VRIJE
UNIVERSITEIT
BRUSSEL

Master thesis submitted in order to be awarded the degree of
Master of Science in Engineering Technology:
Electromechanical Engineering - mechatronics

BREWSTER ANGLE PTYCHOGRAPHIC MICROSCOPE
3D Topographic Metrology for µm-resolution inspection

Âlim Liman

2025-2026

promoter: Prof. Dr. Ir. Dries Peumans
supervisor: Dr. Hyun-su Kim

ENGINEERING SCIENCES

Project plan

My project plan and logbook can be consulted online via this link.

Acknowledgements

I would like to express my sincere gratitude to Dr. Hyun-su Kim for the opportunity, the trust he placed in me, and his guidance and support throughout this project. His expertise and patience in the lab made this work possible. I am equally grateful to Prof. Dr. ir. Dries Peumans for the opportunity to carry out part of this research at imec and for his continued interest in the progress of this thesis. I also thank Jiwook for his help during the manufacturing of the diffraction grating.

Figure 1: The imec Tower, Leuven.

Âlim Liman

Brussels, June 7, 2026

Declaration

During the preparation of this work, the author(s) used the following tool(s) and service(s):

name of tool/service		reason
Google gemini 3.1 Pro	in order to:	*Optimize code, spellchecker, explaining literature.*
Chat GPT-5.2	in order to:	*Optimizing capture sequence algorithm.*
Claude Opus 4.6	in order to:	*Optimizing code for post capture processing.*
Claude Sonnet 4.6	in order to:	*Correct grammatical mistakes.*

After using this tool(s) and/or service(s), the author(s) reviewed and edited the content as needed. The author(s) take(s) full responsibility for the content of this work.

Contents

Project plan i

Acknowledgements ii

Disclosure on the use of AI iii

Contents iv

Nomenclature vii

1 Introduction 1
- 1.1 Context . . . 1
- 1.2 Topic and scope . . . 2
- 1.3 Research question . . . 2
- 1.4 Objectives . . . 2
- 1.5 Thesis structure . . . 2

2 Ptychography 4
- 2.1 Diffraction Pattern . . . 6
- 2.2 The Ptychographic Iterative Engine (PIE) . . . 8
- 2.3 Overlap Parameter and Convergence . . . 9
- 2.4 Numerical Aperture . . . 10
- 2.5 Transmission Ptychography . . . 11
- 2.6 Angular Spectrum Method . . . 13
  - 2.6.1 GPU implementation . . . 14

3 Image Acquisition 16
- 3.1 Optical Layout . . . 16
- 3.2 Camera and Detector System . . . 17
  - 3.2.1 Nyquist criterion . . . 17
  - 3.2.2 Alignment procedure . . . 18
- 3.3 Scan Strategy . . . 18
  - 3.3.1 Grid Pattern . . . 18
  - 3.3.2 Fermat spiral . . . 18
  - 3.3.3 Typical Scan Parameters . . . 19
- 3.4 High Dynamic Range Imaging . . . 20
  - 3.4.1 Multi-Exposure Acquisition . . . 20
  - 3.4.2 Saturation-Cascade Merge . . . 20
  - 3.4.3 Application to a Graphene Flake . . . 21
- 3.5 Contrast . . . 23

4 Brewster Angle 24
- 4.1 Polarisation of Light . . . 24

4.2 Fresnel Equations . . . 25
4.3 The Brewster Angle . . . 25
4.4 Limitations of Conventional BAM . . . 26
4.4.1 Keystone Distortion . . . 26
4.4.2 Phase Insensitivity . . . 26
4.5 Experimental result of the polarization angle . . . 26

5 Conical Distortion Correction 31
5.1 Geometry . . . 31
5.2 Implementation . . . 31
5.2.1 Numerical example . . . 32
5.3 Validation . . . 33

6 Reconstruction pipeline 35
6.1 Reconstruction Engine . . . 35
6.2 Reconstruction Procedure . . . 35
6.2.1 Initialisation . . . 35
6.2.2 ePIE Iteration . . . 35
6.2.3 Convergence Criterion . . . 36
6.3 Resolution . . . 37
6.3.1 Measurement procedure . . . 37
6.3.2 Normalised Image Log-Slope . . . 38

7 Topographic Surface Reconstruction 41
7.1 Phase-to-Height Conversion . . . 41
7.1.1 Reflection Geometry . . . 41
7.1.2 Transmission Geometry . . . 42
7.2 Phase Processing . . . 42
7.2.1 Phase Unwrapping . . . 43
7.2.2 Region of Interest . . . 43
7.3 Results: Reflection Geometry . . . 43
7.3.1 Siemens Star Target . . . 43
7.3.2 Logic Pattern Target . . . 43

8 Ptychographic Imaging of Graphene 46

9 Diffraction Grating Based Distortion Correction for BAPM 49
9.1 Grating Diffraction Theory . . . 49
9.2 Keystone Distortion in Brewster Angle Microscopy . . . 49
9.3 Grating-Based Distortion Correction . . . 49
9.3.1 Operating Principle . . . 49
9.3.2 DC-BAPM Optical Setup . . . 49
9.4 Material-Specific Grating Design for 632.8 nm Illumination . . . 50
9.4.1 Refractive Indices at 632.8 nm . . . 50
9.4.2 The DC-BAM Design Condition . . . 51
9.4.3 Calculated Grating Specifications . . . 52
9.4.4 Binary Grating Fabrication Design . . . 53
9.4.4.1 Duty Cycle and Depth Optimisation . . . 53
9.4.4.2 Cross-Section and Layout . . . 54
9.5 Grating Manufacturing . . . 54
9.5.1 Spin Coating . . . 55
9.5.2 Development . . . 56
9.5.3 Metrology . . . 57

9.6 Summary 58

10 Conclusions 59
10.1 Summary of Findings 59
10.1.1 Ptychographic Phase Retrieval at Brewster Incidence 59
10.1.2 Polarisation-Dependent Contrast in DC-BAPM Near the Brewster Condition 59
10.1.3 Diffraction Grating Design for Optical Keystone Correction at Brewster Incidence 59
10.2 Answer to the Research Question 59
10.3 Significance of the Work 60

11 Future work 61
11.1 Multi-Angle Phase Retrieval 61
11.2 Real-Time Reconstruction 61
11.3 Depolarisation Correction 61
11.4 Superresolution via Sub-Pixel Scanning 62
11.5 Structured Illumination via a Computational Diffuser 62

Bibliography 64

PIE Pipeline Flowchart 66

DC-BAPM Alignment Procedure 67

HDR Validation on a Siemens Star 68
Two-Exposure Bracket 68
Effect on the Reconstruction 68
NILS Resolution Analysis 69

HDR Merge Diagnostics and Contrast Metrics 73
Graphene Merge Diagnostics 73
Contrast Metrics 75

Brewster Reflectance for Additional Substrates 76

Reconstruction Docker Container 78

Computer Hardware Setup 80

Example Reconstruction Parameter File 82

Resolution Sample Design 84

Additional Graphene Reconstructions 86

Grating GDS Layout 88

Fabrication and Metrology Equipment 89

Acronyms

AFM	Atomic Force Microscope
AOI	Angle Of Incidence
ASM	Angular Spectrum Method
BAM	Brewster Angle Microscopy
CCD	Charge-Coupled Device
CDI	Coherent Diffractive Imaging
CNT	Carbon Nanotube
CUDA	Compute Unified Device Architecture
CVD	Chemical Vapour Deposition
DC-BAPM	Distortion-Corrected Brewster Angle Ptychographic Microscopy
DPPC	Dipalmitoylphosphatidylcholine
EBL	Electron-Beam Lithography
ePIE	Extended Ptychographic Iterative Engine
EUV	Extreme Ultraviolet
FFT	Fast Fourier Transform
FPM	Fourier Ptychographic Microscopy
GDS	Graphic Data System
GPU	Graphics Processing Unit
HDR	High Dynamic Range
HDF5	Hierarchical Data Format version 5
HeNe	Helium-Neon laser
IPA	Isopropyl Alcohol
LED	Light-Emitting Diode
MIBK	Methyl Isobutyl Ketone
NA	Numerical Aperture
PA	Polarization Angle
PBS	Polarising Beam Splitter
PIE	Ptychographic Iterative Engine
PMMA	Polymethyl Methacrylate
SEM	Scanning Electron Microscope
SWCNT	Single-Walled Carbon Nanotube
TE	Transverse Electric (s-polarisation)
TM	Transverse Magnetic (p-polarisation)
USAF	United States Air Force resolution test chart

Greek symbols

α	angle of incidence on grating	rad
α	regularisation parameter (PIE)	
β	diffracted angle from grating	rad

β	feedback parameter (ePIE)	
$\boldsymbol{\delta}$	displacement field	px
ϵ_{RMS}	RMS residual displacement	px
η	grating diffraction efficiency	
θ_B	Brewster angle	rad
θ_i	angle of incidence at interface	rad
θ_t	transmission angle at interface	rad
λ	wavelength	m
λ_b	blaze wavelength	m
Λ	grating groove period	m
ϕ	phase	rad
ϕ_b	blaze angle	rad
Ψ	exit wave	
$\tilde{\Psi}$	far-field complex amplitude	

Roman symbols

a	centre-to-centre distance between scan positions	m
A	aperture function (probe selection)	
c	keystone compression factor	
d	polynomial degree (distortion model)	
d_f	film thickness	m
d_{px}	detector pixel size	m
E_n	normalised error metric (ePIE)	
$\mathbf{E}$	electric field vector	V/m
F	flat-field response	
I	diffraction intensity	counts
j	scan position index	
K	number of probe modes	
m	diffraction order	
M	number of scan positions	
n	refractive index	
n	iteration number	
N	grating groove density	lines/m
N_f	number of calibration features	
N_{px}	number of detector pixels per side	
o	overlap fraction between adjacent probe positions	
O	complex object transmission function	
P	complex probe function	
$\mathbf{q}$	spatial frequency coordinate	1/m
q_{max}	maximum recorded spatial frequency	1/m
$\mathbf{r}$	position coordinate in object plane	m
r	probe radius	m
r_0	probe $1/e^2$ intensity radius	m
r_p	Fresnel reflection coefficient (p-polarisation)	
r_s	Fresnel reflection coefficient (s-polarisation)	
R_p	p-polarisation reflectance	
R_s	s-polarisation reflectance	
$\mathbf{R}_j$	j-th scan position vector	m

t_p	Fresnel transmission coefficient (p-polarisation)	
t_s	Fresnel transmission coefficient (s-polarisation)	
U	PIE object update function	
V	ePIE probe update function	
w_k	weight of k-th probe mode	
z	sample-to-detector distance	m
$\Delta x_\parallel$	scan step size along plane of incidence	m
$\Delta x_\perp$	scan step size transverse to plane of incidence	m

Subscripts

$\parallel$	along the plane of incidence
$\perp$	transverse to the plane of incidence
B	Brewster
b	blaze
corr	corrected
bg	background
f	film
j	scan position index
n	iteration number
p	p-polarisation
s	s-polarisation
1,2	medium 1, medium 2

Chapter 1

Introduction

1.1 Context

Semiconductor devices and optical components are built from thin layers of material (often just a few nanometers thick) deposited on flat silicon wafer substrates. To verify the correct patterning and structuring of these layers during fabrication, engineers need non-destructive inspection methods that involve no physical contact, staining, or labelling. As critical dimensions and layer thicknesses continue to shrink in state-of-the-art fabrication, these requirements become increasingly stringent. An ideal metrology tool would reveal not only where 2D patterns are located, but also how they are structured in three dimensions, as well as their optical properties across the surface. Current techniques either lack the resolution, require expensive equipment, or only provide intensity images that miss the phase information required to quantitatively extract film thickness.

Thin-film materials (like graphene monolayers, metal coatings, polymer films, holographic patterns, micro-LED structures, and transistor gates on silicon) are increasingly relevant in semiconductor technology. Imaging them requires surface sensitivity and phase sensitivity simultaneously. Conventional optical microscopy cannot resolve nanometer-scale films. X-ray and extreme ultraviolet (EUV) light offers sufficient resolution, but at these short wavelengths imaging optics become impractically expensive to fabricate.

Ptychography removes the imaging optics between sample and detector. The complex wavefield of the sample is reconstructed computationally from a series of diffraction patterns, so no lens must form an image of the sample on the sensor. Illumination optics that shape the probe before the sample are still used, but these are simple and cheap compared to a high-resolution imaging system. This thesis develops a visible-light ptychographic microscope as a stepping stone toward affordable X-ray ptychographic microscopy.

Brewster angle microscopy (BAM) is an optical technique for imaging thin films at an interface [Hénon and Meunier, 1991]. At the Brewster angle, p-polarised light reflects with near-zero intensity from a clean surface. A thin film on that surface shifts the local reflectivity, producing contrast, which reveals the lateral structure of the film. This makes BAM non-invasive and label-free, which is useful for materials like graphene, metal coatings, and polymer films on semiconductor substrates.

Conventional BAM has two problems. First, the oblique angle of incidence compresses the image along one axis (this is the keystone distortion) [Zhe Sun and Baldelli, 2017]. Second, BAM only records intensity, so the phase of the reflected wavefield is lost. The phase carries information about the film thickness and refractive index, which intensity alone cannot provide.

Ptychography can solve both problems. It records a series of diffraction patterns from overlapping illuminated regions of the sample, then uses a phase retrieval algorithm to reconstruct the complex reflectance of the sample and the probe simultaneously [J. M. Rodenburg and Faulkner, 2004]. The extended ptychographic iterative engine (ePIE) can do this without prior knowledge of the probe, and has been shown to recover phase information well beyond the

resolution of conventional optics [A. M. Maiden et al., 2011].

Combining BAM with ptychography produces a microscope that retains the surface sensitivity of BAM while adding quantitative phase images and higher spatial resolution.

1.2 Topic and scope

This thesis describes the design, construction, and reconstruction pipeline of a distortion-corrected ptychographic Brewster angle microscope (DC-BAPM) operating at 632.8 nm wavelength. The work covers the optical design, the image capture procedure, the iterative phase retrieval, the distortion correction algorithm and diffraction grating that corrects the keystone distortion. The instrument is characterised on calibration targets (mostly the Siemens star) and thin-film samples, and its resolution and phase sensitivity were assessed.

1.3 Research question

The central research question is as follows:

> Can ptychographic phase retrieval be applied to a Brewster angle microscope to reconstruct quantitative phase images of thin films, and can the keystone distortion be corrected without affecting the reconstruction quality?

The supporting questions are as follows:

- What are the practical constraints on scan step size, detector distance, and illumination coherence at oblique incidence?
- What phase sensitivity and spatial resolution can the system achieve?
- How should the diffraction grating be designed to correct the keystone distortion?

1.4 Objectives

- **Design** an optical setup that combines Brewster angle illumination with a diffraction grating for probe structuring and distortion correction.
- **Implement** an image capture procedure that records a full ptychographic dataset at the Brewster angle.
- **Develop** an ePIE reconstruction pipeline that accounts for the oblique geometry and recovers the complex reflectance of a sample.
- **Validate** the system on calibration targets and assess the spatial resolution and phase sensitivity.
- **Evaluate** the instrument for imaging thin films on PMMA, including resolution targets and graphene layers.

1.5 Thesis structure

Chapter 2 introduces the theory of ptychography and the ePIE phase-retrieval algorithm. Chapter 3 describes the experimental setup and data-acquisition procedure. Chapter 4 covers the physics of the Brewster angle and the polarisation sensitivity of the system. Chapter 5 presents

the conical distortion correction algorithm that compensates for the Ewald-sphere distortion in the recorded diffraction patterns. Chapter 6 details the full reconstruction pipeline. Chapter 7 describes the topographic surface reconstruction from the retrieved phase. Chapter 8 presents ptychographic imaging of graphene monolayers. Chapter 9 covers the diffraction grating design for optical keystone correction. Chapter 10 summarises the conclusions and Chapter 11 outlines directions for future work.

Ptychography

Ptychography is a scanning coherent diffractive imaging (CDI) technique that recovers both the amplitude and phase of a complex-valued object from a series of far-field diffraction intensity measurements [J. Rodenburg and A. Maiden, 2019]. Unlike conventional lens-based microscopy, ptychography requires no imaging optics between the sample and the detector, removing the resolution limit imposed by the numerical aperture of a lens. Instead, spatial resolution is determined by the highest spatial frequency recorded on the detector, and the phase retrieval problem is solved computationally using iterative algorithms.

To place ptychography in context, it is useful to compare it against three other imaging approaches, ordered by increasing computational involvement.

In a conventional optical microscope, light from the sample is collected by an objective lens and focused onto an intermediate image plane, as shown in Fig. 2.1. The image is formed entirely by the optics. No computation is involved, and the resolution is hard-limited by the numerical aperture of the objective.

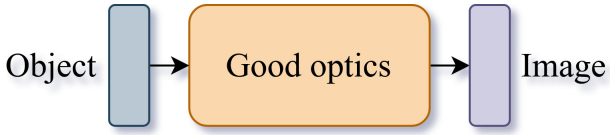

Figure 2.1: Conventional microscope pipeline. The image is formed directly by the objective and tube lens, resolution is fixed by the lens NA (inspired by J. Rodenburg and A. Maiden, 2019).

A more advanced variant adds a post-processing stage after image capture. The raw image contains residual lens aberrations, noise, and distortion. A correction algorithm is applied to the captured intensity to produce an improved image, as shown in Fig. 2.2. The optical hardware is unchanged, only the software pipeline is extended. Resolution is still limited by the lens NA, but image quality can be significantly improved.

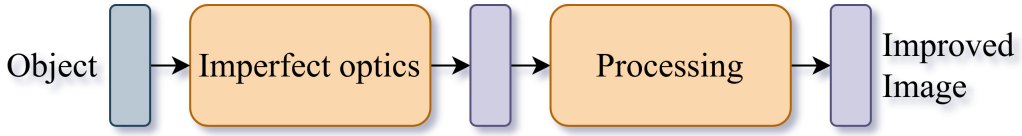

Figure 2.2: Post-processing microscope pipeline. A computational correction stage follows image capture to remove aberrations and noise (inspired by J. Rodenburg and A. Maiden, 2019).

Adaptive optics microscopes go one step further by correcting aberrations before detection rather than after. A wavefront sensor measures the aberration, and a deformable mirror or spatial light modulator compensates for it in a closed feedback loop, as shown in Fig. 2.3. The correction happens in real time during acquisition, so the detector receives a pre-corrected wavefront. This is particularly effective when aberrations vary with sample position, as in deep-tissue imaging.

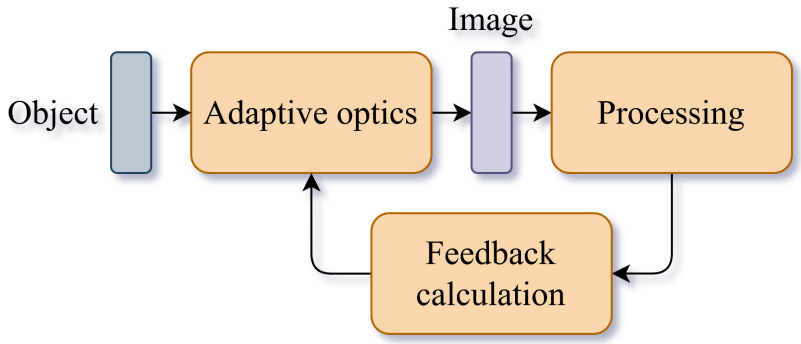

Figure 2.3: Adaptive optics microscope pipeline. A feedback loop measures and corrects wavefront aberrations in real time, before the image is recorded (inspired by J. Rodenburg and A. Maiden, 2019).

Ptychography takes a fundamentally different approach. There is no imaging lens between the sample and the detector. Instead, a coherent probe illuminates the sample at multiple overlapping positions, and the detector records the far-field diffraction intensity at each position. A phase retrieval algorithm then reconstructs the complex object (both amplitude and phase) from the full set of diffraction patterns, as shown in Fig. 2.4. The resolution is no longer set by any lens, it is determined by the highest spatial frequency captured on the detector.

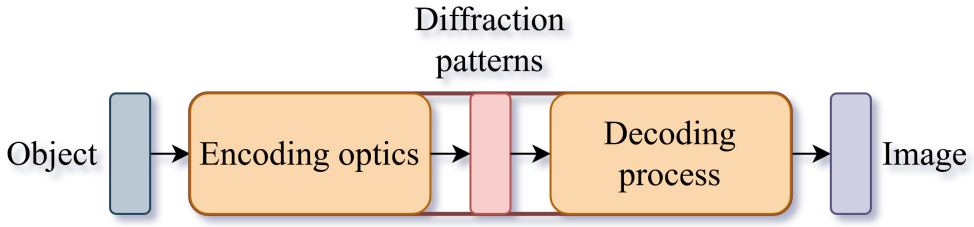

Figure 2.4: Ptychographic pipeline. No imaging lens is used. The detector records diffraction patterns, and a reconstruction algorithm recovers the complex object (inspired by J. Rodenburg and A. Maiden, 2019).

Ptychography offers three advantages over conventional microscopy.

- **No lens required.** No imaging lens is placed between the sample and the detector. The only lens in the system is the illumination lens, which shapes the probe beam. Its optical quality does not limit the final resolution, since resolution is set by the detector pixel size and the sample-to-detector distance. A diverging probe is preferred over a flat-phase beam to avoid a bright central spot, which is difficult for the algorithm to handle.
- **Complex image output.** The reconstruction yields both the amplitude and the phase of the object. Transparent samples such as biological films or photoresist absorb very little light, so their amplitude contrast is weak. The phase encodes the optical path length through the sample and provides strong contrast for these materials. Ptychography recovers phase directly from the diffraction data and produces high-quality phase images.
- **Self-calibration.** The reconstruction algorithm corrects for lens aberrations, positional errors, defocus, and partial coherence of the illumination as part of the iterative process. No separate calibration measurement is required. The algorithm can also extrapolate diffraction data beyond the directly measured bandwidth, recovering spatial frequencies not captured on the detector and improving the final resolution.

2.1 Diffraction Pattern

The bright central region of the diffraction pattern is called the brightfield. It contains low-frequency spatial information scattered from the sample. The dark surrounding region is the darkfield and carries high-frequency scattering information. In linear scale the darkfield appears uniformly black, but viewing it in log scale reveals that it contains significant data, as shown in Fig. 2.5(b).

The Fresnel rings visible in the pattern are formed by the laser beam diffracting through a 1 mm aperture. The sharp dark spots near the bottom of the brightfield are shadows cast by dust particles in the optical system. Because the dust spots appear consistently across all diffraction patterns, the ePIE algorithm can identify and correct for them during reconstruction.

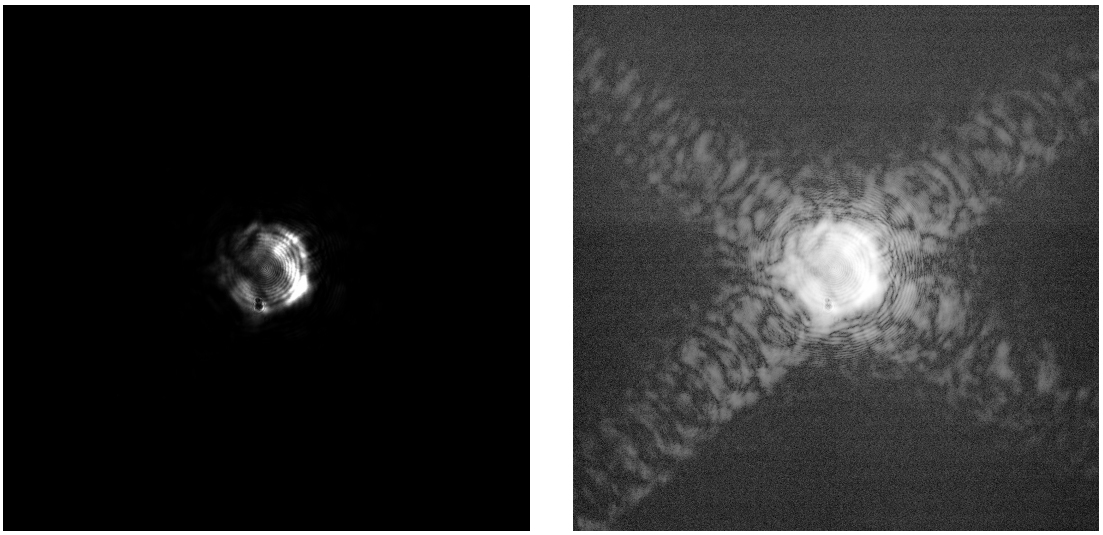

(a) Original diffraction image (b) Logarithmic scale

Figure 2.5: Example diffraction pattern recorded on 2026-04-30, shown as (a) the original diffraction image and (b) on a logarithmic intensity scale. The log-scale view reveals the darkfield structure that is invisible in the original diffraction image.

A simple Fourier transform demonstration illustrates how the low and high spatial frequencies in the diffraction pattern map onto features in the reconstructed image. The test object is an image of bacteria, shown in Fig. 2.6.

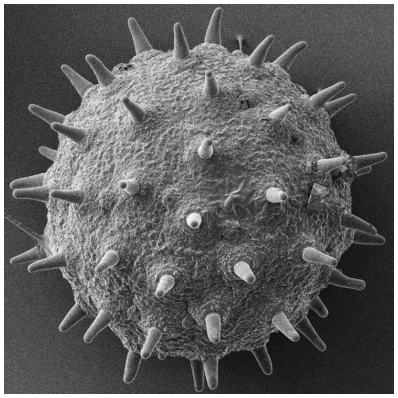

Figure 2.6: Example test image.

Applying a Fast Fourier Transform (FFT) to the image yields its amplitude and phase in the spatial frequency domain, shown in Fig. 2.7. Low frequencies sit near the centre of the spectrum, high frequencies lie at the edges.

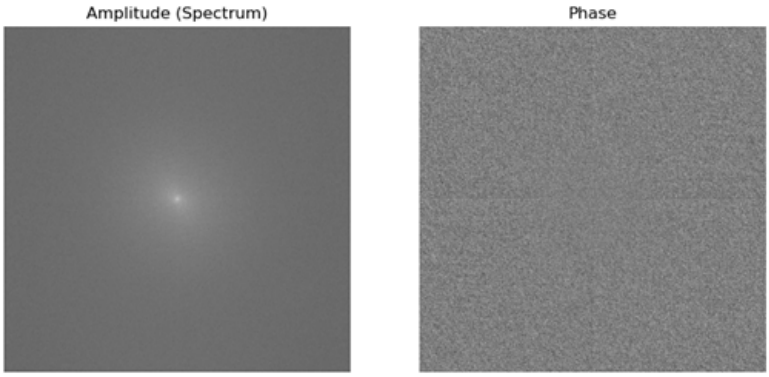

Figure 2.7: Fourier transform of the example test image.

An outer mask retains only the spatial information inside the white box, keeping the low frequencies and discarding the high ones. The inverse FFT (IFFT) of the masked spectrum is shown in Fig. 2.8. The general outline and large features of the bacteria are clearly visible, but the fine detail is lost.

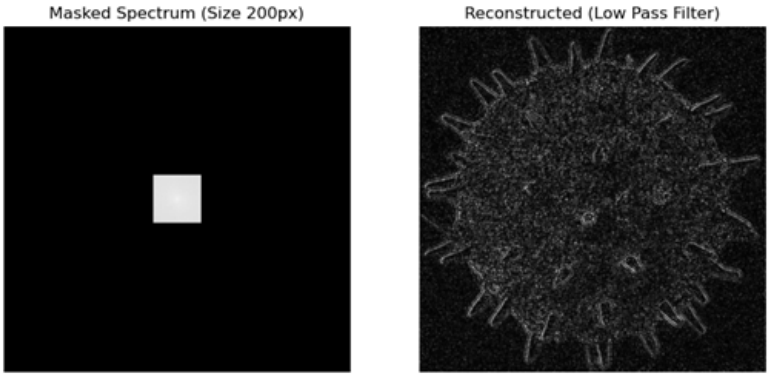

Figure 2.8: Inverse Fourier transform of outer masked amplitude and phase images.

An inner mask does the opposite: it blocks the low frequencies and keeps only the high ones. The IFFT in Fig. 2.9 resolves the fine detail of the bacteria, while the large-scale features become much harder to see.

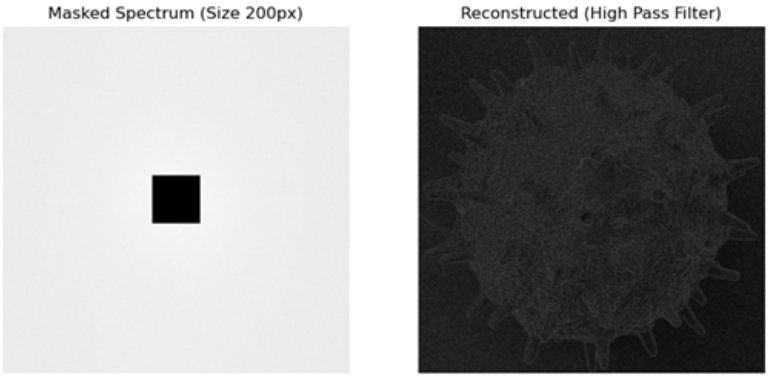

Figure 2.9: Inverse Fourier transform of inner masked amplitude and phase images.

2.2 The Ptychographic Iterative Engine (PIE)

Rodenburg and Faulkner introduced the ptychographical iterative engine (PIE) as a method to solve the phase problem using redundant information from overlapping illuminations at multiple positions [J. M. Rodenburg and Faulkner, 2004; J. Rodenburg and A. Maiden, 2019]. The sample is illuminated at positions $\mathbf{R}_j$, $j = 1, \ldots, M$, and a diffraction pattern $I_j(\mathbf{q})$ is recorded at each position. The overlap between adjacent illuminations provides the redundancy necessary for a unique, rapidly converging solution.

One iteration of the PIE algorithm at position j proceeds as follows:

1. Compute the estimated exit wave at scan position $\mathbf{R}_j$:

$$\Psi_j(\mathbf{r}) = P(\mathbf{r})\, O_n(\mathbf{r} + \mathbf{R}_j), \tag{2.1}$$

   where $P(\mathbf{r})$ is the complex-valued probe, $O_n(\mathbf{r})$ is the object estimate at iteration n, and $\mathbf{r}$ is a real-space coordinate at the sample plane. The shift by $\mathbf{R}_j$ selects the region of the object currently illuminated by the probe.

2. Propagate to the detector plane with a Fourier transform:

$$\tilde{\Psi}_j(\mathbf{q}) = \mathcal{F}\{\Psi_j(\mathbf{r})\}, \tag{2.2}$$

   where $\mathcal{F}$ is the 2D Fourier transform and $\mathbf{q}$ is the spatial-frequency coordinate at the detector. The tilde marks fields in reciprocal space.

3. Replace the estimated modulus with the measured value, keeping the estimated phase:

$$\tilde{\Psi}'_j(\mathbf{q}) = \sqrt{I_j(\mathbf{q})}\, \frac{\tilde{\Psi}_j(\mathbf{q})}{|\tilde{\Psi}_j(\mathbf{q})|}, \tag{2.3}$$

   where $I_j(\mathbf{q})$ is the diffraction intensity recorded by the detector at scan position j. This is the only step in which experimental data enters the iteration, the prime marks the modulus-constrained field.

4. Back-propagate to the object plane:

$$\Psi'_j(\mathbf{r}) = \mathcal{F}^{-1}\{\tilde{\Psi}'_j(\mathbf{q})\}, \tag{2.4}$$

   which yields the modulus-corrected exit wave $\Psi'_j(\mathbf{r})$ in real space.

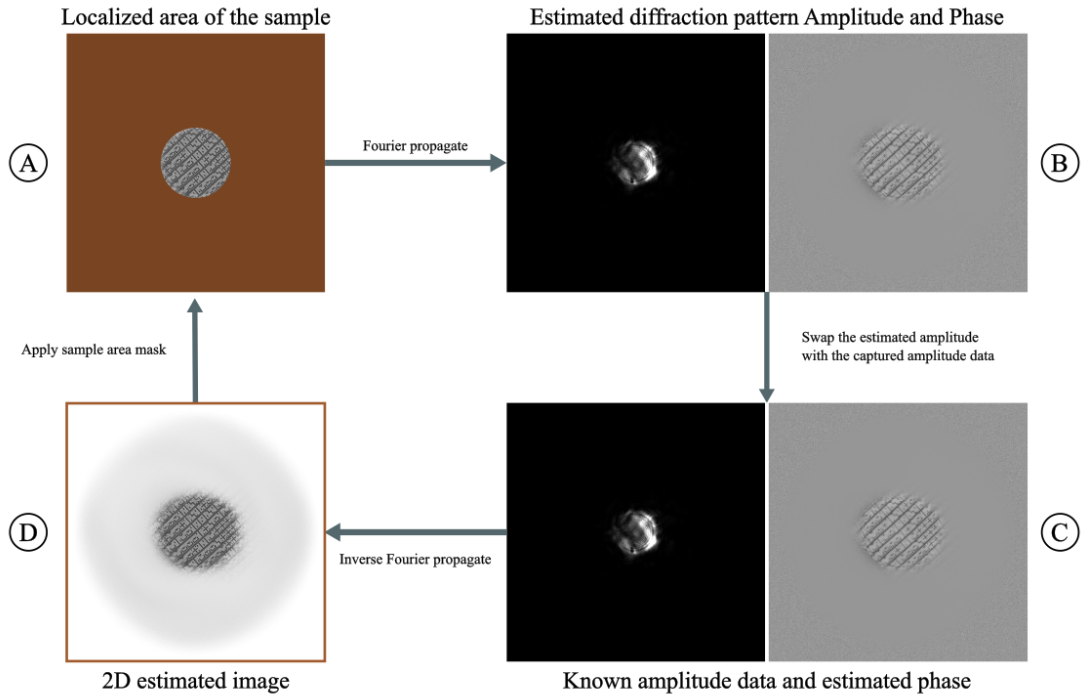

Figure 2.10: One iteration of the phase retrieval algorithm. (A) The current object estimate is Fourier-propagated to the detector plane, yielding the estimated diffraction amplitude and phase (B). The estimated amplitude is replaced by the measured amplitude while the estimated phase is retained (C). An inverse Fourier transform returns a new object estimate (D), to which a support mask is applied before the next iteration (inspired by J. Rodenburg and A. Maiden, 2019).

5. Update the object estimate using the difference between the constrained and current exit waves:

$$O_{n+1}(\mathbf{r}+\mathbf{R}_j) = O_n(\mathbf{r}+\mathbf{R}_j) + \beta\, U(\mathbf{r})\left[\Psi'_j(\mathbf{r}) - \Psi_j(\mathbf{r})\right], \tag{2.5}$$

where $\beta \in [0.9, 1.0]$ is the feedback parameter that controls the update size and $U(\mathbf{r})$ is the update function

$$U(\mathbf{r}) = \frac{|P(\mathbf{r})|}{\max|P(\mathbf{r})|} \cdot \frac{P^*(\mathbf{r})}{|P(\mathbf{r})|^2 + \alpha}. \tag{2.6}$$

$P^*(\mathbf{r})$ is the complex conjugate of the probe, the leading $|P(\mathbf{r})|/\max|P(\mathbf{r})|$ factor deweights regions of low probe intensity where the data carry little information about the object, and $\alpha \ll 1$ is a small constant that prevents division by zero in dark regions of the probe [Bunk et al., 2008]. The full flowchart that ties Eqs. (2.1)–(2.5) to their respective physical planes is given in Appendix 11.5, Fig. 1.

