Online version of T.E.Boult and G.Wolberg. “Local Image Reconstruction and Sub-Pixel Restoration Algorithms.” Computer graphics and image processing: Graphical models and Image processing (CVGIP:GMIP), Vol 55, No. 1. pp. 63-77, Jan. 1993.

Unfortunately the journal printed, in error, the “low-quality” versions of the some of the figures. (They did not print a new version and all the reprints, including the hundreds I ordered , will include those incorrect figures.) Since the paper is about quality of Reconstruction having access to the high quality images is important. Please cite the above journal version, not any “online” version you might find.

LOCAL IMAGE RECONSTRUCTION AND SUB-PIXEL RESTORATION ALGORITHMS

Terrance E. Boult
Dept. of Computer Science
Columbia University
New York, NY 10027

George Wolberg
Department of Computer Science
City College of New York
New York, NY 10031

ABSTRACT

This paper introduces a new class of reconstruction algorithms that are fundamentally different from traditional approaches. We deviate from the standard practice that treats images as point samples. In this work, image values are treated as area samples generated by nonoverlapping integrators. This is consistent with the image formation process, particularly for CCD and CID cameras. We show that superior results are obtained by formulating reconstruction as a two-stage process: image restoration followed by application of the point spread function (PSF) of the imaging sensor. By coupling the PSF to the reconstruction process, we satisfy a more intuitive fidelity measure of accuracy that is based on the physical limitations of the sensor. Efficient local techniques for image restoration are derived to invert the effects of the PSF and estimate the underlying image that passed through the sensor.

The reconstruction algorithms derived herein are local methods that compare favorably to cubic convolution, a well-known local technique, and they even rival global algorithms such as interpolating cubic splines. Evaluations are made by comparing their passband and stopband performances in the frequency domain, as well as by direct inspection of the resulting images in the spatial domain. A secondary advantage of the algorithms derived with this approach is that they satisfy an imaging-consistency property. This means that they exactly reconstruct the image for some function in the given class of functions. Their error can be shown to be at most twice that of the ''optimal'' algorithm for a wide range of optimality constraints.

1. INTRODUCTION

Digital image reconstruction refers to the process of recovering a continuous image from its samples. This problem is of fundamental importance in digital image processing, particularly in applications requiring image resampling, such as image warping, correction for geometric distortions, and image registration. Its role in these applications is to furnish a spatial continuum of image values from discrete pixels so that the input image may be resampled at any arbitrary position, even those at which no data was originally supplied. Despite the great flurry of activity in reconstruction, the subject remains open to new solutions designed to address the tradeoff between reconstruction accuracy and computational complexity. The objective of this paper is to present a new class of algorithms for image reconstruction/restoration.

Whereas reconstruction simply derives a continuous image from its samples, restoration attempts to go one step further. It assumes that the underlying image has undergone some degradation before sampling, and so it attempts to estimate the original continuous image from its corrupted samples. Restoration techniques must therefore model the degradation and invert its effects on the observed image samples. In our application, we consider a limited class of degradation models motivated by point spread functions of commonly available imaging devices.

The use of image restoration techniques permits the work presented in this paper to achieve image reconstruction in a fundamentally different way than traditional approaches. This approach, which we refer to as *imaging-consistent reconstruction*, formulates reconstruction as a two-stage process: functional image restoration followed by blurring according to the sensor model. This approach is in the spirit of the work of Huck et. al (e.g. [8]) where it is argued that sampling and image formation should be considered together. Imaging-consistent algorithms directly combine knowledge of image formation and sampling into the restoration/reconstruction process. The way that knowledge is used, however, is quite different from [8].

Imaging-consistent algorithms treat the image values as area samples, i.e. each sample is a weighted integral. The weighting function is the point spread function (PSF) of the imaging sensor. We assume a single-pixel PSF and use this to derive a functional approximation to the deblurred (restored) image. (This permits the restoration to have local support.) We then blur the functional restoration by the PSF to obtain a reconstruction. By coupling the PSF to the reconstruction process, we satisfy a more intuitive fidelity measure of accuracy that is based on the physical limitations of the

sensor. Reconstruction (resampling) is more accurate in the sense that it is the exact reconstruction for some input function which, given the sensor model, would also produce the measured image data.

The imaging-consistent reconstruction algorithms derived herein are local methods that compare favorably to cubic convolution, a well-known local technique, and they even rival global algorithms such as interpolating cubic splines. Evaluations are made by comparing their passband and stopband performances in the frequency domain, using an error criteria defined in [12], as well as by direct inspection of the resulting images in the spatial domain.

This paper is organized as follows. Section 2 reviews previous work in image reconstruction, introduces the image formation process, and motivates the use of image restoration for reconstruction. That section describes why they should be unified into a common formulation, and describes the class of allowable functions that the restoration stage can estimate. The actual imaging-consistent algorithms are presented in Section 3. Section 4 contains the analysis of the new algorithms, including example imagery. Finally, conclusions and future work are discussed in Section 5.

2. BACKGROUND AND DEFINITIONS

Image reconstruction has received much attention in the literature [1, 9, 15]. It is well-known that the sinc function is the ‘‘ideal’’ reconstruction (in the spatial domain) assuming a band-limited signal and sufficient sampling. There are, however, several properties of this filter that are not desirable. Since edges constitute high frequencies, and since the basis functions are themselves oscillatory (i.e, sine waves), fast edge transitions will be distorted into smooth intensity changes with persisting ringing in the vicinity. Generally it is desired to reduce this ringing. A second difficulty with sinc interpolation is that the sinc function has infinite extent, making it impractical for digital images of finite dimensions. As a result, approximations to the ‘‘ideal’’ reconstructed image are sought by all practical implementations. Popular reconstruction algorithms such as linear interpolation, cubic convolution, and cubic spline only *approximate* the ‘‘ideal’’ sinc reconstruction [13].

The problem of approximate reconstruction of a non-bandlimited image is fundamentally different from the exact reconstruction of a bandlimited image. In these cases, the traditional analysis of aliasing and truncation errors is seldom used because the expected error is often unbounded. Fortunately, the actual reconstructed images do not tend to reflect such pessimistic estimates. The problem with aliasing and truncation errors is that they are not closely related to the visual fidelity of the reconstructed image. Instead, some fidelity measure must be introduced to predict and control image distortion. See [11, 12] for details. To date, the study of adequate fidelity measures for image reconstruction that incorporate the degradations due to imaging sensors is lacking.

This paper will compare the new methods derived herein with cubic convolution. Cubic convolution is a third-degree interpolation algorithm originally suggested by Rifman and McKinnon [17] as an efficient approximation to the theoretically optimum sinc interpolation function. Its interpolation kernel is derived from constraints imposed on the general cubic interpolation formula to share a truncated appearance of the sinc function and also keep the support small. Cubic convolution has a free parameter A which is studied in [13]. That paper discusses the selection of A based on the frequency content of the image. In a recent paper [16], the frequency domain analysis developed in [12] was used to show that the additional parameter of the two-parameter cubic filter [10] beyond that of the one-parameter cubic convolution does not improve the reconstruction fidelity.

2.1. Image Formation

Image formation is generally described as a sequence of filtering operations. We briefly review our model here. More details can be found in the references [1,9,15]. Also, see [8] where sampling and image formation are treated together.

Let $f(x,y)$ be the intensity distribution of a scene at the front aperture of a lens. That distribution is acted upon by $h_1(x,y)$, the blurring component of the lens, yielding $f_1(x,y)$. A geometric distortion function $h_2(x,y)$ may be added to the image by the lens to yield image $f_2(u,v)$. For instance, lens aberrations may cause pincushion or barrel effects in the image. Although the blurring and geometric distortions induced by a real lens are not necessarily decoupled in this manner, we choose to model it this way because it lends itself to conceptual simplification. A similar model was used by others for motion-blur and lens aberrations [18, 19]. This serves to interchange spatially-varying point spread functions (SVPSF) with a cascade of two simpler components: a spatially-invariant blur and a warp. Recent work in efficient geometric transformations [4, 22, 24], as well as spatially invariant blur, permit cost-effective solutions to the general linear spatially varying (LSV) problem posed here.

At this point, f_2 strikes the image sensor where it undergoes further blurring by a point spread function $h_3(u,v)$ to generate image $f_3(u,v)$. This reflects the limitations of the sensor to accurately resolve each point without the influence of neighboring points. We choose to use a simple model wherein this blurring takes place within one pixel because there are physical boundaries between photosites which prevent multi-pixel blur. Thus we ignore charge bleeding and cross-talk. This model of degradation shall prove to be important for deriving local restoration algorithms later in this paper.

Image f_3 undergoes spatial sampling as it strikes the discrete photosites in a CCD or CID camera. In tube or vidicon cameras, there is spatial sampling (by the focus spot) in the vertical direction. The combination of h_3 with sampling is known as area sampling. It reflects the finite size of the sampling area. If h_3 is taken to be an impulse, then we have point sampling. This is an ideal concept that is often assumed to be true for theoretical considerations, but is generally not true in practice. In either case, intensities in the sampled image I_s are now defined only for integer values of u and v. The digital image $I(u,v)$ is obtained via an analog-to-digital converter that quantizes the samples of I_s. This completes the image formation process, leaving $I(u,v)$ as the beginning for subsequent processing, including image reconstruction and restoration.

Traditionally, the area sampling PSF is usually folded into the lens PSF. However, in this model the geometric distortion h_2 makes this a spatially-varying process. Thus, in this paper, we only concern ourselves with inverting the blurring function h_3. The lens blur h_1 may be dealt with using traditional techniques. In [2], we also consider partial compensation for h_2.

2.2. Deblurring and our Sensor Model

There is a considerable body of work dealing with linear space-invariant (LSI) systems [14, 15]. We will assume that h_3 is spatially-invariant. For convenience, we will drop the subscripts for f and h. The LSI systems we consider in this paper are *almost* ideal imagers. In particular, we consider systems where the PSF has spatial extent *equal to or smaller than* the inter-sample distance. This accounts for the sensor blurring due to nonoverlapping area samplers with non-negligible sampling sizes, e.g., the electron beam in vidicons, CCD photosites, and sampling apertures in microdensitometers. Although there is actually some inter-pixel effects in CCD cameras and other discrete photosite devices, these are taken to be secondary effects compared to area sampling. Since all physically realizable imaging systems must have a PSF of finite spatial extent, point sampling, the basis for most reconstruction algorithms, is not physically realizable. Researchers have developed complex models of CCD blurring in terms of MTFs (e.g., [20], [3]). Unfortunately, dealing with multi-pixel blur increases the computational complexity because the algorithm will no longer be local. In this paper, we will assume that a model of the PSF exists and consider only two rather simple models: a Rect filter and a Gaussian filter approximated by a cubic B-spline [25]. According to the model developed by Burns, the Gaussian approximation is more realistic. However, as you can see in Fig. 5 the quality of reconstructions is quite good for both.

2.3. Image Restoration

Given that we know I and h, we seek to solve for f. This problem, known as image restoration, is of considerable interest in image processing [1]. Restoration algorithms are used to invert the degradation that enters into the image formation process. It leaves us with a functional form for f that enables us to perform image resampling at any desired location. We seek to derive an efficient restoration algorithm that inverts the point spread function of imaging sensors. This is related to recent work [8] in which the process of image gathering has been directly incorporated into restoration algorithms.

There is an interesting relationship between reconstruction and restoration, as illustrated in Fig. 1. While both processes start from the image samples in I, reconstruction limits itself to the problem of deriving the continuous function R. Restoration attempts to estimate the original input function f. Obviously, the two problems of reconstruction and restoration are related, and the latter is more difficult. With reasonable assumptions, exact reconstruction of R is at least theoretically possible. On the other hand, exact reconstruction of f requires considerably more tenuous assumptions. Even given the necessary assumptions, the problem of restoration is still considerably more difficult.

Figure 1: Relationship between reconstruction and restoration.

2.4. Input Model

Before we can discuss the restoration algorithms we must specify the class of allowable functions. The most elementary constraints on f are: integrability of the product of f and the PSF, and non-negativity. A reconstruction algorithm based on this model does not need to make the type of assumptions traditionally imposed. However, to get practical models, we may wish to heuristically impose more constraints on the model. Thus we will define the class F_0 as the space of continuous functions with a local analytic expansion almost everywhere and a bounded second derivative almost everywhere. We also define a second class, F_1, as the space of continuously differentiable functions with a local analytic expansion almost everywhere, and a bounded second derivative almost everywhere. The almost everywhere statements allow, on sets of measure zero, discontinuities in the second derivative. We further restrict our implementation by assuming that the number of such discontinuities is less than $N+1$, where N is the number of samples.

Due to the bounded second derivative constraint, there is an implicit limit to the amplitude of high frequency components, and an implicit bandlimit to the allowable functions in the class. We find this type of limit more intuitive, as it does not require us to specify a cut-off frequency. Instead of explicitly defining a passband and a stopband, it says that the function can not have ''too sharp a turn'', a concept that naturally combines frequency and amplitude.

3. NEW RESTORATION/RECONSTRUCTION ALGORITHMS

In this section, we develop the mathematics for new image restoration / reconstruction methods which we refer to as *imaging-consistent algorithms*. We say that an algorithm has the imaging-consistent property if it yields an exact reconstruction for some allowable input. These are, in turn, defined by cascading the PSF model of the imaging system with the assumed mathematical model of input images. Thus, the reconstruction is exactly consistent with some input which is indistinguishable from the actual input given the (degraded) image. As we shall see later, algorithms that satisfy this property enjoy good error characteristics for many definitions of error.

We present only one dimensional image models because higher dimensions will be treated separably. Non-separable multi-dimensional filters can also be defined. To simplify our discussions we will use the following conventions. Let us denote the image values by v_i, and pixel boundaries as k_i with regular spacing m. For our algorithms the intensity value v_i will be centered (located) at $k_i+\frac{1}{2}m$. It will also be convenient to let $x = \frac{t-k_i}{m}$, as t varies over a pixel. In most cases, we will not explicitly deal with the end conditions. However, since the algorithms are all local, the reader's choice for the end conditions will have only local effect.

We will briefly describe the following algorithms:

1) A quadratic imaging-consistent algorithm assuming a Rect filter for the PSF. Note that the Rect function is defined to be a constant value over the specified spatial interval.
2) A quadratic imaging-consistent algorithm assuming a cubic B-spline approximation to a Gaussian PSF.
3) An imaging-consistent algorithm with two cubic pieces per pixel assuming a Rect filter for the PSF.

We have tried numerous other methods, and found these to be among the simplest and the best. Readers are encouraged to try and tailor their favorite local PSF and generate their own imaging-consistent reconstruction algorithms.

3.1. A quadratic imaging-consistent algorithm assuming a Rect PSF filter

Once our definition of information (area samples) is accepted, probably the simplest method to consider is an integrable interpolatory quadratic method. For a Rect PSF filter, this is very easy to derive. Again, we are assuming centered pixels. To ensure the function is in F_0, we need continuity, but we also desire a local method. Hence we define the value of the reconstruction at the pixel boundaries k_i and k_{i+1}, which we will refer to as E_i and E_{i+1}. Any method of reconstruction could be used to compute these values though our examples will only include cubic convolution.

Given the value at the edges of the pixel, an additional constraint that the integral across the pixel equal V_i gives exactly three constraints. They are

$$Q(0) = E_i \; ; \quad Q(1) = E_{i+1} \; ; \quad \int_0^1 Q(x)\,dx = V_i$$

From this, one can determine the following quadratic polynomial:

$$Q(x) = E_i + (6V_i - 2E_{i+1} - 4E_i)\,x + 3(E_{i+1} + E_i - 2V_i)\,x^2$$

where $x = \frac{t-k_i}{m}$. The integral of this quadratic over the interval k_i to k_{i+1} is exactly V_i. Using cubic convolution to derive E_i and E_{i+1}, we have

$$E_i = \frac{1}{8}(AV_{i-2} + (4-A)V_{i-1} + (4-A)V_i + AV_{i+1})$$

$$E_{i+1} = \frac{1}{8}(AV_{i-1} + (4-A)V_i + (4-A)V_{i+1} + AV_{i+2})$$

The parameter A is generally in the range [-3, 0] in order to make the cubic convolution kernel resemble the sinc function [13]. This then defines the quadratic restortation algorithm using a RECT PSF, and a free parameter A. We will abbreviate this as QRsR A with QRsR -.5 and QRsR -1 being studied in the experimental section.

The above restoration method is always defined, though at times it will produce restorations with negative values. In general, the maximum value of the second derivative is on the order of the maximum step change in image intensity values.

To define the reconstruction algorithm, we simply blur the resulting restoration by a Rect filter. One can derive a functional form for this that results in one cubic polynomial that spans from the center of one input pixel to the next:

$$V_i + (E_{i+1} - E_i)\,w + (2E_i - E_{i+2} - E_{i+1} + 3(V_{i+1} - V_i))\,w^2 + (E_{i+2} - E_i - 2(V_{i+1} - V_i))\,w^3$$

where $w = x + .5$. We refer to the resulting (globally) continuously differentiable function as QRR A, where A is the value of parameter used for determining E_i and E_{i+1}. QRR -.5 and QRR -1 are examples we will examine in the experimental section.

3.2. A quadratic imaging-consistent algorithm assuming a Gaussian PSF

A second quadratic method is obtained by assuming the PSF is given by $G(x)$, a cubic B-spline approximation to a Gaussian. This PSF has four cubic segments in the single input pixel which are used to weigh the integral. Again, we are assuming centered pixels. To ensure the function is in F_0, we need continuity, but we also desire a local method. Hence we define the value of the restoration at the pixel boundaries k_i and k_{i+1}, which we will refer to as E_i and E_{i+1}.

Given the value at the edges of the pixel, an additional constraint that the integral across the pixel, weighted by the cubic B-spline, must equal V_i gives exactly three constraints. They are

$$P(0) = E_i \;;\quad P(1) = E_{i+1} \;;\quad \int_0^1 P(x)\,G(x)\,dx = V_i$$

From this, one can determine the quadratic polynomial to be:

$$P(x) = ((((48V_i - 24(E_{i+1} + E_i))2x + (-48V_i + 13E_{i+1} + 35E_i))x + -11\,E_i))/-11;$$

This then defines the quadratic restortation algorithm using a Gaussian-like PSF, and a free parameter A. We will abbreviate this as QRsG A with QRsG -.5 and QRsG -1 being studied in the experimental section. Again, the above method is always defined, though at times it will produce reconstructions with negative values. The maximum value of the second derivative is on the order of twice the maximum step change in image intensity values.

To define the reconstruction algorithm, we simply blur the resulting restoration by a sampled version of the cubic B-spline approximation to a Gaussian PSF. One can derive a functional form for this that results in four cubic polynomials per segment connected as a continuously differentiable function. We refer to the resulting algorithm as QRG A, where A is the value of parameter used for determining E_i and E_{i+1}. QRG -.5 and QRG -1 are examples we will examine in the experimental section.

3.3. An imaging-consistent algorithm with 2-piece cubics assuming a Rect PSF filter

Additional smoothness, in terms of global differentiability, seems to be a natural extension. We first tried a single quartic polynomial per segment, but this resulted in too much ringing. Hence we decided to try two cubic polynomials per segment joined such that the overall function was continuously differentiable. Again to ensure continuity, but still allow the method to be local, we need to define the value of the reconstruction at the pixel boundaries k_i and k_{i+1}, which we will refer to as E_i and E_{i+1}. In addition, we need derivative values which we an obtain from cubic convolution. We refer to the left and right side derivative values as E_0' and E_1', respectively. Note that with the integral constraint, this totals seven constraints, but we have eight parameters. We defined the method with a general free parameter K, but any value other than zero, will introduce significant frequency-sample ripple. Hence all the experiments set the free parameter K to zero. In addition, the computation of E has a free parameter A.

The coefficients for the cubics are obtained by solving the seven equations mentioned above. Setting $K = 0$ yields

$$T_1(x) = x^3\,(53760\,V_0 + 156\,E_1' - 1184\,E_1 - 2924\,E_0' - 12256\,E_0 - 11072\,K)/(-692)$$

$$+ x^2(-40320\,V_0 - 117\,E_1' + 888\,E_1 + 2885\,E_0' + 9192\,E_0 + 5536\,K)/(-692) + x\,E_0' + E_0;$$

$$T_2(x) = x^3(53760\,V_0 + 2924\,E_1' - 12256\,E_1 - 156\,E_0' - 1184\,E_0 + 11072\,K)/692$$

$$+ x^2(-120960\,V_0 - 5887\,E_1' + 27576\,E_1 + 351\,E_0' + 2664\,E_0 - 27680\,K)/692$$

$$+ x\,(80640\,V_0 + 3694\,E_1' - 18384\,E_1 - 234\,E_0' - 1776\,E_0 + 22144\,K)/692;$$

$$+ (-13440\,V_0 - 731\,E_1' + 3756\,E_1 + 39\,E_0' + 296\,E_0 - 5536\,K)/692.;$$

To define the reconstruction algorithm, we simply blur the resulting restoration by a Rect filter. One can derive a functional form for this that results in two quartic polynomials per segment, and results in a continuously twice differentiable function. We refer to the resulting reconstruction algorithm as TPCRR A, where A is the value of parameter used for determining. TPCRR -1 will be discussed in the experimental section.

4. ANALYSIS OF RESTORATION/RECONSTRUCTION ALGORITHMS

We now present the analysis of the imaging-consistent algorithms. Our analysis has three components:

1) Spectral comparison of impulse responses. Comparisons will be made directly on the spectrum and using Park and Schowengerdt's SR blur measure ($e^2(v)$).
2) Reconstruction/restoration of actual images.
3) The relationship between imaging-consistent algorithms, information-based complexity, and linear error measures.

Each of the subsections will begin with a more complete description of the associated analysis. We point out that in all cases we considered a larger set of parameters and classes of imaging-consistent algorithms than those presented here. We have chosen to present the parameters/classes which we subjectively felt were the best of the new algorithms. For a particular application, however, the approach used to derive these algorithms might, with different parameters, yield a ''better'' filter than those presented here.

In addition to our own algorithms, we will be analyzing four other algorithms. These algorithms will be piecewise linear, cubic convolution with $A = -.5$ and $A = -1$ (see [17]), and a cubic spline interpolation using not-a-knot end conditions (see [5]). We will refer to ideal filtering, which for reconstruction would be an infinite sinc filter. We feel these five algorithms provide a good coverage of the reconstruction algorithms currently in use. In addition, the spatial extent of the cubic convolution is approximately equal to that of the new imaging-consistent algorithms.

4.1. Spectral analysis of impulse response

In this section we present the result of our analysis of the impulse response of the algorithms. The impulse responses themselves are not shown (they mostly look similar to that of cubic convolution), but rather a spectral analysis is presented. In general, it is believed that a good reconstruction filter will approximate the infinite sinc filter, except that the transition between the stopband and passband will be smooth. In general then, one wants a filter which will have low amplitude in the stopband, almost unit response in the passband, and a smooth transition between the two regions.

For restoration, the ''ideal'' filter is not simply sinc interpolation. Assuming a bandlimited function sampled above its Nyquist rate, the spectrum of an ''ideal restoration'' filter for our problem is (ignoring noise) the spectrum of the sinc interpolation filter divided by the spectral response of the PSF. Thus, one does not expect unit response in the passband, but one does expect zero response in the stopband. For the PSFs defined by the spatial box-filter and the single pixel wide cubic B-spline, the logarithm (base 10) of the response of the sinc divided by the spectrum of the PSF is shown in Figs. 2a and 2b, respectively. Again, we might ask for a smooth transition between the stopband and passband. A good restoration filter will thus have small amplitude for high frequency, and a passband shaped like those in Figs. 2a and 2b.

Each algorithm is analyzed by means of interpreting the data in two graphs: the logarithm of the frequency response (log MTF), and the Park SR blur error measure. For a given algorithm, these were computed as follows. First, the impulse response was obtained by applying the algorithm to an 65 pixel idealized ''image'', $IN[0..64]$ with $IN[32]$=64 and $IN[i]$=0 otherwise. This image was then reconstructed/resampled to result in an output vector OUT[0..2048] (resulting in 32 new samples per original pixel). The frequency response (MTF) was computed as the sum of the square of the real and imaginary components of the FFT of the impulse response in OUT. The plots show the logarithm (base 10) of the MTF, with all values less than 10^{-6} clipped. The graph is in cycles per original pixel. Since the MTF is symmetric, only positive frequencies are shown. Furthermore, since we magnified by a factor of 32, each horizonal position in the graph represents 1/32 of a cycle in the original spatial domain. Dotted vertical lines are drawn at the cut-off frequency of the 1/2 cycle per original pixel (these separate the stopband and passband), as well as a

Figure 2: Graphs of the log spectral response of inverse filters for restoration. Note how both inverse filters cause an amplification of frequencies near the Nyquist rate. a) Inverse Filter when PSF = Rect, b) Filter when PSF = B-spline approximation to a Gaussian.

horizontal dotted line at unit amplitude. The spectrum of an infinite sinc filter (in the spatial domain) would be the Rect filter contained in these two dotted lines. We note that all of the filters are linear phase, and hence phase spectrums are not shown. (There is a 1/2 cycle phase shift for the imaging-consistent algorithms.)

