

Blur-Aware Image Downsampling

Matthew Trentacoste¹ & Rafal Mantiuk² & Wolfgang Heidrich¹

¹University of British Columbia, Canada

²Bangor University, United Kingdom

Abstract

Resizing to a lower resolution can alter the appearance of an image. In particular, downsampling an image causes blurred regions to appear sharper. It is useful at times to create a downsampled version of the image that gives the same impression as the original, such as for digital camera viewfinders. To understand the effect of blur on image appearance at different image sizes, we conduct a perceptual study examining how much blur must be present in a downsampled image to be perceived the same as the original. We find a complex, but mostly image-independent relationship between matching blur levels in images at different resolutions. The relationship can be explained by a model of the blur magnitude analyzed as a function of spatial frequency. We incorporate this model in a new appearance-preserving downsampling algorithm, which alters blur magnitude locally to create a smaller image that gives the best reproduction of the original image appearance.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Line and curve generation

1. Introduction

One pervasive trend in imaging hardware is the ever-increasing pixel count of image sensors. Today even inexpensive cameras far outperform common display technologies in terms of image resolution. For example, very few cameras remaining on the market, including mobile devices, capture an image at low enough resolution to show on a 1080p HDTV display without resizing. In an extreme case, the preview screen on one Nikon professional digital SLR can only display 1.5% the pixels captured by the sensor. The cell-phone camera owner and the 4K cinematographer face the same problem of getting an accurate depiction of the image when they can't see all the pixels.

While high resolution images are needed for a number of applications such as on-camera previewing, print output or cropping, the image is often previewed on a display of lower resolution. As a result, image downsampling has become a regular operation when viewing images. Conventional image downsampling methods do not accurately represent the appearance of the original image, and lowering the resolution of an image alters the perceived appearance. In particular, downsampling can cause blurred regions to look sharp and the resulting image often appears higher quality than its

full-size counterpart. While the higher quality images can be desirable for purposes such as web publishing, the change is problematic in cases where the downsampled version is to be used to make decisions about the quality of the full-scale image, for example in digital view finders.

In this paper, we aim to develop an image downsampling operator that preserves the appearance of blurriness in the lower resolution image. This is a potentially complex task — the human visual system's ability to differentiate blurs is dependent on spatial frequency, and edges blurred by different amounts may be perceived as different at one scale but equal at another. Additionally, there is potential for content-dependent blur perception where the same amount of blur is perceived differently, depending on the type(s) of object(s) shown.

We approach this problem by conducting a perceptual study to understand the relationship between the amount of blur present in an image, and the perception of blur at different image sizes. Our study determines how much blur must be present in a downsampled image to have the same appearance as the original. We find a complex and mostly image-independent relationship between matching blur levels in images at different resolutions. The relationship can

be explained by a linear model when the blur magnitude is analyzed in terms of spatial frequency.

Using the results of this study, we develop a new image resizing operator that amplifies the blur present in the image while downsampling to ensure it is perceived the same as the original. While our algorithm is compatible with any combination of methods for producing a spatially-variant estimate of image blur and spatially-variant image filtering, we base our implementation around a modified version of the algorithm by Samadani *et al.* [SLT07]. The result is a fully-automatic method for downsampling images while preserving their appearance, the performance of which we verify with another user study.

2. Related Work

The manner in which the human visual system perceives blur is a complex subject and the topic of a number of perceptual studies. Cufflin *et al.* [CMM07] and Chen *et al.* [CCT*09] attempt to quantify blur discrimination, the ability to perceive whether two blurs are different, while Mather and Smith [MS02] investigated how blur discrimination affects depth perception. Held *et al.* [HCOB10] studied how the pattern of blur in an image influenced users' perception of the absolute distance and scale of objects in the scene.

In the context of image resizing, Fattal [Fat07], Kopf *et al.* [KCLU07] and Shan *et al.* [SLJT08] developed techniques for intelligently *upsampling* images. These methods use assumptions about the image statistics to invent additional information based on existing details to provide a more natural rendition than a reconstruction-filter resampling. Also related are seam carving methods such as the one by Avidan and Shamir [AS07]. These methods resize images by inserting or removing pixels in the least important regions of the image, preserving the overall structure. A number of additional methods have been presented, including extensions to video [PKVP09, KLHG09]. However, seam-carving mostly focuses on adjusting aspect ratios and is combined [RSA09] with regular downsampling operations for extreme changes in resolution. Rubinstein *et al.* [RGSS10] provides a comprehensive overview of existing methods, and conducts both perceptual and objective analyses of each, comparing their performance.

Estimating the amount of blur present in images is a well-studied but still not completely solved problem. Blind deconvolution methods, such as those by Lam and Goodman [LG00] and Fergus *et al.* [FSH*06] iteratively estimate the shape of the blur while sharpening the image. While these methods assume that the blur kernel is invariant across the entire image, other methods recover spatially-variant blur. Examples include classification of defocus and motion blur by Liu *et al.* [LLJ08] and scale-space methods such as Elder and Zucker [EZ98].

Most pertinent to our method is the work on increasing the

blur in images. Bae and Durand [BD07] produce a spatially-variant estimate of the amount of blur present at each location in the image and increase the blur to simulate wider apertures. However, the method is computationally very expensive and thus not suitable for many applications, such as a digital viewfinder. Samadani *et al.* [SLT07] developed a method of amplifying specific artifacts present in full-size images to be visible in thumbnail images, including estimating and amplifying blur. In both cases, the amount of blur is increased by single scale factor, specified by the user. The perception of blur is more complex than this relationship and neither method can ensure that the appearance of blur will remain constant if the image is resized.

