

Research Article

A Fast Algorithm for Selective Signal Extrapolation with Arbitrary Basis Functions

Jürgen Seiler (EURASIP Member) and André Kaup (EURASIP Member)

Chair of Multimedia Communications and Signal Processing, University of Erlangen-Nuremberg, Cauerstraße 7, 91058 Erlangen, Germany

Correspondence should be addressed to Jürgen Seiler, seiler@int.de

Received 7 July 2010; Revised 1 December 2010; Accepted 18 January 2011

Academic Editor: Ana Pérez-Neira

Copyright © 2011 J. Seiler and A. Kaup. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Signal extrapolation is an important task in digital signal processing for extending known signals into unknown areas. The Selective Extrapolation is a very effective algorithm to achieve this. Thereby, the extrapolation is obtained by generating a model of the signal to be extrapolated as weighted superposition of basis functions. Unfortunately, this algorithm is computationally very expensive and, up to now, efficient implementations exist only for basis function sets that emanate from discrete transforms. Within the scope of this contribution, a novel efficient solution for Selective Extrapolation is presented for utilization with arbitrary basis functions. The proposed algorithm mathematically behaves identically to the original Selective Extrapolation but is several decades faster. Furthermore, it is able to outperform existent fast transform domain algorithms which are limited to basis function sets that belong to the corresponding transform. With that, the novel algorithm allows for an efficient use of arbitrary basis functions, even if they are only numerically defined.

1. Introduction

The extrapolation of signals is a very important area in digital signal processing, especially in image and video signal processing. Thereby, unknown or not accessible samples are estimated from known surrounding samples. In image and video processing, signal extrapolation tasks arise for example, in the area of concealment of transmission errors as described in [1] or for prediction in hybrid video coding as shown in [2].

In general, signal extrapolation can be regarded as underdetermined problem as there are infinitely many different solutions for the signal to be estimated, based on the known samples. According to [3], sparsity-based algorithms are well suited for solving underdetermined problems as these algorithms are able to cover important signal characteristics, even if the underlying problem is underdetermined. These algorithms can be applied well to image and video signals, as in general natural signals are sparse [4] in certain domains, meaning that they can be described by only few coefficients.

As has been shown in [5, 6], out of the group of sparse algorithms the greedy sparse algorithms are of interest, as these algorithms are able to robustly solve the problem. One algorithm out of this group is for example, Matching Pursuits from [7]. Another powerful greedy sparse algorithm is the Selective Extrapolation (SE) from [8]. SE iteratively generates a model of the signal to be extrapolated as weighted superposition of basis functions. In the past years, this extrapolation algorithm also has been adopted by several others like [9, 10] to solve extrapolation problems in their contexts.

Unfortunately, SE as it exists up to now is computationally very expensive. This holds except for the case that basis function sets are regarded, that emanate from discrete transforms. In such a case, the algorithm can be efficiently carried out in the transform domain. The functions of the Discrete Fourier Transform (DFT) [11] are one example for such a basis function set. Using this set, an efficient implementation in the Fourier domain exists by Frequency Selective Extrapolation (FSE) [8]. If basis function sets that

do not emanate from discrete transforms or overcomplete basis function sets or even only numerically defined basis functions are regarded, such transform domain algorithms cannot exist.

Although Fourier basis functions have proven to form a good set for a wide range of signals, there also exist signals where other basis function sets lead to better extrapolation results. This may happen when the support area on which the extrapolation is based is very unequal or in the case when very steep signal changes occur as for example, in artificial signals. Figure 1 shows three examples for such signals. The left column shows the original signal, the second column shows a distorted signal with the area to be extrapolated marked in black. The signals in the third column result from applying FSE which utilizes Fourier basis functions. In the last column, Selective Extrapolation is carried out with different basis function sets. In the first row, the basis function set results from the union of the functions from the Discrete Cosine Transform (DCT) [12] and the Walsh-Hadamard Transform (WHT) [13]. In the second row, a binarized version of DFT functions is used in order to reconstruct the steep changes in this artificial signal. In the third row, the basis function set emanates from the union of DFT functions and binarized DFT functions. The three examples have in common that the used basis function sets produce significantly better subjective as well as objective results than the Fourier-based extrapolation does. But they have also in common that for such sets no efficient transform domain implementation can exist which would be necessary for a fast implementation.

Within the scope of this contribution we want to introduce a novel spatial domain solution for SE which is called Fast Selective Extrapolation (FaSE). This algorithm is able to generate a model of the signal for arbitrary basis functions in the same way as the original SE, even in the case that the basis function set does not possess any structure and the basis functions are only numerically defined or in the case that an overcomplete basis function set is regarded. But at the same time, the algorithm is very fast as it can efficiently trade computational complexity versus memory consumption. The paper is organized as follows: first, SE will be reviewed for the general case of complex-valued basis functions. With that, an overview of the algorithm is given and the computationally most expensive steps are pointed out. After that, the novel Fast Selective Extrapolation is presented in detail and its complexity is compared to SE. Finally, simulation results are given for proving the abilities of the novel algorithm.

2. Review of Selective Extrapolation

For the presentation of Selective Extrapolation (SE), a scenario as shown in Figure 2 is regarded. There, signal parts which have to be extrapolated are subsumed in loss area \mathcal{B} . For extrapolating the signal, surrounding correctly received signal parts are used. These signal parts form the support area \mathcal{A} . The two areas together form the so-called extrapolation area \mathcal{L} which is of size $M \times N$ samples and

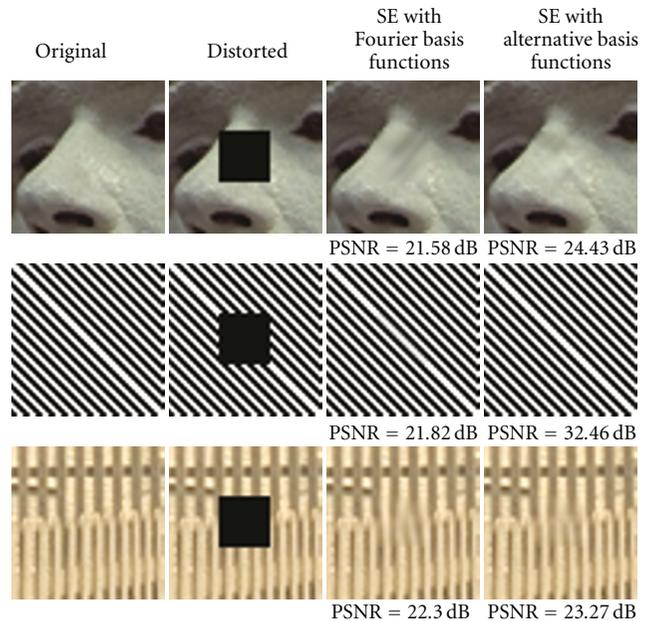


FIGURE 1: Examples for image signals where Fourier basis functions provide insufficient extrapolation quality. In every row, original signal, distorted signal, and extrapolated signals are shown. Extrapolation is carried out either with DFT basis functions or alternative ones. Top row: union of DCT and WHT basis functions. Mid row: binarized DFT basis functions. Bottom row: union of DFT and binarized DFT basis functions.

