Application of Digital Holography to Underwater Imaging

Dr. Nicholas Burns.

Research Assistant.

University of Aberdeen

Brief Recap

There are broadly two modes of recording holograms.

Off-axis recording can capture surface information from an object.

In-line recording only captures object silhouettes (or shadows).

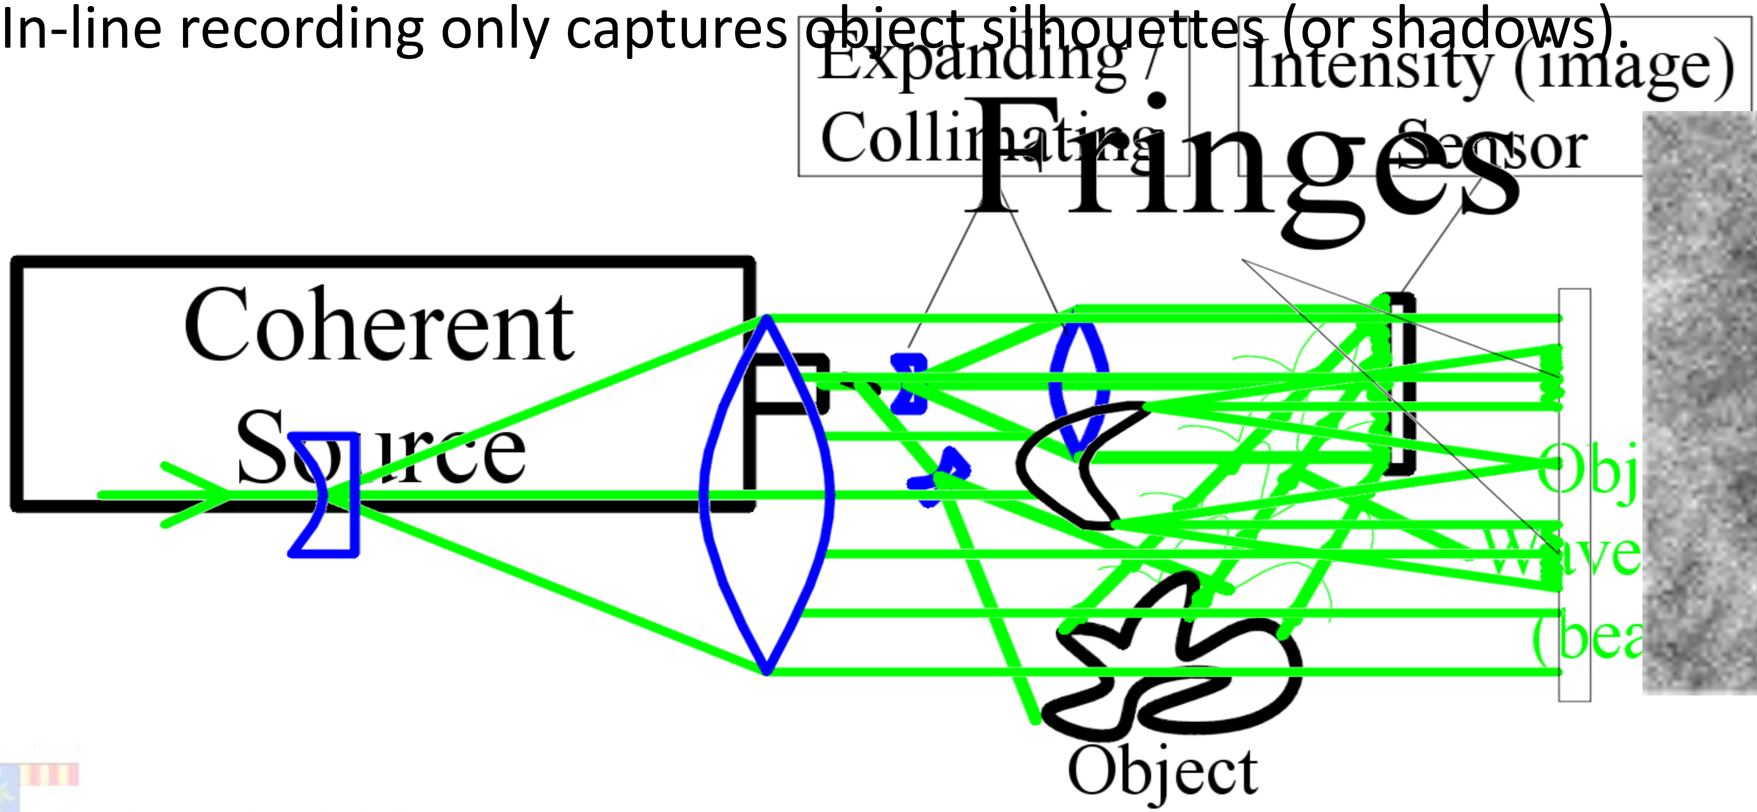

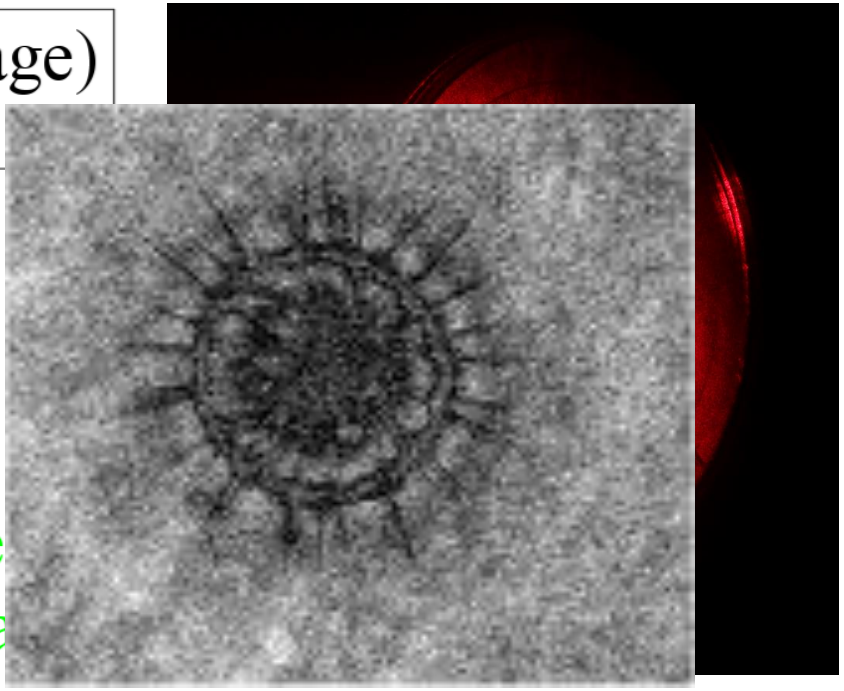

Brief Recap

Whichever mode of recording is used, the hologram itself is formed of interference fringes.

These fringes must be recorded onto a photo-sensitive medium.

In classical holography, we use a photo-sensitive emulsion similar to classic photography.

In digital holography, we use a digital image sensor, much as we do with digital photography.

For classical holograms, the hologram must be replayed using optical diffraction of a laser.

For digital holograms, the hologram is mathematically reconstructed by simulating optical diffraction.

Spatial Frequency

From the perspective of **digital** holography, the important difference between in-line and off-axis is in the attributes of the diffraction pattern.

The diffraction pattern consists of many **fringes** formed by the interference of light reaching the recording medium from different sources.

The most important attribute of all is **Spatial Frequency**, which is the inverse of the fringe spacing in the hologram.

The spacing of the fringes, d_f, can be computed for two intersecting plane waves by the following:

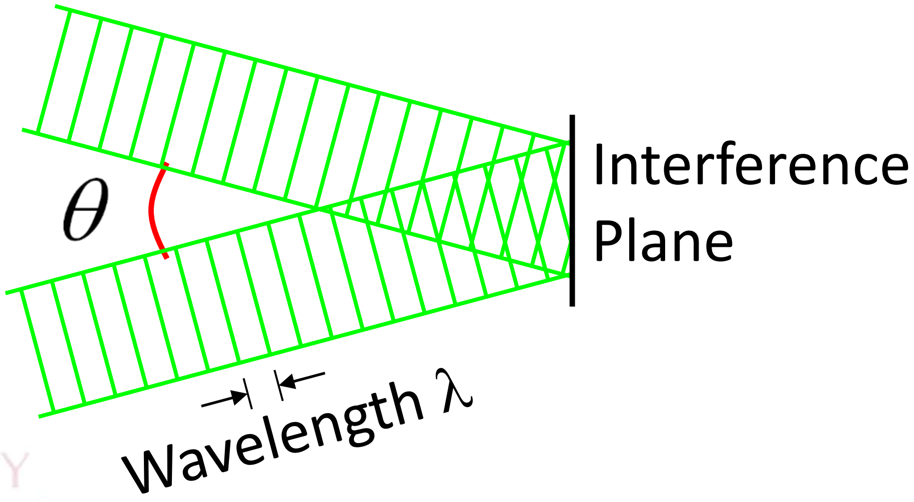

$$d_f = \frac{\lambda}{\sin\theta}$$

Spatial Frequency

As the angle of intersection approaches zero, the spatial frequency also approaches zero.

The general trend can be visualised by the following plot. The distance unit is arbitrary.