As one can easily see, the spectral/spatial response of the new reconstruction filters compares quite favorably with the previous algorithms. If one looks at the regular MTF (not shown), almost all of the new algorithms have very flat response in the passband. What is generally of more concern is the rate of decay in the stop band. The fall-off is more visible in the log plots shown in Figure 3. Also of concern is the smoothness of the transition between the passband and the stopband. To summarize, the spectral response of the new quadratic reconstruction filters (QRR -.5 and QRG -.5) is superior to that of the previous local algorithms and comparable to the global cubic spline. Some of the other new algorithms (e.g. QRR -1, TPCR -1) have much faster fall-offs, but do not have smooth transitions. We also note that spectrum of the the restoration algorithms (QRsR -.5 and QRsR -1) are not quite as good if viewed as reconstructions, but resemble smoothed versions of the inverse filters (compare to Fig. 2a).

In [12], a method of evaluation of reconstruction filters was proposed. Although the general form of the criterion is image-dependent, they derived an important quantity, $e^2(v)$ (SR blur), which provides a scene-independent measure of reconstruction filter quality. This criterion was recently used in [16] to compare cubic convolution with $A=-.5$ to the two parameter family of cubics studied in [10]. (Their two-parameter family includes cubic convolution and cubic B-spline approximation as a special case). They concluded that cubic convolution with $A=-.5$ was the best method, at least according to this criterion. As we shall see, according to this criterion our new filters are even better.

The error function to be used is

$$e^2(v)=|1-h(0)|^2+\sum_{n\neq 0}|h(v-n)|^2$$

where $h(v)$ is the impulse response of the reconstruction filter at frequency v. Both n and v are measured in cycles per pixel. Note that this definition considers the folding of higher frequencies into low (the summation term), and thus can be applied even for non-bandlimited reconstruction filters. See [12] for a better explanation of this error term, as well as an insightful discussion on its use for studying reconstruction from images sampled below the Nyquist rate.

We have approximated (summing over the vector *OUT* defined above) the SR blur ($e^2(v)$) measure for all of the filters studied herein. For reconstruction, a sinc filter is optimal and its error measure, $e^2(v)$ would be 0 below .5 cycles/pixel and 2 for frequencies greater than or equal to .5 cycles/pixel. Since [16] shows that $A=-.5$ is optimal for cubic convolution, we will use that as a standard algorithm. In addition, we will also use cubic spline interpolation as a standard. Figure 4a shows the Park error criterion plots for these two algorithms (dots and * patterns respectively), as well the new quadratic reconsruction algorithm using a RECT and a free parameter of -1 (QRR -1). Note that cubic splines outperform the cubic convolution, but that for most frequencies, the imaging consistent algorithm outperforms the cubic spline. More graphs showing the error properties of the new algoirhtms are shown in Figures 4b-d. To facilitate comparison the measure for the ‘‘optimal’’ cubic convolution (CC -.5) appears in each plot.

As one can readily see, the new algorithms are in general as good, and in some cases always better than cubic convolution. We generally fare well compared to cubic splines, and again at times are significantly better. The ideal spectral response for reconstruction was a Rect filter, and is instrumental in the development of the above error measure. Since the ideal spectral response for restoration is not a Rect function, we are not sure how to interpret these error characteristics for the restoration algorithms, but include them for comparison.

Now that we have introduced the measures and plots, we will discuss the results for the different algorithms. Figure 3a shows the log MTF associated with linear interpolation. The log MFT of the the central impulse from a cubic spline with not-a-knot condition is shown in Figure 3b. Note the poor spectral and error characteristics of linear interpolation. For the cubic spline, the spectrum is reasonably flat in the passband, and small in the stopband. It also performs well according to the SR blur measure (see Figure 4a). Overall, the cubic spline is among the best filters Unfortunately, it is a global algorithm. Linear interpolation, has the worst spectral characteristics of any filter tested, but is included because of its prevalent use and low complexity.

Figure 3c shows the log MTF graphs associated with cubic convolution with $A = -.5$ (CC -.5), and 3d shows the log MTF for it with $A = -1$ (CC -1). For CC -.5, we have good performance according to the SR blur measure, see Figure 4a), and reasonable response in the stop band. Cubic convolution with $A = -1$ appears to have better high frequency damping. It has better SR blur measure for a small range of frequencies, but for most frequencies CC -.5 has a lower value of $e^{2(v)}$. Overall, both are good filters, with CC -.5 having a slight edge.

Figure 3e-o shows the log MTF for various imaging-consistent algorithms. Figure 3e shows the graph associated with our new quadratic restoration method assuming a Rect filter for the imaging system PSF, and using cubic convolution with $A = -.5$ to determine interpolation points (QRsR -.5). Figure 3f shows the performance of the associated reconstruction filter (QRR -.5). The graphs of their SR blur measure are shown in Figure 4b. For reconstruction we have admirable performance according to $e^2(v)$ (better than cubic convolution), flat response in the passband and reasonably fast decay in the stopband. The restoration filter, QRsR -.5, has the requisite amplification of frequencies near the Nyquist rate, and reasonably good performance in the stopband. Overall, both are very good filters; they are local too.

Figure 3g and shows the graphs associated with our new quadratic restoration method assuming a Rect filter for the imaging system PSF, and using cubic convolution with $A = -1$ to determine interpolation points (QRsR -1). Figure 3h is the associated reconstruction method, QRR -1. For reconstruction we have our best performance according to the SR blur measure (see Figure 4a); it alway performs better than the optimal cubic convolution, and generally performs better than cubic splines. The filter QRR -1 has very flat response in the passband and faster decay in the stopband than does QRR -.5. The restoration filter QRsR -1, has a better approximation of the inverse filter than did QRs -.5. It also has better stopband performance. Both are very good local filters, rivaling global filters. The reconstruction algorithm, QRR -1, is overall probably one of the two best filters that we tested.

Figure 3i shows the log MTF plot associated with our new quadratic restortation method assuming a cubic B-spline approximation to a Gaussian filter for the imaging system PSF, and using cubic convolution with $A = -.5$ to determine interpolation points, QRsG -.5. The associated reconstruction method, QRG -.5, generates Figure 3j. The $e^{2(v)}$ performance of both QRsG -.5 and QRG -.5 are shown in Figure 4c. The spectral response of the restoration algorithm, QRsG -.5, does not match the inverse filters, in fact it seems to behave more like a reconstruction algorithm than a restoration algorithm. Note the good SR blur ($e^2(v)$) performance, and the reasonably flat response in the passband. The reconstruction algorithm, QRG -.5, had the best performance in the stopband, very flat passband response, and reasonably good $e^2(v)$ performance. Overall, the reconstruction filter, QRG -.5, is one of the top two filters we tested.

Figure 3k and 3l shows the log MTF graph associated with QRG -1, our new quadratic method assuming a B-spline approximation to a Gaussian filter for the imaging system PSF, and using cubic convolution with $A = -1$ to determine interpolation points (QRsG-1 and QRG -1 for restoration and reconstruction, respectively). Although QRsG -1 results in better restoration performance than QRsG-.5, both are significantly smoother than the inverse filter. The reconstruction algorithm, QRG -1, has among the the best performance in the stopband, very flat passband response, and reasonably good $e^2(v)$ performance (see 4c. It has poorer performance at very low frequencies, which are quite important for relatively large magnifications.

Figure 3m-3n shows the graph associated with our 2-piece cubic method assuming a Rect filter for the imaging system, and using cubic convolution with $A = -1$ to determine interpolation points (TPCRsR -1 and TPCRR -1). We developed this method because we thought that the additional smoothness (everywhere differentiable, i.e., F_1) might yield better performance. While it did result in better stopband performance, it did not yield an overall better algorithm. Given its additional cost, we do not feel it is a competitive filter. The MTF graphs for $A = -.5$ (not shown), are similar but not as good. The SR blur measure for these methods are show in Figure 4d.

We tested other parameter settings for the algorithms, none resulting in overall better performance. The only other noteworthy item found in testing was for our new quadratic method assuming a Rect PSF with the pixel boundaries determined using linear interpolation, (QRR Linear). This gives the restoration method a support of only three pixels and greatly reduces the computational costs for both restoration and reconstrucion. (The associated reconstruction method has a support of 5 pixels). The spectral response of the reconstruction method (see Figure 3o) is almost indistinguishable from that of cubic convolution using $A = -.5$; so too is its error characteristics, see Figure 4b. Reconstructed images using this method are also visually indistinguishable from cubic convolution using $A = -.5$ and have not been included.

4.2. Experimental evaluation using real images

In this section we present a short experimental comparison of the different algorithms using real images.

The image used in these examples were obtained using a Sharp JX-450 color scanner. This is a 300 PPI scanner with 6 bit accuracy (it returns 8 bits, of which only the leading 6 are significant). We have been assured that in full resolution mode (300 PPI) each pixel corresponded to a CCD photo-site [21]. The physical movement of the scanning element was, approximately, in the positive y direction of the images. We converted the RGB scanner output to a monochrome image.

This scanner was used to scan a part of pattern 12 from the IEEE STD 167A-1987 Facsimile Test chart. This pattern is a W. and L.E. Gurley type Pestrecov star pattern. We scanned only a 64 by 64 pixel (.21333 inch square region from this pattern. Because of its size we do not show the original data. However, from the reconstruction the reader should recognize the original input pattern (but not necessarily the degradations imposed by the scanner. In the reconstructions, there is a white band cutting across the star pattern. This circular band corresponds to the 200 lines per inch circle in the original pattern. Just visible in the upper left of imaging-consistent reconstructions is the 100 line per inch circle from the pattern. In the other reconstructions, this artifact is clipped by the display process described below. As the lines on the original pattern converge (toward the lower right), they also approach and then surpass the Nyquist rate for the scanner, thus some aliasing is to be expected. We note that on the left of the original scanned pattern was a slight imperfection which is visible as a white blob in all of the reconstructions.

The reconstructions were obtained by magnifying the original 64 by 64 image up to a 512 by 512 image (i.e., by a factor of 8). The magnification was done separably, first in X, then in Y. The 512 by 512 images were cropped down to 500 by 500 to allow the images to fit in the allotted space at full resolution. In addition, all of the methods have boundary conditions which may cause edge effects. (We handled the boundary conditions by pixel replication.) The 500 by 500 magnified images were Linotyped at 2540 DPI into a space of 3.55 inches. Using Postscript, halftones were generated at 141 lines per inch. This gives approximately 500 pixels with 256 graylevels.

The actual images appear in Figure 5. Figure 5a has 4 panels. The upper left shows the results using CC -.5; upper right shows results using a cubic spline; lower left shows results using QRR -1; and lower right shows results using QRG -1. As one can see, their is considerably more aliasing (ringing) on the right and bottom of the CC based reconstruction. The new algorithms show some slight ringing (QRG a little more) but are nearly as good as the cubic spline. Figure 5b also has 4 panels. The upper left is reconstruction using QRR-.5; the upper right is reconstuction using TPCRR -1; the lower left shows restoration using QRsR -1; the lower right shows restoration using QRsG-1. The reconstructions, while not as good as QRR -1 and QRG-1, are still comparable to cubic convolution.

We point out that the imaging-consistent algorithms all assume that the data is centered within the pixel, while the cubic convolution, linear interpolation and cubic spline methods all assume the data is at the left of the pixel. Thus there is a four-pixel shift in the imaging consistent reconstructions. This could, of course, be corrected if one prefers the left-justified view of the data.

The restorations have significantly more ringing, as they amplify frequencies near the sampling rate. The fact that the restorations do not look too much like what we expect of the orginal signal suggests that the restoration model, either the PSF and/or the functional model, is rather weak. Our algorithm only allowed discontinuities of the second derivative at boundaries of the original samples. To get a better restoration, this needs to be relaxed to allow them to occur at any point within the orginal sample.

4.3. General error properties of imaging-consistent algorithms

The imaging consistent reconstruction approach follows quite naturally from a general approach to algorithm development known as information-based complexity(IBC), see [23]. Because of this relationship, we know that the imaging-consistent algorithms enjoy good error properties for many definitions of error.

In the IBC theory, there are a few key concepts which apply to our situation. First, the only available information is mathematically captured by the *information operator*, which converts an unknown function to input values. In our case the information operator is the area sampling process. Secondly, this operator is applied to an unknown function from an *a priori* known class of functions, which in our case is F_0 or F_1 defined previously. Thirdly, they measure error in another space, e.g. the sum of squares of pointwise distance between functions. Finally, they define the concept of an *interpolatory algorithm* to be one which produces a function from the class of allowable functions such that when the information operator is applied to this function, it returns the measured data. That is, the function interpolates (in a generalized sense) the measured data.

More importantly, it has been proven that such interpolatory algorithms enjoy strong error properties. In particular, interpolatory algorithms have an error at most twice that of any algorithm for *any* error measure defined as a weighted norm on the space of solutions (e.g. L^2, or even a weighted least-squares measure). Note that most image-

quality measures that result in a single number are error measures of this type, e.g. the SQF measure of [7], [6], as well as the extensions of this which include contrast effects or any other weighted integral of the MTF. The measure of [12], when weighted and integrated over frequency v, is also this type of error measure. Thus, an IBC type interpolatory algorithm is, within a factor of 2, an optimal algorithm for each of these image-reconstruction quality measures. Of course, being optimal does not make them any good because the error may be infinite or very large. If however, a good algorithm exists according to the *assumed* mathematical model (functional class, information operator and error measure) then the interpolatory algorithm should be acceptable. In our examples, if there are enough samples in the image so that variation of the second derivative is small, then error of the optimal algorithm is small and the imaging-consistent algorithms should perform well.

By definition, an imaging consistent reconstruction algorithm is an (IBC) interpolatory algorithm if the reconstructed function is from the class (in our case that means ensuring the second derivative is below the prescribed bound). In fact, the imaging consistent property is almost analogous to the IBC concept of interpolatory algorithm. We defined the term ‘‘imaging-consistent algorithm’’ to avoid confusion with the more traditional use, in image processing, of the term interpolatory algorithm.

If the data had been taken as point samples, then an interpolatory algorithm would simply be one that interpolates the data in the traditional sense and has properly bounded second derivative. This is probably one reason why cubic splines are so effective: they interpolate the data and minimize a measure based on the second derivative. If we assume the data was generated by area sampling, an interpolatory algorithm must be one such that when the resulting function is integrated (weighted by the PSF) the results interpolate the data. Our algorithms satisfy this property.

While we started this whole endeavor assuming that area sampling was important, we knew that tremendous progress had been made assuming (properly filtered) point samples. We also knew that existing algorithms performed reasonably well when measured subjectively. To our knowledge, all existing quantitative measures of image reconstruction quality are based on spectral analysis (with the implicit assumption of point sampling) and measure performance with respect to the ‘‘ideal reconstruction’’, a sinc function. Thus we were faced with the question of how to assess the performance of these new filters. Rather than pursuing the arduous task of psychological testing of subjective performance, we decided to use traditional spectral analysis. We felt we would be happy if we came close to traditional filters like cubic convolution. As was shown in the previous two sections, we did more than come close. Using quantitative measures the new imaging consistent reconstruction algorithms generally out-performed cubic convolution, and generally rivaled cubic splines.

5. CONCLUSIONS

This paper introduced a new class of reconstruction algorithms, imaging-consistent reconsructions, that are fundamentally different from traditional approaches. We treat image values as area samples generated by nonoverlapping integrators, which is consistent with the image formation process in CCD and CID cameras. We obtained excellent results by formulating reconstruction as a two-stage process: image restoration followed by application of the point spread function (PSF) of the imaging sensor. Efficient local techniques for image restoration are derived to invert the effects of the PSF and estimate the underlying image that passed through the sensor. We define the imaging-consistency constraint which requires the approximate reconstruction to be the exact reconstruction for some function in the allowable class of functions. This class of functions is defined by bounding the maximum absolute value of the second derivative of the underlying continuous signal. The error of the new algorithms can be shown to be at most twice the error of the optimal algorithm for a wide range of optimality constraints.

The reconstruction algorithms derived herein are local methods that compare favorably to cubic convolution, a well-known local technique, and they even rival global algorithms such as interpolating cubic splines. Evaluations were made by comparing their passband and stopband performances in the frequency domain, as well as by direct inspection of the resulting images in the spatial domain.

The results of this work can be applied to the general problem of nonlinear image warping and reconstruction from nonuniform samples. We have demonstrated the usefulness of the imaging-consistent algorithms in chromatic aberration correction using image warping [2]. Using error measures appropriate for that application, we found the imaging-consistent algorithms significantly out-performed cubic convolution. Future work will extend our technique to deal with other point spread functions, including those that overlap, and more realistic CCD models. (Unfortunately, dealing with overlapping PSFs will almost assuredly require global algorithms).

We can see that while we used simplistic CCD PSF models, the quality of the reconstructions does not appear to be significantly affected by it. We attribute this to the fact that the PSF is used to deblur and then reblur, thus mitigating the effects of errors in that model.

6. REFERENCES

[1] Andrews, H.C. and B.R. Hunt, *Digital Image Restoration*, Prentice-Hall, Englewood Cliffs, NJ, 1977.

[2] Boult, Terrance and George Wolberg, ''Correcting Chromatic Aberrations Using Image Warping,'' *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, June 1992.

[3] Burns, Peter D., ''Signal-To-Noise Ratio Analysis of Charge-Coupled Device Imagers,'' Proc. SPIE vol. 1242, 1990.

[4] Catmull, Edwin and Alvy Ray Smith, ''3-D Transformations of Images in Scanline Order,'' *Computer Graphics*, (SIGGRAPH '80 Proceedings), vol. 14, no. 3, pp. 279-285, July 1980.

[5] de Boor, Carl, *A Practical Guide to Splines*, Springer-Verlag, NY, 1978.

[6] Drago, F.J. and E.M. Granger, ''Optics Requirements for Making Color Fiche Transparencies of Maps, Charts, and Documents,'' *SPIE vol.* 549 *Image Quality*: *An Overview*, pp. 25-33, 1985.

[7] Granger, E.M.''Subjective Assessment and Specification of Color Image Quality,'' E.M. Granger P-86, *SPIE Proc. Image Assessment and Specification*, May 20-22, 1974.

[8] Huck, Friedrich O., Rachel Alter-Gartenberg, and Zia-Ur Rahman, ''Image Gathering and Digital Restoration for Fidelity and Visual Quality,'' *CVGIP*: *Graphical Models and Image Processing*, vol. 53, no. 1, pp. 71-83, 1991.

[9] Jain, Anil K., *Fundamentals of Digital Image Processing*, Prentice-Hall, Englewood Cliffs, NJ, 1989.

[10] Mitchell, Don P. and Arun N. Netravali, ''Reconstruction Filters in Computer Graphics,'' *Computer Graphics*, (SIGGRAPH '88 Proceedings), vol. 22, no. 4, pp. 221-228, August 1988.

[11] Oakley, J.P., and M.J. Cunningham, ''A Function Space Model for Digital Image Sampling and Its Application in Image Reconstruction,'' *Computer Vision*, *Graphics*, *and Image Processing*, vol. 49, pp. 171-197, 1990.

[12] Park, Stephen K. and Robert A. Schowengerdt, ''Image Sampling, Reconstruction, and the Effect of Sample-Scene Phasing,'' *Applied Optics*, vol. 21, no. 17, pp. 3142-3151, September 1982.

[13] Park, Stephen K. and Robert A. Schowengerdt, ''Image Reconstruction by Parametric Cubic Convolution,'' *Computer Vision*, *Graphics*, *and Image Processing*, vol. 23, pp. 258-272, 1983.

[14] Pavlidis, Theo, *Algorithms for Graphics and Image Processing*, Computer Science Press, Rockville, MD, 1982.

[15] Pratt, W.K., *Digital Image Processing*, Wiley, NY, 1990.

[16] Reichenbach, Stephen E. and Stephen K. Park, ''Two-Parameter Cubic Convolution for Image Reconstruction,'' *Proc. SPIE Visual Communications and Image Processing*, vol. 1199, pp. 833-840, 1989.

[17] Rifman, S.S. and D.M. McKinnon, ''Evaluation of Digital Correction Techiques for ERTS Images — Final Report, Report 20634-6003-TU-00, TRW Systems, Redondo Beach, Calif., July 1974.

[18] Robbins, Gregory M. and Thomas S. Huang, ''Inverse Filtering For Linear Shift-Invariant Imaging Systems,'' *Proc. IEEE*, vol. 60, no. 7, pp. 862-871, July 1972.

[19] Sawchuk, Alexander A., ''Space-Variant Image Restoration by Coordinate Transformations,'' *J. Optical Society of America*, vol. 64, no. 2, pp. 138-144, February 1974.

[20] Schreiber, William F., *Fundamentals of Electronic Imaging Systems*, Springer-Verlag, Berlin, 1986.

[21] Sharp Technical Hotline, *Personal Communication*, 1991.

[22] Smith, Alvy Ray, ''Planar 2-Pass Texture Mapping and Warping,'' *Computer Graphics*, (SIGGRAPH '87 Proceedings), vol. 21, no. 4, pp. 263-272, July 1987.

[23] Traub, Joseph F., Gregorz W. Wasilkowski, and Henryk Wozniakowski, *Information-Based Complexity*, Academic Press, New York, 1988.

[24] Wolberg, George and Terrance E. Boult, ''Separable Image Warping With Spatial Lookup Tables,'' *Computer Graphics*, (SIGGRAPH '89 Proceedings), vol. 23, no. 3, pp. 369-378, July 1989.

[25] Wolberg, George, *Digital Image Warping*, IEEE Computer Society Press, Los Alamitos, CA, 1990.

Figures for T.E.Boult and G.Wolberg. "Local Image Reconstruction and Sub-Pixel Restoration Algorithms." Computer graphics and image processing: Graphical models and Image processing (CVGIP:GMIP), Vol 55, No. 1. pp. 63-77, Jan. 1993.

Because the journal printed, in error, the "low-quality" versions of the some of the figures, the fair-use doctrine should allow you to download this paper/figures. (All the reprints, including any new ones ordered, will include those incorrect figures.) Please cite the above journal version.

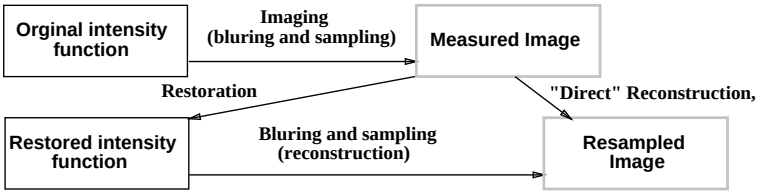

Figure 1: Relationship between reconstruction and restoration.

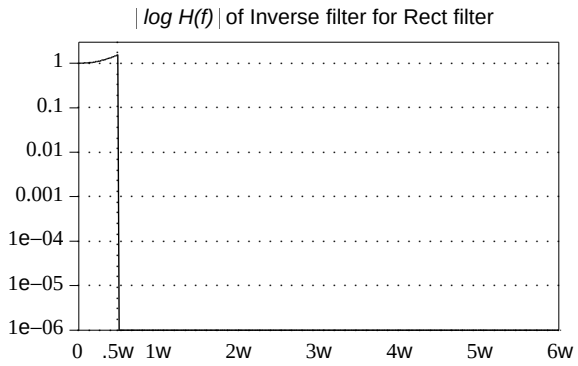

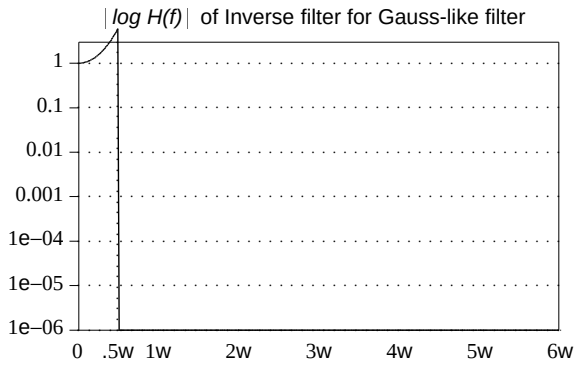

Figure 2: Graphs of the spectral response of inverse filters for restoration. Note how both inverse filters cause an amplification of frequencies near the Nyquist rate. a) Inverse Filter when PSF = Rect, b) Filter when PSF = B-spline approximation to a Gaussian.

