The local low-dimensionality of natural images

Olivier Hénaff

Johannes Ballé

Neil Rabinowitz

Eero Simoncelli

Howard Hughes Medical Institute

New York University


LA CORNOUAILLE

image *filter* *marginal response*

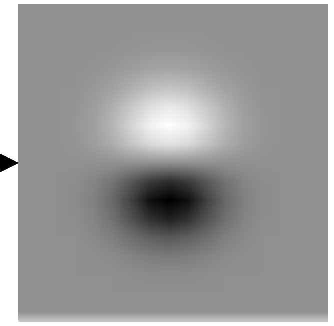

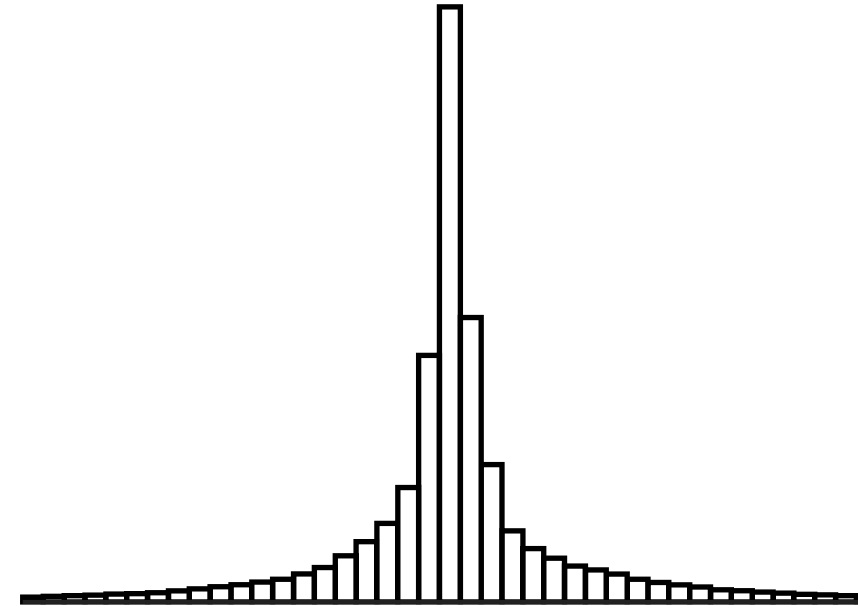

filter response

Burt & Adelson (1982)
Field (1987)

image *filter* *marginal response*

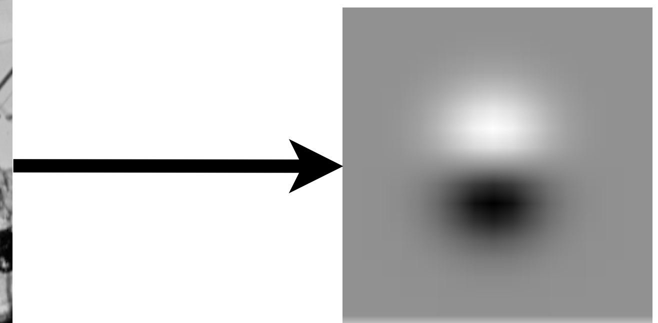

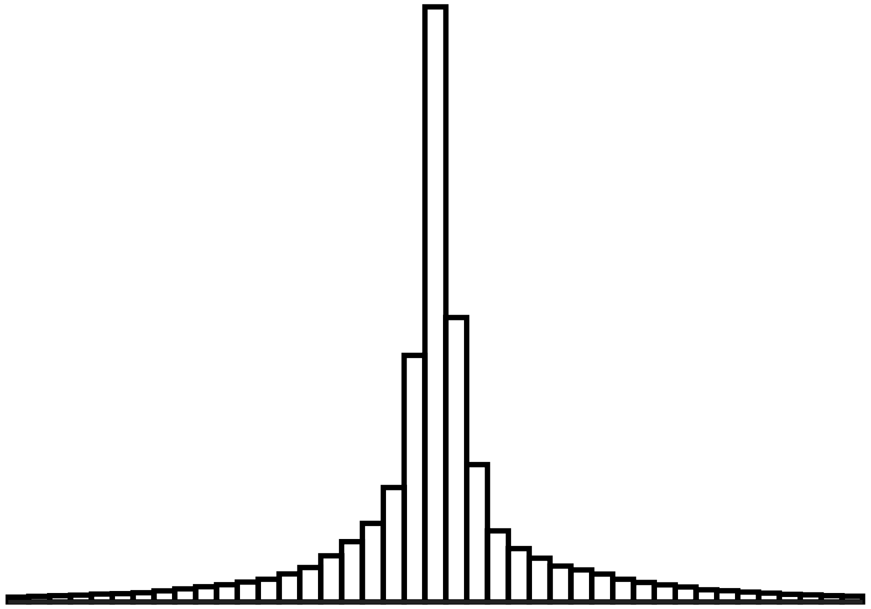

filter response

Burt & Adelson (1982)
Field (1987)
Olshausen & Field (1996)
Bell & Sejnowski (1997)

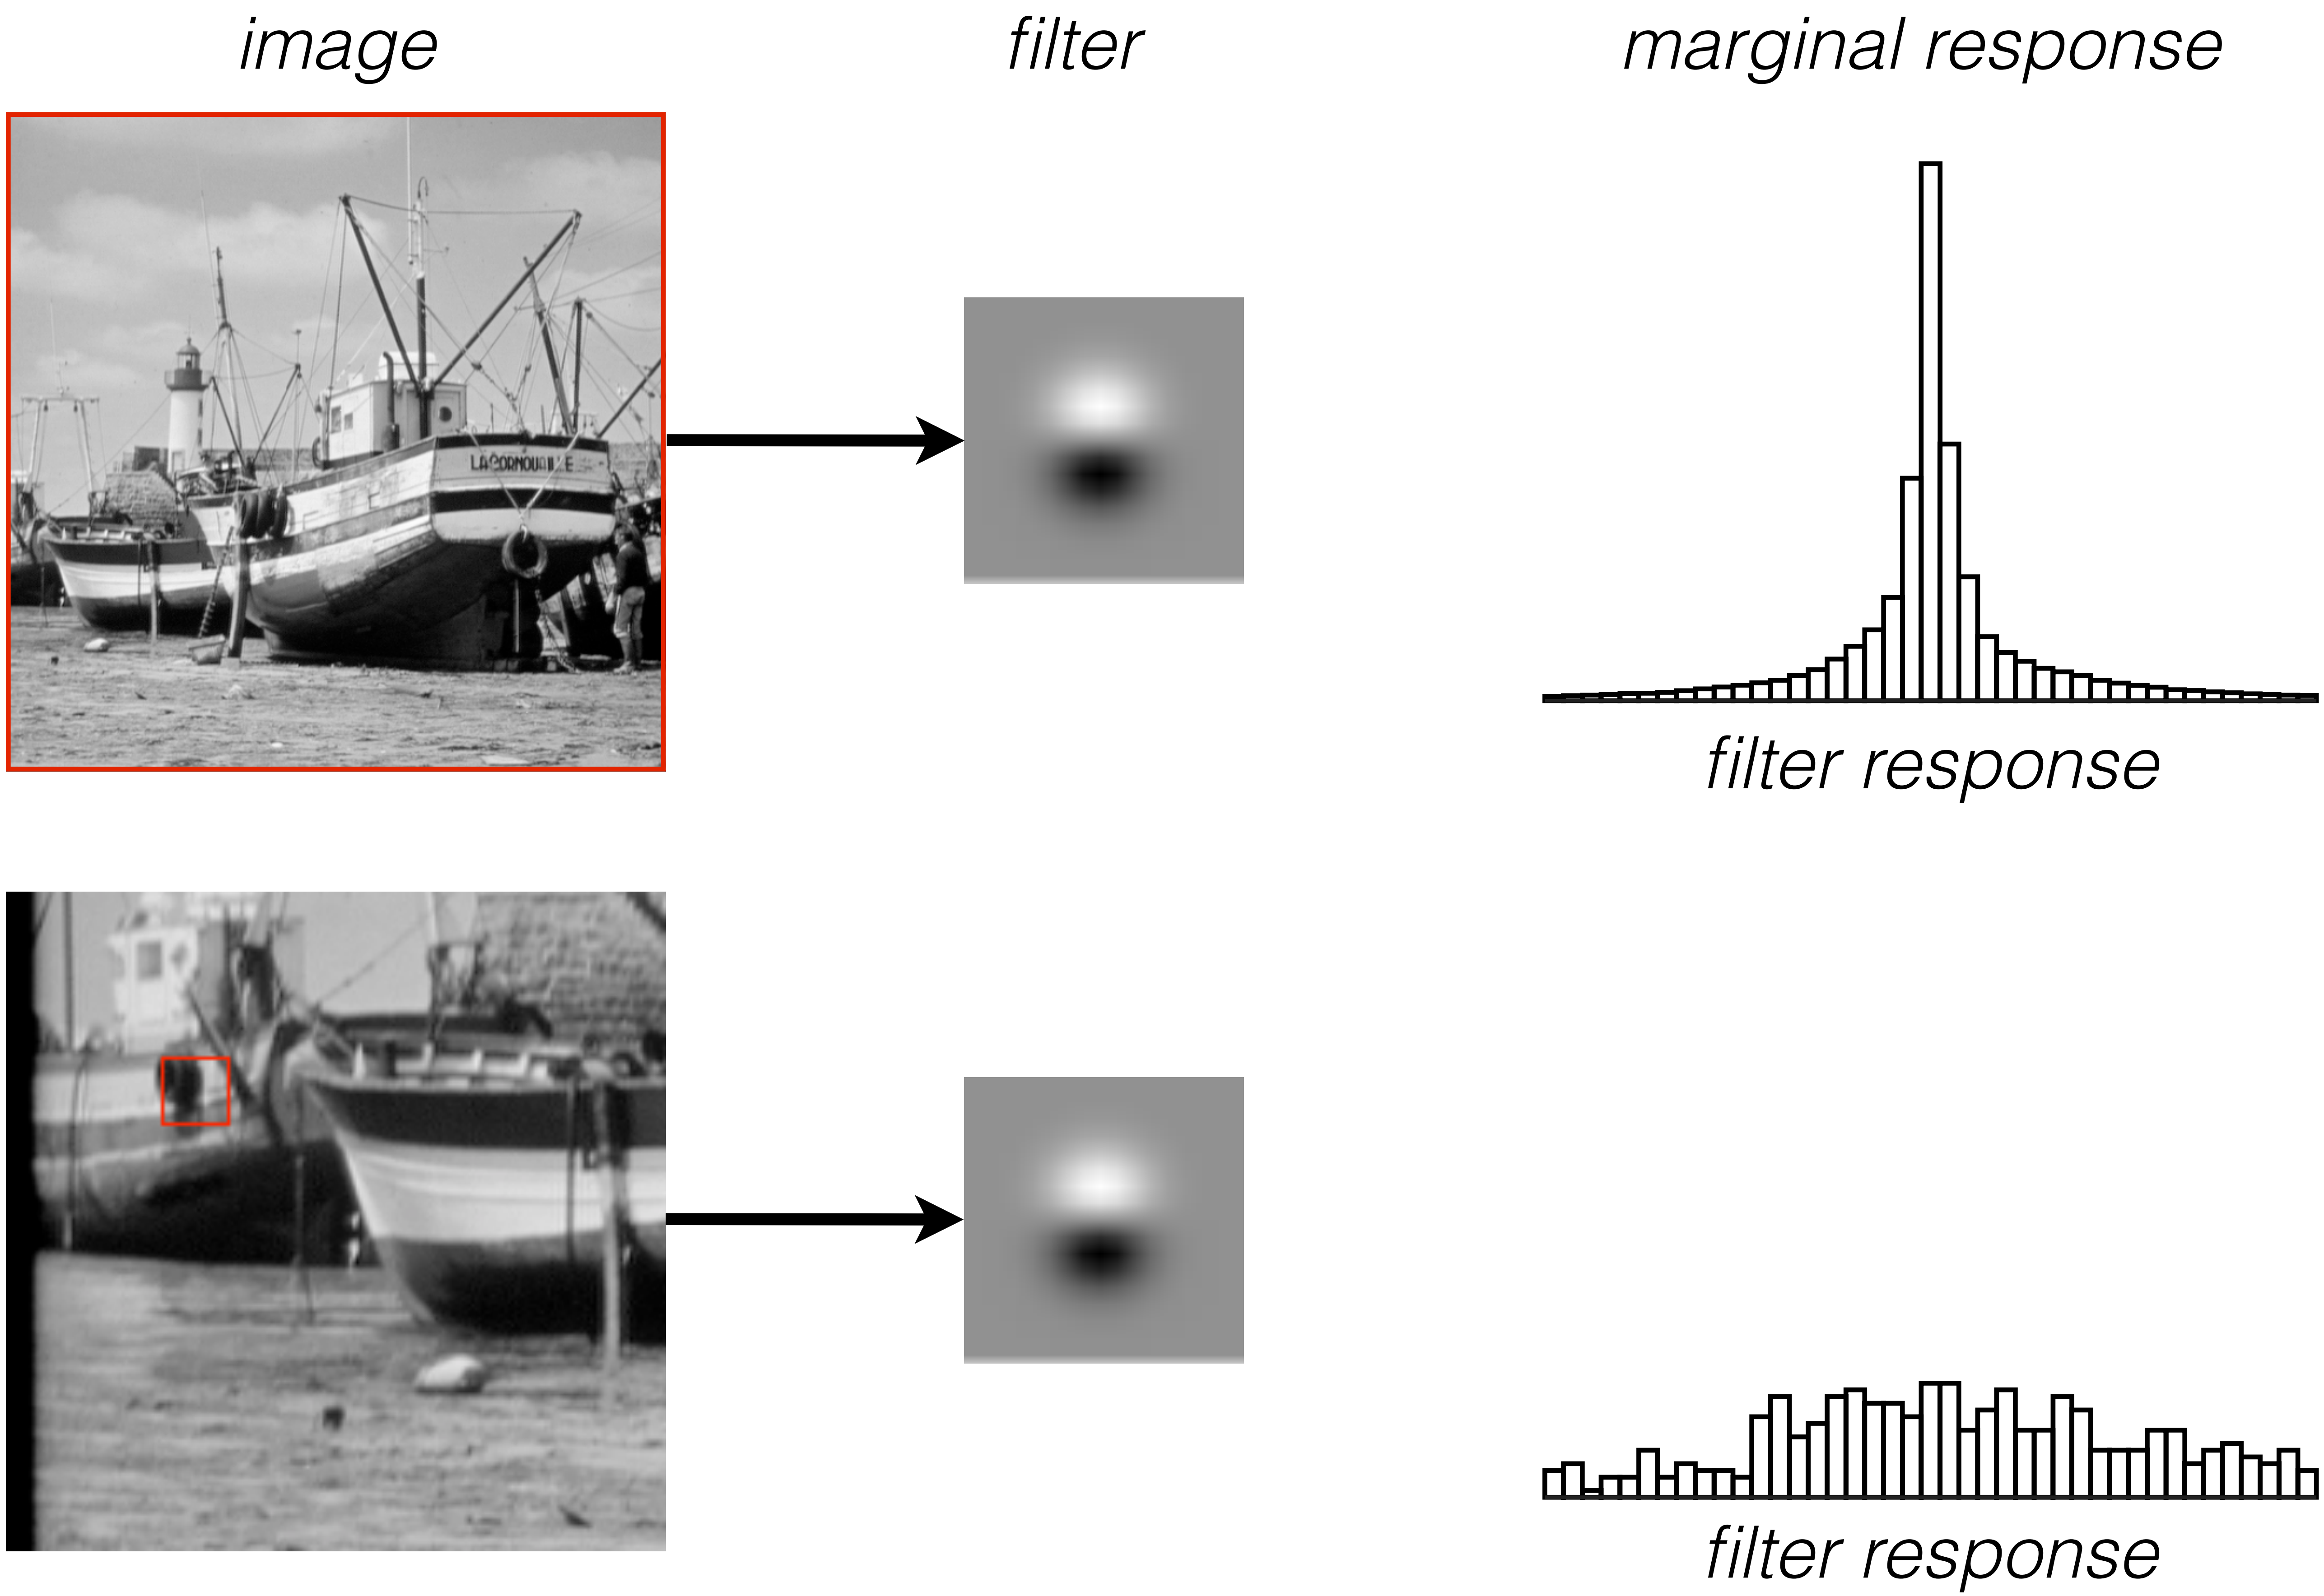
image
filter
marginal response
filter response
filter response

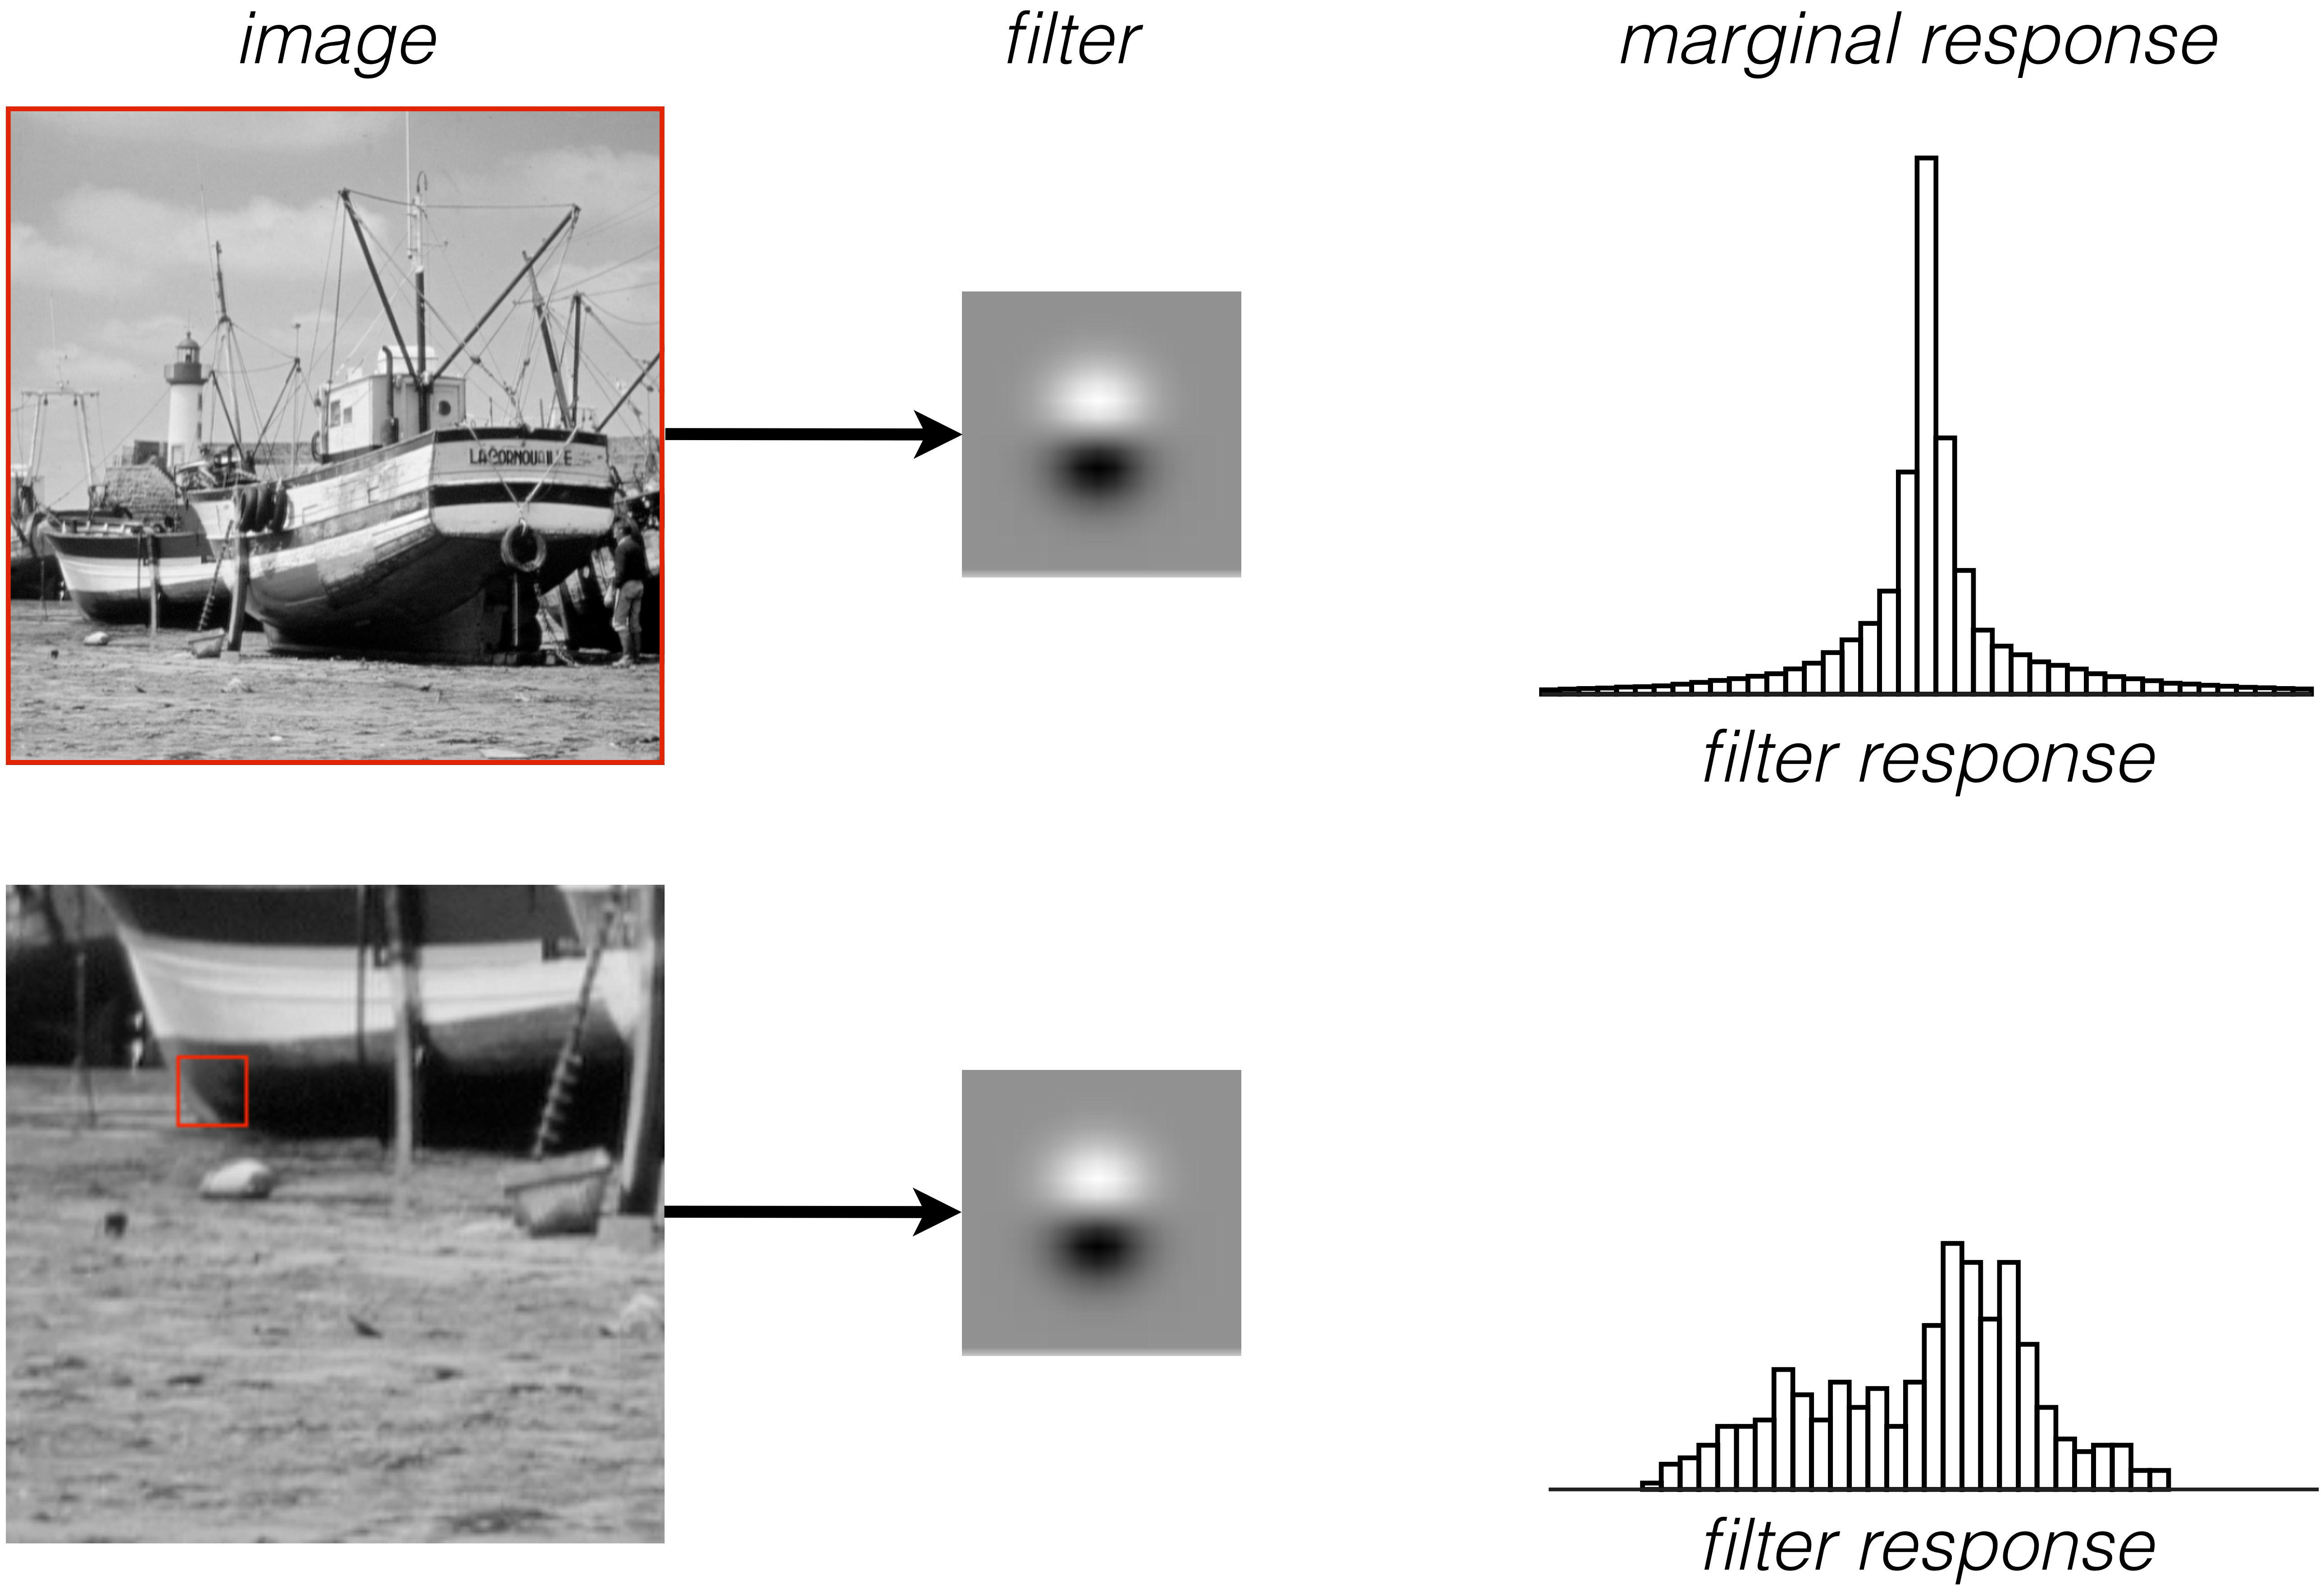
image
filter
marginal response
filter response
filter response

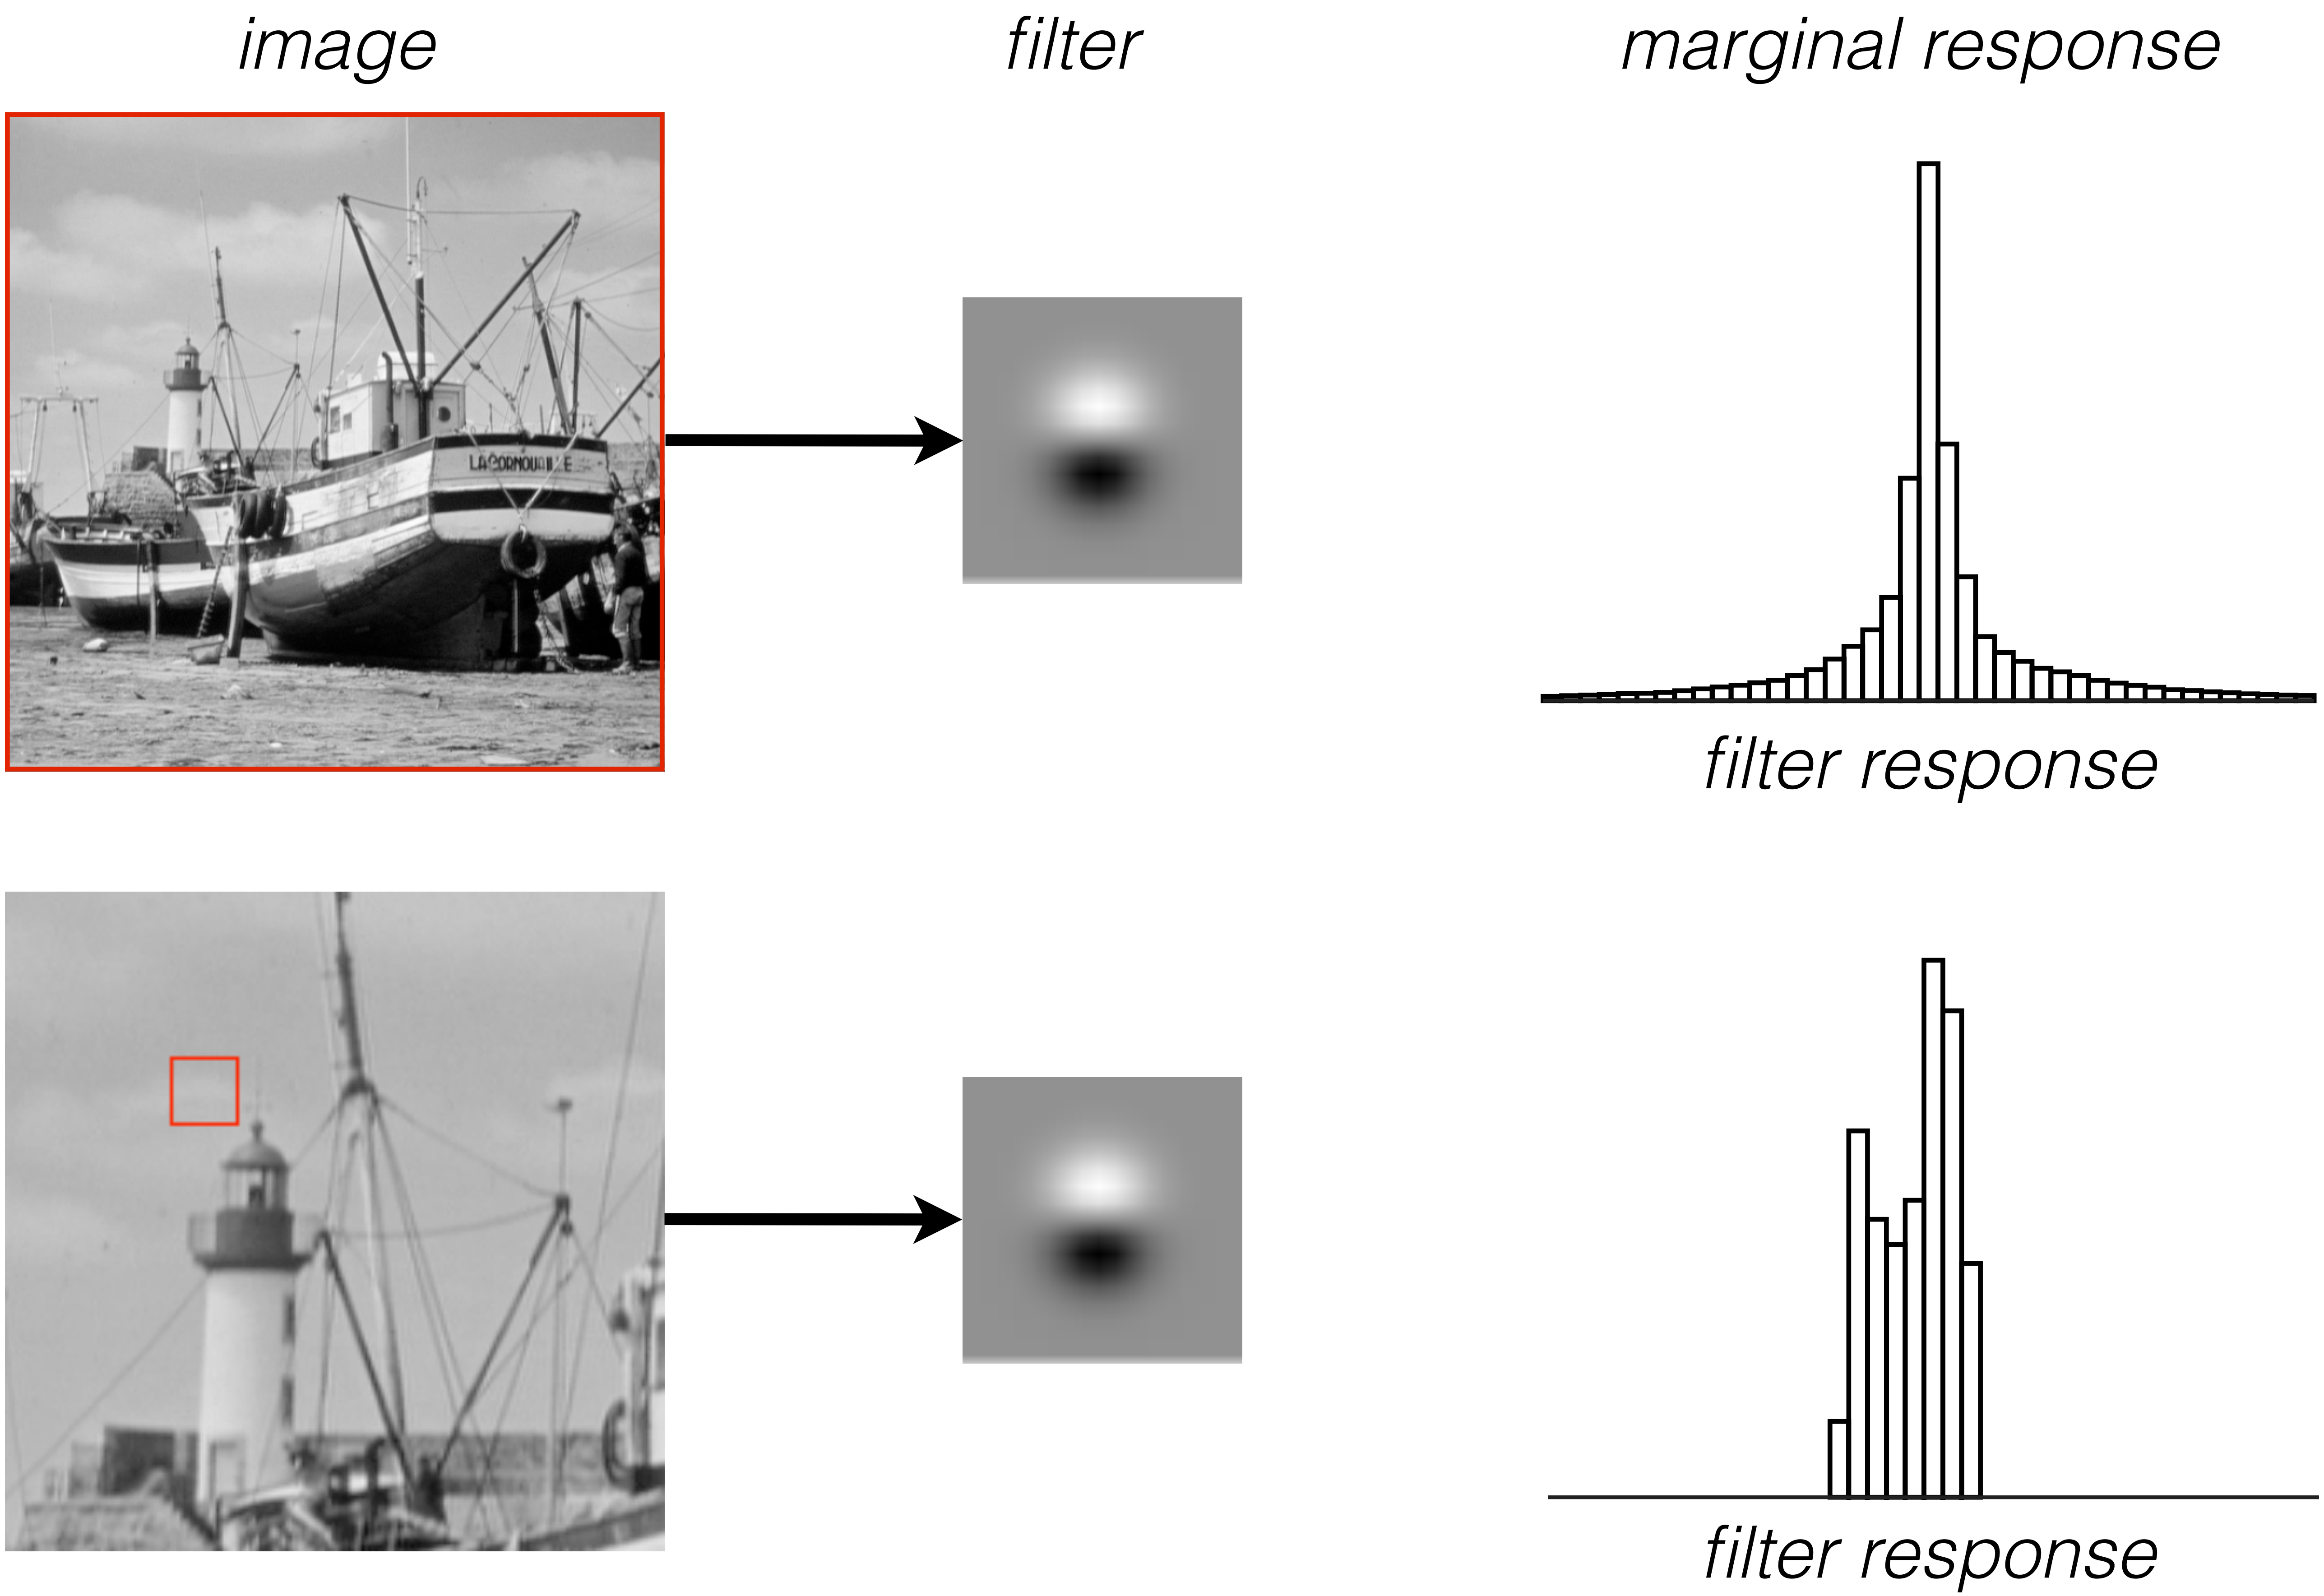

Ruderman (1994)
Wainwright & Simoncelli (2000)

