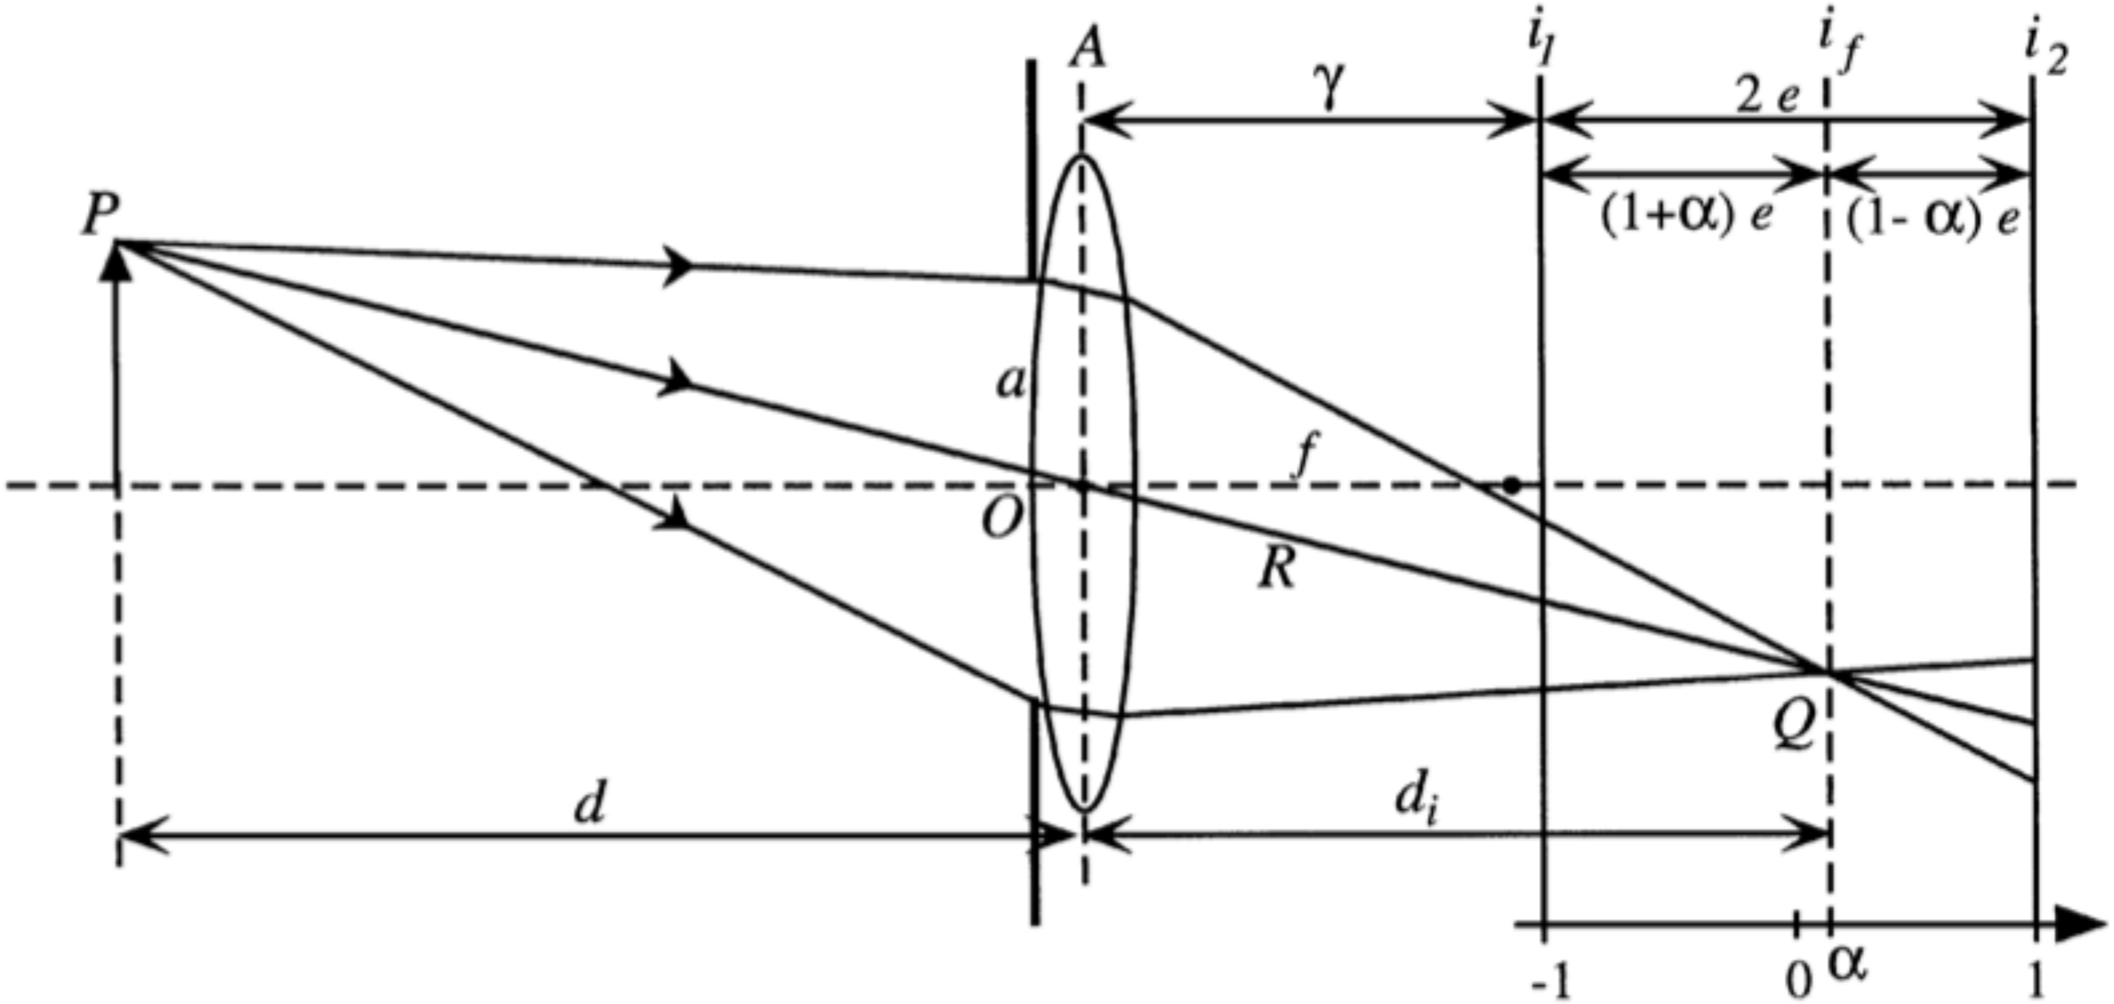

Rational Filters for Passive Depth from Defocus by Watanabe,Nayar

(a) far-focused image

(b) near-focused image

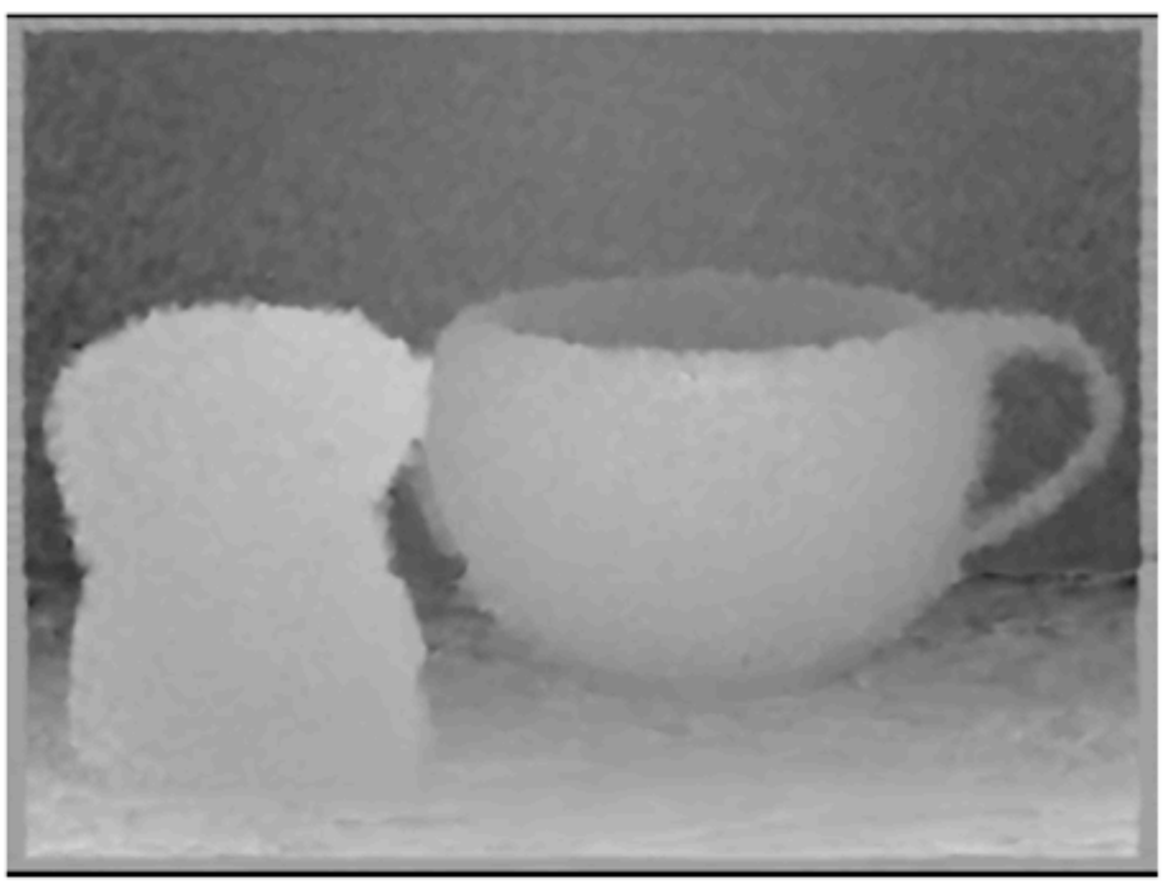

(c) gray-coded depth map without post-filtering

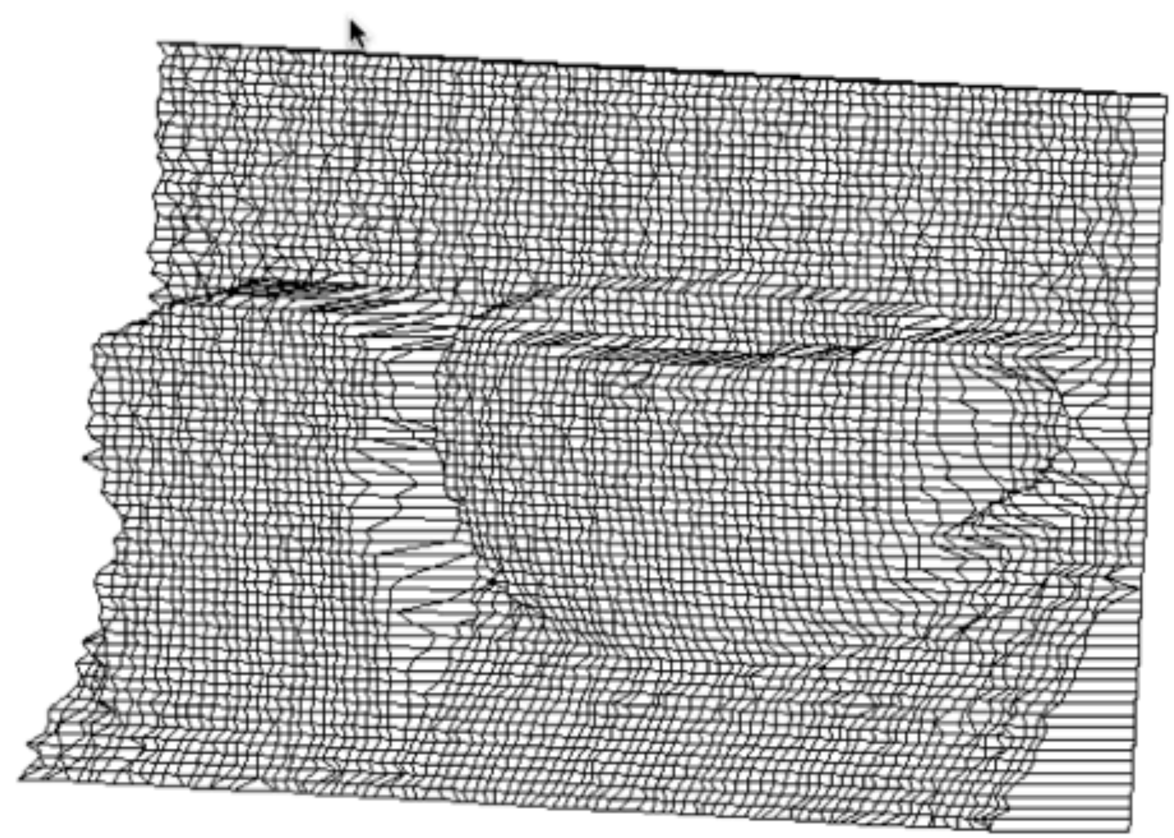

(d) wireframe plot of (c)

(a) far-focused image

(b) near-focused image

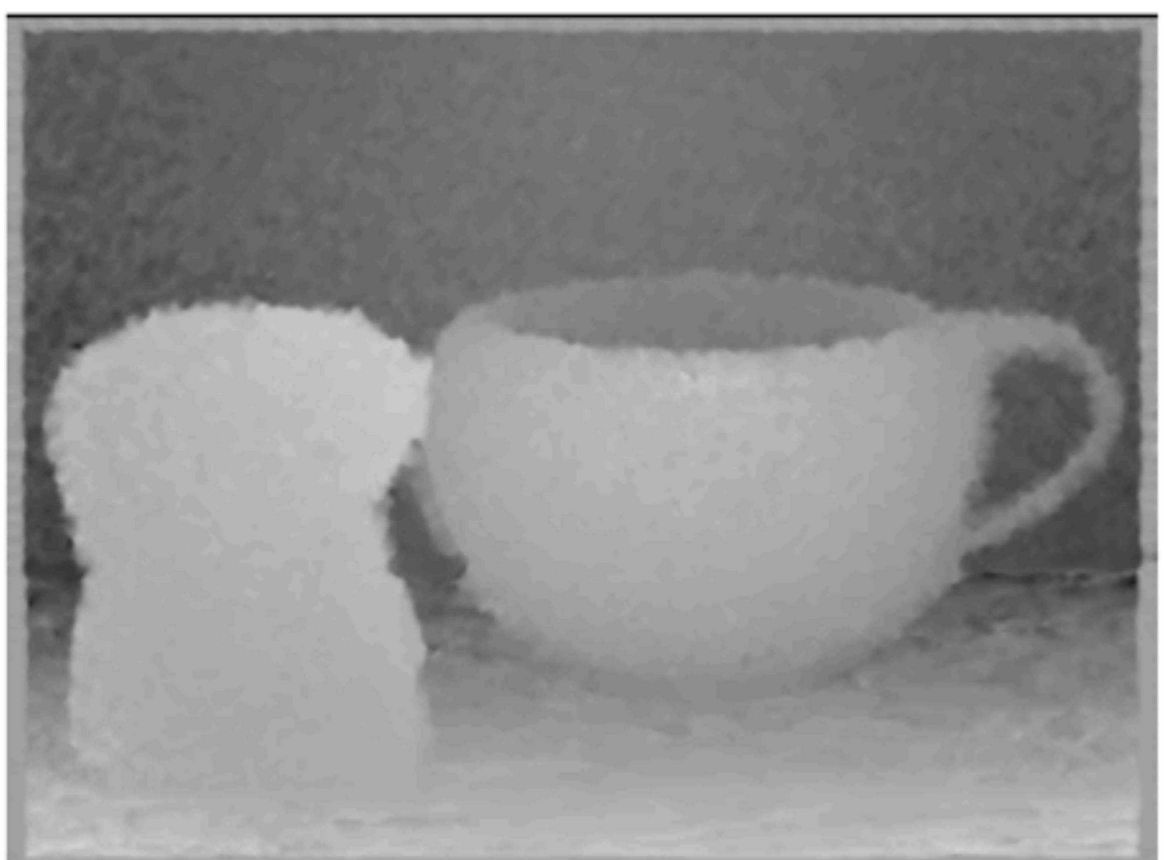

(c) gray-coded depth map without post-filtering

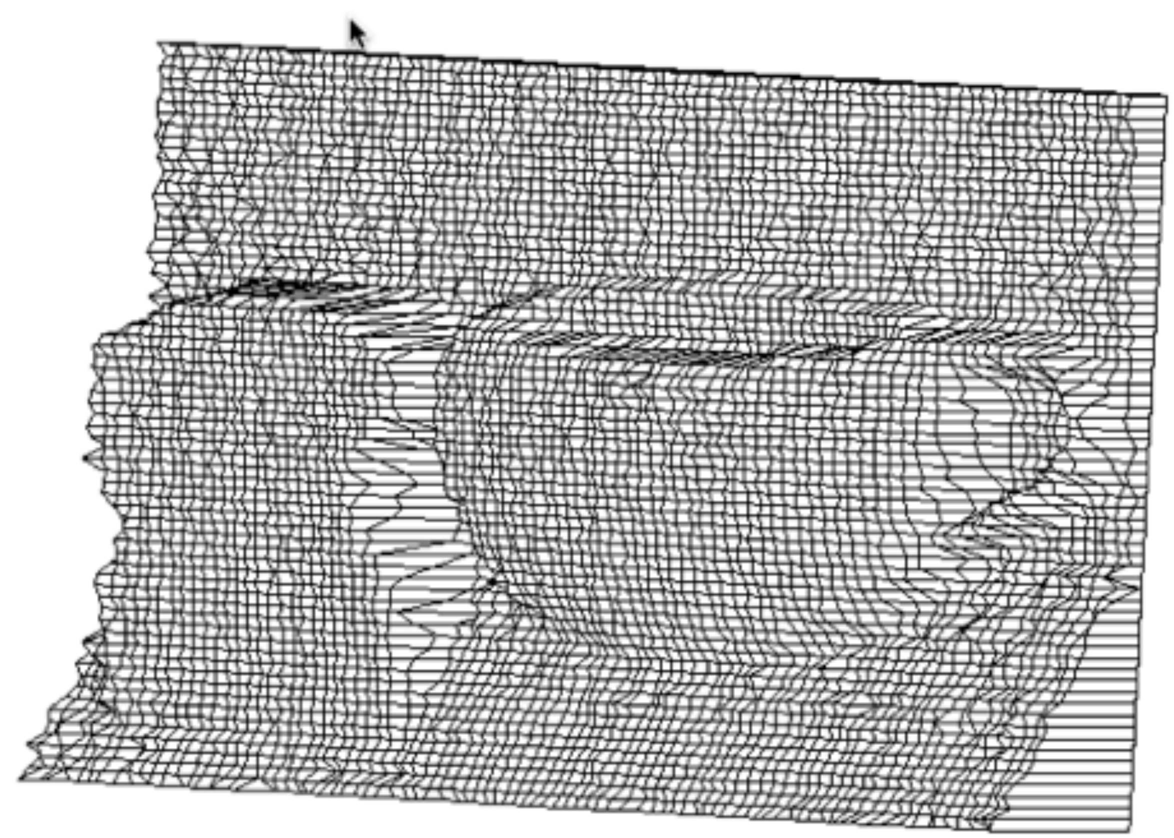

(d) wireframe plot of (c)

(a) far-focused image

(b) near-focused image

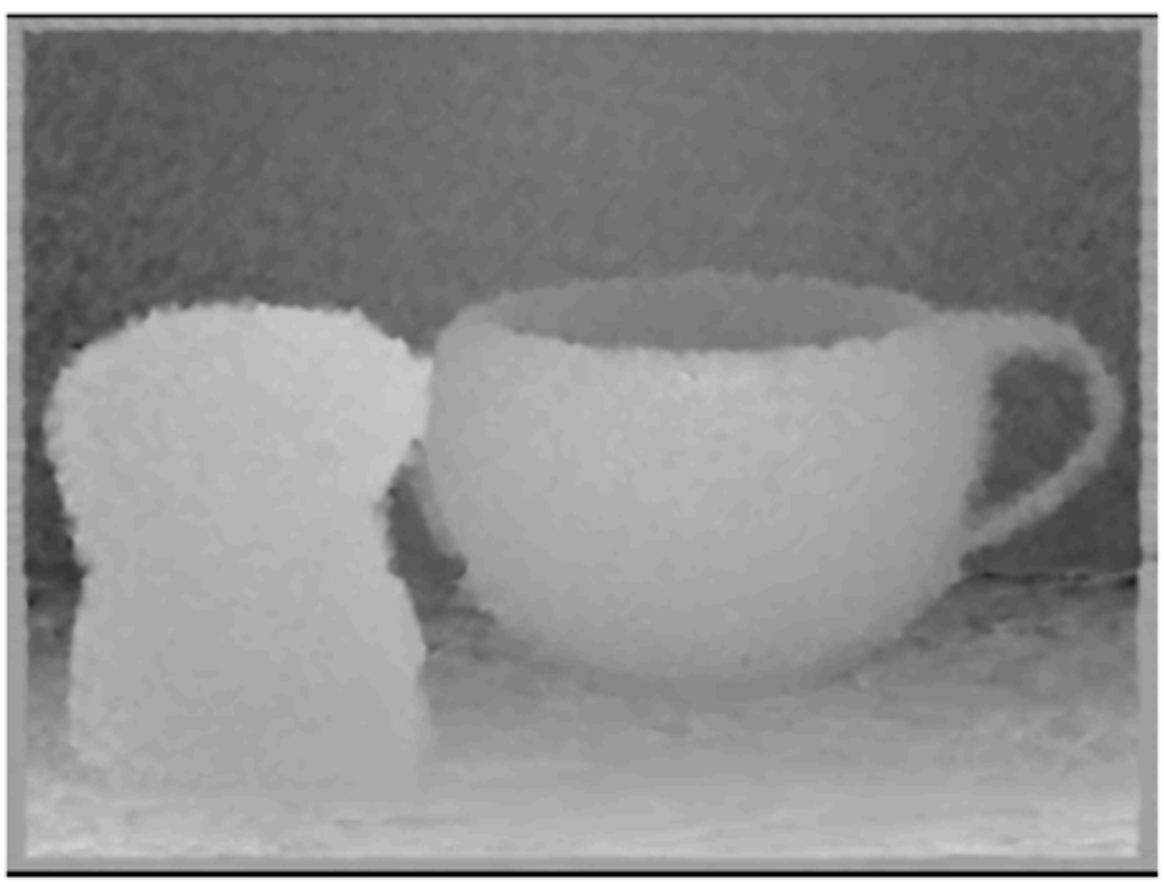

(c) gray-coded depth map without post-filtering

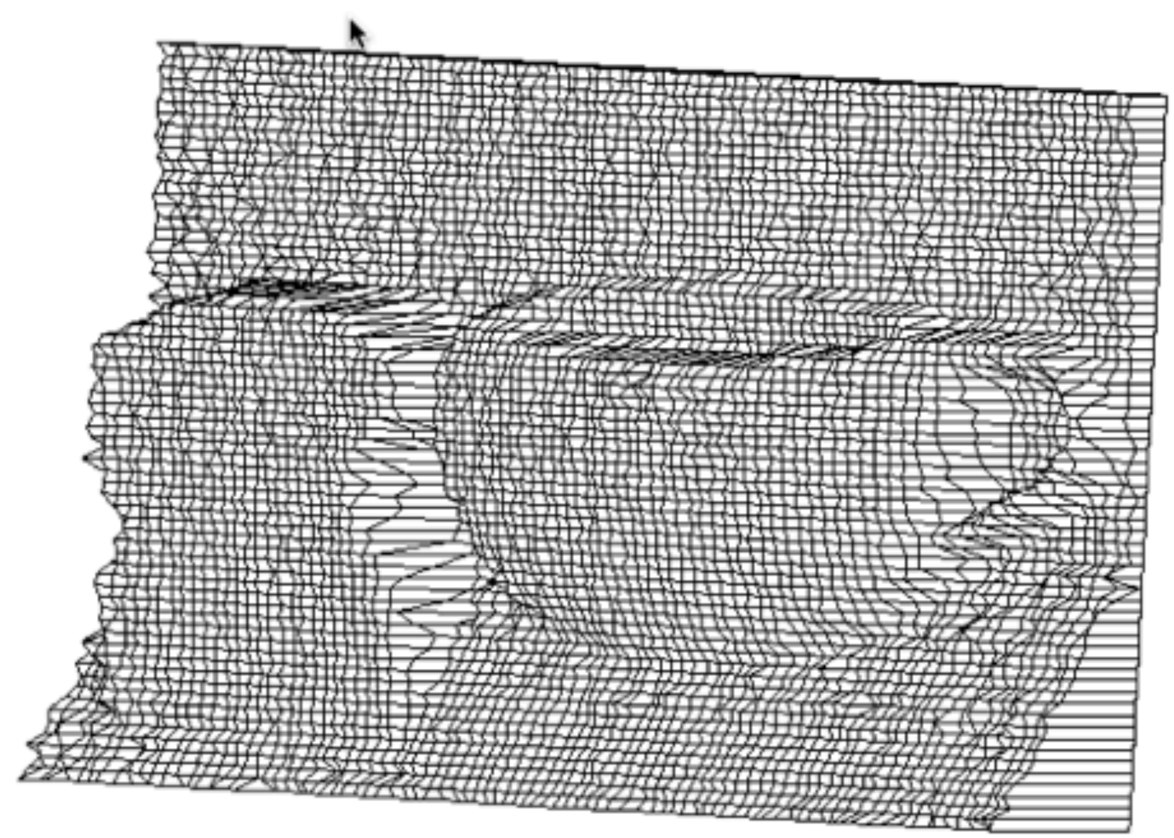

(d) wireframe plot of (c)

(a) far-focused image

(b) near-focused image

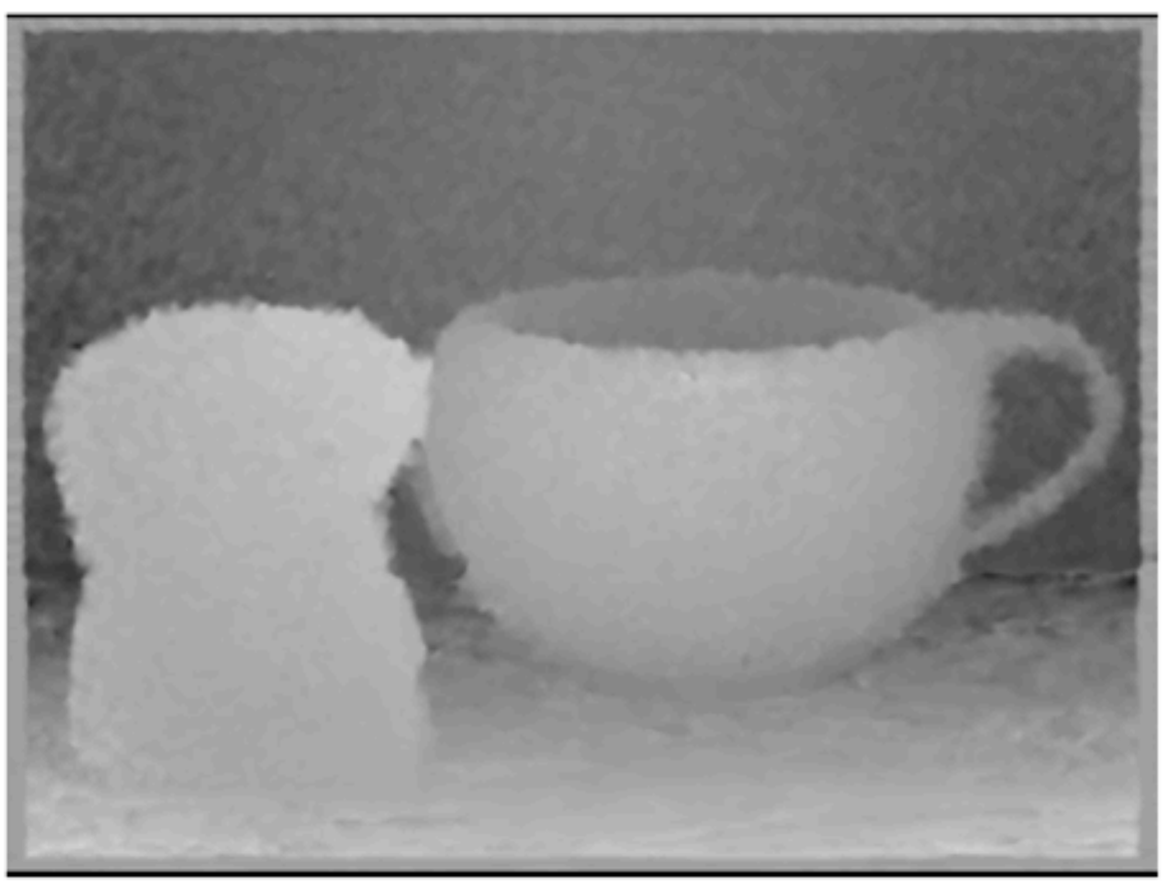

(c) gray-coded depth map without post-filtering

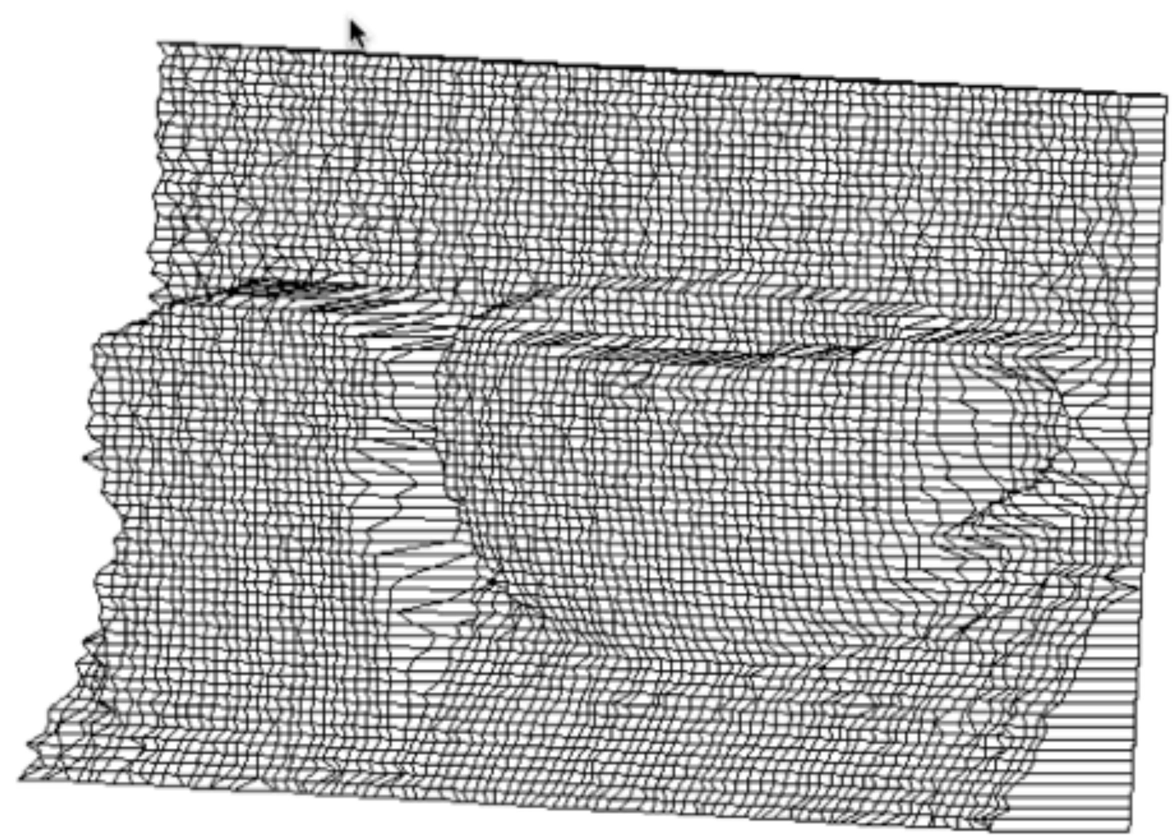

(d) wireframe plot of (c)