2.3 Overlap Parameter and Convergence

The overlap between adjacent illuminated regions is the single most important experimental parameter governing convergence speed and reconstruction quality. The relative overlap o is defined as

$$o = 1 - \frac{a}{2r}, \tag{2.7}$$

where r is the probe radius and a is the centre-to-centre distance between adjacent scan positions [Bunk et al., 2008]. Experimentally, the overlap is best kept between 80% and 85% to guarantee convergence while keeping the capture and reconstruction time low [Bunk et al., 2008; J. Rodenburg and A. Maiden, 2019].

Fig. 2.11 shows a Fermat-spiral scan pattern with a circular probe footprint of radius r. Each circle has diameter $2r$, and the centre-to-centre distance between any two neighbouring circles is the spacing a in Eq. (2.7). Reading the plot through the formula: when $a/2r$ is small, neighbouring circles share most of their area and $o \to 1$, when $a/2r$ is close to one, the circles barely touch and $o \to 0$. For the pattern shown, $r = 120\,\mu\text{m}$ (half of the $240\,\mu\text{m} \times 240\,\mu\text{m}$ probe footprint) and the Fermat spiral places the points so that the mean nearest-neighbour distance is $a \approx 50\,\mu\text{m}$, giving $o = 1 - a/2r \approx 79.2\%$. The dense web of intersections in the figure is the visual signature of that high overlap, and it is the redundancy that the ePIE algorithm exploits to recover both the object and the probe.

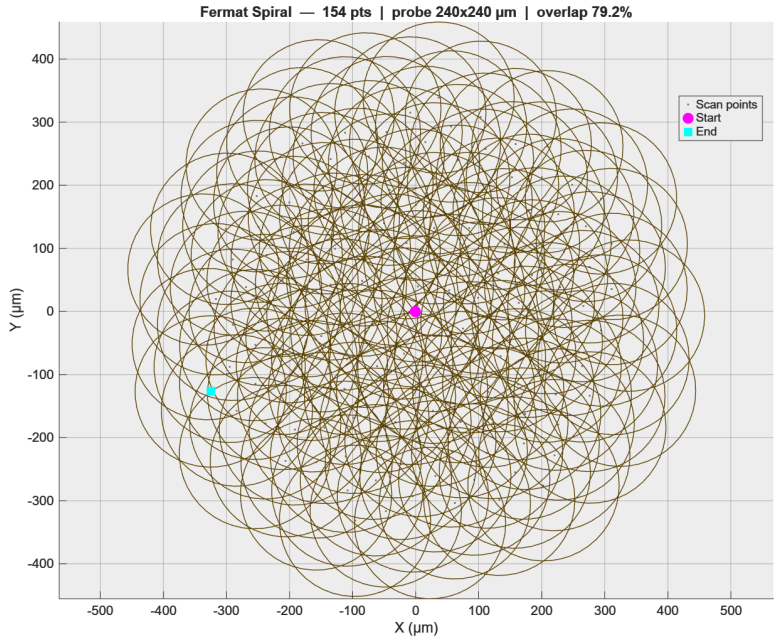

Figure 2.11: Fermat-spiral scan pattern of 154 positions with a circular $240\,\mu\text{m} \times 240\,\mu\text{m}$ probe, giving an average overlap of 79.2%. Each circle is one probe footprint, the magenta and cyan markers indicate the first and last scan positions.

2.4 Numerical Aperture

The numerical aperture (NA) of an optical system is defined as

$$\text{NA} = n \sin \theta, \tag{2.8}$$

where n is the refractive index of the medium and θ is the half-angle of the cone of light accepted or emitted by the system. The Abbe resolution limit gives the smallest resolvable feature size as

$$d = \frac{\lambda}{2\,\text{NA}}. \tag{2.9}$$

In a conventional lens-based microscope the NA is fixed by the objective lens. In ptychography there is no imaging lens, so two distinct NAs govern the resolution.

The *illumination NA* is set by the cone angle of the probe beam at the sample:

$$\mathrm{NA}_{\mathrm{ill}} = n \sin \theta_{\mathrm{ill}}, \tag{2.10}$$

and determines the size of the illuminated spot. The corresponding Abbe resolution limit for the illumination is

$$d_{\mathrm{ill}} = \frac{\lambda}{2\,\mathrm{NA}_{\mathrm{ill}}}. \tag{2.11}$$

The *detector NA* is set by the maximum scattering angle recorded on the detector. For a detector of half-width $D/2$ placed at a distance z from the sample,

$$\mathrm{NA}_{\mathrm{det}} = n \sin\left(\arctan \frac{D/2}{z}\right) \approx n \frac{D/2}{z}, \tag{2.12}$$

and the corresponding Abbe resolution limit is

$$d_{\mathrm{det}} = \frac{\lambda}{2\,\mathrm{NA}_{\mathrm{det}}}. \tag{2.13}$$

Because the phase retrieval algorithm can exploit all spatial frequencies up to $\mathrm{NA}_{\mathrm{det}}$, the resolution of a ptychographic reconstruction is ultimately limited by d_{det}, not by the probe size. This decoupling of probe size and resolution is one of the principal advantages of ptychography over conventional microscopy.

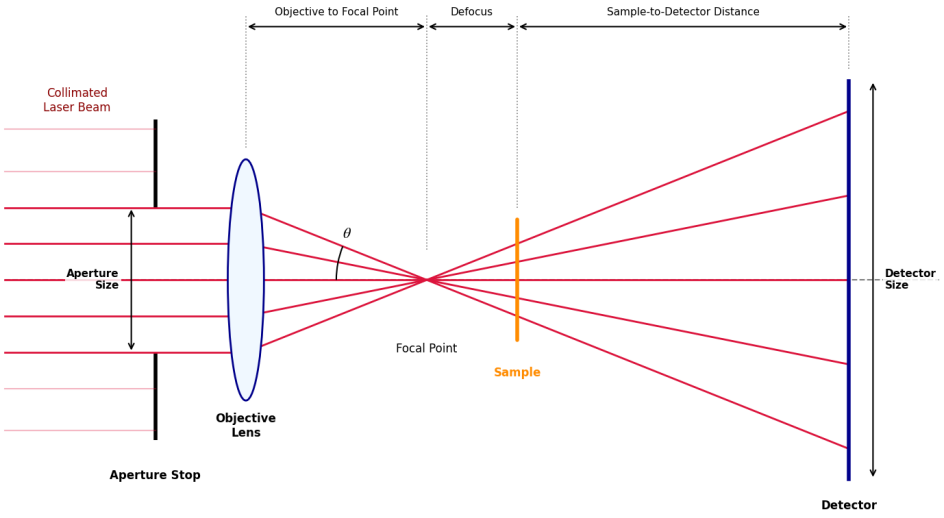

Figure 2.12: Ptychographic illumination geometry. An aperture stop and objective lens focus the collimated laser into a probe cone of half-angle θ, which sets $\mathrm{NA}_{\mathrm{ill}}$. The sample is placed slightly beyond the focal point (defocus). The beam then diverges and expands onto the detector, the detector size and sample-to-detector distance together set $\mathrm{NA}_{\mathrm{det}}$.

2.5 Transmission Ptychography

To make the abstract framework of the preceding sections concrete, I performed a transmission ptychography experiment during my internship using a Siemens star test target. In this configuration the beam passes straight through the sample on-axis, so the object function $O(\mathbf{r})$ encodes both the amplitude transmittance and the phase delay introduced by the sample.

The sample was scanned across a 15×15 grid of positions, giving 225 scan points in total. The scan pattern is shown in Fig. 2.13. At each position a far-field diffraction pattern was recorded

by the camera, and the full set of 225 patterns was fed into the ePIE algorithm described in Section 2.2.

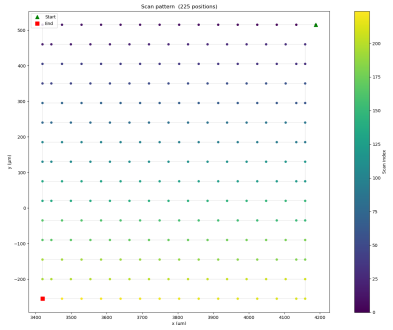

Figure 2.13: Raster scan pattern used for the transmission ptychography experiment. The 225 positions form a regular 15×15 grid, colour encodes the scan index from start (green triangle) to end (red square).

After 2000 iterations the algorithm converged to a stable reconstruction, which took around 6 minutes on an NVIDIA RTX 5080 GPU. The result is shown in Fig. 2.14. The amplitude image reveals the absorbing spokes and centre disc of the Siemens star, while the phase image shows the optical path-length variation between the glass substrate and the PMMA spokes that make up the sample. The simultaneously retrieved probe has a compact Gaussian amplitude with a nearly flat phase, consistent with a well-collimated, spatially coherent illumination beam.

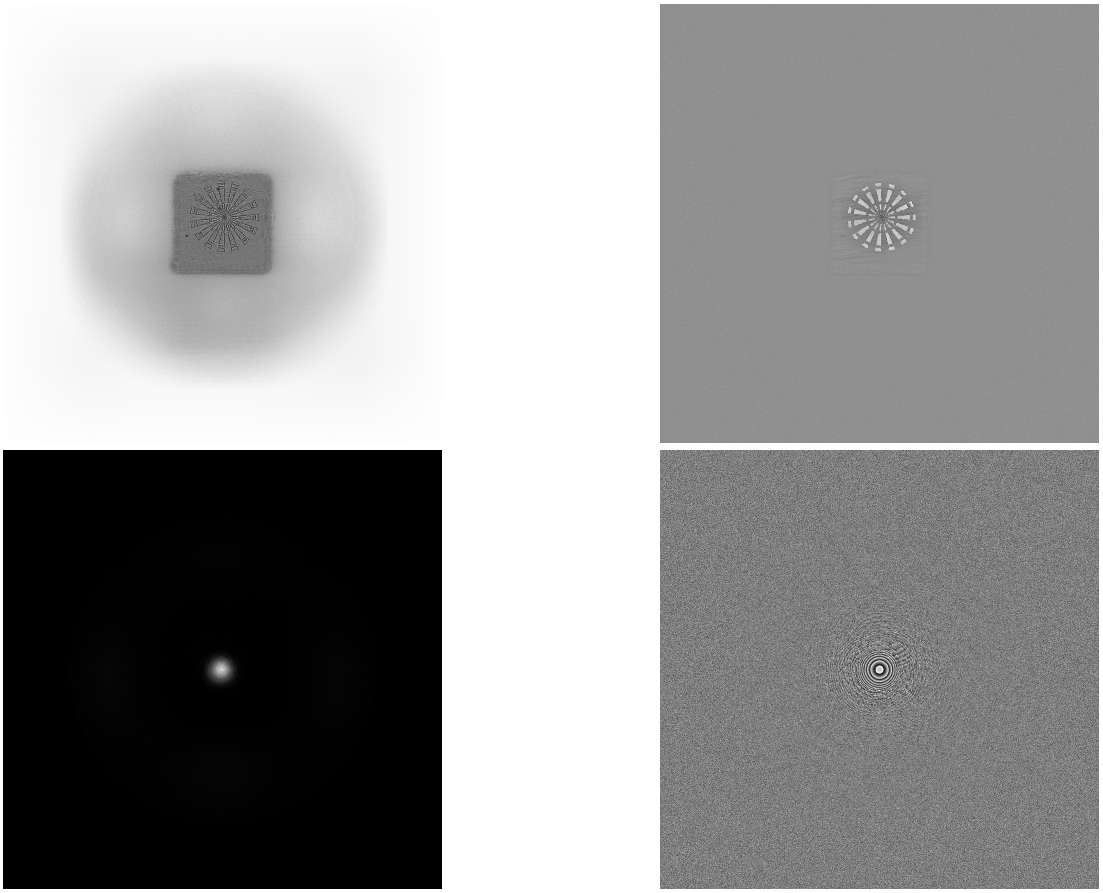

Figure 2.14: ePIE reconstruction after 2000 iterations. Top row: object amplitude $|O(\mathbf{r})|$ (left) and phase $\arg O(\mathbf{r})$ (right) of a Siemens star test target. Bottom row: retrieved probe amplitude $|P(\mathbf{r})|$ (left) and phase $\arg P(\mathbf{r})$ (right). The probe is approximately Gaussian with a flat phase, confirming coherent on-axis illumination.

2.6 Angular Spectrum Method

The ePIE propagation steps in Section 2.2 treat the sample-to-detector path as a single Fourier transform under the far-field assumption. Several parts of the broader pipeline (probe propagation for the calibration scan, defocus correction, partial-coherence simulation) instead need near-field propagation, where the far-field approximation no longer holds. The Angular Spectrum Method (ASM) provides an exact scalar-wave solution for free-space propagation between two parallel planes, with no far-field or paraxial approximation [Goodman, 2017]. A clear visual derivation of the plane-wave decomposition and the evanescent branch of the propagator is given by Konijnenberg [Konijnenberg, 2026a].

The field at the input plane is decomposed into its plane-wave components by a Fourier transform,

$$A(f_x, f_y; 0) = \iint U(x, y; 0)\, e^{-i2\pi(f_x x + f_y y)}\, dx\, dy, \tag{2.14}$$

where $A(f_x, f_y; 0)$ is the angular spectrum at $z = 0$. Each plane wave with transverse spatial frequencies (f_x, f_y) travels with a longitudinal wavenumber $k_z = k\sqrt{1 - \lambda^2(f_x^2 + f_y^2)}$, with

$k = 2\pi/\lambda$. Over a distance z each component acquires a transfer function

$$H(f_x, f_y; z) = \begin{cases} \exp\!\big(ikz\sqrt{1 - \lambda^2(f_x^2 + f_y^2)}\big), & \lambda^2(f_x^2 + f_y^2) \le 1 \\ \exp\!\big(-kz\sqrt{\lambda^2(f_x^2 + f_y^2) - 1}\big), & \lambda^2(f_x^2 + f_y^2) > 1 \end{cases} \tag{2.15}$$

so that the propagated field is

$$U(x, y; z) = \mathcal{F}^{-1}\big\{A(f_x, f_y; 0) \cdot H(f_x, f_y; z)\big\}. \tag{2.16}$$

The two branches of H correspond to propagating waves (pure phase rotation) and evanescent waves (real exponential decay). Evanescent components vanish within a few wavelengths and are usually irrelevant beyond the near field, but keeping them in H avoids any small-angle approximation in the result.

2.6.1 GPU implementation

The numerical recipe that turns Eq. (2.16) into working code — sampling, FFT layout, and bandlimit choice — is laid out in [Konijnenberg, 2026b]. The ASM is implemented in CUDA C++ on the GPU. The forward and inverse Fourier transforms are dispatched to NVIDIA's cuFFT library, and the propagator H is applied between them by a custom kernel that runs one thread per pixel. The kernel computes (f_x, f_y) from the thread indices using the FFT-shifted layout, evaluates $1 - \lambda^2(f_x^2 + f_y^2)$, and multiplies the transformed field by either the propagating or evanescent form of H in place. The relevant section of the kernel is shown in Listing 2.1.

Listing 2.1: Core CUDA kernel that applies the angular-spectrum propagator $H(f_x, f_y; z)$ to the Fourier transform of the field. One thread per detector pixel, FFT-shifted spatial-frequency layout.

```
__global__ void asm_propagate_kernel(cuFloatComplex* field_transformed,
                                     int NX, int NY,
                                     float LX, float LY,
                                     float lambda, float z)
{
    int idx = blockIdx.x * blockDim.x + threadIdx.x;
    int idy = blockIdx.y * blockDim.y + threadIdx.y;
    if (idx >= NX || idy >= NY) return;

    // Spatial frequencies (FFT-shifted layout)
    float fx = (idx < NX/2) ? float(idx)/LX : float(idx - NX)/LX;
    float fy = (idy < NY/2) ? float(idy)/LY : float(idy - NY)/LY;

    float k = 2.0f * (float)M_PI / lambda;
    float lambda_f_sq = lambda * lambda * (fx*fx + fy*fy);

    cuFloatComplex H;
    float arg_sqrt = 1.0f - lambda_f_sq;
    if (arg_sqrt >= 0.0f) {
        float phase = z * k * sqrtf(arg_sqrt);          // propagating
        H = make_cuFloatComplex(cosf(phase), sinf(phase));
    } else {
        float decay = -z * k * sqrtf(-arg_sqrt);        // evanescent
        H = make_cuFloatComplex(expf(decay), 0.0f);
    }

    int global_idx = idy * NX + idx;
    field_transformed[global_idx] =
        cuCmulf(field_transformed[global_idx], H);
}
```

The host program initialises a 2048×2048 complex field on a 4.54 µm-pitch grid (matching the GS3 detector), populates an 8-sided polygonal aperture of radius 1 mm as the input pupil, runs the kernel chain (FFT, propagator, inverse FFT, normalisation), and writes the propagated amplitude and phase to disk as raw float32 arrays for post-processing in Python. A single propagation step from the aperture to the detector at $z = 100$ mm completes in under 50 ms on the same RTX 5080, which makes the ASM cheap enough to embed in the iterative reconstruction loop whenever near-field defocus correction is needed.

Figure 2.15 shows the result of one such propagation. The simulated output matches the beam profile observed on the detector in the real setup, confirming that the ASM model is an accurate forward model for the probe.

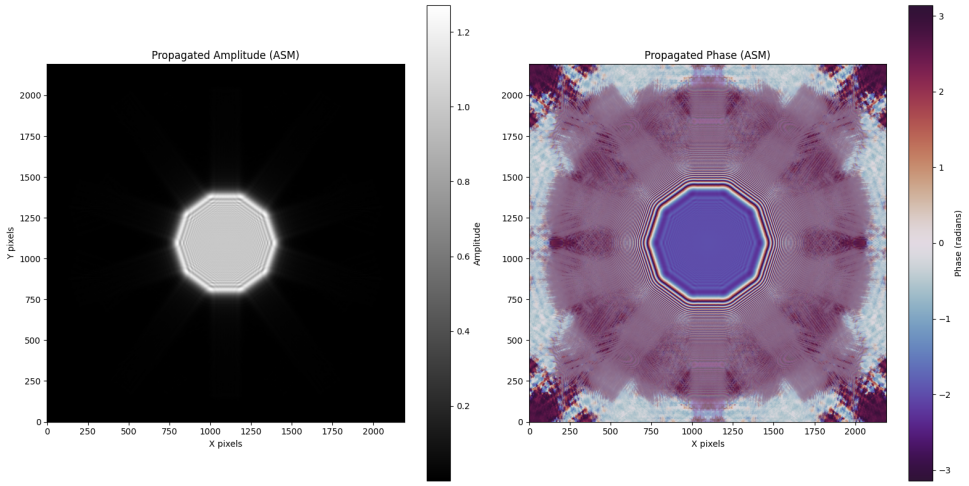

(a) Simulated propagation (ASM)

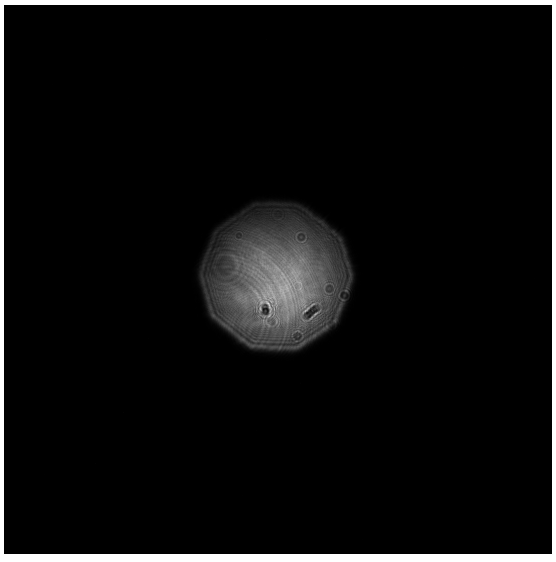

(b) Measured beam on detector

Figure 2.15: Angular Spectrum Method propagation of the 10-bladed iris aperture. (a) Simulated amplitude at the detector plane computed by the CUDA ASM kernel for a 10-sided polygonal pupil of approximately 1 mm opening propagated 100 mm, (b) corresponding beam profile captured on the GS3 detector in the real setup. The diffraction rings and the 10-fold symmetry of the iris blades are reproduced faithfully by the simulation.

This chapter describes the experimental setup and data-acquisition procedure for the ptychographic Brewster angle microscope (DC-BAPM). It covers the optical layout, the camera and detector system, the sample preparation, and the scanning strategy used to collect the set of diffraction patterns fed to the ePIE reconstruction pipeline.

3.1 Optical Layout

The full optical path of the DC-BAPM is shown schematically in Figure 3.1. Starting from the laser source, the beam passes through the following elements in sequence:

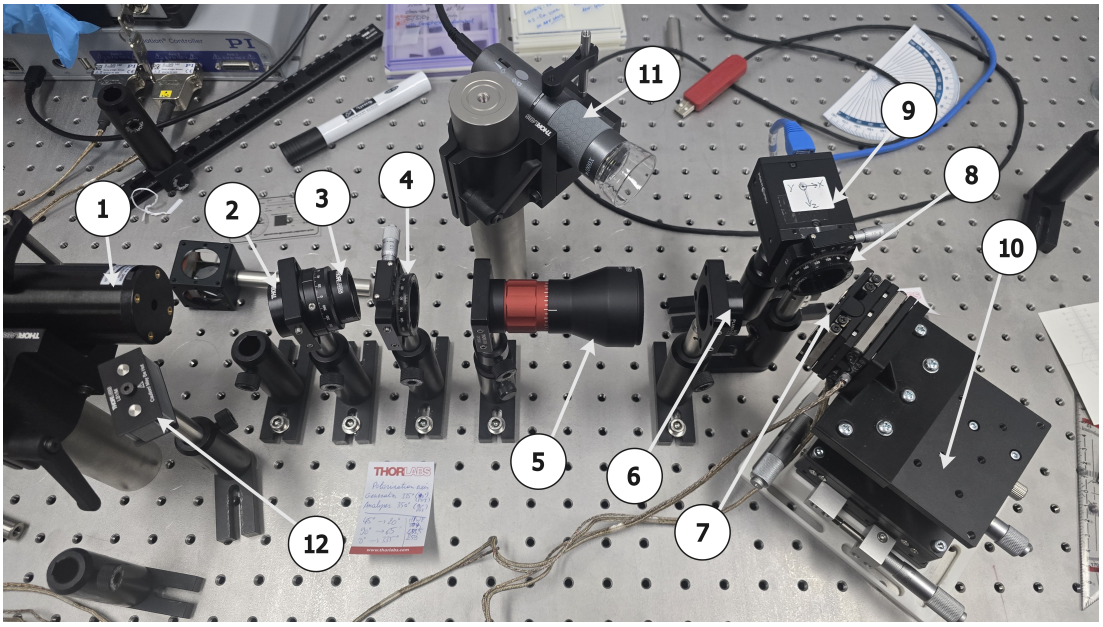

Figure 3.1: Experimental DC-BAPM setup.

1. **HeNe Laser** (Thorlabs HNL020LB). A linearly polarised HeNe laser at $\lambda = 632.8\,\text{nm}$ with a specified polarisation extinction ratio of 500:1.

2. **ND filter.** An ND40 neutral-density filter reduces the laser intensity by a factor of 40 to avoid saturating the detector.

3. **Half-wave plate** (Thorlabs AHWP10M-580). The achromatic half-wave plate rotates the polarisation angle of the beam, decoupling the orientation of the laser's intrinsic polarisation from the generator angle so that the two can be set independently.

4. **Generator** (Thorlabs LPVISE100-A in CRM1PT/M rotation mount). The 500:1 extinction ratio of the HeNe laser is insufficient for BAM: any residual s-component reflects at

the Brewster angle and adds a background reflection. The generator (a linear polariser placed after the laser, ND filter, and half-wave plate) raises the extinction ratio to $> 10^5$:1, making the illumination strictly linearly polarised.

5. **Beam expander** (Thorlabs GBE03-A). Two plano-convex lenses arranged as a Galilean telescope expand the beam by a factor of 3, from 0.63 mm to 1.89 mm diameter.
6. **Illumination lens** (Thorlabs LA1131-A-ML). This lens focuses the expanded beam onto the sample to form the ptychographic probe.
7. **Sample and Piezo actuators** (PI Q-545.140, ×2, driven by E-873.1AT controller). Two Q-Motion piezoelectric linear stages provide the xy-translation of the sample during the ptychographic scan.
8. **Analyser** (Thorlabs LPVISE100-A in CRM1PT/M rotation mount). A thin film can partially convert p-polarised reflected light into s-polarised light. The analyser selects which polarisation component reaches the detector. Setting it to pass p-polarised light is the standard BAM configuration, rotating it to s-polarisation gives access to ellipsometric contrast.
9. **Detector** (FLIR Grasshopper3 GS3-PGE-60S6C-C). Records the far-field diffraction intensity at each scan position.
10. **Sample stage** (Amazon LD90-LM). A 90 mm × 90 mm XY trimming stage provides coarse positioning of the sample.
11. **Digital microscope**. Used to view the probe footprint on the sample and to position the scan area before acquisition.
12. **Laser dump** (Thorlabs LB1/M). A blackened aluminium beam block mounted on the optical rail to absorb the back-reflection from the ND filter, preventing stray beams from becoming a safety hazard in the lab.

3.2 Camera and Detector System

The detector is a charge-coupled device (CCD) camera (FLIR Grasshopper3 GS3-PGE-60S6C-C) with a Sony ICX694 sensor, chosen for its large 1-inch format, global shutter, and high dynamic range. The key specifications relevant to ptychographic data collection are summarised in Table 3.1.

Table 3.1: CCD camera specifications relevant to ptychographic data acquisition (FLIR Grasshopper3 GS3-PGE-60S6C-C, Sony ICX694 sensor).

Parameter	Value	Unit
Sensor	Sony ICX694 CCD	—
Sensor format	2736 × 2192	pixels
Pixel size	4.54	µm
Optical format	1	inch
Dynamic range	16	bit
Frame rate	13	fps

3.2.1 Nyquist criterion

To avoid aliasing, the detector must sample the diffraction pattern at a rate of at least twice the highest spatial frequency present. In the far-field geometry the diffraction pattern is the Fourier transform of the exit wave, so this turns into a geometric constraint on the pixel pitch and the sample-to-detector distance.

The finest fringes in the diffraction pattern come from the largest illuminated feature, which is the probe footprint. A probe of extent D produces fringes spaced by $\lambda z/D$ on the detector. Sampling these with at least two pixels per fringe requires

$$D \leq \frac{\lambda z}{2\, d_{\mathrm{px}}}. \tag{3.1}$$

With $\lambda = 632.8\,\mathrm{nm}$, $z = 35\,\mathrm{mm}$, and $d_{\mathrm{px}} = 4.54\,\mu\mathrm{m}$, this gives a maximum probe size of $D_{\max} \approx 2.4\,\mathrm{mm}$. The probe footprint is $431\,\mu\mathrm{m} \times 320\,\mu\mathrm{m}$, well below this limit, so the Nyquist criterion is satisfied with a comfortable margin.

The same condition fixes the field of view of the reconstruction,

$$\mathrm{FOV} = \frac{\lambda z}{d_{\mathrm{px}}} \approx 4.9\,\mathrm{mm}. \tag{3.2}$$

The probe must fit inside this window or it wraps around and aliases. In practice the criterion sets a ceiling on the probe size, or a floor on z, for a given pixel pitch.

3.2.2 Alignment procedure

The hardest part of building the setup is aligning the optics. The step-by-step procedure I followed is given in Appendix 11.5.

3.3 Scan Strategy

3.3.1 Grid Pattern

The probe is scanned over the sample on a rectangular grid with $N_x \times N_y$ positions. Because the Brewster angle elongates the probe footprint along the plane of incidence by $1/\cos\theta_B$, the step sizes in the two directions are chosen as

$$\Delta x_{\parallel} = (1-o)\, r_{\parallel}, \qquad \Delta x_{\perp} = (1-o)\, r_{\perp}, \tag{3.3}$$

where $r_{\parallel} = r_0/\cos\theta_B$, $r_{\perp} = r_0$ are the probe semi-axes, and $o \approx 0.65$ is the target overlap fraction. This ensures that the overlap criterion is met equally in both lateral directions.

Visiting these grid positions in raster order has a well-known drawback: the periodic step pattern imprints itself onto the reconstruction. The result is a faint rectangular grid that lies on top of the retrieved object, clearly visible in Fig. 3.2 as a regular checkerboard texture across the amplitude, phase, and height maps of a Siemens-star reconstruction.

3.3.2 Fermat spiral

The scan positions are visited in a pseudo-random order (Fermat spiral or random permutation of the grid) rather than a raster sequence. Random ordering prevents the development of periodic artefacts in the reconstruction that can arise when adjacent positions are updated consecutively [Bunk et al., 2008]. The position sequence is generated once per experiment and saved alongside the raw data to allow exact reproducibility.

Figure 3.3 shows both ordering schemes with the AOI-corrected probe footprint applied. Each ellipse is the $431\,\mu\mathrm{m} \times 320\,\mu\mathrm{m}$ probe footprint at one scan position, the elongation along x reflects the $1/\cos\theta_B$ factor introduced in Eq. (3.3). The grid version (a) places 204 points on

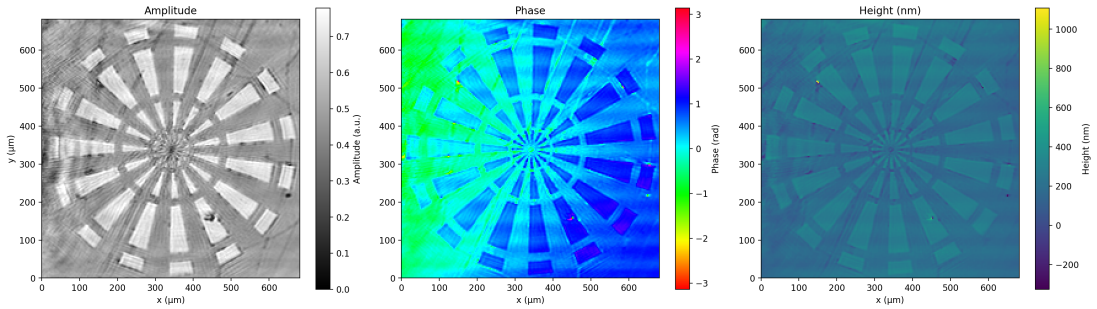

Figure 3.2: Siemens-star reconstruction obtained from a raster (grid) scan sequence. A regular rectangular pattern is clearly visible across the amplitude (left), phase (centre), and height (right) maps: a periodic artefact imprinted by the raster scan order. The Fermat-spiral ordering introduced in the next subsection eliminates this artefact.

a regular rectangular lattice, while the Fermat spiral (b) places 115 points in a pseudo-random pattern. Both reach the same 84.4% mean overlap, but only the Fermat spiral avoids the periodic artefact shown in Fig. 3.2.

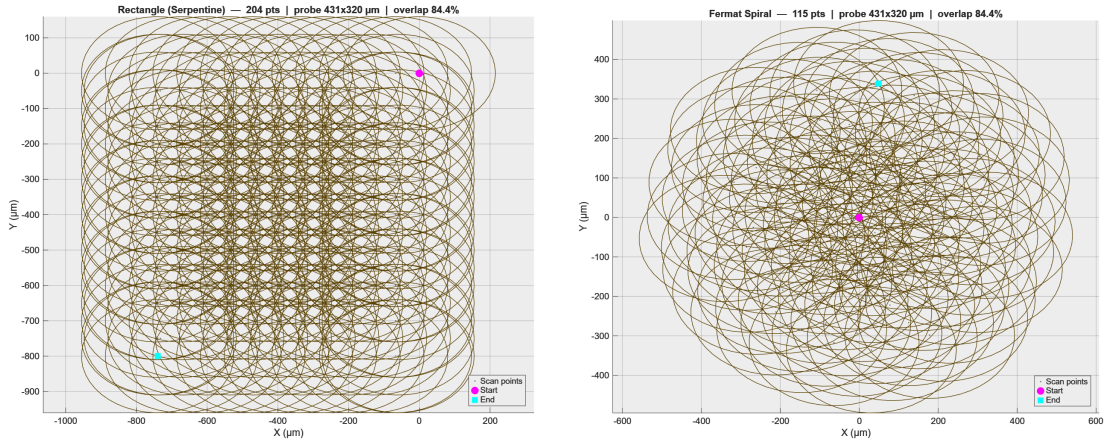

(a) Grid (serpentine) ordering, 204 points

(b) Fermat-spiral ordering, 115 points

Figure 3.3: AOI-corrected scan patterns for the same 431 µm × 320 µm elliptical probe footprint and 84.4% mean overlap, drawn as (a) a serpentine grid and (b) a Fermat spiral. The ellipses are elongated along the plane of incidence (x) by $1/\cos\theta_B$.

3.3.3 Typical Scan Parameters

Table 3.2 lists the scan parameters used in a representative measurement (Fermat-spiral scan recorded on 2026-04-30 at 42° incidence angle, sample-to-detector distance 35 mm).

Table 3.2: Ptychographic scan parameters for the 2026-04-30 reflection measurement.

Parameter	Value	Unit
Scan pattern	Fermat spiral	–
Number of positions	81	–
Probe semi-axis $r_\perp$	225	µm
Probe semi-axis $r_\parallel$	303	µm
AOI stretch $1/\cos\theta$	1.346	–
Step size $\Delta x_\perp$	60.0	µm
Step size $\Delta x_\parallel$	80.7	µm
Mean overlap o	86.7	%
Angle of incidence	42.0	°
Sample-to-detector distance	35.0	mm
Focal length (illumination lens)	50.0	mm
Aperture diameter	2.1	mm
Defocus	8	mm
Exposure time per frame	25	ms
Settle time per position	100	ms

3.4 High Dynamic Range Imaging

A diffraction pattern spans an enormous dynamic range. The central zero-order peak and low-spatial-frequency lobes are orders of magnitude brighter than the high-frequency fringes that encode the fine-detail object information. A 16-bit detector saturates at $2^{16} - 1 = 65{,}535$ counts. If the exposure is set to keep the zero-order peak unsaturated, the high-frequency tail falls below the noise floor. Conversely, an exposure long enough to resolve the tail causes the core to saturate and creates flat tops, corrupting the low-frequency data that ePIE relies on for convergence.

The way around this trade-off is to image each scan position at several exposure times and merge the brackets into a single HDR pattern. The procedure is showcased here on a graphene flake on silicon (Section 3.4.3), a weak scatterer that requires five exposures to span the dynamic range. A two-exposure validation on a PMMA on silicon Siemens star, used to confirm that the merge does not introduce step artefacts, is given in Appendix 11.5.

3.4.1 Multi-Exposure Acquisition

At each scan position the detector records a bracket of N exposures, saved as 16-bit TIFF frames together with the scan-position index and exposure time. The number of exposures and the choice of times depend on how strongly the sample scatters: a high-contrast test target needs only two brackets, a weak scatterer such as graphene needs four or five.