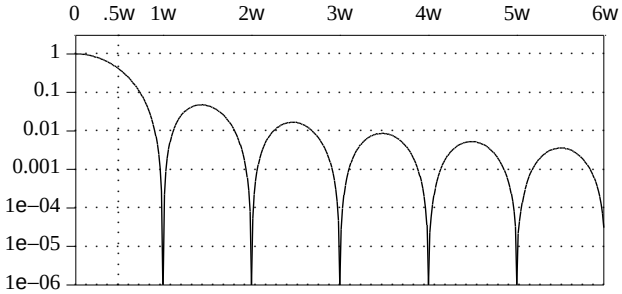

Fig. 3a) $|\log H(f)|$ for Linear Interpolation

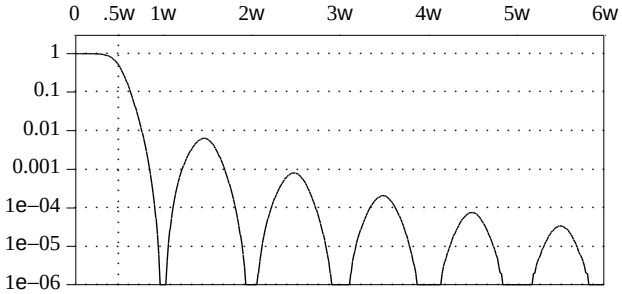

Fig. 3b) $|\log H(f)|$ for Cubic Spline [central impulse]

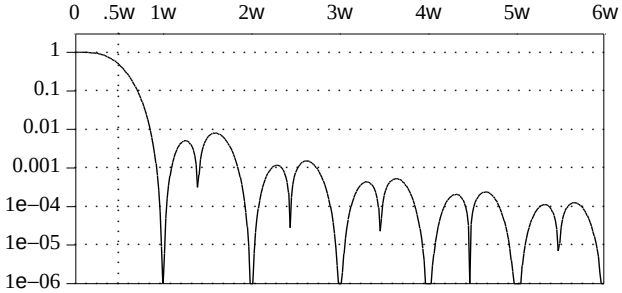

Fig. 3c) $|\log H(f)|$ for Cubic Convolution A -.5

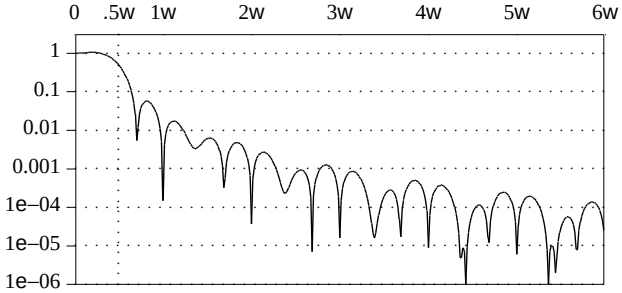

Fig. 3d) $|\log H(f)|$ for Cubic Convolution A -1

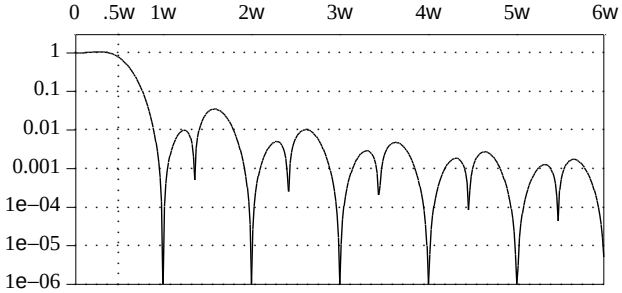

Fig. 3e) $|\log H(f)|$ for Imaging-Consistent Restoration: QRsR -.5

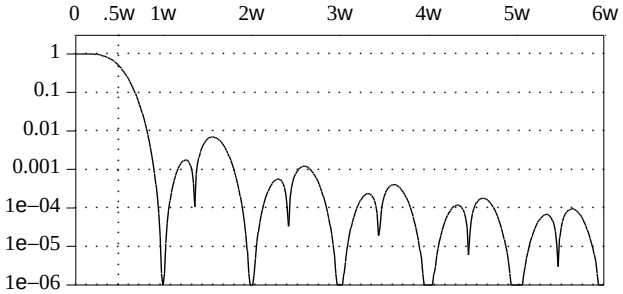

Fig. 3f) $|\log H(f)|$ for Imaging-Consistent Reconstruction: QRR -.5)

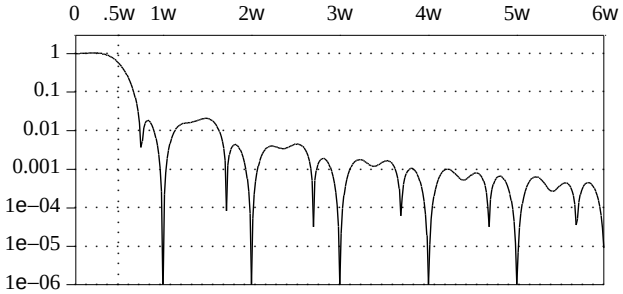

Fig. 3g) $|\log H(f)|$ for Imaging-Consistent Restoration: QRsR -1

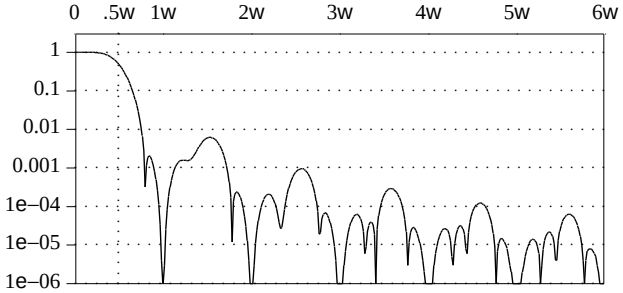

Fig. 3h) $|\log H(f)|$ for Imaging-Consistent Reconstruction: QRR -1

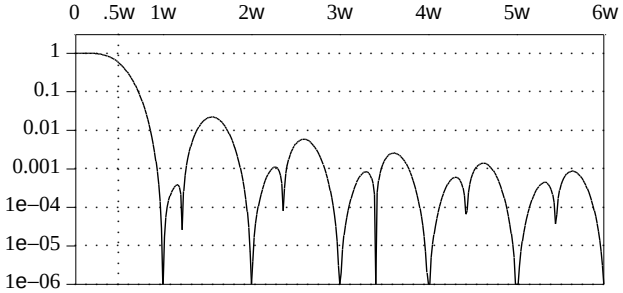

Fig. 3i) $|\log H(f)|$ for Imaging-Consistent Restoration: QRsG -.5

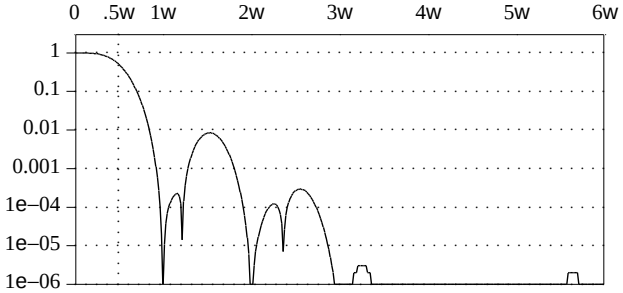

Fig. 3i) $|\log H(f)|$ for Imaging-Consistent Reconstruction: QRG -.5

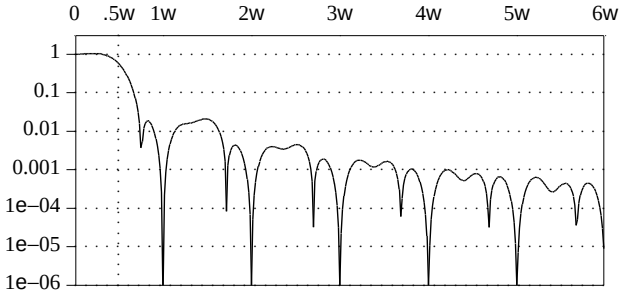

Fig. 3k) | *log H(f)* | for Imaging-Consistent Restoration: QRsG -.5

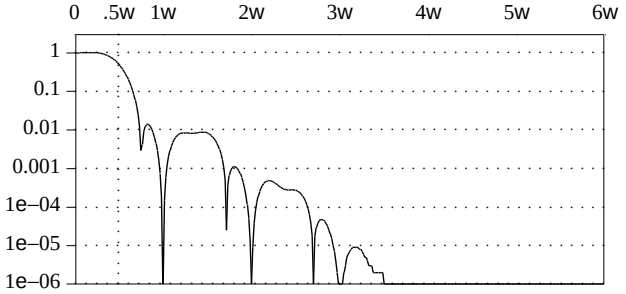

Fig. 3l) | *log H(f)* | for Imaging-Consistent Reconstruction: QRG -.5)

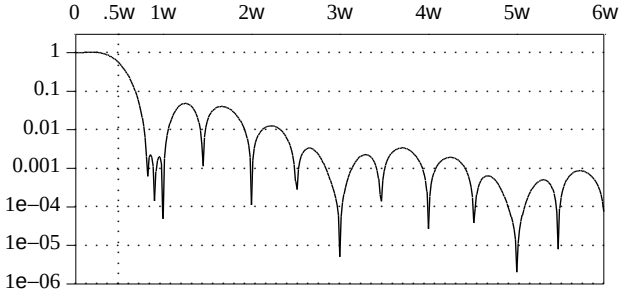

Fig. 3m) | *log H(f)* | for Imaging-Consistent Restoration: TPCRsR -1

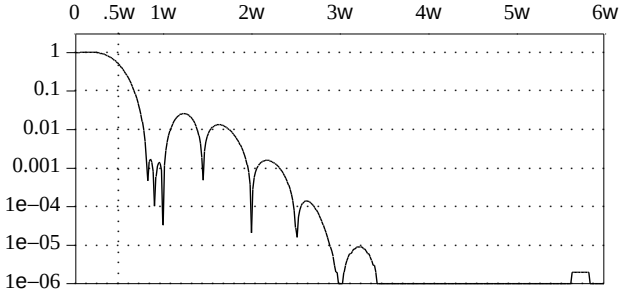

Fig. 3n) | *log H(f)* | for Imaging-Consistent Reconstruction: TPCRR -1

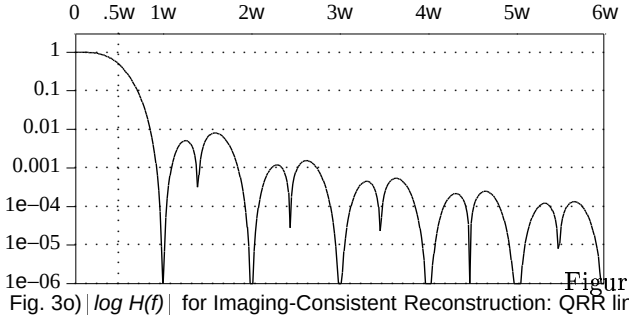

Fig. 3o) | *log H(f)* | for Imaging-Consistent Reconstruction: QRR linear

Figure 3: MTF plots

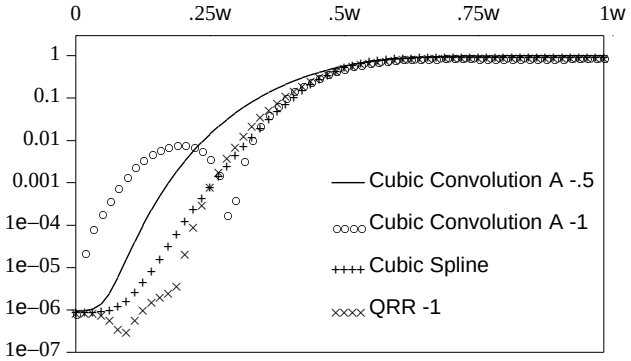

Figure 4a) Error Characteristics of H(f)

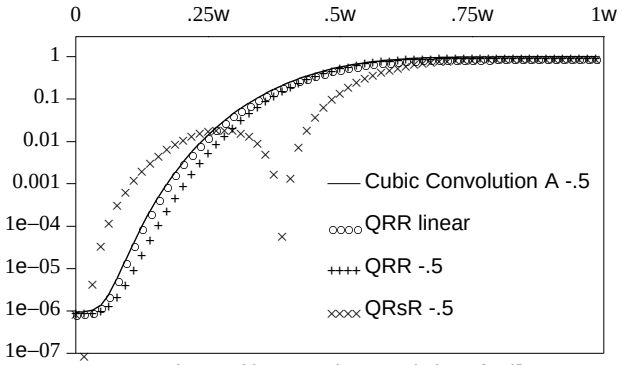

Figure 4b) Error Characteristics of H(f)

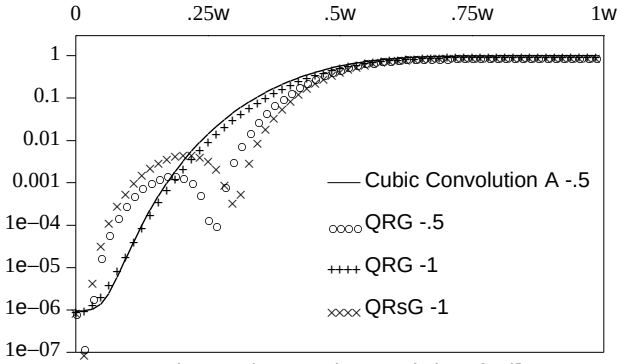

Figure 4c) Error Characteristics of H(f)

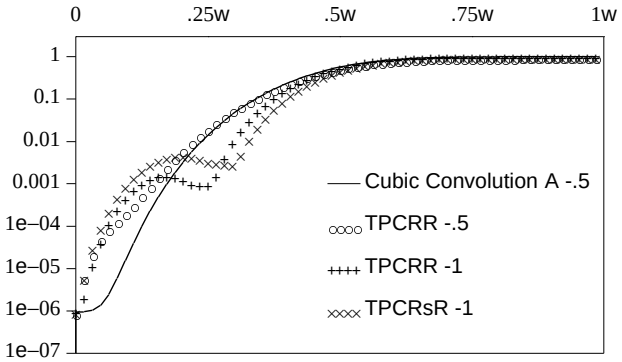

Figure 4d) Error Characteristics of H(f)

Figure 4: SR blur measure

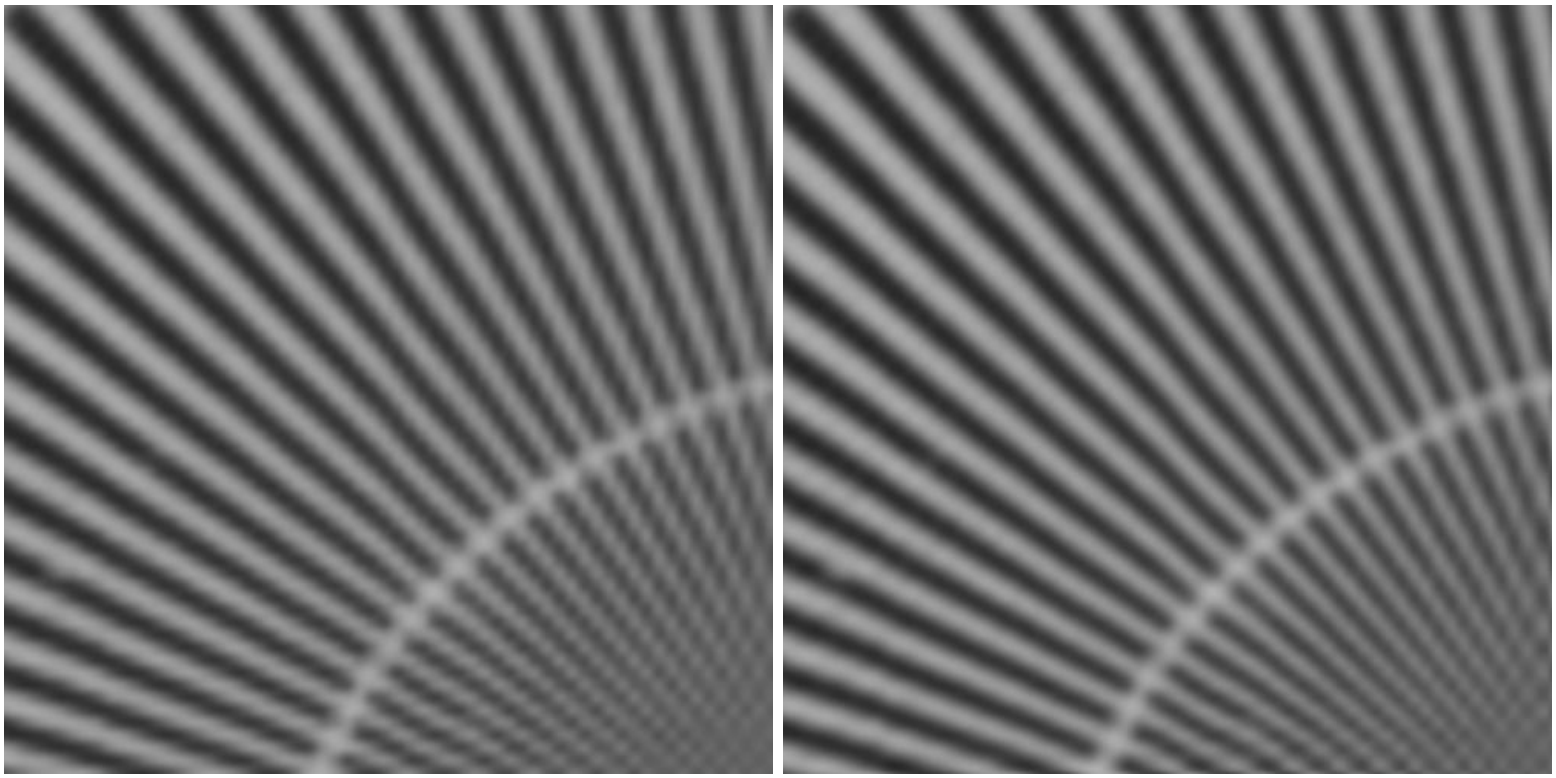

Figure 5: The left image shows the reconstruction, at a sampling rate 8 times the original test pattern, using cubic convolution with A= -.5. Note the aliasing artifacts in the upper right and lower left of the image. The right image shows results using cubic spline interpolation, which is much better.

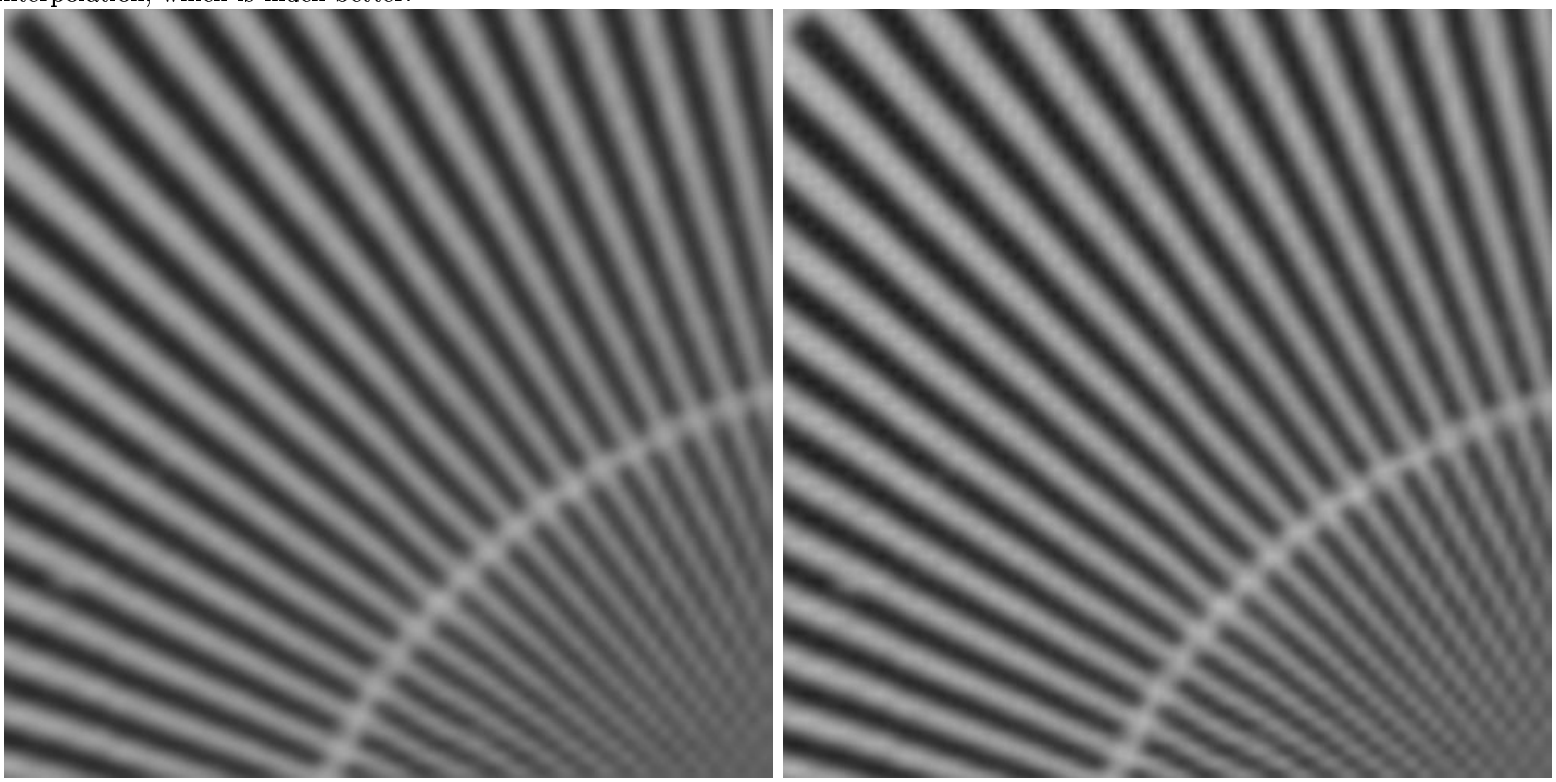

Figure 6: The left image shows reconstruction using the imaging-consistent quadratic-based algorithm with PSF=RECT and a free parameter of -1.0. The right image shows results of the imaging-consistent quadratic-based algorithm with PSF≈Gaussian, and a free parameter of -1.0. The slight shift (4 pixels) is caused by a difference in the interpretation of location of input data. See text for details. Visually, the quality of the new algorithms exceeds cubic convolution, and is about equal to the cubic spline.

