Single-Particle Reconstruction

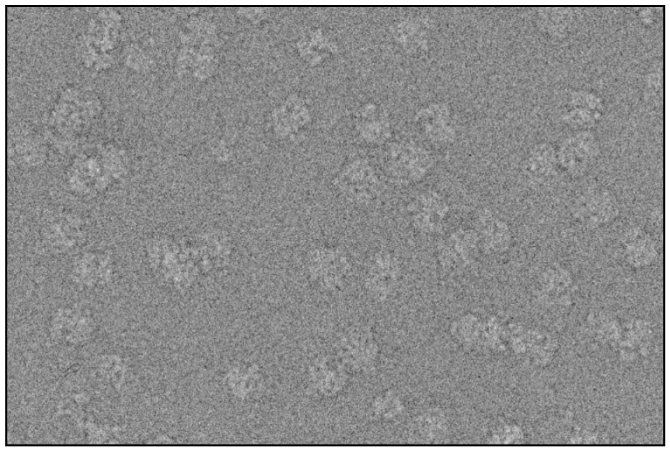

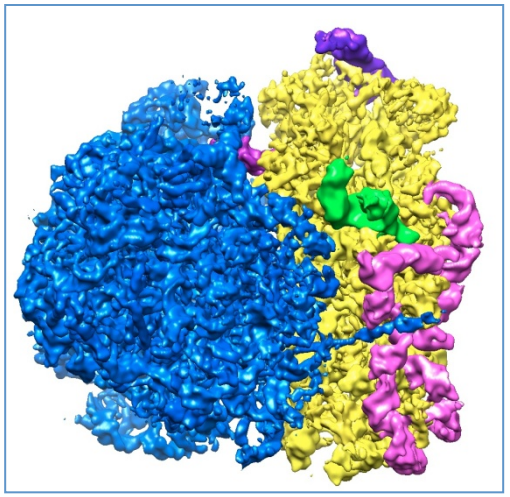

How to get from here

to there

Joachim Frank

Howard Hughes Medical Institute, Department of Biochemistry and Molecular Biophysics, and Department of Biological Sciences, Columbia University

Supported by HHMI, NIH/NCRR, NIH R01 GM55440, and NIH R01 GM29169

Imaging in the Transmission Electron Microscope

Electrons

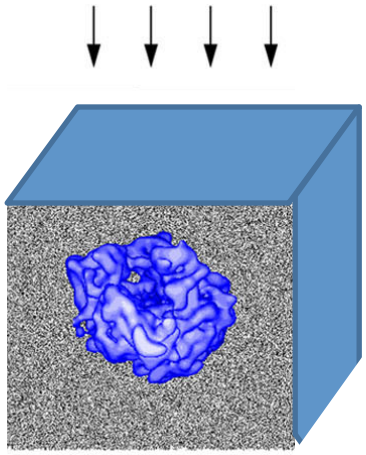

generated in cathode

3D object

molecule embedded in vitreous ice

2D projection

recorded on film or digital camera

Transmission means that the signal is generated from electrons passing through the specimen. We see a 2D projection = line integral over the 3D density along the beam.

Specimen preparation

- Purified sample – standards of purity have changed with the advent of classification ("computational purification"). In some experiments it is even desirable to admit molecules in different conformational and compositional states.
- Apply sample to EM grid as a thin film (~1000 Å) suspended over holes.
- Carefully controlled blotting is a critical step – control blotting force and time
- Coverage with molecules is determined by:

  (1) sample concentration

  (2) geometry and makeup of metal grid -- copper, molybdenum, gold

  copper (traditional), **molybdenum** (match heat expansion of carbon), **gold** (avoid charge-induced vibrations)

  (3) geometry of the overlaid carbon grid (Quantifoil) – size of holes, thickness of film

  (4) (optional) overlaid thin carbon film

Specimen support

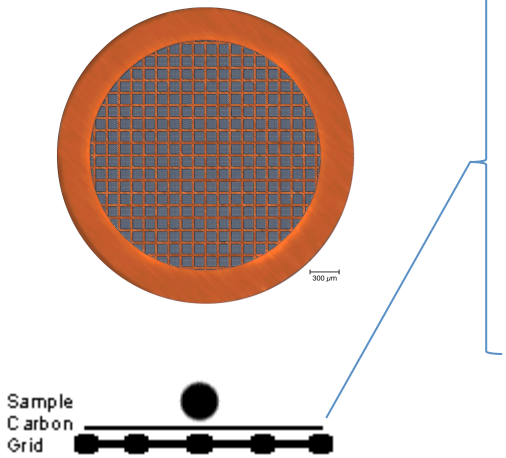

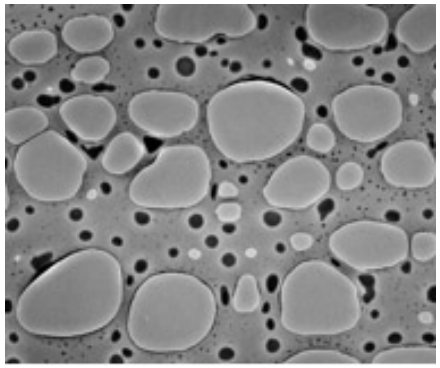

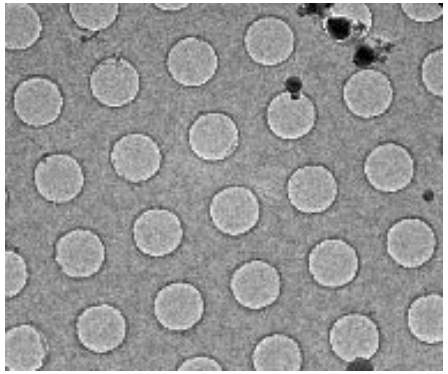

EM grid, Copper, 3 mm

Home-made

Quantifoil or C-flat

Carbon layer on EM grid

Quantifoil vs. C-flat grids: different edges, different thickness, different size of meniscus problem

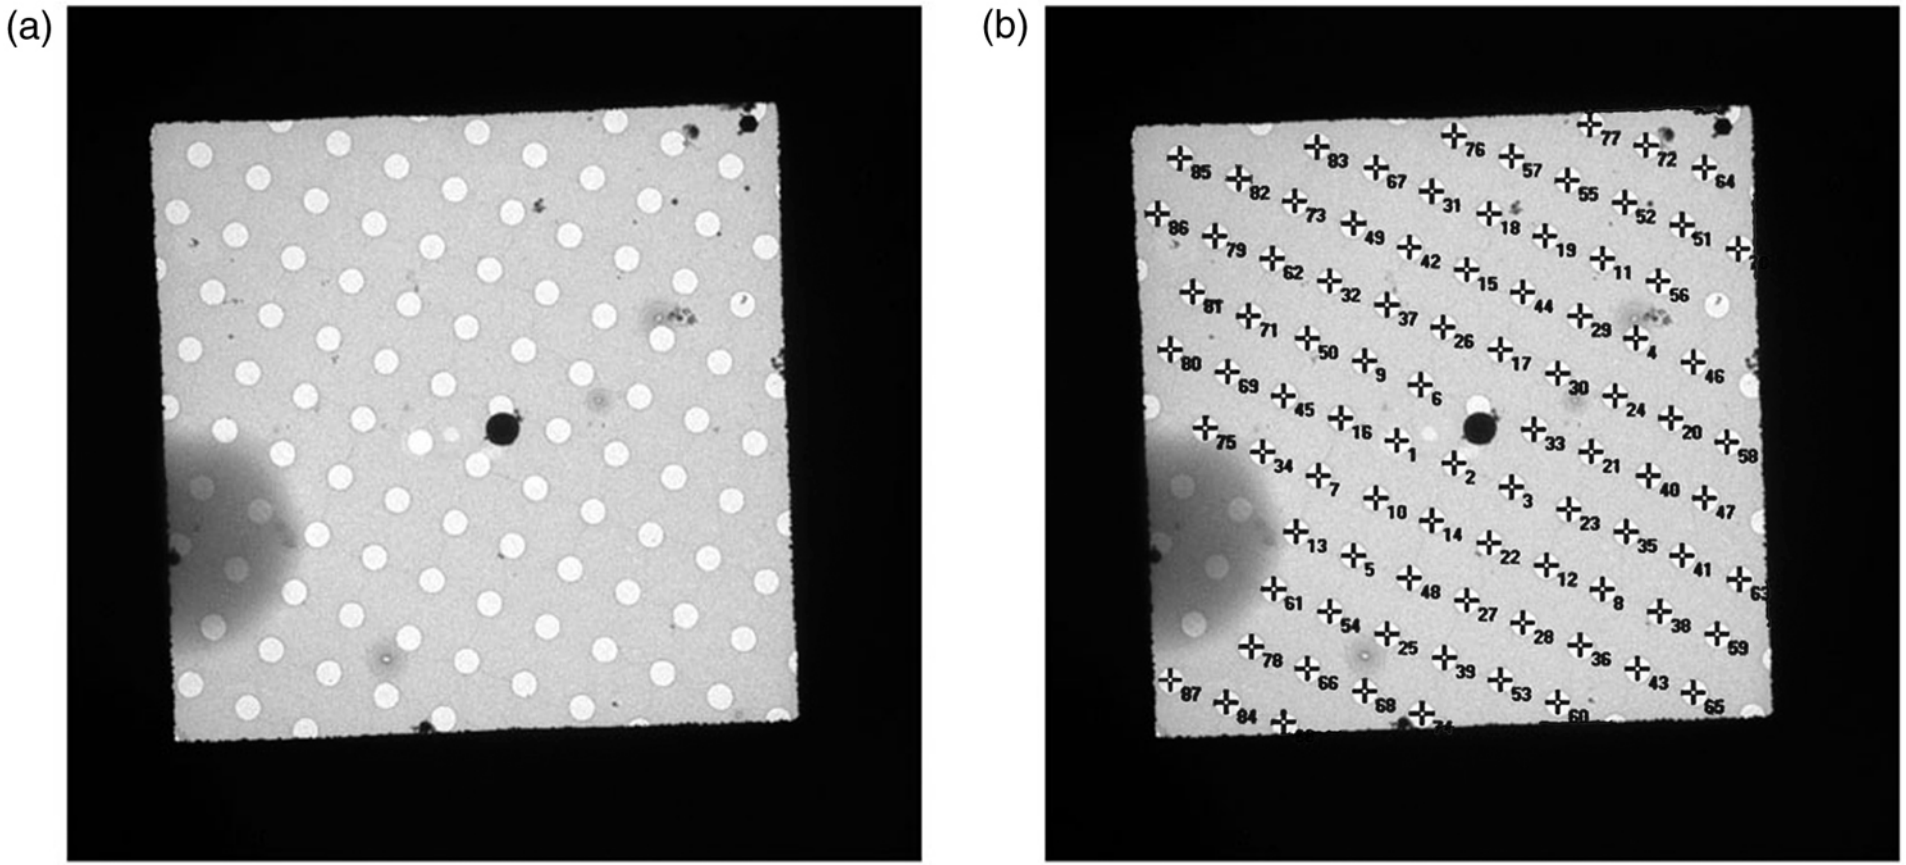

Holes regularly arranged on a Quantifoil grid enable fully automated data acquisition

EXAMPLE OF MENISCUS EFFECT: MOLECULES ACCUMULATE NEAR EDGE OF HOLE

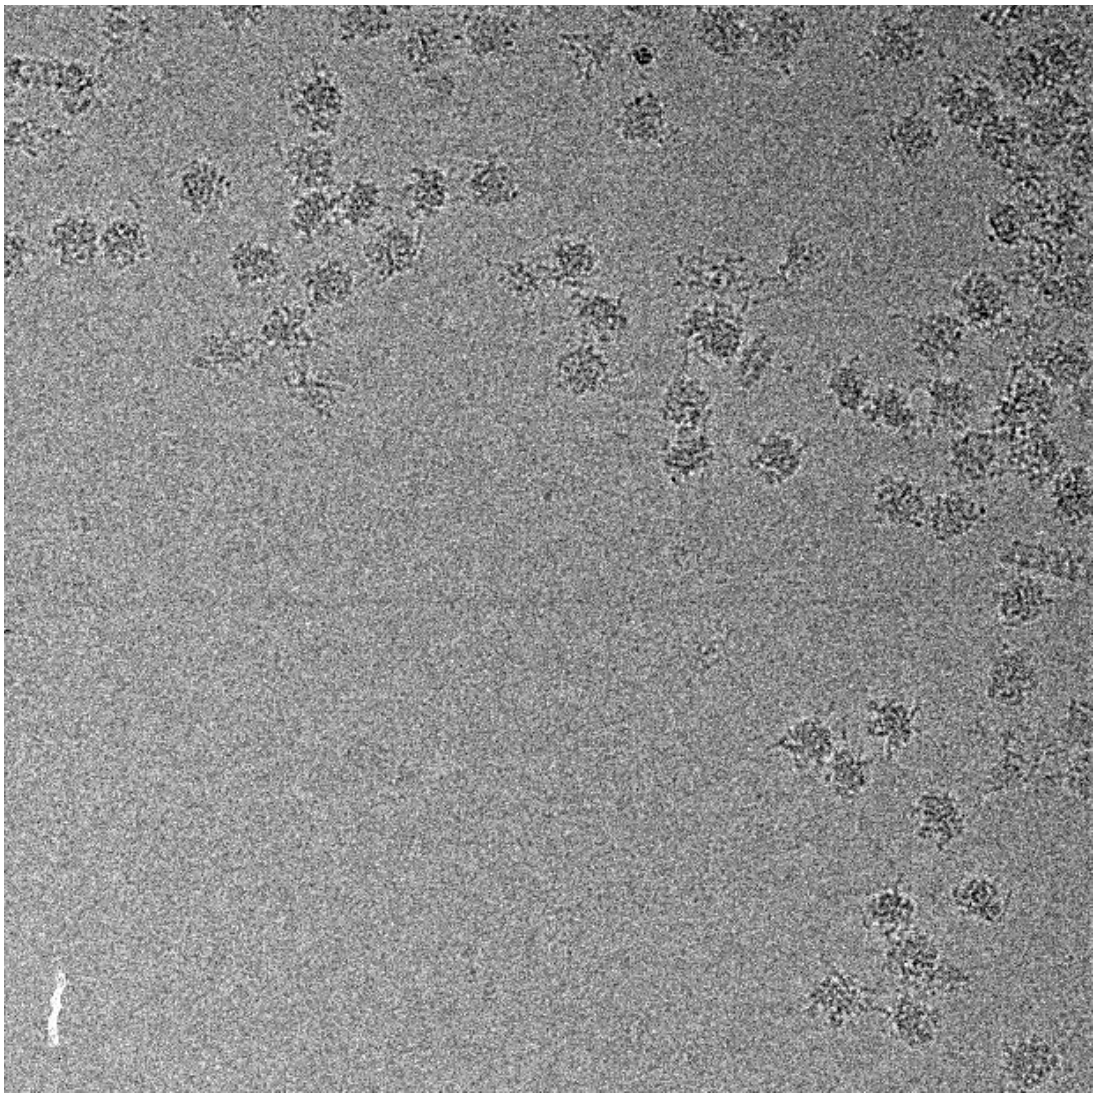

Thin carbon (~100 Å), produced by evaporation on mica, floated onto Quantifoil-coated grid.

(1) Enhance signal of power spectrum, for CTF determination
(2) Induce more even coverage of orientations for some molecules
(e.g., ribosomes)

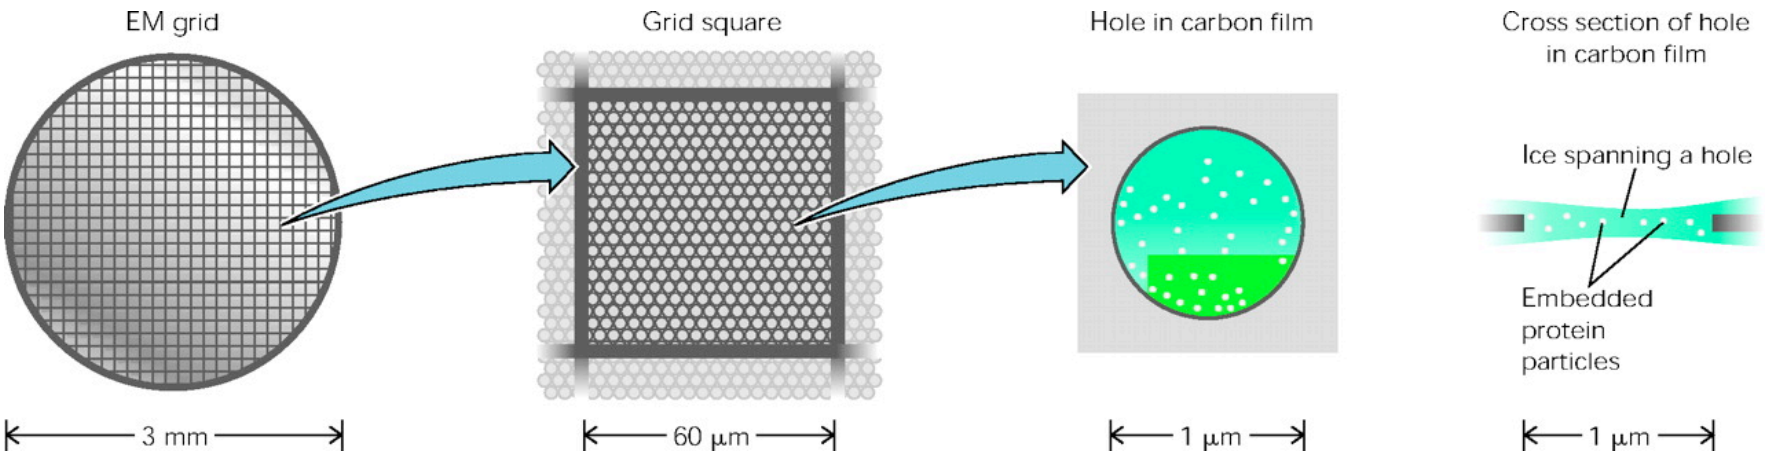

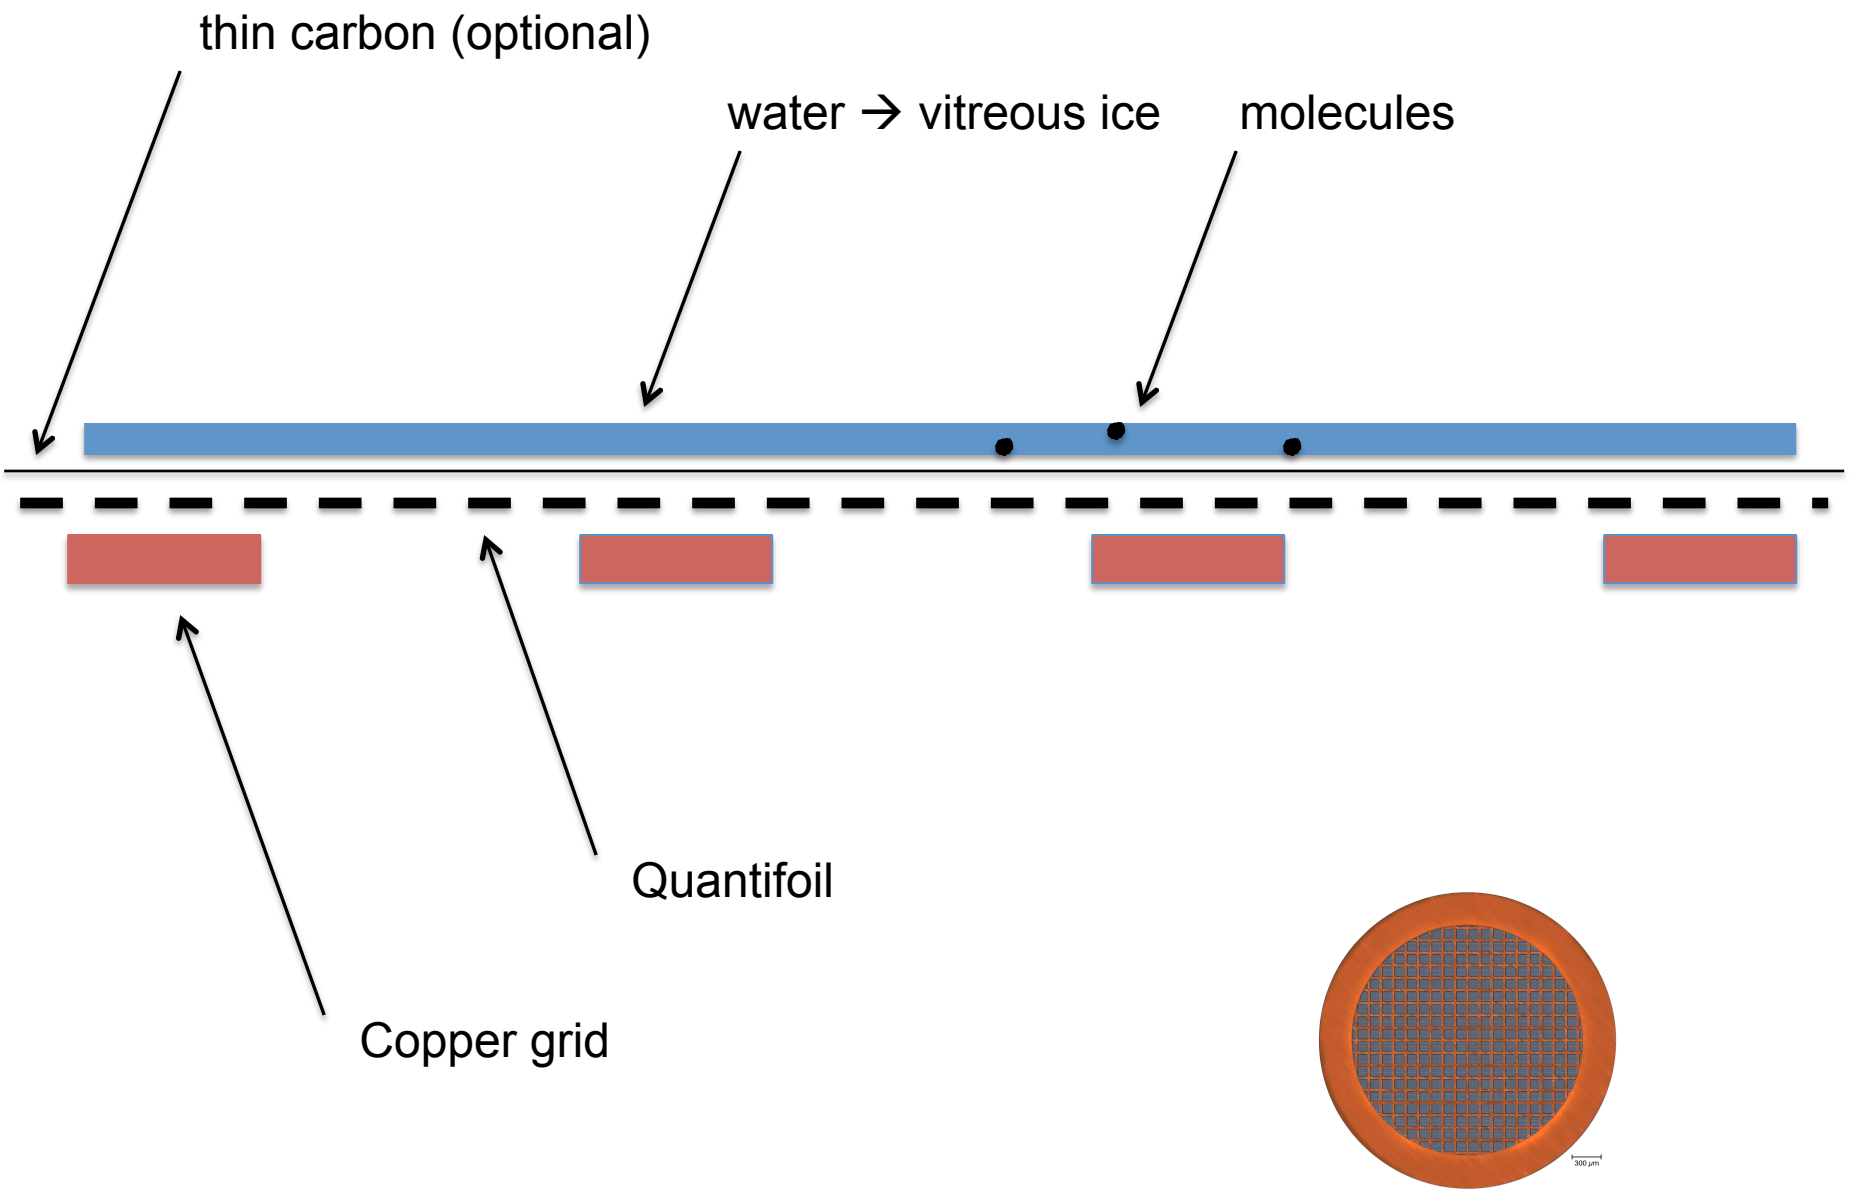
thin carbon (optional)
water → vitreous ice
molecules
Quantifoil
Copper grid

GOLD GRIDS

John Russo and Lori Passmore discovered that the carbon over the grid square oscillates like a drum, moving up and down. There is a sideways component, as well.

Gold grids reduce this effect 50-fold.