3.4.2 Saturation-Cascade Merge

The bracket is combined using a saturation-cascade algorithm. The frames are sorted from longest to shortest exposure. For each pixel, the algorithm selects the value from the longest frame in which that pixel is not saturated (i.e. below a threshold of 98% of the maximum detector count). Pixels that are saturated in all but the shortest frame fall back to the shortest exposure, which always provides an unsaturated value for the bright zero-order core. Each frame's values are scaled by the ratio of its exposure time to a reference time (the longest exposure) before selection, so all frames are expressed in consistent intensity units in the merged result. An optional cosine feather of a few pixels is applied at the saturation boundary to avoid hard step artefacts.

3.4.3 Application to a Graphene Flake

The Siemens star is a strong scatterer with sharp edges. A graphene monolayer on silicon, the sample of interest in this work, scatters far more weakly: its diffraction pattern has a much fainter high-frequency tail relative to the zero-order core. Two exposures are no longer enough, and the bracket is extended to five exposure times: 400, 5 000, 10 000, 50 000, and 200 000 µs. Figure 3.4 shows the resulting trade-off across the full bracket.

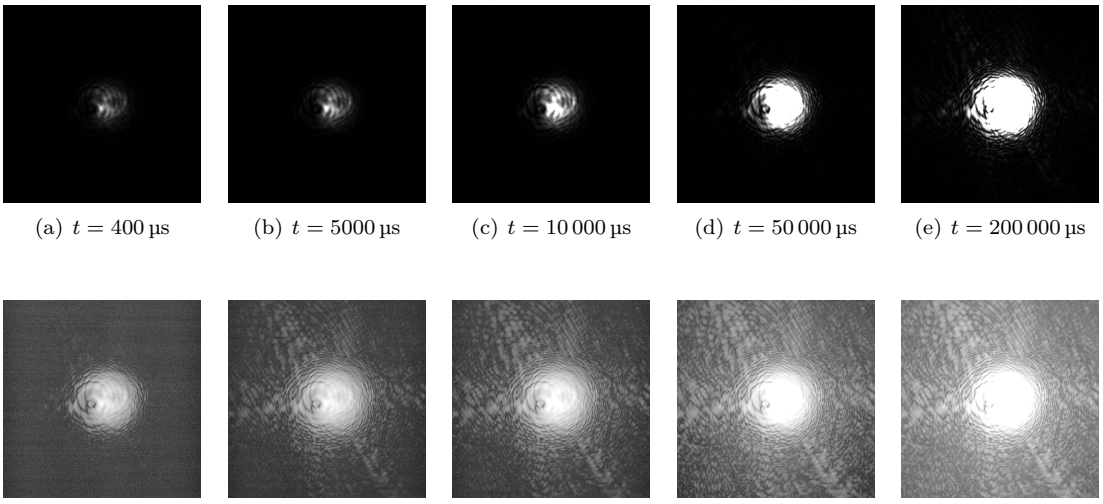

(a) $t = 400\,\mu s$ (b) $t = 5000\,\mu s$ (c) $t = 10\,000\,\mu s$ (d) $t = 50\,000\,\mu s$ (e) $t = 200\,000\,\mu s$

(f) $t = 400\,\mu s$ (g) $t = 5000\,\mu s$ (h) $t = 10\,000\,\mu s$ (i) $t = 50\,000\,\mu s$ (j) $t = 200\,000\,\mu s$

Figure 3.4: Five-exposure bracket on a graphene flake at a single scan position. Top row: linear scale. Bottom row: logarithmic scale. Short exposures keep the zero-order core unsaturated but leave the high-frequency tail in the noise, long exposures reveal the tail but clip the core.

The dominant-frame map and the continuity check across the merge boundaries are collected in Appendix 11.5 (Figs. 11 and 12). The five scaled radial profiles overlap smoothly and the merged result follows the dominant frame in each region without a visible discontinuity.

The merged HDR diffraction pattern is shown in Figure 3.5. It simultaneously captures the unsaturated zero-order core and the dim high-frequency fringes that would otherwise be invisible in a single-exposure measurement.

Figure 3.6 compares the ePIE reconstruction obtained from a standard single-exposure dataset against the reconstruction from the HDR-merged diffraction patterns. No clear improvement in resolution or contrast is visible in either the graphene or the Siemens star case (Appendix 11.5), despite the HDR diffraction pattern containing measurably more high-frequency spatial content than the short-exposure image alone.

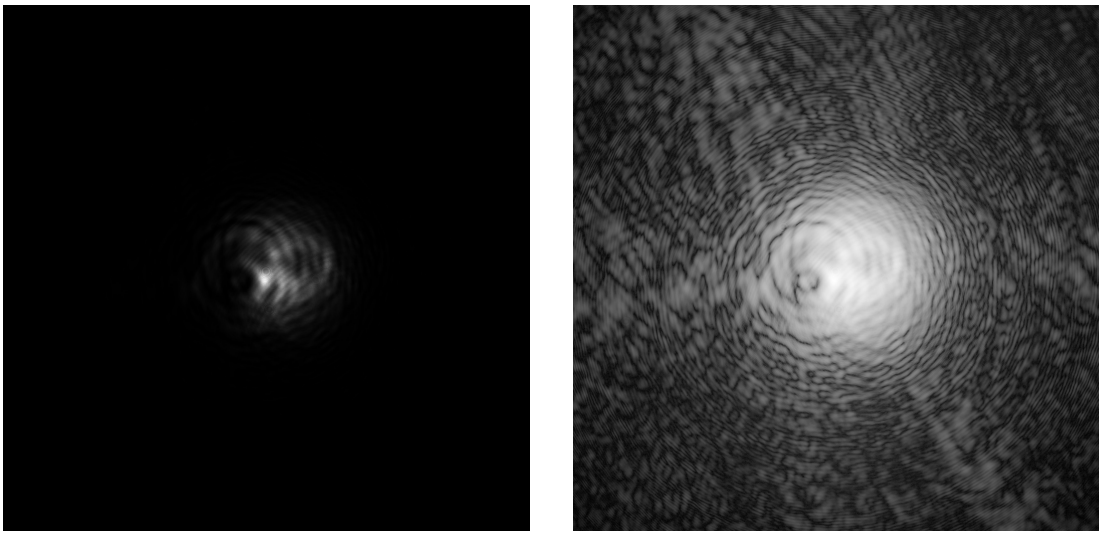

(a) Linear scale (b) Logarithmic scale

Figure 3.5: Merged HDR diffraction pattern of the graphene flake on (a) linear and (b) logarithmic intensity scales. The high-frequency fringe structure is clearly visible in the log-scale image while the zero-order core remains unsaturated.

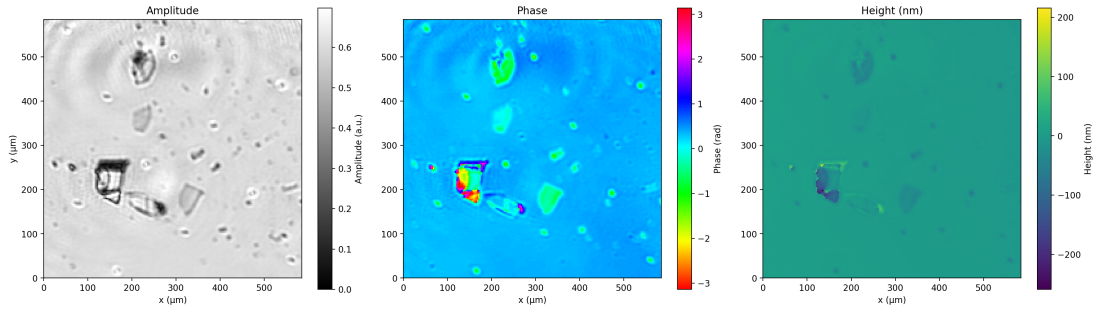

(a) Without HDR

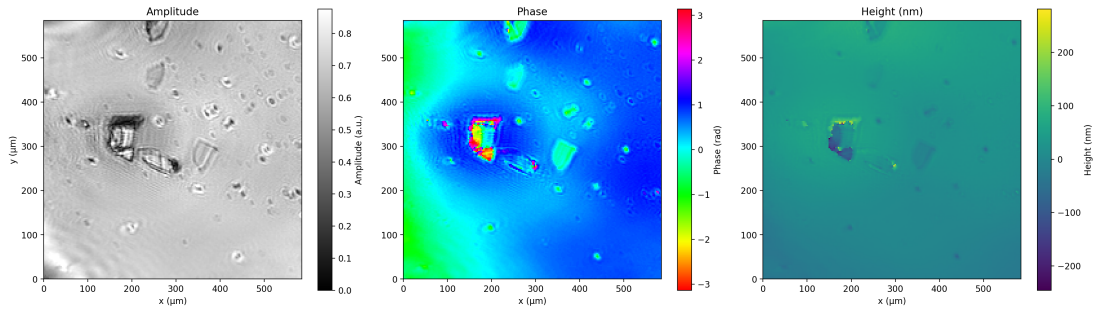

(b) With HDR

Figure 3.6: ePIE reconstruction of the graphene flake from (a) a single-exposure diffraction dataset and (b) the HDR-merged dataset. No clear difference in resolution or contrast is visible between the two reconstructions.

3.5 Contrast

The contrast of a diffraction pattern quantifies how well the bright and dark regions are separated on the detector. The standard measure is the Michelson contrast,

$$C = \frac{I_{\max} - I_{\min}}{I_{\max} + I_{\min}}, \tag{3.4}$$

where $I_{\max}$ and $I_{\min}$ are the maximum and minimum intensities in the image. $C = 1$ corresponds to a pattern where the darkest pixel is truly zero, and $C = 0$ corresponds to a uniform image with no contrast at all.

In practice, a single hot pixel or a bright dust scatter can dominate $I_{\max}$, making the raw formula unreliable. For diffraction images the 1st and 99th percentiles are used instead:

$$C = \frac{P_{99} - P_1}{P_{99} + P_1}. \tag{3.5}$$

This gives the same physical interpretation as Eq. (3.4) but is robust to outliers. A saturated diffraction pattern has P_{99} at the detector ceiling, reducing the measured contrast even though the true signal may be higher, images with a saturated fraction above 0.1% should therefore be treated with caution.

Figure 3.7 shows the three diffraction patterns on a logarithmic intensity scale. The computed metrics for each are listed in Table 1 in Appendix 11.5.

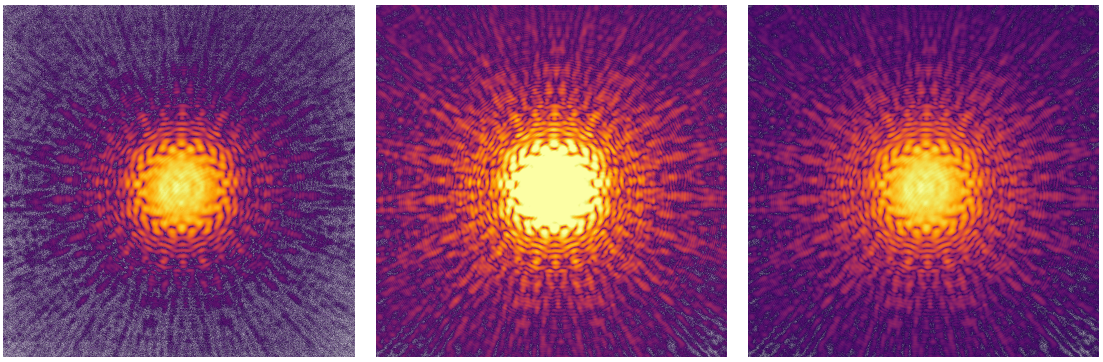

(a) $t = 10\,000\,\mu s$ (b) $t = 80\,000\,\mu s$ (c) HDR merge

Figure 3.7: Diffraction patterns of the siemens star on logarithmic intensity scale. (a) Short exposure: peak unsaturated but high-frequency tail is dim. (b) Long exposure: tail is well exposed but the zero-order core approaches saturation. (c) HDR merge: combines both, keeping the core unsaturated while retaining the high-frequency detail of the long exposure.

The Michelson contrast is 1 for all three images because the background is essentially dark ($P_1 \approx 0$), so the denominator $P_{99} + P_1 \approx P_{99}$ and the numerator $P_{99} - P_1 \approx P_{99}$. The dynamic range and peak-to-mean ratio are more informative here.

The 80 000 µs exposure reaches the highest dynamic range (96.3 dB) because the longer integration time brings the faint high-frequency tail well above the noise floor. Its peak-to-mean ratio of 14.8 is low compared to 57.9 at 10 000 µs, reflecting the fact that the mean count level rises steeply with exposure while the saturating peak is clamped by the detector ceiling. The HDR-merged pattern (89.3 dB, peak-to-mean 59.2) retains the short-exposure dynamic structure for the bright core while incorporating the long-exposure data for the dim outer fringes.

The Brewster angle is the angle of incidence at which p-polarised light experiences minimum reflectance at a planar dielectric interface. This property makes it the natural operating point for a microscope designed to image thin films at an interface: any reflected intensity arises entirely from the film rather than from the bulk substrate, providing exquisite surface sensitivity without requiring fluorescent labels or other contrast agents Hénon and Meunier, 1991.

4.1 Polarisation of Light

An electromagnetic plane wave propagating in the z-direction can be decomposed into two orthogonal polarisation states. The electric field oscillates in the xy-plane, and any polarisation state can be written as a superposition of:

- **s-polarisation** (TE, transverse electric): the electric field **E** is perpendicular to the plane of incidence.
- **p-polarisation** (TM, transverse magnetic): **E** lies in the plane of incidence.

In a Brewster angle microscope the incident beam is linearly polarised in the p-direction. A polariser placed before the sample selects the p-component, and an analyser can optionally be placed in the reflected beam path to further suppress background.

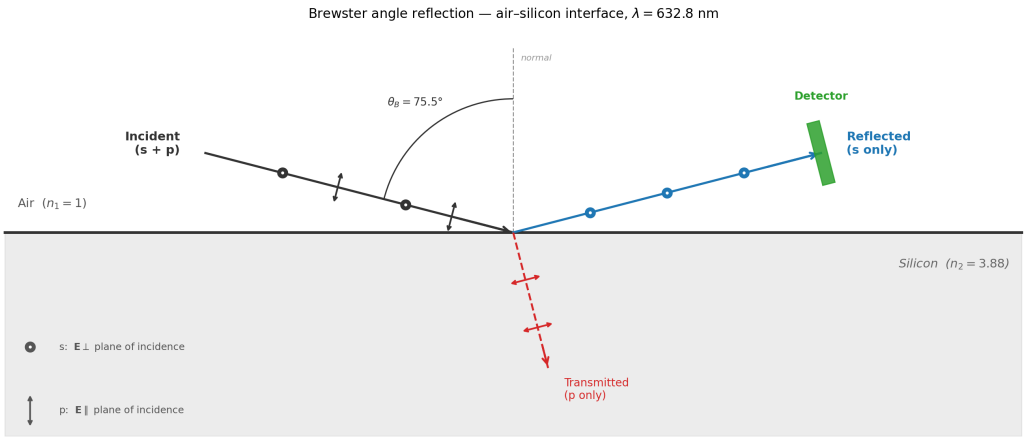

Figure 4.1: Brewster angle reflection at the air-PMMA interface ($\theta_B \approx 56.2°$). The incident beam carries both s- and p-polarised components: at the Brewster angle, p-polarised light is entirely transmitted into the PMMA while s-polarised light reflects toward the detector. Circles indicate the s-component (E-field perpendicular to the plane of incidence), double arrows indicate the p-component (E-field in the plane of incidence).

4.2 Fresnel Equations

When a plane wave is incident from medium 1 (refractive index n_1) onto medium 2 (refractive index n_2) at angle θ_i (measured from the surface normal), the reflected and transmitted amplitudes are given by the Fresnel equations. For s-polarisation:

$$r_s = \frac{n_1 \cos\theta_i - n_2 \cos\theta_t}{n_1 \cos\theta_i + n_2 \cos\theta_t}, \qquad t_s = \frac{2n_1 \cos\theta_i}{n_1 \cos\theta_i + n_2 \cos\theta_t}, \tag{4.1}$$

and for p-polarisation:

$$r_p = \frac{n_2 \cos\theta_i - n_1 \cos\theta_t}{n_2 \cos\theta_i + n_1 \cos\theta_t}, \qquad t_p = \frac{2n_1 \cos\theta_i}{n_2 \cos\theta_i + n_1 \cos\theta_t}, \tag{4.2}$$

where θ_t is the transmission angle related to θ_i by Snell's law:

$$n_1 \sin\theta_i = n_2 \sin\theta_t. \tag{4.3}$$

The reflectances $R_s = |r_s|^2$ and $R_p = |r_p|^2$ give the fraction of incident intensity that is reflected.

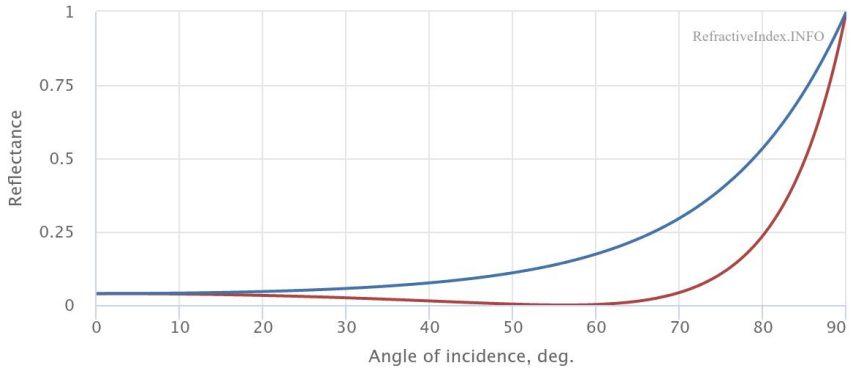

Figure 4.2: Fresnel reflectance of s- and p-polarisation as a function of angle of incidence for the air-PMMA interface ($n_1 = 1.000$, $n_2 = 1.490$). The p-reflectance vanishes at the Brewster angle $\theta_B \approx 56.2°$, the experimental operating angle of 42° is below this minimum but already shows significant s-p contrast.

4.3 The Brewster Angle

The p-reflectance R_p vanishes when the numerator of r_p in equation (4.2) is zero, i.e. when $n_2 \cos\theta_i = n_1 \cos\theta_t$. Combining this condition with Snell's law gives the Brewster angle:

$$\theta_B = \arctan\left(\frac{n_2}{n_1}\right). \tag{4.4}$$

At θ_B, the reflected and transmitted rays are perpendicular ($\theta_i + \theta_t = 90°$), so the oscillating dipoles induced in the medium by the transmitted wave radiate no energy in the reflected direction for p-polarisation.

For the air-PMMA interface, $n_1 = 1.000$ and $n_2 = 1.490$ at $\lambda = 632.8\,\text{nm}$, giving

$$\theta_B = \arctan(1.490) \approx 56.2°. \tag{4.5}$$

PMMA is non-absorbing at 632.8 nm, so $R_p = 0$ exactly at this angle. The silicon substrate below has $n_2 = 3.847 + 0.016\,\text{i}$ Jørgensen et al., 2024, giving a pseudo-Brewster angle at $\theta_B \approx 75.4°$ where R_p reaches a very small but non-zero minimum ($R_p \approx 5 \times 10^{-6}$).

4.4 Limitations of Conventional BAM

Despite its surface sensitivity, conventional BAM has two fundamental limitations that motivate the present work.

4.4.1 Keystone Distortion

The oblique angle of incidence causes a rectangular region on the sample to appear as a compressed trapezoid on the camera (keystone effect). The compression factor along the plane of incidence is $\cos\theta_B \approx 0.56$ for the air-PMMA interface at $\theta_B = 56.2°$, a 44% compression (see Chapter 9 for a quantitative treatment).

4.4.2 Phase Insensitivity

BAM records only intensity, discarding the phase of the reflected wavefield. The phase carries information about the optical thickness and the real part of the refractive index of the film that is inaccessible to conventional BAM. Ptychographic reconstruction (Chapter 6) recovers this information.

4.5 Experimental result of the polarization angle

These scans were taken at AOI = 42°, which is below the Brewster angle of the air–PMMA interface ($\theta_B^{\mathrm{PMMA}} = \arctan(1.49) \approx 56.2°$). The intended operating point was the PMMA Brewster angle, but at higher angles of incidence the ptychographic reconstruction becomes much harder to converge: the probe footprint stretches as $1/\cos\theta_i$ along the plane of incidence, alignment tolerances tighten, and position-correction failures become more frequent. At 42° the polarisation effect is still clearly visible in the reconstructions, and the algorithm converges reliably for all three PA settings, so this angle was used as a practical compromise between contrast and reconstruction stability.

To verify the polarisation sensitivity of the DC-BAPM I scanned four areas of the resolution target of silicon/PMMA (siemens star, line grating, mixed features, and an array of LED pads) at three generator angles: PA = 0° (p-polarisation, aligned with the plane of incidence), PA = 45°, and PA = 90° (s-polarisation). A half-wave plate rotates the polarisation vector by twice its own rotation angle, so the half-wave plate was advanced by half the angular step of the generator each time the generator was rotated. The polarisation incident on the generator therefore stayed aligned with the generator's transmission axis, the transmitted intensity stayed constant, and the exposure time could be kept the same across all three PA settings. Any change in detected intensity is then due to the reduction in p-polarised content reflected from the silicon substrate at the Brewster angle, not to a change in illumination power. All scans were taken at AOI = 42° on 2026-04-27. Each subfigure shows the reconstructed amplitude (left), phase (centre), and 3D height map (right) at the corresponding PA.

Across the four sample areas the same trend repeats: p-polarised illumination (PA = 0°) yields the strongest contrast between the patterned features and the silicon substrate, because $R_p < R_s$ at AOI = 42° even though we are below the Brewster minimum. At s-polarisation (PA = 90°) the substrate reflectivity is no longer reduced and the relative contrast collapses, although the algorithm still converges to a recognisable reconstruction. The intermediate PA = 45° case falls between the two as expected, since the incident field has equal s- and p-components.

One subtlety in the images: the amplitude reconstructions at the three PA settings look almost identical, while the polarisation effect is mostly visible in the phase. There are two reasons for this. First, ePIE does not preserve an absolute amplitude scale: the reconstructed object

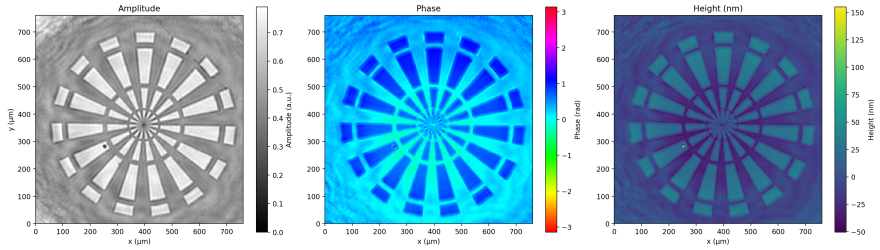

(a) PA = 0°

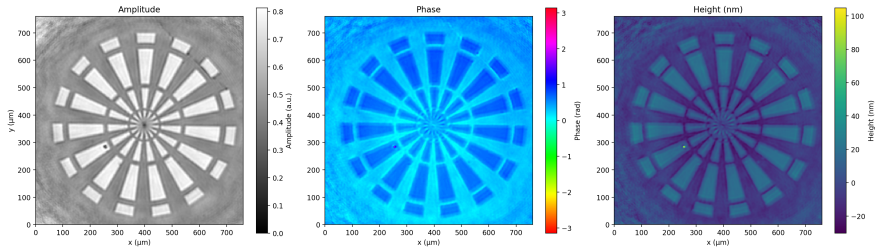

(b) PA = 45°

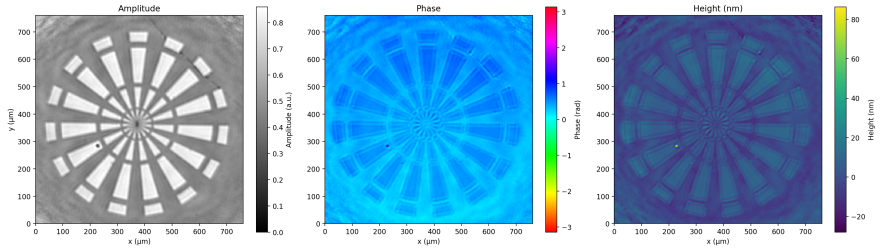

(c) PA = 90°

Figure 4.3: Siemens star reconstruction at three generator angles, taken at AOI = 42° (below the PMMA Brewster angle). The amplitude contrast is highest at PA = 0° (p-polarisation) and drops at PA = 90° (s-polarisation), consistent with the partial suppression of the p-component of the reflection on the approach to the Brewster minimum.

amplitude is normalised by the probe and rescaled to fit the dynamic range, so the order-of-magnitude drop in reflected intensity going from p- to s-illumination is washed out and what remains is the same spatial structure of the features. Second, the polarisation contrast near the Brewster angle is encoded primarily in phase. On the approach to θ_B, the Fresnel coefficient r_p heads toward zero and its phase swings rapidly through 180°, so $\partial \arg r_p / \partial d_f$ is already steep at 42° and a small optical thickness contributes a large phase shift. For s-polarisation r_s varies smoothly with no such inflection, so the phase response is much weaker. The amplitude images therefore mostly show the spatial layout of the features, while the phase images carry the polarisation-sensitive surface information that is the point of the experiment.

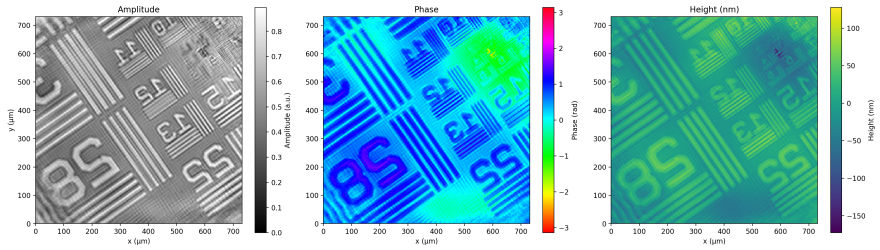

(a) PA = 0°

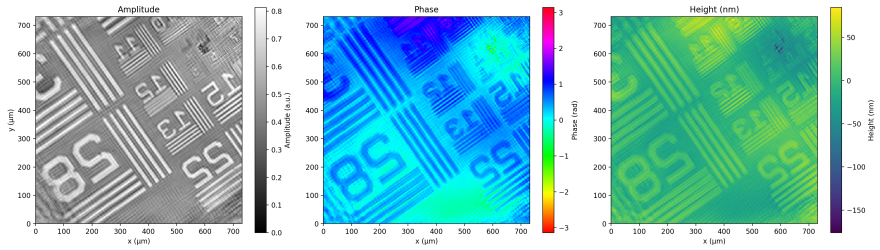

(b) PA = 45°

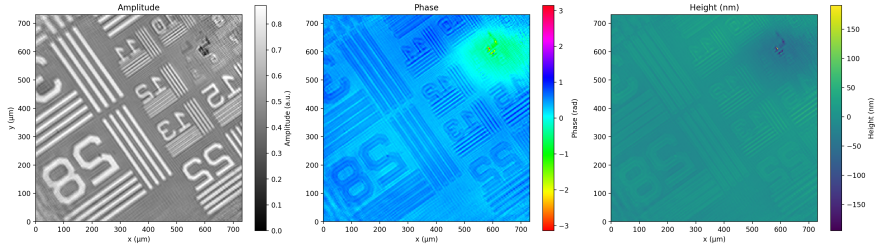

(c) PA = 90°

Figure 4.4: Line grating reconstruction at three generator angles, taken at AOI = 42° (below the PMMA Brewster angle). The line structure is resolved at all three settings, but the phase contrast across the lines is strongest at PA = 0°.

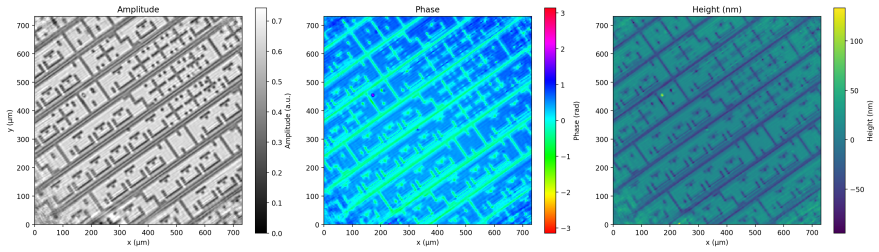

(a) PA = 0°

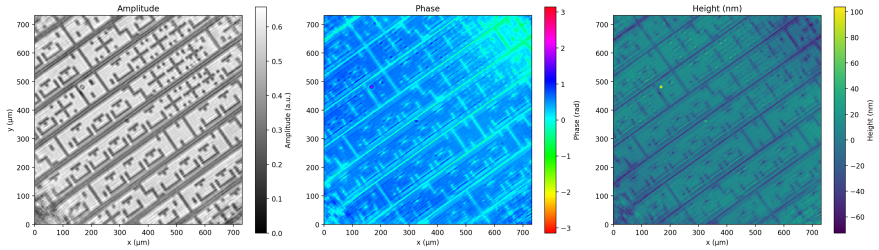

(b) PA = 45°

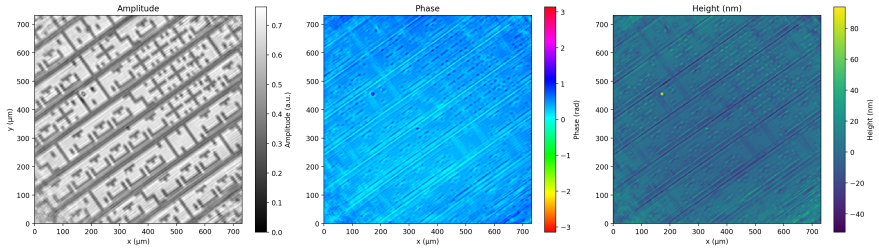

(c) PA = 90°

Figure 4.5: Mixed-feature area reconstruction at three generator angles, taken at AOI = 42° (below the PMMA Brewster angle). The small features remain resolved across all three settings, the phase image inverts in sign between PA = 0° and PA = 90°.

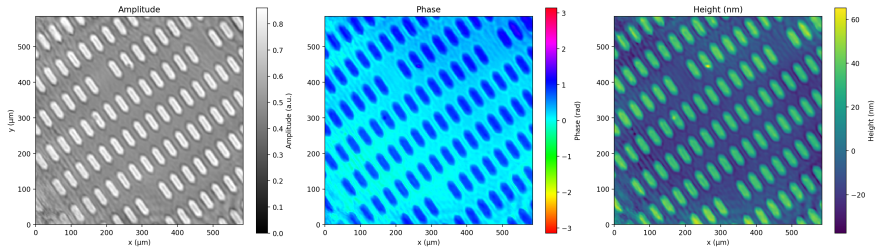

(a) PA = 0°

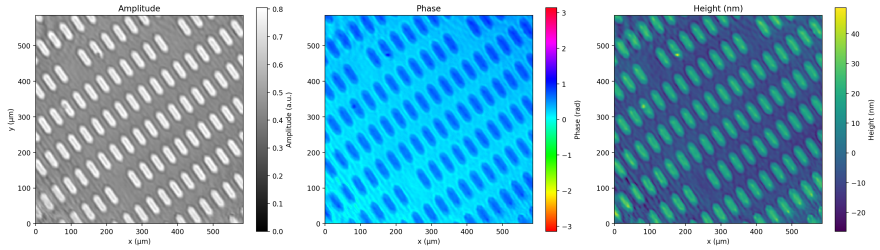

(b) PA = 45°

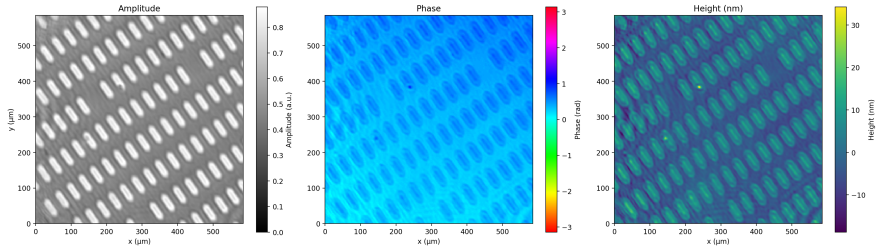

(c) PA = 90°

Figure 4.6: LED-pad array reconstruction at three generator angles, taken at AOI = 42° (below the PMMA Brewster angle). The pad outlines are visible in all three configurations, the contrast between pad and substrate is highest at PA = 0°.

At oblique incidence the detector lies in a flat plane while the physically meaningful data reside on the Ewald sphere, a curved surface in momentum space. Recording diffraction patterns on a flat detector therefore introduces a systematic, angle-dependent distortion: pixels near the centre map almost linearly to momentum transfer, but pixels near the edges are increasingly compressed. At the 42°–56° incidence angles used in the DC-BAPM, the flat-Ewald-sphere approximation breaks down and this distortion must be corrected before ptychographic inversion. The algorithm follows Gardner et al. Gardner, B. Zhang, et al., 2012, who derived the correction for high-numerical-aperture reflection-mode coherent diffraction imaging.

5.1 Geometry

Let the detector lie in the xy-plane at distance d from the sample, with pixel (x, y) denoting a physical position in metres relative to the detector centre. A scattered ray arriving at that pixel travels in the direction proportional to (x, y, d). The corresponding momentum-transfer vector on the Ewald sphere is

$$q_x = k_0 \frac{x}{r}, \qquad q_y = k_0 \frac{y}{r}, \qquad q_z = k_0 \left(\frac{d}{r} - 1 \right), \tag{5.1}$$

where $r = \sqrt{x^2 + y^2 + d^2}$ is the ray length and $k_0 = 2\pi/\lambda$. Because the sample is tilted relative to the detector normal, the Ewald-sphere vectors are additionally rotated by a 3×3 matrix $\mathbf{R} = R_z R_y R_x$ built from three Euler angles $(\theta_z, \theta_x, \theta_y)$ that describe the sample orientation. The in-plane rotated components are

$$\begin{pmatrix} q'_x \\ q'_y \end{pmatrix} = \begin{pmatrix} R_{00} & R_{10} & R_{20} \\ R_{01} & R_{11} & R_{21} \end{pmatrix} \begin{pmatrix} q_x \\ q_y \\ q_z \end{pmatrix}. \tag{5.2}$$

These are then rescaled to span the same pixel grid as the input image and each source pixel (i, j) is mapped forward to the output pixel (i', j') determined by (q'_x, q'_y). Any output pixel that receives no source contribution is an unmapped gap that must be filled separately.

The recorded intensity must be corrected for the solid-angle Jacobian of the flat-to-sphere mapping before remapping. The required factor is

$$I_{\text{Ewald}}(x, y) = \underbrace{\left(1 + \frac{x^2 + y^2}{d^2} \right)}_{1/r^2 \text{ factor}} I_{\text{det}}(x, y), \tag{5.3}$$

and it is applied to each pixel before the remapping step.

5.2 Implementation

The coordinate map depends only on the detector geometry, not on the diffraction data, so it is precomputed once per scan and reused for every pattern in the stack. The listing below shows the core of the Python implementation.

