Denoising with filters

A crucial property of images

Pixels are like their neighbors

mostly, for most pixels

Imagine you wish to denoise an image. You could do so by averaging neighbors (a filter!).

Smoothing with box filter revisited

- What's wrong with this picture?

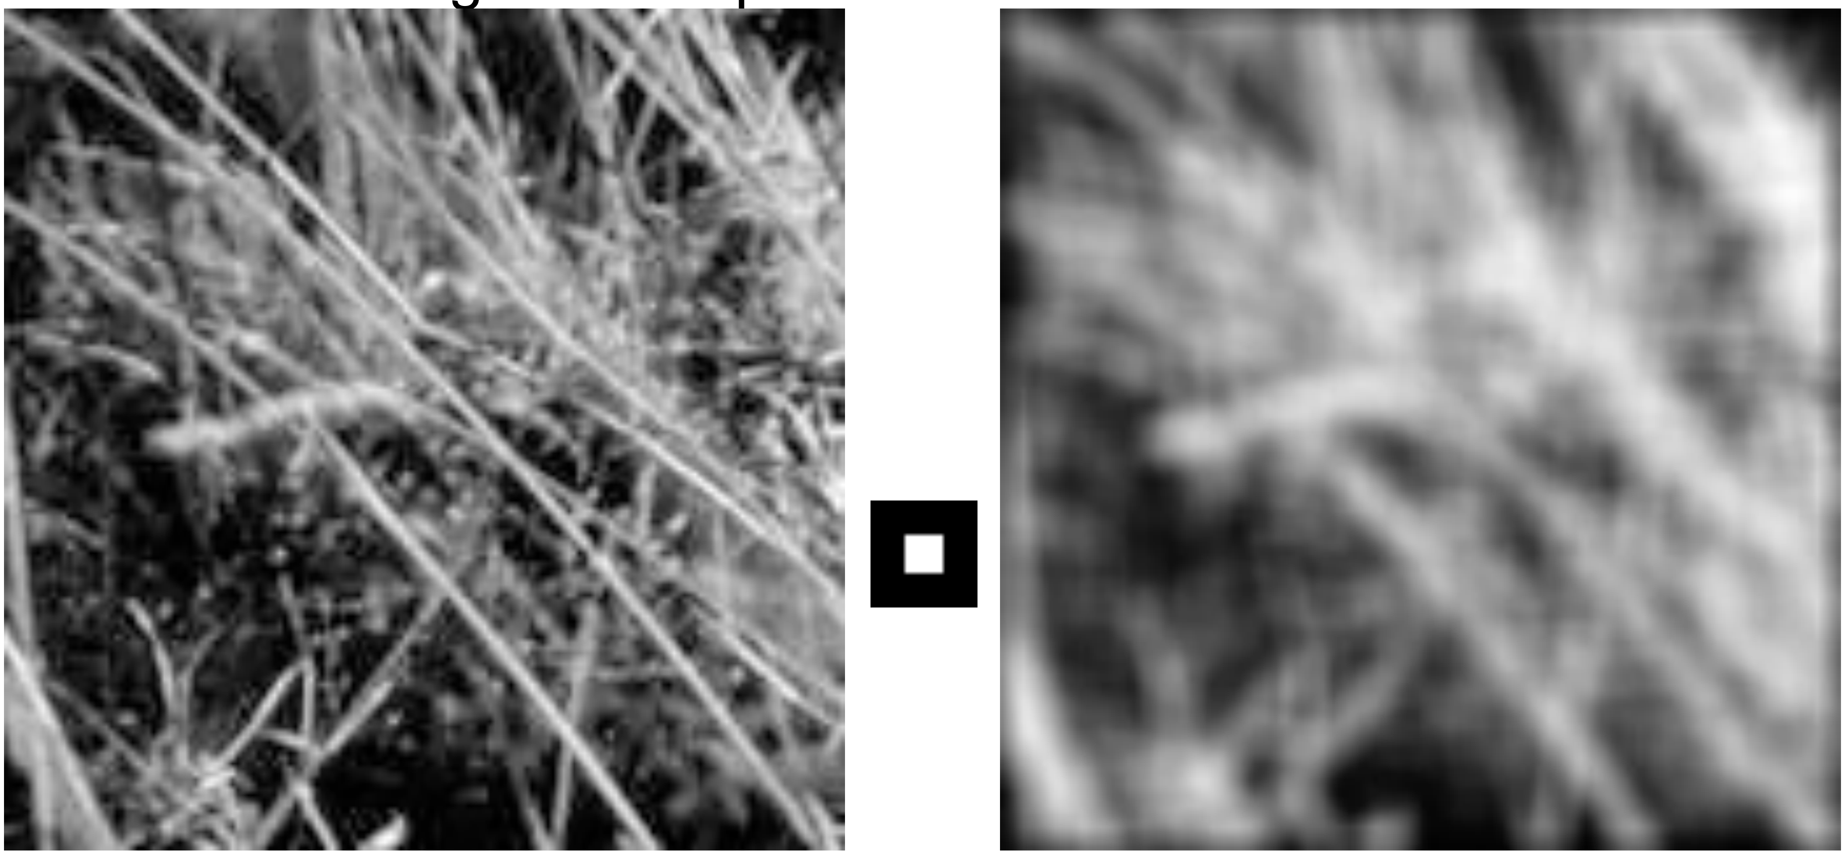

Source: D. Forsyth

A crucial property of images

Pixels are like their neighbors

mostly, for most pixels

the closer the neighbor the more alike

Imagine you wish to denoise an image. You could do so by averaging neighbors (a filter!). Weighting the neighbors so nearby neighbors get heavier weights is a good move.

