

Algorithms for Piecewise Constant Signal Approximations

Leif Bergerhoff

Mathematical Image Analysis Group
Saarland University
66041 Saarbrücken, Germany
bergerhoff@mia.uni-saarland.de

Joachim Weickert

Mathematical Image Analysis Group
Saarland University
66041 Saarbrücken, Germany
weickert@mia.uni-saarland.de

Yehuda Dar

Computer Science Department
Technion – Israel Institute of Technology
Haifa 3200003, Israel
ydar@cs.technion.ac.il

Abstract—We consider the problem of finding optimal piecewise constant approximations of one-dimensional signals. These approximations should consist of a specified number of segments (samples) and minimise the mean squared error to the original signal. We formalise this goal as a discrete nonconvex optimisation problem, for which we study two algorithms. First we reformulate a recent adaptive sampling method by Dar and Bruckstein in a compact and transparent way. This allows us to analyse its limitations when it comes to violations of its three key assumptions: signal smoothness, local linearity, and error balancing. As a remedy, we propose a direct optimisation approach which does not rely on any of these assumptions and employs a particle swarm optimisation algorithm. Our experiments show that for nonsmooth signals or low sample numbers, the direct optimisation approach offers substantial qualitative advantages over the Dar–Bruckstein method. As a more general contribution, we disprove the optimality of the principle of error balancing for optimising data in the ℓ^2 norm.

Index Terms—Adaptive Signal Processing, Nonuniform Sampling, Nonconvex Optimisation, Particle Swarm Optimisation, Segmentation

I. INTRODUCTION

One of the essential concepts in signal processing is the sampling and reconstruction of continuous signals. The classical sampling theory (see e.g. [1] for a review) considers uniform sampling and provides conditions under which the sampling allows lossless signal reconstruction. These concepts have been extended to nonuniform sampling where one can adapt the sampling rate to the local signal bandwidth [2]–[5].

If one aims at lossy signal representations to achieve higher compression rates, other options can be preferable. Recently Dar and Bruckstein [6] have introduced a simple and efficient adaptive sampling strategy for approximating 1-D signals by piecewise constant functions. It involves three assumptions: smoothness, local linearity, and error balancing. In practice, however, signals can be nonsmooth, they can violate local linearity, and the optimality of error balancing is unclear. Thus, finding the optimal approach for the general case remains an open problem.

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 741215). It has been initiated at a joint visit of Alfred Bruckstein and Joachim Weickert to the Isaac Newton Institute, which was supported by EPSRC Grant Number EP/K032208/1 and by a Rothschild Distinguished Visiting Fellowship.

A. Our Contributions

To address this problem, we first revisit the Dar–Bruckstein model by deriving it in a simpler, alternative way which is inspired by the work of Belhachmi et al. [7]. Its analysis enables us to quantify the effects of violating local linearity and to disprove the optimality of error balancing. As a remedy, we propose an energy minimisation model that does not rely on smoothness, local linearity, or error balancing. It favours globally optimal piecewise constant signal approximations which minimise the mean squared error (MSE). This requires to solve a nonconvex optimisation problem, for which a particle swarm optimisation (PSO) algorithm [8] performs well. Experiments show that the quality of our novel approach can exceed the one of the Dar–Bruckstein method.

B. Structure of the Paper

In Section II we introduce the underlying approximation problem, reformulate and analyse the Dar–Bruckstein model, and we propose our novel method along with the PSO algorithm. Section III presents experimental comparisons of both approaches with a smooth, a nonsmooth, and a noisy signal. We conclude with a summary and an outlook in Section IV.

II. MODELLING

In this section we formalise our approximation problem and discuss two different solution strategies: 1. The Dar–Bruckstein approach for which we give a new and compact derivation and an analysis of its limitations. 2. Our novel direct optimisation method which aims at overcoming these limitations by renouncing assumptions of the Dar–Bruckstein model that may be violated or cause suboptimal solutions.

