Intel® RealSense™ LiDAR Camera L515

Intel® RealSense™ LiDAR Camera L515 user guide

Revision 1.0

July 2020

Contents

Figures

Introduction

The Intel® RealSense™ LiDAR Camera L515 is Intel's first release of a LiDAR camera enabling highly accurate depth sensing in a small form factor.

Small enough to fit in the palm of your hand, the L515 is 61mm in diameter and 26mm in height. At approximately 100g, it is designed to be easily situated on any system or attached to a tablet or phone. It also runs at less than 3.5W, considerably lower than competing time-of-flight (TOF) solutions covering the same max object distance. As all the depth calculations run on the device, it also has low host compute requirements, resulting in true platform independence.

With a short exposure time of <100ns per depth point, even rapidly moving objects can be captured with minimal motion blur. Optimized for indoor lighting, the L515 processes over 23 million depth points per second via a custom-made ASIC. The product has been designed for use case flexibility with the inclusion of an RGB camera and an inertial measurement unit.

The product includes the following features:

- Depth Capture from 0.25 to 9m[1]
- 2MP RGB Camera
- Inertial Measurement Unit (IMU)
- Up to 30FPS Depth at 1024x768 (XGA)
- Up to 30FPS Color at 1920x1080 (FHD)
- Class 1 Laser Compliant

[1] For 95% reflective target

Getting started

System requirements

Minimum system requirements are:

- Ubuntu 16.xx/18.04 LTS
- Windows 10 (build 15063 or later)

Connecting the camera

The camera is connected to a host PC via a USB3.0 type C cable.

Figure 1 Connecting the camera to host PC

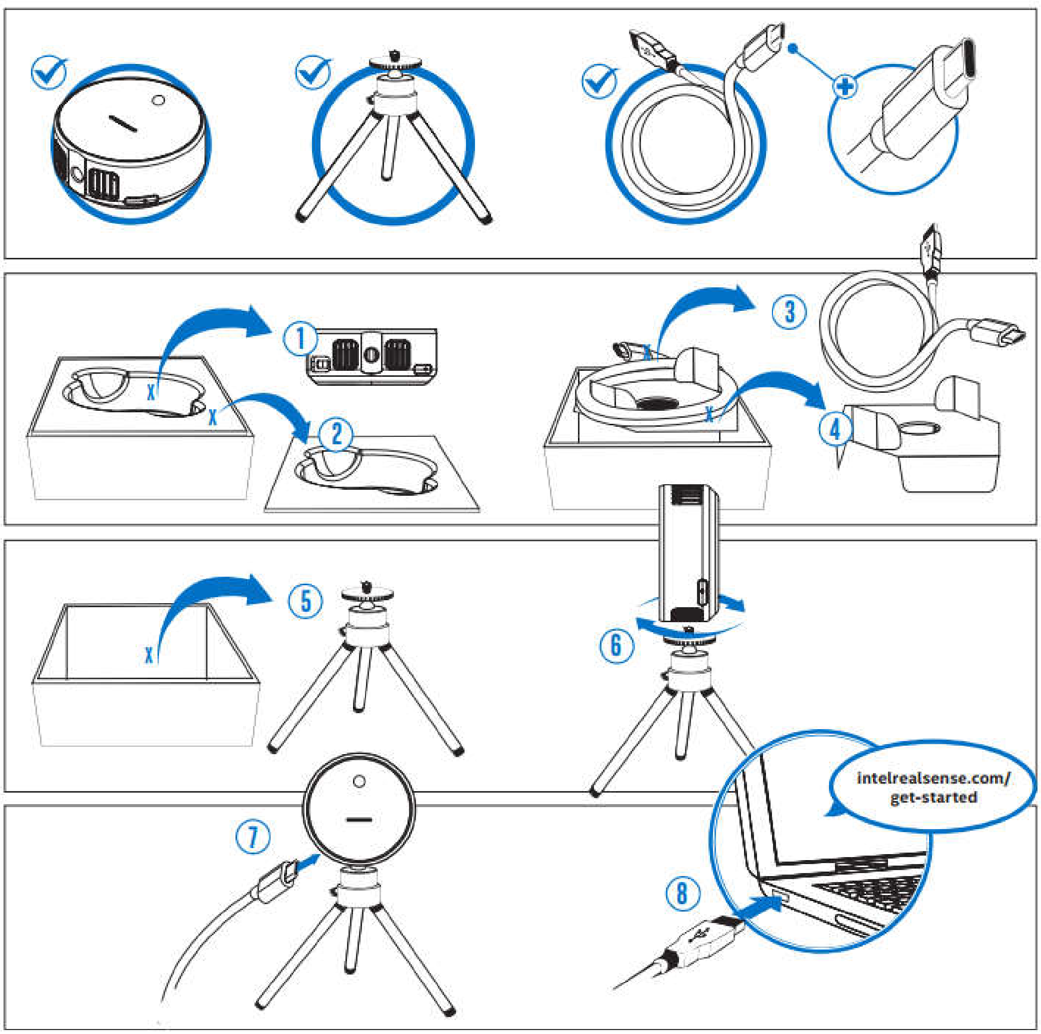

Viewing the image streams

In order to run the camera, you will need to download the latest Intel RealSense SDK from this location: https://www.intelrealsense.com/sdk-2/

We recommend that you only download Intel RealSense Viewer, as opposed to the entire SDK, unless you are a developer. This application can be used to view, record and playback depth streams, set camera configurations and other controls.

Available streams

L500 Depth Sensor: The 3 available streams are Depth, Infrared, and Confidence.

Depth stream displays the depth map, as computed by the device.

Infrared stream represents the intensity of the reflected laser light off the objects in the scene.

Confidence stream shows the relative confidence level for each computed depth value. Brighter pixels mean high confidence in the calculated depth, darker pixels mean low confidence in the calculated depth. These confidence values are used to filter out the noisy (low confidence) pixels from the depth map.

Figure 2 Depth sensor streams, from left: Infrared, Depth, Confidence

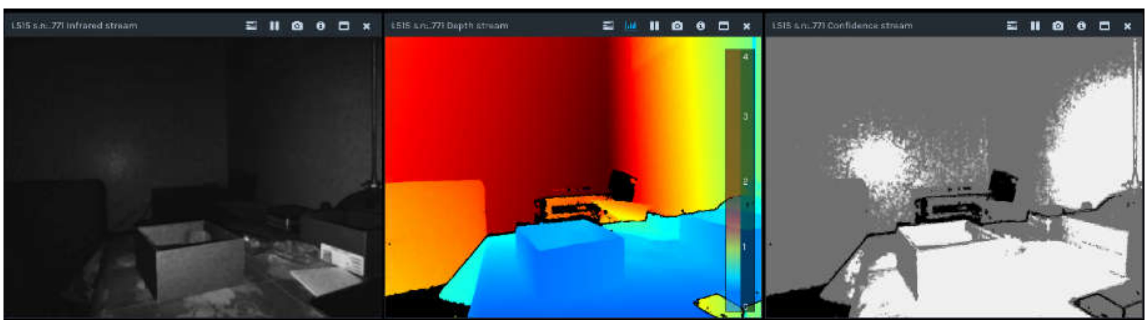

RGB camera: There are multiple resolutions and frame rates available to choose from.