is depicted by the spatial coordinates m and n . The signal in \mathcal{L} is denoted by $s[m, n]$ but is only available in the support area \mathcal{A} . The extrapolation of square blocks is used for presentational reasons at this point only. In general, arbitrarily shaped regions can be extrapolated. In addition to that, in general, the used basis functions can as well be larger than the regarded extrapolation area. In such a case, the extrapolation area has to be padded with zeros to be of the same size as the basis functions. But, for presentational reasons we also assume that the extrapolation area and the basis functions have the same size subsequently.

As described in [8], SE aims at generating a parametric model $g[m, n]$ for signal $s[m, n]$ in whole area \mathcal{L} . The model

$$g[m, n] = \sum_{k \in \mathfrak{K}} \hat{c}_k \varphi_k[m, n] \quad (1)$$

emanates from a weighted superposition of the basis functions $\varphi_k[m, n]$ which are defined over complete \mathcal{L} and are indexed by k . Set \mathfrak{K} contains the indices of all basis functions used for model generation. As not all possible basis functions are used for the model, set \mathfrak{K} is a subset of dictionary \mathcal{D} which holds all basis functions. In order to control the weights of the individual basis functions, one expansion coefficient \hat{c}_k is assigned to each basis function $\varphi_k[m, n]$. The challenge is to determine which basis functions to use for the model and to calculate the corresponding weights. SE solves this problem iteratively, at which in every iteration one basis function is selected and the corresponding weight is estimated.

This is achieved by successively approximating signal $s[m, n]$ in support area \mathcal{A} and identifying the dominant basis functions of the signal. In doing so, the signal can be continued well into area \mathcal{B} if an appropriate set of basis functions is used.

Initially, model $g^{(0)}[m, n]$ is set to zero and with that the initial approximation residual

$$r^{(0)}[m, n] = s[m, n] \quad (2)$$

is equal to the original signal. At the beginning of each iteration, in general the ν -th iteration, a weighted projection of the residual onto each basis function is conducted. For every basis function, this leads to the projection coefficient

$$p_k^{(\nu)} = \frac{\sum_{(m,n) \in \mathcal{L}} r^{(\nu-1)}[m, n] \varphi_k^*[m, n] w[m, n]}{\sum_{(m,n) \in \mathcal{L}} \varphi_k^*[m, n] w[m, n] \varphi_k[m, n]}, \quad \forall k, \quad (3)$$

which results from the quotient of the weighted scalar product between the residual and the basis function and the weighted scalar product between the basis function and itself. In this context, the weighting function

$$w[m, n] = \begin{cases} \rho[m, n], & \text{for } (m, n) \in \mathcal{A} \\ 0, & \text{for } (m, n) \in \mathcal{B} \end{cases} \quad (4)$$

has two tasks. Firstly, it is used to mask area \mathcal{B} from the calculation of the scalar product as there is no information available about the signal. Secondly, using the function $\rho[m, n]$ can control the influence different samples have on the model generation depending on their position. For instance, samples far away from loss area \mathcal{B} can get a smaller weight and due to this weaker influence on the model generation compared to the samples close to area \mathcal{B} . In [14], an exponentially decreasing weight

$$\rho[m, n] = \hat{\rho}^{\sqrt{(m - ((M-1)/2))^2 + (n - ((N-1)/2))^2}} \quad (5)$$

is proposed with $\hat{\rho}$ controlling the decay.

After the projection coefficients have been calculated for all basis functions, one basis function has to be selected to be added to the model in the actual iteration. The choice falls on the basis function that minimizes the weighted distance

$$e_k^{(\nu)} = \sum_{(m,n) \in \mathcal{L}} \left(\left| r^{(\nu-1)}[m, n] - p_k^{(\nu)} \varphi_k[m, n] \right|^2 w[m, n] \right) \quad (6)$$

between the approximation residual $r^{(\nu-1)}[m, n]$ and the projection $p_k^{(\nu)} \varphi_k[m, n]$ onto the according basis function. In this process, again weighting function $w[m, n]$ from above is used. Hence, the index $u^{(\nu)}$ of the basis function to be added in the ν -th iteration is

$$\begin{aligned} u^{(\nu)} &= \arg \min_k \left(e_k^{(\nu)} \right) \\ &= \arg \max_k \left(\left| p_k^{(\nu)} \right|^2 \sum_{(m,n) \in \mathcal{L}} \varphi_k^*[m, n] w[m, n] \varphi_k[m, n] \right). \end{aligned} \quad (7)$$

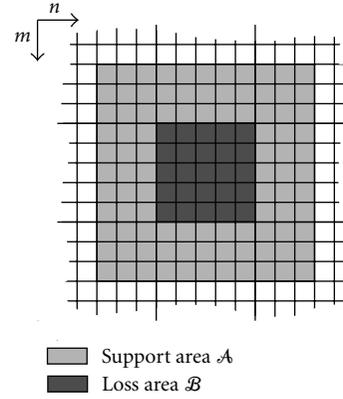


FIGURE 2: Extrapolation area \mathcal{L} consisting of loss area \mathcal{B} and support area \mathcal{A} .

Subsequent to the basis function selection, the corresponding weight has to be determined. In this process, it has to be noted that although the basis functions may have been orthogonal with respect to the complete extrapolation area \mathcal{L} they cannot be anymore if the scalar products are evaluated in combination with the required weighting function. This effect which has not been considered in the original paper from [8] is called orthogonality deficiency and is described in detail in [15]. In [16], fast orthogonality deficiency compensation is proposed to efficiently estimate the expansion coefficient by taking only the fraction γ of the projection coefficient:

$$\hat{c}_{u^{(\nu)}} = \gamma \cdot p_{u^{(\nu)}}^{(\nu)}. \quad (8)$$

The factor γ is between zero and one and depends on the extrapolation scenario, as described in detail in [16].