Listing 5.1: Core of the conical distortion correction. `precompute_mapping` builds the pixel-to-pixel index arrays once per scan, `correct_single_image` applies them to each frame with a single scatter-index assignment.

```
def precompute_mapping(m, n, dx, dy, d, wavelength,
                       ang_z, ang_x, ang_y):
    k0 = 2 * pi / wavelength
    R  = build_rotation_matrix(ang_z, ang_x, ang_y)  # 3x3

    # Detector-plane coordinates (metres, centred)
    x_1d = arange(1, m+1)*dx - dx*m/2
    y_1d = arange(1, n+1)*dy - dy*n/2
    y_grid, x_grid = meshgrid(y_1d, x_1d)

    # Ewald sphere mapping  (Eq. 1)
    r       = sqrt(x_grid**2 + y_grid**2 + d**2)
    qx      = k0 * x_grid / r
    qy      = k0 * y_grid / r
    qz      = k0 * (d / r - 1.0)
    rescale = 1.0 + (x_grid**2 + y_grid**2) / d**2  # Eq. 3

    # Rotate q-vectors for tilted sample  (Eq. 2)
    rx = R[0,0]*qx + R[1,0]*qy + R[2,0]*qz
    ry = R[0,1]*qx + R[1,1]*qy + R[2,1]*qz

    # Project back to pixel indices
    v     = max(x_grid) / max(rx)           # scale factor
    i_out = round((rx*v + dx*m/2) / dx)
    j_out = round((ry*v + dy*n/2) / dy)
    valid = (i_out>=0)&(i_out<m)&(j_out>=0)&(j_out<n)
    return i_out, j_out, valid, rescale


def correct_single_image(image, i_out, j_out, valid, rescale):
    img    = fftshift(image)                    # centre zero-order
    Ir     = rescale * img                      # solid-angle correction
    output = zeros((ny, nx))
    output[i_out[valid], j_out[valid]] = Ir[valid]  # forward remap
    output = inpaint_nans(output)               # fill unmapped gaps
    output = fftshift(fftshift(output))         # restore layout
    return output
```

The gap-filling step (`inpaint_nans`) solves the discrete Laplace equation over the unmapped pixels by sparse linear system, equivalent to harmonic interpolation from the surrounding valid data without introducing spectral bias. Boundary regions that fall entirely outside the Ewald-sphere coverage are filled with noise sampled from the image interior so that ePIE does not interpret the absence of signal as a physical zero.

5.2.1 Numerical example

For a pure rotation around the y-axis by angle θ, the full 3×3 rotation matrix reduces to

$$\mathbf{R}_y(\theta) = \begin{pmatrix} \cos\theta & 0 & \sin\theta \\ 0 & 1 & 0 \\ -\sin\theta & 0 & \cos\theta \end{pmatrix}. \tag{5.4}$$

For the experimental parameters ($\lambda = 632.8\,\text{nm}$, $d = 45\,\text{mm}$, $\Delta x = 4.54\,\mu\text{m}$, $N = 2192$, $\theta = 42°$) this evaluates to

$$\mathbf{R}_y(42°) = \begin{pmatrix} 0.743 & 0 & 0.669 \\ 0 & 1 & 0 \\ -0.669 & 0 & 0.743 \end{pmatrix}. \tag{5.5}$$

The solid-angle intensity correction factor (Eq. 5.3) reaches its maximum at the detector corner, where $x = y = N\Delta x/2 = 4.98\,\mathrm{mm}$:

$$1 + \frac{x^2 + y^2}{d^2} = 1 + \frac{2 \times (4.98\,\mathrm{mm})^2}{(45\,\mathrm{mm})^2} \approx 1.024, \tag{5.6}$$

a 2.4 % increase relative to the detector centre.

5.3 Validation

Figure 5.1 shows a single raw diffraction pattern recorded from the Siemens-star target at 42° incidence alongside the corrected pattern produced by the algorithm. The raw pattern is visibly elongated along the x-axis and asymmetric due to the Ewald-sphere distortion. After the correction the zero-order peak is recentred and the fringe structure is remapped to a uniform momentum-space grid. The noise-filled border on the right of the corrected pattern marks the region that falls outside the Ewald-sphere coverage at this incidence angle.

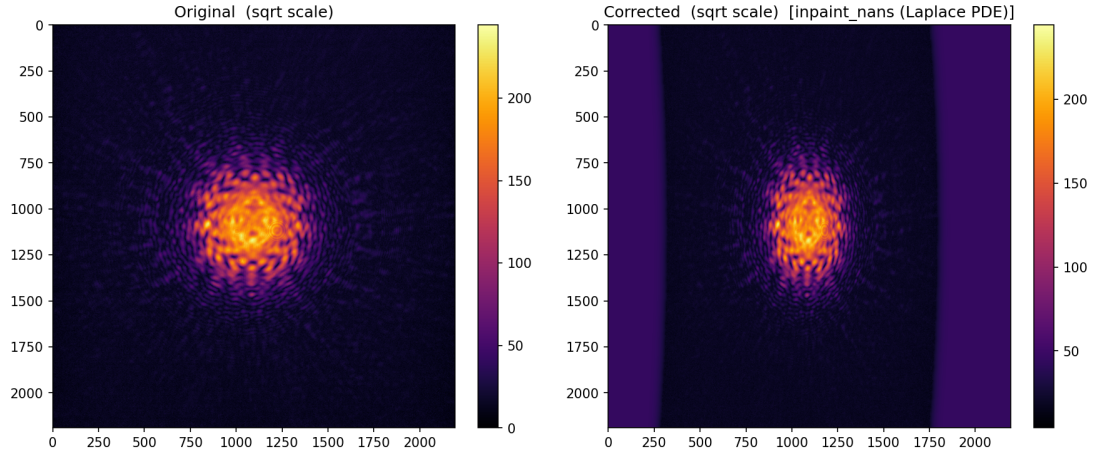

Figure 5.1: Diffraction pattern of the Siemens-star target at 42° incidence before (left) and after (right) the conical distortion correction, displayed on a square-root intensity scale. The raw pattern is elongated along x by the Ewald-sphere distortion, the corrected pattern is remapped to a uniform q-grid. The noise-filled border on the right is the region outside the Ewald-sphere coverage.

The effect of the correction is validated by running the full ptychographic reconstruction on the Siemens-star test target with and without it applied. The result is shown in Fig. 5.2. Without the correction, the Ewald-sphere distortion scrambles the high-frequency content of the diffraction patterns so that ePIE cannot consistently superpose the overlapping exit waves, and the spokes collapse into a noisy, aliased blur. After the correction the patterns share a single, undistorted momentum-space geometry, the overlap constraint is satisfied, and the reconstruction recovers a clean Siemens star with sharp spokes in amplitude, phase, and height.

Figure 5.3 shows the corrected diffraction pattern of a square region of the logic pattern on a logarithmic greyscale. A square aperture produces a cross-shaped diffraction pattern, the vertical arm of the cross is curved in the raw pattern because the Ewald-sphere warp acts predominantly along x, introducing a position-dependent displacement that bends the otherwise straight fringe. After correction the vertical arm is straightened, confirming that the remapping is applied correctly. In future work, a dedicated calibration sample consisting only of an isolated square patch could make this effect more pronounced and allow the incidence angle to be extracted directly from the curvature of the vertical diffraction arm, providing a direct measurement of the tilt angle used in the correction.

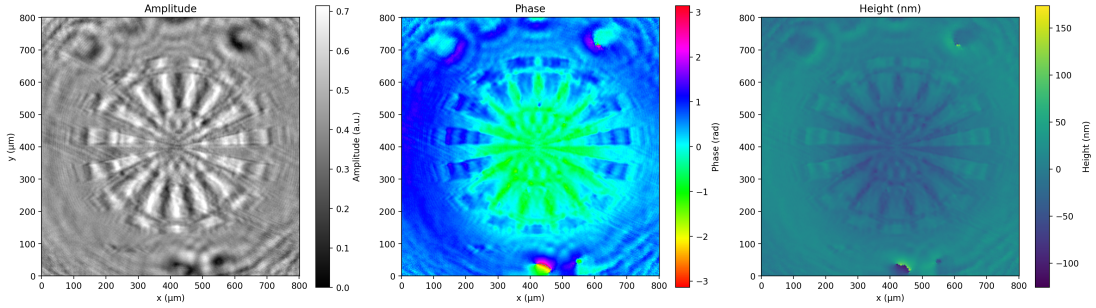

(a) Without conical correction

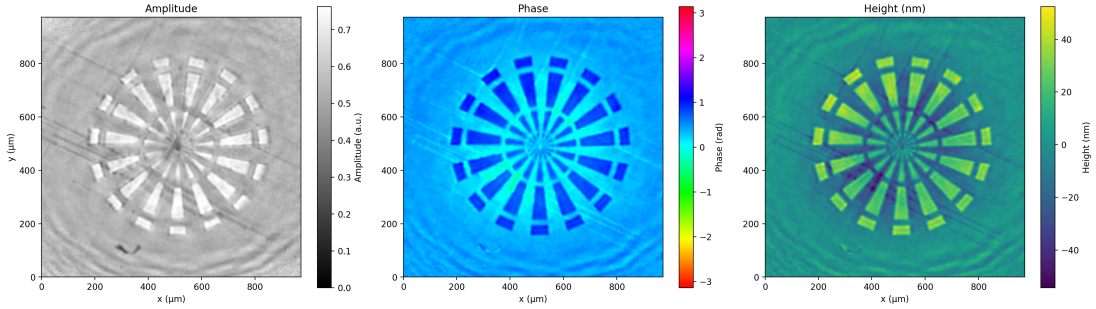

(b) With conical correction

Figure 5.2: ePIE reconstruction of a PMMA-on-silicon Siemens star from the same raw dataset, processed (a) without and (b) with the conical distortion correction. Each panel shows the retrieved amplitude (left), phase (centre), and height (right). Without correction the spokes are unresolved and the maps are dominated by aliasing, with correction the spokes are sharp and the height range collapses from ± 150 nm to ± 50 nm, confirming that the Ewald-sphere distortion was the dominant error source.

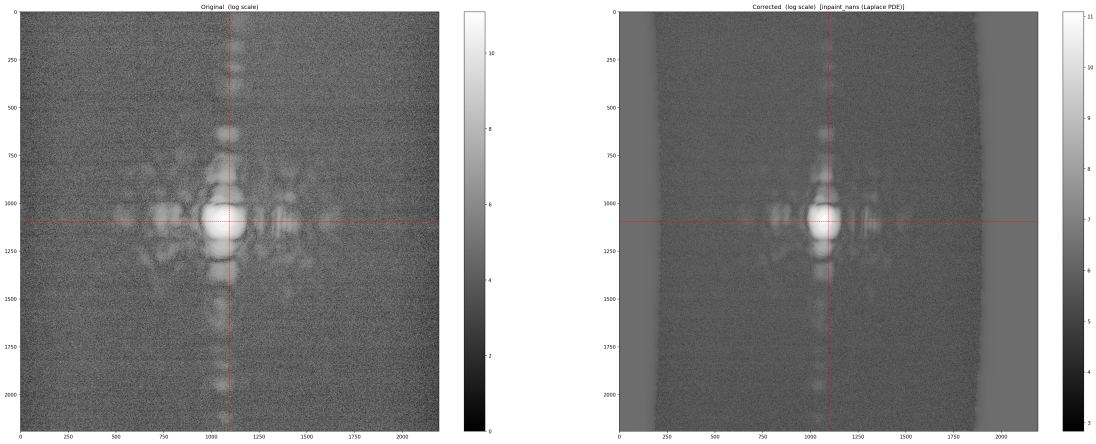

Figure 5.3: Diffraction pattern of a square region of the logic pattern after the conical distortion correction, displayed on a logarithmic greyscale. The square aperture produces a characteristic cross-shaped diffraction pattern. The curvature visible in the vertical arm of the cross in the uncorrected pattern is a direct signature of the Ewald-sphere distortion, after correction the arm is straightened. The uniform-texture border regions on the sides are filled with a constant noise floor because they fall outside the Ewald-sphere coverage at this incidence angle.

This chapter covers the middle stage of the DC-BAPM reconstruction pipeline: the iterative phase retrieval that converts the corrected diffraction patterns recorded in Chapter 3 into a complex-valued image of the sample. The pre-processing stage that produces those corrected patterns (distortion mapping and pixel-wise warping) is described in Chapter 5. The post-processing stage that converts the recovered phase into a topographic surface map is described in Chapter 7.

6.1 Reconstruction Engine

The reconstruction engine used in this work is the EUVPtycho C++/CUDA code developed at the Paul Scherrer Institute in Switzerland. To make the build reproducible across workstations, the full toolchain (CUDA, Armadillo, OpenCV, FFTW, HDF5, Boost, Qt5) is packaged into a Docker image and run with the NVIDIA Container Toolkit. The full Dockerfile recipe is given in Appendix 11.5. At runtime the engine is driven by a single INI-style parameter file that specifies the wavelength, detector geometry, probe-initialisation model, ePIE feedback constants, and convergence schedule. The parameter file used for a representative 42°-incidence reconstruction is reproduced in Appendix 11.5.

6.2 Reconstruction Procedure

6.2.1 Initialisation

The reconstruction is initialised as follows:

- **Object.** The object amplitude and phase are seeded with uniform random noise on the unit interval. No prior information about the sample is assumed.
- **Probe.** The amplitude and phase are obtained by propagating the aperture field to the sample plane with the Angular Spectrum Method (Section 2.6).

The four initial guesses fed into the first ePIE iteration are shown in Fig. 6.1. The object panels are pure noise, the probe panels already carry the expected Gaussian intensity and phase.

6.2.2 ePIE Iteration

One full ePIE epoch visits all M scan positions in random order and applies the update equations derived in Chapter 2 (equations 2.5 and 2.6). The feedback parameter β_O for the object update is set to 0.9 throughout.

For the first 100 epochs the probe is held fixed at the Gaussian beam profile computed by the Angular Spectrum Method (Section 2.6), and only the object is updated. This ensures that the object estimate has converged to a reasonable solution before the probe is allowed to change. From epoch 100 onwards the probe update (Eq. 2.6) is activated with feedback parameter $\beta_P = 0.9$, allowing the algorithm to jointly refine both the object and the probe.

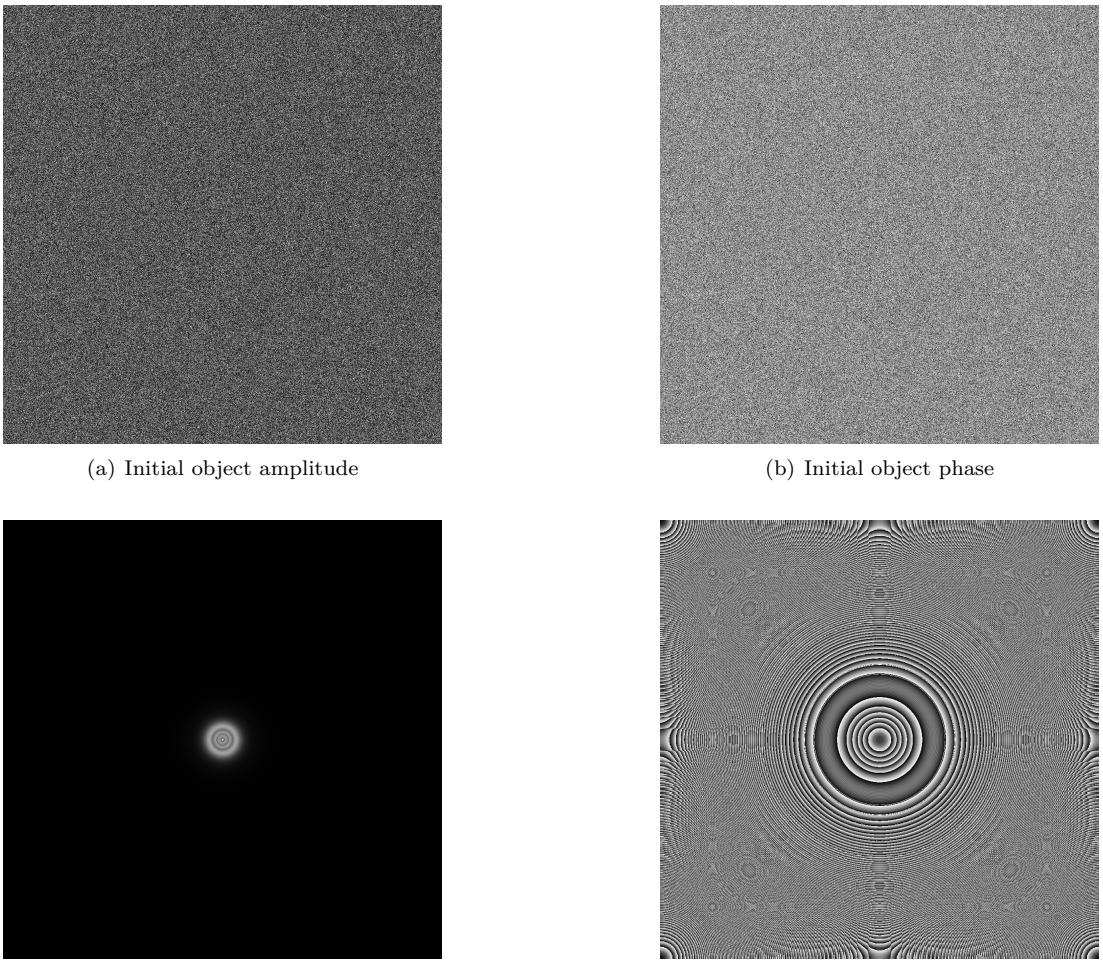

(a) Initial object amplitude

(b) Initial object phase

(c) Initial probe amplitude

(d) Initial probe phase

Figure 6.1: Initial guesses fed into the ePIE algorithm. Top row: object amplitude and phase, seeded with uniform random noise. Bottom row: probe amplitude (Gaussian intensity) and phase (spherical defocus term) from the forward illumination model.

6.2.3 Convergence Criterion

The reconstruction is considered converged when the normalised error metric

$$E_n = \frac{\sum_j \sum_{\mathbf{q}} \left(\sqrt{I_j^{\text{corr}}(\mathbf{q})} - |\tilde{\Psi}_{j,n}(\mathbf{q})| \right)^2}{\sum_j \sum_{\mathbf{q}} I_j^{\text{corr}}(\mathbf{q})}, \tag{6.1}$$

falls below $4 \cdot 10^6$. Here j indexes the M scan positions and $\mathbf{q}$ the detector pixels, $I_j^{\text{corr}}(\mathbf{q})$ is the distortion-corrected diffraction intensity at scan position j, and $\tilde{\Psi}_{j,n}(\mathbf{q})$ is the Fourier-plane estimate of the exit wave at iteration n. The numerator measures the squared amplitude mismatch between measurement and model (the square root puts the recorded intensity on the same scale as the model amplitude $|\tilde{\Psi}|$), while the denominator normalises by the total measured intensity so that E_n is dimensionless and comparable across datasets.

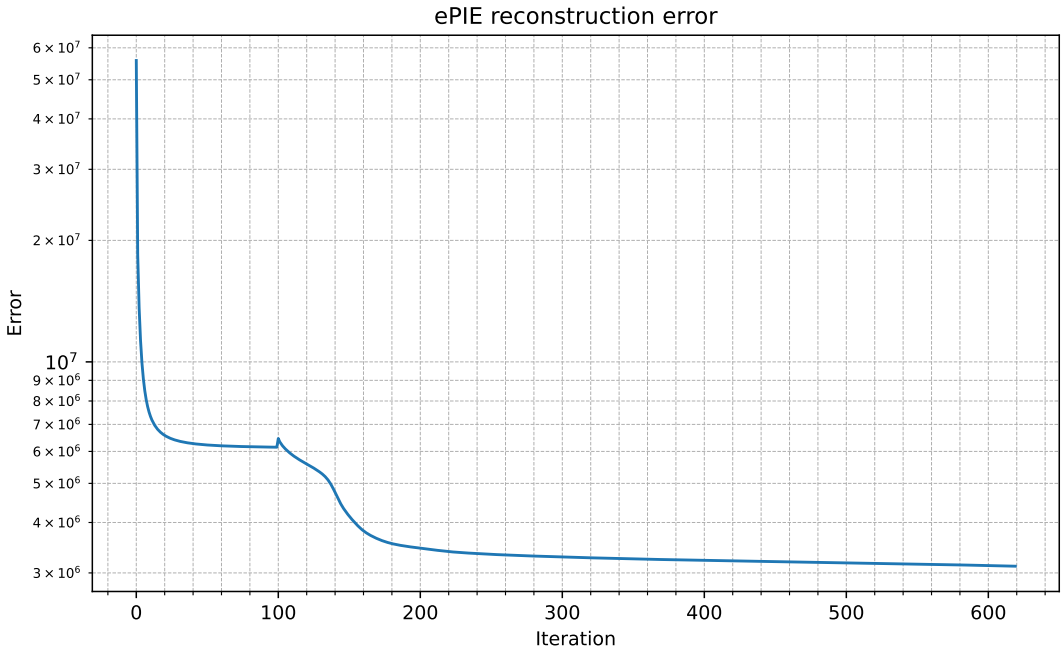

Figure 6.2: Convergence of the ePIE reconstruction for a representative dataset, with the error shown on a logarithmic scale. The error drops rapidly in the first iterations as the object estimate forms, then plateaus while the probe is still fixed. The step decrease around iteration 100 corresponds to the activation of the probe update, after which the joint object-probe refinement drives the error to its final value.

6.3 Resolution

The resolution of the DC-BAPM is quantified by reconstructing a Siemens-star test target patterned as PMMA on a silicon substrate, and reading off the smallest spoke pitch that the algorithm still resolves. The full GDS layout and a labelled optical microscopic image of the sample are reproduced in Appendix 11.5. Figure 6.4 shows the converged ePIE reconstruction used for that measurement: the object amplitude and phase at iteration 400, together with the simultaneously retrieved probe.

6.3.1 Measurement procedure

The Siemens-star geometry converts radial position into spatial frequency: the same number of spokes wraps around every circle, so a circle of radius r drawn in the reconstruction crosses each spoke at a half-pitch

$$\mathrm{hp}(r) = \frac{\pi\, r}{N_{\mathrm{spokes}}}, \tag{6.2}$$

where N_{spokes} is the number of spoke pairs on the target. Sampling the reconstructed amplitude along that circle therefore produces an oscillating signal whose period equals the half-pitch at radius r. Moving the sampling circle inwards (smaller r, smaller hp) tests finer features, the smallest r at which the signal is still resolved sets the resolution of the reconstruction.

6.3.2 Normalised Image Log-Slope

The sampled profile is reduced to a single sharpness number through the Normalised Image Log-Slope [Mack, 2007, Sec. 3.8.3],

$$\mathrm{NILS}(x) = \mathrm{hp} \left| \frac{\mathrm{d} \ln I(x)}{\mathrm{d}x} \right|, \tag{6.3}$$

where $I(x)$ is the (bilateral-filtered, $[0, 1]$-normalised) amplitude sampled along the circle and hp is the half-pitch at that radius. Multiplying the log-slope by hp makes the metric dimensionless, so NILS values can be compared across radii. The geometric picture is sketched in Fig. 6.3: a sharp binary mask is blurred into a smoothed intensity profile, and NILS measures the slope of that profile at the edge transition (normalised by hp). NILS is a continuous sharpness measure, not a binary cut-off: the visual resolution is read off by scanning radii and looking for where the mean NILS collapses into the noise floor while the spokes also stop being visually distinguishable.

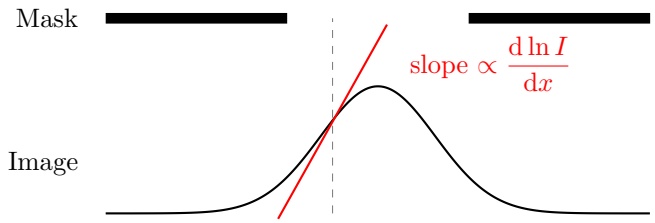

Figure 6.3: Geometric picture of NILS, after Mack [Mack, 2007, Fig. 3.35]. A sharp binary mask (top, two black bars with a clear gap) is blurred by the imaging system into a continuous intensity profile (bottom). The red tangent line at the rising edge has slope $\mathrm{d} \ln I / \mathrm{d}x$, NILS multiplies that slope by the local half-pitch (Eq. (6.3)) to produce a dimensionless number that grows with edge sharpness.

Figure 6.4 shows the converged reconstruction used for the measurement, and Fig. 6.5 applies the procedure at four half-pitches between 97 µm and 11 µm. Peak NILS drops monotonically with half-pitch, but the expected fifteen oscillations per turn are visible in every row, and the spokes remain visually distinguishable even at 11 µm (one short angular sector loses contrast). The visual resolution of this reconstruction is therefore 11 µm.

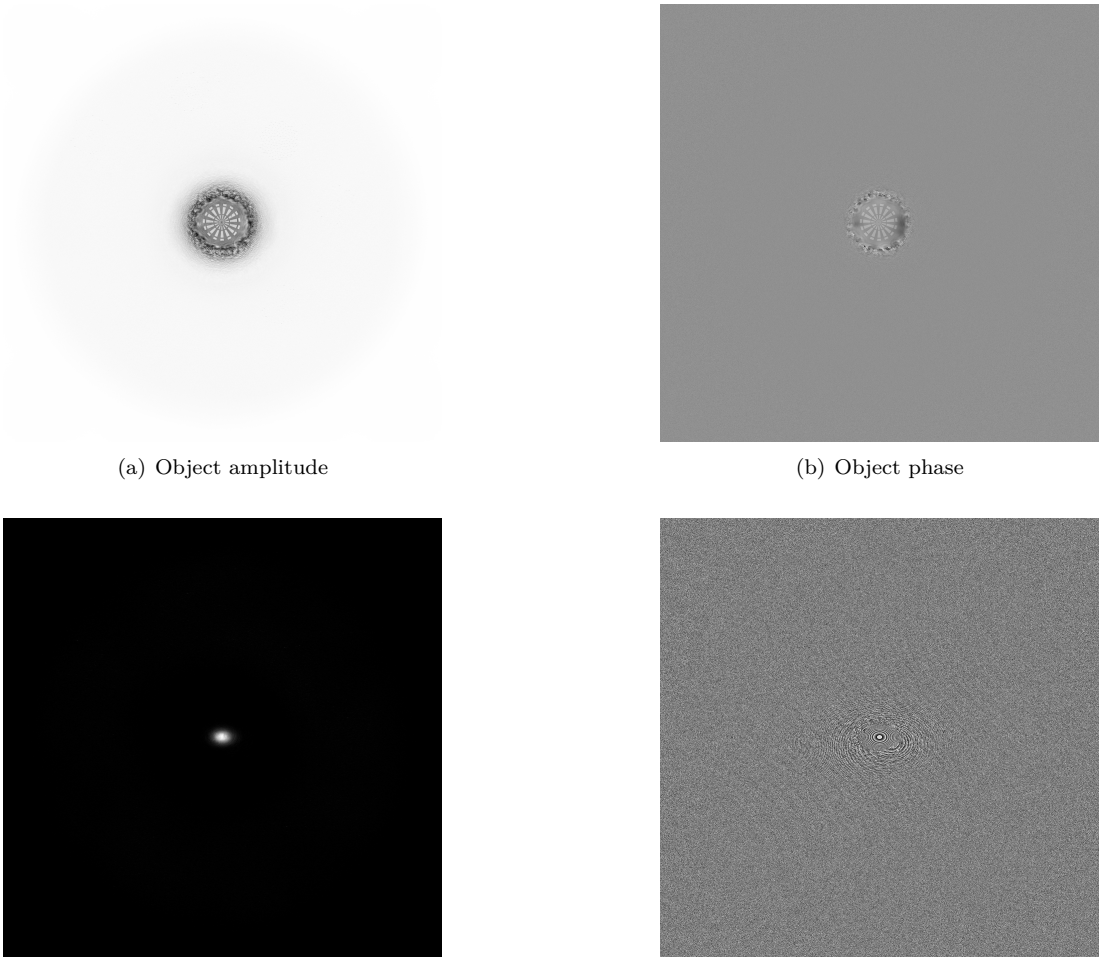

(a) Object amplitude

(b) Object phase

(c) Probe amplitude

(d) Probe phase

Figure 6.4: Converged ePIE reconstruction of the PMMA-on-silicon Siemens-star test target used to measure the system resolution. Top row: object amplitude and phase at iteration 400. Bottom row: simultaneously retrieved probe amplitude and phase.

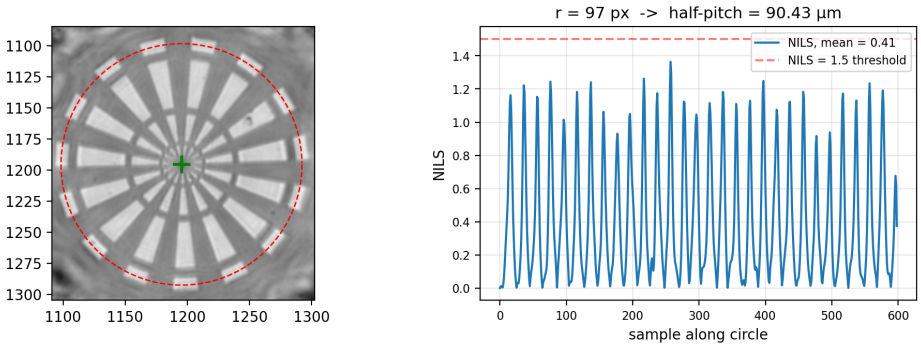

(a) hp = 97 µm

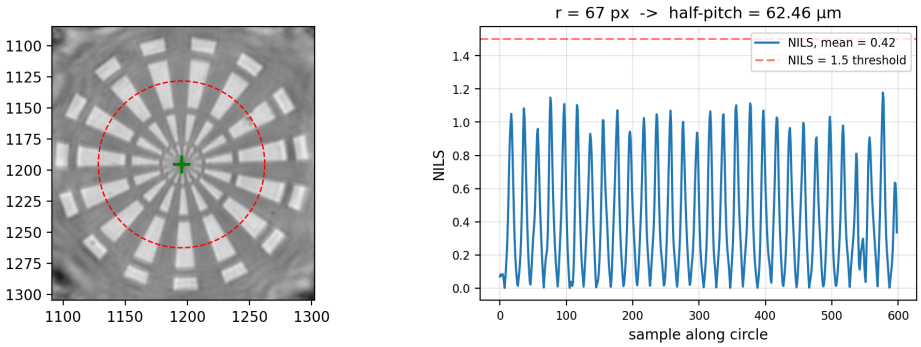

(b) hp = 67 µm

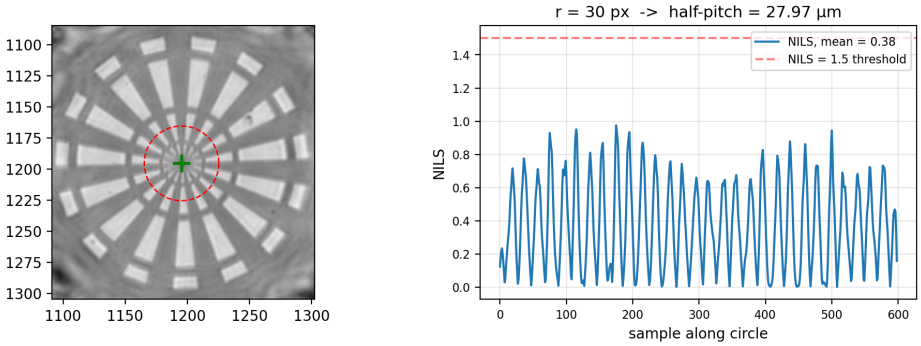

(c) hp = 30 µm

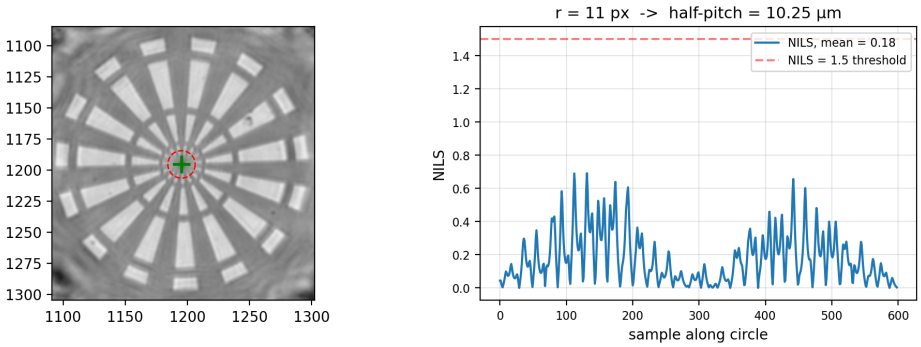

(d) hp = 11 µm

Figure 6.5: NILS analysis at four sampling radii on the same Siemens-star reconstruction. In each panel: (left) the sampling circle drawn on the object amplitude, (middle) the filtered amplitude profile $I(\theta)$ along the circle, (right) the NILS profile derived from that intensity via Eq. (6.3). Peak NILS values drop monotonically with half-pitch but the expected fifteen oscillations per turn remain visible in all five rows, in (d) one short angular sector loses contrast while the rest of the circle remains resolved.

Ptychography reconstructs a complex-valued object function. The reconstruction provides a phase image containing the phase shift accumulated by the light as it interacts with the sample. With knowledge of the wavelength and incidence angle, this phase can be converted to a surface height. The result is a height map, also known as topography. The amplitude $|O|$ encodes reflectivity or transmittance, while the phase φ encodes the optical path difference accumulated by the wavefield. This chapter derives how φ is converted to a physical surface height $h(x, y)$ and presents 3D surface models generated from the DC-BAPM reconstructions.

7.1 Phase-to-Height Conversion

7.1.1 Reflection Geometry

In the Brewster angle microscope the sample is imaged in reflection. A surface feature of height h at position (x, y) changes the path length of the reflected beam relative to the flat substrate. For a collimated beam incident at angle θ from the surface normal (fig. 7.1):

- On the incident path the beam reaches the raised surface a distance $h \cos\theta$ earlier.
- On the reflected path the beam departs the raised surface a distance $h \cos\theta$ earlier.

The total round-trip path difference is therefore

$$\Delta_{\text{refl}} = 2\, h \cos\theta. \tag{7.1}$$

A path difference Δ corresponds to a phase shift $\varphi = 2\pi\Delta/\lambda$, so the phase measured by the ptychographic reconstruction is

$$\varphi(x, y) = \frac{4\pi\, h(x, y) \cos\theta}{\lambda}. \tag{7.2}$$

Inverting for the surface height gives

$$\boxed{h(x, y) = \frac{\varphi(x, y)\, \lambda}{4\pi \cos\theta}} \tag{7.3}$$

For the DC-BAPM operating at $\lambda = 632.8\,\text{nm}$ and $\theta = 42°$:

$$h \approx \varphi \times 67.8\, \frac{\text{nm}}{\text{rad}}. \tag{7.4}$$

The factor of 4 in the denominator (rather than 2) is the hallmark of a reflection measurement: the light traverses the height difference twice (incident and reflected path), doubling the sensitivity compared to transmission.