The interference pattern intensity at an intersection angle of 40° is such:

Interference of two beams at wavelenth 20

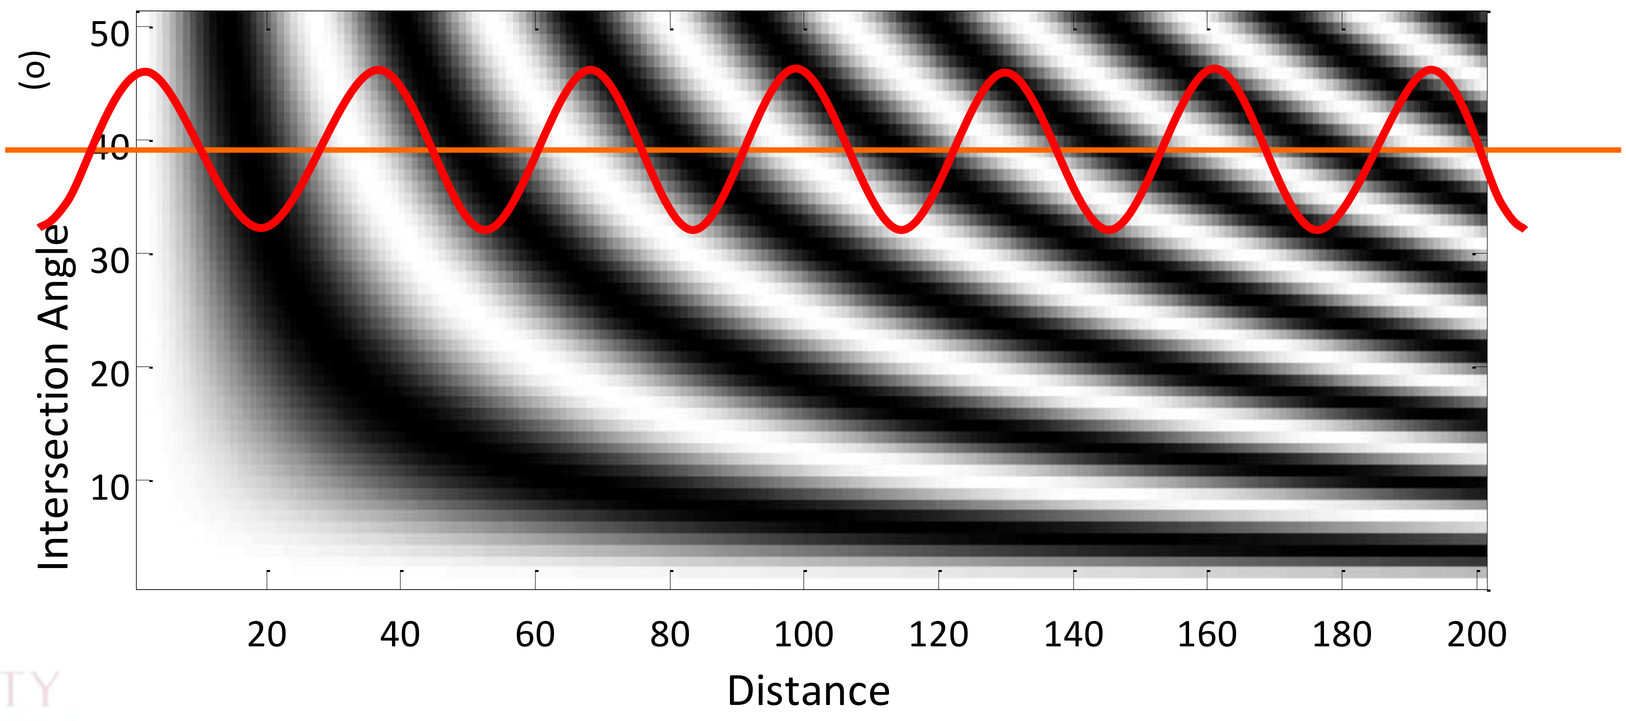

Spatial Frequency Requirements

The spatial frequency has minimal impact in classical holography where the fringes are recorded on photo-sensitive emulsion.

This is because the **grain size** of a holographic emulsion is in the range of some tens of nanometres. A typical grain size might be, e.g. 40 nm.

Referring to Nyquist Sampling Theorem, there must be two samples per full wave of a signal.

With a 40 nm grain size, the minimum fringe spacing that could be recorded is then 80 nm.

Compare this to the expected worst-case fringe spacing for two plane waves at 532 nm intersecting at 90°:

$$d_f = \frac{\lambda}{\sin\theta} = \frac{532\times10^{-9}}{1} = 532\times10^{-9}\ \text{nm}$$

No Problem here.

Spatial Frequency Requirements

In digital holography, our equivalent **grain size** is represented by the **pixel size** of the image sensor used.

Modern image sensors have pixel sizes down to approximately 2 μm square.

This limits spatial frequencies in *x* and *y* axes to 4 μm.

Reversing our spatial frequency equation and feeding in this value, we can arrive at a maximum angle of intersection that can be successfully recorded by the image sensor:

$$\theta = \sin^{-1}\left(\frac{\lambda}{d_f}\right) = \sin^{-1}\left(\frac{532\times10^{-9}}{4\times10^{-6}}\right) = 7.64^\circ$$

Spatial Frequency Requirements

This maximum intersection angle governs the maximum possible field of view in the hologram.

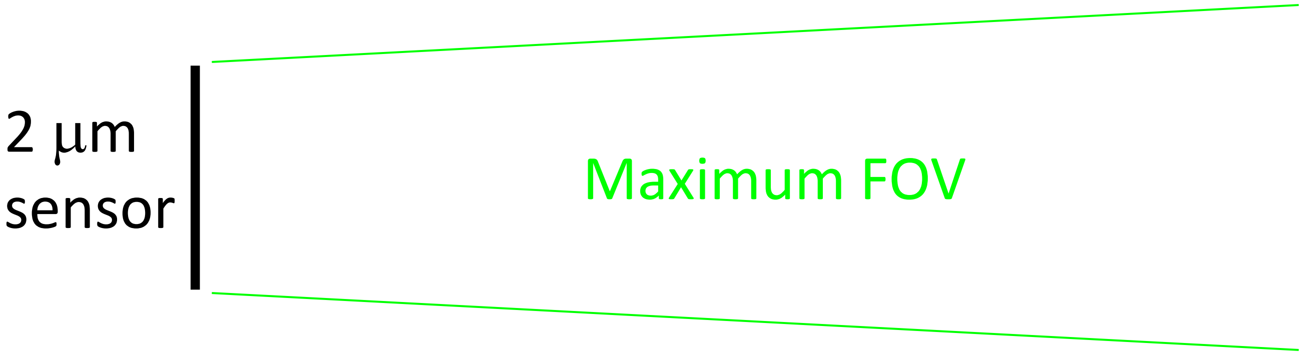

This limitation makes digital holography poorly suited for imaging large objects.

In order to successfully image large objects we could…

Place the object several metres from the sensor (requiring strong illumination).

Use a negative lens to extend our FOV, introducing scaling.

In-line digital holography

The spatial frequency limitation does not impact on the recording of in-line holograms.

This is because an in-line hologram is intrinsically FOV-limited to within the recording beam.

Returning to the diagram of earlier:

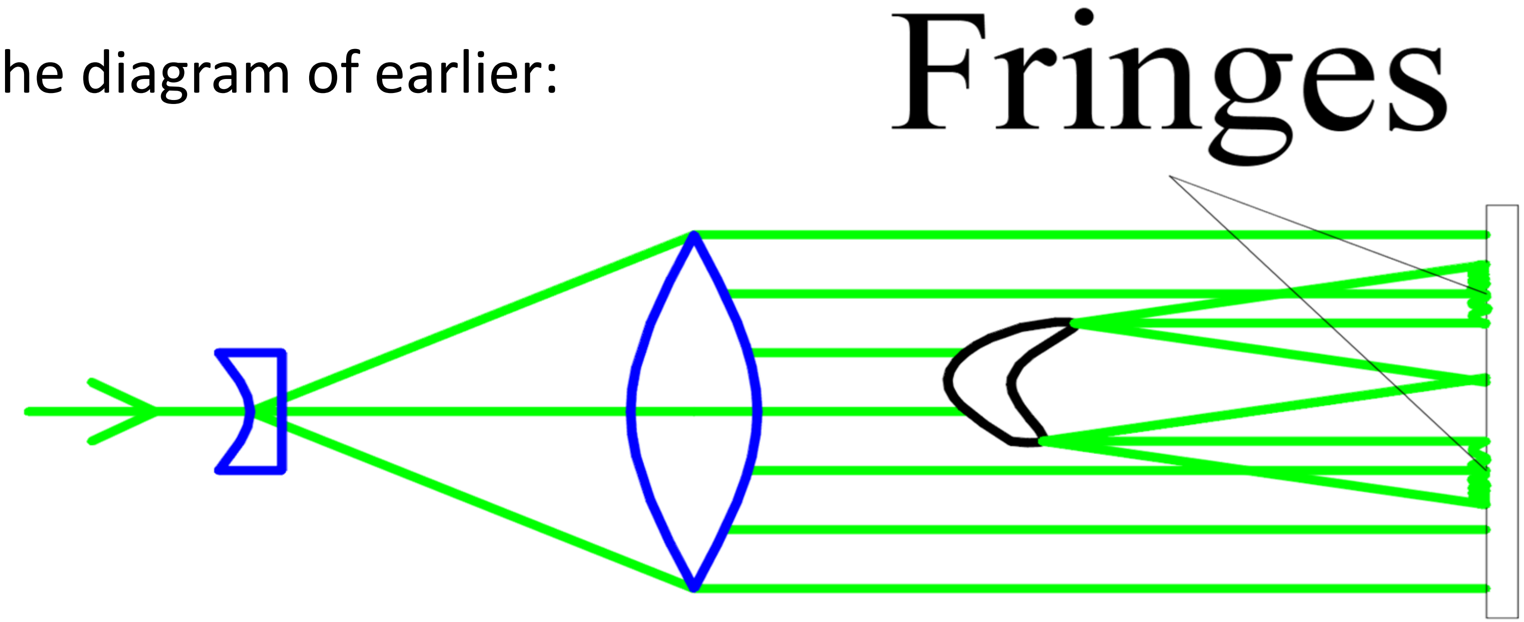

With this configuration we sacrifice the ability to see surface information, and instead only obtain transmitted light.

