University of Glasgow | College of Science & Engineering

Full-Colour, Computational Ghost Video

Miles Padgett
Kelvin Chair of Natural Philosophy

A Quantum Ghost Imager

- Generate random photon pairs, illuminate both object and camera

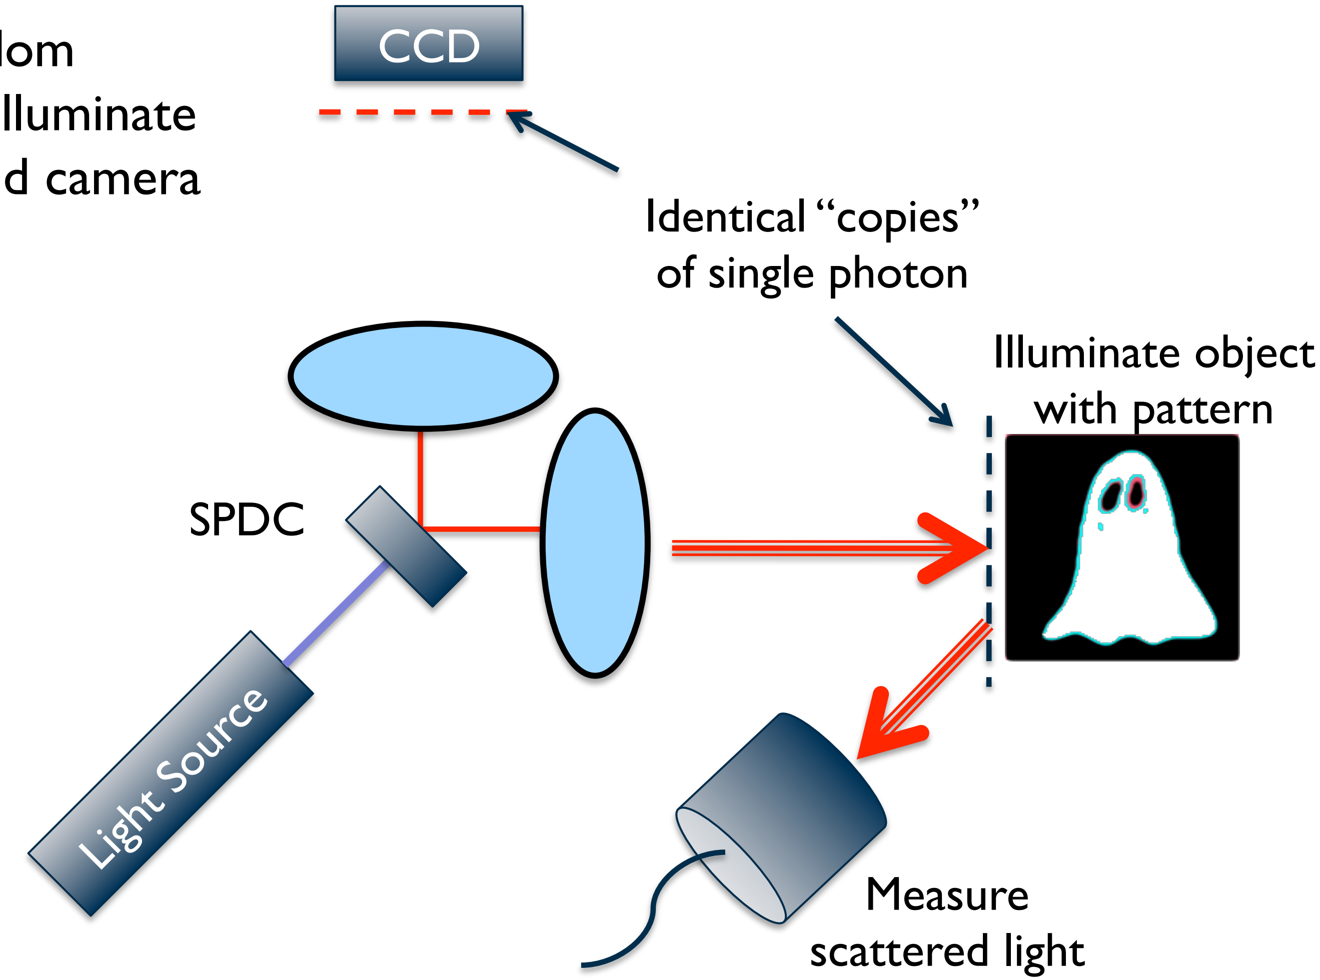

A Classical Ghost Imager

- Generate random speckle, illuminate both object and camera

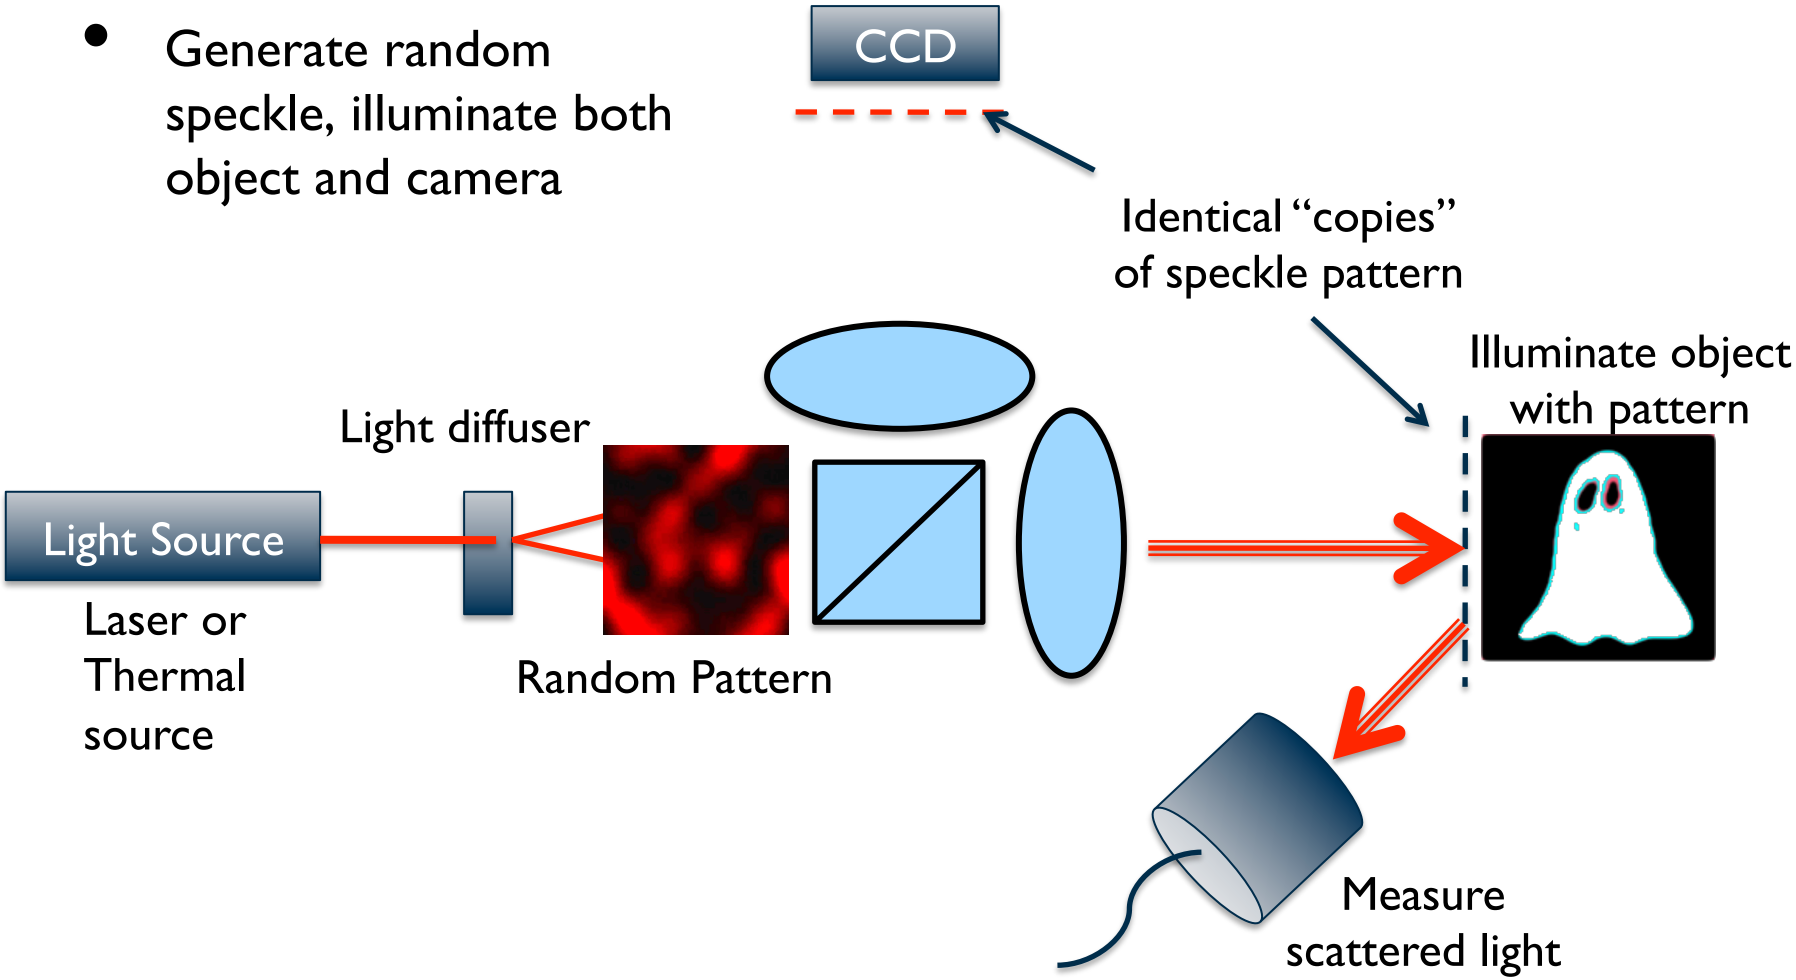

A Computational Ghost Imager

- Generate deterministic speckle using spatial light modulator, no need for CCD – the computer already knows!

