



IRWIN AND JOAN JACOBS
CENTER FOR COMMUNICATION AND INFORMATION TECHNOLOGIES

A Hybrid Vector Wiener Filter Approach to Translational Super-Resolution

**Tomer Michaeli and
Yonina C. Eldar**

CCIT Report # 756
February 2010

■ ■ ■ ■ Electronics
■ ■ ■ ■ Computers
■ ■ ■ ■ Communications

DEPARTMENT OF ELECTRICAL ENGINEERING
TECHNION - ISRAEL INSTITUTE OF TECHNOLOGY, HAIFA 32000, ISRAEL



A Hybrid Vector Wiener Filter Approach to Translational Super-Resolution

Tomer Michaeli, *Student Member, IEEE* and Yonina C. Eldar, *Senior Member, IEEE*

Abstract—We address the problem of purely-translational super-resolution (SR) for signals in arbitrary dimensions. We show that discretization, a key step in many SR algorithms, inevitably leads to inaccurate modeling. Instead, we treat the problem entirely in the continuous domain by modeling the signal as a continuous-space random process and deriving its linear minimum mean-squared error (LMMSE) estimate given the low-resolution discrete-space observations. We derive a closed form expression for the resulting mean-squared error and use it to analyze the emergence of periodic artifacts in the super-resolved signal. We also provide three efficient implementation schemes of the LMMSE estimate, one of which specialized for 1D applications. These methods constitute a natural generalization of several well known single-image recovery algorithms, such as spline interpolation, to the multichannel SR setting. Experiments on real-world images demonstrate the advantage of our approach with respect to several prominent SR techniques that rely on discretization.

Index Terms—Super-resolution, nonuniform interpolation, hybrid Wiener filter.

I. INTRODUCTION

SUPER-RESOLUTION refers to the process of combining several low-resolution descriptions of a signal to form one higher resolution version of it. In the field of image processing, there are 1D, 2D and 3D variants of this task. Perhaps the most commonly treated scenario is that of purely-translational spatial (2D) super-resolution (SR) [1], [2]. Here, several low-resolution noisy images of a scene are captured by a camera, each with a different translation, and the goal is to produce one high-resolution image of the same scene. Three-dimensional scenarios arise in space-time SR applications [3]. There, several video sequences of the same scene are fused into one higher-resolution video stream at a higher frame-rate. If the video cameras are spatially co-calibrated, then processing can be carried out only along the time dimension (temporal SR), rendering the problem one-dimensional. In this paper, we collectively refer to all problems of this type as SR, and develop a theory for arbitrary d -dimensional signals.

The physical model underlying purely-translational SR scenarios can be described mathematically as [4]

$$c_k[\mathbf{n}] = [(s * x)(\mathbf{t})]_{\mathbf{t}=\mathbf{n}-\mathbf{t}_k} + u_k[\mathbf{n}], \quad k = 1, \dots, K, \quad (1)$$

This work was supported in part by the Israel Science Foundation under Grant no. 1081/07 and by the European Commission in the framework of the FP7 Network of Excellence in Wireless COMmunications NEWCOM++ (contract no. 216715).

The authors are with the department of electrical engineering, Technion-Israel Institute of Technology, Haifa, Israel (phone: +972-4-8294682, fax: +972-4-8295757, e-mail: tomermic@tx.technion.ac.il, yonina@ee.technion.ac.il).

where $c_k[\mathbf{n}]$, $\mathbf{n} \in \mathbb{Z}^d$, is the k th discrete-space observation of the continuous-space signal $x(\mathbf{t})$, $\mathbf{t} \in \mathbb{R}^d$, acquired with translation \mathbf{t}_k and additive noise $u_k[\mathbf{n}]$. The filter $s(\mathbf{t})$ is associated with the imaging device. It corresponds to the point-spread-function (PSF) of the lens in 2D SR, to the temporal integration profile of the sensor in 1D SR, and to a combination of both in 3D space-time SR. The goal of an SR algorithm is to produce a discrete-space high-resolution signal $\mathbf{x}_{\text{HR}}[\mathbf{n}]$ which corresponds to the samples of $x(\mathbf{t})$ on a dense grid.

In much of the recent literature, SR is modeled via the discrete-space relations [1], [2]:

$$\mathbf{c}_k = \mathbf{S}_k \mathbf{x}_{\text{HR}} + \mathbf{u}_k, \quad k = 1, \dots, K. \quad (2)$$

Here \mathbf{c}_k , \mathbf{u}_k , and \mathbf{x}_{HR} are column vectors comprising the elements of $c_k[\mathbf{n}]$, $u_k[\mathbf{n}]$, and $\mathbf{x}_{\text{HR}}[\mathbf{n}]$ respectively, and the matrix \mathbf{S}_k accounts for the filtering by $s(\mathbf{t})$ and the sampling. This discrete formulation has several advantages over the continuous model (1). Most noteworthy, it allows the construction of a finite-dimensional optimization problem incorporating complicated realistic prior-knowledge assumptions on the unknown \mathbf{x}_{HR} [2]. Such optimization problems may be solved using standard optimization methods. Nevertheless, as we show in this paper, there are many situations in which the continuous-space equations (1) cannot be represented in the discrete form (2). In other words, we show that in certain settings SR must be regarded as a continuous-space interpolation problem. Algorithms which do not treat the problem from this viewpoint, inevitably suffer from a model miss-match error.

Prior SR work falling into the interpolation category (also termed frequency-domain methods [4]) include [5], [6], [7], [8], [9], [10], [11], [12]. In these algorithms, the continuous-space scene $x(\mathbf{t})$ is assumed to be bandlimited. In the single-channel interpolation literature, recent work has shown that this type of prior knowledge is often not in agreement with the typical behavior of natural images [13], [14]. Furthermore, it has been experimentally demonstrated that interpolation algorithms relying on non-bandlimited models, such as the Matérn prior, can lead to improved reconstruction results [15], [16], [17]. In SR applications, the bandlimited assumption translates to a stringent limitation on the required number of measurements. Specifically, the number of observations K required for increasing the resolution by a factor of Δ must be at least Δ^d in the methods mentioned above. Thus, to increase resolution by a factor of 4 in every dimension, for example, at least 16 still images are required in a spatial SR task ($d = 2$), and at least 64 video sequences are needed in a space-time SR scenario ($d = 3$).

An additional disadvantage of existing continuous-domain SR methods is that they do not explicitly incorporate statistical assumptions on $x(\mathbf{t})$. Such assumptions have been shown to greatly improve SR performance when using the discrete model (2) [2].

In [18], the authors proposed an SR approach, which can (in theory) be adapted to work in the continuous-domain, and which does not rely on the bandlimited assumption. There, prior knowledge on $x(\mathbf{t})$ is implicitly defined by a user-chosen “back-projection kernel”. However, the effect of this kernel on the reconstruction is not clear. Moreover, by adapting this algorithm to the continuous-domain, we end up with the need to iteratively update the continuous-space recovery $\hat{x}(\mathbf{t})$. This, of course, cannot be done in a precise manner. In practice, only the densely sampled version $\hat{x}_{\text{HR}}[\mathbf{n}]$ is maintained throughout the iterations. Therefore, practically, this method belongs to the category of algorithms relying on the incorrect discrete formulation (2).

In this paper we treat the SR interpolation problem without resorting to the commonly used bandlimited assumption. Motivated by the good single-channel image interpolation results reported in [13], [14], [15], [16], [17], here we model $x(\mathbf{t})$ as a wide-sense-stationary (WSS) continuous-space random signal with known power-spectral-density (PSD). We then derive the linear minimum mean-squared error (MSE) estimator of $x(\mathbf{t})$ given the discrete-space observations $c_1[\mathbf{n}], \dots, c_K[\mathbf{n}]$. We term this estimator the *vector hybrid Wiener filter* as it operates on a discrete-space vector process $\mathbf{c}[\mathbf{n}] = (c_1[\mathbf{n}], \dots, c_K[\mathbf{n}])^T$, and outputs a continuous-space recovery $\hat{x}(\mathbf{t})$.

We show that in situations where both the PSF and the signal’s autocorrelation function are compactly supported, the resulting estimator can be implemented efficiently. Specifically, for 1D SR, we provide an implementation which consists of digital filtering of the discrete-space measurements $c_k[\mathbf{n}]$, followed by a simple interpolation stage. For SR in higher dimensions (e.g., spatial or space-time SR), we propose two algorithms for efficient implementation of the resulting estimator. Both approaches include an iterative stage which operates on the measurements $c_1[\mathbf{n}], \dots, c_K[\mathbf{n}]$, followed by an interpolation step. An important feature of our iterative approaches, is that iterations are performed in the digital domain so that the implementation stays loyal to the continuous formulation (1). In other words, our approach does not rely on approximations of the type required to adapt [18] to work in the continuous-domain.

When the autocorrelation function of the scene is chosen to be a B-spline, our approach constitutes a natural generalization of the widely used (single-channel) spline interpolation method to the multichannel SR setting. We term the resulting algorithm *spline super resolution*. As we demonstrate through experiments on real-world images, our multichannel spline interpolation method often yields recoveries, which are better than those produced by current state-of-the-art algorithms.

The paper is organized as follows. In Section II, we show mathematically why discretization is generally not loyal to the true physical setting. In Section III, we present our assumed continuous-space statistical model and derive the

corresponding LMMSE estimator. Three efficient implementation schemes for the LMMSE estimator are provided in Sections IV and V. In Section VI, we analyze the reconstruction error patterns and discuss the inherent tradeoff between a sharp recovery and periodic artifacts. To obtain a practical SR algorithm, we specialize our approach in Section VII to spline autocorrelation models, which leads to the spline super resolution algorithm. Finally, in Section VIII we present several experiments on real-world images, demonstrating the advantage of our approach over SR algorithms which rely on discretization.

II. THE NEED FOR CONTINUOUS-SPACE TREATMENT

In this section we analyze the situations in which purely translational SR cannot be accurately modeled via the discrete form (2).

A. Motivation

To develop intuition, we begin by examining a simple case of an imaging device whose lens’ PSF is negligible with respect to pixel size. In this situation, the pixels in the image correspond to integrals of the continuous-space scene $x(\mathbf{t})$ over non-overlapping rectangular environments. Therefore, the filter $s(\mathbf{t})$ of (1) in this case is a rectangular kernel:

$$s(\mathbf{t}) = \begin{cases} 1 & \mathbf{t} \in [0, 1]^d \\ 0 & \text{else.} \end{cases} \quad (3)$$

Now, assume we wish to construct a representation $x_{\text{HR}}[\mathbf{n}]$ whose resolution is Δ times higher in each dimension. In other words, we want the pixels in $x_{\text{HR}}[\mathbf{n}]$ to correspond to integrations of $x(\mathbf{t})$ over contiguous non-overlapping environments of size $1/\Delta$. This setting is demonstrated in Fig. 1 for a uni-dimensional signal ($d = 1$), with $t_1 = 0$, $t_2 = -1/2$ and $\Delta = 2$. As can be seen, in this case each of the elements in the low-resolution sequences $c_k[\mathbf{n}]$ can be expressed as a linear combination of the high-resolution pixels $x_{\text{HR}}[\mathbf{n}]$. For example, $c_1[1] = x_{\text{HR}}[1] + x_{\text{HR}}[2]$, and similarly $c_2[3] = x_{\text{HR}}[4] + x_{\text{HR}}[5]$.

The above example demonstrates that in some situations, the discrete formulation (2) is loyal to the physical settings. However, what made discretization possible in this scenario was the fact that the magnification Δ was an integer (namely $\Delta = 2$), and that the translations t_k were integer multiples of $1/\Delta$ (namely 0 and $-1/2$). It is easily verified that unless these two conditions are met, discretization is impossible in this setting.

Typically, the magnification factor Δ is indeed chosen to be an integer. Furthermore, one may argue that when Δ is sufficiently large, the translation t_k can be quite accurately approximated by ℓ_k/Δ for some $\ell_k \in \mathbb{Z}^d$, so that use of the discrete model (2) is justified. However, as we show next, even under these two assumptions, discretization is generally *impossible* for non-rectangular PSF functions.

B. Discretization with General PSFs

To study discretization with arbitrary PSF functions, we need to specify explicitly the relation between the discrete-space high resolution signal $x_{\text{HR}}[\mathbf{n}]$ and the continuous-space




Fig. 1: Increasing resolution by a factor of 2 with a rectangular PSF. The low resolution sequences $c_1[n]$ and $c_2[n]$ were acquired with translations $t_1 = 0$ and $t_2 = -1/2$ respectively, and correspond to linear combinations of elements from the high resolution sequence $x_{HR}[n]$.

signal $x(t)$. Here we assume that $x_{HR}[n]$ corresponds to the samples of $x(t)$ on a grid of points $\{n/\Delta : n \in \mathbb{Z}^d\}$ for some magnification $\Delta > 1$, after having been convolved with a PSF $w(t)$:

$$x_{HR}[n] = [(w * x)(t)]_{t=n/\Delta}. \quad (4)$$

This formulation is general enough to capture a wide range of interesting scenarios. For example, by choosing the target PSF $w(t) = \delta(t)$, we get that $x_{HR}[n]$ corresponds to the point-wise samples $x(n/\Delta)$. The PSF $w(t)$ can also be chosen to provide $x_{HR}[n]$ with the realistic impression that it was obtained by moving the imaging device towards the scene. This is done by setting $w(t) = s(\Delta t)$, as in the example of Fig. 1. Finally, if the profile of $s(t)$ is too wide with respect to the low-resolution grid, then $w(t)$ can be set to be narrower than $s(\Delta t)$, in order to obtain a deblurring effect in the high-resolution domain.