3. Experiment Design

The basic premise of our paper is that the blur in an image is perceived differently when that image is downsampled. In order to create a downsampled image that preserves the appearance of the original image, we must quantify that change in perception. The experiment was intended to measure the amount of blur that needs to be present in a thumbnail image in order to match the appearance of blur in a full-size version of the same image. This relationship was measured in a blur-matching experiment.

Observers were presented a full-size image, as well as two thumbnail images. They were asked to adjust the amount of blur in both thumbnail images such that the first thumbnail was just noticeably blurrier than the full-size image, and the second thumbnail was just noticeably sharper. Example stimulus is shown in Figure 1. We found that this 'bracketing' procedure resulted in more accurate measurements than direct matching and was necessary due to the relatively wide range of blur parameters that result in approximately equal appearance. Such variation of the method of adjustment was used before to measure a just noticeable blur in the context of the depth of focus of the eye [YIC10] and the brightness of the glare illusion [YIMS08]. The matching blur amount was computed as the mean of the 'less-' and 'more-blurry' measurements.

An alternative experiment design that would produce more accurate results, is the 2-alternative-forced-choice procedure, in which the observers are asked to select a blurrier/sharper image when presented the original and downsampled version and the amount of blur is randomly added or removed from the smaller image. Such procedure, although more accurate, consumes much more time (is on average 5–20 longer) and thus is not effective with a larger group of observers. The objective of our study was to gather data for an 'average' observer, thus it was more important to collect a larger number of measurements for a larger population, rather than fewer but more accurate measurements.

Viewing conditions. The images were presented on a 27" Dell 2707WFPc display with 1920×1200 resolution. The

experiment was run in a dimmed room with no visible display glare. The viewing distance was 1 m, resulting in a pixel Nyquist frequency of 30 cycles per visual degree.

Image selection. A pilot study was run to observe how the blur estimates differ between images, and in order to identify a possibly small set of images that would still reflect image-dependent effects. For the pilot experiment we selected 20 images containing people, faces, animals, man-made objects, indoor and outdoor scenes. The pilot experiment was run with 10 blur-levels and only 7 observers. The results were averaged for each test condition (blur level \times downsampling level) to form a vector value. Then, the Euclidean distance was computed between vectors for each pair of images, to build a difference matrix. The difference data was then projected onto a 2D space using multi-dimensional scaling [KW78] in order to produce the plot in Figure 2. The plot reflects image-dependent differences in blur perception for the same blur parameters. To maximize diversity in image content, we selected five images which were located far apart on the plot and thus were likely to be the most different in terms of produced results.

Stimuli. For both the pilot study and the full experiment, differently blurred versions of images were generated by introducing synthetic blur to full-size images with no noticeable blur in them. Since we could not control where in the image users were looking to make their judgements, we introduced uniform blur to completely in-focus images to avoid any ambiguities in response. For this purpose we used a Gaussian kernel of a specified standard deviation ζ . Thumbnail images were produced by the same process, except that the convolution of a full-size image was followed by nearest-neighbor resampling. We chose a nearest neighbor filter for this step in order to not distort the experiments by introducing additional low-pass filtering. However, as a result, some amount of aliasing was present for small blur kernels under large downsampling factors (also see Section 4). The reported ζ -values are given in visual degrees to make them display-independent (in this paper, we use ζ



Figure 1: Screen capture of the stimuli used in the experiment. Subjects adjust the blur in the small images on the right to match the blur in the large image on the left.

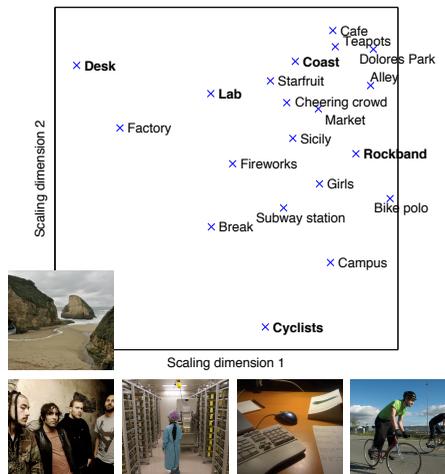


Figure 2: The result of multi-dimensional scaling on the differences between per-image results collected in a pilot study. The pilot study was intended to identify the representative images that had the highest potential to reveal any image-dependencies in the study. The images selected for use in the main study are shown at the bottom.

to denote standard deviations of blur kernels expressed in visual degrees, and σ for blur kernels expressed in pixels). The five selected images were shown at 10 blur levels, ranging from 0 to 0.26 visual degrees, and at three downsampling factors: 2, 4, and 8.

Observers. 24 observers (14 male and 10 female) participated in the study. They were paid and unaware of the purpose of the experiment. The observer age varied from 21 to 38 with the average 28. All observers had normal or corrected-to-normal vision.

Experimental procedure. Given a reference image with blur ζ_r , the observers were asked to adjust the matching blur to be just noticeably stronger in one and just noticeably weaker in the other thumbnail image. Each observer repeated the measurement for each condition three times, but each observer was assigned a random subset of 30 out of 150 conditions to reduce workload (150 = 3 downsampling factors \times 10 blur levels \times 5 images). In total over 2,100 measurements were collected. The experiment was preceded with a training session during which no data was recorded, followed by three main sessions with voluntary breaks between them. The breaks were scheduled so that each session lasted less than 30 minutes.

4. Experiment Results

The results of the experiment, averaged over the five selected images and for each image separately, are shown in Figure 3.

