TUNING D435 AND D415 CAMERAS FOR OPTIMAL PERFORMANCE

Presenter: Anders Grunnet-Jepsen, Ph.D.
CTO for Intel RealSense Group
May 2018

WHAT YOU WILL LEARN TODAY

- Choose the best camera for the job

- Operate at the optimal Camera Resolution

- Make Sure Image is Properly Exposed

- Verify Performance regularly on a Flat Wall or Target

- Play with the Depth Settings

- Use Post Processing

- Change the depth units

- Use an external projector

- Use the left color camera

- Change the Background

- Change IR Filter or background illumination

- Avoid Repetitive Structures

- Use Sunlight, but avoid glare

- Use Multiple Cameras

GET THE MOST OUT OF YOUR REALSENSE D415 AND D435 DEPTH CAMERA

(intel)

CHOOSE BEST REALSENSE™ DEPTH CAMERA FOR THE JOB

D415 and D435 represent two design points in a vast product landscape.

Image Sensor	OV2740	OV9282
Active Pixels	1920 × 1080 Color	1280 X 800 Monochrome
Sensor Aspect Ratio	16:9	8:5
Baseline	55mm	50mm
F Number	f/2.0	f/2.0
Focal Length	1.88mm	1.93mm
Filter Type	IR Cut – D400, None – D410	None
Focus	Fixed	Fixed
Shutter Type	Rolling Shutter	Global Shutter
Signal Interface	MIPI CSI-2, 2X Lanes	MIPI CSI-2, 2X Lanes
Horizontal Field of View	69.4	91.2
Vertical Field of View	42.5	65.5
Diagonal Field of View	77	100.6
Distortion	<=1.5%	<=1.5%

- **D415**
  - Rolling shutter, 2MP color stereo
  - Narrower FOV
  - Offers high-depth resolution when the object size is small and more precise measurements are required.
- **D435**
  - Global image shutter, 1MP monochrome stereo
  - Wide field of view
  - Accurate depth perception when the object is moving or the device is in motion, and it covers more field of view, minimizing blind spots.

(intel)

# DEMO DEPTH CAMERAS

UPDATE TO LATEST FW AND SW

D415   D435

1. **Update FW:**

    1. https://realsense.intel.com/intel-realsense-downloads/#firmware

2. **Get latest version of RealSense SDK (aka "LibRS")**

    1. https://realsense.intel.com/sdk-2/#install

    2. Start with RealSense Viewer

    3. Look at all sample code

3. **Make sure you have a PC with USB3 port**

    1. The RealSense cameras will support USB2 but with much diminished performance
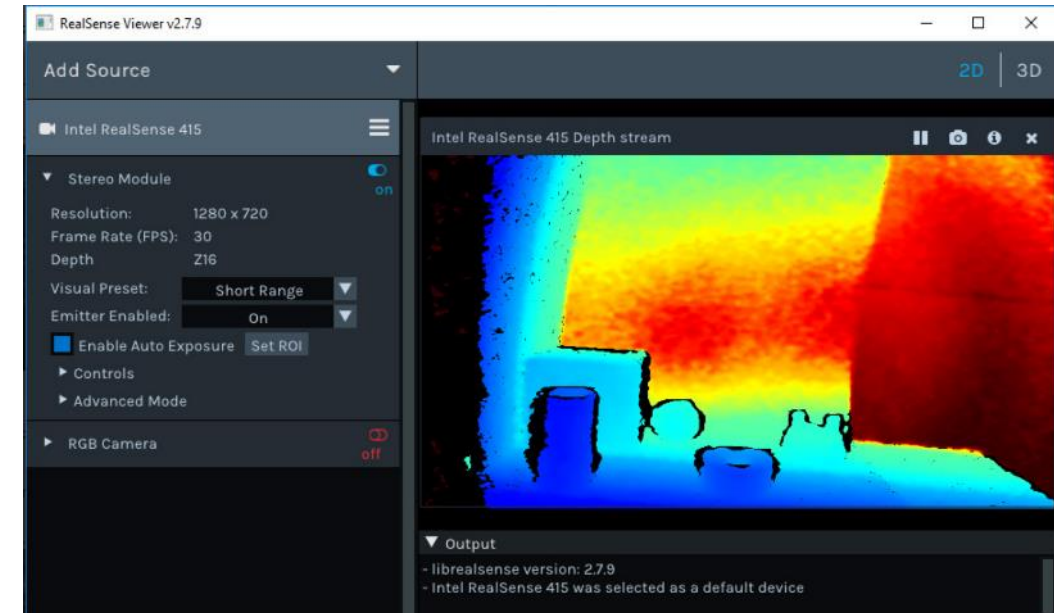
OPERATE AT THE OPTIMAL CAMERA RESOLUTION

1. **Stereo depth sensors derive their depth ranging performance from the ability to match positions of objects in the left and right images.**

    1. Higher input resolution + better input image → the better the depth precision.
    2. Lower resolutions can be used but will degrade the depth precision.

2. **Sometimes a lower resolution is required by your application for computational reasons**

    1. Use post-processing algorithms to immediately subsample (decimate) the higher resolution depth map and color image AFTER the data has been received.
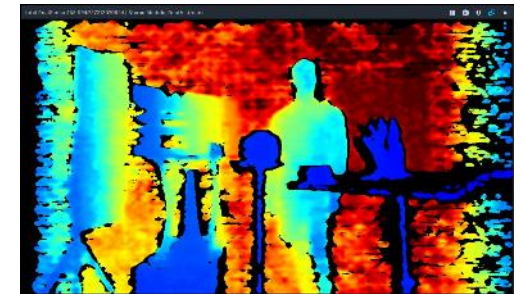
3. **Main Reasons you may choose to set the D435/D415 Cameras into a lower resolution modes:**