After one basis function has been selected and the corresponding weight has been determined, the model and the residual have to be updated by adding the selected basis function to the model generated so far:

$$g^{(\nu)}[m, n] = g^{(\nu-1)}[m, n] + \hat{c}_{u^{(\nu)}} \varphi_{u^{(\nu)}}[m, n]. \quad (9)$$

The approximation residual can be updated in the same way and results in

$$r^{(\nu)}[m, n] = r^{(\nu-1)}[m, n] - \hat{c}_{u^{(\nu)}} \varphi_{u^{(\nu)}}[m, n]. \quad (10)$$

The above described iterations are repeated until the pre-defined number of I iterations is reached. Finally, area \mathcal{B} is cut out of the model and is used for replacing the lost signal.

Algorithm 1 shows the pseudocode of SE for giving a compact overview of this algorithm. Regarding this code and taking into account the equations above, the weighted projection onto all the basis functions in every iteration can be identified as computationally most expensive step. To obtain the projection, a weighted scalar product between the residual and every basis function has to be carried out, leading to a large number of multiplications and additions. Compared to this, the actual basis function selection, the expansion coefficient estimation, and the model and residual update have a very small complexity.

```

input: distorted signal  $s[m, n]$ , weighting function  $w[m, n]$ , basis
functions  $\varphi_k[m, n]$ 
/* Initial residual is equal to original signal */
 $r[m, n] = s[m, n], \forall (m, n)$ 
for all  $\nu = 1, \dots, I$  do
  /* Projection onto basis functions */
  for all  $k = 0, \dots, |\mathcal{D}| - 1$  do
     $p_k = \frac{\sum_{(m,n) \in \mathcal{L}} r[m, n] \varphi_k^*[m, n] w[m, n]}{\sum_{(m,n) \in \mathcal{L}} \varphi_k^*[m, n] w[m, n] \varphi_k[m, n]}$ 
  end for
  /* Basis function selection */
   $u = \arg \max_k (|p_k|^2 \sum_{(m,n) \in \mathcal{L}} \varphi_k^*[m, n] w[m, n] \varphi_k[m, n])$ 
  /* Expansion coefficient estimation */
   $\hat{c} = \gamma p_u$ 
  /* Model and residual update */
   $g[m, n] = g[m, n] + \hat{c} \varphi_u[m, n], \forall (m, n)$ 
   $r[m, n] = r[m, n] - \hat{c} \varphi_u[m, n], \forall (m, n)$ 
end for
/* Replace distorted signal parts */
for all  $(m, n) \in \mathcal{B}$  do
   $s[m, n] = g[m, n]$ 
end for
output: extrapolated signal  $s[m, n]$ 

```

ALGORITHM 1: Selective Extrapolation for arbitrary basis functions.

3. Fast Selective Extrapolation

In order to solve the dilemma of the huge computational complexity of SE, we propose a novel formulation of this algorithm that also operates in the spatial domain but is as fast as transform domain algorithms which have been mentioned at the beginning. With that, the advantages of both approaches are combined: the high speed of transform domain algorithms and the independence from certain basis function sets, offered by the spatial domain SE algorithm. The high speed of the novel algorithm results from the fact that the weighted scalar products only have to be evaluated once, prior to the first iteration. In the successive iterations they can be replaced by a recursive calculation. The novel algorithm is called Fast Selective Extrapolation (FaSE) and is outlined in detail for the general complex-valued scenario subsequently. If only real-valued signals and basis functions are regarded, the conjugate complex operations can just be discarded.

Although the principal behavior of FaSE is similar to SE, the residual $r[m, n]$ in the spatial domain is not regarded, but rather the weighted scalar products between the residual and the basis functions. This yields

$$R_k^{(\nu)} = \sum_{(m,n) \in \mathcal{L}} r^{(\nu)}[m, n] \varphi_k^*[m, n] w[m, n], \quad \forall k \quad (11)$$

for depicting the weighted scalar product between the residual and the basis function with index k in the ν -th iteration. This scalar product has to be evaluated only once

explicitly. This has to be done for the initial step, where the residual is equal to the original signal, leading to

$$R_k^{(0)} = \sum_{(m,n) \in \mathcal{L}} s[m, n] \varphi_k^*[m, n] w[m, n], \quad \forall k. \quad (12)$$

After the initial $R_k^{(0)}$ have been determined, all subsequent calculations can be carried out with respect to the weighted scalar products and no explicit evaluation of the scalar products is necessary anymore.

Using $R_k^{(\nu)}$ and exploiting the fact that the square root is a monotonic increasing function for positive arguments, the basis function selection from (7) can be simplified to

$$u^{(\nu)} = \arg \max_k \frac{|R_k^{(\nu-1)}|}{\sqrt{\sum_{(m,n) \in \mathcal{L}} \varphi_k^*[m, n] w[m, n] \varphi_k[m, n]}}. \quad (13)$$

Using expression $R_{u^{(\nu)}}^{(\nu-1)}$ for the weighted scalar product between the selected basis function and the residual from the previous iteration, the estimate for the expansion coefficient results in

$$\hat{c}_{u^{(\nu)}} = \gamma \frac{R_{u^{(\nu)}}^{(\nu-1)}}{\sum_{(m,n) \in \mathcal{L}} \varphi_{u^{(\nu)}}^*[m, n] w[m, n] \varphi_{u^{(\nu)}}[m, n]}. \quad (14)$$

Here, again fast orthogonality deficiency compensation is used to derive the estimate for the expansion coefficient from the projection coefficient. Finally, the update of the model in every iteration can be carried out according to (9).

For the subsequent iterations, the weighted scalar products can be updated by applying definition (11) on the residual update from (10), yielding

$$\begin{aligned} R_k^{(y)} &= \sum_{(m,n) \in \mathcal{L}} \left(r^{(y-1)}[m,n] - \hat{c}_{u^{(y)}} \varphi_{u^{(y)}}[m,n] \right) \\ &\quad \times \varphi_k^*[m,n] w[m,n] \\ &= R_k^{(y-1)} - \hat{c}_{u^{(y)}} \sum_{(m,n) \in \mathcal{L}} \varphi_k^*[m,n] w[m,n] \varphi_{u^{(y)}}[m,n]. \end{aligned} \quad (15)$$

Obviously, the weighted scalar product between the residual and a certain basis function can be easily updated from one iteration to the other by subtracting the weighted scalar product between the actual basis function and the selected one, further weighted by the estimated expansion coefficient. Since the update only incorporates the weighted scalar product between two basis functions and is independent of the actual residual, it can be carried out very fast by calculating the different weighted scalar products of all basis functions in advance.