A. Problem Statement

Consider a signal domain $[a, b] \subset \mathbb{R}$ and some integrable 1-D signal $f : [a, b] \rightarrow \mathbb{R}$. We want to approximate f by a piecewise constant signal $u : [a, b] \rightarrow \mathbb{R}$ that has N segments (samples) and minimises the MSE w.r.t. f . This requires to find $N-1$ segment boundaries x_1, x_2, \dots, x_{N-1} with

$$a < x_1 < x_2 < \dots < x_{N-1} < b. \quad (1)$$

It is convenient to introduce also $x_0 := a$ and $x_N := b$. Then we approximate f by a piecewise constant signal of type

$$u(x) := \begin{cases} u_i, & \text{for } x \in [x_i, x_{i+1}) \text{ and } i < N-1, \\ u_{N-1}, & \text{for } x \in [x_{N-1}, x_N]. \end{cases} \quad (2)$$

Since we aim at an ℓ^2 -optimal approximation, each constant u_i is given by the mean value of f in $[x_i, x_{i+1}]$:

$$u_i := \frac{1}{x_{i+1} - x_i} \int_{x_i}^{x_{i+1}} f(y) dy. \quad (3)$$

We observe that u is completely determined by f and the segment boundary vector $\mathbf{x} := (x_1, x_2, \dots, x_{N-1})^\top$. Thus, our problem comes down to minimising the discrete energy

$$E(\mathbf{x}) = \frac{1}{b-a} \sum_{i=0}^{N-1} \int_{x_i}^{x_{i+1}} (f(y) - u_i)^2 dy. \quad (4)$$

Depending on the application, this may be interpreted as function approximation, adaptive sampling, segmentation, or lossy signal compression.

Although the energy (4) does not look very complicated, in general it is nonsmooth, nonconvex, and may have many local minima. This is illustrated in Fig. 1 for a piecewise constant signal $f(x)$, for which we want to find an approximation $u(x)$ with a single jump position x_1 . This corresponds to the simplest nontrivial scenario $N = 2$.

B. Compact Reformulation of the Dar–Bruckstein Method

Recently Dar and Bruckstein [6] have proposed an approach to solve the problem in Section II-A very efficiently. To explain the underlying ideas and assumptions in a simple and transparent way, we reformulate its derivation. This reformulation is inspired by work of Belhachmi et al. [7, Section 6].

Denoting the squared error in the interval $[x_i, x_{i+1}]$ by

$$e_i := \int_{x_i}^{x_{i+1}} (f(y) - u_i)^2 dy, \quad (5)$$

we can write the energy function (4) as

$$E(\mathbf{x}) = \frac{1}{b-a} \sum_{i=0}^{N-1} e_i. \quad (6)$$

Dar and Bruckstein assume that the input signal f is a continuously differentiable (C^1) function and that N is large enough such that f can be approximated well by a linear function within each interval $[x_i, x_{i+1}]$ for $i = 0, 1, \dots, N-1$. Thus, in (x_i, x_{i+1}) we have