    1. Reduce the minimum operating range, which will scale with the x-y resolution.
    2. Reduce the bandwidth on the USB3.0 bus, to allow, for example, more cameras to operate at the same time, if the bandwidth limit has been reached.
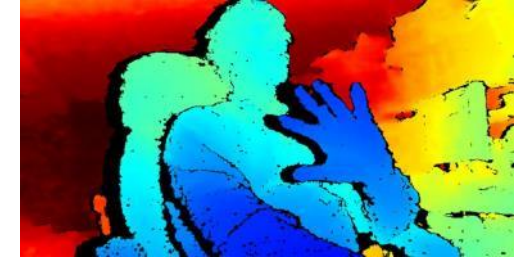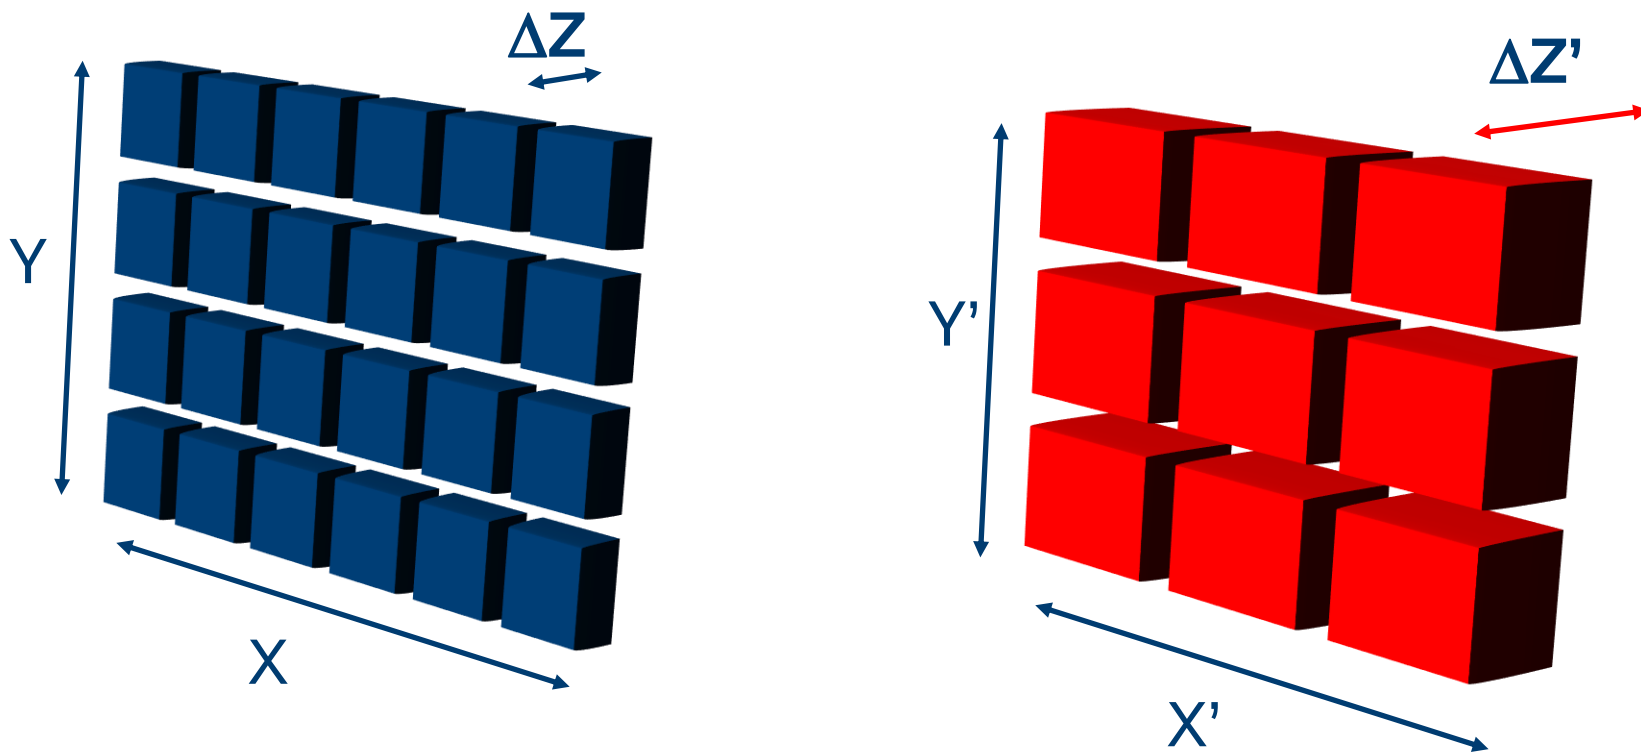

D435: 848x480
MinZ=17cm

D415: 1280x720
minZ=45cm

OPERATE AT THE OPTIMAL CAMERA RESOLUTION

DECREASING XY Resolution will INCREASE Z error

DEMO OF MIN-Z

UNDERSTAND MIN-Z

$$MinZ(mm) = focal\ length(pixels)\ x\ Baseline(mm)/126$$

$$where\ focal\ length(pixels) = \frac{1}{2}\frac{Xres(pixels)}{tan(\frac{HFOV}{2})}$$

1. **Minimum Operating Distance (MinZ) changes with resolutions**
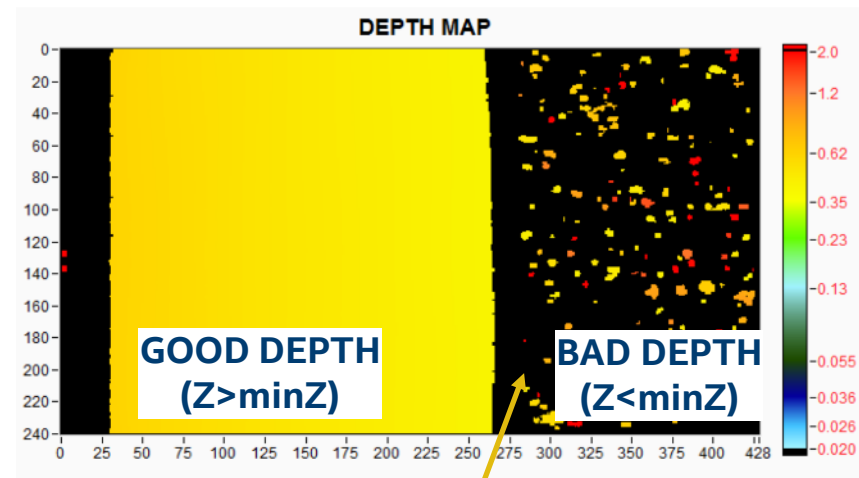   1. It scales linearly with the output Xres

2. **Shorter than MinZ there is depth noise**
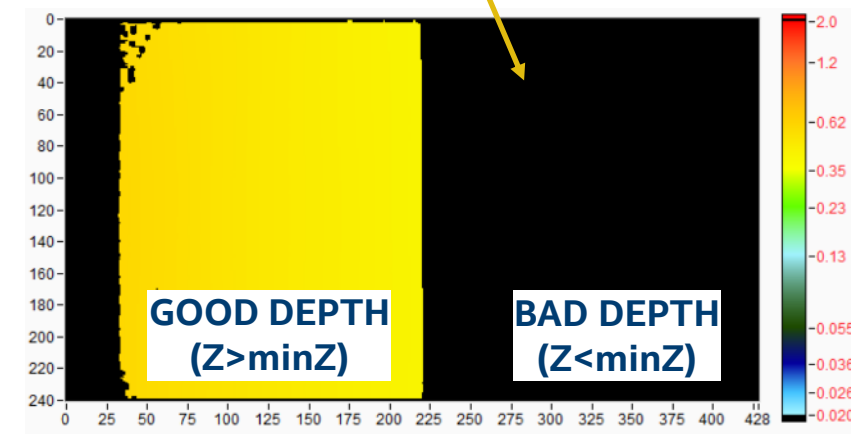   1. This can be mitigated with different depth settings

3. **MinZ can be shifted using "Disparity Shift":**
   1. Increasing Disparity shift above 0, will create a MaxZ
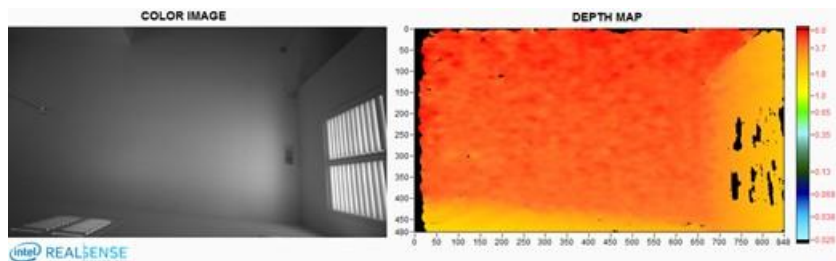
Looking at Slanted Plane

DEPTH MAP

GOOD DEPTH (Z>minZ)   BAD DEPTH (Z<minZ)

"MinZ Noise" can be reduced using depth settings.
Above= "High Density", Bottom="High Accuracy" preset

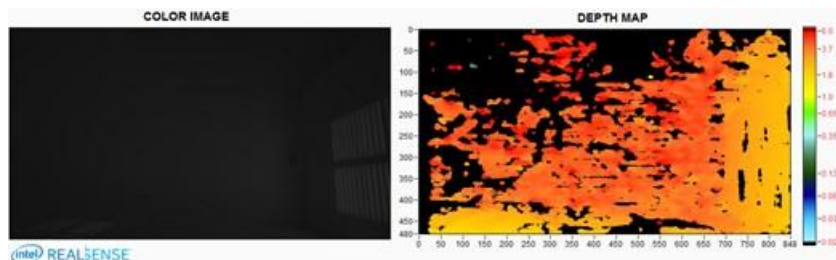GOOD DEPTH (Z>minZ)   BAD DEPTH (Z<minZ)

D415 at 1280x720 → MinZ is ~44cm
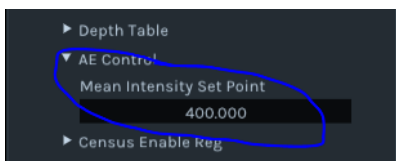**D435 at 848x480 → MinZ is ~17cm**

(intel)

MAKE SURE IMAGE IS PROPERLY EXPOSED

Example: Good Exposure

Example: Bad (under exposed)

Poor exposure is the number one reason for bad performance.