image *filters* *joint responses*

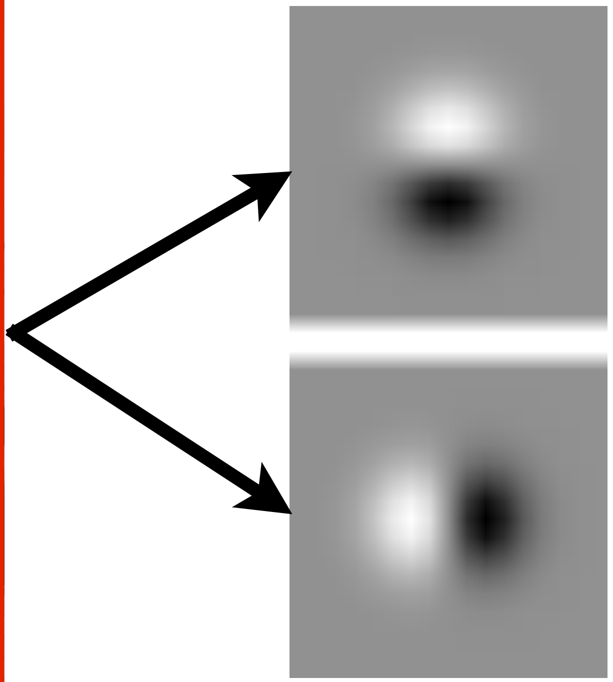

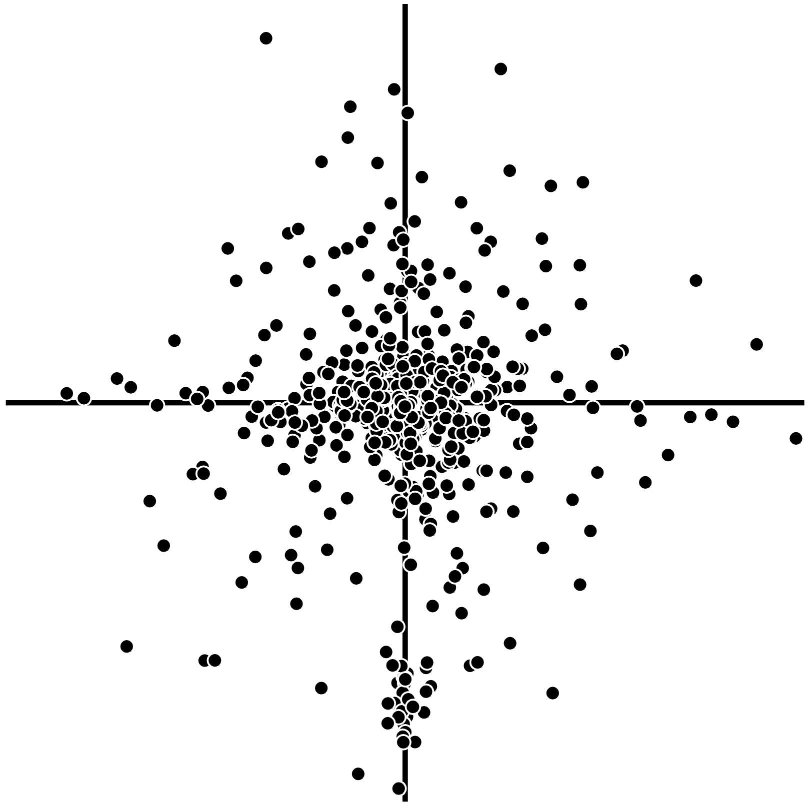

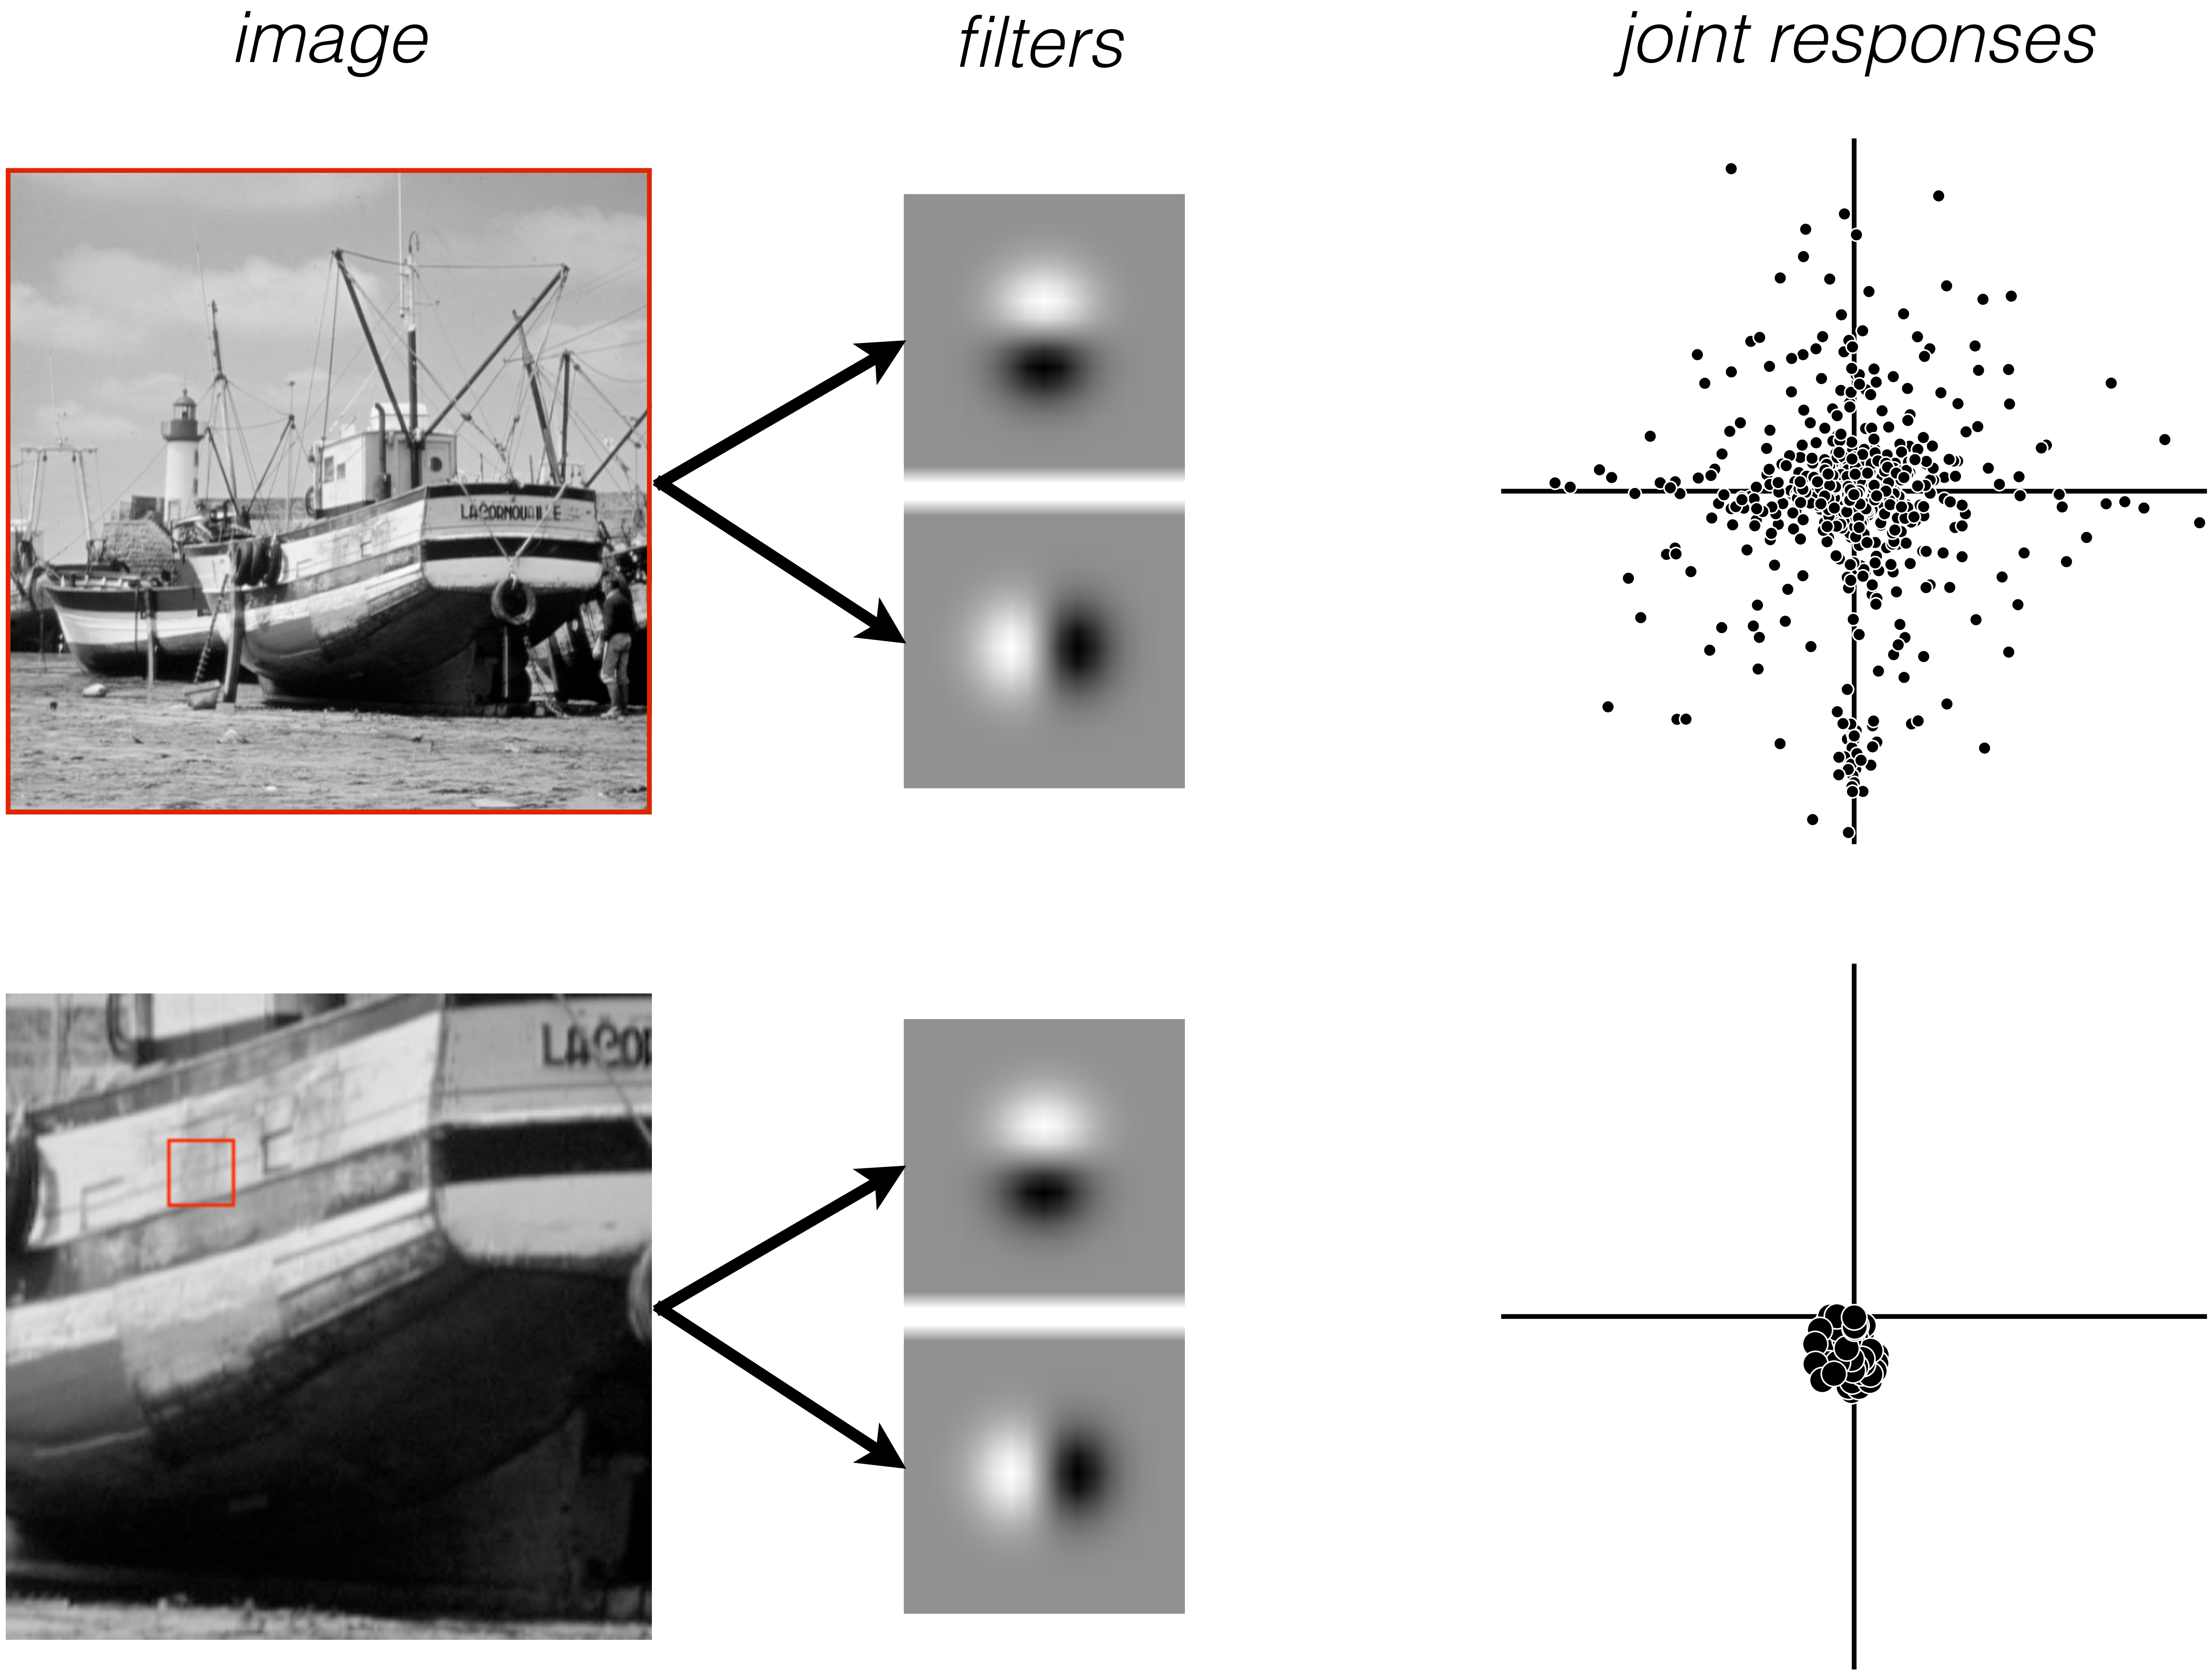
image
filters
joint responses

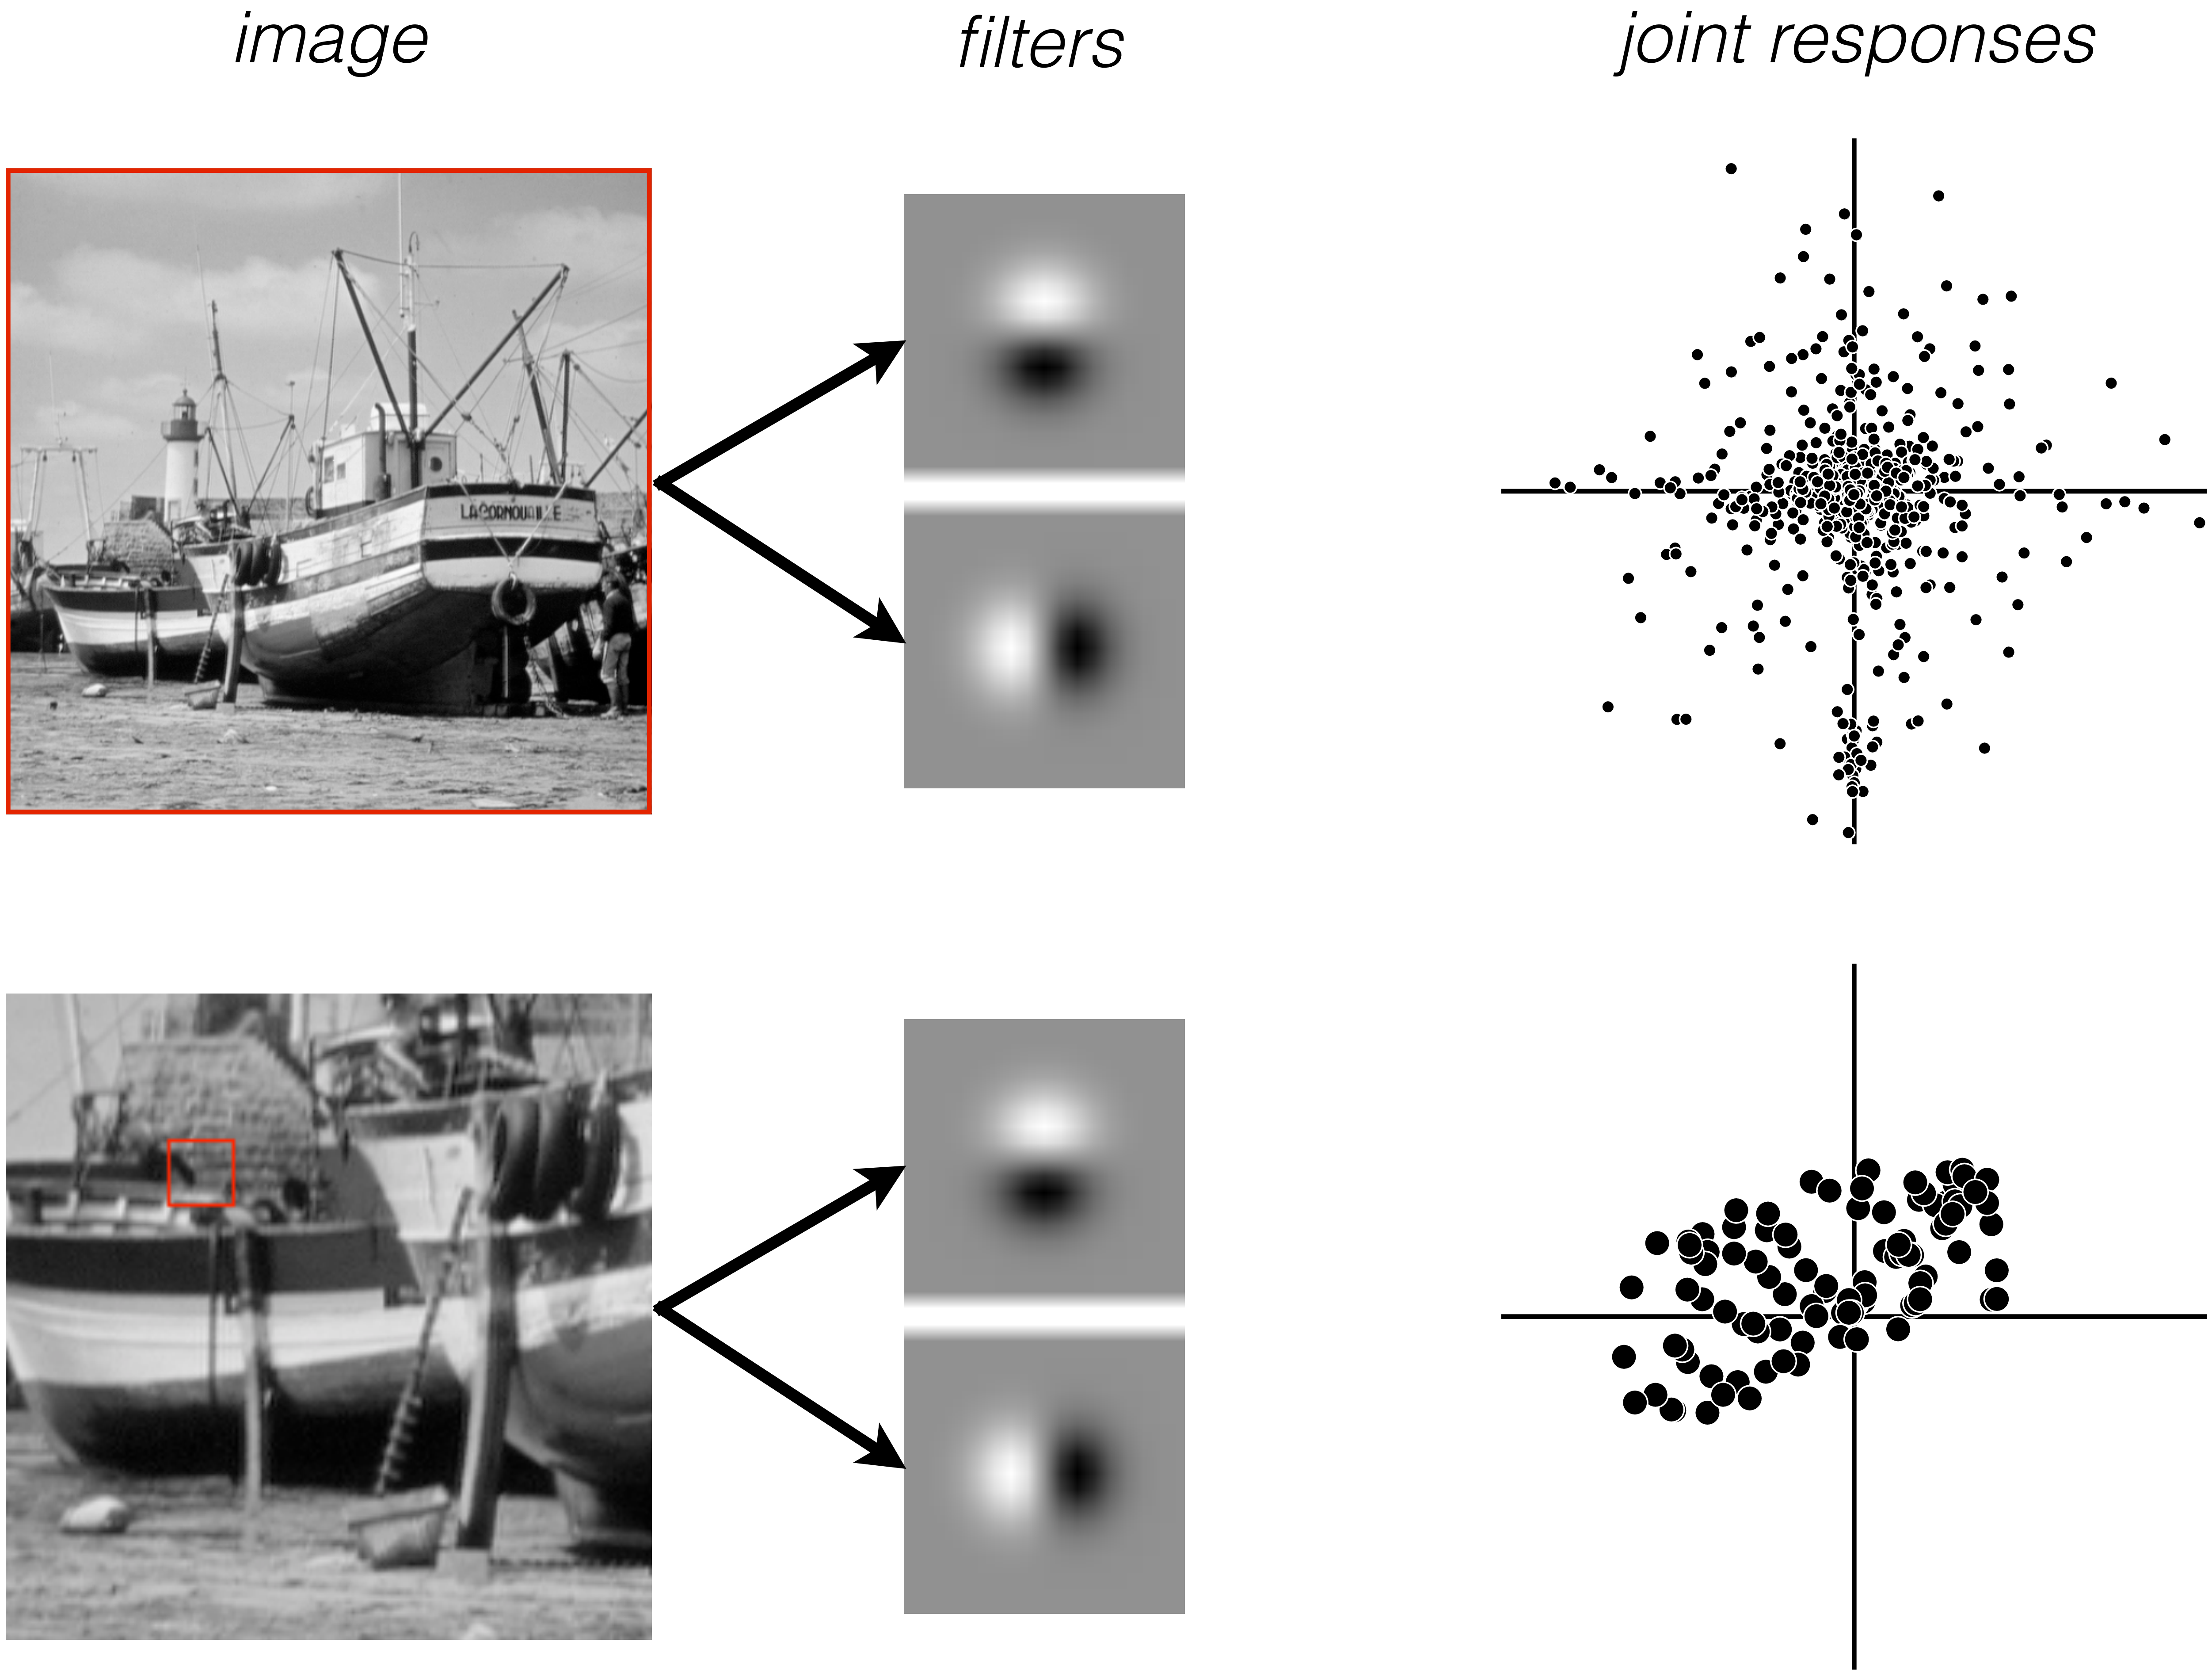
image
filters
joint responses

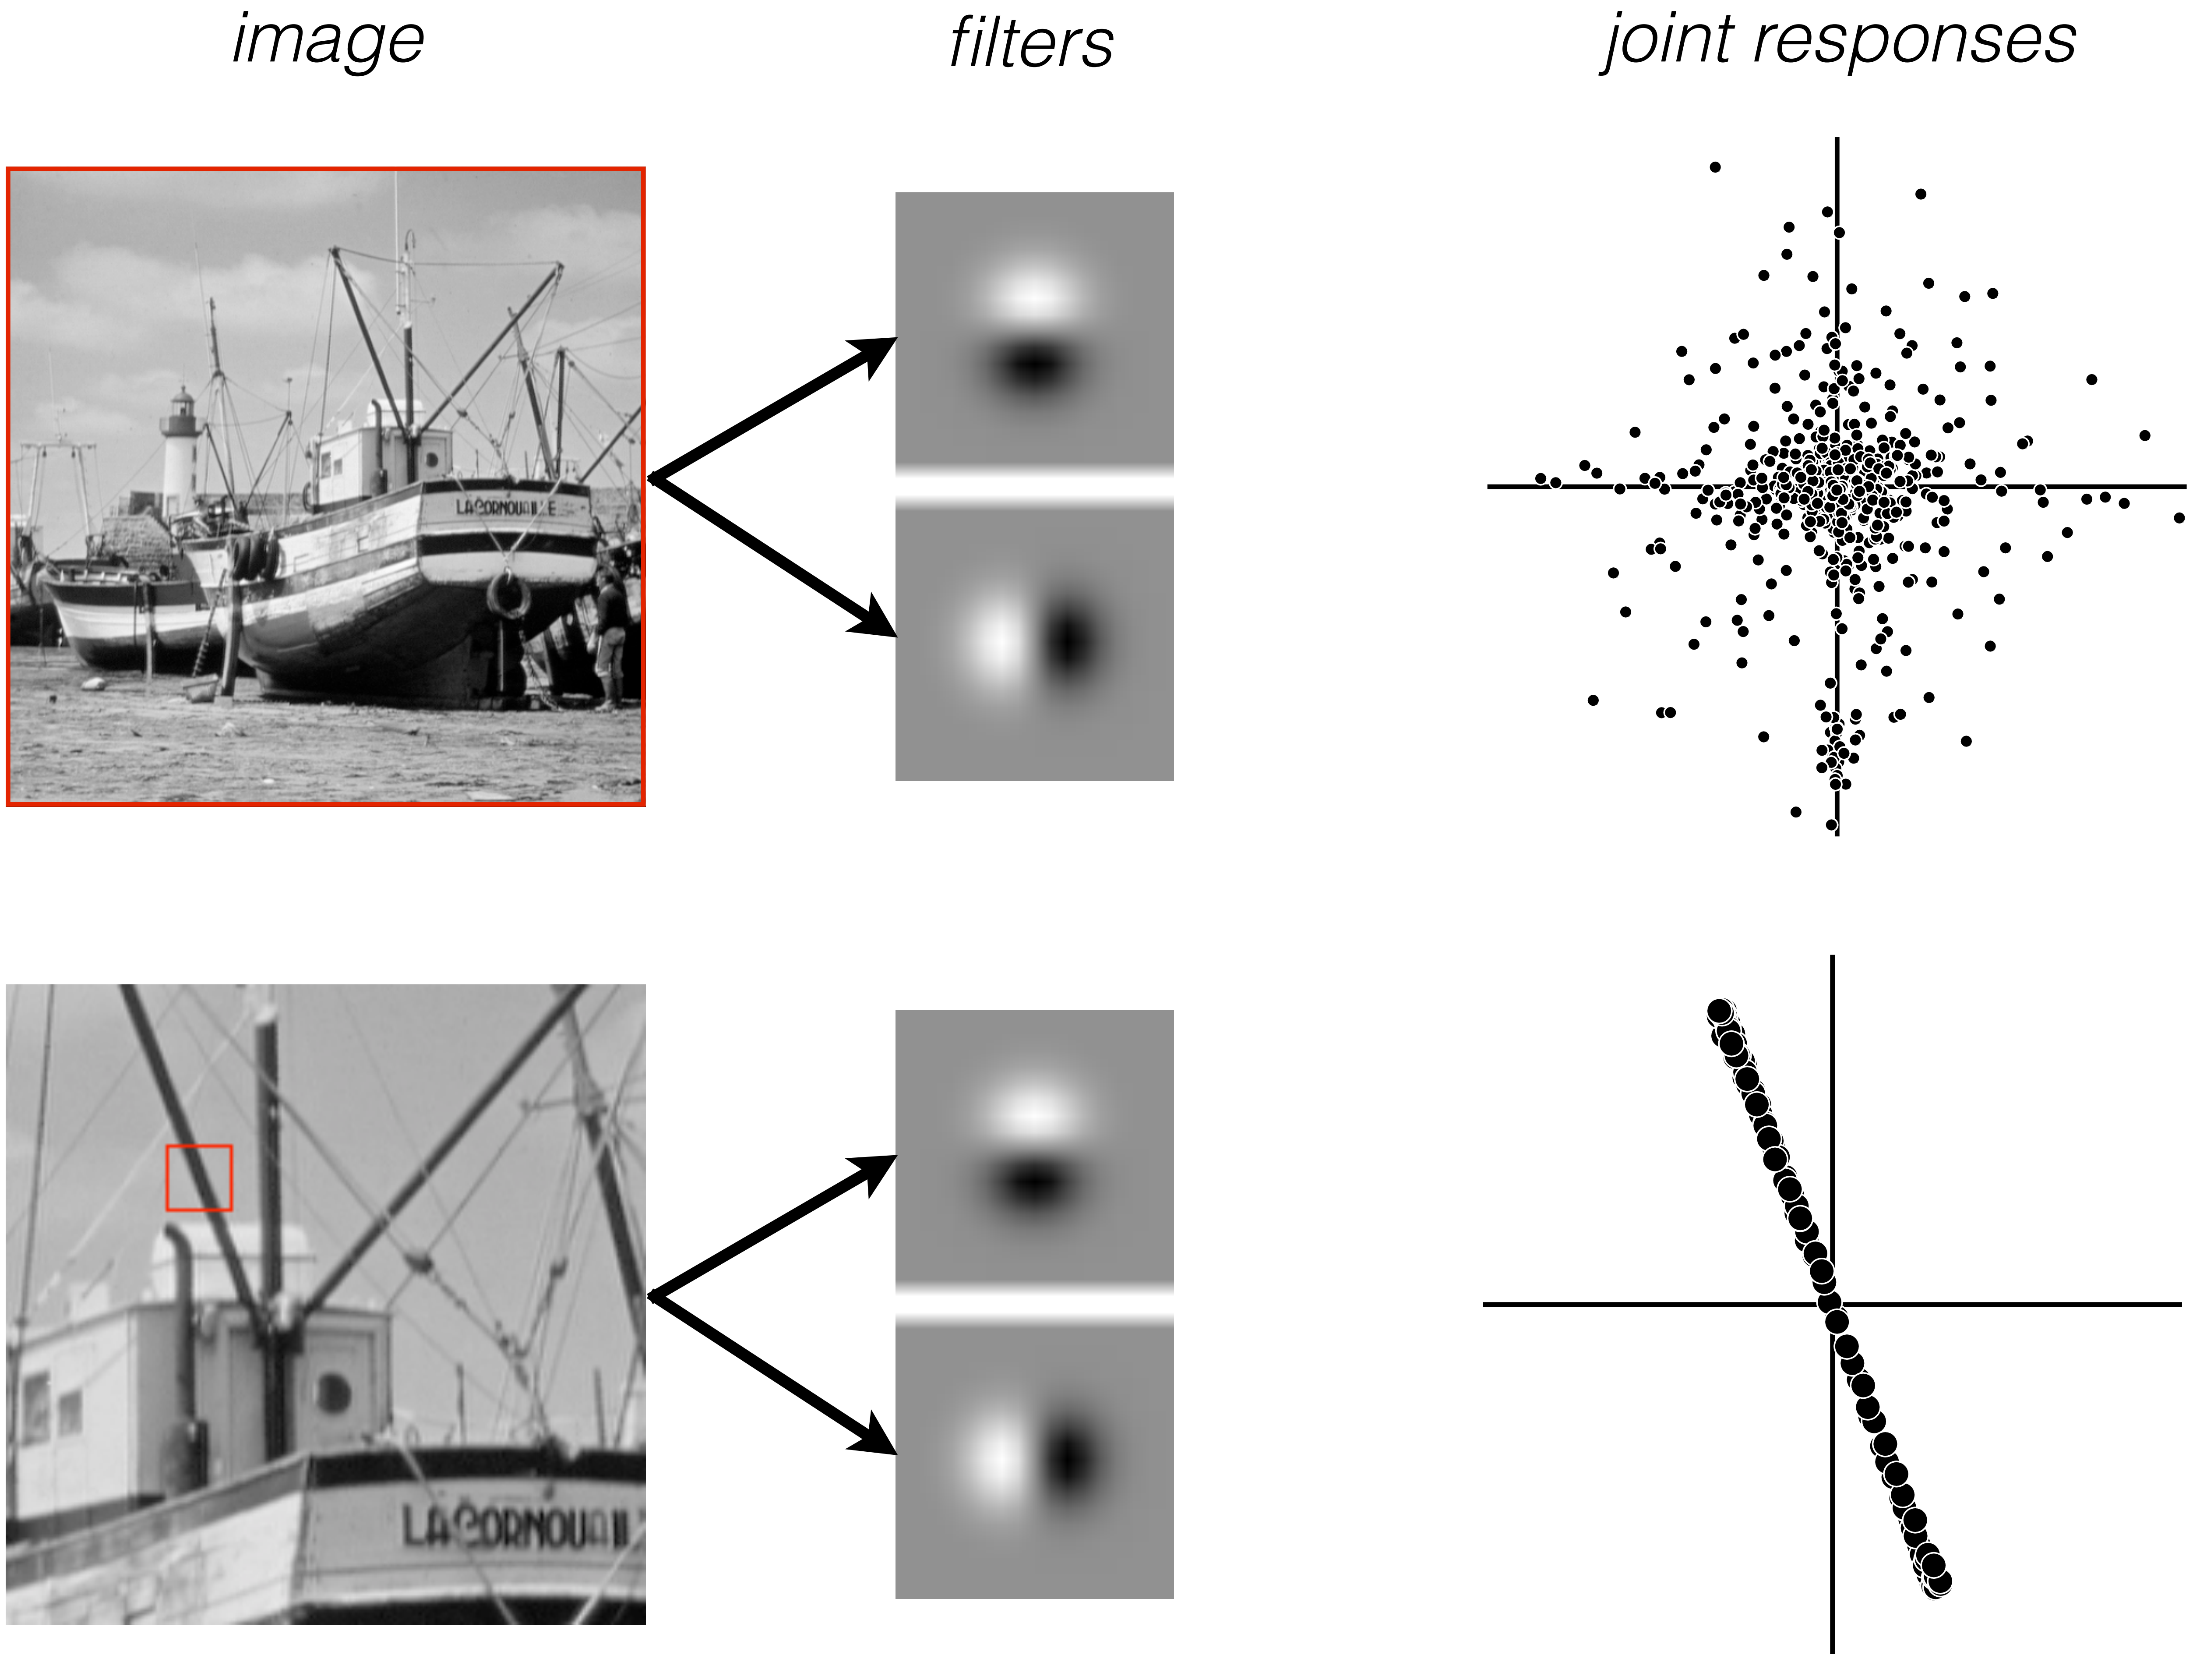

image
filters
joint responses

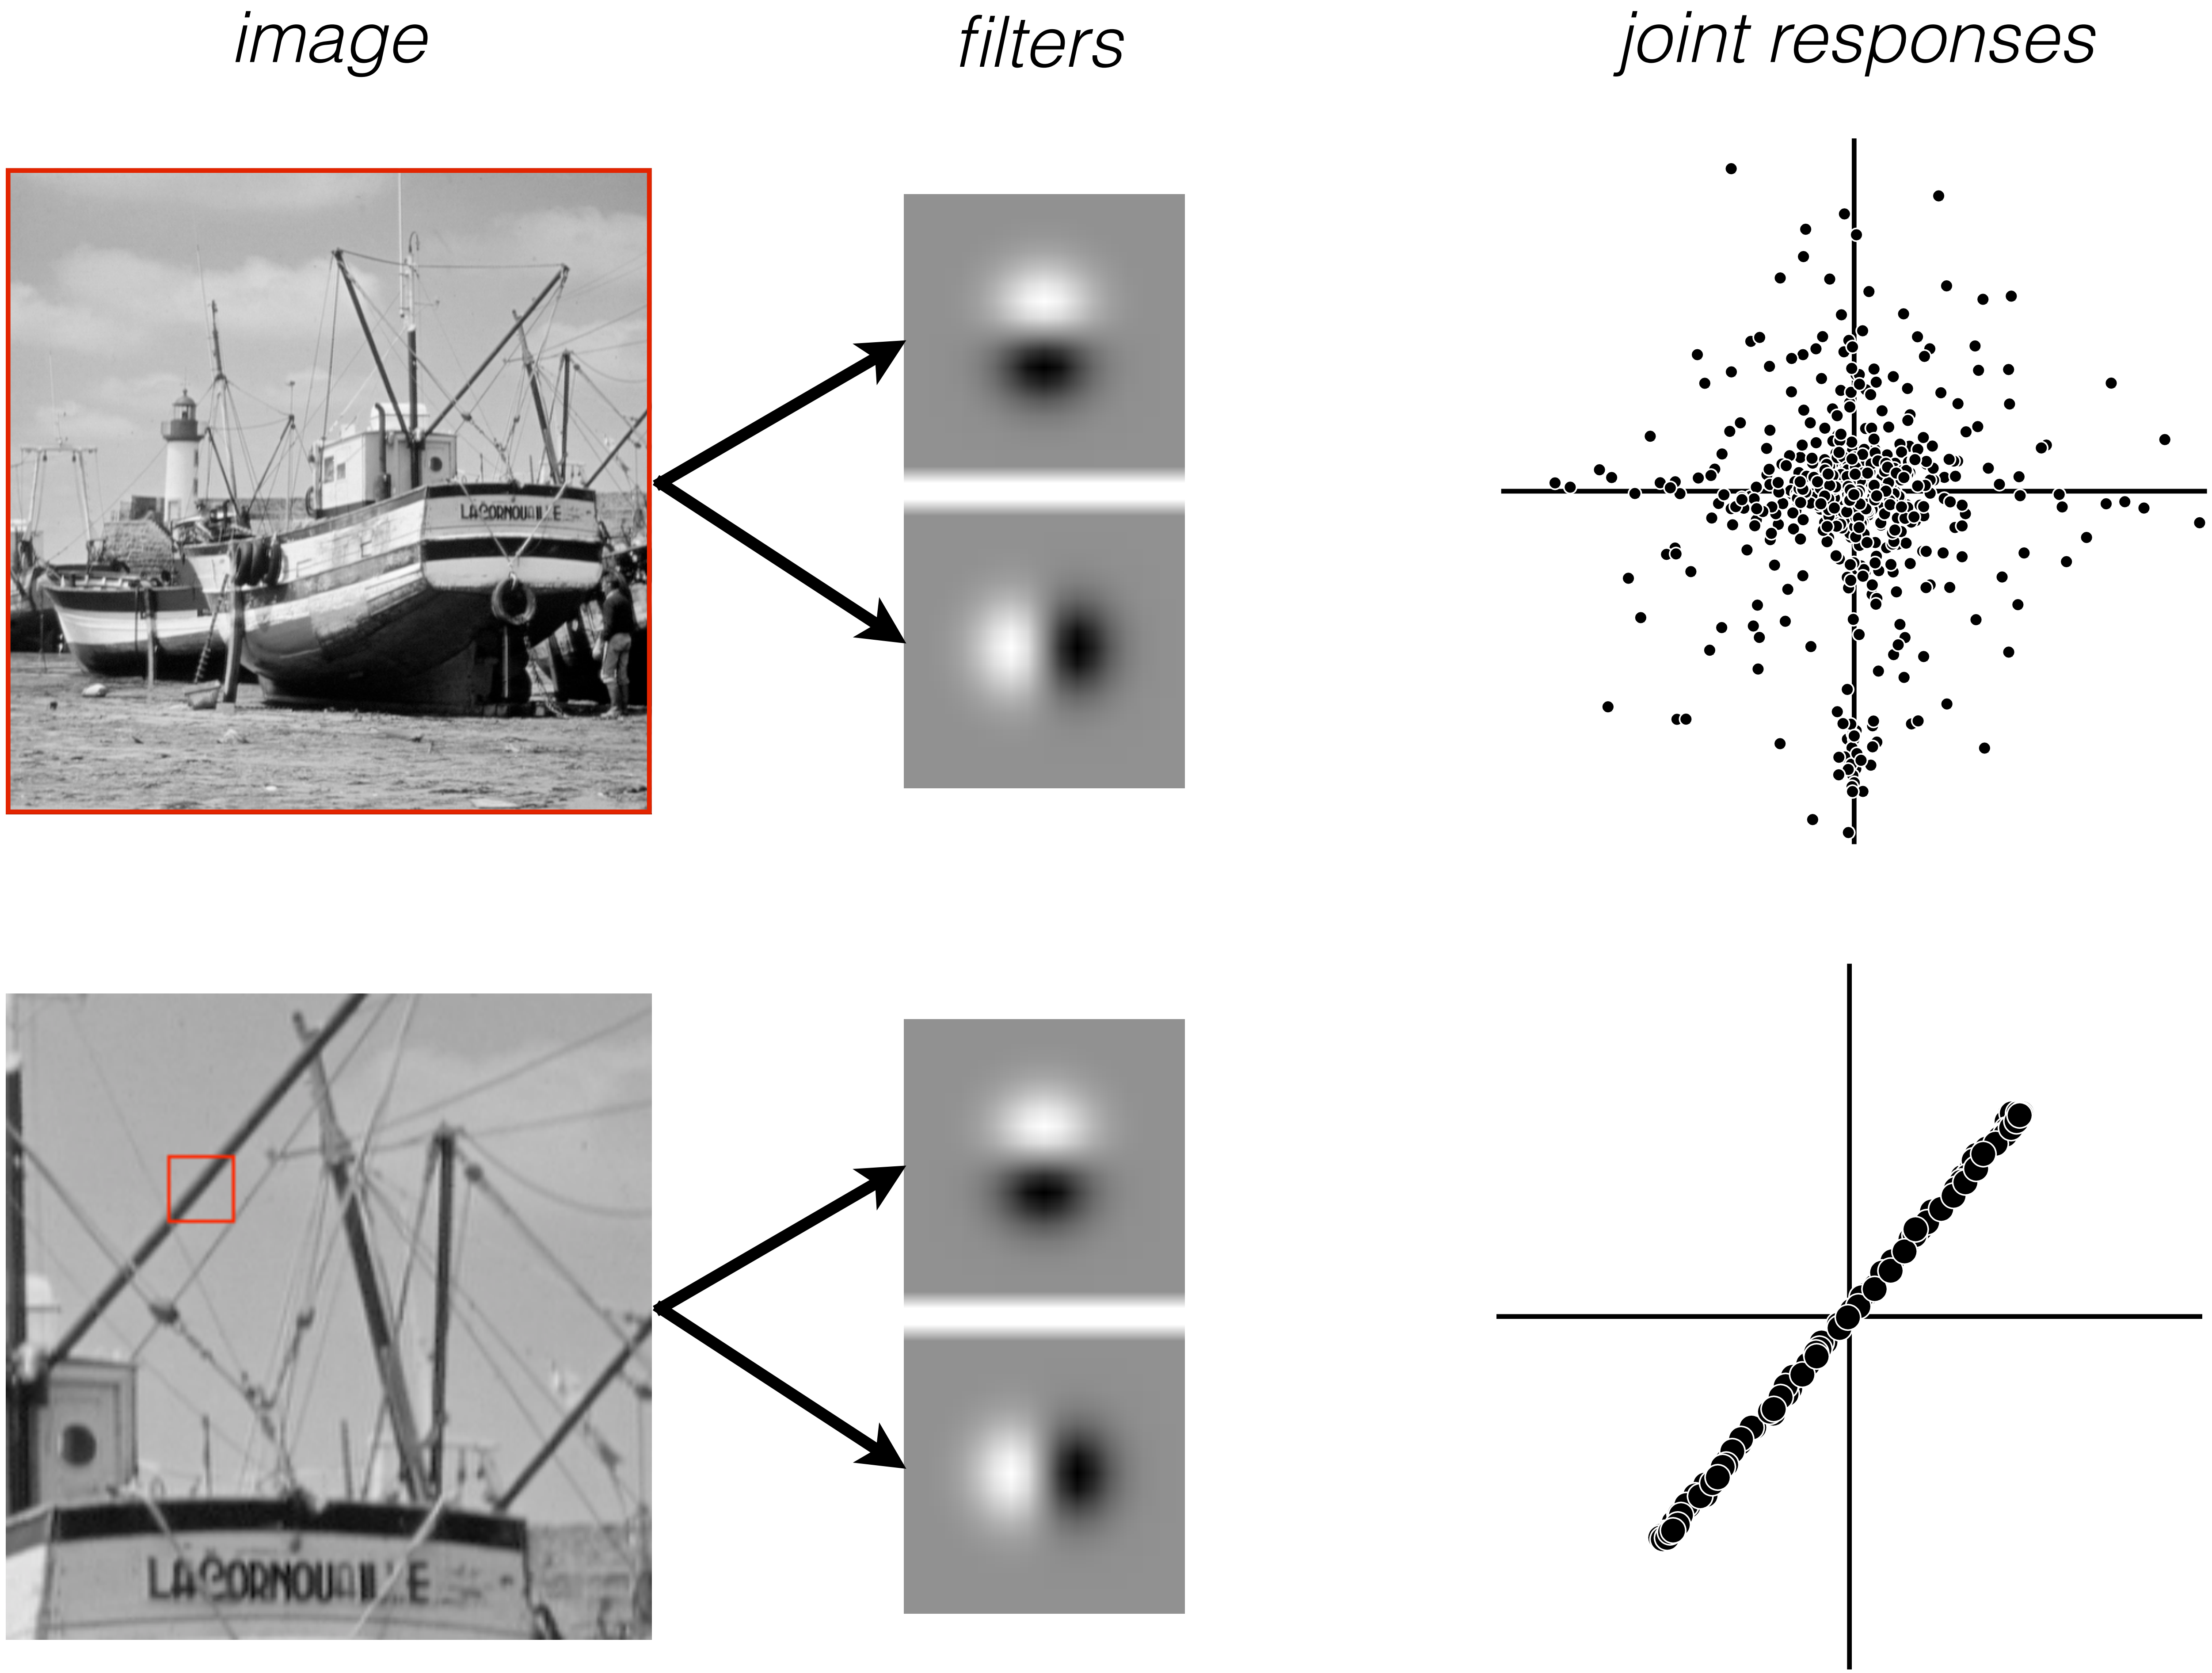
image
filters
joint responses
LACORNOUAILLE

