CTF Correction with IMOD

CTF Correction

- When microscope is operated in underfocus to produce phase contrast, the contrast is inverted in some spatial frequency ranges

We See Only Amplitudes, Not Phases, in these Transforms

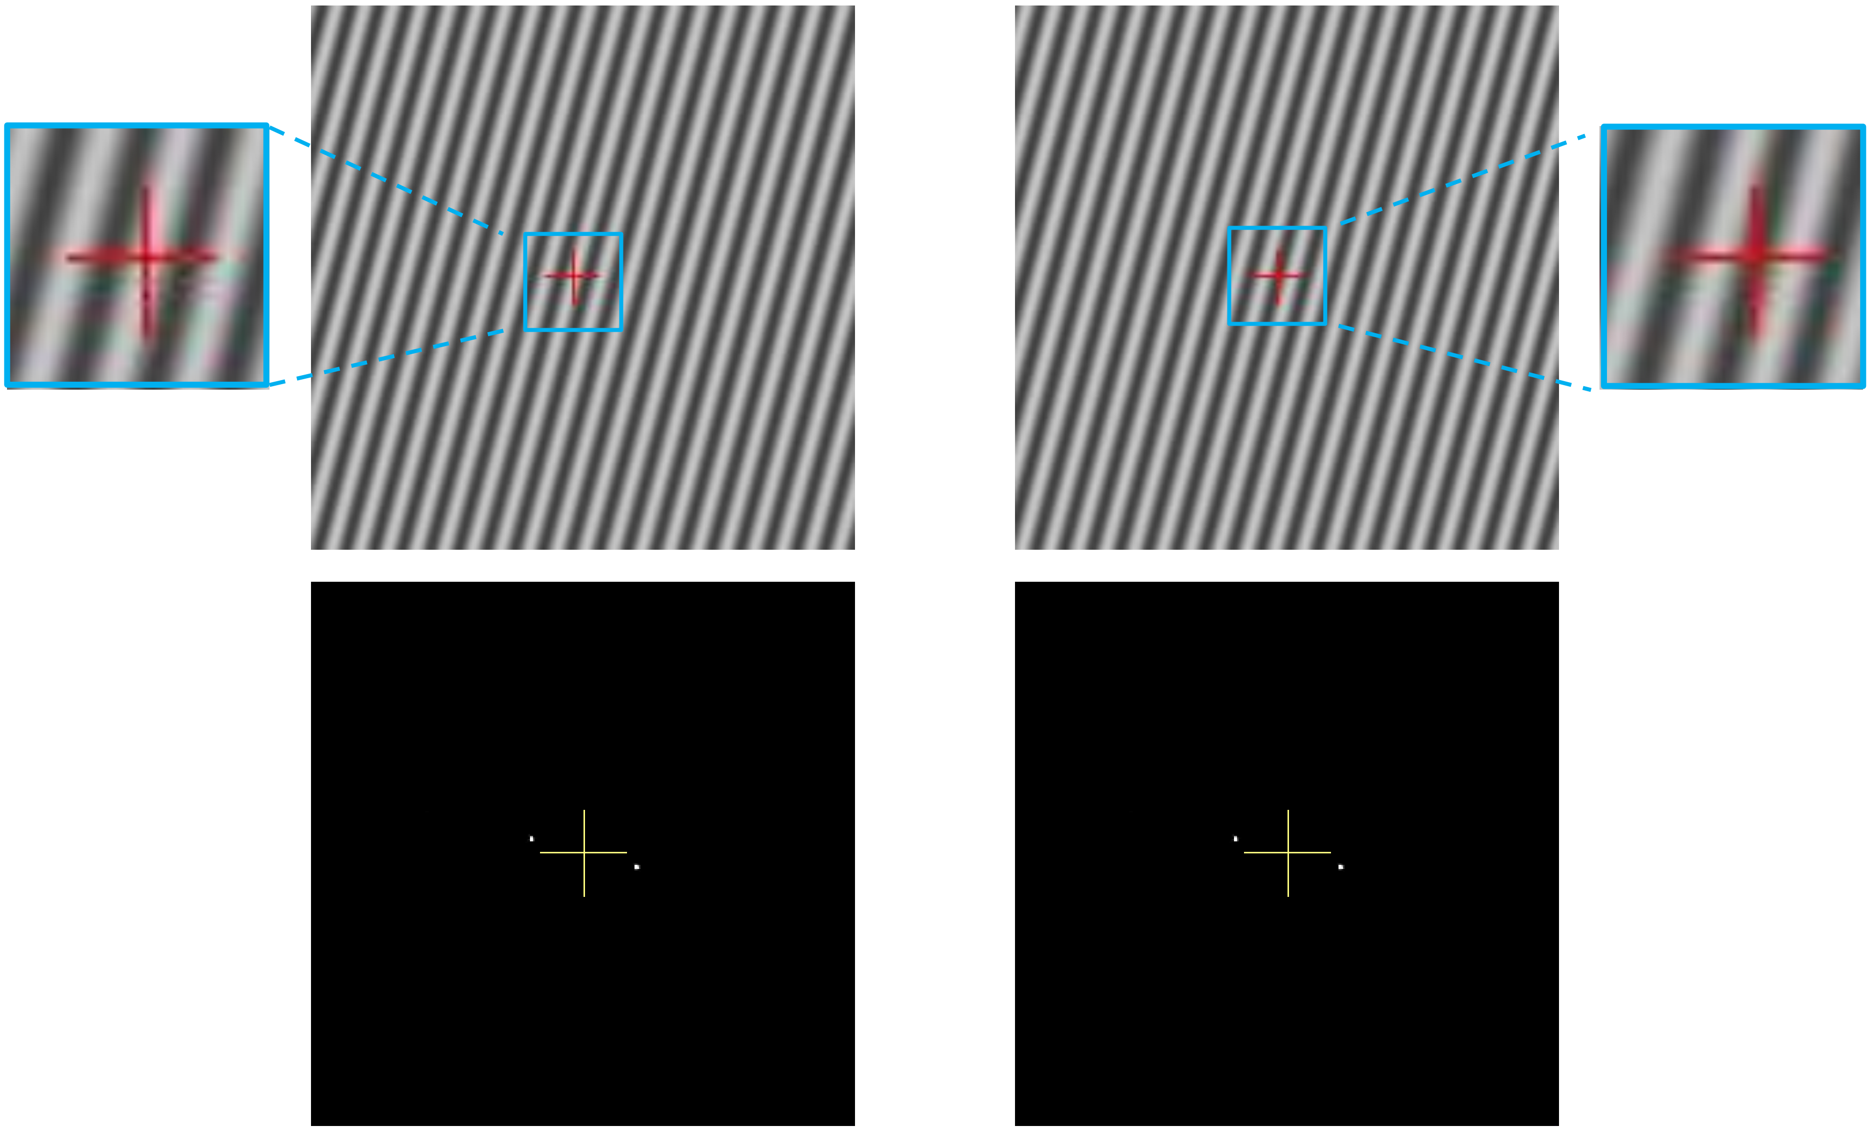

Phase-shifted images with the same power spectrum

Underfocus Enhances Phase Contrast of Low Frequencies

- Most of what we perceive in an image consists of low frequencies
- Thus underfocus enhances the contrast that we see

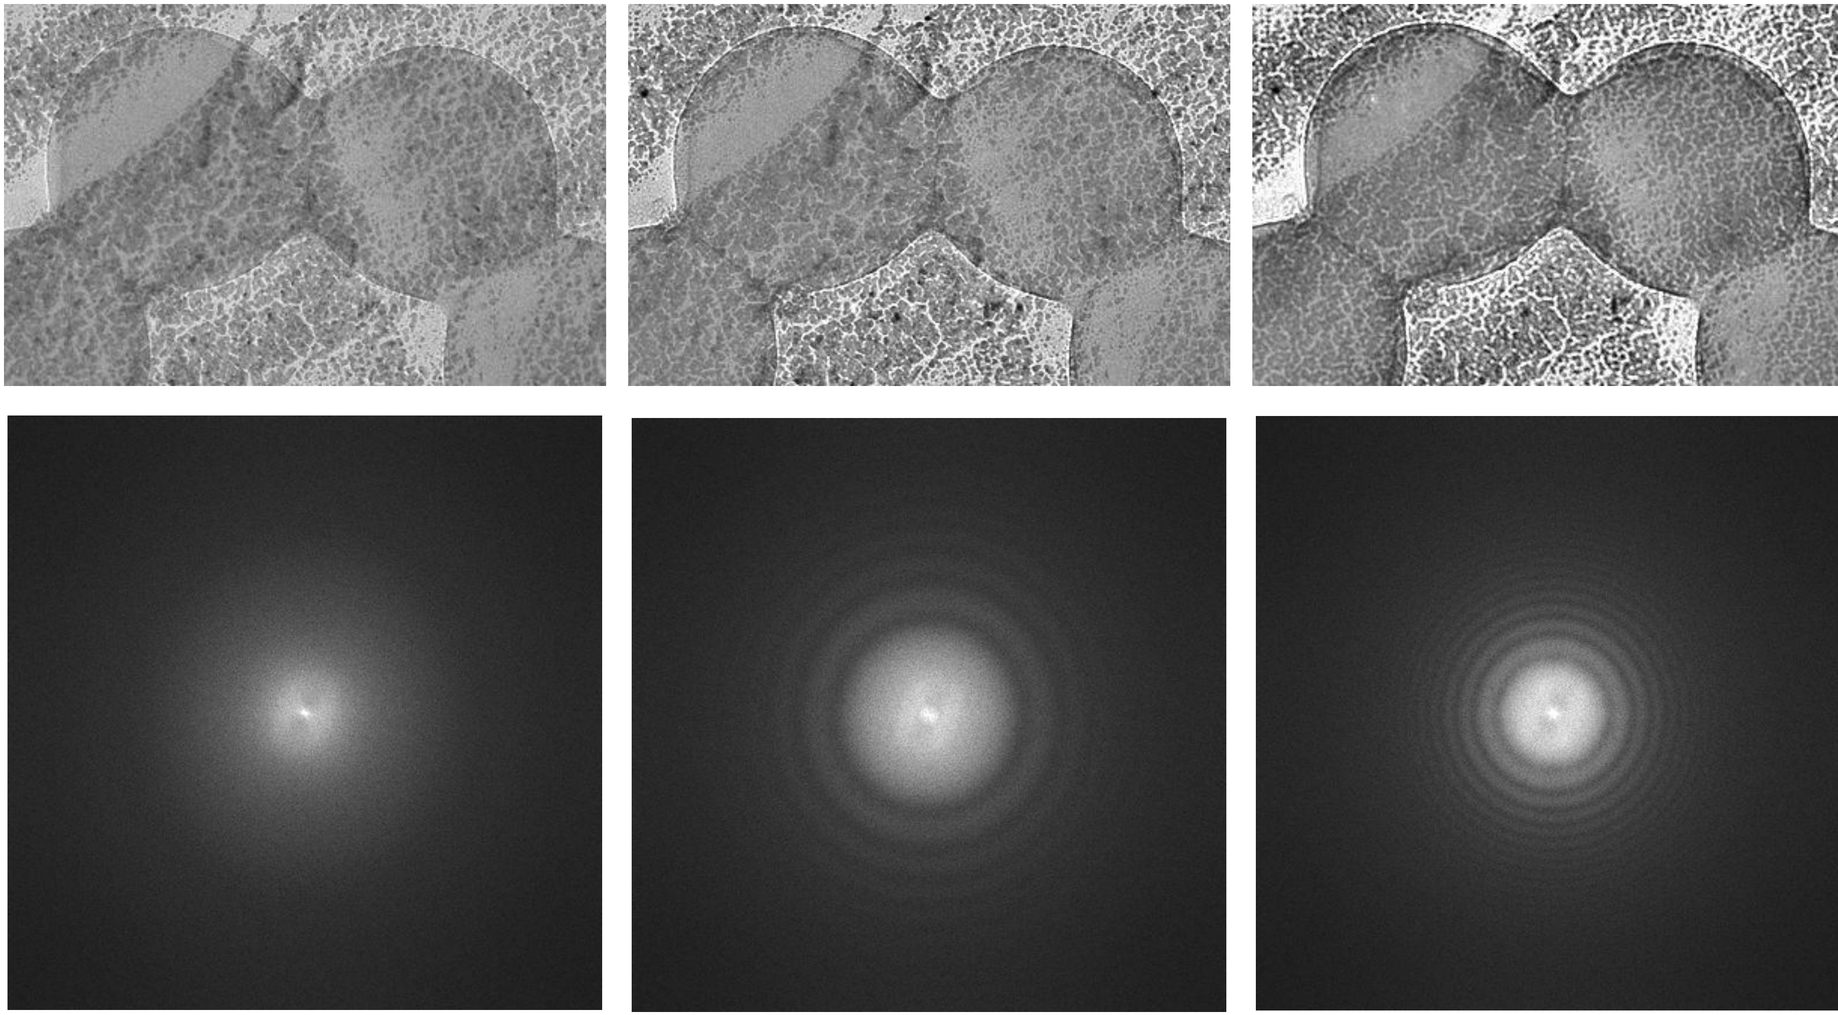

0.4 µm underfocus

2 µm underfocus

6 µm underfocus

Underfocus Also Produces Thon Rings in Power Spectra

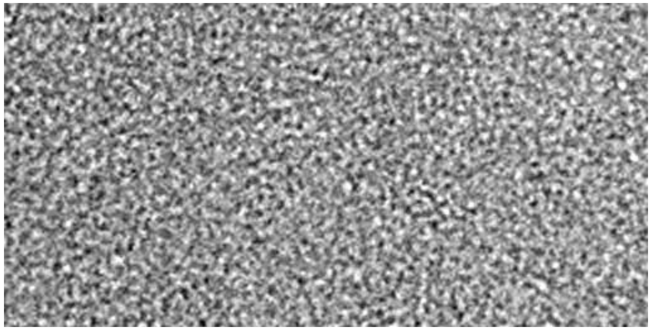

Carbon film

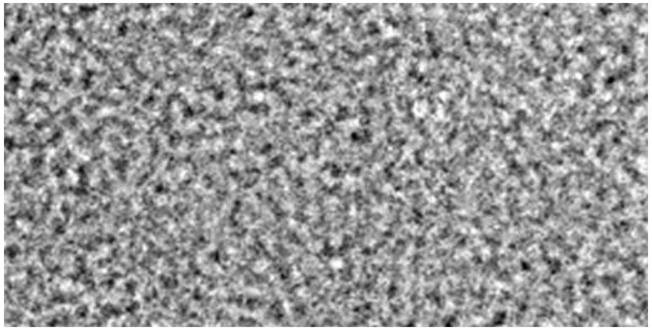

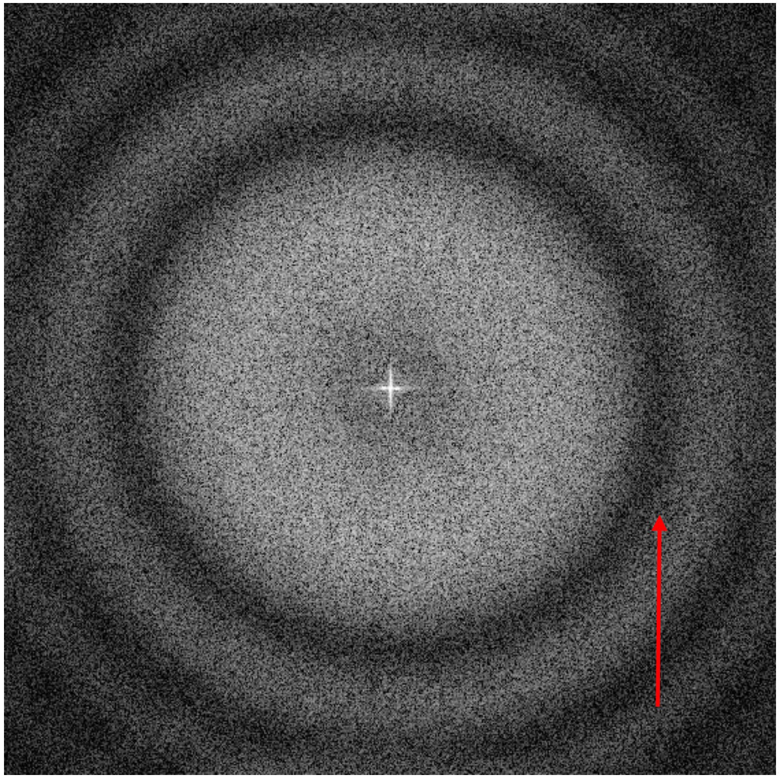

1.5 µm underfocus
First zero at 0.56/nm (1/1.8 nm)

