How to Choose a 3D Vision Technology

Chris Osterwood
Chief Technology Officer, Carnegie Robotics
osterwood@carnegierobotics.com

Carnegie Robotics, LLC

- Supplier of reliable robotic components and systems.
- We have particular focus on and capability in autonomous and semi-autonomous mobile ground robots. This includes perception, autonomy, positioning, and safety – for size, weight, power, and cost constrained applications.

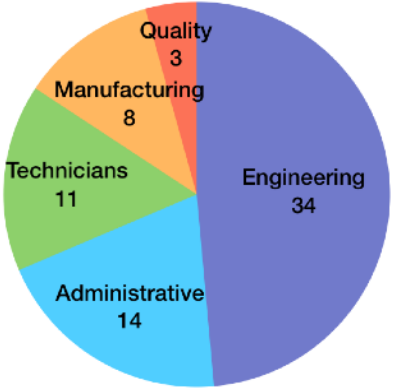

Outline

Carnegie Robotics

- Criteria & Definitions
- Overview of 3D Sensor Modalities : **LIDAR, ToF Camera, Stereo**
    - How do they work?
    - What are their general advantages and disadvantages?
    - Sample Data
- 3D Sensor Testing & Edge Cases
- Not covered: Laser Triangulation, Structured Light, & Computed Tomography

Evaluation Criteria - Performance

Carnegie Robotics

- **Field of View (FOV):** Defines angular area of perceivable field of sensor.
- **Density:**
  - Angular step size between sample points.
  - Can be different horizontally and vertically.
- **Resolution:** Generally ***FOV x Density***.

- **Depth Accuracy:** Difference between measured range and actual range.
- **Depth Resolution:** Step size between possible measurement ranges (along measurement axis)
- **Minimum and Maximum Range:** Defines distances that are perceivable to the sensor. May vary by object material, brightness, reflectivity, etc.

- **Rate:** Specified in “frames” or “points” per second, depending on modality.

Evaluation Criteria - Non-Performance

- **Size, Weight, & Power**
- **Cost**

- **Sealing:** Dust proof, Water splash, Water streams, Immersion
- **Shock & Vibration Rating**

- **Communication Interface:** Ethernet, USB, Serial, CANBus, Firewire
- **Synchronization:** None, Hardware, Software (broadcast trigger or network time sync)
- **Software Interface:**
    - Published wire protocol?
    - Drivers. Platform? Language? License?
    - Example usage software?
- **Other:** Trigger, GPIO, Lighting, IMU, Temperature & Health Reporting

3D Sensor Technologies

3D LIDAR : How do they work?

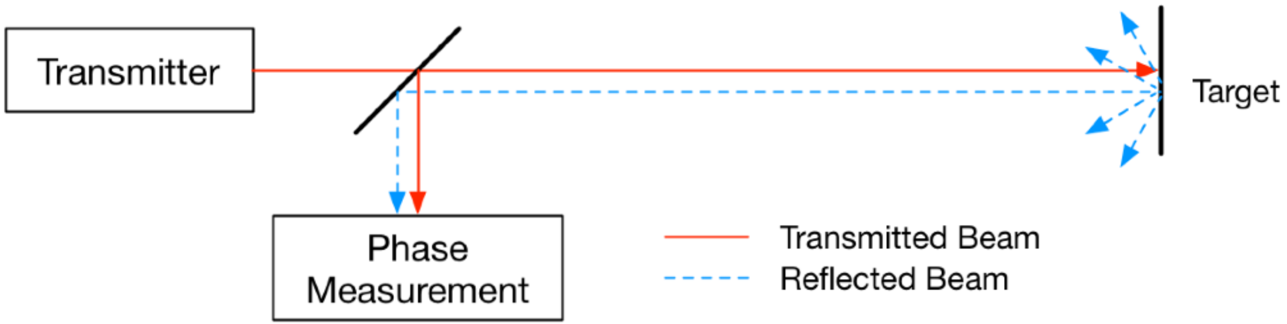

- Building block is single-beam system above.
- 3D LIDARS can have:
  - 1 beam steered mechanically in 2 DOF
  - Multiple beams steered mechanically in 1 DOF
- Steering can be done with rotating mirror, galvanometer, or directly upon the emitter & detector.

3D LIDAR

Carnegie Robotics

- Advantages:
  - **Constant error model —** 1 sigma can be +/- 1cm.
  - **Long range.** 10m to 2km possible depending on laser power and pulse duration.
  - **General wide horizontal FOV** due to rotating mirror or rotating emitters and detectors.
  - Can have large vertical FOV. Dependent on laser and mechanism details.

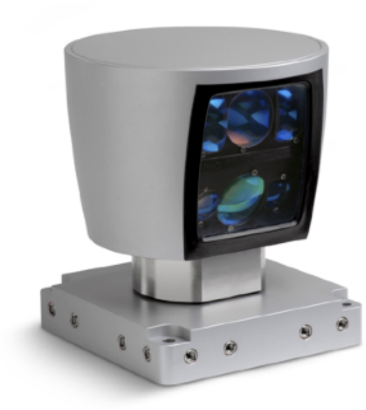

- Downsides:
  - **Cost.** Trends are encouraging here, but low cost systems are not yet available.
  - **Scan time** – no instantaneous capture of data. System has to compensate for platform (or mechanism) motion during scan.
  - **Limited angular density & asymmetric density.**
  - **Moving parts** & non-planar windows are difficult to seal.
  - Affected by **obscurants** in the environment.

ToF Cameras : How do they work?

Carnegie Robotics

- Like LIDAR, but every pixel is a pixel is a measurement point.
- Results in limited range or more transmission power.

- Accuracy increases with modulation frequency.
- Maximum range decreases with modulation frequency.
- Systems generally use multiple frequencies to allow for long range and non-ambiguous range.

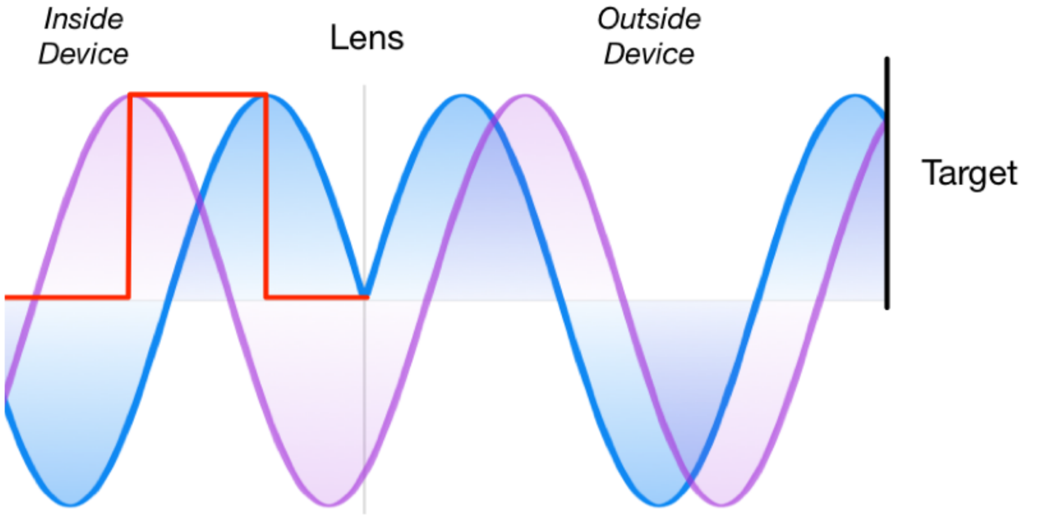

$$z = \frac{c}{4\pi f}\theta$$

$$z_{amb} = \frac{c}{2f}$$

θ Phase Shift (0 to 2π)
c Speed of Light
f Modulation Frequency

Transmitted Beam
Reflected Beam
Measured Phase Shift

30 Mhz wavelength is 10m, therefore z_{amb} is 5m

ToF Cameras

Carnegie Robotics

- Advantages:
  - **Dense data** regardless of scene texture.
  - **Instantaneous** horizontal & vertical FOV.
  - **Linear error** with respect to distance.
  - **Narrower shadow** behind positive obstacles.
  - **Solid state** – no moving parts.
  - Fewer “camera” parameters.