Smoothing with box filter revisited

- What's wrong with this picture?
- What's the solution?
  - To eliminate edge effects, weight contribution of neighborhood pixels according to their closeness to the center

"proportional to"
(renormalize values to sum to 1)

$$G(x, y) \propto \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right)$$

standard deviation
(determines size of "blob")

Gaussian filter

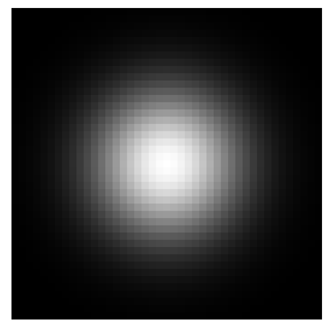

Adapted from [D. Fouhey and J. Johnson]

Gaussian filters

$\sigma = 1$ $\sigma = 2$ $\sigma = 4$ $\sigma = 8$

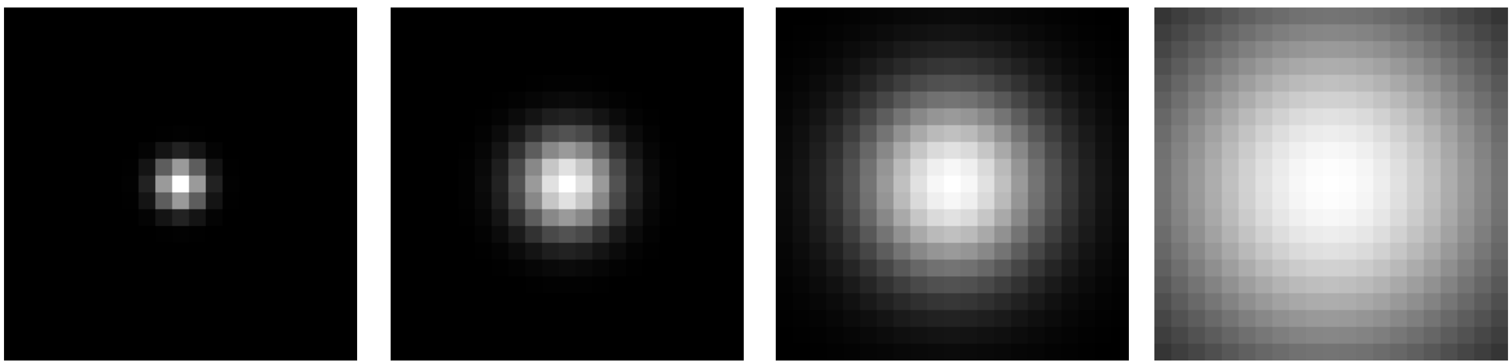

Filter size: 21×21

Adapted from [D. Fouhey and J. Johnson]

Choosing filter size

- Rule of thumb: set filter width to about 6σ (captures 99.7% of the energy)

$\sigma = 8$
Width = 21

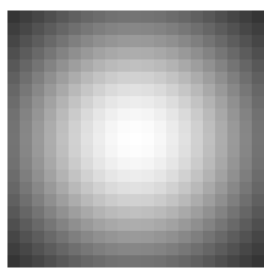

Too small!

$\sigma = 8$
Width = 43

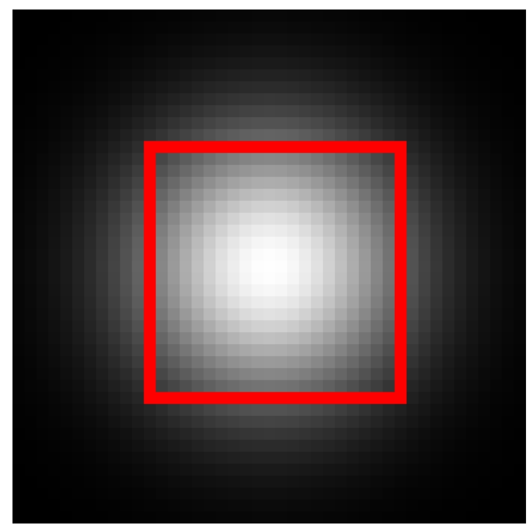

A bit small (might be OK)