Using (4), $x_{HR}[n]$ can be written as

$$x_{HR}[n] = \int_{\mathbb{R}^d} x(t) w\left(\frac{n}{\Delta} - t\right) dt. \quad (5)$$

Similarly, (1) implies that the low-resolution samples $c_k[n]$ are given by

$$c_k[n] = \int_{\mathbb{R}^d} x(t) s(n - t - t_k) dt. \quad (6)$$

The conjecture that $c_k[n]$ is a linear function of $x_{HR}[n]$ corresponds to the existence of a set of coefficients $S_k[m, n]$,

$m, n \in \mathbb{Z}^d$, such that

$$\begin{aligned} c_k[n] &= \sum_{m \in \mathbb{Z}^d} S_k[n, m] x_{HR}[m] \\ &= \sum_{m \in \mathbb{Z}^d} S_k[n, m] \int_{t \in \mathbb{R}^d} x(t) w\left(\frac{m}{\Delta} - t\right) dt \\ &= \int_{t \in \mathbb{R}^d} x(t) \left(\sum_{m \in \mathbb{Z}^d} S_k[n, m] w\left(\frac{m}{\Delta} - t\right) \right) dt, \end{aligned} \quad (7)$$

where we used (5). Discretization is possible if and only if (7) equals (6) for every signal $x(t)$ in $L_2(\mathbb{R}^d)$. This happens only if

$$s(n + t - t_k) = \sum_{m \in \mathbb{Z}^d} S_k[n, m] w\left(\frac{m}{\Delta} + t\right) \quad (8)$$

for every $n \in \mathbb{Z}^d$ and for almost every $t \in \mathbb{R}^d$.

Taking the d -dimensional continuous-space Fourier transform (CSTF) of both sides of (8) with respect to t , yields the condition

$$S(\omega) e^{i\omega^T(n - t_k)} = W(\omega) A_{k,n}(e^{-i\omega/\Delta}). \quad (9)$$

Here ω denotes the d -dimensional frequency vector, $S(\omega)$ and $W(\omega)$ are the CSTFs of $s(t)$ and $w(t)$ respectively, and $A_{k,n}(e^{i\omega})$ is the discrete-space Fourier transform (DSFT) of $S_k[n, m]$ with respect to the index m . As motivated in Section II-A, we now make the simplifying assumption that the magnification factor Δ is an integer, and that the shift vector t_k comprises only integer multiples of $1/\Delta$. We therefore substitute $t_k = \ell_k/\Delta$, where ℓ_k is some vector in \mathbb{Z}^d , and get that $S(\omega)$ and $W(\omega)$ must satisfy

$$S(\omega) = W(\omega) A_{k,n}(e^{-i\omega/\Delta}) e^{i\omega^T(\ell_k - \Delta n)/\Delta}. \quad (10)$$

Finally, noting that $A_{k,n}(e^{-i\omega/\Delta}) e^{i\omega^T(\ell_k - \Delta n)/\Delta}$ is a function which is $2\pi\Delta$ -periodic in each of the elements of the frequency vector ω , we arrive at the following conclusion.

Proposition 1: Let $S(\omega)$ and $W(\omega)$ denote the CSFTs of the two PSFs $s(t)$ and $w(t)$ respectively. Furthermore, let $\Delta \geq 1$ be an integer magnification factor and t_k be a shift vector whose elements are integer multiples of $1/\Delta$. Assume that a continuous-space scene $x(t)$ is sampled on the grid $\{t = n - t_k : n \in \mathbb{Z}^d\}$, after having been convolved with $s(t)$, to yield the discrete-space signal $c_k[n]$. Similarly, assume that $x(t)$ is sampled on the grid $\{t = n/\Delta : n \in \mathbb{Z}^d\}$ after having been convolved with $w(t)$, to yield the discrete-space signal $x_{HR}[n]$. Then $c_k[n]$ can be written as a linear function of $x_{HR}[n]$ for every scene $x(t)$ if and only if there exists a $2\pi\Delta$ -periodic function $B_k(e^{i\omega/\Delta})$ such that

$$S(\omega) = W(\omega) B_k(e^{i\omega/\Delta}) \quad (11)$$

for almost every $\omega \in [-\pi, \pi]^d$.

We see that even in the simplistic case of integer magnification and quantized shifts, the target PSF $w(t)$ cannot be arbitrary. To demonstrate the implications of Proposition 1,




Fig. 2: Increasing resolution by a factor of 2 with the PSF $s(t) = \text{sinc}(t)$. Here $S(\omega) = \text{rect}(\omega/(2\pi))$ (top) and consequently $S(\omega/2)$ (middle) can be expressed as the product of $S(\omega)$ and a rectangular wave of period 4π (bottom).

consider the case in which $w(t) = s(\Delta t)$. In this situation, condition (12) becomes

$$S(\omega) = S\left(\frac{\omega}{\Delta}\right) B_k(e^{i\omega/\Delta}). \quad (12)$$

As we have seen in Section II-A, this condition holds for a rectangular PSF. As another example, it is easy to verify that (12) is satisfied when the PSF $s(t)$ is an ideal low-pass filter $\text{sinc}(t)$, as demonstrated in Fig. 2. However, in this scenario, the entire frequency content above $\omega = \pi$ is zeroed out in each dimension, rendering super-resolution ill-posed. A more realistic situation corresponds to the Gaussian PSF model, which is widely used in the SR literature. In this case,

$$S(\omega) = \exp\left\{\frac{-\|\omega\|^2}{2\sigma_s^2}\right\}. \quad (13)$$

for some $\sigma_s > 0$. Direct computation shows that

$$\frac{S(\omega)}{S(\frac{\omega}{\Delta})} = \exp\left\{\frac{-\|\omega\|^2}{2\sigma_s^2} \left(1 - \frac{1}{\Delta^2}\right)\right\}, \quad (14)$$

which is not a $2\pi\Delta$ -periodic function. Therefore this kind of PSF does not conform with (12), and cannot be treated via the discrete model (2). We conclude that to be loyal to the physical setting when the conditions of Proposition 1 are not met, SR must be addressed in the continuous domain.

III. THE VECTOR HYBRID WIENER FILTER

The discussion above revealed that the discrete model (2) is generally invalid. Therefore, we now turn to the problem of recovering the high resolution description $x_{\text{HR}}[\mathbf{n}]$ from the low resolution sequences $c_k[\mathbf{n}]$, $k = 1, \dots, K$, based on the continuous-space model (1).

A. Statistical Model

Ultimately, our goal is to devise an algorithm which produces an estimate $\hat{x}_{\text{HR}}[\mathbf{n}]$ minimizing the MSE

$$\epsilon_x^2[\mathbf{n}] = E\left[(\hat{x}_{\text{HR}}[\mathbf{n}] - x_{\text{HR}}[\mathbf{n}])^2 | x(t)\right], \quad (15)$$

for every pixel $\mathbf{n} \in \mathbb{Z}^d$, where the expectation is over realizations of the noise sequences $u_k[\mathbf{n}]$, $k = 1, \dots, K$ in (1). Unfortunately, the MSE depends on the underlying scene because $x_{\text{HR}}[\mathbf{n}]$ is a function of $x(t)$, as implied by (5). Minimizing $\epsilon_x^2[\mathbf{n}]$ uniformly over all signals $x(t)$ is impossible. Consequently, comparison between different SR approaches is not a well defined problem since one method may be better than another for some signals $x(t)$ and worse for others. To overcome this obstacle, we propose modeling the signal $x(t)$ as the realization of a random process with known statistics. This allows the replacement of the signal-dependent MSE $\epsilon_x^2[\mathbf{n}]$ by its expectation $\epsilon^2[\mathbf{n}] = E[\epsilon_x^2[\mathbf{n}]]$ over all possible signal realizations, resulting in the criterion

$$\epsilon^2[\mathbf{n}] = E\left[(\hat{x}_{\text{HR}}[\mathbf{n}] - x_{\text{HR}}[\mathbf{n}])^2\right]. \quad (16)$$

In the SR literature, various statistical priors have been proposed for describing the typical behavior of natural images. These include, for example, Gaussian and Huber random Markov fields [19], [20], [21]. However, these studies modeled the statistics of the desired discrete-space signal $x_{\text{HR}}[\mathbf{n}]$ rather than that of the continuous-space signal $x(t)$. In the single image interpolation literature, several recent works have used the assumption that continuous-space natural scenes are realizations of some stationary process [13], [14], [15], [16], [17]. It has been demonstrated that when the power-spectral density (PSD) of $x(t)$ is chosen to have a polynomial decay in frequency, the resulting recovery is superior to the Shannon interpolation, which relies on the traditional bandlimited model. Based on these findings, in this paper we adopt the stationarity assumption and extend the Bayesian single image recovery techniques to the multi-frame SR scenario.

B. Optimal Linear Recovery

Our goal is to *linearly* estimate $x_{\text{HR}}[\mathbf{n}]$ given the measurements $c_1[\mathbf{n}], \dots, c_K[\mathbf{n}]$. Due to the linearity of the convolution operation in (4), the linear minimum MSE (LMMSE) estimate of $x_{\text{HR}}[\mathbf{n}]$ is given by

$$\hat{x}_{\text{HR}}[\mathbf{n}] = [(w * \hat{x})(t)]_{t=\mathbf{n}/\Delta}, \quad (17)$$

where $\hat{x}(t)$ is the LMMSE estimate of $x(t)$ given the measurements. Therefore, to obtain a closed form expression for $\hat{x}_{\text{HR}}[\mathbf{n}]$, we will first compute the continuous-space LMMSE recovery $\hat{x}(t)$ and then sample it on the high-resolution grid using (17).

Before analyzing our estimation problem in detail, we would like to note that the purely translational SR setting treated in this paper can be viewed as a special case of multichannel sampling, as schematically shown in Fig. 3. In this setting, a signal $x(t)$ passes through K filters, which in our case correspond to $s_k(t) = s(t - t_k)$, $k = 1, \dots, K$, contaminated

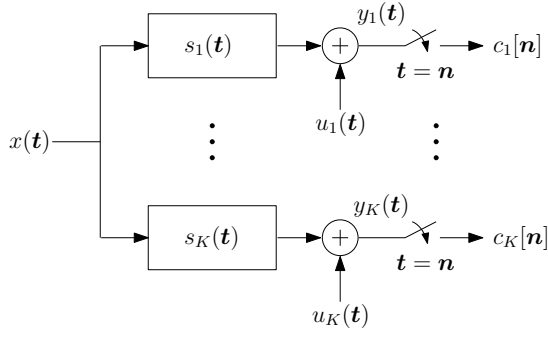


Fig. 3: Multichannel sampling scheme.

by continuous-space noise processes¹ $u_1(t), \dots, u_K(t)$, and then sampled on the grid $\{t = n : n \in \mathbb{Z}^d\}$ to yield the observed measurements $c_1[n], \dots, c_K[n]$. Thus, denoting

$$y_k(t) = (x * s_k)(t) + u_k(t), \quad k = 1, \dots, K, \quad (18)$$

our goal can be more generally described as estimating a continuous-space signal $x(t)$ based on equidistant samples of a set of continuous-space processes $y_1(t), \dots, y_K(t)$, which are statistically related to $x(t)$, such that the MSE $E[(x(t) - \hat{x}(t))^2]$ is minimized for every $t \in \mathbb{R}^d$.

The case in which only one measurement channel is available (*i.e.*, $K = 1$) was treated in [22]. We refer to the resulting reconstruction formula as the scalar hybrid Wiener filter since its input is the (scalar) discrete-space signal $y_1(n)$, $n \in \mathbb{Z}^d$, whereas its output is a continuous-space signal $\hat{x}(t)$, $t \in \mathbb{R}^d$. Consequently, we refer to our multichannel setup (for $K > 1$) as the *vector hybrid Wiener filter*.

