Machine Learning for Signal Reconstruction from Streaming Time-series Data

Emilio Ruiz-Moreno

Machine Learning for Signal Reconstruction from Streaming Time-series Data

Emilio Ruiz-Moreno

Machine Learning for Signal Reconstruction from Streaming Time-series Data

Doctoral Dissertation for the Degree *Philosophiae Doctor (Ph.D.)* at the Faculty of Engineering and Science, Specialisation in ICT

University of Agder
Faculty of Engineering and Science
2024

Preface

This Ph.D. thesis results from the research conducted at the the WISENET Center, Department of Information and Communication Technology, Universitetet i Agder (UiA), Grimstad, Norway, from August 2018 to September 2023. The corresponding Ph.D. degree has been supervised by Professor Baltasar Beferull-Lozano, and co-supervised by Dr. Luis Miguel López-Ramos, and Professor Jing Zhou.

During my Ph.D. program, I have combined research work at the SFI Offshore Mechatronics Centre for Research-based Innovation, from September 2018 to August 2022, under grant 237896/O30, with teaching duties at UiA from January 2020 to August 2022. Moreover, in addition to the Ph.D. coursework components, I attended the machine learning crash course (MLCC) at Università di Genova, Italy, in 2019; the DeepLearn 2021 Summer at Las Palmas de Gran Canaria, Spain, in 2021; and the finance track of the Oxford Machine Learning Summer School (OxML) at the University of Oxford, England, in 2022.

Acknowledgments

My doctoral journey has been nothing short of a thrilling quest, filled with new experiences and challenges that have enriched my understanding of the world and field of study. Along this path, I have been fortunate to encounter mentors, colleagues, and friends who have aided me in my academic endeavors and have played pivotal roles in shaping the essence of this thesis. So, let me take this opportunity to express my heartfelt gratitude to all those who have journeyed with me.

First and foremost, I would like to extend my gratitude to my supervisor Baltasar for his exceptional guidance and unwavering support. His mentorship has been a cornerstone of my academic growth, and I am truly fortunate to have had the privilege of working with him. I would also like to thank my co-supervisor Luismi, whose profound domain knowledge and trust in my academic skills have been instrumental in the research presented in this thesis. Beyond the academic realm, Baltasar and Luismi's kindness has even made this stiff scholarly pursuit enjoyable.

Thanks to everyone at the WISENET Center, old and new colleagues, for providing the warmest working environment I could have asked for. Especially to my core office mates Siavash, Ravi, and Juan Diego for all the shared moments. I will keep filling whiteboards with souvenir magnets for old times' sake.

Many thanks to all the friends I have made along the way. All the board games, nights at one of the three pubs in Grimstad, and moments lived together hold a special place in my memories and have made these past years unforgettable.

I want to thank my friend Alfredo; I am sure I could not have had a better housemate.

I want also to thank my family and friends back home. To my parents and sisters for their immeasurable love and support, to Cheto and Bambi for their dog affection, and to my lifelong friends for always being there for me despite time and distance.

Last but not least, to my partner Eleni. Thank you for your patience, continuous encouragement, and help throughout the last few years.

<div align="center">

Emilio Ruiz-Moreno
Oslo, Norway
January 24, 2024

</div>

Abstract

Nowadays, deploying cyber-physical networked systems generates tremendous streams of data, with data rates increasing as time goes by. This trend is especially noticeable in several fairly automated sectors, such as energy or telecommunications. Compared to the last decades, this represents not only an additional large volume of data to explore and the need for more efficient and scalable data analysis methods but also raises additional challenges in the design and analysis of real-time streaming data processing algorithms. In many applications of interest, it is required to process a sequence of samples from multiple, possibly correlated, data time series that are acquired at different sampling rates and which may be quantized in amplitude at different resolutions. A commonly sought goal is to obtain a low-error signal reconstruction that can be uniformly resampled with a temporal resolution as fine as desired, hence facilitating subsequent data analyses.

This Ph.D. thesis consists of a compendium of four papers that incrementally investigate the task of sequentially reconstructing a signal from a stream of multivariate time series of quantization intervals under several requirements encountered in practice and detailed next.

First, we investigate how to track signals from streams of quantization intervals while enforcing low model complexity in the function estimation. Specifically, we explore the use of reproducing kernel Hilbert space-based online regression techniques expressly tailored for such a task. More specifically, the core techniques we devise and employ are influenced by the abundant theoretical and practical benefits in the literature about proximal operators and multiple kernel approaches.

Second, we require the signal to be sequentially reconstructed, subject to smoothness constraints, and as soon as a data sample is available (zero-delay response). These well-motivated requirements appear in many practical problems, including online trajectory planning, real-time control systems, and high-speed digital-to-analog conversion. We address this challenge through a novel spline-based approach underpinned by a sequential decision-making framework and assisted with deep learning techniques. Specifically, we use recurrent neural networks to capture the temporal dependencies among data, helping to reduce the roughness of the reconstruction on average.

Finally, we analyze the requirement of consistency, which amounts to exploiting all available information about the signal source and acquisition system to optimize some figure of reconstruction merit. In our context, consistency means guaranteeing that the reconstruction lies within the acquired quantization intervals. Consistency has been proven to entail a profitable-in-practice asymptotic error-rate decay as the sampling rate increases. Particularly, we investigate the impact of consistency on zero-delay reconstruction and also incorporate the idea of exploiting the spatiotemporal dependencies among multivariate signals.

Sammendrag

I dag genererer utplassering av cyberfysiske nettverkssystemer enorme datastrømmer, med datahastigheter som øker etter hvert som tiden går. Denne trenden er spesielt merkbar i flere ganske automatiserte sektorer, som energi eller telekommunikasjon. Sammenlignet med de siste tiårene representerer dette ikke bare et ekstra stort volum av data å utforske og et behov for mer effektive og skalerbare dataanalysemetoder, men reiser også ytterligere utfordringer i utformingen og analysen av sanntids strømnings databehandling algoritmer. I mange applikasjoner av interesse er det nødvendig å behandle en sekvens av prøver fra flere, muligens korrelerte, datatidsserier som er innhentet ved forskjellige samplingshastigheter og som kan kvantiseres i amplitude ved forskjellige oppløsninger. Et vanlig mål er å oppnå en rekonstruksjon med lav feilsignal som kan uniformt gjensamples med en så fin temporal oppløsning som ønsket, og dermed fasilitere etterfølgende dataanalyser.

Denne Ph.D. avhandlingen består av et kompendium av fire artikler som inkrementelt undersøker oppgaven med å sekvensielt rekonstruere et signal fra en strøm av multivariate tidsserier av kvantiseringsintervaller under flere krav som møtes i praksis og detaljert deretter.

Først undersøker vi hvordan vi kan spore signaler fra strømmer av kvantiseringsintervaller mens vi håndhever lav modellkompleksitet i funksjonsestimeringen. Spesielt utforsker vi bruken av reproducing kernel Hilbert-rombaserte online regresjonsteknikker som er spesielt skreddersydd for en slik oppgave. Mer spesifikt er kjerneteknikkene vi utarbeider og bruker påvirket av de mange teoretiske og praktiske fordelene i litteraturen om proksimale operatører og flere kjernetilnærminger.

For det andre krever vi at signalet rekonstrueres sekvensielt, underlagt jevnhetsbegrensninger, og så snart en dataprøve er tilgjengelig (null-forsinkelsesrespons). Disse godt motiverte kravene dukker opp i mange praktiske problemer, inkludert online baneplanlegging, sanntidskontrollsystemer og høyhastighets digital-til-analog konvertering. Vi takler denne utfordringen gjennom en ny spline-basert tilnærming underbygget av et sekvensielt beslutningsrammeverk og assistert med dyplæringsteknikker. Spesifikt bruker vi recurrent nevrale nettverk for å fange opp de tidsmessige avhengighetene mellom data, og bidrar til å redusere grovheten til rekonstruksjonen i gjennomsnitt.

Til slutt analyserer vi kravet om konsistens, som utgjør å utnytte all tilgjengelig informasjon om signalkilden og innsamlingssystemet for å optimalisere noen form for rekonstruksjonsverdi. I vår sammenheng betyr konsistens å garantere at rekonstruksjonen ligger innenfor de innhentede kvantiseringsintervallene. Konsistens har vist lønnsom praktisk asymptotisk fofall av feilraten når samplingsfrekvensen øker. Spesielt undersøker vi virkningen av konsistens på null-forsinkelsesrekonstruksjon og inkluderer også ideen om å utnytte de spatiotemporale avhengighetene blant multivariate signaler.

Publications

Paper A **E. Ruiz-Moreno** and B. Beferull-Lozano, "Tracking of quantized signals based on online kernel regression," in *2021 IEEE 31st International Workshop on Machine Learning for Signal Processing (MLSP), pp. 1-6, IEEE, 2021.*

Paper B **E. Ruiz-Moreno** and B. Beferull-Lozano, "An online multiple kernel parallelizable learning scheme," *IEEE Signal Processing Letters*, 2023.

Paper C **E. Ruiz-Moreno**, L. M. López-Ramos and B. Beferull-Lozano, "A trainable approach to zero-delay smoothing spline interpolation," *IEEE Transactions on Signal Processing*, vol. 71, pp. 4317-4329, 2023.

Paper D **E. Ruiz-Moreno**, L. M. López-Ramos and B. Beferull-Lozano, "Consistent signal reconstruction from streaming multivariate time series," submitted to *IEEE Transactions on Signal Processing.*

The following paper was published during the Ph.D. degree but is not included in the thesis.

S. Mollaebrahim, B. Beferull-Lozano and **E. Ruiz-Moreno**, "Decentralized subspace projection for asymmetric sensor networks," in *2020 IEEE 92nd Vehicular Technology Conference (VTC2020-Fall).*

Contents

List of Figures

List of Tables

Abbreviations

Adam   Adaptive moments

ADC   Analog-to-digital converter

AR   Autoregressive

BPTT   Backpropagation through time

CC   Cumulative cost

CFA   Cost function approximation

DCOL   Differentiable convex optimization layer

DP   Dynamic programming

ERM   Empirical risk minimization

GRU   Gated recurrent unit

KAPSM   Kernel adaptive projected subgradient method

LSTM   Long short-term memory

MAE   Mean absolute error

MSE   Mean square error

NORMA   Naive online R_{reg} minimization algorithm

RERM   Regularizaed empirical risk minimization

RK   Reproducing kernel

RKHS   Reproducing kernel Hilbert space

VAR   Vector autoregressive

WORM   Windowed online regularized cost minimization

Chapter 1

Introduction

The study of temporal measurements is a matter of great importance in many domains of applied science and engineering, such as statistics or signal processing. In this context, *time series* are the object of study, consisting of a sequence of data samples or observations indexed in time order. Accordingly, the study techniques are usually referred to as methods for time-series analysis.

This chapter motivates the importance of methods for time series analysis, reveals some related research gaps, and discusses possible research avenues to address them.

1.1   Motivation

Streaming data can be understood as a time series whose terms are continuously generated, possibly from different sources. Streaming data is of paramount value and necessity in the present day, being applied in a broad range of industries, including finance, transportation, or the industrial Internet of Things, among others [15].

In many practical tasks involving streaming data, the ground truth behind the process being measured is unknown. Instead, one only has access to the received time series observations. These observations may be very different in nature for each task, for instance, real-valued, discrete, or even symbolic data.
When dealing with physical phenomena, the observed data samples are collected through a data *acquisition system*, usually comprised of a set of networked sensors and a computing unit. In this way, and due to the acquisition system limitations, the resulting observations typically consist of signal values sampled in time and quantized in amplitude, not necessarily uniformly for both.
A common goal for these tasks is to recover the ground truth process from the sequence of observations, for example, via some function approximation method, with the lowest possible error.

In most cases of interest, the required methods must account for large data volumes and, hence, be able to run online, sometimes under real-time constraints. Such constraints become prevalent when a safety-critical application requires instantaneous decisions based on the latest received data. Also, highly-competitive settings, such as algorithmic trading [16], give an advantage to those methods that minimize the latency between signal acquisition and order execution. In user-facing applications, avoiding delay is positive for user experience and system usability. For instance, virtual reality devices rely on real-time signal processing to ensure a seamless and immersive experience [17]. A fast response of actuators based on sensor

input is also beneficial on board drones, power grids, or wind turbines.

On the other hand, these methods must be suitable for time series, possibly multiple ones concurrently, of misaligned, non-uniformly sampled, and quantized intervals. This is because patterns from connected systems can improve processing results. Moreover, misaligned samples occur frequently, for example, in information fusion from different sensors [18] or when synchronization between devices is challenging.

Finally, and whenever possible, it is convenient that these methods exploit prior information about the underlying physical process, such as the smoothness and low model complexity shared by most physical signals and their spatiotemporal dynamics, to reduce the reconstruction error. Keeping the smooth nature and complexity of the signal reconstruction low is important to avoid artifacts or misleading interpretations that worsen its generality over unseen data.

1.2 Summarized State-of-the-art

Online optimization methods for signal reconstruction problems generating a series of function estimates from streaming data are useful to reduce computational complexity in large-scale problems [19], to dynamically adapt to new patterns in the data [20], and also to enable acting under real-time constraints [21].

Among them, online kernel-based penalized regression approaches stand out in popularity, mostly because they allow estimating functions of arbitrary complexity by finding non-linear patterns at a moderate computational cost [12, 22]. Within the considered framework of streamed time series, several sliding window kernel-based methods have been proposed to increase the reliability of the signal estimation by processing consecutive observations concurrently.

Arguably, the most suited approach to tackle our posed problem is the one presented in [23], which can provide a low-delay response by running a single processing iteration over the quantized data samples within a sliding window every time an observation is received.

However, it cannot guarantee that the estimated signal passes through all the observed quantization intervals, thus yielding function estimates inconsistent with the acquisition system. Another issue is that explicit regularization is not included, and therefore, the reduced model complexity associated with smooth and bounded underlying physical signals is not fully promoted. Lastly, function estimates modeled as a kernel expansion make it challenging to guarantee smoothness in a sequentially reconstructed signal; thus, using them as tracking systems rather than for smooth signal reconstructions is preferable.

On the other hand, online interpolation methods can be suitable candidates for overcoming the aforementioned issues. These methods typically use piecewise-defined functions to model a sequence of local signal estimates [24, 6]. Piecewise-modeled signal estimates can be updated every time an observation is received by assembling a new function piece while guaranteeing smoothness and a low-complexity model.

Several studies have explored the use of *spline*-based online interpolation methods from streaming data under real-time constraints [25, 26]. However, most of the existing spline-based online interpolation methods require a multi-step lookahead or sliding window mechanism, leading to a delay.

To the best of our knowledge, the only existing spline-based *zero-delay* interpolation method in the literature is the myopic approach (referred to as the "classical greedy approach" in [26]), a purely local method; thus, it does not exploit nonlocal information for improved signal reconstruction. For instance, it cannot accommodate the possible long-term spatiotemporal relationships within any multivariate time series.

Finally, consistent signal reconstruction methods exploit all available knowledge about the underlying physical signal, such as the signal class (e.g., bandlimited signals) and the acquisition process (e.g., analog-to-digital converters), to improve the error-rate decay as the sampling rate increases. These have been proven useful under offline settings [27, 28]. However, there is a research gap regarding online consistent signal reconstruction methods.

1.3   Research Questions

The motivation for this thesis work can be summarized in the following research questions derived from the current State-of-the-art:

RQ1 Can we design an online kernel-based method to reliably track the underlying physical signal from a univariate time series of quantization intervals?

RQ2 In addition to satisfying the low model complexity (implicitly imposed) in **RQ1**, if we require a more strict zero-delay response constraint and smoothness in the reconstructed signal, how can we exploit the temporal dependencies of univariate time series of data points to achieve an accurate spline-based signal reconstruction?

RQ3 Besides the previous requirements in **RQ1** and **RQ2**, how can we ensure a consistent signal reconstruction while exploiting the spatiotemporal relations of multivariate time series of quantization intervals to achieve an accuracy-improved spline-based signal reconstruction?

1.4   Outline

This Ph.D. thesis consists of a compendium of four papers. These papers have been either published or submitted for publication in peer-reviewed international conference proceedings and journals. They can be found in the appendices of this thesis. The thesis is divided into five chapters, including the present one.

Chapter 1 introduces the thesis topic, motivates its importance, summarizes the state-of-the-art, and exposes several research gaps and possible research paths to address them. The key theoretical concepts in the compendium of papers are discussed in **Chapter 2**. Next, **Chapter 3** summarizes papers A and B, which address the **RQ1**. Similarly, **Chapter 4** summarizes papers C and D, addressing the **RQ2** and the **RQ3**, respectively. Finally, **Chapter 5** concludes the paper and presents the future work.

Chapter 2

Background Theory

This thesis work makes use of several machine-learning techniques as auxiliary tools to attain real-time signal reconstruction from streaming data. This background theory chapter is meant to cover the basic idea behind those techniques, serving as a prelude for the compendium of papers attached in the appendices.

Accordingly, the connection, in the context of this Ph.D. thesis, between the concepts discussed next is purposely omitted for the sake of conciseness and is instead established and detailed in the ensuing thesis chapters.

2.1 Definitions and Terminology

This section covers key definitions and technical terminology used in this Ph.D. thesis. Specifically, some concepts related to acquisition systems, such as quantization or, more generally, uncertainty intervals, and some others related to signal reconstruction methods, such as delay, smoothness, and spline models.

2.1.1 Quantization

In signal processing, *quantization* is the process of partitioning the image domain of a given signal by mapping measured signal values, often from a continuous set, to values in a countable smaller set, possibly finite [29]. It is commonly used to reduce signal data size for more efficient storing, faster processing, and transmission.

   Any device or algorithm that performs quantization can be referred to as a quantizer. For example, a uniform mid-tread quantizer, which quantizes the signal values at any tth time stamp $\psi(x_t)$, with a quantization stepsize Δ, as

$$y_t = \text{round}\left(\frac{\psi(x_t)}{\Delta}\right)\Delta, \tag{2.1}$$

where round() denotes the round function. In this case, the tuple x_t, y_t, and Δ can describe any of the resulting quantized signal values.

2.1.2 Uncertainty intervals

An *uncertainty interval* can be defined as a continuous region in the image domain of the underlying signal being acquired that reflects our confidence, or knowledge,
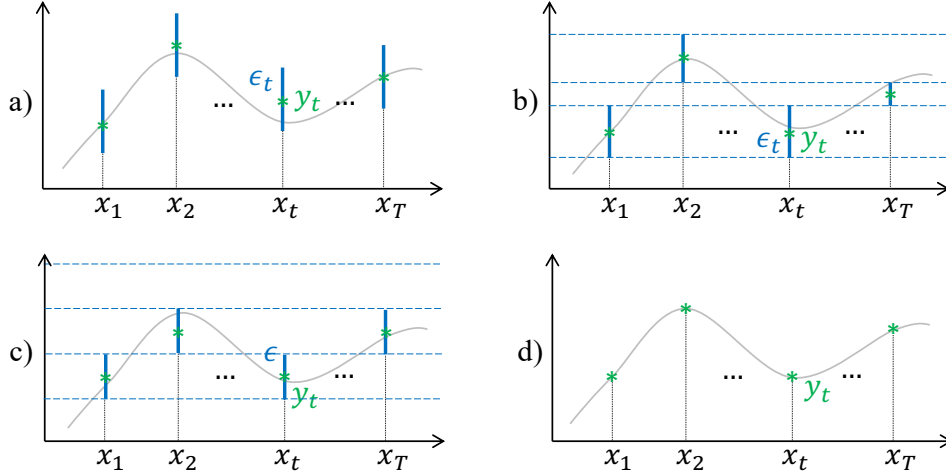
Figure 2.1: Example streams of uncertainty intervals. In a), we have intervals of different sizes that can overlap (in the image domain), thus not necessarily partitioning the image domain; an example of such data is safety regions in a trajectory curve. In b) and c), we respectively show a stream of non-uniform and uniform ($\epsilon_t = \epsilon$ for all time steps) quantization intervals. Finally, in d), we present a stream of data points, where $\epsilon_t = 0$, for all time steps.

about the signal acquisition at a given time stamp. It is also the primary data type considered in this Ph.D. work, which we sometimes refer to as a data sample.

For a formal definition, consider a stream of uncertainty intervals; then, the tth observed uncertainty interval is expressed as $\boldsymbol{o}_t = [x_t, y_t, \epsilon_t]^\top$, where $x_t \in \mathbb{R}$ denotes its time stamp, $y_t \in \mathbb{R}$ is the center of the interval, and $\epsilon_t \in \mathbb{R}_+$ stands for the interval half step. This definition encompasses many data types encountered in practice, such as safety regions, quantization intervals, and data points (data samples with no uncertainty in the image domain), as illustrated in Fig. 2.1.

Even though the methods proposed in Paper A, Paper B, and Paper D are designed to work with any uncertainty interval, in this Ph.D. thesis, we restrict to (not necessarily uniform) quantization intervals as a common use case.

2.1.3 Zero-delay Response

In machine learning, online learning is a well-established paradigm devised for sequential data [30, 31]. Online learning methods are typically used to alleviate the computational complexity of "offline" algorithms, i.e., those techniques designed to work over the entire data set at once, usually over huge data sets in practice. Similarly, they are also used in tasks where adapting to new patterns in the data or satisfying real-time constraints is required.

In the context of online learning for signal reconstruction, delay can be understood as i) the difference between the time a sample is received and the time its corresponding signal reconstruction is delivered and ii) the number of samples that are necessary to process before a signal reconstruction is proposed.
This thesis adopts the concept of delay in terms of the number of samples. Therefore, an algorithm sequentially reconstructing a signal without waiting for the next sample is deemed as operating under a *zero-delay* response.

2.1.4 Smoothness

Given a univariate function, its *smoothness* refers to the number of continuous derivatives it has over the interior of its domain. At the very minimum, a function could be considered smooth if it is differentiable (and hence continuous) everywhere. At the other end, it might also possess derivatives of all orders in its domain, in which case it is said to be infinitely differentiable.