In-line digital holography

The inline configuration is what is implemented in the digital holographic camera, eHoloCam.

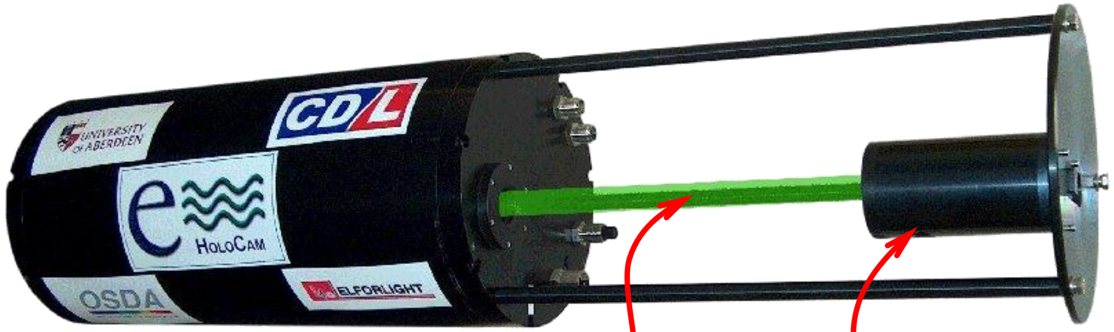

The camera uses a collimated beam, transmitted through the sample volume and imaged by the sensor.

Collimated beam

Sensor Housing

eHoloCam has a sensor with 3.5 μm pixel pitch.

The sensor can record at full-resolution with a framerate of 5 Hz.

The framerate can be increased to 25 Hz if the effective resolution is reduced by a factor of 3.

This system allows recording of holographic videos. Almost impossible in classical holography.

Recording in a Natural Environment

To successfully record holograms, the subject must not move during the recording process.

Any movement causes the interference fringes to shift, causing blurring of the hologram.

In a natural environment, it is next to impossible to keep the subjects from moving, therefore the illumination must be controlled to obtain short pulses.

This is achieved with eHoloCam by using a pulsed solid-state laser (FD Nd-YAG).

The pulse duration is in the order of 5 ns. The subject must not move by more than one pixel pitch in this time.

Taking Distance and dividing by time:

$$\frac{3.5\times10^{-6}}{5\times10^{-9}} = 700\mathrm{ms}^{-1}$$

Recording Set-up

The following video illustrates the recording set-up used in eHoloCam:

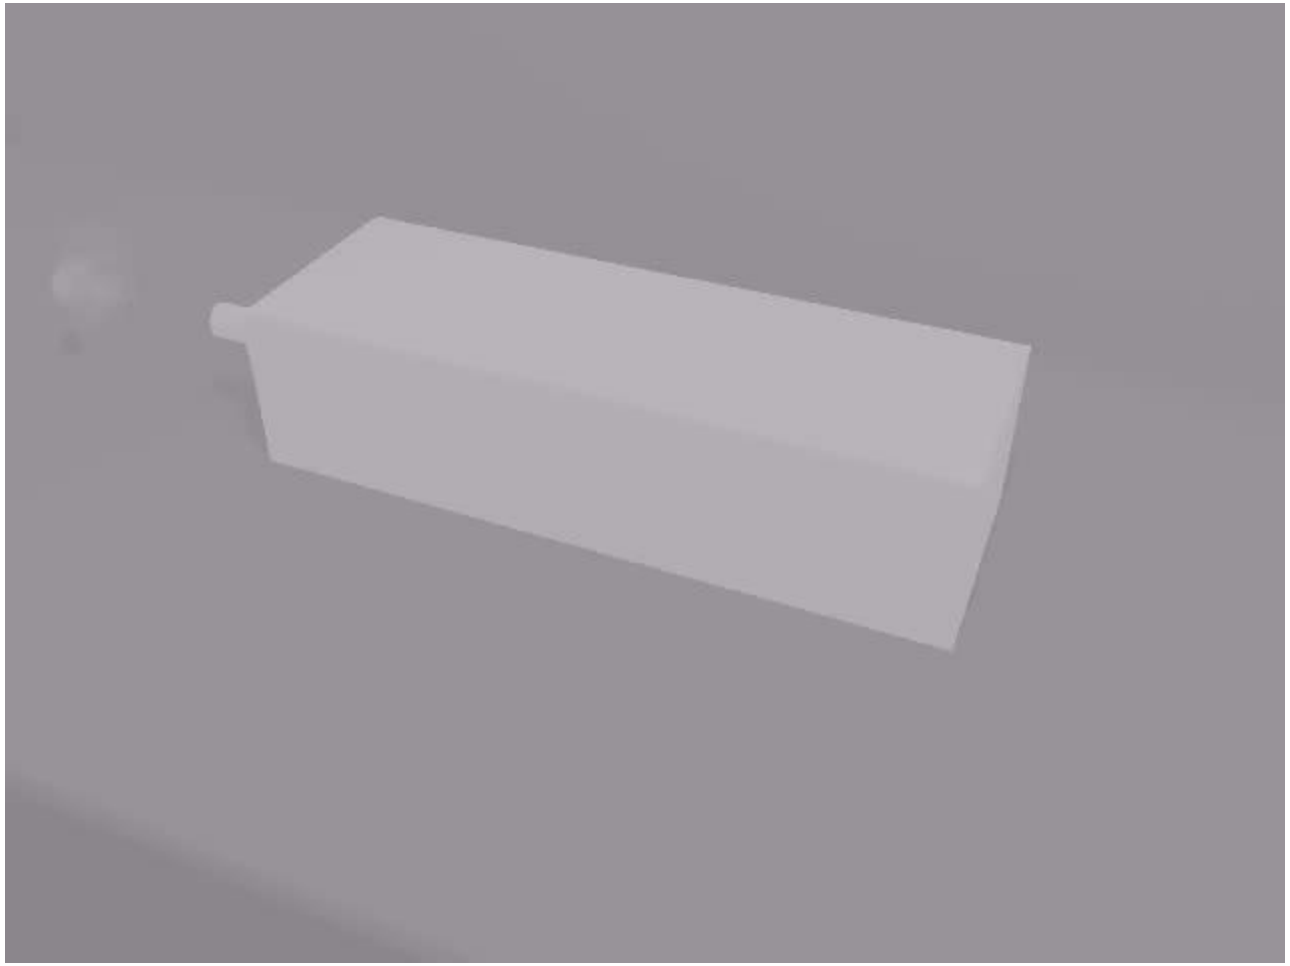

UNIVERSITY OF ABERDEEN

STRONG MAR

Hologram Storage

Storage of digital holograms is the same as for any digital information.

In the case of holograms from eHoloCam, they are stored as raw-format videos.

The format used is 256-level greyscale (8 bits per pixel), and each data block directly represents up to 2208 x 3000 pixels as per the sensor resolution.

This represents some 6.5 MB per holovideo frame, resulting in about 32 MB/s data throughput.

Modern computer systems are quite capable of maintaining this data throughput via USB or IEEE1394 (Firewire) and storing to hard-disk.

Hologram Storage

It is important **not** to apply typical video compression methods to the holovideos produced.

Most video compression techniques remove information to reduce space requirements, and will corrupt the hologram fringes.

JPEG compression, for instance, removes elements that are not noticeable to the human eye.

These elements are however important in the reconstruction of the hologram.

Lossless compression may be used to reduce hologram size, however compression ratios are likely to be poor.

Holograms contain a great deal of speckle noise, resulting in high information content.

eHoloCam Block Diagram

The design of eHoloCam illustrates a typical approach to designing a system to record digital holograms.

The laser has been configured to run from a 24 V Supply, and accept RS232 control and a TTL Trigger. The power supply can be controlled by a further input.