nuclear norm

$$\left|\left|Y\right|\right|_* = \sum_i \sigma_i$$

L1 norm of singular values

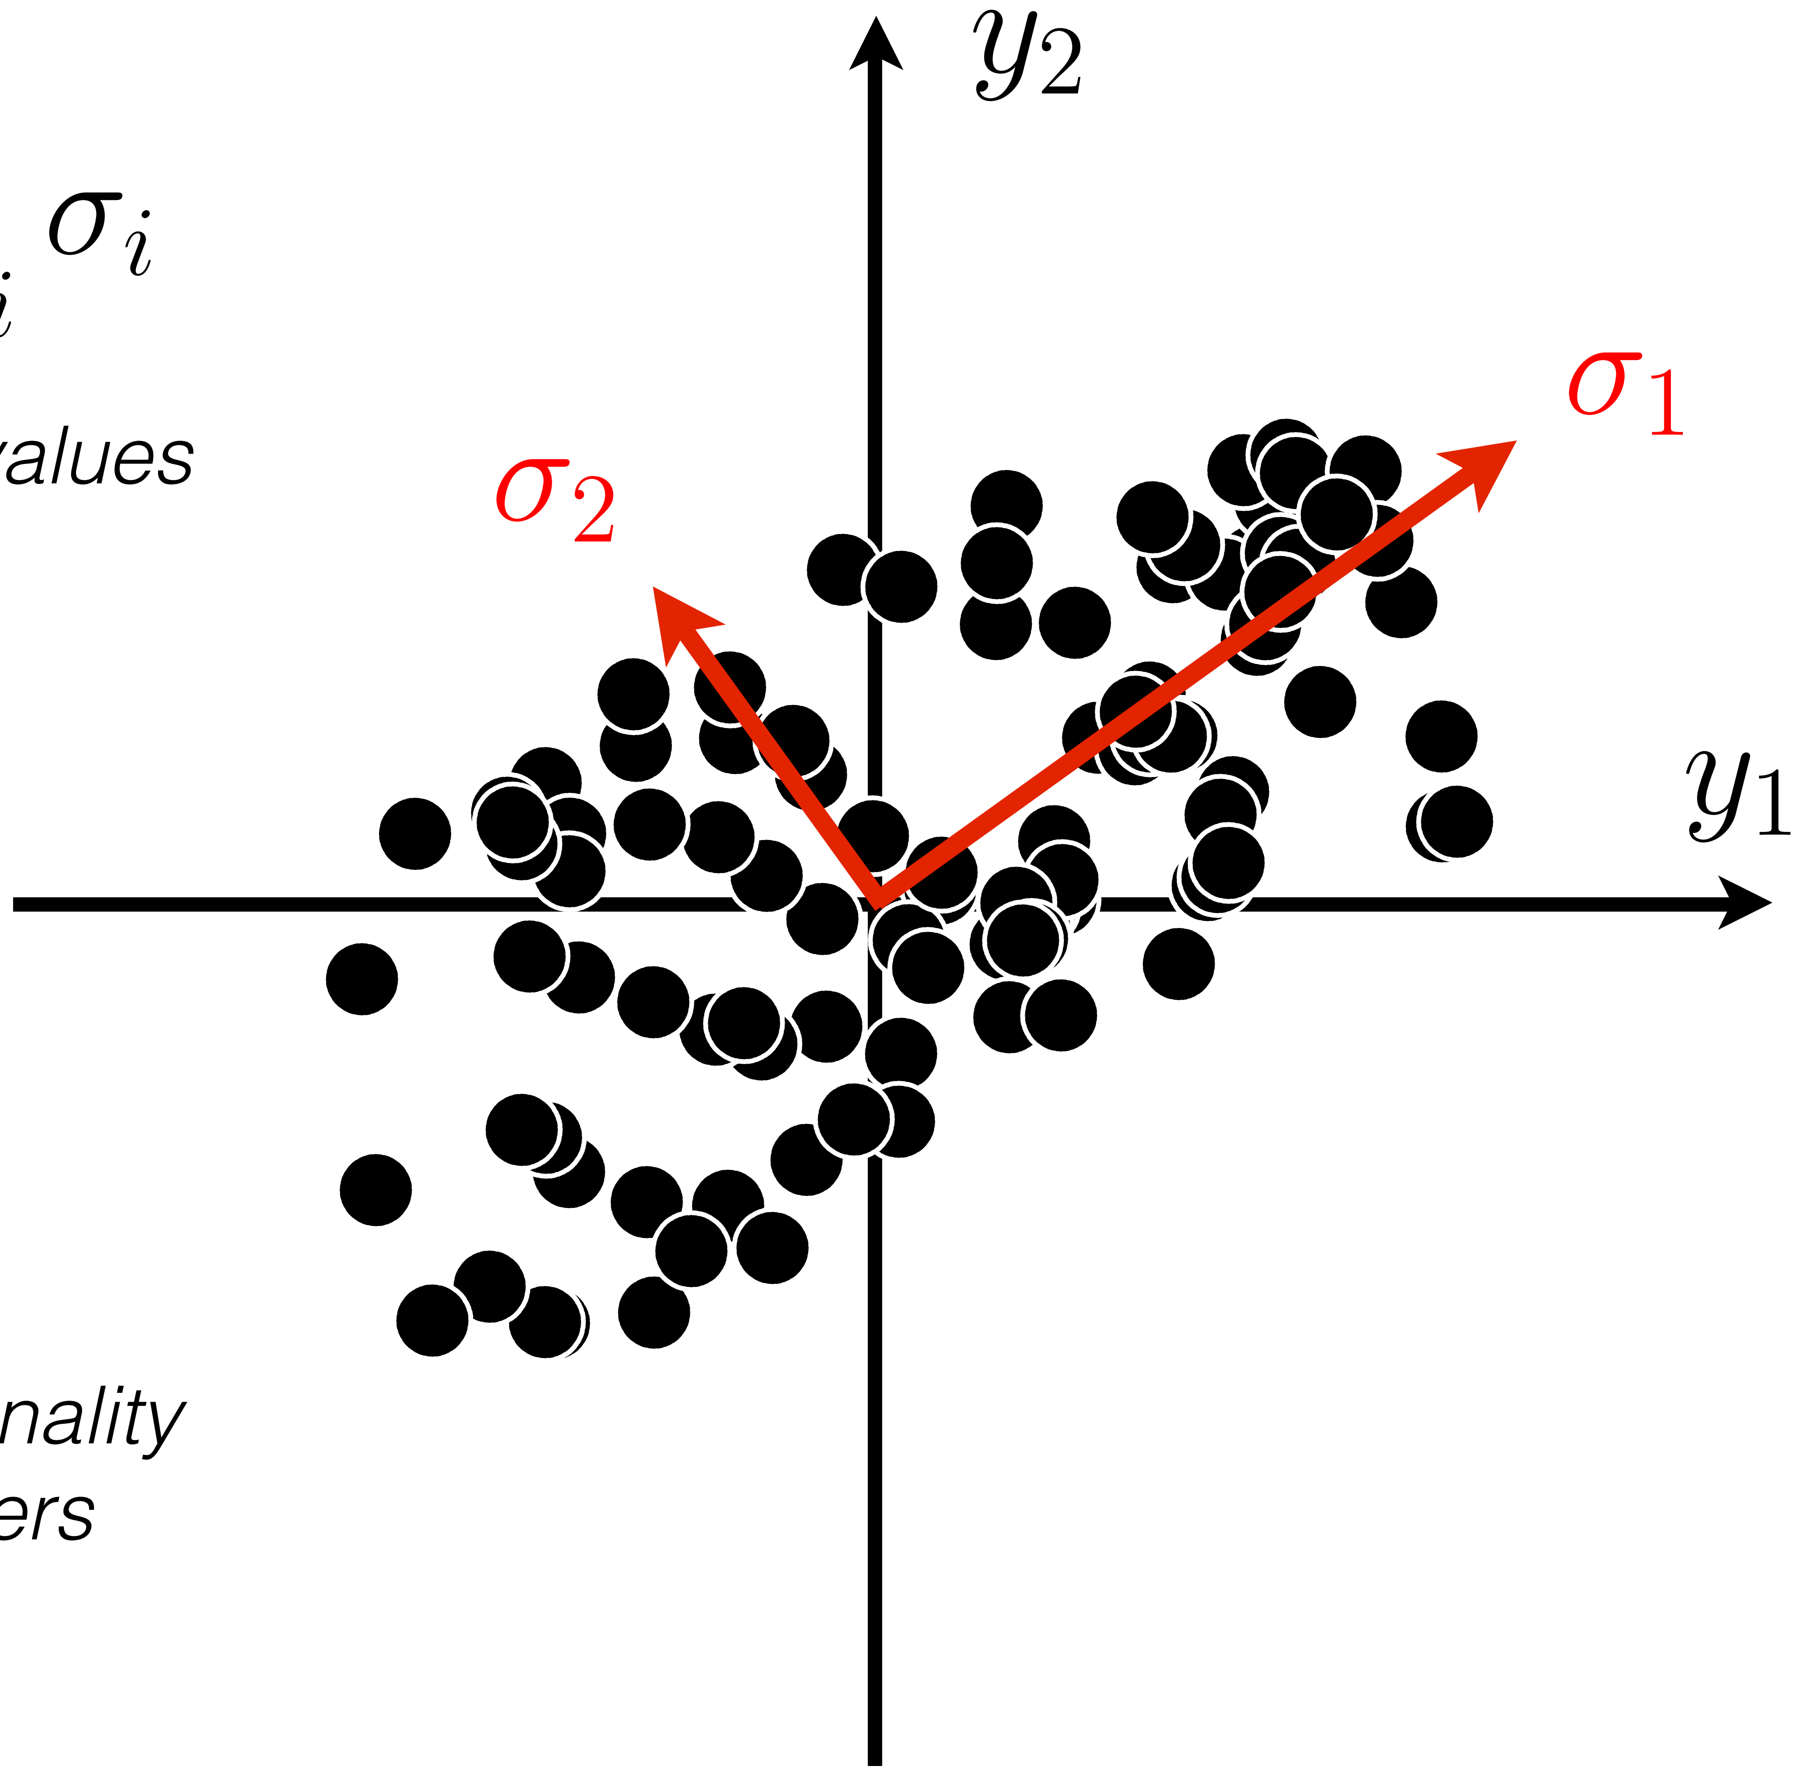

measure of dimensionality
for any number of filters

filters

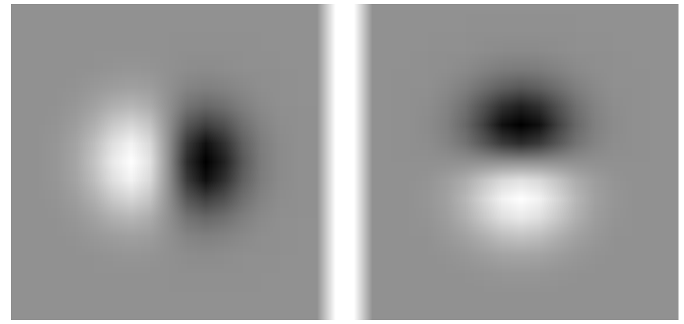

images

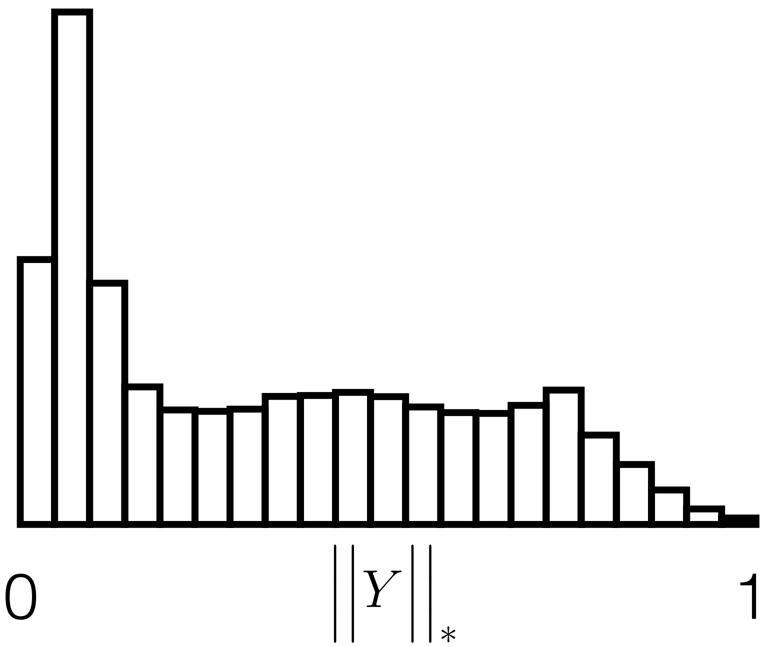

filters

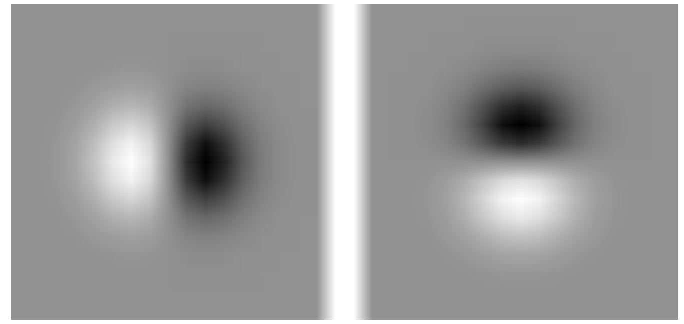

phase randomized

images

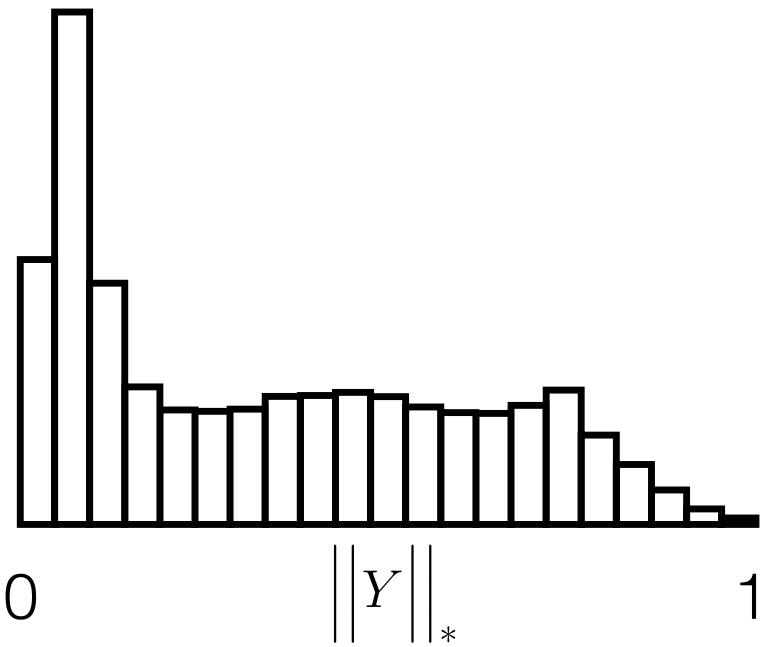

filters

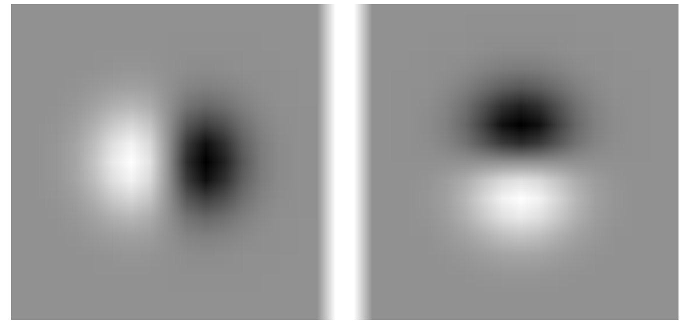

images

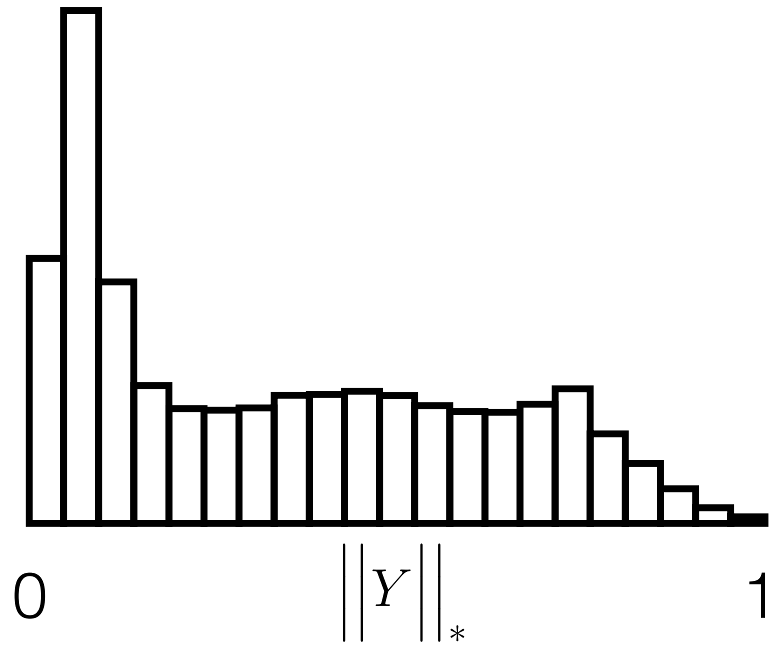

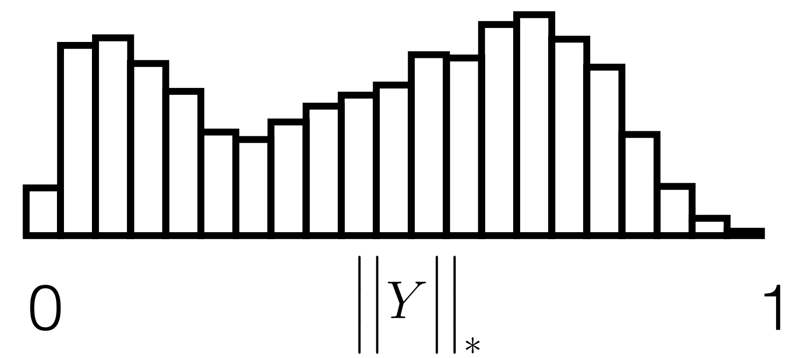

filters

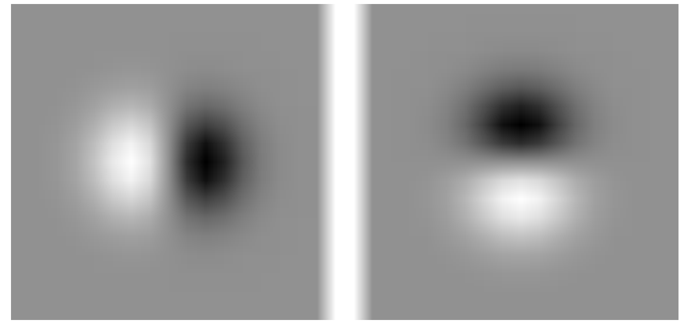

images

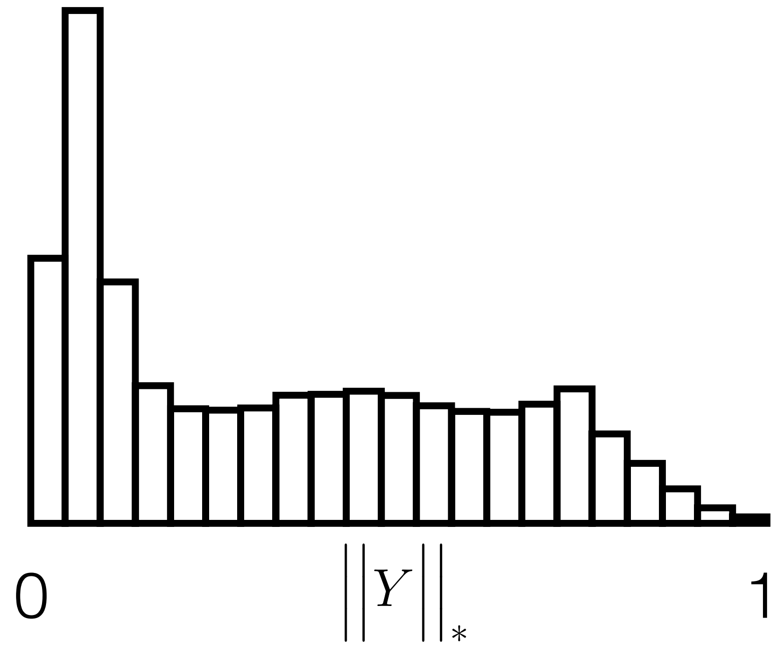

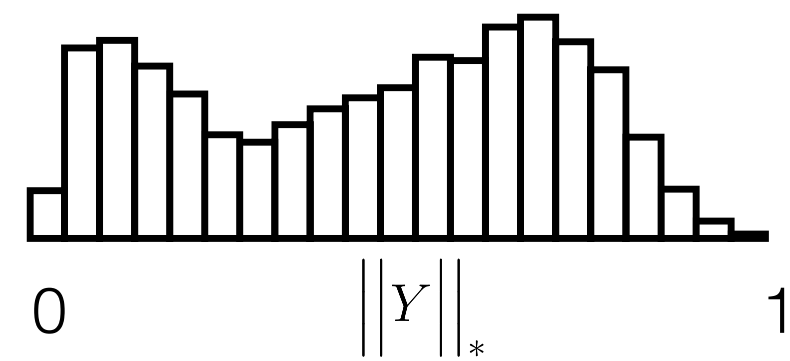

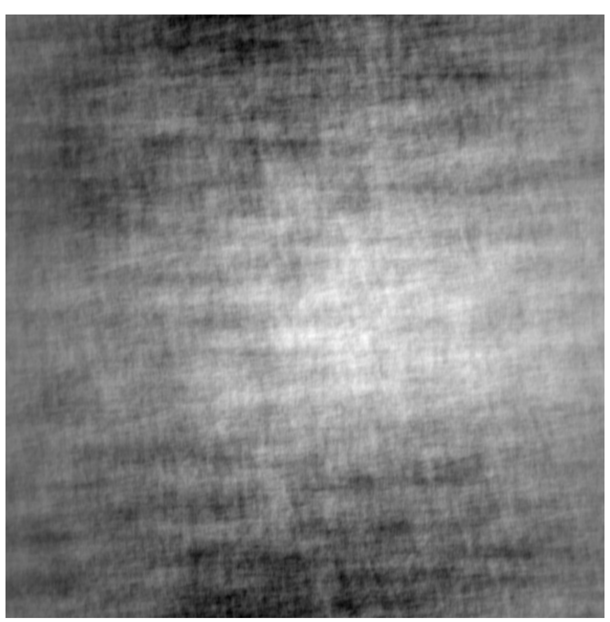

phase randomized

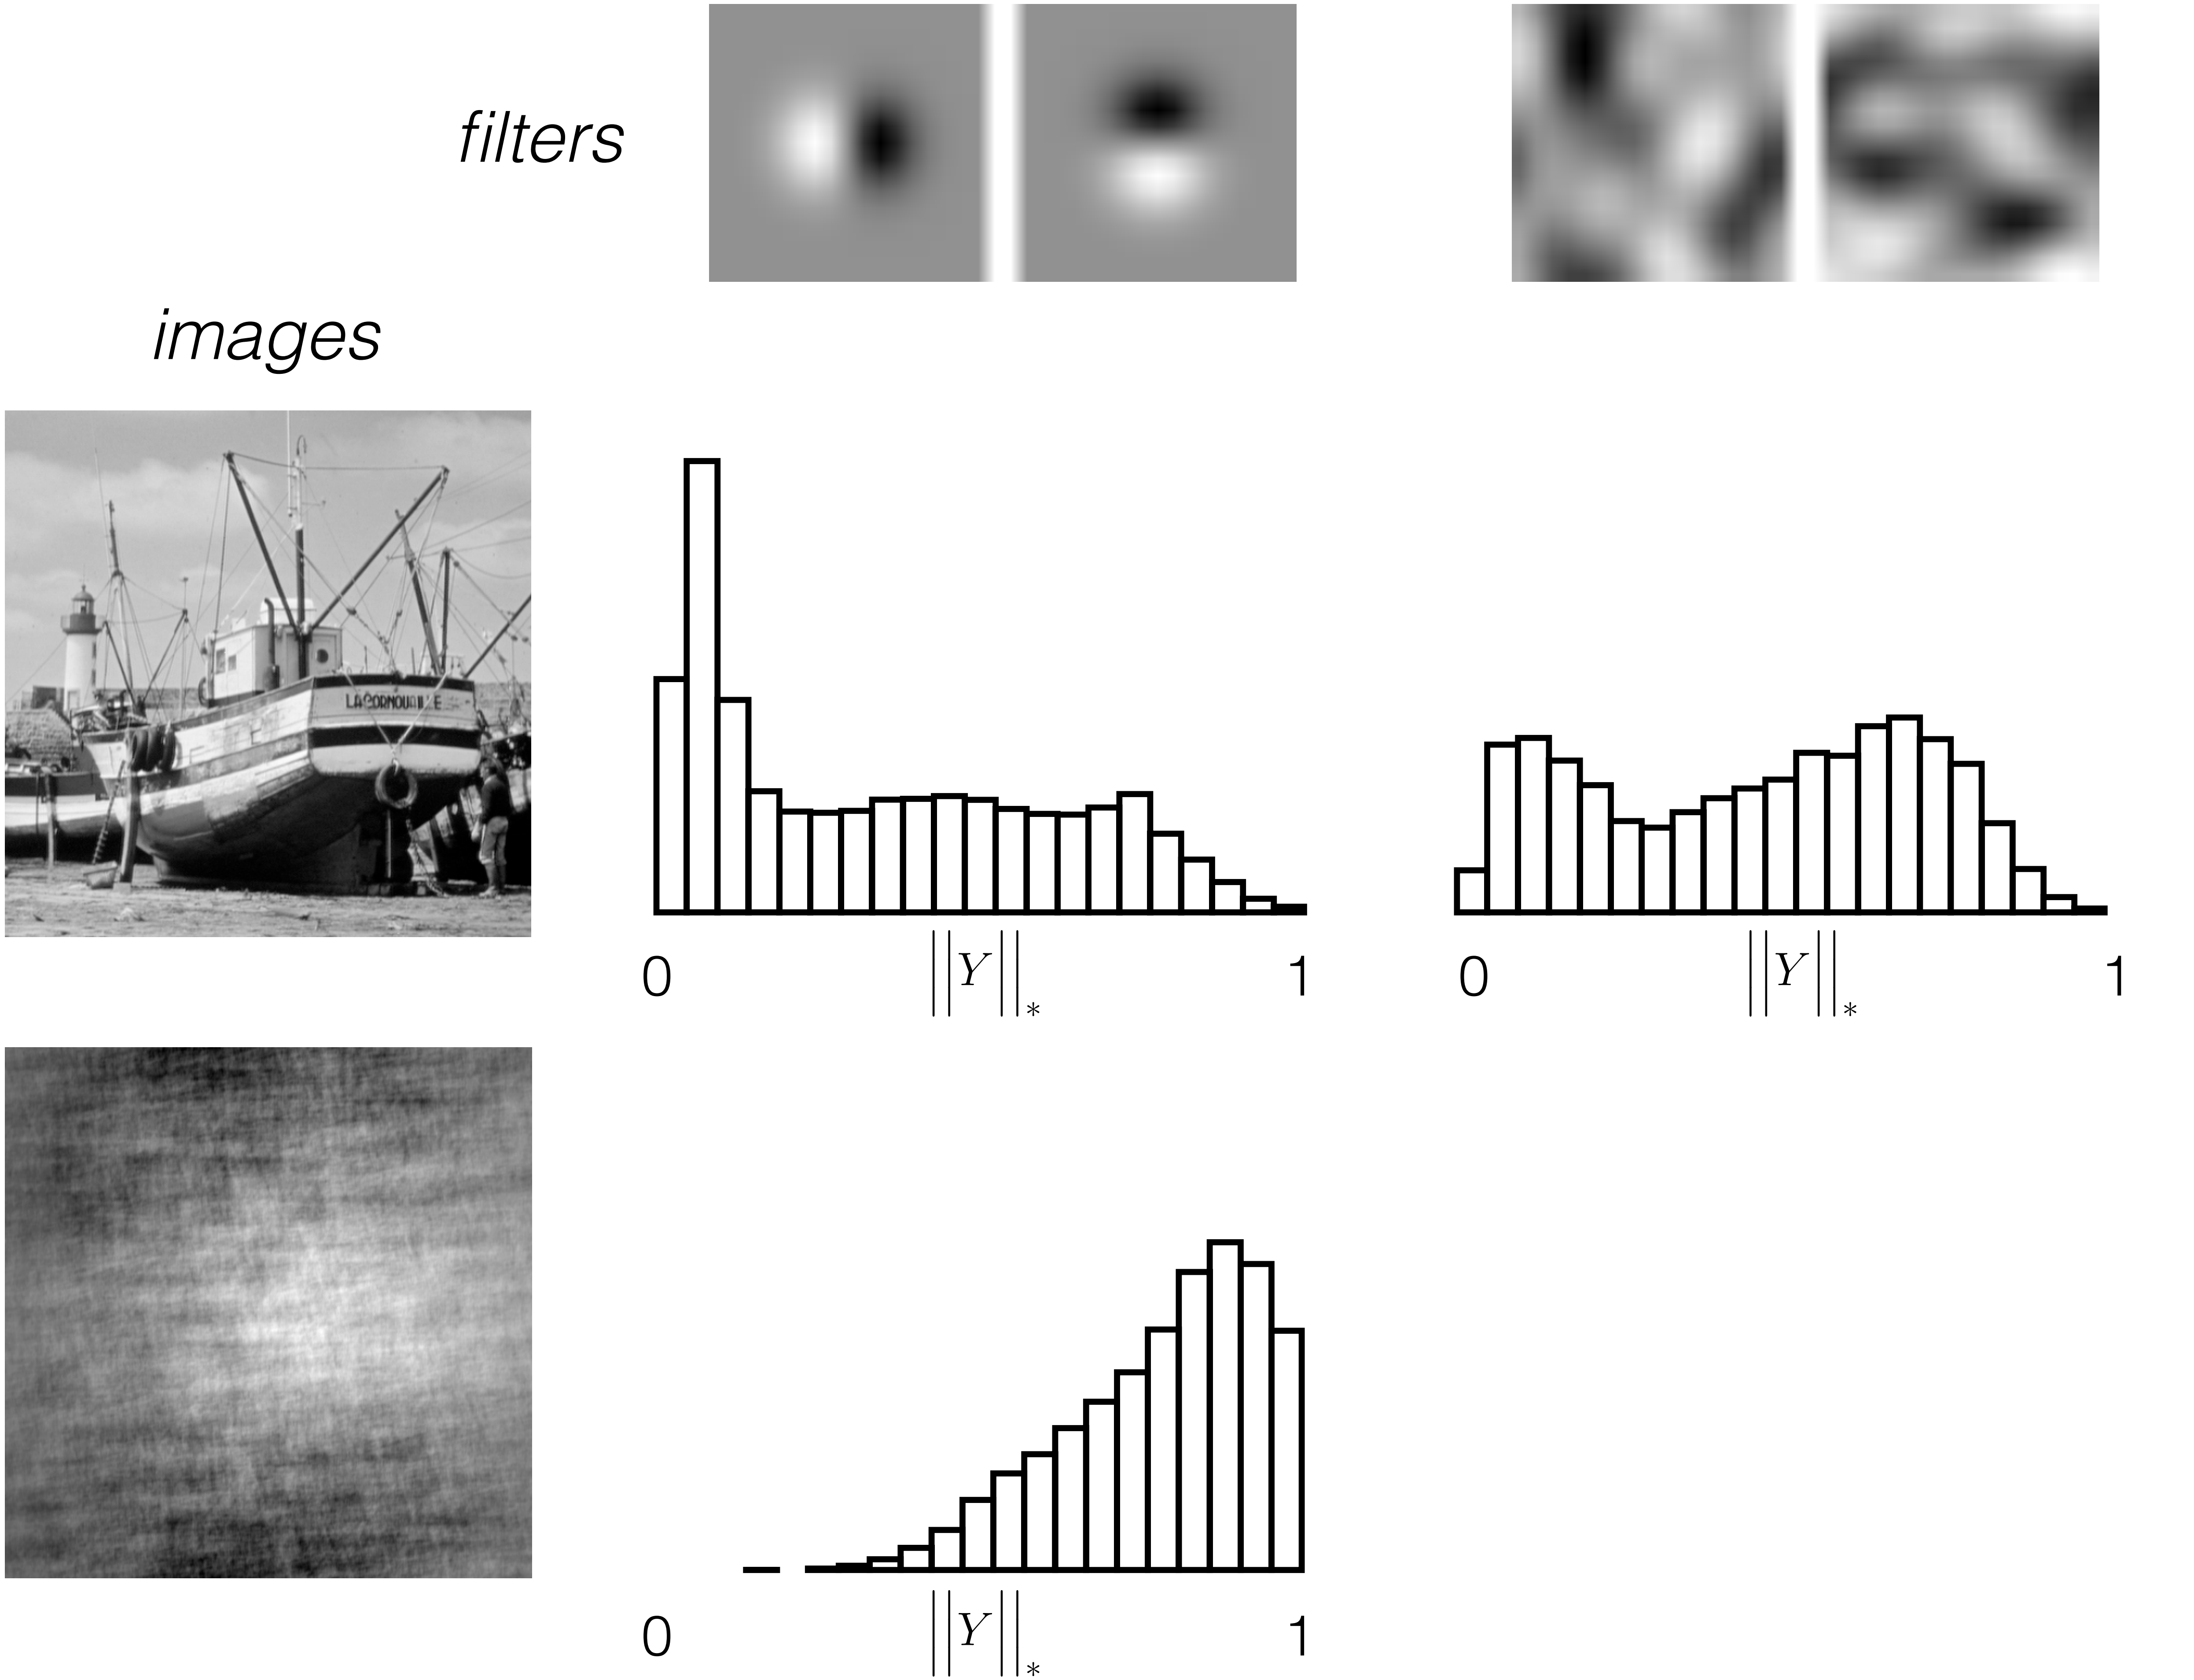
filters
images
0
$\|Y\|_*$
1
0
$\|Y\|_*$
1
0
$\|Y\|_*$
1

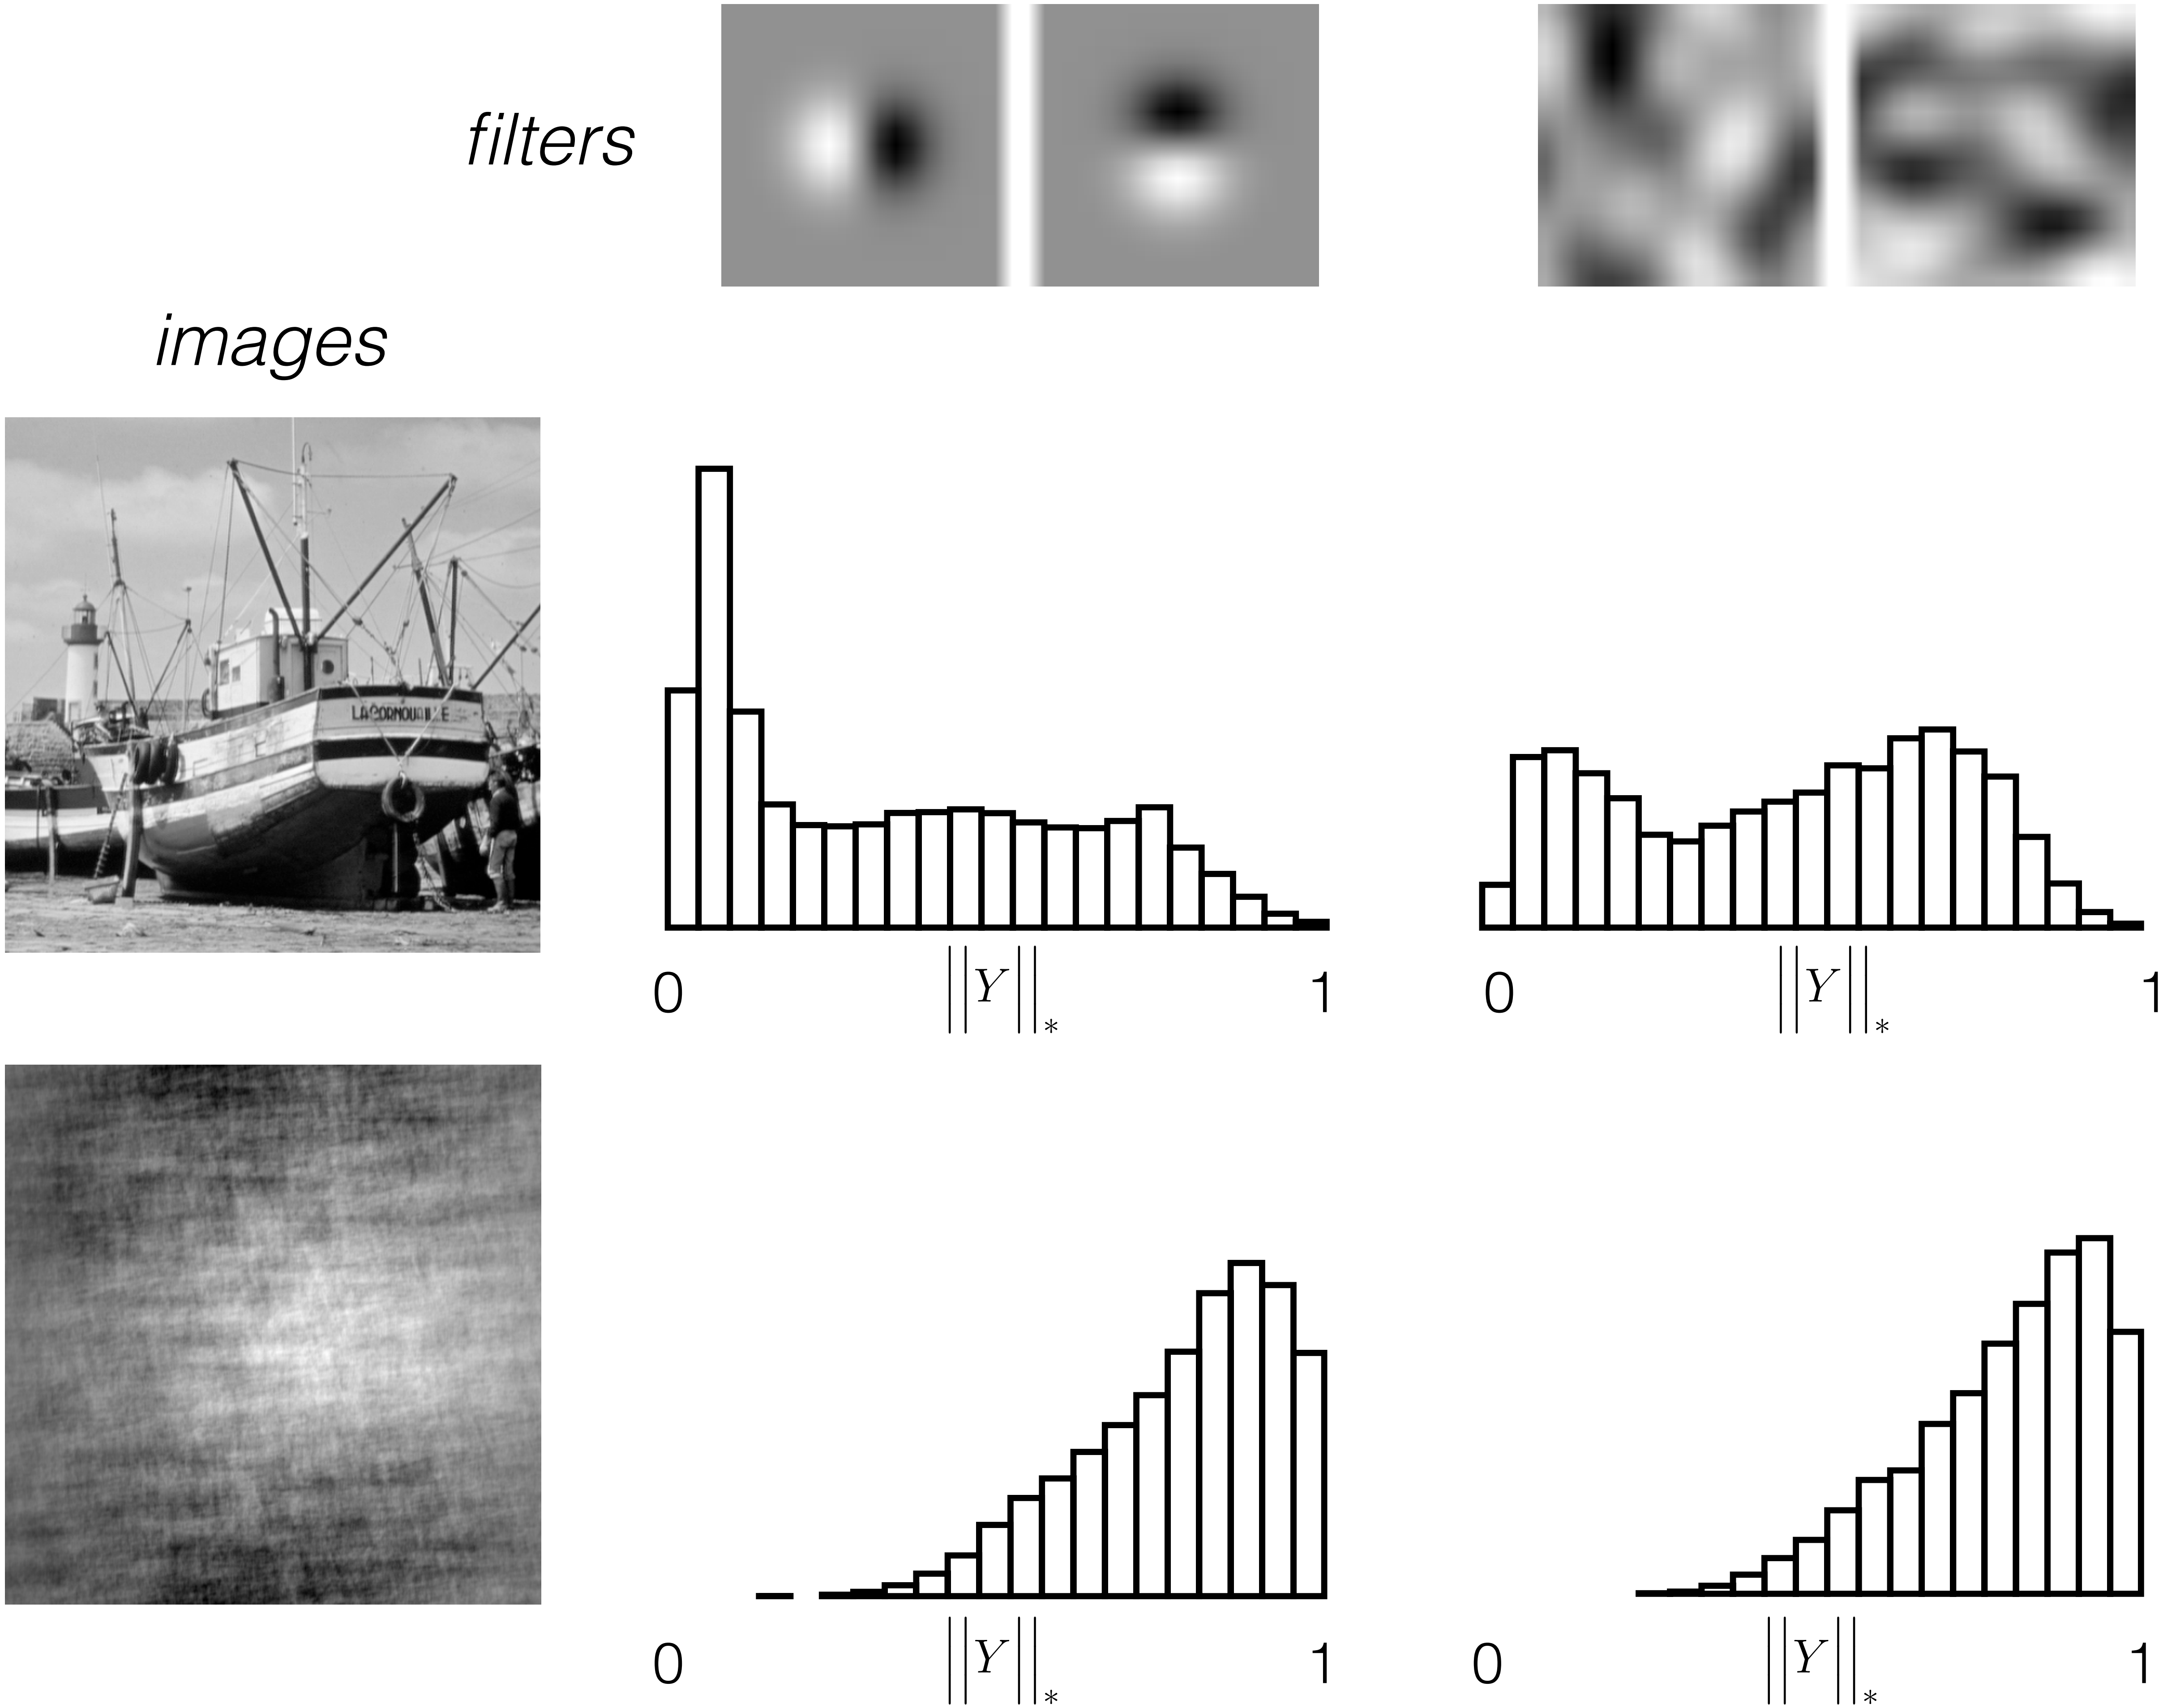
filters
images
0
$||Y||_*$
1
0
$||Y||_*$
1
0
$||Y||_*$
1
0
$||Y||_*$
1