In addition, the term smoothness is usually used as the opposite of the concept of *roughness*, a derivative-based metric that quantifies the model complexity of certain smooth functions [32]. Mathematically, given a space of functions $\mathcal{F}_\rho$ over the domain $\mathrm{dom}(f)$ with $\rho - 1$ absolutely continuous derivatives and with a squared integrable ρ-th derivative, their roughness $r_\rho : \mathcal{F}_\rho \to \mathbb{R}$ is computed as

$$r_\rho(f) = \int_{\mathrm{dom}(f)} \left(D_x^\rho f(x)\right)^2 \, dx. \tag{2.2}$$

In summary, the word "smoothness" may refer to the fact that the roughness of a function is minimized or to the fact that the function has a certain number of continuous derivatives. This Ph.D. thesis deals with both these distinct but related concepts since, to minimize the ρ-roughness, we need the function to be at least $(\rho - 1)$-smooth. Whenever the meaning is unclear from the context, the latter will be referred to as the degree of smoothness.

2.1.5 Spline models

By convention, a spline model is defined as a piecewise polynomial function. In this thesis work, we denote any one-dimensional spline composed of T function pieces as

$$f_T(x) = \begin{cases} g_1(x), & \text{if } x_0 < x \le x_1, \\ g_2(x), & \text{if } x_1 < x \le x_2, \\ \vdots \\ g_T(x), & \text{if } x_{T-1} < x \le x_T, \end{cases} \tag{2.3}$$

where every tth function piece, or function section, $g_t : (x_{t-1}, x_t] \to \mathbb{R}$ is a polynomial of the form

$$g_t(x) = \boldsymbol{a}_t^\top \boldsymbol{p}_t(x), \tag{2.4}$$

with combination coefficients $\boldsymbol{a}_t \in \mathbb{R}^{d+1}$ and basis vector function $\boldsymbol{p}_t : (x_{t-1}, x_t] \to \mathbb{R}^{d+1}$ given by

$$\boldsymbol{p}_t(x) = [1, (x - x_{t-1}), \dots, (x - x_{t-1})^d]^\top, \tag{2.5}$$

where the positive integer d denotes the order of the spline.
On the other hand, a spline model is said to have a degree of smoothness $\varphi \in \mathbb{N}$ if it has φ continuous derivatives over the interior of its domain $\mathrm{dom}(f_T) = \bigcap_{t=1}^{T}(x_{t-1}, x_t]$. Some spline configurations with different order and degree of smoothness are illustrated in Fig. 2.2.

It is worth mentioning that there are many alternative ways of equivalently representing the spline model in (2.3); for instance, by using overlapping basis functions implicitly accommodating the smoothing constraints, such as in B-splines or cardinal splines models [33]. Although these alternative representations are equivalent,

(a) Spline of order d with discontinuous first derivative.

(b) Spline of order d with at least the zeroth and first derivatives continuous.

Figure 2.2: Example of interpolation splines. The asterisks represent data points; in this case, they also coincide with the spline knots (joint points).
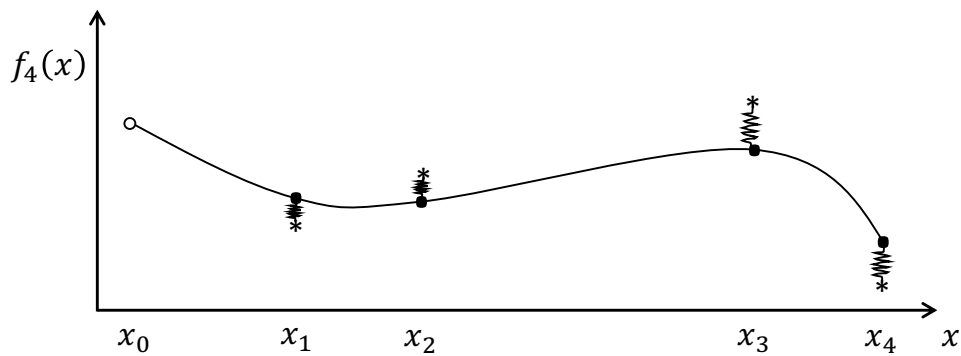
Figure 2.3: Physical analogy of a smoothing spline [6]. The asterisks represent data points, while the black points represent the spline knots. Here, the smoothing spline can be viewed as a strip hooked up to the data points by ideal springs (for small elongations).

it is preferable to use the piecewise representation for some tasks since we can directly impose continuity constraints into the model representation. Otherwise, the continuity constraints are embedded into the basis expansion, hence making such a task cumbersome, as further discussed in Paper C.

2.1.5.1   Smoothing spline interpolation

Let us analyze the naming conventions that constitute the notion of "smoothing spline interpolation" by parts. The term "smoothing" refers to a controlled trade-off between fitting data points and proposing a smooth (low complexity) signal estimate. By "interpolation," one usually refers to estimating new data within the time domain delimited by the range of known data points rather than simply connecting them.

On the other hand, "interpolation splines" and "smoothing splines" may seem completely different constructs, especially from an implementation perspective, since the former passes through the known data points and the latter does not necessarily. However, by construction, "interpolation splines" can be seen as a particular limit case of "smoothing splines."

This is because the smoothing spline is defined as

$$f_T(x) = \arg \min_{f \in \mathcal{F}_\rho} \sum_{t=1}^{T} (f(x_t) - y_t)^2 + \frac{\eta}{2} \, r_\rho(f), \tag{2.6}$$

where the class of functions $\mathcal{F}_\rho$ and roughness metric r_ρ, are the same as introduced in Sec. 2.1.4. Thus, notice that as the parameter $\eta \to 0$, the smoothing spline converges to the interpolating spline. Some intuition behind the smoothing spline, defined in (2.6), is provided in Fig. 2.3.

2.2   Proximal Algorithms

Suppose we are given a differentiable convex function $f : \mathbb{R}^D \to \mathbb{R}$ whose minimum value(s) cannot be found in closed form or that finding them is computationally prohibitive. In this case, one can resort to iterative methods [34].

Roughly, an iterative method is an algorithm that generates a sequence of iterates, e.g., $\{\boldsymbol{x}_1, \boldsymbol{x}_2, \dots\}$. These algorithms are typically devised so that the sequence of iterates converges to an optimal solution, denoted as $\boldsymbol{x}^* = \arg \min_{\boldsymbol{x} \in \mathbb{R}^D} f(\boldsymbol{x})$, as the number of iterations, or steps, increases.

Arguably, the most emblematic iterative algorithm is the *gradient descent* algorithm [35], whose core program is exposed next in **Algorithm 1**.

Algorithm 1 Gradient Descent

Input: Initial $\boldsymbol{x}_1$, a convergence criterion and a step-size sequence $\{\eta_j\}_j$ or rule.
  1: **for** $i = 1, 2, \dots$ **do**
  2:    $\boldsymbol{x}_{i+1} = \boldsymbol{x}_i - \eta_i \nabla f(\boldsymbol{x}_i)$.
  3:    Break if the convergence criterion is satisfied.
  4: **end for**
Output: The last iterate.

Figure 2.4: Gradient descent as the minimization of successive quadratic local approximations for $D = 1$.

Figure 2.5: Example of subgradients for $D = 1$. At x_A the function f is differentiable, and g_A is the unique derivative. On the other hand, at the point x_B the function f is non-differentiable and has many subgradients (for convenience, only two subgradients $g_{B,1}$ and $g_{B,2}$ are shown).

The descending step in **Algorithm 1** (line 2) admits several interpretations. One of them, which allows for a more in-depth comparison analysis later on, is the following: From the second-order Taylor approximation around $\boldsymbol{x}_i$, i.e.,

$$f(\boldsymbol{x}) \simeq \tilde{f}(\boldsymbol{x}) = f(\boldsymbol{x}_i) + \nabla f(\boldsymbol{x}_i)^\top(\boldsymbol{x} - \boldsymbol{x}_i) + \frac{1}{2}(\boldsymbol{x} - \boldsymbol{x}_i)^\top \mathbf{H}(f(\boldsymbol{x}_i))(\boldsymbol{x} - \boldsymbol{x}_i), \quad (2.7)$$

we can replace the Hessian $\mathbf{H}(f(\boldsymbol{x}_i))$ for a more computationally friendly value $\frac{1}{\eta_i}\mathbf{I}_D$, where $\mathbf{I}_D$ is the identity matrix of size D and η_i is a positive real parameter closely related to the curvature of a quadratic approximation, obtaining

$$f(\boldsymbol{x}) \simeq \tilde{f}_{\eta_i}(\boldsymbol{x}) = f(\boldsymbol{x}_i) + \nabla f(\boldsymbol{x}_i)^\top(\boldsymbol{x} - \boldsymbol{x}_i) + \frac{1}{2\eta_i}\|\boldsymbol{x} - \boldsymbol{x}_i\|_2^2. \quad (2.8)$$

Then the next iterate is given by $\boldsymbol{x}_{i+1} = \arg\min_{\boldsymbol{x} \in \mathbb{R}^D} \tilde{f}_{\eta_i}(\boldsymbol{x})$, which can be found as the stationary point satisfying that

$$\nabla \tilde{f}_{\eta_i}(\boldsymbol{x}_{i+1}) = 0, \quad (2.9a)$$

$$\text{or, equivalently, } \nabla f(\boldsymbol{x}_i) + \frac{1}{\eta_i}(\boldsymbol{x} - \boldsymbol{x}_i) = 0. \quad (2.9b)$$

The corresponding visual representation is provided in Fig. 2.4.

In many applications, the objective function f may be non-differentiable. In such cases, it is not possible to compute the gradient everywhere; thus the gradient descent method cannot be used. On the contrary, the *subgradient* method is a simple algorithm for minimizing non-differentiable convex functions [36].

The main difference with respect to gradient descent lies in the "descending" step in **Algorithm 1** (line 2). Instead, the subgradient method iterates as

$$\boldsymbol{x}_{i+1} = \boldsymbol{x}_i - \eta_i \, \boldsymbol{g}_i, \quad (2.10)$$

where $\boldsymbol{g}_i$ is any subgradient of f at $\boldsymbol{x}_i$, that is, any vector in $\mathbb{R}^D$ such that

$$f(\boldsymbol{x}) \geq f(\boldsymbol{x}_i) + \boldsymbol{g}_i^\top(\boldsymbol{x} - \boldsymbol{x}_i), \quad (2.11)$$

is satisfied, as illustrated in Fig. 2.5.

However, the subgradient method is not a descent method; that is, the function value at the next iterate can increase, hence slowing the convergence compared to gradient descent methods.

Differently, *proximal* algorithms can deal with a non-smooth function f while matching the gradient descent algorithm convergence rate [37]. They make use of the proximal operator, in this case, a mapping $\text{prox}_f^\eta : \boldsymbol{v} \in \mathbb{R}^D \mapsto \text{prox}_f^\eta(\boldsymbol{v}) \in \mathbb{R}^D$ that comprises between minimizing f and being near to a given point $\boldsymbol{v}$, where η can be seen as a trade-off parameter between these terms.

In proximal algorithms, the equivalent to the descending step in **Algorithm 1** (line 2) at each iteration consists of evaluating the proximity (prox) operator of a function, which involves solving a convex optimization subproblem. This subproblem can be solved with standard optimization methods, but it often admits closed-form representations or can be solved very quickly with specialized methods [37].

As an example, consider a convex non-differentiable function f that can be decomposed as $f(\boldsymbol{x}) = g(\boldsymbol{x}) + h(\boldsymbol{x})$, where g is convex and differentiable and h is also

convex but not necessarily differentiable.

Similarly to the gradient descent method, we can perform a Taylor expansion of only the differentiable term in the objective, i.e., we expand g around the iterate $\boldsymbol{x}_i$ and use the stationary point of the resulting approximation as the next iterate. Mathematically,

$$\boldsymbol{x}_{i+1} = \arg \min_{\boldsymbol{x} \in \mathbb{R}^D} \left\{ g(\boldsymbol{x}_i + \nabla g(\boldsymbol{x}_i)^\top (\boldsymbol{x} - \boldsymbol{x}_i) + \frac{1}{2\eta} \|\boldsymbol{x} - \boldsymbol{x}_i\|_2^2 + h(\boldsymbol{x}) \right\} \qquad (2.12a)$$

$$= \arg \min_{\boldsymbol{x} \in \mathbb{R}^D} \left\{ \frac{1}{2\eta} \|\boldsymbol{x} - (\boldsymbol{x}_i - \eta \nabla g(\boldsymbol{x}_i))\|_2^2 + h(\boldsymbol{x}) \right\} \qquad (2.12b)$$

$$\triangleq \mathrm{prox}_h^\eta (\boldsymbol{x}_i - \eta \nabla g(\boldsymbol{x}_i)), \qquad (2.12c)$$

where the algebraic step (2.12c) applies the definition of the prox operator.

This algorithm is known in the literature as the proximal gradient descent method or the proximal forward-backward iterative scheme [38].

2.3 Learning with Kernels

Empirical risk minimization (ERM) is arguably one of the most popular supervised approaches to designing learning algorithms [39]. The general idea consists of minimizing the empirical loss

$$\frac{1}{T} \sum_{t=1}^T \ell(f(x_t), y_t), \qquad (2.13)$$

as a proxy for the expected loss

$$\mathbb{E}_{x,y \sim p(x,y)} [\ell(f(x), y)], \qquad (2.14)$$

since, in most cases, the expectation cannot be obtained due to a prohibitive computational cost or because the joint probability $p(x, y)$ is unknown.

To put the above algorithmic idea into practice, one needs to fix a suitable function space and then minimize the empirical loss over it. Additionally, the function space should allow feasible computations and be rich (encompassing a large number of candidate functions) when the model complexity is unknown a priori. Here, the space of linear functions is conceivably the simplest example.

On the other hand, if the function space is too rich, solving an ERM problem can cause overfitting, i.e., poor generalization to unseen data. Fortunately, regularization techniques can enforce stability (by reducing the model complexity) and promote generalization [40]. One possible way to do this is by adding a regularization term (penalizing a metric of the complexity of the learned function) to the empirical loss.

2.3.1 Beyond Linear Models

Suppose we have a set of data points $\mathcal{S} = \{(x_t, y_t)\}_{t=1}^T$, where $x_t, y_t \in \mathbb{R}$, to which we want to fit a linear model $f : x \in \mathbb{R} \mapsto wx \in \mathbb{R}$, with $w \in \mathbb{R}$. Then, if certain conditions, such as a linear relationship between the independent and dependent variables in the data set, are not met, the linear model provides inaccurate solution estimates, as shown in Fig. 2.6.

(a) Accurate fit.

(b) Underfit. The space of linear functions is not rich enough to properly capture the data relationships.

Figure 2.6: Linear model fit via ridge regression for two different data sets.

Figure 2.7: Geometrical interpretation of the feature map. On the left is the input space, where the model f is seen as a non-linear function. On the right is the feature space, where the model is seen as linear with respect to the coefficients in $\boldsymbol{w}$.

Figure 2.8: Fit of a quadratic function as a linear combination of features via ridge regression.

At this point, one may wonder whether it is possible to learn a linear model from a richer function space to overcome underfitting. The answer is yes.

In general, we can use any feature map $\boldsymbol{\phi}(\boldsymbol{x}) = [\varphi_1(\boldsymbol{x}), \ldots, \varphi_F(\boldsymbol{x})]^\top$, where $\varphi_i : \mathbb{R}^D \to \mathbb{R}$ and $\boldsymbol{\phi} : \mathbb{R}^D \to \mathbb{R}^F$ (usually for $F \gg D$). This leads to non-linear models in the input space but linear in the feature space of the form

$$f(\boldsymbol{x}) = \boldsymbol{w}^\top \boldsymbol{\phi}(\boldsymbol{x}) = \sum_{i=1}^{F} w_i \varphi_i(\boldsymbol{x}), \tag{2.15}$$

where $\boldsymbol{w} \in \mathbb{R}^F$. This is illustrated in Fig. 2.7.

As an example, the regression problem posed before in Fig. 2.6b corresponds to $D = 1$ and can be solved with a feature map in $F = 3$ of the form $\boldsymbol{\phi}(x) = [1, x, x^2]^\top$, as shown in Fig. 2.8.

On the other hand, thanks to the Representer Theorem [41], we know that for any given loss $\ell : \mathbb{R}^2 \to \mathbb{R}$ the solution to the optimization problem

$$\boldsymbol{w}^* = \arg \min_{\boldsymbol{w} \in \mathbb{R}^F} \frac{1}{T} \sum_{t=1}^{T} \ell(\boldsymbol{\phi}(\boldsymbol{x}_t)^\top \boldsymbol{w}, y_t) + \frac{\eta}{2} \|\boldsymbol{w}\|_2^2, \tag{2.16}$$

can be written as $\boldsymbol{w}^* = \boldsymbol{\alpha}^\top \boldsymbol{\Phi}$, where $\boldsymbol{\Phi} = [\boldsymbol{\phi}(\boldsymbol{x}_1), \ldots, \boldsymbol{\phi}(\boldsymbol{x}_T)]^\top \in \mathbb{R}^{T \times F}$ is the data matrix in the feature space, and $\boldsymbol{\alpha} \in \mathbb{R}^T$ is some real-valued vector of coefficients. Then, the solution model is of the form

$$f(\boldsymbol{x}) = \boldsymbol{\phi}(\boldsymbol{x})^\top \boldsymbol{w}^* \tag{2.17a}$$

$$= \boldsymbol{\phi}(\boldsymbol{x})^\top \boldsymbol{\Phi}^\top \boldsymbol{\alpha} \tag{2.17b}$$

$$= \sum_{t=1}^{T} \alpha_t \boldsymbol{\phi}(\boldsymbol{x})^\top \boldsymbol{\phi}(\boldsymbol{x}_t), \tag{2.17c}$$

that is, $f(\boldsymbol{x})$ is expressed by using inner products between feature vectors.

In fact, we do not even need to explicitly define a feature map $\boldsymbol{\phi}$, but only a reproducing kernel function $k(\boldsymbol{x}, \boldsymbol{x}_t) = \boldsymbol{\phi}(\boldsymbol{x})^\top \boldsymbol{\phi}(\boldsymbol{x}_t)$. This procedure is commonly referred to in the literature as the *kernel trick* [42].

Finally, every kernel is associated with a unique reproducing kernel Hilbert space (RKHS), and conversely. The RKHS theory is outstanding; the underlying definitions are conceptually simple yet powerful and broadly applicable in practice [43].

2.4 Finite Horizon Stochastic Dynamic Programming

There exist several classes of techniques for optimal decision-making in dynamic state spaces. Classical control theories, discrete-time dynamic programming (DP), and machine learning-based approaches, often referred to as reinforcement learning [44], aim to determine the best possible action or sequence of actions in a given situation, considering the evolution of states of a dynamic environment over time. In this section, we will delve into deterministic and stochastic DP.

Deterministic dynamic programming problems involve a discrete-time dynamic system of the form

$$\boldsymbol{s}_{t+1} = F_t(\boldsymbol{s}_t, \boldsymbol{a}_t), \text{ for } t = 0, \dots, T-1, \tag{2.18}$$

where t is the time index, $\boldsymbol{s}_t$ is the tth state, $\boldsymbol{a}_t$ is the tth action, F_t describes the state update mechanism at time step t, and T is the total number of steps or horizon.

The set of all possible states $\boldsymbol{s}_t$ is referred to as the state space at time t. It can be any set and can depend on t; this generality is one of the great strengths of the DP methodology [45]. Similarly, the set of all possible actions $\boldsymbol{a}_t$ is called the action space at time t. Again, it can be any set and can depend on t. The problem also involves an additive real-valued cost function in the sense that the cost incurred at time t, denoted by $g_t(\boldsymbol{s}_t, \boldsymbol{a}_t)$, accumulates over time, and may depend on t. For a given initial state $\boldsymbol{s}_0$, the total cost of an action sequence $\{\boldsymbol{a}_0, \dots, \boldsymbol{a}_{T-1}\}$ is

$$J(\boldsymbol{s}_0; \boldsymbol{a}_0, \dots, \boldsymbol{a}_{T-1}) = g_T(\boldsymbol{s}_T) + \sum_{t=0}^{T-1} g_t(\boldsymbol{s}_t, \boldsymbol{a}_t), \tag{2.19}$$

where $g_T(\boldsymbol{s}_T)$ is a terminal cost incurred at the end of the process.

The objective is to minimize the total cost (2.19) over all sequences of feasible actions.

The finite-horizon stochastic DP problem differs from the deterministic version primarily in the nature of the discrete-time dynamic system that governs the evolution of the states. This system includes a random variable w_t, which accounts for the uncertainty and is characterized by a probability distribution $p_t(\cdot|\boldsymbol{s}_t, \boldsymbol{a}_t)$ that may depend explicitly on $\boldsymbol{s}_t$ and $\boldsymbol{a}_t$, but is independent of prior values of $w_{t-1}, \dots, w_0$. In this case, the state update mechanism is of the form

$$\boldsymbol{s}_{t+1} = F_t(\boldsymbol{s}_t, \boldsymbol{a}_t, w_t), \text{ for } t = 0, \dots, T-1, \tag{2.20}$$

and the cost per time step g_t also depends on the corresponding random uncertain variable w_t. Additionally, one optimizes the expected cost not over sequences of actions but rather over policies or decision strategies that consist of a sequence of functions $\pi = \{\mu_0, \dots, \mu_{T-1}\}$ mapping states into actions.

Policies are more general objects than sequences of actions. In the presence of stochastic uncertainty, policies often incur a lower cost than any fixed sequence of actions since they allow choices of actions that incorporate knowledge of the current state. Without this knowledge, the agent may not act appropriately, and the cost can be adversely affected. In other words, policies encode the necessary knowledge

Random transition
$$\boldsymbol{s}_{t+1} = F_t(\boldsymbol{s}_t, \mu_t(\boldsymbol{s}_t), w_t)$$

Random cost
$$g_t(\boldsymbol{s}_t, \mu_t(\boldsymbol{s}_t), w_t)$$

Terminal cost
$$g_T(\boldsymbol{s}_T)$$

Figure 2.9: Illustration of a T-steps finite horizon stochastic dynamic programming problem.

to adapt to any state of the environment and minimize the expected cost. The general idea is visually summarized in Fig. 2.9.