Check whether auto-exposure works well, or switch to manual exposure to make sure you have good color or monochrome left and right images.

Adjust EXPOSURE first for best results, with GAIN constant

1. Adjust the exposure (Ex: 33000 is 33ms.)

2. Keep the GAIN=16 (its lowest setting). Increasing gain tends to introduce electronic noise and, while the color image may look better, the depth quality will degrade.

3. Overexposing the images can be just as bad as under-exposing them, so be careful to find the right exposure.

Two Considerations when using Auto-Exposure feature: Region-of-Interest & Setpoint

1. When autoexposure is turned on, it will average the intensity of all the pixels inside of a predefined Region-Of-Interest (**ROI**) and will try to maintain this value at a predefined **Setpoint.**

2. Both the ROI and the Setpoint can be set in through API and RealSense Viewer:

3. Setpoint can be found under the Advanced Controls/AE Control.

4. ROI will only appear after streaming has been turn on. You will then be able to drag a ROI on the image.

DEMO EXPOSURE CONTROL

VERIFY PERFORMANCE REGULARLY ON A FLAT WALL OR TARGET

- **Test on Textured flat wall:** *This is the best way to verify performance.* Use an Intel recommended target like the one shown.

  - The pattern may be projected onto the wall or printed and attached to a flat, rigid surface.

  - Turn OFF the projector. Point to the wall and measure the RMS error of a plane-fit.

  - The **RealSense Depth Quality Tool** will provide metrics including RMS error. Read white paper:
    - https://www.intel.com/content/dam/support/us/en/documents/emerging-technologies/intel-realsense-technology/RealSense_DepthQualityTesting.pdf

- **Test on White wall:**

  - Turn ON projector and adjust laser power (ex: 150mW). If you see laser point saturation, reduce laser power; if depth is sparse, try increasing laser power.

- **Recalibrate:** If the flat target test shows a degradation in performance (Central ROI RMS subpixel>0.2 pixels, bad fill rate) then it may be necessary to recalibrate the camera.

  - https://realsense.intel.com/intel-realsense-downloads/#cal

https://realsense.intel.com/intel-realsense-downloads
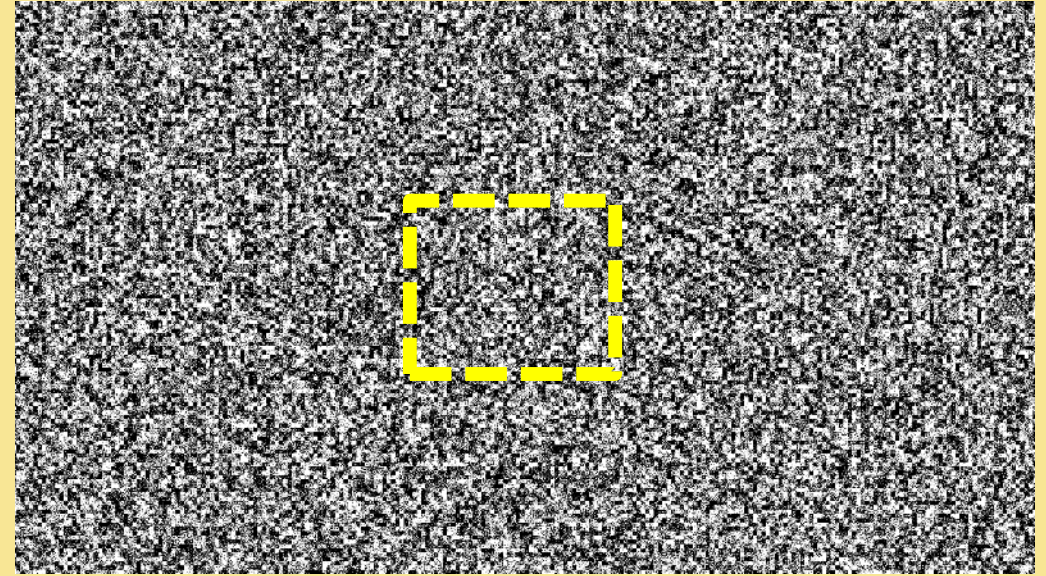
- This specific pattern was generated to show texture at many different scales, so should work across a fairly large range of distances.

- A well-calibrated camera should across central area produce **Subpixel RMS error < 0.1**
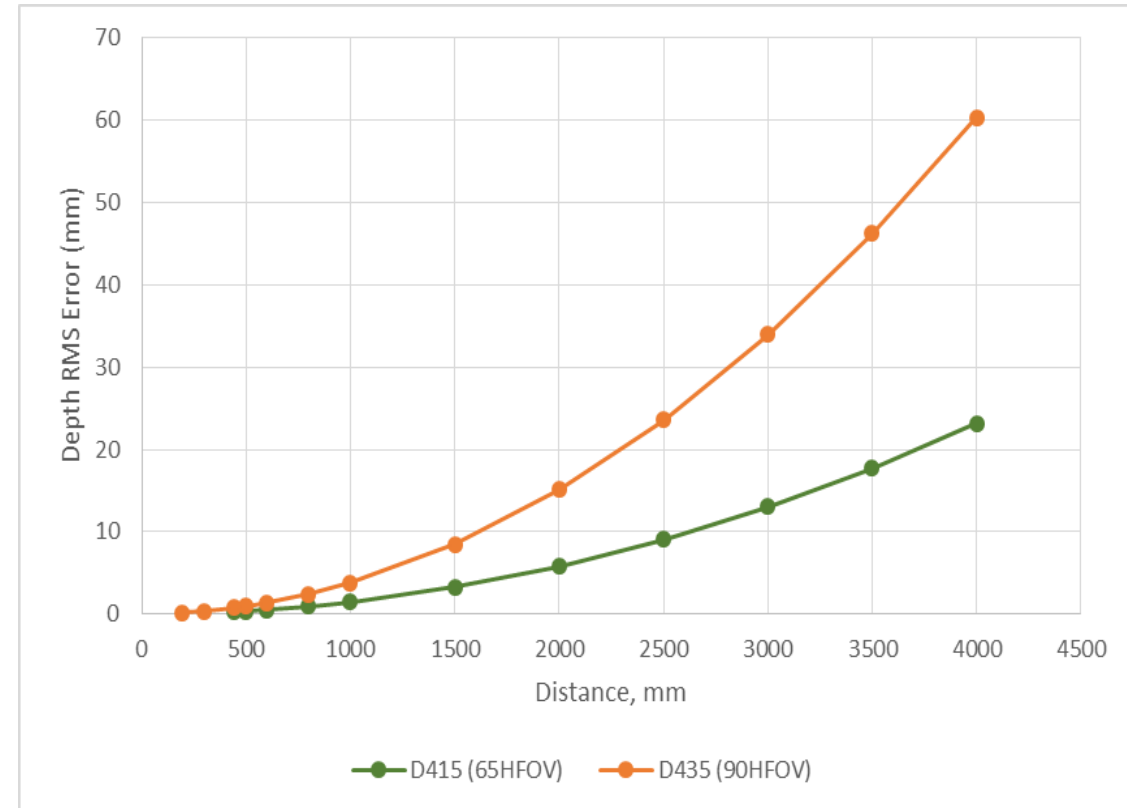
UNDERSTAND THEORETICAL LIMIT

$$\text{Depth RMS error(mm)} = \frac{\text{Distance(mm)}^2 \text{ x Subpixel}}{\text{focal length(pixels)x Baseline(mm)}}$$

$$\text{where focal length(pixels)} = \frac{1}{2}\frac{\text{Xres(pixels)}}{\tan\left(\frac{\text{HFOV}}{2}\right)}$$

- **Understand theoretical limit:** The RMS error represents the depth noise for a localized plane fit to the depth values.