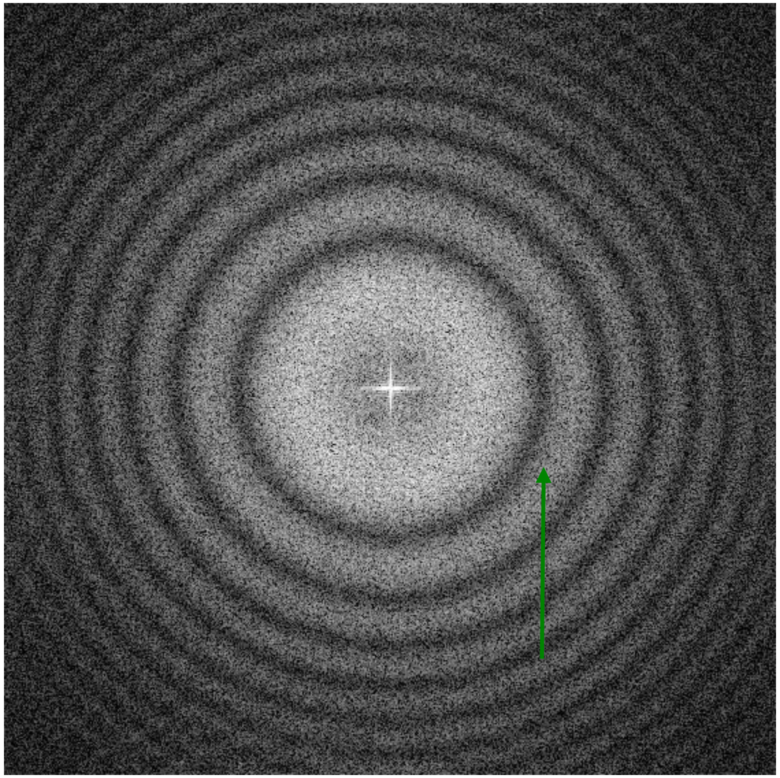

5 µm underfocus
First zero at 0.32/nm (1/3.1 nm)

The CTF Produces Contrast Inversions

- Thon rings from oscillations in the CTF are an intrinsic result of the physics of phase image formation with underfocus
- In fact the CTF goes THROUGH zero, so contrast is inverted for frequencies between the first and second zero, third and fourth zero, etc.

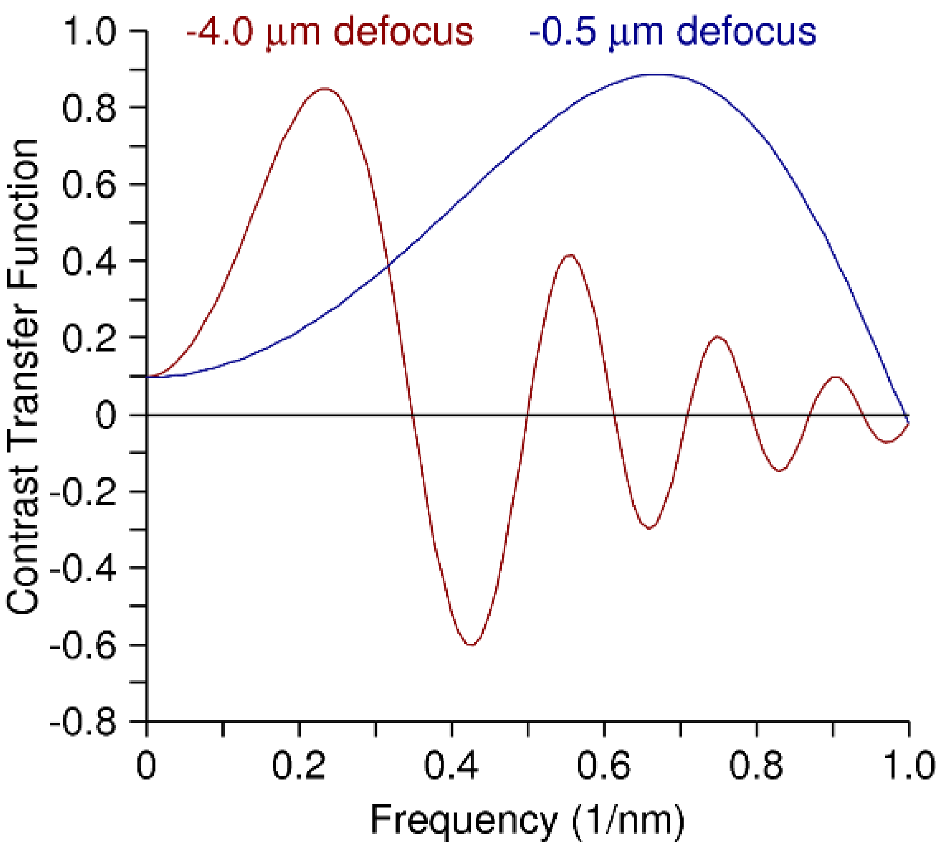

Example:

300 kV FEG

0.5 nm pixel size

Amplitude contrast 0.1

First inversion for -4 µm:

1/2.8 nm to 1/2.0 nm

Phases Must Be Right to See a Correct Image

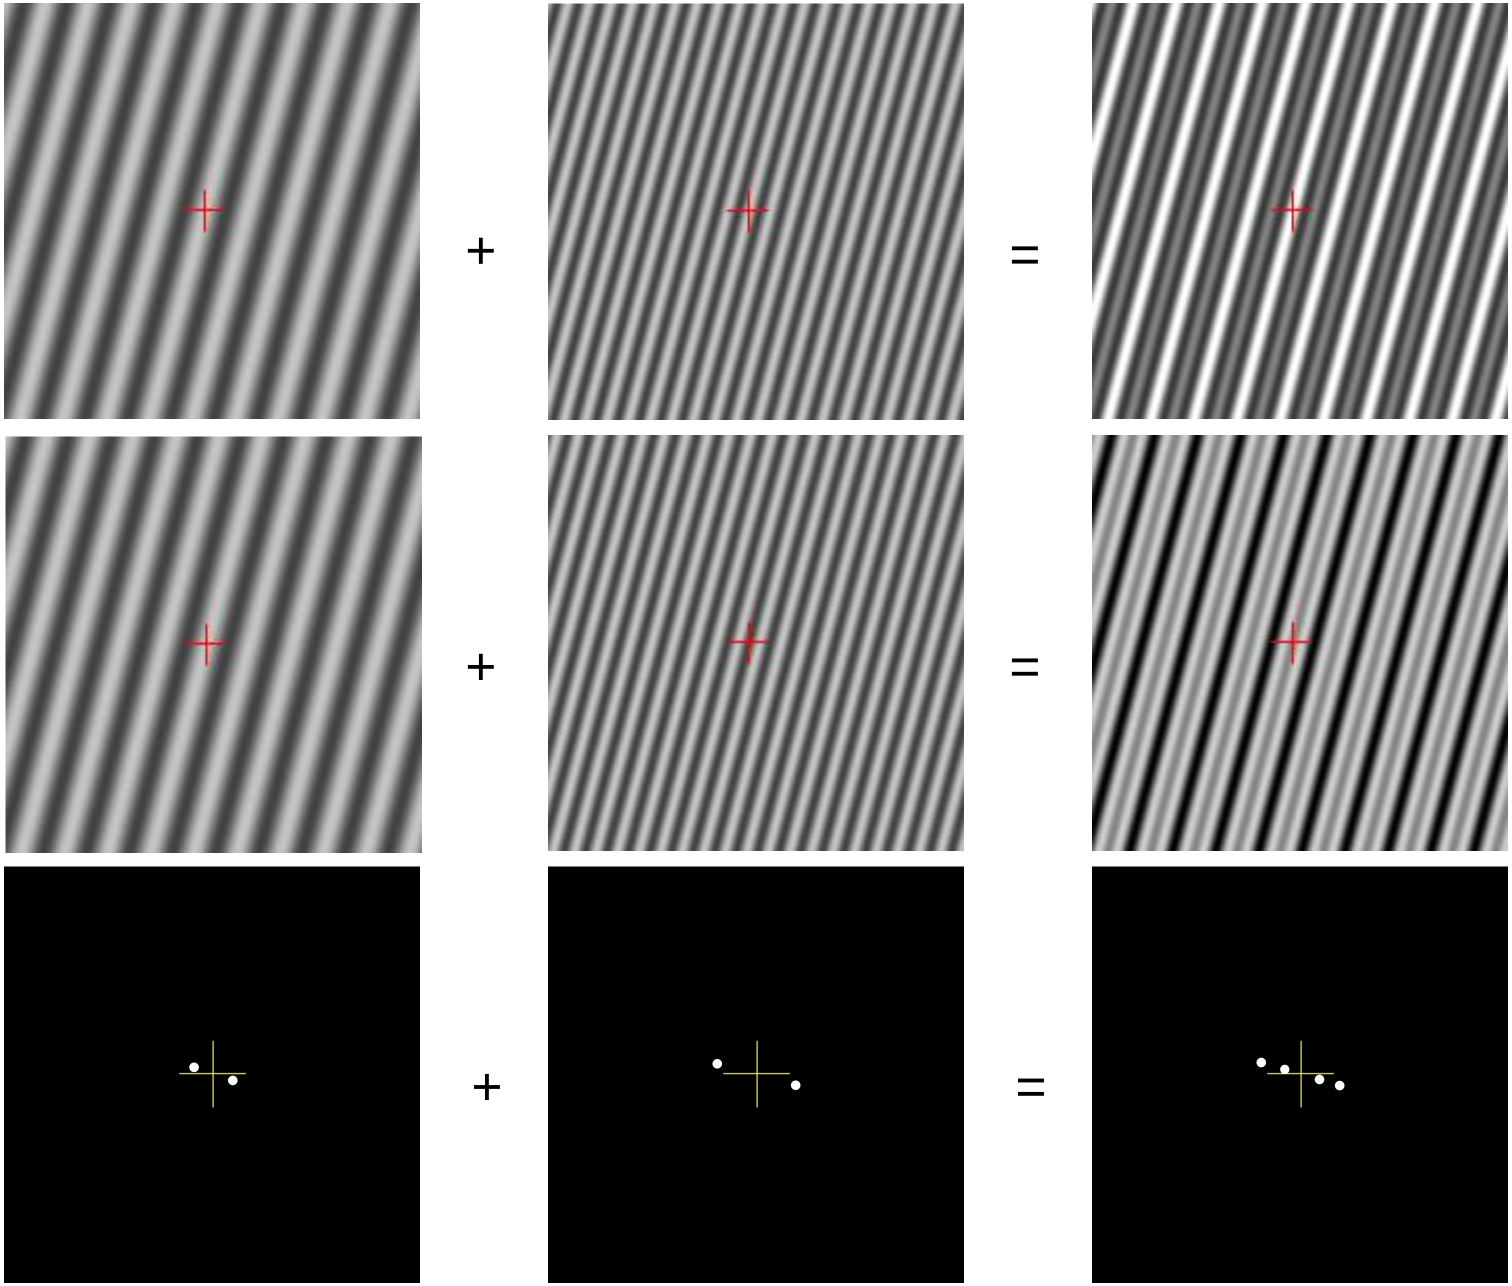

CTF Correction Matters for Subvolume Averaging

- Without CTF correction, information past the first zero is incorrect and will give the wrong structure
- Example from a microtubule decorated with Eg5, taken at -8 µm defocus; the 4-nm tubulin repeat is between 1st and 2nd zero
- Without correction, Eg5 heads have the wrong shape and tubulin is almost 180º out of phase from where it should be

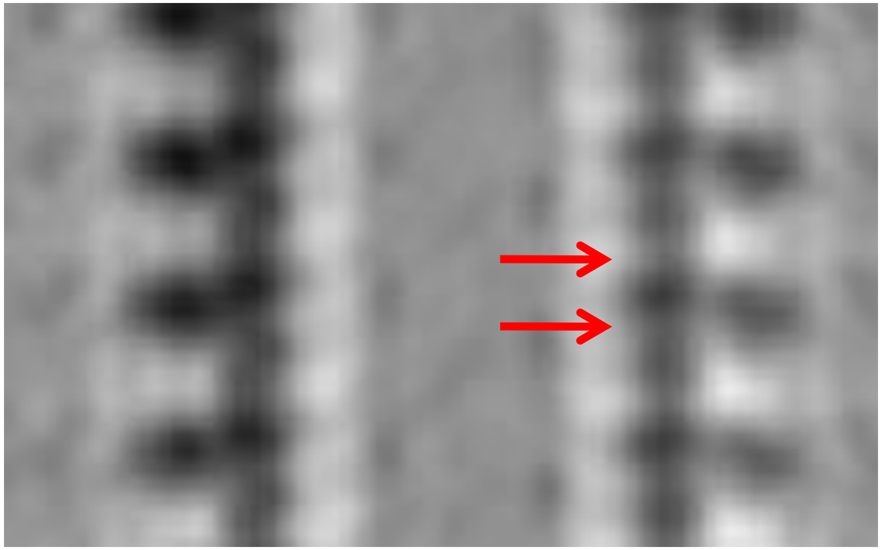

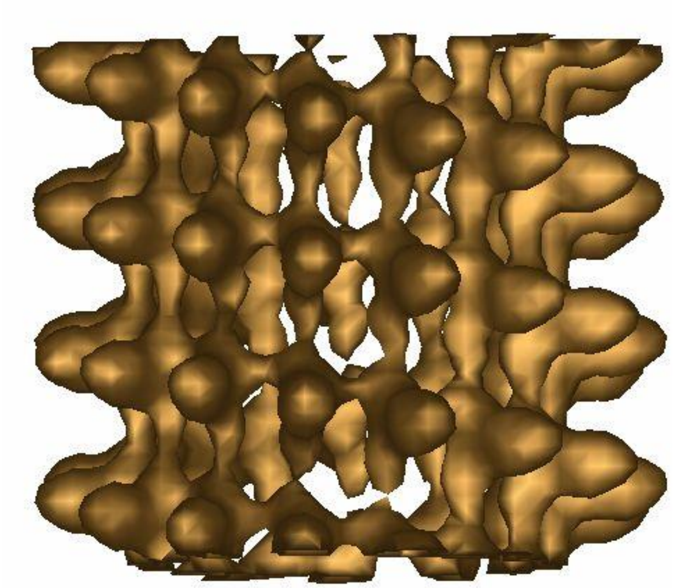

Average from Eg5-decorated microtubule, not corrected

CTF Correction Matters for Subvolume Averaging

- Without CTF correction, information past the first zero is incorrect and will give the wrong structure
- Example from a microtubule decorated with Eg5, taken at -8 µm defocus; the 4-nm tubulin repeat is between 1st and 2nd zero
- Without correction, Eg5 heads have the wrong shape and tubulin is almost 180º out of phase from where it should be

