Blur-Aware Image Downsampling

Matthew Trentacoste
Rafał Mantiuk
Wolfgang Heidrich

University of British Columbia
Bangor University

Is the photograph blurry?

Is the photograph blurry?

Is the photograph blurry?

Is the photograph blurry?

Motivation

- Sensors higher resolution than displays
- Image display implies image downsizing
- Conventional downsizing doesn't accurately represent image appearance and perception of the image changes

- Users can make inaccurate quality assessments when not viewing image pixels 1-to-1 with display pixels

- HDTV only 2mp, even mobile phones 3+mp
- Specifically, lowering the resolution of the image can cause blurred regions to seem sharp
- Downsampled appears higher quality than original counterpart

Motivation

2 Mp

- Sensors higher resolution than displays
- Image display implies image downsizing
- Conventional downsizing doesn't accurately represent image appearance and perception of the image changes

- Users can make inaccurate quality assessments when not viewing image pixels 1-to-1 with display pixels

- HDTV only 2mp, even mobile phones 3+mp
- Specifically, lowering the resolution of the image can cause blurred regions to seem sharp
- Downsampled appears higher quality than original counterpart

Motivation

2 Mp

- Sensors higher resolution than displays
- Image display implies image downsizing
- Conventional downsizing doesn't accurately represent image appearance and perception of the image changes

3-22 Mp

- Users can make inaccurate quality assessments when not viewing image pixels 1-to-1 with display pixels

- HDTV only 2mp, even mobile phones 3+mp
- Specifically, lowering the resolution of the image can cause blurred regions to seem sharp
- Downsampled appears higher quality than original counterpart

Motivation

- Want to preserve appearance of blur when downsampling
- Perceptual experiment: relation between blur present and perception at different sizes
- New resizing operator that amplifies blur present to ensure the result is perceived the same as the original

- Compatible with any spatially-variant blur estimation
- We chose to base our work off that of Samadani et al.

Organization

- Related work
- Experiment design + results
- Model of perceived blur
- Blur estimation
- Accurate blur synthesis
- Evaluation + conclusion

Related work

- Blur perception
  [Cufflin 2007][Chen 2009]
  [Mather 2002][Held 2010]

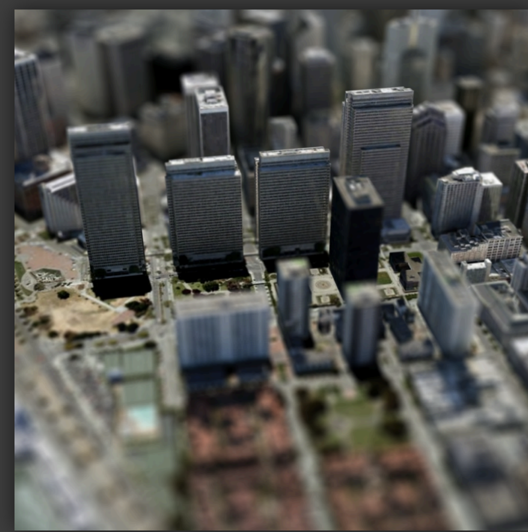

- Intelligent upsampling
  [Fattal 2007][Kopf 2007][Shan 2008]

- Seam carving
  [Avidan 2007]
  [Rubenstein 2009,2010]

- Blur discrimination: Cufflin / Chen
- Blur discrimination + depth perception: Mather
- Using blur patterns to affect perception of distance and scale: Held

- Intelligent upsampling - use image statistics to hallucinate information reconstruction filter can't

- Seam carving
- Remove column or row of pixels that change the image the least
- Mostly change aspect ratio

Related work

- Blind deconvolution
  [Lam 2000][Fergus 2006]
- Spatially-variant blur estimation
  [Elder 1998][Liu 2008]
- Blur magnification
  [Bae 2007][Samadani 2007]

- Blind deconvolution
- Estimate the PSF while deconvolution, assume spatially invariant PSF (motion blur)

- Spatially variant blur estimation
- Use simpler (Gaussian) PSF model but change it per pixel

- Bae is computationally expensive and not suitable for applications such as a digital viewfinders
- Amount of blur increased by single scale factor, specified by the user
- Blur perception more complex and neither method can ensure that the appearance of blur will remain constant if the image is resized.

Perceptual study

- Blur-matching experiment
- Given large image with reference amount of blur present ς_r
- Need to adjust blur in smaller images to match appearance of large
- Repeated for between 0 and .26 visual degrees and downsamples of 2x 4x 8x

- We have noted that images appear sharper as they are downsampled
- And we want to correct for this
- In order to do so, we need to know how much sharper images appear
- when downsampled by a given amount
- Put another way, we want to know how much blur we need to add to small image
- To match the original

- One just sharper, one just blurrier -- JND of blur

- .26 visual degrees approx Gaussian blur of 15px (1m display distance)

- Use alternate sigma for blurs in visual degrees, use conventional sigma for blurs in pixels

Perceptual study

- Blur-matching experiment
- Given large image with reference amount of blur present ς_r
- Need to adjust blur in smaller images to match appearance of large
- Repeated for between 0 and .26 visual degrees and downsamples of 2x 4x 8x

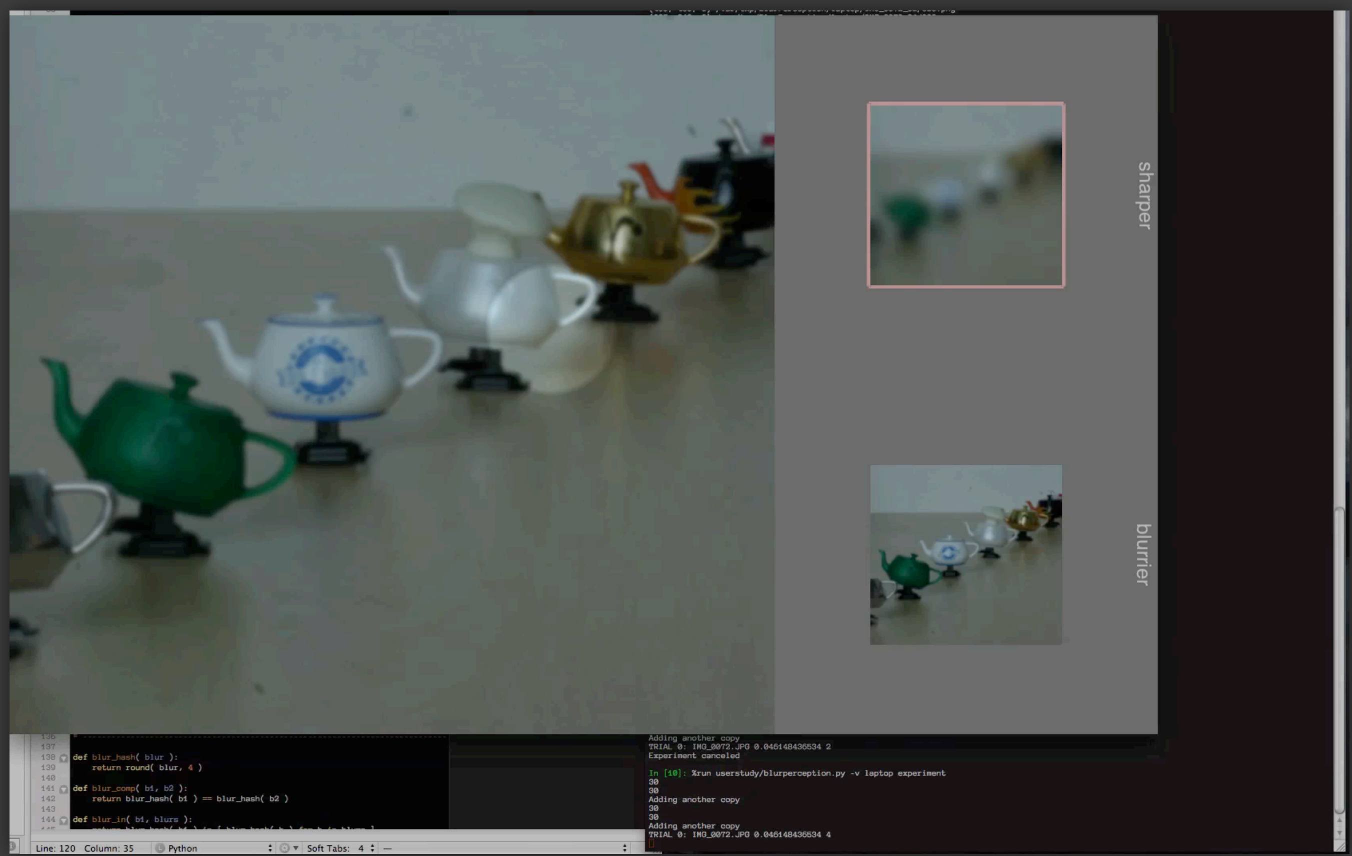

- We have noted that images appear sharper as they are downsampled
- And we want to correct for this
- In order to do so, we need to know how much sharper images appear
- when downsampled by a given amount
- Put another way, we want to know how much blur we need to add to small image
- To match the original

- One just sharper, one just blurrier -- JND of blur

- .26 visual degrees approx Gaussian blur of 15px (1m display distance)

- Use alternate sigma for blurs in visual degrees, use conventional sigma for blurs in pixels

Perceptual study

- Add uniform synthetic blur to full-size images with no noticeable blur present
- Same process for thumbnails, with nearest neighbor sampling
- 5 images selected from pre-study of 20 -- 150 conditions, trial subset of 30, 3x each
- 24 observers participated in over 2100 trials


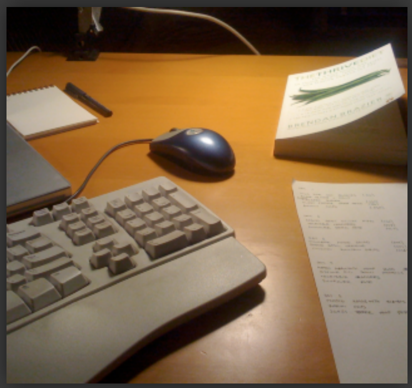


- Because we couldn’t control where the subjects were looking to make their judgments
- Nearest neighbor implies anti-aliasing for small blurs at large downsamples
- Conditions = 3 downsamples x 10 blurs x 5 images

Matching results

- Matching blur larger than reference blur, smaller images appear sharper
- Curves level off with larger blur, downsample -- blur sufficient to covey appearance
- Reported values include any blur needed to remove aliasing artifacts
- Viewing setup had Nyquist limit of 30 cpd - results not due to limited resolution in terms of pixels, but visual angle

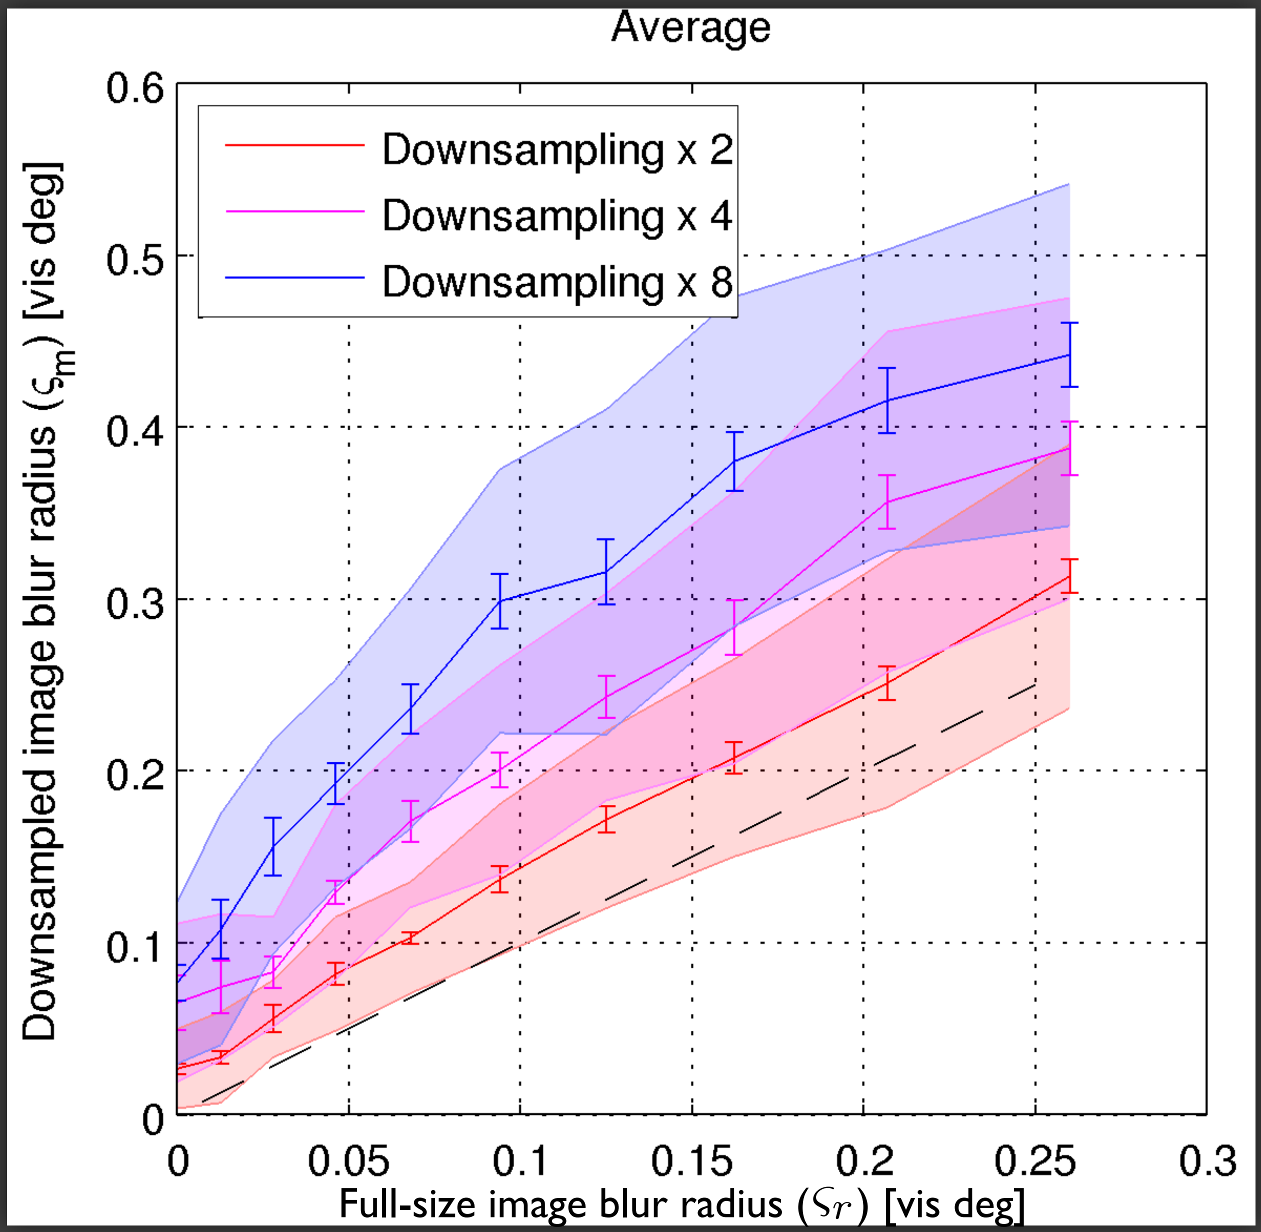

- Shaded regions denote blur chosen for sharper/blurrer image
- Error bars - 95% confidence interval

- All curves above x=y dashed line

- If blur not sufficient to remove aliasing, downsampled appeared sharper
- Subjects were instructed to match blur
- Ended up setting the amount of blur to a value close to optimal low-pass filter for given downsample

- Results are dependent on the scale of the image on the retina
- So in addition to how large the image is on the screen, viewing distance matters

Matching results

- Matching blur larger than reference blur, smaller images appear sharper
- Curves level off with larger blur, downsample -- blur sufficient to covey appearance
- Reported values include any blur needed to remove aliasing artifacts
- Viewing setup had Nyquist limit of 30 cpd - results not due to limited resolution in terms of pixels, but visual angle