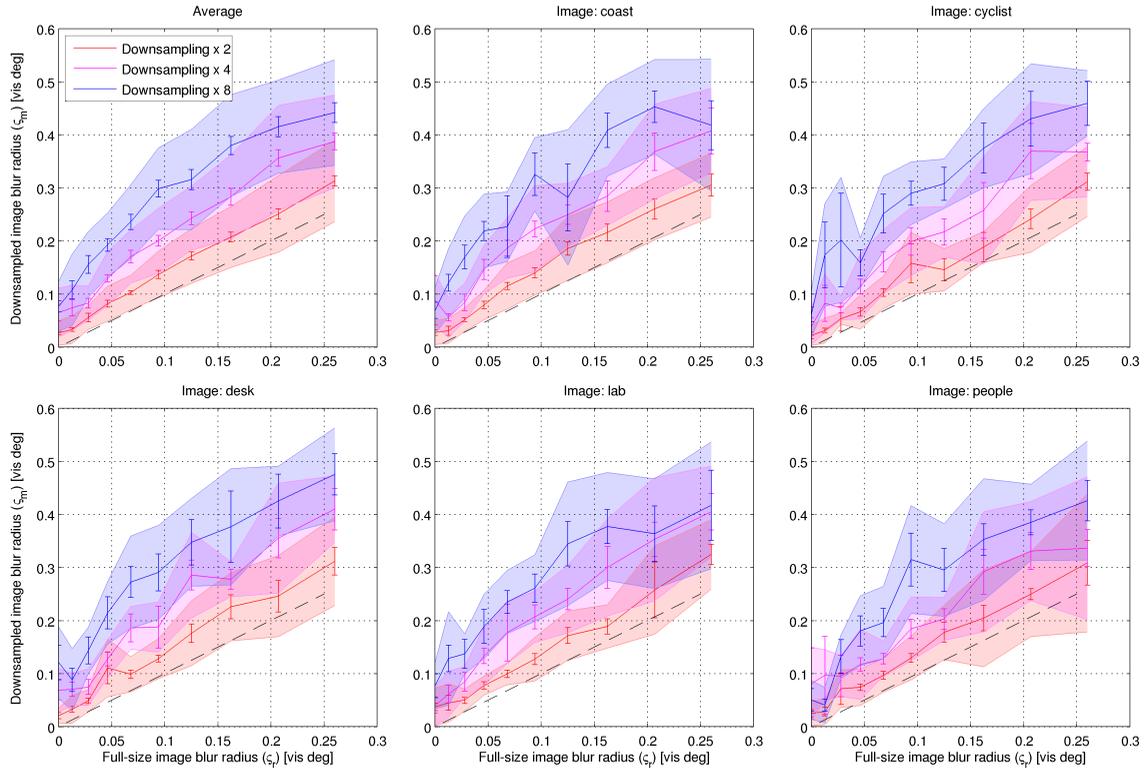


Figure 3: The results of the blur matching experiments, plotted separately for averaged data ζ_m (top-left) and for each individual image. The continuous lines are the expected magnitudes of matching blurs found by computing the average between two measurements for ‘more blurry’ and ‘less blurry’. The error bars represent a 95% confidence interval. The edges of the shaded region correspond to the mean measurement for ‘more blurry’ and ‘less blurry’.

The results are very consistent regardless of the image content, but averaging across all images is necessary to reduce variability in the data. Because both ζ values are reported with respect to the blur in the full-size image, the $y = x$ line (dashed black line in the plot) is equivalent to retaining the blur of the original image. For all data points the matching blur is larger than the blur in the original images (points above the dashed black line). This is because images look sharper after downsampling and they need to be additionally blurred to match the appearance of full-size images.

The experimental curves also level off at higher downsampling levels and for larger blur amounts. This effect is easy to explain after inspecting actual images, in which the amount of blur is so large that it sufficiently conveys the appearance of the full-size image and no additional blurring is needed.

It is important to note that the reported values also include the blurring necessary to remove aliasing artifacts. As mentioned in the previous section, we used a simple nearest-neighbor filter to resample the blurred high-resolution images so that the results are not confounded with an anti-aliasing filter. If the blur was not sufficient to prevent aliasing

in the downsampled image, the result appeared sharper than the original. We observed that when no blur was present in the large image, subjects adjusted the amount of blur in the thumbnail to a value close to the optimal low pass filter for the given downsampling factor.

5. Model for Matching Blur Appearance

In this section we introduce a model that can predict our experimental results. The plot curves in Figure 3 suggest a non-linear relation for matching blur in original and downsampled images. However, we show that the averaged measured data ζ_m is well explained by the combination of an anti-aliasing filter ζ_d and a model \mathcal{S} , which is linear in *spatial* frequencies (measured in cycles per visual degree):

$$\hat{\zeta}_m = \sqrt{\zeta_d^2 + \mathcal{S}^2}. \quad (1)$$

The $\hat{\zeta}_m$ is the model prediction of the experimental blur-matching data for an average observer. The term ζ_d approximates the effect of an ideal anti-aliasing filter. The standard deviation ζ_d of the Gaussian filter that provides a least-

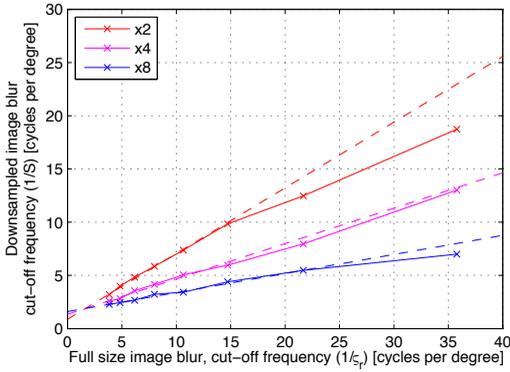


Figure 4: Average matching blur data (Figure 3 first panel), with the anti-aliasing component ζ_d removed, replotted as the roll-off frequency. The matching blur follows straight lines, except for the small blur amounts (high frequency roll-off), where aliasing dominates. The two lowest value ζ_r points were omitted from the plot as the values were excessively large due to the $1/\zeta$ transform.

squares fit of the sinc filter is

$$\zeta_d = d \frac{\sqrt{3 \log 2}}{\pi \cdot p}, \quad (2)$$

where d is the downsampling factor, while p is a conversion factor that maps from image units (pixels) to visual degrees, which is equal to the number of pixels per visual degree. In our experiments, we had subjects sit further away from the screen than usual, to prevent limitations in screen resolution from affecting the results. As a result, $p = 60$ for our experimental data.