- Downsides:
  - **Low resolution.**
  - **Long integration time** (to increase SNR and resolve range ambiguity) causes motion blur if platform or objects are moving.
  - Susceptible to **multi-echo returns** distorting data.
  - Affected by **obscurant** in the environment.
  - **Limited FOV** – generally both horizontal and vertical.
  - May not have an “image” output.

Stereo Camera : How do they work?

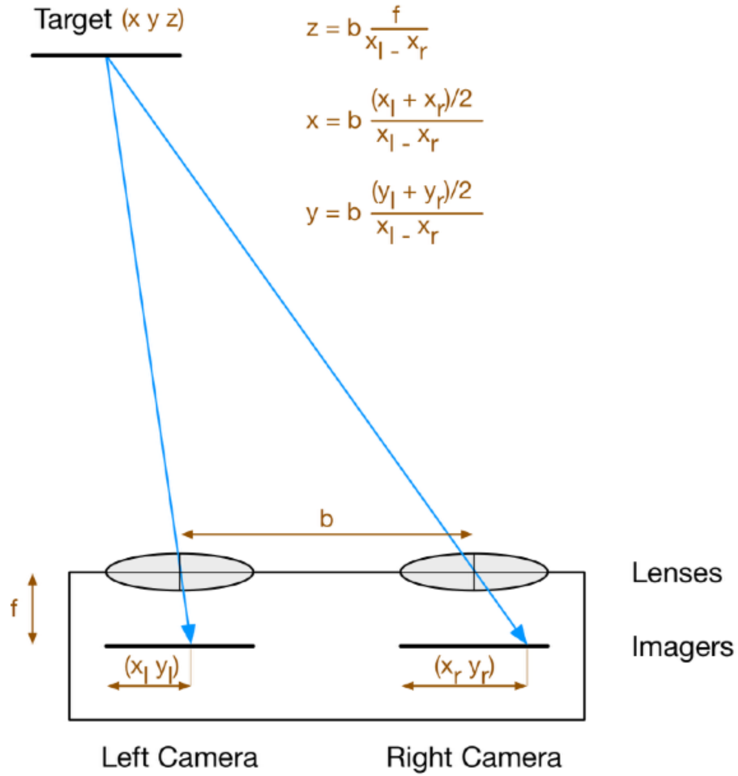

- Features matches are found in left and right cameras. Different in lateral position (disparity) is inversely proportional to distance.
- Matching process (correspondence problem) is critical to data accuracy and density.
- “Disparity search range” = number of horizontal pixels in right camera searched before moving onto the next left pixel. A larger search range allows seeing objects closer to the camera, but generally reduces 3D throughput.

Active & Passive Stereo

- *Active Stereo: Correspondence problem aided through projection of arbitrary texture into the scene.*

- Advantages:
  - **Very high resolution and density.**
  - **Instantaneous** horizontal & vertical FOV.
  - **Flexible FOV** options (from 130° to 10°)
  - **Monochrome or color images** are co-registered with 3D data.
  - Many "camera" parameters.
  - **Passive solution** is immune to obscurant in the environment.
  - **Solid state** – no moving parts.
- Downsides:
  - Data density **dependent on scene texture** or pattern projection.
  - **Non-linear error model** with respect to distance.
  - **Double shadow** behind positive obstacles.
  - **Computationally expensive** process increases BOM cost.