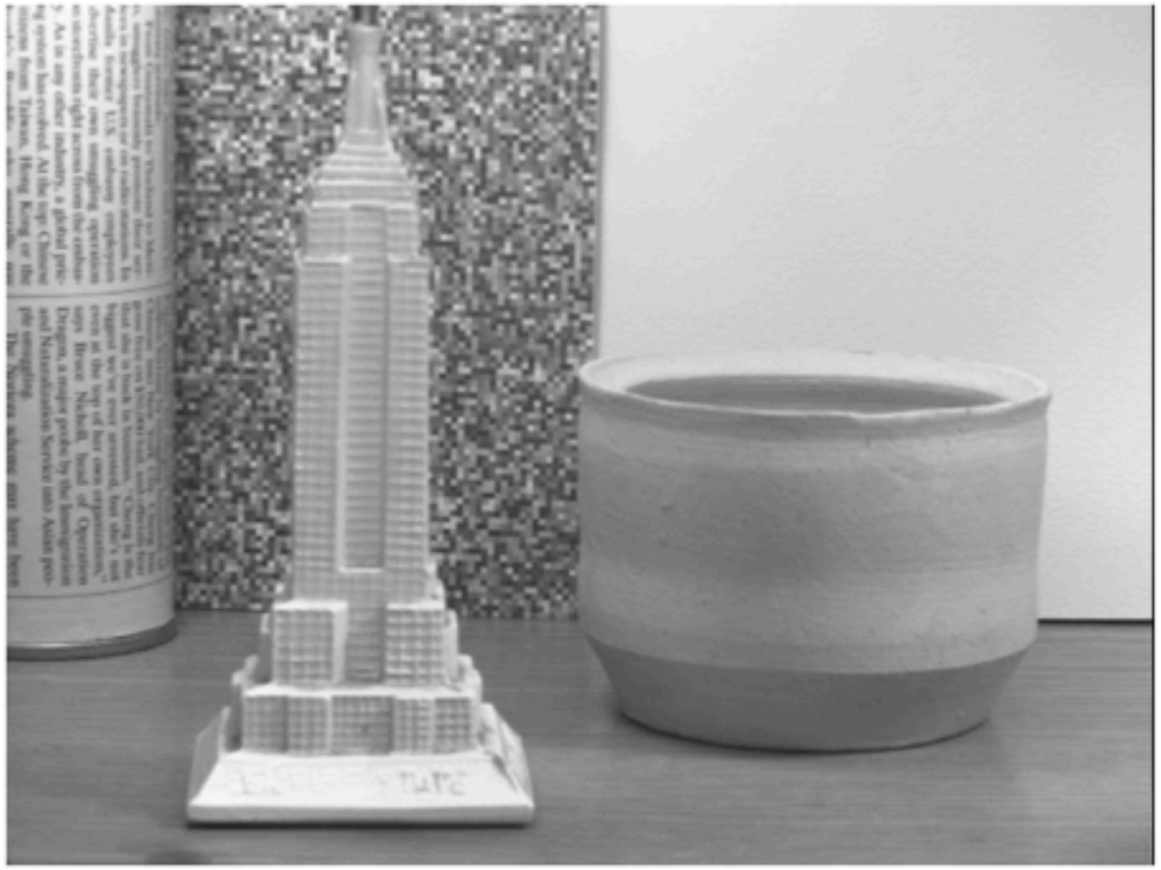

(a) far-focused image

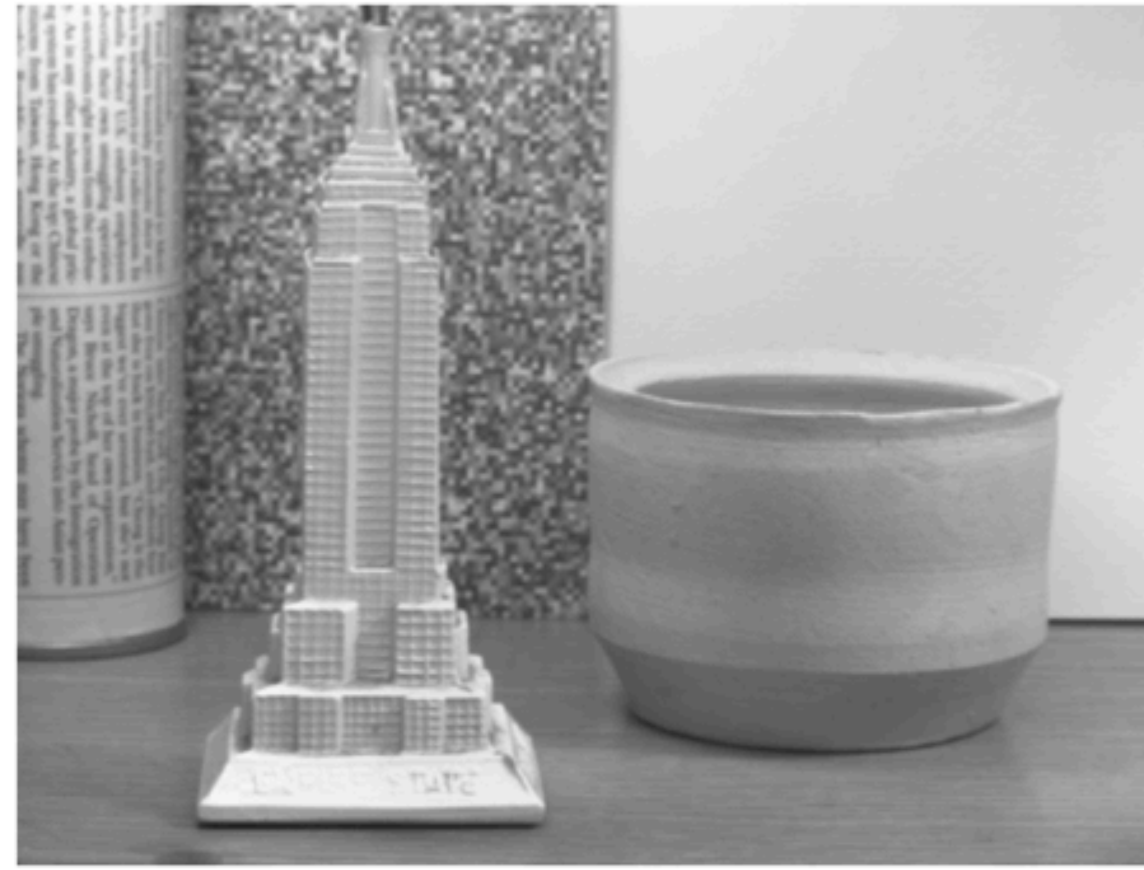

(b) near-focused image

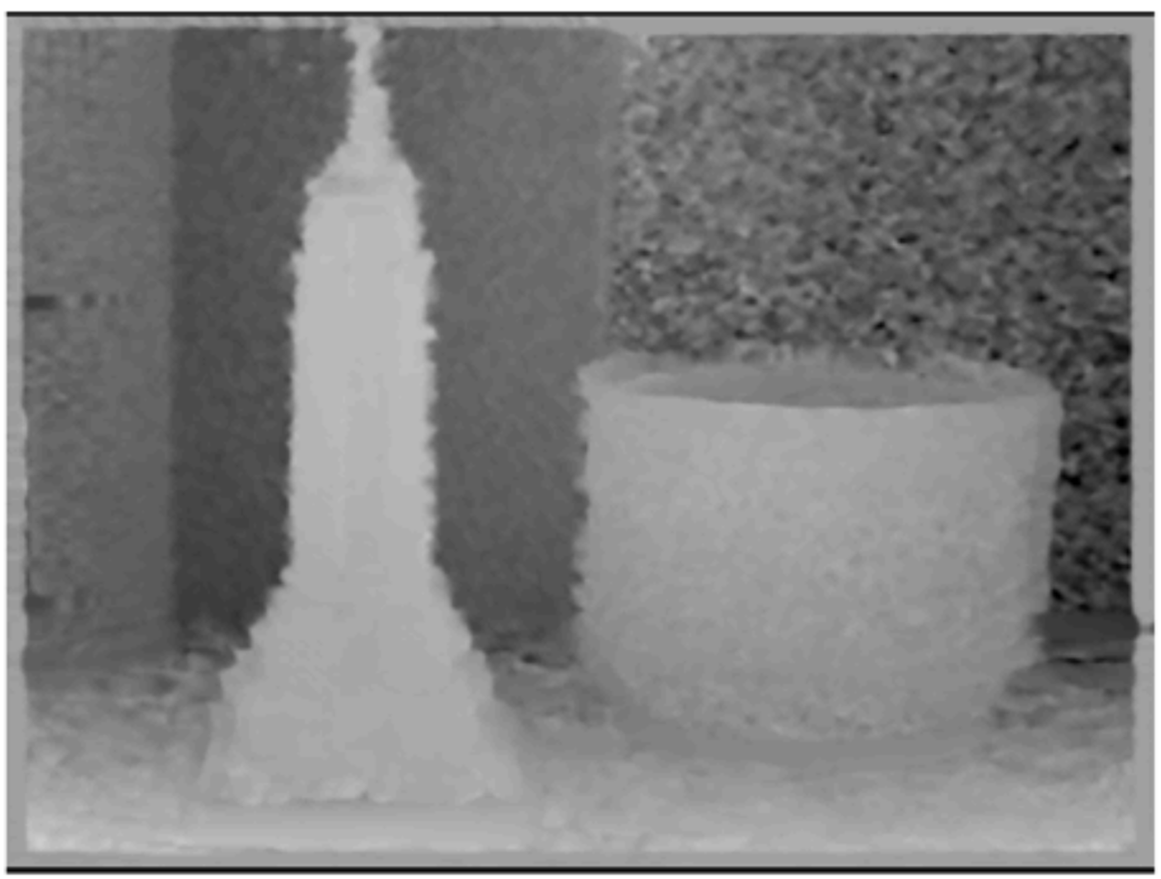

(c) gray-coded depth map

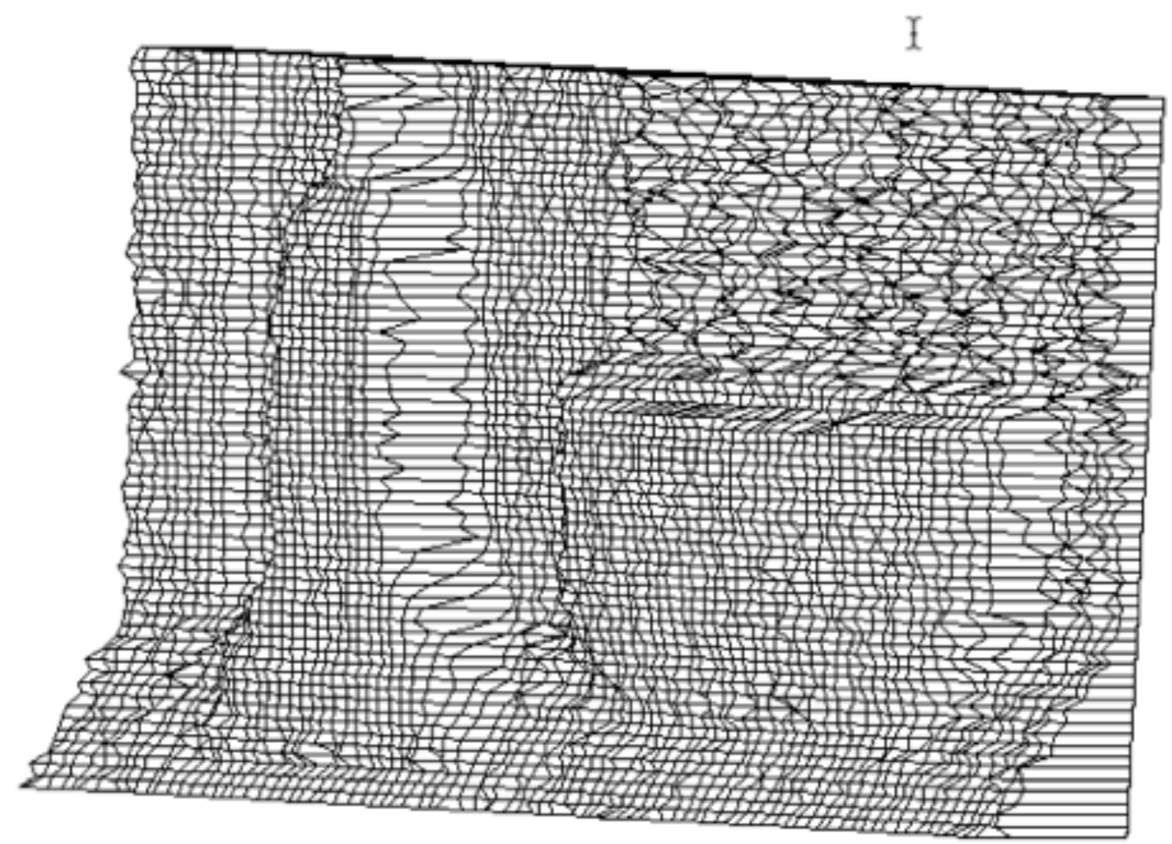

(d) wireframe of the depth map

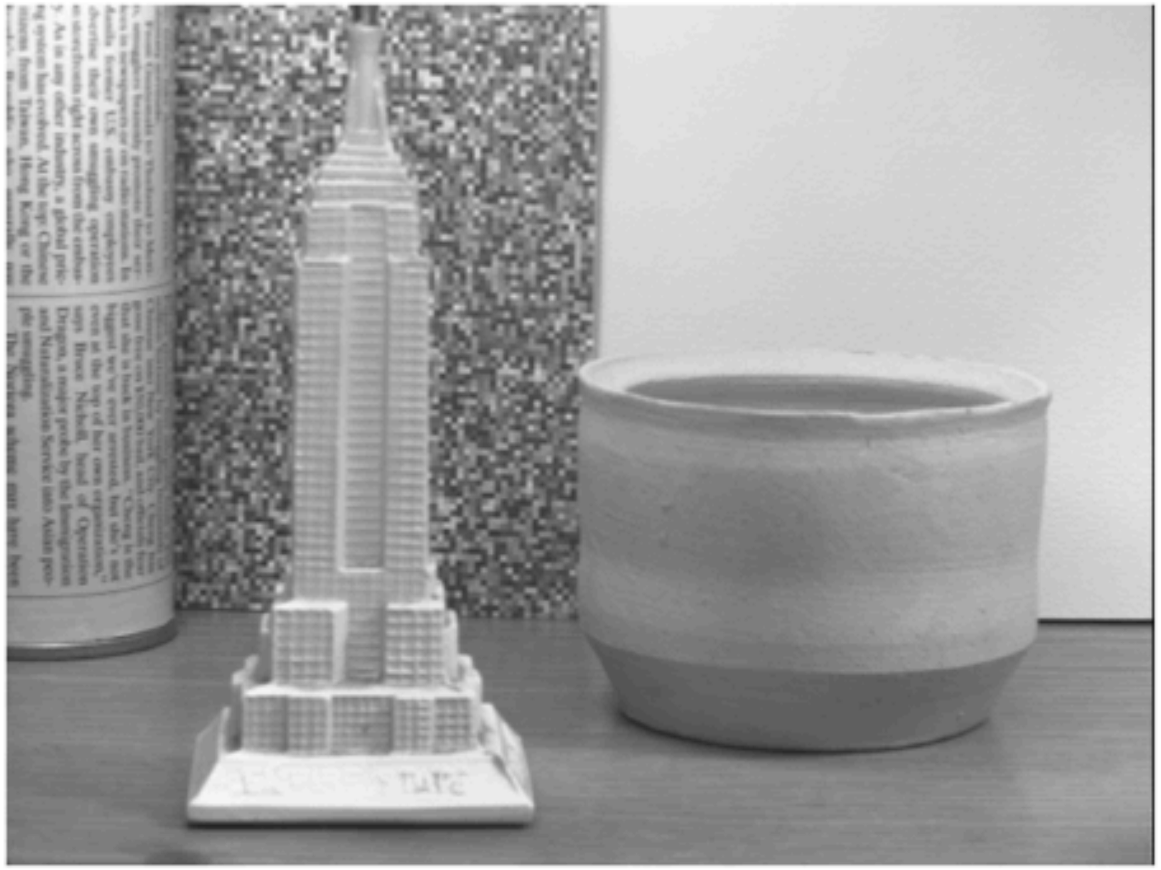

(a) far-focused image

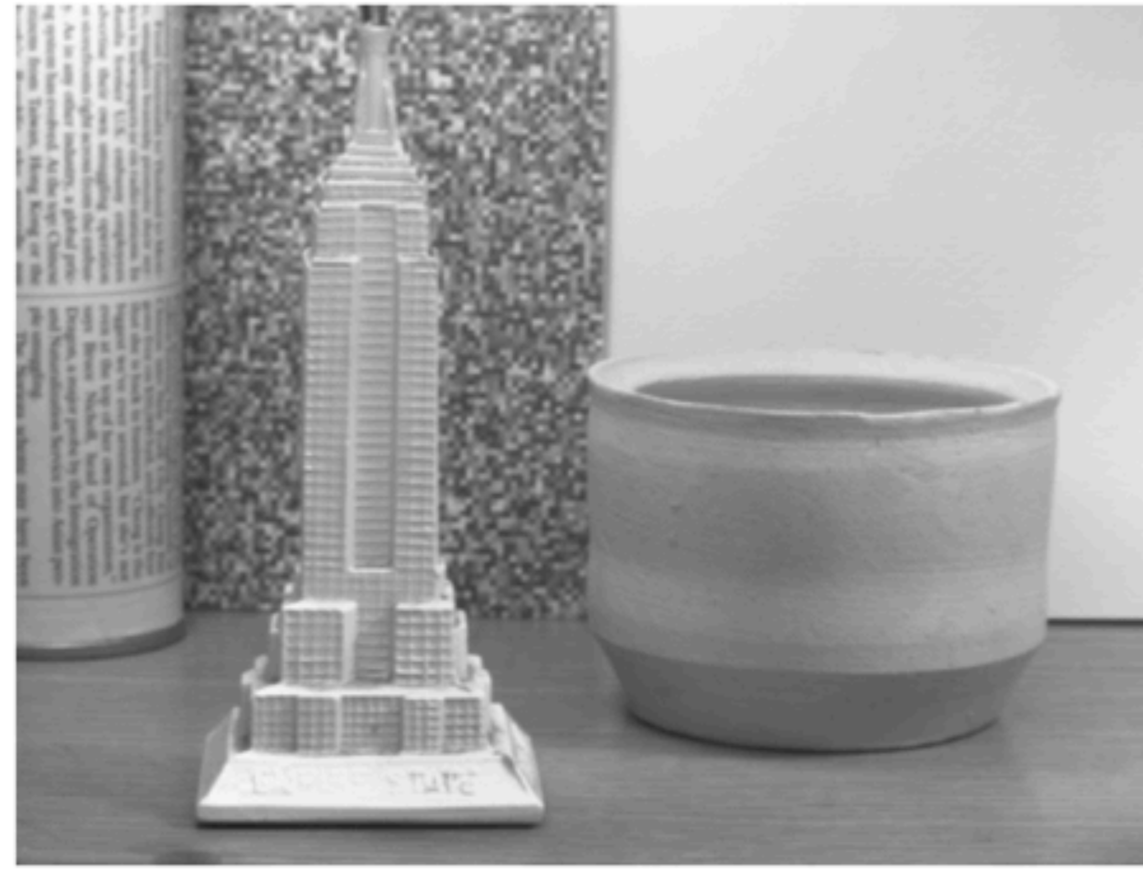

(b) near-focused image

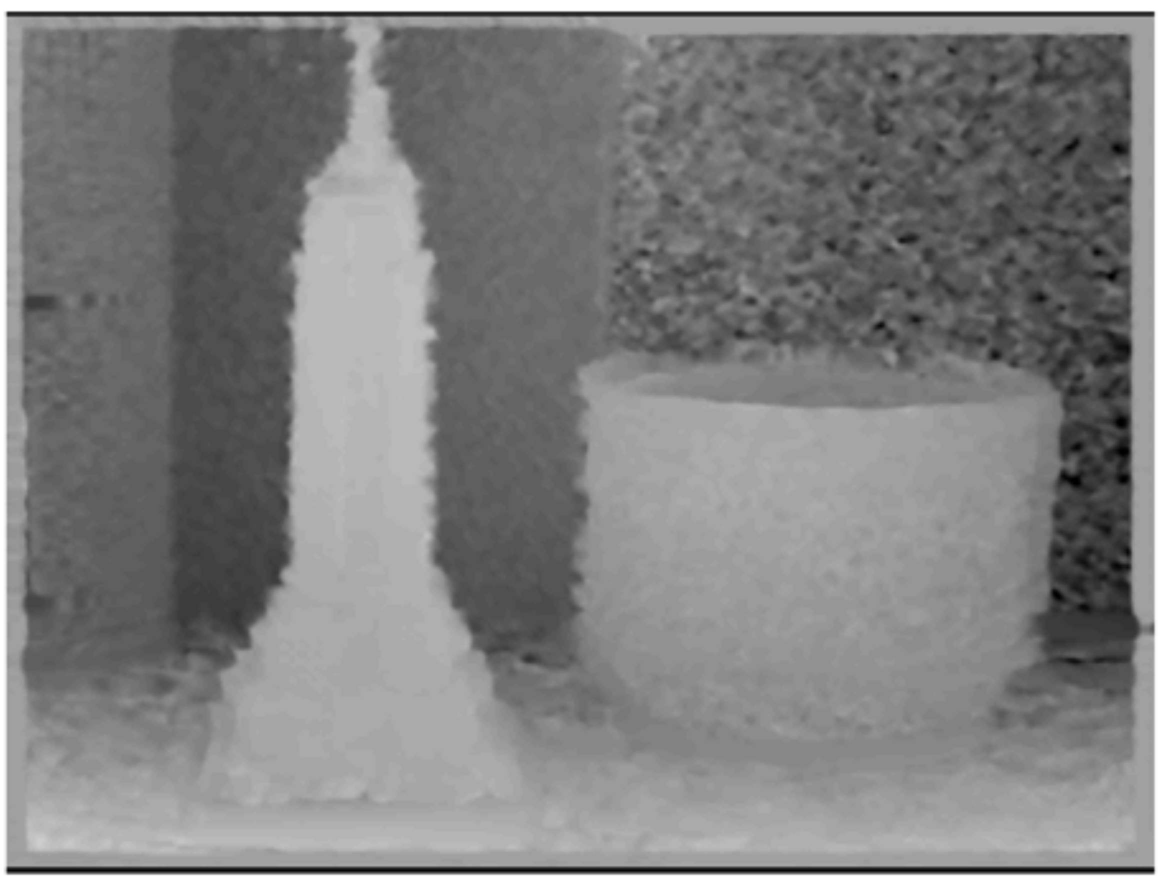

(c) gray-coded depth map

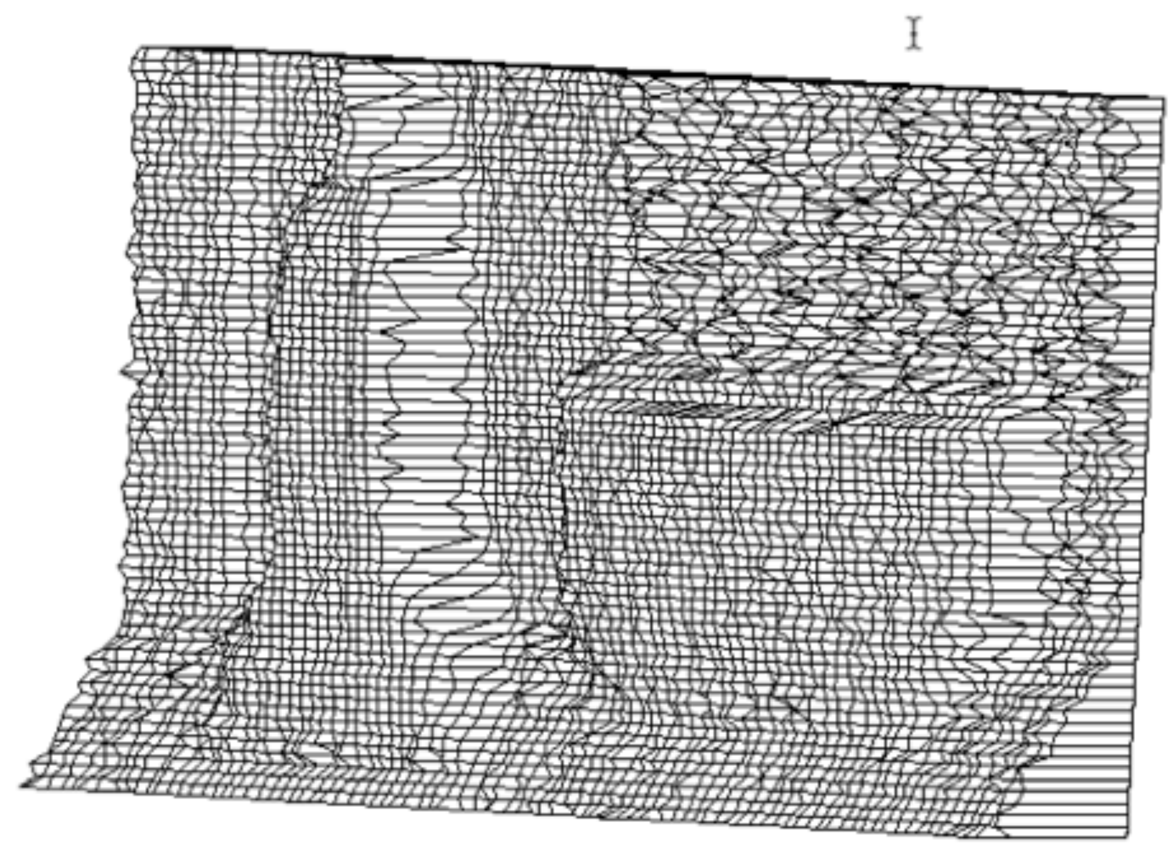

(d) wireframe of the depth map

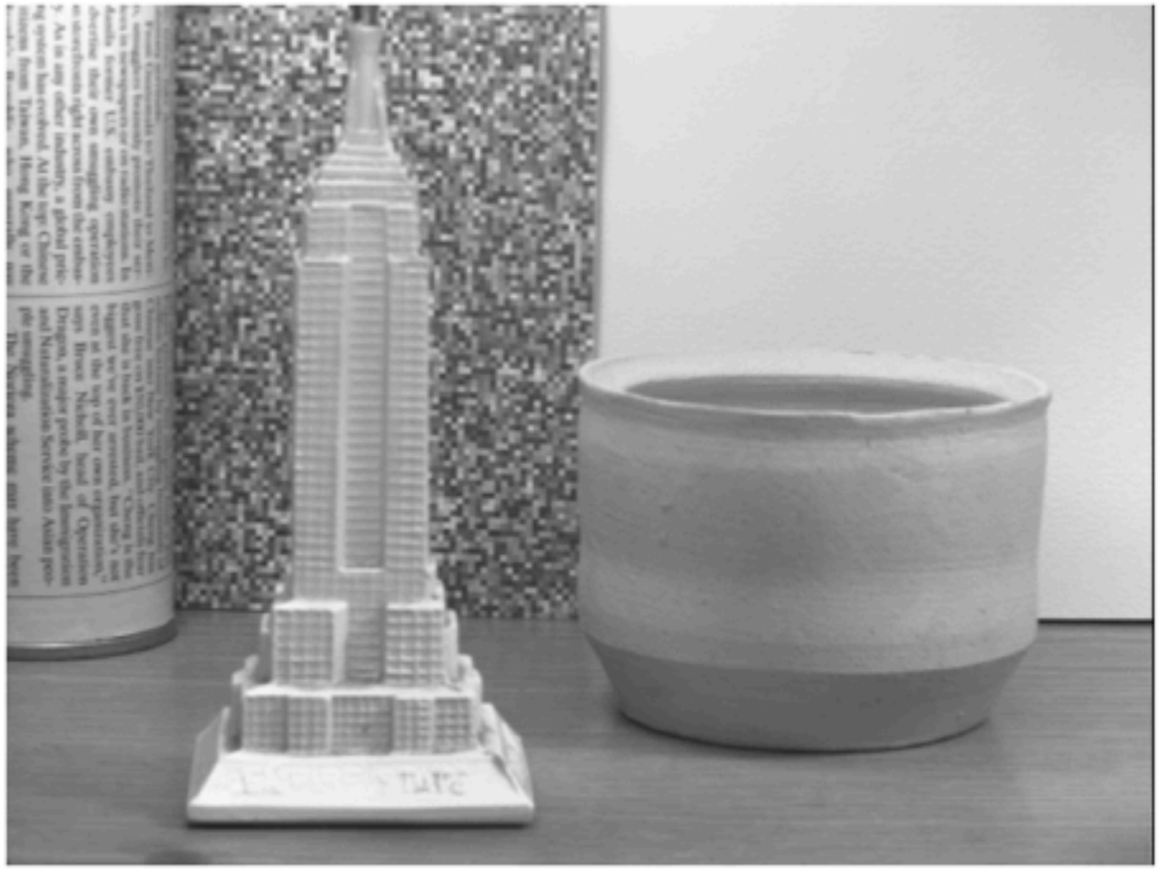

(a) far-focused image

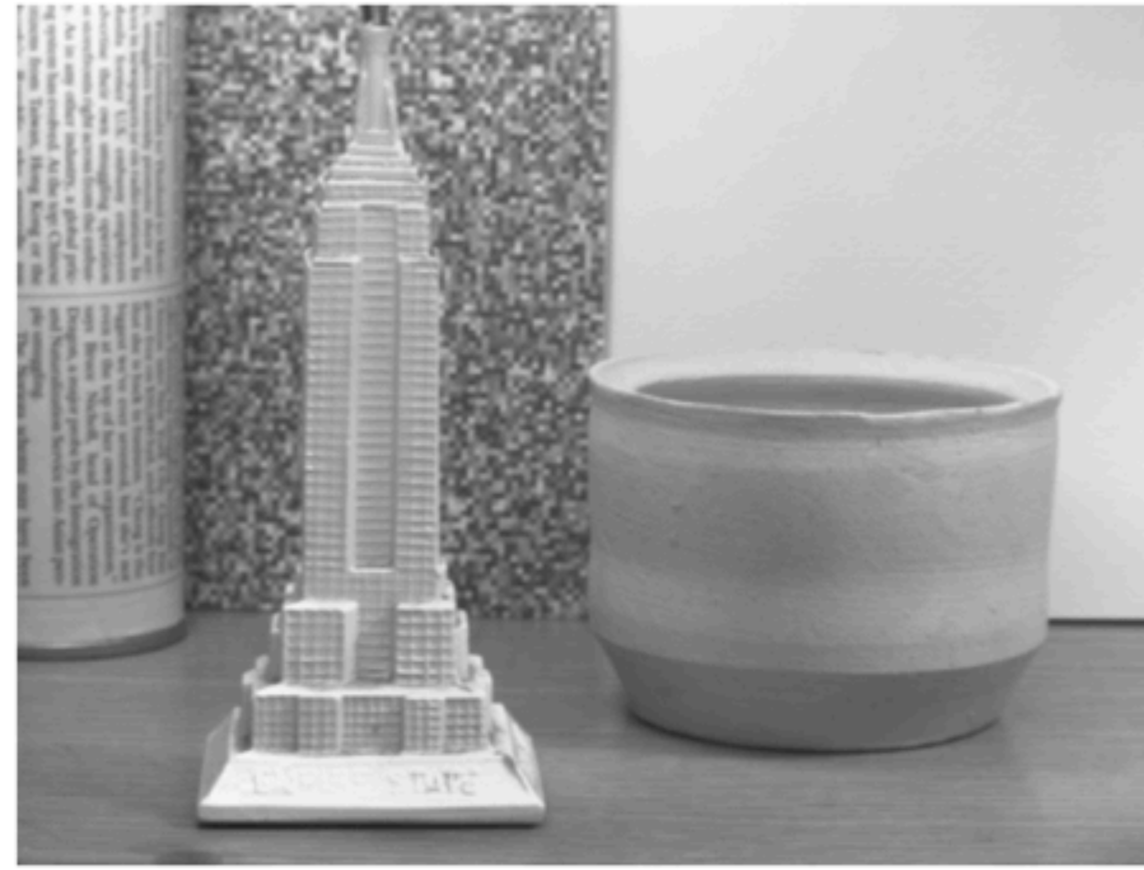

(b) near-focused image

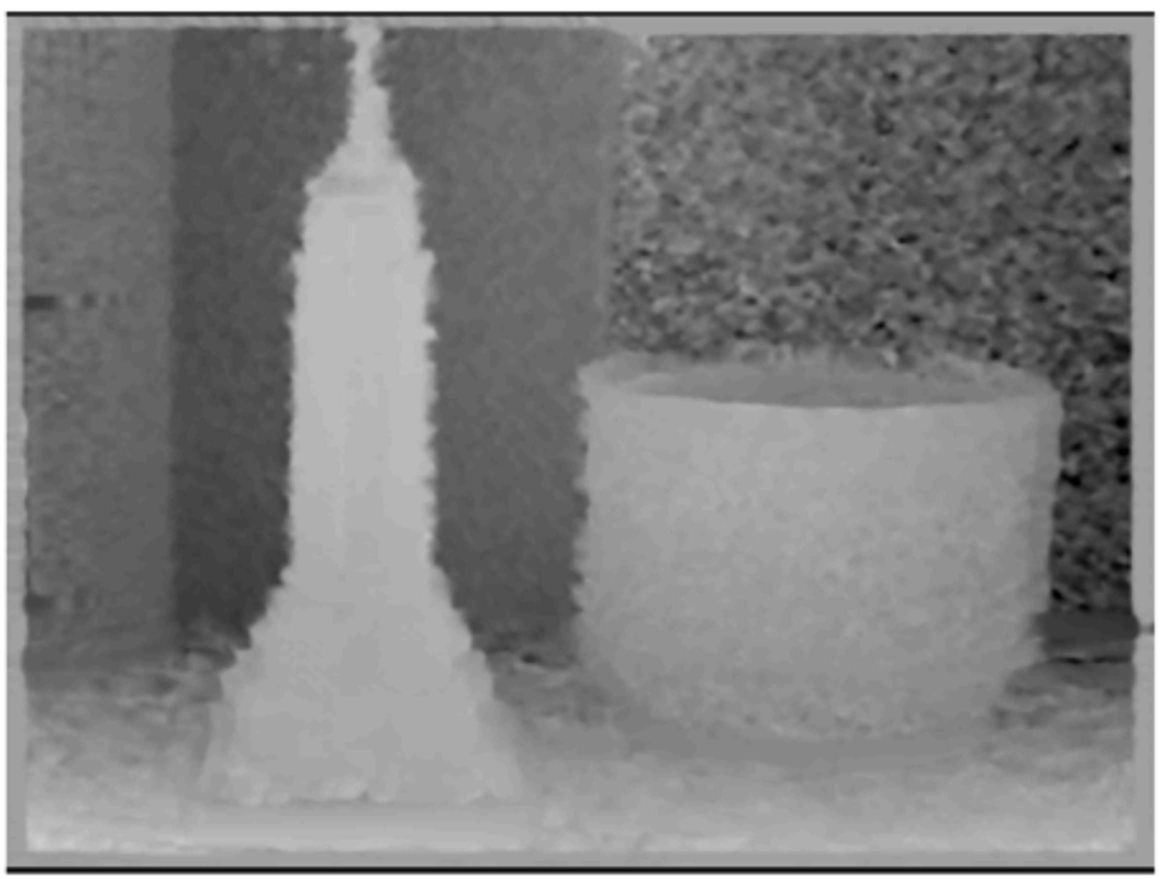

(c) gray-coded depth map

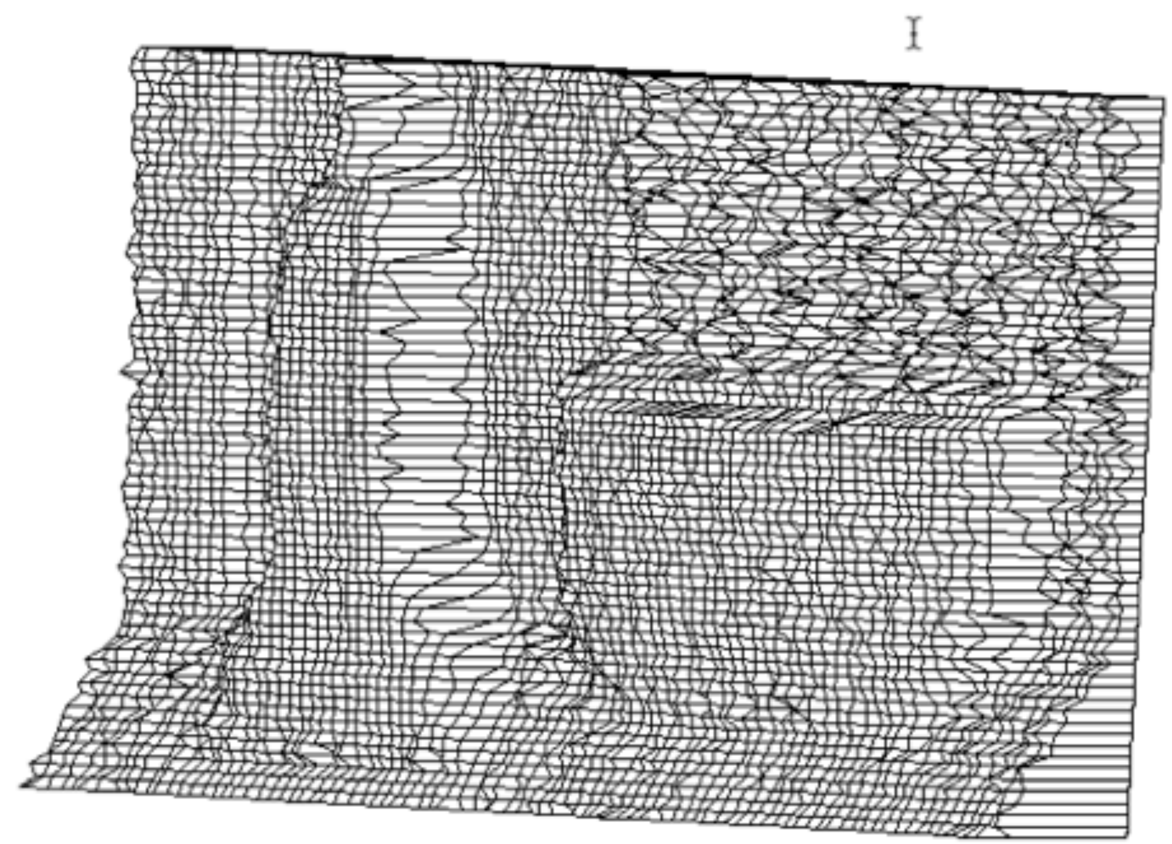

(d) wireframe of the depth map

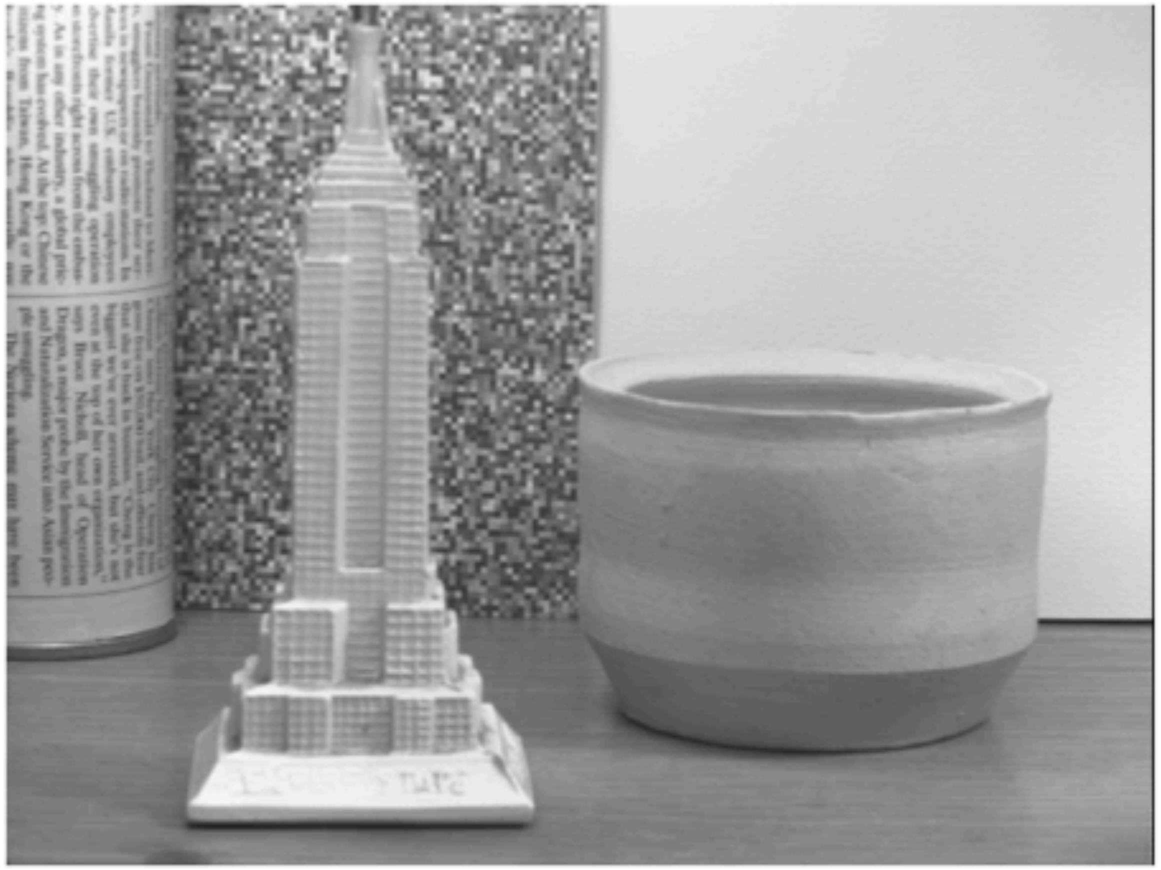

(a) far-focused image

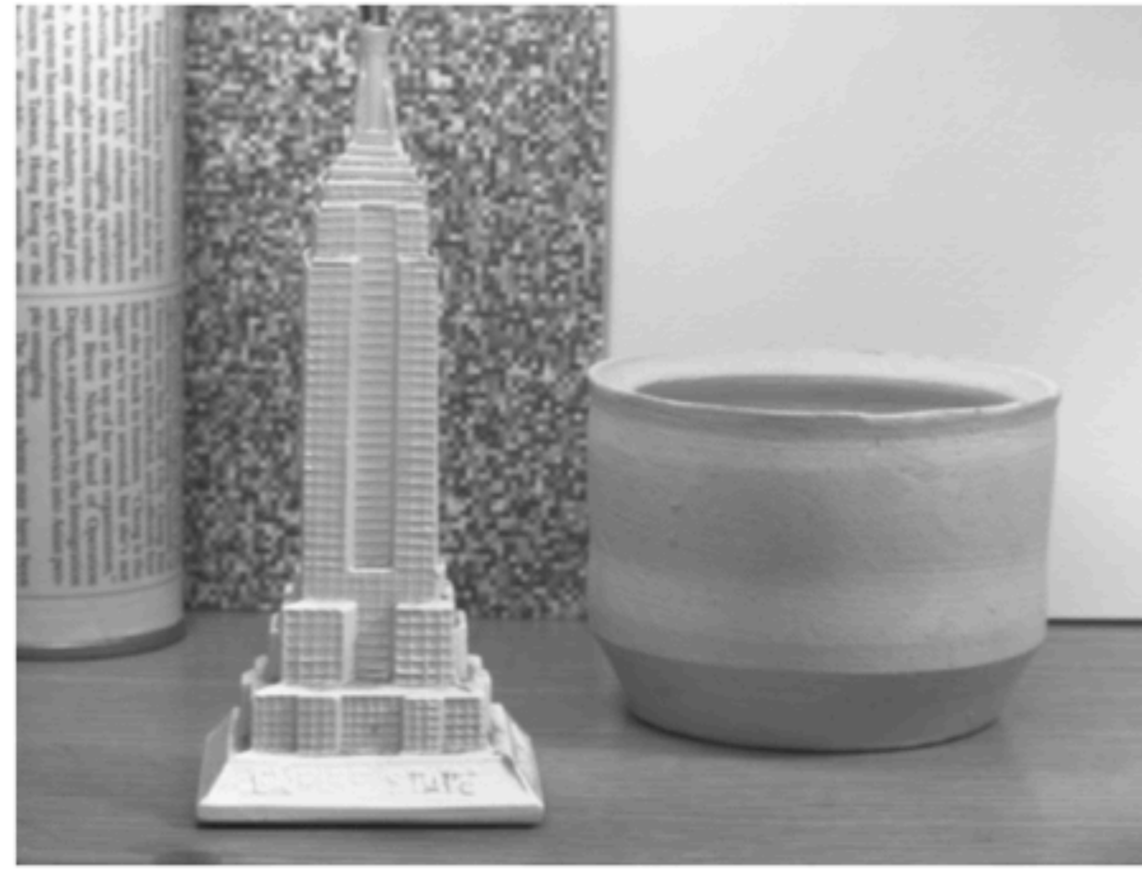

(b) near-focused image

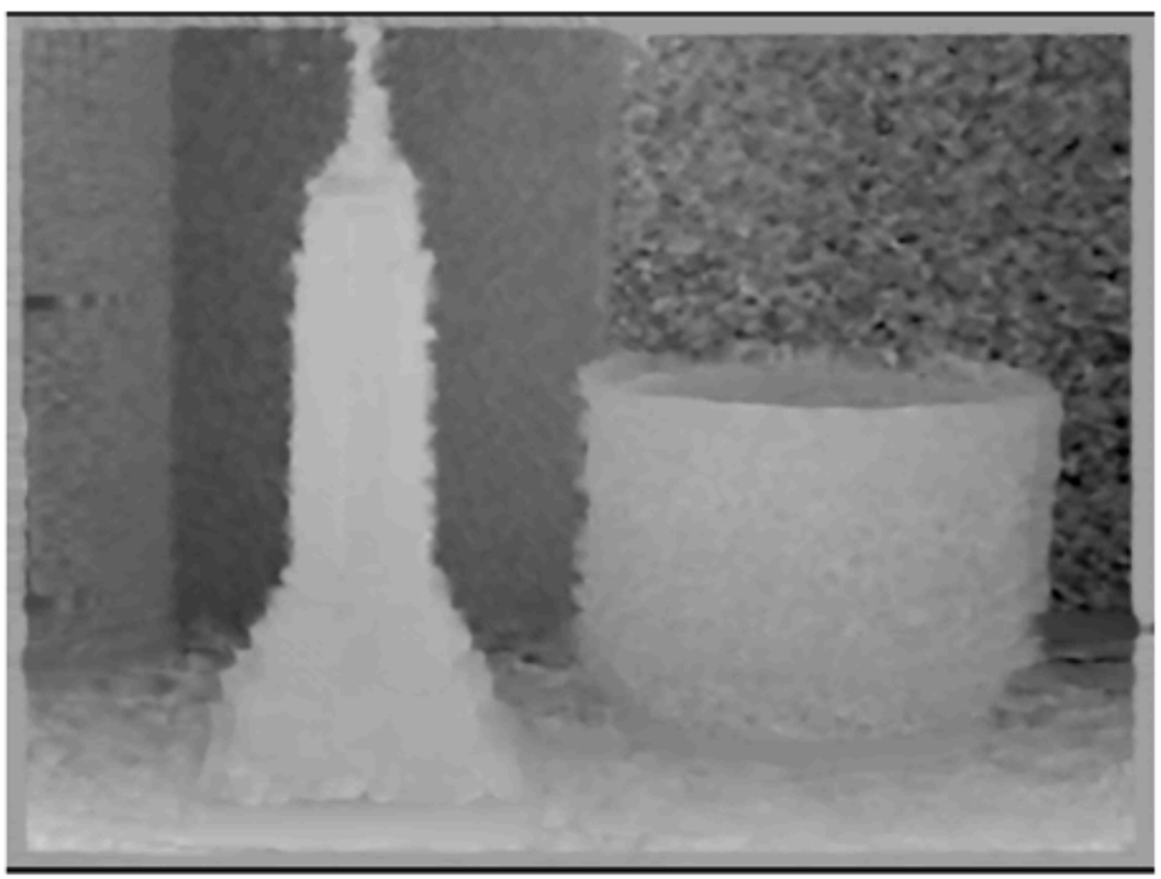

(c) gray-coded depth map

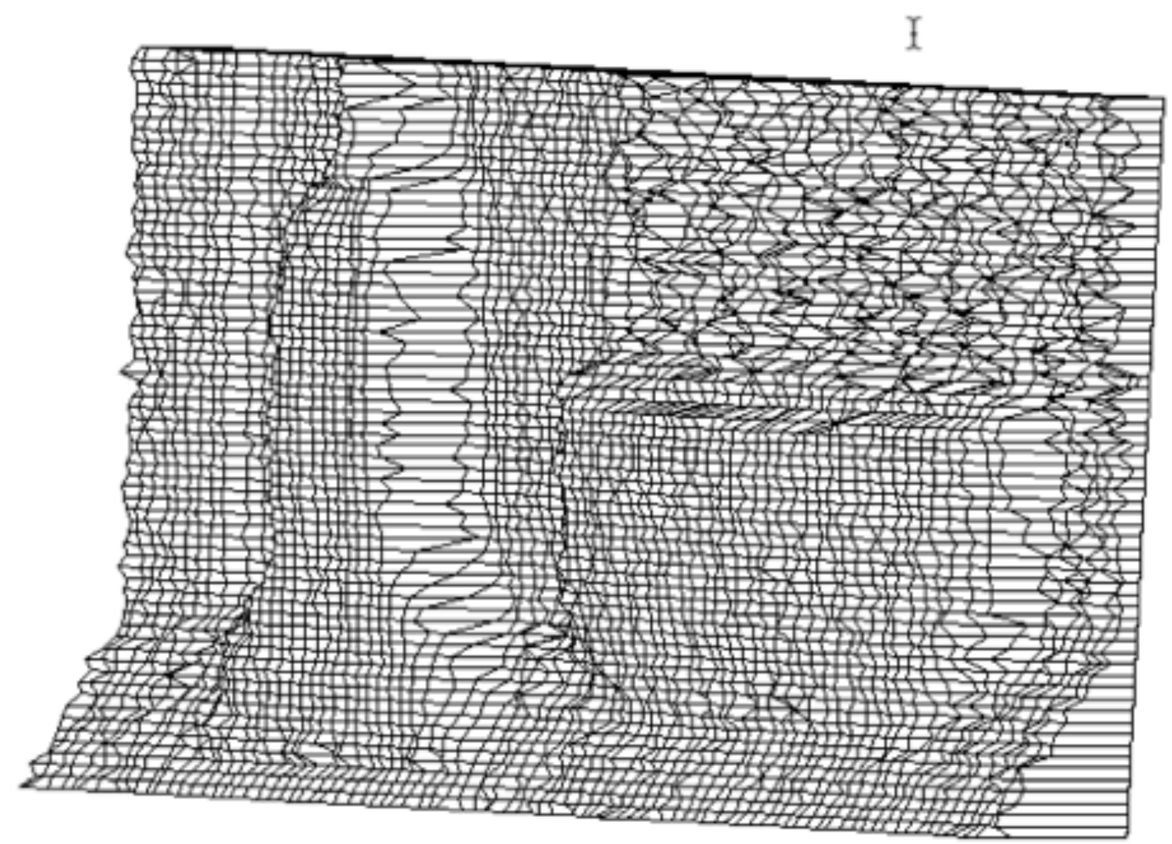

(d) wireframe of the depth map

How does the Kinect Gesture Tracking Work

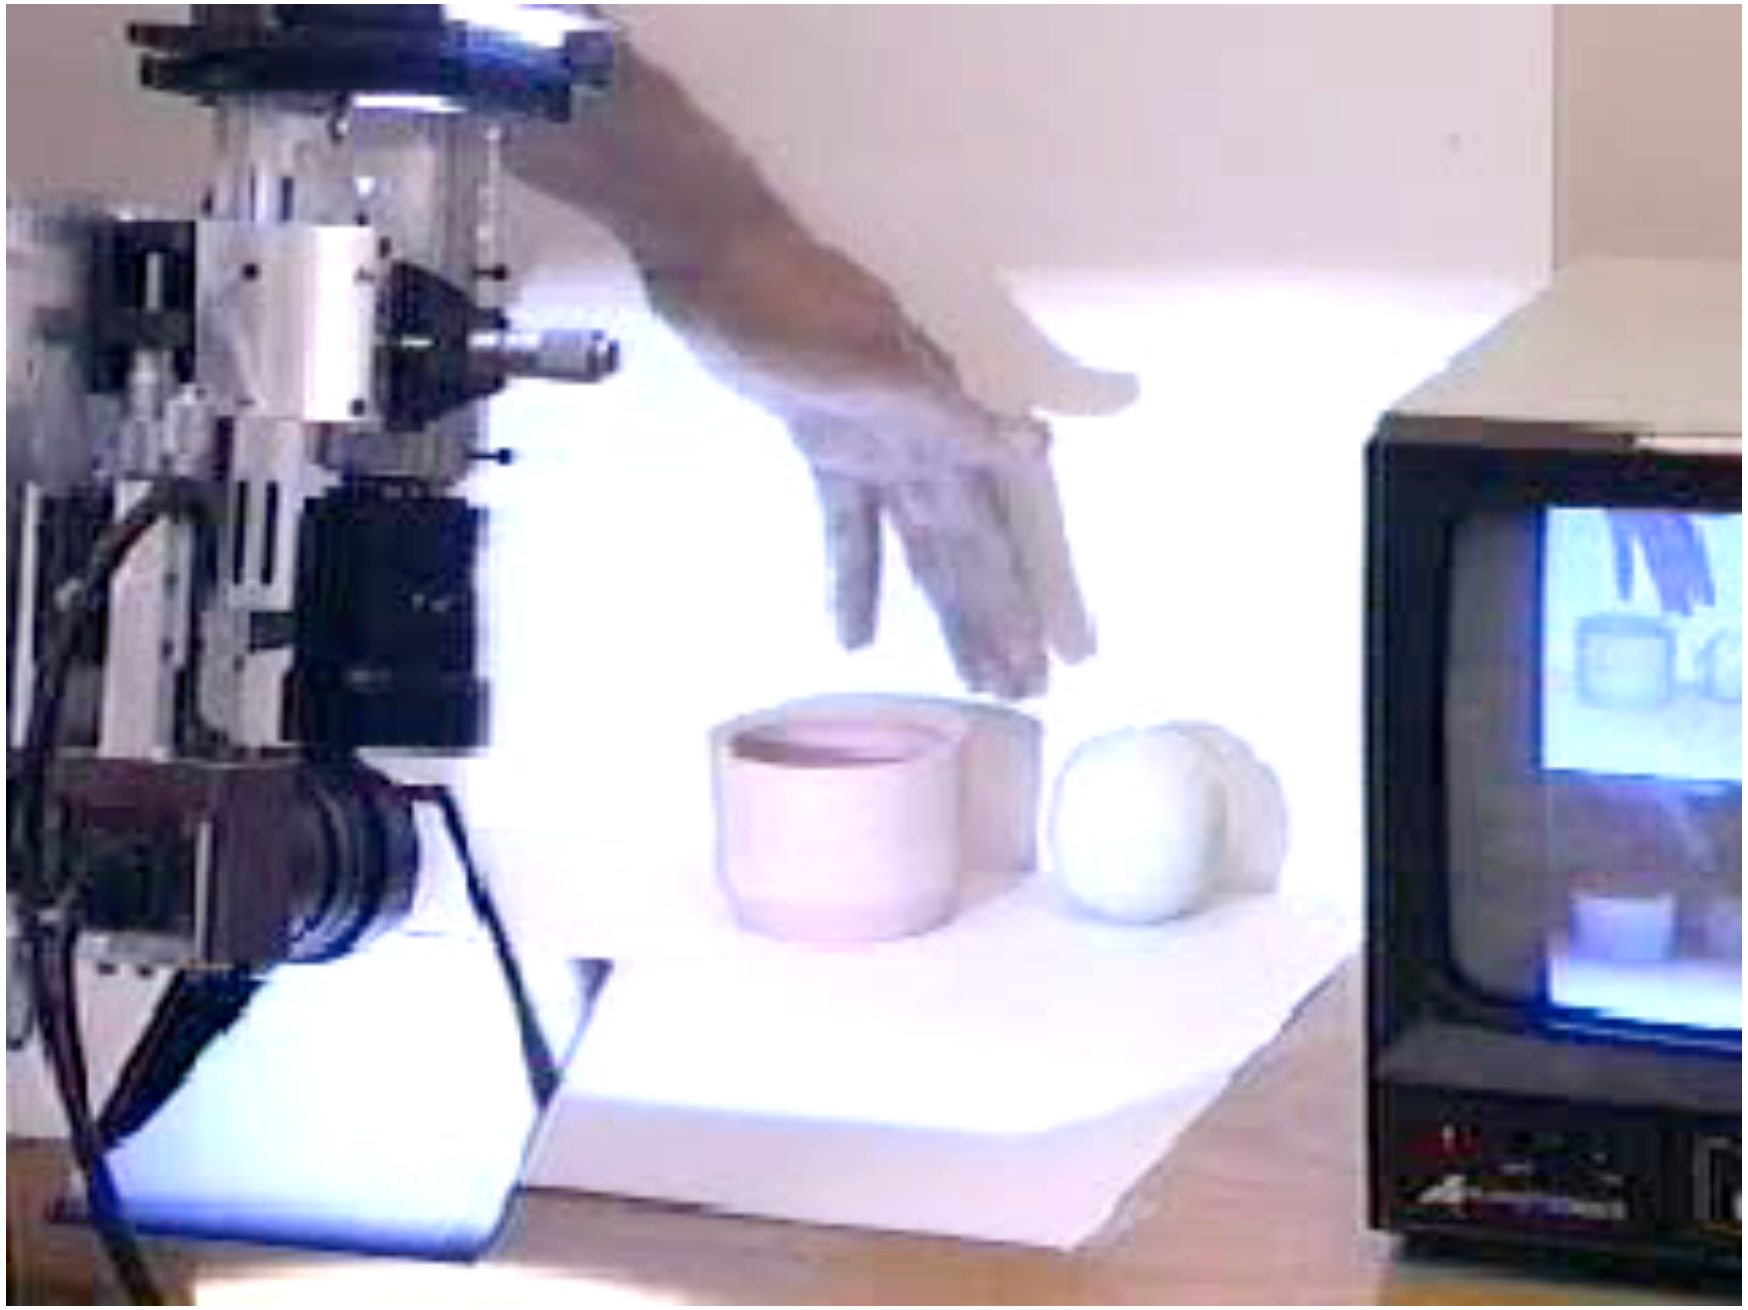

How does the Kinect Gesture Tracking Work

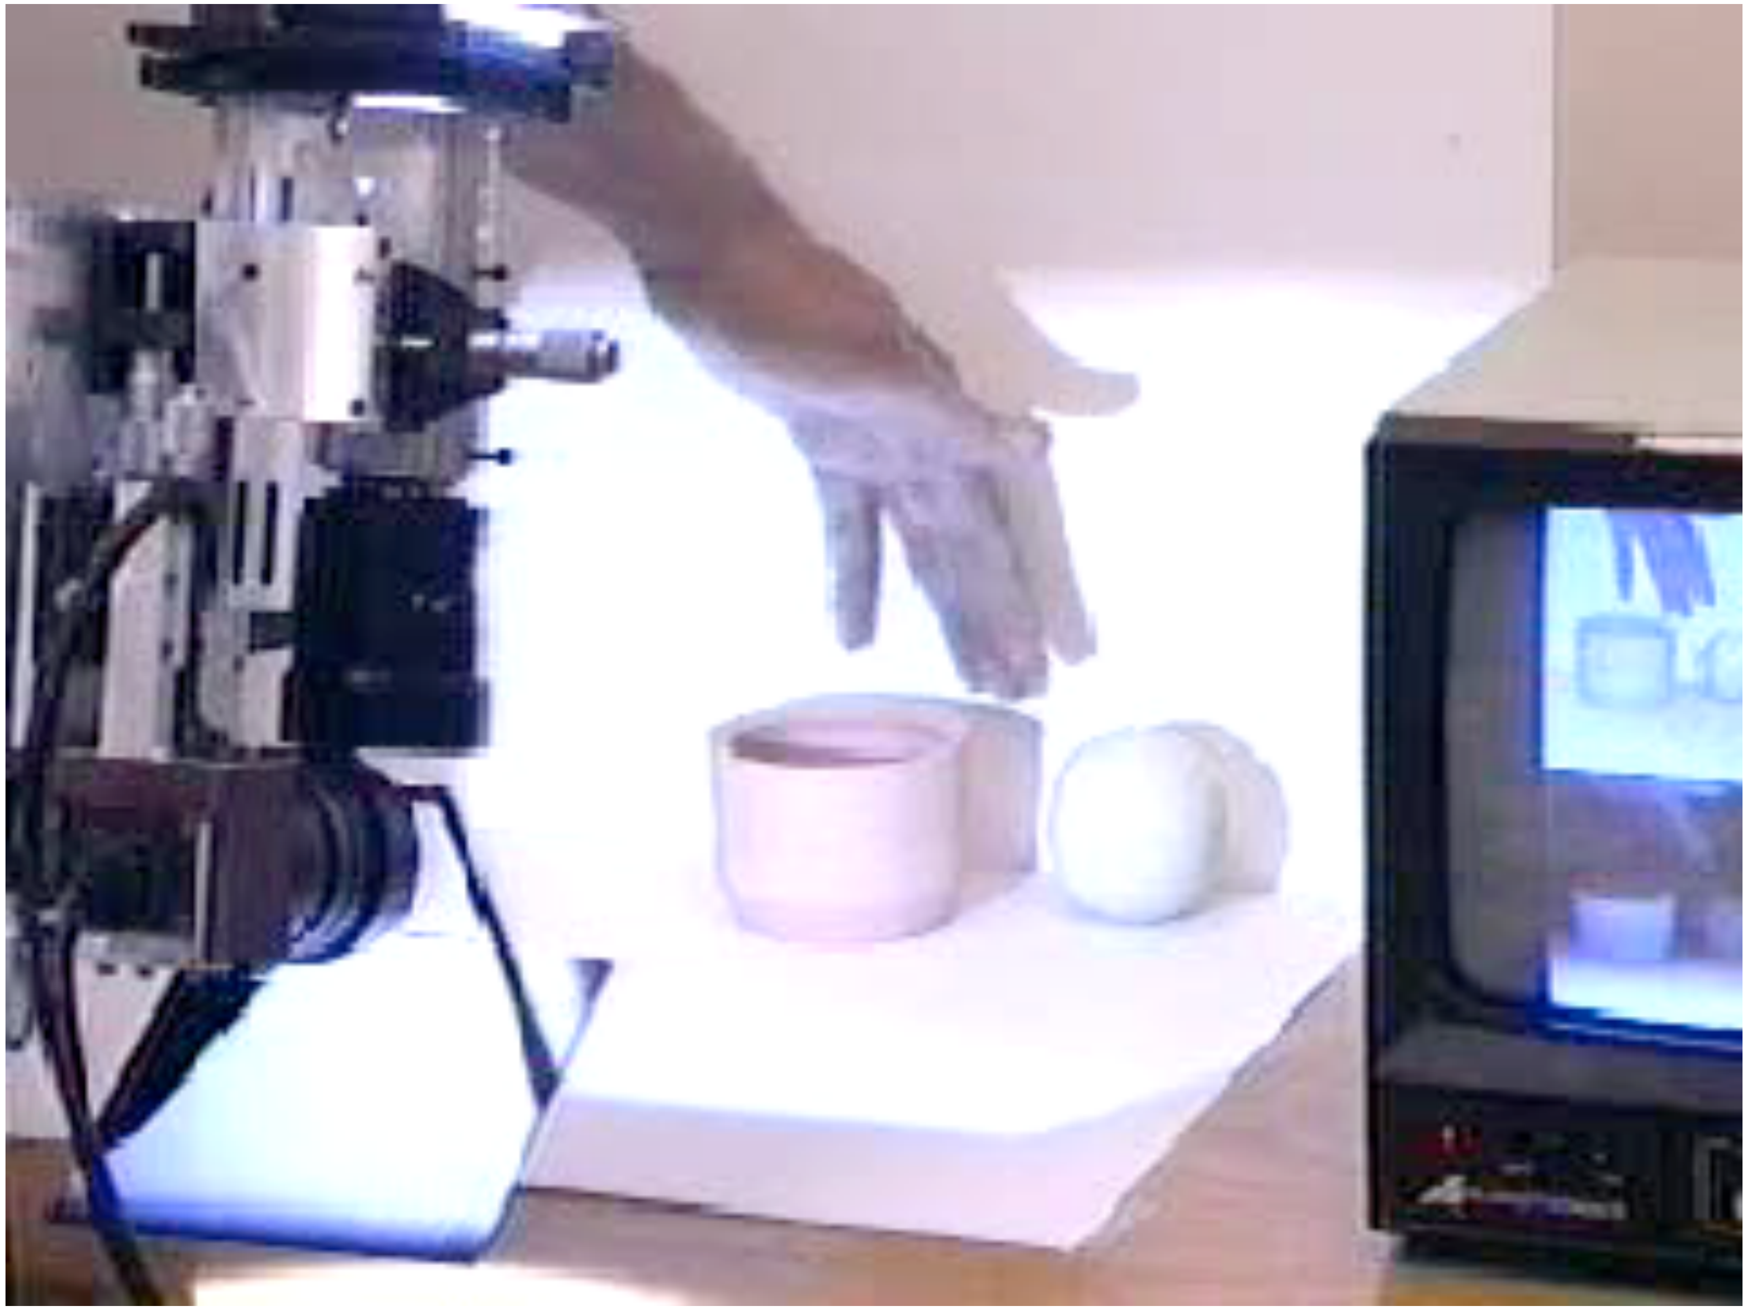

http://www.cs.columbia.edu/CAVE/projects/depth_defocus/images/defocus_ball.mpg

How does the Kinect Gesture Tracking Work

- Depth from focus