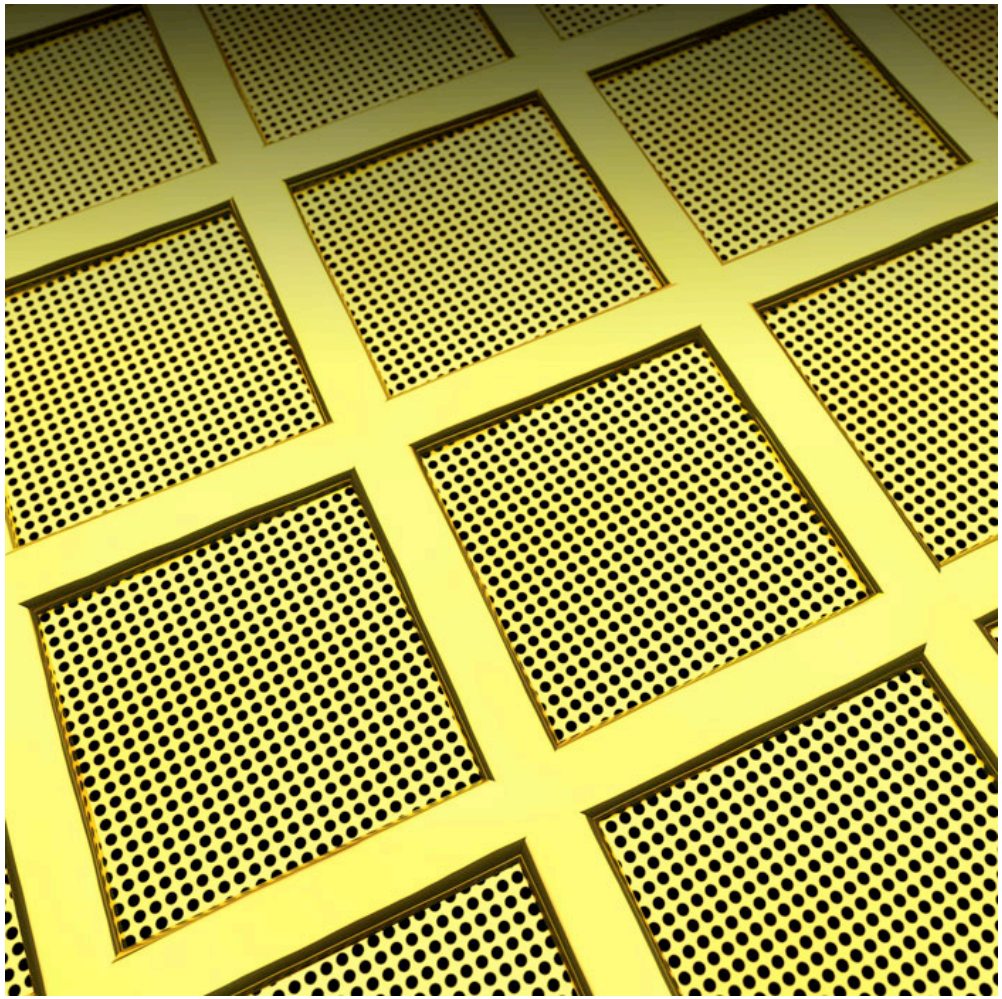

Russo and Passmore, Science 2014

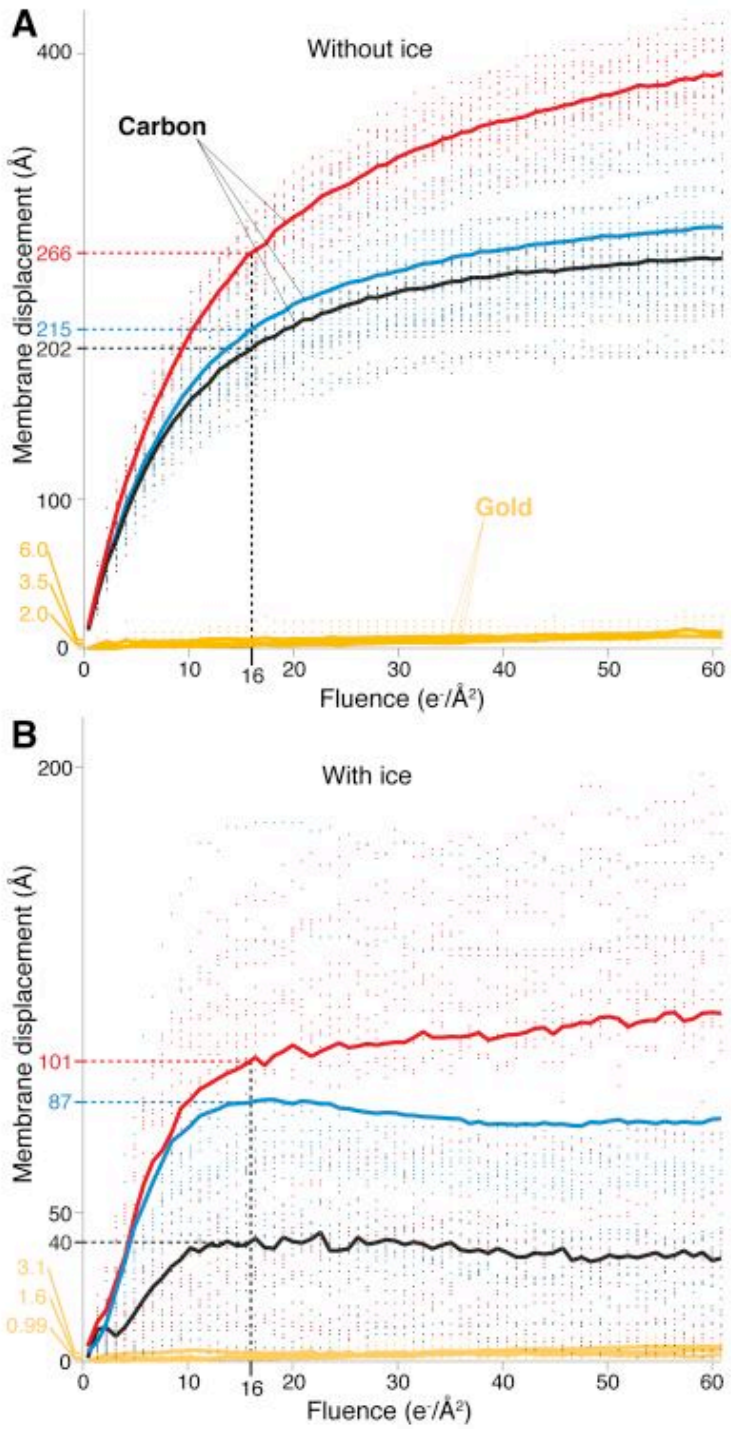

with ice

Russo and Passmore,
Science 2014

Plunge-freezer to prepare samples for cryo-EM

Manual automated, climatized

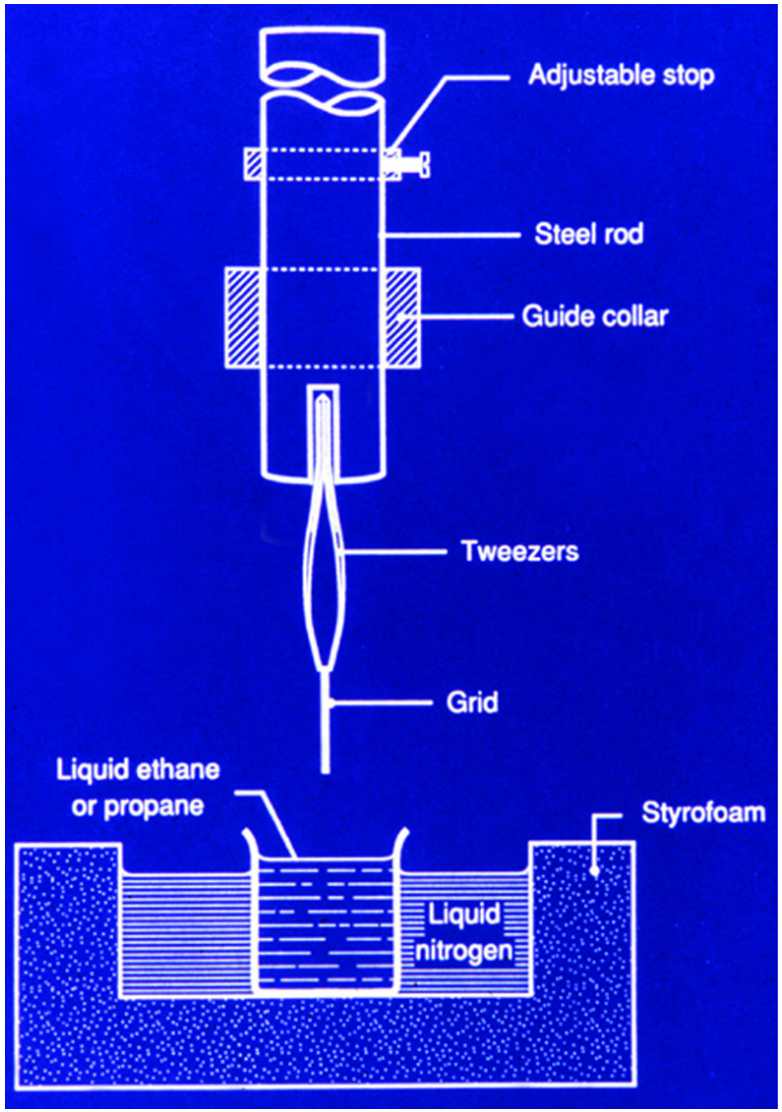

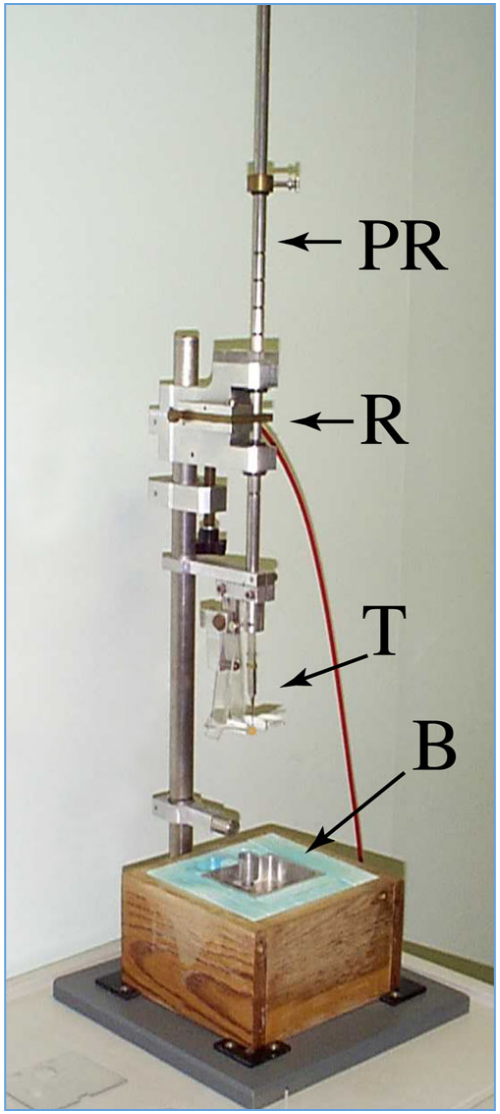

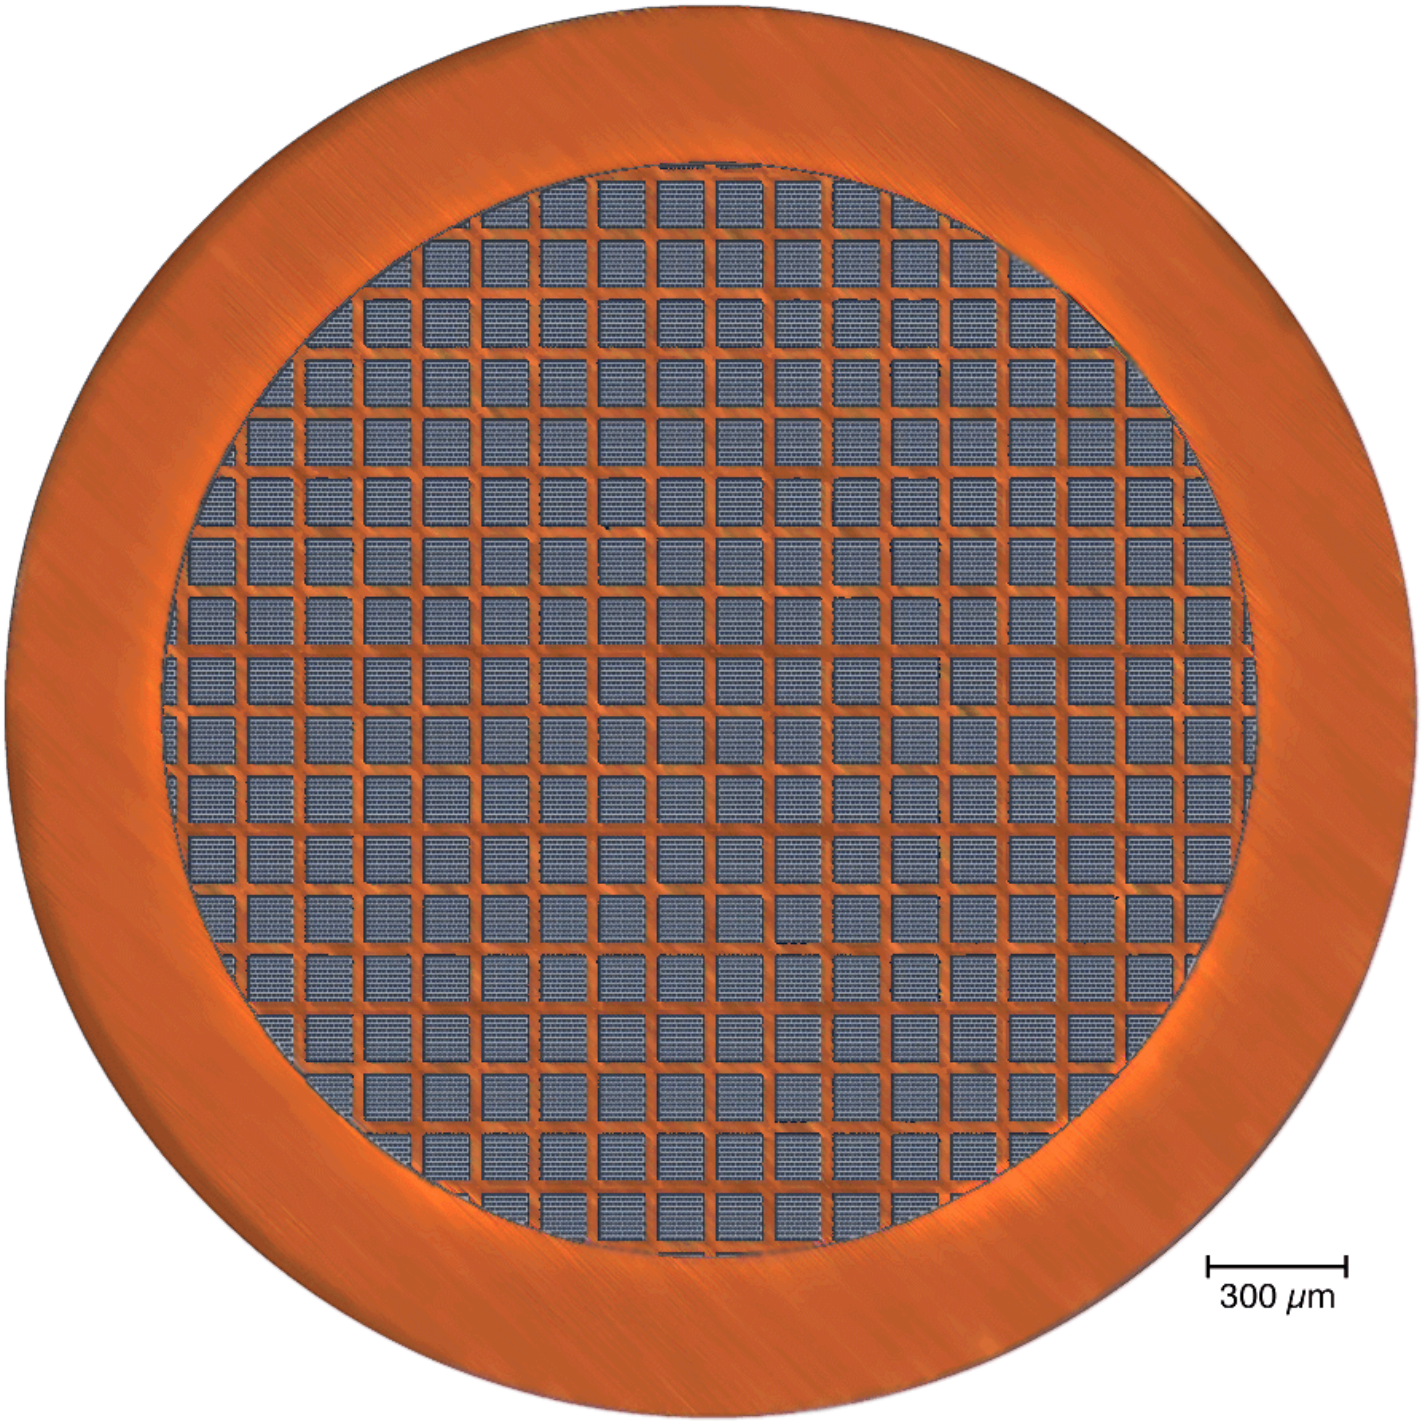
300 μm

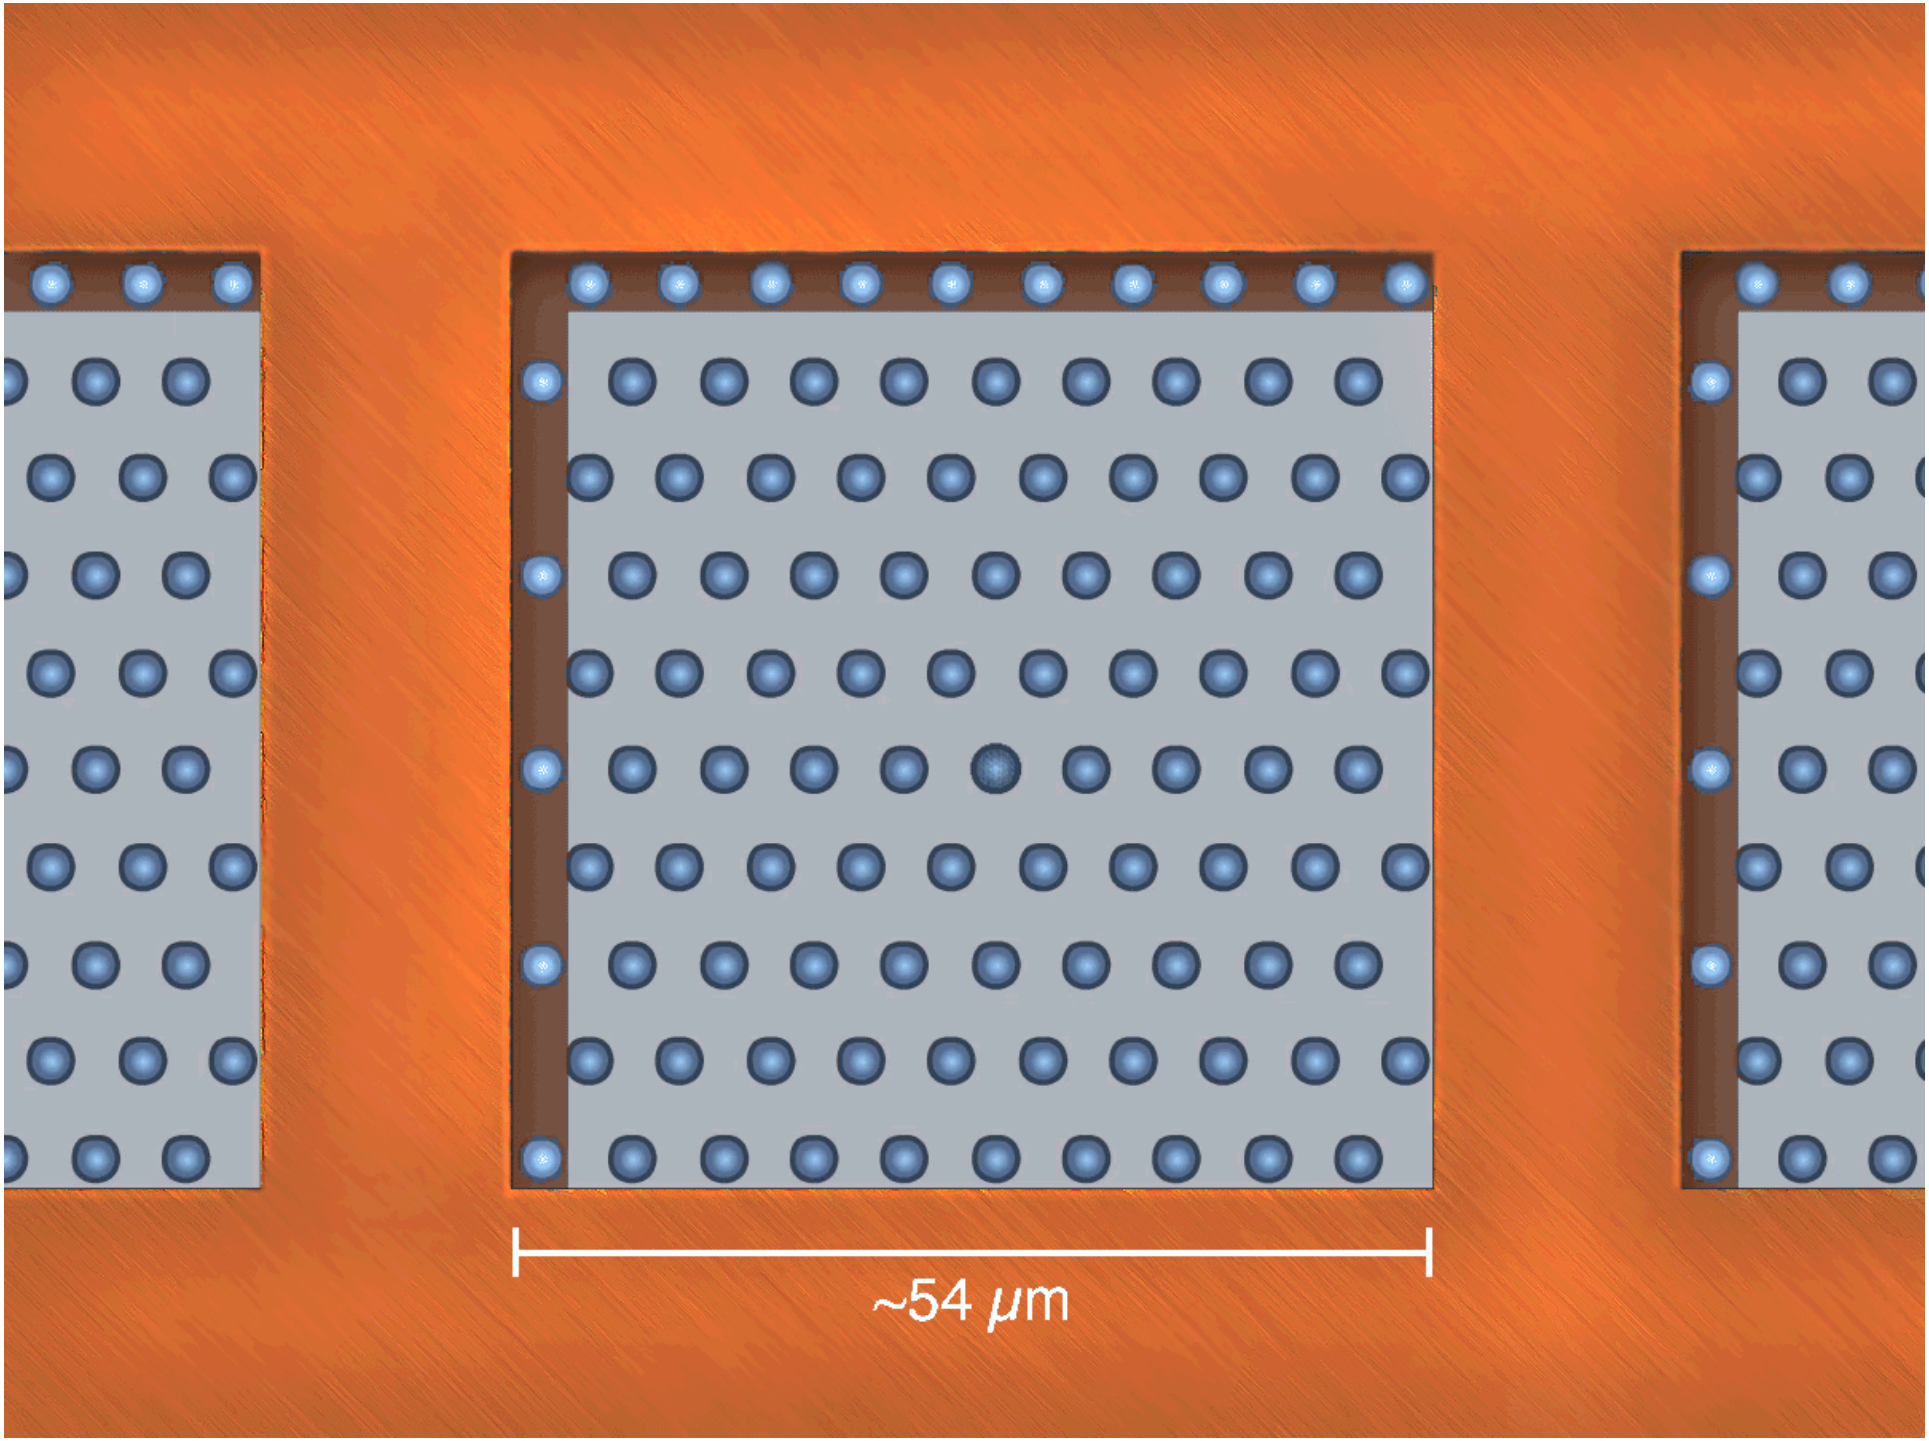
~54 μm

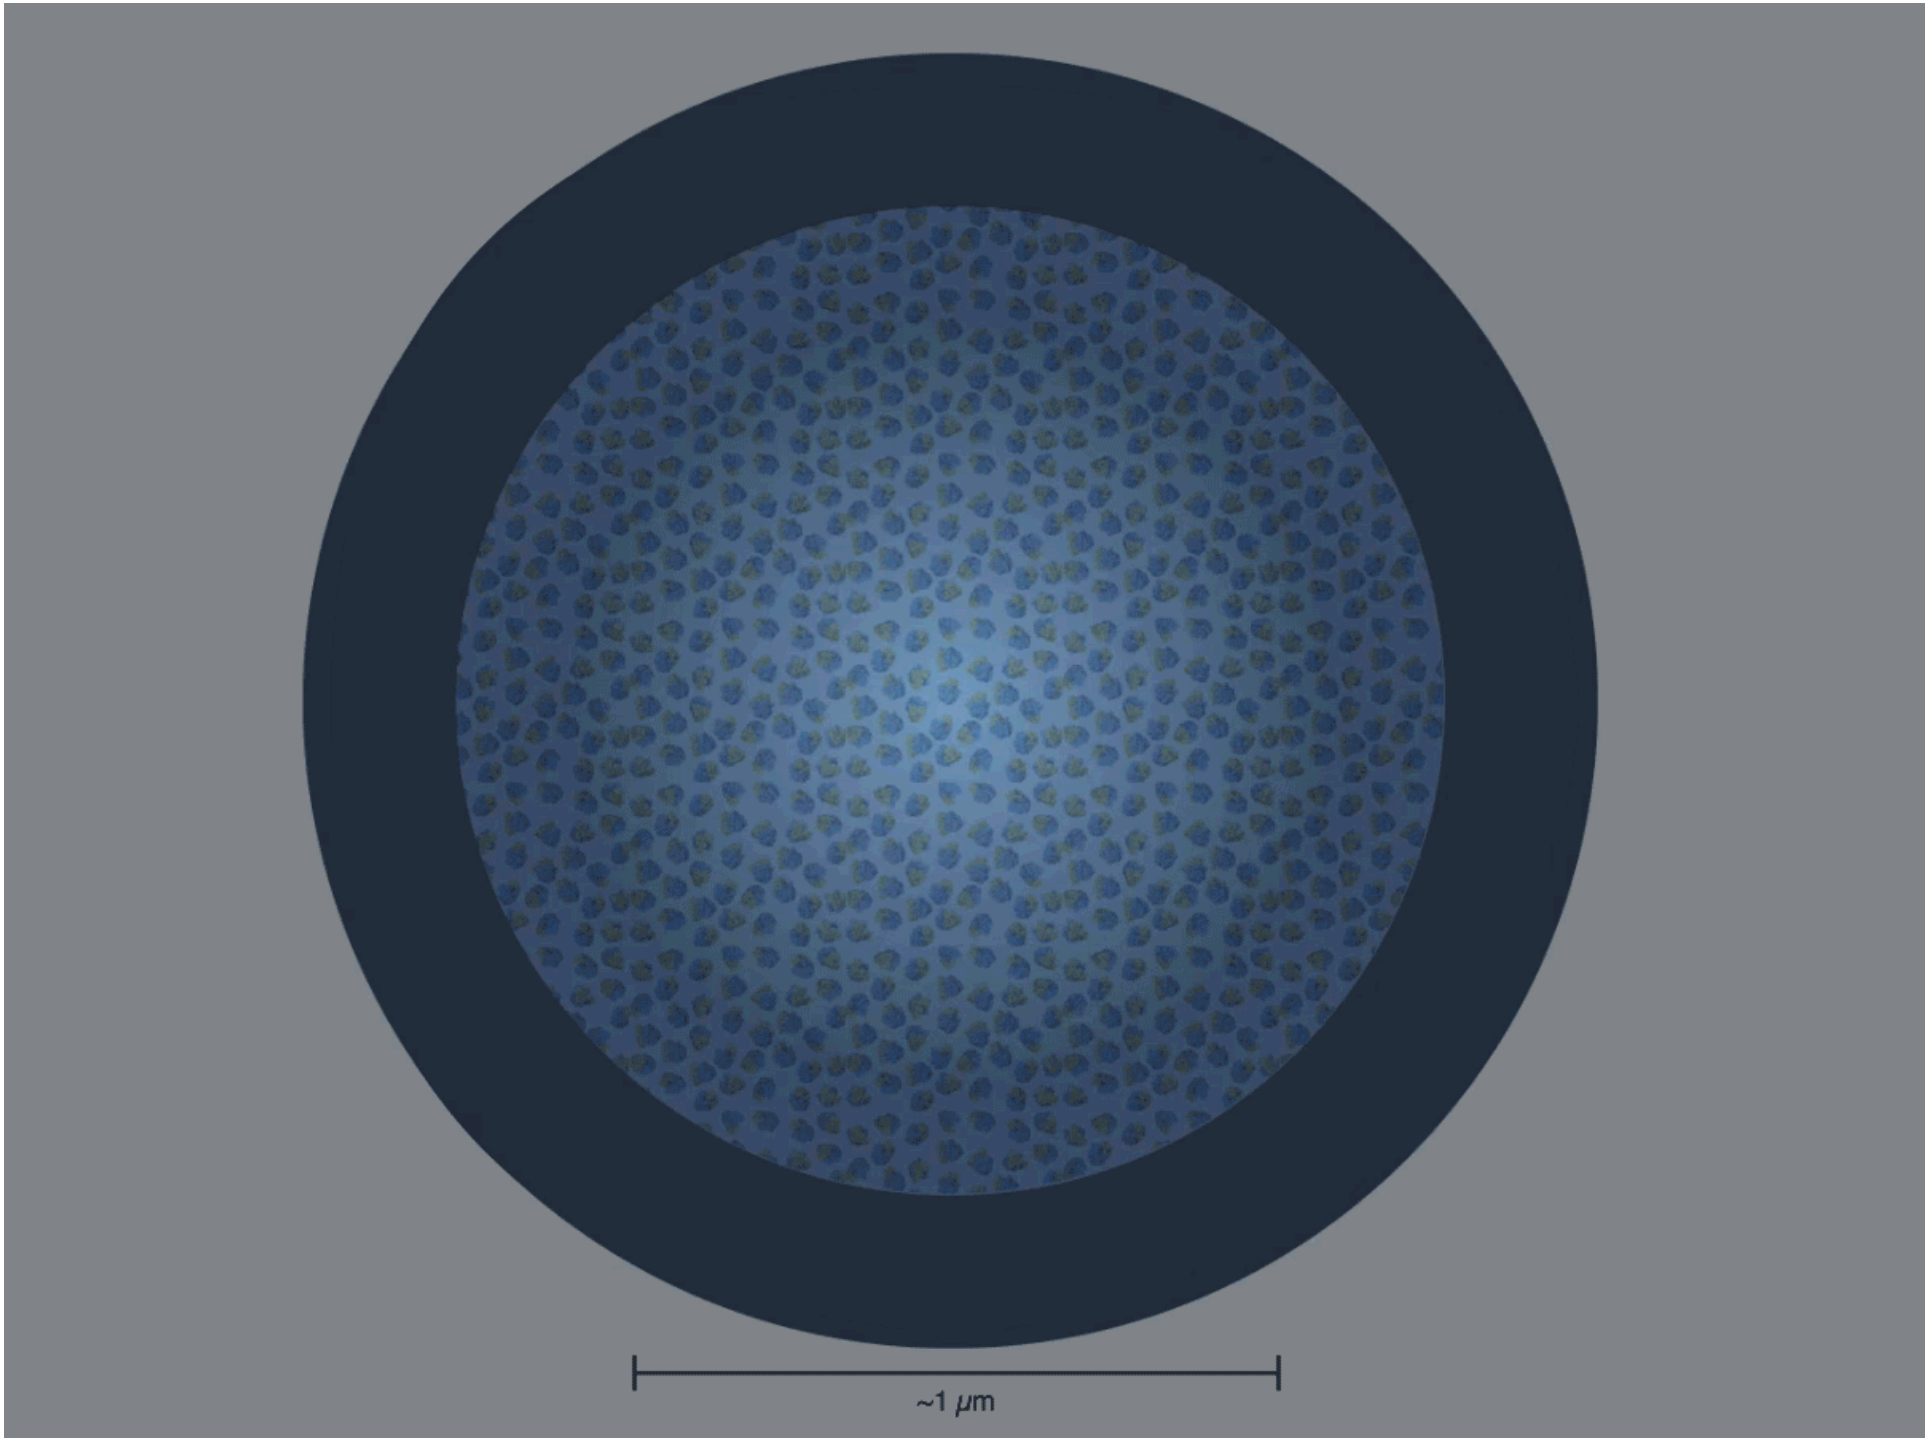
~1 μm

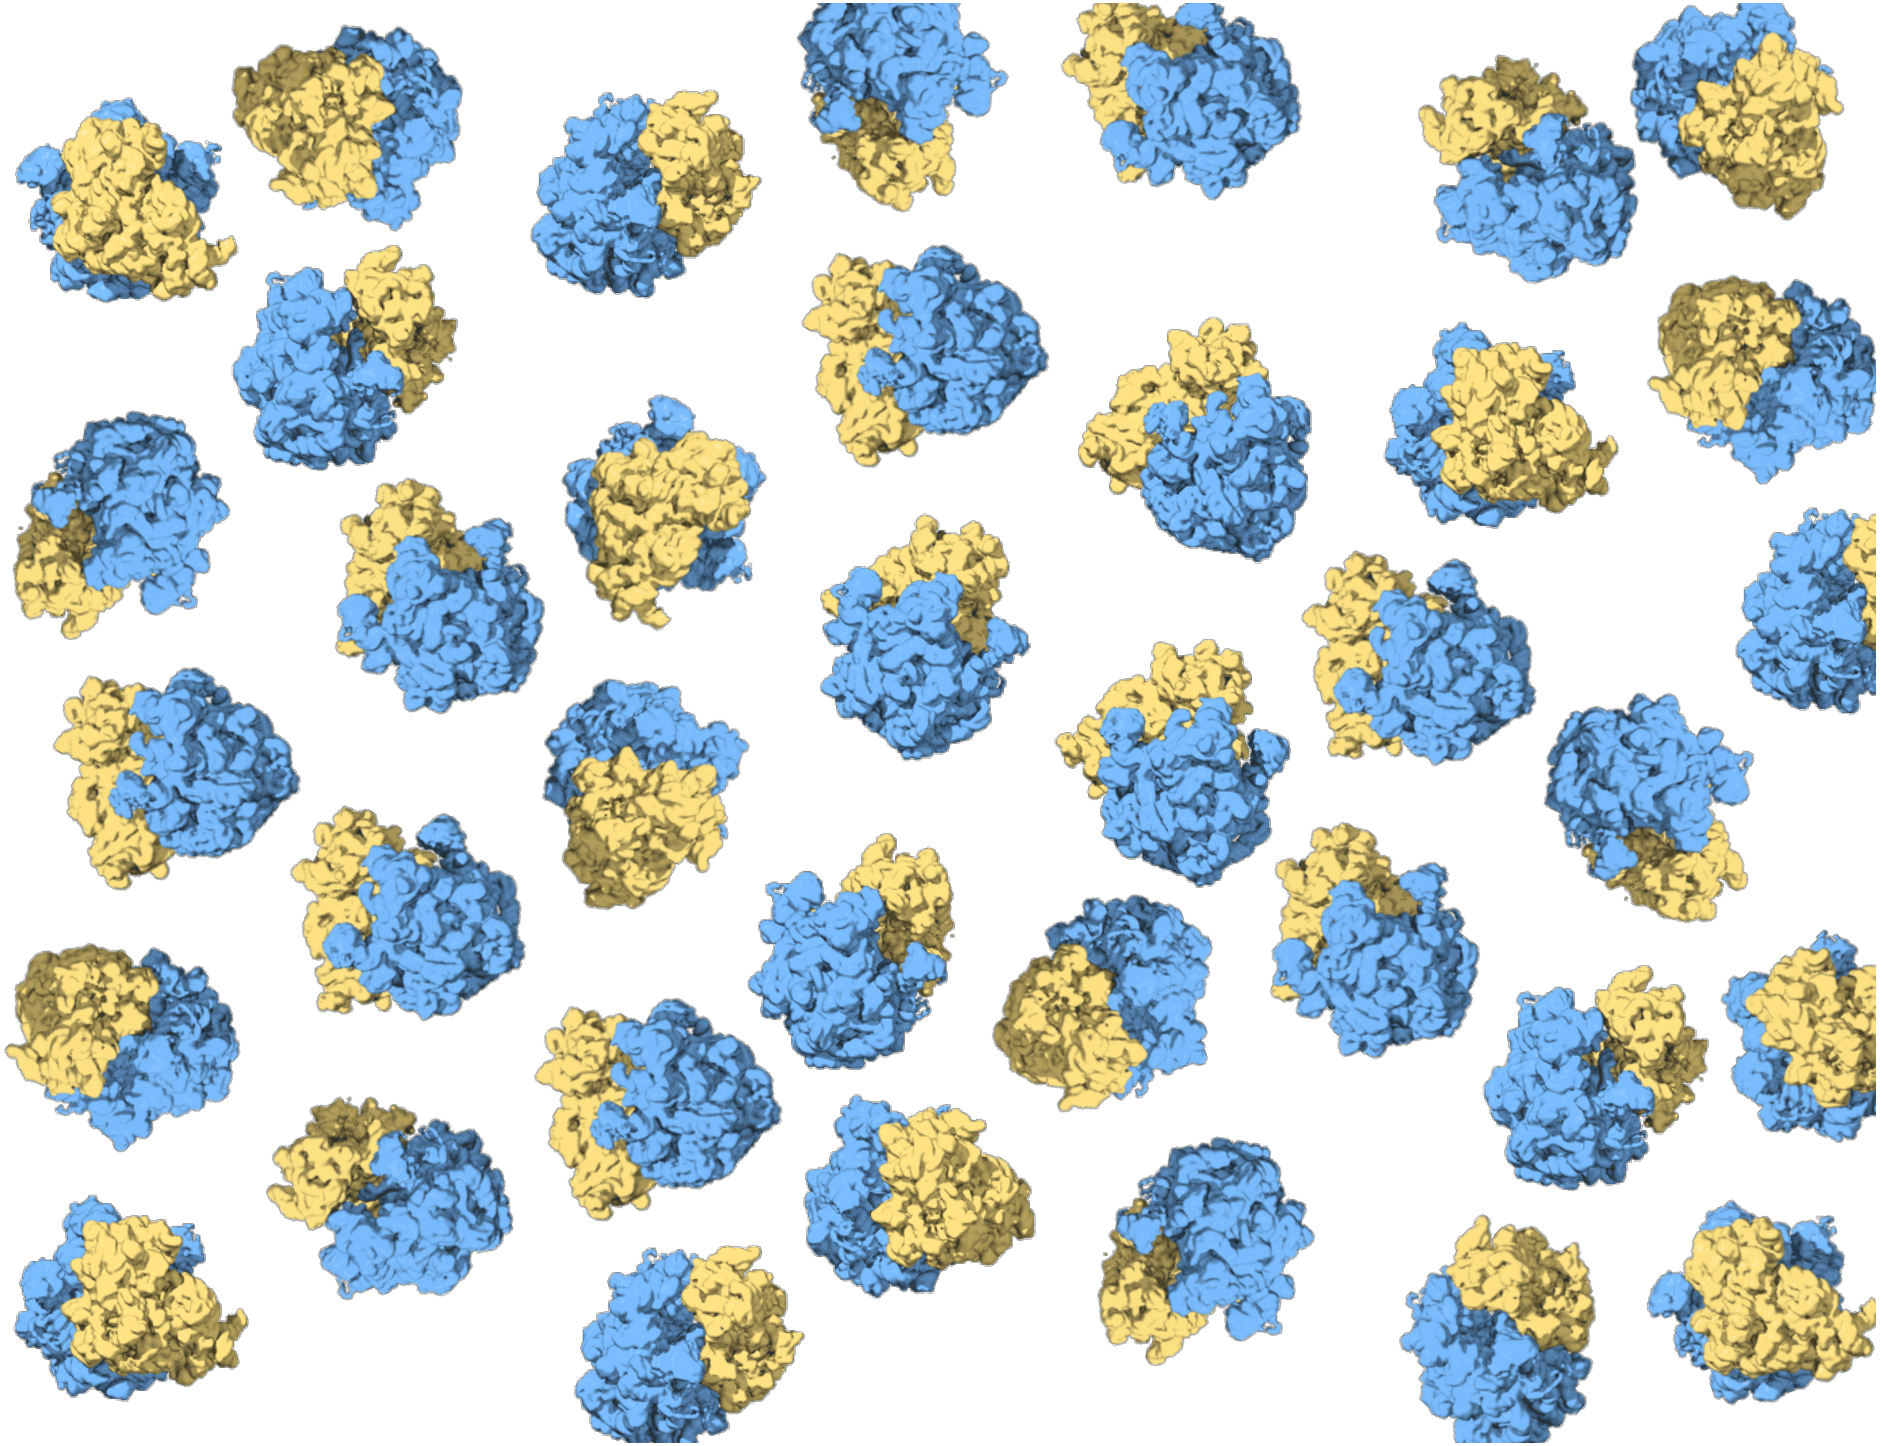

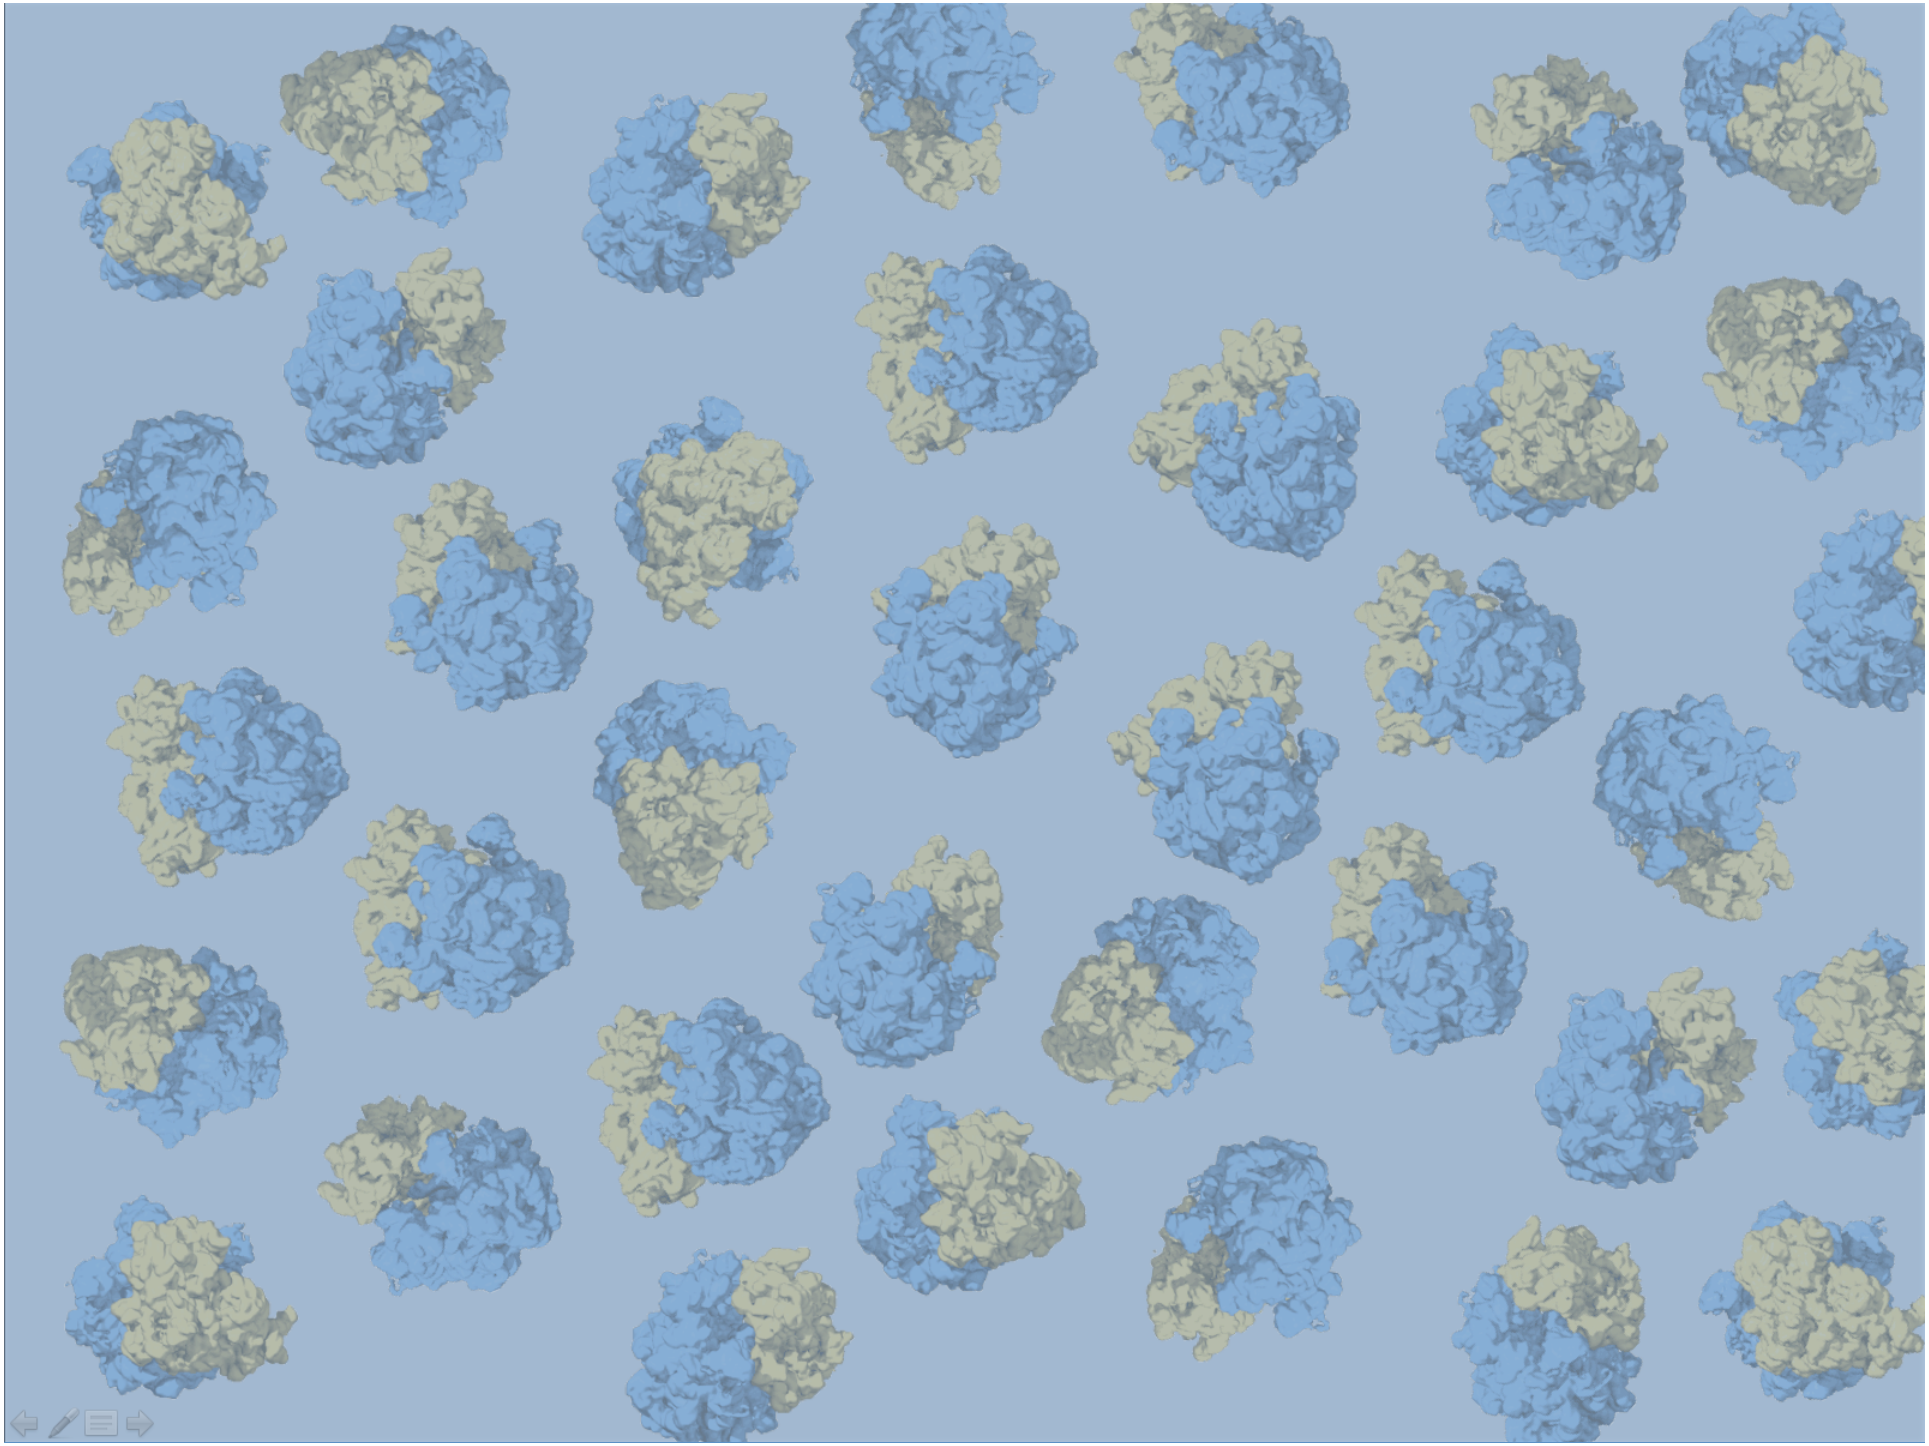

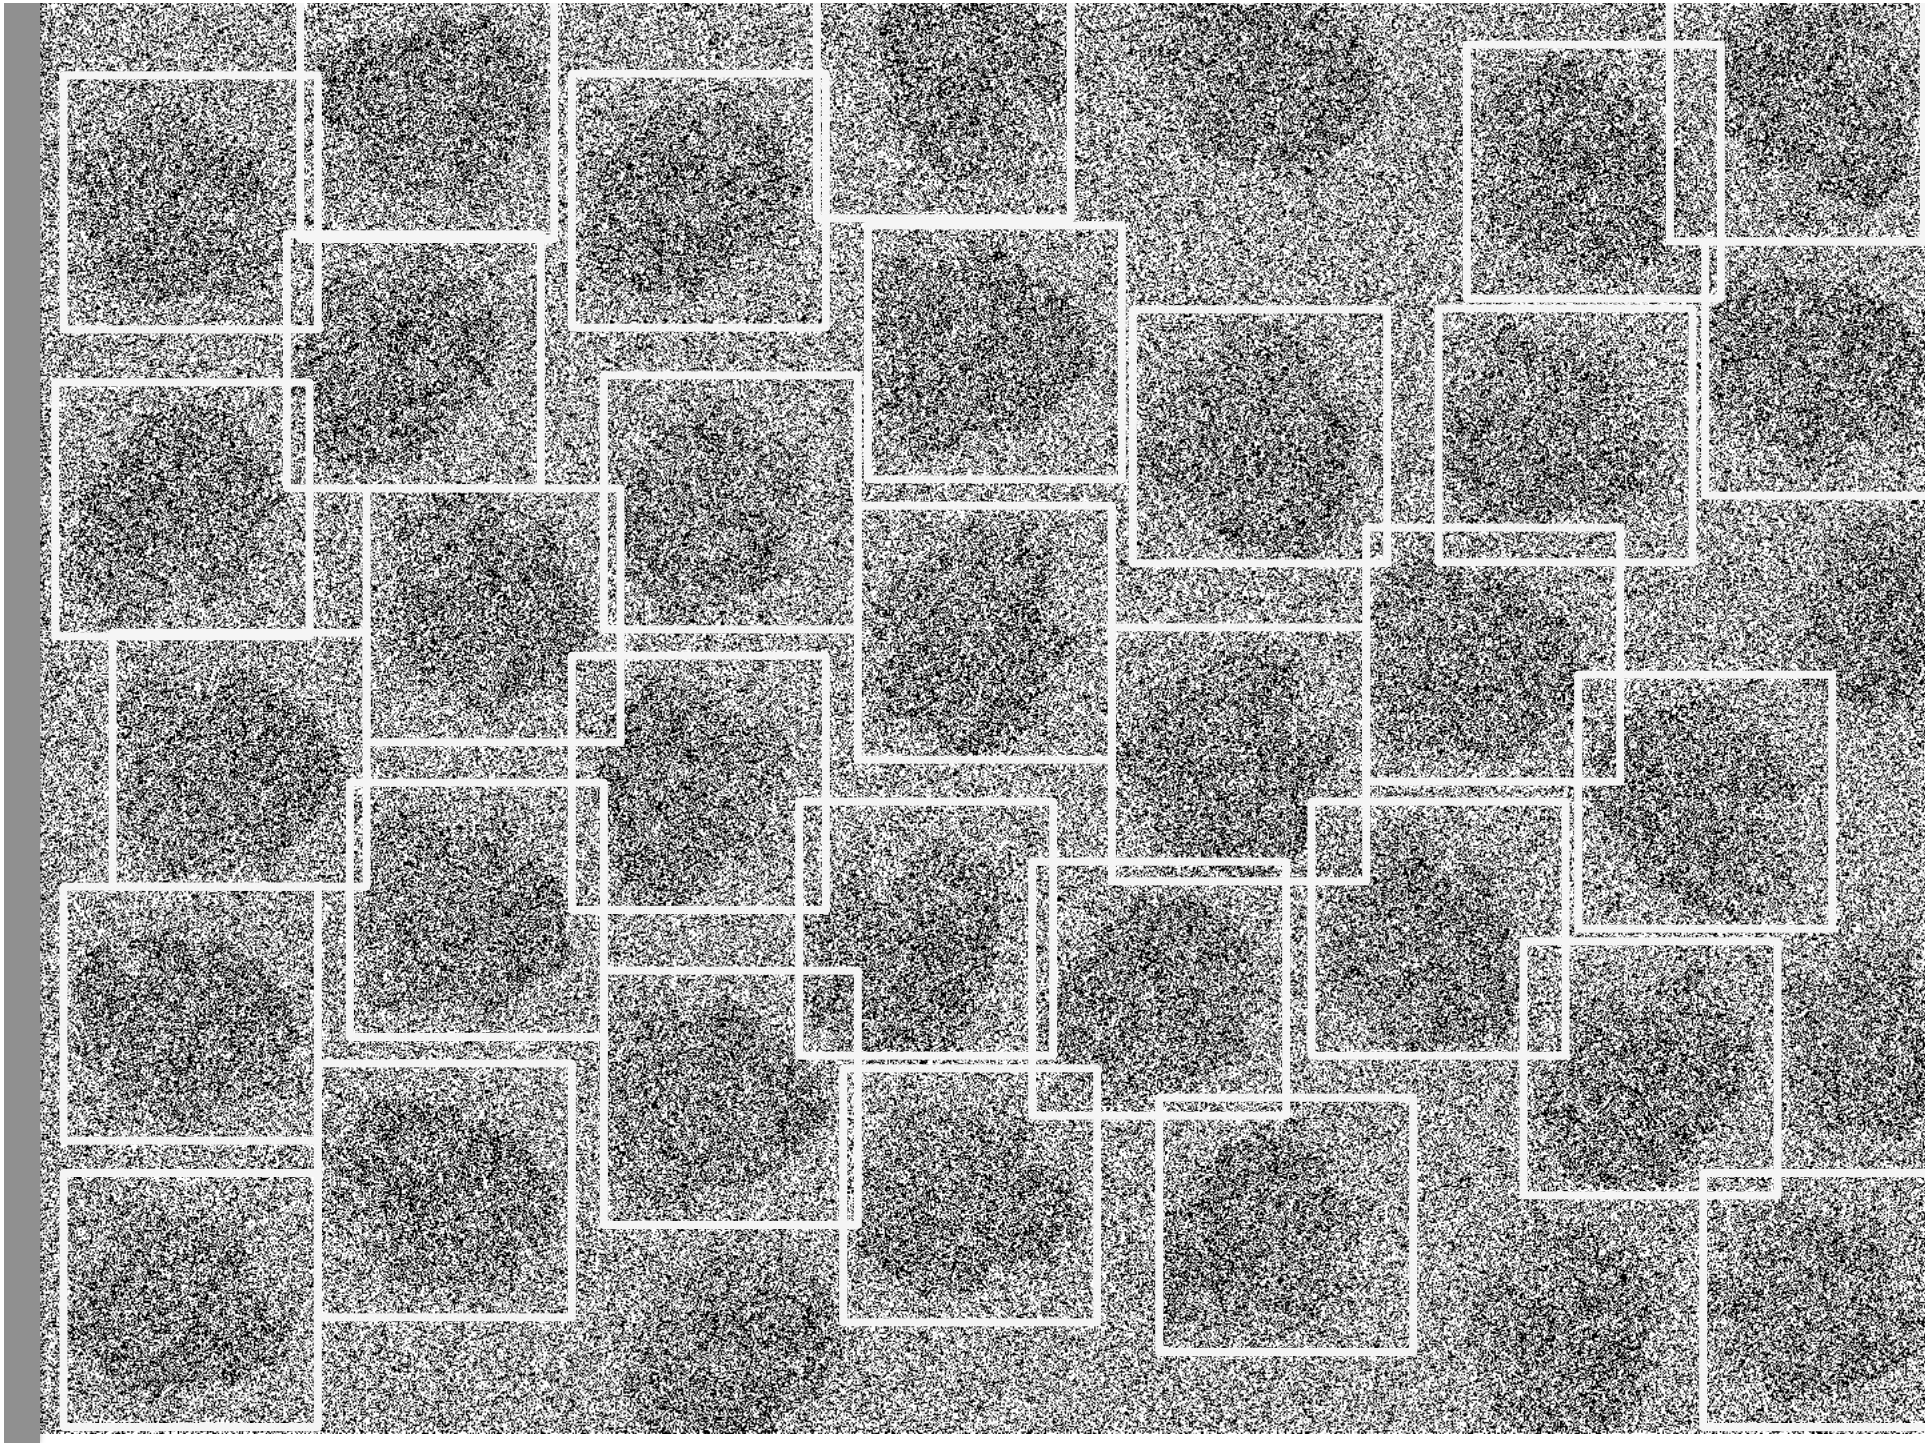

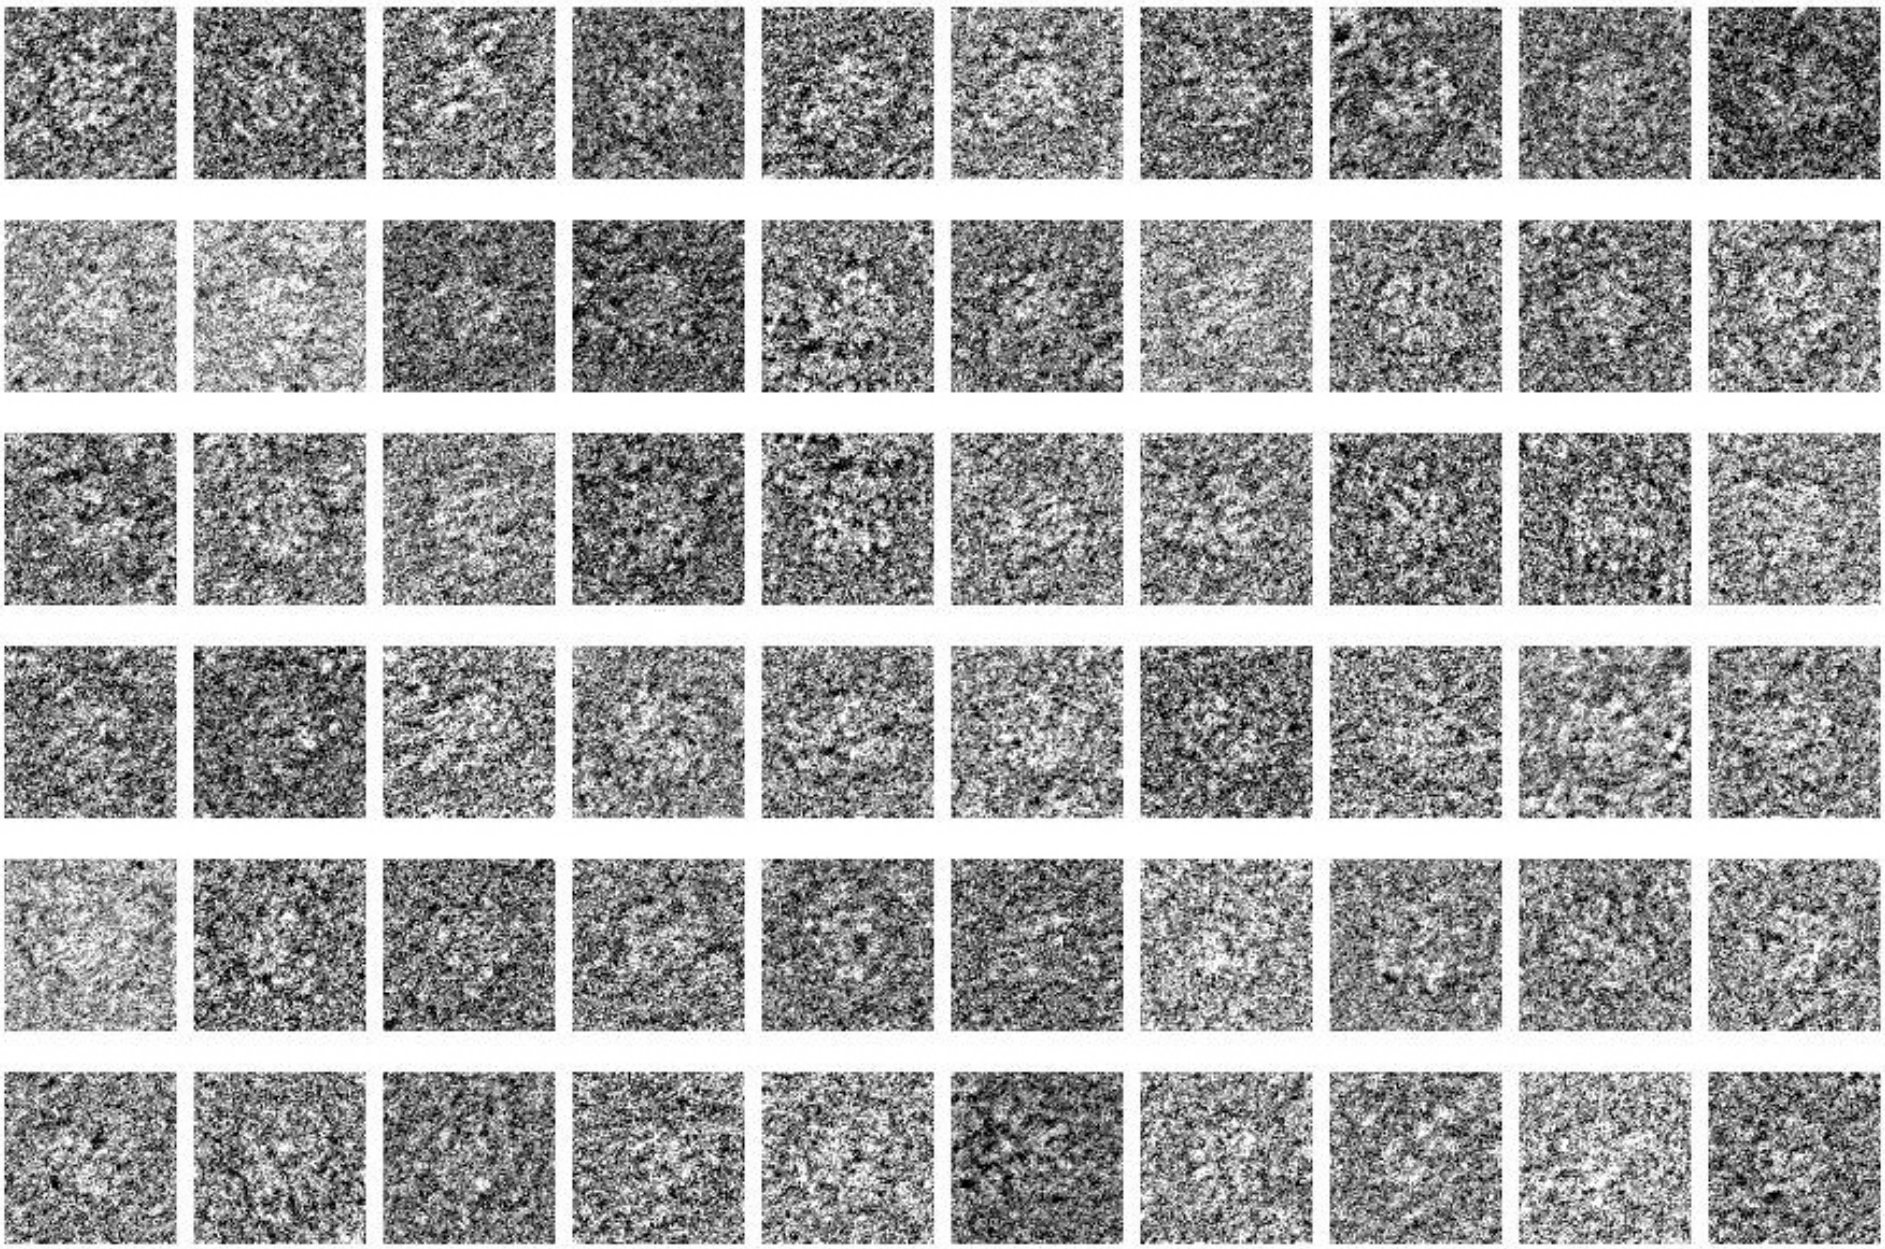

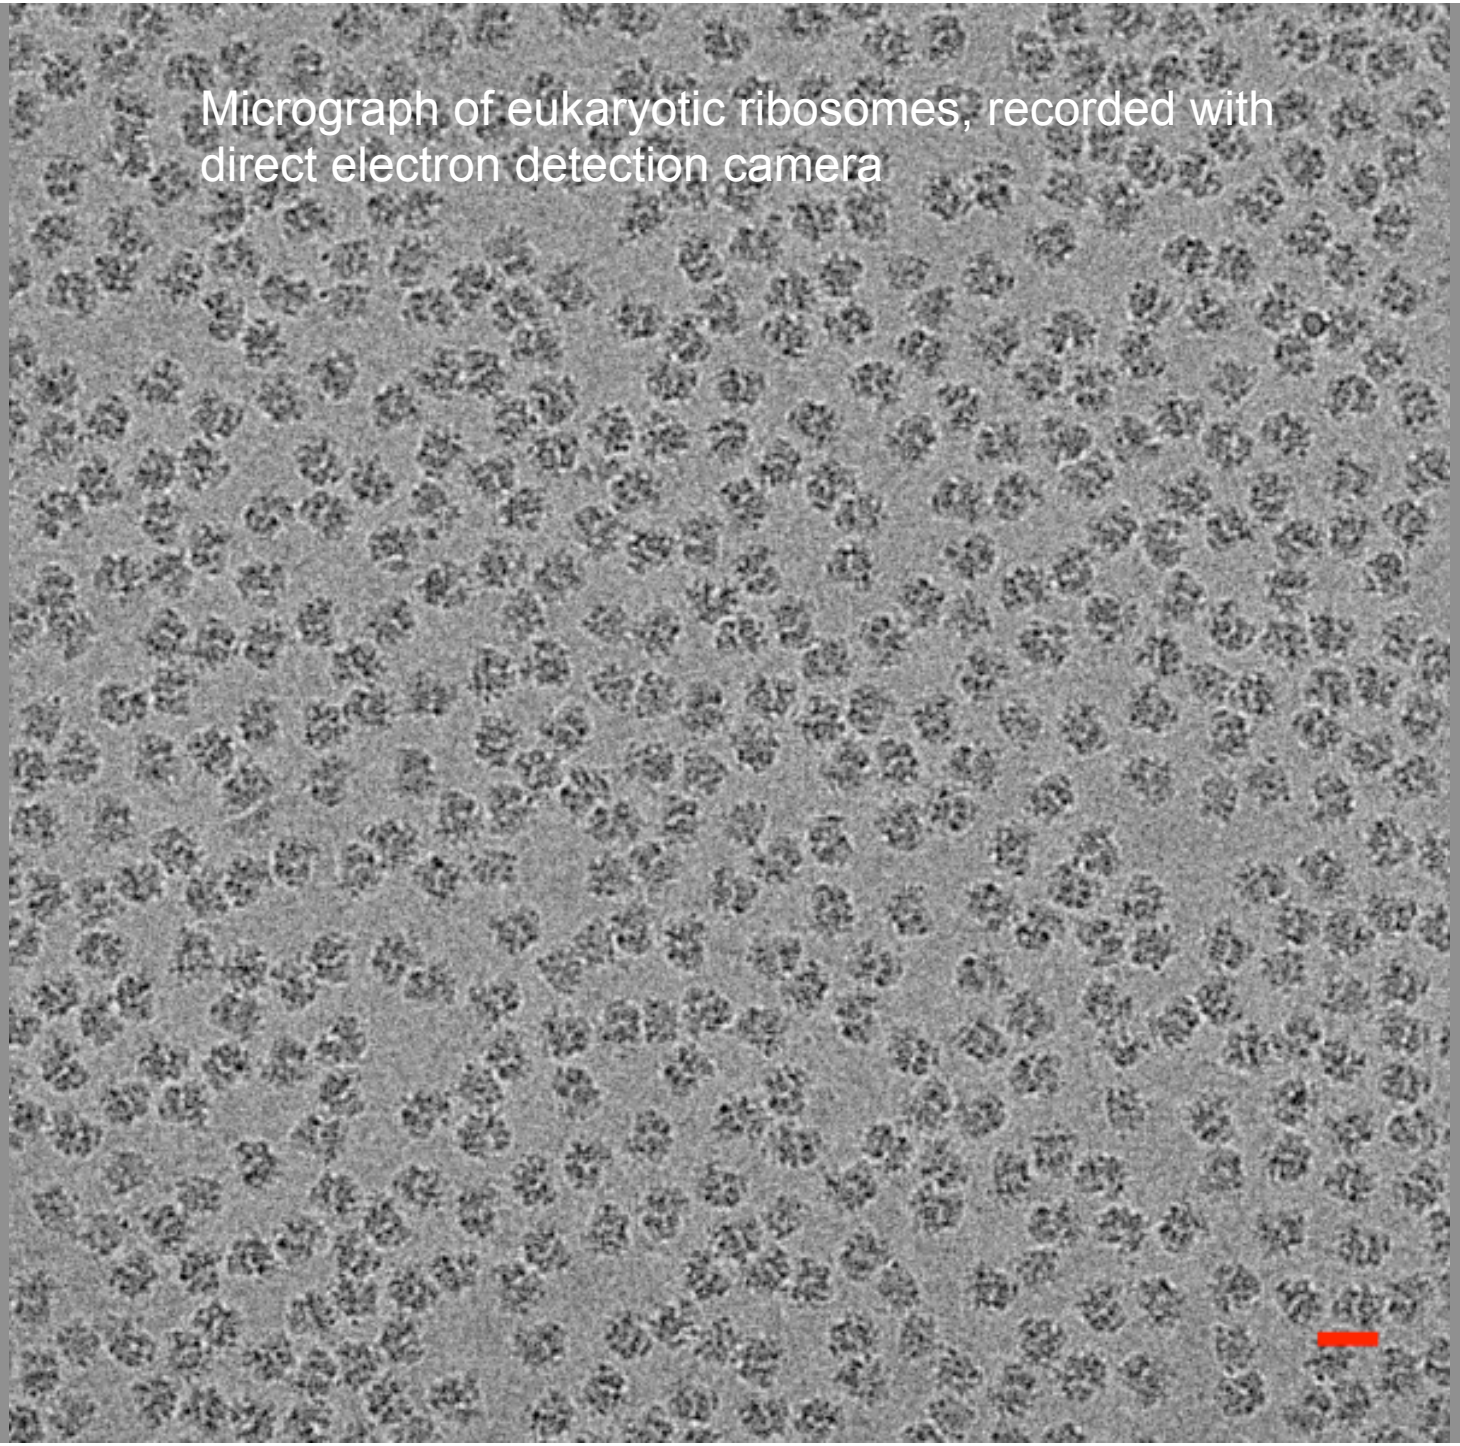
Micrograph of eukaryotic ribosomes, recorded with
direct electron detection camera

How to get from 2D to 3D: The Projection Theorem

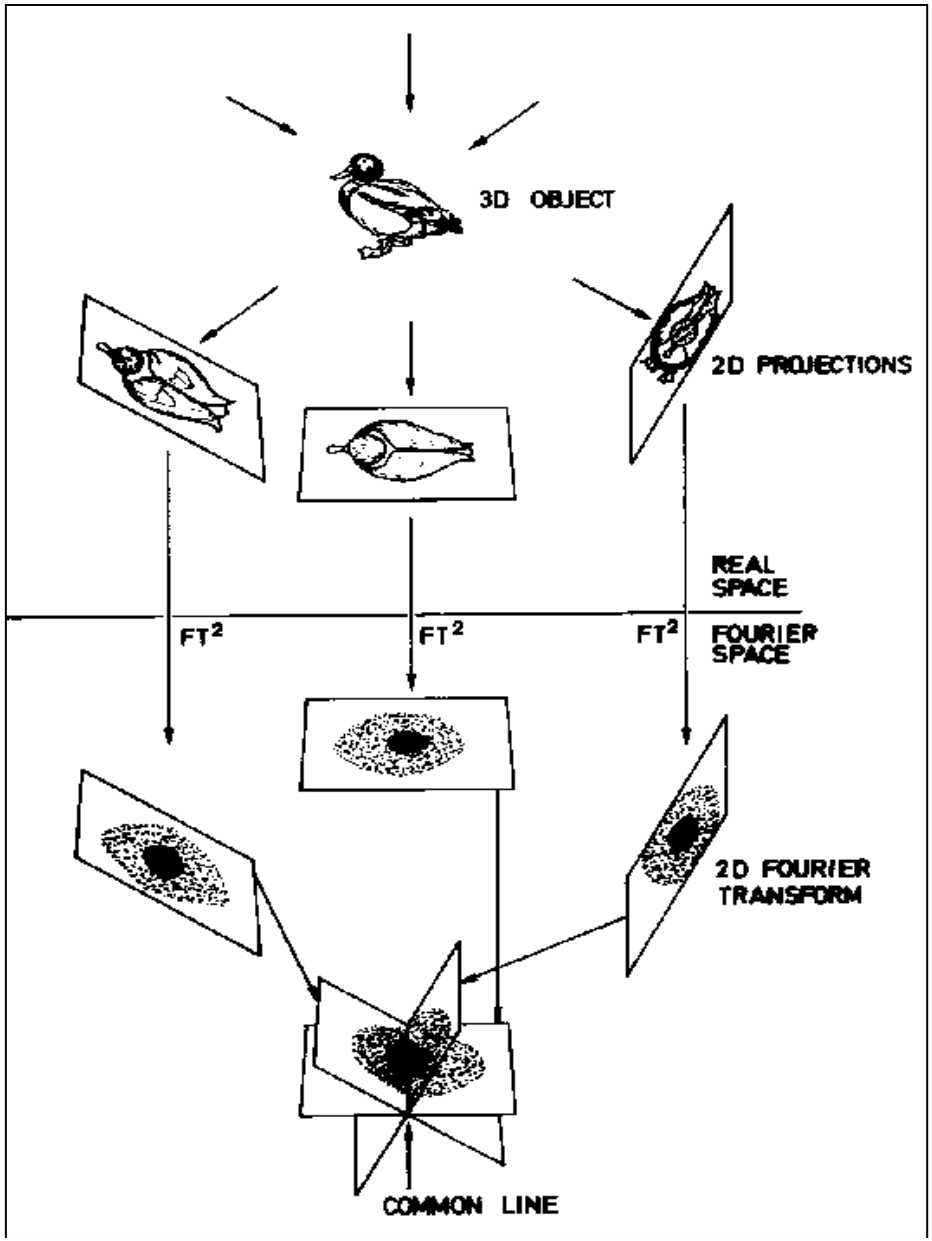

1) The transmission electron microscope forms projections of the 3D object.