This novel formulation of the SE algorithm has two advantages. First of all, the residual now does not have to be calculated explicitly in every iteration step anymore, and the weighted scalar products between the residual and the basis functions are updated. But more important is the fact that the most complex calculations can be carried out in advance and can be tabulated. Namely, these are the weighted scalar products between every two basis functions and one over the square root of the weighted scalar product between a basis function and itself. This leads to the matrix

$$C_{(k,l)} = \sum_{(m,n) \in \mathcal{L}} \varphi_k^*[m,n] w[m,n] \varphi_l[m,n], \quad \forall k, l \quad (16)$$

containing the weighted scalar products between every two basis functions and the vector

$$D_k = \frac{1}{\sqrt{\sum_{(m,n) \in \mathcal{L}} \varphi_k^*[m,n] w[m,n] \varphi_k[m,n]}} = \frac{1}{\sqrt{C_{(k,k)}}}, \quad \forall k \quad (17)$$

holding the inverse of the square root of the weighted scalar products. Obviously, $C_{(k,l)}$ and D_k are independent of the input signal and the residual. Hence, they only have to be calculated once and do not have to be calculated for every extrapolation process. Thus, they can either be computed at the beginning of the extrapolation process or read from storage. During the whole computation, they are kept in memory. Furthermore, as $C_{(k,l)}$ is of size $|\mathcal{D}|^2$ and D_k has length $|\mathcal{D}|$, the memory consumption is manageable without any problems. Here, the expression $|\mathcal{D}|$ expresses the cardinality of dictionary \mathcal{D} that contains all possible basis functions. Regarding the two equations above, one can see that they both depend on the weighting function. If different weighting functions are used, $C_{(k,l)}$ and D_k have to be adapted according to the weighting function. But, regarding typical signal extrapolation tasks as for example, error concealment

or prediction, the same patterns or only a small number of different patterns occur. Therefore, this also is not a problem, as $C_{(k,l)}$ and D_k can be calculated for the different patterns in advance as well. During the generation of $C_{(k,l)}$, the complex symmetry of this matrix can be exploited and only $(|\mathcal{D}|^2 + |\mathcal{D}|)/2$ weighted scalar products have to be actually calculated.

Using these precalculated and tabulated values, the basis function selection from (13) can be rewritten as

$$u^{(y)} = \arg \max_k \left| R_k^{(y-1)} \right| \cdot D_k. \quad (18)$$

In addition to that, the estimation of the expansion coefficient from (14) can also be expressed very compactly by

$$\hat{c}_{u^{(y)}} = \gamma R_{u^{(y)}}^{(y-1)} D_{u^{(y)}}^2. \quad (19)$$

Furthermore, the update of the weighted scalar products between the residual and all possible basis functions from (15) can also be formulated very efficiently by

$$R_k^{(y)} = R_k^{(y-1)} - \hat{c}_{u^{(y)}} C_{(k,u^{(y)})}, \quad \forall k. \quad (20)$$

Regarding the three equations above, one can recognize that instead of evaluating the weighted scalar products in every iteration step explicitly, only one value has to be read from memory for every calculation. Thus, the very high computational load from the original spatial domain SE is traded against an increased memory consumption. But as the memory consumption still is easily manageable this is a quite reasonable exchange.

The novel FaSE implementation has the further advantage that no divisions are required. With that, this implementation is suited more for fixed point or integer implementations than the original SE. In such a scenario, D_k could be calculated with high accuracy and then quantized to integer or fixed point values. Thus, no expensive divisions have to be carried out within the iteration loop and the effect of error propagation due to a restricted word length can be reduced. Depending on the architecture on which the extrapolation is carried out and the regarded application, it may be preferable to store $1/C_{(k,k)}$ instead of $1/\sqrt{C_{(k,k)}}$ and to calculate $|\cdot|^2$ instead of $|\cdot|$. By using this modification, the complexity could be reduced a little bit more if the platform on which the extrapolation runs directly supports the relevant operations. Nevertheless, at this point a sufficiently high computational accuracy is assumed for the above outlined calculations. For a hardware implementation or an implementation on a digital signal processor, finite-word length effects have to be considered and further research is necessary for determining the required bit-depth of the tables and the impact of fixed-point arithmetic.

In order to give a final overview of FaSE, Algorithms 2 and 3 show the pseudo code for generating the tabulated values and for the actual model generation. The table generation is separated from the model generation for emphasizing again that the generation of the tables only has to be carried out once. Regarding the operations that have to be carried out within the iteration loop, one can recognize

```

input: basis functions  $\varphi_k[m, n]$ , weighting function  $w[m, n]$ 
for all  $k = 0, \dots, |\mathcal{D}| - 1$  do
  for all  $l = k, \dots, |\mathcal{D}| - 1$  do
     $C_{(k,l)} = \sum_{(m,n) \in \mathcal{L}} \varphi_k^*[m, n] w[m, n] \varphi_l[m, n]$ 
     $C_{(l,k)} = C_{(k,l)}^*$ 
  end for
   $D_k = \frac{1}{\sqrt{C_{(k,k)}}}$ 
end for
output: tabulated values  $C_{(k,l)}$  and  $D_k$ 

```

ALGORITHM 2: Generation of the tabulated values $C_{(k,l)}$ and D_k .

that only very simple operations have to be performed which can furthermore be processed very fast. The only computational expensive operation is the initial calculation of $R_k^{(0)}$, but compared to the original SE, this complex step has to be carried out only once instead of in every iteration.

4. Complexity Evaluation

Regarding the two previous sections, one can recognize that FaSE is able to outperform the original SE since the computational complexity within the iteration loop is reduced and since as many calculations as possible are carried out in advance and are tabulated. To quantify the complexity of SE and FaSE, the number of operations is regarded that is necessary for generating the model by each of the algorithms. In Table 1 for SE, FaSE, and the table generation for FaSE, the number of operations is listed, depending on the extent M, N of extrapolation area \mathcal{L} , dictionary size $|\mathcal{D}|$, and the number of iterations I to be carried out. Here, the operations are separated into three groups, the number of multiplications (MUL), the number of additions (ADD), and the number of other operations (OTHER) like divisions, comparisons or the calculation of a square root. As the general case of complex-valued signals and basis functions is regarded, MUL and ADD describe complex-valued multiplications and additions. For presentational reasons, a further separation of these operations into real-valued operations is omitted.