  - For a well textured target you can expect to see Subpixel RMS < 0.1

  - On passive textured target, expect you ~30% better RMS with the laser turned OFF (due to residual laser speckle).

- **The depth noise**:

  - Tends to manifest itself as "bumps" on a plane, looking a bit like en egg carton. This is expected.

The graph is obtained using subpixel=0.08:
D415 with HFOV=65deg, Xres=1280, and baseline=55mm,
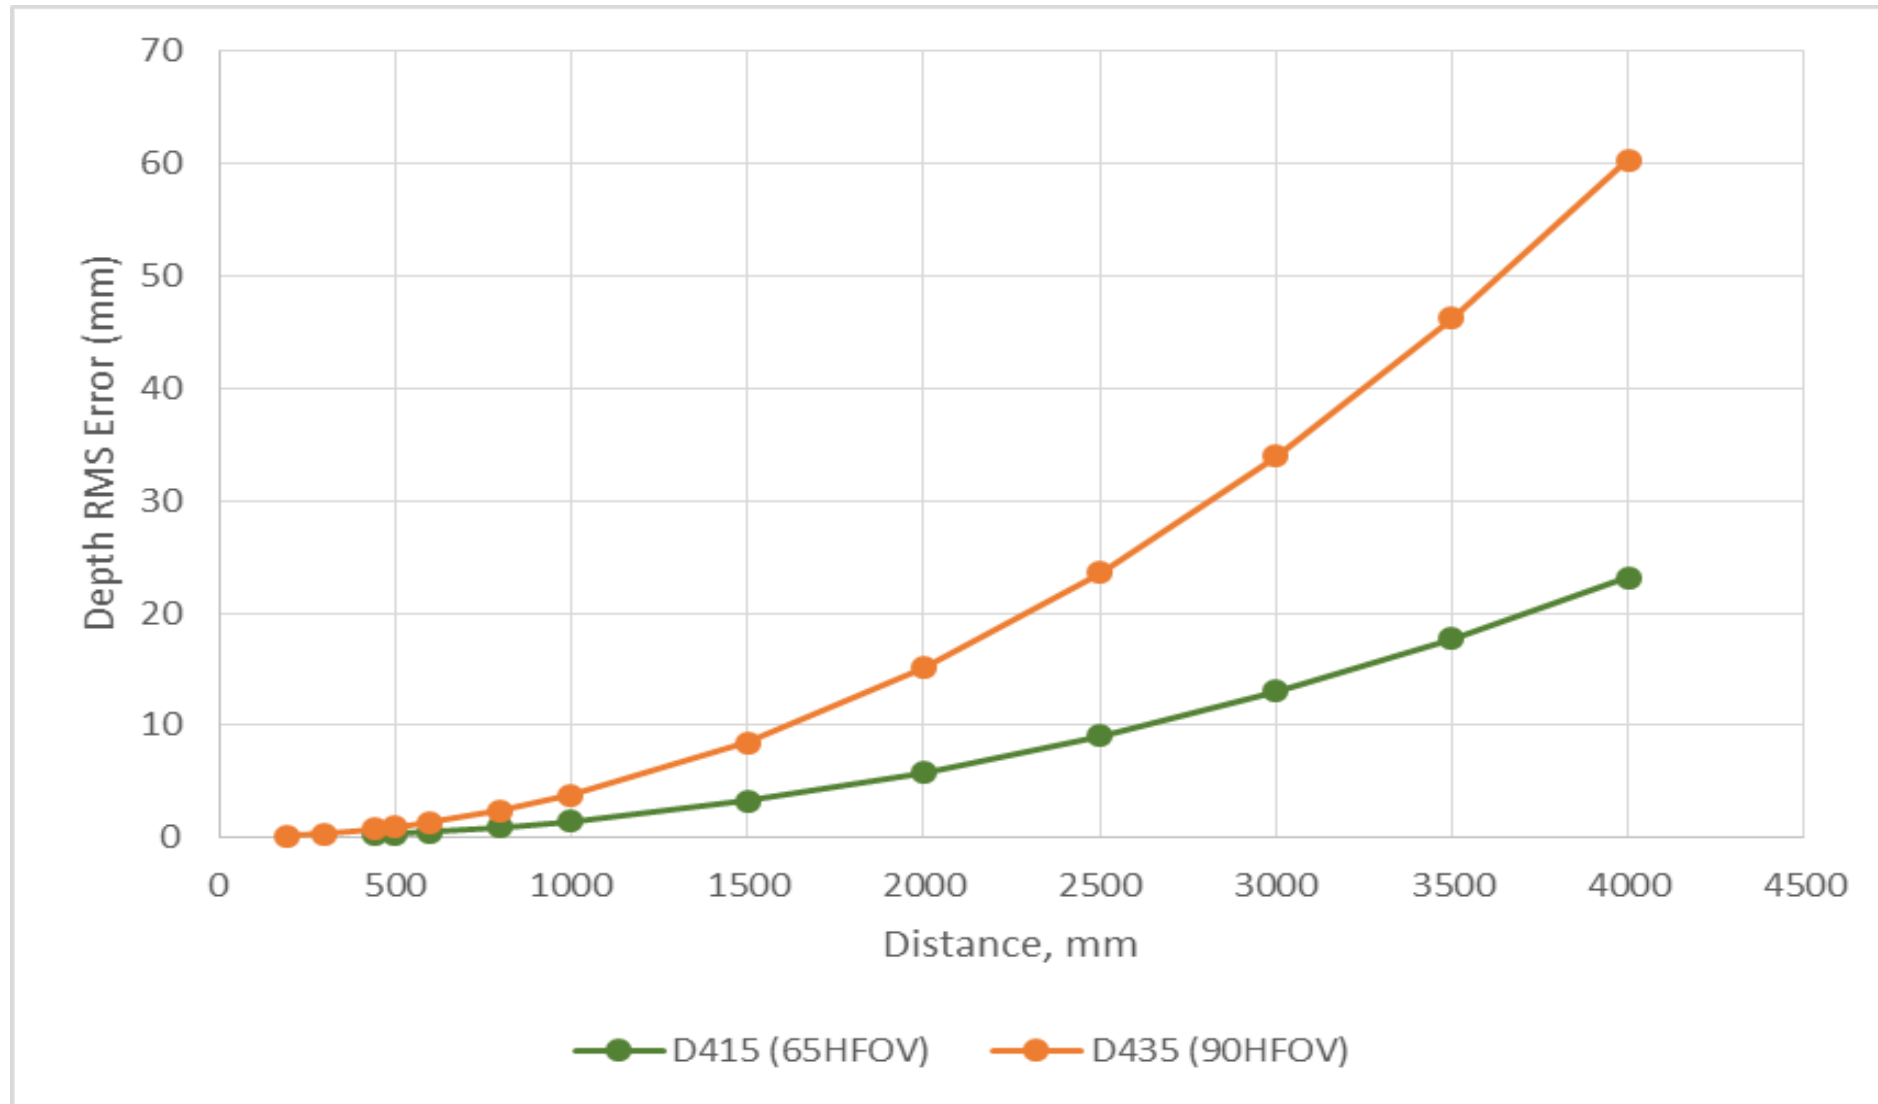D435 with HFOV=90deg, Xres=848, and baseline=50mm

UNDERSTAND THEORETICAL LIMIT

MOVE AS CLOSE AS POSSIBLE TO YOUR OBJECT

1. Cameras works best near!

   - Depth error scales as the *square of the distance away*.