2) The Theorem:

"The 2D Fourier transform of the projection of a 3D density is a *central section* of the 3D Fourier transform of the density, *perpendicular* to the direction of projection."

3) It is necessary to collect a sufficient number of projections over a large angular range. From these projections, the object's density distribution can be reconstructed.

First 3D reconstruction from EM images:
3D Reconstruction of Bacteriophage
Tail Using the Fourier-Bessel Approach

1968

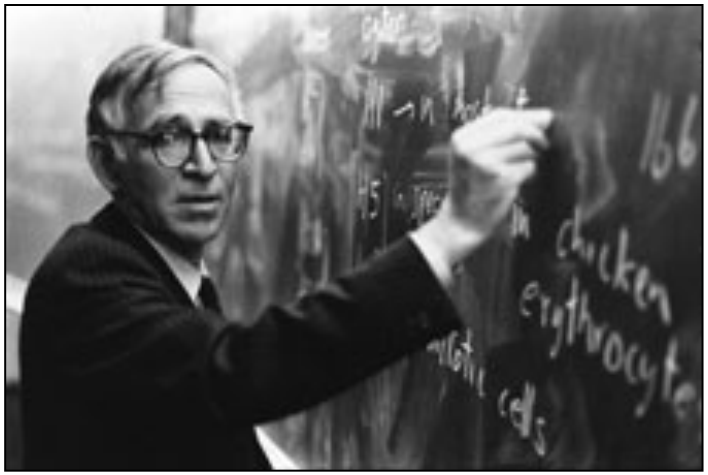

Aaron Klug and David DeRosier, LMB, MRC

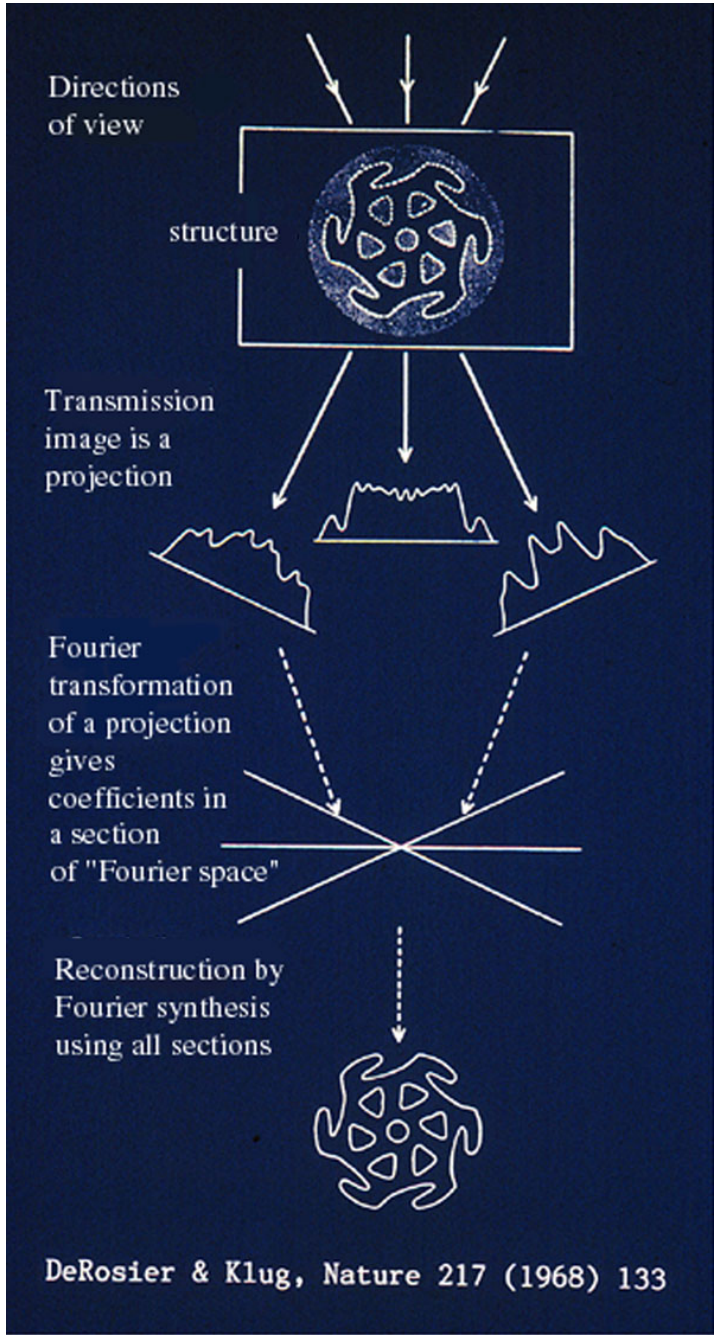

How do we collect the projections?

Three data collection strategies for 3D reconstruction:

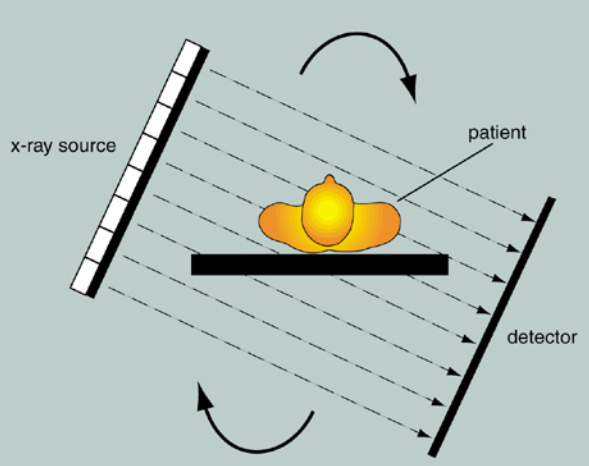

CAT - scan

- beam rotating
- patient stationary

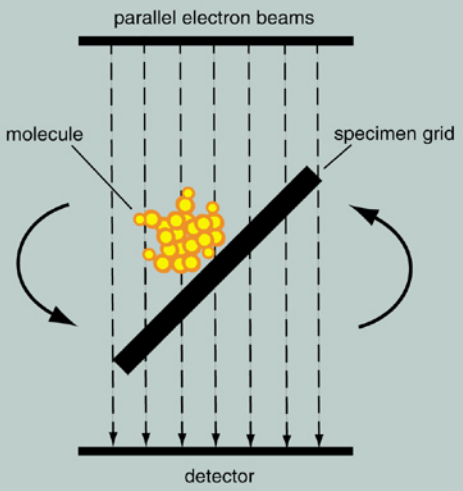

Electron Tomography

- molecule rotating
- beam stationary

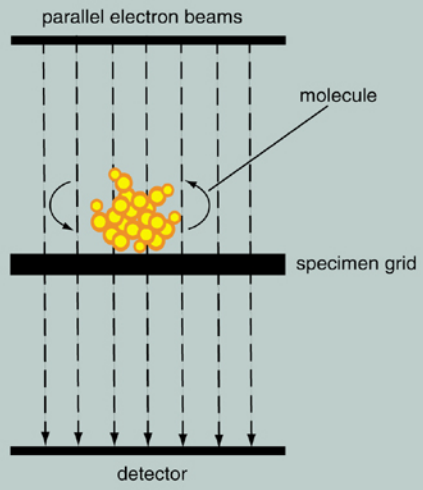

Single particle reconstruction

- molecule "rotating"
- beam stationary

Determine orientations, 3D Reconstruction

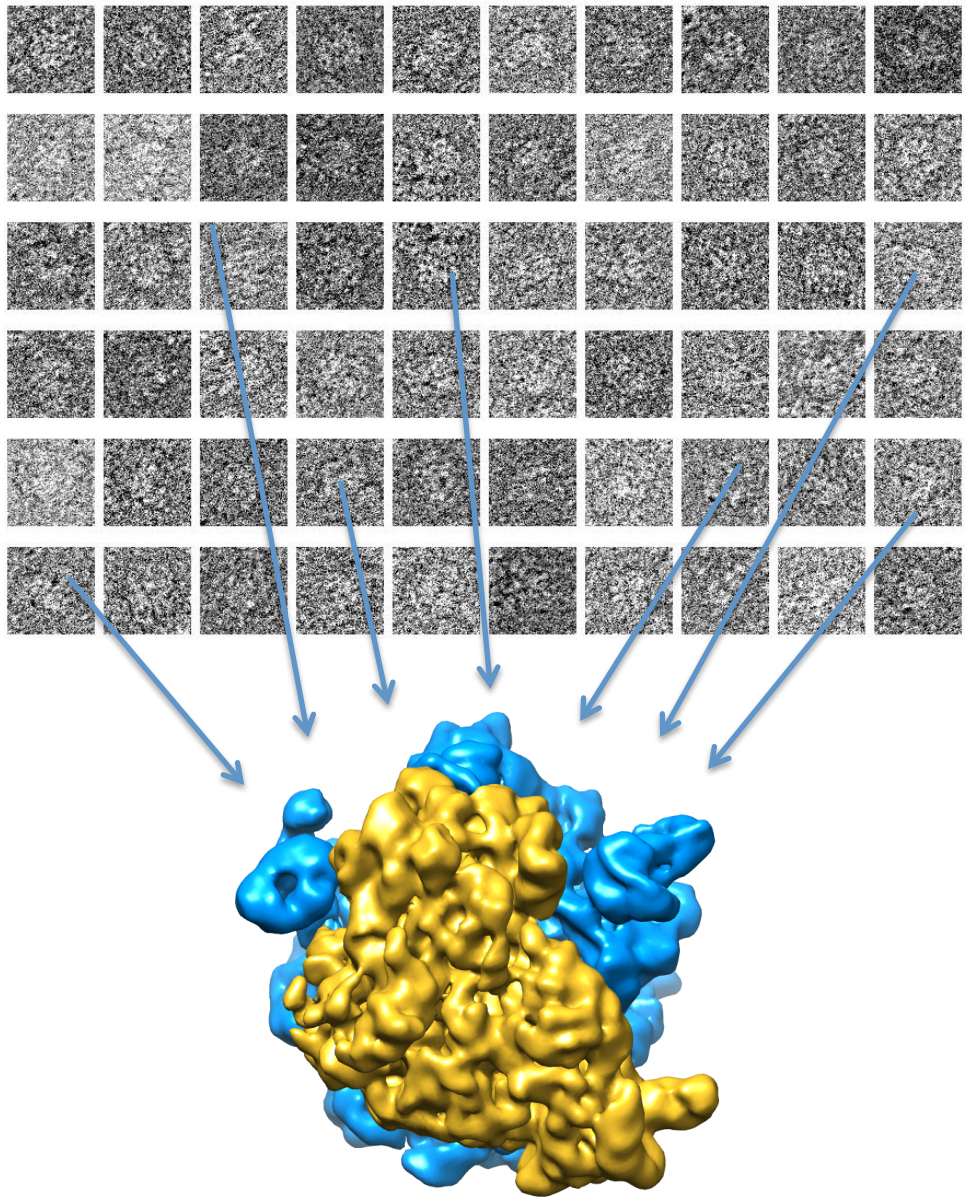

Image classification

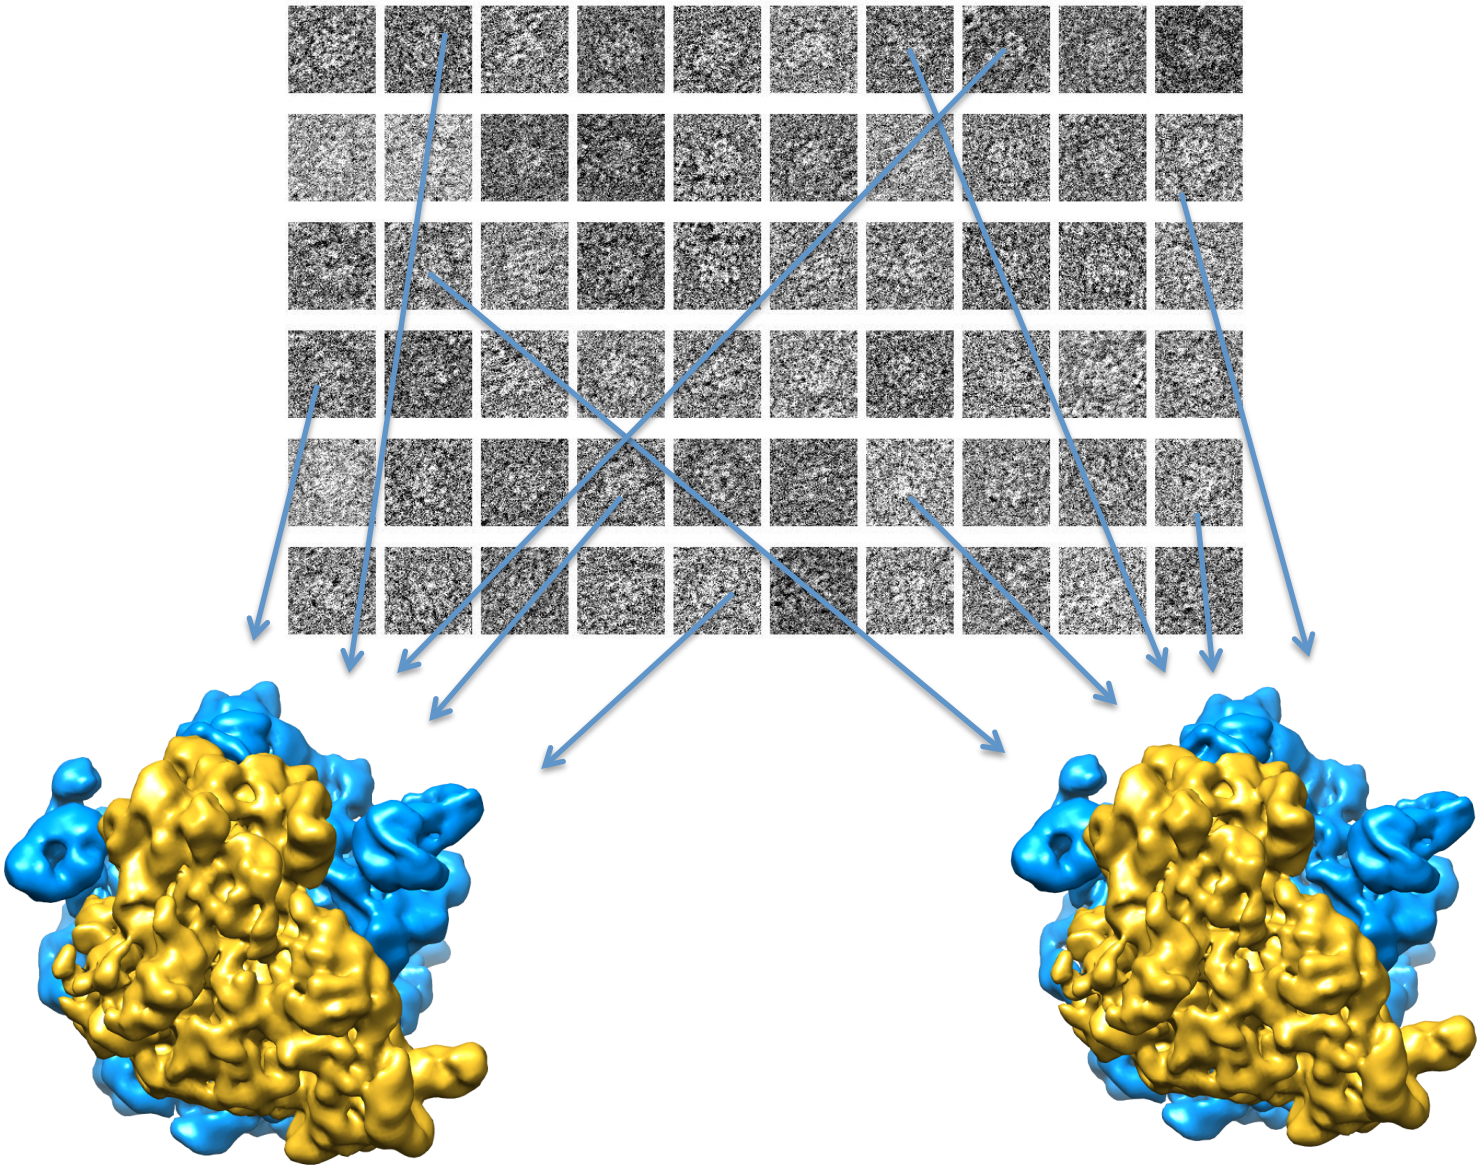

The Single-Particle Approach to Averaging and Reconstruction in EM of Macromolecules

"Single" = unattached, free from contacts with other molecules.
This affects methodology of specimen preparation, electron microscopy, and image processing.

Why single particles?

Advantages:

- native conformation, unaffected by crystal packing
- functionally meaningful states can be visualized
- no part of the molecule needs to be chopped off for visualization
- multiple states visualized from the same sample
- ideal for looking at the dynamics of a molecular machine

Disadvantages (up to 2012):

- large computational challenges
- atomic resolution difficult to achieve for particles lacking symmetries

The Single-Particle Approach to Averaging and Reconstruction in EM of Macromolecules

"Single" = unattached, free from contacts with other molecules.
This affects methodology of specimen preparation, electron microscopy, and image processing.

Why single particles?

Advantages:

- native conformation, unaffected by crystal packing
- functionally meaningful states can be visualized
- no part of the molecule needs to be chopped off for visualization
- multiple states visualized from the same sample
- ideal for looking at the dynamics of a molecular machine

Disadvantages (2012 onwards):

- large computational challenges
- ~~atomic resolution is as yet difficult to achieve without symmetries~~ (DED cameras)

Single-Particle Reconstruction
Main initial assumptions:

1) All particles in the specimen have (approx.) identical structure
2) All are linked by 3D rigid body transformations (rotations, translations)
3) Particle images are interpreted as a "signal" part (= the projection of the common structure) plus "noise"

Important requirement:

even angular coverage, without major gaps.

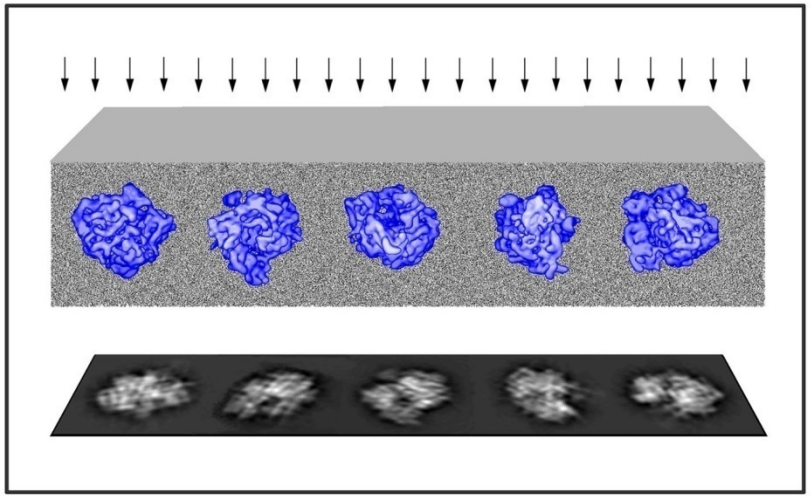

Single-Particle Methods
are also used for ordered assemblies

Ordered assemblies never strictly follow the symmetries, so image processing relying entirely on the exact validity of the symmetries will fail to retrieve the high-resolution information. Applications of single particle methods:

- Helical order (e.g., acetylcholine receptor, actin-myosin fiber) – Ed Egelman's (and Pawel Penczek's) newer methods.
- 2D crystalline order (e.g., purple membrane protein) – "unbending" approach vs "patch averaging" method.
- Icosahedral order (e.g., adenovirus) – radical departure from initial MRC-based analytical methods.

General principle: (1) use symmetry to roughly locate the repeats, then (2) refine position by using the cross-correlation function (CCF)

“Shot Noise”

At the low exposure settings (e.g., 10 e/Å^2) required to avoid radiation damage, the fluctuations of the electron distribution is a serious source of noise, called shot noise. Low-exposure images typically have an SNR of 0.1

Only by averaging over a sufficient number of particle projections can the original signal be retrieved.

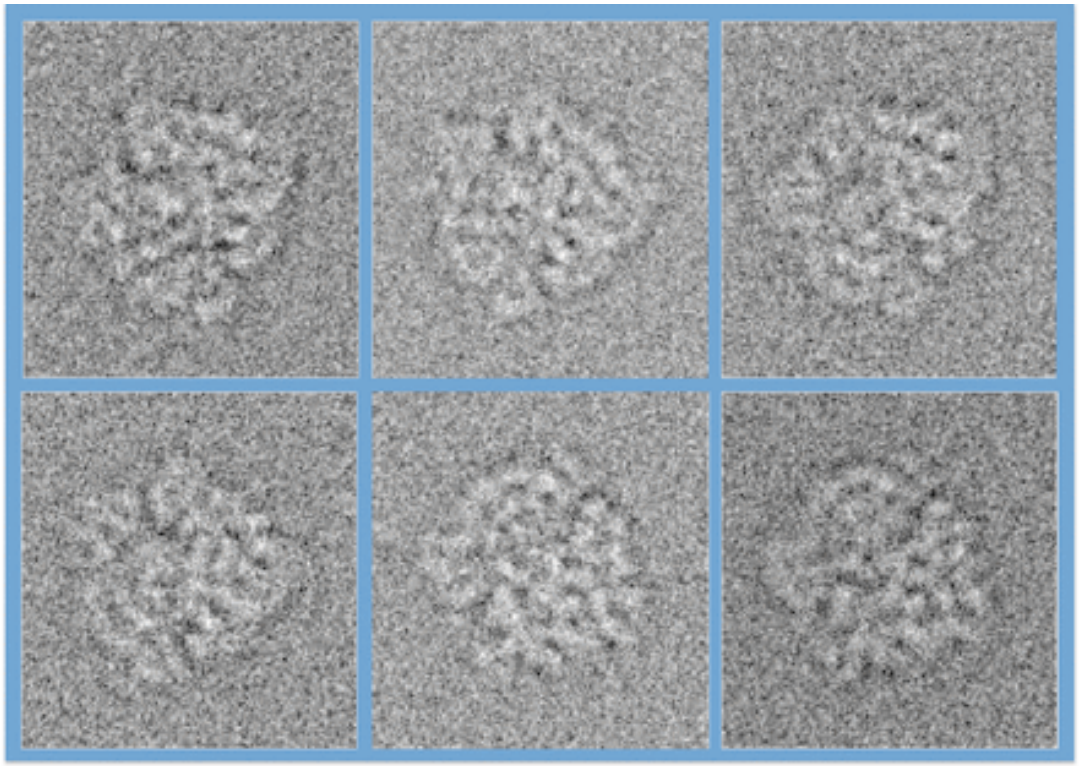

Simulated images of ribosome at SNR ~ 0.1

Schwander et al., Phil. Trans. Roy. Soc. 2014

“Structural Noise”

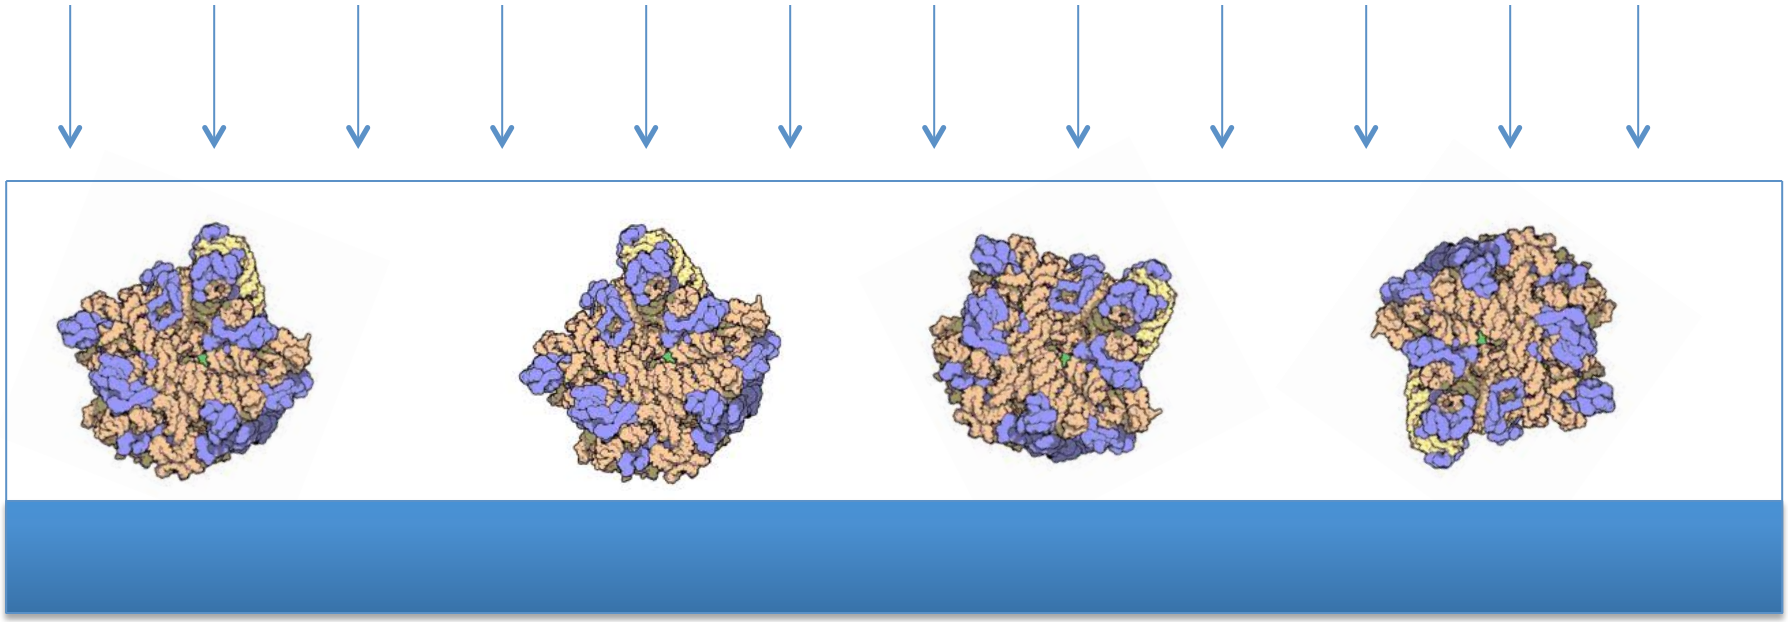

The matrix of ice, and carbon deposit (if used) has a unique structure which is superimposed when a projection image is formed. When images of particles are averaged, the superimposed structure of the surrounding must be considered “noise” since it is not reproducible from one particle to the next..

Signal and noise; definition of SNR

- Signal $s(\mathbf{r})$ (predictable, deterministic, originating from the object)

  versus

- Noise $n(\mathbf{r})$ (stochastic; unrelated to the signal; aperiodic [no two realisations are the same])
- Signal-to-noise ratio (SNR) = signal variance/noise variance

- Averaging over N noisy realizations of a signal increases the SNR by a factor of N
- *Note that what is signal and what is noise in a given experiment depends on the way the experiment is designed.*

Two-dimensional processing: averaging of images to eliminate noise

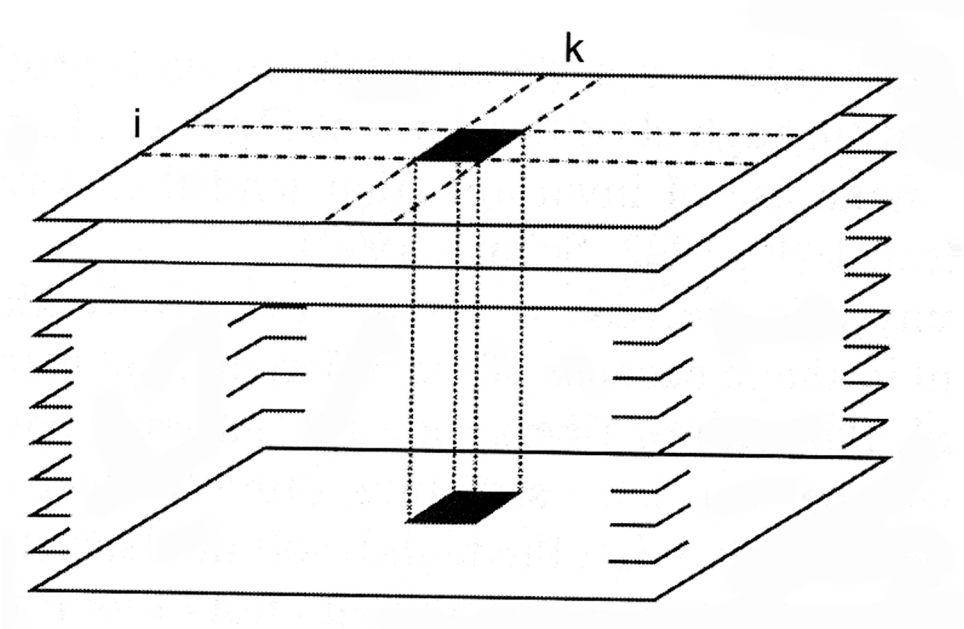

Averaging requires the images to be aligned: in the array, each image element, or pixel, must refer to the same point of the object.

Note: averaging in 2D only makes sense for molecules presenting the exact same view

Averaging, to improve the SNR

Averaging can be done in one of two ways:

-- either --

make use of order or symmetries to locate repeats

– or --

make use of cross-correlation search to locate repeats

In the first case, Fourier methods can be used instead of real-space averaging

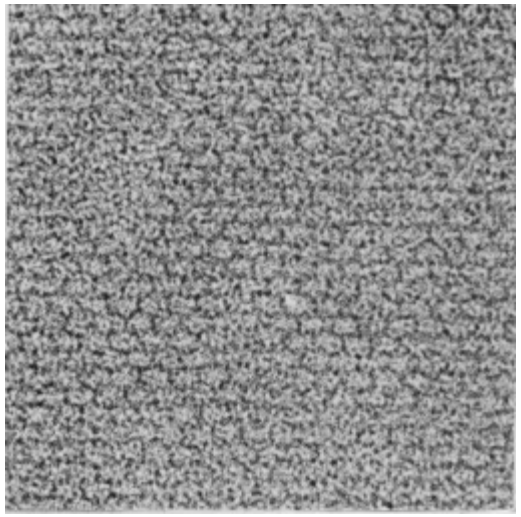

Projection image of molecule ordered on a periodic lattice.
Noise is intermixed.

2D Fourier transform

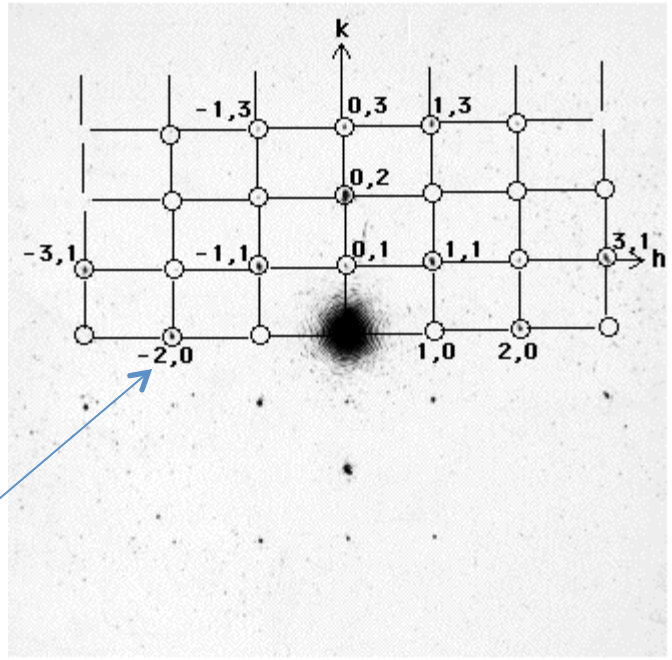

Fourier transform is concentrated in spots on the reciprocal lattice. Noise is separated.

Indices define position on reciprocal grid
Fourier transform is complex, each spot described by an amplitude and a phase of a component wave

Inverse Fourier transform

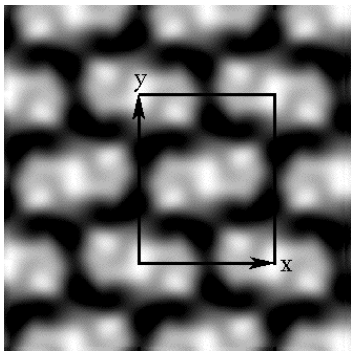

Noise-free projection image

1) make use of order or symmetries to locate repeats

2) make use of cross-correlation search to locate repeats.

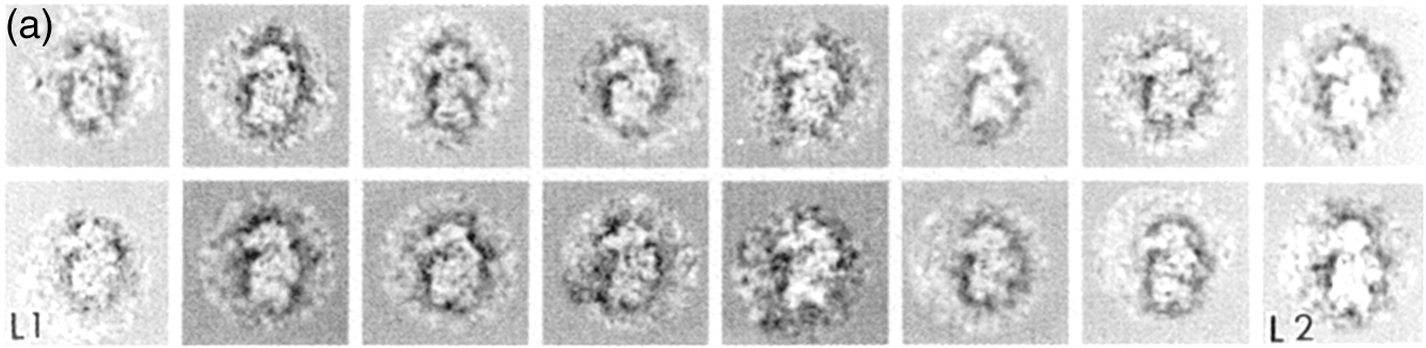

40S ribosomal subunits of HeLa cells, negatively stained

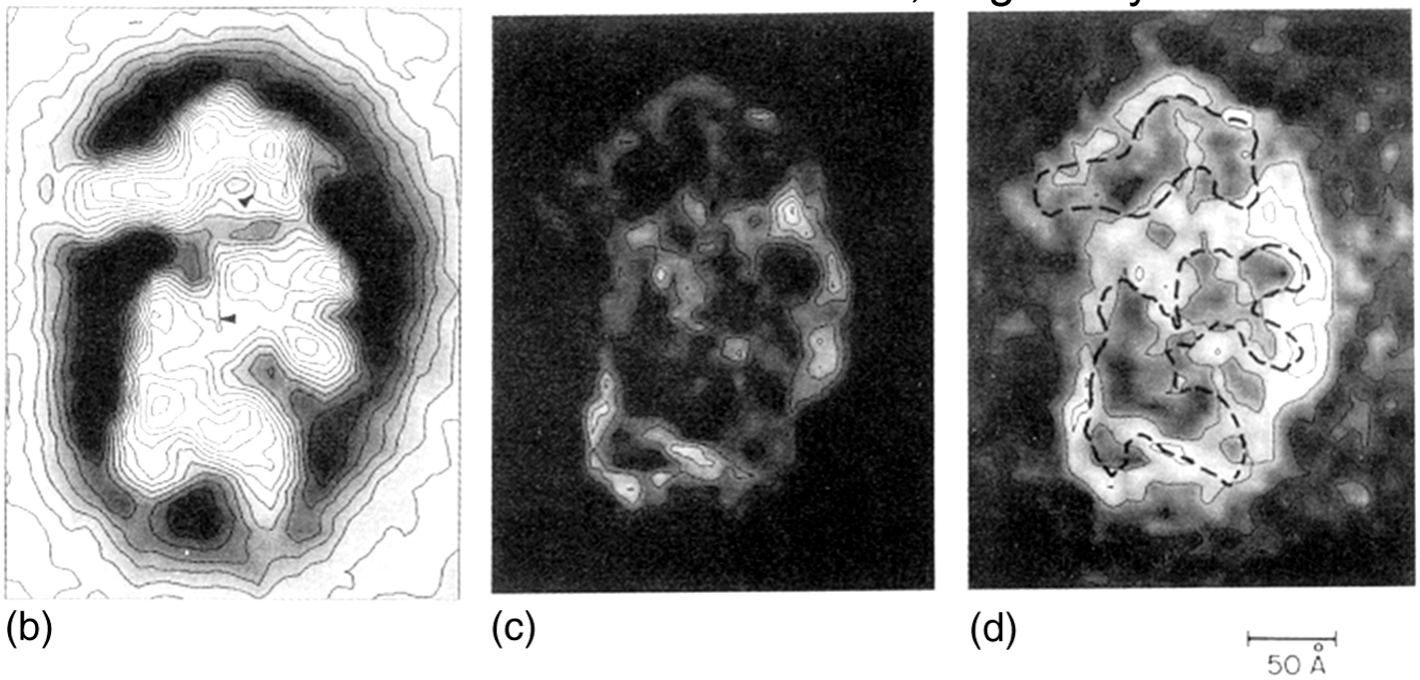

average variance map s.d. map

Frank et al., Science 1981

Alignment of single-particle projections ("particles") is achieved by cross-correlation

- Translational cross-correlation function (CCF)

  Discrete, unnormalized version:

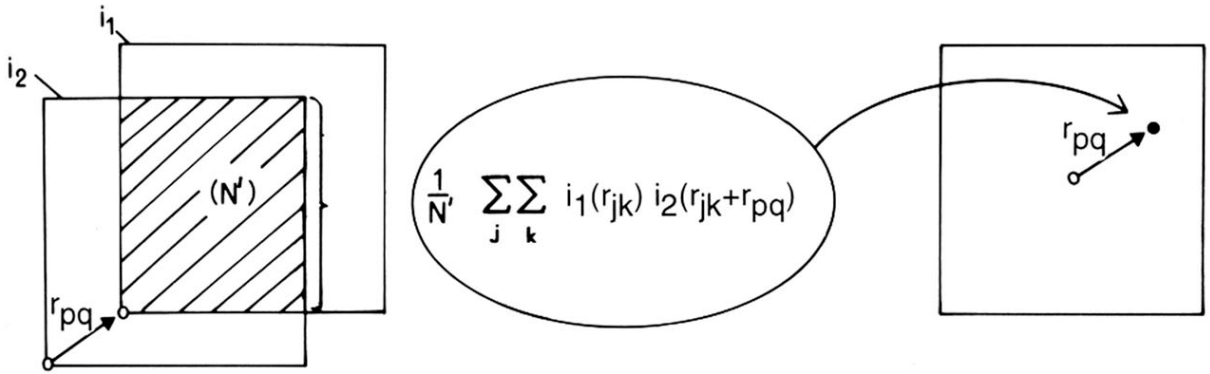

Fig. 3.8. Definition of the cross-correlation function. Image 1 is shifted with respect to image 2 by vector $\mathbf{r}_{pq}$. In this shifted position, the scalar product of the two images arrays is formed and put into the CCF matrix at position (p, q). The vector $\mathbf{r}_{pq}$ is now allowed to assume all positions on the sampling grid. In the end, the CCF matrix has an entry in each position. From Frank (1980). Reproduced with permission of Springer-Verlag, New York.

- Rotational CCF

2D Fourier transform

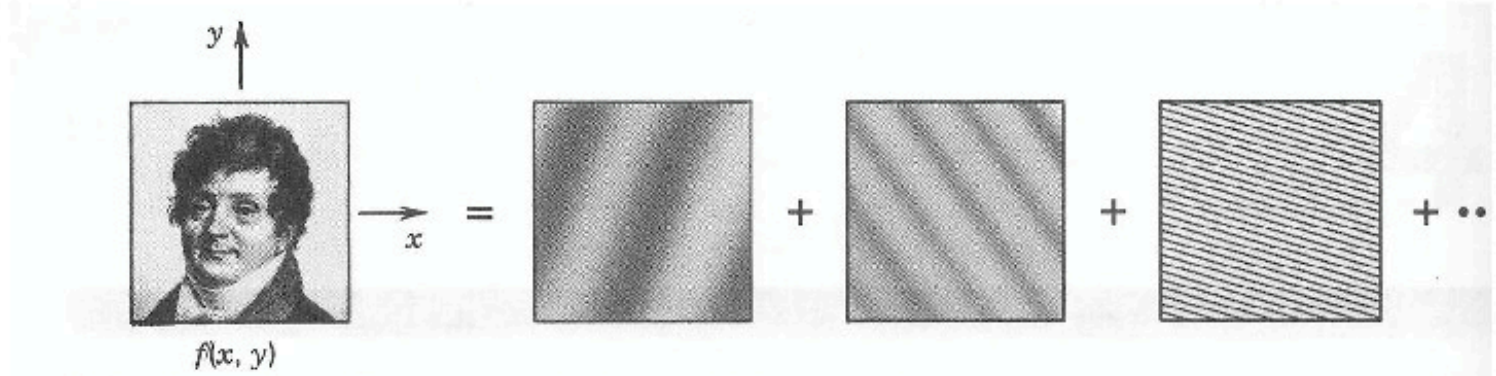

$$F(K_x, K_y) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} f(x,y) \exp(i(K_x x + K_y y)) dx dy$$

An image can be considered a superposition of sine waves of different spatial frequencies running in different directions. Each sine wave is characterized by an **amplitude** and a **phase.**
Alternative representation (as in this diagram) employ complex exponential functions with complex coefficients.

The Discrete Fourier Transform (DFT) in 2D can be defined as:

$$F(u,v) = \frac{1}{MN}\sum_{m=0}^{M-1}\sum_{n=0}^{N-1} f(m,n)\, exp\left[-2\pi i\left(\frac{mu}{M}+\frac{nv}{N}\right)\right]$$

$$u = 0,1,\dots,M-1, \qquad v = 0,1,\dots,N-1$$

The inverse 2D DFT is given by:

$$f(m,n) = \sum_{u=0}^{M-1}\sum_{v=0}^{N-1} F(u,v)\, exp\left[2\pi i\left(\frac{mu}{M}+\frac{nv}{N}\right)\right]$$

$$m = 0,1,\dots,M-1, \qquad n = 0,1,\dots,N-1$$

An image f(m, n) is represented as a finite series of 2D exponentials with complex coefficients F(u,v).