$$f'(x) = \frac{f(x_{i+1}) - f(x_i)}{x_{i+1} - x_i} =: f'_i, \quad (7)$$

$$f(x) = f(x_i) + (x - x_i) f'_i, \quad (8)$$

$$u_i = \frac{1}{2} (f(x_i) + f(x_{i+1})). \quad (9)$$

Using this in (5) and applying some simple calculations yields

$$e_i = \frac{1}{12} h_i^3 f_i'^2 \quad (10)$$

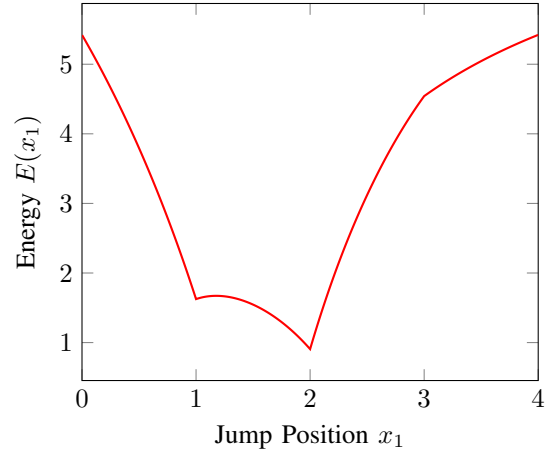
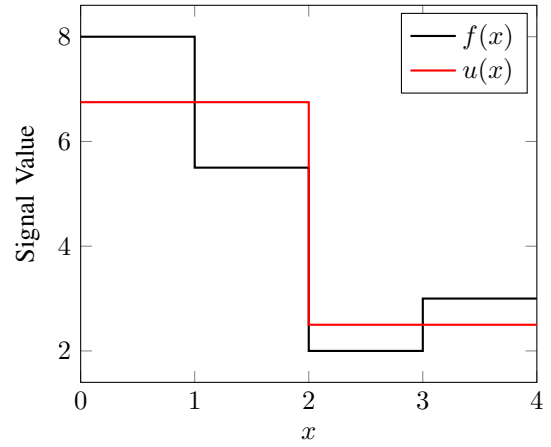


Fig. 1. (a) **Top:** A piecewise constant signal $f(x)$. Its optimal piecewise constant approximation $u(x)$ with $N = 2$ segments is obtained for $x_1 := 2$. (b) **Bottom:** The corresponding energy / MSE curve as a function of x_1 is nonsmooth and nonconvex, and it has two local minima.

where $h_i := x_{i+1} - x_i$ denotes the interval width.

As a heuristic for minimising the global energy (6), one may assume that \mathbf{x} is optimal if all local errors e_i are balanced. Using $e_0 = e_1 = \dots = e_{N-1} = \text{const.}$ with (10) gives the following proportionalities:

$$f_i'^2 \sim \frac{1}{h_i^3} \implies \frac{1}{h_i} \sim \sqrt[3]{f_i'^2}. \quad (11)$$

Since $1/h_i$ can be seen as a measure for the local density of the sampling points, one should choose the interval boundaries for optimal sampling proportional to $\sqrt[3]{f_i'^2}$. Consequently, Dar and Bruckstein select \mathbf{x} such that every segment $[x_i, x_{i+1}]$ contains the same amount of the cube root of the squared signal derivative. More precisely:

$$\int_{x_i}^{x_{i+1}} \sqrt[3]{(f'(y))^2} dy = \frac{1}{N} \int_a^b \sqrt[3]{(f'(y))^2} dy =: T_{opt} \quad (12)$$

for $i = 0, 1, \dots, N-1$. The threshold T_{opt} is computed a priori. Thus, the analytical formula (12) allows to estimate the interval boundaries \mathbf{x} in a simple and efficient way.

C. Limitations of the Dar–Bruckstein Method

We have seen that the Dar–Bruckstein approach relies on three assumptions: C^1 -smoothness, local linearity, and error balancing. Let us now analyse the impact of these assumptions on the optimality of the method in detail.

- Obviously the smoothness assumption on f is violated if the signal is nondifferentiable or noisy.
- To quantify inaccuracies caused by violations of the local linearity assumption, we derive a formula for e_i that does not use this assumption. We can rewrite (5) as

$$e_i = \int_{x_i}^{x_{i+1}} (f(y) - f(\xi_i))^2 dy, \quad (13)$$

where we have used the continuity of f , which guarantees that there exists a $\xi_i \in [x_i, x_{i+1}]$ with $f(\xi_i) = u_i$. With the mean value theorem, Equation (13) becomes

$$e_i = \int_{x_i}^{x_{i+1}} (\xi_i - y)^2 (f'(\theta_i))^2 dy \quad (14)$$

$$= \frac{1}{3} ((x_{i+1} - \xi_i)^3 + (\xi_i - x_i)^3) (f'(\theta_i))^2 \quad (15)$$

for a suitable $\theta_i \in [x_i, x_{i+1}]$. Using $h_i = x_{i+1} - x_i$ and defining $\eta_i := (x_{i+1} - \xi_i)/h_i$ allows to rewrite (15) as

$$e_i = \frac{1}{3} (1 - 3\eta_i + 3\eta_i^2) h_i^3 (f'(\theta_i))^2. \quad (16)$$