Adapted from D. Fouhey and J. Johnson

Gaussian vs. box filtering

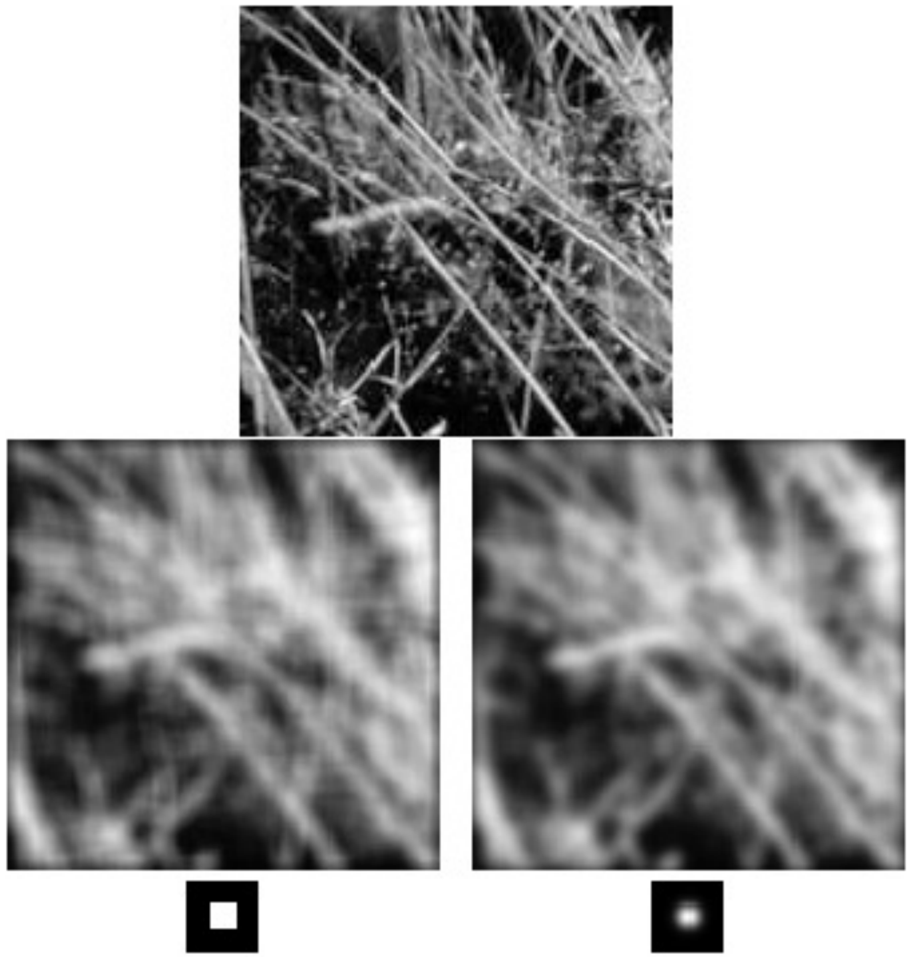

Gaussian noise

The simplest model of image noise is the *additive stationary Gaussian noise* (or *Gaussian noise*) model, where each pixel has added to it a value chosen independently from the same normal (Gaussian – same Gauss, different sense) probability distribution. This distribution almost always has zero mean. The standard deviation is a parameter of the model. Figure 4.6 shows some examples of additive stationary Gaussian noise.