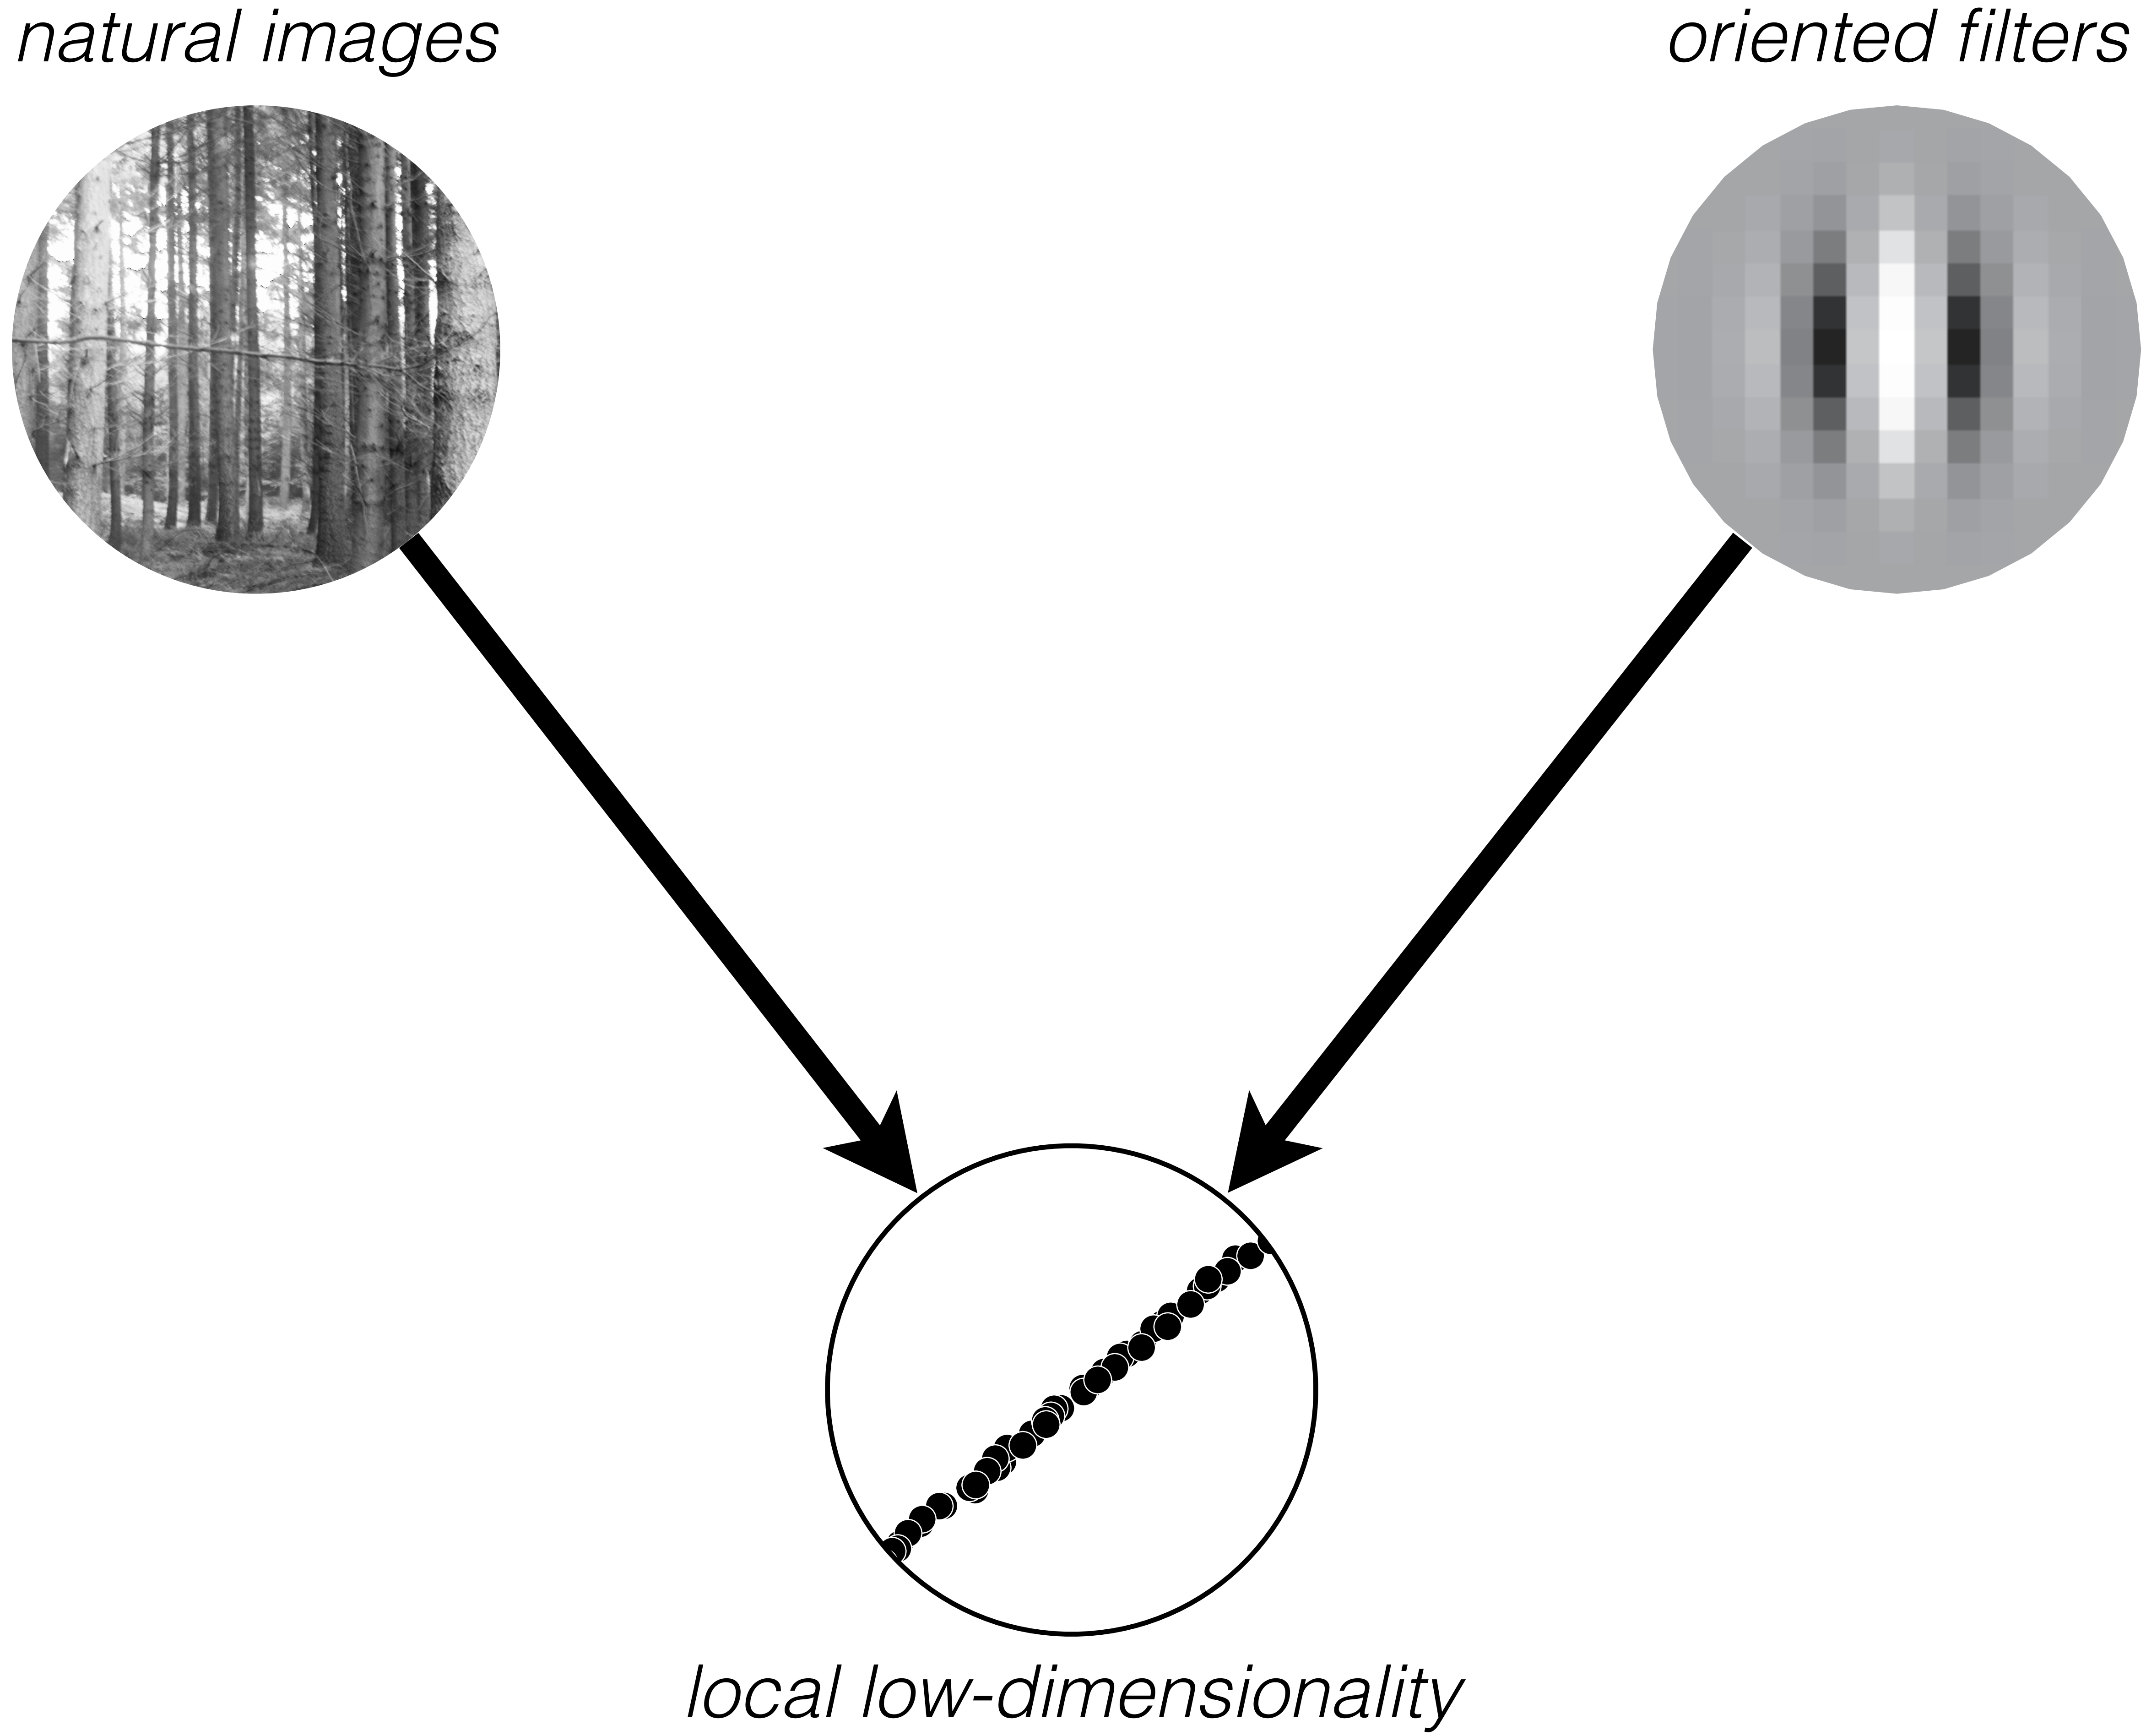
natural images
oriented filters
local low-dimensionality