Discrete Fourier representation implies infinite repetition of the image

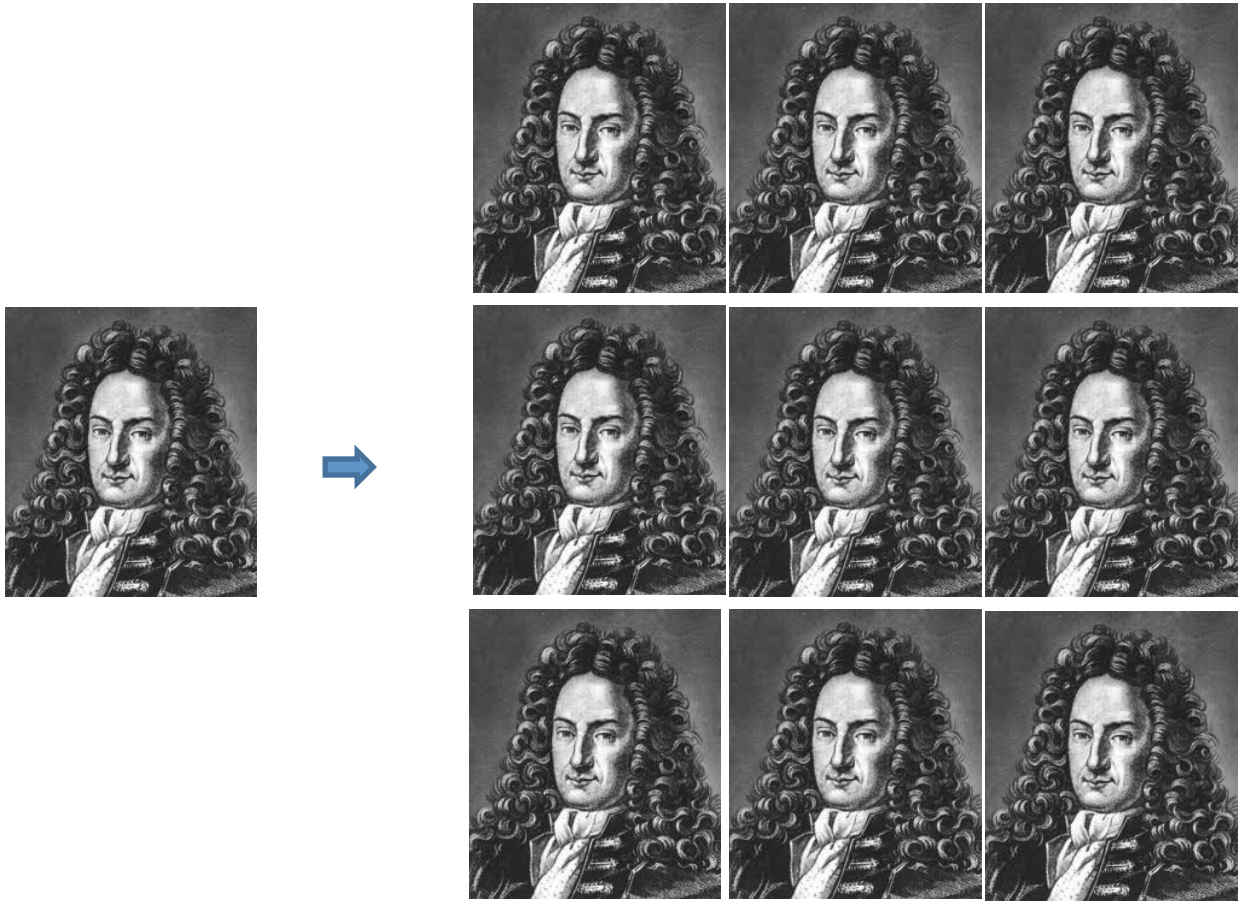

notations:

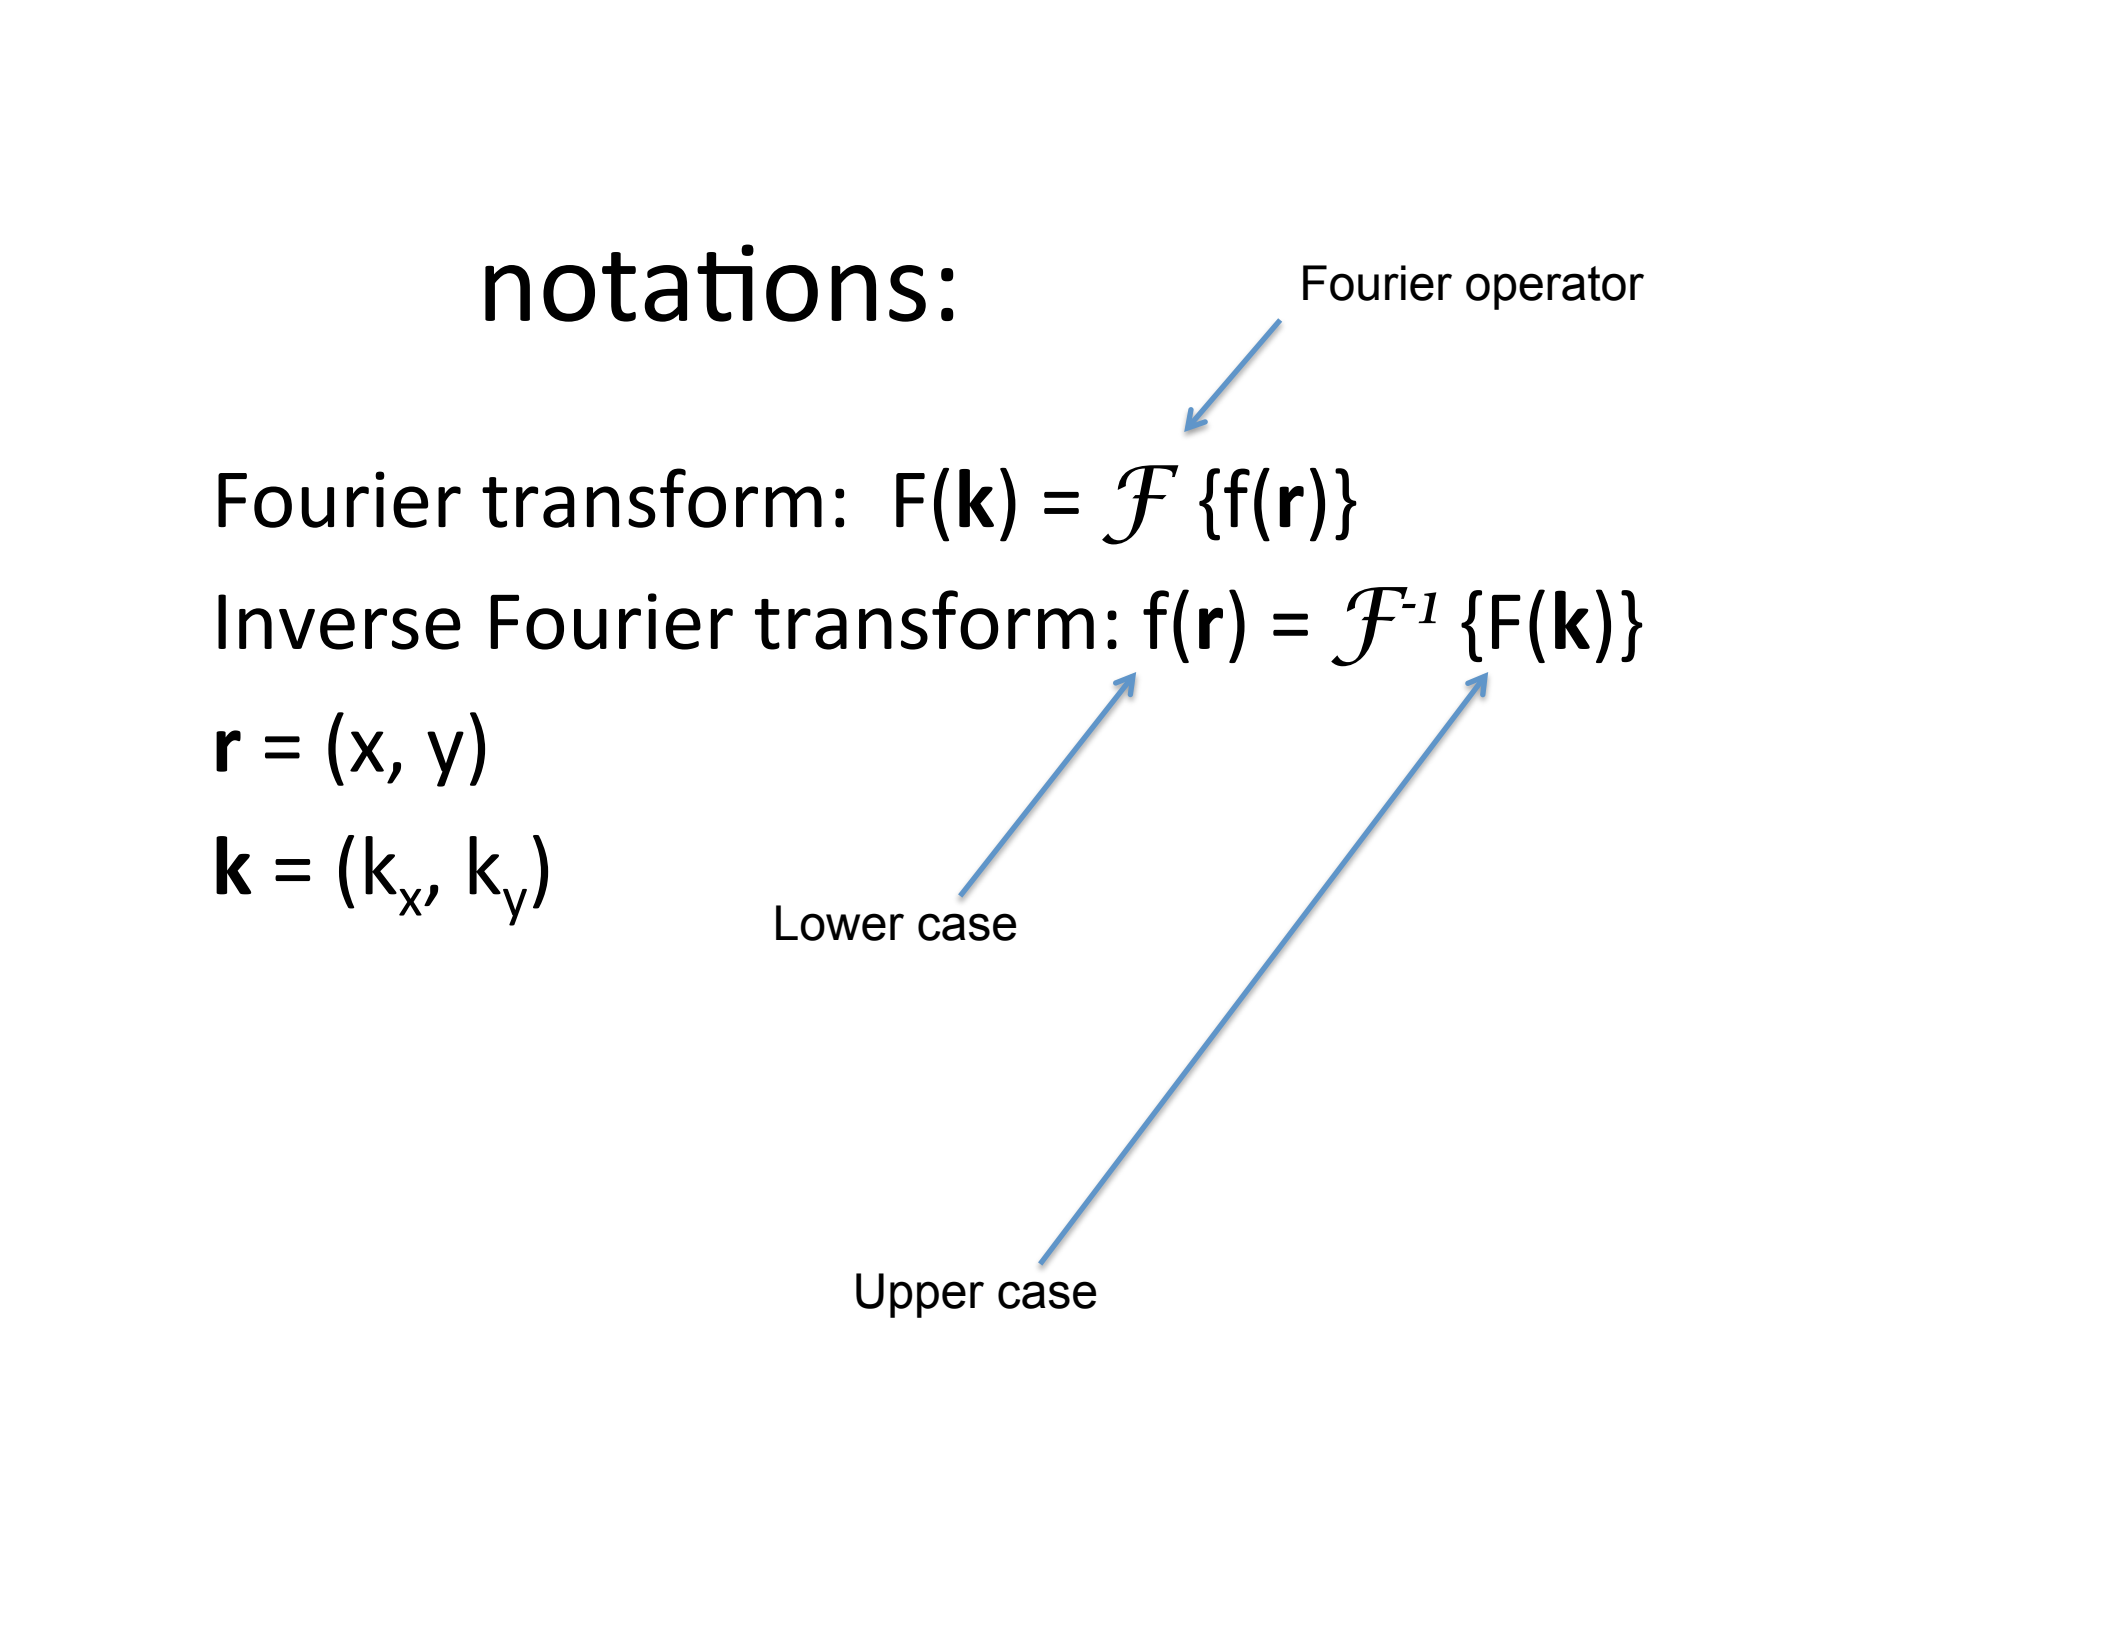
Fourier operator
Fourier transform: F(**k**) = $\mathcal{F}$ {f(**r**)}
Inverse Fourier transform: f(**r**) = $\mathcal{F}^{-1}$ {F(**k**)}
r = (x, y)
k = (k_x, k_y)
Lower case
Upper case

Parseval's Theorem -- Conservation of Power

$F(\mathbf{k}) = \mathcal{F}\{i(\mathbf{r})\}$

definition: $P(\mathbf{k}) = |F(\mathbf{k})|^2$ is the "Power spectrum"

Total power is the same in real and Fourier space:

$$\int_{\emptyset} |F(\mathbf{k})|^2 \, d\mathbf{k} = \int |i(\mathbf{r}) - avrg|^2 \, d\mathbf{r}$$

where $avrg = 1/area \times \int i(\mathbf{r}) d\mathbf{r}$

and subscript $_{\emptyset}$ indicates "exclude origin in the integration"

Application: Signal-to-Noise ratio can be computed in Fourier space:

$$SNR = \int_{\emptyset} |S(\mathbf{k})|^2 \, d\mathbf{k} / \int_{\emptyset} |N(\mathbf{k})|^2 \, d\mathbf{k}$$

Point spread function and Contrast transfer function

In an optical instrument, the aperture limit, the aberrations of the lens and other imperfections have the effect that a single point in the object is imaged as an extended 2D function, the so-called Point Spread Function (PSF)

The FT of the PSF is known as the *Contrast Transfer Function* (CTF).

In the EM, the CTF is given by

$$CTF = sin(\gamma(\mathbf{k}))$$

where

$$\gamma(\mathbf{k}) = -\pi\lambda\left[\Delta z + \frac{z_a}{2} sin2(\phi - \phi_0)\right] k^2 + \frac{1}{4}\lambda^3 C_s k^4$$

defocus ax. astigmatism spherical aberration

k = spatial frequency vector; *k* = length of this vector

Point-Spread Function =
Response of the optical instrument to a point object

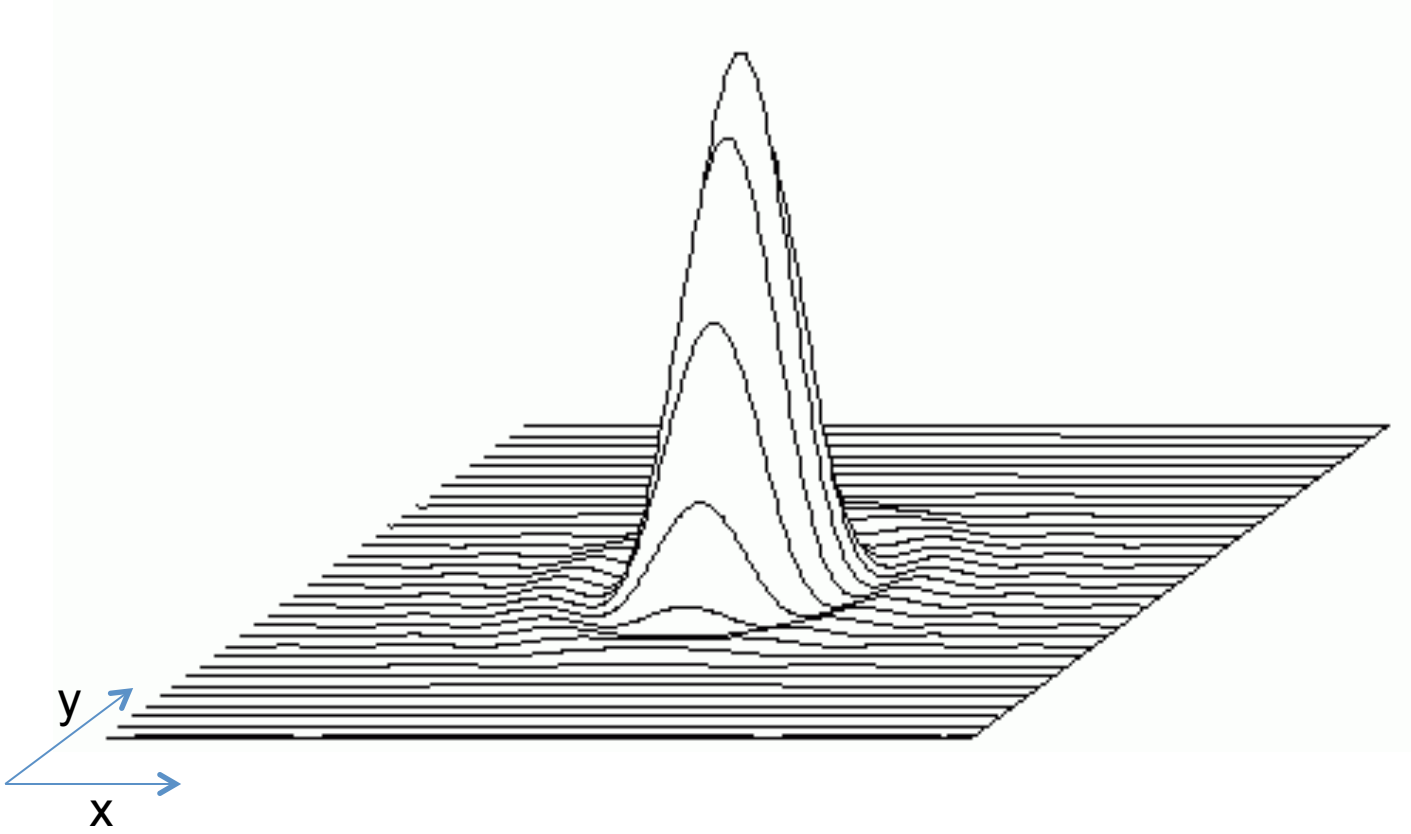

The point spread function has finite width, and is centered at the location that the point would have in the image formed by an ideal instrument.

Contrast transfer function

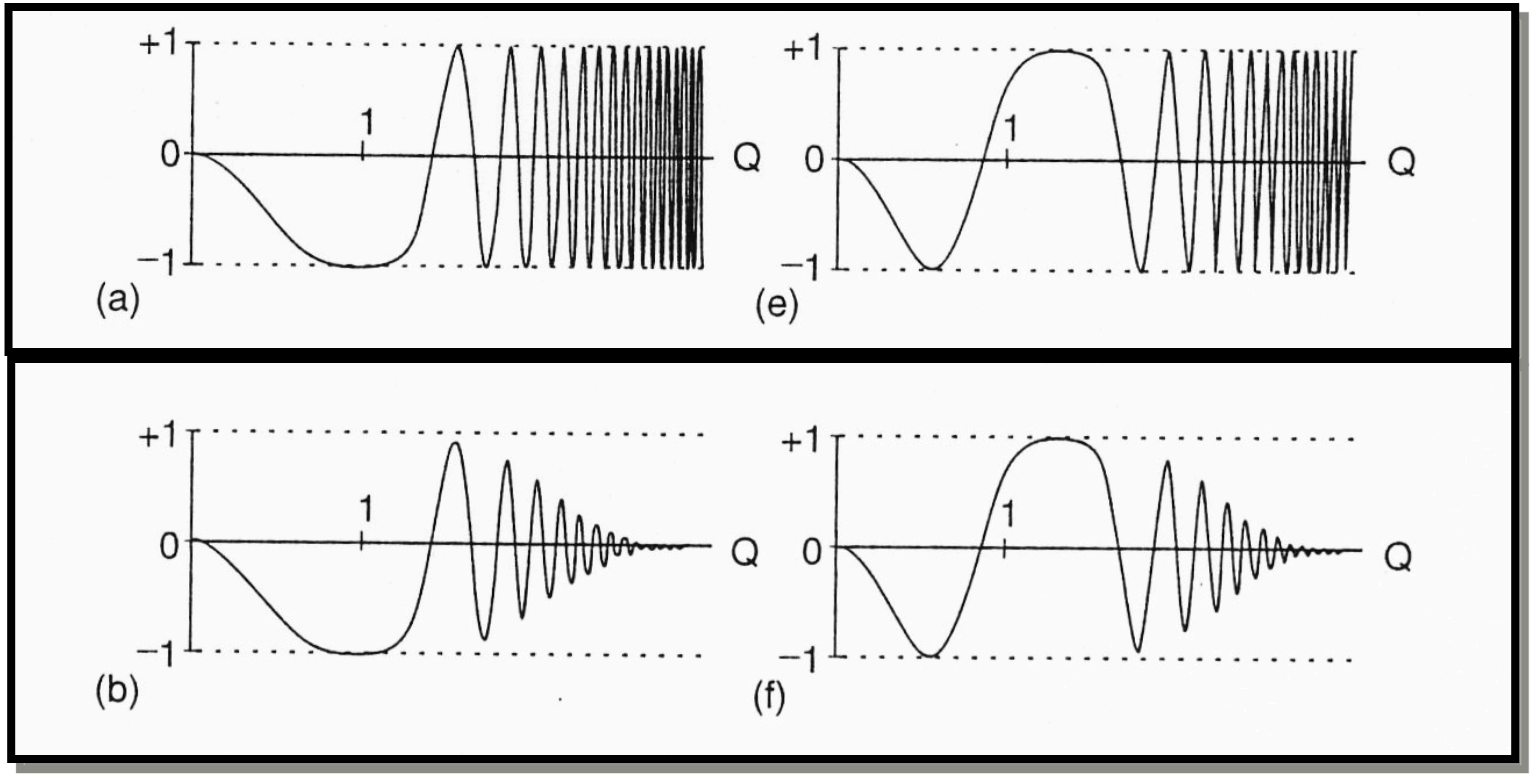

Coherent

Partially coherent

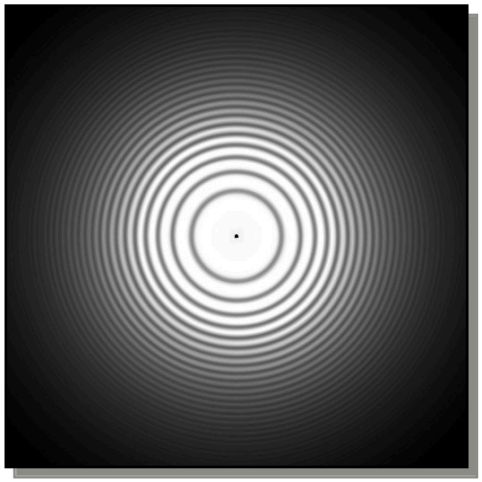

Power spectrum
“Thon rings”

after Fritz Thon, a pioneer in optical diffraction analysis

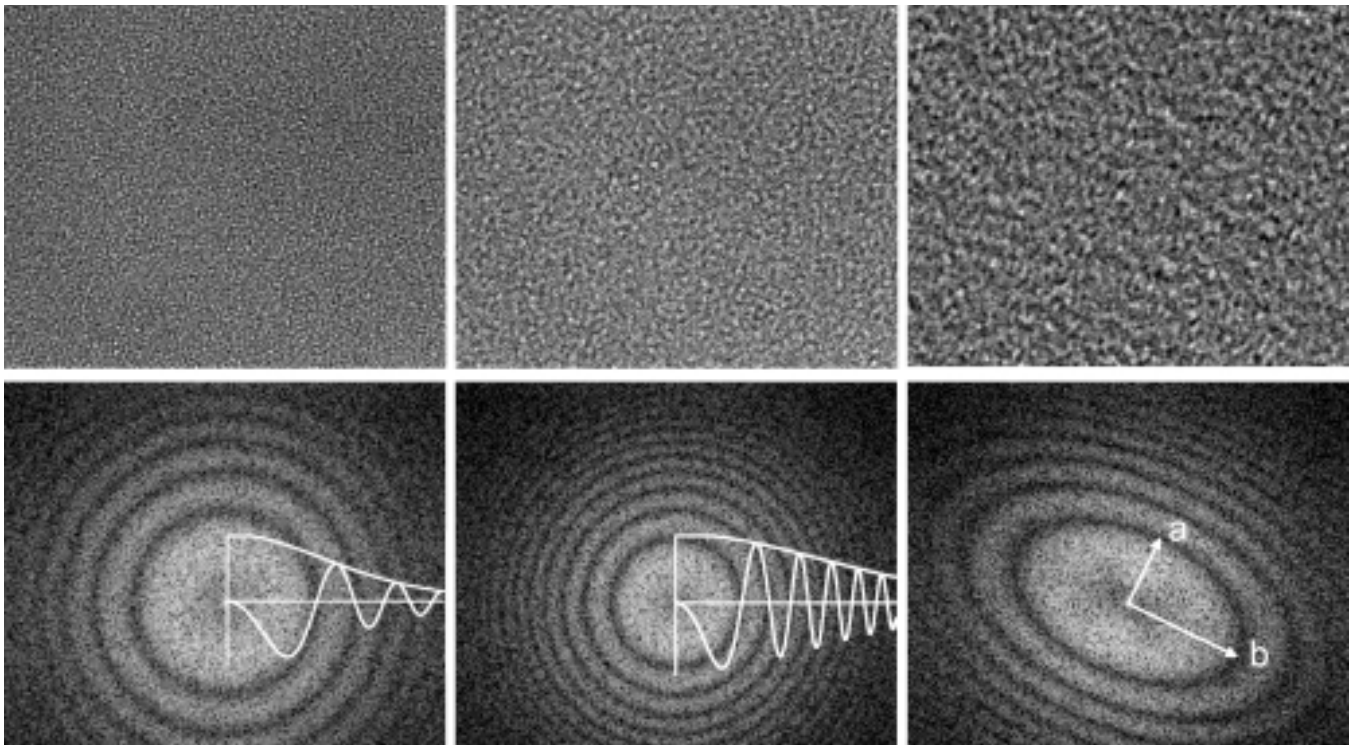

What Thon rings show:

1) how far the information transmitted ranges in Fourier space
2) whether the lens is astigmatic (CTF depends on angle in the plane)

Why do we see rings? Because for an amorphous object, such as carbon, the amplitudes of Fourier components are roughly the same throughout Fourier space. Without CTF, we would see a uniform (white) disk up to the radius that corresponds to the resolution limit.
Instead we see concentric white rings separated by black lines (zero transitions).

Defocus (and hence the CTF) is affected by the particle's z-position within the ice layer.

Ideally, defocus should be measured for each particle separately, but the signal is often not stron enough.

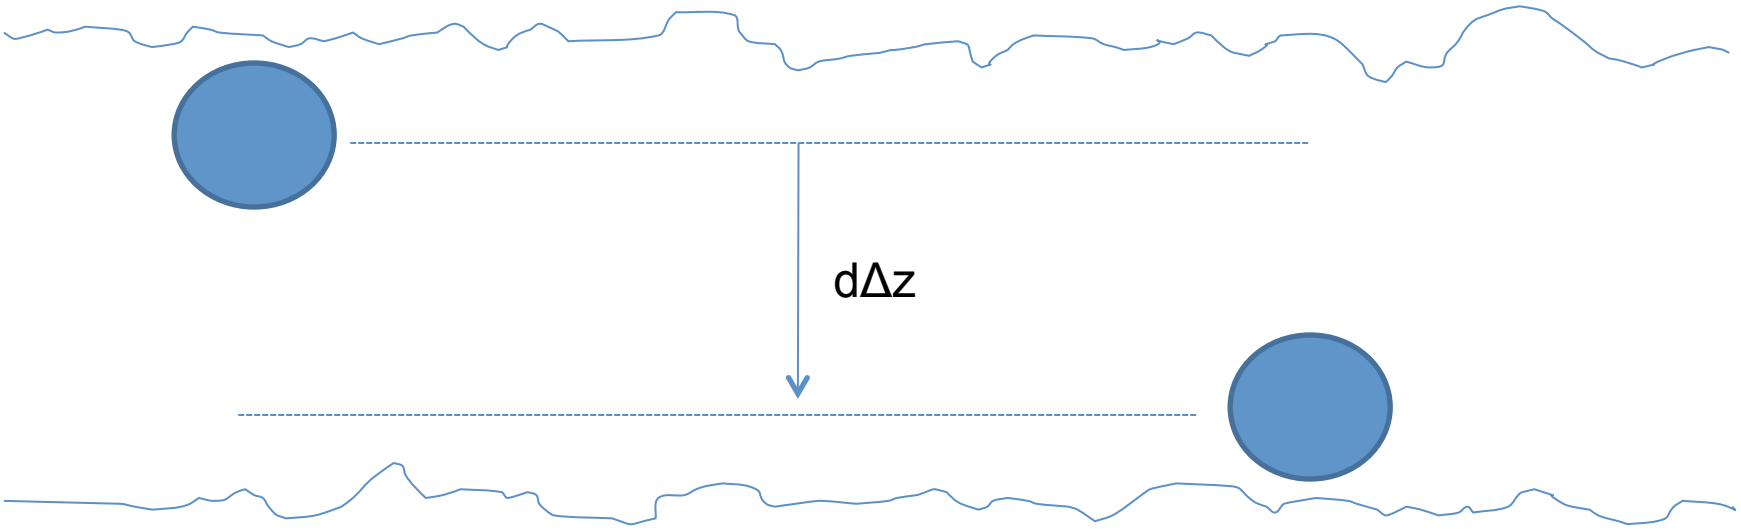

Ice layer large compared with particle diameter

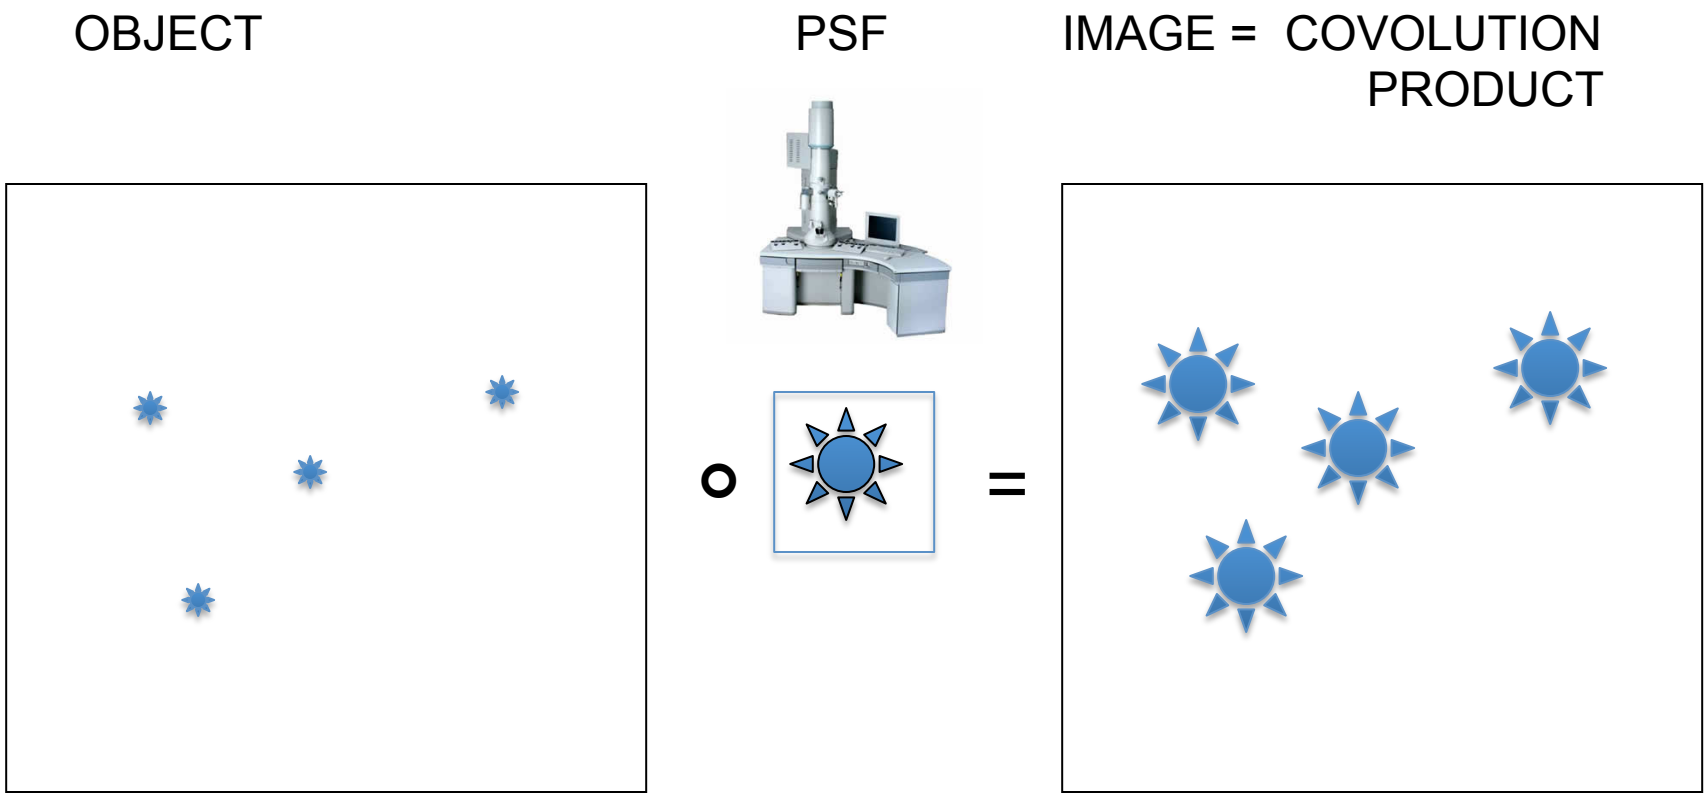

An object consisting of points, convoluted with the point spread function of the optical instrument, results in an image in which each point is replaced by the PSF..

Convolution theorem

Convolution theorem

$o(\mathbf{r})$ = "2D object" or "2D projection of a 3D object"
$s(\mathbf{r})$ = "signal resulting from EM imaging"
$h(\mathbf{r})$ = "point-spread function"

$s(\mathbf{r}) = o(\mathbf{r}) \circ h(\mathbf{r})$ = "convolution product of $o(\mathbf{r})$ with $h(\mathbf{r})$"

stands for

$$s(x) = \iint o(x', y')h(x - x', y - y')dx'dy'$$

or

$$s(\mathbf{r}) = \int o(\mathbf{r}')h(\mathbf{r} - \mathbf{r}')d\mathbf{r}'$$

○	Convolution
☼	Correlation
x	Scalar product

Special case: $o(\mathbf{r}) \circ \delta(r) = s(\mathbf{r})$ convolution with a delta function is an identity operation

Convolution Theorem says: $S(\mathbf{k}) = \mathcal{F}\{s(\mathbf{r})\} = O(\mathbf{k})H(\mathbf{k})$, where $S(\mathbf{k}) = \mathcal{F}\{o(\mathbf{r})\}$, $O(\mathbf{k}) = \mathcal{F}\{o(\mathbf{r})\}$, *and* $H(\mathbf{k}) = \mathcal{F}\{h(\mathbf{r})\}$

"Correlation theorem"

$s(\mathbf{r})$ = "signal resulting from EM imaging" – two versions: $s_1(\mathbf{r})$ and $s_2(\mathbf{r})$

CCF($\mathbf{r}$) = $s_1(\mathbf{r})$ ☼ $s_2(\mathbf{r})$ = "correlation function of $s_1(\mathbf{r})$ with $s_2(\mathbf{r})$"

stands for

$$CCF(x,y) = \iint s_1(x',y')s_2(x+x',y+y')dx'dy'$$

or

$$CCF(\mathbf{r}) = \int s(\mathbf{r}')s(\mathbf{r}+\mathbf{r}')d\mathbf{r}'$$

○ Convolution

☼ Correlation

x Scalar product

Special case: $s_1(\mathbf{r})$ ☼ $s_1(\mathbf{r})$ <u>autocorrelation function</u>

<u>Correlation Theorem says</u>: $\Phi(\mathbf{k}) = \mathcal{F}\{\text{CCF}^{\circledR}\} = S_1(\mathbf{k})S_2^*(\mathbf{k})$ where $S_1(\mathbf{k}) = \mathcal{F}\{s_1(\mathbf{r})\}$, $S_2(\mathbf{k}) = \mathcal{F}\{s_2(\mathbf{r})\}$

Translational alignment using the CCF – a practical example

Padding is needed when CCF is computed via Fourier methods:

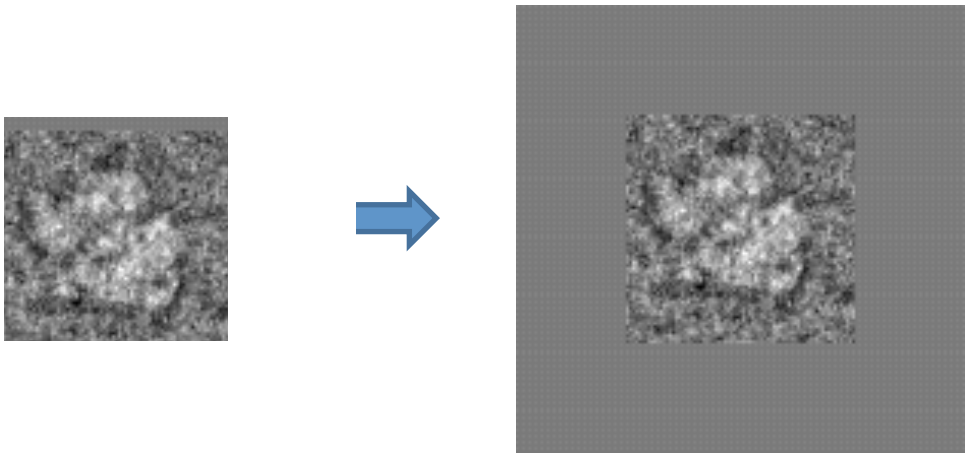

Need for padding follows from the discrete Fourier representation:

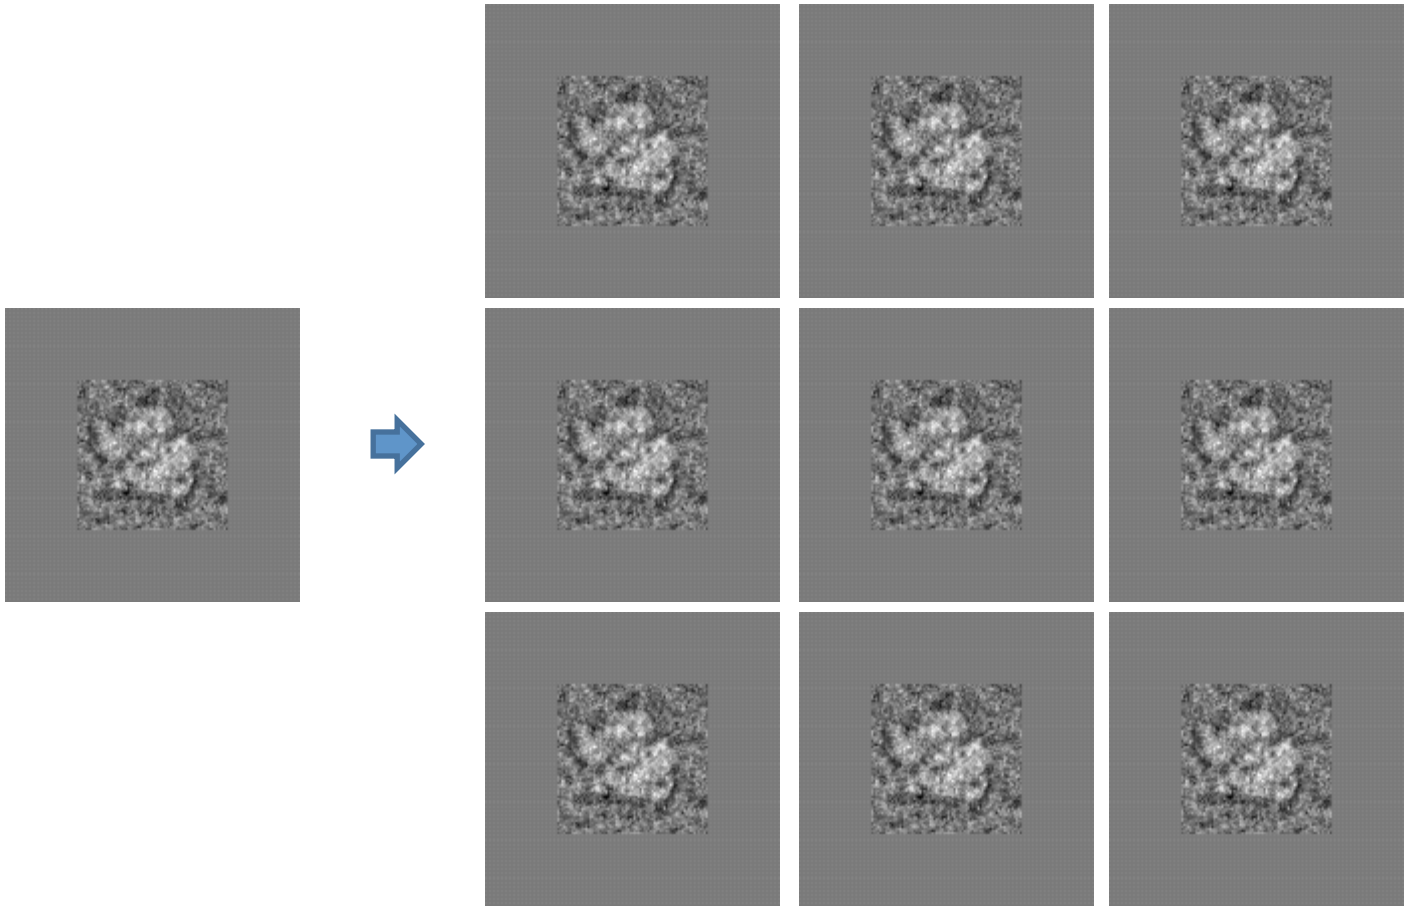

Nonpadded image would get superimposed on a copy of itself

Translational alignment using the CCF – a practical example

The peak indicates postion of perfect alignment of two images of the same molecule

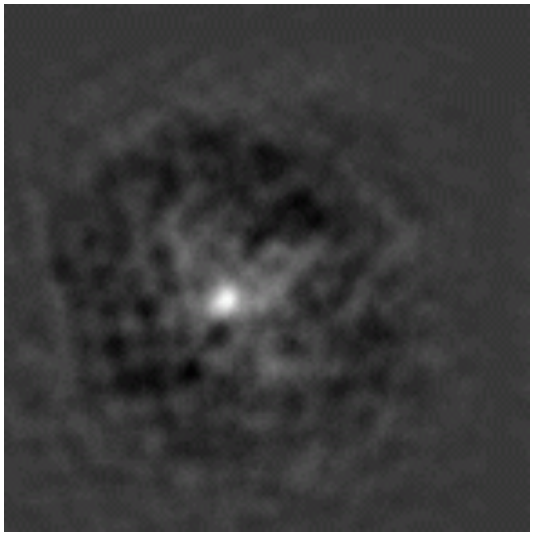

CCF

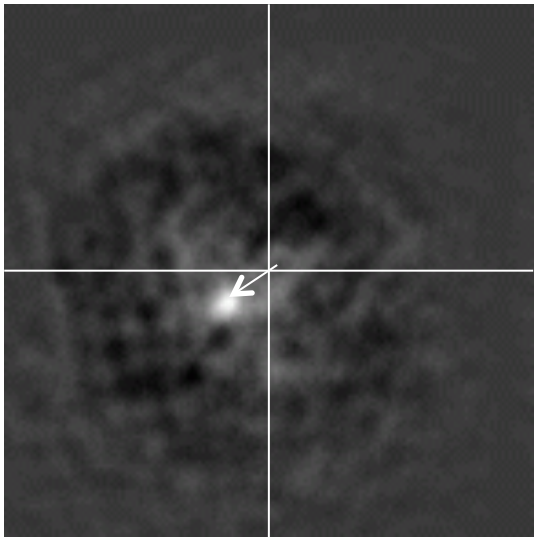

shift of peak from origin
indicates relative shift of images

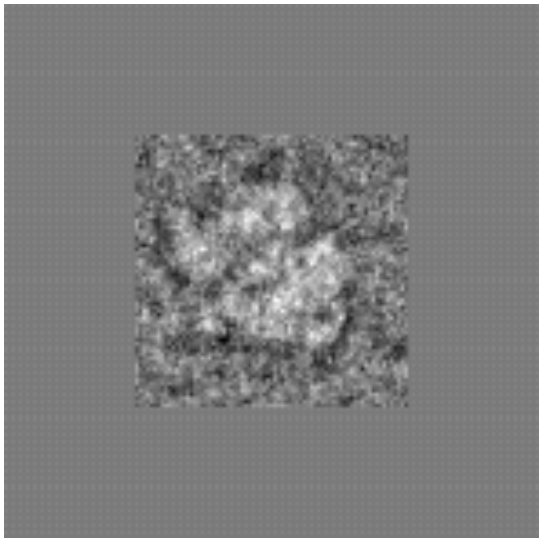

one of the images,
padded

Auto-correlation function = cross-correlation of an image with itself

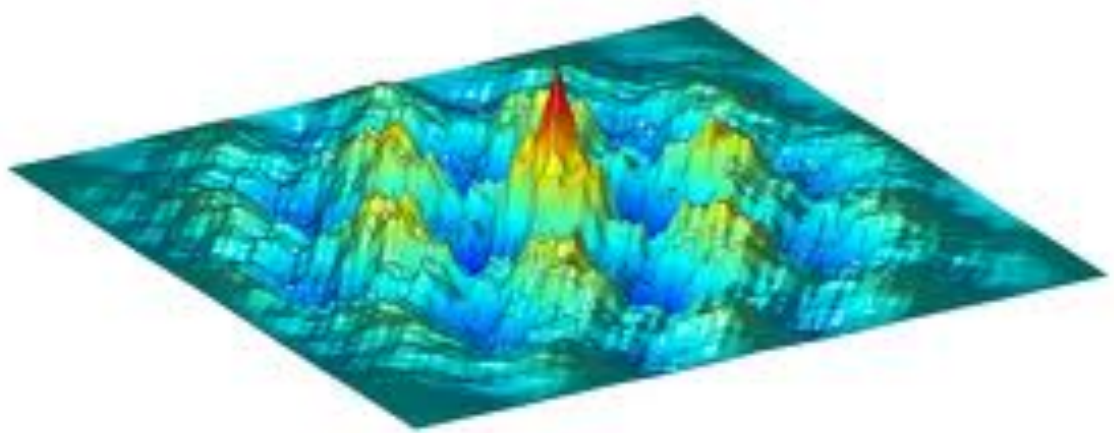

(we need to know this since it's used in the next slide)

CCF of EM Images with different CTFs

CCF of two images s_1 (**r**), s_2 (**r**) of the same object signal o(**r**) with different CTFs, H_1(**k**) and H_2(**k**)

(using *correlation* and *convolution theorems*).

s_1(**r**) = o(**r**) ○ h_1(**r**); s_2 = o(**r+Δr**) ○ h_2(**r**)

CCF(**r**) = s_1(**r**) ☼ s_2(**r**)

S_1(**k**) x S_2*(**k**) = O(**k**) x H_1 (**k**)] x {O(**k**) x exp[2πi**kΔr**] x H_2(**k**)}*

= O(**k**) x O*(**k**) x exp[-2πi**kΔr**] x H_1(**k**) x H_2*(**k**)

○ Convolution

☼ Correlation

x Scalar product

Now back in real space:

CCF(**r**) = [o(**r**) ☼ o(**r**)] ○ [h_1 (**r**) ☼ h_2 (**r**)] ○ δ(r-Δr)

Shifted auto-correlation function of signal CCF of point spread functions

Scherzer focus (Δ =1)

Δ=10

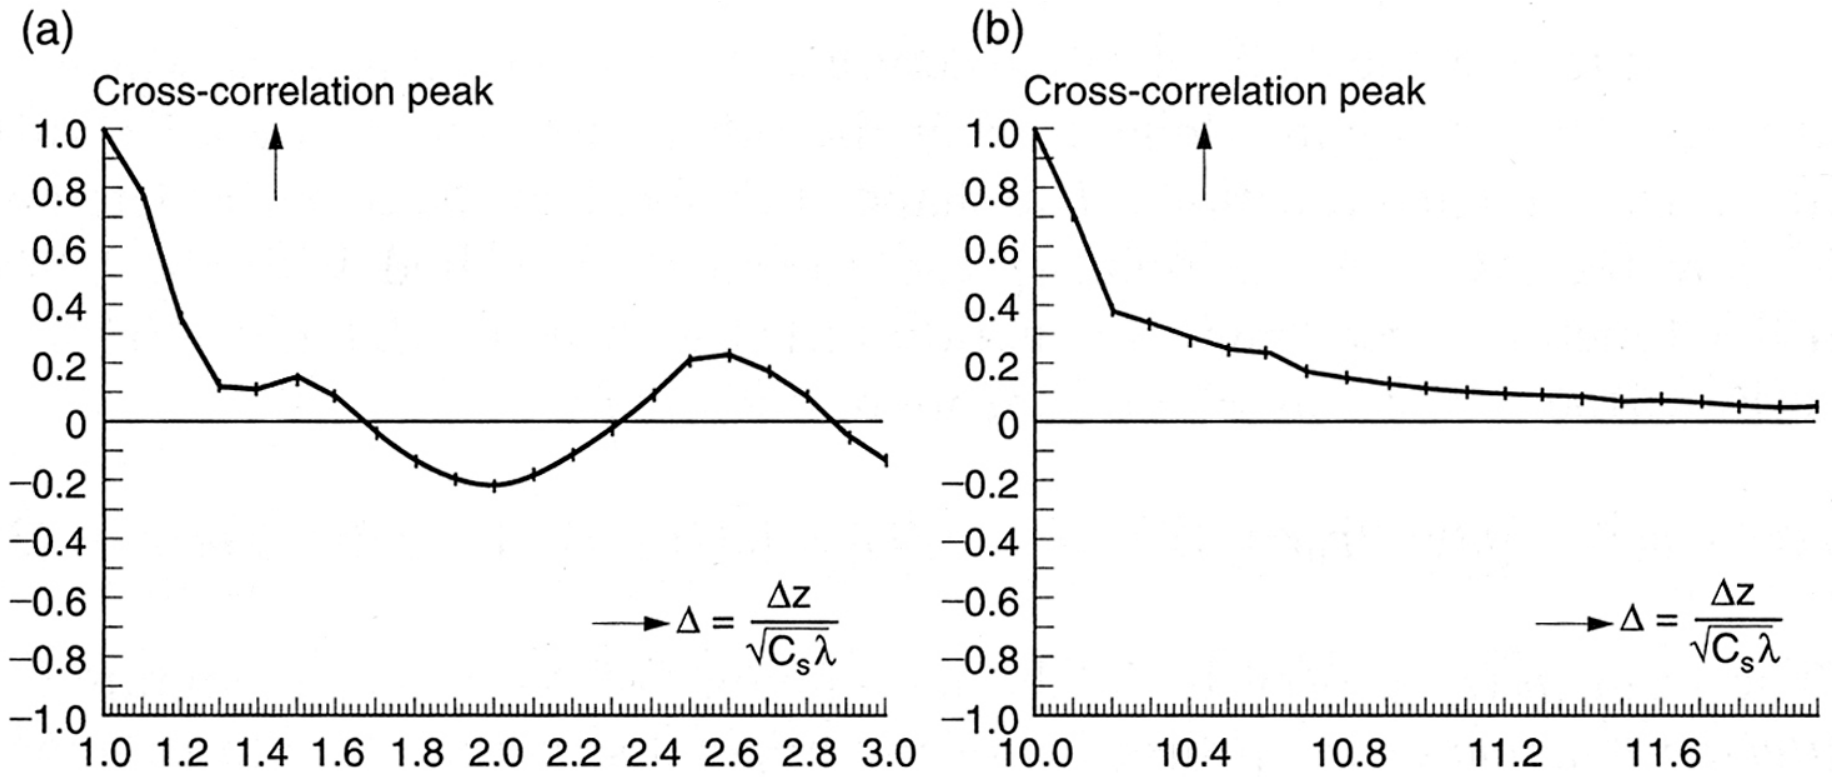

Value of normalized CCF peak as a function of the difference in focus

Tools: CCFs of micrographs of the same specimen, as a function of $\delta\Delta z$

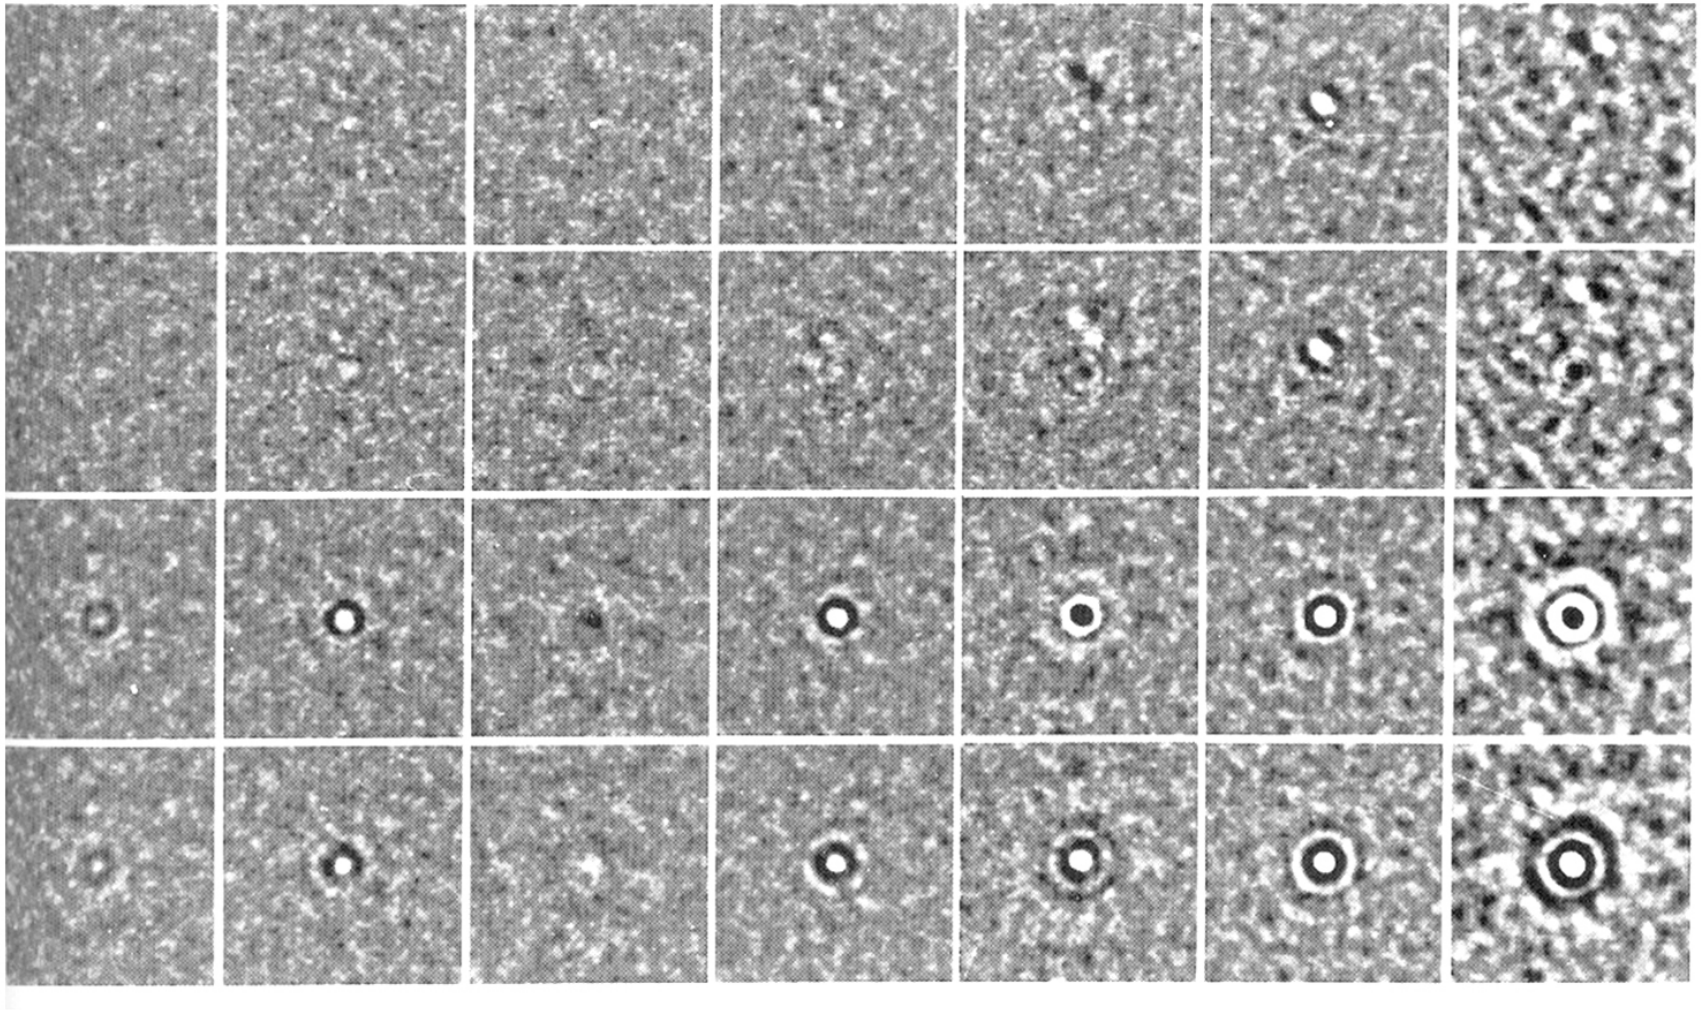

The cross-correlation function has a peak standing out from a noisy background

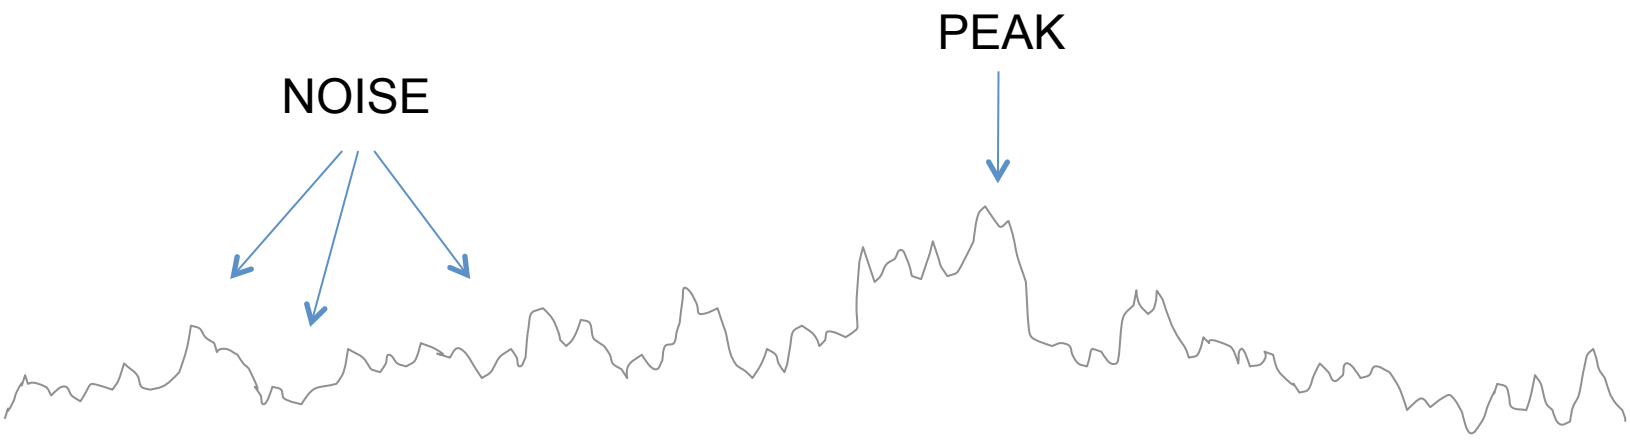

The SNR in the CCF determines whether we will be successful in finding the correct peak and its x,y coordinates. What factors affect the size of the SNR?