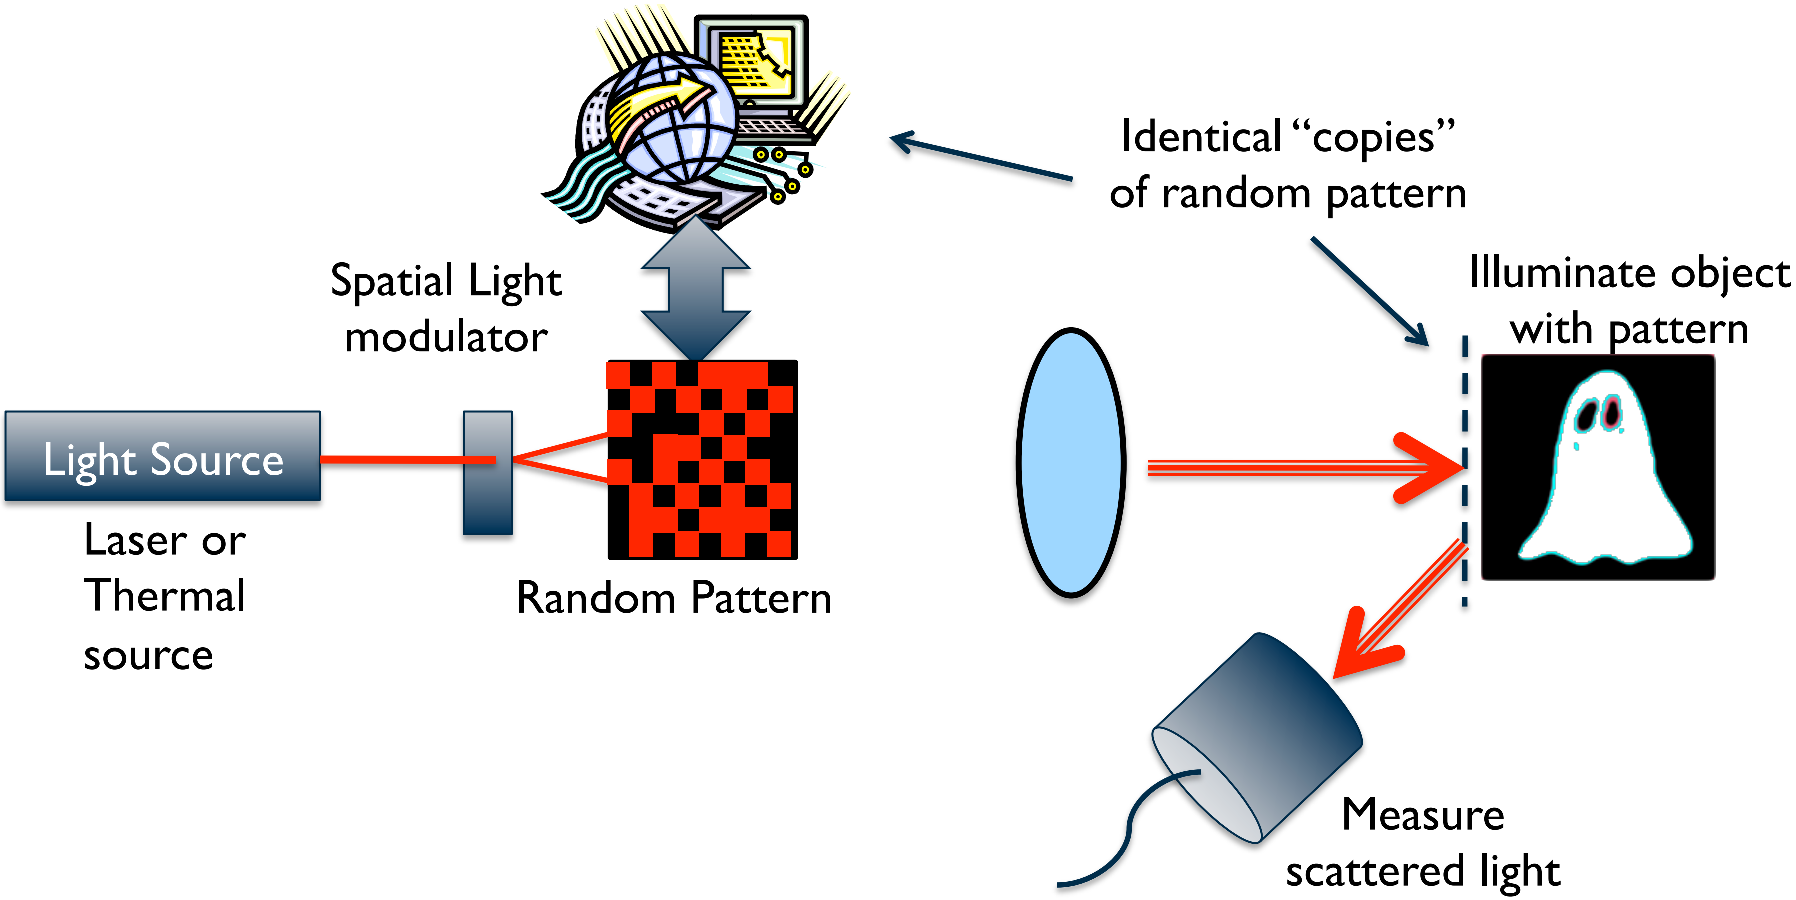

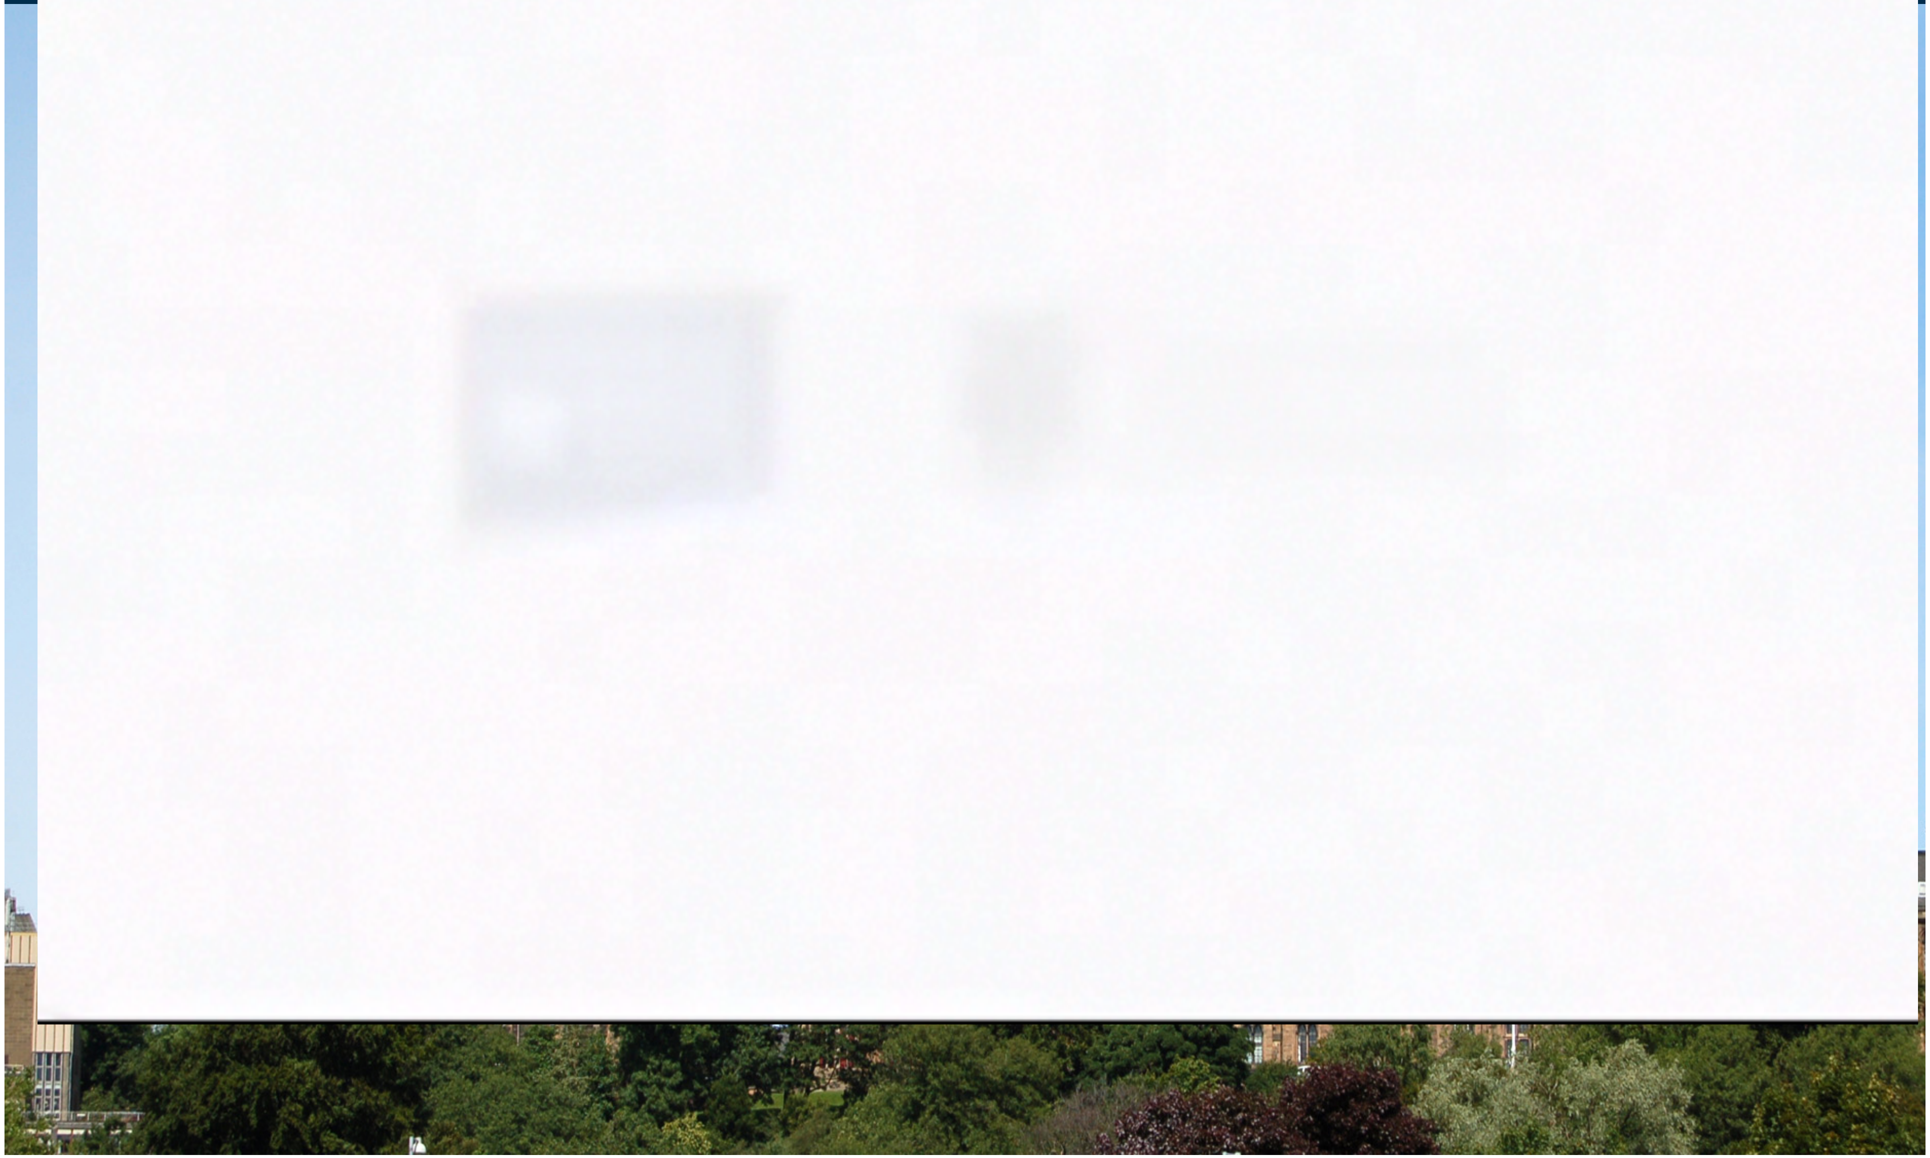

3D Computational Ghost Imaging

Dr. Matthew Edgar

Mr. Baoqing Sun

Mr. Stephen Welsh

Dr. Richard Bowman

and L E Vittert

Prof A Bowman (Statistics)