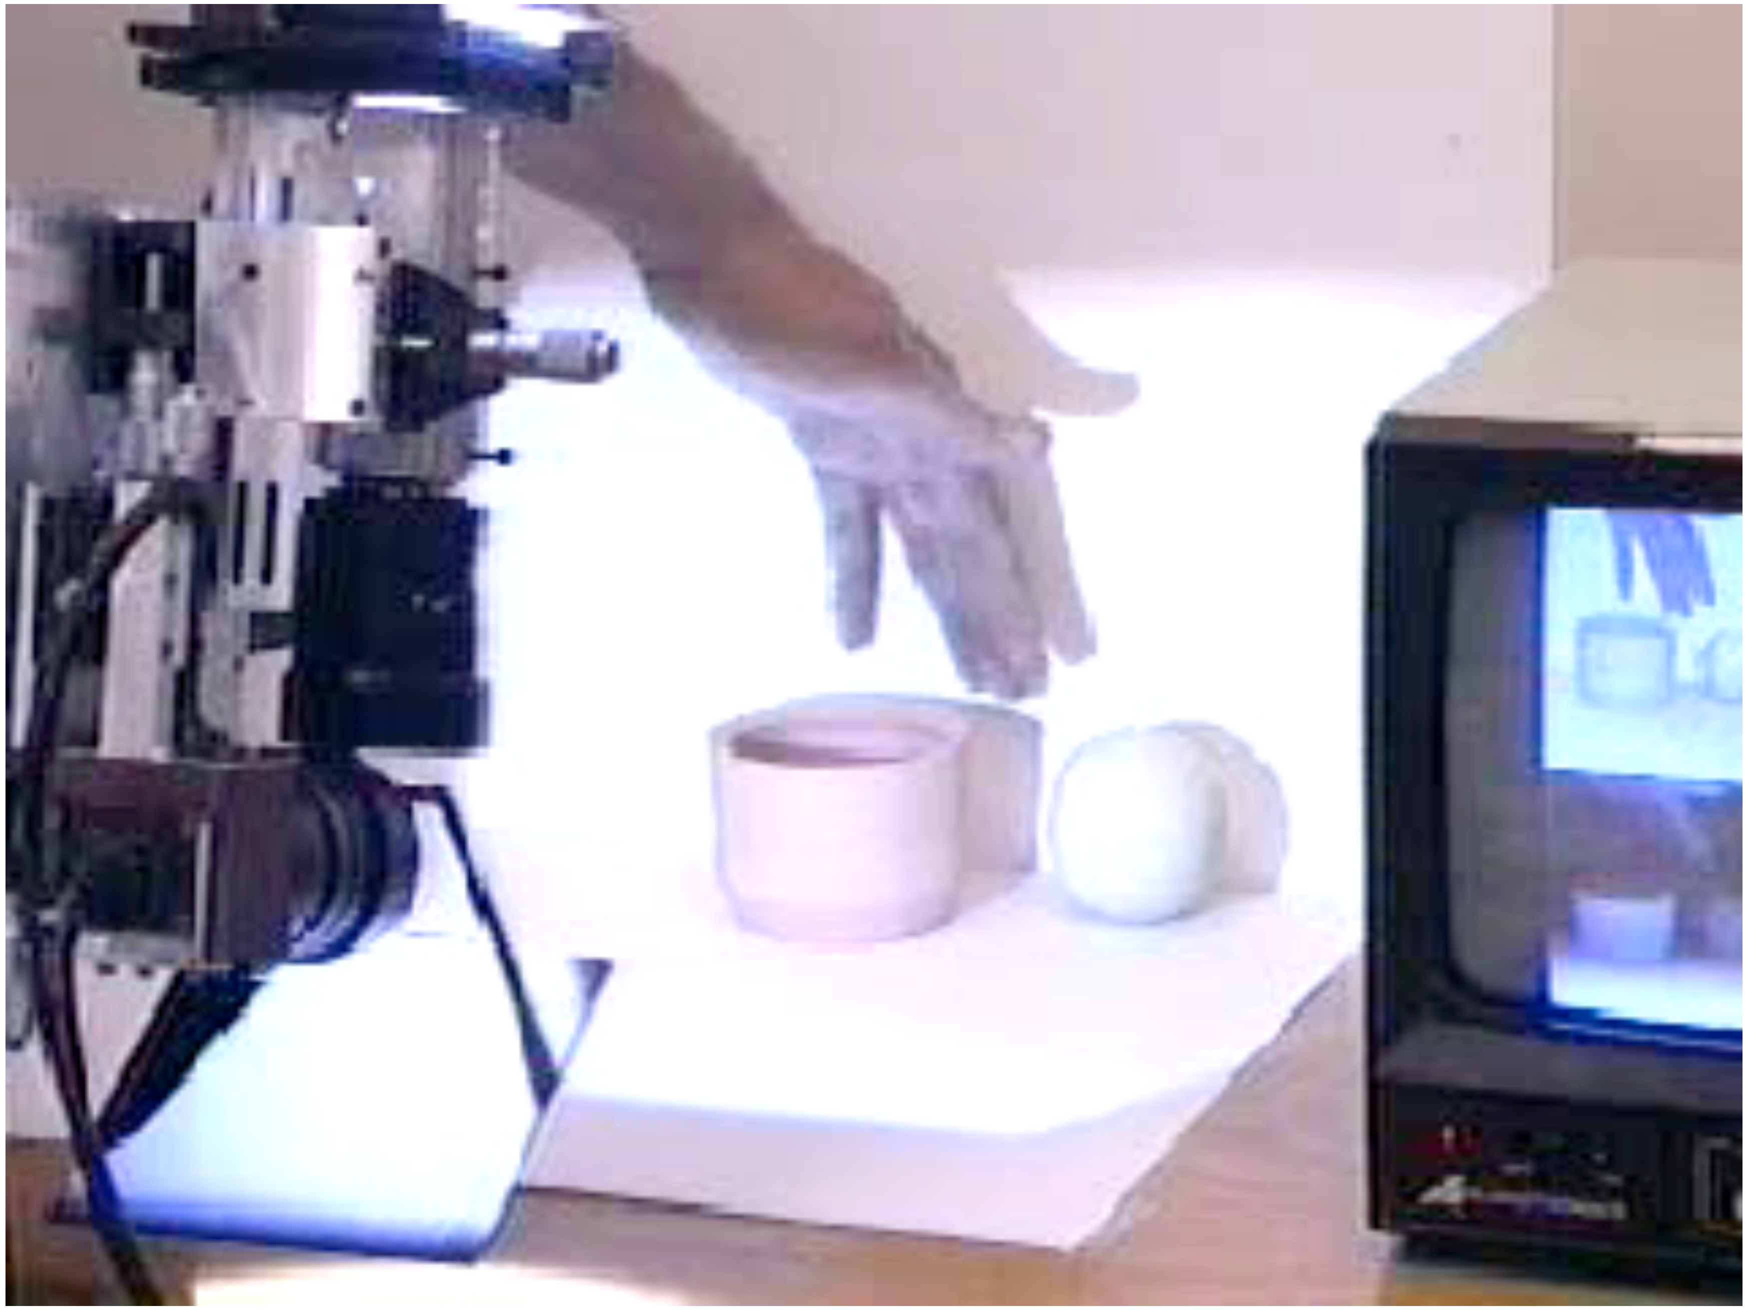

http://www.cs.columbia.edu/CAVE/projects/depth_defocus/images/defocus_ball.mpg

How does the Kinect Gesture Tracking Work

- Depth from focus
  - Blurry things are further away

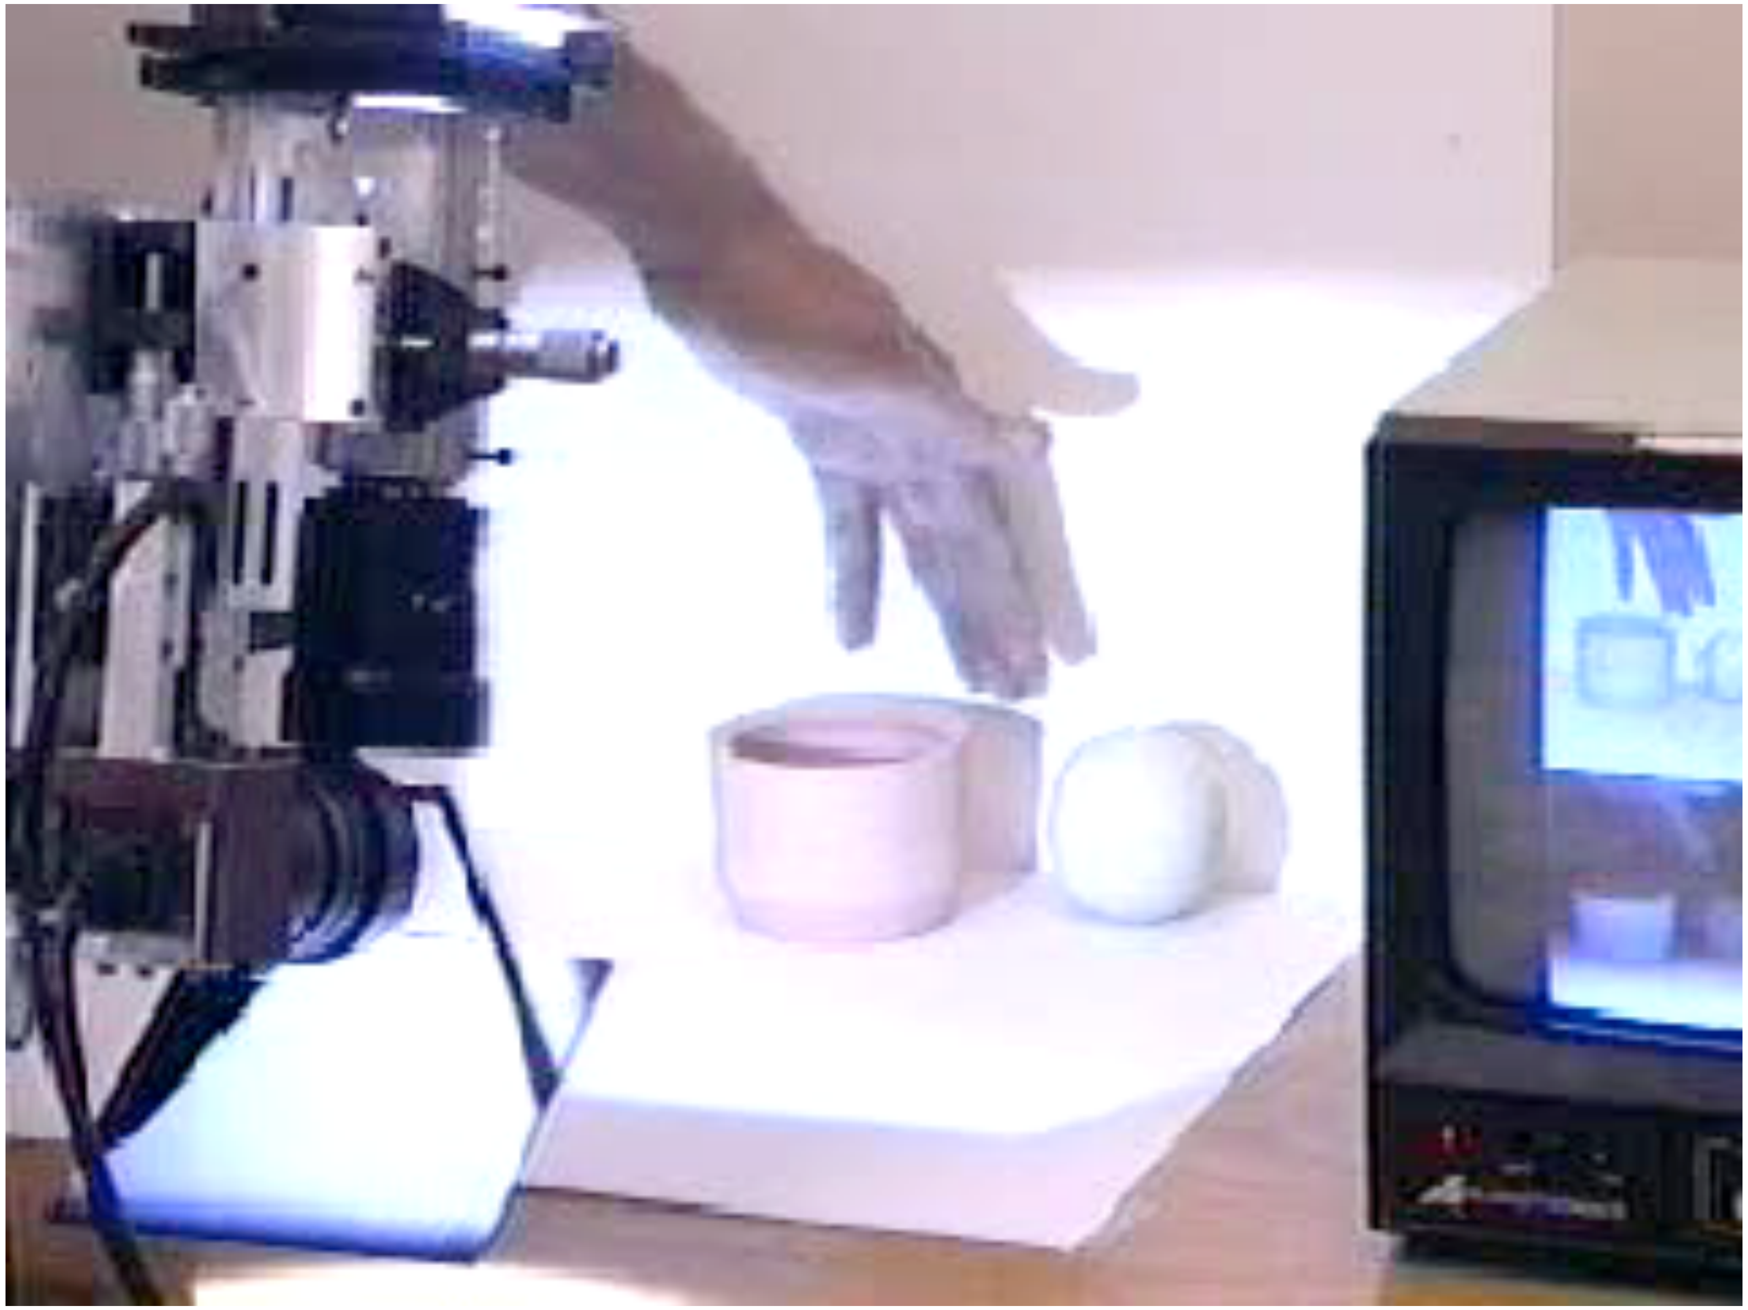

http://www.cs.columbia.edu/CAVE/projects/depth_defocus/images/defocus_ball.mpg

How does the Kinect Gesture Tracking Work

How does the Kinect Gesture Tracking Work

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

How does the Kinect Gesture Tracking Work

- Depth from focus

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

How does the Kinect Gesture Tracking Work

- Depth from focus
  - The Kinect dramatically improves the accuracy of traditional depth from focus

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

How does the Kinect Gesture Tracking Work

- Depth from focus
  - The Kinect dramatically improves the accuracy of traditional depth from focus
  - The Kinect uses a special (“astigmatic”) lens with different focal length in x and y directions

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

How does the Kinect Gesture Tracking Work

- Depth from focus
  - The Kinect dramatically improves the accuracy of traditional depth from focus
  - The Kinect uses a special (“astigmatic”) lens with different focal length in x and y directions
  - A projected circle then becomes an ellipse whose orientation depends on depth