Gaussian smoothing of Gaussian noise

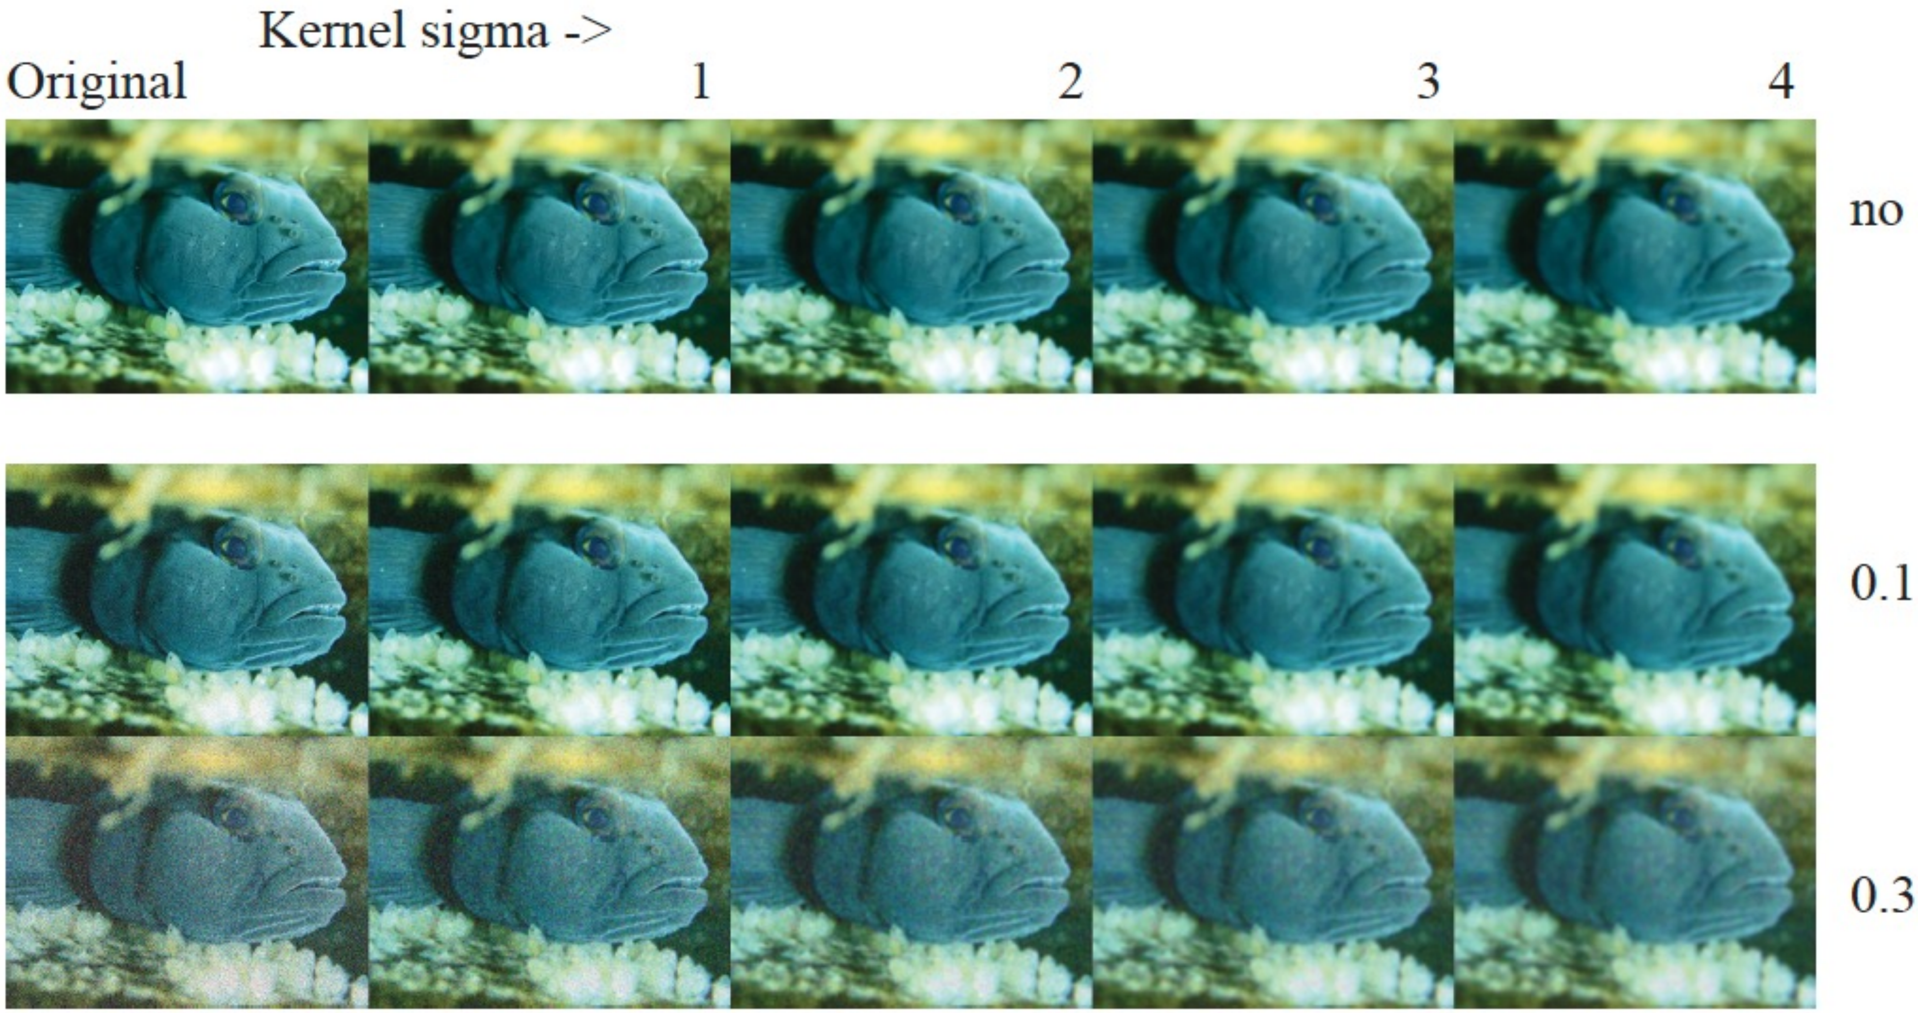

Smoothing by how much?

The choice of σ (or scale) for the Gaussian follows from the following considerations:

- If the standard deviation of the Gaussian is very small—say, smaller than one pixel—the smoothing will have little effect because the weights for all pixels off the center will be very small.
- For a larger standard deviation, the neighboring pixels will have larger weights in the weighted average, which in turn means that the average will be strongly biased toward a consensus of the neighbors. This will be a good estimate of a pixel's value, and the noise will largely disappear at the cost of some blurring.
- Finally, a kernel that has a large standard deviation will cause much of the image detail to disappear, along with the noise.