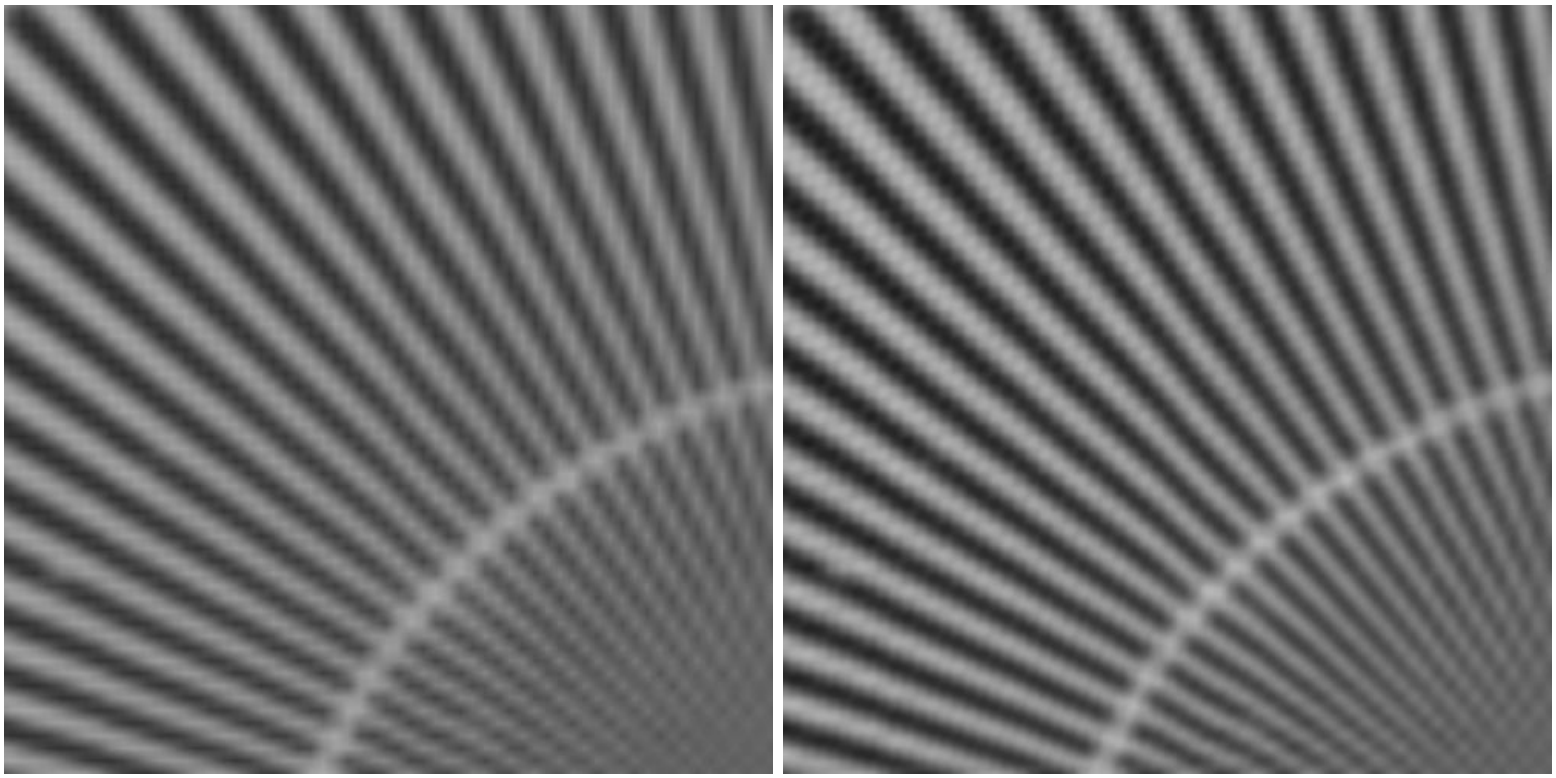

Figure 7: The left shows reconstruction using bi-linear interpolation. The right shows the reconstruction using cubic convolution with parameter -1. While very inexpensive, bi-linear interpolation is noticeably inferior to the others. Cubic Convolution with -1 is not as smooth as with -.5

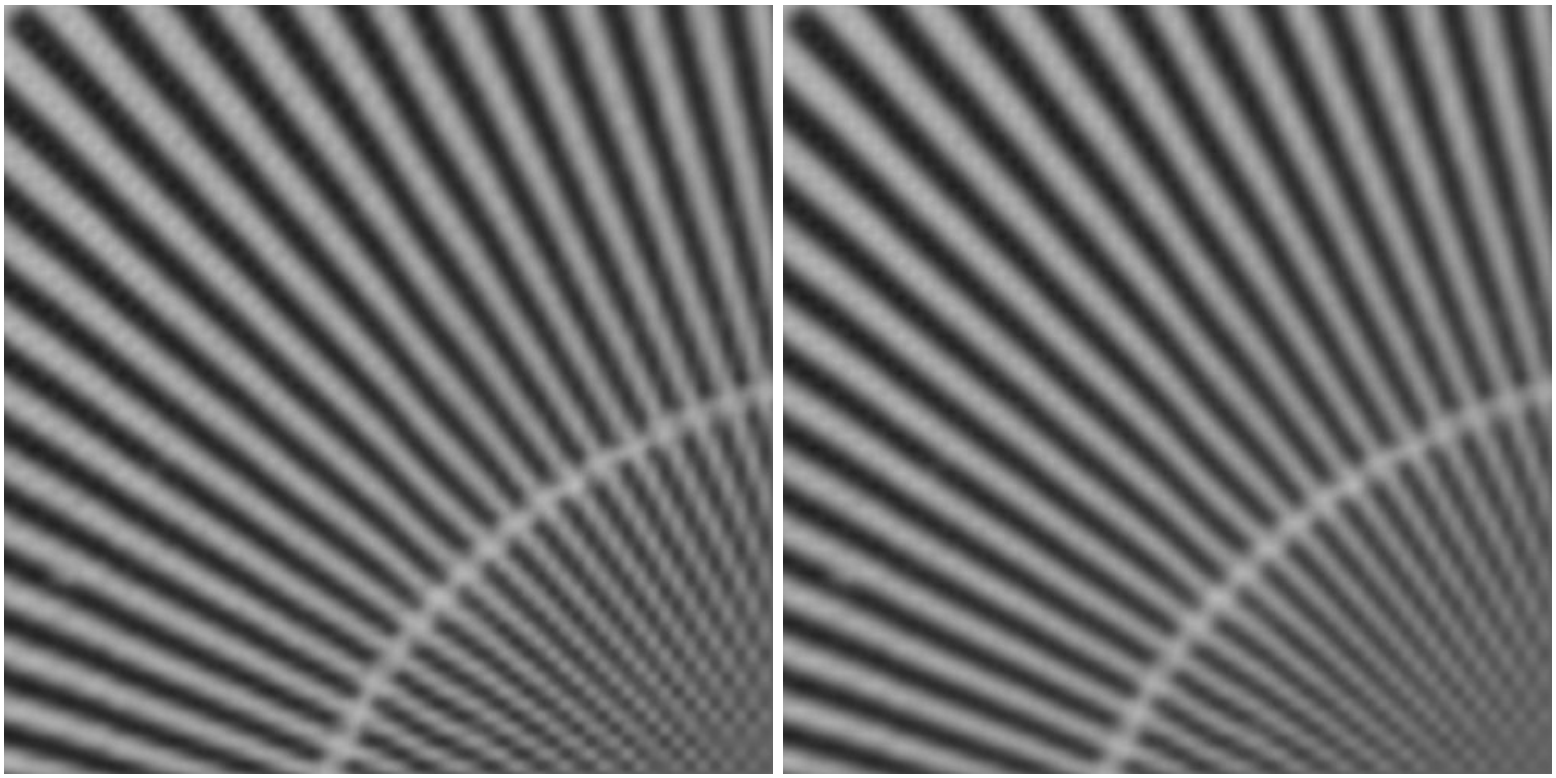

Figure 8: The left shows *restoration* using the imaging-consistent quadratic-based algorithm with PSF=RECT and a free parameter of -1. The right shows *restoration* using the imaging-consistent quadratic-based algorithm with PSF≈GAUSS and free parameter = -1.0. The reconstruction behavior for these setting was show in figure6. While they were very good reconstruction they did not really capture the expected sub-pixel restoration. This is likely the result of overly simplified PSF models.

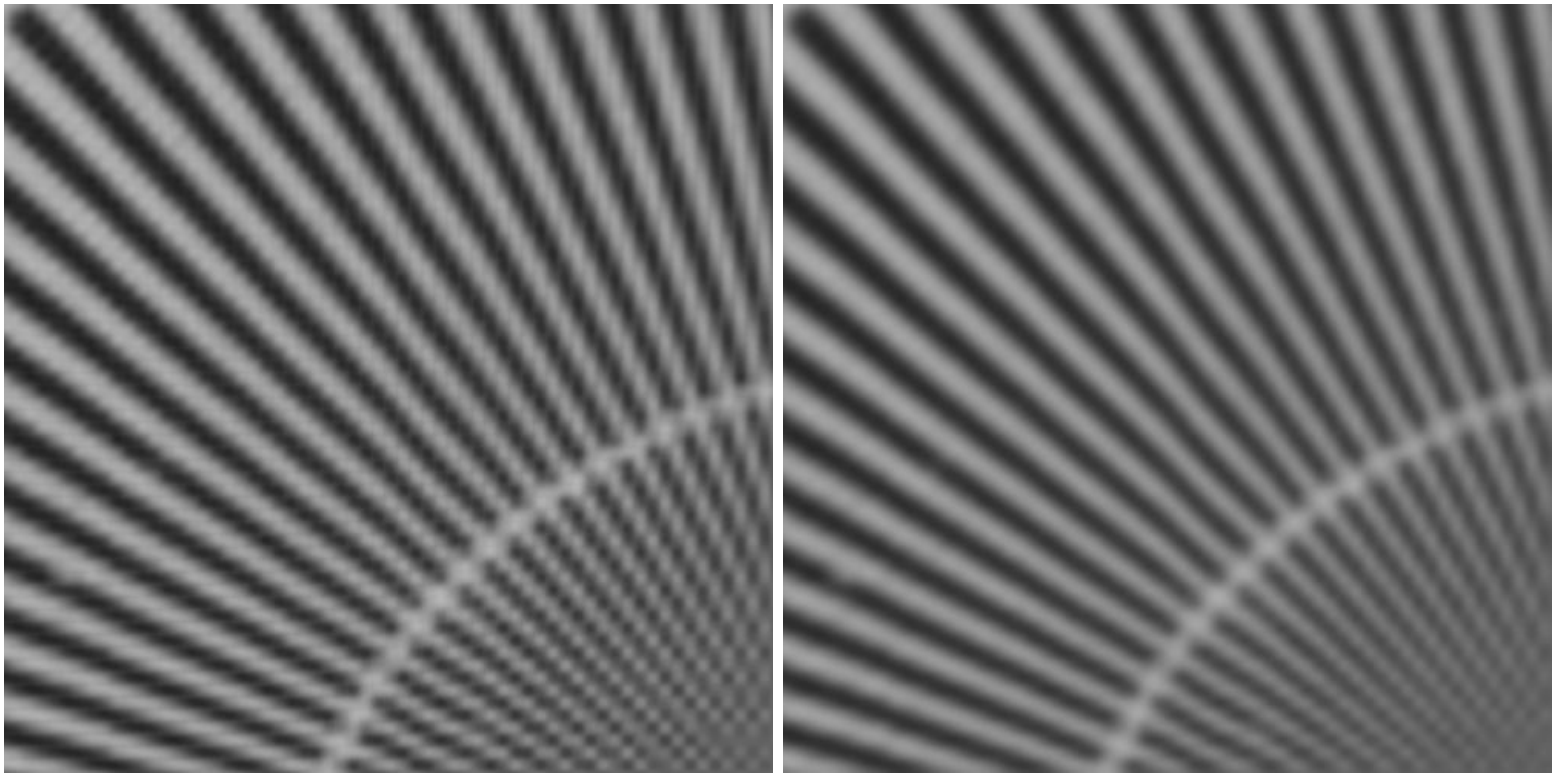

Figure 9: The left shows *restoration* using imaging-consistent quadratic-based algorithm with PSF=RECT and free parameter = -.5. The right shows the reconstruction for the same parameters. While not as good as when the free parameter is -1, it is still quite good.

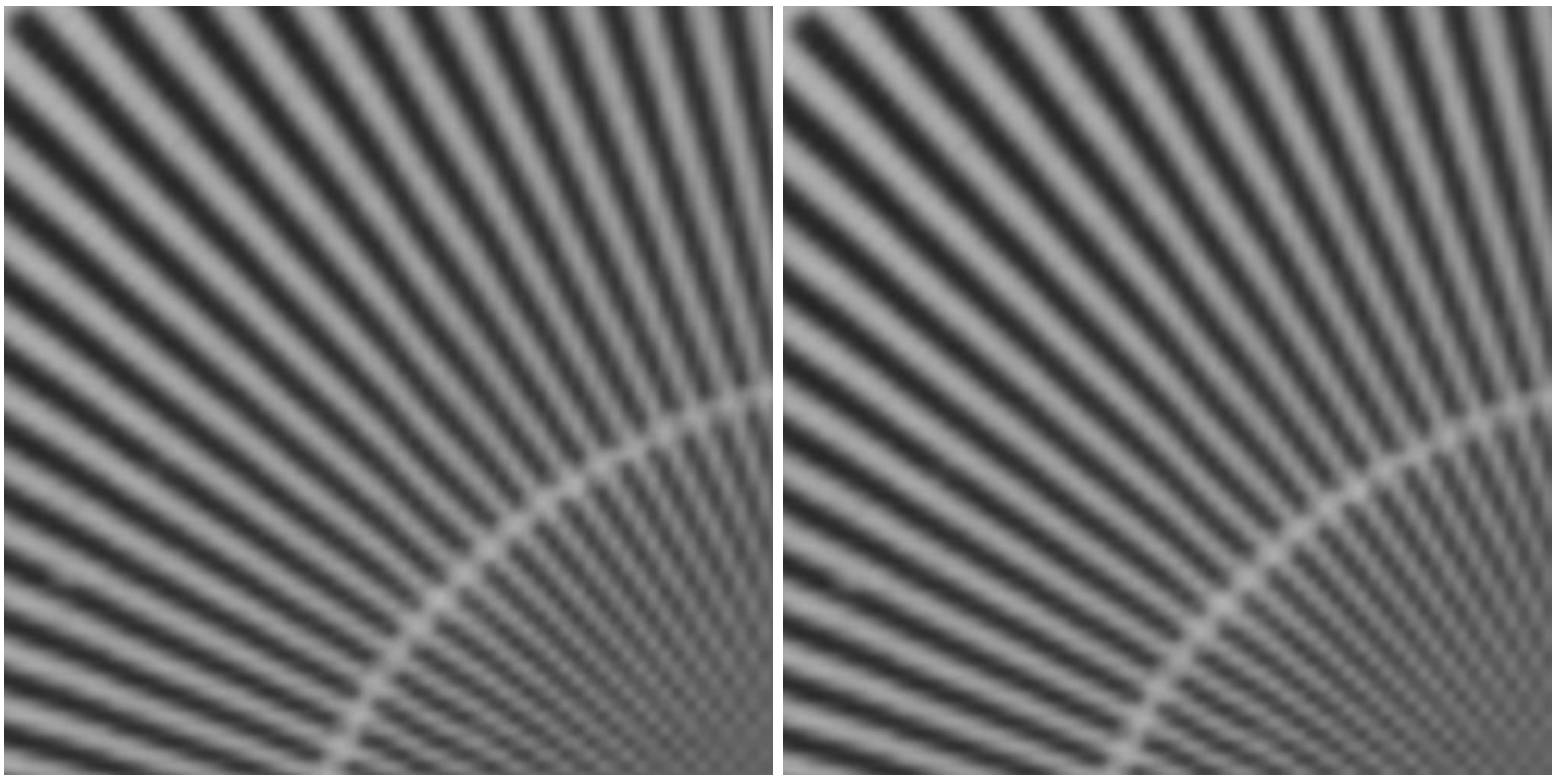

Figure 10: The left shows *restoration* using imaging-consistent quadratic-based algorithm with PSF≈GAUSS and free parameter = -.5. The right shows the reconstruction for the same parameters.

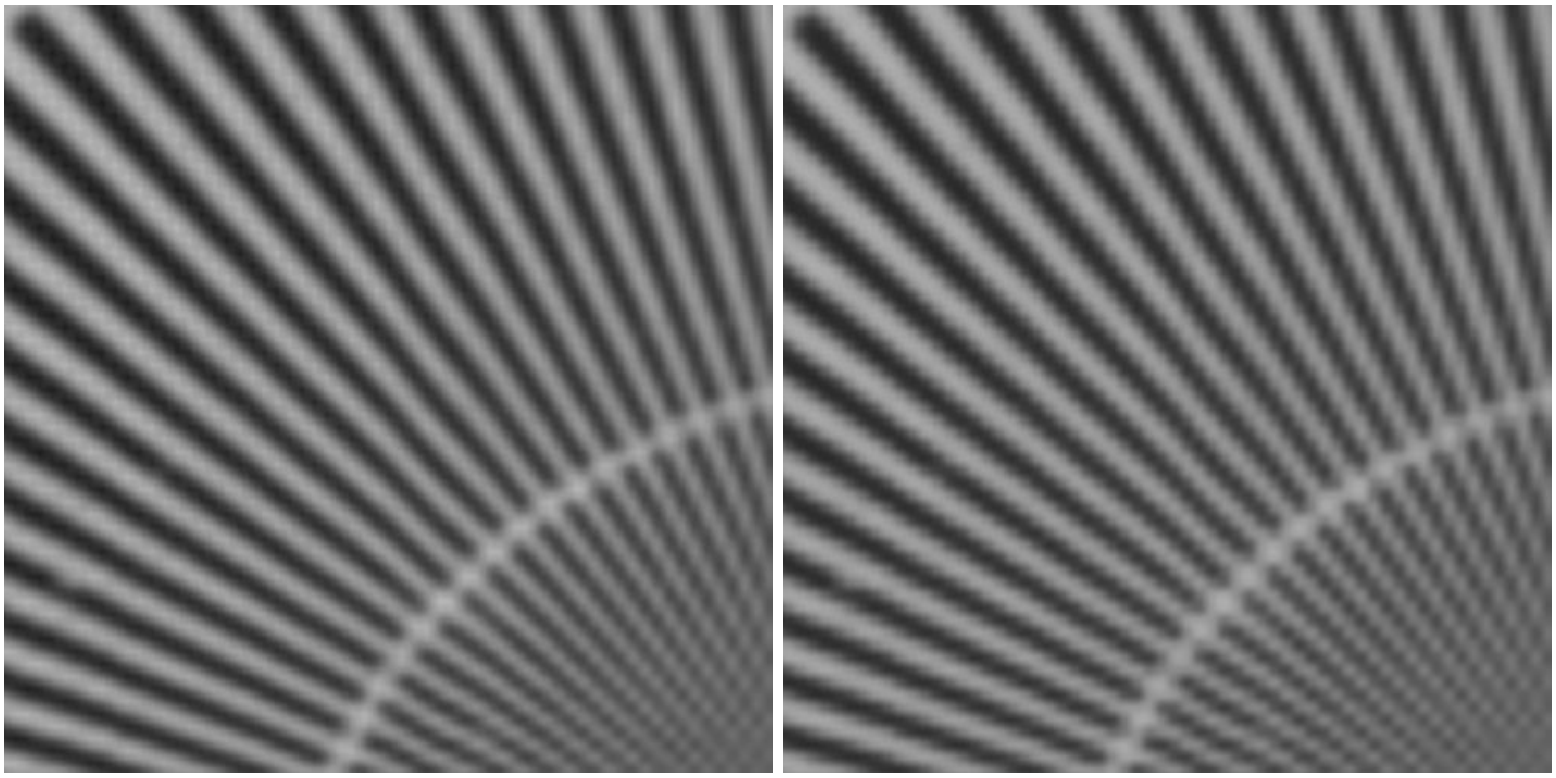

Figure 11: The left shows reconstruction using the imaging-consistent two pieced cubic algorithm with PSF=RECT and free parameter = -.1. The right shows the same algorithm with free parameter -.5. This algorithm has a more complex underlying model of functions, has noticeably better spectral response at higher frequencies, and is more flexibility (as a two piece cubic). However, its visual performance is not significantly better and the quadratic-based algorithms.

To really see the quality of the reconstuction figures you need 2400 dpi printing (or at least 1200dpi). They are fine when viewed on a computer screen at 100dpi (where the added bit depth makes up the difference). So make it easier for people to "print" this on the local printer we include copies of each sub figure rendered a halftone screen of about 53 pixels per inch (so 300-600dpi will be acceptable). To do this we have reproduced all of the subfigures at twice the size (thus providing closer to 1 halftone cell per pixel (the average is .85 cells per pixel). This is still not as good as the actual journal final version. We suggest you hold the following pages at at least twice the normal reading distance (about 2.5 times is best)

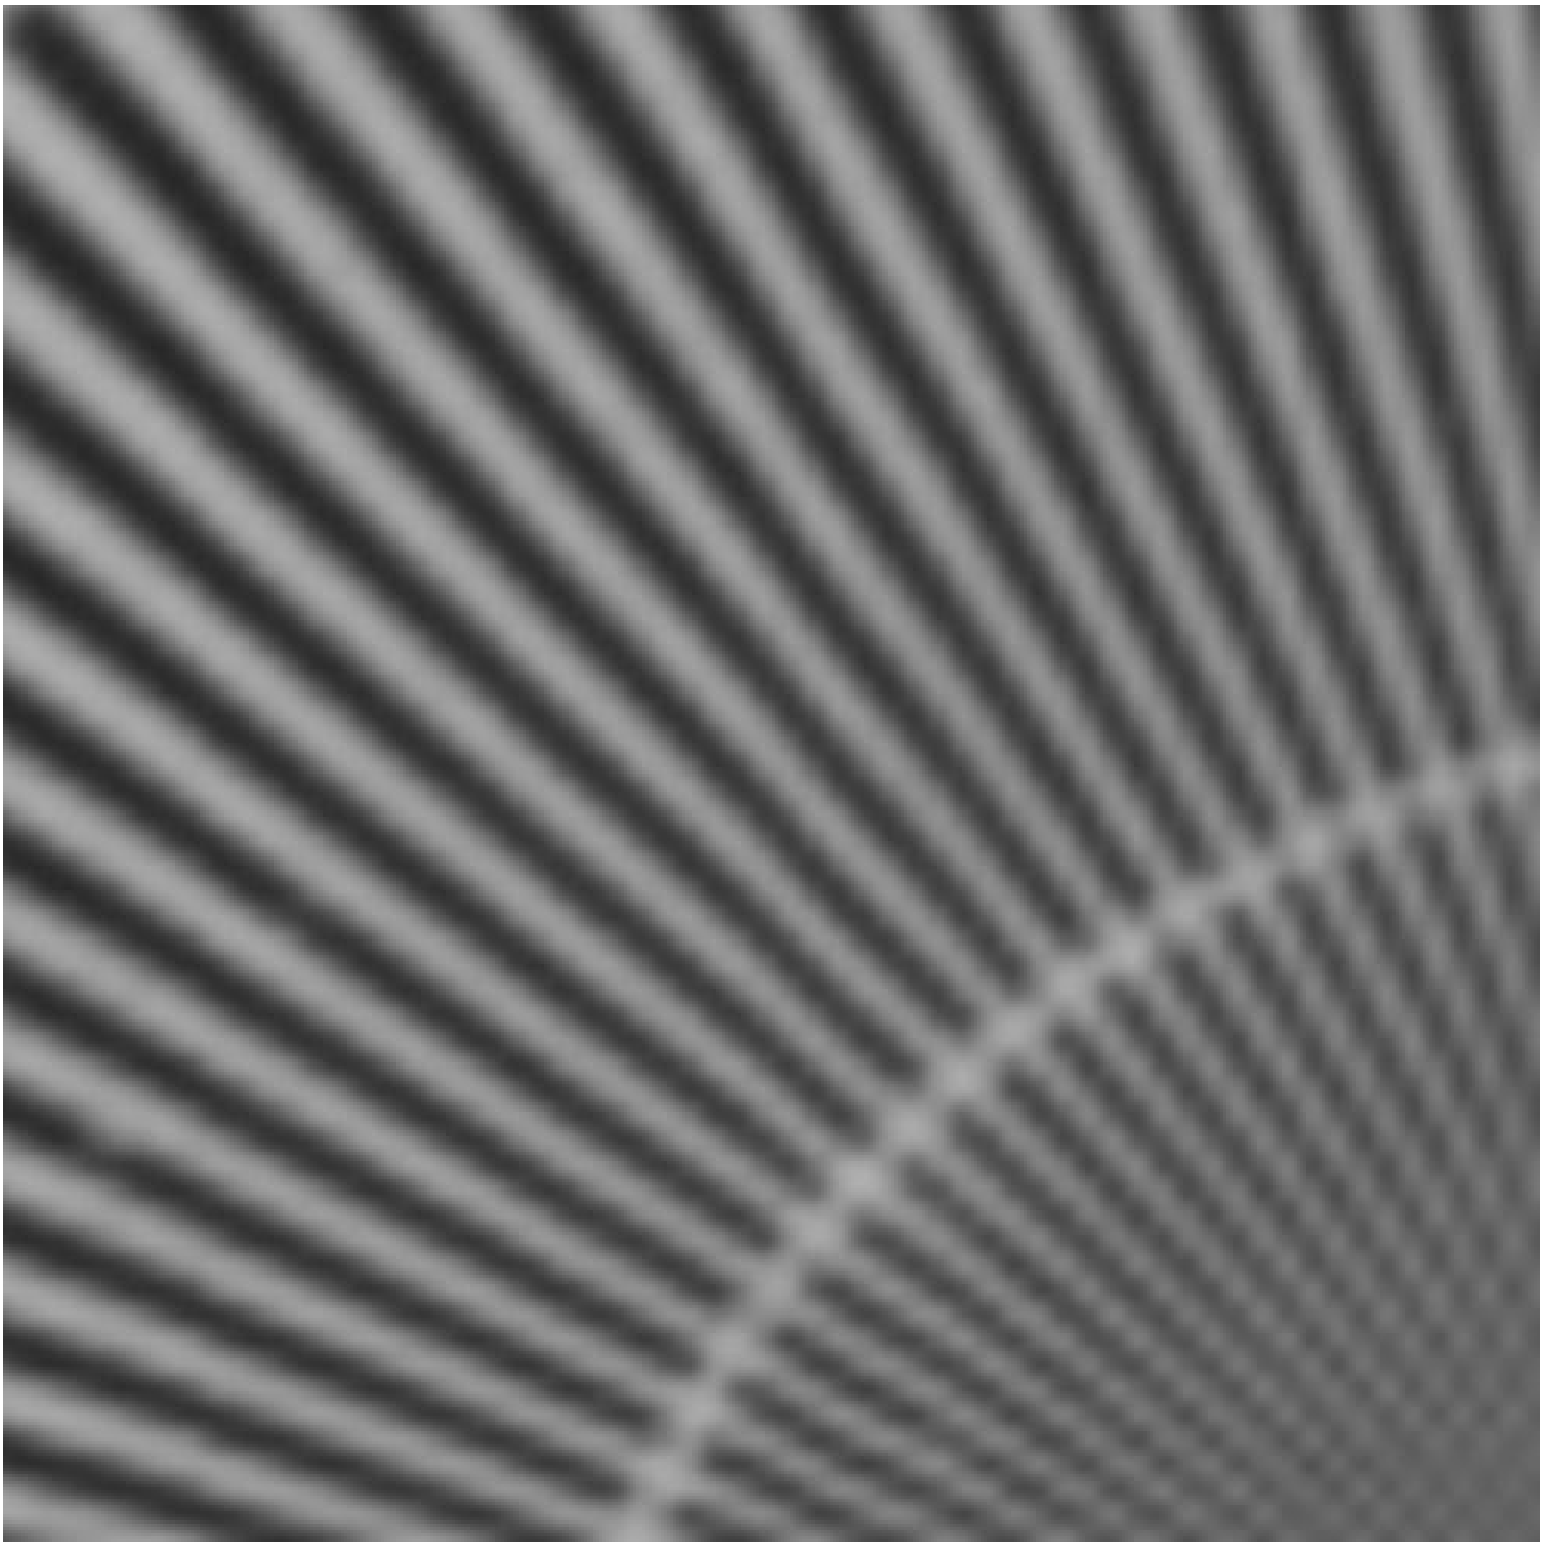

Figure 12: Left part of figure5 enlarged. Note that the final versions would be printed at the quality of the first figure, and this page would not appear.