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

How does the Kinect Gesture Tracking Work

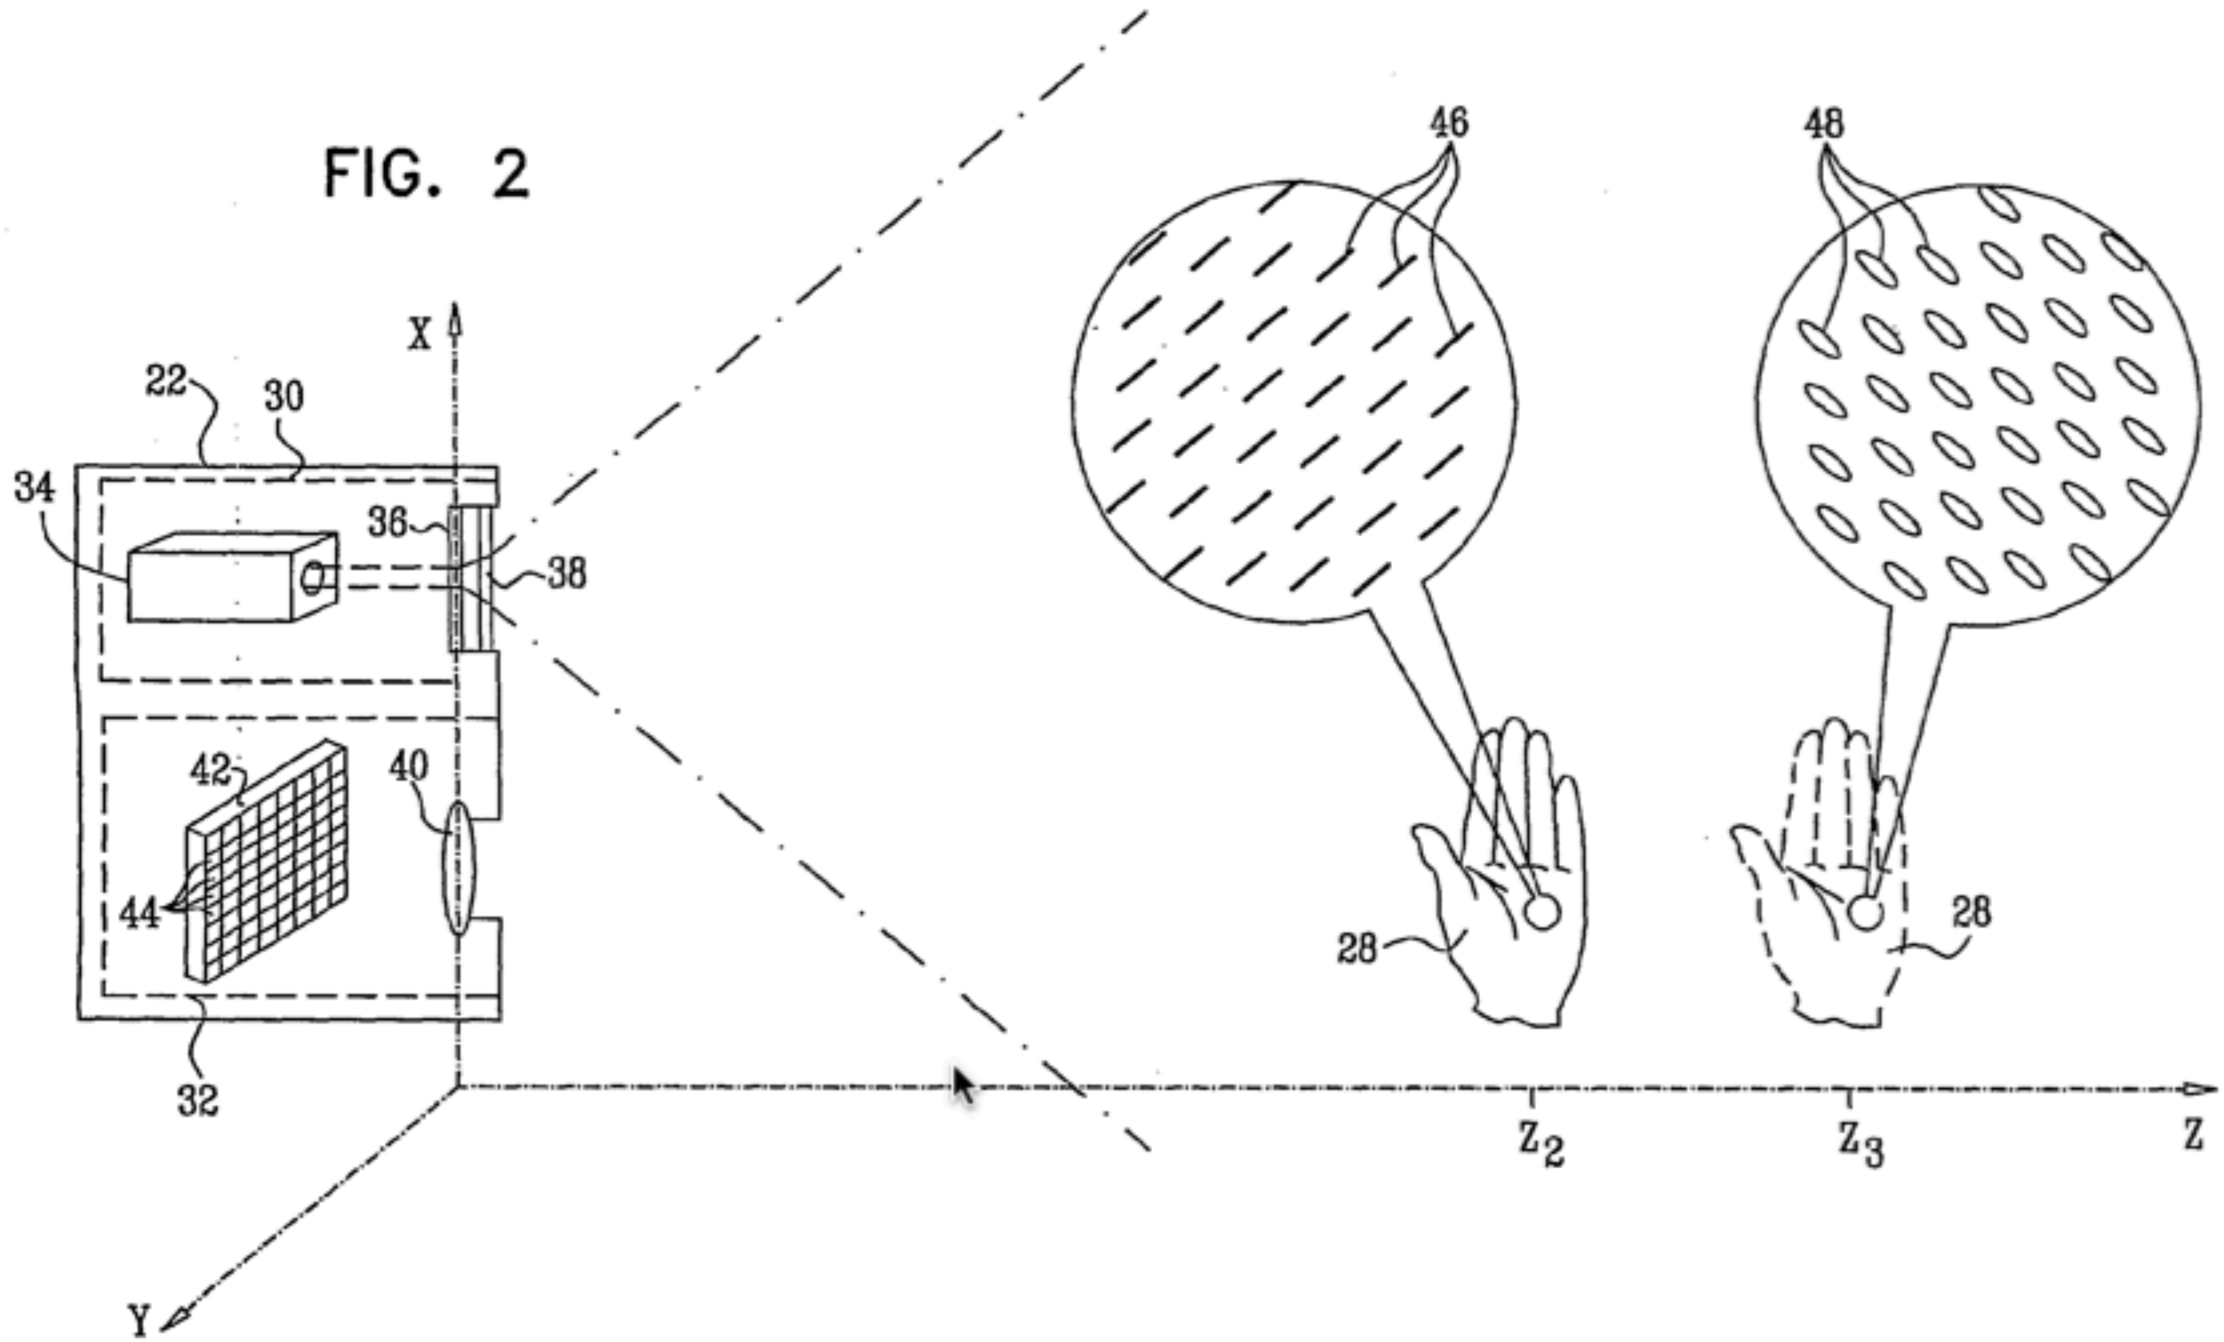

How does the Kinect Gesture Tracking Work

FIG. 2

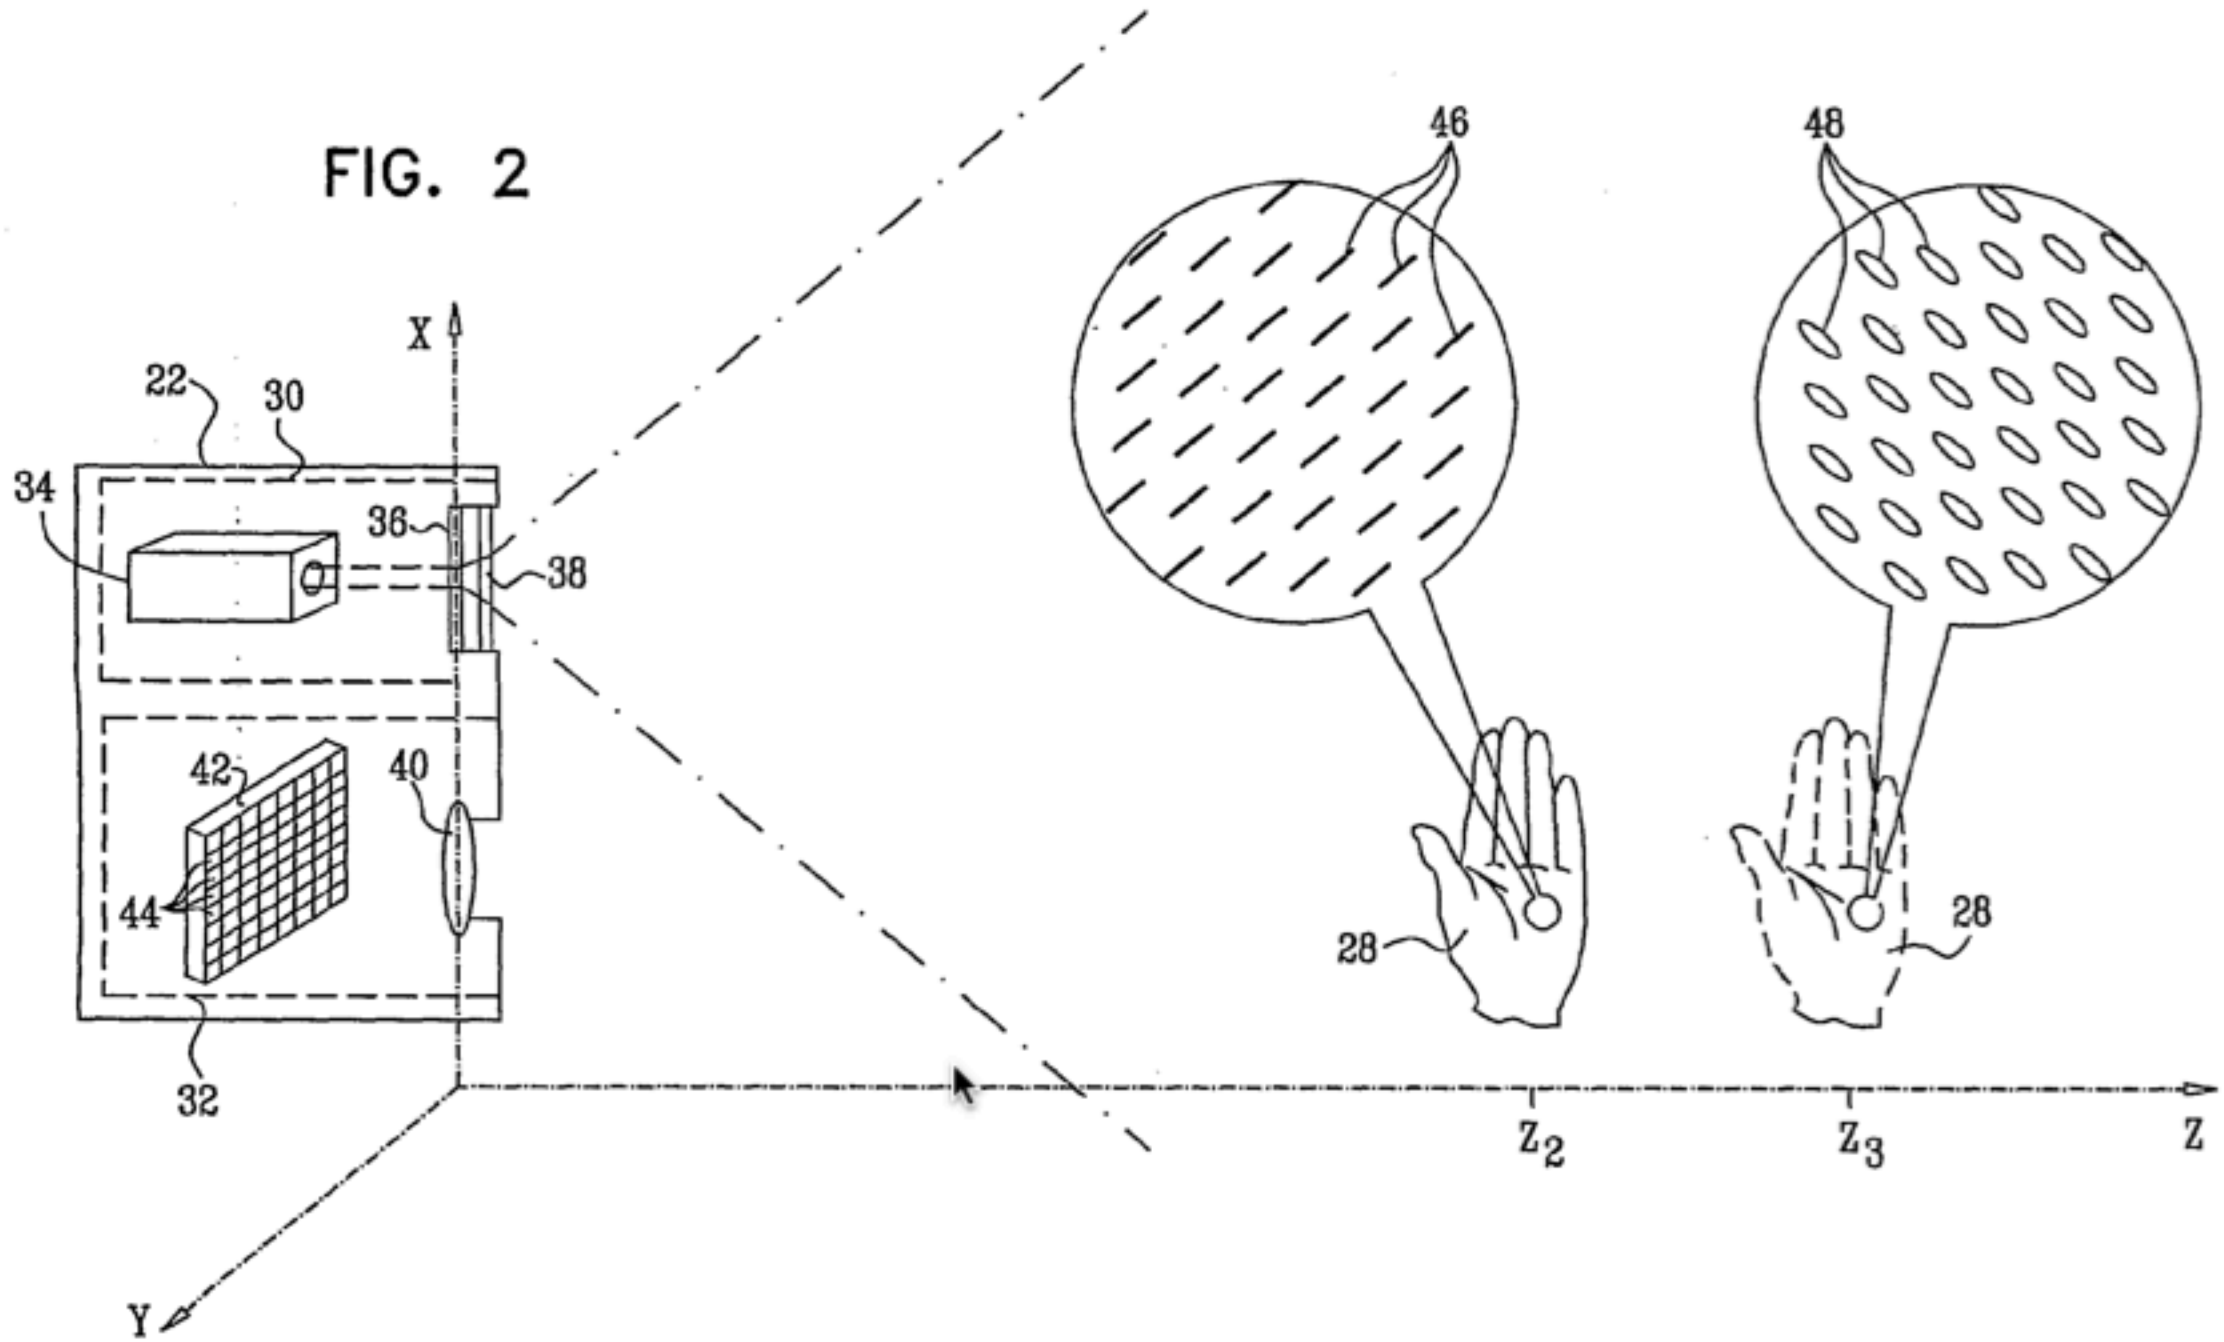

Patent: US 2010/0290698 A1

How does the Kinect Gesture Tracking Work

How does the Kinect Gesture Tracking Work

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

How does the Kinect Gesture Tracking Work

- Depth from stereo

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

How does the Kinect Gesture Tracking Work

- Depth from stereo
  - Based on parallax

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

How does the Kinect Gesture Tracking Work

- Depth from stereo
  - Based on parallax
  - If you look at the scene from two angles, objects that are close get shifted to the side more than objects that are far away

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

How does the Kinect Gesture Tracking Work

- Depth from stereo
  - Based on parallax
  - If you look at the scene from two angles, objects that are close get shifted to the side more than objects that are far away
  - The Kinect analyzes the shift of the speckle pattern by projecting from one location and observing from another

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

How does the Kinect Gesture Tracking Work

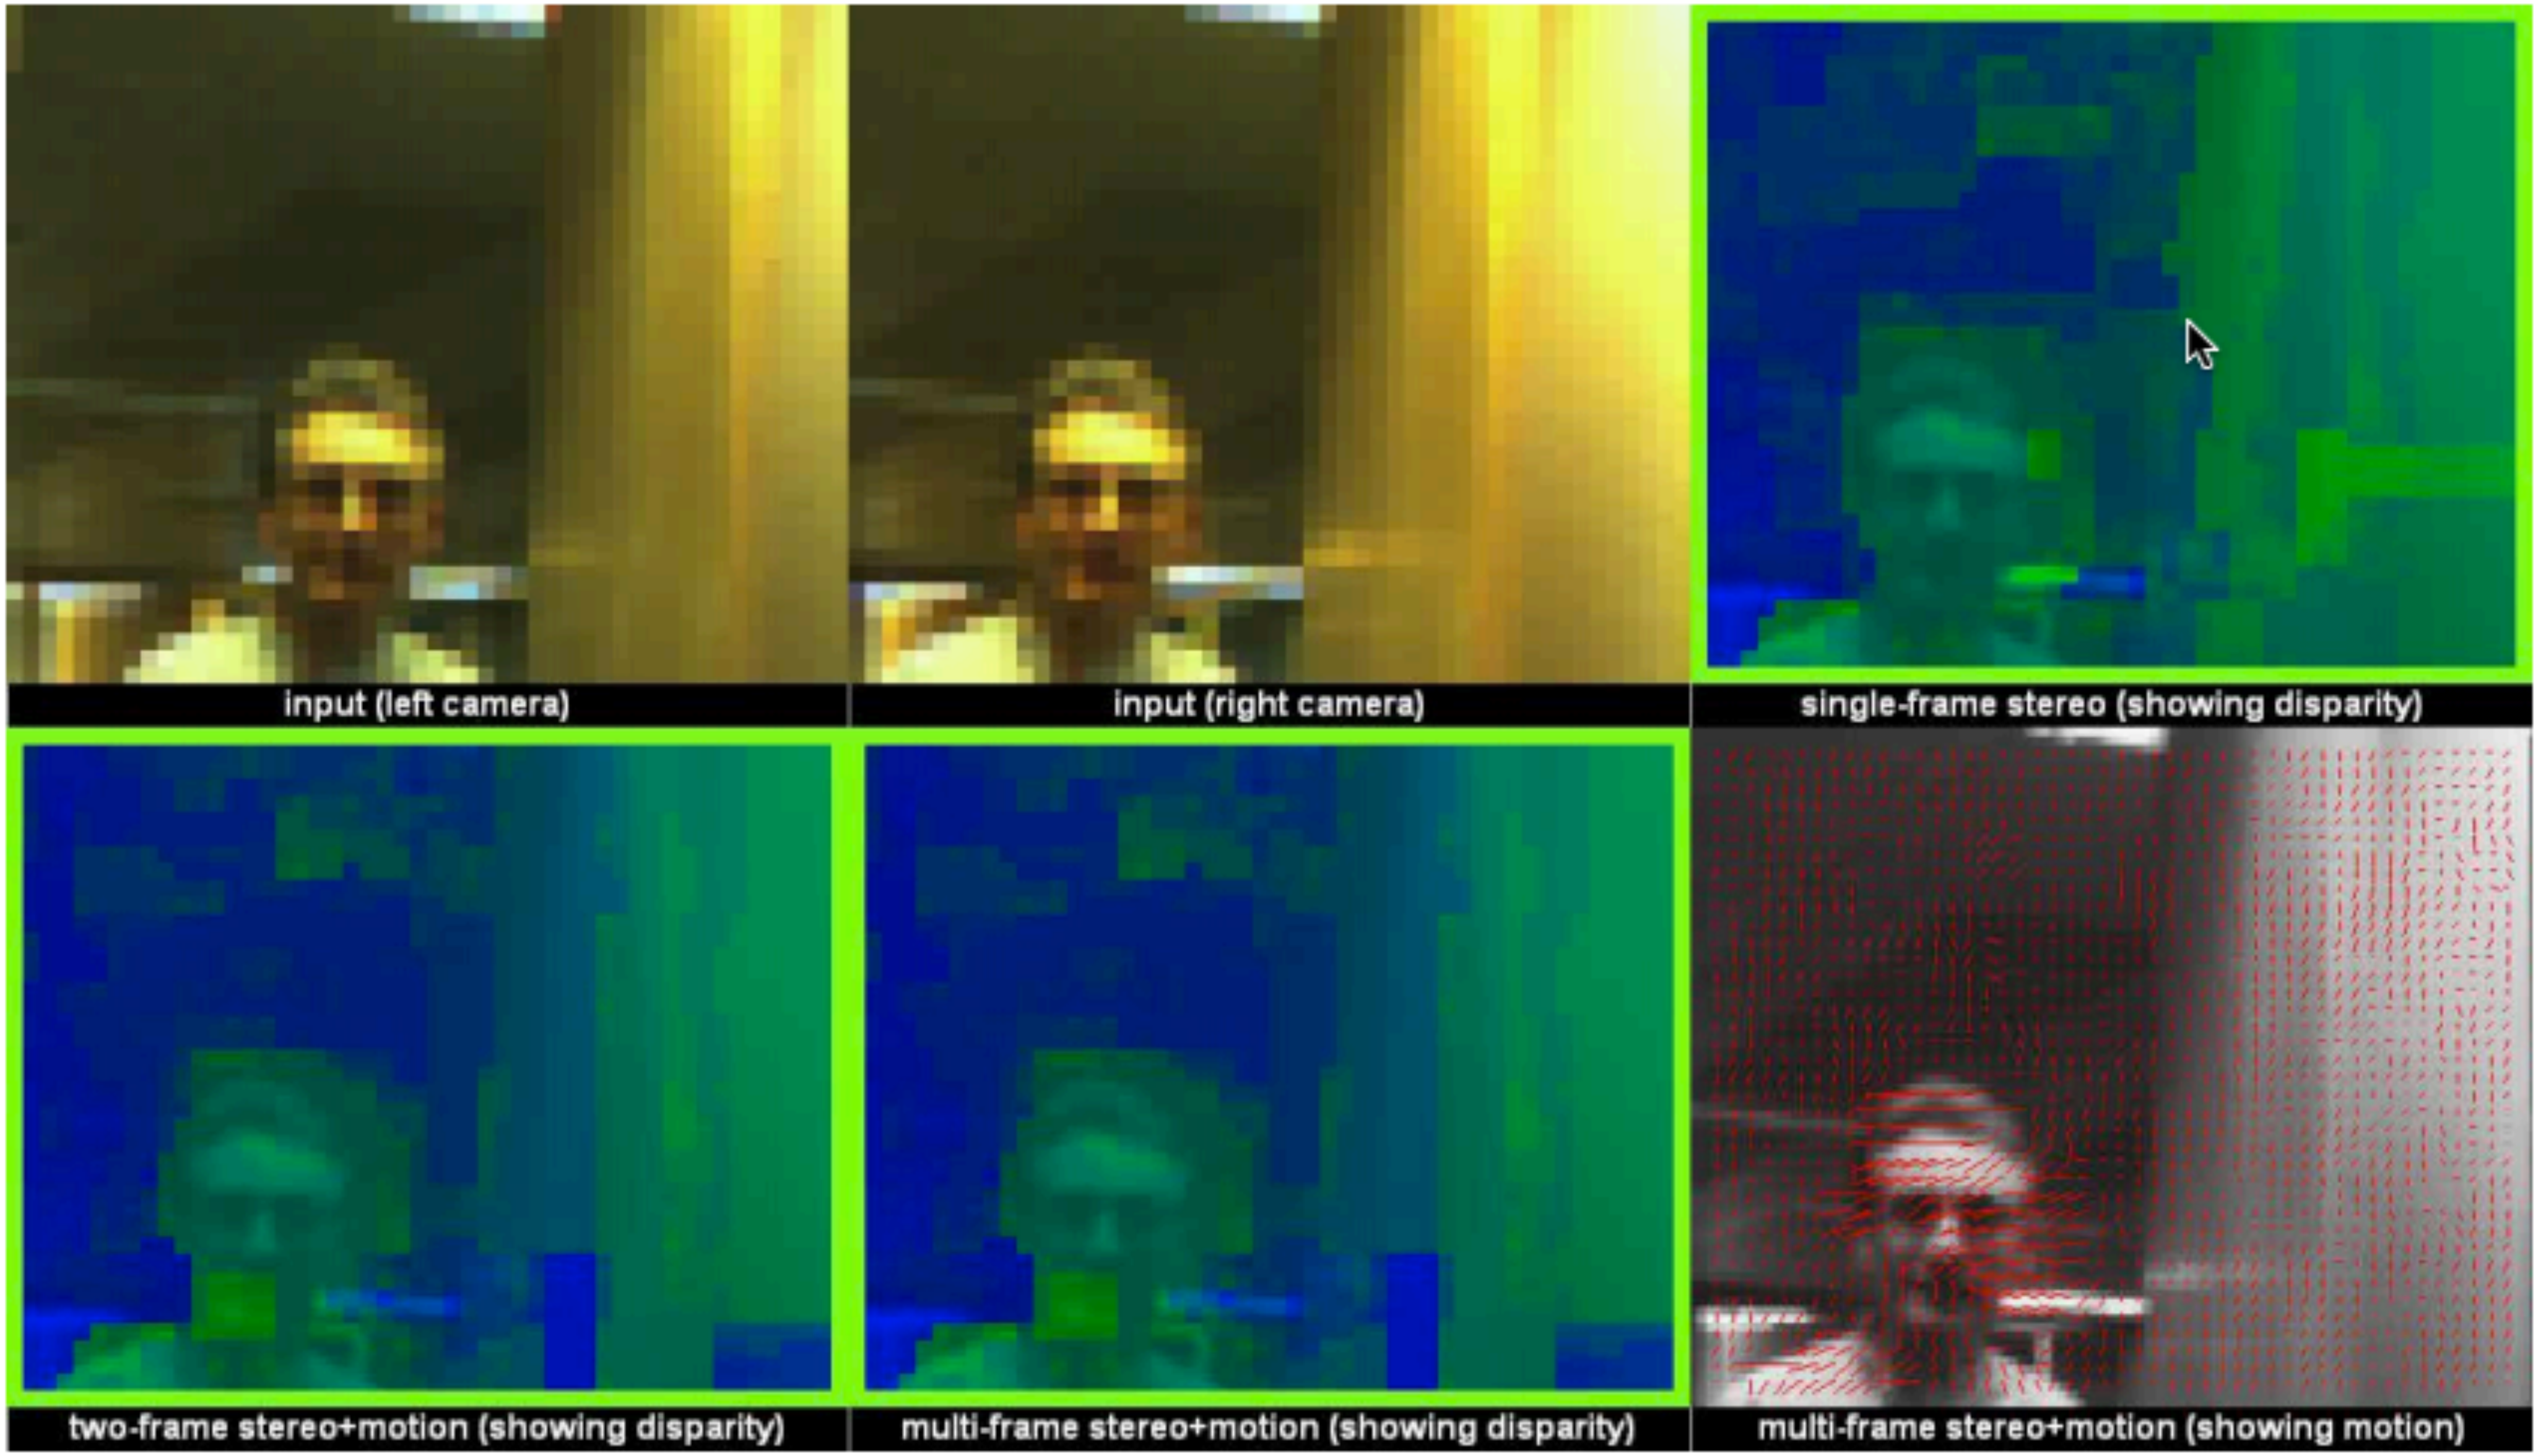

Isard and M, ACCV (2006)

How does the Kinect Gesture Tracking Work

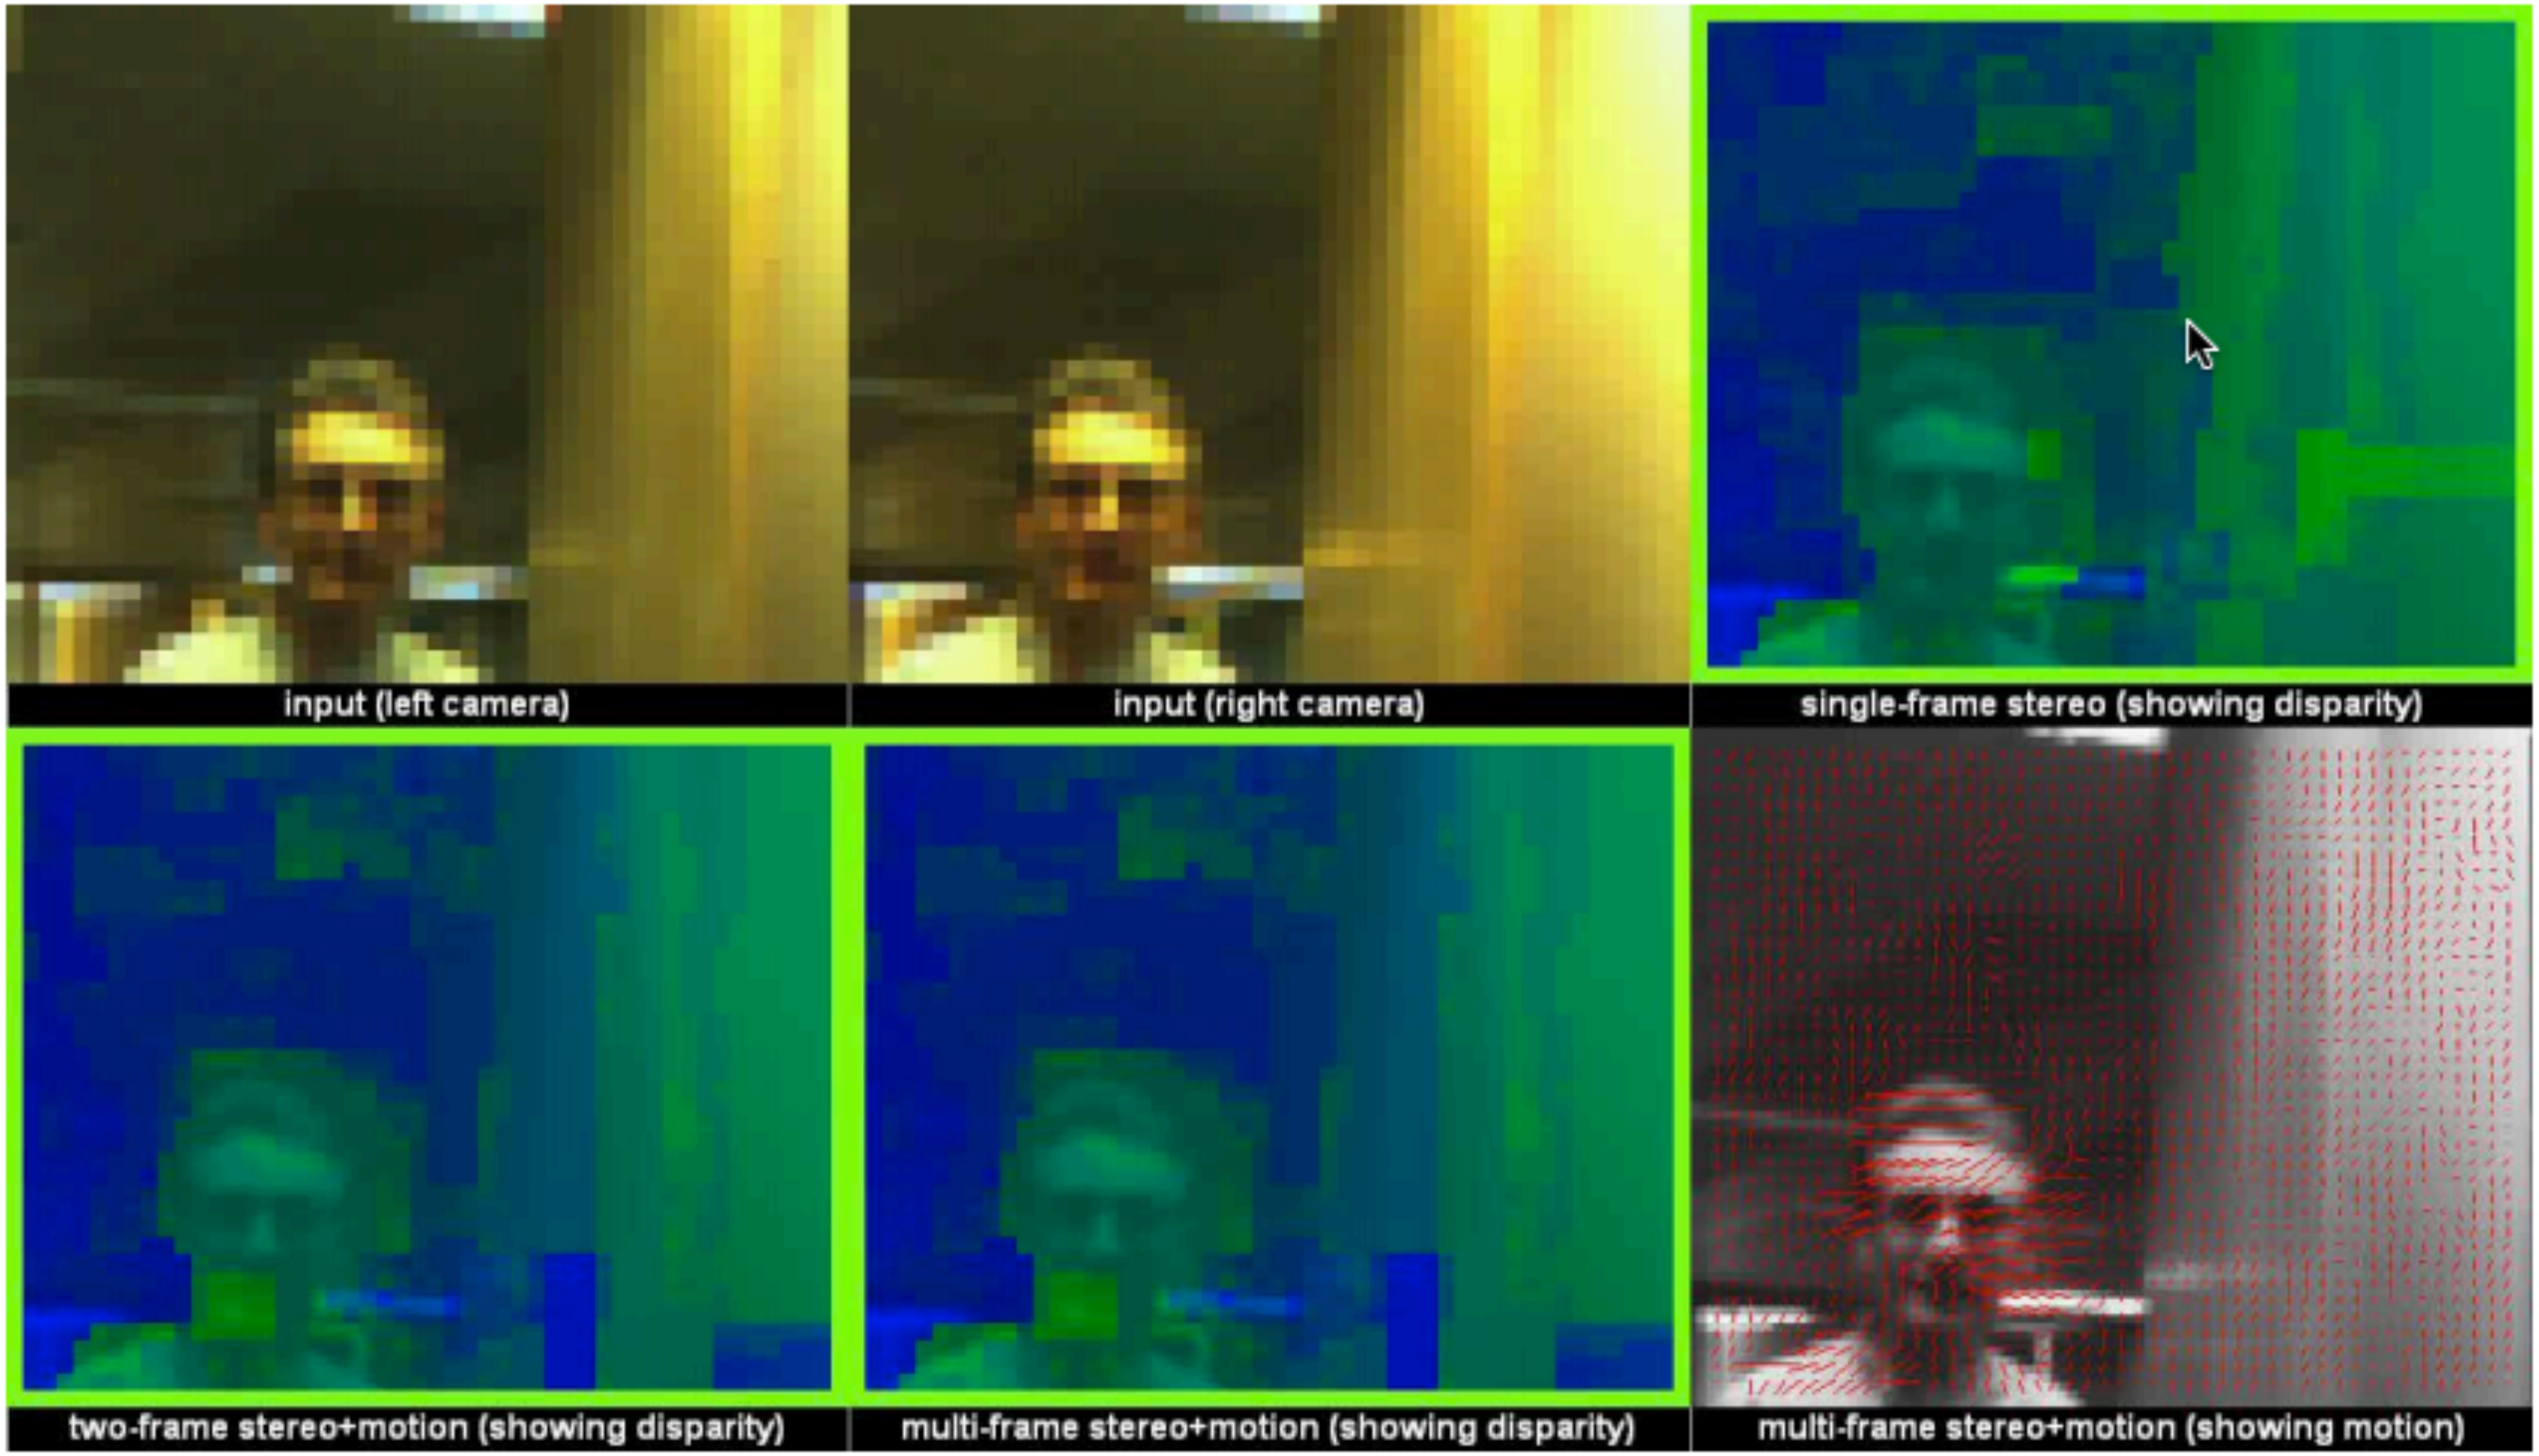

Isard and M, ACCV (2006)