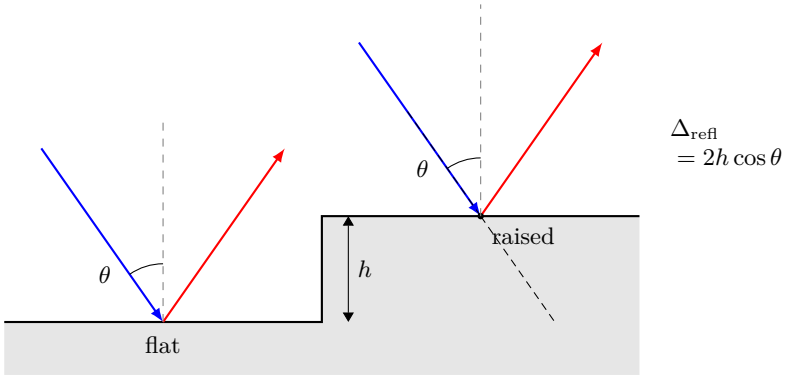

Figure 7.1: Geometry of the reflection phase-to-height conversion. Two parallel rays incident at angle θ from the surface normal bounce off a step of height h: one off the flat substrate at the left, the other off the raised surface at the right. The raised bounce shortens the optical path by $h\cos\theta$ on the incident leg and again by $h\cos\theta$ on the reflected leg, giving the total round-trip path difference $\Delta_{\text{refl}} = 2h\cos\theta$ used in Eq. (7.1).

7.1.2 Transmission Geometry

In the transmission ptychographic microscope the beam passes through the sample once at normal incidence, so $\theta = 0$ and $\cos\theta = 1$. For a sample of thickness h and refractive index n the optical path through the material is $n\,h$, while the same thickness of air contributes only h. The single-pass path difference is therefore

$$\Delta_{\text{trans}} = (n-1)\,h, \tag{7.5}$$

so the refractive-index contrast $(n-1)$ replaces the geometric $2\cos\theta$ factor of the reflection case.

The corresponding phase shift is

$$\varphi(x,y) = \frac{2\pi\,(n-1)\,h(x,y)}{\lambda}, \tag{7.6}$$

and the height is recovered as

$$\boxed{h(x,y) = \frac{\varphi(x,y)\,\lambda}{2\pi\,(n-1)}} \tag{7.7}$$

For a PMMA target at $\lambda = 632.8\,\text{nm}$ with $n \approx 1.49$:

$$h \approx \varphi \times 206\,\frac{\text{nm}}{\text{rad}}. \tag{7.8}$$

Compared to Eq. (7.4) ($67.8\,\frac{\text{nm}}{\text{rad}}$ for the reflection geometry at $\theta = 42°$), each radian of phase in transmission corresponds to roughly three times the height. Transmission is therefore noticeably less sensitive to surface topography than reflection at this wavelength.

7.2 Phase Processing

The raw phase $\varphi = \arg\big(O(\mathbf{r})\big)$ returned by ePIE requires three processing steps before it can be converted to a meaningful height map.

7.2.1 Phase Unwrapping

The arg function returns values wrapped to $(-\pi, \pi]$. If the true phase exceeds this range, 2π discontinuities appear in the map. For the thin-film samples studied in this thesis the phase excursion is typically small and unwrapping is rarely needed. When it is, a 2D quality-guided unwrapping algorithm is applied to remove the discontinuities.

7.2.2 Region of Interest

The ePIE object array is larger than the illuminated area: the edges contain uninitialised or noisy pixels. Before height conversion a square region of interest centred on the sample is selected. The size of this region determines the field of view of the resulting 3D model.

7.3 Results: Reflection Geometry

The reflection-mode DC-BAPM was used to image a Siemens star resolution target and a USAF 1951 snake grid target at $\theta = 42°$. The ePIE reconstruction converged in 400 iterations. The phase-to-height conversion was applied using equation (7.3).

7.3.1 Siemens Star Target

Figure 7.2 shows the reconstructed amplitude and phase of the Siemens star target. The amplitude image clearly resolves the radial spokes, and the phase image shows a uniform phase shift between the PMMA features and the silicon substrate.

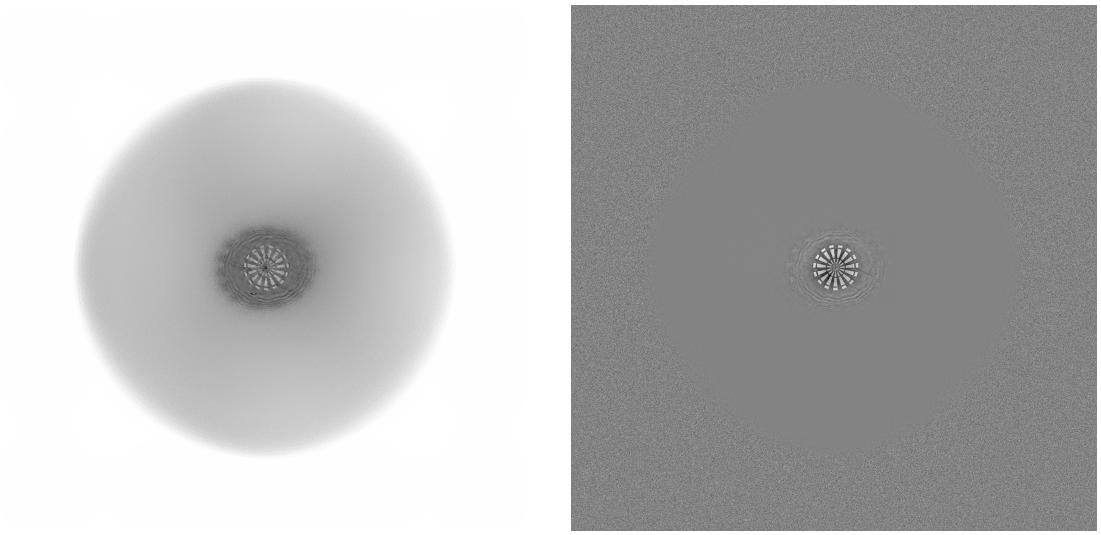

Figure 7.2: Reconstructed (a) amplitude and (b) phase of the Siemens star target in reflection at $\theta = 42°$, after 400 ePIE iterations.

Figure 7.3 presents the amplitude, phase, and computed height map side by side. The height map shows a clear distinction between the raised PMMA features and the substrate.

Figure 7.4 shows the corresponding 3D surface model. The Siemens star spokes are visible as raised plateaux on the substrate.

7.3.2 Logic Pattern Target

The same procedure was applied to a logic pattern target. Figure 7.5 shows the amplitude, phase, and height map. The horizontal and vertical bar patterns are clearly resolved in all three representations.

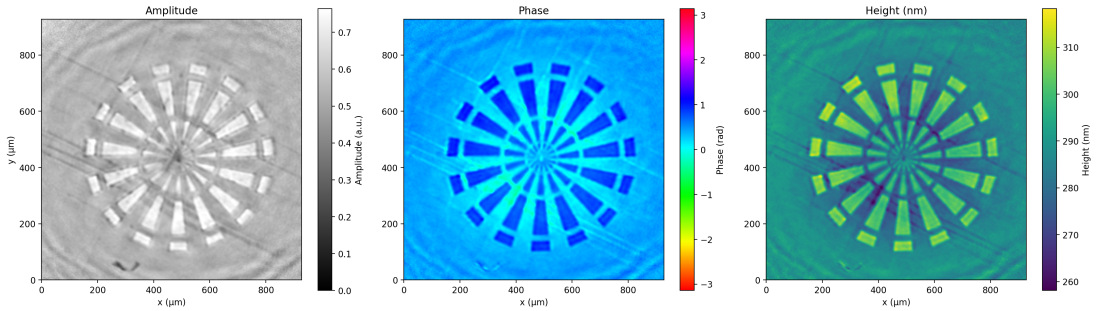

Figure 7.3: Siemens star in reflection: amplitude, unwrapped phase, and height map computed via equation (7.3). The colour bar on the right indicates the surface height in nanometres.

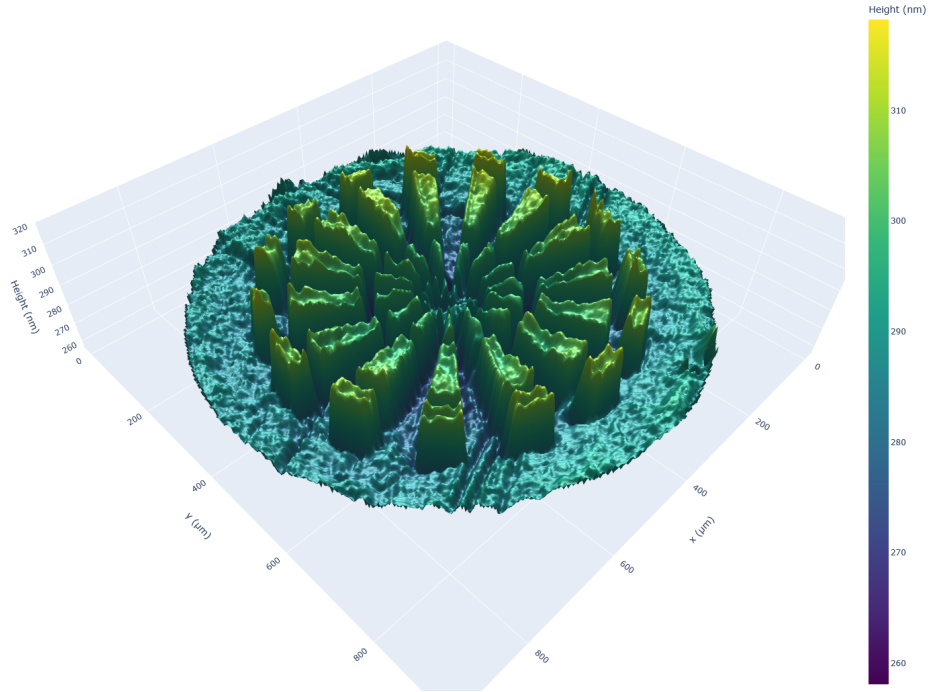

Figure 7.4: 3D surface model of the Siemens star target reconstructed in reflection geometry. The height is derived from the ptychographic phase via equation (7.3).

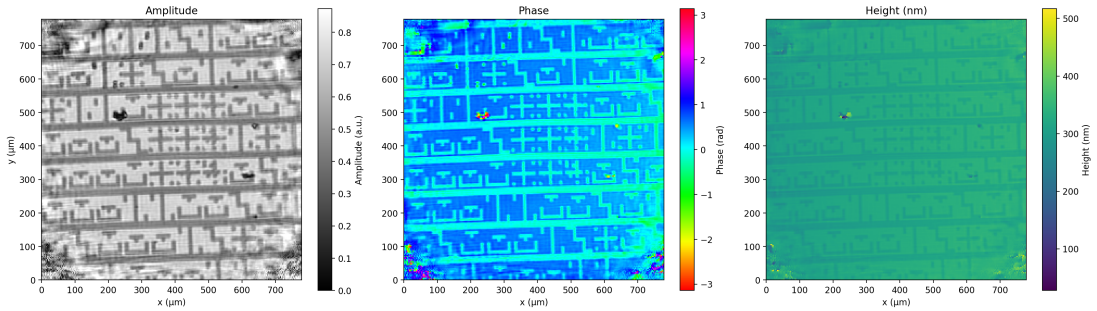

Figure 7.5: Logic pattern target in reflection: amplitude, unwrapped phase, and height map. The bar-pattern elements produce distinct height steps.

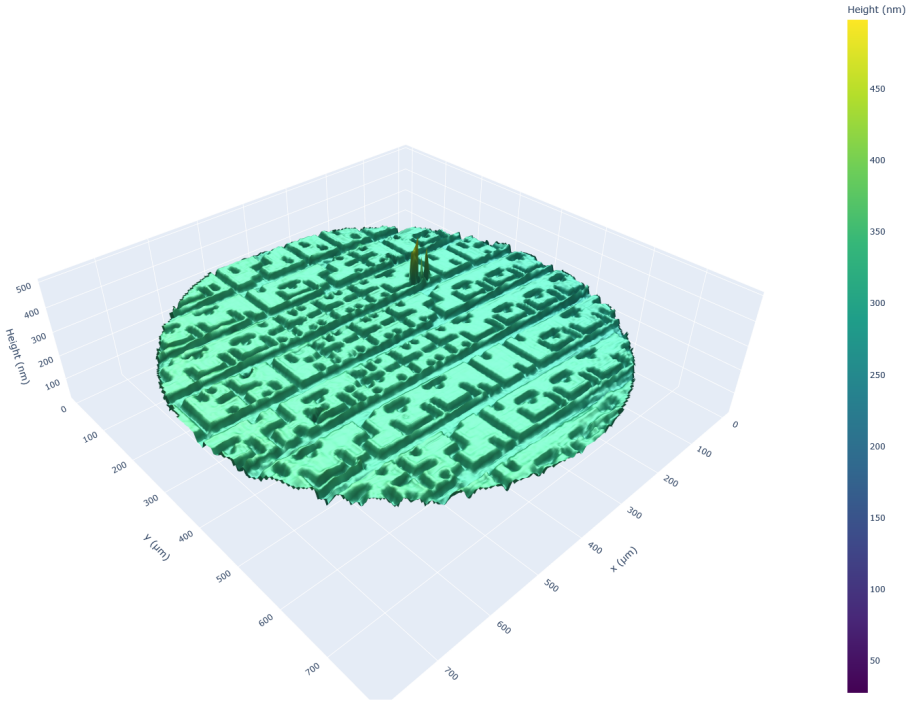

Figure 7.6: 3D surface model of the logic pattern target reconstructed in reflection geometry. The height is derived from the ptychographic phase via equation (7.3).

Graphene is a single atomic layer of carbon atoms in a hexagonal lattice, 0.335 nm thick. Figure 8.1 shows a scanning probe microscope image of a graphene layer. Its properties depend sensitively on the layer count and crystalline quality, so reliable, fast metrology of layer count is one of the practical bottlenecks for any process that uses Chemical Vapour Deposition (CVD) or exfoliated graphene.

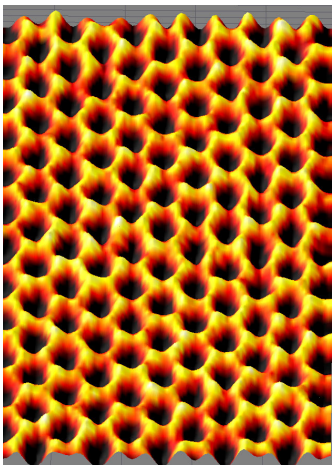

Figure 8.1: A scanning probe microscopy image of graphene. Source: https://www.flickr.com/photos/armymaterielcommand/6795812766.

Conventional optical metrology of graphene is hard. A monolayer absorbs only 2.3% of incident visible light, so the amplitude contrast against the substrate is weak. The standard workaround is to deposit the graphene on a Si/SiO_2 wafer with a carefully chosen oxide thickness (typically 90 nm or 285 nm) that produces a thin-film interference contrast strong enough to see the flake under a microscope. The phase shift introduced by the layer is small but well-defined: with a refractive index of $n \approx 2.6+1.3i$ at 632.8 nm, a 0.335 nm monolayer introduces a few millirad of phase shift on reflection. Conventional intensity-only imaging discards this phase entirely.

DC-BAPM addresses this directly. The ptychographic reconstruction gives both the amplitude and the phase of the reflected wavefield (Chapter 2), with the phase channel sensitive to changes well below one radian. The Brewster-angle geometry suppresses the bare-substrate reflection (Chapter 4), and the high-dynamic-range acquisition captures the full intensity tail (Section 3.4). Together, these allow layer-count metrology on graphene.

This chapter presents DC-BAPM reconstructions of graphene transferred onto a Si/SiO_2 substrate with an oxide thickness of 278 nm, taken on 2026-04-28 and 2026-04-29 at AOI = 42°. Preliminary scans from 2026-04-28 are collected in Appendix 11.5.

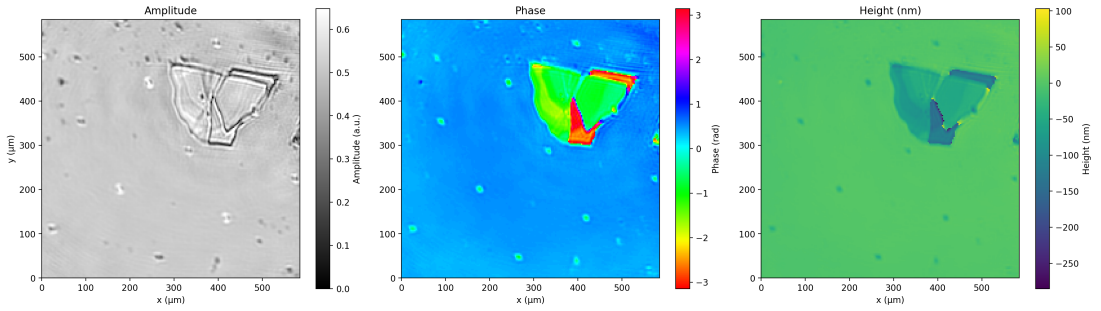

Figure 8.2: Graphene sample.

To quantify the step height across a graphene edge, a line is drawn across the reconstructed height map and the profile along that line is shown on the plot to the right. Figures 8.3 and 8.4 show two such cross-sections taken at different positions on the same flake.

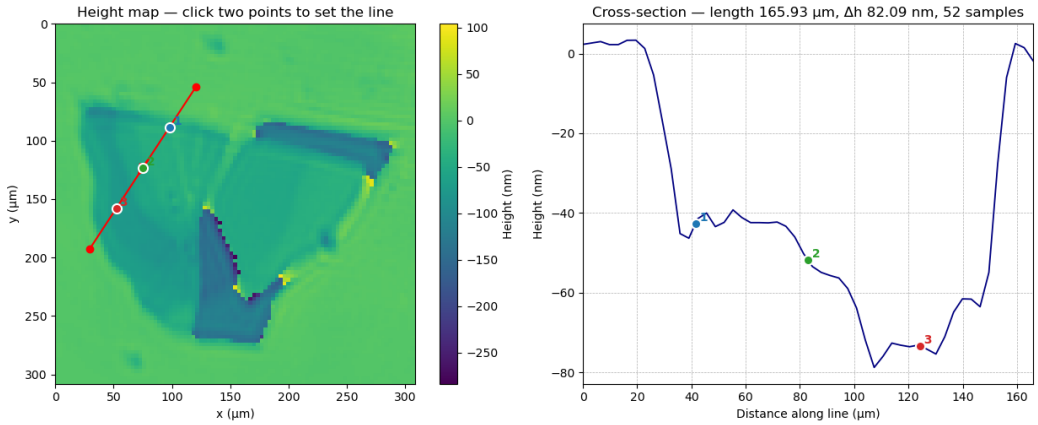

Figure 8.3: Cross-section 1: height profile along a line drawn across a graphene edge in the reconstruction of Fig. 8.2.

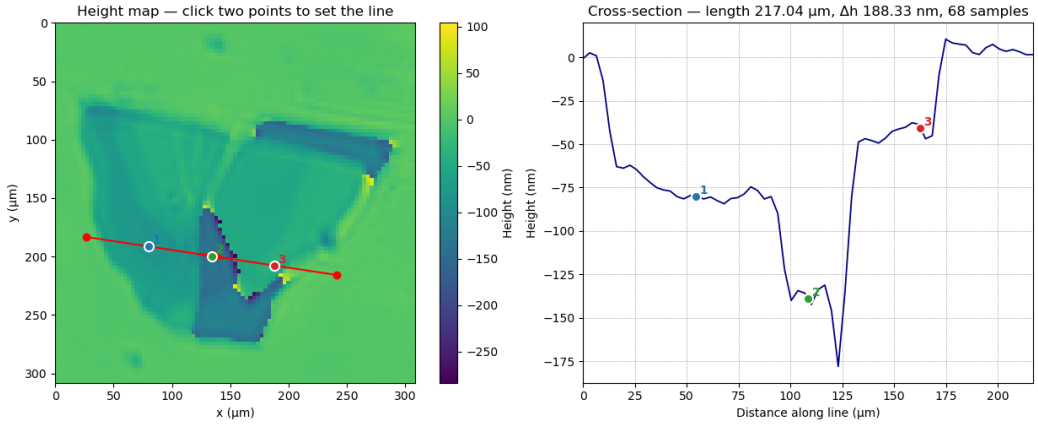

Figure 8.4: Cross-section 2: height profile along a second line drawn across the same flake.

The peak-to-valley depths read off these profiles (≈ 82.09 nm at point 3 of Fig. 8.3, ≈ 188.33 nm at point 2 of Fig. 8.4) are more than two orders of magnitude larger than the 0.34 nm thickness of a graphene monolayer. These don't represent monolayer steps.

To resolve the 0.34 nm step expected of a monolayer, the leveling has to be tightened on a known-substrate region adjacent to the flake before the profile is read. With the current two-point leveling the height channel is dominated by $\sim 10^2$ nm features that are not the graphene layer itself.

To test whether a region existed where the height channel was not dominated by this baseline drift, a second reconstruction was examined. Figure 8.5 shows a cross-section across a small feature in that reconstruction.

The peak-to-valley depth is 12 nm, equivalent to roughly 35 graphene layers (12 nm/0.34 nm). This is still far above the 0.34 nm step of a true monolayer. None of the three cross-sections resolve a monolayer step, all measured features correspond to multilayer regions or residue.

To image a genuine monolayer with DC-BAPM, the scan position must be chosen deliberately.

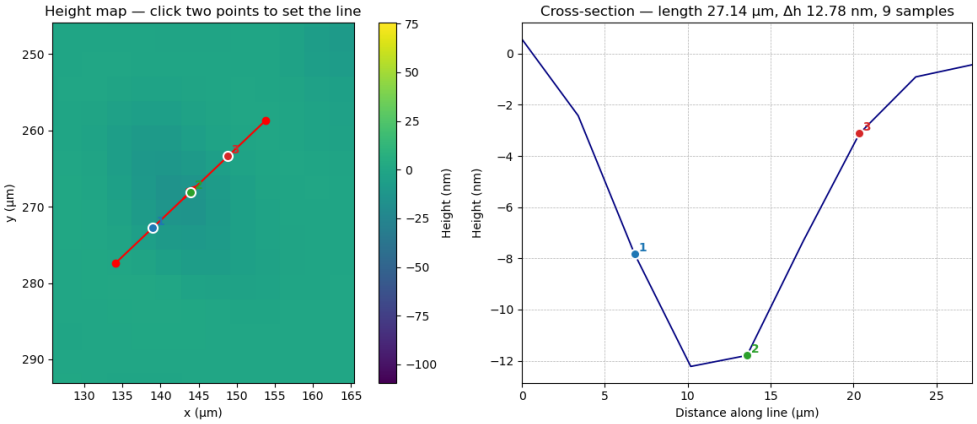

Figure 8.5: Cross-section 3: height profile along a short line drawn across a feature in the 2026-04-29 reconstruction. Peak-to-valley depth is 12 nm, an order of magnitude smaller than the cross-sections taken on the 2026-04-28 reconstruction.

A practical approach would be to first locate a confirmed monolayer region using an optical microscope, for example on a 278 nm SiO_2 substrate where interference contrast makes monolayer graphene visible, and then register that position on the DC-BAPM stage so the ptychographic scan targets the same area.

A reflection diffraction grating redistributes incident light into discrete orders at angles set by the grating equation [Thorlabs, 2026]. This chapter derives the groove density and blaze angle required for DC-BAPM distortion correction and describes the EBL fabrication attempt.

9.1 Grating Diffraction Theory

The m-th diffraction order leaves a reflection grating at angle β_m given by

$$\sin\alpha + \sin\beta_m = mN\lambda, \tag{9.1}$$

where α is the angle of incidence and N is the groove density. For the monochromatic HeNe source used in our setup, λ is constant, so dispersion is not a constraint, efficiency in the working order is discussed in Section 9.4.4.1.

9.2 Keystone Distortion in Brewster Angle Microscopy

The tilted sample plane (Chapter 4) compresses the image along the plane of incidence by a factor

$$c = \cos\theta_B, \tag{9.2}$$

and simultaneously tilts the focal plane [Zhe Sun and Baldelli, 2017]. For silicon ($\theta_B \approx 75.4°$) this is 75%, for PMMA ($\theta_B = 56.1°$) it is 44%. Both effects must be corrected before ptychographic reconstruction.

9.3 Grating-Based Distortion Correction

9.3.1 Operating Principle

A reflection grating placed at the intermediate image plane redirects each image ray by an angle that depends on its position along the dispersion direction, stretching the compressed axis by $1/\cos\theta_B$ [Zhe Sun and Baldelli, 2017]. Because the grating surface coincides with the intermediate image, all object points focus there regardless of field position, eliminating the focus tilt. The relay lens between the sample and the grating focuses the diverging diffraction pattern onto the grating surface, ensuring that high-spatial-frequency rays do not fall outside the grating aperture before being redirected toward the detector. The numerical aperture of the relay lens therefore limits the maximum spatial frequency collected, making it a key design parameter. The groove density N that achieves the distortion correction is derived in Section 9.4.2. Residual higher-order geometric distortions are characterised in Chapter 5.

9.3.2 DC-BAPM Optical Setup

Figure 9.1 shows the DC-BAPM optical layout with the distortion-correcting grating in place. The grating sits at the intermediate image plane between the sample and the camera, with

its grooves perpendicular to the plane of incidence. A relay lens between the sample and the grating focuses the diverging diffraction pattern reflected from the sample onto the grating surface, the grating then redirects the first-order diffracted light along its surface normal into the CCD. The first-order diffracted beam exits the grating along the surface normal ($\beta_1 = 0°$, see Section 9.4.2), and the detector surface is mounted perpendicular to that direction. This places the camera and the grating in parallel planes, eliminating both the keystone compression and the focus tilt of the Brewster-angle imaging path.

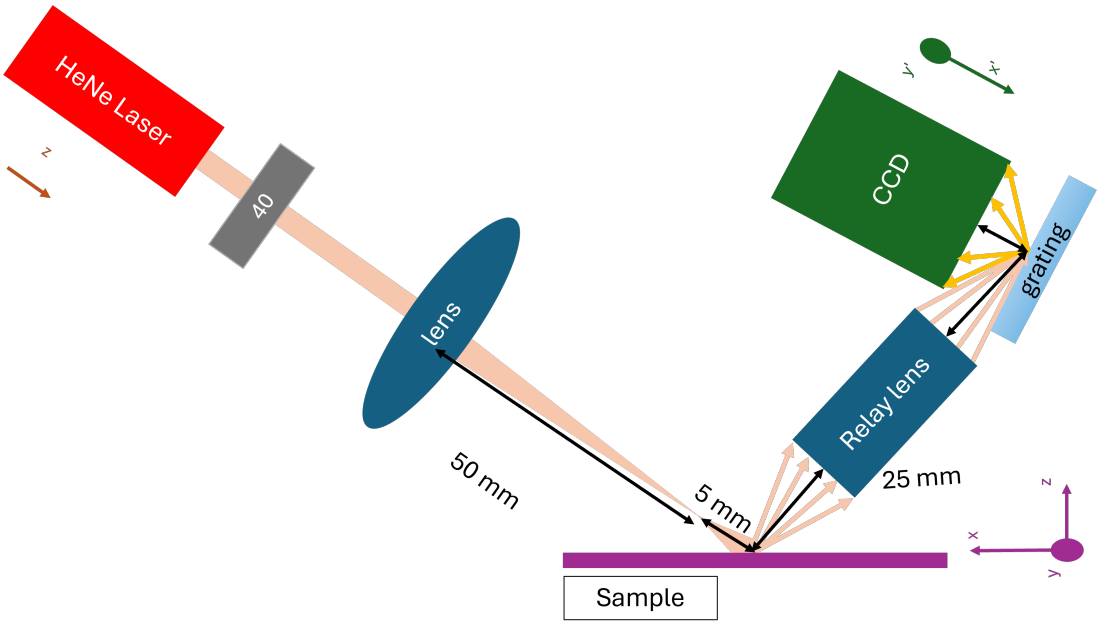

Figure 9.1: Optical layout of the DC-BAPM with the distortion-correcting reflection grating. The grating is placed at the intermediate image plane with its grooves perpendicular to the plane of incidence. The first-order diffracted beam exits normal to the grating surface and is collected by the camera, which lies parallel to the grating. The full optical-path schematic with all labelled elements is given in Fig. 3.1.

9.4 Material-Specific Grating Design for 632.8 nm Illumination

The groove density N and blaze angle ϕ_b of the correcting grating are functions of the substrate Brewster angle, which is in turn set by the substrate refractive index at the laser wavelength. This section derives the complete set of grating parameters for five materials relevant to Brewster angle microscopy: water, PMMA, carbon, silicon, and aluminium, all at the HeNe laser wavelength $\lambda = 632.8\,\text{nm}$, with air ($n_1 = 1.000$) as the incident medium. The design condition follows the approach of Sun *et al.* [Zhe Sun and Baldelli, 2017].

9.4.1 Refractive Indices at 632.8 nm

The complex refractive index $\tilde{n} = n + \mathrm{i}k$ for each material was obtained from the refractiveindex.info optical-constants database, the specific datasets and direct URLs are noted in Table 9.1.

For transparent dielectrics (water, PMMA, carbon with $k \approx 0$), the p-polarisation reflectance vanishes exactly at the Brewster angle

$$\theta_B = \arctan\left(\frac{n_2}{n_1}\right), \tag{9.3}$$

so the substrate appears completely dark at the bare interface. For absorbing media ($k > 0$: Si, Al), equation (9.3) no longer applies, instead a *pseudo-Brewster angle* θ_{pB} exists at which the p-reflectance reaches a minimum (but not zero). For silicon, where $k \ll n$ ($n = 3.847$, $k = 0.016$), the approximation $\theta_{\mathrm{pB}} \approx \arctan(n)$ is accurate to within $0.1°$. For aluminium, which has a large imaginary part ($n = 1.20$, $k = 6.92$), θ_{pB} must be found by numerical minimisation of $|r_p(\theta)|^2$ and is approximately $80°$, the minimum $R_p \approx 70\%$, so aluminium substrates give intrinsically low BAM contrast.

Table 9.1: Complex refractive index and Brewster (or pseudo-Brewster) angle for five substrates at $\lambda = 632.8\,\mathrm{nm}$. Sources: Water — Gupta *et al.* 2022 (https://refractiveindex.info/?shelf=other&book=H2O-C3H5_OH_3&page=Gupta-0), PMMA — Beadie *et al.* 2015 (https://refractiveindex.info/?shelf=organic&book=poly_methyl_methacrylate&page=Beadie), Carbon — Peter 1923 (https://refractiveindex.info/?shelf=main&book=C&page=Peter), Si — Aspnes & Studna 1983 (https://refractiveindex.info/?shelf=main&book=Si&page=Aspnes), Al — Rakić 1995 (https://refractiveindex.info/?shelf=main&book=Al&page=Rakic).

Material	n	k	θ_B (°)	$\cos\theta_B$	$1/\cos\theta_B$
Water (H_2O)	1.3309	0	53.07	0.6012	1.663
PMMA	1.4909	0	56.12	0.5581	1.792
Carbon (C)	2.42	0	67.5	0.3820	2.618
Silicon (Si)	3.847	0.016	75.4	0.2514	3.978
Aluminium (Al)	1.20	6.92	≈ 80	0.174	≈ 5.75

Figure 9.2 shows the Fresnel reflectance of PMMA as a function of angle of incidence. The p-polarisation reflectance R_p drops to zero at $\theta_B = 56.1°$, while the s-polarisation reflectance R_s remains comparatively high throughout. At the experimental operating angle of $42°$ the difference between R_s and R_p is already substantial: R_p is significantly suppressed while R_s is still rising, so the polarisation selectivity that underlies BAM contrast remains effective at this angle. This makes $\alpha = 42°$ a viable operating point even though it does not coincide with the Brewster angle.

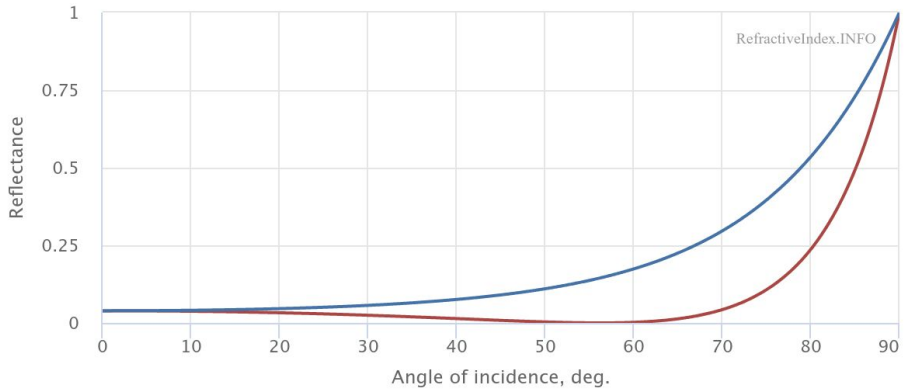

Figure 9.2: Fresnel reflectance of PMMA ($n = 1.491$, $k = 0$) at $\lambda = 632.8\,\mathrm{nm}$. The p-polarisation reflectance vanishes at the Brewster angle $\theta_B = 56.1°$, the s-polarisation reflectance remains high across the full angular range. At the experimental operating angle of $42°$ the s-p contrast is already significant, confirming that the Brewster-angle polarisation selectivity is operative at this angle.

9.4.2 The DC-BAM Design Condition

The key design insight of Sun *et al.* [Zhe Sun and Baldelli, 2017] is that two conditions must be satisfied simultaneously:

1. The incidence angle on the grating equals the substrate Brewster angle: $\alpha = \theta_B$.

2. The first-order diffracted beam exits *perpendicular* to the grating surface: $\beta_1 = 0°$.

Condition 2 ensures that the entire image carried by the first-order beam travels along the grating normal, so the collection objective (aligned with the grating normal) images the flat grating surface onto the flat camera without any focus tilt and without keystone distortion. The relay lens between sample and grating maps the tilted sample surface to the grating surface at 1:1 scale *along the surface*, the grating then redirects all image rays perpendicular to its surface, and the collection objective images this flat, correctly-scaled surface undistorted onto the camera.

Substituting $\alpha = \theta_B$ and $\beta_1 = 0°$ into the grating equation (equation 9.1) with $m = 1$:

$$\sin\theta_B + \sin 0° = N\lambda \quad \Longrightarrow \quad \boxed{N = \frac{\sin\theta_B}{\lambda}}. \tag{9.4}$$

This is the central design formula: the required groove density is determined solely by the substrate Brewster angle and the laser wavelength.

Cross-check. For $\lambda = 532\,\text{nm}$ on water ($\theta_B = 53.06°$), equation (9.4) gives $N = 1503$ l/mm, consistent with the $N = 1503$ l/mm grating calculated by Sun *et al.* [Zhe Sun and Baldelli, 2017].

Blaze condition. Maximum first-order efficiency is achieved when the groove facet specularly reflects the incident beam into the first order. For $\alpha = \theta_B$ and $\beta_1 = 0°$ the blaze angle is the bisector of the incident and diffracted chief rays:

$$\phi_b = \frac{\alpha + \beta_1}{2} = \frac{\theta_B + 0°}{2} = \frac{\theta_B}{2}. \tag{9.5}$$

This is the *Littrow condition* for this geometry: the blaze facet normal bisects the angle between the illumination and the collected diffraction order, maximising the diffraction efficiency at the operating wavelength.

9.4.3 Calculated Grating Specifications

Table 9.2 gives the groove density and blaze angle computed from equations (9.4) and (9.5) for the five substrates at $\lambda = 632.8\,\text{nm}$.

Table 9.2: Grating specifications for the DC-BAM at $\lambda = 632.8\,\text{nm}$ using the design condition $\alpha = \theta_B$, $\beta_1 = 0°$ (Sun *et al.* [Zhe Sun and Baldelli, 2017]). $N = \sin\theta_B/\lambda$, $d = 1/N$: groove period, $\phi_b = \theta_B/2$: Littrow blaze angle. The column N_{532} shows the equivalent groove density for a 532 nm laser, for reference.