The computationally most expensive step in SE is the projection onto the basis functions. For the weighted projection of the residual onto a single basis function, $4MN$ complex-valued multiplications, $2MN$ additions and one division are required. Since SE has to project the residual in each iteration onto every of the $|\mathcal{D}|$ basis functions, these numbers have to be further multiplied by $I \cdot |\mathcal{D}|$. For selecting the basis function to be added, in every iteration $(2MN + 1)|\mathcal{D}|$ multiplications, $MN|\mathcal{D}|$ additions and $|\mathcal{D}|$ comparisons and absolute value calculations are required and the model and residual update consumes $2MN + 1$ multiplications and $2MN$ additions. Due to this, the overall complexity of SE with respect to the number of iterations is proportional to $\mathcal{O}(I \cdot MN \cdot |\mathcal{D}|)$. In contrast to this, FaSE has to evaluate the weighted scalar product between the input signal and the basis functions only once, prior

TABLE 1: Number of required operations for model generation by SE and FaSE and for generating the tables.

	SE
MUL	$I \cdot (6MN \cdot \mathcal{D} + \mathcal{D} + 2MN + 1)$
ADD	$I \cdot (3MN \cdot \mathcal{D} + 2MN)$
OTHER	$3I \cdot \mathcal{D} $
	FaSE
MUL	$2MN \cdot \mathcal{D} + I \cdot (2 \mathcal{D} + MN + 1)$
ADD	$MN \cdot \mathcal{D} + I \cdot (\mathcal{D} + MN)$
OTHER	$2I \cdot \mathcal{D} $
	FaSE table generation
MUL	$(\mathcal{D} ^2 + \mathcal{D}) \cdot MN$
ADD	$(\mathcal{D} ^2 + \mathcal{D}) \cdot MN/2$
OTHER	$(\mathcal{D} ^2 + \mathcal{D}) \cdot MN/2 + \mathcal{D} $

to the iterations. This calculation requires only $2MN \cdot |\mathcal{D}|$ complex-valued multiplications and $MN \cdot |\mathcal{D}|$ additions. Within every iteration, only $2|\mathcal{D}| + MN + 1$ complex-valued multiplications, $|\mathcal{D}| + MN$ additions, and $|\mathcal{D}|$ comparisons and absolute value calculations have to be carried out. As $|\mathcal{D}|$ and MN are of the same magnitude, the computational complexity of FaSE increases proportional to $\mathcal{O}(I \cdot |\mathcal{D}|)$ with respect to the number of iterations. For generating the tables, one has to consider that the weighted scalar products between every two basis functions have to be evaluated, resulting in a complexity that is proportional to $\mathcal{O}(MN \cdot |\mathcal{D}|^2)$, as shown in Table 1.

Figure 3 shows the number of operations with respect to the number of iterations for $M = N = 64$ and $|\mathcal{D}| = 4096$. This plot only shows the overall number of operations, that is, the sum of MUL, ADD, and OTHER, in order to give a rough impression of the overall complexity and compare the different algorithms. The fact that complex operations like divisions require more processing time than a simple multiplication is omitted for this plot. It can be easily recognized that the number of operations that is necessary for generating the model by SE is several decades larger than for FaSE. The plot further shows the number of operations that is required for generating the tabulated $C_{(k,l)}$ and D_k , indicated by a rhomb. In addition to that, the number of operations for the table generation is displayed as dashed line over the complete iteration range. It has to be noted that the table generation is independent of the iterations and this illustration is only chosen for comparing the complexity of the table generation with SE. Therewith, it can be recognized that the table generation requires roughly as many operations as 1000 iterations of SE would require. Since the number of iterations for generating the model can easily reach values larger than 200 as has been shown in [16], the expense for generating the tables amortize even after a small number of blocks. Taking into account that in typical scenarios a large number of blocks are extrapolated with the same weighting function, the complexity for generating the tables very soon becomes negligible.

```

input: distorted signal  $s[m, n]$ , weighting function  $w[m, n]$ , basis
functions  $\varphi_k[m, n]$ , tabulated values  $C_{(k,l)}$  and  $D_k$ 
/* Calculation of the initial weighted scalar product */
for all  $k = 0, \dots, |\mathcal{D}| - 1$  do
     $R_k = \sum_{(m,n) \in \mathcal{L}} s[m, n] \varphi_k^*[m, n] w[m, n]$ 
end for
for all  $\nu = 1, \dots, I$  do
    /* Basis function selection */
     $u = \arg \max_k |R_k| D_k$ 
    /* Expansion coefficient estimation */
     $\hat{c} = \gamma R_u D_u^2$ 
    /* Model update */
     $g[m, n] = g[m, n] + \hat{c} \varphi_u[m, n], \quad \forall (m, n)$ 
    for all  $k = 0, \dots, |\mathcal{D}| - 1$  do
         $R_k = R_k - \hat{c} C_{(k,u)}$ 
    end for
end for
/* Replace distorted signal parts */
for all  $(m, n) \in \mathcal{B}$  do
     $s[m, n] = g[m, n]$ 
end for
output: extrapolated signal  $s[m, n]$ 
    
```

ALGORITHM 3: Fast Selective Extrapolation for arbitrary basis functions.

5. Results for Arbitrary Basis Functions

In order to support the complexity evaluation from the previous section, the processing time for SE and FaSE is further examined. The first results presented are for arbitrary two-dimensional basis functions. In this case, only the original SE and the novel FaSE can be used, as transform domain algorithms like FSE cannot deal with arbitrary basis functions. For the runtime evaluation, the model generation has been implemented in C, compiled with gcc 4.3.2 and optimizations -O3, and the simulations have been carried out on an Intel Core2 Quad @ 2.83 GHz, equipped with 8 GB RAM. In order to reduce the influence from the operating system, multiple runs of the simulations have been conducted and the computation has been limited to the usage of only one single core.