To make the derivation general, all we assume in this section is that $x(t)$ and $y_1(t), \dots, y_K(t)$ are jointly WSS. We denote their cross-correlation functions by

$$R_{xy}^k(\tau) = E[x(t)y_k(t - \tau)], \quad k = 1, \dots, K, \quad (19)$$

$$R_{yy}^{k,\ell}(\tau) = E[y_\ell(t)y_k(t - \tau)], \quad k, \ell = 1, \dots, K. \quad (20)$$

The cross-spectra are given by the CSFTs

$$\Gamma_{xy}^k(\omega) = \mathcal{F}\{R_{xy}^k\}(\omega), \quad k = 1, \dots, K, \quad (21)$$

$$\Gamma_{yy}^{k,\ell}(\omega) = \mathcal{F}\{R_{yy}^{k,\ell}\}(\omega), \quad k, \ell = 1, \dots, K. \quad (22)$$

At any given location $t \in \mathbb{R}^d$, the estimate $\hat{x}(t)$ corresponds to a linear combination of the measurements $y_1(n), \dots, y_K(n)$, $n \in \mathbb{Z}^d$. Therefore, $\hat{x}(t)$ can be expressed as

$$\hat{x}(t) = \sum_{k=1}^K \sum_{n \in \mathbb{Z}^d} y_k(n) \tilde{v}_k(t - n), \quad (23)$$

for some set of functions $\tilde{v}_1(t), \dots, \tilde{v}_K(t)$, which we call *interpolation kernels*. The following theorem provides a closed form for the set of interpolation kernels minimizing the MSE.

Theorem 2 (Vector hybrid Wiener filter): *The interpolation kernels $\tilde{v}_1(t), \dots, \tilde{v}_K(t)$ in (23) that minimize the MSE*

¹To comply with (1), the continuous-space processes $u_1(t), \dots, u_K(t)$ should be chosen such that, when sampled on \mathbb{Z}^d , produce processes whose statistics are the same as that of the discrete-space noise signals $u_1[n], \dots, u_K[n]$ of (1).

$E[(x(t) - \hat{x}(t))^2]$ for every $t \in \mathbb{R}^d$, are given in the frequency domain by

$$\tilde{V}(\omega) = \left(\sum_{n \in \mathbb{Z}^d} \Gamma_{yy}(\omega - 2\pi n) \right)^{-1} \Gamma_{xy}(\omega), \quad (24)$$

where $\Gamma_{xy}(\omega) = (\Gamma_{xy}^1(\omega), \dots, \Gamma_{xy}^K(\omega))^T$, $\tilde{V}(\omega) = (\tilde{V}_1(\omega), \dots, \tilde{V}_K(\omega))^T$ is a $K \times 1$ vector comprising the CSFTs of the interpolation kernels, and $\Gamma_{yy}(\omega)$ is a $K \times K$ matrix whose (k, ℓ) th entry is $\Gamma_{yy}^{k,\ell}(\omega)$.

Proof: See Appendix A. ■

In the case $K = 1$, (24) reduces to the scalar hybrid Wiener filter [22], which was used for the task of single-image zooming in [15], [16], [17].

In the multichannel setting $K > 1$, the interpolation formula (23) can be broken up into K single-channel interpolation procedures. Specifically, (23) implies that each of the K low-resolution signals has to be interpolated using an appropriate reconstruction kernel and then the resulting K high-resolution estimates should be summed up to produce $\hat{x}(t)$. However, this naive implementation of (23) is typically not efficient since the optimal interpolation kernels $\tilde{v}_1(t), \dots, \tilde{v}_K(t)$ are generally not compactly supported. This means that in each of the single-channel recoveries, every pixel depends on all pixels of the corresponding low-resolution image. In the next sections, we provide several efficient implementations for the multichannel hybrid Wiener filter.

IV. EFFICIENT IMPLEMENTATION SCHEMES

We now specialize the hybrid Wiener filter to the SR setting and provide efficient implementation schemes.

We begin by noting that for any given sequence $c[n]$, $n \in \mathbb{Z}^d$, and function $h(t)$, $t \in \mathbb{R}^d$, we have the CSFT relation

$$\sum_{n \in \mathbb{Z}^d} c[n] h(t - n) \xrightarrow{\mathcal{F}} C(e^{j\omega}) H(\omega), \quad (25)$$

where $C(e^{j\omega})$ is the DSFT of $c[n]$ and $H(\omega)$ is the CSFT of $h(t)$. Using this relation together with (24), the reconstruction formula (23) can be expressed in the frequency domain:

$$\begin{aligned} \hat{X}(\omega) &= \sum_{k=1}^K C_k(e^{j\omega}) \tilde{V}_k(\omega) \\ &= \tilde{V}^T(\omega) C(e^{j\omega}) \\ &= \Gamma_{xy}^T(\omega) \left(\sum_{n \in \mathbb{Z}^d} \Gamma_{yy}(\omega - 2\pi n) \right)^{-T} C(e^{j\omega}) \\ &= \Gamma_{xy}^T(\omega) D(e^{j\omega}), \end{aligned} \quad (26)$$

where we denoted $C(\omega) = (C_1(\omega), \dots, C_K(\omega))^T$ and defined

$$D(e^{j\omega}) = \left(\sum_{n \in \mathbb{Z}^d} \Gamma_{yy}(\omega - 2\pi n) \right)^{-T} C(e^{j\omega}). \quad (27)$$

The term in the parentheses is a matrix-valued function, which is 2π -periodic in each of the elements of the frequency vector

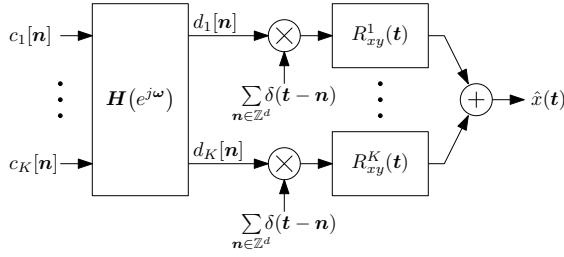


Fig. 4: Multichannel reconstruction scheme.

ω . Therefore, it corresponds to a multiple-input multiple-output (MIMO) digital filtering operation. Consequently, formula (23) can be implemented in two stages. First, the vector process $c[n] = (c_1[n], \dots, c_K[n])^T$ is fed into the $K \times K$ MIMO digital filter

$$H(e^{j\omega}) = \left(\sum_{n \in \mathbb{Z}^d} \Gamma_{yy}^T(\omega - 2\pi n) \right)^{-1} \quad (28)$$

to obtain a “corrected” process $d[n] = (d_1[n], \dots, d_K[n])^T$. Then, $\hat{x}(t)$ is formed using

$$\hat{x}(t) = \sum_{k=1}^K \sum_{n \in \mathbb{Z}^d} d_k[n] R_{xy}^k(t - n). \quad (29)$$

This scheme is depicted in Fig. 4.

Finally, to obtain the desired discrete-space high-resolution estimate we have to convolve $\hat{x}(t)$ with the target PSF $w(t)$ and sample it on a dense grid, as implied by (17). Therefore,

$$\hat{x}_{\text{HR}}[m] = \sum_{k=1}^K \sum_{n \in \mathbb{Z}^d} d_k[n] v_k\left(\frac{m}{\Delta} - n\right), \quad (30)$$

with the kernels $v_k(t) = (w * R_{xy}^k)(t)$, $k = 1, \dots, K$.

As we show next, this alternative scheme can be implemented efficiently under several conditions, which are commonly satisfied in practice. However, before discussing the implementation details, we note that the splitting of the recovery into a digital correction stage and an interpolation stage carries another important advantage. Specifically, the process of constructing the corrected images $d_1[n], \dots, d_K[n]$ from the raw images $c_1[n], \dots, c_K[n]$ via (27) does not depend on the magnification factor Δ . Thus, once $d_1[n], \dots, d_K[n]$ are computed, we can construct $\hat{x}_{\text{HR}}[m]$ using (30) with any desired magnification. In contrast, in most existing SR techniques, if the user wishes to change the magnification factor, then the algorithm has to be run from the beginning again.

We now proceed in making the following assumptions:

- 1) The noise processes $u_1(t), \dots, u_K(t)$ are independent of the scene $x(t)$ and are characterized by $R_{uu}^{k,\ell}(\tau) = E[u_\ell(t)u_k(t - \tau)] = \sigma_u^2 \delta_{k,\ell} \text{sinc}(\tau)$,
- 2) $\text{supp}\{s(t)\} \subseteq [-L_s, L_s]^d$,
- 3) $\text{supp}\{w(t)\} \subseteq [-L_w, L_w]^d$, and
- 4) $\text{supp}\{R_{xx}(t)\} \subseteq [-L_x, L_x]^d$.

Assumption 1 is equivalent to replacing the analog noise signals $u_1(t), \dots, u_K(t)$ by digital white noise processes

$u_1[n], \dots, u_K[n]$ with variance σ_u^2 , which are added after the sampling occurs. Assumptions 2, 3 and 4 are required in order to make the recovery algorithm efficient, as we detail next.

For simplicity, we also assume in the sequel that the translations t_1, \dots, t_K are in the range $[-0.5, 0.5]^d$. This does not pose any limitation since a shift $|t_k| > 0.5$ can always be compensated for by introducing an integer shift of $-\lfloor t_k + 0.5 \rfloor$ samples to $c_k[n]$ prior to reconstruction.

A. Interpolation Stage

We first examine the interpolation stage, namely the computation of $\hat{x}_{\text{HR}}[m]$ via (30). Using Assumption 1, the reconstruction kernels $R_{xy}^1(t), \dots, R_{xy}^K(t)$ appearing in Fig. 4 are given by

$$\begin{aligned} R_{xy}^k(t) &= E[x(t)y_k(0)] \\ &= E\left[x(t) \left(\int_{\mathbb{R}^d} s(-\tau)x(\tau - t_k)d\tau + u_k[0] \right)\right] \\ &= \int_{\mathbb{R}^d} \bar{s}(\tau) R_{xx}(t + t_k - \tau)d\tau \\ &= (\bar{s} * R_{xx})(t + t_k), \end{aligned} \quad (31)$$

where we denoted $\bar{s}(t) = s(-t)$. Consequently, the kernels $v_1(t), \dots, v_K(t)$ in (30), which are used to reconstruct $\hat{x}_{\text{HR}}[m]$, are given by

$$v_k(t) = (w * R_{xy}^k)(t) = (w * \bar{s} * R_{xx})(t + t_k). \quad (32)$$

Assumptions 2, 3 and 4 imply that the supports of the kernels $v_1(t), \dots, v_K(t)$ are finite, making the interpolation practical. Specifically, for any $m_0 \in \mathbb{Z}^d$, the computation of $\hat{x}_{\text{HR}}[m_0]$ involves at most $[2(L_w + L_s + L_x)]^d$ multiplications in processing each of the corrected channels $d_1[n], \dots, d_K[n]$. This sums up to a total of $K[2(L_w + L_s + L_x)]^d$ multiplications per high-resolution pixel $\hat{x}_{\text{HR}}[m_0]$.

This reduction in computational load in the interpolation stage comes at the cost of the MIMO digital filtering stage. Specifically, the optimal interpolation kernels $\tilde{v}_1(t), \dots, \tilde{v}_K(t)$ in the naive implementation (23) are generally not compactly supported. Therefore, the emergence of compactly supported kernels $v_1(t), \dots, v_K(t)$ in our alternative scheme means that the MIMO digital pre-filter $H(e^{j\omega})$ must have an infinite impulse response (IIR). However, as we show next, Assumptions 2 and 4 can be employed to implement the digital filtering stage efficiently as well.

V. EFFICIENT DIGITAL PRE-FILTERING

We now present three methods for efficiently applying the IIR MIMO filter $H(e^{j\omega})$ on the samples $c[n]$ to obtain the corrected low-resolution signals $d[n]$, which are used in the interpolation stage. The first approach we study is extremely efficient, but is valid only for 1D SR applications. The other two methods are iterative, however they can be applied in arbitrary dimensions. In particular, they are applicable to spatial and spatiotemporal SR scenarios.

A. Digital Correction Stage for 1D Signals

We begin by addressing the case of uni-dimensional signals (*i.e.*, $d = 1$). In this case $x(t)$ is a signal defined over $t \in \mathbb{R}$ and the low resolution measurements $c_k[n]$ are sequences defined over $n \in \mathbb{Z}$.

A close inspection of (28), reveals that $\mathbf{H}(e^{j\omega})$ corresponds to the convolutional inverse of the matrix sequence $\mathbf{Q}[n]$, whose (k, ℓ) entry is $R_{yy}^{\ell, k}(n)$. In our setting,

$$\begin{aligned} \mathbf{Q}_{k, \ell}[n] &= R_{yy}^{\ell, k}(n) \\ &= E[y_k(n)y_\ell(0)] \\ &= E\left[\left(\int_{\mathbb{R}^d} s(n-\eta)x(\eta-t_k)d\eta + u_k[n]\right) \times \left(\int_{\mathbb{R}^d} s(-\tau)x(\tau-t_\ell)d\tau + u_\ell[0]\right)\right] \\ &= \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} s(n-\eta)s(-\tau)R_{xx}(\eta-\tau+t_\ell-t_k)d\tau d\eta \\ &\quad + \sigma_u^2 \delta_{k, \ell} \delta[n] \\ &= (\bar{s} * s * R_{xx})(n+t_\ell-t_k) + \sigma_u^2 \delta_{k, \ell} \delta[n]. \end{aligned} \quad (33)$$

Assumptions 2 and 4 imply that the sequence $\mathbf{Q}[n]$ is compactly supported. This, in turn, means that the Z -transform of $\mathbf{Q}[n]$ can be written as

$$\mathbf{Q}(z) = \mathbf{A}_p^T z^{-p} + \dots + \mathbf{A}_1^T z^{-1} + \mathbf{A}_0 + \mathbf{A}_1 z + \dots + \mathbf{A}_p z^p, \quad (34)$$

where $p = \lceil 2L_s + L_x \rceil$.