IMU: Accelerometer and gyroscope data are displayed.

Camera controls

In this section we explain the various controls available to the user to control the depth camera configuration.

Before turning the camera on, let's make sure that the settings are selected properly for the user's environment and use case.

Presets

The main knobs to control the camera are the laser power and the receiver gain.

Laser power

The higher the laser power the longer the camera range will be. The longer range comes at the expense of the shorter range detection, meaning that if the laser power is too high, closer objects will appear too bright (saturated) and a proper depth measurement may not be possible.

Receiver gain

High receiver gain is useful in situations where the signal coming back from an object is too low (distant objects and/or dark objects). However, selecting a high receiver gain also means amplifying the noise. This could include the interfering light from other sources, such as the sunlight. Therefore, a careful selection of the receiver gain is needed, depending on the presence of ambient light.

Controlling laser power and receiver gain can help improve depth performance in some situations. However, there are tradeoffs that should be understood when adjusting these settings.

	Pro	Con
Increase laser power	Easier to measure poorly reflective objects Extends depth range	Near objects may oversaturate receiver. Highly reflective objects my oversaturate receiver.
Decrease laser power	Minimum distance from camera to object is reduced.	Maximum range is reduced.
Increase receiver gain	Easier to measure poorly reflective objects Extends depth range	Increase in noise from ambient light.
Decrease receiver gain	Minimum distance from camera to object is reduced. Lowers noise from ambient light.	Maximum range is reduced.

Now that we know the trade-offs involved in setting the laser power and the receiver gain, let's examine the various presets that are available to the user. These presets control 3 parameters: laser power, receiver gain, and min distance. It should be noted that these parameters can be controlled manually as well, so these presets only represent suggestions for the user depending on their use case.

Figure 3 Depth camera presets

<u>No Ambient Light</u>: This preset is useful when there is no ambient sunlight in the scene (fully indoors use case, with no light coming through windows). Although sunlight is the most common light source containing near IR at 860nm, other light sources that can reduce L515 performance include halogen and some LED. Minimum object distance is 50cm. The main difference between No Ambient and Max Range is the laser power which is lower on this preset to avoid false depth on objects that are on longer distances than the ambiguity range (10m-VGA, 6.5m-XGA).

<u>Low Ambient Light</u>: This preset is recommended for environments where there may be a low amount of ambient light present. Similar to Max Range preset, the laser power is set to maximum but the receiver gain is reduced to avoid saturation of the camera due to ambient light. The preset is also recommended for cases that the user wants to detect close objects (<50cm).

<u>Max range</u>: With this preset, the laser power is set to maximum as well as the receiver gain which optimizes the depth quality in indoor conditions. This is a good setting for use cases where maximum range is needed, and there are no objects closer than 50cm that need to be detected.

<u>Short range</u>: This preset lowers the laser power and gain so that close objects do not oversaturate the receiver. This allows operation at a close distance to objects. This setting may not be good if objects further away in the scene also need to perform well.

Figure 4 Difference between "No Ambient Light" and "Low Ambient Light" presets in the presence of ambient sunlight. Left: No Ambient Light, Right: Low Ambient Light

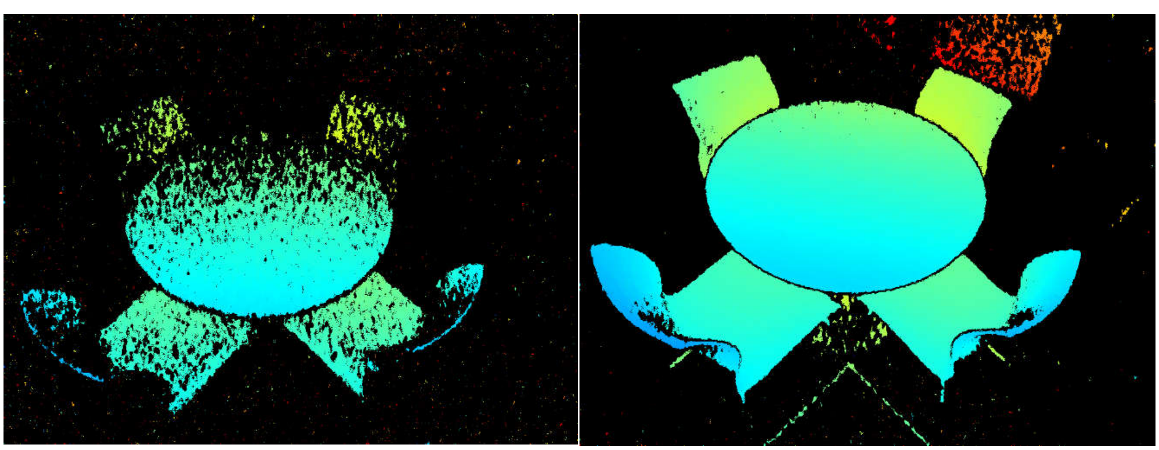

Depth sensor resolution, frame rate, and available streams

Resolution: Two resolutions are available to choose from: 640x480 (VGA) and 1024x768 (XGA). In order to achieve longer range detection, we recommend choosing the VGA resolution. If you would like to have a higher lateral resolution (better edge fidelity), the XGA resolution is recommended.

Frame rate: The L515 offers one fixed frame rate of 30 frames per second (fps).

Available streams: You can choose to stream any combination of Depth, Infrared, and Confidence streams.

Controls

Figure 5 Depth camera controls

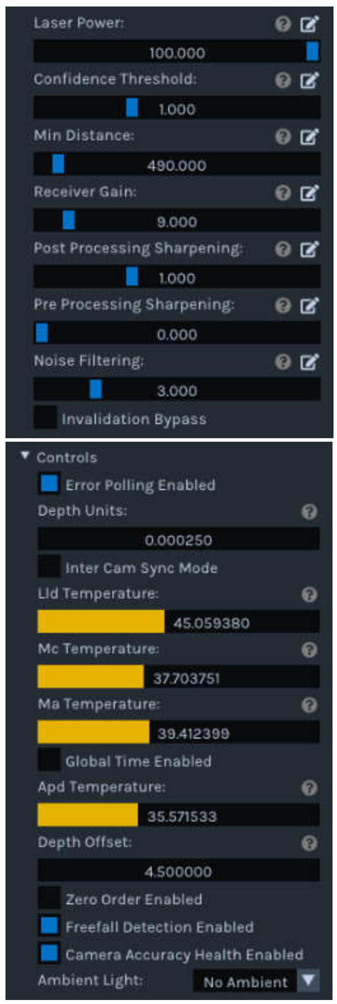

Freefall Detection: When enabled, the camera uses information from the onboard IMU to determine if the camera is falling. When a fall is detected the MEMS stops moving which is the better state for surviving high impact.