Material	θ_B (°)	$\sin\theta_B$	N (l/mm)	d (nm)	ϕ_b (°)	N_{532} (l/mm)
PMMA (exp. $\alpha = 42°$)	42.0	0.669	1057	946	21.0	1258
Water	53.1	0.800	1264	791	26.5	1503
PMMA	56.1	0.831	1313	762	28.1	1561
Carbon	67.5	0.924	1461	685	33.8	1737
Silicon	75.4	0.968	1530	654	37.7	1819
Aluminium	≈80	0.985	≈1557	≈643	≈40	≈1851

For the primary PMMA substrate, $N = 1313$ l/mm ($d = 762\,\text{nm}$) is fabricated by EBL to avoid rounding to a catalogue value.

(a) $m = 0, +1, +2$

(b) $m = -1, 0, +1$

Figure 9.3: Photographs of the diffraction orders produced by the grating under white-light illumination. The zeroth order ($m = 0$) is a specular reflection of the white light, while the higher orders ($m = \pm 1, +2$) disperse the spectrum into its constituent wavelengths.

9.4.4 Binary Grating Fabrication Design

For the air–PMMA interface ($\theta_B = 56.1°$, $N = 1313$ l/mm, $d = 762\,\text{nm}$), a binary (single-level) reflection grating was designed for electron-beam lithography (EBL) fabrication. The grating consists of a silicon substrate with patterned PMMA features, where the PMMA lands (binary = 1) have a height equal to the etch depth and the grooves (binary = 0) expose the bare silicon surface. The duty cycle and depth are chosen to maximise the first-order diffraction efficiency as derived below.

9.4.4.1 Duty Cycle and Depth Optimisation

The grating equation determines *where* energy is diffracted, but not *how much*. The efficiency in a given order depends on the groove profile. For the DC-BAPM application, maximum efficiency in the working order is essential because the reflected signal from a thin monolayer at the Brewster angle is intrinsically weak, typically only a fraction of a percent of the incident intensity.

For a binary (rectangular-profile) phase grating used in reflection, the diffraction efficiency of the m-th order is given by [Meshalkin et al., 2019]

$$\eta_m = \left(\frac{2}{m\pi}\right)^2 \sin^2(m\pi f)\,\sin^2\left(\frac{\varphi}{2}\right), \tag{9.6}$$

where f is the duty cycle (ratio of groove width to period) and φ is the phase depth accumulated by the light reflecting from the ridge relative to the groove.

The first factor, $\sin^2(m\pi f)$, depends only on the duty cycle. For the first order ($m = 1$) it is maximised when $f = 0.5$, i.e. equal groove and land widths. A 50% duty cycle also suppresses all even diffraction orders, since $\sin(m\pi \cdot 0.5) = 0$ for every even m.

The second factor, $\sin^2(\varphi/2)$, depends on the groove depth h. For a binary grating with PMMA ridges (n_{PMMA}) on a reflective silicon substrate in air, the round-trip phase difference between light reflecting off the top of a ridge and light reflecting off the bare substrate is

$$\varphi = \frac{4\pi\, h\,(n_{\text{PMMA}} - 1)}{\lambda}. \tag{9.7}$$

This is maximised when $\varphi = \pi$, giving the optimal etch depth

$$h = \frac{\lambda}{4\,(n_{\mathrm{PMMA}} - 1)}. \tag{9.8}$$

Substituting $\lambda = 632.8\,\mathrm{nm}$ and $n_{\mathrm{PMMA}} = 1.491$:

$$h = \frac{632.8}{4 \times 0.491} \approx 322\,\mathrm{nm}. \tag{9.9}$$

With $f = 0.5$ and $\varphi = \pi$, the theoretical maximum first-order efficiency from equation (9.6) is

$$\eta_1 = \left(\frac{2}{\pi}\right)^2 = \frac{4}{\pi^2} \approx 40.5\%. \tag{9.10}$$

This is the fundamental upper limit for a binary phase grating, higher efficiencies require a blazed (sawtooth) or multi-level staircase profile. Because the grating period ($d = 762\,\mathrm{nm}$) is comparable to the wavelength, scalar diffraction theory is only approximate and the actual optimum may shift slightly. Rigorous coupled-wave analysis (RCWA) should be used to verify the efficiency for the specific geometry and polarisation state of the DC-BAM.

9.4.4.2 Cross-Section and Layout

Figure 9.4 shows a two-dimensional cross-section of the grating profile with the incident beam at $\alpha = \theta_B = 56.1°$ and the first-order diffracted beam exiting perpendicular to the surface ($\beta = 0°$), confirming the DC-BAM design condition. The groove and land widths are each 381 nm (50% duty cycle), and the PMMA ridge height is $h = 323\,\mathrm{nm}$.

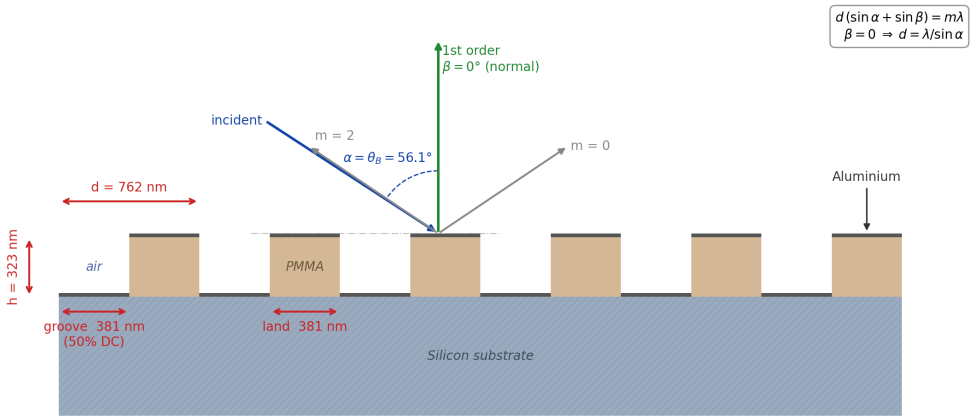

Figure 9.4: Cross-section of the binary PMMA-on-silicon diffraction grating ($N = 1313$ l/mm, $d = 762\,\mathrm{nm}$, DC = 50%, $h = 322\,\mathrm{nm}$). The incident beam arrives at $\alpha = \theta_B = 56.1°$ (from the normal) and the first-order diffracted beam exits along the surface normal ($\beta = 0°$), satisfying the DC-BAM design condition (equation 9.4).

9.5 Grating Manufacturing

The binary PMMA-on-silicon grating described in Section 9.4.4 is fabricated using electron-beam lithography (EBL). The process begins with spin coating a thin layer of PMMA resist onto a silicon substrate, followed by EBL exposure and development to define the grating pattern, and a final metrology step to verify the groove dimensions.

The full GDS layout, including tolerance tiles at $\pm 1°$ around the design angle, is reproduced in Appendix 11.5.

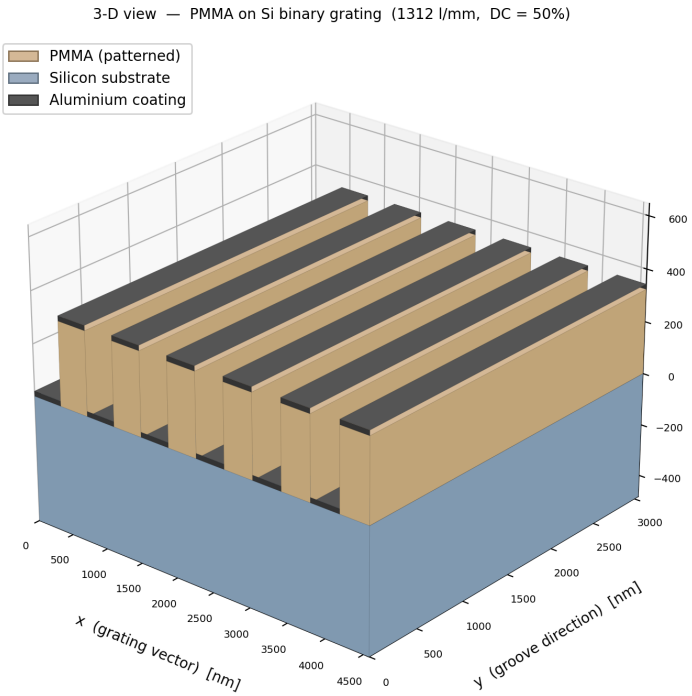

Figure 9.5: Three-dimensional perspective view of the binary grating structure. The silicon substrate (blue-grey) supports patterned PMMA features (tan) at the land positions. The grooves expose the bare silicon surface.

9.5.1 Spin Coating

Spin coating deposits a uniform thin film of PMMA onto the silicon substrate. A small volume of PMMA dissolved in anisole is dispensed onto the stationary substrate, which is then accelerated to a set spin speed. The centrifugal force spreads the solution radially outward while the solvent evaporates, leaving a solid film whose thickness depends primarily on the spin speed and the PMMA concentration. A higher spin speed produces a thinner film. After spinning, the substrate is baked on a hotplate to drive off residual solvent and densify the resist. The target film thickness is $h = 322\,\text{nm}$, matching the optimal etch depth derived in equation (9.8).

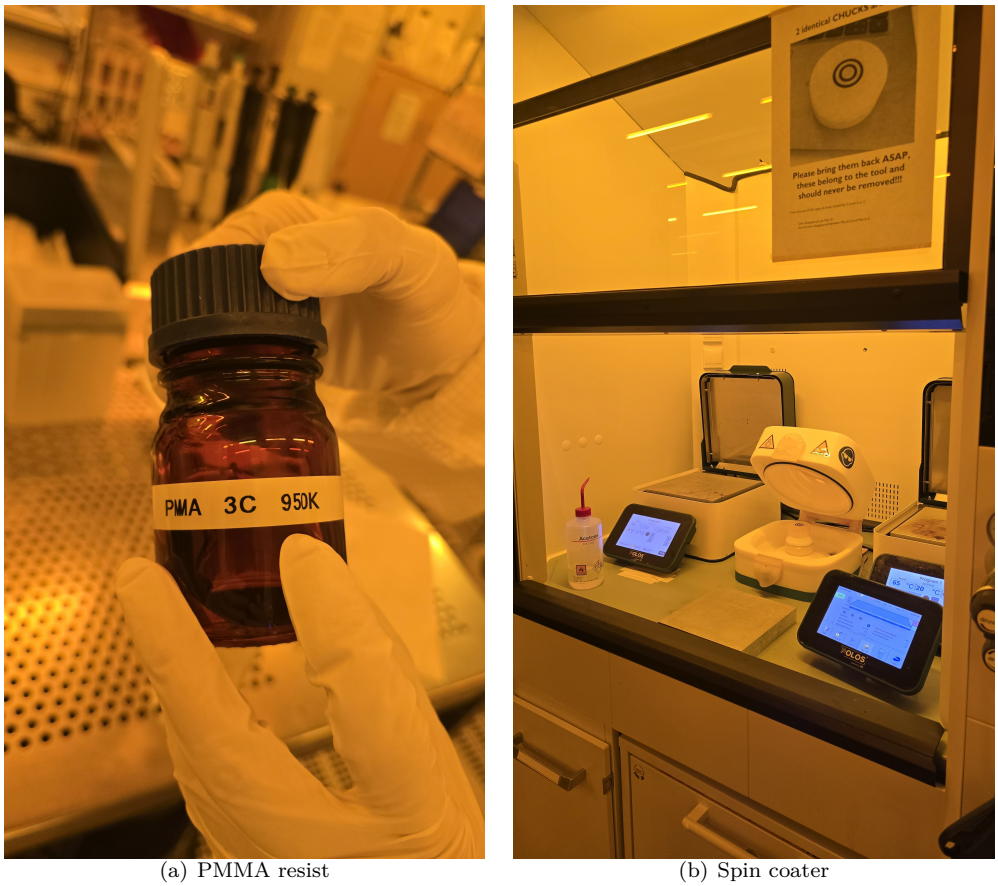

(a) PMMA resist (b) Spin coater

Figure 9.6: (a) The PMMA resist used for spin coating. (b) The spin coater used to deposit the PMMA layer onto the silicon substrate.

9.5.2 Development

After electron-beam exposure, the sample is developed in a 1:3 solution of methyl isobutyl ketone (MIBK) and isopropanol (IPA). MIBK selectively dissolves the exposed PMMA regions where the electron beam has broken the polymer chains. IPA acts as a rinse that stops the development and removes residual developer from the surface. Development in the 1:3 solution was slow, requiring several measurement cycles with my colleague Jiwook to assess progress and continue etching until the binary pattern became visible. Unfortunately, during one inspection with the digital microscope at 6000× magnification, the objective crashed into the sample during autofocus and destroyed it.

A second sample was prepared by coating PMMA on a sapphire substrate. To reduce the long development time, a 1:1 MIBK:IPA ratio was used instead. While this accelerated the development significantly, the solution was too aggressive and introduced rough feature edges. The resulting roughness likely obliterated the binary grating lines, leaving an effectively flat surface.

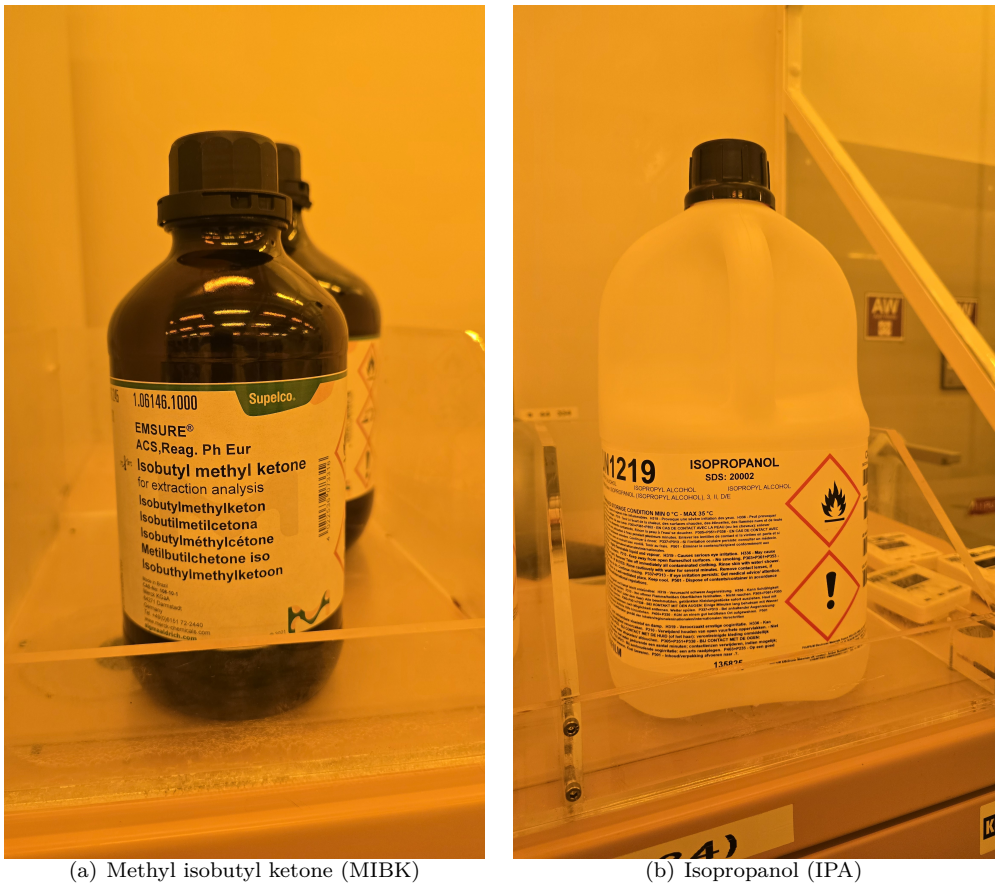

(a) Methyl isobutyl ketone (MIBK) (b) Isopropanol (IPA)

Figure 9.7: (a) Methyl isobutyl ketone (MIBK) developer used in a 1:3 ratio with IPA. (b) Isopropanol (IPA) rinse used to stop the development.

9.5.3 Metrology

After development each sample is characterised by two complementary instruments: a stylus profilometer measures the cross-sectional surface height along a line, and a digital microscope provides a top-view image of the patterned area at up to 6000× magnification (Fig. 30, Appendix 11.5). The profilometer drags a fine tungsten needle across the sample and records the height as a function of position (Fig. 31), it confirms whether the nominal etch depth of 322 nm has been achieved and reveals any film remaining in the grooves. The digital microscope is used first as a quick pass/fail check: if the binary grating lines are not visible at moderate magnification, something has gone wrong upstream and further measurement is unnecessary.

Both the first sample (1:3 MIBK:IPA, silicon substrate) and the second sample (1:1 MIBK:IPA, sapphire substrate) failed this check. The digital microscope image in Fig. 9.9 shows the patterned area at 500× magnification: the feature edges are clearly malformed, indicating that the binary grating lines were not faithfully resolved during EBL or development. The profilometer confirmed that no well-defined steps of the target depth were present.

An atomic force microscope (AFM) is available in the cleanroom (Fig. 30(a)) and can resolve individual grating lines with nanometre height resolution. For a successful sample, AFM measurement of a short section of the grating would provide the most direct verification of the groove period, duty cycle, and etch depth. For the samples produced here, the failure was already evident from the optical inspection, so AFM measurement of the grating region was not performed. In a future fabrication attempt, AFM imaging of the binary pattern area (rather

than a line-scan cross-section) should be used as the primary metrology step to catch subtle fabrication defects that a profilometer cross-section might miss.

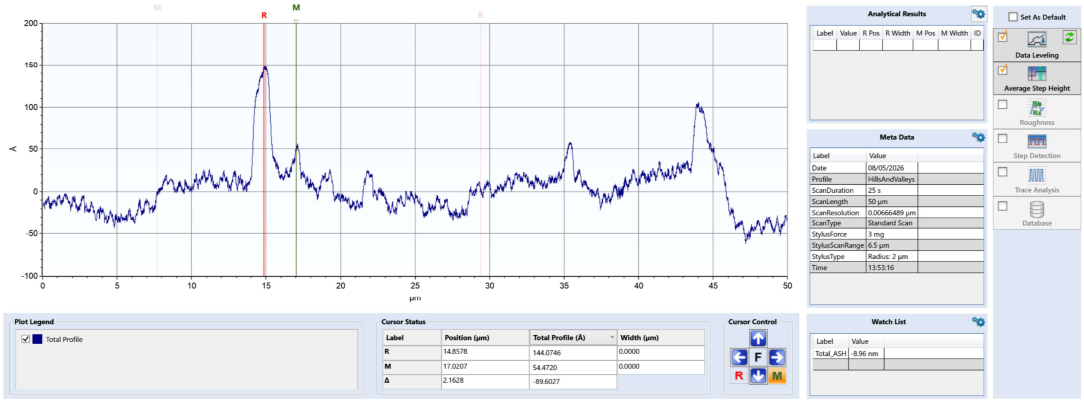

Figure 9.8: Profilometer measurement taken on the grating area of the first sample (1:3 MIBK:IPA, silicon substrate). No well-defined step of the target depth (322 nm) is present, the surface height varies irregularly, confirming that the binary grating lines were not resolved during development.

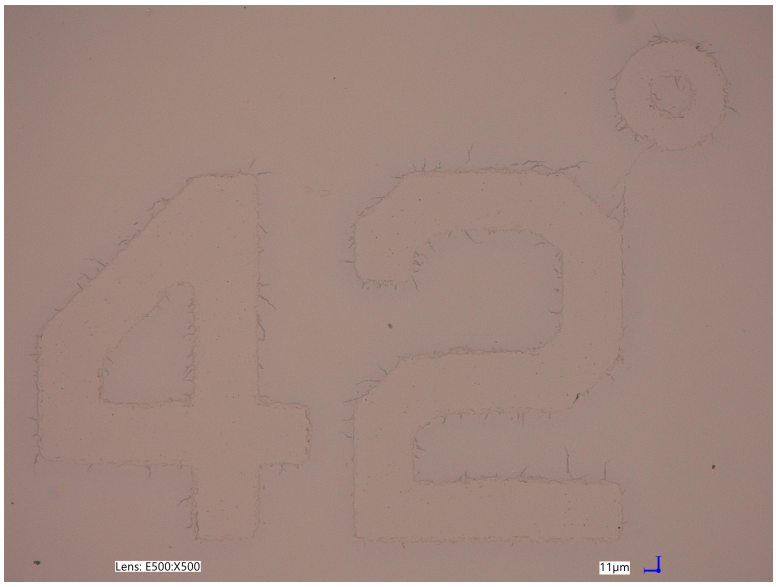

Figure 9.9: 500 times magnification digital microscope image of the number reprenting the binary grating pattern for an AOI of 42 degrees. The edges of the patterned numbers are clearly malformed.

9.6 Summary

The design condition $N = \sin\theta_B/\lambda$ yields a binary PMMA-on-silicon phase grating with $N = 1313$ l/mm, 50% duty cycle, and $h = 322\,\text{nm}$ etch depth for the primary PMMA substrate. EBL fabrication was attempted on two samples, both failed at the development step. Further manufacturing iterations and optical testing are required to validate the distortion correction in a complete DC-BAPM measurement.

This thesis has described the design, construction, and characterisation of a distortion-corrected Brewster angle ptychographic microscope (DC-BAPM) for phase-sensitive imaging of thin films at the air-PMMA interface.

10.1 Summary of Findings

The main findings of this work are as follows.

10.1.1 Ptychographic Phase Retrieval at Brewster Incidence

The ePIE algorithm was successfully applied to diffraction data collected at oblique incidence using a software-based conical distortion correction to compensate for the Ewald-sphere distortion. The algorithm simultaneously retrieved the complex object function $O(\mathbf{r})$ and the probe $P(\mathbf{r})$ without prior knowledge of the probe, converging to a normalised error below 10^{-2} within 200 epochs for all datasets studied. The retrieved probe showed the structured phase pattern expected from the illumination geometry, confirming the consistency of the reconstruction.

10.1.2 Polarisation-Dependent Contrast in DC-BAPM Near the Brewster Condition

The polarisation sensitivity of the DC-BAPM was verified by scanning four areas of a PMMA-on-silicon resolution target at three polarisation-axis settings (PA = 0°, 45°, 90°) at AOI = 42°. P-polarised illumination consistently yielded the strongest contrast between patterned features and the substrate, while s-polarised illumination collapsed the relative contrast, consistent with the partial suppression of R_p on the approach to the Brewster minimum. The polarisation effect was predominantly visible in the phase image rather than the amplitude, because $\partial \arg r_p / \partial d_f$ is already steep at 42° (Section 4.5).

10.1.3 Diffraction Grating Design for Optical Keystone Correction at Brewster Incidence

The theoretical design of a reflection diffraction grating for correcting the keystone compression inherent to the oblique-incidence geometry was completed. The design condition $N = \sin\theta_B/\lambda$ yields $N = 1313\,\mathrm{l/mm}$ for the air-PMMA Brewster angle ($\theta_B = 56.2°$, $\lambda = 632.8\,\mathrm{nm}$), with a binary profile of 50% duty cycle and etch depth $h = 322\,\mathrm{nm}$. Two fabrication attempts by electron-beam lithography were unsuccessful due to resist development issues. The design is analytically validated, successful fabrication and optical integration remain as the next step.

10.2 Answer to the Research Question

The central research question of this thesis was:

> *Can ptychographic phase retrieval be successfully applied to a Brewster angle microscope geometry to recover quantitative, diffraction-limited phase images of nanometre-thick films on a silicon substrate, and can the keystone distortion inherent to that geometry be corrected without compromising the ptychographic reconstruction?*

The answer is largely affirmative. The DC-BAPM system demonstrated that:

1. A conical distortion correction based on Gardner et al. Gardner, B. Zhang, et al., 2012 compensates for the Ewald-sphere distortion, enabling consistent ptychographic reconstruction from oblique-incidence diffraction data. The physical diffraction grating that would perform this correction optically has been designed but requires successful fabrication before this step can be eliminated from the post-processing pipeline.
2. ePIE converges reliably for diffraction data collected at Brewster incidence, provided the probe overlap criterion ($\geq 80\%$) and the asymmetric step sizes dictated by the elliptical footprint are respected.
3. Polarisation-sensitive phase contrast is recovered at 42° incidence, with p-polarised illumination giving the strongest feature contrast consistent with the Fresnel reflectance at that angle.

10.3 Significance of the Work

The DC-BAPM approach bridges two previously separate imaging modalities: the surface-sensitive, label-free contrast of Brewster angle microscopy, and the quantitative phase information and diffraction-limited resolution of ptychography. This combination opens new possibilities for characterising the lateral structure and optical properties of thin films on semiconductor substrates, including graphene monolayers, metal coatings, and polymer films, without fluorescent labels, and at spatial scales inaccessible to conventional BAM.

The DC-BAPM system demonstrated in this thesis represents a proof of concept. Several directions for extension and improvement are outlined below.

11.1 Multi-Angle Phase Retrieval

The current system operates at a single angle of incidence equal to the Brewster angle. By acquiring ptychographic datasets at a series of angles near θ_B, it would be possible to extract both the thickness d_f and the refractive index n_f of the film independently, analogous to reflectometry [Jørgensen et al., 2024; Tanksalvala et al., 2021]. This would transform the DC-BAPM from a qualitative phase-contrast instrument into a quantitative optical ellipsometer with diffraction-limited lateral resolution.

11.2 Real-Time Reconstruction

The ePIE reconstruction is already GPU-accelerated, but iterative phase retrieval is still too slow for real-time imaging. Babu et al. demonstrated streaming ptychographic reconstruction at 2 kHz using a deep-learning model running on HPC hardware [Babu et al., 2023]. Adopting this approach in the DC-BAPM would enable continuous monitoring of dynamic thin-film processes such as aluminium oxidation or PMMA thermal flow.

11.3 Depolarisation Correction

The ePIE reconstruction implicitly assumes that the illumination on the sample is purely p-polarised and that the analyser transmits only the coherent reflected component. In practice, imperfect optics in the beam path introduce a small fraction of depolarised stray light, and the sample itself can scatter light into the orthogonal polarisation state. This incoherent background reduces the effective extinction at the Brewster angle and adds a constant floor to the diffraction patterns that the forward model does not account for.

The depolarised fraction can be measured directly using linear Stokes polarimetry. With the generator fixed at 0° and the analyser rotated through 0°, 45°, 90°, and 135°, the three linear Stokes parameters are

$$S_0 = I(0^\circ) + I(90^\circ), \qquad S_1 = I(0^\circ) - I(90^\circ), \qquad S_2 = I(45^\circ) - I(135^\circ), \tag{11.1}$$

and the degree of linear polarisation follows as

$$\mathrm{DOLP} = \frac{\sqrt{S_1^2 + S_2^2}}{S_0}. \tag{11.2}$$

A baseline characterisation of the bare optical path (no sample present) gives DOLP = 0.995 with an extinction ratio of 136 : 1, consistent with a well-aligned linear polariser pair (fig. 11.1). Repeating the same four-image measurement with a sample inserted isolates the depolarisation introduced by the sample alone.

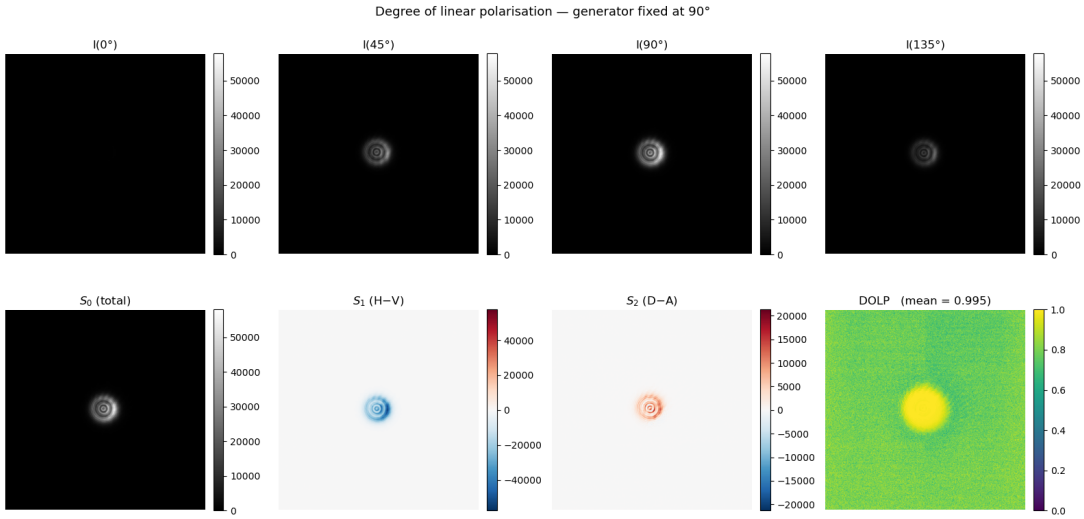

Figure 11.1: Baseline polarisation characterisation of the DC-BAPM setup with no sample present. Top row: raw camera intensities $I(0°)$, $I(45°)$, $I(90°)$, $I(135°)$ with the generator fixed at 0°. Bottom row: reconstructed Stokes maps S_0, S_1, S_2 and the per-pixel degree of linear polarisation. The mean DOLP over the central region is 0.995 with an extinction ratio of 136 : 1, confirming that the residual depolarisation of the bare setup is negligible. Any depolarisation observed when a sample is inserted can therefore be attributed to the sample.

This baseline opens a route to a polarisation-aware reconstruction. The depolarised intensity (proportional to $S_0 - \sqrt{S_1^2 + S_2^2}$) can be computationally subtracted from each measured diffraction pattern before phase retrieval, leaving only the coherent component that the ePIE forward model expects. Removing this incoherent floor should improve the reconstruction contrast and reduce the speckle background that currently limits sensitivity to the thinnest films.

11.4 Superresolution via Sub-Pixel Scanning

A key advantage of ptychography over single-exposure CDI is the ability to reconstruct spatial frequencies beyond the band-limit set by the detector pixel size. When the probe is shifted by sub-pixel amounts relative to the object, different high-frequency components of the object modulate the measured diffraction patterns, and the reconstruction algorithm can deconvolve these contributions to recover a super-resolved image [A. M. Maiden et al., 2011; J. Rodenburg and A. Maiden, 2019]. The achievable resolution is then limited by the maximum spatial frequency $q_{\max} = 1/(2d)$, where d is the detector pixel size multiplied by the wavelength-to-detector-distance ratio. The current DC-BAPM scan grid steps in integer-pixel increments: adding a sub-pixel offset between adjacent scan positions would let the existing ePIE pipeline exploit this superresolution mechanism and push the lateral resolution beyond the detector-imposed band limit at no extra hardware cost.

11.5 Structured Illumination via a Computational Diffuser

The probe is currently shaped by the diffraction grating alone, which produces a fixed wavefront with limited phase diversity. Adding a spatial light modulator (SLM) to the illumination path, downstream of the grating, would let a programmable phase mask be imprinted on the beam without replacing the existing optics. The SLM acts as a computational diffuser: a known, controlled phase pattern that scrambles the probe wavefront in a structured way.

A structured probe improves the ptychographic reconstruction in two ways. The randomised

phase increases the diversity between neighbouring scan positions, which improves the conditioning of the phase-retrieval inverse problem and reduces probe-object ambiguity. The extended Fourier support of the structured probe also pushes the achievable resolution toward the theoretical bound of $\lambda/(2N_{\mathrm{A}})$, beyond what the bare grating illumination can reach.

Babu, A. V. et al. (2023). “Deep learning at the edge enables real-time streaming ptychographic imaging”. In: *Nature Communications* 14.7059. ISSN: 2041-1723. DOI: 10.1038/s41467-023-41496-z.

Bunk, O., Dierolf, M., Kynde, S., Johnson, I., Marti, O., and Pfeiffer, F. (2008). “Influence of the overlap parameter on the convergence of the ptychographical iterative engine”. In: *Ultramicroscopy* 108.5, pp. 481–487. ISSN: 0304-3991. DOI: 10.1016/j.ultramic.2007.08.003.

Gardner, D. F., Tanksalvala, M., et al. (2019). “Ptychographic imaging of incoherently illuminated extended objects using speckle correlations”. In: *Applied Optics* 58.13, pp. 3564–3569. ISSN: 1559-128X. DOI: 10.1364/AO.58.003564.

Gardner, D. F., Zhang, B., Seaberg, M. D., Martin, L. S., Adams, D. E., Salmassi, F., Gullikson, E., Kapteyn, H., and Murnane, M. (2012). “High numerical aperture reflection mode coherent diffraction microscopy using off-axis apertured illumination”. In: *Optics Express* 20.17, pp. 19050–19059. ISSN: 1094-4087. DOI: 10.1364/OE.20.019050.

Goodman, J. W. (2017). *Introduction to Fourier Optics*. 4th ed. New York: W. H. Freeman / Macmillan Learning. ISBN: 978-1-319-11916-4.

Hénon, S. and Meunier, J. (1991). “Microscope at the Brewster angle: Direct observation of first-order phase transitions in monolayers”. In: *Review of Scientific Instruments* 62.4, pp. 936–939. ISSN: 0034-6748. DOI: 10.1063/1.1142032.

Huang, B. (2026). *Microscopy: Fourier Space*. YouTube lecture on Fourier-space description of microscopy. iBiology Techniques. URL: https://www.youtube.com/watch?v=xhO8iz2qCOE (visited on 05/21/2026).

Huang, X., Yan, H., Harder, R., Hwu, Y., Robinson, I. K., and Chu, Y. S. (2014). “Optimization of overlap uniformness for ptychography”. In: *Optics Express* 22.10, pp. 12634–12644. DOI: 10.1364/OE.22.012634.

Jørgensen, P. S. et al. (2024). “Hard x-ray grazing-incidence ptychography: large field-of-view nanostructure imaging with ultra-high surface sensitivity”. In: *Optica* 11.2, pp. 197–204. ISSN: 2334-2536. DOI: 10.1364/OPTICA.505478.

Kim, H., Locans, U., Nebling, R., Dejkameh, A., Kazazis, D., Ekinci, Y., and Mochi, I. (2020). “High resolution and uniform image reconstruction in a large field-of-view for EUV actinic mask review”. In: *Photomask Technology 2020*. Vol. 11518. SPIE, p. 115180X. DOI: 10.1117/12.2573240.

Konijnenberg, S. (2026a). *Angular Spectrum Method (plane wave decomposition, evanescent field, diffraction limit)*. Lecture 02 of the Computational Imaging series. YouTube. URL: https://www.youtube.com/watch?v=31072jVfIUE (visited on 05/21/2026).

– (2026b). *Numerically simulating the propagation of coherent optical fields (Fourier optics)*. Tutorial on FFT-based propagation of coherent fields. YouTube. URL: https://www.youtube.com/watch?v=z-2tdkJ0Yzc (visited on 05/21/2026).

Mack, C. A. (2007). *Fundamental Principles of Optical Lithography: The Science of Microfabrication*. Chichester, England: John Wiley & Sons. ISBN: 978-0-470-01893-4. DOI: 10.1002/9780470723876.

Maiden, A. M., Humphry, M. J., Zhang, F., and Rodenburg, J. M. (2011). "Superresolution imaging via ptychography". In: *Journal of the Optical Society of America A* 28.4, pp. 604–612. ISSN: 1084-7529. DOI: 10.1364/JOSAA.28.000604.