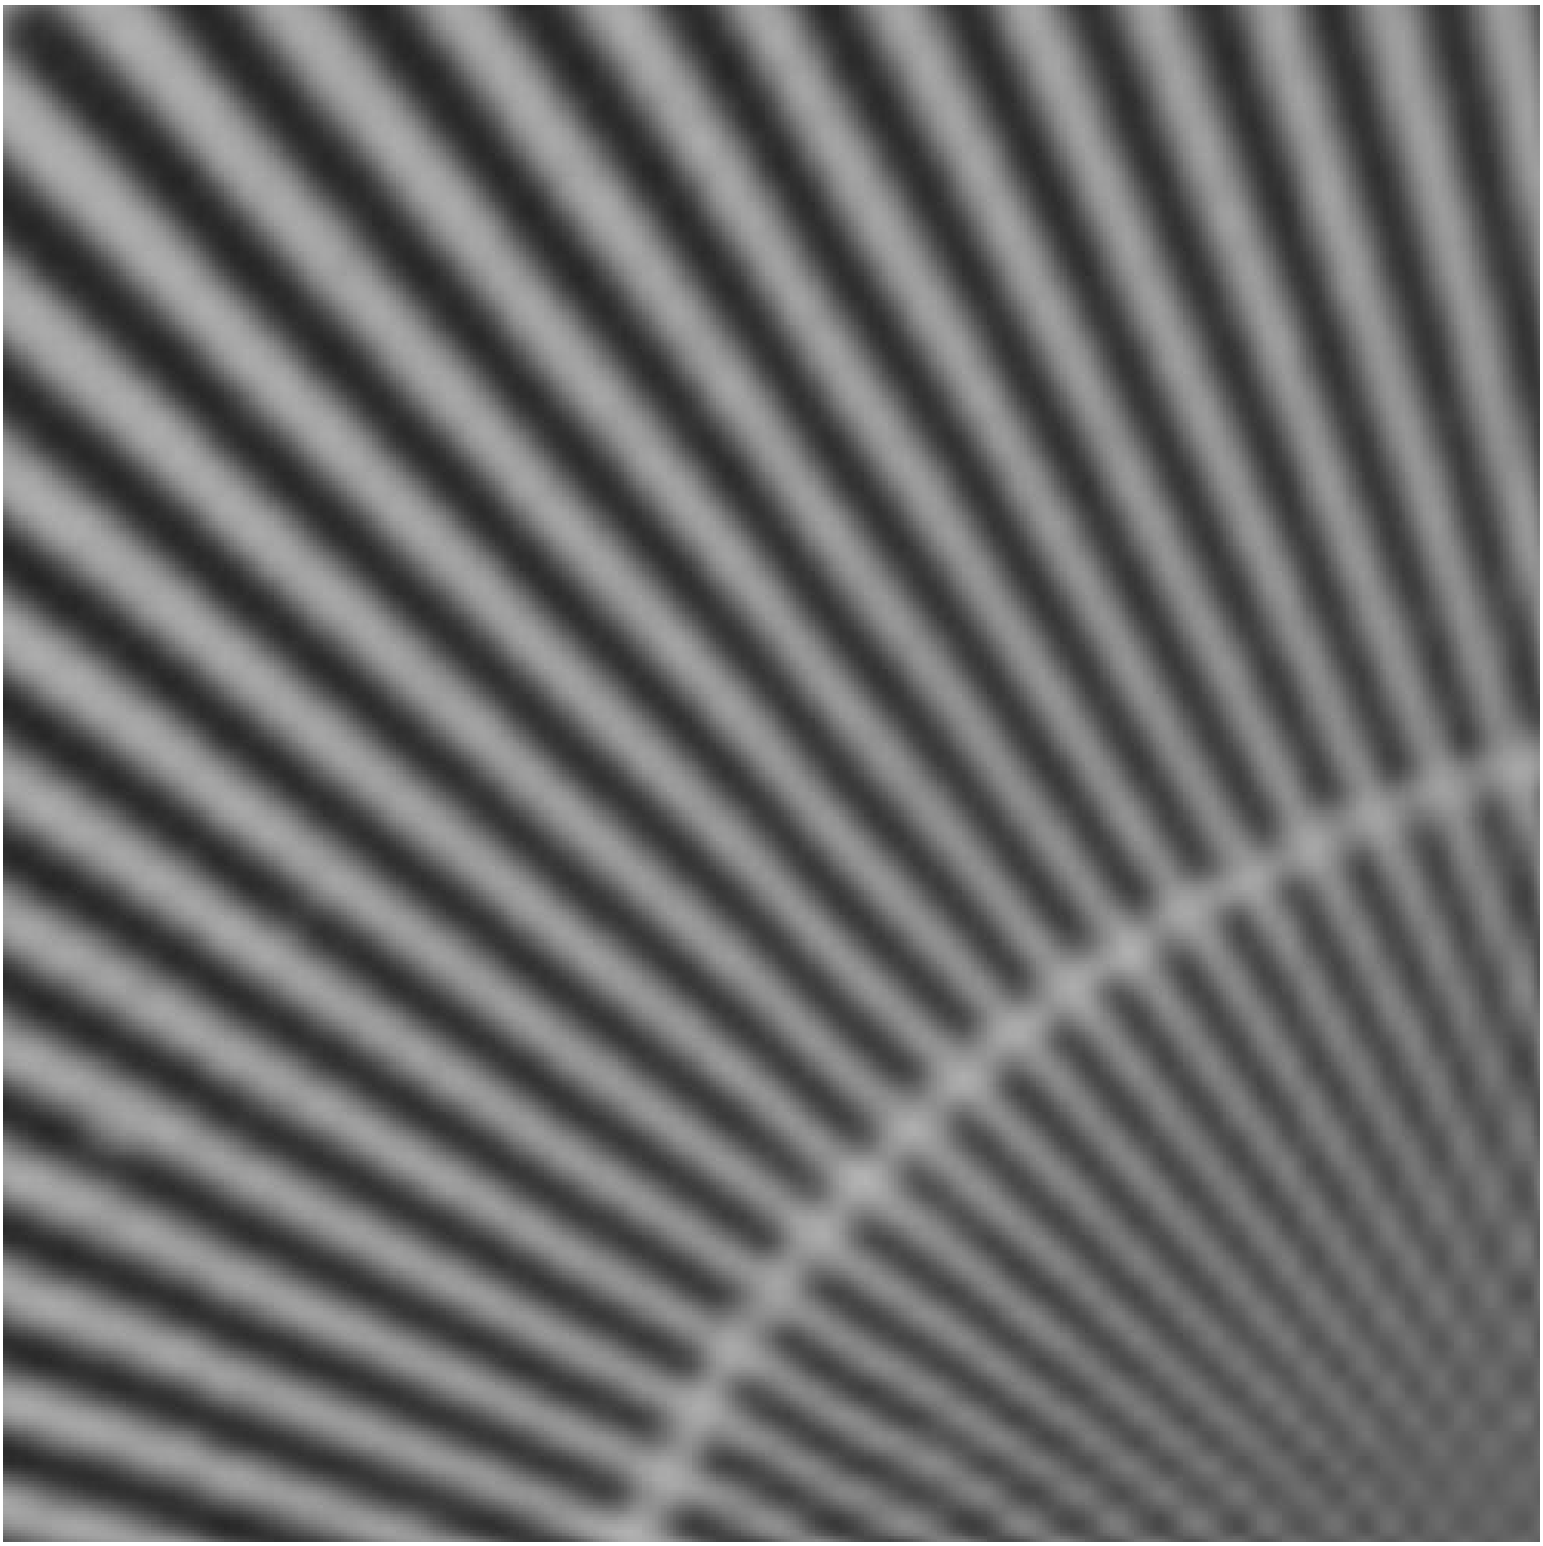

Figure 13: Right part of figure 5 enlarged. Note that the final versions would be printed at the quality of the first figure, and this page would not appear.

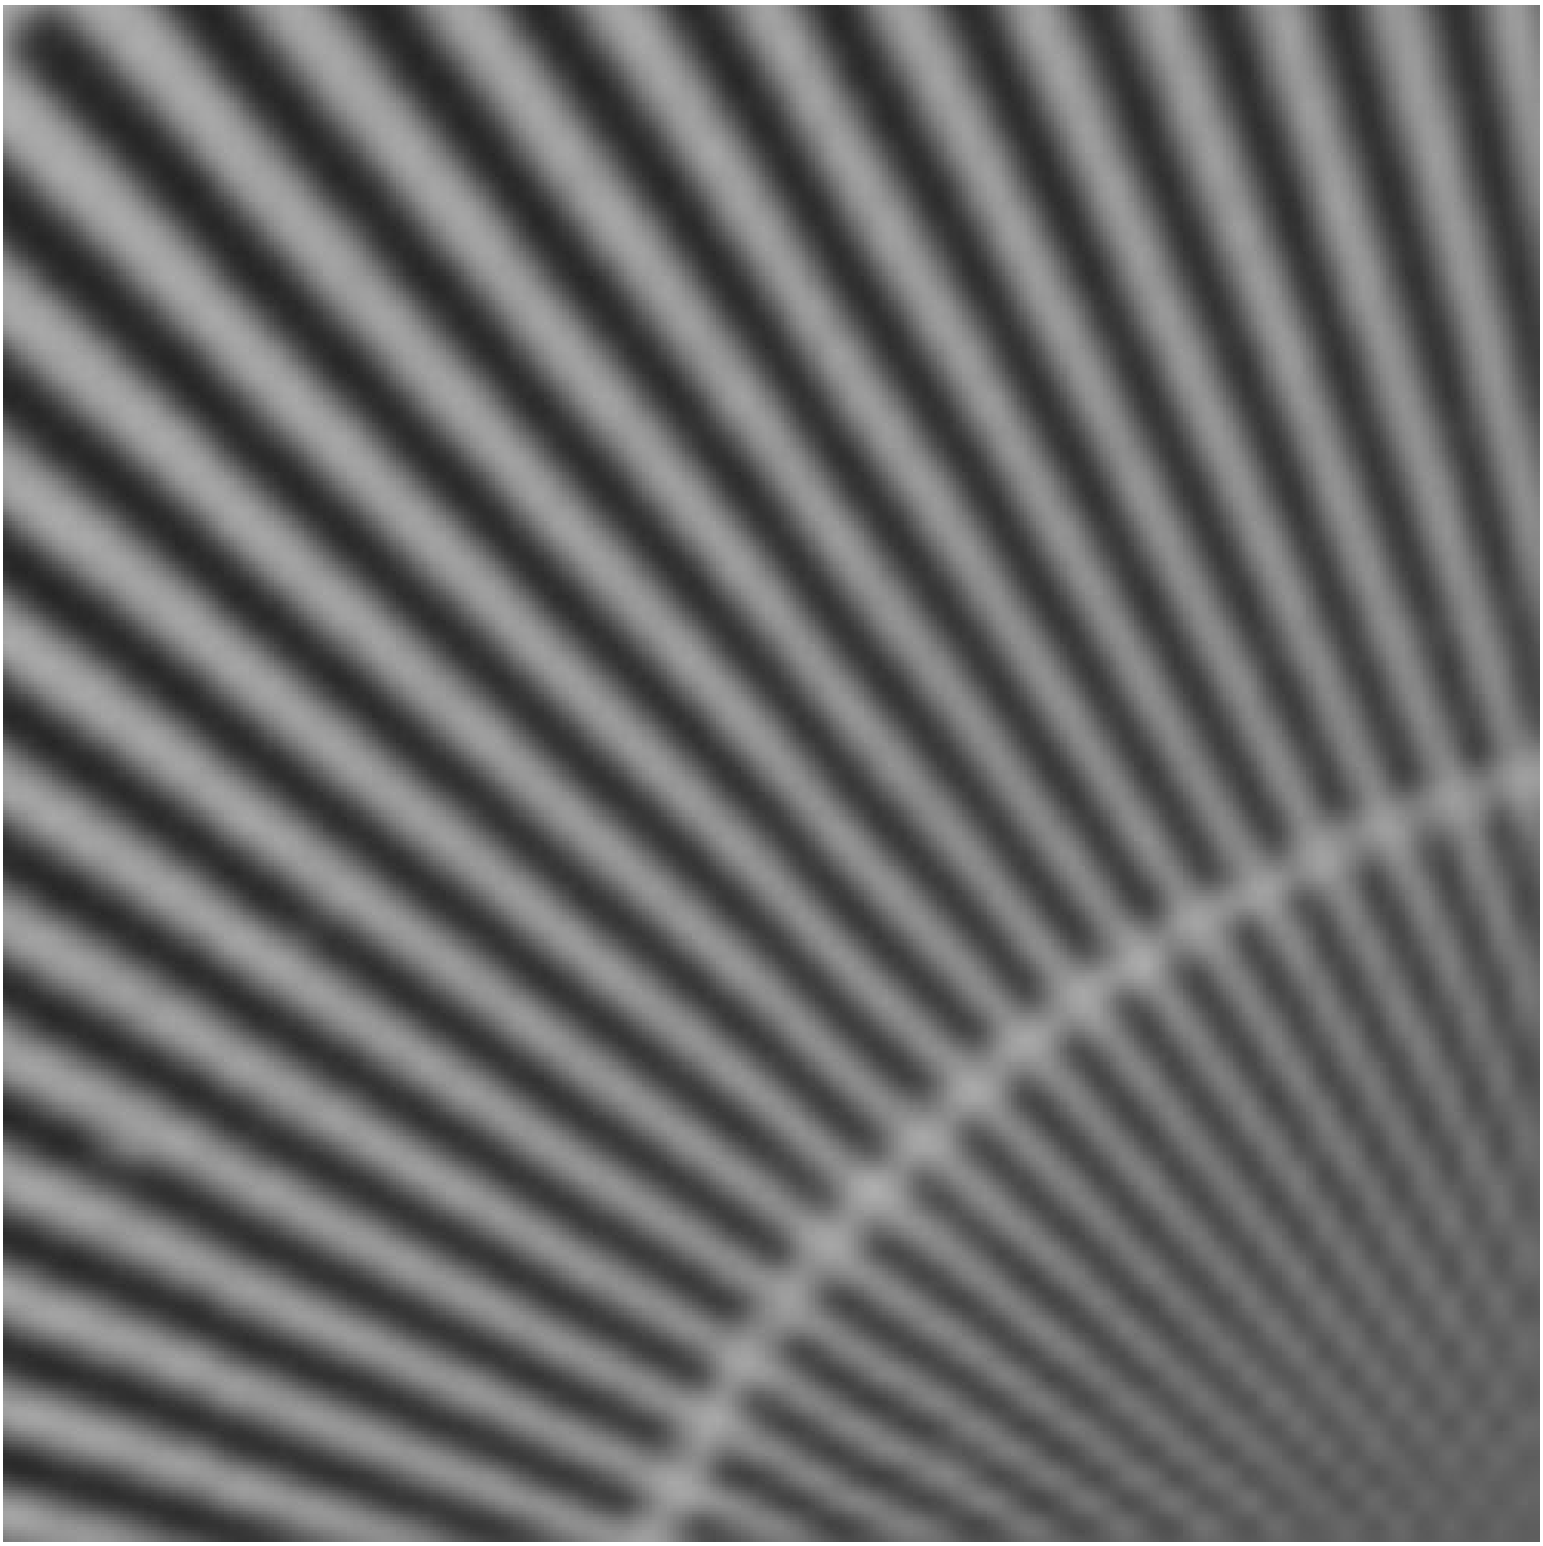

Figure 14: Left part of figure6 enlarged. Note that the final versions would be printed at the quality of the first figure, and this page would not appear.

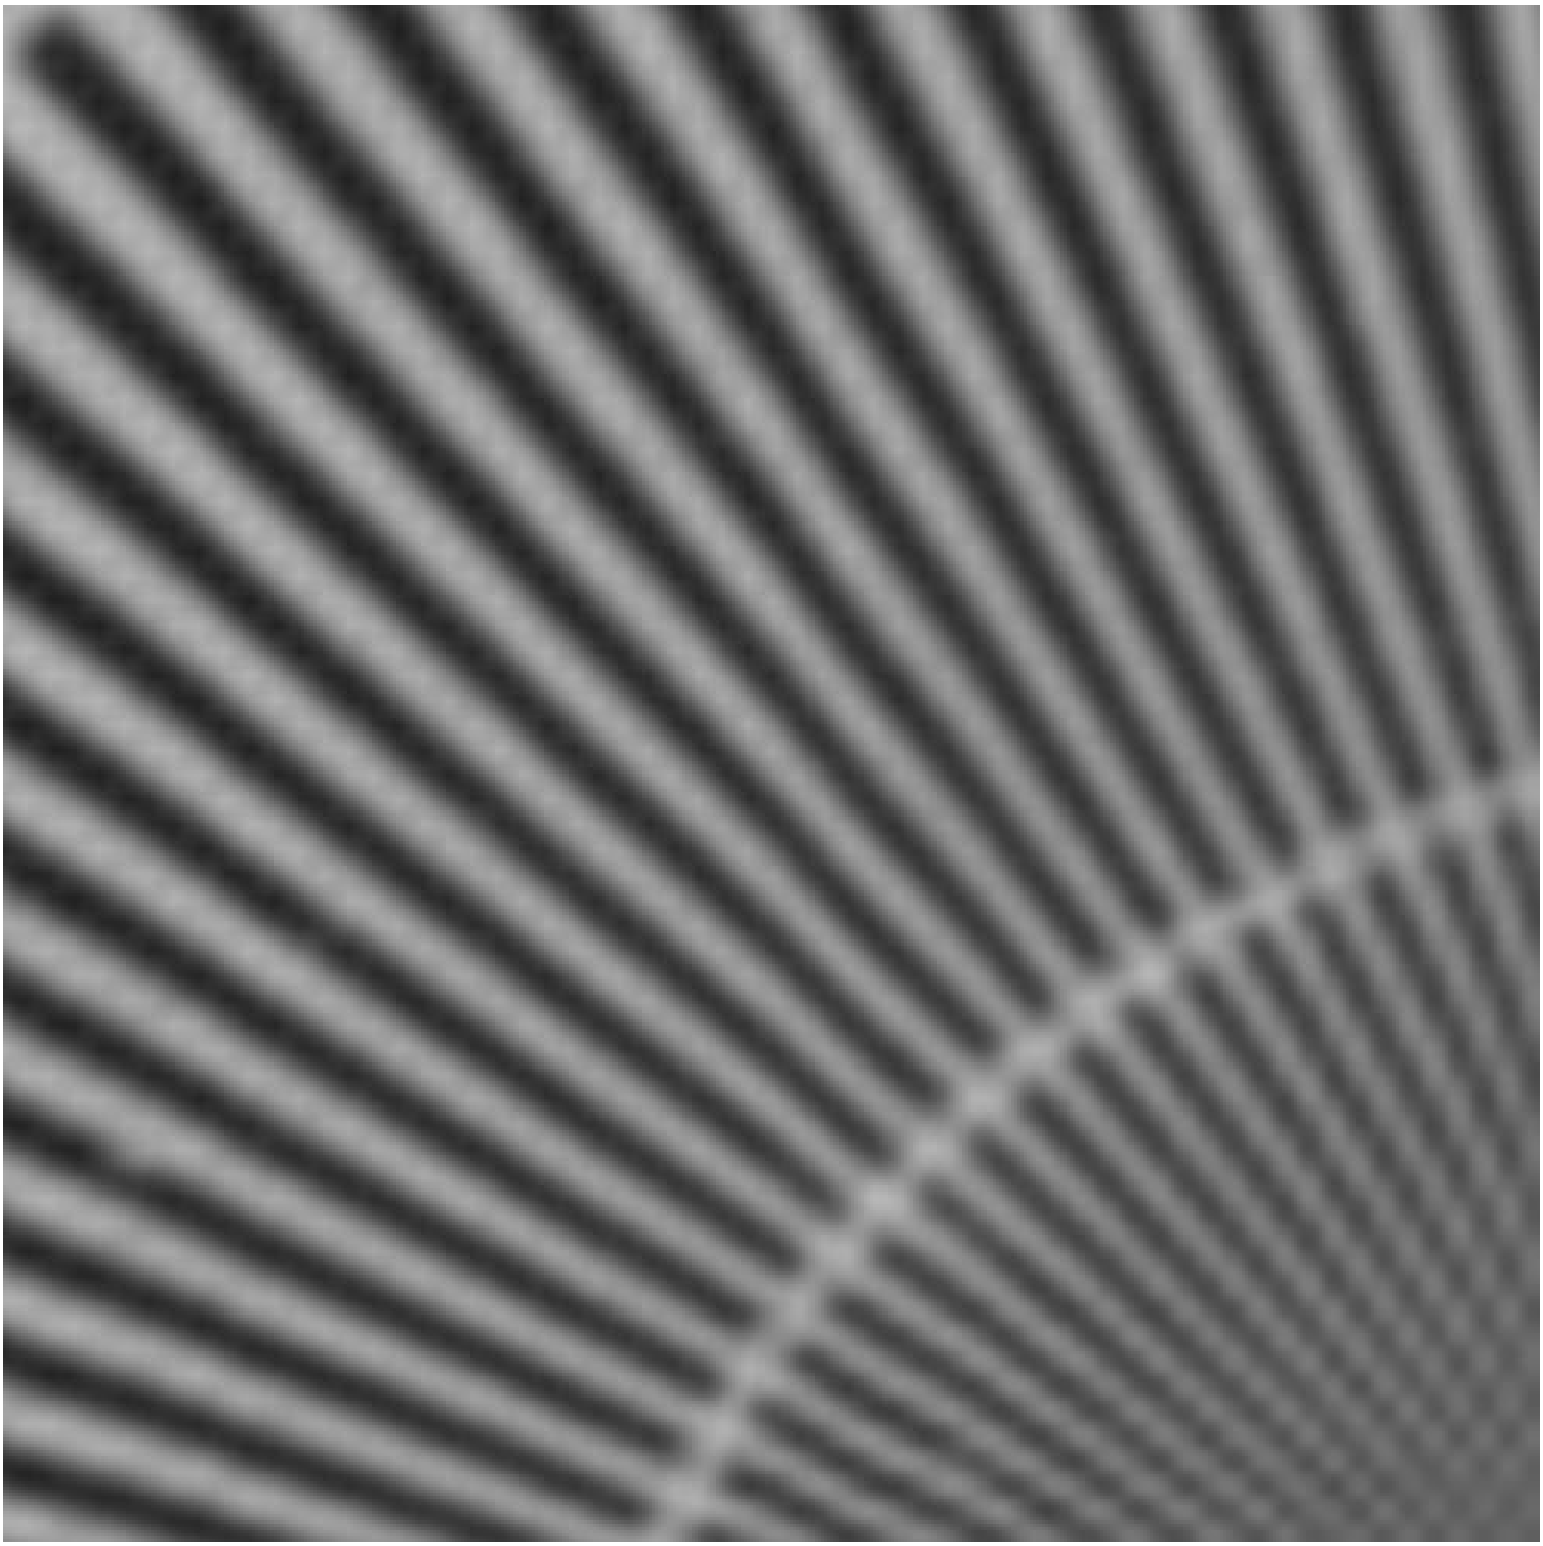

Figure 15: Right part of figure 6 enlarged. Note that the final versions would be printed at the quality of the first figure, and this page would not appear.