Laser Power: Adjustable between 0-100%.

<u>Confidence Threshold</u>: Adjustable between 0-3. The depth map is filtered by the confidence value per depth point. If this slider is set to 0, no filtering will be performed, thus showing all the calculated depth values, even when the noise is too high and the measurement is most likely wrong. Setting this value to max (3) only passes the very high confidence measurements through. The default setting is 1. High confidence threshold is recommended for applications that rely on "single-shot" performance and can't afford to handle false depth values via temporal averaging of multiple frames. On the other hand, applications that can allow temporal filtering and tracking can operate with low confidence as it increases the data fill rate.

Figure 6 An example scene to illustrate the effect of confidence threshold

Figure 7 Confidence is disabled. False depth values appear in low signal regions (top-right and bottom-left corners) but the edges of the objects are sharp (magnified left ear of Mickey Mouse shown on the right).

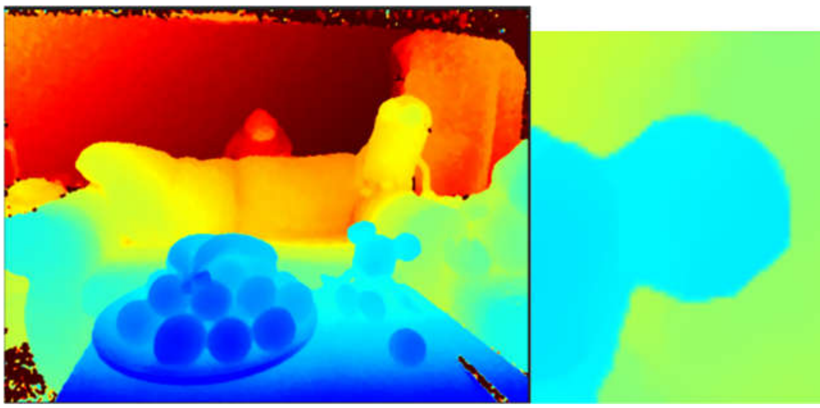

Figure 8 Confidence is set to 1 (default). False depth probability in low signal regions is less than 0.1% but the edges are not as sharp (they have been invalidated due to low confidence).

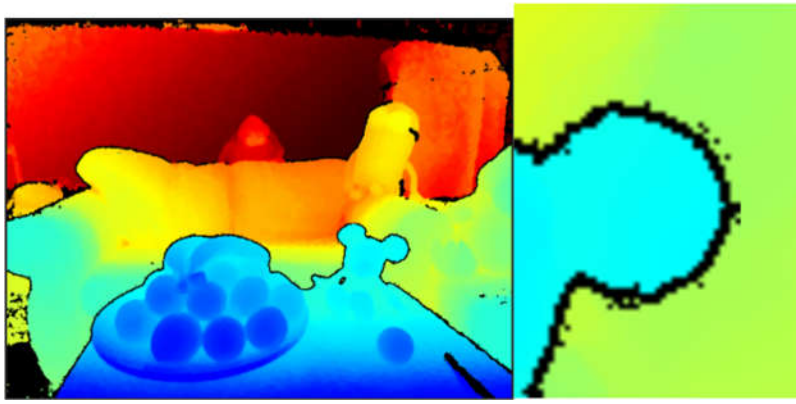

Figure 9 Confidence is set to 3 (max). No false depth but fill rate is also impacted, as there are fewer depth values that pass the high confidence threshold.

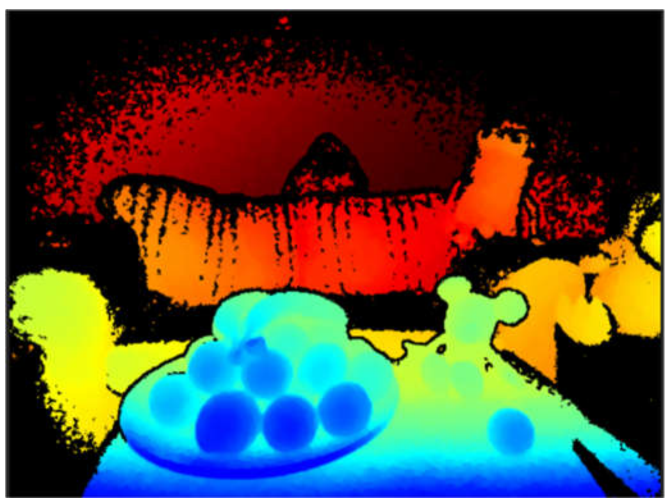

Min Distance: This sets the minimum object distance to the camera. Below this value, all calculated depth values will be set to 0 (unknown).

Receiver Gain: Receiver Gain is adjusted to accommodate different depth ranges and lighting conditions. Increasing gain has the tradeoff of also increasing noise.

Invalidation Bypass: This has the same effect as setting the Confidence Threshold to 0. Enabling invalidation bypass allows all the calculated depth values through.

Zero Order Enabled: This filter can be enabled when there are objects closer than 50cm to the camera, under certain scenarios. It removes the glow around the close objects when the scene behind the objects is far away. Off by default.

Figure 10 Effect of zero order filter on the depth map. Left: Disabled, Right: Enabled, the noise has been removed from around the fingers.

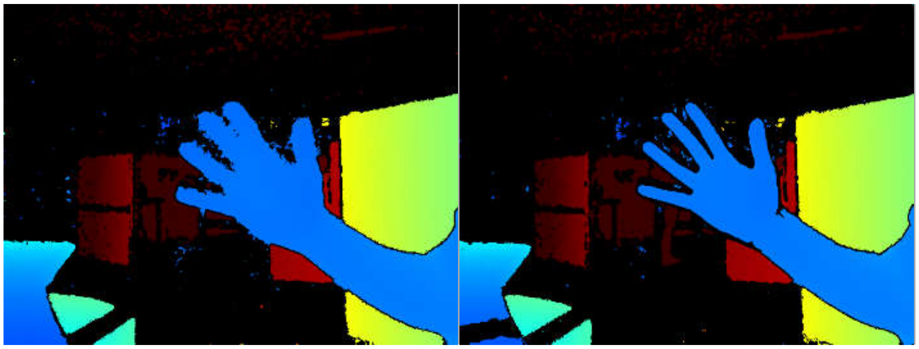

Post processing sharpening: The sharpening control impacts the smoothing of the edges in the depth map. High sharpening value gives sharp edges, while a low sharpening value causes smoothing of the edges. A byproduct of this smoothing is reduction in the noise

Figure 12 Color image of an object with sharp edges to illustrate the effect of sharpening filter

Figure 11 Sharpening set to max. The edges are sharp.