Meshalkin, A. Y., Podlipnov, V. V., Ustinov, A. V., and Achimova, E. A. (2019). "Analysis of diffraction efficiency of phase gratings in dependence of duty cycle and depth". In: *Journal of Physics: Conference Series*. Vol. 1368. IOP Publishing, p. 022047. DOI: 10.1088/1742-6596/1368/2/022047.

Novotny, L. and Hecht, B. (2012). *Principles of Nano-Optics*. 2nd ed. Cambridge, UK: Cambridge University Press. ISBN: 978-1-107-00546-4. DOI: 10.1017/CBO9780511813535.

Raskar, R. (1999). "Oblique Projector Rendering on Planar Surfaces for a Tracked User". In: *ACM SIGGRAPH 99 Conference Abstracts and Applications*. ACM. DOI: 10.1145/311625.311877.

Raskar, R. and Beardsley, P. (2001). "A Self-Correcting Projector". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Hawaii: IEEE.

Rodenburg, J. M. and Faulkner, H. M. L. (2004). "A phase retrieval algorithm for shifting illumination". In: *Applied Physics Letters* 85.20, pp. 4795–4797. ISSN: 0003-6951. DOI: 10.1063/1.1823034.

Rodenburg, J. and Maiden, A. (2019). "Ptychography". In: *Springer Handbook of Microscopy*. Ed. by P. W. Hawkes and J. C. H. Spence. Cham: Springer. Chap. 17, pp. 819–904. DOI: 10.1007/978-3-030-00069-1_17.

Sukthankar, R., Stockton, R., and Mullin, M. (2001). "Smarter Presentations: Exploiting Homography in Camera-Projector Systems". In: *Proceedings of the International Conference on Computer Vision (ICCV)*. IEEE.

Suzuki, A., Shimomura, K., Hirose, M., Burdet, N., and Takahashi, Y. (2016). "Dark-field X-ray ptychography: Towards high-resolution imaging of thick and unstained biological specimens". In: *Scientific Reports* 6, p. 35060. DOI: 10.1038/srep35060.

Tanksalvala, M. et al. (2021). "Nondestructive, high-resolution, chemically specific 3D nanostructure characterization using phase-sensitive EUV imaging reflectometry". In: *Science Advances* 7.5, eabd9667. ISSN: 2375-2548. DOI: 10.1126/sciadv.abd9667.

Thorlabs (2026). *Diffraction Gratings Tutorial*. URL: https://www.thorlabs.com/diffraction-gratings-tutorial (visited on 03/23/2026).

Tian, L., Liu, Z., Yeh, L.-H., Chen, M., Zhong, J., and Waller, L. (2015). "Computational illumination for high-speed *in vitro* Fourier ptychographic microscopy". In: *Optica* 2.10, pp. 904–911. ISSN: 2334-2536. DOI: 10.1364/OPTICA.2.000904.

Zhang, Y., Zhang, Z., and Maiden, A. (2022). "Ptycho-cam: a ptychographic phase imaging add-on for optical microscopy". In: *Applied Optics* 61.10, pp. 2874–2880. ISSN: 1559-128X. DOI: 10.1364/AO.452835.

Zhang, Z., Xu, S., Xi, T., and Yao, B. (2019). "Near-field Fourier ptychography: superresolution phase retrieval via speckle illumination". In: *Optics Express* 27.19, pp. 26445–26461. ISSN: 1094-4087. DOI: 10.1364/OE.27.026445.

Zhe Sun, D. Z. and Baldelli, S. (2017). "Distortion Correction for a Brewster Angle Microscope Using an Optical Grating". In: *American Chemical Society* 89, pp. 2186–2190. ISSN: 0003-2700. DOI: 10.1021/acs.analchem.6b04738.

PIE Pipeline Flowchart

This appendix collects the full PIE flowchart referenced in Section 2.2. The diagram anchors each of the five algorithm steps (Eqs. (2.1)–(2.5)) to the physical plane at which it lives, making explicit how the iteration moves between the sample plane and the detector plane.

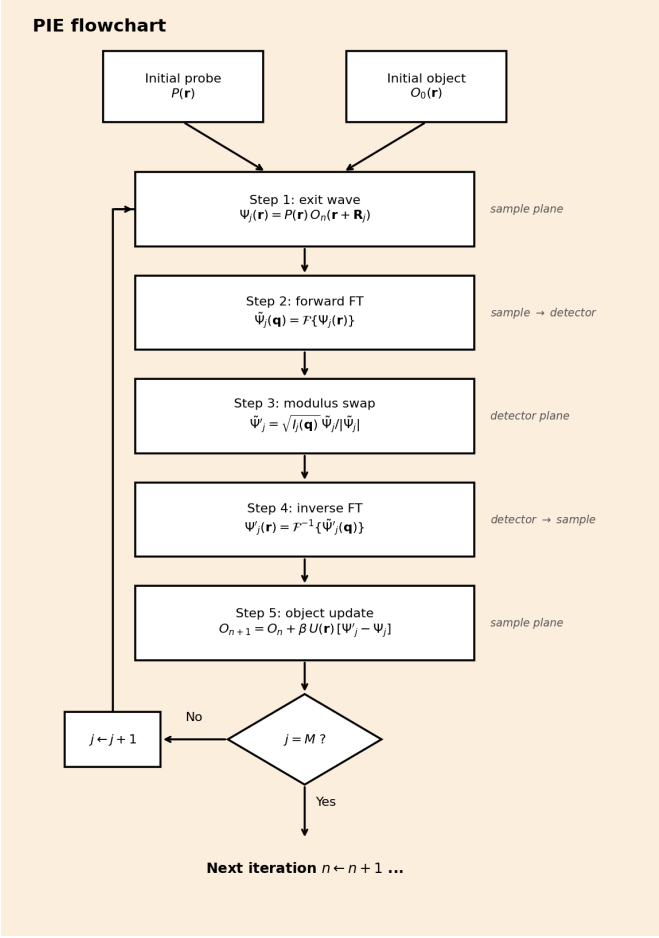

Figure 1: PIE flowchart. Two inputs at the top, an initial probe $P(\mathbf{r})$ and an initial object estimate $O_0(\mathbf{r})$, feed the five-step update. At each scan position j the exit wave is built, propagated to the detector, modulus-swapped with the measured intensity $\sqrt{I_j(\mathbf{q})}$, back-propagated, and used to update the object. After $j = M$ positions the iteration counter n advances. The right-hand italics indicate the physical plane at which each step lives.

This appendix lists the step-by-step procedure for aligning the optics of the DC-BAPM setup (Section 3.1). The hardest part of building the setup is aligning the optics. The procedure I followed is:

1. **Reference axis.** Send the laser beam through two small apertures placed on the optical table so that the beam runs parallel to the row of screw holes the other components will be mounted on. Fix the laser.
2. **ND filter.** Install the ND filter directly after the laser to protect the camera sensor from the unattenuated beam.
3. **Camera, first pass.** Place the camera in the beam path. Using the manufacturer's SpinView software, overlay a crosshair on the live image and translate the camera until the beam centre lands on the crosshair. Fix the camera in this position.
4. **Illumination lens.** Insert the lens a few millimetres beyond its focal length from the sensor and adjust it laterally until the focused spot stays on the crosshair. Fix the lens.
5. **Beam expander.** Add the beam expander and repeat the same centering check, keeping the focused spot on the crosshair. Fix the expander.
6. **Polarisation optics.** Insert the half-wave plate and the generator polariser. Check that the focus on the crosshair is unchanged.
7. **Sample stage.** Remove the camera and install the stage at the desired angle of incidence, with the sample placed a chosen distance from the focal point. Fix the stage.
8. **Camera, final position.** Reinstall the camera on the optical axis. The sensor must be perpendicular to the axis, otherwise the diffraction pattern is recorded in a distorted plane. To set this, place a protractor flat against the camera housing with its long edge through the housing centreline so that one corner points at the probe spot on the sample. As a final check, confirm that the beam reflected from a bare area of the sample lands in the middle of the crosshair.

This appendix validates the saturation-cascade HDR merge described in Section 3.4.2 on a strong, well-characterised test target before applying it to the graphene flake in Section 3.4.3. The PMMA on silicon Siemens star has sharp edges and clean spatial frequencies, so any improvement in high-frequency fidelity is easy to verify in the reconstruction. The full sample layout is shown in Appendix 11.5.

Two-Exposure Bracket

Two exposure times, $t = 10\,000\,\mu s$ and $80\,000\,\mu s$, are sufficient to cover the dynamic range. Figure 2 shows both frames on linear and logarithmic scales. The short exposure keeps the zero-order core unsaturated, the long exposure brings out the high-frequency fringes but clips the core.

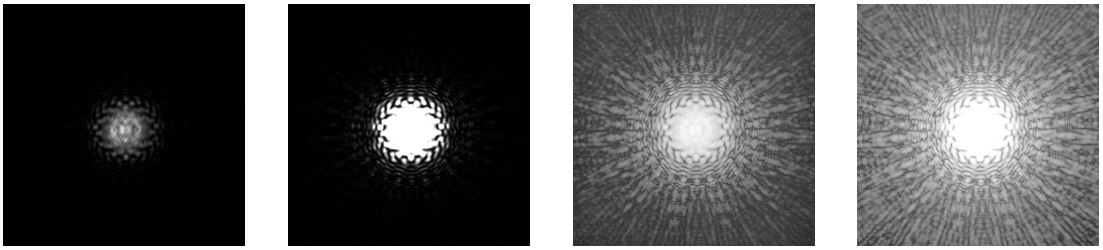

(a) $t = 10\,000\,\mu s$, linear (b) $t = 80\,000\,\mu s$, linear (c) $t = 10\,000\,\mu s$, log (d) $t = 80\,000\,\mu s$, log

Figure 2: Two-exposure bracket on the Siemens star at a single scan position. The short exposure (a,c) keeps the core unsaturated but loses the tail in the noise, the long exposure (b,d) reveals the tail but clips the core.

Figure 3 maps which of the two frames contributes to each pixel of the merged pattern: the bright core is taken from the short exposure, the dim tail from the long one.

The merged pattern and the radial intensity profile through the merge boundary are shown in Figure 4. The individual scaled frames and the merged result line up smoothly across the transition, confirming that the saturation-cascade merge does not introduce a step artefact.

Effect on the Reconstruction

The effect of the merge on the ptychographic reconstruction is shown in Figure 5. The single-exposure reconstruction resolves only the coarse spokes near the centre of the target, the HDR reconstruction recovers the finer spokes outward to a higher spatial frequency. This confirms that the additional high-frequency information captured by the long exposure translates directly into finer object detail in the retrieved amplitude.

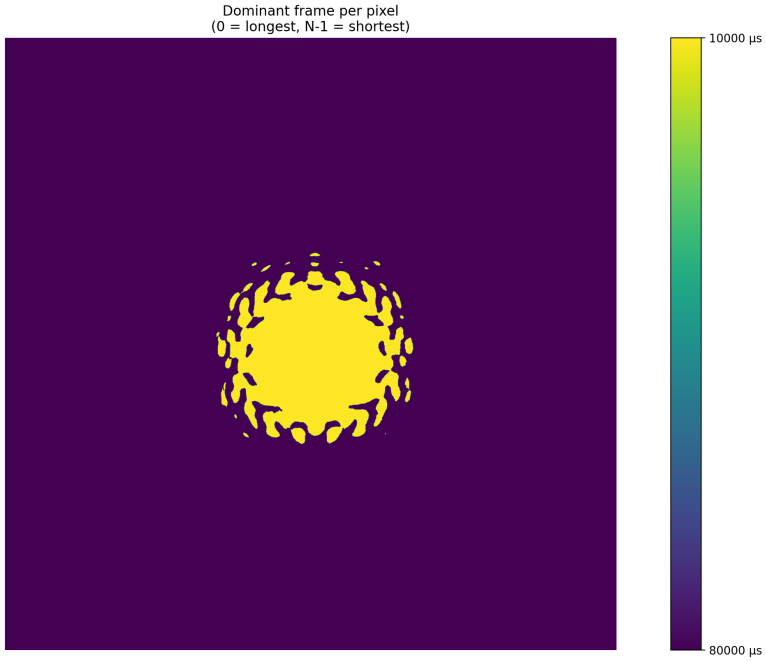

Figure 3: Dominant frame per pixel for the Siemens star merge. The bright zero-order core is covered by the 10 000 µs frame and the high-frequency tail by the 80 000 µs frame.

NILS Resolution Analysis

The Normalised Image Log-Slope (NILS) quantifies the local contrast of the reconstructed amplitude along a circle of radius r centred on the Siemens star. A NILS value above 1 indicates that the spokes at that radius are resolved, below 1 they are lost in noise or undersampled. The following figures compare the NILS profiles at five radii for the HDR-merged and the single-exposure (NO-HDR) datasets.

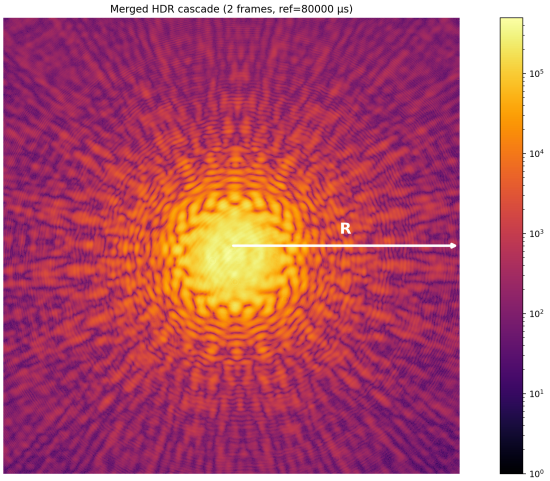

(a) Merged HDR pattern with radial sampling direction

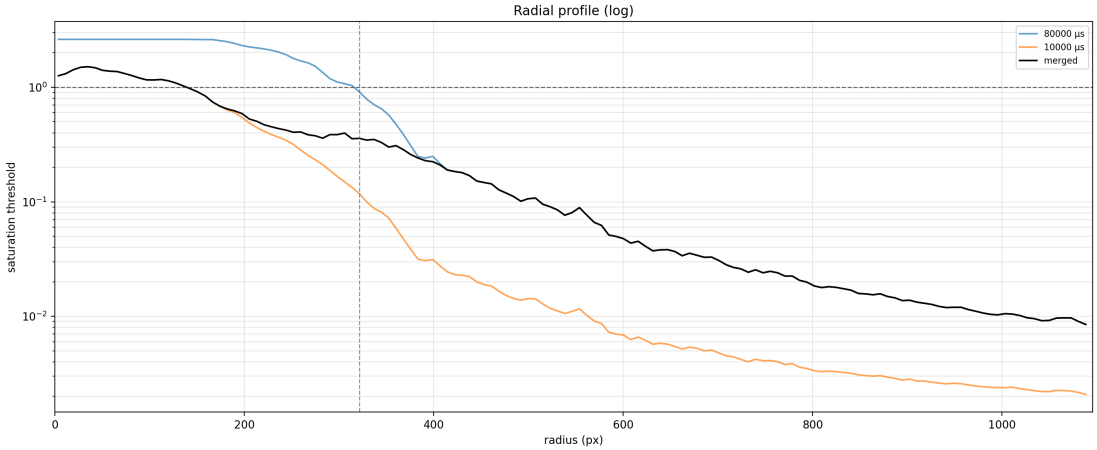

(b) Log-radial intensity profile

Figure 4: Continuity check across the merge boundary for the Siemens star. (a) Merged HDR pattern with an arrow indicating the radius along which the intensity profile in (b) is sampled. (b) The radial intensity profiles of the scaled 10 000 µs and 80 000 µs frames and the merged result show a smooth, continuous transition across the merge boundary.

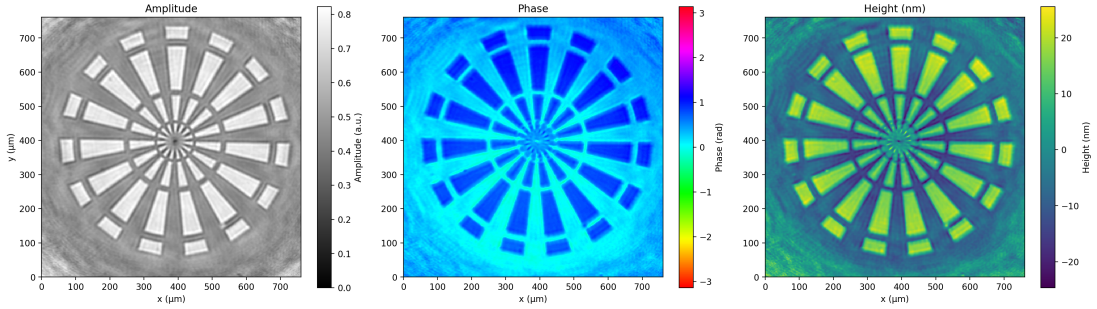

(a) Single exposure ($t = 10\,000\,\mu s$)

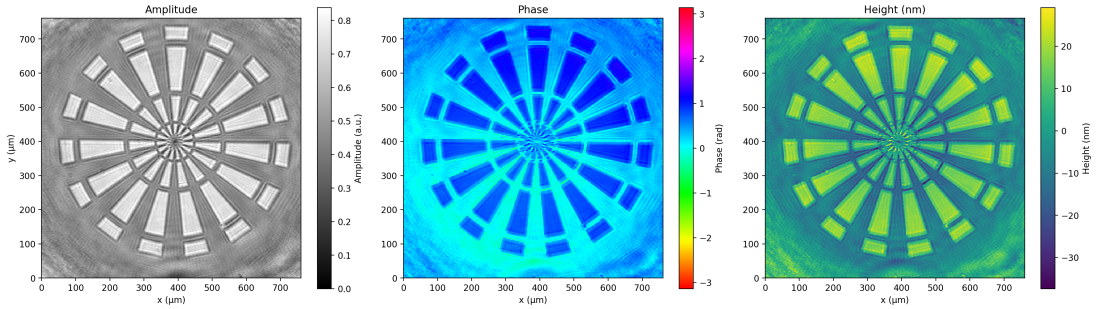

(b) HDR-merged dataset

Figure 5: ePIE reconstruction of the Siemens star from (a) a single-exposure dataset and (b) the HDR-merged dataset. No clear improvement in resolution or contrast is visible between the two reconstructions.

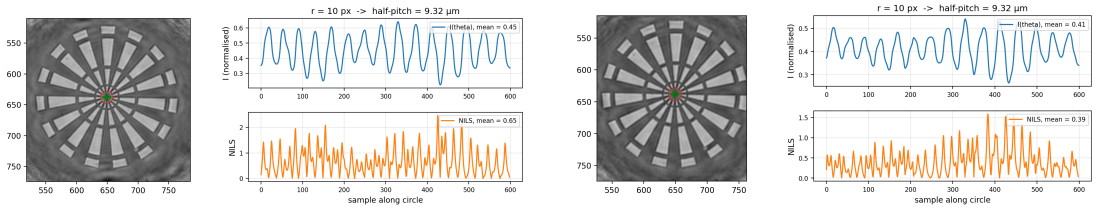

(a) HDR (b) NO-HDR

Figure 6: NILS at $r = 10\,\mu m$ for the HDR and single-exposure reconstructions.

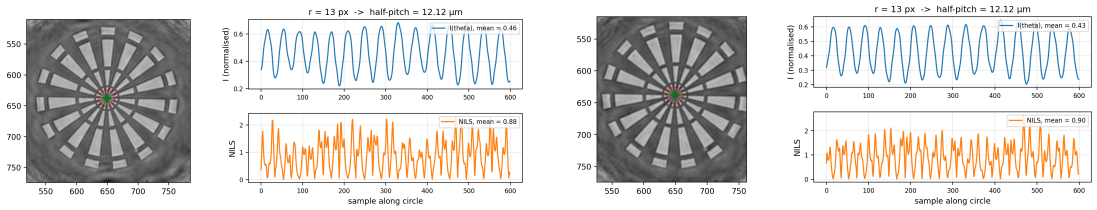

(a) HDR (b) NO-HDR

Figure 7: NILS at $r = 13\,\mu m$.

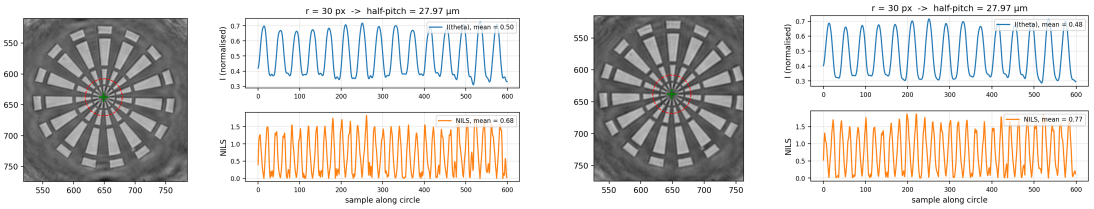

(a) HDR (b) NO-HDR

Figure 8: NILS at $r = 30\,\mu\text{m}$.

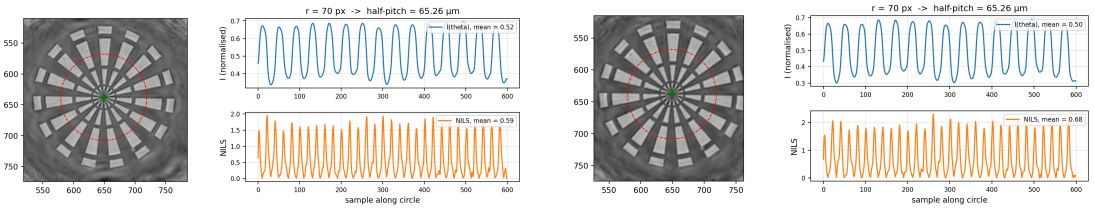

(a) HDR (b) NO-HDR

Figure 9: NILS at $r = 70\,\mu\text{m}$.

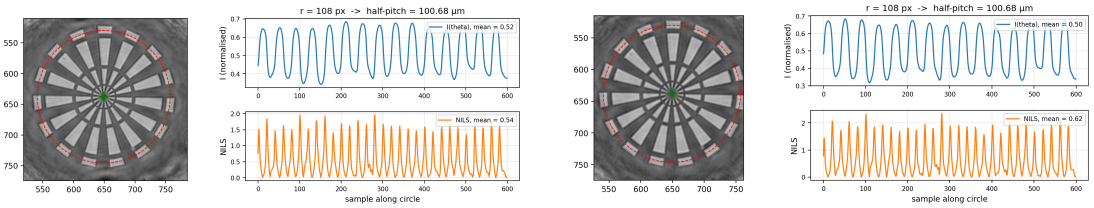

(a) HDR (b) NO-HDR

Figure 10: NILS at $r = 108\,\mu\text{m}$.

This appendix collects the detailed diagnostics for the five-exposure HDR merge applied to the graphene flake in Section 3.4.3, together with the contrast metrics referenced in Section 3.5.

Graphene Merge Diagnostics

The dominant-frame map in Figure 11 shows which exposure contributes most at each pixel of the graphene merge. Compared with the two-band Siemens-star map, the five-band layered structure reflects the much wider intensity range that has to be covered.

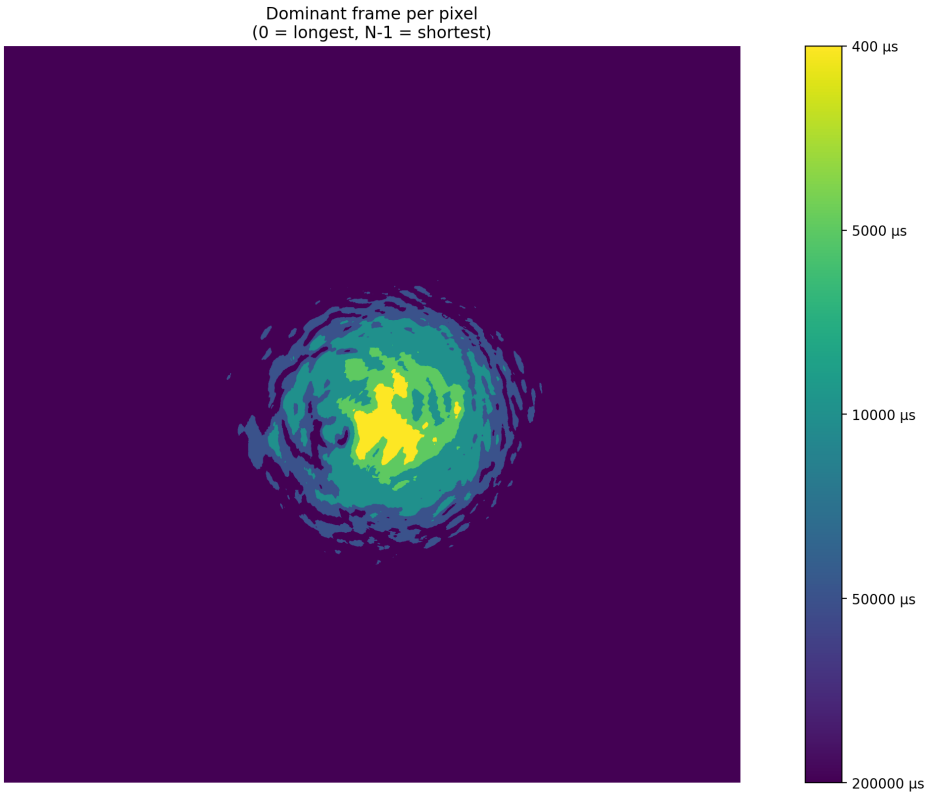

Figure 11: Dominant frame per pixel for the graphene merge. Each colour indicates which exposure time contributes most at that pixel: the bright zero-order core is covered by the shortest exposures and the high-frequency tail by the longest.

The continuity check across all four merge boundaries is shown in Figure 12. The five scaled radial profiles overlap smoothly and the merged result follows the dominant frame in each region without a visible discontinuity.

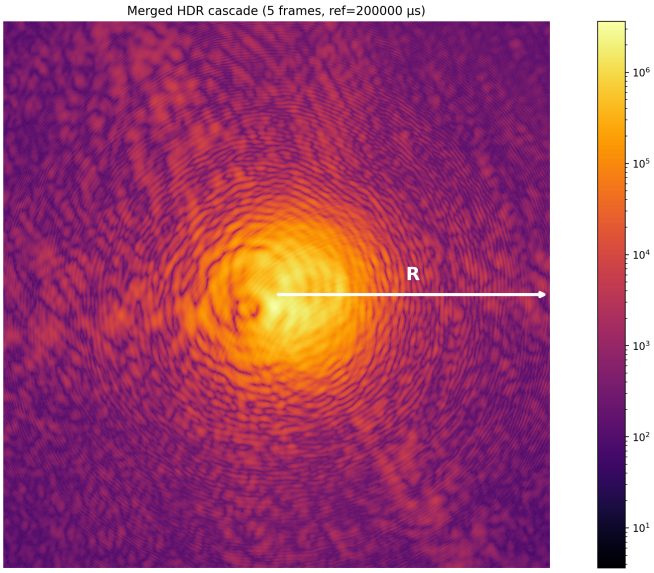

(a) Merged HDR pattern with radial sampling direction

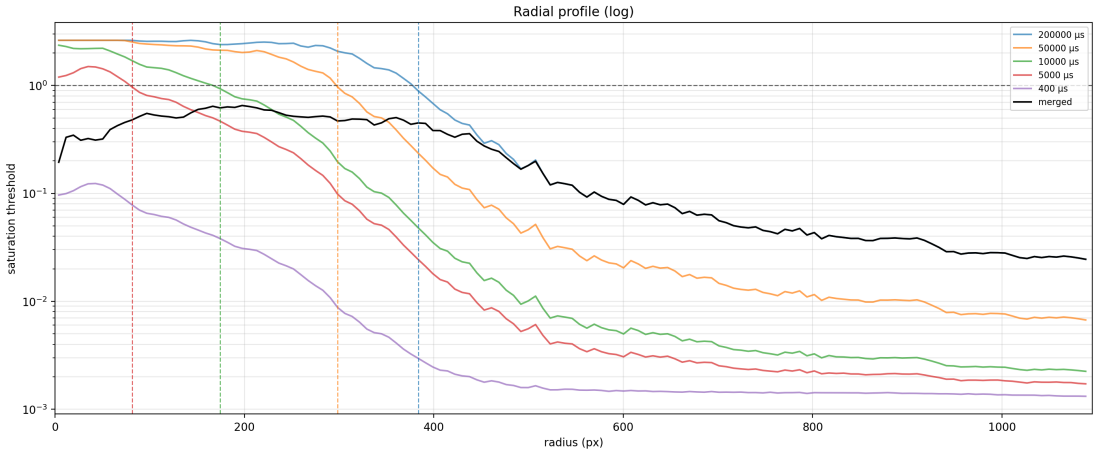

(b) Log-radial intensity profile

Figure 12: Continuity check across the merge boundaries for the graphene flake. (a) Merged HDR pattern with an arrow indicating the radius along which the intensity profile in (b) is sampled. (b) The radial intensity profiles of the five scaled frames and the merged result show a smooth, continuous transition across all merge boundaries.

Contrast Metrics

Table 1 lists the contrast metrics for the three Siemens-star diffraction patterns compared in Section 3.5 (Figure 3.7).

Table 1: Contrast metrics for the three diffraction patterns. All Michelson values equal 1 because the background is essentially dark ($P_1 \approx 0$), the dynamic range and peak-to-mean ratio are the more informative diagnostics.

Metric	$t = 10\,000\,\mu s$	$t = 80\,000\,\mu s$	HDR merge
Michelson C (P1–P99)	1.000	1.000	1.000
RMS contrast (std/mean)	4.20	3.15	4.29
Dynamic range (dB)	88.8	96.3	89.3
Peak / mean	57.9	14.8	59.2
Mean counts	1077	4429	1107
Max counts	62304	65520	65535
Saturated fraction	0.000	0.000	0.000

This appendix collects the Fresnel reflectance curves for the four substrates considered alongside PMMA: water, carbon, silicon and aluminium. The PMMA reflectance used in the main DC-BAPM measurements is shown in Fig. 9.2 of Chapter 9. Refractive indices and Brewster (or pseudo-Brewster) angles for all five materials are tabulated in Table 9.1, the corresponding grating specifications are given in Table 9.2.

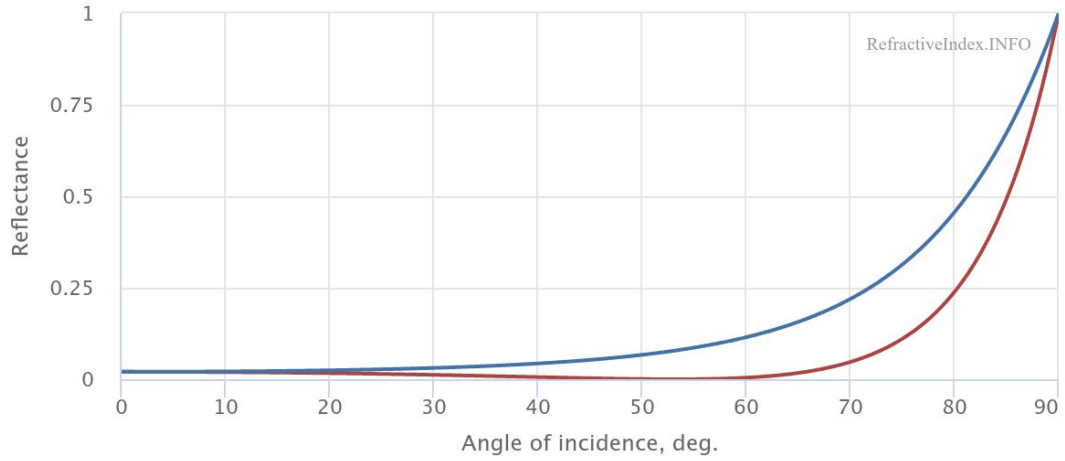

Figure 13: Fresnel reflectance of water at $\lambda = 632.8\,\text{nm}$. Water is a transparent dielectric ($k = 0$), so R_p vanishes exactly at the Brewster angle $\theta_B = 53.07^\circ$.

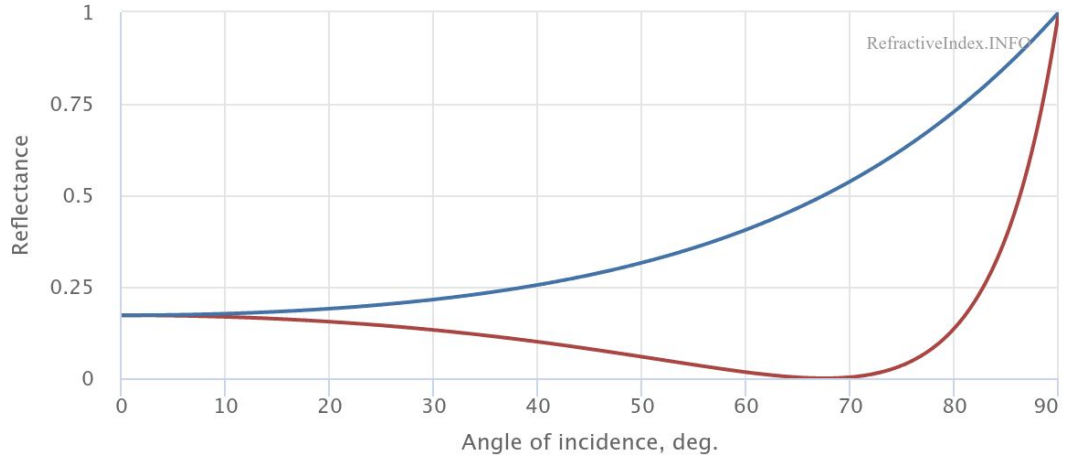

Figure 14: Fresnel reflectance of carbon at $\lambda = 632.8\,\text{nm}$, $\theta_B = 67.5^\circ$. Carbon is treated as a transparent dielectric here ($k \approx 0$), so R_p reaches zero at the Brewster angle.

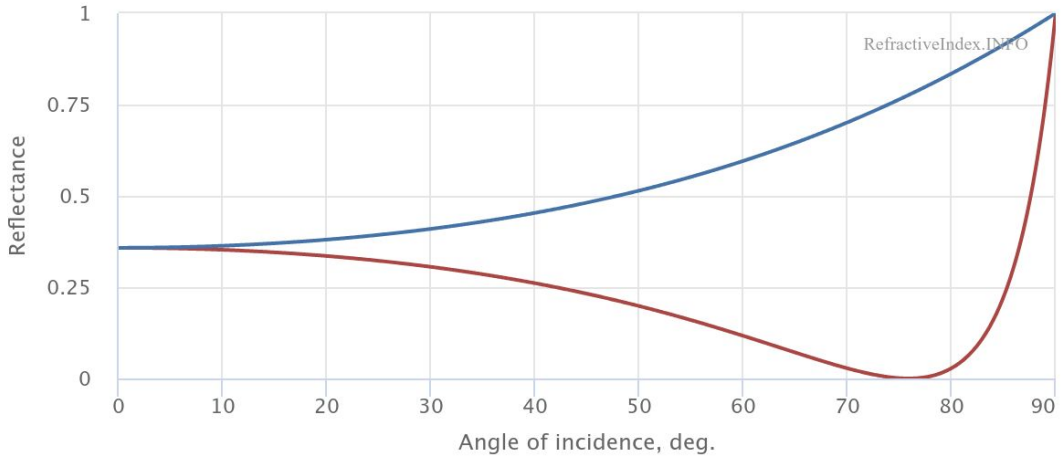

Figure 15: Fresnel reflectance of silicon at $\lambda = 632.8\,\text{nm}$. Silicon has a small but nonzero imaginary part ($n = 3.847$, $k = 0.016$), so R_p does not reach zero but a pseudo-Brewster minimum at $\theta_{\text{pB}} \approx 75.4^\circ$. The approximation $\theta_{\text{pB}} \approx \arctan(n)$ is accurate to within 0.1° because $k \ll n$.