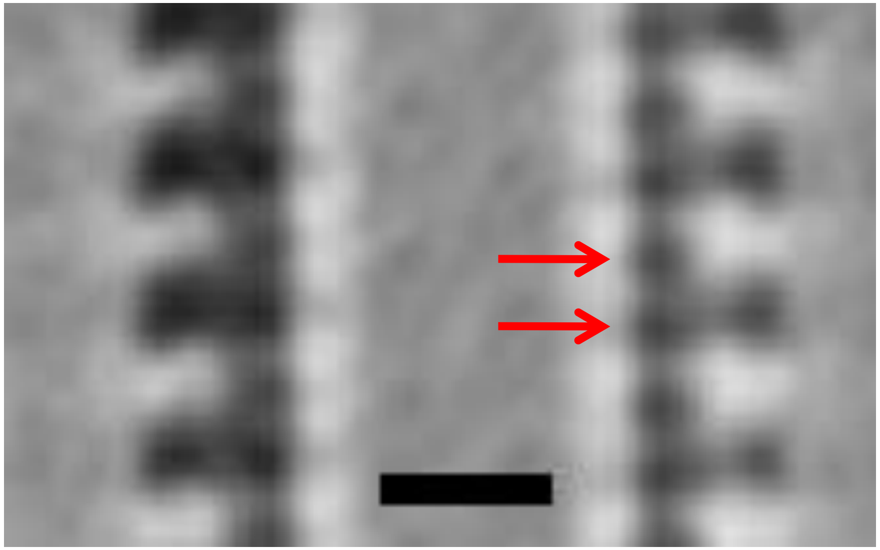

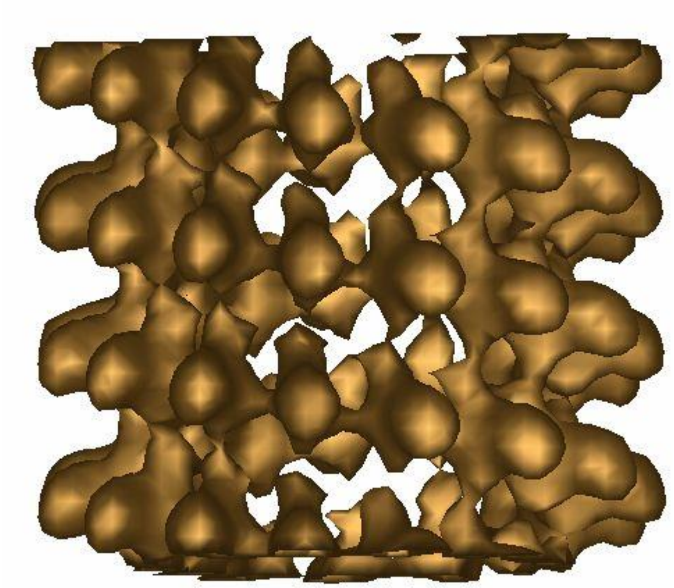

Average from Eg5-decorated microtubule, CTF corrected

What Makes CTF Correction of Tilt Series Hard?

- There is a different defocus gradient across each tilt series image
  - Not straightforward to invert phases in an FFT
  - Adding together data from different defoci will blur the power spectrum, make it hard to detect zeros

Defocus and Thon Ring Variation in Tilted Sample

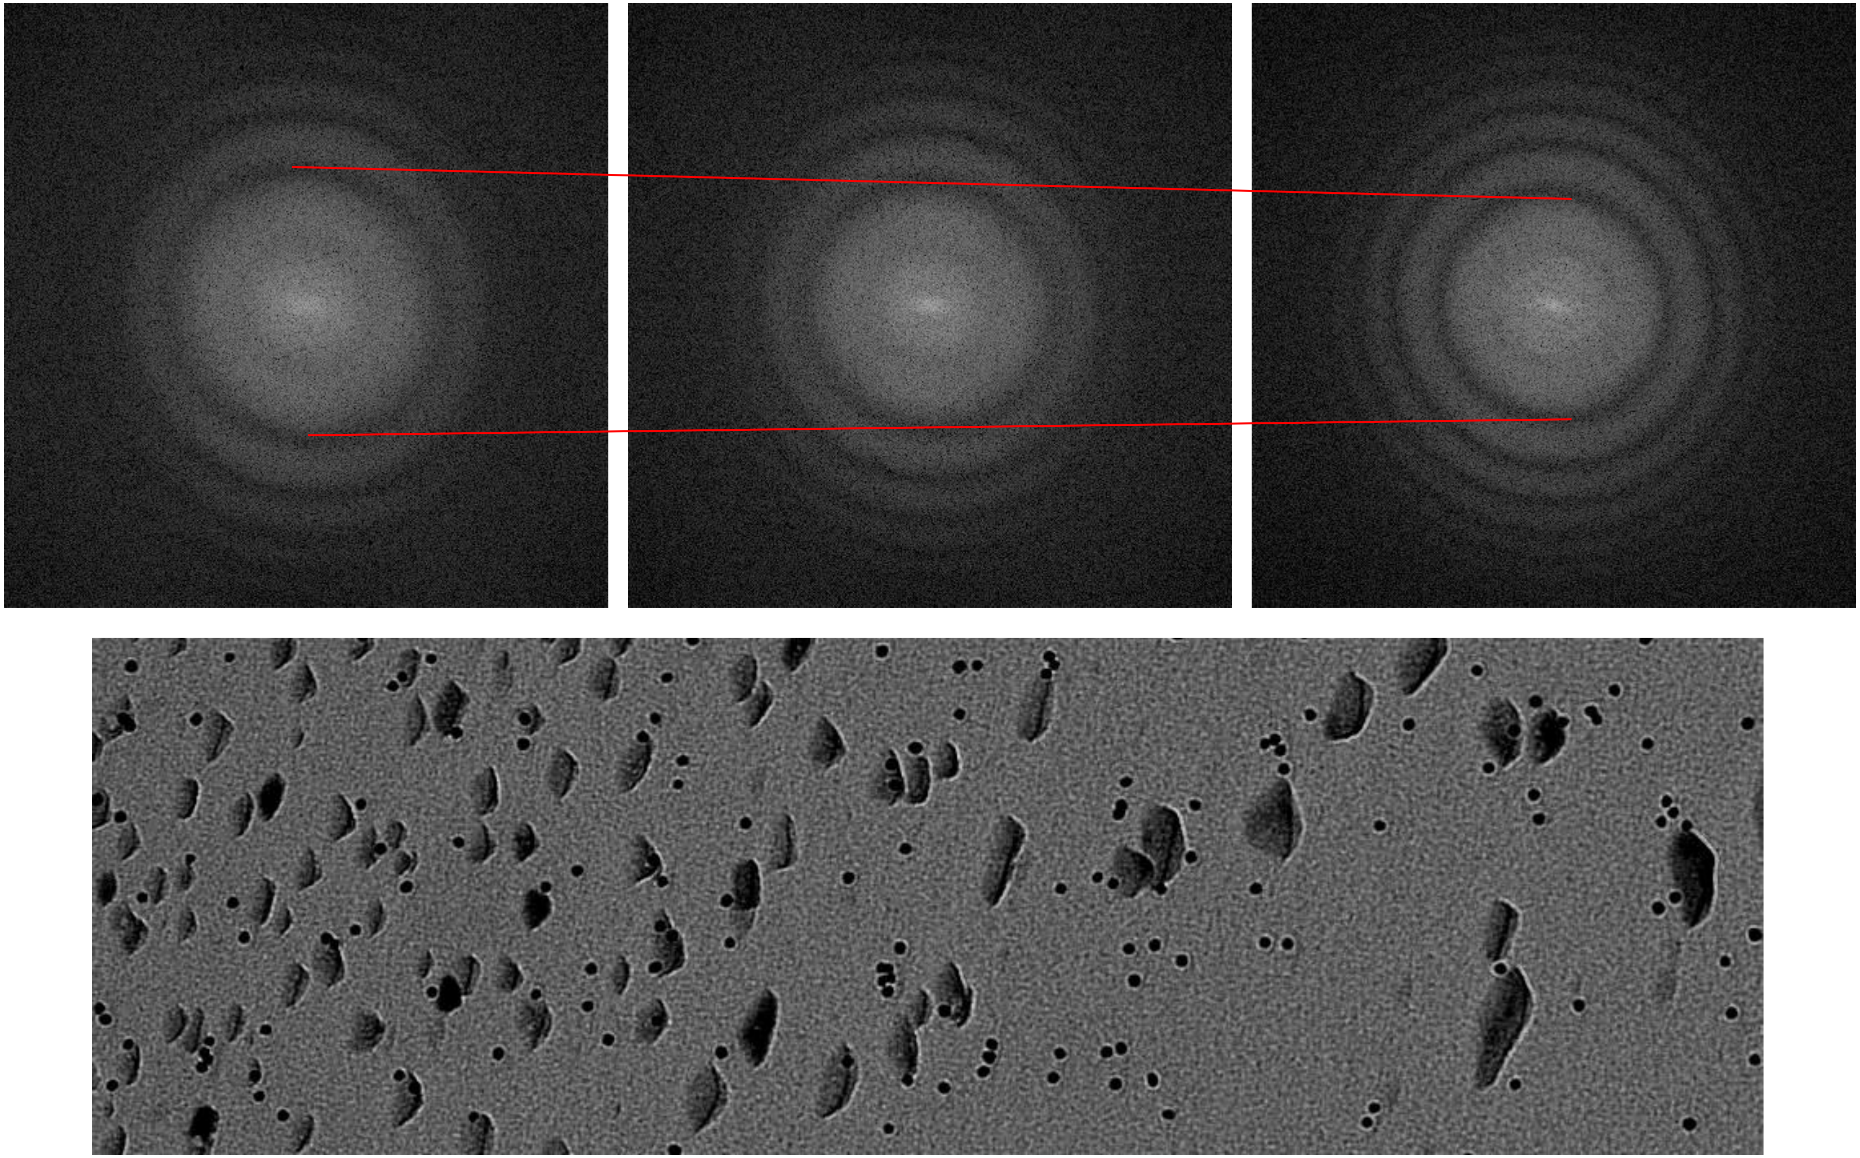

5 nm gold on carbon film, tilted to 65°

4.9 µm underfocus | 6 µm underfocus | 7.1 µm underfocus

What Makes CTF Correction of Tilt Series Hard?

- There is a different defocus gradient across each tilt series image
  - Not straightforward to invert phases in an FFT
  - Adding together data from different defocuses will blur the power spectrum, make it hard to detect zeros
- The dose is ~10-20 fold lower per micrograph than for single-particle averaging
  - Before direct detectors, data had to be averaged from multiple images to see CTF effect in power spectrum, unless the defocus was large
- There may be little information past the first few zeros
  - This depends on relationship between pixel size and defocus
  - For cryoET with relatively low defocus, there may be few zeros in spectrum, unless very high resolution is being targeted

Rotational Averaging and Background Subtraction Are Essential for Seeing CTF Zeros in Tilt Series

(An extreme case, CCD camera data)

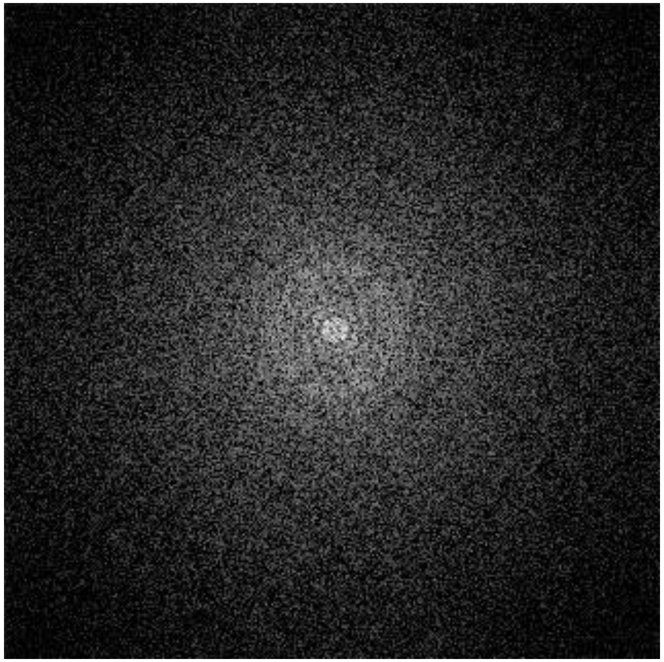

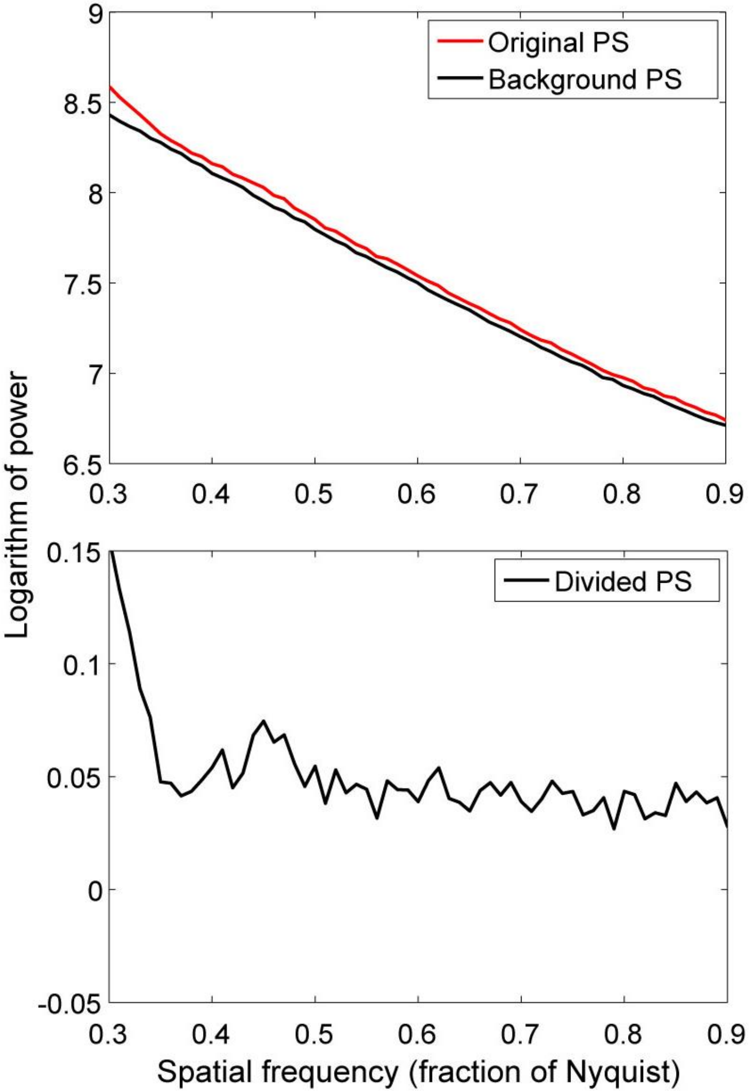

Rotational Averaging and Background Subtraction Are Essential for Seeing CTF Zeros in Tilt Series

- The Ctfplotter program originally used noise images to estimate the background
  - See Ctfplotter man page for instructions
- These should be unnecessary with electron counting data

Background Subtraction Is Possible without Noise Images using a Two-stage Procedure