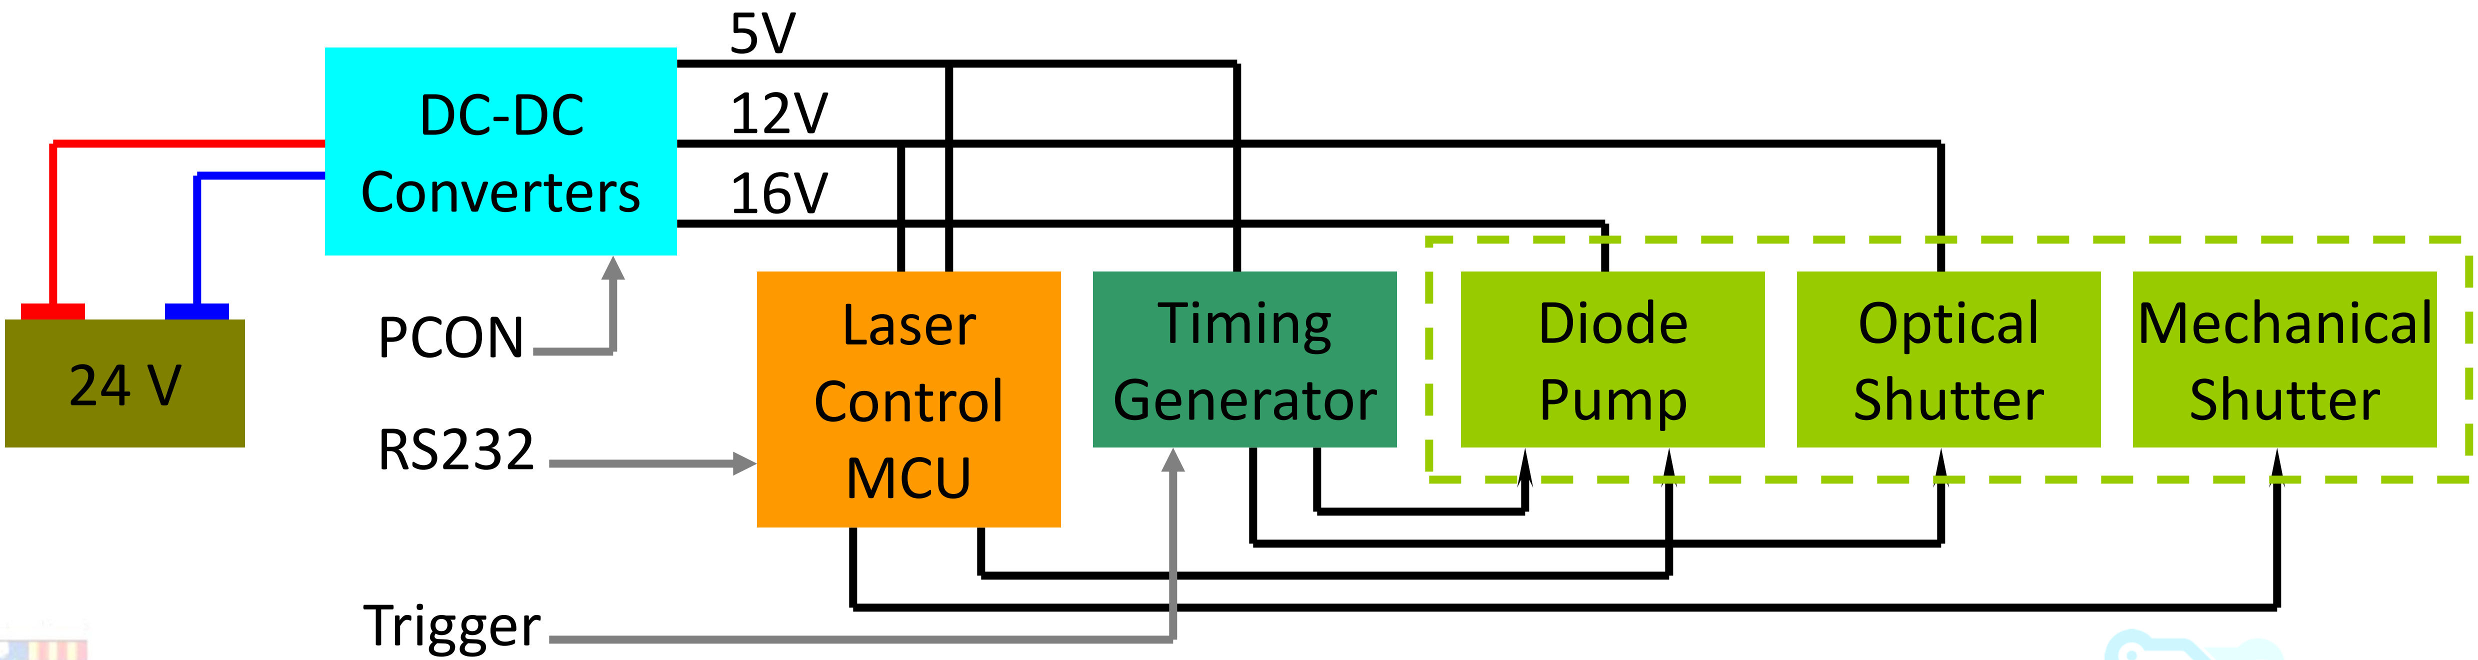

eHoloCam Block Diagram

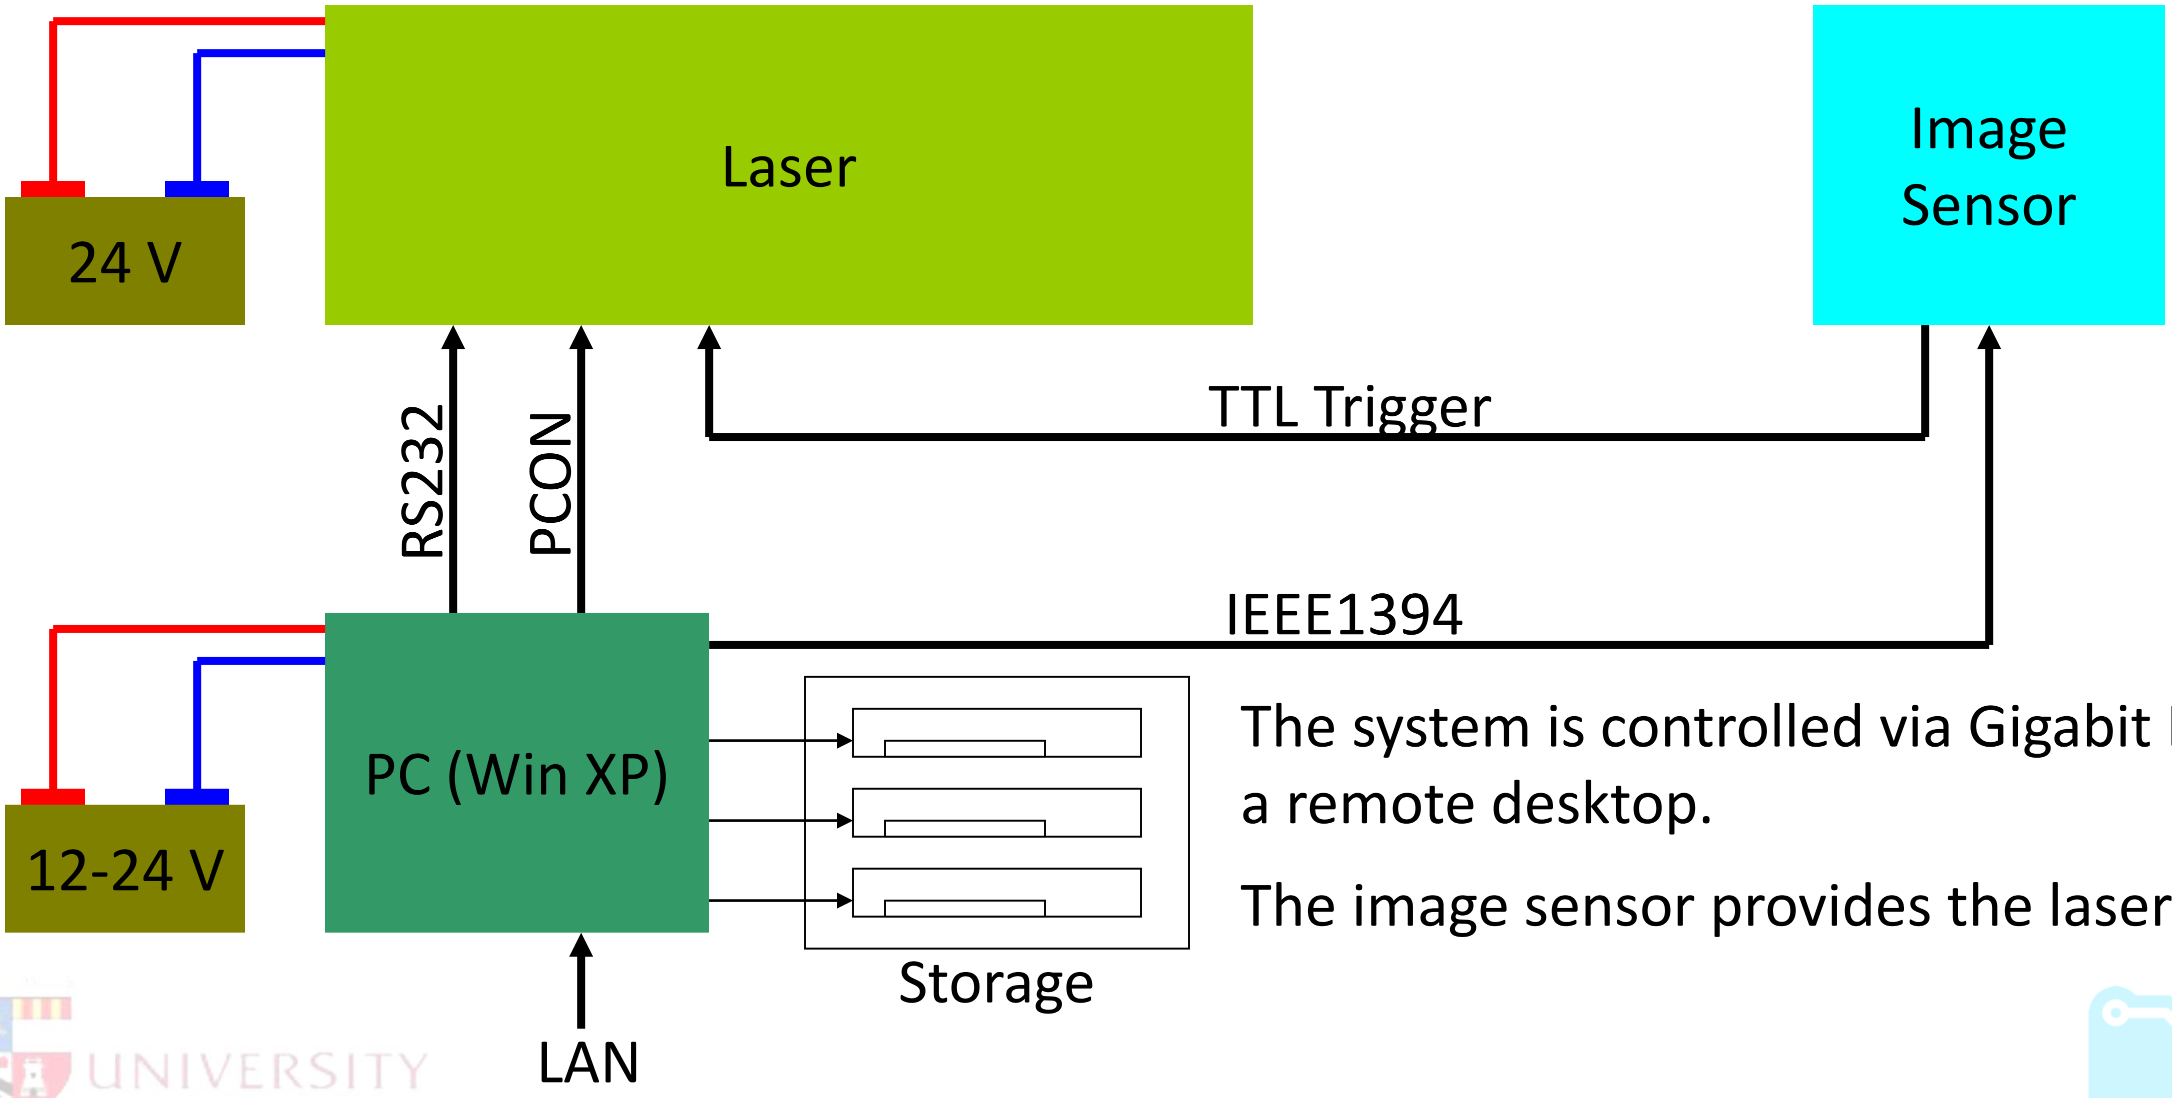

The system is controlled via Gigabit LAN with a remote desktop.

The image sensor provides the laser trigger.

Analysis of Holograms

To anaylse holograms, it is necessary to apply a reconstruction algorithm.

The algorithms used are based on the Fresnel-Kirchhoff diffraction integral.

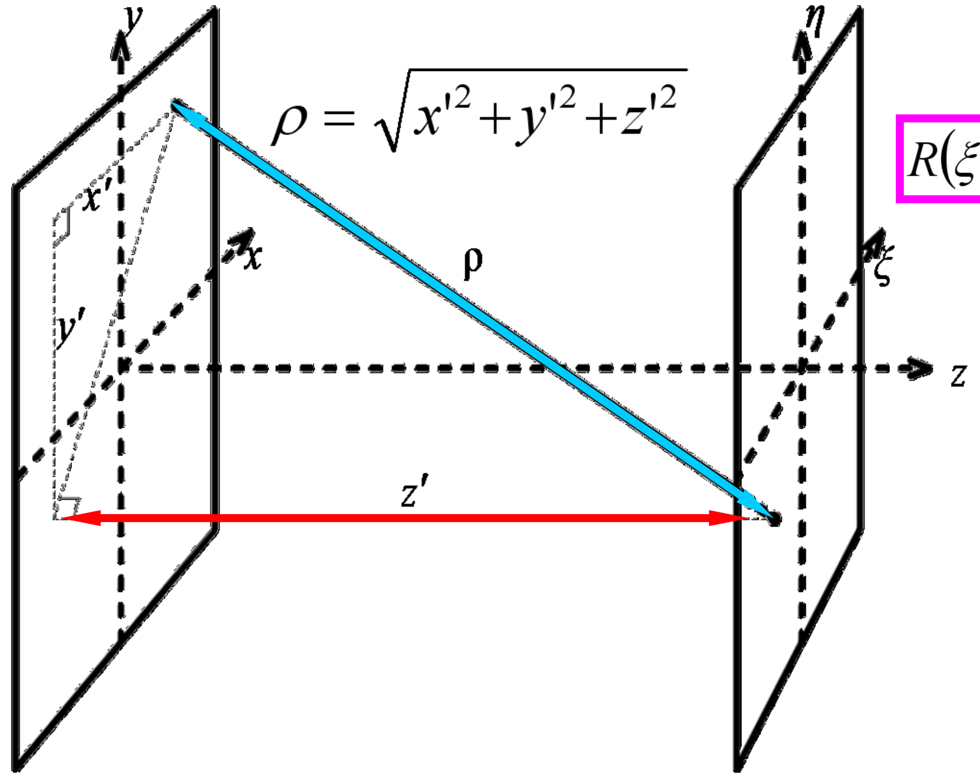

Hologram Plane (x, y) Reconstruction Plane (ξ, η)

$$R(\xi,\eta) = \frac{i}{\lambda}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} H(x,y) \frac{\exp\left(-i\frac{2\pi}{\lambda}\rho\right)}{\rho} Q\,dx\,dy$$

Given a particular reconstruction distance, *z*…

The phase of the wave from a hologram pixel can be found at a pixel on the reconstruction plane based on the distance ρ.

This can be represented as the complex number *a+jb* = exp… as above.

The intensity reduces with increasing distance, ρ.

And the result is summed for all hologram pixels.

Analysis of Holograms

Applying the Fresnel-Kirchhoff diffraction integral directly is possible, but very slow.

It requires each (ξ, η) pixel to be individually computed according to:

$$R(\xi,\eta) = \frac{i}{\lambda}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} H(x,y)\frac{\exp\left(-i\frac{2\pi}{\lambda}\rho\right)}{\rho}Q\,dx\,dy$$

The above does however share similarities with the Fourier transform.