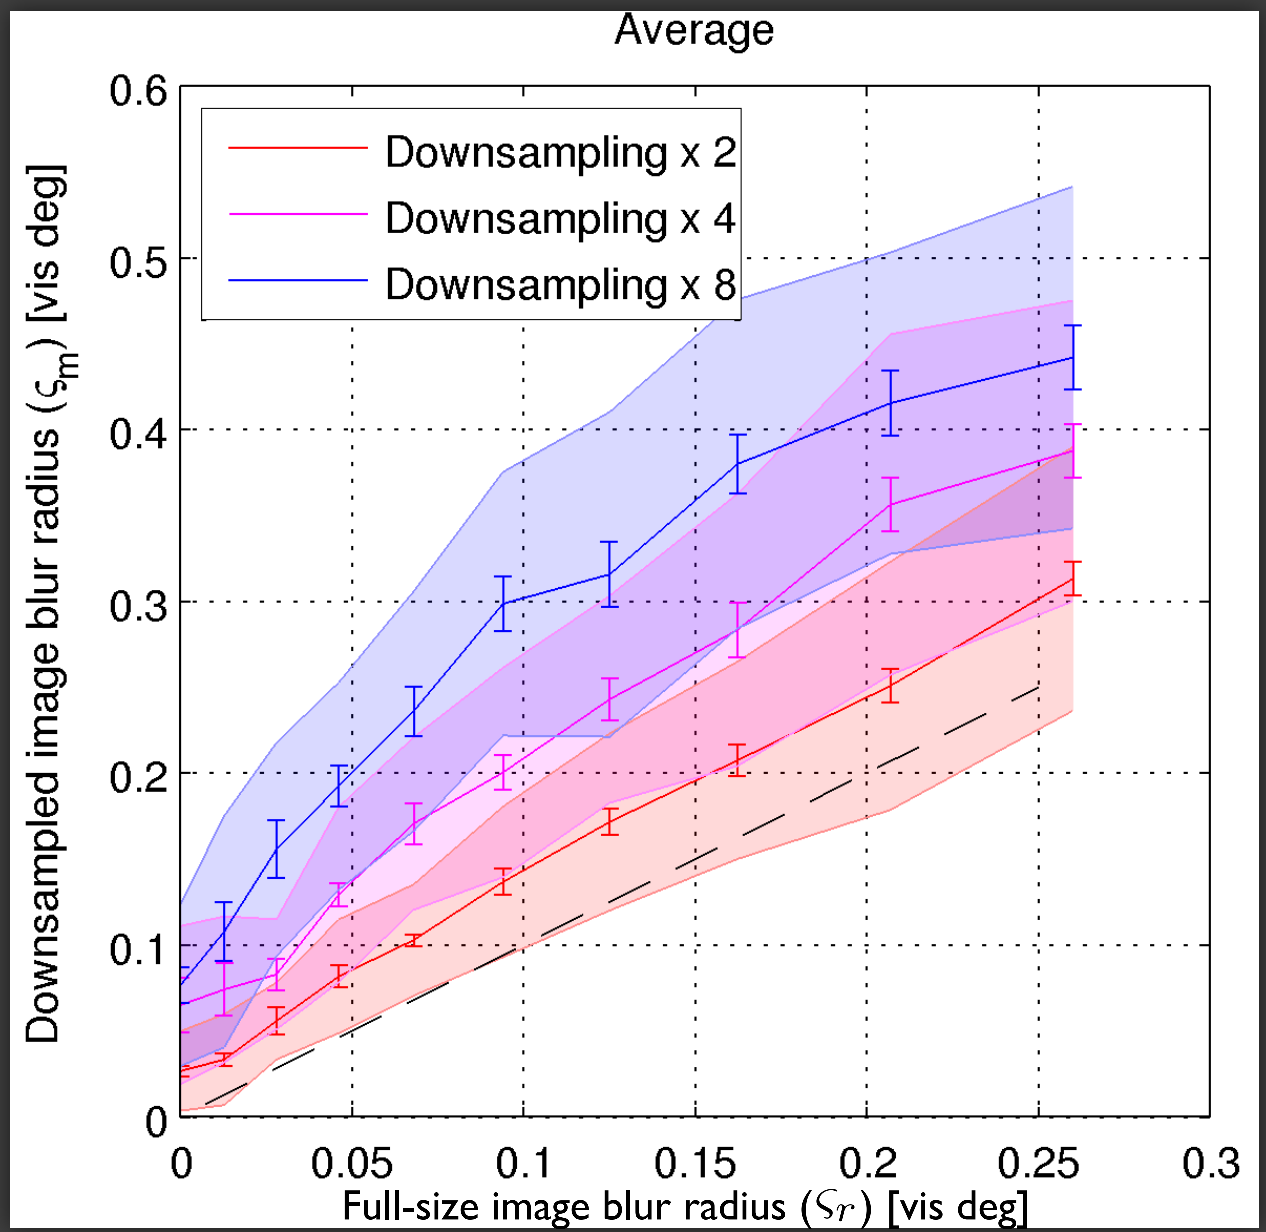

11

- Shaded regions denote blur chosen for sharper/blurrer image
- Error bars - 95% confidence interval

- All curves above x=y dashed line

- If blur not sufficient to remove aliasing, downsampled appeared sharper
- Subjects were instructed to match blur
- Ended up setting the amount of blur to a value close to optimal low-pass filter for given downsample

- Results are dependent on the scale of the image on the retina
- So in addition to how large the image is on the screen, viewing distance matters

Blur appearance model

- Measured data ς_m well predicted by anti-aliasing filter ς_d and model $\mathcal{S}$ in spatial frequencies

$$\widehat{\varsigma}_m = \sqrt{\varsigma_d^2 + \mathcal{S}^2}$$

- After removing ς_d, we model $\mathcal{S}$ as a linear function in $1/\varsigma$ spatial frequencies
- Full model provides accurate and plausible fit of the measured data in the spatial domain

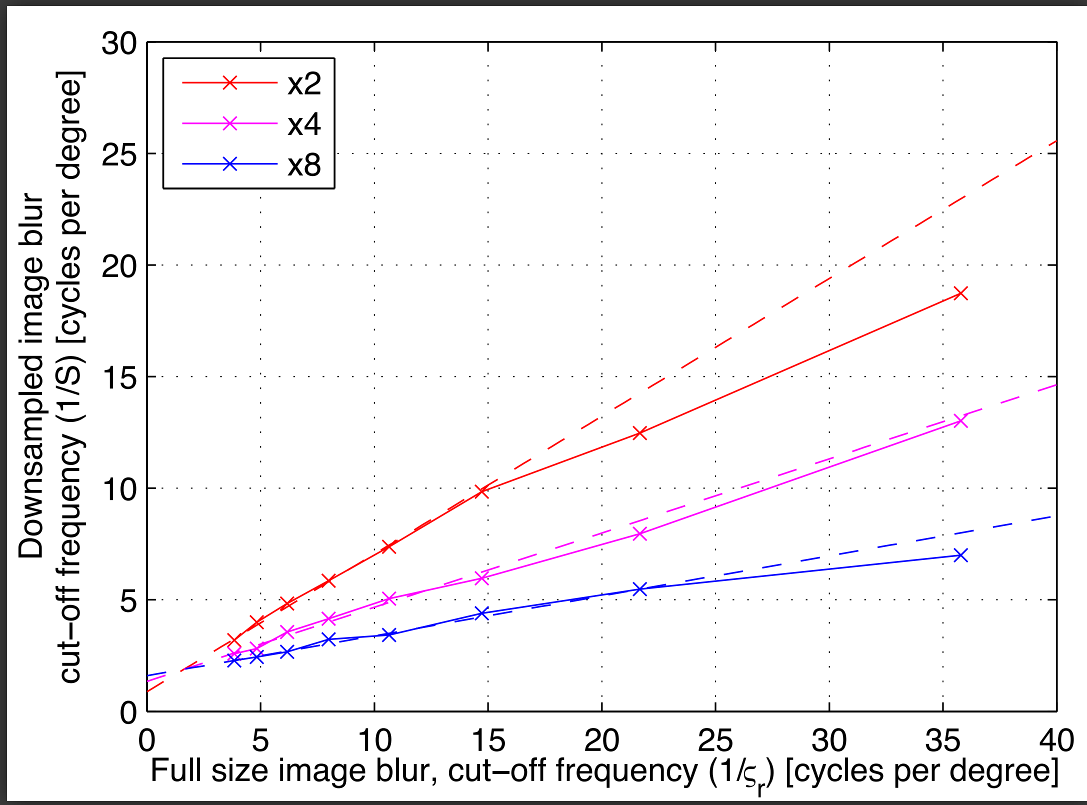

- Derive a model from this data
- Allows us to interpolate and extrapolate to cover cases not in our experiments

- sigma_d approximates ideal anti-aliasing filter
- is represented as the least squares fit of a Gaussian to the sinc function in cycles per degree

- Well aligned besides several high frequency measurements in 2x downsample
- Attribute to measurement error magnified by 1/sigma

- Have supplementary materials to demonstrate model on a number of images not included in the study

- Use this model to determine the desired about of blur in downsampled image
- But first need to determine how much blur is already present

Blur appearance model

- Measured data ς_m well predicted by anti-aliasing filter ς_d and model $\mathcal{S}$ in spatial frequencies

$$\widehat{\varsigma}_m = \sqrt{\varsigma_d^2 + \mathcal{S}^2}$$

- After removing ς_d, we model $\mathcal{S}$ as a linear function in $1/\varsigma$ spatial frequencies

- Full model provides accurate and plausible fit of the measured data in the spatial domain

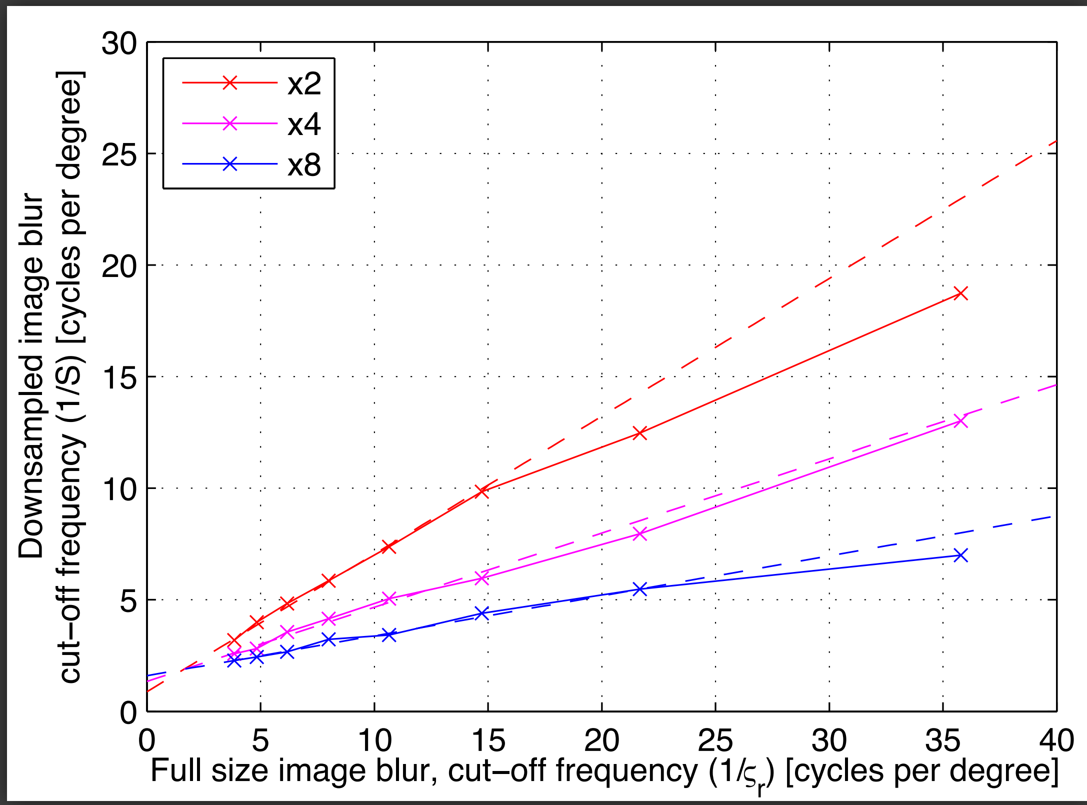

$$\mathcal{S}(\varsigma_r, d) = \frac{1}{2^{-0.893\log_2(d)+0.197}(\frac{1}{\varsigma_r} - 1.64) + 1.89}$$

- Derive a model from this data
- Allows us to interpolate and extrapolate to cover cases not in our experiments

- sigma_d approximates ideal anti-aliasing filter
- is represented as the least squares fit of a Gaussian to the sinc function in cycles per degree

- Well aligned besides several high frequency measurements in 2x downsample
- Attribute to measurement error magnified by 1/sigma

- Have supplementary materials to demonstrate model on a number of images not included in the study

- Use this model to determine the desired about of blur in downsampled image
- But first need to determine how much blur is already present

Blur appearance model

- Measured data ς_m well predicted by anti-aliasing filter ς_d and model $\mathcal{S}$ in spatial frequencies

$$\widehat{\varsigma}_m = \sqrt{\varsigma_d^2 + \mathcal{S}^2}$$

- After removing ς_d, we model $\mathcal{S}$ as a linear function in $1/\varsigma$ spatial frequencies

- Full model provides accurate and plausible fit of the measured data in the spatial domain

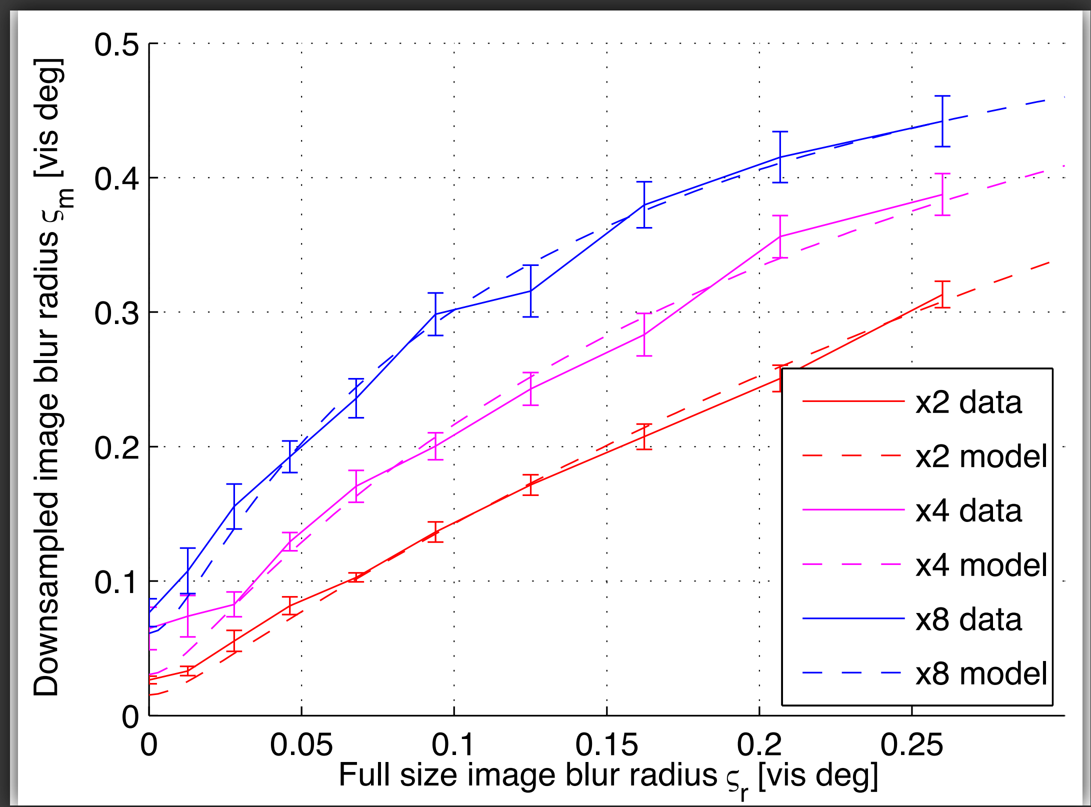

$$\mathcal{S}(\varsigma_r, d) = \frac{1}{2^{-0.893\log_2(d)+0.197}\left(\frac{1}{\varsigma_r} - 1.64\right) + 1.89}$$

- Derive a model from this data
- Allows us to interpolate and extrapolate to cover cases not in our experiments

- sigma_d approximates ideal anti-aliasing filter
- is represented as the least squares fit of a Gaussian to the sinc function in cycles per degree

- Well aligned besides several high frequency measurements in 2x downsample
- Attribute to measurement error magnified by 1/sigma

- Have supplementary materials to demonstrate model on a number of images not included in the study

- Use this model to determine the desired about of blur in downsampled image
- But first need to determine how much blur is already present

Blur estimation

- Spatially-variant estimate of the blur present at each pixel of image
- Calibrate method of Samadani *et al.* to provide estimate of blur in absolute units
- Downsampling approximates a blur-free image
- Relation between width of a Gaussian profile and the peak value of its derivative

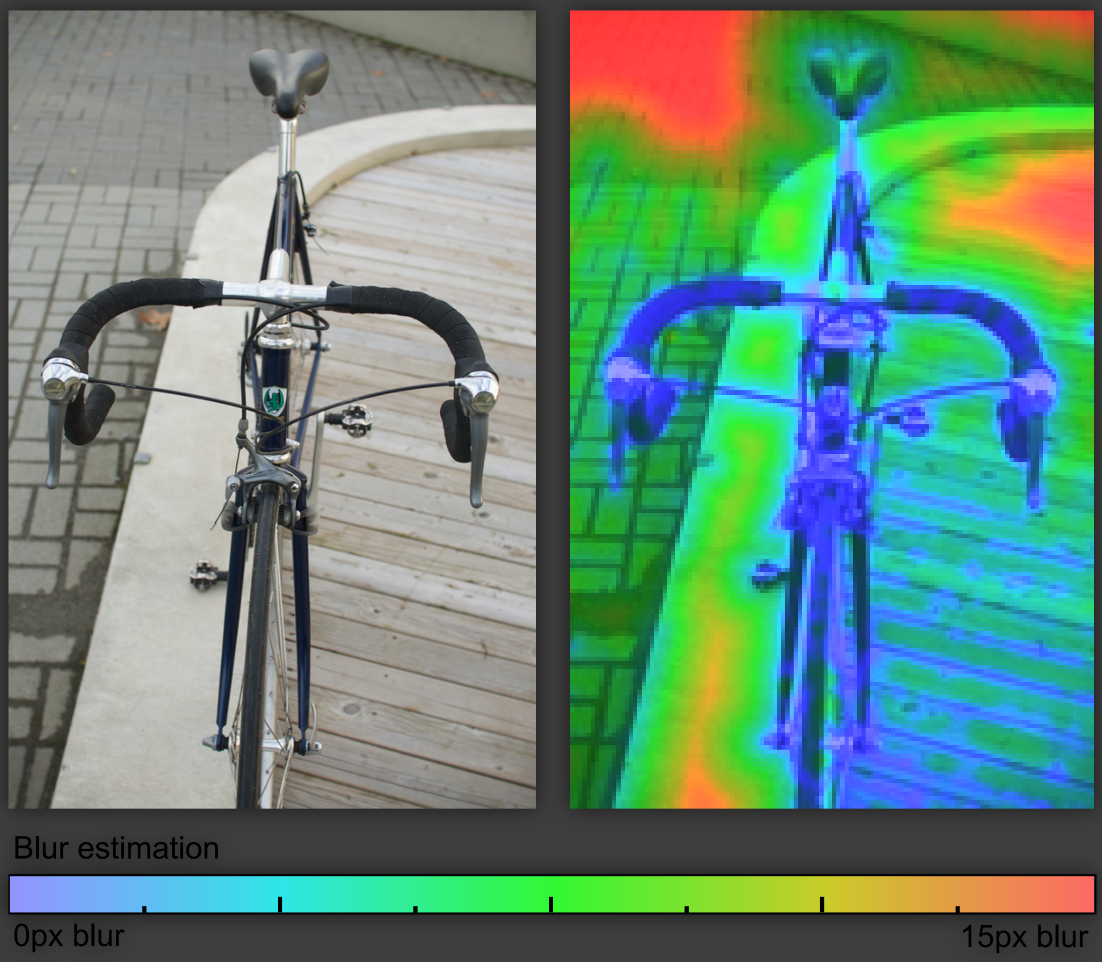

Chose their method because of its efficiency and the potential of an in-camera implementation