- First stage: the program fits two straight lines to the low and high frequency ends of the spectrum
  - It finds the frequency where the slopes of the lines fit before and after that point differ the most
  - It subtracts the high-frequency line from the whole spectrum to make the high-frequency part flat on average

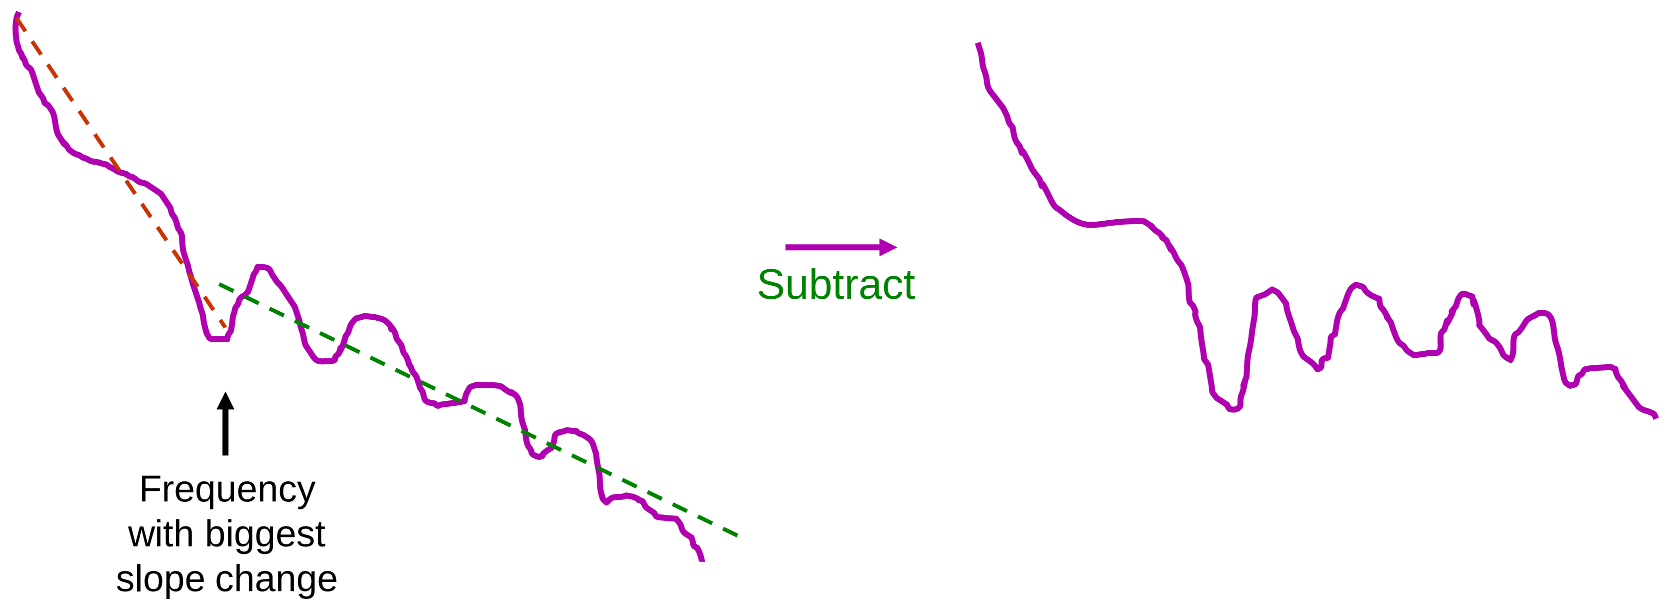

Background Subtraction Is Possible without Noise Images using a Two-stage Procedure

- Second stage: a concave polynomial (concave up or down) is fit to the spectrum from the first zero onwards
  - It fits to minima if it can find them
  - Polynomial is subtracted to flatten the spectrum
  - This “baseline fitting” is available after noise subtraction too
- The two stages work well together with electron counting data, may work for other direct detector data

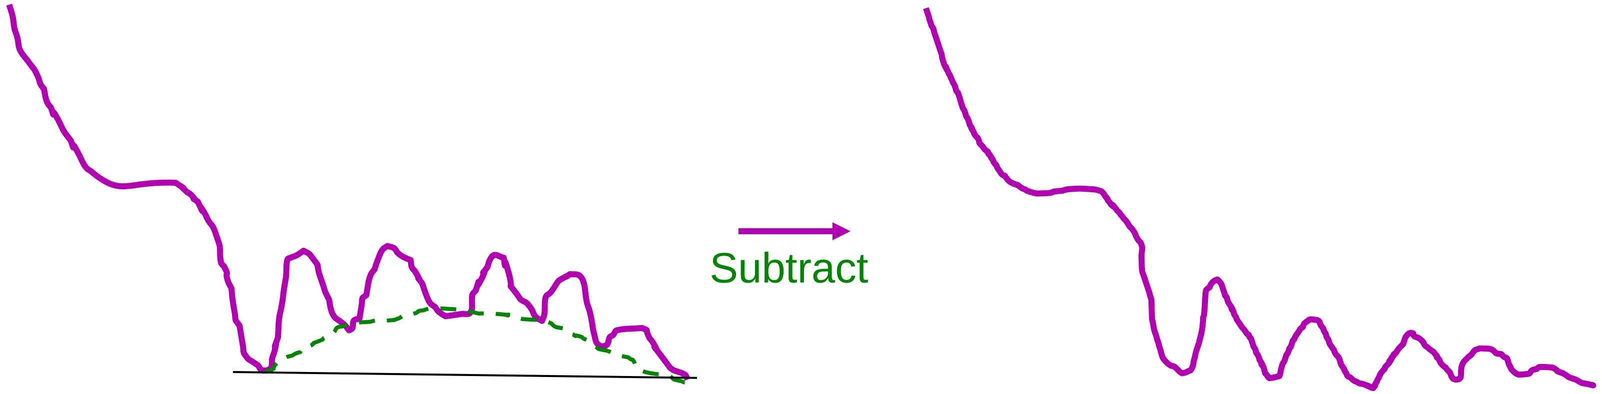

To Compute Power Spectra, Each Image Is Divided into Overlapping Tiles

- This is periodogram averaging, standard method used to get power spectra for CTF correction
  - Tiles are typically 256 pixels square and overlap by 50%
  - Fourier transforms are taken separately and averaged

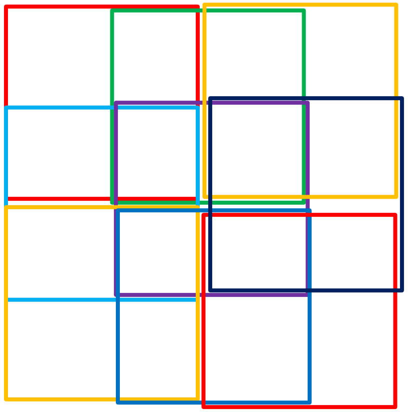

Overlapping tiles

Adding Together Spectra from Different Defoci

- Power spectra from central tiles are simply added together
- Spectra from off-center tiles are scaled so that the zeros line up when they are added into spectra from central tiles
  - Without this, the zeros will blur out
  - And stronger signal from one side could bias the defocus

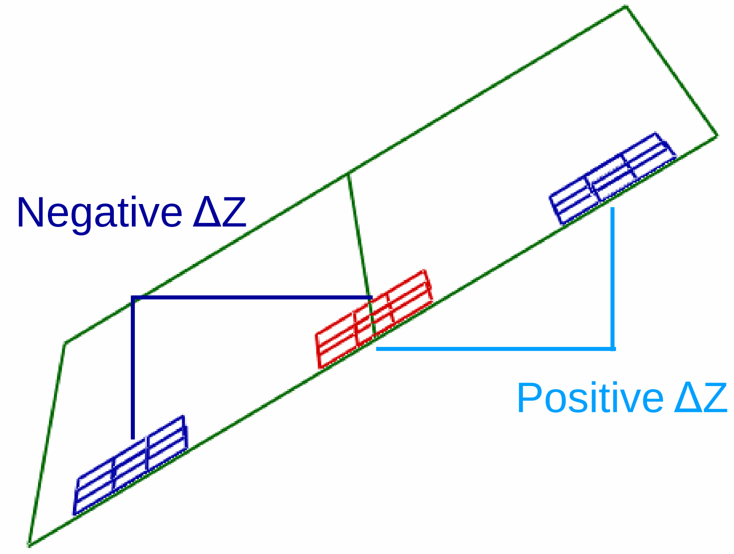

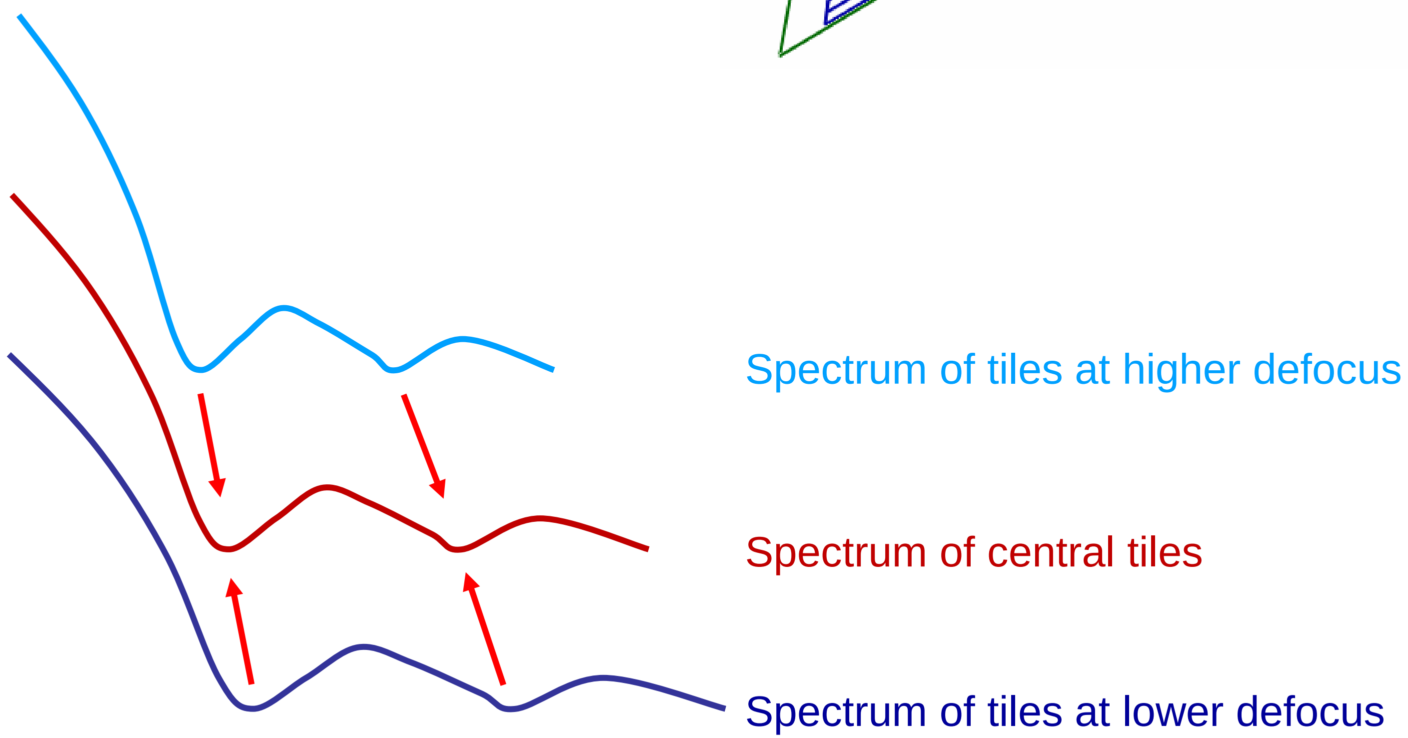

Philosophy of Ctfplotter

- Fit automatically when possible, but use interactive tools to improve the result when necessary
- Fit to data from a single view when this is accurate enough, but fit to data summed over multiple views when necessary
  - If signal is not strong enough to allow reliable fitting, fitting to every view will do more harm than good
  - Fitting to multiple views can correct for a trend in defocus through the series, which can be up to ~1 µm, even if view-to-view variations cannot be corrected
- Apply summing differently for different parameters being found

Philosophy of Ctfplotter

- Fit automatically when possible, but use interactive tools to improve the result when necessary
- Fit to data from a single view when this is accurate enough, but fit to data summed over multiple views when necessary
  - If signal is not strong enough to allow reliable fitting, fitting to every view will do more harm than good
  - Fitting to multiple views can correct for a trend in defocus through the series, which can be up to ~1 µm, even if view-to-view variations cannot be corrected
- Apply summing differently for different parameters being found
- What parameters?
  - Defocus
    - Always highest priority for fitting to single views
  - Astigmatism – variation in defocus with direction in image
    - Expected to vary slowly through series, summing over views much more acceptable
  - Phase shift when imaging with phase plate
    - Will build up slowly with exposure

Example of Fitting to Single Image at -41° Tilt

- The IMOD tutorial data set from a K2 camera is useful to work with, despite being unusual
  - Dose here was only ~0.5 $e^-/Å^2$

Tilt series of Giardia flagellum, taken by Cindi Schwartz on Krios at Janelia Farm
±60° @ 2°,
~30 $e^-/Å^2$ total dose
0.5-sec Super-resolution exposures reduced by 4 to 0.67 nm pixel

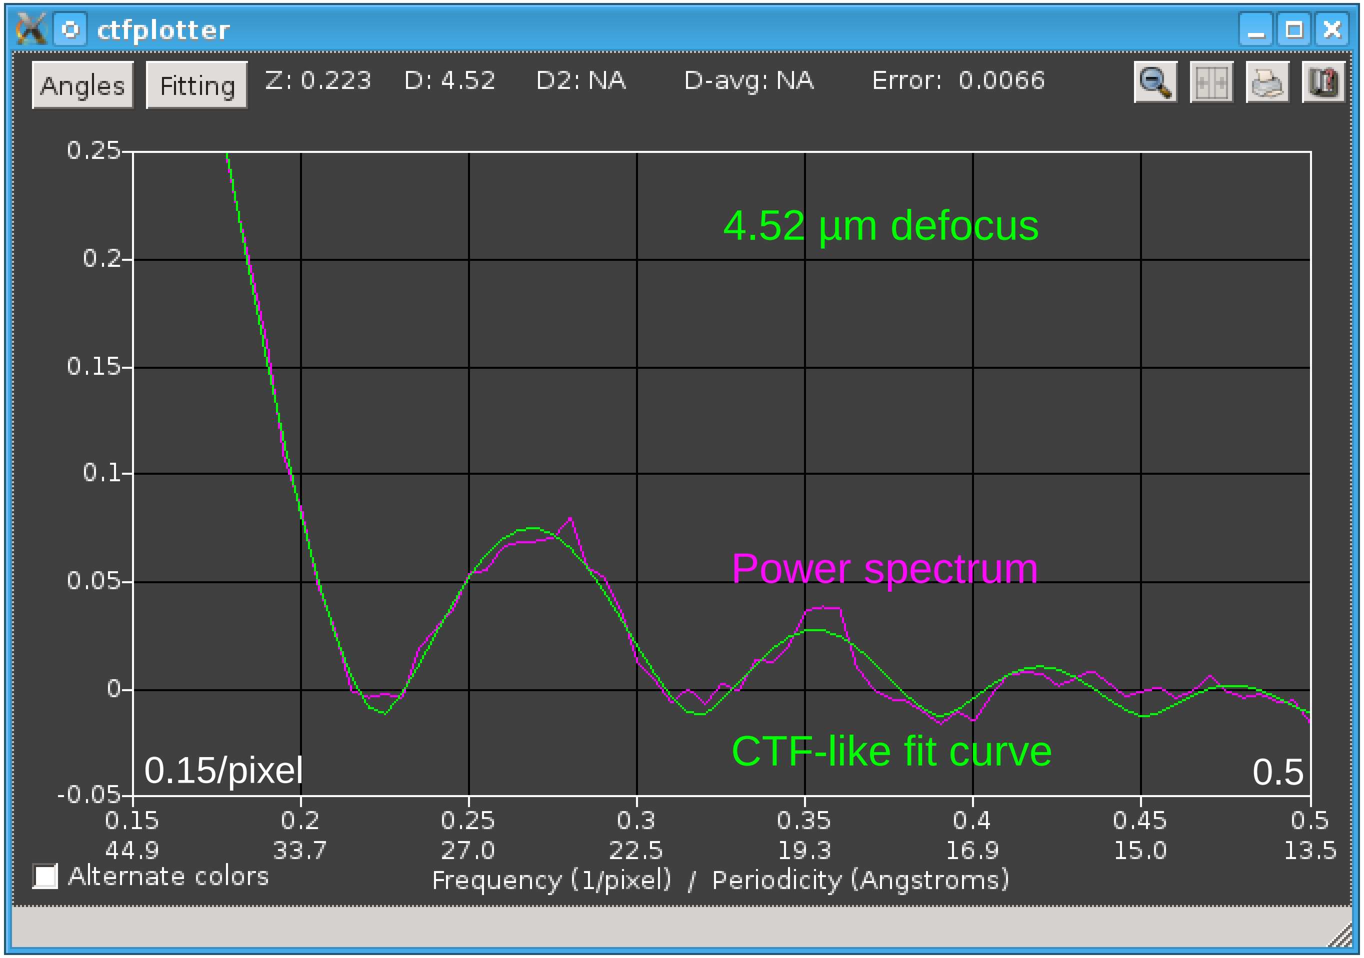

Rotational Averaging from Wedges in Spectrum

- Averaging over restricted angular range (60°) shows clear changes in defocus that can still be measured well