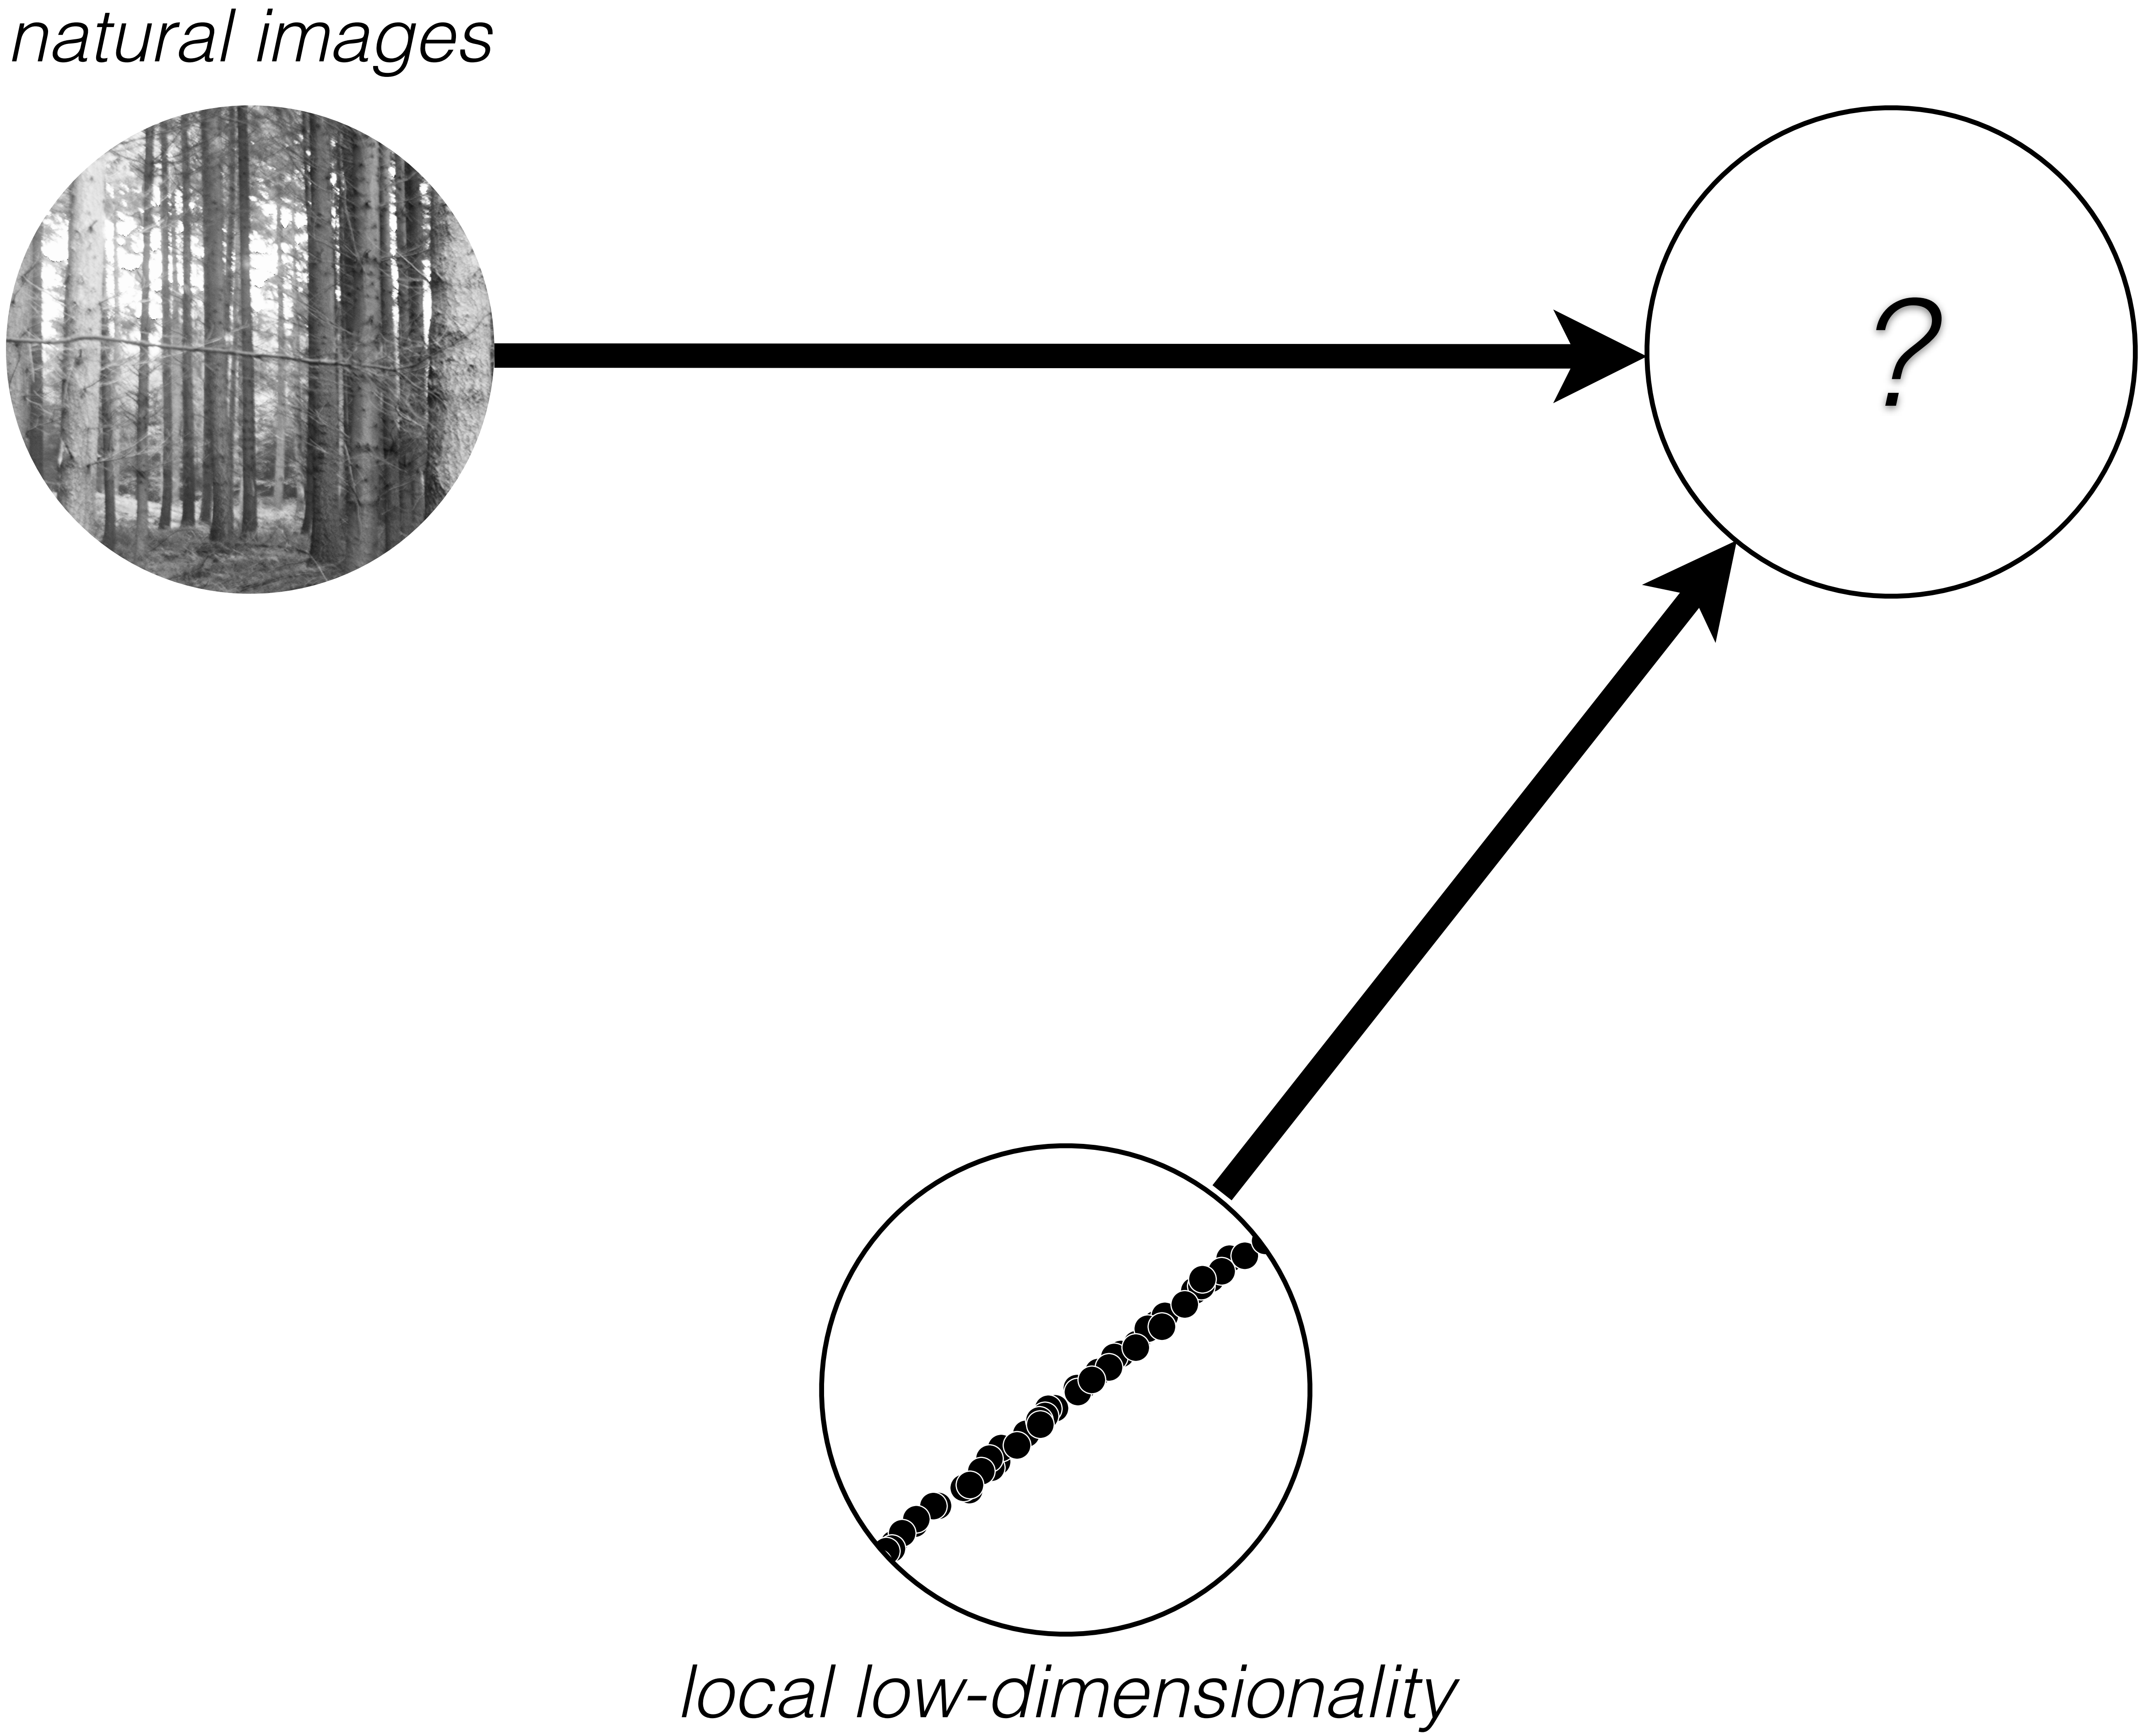

natural images
?
local low-dimensionality

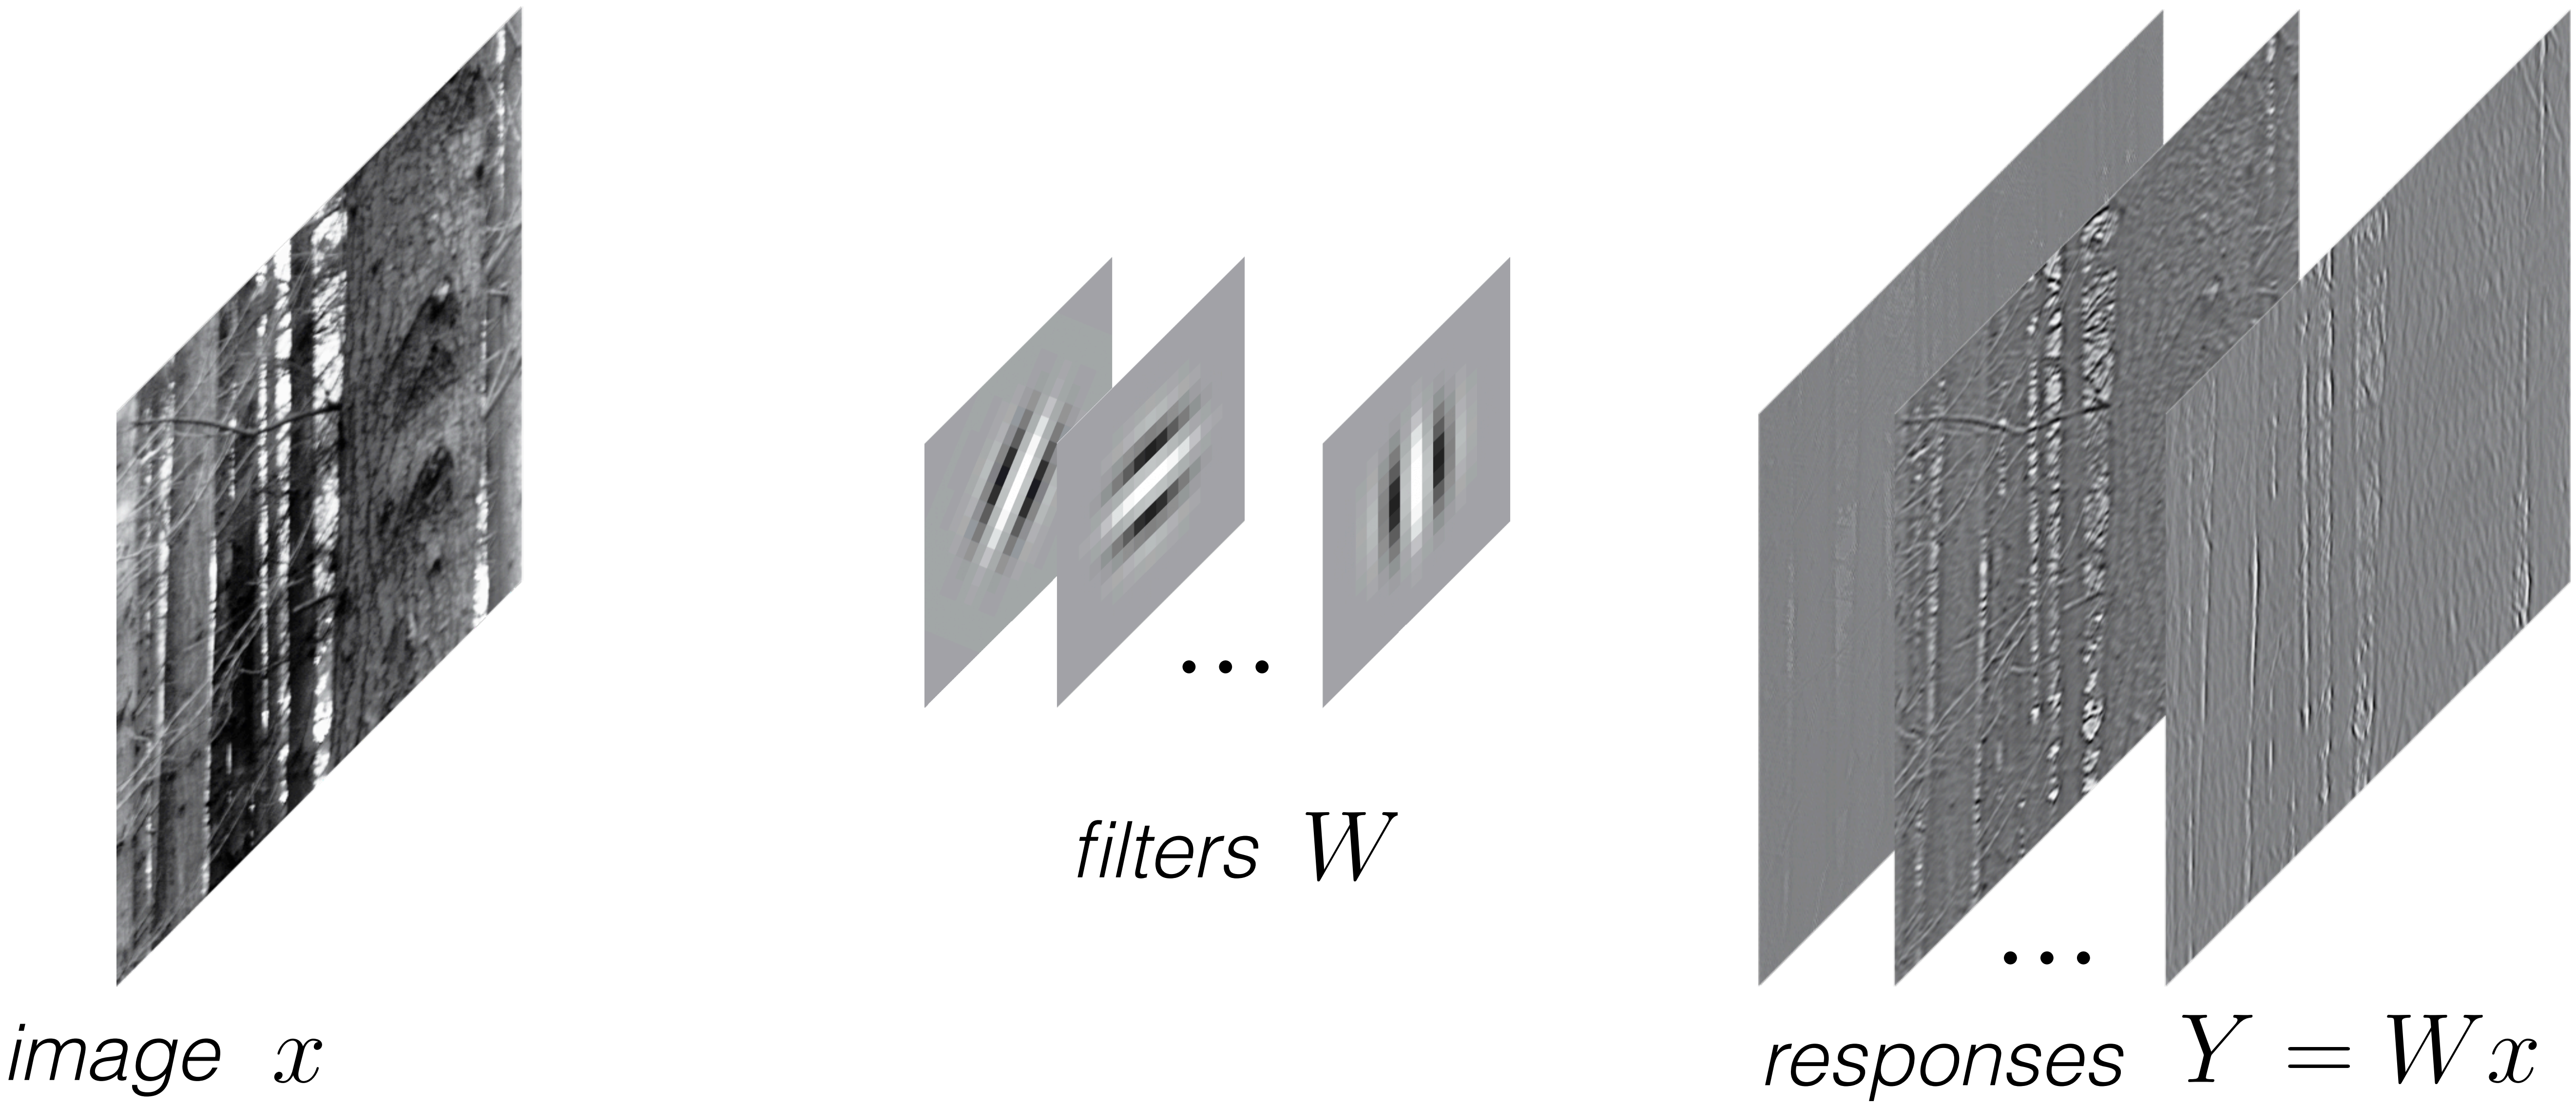
image x
filters W
responses Y = Wx

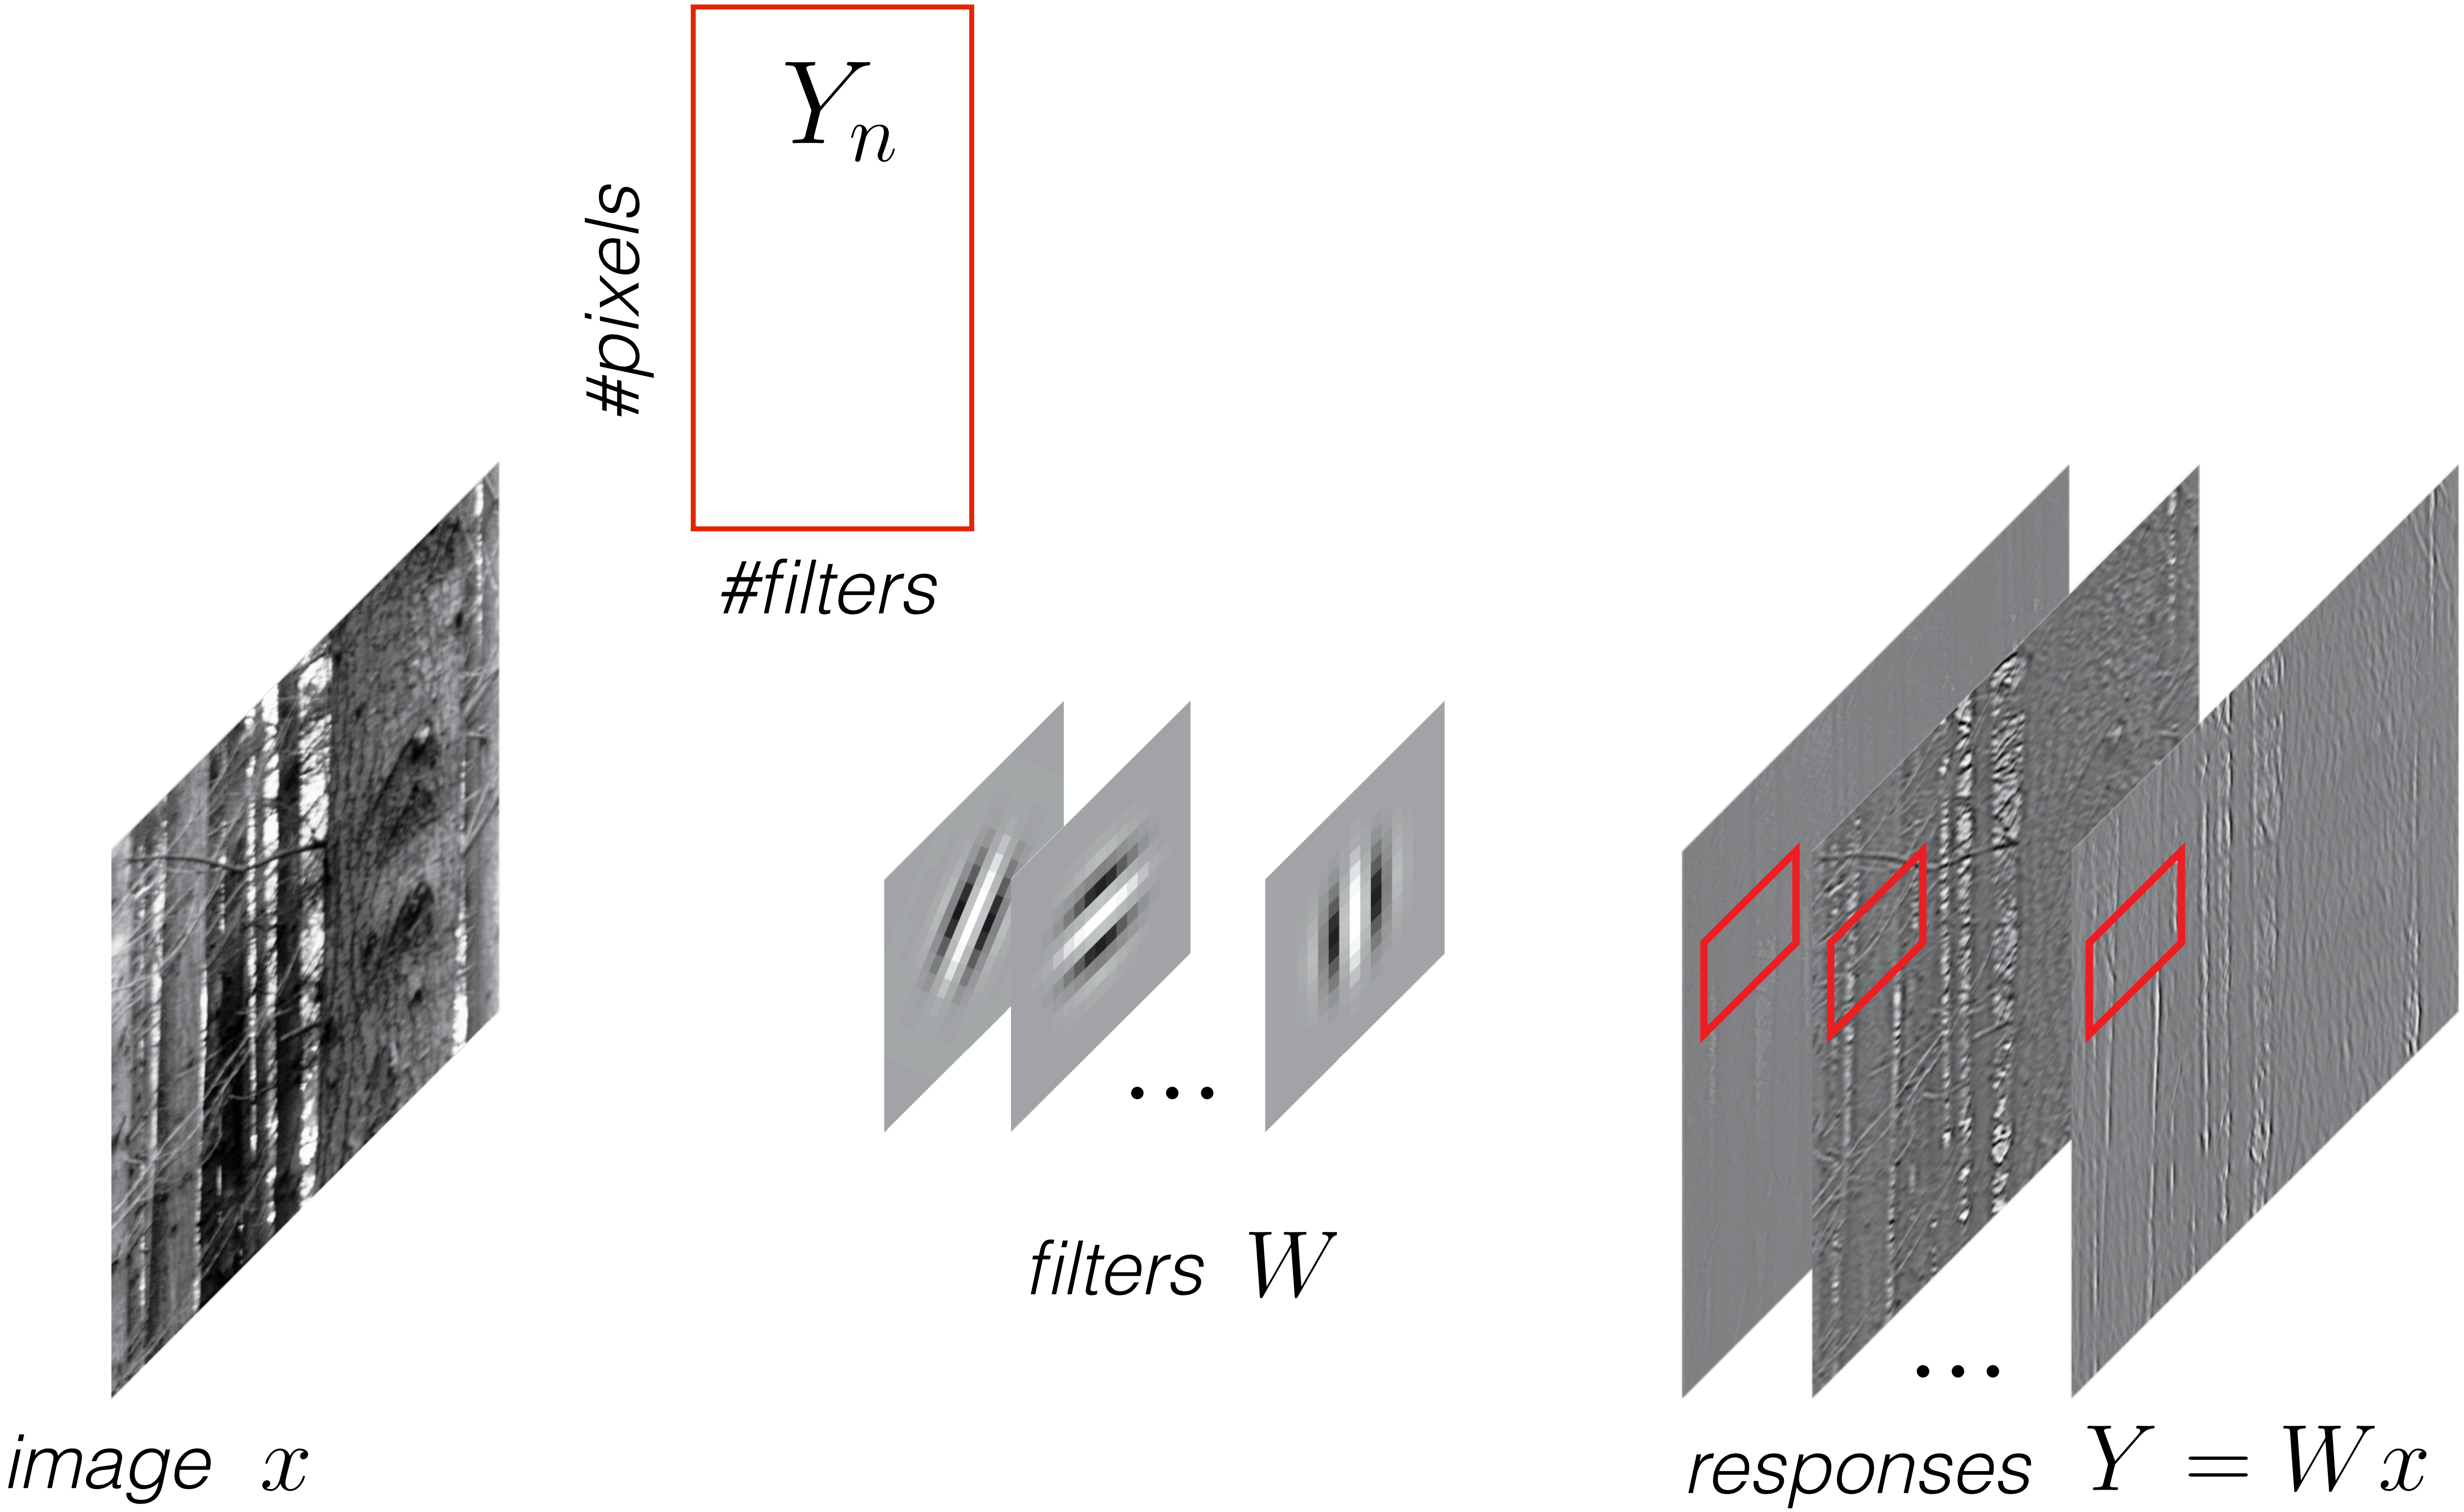

Y_n
#pixels
#filters
filters W
image x
responses $Y = Wx$

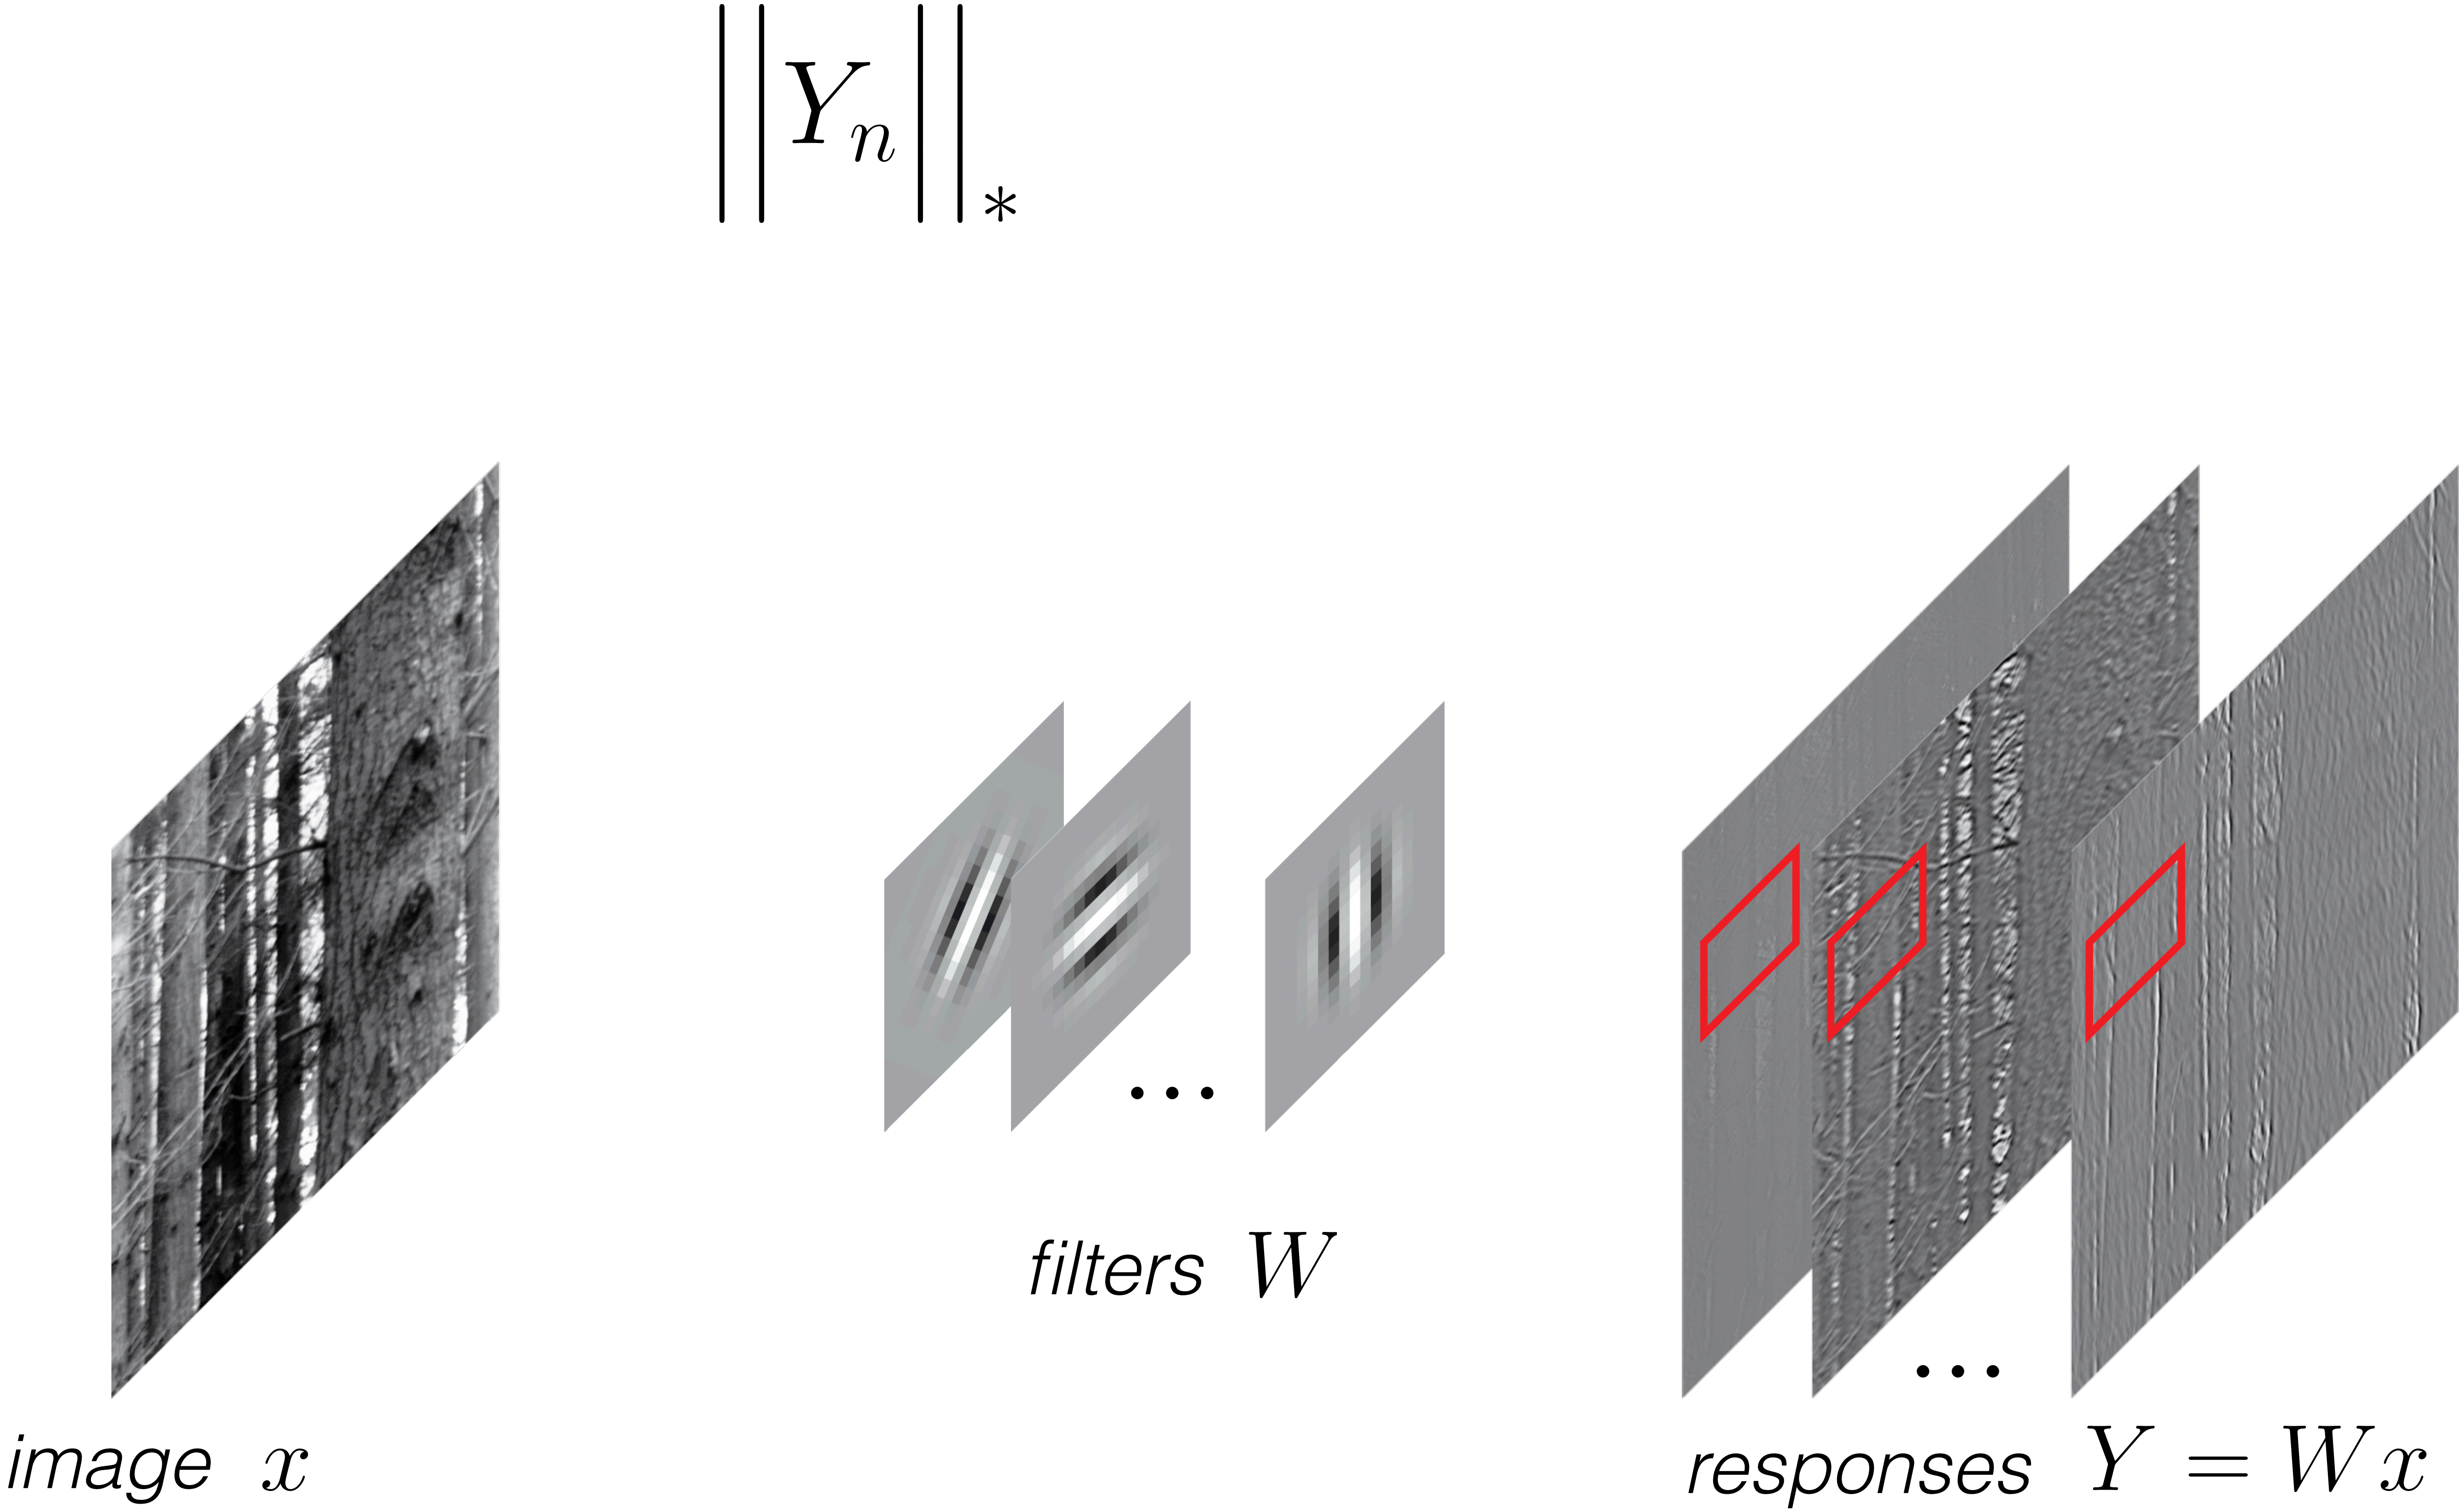

$\left\|\left|Y_n\right|\right\|_*$
filters W
image x
responses $Y = Wx$

$$E(W) = \sum_n \left\| Y_n \right\|_*$$

local low-dimensionality

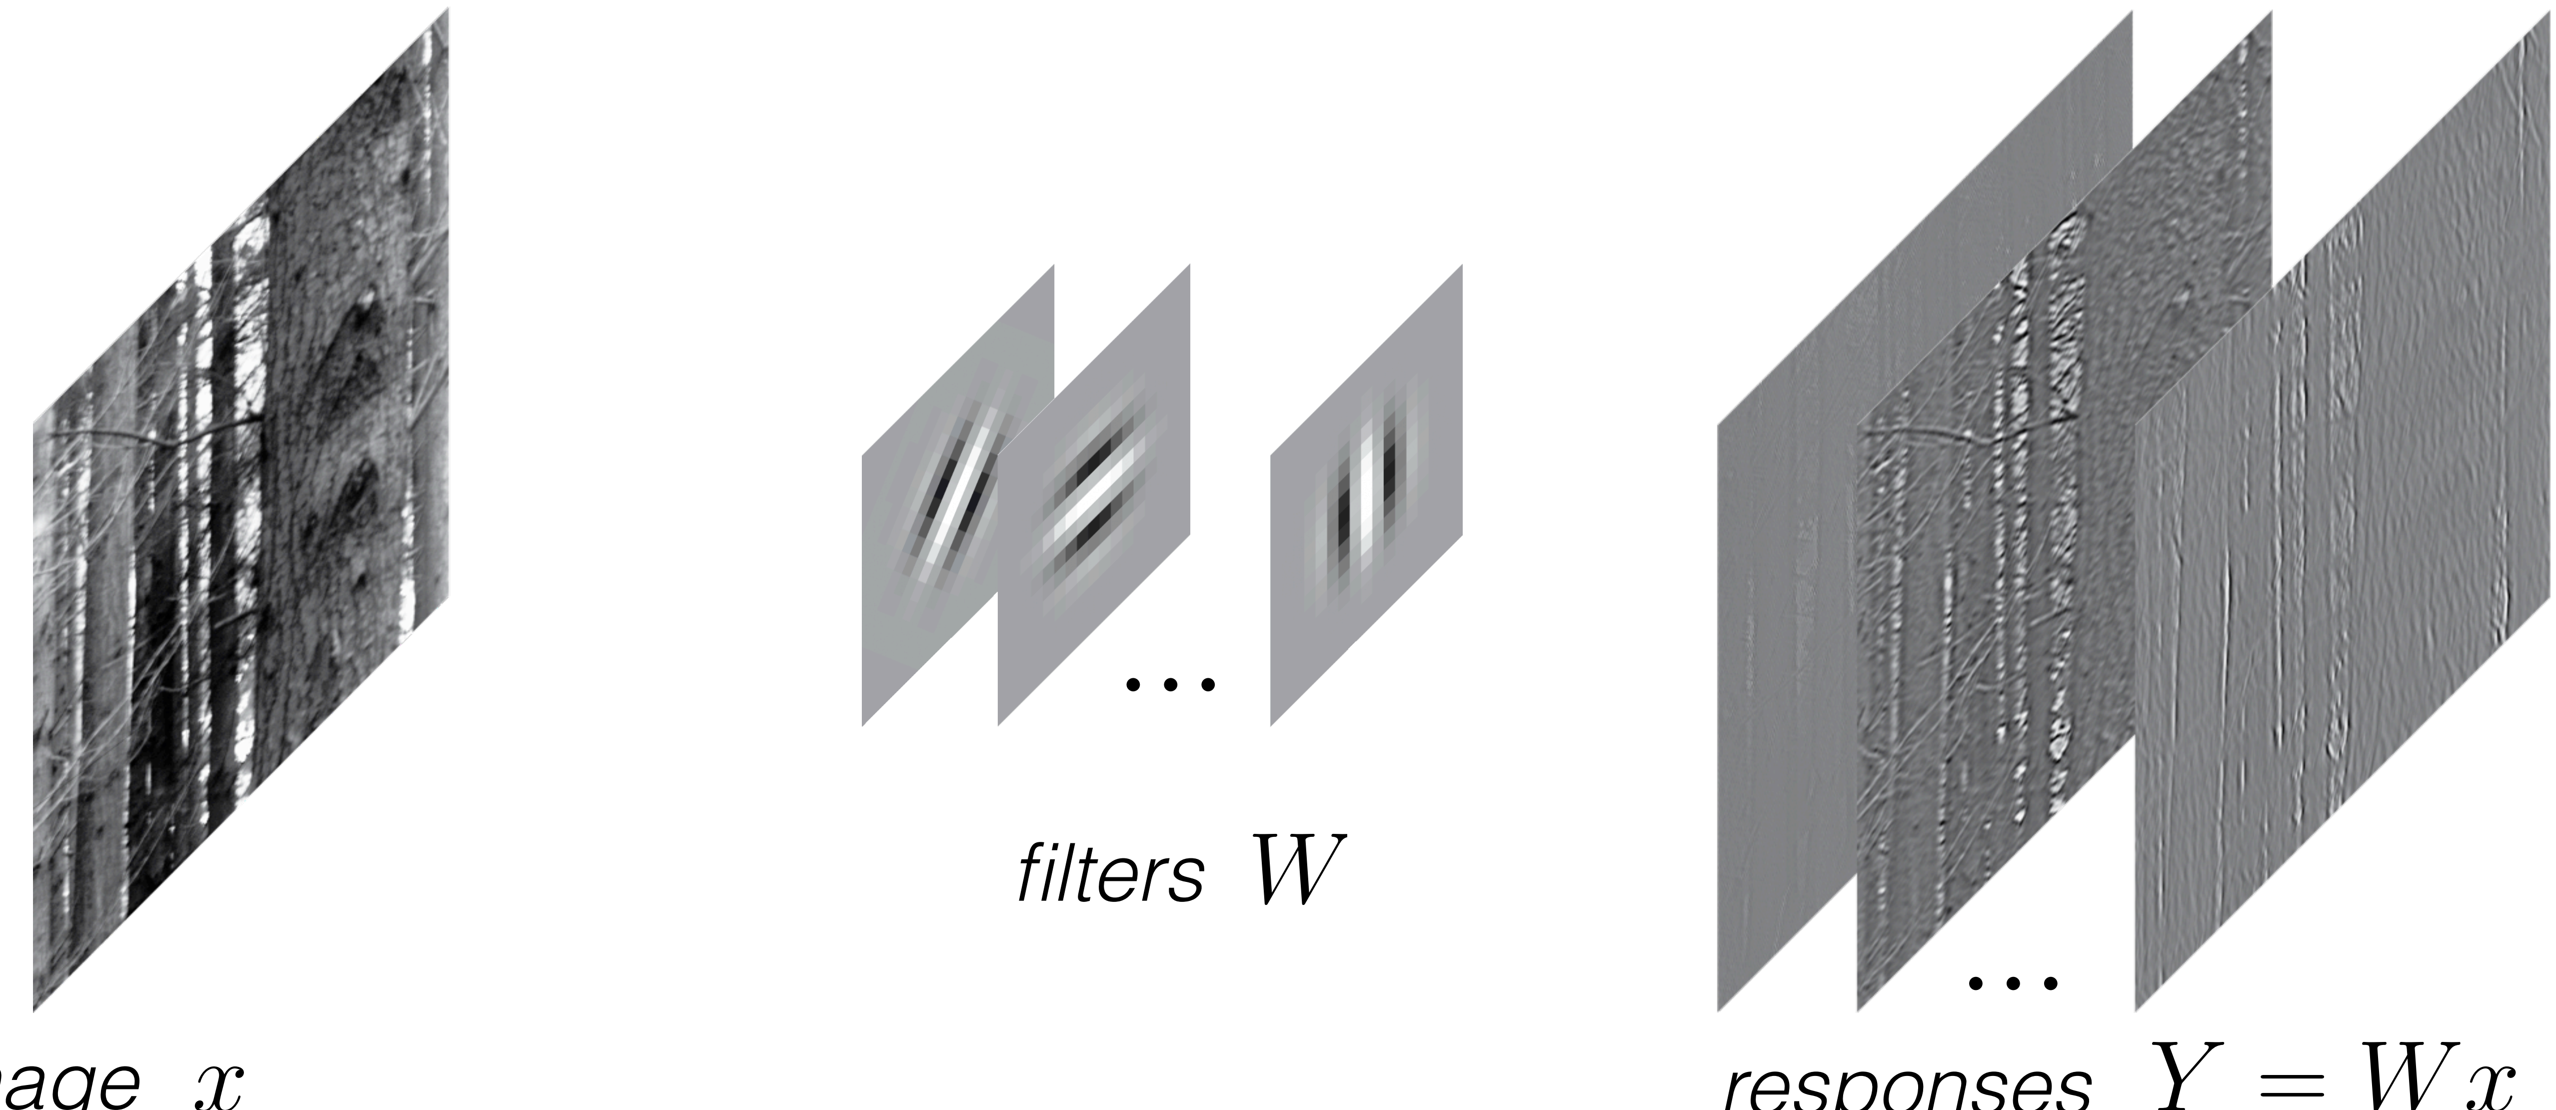

$$E(W) = \sum_n \left\| Y_n \right\|_*$$

local low-dimensionality

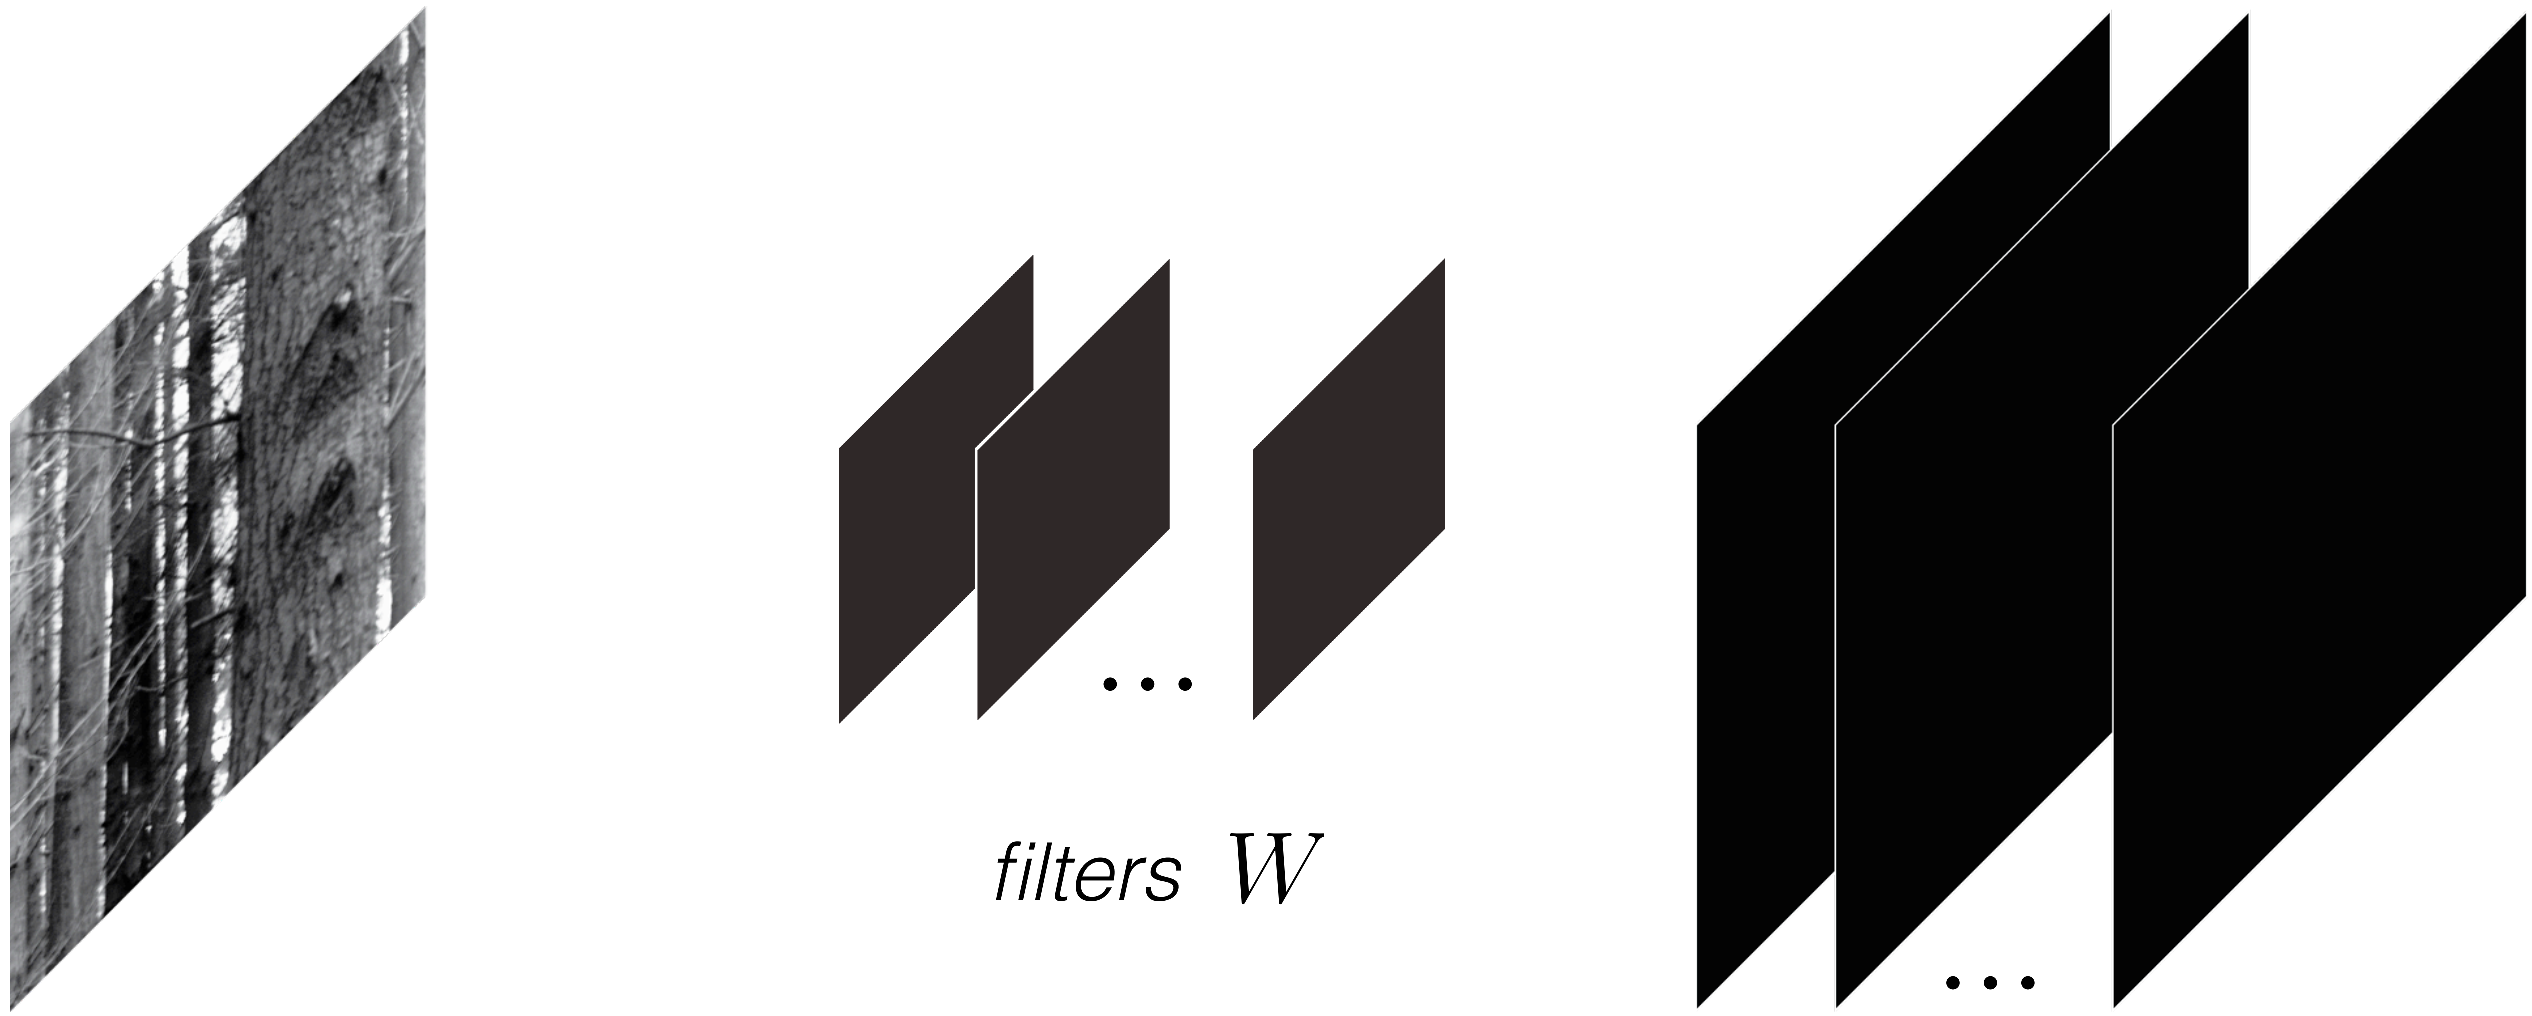

image x

responses $Y = Wx$

$$E(W) = \sum_n \left\| Y_n \right\|_* + \alpha \left\| \hat{x} - x \right\|_2^2$$

local low-dimensionality

information preservation

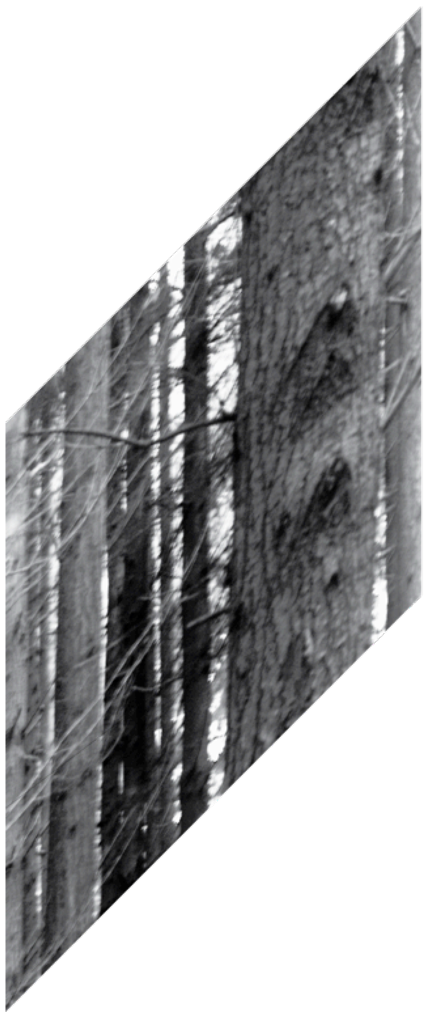

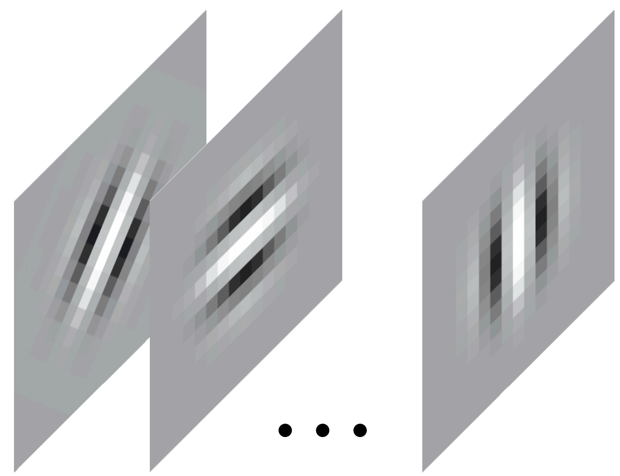

filters W

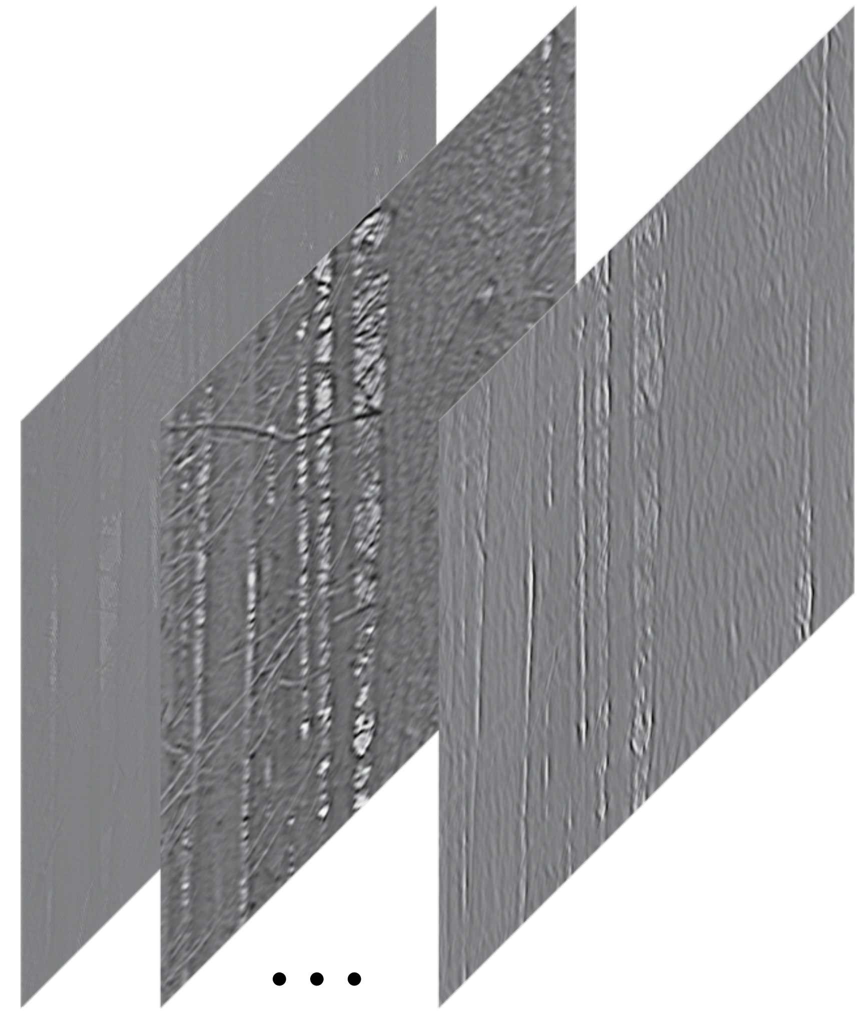

image x

reconstruction $\hat{x} = W^T Y$

responses $Y = Wx$

$$E(W) = \sum_n \left\| Y_n \right\|_* + \alpha \left\| \hat{x} - x \right\|_2^2$$

local low-dimensionality — *information preservation*

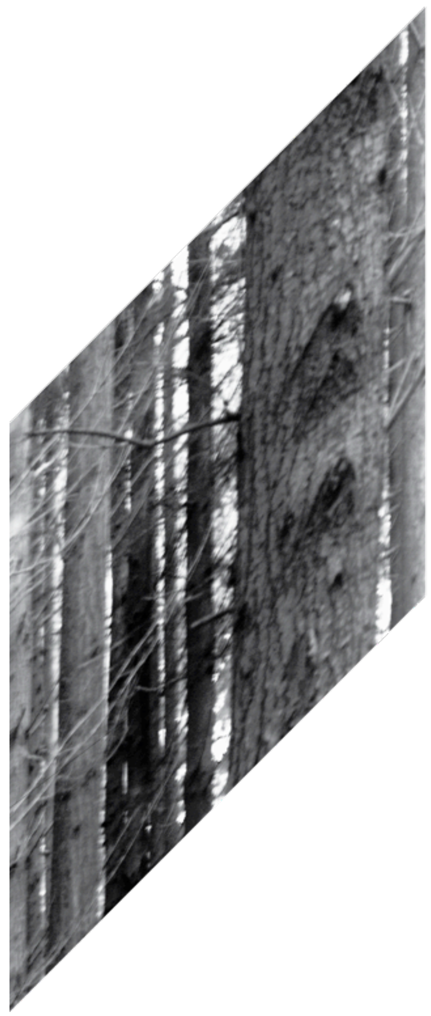

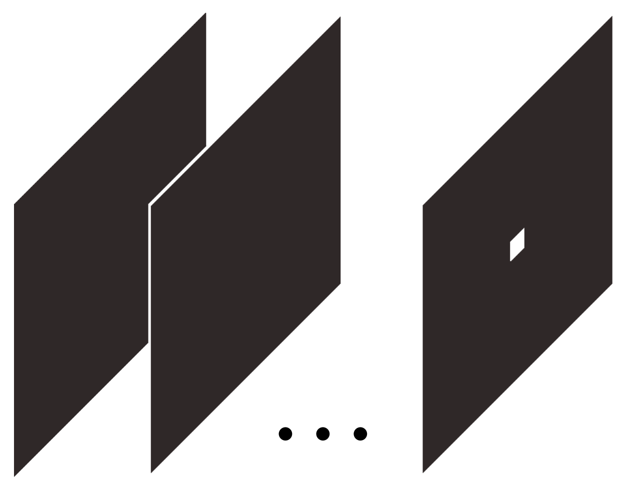

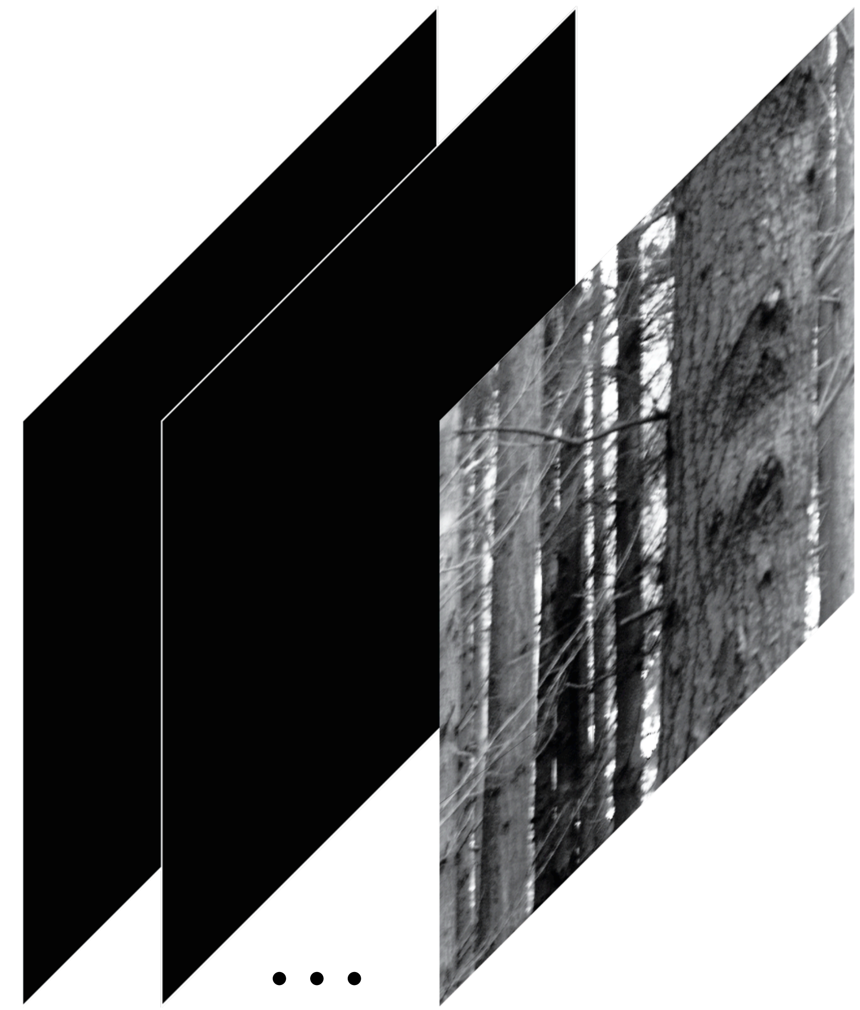

filters W

image x

reconstruction $\hat{x} = W^T Y$

responses $Y = Wx$

$$E(W) = \sum_n \left\| Y_n \right\|_* + \alpha \left\| \hat{x} - x \right\|_2^2$$