Sensor Error Models

Carnegie Robotics

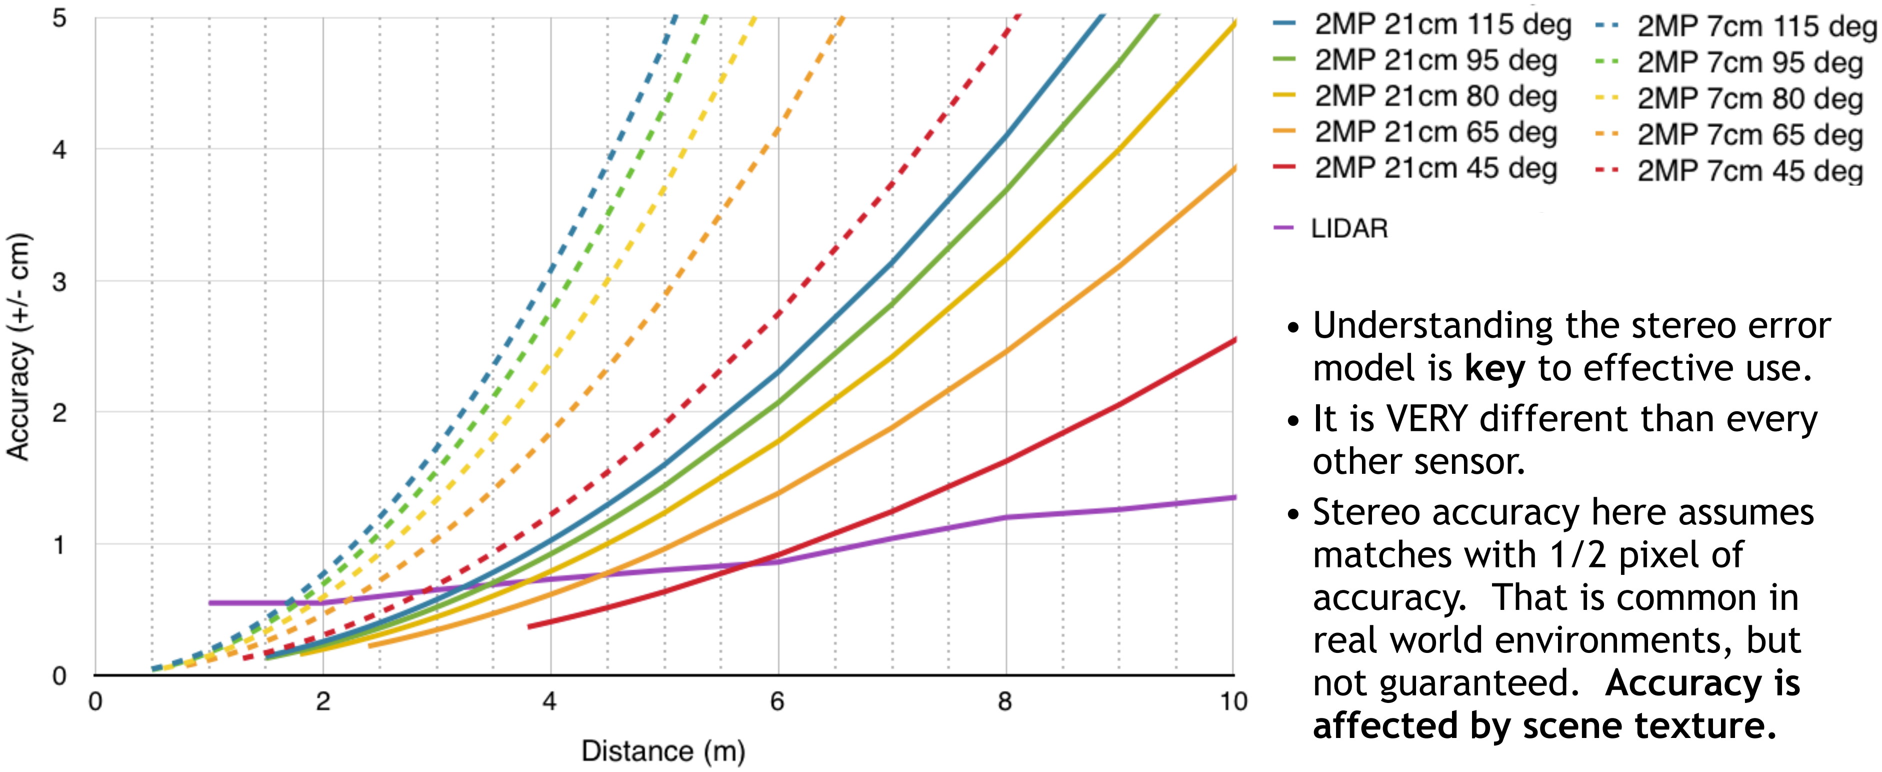

- Understanding the stereo error model is **key** to effective use.
- It is VERY different than every other sensor.
- Stereo accuracy here assumes matches with 1/2 pixel of accuracy. That is common in real world environments, but not guaranteed. **Accuracy is affected by scene texture.**

General Sweet Spots

Carnegie Robotics

- **LIDAR**
  - High speed vehicles — well served by their long range and high accuracy at range.
  - Vehicles with high turning rates — require wide horizontal FOV.
- **ToF Camera**
  - Indoor environments.
  - Short range object scanning & gesture recognition.
- **Stereo Camera**
  - Outdoor applications with high levels of ambient obscurants.
  - Multi-view applications (overlapping FOVs).
  - Applications which require a non-emitting solution.

	LIDAR	ToF	Stereo
HFOV			
VFOV			
Density			
Range Accuracy			
Min Range			
Max Range			
Data Rate			
Obscurant			
Cost			
Sealing			
Shock / Vibe			

Qualitative Comparison

3D LIDAR Sample Data

Carnegie Robotics

- 1: 8 seconds of persisted data.
- 2: 1 second of persisted data, high rotation rate.

- Notice:
  - Wide FOV.
  - Difficulty in seeing / identifying person.
  - Density differences horizontally & vertically.
  - Mixed pixel returns.

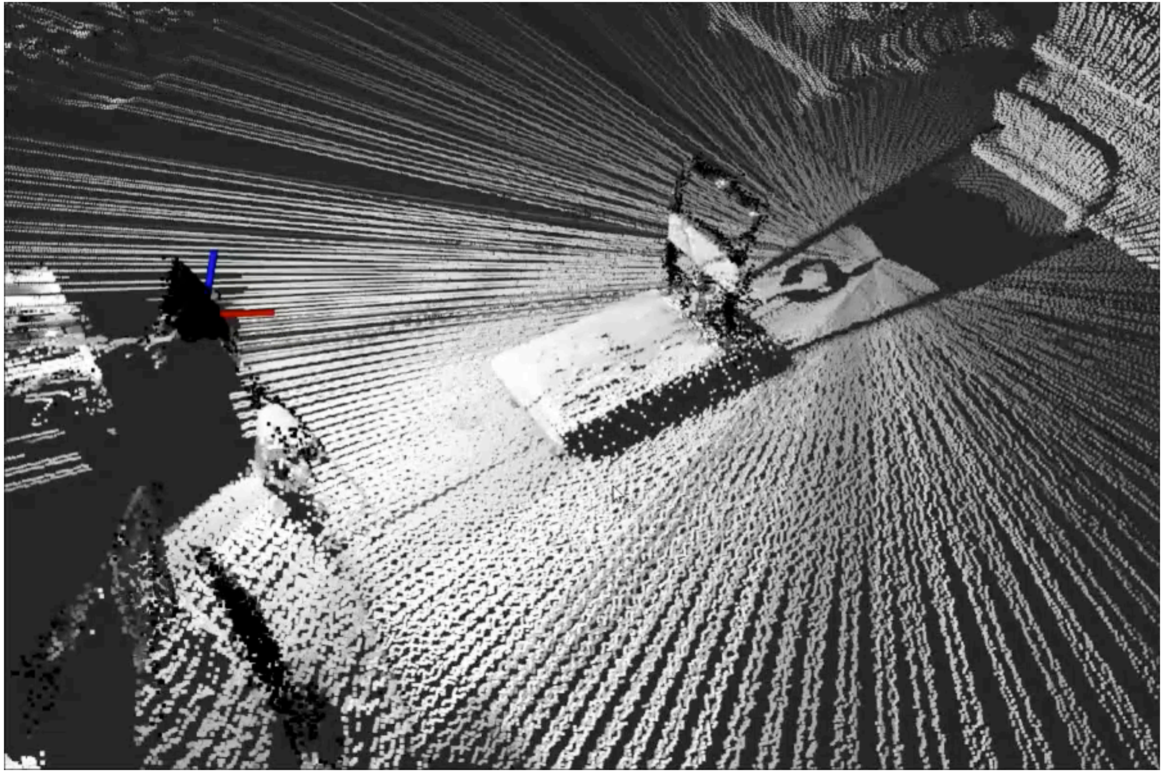

ToF Camera Sample Data : Low Cost

- 1: Axis Colored Scene
- 2: Intensity Colored Scene

- Notice:
  - Low resolution / Low Rate.
  - Narrow FOV.
  - Invisible / bent cart handle.
  - “Flashlight effect”.
  - Mid-air returns.
  - 1/2 of floor missing.