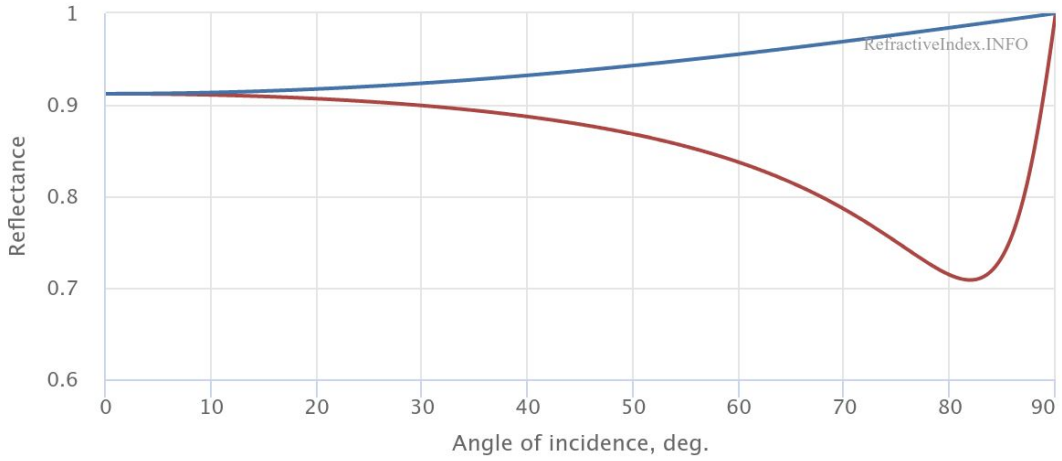

Figure 16: Fresnel reflectance of aluminium at $\lambda = 632.8\,\text{nm}$. Aluminium has a large imaginary part ($k = 6.92$), so there is no true Brewster zero, the pseudo-Brewster angle near $\theta_{\text{pB}} \approx 80^\circ$ leaves a residual $R_p^{\min} \approx 70\%$, which gives Al substrates intrinsically low BAM contrast.

The reconstruction engine used in this work is the EUVPtycho C++/CUDA code developed at the Paul Scherrer Institute in Switzerland. Reproducing its build environment from scratch is non-trivial: the engine depends on a specific CUDA toolkit version, on Armadillo built from source, on OpenCV, FFTW, HDF5, Boost, and Qt5, and on a particular Ubuntu base system. To make the reconstruction reproducible across workstations I packaged the full build in a Docker image. The container is built from a single `Dockerfile` in seven steps.

1. **Base image.** The image extends `nvidia/cuda:12.4.1-cudnn-devel-ubuntu22.04`. This provides CUDA 12.4, cuDNN, and the development headers on top of Ubuntu 22.04. At runtime the host GPU driver is forwarded into the container by the NVIDIA Container Toolkit, so the image stays portable between machines as long as their drivers support CUDA 12.4.

2. **System dependencies.** A single `apt-get install` call pulls in the build tools (`build-essential`, `cmake`, `git`, `wget`), the scientific libraries used by the engine (`libopencv-dev`, `libfftw3-dev`, `libhdf5-dev`), the Boost components (`filesystem`, `system`, `serialization`, `regex`, `iostreams`), and the Qt5 build system (`qtbase5-dev`, `qt5-qmake`) together with `gnuplot` for diagnostic plotting. The apt cache is removed immediately afterward to keep the image small.

3. **Armadillo.** The Ubuntu repositories do not carry a recent enough Armadillo release, so version 12.8.2 is downloaded from SourceForge, built with `cmake; make`, installed system-wide, and the source tree is deleted. A final `ldconfig` refreshes the dynamic linker cache.

4. **CUDA path symlink.** EUVPtycho's build scripts expect to find the CUDA toolkit at `/opt/cuda`. A symlink from `/usr/local/cuda-12.4` (the toolkit location in the base image) to `/opt/cuda` satisfies that expectation without touching the build scripts.

5. **Header and library paths.** `CPATH` is set to include `/usr/include/opencv4` and `/usr/include/hdf5/serial`, and `LIBRARY_PATH` is set to the matching HDF5 library directory. This avoids edits to the project's qmake file for the Debian-packaged include layout.

6. **Source copy and compile.** The source tree on the host is copied into `/usr/src/app/EUVPtycho/euvptycho/src` inside the image. `qmake EUVPtycho.pro` generates the Makefile, and `make clean && make` compiles the binaries.

7. **Runtime PATH.** The compiled executables in `/usr/src/app/EUVPtycho/bin` are appended to `PATH`, and the container's default command is set to `bash` so the user lands in an interactive shell with the engine on the path.

The image is built and run with

```
docker build -t euvptycho .
docker run --gpus all -it -v /path/to/data:/data euvptycho
```

The `-gpus all` flag exposes the host GPU to the container, and the bind mount makes the diffraction datasets and reconstruction outputs available on the host filesystem. Two prerequisites apply on the host: the GPU must be an NVIDIA card (CUDA is vendor-specific and will not run on AMD or Intel graphics), and the NVIDIA Container Toolkit must be installed alongside Docker. The toolkit provides the runtime hook that forwards the host GPU device and driver into the container, without it, `-gpus all` is silently ignored and the engine falls back to a CPU code path or fails outright. Building from a clean cache takes a few minutes, dominated by the Armadillo compile. Subsequent builds from the cached image are much faster and take only a few seconds.

This appendix shows photographs of the physical computer hardware used to drive the DC-BAPM acquisition and run the ePIE reconstruction. The reconstruction engine requires an NVIDIA GPU (see Appendix 11.5), since the laptop used in the lab has a slower NVIDIA card, an external GPU enclosure was connected over Thunderbolt and used for the reconstruction.

Figure 17: Full DC-BAPM setup on the optical table, showing the optical components (laser, polarisation optics, sample stage, grating, and camera) together with the laptop and external GPU enclosure used to acquire and process the data.

Figure 18: Close-up of the laptop and the external NVIDIA GPU enclosure connected over Thunderbolt.

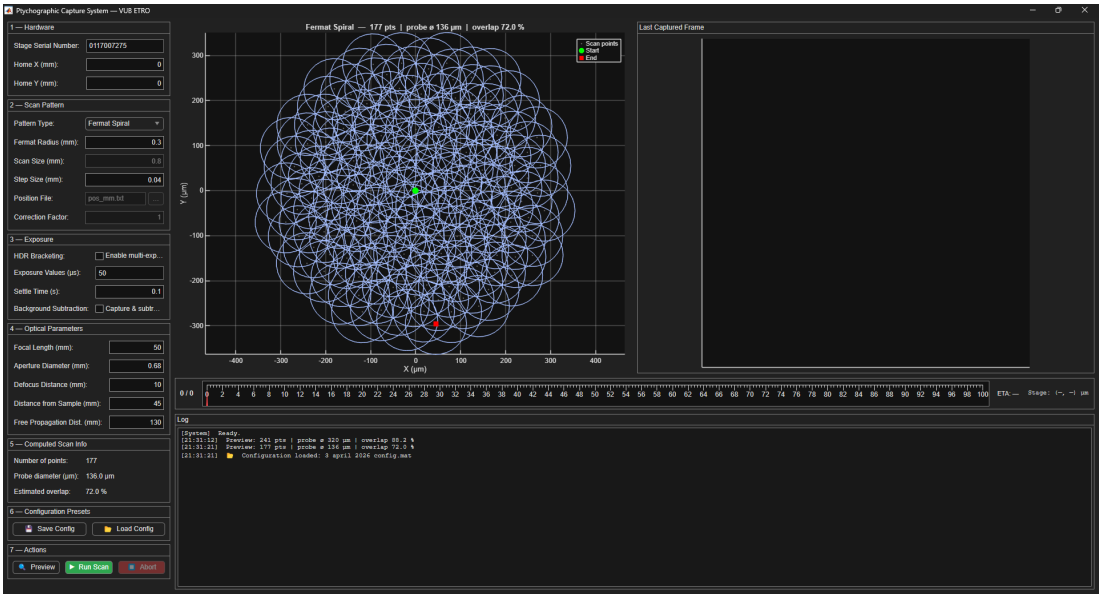

Figure 19: Graphical user interface used to launch and monitor a capture.

The EUVPtycho engine described in Appendix 11.5 is driven by a plain-text parameter file passed as its only command-line argument. The file groups parameters into INI-style sections (`[General]`, `[Detector]`, `[Probe]`, `[Reconstruction]`, `[Refinement]`, `[Simulation]`, `[GPU]`, `[Data]`) and is consumed by the engine's parser at start-up. Listing 1 reproduces the parameter file used for one of the reconstructions in this work (42° incidence, p-polarised, no HDR).

Listing 1: Reconstruction parameter file passed to the EUVPtycho engine for a representative 42°-incidence dataset.

```
[General]
run_name          = ''
output_name       = 'ePIE-corrected_42deg'
wavelength        = 632.8e-9
visualize         = False
nr_photons_arr    = [1e2]
nr_threads        = 16
save_params       = True

[Detector]
nr_pixels         = 2192
det_pixel_size    = 4.54e-6
det_rot_ang       = 0.0
det_tilt_angX     = 0
flip_xy           = False
invert_x          = True
invert_y          = False
bin               = 2

[Probe]
probe_type        = 'focus'
aperture_diam     = 2.1000e-03    # 2.10 mm
aperture_type     = 'circ'
edgewidth         = 2
prop_free_dist    = 7.0000e-02    # 70.00 mm
dist_obj_screen   = 5.2000e-02    # 52.00 mm
focal_dist        = 5.0000e-02    # 50.00 mm
dist_from_focus   = 9.0000e-03    #  9.00 mm

[Reconstruction]
algorithm         = ['EPIE']
nr_iterations     = [400]
dump_interval     = [10, 100, 150, 200, 300, 400,
                     600, 800, 1000, 1400, 1700, 2000]
alpha             = 0.7
beta              = 0.9
epie_poscorr      = False
poscorr_x         = 1.0
poscorr_y         = 1.0
poscorr_beta      = 2.0
plot_interval     = 20
dump_start        = 10
count_bound       = 0.001
```

```
probe_change_start    = 100
probe_change_stop     = 1500
probe_regularization = 0.1
clip_object       = True
clip_max          = 1.0
clip_min          = 0.0
doParamSweep      = False
probe_constrain       = True
probe_keep_ratio      = 0.94
probe_renorm2data     = False
median_filter_use     = False
filter_until_it       = 4
dynSatThreshRatio     = 0.55
nr_probes             = 1
nr_probe_modes        = 1
probe_mode_power      = [0.6, 0.2, 0.08, 0.02]

[Refinement]
ml_refinement      = False
use_existing_data = False
ml_it              = 300
mue                = 30.0

[Simulation]
is_simulation     = False

[GPU]
enable_full_gpu   = True
```

The key physical parameters and their role in the reconstruction are:

- `wavelength`: the HeNe wavelength (632.8×10^{-9} m) that fixes the spatial-frequency calibration of the detector.
- `nr_pixels`, `det_pixel_size`, `bin`: detector geometry. The 2192×2192 array is binned 2×2 to reduce memory bandwidth without sacrificing the highest spatial frequency carried by the diffraction pattern.
- `aperture_diam`, `focal_dist`, `dist_from_focus`, `dist_obj_screen`: the forward illumination model that produces the initial probe guess shown in Fig. 6.1.
- `algorithm`, `nr_iterations`: select ePIE and run it for 400 outer iterations.
- `alpha`, `beta`: the object regularisation constant and feedback parameter of Chapter 2.
- `probe_change_start`, `probe_change_stop`: the iteration window in which the probe is allowed to update, outside this window the probe is held fixed.
- `epie_poscorr`, `poscorr_x`, `poscorr_y`, `poscorr_beta`: scan-position refinement options.
- `nr_probes`, `nr_probe_modes`, `probe_mode_power`: Set to a single mode for the runs in this work.

The resolution test sample used throughout this work is a PMMA on silicon resolution target fabricated by IMEC. It contains four distinct regions, each targeting a different aspect of the system's imaging capability. Figure 20 shows a wide-field optical photograph of the sample with each region labelled, and Figure 21 shows the corresponding GDS layout from which the sample was patterned.

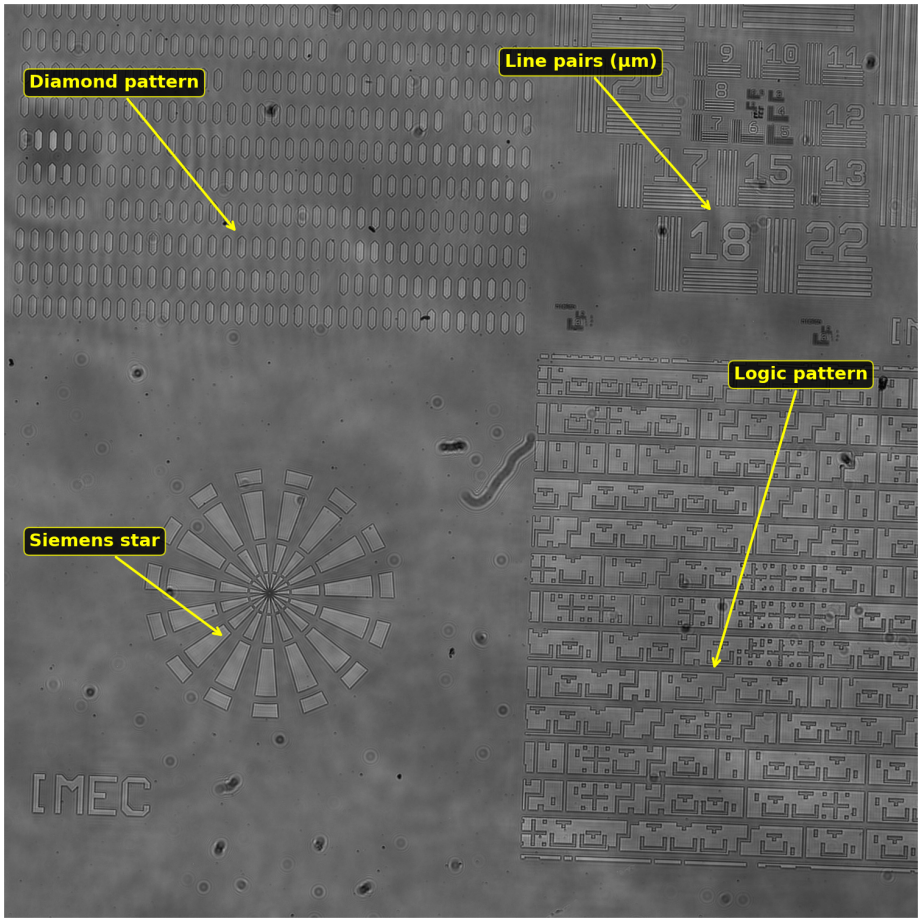

Figure 20: Wide-field optical microscope image of the IMEC resolution target. The four functional regions are indicated: the Siemens star (bottom-left), the diamond pattern (top-left), the micrometer line pairs (top-right), and the logic pattern (bottom-right).

Siemens star. A radial spoke target with fifteen spoke pairs. The spokes converge at the centre so that radial position maps directly to spatial frequency (Eq. (6.2)). This region is used for quantitative resolution measurements via the NILS metric (Section 6.3.2).

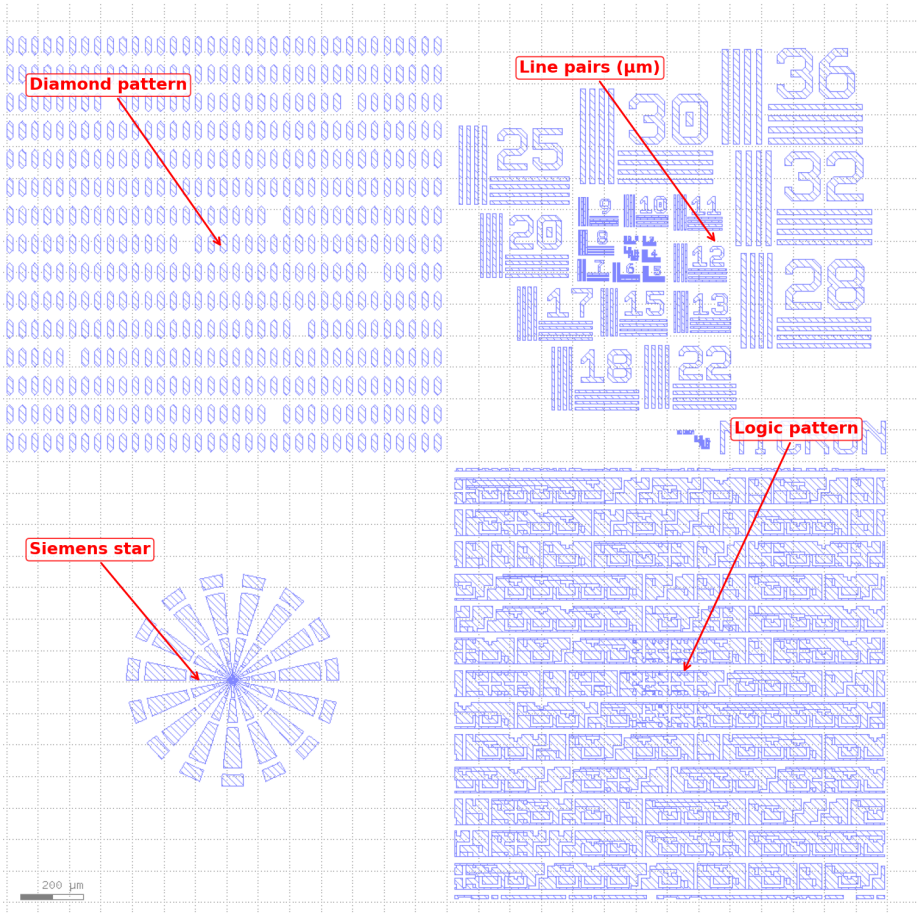

Figure 21: GDS layout of the resolution target used to pattern the PMMA layer. The four regions correspond directly to those visible in Fig. 20: the Siemens star (bottom-left), the diamond pattern (top-left), the micrometer line pairs (top-right), and the logic pattern (bottom-right). The scale bar indicates 200 µm.

Diamond pattern. A periodic array of diamond-shaped PMMA features arranged in a regular grid. The array provides a large uniform region for verifying distortion correction and stitching consistency over a wide field.

Line pairs (µm). Isolated groups of parallel PMMA lines with widths ranging from approximately 1 µm to 36 µm, each group labelled with its nominal pitch in micrometres. These features serve as a direct spatial calibration check.

Logic pattern. A dense array of logic-gate-like PMMA structures repeated across a large block. The pattern mimics the irregular, high-spatial-frequency content typical of a semiconductor circuit and tests the system's ability to reconstruct aperiodic, broadband objects.

The following reconstructions were recorded on 2026-04-28 and 2026-04-29 at AOI = 42° as additional scans of the graphene sample. They are included here for completeness, the main analysis is presented in Chapter 8.

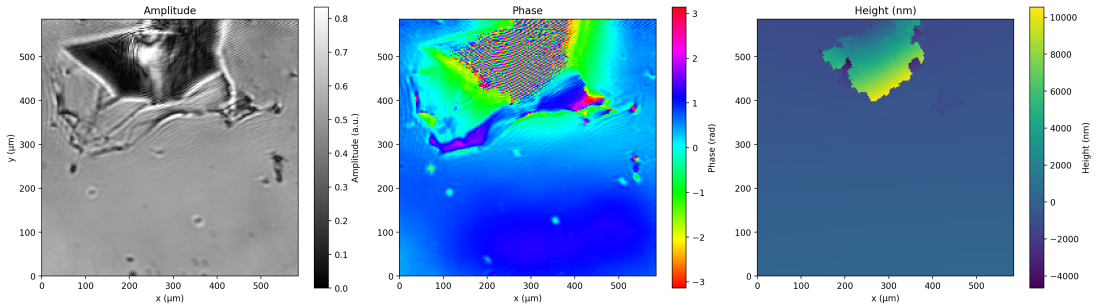

Figure 22: Graphene sample, scan 1 (2026-04-28, 17:00).

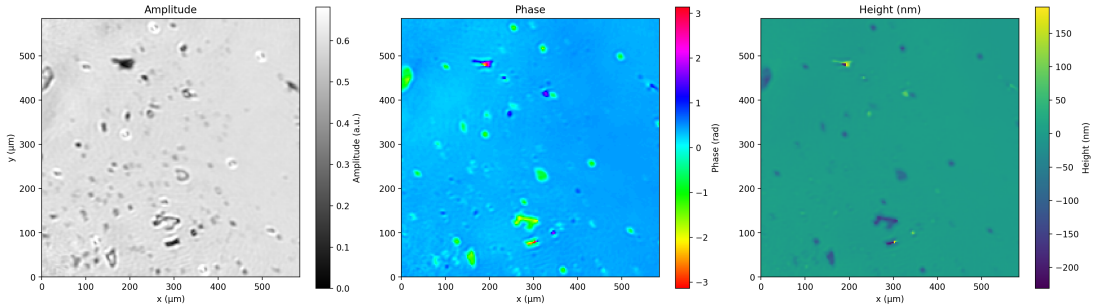

Figure 23: Graphene sample, scan 2 (2026-04-28, 19:26).

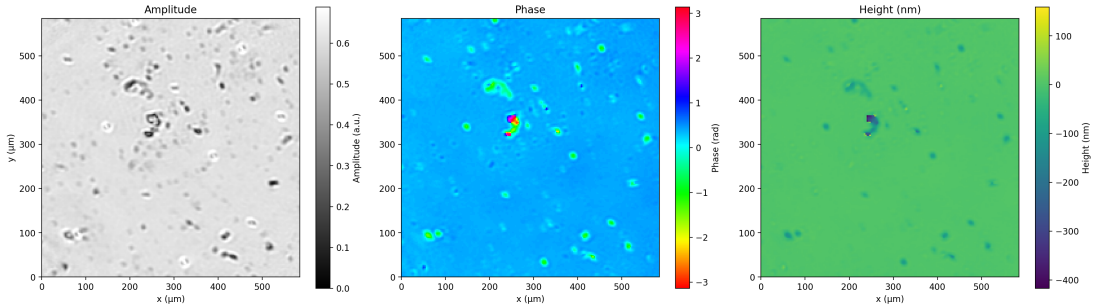

Figure 24: Graphene sample, scan 3 (2026-04-28, 19:45).

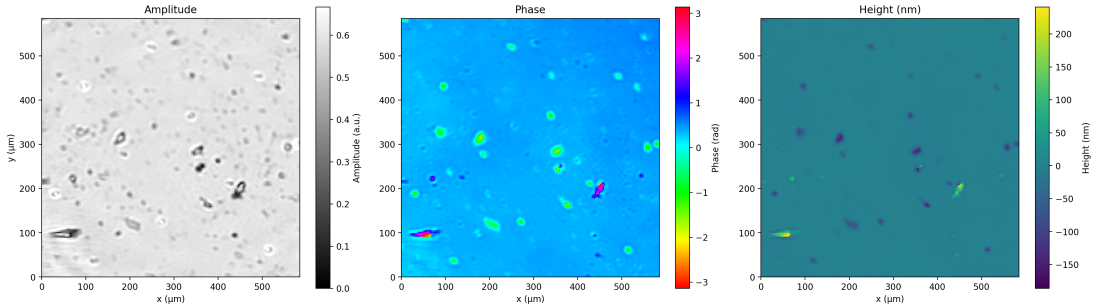

Figure 25: Graphene sample, scan 4 (2026-04-28, 19:57).

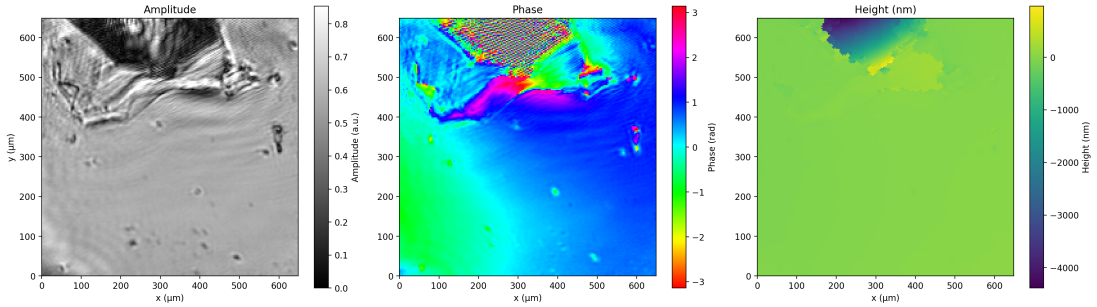

Figure 26: Graphene sample, scan 5 (2026-04-29, 10:02).

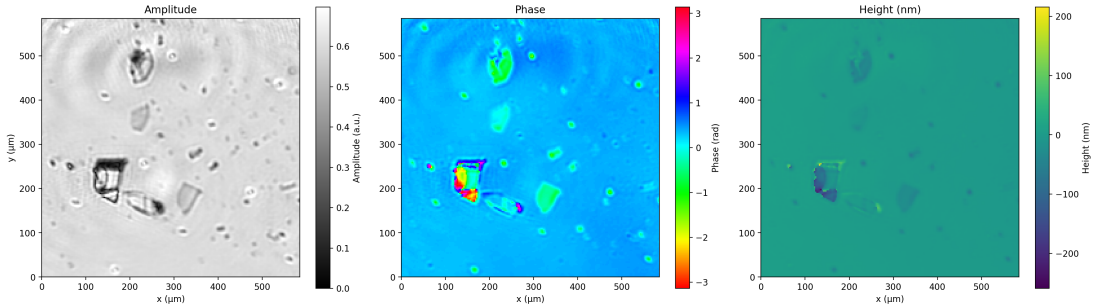

Figure 27: Graphene sample, scan 6 (2026-04-29, 10:19).

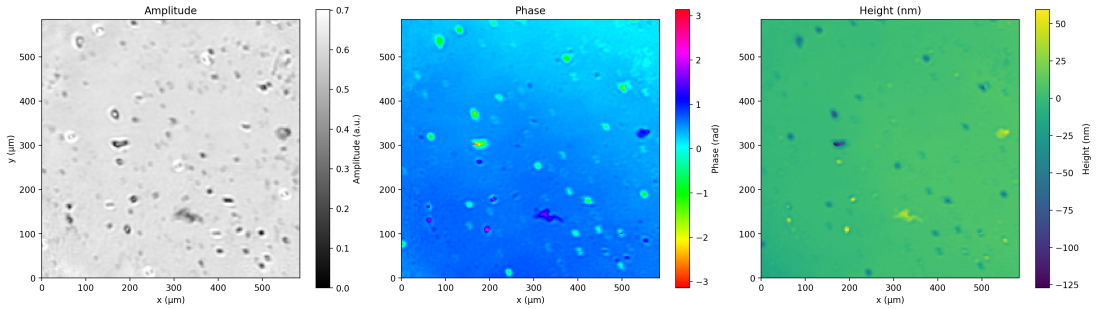

Figure 28: Graphene sample, scan 7 (2026-04-29, 17:27).

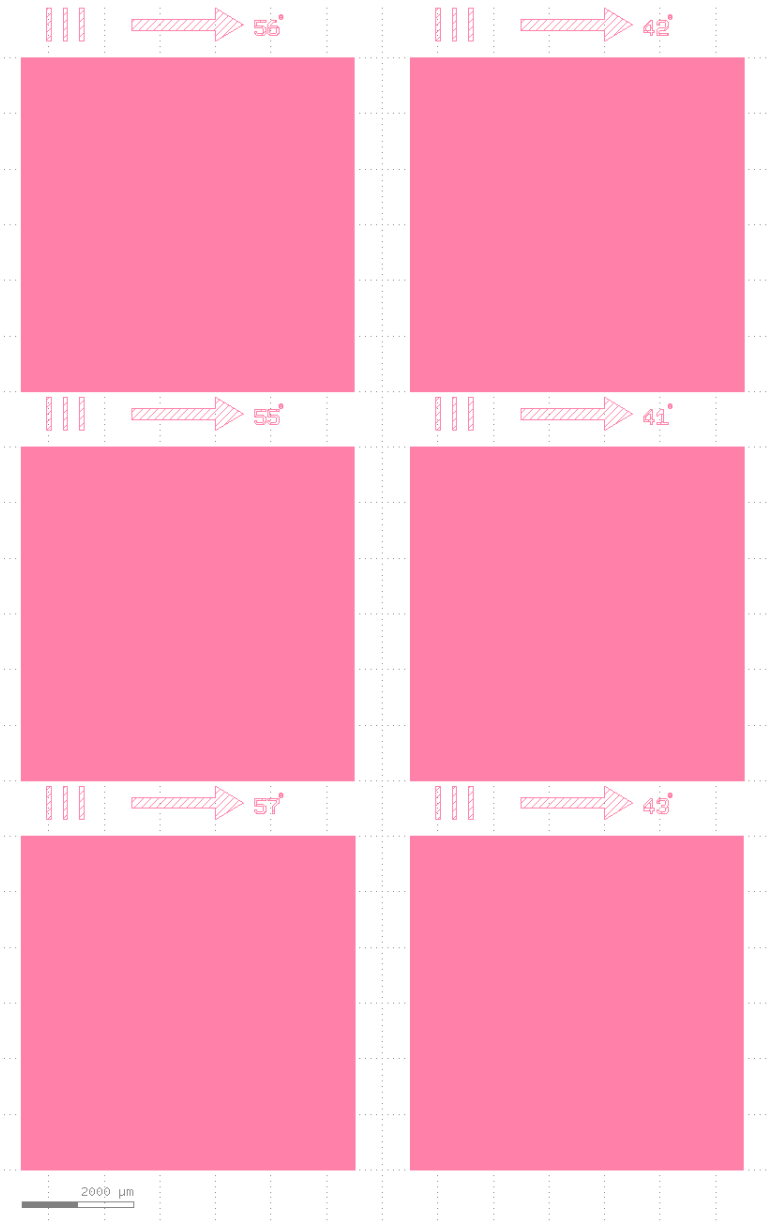

Figure 29: GDS layout of the binary PMMA-on-silicon diffraction gratings. Six tiles are patterned on the same chip for incidence angles of 41°, 42°, 43°, 55°, 56°, and 57°. The primary design tile (56°, top-left) targets the air-PMMA Brewster angle, the 42° tile (top-right) targets the experimental operating angle.

This appendix collects photographs of the cleanroom instruments used in the grating fabrication and metrology described in Chapter 9.

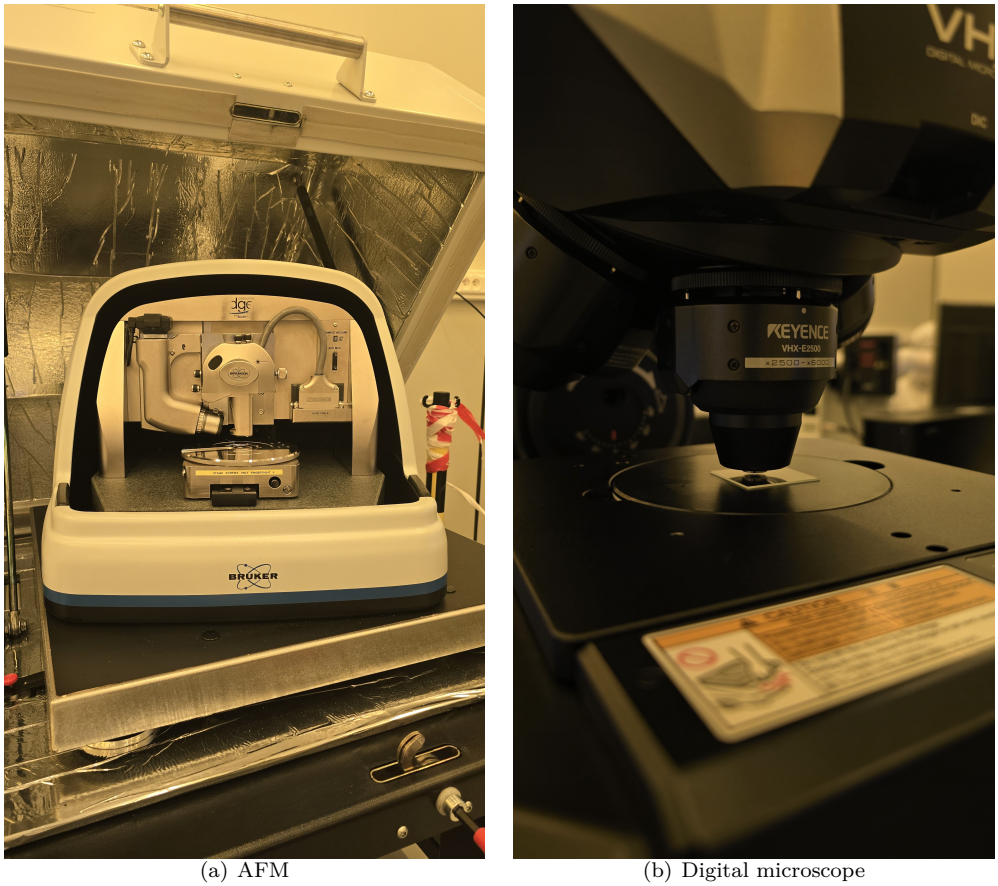

(a) AFM (b) Digital microscope

Figure 30: (a) Atomic force microscope (AFM). (b) Digital microscope used to inspect the grating surface.

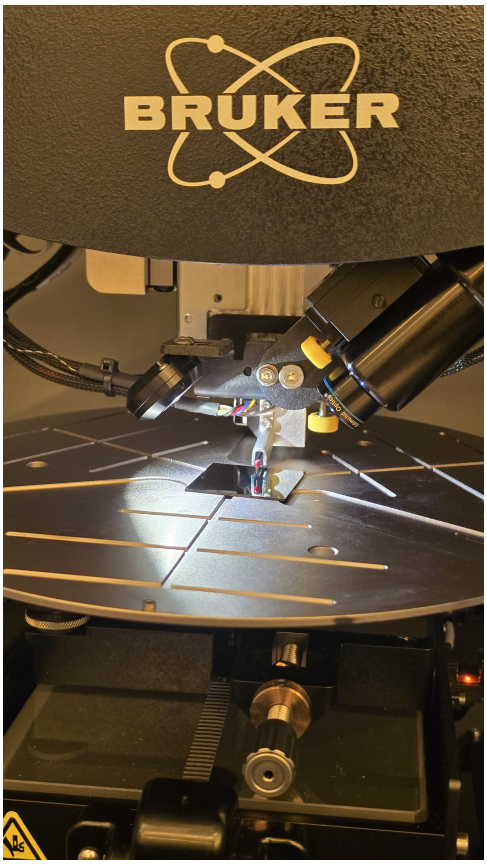

Figure 31: The stylus profilometer used for surface-height measurement of the diffraction grating.