To motivate our choice of model, we remove the anti-aliasing component ζ_d from the experimental data and plot it in terms of spatial frequency $1/\zeta$ in Figure 4. The plot shows the experimental data expressed as the \mathcal{S} component of Equation 1. All data points are now well aligned and mostly in linear relation, except several measurements at high frequencies and for the $2\times$ downsampling factor. We attribute these inaccuracies to the measurement error, which is magnified in this plot because of the $f = 1/\zeta$ transform. The plot demonstrates that the remaining term \mathcal{S} can be modeled as a set of straight lines when expressed in terms of spatial frequencies. Moreover, the lines cross at approximately the same point. The model that provides the best least-squares fit of the experimental data in terms of ζ -values is

$$\mathcal{S}(\zeta_r, d) = \frac{1}{2^{-0.893 \log_2(d) + 0.197} \left(\frac{1}{\zeta_r} - 1.64 \right) + 1.89}, \quad (3)$$

where d is the downsample amount and ζ_r is the amount of the reference blur in the original image.

Figure 5 plots the combined blur model $\widehat{\zeta}_m$ as compared to the results from our experiments. The figure shows that

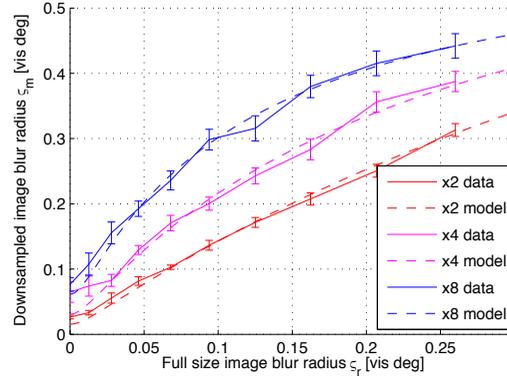


Figure 5: Blur model $\widehat{\zeta}_m$ (dashed lines) compared with the experiment results ζ_m (continuous lines and error bars).

the fitting error is quite acceptable, even for low- ζ (high-frequency) points, which did not follow the linear relation in Figure 4. When comparing plots, note that the large frequency values correspond to small ζ -values. While a higher-order function could provide a better fit, our experimental data does not provide enough evidence to justify such a step. Moreover, we believe that a linear model in terms of spatial frequency is more plausible than a higher-order function. In the supplementary materials we include a number of examples in which the model is tested on images that were not used in the experiment.

Note that the combined model of matching blur $\widehat{\zeta}_m$ is the absolute amount of blur that needs to be present in the full-size image before downsampling and is expressed in units of visual degrees. In Section 6.3 we explain how to compute the amount of blur that needs to be added to a downsampled image.

6. Resizing Algorithm

The goal of our algorithm is to use the results of our experiment to automatically produce a downsampled image that preserves the appearance of the original blur. We first compute a spatially-varying estimate of the amount of blur present in the full-size image. Given that estimate, we use the results of our study to determine how much additional blur is needed for the specified downsample. Finally, we synthesize a new downsampled image with the amount of blur required to preserve the appearance of the image.

Our overall approach can work with any method that provides a spatially variant estimate of image blur. We considered the method by Elder and Zucker [EZ98], but it only produces estimates at edge locations and requires the work of Bae and Durand [BD07] to provide a robust estimate of the blur at all pixels. While the approach produces high-quality results, it operates at the resolution of the original image and is computationally intensive.

Instead, we chose to base our method on the algorithm of Samadani *et al.* [SMBC10] because of its simplicity and computational efficiency, which is a result of performing most work at the resolution of the downsampled image. We do not use the final resulting thumbnail of Samadani *et al.*, instead making use of the spatially-variant blur estimate they compute as an intermediate result. In Section 6.2 we extend that method so that it provides blur estimates at each pixel location in terms of the Gaussian kernel σ before synthesizing the final image using our model of perceived blur.

The Gaussian model differs from the geometric model for defocus blur. However, it has been argued that Gaussian blur better accounts for artifacts in actual cameras, and it has been used widely in computer vision [Pen87, Sub92]. Note that the following considerations assume σ values expressed in pixels rather than visual degrees.

6.1. Base Algorithm

To preface our work, we summarize the blur estimation method of Samadani *et al.*, which produces a spatially-varying map at the resolution of the thumbnail. The algorithm first generates a standard thumbnail, t_s , and produces a scale-space [Lin94b] of thumbnails blurred by different amounts. Image features are computed for the high resolution original as well as for each of the thumbnail images in the scale-space. The amount of blur is determined by the level of the scale-space with feature values most similar to those of the original image features.

These features are computed as the maximum absolute difference between a pixel and its eight neighbors. In the case of the original image, the feature values are downsampled using a maximum filter to produce a thumbnail resolution *range map*, r_o . The levels of the thumbnail scale-space l_{σ_j} are created by convolving the standard thumbnail with a set of Gaussian kernels of standard deviations σ_j , where l_{σ_0} is the unblurred, original thumbnail. For each of these images l_{σ_j} , a low resolution range map r_{σ_j} is generated.