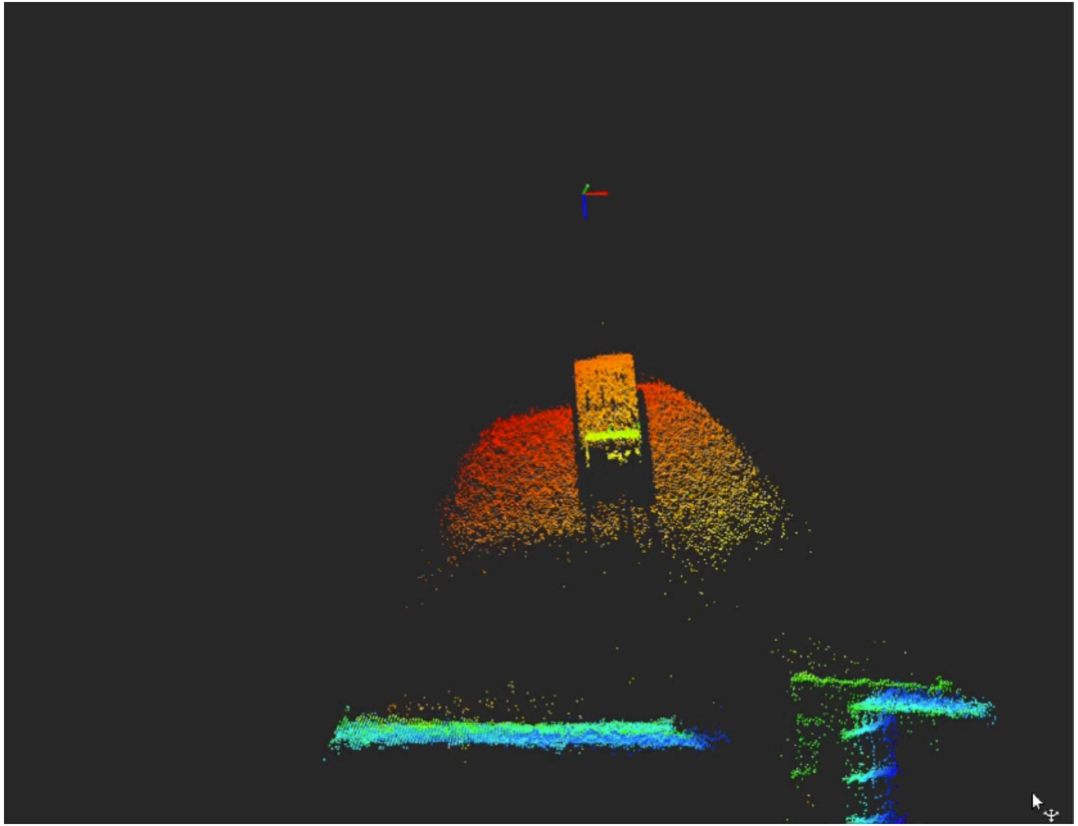

ToF Camera Sample Data : Industrial

Carnegie Robotics

- 1 & 2: Axis colored scene.
- Notice:
  - High resolution / High rate.
  - Distortion with near-field object.
  - Distortion to rear wall with foreground object.
  - Multi-path distortion in floor.
  - Mixed-pixel returns.

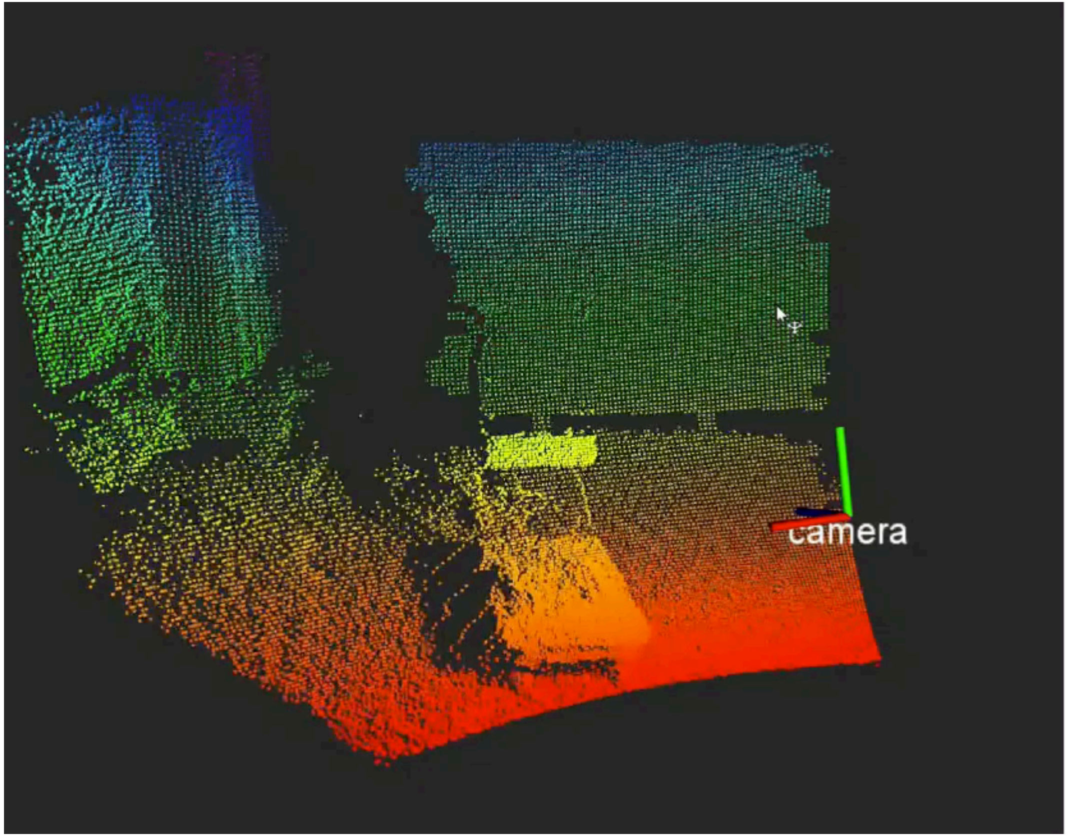

Stereo Sample Data

Carnegie Robotics

- 1: Overview of scene.
- 2: Close-up of chess board.
  - Post-matching cost threshold starts high (all data), is lowered (only high confidence data), and is raised again.

- Notice:
  - High density — can see pieces, can see cart handle.
  - High Rate & Wider FOV.
  - Correlated scene color.
  - Data below ground caused by specular reflections.
  - False “spike” in rear wall caused by cart handle.

Active versus Passive Stereo

Carnegie Robotics

- IR Pattern Projector turned on and off throughout scene.
- Areas of low-confidence have higher-confidence with projector on. This results in lower-noise and more accurate data.
- Scene has RGB colormap of monochrome intensity.