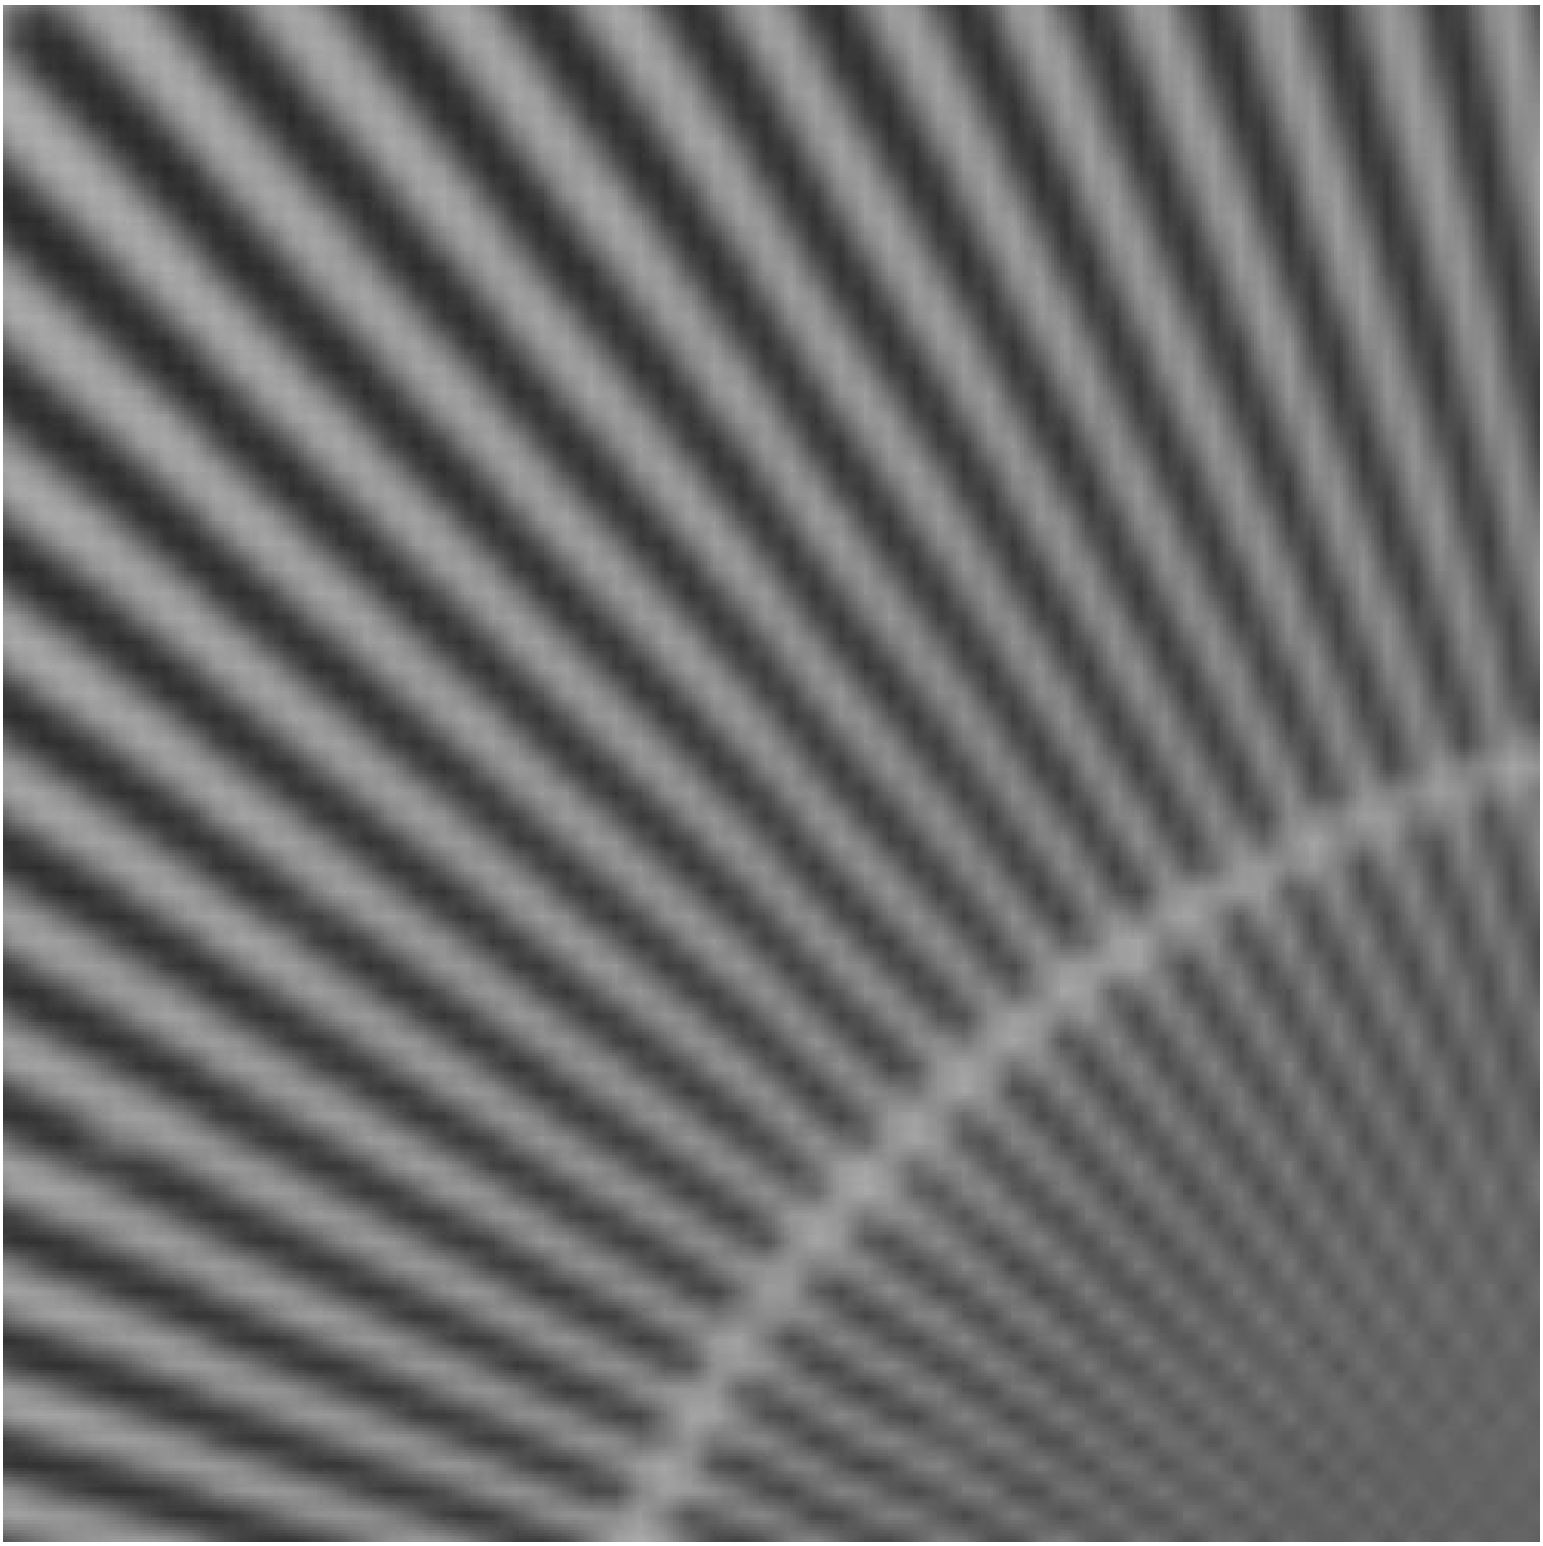

Figure 16: Left part of figure7 enlarged. Note that the final versions would be printed at the quality of the first figure, and this page would not appear.

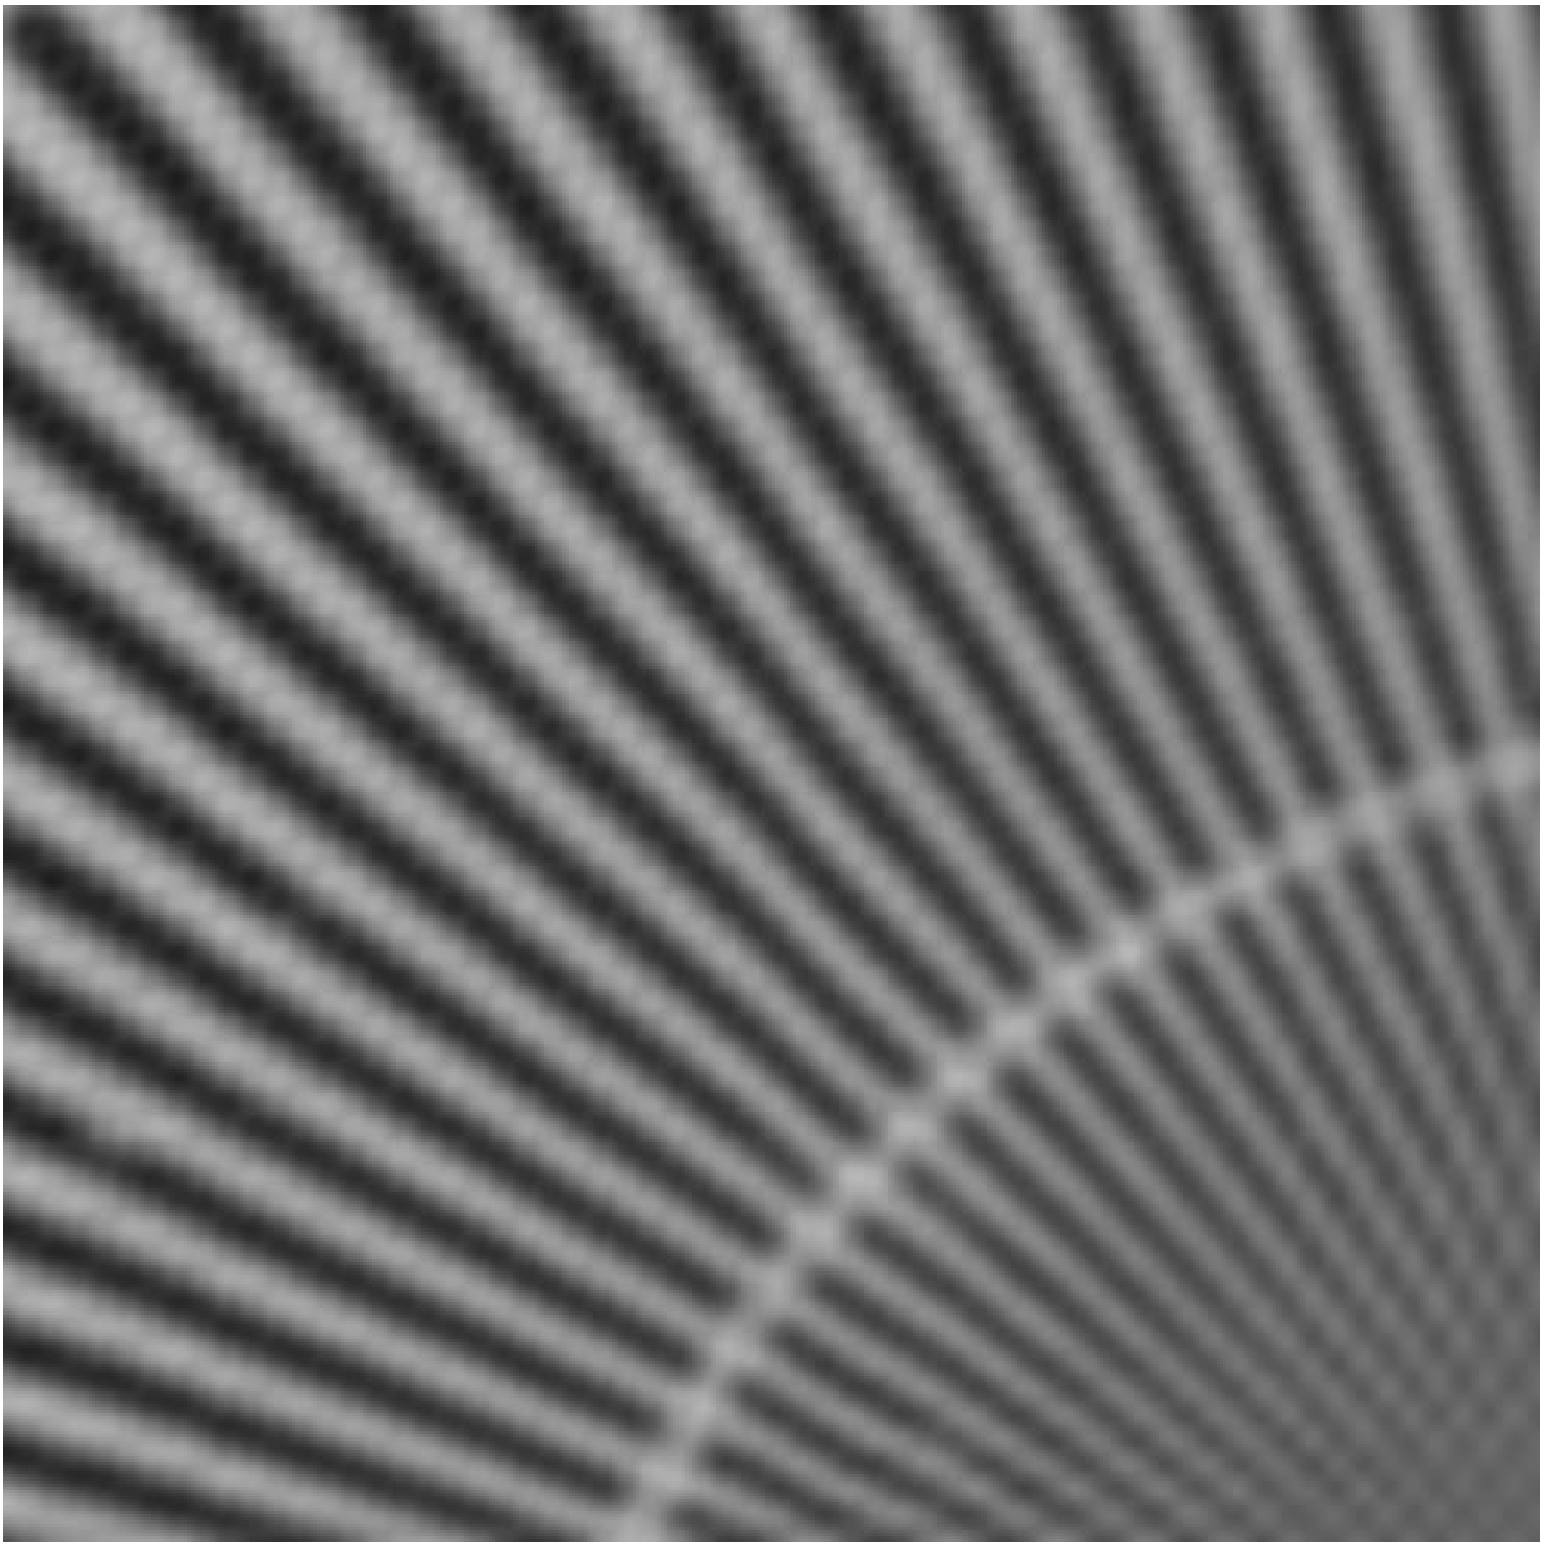

Figure 17: Right part of figure 7 enlarged. Note that the final versions would be printed at the quality of the first figure, and this page would not appear.

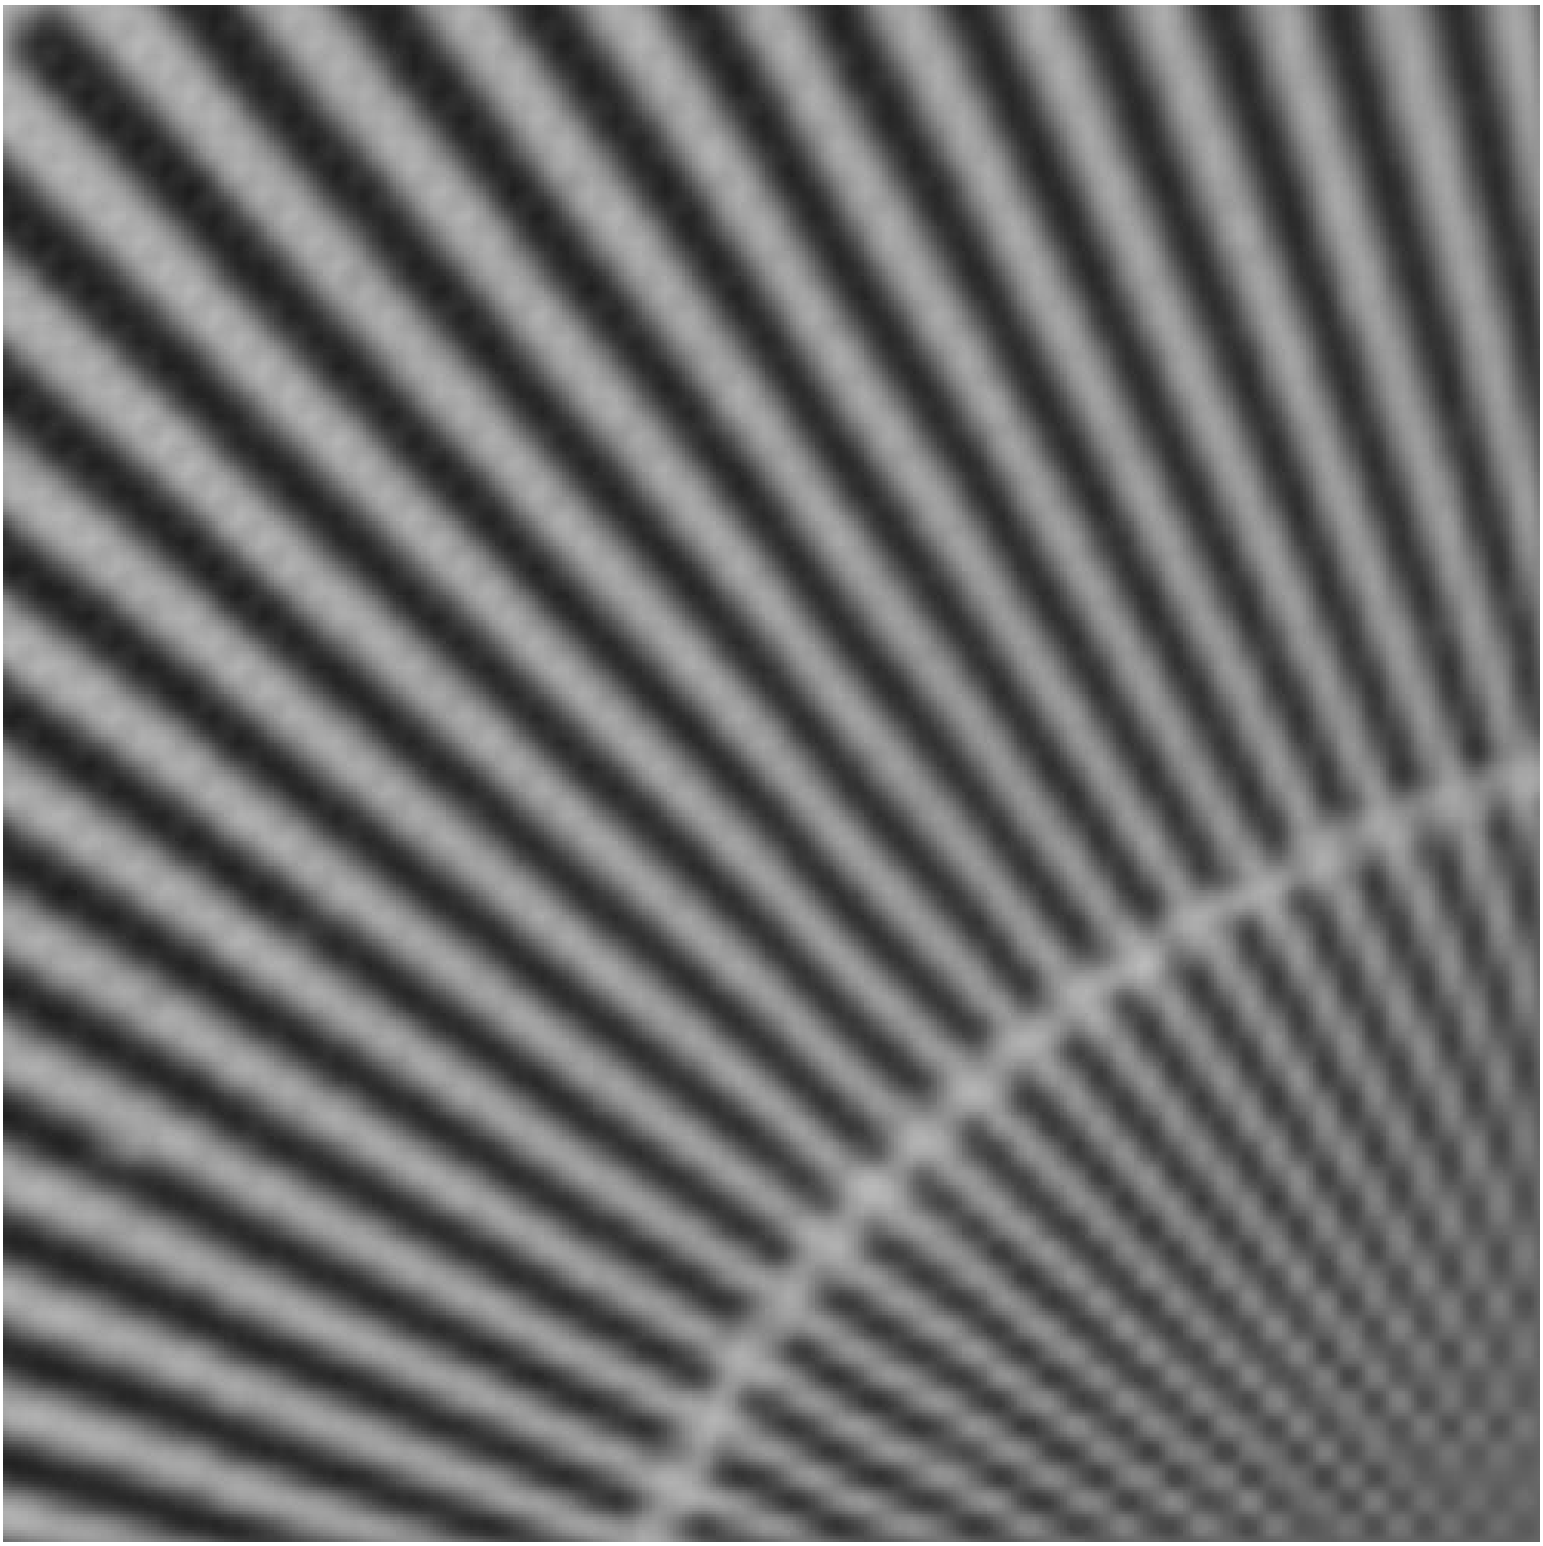

Figure 18: Left part of figure8 enlarged. Note that the final versions would be printed at the quality of the first figure, and this page would not appear.