Comparing the exact error (16) with the error (10) that exploits local linearity shows the following: Since $\xi_i \in [x_i, x_{i+1}]$, we know that $\eta_i \in [0, 1]$. However, only for $\eta_i = \frac{1}{2}$, we obtain $\frac{1}{3} (1 - 3\eta_i + 3\eta_i^2) = \frac{1}{12}$. In the worst case with $\eta_i = 0$ or 1, this factor becomes $\frac{1}{3}$. Moreover, since $f \in C^1[a, b]$, there exist constants $m := \min_{[a,b]} f'$ and $M := \max_{[a,b]} f'$. Thus, $(f'(\theta_i))^2$ can attain any value between m^2 and M^2 , which can differ substantially from $f_i'^2$. This shows that without local linearity, (10) can be violated severely. Moreover, (12) does no longer balance the errors then.

- While the principle of error balancing sounds plausible, one cannot prove that it is fulfilled for the globally optimal u which minimises the MSE. Actually, already Fig. 1(a) serves as counterexample: The error in the left segment of u is clearly larger than in the right segment.

D. A Novel Direct Optimisation Approach

The preceding discussion shows that it can be desirable to renounce all three assumptions of the Dar–Bruckstein model. Interestingly, there is a surprisingly simple solution: We can rely directly on the discrete model (4), which is perfect from a modelling viewpoint. However, we have to deal with a challenging nonsmooth and nonconvex optimisation problem with numerous local minima. Since we cannot expect to find an efficient algorithm with formal convergence guarantees to a global minimum, we use a nature-inspired metaheuristic that ends up in a good local minimum. Based on our tests, we recommend to minimise (4) by a Particle Swarm Optimisation

(PSO) approach. PSO is an iterative global optimisation technique for nonlinear functions. It emerged from a simulation of social behaviour [8]. The term “swarm” refers to a specified number of n virtual particles which explore the solution space while interacting among each other. The i -th particle represents one solution \mathbf{x}_i with corresponding energy $E(\mathbf{x}_i)$.

In our work we apply the Standard Particle Swarm Optimisation 2011 algorithm (SPSO-2011), which uses an adaptive random particle neighbourhood topology and features rotation invariance [9]. The initial particle positions \mathbf{x}_i^0 with $i = 1, 2, \dots, n$ are uniformly distributed in the solution space. In iteration step k , the particle position \mathbf{x}_i^k is updated as follows:

$$\mathbf{p}_i^k = \mathbf{x}_i^k + c_1 \mathbf{u}_1^k \odot (\mathbf{P}_i^k - \mathbf{x}_i^k), \quad (17)$$

$$\mathbf{l}_i^k = \mathbf{x}_i^k + c_2 \mathbf{u}_2^k \odot (\mathbf{L}_i^k - \mathbf{x}_i^k), \quad (18)$$

$$\mathbf{g}_i^k = \frac{1}{3} (\mathbf{x}_i^k + \mathbf{p}_i^k + \mathbf{l}_i^k), \quad (19)$$

$$\mathbf{v}_i^{k+1} = \omega \mathbf{v}_i^k + \mathcal{H}_i(\mathbf{g}_i^k, |\mathbf{g}_i^k - \mathbf{x}_i^k|) - \mathbf{x}_i^k, \quad (20)$$

$$\mathbf{x}_i^{k+1} = \mathbf{x}_i^k + \mathbf{v}_i^{k+1}, \quad (21)$$

where \mathbf{P}_i^k denotes the previously best position of particle i w.r.t. the energy E , and \mathbf{L}_i^k is the best previous position within its neighbourhood of size K . The symbol \odot denotes element-wise vector multiplication, and \mathbf{u}_1^k and \mathbf{u}_2^k are independent and uniformly distributed random vectors with components in $[0, 1]$. The scalars c_1 , c_2 , and ω are nonnegative weights. $\mathcal{H}_i(\mathbf{g}_i^k, |\mathbf{g}_i^k - \mathbf{x}_i^k|)$ selects a random point from the hypersphere around \mathbf{g}_i^k with radius $|\mathbf{g}_i^k - \mathbf{x}_i^k|$ in the Euclidean norm $|\cdot|$.