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

How does the Kinect Gesture Tracking Work

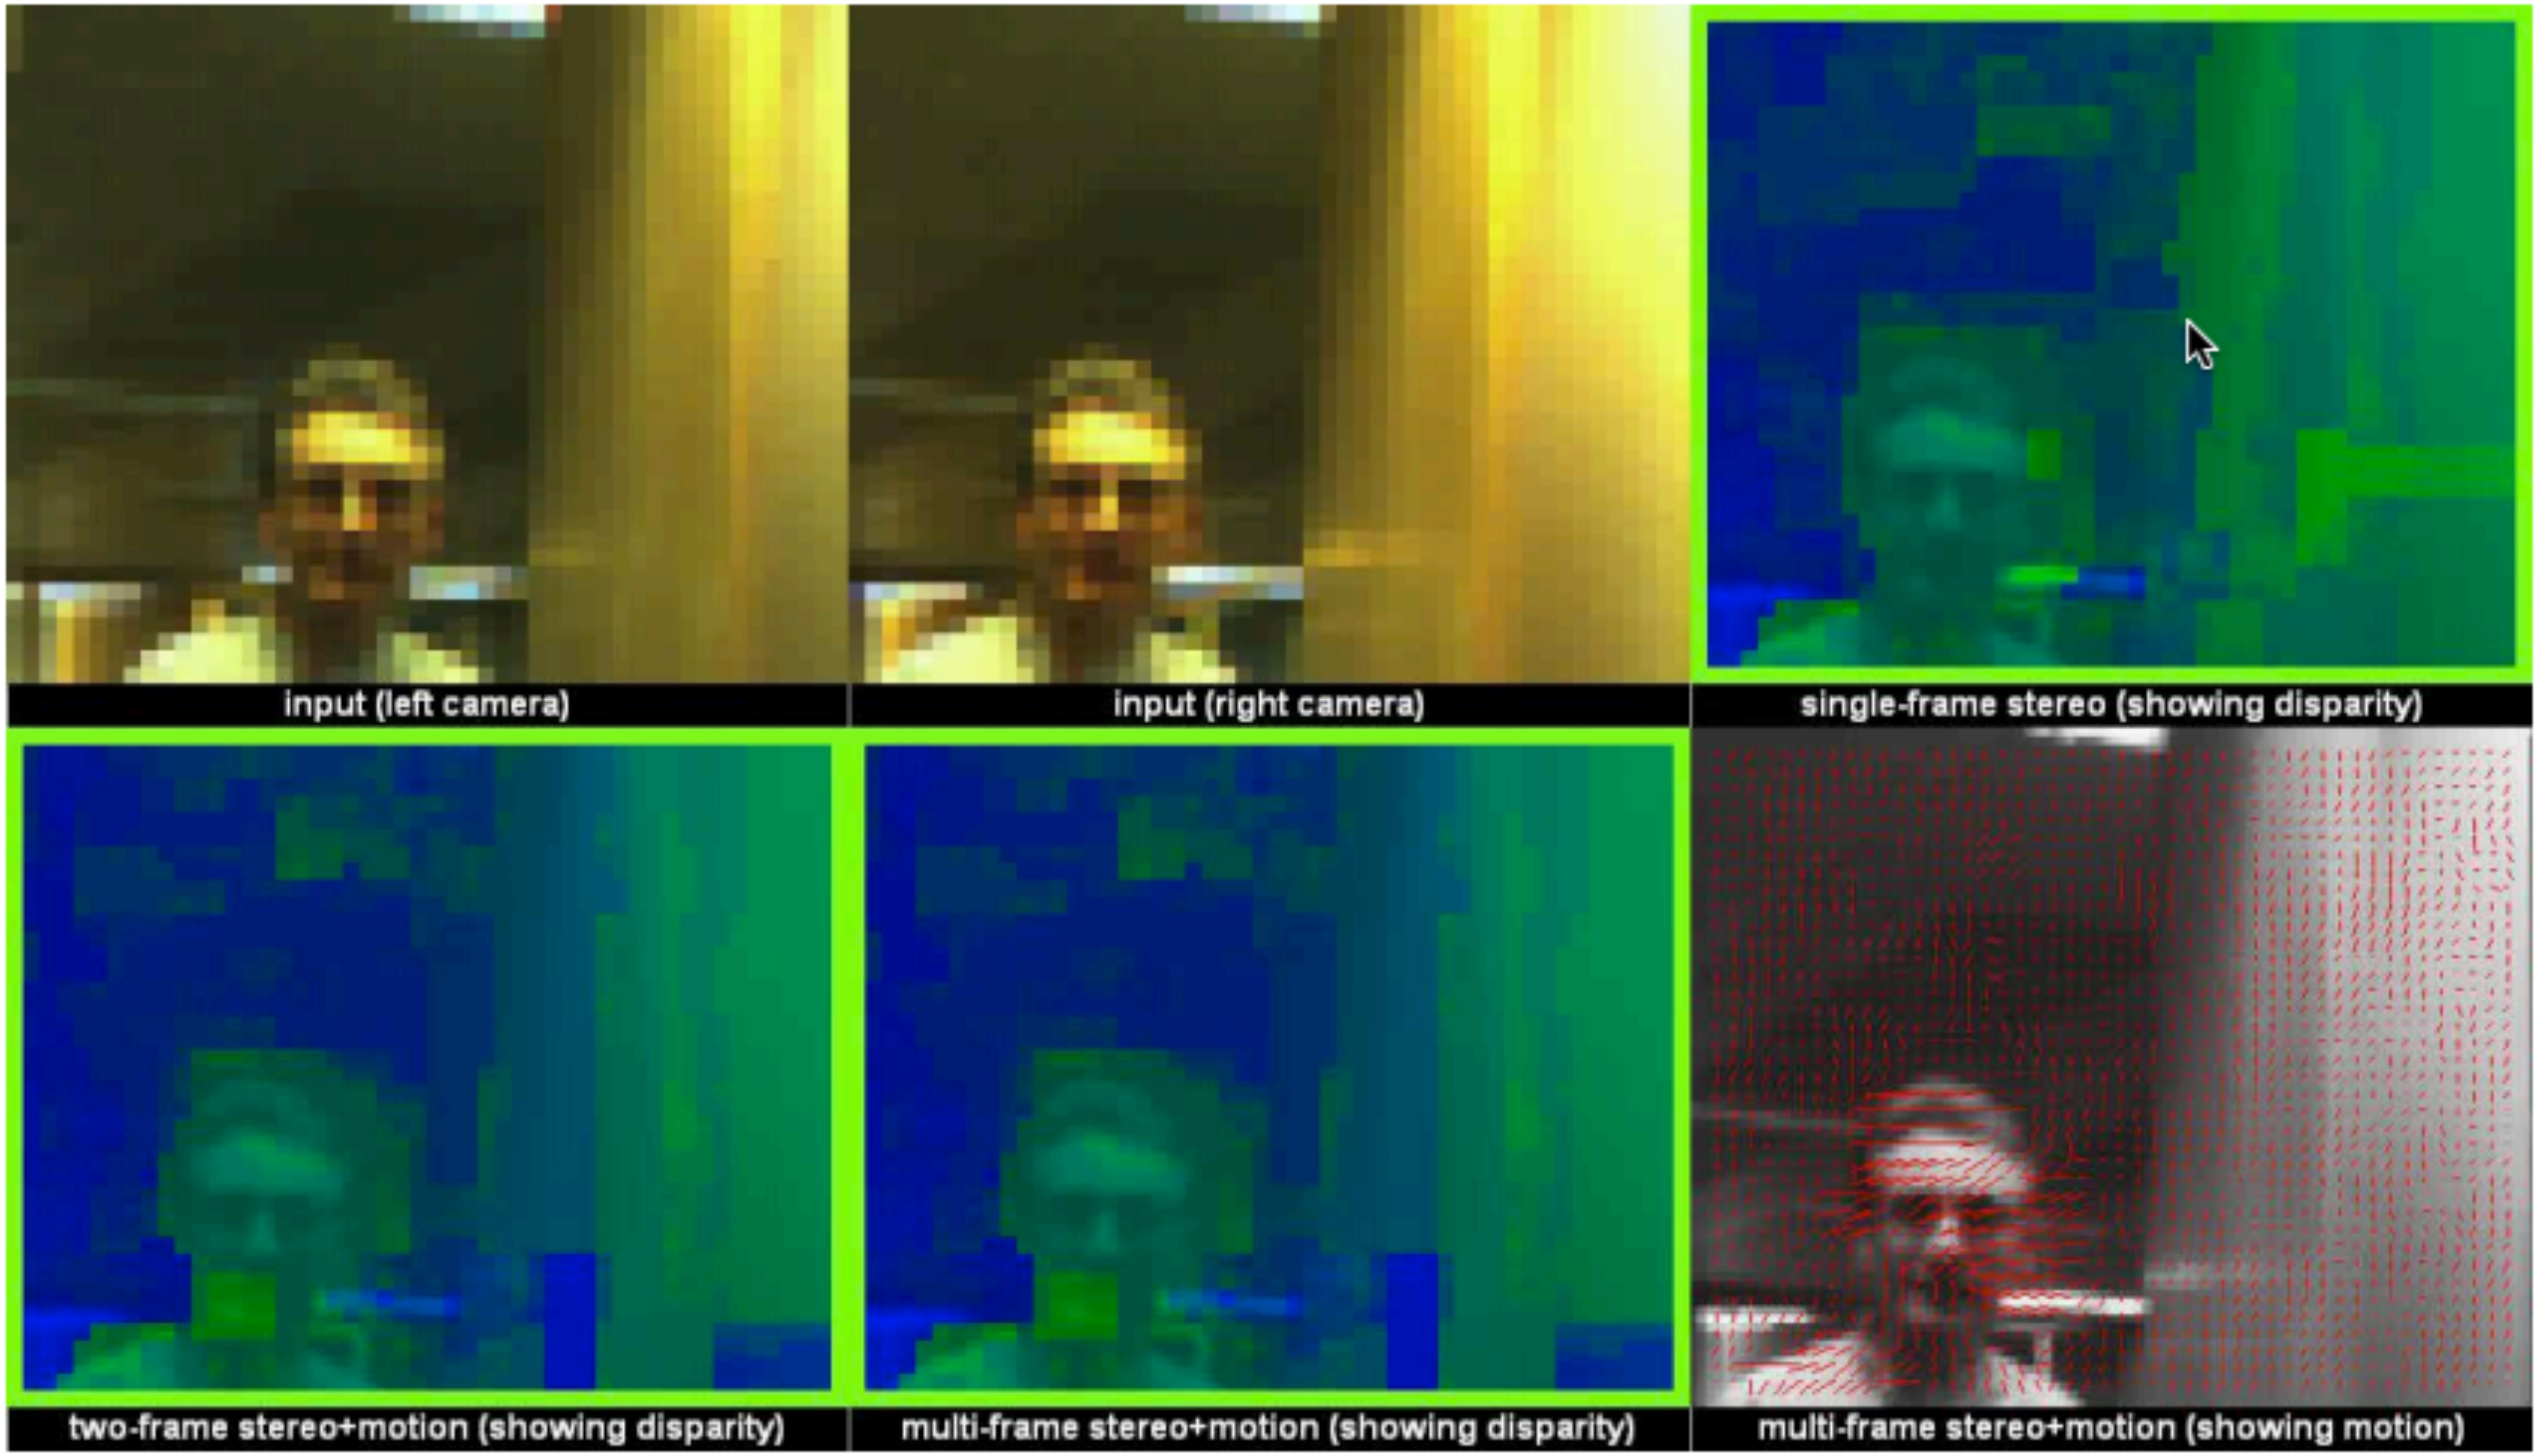

Isard and M, ACCV (2006)

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

How does the Kinect Gesture Tracking Work

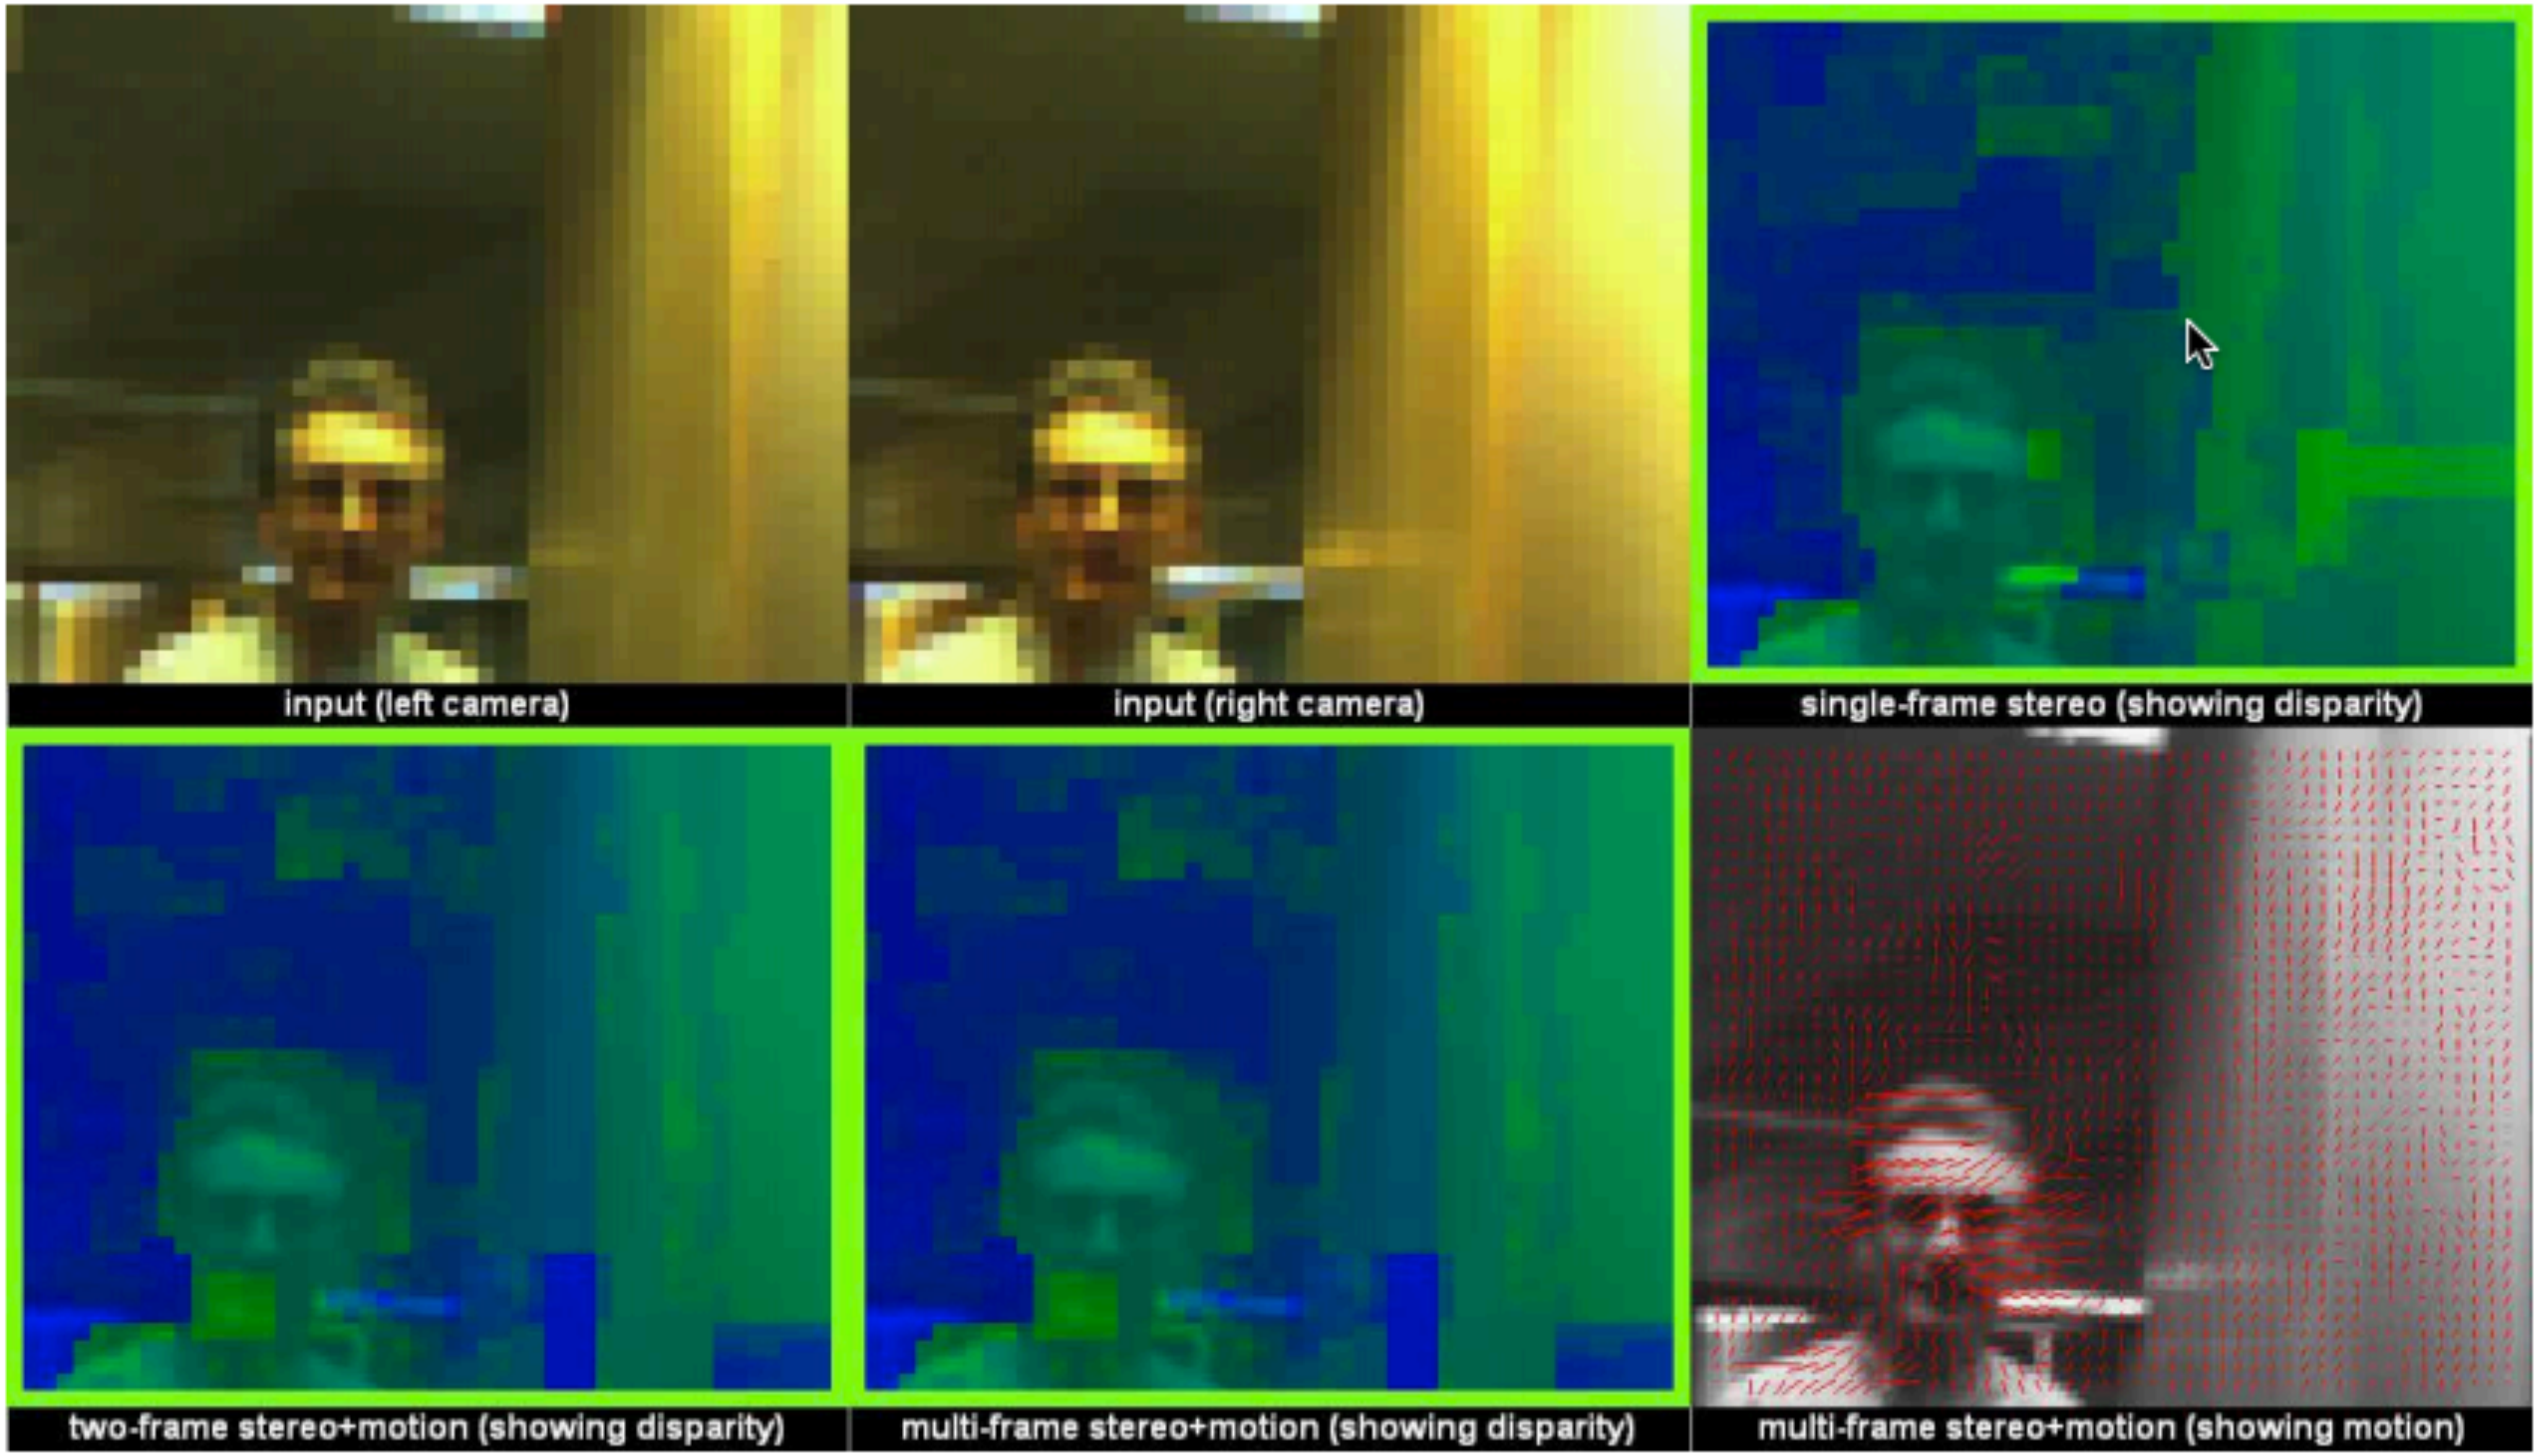

Isard and M, ACCV (2006)

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

How does the Kinect Gesture Tracking Work

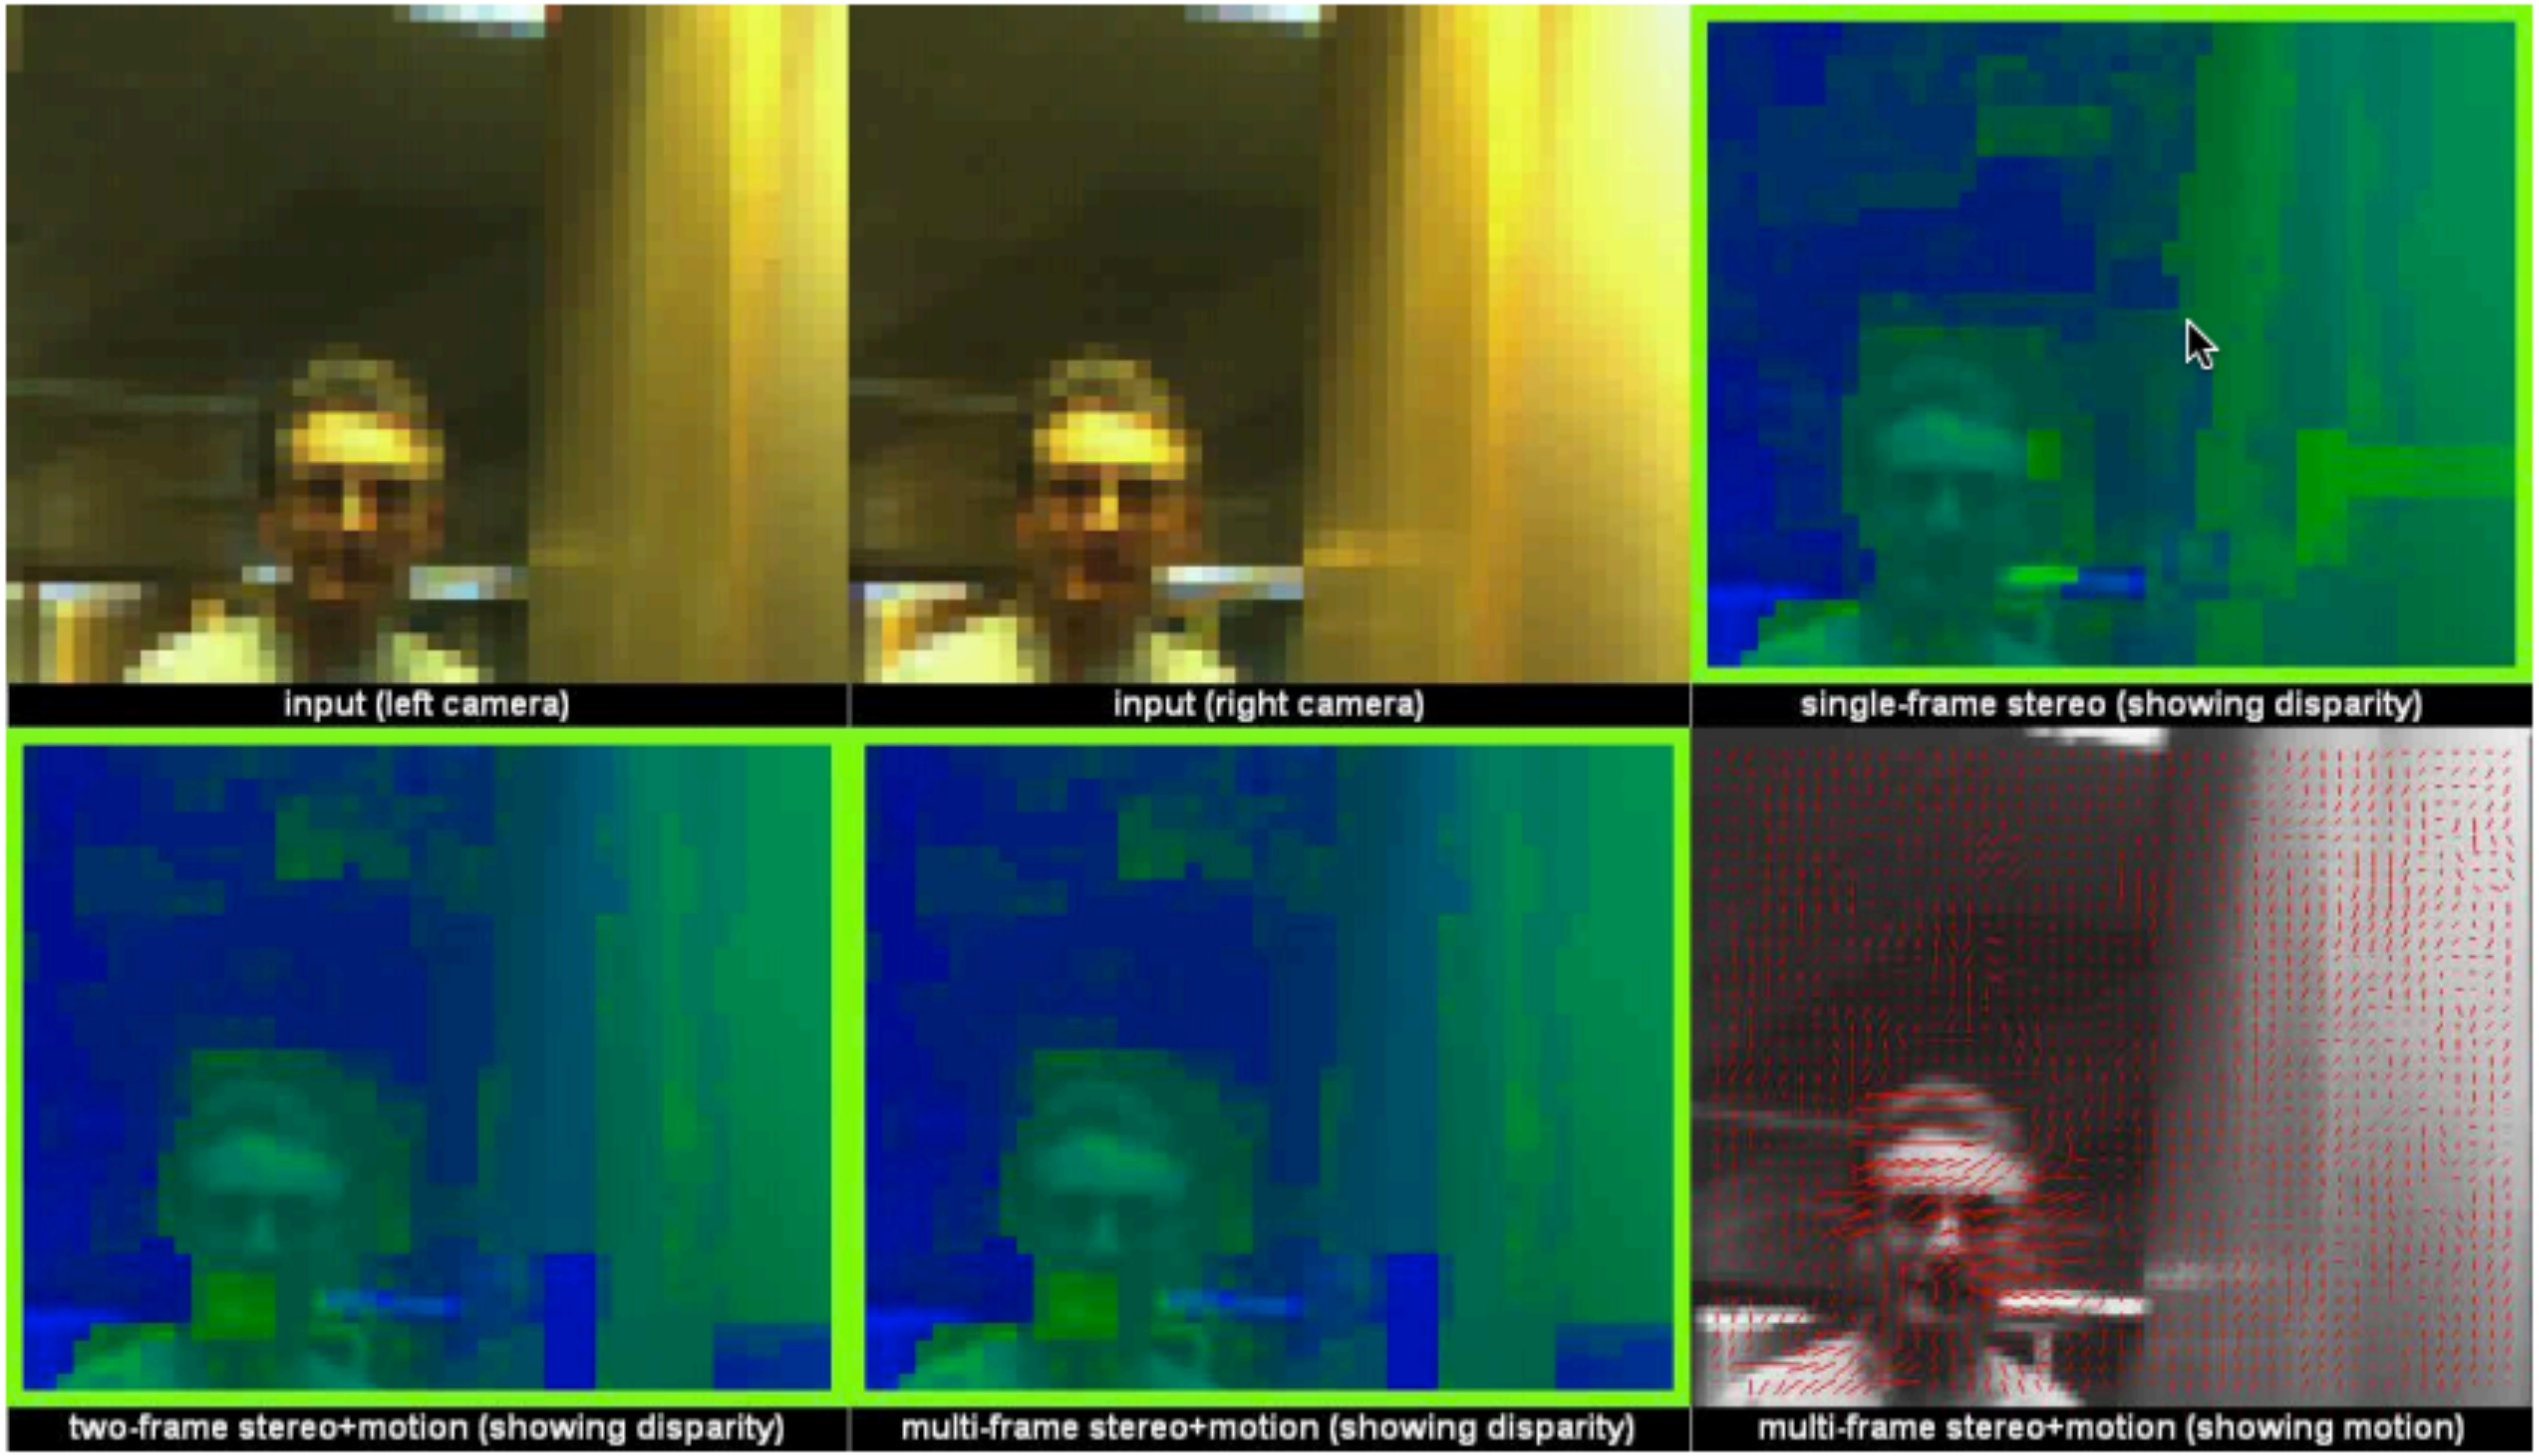

Isard and M, ACCV (2006)

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

How does the Kinect Gesture Tracking Work

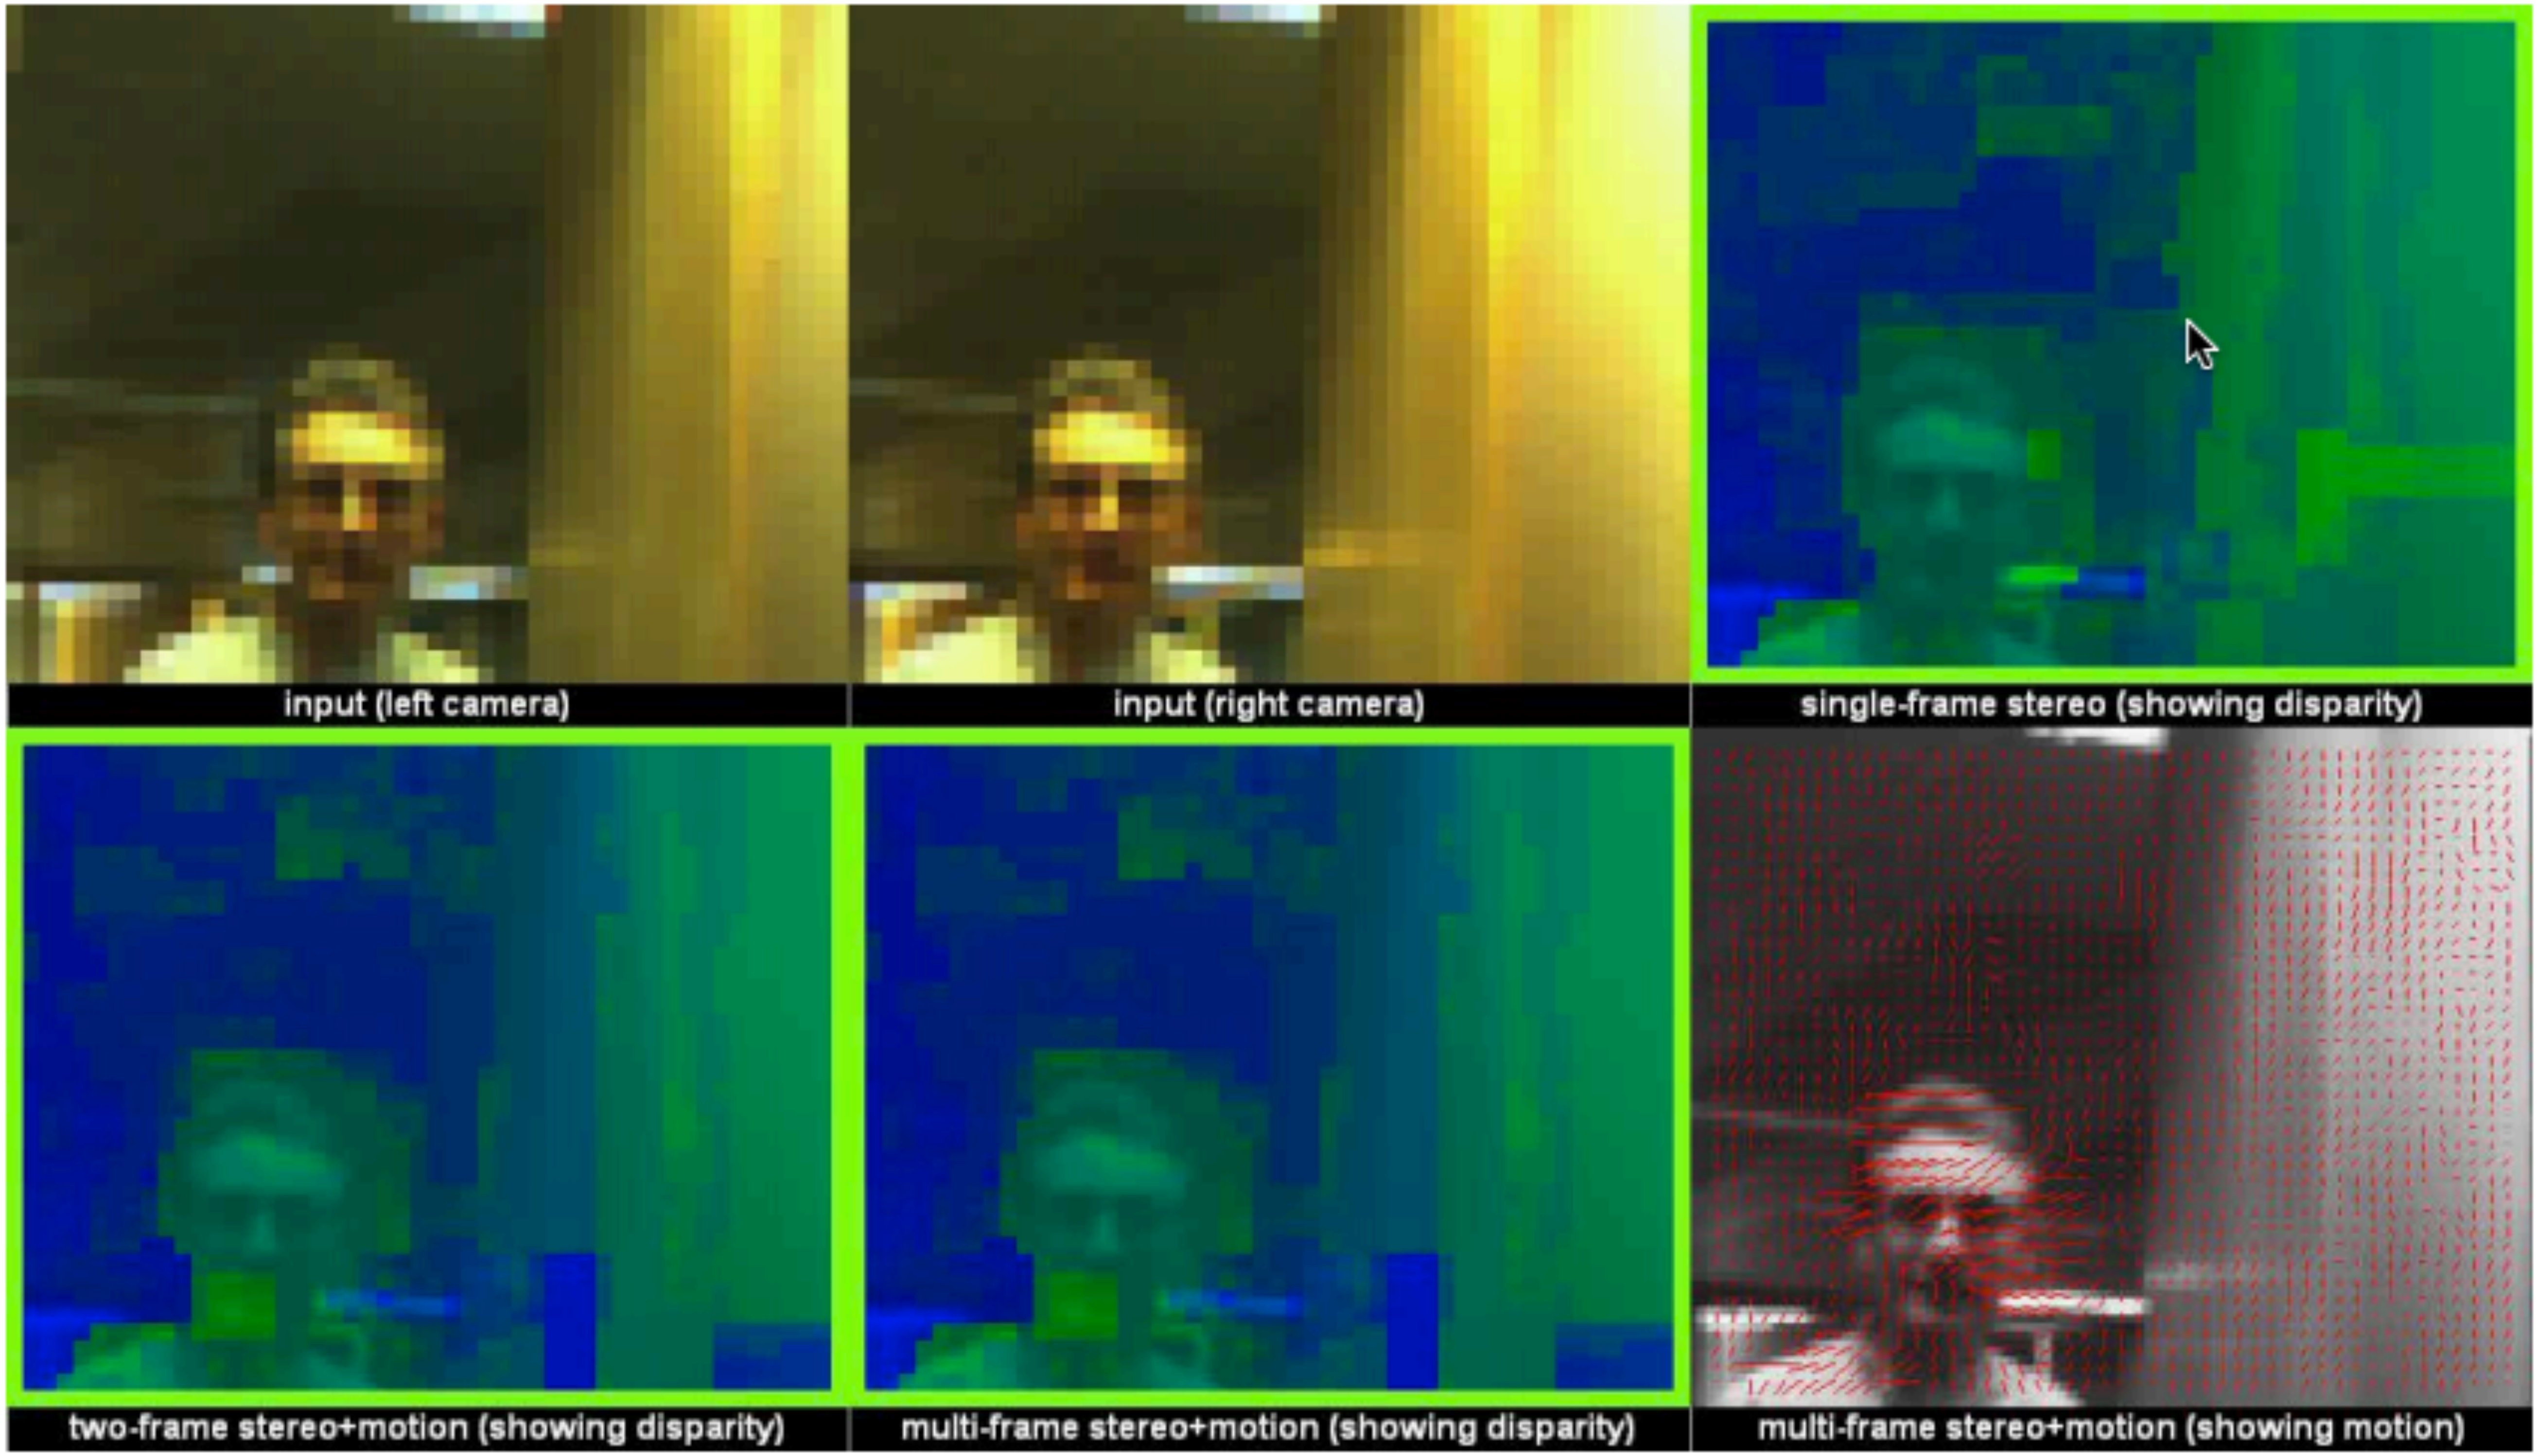

Isard and M, ACCV (2006)