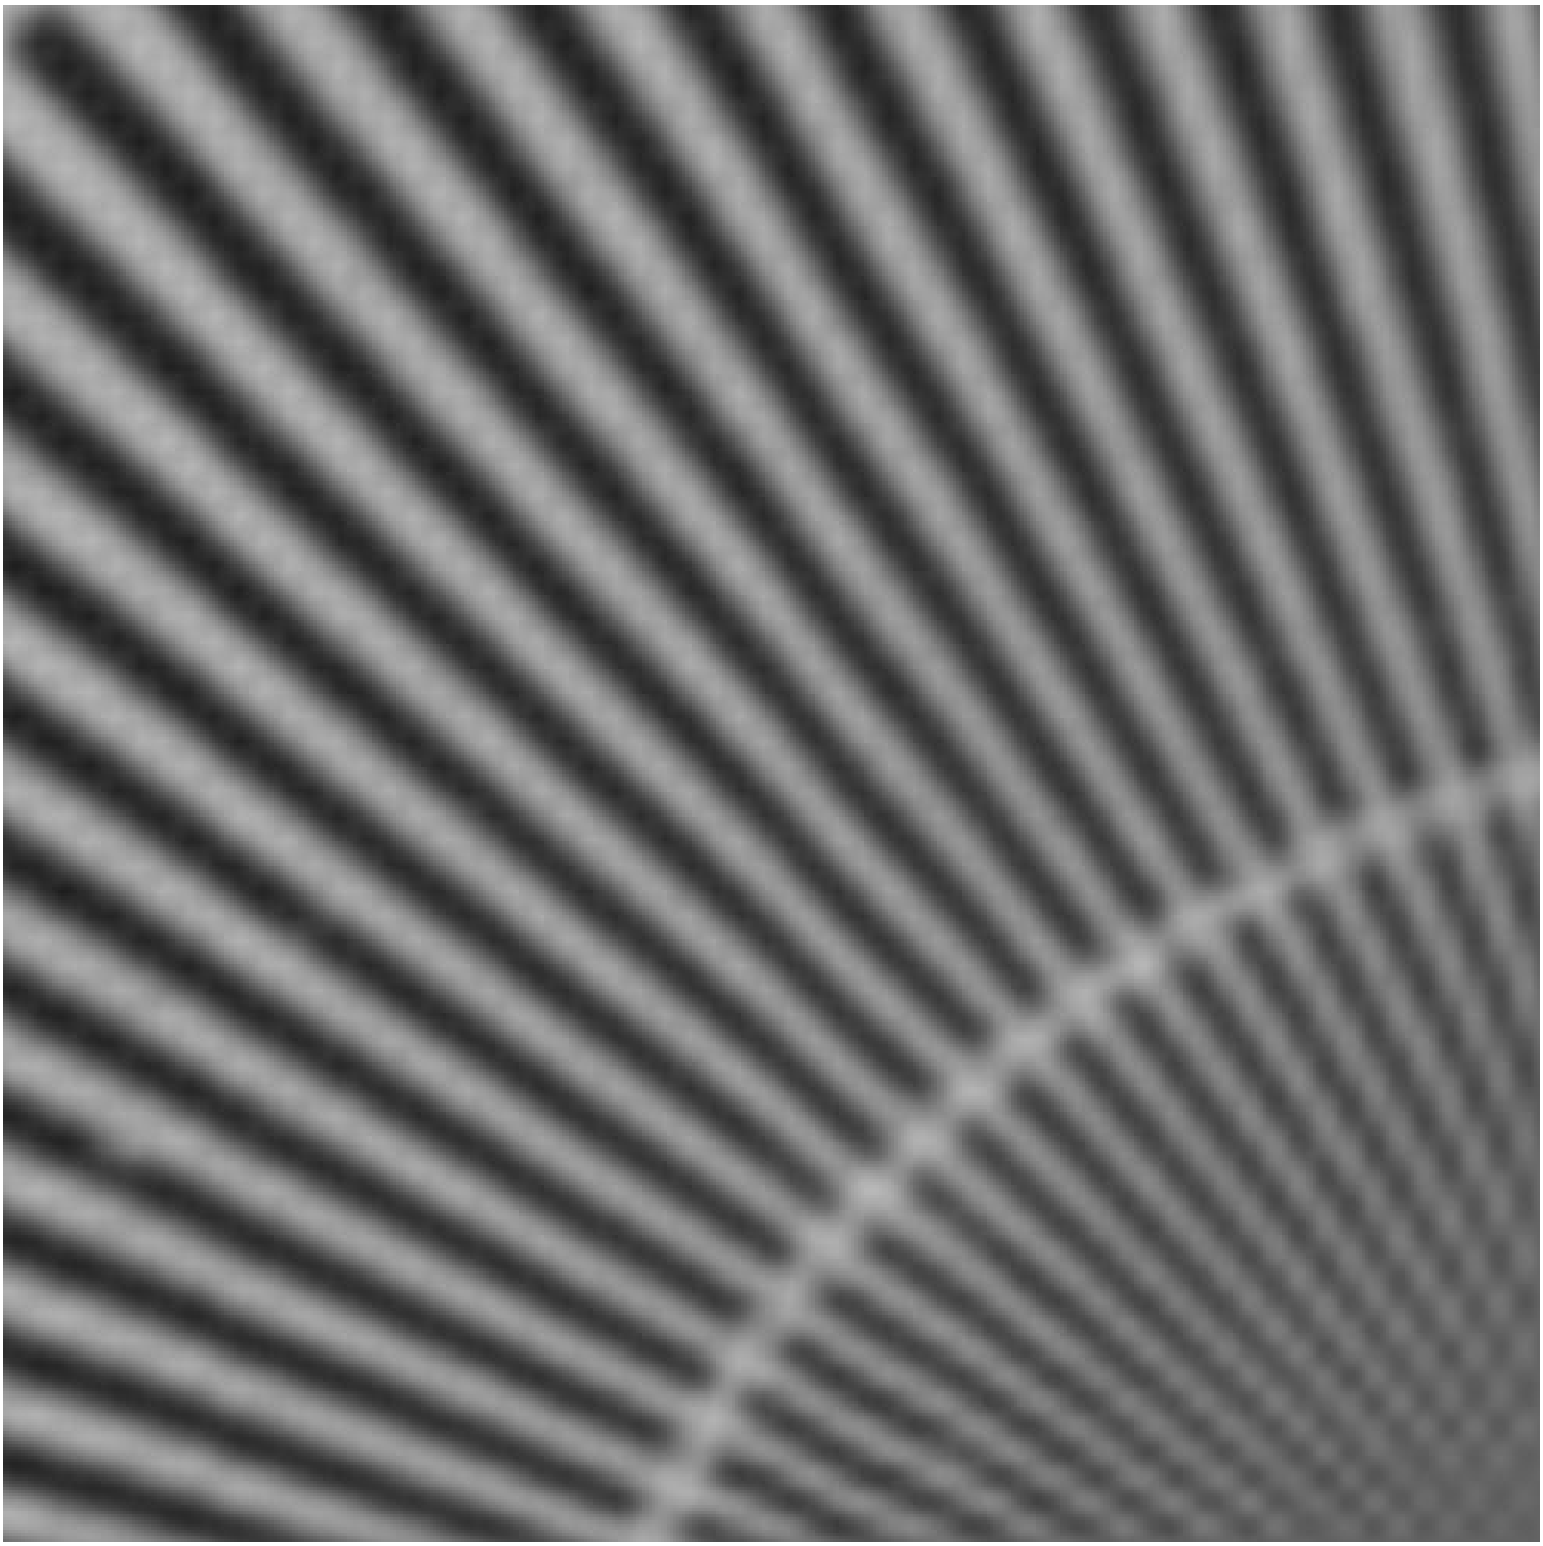

Figure 19: Right part of figure 8 enlarged. Note that the final versions would be printed at the quality of the first figure, and this page would not appear.

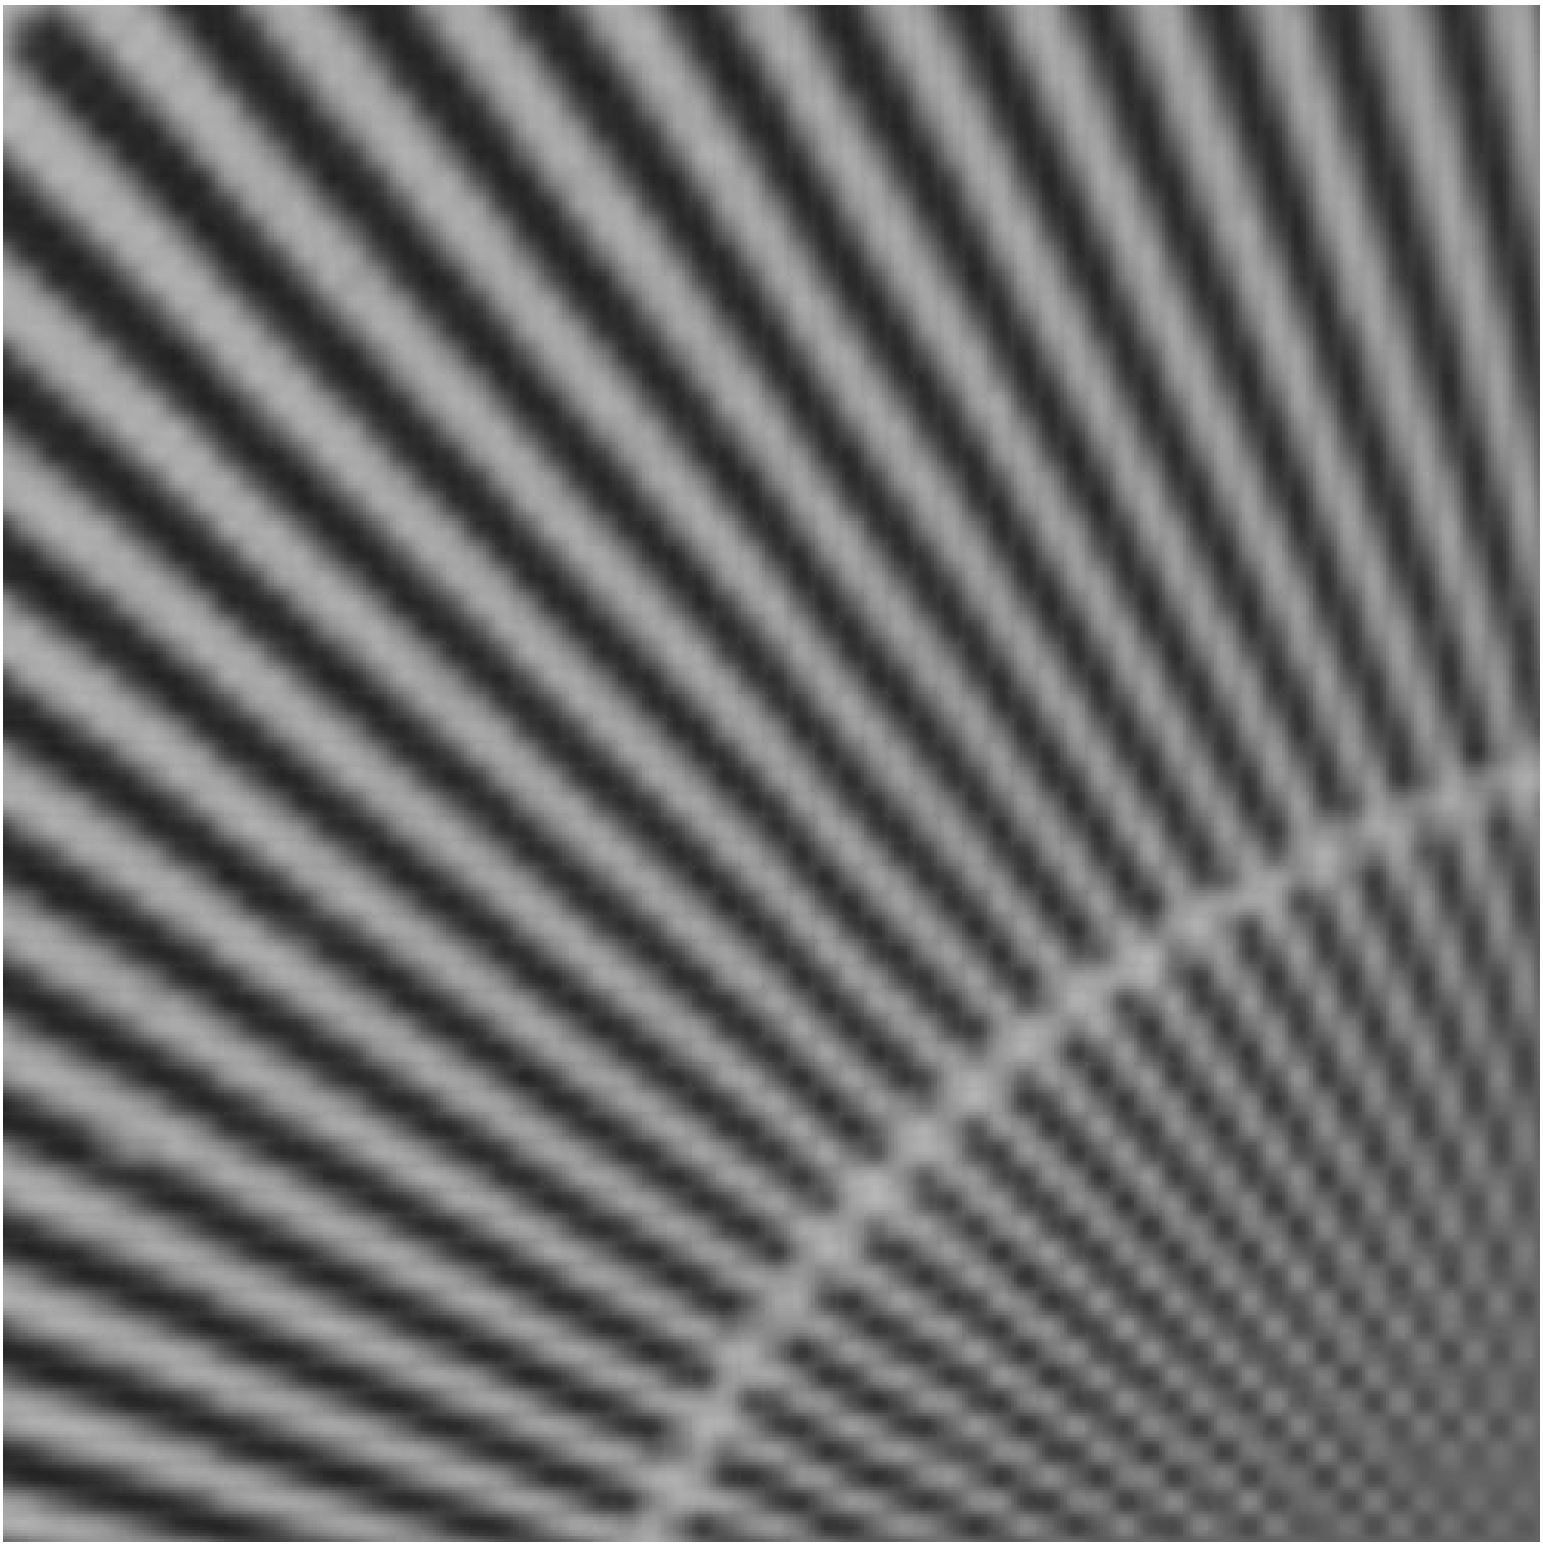

Figure 20: Left part of figure9 enlarged. Note that the final versions would be printed at the quality of the first figure, and this page would not appear.

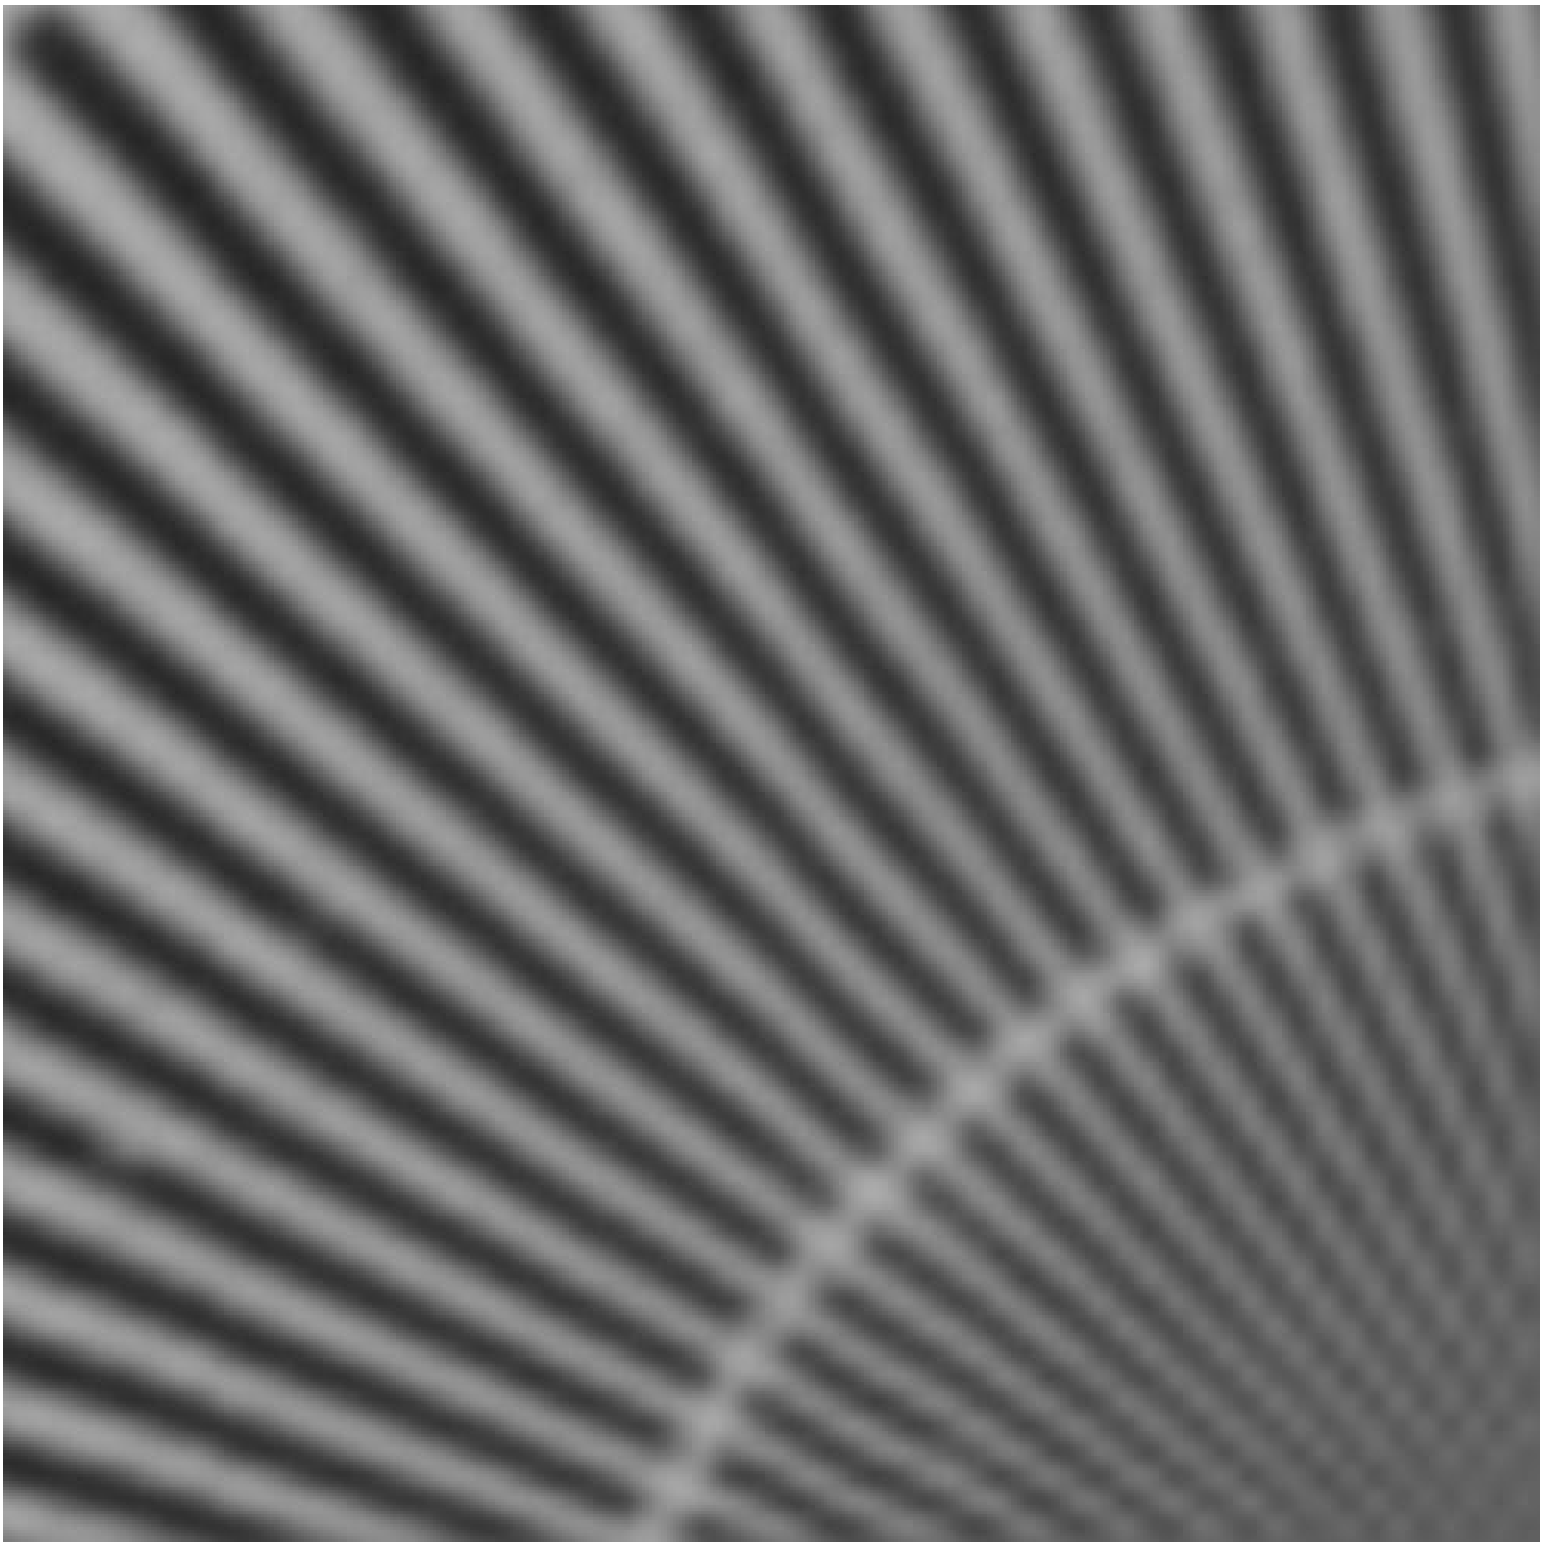

Figure 21: Right part of figure 9 enlarged. Note that the final versions would be printed at the quality of the first figure, and this page would not appear.

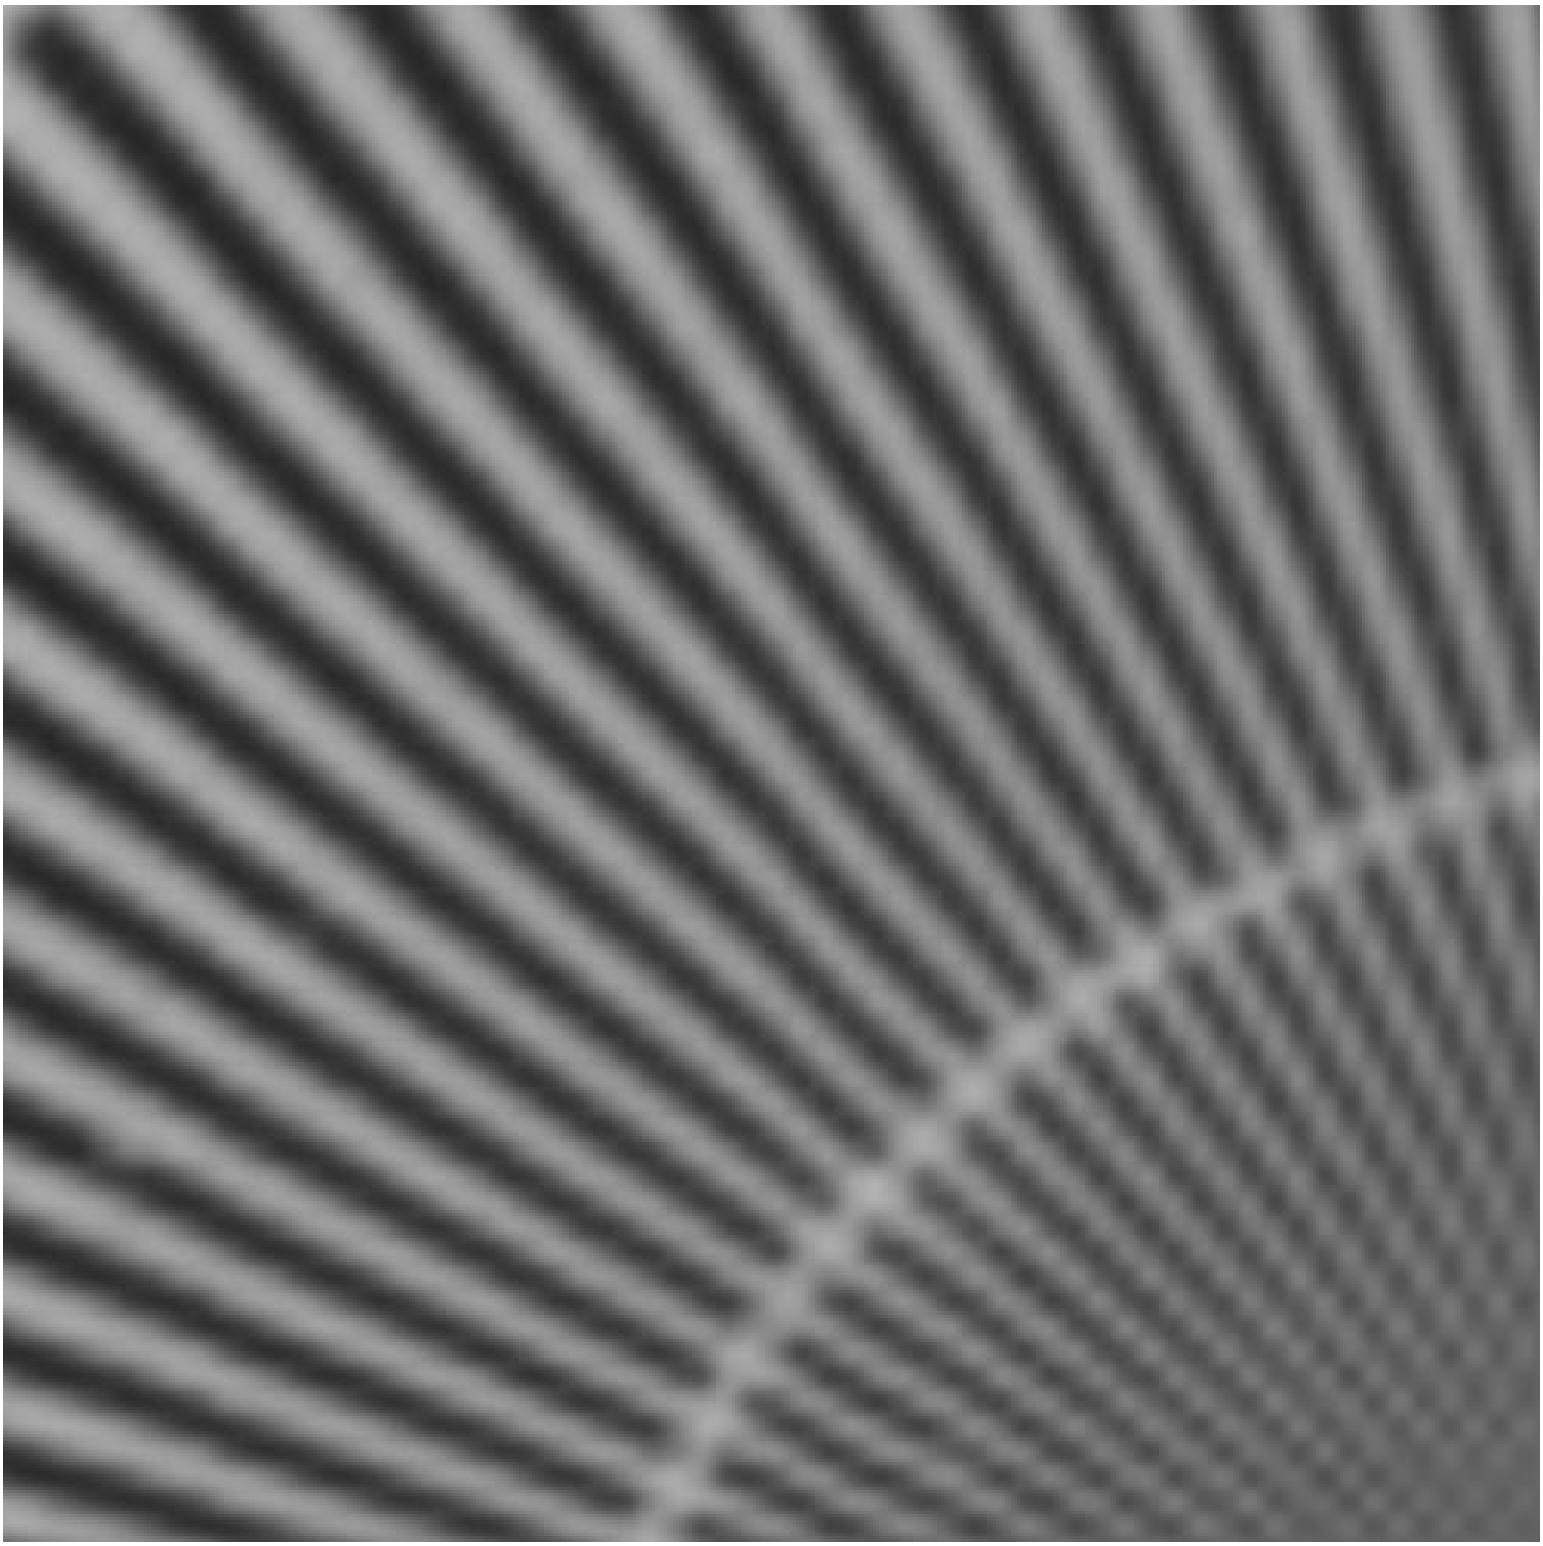

Figure 22: Left part of figure10 enlarged. Note that the final versions would be printed at the quality of the first figure, and this page would not appear.