This algorithm can be understood as follows. Equation (17) incorporates the successful history of particle i by defining a point \mathbf{p}_i^k near \mathbf{P}_i^k . Equation (18) expresses the knowledge of its neighbours by specifying a point \mathbf{l}_i^k near \mathbf{L}_i^k . In (19) we compute the centre of gravity \mathbf{g}_i^k of \mathbf{x}_i^k , \mathbf{p}_i^k , and \mathbf{l}_i^k . Equation (20) involves \mathbf{g}_i^k to update the velocity \mathbf{v}_i^{k+1} of particle i . This velocity is used in (21) to move \mathbf{x}_i^k to its new position \mathbf{x}_i^{k+1} . If the global optimum shows no improvement then each particle randomly selects K new neighbours. The algorithm stops if a maximum number of iterations or a tolerable energy threshold is reached. For more details, we refer to [9].

III. EXPERIMENTS

Let us now evaluate the approximation quality of the Dar–Bruckstein method and our direct optimisation approach.

We have implemented the Dar–Bruckstein model as is proposed in [6, Subsection II. A.] using (12). For the PSO algorithm for optimising our direct model, we adhere to [9]. As PSO parameters we use a maximum number of 10000 iterations, a swarm size of $n = 1000$, and a neighbourhood size of $K = 20$. We reset the neighbourhood structures after 15 iterations with no change in the global minimum. Following [9], we choose $c_1 = c_2 = 0.5 + \ln(2)$ and $\omega = 1/(2 \ln(2))$. Since the PSO algorithm involves randomisation, the quality of multiple program runs with identical parameters may differ somewhat. Thus, for every sample number N , we run the PSO algorithm 50 times and report the mean μ_{MSE} , the standard deviation σ_{MSE} , the minimum \min_{MSE} , and the maximum \max_{MSE} of our MSE computations. These MSE results are

listed in Tab. I, and Fig. 2 displays the signals, error graphs, and approximations.

In our first experiment we consider the chirp signal $f(x) = 255 \cos(2\pi x(1 + 5x))$ within the interval $[0, 1]$; see Fig. 2, top left. It was also studied in [6]. It constitutes a prototype for a smooth signal, which is nevertheless challenging in its high frequent right part. We observe that for $N \leq 50$, the Dar–Bruckstein approach performs consistently worse than the direct optimisation method. It appears to suffer from violations of the local linearity assumption. For larger values of N , this effect vanishes more and more, and the Dar–Bruckstein algorithm reaches a comparable quality as the direct approach. It is even slightly ahead for very large N . This is caused by the nonoptimality of the PSO algorithm: Its performance deteriorates somewhat for the more complex optimisation problems that arise for large N . We expect that this can be solved with more advanced optimisation techniques that we will consider in our future research.

In our second experiment, we use the nonsmooth real-world signal shown in Fig. 2, top centre. It represents line 51 of the 8-bit test image *trui*, which depicts a lady with a scarf. Within each of its 256 pixels, the function values are regarded constant. Having many jump discontinuities, this signal violates both the smoothness and the local linearity assumption. Thus, it is not surprising that it poses substantial challenges for the Dar–Bruckstein model: We observe that the Dar–Bruckstein MSE is about twice as large as the MSE of the direct optimisation method, and that this factor remains also for large values of N . The direct optimisation approach clearly benefits from its absence of any smoothness or local linearity requirements.