It is possible, therefore, to rearrange and approximate and arrive at a number of variations of the algorithm that utilise Fourier transforms to obtain reconstructed images.

The variation we apply in our work is known as the “Angular Spectrum” approach.

Angular Spectrum

$$R_{(\xi,\eta)} = \Im^{-1}\left(\Im\left(E_{(x,y)}H_{(x,y)}\right)e^{(ik_d d)}\right)$$

Where

(ξ,η)	Reconstruction Plane Coordinates (at distance d)
(x,y)	Hologram Plane Coordinates
k_d	Propagation Matrix
$\Im$	Denotes Fourier Transform

The matrix k_d is constant for a particular hologram size and recording wavelength. This is calculated once for a particular hologram, and reused for each reconstruction.

$$k_d = \sqrt{k^2 - k_x^2 - k_y^2} \quad k = \frac{2\pi}{\lambda} \quad k_x = \frac{2\pi}{\Delta X}\left(\frac{x}{X} - \frac{1}{2}\right) \quad k_y = \frac{2\pi}{\Delta Y}\left(\frac{y}{Y} - \frac{1}{2}\right)$$

An interesting feature of digital holography is that the phase is reconstructed as well as the intensity. This is impossible in classical holography.

Reconstruction Approach

The nature of the reconstruction algorithms is such that holograms are reconstructed slice-by-slice.

Particles can be seen to start defocused, approach focus, focus, then leave focus as we scan.

This is synonymous with focusing an image formed by a lens by moving the image plane forward and backward.

The appearance can be seen in the video:

UNIVERSITY OF ABERDEEN

STRONG MAR

Reconstruction Approach

The nature of the reconstruction algorithms is such that holograms are reconstructed slice-by-slice.