Sharpening

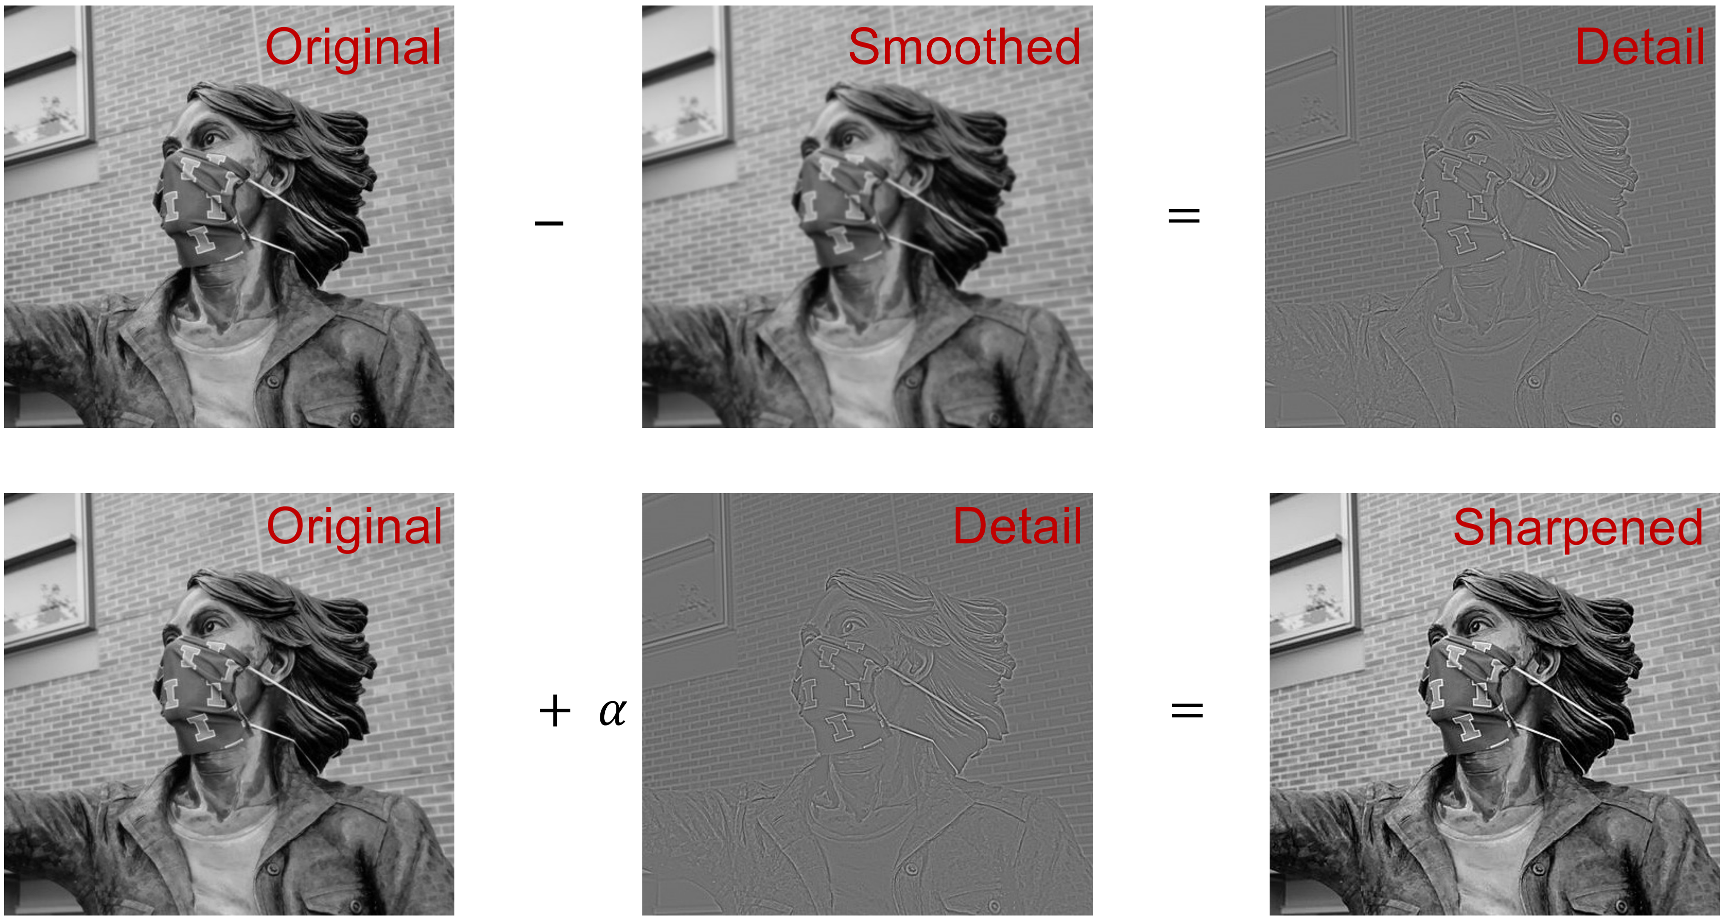

Source:
S. Gupta

“Detail” filter follows from linearity

$$I - I * g = I * (e - g)$$

unit impulse
(identity)