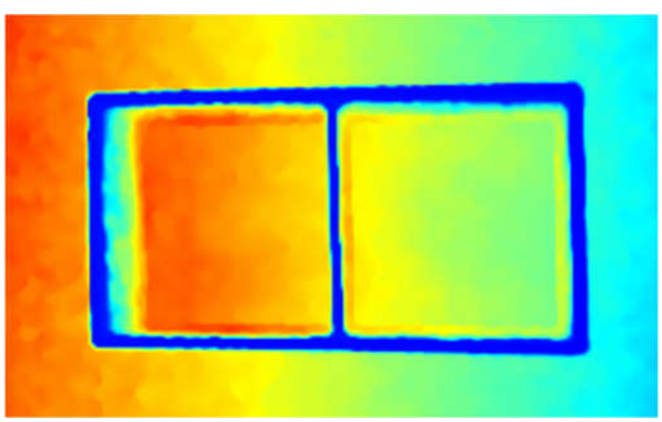

Figure 13 Sharpening set to min. The edges are smoothed, but the noise has been reduced (smoothed) as well.

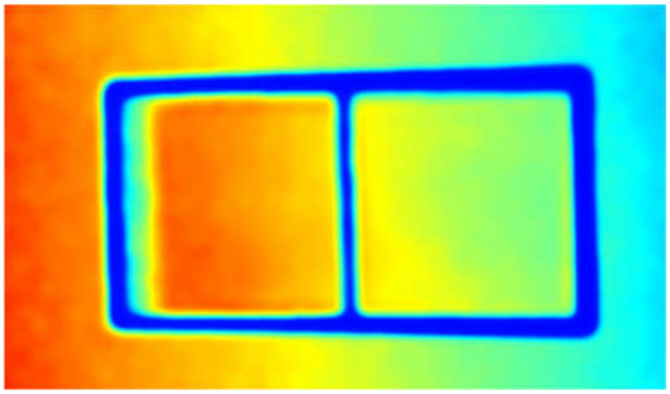

Freefall Detection Enabled: On by default, this feature enables an algorithm that detects when the camera is falling, utilizing the internal IMU. If a free-fall is detected, the scanning mirror will be automatically turned off in order to preserve the most sensitive part of the camera from breakage.

Depth Visualization

These controls set how the depth map is displayed on the screen. There are several Visual Presets and Color Schemes to choose from, depending on the use case and the color preference.

Figure 14 Depth visualization controls

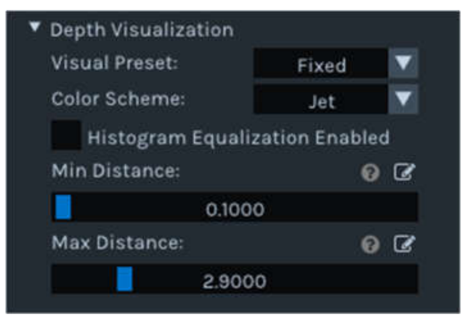

Histogram Equalization allows the Viewer to choose the min & max distance in the scene in order to optimize the visualization. The user can also elect to define the min & max values manually, which could be useful if they would like to focus on a certain range of distances in the scene.

3D viewer

The 3D viewer allows manipulation of the point cloud in 3D space. For instance, in the below examples the 2D and 3D views of a scene is shown. The user can also select to overlay the IR stream or the color stream on the 3D view, by selecting the Texture Source from the drop-down menu.

Figure 15 3D view of a scene. In order from top to bottom: Color, 2D depth, 3D view, 3D view with Color as texture source