Ghost imaging with classical light

Experimental setup for 3D computational ghost imaging

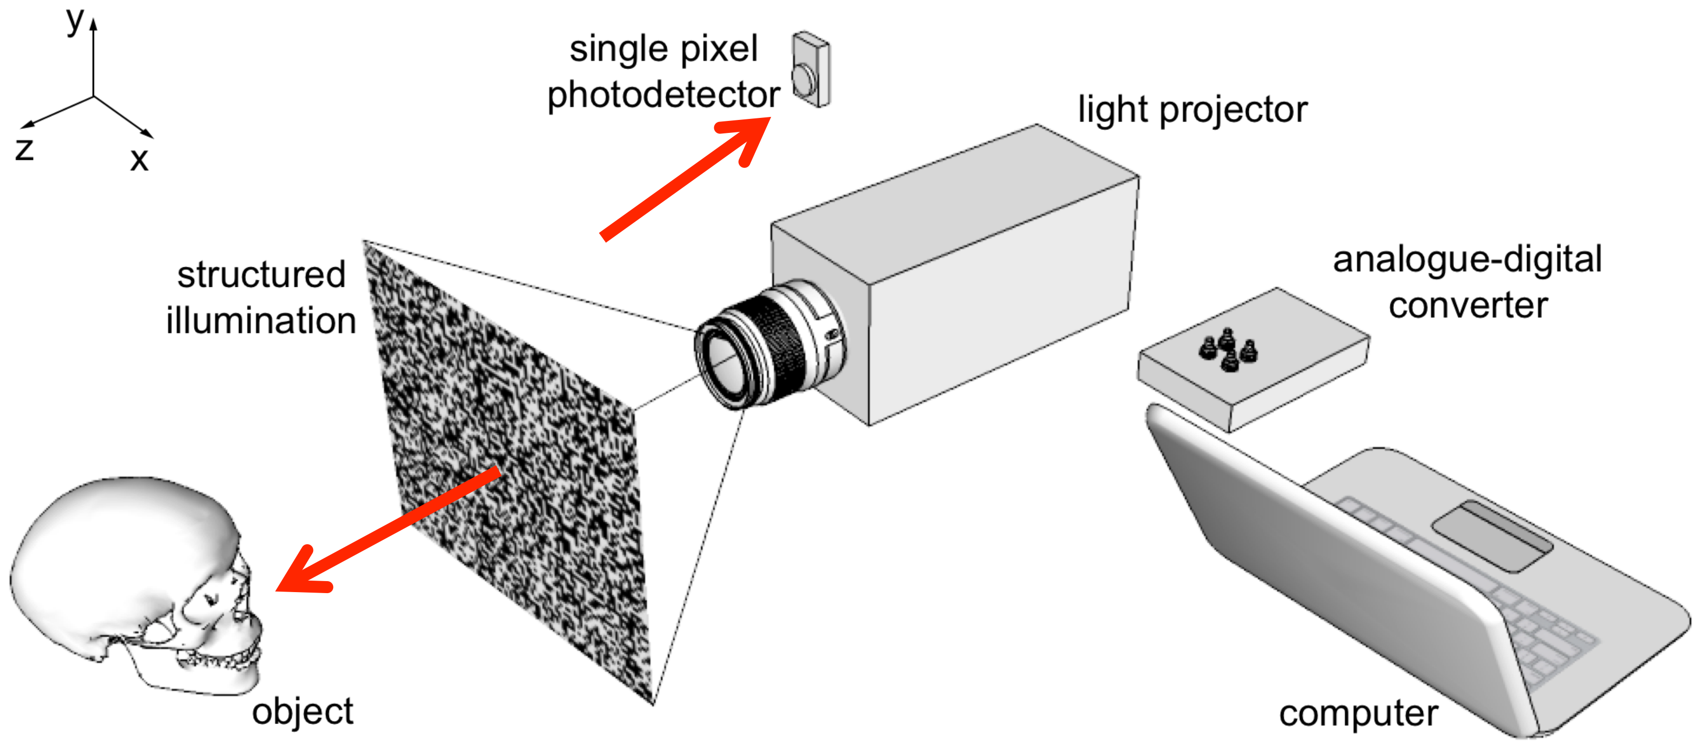

Projecting a series of random pattern

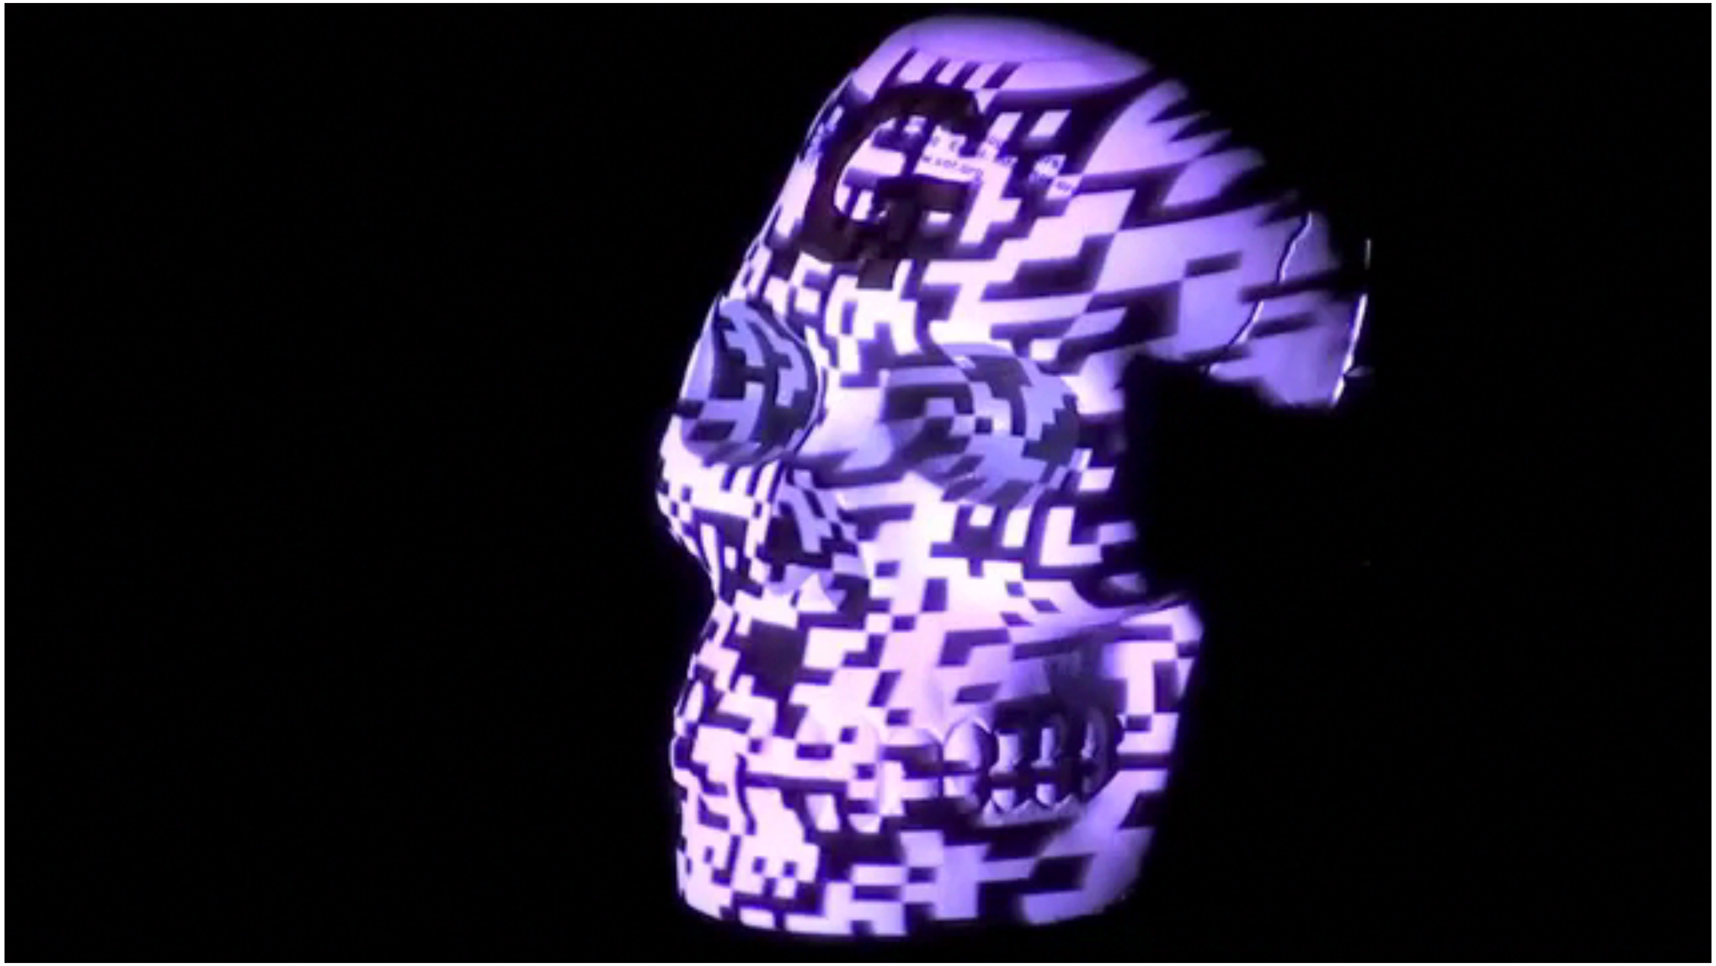

University of Glasgow

Traditional Ghost Imaging

$$\sum_{N}$$

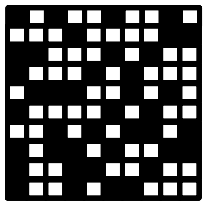

random pattern X

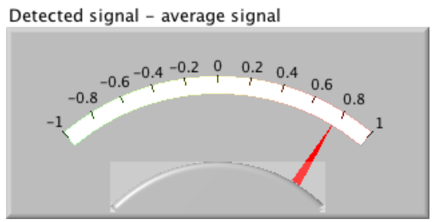

(detector signal - average detector signal)

≈ Need N different patterns to give N pixel image

Or use "compressive" techniques (c.f. JPEG) to do better!

Traditional Ghost Imaging

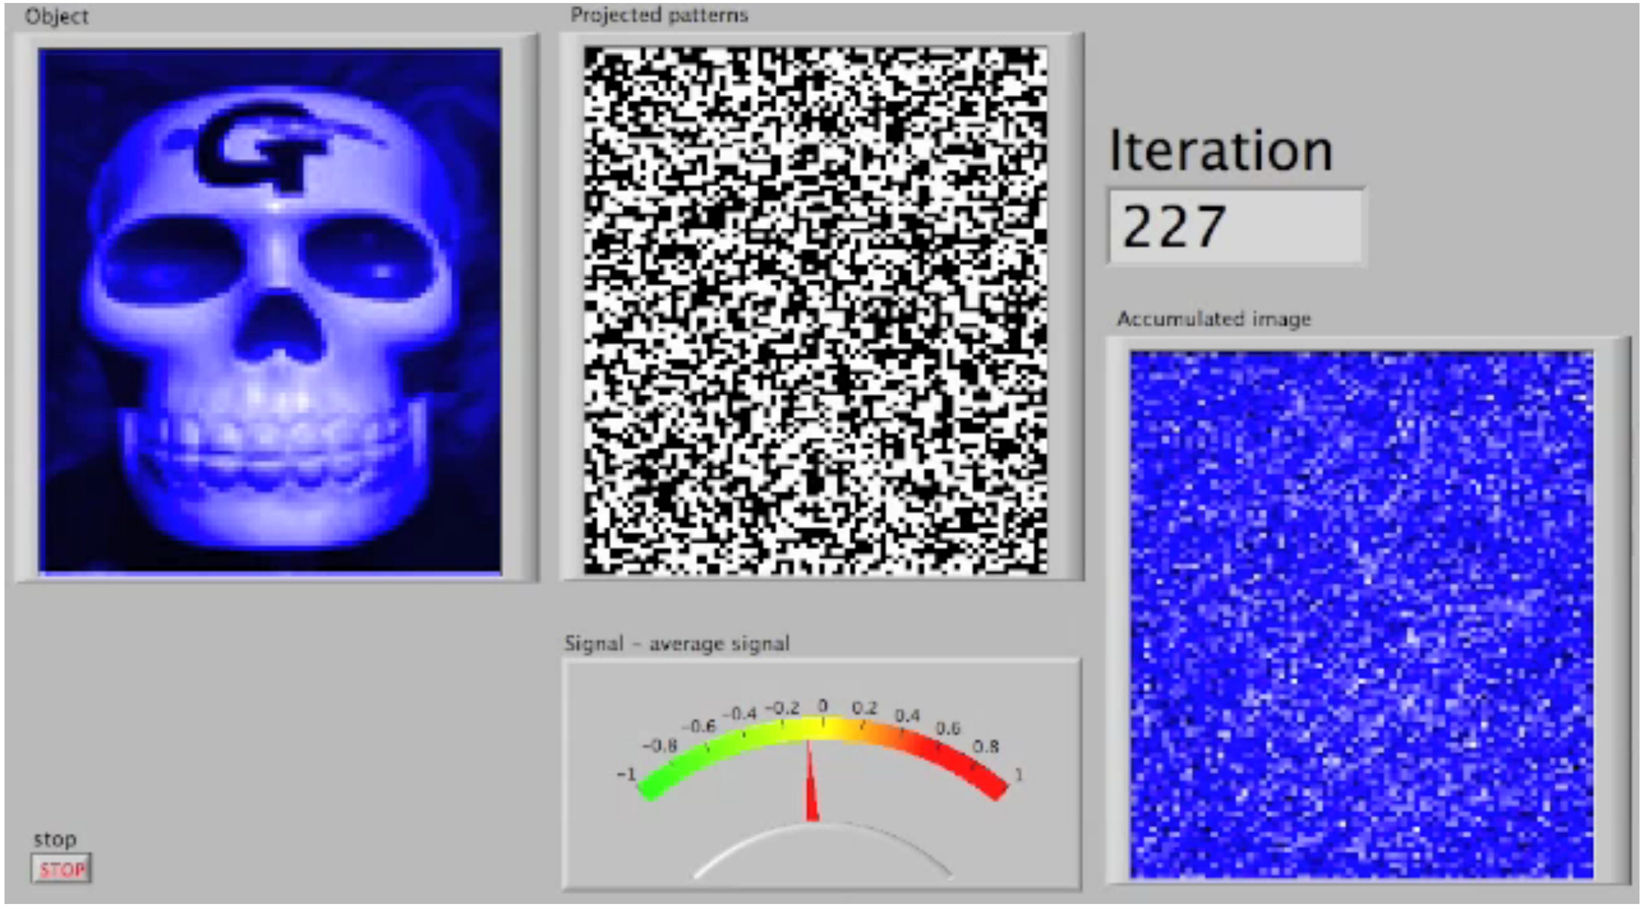

University of Glasgow

Ghost imaging with classical light

Iterative reconstruction of 2D image

Test object (toy skull)

How many measurement does an image require?