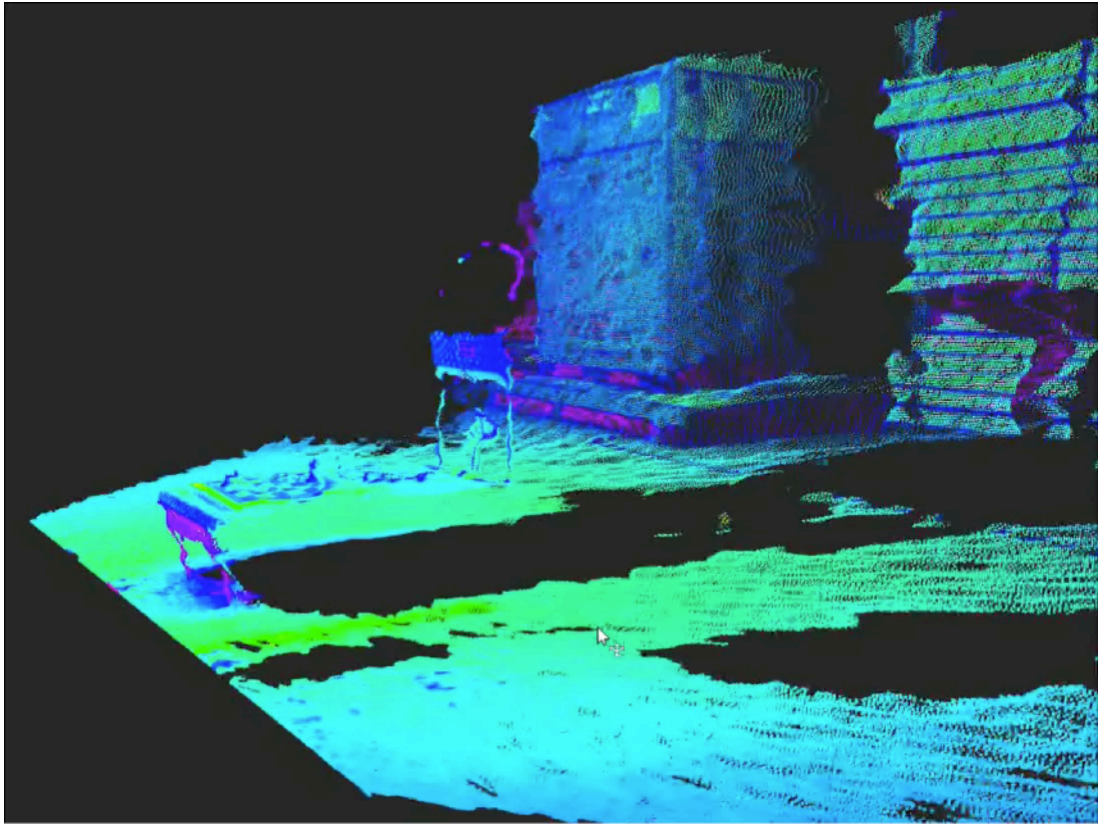

Stereo from Passive Thermal

Carnegie Robotics

Optical Stereo

Thermal Stereo

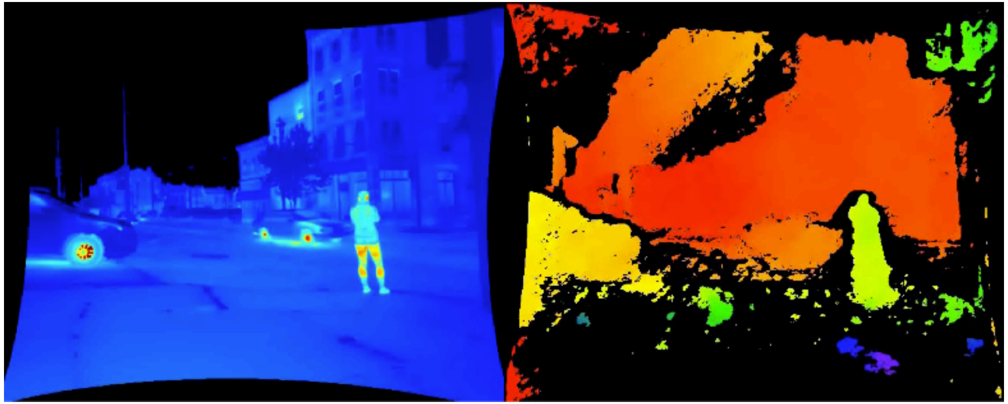

Testing 3D Sensors

3D LIDAR : Lighting : Vendor A

Carnegie Robotics

300 lux : Fluorescent

8	3	6	9	7	5

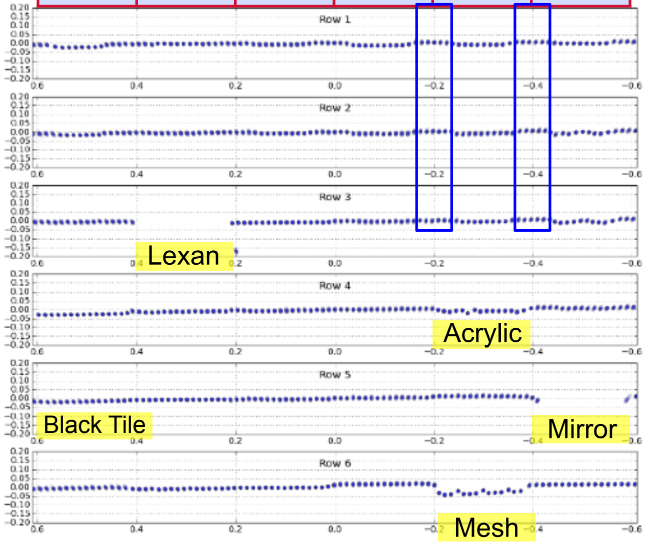

5000 lux : Halogen

4	9	5	4	7	6

Std. dev. of first row (mm)

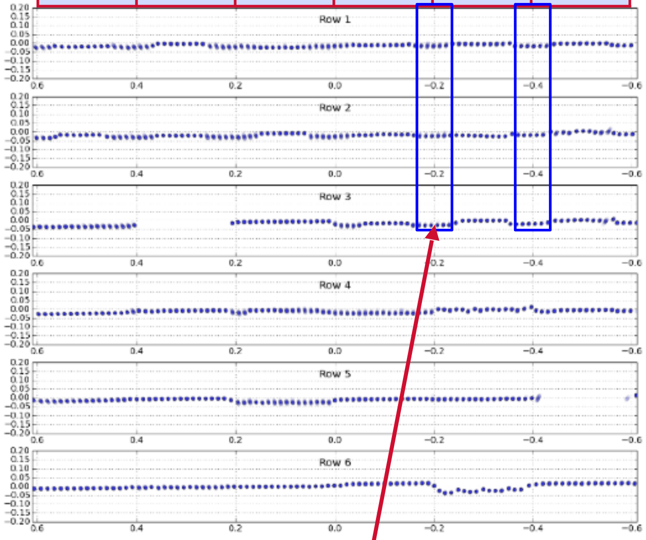

White areas shifted 2.5cm closer to sensor

3D LIDAR : Lighting : Vendor B

300 lux : Fluorescent

10	7	7	9	10	8

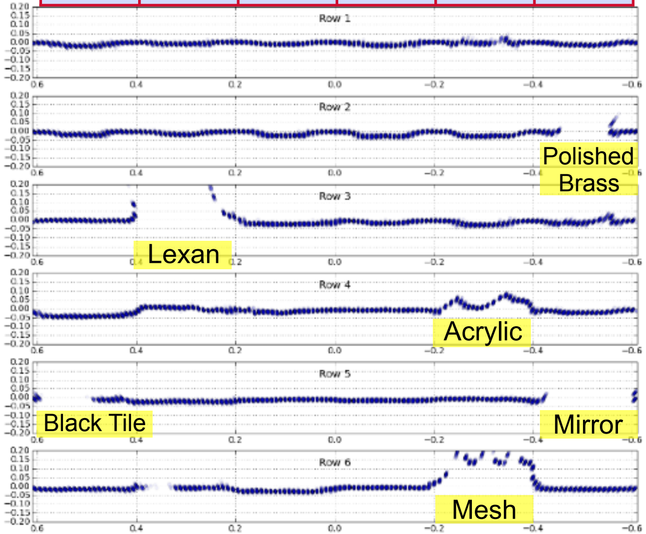

50% more noise than Vendor A

5000 lux : Halogen

11	9	10	12	13	12

Std. dev. of first row (mm)