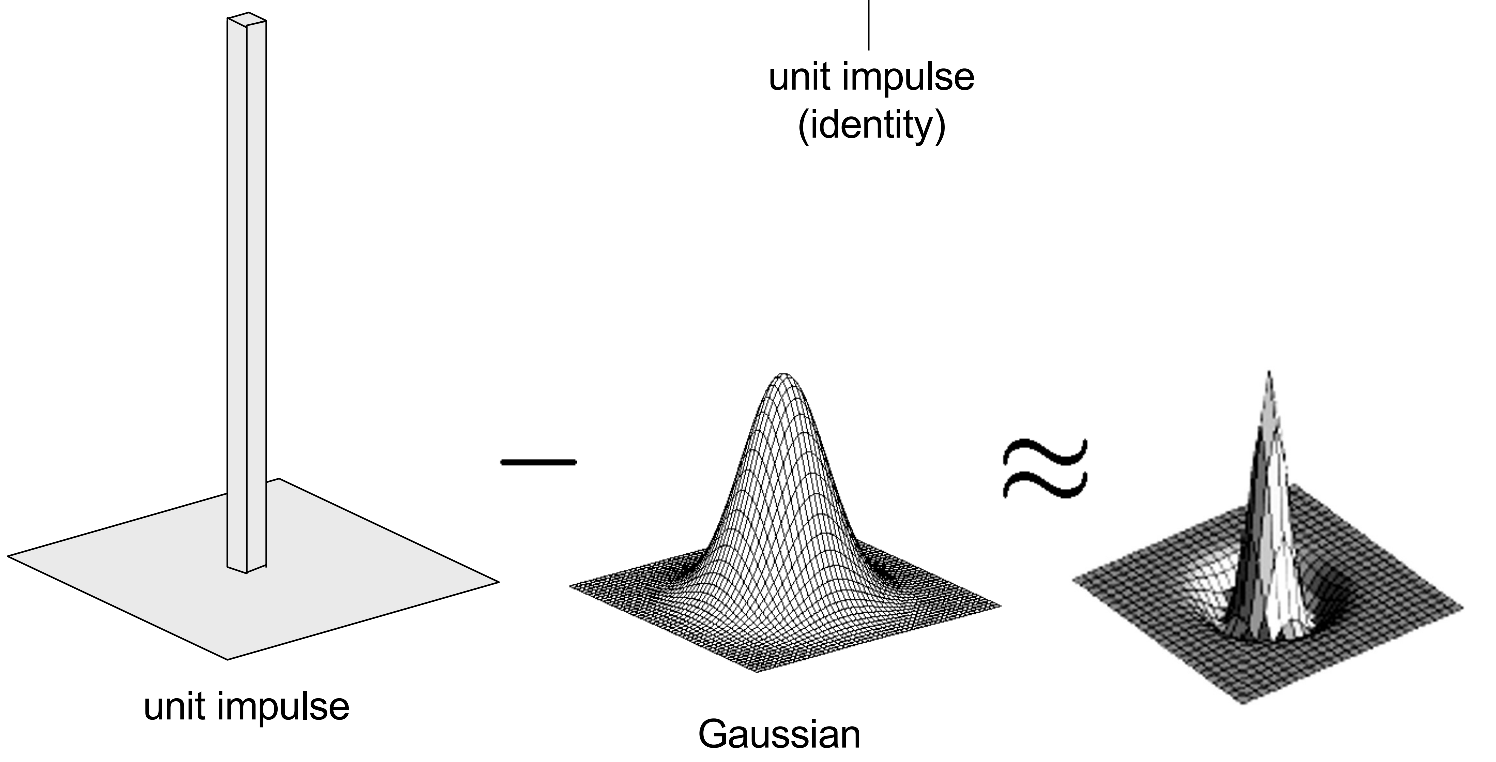

Poisson noise

For each pixel location, flip a biased coin

- if it comes up heads, move on
- if it comes up tails, flip a fair coin
  - if that is heads, pixel -> full bright
  - tails, pixel -> full dark

Models device damage, manufacturing failures, some kinds of transmission error, etc.