2. *Where at all possible*, try to get as close to the object as possible (but >minZ)

3. Consider using **Multiple Cameras** to increase effective FOV.

4. To reduce MinZ, you will need to

   1. Reduce Resolution of the Depth imagers → minZ scales linearly with Xres, or

   2. Try using the **"Disparity shift"** parameter

      1. It is best to ensure that **no objects farther be away than MaxZ**. Ex: depth camera pointing at the table surface.

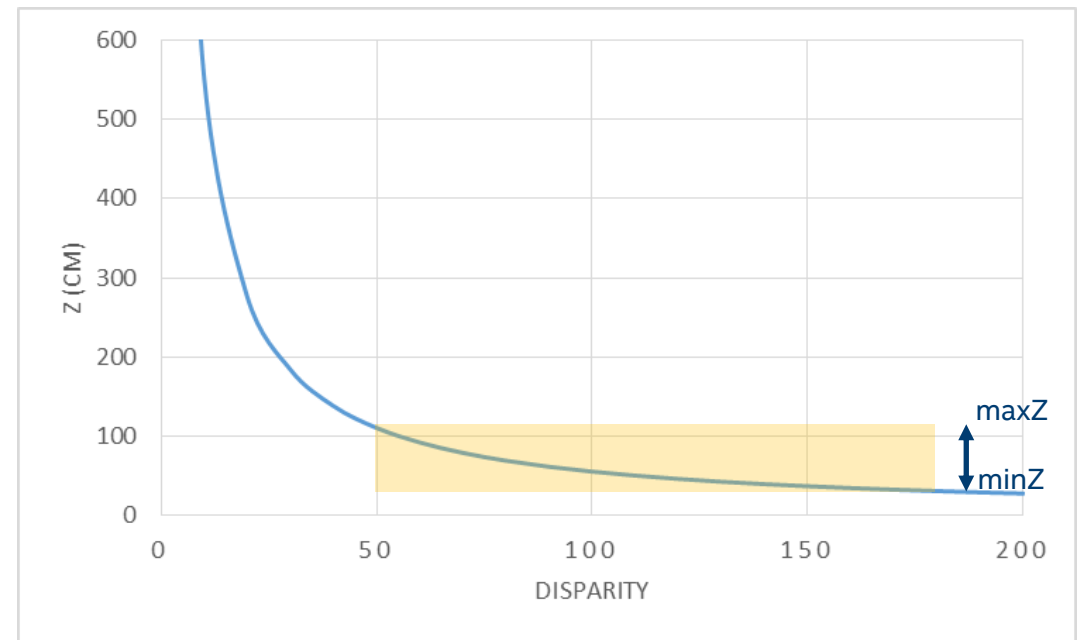   2. Using Disparity Shift will add a **MaxZ** (see next page)
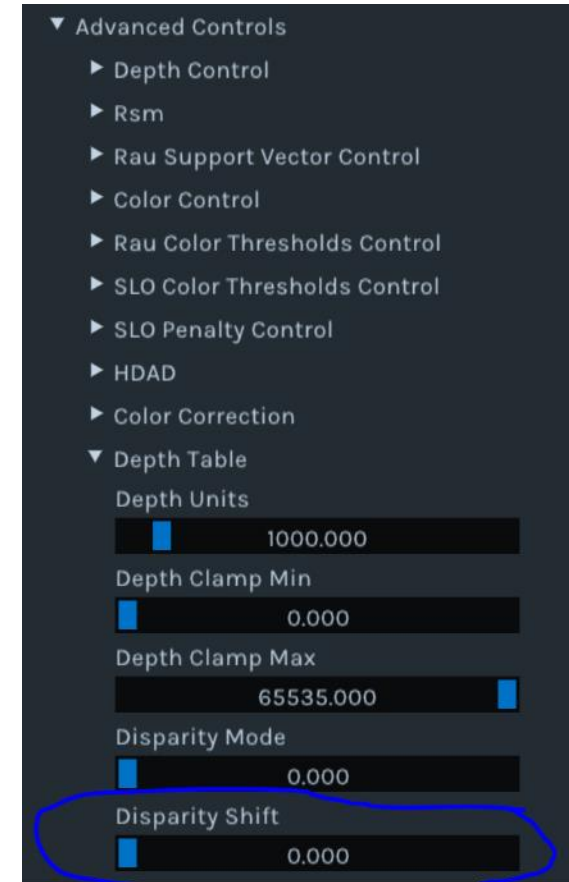
RANGE LIMITS

1. Stereo cameras will see INFINITELY far...provided disparity shift=0

2. The RealSense Camera can search a **126 disparity range** → MinZ and MaxZ can change with disparity shift

Disparity Shift=0 → **MinZ=45cm, MaxZ=Infinity**
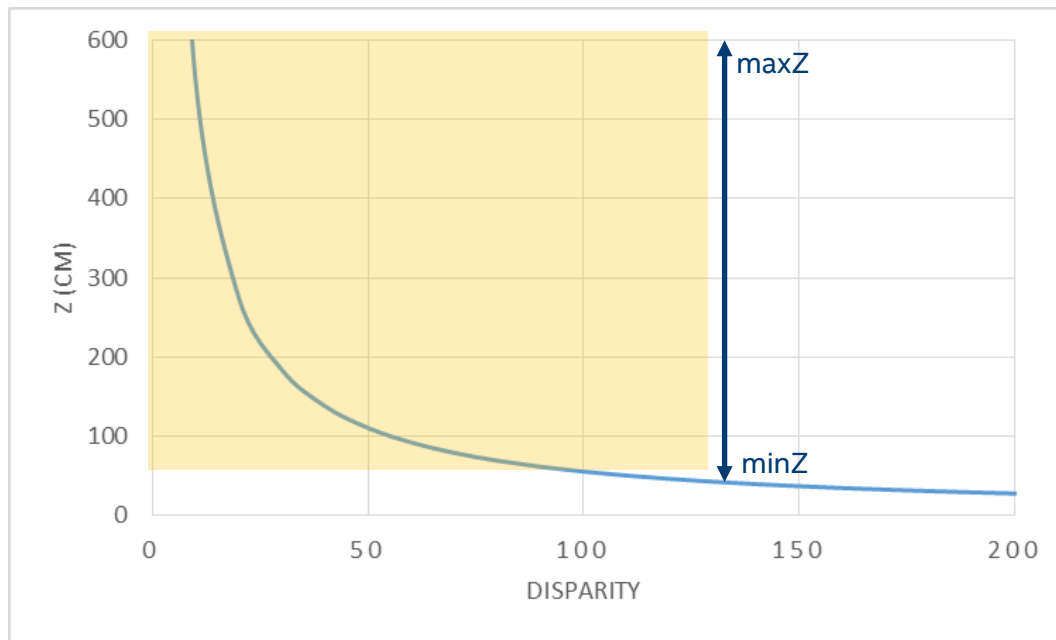
Disparity Shift=50 → **MinZ= 30cm, MaxZ=110cm**

Examples for D415

PLAY WITH THE DEPTH SETTINGS

Tip!

Preset	Use Cases recommended for usage
High Density	Higher Fill factor, sees more objects. (Ex. BGS and 3D Enhanced Photography, Object recognition)
Medium Density	Balance between Fill factor and accuracy.
High Accuracy	High confidence threshold value of depth, lower fill factor. (Ex. Object Scanning, Collision Avoidance, Robots)
Hand	Good for Hand Tracking, Gesture recognition, good edges
Default	Best Visual appeal, Clean edges, Reduced PointCloud Spraying

Note: Some of the presets are different based on the resolution and some are the same for all resolutions. The resolution range is:
•High Res - 1280x720 or 848x480
•Med Res - 640x480 or 640x360
•Low Res - 480x270 or 424x240

Note: Given that there is a complex interplay between all parameters, we currently use machine learning to globally optimize for different usages. For this reason we *do not* plan on providing any future descriptions of *"what each parameter does"*.