Thus, for given cost functions g_t for $t = 0, \dots, T$, the expected total incurred cost by following a given policy π starting at $\boldsymbol{s}_0$ is

$$J_\pi(\boldsymbol{s}_0) = \mathbb{E}\left[g_T(\boldsymbol{s}_T) + \sum_{t=0}^{T-1} g_t(\boldsymbol{s}_t, \mu_t(\boldsymbol{s}_t), w_t)\right], \qquad (2.21)$$

where the expectation is taken over all random variables $\boldsymbol{s}_t$ and w_t. An optimal policy, denoted as π^*, is the one that minimizes (2.21), that is, $\pi^* = \arg\min_{\pi \in \Pi} J_\pi(\boldsymbol{s}_0)$, where Π is the set of all policies.

On the other hand, Bellman's principle of optimality says that the tail of an optimal sequence is optimal for the tail subproblem [46]. That is, the tail $\{\boldsymbol{a}_t^*, \cdots, \boldsymbol{a}_{T-1}^*\}$ of an optimal sequence of actions $\{\boldsymbol{a}_0^*, \cdots, \boldsymbol{a}_{T-1}^*\}$ is optimal for the tail subproblem that starts at the state $\boldsymbol{s}_t^*$ of the optimal trajectory of states $\{\boldsymbol{s}_1^*, \cdots, \boldsymbol{s}_T^*\}$, as shown in Fig. 2.10.
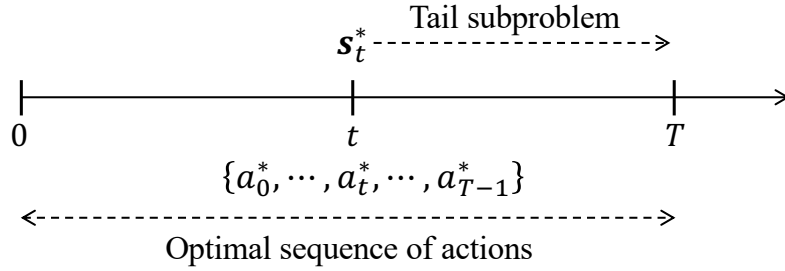
Figure 2.10: Illustration of the principle of optimality.

Based on the principle of optimality, we can compute the tth optimal action $\boldsymbol{a}_t^* = \mu_t^*(\boldsymbol{s}_t)$ by a backward induction process, as

$$\mu_t^*(\boldsymbol{s}_t) \in \arg\min_{\boldsymbol{a} \in \mathcal{A}(\boldsymbol{s}_t)} \mathbb{E}\left[g_t(\boldsymbol{s}_t, \boldsymbol{a}, w_t) + J_{t+1}^*(F_t(\boldsymbol{s}_t, \boldsymbol{a}, w_t))\right], \qquad (2.22)$$

where $\mathcal{A}(\boldsymbol{s}_t)$ denotes the set of feasible actions from the state $\boldsymbol{s}_t$, and J_{t+1}^* is the optimal cost-to-go from the $(t+1)$th step.

However, computing the optimal cost-to-go can be computationally expensive. This motivates (suboptimal) approximation techniques, where, for example, one replaces J_t^* with an approximate cost-to-go function, denoted as $\tilde{J}_t$, that is less complex to obtain.

2.4.1 Policy Parametrization Through Cost Parametrization

Suppose we start a backward induction procedure with a parametric cost-to-go function approximation $\tilde{J}_t(\boldsymbol{s}_t; \boldsymbol{\theta})$, for every step t, where $\boldsymbol{\theta}$ is a vector of P trainable parameters. Then, we can obtain an approximation in policy space of parametric policies $\tilde{\pi} = \{\tilde{\mu}_0, \ldots, \tilde{\mu}_{T-1}\}$ defined through the following optimization problem

$$\tilde{\mu}_t(\boldsymbol{s}_t; \boldsymbol{\theta}) \in \arg \min_{\boldsymbol{a} \in \mathcal{A}(\boldsymbol{s}_t)} \mathbb{E}\left[g_t(\boldsymbol{s}_t, \boldsymbol{a}, w_t) + \tilde{J}_{t+1}\left(F_t(\boldsymbol{s}_t, \boldsymbol{a}, w_t); \boldsymbol{\theta}\right)\right]. \tag{2.23}$$

An important issue here is the computation of the expected value and the minimization over $\boldsymbol{a} \in \mathcal{A}(\boldsymbol{s}_t)$. Both of these operations may involve substantial computational load, which is of particular concern when the minimization is to be performed online.

One possibility to eliminate the expected value from the expression (2.23) is to choose a typical value (depending on the application) of uncertain variable $\tilde{w}_t$, and use the action that solves the resulting deterministic problem. On top of this, both the cost functions g_t and cost-to-go function approximations $\tilde{J}_t$ can be chosen so the optimization problem (2.23) admits a closed-form solution. Moreover, g_t and $\tilde{J}_t$ can also be chosen to be stationary, i.e., $g \equiv g_t = g_{t+1}$ and $\tilde{J} \equiv \tilde{J}_t = \tilde{J}_{t+1}$, for all time steps, which is a more suitable choice for tasks of variable horizon (undetermined number of time steps), such as those involving streaming data.

2.5 Recurrent Neural Networks

Recurrent neural networks (RNNs) are a family of neural networks designed to process sequential data [47]. RNNs find applications in domains ranging from time series forecasting and natural language processing to multimedia analysis and control systems. They specialize in handling data sequences with a length that would be impractical for other types of neural networks. Furthermore, most RNNs have the ability to process sequences of varying length.

The key idea behind RNNs is sharing parameters across different parts of the model. If an RNN had separate parameters for each sequence data term, it would neither be able to generalize to sequences of arbitrary length nor to share learned behaviors across different terms. Such sharing is particularly important when a specific piece of information can occur at multiple positions within the data sequence. There is a wide variety of RNNs; some representative examples are shown in Fig. 2.11.

Computing the gradient through an RNN is straightforward. One can apply the backpropagation algorithm to the unrolled computational graph of the RNN. This is known as backpropagation through time (BPTT) [48]. Gradients obtained by BPTT may then be used together with any gradient-based optimizer to train an RNN.

The mathematical challenge of learning long-term dependencies in RNNs is that gradients propagated over many steps tend to either vanish or explode [49]. In the scalar case, this implies multiplying a weight w by itself many times. The product $w \cdot w \cdots w$ will either vanish or explode depending on the magnitude of w.

On the other hand, even if we assume that the parameter values are such that the RNN is stable, the difficulty with long-term dependencies arises from the exponentially smaller weights given to long-term interactions (involving the multiplication of

(a) Many-to-many RNN   (b) Many-to-one RNN   (c) One-to-many RNN

Figure 2.11: Common RNN architectures. From darker blue (bottom) to lighter (top), we have the input, the hidden layer, and the output. The dashed empty cell represents an initial hidden state. Some practical applications for the architectures in Fig. 2.11a, Fig. 2.11b, and Fig. 2.11c are time-series forecasting [7], sentiment analysis [8], and image captioning [9], respectively.

many Jacobians) compared to short-term ones. Whenever the model can represent long-term dependencies, the gradient related to a long-term interaction has an exponentially smaller magnitude than the gradient related to a short-term interaction. This means not that it is impossible to learn but that it might take a very long time to learn long-term dependencies because these dependencies will tend to be hidden by the fluctuations arising from short-term dependencies.

Nowadays, one of the most effective models for data sequences used in practical applications is the gated RNNs. They are based on creating paths through time with derivatives that seldom vanish nor explode. This is done through a mechanism that adds or subtracts (instead of multiplying) relevant or irrelevant information to the shared state. Noticeable examples of gated RNNs are the long short-term memory (LSTM) [50] architecture and the gated recurrent unit (GRU) [51].

2.6 Consistent Signal Estimates

Given a bandlimited signal ψ, it is well known that we can ensure a reversible discretization in time by sampling at, or above, its Nyquist sampling rate [52, 53]. In practice, one cannot avoid an additional discretization in amplitude due to the physical limitations of acquisition systems, causing an irreversible loss of information and, therefore, making a perfect signal reconstruction no longer possible. These physical limitations of acquisition systems are typically modeled through quantization mappings.

Even though quantization is a deterministic operation, the signal reconstruction error can be analyzed effectively from a stochastic point of view when certain assumptions are met [54, 55]. Under these assumptions, low-pass filtering reduces the mean squared error (MSE) of a reconstructed signal f by a factor proportional to the squared quantization step size Δ^2 and inversely proportional to the oversampling

Figure 2.12: Geometric representation of the set of consistent signal estimates $\mathcal{Q}(\boldsymbol{y}) \cap \mathcal{V}$.

ratio R. That is,

$$\mathbb{E}\left[\|f - \psi\|_2^2\right] \propto \frac{\Delta^2}{R}, \tag{2.24}$$

where the expectation is taken over the amplitude of the reconstruction error, modeled as a uniform random variable [54].

On the other hand, deterministic analyses, based on deterministic bounds over the quantization levels, show that it is possible to reconstruct quantized signals with a squared norm error that is asymptotically proportional to the square of the quantization step size and inversely proportional to the square of the oversampling ratio, that is

$$\|f - \psi\|_2^2 = \mathcal{O}\left(\frac{\Delta^2}{R^2}\right), \tag{2.25}$$

via *consistent signal estimates* [27, 56].
In this case, a consistent signal estimate f is a bandlimited function that passes through all quantization intervals obtained after acquiring ψ.
More rigorously, the acquired signal $\boldsymbol{\psi} = [\psi(x_1), \ldots, \psi(x_T)]^\top \in \mathbb{R}^T$ is known only through its quantized version $\boldsymbol{y} \triangleq q(\boldsymbol{\psi}) = [q(\psi(x_1)), \ldots, q(\psi(x_T))]^\top$. Thus, the full information available about the acquisition of ψ is that $\boldsymbol{\psi} \in \mathcal{Q}(\boldsymbol{y})$, where $\mathcal{Q}(\boldsymbol{y})$ is the set of possible input signals with quantized version $\boldsymbol{y}$. As an example, when q is a uniform mid-tread quantizer with quantization step Δ, i.e., $q(\psi(x_t)) = \text{round}(\psi(x_t)/\Delta) \cdot \Delta$, $\mathcal{Q}(\boldsymbol{y})$ is a hypercube in $\mathbb{R}^T$ of side Δ, see Fig. 2.12. On the other hand, we additionally know that ψ belongs to the space of bandlimited signals $\mathcal{V}$. Hence, the elements of $\mathcal{Q}(\boldsymbol{y}) \cap \mathcal{V}$ (illustrated in Fig. 2.12) are referred to as consistent signal estimates of ψ.

The behavior presented in (2.25) indicates that when a uniformly quantized and oversampled signal is reconstructed using consistent estimates, the reconstruction error can be reduced at the same asymptotic rate by either increasing the oversampling ratio R or decreasing the quantization step size Δ. This balanced signal reconstruction error-rate decay has important practical implications. For instance, in practice, it is more convenient to increase the oversampling ratio because the implementation cost (including monetary cost) of the analog circuitry necessary for

halving the quantization step size is much higher than that for doubling the over-sampling ratio [57].

In more general settings, one can consider non-uniform quantization (the hyper-cube in Fig. 2.12 becomes a rectangular hyper-parallelepiped), and a different signal space than $\mathcal{V}$; for example, the space of non-bandlimited signals with finite rate of innovation [58], such as signals modeled by a train of delta functions or splines. This can be done so that the trend behavior in (2.25) is met [59, 28].

Chapter 3

Kernel-based Signal Tracking from Univariate Time Series of Quantization Intervals

This chapter summarizes Paper A and Paper B and delves into the synergies between them for the task of tracking univariate time series of quantization intervals via kernel-based online regression. The background theory relevant to this chapter can be found in Sec. 2.1, 2.2, and 2.3.

3.1   Motivation

Regression techniques based on reproducing kernels have shown remarkable success under stochastic optimization frameworks [60, 22]. However, most kernel-based methods in the literature are unsuited for tracking streamed quantized data via regression under dynamic environments. This shortcoming can primarily be attributed to the lack of flexibility in selecting the best-suited functional cost or to the mathematical structure of the regression problem not being adequately exploited.
In Paper A, we overcome these issues by introducing a novel method consisting of an online regression problem whose functional cost can accommodate quantized data and a stochastic proximal method that leverages its mathematical structure.

On the other hand, the choice of the reproducing kernel can significantly affect the performance of the proposed method in Paper A. Selecting an appropriate reproducing kernel can result in a computational challenge [61], especially when dealing with data-rich tasks without prior information about the solution domain.
We address this challenge in Paper B by devising a learning scheme combining several single kernel-based online methods, including the one proposed in Paper A. Our learning scheme uses a multi-kernel learning formulation [62], which can expand the solution space, thereby increasing the plausibility of finding higher-performance solutions while reducing the kernel-selection bias. Additionally, our learning scheme is parallelizable, allowing for the distribution of computational load across multiple computing units.

Figure 3.1: Visual representation of the online setting in Paper A and Paper B.

3.2   Problem Formulation

Online settings can be adopted to achieve low run-time complexity in exchange for a tolerable loss of accuracy [63, 64], e.g., by processing only a few data samples every iteration or discarding data samples after a few processing steps.

Paper A and Paper B share the same online setting core in their problem formulation, as summarized next.

Given a possibly infinite sequence $\mathcal{S}$ of quantization intervals, acquired at strictly increasing and non-necessarily equispaced time stamps, consider the following online setting, where the quantization intervals become available sequentially:

At each step n, the learner proposes a function estimate, which we denote as $f^{(n)}$. In response, an adversary penalizes the proposed function estimate with an incurred cost $\mathcal{C}(f^{(n)}; \mathcal{S}_L^{(n)})$, where $\mathcal{S}_L^{(n)} \subseteq \mathcal{S}$ denotes a sliding data window of (at most) length L. Then, the adversary reveals the nth data window $\mathcal{S}_L^{(n)}$ to the learner, which can be used at the next step $n + 1$. This online setting is illustrated in Fig. 3.1.

It is worth mentioning that, in this case, both the learner and the adversary know the "shape" of the functional cost $\mathcal{C}$. Moreover, at step n, the learner uses the previously proposed function estimate $f^{(n-1)}$ to compute $f^{(n)}$, except at $n = 1$ for which no quantization interval has been received yet and thus, $f^{(1)}$ is chosen arbitrarily.

3.3   Proposed Solution

The functional cost considered in Paper A penalizes a given trade-off between a weighted sum of distances to the quantization intervals within a sliding data window and the model complexity of the function estimate.

Particularly, each of the distances to the quantization intervals are computed through the distance functional

$$d(f; H_i) = \inf_{h \in H_i} \|f - h\|_{\mathcal{H}}, \tag{3.1}$$

where $\mathcal{H}$ is the RKHS under consideration and

$$H_i \triangleq \{h \in \mathcal{H} : |h(x_i) - y_i| \leq \epsilon_i\}, \tag{3.2}$$

is the hyperslab associated with the ith quantization interval, i.e., the set of all functions $h \in \mathcal{H}$ that pass through the ith quantization interval.

Regarding the term accounting for the model complexity, we use the squared Hilbert norm, i.e., $\|f\|_{\mathcal{H}}^2$.

As a result, the functional cost in Paper A consists of a composite of a non-smooth and a smooth function term. Motivated by this mathematical structure, our proposed Windowed Online Regularized cost Minimization (WORM) algorithm uses the stochastic proximal average functional gradient descent method due to its favorable convergence performance [65, 37].

3.3.1 Proposed Learning Scheme

Figure 3.2: Visual representation, at the nth step, of the learning scheme proposed in Paper B for a set of P WORM methods.

In Paper B, we propose a learning scheme that can be seen as a higher-level learner that iteratively chooses the lowest incurring "learning" cost[1] convex combination of function estimates provided by lower-level learners, e.g., several P WORM algorithms, over different reproducing kernels or kernel hyperparameters, running in parallel. The complexity of the combination step we devise can scale down to linearly with P on average, i.e., $\mathcal{O}(P)$. Fig. 3.2 illustrates this learning scheme.

3.4 Performance Analyses

Under an online setting, as described in Sec. 3.2, the incurred cost accumulated over time steps receives the name of *cumulative cost*. As an illustration, suppose that the sequence $\mathcal{S}$ contains N terms; then, the cumulative cost is computed as $\sum_{n=1}^{N} \mathcal{C}(f^{(n)}, \mathcal{S}_L^{(n)})$.

Notice that the cumulative cost is a measure of performance against overfitting. This is because every nth function estimate $f^{(n)}$ must be proposed before the nth sliding data window $\mathcal{S}_L^{(n)}$ becomes available to the learner, see Fig. 3.1.

Popular performance metrics are constructed from the cumulative cost [66]. For instance, the *dynamic regret* metric is defined as

$$\mathbf{Reg}(f^{(1)}, \ldots, f^{(N)}) \triangleq \sum_{n=1}^{N} \mathcal{C}(f^{(n)}; \mathcal{S}_L^{(n)}) - \mathcal{C}(f_*^{(n)}; \mathcal{S}_L^{(n)}), \qquad (3.3)$$

---

[1]Learning cost refers to a cost incurred while "learning" over a revealed sliding data window.
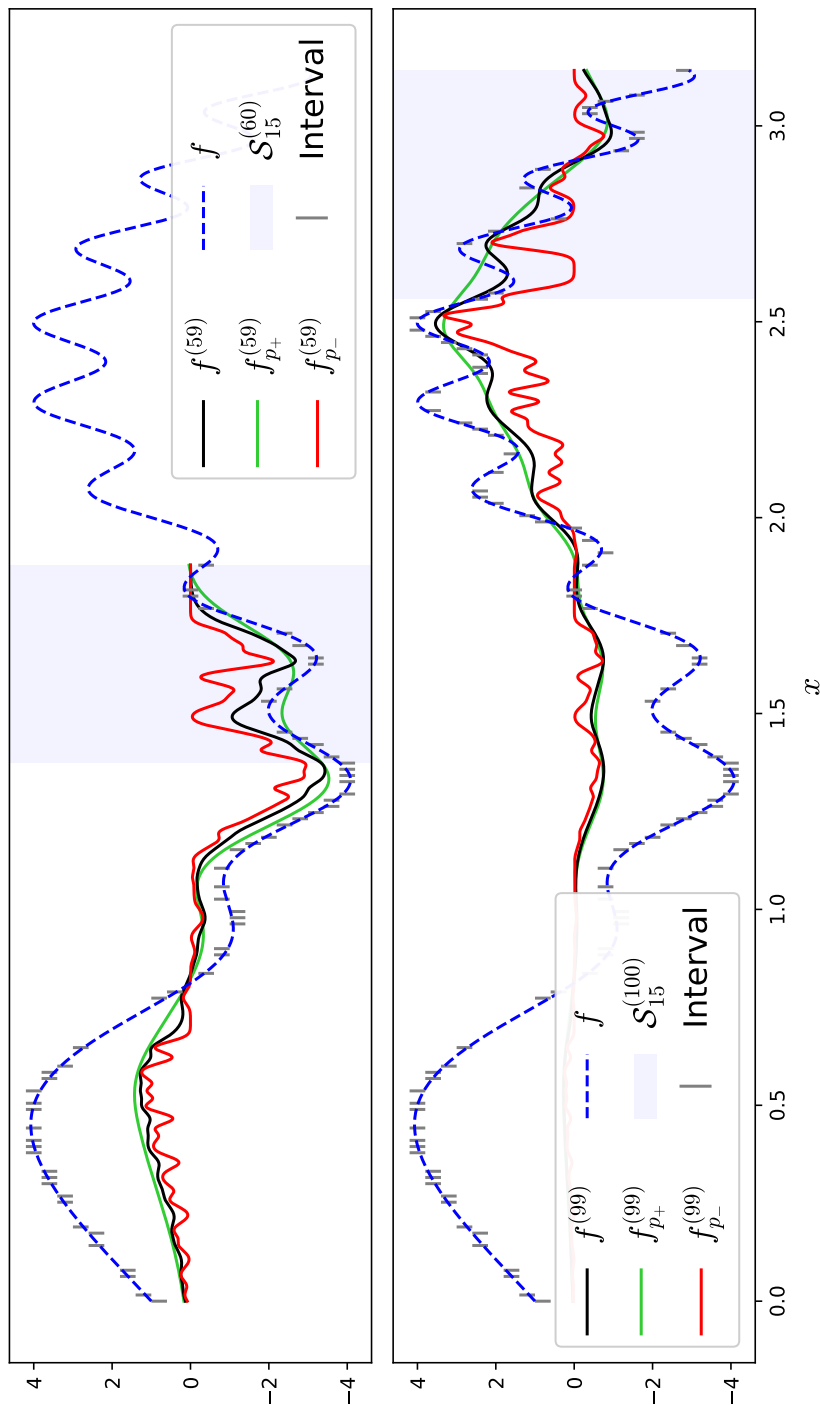
Figure 3.3: Signal tracking using learning scheme proposed in Paper B and $P = 10$ WORM algorithms with Gaussian kernels, i.e., $k(x, t) = \exp(-(x - t)^2/\sigma^2)$, with different widths σ spanning from 0.015 to 0.1, with a window of length $L = 15$, a regularization parameter of 0.2 and a learning rate of 2. The best and the worst individual WORM algorithms, in terms of the cumulative cost up to time step n (as shown in Fig. 3.4), are denoted by the indices p_+ and p_-, respectively. The data samples have been randomly sampled from $f(x) = \sin(2\pi x^2) + \cos(\pi x) + 3\sin(\pi x)$ and uniformly quantized with a half stepsize $\epsilon = 0.2$.

Figure 3.4: Cummulative cost up to time step n, i.e., $\sum_{i=1}^{n} \mathcal{C}(f^{(n)}; \mathcal{S}_L^{(n)})$, for the functional cost proposed in Paper A, and introduced in Sec. 3.3, for the $P = 10$ WORM algorithms and proposed learning scheme discussed in Fig. 3.3. As we can see, the proposed learning scheme in Paper B outperforms the best of the individual WORM algorithms used eventually.

where $f_*^{(n)} = \arg\min_f \mathcal{C}(f; \mathcal{S}_L^{(n)})$, and it indicates how well the sequence of function estimates $\{f^{(n)}\}_{n=1}^{N}$ matches the sequence of optimal decisions $\{f_*^{(n)}\}_{n=1}^{N}$.

In Paper A, we derive a theoretical upper bound for the dynamic regret incurred by the WORM algorithm.