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

How does the Kinect Gesture Tracking Work

How does the Kinect Gesture Tracking Work

- Structured light
    - Kinect Style
        - The Kinect uses an infrared projector and sensor
        - It does not use the RGB camera for depth computation
    - Depth from focus
    - Depth from parallax
    - This creates a depth map

How does the Kinect Gesture Tracking Work

How does the Kinect Gesture Tracking Work

- Demo
    - Explain “point cloud”

How does the Kinect Gesture Tracking Work

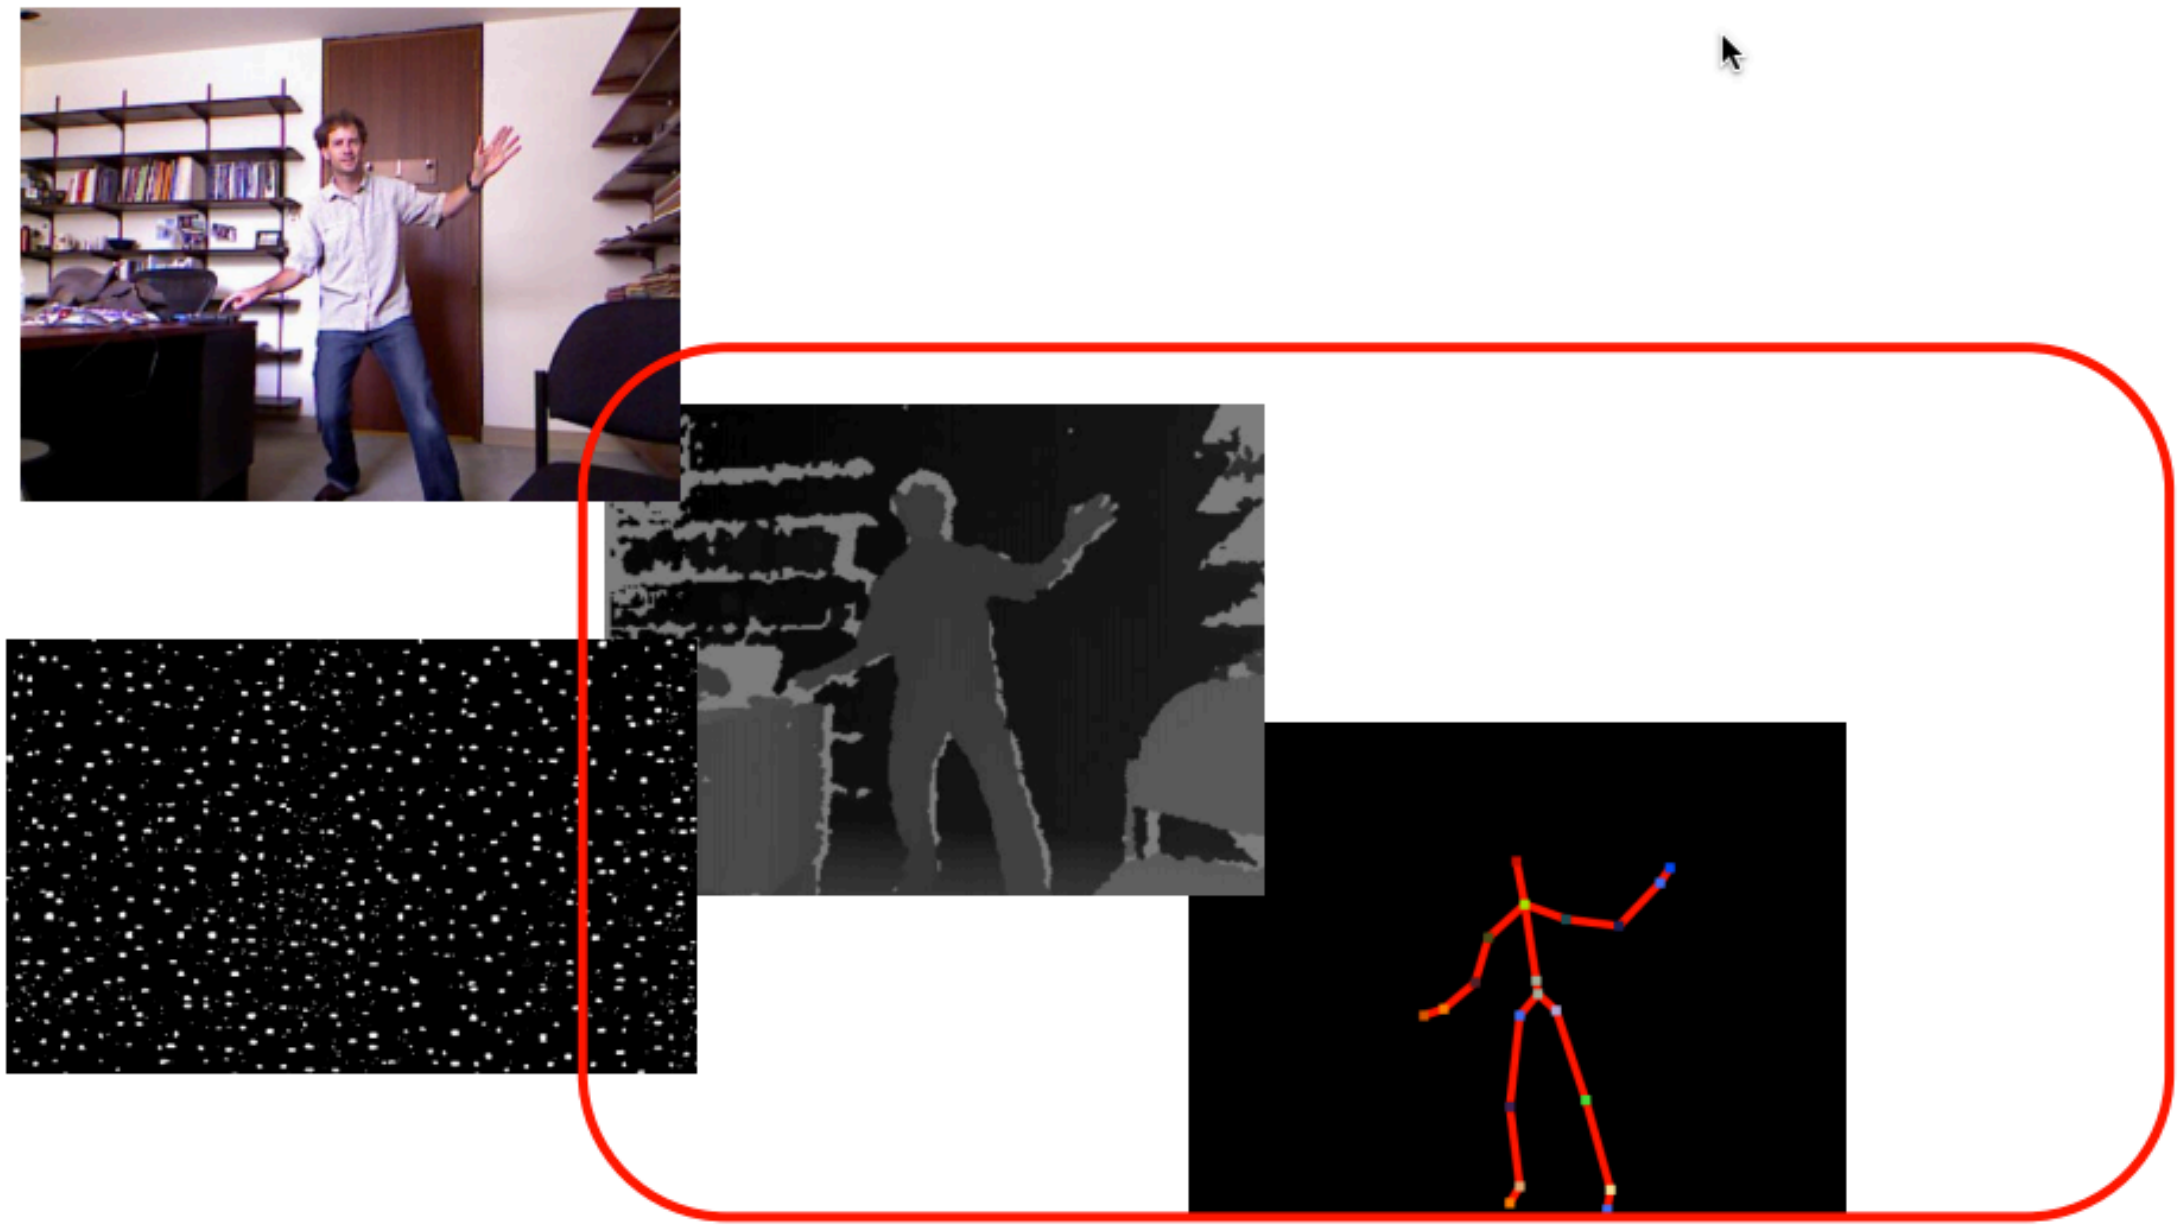

How does the Kinect Gesture Tracking Work

- Inferring body position is a two-stage process
  - First, Depth-map
  - Then Body Position

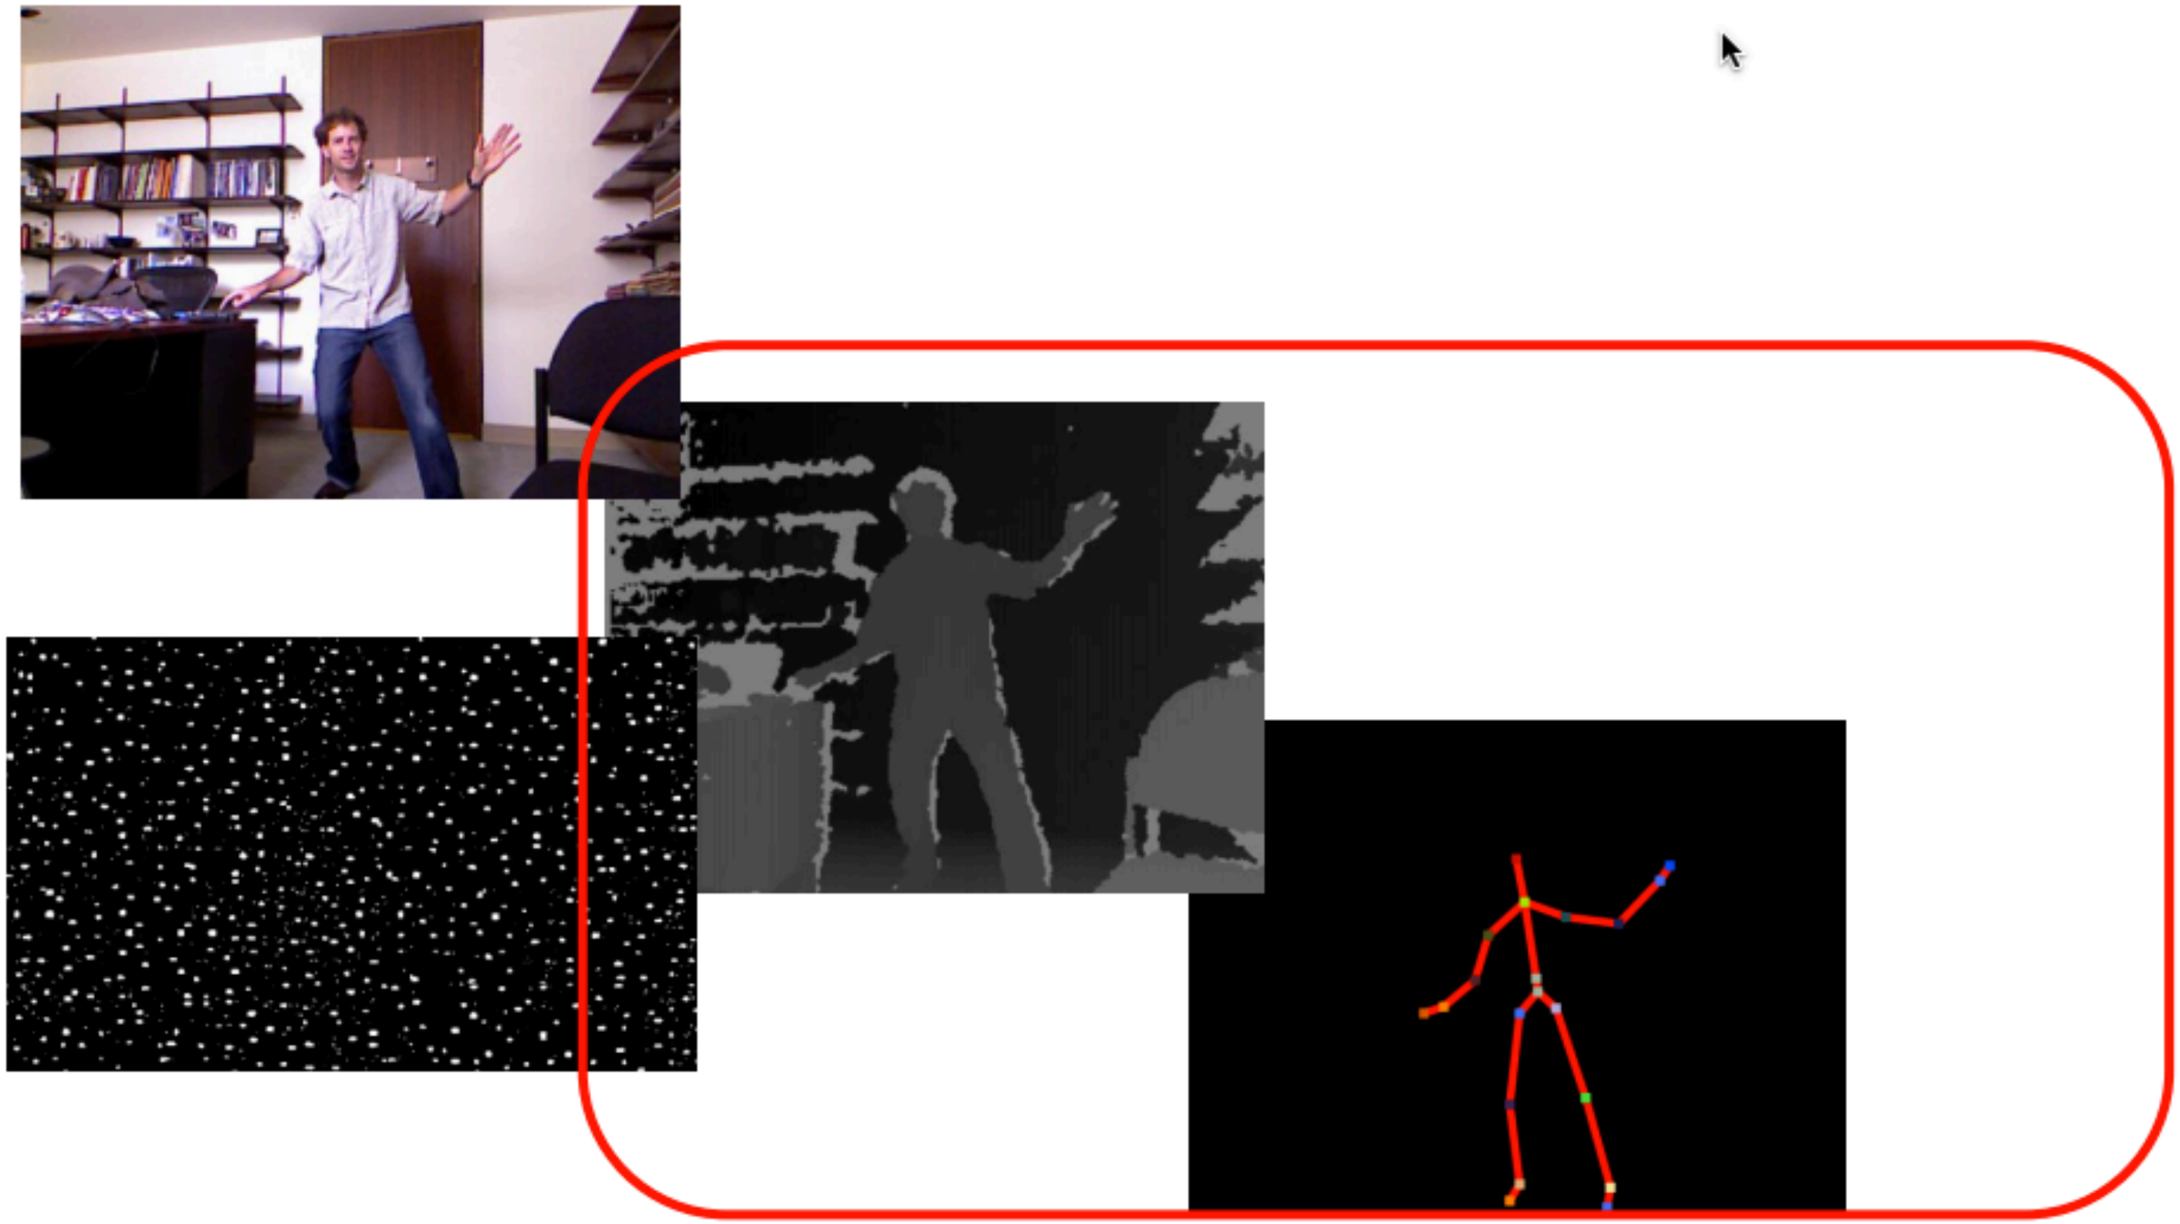

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

How does the Kinect Gesture Tracking Work

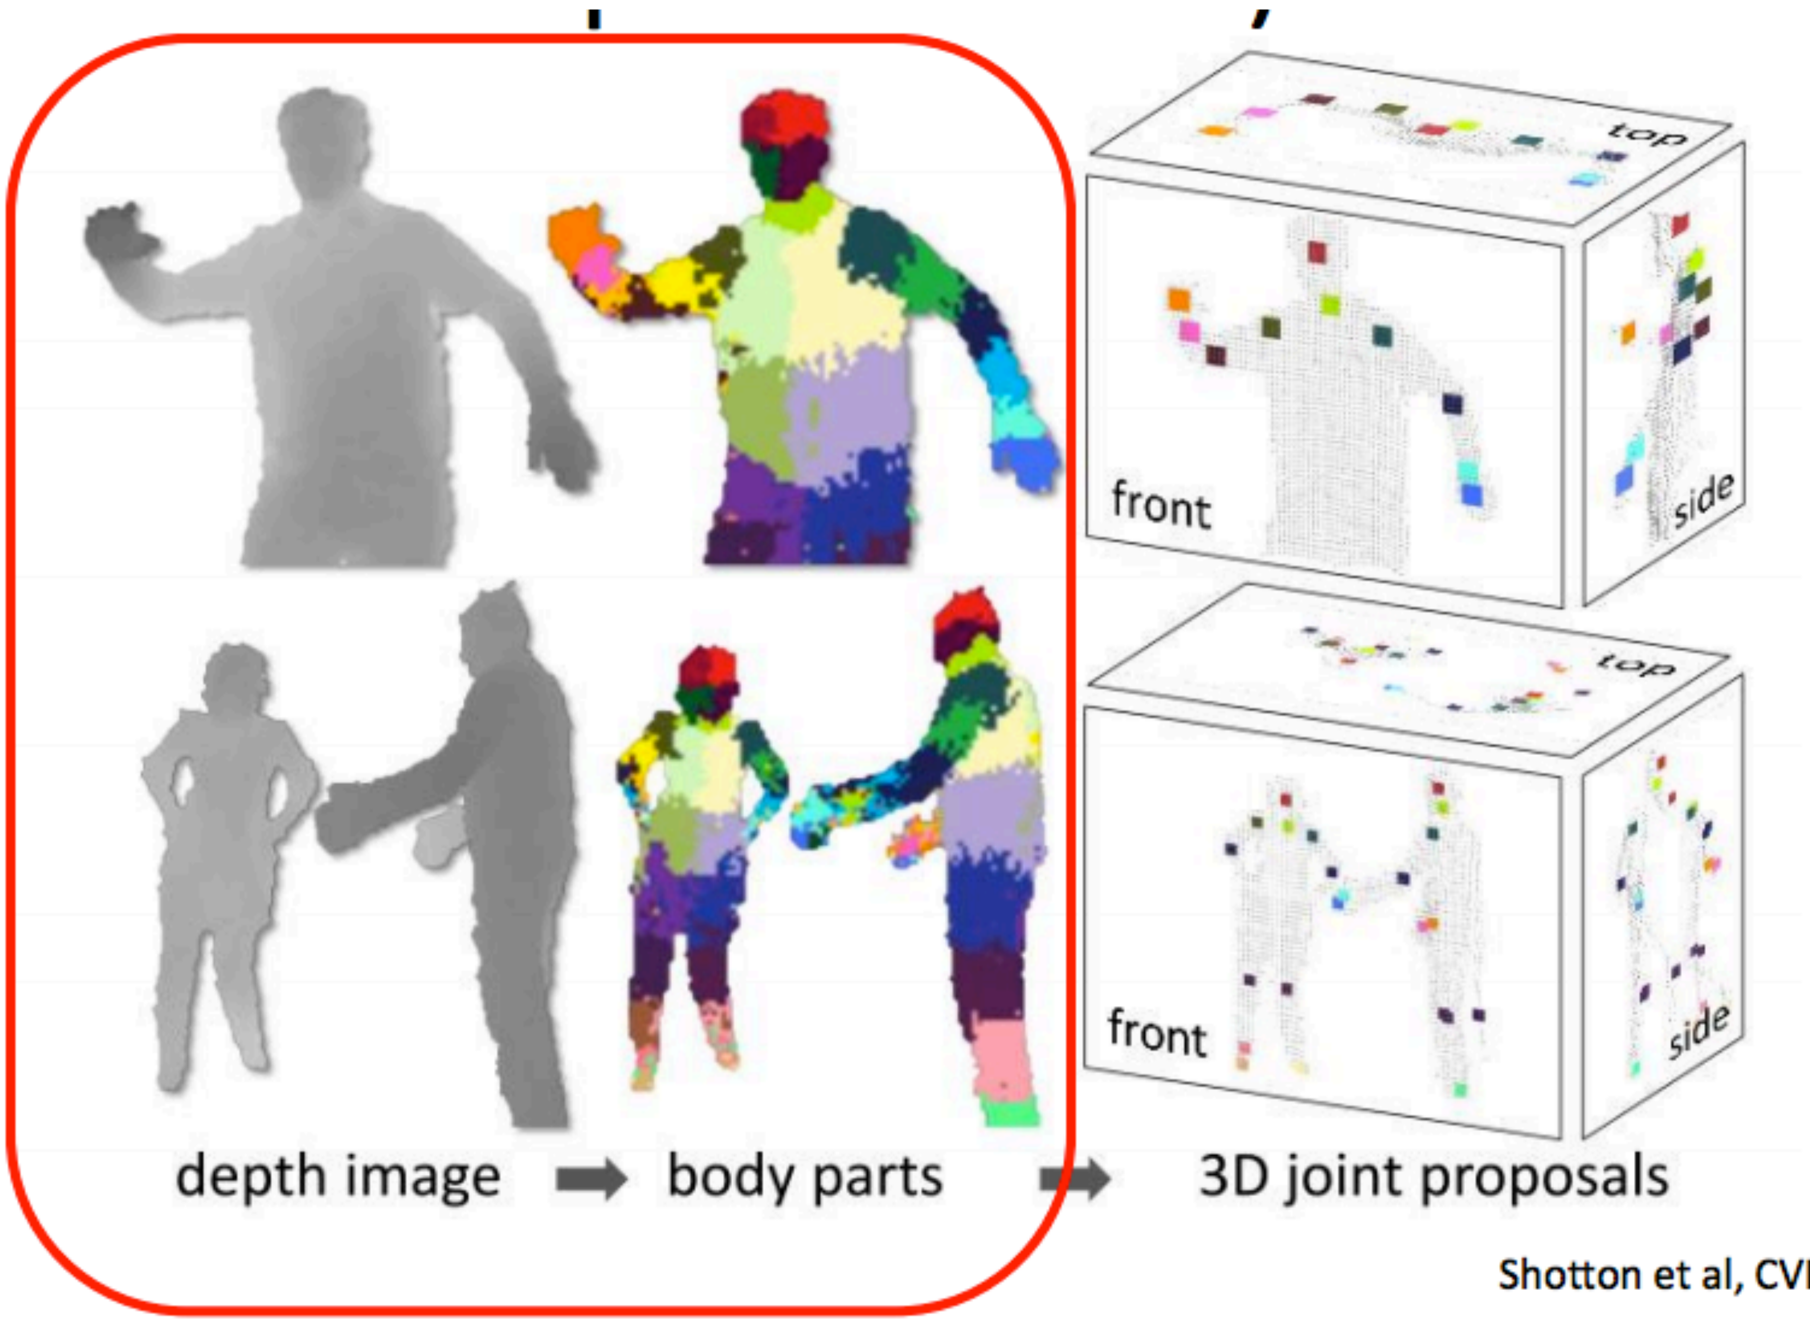

Shotton et al, CVPR(2011)

How does the Kinect Gesture Tracking Work

- Stage 2 has 2 substages
  - Map the depth-map to body-parts
  - map body-parts to 3d Joint proposals

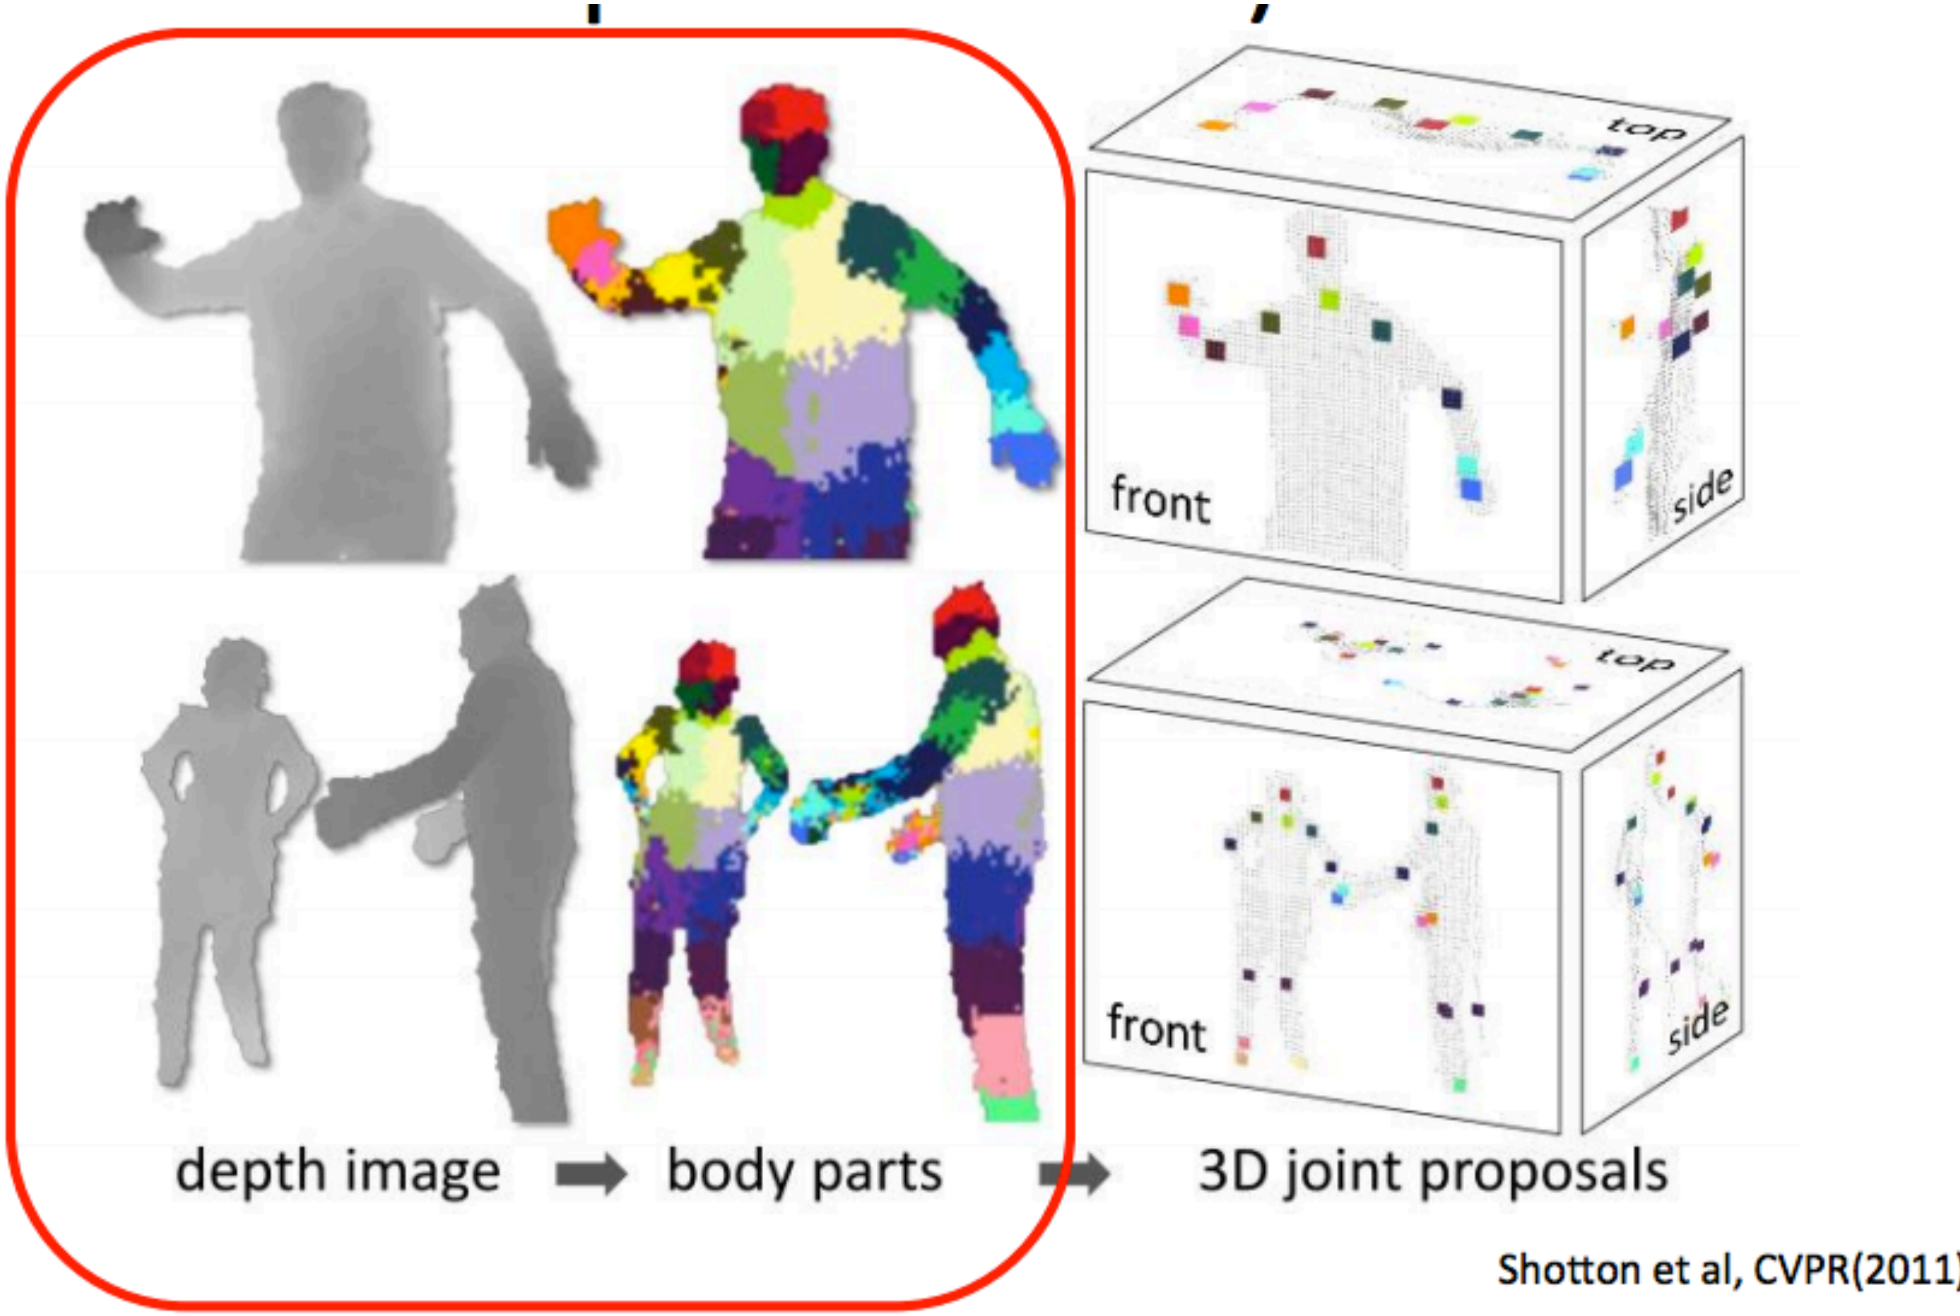

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

How does the Kinect Gesture Tracking Work

Stage 2.1 starts with 100,000 depth images with known skeletons (from a motion capture system)

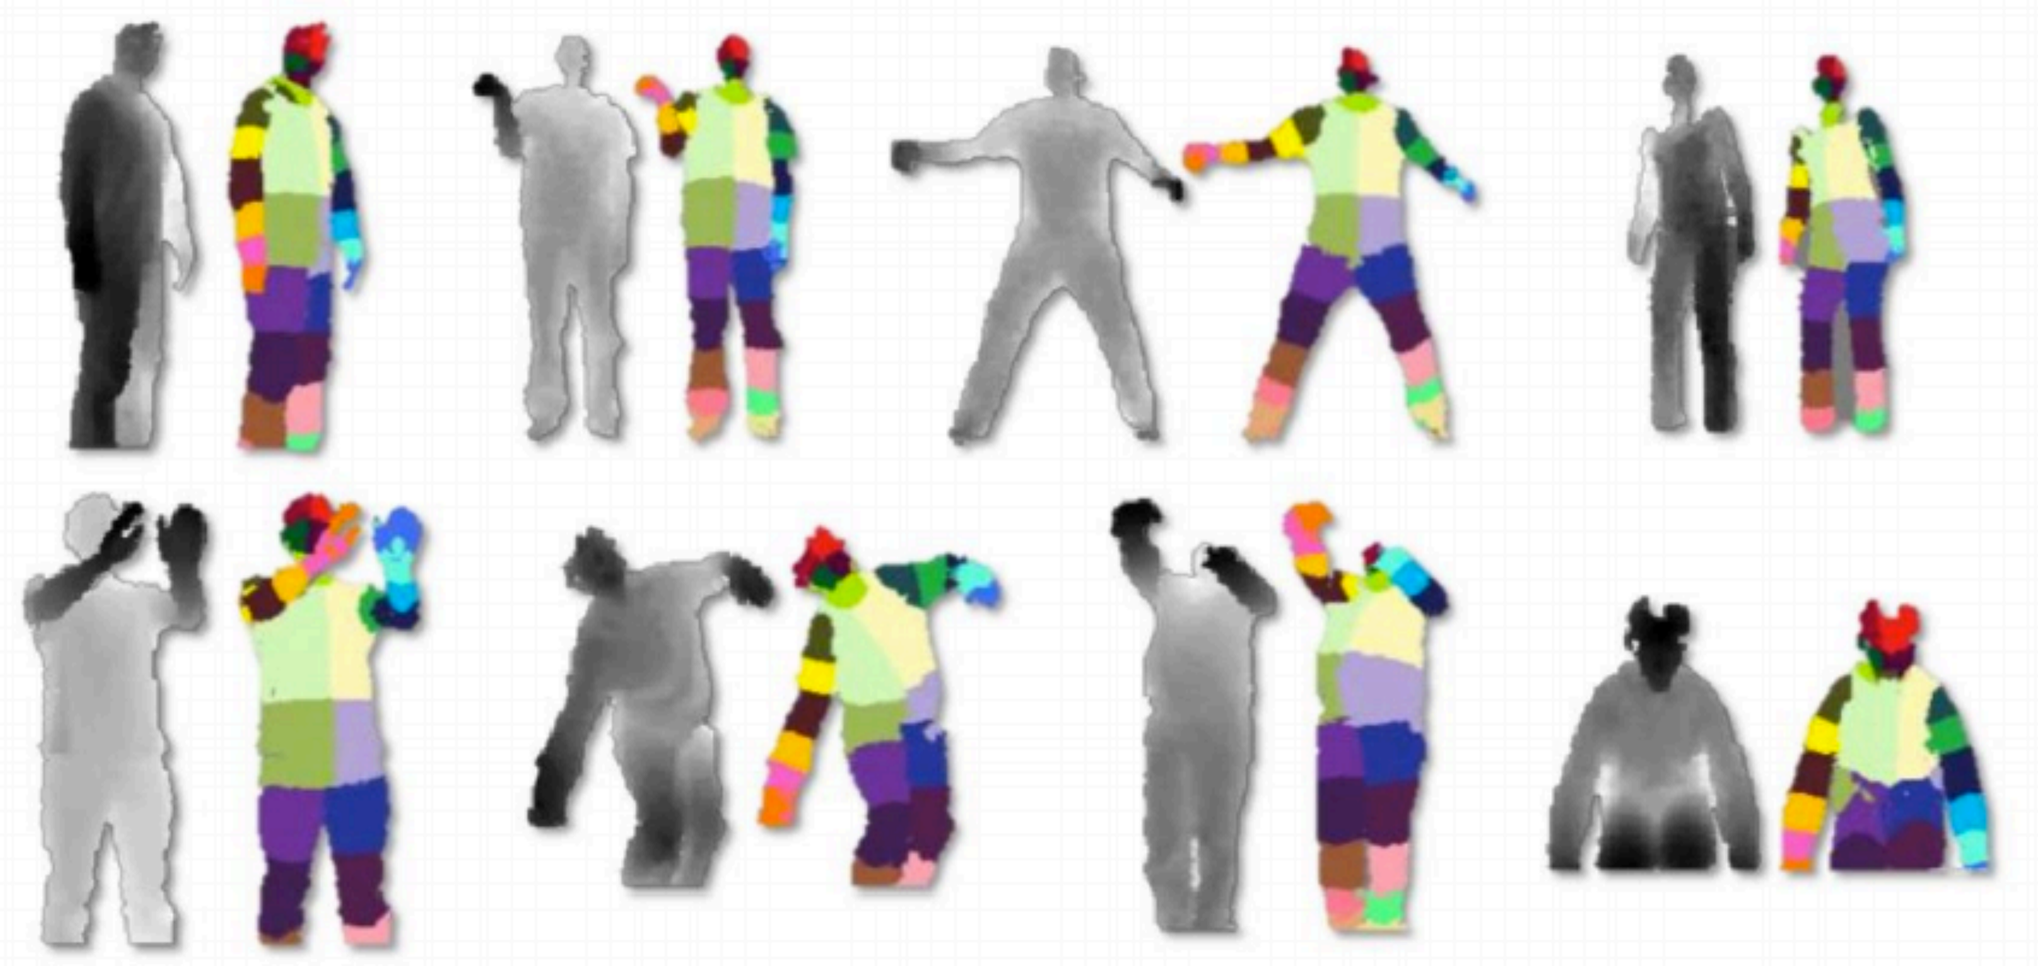

Shotton et al, CVPR(2011)