The depth calculations in the D4 VPU are influenced by over 40 different parameters.

There are currently 5 different Depth Presets or configurations built into the RealSense Viewer.

- "Default" setting is recommend to start

- "Higher density" is good when more fill is required

- "High accuracy" is good for minimizing any hallucinations.

The parameters or Depth Presets can be directly accessed through the "Advanced Mode" api.

(intel)

USE POST PROCESSING

		Benefit
Intelligent Sub-sampling	We recommend using "non-zero mean" or "non-zero median" for a pixel and its neighbors. (ex only look at non-zero values) All stereo algorithms involve some convolution operations, so reducing the (X, Y) resolution after capture is usually very beneficial for reducing the required compute for higher-level apps. Ex: Reducing by X and Y by 3x, will and general decrease subsequent compute load 9x.	• Less points→Less CPU • Rudimentary hole filling • Some smoothing of data • Can help point cloud visualizing
Temporal filtering	*Whenever possible,* use some amount of time averaging to improve the depth, making sure not to take "holes" (depth=0) into account. We recommend using an IIR filter, but ONLY if noise is BELOW a distance threshold (or disparity step size) In some cases it may also be beneficial to use "**persistence**", where the last valid value is retained indefinitely or within a certain time frame.	• Improves depth noise • Persistence can help "see farther".
Edge-preserving filtering	This will smooth the SPATIAL depth noise, retaining edges while making surfaces flatter. However, again care should be taken to use parameters that do not over-aggressively remove features. We recommend doing this processing in the disparity domain (i.e. the depth scale is 1/distance), and experimenting by gradually increasing the step size threshold until it looks best for the intended usage. Another successful post-processing technique is to use a Domain-Transform filter guided by the RGB image or a bilinear filter. This can help sharpen up edges for example.	• Smooth spatial depth noise ("bumps"), • Helps sharpen edges
Hole-filling	Simple fill algorithms can be used to fill in the "holes" with best guesses based on neighboring values, or the RGB image. This should be used sparingly, as it often leads to wrong values.	For applications are intolerant of holes in the depth – fills holes with best guesses (ex: for depth enhanced photography)

DEMO OF DEPTH SETTINGS & POST-PROCESSING

# D415 INDOORS (AT ~1M) WITH SOME POST-PROCESSING

Click on video to play

(intel)

CHANGE THE DEPTH STEP SIZE

By default the D4 VPU provides **16bit depth** with depth unit =1000um (1mm). This means the max *reported* range will be ~65m.

It may be necessary to change the **"Depth Units"** (Advanced Mode), to see higher resolution (smaller steps), or longer (larger steps).

- With DepthUnit=5000um, Cameras will report range to ~325m.

- With DepthUnit=100um, Cameras will report depth to ~6.5m

  - When operating at very close range, the D4xx cameras can inherently deliver *depth resolution well below 1mm*. Reduce depth units to avoid quantization effects.
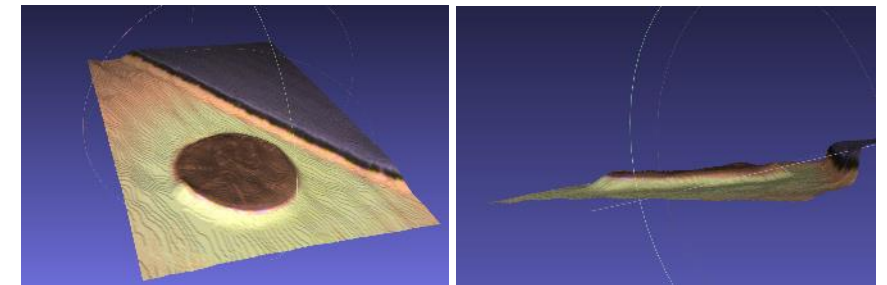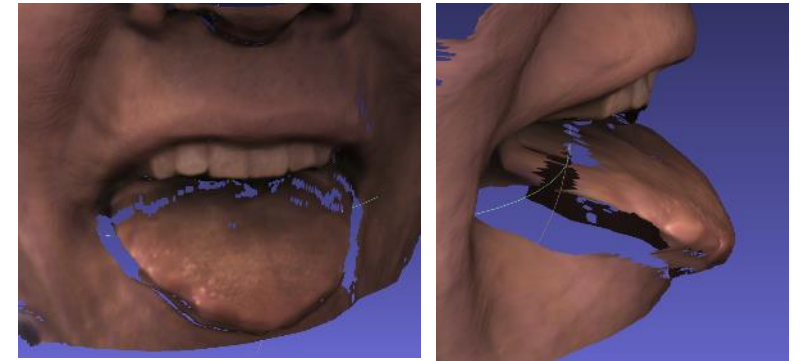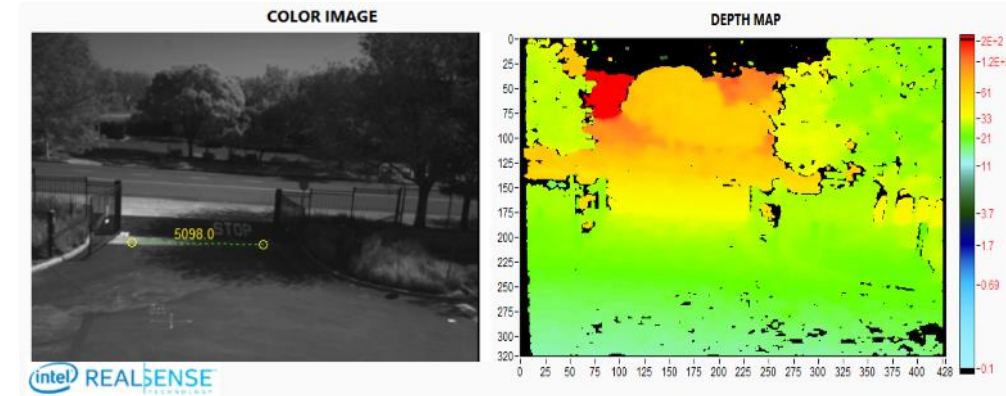
Near Range: While it is possible to get depth for really close objects by using disparity shift & reducing the X-resolution & reducing the depth units, there will be other factors that will ultimately limit near-range performance.
  i. The D4xx cameras are currently focused at 50cm to infinity, so will start to go out of focus well below 20cm.
  ii. Also, for really near objects that are not flat, like a wall, the views from the left and right cameras may be different at really close range and will therefore not give any depth range.

COLOR IMAGE DEPTH MAP

CONSIDER USING AN EXTERNAL PROJECTOR

When to consider **external** high-power projector:

- ✓ For large rooms that could benefit from additional illumination

- ✓ Add even more inherent texture (ex: to flat white walls)

- ✓ For fixed installations where low power consumption is not critical

- ✓ For objects that are far away

- ✓ Option to use non-laser illumination can improve depth

Tip!

Choosing a Projector and a Projector Pattern

1. **Projection pattern:** use somewhere between 5000 and 300,000 semi-random dots. We normally recommend patterns that are scale invariant. This means that if the separation between the projector and the depth sensor changes, then the depth sensor will still see texture across a large range of distances.

2. **Wavelength:** It is possible to use light in visible or near-infrared (ranging from ~400-1000nm). Ex: 850nm is invisible to the human eye while still being visible to the IR sensor

3. **Reduce laser speckle**: It is extremely important to take all means necessary to *reduce the laser-speckle* or it will adversely affect the depth. Non-coherent light sources will lead to best results in terms of depth noise, but cannot normally generate as high contrast pattern at long range.

4. **Location:** Projectors do not need to be co-located with the depth sensor. It is also acceptable to move or shake the projector if desired, but it is not needed.

5. **Eye-safe:** Make sure the projector is eye-safe.