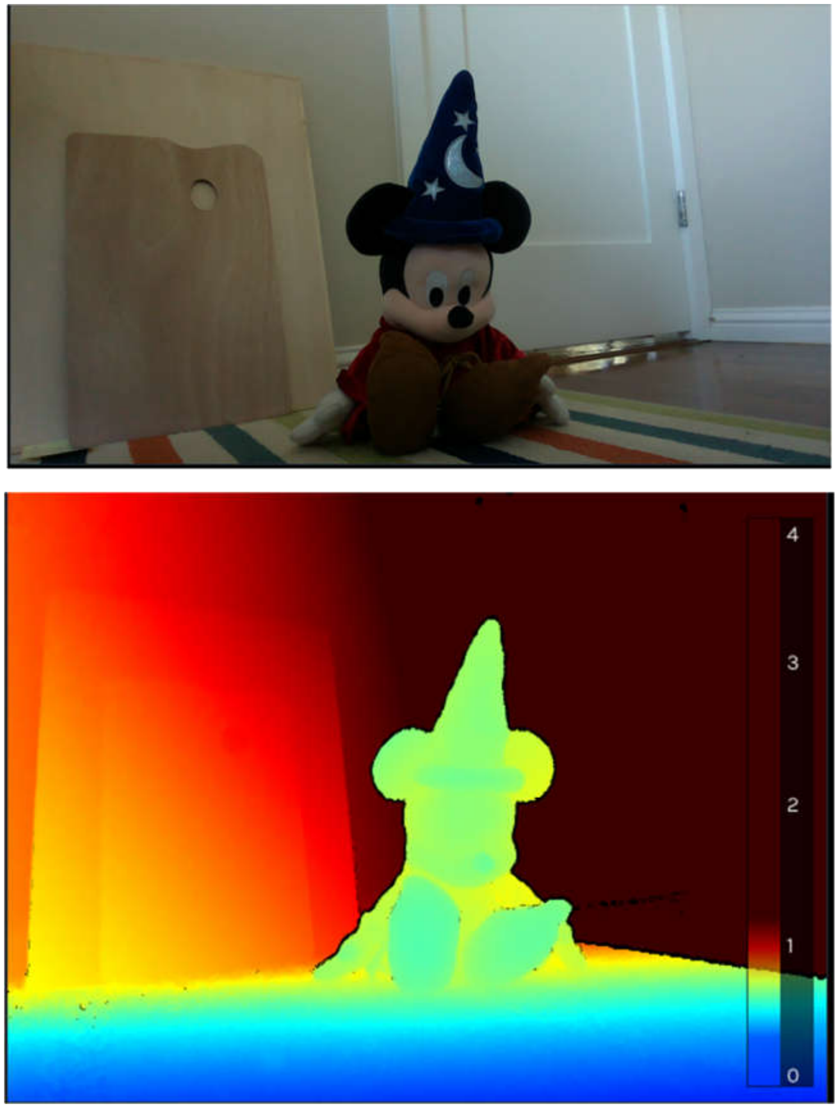

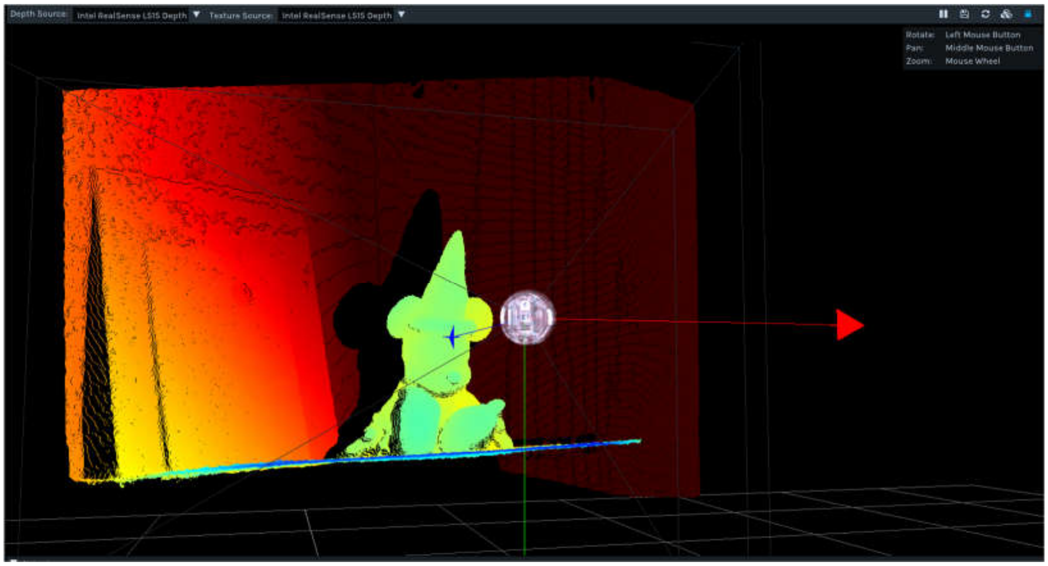

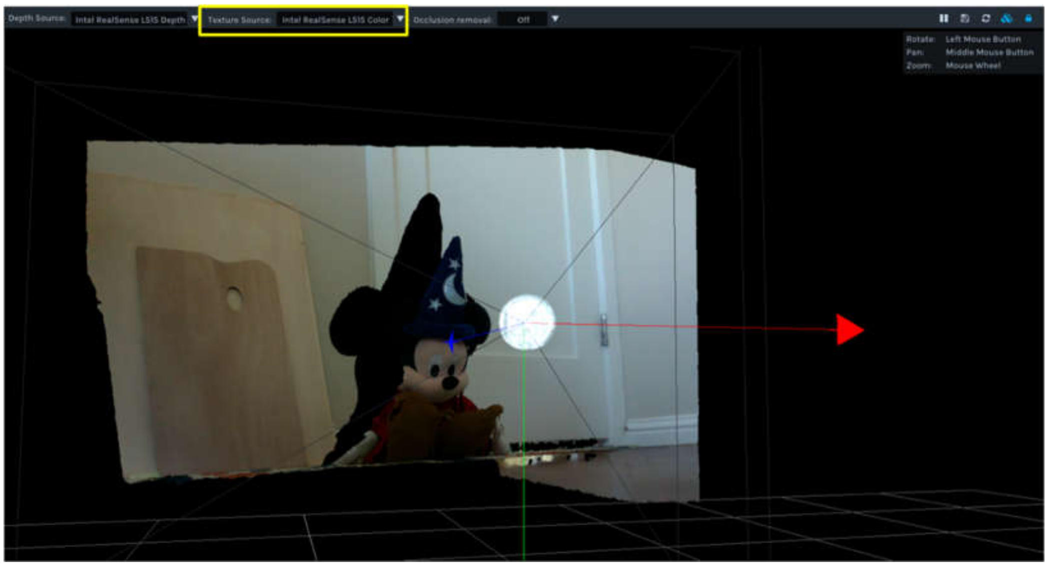

Post Processing

The post processing filters run on the host computer on the acquired frames.

Figure 16 Post processing filters

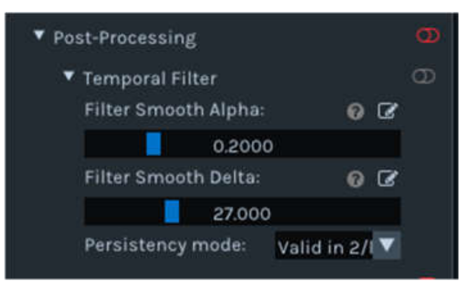

Temporal Filter: Averages a few frames in order to reduce the noise and to fill the small holes in the depth map. Its suggested use is for static scenes. Using the temporal filter on moving objects will cause blur.

For more details on the temporal filter and its parameter, please refer to our white paper at: https://dev.intelrealsense.com/docs/depth-post-processing (sub-section 4, under Simple Post Processing)

Useful tips

Tips for optimal performance

- To get the best out of the camera, we recommended that you start with the “Low Ambient” preset, choose the sensor resolution (VGA/XGA), and leave the rest of the sliders at their default positions.
- Avoid placing the camera directly on top of a highly reflective surface. The reflections off the surface back into the camera can corrupt the depth map. Use the tripod if you would like to set the camera on a tabletop.
- The depth measurement works best on rough surfaces with diffuse reflection, where most of the reflected laser light will be collected by the receiver. Highly reflective objects such as mirrors and other smooth surfaces have specular reflection. With specular reflection the laser light may not get reflected back into the receiver for detection, and no depth value will be registered. In order to successfully measure the depth on smooth surfaces the camera would have to be above the surface, looking at a large angle of incidence toward the plane.

Figure 17 Reflection of laser light off smooth and rough surfaces

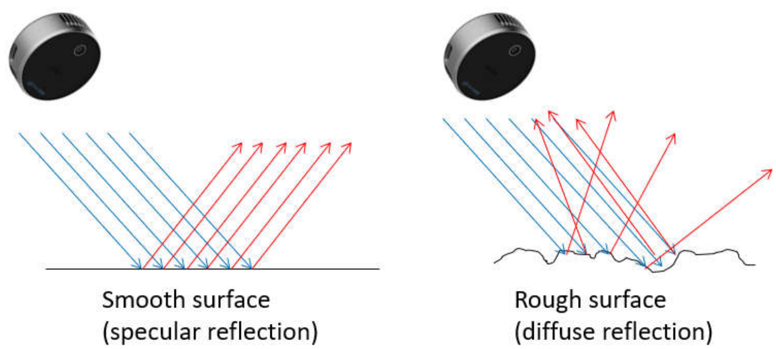

An example of a specular reflection can be seen below, where the camera is placed at a low angle with respect to the mobile phone, and thus no depth value can be registered for the glass surface.

Figure 18 When the light is incident upon a smooth surface at low angles, there will be no reflected light back into the receiver. The light bounces off the phone surface without ever coming back into the photodiode.