Blur estimation

Edge

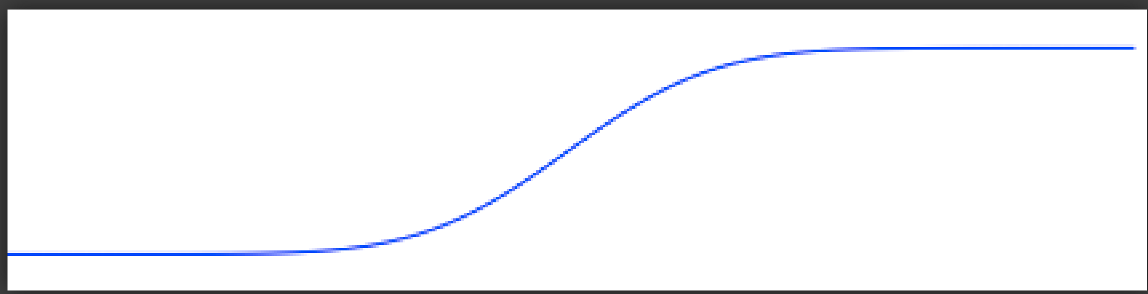

Derivative

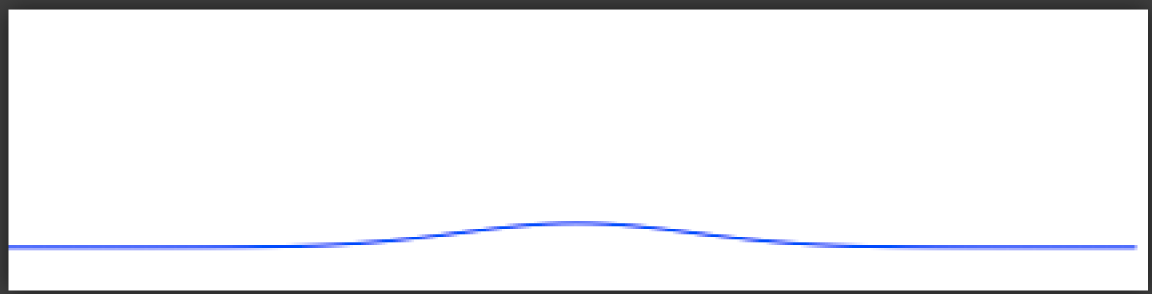

- Going to use the problem we hope to solve to help us estimate the blur
- Algo assumes that thumbnail provides a nearly blur-free approximation of image

- Demonstrate using 1D Gaussian profile
- Blur is reduced as the image is downsampled
- Thumbnail blurred edge approximates a step edge

Blur estimation

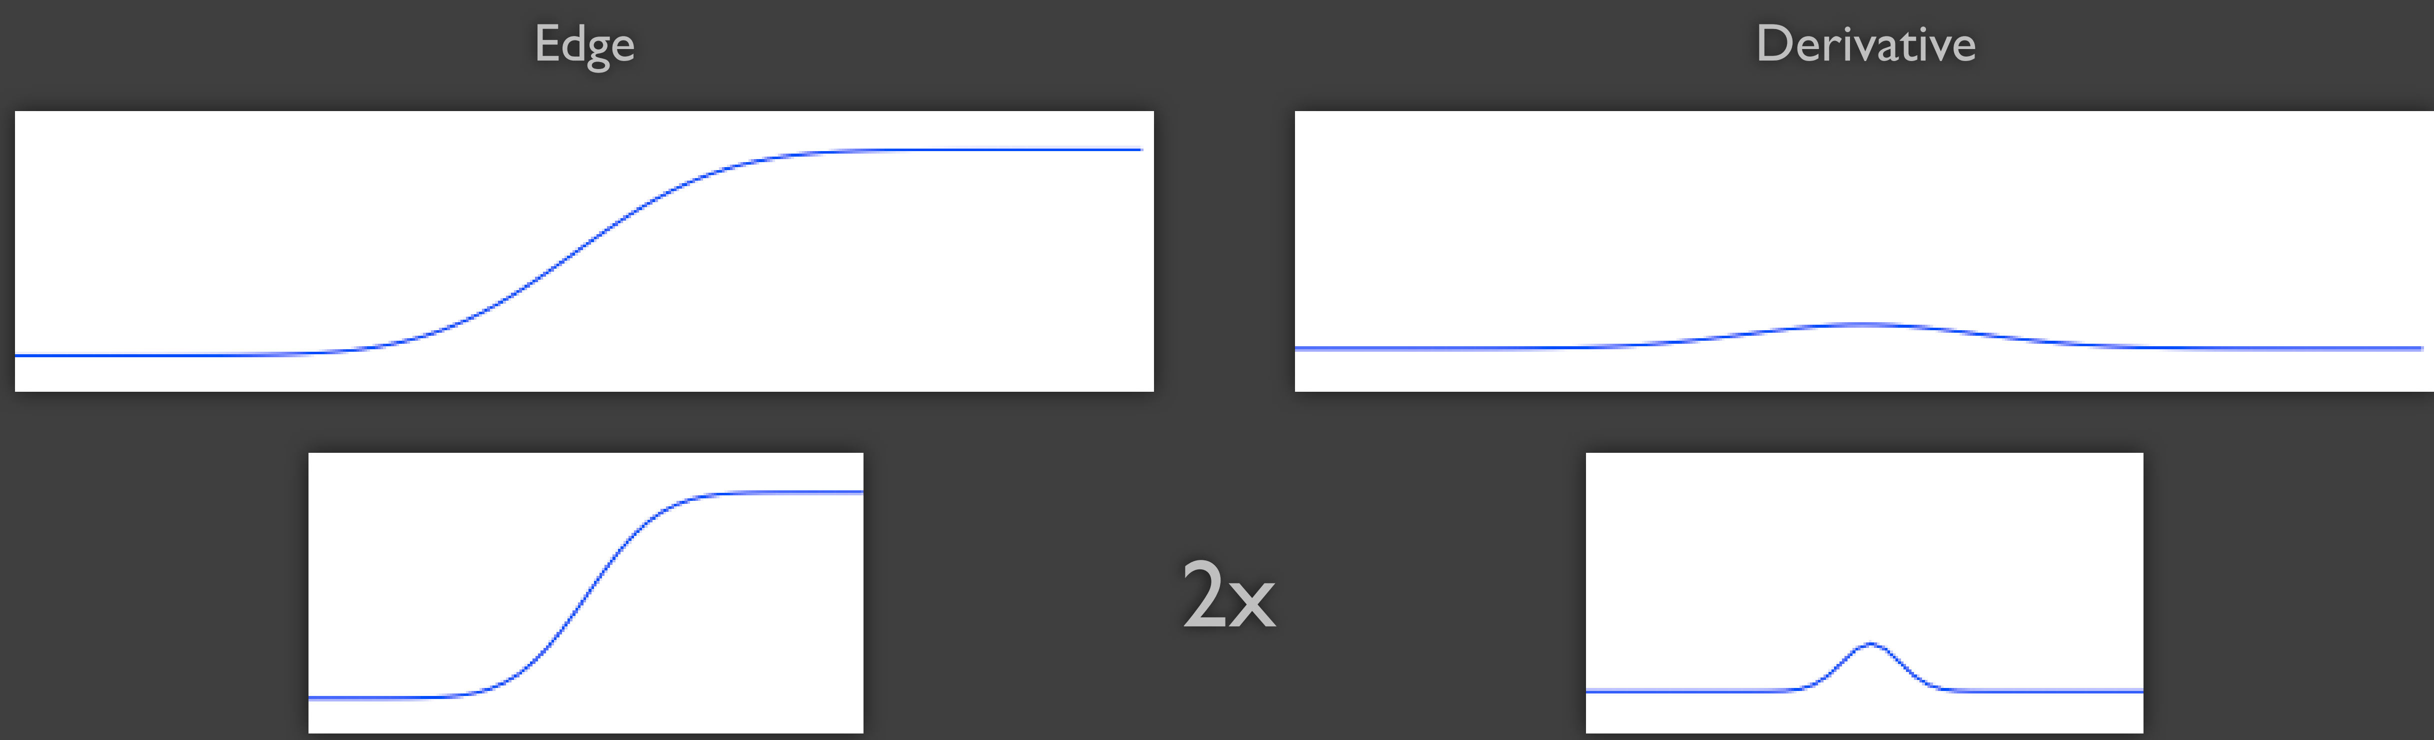

- Going to use the problem we hope to solve to help us estimate the blur
- Algo assumes that thumbnail provides a nearly blur-free approximation of image

- Demonstrate using 1D Gaussian profile
- Blur is reduced as the image is downsampled
- Thumbnail blurred edge approximates a step edge

Blur estimation

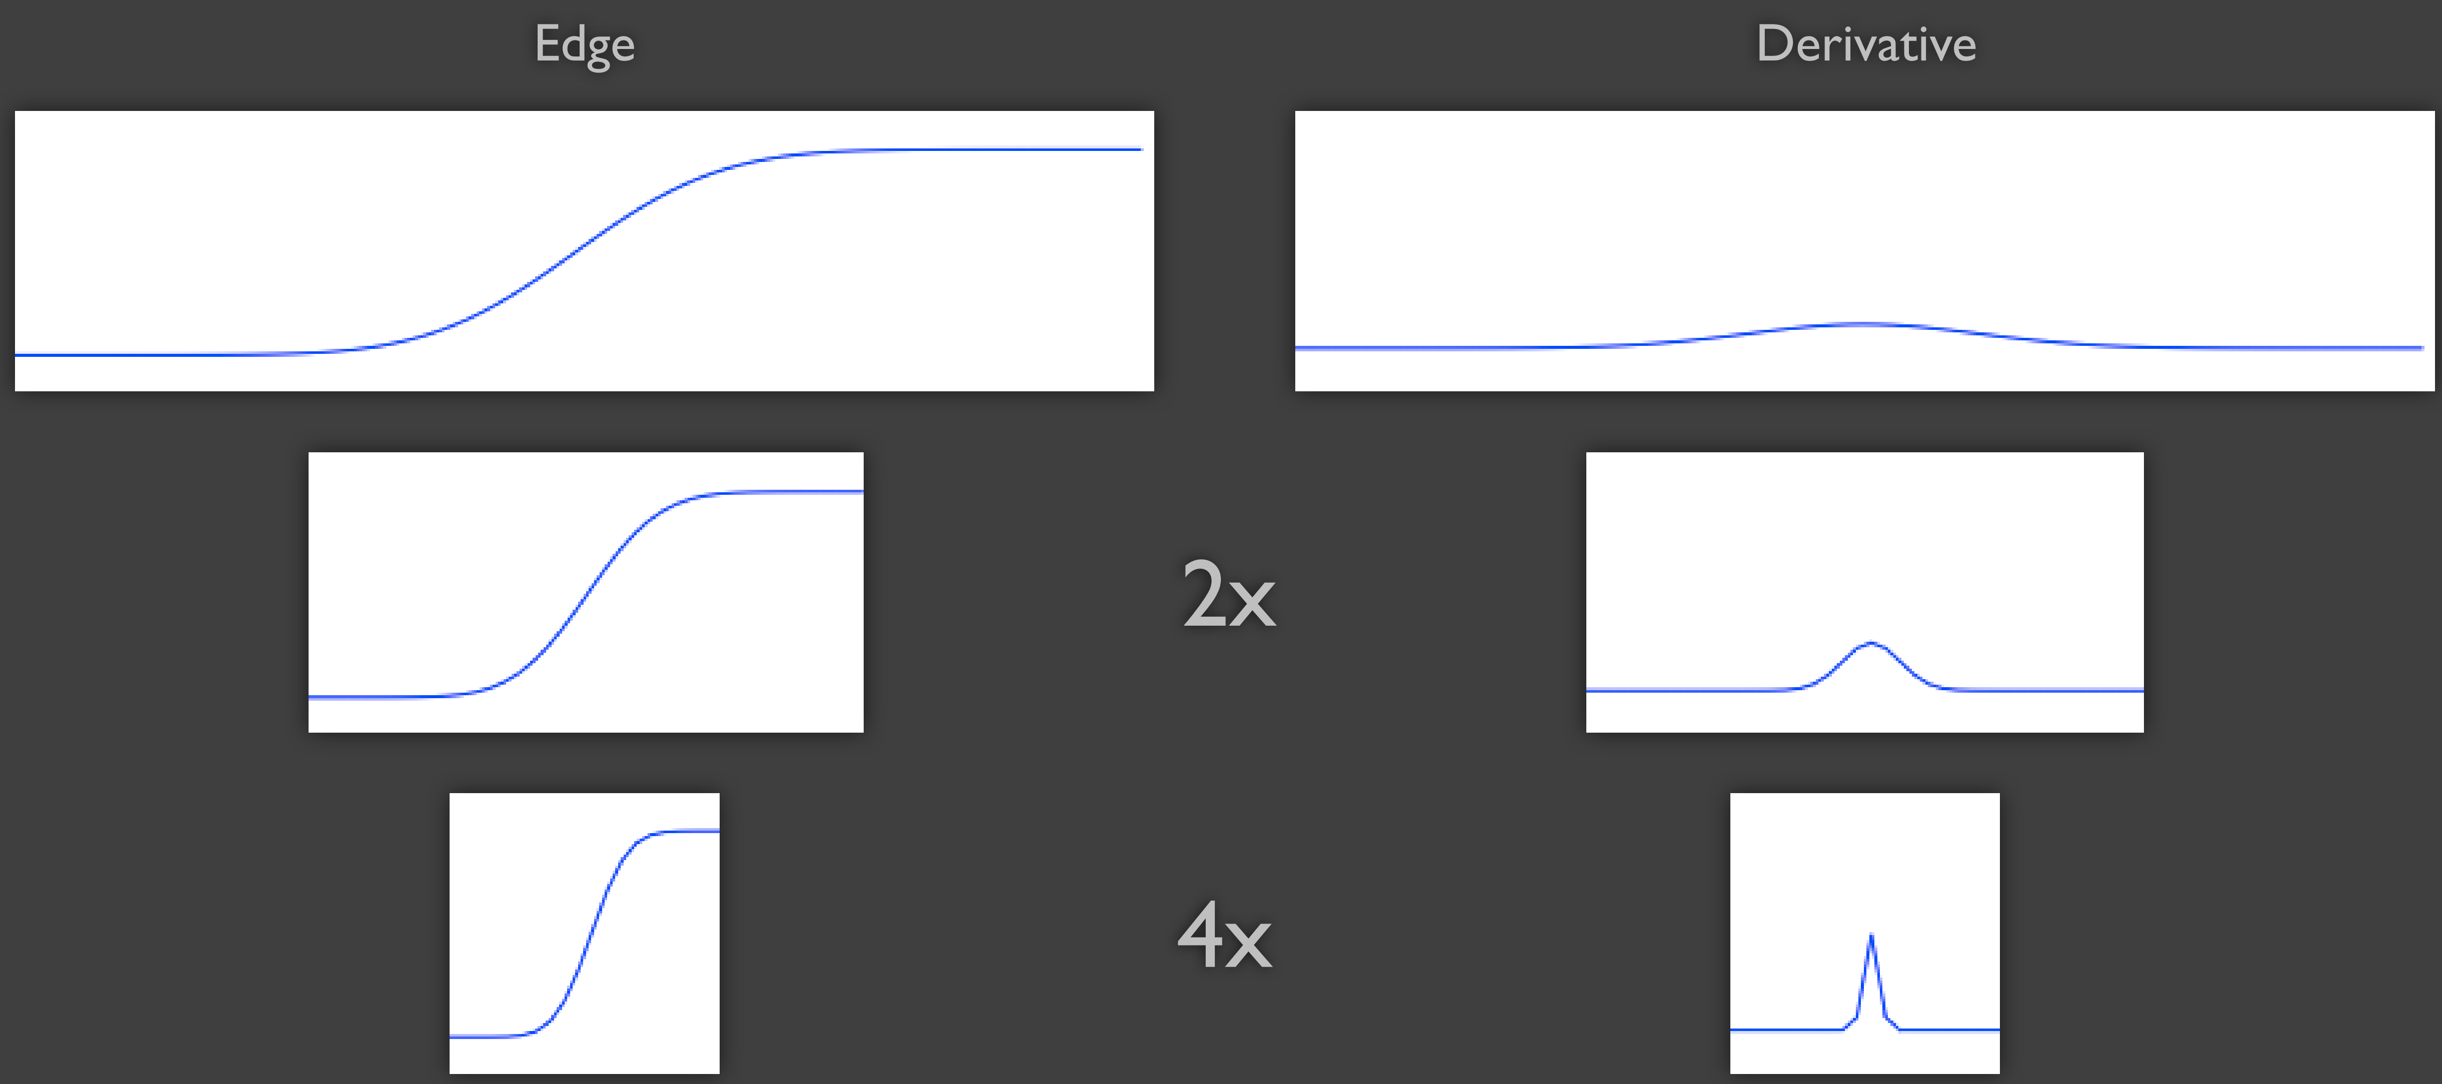

- Going to use the problem we hope to solve to help us estimate the blur
- Algo assumes that thumbnail provides a nearly blur-free approximation of image