Our key observation is that if $\mathbf{Q}(z)$ is positive definite on the unit circle $|z| = 1$, then it can be factored as

$$\mathbf{Q}(z) = \mathbf{B}(z)\mathbf{B}^T(z^{-1}), \quad (35)$$

where

$$\mathbf{B}(z) = \mathbf{B}_0 + \mathbf{B}_1 z + \dots + \mathbf{B}_p z^p \quad (36)$$

is a matrix whose determinant does not vanish inside the unit circle $|z| \leq 1$. Such a factorization can be obtained *e.g.*, by any one of the methods surveyed in [23]. Using (35), we can now carry out the filtering by $\mathbf{H}(z) = \mathbf{Q}^{-1}(z)$ of (28) in two stages. Specifically, the “corrected” vector sequence $\mathbf{d}[n]$ is given in the frequency domain by

$$\begin{aligned} \mathbf{D}(z) &= \mathbf{H}(z)\mathbf{C}(z) \\ &= \mathbf{Q}^{-1}(z)\mathbf{C}(z) \\ &= \mathbf{B}^{-T}(z^{-1})\mathbf{B}^{-1}(z)\mathbf{C}(z). \end{aligned} \quad (37)$$

Thus, we first form an auxiliary vector sequence $\mathbf{c}'[n]$, by applying the MIMO filter $\mathbf{B}^{-1}(z)$ on the samples $\mathbf{c}[n]$. Then, we obtain $\mathbf{d}[n]$ by passing $\mathbf{c}'[n]$ through the MIMO filter $\mathbf{B}^{-T}(z^{-1})$.

The input-output relation of the first filtering operation is given by

$$\begin{aligned} \mathbf{C}(z) &= \mathbf{B}(z)\mathbf{C}'(z) \\ &= (\mathbf{B}_0 + \mathbf{B}_1 z + \dots + \mathbf{B}_p z^p) \mathbf{C}'(z), \end{aligned} \quad (38)$$

which shows that

$$\mathbf{C}'(z) = \mathbf{B}_0^{-1} (\mathbf{C}(z) - (\mathbf{B}_1 z + \dots + \mathbf{B}_p z^p) \mathbf{C}'(z)). \quad (39)$$

In the time domain, this filter can be implemented by the recursive formula

$$\mathbf{c}'[n] = \mathbf{B}_0^{-1} (\mathbf{c}[n] - \mathbf{B}_1 \mathbf{c}'[n+1] - \dots - \mathbf{B}_p \mathbf{c}'[n+p]) \quad (40)$$

running from right to left. Similarly, the input-output relation of the second filtering operation is given by

$$\begin{aligned} \mathbf{C}'(z) &= \mathbf{B}^T(z^{-1})\mathbf{D}(z) \\ &= (\mathbf{B}_0^T + \mathbf{B}_1^T z^{-1} + \dots + \mathbf{B}_p^T z^{-p}) \mathbf{D}(z), \end{aligned} \quad (41)$$

which shows that

$$\mathbf{D}(z) = \mathbf{B}_0^{-T} \left(\mathbf{C}'(z) - (\mathbf{B}_1^T z^{-1} + \dots + \mathbf{B}_p^T z^{-p}) \mathbf{D}(z) \right). \quad (42)$$

In the time domain, this filter can be implemented by the recursive formula

$$\mathbf{d}[n] = \mathbf{B}_0^{-T} \left(\mathbf{c}'[n] - \mathbf{B}_1^T \mathbf{d}[n-1] - \dots - \mathbf{B}_p^T \mathbf{d}[n-p] \right) \quad (43)$$

running from left to right.

Due to the spectral factorization properties, the determinant of the anti-causal filter $\mathbf{B}^{-1}(z)$ has no poles inside the complex unit circle $|z| \leq 1$ and therefore (40) is a stable operation. Similarly, the determinant of the causal filter $\mathbf{B}^{-T}(z^{-1})$ has no poles outside the unit circle, implying that the recursive formula (43) is also stable. Algorithm 1 summarizes the proposed 1D SR interpolation scheme.

The technique outlined above follows similar ideas as those in the direct B-spline transform introduced in [24], which is used for single-channel spline interpolation from uniformly spaced samples. There are, however, several major differences a practitioner must be aware of, which are caused by the fact that the SR setting is more complicated. First, in contrast to the scalar interpolation scenario, here the order in which (40) and (43) are performed is important. Second, this scheme cannot be extended to multiple dimensions by operating along each dimension separately, as done in the scalar case. Specifically, even if the PSF $s(\mathbf{t})$ and the autocorrelation $R_{xx}(\mathbf{t})$ are separable functions of the coordinates \mathbf{t} , the digital filtering by $\mathbf{H}(e^{j\omega})$ of (28) is generally not equivalent to applying (40) and (43) sequentially on each dimension. Last, we note that whereas in scalar interpolation the filter coefficients are computed in advance, in typical SR applications the translations t_1, \dots, t_K are estimated from the data and thus the matrices $\mathbf{B}_0, \dots, \mathbf{B}_p$ are not known beforehand. This may complicate potential hardware implementations of the filtering stage.

B. Digital Correction Stage via the Neumann Series

Next, we propose an efficient iterative technique for implementing the filtering by $\mathbf{H}(e^{j\omega})$ of (28) in arbitrary dimensions.

We begin by choosing a constant $\alpha > 0$ such that

$$\|\mathbf{I} - \alpha \mathbf{Q}(e^{j\omega})\| < 1, \quad \forall \omega \in \mathbb{R}^d, \quad (44)$$

where $\mathbf{Q}(e^{j\omega}) = \mathbf{H}^{-1}(e^{j\omega})$, \mathbf{I} is the identity matrix and $\|\mathbf{A}\|$ denotes the spectral norm of a matrix \mathbf{A} . Below, we discuss how such a constant can be chosen. Now, the filtering

Algorithm 1 Fast 1D hybrid Wiener super-resolution.

Input: Low resolution sequences $c_1[n], \dots, c_K[n]$, $n \in \mathbb{Z}$, noise variance σ_u^2 , shifts t_1, \dots, t_K in the range $[-0.5, 0.5]$, imaging PSF $s(t)$ supported on $[-L_s, L_s]$, target PSF $w(t)$ supported on $[-L_w, L_w]$, signal's autocorrelation $R_{xx}(t)$ supported on $[-L_x, L_x]$, magnification factor Δ .

Output: High resolution recovery $\hat{x}_{\text{HR}}[n]$, $n \in \mathbb{Z}$.

- 1: **PSD computation:** Set $(A_0)_{k,\ell} = (\bar{s} * s * R_{xx})(t_\ell - t_k) + \sigma_u^2 \delta_{k,\ell}$ and $(A_n)_{k,\ell} = (\bar{s} * s * R_{xx})(n + t_\ell - t_k)$, $n = 1, \dots, p$, where $p = \lceil 2L_s + L_x \rceil$.
 - 2: **Spectral factorization:** Given the matrices A_0, \dots, A_p , compute matrices B_0, \dots, B_p satisfying (35) using any matrix spectral factorization algorithm, *e.g.*, one of the methods surveyed in [23].
 - 3: **Anti-causal MIMO filtering:** Apply (40) on the vector process $\mathbf{c}[n] = (c_1[n], \dots, c_K[n])^T$ to obtain an auxiliary vector process $\mathbf{c}'[n] = (c'_1[n], \dots, c'_K[n])^T$.
 - 4: **Causal filtering:** Apply (43) on $\mathbf{c}'[n]$ to obtain the vector process $\mathbf{d}[n] = (d_1[n], \dots, d_K[n])^T$.
 - 5: **Interpolation:** Compute $\hat{x}_{\text{HR}}[m] = \sum_{k=1}^K \sum_{n \in \mathbb{Z}} d_k[n] v_k(m/\Delta - n)$, where $v_k(t) = (w * \bar{s} * R_{xx})(t + t_k)$, $k = 1, \dots, K$.
-

operation (27) with $\mathbf{H}(e^{j\omega})$ of (28) can be expressed via the Neumann series [25]

$$\begin{aligned} \mathbf{D}(e^{j\omega}) &= \mathbf{H}(e^{j\omega}) \mathbf{C}(e^{j\omega}) \\ &= \mathbf{Q}^{-1}(e^{j\omega}) \mathbf{C}(e^{j\omega}) \\ &= \alpha (\mathbf{I} - (\mathbf{I} - \alpha \mathbf{Q}(e^{j\omega})))^{-1} \mathbf{C}(e^{j\omega}) \\ &= \alpha \sum_{r=0}^{\infty} (\mathbf{I} - \alpha \mathbf{Q}(e^{j\omega}))^r \mathbf{C}(e^{j\omega}). \end{aligned} \quad (45)$$

As we have seen in Section V-A, Assumptions 2 and 4 imply that $\mathbf{Q}(e^{j\omega})$ is a finite impulse response (FIR) filter. Consequently, $\mathbf{I} - \alpha \mathbf{Q}(e^{j\omega})$ is also an FIR filter. This shows that filtering the samples $\mathbf{c}[n]$ with the IIR filter $\mathbf{H}(e^{j\omega})$ can be performed by a sequence of FIR filtering operations.

To obtain a simple recursive implementation, we let

$$\hat{\mathbf{D}}_m(e^{j\omega}) = \alpha \sum_{r=0}^m (\mathbf{I} - \alpha \mathbf{Q}(e^{j\omega}))^r \mathbf{C}(e^{j\omega}) \quad (46)$$

denote the approximation of $\mathbf{D}(e^{j\omega})$ by the first $m+1$ terms of the Neumann series. Then $\hat{\mathbf{D}}_{m+1}(e^{j\omega})$ can be calculated from $\hat{\mathbf{D}}_m(e^{j\omega})$ as follows:

$$\begin{aligned} \hat{\mathbf{D}}_{m+1}(e^{j\omega}) &= \alpha \sum_{r=0}^{m+1} (\mathbf{I} - \alpha \mathbf{Q}(e^{j\omega}))^r \mathbf{C}(e^{j\omega}) \\ &= \alpha \left(\mathbf{C}(e^{j\omega}) + \sum_{r=1}^{m+1} (\mathbf{I} - \alpha \mathbf{Q}(e^{j\omega}))^r \mathbf{C}(e^{j\omega}) \right) \\ &= \alpha \left(\mathbf{C}(e^{j\omega}) + (\mathbf{I} - \alpha \mathbf{Q}(e^{j\omega})) \frac{1}{\alpha} \hat{\mathbf{D}}_m(e^{j\omega}) \right) \\ &= \hat{\mathbf{D}}_m(e^{j\omega}) + \alpha (\mathbf{C}(e^{j\omega}) - \mathbf{Q}(e^{j\omega}) \hat{\mathbf{D}}_m(e^{j\omega})). \end{aligned} \quad (47)$$

As $\hat{\mathbf{D}}_m(e^{j\omega})$ approaches $\mathbf{Q}^{-1}(e^{j\omega}) \mathbf{C}(e^{j\omega})$, the increment $\mathbf{C}(e^{j\omega}) - \mathbf{Q}(e^{j\omega}) \hat{\mathbf{D}}_m(e^{j\omega})$ approaches zero. In practice, when this increment is sufficiently small, the iterations may be terminated.

To choose a step size $\alpha > 0$ that guarantees convergence, we note that the matrix $\mathbf{Q}(e^{j\omega})$ is positive semi-definite.

Consequently,

$$\begin{aligned} \|\mathbf{I} - \alpha \mathbf{Q}(e^{j\omega})\| &= \\ &= \max \{1 - \alpha \lambda_{\min}(\mathbf{Q}(e^{j\omega})), \alpha \lambda_{\max}(\mathbf{Q}(e^{j\omega})) - 1\}, \end{aligned} \quad (48)$$

where $\lambda_{\min}(\cdot)$ and $\lambda_{\max}(\cdot)$ denote respectively the minimal and maximal eigenvalues of their matrix argument. Therefore, for (44) to hold, we need that

$$1 - \alpha \lambda_{\min}(\mathbf{Q}(e^{j\omega})) < 1, \quad (49)$$

$$\alpha \lambda_{\max}(\mathbf{Q}(e^{j\omega})) - 1 < 1, \quad (50)$$

for all $\omega \in [-\pi, \pi]^d$. Condition (50) is satisfied if

$$\alpha < \frac{2}{\lambda_{\max}(\mathbf{Q}(e^{j\omega}))}, \quad \forall \omega \in [-\pi, \pi]^d. \quad (51)$$

Condition (49) can be fulfilled only if $\lambda_{\min}(\mathbf{Q}(e^{j\omega})) > 0$ for every $\omega \in [-\pi, \pi]^d$. In this case, it is satisfied if $\alpha > 0$.