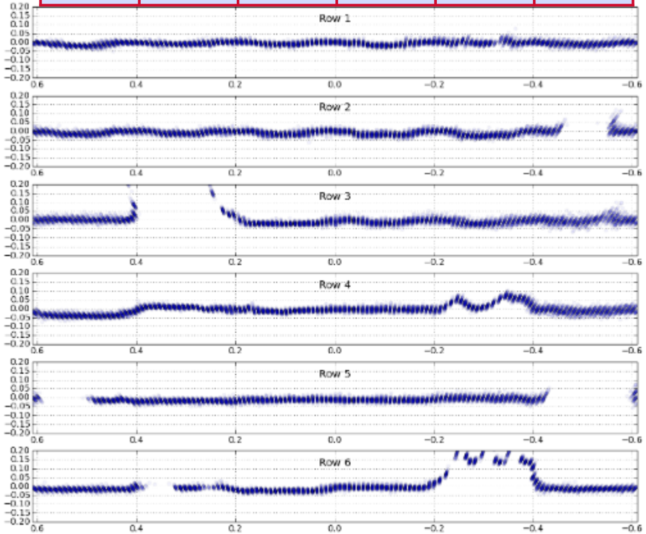

High IR content of Halogen increase noise additional 40%

ToF Camera : Lighting

300 lux : Fluorescent

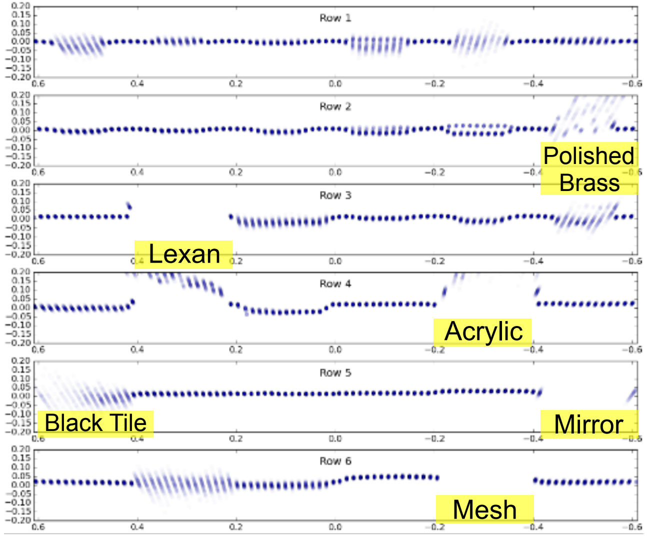

Higher range noise on dark objects

5000 lux : Halogen

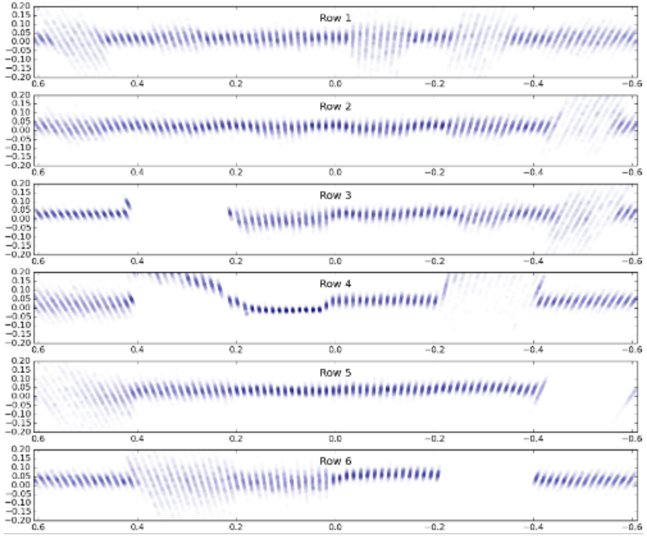

High IR content of Halogen increases noise by 2x to 5x

Stereo Camera : 7cm Baseline Range Accuracy

Carnegie Robotics

0.4 meter

4	4	4	4	4	4

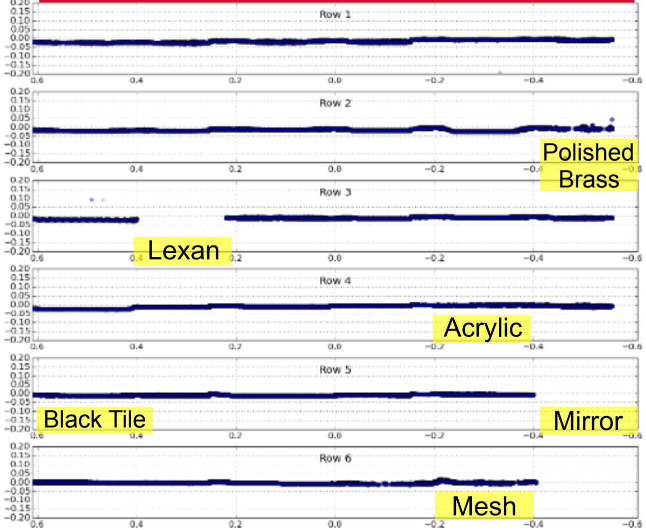

1.2 meter

9	9	12	10	8	10

Std. dev. of first row (mm)

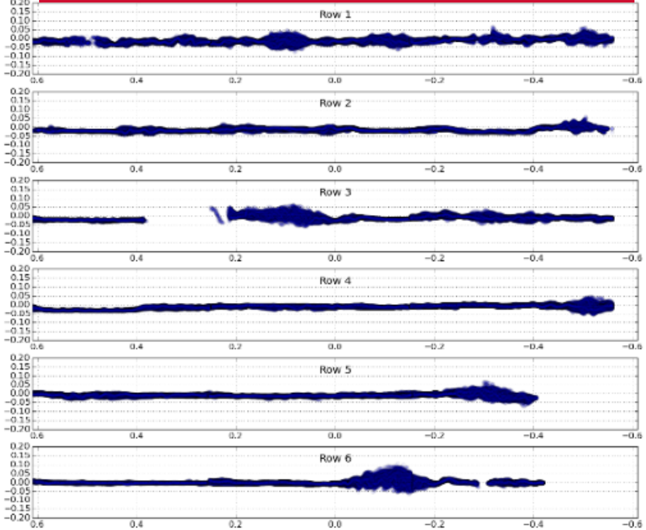

Higher range noise on low-texture objects

3D LIDAR : Error Changes Over Time

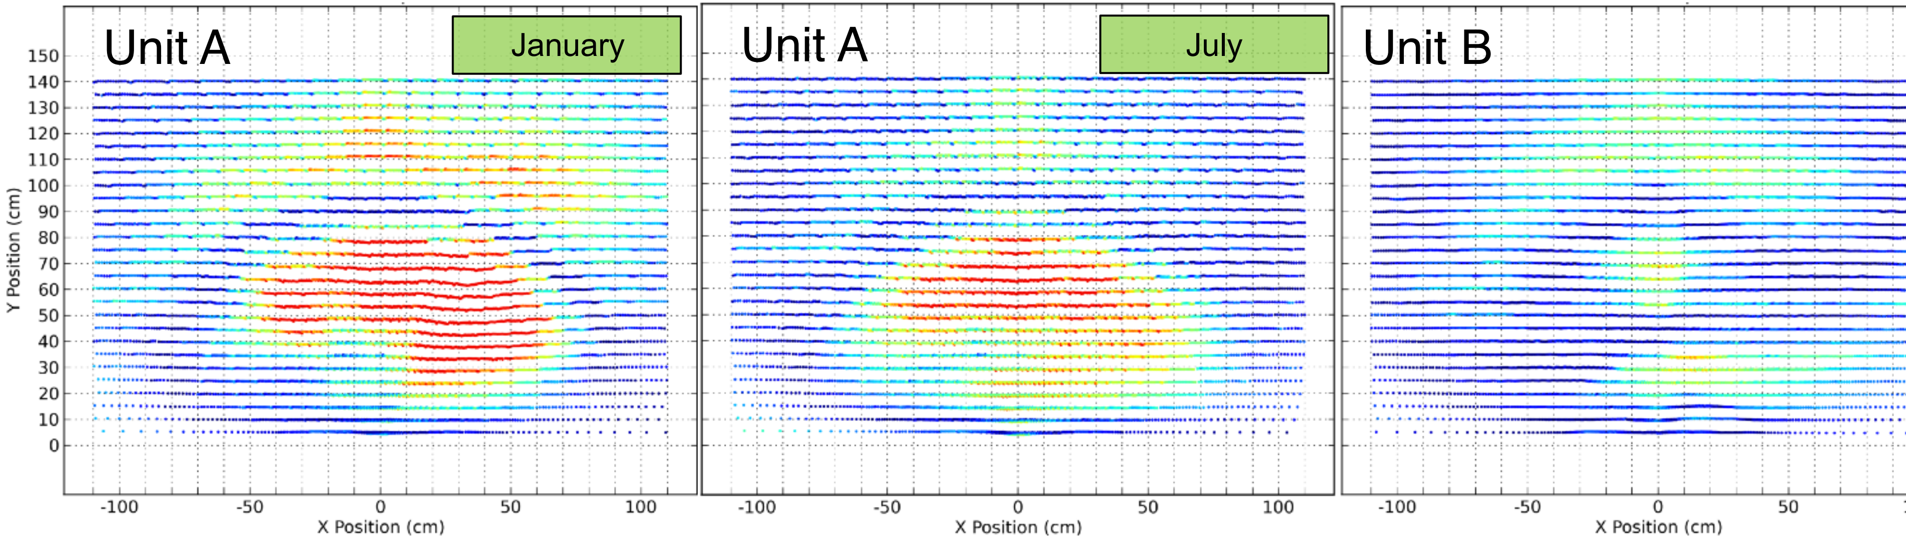

Error to Color Mapping

0.00 cm : Blue
0.50 cm : Green
1.00 cm : Yellow
1.25 cm: Orange
1.50+ cm: Red

- Every laser (even same SKU) is different.
- Lasers change over time.
- Error car vary by emission angle & range.