The estimate of the blur present in an image is represented by the blur map m . Each pixel i of the blur map is determined as

$$m(i) = \min_j \{j \mid r_{\sigma_j}(i) \leq \gamma r_o(i)\}, \quad (4)$$

where γ is a user-specified parameter that controls which r_{σ_j} most closely matches r_o and in turn adjusts the amount of blur added. An example of the blur map is shown in Figure 6. The final image is synthesized by selecting the pixel from l_{σ_j} that corresponds to $m(i)$.

While Samadani *et al.* recover a blur map for the image, it is worth noting that this blur map does not necessarily represent the defocus or depth of pixels. It can better be understood as a map of relative gradient magnitude per image region. While rapidly changing derivatives correspond to sharp

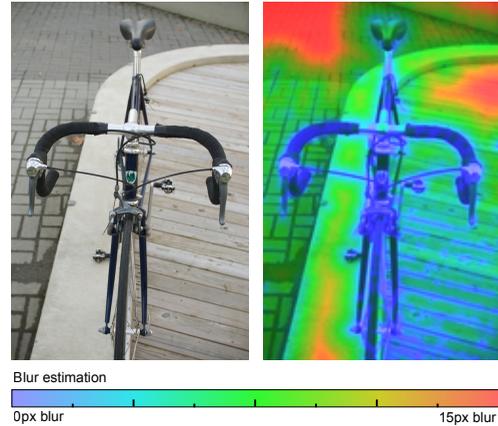


Figure 6: Input image (left) and the associated blur map (right). Note that while the curb is the same distance as the wooden boards, it is estimated as blurrier due to the lack of detail in that area. The blur map is visible only in color.

regions, there is an ambiguity with slowly changing derivatives. If an image region is a flat color, we cannot determine whether that is a detailed region that is out of focus or it is in focus but lacks any detail. Both situations are equivalent from the viewpoint of this algorithm.

6.2. Blur Estimation

While Samadani *et al.*'s blur estimation provides a means of controlling the relative increase in the amount of blur in the resulting thumbnail, it does not provide an absolute measure of the blur present in the large image. In order to make use of the results of the model from Section 5, we need to know the scale of the blur present in the original image. To do so, we extend their local image features to a general relationship between the width of a blurred edge and the corresponding derivatives at different resolutions, which we can use to recover the scale of the original image blur.

In the case of a 1D Gaussian blurred edge of normalized contrast, the edge profile is the integral of the Gaussian function. The derivatives of this profile follow the Gaussian function, with the peak lying at the center point of the edge. For a Gaussian profile with standard deviation σ to have a contrast of 1, the derivative of the edge cross-section will be:

$$g(\sigma, x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{x^2}{2\sigma^2}}. \quad (5)$$

This scaling factor establishes a relationship between the width of the Gaussian profile and the scale of its derivatives. If the width of the edge profile changes by a factor k , the derivatives must change by a factor of $1/k$ to retain the same contrast.

The range map operator in Samadani *et al.* approximates

the gradient magnitude. For an edge with blur σ_o , the corresponding range map will equal a Gaussian distribution at the edge location with an amplitude of $1/\sqrt{2\pi\sigma_o^2}$. After downsampling that image to obtain r_o , the value at the edge location is still equal to $1/\sqrt{2\pi\sigma_o^2}$.

Due to downsampling, the effective width of that edge in the thumbnail l_{σ_o} will differ by the downsample factor d . The pixel corresponding to the edge location in the thumbnail range map r_{σ_j} will be

$$\frac{1}{\sqrt{2\pi\left(\frac{\sigma_o}{d}\right)^2}} \quad (6)$$

due to the relationship between the width and scale of a Gaussian mentioned above. Additionally, that σ_o/d will be further altered by the Gaussian filtering that generates the scale-space images l_{σ_j} . Using the convolution formula for Gaussian functions:

$$g(n_1, \sigma_1) \otimes g(n_2, \sigma_2) = g\left(n_1 + n_2, \sqrt{\sigma_1^2 + \sigma_2^2}\right), \quad (7)$$

the width of the edge in thumbnail scale-space image l_{σ_j} will thus be

$$\sqrt{\left(\frac{\sigma_o}{d}\right)^2 + \sigma_j^2}. \quad (8)$$

We construct the scale-space from a series of blurs with a uniform spacing of β , which implies $\sigma_j = \beta j$. This way, the choice of β along with the maximum j determines the range and quantization of the scale-space. The end result is two different values for corresponding pixels of the two range maps:

$$r_o = \frac{1}{\sqrt{2\pi\sigma_o^2}} \text{ vs. } r_{\sigma_j} = \frac{1}{\sqrt{2\pi\left[\left(\frac{\sigma_o}{d}\right)^2 + (\beta j)^2\right]}}. \quad (9)$$

Figure 7 depicts the relationship between r_o and the scale-space r_{σ_j} for an edge of increasing blur.

For the algorithm to select the correct value for the blur map $m(i)$, the two values of Equation 9 must be equal. In complex images, adjacent image features alter the derivative values at this edge location and a direct solution would misestimate σ_o . Our approach determines σ_o based on the correspondence between the range map and the levels of the scale-space. Adjacent features alter the gradient magnitude in both range maps in the same fashion, and the correspondence between them is preserved. Additionally, image features smaller than $\sqrt{\sigma}$ are suppressed on the scale-space level with a Gaussian blur of σ , eliminating some of overlapping features [Lin94a].