In practice, it may be time-consuming to seek the worst-case $\omega \in [-\pi, \pi]^d$ in (51) such that convergence is ensured for all frequencies. However, recall from (28) that $\mathbf{Q}(e^{j\omega})$ corresponds to a sum of shifted replicas of the measurement spectrum $\Gamma_{yy}(\omega)$. In typical SR applications, this spectrum rapidly decays as a function of frequency. Thus, it is reasonable to expect that $\lambda_{\max}(\mathbf{Q}(e^{j\omega}))$ attains its maximum at $\omega = \mathbf{0}$, so that (51) can be verified only there. Below, we also discuss how the step-size α can be varied from one iteration to the next, in order to achieve a faster convergence. Algorithm 2 summarizes the proposed SR interpolation scheme.

C. Digital Correction Stage via Steepest-Descent Iterations

Algorithm 2 is very efficient, since every iteration consists only of applying a MIMO FIR filter with $(2p+1)^d$ coefficients. Its disadvantage, though, is that it is commonly impossible to choose a step size α for which the convergence at all frequencies is fast. We therefore now propose a method in which the step size is updated from one iteration to the next. As we will see, this makes each iteration slower, but guarantees a rapid convergence at all frequencies.

As we have seen, the corrected signal $\mathbf{d}[n]$ is obtained by applying the filter $\mathbf{H}(e^{j\omega}) = \mathbf{Q}^{-1}(e^{j\omega})$ on the samples $\mathbf{c}[n]$.

Algorithm 2 Hybrid Wiener super-resolution via Neumann iterations.

Input: Low resolution signals $c_1[\mathbf{n}], \dots, c_K[\mathbf{n}], \mathbf{n} \in \mathbb{Z}^d$, noise variance σ_u^2 , shifts $\mathbf{t}_1, \dots, \mathbf{t}_K$ in the range $[-0.5, 0.5]^d$, imaging PSF $s(\mathbf{t})$ supported on $[-L_s, L_s]^d$, target PSF $w(\mathbf{t})$ supported on $[-L_w, L_w]^d$, signal's autocorrelation $R_{xx}(\mathbf{t})$ supported on $[-L_x, L_x]^d$, magnification factor Δ .

Output: High resolution recovery $\hat{x}_{\text{HR}}[\mathbf{n}], \mathbf{n} \in \mathbb{Z}^d$.

- 1: **PSD computation:** Set $p = \lceil 2L_s + L_x \rceil$ and calculate the matrices $(\mathbf{Q}[\mathbf{n}])_{k,\ell} = (\bar{s} * s * R_{xx})(\mathbf{n} + \mathbf{t}_\ell - \mathbf{t}_k) + \sigma_u^2 \delta_{k,\ell} \delta[\mathbf{n}]$ for $\mathbf{n} \in [-p, p]^d$.
 - 2: **Step size calculation:** Choose a step size α satisfying $0 < \alpha < 2/\lambda_{\max}(\sum_{\mathbf{n} \in [-p, p]^d} \mathbf{Q}[\mathbf{n}])$.
 - 3: **Initialization:** Set $m = 0$, $\mathbf{d}_m[\mathbf{n}] = \mathbf{0}$, for all $\mathbf{n} \in \mathbb{Z}^d$.
 - 4: **repeat**
 - 5: Compute $\varepsilon_m[\ell] = c[\ell] - \sum_{\mathbf{r} \in [-p, p]^d} \mathbf{Q}[\mathbf{r}] \mathbf{d}_m[\ell - \mathbf{r}]$ for every $\ell \in \mathbb{Z}^d$.
 - 6: To use the optimal step-size, compute $\tilde{\varepsilon}_m[\ell] = \sum_{\mathbf{r} \in [-p, p]^d} \mathbf{Q}[\mathbf{r}] \varepsilon_m[\ell - \mathbf{r}]$ for every $\ell \in \mathbb{Z}^d$ and set $\alpha_m = (\sum_{\mathbf{n} \in \mathbb{Z}^d} \varepsilon_m^T[\mathbf{n}] \tilde{\varepsilon}_m[\mathbf{n}]) / (\sum_{\mathbf{n} \in \mathbb{Z}^d} \|\tilde{\varepsilon}_m[\mathbf{n}]\|^2)$. Otherwise set $\alpha_m = \alpha$.
 - 7: Set $\mathbf{d}_{m+1}[\mathbf{n}] = \mathbf{d}_m[\mathbf{n}] + \alpha_m \varepsilon_m[\mathbf{n}]$ for every $\mathbf{n} \in \mathbb{Z}^d$.
 - 8: Set $m \leftarrow m + 1$.
 - 9: **until** $\sum_{\mathbf{n} \in \mathbb{Z}^d} \|\varepsilon_m[\mathbf{n}]\|^2$ ceases to decrease.
 - 10: **Interpolation:** Compute $\hat{x}_{\text{HR}}[\ell] = \sum_{k=1}^K \sum_{\mathbf{n} \in \mathbb{Z}^d} d_k[\mathbf{n}] v_k(\ell/\Delta - \mathbf{n})$, where $v_k(\mathbf{t}) = (w * \bar{s} * R_{xx})(\mathbf{t} + \mathbf{t}_k)$, $k = 1, \dots, K$.
-

Therefore, for every $\omega \in [-\pi, \pi]^d$, the vector $\mathbf{D}(e^{j\omega})$ is the minimizer of the objective

$$f(\mathbf{D}(e^{j\omega})) = \|\mathbf{Q}(e^{j\omega}) \mathbf{D}(e^{j\omega}) - \mathbf{C}(e^{j\omega})\|^2. \quad (52)$$

The gradient of $f(\mathbf{D}(e^{j\omega}))$ with respect to $\mathbf{D}(e^{j\omega})$ is given by

$$\nabla_{\mathbf{D}} f(e^{j\omega}) = 2\mathbf{Q}^H(e^{j\omega}) (\mathbf{Q}(e^{j\omega}) \mathbf{D}(e^{j\omega}) - \mathbf{C}(e^{j\omega})). \quad (53)$$

Noting that in our setting $\mathbf{Q}^H(e^{j\omega}) = \mathbf{Q}(e^{j\omega})$, the minimizer of $f(e^{j\omega})$ can be found using the steepest descent iterations

$$\begin{aligned} \hat{\mathbf{D}}_{m+1}(e^{j\omega}) &= \hat{\mathbf{D}}_m(e^{j\omega}) + \\ &\quad \alpha_m \mathbf{Q}(e^{j\omega}) \left(\mathbf{C}(e^{j\omega}) - \mathbf{Q}(e^{j\omega}) \hat{\mathbf{D}}_m(e^{j\omega}) \right) \end{aligned} \quad (54)$$

Note that (54) is the same as the Neumann update rule (47), besides the factor $\mathbf{Q}(e^{j\omega})$ and the fact that the step-size α_m may change from one iteration to the next. If α_m is properly chosen at each iteration, then the sequence of approximations $\hat{\mathbf{D}}_m(e^{j\omega})$ will converge to $\mathbf{D}(e^{j\omega})$ [26].

We would like to choose a step size α_m , such that the energy of the error is minimized. Specifically, we would like to minimize

$$\begin{aligned} \int_{\omega \in [-\pi, \pi]^d} f(\hat{\mathbf{D}}_{m+1}(e^{j\omega})) d\omega &= \\ \int_{\omega \in [-\pi, \pi]^d} f(\hat{\mathbf{D}}_m(e^{j\omega}) &+ \\ + \alpha_m \mathbf{Q}(e^{j\omega}) \left(\mathbf{C}(e^{j\omega}) - \mathbf{Q}(e^{j\omega}) \hat{\mathbf{D}}_m(e^{j\omega}) \right)) d\omega. \end{aligned} \quad (55)$$

Substituting the expression for $f(\mathbf{D}(e^{j\omega}))$ from (52) and equating the derivative with respect to α_m to zero, we find

that the optimal step size is given by

$$\begin{aligned} \alpha_m &= \frac{\int_{\omega \in [-\pi, \pi]^d} \|\tilde{\mathbf{E}}_m(e^{j\omega})\|^2 d\omega}{\int_{\omega \in [-\pi, \pi]^d} \|\mathbf{Q}(e^{j\omega}) \tilde{\mathbf{E}}_m(e^{j\omega})\|^2 d\omega} \\ &= \frac{\sum_{\mathbf{n} \in \mathbb{Z}^d} \|\tilde{\varepsilon}_m[\mathbf{n}]\|^2}{\sum_{\mathbf{n} \in \mathbb{Z}^d} \left\| \sum_{\ell \in [-p, p]^d} \mathbf{Q}[\ell] \tilde{\varepsilon}_m[\mathbf{n} - \ell] \right\|^2}, \end{aligned} \quad (56)$$

where $\tilde{\mathbf{E}}_m(e^{j\omega}) = \mathbf{Q}(e^{j\omega}) (\mathbf{C}(e^{j\omega}) - \mathbf{Q}(e^{j\omega}) \hat{\mathbf{D}}_m(e^{j\omega}))$ and $\tilde{\varepsilon}_m[\mathbf{n}]$ is its inverse DSFT. Algorithm 3 summarizes the proposed SR interpolation scheme.

It is important to note that each steepest-descent iteration comprises three applications of the FIR MIMO filter $\mathbf{Q}[\mathbf{n}]$, as opposed to the single filtering operation required in the Neumann method. Therefore, for Algorithm 3 to be preferable to Algorithm 2 in terms of running time, it has to converge to the optimum at least three times faster. Unfortunately, this is usually not the case. In typical situations, its first few iterations lead to a dramatic increase in the error, but its convergence becomes very slow near the optimum. This phenomenon of the steepest-descent method is well known [26].

Interestingly, the computational complexity of Algorithm 3 can be reduced, while retaining the benefits of a varying step-size. To do this, we note that the direction of the Neumann update (47) is related to that of the steepest-descent (54), via multiplication by the positive-definite matrix $\mathbf{Q}^{-1}(e^{j\omega})$. Thus, the objective is guaranteed to decrease in the direction of the Neumann update (47). Now, we can compute the step-size α_m leading to the maximal reduction in the error for the Neumann iterations, similar to equations (55) and (56), which were relevant for the steepest-descent method. The resulting optimal Neumann step-size is given by

$$\alpha_m = \frac{\sum_{\mathbf{n} \in \mathbb{Z}^d} \varepsilon_m^T[\mathbf{n}] \tilde{\varepsilon}_m[\mathbf{n}]}{\sum_{\mathbf{n} \in \mathbb{Z}^d} \|\tilde{\varepsilon}_m[\mathbf{n}]\|^2}, \quad (57)$$

where $\varepsilon_m[\ell] = c[\ell] - \sum_{\mathbf{r} \in [-p, p]^d} \mathbf{Q}[\mathbf{r}] \mathbf{d}_m[\ell - \mathbf{r}]$, and $\tilde{\varepsilon}_m[\mathbf{n}]$ is as in (56). This option of using a varying step-size in the Neumann iterations is presented in Algorithm 2. It requires

Algorithm 3 Hybrid Wiener super-resolution via steepest descent iterations.

Input: Low resolution signals $c_1[\mathbf{n}], \dots, c_K[\mathbf{n}]$, $\mathbf{n} \in \mathbb{Z}^d$, noise variance σ_u^2 , shifts $\mathbf{t}_1, \dots, \mathbf{t}_K$ in the range $[-0.5, 0.5]^d$, imaging PSF $s(\mathbf{t})$ supported on $[-L_s, L_s]^d$, target PSF $w(\mathbf{t})$ supported on $[-L_w, L_w]^d$, signal's autocorrelation $R_{xx}(\mathbf{t})$ supported on $[-L_x, L_x]^d$, magnification factor Δ .

Output: High resolution recovery $\hat{x}_{\text{HR}}[\mathbf{n}]$, $\mathbf{n} \in \mathbb{Z}^d$.