For our third experiment, we degrade the previous image signal by additive Gaussian noise with zero mean and standard deviation 20; see Fig. 2, top right. This takes the violation of the smoothness and local linearity assumption to the extremes. While the approximation quality of both approaches deteriorates in comparison to the noise-free scenario, the MSE of the Dar–Bruckstein method remains almost twice as large as the one of our direct optimisation approach.

Without going into details, we remark that in all experiments and for both methods, we found substantial deviations from perfect error balancing. For the Dar–Bruckstein method this shows that the local linearity assumption is still not fully met. For our direct optimisation approach it indicates that the optimal solution does not satisfy ideal error balancing. This is in agreement with the counterexample in Fig. 1(a).

IV. CONCLUSIONS AND OUTLOOK

We have studied piecewise constant signal approximations that possess a specified number of samples and aim at minimising the MSE to the original signal.

As a first contribution, we have provided a simple alternative derivation of the recent Dar–Bruckstein approach. It enabled us to analyse the limitations of the method in detail. We have shown that the quality of the Dar–Bruckstein approach

TABLE I
APPROXIMATION QUALITY OF THE DAR–BRUCKSTEIN (DB) APPROACH
AND OUR DIRECT OPTIMISATION METHOD

	N	DB	Our Method			
		MSE	μ_{MSE}	σ_{MSE}	minMSE	maxMSE
Chirp Signal	5	32166.88	19055.49	0.00	19055.49	19055.49
	10	23655.94	12014.77	459.47	11131.75	12431.60
	20	5661.01	4556.37	632.89	3403.23	5998.27
	30	2590.87	2037.90	67.90	1906.95	2285.76
	40	1477.48	1218.49	37.17	1177.55	1357.26
	50	975.93	854.12	20.76	823.13	907.48
	60	686.30	624.44	13.37	601.65	665.50
	70	510.96	479.90	9.26	460.10	513.00
	80	377.20	383.08	6.95	367.62	400.07
	90	307.64	311.04	5.11	301.63	328.83
100	247.84	258.05	4.25	249.73	269.60	
Image Signal	5	674.31	347.36	0.00	347.36	347.36
	10	436.22	126.37	3.51	123.47	130.96
	20	115.28	47.83	3.75	43.05	60.15
	30	54.46	26.28	1.28	24.03	29.48
	40	36.58	17.39	0.79	15.15	18.83
	50	26.34	12.63	0.60	11.41	14.28
	60	20.90	10.23	0.42	8.88	11.20
	70	16.25	8.34	0.40	7.54	9.28
	80	15.47	7.01	0.43	6.09	8.29
	90	11.13	5.87	0.31	5.21	6.67
100	10.20	5.09	0.28	4.33	5.64	
Noisy Image Signal	5	1532.52	709.29	0.00	709.29	709.29
	10	1130.94	454.99	6.54	451.57	474.21
	20	588.69	332.17	8.15	312.44	350.30
	30	435.49	286.25	4.90	277.01	300.15
	40	391.56	262.20	5.33	244.90	273.72
	50	391.27	242.26	8.45	219.93	262.93
	60	330.87	212.69	7.68	196.22	229.98
	70	301.52	193.08	8.39	174.53	212.21
	80	288.86	170.87	6.63	157.66	189.61
	90	258.99	156.27	7.25	139.10	172.14
100	250.82	142.21	6.97	128.11	164.89	

is compromised, if smoothness or local linearity are violated, or if error balancing leads to suboptimal solutions.

In a second step, these insights have triggered us to consider a direct optimisation model that renounces all three assumptions. It gives a globally optimal solution if one can solve the corresponding nonconvex optimisation problem exactly. Our experiments show that already a suboptimal particle swarm optimisation algorithm yields good local minima in practice.

Evaluating the quality of both approaches for a smooth, a nonsmooth, and a noisy signal has demonstrated that the direct method offers better quality if the signal lacks smoothness or the number of samples is low. For smooth signals with many samples, the Dar–Bruckstein algorithm remains a simple and efficient alternative with comparable quality.