Criterion for detection of CCF peak: feasibility of alignment

Because of the low signal-to-noise ratio in the images, there exists a critical threshold for the feasibility of alignment of two raw images of a molecule.

The critical parameters are:

p_{crit} -- maximum exposure [electrons/unit area] the molecule can tolerate

D -- particle size

c – contrast

d – resolution (distance)

Particle size D should satisfy $$D \geq \frac{3}{c^2 dp_{crit}}$$

Saxton and Frank (Ultramicroscopy 1976)

Rule of thumb: alignment easy when mol. mass > 400 kD BUT: needs to be revised in view of recent developments

Henderson (Quart. Rev. Biophys. 1995): number of molecules of a given size required to reach 3A resolution, based on scattering data for electrons, X-rays, and neutrons.

2D alignment strategies:

reference-based
versus reference-free

A variant of the reference-based method updates the reference as it goes along, and becomes thus less dependent on the initial choice of reference.

How to combine rotation with translation:
(1) Iterative (alternating translational with rotational search)
(2) Use of invariants

Multi-Step Reference-Based Alignment

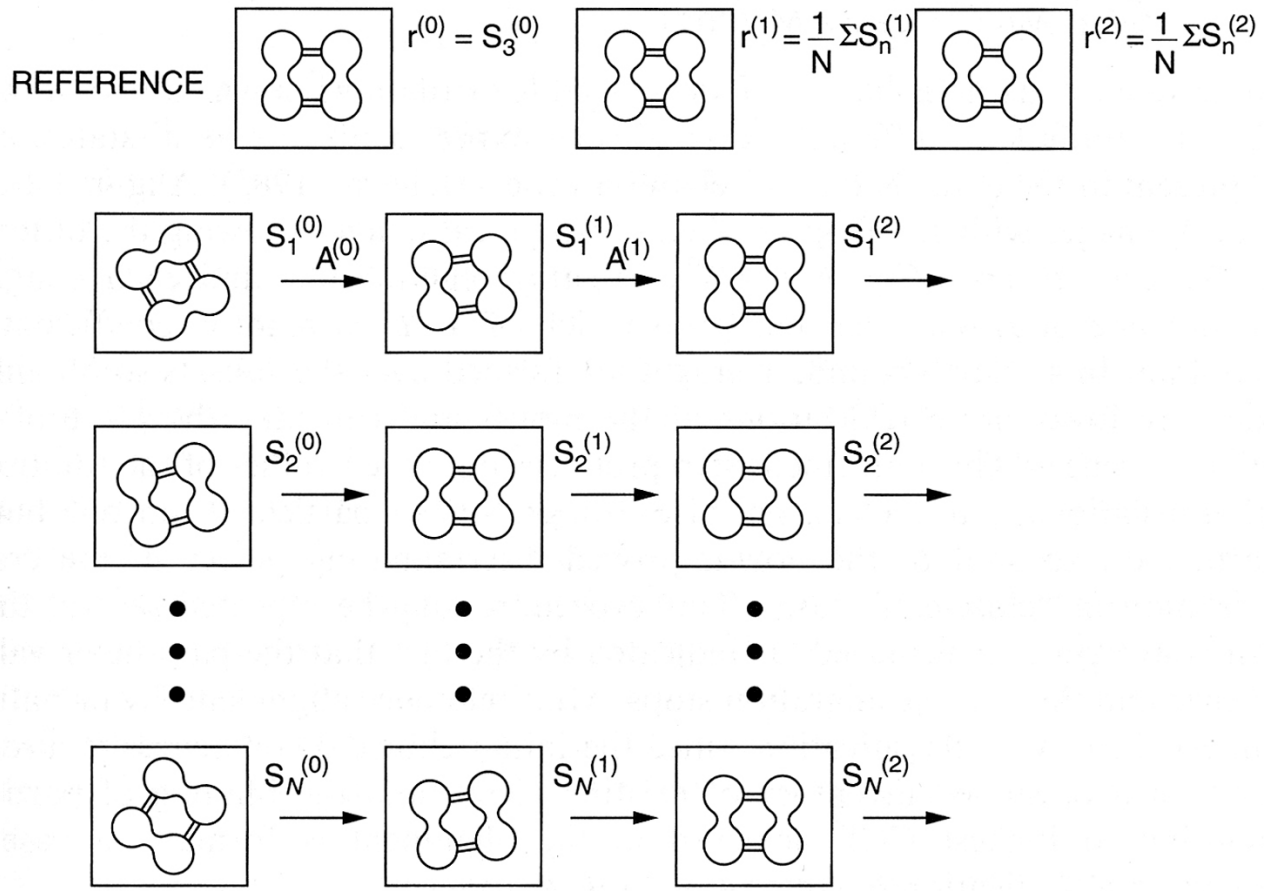

Noise-free molecule and its ACF

The molecule in arbitrary positions

ACFs of molecules above

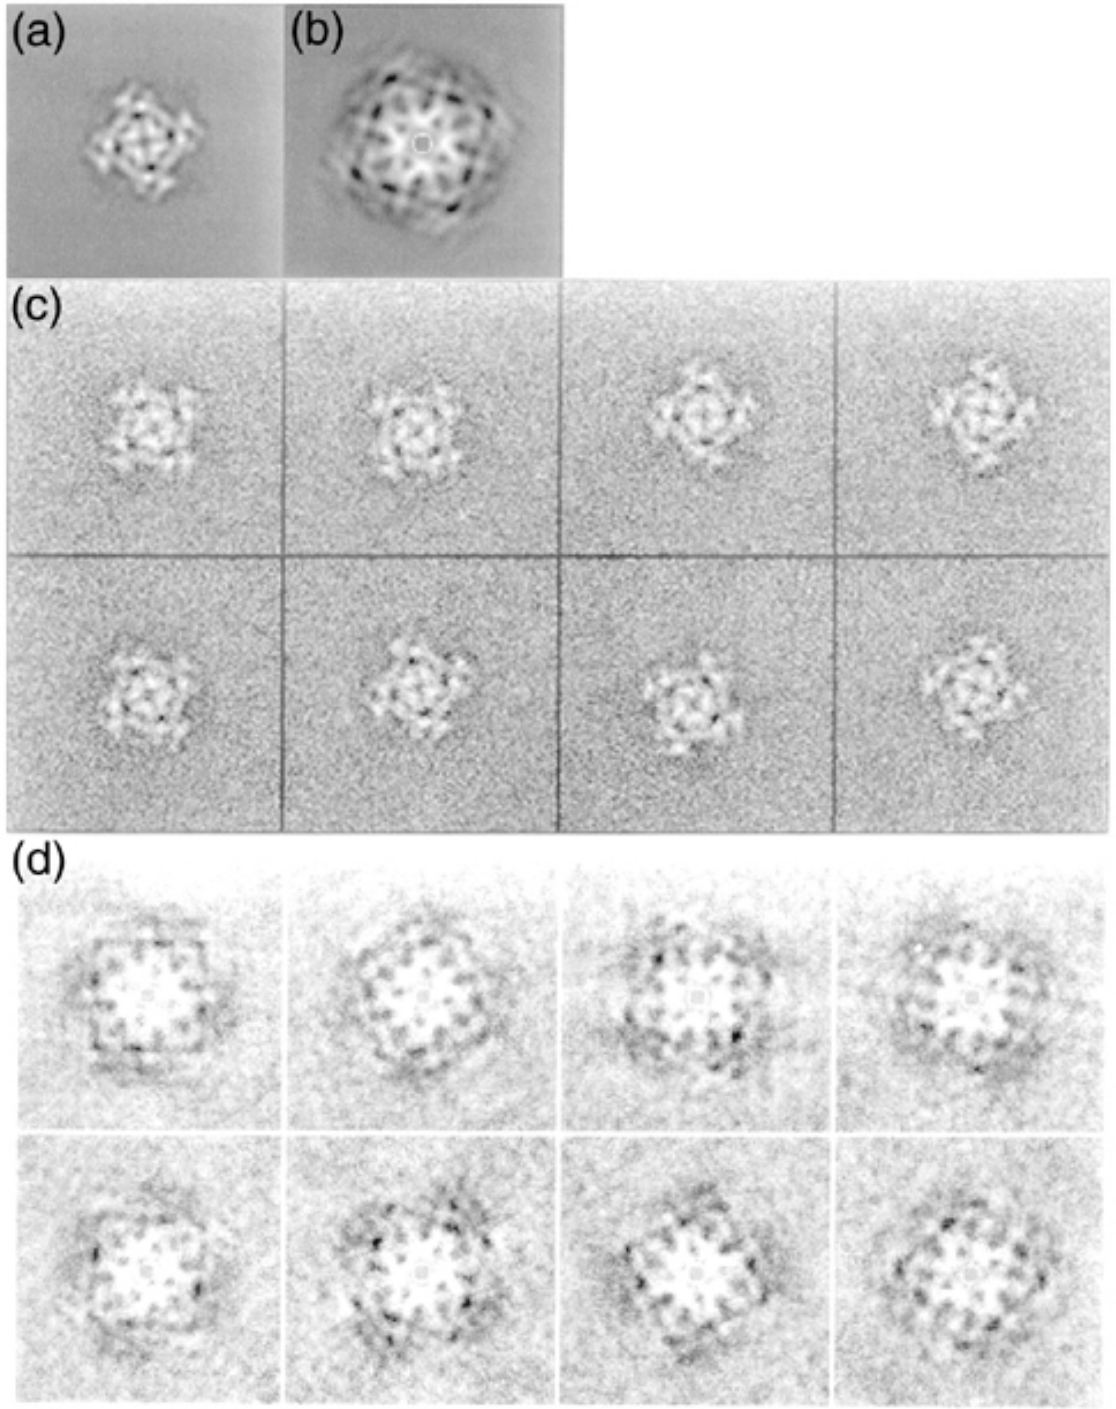

ACF – based alignment method

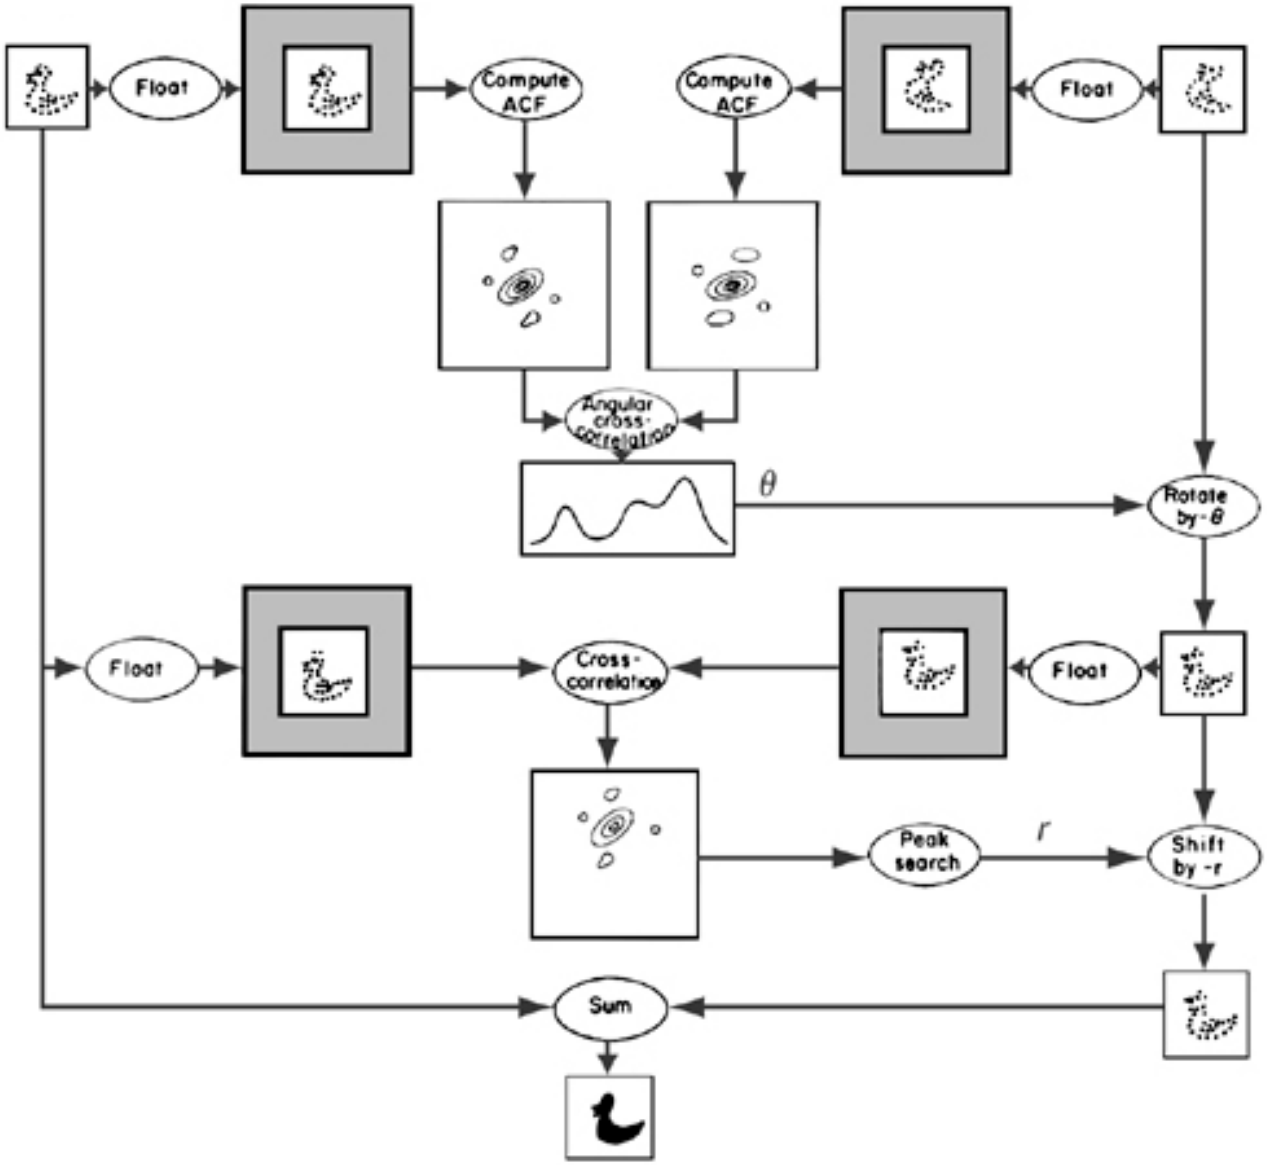

Variance map

- The variance map is a "byproduct" of the averaging. It can be used to find the regions where the images, on average, differ maximally.

- It is also the yardstick that helps determine whether or not a density in a difference map is significant.

$$v_{(N)}(\mathbf{r}_j) = \frac{1}{(N-1)} \sum_{i=1}^{N} [p_i(\mathbf{r}_j) - \bar{p}_{(N)}(\mathbf{r}_j)]^2$$

particle image average image

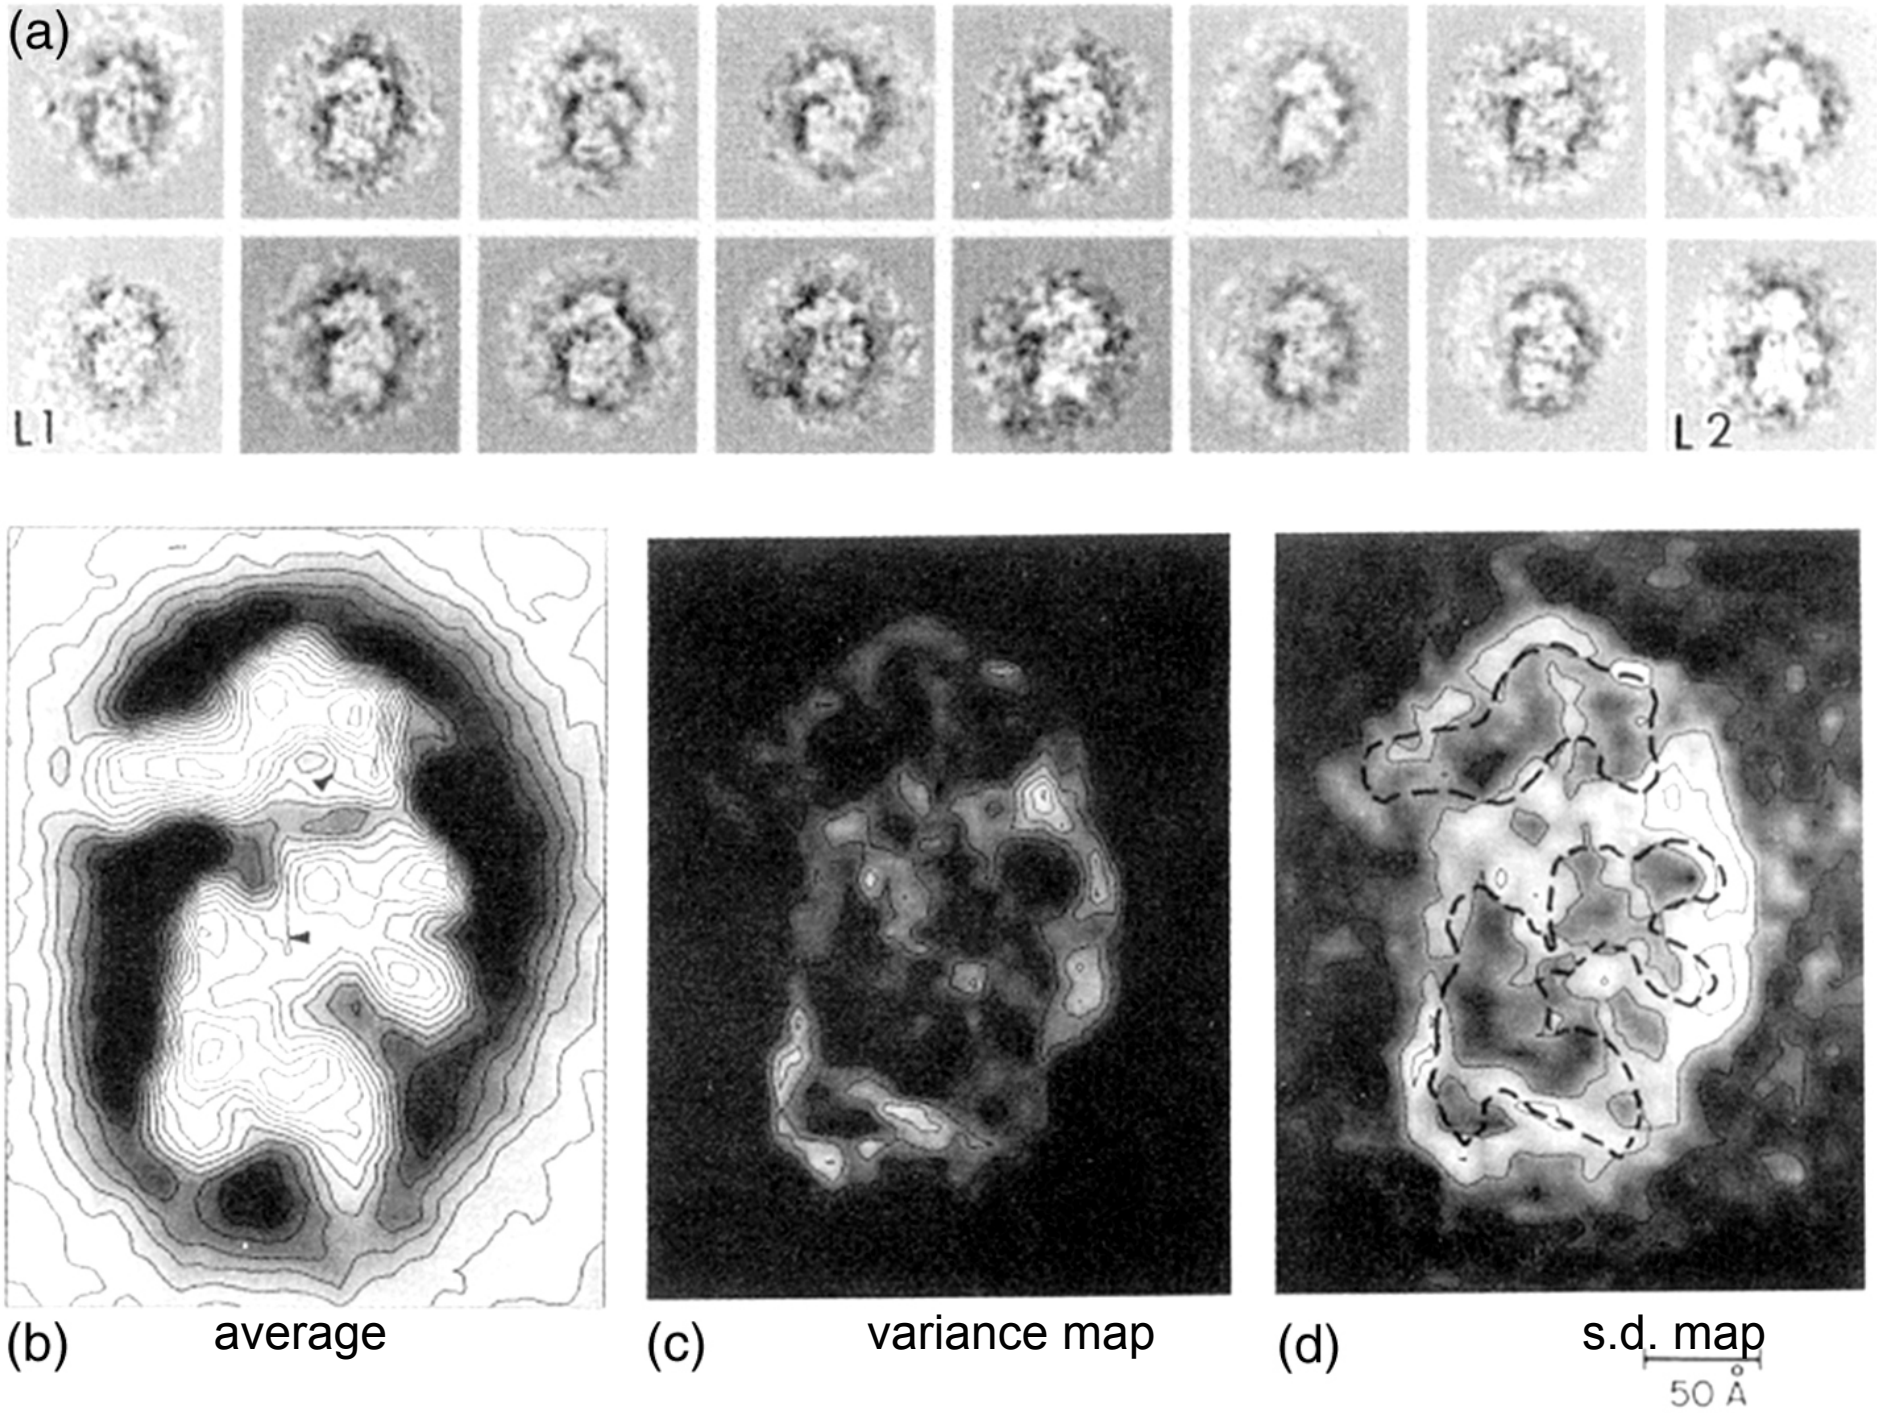
(a)
L1
L2
(b) average
(c) variance map
(d) s.d. map
50 Å

Alignment of frames in movie mode of Direct Electron Detection cameras

Most DED cameras allow data to be collected in multiple frames. This makes it possible to correct for drift, even on the level of single particles.

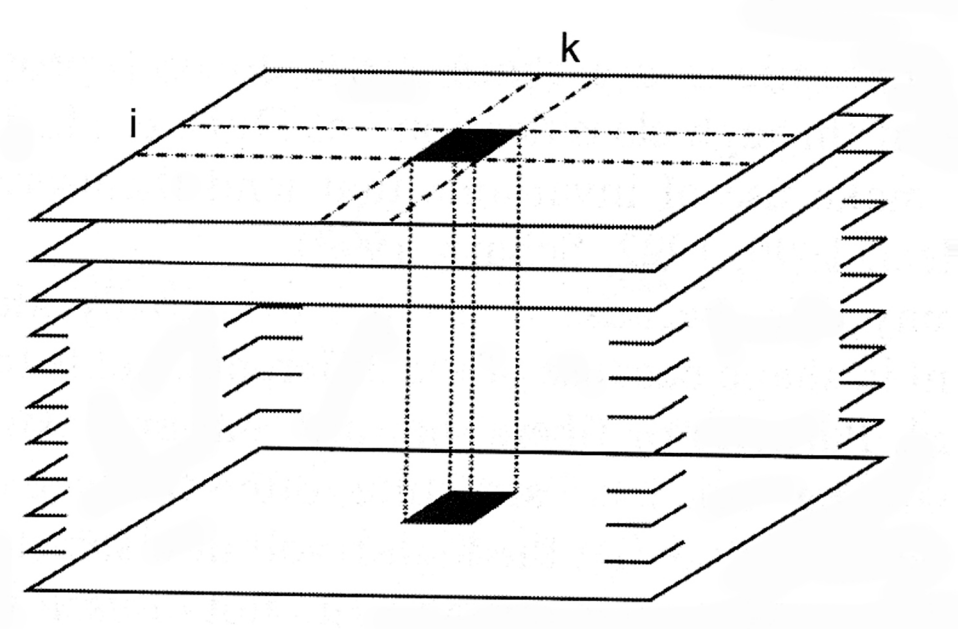

γ-secretase

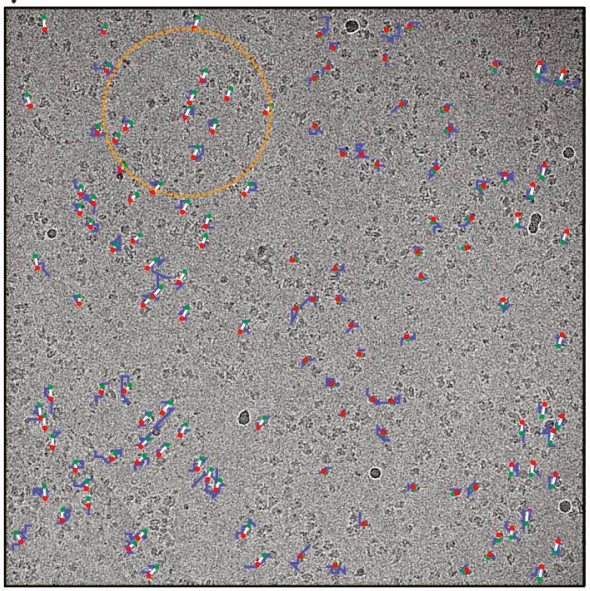

β-galactosidase

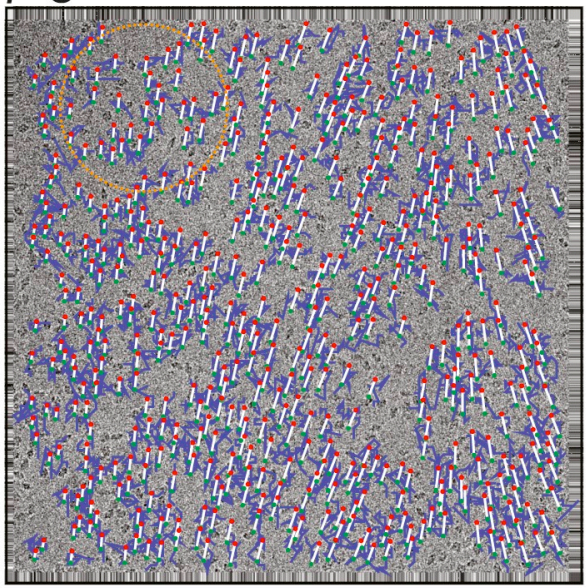

complex-I

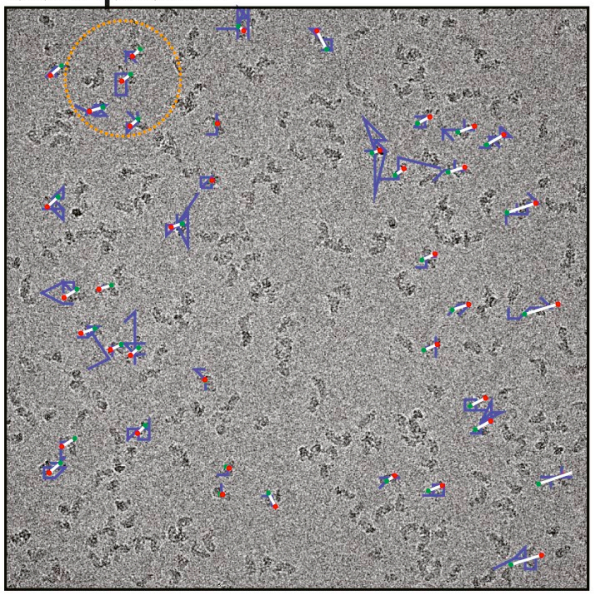

mitoribosome

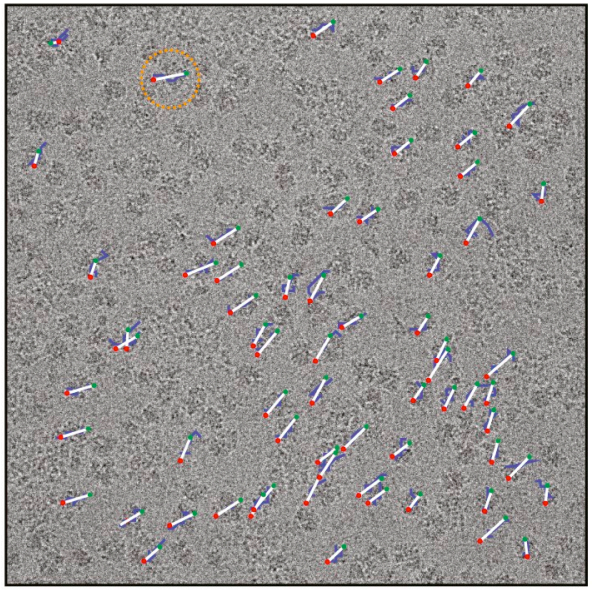

(vectors of motion are exaggerated)

S. Scheres, eLife 2014

Resolution criterion: Images of two points, as function of their separation

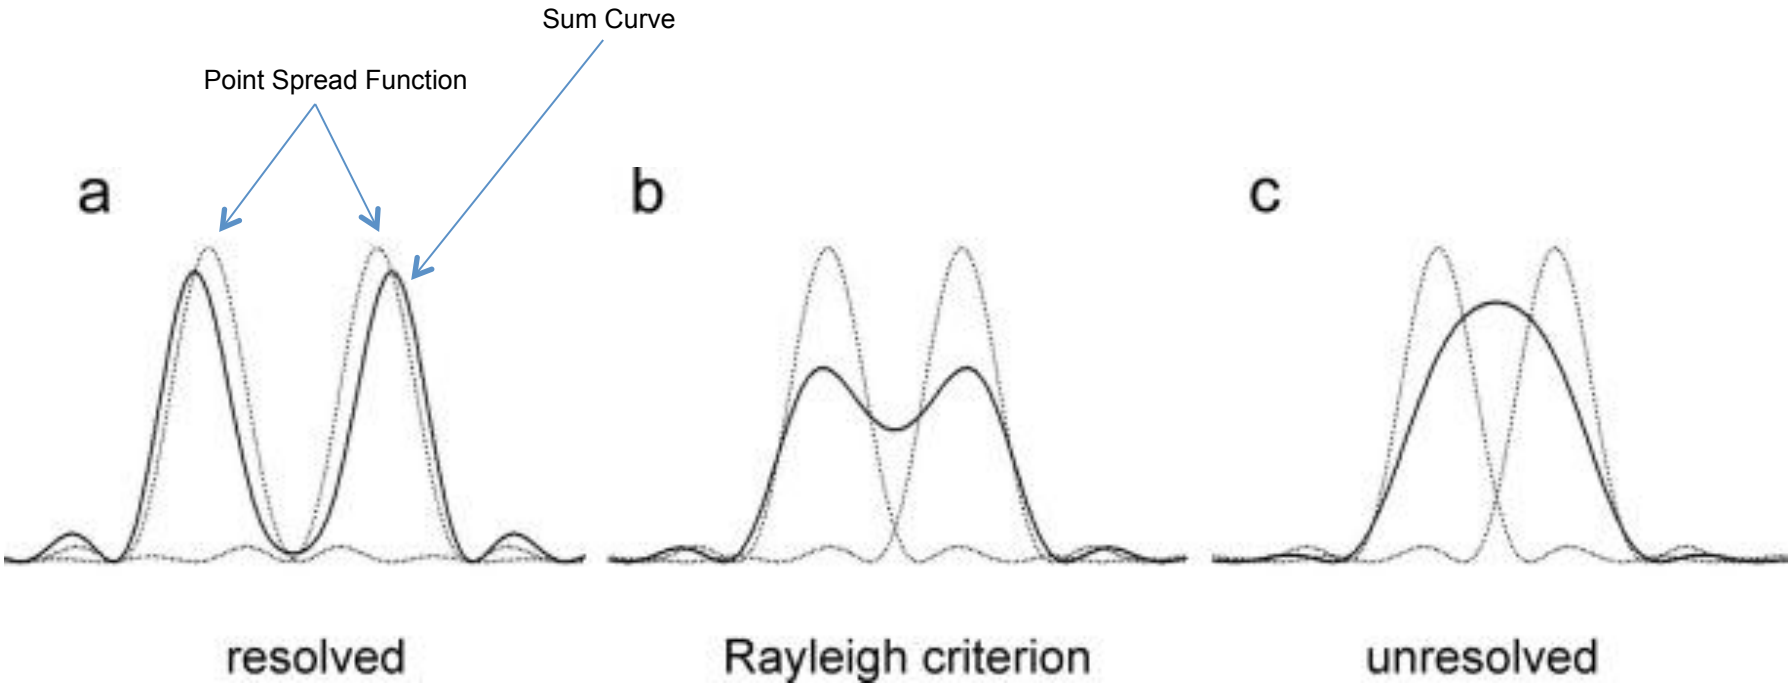

resolved	Rayleigh criterion	unresolved
Peaks are distinguishable	Peaks are just resolved	Peaks blend into a single one
when distance is > $d_{rayleigh}$	when distance is = $d_{rayleigh}$	when distance < $d_{rayleigh}$

E. V. Orlova and H. R. Saibil, in *Chemical Reviews* 2011

Resolution definition, determination in Fourier space

- Resolution is a reciprocal quantity, measured in Fourier space
- Defined as the spatial frequency [1/Å] up to which information is reproducible, by some measure of reproducibility
- Decomposition of information, by Fourier rings
- Randomly picked halfsets (e.g., odd- vs. even-numbered images)
- Compare averages [reconstructions] from halfsets over rings (shells) in Fourier space

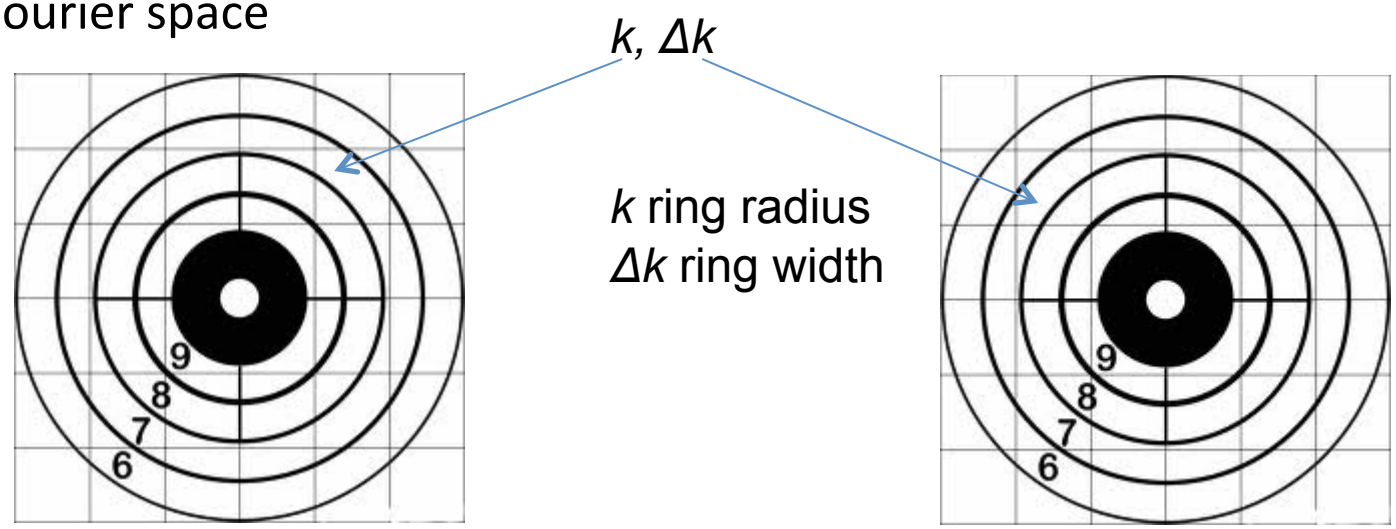

Resolution measures & criteria: Fourier ring/shell correlation

$F_1(\mathbf{k}), F_2(\mathbf{k})$ Fourier transforms of halfset averages (or halfset reconstructions)

$$FSC(k,\Delta k) = \frac{\mathrm{Re}\{\sum_{(k,\Delta k)} F_1(\mathbf{k})F_2(\mathbf{k})\}}{\{\sum_{(k,\Delta k)} |F_1(\mathbf{k})|^2 |F_2(\mathbf{k})|^2\}^{1/2}}$$

k = *spatial frequency vector*

$k = |\mathbf{k}|$ *abs.* size of *spatial frequency*

Δk = *ring width* or (in 3D) *shell thickness*

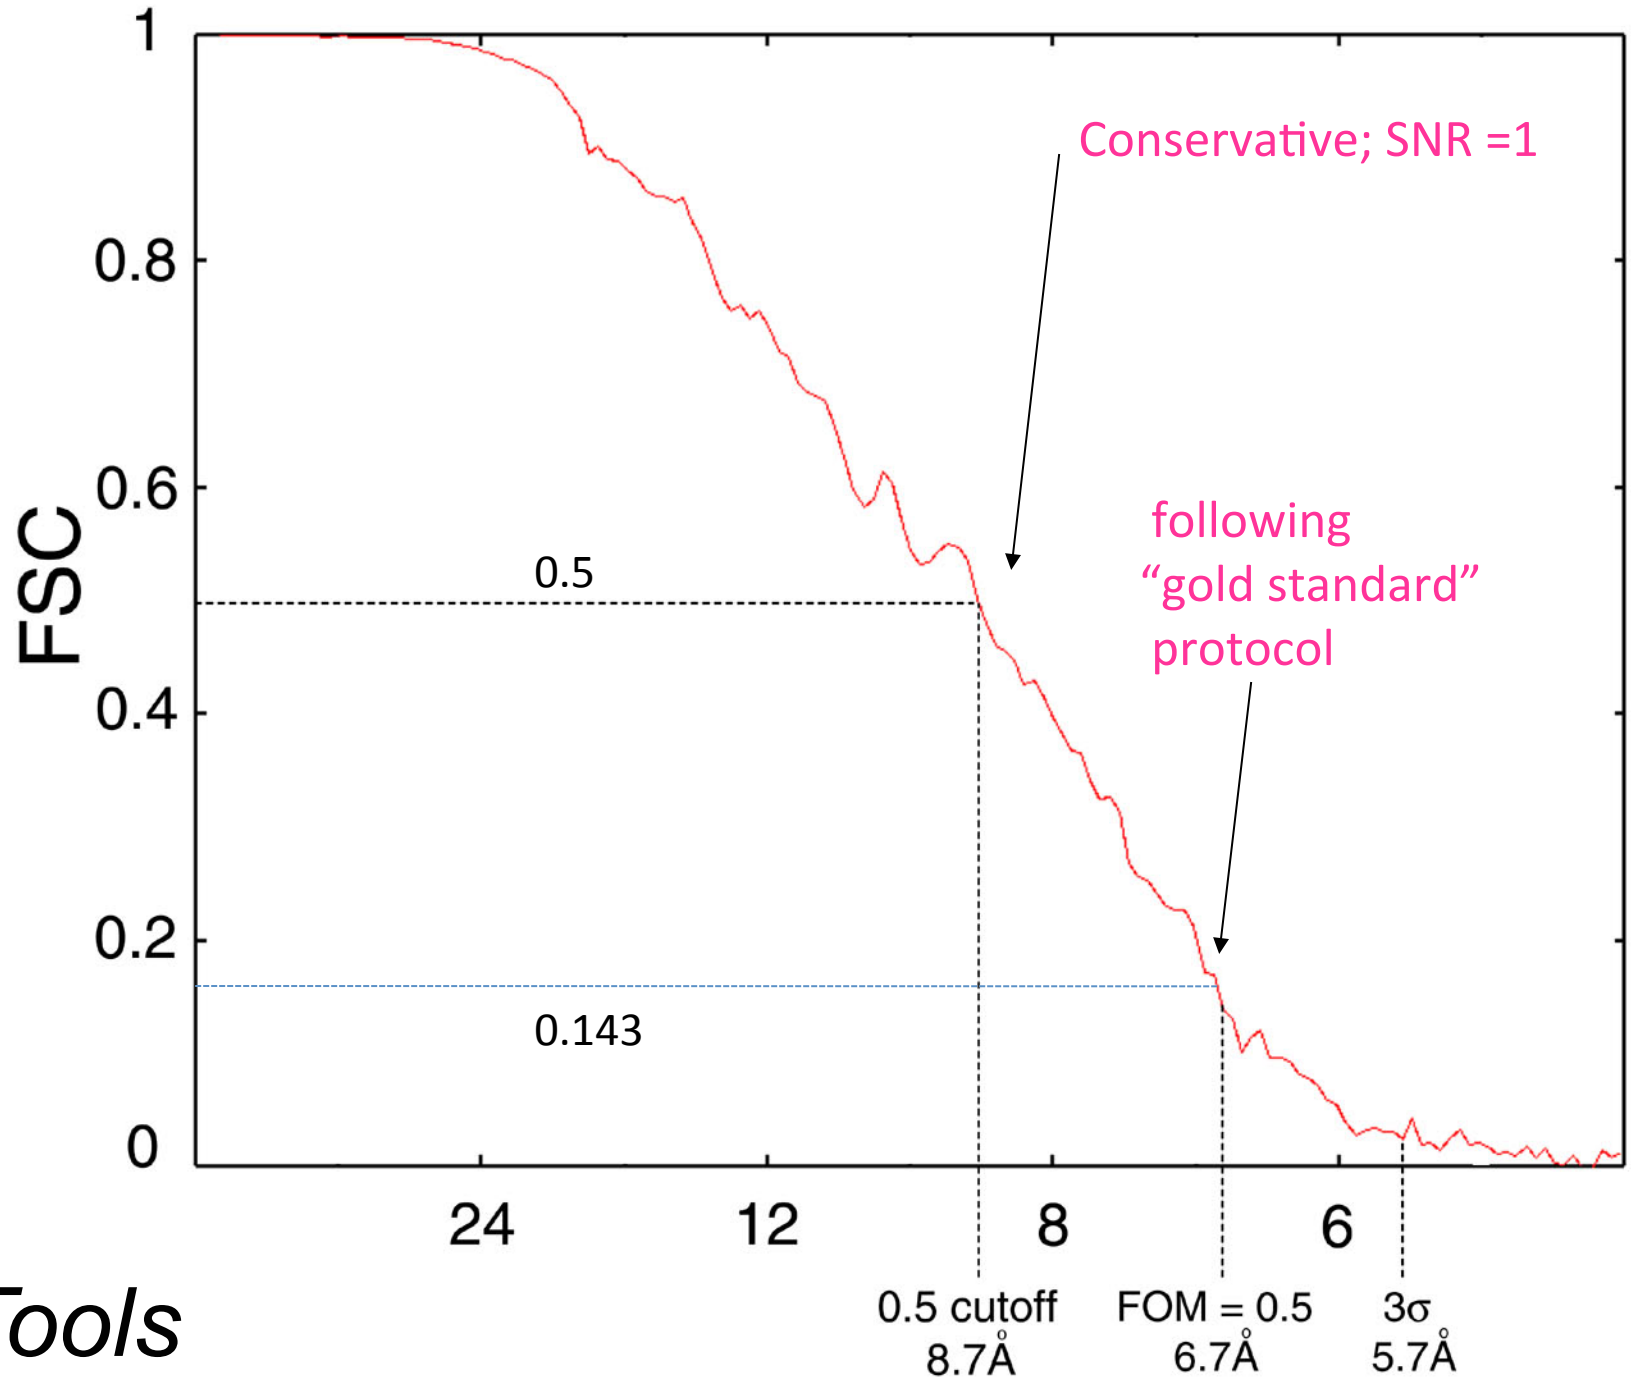
1
0.8
0.6
0.4
0.2
0
FSC
0.5
0.143
Conservative; SNR =1
following
“gold standard”
protocol
24
12
8
6
0.5 cutoff
8.7Å
FOM = 0.5
6.7Å
3σ
5.7Å

Tools

Multivariate Data Analysis and Classification

- Images often need to be sorted into classes
- Heterogeneity is due to (1) different viewing angle and (2) different conformations of the molecules
- Sorting them visually only works in the simplest cases
- Multivariate analysis reduces the dimensionality of the classification problem

Classification in 2D

RATIONALE:

Inventory of existing views

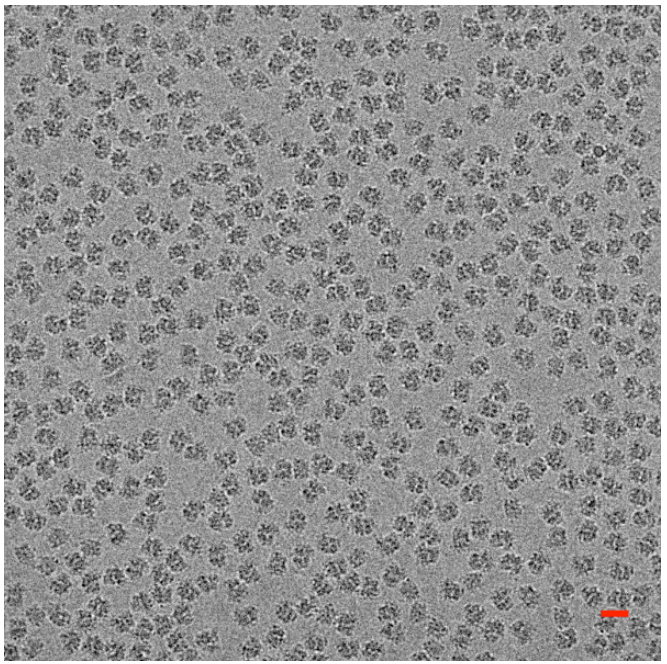

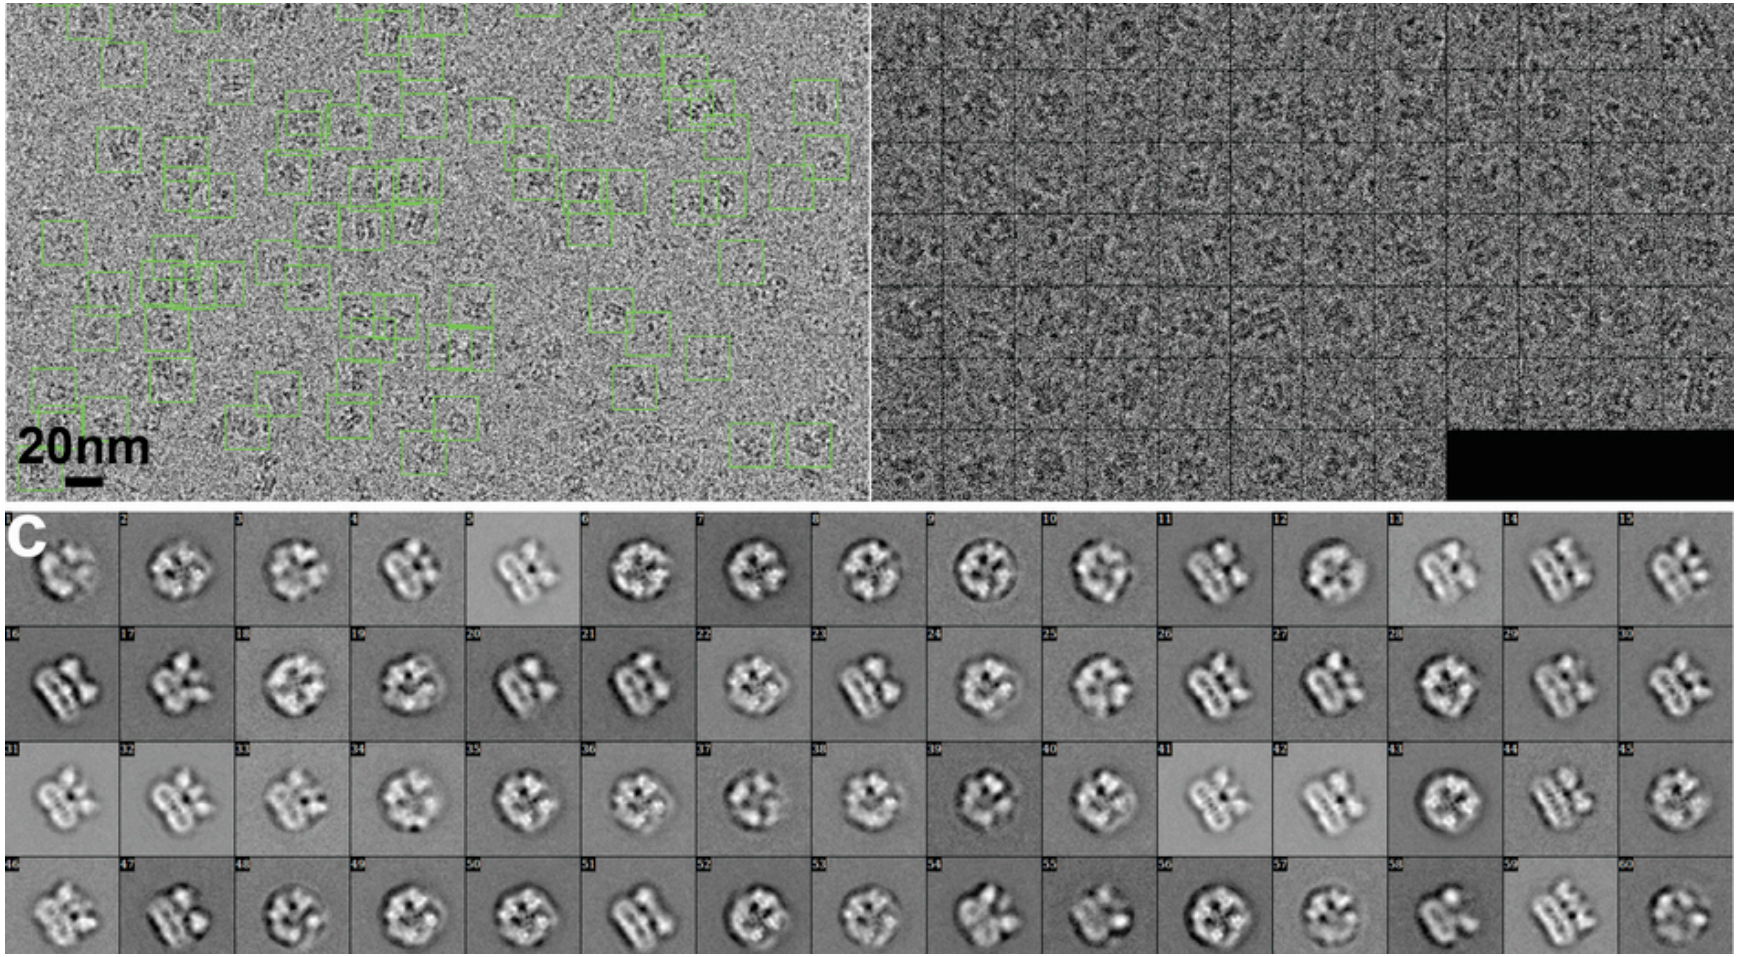

Liao et al., Nature 2013

An image represented in a high-dimensional Euclidean space.