On the other hand, Paper B shows experimentally how the proposed learning scheme outperforms the best combination of some single-kernel methods in terms of the regularized least squares cumulative cost.

For an example explicitly combining our proposed WORM algorithm and our proposed learning scheme, see Fig. 3.3 and Fig. 3.4.

3.5 Contributions

This section summarizes the main contributions of Paper A ($\heartsuit$) and Paper B ($\clubsuit$).

$\heartsuit$ We design a novel online algorithm called WORM to minimize a non-differentiable windowed cost assisted by the proximal average functional gradient descent method. We consider this algorithm for the problem of regression-based tracking of quantized signals.

$\heartsuit$ We provide tracking performance guarantees for WORM through a dynamic regret theoretical upper bound and experimental performance analyses using synthetic data.

$\clubsuit$ We present a multi-kernel learning scheme that combines online single kernel-based methods, outperforming the best of them individually in terms of the

cumulative regularized least squares cost metric, with a comparable computational load per computing unit. This corroborates the ability of the proposed scheme to effectively accommodate a larger function space (from which to draw function estimates) of multi-kernel methods while keeping the lower computational complexity of online single-kernel-based methods.

Chapter 4

Zero-delay Spline Interpolation Policies

This chapter summarizes Paper C and Paper D and delves into the common aspects between them for the task of zero-delay smoothing spline interpolation from streaming time-series data. The background theory relevant to this chapter can be found in Sec. 2.1, 2.4, 2.5 and 2.6.

4.1 Motivation

In the previous Ch. 3, we address the task of tracking signals from univariate time series of quantization intervals. However, some tasks may impose more demanding requisites, such as zero-delay response and smoothness in the reconstruction; for instance, online trajectory planning [67, 68], real-time control systems [69, 70], and high-speed digital to analog conversion [71], among others.

Unlike RKHS-based online methods, online interpolation methods can be suitable candidates for smooth signal reconstruction under zero-delay response and low model complexity requirements. This is because interpolation methods allow shaping piecewise-modeled signal estimates by sequentially assembling new function pieces. Piecewise polynomial functions, or splines, are arguably the most widely used ones [72, 73, 74, 75], presumably because of their model simplicity and approximation capabilities over functions of arbitrary complexity.

Based on this, Paper C presents a novel approach to zero-delay smoothing spline interpolation underpinned by a sequential decision-making framework [76]. This approach allows us to model the impact of each decision (i.e., each interpolated function piece) on future interpolations and, therefore, on the cumulative average cost, in this case, a roughness-based metric, resulting in a more accurate reconstruction. Then, an interpolation technique is proposed, assisted by an RNN that exploits the temporal dependencies between streamed data points to minimize the accumulated cost on average.

On the other hand, the method proposed in Paper C is unsuitable for quantized data or multivariate time series; these limitations are addressed in Paper D.

As a result, the proposed method in Paper D achieves one of the main goals of this thesis, namely, the zero-delay smooth signal reconstruction from streamed multivariate time series of quantization intervals.

The rest of this section presents the state-of-the-art on the zero-delay smooth

(a) Signal estimates and reconstruction proposed by NORMA. The reconstructed signal has inevitable discontinuities even though each of the successive signal estimates is continuous (and infinitely smooth).

(b) Every successive signal estimate, in this case, the Tth signal estimate, is constructed as a kernel expansion of overlapping terms. Each one of the expansion terms is plotted with a continuous colored line.

Figure 4.1: Example of the limitation of an online RKHS-based method, in this case, the naive online R_{reg} (regularized risk) minimization algorithm (NORMA) with a Gaussian kernel, for guaranteeing the smoothness of the reconstruction.
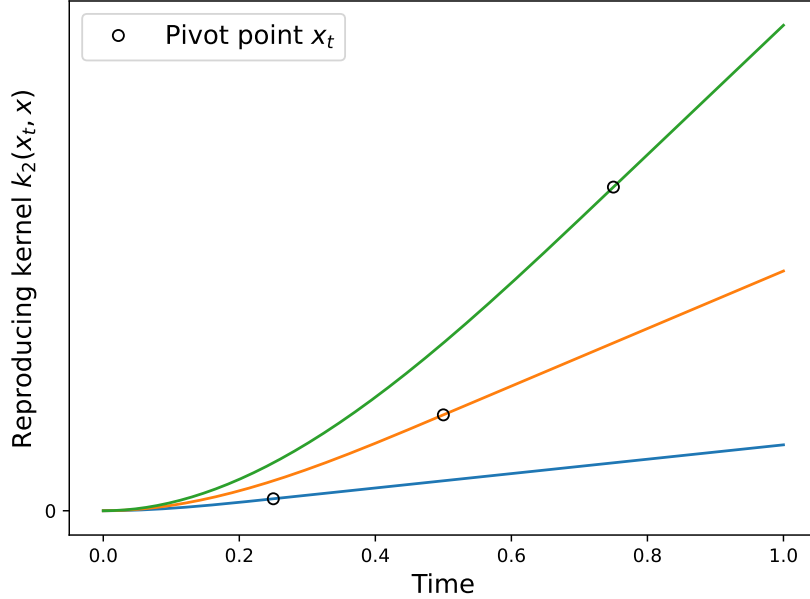
Figure 4.2: The reproducing kernel shown in (4.4) "centered" at different pivot points x_t.

signal reconstruction from streaming data points problem and further describes the limitations of RKHS-based methods for such tasks.

4.1.1 State-of-the-art

Some works have explored how to sequentially interpolate streaming data under real-time requirements using splines [25, 26]. However, the online methods developed in these works involve a sliding data window, which introduces a delay. This is because a delayed response allows them to successively correct the signal estimates as long as they are updated within the delay limits.

To the best of our knowledge, the myopic approach introduced in Ch. 1 is the only spline interpolation method that operates under a zero-delay response. However, this approach completely disregards any information that the past samples contain about the forthcoming data samples. As a result, signal reconstruction tasks under zero-delay response and smoothness requirements may be adversely affected since any of the current reconstructed signal portions can confine the future portions through the continuity constraints, thus affecting their accuracy through suboptimal results, possibly resulting in instabilities. This drawback is further explained in Paper C.

4.1.2 Limitation of RKHS-based Approaches

Online RKHS-based methods aim at yielding a sequence of signal estimates with regret guarantees [77]. To this end, they initially propose a signal estimate that is updated, possibly globally (along the whole timespan of the received data samples), as new data samples arrive. Their goal is to refine the signal estimate rather than to reconstruct new signal pieces smoothly. Therefore, neither smoothness nor continuity of the sequentially reconstructed signal is guaranteed.

As shown in Fig. 4.1, any spline-based signal estimate constructed as a reproducing kernel (basis) expansion suffers from overlapping between kernels. That is, if the tth expansion term $\alpha_t k(x_t, x)$ is non-zero, then, when added (or updated), it will affect the output of the overall expansion (signal estimate) and, hence, does not guarantee the continuity of the reconstruction. This happens for any valid kernel, and Fig. 4.1 illustrates it for the case of a Gaussian reproducing kernel.

As a closely related example, the smoothing spline interpolation problem introduced in Sec. 2.1.5.1 can be equivalently formulated using an RKHS framework [24]. That is,

$$\underset{f \in \mathcal{H}_\rho}{\text{minimize}} \ \sum_{t=1}^{T} \left(f(x_t) - y_t \right)^2 + \eta \|f\|_{\mathcal{H}_\rho}^2, \tag{4.1}$$

where (x_t, y_t) is the tth data point of a given sequence of T terms, $\eta \geq 0$ is the regularization hyperparameter, and $\mathcal{H}_\rho$ is the RKHS induced by the inner product

$$\langle f, h \rangle_{\mathcal{H}_\rho} = \int_0^1 D_x^\rho f(x) \, D_x^\rho h(x) \, dx, \tag{4.2}$$

with corresponding RKHS norm $\|f\|_{\mathcal{H}_\rho} = \sqrt{\langle f, f \rangle_{\mathcal{H}_\rho}}$. It is well known, thanks to the Representer Theorem [60], that the solution to (4.1) is of the form

$$f_T(x) = \sum_{t=1}^{T} \alpha_t k_\rho(x_t, x), \tag{4.3}$$

where $\alpha_t \in \mathbb{R}$, for $t = 1, \ldots, T$, are the kernel expansion coefficients and $k_\rho : [0, 1]^2 \to \mathbb{R}$ is the reproducing kernel associated with $\mathcal{H}_\rho$. For convenience, let us choose $\rho = 2$. Then, the reproducing kernel associated to $\mathcal{H}_2$ [24, 78] is given by

$$k_2(x_t, x) = \frac{1}{2} \max\{x_t, x\} \min^2\{x_t, x\} - \frac{1}{6} \min^3\{x_t, x\}. \tag{4.4}$$

The "shape" of the reproducing kernel in (4.4) is shown in Fig. 4.2. From the figure, one can notice that there is an overlap among different expansion terms over the whole time domain with non-zero values.

4.2   Problem Formulation

Paper C shows that interpolating a sequence of data points under a zero-delay response and smoothness requirements can be done with a given decision strategy or policy, where the resulting action takes the form of spline coefficients. From this observation, we pose the problem of finding the policy that achieves the lowest incurred cost on average from any time series of data points acquired from a given signal source, formulated as a policy search by cost minimization [45].

More specifically, the cost mentioned above consists of a weighted sum between the sum of squared residuals and a derivative-based measure of roughness, as in the smoothing spline interpolation problem described in (2.6), but from an action-state representation. All the remaining details are fully explained in Paper C.

Paper D formulates a similar problem to Paper C but is oriented to multivariate time series of uncertainty intervals and subject to consistency in the signal reconstruction. In this way, an additional component of the problem formulation in Paper D is how to properly extend the concept of consistency to an online setting.

Figure 4.3: Action-state interpretation of the smoothing spline interpolation problem. The tth state $\boldsymbol{s}_t$ contains the necessary information to continue the interpolation task from time stamp x_{t-1}, that is, $\boldsymbol{s}_t = [x_{t-1}, x_t, y_t, \boldsymbol{e}_{t-1}^\top]^\top$. The tth action $\boldsymbol{a}_t$ are the coefficients of the tth function section g_t, and in our case the are computed through a parametric policy as $\boldsymbol{a}_t = \boldsymbol{\mu}_{\boldsymbol{\theta}}(\boldsymbol{s}_t)$.

In the ensuing sections, we use the content of Paper C as the common thread. Nevertheless, we include and expand the main contributions in Paper D when convenient.

4.3   Proposed Solution

The problem of finding the most suitable policy inside the space of candidate policies can be overwhelmingly complex due to the diversity and large number of available policies within [45]. Policy approximation techniques can help to reduce the policy space effectively. These techniques tend to work best (in the sense of providing an adequate policy) when the problem has a clear structure that can be accommodated into the policy. In our case, we aim to incorporate the temporal dependencies across the observations and the smoothness requirement.

Specifically, in Paper C, we propose a family of parametrized policies through cost parametrization. The temporal dependencies are captured thanks to an RNN incorporated into the policy design, and the smoothness requirement is imposed as a constraint in a differentiable convex optimization layer (DCOL) [14].

Mathematically, the proposed policy takes the following form:

$$\boldsymbol{\mu}_{\boldsymbol{\theta}}(\boldsymbol{s}_t) = \arg \min_{\boldsymbol{a} \in \mathcal{A}(\boldsymbol{s}_t) \subseteq \mathbb{R}^{d+1}} \left\{ \kappa_\eta(\boldsymbol{s}_t, \boldsymbol{a}) + \boldsymbol{J}_{\boldsymbol{\theta}}(\boldsymbol{s}_t, \boldsymbol{a}; \boldsymbol{h}_t) \right\}, \qquad (4.5)$$

where $\boldsymbol{s}_t$ represents the tth state representation of the interpolation task; in this case, $\boldsymbol{s}_t = [x_{t-1}, x_t, y_t, \boldsymbol{e}_{t-1}^\top]^\top$, where each vector $\boldsymbol{e}_{t-1} \in \mathbb{R}^{\varphi+1}$ contains $\varphi + 1$ continuity constraint values. Accordingly, $\mathcal{A}(\boldsymbol{s}_t)$ is the set of spline coefficients $\boldsymbol{a} \in \mathbb{R}^{d+1}$ that satisfy the continuity constraints in $\boldsymbol{e}_{t-1}$ at time step t. A more visual explanation is provided in Fig. 4.3.

On the other hand, κ_η is the action-state representation of the smoothing spline interpolation cost in (2.6), and computing $\boldsymbol{J}_{\boldsymbol{\theta}}$ involves evaluating an RNN with latent state at time step t denoted by $\boldsymbol{h}_t \in \mathbb{R}^H$. Moreover, κ_η, $\boldsymbol{J}_{\boldsymbol{\theta}}$ and $\mathcal{A}(\cdot)$ are chosen such that evaluating the policy in (4.5) consists in solving a convex optimization problem that admits a closed-form solution that can be differentiated.

In this way, we can understand the evaluation of $\boldsymbol{\mu}_{\boldsymbol{\theta}}$ as a forward pass of a DCOL.

Figure 4.4: Snapshot at time x_{25} of a zero-delay smoothing spline interpolation using the policy proposed in Paper C. The signal reconstruction is done using only the center of the quantization intervals. Consequently, the policy is unaware of the quantization step sizes and thus cannot enforce consistency. The image is borrowed from Paper D for the sake of illustration.

4.3.1 Enforcing Consistency

The family of policies proposed in Paper C is only suitable for data points and not for any other uncertainty interval, including quantized data. As a result, the information contained within the quantization intervals is ignored, as shown in Fig. 4.4.

On the other hand, and as discussed in Sec. 2.6, consistent signal reconstruction methods [27, 28] exploit such information and have proved profitable (in the sense of implementation cost) in practice [57].

Paper D successfully extends the family of policies proposed in Paper C to deal with quantized data while ensuring consistency. This is done by formalizing the concept of consistent signal reconstruction from streaming data under zero-delay response requirements. In this way, our formulation generalizes the concept of consistency beyond offline settings rather than establishing a new one.

4.3.2 Multivariate Case

The family of candidate policies proposed in Paper C successfully exploits temporal relationships within univariate time series. However, they ignore any possible spatial relationships among multivariate time series. Paper D fixes this limitation by adequately extending the family of policies proposed in Paper C again. The overall idea is shown in Fig. 4.5.

4.4 Performance Analyses

Both of the proposed methods in Paper C and Paper D are compared with respect to the state-of-the-art, i.e., the myopic approach, and the best possible solution (obtained under an offline setting in practice) in terms of a roughness-based cost. The results show a considerable improvement with respect to the state-of-the-art. On the other hand, the experimental results in paper D, describing the benefits

Figure 4.5: Snapshot at time x_t of a multivariate time series of quantization intervals (aligned in time across series) being reconstructed. The policy that achieves a zero-delay consistent signal reconstruction while accounting for the spatiotemporal dependencies among data samples is represented visually as a microprocessor. The upper indices (n) on the signal estimates refer to the series they reconstruct. The green shaded areas indicate the currently reconstructed portion of the multivariate signal, while the red ones represent the future. The image is borrowed from paper D for illustration purposes.

of zero-delay consistently reconstructed signals, agree with the expected behavior known for offline settings.

Specifically, we experimentally find that the MSE of a consistent zero-delay reconstruction f_{cons} decreases roughly at the same rate regardless of whether we reduce the quantization step size Δ or increase the oversampling ratio R (possible thanks to the finite rate of innovation of spline-based signals), as

$$\text{MSE}(f_{\text{cons}}) \propto \frac{\Delta^{0.99 \pm 0.09}}{R^{0.812 \pm 0.008}}. \tag{4.6}$$

Moreover, the error-rate decay in (4.6) roughly doubles (in logarithmic scale) the one achieved by an inconsistent zero-delay reconstruction f_{inc}, such as the one proposed in Paper C and illustrated in Fig. 4.4, with respect to the oversampling ratio R. Expressly,

$$\text{MSE}(f_{\text{inc}}) \propto \frac{1}{R^{0.435 \pm 0.002}}. \tag{4.7}$$

Finally, we observe that the proposed algorithms could not find stable policy configurations for certain polynomial orders and degrees of smoothness of the spline. However, these instabilities also manifest in the myopic configuration (existing before our method), hinting at a more fundamental problem that requires further research.

4.5 Contributions

This section summarizes the main contributions of Paper C ($\diamondsuit$) and Paper D ($\spadesuit$).

$\diamondsuit$ We rigorously formulate the problem of smoothing spline interpolation from streaming data, where each spline section has to be determined as soon as a data sample is available and without having access to subsequent data (zero-delay response requirement). Due to its nature, it is formulated as a sequential decision-making problem.

◇ Unlike previously proposed zero-delay methods (e.g., the myopic approach, which is not trainable), our method trains a policy to minimize the smoothing interpolation problem cost metric on average. To capture the temporal, possibly long-term, dependencies between the streamed data samples and exploit them to reduce the cost metric on average, we incorporate an RNN able to capture the temporal signal dynamics.

◇ The proposed policy guarantees that the reconstructed signal is smooth. This is achieved by adding a DCOL at the output of the RNN and imposing a set of continuity constraints at each interpolation step. In addition, such a layer admits a closed-form evaluation, resulting in improved computational efficiency with respect to off-the-shelf DCOL libraries.

◇ We present extensive experimental results that validate the zero-delay smoothing spline interpolation approach over synthetic and real data. Additionally, we show how such an approach outperforms the state-of-the-art (namely, the myopic approach) zero-delay methods in terms of the smoothing spline interpolation problem cost metric on average.

♠ We formalize the concept of consistent signal reconstruction from streaming data under zero-delay response requirements. The resulting formulation generalizes the concept of consistency beyond offline settings.

♠ To the best of our knowledge, Paper D devises the first trainable method for zero-delay signal reconstruction from multivariate time series of quantization intervals, which can also enforce consistency.

♠ We show experimentally that the reconstruction error incurred by the zero-delay consistent signal reconstruction proposed method decays at the same rate by decreasing the quantization step size or increasing the oversampling ratio. Moreover, the error-rate decay slope with respect to the oversampling ratio doubles (in logarithmic scale) the one obtained with a similar zero-delay smooth signal reconstruction method that does not enforce consistency.

Chapter 5

Concluding Remarks

5.1 Conclusion

In this Ph.D. thesis, we address the problem of signal reconstruction from streamed multivariate time series of quantization intervals subject to zero-delay and smoothness requirements.

Although most existing algorithms can manage the problem requirements of estimating smooth signals or delivering a zero-delay response separately, addressing them together becomes challenging. This challenge is the core concept of this dissertation, and it has been solved incrementally through Papers A to D. A summary of the requirements satisfied by the proposed methods is provided in Table 5.1.

Table 5.1: Scope of the methods proposed in the compendium of papers.

	Zero-delay	Quantization	Smooth	Consistency	Multivariate
Paper A, B[1]	✔ [2]	✔	✘	✘	✘
Paper C	✔	✘	✔	✘	✘
Paper D	✔	✔	✔	✔	✔

5.2 Future Work

The work discussed in this dissertation can serve as the foundation and motivation for some noteworthy research directions.

5.2.1 Fundamental Research on Zero-delay Smoothing Spline Instability Issues

We experimentally observe that the myopic policy described in Paper C is not stable for a degree of smoothness greater than 2, that is, smooth above the second derivative regardless of the spline order. We also observe instability under the myopic policy

---

[1]Considering that the learning scheme is applied to the algorithm proposed in Paper A.

[2]The sliding window can be shrunk to the length of one data sample.

(a) Batch solution with a spline of order 3 and degree of smoothness 2. No observed instability issues.

(b) Myopic policy with a spline of order 3 and degree of smoothness 2. Instability is observed.

(c) Myopic policy with a spline of order 4 and degree of smoothness 2. No observed instability issues.

Figure 5.1: Instability issues suffered by the myopic policy against a stable configuration and the batch solution, for a degree of smoothness 2.

for a spline signal estimate of order 3 and degree of smoothness 2. This is illustrated in Fig. 5.1.

Our proposed methods in Paper C and Paper D are unstable for the same spline configurations since they implicitly use the myopic policy as a guided starting point.

This motivates further theoretical instability studies, the design of alternative policy architectures, or the exploration of low-delay approaches by relaxing the zero-delay constraint.

5.2.2 Real-time Safe Trajectory Planning

Suppose a trajectory planning problem is subject to passing through safety boxes of any dimension, e.g., spatial dimensions, angles, etc., at a given set of time instants. Then, the safety boxes can be seen as a set of temporarily aligned intervals, one for each dimension separately. In other words, a trajectory planning problem subject to safety boxes at certain time instants can be formulated as a signal reconstruction from a multivariate time series of uncertainty intervals that can overlap, as shown in Fig. 5.2.

For this reason, real-time safe trajectory planning tasks can benefit from the method proposed in Paper D.

This is because, in addition to passing through all safety boxes, even when the next safety box is revealed only one time step ahead, the proposed method can ensure smoothness, e.g., a continuous acceleration, and low-model complexity, e.g., damped oscillations.

5.2.3 Compression for Reconstruction and Inference

Piecewise polynomial approximation techniques can be used for time-series compression [79]. Normally, they divide a time series into several segments of variable length and find the best polynomial approximating them (according to some metric). One of their main advantages is that, despite the compression being lossy, a maximum deviation from the original (uncompressed) data can be fixed in advance to enforce a certain reconstruction accuracy.

We believe our proposed technique in Paper D can be readily adapted to compress streaming time-series data under a maximum reconstruction deviation and real-time constraints once the stability issue is fixed. The compressed space would consist of a set of interdependent spline coefficients due to the smoothness constraints (fixing some degrees of freedom); therefore, only a few coefficients must be stored.

For example, suppose we want to use a quadratic spline with tth function section given by $g_t(x) = a_{t,1} + a_{t,2}(x - x_{t-1}) + a_{t,3}(x - x_{t-1})^2$ and continuity up to the first derivative, i.e., the following recursive relations

$$g_t(x_t) = g_{t+1}(x_t) \rightarrow a_{t+1,1} = a_{t,1} + a_{t,2}u_t + a_{t,3}u_t^2, \tag{5.1a}$$

$$g_t'(x_t) = g_{t+1}'(x_t) \rightarrow a_{t+1,2} = a_{t,2} + 2a_{t,3}u_t, \tag{5.1b}$$

are satisfied, where $u_t \triangleq x_t - x_{t-1}$. Then the coefficients $\{a_{t,1}\}_t$ and $\{a_{t,2}\}_t$ can be obtained from an initial tuple $\{a_{0,1}, a_{0,2}, a_{0,3}\}$, and the sequences $\{x_t\}_t$ and $\{a_{t,3}\}_t$ through the recursive relations in (5.1).