In our ongoing work, we are investigating alternative optimisation approaches for the nonconvex model. Moreover, we will extend our research on piecewise constant approximations to the multidimensional case, and we are going to study suitable applications, e.g. in data compression and quantisation.

Apart from its specific contributions to signal approximation, our paper may also be of more general relevance by disproving the optimality of the principle of error balancing for optimisation problems that involve the ℓ^2 -norm. This principle is

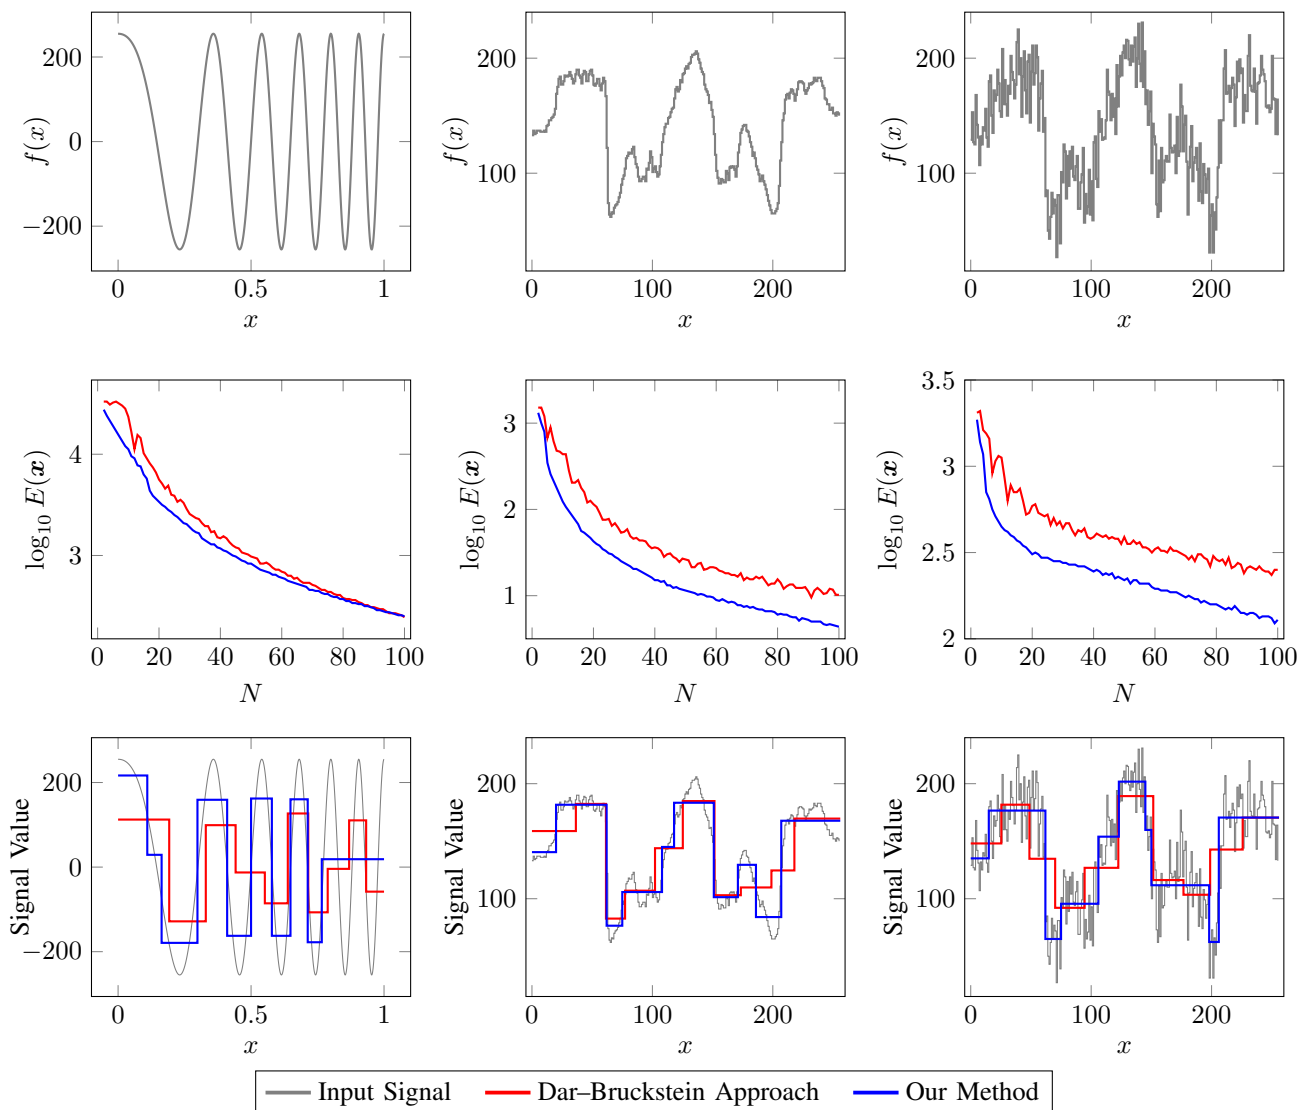


Fig. 2. **Left column:** Results for the chirp signal. **Central column:** Results for the image signal. **Right column:** Results for the image signal degraded by additive Gaussian noise. **Top row:** Original signals. **Middle row:** MSE for the Dar-Bruckstein approach and \min_{MSE} for our method as a function of the sample number N . **Bottom row:** Signal approximations for 10 samples. We present the results with lowest MSE from all 50 runs of the PSO algorithm.

omnipresent in numerous applications, ranging from scientific computing to computer graphics. Thus, understanding its sub-optimality in more detail may also pave the road to better algorithms beyond the field of signal processing.

V. ACKNOWLEDGEMENT

We thank Tobias Alt and Matthias Augustin for valuable comments and fruitful discussions.

REFERENCES

- [1] A. J. Jerri, "The Shannon sampling theorem – Its various extensions and applications: A tutorial review," *Proceedings of the IEEE*, vol. 65, no. 11, pp. 1565–1596, Nov. 1977.