- Demonstrate using 1D Gaussian profile
- Blur is reduced as the image is downsampled
- Thumbnail blurred edge approximates a step edge

Blur estimation

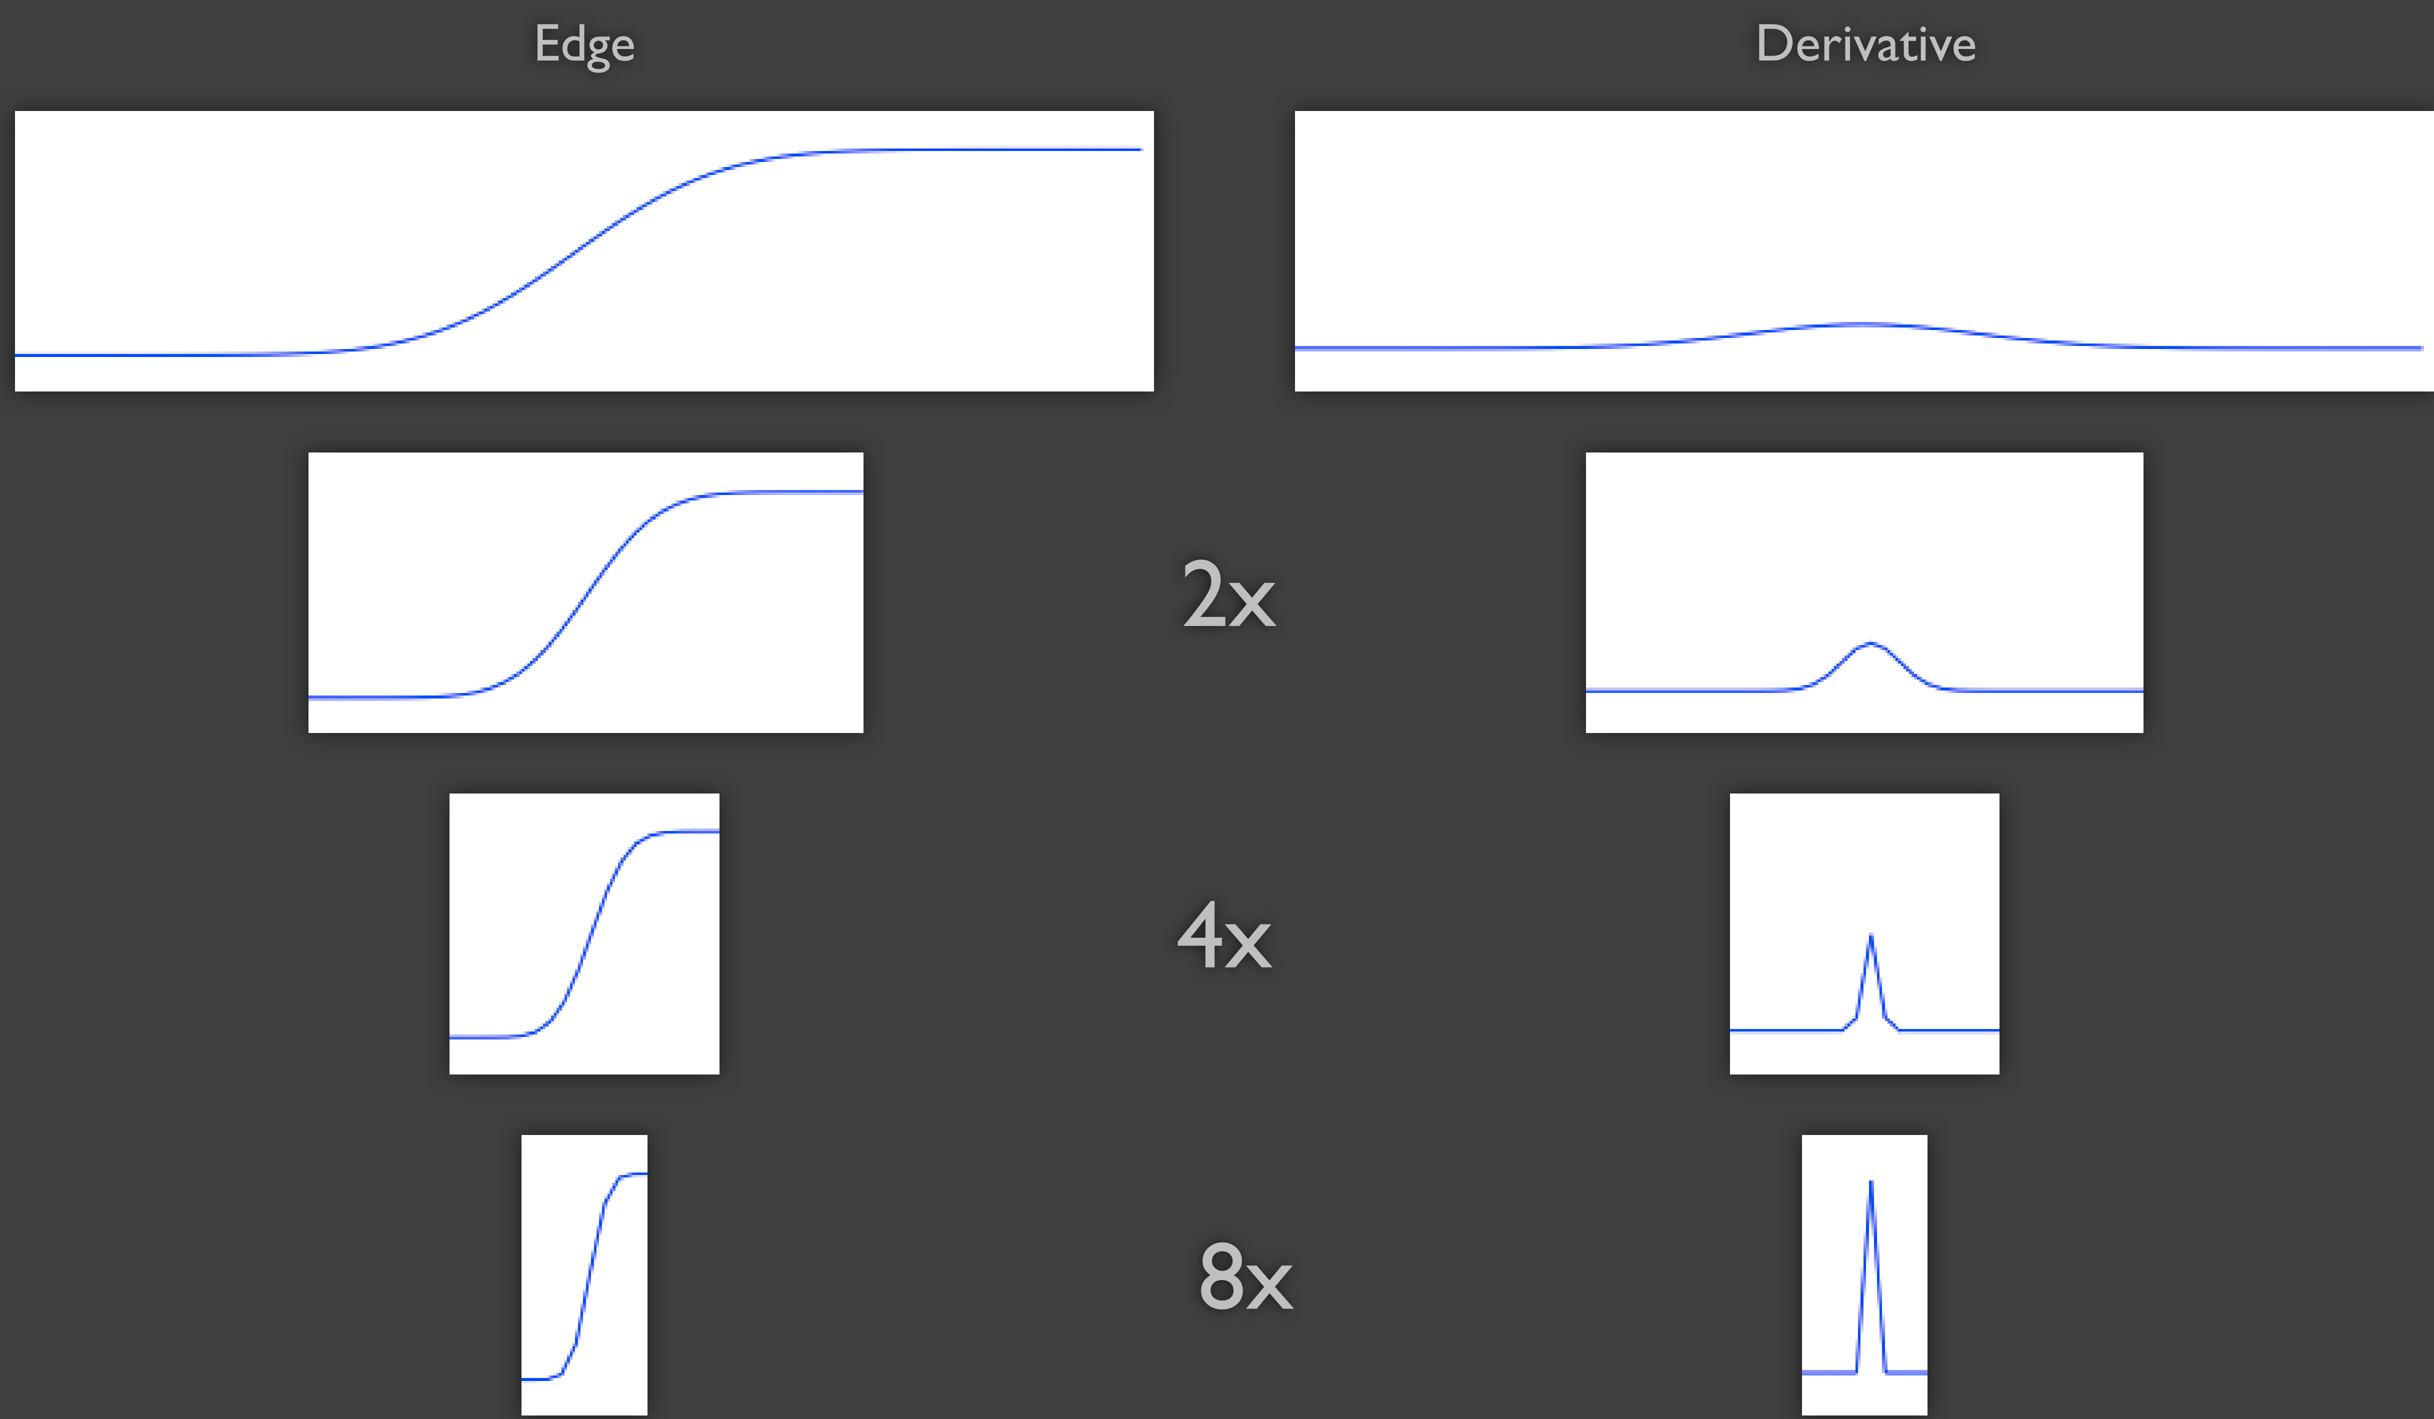

- Going to use the problem we hope to solve to help us estimate the blur
- Algo assumes that thumbnail provides a nearly blur-free approximation of image

- Demonstrate using 1D Gaussian profile
- Blur is reduced as the image is downsampled
- Thumbnail blurred edge approximates a step edge

Blur estimation

Edge

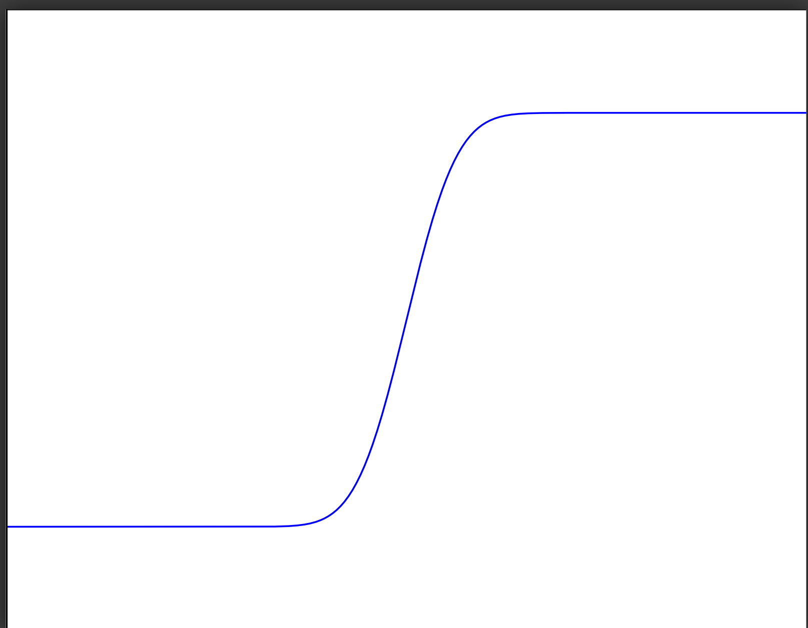

width: σ

Gradient magnitude

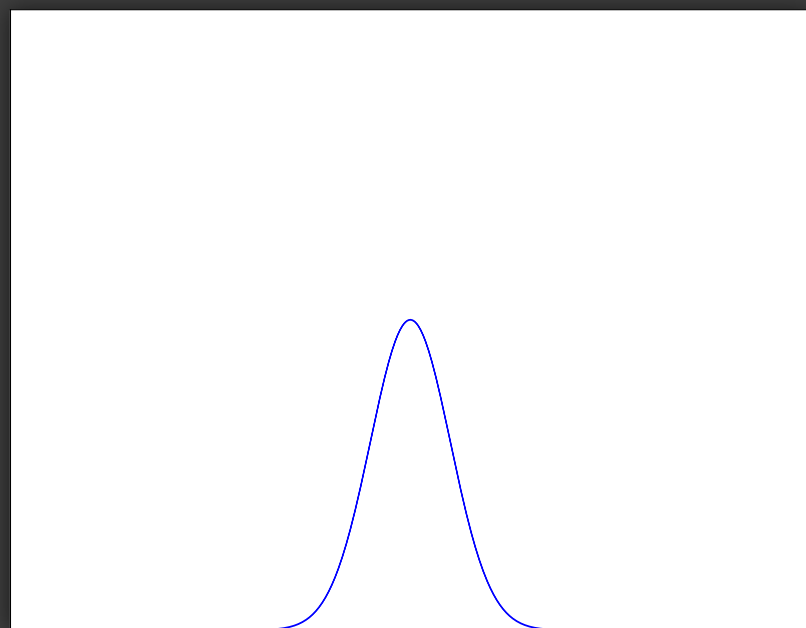

$$g\left(x,\sigma\right)=\frac{1}{\sqrt{2\pi\sigma^2}}e^{-\frac{x^2}{\sqrt{2\sigma^2}}}$$

- Correspondence between blur present and gradient magnitude

- Compare gradient magnitude of original image with the stronger gradients in our thumbnail
- Blur the thumbnail to have its gradients match those of original image

- Construct a scalespace of increasing blurs
- The thumbnail with the gradient magnitude closest to that of our original image
- Tells us how much blur is in original

Blur estimation

Edge

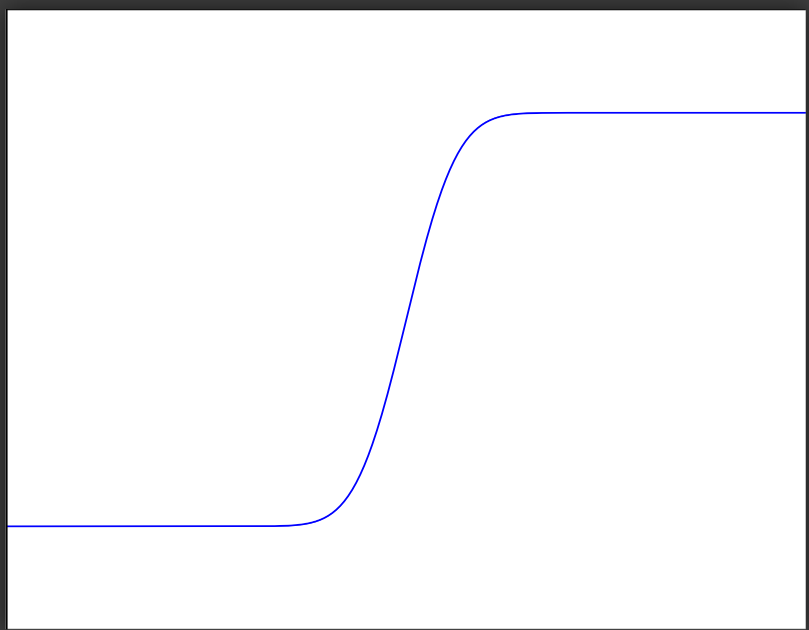

width: σ

Gradient magnitude

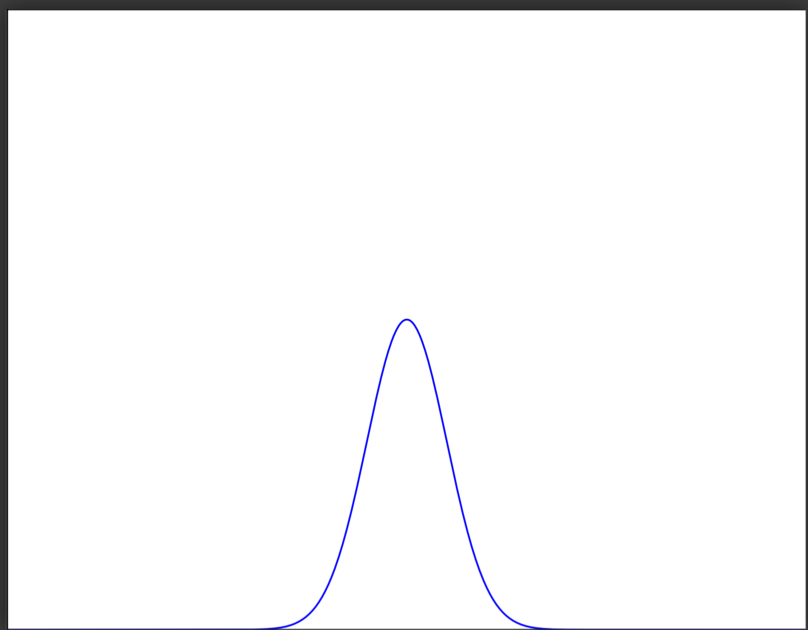

$$g\left(x,\sigma\right)=\frac{1}{\sqrt{2\pi\sigma^2}}e^{-\frac{x^2}{\sqrt{2\sigma^2}}}$$