Smoothing Poisson noise with a gaussian filter

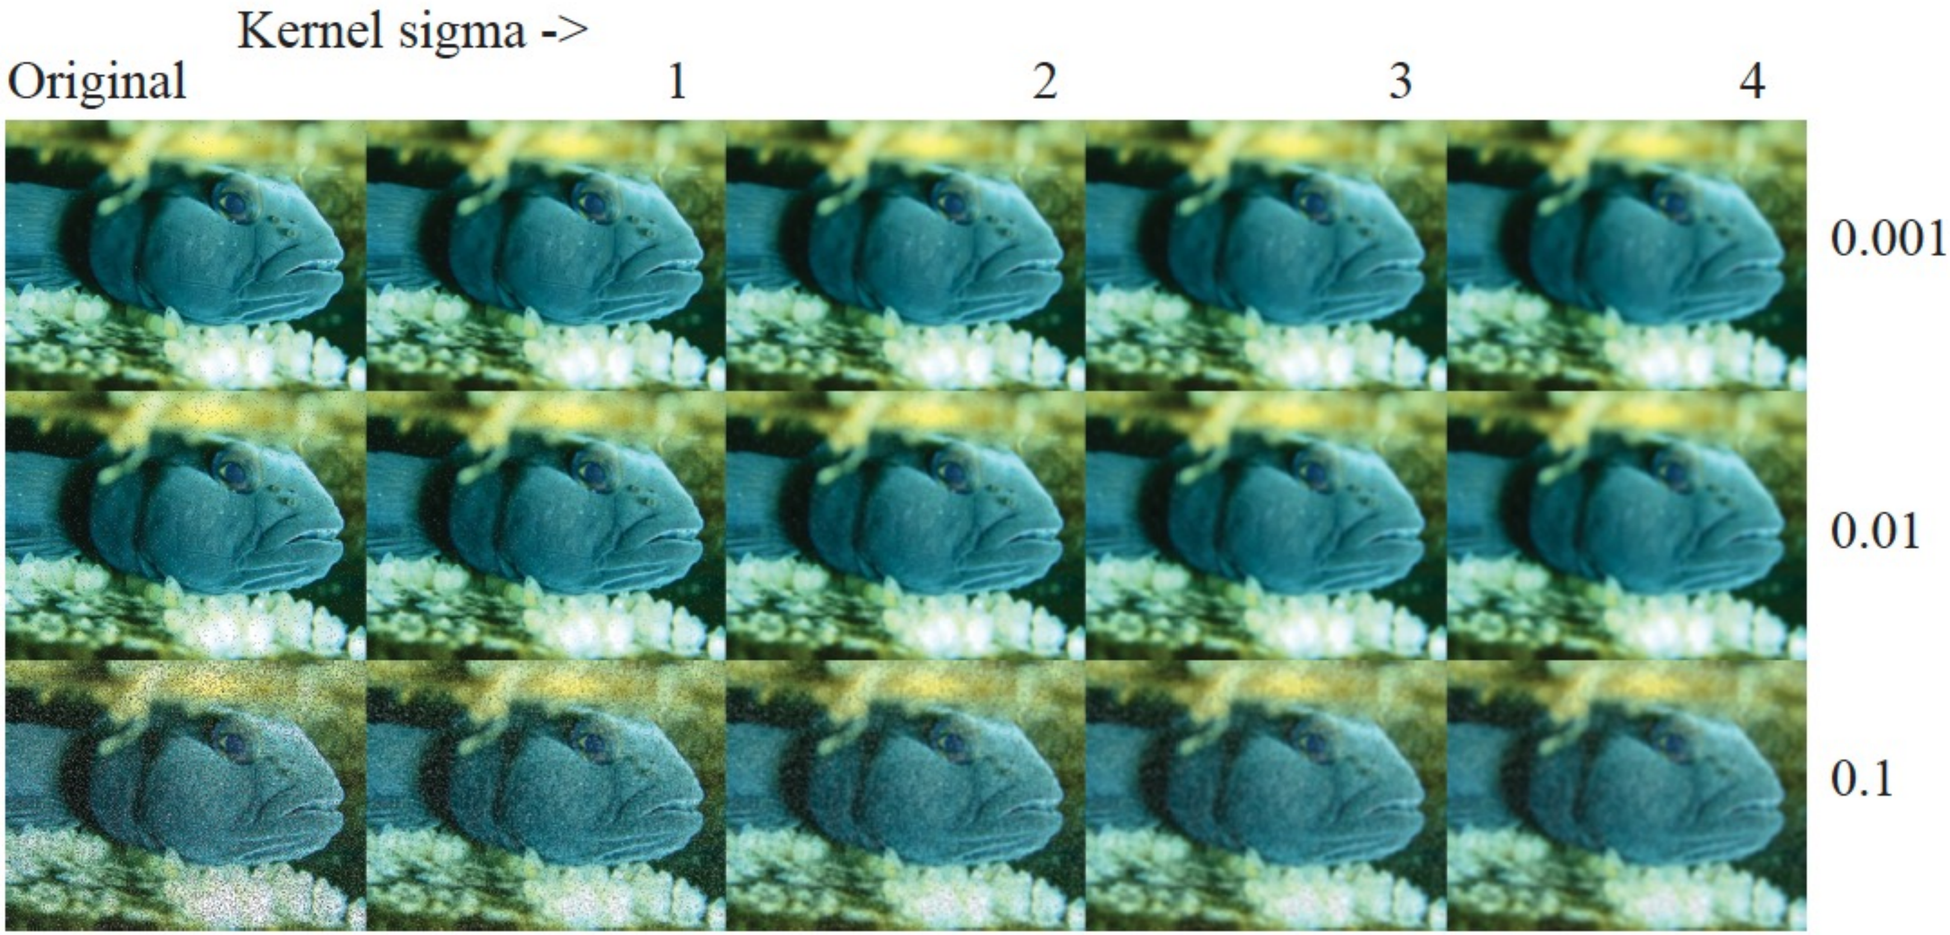

The median filter

$N_{ij} = \text{median}(\text{Neighborhood}(O_{ij}))$

THIS ISN’T LINEAR!

(check you’re sure of this)