- [2] K. Horiuchi, "Sampling principle for continuous signals with time-varying bands," *Information and Control*, vol. 13, no. 1, pp. 53–61, Jul. 1968.
- [3] J. J. Clark, M. R. Palmer, and P. D. Lawrence, "A transformation method for the reconstruction of functions from nonuniformly spaced samples," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 33, no. 5, pp. 1151–1165, Oct. 1985.
- [4] N. N. Brueller, N. Peterfreund, and M. Porat, "Non-stationary signals: optimal sampling and instantaneous bandwidth estimation," in *Proc. IEEE-SP International Symposium on Time-Frequency and Time-Scale Analysis*, Pittsburgh, PA, USA, Oct. 1998, pp. 113–115.
- [5] D. Wei and A. V. Oppenheim, "Sampling based on local bandwidth," in *2007 Conference Record of the Forty-First Asilomar Conference on Signals, Systems and Computers*, Nov. 2007, pp. 1103–1107.
- [6] Y. Dar and A. M. Bruckstein, "On high-resolution adaptive sampling of deterministic signals," arXiv:1611.01850 [cs.IT], Apr. 2018.
- [7] Z. Belhachmi, D. Bucur, B. Burgeth, and J. Weickert, "How to choose interpolation data in images," *SIAM Journal of Applied Mathematics*, vol. 70, no. 1, pp. 333–352, 2009.
- [8] J. Kennedy and R. Eberhart, "Particle swarm optimization," in *Proc. 1995 International Conference on Neural Networks*, vol. 4, Perth, WA, Australia, Nov. 1995, pp. 1942–1948.
- [9] M. Zambrano-Bigiarini, M. Clerc, and R. Rojas-Mujica, "Standard particle swarm optimisation 2011 at CEC-2013: A baseline for future PSO improvements," in *Proc. IEEE Congress on Evolutionary Computation*, Cancun, Mexico, Jun. 2013, pp. 2337–2344.