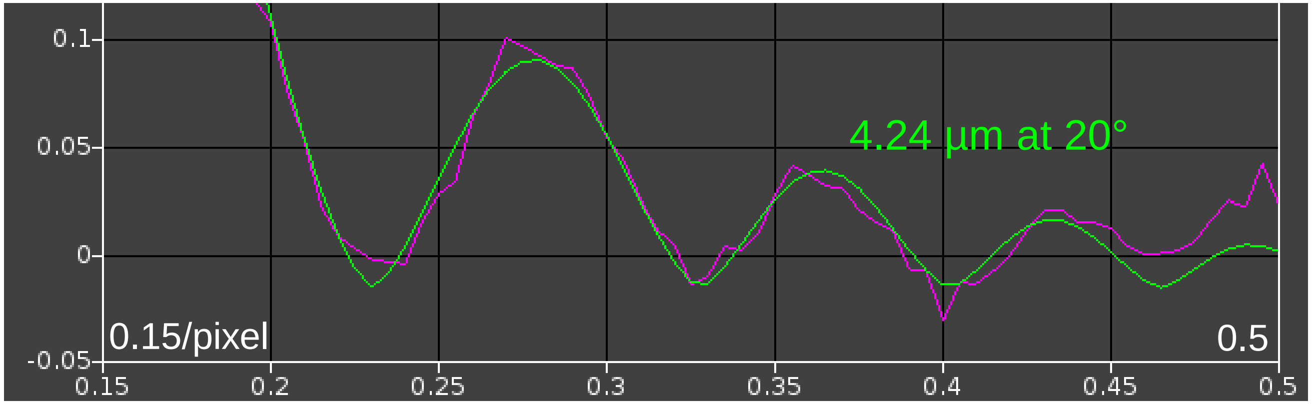

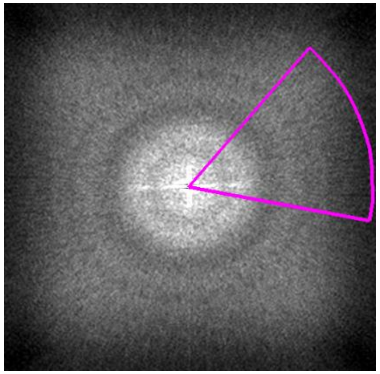

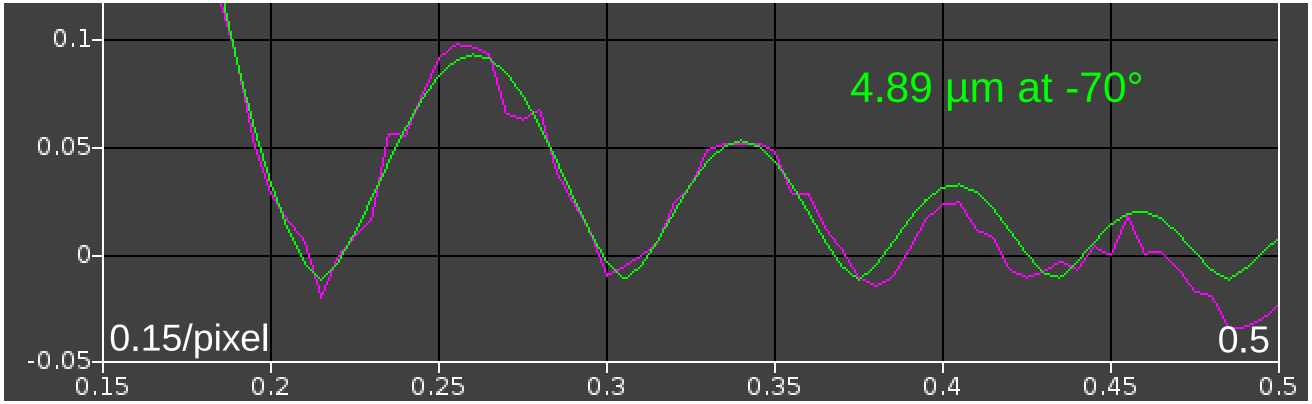

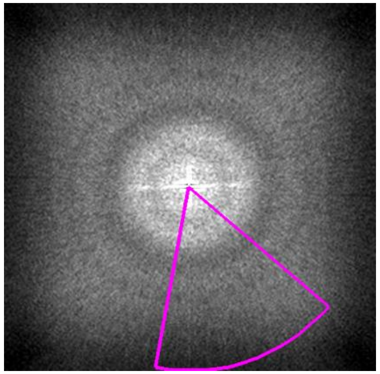

Astigmatism Can Be Determined from Wedge Spectra

- Measured defocus depends on wedge extent
- Use search to find underlying defocus and astigmatism that reproduce the measurements when averaged over the wedge extent
- This works: these wedge sizes give the same estimate

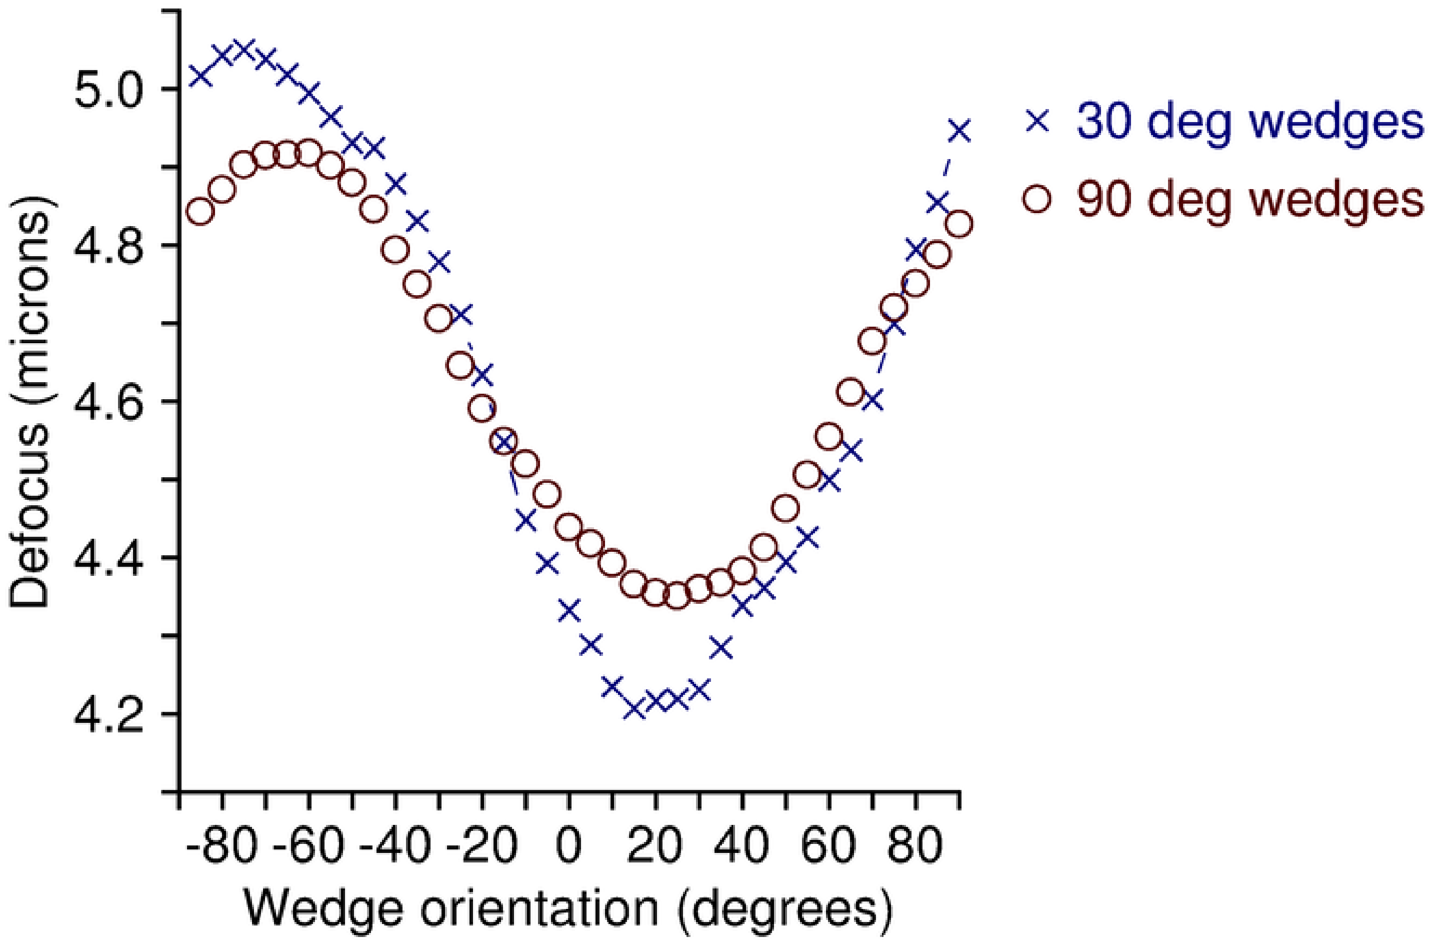

Adjusting for Astigmatism Gives Better Spectrum

- Astigmatism is found by measuring defocus in 36 overlapping wedge spectra
- To get sum for final defocus estimate, narrower wedge spectra are scaled based on their expected defocus from the astigmatism estimate

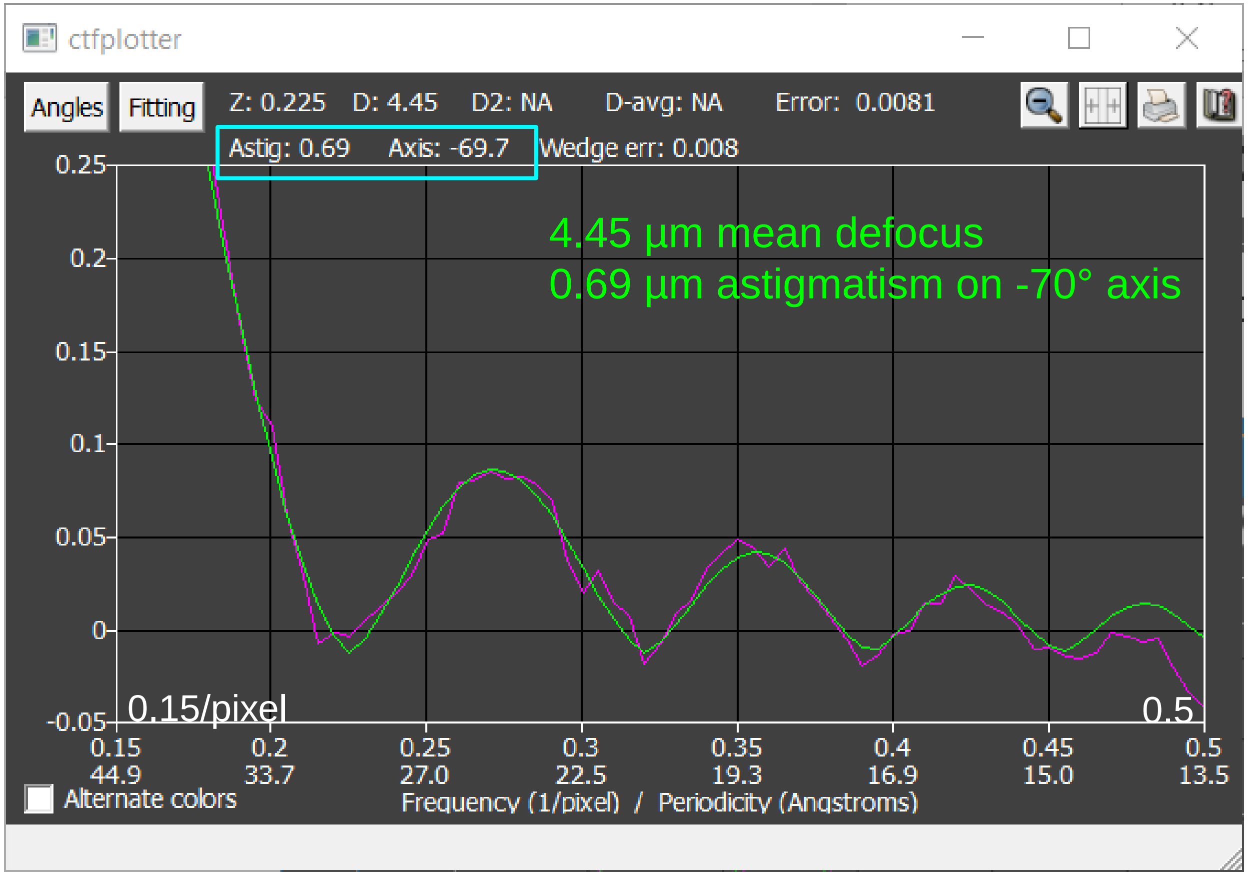

Adjusted for astigmatism

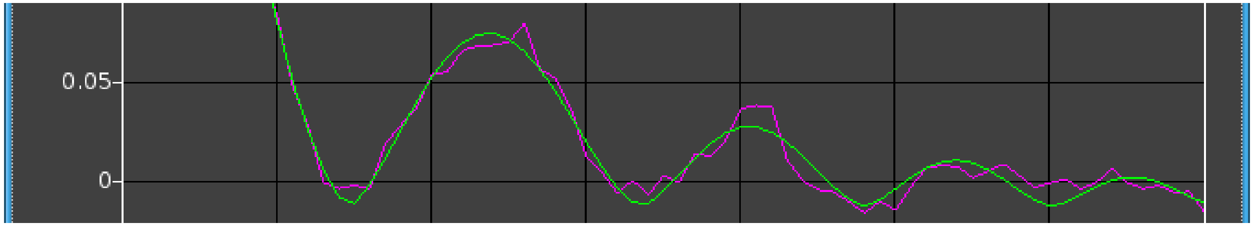

Unadjusted

Measuring Astigmatism from Wedge Analysis is More Reliable than Fitting to 2-D Spectra with Ctffind4