- 1: **PSD computation:** Set $p = \lceil 2L_s + L_x \rceil$ and calculate the matrices $(\mathbf{Q}[\mathbf{n}])_{k,\ell} = (\bar{s} * s * R_{xx})(\mathbf{n} + \mathbf{t}_\ell - \mathbf{t}_k) + \sigma_u^2 \delta_{k,\ell} \delta[\mathbf{n}]$ for $\mathbf{n} \in [-p, p]^d$.
 - 2: **Initialization:** Set $m = 0$, $\mathbf{d}_m[\mathbf{n}] = \mathbf{0}$, for all $\mathbf{n} \in \mathbb{Z}^d$.
 - 3: **repeat**
 - 4: Compute $\boldsymbol{\varepsilon}_m[\ell] = \mathbf{c}[\ell] - \sum_{\mathbf{r} \in [-p, p]^d} \mathbf{Q}[\mathbf{r}] \mathbf{d}_m[\ell - \mathbf{r}]$ for every $\ell \in \mathbb{Z}^d$.
 - 5: Compute $\tilde{\boldsymbol{\varepsilon}}_m[\ell] = \sum_{\mathbf{r} \in [-p, p]^d} \mathbf{Q}[\mathbf{r}] \boldsymbol{\varepsilon}_m[\ell - \mathbf{r}]$ for every $\ell \in \mathbb{Z}^d$.
 - 6: Compute $\alpha_m = (\sum_{\mathbf{n} \in \mathbb{Z}^d} \|\tilde{\boldsymbol{\varepsilon}}_m[\mathbf{n}]\|^2) / (\sum_{\mathbf{n} \in \mathbb{Z}^d} \|\sum_{\ell \in [-p, p]^d} \mathbf{Q}[\ell] \tilde{\boldsymbol{\varepsilon}}_m[\mathbf{n} - \ell]\|^2)$.
 - 7: Set $\mathbf{d}_{m+1}[\mathbf{n}] = \mathbf{d}_m[\mathbf{n}] + \alpha_m \tilde{\boldsymbol{\varepsilon}}_m[\mathbf{n}]$ for every $\mathbf{n} \in \mathbb{Z}^d$.
 - 8: Set $m \leftarrow m + 1$.
 - 9: **until** $\sum_{\mathbf{n} \in \mathbb{Z}^d} \|\boldsymbol{\varepsilon}[\mathbf{n}]\|^2$ ceases to decrease.
 - 10: **Interpolation:** Compute $\hat{x}_{\text{HR}}[\ell] = \sum_{k=1}^K \sum_{\mathbf{n} \in \mathbb{Z}} d_k[\mathbf{n}] v_k(\ell/\Delta - \mathbf{n})$, where $v_k(\mathbf{t}) = (w * R_{xy}^k)(\mathbf{t})$, $k = 1, \dots, K$.
-

two filtering operations in each iteration, as opposed to the three required in Algorithm 3.

D. Comparison

For 1D applications, Algorithm 1 seems to be advantageous over Algorithms 2 and 3 due to its non-iterative form. Specifically, it does not suffer from error due to a finite number of iterations. Furthermore, its time-consumption is known in advance and does not depend on the underlying signal.

For SR in two or more dimensions, one must use either Algorithm 2 (with a fixed or varying step-size) or Algorithm 3. The rate of convergence of the steepest descent method (Algorithm 3) is known to be very slow near the optimum, especially whenever the ratio $\lambda_{\max}(\mathbf{Q}(e^{j\omega})) / \lambda_{\min}(\mathbf{Q}(e^{j\omega}))$ is large [26]. The latter is controlled by the noise variance σ_u^2 , which appears in the diagonal elements of $\mathbf{Q}(e^{j\omega})$. The condition number of $\mathbf{Q}(e^{j\omega})$ also affects the convergence rate of the Neumann method in a similar manner.

Figure 5 depicts the root MSE (RMSE) as a function of the number of iterations in the constant and the varying step-size Neuman methods, and in the steepest descent algorithm. Figures 5a and 5b differ by the noise level σ_u^2 . The former was run with a signal-to-noise ratio (SNR) of 0dB whereas the latter with 8dB. These experiments validate that the convergence rate of all algorithms is indeed higher for low SNRs. Furthermore, it is evident that the steepest-descent method has the slowest convergence rate. That, together with the fact that its iterations are more computational demanding than the other methods, make it the least favorable of all three approaches.

The varying step-size Neumann iterations provide the fastest convergence. However, the gain with respect to using a constant step-size is rather small. On the other hand, each iteration of the varying step-size Neuman approach requires twice the amount of computations as its fixed step-size counterpart. Therefore, we conclude that the best approach in terms of overall time consumption, is Algorithm 2 with a constant step-size.

VI. ERROR ANALYSIS AND PATTERNS

In [27], [28], the authors investigated bounds on the achievable MSE in SR scenarios. These analyses, however, did not incorporate any assumed prior knowledge on the underlying scene. As we show next, one of the benefits of our approach, is that we can obtain closed form expressions for the MSE in our setting. This is due to the fact that we assume knowledge of the statistics of $x(\mathbf{t})$. The error analysis we provide in this section is particularly useful in predicting the appearance and the extent to which periodic structures appear in the reconstructed high-resolution recovery $\hat{x}_{\text{HR}}[\mathbf{n}]$.

Our goal is to obtain an expression for the MSE $\epsilon^2[\mathbf{m}]$ defined in (16), which is the error in estimating $x_{\text{HR}}[\mathbf{m}]$ by $\hat{x}_{\text{HR}}[\mathbf{m}]$. Substituting (4) and (17), $\epsilon^2[\mathbf{m}]$ can be expressed as

$$\begin{aligned}
 \epsilon^2[\mathbf{m}] &= E[(\hat{x}_{\text{HR}}[\mathbf{n}] - x_{\text{HR}}[\mathbf{n}])^2] \\
 &= E\left[\left(\int_{\mathbf{t} \in \mathbb{R}^d} (x(\mathbf{t}) - \hat{x}(\mathbf{t})) w\left(\frac{\mathbf{m}}{\Delta} - \mathbf{t}\right) d\mathbf{t}\right)^2\right] \\
 &= E\left[\int_{\mathbf{t} \in \mathbb{R}^d} \int_{\boldsymbol{\eta} \in \mathbb{R}^d} (x(\mathbf{t}) - \hat{x}(\mathbf{t}))(x(\boldsymbol{\eta}) - \hat{x}(\boldsymbol{\eta})) \right. \\
 &\quad \left. \times w\left(\frac{\mathbf{m}}{\Delta} - \mathbf{t}\right) w\left(\frac{\mathbf{m}}{\Delta} - \boldsymbol{\eta}\right) d\mathbf{t} d\boldsymbol{\eta}\right] \\
 &= \int_{\mathbf{t} \in \mathbb{R}^d} \int_{\boldsymbol{\eta} \in \mathbb{R}^d} e(\mathbf{t}, \boldsymbol{\eta}) w\left(\frac{\mathbf{m}}{\Delta} - \mathbf{t}\right) w\left(\frac{\mathbf{m}}{\Delta} - \boldsymbol{\eta}\right) d\mathbf{t} d\boldsymbol{\eta},
 \end{aligned} \tag{58}$$

where $\hat{x}(\mathbf{t})$ denotes the LMMSE estimate of the continuous scene $x(\mathbf{t})$, and

$$e(\mathbf{t}, \boldsymbol{\eta}) = E[(x(\mathbf{t}) - \hat{x}(\mathbf{t}))(x(\boldsymbol{\eta}) - \hat{x}(\boldsymbol{\eta}))] \tag{59}$$

is the covariance function of the error in estimating $x(\mathbf{t})$ by $\hat{x}(\mathbf{t})$. Therefore, the discrete-space error pattern $\epsilon^2[\mathbf{m}]$ is associated with the continuous-space error covariance function of the vector hybrid Wiener filter.

Theorem 3 (Error of the vector hybrid Wiener filter): The error covariance (59) of the LMMSE estimate $\hat{x}(\mathbf{t})$ of $x(\mathbf{t})$

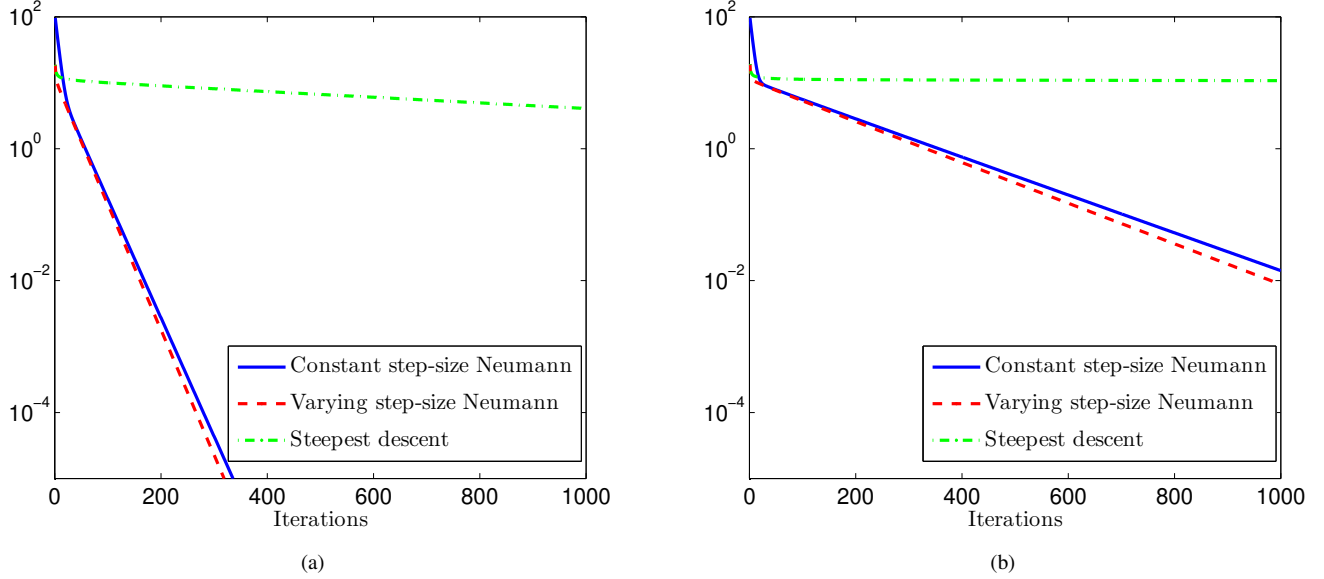


Fig. 5: RMSE as a function of iterations. (a) SNR of 0dB. (b) SNR of 8dB.

based on the measurements $y_1(\mathbf{n}), \dots, y_K(\mathbf{n})$, $\mathbf{n} \in \mathbb{Z}^d$, is given by

$$e(\mathbf{t}, \boldsymbol{\eta}) = R_{xx}(\mathbf{t} - \boldsymbol{\eta}) - \sum_{k=1}^K \sum_{\mathbf{n} \in \mathbb{Z}^d} \tilde{v}_k(\mathbf{t} - \mathbf{n}) R_{xy}^k(\boldsymbol{\eta} - \mathbf{n}), \quad (60)$$

where $\tilde{v}_1(\mathbf{t}), \dots, \tilde{v}_K(\mathbf{t})$ are given by (24).

Proof: See Appendix B. ■

Substituting (60) into (58), we have

$$\begin{aligned} \epsilon^2[\mathbf{m}] &= \int_{\mathbf{t} \in \mathbb{R}^d} \int_{\boldsymbol{\eta} \in \mathbb{R}^d} R_{xx}(\mathbf{t} - \boldsymbol{\eta}) w\left(\frac{\mathbf{m}}{\Delta} - \mathbf{t}\right) w\left(\frac{\mathbf{m}}{\Delta} - \boldsymbol{\eta}\right) d\mathbf{t} d\boldsymbol{\eta} \\ &\quad - \sum_{k=1}^K \sum_{\mathbf{n} \in \mathbb{Z}^d} \left(\int_{\mathbf{t} \in \mathbb{R}^d} \tilde{v}_k(\mathbf{t} - \mathbf{n}) w\left(\frac{\mathbf{m}}{\Delta} - \mathbf{t}\right) d\mathbf{t} \right) \\ &\quad \times \left(\int_{\mathbf{t} \in \mathbb{R}^d} R_{xy}^k(\boldsymbol{\eta} - \mathbf{n}) w\left(\frac{\mathbf{m}}{\Delta} - \boldsymbol{\eta}\right) d\boldsymbol{\eta} \right) \\ &= \int_{\boldsymbol{\eta} \in \mathbb{R}^d} (R_{xx} * w)\left(\frac{\mathbf{m}}{\Delta} - \boldsymbol{\eta}\right) w\left(\frac{\mathbf{m}}{\Delta} - \boldsymbol{\eta}\right) d\boldsymbol{\eta} \\ &\quad - \sum_{k=1}^K \sum_{\mathbf{n} \in \mathbb{Z}^d} (\tilde{v}_k * w)\left(\frac{\mathbf{m}}{\Delta} - \mathbf{n}\right) (R_{xy}^k * w)\left(\frac{\mathbf{m}}{\Delta} - \mathbf{n}\right). \end{aligned} \quad (61)$$

With a change of variables $\boldsymbol{\eta} \leftarrow \mathbf{m}/\Delta - \boldsymbol{\eta}$, and using the notation $v_k(\mathbf{t}) = (R_{xy}^k * w)(\mathbf{t})$ introduced in (30), this expression simplifies to

$$\begin{aligned} \epsilon^2[\mathbf{m}] &= \int_{\boldsymbol{\eta} \in \mathbb{R}^d} (R_{xx} * w)(\boldsymbol{\eta}) w(\boldsymbol{\eta}) d\boldsymbol{\eta} \\ &\quad - \sum_{k=1}^K \sum_{\mathbf{n} \in \mathbb{Z}^d} (\tilde{v}_k * w)\left(\frac{\mathbf{m}}{\Delta} - \mathbf{n}\right) v_k\left(\frac{\mathbf{m}}{\Delta} - \mathbf{n}\right). \end{aligned} \quad (62)$$

The first term in this expression is constant. The second term, however, is a function of \mathbf{m} . It can be seen that if the magnification factor Δ is an integer, then the error pattern $\epsilon^2[\mathbf{m}]$ is periodic with period Δ .