How does the Kinect Gesture Tracking Work

Stage 2.1 starts with 100,000 depth images with known skeletons (from a motion capture system)

Shotton et al, CVPR(2011)

How does the Kinect Gesture Tracking Work

For each real image, render dozens more using computer graphics techniques

- Use computer graphics to render all sequences for 15 different body types, and vary several other parameters
- Thus obtain over a million training examples

How does the Kinect Gesture Tracking Work

For each real image, render dozens more using computer graphics techniques

- Use computer graphics to render all sequences for 15 different body types, and vary several other parameters
- Thus obtain over a million training examples

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

How does the Kinect Gesture Tracking Work

Shotton et al, CVPR(2011)

How does the Kinect Gesture Tracking Work

Shotton et al, CVPR(2011)

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

Stage 2.1 transforms depth image to body part image

- Start with 100,000 depth images with known skeleton (from a motion capture system)
- For each real image, render dozens more using computer graphics techniques.
- Learn a *randomized decision forest*, mapping depth images to body parts

Stage 2.1 transforms depth image to body part image

- Start with 100,000 depth images with known skeleton (from a motion capture system)
- For each real image, render dozens more using computer graphics techniques.
- Learn a *randomized decision forest*, mapping depth images to body parts

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

How does the Kinect Gesture Tracking Work

A randomized decision forest is a more sophisticated version of the classic *decision tree*

How does the Kinect Gesture Tracking Work

A randomized decision forest is a more sophisticated version of the classic *decision tree*

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

A decision tree is like a pre-planned game of "twenty questions"

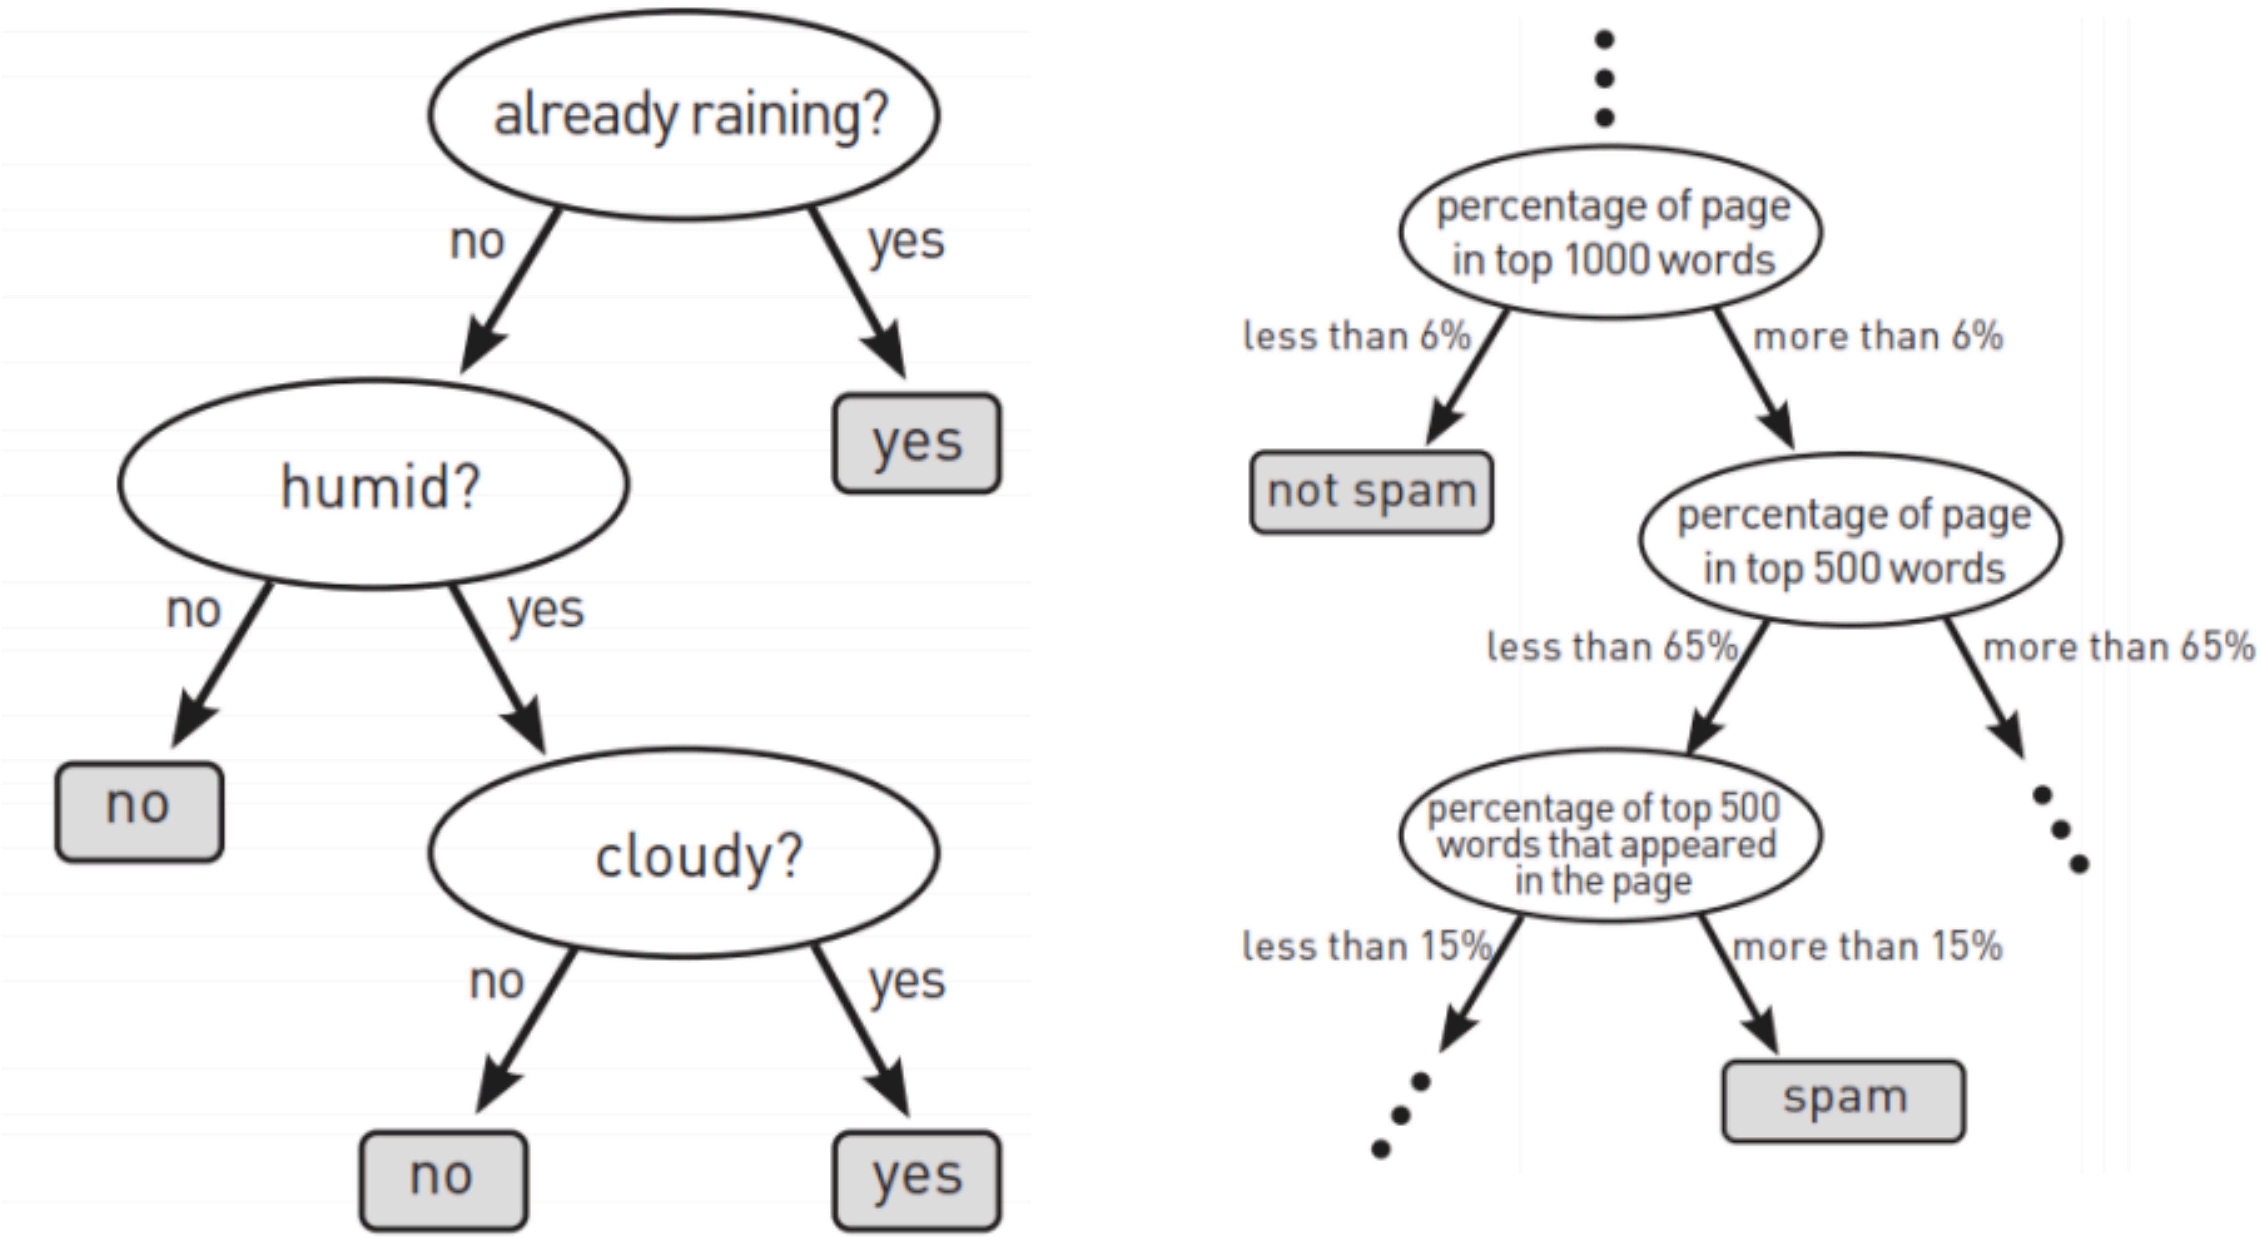

Ntoulas et al, WWW (2006)

A decision tree is like a pre-planned game of "twenty questions"

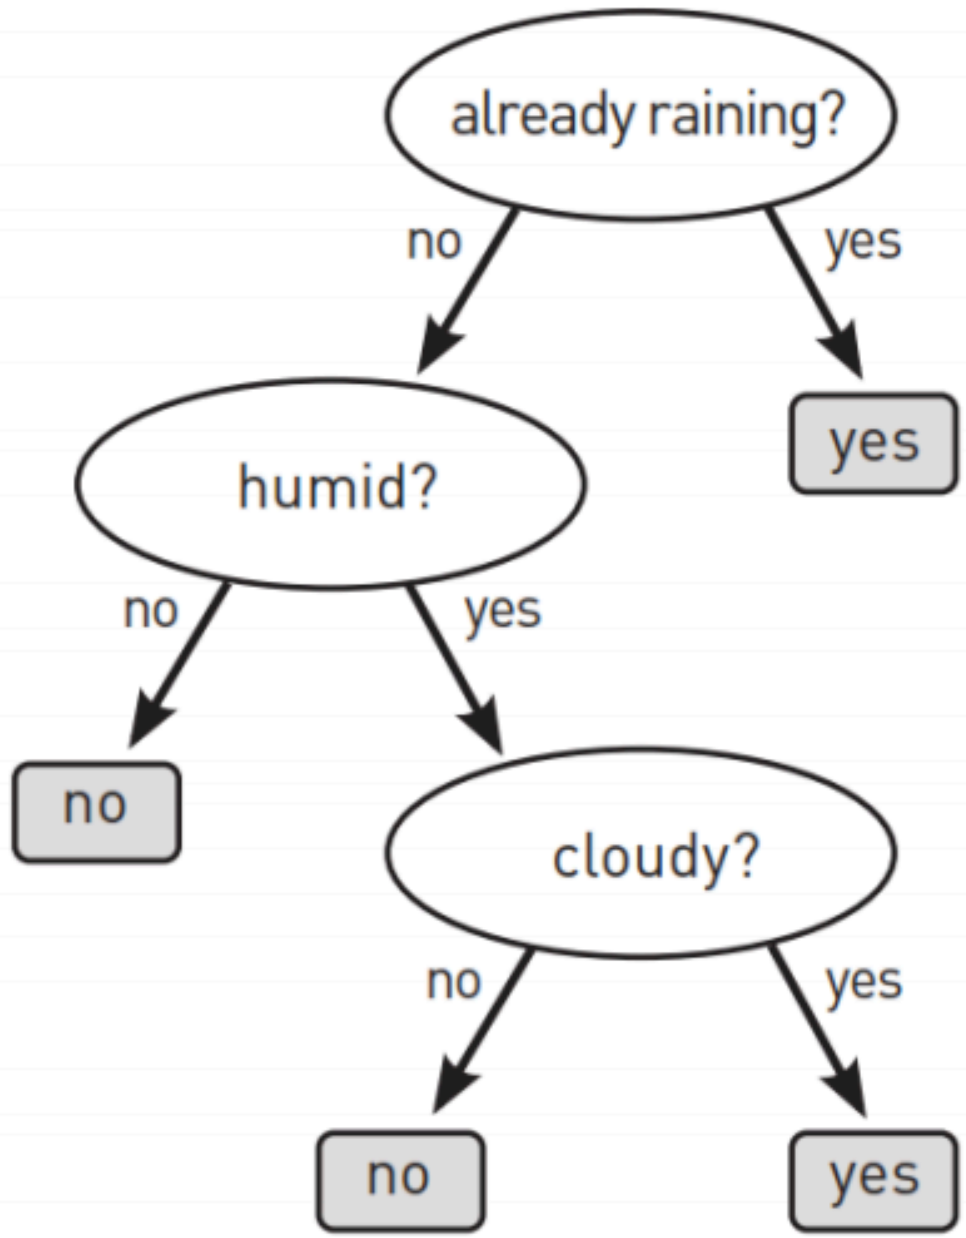

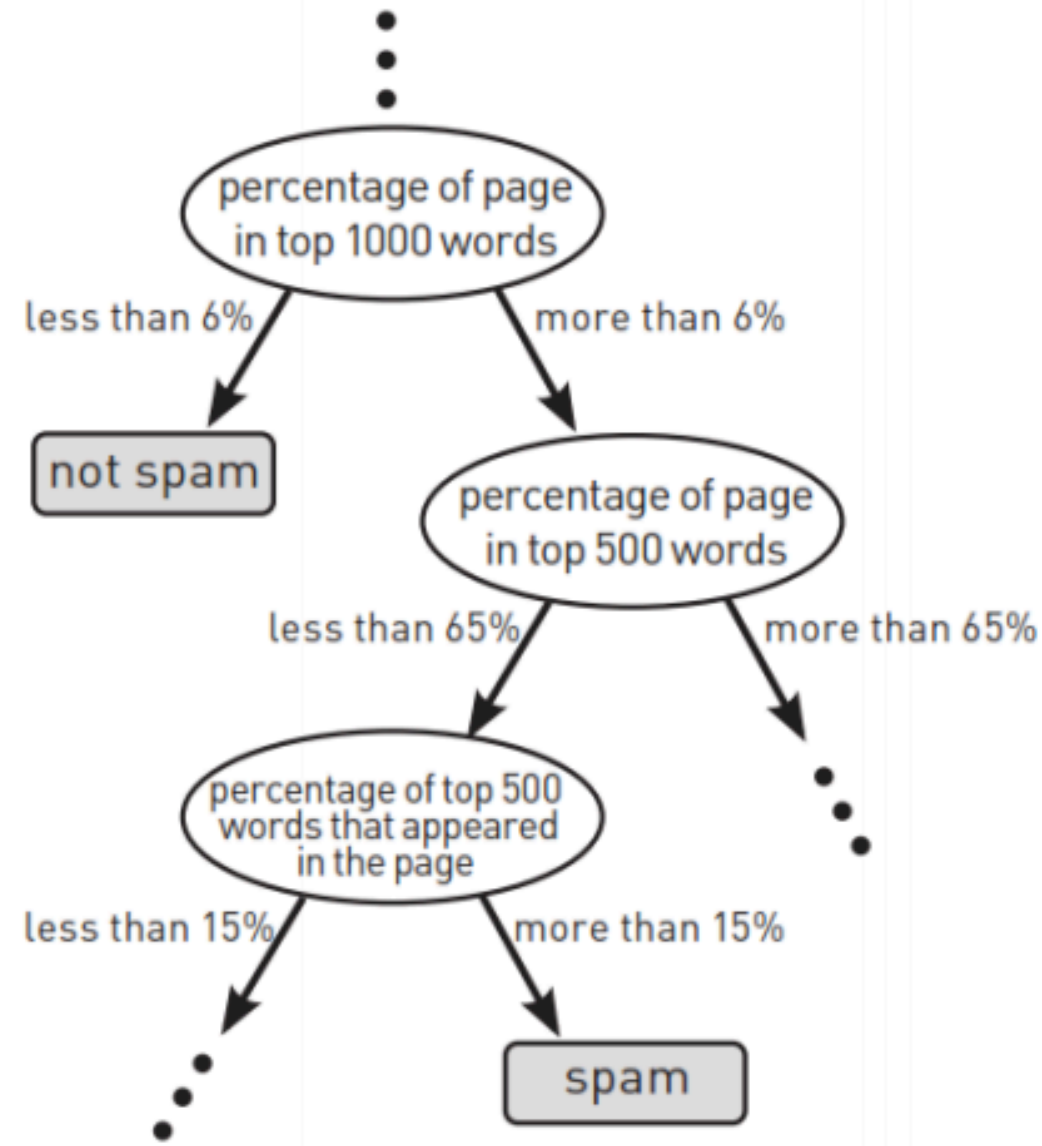

Ntoulas et al, WWW (2006)

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

What kind of “questions” can the Kinect ask in its twenty questions?

- Simplified version:
  - “is the pixel at *that* offset in the background?”
- Real version:
  - “how does the (normalized) depth at *that* pixel compare to *this* pixel?” [see Shotton et al, equation 1]

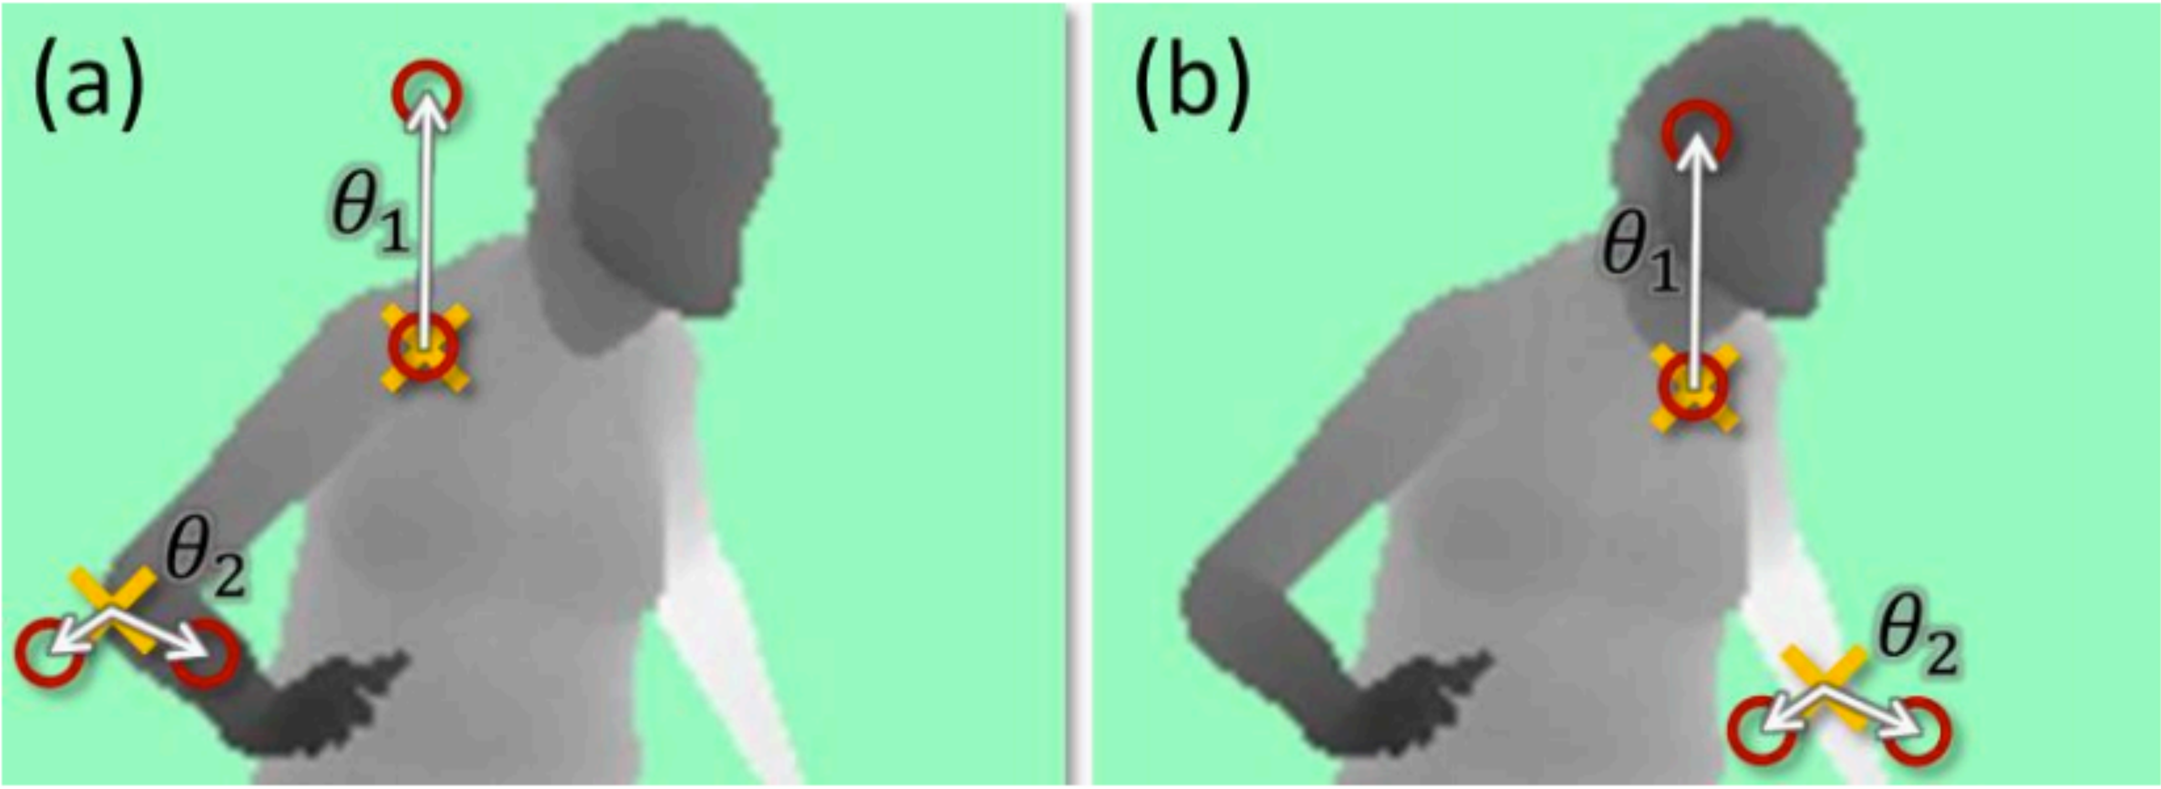

Shotton et al, CVPR(2011)

What kind of "questions" can the Kinect ask in its twenty questions?

- Simplified version:
  - "is the pixel at *that* offset in the background?"
- Real version:
  - "how does the (normalized) depth at *that* pixel compare to *this* pixel?" [see Shotton et al, equation 1]

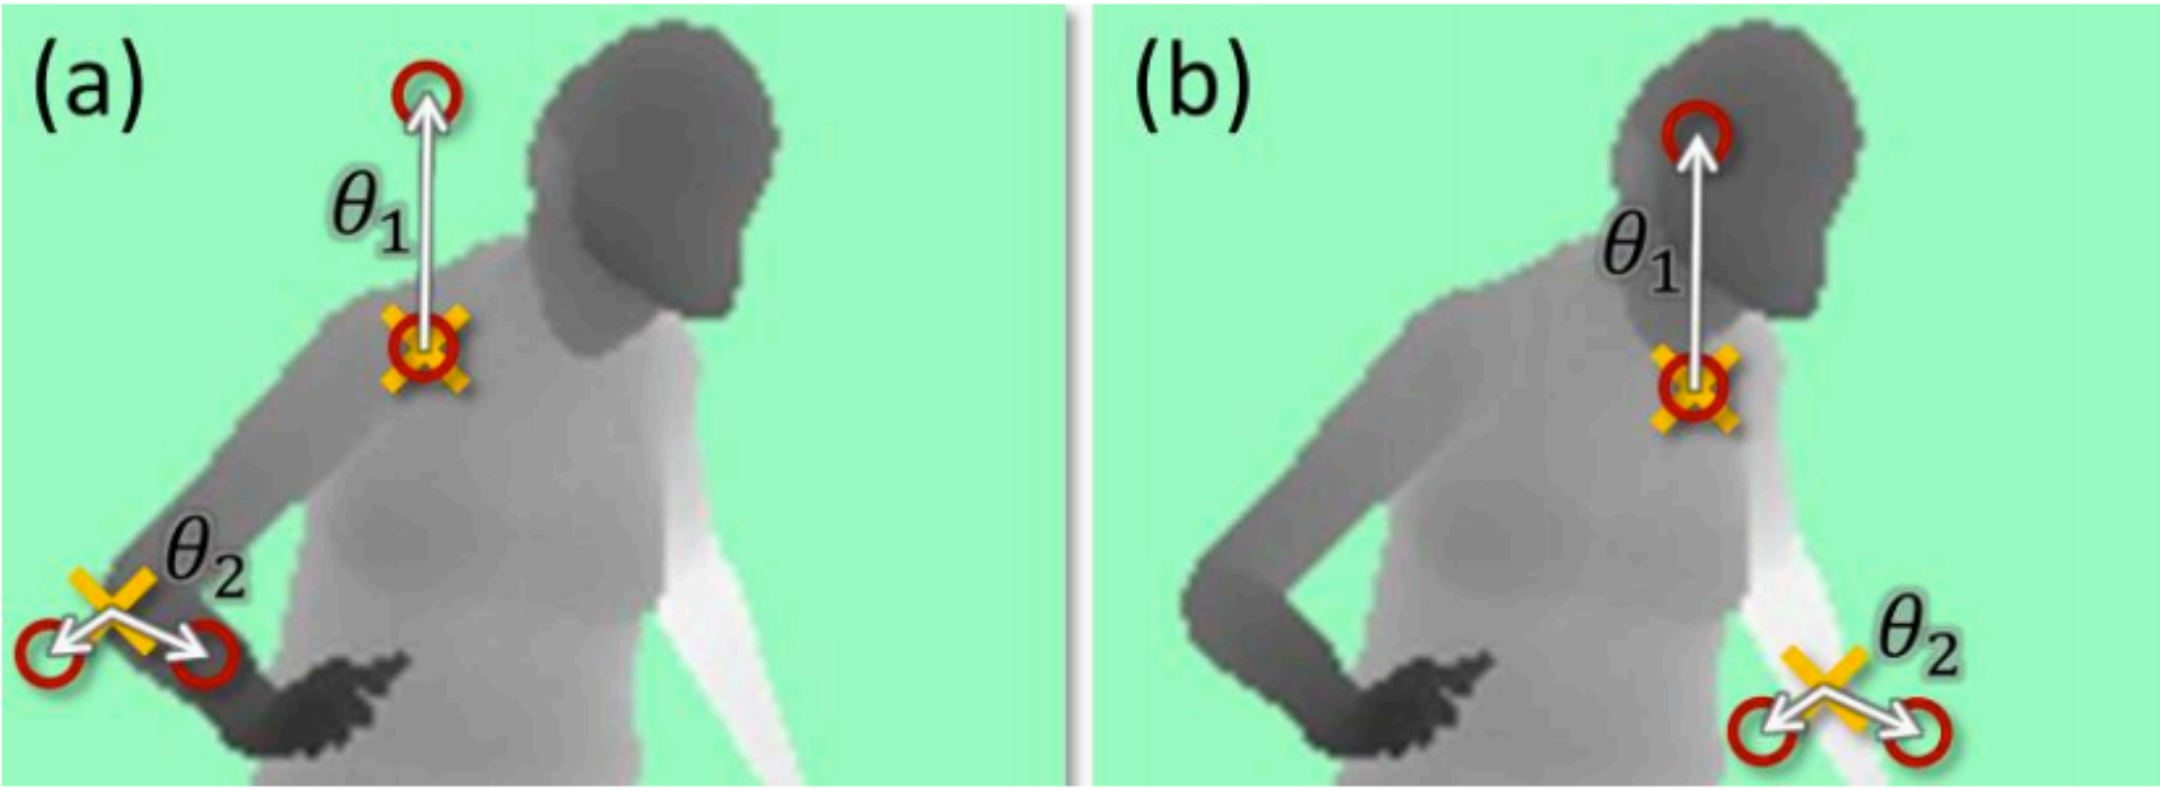

Shotton et al, CVPR(2011)

How does the Kinect Gesture Tracking Work

To learn a decision tree, you choose as the next question the one that is most "useful" on (the relevant part of) the training data

- E.g. for umbrella tree, is "raining?" or "cloudy?" more useful?
- In practice, "useful" = *information gain G* (which is derived from *entropy H*):

$$G(\phi) = H(Q) - \sum_{s\in\{\mathrm{l},\mathrm{r}\}} \frac{|Q_s(\phi)|}{|Q|} H(Q_s(\phi))$$

Shotton et al, CVPR(2011)

How does the Kinect Gesture Tracking Work

To learn a decision tree, you choose as the next question the one that is most "useful" on (the relevant part of) the training data

- E.g. for umbrella tree, is "raining?" or "cloudy?" more useful?
- In practice, "useful" = *information gain G* (which is derived from *entropy H*):

$$G(\phi) \quad = \quad H(Q) - \sum_{s \in \{\mathrm{l},\mathrm{r}\}} \frac{|Q_s(\phi)|}{|Q|} H(Q_s(\phi))$$

Shotton et al, CVPR(2011)

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

Kinect actually uses a *randomized* decision *forest*

- Randomized:
  - Too many possible questions, so use a random selection of 2000 questions each time
- Forest:
  - learn multiple trees
  - to classify, add outputs of the trees
  - outputs are actually probability distributions, not single decisions

Kinect actually uses a *randomized* decision *forest*

- Randomized:
  - Too many possible questions, so use a random selection of 2000 questions each time
- Forest:
  - learn multiple trees
  - to classify, add outputs of the trees
  - outputs are actually probability distributions, not single decisions

Kinect actually uses a *randomized* decision *forest*

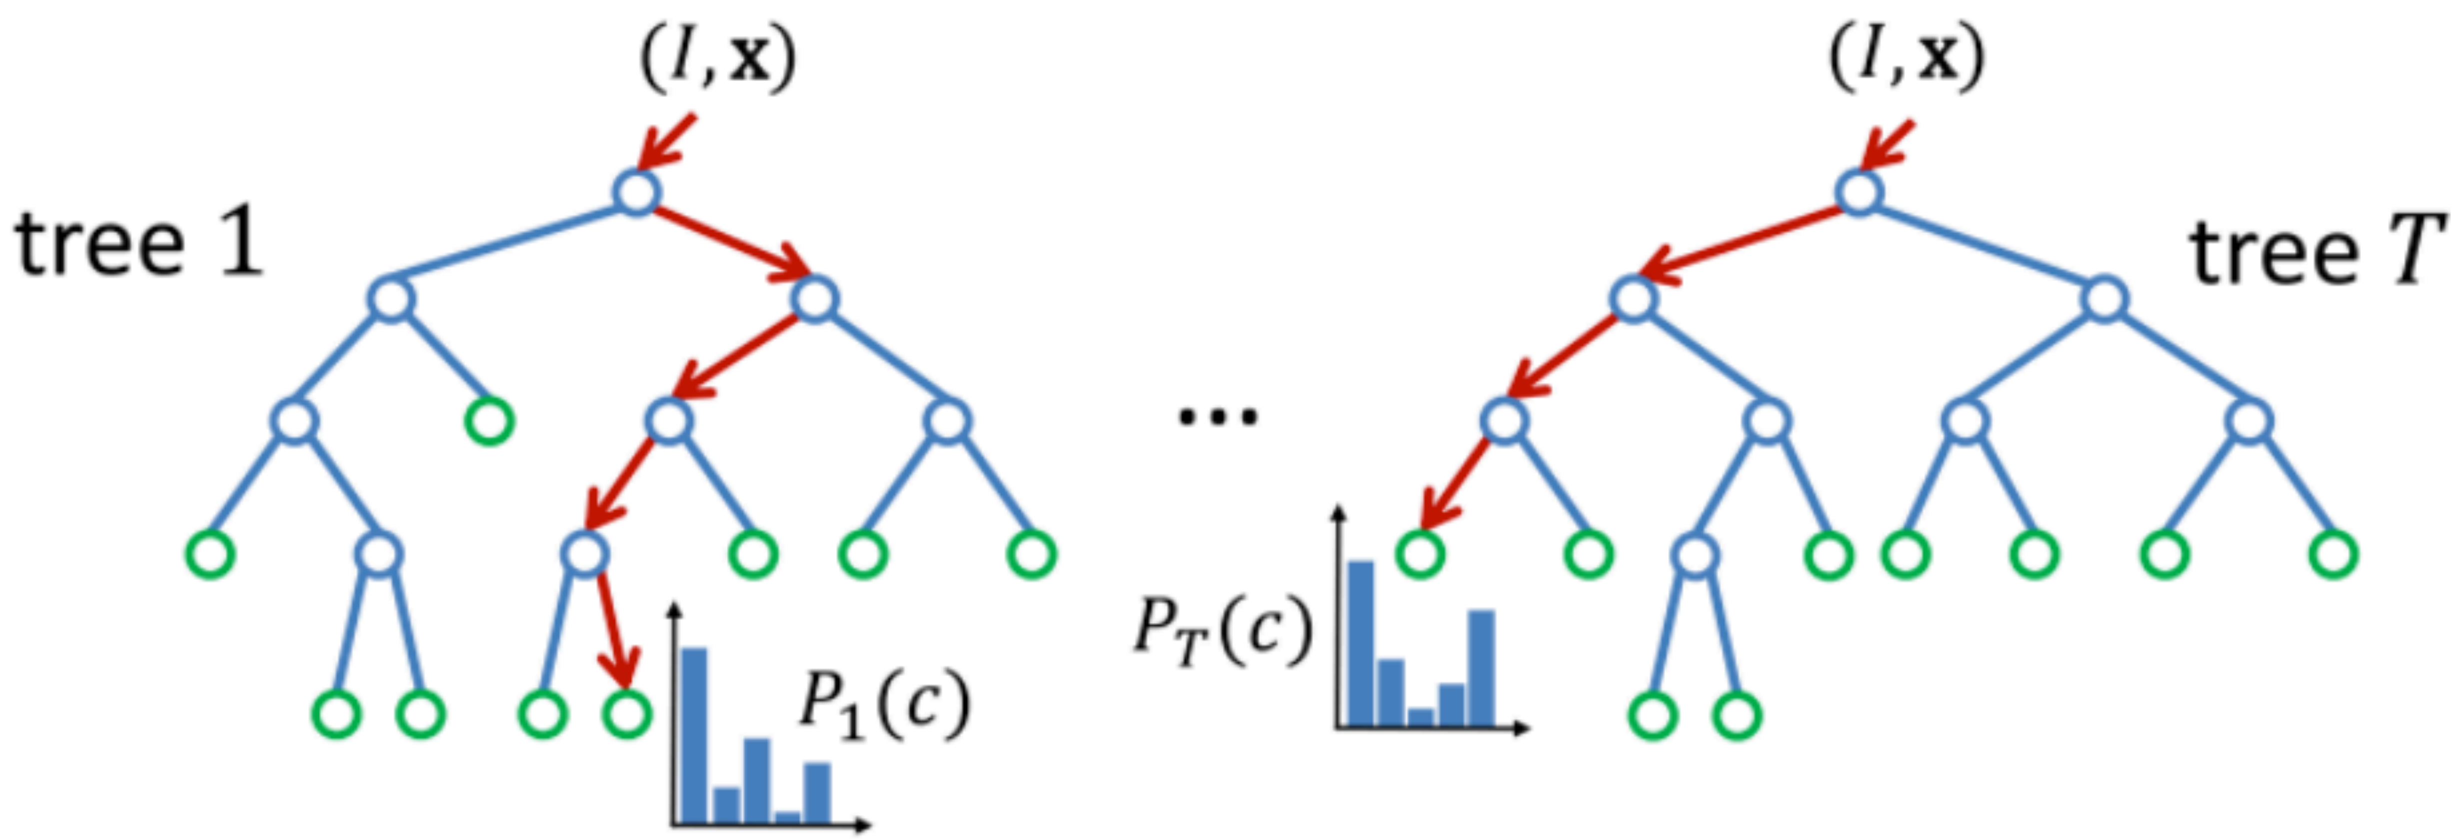

Shotton et al, CVPR(2011)

Kinect actually uses a *randomized* decision *forest*

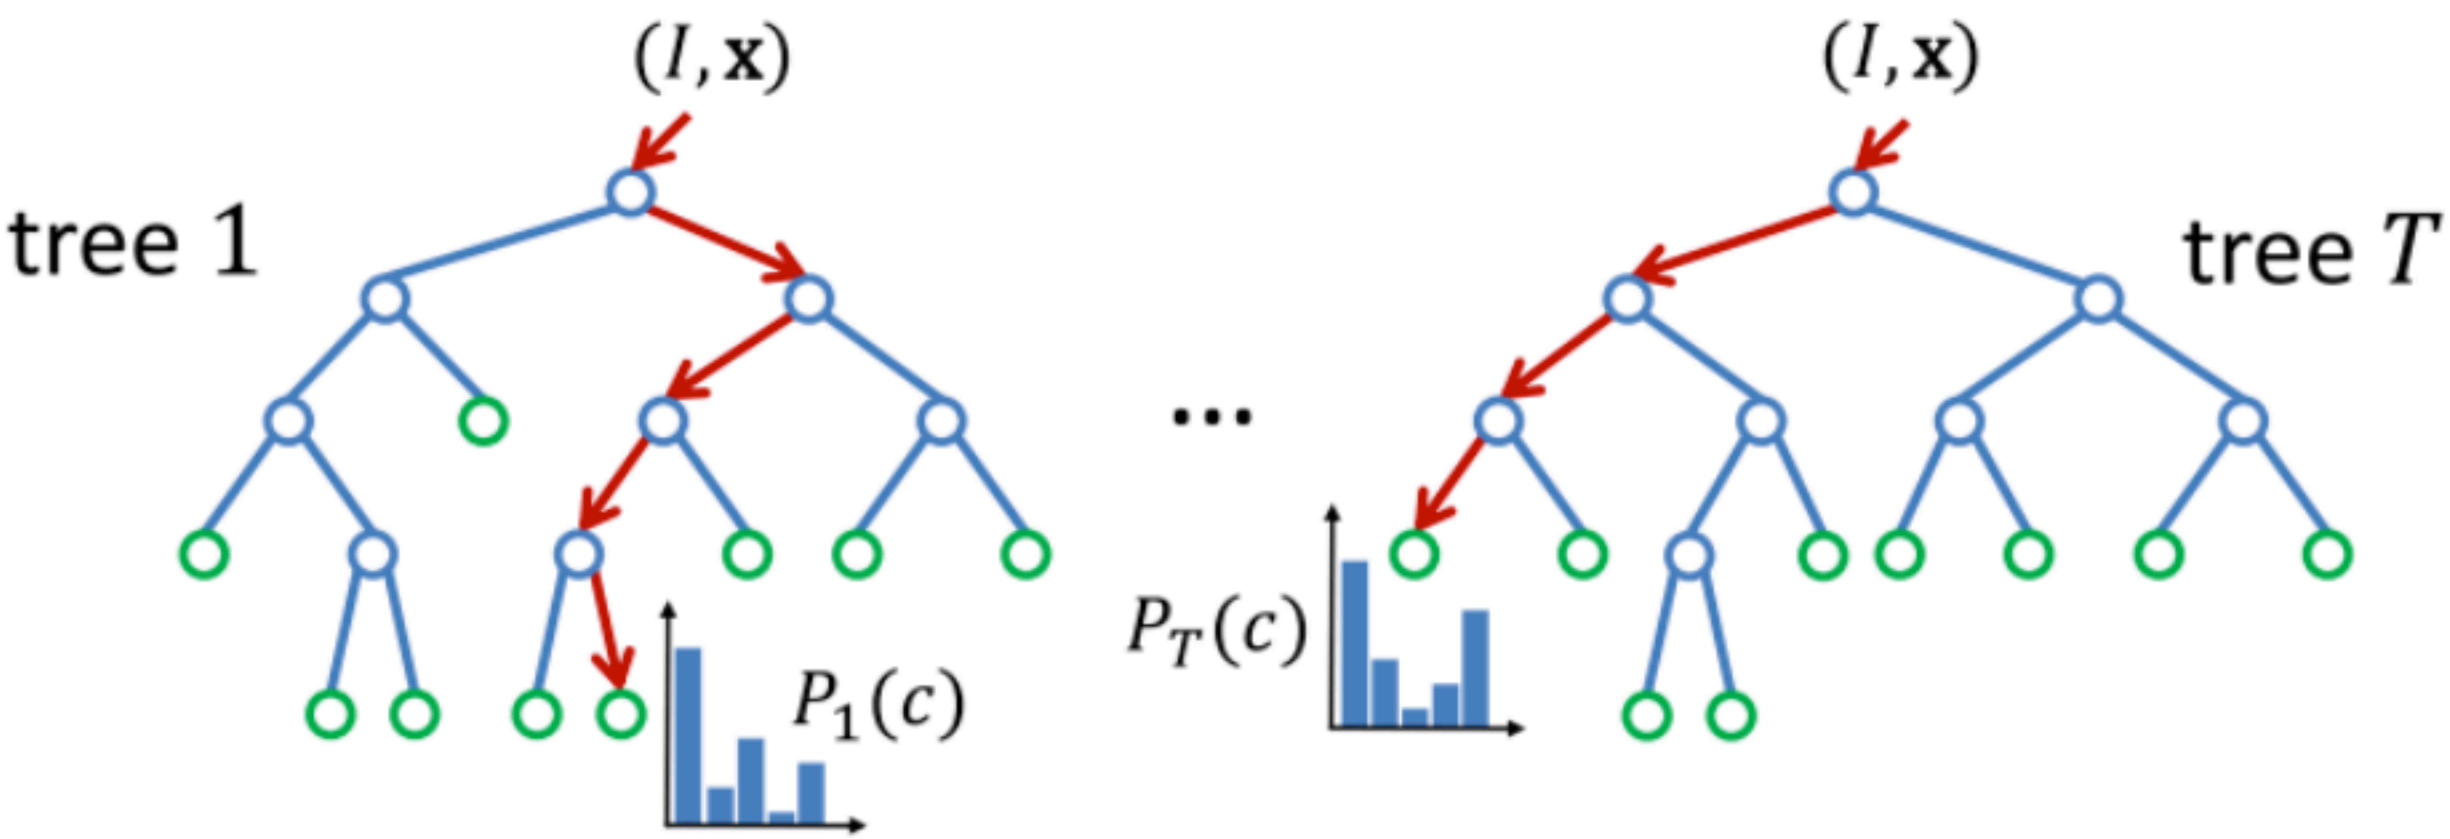

Shotton et al, CVPR(2011)

How does the Kinect Gesture Tracking Work

Learning the Kinect decision forest requires 24,000 CPU-hours, but takes only a day using hundreds of computers simultaneously

"To keep the training times down we employ a distributed implementation. Training 3 trees to depth 20 from 1 million images takes about a day on a 1000 core cluster."