From here, notice that if most of the stored coefficients $\{a_{t,3}\}_t$ are zero-valued, we can achieve competitive compression ratios. This can be done, for instance, by

(a) Two-dimensional trajectory. The blue squared boxes represent safety regions that must be visited. The red rounded boxes represent dangerous areas that should be avoided.

(b) Decomposition of the 2-dimensional trajectory from Fig. 5.2a in a multivariate time series. The safety boxes become uncertainty intervals that can overlap.

Figure 5.2: Two-dimensional representation of a zero-delay safe trajectory planning task. The method proposed in Paper D can ensure that the trajectory passes through the safety regions at a given instant and promotes low-model complexity and smoothness.

promoting sparsity through a cost term.

In this case, the reconstruction step would recover all coefficients via (5.1) and then evaluate the spline model. Moreover, it would not only recover the original data but an estimate of intermediate points with a time resolution as high as desired.

On the other hand, some tasks may benefit from working directly in the compressed state (without the need for the reconstruction step). For those tasks, we might be interested in promoting a certain task-dependent performance metric and not only sparsity. We refer to this type of compression as compression for inference, and Fig. 5.3 illustrates how the architecture of our model can be adapted for such purposes.
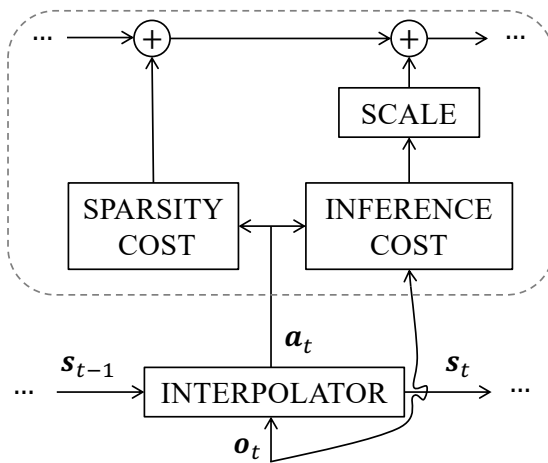
Figure 5.3: Architecture prototype of the proposed spline-based compressor for inference. The INTERPOLATOR cell may contain trainable parameters, and it ensures a spline reconstruction under a given maximum deviation. The SCALE cell acts as a user-defined trade-off hyperparameter. The notations s_t, a_t, and o_t refer to the tth state, action, and observation, respectively. The components in the gray dashed rounded box are used only during the training phase.

Appendix A

Paper A

Title:	Tracking of quantized signals based on online kernel regression
Authors:	**Emilio Ruiz-Moreno** and Baltasar Beferull-Lozano
Conference:	IEEE International Workshop on Machine Learning for Signal Processing

Tracking of quantized signals based on online kernel regression

Emilio Ruiz-Moreno and Baltasar Beferull-Lozano

Abstract — **Kernel-based approaches have achieved noticeable success as non-parametric regression methods under the framework of stochastic optimization. However, most of the kernel-based methods in the literature are not suitable to track sequentially streamed quantized data samples from dynamic environments. This shortcoming occurs mainly for two reasons: first, their poor versatility in tracking variables that may change unpredictably over time, primarily because of their lack of flexibility when choosing a functional cost that best suits the associated regression problem; second, their indifference to the smoothness of the underlying physical signal generating those samples. This work introduces a novel algorithm constituted by an online regression problem that accounts for these two drawbacks and a stochastic proximal method that exploits its structure. In addition, we provide tracking guarantees by analyzing the dynamic regret of our algorithm. Finally, we present some experimental results that support our theoretical analysis and show that our algorithm has a favorable performance compared to the state-of-the-art.**

A.1   Introduction

Regression problems are some of the most important problems due to their numerous applications and relevance in a wide range of fields. In practice, regression problems are usually formulated as convex optimization problems with strongly convex objectives over feasible convex sets. Besides being one of the most benign settings, this formulation includes significant instances of interest, such as those arising in regularized regression [19], for example, to reduce the complexity of the reconstruction by promoting smoothness. Because of this, such strongly convex objectives are commonly set as the sum of a convex loss, which reflects how far the solution lies from the data samples, and a strongly convex regularizer, which controls the complexity of the solution. On the other hand, most real-world scenarios where regression techniques may be useful occur under dynamic environments. This fact motivates the design of online methods. They allow tracking over time the underlying target signals in a recursive manner with reduced memory and computational needs.

In particular, this paper focuses on sequentially streamed quantized signals. When the underlying physical process generating the signal data samples is unknown, as usual in practice, instead of blindly selecting a certain ad-hoc parametric regression model, the target signal can be estimated from the data samples. This can be done by means of non-parametric regression methods at the expense of a certain memory and computational cost that can be controlled.

Under the mathematical framework of Reproducing Kernel Hilbert Spaces (RKHSs) and thanks to the Representer Theorem [80], such a non-parametric estimate can be constructed from a pre-selected reproducing kernel with a complexity that grows

linearly with the number of data samples. Regression with kernels and its online variants have been widely studied in the literature [60, 22]. Their main strength is that they are able to find non-linear patterns at a reasonable computational cost. The Naive Online regularized Risk Minimization Algorithm (NORMA) [12] is arguably the most representative algorithm from the stochastic approximation kernel-based perspective. In its standard form, it concentrates all the novelty in the new expansion coefficient of the signal estimate. However, intuitively, it seems reasonable to distribute the novelty among several expansion coefficients that contribute to the signal estimate instead. In this way, the novelty and correction of previous estimate errors are integrated, more ergonomically, in the signal estimate.

To the best of our knowledge, most of the existing literature has focused on controlling the signal estimate complexity rather than focusing on strategies to control the error in the estimates. Examples of research works controlling complexity are truncation [12] and model-order control via dictionary refining [81], among others [82, 83]. Only some works have studied reducing the signal estimate errors by means of a sliding window scheme [84, 23, 85]. However, in [84], the *selection criterion* to choose among all possible function estimates is least squares making it unsuitable for more general settings, such as incorporating quantization intervals instead of signal values. Similarly, in [23], even though its selection criterion allows certain freedom, regularization is not encouraged and therefore, the smoothness of the underlying physical signal is not fully promoted. Lastly, in [85], the selection criterion is constructed as a regularized augmentation of instantaneous loss-data pairs. As a result, it naturally extends NORMA in a sliding window scheme. Nonetheless, in this work, we present a novel algorithm constituted by a robust selection criterion alongside a conveniently engineered optimization method that outperforms all these algorithms for the task of regression-based tracking of quantized signals.

The paper is structured as follows: Sec. A.2 presents the *windowed cost* and formulates the problem from a learner-adversary perspective. Then, in Sec. A.3, we provide our main contribution: a novel method to minimize the windowed cost via proximal average functional gradient descent. The resulting approach, a novel algorithm called WORM, is used for the practical use case of regression-based tracking of quantized signals. Next, in Sec. A.4, we provide its tracking guarantees through a dynamic regret analysis. Finally, in Sec. A.5, we analyze the experimental performance of our algorithm using synthetic data, and Sec. A.6 concludes the paper.

A.2 Problem formulation

Given a possibly endless sequence $\mathcal{S} = \{(x_i, y_i)\}_{i=1}^N$ of data samples $(x_i, y_i) \in \mathcal{X} \times \mathcal{Y} \subseteq \mathbb{R}^2$, with strictly increasing and non necessarily equispaced timestamps $\{x_i\}_{i=1}^N$, consider an online learning setting where the data samples become available sequentially. Under this scenario, online regression problems can be understood as an interplay between an algorithm (or learner) and an adversary (or environment) [86, 87]. At each step $n \in \mathbb{N}$, an algorithm proposes a function estimate, which we denote as f_n, from an RKHS $\mathcal{H}$. In response, an adversary selects a functional cost $\mathcal{C}_n : \mathcal{H} \to \mathbb{R}$ and penalizes the proposed function estimate with the incurred cost $\mathcal{C}_n(f_n)$. Then, the adversary reveals relevant information about the form of $\mathcal{C}_n$ that is used by the algorithm at the next step.

Unlike most of the previous work, which uses an instantaneous functional cost,

i.e., a functional evaluated over one data sample, we formulate the *hypothesis* that a concurrent functional cost, i.e., a functional that considers up to $L \in \mathbb{N}$ data samples simultaneously, may lead to better performance at the expense of a higher but bounded computational cost.

In order to test our hypothesis, we first consider a proper convex *instantaneous loss* $\ell_n : \mathcal{H} \to \mathbb{R} \cup \{\infty\}$ given by

$$\ell_n(f) \triangleq \ell(\langle f, k(x_n, \cdot) \rangle_{\mathcal{H}}, y_n) = \ell(f(x_n), y_n), \tag{A.1}$$

where $k(x_n, \cdot)$ is the reproducing kernel associated with the RKHS $\mathcal{H}$ centered at x_n. Notice that the equality in (A.1) holds thanks to the reproducing property [12]. Consequently, we define the so-called windowed cost as a composite of a weighted arithmetic mean of an instantaneous loss as in (A.1), computed over L consecutive data samples and the squared Hilbert norm associated to $\mathcal{H}$ as the regularizer, i.e.,

$$\mathcal{C}_n(f) \triangleq \mathcal{L}_n(f) + \frac{\lambda}{2} \|f\|_{\mathcal{H}}^2, \tag{A.2}$$

with regularization parameter $\lambda > 0$ and where the *windowed loss* $\mathcal{L}_n : \mathcal{H} \to \mathbb{R} \cup \{\infty\}$ is given by

$$\mathcal{L}_n(f) = \sum_{i=l_n}^{n} \omega_i^{(n)} \ell_i(f), \tag{A.3}$$

where $l_n = \max\{1, n - L + 1\}$ and $\sum_{i=l_n}^{n} \omega_i^{(n)} = 1$ with $\omega_i^{(n)} \geq 0$. Finally, the RKHS $\mathcal{H}$, the instantaneous loss ℓ, the regularizer parameter λ and the tuning routine of the convex weights $\{\omega_i^{(n)}\}_{i=l_n}^{n}$ are specified by the user.

A.2.1 Performance analysis

The performance of an online algorithm can be measured by comparing the total cost incurred by the algorithm, given by $\sum_{n=1}^{N} \mathcal{C}_n(f_n)$, and the total corresponding cost incurred by a genie that knows all the costs in advance, that is, $\sum_{n=1}^{N} \mathcal{C}_n(f_n^*)$, where $f_n^* = \arg\min_{f \in \mathcal{H}} \mathcal{C}_n(f)$. Such a metric, referred to as *dynamic regret*, is defined as

$$\mathbf{Reg}_N \triangleq \sum_{n=1}^{N} \mathcal{C}_n(f_n) - \mathcal{C}_n(f_n^*). \tag{A.4}$$

The dynamic regret captures how well the sequence of function estimates $\{f_n\}_{n=1}^{N}$ matches the sequence of optimal decisions in environments that may change unpredictably over time. In general, obtaining a bound on the dynamic regret may not be possible [66]. However, under some mild assumptions on the sequence of functional costs, it is possible to derive worst-case bounds in terms of the cumulative variation of the optimal function estimates

$$\mathrm{C}_N = \sum_{n=2}^{N} \|f_n^* - f_{n-1}^*\|_{\mathcal{H}}. \tag{A.5}$$

In fact, some interesting bounds can be derived if we consider specific rates of variability [86], namely, from zero cumulative variation to a steady tracking error.

A.3   Proposed solution

The smoothness of a windowed cost, as in (A.2), depends on whether or not its instantaneous loss is smooth. Steepest descent methods have been traditionally used for differentiable problems. As for non-differentiable problems, they can be handled, in principle, via the subgradient method and its variants. However, when the cost consists of a composite of a smooth and a non-smooth term, proximal gradient descent methods are preferable because they provide faster convergence as compared to subgradient methods [37]. Therefore, if the instantaneous loss is a proximable[1] non-smooth functional, we propose to minimize the windowed cost via proximal average functional gradient descent due to its favorable convergence performance [88].

A.3.1   Stochastic proximal average functional gradient descent

Our proposed algorithm, the Windowed Online Regularized cost Minimization (WORM), makes use of the stochastic proximal average functional gradient descent. Encouraged by the windowed loss in (A.3), it exploits the concept of the so-called *proximal average functional*. Let us first to introduce some definitions to motivate our algorithm:

Given an RKHS $\mathcal{H}$, a closed proper convex functional $\ell : \mathcal{H} \rightarrow \mathbb{R} \cup \{\infty\}$ and a real parameter $\eta > 0$, the Moreau envelope of ℓ with smoothing parameter η, is defined as

$$\mathrm{M}_\ell^\eta(h) \triangleq \inf_{g \in \mathcal{H}} \left\{ \ell(g) + \frac{1}{2\eta} \|g - h\|_{\mathcal{H}}^2 \right\}, \tag{A.6}$$

for all $h \in \mathcal{H}$. The Moreau envelope is a smooth functional that is continuously differentiable (even if ℓ is not), and such that the set of minimizers of ℓ and M_ℓ^η are the same. Thus, the problems of minimizing ℓ and M_ℓ^η can be shown to be equivalent [37]. In addition, a derivative step with respect to the Moreau envelope corresponds to a proximal step with respect to the original function, i.e.,

$$\partial \mathrm{M}_\ell^\eta = \eta^{-1} \left(\mathrm{I} - \mathrm{prox}_\ell^\eta \right), \tag{A.7}$$

where $\mathrm{I} : \mathcal{H} \rightarrow \mathcal{H}$ is the identity operator and $\mathrm{prox}_\ell^\eta : \mathcal{H} \rightarrow \mathcal{H}$ is the proximal operator defined as

$$\mathrm{prox}_\ell^\eta(h) \triangleq \arg \min_{g \in \mathcal{H}} \left\{ \ell(g) + \frac{1}{2\eta} \|g - h\|_{\mathcal{H}}^2 \right\}, \tag{A.8}$$

for all $h \in \mathcal{H}$. Notice that since the objective in (A.8) is strongly convex, the proximal map is single-valued.

Next, we denote by $\mathcal{L}_n^\eta$ the so-called proximal average functional of the windowed loss in (A.3) at instant n with real parameter $\eta > 0$, as the unique closed proper convex functional such that

$$\mathrm{M}_{\mathcal{L}_n^\eta}^\eta = \sum_{i=l_n}^n \omega_i^{(n)} \mathrm{M}_{\ell_i}^\eta, \tag{A.9}$$

---

[1]Its proximal operator can be computed efficiently.

where $\ell_i \triangleq \ell(f(x_i), y_i)$ for all $f \in \mathcal{H}$. Even though it is possible to derive an explicit expression for the proximal average functional from its definition (definition 4.1, [65]), for the sake of clarity, and since only its existence is needed for the algorithm, we do not include its explicit form here.

At each iteration n, our algorithm executes the steps:

$$\bar{f}_n = f_n - \eta \partial_f \left. \frac{\lambda}{2} \|f\|_{\mathcal{H}}^2 \right|_{f=f_n}, \tag{A.10a}$$

$$f_{n+1} = \mathrm{prox}_{\mathcal{L}_n^\eta}^\eta(\bar{f}_n), \tag{A.10b}$$

with $0 \le \eta < \lambda^{-1}$. The first algorithm step (A.10a) is equivalent to $\bar{f}_n = \rho f_n$ with $\rho \triangleq (1 - \eta\lambda) \in [0, 1)$. The proximal operator $\mathrm{prox}_{\mathcal{L}_n^\eta}^\eta : \mathcal{H} \to \mathcal{H}$, can be readily computed by differentiating both sides of the definition in (A.9) while applying the Moreau envelope property given by (A.7), getting

$$\mathrm{prox}_{\mathcal{L}_n^\eta}^\eta = \sum_{i=l_n}^{n} \omega_i^{(n)} \mathrm{prox}_{\ell_i}^\eta. \tag{A.11}$$

The remaining steps depend on the choice of the instantaneous loss. In particular, since we are interested in quantized signals, an adequate functional instantaneous loss must not penalize the function estimates that pass through the intervals. We develop further this reasoning in Sec. A.3.2.

A.3.2 Application to online regression of quantized signals

Consider the sequence of quantization intervals, where the ith quantization interval is given by its timestamp $x_i \in \mathbb{R}$, center $y_i \in \mathbb{R}$ and quantization half step-size $\epsilon \in \mathbb{R}_+$. Subsequently, we can construct its associated sequence of closed hyperslabs, each one of them defined as the convex set

$$H_i \triangleq \{f \in \mathcal{H} : |f(x_i) - y_i| \le \epsilon\}, \tag{A.12}$$

that contains all the functions in $\mathcal{H}$ passing through the ith quantization interval, and use the metric distance functional to the ith hyperslab

$$d_i(f) \triangleq \inf_{h \in H_i} \|f - h\|_{\mathcal{H}} = \|f - P_{H_i}(f)\|_{\mathcal{H}}, \tag{A.13}$$

as an instantaneous loss to discern between all possible function candidates $f \in \mathcal{H}$. The mapping $P_{H_i} : \mathcal{H} \to H_i$ stands for the metric projection onto H_i and can be expressed as $P_{H_i}(f) = f - \beta_i k(x_i, \cdot)$ (example 38, [23]), where every coefficient β_i is computed as

$$\beta_i = \begin{cases} \frac{f(x_i) - y_i - \epsilon}{k(x_i, x_i)}, & \text{if } f(x_i) > y_i + \epsilon, \\ 0, & \text{if } |f(x_i) - y_i| \le \epsilon, \\ \frac{f(x_i) - y_i + \epsilon}{k(x_i, x_i)}, & \text{if } f(x_i) < y_i - \epsilon. \end{cases} \tag{A.14}$$

For practical purposes, the relation in (A.13) can be equivalently computed as $d_i(f) = |\beta_i| k(x_i, x_i)^{\frac{1}{2}}$.

Regarding the tuning routine of the convex weights in (A.3), recall that if the set $\{i \in [l_n, n] : \bar{f}_n \notin H_i\} = \emptyset$, any choice of convex weights incurs zero windowed loss. If not, each convex weight is tuned as

$$\omega_i^{(n)} = \frac{d_i(\bar{f}_n)^m}{\sum_{j=l_n}^n d_j(\bar{f}_n)^m} = \frac{|\bar{\beta}_i^{(n)}|^m k(x_i, x_i)^{\frac{m}{2}}}{\sum_{j=l_n}^n |\bar{\beta}_j^{(n)}|^m k(x_j, x_j)^{\frac{m}{2}}}, \qquad (A.15)$$

where $\bar{\beta}_i^{(n)}$ comes from the metric projection map $P_{H_i}(\bar{f}_n)$ and m is a user predefined non-negative real power. In this way, if $m = 0$ the convex weights are all equal. On the other hand, when m tends to infinity, only the weight associated to the largest distance is considered. Thus, the power m allows, with a range of flexibility, to weigh more those windowed loss terms in which the intermediate update $\bar{f}_n$ incurs a larger loss.

Accordingly, from the proximal operator of the metric distance (Chapter 6, [89]) with parameter η, i.e.,

$$\mathrm{prox}_{d_i}^\eta(\bar{f}_n) = \bar{f}_n + \min\left\{1, \frac{\eta}{d_i(\bar{f}_n)}\right\}(P_{H_i}(\bar{f}_n) - \bar{f}_n) \qquad (A.16)$$

and the proximal average decomposition in (A.11), we can rewrite the algorithm step (A.10b) as

$$f_{n+1} = \bar{f}_n - \sum_{i=l_n}^n \omega_i^{(n)} \min\left\{1, \frac{\eta}{d_i(\bar{f}_n)}\right\} \bar{\beta}_i^{(n)} k(x_i, \cdot). \qquad (A.17)$$

Finally, assuming that the algorithm does not have access to any a priori information when it encounters the first data sample, we can set $f_1 = 0$. Then, from the algorithm step (A.10a), substituting each function estimate by its kernel expansion, i.e., $f_n = \sum_{i=1}^{n-1} \alpha_i^{(n)} k(x_i, \cdot)$ and identifying terms in (A.17), we obtain the following closed-form update rule for the non-parametric coefficients

$$\alpha_i^{(n+1)} = \begin{cases} \rho\alpha_i^{(n)} - \omega_i^{(n)}\Gamma_{\eta,i}^{(n)} & \text{if } i \in [1, n-1], \\ -\omega_i^{(n)}\Gamma_{\eta,i}^{(n)} & \text{if } i = n, \end{cases} \qquad (A.18)$$

where $\Gamma_{\eta,i}^{(n)} \triangleq \min\left\{|\bar{\beta}_i^{(n)}|, \eta k(x_i, x_i)^{-\frac{1}{2}}\right\} \mathrm{sign}(\bar{\beta}_i^{(n)})$ if $i \in [l_n, n]$ and equals zero otherwise.

A.3.2.1 Sparsification

The WORM algorithm, like many other kernel-based algorithms, suffers from the curse of kernelization [82], i.e., unbounded linear growth in model size and update time with the amount of data. For the considered application in Sec. A.3.2, a simple complexity control mechanism as kernel series truncation allows to preserve, to some extent, both performance as well as theoretical tracking guarantees, as we show in Secs. A.4 and A.5. Thus, given a user-defined truncation parameter $\tau \in \mathbb{N}$, such that $\tau > L$, if the number of effective coefficients constituting the function estimate f_n exceeds τ, we remove the older expansion term, i.e.,

$$e_n = \alpha_{n-\tau}^{(n)} k(x_{n-\tau}, \cdot), \qquad (A.19)$$

where $\alpha_{n-\tau}^{(n)} = \rho^{n-\tau+L}\alpha_{n-\tau}^{(n-\tau+L)}$. For the sake of illustration, consider a Gaussian reproducing kernel, i.e., $k(x,t) = \exp(-(x-t)^2/(2\sigma^2))$ with positive width σ. Then, the contribution of the truncated term e_n at timestamp x_n depends on the ratio $(x_{n-\tau} - x_n)^2/\sigma^2$; hence, a bigger ratio implies a smaller *truncation error*. Truncating before the algorithm step (A.10a) allows distributing the effects of the truncation error among the elements within the window. **Algorithm 2** describes in pseudocode our truncated WORM algorithm.