For the simulations, a block of size 16×16 samples is extrapolated from its surrounding samples. Furthermore, different sizes of extrapolation area \mathcal{L} between 48×48 and 96×96 samples are regarded. Figure 4 shows the extrapolation time per block for different numbers of candidate basis functions and for 250 iterations performed for model generation. For this plot, the cardinality of the dictionary is selected to be of the same size as the extrapolation area. Thus, \mathcal{D} varies between $|\mathcal{D}| = 2304$ basis functions of size 48×48 and $|\mathcal{D}| = 9216$ basis functions of size 96×96 . Comparing the two curves of SE and FaSE, one can easily recognize that FaSE is about 250 times faster than the original SE, independently of the problem size. This is due to the fact that for FaSE the computationally expensive weighted scalar products only have to be evaluated once, namely, prior to the first iteration. In the later iterations, the expensive steps can be avoided by making use of the tabulated values and

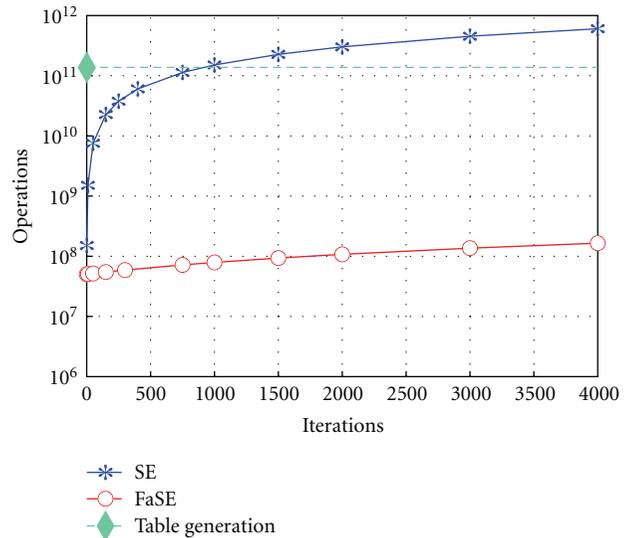


FIGURE 3: Operations per block for model generation by SE and FaSE and operations necessary for generating tabulated $C_{(k,l)}$ and D_k . For comparison, the operations for generating the tables are drawn over the complete iterations range although they have to be calculated only once. Spatial sizes $M = 64$ and $N = 64$ and dictionary size $|\mathcal{D}| = 4096$.

avoiding the update of the residual. For these evaluations, the calculation time for generating the tabulated values is not considered, as they only have to be computed once and can be stored. The very high computational cost of the weighted scalar products can also be recognized by regarding Figure 5 that shows the extrapolation time per block over iterations

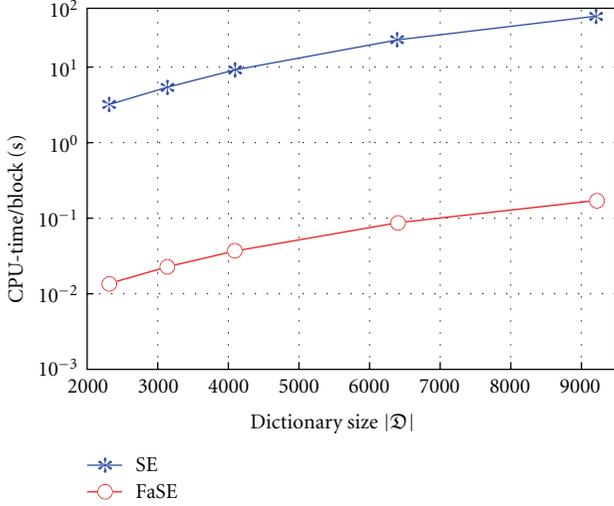


FIGURE 4: Processing time over dictionary size for 2D model generation with arbitrary real-valued basis functions and 250 iterations. The size of the extrapolation area is chosen so that $MN = |\mathcal{D}|$ holds for every data point.

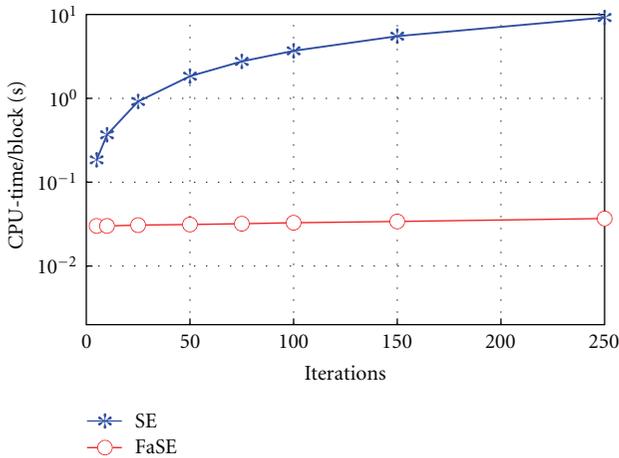


FIGURE 5: Processing time over iterations for 2D model generation with arbitrary real-valued basis functions of size 64×64 and dictionary size $|\mathcal{D}| = 4096$.

for an extrapolation scenario of size 64×64 samples. Taking into account the logarithmic axis, one can recognize that the processing time per block more or less linearly increases for SE, whereas for FaSE the processing time per block increases only very slowly. The results correspond well to the analytical complexity evaluation, and the speed gain of FaSE over SE is of the same magnitude as shown in the previous section. Apparently, Figure 3 cannot be directly translated into the processing time shown in Figure 5, since not all regarded operations consume the same processing time and since the analytical evaluation cannot account for optimizations introduced by the compiler.

Figure 6 shows the processing time for generating the tables for different dictionary sizes $|\mathcal{D}|$ and for different

TABLE 2: Average results for extrapolation of 126 blocks of size 16×16 samples in every image of the Kodak test image database.

Algorithm	PSNR	Processing time per block
TV [17]	22.40 dB	0.54 sec
SFG [18]	23.63 dB	15.29 sec
SDI [19]	21.60 dB	0.0003 sec
FaSE	23.82 dB	0.38 sec

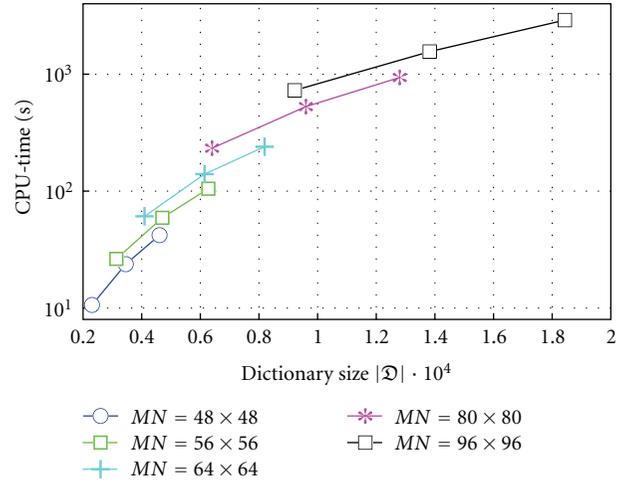


FIGURE 6: Processing time for generating tables over dictionary size $|\mathcal{D}|$ for extrapolation areas of different sizes MN .

sizes of extrapolation area \mathcal{L} . Comparing these results with the ones shown in Figure 5 one can recognize that for an extrapolation area of size 64×64 , a dictionary size of $|\mathcal{D}| = 4096$ and 250 iterations, the table generation only takes as long as SE would roughly need for extrapolating 6 blocks. This corresponds well to the theoretical results presented in the analytical evaluation. The discrepancy follows from the fact that different operations consume unequal amounts of processing time while in the analytical evaluation only the absolute number of operations has been counted.