To determine the value of the blur map $m(i)$, Samadani *et al.* employ the user-specified parameter γ to bias the selection of values for $m(i)$ towards more or less blurred levels of r_{σ_j} . Noting that the relation between the two range maps depends on the downsample amount, we instead determine the

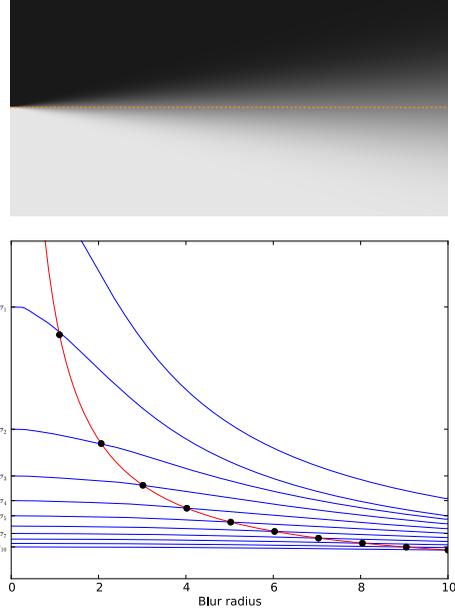


Figure 7: Top: image of a step edge with blur increasing from 0...10 along the x-axis. Bottom: range map values along dotted line for the original image r_o (red), and the downsampled scale-space r_{σ_j} (blue). Note that the intersection between r_o and r_{σ_j} (black dots) happens at $x = j$.

value of γ that correctly scales the range map of the original image and the range maps of the downsampled scale-space. We solve the following equation between r_o and r_{σ_j}

$$\gamma \frac{1}{\sqrt{2\pi\sigma_o^2}} = \frac{1}{\sqrt{2\pi\left[\left(\frac{\sigma_o}{d}\right)^2 + (\beta j)^2\right]}} \quad (10)$$

for the value of γ that ensures the index of the blur map selected with Equation 4 is equal to the width of the blur in the original large image, $m(i) = j$ if $\sigma_o = j$. Canceling terms and solving for γ yields:

$$\gamma = \frac{1}{\sqrt{\left(\frac{1}{d}\right)^2 + \beta^2}}. \quad (11)$$

The result is a value of γ automatically chosen for a given downsample factor and scale-space resolution. In our method, we use 25 levels of blur ($j = 1, \dots, 25$) and a value of $\beta = 0.4$. The resulting values for downsamples of 2, 4, and 10 are $\gamma_{d=2} = 1.55$, $\gamma_{d=4} = 2.11$ and $\gamma_{d=10} = 2.48$.

6.3. Perceptually Accurate Blur Synthesis

With an accurate estimate of the blur present at each pixel of the large image, we use our model from Section 5 to compute the amount of blur desired in the downsampled image. To produce the appearance-matching image, we reduce its



Figure 8: Comparison of conventional downsampling and our method for two images. The bottom row contains cropped portions of the images at the original resolution (see pink boxes in conventional thumbnails). Note the blur present in the eye of the robot sculpture and cardboard box is visible in our result, but appears sharp in the conventional thumbnail.

resolution by downsampling it by a factor of d using the standard technique with an antialiasing filter. Because the anti-aliasing is now accounted for, we use the aliasing-free component of the model $\mathcal{S}(\zeta_r, d)$ from Equation 3, rather than the complete model $\hat{\zeta}_m$. Given the existing blur in the full-size image σ_o , the amount of blur that needs to be added to a *downsampled* image is expressed as

$$\sigma_a = \sqrt{\left(\frac{\mathcal{S}(\sigma_o \cdot p^{-1}, d) \cdot p}{d}\right)^2 - \sigma_o^2}, \quad (12)$$

The downsampling factor d reduces the blur amount as we work on a lower-resolution downsampled image. The conversion factor p , which is equal to the number of pixels per one visual degree, converts visual degrees used in the model to pixels used in the blur estimation. For a computer monitor seen from a typical distance, p is approximately 30 pix/deg .

To produce the final image, for each level of the scale-space σ_j we blur the downsampled image by the corresponding amount of additional σ_a , then linearly blend sequential pairs of those blurred images together to approximate non-integer values of σ_a . While more accurate spatially-variant blur synthesis is possible, such as [PCH10], we haven't noticed any artifacts requiring such methods.

7. Evaluation

In this section, we provide results of our method and compare our approach to that of Samadani *et al.* [SMBC10]. We encourage the reader to look at the electronic versions of the images, which represent the fine details better than prints. We also provide more examples in the video and the supplemental materials.

Figure 8 compares the results of our algorithm to those of a conventional downsampling method of low-pass filtering the image followed by nearest-neighbor sampling. In both the example of the robot sculpture and the art supplies, objects that appear in focus (such as the head of the robot or cardboard box) in the conventionally-downsampled image are in fact blurry, as can be seen in the zoomed portions. Our method accurately detects this blur and preserves the appearance in the downsampled image.

Figure 9 demonstrates the effectiveness of our algorithm at preserving the appearance of blur in images across multiple downsample factors. In this image, the original images are downsampled by a factor of 2 and 4, and the smaller versions retain the same impression of the depth of field.

Figure 10 compares the results of our method to those of the original method of Samadani *et al.* [SMBC10]. If the value of γ is manually chosen for the image, their method can approximate our own. However, if the value of γ is incorrectly chosen, their method will either introduce too much blur and remove detail from the branches in the upper left or not introduce enough blur and retain all the details in the flowers. Even with a correctly chosen value of γ their method can only linearly scale the amount of blur, and cannot model the more complex relationship between existing blur and desired blur observed in the user study.