local low-dimensionality *information preservation*

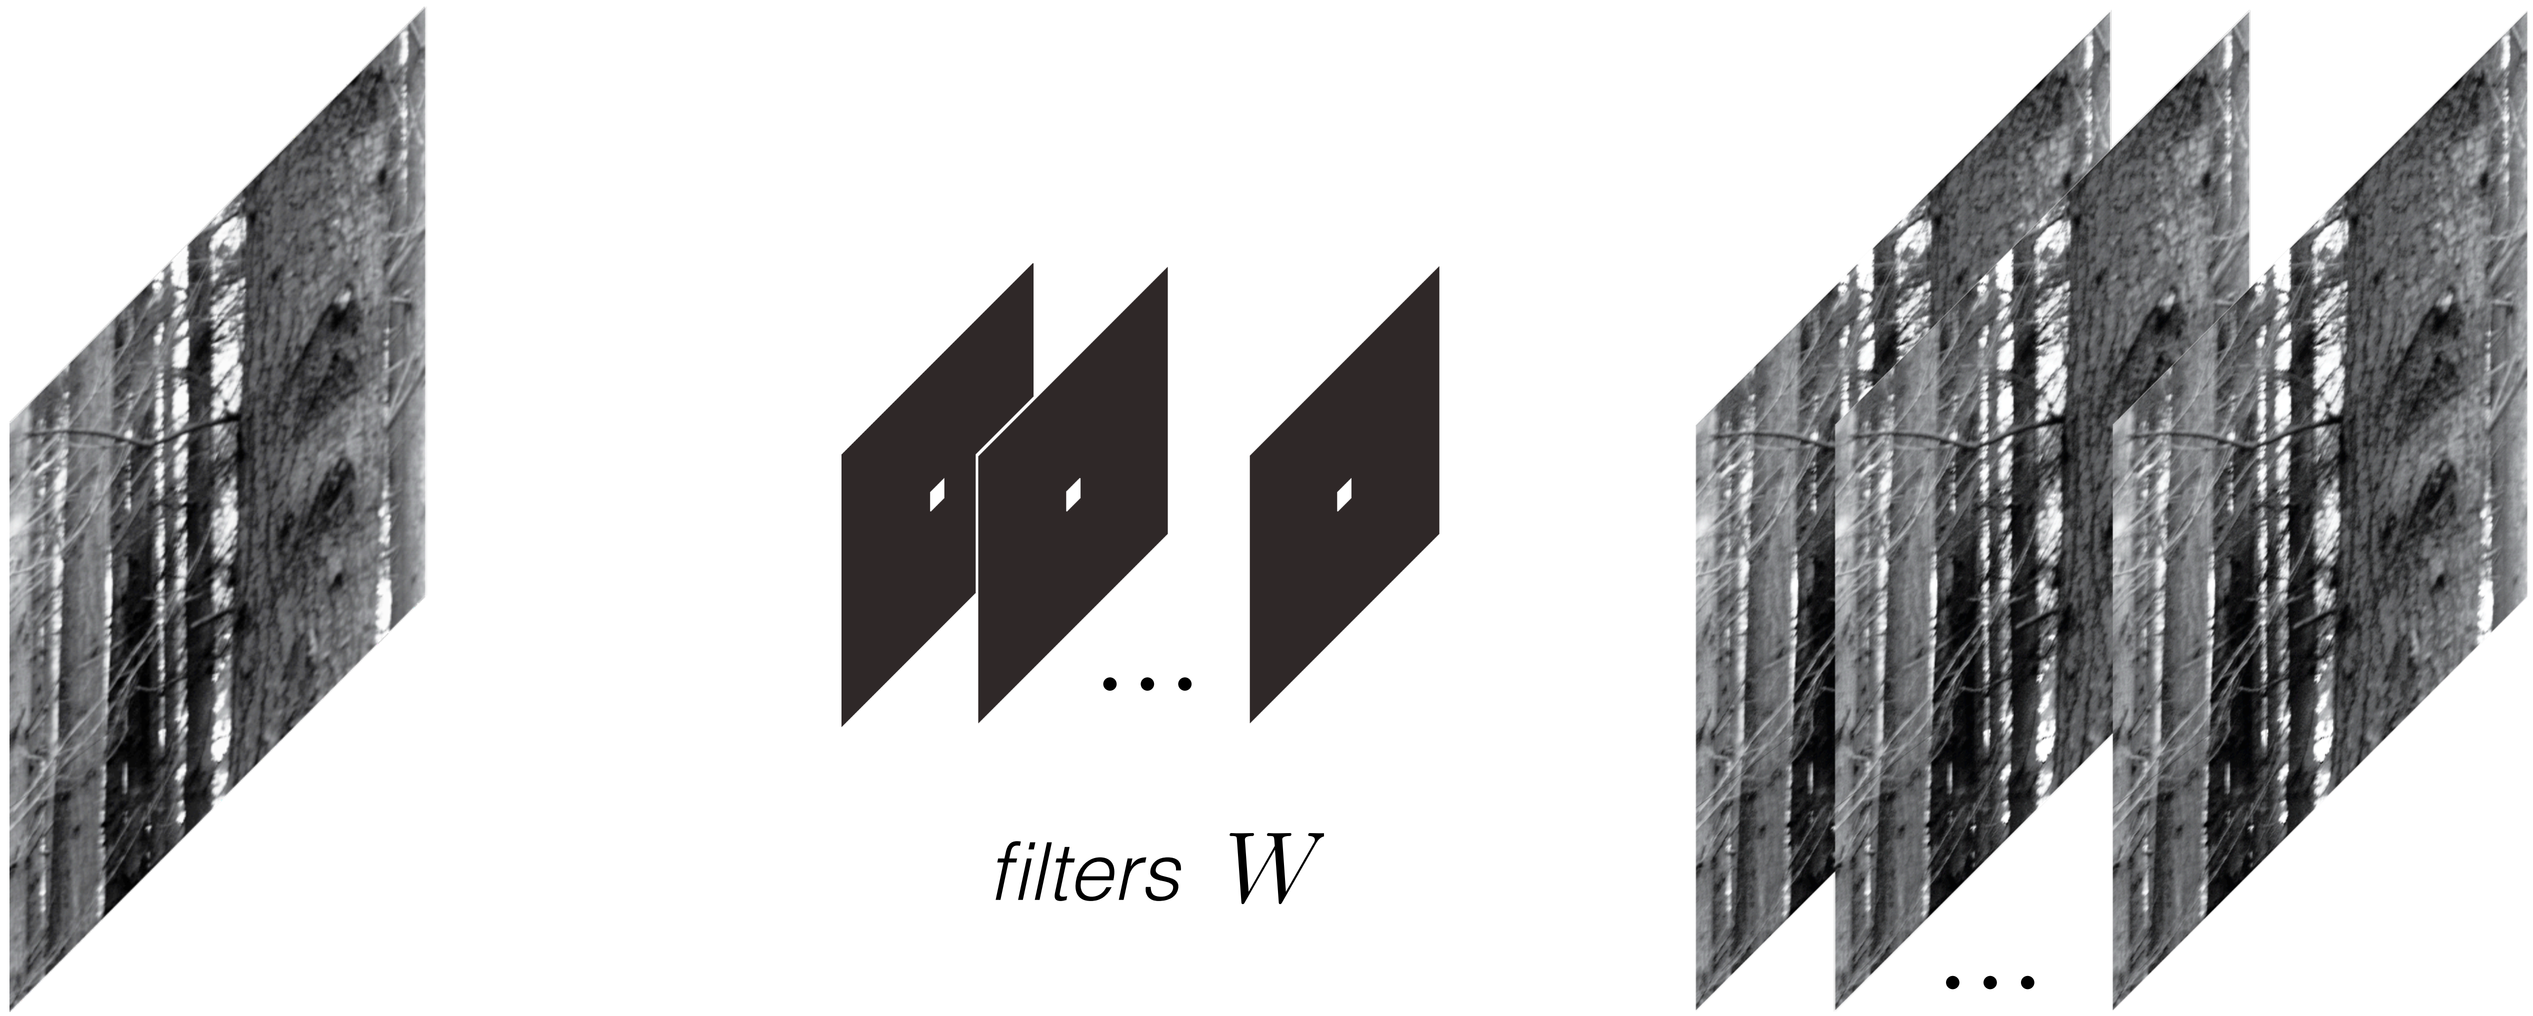

image x

reconstruction $\hat{x} = W^T Y$

responses $Y = Wx$

$$E(W) = \sum_n \left\|Y_n\right\|_* + \alpha \left\|\hat{x} - x\right\|_2^2 - \beta \left\|Y\right\|_*$$

local low-dimensionality | *information preservation* | *global high-dimensionality*

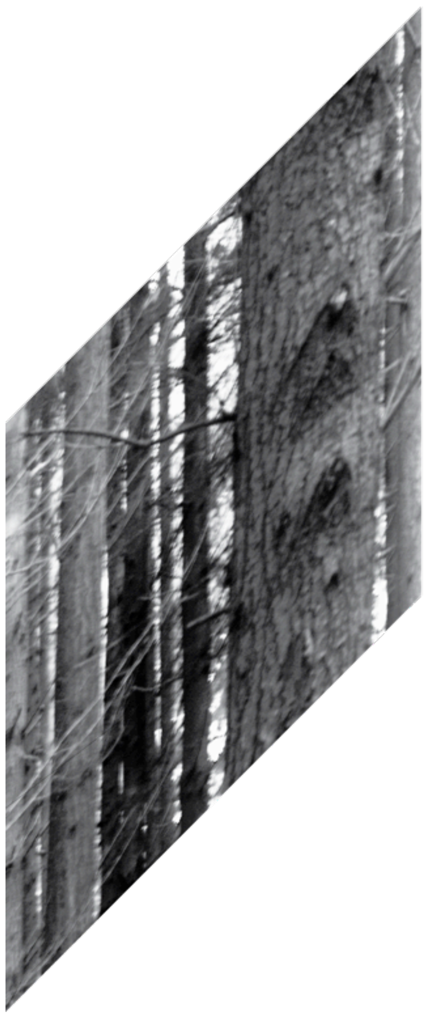

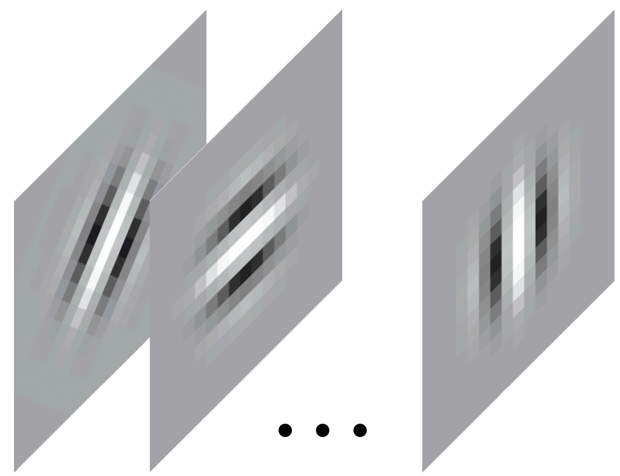

filters W

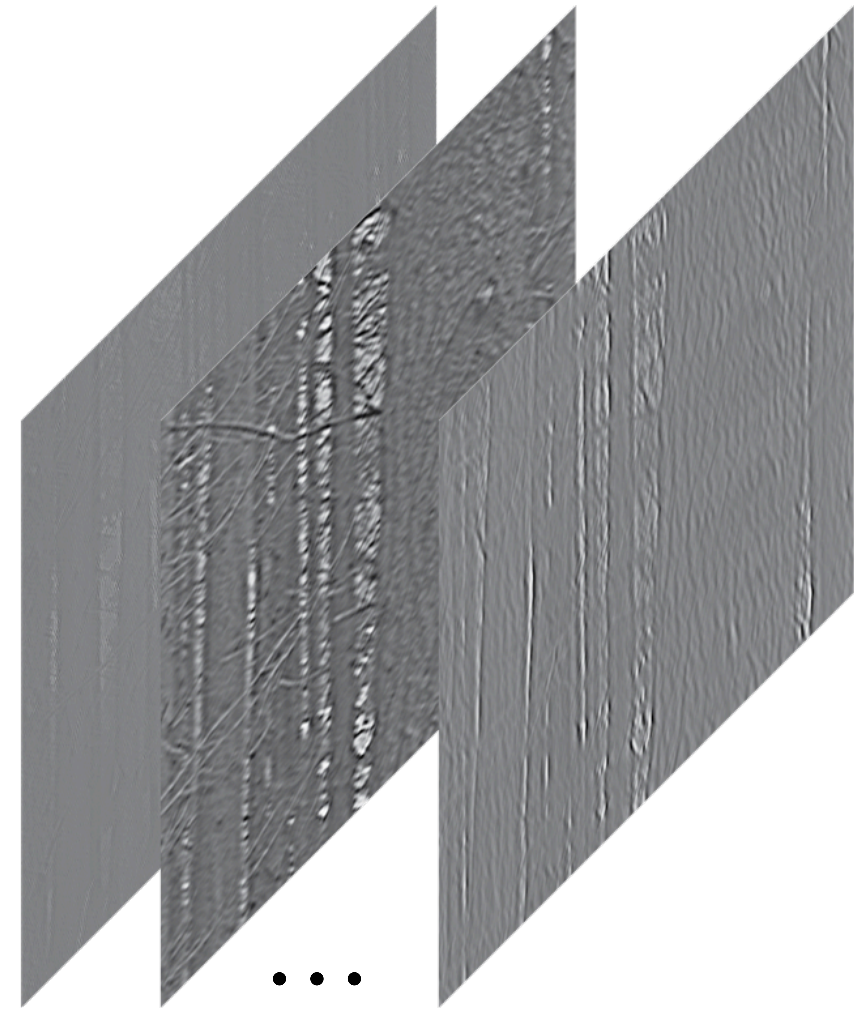

image x

reconstruction $\hat{x} = W^T Y$

responses $Y = Wx$

$$E(W) = \sum_n \left\| Y_n \right\|_* + \alpha \left\| \hat{x} - x \right\|_2^2 - \beta \left\| Y \right\|_*$$

local low-dimensionality *information preservation* *global high-dimensionality*

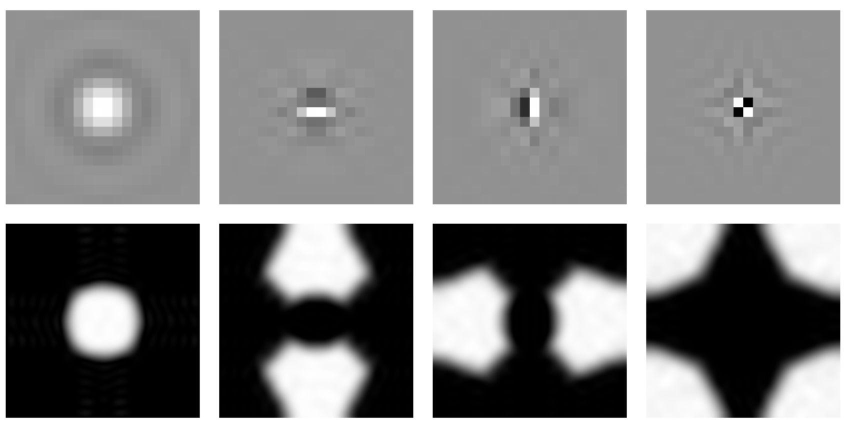

spatial filters

Fourier magnitudes

$$E(W) = \sum_n \left\|Y_n\right\|_* + \alpha\left\|\hat{x} - x\right\|_2^2 - \beta\left\|Y\right\|_*$$

local low-dimensionality | *information preservation* | *global high-dimensionality*

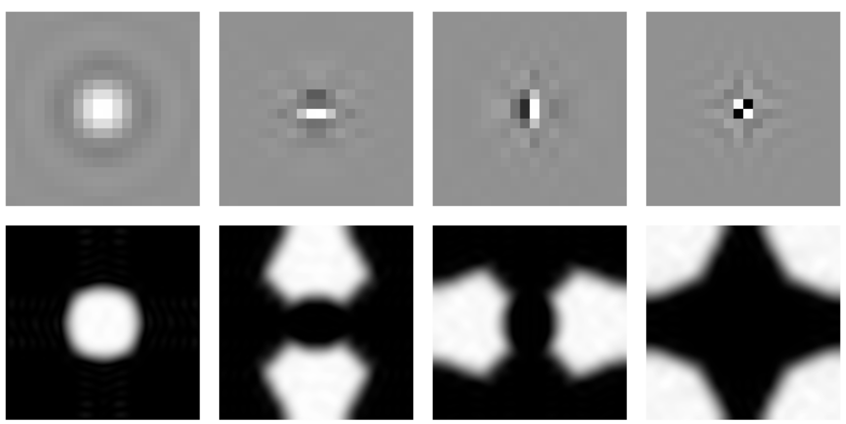

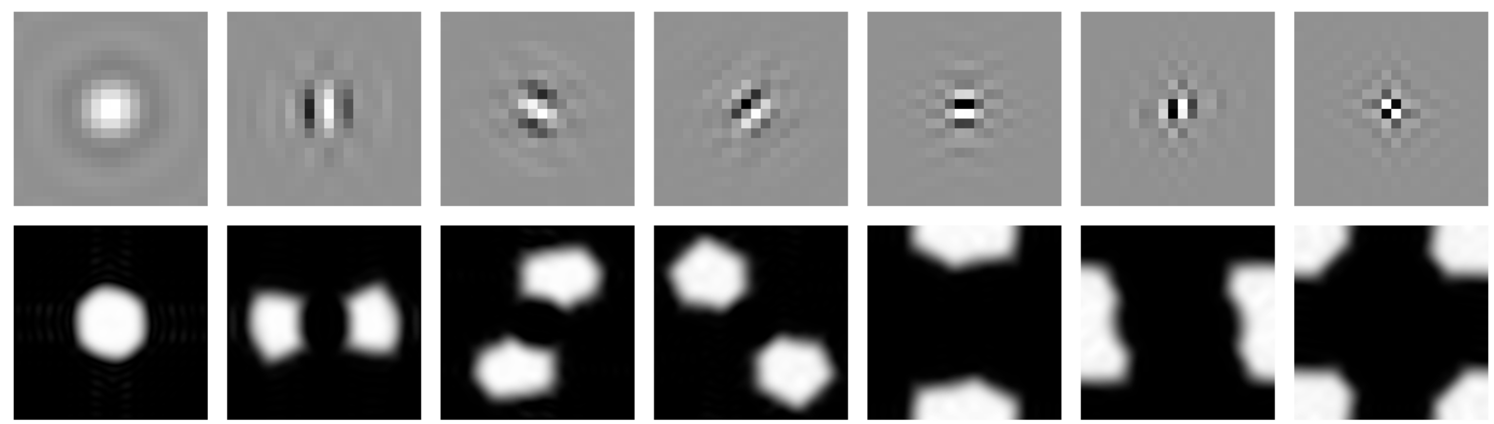

$$E(W) = \sum_n \left\| Y_n \right\|_* + \alpha \left\| \hat{x} - x \right\|_2^2 - \beta \left\| Y \right\|_*$$

local low-dimensionality *information preservation* *global high-dimensionality*

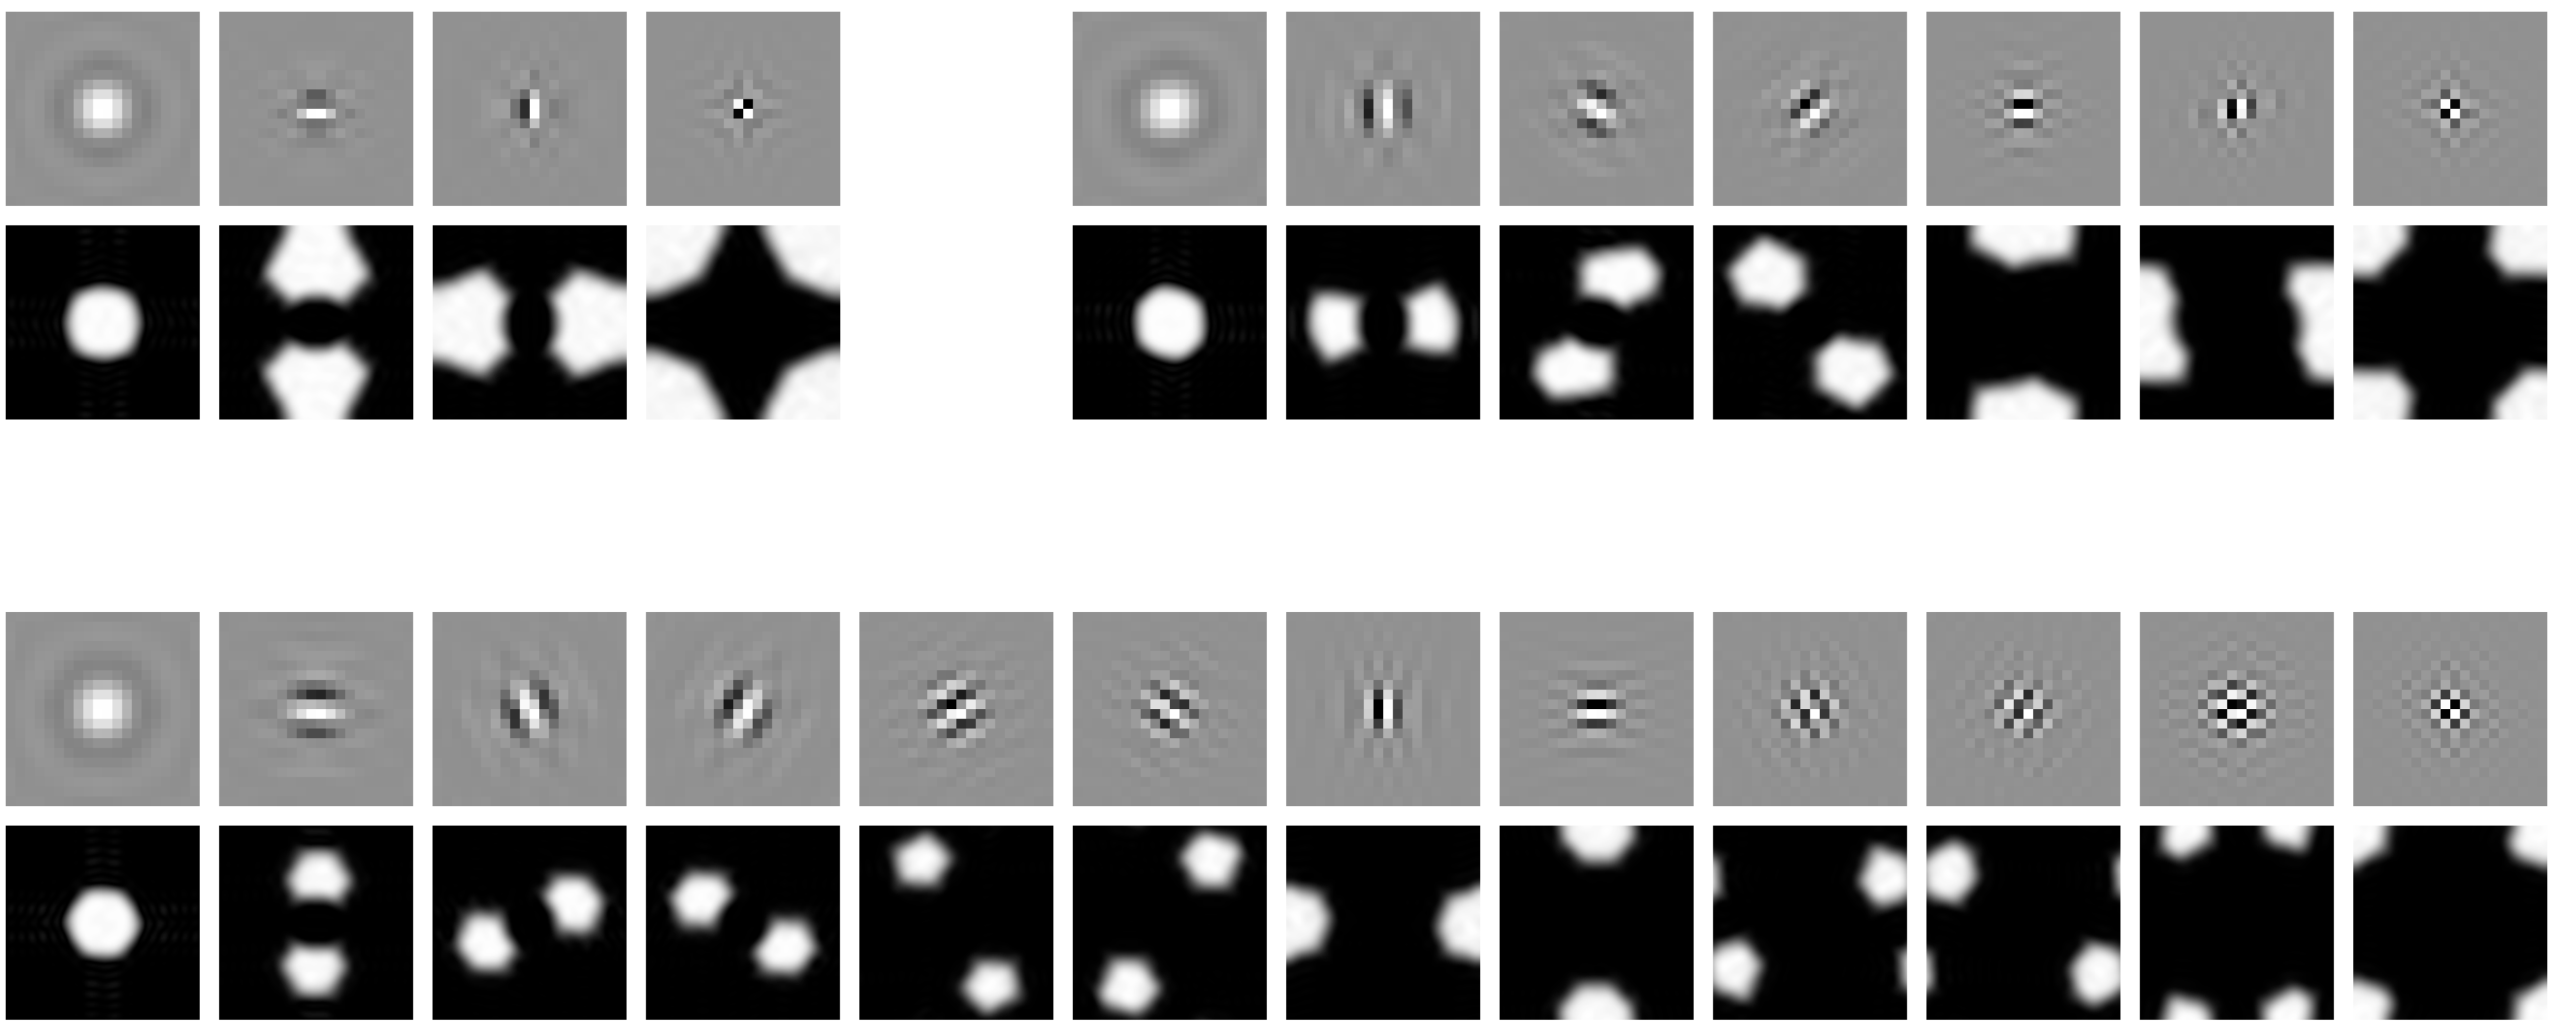

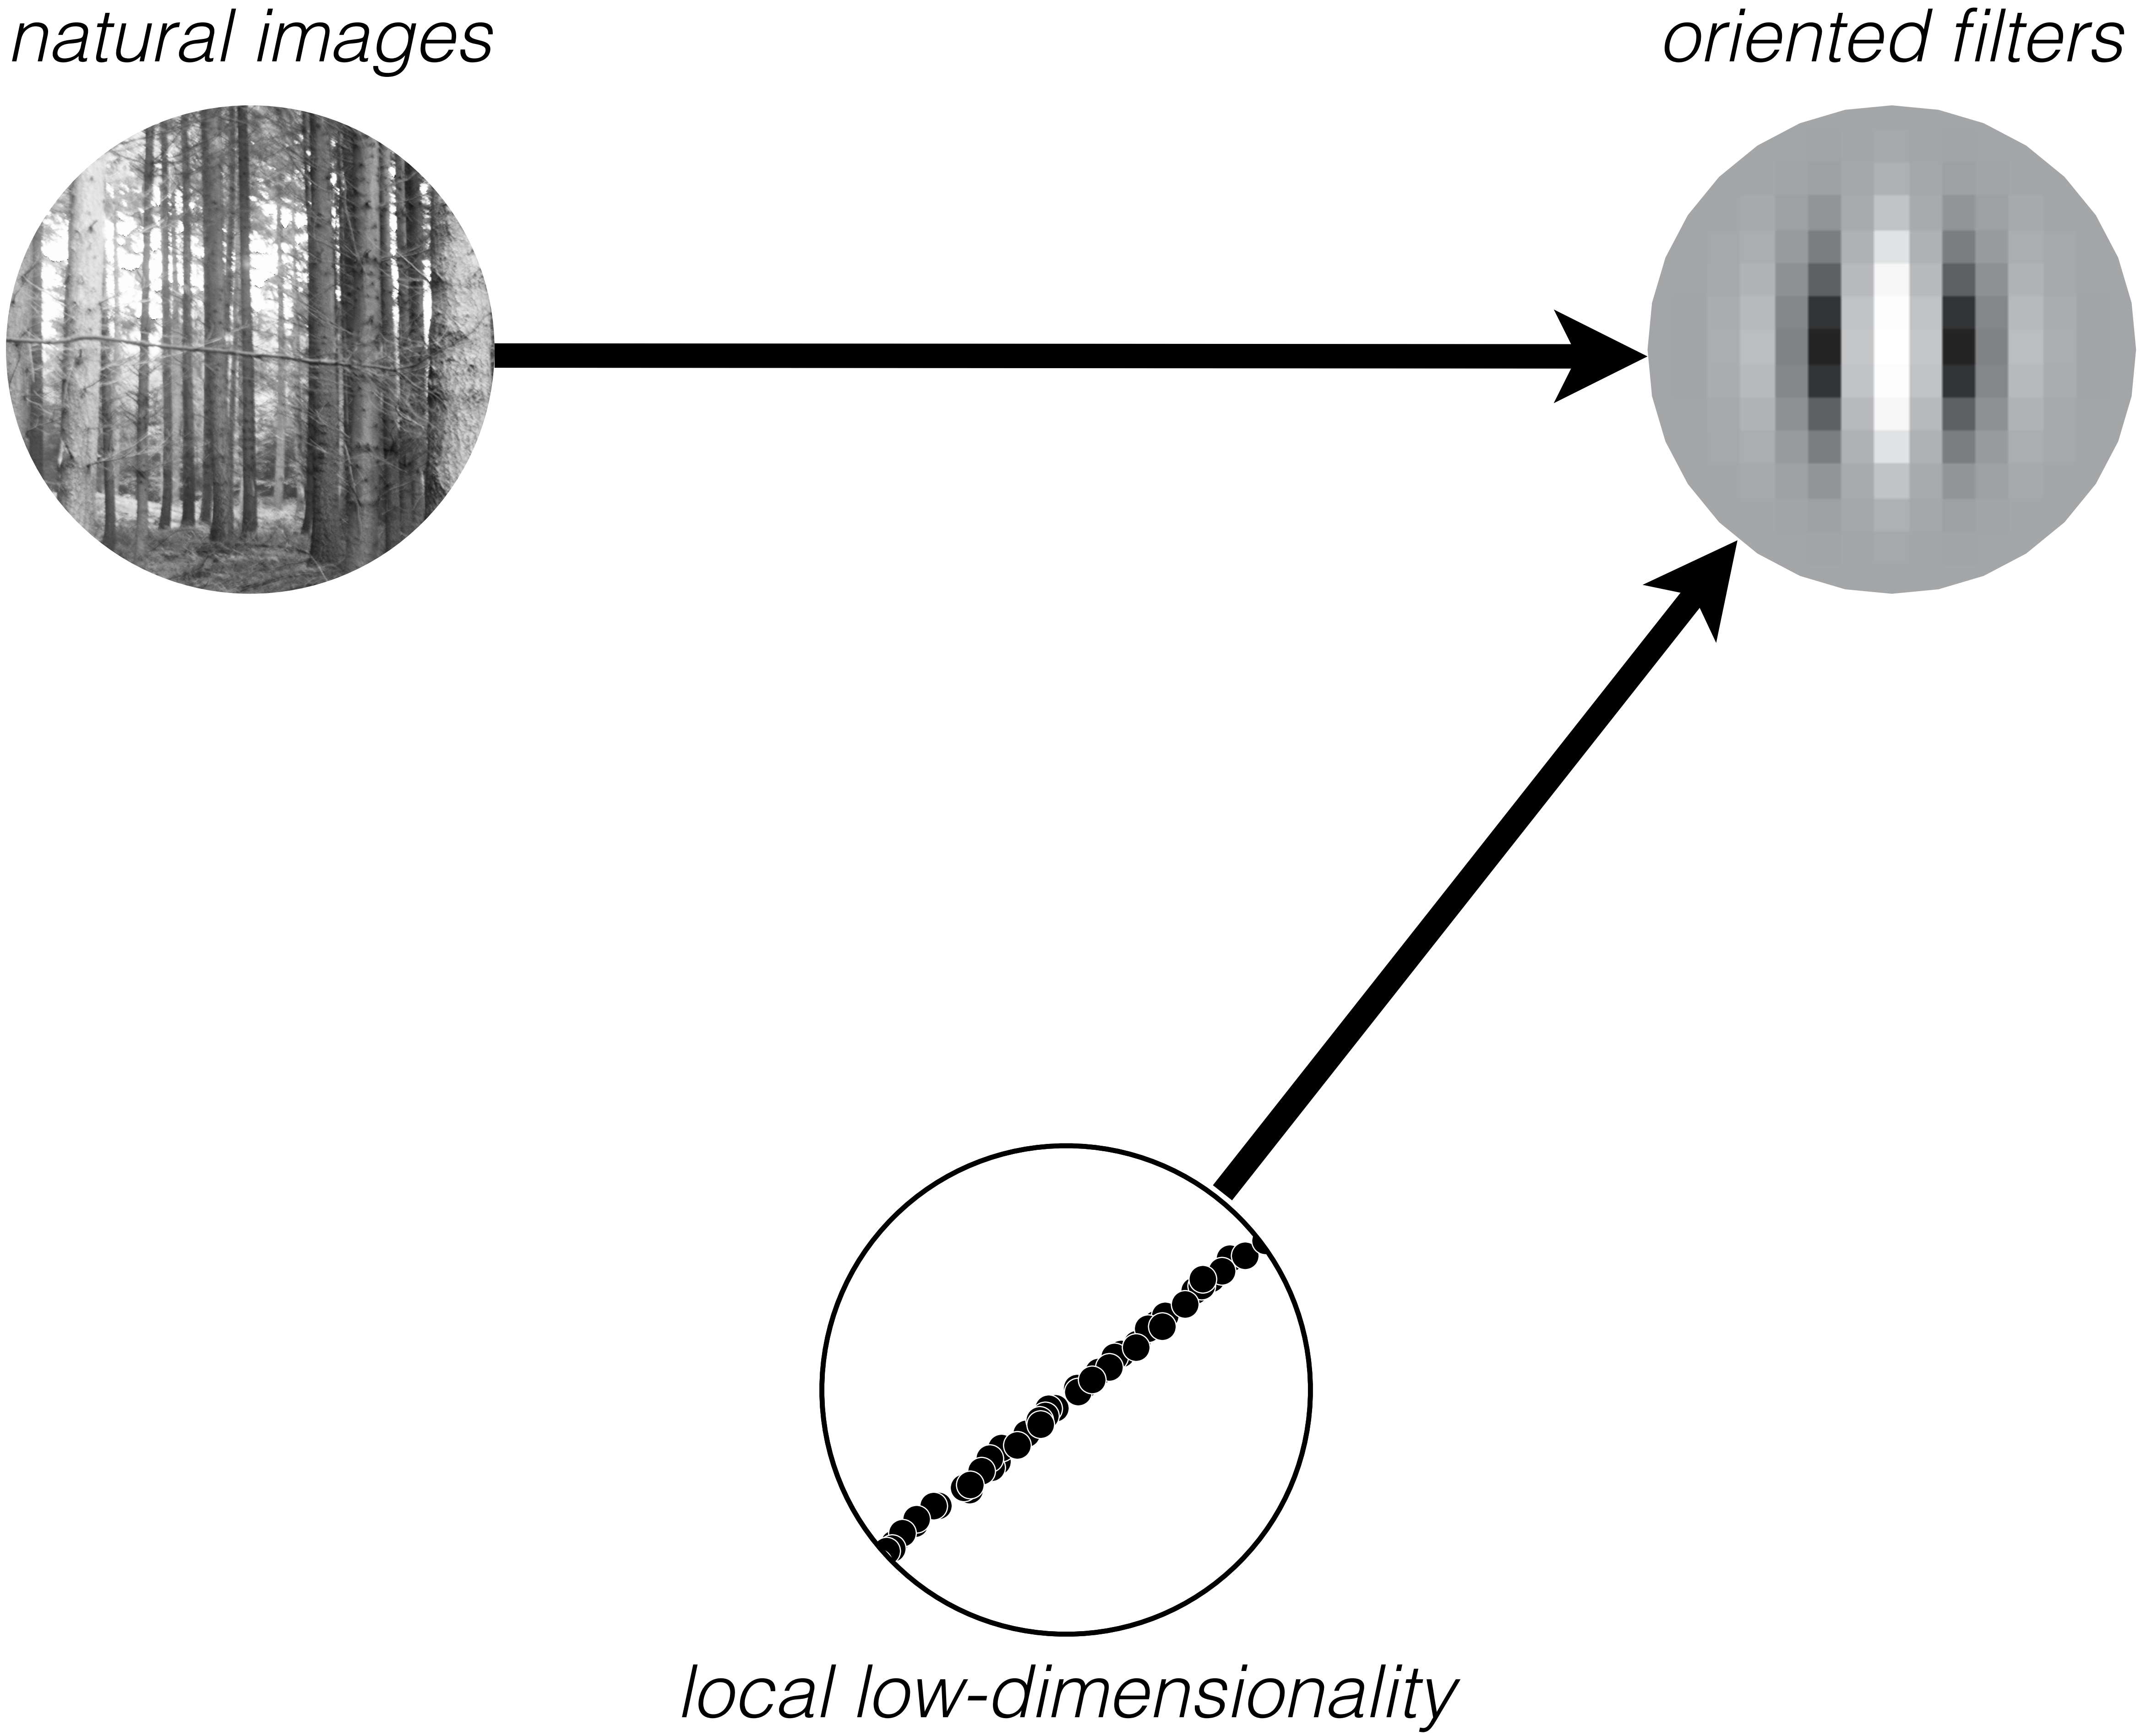
natural images
oriented filters
local low-dimensionality