Since the proposed novel spatial domain solution does not affect the model generation principle of SE, still a very high extrapolation quality can be achieved. Due to the acceleration of the algorithm, now very good extrapolation results can be achieved at a manageable complexity for arbitrary basis functions. To prove this, Table 2 shows the average extrapolation quality in terms of PSNR and the processing time for extrapolating 126 blocks of size 16×16 samples in every image from the Kodak image database. For comparison, the Total Variation Image Reconstruction (TV) algorithm from [17], the patch-based algorithm from [18] that uses Stochastic Factor Graphs (SFG), and the simple but very fast Spatial-Domain Interpolation (SDI) from [19] are regarded. The comparison has been carried out in MATLAB R2008b, and again only one core of the above-mentioned computer has been used. Apparently,

FaSE provides the highest extrapolation quality among the considered algorithms only with SFG coming close. But at the same time, it is the second fastest algorithm.

6. Modifications for Transform-Based Basis Function Sets

As aforementioned, for FaSE the weighted scalar products only have to be evaluated prior to the first iteration. In the case that the regarded basis function set contains a subset of basis functions that emanate from a discrete transform as for example, functions of the DCT or the DFT, the explicit evaluation of the weighted scalar products can be simplified by replacing the summation over the product between the weighted signal and the basis function by the corresponding transform coefficient of the weighted signal which can be achieved through a fast transform. To give an example, the idea that the basis function set contains some basis functions which emanate from the DFT will be extended. In this case, a basis function is defined by

$$\varphi_k[m, n] = e^{j(2\pi/M)\mu_k m} e^{j(2\pi/N)\eta_k n} \quad (21)$$

with vertical frequency μ_k and horizontal frequency η_k . Then, the summation from (12) can be expressed by the DFT

$$\sum_{(m,n) \in \mathcal{L}} s[m, n] \varphi_k^*[m, n] w[m, n] = \text{DFT} \{s[m, n] w[m, n]\} |_{\mu_k, \eta_k} \quad (22)$$

at frequency μ_k, η_k . Thus, the weighted scalar products for many basis functions can be efficiently evaluated simultaneously by making use of fast transforms like the Fast Fourier Transform [20] or, respectively, a fast transform that is appropriate to the regarded basis functions. It has to be noted that the utilization of fast transforms is only reasonable if a large number of transform domain coefficients has to be calculated at the same time. The fast transforms only speed up the parallel calculation of many coefficients. The calculation of just a single coefficient would take as long as the explicit evaluation of the weighted scalar product. The above-described property could also be used for speeding up the table generation in (16). Regarding again the example of a subset of DFT basis functions, the product between a basis function and a conjugate complex second one is equal to a basis function where the horizontal and vertical frequency results from the difference of the original frequencies:

$$\begin{aligned} \varphi_k^*[m, n] \varphi_l[m, n] &= e^{-j(2\pi/M)\mu_k m} e^{-j(2\pi/N)\eta_k n} e^{j(2\pi/M)\mu_l m} e^{j(2\pi/N)\eta_l n} \\ &= e^{j(2\pi/M)(\mu_l - \mu_k)m} e^{j(2\pi/N)(\eta_l - \eta_k)n}. \end{aligned} \quad (23)$$

Hence, (16) can also be expressed by the corresponding coefficients from the DFT. For other transform-based basis function sets, similar properties exist.

In addition to the results for arbitrary basis functions shown in Section 5, the performance of FaSE and SE is compared to a transform domain algorithm. For this, FSE

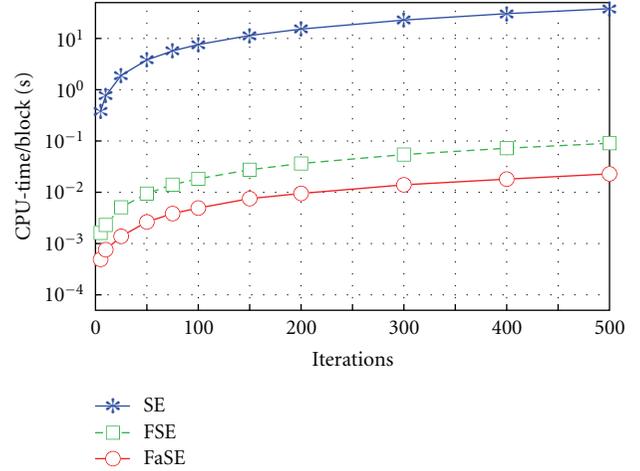


FIGURE 7: Processing time over iterations for 2D model generation with DFT basis functions, $|\mathcal{D}| = 4096$.

is regarded that utilizes Fourier functions for extrapolation. Here the circumstance has to be considered, that, as described in [8], FSE does not generate a complex-valued model. FSE selects in every iteration step one basis function and its corresponding conjugate complex one, in such a way that the model is always real-valued. Hence, in most cases two basis functions are selected in an iteration, with the exception of the real-valued constant basis function and the function with the highest possible alternation. Thus, the number of iterations has to be doubled for SE and FaSE for a fair comparison as they select only one basis function per iteration. Figure 7 shows the processing time per block for the different approaches with $|\mathcal{D}| = 4096$ Fourier basis functions of size 64×64 . For these simulations, the initial scalar products for FaSE are expressed by the transform coefficients according to (22). Although FaSE needs twice the number of iterations as FSE for generating the model, it is still significantly faster than FSE and furthermore several magnitudes faster than the original spatial domain SE.

Taking all the results from the two previous sections into account, the following recommendations can be given. In the case that the Selective Extrapolation is carried out with Fourier basis functions or other basis function sets that are based on a discrete transform, one can decide either to use a transform domain algorithm or the novel FaSE. If the same extrapolation scenario always is considered, the tables only have to be calculated once and the time gain of FaSE prevails, otherwise the transform domain algorithm is the better choice as no calculation of the tables is necessary. If the extrapolation process is carried out with basis functions for which no transform domain implementation is possible, FaSE should be preferred over the original SE. FaSE is able to efficiently trade computational complexity versus memory consumption as the expensive operations only have to be carried out once. Thus, the actual iterations for generating the model become very simple and very fast.