- Correspondence between blur present and gradient magnitude

- Compare gradient magnitude of original image with the stronger gradients in our thumbnail
- Blur the thumbnail to have its gradients match those of original image

- Construct a scalespace of increasing blurs
- The thumbnail with the gradient magnitude closest to that of our original image
- Tells us how much blur is in original

Blur estimation

Edge

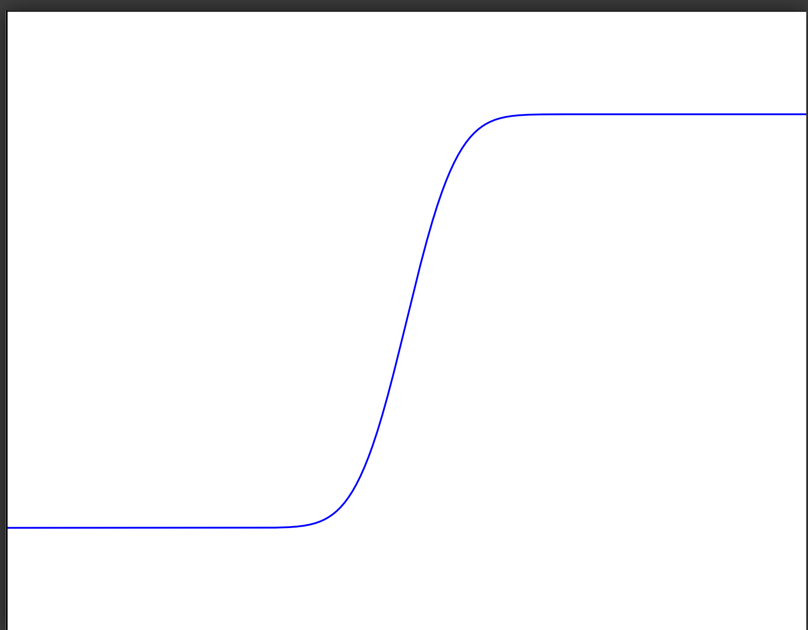

width: σ

Gradient magnitude

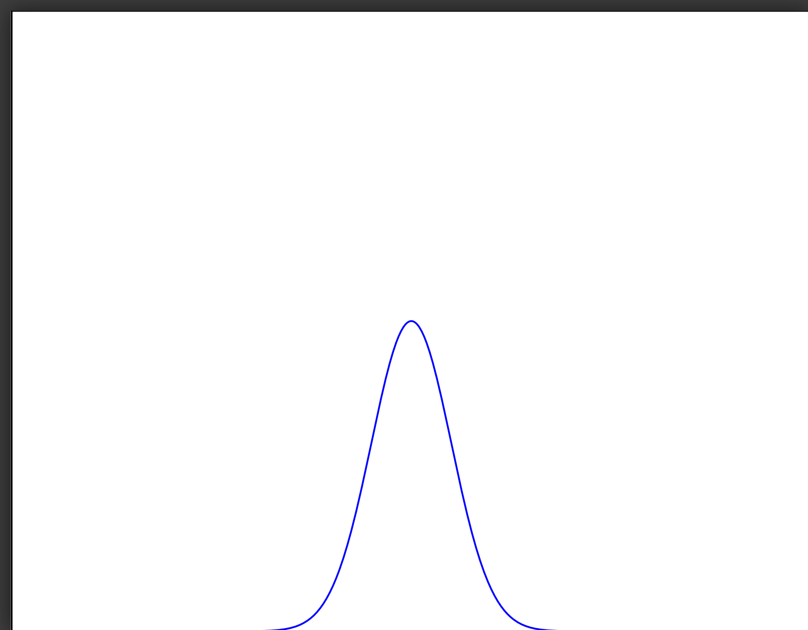

$$g(x, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{x^2}{\sqrt{2\sigma^2}}}$$

Downsampled scale space

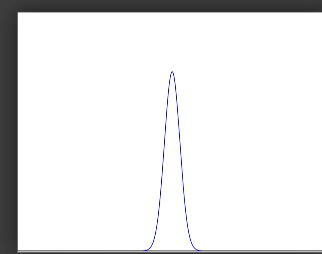

- Correspondence between blur present and gradient magnitude

- Compare gradient magnitude of original image with the stronger gradients in our thumbnail
- Blur the thumbnail to have its gradients match those of original image

- Construct a scalespace of increasing blurs
- The thumbnail with the gradient magnitude closest to that of our original image
- Tells us how much blur is in original

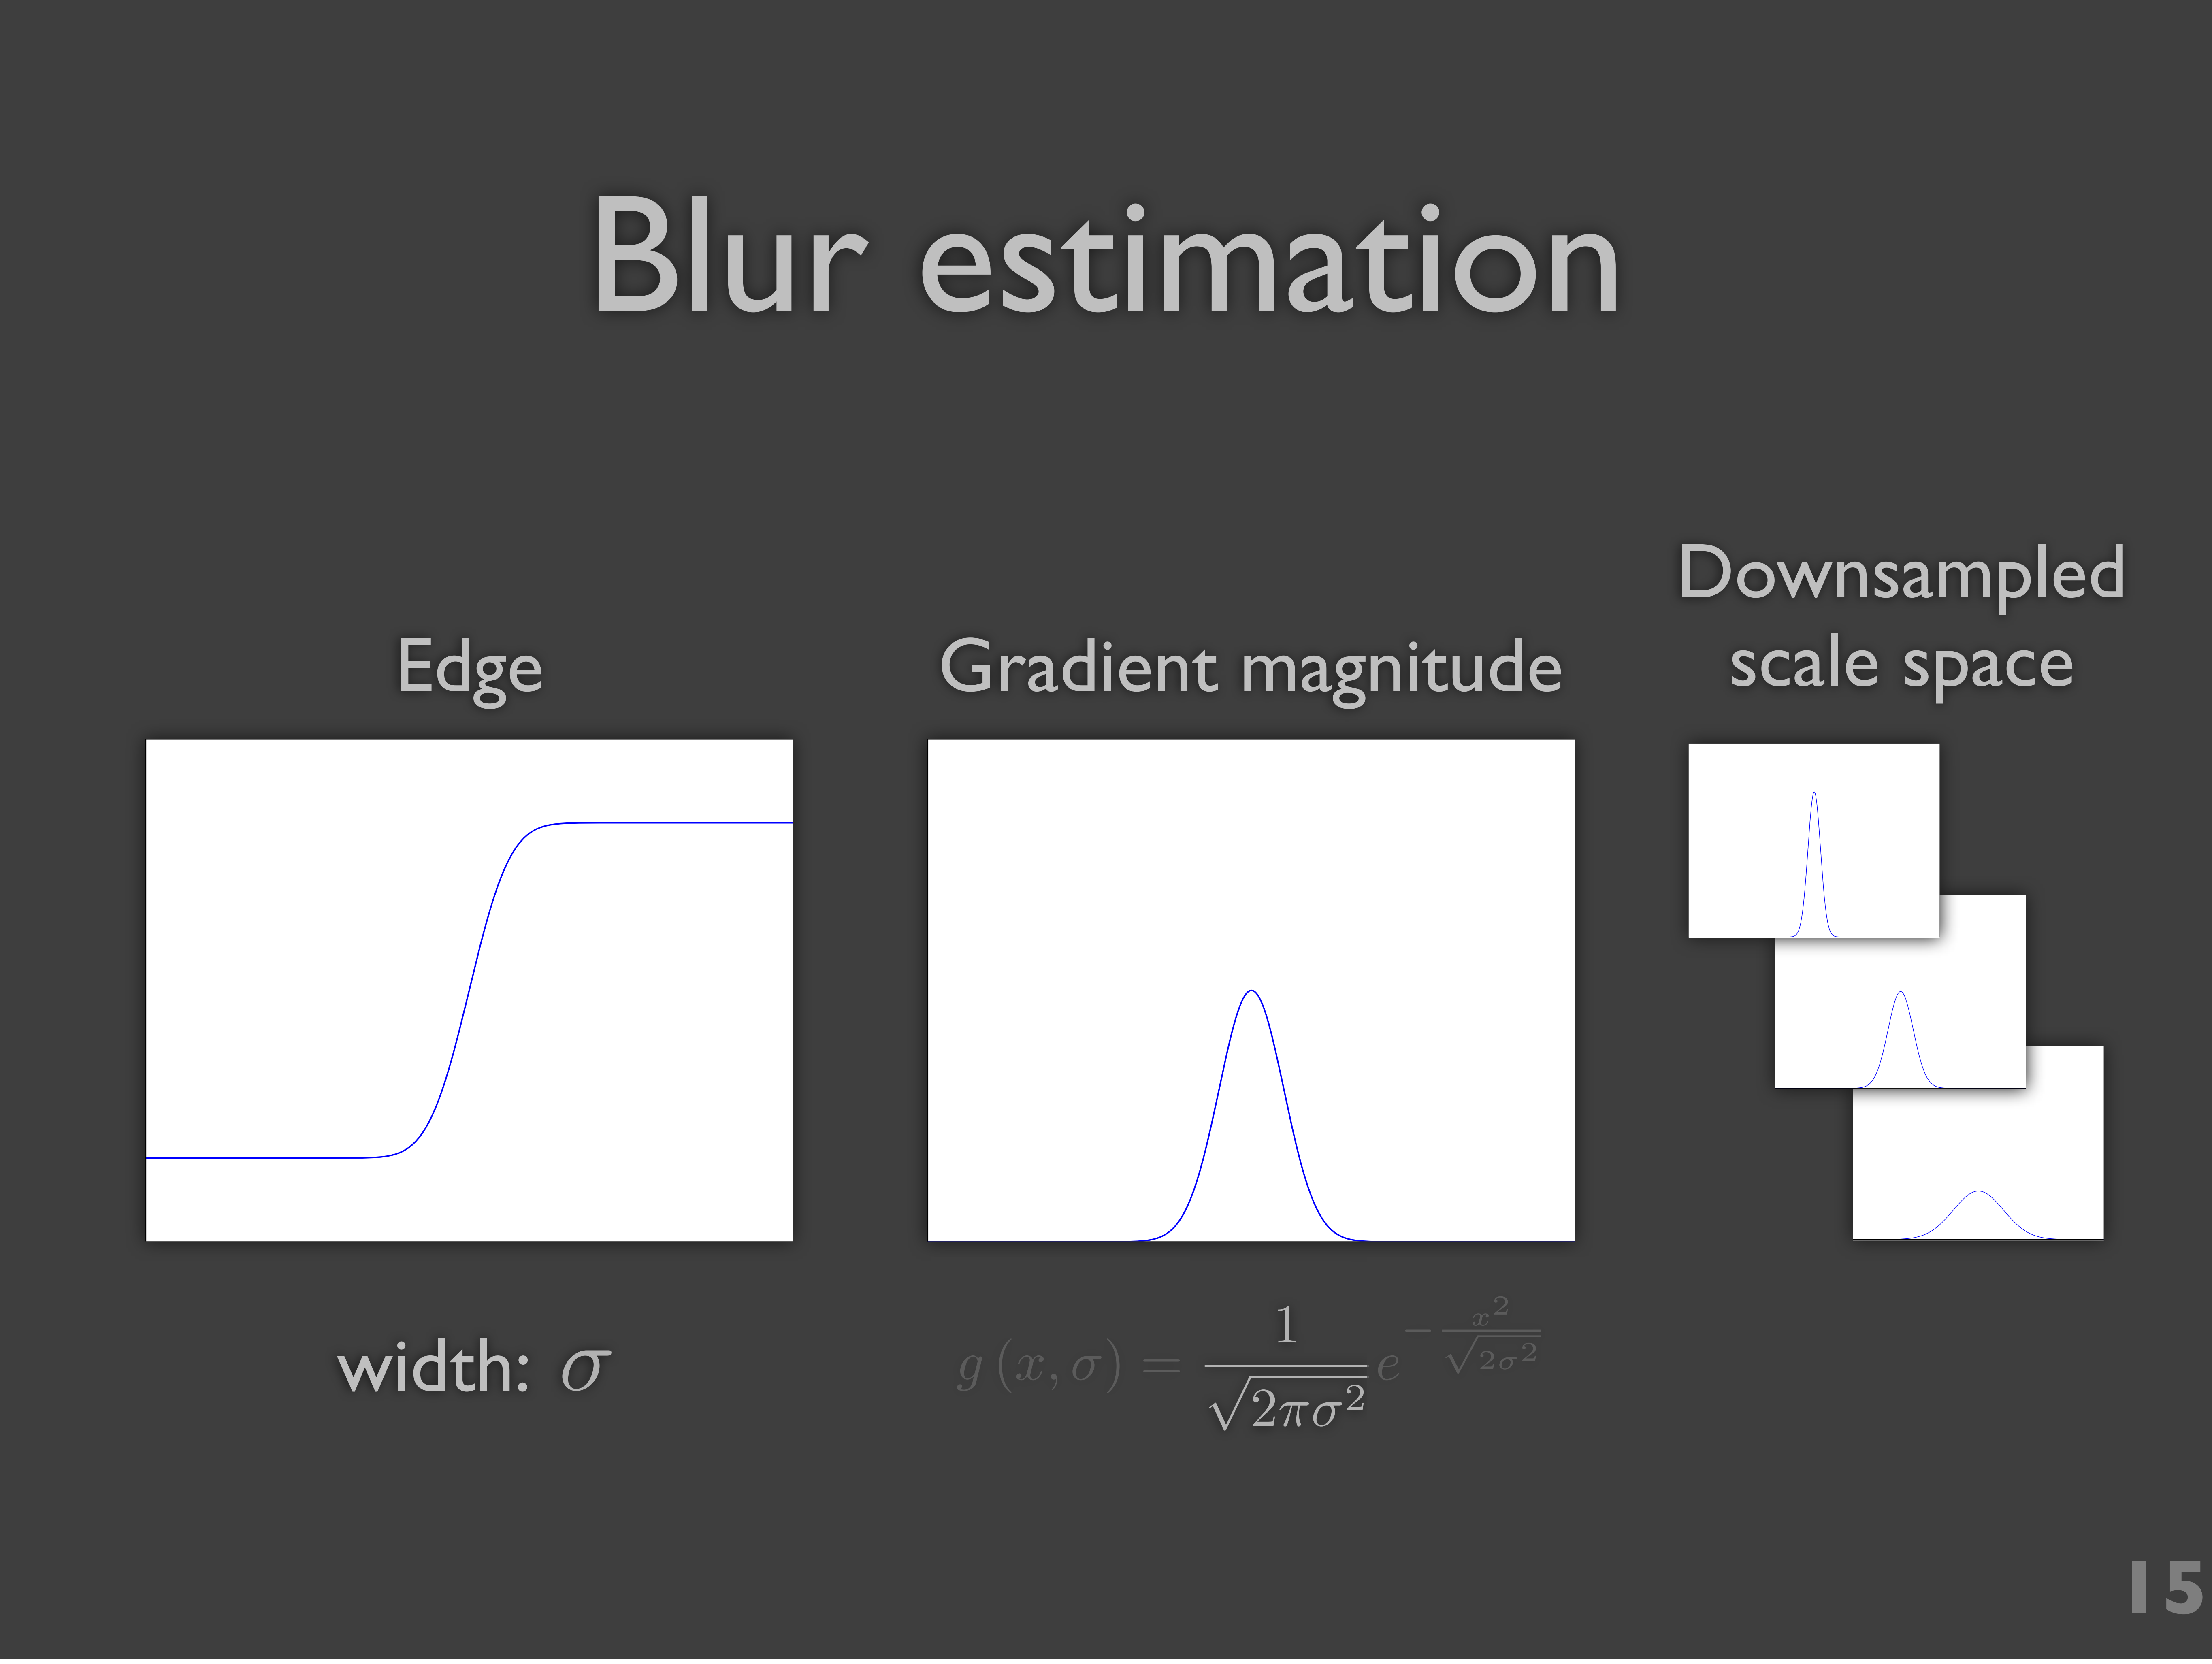

- Correspondence between blur present and gradient magnitude

- Compare gradient magnitude of original image with the stronger gradients in our thumbnail
- Blur the thumbnail to have its gradients match those of original image

- Construct a scalespace of increasing blurs
- The thumbnail with the gradient magnitude closest to that of our original image
- Tells us how much blur is in original