Algorithm 2 truncated WORM

Input: The data tuples $\{(x_n, y_n)\}_n^N$, the quantization half step-size ϵ, an RKHS $\mathcal{H}$, the window length L, the regularization parameter λ, the learning rate η, the power m and the truncation parameter τ.

1: Set $\boldsymbol{\alpha} := \text{queue}([\,], \text{maxlen} = \tau)$.
2: **for** $n = 1, 2, \ldots$ **do**
3:     Append one zero to the queue $\boldsymbol{\alpha}$.
4:     Set $\boldsymbol{\alpha} := (1 - \eta\lambda)\boldsymbol{\alpha}$.
5:     Set $\zeta_L := \{\max\{1, n - L + 1\}, \ldots, n\}$ and
    $\zeta_\tau := \{\max\{1, n - \tau + 1\}, \ldots, n\}$.
6:     **for** i in ζ_L **do**
7:         Compute $\bar{f}(x_i) := \sum_{j \in \zeta_\tau} \alpha_{j-\max\{n-\tau,0\}} k(x_j, x_i)$
8:         Compute $\bar{\beta}_i$ w.r.t. $\bar{f}$ as in (A.14).
9:     **end for**
10:     Set $\zeta_{\bar{f}} := \{i \in \zeta_L : \bar{\beta}_i \neq 0\}$.
11:     **for** i in $\zeta_{\bar{f}_n}$ **do**
12:         Compute the convex weights ω_i as in (A.15).
13:         Compute $\Gamma_{\eta,i} := \min\{|\bar{\beta}_i|, \eta k(x_i, x_i)^{-\frac{1}{2}}\}\text{sign}(\bar{\beta}_i)$.
14:         Update $\alpha_i := \alpha_i - \omega_i\Gamma_{\eta,i}$.
15:     **end for**
Output: The vector $\boldsymbol{\alpha}$, which yields the function estimate
    $f_n = \sum_{i=\max\{1, n-\tau+1\}}^{n-1} \alpha_{i-\max\{n-\tau,0\}} k(x_i, \cdot)$.
16: **end for**

A.4   Dynamic regret analysis

In this section, we derive a theoretical upper bound for the dynamic regret incurred by the truncated WORM algorithm. As a standard assumption [87], suppose that the norms $\|\mathcal{C}'_n(f_n)\|_{\mathcal{H}}$ are bounded by a positive constant G, i.e.,

$$\sup_{f_n \in \mathcal{H}, n \in [1,N]} \|\mathcal{C}'_n(f_n)\|_{\mathcal{H}} \leq G. \tag{A.20}$$

For the sake of notation, we omit the sub-index $\mathcal{H}$ in inner products and norms since the RKHS is clear by context. Considering the assumption in (A.20) and the

first order convexity condition of the windowed cost,

$$
\begin{aligned}
\mathbf{Reg}_N = \sum_{n=1}^{N} \mathcal{C}_n(f_n) - \mathcal{C}_n(f_n^*) &\leq \sum_{n=1}^{N} \langle \mathcal{C}_n'(f_n), f_n - f_n^* \rangle \\
&\leq \sum_{n=1}^{N} \|\mathcal{C}_n'(f_n)\| \|f_n - f_n^*\| \leq G \sum_{n=1}^{N} \|f_n - f_n^*\|,
\end{aligned}
\tag{A.21}
$$

it is clear that the dynamic regret is bounded above.

Consider the distance between the function estimate f_{n+1} and the optimal estimate $f_n^* = \arg\min_{f \in \mathcal{H}} \mathcal{C}_n(f_n)$, i.e.,

$$
\|f_{n+1} - f_n^*\| = \|\mathrm{prox}_{\mathcal{L}_n^\eta}^{\eta}(\bar{f}_n) - \mathrm{prox}_{\mathcal{L}_n^\eta}^{\eta}(\bar{f}_n^*)\|.
\tag{A.22}
$$

Hence, from the relation (A.22), the firmly non-expansiveness of the proximal operator [37], and the method step (A.10a) with truncation, we achieve the following inequality

$$
\|f_{n+1} - f_n^*\| \leq \rho \|f_n - e_n - f_n^*\| \leq \rho \|f_n - f_n^*\| + \rho \|e_n\|,
\tag{A.23}
$$

with coefficient $\rho \triangleq (1 - \eta\lambda) \in [0, 1)$. Finally, we can rewrite

$$
\begin{aligned}
\sum_{n=1}^{N} \|f_n - f_n^*\| &= \|f_1 - f_1^*\| + \sum_{n=2}^{N} \|f_n - f_n^*\| \\
&= \|f_1 - f_1^*\| + \sum_{n=2}^{N} \|f_n - f_{n-1}^* + f_{n-1}^* - f_n^*\| \\
&\leq \|f_1 - f_1^*\| + \sum_{n=2}^{N} \|f_n - f_{n-1}^*\| + \sum_{n=2}^{N} \|f_n^* - f_{n-1}^*\| \\
&\leq \|f_1 - f_1^*\| + \rho \sum_{n=2}^{N} \|f_{n-1} - f_{n-1}^*\| + \mathrm{C}_N + \rho \mathrm{E}_N \tag{A.24a} \\
&\leq \|f_1 - f_1^*\| + \rho \sum_{n=1}^{N} \|f_n - f_n^*\| + \mathrm{C}_N + \rho \mathrm{E}_N, \tag{A.24b}
\end{aligned}
$$

where the step (A.24a) comes after using the relation (A.23), the definition of cumulative variation in (A.5), and renaming the *cumulative truncation error* $\mathrm{E}_N \triangleq \sum_{n=2}^{N} \|e_n\|$. In step (A.24b), we rename the summation index and add the positive term $\rho \|f_N - f_N^*\|$ to the right hand-side of the inequality.

Regrouping the terms in (A.24), leads to

$$
\sum_{n=1}^{N} \|f_n - f_n^*\| \leq \frac{1}{1 - \rho} \left(\|f_1 - f_1^*\| + \mathrm{C}_N + \rho \mathrm{E}_N \right)
\tag{A.25}
$$

and substituting the relation obtained in (A.25) into the inequality (A.21), allows to upper-bound the dynamic regret as

$$
\mathbf{Reg}_N \leq \frac{G}{1 - \rho} \left(\|f_1 - f_1^*\| + \mathrm{C}_N + \rho \mathrm{E}_N \right).
\tag{A.26}
$$

Figure A.1: Average q-inconsistency of the sequence of function estimates $\{f_n\}_{n=1}^{100}$ over 500 different quantized signals.

This result explicitly shows the trade-off between tracking accuracy and model complexity [85]. In other words, without truncation, the dynamic regret reduces to $\mathbf{Reg}_N \leq \mathcal{O}(1 + \mathrm{C}_N)$, depending entirely on the environment. On the other hand, if we control the complexity of the function estimates via any truncation strategy such that the norm of the truncation error is upper bounded by a positive constant, i.e., $\sup_{n \in [1,N]} \|e_n\| \leq \delta$, the dynamic regret reduces to $\mathbf{Reg}_N \leq \mathcal{O}(\mathrm{C}_N + \delta T)$, leading to a steady tracking error in well-behaved environments.

A.5   Experimental results

As suggested in Sec. A.1, we compare the performance of our algorithm WORM with the KAPSM algorithm [23] and the augmented version of NORMA proposed in [85]. Moreover, since complexity control methods aim to limit the model order of the function estimate by lower-order approximations, we do not consider here any of them in order to isolate their effects on the performance of the algorithms. However, we have considered the truncated version of the WORM algorithm in our experiments to show that even a low complexity control technique such as truncation may lead to competitive performance.

Considering the application described in Section A.3.2, we have generated quantized versions of 500 realizations of a given AR(1) process. Each realization has been carried out for 100 data samples. In turn, each sample has been computed recurrently via $z_n = \varphi z_{n-1} + u_n$ with $z_0 = 0$, parameter $\varphi = 0.9$ and Gaussian noise $u_n \sim \mathcal{N}(0, 1)$. The center of the quantization intervals is computed by means of $y_n = \mathrm{round}(z_n/\epsilon) \cdot \epsilon$ with quantization half-step $\epsilon = 0.5$. The corresponding times-

Figure A.2: Average complexity of the sequence of function estimates $\{f_n\}_{n=1}^{100}$ over 500 different quantized signals.

Figure A.3: Comparison of regression plots for the last function estimate f_{100}, over the last 45 data samples of a synthetically generated quantized signal.

tamps $\{x_n\}_{n=1}^{100}$ are uniformly arranged. For the sake of illustration, we use a Gaussian reproducing kernel, i.e., $k(x,t) = \exp(-(x-t)^2/(2\sigma^2))$, with $\sigma = 3$. All four algorithms use the same window length $L = 10$. As to the augmented NORMA, the WORM algorithm and its truncated version, all use the same learning rate $\eta = 1.5$ and regularization parameter $\lambda = 0.005$. We restrict the truncated WORM function estimates expansion to a maximum of 30 terms, i.e., $\tau = 30$. Both versions of WORM use the power $m = 2$. For the augmented NORMA, the instantaneous loss terms within the nth window are equally weighted with the weight $\min\{n, L\}^{-1}$ and $\partial_f d_i(f_n) = \text{sign}(\beta_i^{(n)}) k(x_i, x_i)^{-\frac{1}{2}} k(x_i, \cdot)$ is used as a valid functional subgradient. We also define the *q-inconsistency*, i.e., $\sum_{i=q_n}^{n} d_i(f_n)$, with $q_n = \max\{n - q + 1, 1\}$ and $q = 20$, and use the squared Hilbert norm, $\|f_n\|_{\mathcal{H}}^2 = \sum_{i,j=\tau_n}^{n} \alpha_i^{(n)} \alpha_j^{(n)} k(x_i, x_j)$, with $\tau_n = \max\{n - \tau + 1, 1\}$, as performance metrics for the function estimates. The first metric measures how far is the function estimate of falling into the last q received quantization intervals. The second metric measures the function estimate complexity.

As shown in Fig. A.1 and Fig. A.2, there is a trade-off between q-inconsistency and complexity. The WORM algorithm successfully balances both altogether. As to its truncated version, the same experimental results show that the complexity can be successfully controlled at the expense of little accuracy. Finally, Fig. A.3 shows a snapshot of the last function estimate f_{100} for each algorithm.

A.6 Conclusion

In this paper, we propose a novel algorithm, WORM, for regression-based tracking of quantized signals. We derive a theoretical dynamic regret bound for WORM that ensures tracking guarantees. Our experiment shows that WORM provides better signal reconstruction in terms of consistency and smoothness altogether compared to the state-of-the-art.

Appendix B

Paper B

Title:	An Online Multiple Kernel Parallelizable Learning Scheme
Authors:	**Emilio Ruiz-Moreno** and Baltasar Beferull-Lozano
Journal:	IEEE Signal Processing Letters

An Online Multiple Kernel Parallelizable Learning Scheme

Emilio Ruiz-Moreno and Baltasar Beferull-Lozano

Abstract — The performance of reproducing kernel Hilbert space-based methods is known to be sensitive to the choice of the reproducing kernel. Choosing an adequate reproducing kernel can be challenging and computationally demanding, especially in data-rich tasks without prior information about the solution domain. In this paper, we propose a learning scheme that scalably combines several single kernel-based online methods to reduce the kernel-selection bias. The proposed learning scheme applies to any task formulated as a regularized empirical risk minimization convex problem. More specifically, our learning scheme is based on a multi-kernel learning formulation that can be applied to widen any single-kernel solution space, thus increasing the possibility of finding higher-performance solutions. In addition, it is parallelizable, allowing for the distribution of the computational load across different computing units. We show experimentally that the proposed learning scheme outperforms the combined single-kernel online methods separately in terms of the cumulative regularized least squares cost metric.

B.1   Introduction

Reproducing kernel Hilbert space (RKHS)-based methods allow modeling highly non-linear relationships at a moderate computational cost [60]. Thanks to their simplicity and generality, they have been successfully adopted in a wide range of signal-processing applications [90, 91].

The performance of any RKHS-based method strongly relies on a preselected reproducing kernel (RK). The efficient selection of an adequate RK presumes some task-specific prior information, such as knowledge about the data domain, invariant data transformations, geometrical data structures, or some properties of the underlying data generating process [92]. For example, spline interpolation RKs are best suited for smooth data [24, 78]. Similarly, radial basis RKs can perform poorly if their associated hyperparameters are not properly tuned to the task. The kernel-selection issue cannot be easily mitigated via cross-validation [93] because the associated computational load grows prohibitively with the number of RKs. On the other hand, efficiently computable approximations of the leave-one-out error [61] or hyperparameter optimization techniques [94] usually involve non-convexity and may lead to undesirable local minima.

Multi-kernel methods compensate for the lack of task-specific prior information using a predefined set of RKs known as *dictionary*. The dictionary can be formed by integrating different types of RKs, the same RK with different hyperparameter values, or a mix of both. Typically, the preselected RK is constructed as a combination of several RKs from the dictionary [62]. Therefore, how the dictionary is formed and how the preselected RK is constructed have a pivotal impact on the resulting accuracy and complexity of the method. For instance, the larger the dictionary

is, the more likely it is to reduce the kernel-selection bias compared to a particular RK or hyperparameter choice. In addition, larger dictionaries allow for greater adaptability when learning from data samples since, in practice, these samples may come from a combination of different sources. On the other hand, increasing the dictionary size becomes computationally demanding or even prohibitive. For this reason, a commonly sought goal for multi-kernel methods is to find a compromise between performance and a computationally light and compact representation of the proposed solution in terms of the dictionary elements [95].

Another related scaling issue that also applies to single-kernel methods is the curse of kernelization, i.e., potentially unbounded linear growth in model size with the amount of data [82]. This drawback is generally addressed through online approaches, which may rely on sparsification procedures [96, 97, 98, 99, 100] or dimensionality reduction approximations [101, 102], among others [103].

In this context, some works [104, 105, 106] have explored the use of online methods for determining the optimal solution associated with each single RK within the dictionary, as well as the best combination of these single-kernel solutions under a given task. Following a conceptually similar approach, this paper proposes a novel multi-kernel learning scheme that can be parallelized across RKs by efficiently combining the solution of several single RK-based online methods running concurrently. This provides scalability with respect to the number of data samples and adaptability across different data patterns. Moreover, it allows to distribute the computational load across different computing units as the dictionary size increases. Our proposed scheme applies to any task that can be formulated as a regularized empirical risk minimization (RERM) convex problem [39, 107]. Finally, the performance of the proposed learning scheme is experimentally validated in terms of the cumulative regularized least squares cost metric.

B.2   Problem formulation

Supervised learning is arguably one of the core topics in machine learning [108]. Many supervised learning tasks can be formulated as RERM convex problems whose solution admits a kernel representation. That is, given a set of N data samples $\mathcal{S} = \{(x^{(n)}, y^{(n)})\}_{n=1}^{N} \subseteq \mathcal{X} \times \mathcal{Y}$, and an RKHS[1] $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$, the goal is to find a function estimate $f \in \mathcal{H}$ minimizing the following regularized functional cost

$$\mathcal{C}_{\eta}(f;\mathcal{S}) = \sum_{n=1}^{N} \ell\left(f(x^{(n)}), y^{(n)}\right) + \frac{\eta}{2}\|f\|_{\mathcal{H}}^{2}, \tag{B.1}$$

where the loss $\ell : \mathcal{Y}^2 \to \mathbb{R} \cup \{\infty\}$ is a proper convex function used as a goodness-of-fit metric, the regularizer $\|\cdot\|_{\mathcal{H}}^{2} : \mathcal{H} \to \mathbb{R}$ is the squared RKHS $\mathcal{H}$ induced norm, and the hyperparameter $\eta \in \mathbb{R}_{+}$ controls the model complexity of the solution.

Under a multi-kernel learning framework, one typically constructs a valid RK by adequately combining the RKs within a preselected dictionary [62]. Particularly, finding the RK within a convex hull of P positive definite RKs that yields the function estimate incurring the lowest functional cost (B.1) is equivalent to obtaining a solution from $\mathcal{H}$ built as the RKHS direct sum $\mathcal{H}_1 \oplus \cdots \oplus \mathcal{H}_P$, where each pth

---

[1]The notation $\mathcal{Y}^{\mathcal{X}}$ refers to the set of functions from $\mathcal{X}$ to $\mathcal{Y}$.

Figure B.1: Visual description of the online setting considered.

RKHS $\mathcal{H}_p = \overline{\text{span}}\{k_p(x, \cdot) : x \in \mathcal{X}\}$, being $k_p : \mathcal{X}^2 \to \mathbb{R}$ its associated RK [109]. The solution $f \in \mathcal{H}$ that minimizes (B.1), can be expressed without loss of generality as $f = \boldsymbol{\theta}^\top \boldsymbol{f}$, where $\boldsymbol{f} = [f_1, \ldots, f_P]^\top \in \mathcal{H}_{1:P}$, with $\mathcal{H}_{1:P} \triangleq \mathcal{H}_1 \times \cdots \times \mathcal{H}_P$ and $\boldsymbol{\theta} = [\theta_1, \ldots, \theta_P]^\top \in \Delta^{P-1}$, with $\Delta^{P-1} \triangleq \{\boldsymbol{\beta} \in \mathbb{R}^P : \boldsymbol{\beta} \succeq \mathbf{0} \text{ and } \mathbf{1}^\top \boldsymbol{\beta} = 1\}$ denoting a simplex [110]. Thus, the RERM problem posed before becomes

$$\min_{\boldsymbol{\theta} \in \Delta^{P-1}, \boldsymbol{f} \in \mathcal{H}_{1:P}} \mathcal{C}_\eta (f; \mathcal{S}) \tag{B.2a}$$

$$\text{subject to: } f = \boldsymbol{\theta}^\top \boldsymbol{f}. \tag{B.2b}$$

Optimization problem (B.2) is bi-convex, meaning that it is convex in $\boldsymbol{\theta}$ for a fixed $\boldsymbol{f}$ and vice-versa. Still, it is not jointly convex in both optimization variables. It can be tackled via specialized methods that primarily exploit the convex substructures of the problem [111]. However, these methods do not scale well with the number of RKs and data samples, denoted as P and N, respectively.

This paper presents a method to solve efficiently (B.2) for large P and N.

B.3 Proposed solution

This section describes how to synergize an online formulation and an upper bound on the objective (B.2a) to solve (B.2) scalably with respect to N and P.

B.3.1 Online setting

Online settings [63, 64] can be adopted to solve (B.2) achieving low run-time complexity with respect to N while incurring a certain tolerable (cumulative) cost. They usually trade-off solution accuracy for speed, e.g., by processing only a few samples every iteration, for low memory complexity, e.g., by discarding samples after a few processing steps, or for model complexity control, i.e., bounded model size regardless of whether N increases.

The online setting considered in this work can be cast as a method-environment iterative game [112]. The data samples in $\mathcal{S}$ are assumed to be available sequentially. Then, at each iteration step n, the method chooses a function estimate $f^{(n)} \in \mathcal{H}_1 \oplus \cdots \oplus \mathcal{H}_P$ expressible as $f^{(n)} = \boldsymbol{\theta}^{(n)\top} \boldsymbol{f}^{(n)}$ with $\boldsymbol{\theta}^{(n)} \in \Delta^{P-1}$ and $\boldsymbol{f}^{(n)} \in \mathcal{H}_{1:P}$. In response, the environment penalizes the proposed function estimate $f^{(n)}$ with the incurred cost $\mathcal{C}_\eta(f^{(n)}; \mathcal{S}_L^{(n)})$, where $\mathcal{S}_L^{(n)} = \{(x^{(i)}, y^{(i)})\}_{i=n_L}^n \subseteq \mathcal{S}$ is a sliding window of L data samples and $n_L \triangleq \max\{n - L + 1, 1\}$. Finally, once the nth function

---

Algorithm 3 (Quadratic program) Projection onto simplex.

Input: The function estimate components $\{f_p^{(n)}\}_{p=1}^P$, the data window $\mathcal{S}_L^{(n-1)}$ and the loss ℓ.

1: Set the components of $\boldsymbol{a}$ and $\boldsymbol{b}$ as in (B.5).

2: Sort $\boldsymbol{b}$ in ascending order, denoted as $\boldsymbol{u}$: $u_1 \leq u_2 \leq \cdots \leq u_P$ and rearrange $\boldsymbol{a}$ accordingly into $\boldsymbol{v}$.

3: Find $\rho = \max \left\{ 1 \leq j \leq P : u_j - \frac{\sum_{i=1}^{j} \frac{u_i}{v_i} + 2}{\sum_{i=1}^{j} \frac{1}{v_i}} < 0 \right\}$.

4: Define $\mu = -\frac{2 + \sum_{i=1}^{\rho} \frac{u_i}{v_i}}{\sum_{i=1}^{\rho} \frac{1}{v_i}}$.

5: Compute $\theta_p = \max \left\{ -\frac{1}{2a_p}(b_p + \mu), 0 \right\}$ for $p = 1, \ldots, P$.

Output: $\boldsymbol{\theta}$.

---

estimate $f^{(n)}$ is chosen, the method receives the nth data window $\mathcal{S}_L^{(n)}$, which can be used at the next iteration[2] step $n+1$. Fig. B.1 visually describes the procedure.