Reasonable to assume no. of measurements = no. of pixels

When the number of unknowns exceeds the amount of data then many different solutions fit the data perfectly!   i.e. $\chi^2/N = 0$

But (in the presence of noise) it is very unlikely that your measured all data was perfect.   Much more likely is that $\chi^2/N \approx 1$

So of all these possible solutions (images) which one should you pick?

Real images
- Only have positive intensities
- Can be JPEG compressed (they are sparse in spatial frequency)

"Least squares fitting" is a necessary, but not sufficient, strategy

Compressive Ghost Imaging

64x 84 (=5376) pixels

Object & example of random illumination

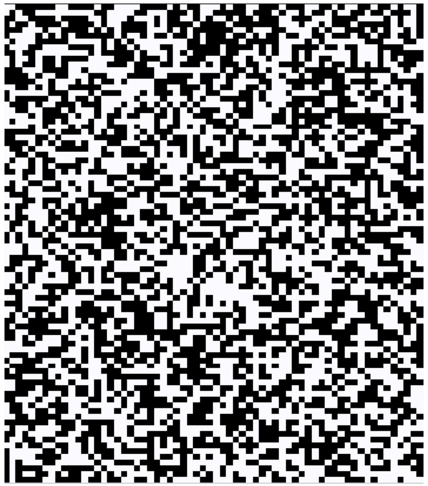

2700 patterns

Recon by Chi-squared minimization

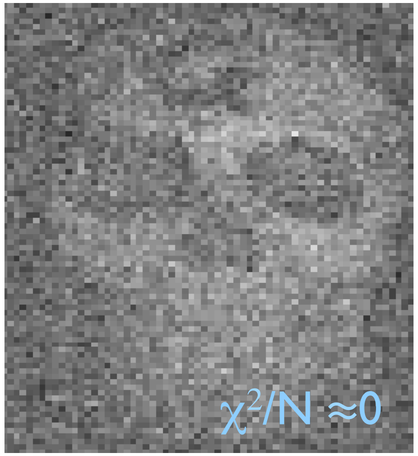

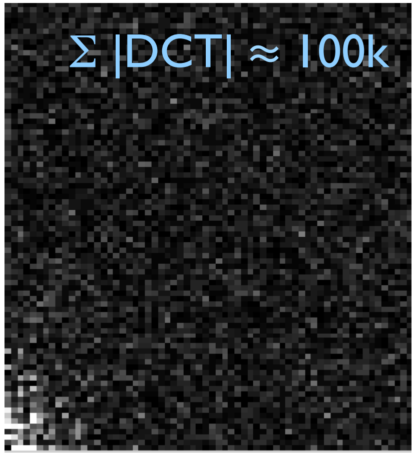

Chi-squared solution + regularization

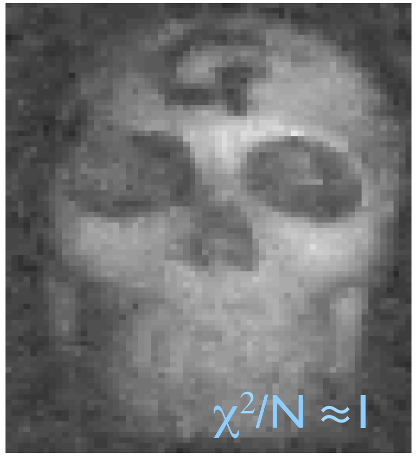

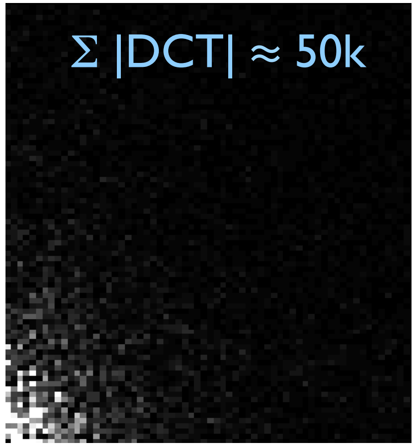

How to optimise (Random Search)

Calculate Cost of "solution"

$\chi^2/N + \lambda \ \times R = \text{cost}$

Keep making small changes to image

R is a measure of Entropy or DCT sparsity, or total variation etc.

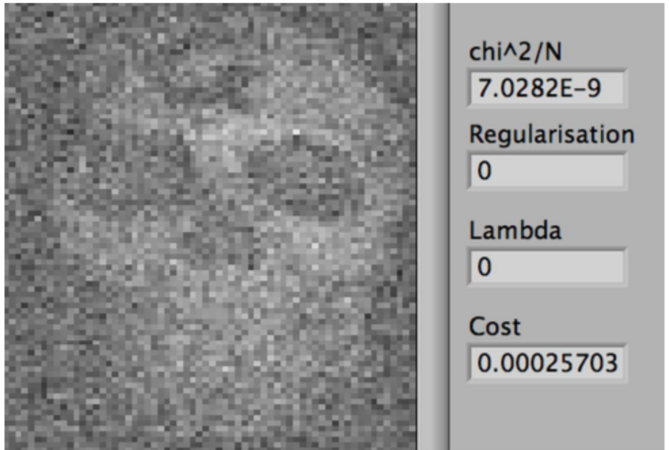

Recalculate Cost

$\chi^2/N + \lambda \ \times R = \text{cost}$

Repeat, keeping only those changes which reduce cost

Many possible solutions have an acceptable, χ^2 so pick the one of these that has other properties too…..