Blur estimation

Edge

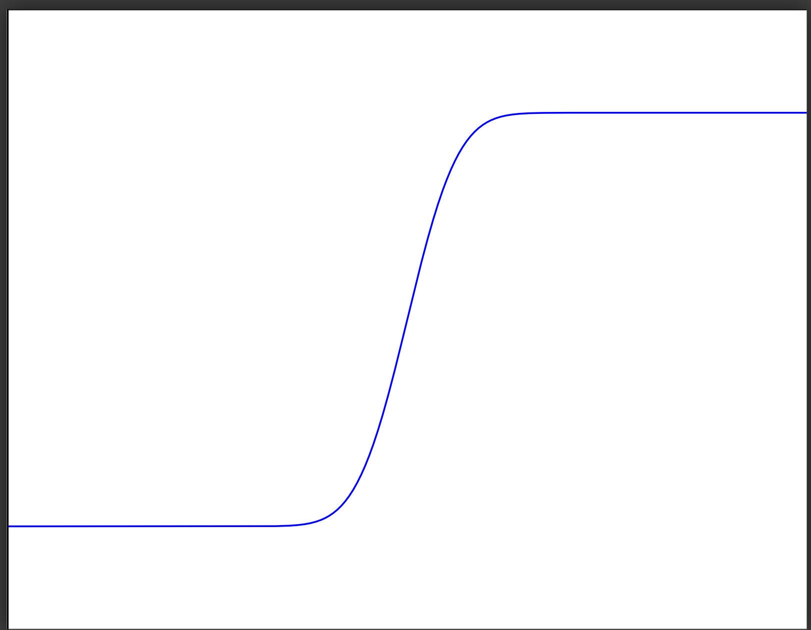

width: σ

Gradient magnitude

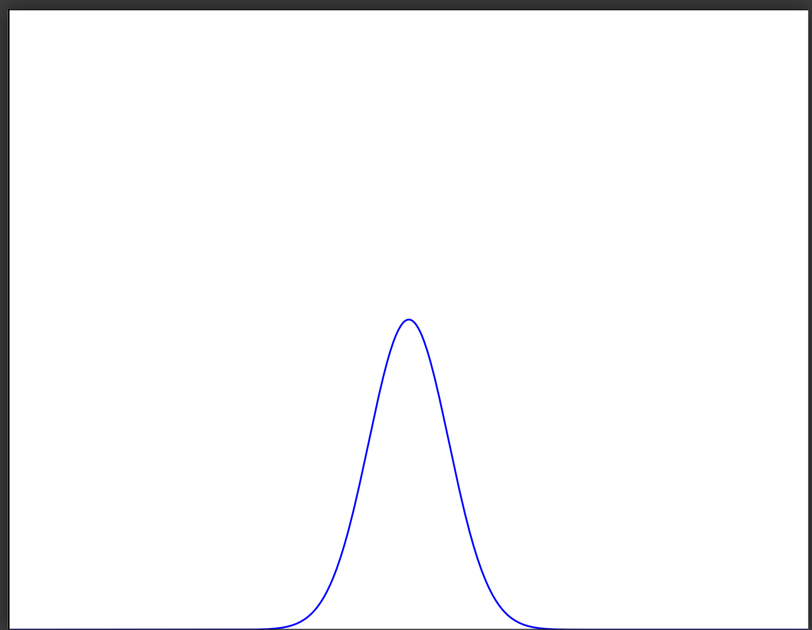

$$g(x, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{x^2}{\sqrt{2\sigma^2}}}$$

Downsampled scale space

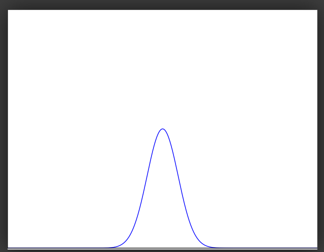

- Correspondence between blur present and gradient magnitude

- Compare gradient magnitude of original image with the stronger gradients in our thumbnail
- Blur the thumbnail to have its gradients match those of original image

- Construct a scalespace of increasing blurs
- The thumbnail with the gradient magnitude closest to that of our original image
- Tells us how much blur is in original

Blur estimation

original gradients

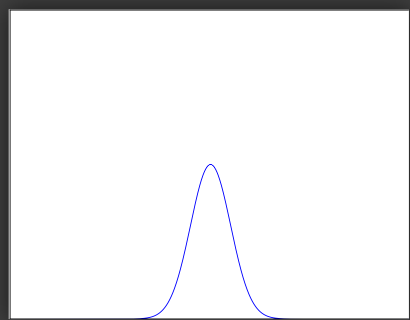

thumbnail gradients

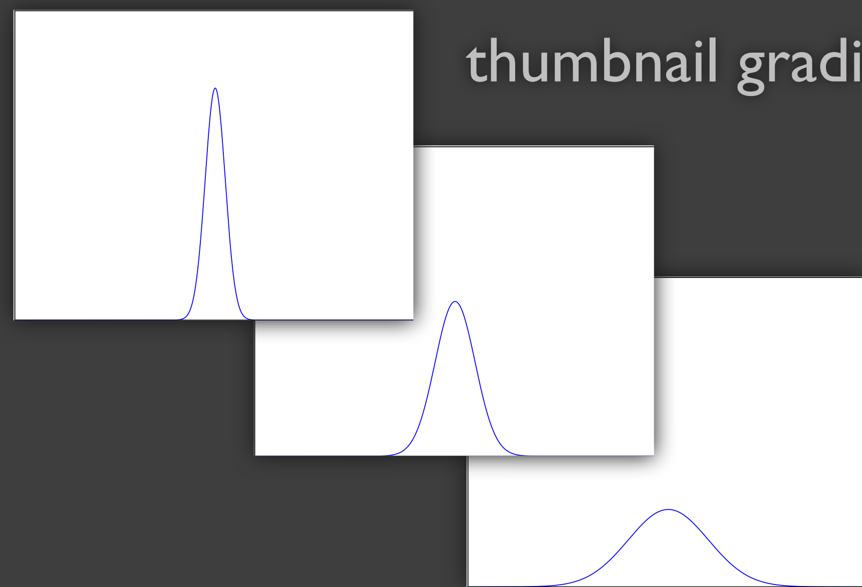

$$\frac{1}{\sqrt{2\pi\sigma_o^2}} \approx \frac{1}{\sqrt{2\pi\left[\left(\frac{\sigma_o^2}{d}\right)+(\beta j)^2\right]}}$$

d downsample

j scale space level

β quantization term

- Perform the estimation at the resolution of the downsampled image
- Downsample the gradients of the original image to the output resolution

- Define that If the original image blur is j
- We want the jth level of the scalespace to be equal to original gradients

- So, substitute j for sigma
- Introduce a scaling term gamma to correct for the difference
- Solve for the value of gamma in terms of downsample and quantization of scalespace
- To correctly align original and scalespace gradients

- Thus determining the appropriate level of the scalespace

Blur estimation

original gradients

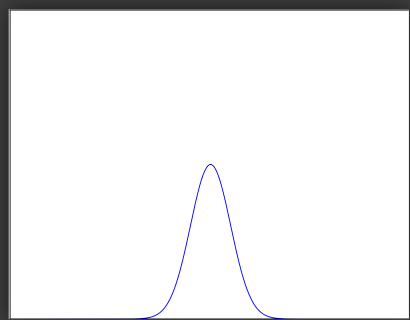

thumbnail gradients

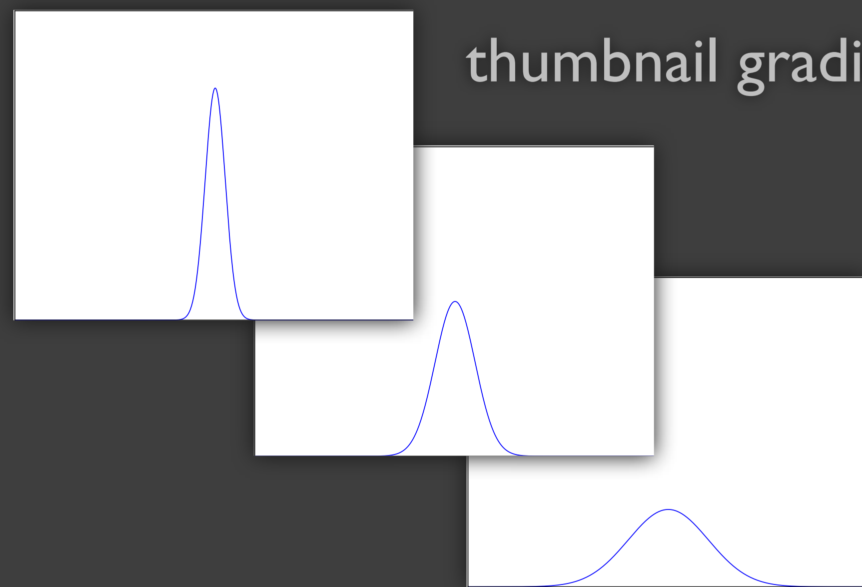

$$\frac{1}{\sqrt{2\pi j^2}} \approx \frac{1}{\sqrt{2\pi\left[\left(\frac{j^2}{d}\right)+(\beta j)^2\right]}}$$

d downsample

j scale space level

β quantization term

- Perform the estimation at the resolution of the downsampled image
- Downsample the gradients of the original image to the output resolution

- Define that If the original image blur is j
- We want the jth level of the scalespace to be equal to original gradients

- So, substitute j for sigma
- Introduce a scaling term gamma to correct for the difference
- Solve for the value of gamma in terms of downsample and quantization of scalespace
- To correctly align original and scalespace gradients

- Thus determining the appropriate level of the scalespace

Blur estimation

original gradients

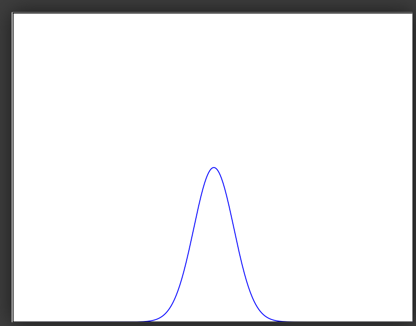

thumbnail gradients

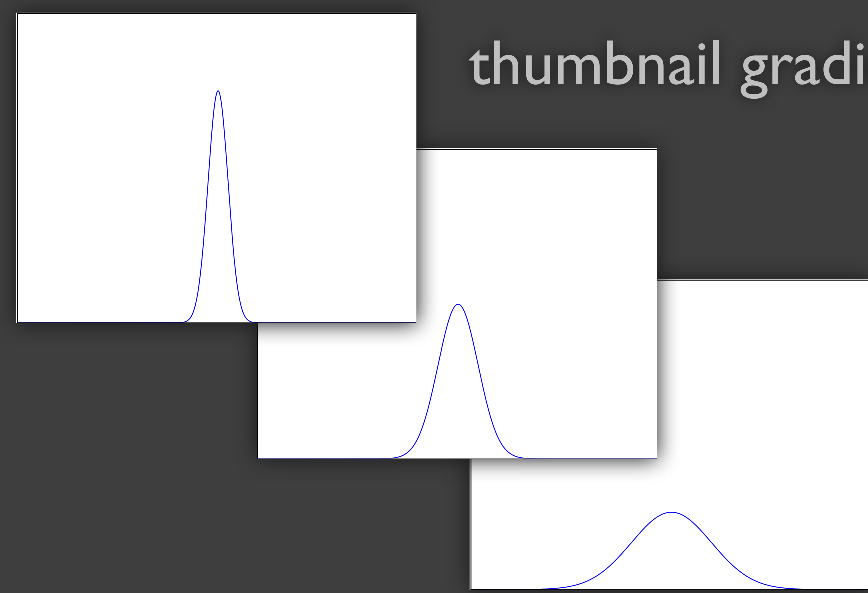

$$\gamma \frac{1}{\sqrt{2\pi j^2}} = \frac{1}{\sqrt{2\pi\left[\left(\frac{j^2}{d}\right)+(\beta j)^2\right]}}$$

d downsample
j scale space level
β quantization term

- Perform the estimation at the resolution of the downsampled image
- Downsample the gradients of the original image to the output resolution

- Define that If the original image blur is j
- We want the jth level of the scalespace to be equal to original gradients

- So, substitute j for sigma
- Introduce a scaling term gamma to correct for the difference
- Solve for the value of gamma in terms of downsample and quantization of scalespace
- To correctly align original and scalespace gradients

- Thus determining the appropriate level of the scalespace

Blur estimation

original gradients

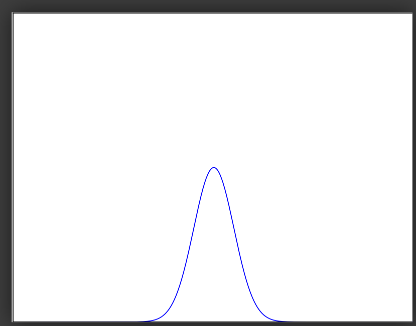

thumbnail gradients

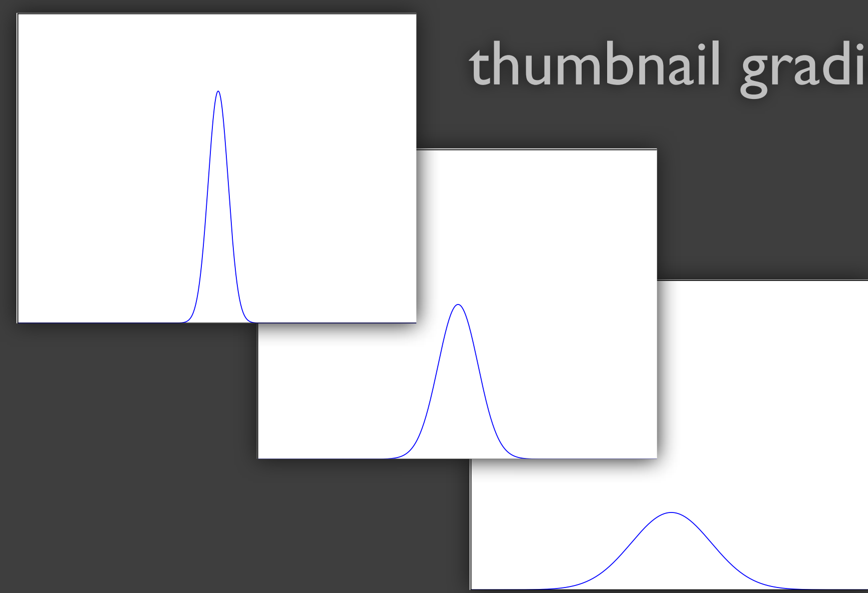

$$\gamma \frac{1}{\sqrt{2\pi j^2}} = \frac{1}{\sqrt{2\pi\left[\left(\frac{j^2}{d}\right) + (\beta j)^2\right]}}$$

$$\gamma = \frac{1}{\sqrt{\left(\frac{1}{d}\right)^2 + \beta^2}}$$

d downsample
j scale space level
β quantization term

- Perform the estimation at the resolution of the downsampled image
- Downsample the gradients of the original image to the output resolution

- Define that If the original image blur is j
- We want the jth level of the scalespace to be equal to original gradients

- So, substitute j for sigma
- Introduce a scaling term gamma to correct for the difference
- Solve for the value of gamma in terms of downsample and quantization of scalespace
- To correctly align original and scalespace gradients

- Thus determining the appropriate level of the scalespace

Blur estimation

original gradients

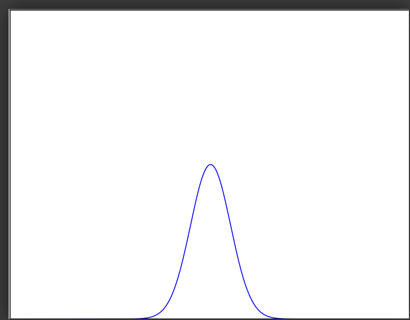

thumbnail gradients

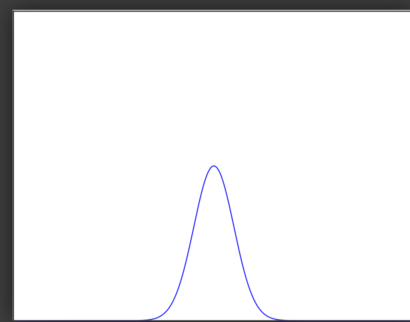

$$\gamma \frac{1}{\sqrt{2\pi j^2}} = \frac{1}{\sqrt{2\pi\left[\left(\frac{j^2}{d}\right)+(\beta j)^2\right]}}$$

$$\gamma = \frac{1}{\sqrt{\left(\frac{1}{d}\right)^2+\beta^2}}$$

d downsample
j scale space level
β quantization term

- Perform the estimation at the resolution of the downsampled image
- Downsample the gradients of the original image to the output resolution

- Define that If the original image blur is j
- We want the jth level of the scalespace to be equal to original gradients

- So, substitute j for sigma
- Introduce a scaling term gamma to correct for the difference
- Solve for the value of gamma in terms of downsample and quantization of scalespace
- To correctly align original and scalespace gradients

- Thus determining the appropriate level of the scalespace

Blur estimation

- Scaled original image gradients by gamma to align with scalespace
- If jth level is the closest match to r_o, implies a blur of j pixels in the original image
- Thus ensuring the estimate blur corresponds to some absolute measure of pixels