3D Resolution Testing

- In some applications XYZ resolution matters as much as depth accuracy.

- How do you measure 3D resolution? Take cues from 2D camera testing.
  - Raised Relief 3-D Resolution Target
  - Derived from USAF 1951 resolution target
  - THE standard image quality test for 60 years

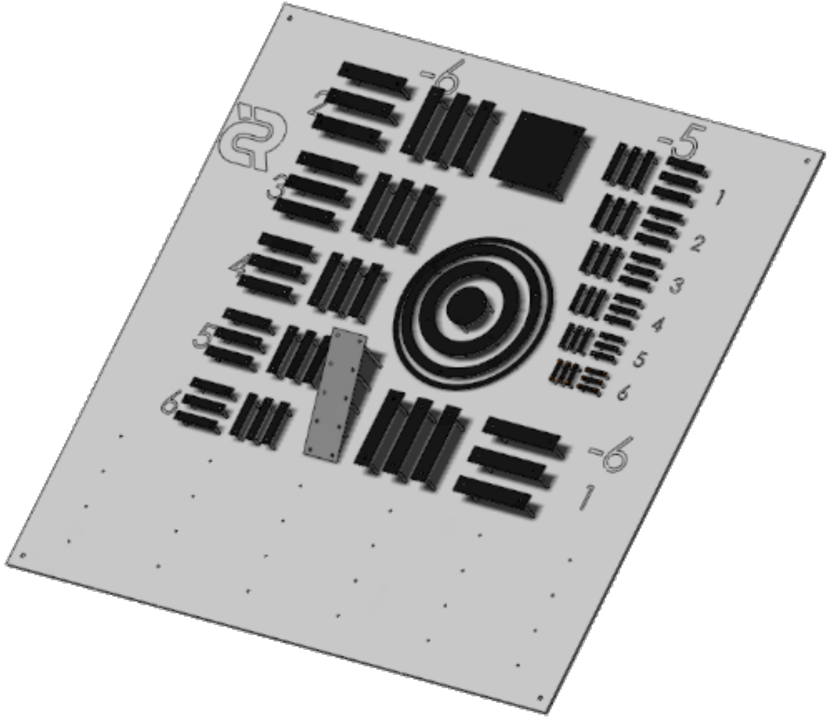

3D Resolution Testing Process

- Capture Image
- Identify and Isolate Target

- Outputs:
    - Colorized Depth Image
    - Colorized Depth Error Image
    - RMS error on features and between features for each grouping size