Smoothing Poisson noise with a median filter

Window size ->

Original 3 5 7 9

0.001

0.01

0.1

Median filters vs Gaussian noise

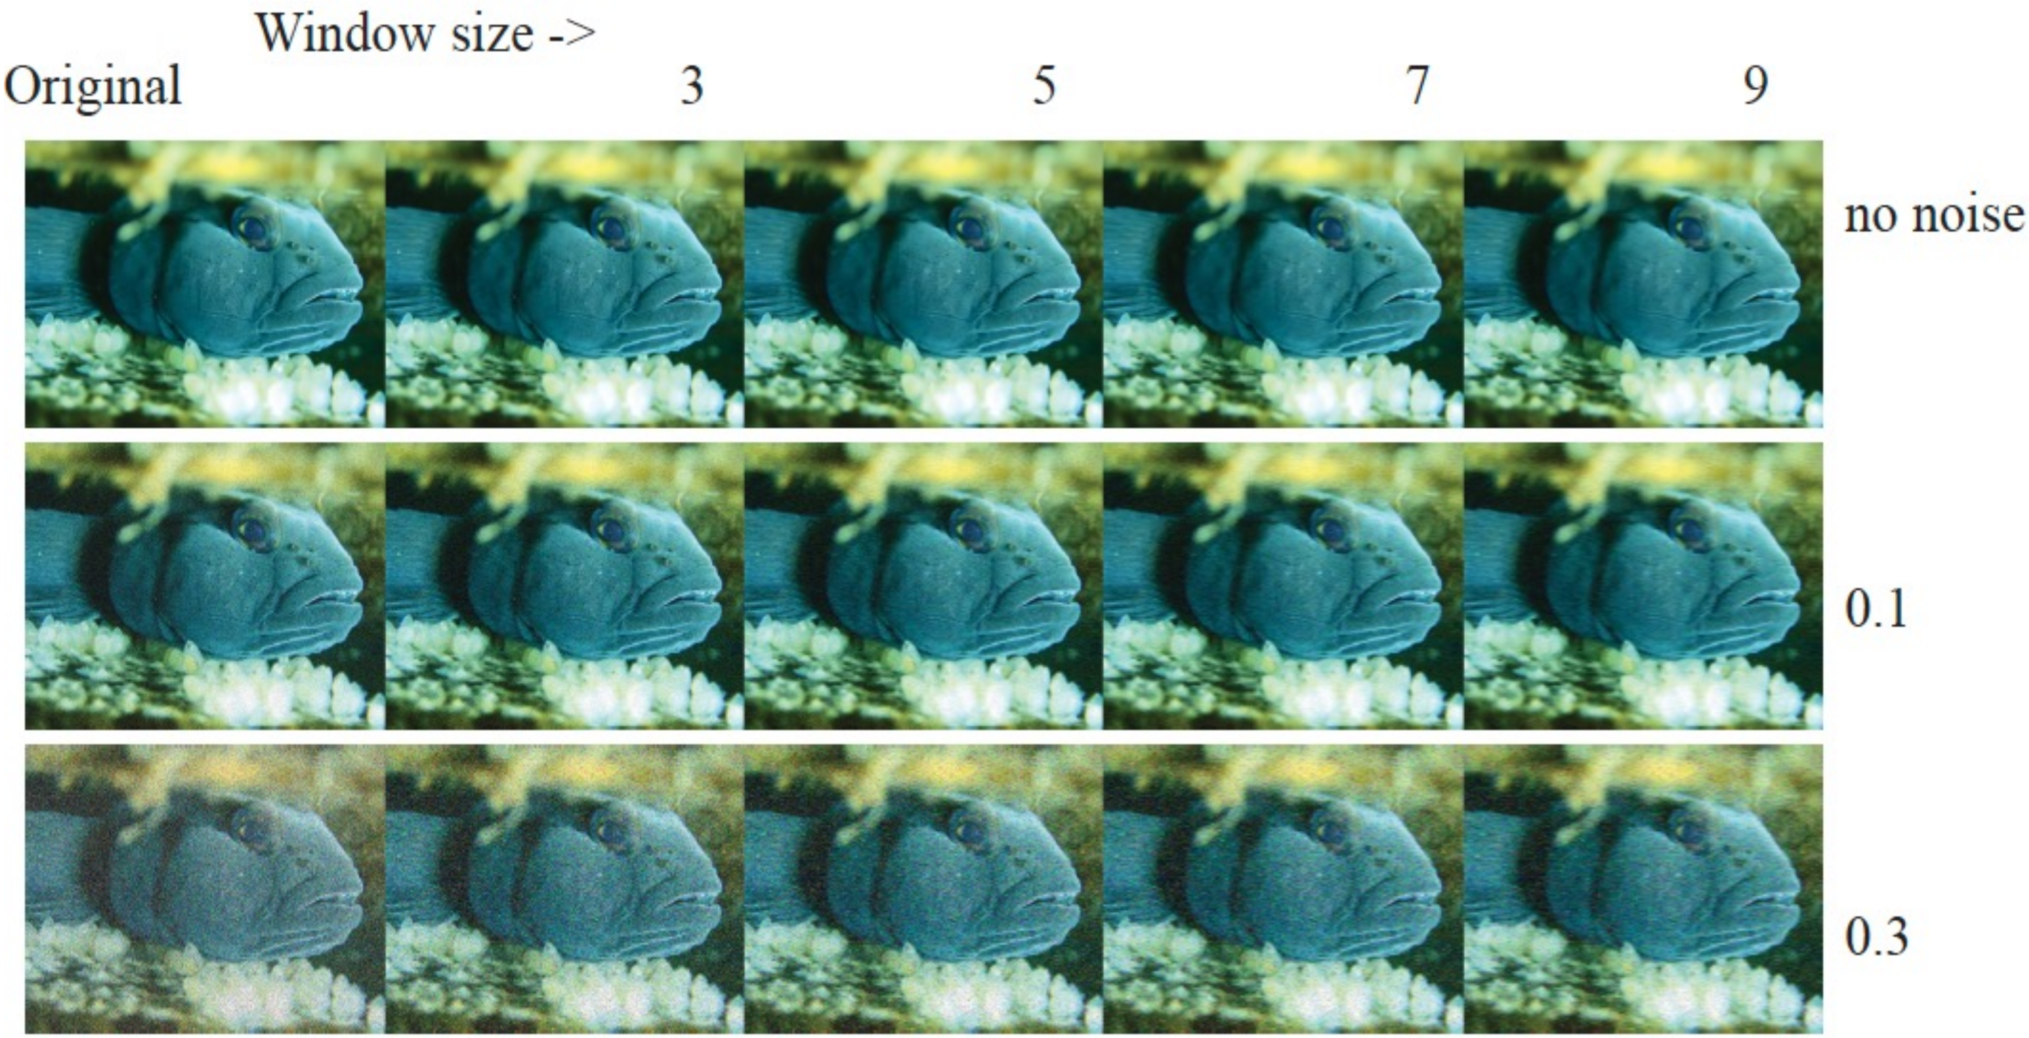

Gaussian smoothing of Gaussian noise