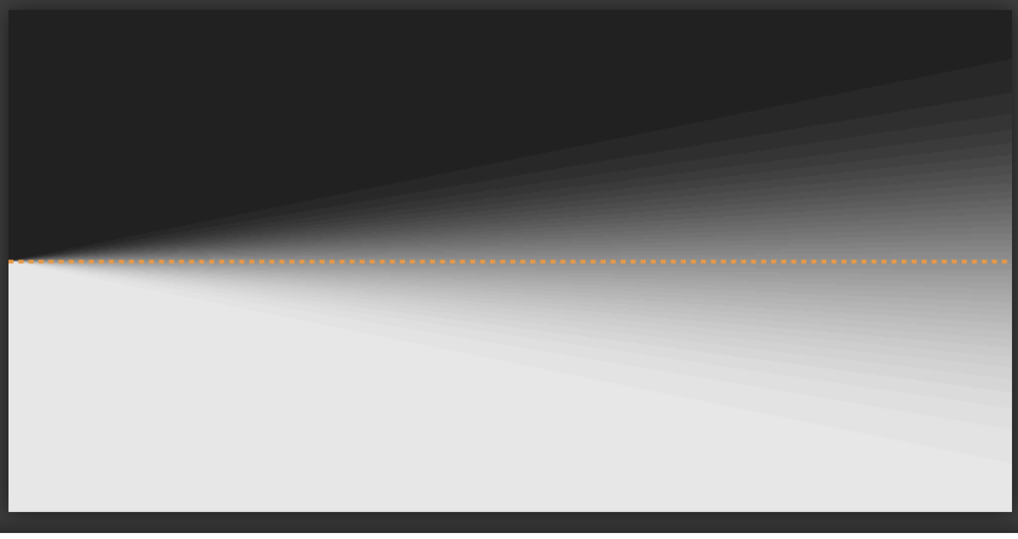

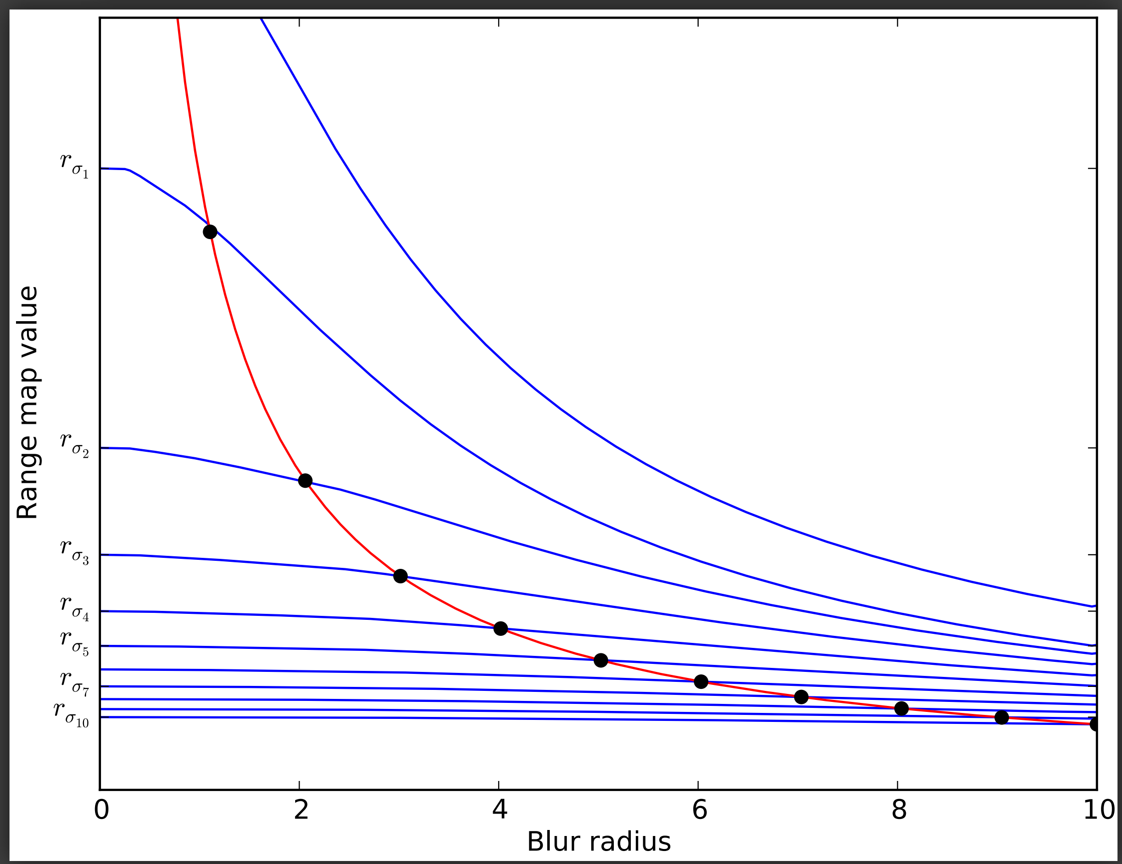

- Top image shows a black/white edge
- Increasing in blur of 0 to 10 pixels, left to right

- Bottom plot shows the original image gradients in red
- And different scalespace levels in blue
- Red intersects the jth blue level at x=j and we get the absolute blur

Blur synthesis

- Model specifies desired blur, give blur present determine how much to add
- Created thumbnail by standard downsample -- already includes anti-aliasing, so use model $\mathcal{S}$ instead of $\widehat{\varsigma}_m$
- Given existing blur σ_o compute blur to add σ_a

$$\sigma_a = \sqrt{\left(\frac{\mathcal{S}(\sigma_o \cdot p^{-1}, d) \cdot p}{d}\right)^2 - \sigma_o^2}$$

- d is downsample
- p is conversion between pixels and angular visual degrees
- 30 pixels per degree in a standard configuration

- Convert from from pixels to visual degrees
- Compute result of model (in visual degrees of the full image)
- Convert back to pixels and account for downsample
- Compute amount required given existing blur sigma_o using convolution of Gaussians theorem

Blur synthesis

- Model specifies desired blur, give blur present determine how much to add
- Created thumbnail by standard downsample -- already includes anti-aliasing, so use model $\mathcal{S}$ instead of $\widehat{\varsigma}_m$
- Given existing blur σ_o compute blur to add σ_a

$$\sigma_a = \sqrt{\left(\frac{\mathcal{S}(\sigma_o \cdot p^{-1}, d) \cdot p}{d}\right)^2 - \sigma_o^2}$$

- d is downsample
- p is conversion between pixels and angular visual degrees
- 30 pixels per degree in a standard configuration

- Convert from from pixels to visual degrees
- Compute result of model (in visual degrees of the full image)
- Convert back to pixels and account for downsample
- Compute amount required given existing blur sigma_o using convolution of Gaussians theorem

Blur synthesis

- Model specifies desired blur, give blur present determine how much to add
- Created thumbnail by standard downsample -- already includes anti-aliasing, so use model $\mathcal{S}$ instead of $\widehat{\varsigma}_m$
- Given existing blur σ_o compute blur to add σ_a

$$\sigma_a = \sqrt{\left(\frac{\mathcal{S}(\sigma_o \cdot p^{-1}, d) \cdot p}{d}\right)^2 - \sigma_o^2}$$

- d is downsample
- p is conversion between pixels and angular visual degrees
- 30 pixels per degree in a standard configuration

- Convert from from pixels to visual degrees
- Compute result of model (in visual degrees of the full image)
- Convert back to pixels and account for downsample
- Compute amount required given existing blur sigma_o using convolution of Gaussians theorem

Blur synthesis

- Model specifies desired blur, give blur present determine how much to add
- Created thumbnail by standard downsample -- already includes anti-aliasing, so use model $\mathcal{S}$ instead of $\widehat{\varsigma}_m$
- Given existing blur σ_o compute blur to add σ_a

$$\sigma_a = \sqrt{\left(\frac{\mathcal{S}(\sigma_o \cdot p^{-1}, d) \cdot p}{d}\right)^2 - \sigma_o^2}$$

- d is downsample
- p is conversion between pixels and angular visual degrees
- 30 pixels per degree in a standard configuration

- Convert from from pixels to visual degrees
- Compute result of model (in visual degrees of the full image)
- Convert back to pixels and account for downsample
- Compute amount required given existing blur sigma_o using convolution of Gaussians theorem

Blur synthesis

- Model specifies desired blur, give blur present determine how much to add
- Created thumbnail by standard downsample -- already includes anti-aliasing, so use model $\mathcal{S}$ instead of $\widehat{\varsigma}_m$
- Given existing blur σ_o compute blur to add σ_a

$$\sigma_a = \sqrt{\left(\frac{\mathcal{S}(\sigma_o \cdot p^{-1}, d) \cdot p}{d}\right)^2 - \sigma_o^2}$$

- d is downsample
- p is conversion between pixels and angular visual degrees
- 30 pixels per degree in a standard configuration

- Convert from from pixels to visual degrees
- Compute result of model (in visual degrees of the full image)
- Convert back to pixels and account for downsample
- Compute amount required given existing blur sigma_o using convolution of Gaussians theorem

Blur synthesis

- Model specifies desired blur, give blur present determine how much to add
- Created thumbnail by standard downsample -- already includes anti-aliasing, so use model $\mathcal{S}$ instead of $\widehat{\varsigma}_m$
- Given existing blur σ_o compute blur to add σ_a

$$\sigma_a = \sqrt{\left(\frac{\mathcal{S}(\sigma_o \cdot p^{-1}, d) \cdot p}{d}\right)^2 - \sigma_o^2}$$

- d is downsample
- p is conversion between pixels and angular visual degrees
- 30 pixels per degree in a standard configuration

- Convert from from pixels to visual degrees
- Compute result of model (in visual degrees of the full image)
- Convert back to pixels and account for downsample
- Compute amount required given existing blur sigma_o using convolution of Gaussians theorem

Blur synthesis

- To produce final image blur each level scalespace l_{σ_j} by corresponding σ_a, linearly blend for non-integer σ_a

Scalespace

+

Blur map

=

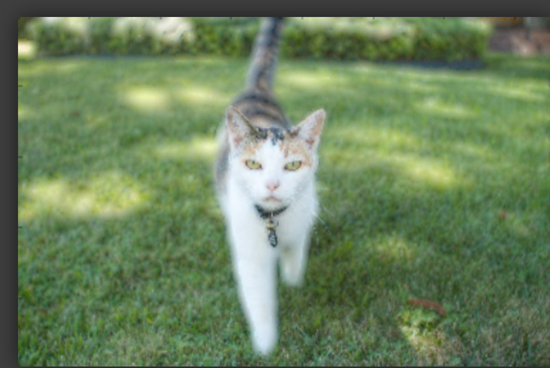

Final result

Evaluation

Naive

Samadani gamma=4

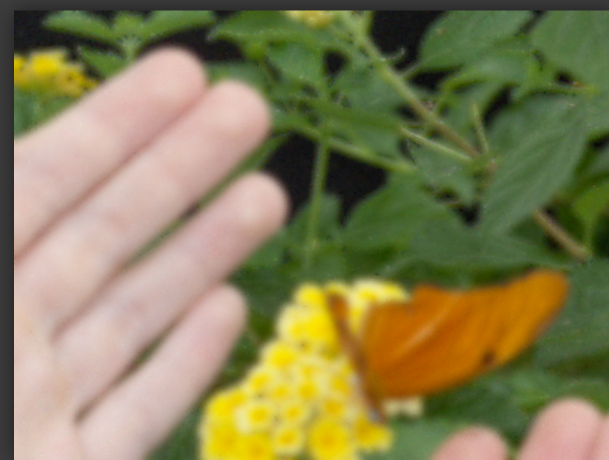

Samadani gamma=.5

Blur-Aware

- Another example is this poorly focused image, where the foreground is out of focus

- Samadani gamma = 4 is too sharp for hand and butterfly
- Samadani gamma = .5 is too blurry for leaves at top
- Our method blurs hand and flowers but leaves still in focus

- Again, viewing distance matters
- Depending on where you are in the hall, this will be more or less apparent