B.3.2 Upper bound on the functional cost

As we show next, we can improve scalability with respect to the number of RKs, by making use of the following upper bound:

$$
\mathcal{C}_\eta(f; \mathcal{S}) = \sum_{n=1}^{N} \ell\left(\boldsymbol{\theta}^\top \boldsymbol{f}(x^{(n)}), y^{(n)}\right) + \frac{\eta}{2}\|\boldsymbol{\theta}^\top \boldsymbol{f}\|_{\mathcal{H}}^2 \tag{B.3a}
$$

$$
\leq \sum_{n=1}^{N} \sum_{p=1}^{P} \theta_p \ell\left(f_p(x^{(n)}), y^{(n)}\right) + \theta_p^2 \frac{\eta}{2}\|f_p\|_{\mathcal{H}_p}^2 \tag{B.3b}
$$

$$
\triangleq \breve{\mathcal{C}}_\eta(\boldsymbol{\theta}, \boldsymbol{f}; \mathcal{S}). \tag{B.3c}
$$

The first upper-bounded term in (B.3b) follows directly from Jensen's inequality [113], whereas the second term is obtained by invoking the definition of the RKHS direct sum norm [114]. Even though the second term in (B.3b) could have also been upper bounded through Jensen's inequality, as in [110], exploiting the definition of the RKHS direct sum norm, leads to a tighter bound because $\theta_p^2 \leq \theta_p$ for $\theta_p \in [0,1]$.

The key advantage of the upper bound cost (B.3c) is that it is separable across the P RKs within the dictionary, hence allowing for parallelization at the expense of some loss in optimality, albeit with still satisfactory performance.

B.3.3 Parallelizable learning scheme

Our proposed learning scheme consists of executing at each iteration step n the following consecutive operations:

1) Every pth function estimate component $f_p^{(n)}$ is chosen through a single-kernel online method operating over the pth RK within the dictionary. The methods are selected by the user, and they can be different for each RK as long as all of

---

[2]At the first iteration step $n = 1$, the method has not received any data sample, thus $f^{(1)}$ is set as some arbitrary initial function estimate.

them adopt the online setting described in Sec. B.3.1. For example, stochastic gradient descent methods for function estimation [12], and associated variants [81, 1], can be readily used. Since the function estimate components of $\boldsymbol{f}^{(n)}$ can be computed in parallel across P different computing units, the computational cost can be distributed.

2) Next, the convex weights in $\boldsymbol{\theta}^{(n)}$ are chosen as the ones that minimize the partially evaluated upper bound cost (B.3c) at $\boldsymbol{f}^{(n)}$ and $\mathcal{S}_L^{(n-1)}$, referred from now on as learning cost. Mathematically,

$$\boldsymbol{\theta}^{(n)} = \ \arg\min_{\boldsymbol{\theta}\in\Delta^{P-1}} \ \check{\mathcal{C}}_\eta\left(\boldsymbol{\theta}, \boldsymbol{f}^{(n)}; \mathcal{S}_L^{(n-1)}\right) \tag{B.4a}$$

$$= \ \arg\min_{\boldsymbol{\theta}\in\Delta^{P-1}} \ \boldsymbol{\theta}^\top \boldsymbol{A}^{(n)}\boldsymbol{\theta} + \boldsymbol{b}^{(n)^\top}\boldsymbol{\theta}, \tag{B.4b}$$

where $\boldsymbol{A}^{(n)} \triangleq \mathrm{diag}(\boldsymbol{a}^{(n)}) \in \boldsymbol{S}_{++}^P$, with $\boldsymbol{a}^{(n)} \in \mathbb{R}_+^P$, and $\boldsymbol{b}^{(n)} \in \mathbb{R}^P$ whose components are computed[3] as

$$a_p^{(n)} = \frac{\eta}{2}\left\|f_p^{(n)}\right\|_{\mathcal{H}_p}^2, \tag{B.5a}$$

$$b_p^{(n)} = \sum_{i\in\mathcal{I}_L^{(n-1)}} \ell\left(f_p^{(n)}(x^{(i)}), y^{(i)}\right), \tag{B.5b}$$

for all $p \in \mathbb{N}^{[1,P]}$, where $\mathcal{I}_L^{(n-1)}$ corresponds to the index set associated with the data samples in $\mathcal{S}_L^{(n-1)}$. We adapt the projection onto the simplex algorithm discussed in [115, 116, 117] by extending its applicability to any quadratic problem described by a diagonal positive definite matrix with simplex constraints. As a result, the proposed **Algorithm 3** can solve (B.4) exactly. Its computational complexity is bottlenecked by a sorting step; that is, an asymptotic average complexity $\mathcal{O}(P\log P)$ [118]. It should be mentioned that this complexity can be further reduced to $\mathcal{O}(P)$ on average by using a randomized pivot algorithm variation that identifies the parameter ρ (**Algorithm 3**, line 3) using a divide and conquer procedure instead of sorting [119], but this is out of the scope of the present paper.

3) Finally, the function estimate $f^{(n)} = \boldsymbol{\theta}^{(n)^\top}\boldsymbol{f}^{(n)}$ is proposed.

In summary, our scheme can be seen as a higher-level learner that iteratively chooses the lowest incurring learning cost combination of function estimates provided by lower-level learners, namely, single RK methods.

B.4 Performance analysis

Under an online setting, as the one described in Sec. B.3.1, the incurred cost accumulated over time receives the name of *cumulative* cost (CC). In our case, the CC up to the nth time step is given by $\sum_{i=1}^n \mathcal{C}_\eta(f^{(i)}; \mathcal{S}_L^{(i)})$. From here, recall that every ith function estimate $f^{(i)}$ is proposed via Sec. B.3.3, before the ith data window $\mathcal{S}_L^{(i)}$ becomes available; thus, the CC is a measure of performance protecting against overfitting. Intuitively, the lower the growth of the incurred CC with respect to n,

---

[3]Notice that nothing prevents $f_p^{(n)}$ to be zero-valued and thus $\boldsymbol{A}^{(n)}$ from being singular and positive semi-definite. However, we can always set $\boldsymbol{A}^{(n)} = \mathrm{diag}(\boldsymbol{a}^{(n)}) + \delta\boldsymbol{I}_P$ where δ is an arbitrarily small positive value.

the better the expected performance over unseen data. In fact, popular measures of performance, such as the dynamic regret, are constructed as the difference between the CC incurred by the sequence of function estimates proposed by a method and a sequence of comparators [66].

In order to validate our scheme experimentally, we pose a signal reconstruction online problem from synthetically generated streaming data. Specifically, we use the squared loss, i.e., $\ell(f(x), y) = (f(x) - y)^2$, and a dictionary of $P = 20$ Gaussian kernels, i.e., $k_p(x, t) = \exp\left(-\frac{1}{2}(x-t)^2/\sigma_p^2\right)$, with different widths σ_p linearly spaced between 0.1 and 10. The method associated with each RK is an augmented naive online R_{reg} minimization algorithm (NORMA) [12] with a window length of $L = 10$ data samples, a budget of $\tau = 100$ kernel expansion terms (beyond the allowed budget we truncate the oldest terms of the kernel expansion), and a fixed learning rate $\lambda_{\text{NORMA}} = 0.05$. That is, before any possible truncation, each nth function estimate associated with the pth RK is constructed as $f_p^{(n)} = \sum_{i=1}^{n-1} \alpha_{p,i}^{(n)} k_p(x^{(i)}, \cdot)$, where each $\alpha_{p,i}^{(n)} \in \mathbb{R}$ denotes a kernel-expansion coefficient obtained from the following NORMA update:

$$f_p^{(n)} = f_p^{(n-1)} - \lambda_{\text{NORMA}} \, \partial_f \, \mathcal{C}_\eta \left(f; \mathcal{S}_L^{(n-1)}\right)\bigg|_{f=f_p^{(n-1)}}, \qquad (B.6)$$

which, after some algebraic steps, leads to the next closed-form update rule:

$$\alpha_i^{(n)} = \begin{cases} -\lambda_{\text{NORMA}} \, \ell_{p,i}'^{(n-1)} & \text{if } i = n - 1, \\ \gamma \alpha_i^{(n-1)} - \lambda_{\text{NORMA}} \, \ell_{p,i}'^{(n-1)} & \text{if } i \in \mathcal{I}_L^{(n-1)} \backslash \{n - 1\}, \\ \gamma \alpha_i^{(n-1)} & \text{otherwise}, \end{cases} \qquad (B.7)$$

where $\ell_{p,i}'^{(n)} \triangleq \ell'\left(f_p^{(n)}(x^{(i)}), y^{(i)}\right) = 2\left(f^{(n)}(x^{(i)}) - y^{(i)}\right)$ and $\gamma \triangleq (1 - \lambda_{\text{NORMA}}\eta) \in \mathbb{R}_+^{(0,1)}$. The regularization parameter is chosen as $\eta = 0.01$. Lastly, the data samples have been generated via a stable AR(1) process $y^{(n)} = \varphi y^{(n-1)} + u^{(n)}$, with $\varphi = 0.5488135$, $u^{(n)} \overset{\text{iid}}{\sim} \mathcal{N}(0, 0.71519837)$, $y^{(0)} = 0$ and unit time stamps uniformly arranged in time, i.e., $x^{(n)} = n$.

Additionally, we compare our scheme with the online multiple kernel regression (OMKR) algorithm [104], arguably the closest approach conceptually. More specifically, we compare against the budget OMKR gradient-based variant method over the same experimental setting described above. Briefly, the considered OMKR method can be described, at each iteration step n, by the following three-stage scheme:

1) The set of function estimates, in this case, regressors proposed by each one of the P NORMAs, is updated as (B.7) and collected in $\boldsymbol{f}^{(n)} \in \mathcal{H}_{1:P}$.

2) Then, the P weights for combining the multiple regressors are updated as

$$\boldsymbol{w}^{(n)} = \boldsymbol{w}^{(n-1)} - \lambda_{\text{OMKR}}^{(n)} \, \nabla_{\boldsymbol{w}} \, \mathcal{C}_\eta \left(\boldsymbol{w}^\top \boldsymbol{f}^{(n)}; \mathcal{S}_L^{(n-1)}\right)\bigg|_{\boldsymbol{w}=\boldsymbol{w}^{(n-1)}}. \qquad (B.8)$$

In this case, we use an initial learning rate $\lambda_{\text{OMKR}}^{(1)} = 8 \cdot 10^{-4}$ that is halved every 50 steps until a minimum value of 10^{-5}. After some algebraic manipulations and making use of the definition of the RKHS direct sum norm [114], the evaluated gradient in (B.8) equals to

$$\sum_{i \in \mathcal{I}_L^{(n-1)}} \boldsymbol{f}^{(n)}(x^{(i)}) \ell'\left(\boldsymbol{w}^{(n-1)\top} \boldsymbol{f}^{(n)}(x^{(i)}), y^{(i)}\right) + \eta \boldsymbol{A}^{(n)} \boldsymbol{w}^{(n-1)}. \qquad (B.9)$$

Figure B.2: Cumulative cost up to time step n incurred by our learning scheme, the OMKR algorithm, and the combined single RK NORMA regressors individually. The dashed-dotted line $L \cdot n$ is shown as a reference.

The matrix $\boldsymbol{A}^{(n)}$ corresponds to the one introduced in (B.5a), and the initial combination weights are set as $\boldsymbol{w}^{(1)} = \boldsymbol{0}_P$.

3) Finally, the function estimate $f^{(n)} = \boldsymbol{w}^{(n)\top} \boldsymbol{f}^{(n)}$ is proposed.

Unlike our scheme, the OMKR algorithm can eventually learn any linear combination of single-kernel function estimates. However, due to the additive nature of the update step in (B.8), the OMKR algorithm usually suffers from slow convergence rates. Moreover, it requires a sequence of learning rates $\lambda_{\mathrm{OMKR}}^{(n)}$ whose tuning involves optimization techniques or task-specific knowledge, hence adversely affecting performance, e.g., poor learning or instabilities, if not carried out adequately.

Our experimental results in Fig. B.2 show that the CC incurred by our proposed learning scheme outperforms the lowest CC incurred by any of the combined single RK NORMA regressors separately. In the same figure, it can also be observed that our scheme incurs a CC that increases at a lower rate than the one incurred by the OMKR algorithm, thus allowing our scheme to outperform the best NORMA regressor much sooner. The reason behind this observation is arguably the additive nature of the OMKR algorithm, which requires numerous updates to completely remove the residual contributions of irrelevant regressors. For example, see in Fig. B.3 the "spiky" shape of the OMKR signal estimate due to some regressors constructed with a narrow-valued σ_p RK.

Regarding computational resources, both the OMKR algorithm and our scheme can be parallelized across the RKs within the dictionary, which in our experiments means a constant time complexity $\mathcal{O}(\tau)$, and a combination step of complexity $\mathcal{O}(P)$ and $\mathcal{O}(P \log P)$, respectively. However, as mentioned in Sec. B.3.3, the complexity of our scheme can be further reduced to $\mathcal{O}(P)$ on average, making both of the compared approaches computationally equivalent.

Figure B.3: Snapshot of the 42nd signal estimate obtained by the OMKR algorithm, our learning scheme, and the best and the worst of the combined single RK NORMA regressors (denoted by the indices p_+ and p_-, respectively) in terms of the so-far $(n = 42)$ incurred cumulative cost.

B.5   Conclusion

We present a multi-kernel learning scheme that experimentally outperforms the best of the combined single RK methods, in terms of the cumulative regularized least squares cost metric, with a comparable computational load per computing unit. This corroborates the ability of the proposed scheme to effectively accommodate a larger function space (from which to draw function estimates) of multi-kernel methods while keeping the lower computational complexity of online single RK methods. Furthermore, although **Algorithm 3** has been expressly designed for the task discussed in this paper, it can be used to solve any other problem that accepts a formulation as in (B.4b).

Correctness of Algorithm 1

A Lagrangian of problem (B.4) is

$$\mathrm{L}(\boldsymbol{\theta}, \boldsymbol{\lambda}, \mu) = \boldsymbol{\theta}^\top \mathrm{diag}(\boldsymbol{a}^{(n)})\boldsymbol{\theta} + \boldsymbol{b}^{(n)^\top}\boldsymbol{\theta} - \boldsymbol{\lambda}^\top\boldsymbol{\theta} + \mu(\mathbf{1}^\top\boldsymbol{\theta} - 1), \qquad (\text{B.10})$$

being $\mu \in \mathbb{R}$ and $\boldsymbol{\lambda} \in \mathbb{R}^P$ the Lagrange multipliers associated with the equality and inequality constraints, respectively. At the optimal solution $\boldsymbol{\theta}^{(n)}$, the following KKT conditions [120] hold:

Machine Learning for Signal Reconstruction from Streaming Time-series Data

$$2\theta_p^{(n)}a_p^{(n)} + b_p^{(n)} - \lambda_p + \mu = 0, \quad p = 1, \dots, P \tag{B.11a}$$

$$\theta_p^{(n)} \geq 0, \quad p = 1, \dots, P \tag{B.11b}$$

$$\lambda_p \geq 0, \quad p = 1, \dots, P \tag{B.11c}$$

$$\lambda_p \theta_p^{(n)} = 0, \quad p = 1, \dots, P \tag{B.11d}$$

$$\sum_{p=1}^{P} \theta_p^{(n)} = 1. \tag{B.11e}$$

From the complementary slackness, stated in (B.11d), we can deduce that if the primal inequality constraint in (B.11b) is slacked, i.e., greater than zero, then $\lambda_p = 0$ and from the stationarity condition (B.11a), the solution fulfils

$$\theta_p^{(n)} = -\frac{1}{2a_p^{(n)}}(b_p^{(n)} + \mu) > 0. \tag{B.12}$$

On the other hand, if the primal inequality constraint is tight, i.e., $\theta_p^{(n)} = 0$, then the dual constraint (B.11c) is not binding. Again, from the stationarity condition in (B.11a), we can identify those non-binding constraints as those that satisfy the following expression:

$$b_p^{(n)} + \mu = \lambda_p \geq 0. \tag{B.13}$$

In this way, it is clear from (B.13) that the components of the optimal solution that are zero, if any, correspond to the larger components of $\boldsymbol{b}^{(n)}$. Without loss of generality, we can assume that the components of $\boldsymbol{b}^{(n)}$ are sorted in ascending order as long as the components of $\boldsymbol{a}^{(n)}$ are rearranged accordingly. Thus, by comparing $\boldsymbol{b}^{(n)}$ with the solution as follows:

$$b_1^{(n)} \leq \cdots \leq b_\rho^{(n)} \leq b_{\rho+1}^{(n)} \leq \cdots \leq b_P^{(n)},$$
$$\theta_{\rho+1}^{(n)} = \cdots = \theta_P^{(n)} = 0, \tag{B.14}$$

it can be concluded that the index $\rho \in \mathbb{N}^{[1,P]}$ determines the number of components in the solution that are nonzero. From here, and rewriting the equality primal constraint (B.11e) as

$$\sum_{p=1}^{P} \theta_p^{(n)} = \sum_{p=1}^{\rho} \theta_p^{(n)} = -\frac{1}{2}\sum_{p=1}^{\rho}\frac{1}{a_p^{(n)}}(b_p^{(n)} + \mu) = 1, \tag{B.15}$$

the Lagrangian multiplier associated with the equality constraint can be isolated and computed as

$$\mu = -\frac{2 + \sum_{p=1}^{\rho}\frac{b_p^{(n)}}{a_p^{(n)}}}{\sum_{p=1}^{\rho}\frac{1}{a_p^{(n)}}}, \tag{B.16}$$

as long as the index ρ is known.

Theorem 1: Let ρ be the number of positive components in the solution of optimization problem (4), then

$$\rho = \max\left\{1 \leq j \leq P : b_j - \frac{2 + \sum_{i=1}^{j}\frac{b_i}{a_i}}{\sum_{i=1}^{j}\frac{1}{a_i}} < 0\right\}, \tag{B.17}$$

where $\boldsymbol{b}$ is obtained by sorting $\boldsymbol{b}^{(n)}$ components in ascending order and $\boldsymbol{a}$ corresponds to $\boldsymbol{a}^{(n)}$ rearranged accordingly.

Proof. Let us first define the quantities $\varphi_j \triangleq b_j - (2 + \sum_{i=1}^{j} \frac{b_i}{a_i})/\sum_{i=1}^{j} \frac{1}{a_i}$ and $s_{j:k} \triangleq \sum_{i=j}^{k} \frac{1}{a_i}$. Then, the goal is to show that $j = \rho$ is the largest index in $\{1, \ldots, P\}$ for which φ_j remains negative.

For $j < \rho$, we have that

$$\varphi_j = \frac{1}{s_{1:j}} \left(s_{1:j} b_j - \left(2 + \sum_{i=1}^{j} \frac{b_i}{a_i} \right) \right) \tag{B.18a}$$

$$= \frac{1}{s_{1:j}} \left(s_{1:j} b_j - 2 - \sum_{i=1}^{\rho} \frac{b_i}{a_i} + \sum_{i=j+1}^{\rho} \frac{b_i}{a_i} \right) \tag{B.18b}$$

$$= \frac{1}{s_{1:j}} \left(s_{1:j} b_j + s_{1:\rho} \mu + \sum_{i=j+1}^{\rho} \frac{b_i}{a_i} \right) \tag{B.18c}$$

$$= b_j + \mu + \frac{1}{s_{1:j}} \sum_{i=j+1}^{\rho} \frac{1}{a_i} (\mu + b_i) < 0, \tag{B.18d}$$

where in step (B.18b) we use the equivalence $\sum_{i=1}^{j} \frac{b_i}{a_i} = \sum_{i=1}^{\rho} \frac{b_i}{a_i} - \sum_{i=j+1}^{\rho} \frac{b_i}{a_i}$. Next, in step (B.18c), we make use of the relation in (B.16). Finally, the step (B.18d) holds thanks to the relation in (B.12) and because $s_{1:j}, a_i \geq 0 \ \forall i, j$.

For $j = \rho$, and thanks to (B.12), we have $\varphi_\rho = b_\rho + \mu < 0$. Then, using the relation in (B.16), we can verify that (B.17) holds.

For $j > \rho$, we can follow similar algebraic steps as in (B.18) to obtain

$$\varphi_j = \frac{1}{s_{1:j}} \left(s_{1:\rho} (b_j + \mu) + \sum_{i=\rho+1}^{j} \frac{1}{a_i} (b_j - b_i) \right) \geq 0. \tag{B.19}$$

The inequality in (B.19) holds thanks to the relation in (B.13) and the fact that $b_j \geq b_i \ \forall i$.

Appendix C

Paper C

Title: A Trainable Approach to Zero-delay Smoothing Spline Interpolation

Authors: **Emilio Ruiz-Moreno**, Luis Miguel López-Ramos and Baltasar Beferull-Lozano

Journal: IEEE Transactions on Signal Processing

A Trainable Approach to Zero-delay Smoothing Spline Interpolation

Emilio Ruiz-Moreno, Luis Miguel López-Ramos and Baltasar Beferull-Lozano

Abstract — The task of reconstructing smooth signals from streamed data in the form of signal samples arises in various applications. This work addresses such a task subject to a zero-delay response; that is, the smooth signal must be reconstructed sequentially as soon as a data sample is available and without having access to subsequent data. State-of-the-art approaches solve this problem by interpolating consecutive data samples using splines. Here, each interpolation step yields a piece that ensures a smooth signal reconstruction while minimizing a cost metric, typically a weighted sum between the squared residual and a derivative-based measure of smoothness. As a result, a zero-delay interpolation is achieved in exchange for an almost certainly higher cumulative cost as compared to interpolating all data samples together. This paper presents a novel approach to further reduce this cumulative cost on average. First, we formulate a zero-delay smoothing spline interpolation problem from a sequential decision-making perspective, allowing us to model the future impact of each interpolated piece on the average cumulative cost. Then, an interpolation method is proposed to exploit the temporal dependencies between the streamed data samples. Our method is assisted by a recurrent neural network and accordingly trained to reduce the accumulated cost on average over a set of example data samples collected from the same signal source generating the signal to be reconstructed. Finally, we present extensive experimental results for synthetic and real data showing how our approach outperforms the abovementioned state-of-the-art.

C.1 Introduction

Online learning has been studied and applied in a broad range of research fields, including optimization theory [30, 63, 64], signal processing [121], and machine learning [122, 45, 44]. Within these fields, online methods generating a series of estimates from sequentially streamed data are especially useful to reduce complexity in large-scale problems [19], to dynamically adapt to new patterns in the data [20], and to enable acting under real-time requirements [21].