- An image represented by an array of N x M pixels can be thought as a vector in a (generalized) Euclidean space with N x M dimensions.
- For example, an image of 64 x 64 pixels is a vector in a 4096-dimensional space.
- If two images are "similar" it means the distance between the vectors representing them is small. That is, the vector end points lie close together.
- Groups of similar images form clusters in the generalized Euclidean space.
- To show the concept, and introduce an important tool for classification, I will use a simplistic image containing only two pixels.

Generalized Euclidean distance

Euclidean distance between two images f_1 and f_2:

$$E^2_{12}(\alpha, r') = \sum_{j=1}^{J} | f_1(r_j) - f_2(R_\alpha r_j + r') |^2$$

$$= \sum_{j=1}^{J} | f_1(r_j) |^2 + \sum_{j=1}^{J} | f_2(R_\alpha r_j + r') |^2 - \sum_{j=1}^{J} | f_1(r_j) f_2(R_\alpha r_j + r') |$$

const. const. cross-correlation

Similarity = closeness in high-dimensional Euclidean space = small E-distance
→ large value of CCF peak at matching position

Introducing: a set of images, each consisting of 2 pixels

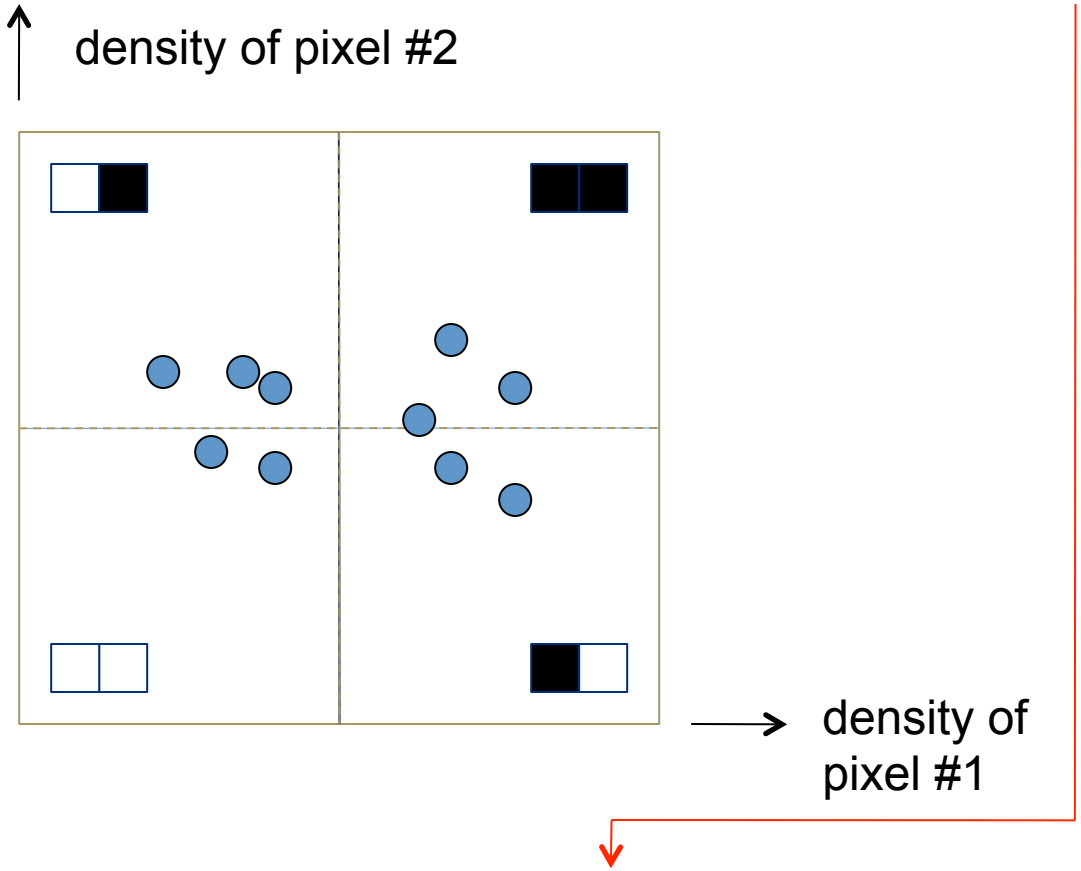

Similarity = closeness in 2-D Euclidean space

Two images are similar if their (generalized) Euclidean distance is small

A set of images consisting of two pixels: Intro into classification

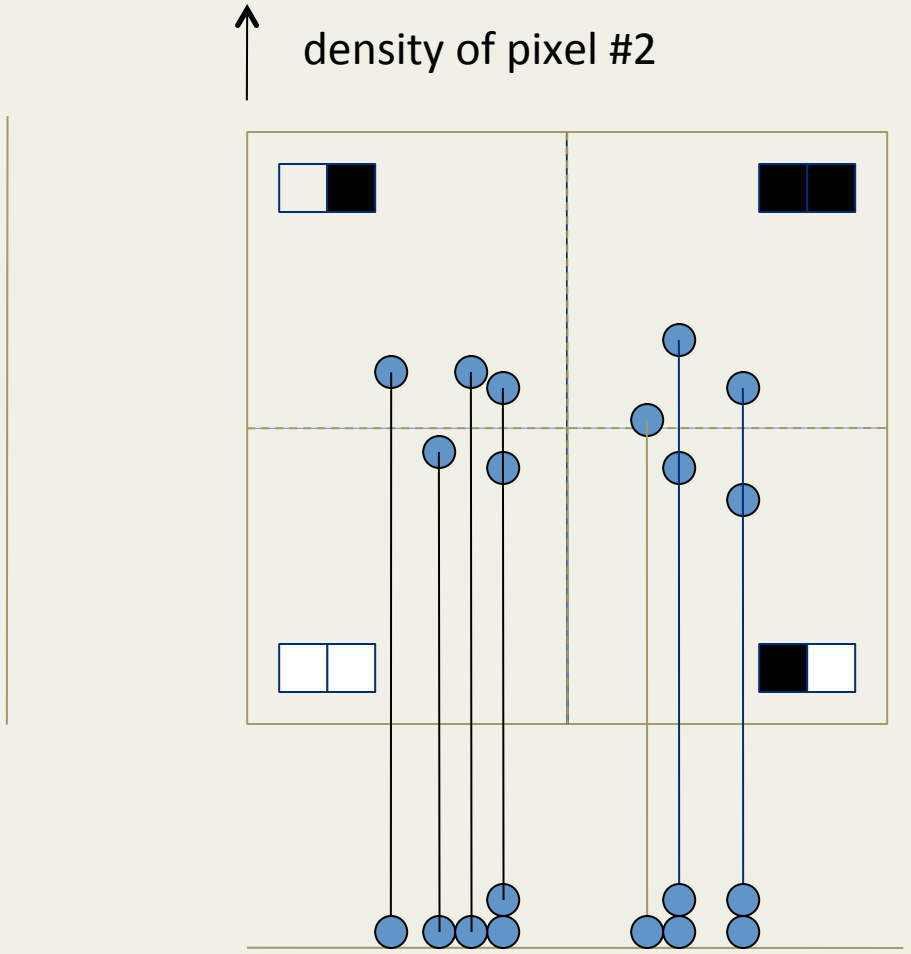

Fortuitous case: direction of grouping happens to coincide with a primary axis.

Projection onto axis 1

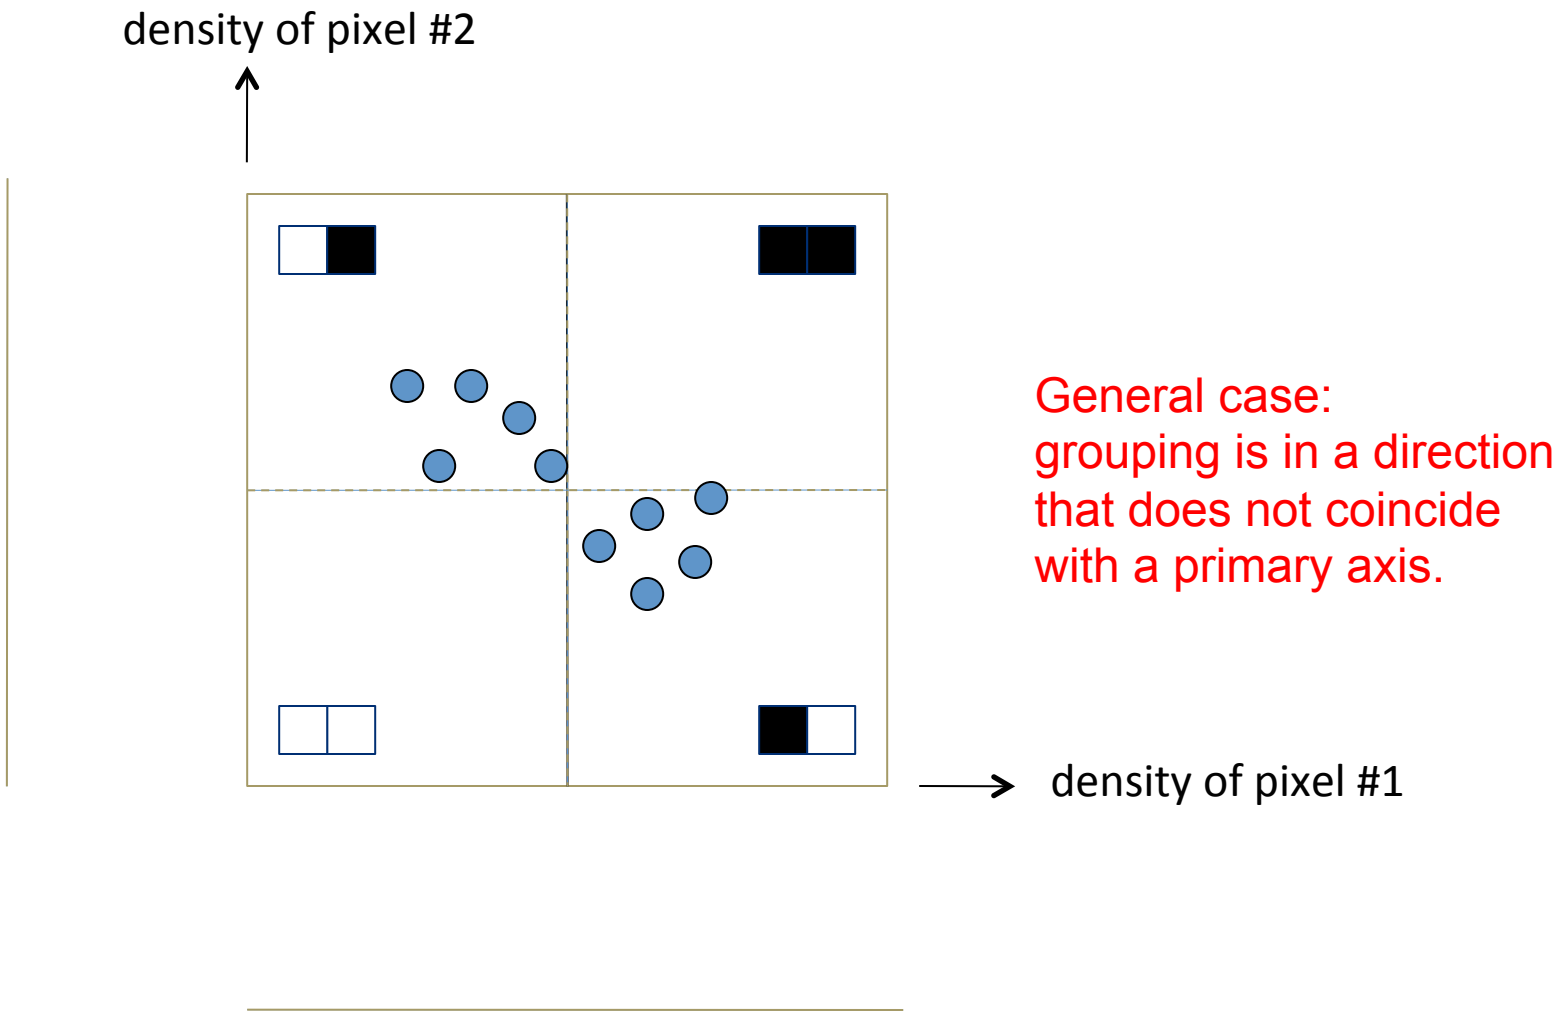

General case:
grouping is in a direction
that does not coincide
with a primary axis.

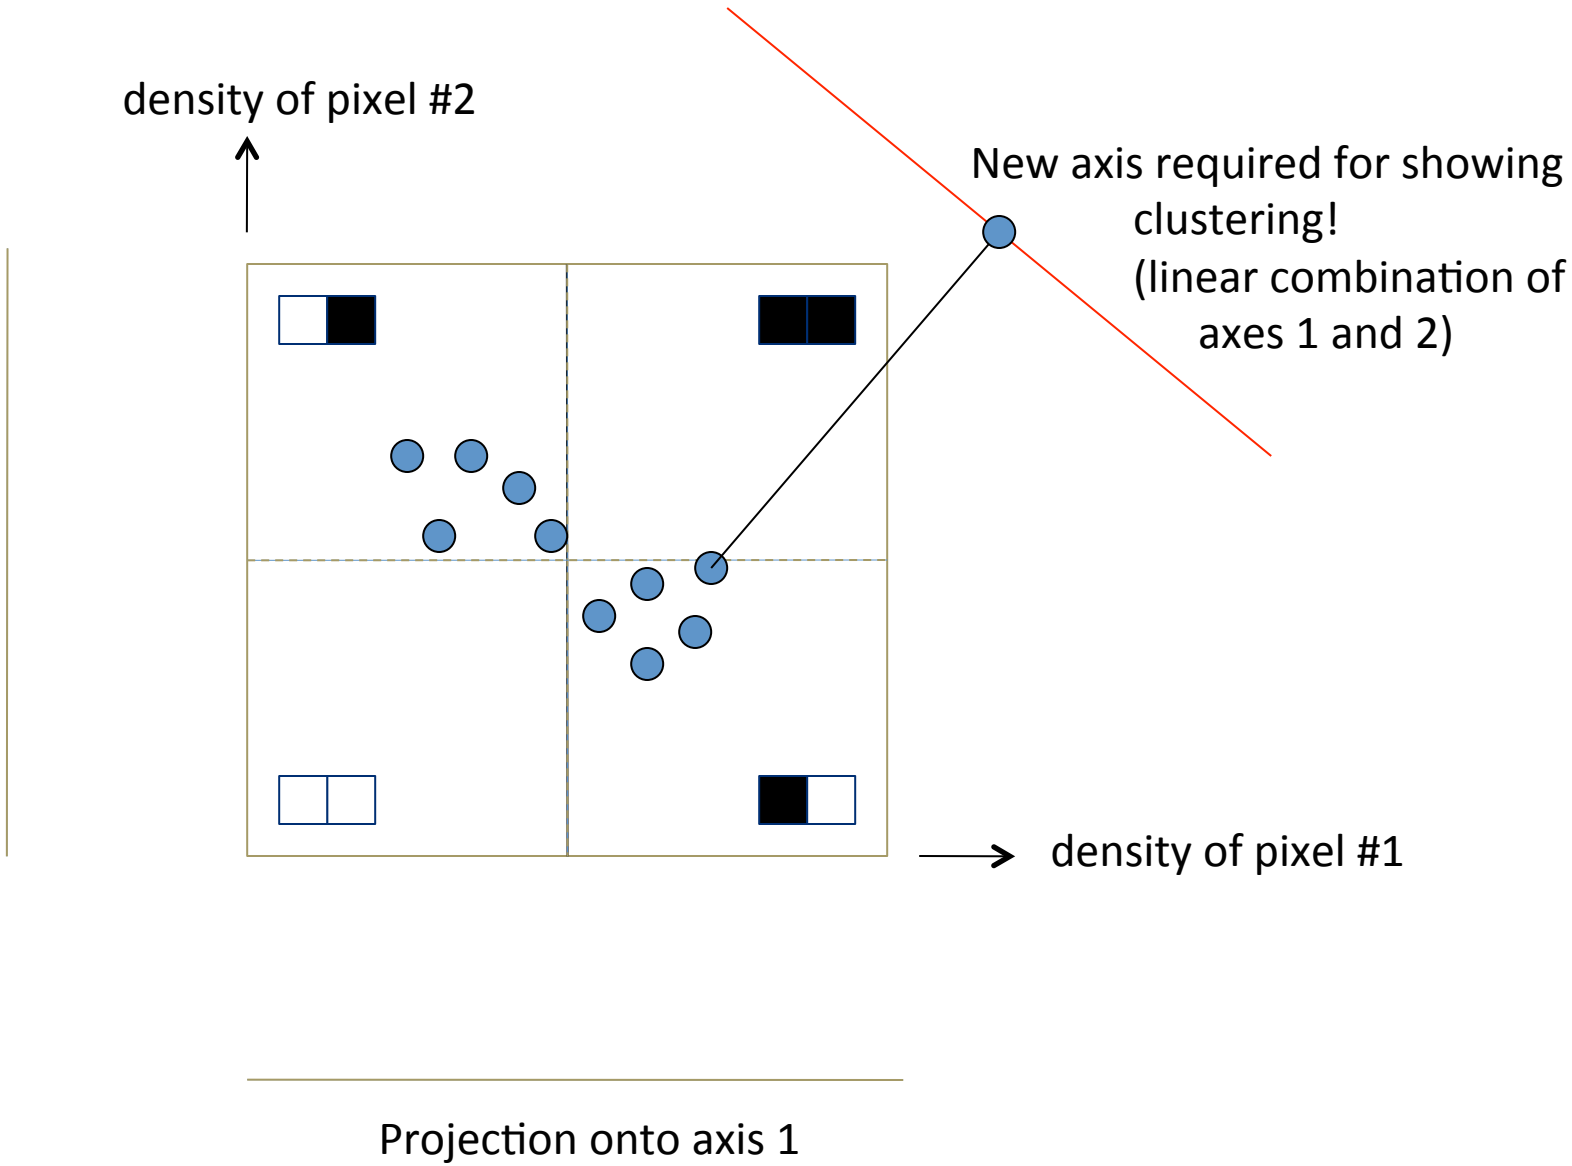
density of pixel #2
New axis required for showing clustering!
(linear combination of axes 1 and 2)
density of pixel #1
Projection onto axis 1

Tools: Classification, and the Role of MDA

- Classification deals with "objects" in the space in which they are represented.
- For instance, a 64x64 image is an "object" in a 4096-dimensional space since, in principle, each of its pixels can vary independently.

  Let's say we have 8000 such images. They would form a cloud with 8000 points in this space. This is an unwieldy problem.
- Unsupervised classification is a method that is designed to find clusters (regions of cohesiveness) in such a point cloud.
- Role of Multivariate Data Analysis (MDA): find a space ("factor space") with ***reduced*** dimensionality for the representation of the "objects". This greatly simplifies classification.
- Reasons for the fact that the space of representation can be ***much smaller*** than the original space: ***resolution limitation*** (neighborhoods behave the same), and ***lateral correlations*** due to the physical origin of the variations (e.g., movement of a structural component is represented by correlated additions and subtractions at the leading and trailing boundaries of the component).

Principle of MDA:

Find new coordinate system, tailored to the data

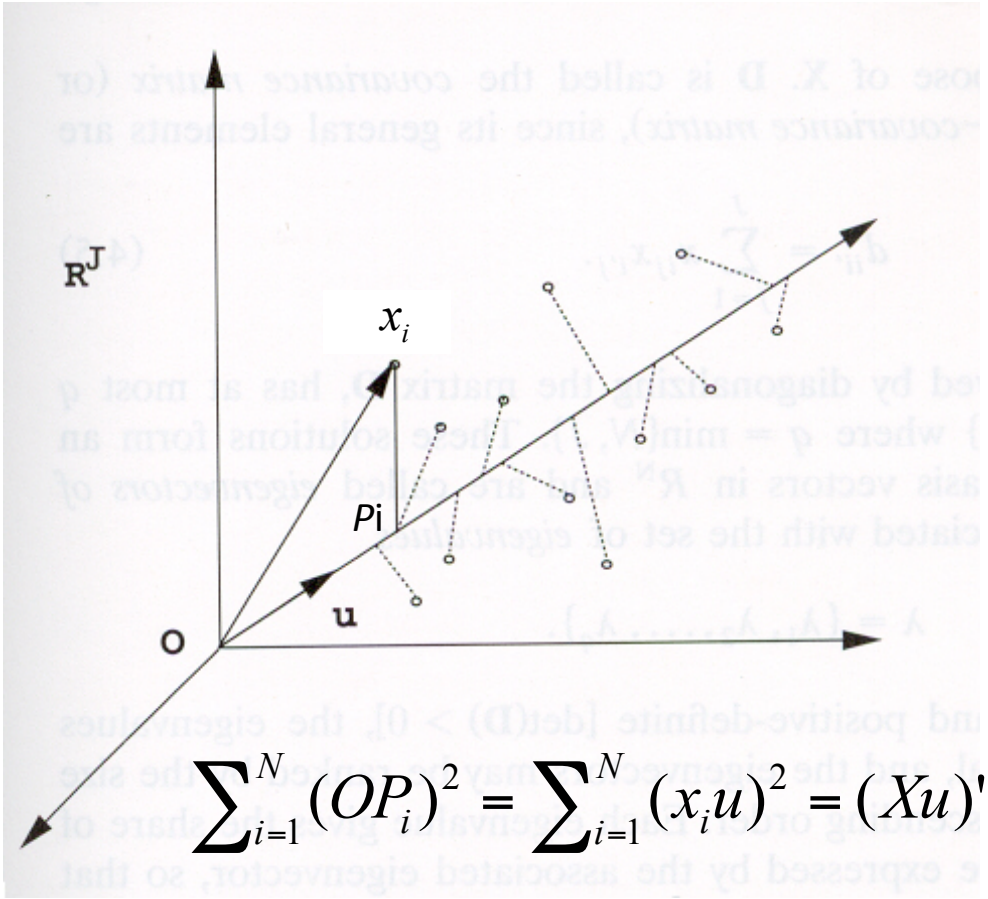

X= matrix containing N image vectors (each with J elements) as rows

$$\sum_{i=1}^{N}(OP_i)^2 = \sum_{i=1}^{N}(x_i u)^2 = (Xu)'X'u = u'X'Xu \xrightarrow{!} \max$$

p. 151 [note error in book figure!]

32 x 32 phantom images in 8 (= 2^3) varieties

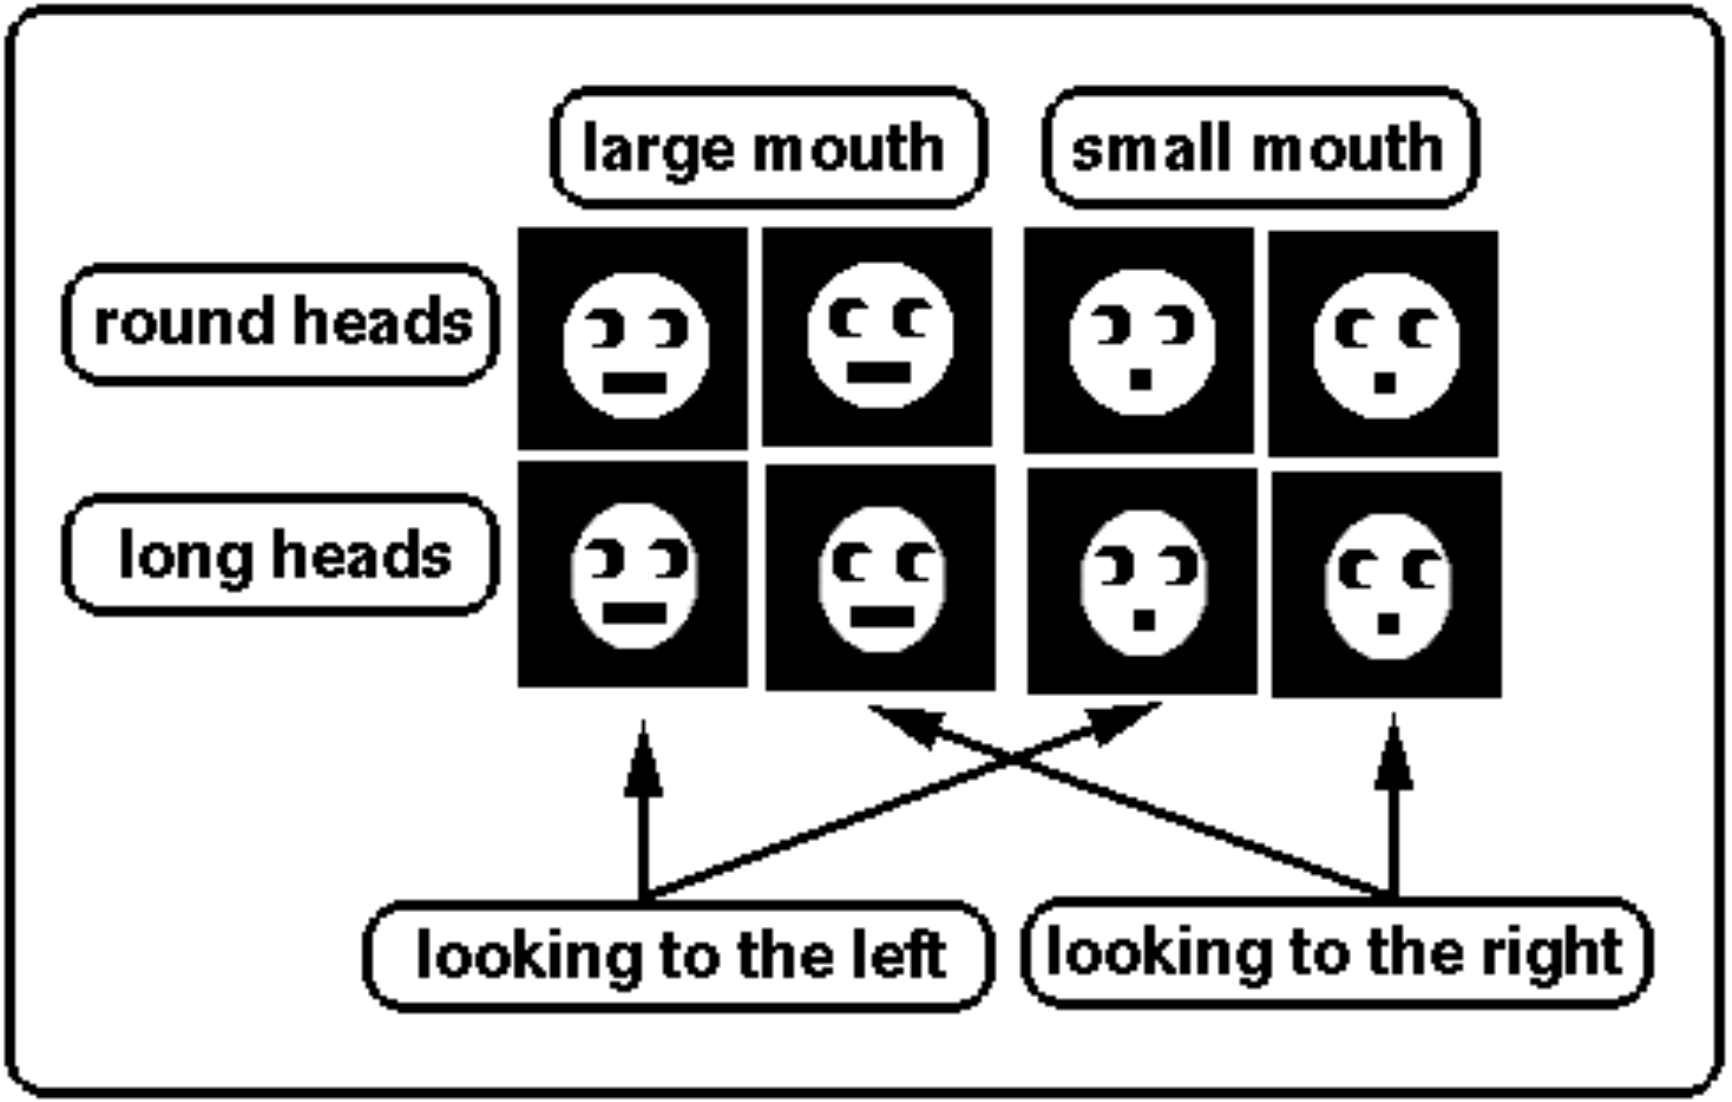

Brétaudière JP and Frank J (1986) *J. Microsc.* **144**, 1-14

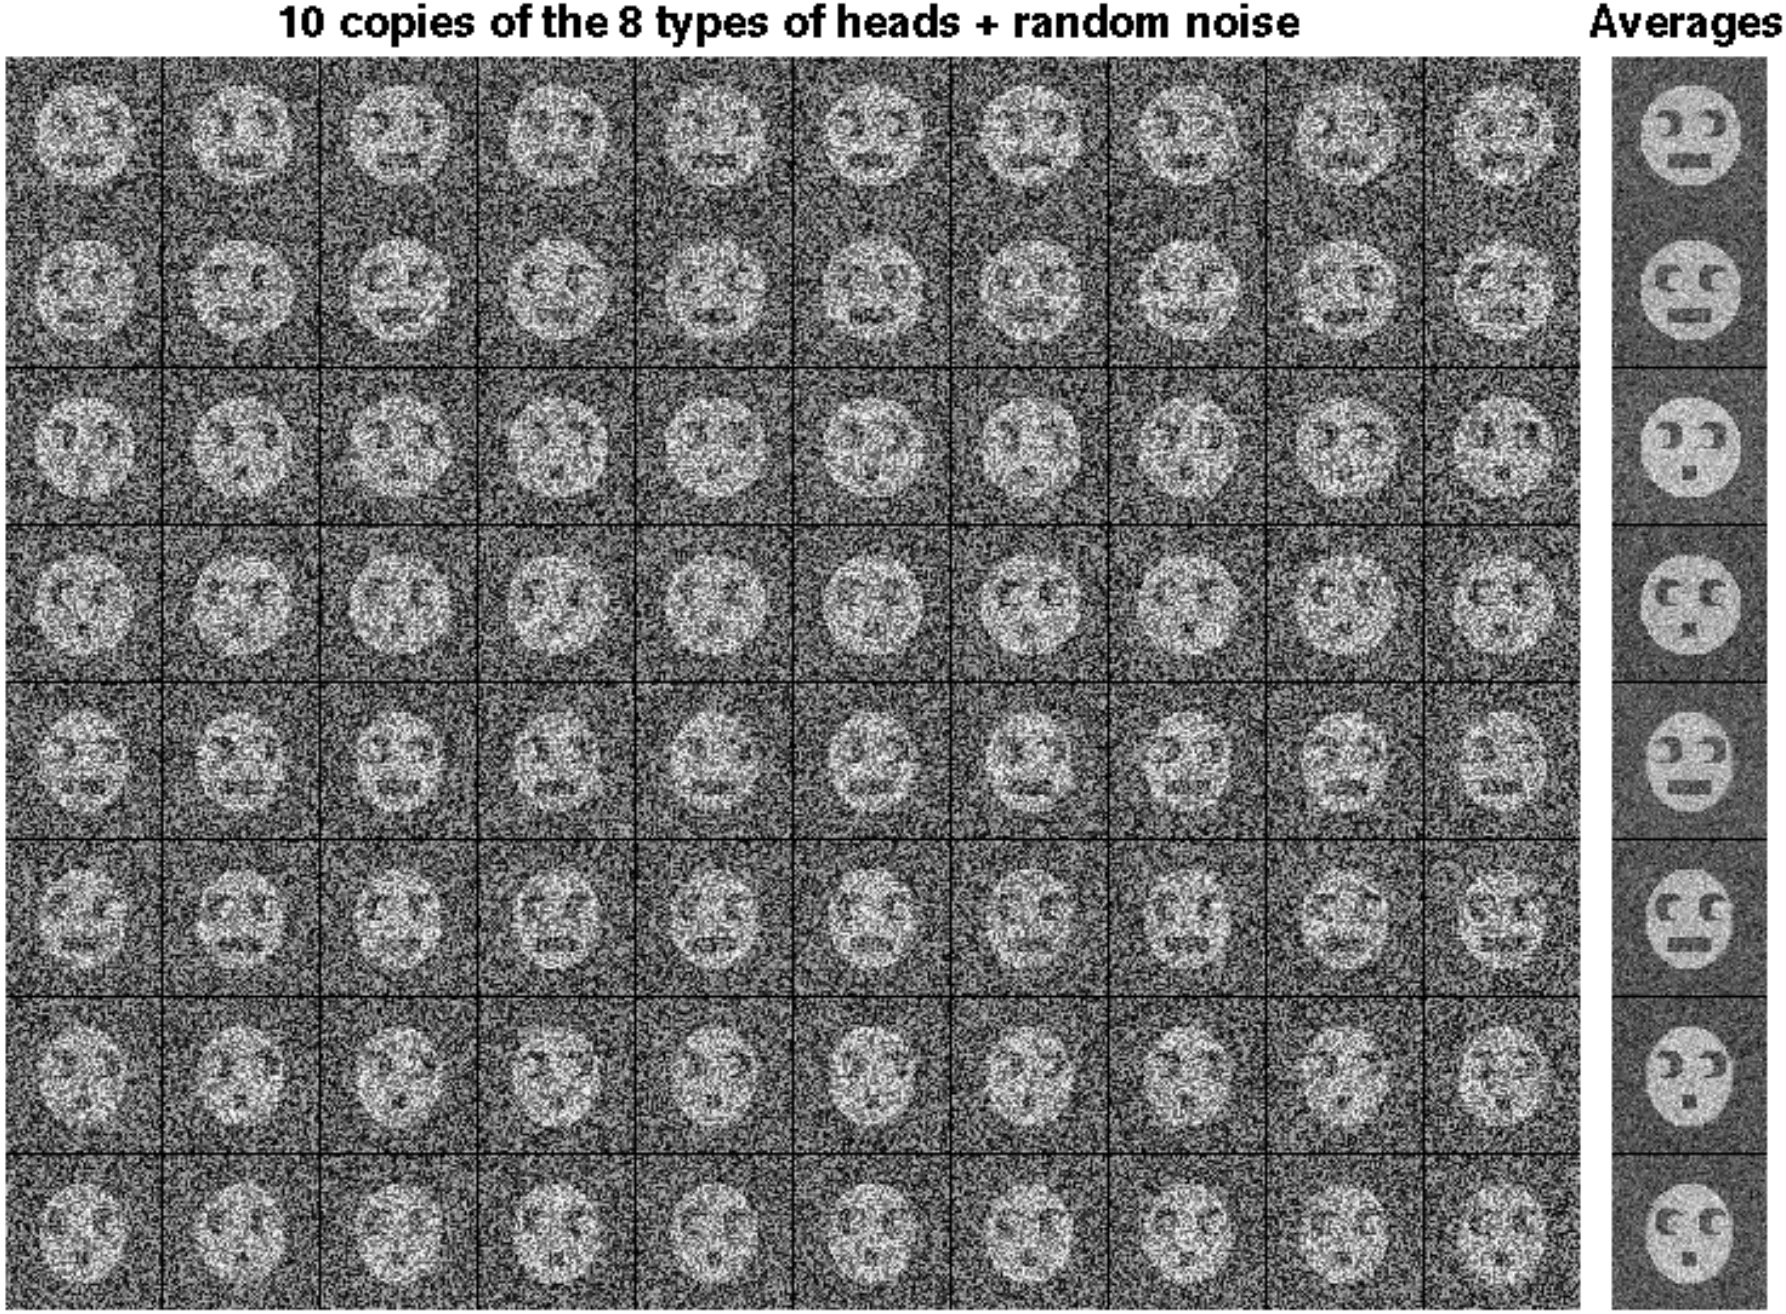
10 copies of the 8 types of heads + random noise
Averages

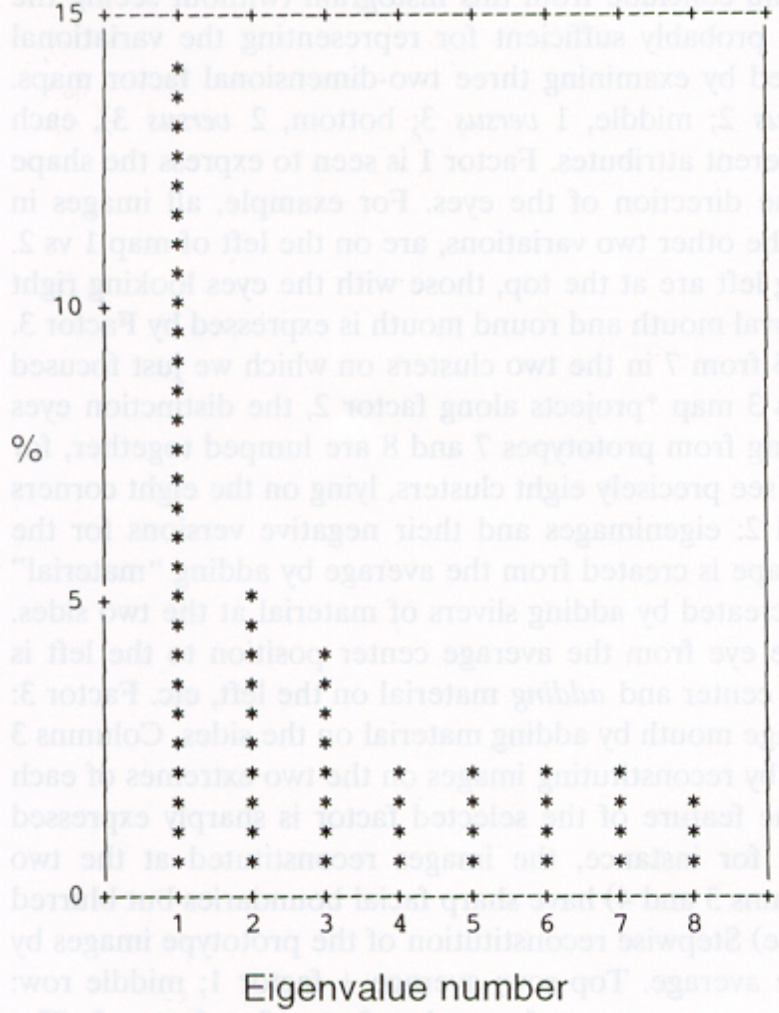

eigenvalue histogram

3 stand out,
i.e., 3 factors are sufficient.!

1 vs 2

Data are clustered in the eight corners of a 3D factor space formed by the first three factors.

1 vs 3

2 vs 3

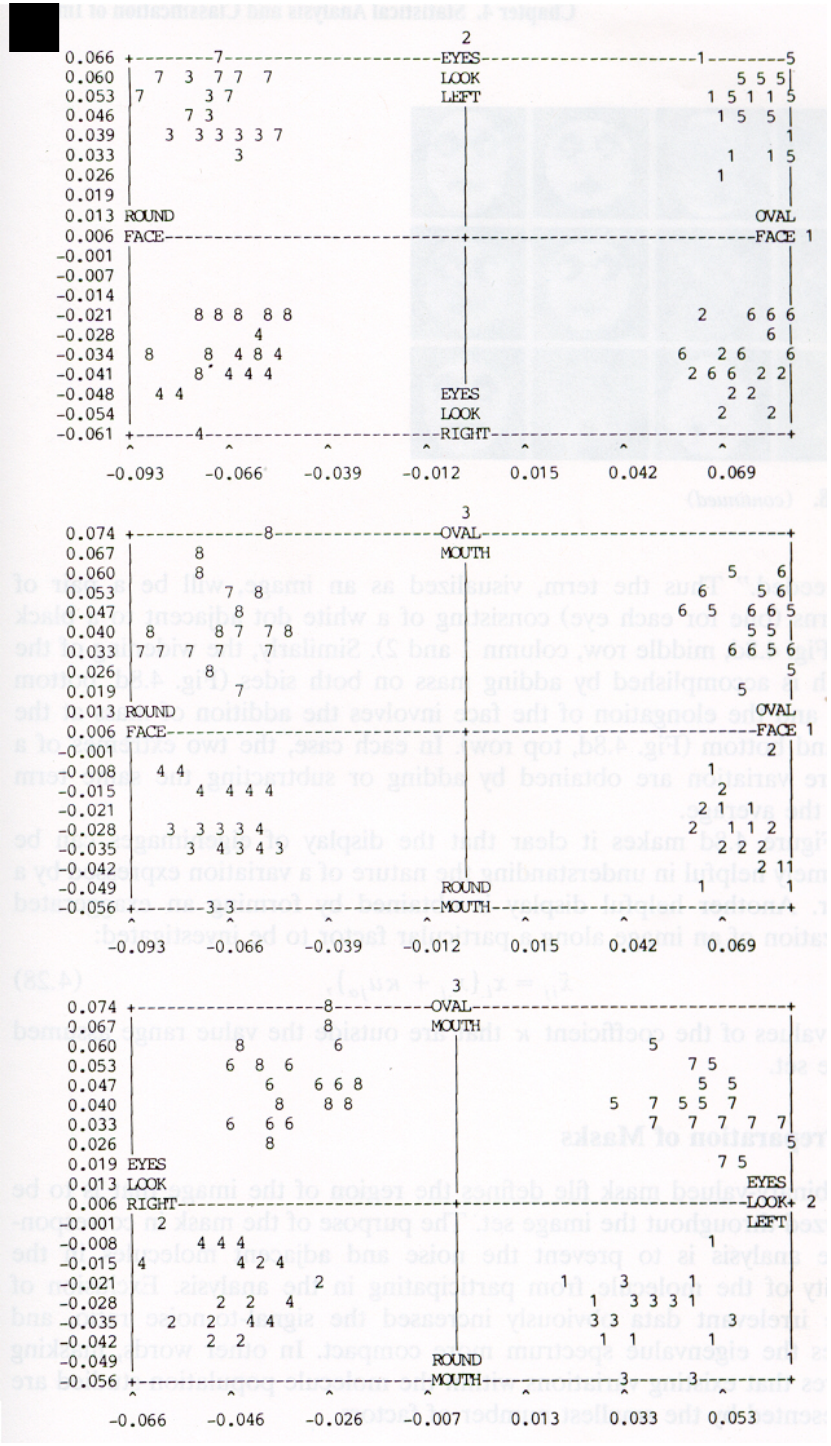

MDA: eigenimages

- Factor 1
- Factor 2
- Factor 3

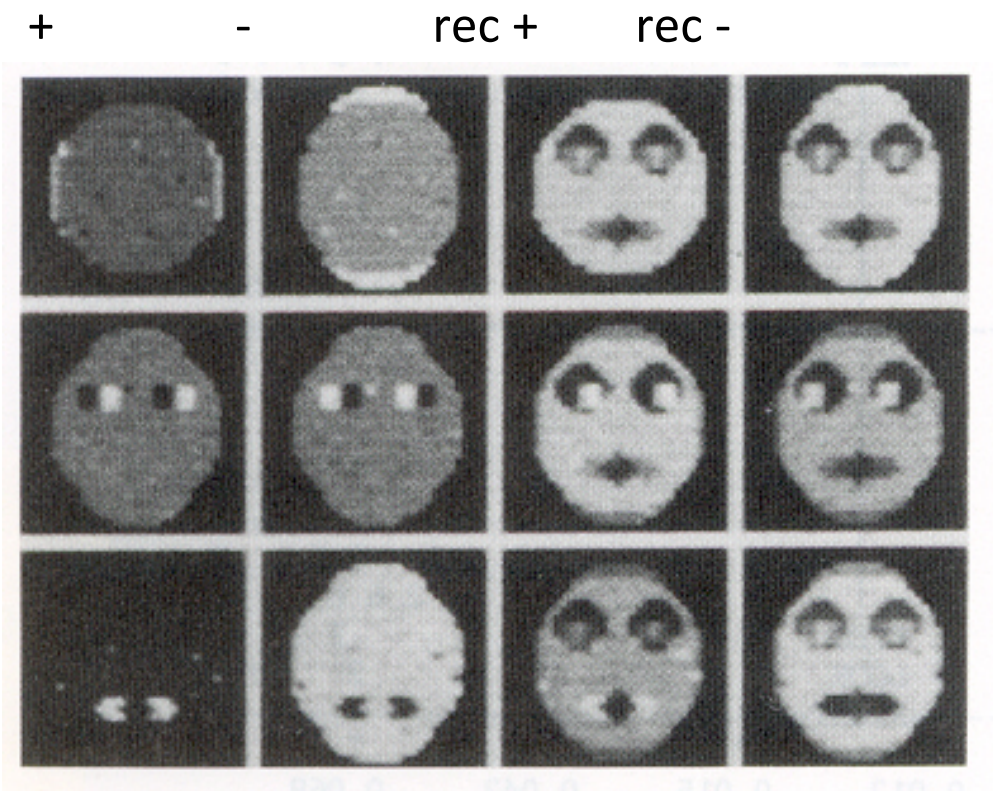

p. 175

Example Factor 2: position of eyes. Average contains centered oval (from superposition of all images containing eyes shifted to right with those having eyes shifted to the left).