- Ctffind4 4.1.8 badly underestimates astigmatism in these images
  - The consistent error in defocus seems to be a side-effect of that

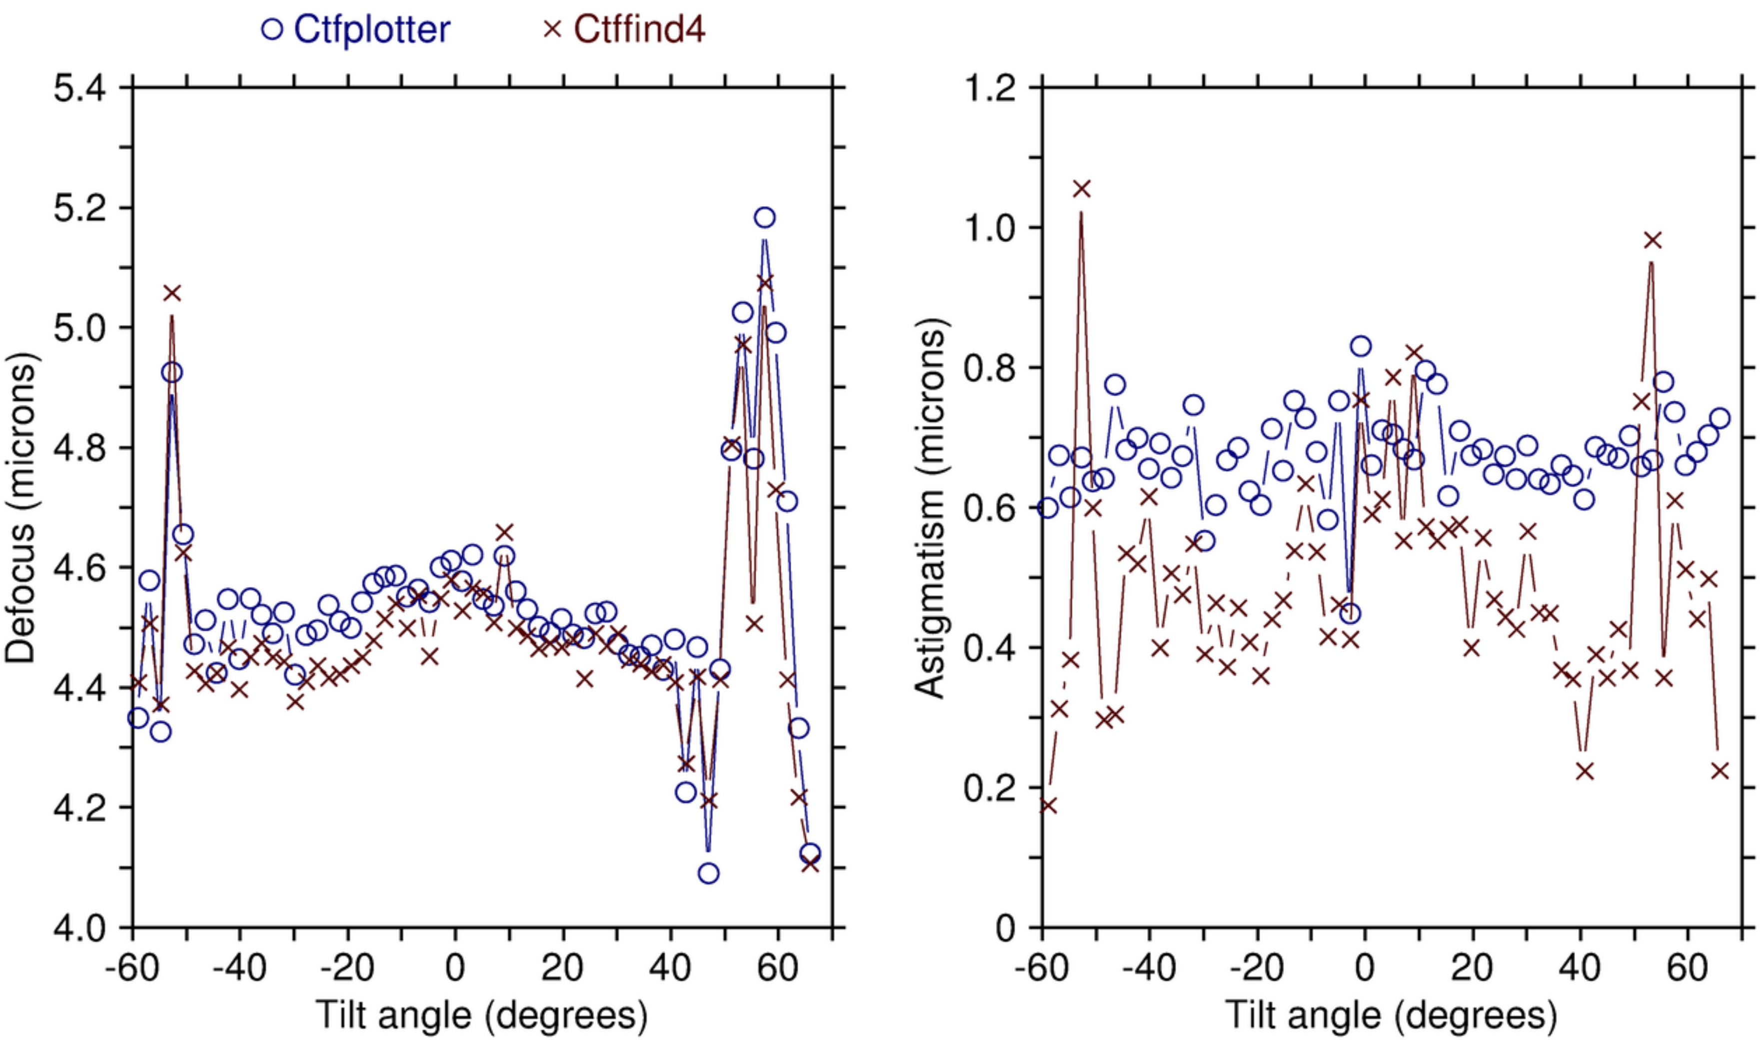

Measuring Astigmatism from Wedge Analysis is More Reliable than Fitting to 2-D Spectra with Gctf

- Gctf does better than Ctffind4 on astigmatism but still has a lot of variability
- Defocus is also closer, but seems to have a tilt-dependent bias

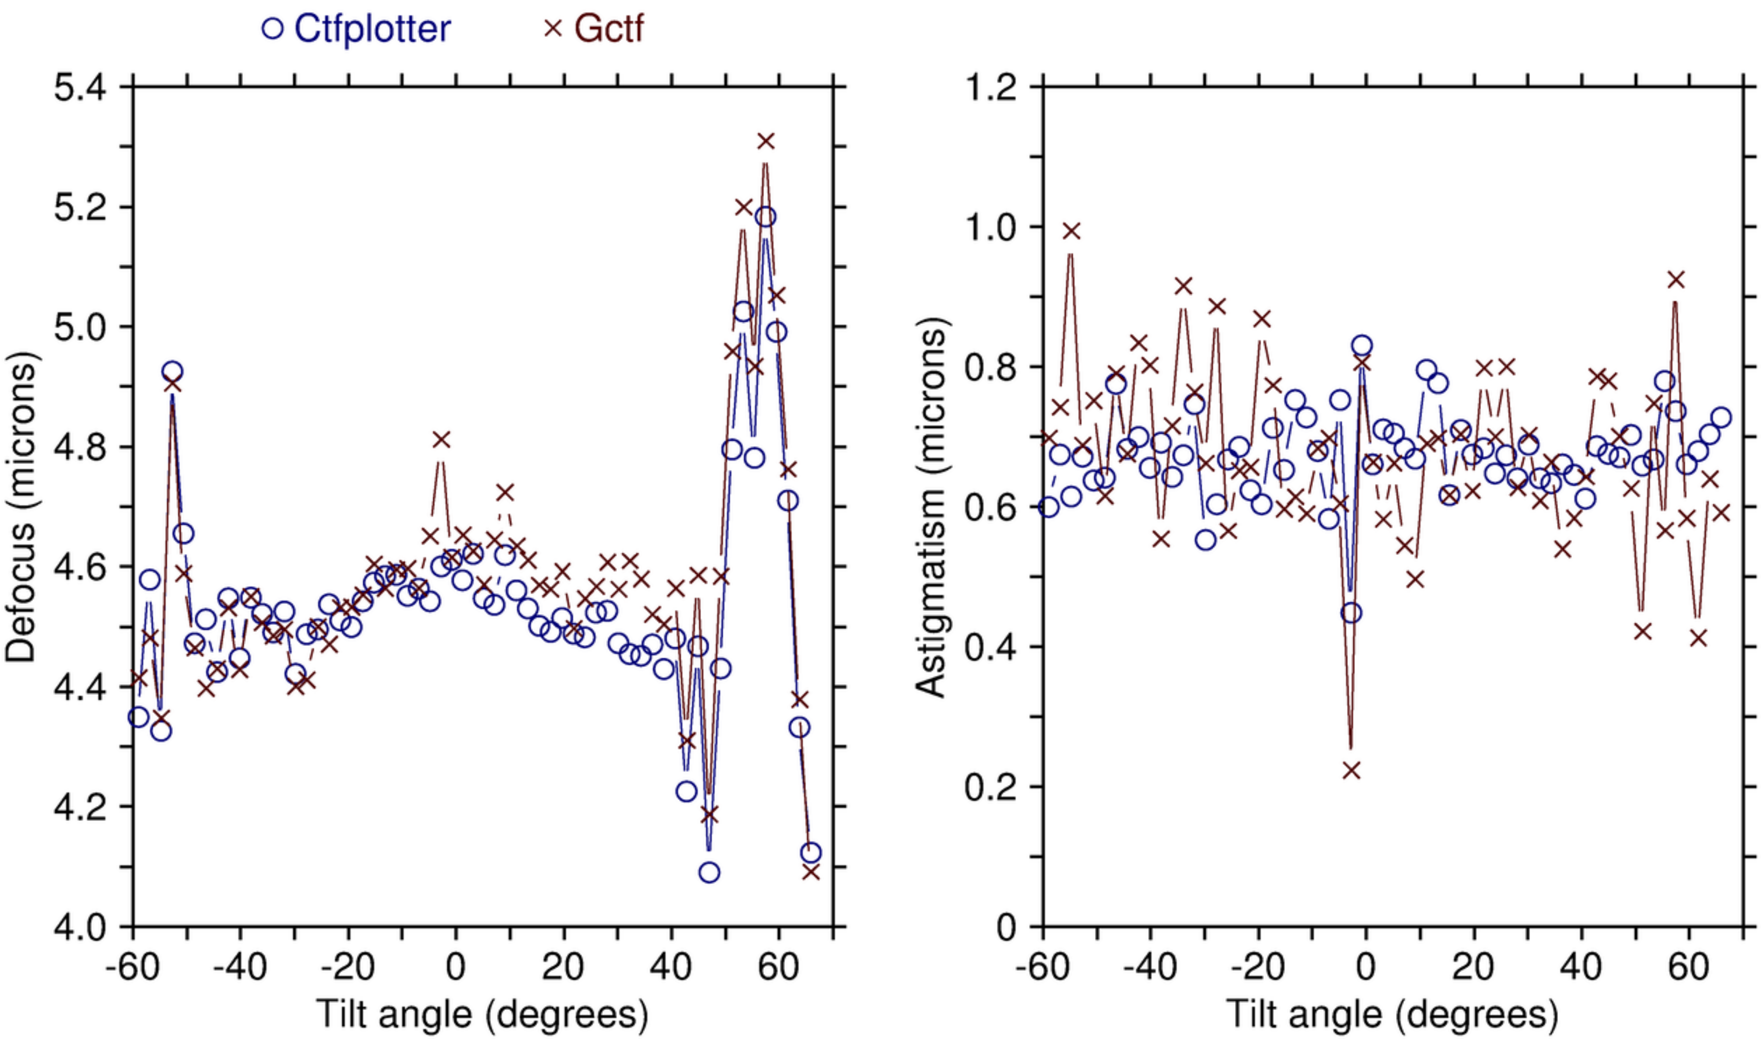

Summing Spectra over 3 or 5 Views Eliminates Implausible Variations in Astigmatism Estimates

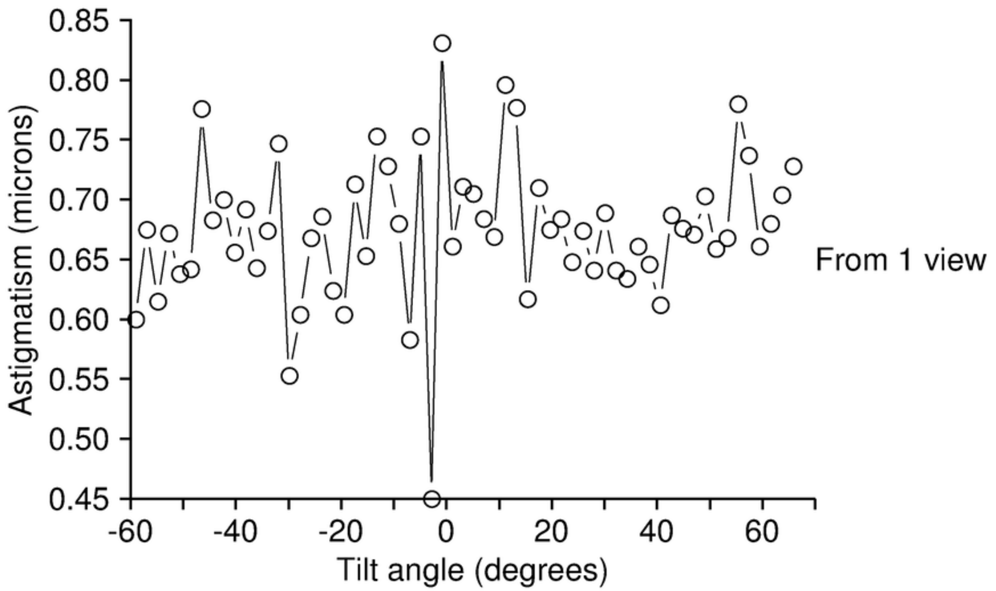

Summing Spectra over 3 or 5 Views Eliminates Implausible Variations in Astigmatism Estimates