—Shotton et al, CVPR(2011)

How does the Kinect Gesture Tracking Work

Learning the Kinect decision forest requires 24,000 CPU-hours, but takes only a day using hundreds of computers simultaneously

"To keep the training times down we employ a distributed implementation. Training 3 trees to depth 20 from 1 million images takes about a day on a 1000 core cluster."

—Shotton et al, CVPR(2011)

How does the Kinect Gesture Tracking Work

Stage 2 has 2 substages (use intermediate "body parts" representation)

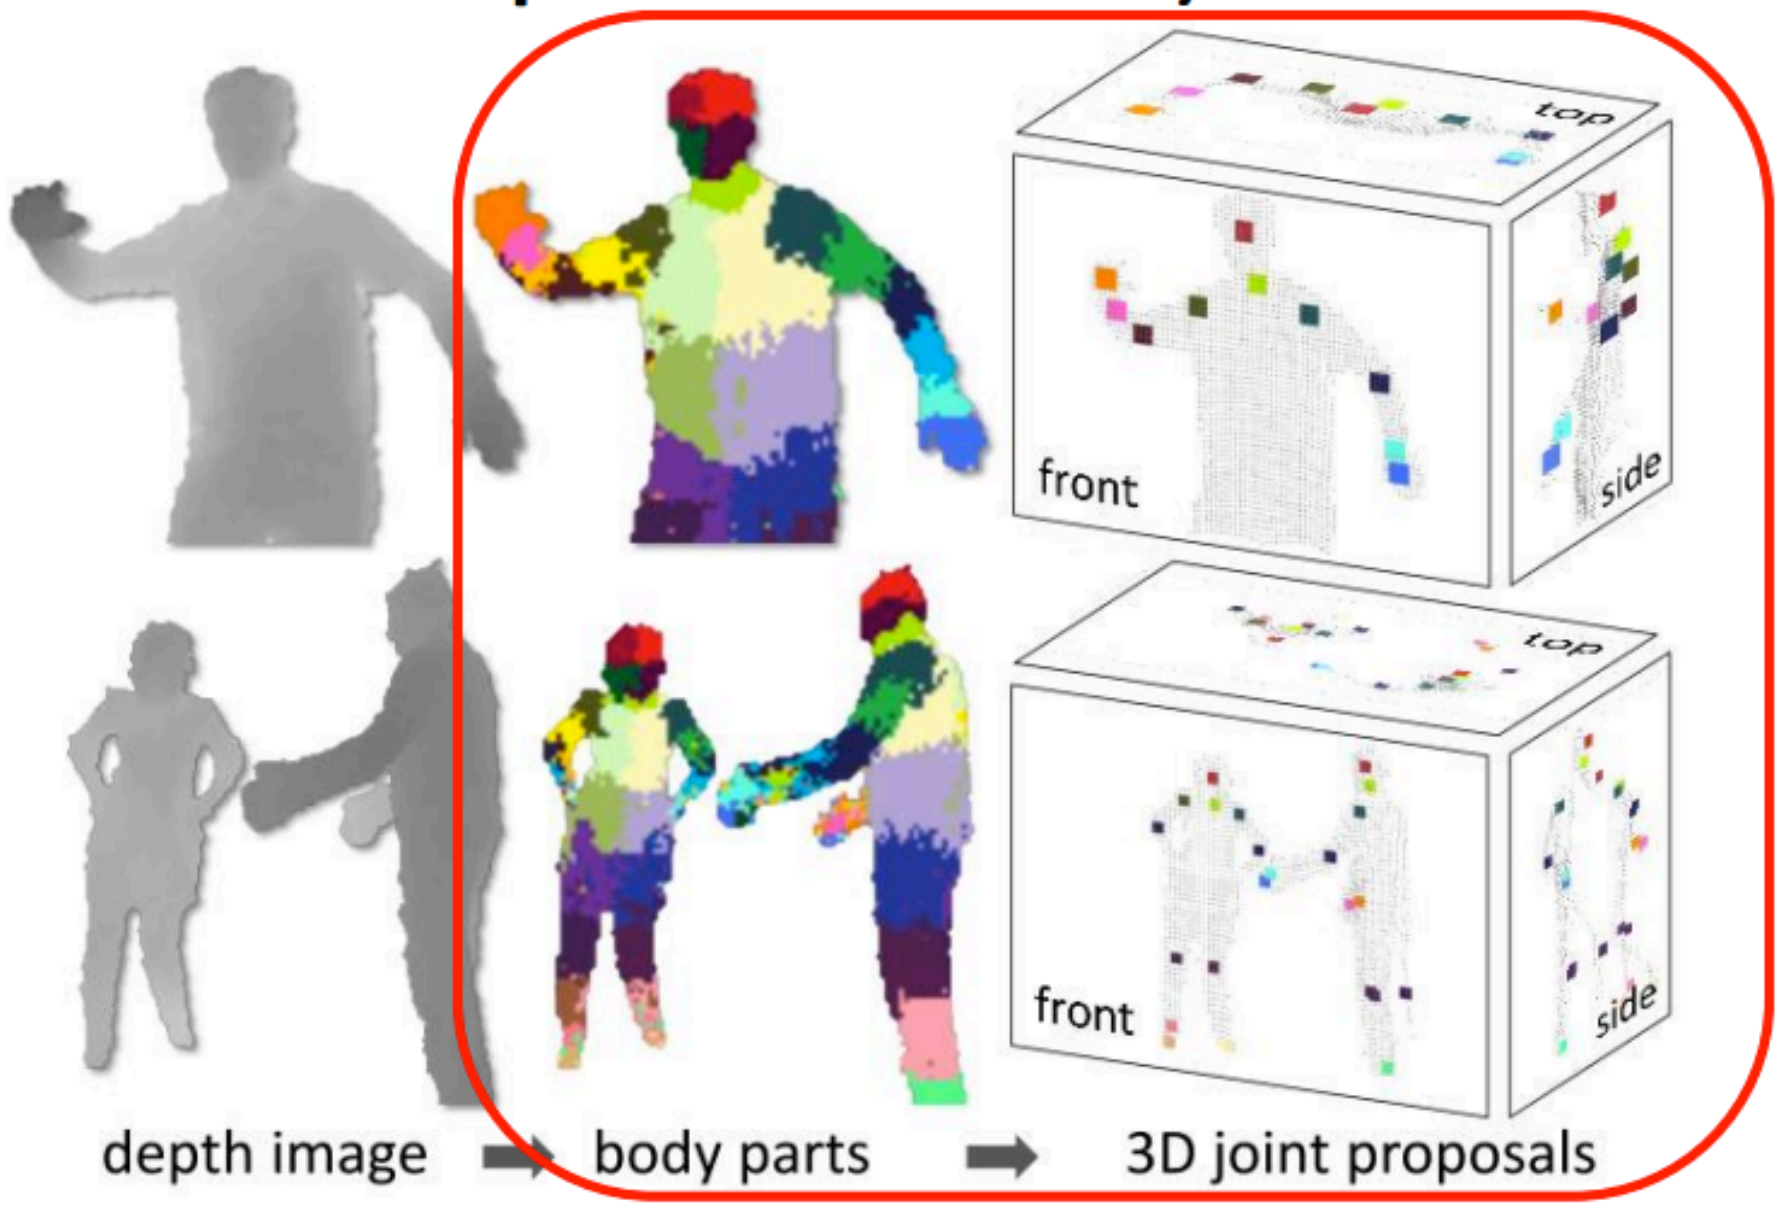

Shotton et al, CVPR(2011)

How does the Kinect Gesture Tracking Work

Stage 2 has 2 substages (use intermediate “body parts” representation)

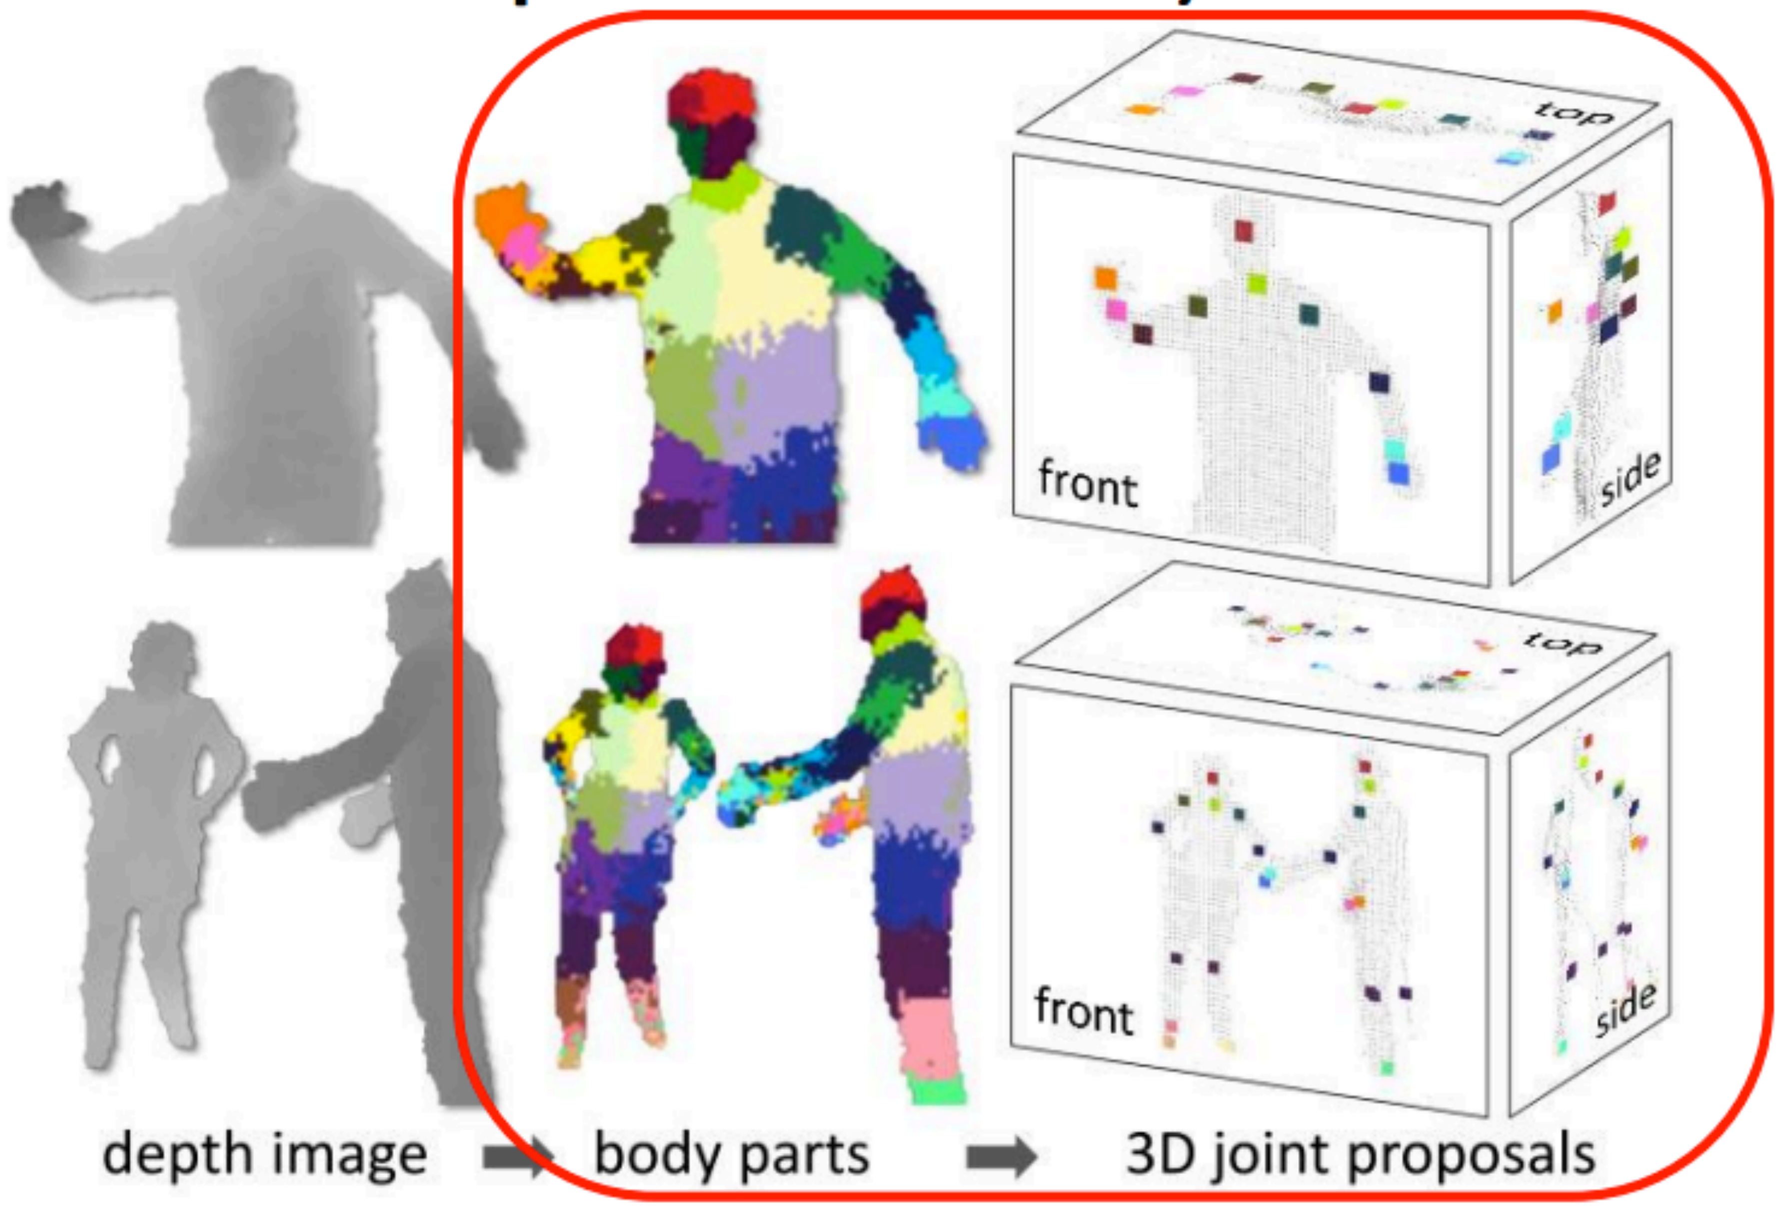

Shotton et al, CVPR(2011)

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

Stage 2.2 transforms the body part image into a skeleton

- The *mean shift* algorithm is used to robustly compute modes of probability distributions
- Mean shift is simple, fast, and effective

How does the Kinect Gesture Tracking Work

Stage 2.2 transforms the body part image into a skeleton

- The *mean shift* algorithm is used to robustly compute modes of probability distributions
- Mean shift is simple, fast, and effective

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

How does the Kinect Gesture Tracking Work

Intuitive Description

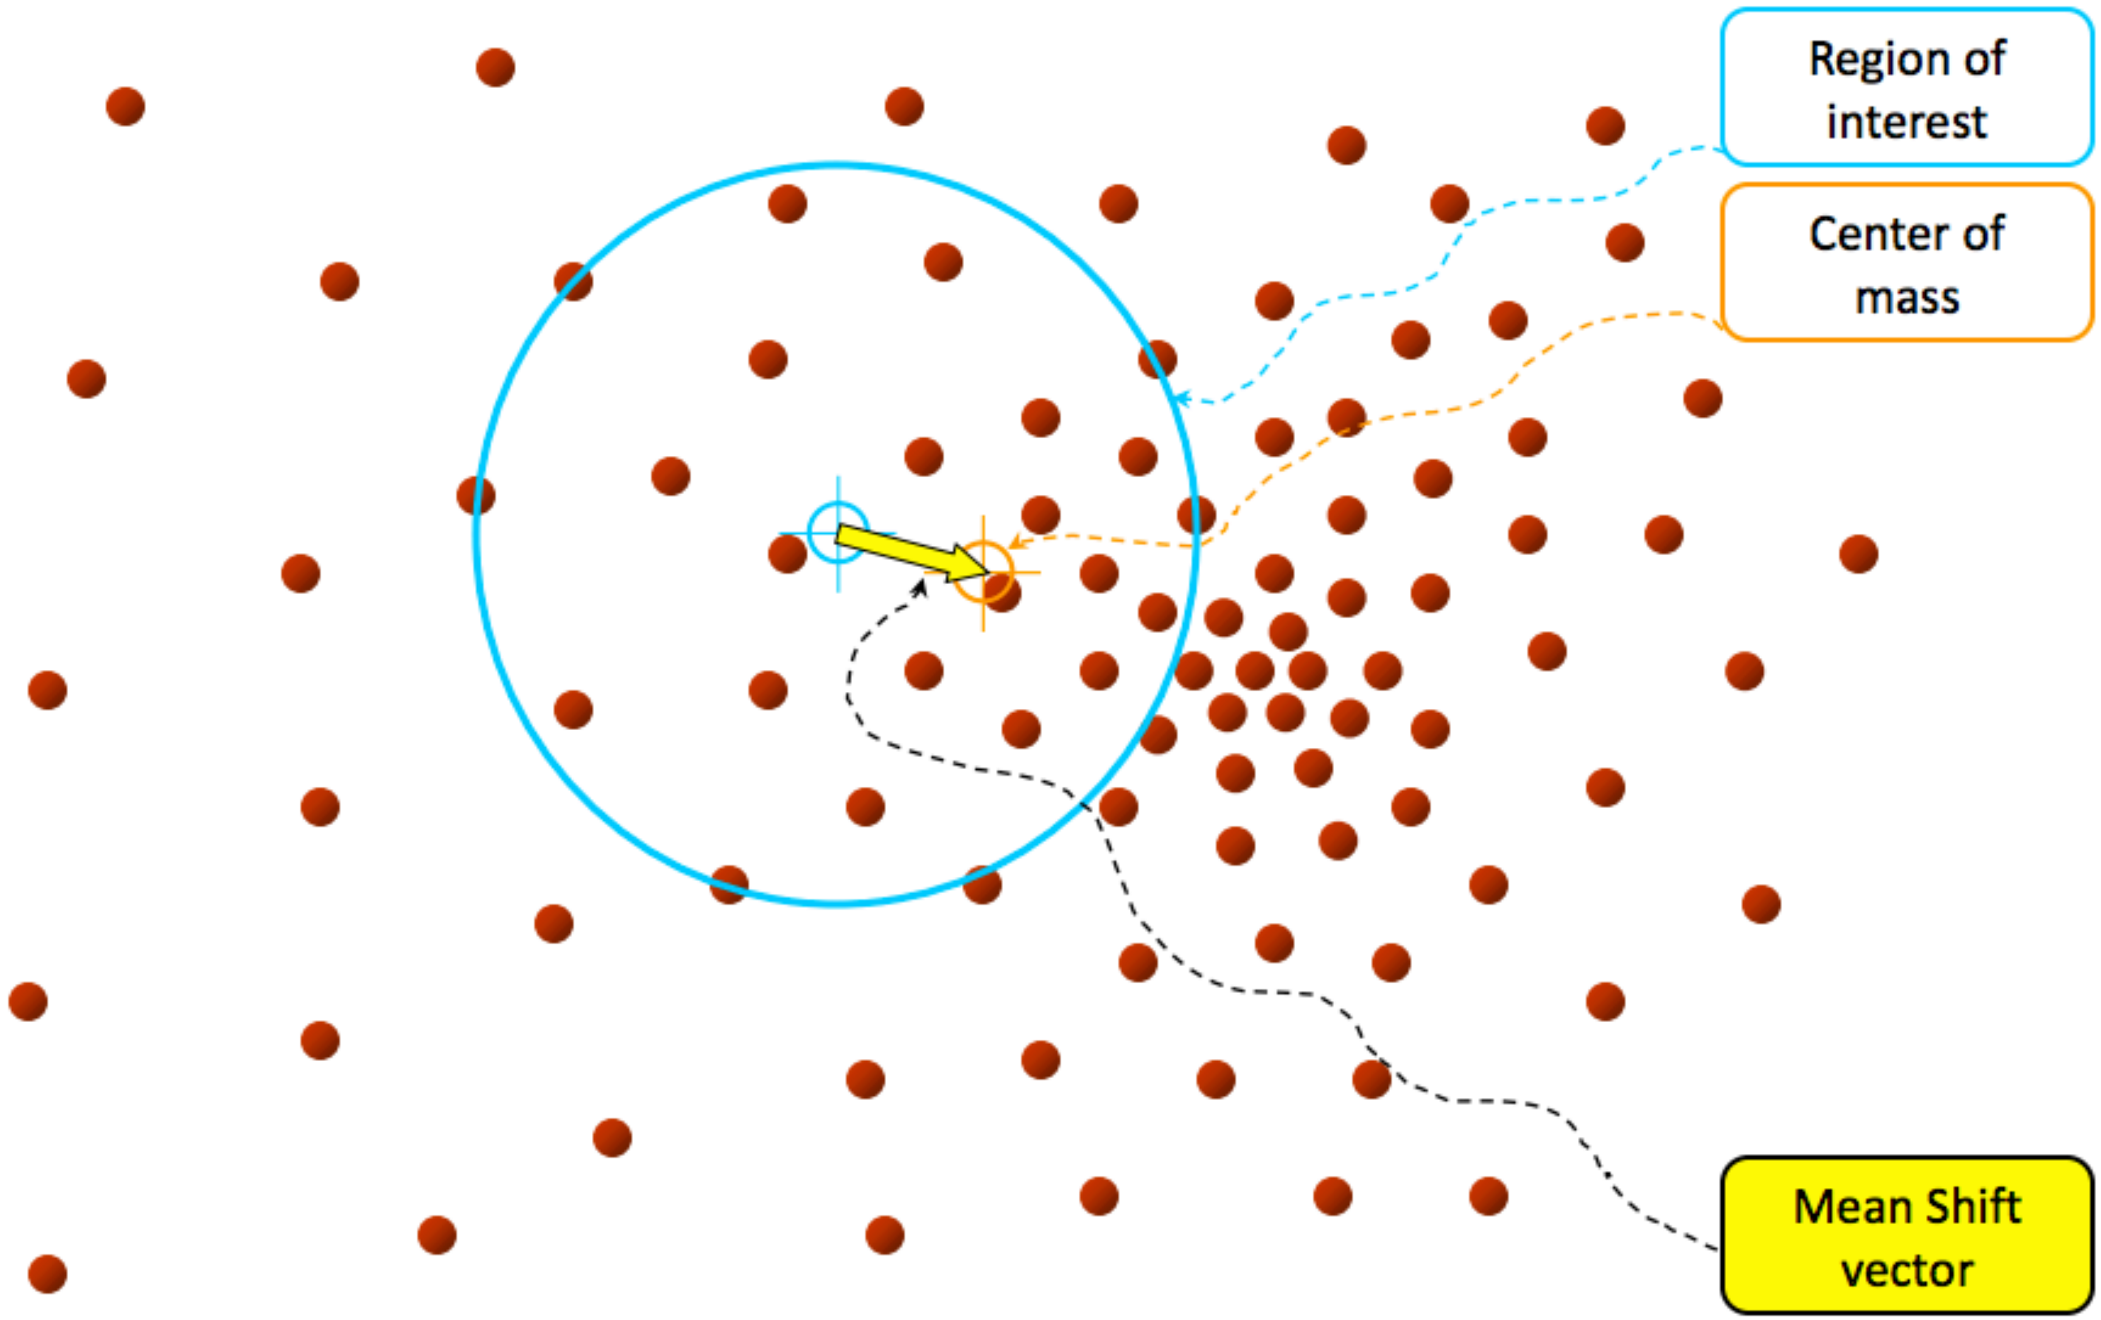

Objective : Find the densest region

Distribution of identical billiard balls

Slide taken from Ukrainitz & Sarel

How does the Kinect Gesture Tracking Work

Intuitive Description

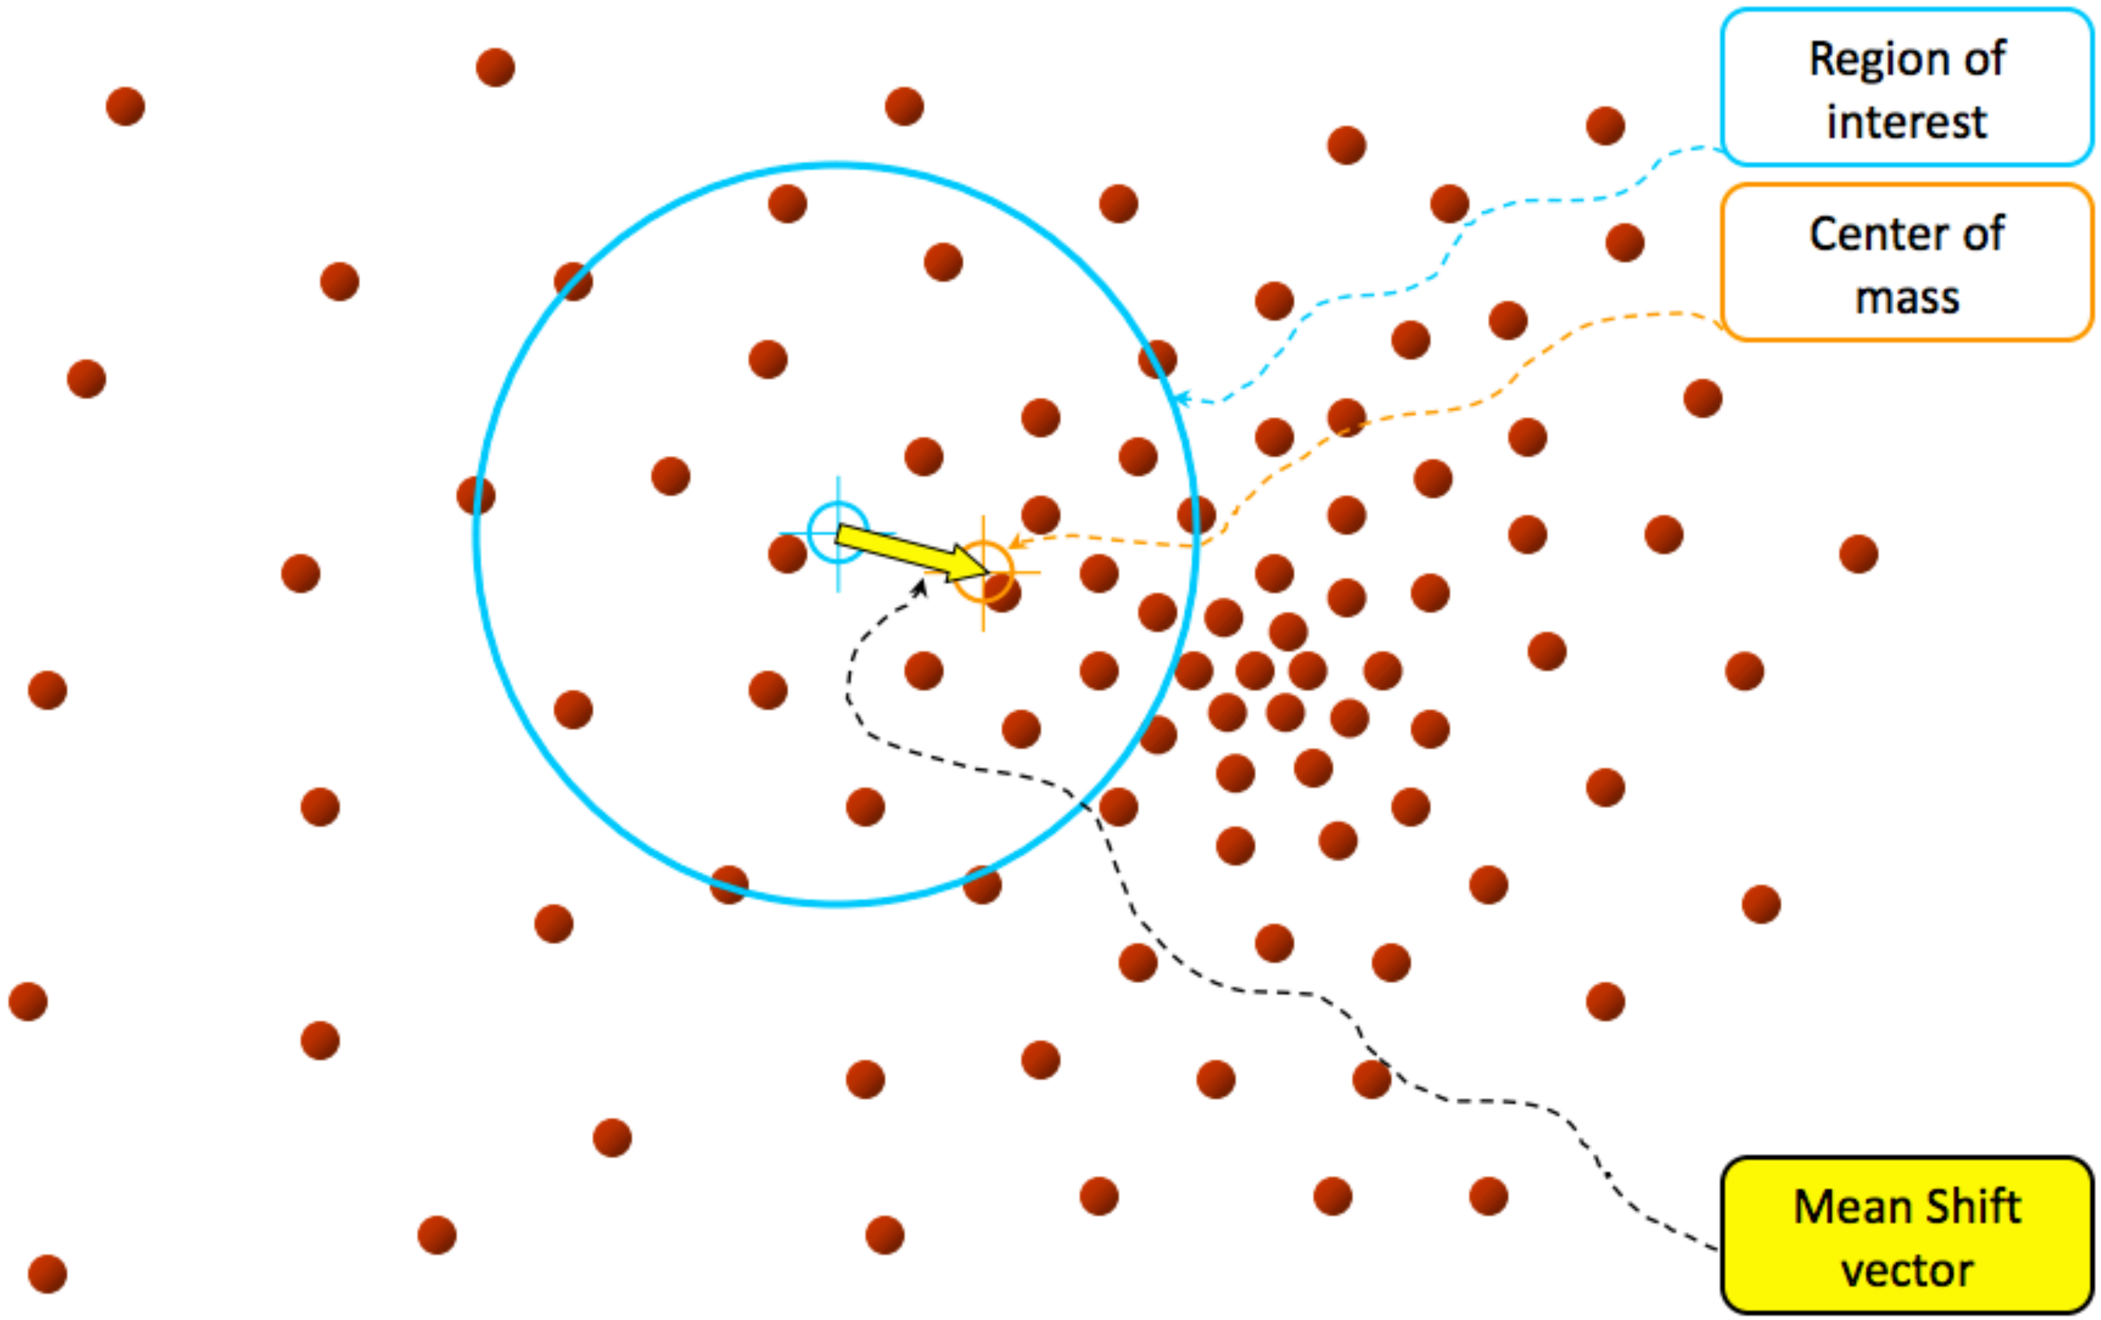

Objective : Find the densest region

Distribution of identical billiard balls

Slide taken from Ukrainitz & Sarel

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

How does the Kinect Gesture Tracking Work

The Kinect uses *structured light* and *machine learning*

- Inferring body position is a two-stage process: first compute a depth map (using structured light), then infer body position (using machine learning)
- The results are great!
- The system uses many college-level math concepts, and demonstrates the remarkable advances in computer vision in the last 20 years

The Kinect uses *structured light* and *machine learning*

- Inferring body position is a two-stage process: first compute a depth map (using structured light), then infer body position (using machine learning)
- The results are great!
- The system uses many college-level math concepts, and demonstrates the remarkable advances in computer vision in the last 20 years

How does the Kinect Gesture Tracking Work

How does the Kinect Gesture Tracking Work

http://youtu.be/abS7G5ZT17c

http://youtu.be/a0mdgdQfa-Q

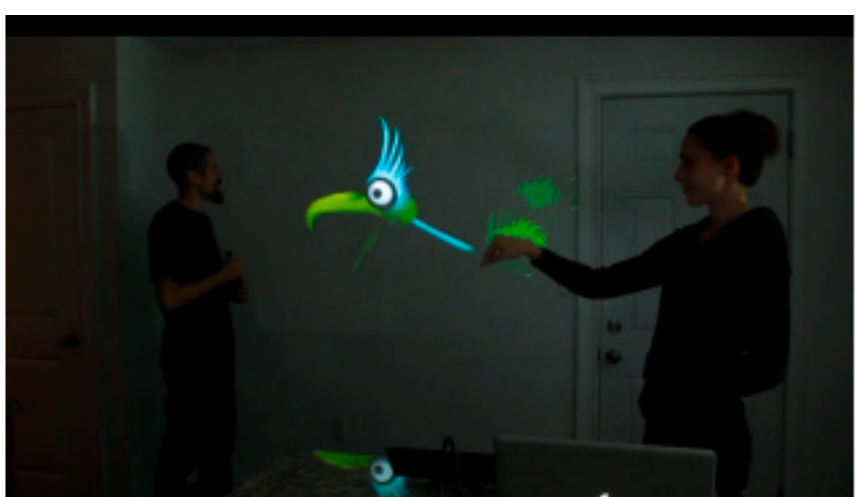

http://vimeo.com/16985224

How does the Kinect Gesture Tracking Work

How does the Kinect Gesture Tracking Work

http://users.dickinson.edu/~jmac/selected-talks/kinect.pdf

LUCI