To express the eye-related features of the images, the eigenimage must either subtract density on the left and add it on the right, or vice-versa.

- Factor 1
- Factor 2
- Factor 3

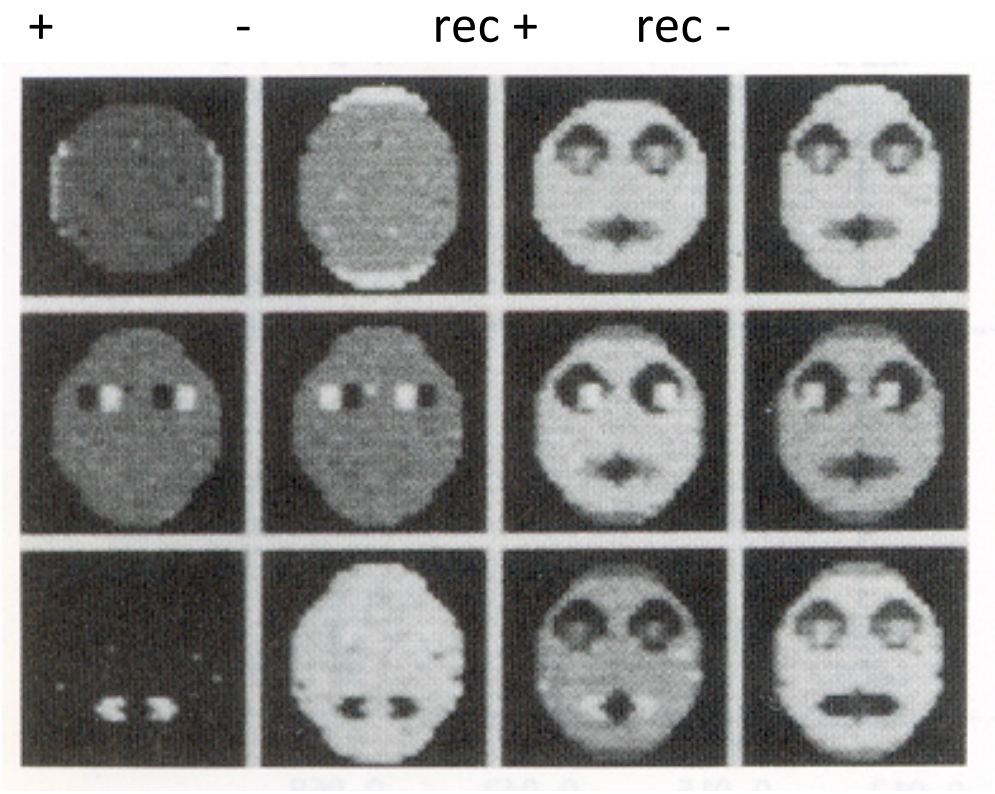

Class Averages

Instances

Total Avrg + F1

Total Avrg + F1+F2

Total Avrg + F1+F2+F3

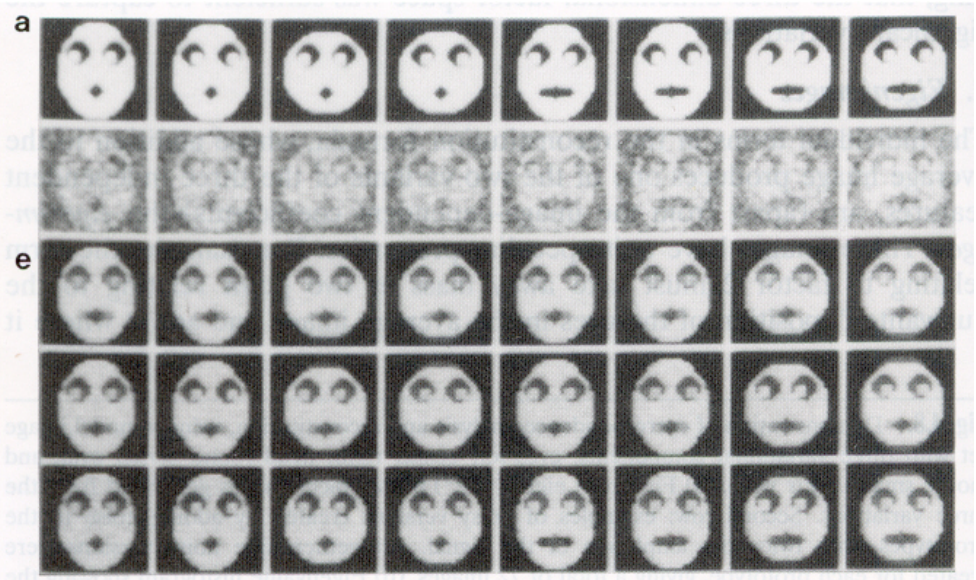

p. 176

+ -

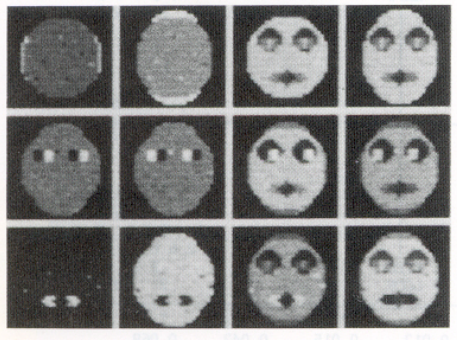

p. 175

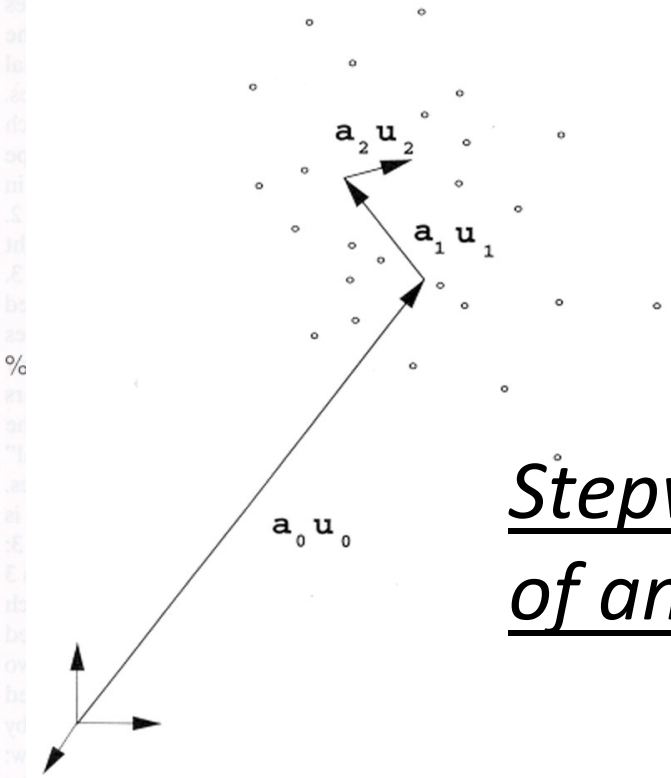

Stepwise reconstitution of an image

p. 159

3D reconstruction -- preliminaries

- Under what conditions are projections of an object similar to one another?
- Similarity ←→ closeness in high-dim E-space
  ←→ belonging to the same cluster
  ←→ high correlation

Shape Transform

- The Shape Transform is the Fourier transform of a binary mask function (1 inside, 0 outside) whose shape is the shape of an object in 2D or 3D
- It indicates the size and shape of the local region in Fourier space within which Fourier coefficients are correlated/ dependent.

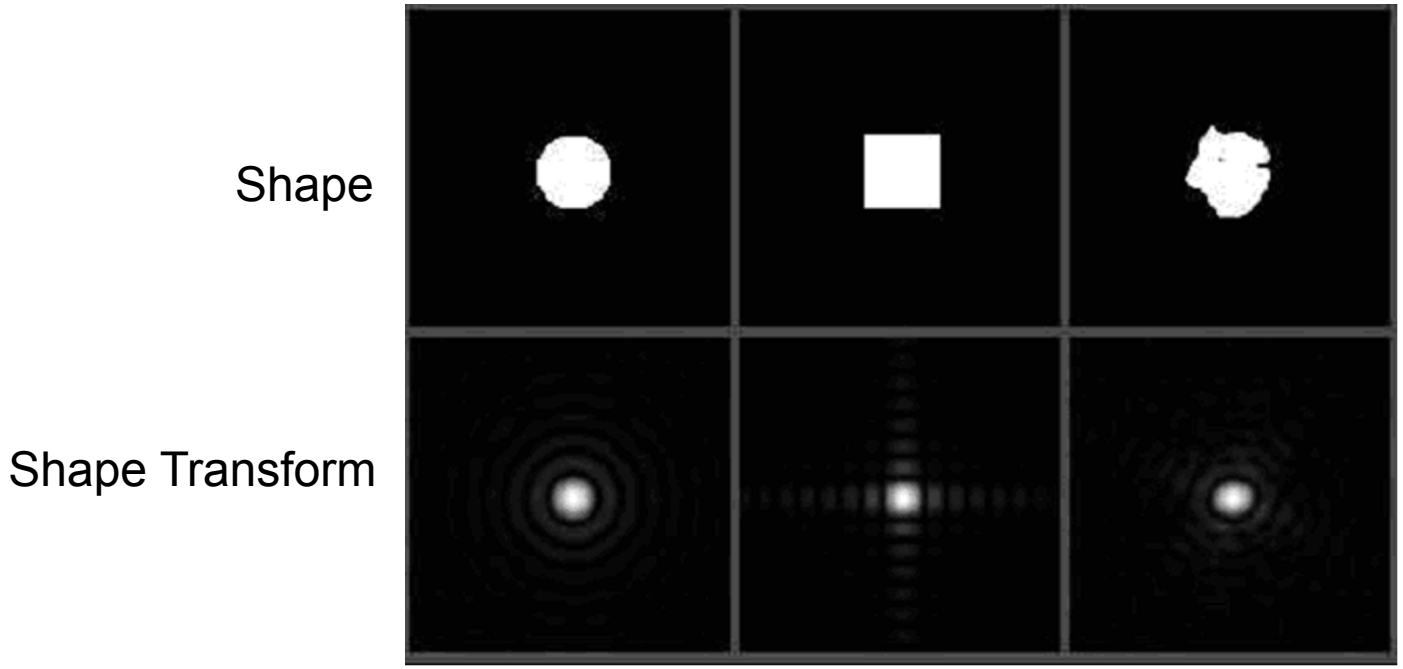

Shape Transforms

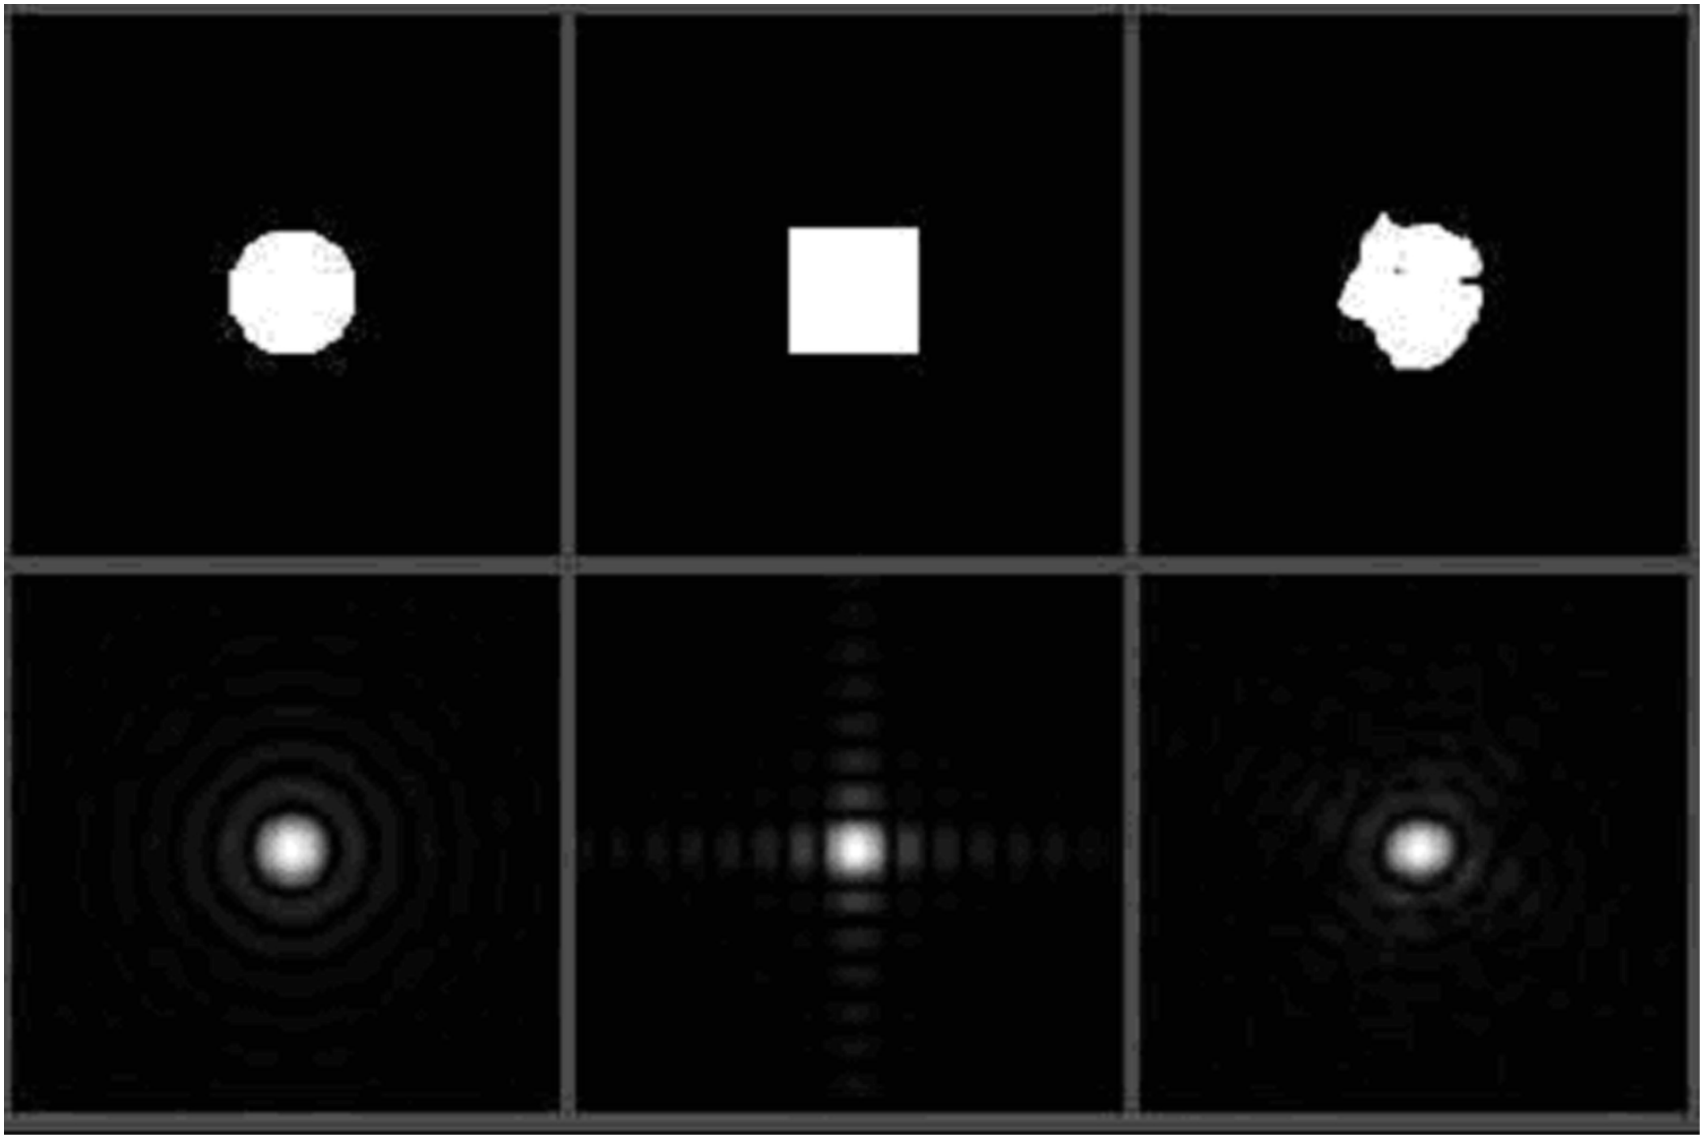

Similarity of projections, condition for 3D reconstruction, and “kissing” shape transforms

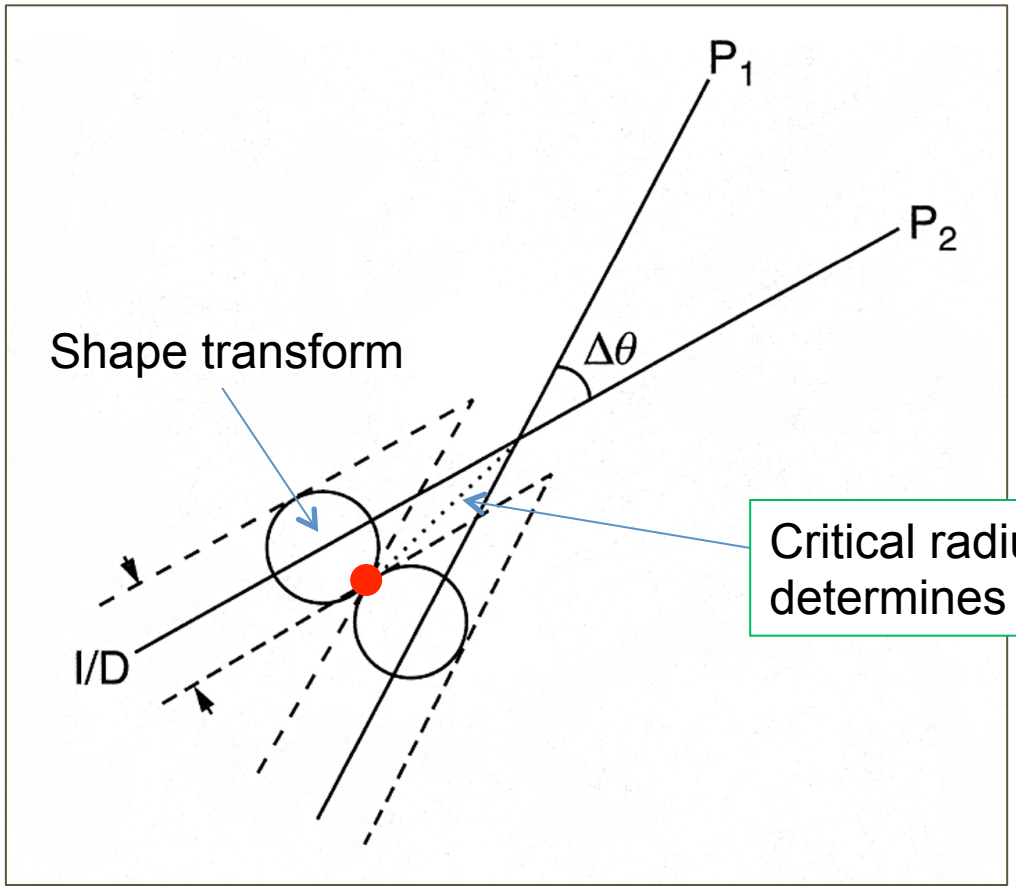

P1, P2 central sections in Fourier space.
Δθ angle subtended by P1, P2
D particle diameter

Determination of Particle Orientations

(A) unknown structure -- bootstrap
- (1) Random-conical (uses unsupervised classification)
- (2) Common lines/ angular reconstitution (uses unsupervised classification)

(B) known structure – low-res map available
- (1) reference-based (3D projection matching = a form of supervised classification)
- (2) common lines/ angular reconstitution

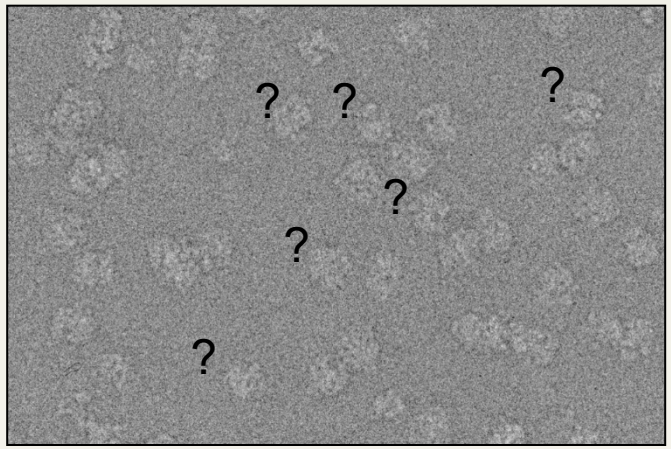

RANDOM CONICAL RECONSTRUCTION

0-degree view

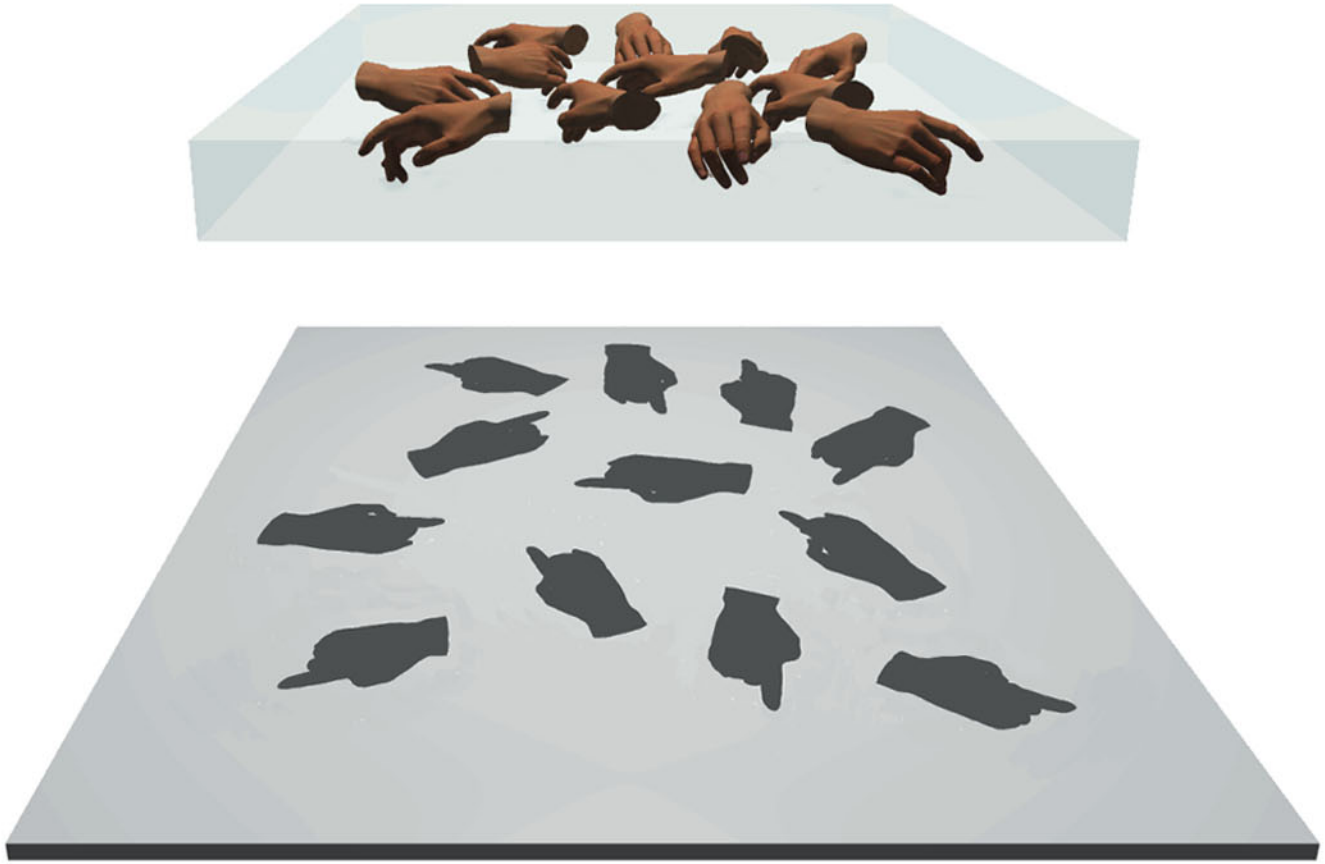

50-degree view

Equivalent geometry
in the coo system of the particle

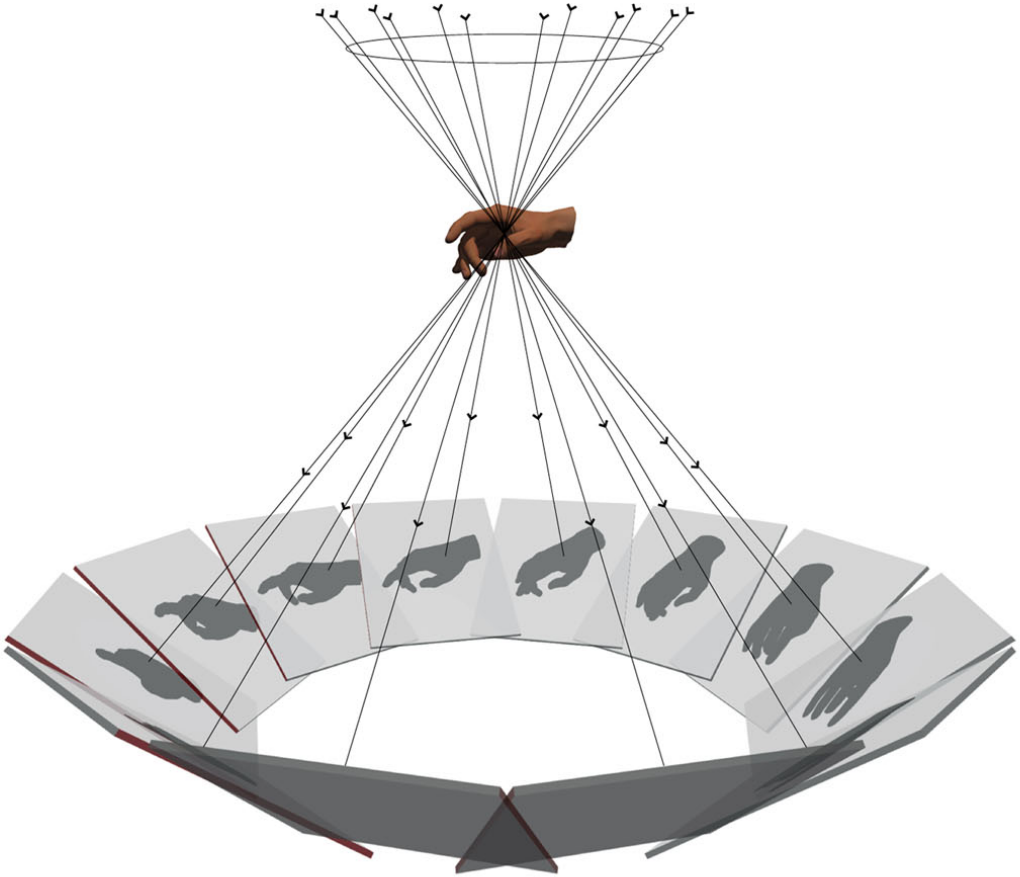

Conical Data Collection Geometry in Fourier Space

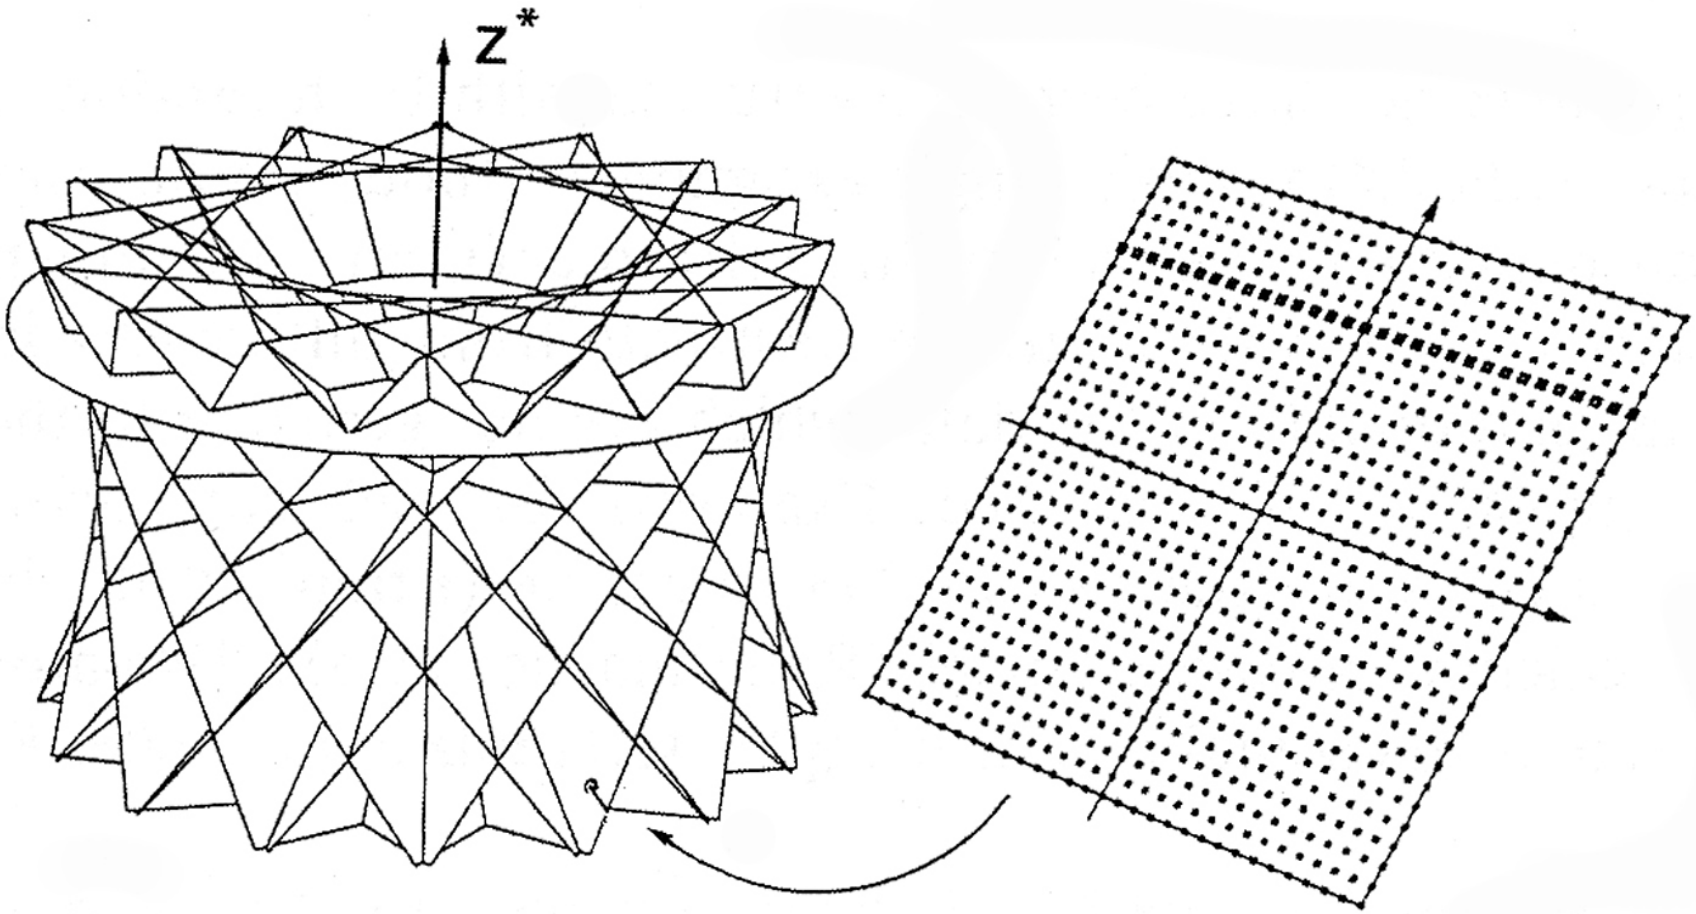

Lanzavecchia et al.

COMMON LINES APPROACH TO DETERMINING ORIENTA

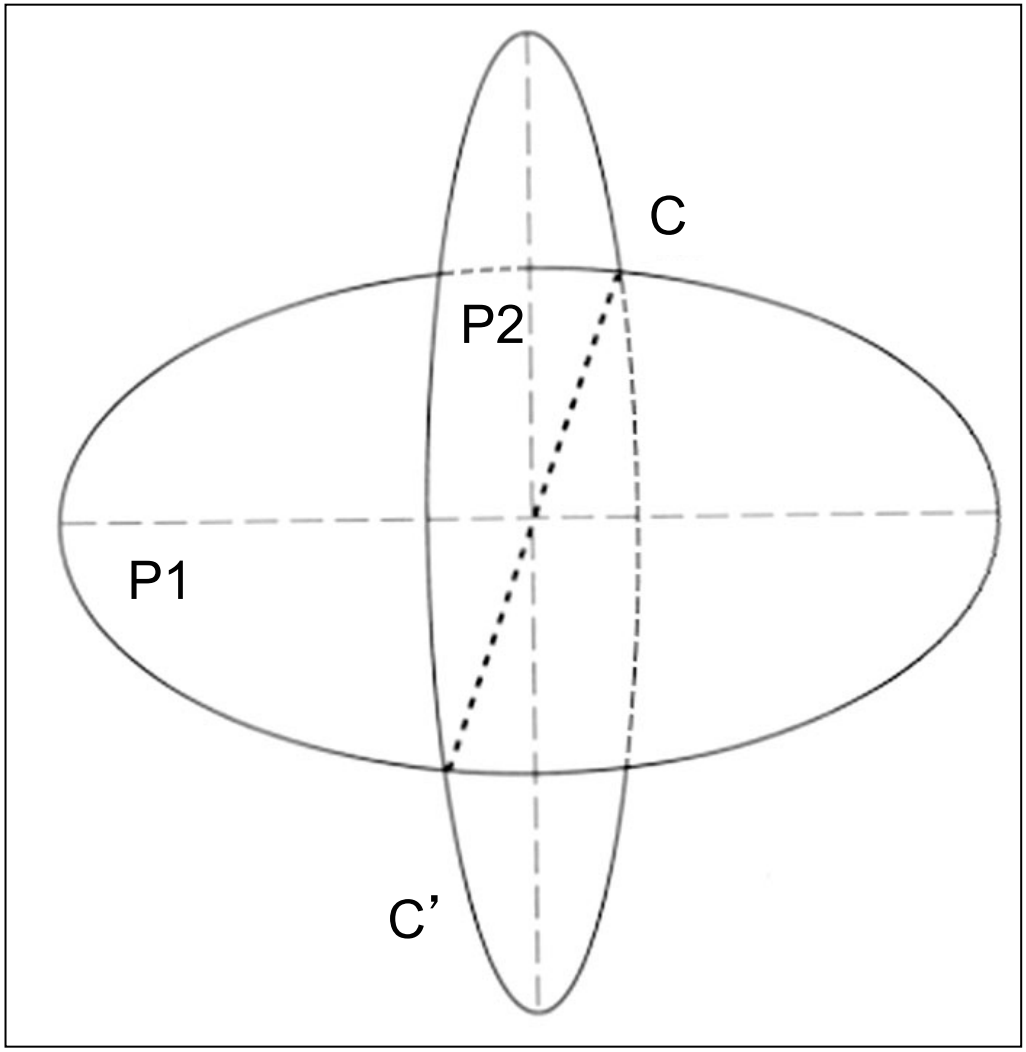

Common line C-C’ of two projections represented by central sections P1 and P2

Two 2D projections of the same 3D object have in common:

in Fourier space: Fourier coefficients along the line of intersection: “common line”

in real space: 1D projections in particular directions

The Sinogram (or Radon transform) of an image is an exhaustive ordered display of all of its1-D projections

Comparing two sinograms, one can find the angles for which maximum agreement is reached

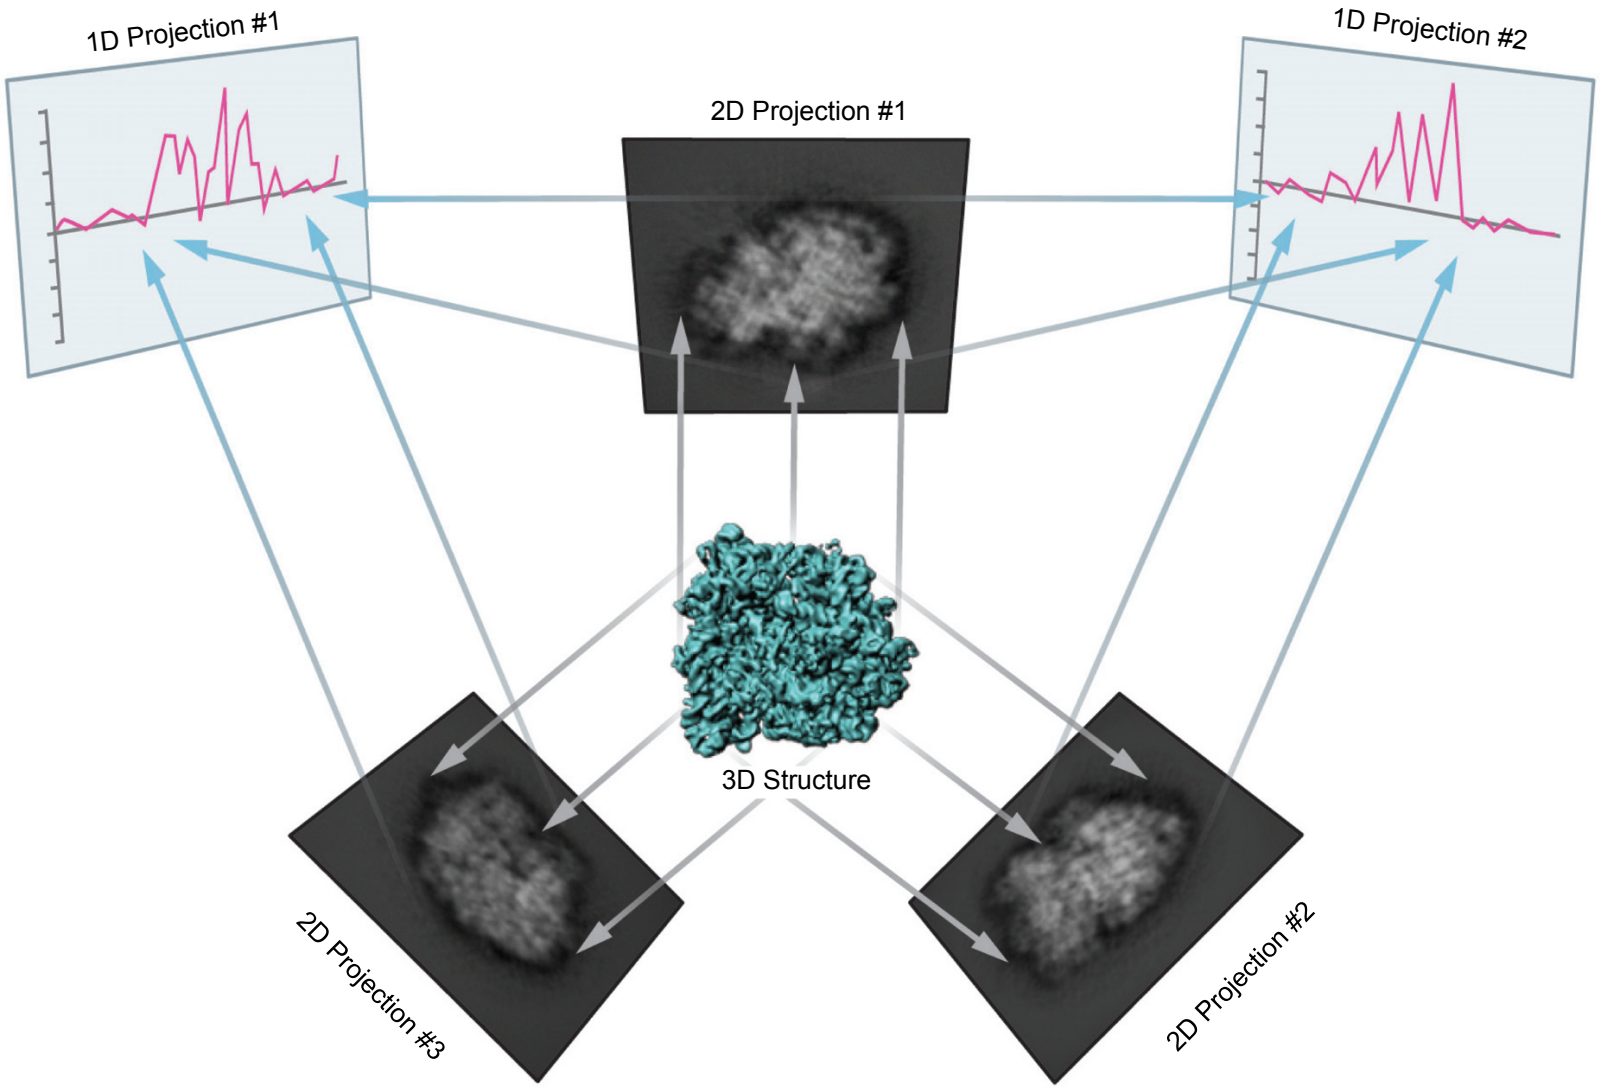

J. Frank, in *Molecular Machines in Biology* 2011

Determination of orientation by projection matching

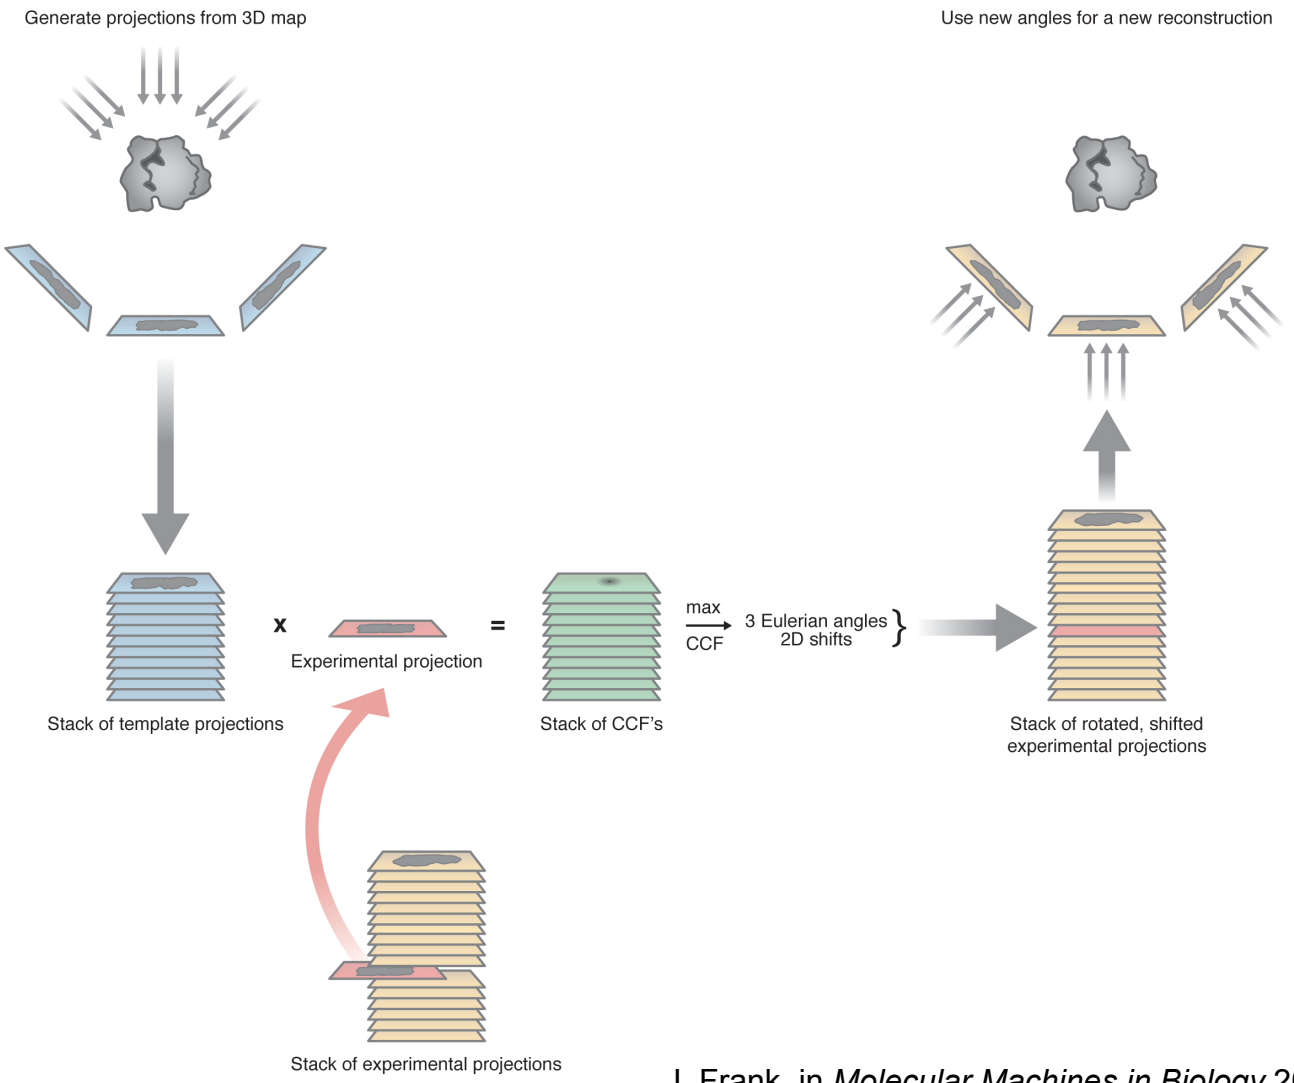

J. Frank, in *Molecular Machines in Biology* 2011

Initial Angular Grid

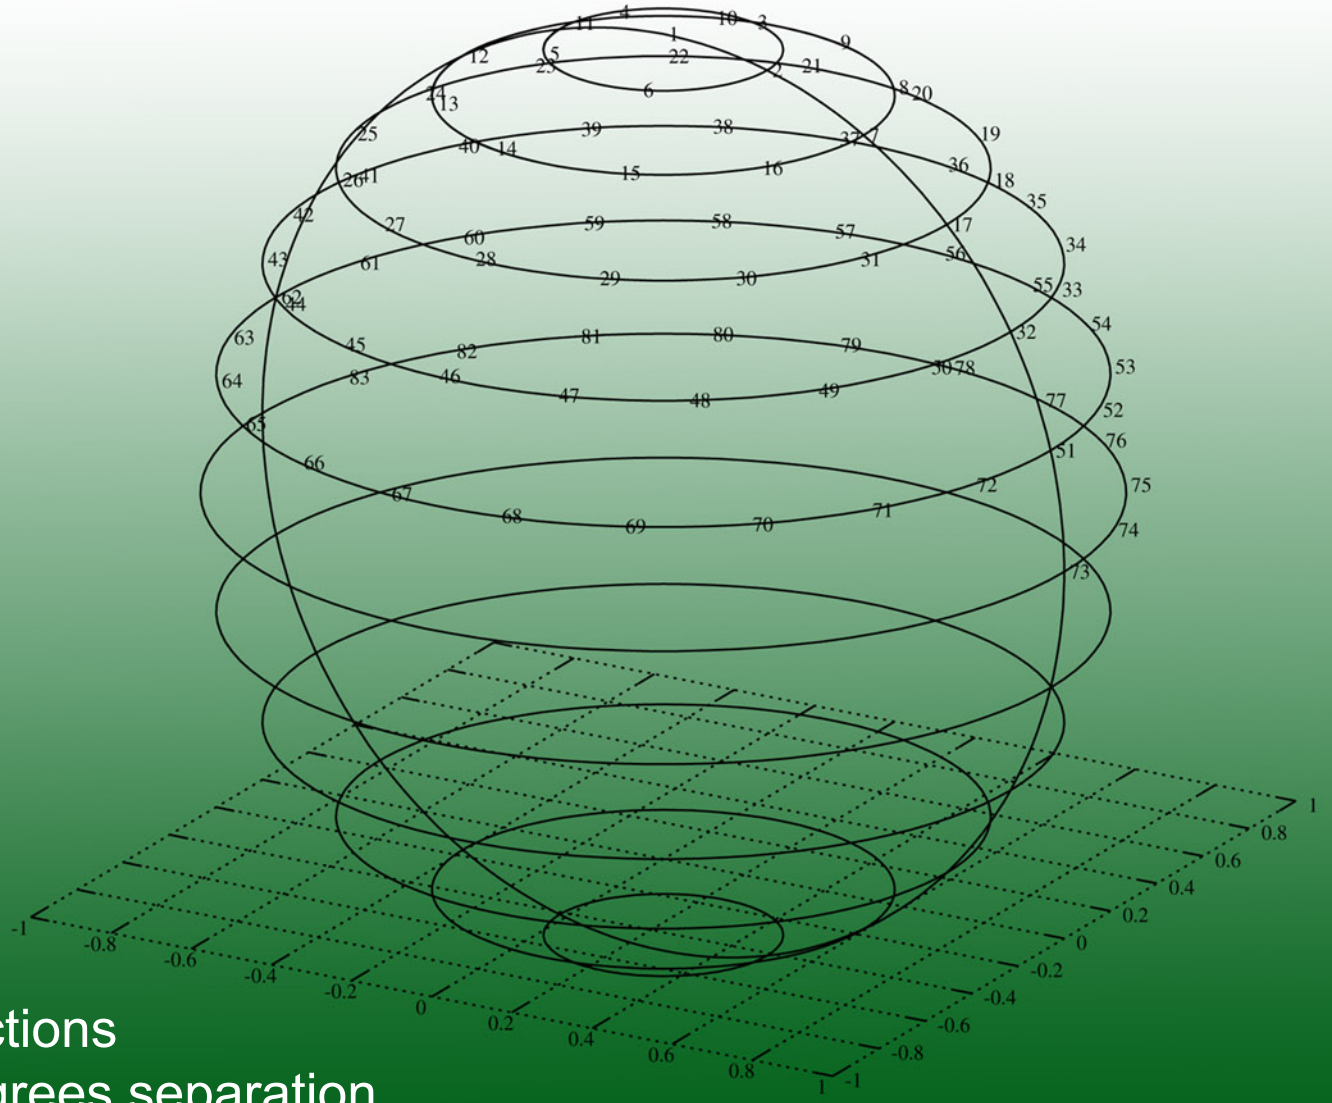

83 directions
~15 degrees separation

83-projection grid

averages of particles classified

Iterative Angular Refinement

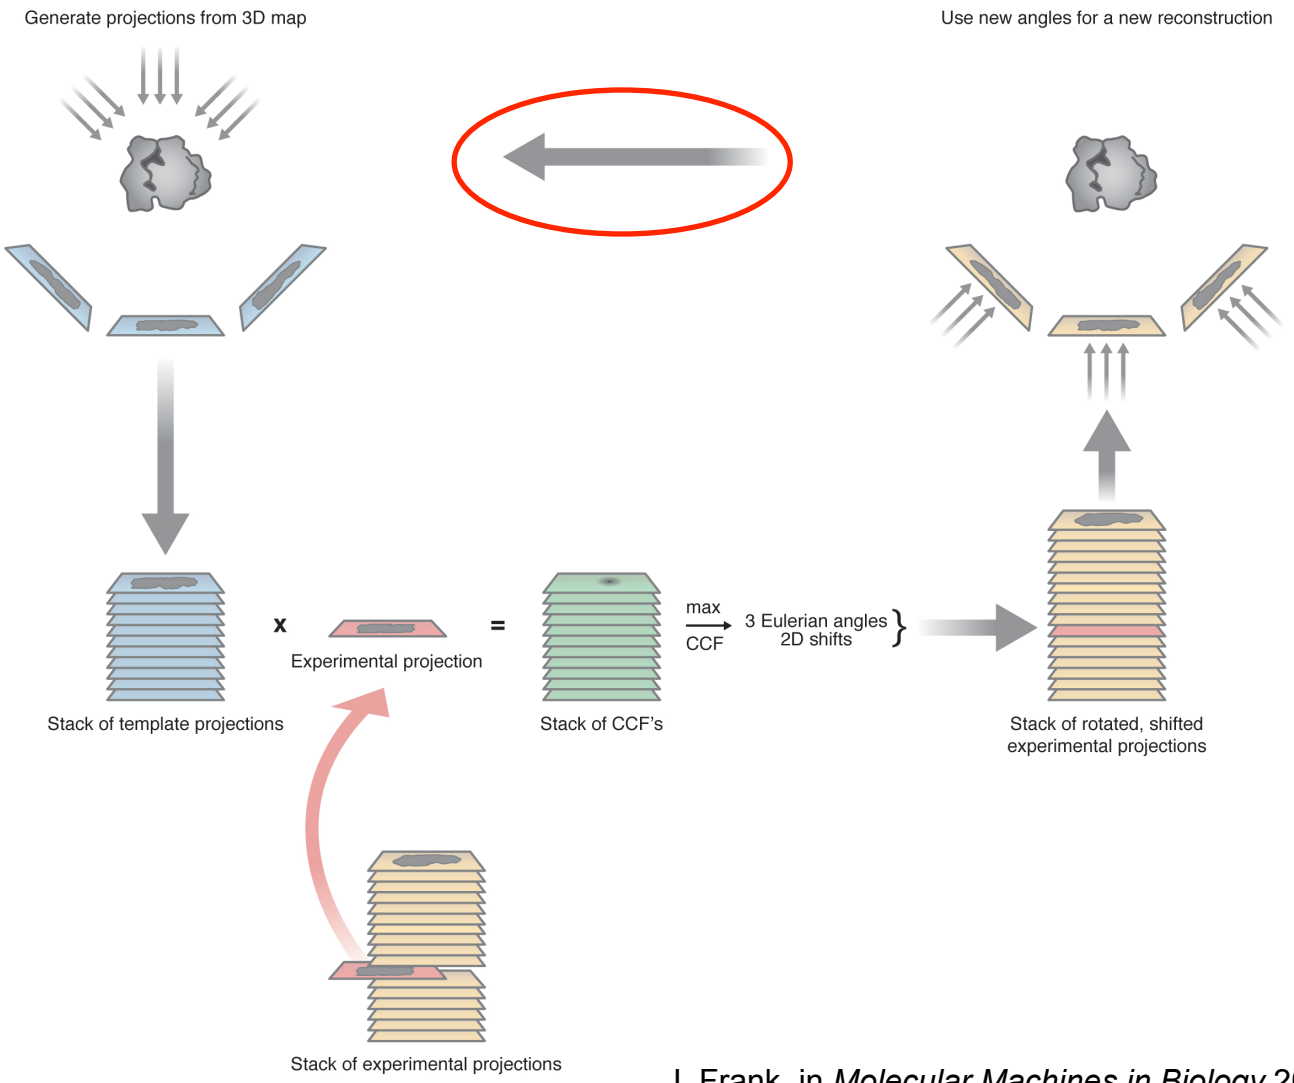

J. Frank, in *Molecular Machines in Biology* 2011

Angular Refinement

Given an initial 3D reference,

Iterate the steps {3D projection matching + reconstruction}
Decrease angular grid size as you go on (range: 15° ➔ 0.5°)

Convergence criteria:
(1) convergence of particle angles
(2) convergence of resolution (monitor progress with FSC)

"Rule of neighborhood" saves computing time

- beware of problem of reference-dependence!