The above exposition is important as it reveals the fact that some pixels in $x_{\text{HR}}[\mathbf{m}]$ are harder to estimate than others. As a consequence, the recovery $\hat{x}_{\text{HR}}[\mathbf{m}]$ will typically exhibit some sort of periodic structure. This effect is very well known in single-image interpolation. The most common example is the nearest-neighbor interpolation method, which produces block-artifacts, but the effect is also apparent in bilinear interpolation, bicubic recovery, etc. [14]. As discussed in [14], the narrower the interpolation kernel is, the stronger the effect of periodic structure appearing in the restored image. In our multichannel setting, the interpolation kernels are determined by the cross-correlation functions $R_{xy}^1(\mathbf{t}), \dots, R_{xy}^K(\mathbf{t})$, which correspond to shifted versions of $R_{xx}(\mathbf{t})$ when the noise is independent of the scene. Therefore, the narrower $R_{xx}(\mathbf{t})$ is, the stronger the periodic artifacts in $\hat{x}_{\text{HR}}[\mathbf{m}]$ will be.

It is important to note that the periodic structure in $\hat{x}_{\text{HR}}[\mathbf{m}]$ can be eliminated, for example by using bandlimited interpolation kernels [14]. However, while suppressing artifacts, this would lead to a large reconstruction error. Thus, there is a tradeoff between periodic artifacts and reconstruction error. By tuning the width of $R_{xx}(\mathbf{t})$ we may produce a smooth recovery which substantially deviates from $x_{\text{HR}}[\mathbf{m}]$, or one that contains periodic artifacts but is more loyal to $x_{\text{HR}}[\mathbf{m}]$. This effect will be demonstrated in Section VIII, where we present SR results on images.

VII. SPLINE SUPER-RESOLUTION

In practical imaging scenarios, the PSD $\Gamma_{xx}(\boldsymbol{\omega})$ should be chosen to roughly match the typical frequency content of the scenes encountered in a specific application (e.g., medical images, natural images, etc.). In the single-image

interpolation literature, several authors have used the Matérn class of PSD functions [15], [16], which has the form $\Gamma_{xx}(\omega) = \alpha(\|\omega\|^2 + \beta^2)^{-(\gamma+d/2)}$ for some constants α , β and γ . It has been demonstrated that this type of prior is suitable for natural images and leads to better reconstruction results than the commonly used bilinear interpolation, Keys' bicubic interpolation [29], and Shannon's sinc interpolation. Nevertheless, the drawback of the Matérn prior is that, as shown in [30], its associated autocorrelation function $R_{xx}(t)$ is not compactly supported. This violates our Assumption 4 in Section IV.

As an alternative class of autocorrelation functions, whose frequency content decays polynomially, we now consider B-splines. A B-spline of degree N , denoted $\beta^N(t)$ is the function obtained by the $(N+1)$ -fold convolution of the unit square

$$\beta^0(t) = \begin{cases} 1 & t \in [-\frac{1}{2}, \frac{1}{2}]^d \\ 0 & \text{else.} \end{cases} \quad (63)$$

In single-image recovery, if both the PSF $s(t)$ and the autocorrelation $R_{xx}(t)$ are B-splines of some degree, then the (scalar) hybrid Wiener filter corresponds to the widely used spline interpolation technique [30]. Consequently, in our multi-channel setting, we refer to the recovery resulting from the vector hybrid Wiener filter as *spline super-resolution* (SSR). Since B-splines all have a bell-like shape whose width is controlled by their degree, modeling the PSF $s(t)$ as a B-spline is adequate for many imaging devices. In particular, if the blur due to the lens is negligible with respect to pixel size, then $s(t)$ can be quite accurately modeled as a B-spline of degree 0. Higher-order B-splines can be used when the lens' blur is more significant.

Assume that the imaging PSF $s(t)$ is a B-spline of degree D_s , the autocorrelation $R_{xx}(t)$ is a B-spline of degree D_x , and the target PSF is a delta function $w(t) = \delta(t)$. Then from (32), we see that the interpolation kernels $v_1(t), \dots, v_K(t)$ are shifted versions of B-splines of degree $D_x + D_s + 1$. Thus, the interpolation stage in our approach reduces to performing standard spline interpolation², with appropriate shift, on each of the low-resolution images, and then summing the results. This can be done efficiently, by relying on the fact that d -dimensional B-spline functions are separable and compactly supported. In particular, if $D_x = D_s = 0$, then the interpolation stage comprises K bilinear interpolation procedures.

The FIR filter $Q[n]$ needed for the digital correction stage can also be computed simply in the SSR method. Specifically, the elements of the MIMO filter $Q[n]$, appearing in Algorithms 2 and 3, involve samples of the function $(\bar{s} * s * R_{xx})(t)$. In our SSR approach, this function is a B-spline of degree $2D_s + D_x + 2$, and thus possesses a closed form expression.

As discussed in Section VI and will be demonstrated in the next section, the choice of the order D_x of the autocorrelation B-spline is guided by two conflicting desires. On one hand, a small D_x corresponds to a narrow autocorrelation function $R_{xx}(t)$, which leads to a sharp recovered image, but with strong periodic artifacts. A large degree D_x corresponds to a wide $R_{xx}(t)$, leading to a smooth recovery, but with a large

²Without the associated pre-filtering, which is used in spline interpolation.




Fig. 6: 1D interpolation using the SSP algorithm.

reconstruction error. In practice, we found that D_x should be chosen to be 1 or 2 in order to obtain satisfying results.

VIII. EXPERIMENTS

We now demonstrate our SR interpolation algorithm via simulations and experiments on real-world data.

A. 1D Super-Resolution

We begin by demonstrating the behavior of our SSR algorithm in the task of recovering 1D signals. Figure 6 depicts two 1D interpolation examples produced by Algorithm 1. Here, the PSF $s(t)$ is a B-spline of degree 0 (namely a rectangular window of width 1), the auto-correlation $R_{xx}(t)$ is a B-spline of degree 2, and the shifts are $(t_1, t_2, t_3) = (-0.1, 0, 0.1)$. The solid curve is the recovery $\hat{x}(t)$ obtained when setting $\sigma_u^2 = 0$. This reconstruction is consistent in the sense that if we were to pass $\hat{x}(t)$ through $s(t)$ and sample it, then the resulting samples would coincide with the measured sequences $c_1[n]$, $c_2[n]$, and $c_3[n]$.

The dashed curve is the estimate obtained when setting the noise variance σ_u^2 to be equal to the variance of the clean measurement sequences. In other words, the noise level is assumed to be 0dB. As can be seen, in this case $\hat{x}(t)$ tends to be smoother at the cost of deviating from the measured samples. This is a result of the fact that the samples are assumed to be contaminated with white noise, whereas the signal is known to have little energy in the high frequencies. Therefore, in this setting the high frequencies of the LMMSE estimate are attenuated with respect to the noiseless case.

B. 2D Super-Resolution

We now present several 2D experiments with the SSR algorithm. In all these experiments, the PSF $s(t)$ was taken to be a rectangular kernel, namely a B-spline of degree $D_s = 0$. We applied the Neumann method (Algorithm 2) with a constant step-size.

We begin by visualizing the general flow of our approach. Figure 7a shows (a portion of) six low-resolution images

of a pine tree, taken with an Olympus C-60 camera. The translations $\mathbf{t}_1, \dots, \mathbf{t}_6$ were estimated from the images. Here, we set the autocorrelation B-spline degree to $D_x = 0$, and tuned σ_u^2 to get an assumed SNR of 0dB. In this setting, the reconstruction kernels are B-splines of degree 1, so that the recovery stage involves applying bilinear interpolation to each of the corrected images $d_1[\mathbf{n}], \dots, d_6[\mathbf{n}]$.

A naive SR approach would be to apply bilinear interpolation to the original images $c_1[\mathbf{n}], \dots, c_6[\mathbf{n}]$ without performing the correction stage. However, as can be seen in Fig. 7b, the corrected images $d_1[\mathbf{n}], \dots, d_6[\mathbf{n}]$ are very different from the original ones. In particular, the correction process is highly dependent on the shifts between the images. Figures 8a and 8b depict respectively the restoration obtained with and without the correction stage. Evidently, the recovery without the pre-filter is blurred with respect to the one with the pre-filter.

Next, we compare the performance of the SSR algorithm to that of several prominent SR techniques, which rely on the discrete model (2). A common method for quantitative comparison between SR methods is to produce a set of low-resolution images $c_1[\mathbf{n}], \dots, c_K[\mathbf{n}]$ from a known high-resolution image $\mathbf{x}_{\text{HR}}[\mathbf{n}]$, and then measure the MSE $E[(x_{\text{HR}}[\mathbf{n}] - \hat{x}_{\text{HR}}[\mathbf{n}])^2]$ attained by each of the studied methods. However, by doing that we would actually conform to the incorrect discrete model (2). As we have seen, in real-world scenarios $c_1[\mathbf{n}], \dots, c_K[\mathbf{n}]$ cannot be produced from $x_{\text{HR}}[\mathbf{n}]$. Therefore, in the sequel we focus on qualitative comparisons on real-world images.

The discrete SR algorithms we compare against are shift-and-add (S&A) with Wiener deconvolution [31], and the methods reported in [32], namely Median S&A with Wiener deconvolution, ℓ_1 -norm regularization, ℓ_2 -norm regularization, and robust recovery (ℓ_1 -norm fidelity) with ℓ_1 -norm regularization. In all our experiments, we supplied these algorithms with the discrete-space PSF, which is the counterpart of $s(\mathbf{t}) = \beta^0(\mathbf{t})$ used in our algorithm. These algorithms are compared with the SSR approach proposed here, with autocorrelation B-spline degrees 0, 1, and 2, which we term respectively SSR0, SSR1, and SSR2. The noise level σ_u^2 was set such that the resulting SNR was 0dB, 8dB and 10dB in the SSR0, SSR1 and SSR2 methods respectively. We used the same estimate for the translations in all algorithms.

Figure 9 shows the SR recoveries obtained from the first $K = 20$ frames of the sequence Disk taken from [33]. As can be seen, the choice of the autocorrelation B-spline degree controls the tradeoff between the recovery of fine features and artifacts. Algorithm SSR1 seems to achieve a good balance between these properties, and together with the ℓ_2 -norm regularization approach, yields the best performance (note in particular the recovery of the letters ‘R’ ‘E’ and ‘F’). A similar behavior can be observed in Fig. 10, where we used the 28-frame sequence Text taken from [33], and in Fig. 11, in which the 6 images from Fig. 7a were processed.

IX. CONCLUSION

In this paper, we analyzed the validity of the widely used discrete-space formulation of the SR problem. We showed

that discrete modeling generally cannot be loyal to the actual physical setting. As a result, algorithms which do not treat the problem from a continuous-space interpolation viewpoint, inevitably suffer from a model miss-match error. This motivated us to model the scene as a continuous-space random process with known statistics, and seek its LMMSE estimate given the observed discrete-space measurements. We derived a closed form expression for the LMMSE estimator, which we coin the *vector hybrid Wiener filter*. We also provided an analytic formula for the MSE of our method, which shed light on the emergence of periodic artifacts in the recovered signal. We presented three efficient implementations of the vector hybrid Wiener filter and showed that when the signal’s autocorrelation is a B-spline function, these methods constitute a natural extension of the widely used single-channel spline interpolation to the multi-channel setup. Experiments on real-world data demonstrated the effectiveness of our proposed approach.