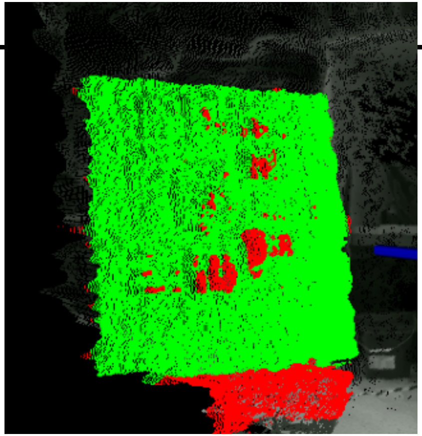

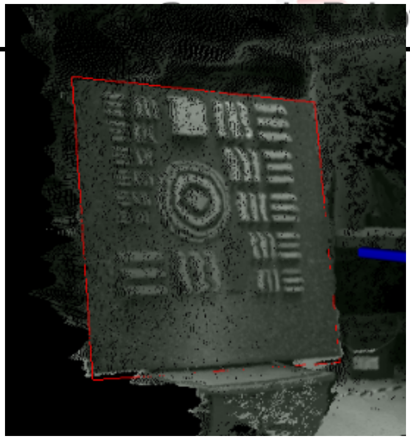

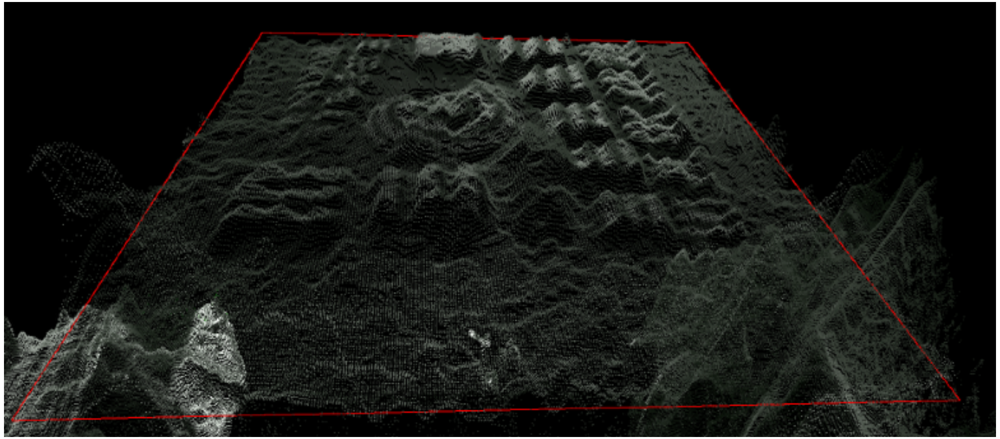

3D Resolution : Industrial ToF Camera

Carnegie Robotics

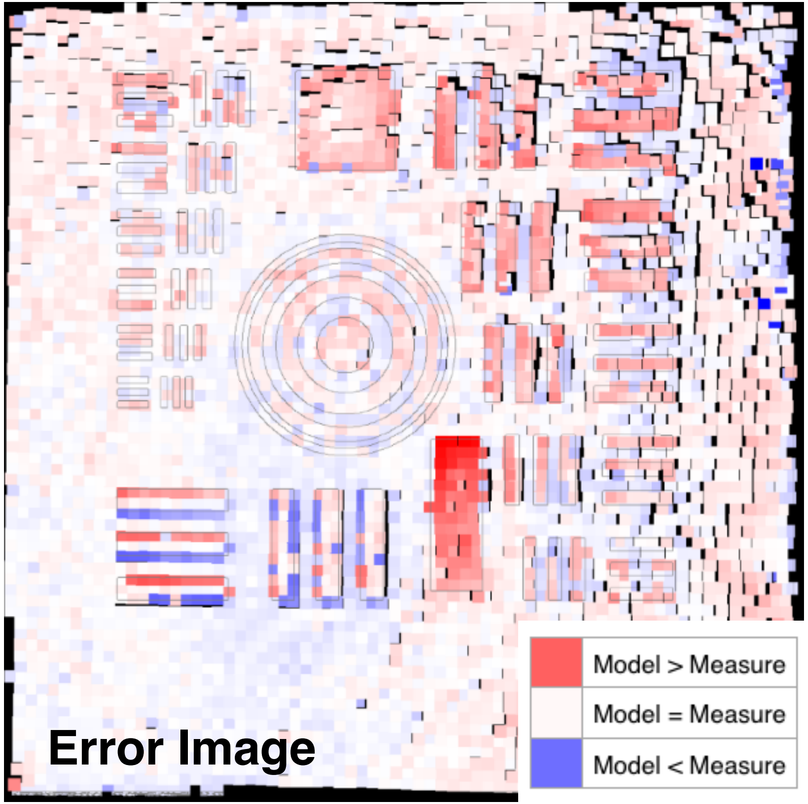

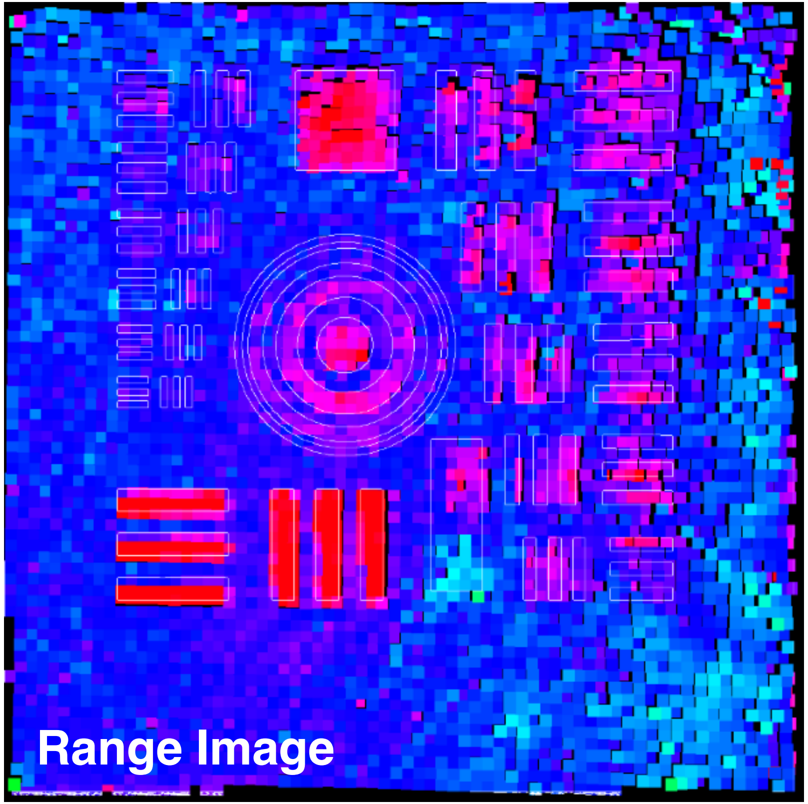

3D Resolution : Stereo

Carnegie Robotics

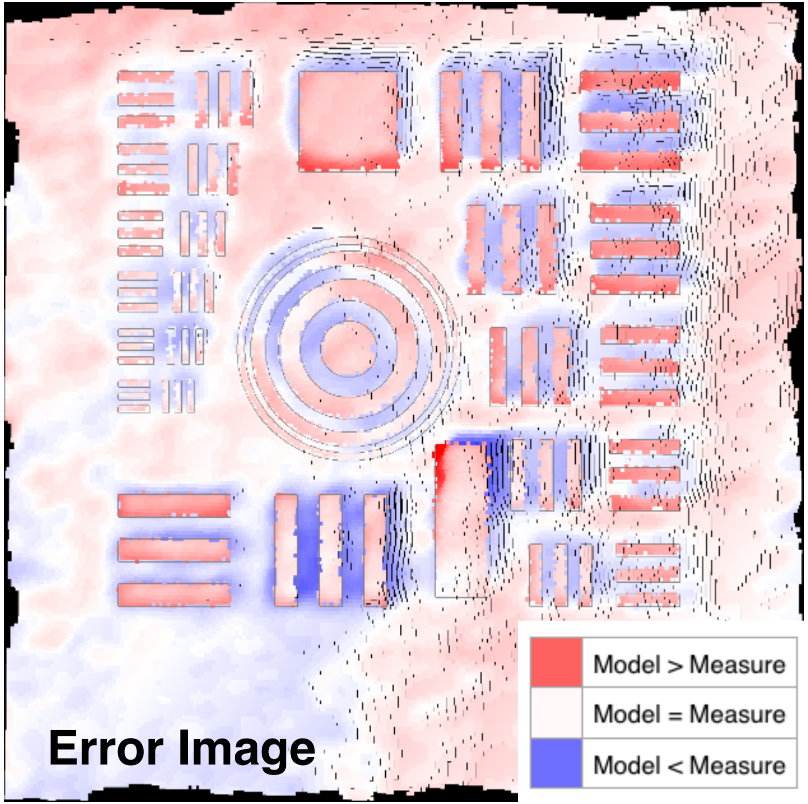

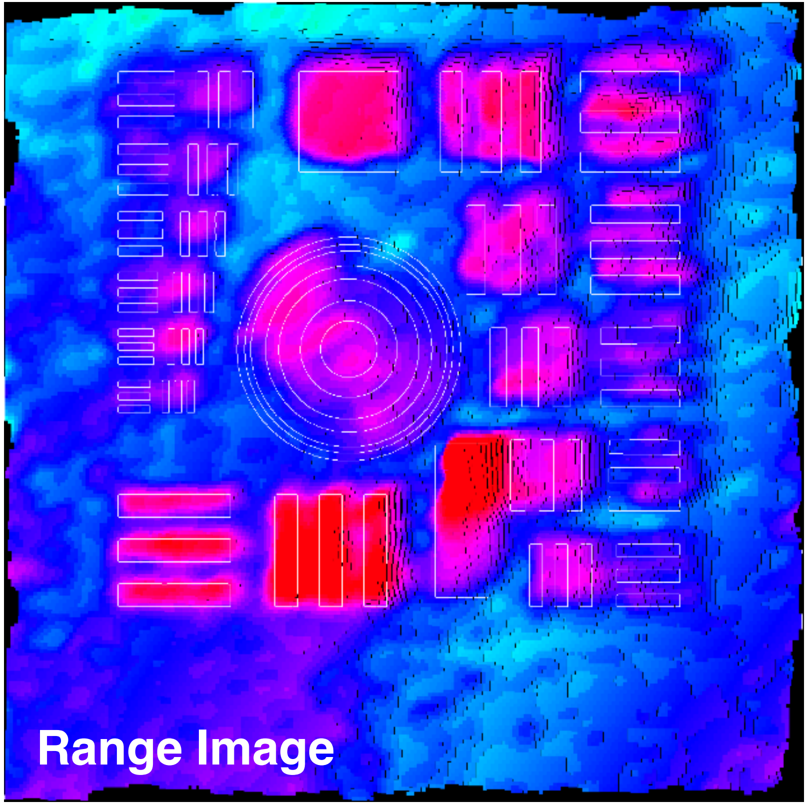

Take Aways

Carnegie Robotics

- There is no perfect sensor. They **all** have limitations and edge cases.

- It is critical to understand the performance & non-performance criteria that matter for your application.
  - Those weights must drive your sensor selection process.
  - You may have to sacrifice one or more lesser characteristics to meet the ones which really matter to you.

- Look for published test data. If none is available:
  - Ask the sensor manufacture,
  - Collect it yourself, or
  - Turn to those who have past experience in integration and downstream software.

Resources

Carnegie Robotics

- Carnegie Robotics:
  - http://carnegierobotics.com/roscon-2017
  - http://carnegierobotics.com/support

- Companies & Products:
  - Carnegie Robotics MultiSense
  - IFM O3D
  - SICK 3vistor-T
  - pmd pico
  - ZED Camera
  - Orbbec Persee
  - Structure Sensor
  - SICK TiM / LMS LIDARs
  - Velodyne 3D LIDARs
  - Hokuyo 2D & 3D LIDAR
  - Quanergy 3D LIDAR