- Start with coarse grid (15 degrees)
- Decrease angular separation, down to 0.5 degrees
- At some point, switch from global coverage to local coverage of previously determined angles

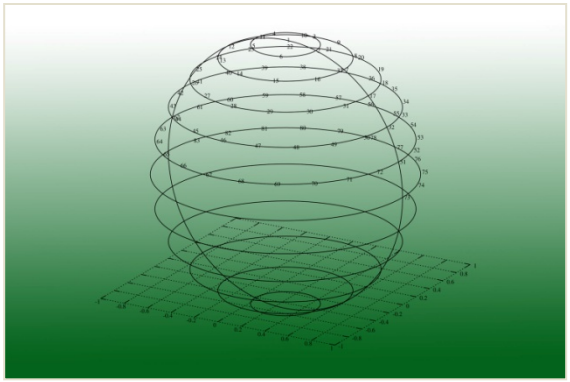

Increasingly finer angular increments

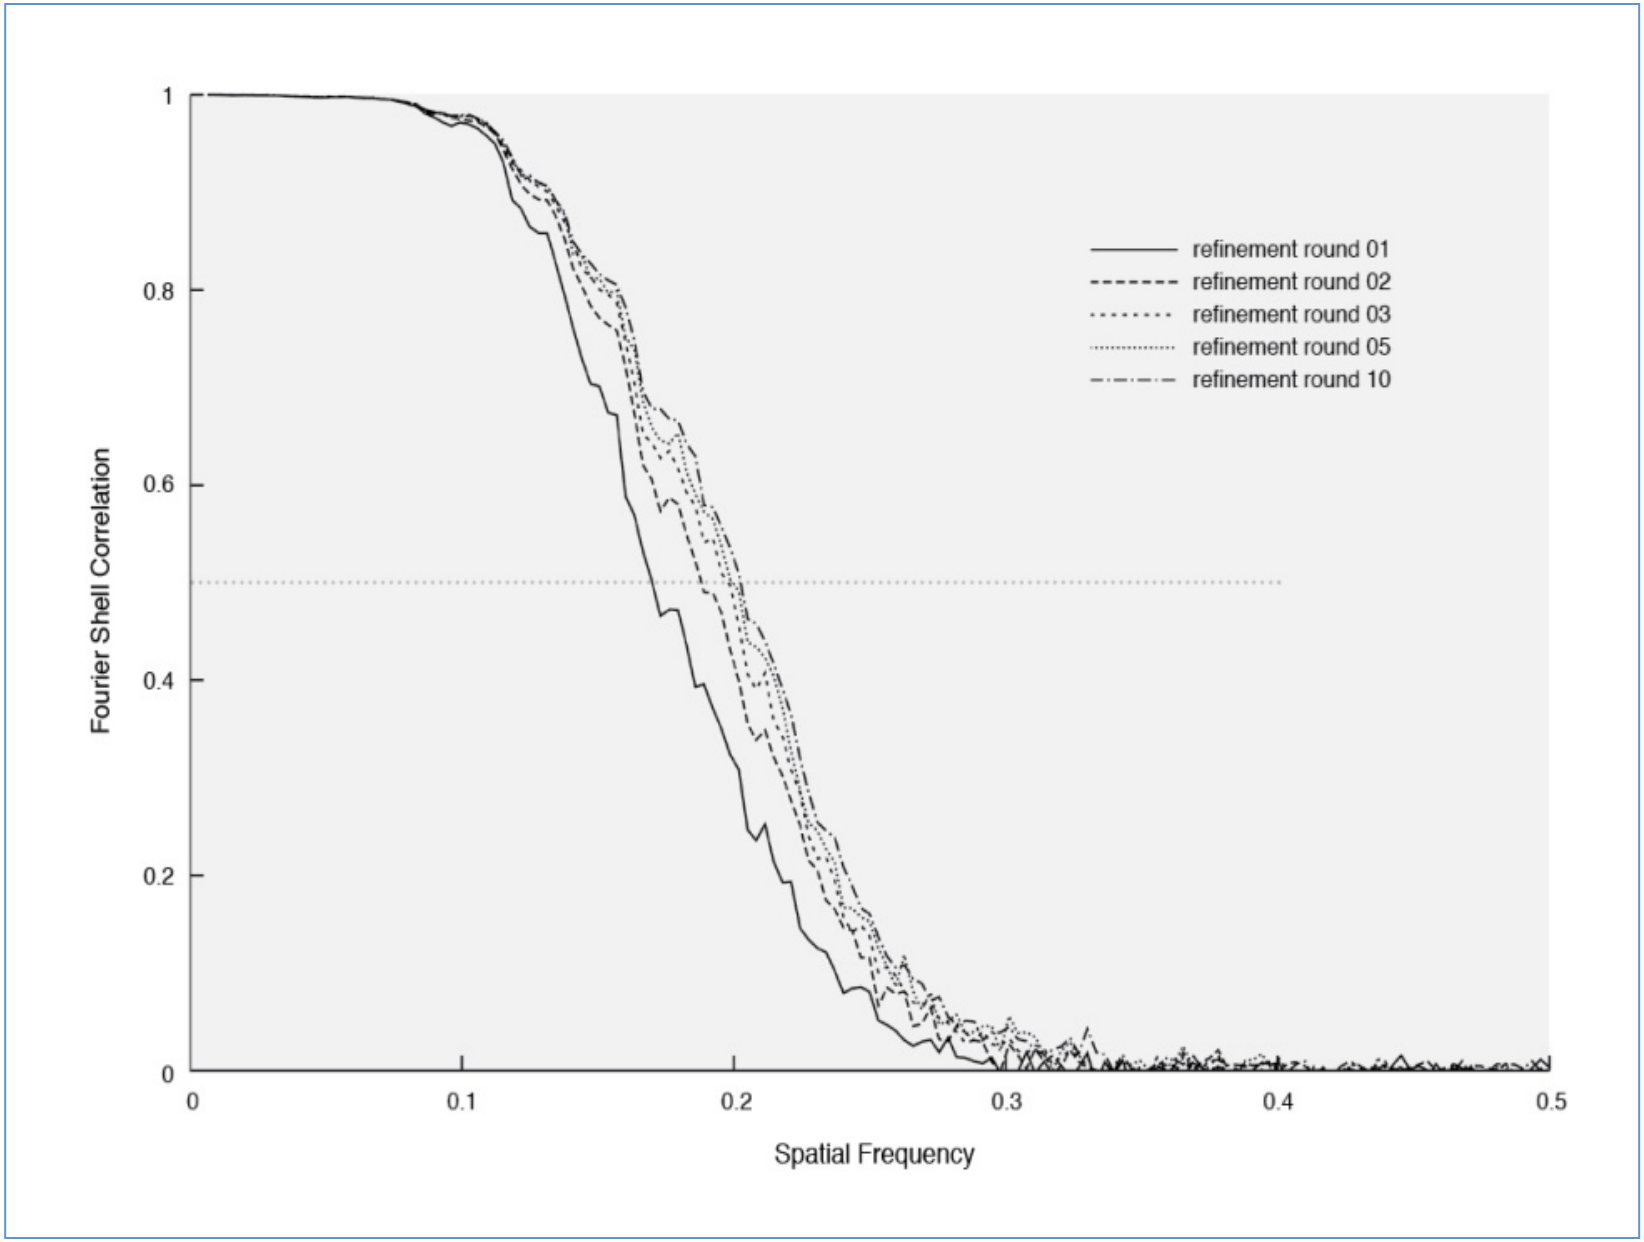

FSC following progress of refinement

3D Unsupervised Classification

Statistical model:
each image is a projection of one of *K* underlying 3D objects, *k.*

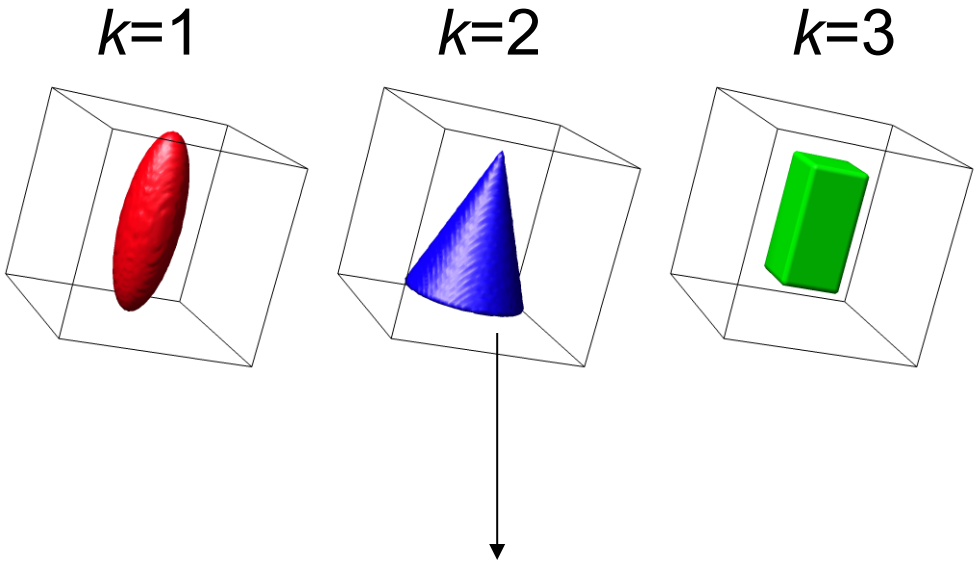

with addition of
white Gaussian noise

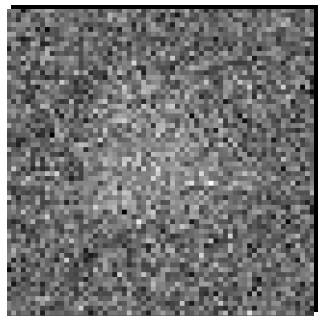

Unknowns: class numbers *k*, rotations, translations

Statistical model: the probability that X_j is observed at pixel j, given the data model A_j, has Gaussian distribution centered on A_j, with halfwidth σ

data

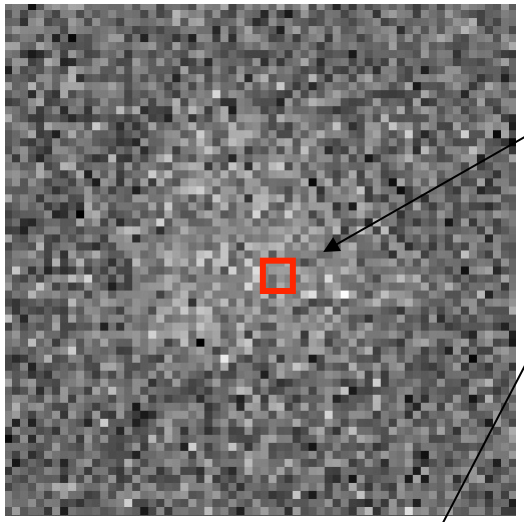

model

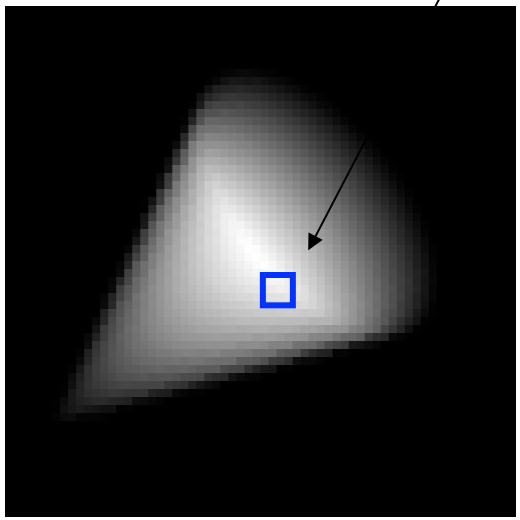

For each pixel j:

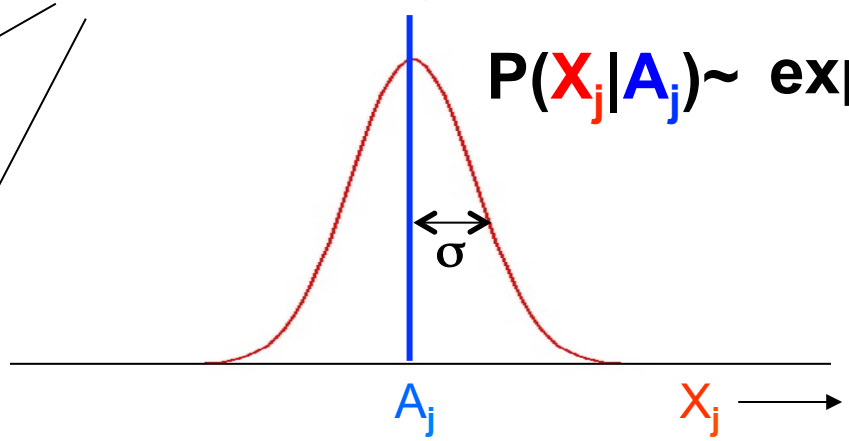

$$P(X_j|A_j) \sim \exp\left(\frac{(X_j - A_j)^2}{-2\sigma^2}\right)$$

white noise =
independence between pixels!

P(data image|model image) ~

$$\prod_j P(X_j|A_j)$$

Likelihood

- Find a model Θ that optimizes the log-likelihood of observing the entire dataset:

class

$$L(\Theta) = \sum_{i=1}^{N} \ln \sum_{k=1}^{K} \sum_{rot} \sum_{trans} P(image_i \mid k, rot, trans, \Theta) P(k, rot, trans \mid \Theta)$$

integrate over all unknowns!

The model Θ comprises: estimates for 3D objects, σ, …

Optimization algorithm: Expectation Maximization

ML3D

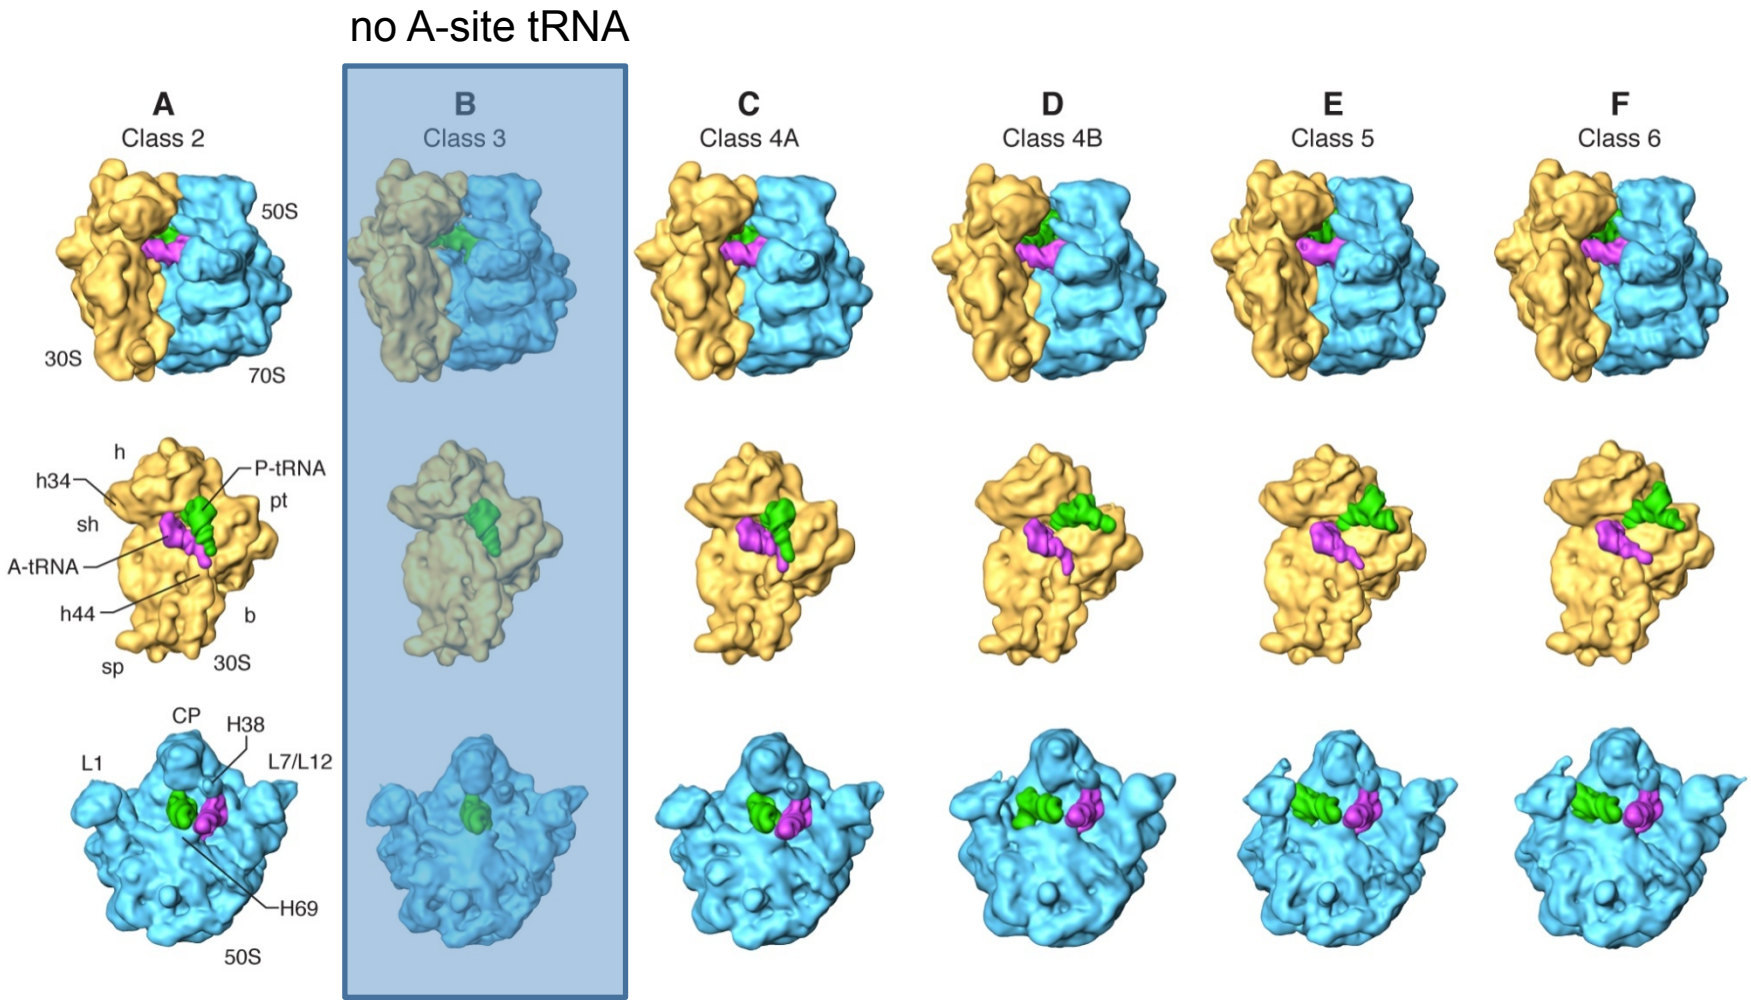

Pre-translocational states of wt 70S *E. coli* ribosome

Agirrezabala et al., PNAS 2012

QUESTIONS?

Generalized Euclidean distance

Euclidean distance between two images f_1 and f_2:

$$E^2_{12}(\alpha, r') = \sum_{j=1}^{J} | f_1(r_j) - f_2(R_\alpha r_j + r') |^2$$

$$= \sum_{j=1}^{J} | f_1(r_j) |^2 + \sum_{j=1}^{J} | f_2(R_\alpha r_j + r') |^2 - \sum_{j=1}^{J} | f_1(r_j) f_2(R_\alpha r_j + r') |$$

const. const. cross-correlation

Similarity = closeness in high-dimensional Euclidean space = small E-distance
→ large value of CCF peak at matching position

It's easily verified that the Fourier transform of any real-valued image has the following property:

$$F(k_x,k_y) = F^*(-k_x,-k_y) \qquad \text{(Friedel's Law)}$$

Examples for Fourier transforms of simple functions:

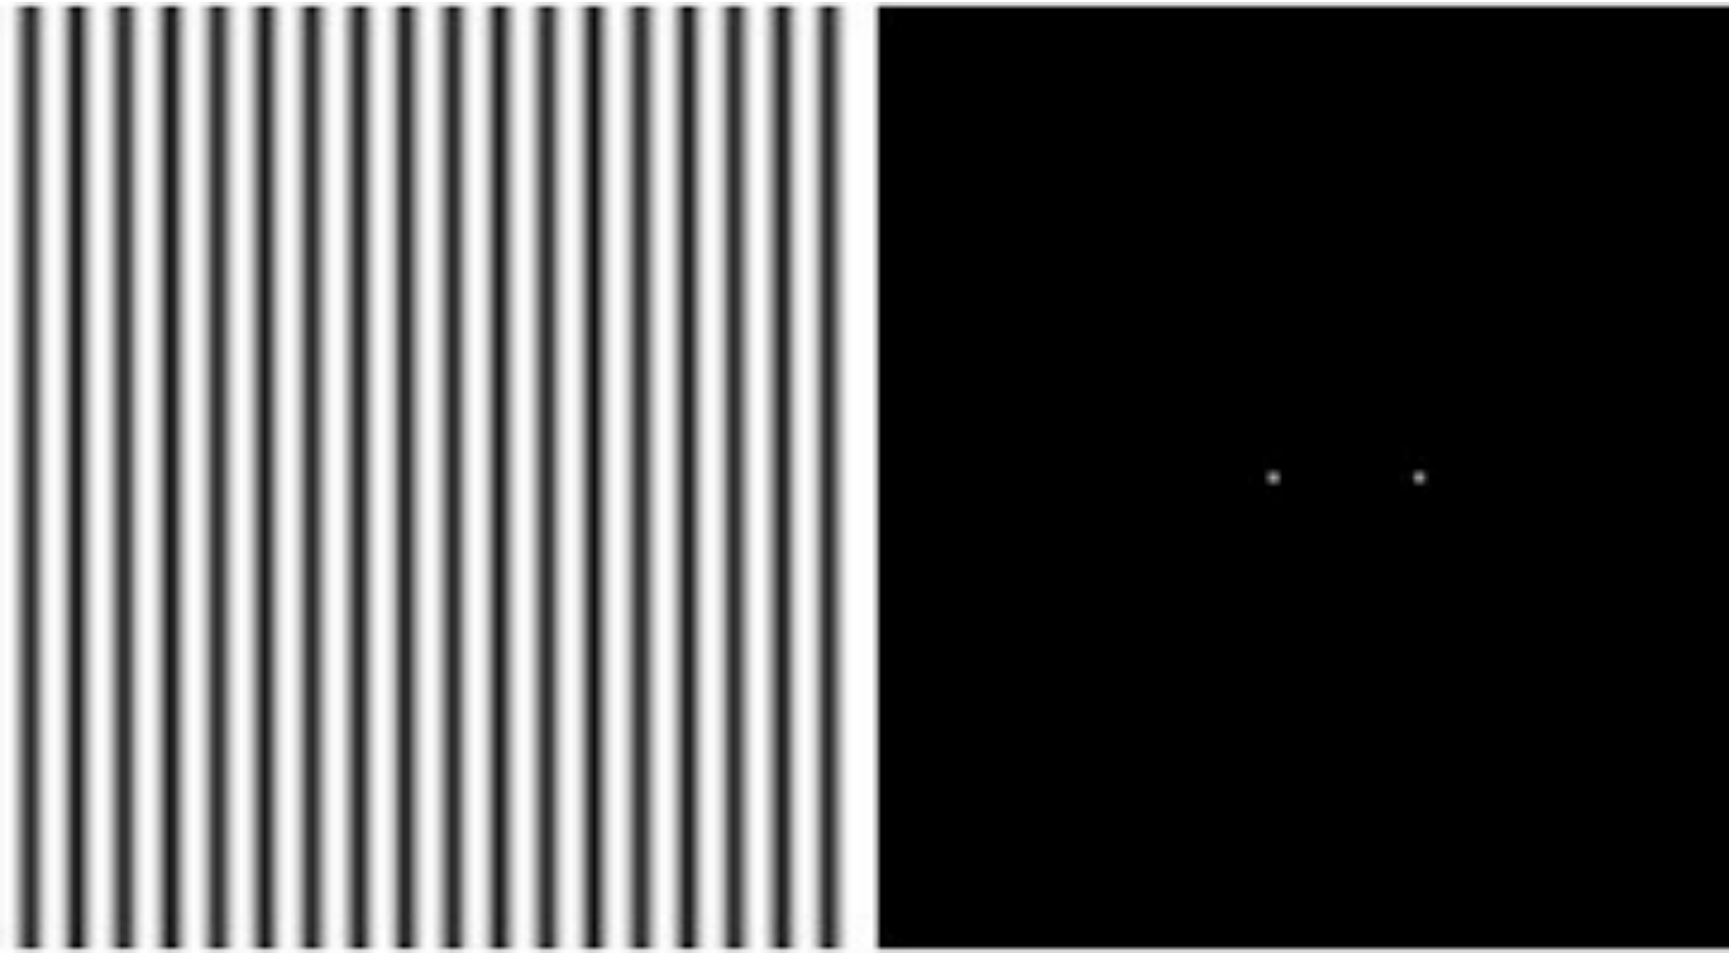

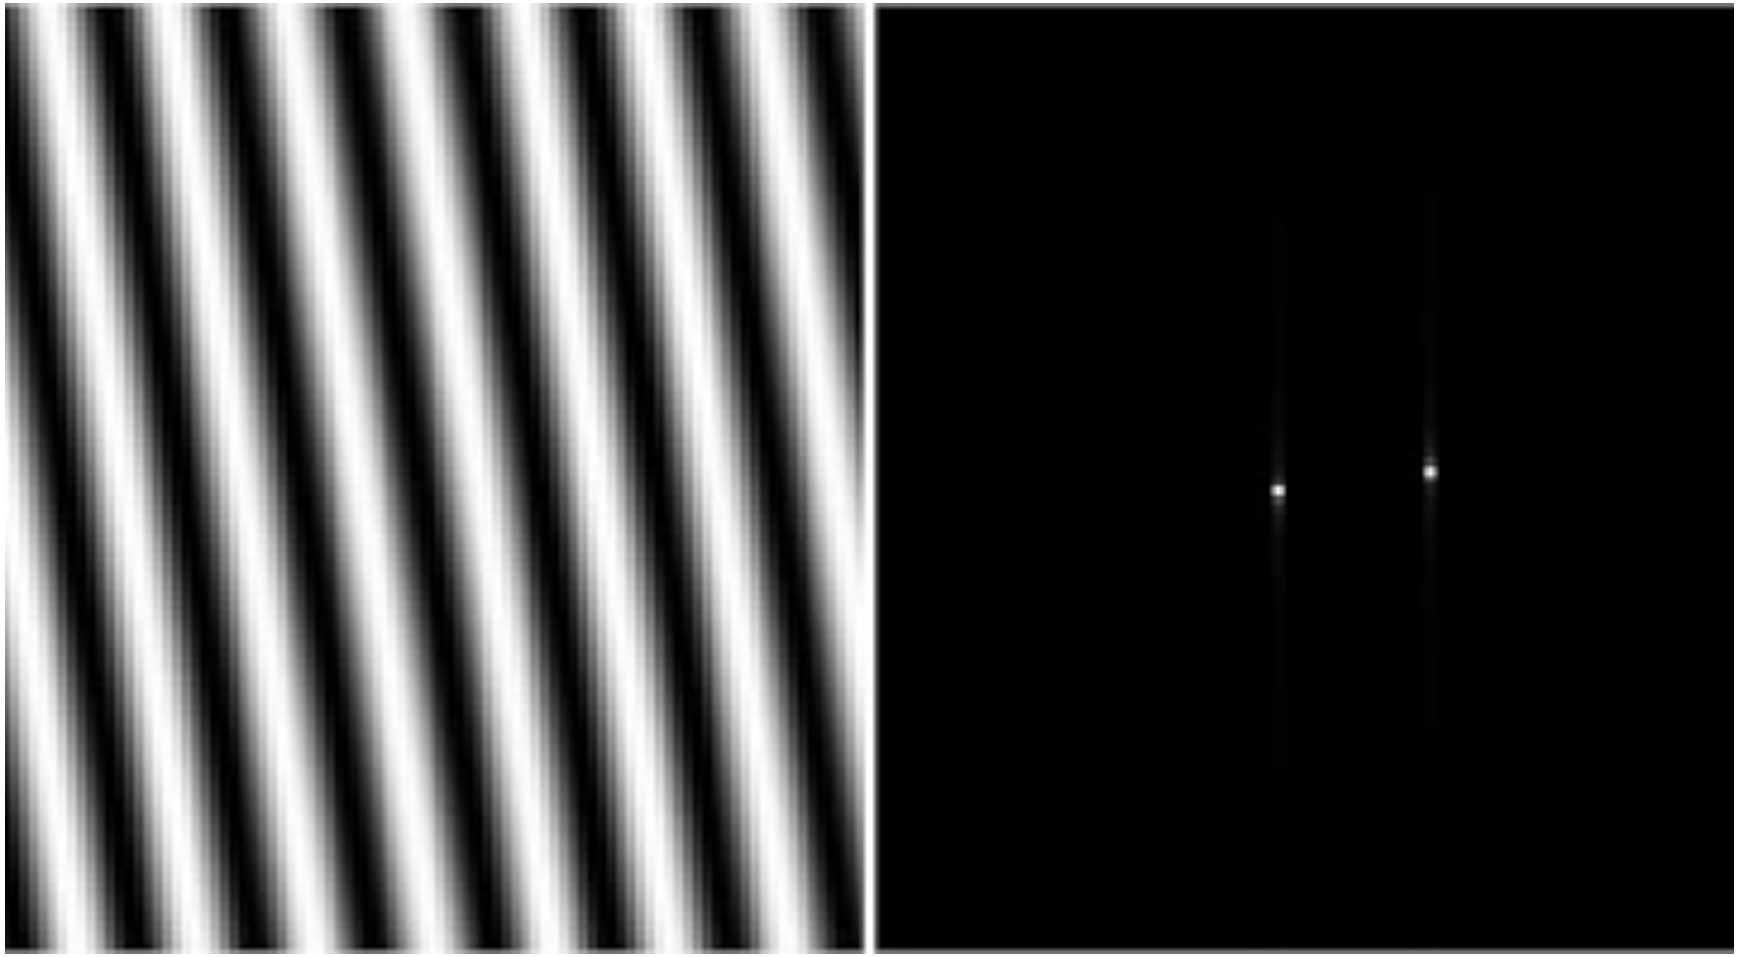

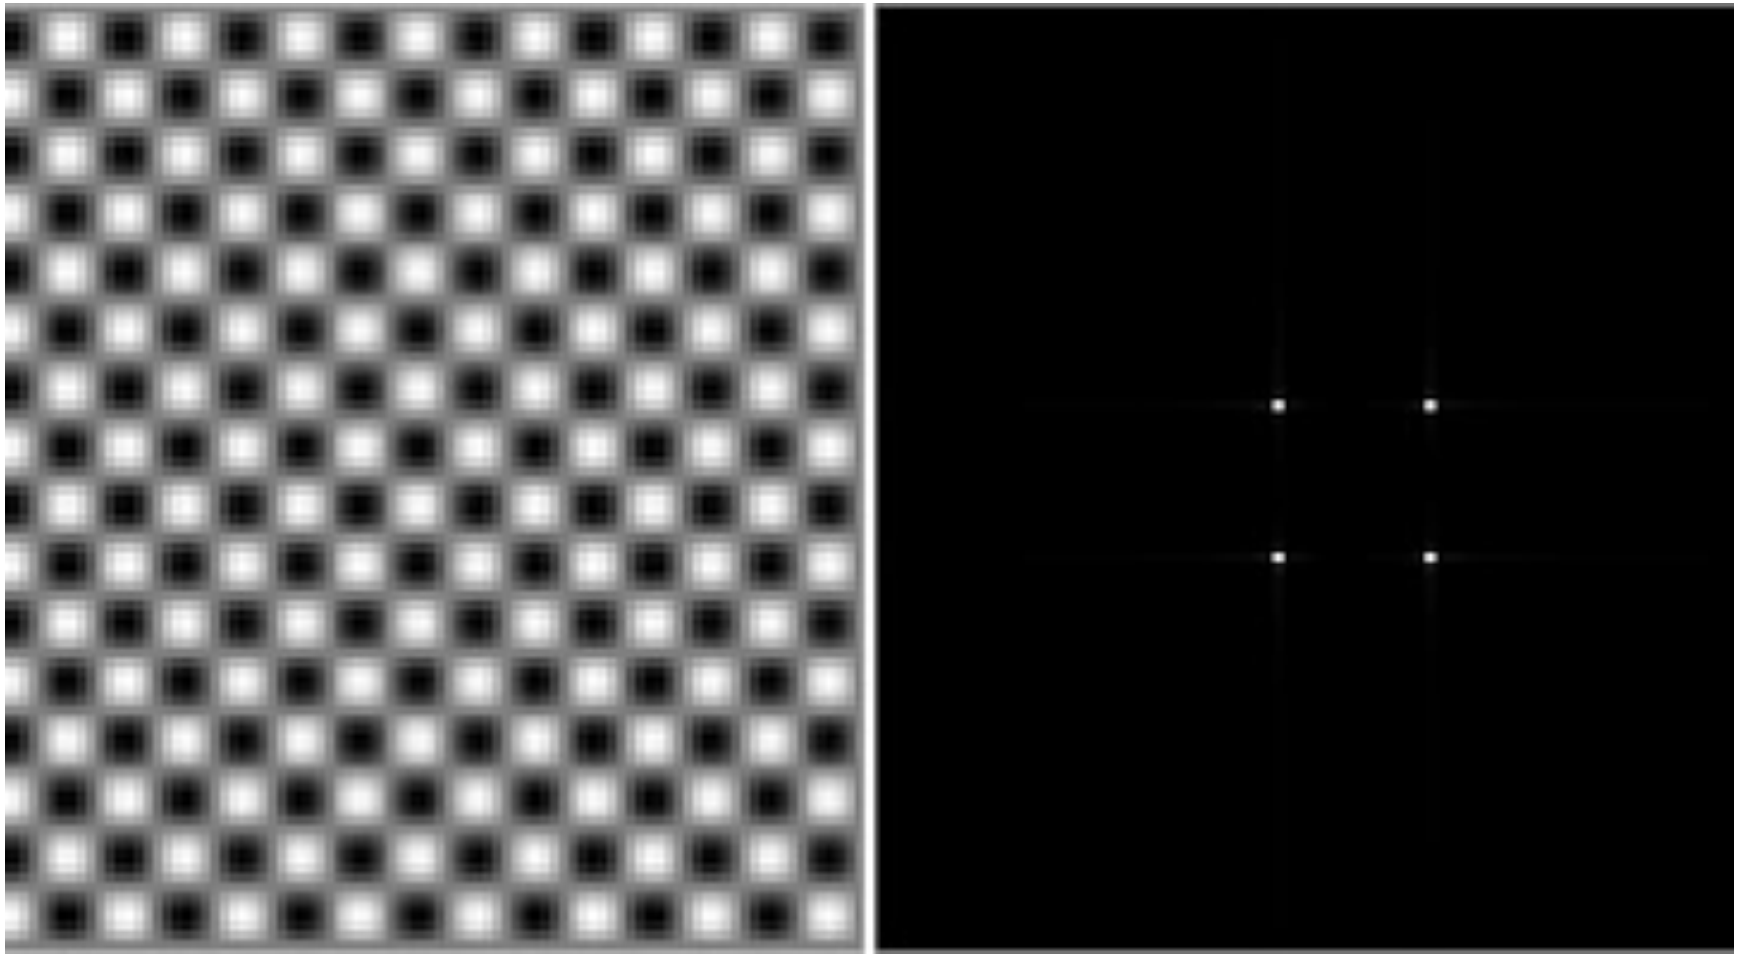

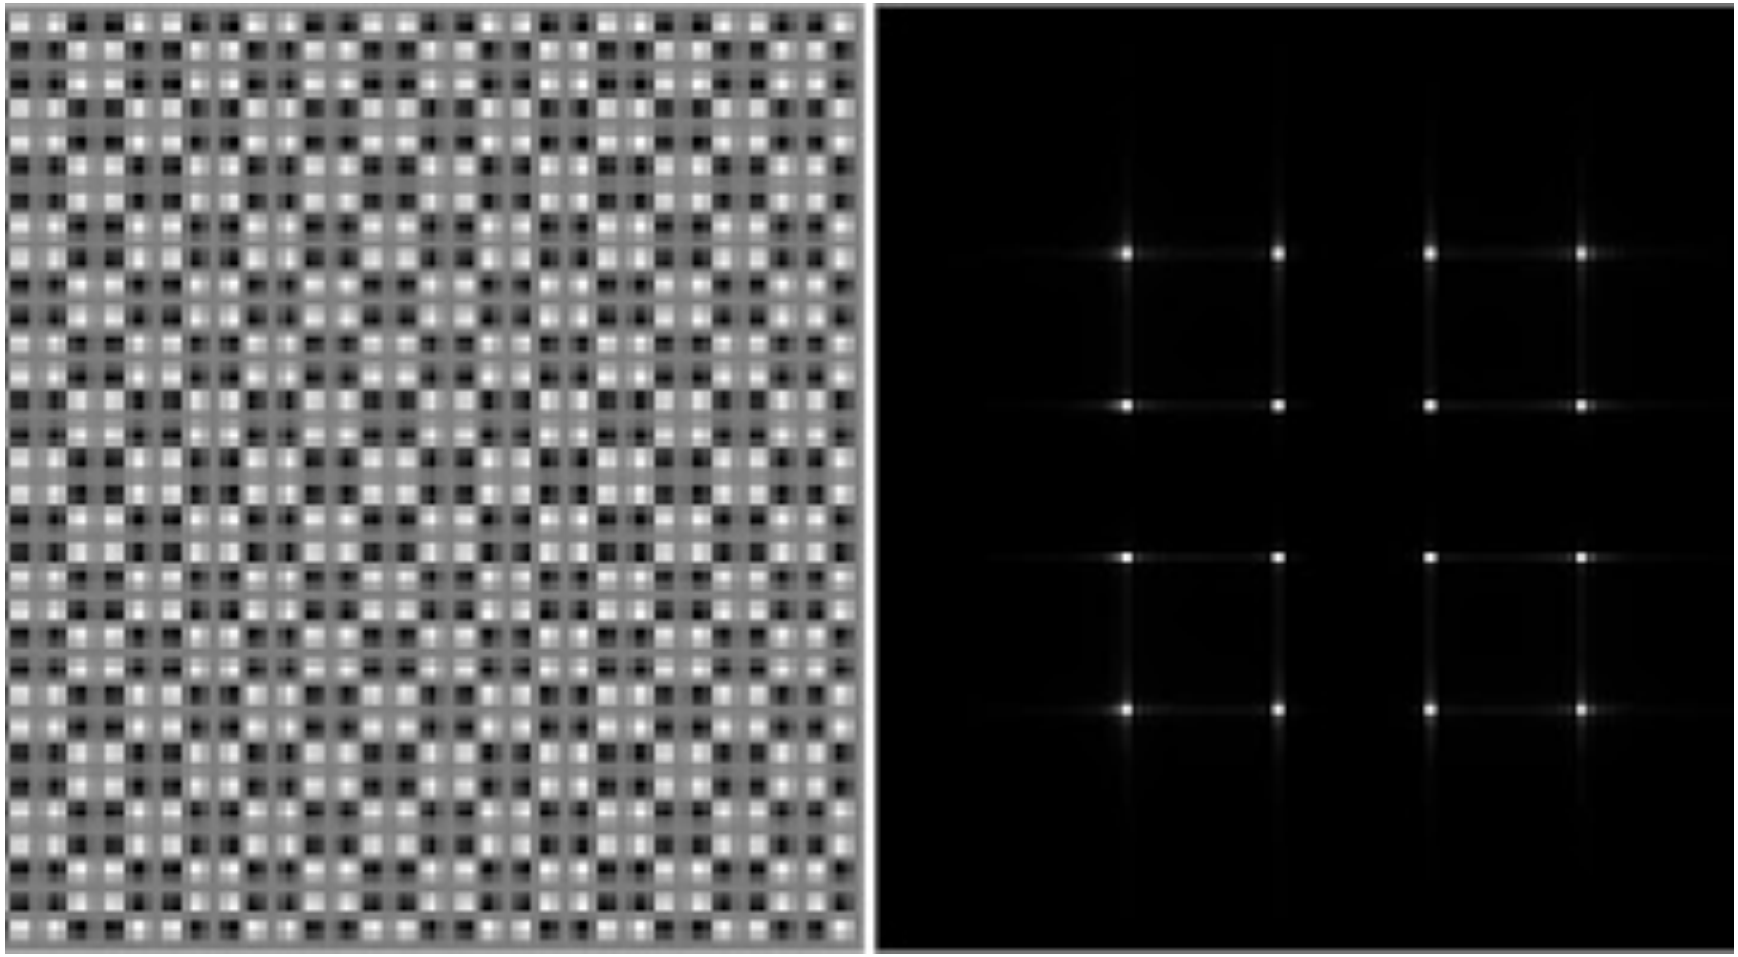

Units of spatial frequency in 2D Fourier space

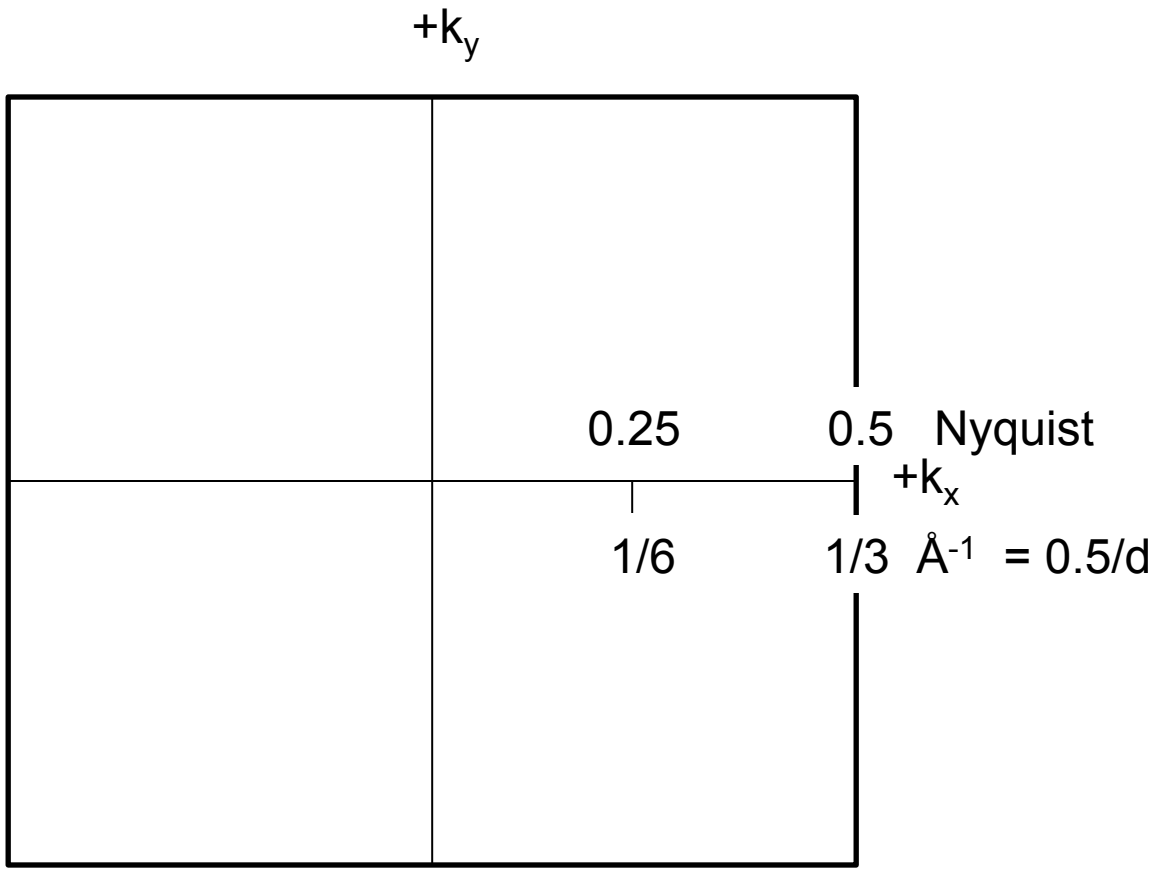

Spatial frequency is either in Nyquist units (0… 0.5) or in physical units 0.5/d relating to the sampling step d. In above example, sampling step is d = 1.5 Å