- The L515, as well as any depth camera based on TOF concept, relies on the reflection of light off the surfaces. For transparent objects, such as the wine glass in the example below, the laser light goes through without coming back to the camera and therefore no depth value will be measured (the depth of the back wall is measured instead).

Figure 19 Transparent objects don't reflect the light back into the camera receiver. The light goes through the empty half of the glass and gets reflected back off the next opaque surface.

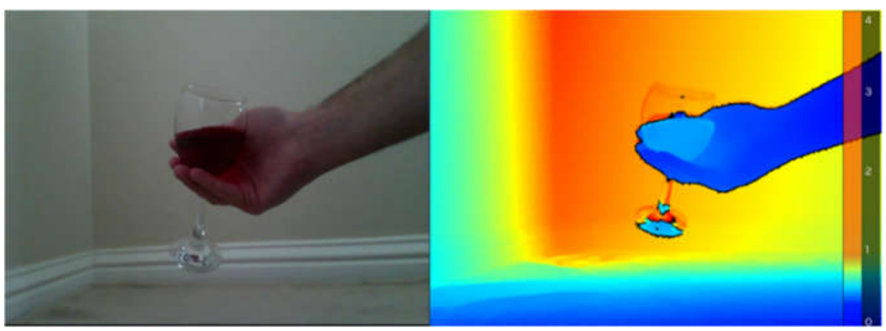

- Avoid placing the objects very close to the camera (<50cm) when there is nothing behind the object that can be detected in the depth map, otherwise you may see artifacts in the depth map.
- Do not place the camera inside an enclosure with no air flow, in order to avoid heat build-up. See datasheet for thermal guidelines.
- While the camera is safe to touch, even at its hottest point, we still recommend using the tripod.
- This is a LiDAR designed for indoor use. The performance will be degraded if the user points the camera towards the window during the day. In general, avoid brightly sunlit locations. Indoor residential/commercial lighting is perfectly fine, as is complete darkness. When there is too much interfering IR light in the environment, the intensity image will look washed out and the depth map will have many holes (unknown depth values).
- When reducing the min distance below 10cm, pay attention to the intensity image. If the image is getting saturated the camera will temporarily stop streaming (the duration is a few frames) to

protect the internal components. In such cases, reducing the laser power can mitigate the situation.

Figure 20 When there is too much ambient light in the scene, the IR image will look washed out (low dynamic range).

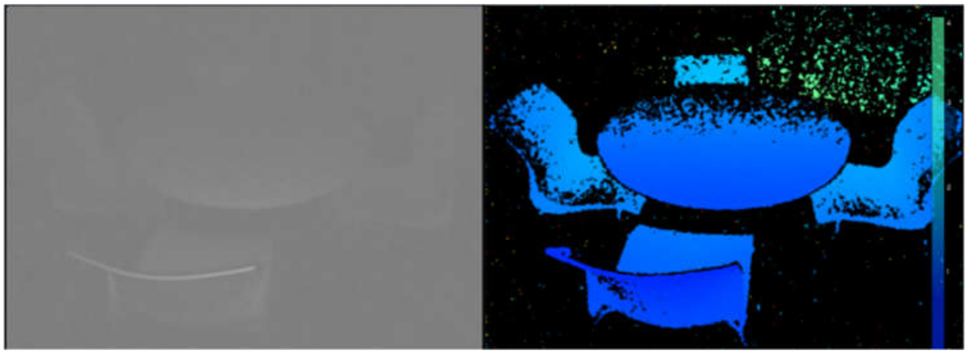

- The max range of the depth camera is dependent on the reflectivity of the target. Darker objects (dark hair, certain carpets, etc) may absorb most of the laser light, thereby reducing the effective range of the device. Increasing the laser power, when possible, will help in such cases.

Infrared Projector Interference and Hardware Synchronization

Interference can occur if the L515 is capturing infrared patterns projected from multiple infrared projectors simultaneously. Low interference assumes a user is in front of each L515 and comfortably spaced apart

Figure 21 Example configurations of multi-camera interference

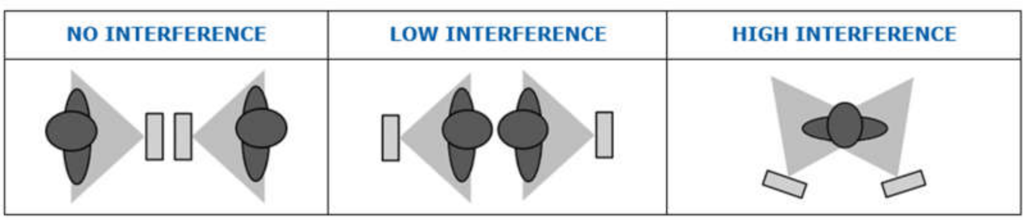

To tackle high interference, L515 supports hardware synchronization between devices. By alternating times when the camera projectors are used, interference can be mitigated or eliminated. When enabling the Hardware synchronization over LibRS (Inter_Cam_Sync_Mode) the L515 is operating in slave mode waiting for the hardware port to be high before enabling the laser and capturing an image.

Example: Body scanning booth with 12 Lidars, controlled by central synchronization unit and enable 360 body scanning in less than 2 seconds.

Instructions on how to enable and use the L515 hardware synchronization feature can be found in the L515 Hardware Synchronization Whitepaper(https://dev.intelrealsense.com/docs/whitepapers).

Achieving maximum range

The L515 is capable of measuring depth at up to 9 meters distance. To achieve usable depth at that distance it is necessary to configure the camera correctly and have an appropriate environment.

In the example image below, the center of depth frame, colored dark red, is at a distance of 9.3 meters from the camera.

Figure 22 Example scene with object at 9m

To achieve maximum range results from the L515 the following conditions are recommended.

- Set camera to Max Range preset and use VGA (640x480) resolution
- The scene should have no ambient light
- Target must have 95% reflectivity for IR light at 860nm and be in the center of the frame

Ambient sunlight will have a significant effect on range. (Although sunlight is the most common source of ambient IR light, any light source operating at 860nm such as halogen and some LED will also effect L515 performance.) An increase in the ambient light requires the camera to work harder in separating the ambient light noise from the light generated by the camera. The strength and quality of the light generated by the camera decreases with range. This is true of any light source and can be easily observed as light is strongest nearest a lamp and becomes weaker the further one moves away from the lamp. This is also true of the light projected by the L515 and therefore, the further the light travels, the weaker it becomes and the smaller the received signal will be.

The L515 is projecting the laser IR beam outwards in a spherical shape, however, displays depth as a plane. Let's consider for example an application that is scanning a wall at 9m distance from the camera. In this case, the distance from the camera at the center of the sphere is indeed 9m however the light that travels to the edge of the field of view horizontal to the center of projection, can be calculated as 9m divided by $\cos 35^0$ (horizontal field of view is 70^0), that's 10.98m. The corner of the FOV is even further away. Since the max range specified for the camera is 9m, one should expect reduction in plane projected FOV when approaching that range. L515 supports 10% ROI at 9m due to that specific reason.