7. Conclusion

Within the scope of this contribution, we presented Fast Selective Extrapolation for image and video signal extrapolation. For this, Selective Extrapolation, a powerful signal extrapolation algorithm has been reviewed and its most complex parts have been identified. The novel algorithm behaves mathematically identical to the original algorithm but is able to outspeed the original algorithm by several decades by effectively trading memory consumption versus processing time. Furthermore, the novel algorithm is able to outperform existent fast transform domain extrapolation algorithms which are even limited to certain basis function sets. With that, it opens the door for further research on carrying out the extrapolation with different basis function sets. Up to now, the extrapolation only has been computationally manageable for special basis function sets that are based on discrete transforms. But by using Fast Selective Extrapolation, the extrapolation can be carried out for arbitrary basis functions which may even be only numerically defined. This ability allows for further research on extrapolation with signal adapted basis functions, obtained through the Karhunen-Loève Transform [21, 22], which has not been computationally feasible up to now.

Although the algorithm has been introduced only for two-dimensional data sets, it can be extended straightforwardly to three dimensions by making use of the ideas from [23] and four dimensions by using [24]. There, a three-dimensional or, respectively, a four-dimensional model is generated in the same way as described above for two dimensions.

References

- [1] T. Stockhammer and M. M. Hannuksela, "H.264/AVC video for wireless transmission," *IEEE Wireless Communications*, vol. 12, no. 4, pp. 6–13, 2005.
- [2] I. Richardson, *H.264 and MPEG-4 Video Compression*, John Wiley & Sons, West Sussex, UK, 2003.
- [3] B. A. Olshausen and D. J. Field, "Sparse coding with an overcomplete basis set: a strategy employed by V1?" *Vision Research*, vol. 37, no. 23, pp. 3311–3325, 1997.
- [4] E. Candès, N. Braun, and M. Wakin, "Sparse signal and image recovery from compressive samples," in *Proceedings of the 4th IEEE International Symposium on Biomedical Imaging (ISBI '07)*, pp. 976–979, April 2007.
- [5] J. A. Tropp, "Greed is good: algorithmic results for sparse approximation," *IEEE Transactions on Information Theory*, vol. 50, no. 10, pp. 2231–2242, 2004.
- [6] V. N. Temlyakov, "Weak greedy algorithms," *Advances in Computational Mathematics*, vol. 12, no. 2-3, pp. 213–227, 2000.
- [7] S. G. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Transactions on Signal Processing*, vol. 41, no. 12, pp. 3397–3415, 1993.
- [8] A. Kaup, K. Meisinger, and T. Aach, "Frequency selective signal extrapolation with applications to error concealment in image communication," *International Journal of Electronics and Communications*, vol. 59, no. 3, pp. 147–156, 2005.
- [9] J. L. Herraiz, S. España, E. Vicente et al., "Frequency selective signal extrapolation for compensation of missing sata in sinograms," in *Proceedings of the IEEE Nuclear Science Symposium Conference Record (NSS/MIC '08)*, pp. 4299–4302, October 2008.
- [10] M. Friebe and A. Kaup, "Fading techniques for error concealment in block-based video decoding systems," *IEEE Transactions on Broadcasting*, vol. 53, no. 1, pp. 286–295, 2007.
- [11] J. Cooley, P. Lewis, and P. Welch, "The finite Fourier transform," *IEEE Transactions on Audio and Electroacoustics*, vol. 17, no. 2, pp. 77–85, 1969.
- [12] N. Ahmed, T. Natarajan, and K. R. Rao, "Discrete cosine transform," *IEEE Transactions on Computers*, vol. 23, pp. 90–93, 1974.
- [13] J. L. Walsh, "A closed set of orthogonal functions," *American Journal of Mathematics*, vol. 55, pp. 5–24, 1923.
- [14] K. Meisinger and A. Kaup, "Minimizing a weighted error criterion for spatial error concealment of missing image data," in *Proceedings of the International Conference on Image Processing (ICIP '04)*, vol. 5, pp. 813–816, 2004.
- [15] J. Seiler, K. Meisinger, and A. Kaup, "Orthogonality deficiency compensation for improved frequency selective image extrapolation," in *Proceedings of the Picture Coding Symposium (PCS '07)*, Lisboa, Portugal, 2007.
- [16] J. Seiler and A. Kaup, "Fast orthogonality deficiency compensation for improved frequency selective image extrapolation," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP '08)*, pp. 781–784, 2008.
- [17] J. Dahl, P. C. Hansen, S. H. Jensen, and T. L. Jensen, "Algorithms and software for total variation image reconstruction via first-order methods," *Numerical Algorithms*, vol. 53, no. 1, pp. 67–92, 2010.
- [18] X. Li, "Variational bayesian image processing on stochastic factor graphs," in *Proceedings of the IEEE International Conference on Image Processing (ICIP '08)*, pp. 1748–1751, October 2008.
- [19] Z. Alkachouh and M. G. Bellanger, "Fast DCT-based spatial domain interpolation of blocks in images," *IEEE Transactions on Image Processing*, vol. 9, no. 4, pp. 729–732, 2000.
- [20] J. Cooley and J. Tukey, "An algorithm for the machine calculation of complex Fourier series," *Mathematics of Computation*, vol. 19, no. 90, pp. 297–301, 1965.
- [21] K. Karhunen, "Über lineare Methoden in der Wahrscheinlichkeitsrechnung," *Annales Academiae Scientiarum Fennicae*, vol. 37, pp. 3–79, 1947.
- [22] M. Loève, "Aléatoires de second order," in *Processus Stochastiques et Moevement Brownien*, P. Lévy, Ed., Hermann, Paris, France, 1948.
- [23] K. Meisinger and A. Kaup, "Spatiotemporal selective extrapolation for 3-D signals and its applications in video communications," *IEEE Transactions on Image Processing*, vol. 16, no. 9, pp. 2348–2360, 2007.
- [24] U. Fecker, J. Seiler, and A. Kaup, "4-D frequency selective extrapolation for error concealment in multi-view video," in *Proceedings of the IEEE 10th Workshop on Multimedia Signal Processing (MMSP '08)*, pp. 267–272, 2008.