APPENDIX A PROOF OF THEOREM 2

The optimal interpolation kernels $\tilde{v}_1(\mathbf{t}), \dots, \tilde{v}_K(\mathbf{t})$ must satisfy the orthogonality principle

$$E[(x(\mathbf{t}) - \hat{x}(\mathbf{t}))y_k(\mathbf{m})] = 0, \quad \forall \mathbf{m} \in \mathbb{Z}^d, \forall k \in [1, K]. \quad (64)$$

Substituting (23), (64) becomes

$$\sum_{\ell=1}^K \sum_{\mathbf{n} \in \mathbb{Z}^d} R_{yy}^{k,\ell}(\mathbf{n} - \mathbf{m}) \tilde{v}_\ell(\mathbf{t} - \mathbf{n}) = R_{xy}^k(\mathbf{t} - \mathbf{m}), \quad (65)$$

for all $\mathbf{m} \in \mathbb{Z}^d$ and $k \in [1, K]$. With a change of variables $\boldsymbol{\tau} = \mathbf{t} - \mathbf{m}$ and $\mathbf{n}' = \mathbf{n} - \mathbf{m}$, (65) can be written as

$$\sum_{\ell=1}^K \sum_{\mathbf{n}' \in \mathbb{Z}^d} R_{yy}^{k,\ell}(\mathbf{n}') \tilde{v}_\ell(\boldsymbol{\tau} - \mathbf{n}') = R_{xy}^k(\boldsymbol{\tau}), \quad \forall k \in [1, K]. \quad (66)$$

Taking the CSFT of both sides, we get

$$\sum_{\ell=1}^K \tilde{V}_\ell(\boldsymbol{\omega}) \sum_{\mathbf{n} \in \mathbb{Z}^d} \Gamma_{yy}^{k,\ell}(\boldsymbol{\omega} - 2\pi\mathbf{n}) = \Gamma_{xy}^k(\boldsymbol{\omega}), \quad \forall k \in [1, K]. \quad (67)$$

Put in matrix form, (67) leads to (24).

APPENDIX B PROOF OF THEOREM 3

Using the orthogonality principle, the error (59) of the LLMSE estimate $\hat{x}(\mathbf{t})$ is given by

$$\begin{aligned} e(\mathbf{t}, \boldsymbol{\eta}) &= E[(x(\mathbf{t}) - \hat{x}(\mathbf{t}))(x(\boldsymbol{\eta}) - \hat{x}(\boldsymbol{\eta}))] \\ &= E[(x(\mathbf{t}) - \hat{x}(\mathbf{t}))x(\boldsymbol{\eta})]. \end{aligned} \quad (68)$$

Substituting the expression (23) for $\hat{x}(\mathbf{t})$, (68) becomes

$$\begin{aligned} e(\mathbf{t}, \boldsymbol{\eta}) &= E \left[\left(x(\mathbf{t}) - \sum_{k=1}^K \sum_{\mathbf{n} \in \mathbb{Z}^d} y_k(\mathbf{n}) \tilde{v}_k(\mathbf{t} - \mathbf{n}) \right) x(\boldsymbol{\eta}) \right] \\ &= R_{xx}(\mathbf{t} - \boldsymbol{\eta}) - \sum_{k=1}^K \sum_{\mathbf{n} \in \mathbb{Z}^d} \tilde{v}_k(\mathbf{t} - \mathbf{n}) R_{xy}^k(\boldsymbol{\eta} - \mathbf{n}). \end{aligned} \quad (69)$$

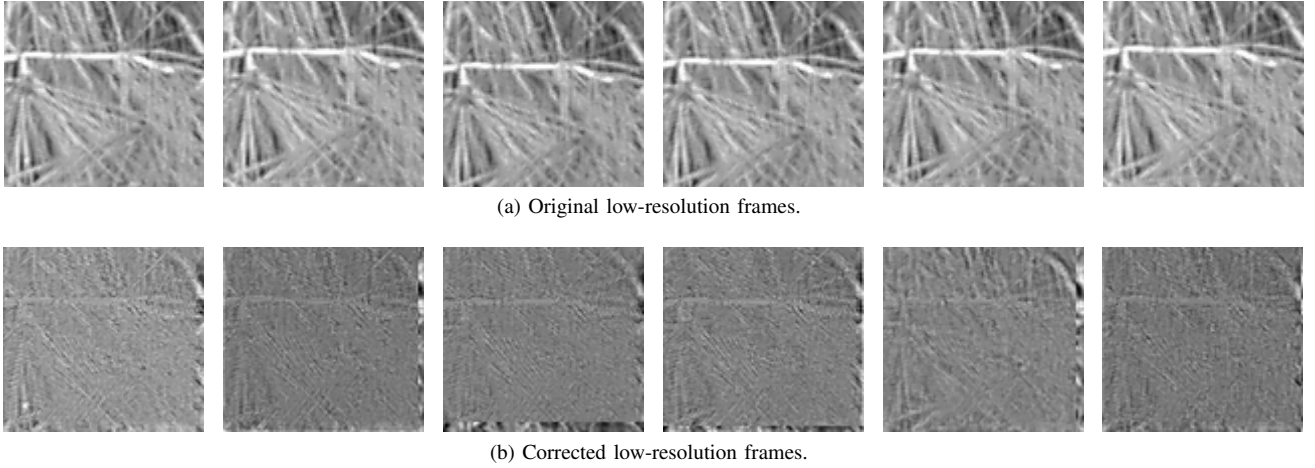


Fig. 7: (a) A set of low-resolution images $c_1[\mathbf{n}], \dots, c_6[\mathbf{n}]$, $\mathbf{n} \in \mathbb{Z}^2$. (b) The corresponding set of corrected images $d_1[\mathbf{n}], \dots, d_6[\mathbf{n}]$, $\mathbf{n} \in \mathbb{Z}^2$.




Fig. 8: Resolution enhancement by a factor of $\Delta = 3$ with the SSR method ($D_x = 0$, $\text{SNR} = 10\text{dB}$). (a) Without digital correction. (b) With digital correction.

The kernels $\tilde{v}_1(\mathbf{t}), \dots, \tilde{v}_K(\mathbf{t})$ are those minimizing the MSE, namely those given in (24).

ACKNOWLEDGMENT

The authors would like to thank Ms. Shira Avidar and Ms. Shlomit Argov for implementation and experimentation with the SSR algorithm.

REFERENCES

- [1] S. C. Park, M. K. Park, and M. G. Kang, "Super-resolution image reconstruction: A technical overview," *IEEE Signal Processing Magazine*, vol. 20, no. 3, pp. 21–36, 2003.
- [2] S. Farsiu, D. Robinson, M. Elad, and P. Milanfar, "Advances and challenges in super-resolution," *Int. J. Imaging Syst. Technol.*, vol. 14, no. 2, 2004.
- [3] E. Shechtman, Y. Caspi, and M. Irani, "Increasing space-time resolution in video," *Lecture Notes in Computer Science*, pp. 753–768, 2002.
- [4] S. Borman and R. L. Stevenson, "Super-resolution from image sequences—A review," in *1998 Midwest Symp. Circuits and Systems*, 1998, pp. 374–378.
- [5] K. Sauer and J. Allebach, "Iterative reconstruction of bandlimited images from nonuniformly spaced samples," *IEEE Trans. Circuits Syst.*, vol. 34, no. 12, pp. 1497–1506, 1987.
- [6] S. P. Kim, N. K. Bose, and H. M. Valenzuela, "Recursive reconstruction of high resolution image from noisy undersampled multiframes," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 38, no. 6, pp. 1013–1027, 1990.
- [7] S. P. Kim and W. Y. Su, "Recursive high-resolution reconstruction of blurred multiframe images," in *IEEE Int. Conf. on Acoustics, Speech and Signal Processing*, 1991, pp. 2977–2980.
- [8] K. Aizawa, T. Komatsu, and T. Saito, "A scheme for acquiring very high resolution images using multiple cameras," in *IEEE Int. Conf. on Acoustics, Speech and Signal Processing.*, vol. 3, 1992.
- [9] H. Ur and D. Gross, "Improved resolution from subpixel shifted pictures," *CVGIP: Graphical Models and Image Processing*, vol. 54, no. 2, pp. 181–186, 1992.
- [10] N. K. Bose, H. C. Kim, and H. M. Valenzuela, "Recursive total least

- squares algorithm for image reconstruction from noisy, undersampled frames,” *Multidimensional Systems and Signal Processing*, vol. 4, no. 3, pp. 253–268, 1993.
- [11] H. Shekarforoush and R. Chellappa, “Data-driven multichannel super-resolution with application to video sequences,” *J. Opt. Soc. Am. A*, vol. 16, no. 3, pp. 481–492, 1999.
 - [12] A. M. Tekalp, M. K. Ozkan, and M. I. Sezan, “High-resolution image reconstruction from lower-resolution image sequences and space-varying image restoration,” in *IEEE Int. Conf. on Acoustics, Speech and Signal Processing*, vol. 3, 1992.
 - [13] S. Ramani, D. Van De Ville, T. Blu, and M. Unser, “Nonideal sampling and regularization theory,” *IEEE Transactions on Signal Processing*, vol. 56, no. 3, p. 1055, 2008.
 - [14] Y. C. Eldar and T. Michaeli, “Beyond bandlimited sampling,” *IEEE Signal Process. Mag.*, no. 3, pp. 48–68, May 2009.
 - [15] C. A. Glasbey, “Optimal linear interpolation of images with known point spread function,” *Scandinavian Image Analysis Conference–SCIA-2001*, pp. 161–168, 2001.
 - [16] S. Ramani, D. Van De Ville, and M. Unser, “Non-ideal sampling and adapted reconstruction using the stochastic Matérn model,” *Proc. Int. Conf. Acoust., Speech, Signal Processing*, vol. 2, 2006.
 - [17] T. Michaeli and Y. C. Eldar, “High rate interpolation of random signals from nonideal samples,” *IEEE Transactions on Signal Processing*, vol. 57, pp. 977–992, March 2009.
 - [18] M. Irani and S. Peleg, “Improving resolution by image registration,” *CVGIP: Graphical Models and Image Processing*, vol. 53, no. 3, pp. 231–239, 1991.
 - [19] P. Cheeseman, B. Kanefsky, R. Kraft, J. Stutz, and R. Hanson, “Super-resolved surface reconstruction from multiple images,” Tech. Rep., December FIA-94-12, NASA, Ames Research Center, Moffett Field, CA, 1994.
 - [20] R. C. Hardie, K. J. Barnard, and E. E. Armstrong, “Joint MAP registration and high-resolution image estimation using a sequence of undersampled images,” *IEEE Transactions on Image Processing*, vol. 6, no. 12, pp. 1621–1633, 1997.
 - [21] R. R. Schultz and R. L. Stevenson, “Extraction of high-resolution frames from video sequences,” *IEEE Transactions on Image Processing*, vol. 5, no. 6, pp. 996–1011, 1996.
 - [22] M. B. Matthews, “On the linear minimum-mean-squared-error estimation of an undersampled wide-sense stationary random process,” *IEEE Trans. Signal Process.*, vol. 48, no. 1, pp. 272–275, 2000.
 - [23] V. Kucera, “Factorization of rational spectral matrices: a survey of methods,” in *Proc. IEE Int. Conf. Contr., Edinburgh, U.K.*, 1991, pp. 1074–1078.
 - [24] M. Unser, A. Aldroubi, and M. Eden, “Fast B-spline transforms for continuous image representation and interpolation,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 13, no. 3, pp. 277–285, 1991.
 - [25] E. E. Tyrtshnikov and E. E. Tyrtshnikov, *A brief introduction to numerical analysis*. Birkhauser, 1997.
 - [26] J. Nocedal and S. J. Wright, *Numerical optimization*. Springer, 2000.
 - [27] S. Baker and T. Kanade, “Limits on super-resolution and how to break them,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1167–1183, 2002.
 - [28] D. Robinson and P. Milanfar, “Statistical performance analysis of super-resolution,” *IEEE Transactions on Image Processing*, vol. 15, no. 6, pp. 1413–1428, 2006.
 - [29] R. G. Keys, “Cubic convolution interpolation for digital image processing,” *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 29, no. 6, pp. 1153–1160, 1981.
 - [30] S. Ramani, D. Van De Ville, T. Blu, and M. Unser, “Nonideal Sampling and Regularization Theory,” *IEEE Trans. Signal Processing*, vol. 56, no. 3, pp. 1055–1070, 2008.
 - [31] M. Elad and Y. Hel-Or, “A fast super-resolution reconstruction algorithm for pure translation motion and common space-invariant blur,” *IEEE Trans. Image Process.*, vol. 10, no. 8, pp. 1187–1193, 2001.
 - [32] S. Farsiu, M. D. Robinson, M. Elad, and P. Milanfar, “Fast and robust multiframe super resolution,” *IEEE Trans. Image Process.*, vol. 13, no. 10, pp. 1327–1344, 2004.
 - [33] “<http://www.soe.ucsc.edu/~milanfar/software/sr-datasets.html>.”



(a) Low-resolution image



(b) S&A + Wiener deconvolution



(c) Median S&A + Wiener deconvolution

(d) ℓ_1 regularization(e) ℓ_2 regularization(f) ℓ_1 fidelity with ℓ_1 regularization

(g) SSR0



(h) SSR1



(i) SSR2

Fig. 9: Resolution enhancement by a factor of $\Delta = 4$ from a set of 20 images.



Fig. 10: Resolution enhancement by a factor of $\Delta = 4$ from a set of 28 images.




Fig. 11: Resolution enhancement by a factor of $\Delta = 4$ from a set of 6 images.