PROJECTION PATTERNS

D435
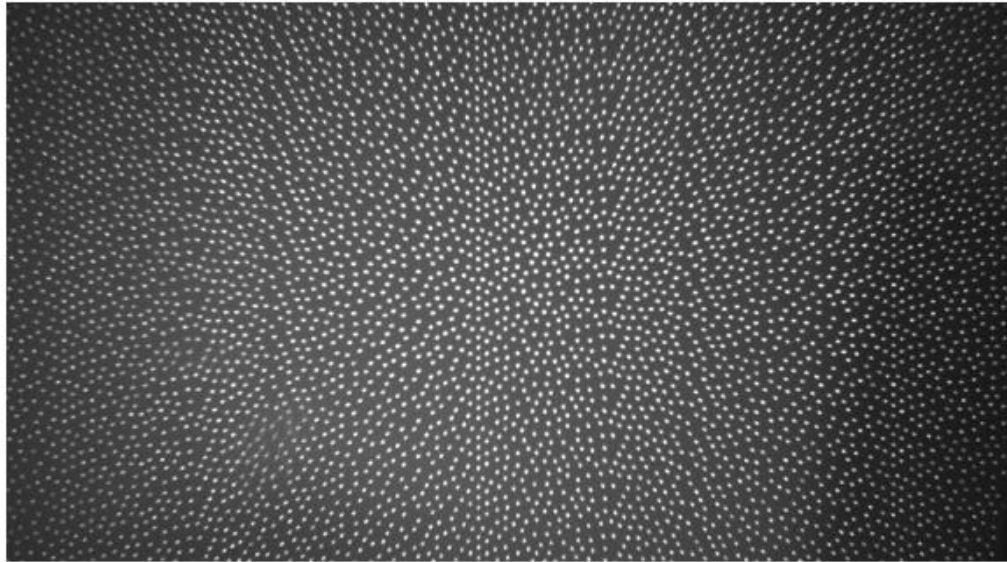
91x65 deg

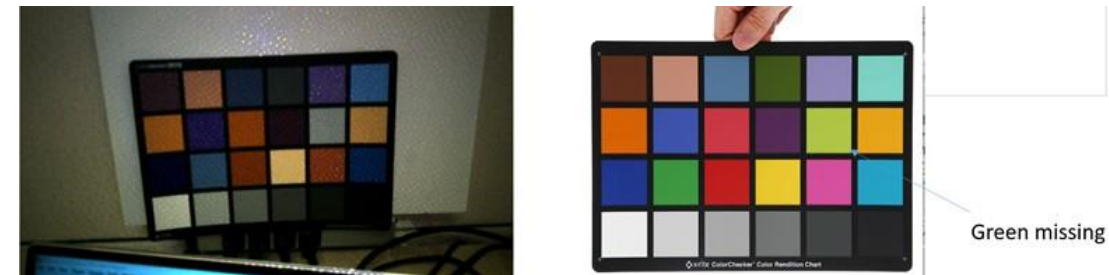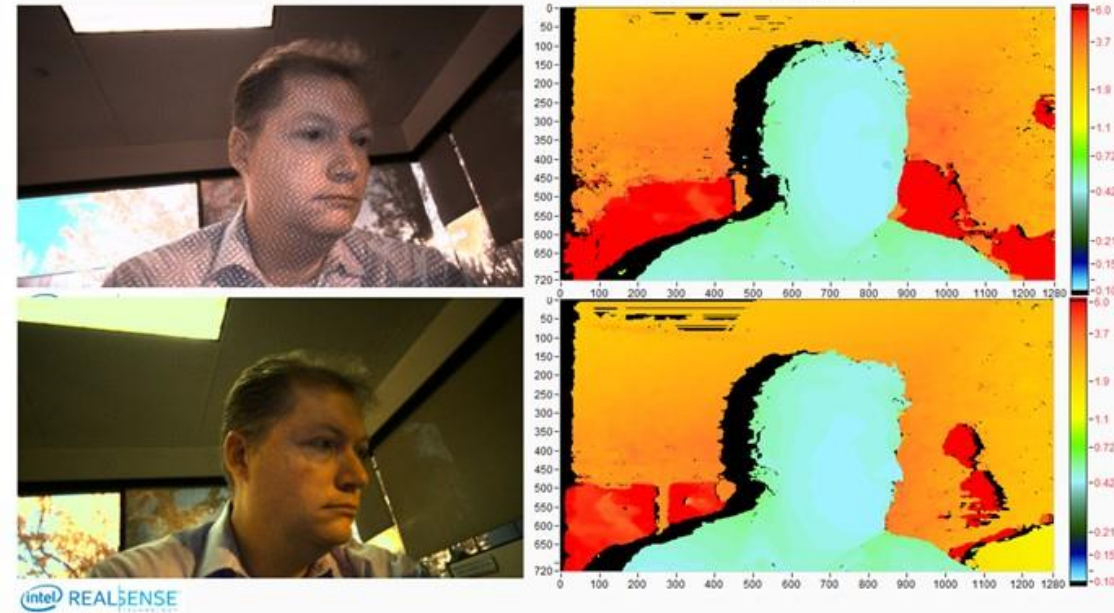**D415**

68x43 deg

USE THE LEFT COLOR CAMERA

1. **The Left Color:** The D415 camera has a built-in 3rd color camera that is processed through an ISP to give good color independent of the depth, BUT there are many reasons to use the *left camera* of the stereo pair instead (or in addition).

   1. Pixel-perfect alignment & calibration to depth

   2. Perfectly time-synchronized

   3. Requires no additional computational overhead to align color-to-depth

   4. Gives no additional occlusion artifacts, because it is exactly co-aligned. The main draw-back is that it will normally show the projector pattern if the projector is turned on.

2. **Removing the projector pattern:**

   1. **Turn on Projector:**

   2. **Use on-chip REMOVE IR feature**: For 1280x720 resolution, apply "color-correction" numbers shown here.


Reducing IR through Color Correction Matrix

intel REALSENSE

Green missing

D415 "REMOVE IR" ColorCorrection={0.520508, 1.99023, 1.50684, -2, -2, -0.0820312, 1.12305, 1.01367, 1.69824, -2, 0.575195, -0.411133}

(intel)

DEMO OF LASER PATTERN REMOVAL FOR D415

(will not work on Monochrome systems, like the D435)