Additionally, we conducted a second user study to verify the effectiveness of our method. Previously, Samadani *et al.* [SMBC10] performed a preference study to determine whether subjects felt their method was more representative of the original image than a conventional thumbnail. This study showed that users did prefer the method of Samadani *et al.* over standard thumbnails. We instead chose to con-

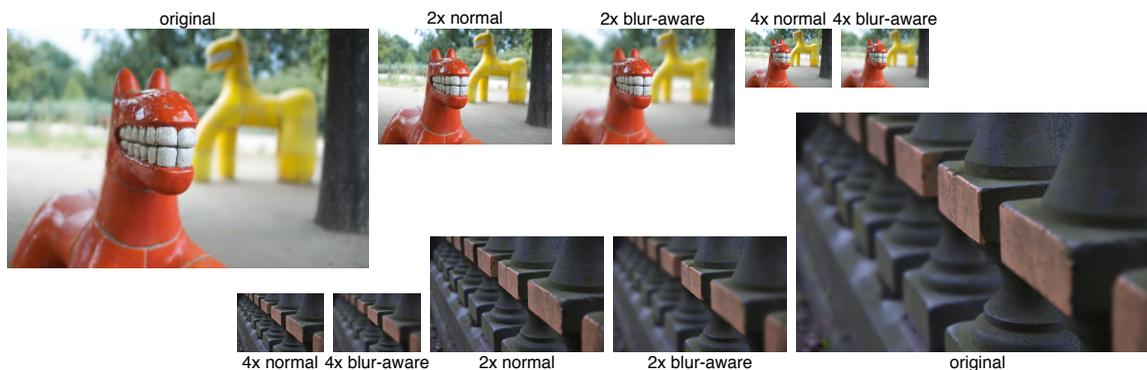


Figure 9: Comparison of appearance of blur at multiple downsampling levels. All of our results retain roughly the same amount of blur as the original while the conventionally downsampled appear to get progressively sharper.

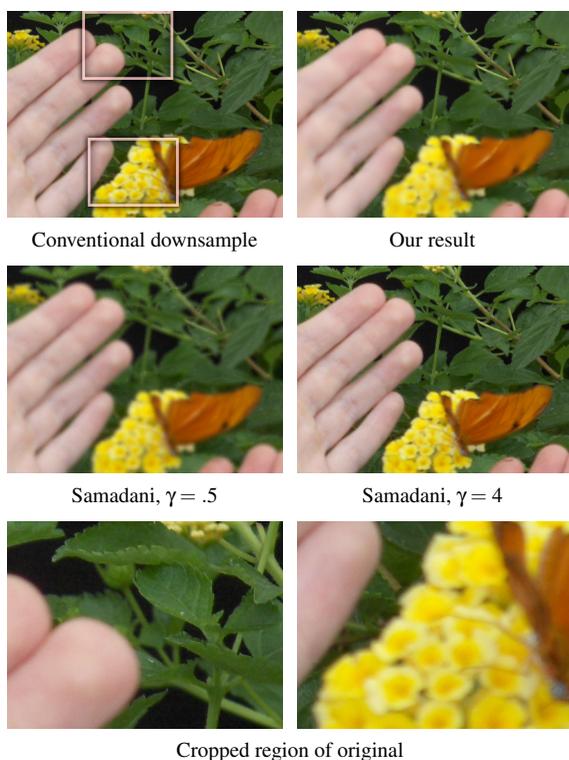


Figure 10: Comparison of naive downsampling and our method (top row) to Samadani et al. with too much blur ($\gamma = .5$) and too little blur ($\gamma = 4$) from incorrect choices of γ . The bottom row contains a cropped region of the original image (see pink boxes) for comparison.

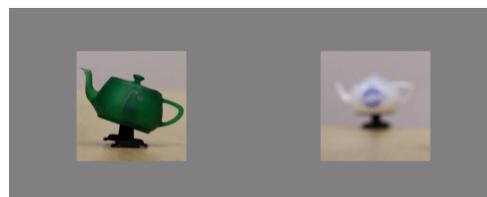


Figure 11: A screenshot of the verification study containing a pair of thumbnails with different amounts of defocus blur. Subjects had to choose which image contained the object more in focus.

duct a task-based survey to determine the extent to which our method improves users' ability to make accurate comparisons of how objects in a scene are blurred.

In this 2-alternative-forced-choice study, we photographed a series of objects with increasing amounts of defocus blur. Thumbnail versions of these images were created using both our algorithm and the conventional downsample process to downsample by a factor of 8. Subjects were shown pairs of images with different amounts of defocus blur and asked to specify in which of the thumbnails the object appeared sharper. Figure 11 contains an example stimuli. A total of 5 observers participated in the study, performing a total of 240 trials for each of the downsampling algorithms.

Overall, subjects correctly identified the sharper object 67% of the time when viewing conventional thumbnails, while they correctly identified the sharper object 83% of the time when viewing the results of our method.

Our method outperforms conventional downsampling when the blur is small enough that the object will appear sharp in a standard thumbnail and blurred in our result. However, both methods exhibit similar performance if the blur is small enough for the object to appear sharp at the orig-

inal resolution, and thus in both thumbnail versions. Likewise, the performance of the two methods will be the same if the blur is large enough that the object will appear blurred in both thumbnails. We used a uniform distribution of blur amounts, so our experiment covers all three of these cases.

8. Conclusion

In this paper, we have presented a perceptually-based model of how the perception of blur in an image changes as the size of that image is reduced. This model is based on a linear relationship between the perceived blur magnitude and the blur present in the image, when analyzed in terms of spatial frequency.

We have used that model to create a new image-resizing operator that preserves the perception of blur in images as they are downsampled, ensuring that the new image appears the same as the original. To do so, we modified an existing blur estimation algorithm by Samadani *et al.* [SMBC10] to provide estimates of the original image in absolute units.

Future work includes extending the concept of preserving the perceived appearance to other image attributes. From an image-quality perspective, accurately preserving the appearance of noise when downsampling can be as important as blur. Additionally, we would like to investigate the relationship between the perceived contrast of an image and change in blur or size and see if there is a similar relationship.