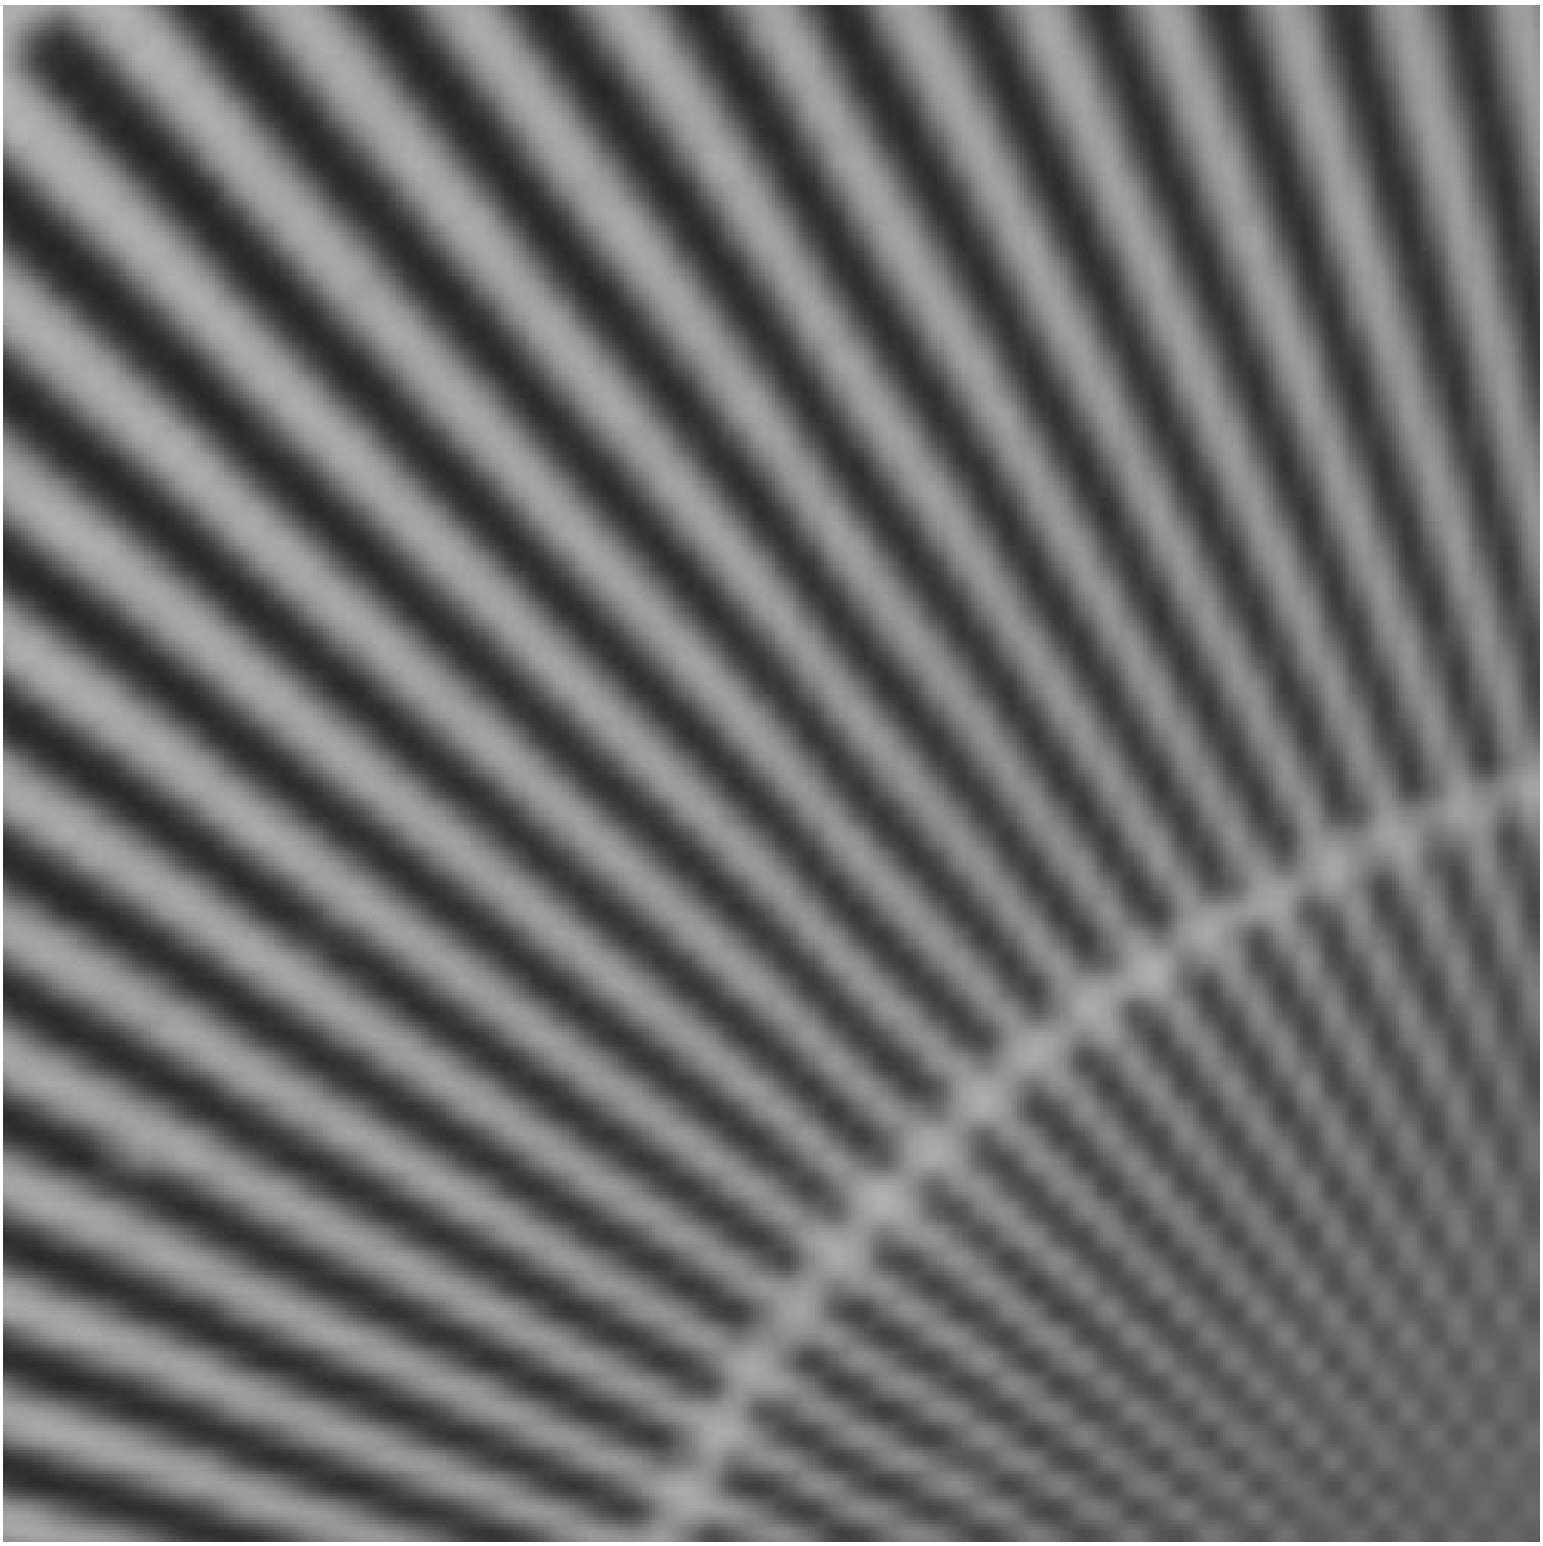

Figure 23: Right part of figure 10 enlarged. Note that the final versions would be printed at the quality of the first figure, and this page would not appear.

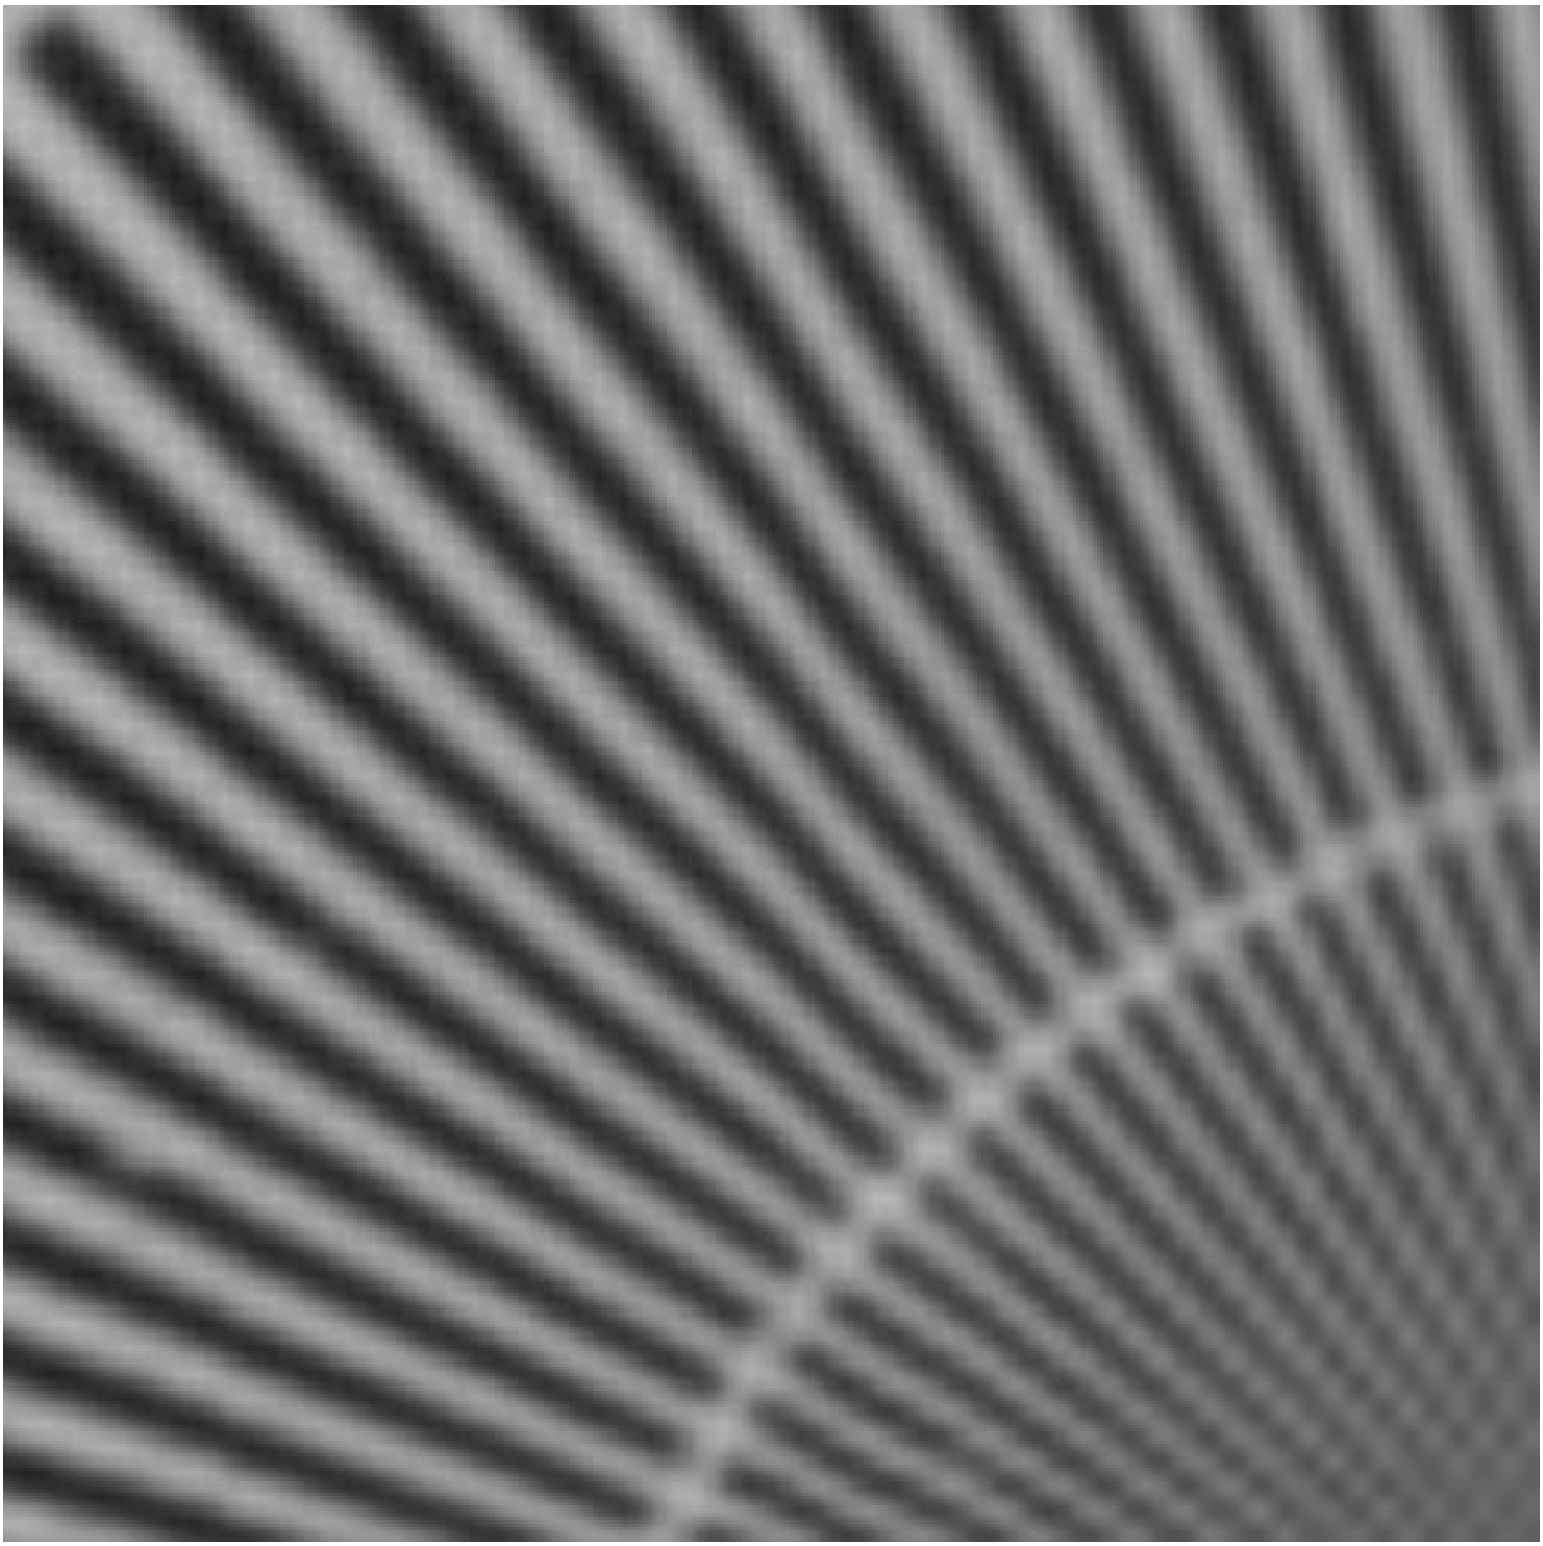

Figure 24: Left part of figure11 enlarged. Note that the final versions would be printed at the quality of the first figure, and this page would not appear.

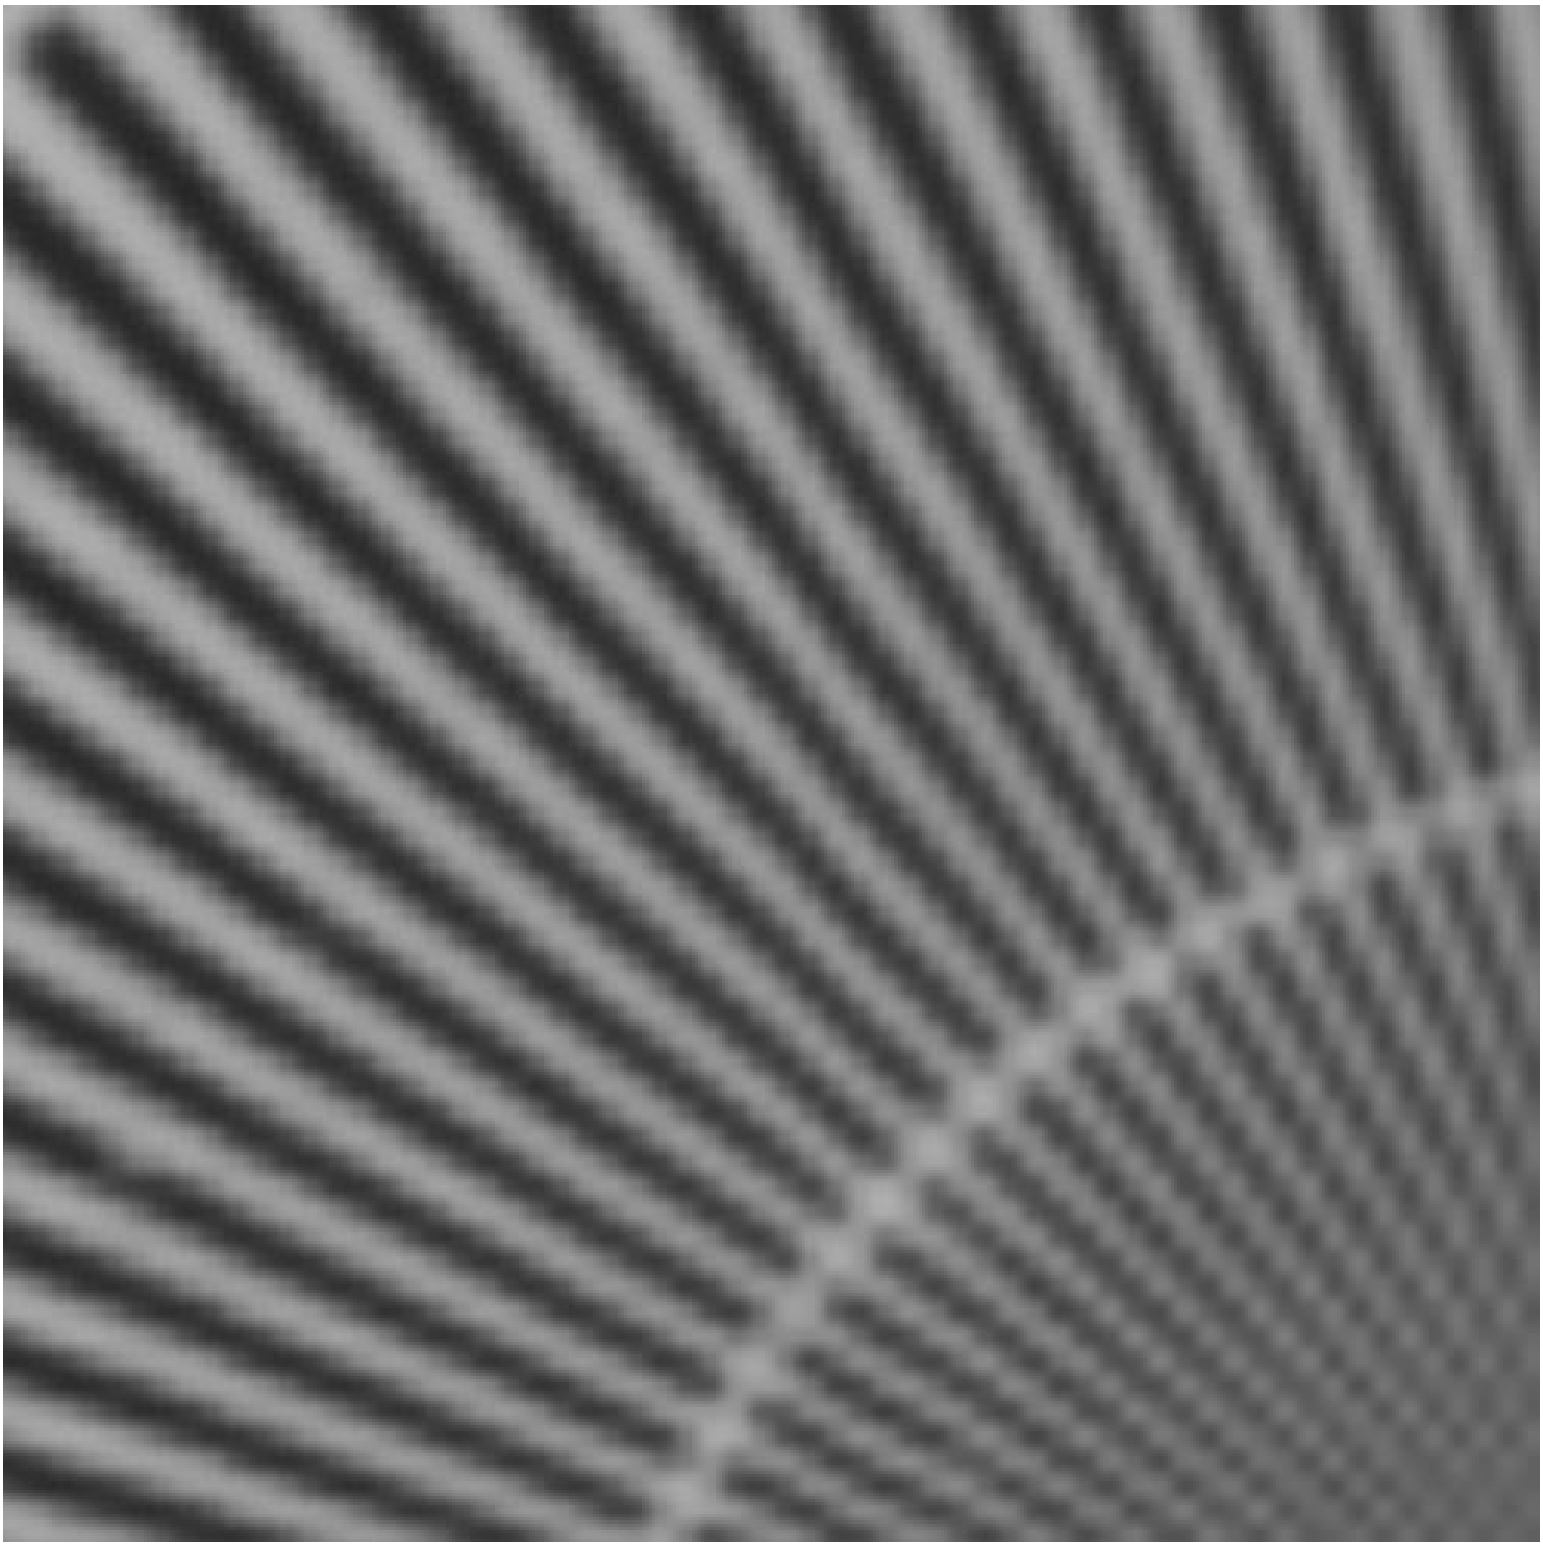

Figure 25: Right part of figure 11 enlarged. Note that the final versions would be printed at the quality of the first figure, and this page would not appear.