Evaluation

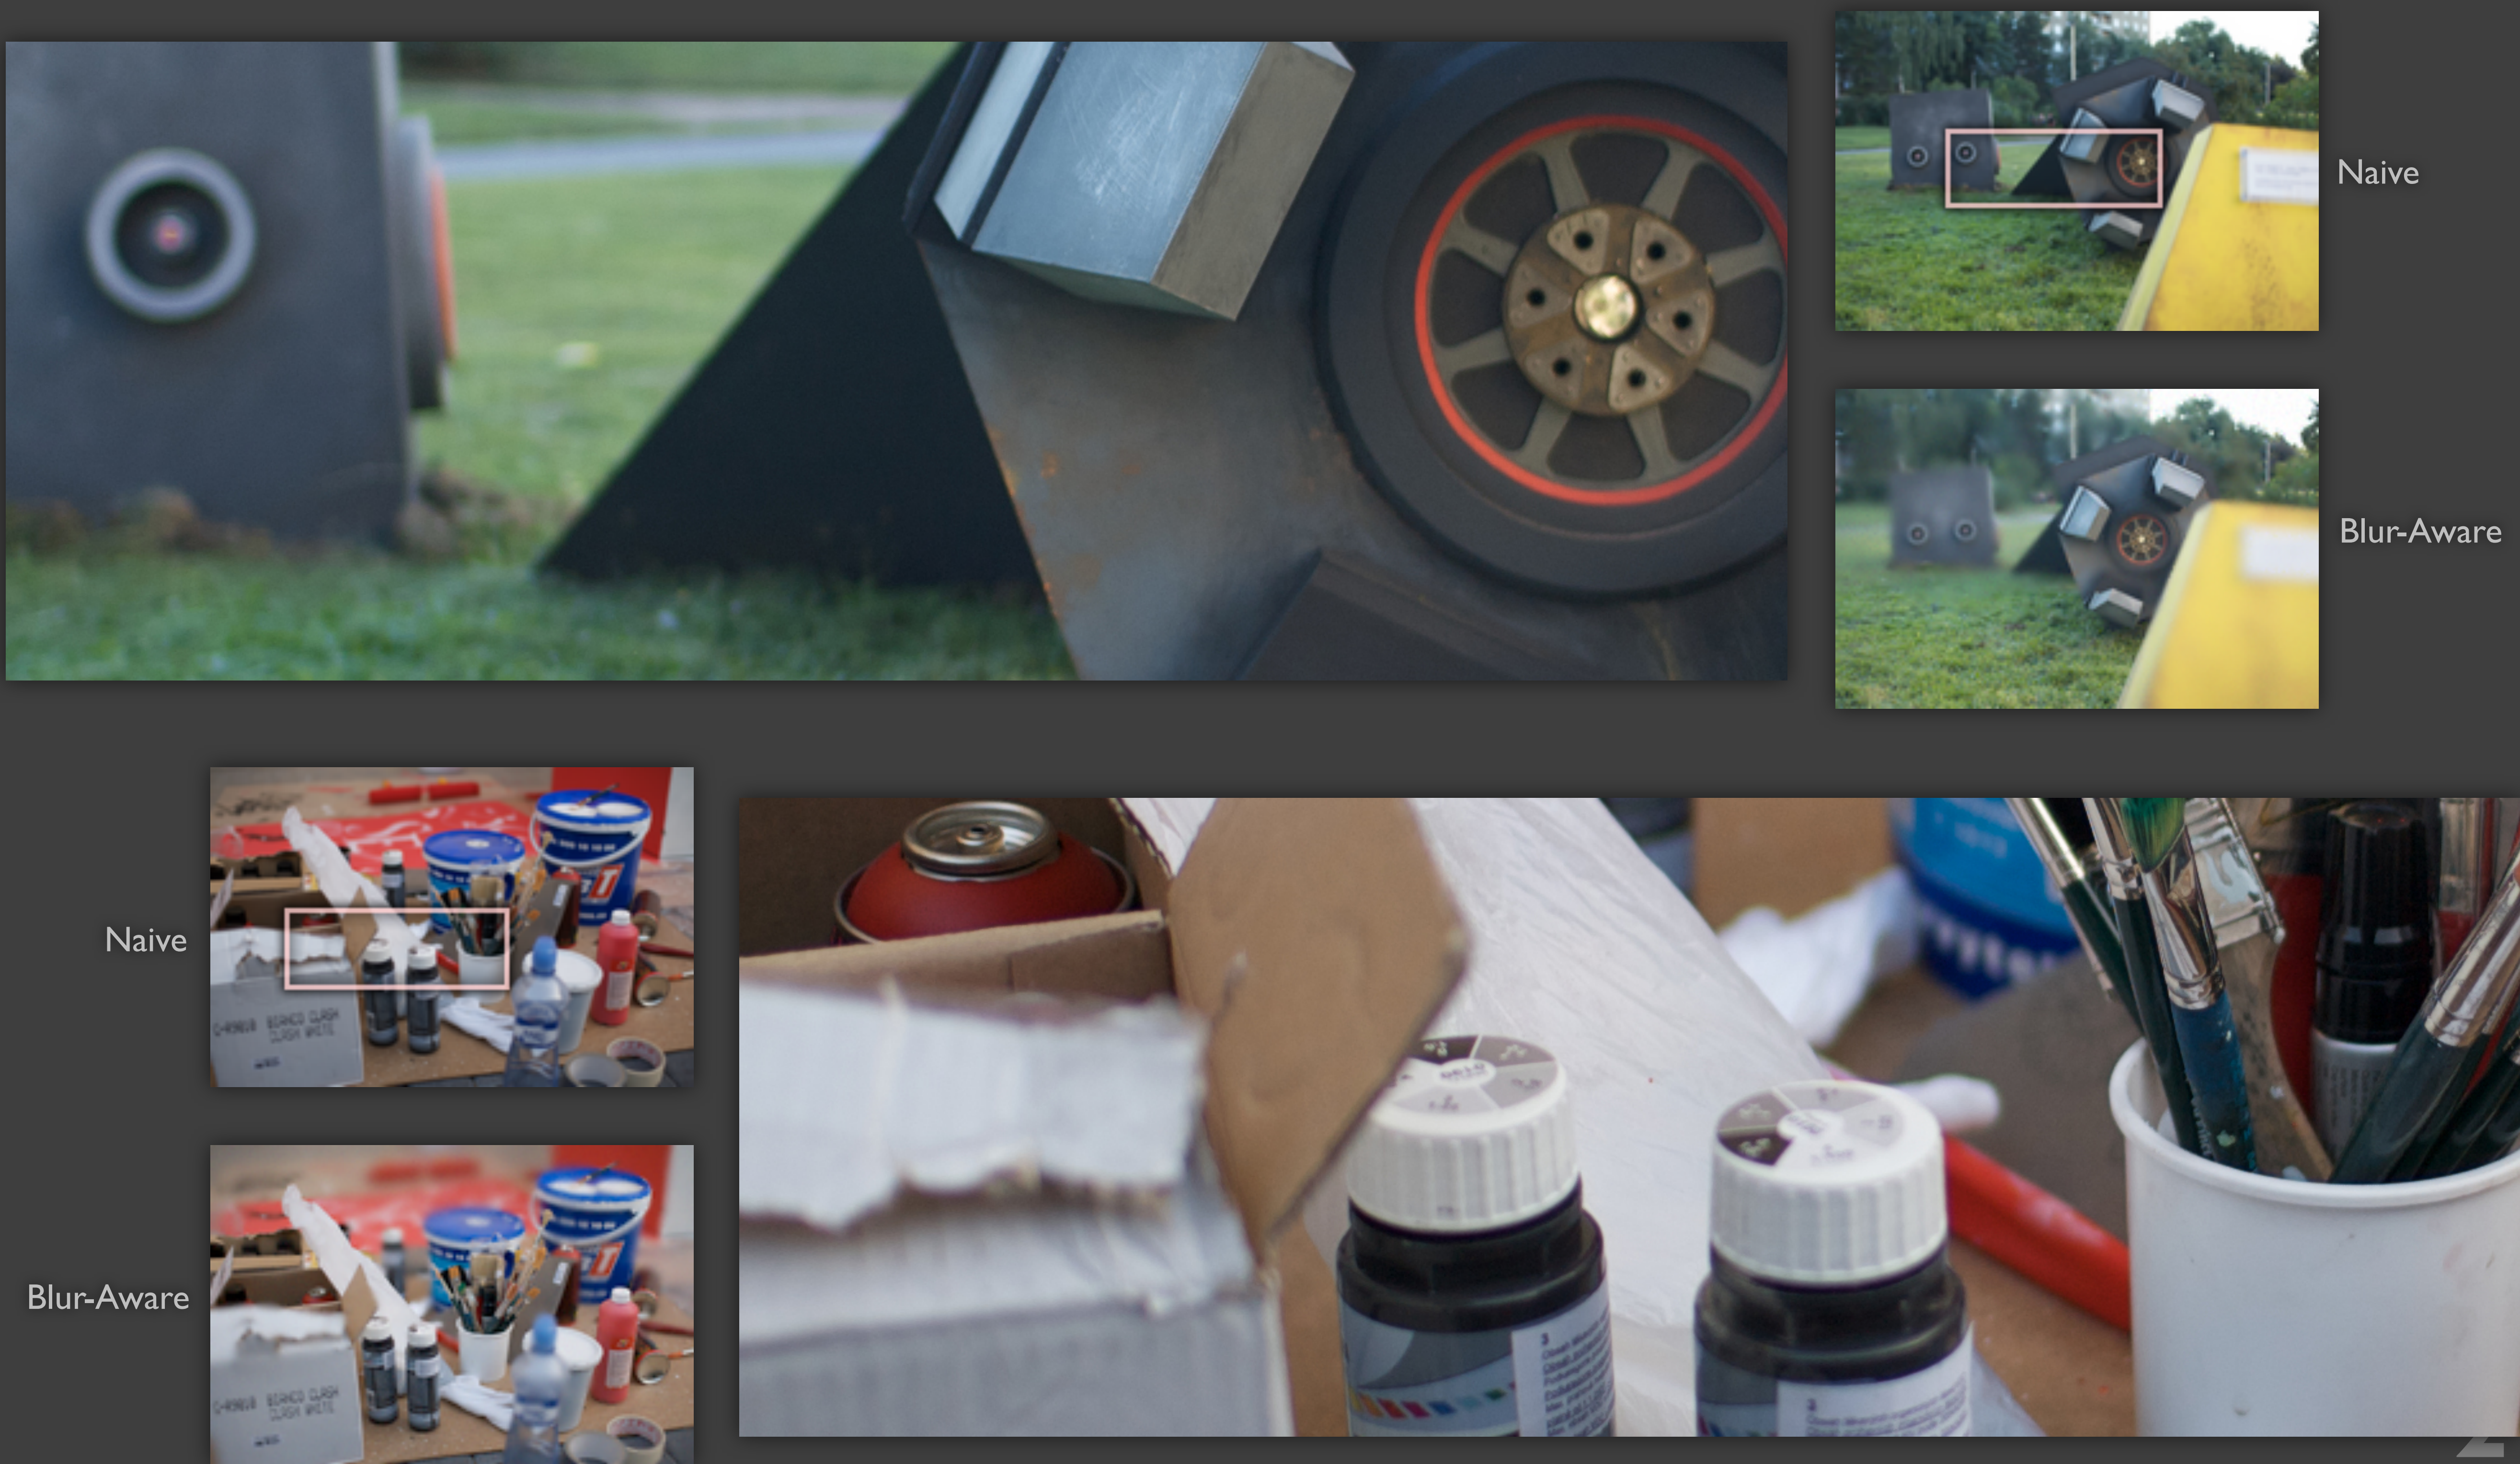

- Two more examples
- Hopefully it should be apparent that the hand of the robot is in focus, while the head is not
- Same with the art supplies in the background
- Our method preserves this while the normal thumbnail appears sharp

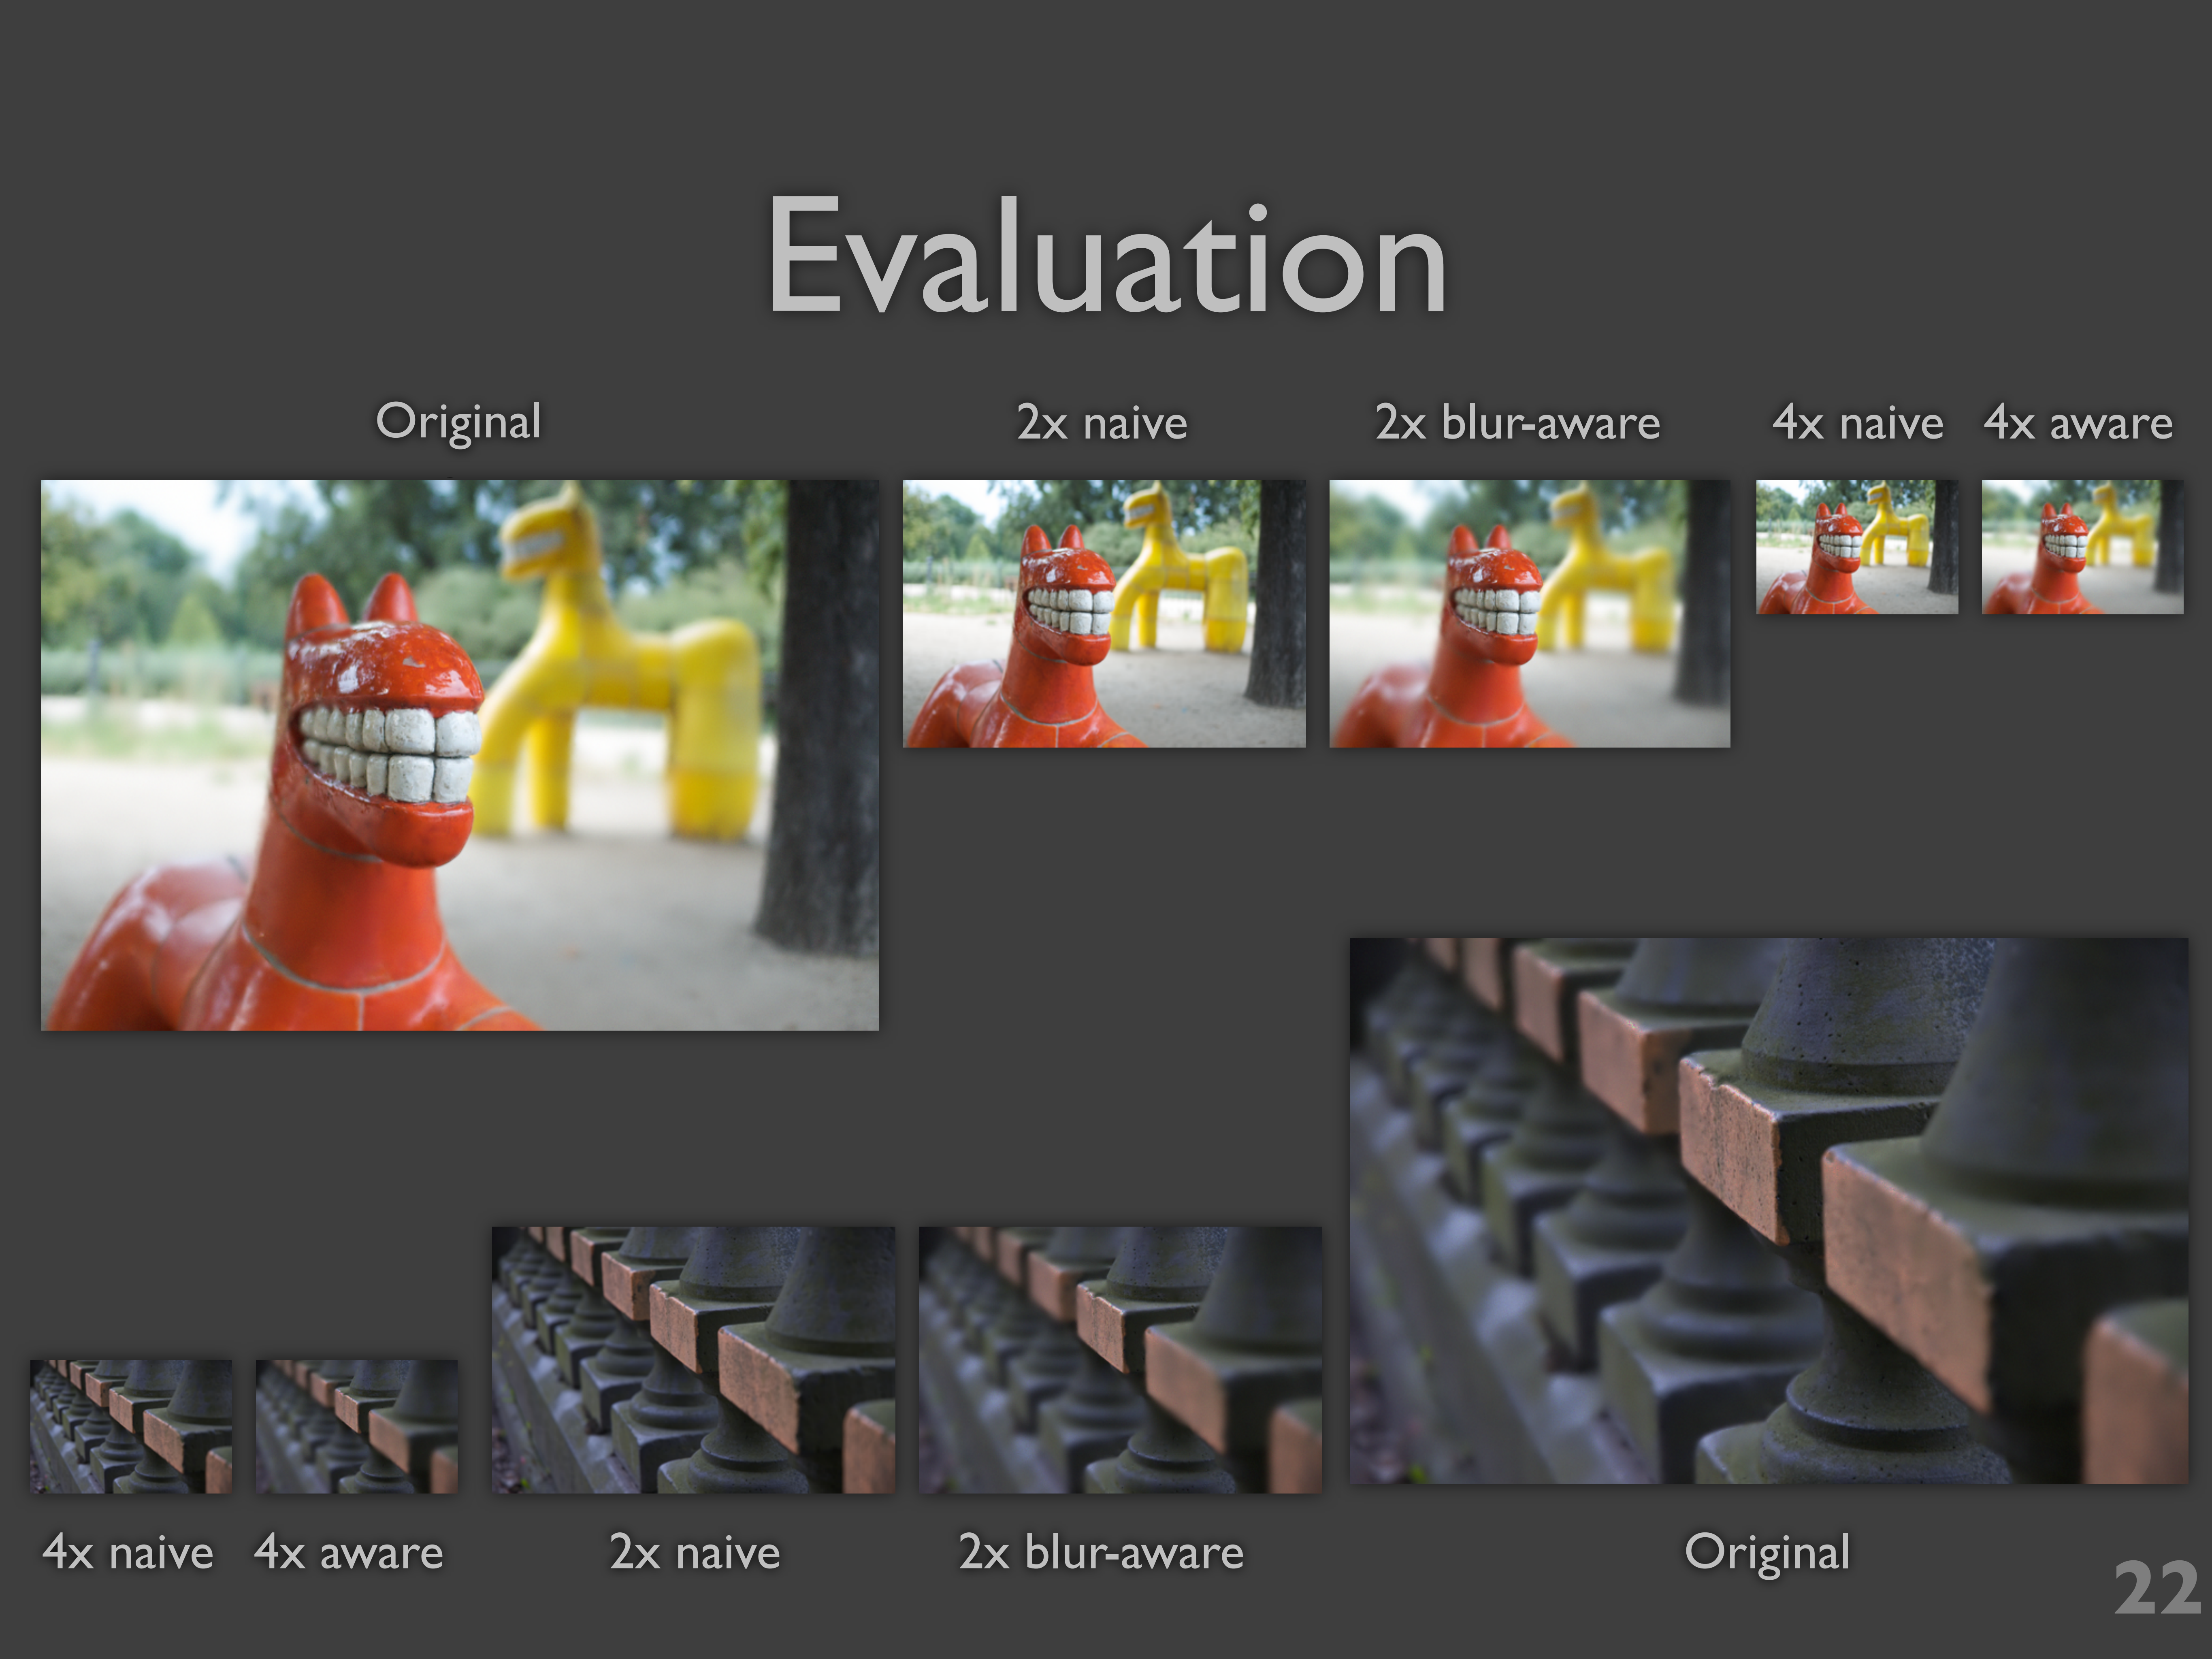

– Reduction in the depth of field in the conventional thumbnails of banister

Conclusion

- Fully automatic image resizing operator that uses a perceptual metric to preserve image appearance
- Effect due to HVS:
  The same metric can account for changes in appearance due to viewing distance
- Future work:
  Other models like camera optics to enhance blur
  Extending principle to other attributes such as noise or contrast

- Relationship between the viewer and the display matters
- Move towards a model of image display that accounts for this relationship
- Either factorizing these distance-dependent effects for lightfield displays
- Or having displays that sense the viewer and display the appropriate content



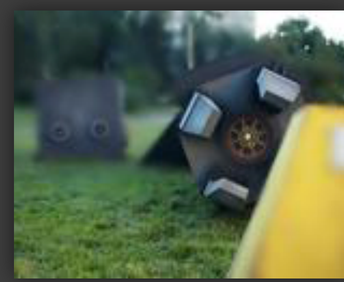

Thanks!

(you and our sponsors)

DOLBY

Research Chair