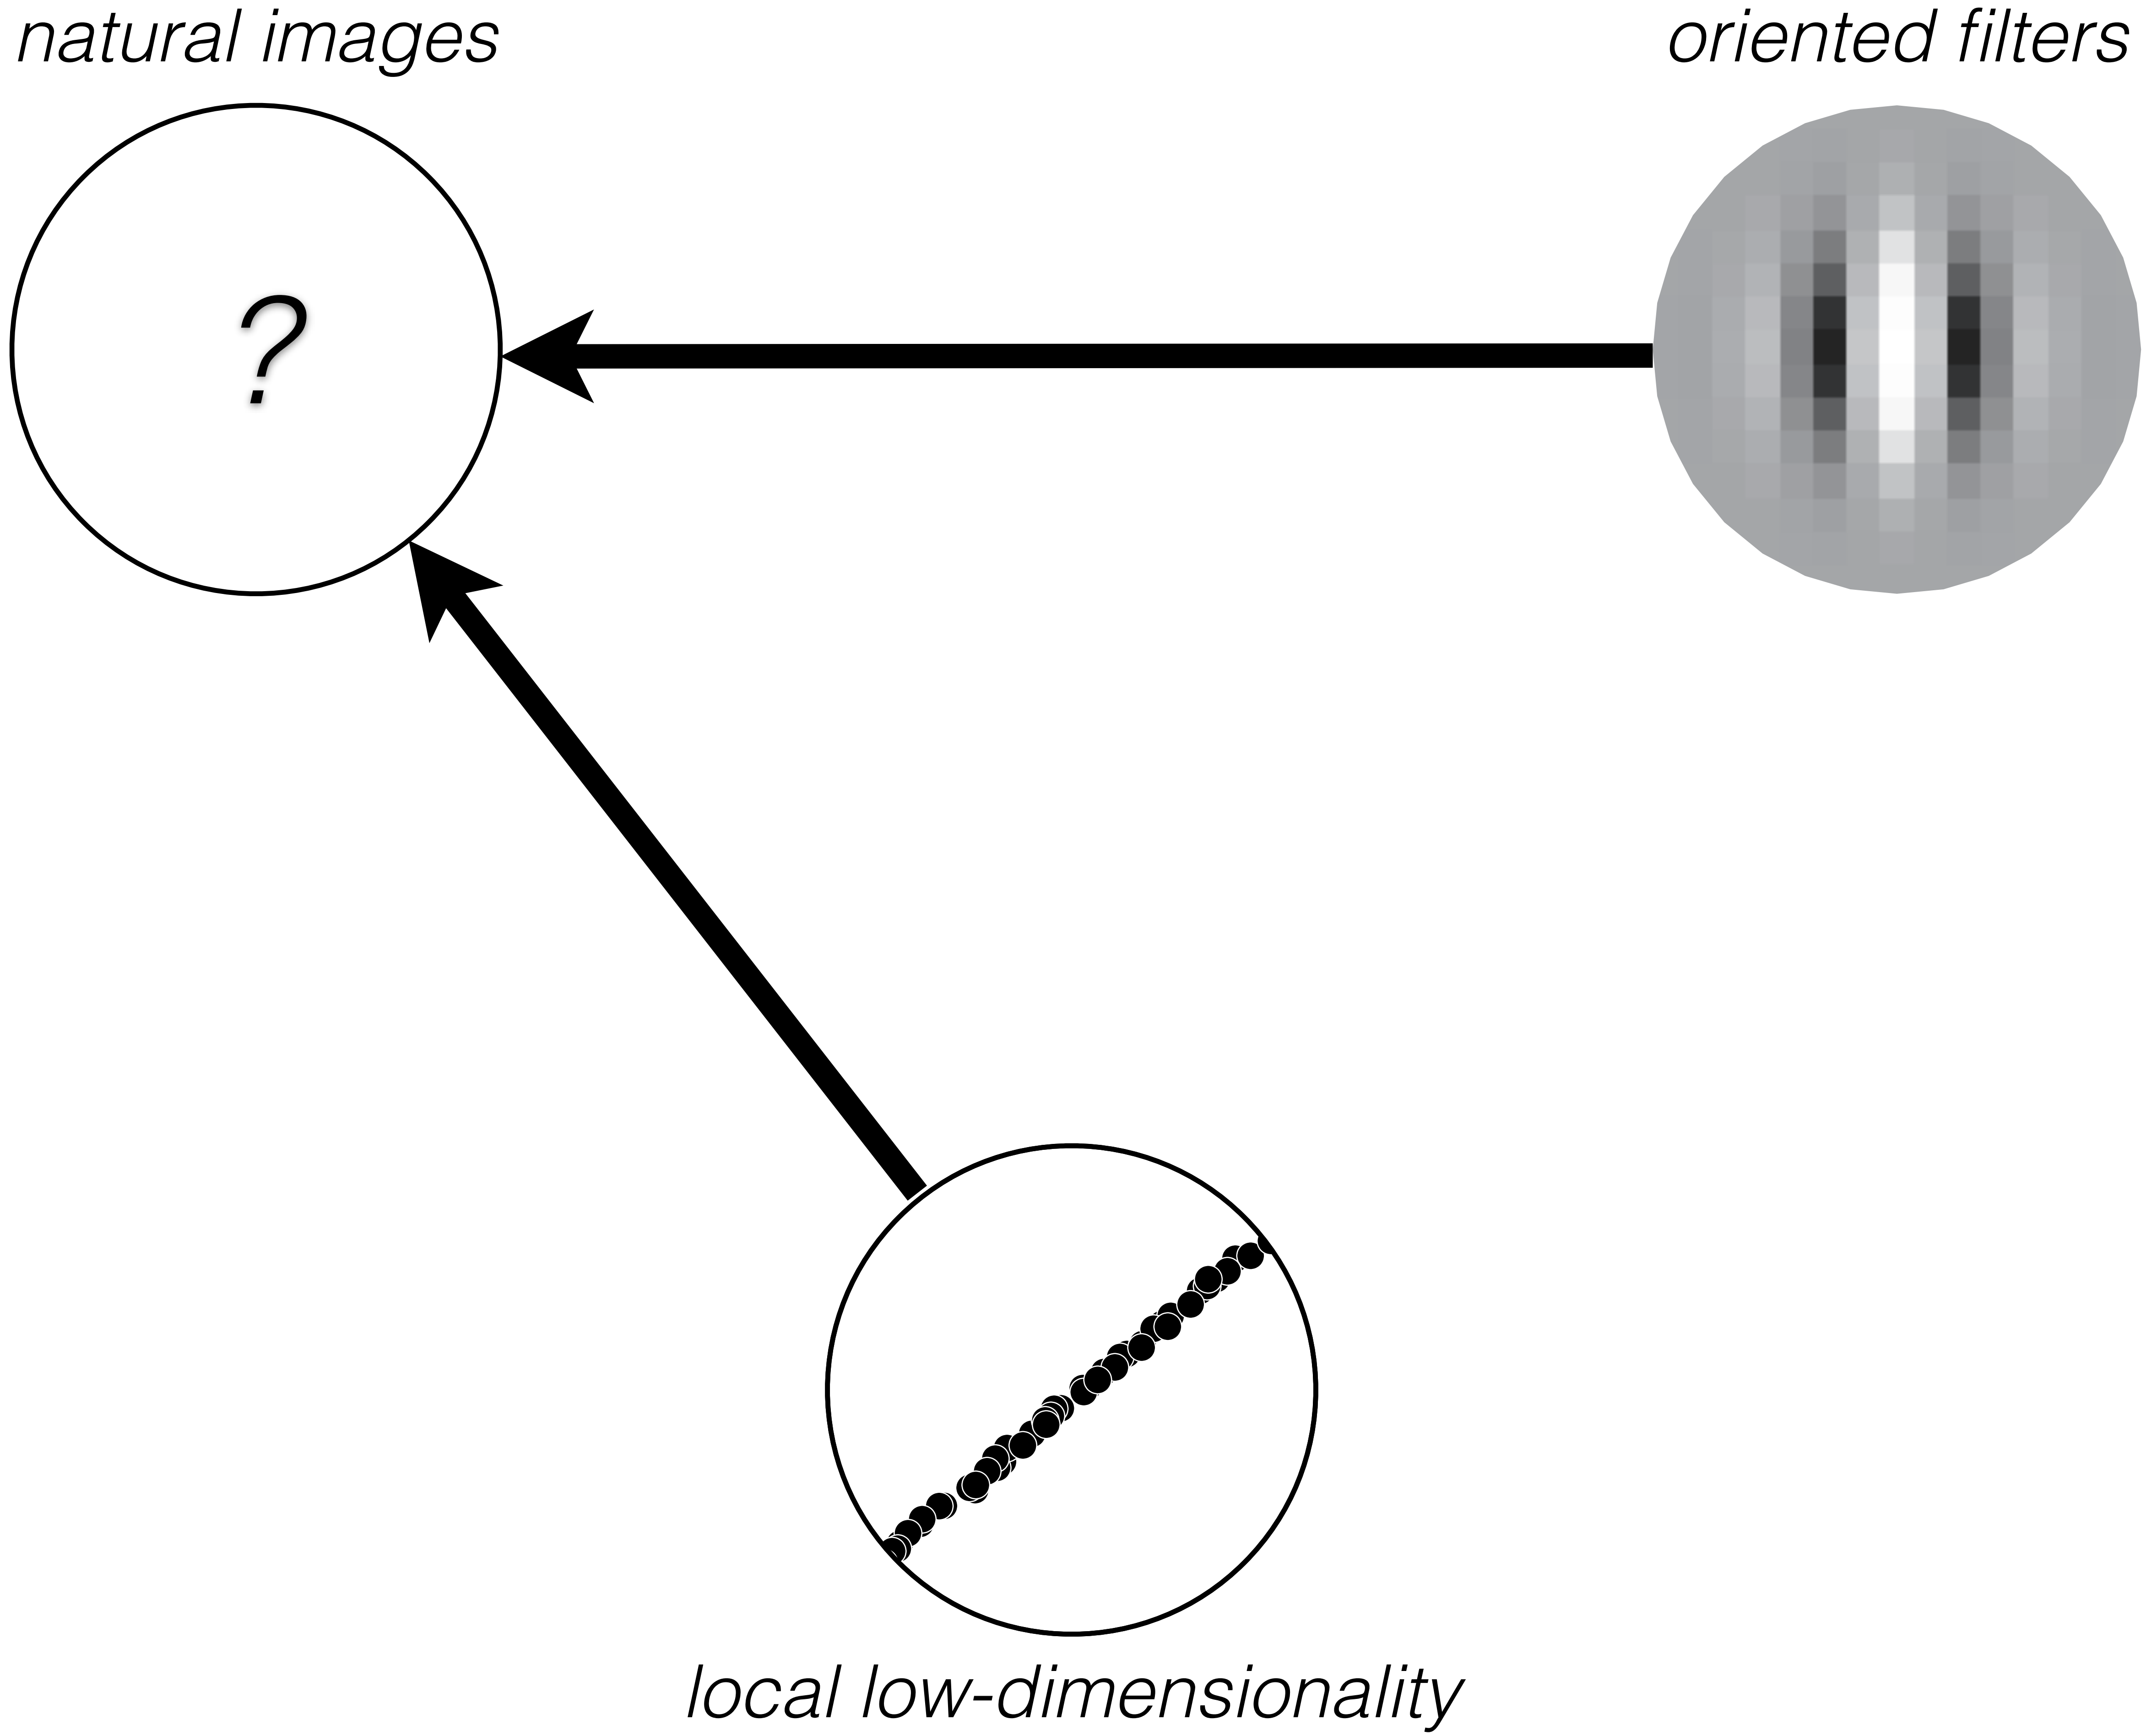
natural images
oriented filters
?
local low-dimensionality