Normal Imaging with “off-axis” illumination

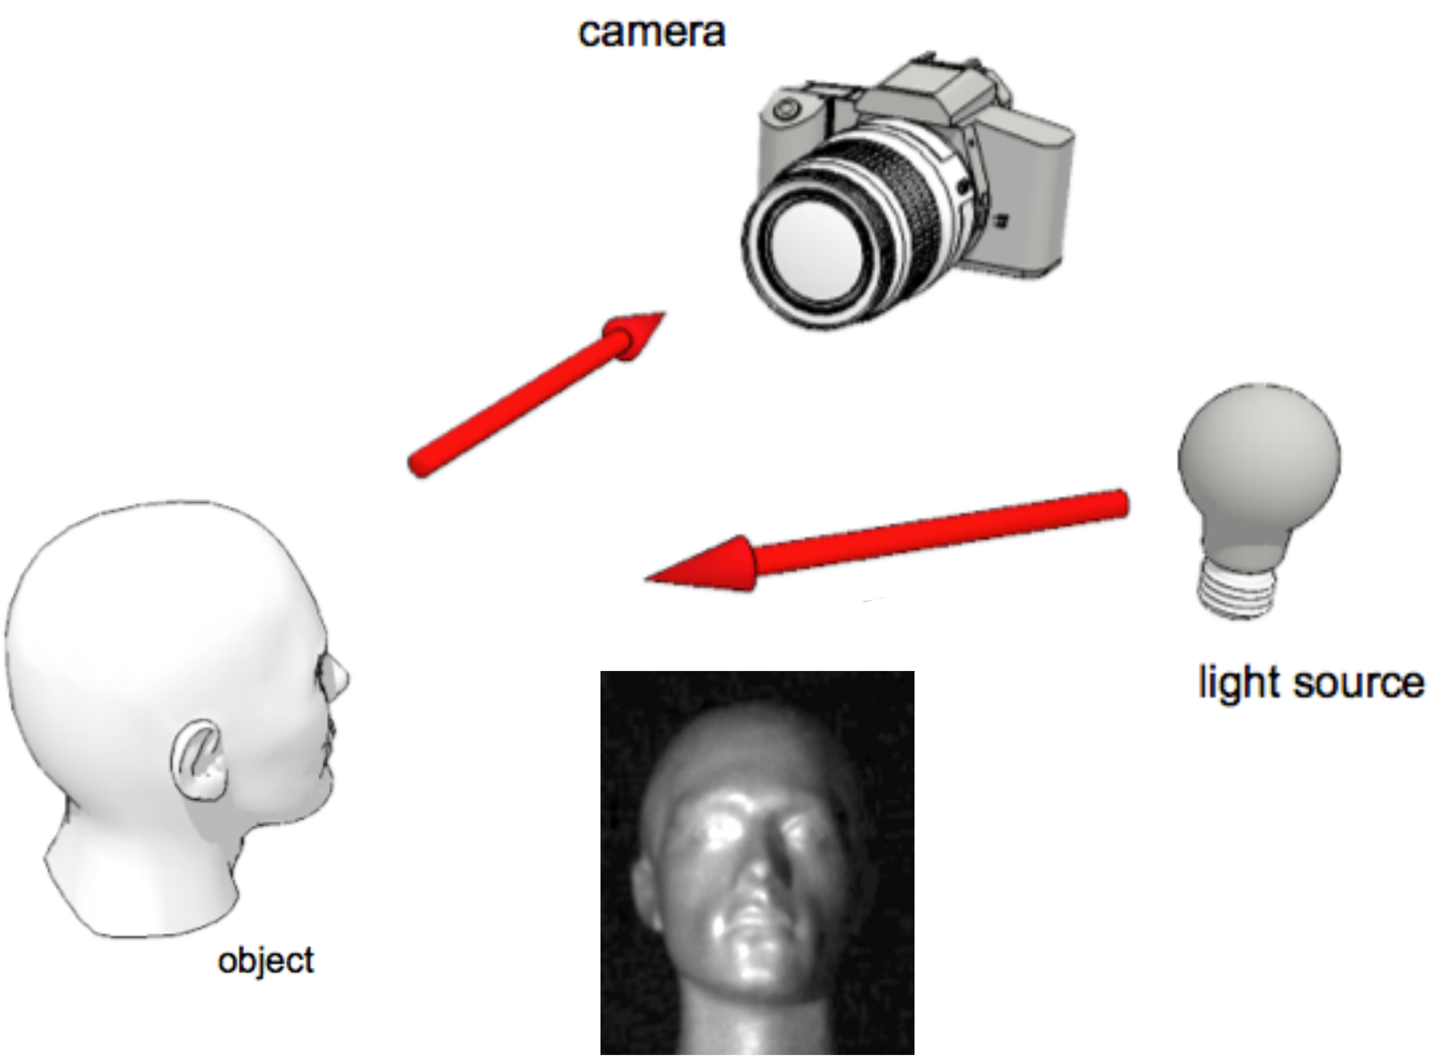

Ghost Imaging with “off-axis” detection

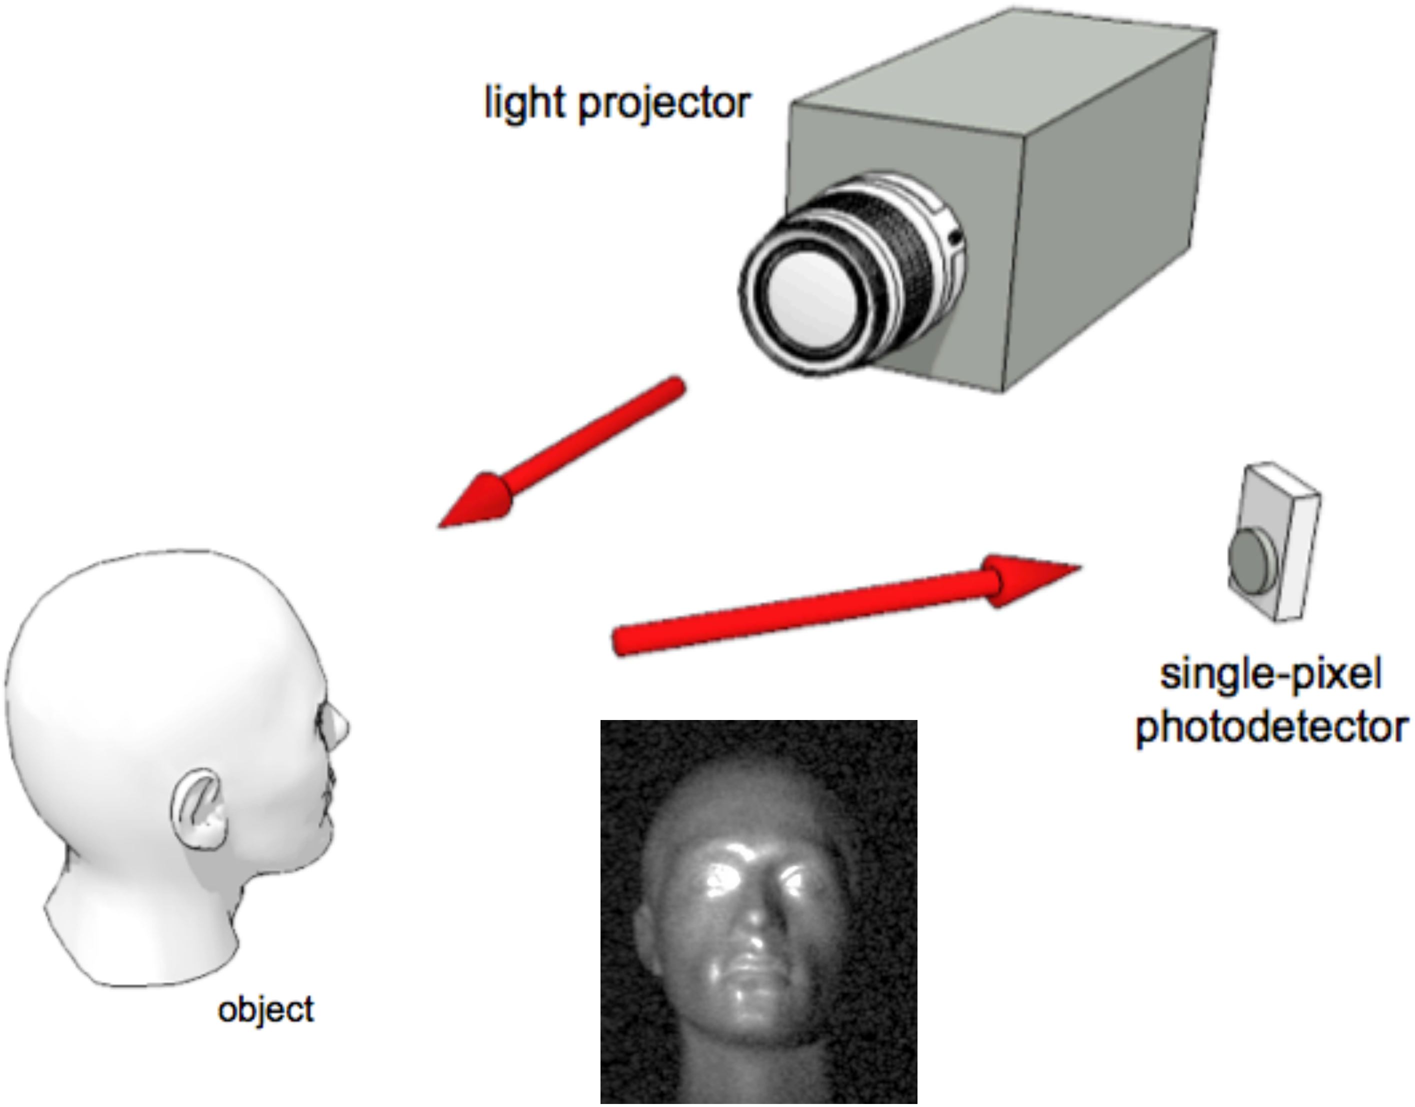

3D Ghost Imaging with classical light

Experimental setup for 3D computational ghost imaging

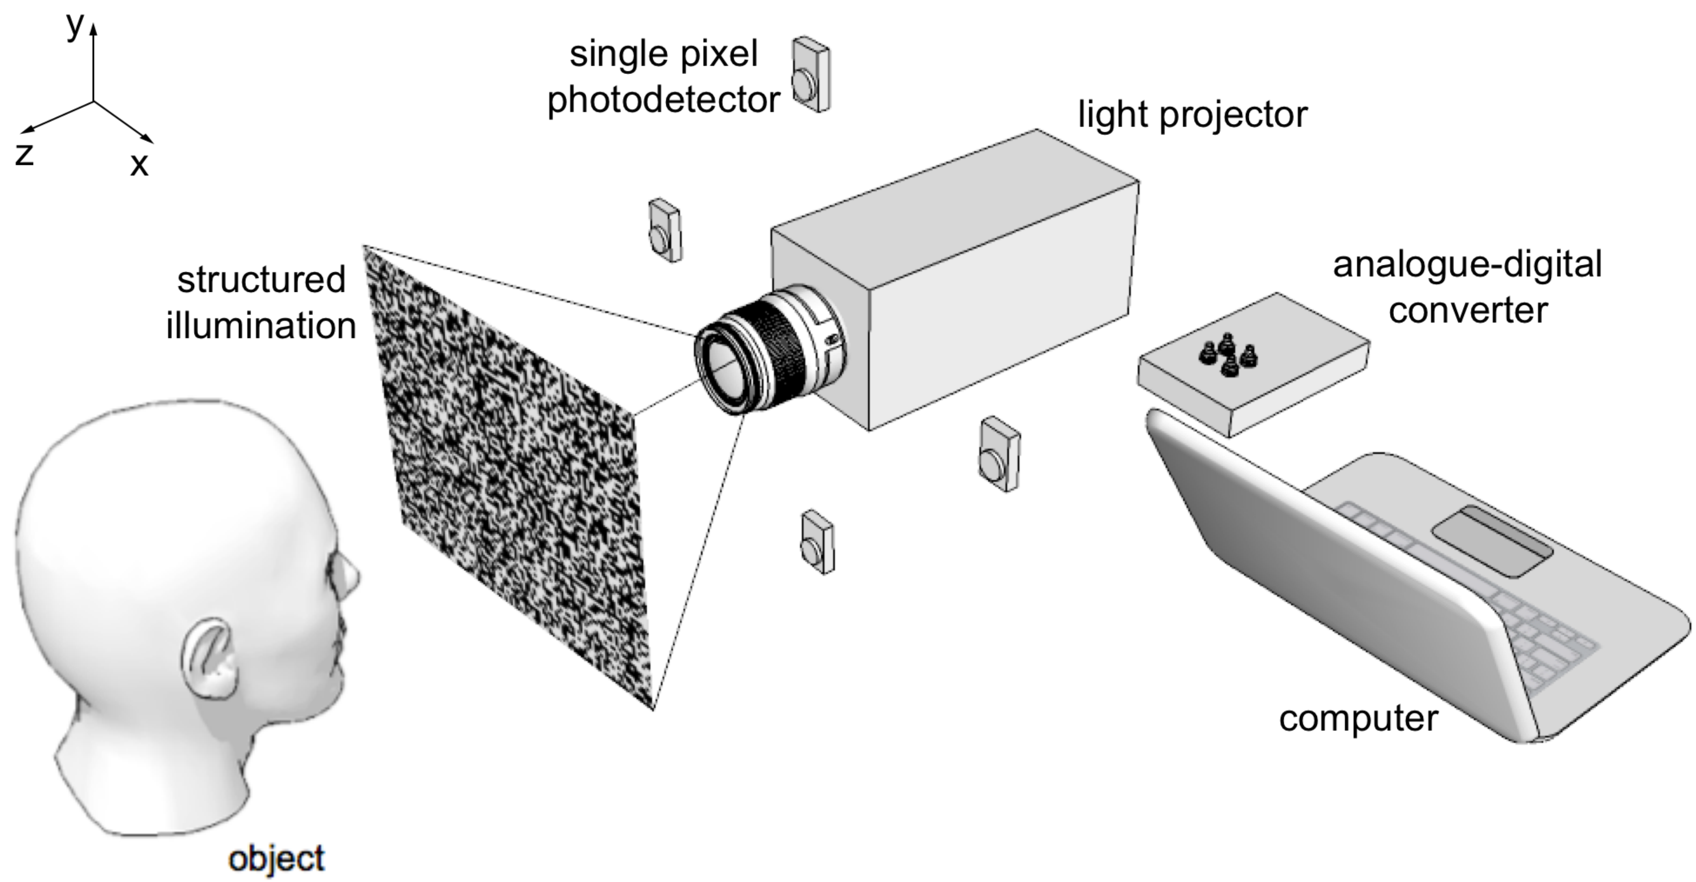

3D Ghost Imaging with classical light

Display random patterns at 660Hz

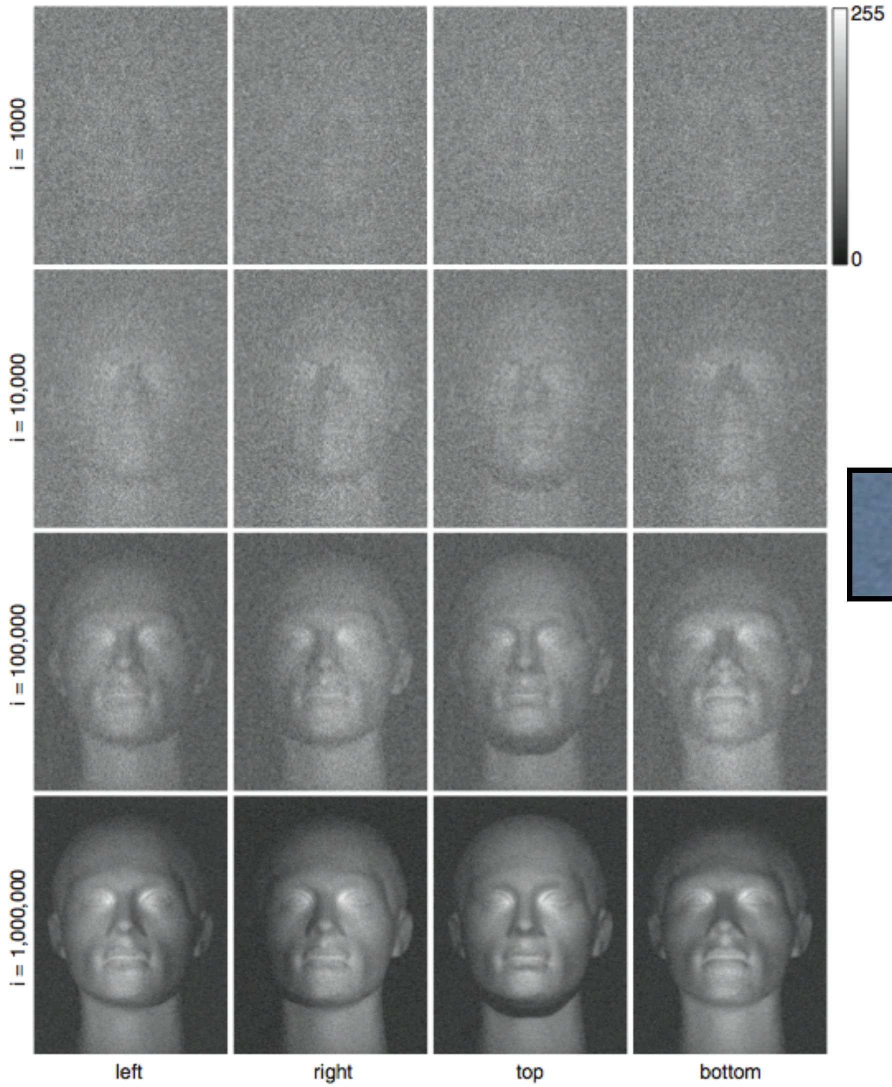

Surface Gradients

Integrate gradients (and optimize) to give surface profile

3D Computational Ghost Imaging

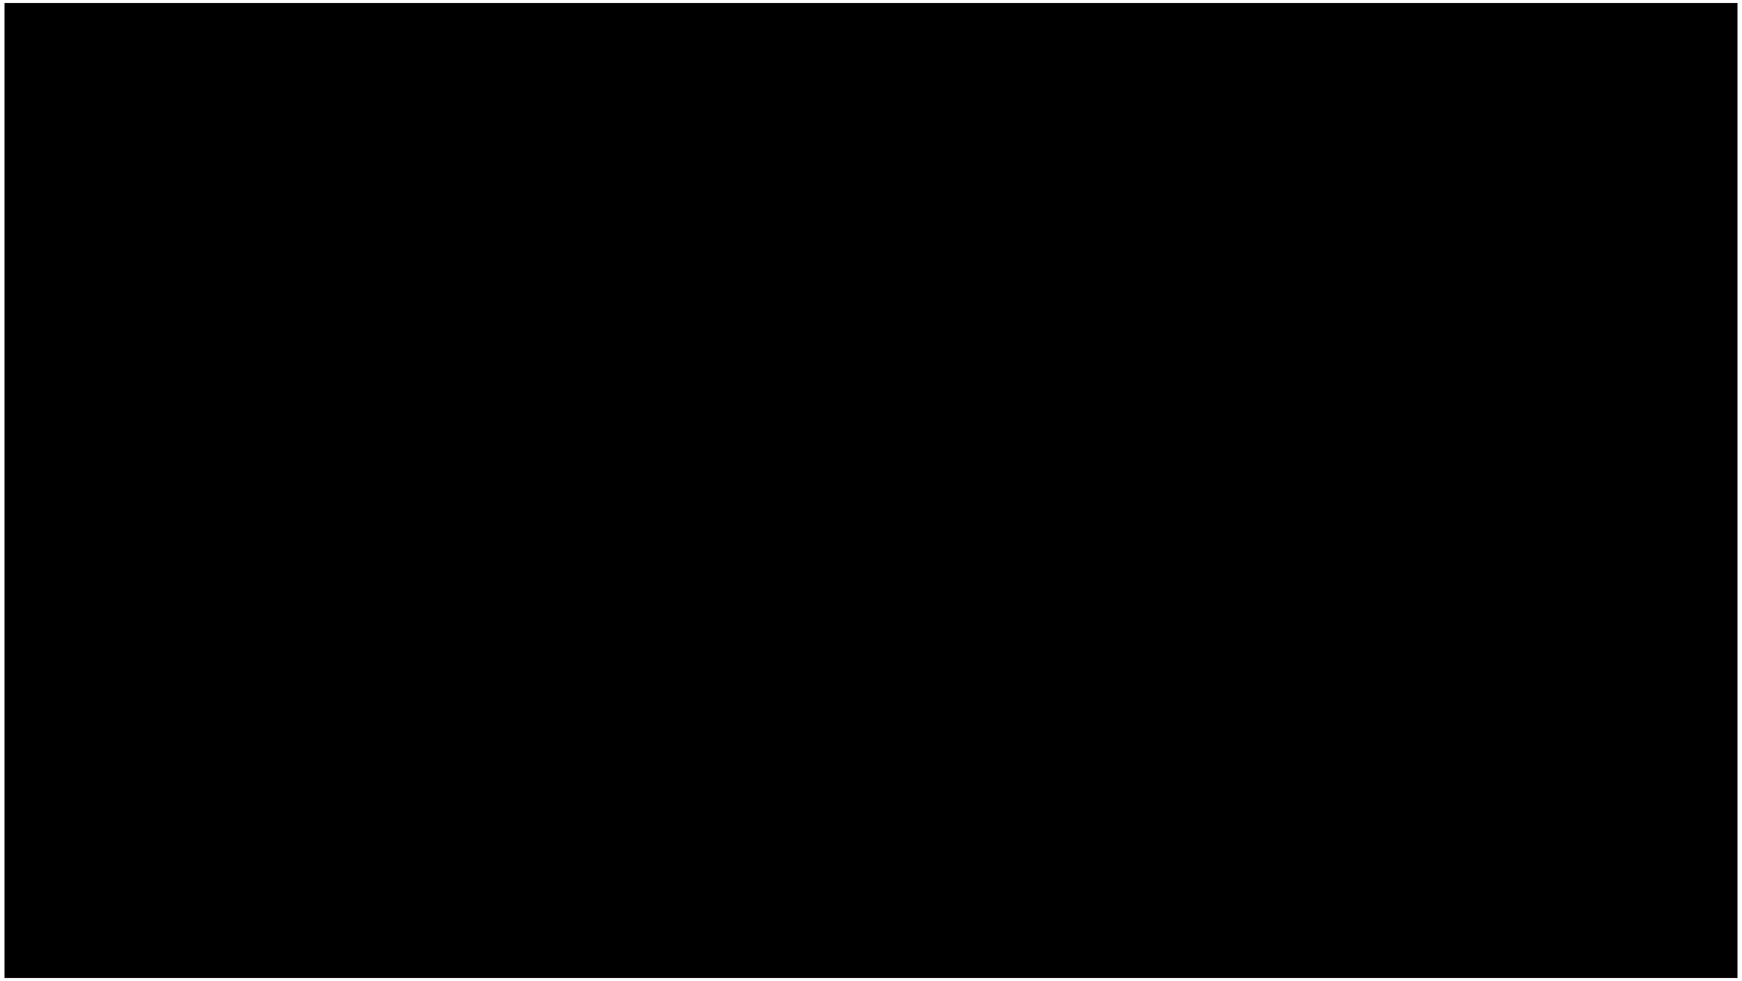

Calculating surface gradients

Right

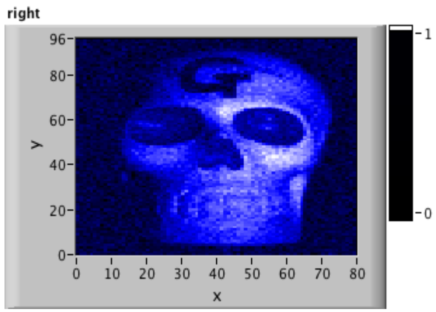

-

Left

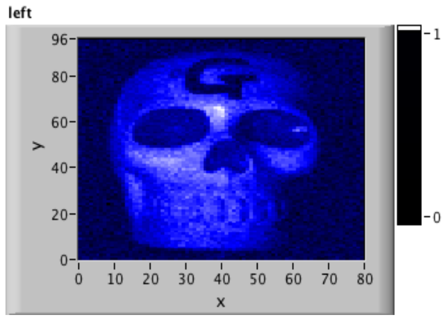

Right

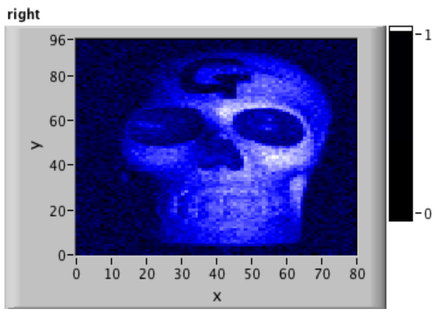

+

Left

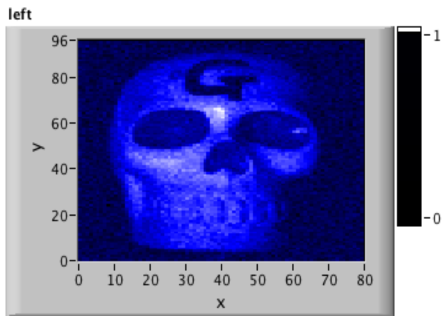

≈

dz/dx
Surface gradient

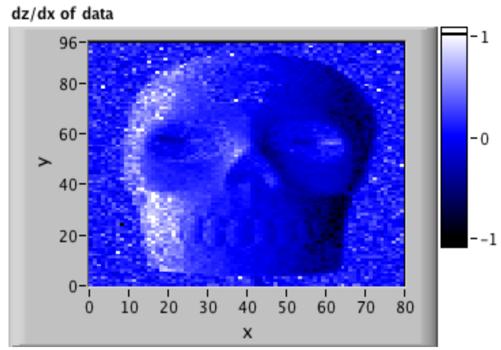