CHANGE THE BACKGROUND

While many applications do not allow for changing the environment, some usages do. For example if the purpose is to examine parcels placed on a desk, it would greatly help the measurements if the desktop had a matt textured pattern. This is preferable to a white glossy background, like a white tabletop.
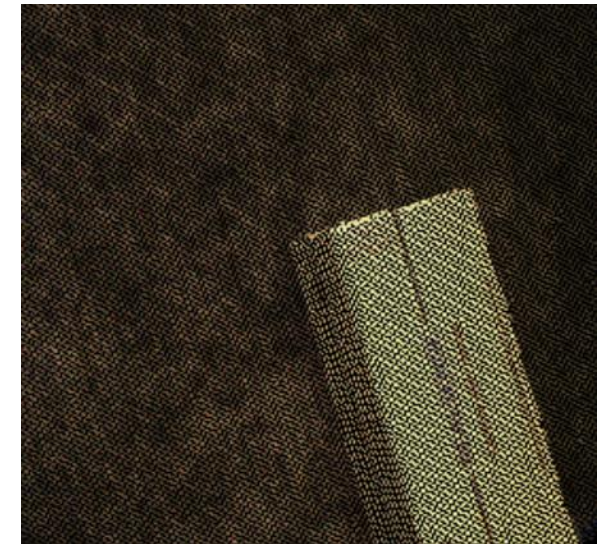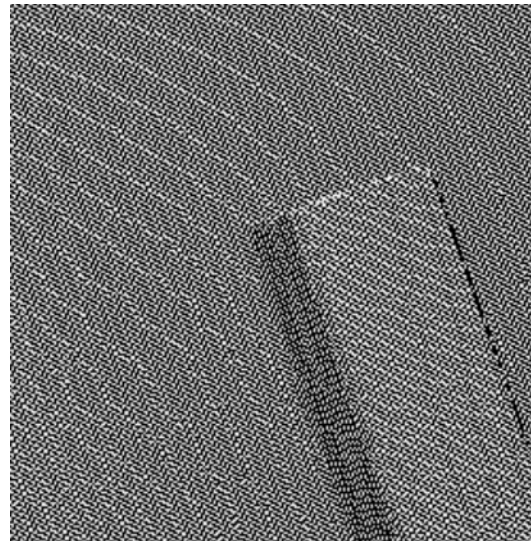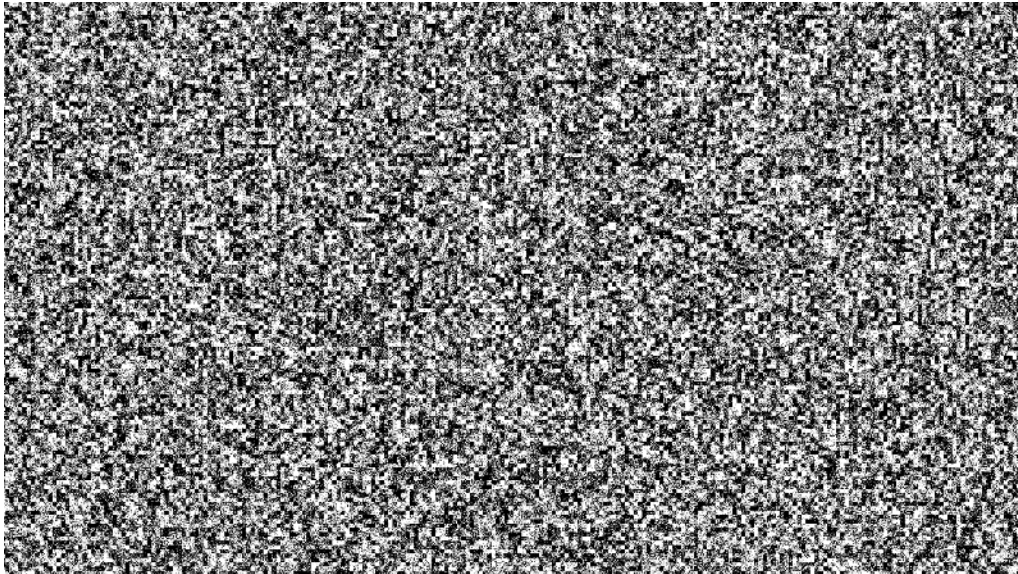
CHANGE IR FILTER OR BACKGROUND ILLUMINATION

- Both the D415 and D435 **see all light** in the visible and IR.

  - RealSense D4 VPU does a very good job of extracting depth from purely passive scenes, without any projector and actually benefits from using Red, Green, Blue images.

- **IR bandpass filter (Block visible):** If the purpose is to see white walls indoors at a long distance, it may be necessary to enhance the SNR of the projected laser pattern.

  - Increase laser power, or

  - Place an IR-bandpass filter in front of the Left+Right imager.

- **Pre-test:** Try turning off all room lights while leaving the projector on (emulating IR bandpass filter)

  - The "white wall" range can increase from ~3m to 10m this way.

Caution: When placing any optical windows in front of the D4xx cameras, it is important to make sure that no laser light from the projector reflects back into the imagers as this will adversely affect the depth. We recommend using separate windows for the left and right imagers and nothing in front of IR projector.

(intel)

AVOID REPETITIVE STRUCTURES

Stereo depth computation depends on determining matches between left and right images. If there are repetitive structures in the scene, like fences or wire-grids, this determination can become ambiguous.
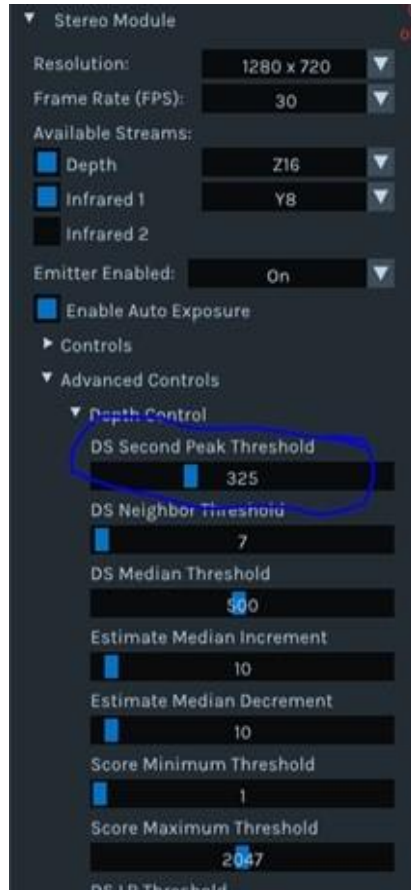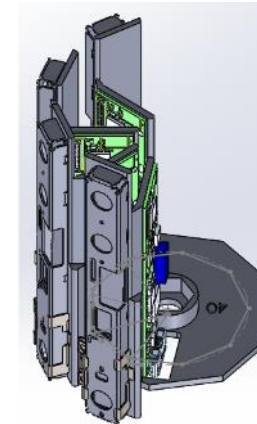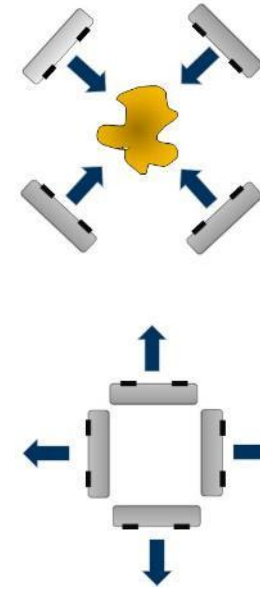
- **"DSSecondPeakThreshold"** - Advanced API setting that can be used to mitigate results of aliasing from repetitive structure. Increase its value and it will reduce the number of "false" matches resulting from aliasing.

- **Tilt:** Another mitigation strategy that may help is to tilt the camera a few degrees (e.g., 20-30 deg) from the horizontal.

USE SUNLIGHT, BUT AVOID GLARE

- **Sunlight can be good:** Most depth cameras degrade dramatically in sunlight. By contrast, both the D415 and D435 tend *to perform even better in bright light*.

  - Depth quality is directly related to the quality of the input images. Sunlight → better images. Sunlight reduces the sensor noise and tend to "brings out" the texture in objects.

  - Moreover, the exposure can be reduced to near 1ms which reduces motion artifacts as well.

- **Avoid glare:** Be careful to avoid lens glare when pointing at or near the sun. It is recommended that the lenses are shielded by baffles outdoors.

- **Exposure setting:** When operating a D4xx camera outside, it is important to be especially careful in regards to the auto-exposure, especially if the sun or reflections of sun are visible in the image.

  - By default the auto-exposure algorithm tries to keep the average intensity of the **whole image** within a certain range. If the sun is in that image, everything else will suddenly become black. For many applications, like autonomous robots or cars, it helps to simply change the Region-of-Interest of the auto-exposure algorithm to a smaller size, or specifically to the lower half of the image.

USE MULTIPLE DEPTH CAMERAS

- **Outward facing:** Use this configuration to increase Field-of-view, FOV

- **Inward facing:** Use this configuration to capture full 3D volume, ie back and front of objects simultaneously.

- **No interference:** D4xx cameras do not interfere with each other.

- **Simultaneous capture:**

  - **Time stamps:** The API provides timestamps and frame counters that can be used to capture within 1 frame

  - **HW sync:** To get to few microseconds sync, it is possible to link the cameras together via an external sync cable and tell one to be master and the others to be slave. (See separate white paper on this. The Multi Camera sync feature targeted to be ready in Q2 2018).

MULTI-CAMERA DEMO