$$\phi(x)_n = Y_n^T Y_n$$

local covariance

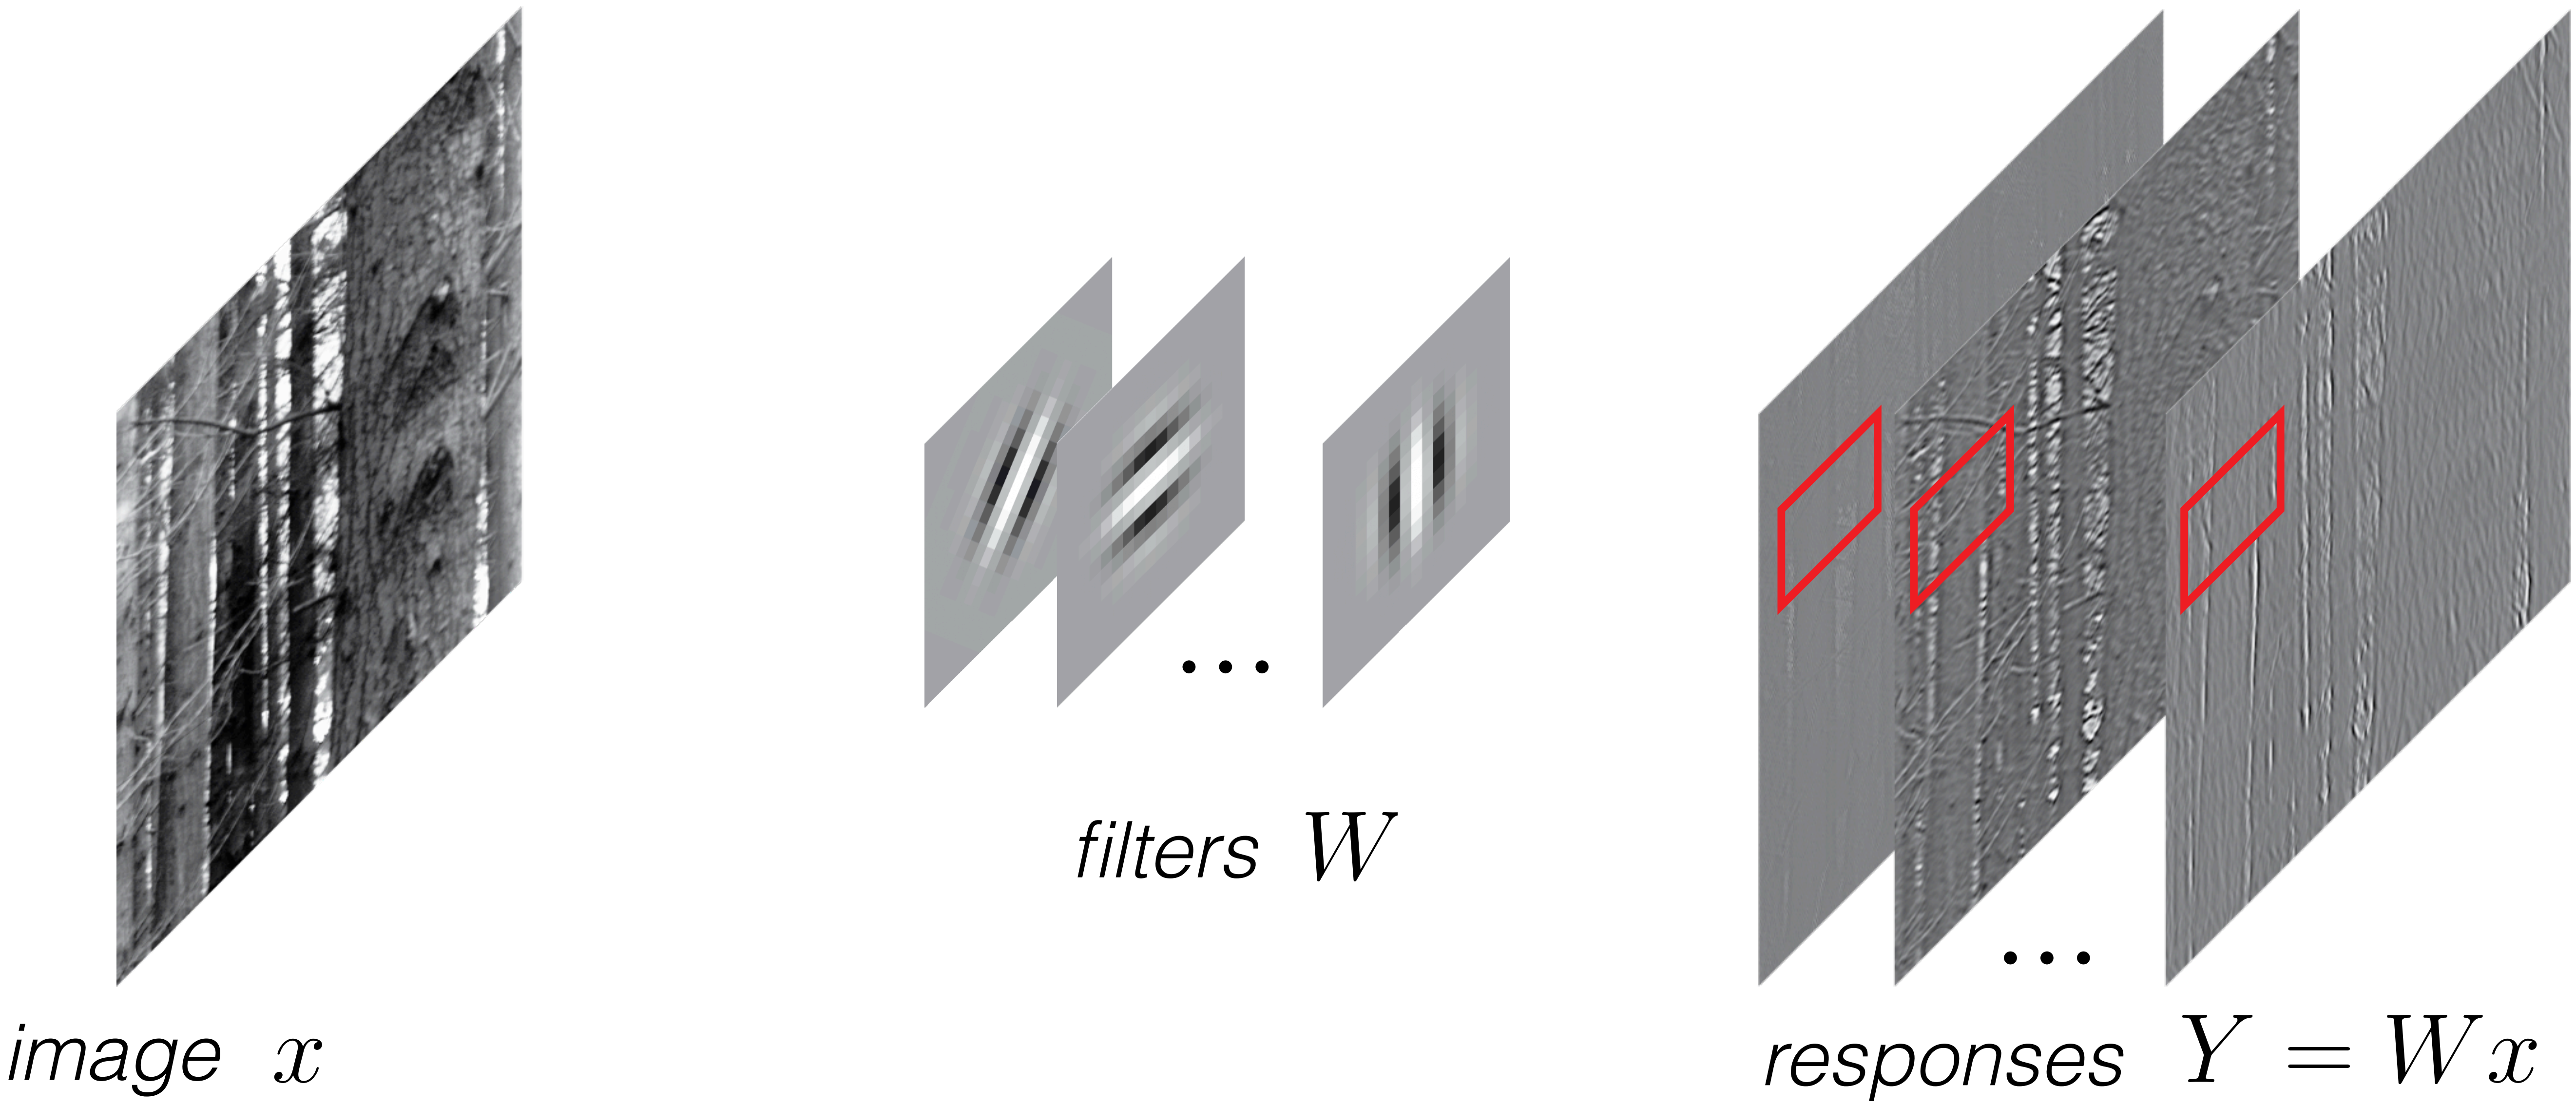

$$\phi(x)_n = Y_n^T Y_n$$

local covariance

$$\phi(x) = [\phi(x)_n]_n$$

covariance map

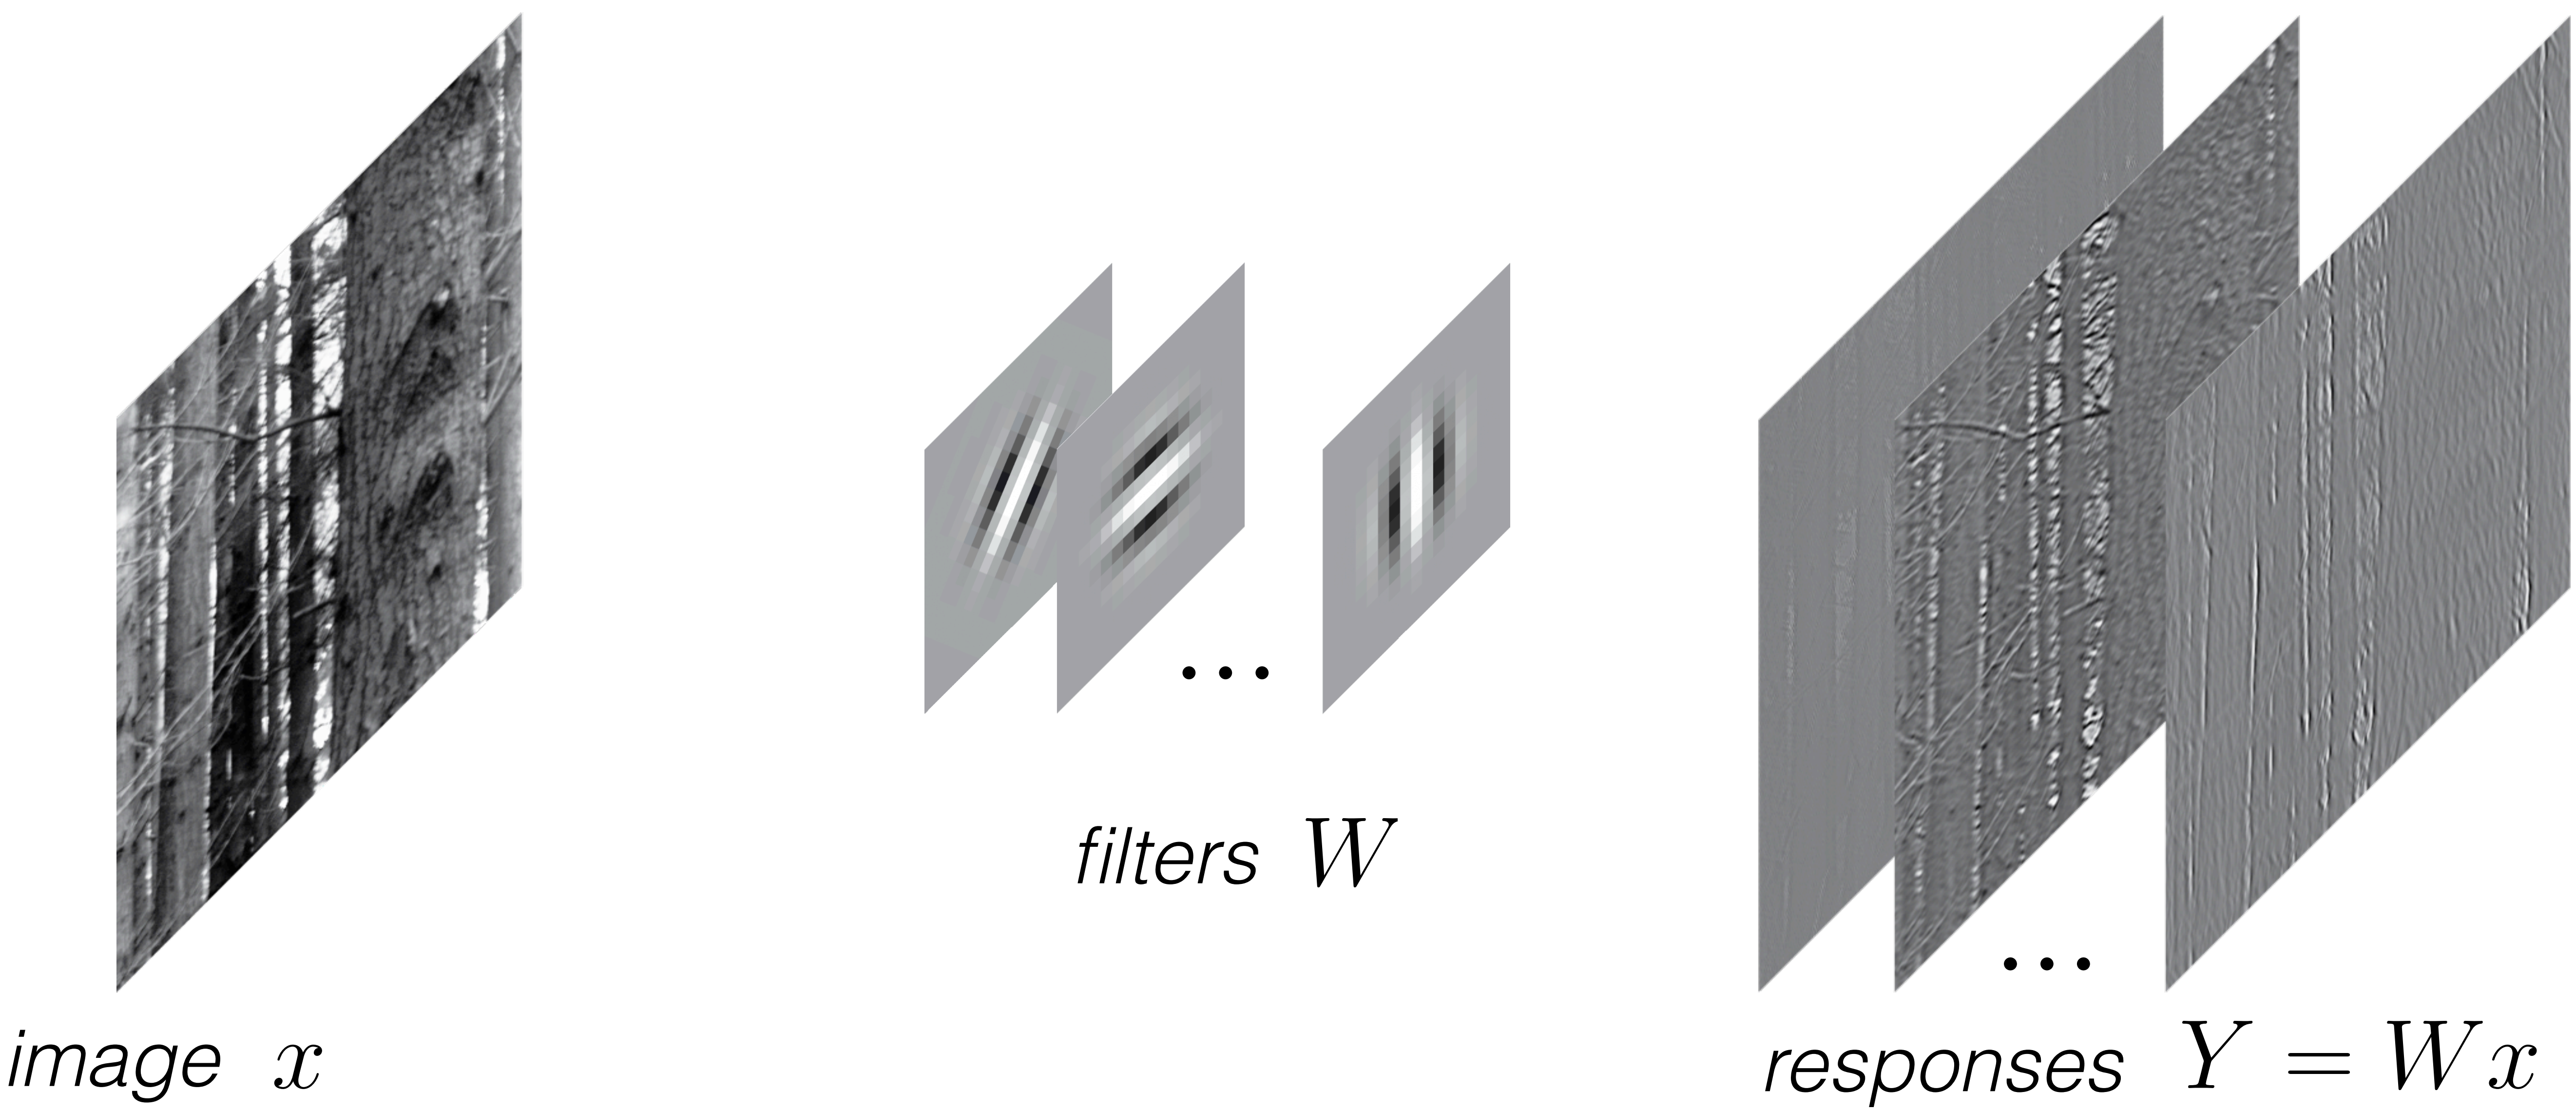

target image

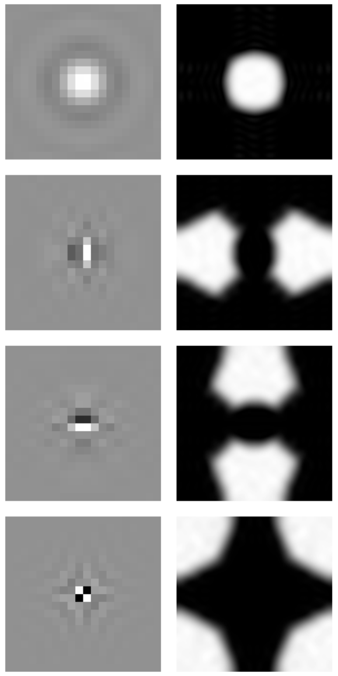

filters

target image

white noise

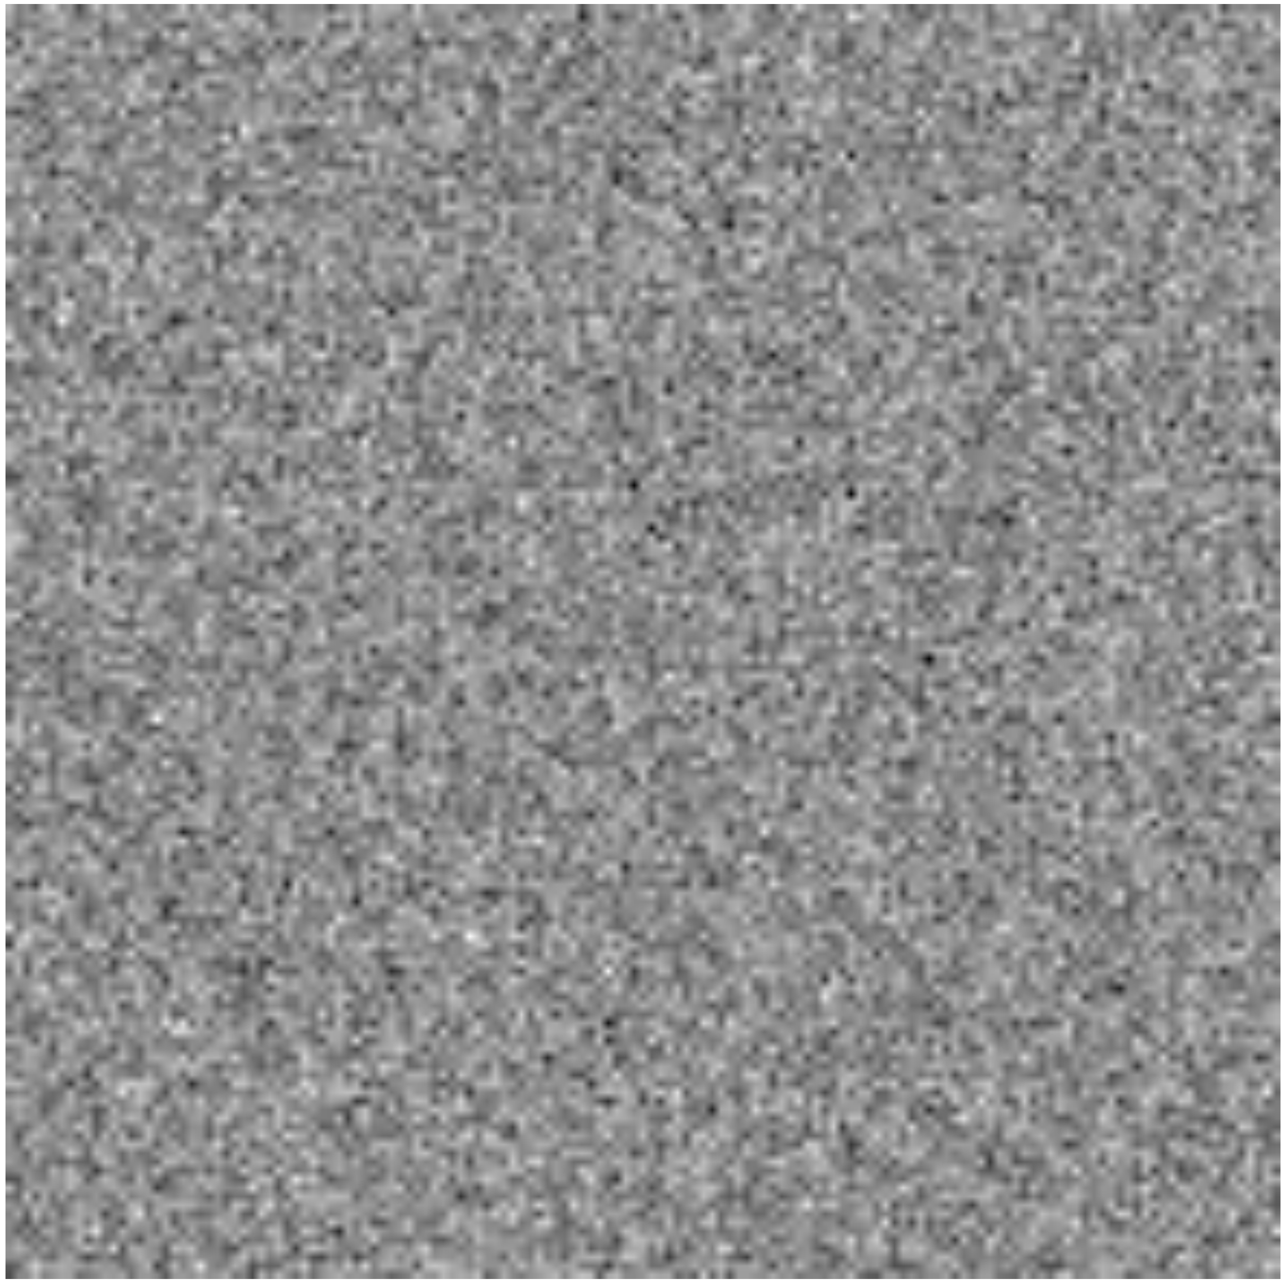

target image | *synthesized image from covariance map*

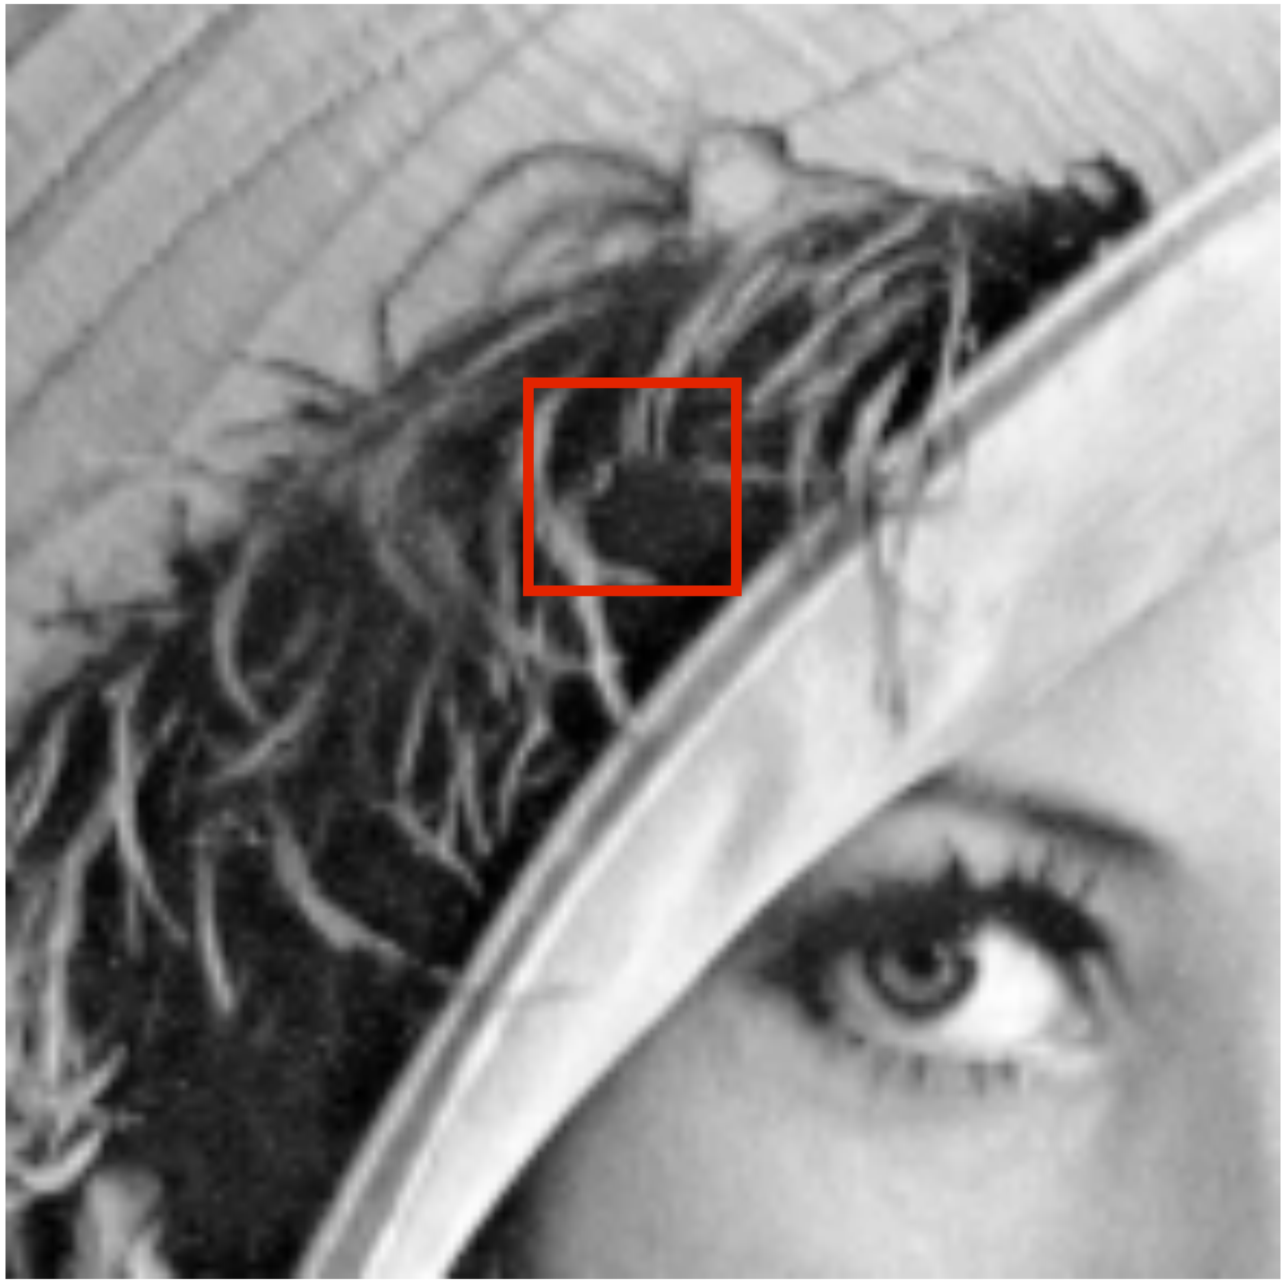

$$\|\Delta x\|_2 \,/\, \|x\|_2 = 11.1\%$$

24 x 24 *neighborhoods*

4 x 4 *subsampling:* 0.6x *overcomplete*

target image

+ *white noise*

synthesized image

from covariance map

$$\|\Delta x\|_2 \,/\, \|x\|_2 = 11.1\%$$

synthesized image from variance map

synthesized image from covariance map

filters

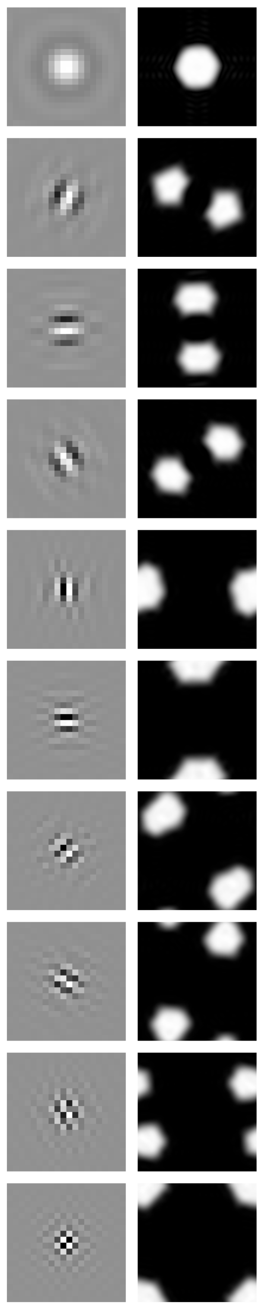

$$\|\Delta x\|_2 \,/\, \|x\|_2 = 11.1\%$$

synthesized image from variance map

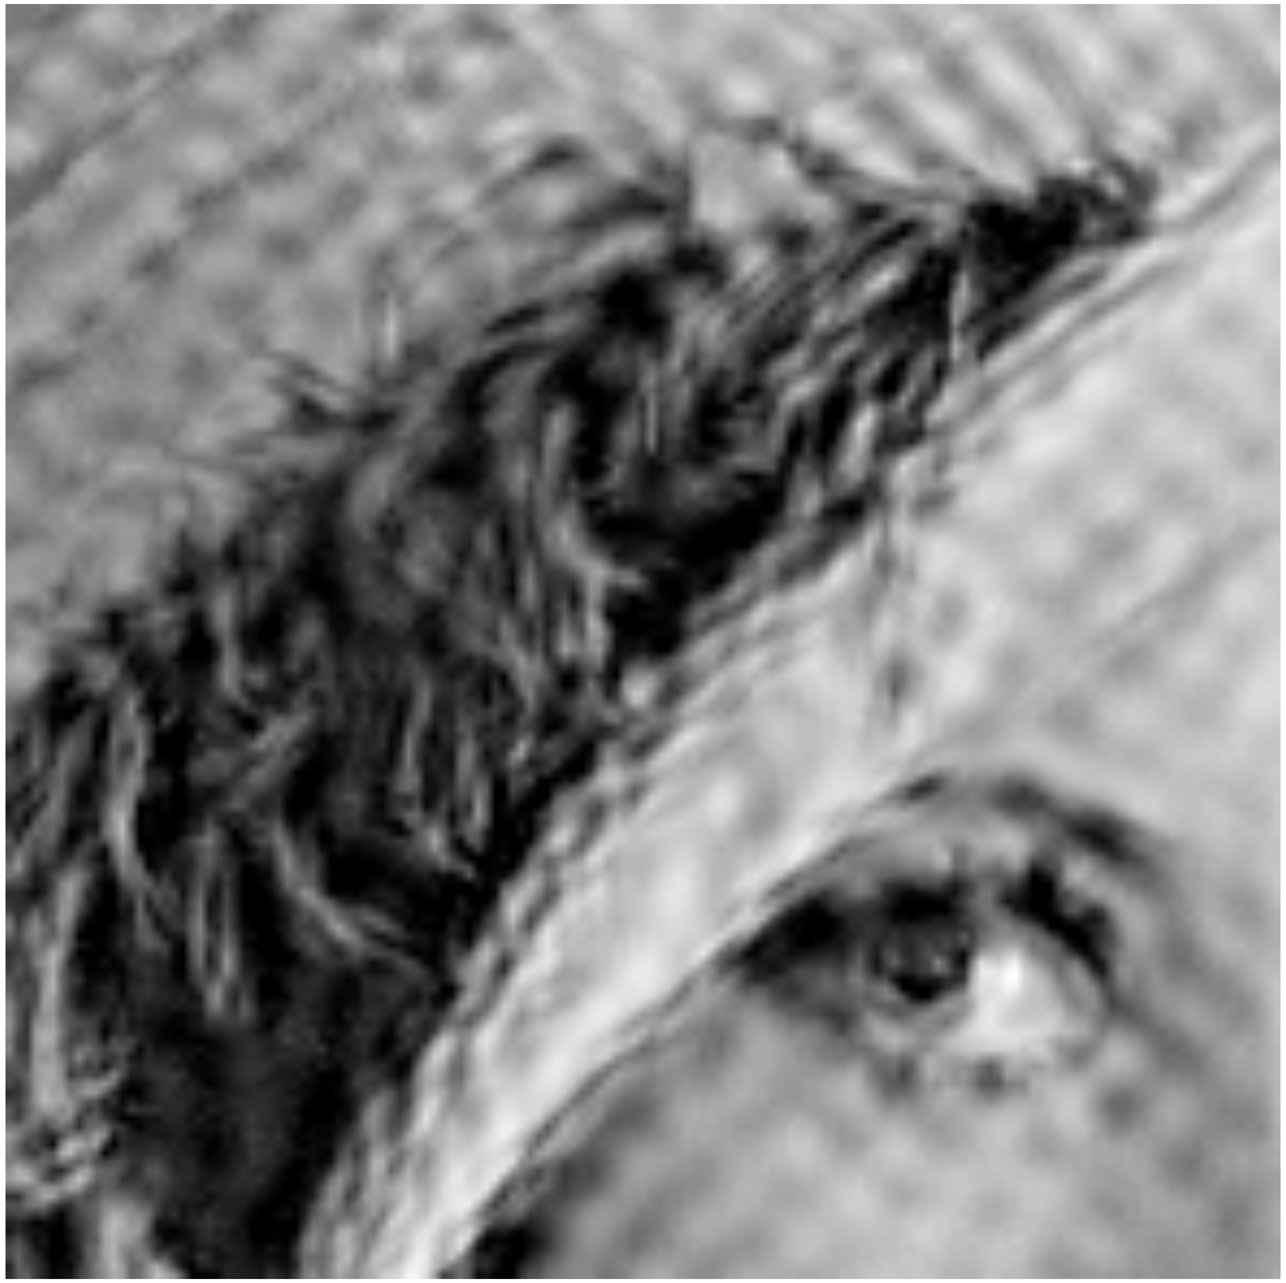

$$\|\Delta x\|_2 \;/\; \|x\|_2 = 20.7\%$$

synthesized image from covariance map

$$\|\Delta x\|_2 \;/\; \|x\|_2 = 11.1\%$$

target image

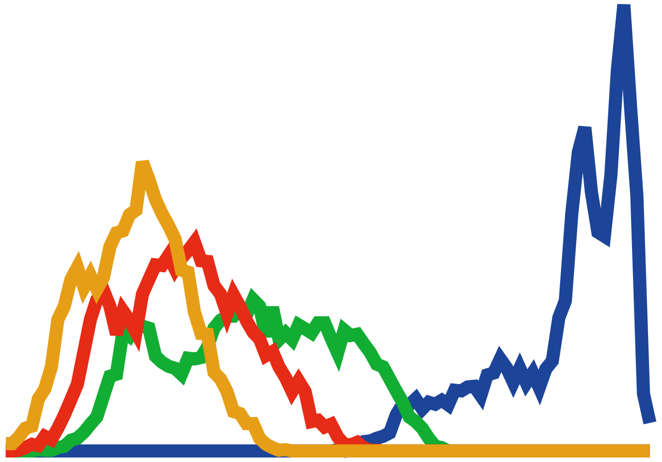

histogram of local eigen values

target image

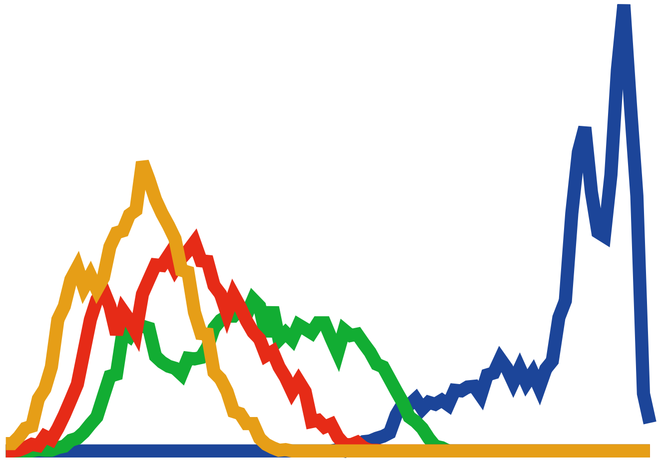

threshold

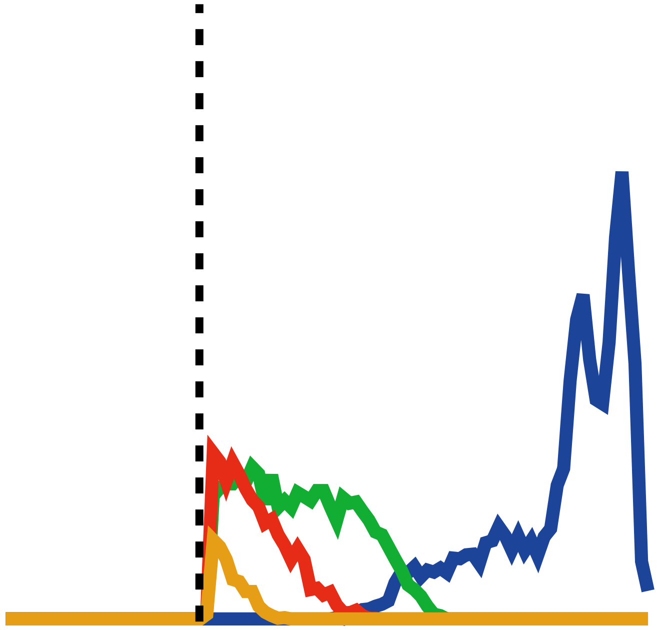

target image

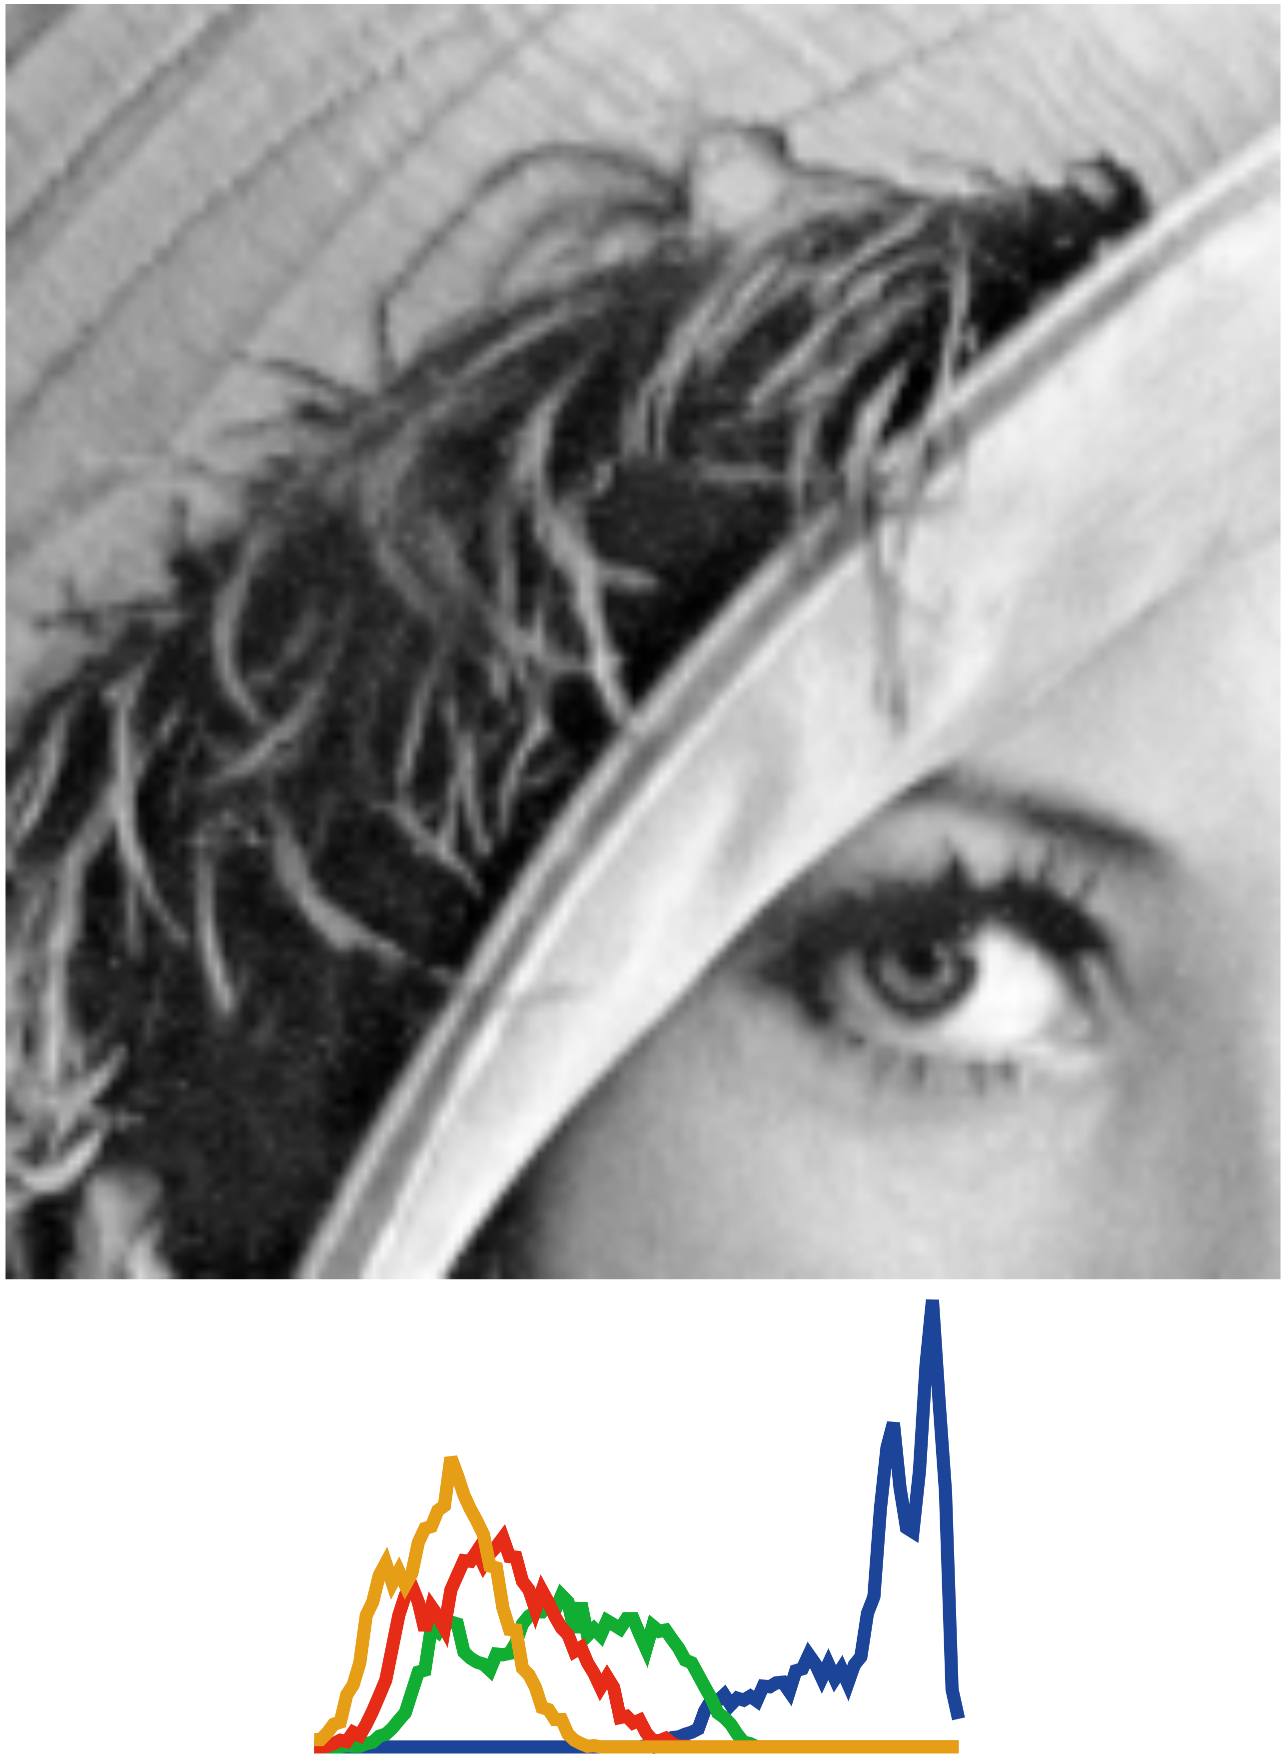

synthesized image

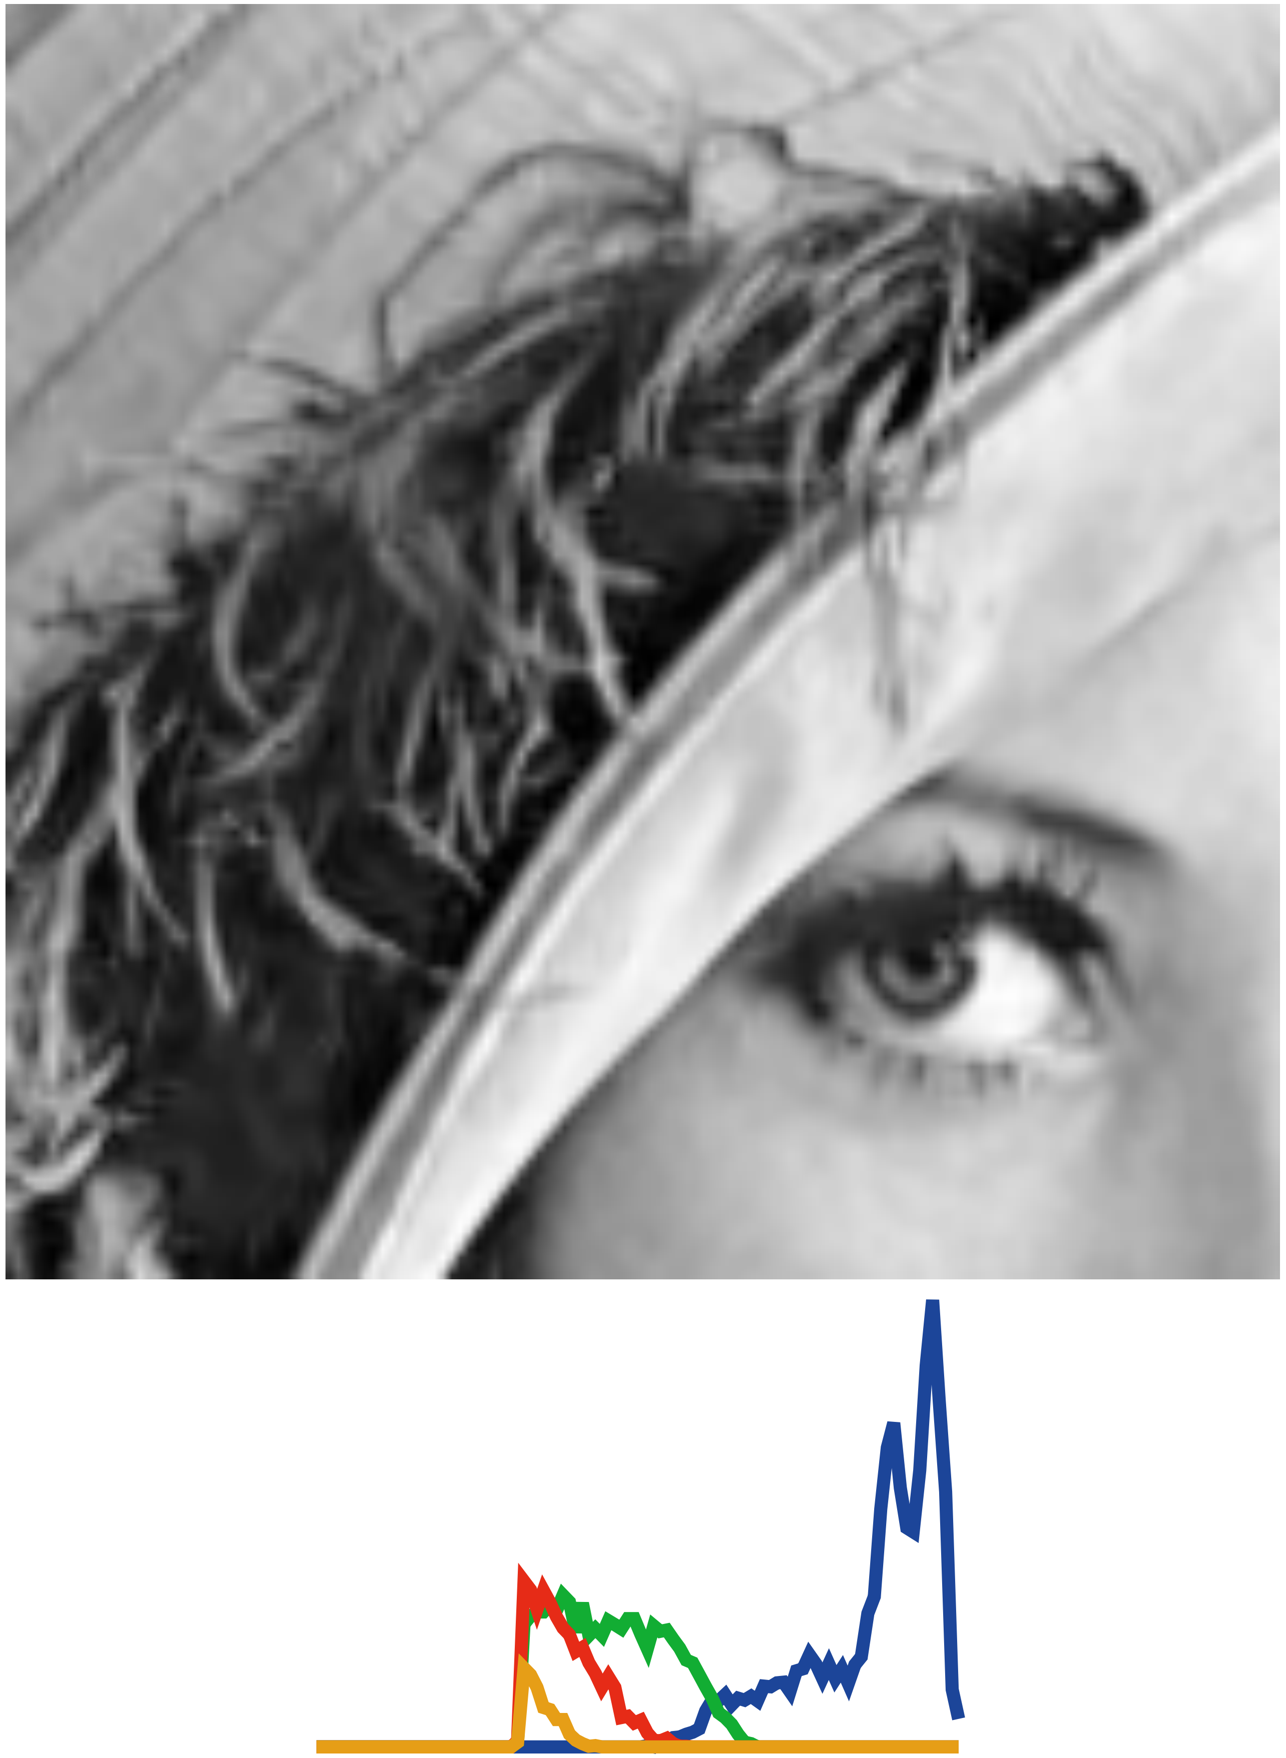

target image

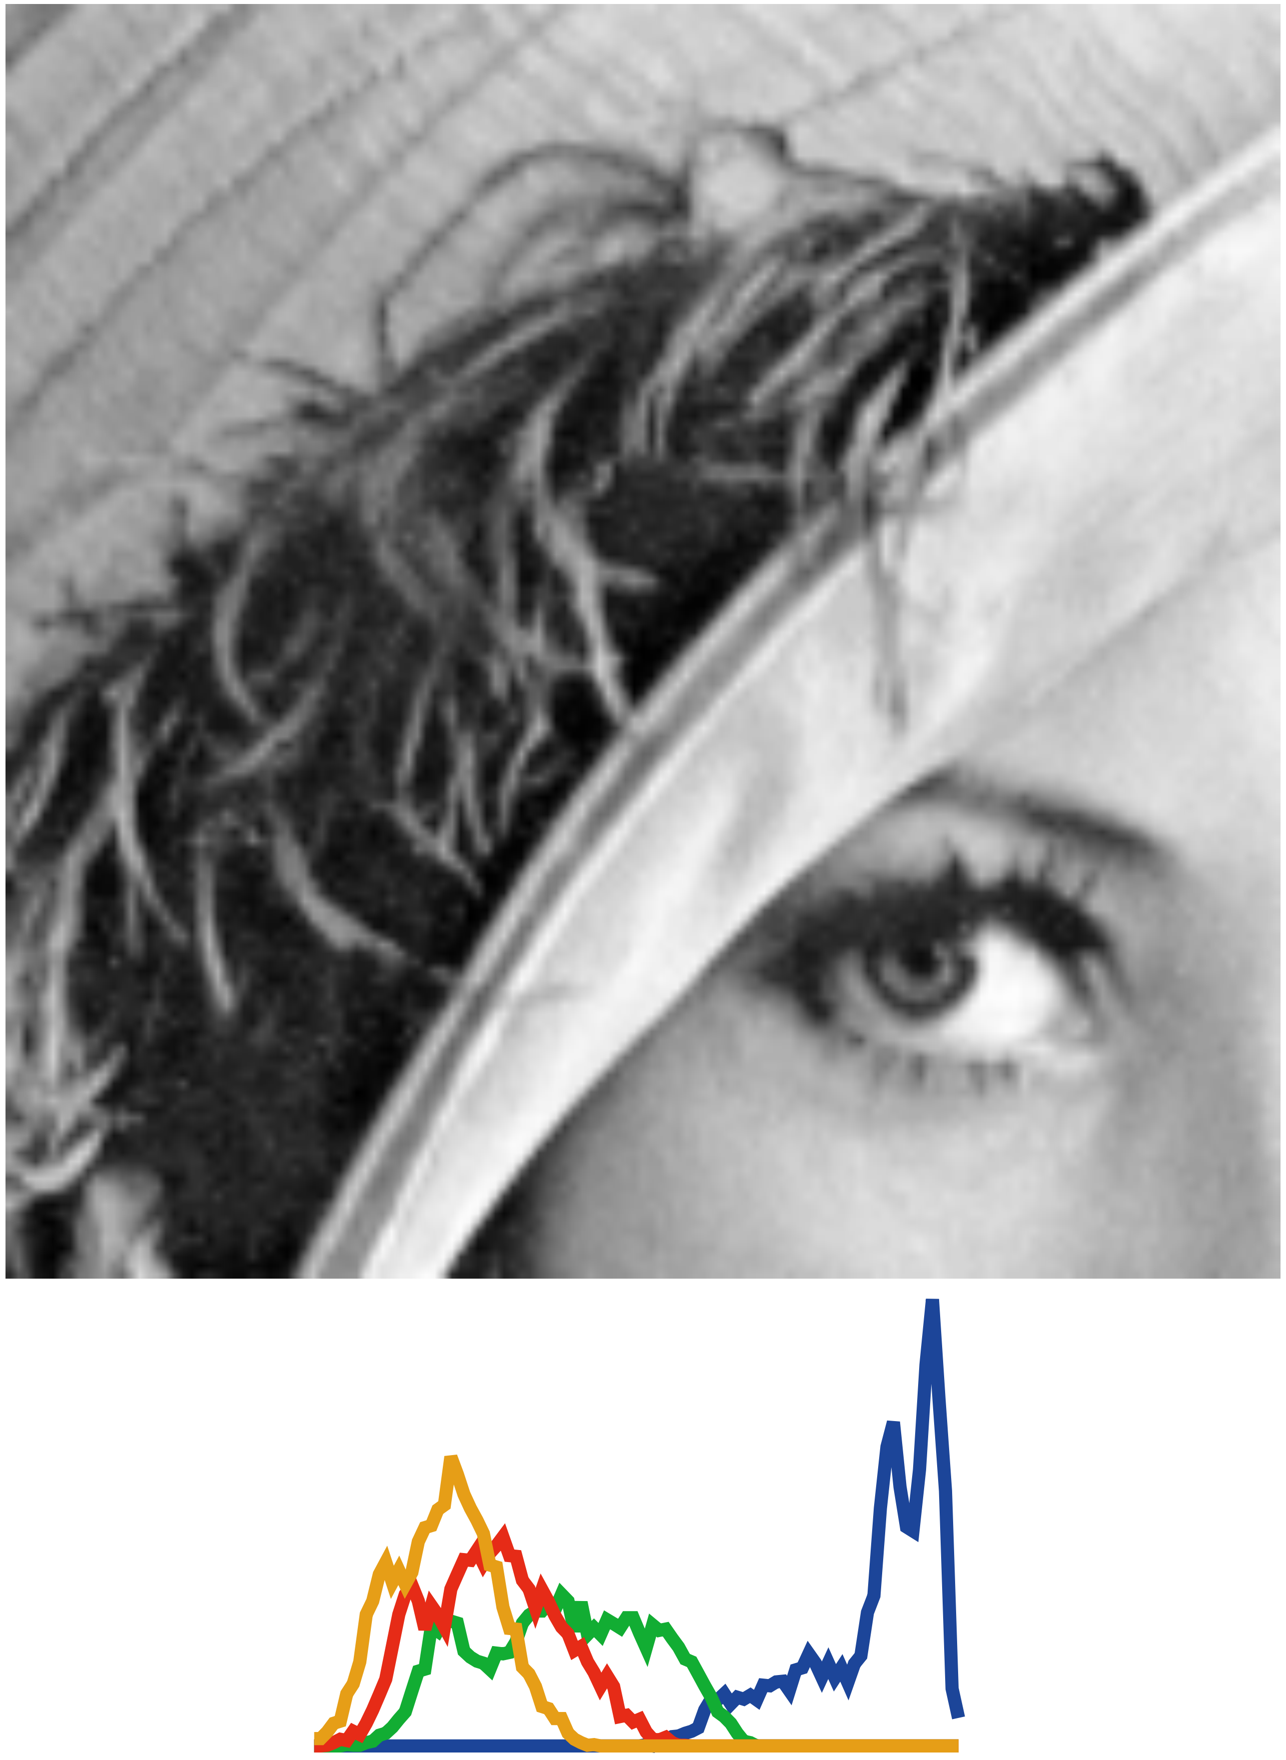

synthesized image

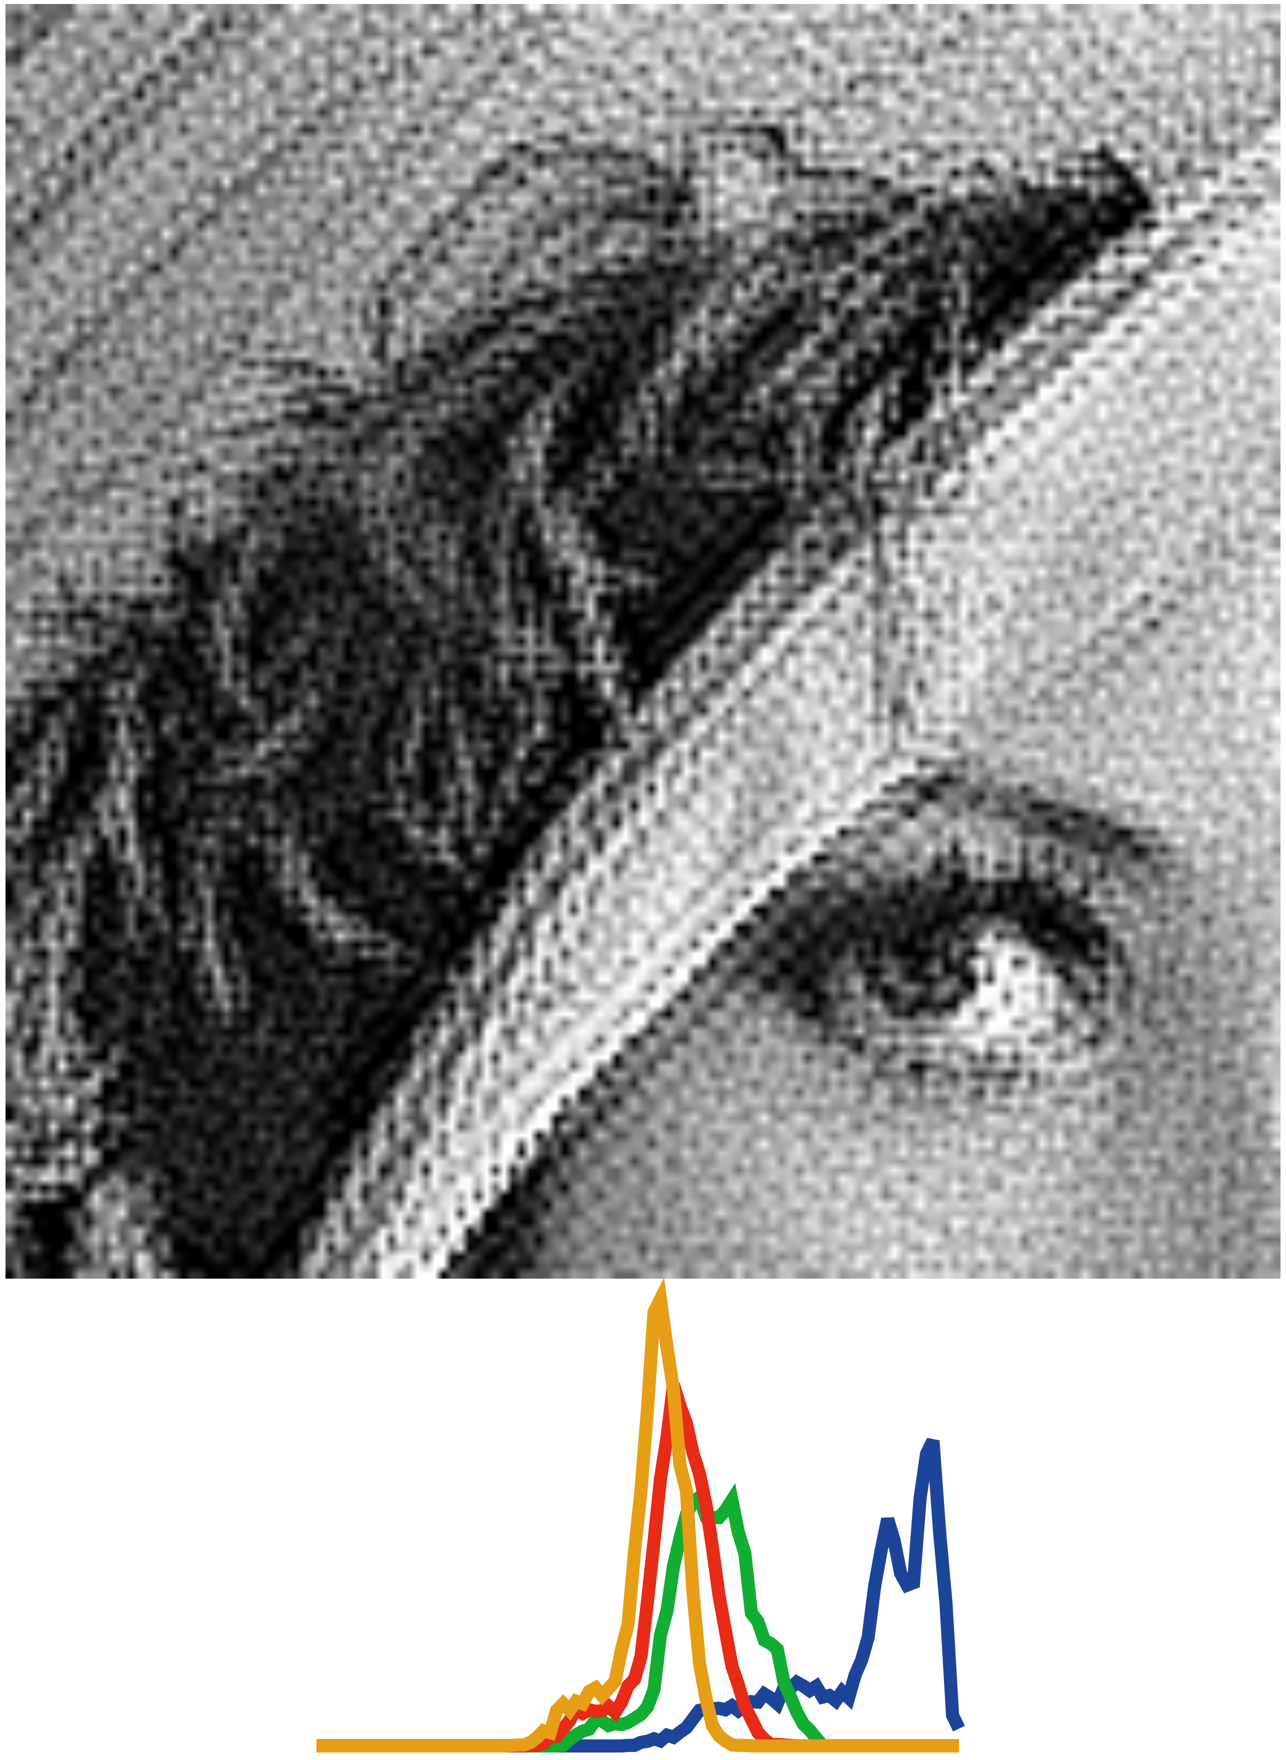

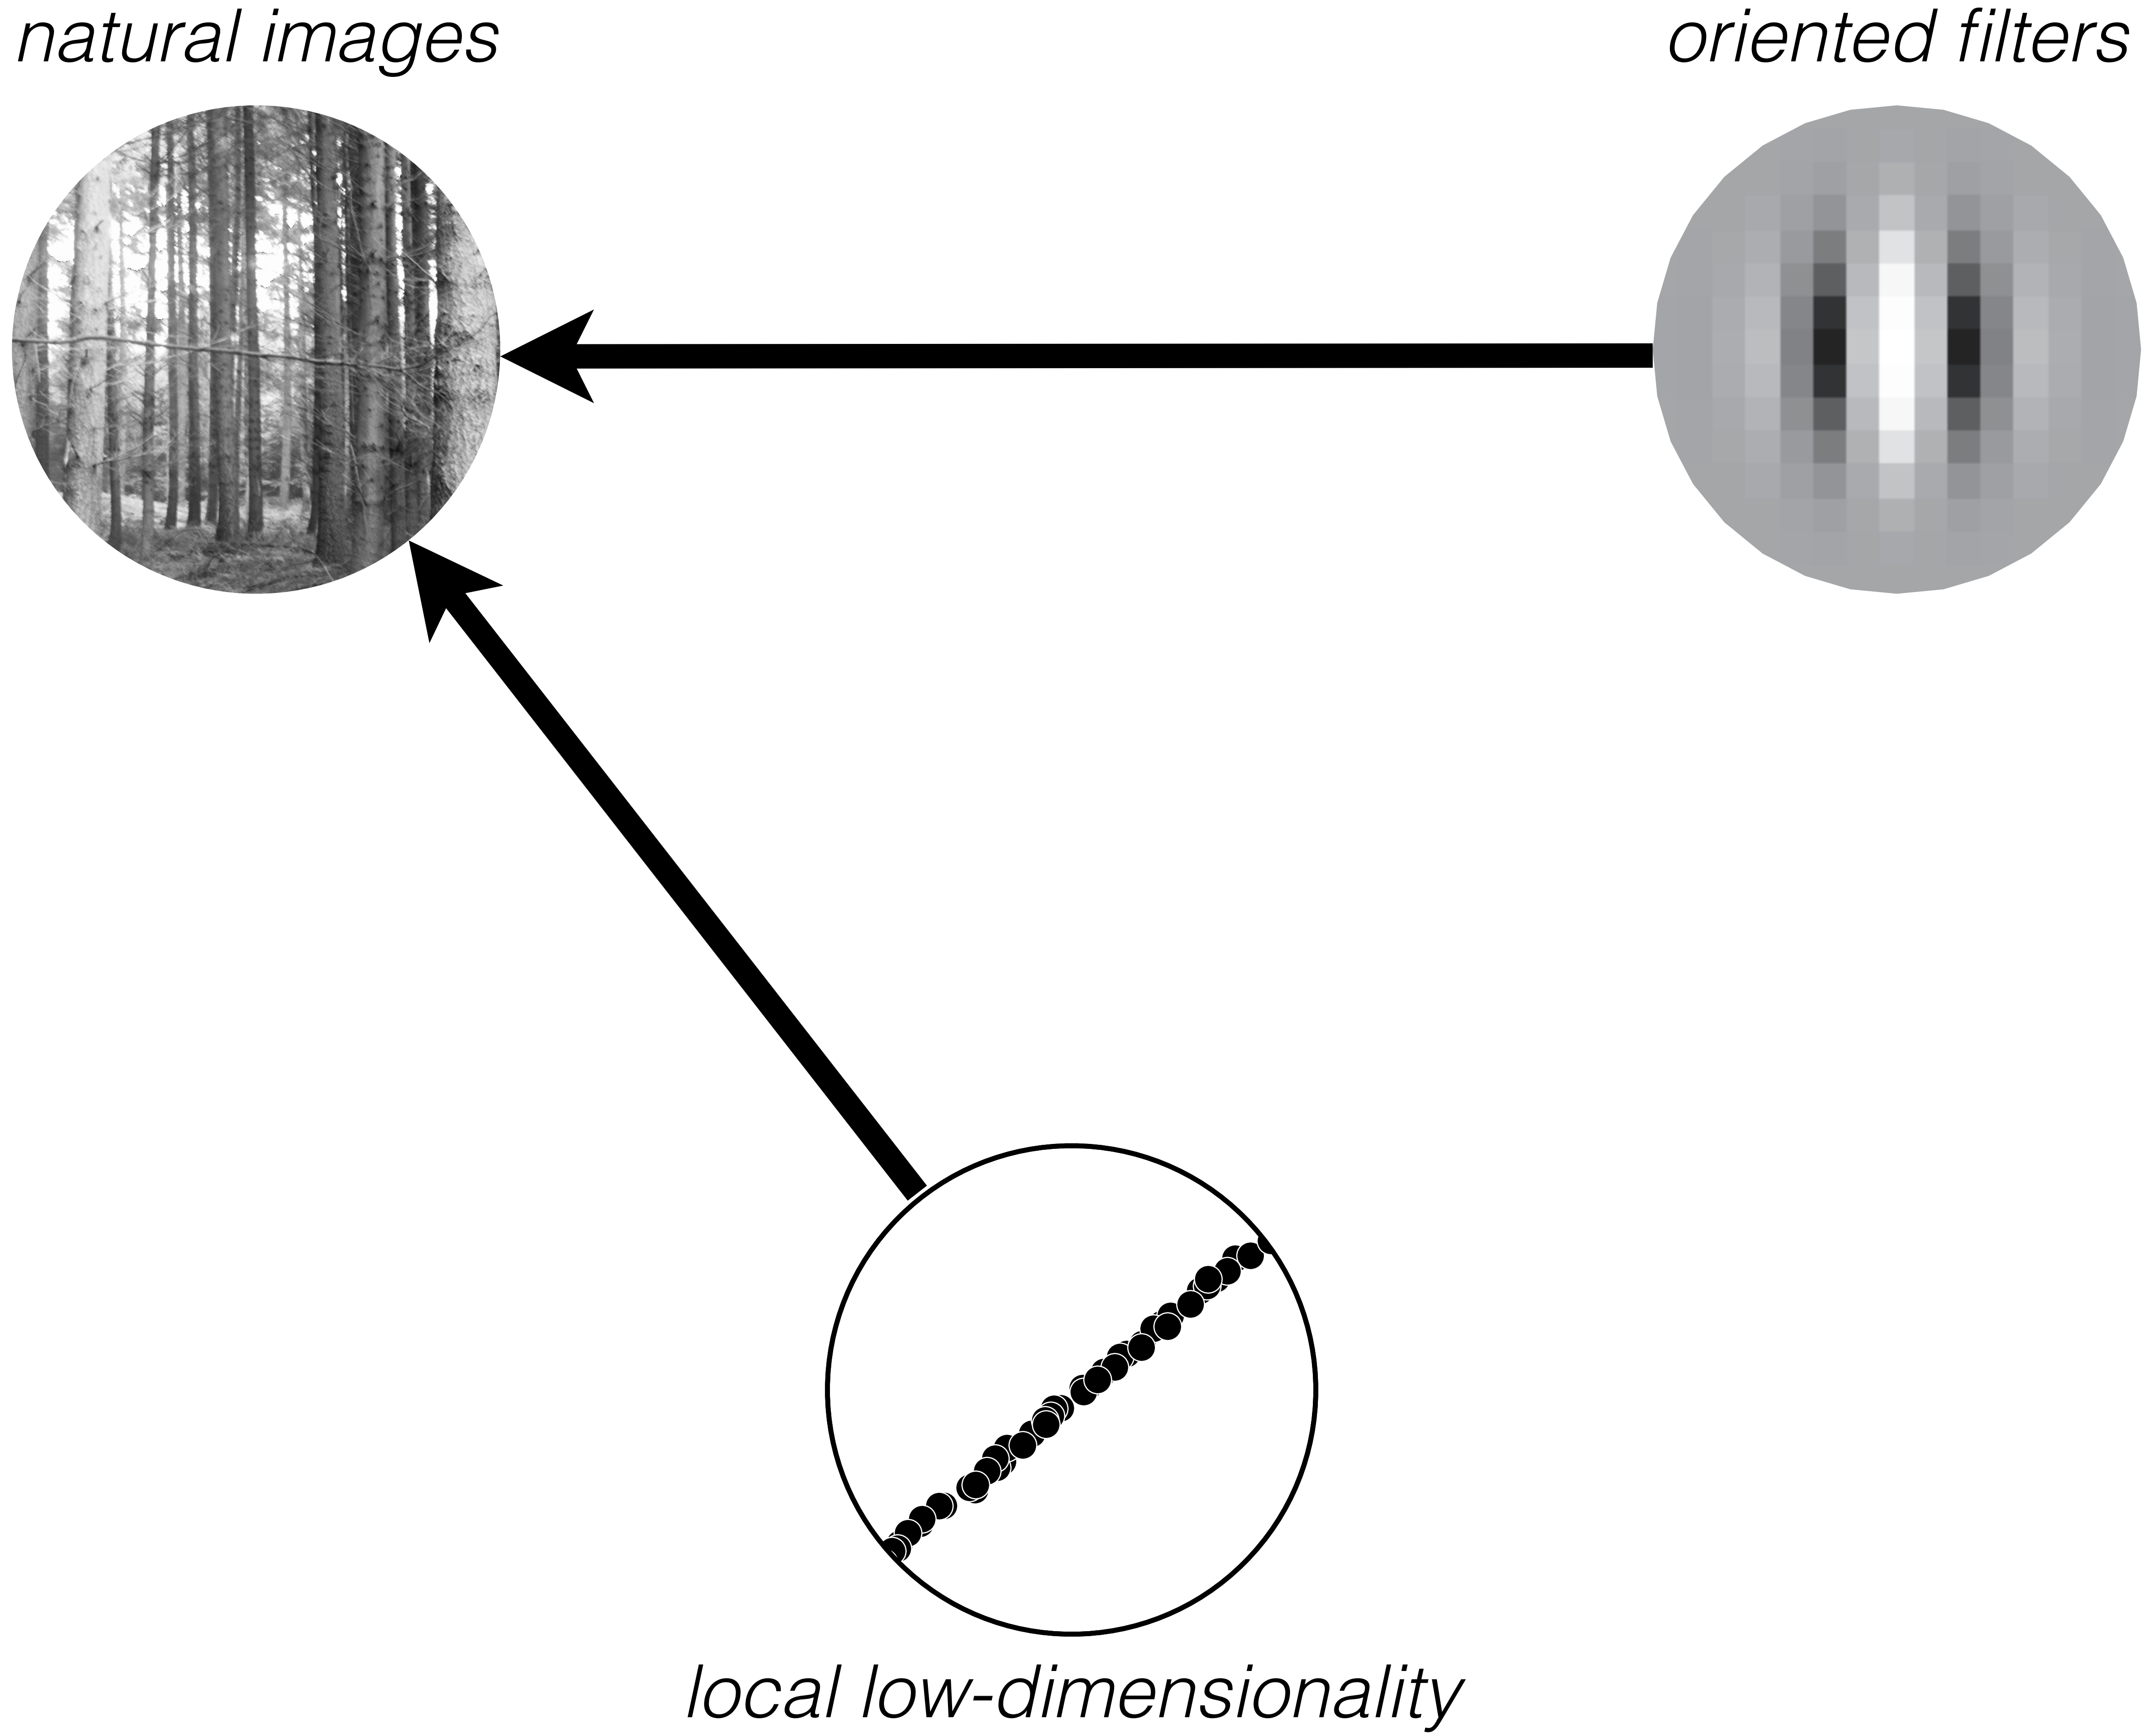
natural images
oriented filters
local low-dimensionality

Conclusion

responses of oriented, band-pass filters are locally low-dimensional

we optimized a bank of filters for local low-dimensionality

representing natural images as a map of low-dimensional covariances captures the perceptually relevant structure

future directions: stacked covariance maps

Thanks!