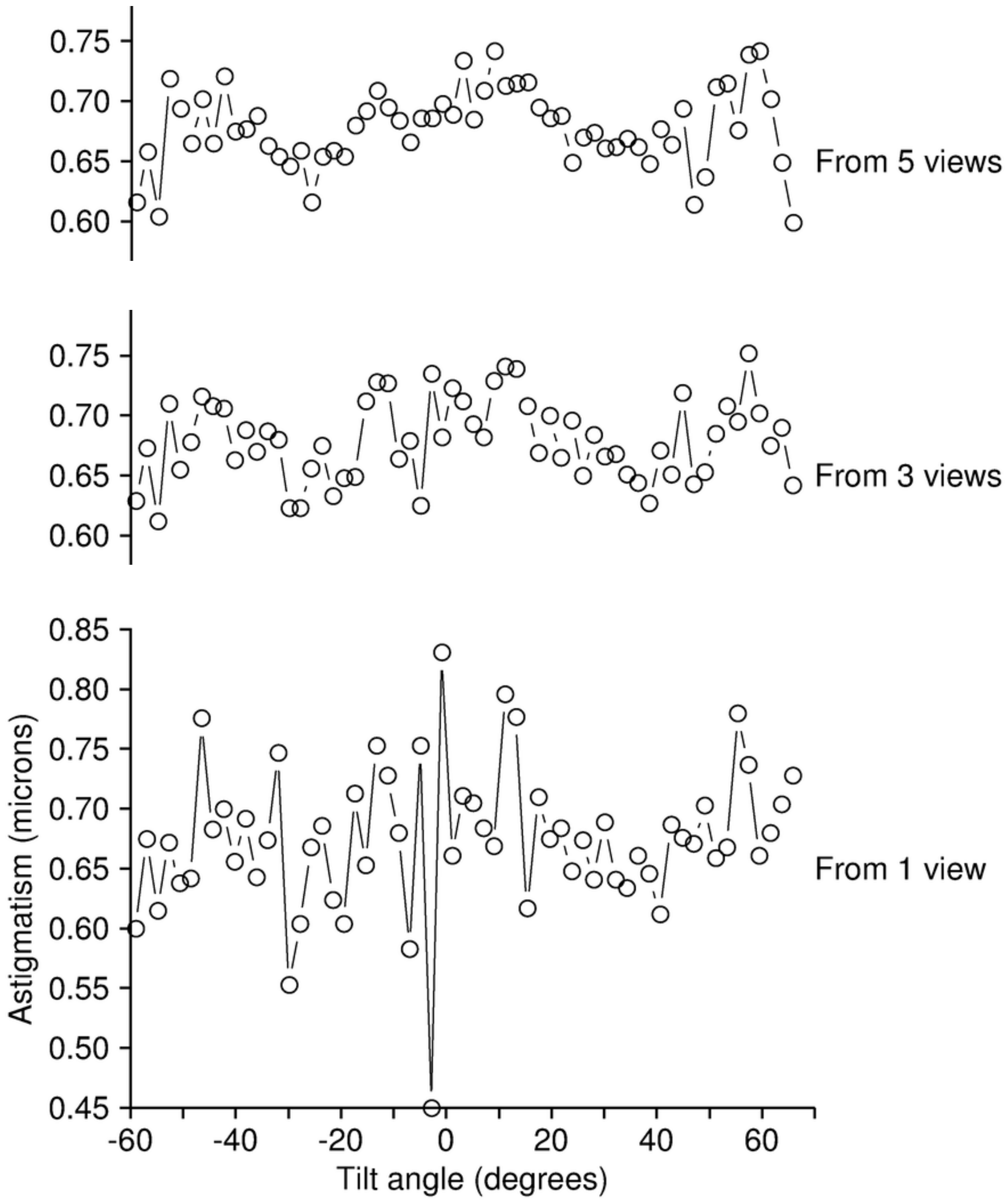

Summing over Views Gives Less Noisy Estimate of Astigmatism Angle Also

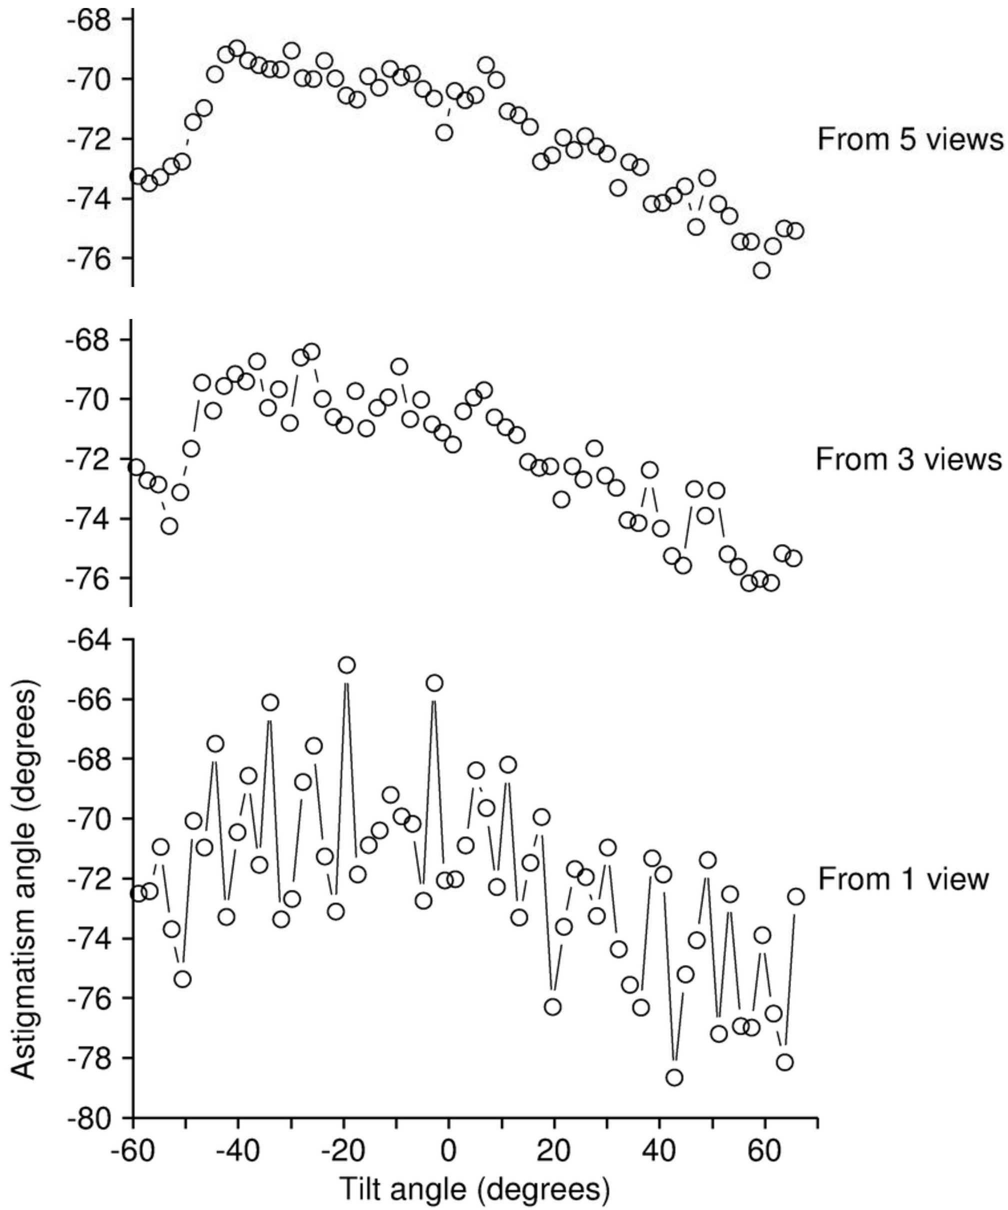

"Cropping" Spectrum to Spread Out Low Frequencies

- With a small pixel size, there are too few frequency bins (data points) between zeros for good fitting
- Cropping the spectrum before binning spreads out the zeros
- This spectrum then has good signal to the fifth zero

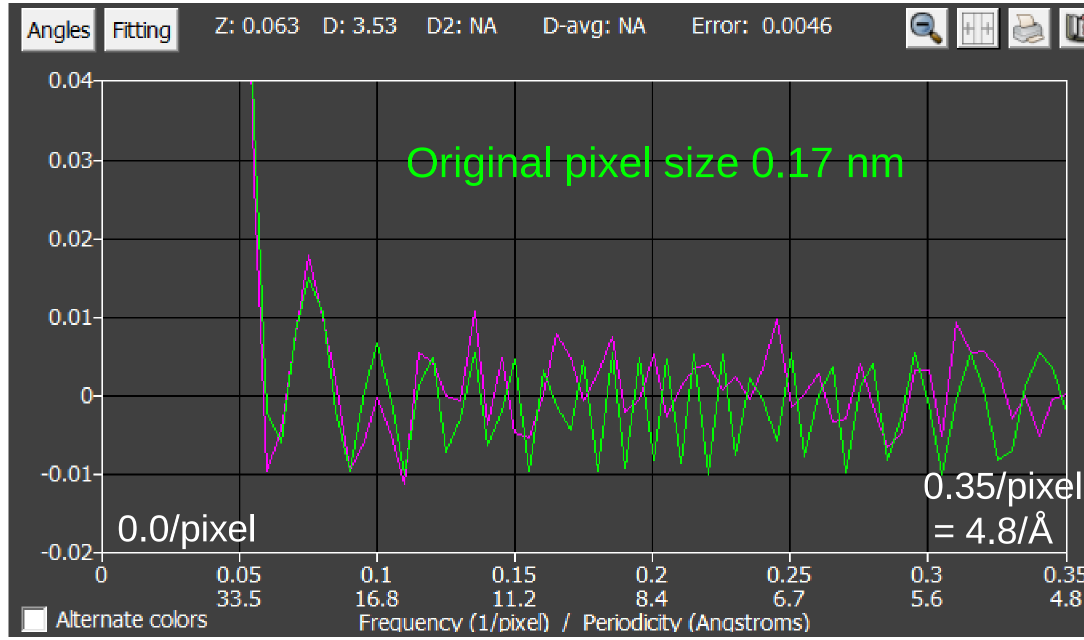

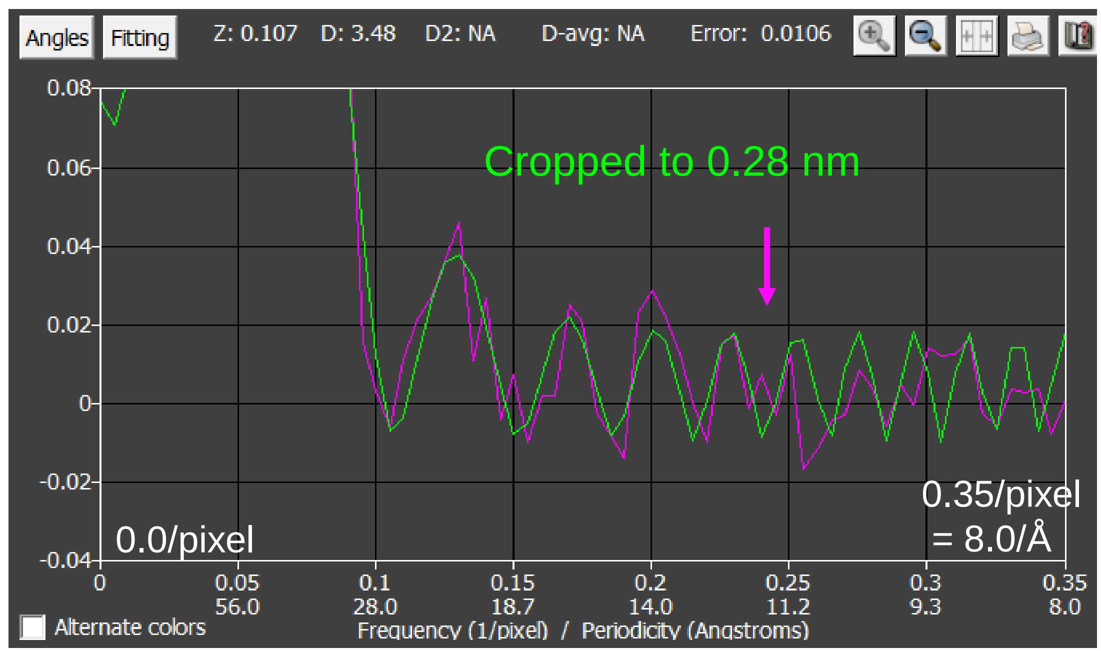

Tilt series of rabies virus particles, from Christiane Riedel,
±45° @ 3°,
6° view, 2.0 $e^{-}/Å^{2}$
K2 camera

“Cropping” Spectrum to Spread Out Low Frequencies

- With a small pixel size, there are too few frequency bins (data points) between zeros for good fitting
- Cropping the spectrum before binning spreads out the zeros
- This spectrum then has good signal to the 5th zero
- Finding astigmatism from 5-11 views extends signal to the 7th zero

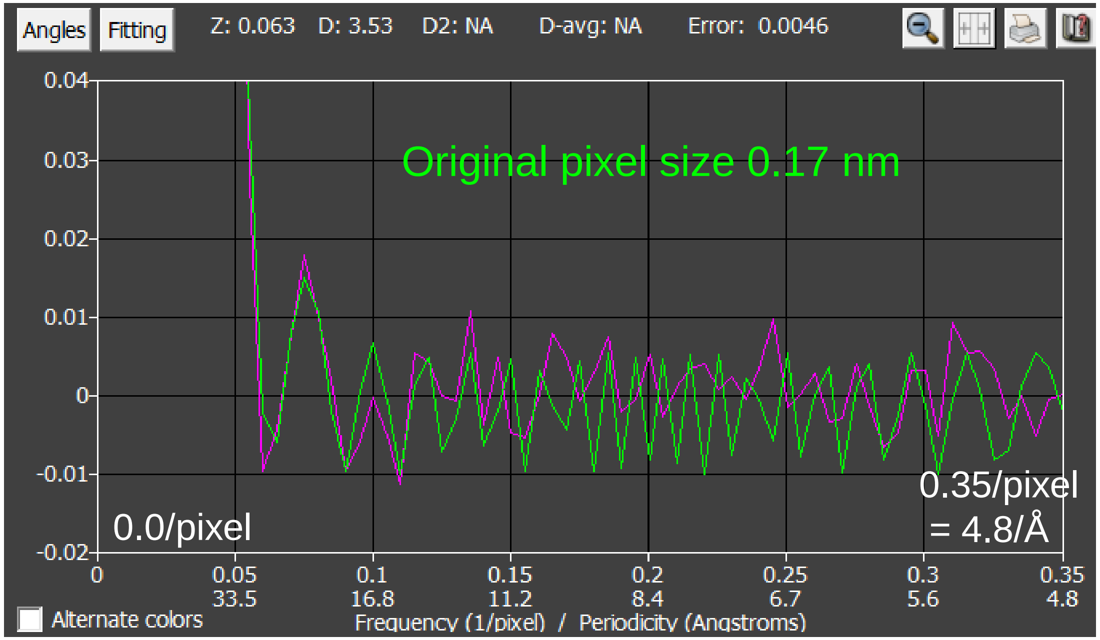

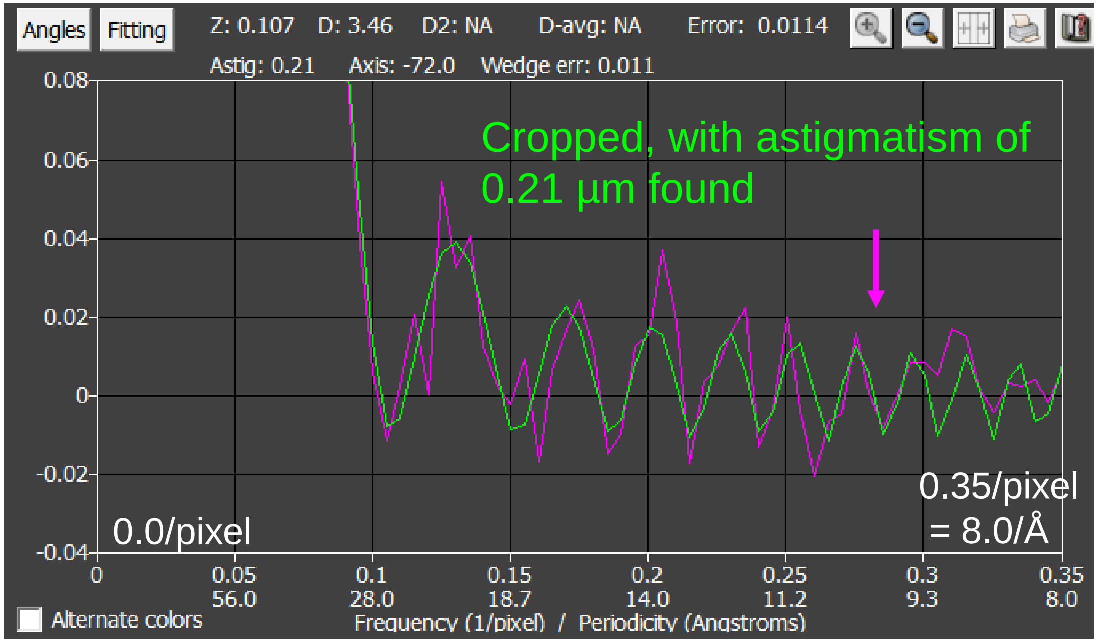

Tilt series of rabies virus particles,
from Christiane Riedel,
±45° @ 3°,
6° view, 2.0 $e^-/Å^2$
K2 camera