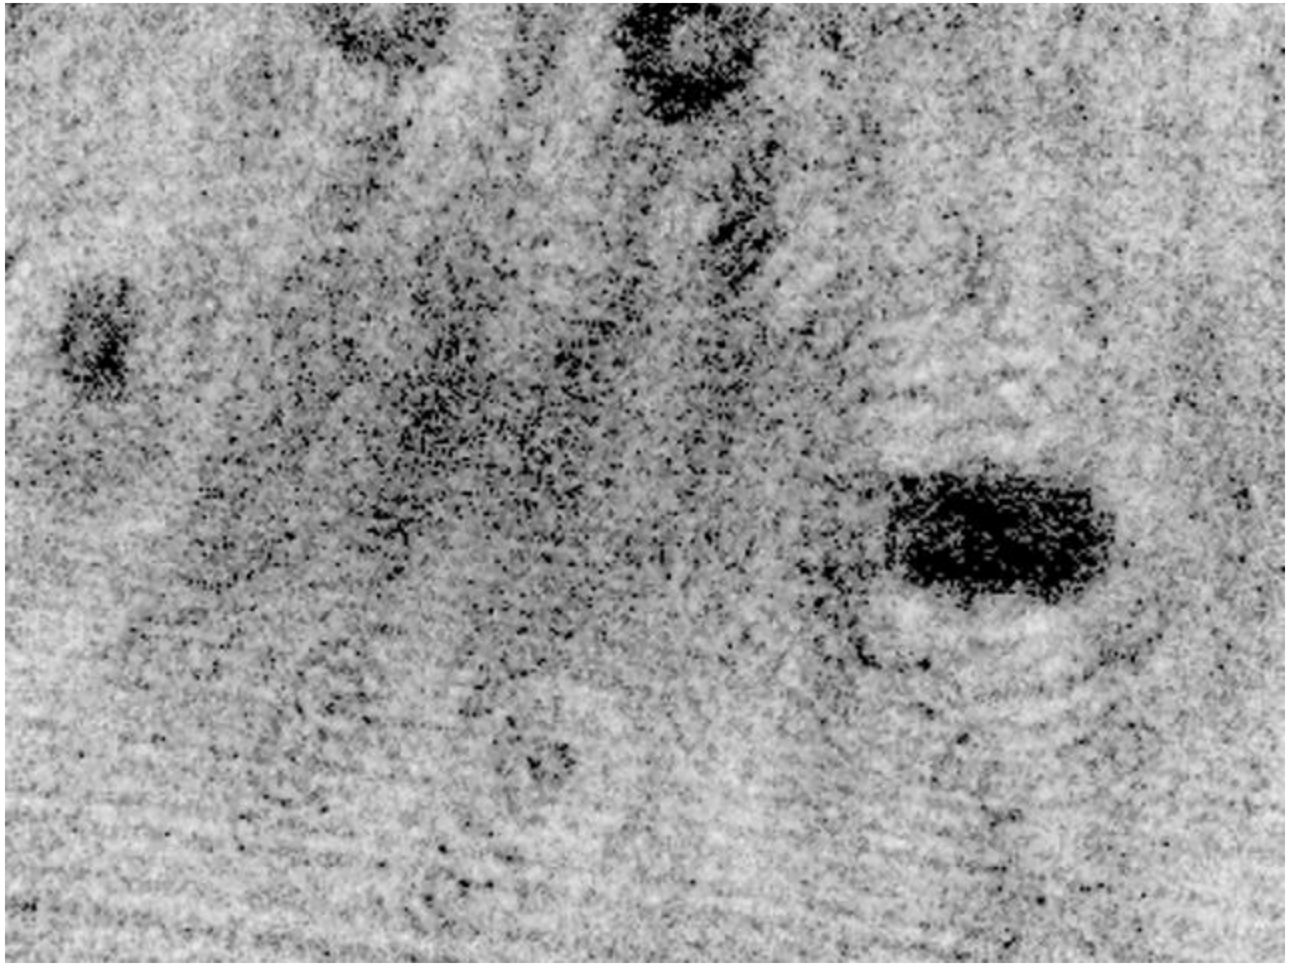

Automated Focusing

One of the problems with analysing digital holograms is that particles must be visually located and brought into focus.

A single holovideo may contain over 1000 holograms, and each hologram spans 400 mm.

Manually identifying particles and focusing them in such a large dataset is extremely time consuming!

The solution is to find mechanisms by which a computer can perform the laborious task and output a set of focused images that can be analysed easily.

Automated Focusing

Several algorithms have been developed for auto-focusing of standard imaging systems.

An obvious example is the auto-focus built into most digital cameras.

When applied to holography, many of these algorithms fail to operate correctly.

This is a result of the speckle noise in the hologram, which causes noise in the focus measure.

A further difficulty is that different particles may come into focus at different times.

It is therefore necessary to apply focusing algorithms to areas of the holograms limited to individual particles, and to find the point when the focus maximises for each area.

This requires some form of inter-plane particle-tracking.

Regions of Interest (ROIs)

To overcome the problem of multiple particles within a single hologram frame, ROIs are detected by identifying possible particles and enclosing them in a rectangle.

When analysing hologram reconstructions, ROIs define the target areas in which to maximise the focus measure.

Binarised Phase Deviation

ROIs can be generated by binarising images, and enclosing the larger black areas in rectangles.

This approach proved unreliable, as intensity variations in the recording beam can cause excess ROIs.

If the reconstructed *phase* is converted to an absolute phase deviation (from mean phase), the effects of illumination diminish.

Binarising the phase image gives sparser collections of pixels, making ROI generation more difficult.

A per-pixel OR over several reconstructions reinforces pixel groups, allowing ROIs to be generated as before.

After ROI generation, focus metrics are applied to ROIs.

ROI Tracking

ROIs generated on each plane are compared with those on the previous plane by area of **non-**overlap.

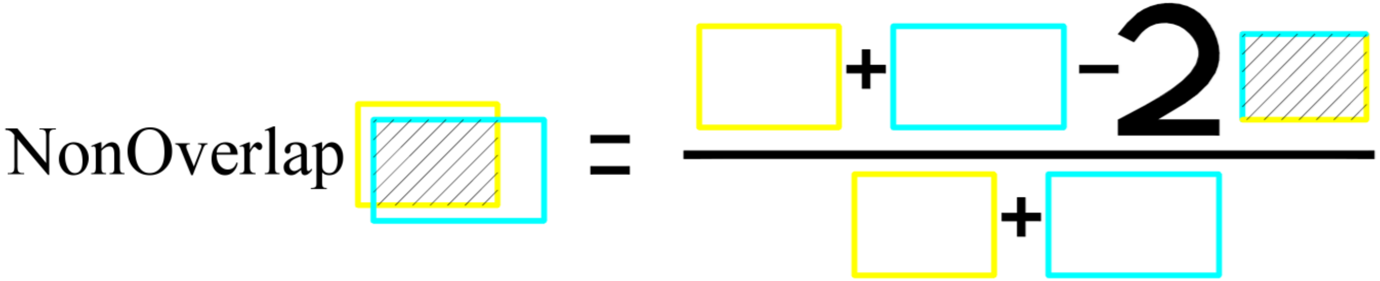

The area of non-overlap is characterised as a percentage value.

This measures the percentage of total area (both ROIs) that does not overlap.

For perfectly matched ROIs, the result is 0%

A ROI is assumed to enclose the same particle between reconstruction distances if the nonoverlap value is small.

Focus Measurement – Gradient

Focused images contain sharp edges: sudden discontinuities between neighbouring pixels. A 1-dimensional example is shown below:

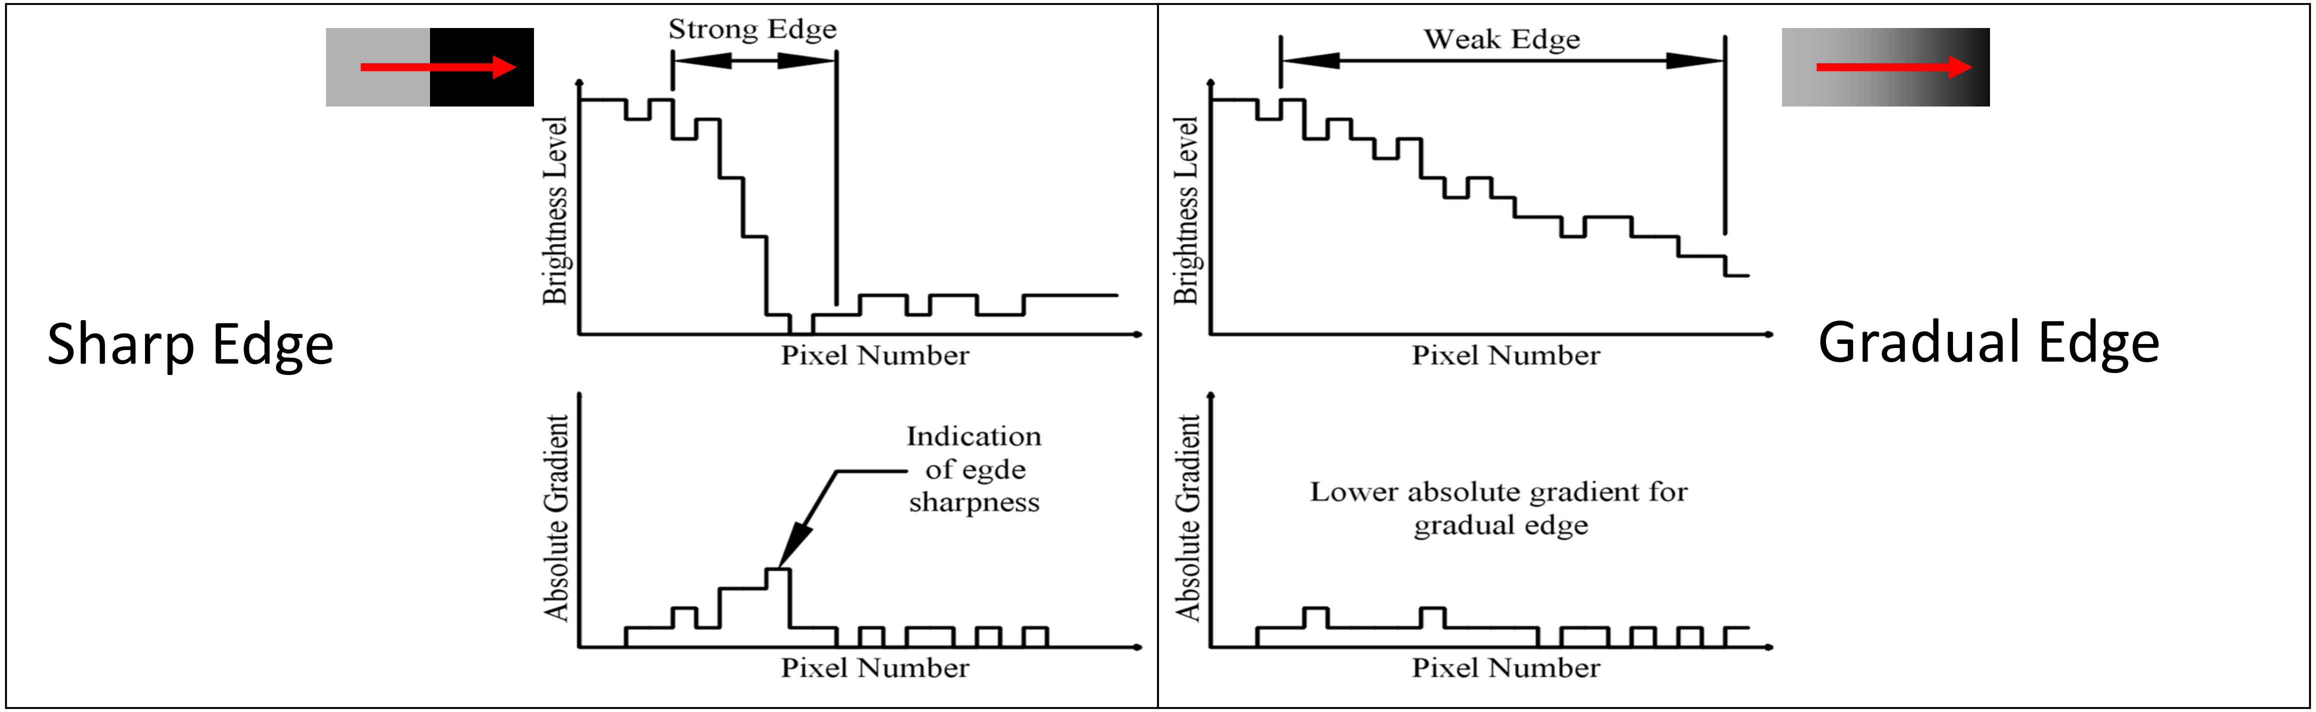

Gradient measures for 2D images must quantify gradient in both X and Y directions.