Figure 23 Illustration showing max range of camera compared to full FOV range

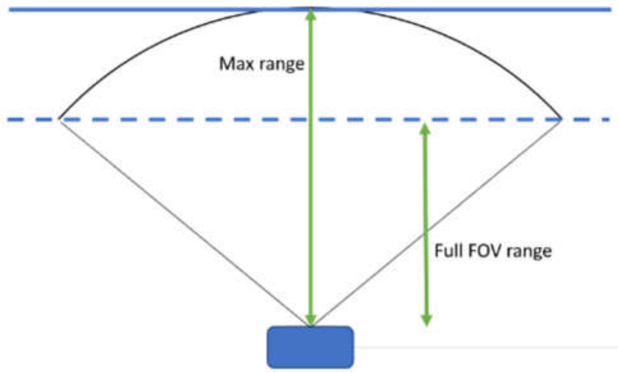

Another point to consider is SNR. L515 can work in very low SNR situations as it operates using coded IR light, however, ambient light or reduction in reflectivity will have an impact the SNR. The SNR to range calculation can be simplified as the square root of the SNR reduction. For example, our video below can demonstrate 9m range, indoor, at white wall. Assuming the wall has 90% reflectivity, if the camera will be pointed at a black wall that has 15% reflectivity, the returned light signal strength will be 6 times lower, that means the SNR is 6 times higher with the 90% reflective wall. To calculate the range at the 15% reflective wall we should take square root of 6 which is 2.44 (square root(6) = 2.44) and divide the 9m distance by 2.44 to get the expected range. In result, the camera will be able to detect the black wall at 3.6m (9/2.44 =3.6).

In the case of high frequency light sources (e.g. >1MHz), which may be found in some time of flight (ToF) cameras, the result is the same as lowering reflectivity in the scene. Therefore it is suggested to not operate ToF cameras near the L515 LiDAR camera.

Depth data quality

Lateral resolution

L515 provides depth map resolution of up to 1024x768 pixels, providing a very high level of edge fidelity.

Figure 25 Depth map of bars of varying gaps, from 5mm to 5cm, with L515, showcasing the high resolution of the depth camera.

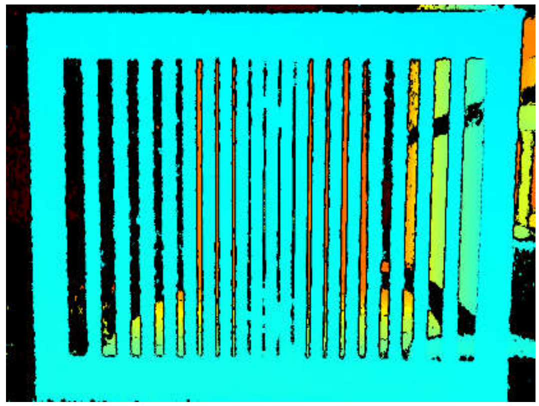

Figure 24 Depth map of a hand showing that the high resolution of the depth camera generates distinctly separated fingers

Compared to other depth sensing methods, such as stereo cameras, L515 has better edge resolution since the depth at each pixel is measured directly.

Figure 28 Depth of the same bar target with a stereo camera.

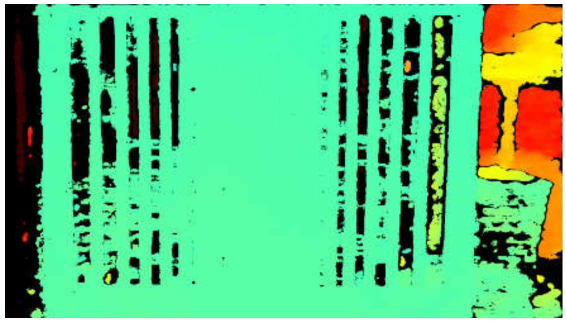

Figure 27 Depth map of a hand taken with a stereo camera.

A high-resolution depth map gives the user the chance to detect small objects at longer distances. Minimum detectable object size (MOS) directly depends on the number of depth pixels, as well as the fidelity of the depth camera in seeing the edges. L515 is capable of detecting objects as small as 9mm diameter x 11mm height from 1m away.

Figure 26 A target with cylinders of various heights and diameters, used to subjectively determine MOS. This example is taken with L515. The smallest observable cylinder is ~9mm diameter x 11mm height.

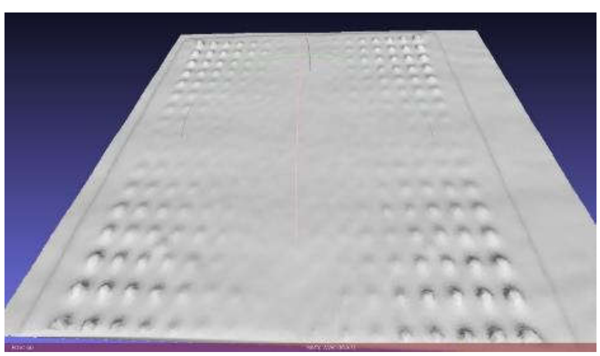

Robustness against motion blur

It takes <100ns to capture the reflected signal for each pixel. This instantaneous detection results in sharp edges in the depth map for moving objects.

Figure 29 Picture of a rotating blade (~100rpm). Left: one frame of a video taken with a cell phone which captures the blade blurry. Right: Depth from L515 (with overlaid intensity values) showing crisp edges for the blade.