More generally, any form of downsampling involves discarding information present in the image. The convention in graphics and image processing is to attempt to produce the highest quality result, which usually involves throwing away higher frequency detail to avoid any aliasing artifacts. Due to the disparity between sensor resolution and display resolution, users often view images and make image assessments based on lower-resolution versions that might not represent their full-size counterparts. We have proposed an approach that considers how the image is perceived, and preserves that appearance rather than producing a higher-quality but less representative result.

References

- [AS07] AVIDAN S., SHAMIR A.: Seam carving for content-aware image resizing. *ACM Trans. Graph.* 26 (July 2007).
- [BD07] BAE S., DURAND F.: Defocus magnification. *Computer Graphics Forum* 26, 3 (2007), 571–579.
- [CCT*09] CHEN C.-C., CHEN K.-P., TSENG C.-H., KUO S.-T., WU K.-N.: Constructing a metrics for blur perception with blur discrimination experiments. In *Proc. SPIE, Image Quality and System Performance VI* (2009), no. 724219.
- [CMM07] CUFFLIN M. P., MANKOWSKA A., MALLIN E. A. H.: Effect of blur adaptation on blur sensitivity and discriminations in emmetropes and myopes. *Investigative Ophthalmology & Visual Science* 48, 6 (2007), 2932–2939.
- [EZ98] ELDER J., ZUCKER S.: Local scale control for edge detection and blur estimation. *IEEE PAMI* 20, 7 (1998), 699–716.
- [Fat07] FATTAL R.: Image upsampling via imposed edge statistics. *ACM Trans. Graph.* 26 (July 2007).
- [FSH*06] FERGUS R., SINGH B., HERTZMANN A., ROWEIS S. T., FREEMAN W. T.: Removing camera shake from a single photograph. *ACM Trans. Graph.* 25, 3 (2006), 787–794.
- [HCOB10] HELD R. T., COOPER E. A., O'BRIEN J. F., BANKS M. S.: Using blur to affect perceived distance and size. *ACM Trans. on Graph.* 29, 2 (2010), 19:1–16.
- [KCLU07] KOPF J., COHEN M. F., LISCHINSKI D., UYTENDAELE M.: Joint bilateral upsampling. *ACM Trans. Graph.* 26 (July 2007).
- [KLG09] KRÄHENBÜHL P., LANG M., HORNING A., GROSS M.: A system for retargeting of streaming video. In *ACM SIGGRAPH Asia 2009 papers* (New York, NY, USA, 2009), SIGGRAPH Asia '09, ACM, pp. 126:1–126:10.
- [KW78] KRUSKAL J. B., WISH M.: *Multidimensional Scaling*. Sage Publications, 1978.
- [LG00] LAM E. Y., GOODMAN J. W.: Iterative statistical approach to blind image deconvolution. *J. Opt. Soc. Am. A* 17, 7 (2000), 1177–1184.
- [Lin94a] LINDBERG T.: Scale-space theory: A basic tool for analyzing structures at different scales. *Journal of applied statistics* 21, 1 (1994), 225–270.
- [Lin94b] LINDBERG T.: *Scale-space theory in computer vision*. Kluwer Academic Publishers, 1994.
- [LLJ08] LIU R., LI Z., JIA J.: Image partial blur detection and classification. In *Proc. CVPR* (2008), pp. 1–8.
- [MS02] MATHER G., SMITH D.: Blur discrimination and its relation to blur-mediated depth perception. *Perception* 31, 10 (2002), 1211–1219.
- [PCH10] POPKIN T., CAVALLARO A., HANDS D.: Accurate and efficient method for smoothly space-variant Gaussian blurring. *IEEE Trans. Img. Proc.* 19, 5 (2010), 1362–1370.
- [Pen87] PENTLAND A. P.: A new sense for depth of field. *IEEE Trans. Pattern Anal. Mach. Intell.* 9 (July 1987), 523–531.
- [PKVP09] PRITCH Y., KAV-VENAKI E., PELEG S.: Shift-map image editing. In *ICCV'09* (Kyoto, Sept 2009), pp. 151–158.
- [RGSS10] RUBINSTEIN M., GUTIERREZ D., SORKINE O., SHAMIR A.: A comparative study of image retargeting. *ACM Trans. Graph.* 29 (December 2010), 160:1–160:10.
- [RSA09] RUBINSTEIN M., SHAMIR A., AVIDAN S.: Multi-operator media retargeting. *ACM Trans. Graph.* 28 (July 2009), 23:1–23:11.
- [SLJT08] SHAN Q., LI Z., JIA J., TANG C.-K.: Fast image/video upsampling. vol. 27, ACM, pp. 153:1–153:7.
- [SLT07] SAMADANI R., LIM S. H., TRETTER D.: Representative image thumbnails for good browsing. In *Proc. ICIP* (2007), vol. 2, pp. II–193–II–196.
- [SMBC10] SAMADANI R., MAUER T. A., BERFANGER D. M., CLARK J. H.: Image thumbnails that represent blur and noise. *IEEE Trans. Img. Proc.* 19, 2 (2010), 363–373.
- [Sub92] SUBBARAO M.: Radiometry. Jones and Bartlett Publishers, Inc., USA, 1992, ch. Parallel depth recovery by changing camera parameters, pp. 340–346.
- [YIC10] YI F., ISKANDER D. R., COLLINS M. J.: Estimation of the depth of focus from wavefront measurements. *Journal of Vision* 10, 4 (2010).
- [YIMS08] YOSHIDA A., IHRKE M., MANTIUK R., SEIDEL H.: Brightness of the glare illusion. In *Proc. of Symposium on Applied Perception in Graphics and Visualization* (2008), ACM, pp. 83–90.