To obtain this measure, one of the best approaches is ‘Tenengrad’.

Focus Metrics - Tenengrad

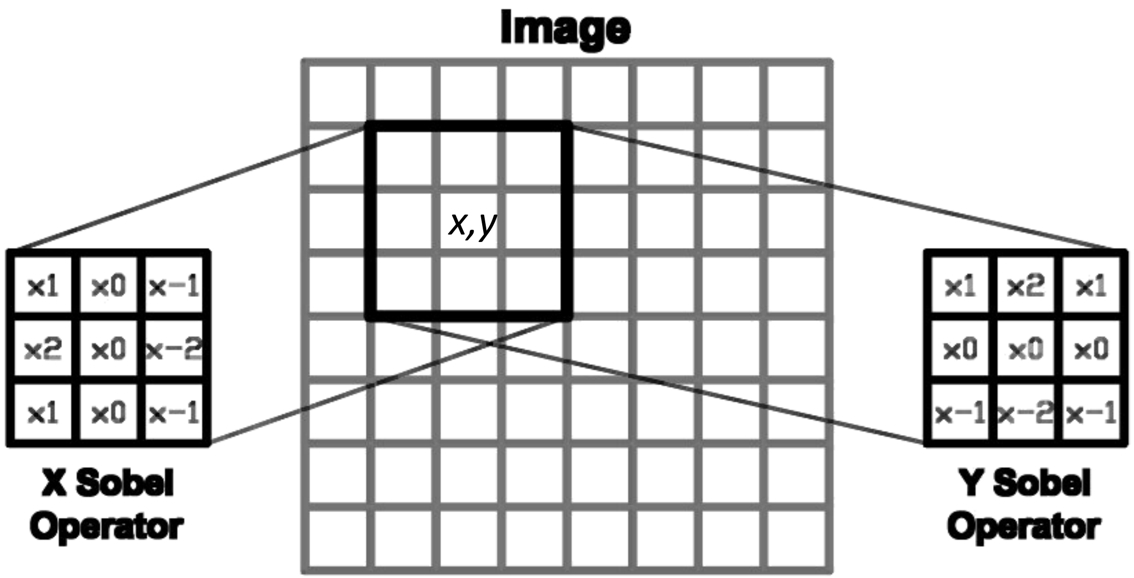

Tenengrad generates a gradient measure using X and Y Sobel convolution operators.

The absolute gradient is computed by the square-root term:

$$G_{xy} = \sqrt{S_{x(x,y)}^2 + S_{y(x,y)}^2}$$

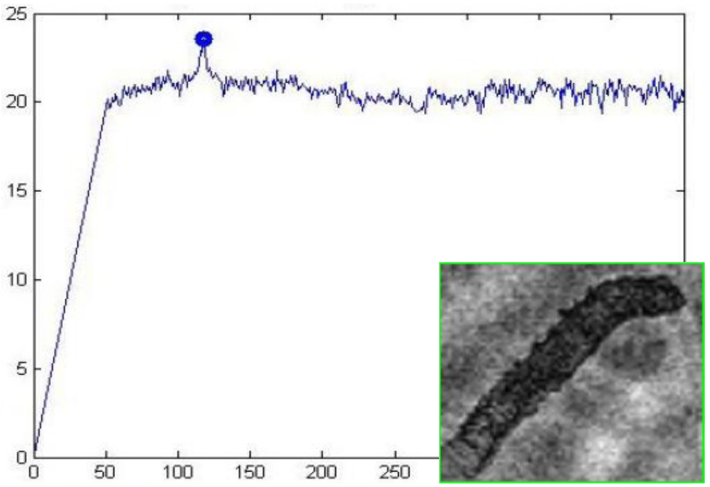

The Tenengrad measure is then computed as the sum of all gradients divided by the image size:

$$T = \frac{1}{XY} \sum_x^X \sum_y^Y G_{xy}$$

For many particles, this approach fails to work.

Tenengrad Improvement

The perfomance of the Tenengrad measure can be improved by introducing a threshold.

The threshold is applied to the square-root term:

$$G_{xy} = \sqrt{S^2_{x(x,y)} + S^2_{y(x,y)}}$$

If the gradient value is below the specified threshold, then it is taken as zero.

This helps to filter out the lower-level constant gradients present due to image speckle.

Unfortunately, the ideal threshold varies from hologram to hologram, and it is difficult to find a value that works well on all holograms!

An Alternative Autofocus

Despite improvements from thresholding, Tenengrad still often performs poorly.

To overcome this problem, a variation to Tenengrad was developed: Contour Gradient.

The new approach utilises contours to localise the application of the Sobel operators to where the particle edges are believed to be located.

The reconstructed image is first analysed with a contour algorithm to enclose any likely particles.

The Sobel operators are then applied along the contour and the summed gradient divided by the length of the contour.

$$F_{CGrad} = \frac{1}{K}\sum_{k=0}^{K} Gx_{P_k}^2 + Gy_{P_k}^2$$

Contour Gradient

A video showing the scanning and extraction processes can be seen below:

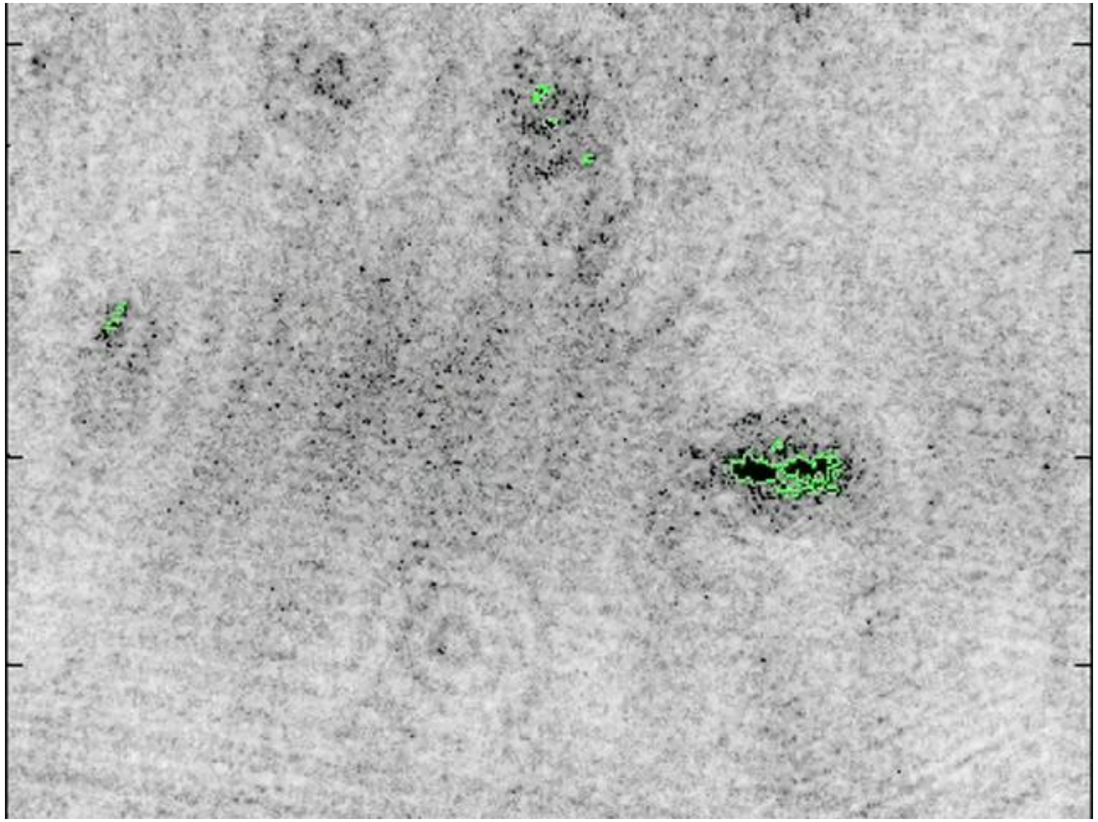

Scan

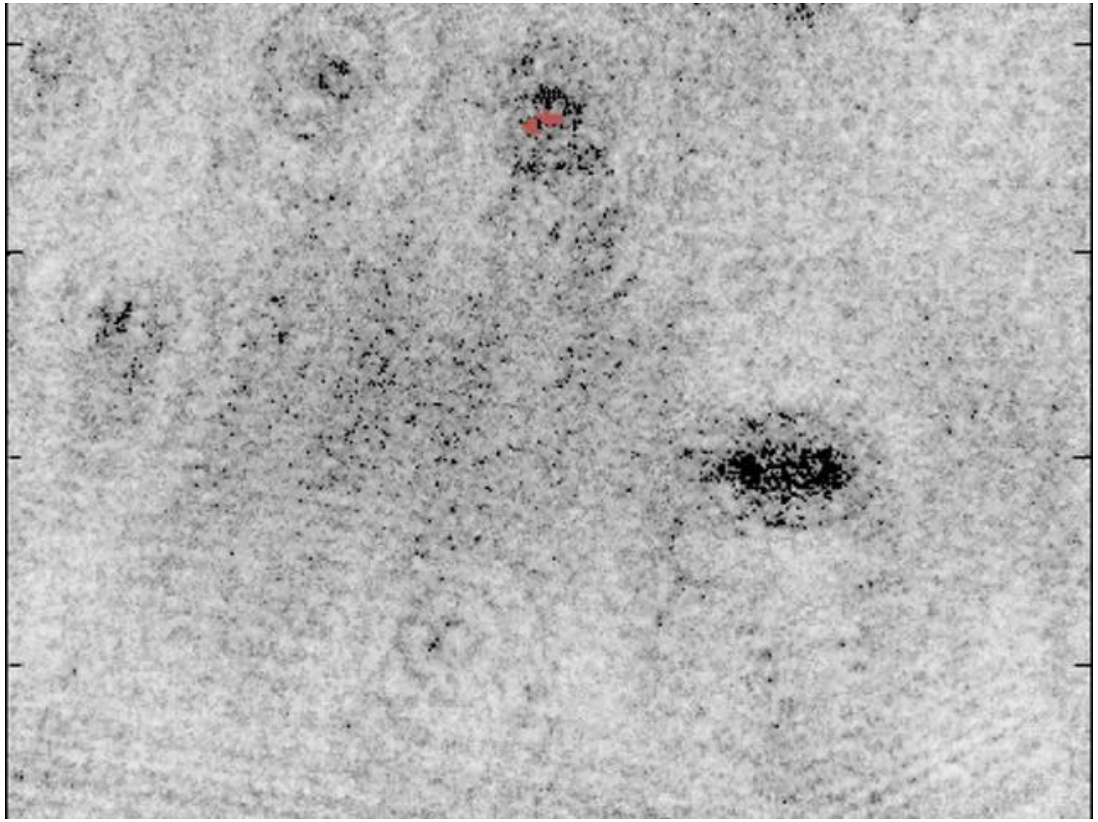

Extract

Contour Grad vs ROI-Tenengrad

Sample Images

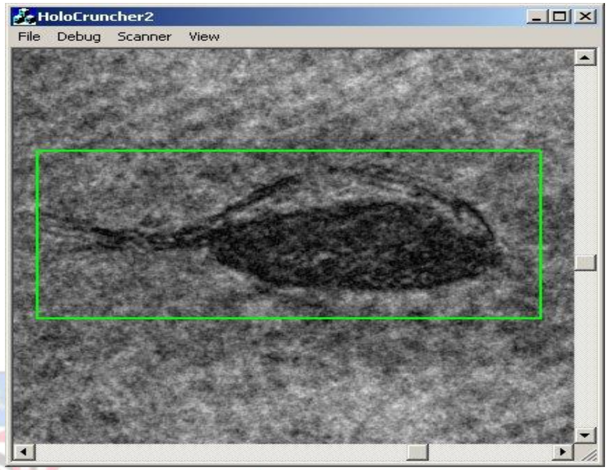

ROI-Tenengrad Metric

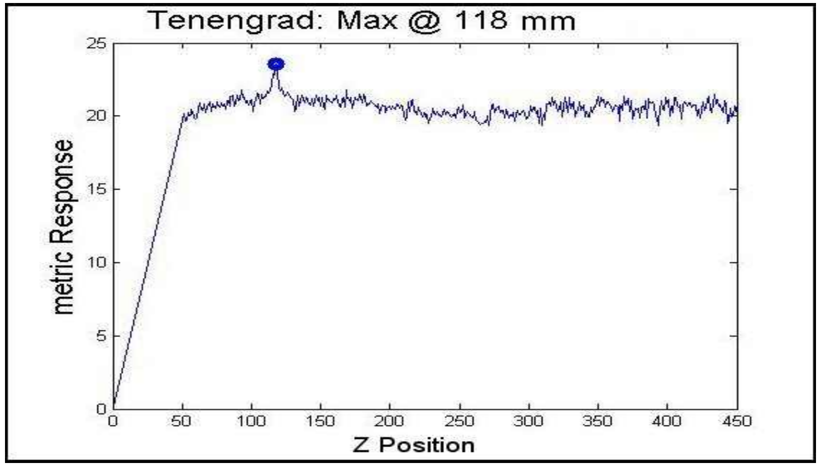

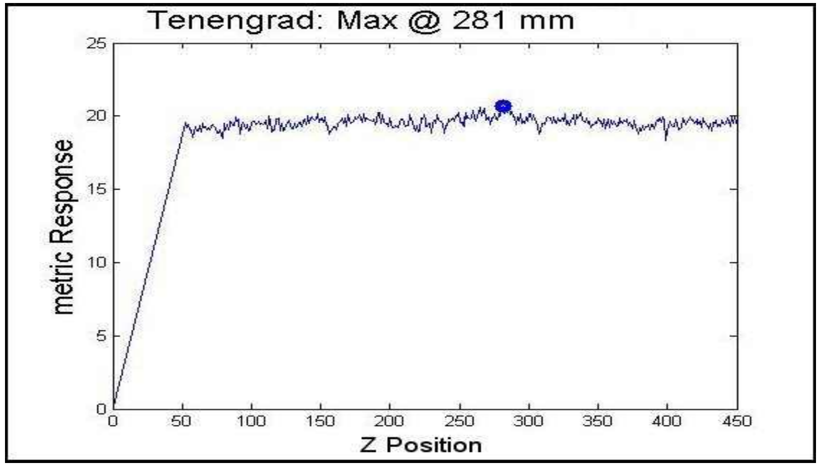

Contour Gradient Metric

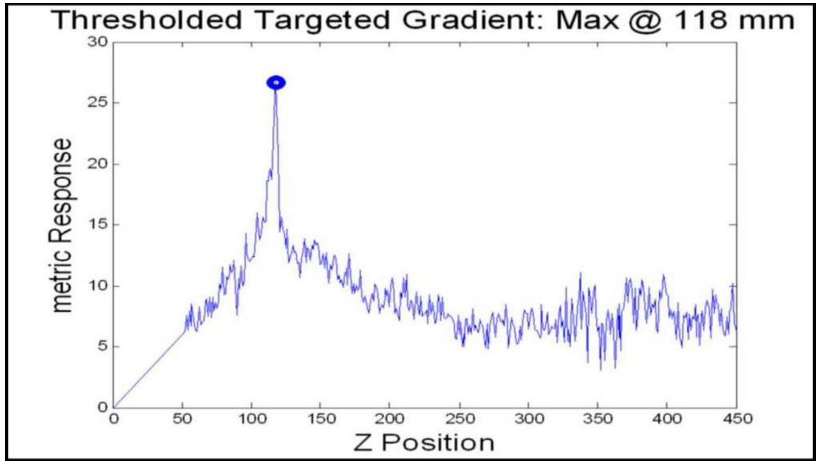

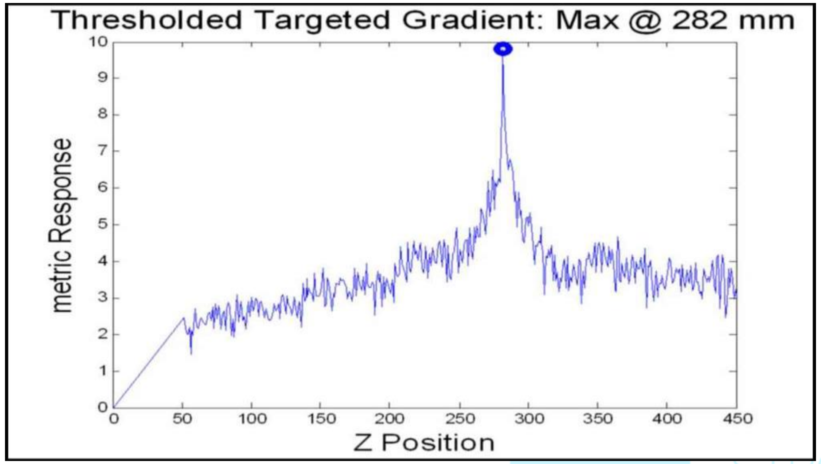

Continuing Research

Even with automated extraction in place, there remain problems to be overcome.

The reconstruction algorithms in use all rely on the Fast Fourier Transform.

While this is indeed quite fast, it still takes several weeks to scan a single holovideo of ~1000 frames.

Reconstructing a single slice of a full-size (2208 x 3000) hologram takes in the order of 3 seconds. The full volume is ~400 slices, and there are ~1000 holograms per video.

Even without accounting for extra processing overhead, this amounts to about 14 days of processing! We need faster reconstruction methods.

After analysis, many hundreds of images may be outputted and require manual identification. Further research is required to automate the identification of plankton species.

Thank you!

Any Questions?