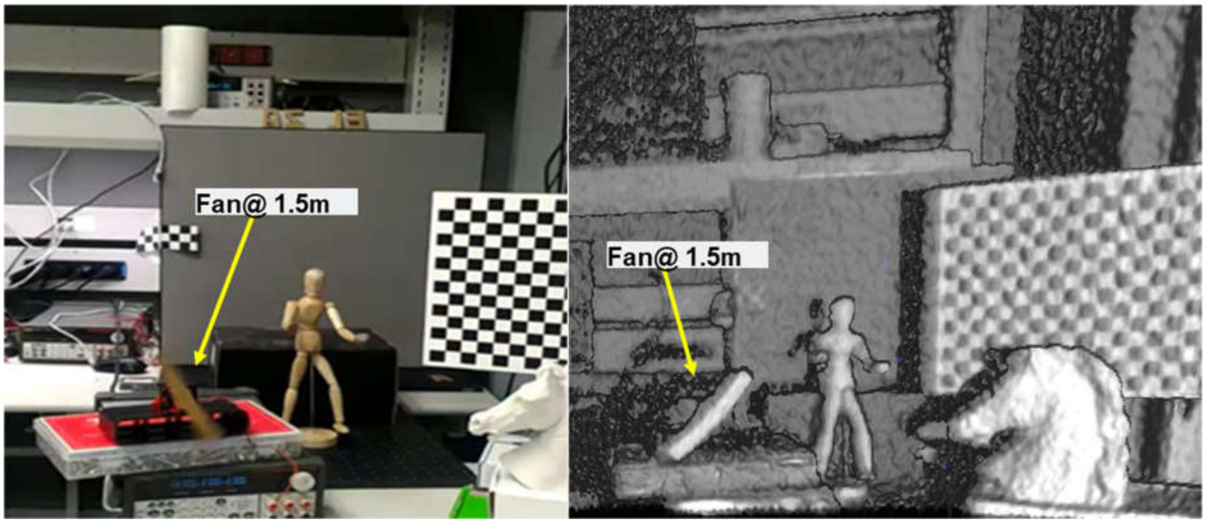

The MEMS mirror scans the laser beam starting from one corner of the FOV to the opposite corner. This whole scan takes slightly less than a frame time. Therefore, any movement in the scene may cause rolling shutter type of distortion.

Multipath artifacts

One common issue with TOF sensors is that the light from the illumination source can bounce off multiple surfaces before getting back into the receiver. This causes the total path of the light to be calculated wrong, and often manifests itself in the edge rounding of the intersecting planes. L515 has a benefit over conventional TOF systems (with flash illumination) in that only one beam of light experiences multipath trip at a given time, and as such the multipath effect is smaller in magnitude. In the example below, the L515 and a conventional TOF camera are pointed at the intersection of two walls. L515 shows straight walls with a small amount of curvature at the intersection of the walls, while the other camera shows bowed walls and a much larger curvature at the intersection.

Figure 30 Illustration of the multipath artifact. The incident light travels a distance of d1 before reaching the first surface. A portion of that light reflects onto the next surface and gets reflected off that surface back to the receiver. The measured distance by the depth camera becomes d1+d2+d3 as opposed to the correct value of 2*d1. This manifests as rounding of the corner in the depth data.

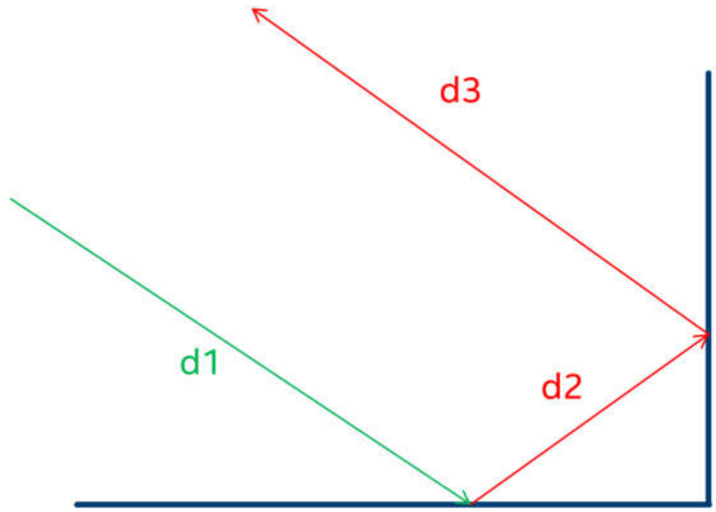

Figure 31 3D view of the intersection of the walls, Left: L515, Right: a conventional TOF camera

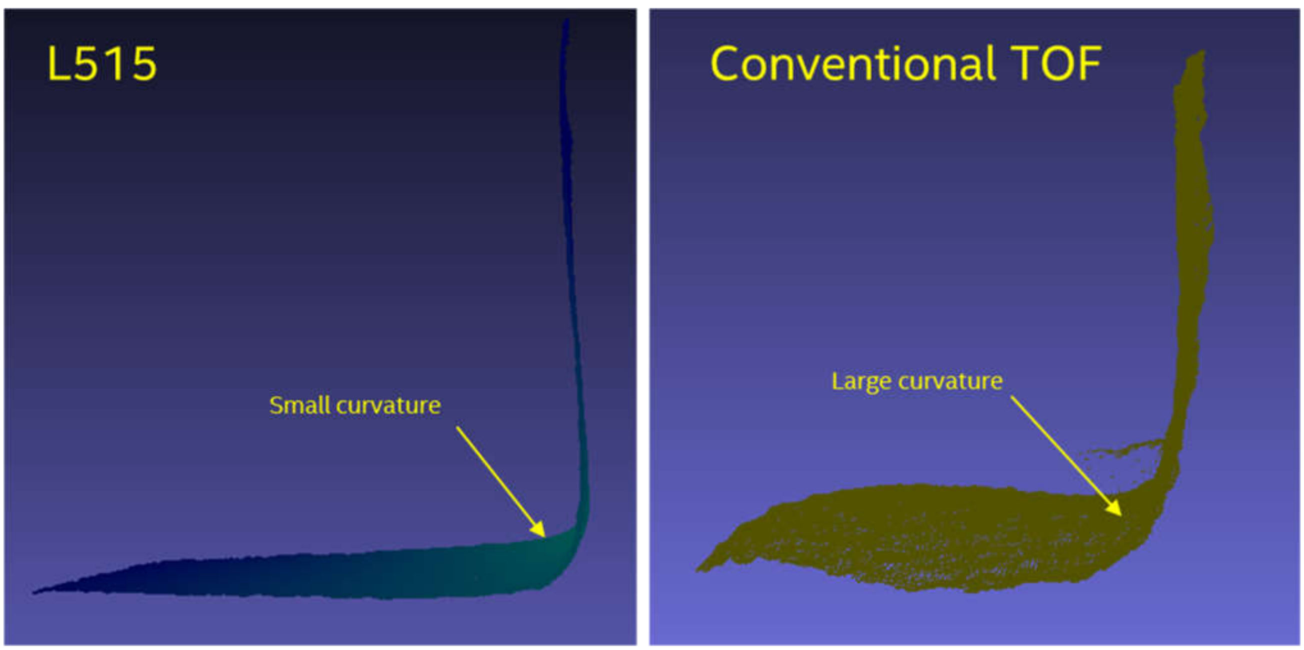

Troubleshooting

For a list of frequently asked questions please visit:
https://github.com/IntelRealSense/librealsense/wiki/Troubleshooting-Q&A

Safety and Handling Instructions

1) Do not power on the camera if any external damage is observed.
2) The camera should be operated in the specified ambient temperature range only.
3) No magnifying optical elements, such as eye loupes and magnifiers, are allowed.
4) Do not modify or service the hardware of this unit in any way. Modification or service of the hardware might cause the emissions to exceed the Class 1 level and result in hazardous radiation exposure.
5) Do not operate the camera in wet conditions.
6) Do not operate the camera in dusty environments.