Works with real images too, "shape from shade"

Calculating Surface height

dz/dx
Surface gradient

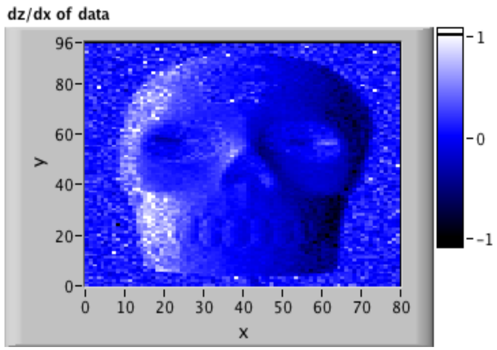

dz/dy
Surface gradient

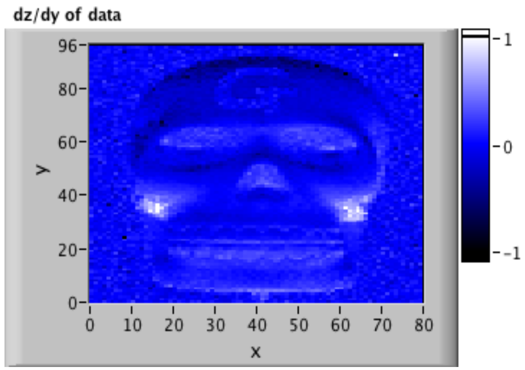

Integrate to give surface height (z) – but with what boundary condition?

z(x=0)=0 , or z(y=0)=0 etc

How to optimise

Make random change to surface

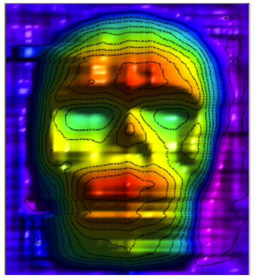

Recalculate gradients

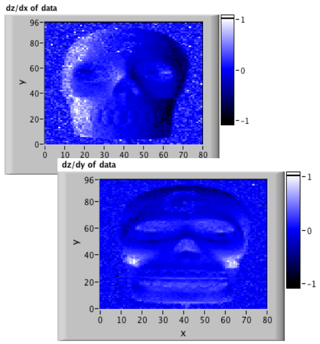

Compare to measured gradients
Calculate χ^2 and add regularisation to give total cost of solution

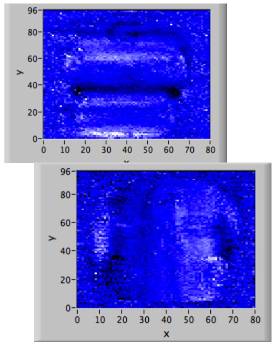

$\chi^2 + \lambda \times R = \text{cost}$

Repeat, keeping only those changes which reduce cost

Regularisation is properties of "solution" we'd like to "encourage" e.g. flatness (i.e. minimise sum of z) and/or smoothness (i.e. minimise sum of d^2z/dx^2)

Set λ at sensible value….