Phase Plate Imaging for Tilt Series

- The motivation for using a phase plate with underfocus is to avoid going overfocus in tilted images
- But if enough defocus is used to require CTF correction, even more defocus may be needed to get enough zeros for accurate defocus and phase determination

Phase Shift's Effect on the Spectrum

- A phase shift changes the first zero location and zero spacings for a given defocus

Tilt series of FIB lamella of yeast cell with Pt layer, from Julia Mahamid, EMBL
±50° @ 2° by Hagen scheme
0° view, 2.4 $e^-/Å^2$
K2 camera

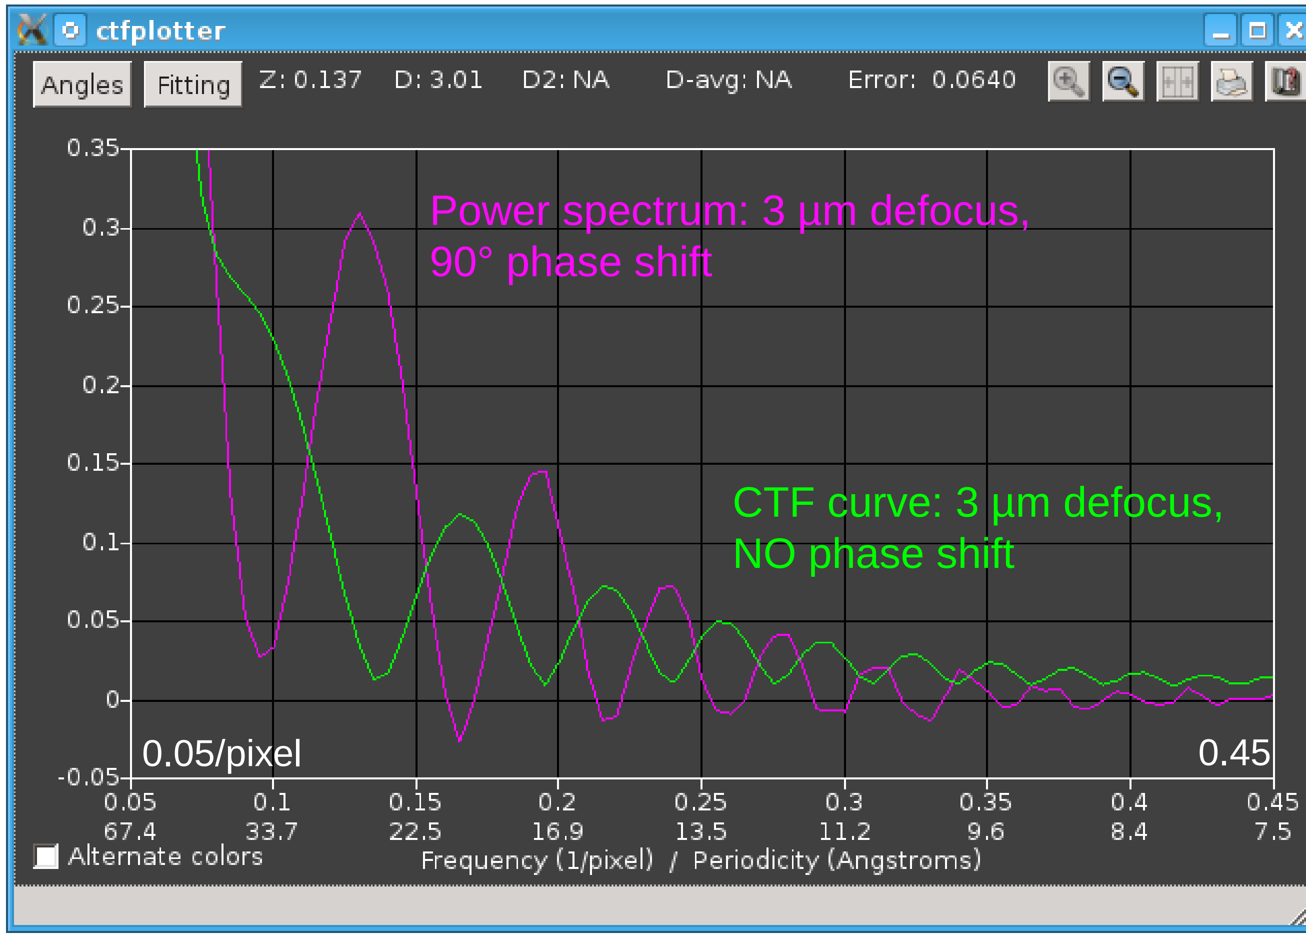

No fitting: CTF curve placed by clicking zero position

Phase Can Also Be Found in Ctfplotter

- It is done with a separate search for phase that minimizes error from CTF-like fit to spectrum
- But the solution can get off when there are many zeros – there is another parameter that can correct for this second-order effect

Tilt series of FIB lamella of yeast cell with Pt layer, from Julia Mahamid, EMBL
±50° @ 2° by Hagen scheme
0° view, 2.4 $e^{-}/Å^{2}$
K2 camera

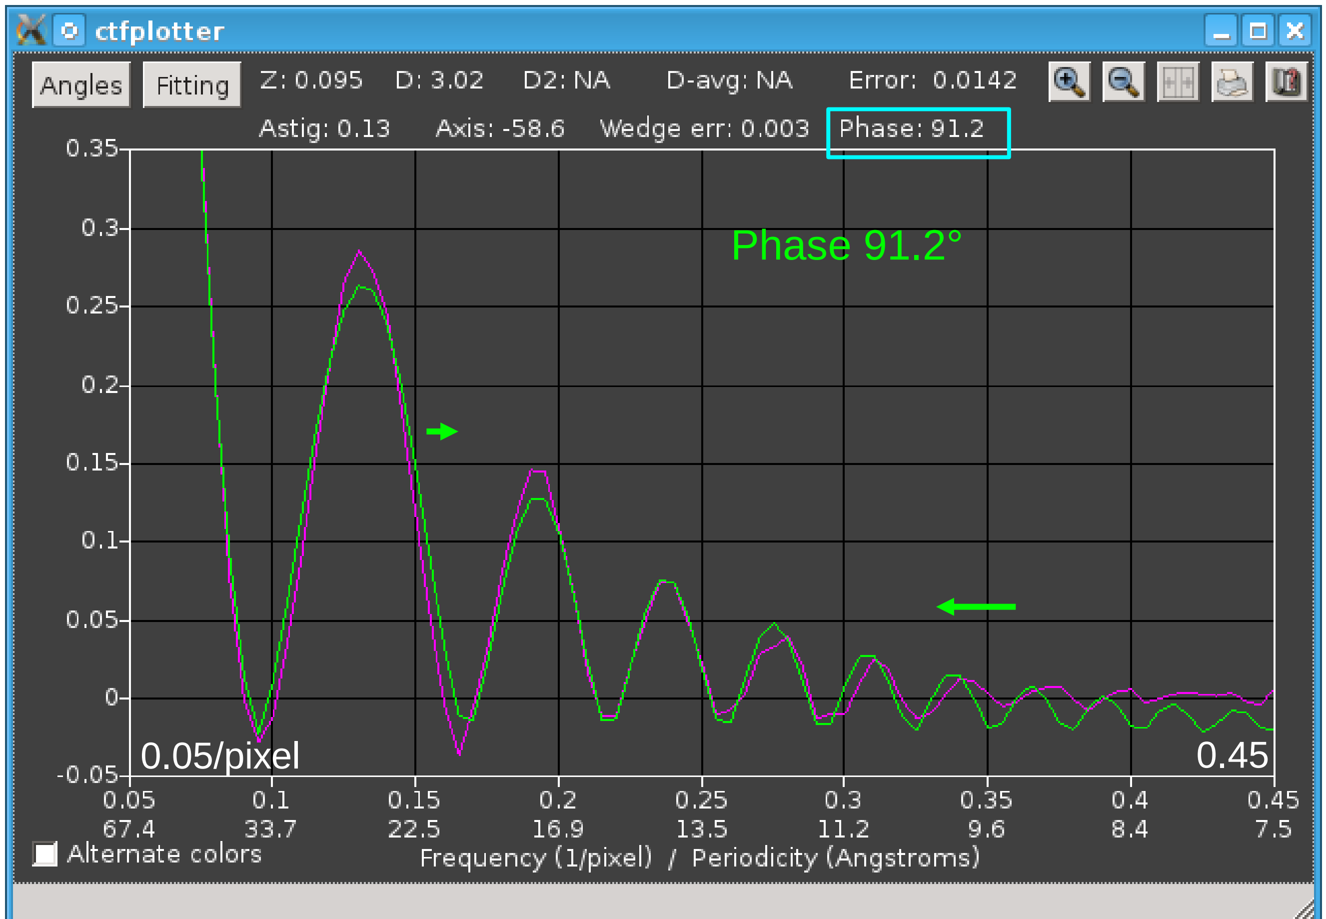

Phase Can Be Found from Sums of Multiple Views Even When Finding Defocus for Every View

- This will give less variability in the phase estimate
- Even if it introduces a small error in phase, the fitted defocus can vary from its true value to compensate and fit the zero positions well, which is what really matters

Starting Ctfplotter

- Be sure to check these parameters when you reach CTF correction panel:
  - Microscope voltage
  - Spherical aberration – a constant for model of microscope and pole piece. 0 is allowed
  - Expected defocus
- Select the noise configuration file (text file listing individual noise image files) or noise image stack if using one
- Start Ctfplotter and zoom the graph
  - Use left mouse button to draw rubber band around region to magnify
  - Place rubber band inside desired region – e.g., to get X to range from 0.1 to 0.5, draw band from -0.08 to 0.48
  - Zoom again vertically if necessary to make hump after first zero occupy a good fraction of the vertical extent
  - Don't be misled by a big hump before the first zero when there is lots of gold

Ctfplotter Steps 1

- Make sure expected defocus is correct to within 10%
  - Click on minimum of dip with left mouse if the fit is not close
  - If necessary, either adjust expected defocus or switch to use “Current defocus estimate”
- Crop spectrum to larger pixel size if zeros are too closely spaced
- Reduce the number of views to fit to find minimum number needed for reliable fit
  - *Decrease* end of fitting range if necessary
- Do initial adjustment of fitting parameters
  - Adjust start of range to be on the steeply falling curve, not far before the first zero
    - The **Start** button will make this adjustment based on current fitted defocus
  - Adjust end of range to wherever is appropriate
    - Type in value, use End button after setting number of zeros, or double-click with right mouse button
    - Crop spectrum (more) to handle warning
  - Adjust baseline fitting order if curvature of the baseline interferes with the region that you are trying to fit

Ctfplotter Steps 2

- Switch to use “Current defocus estimate” to get PS that are most consistent with the defocus
- Try finding astigmatism
  - Increase “Min views for astigmatism” if some fits are off
  - *Increase* end of fitting range if appropriate (i.e., spectrum has more good zeros)
- Step through the series at 5-10 view interval and see if parameters work though the tilt range
  - Just increase “Min views for astigmatism” if that is needed in some parts of range
  - You will probably need to decrease the end of the fitting range to accommodate high tilt, or else decide to use different parameters in subsets of the range

Autofitting in Ctfplotter

- Return to 0 degrees
- If fitting defocus to multiple views, set step size to half the # of views to fit
- Otherwise check “Fit each view separately” or set step size to 1
- Press **Autofit All Steps** and watch it
  - If finding astigmatism, you can turn off “Show wedge fits” when you have seen enough
  - You can press **STOP** to stop the fitting
- Press Graph Values to see graph(s) versus tilt angle
  - Double click a table row to load it in and check the fitting
  - If you improve a fit, press Store in Table to replace value

3D CTF Correction

- Turonova et al. showed that a CTF correction that accounts for Z-height in the specimen is important for high resolution
  - NovaCTF can be used to make a 3D CTF corrected tomogram
- A 3D CTF corrected tomogram can be made in Etomo, and subtomograms can be made with Subtomosetup (which has an interface in Etomo too)

Standard backprojection parameters

Parameters for 3D CTF operations

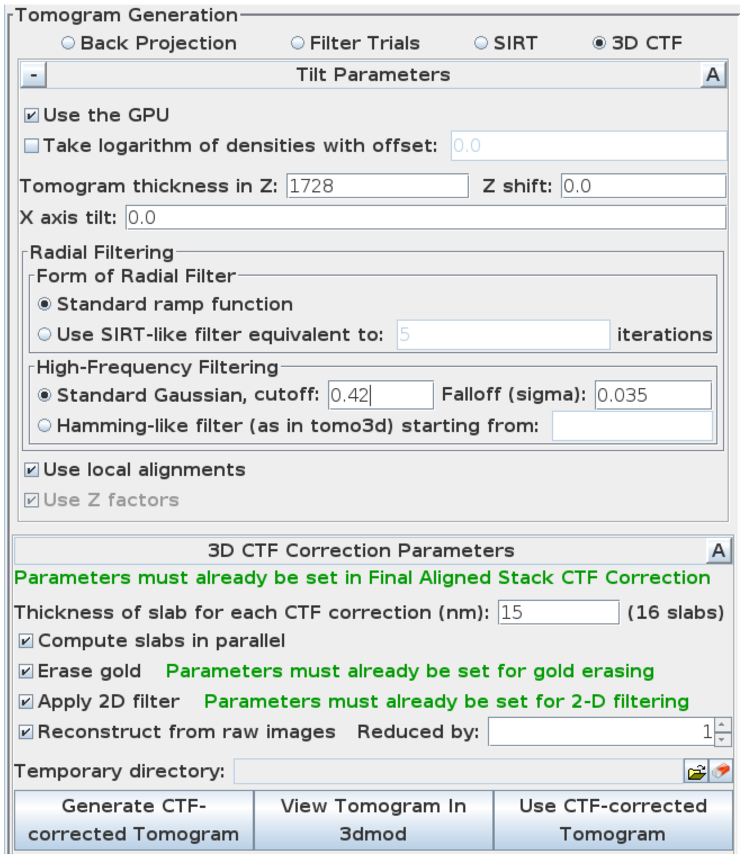

3D CTF Correction in IMOD

- The method is to compute thin slabs, each with its own separate input stack CTF-corrected for that Z height
- It includes optional steps to erase gold after correction and apply dose weighting
- It can reconstruct from an aligned stack as usual or directly from the raw stack
- You have to get the specimen level for this to be accurate