Calculating Surface height

Average over several possible boundary condition

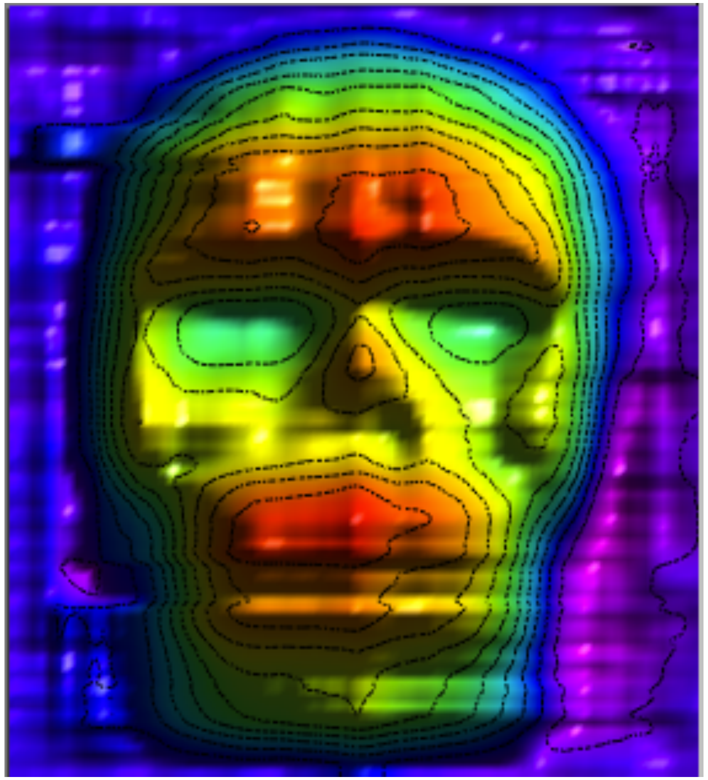

Apply optimisation

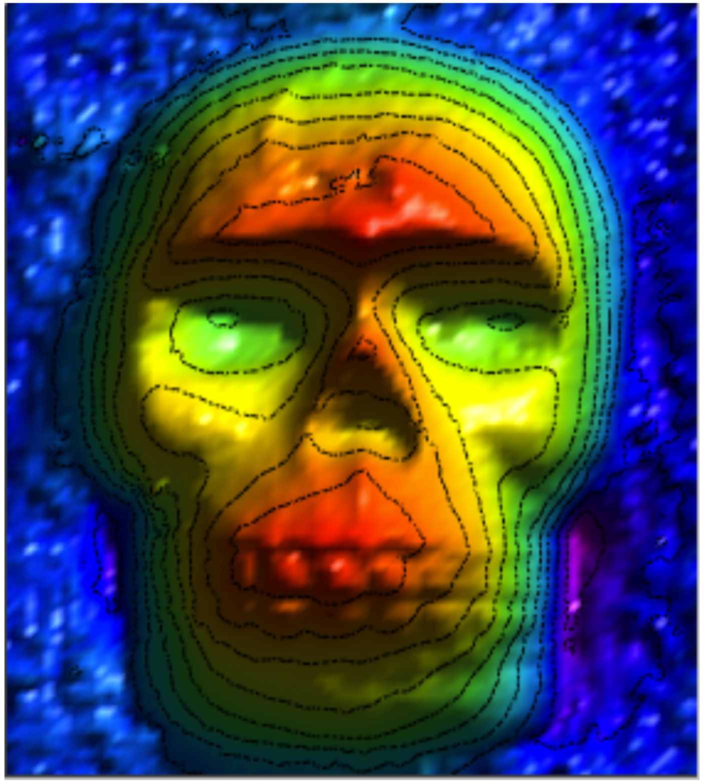

3D Computational Ghost Imaging

Colour Computational Ghost Imaging

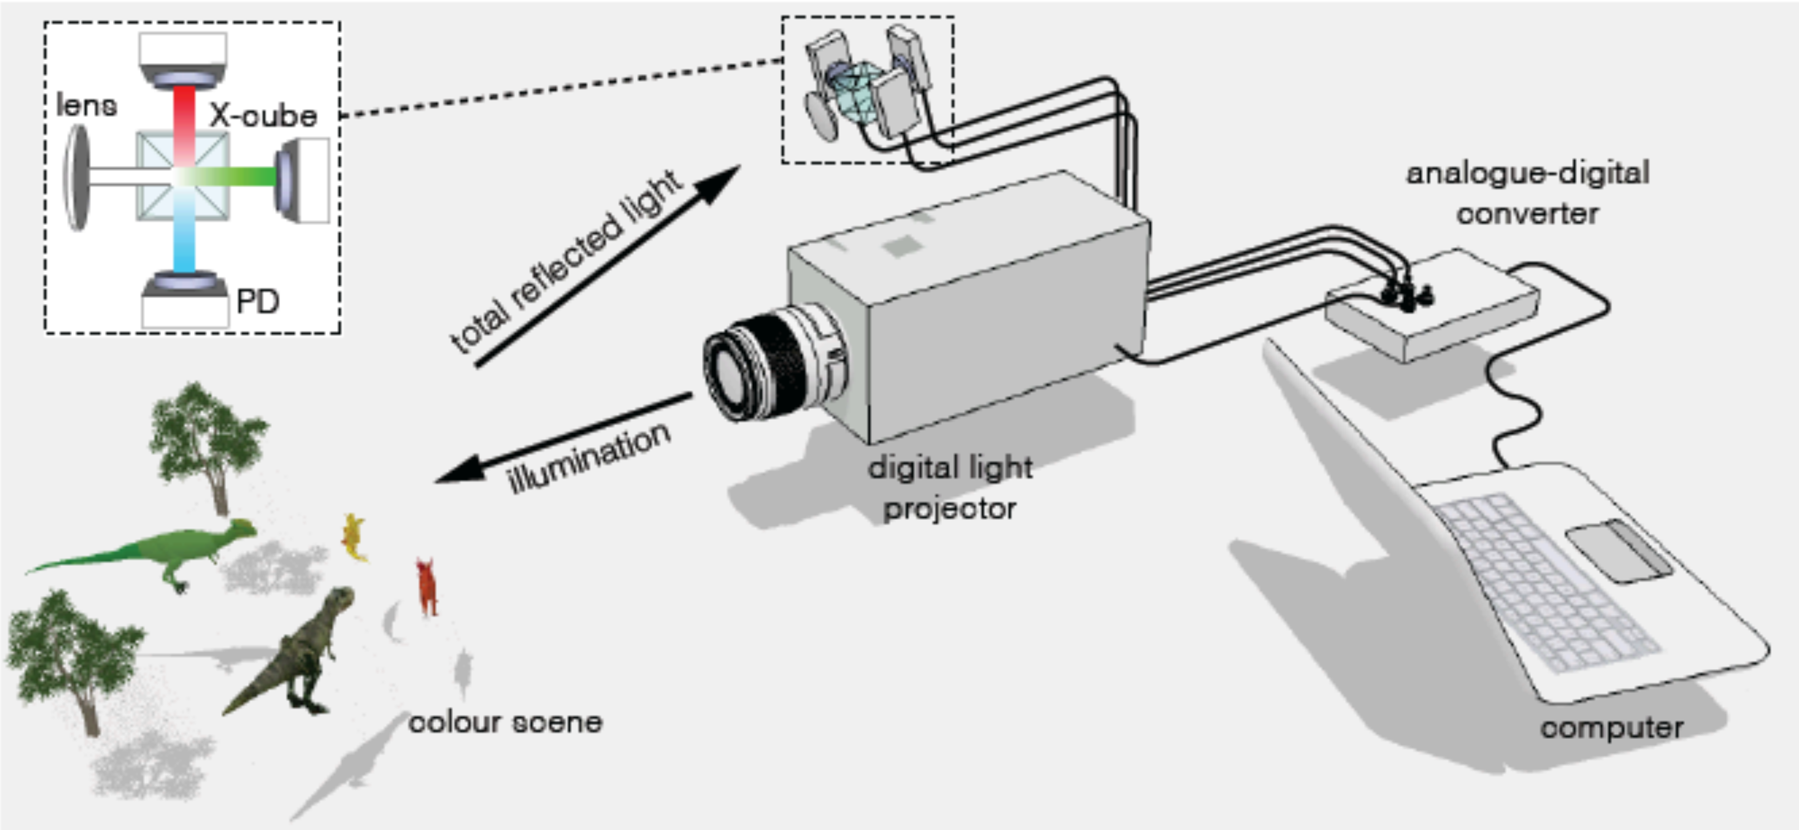

Colour Computational Ghost Imaging

A Computational Ghost Imager

- Generate deterministic speckle using spatial light modulator, no need for CCD – the computer already knows!

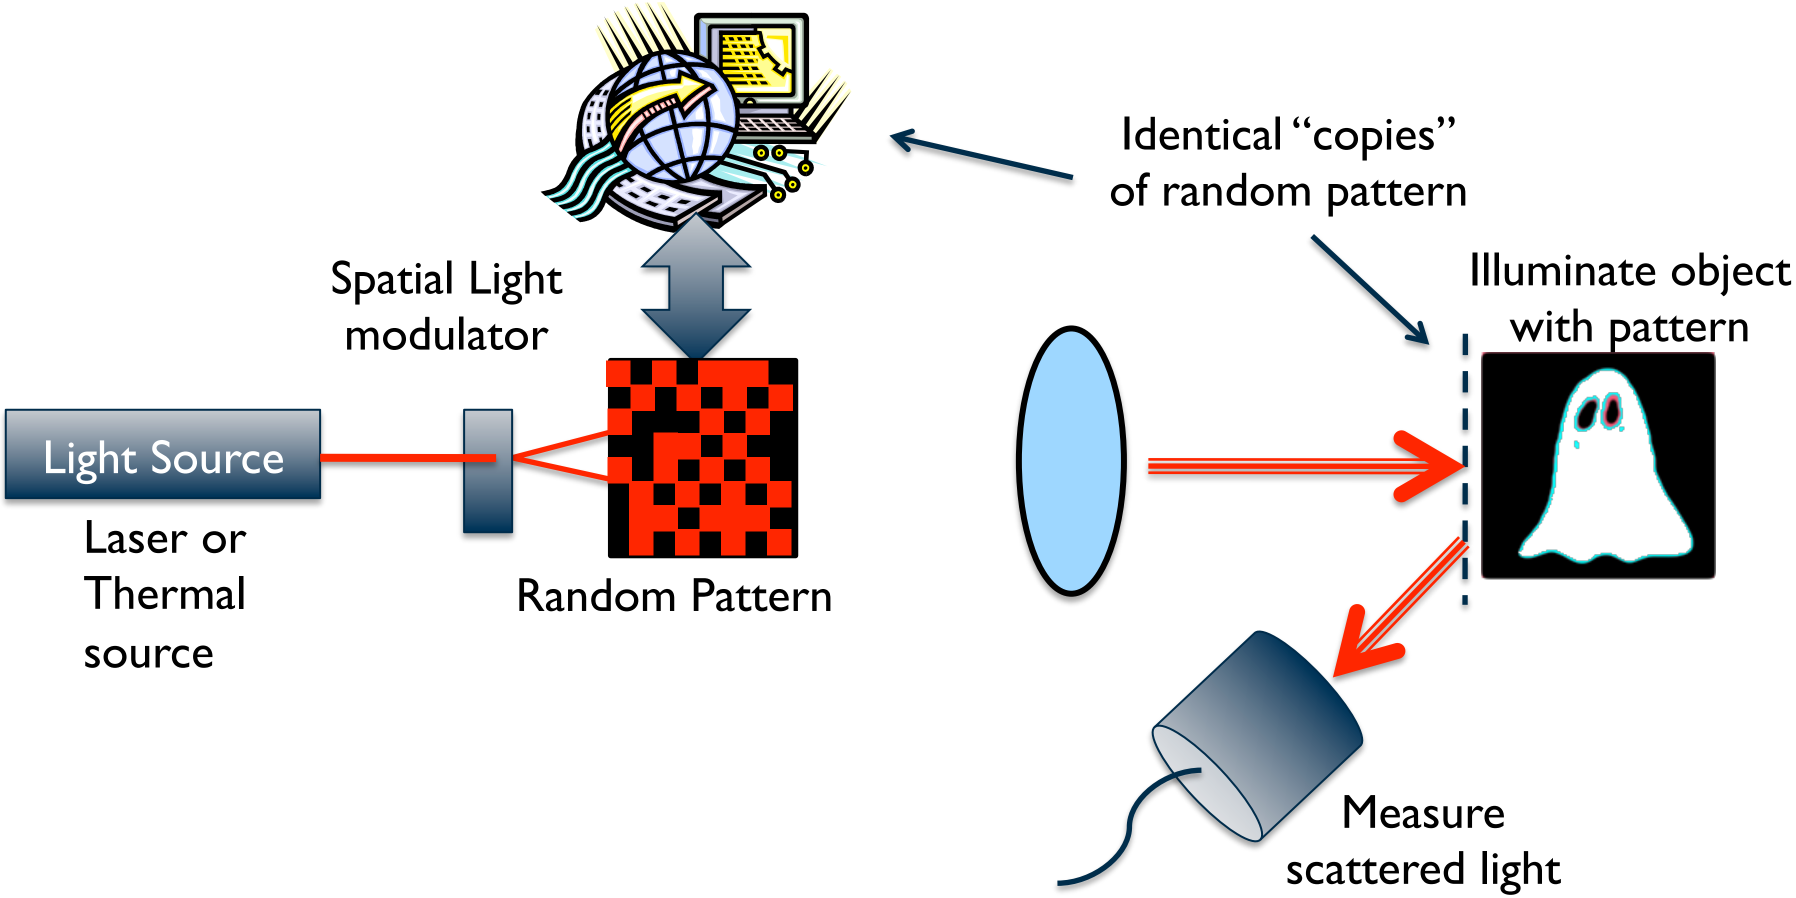

A single-pixel Imager

- Filter the scattered light – use all the same algorithms as a computational ghost imager

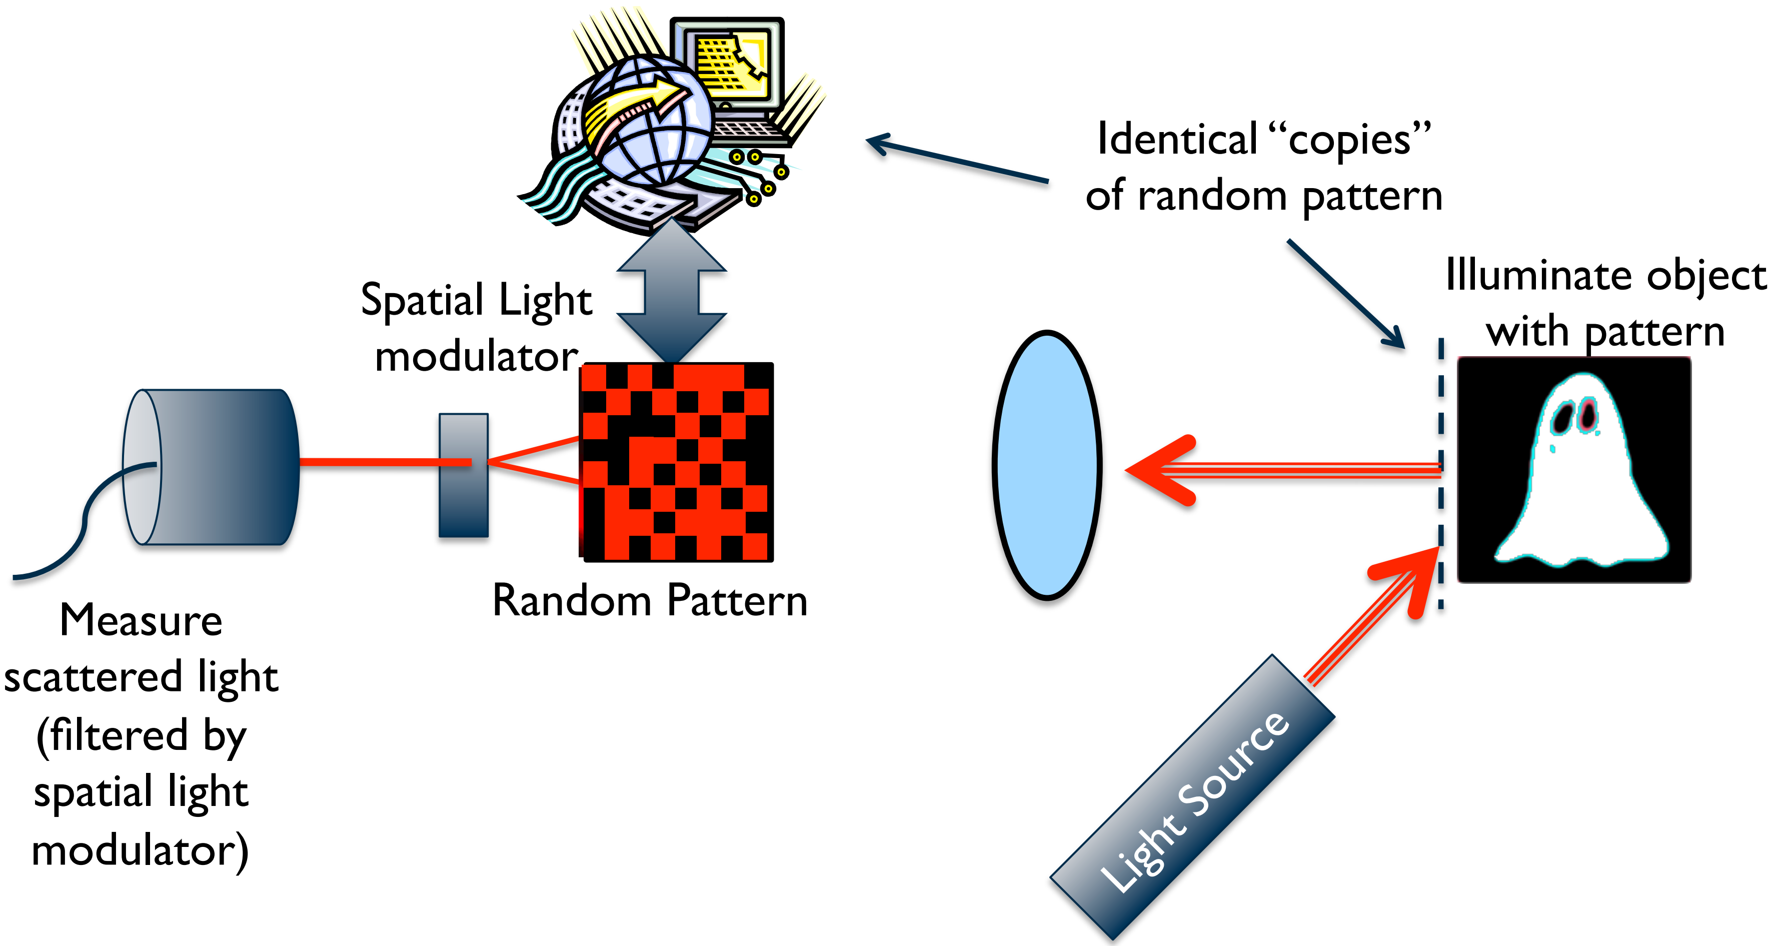

Tips of the trade

Don’t use random patterns use Hadamard patterns

Hadmards are orthogonal to each other (unlike random)

Many Hadamards are redundant within any real image (unlike random)

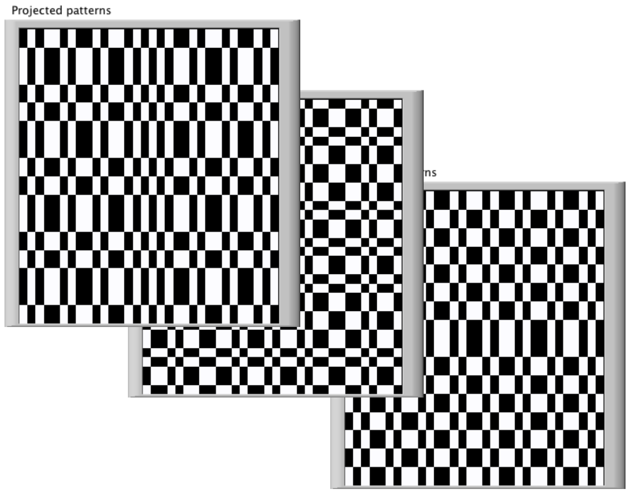

Display every pattern and a +ve and –ve pair (common mode rejection)

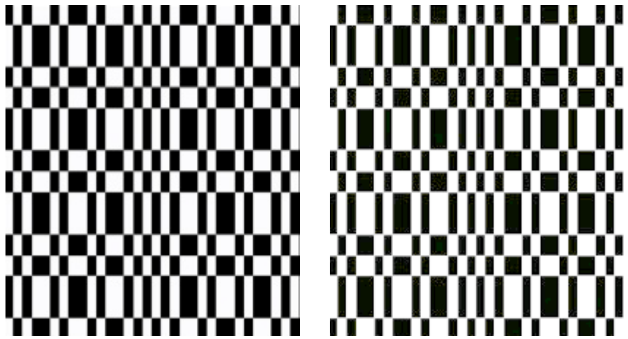

Single (RGB) pixel video

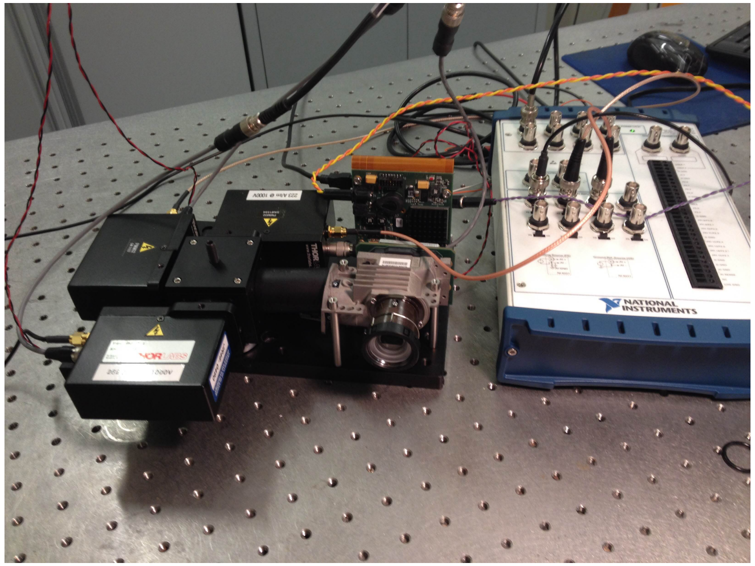

Single (RGB) pixel video

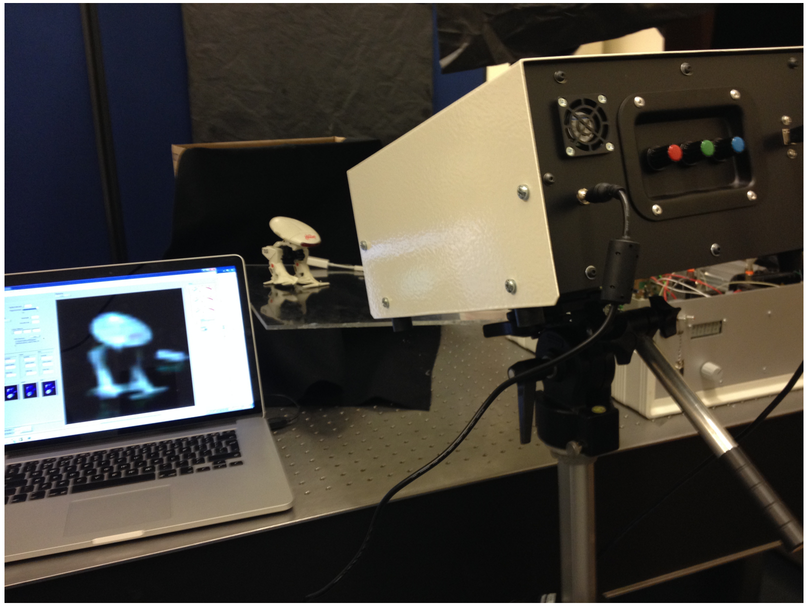

What next?

Imaging at non visible and/or many wavelengths

Imaging of Gas Emissions

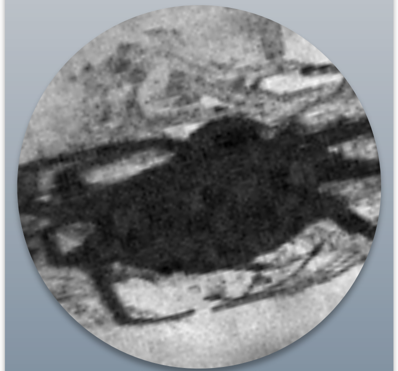

Single Pixel

Microscopy/ video

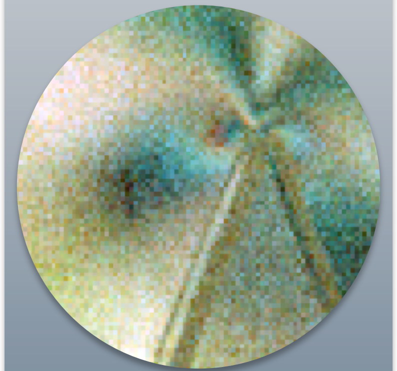

Imaging of polarisation anomalies

Single pixel (Ghost) Imaging

Visible vs. SWIR video

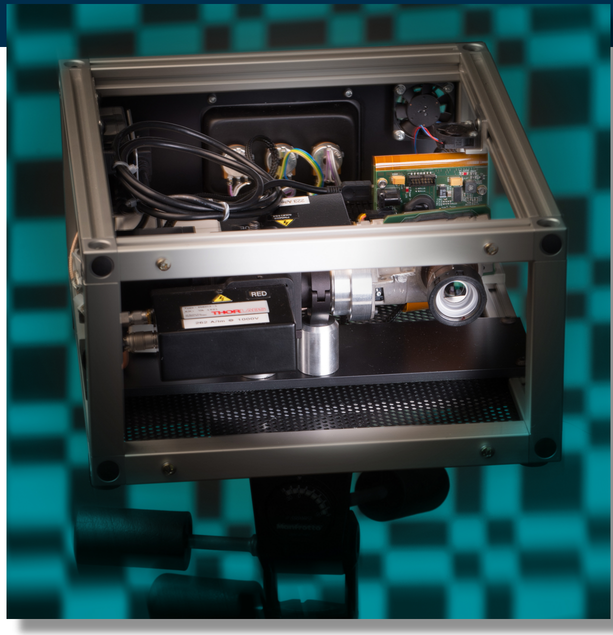

University of Glasgow

Thank you

http://www.gla.ac.uk/schools/physics/research/groups/optics/

3D Computational Imaging with Single-Pixel Detectors

B. Sun,[1]* M. P. Edgar,[1] R. Bowman,[1,2] L. E. Vittert,[3] S. Welsh,[1] A. Bowman,[3] M. J. Padgett[1]

17 MAY 2013 VOL 340 SCIENCE www.sciencemag.org

Fast full-color computational imaging with single-pixel detectors

Stephen S. Welsh,[1]* Matthew P. Edgar,[1] Richard Bowman,[2]
Phillip Jonathan,[3] Baoqing Sun,[1] and Miles J. Padgett[1]

7 October 2013 | Vol. 21, No. 20 | DOI:10.1364/OE.21.023068 | OPTICS EXPRESS 23068