This work addresses the last one of the previous use cases in the context of signal reconstruction. Specifically, it investigates the use of online methods with *zero-delay* response for *smooth* signal reconstruction. First, most physical signals are bounded and smooth due to energy conservation [123]; hence it is beneficial to maintain smoothness as a property during signal reconstruction. Second, the zero-delay requirement demands new portions of the smooth signal to be reconstructed as soon as a new data sample is available. Consequently, a reduced constant complexity per iteration [124] is required so that the online method is executed at a higher speed than the transmission rate at which the data samples are received. These requisites are well-motivated since they appear in many practical problems, such as

Figure C.1: A signal reconstructed from stream data points by different methods and models. The symbols x_t and f_t represent the current time and signal estimate, respectively. a) Recursive least squares [10, 11] with a linear model. b) NORMA [12] with Gaussian kernels. c) Smoothing myopic interpolation with cubic Hermite splines [13].

online trajectory planning [67, 68], real-time control systems [69, 70], and high-speed digital to analog conversion [71], among others. Although the tasks of estimating smooth signals or delivering a zero-delay response are separately managed by most online methods, handling them together becomes challenging, as we expound next.

Some popular online methods that can be used for smooth signal estimation are online kernel methods [60, 1] and online Gaussian processes [125, 126]. They aim at yielding a sequence of signal estimates with convergence guarantees or sublinear regret [77]. To this end, they initially propose a signal estimate which is updated (or modified) possibly globally as new data samples arrive. Their goal is to refine the signal estimate rather than reconstruct new portions of the smooth signal. Therefore, neither smoothness nor even continuity of the sequentially reconstructed signal is guaranteed. In fact, any online method not ensuring the smoothness of the reconstruction during the signal estimate update (even when the signal estimate is modeled by smooth functions) suffers from this issue, as illustrated in Fig. C.1a) and C.1b). On the other hand, online interpolation methods can be suitable candidates for the task of smooth signal reconstruction with a zero-delay response. These methods use piecewise-defined functions to model a sequence of local signal estimates. Some of these functions allow shaping piecewise-modeled signal estimates that can be updated by assembling a new section (or piece) while guaranteeing the smoothness of the overall sequentially reconstructed signal, as shown in Fig. C.1c). Among them, piecewise polynomial functions, also known as *splines*, are arguably the most representative ones [73, 75]. Actually, splines have been used since ancient times [74], long before their mathematical foundations were even established [72],

presumably because of their approximation capabilities over functions of arbitrary complexity and ease of use.

It should be noted that most recursive signal estimation methods modeling function estimates by splines as a basis expansion [24, 33, 127], suffer from the same issues exposed before. This is mainly because the smoothness of their signal estimate is directly incorporated into the basis representation and not treated as a set of continuity constraints. On the contrary, some works [25, 26] have explored the task of interpolating sequentially streamed data under real-time requirements by means of splines subject to continuity constraints. However, the online methods they use involve a multi-step lookahead or shifting window mechanism, which introduces a delay. Indeed, most online methods for spline interpolation work with local information, for instance, a subset constituted by the last sequentially received data samples. In this case, a delayed response allows them to use a larger subset of sequentially received data samples and correct the signal estimates as long as they are updated within the delay limits. In brief, they can expand the extent of available information at the expense of some delay. On the other hand, and to the best of our knowledge, the only zero-delay spline interpolation method in the literature is the *myopic* approach, referred to as the "classical greedy approach" in [26], which reduces the delay response to zero by totally ignoring any source of forthcoming information, i.e., a purely local method. Clearly, there is a research gap on zero-delay spline interpolation methods exploiting additional nonlocal information to achieve a better reconstruction. This motivates us to work on the research question of whether it is possible to maintain the zero-delay requirement while efficiently using more information than the myopic approach.

In this paper, we answer affirmatively to the above research question by introducing a novel method for zero-delay smoothing[1] spline interpolation that incorporates a priori information about the dynamics of the signal being reconstructed. To this end, we identify the elements of a state space-based sequential decision-making process [128] in the context of zero-delay smoothing spline interpolation. The proposed method relies on a *policy*, i.e., a strategy, that yields a section of the spline (*action*) as a function of the current condition of the so far reconstructed signal and the last received data sample (*state*). Such a policy consists of a differentiable convex optimization layer (DCOL) [14] on top of a recurrent neural network (RNN) [129, 130]. The DCOL allows managing continuity constraints (for any differentiability class) at each interpolation step, thus guaranteeing the smoothness of the signal reconstruction. The RNN assists the signal estimate update mechanism when appending a new spline section by taking into account the effect of each interpolated section on future interpolation steps. This aid comes in the form of global data-driven knowledge, and it is tailored to minimize the global cost of the smoothing interpolation problem, on average. The cost is, in this case, the residual sum of squares plus a weighted derivative-based measure of smoothness. Lastly, our method is *trainable* in the sense that it uses example time series, i.e., time series sampled from the same signal source generating the signal to be reconstructed, to customize the policy to the temporal dependencies (dynamics) of the signal at hand.

The main contributions of this paper can be summarized as follows:

- We rigorously formulate the problem of smoothing spline interpolation from

---

[1]Here, the term smoothing refers to a controlled trade-off between fitting the data samples and proposing a smooth signal estimate.

sequentially streamed data, where each spline section has to be determined as soon as a data sample is available and without having access to subsequent data (zero-delay requirement). Due to its nature, it is formulated as a sequential decision-making problem.

- As opposed to previously proposed (myopic and not trainable) zero-delay methods, our method trains a policy that aims at minimizing the smoothing interpolation cost metric on average. In order to capture the temporal, possibly long-term, dependencies between the streamed data samples and exploit them to reduce further the average cost metric, an RNN able to capture the signal dynamics is incorporated.

- The proposed policy guarantees that the reconstructed signal is smooth (a certain number of derivatives are continuous over the interior of the signal domain). This is achieved by adding a DCOL at the output of the RNN and imposing a set of continuity constraints at each interpolation step. In addition, such a layer admits a closed-form evaluation, resulting in improved computational efficiency with respect to off-the-shelf DCOL libraries.

- We present extensive experimental results that validate our approach over synthetic and real data. Additionally, we show how our approach outperforms the state-of-the-art (namely, myopic) zero-delay methods in terms of the smoothing interpolation average cost metric.

The rest of the paper is structured as follows: Sec. C.2 introduces the notation and presents some basic concepts and definitions. Then, in Sec. C.3, we provide our problem formulation. Next, in Sec. C.4 and Sec. C.5, we respectively provide a solution, a benchmark, and a baseline. Thereafter we experimentally validate our solution in Sec. C.6. Finally, Sec. C.7 concludes the paper.

C.2 Preliminaries

In this section, we present the notation and introduce the type of data used in the paper. Afterward, we address the description of spline-based signal estimates as well as related concepts recurrently appearing in this work. Finally, we formally describe the smoothing spline interpolation problem, which will be used as a starting point for the formulation of our problem.

C.2.1 Notation

Vectors and matrices are denoted by bold lowercase and capital letters, respectively. Given a vector $\boldsymbol{v} = [v_1, \ldots, v_C]^\top$, its cth component is indicated as $[\boldsymbol{v}]_c \triangleq v_c$. Similarly, given a matrix $\boldsymbol{M} \in \mathbb{R}^{R \times C}$, the element in the rth row and cth column is indicated as $[\boldsymbol{M}]_{r,c}$. The notation $[\boldsymbol{v}]_{i:j}$ refers to the sliced vector $[v_i, \ldots, v_j]^\top \in \mathbb{R}^{j-i+1}$. We use Euler's notation for the derivative operator; thus, D_x^k denotes the kth derivative over the variable x.

C.2.2 Problem data

The data considered in this paper consist of discrete time series, or *series* for short, of T terms each. We interchangeably refer to the terms of the series as *observations*. Each tth observation $\boldsymbol{o}_t \in \mathbb{R}^2$ is described by its time stamp $x_t \in \mathbb{R}$ and its value $y_t \in \mathbb{R}$, i.e., $\boldsymbol{o}_t = [x_t, y_t]^\top$. The observation-associated time stamps are set in strictly monotonically increasing order, i.e., $x_{t-1} < x_t$ for all terms in the series. Any two consecutive time stamps define a time section $\mathcal{T}_t = (x_{t-1}, x_t]$. Finally, the initial time stamp x_0 is set by the user.

C.2.3 Spline-based signal estimates

A spline is defined as a piecewise polynomial function. We denote any spline composed of T piecewise-defined functions, or *function sections*, as

$$
f_T(x) = \begin{cases} g_1(x), & \text{if } x_0 < x \le x_1 \\ g_2(x), & \text{if } x_1 < x \le x_2 \\ \vdots \\ g_T(x), & \text{if } x_{T-1} < x \le x_T \end{cases} \tag{C.1}
$$

where every tth function section $g_t : \mathcal{T}_t \to \mathbb{R}$ is a linear combination of polynomials of the form

$$
g_t(x) = \boldsymbol{a}_t^\top \boldsymbol{p}_t(x), \tag{C.2}
$$

with combination coefficients $\boldsymbol{a}_t \in \mathbb{R}^{d+1}$ and basis vector function $\boldsymbol{p}_t : \mathcal{T}_t \to \mathbb{R}^{d+1}$ defined as

$$
\boldsymbol{p}_t(x) = \left[1, (x - x_{t-1}), \dots, (x - x_{t-1})^d \right]^\top, \tag{C.3}
$$

The integer d denotes the *order* of the spline. A spline f_T is said to have a *degree of smoothness* φ if it has φ continuous derivatives over the interior of its domain $\mathrm{dom}(f_T) = \bigcup_{t=1}^T \mathcal{T}_t$. Next, **Proposition 1** shows how to enforce continuity up to degree $\varphi \le d$ in a spline-based signal estimate of order d.

Proposition 1: Given a spline expressed as in (C.1), we can enforce its degree of smoothness to be $\varphi \le d$ by imposing the following equality constraint

$$
[\boldsymbol{a}_t]_{1:\varphi+1} = \boldsymbol{e}_{t-1}, \tag{C.4}
$$

for every $t \in \mathbb{N}^{[1,T]}$, where $\boldsymbol{e}_t \in \mathbb{R}^{\varphi+1}$ is a vector such that each of its elements is computed as

$$
[\boldsymbol{e}_t]_i = \frac{1}{(i-1)!} \sum_{j=1}^{d+1} [\boldsymbol{a}_t]_j \, u_t^{j-i} \prod_{k=1}^{i-1} (j-k), \tag{C.5}
$$

with $u_t \triangleq x_t - x_{t-1}$, and with the exception of $\boldsymbol{e}_0$, which determines the initial conditions of the reconstruction and can be either calculated or set by the user.

   Proof: see Appendix C.8.

C.2.4 Smoothing spline interpolation

Consider the space $\mathcal{W}_\rho$ of functions defined over the domain $(x_0, x_T] \subseteq \mathbb{R}$ with $\rho - 1$ absolutely continuous derivatives and with the ρth derivative square integrable. Then, given a whole series of observations $\{\boldsymbol{o}_t\}_{t=1}^T$ with $T \geq \rho$ and a positive hyperparameter η, we can formulate the following batch optimization problem

$$\min_{f \in \mathcal{W}_\rho} \sum_{t=1}^T (f(x_t) - y_t)^2 + \eta \int_{x_0}^{x_T} (D_x^\rho f(x))^2 \, dx \tag{C.6}$$

known as smoothing spline interpolation [131, 6]. The name is due to the unique solution to the optimization problem (C.6) being a spline conformed of T function sections, as in (C.1). More specifically, the solution of (C.6) is a spline of order $2\rho - 1$ with $2\rho - 2$ continuous derivatives and natural boundary conditions [24]. The hyperparameters η and ρ control the smoothness of such a solution. Particularly, the integer ρ dictates the minimum required degree of smoothness of the search function space $\mathcal{W}_\rho$ and the type of regularization[2] (second term in (C.6)). Regarding η, it controls the trade-off between the squared sum of vertical deviations of the signal estimate from the data and the regularization term. Notice that as $\eta \to 0$, the solution of (C.6) approaches the interpolation spline while as $\eta \to \infty$, it tends to the polynomial of order $\rho - 1$ that best fits the observations in the least-squares sense.

On the other hand, note that the structure of the solution of the problem (C.6), being a natural spline, arises organically rather than being imposed in advance. This is a direct consequence of its batch formulation allowing us to delimit the search space $\mathcal{W}_\rho$ to splines of order d and degree of smoothness φ satisfying $2\rho - 1 \leq d$ and $\rho - 1 \leq \varphi \leq 2\rho - 2$ without loss of optimality. From a practical perspective, it is sufficient to choose the minimum required order and degree of smoothness, thus reducing the model's complexity. However, this trait is not necessarily present in online settings. That is, the smoothness of the solution does not arise naturally using online methods, and it has to be enforced. So here, the choice of the spline order and degree of smoothness is rather user-defined or task-oriented.

C.3 Problem formulation

Once the problem data, the description of spline-based signal estimates, and the smoothing spline interpolation problem have been introduced in Sec. C.2, we are ready to formalize the main task of this paper, namely the trainable zero-delay smoothing spline interpolation problem. This section fully describes the aforementioned task from a data-driven sequential decision-making perspective by introducing a suitable dynamic programming (DP) [46] framework. To this end, we first model the *environment*, define the state space and action space, and delimit a suitable family of candidate policies. Then we introduce the *total cost* and formulate the above task as the problem of finding the policy incurring the lowest total cost on average.

---

[2]Our experimental setup focuses on $\rho = 2$, a common choice in practice, which penalizes excessive curvature in the spline. Applications with $\rho > 2$ can also be found, e.g., trajectory planning tasks [132]. However, they are out of the scope of this paper, as we justify in the ensuing Sec. C.6.2.

C.3.1 Characterization of the problem data

In our problem, the data described in Sec. C.2.2 are observed sequentially. Before every tth time step, the observation about to be received $\boldsymbol{o}_t$ remains undetermined but still governed by the dynamics of the environment. In this work, we model the dynamics of the environment as a random process $Y(\omega, x)$, where ω is a sample point from a sample space Ω, and x is a value within an index set $\mathcal{X} \subseteq \mathbb{R}$, in this case, time. At time x_t, all possible outcomes form a random variable $Y(\omega, x_t)$ or y_t for short. If the mth sample is considered at time x_t, the outcome has a value denoted by $Y(\omega_m, x_t)$ or simply $y_{m,t}$. Consequently, if a discrete set of time stamps is chosen, i.e., $\mathcal{X} = \{x_1, \ldots, x_T\}$, T random variables can be formed, and all the information about the discrete random process $Y_{\mathcal{X}}$ is contained in the joint probability density function $P_{Y_{\mathcal{X}}}$.

C.3.2 State space

At every time step t, we encode a snapshot of the observable environment and the condition of the so-far reconstructed signal in a vector-valued variable called state. With $\mathcal{S}$ denoting the state space, each tth state $\boldsymbol{s}_t \in \mathcal{S}$ is constituted by the corresponding observation $\boldsymbol{o}_t$, and the condition at which the reconstruction was left, which is specified by the vector $\boldsymbol{e}_{t-1}$ whose components are given as in (C.5), and the time instant x_{t-1}. Formally, every tth state $\boldsymbol{s}_t$ is expressed as $\boldsymbol{s}_t = [\boldsymbol{o}_t^\top, \boldsymbol{e}_{t-1}^\top, x_{t-1}]^\top$. Since every state $\boldsymbol{s}_t$ is uniquely determined once the spline coefficients $\boldsymbol{a}_{t-1}$ are fixed, we can explicitly describe the state update mechanism, by means of a deterministic mapping, as

$$\boldsymbol{s}_{t+1} = F(\boldsymbol{s}_t, \boldsymbol{a}_t, \boldsymbol{o}_{t+1}). \tag{C.7}$$

Formalizing the state update mechanism in (C.7) allows us to identify all visitable states seamlessly.

C.3.3 Action space

Immediately after receiving the tth observation, we propose a function section as in (C.2), and we implicitly select the spline coefficients $\boldsymbol{a}_t$. This is because the function section is determined as soon as $\boldsymbol{a}_t$ is chosen (the basis vector defined in (C.3) is given). From this point of view, selecting the spline coefficients of a function section can be understood as an action. Any valid action generates a function section of the same order d as the spline reconstruction. Formally, $\boldsymbol{a}_t \in \mathcal{A} \subseteq \mathbb{R}^{d+1}$ for all the tth terms, where $\mathcal{A}$ denotes the action space. However, if we want a reconstructed spline that is continuous up to the φth derivative, not all valid actions are appropriate. In our context, for any tth action to be deemed *admissible* (or feasible), it must satisfy the constraint in (C.4). Notice that the set of admissible actions depends on the current state. Therefore, we accordingly denote the admissible action space as $\mathcal{A}(\boldsymbol{s}_t)$.

C.3.4 Policy space

A policy $\pi = \{\boldsymbol{\mu}_t : \mathcal{S} \rightarrow \mathcal{A}\}_{t \in \{1,2,\ldots\}}$ consists of a sequence of functions that map states into actions. Policies are more general than actions because they incorporate the knowledge of the state. However, notice that not all policies return admissible

actions. Only the policies that satisfy $\pi(\boldsymbol{s}_t) \in \mathcal{A}(\boldsymbol{s}_t)$ for all time steps are termed *admissible policies*. Separately, *stationary policies* are policies that do not change over time, i.e., $\boldsymbol{\mu} \equiv \boldsymbol{\mu}_t = \boldsymbol{\mu}_{t+1}$ for all time steps. Hence, a stationary policy is unequivocally defined by the mapping $\boldsymbol{\mu}$. Stationary policies are suitable for making decisions in problems with a varying horizon (varying number of time steps), assuming usually stationary environments.

These arguments motivate the use of admissible stationary policies. However, the space of admissible stationary policies is huge, and therefore, the problem of finding the most adequate policy within it can be overwhelmingly complex. Policy approximation techniques help reduce the pool of candidate policies by restricting them to a certain family of policies. These techniques tend to work best (in the sense of providing an adequate policy) when the problem has a clear structure that can be accommodated into the policy. In our case, we aim to incorporate the temporal dependencies across the observations into the policy, as well as the notion of smoothness discussed in Sec. C.2.4. To this end, we resort to parametric policy approximation [45] denoting any approximated stationary policy as $\boldsymbol{\mu}_{\boldsymbol{\theta}}$, where the vector $\boldsymbol{\theta} \in \mathbb{R}^P$ contains the P parameters constituting the aforenamed policy. The set Π of parametric stationary policies that return admissible actions is, therefore, the space of policies of interest to this work.

C.3.5   Total expected cost

The following **Proposition 2** shows that the smoothing spline interpolation objective introduced in Sec. C.2.4, equation (C.6), can be expressed as a summation where each term depends on a single action, resembling the sequence of instantaneous costs in a typical DP formulation.

Proposition 2: The objective of the optimization problem (C.6) can be equivalently computed additively as

$$\sum_{t=1}^{T} \left(\boldsymbol{a}_t^\top \boldsymbol{p}_t(x_t) - y_t \right)^2 + \eta \, \boldsymbol{a}_t^\top \boldsymbol{M}_t \boldsymbol{a}_t, \tag{C.8}$$

where $\boldsymbol{M}_t \in \mathbf{S}_+^{d+1}$ with elements given by

$$[\boldsymbol{M}_t]_{i,j} = \begin{cases} 0 & \text{if } i \leq \rho \text{ or } j \leq \rho \\ \frac{u_t^{i+j-2\rho-1}}{i+j-2\rho-1} \prod_{k=1}^{\rho} (i-k)(j-k) & \text{otherwise,} \end{cases} \tag{C.9}$$

being $u_t \triangleq x_t - x_{t-1}$.

Proof: See Appendix C.9.

Based on **Proposition 2**, we can express the objective of the smoothing spline interpolation problem (C.6), as the total cost $\sum_{t=1}^{T} \kappa(\boldsymbol{s}_t, \boldsymbol{a}_t)$, with *cost* $\kappa : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$ given by

$$\kappa(\boldsymbol{s}_t, \boldsymbol{a}_t) = \left(\boldsymbol{a}_t^\top \boldsymbol{p}_t(x_t) - y_t \right)^2 + \eta \, \boldsymbol{a}_t^\top \boldsymbol{M}_t \boldsymbol{a}_t, \tag{C.10}$$

where $\boldsymbol{M}_t$ is constructed as in (C.9). This is because each tth state-action pair contains all necessary information. From here and under a given policy of interest $\boldsymbol{\mu}_{\boldsymbol{\theta}}$, as described in Sec. C.3.4, the metric

$$\underset{y_t \sim P_{Y_{\mathcal{X}}}}{\mathbb{E}} \left[\sum_{t=1}^{T} \kappa \left(\boldsymbol{s}_t, \boldsymbol{\mu}_{\boldsymbol{\theta}}(\boldsymbol{s}_t) \right) \right], \tag{C.11}$$

denotes the *total expected cost* incurred by following such a policy from a given initial state s_0, and traveling all the remaining states $s_t \in \mathcal{S}_t$ via (C.7). The expectation in (C.11) is performed over the random process modeling the dynamics of the environment through the observations within the states.

C.3.6   Policy search by cost optimization

Computing the expectation in (C.11) is computationally expensive or even intractable when the underlying random process generating the series of observations is unknown. Instead, we can rely on sample average approximation of example series collected from past realizations of the process. The sample average approaches the expectation as the number of examples grows. In this way, we can determine a data-driven policy by solving the following optimization problem

$$\arg\min_{\boldsymbol{\mu_\theta} \in \Pi} \sum_{m=1}^{M} \sum_{t=1}^{T} \kappa\left(s_{m,t}, \boldsymbol{\mu_\theta}(s_{m,t})\right) \tag{C.12a}$$

$$\text{s. to: } s_{m,t} = F\left(s_{m,t-1}, \boldsymbol{\mu_\theta}(s_{m,t-1}), o_{m,t}\right), \forall m, t, \tag{C.12b}$$

$$\boldsymbol{\mu_\theta}(s_{m,t}) \in \mathcal{A}(s_{m,t}), \forall m, t, \tag{C.12c}$$

where the integer M denotes the number of example series, indexed by m, and where all the initial states $s_{m,0}$ as well as all observations $o_{m,t}$ are given.

C.4   Proposed solution

The previous Sec. C.3 has provided the necessary definitions and considerations to arrive at a rigorous problem formulation. An exact solution to the problem (C.12) is probably impossible to obtain in practice, mainly due to the complexity of the search space Π. There are multiple possibilities regarding the policy approximation and optimization techniques that can be taken towards obtaining a near-optimal solution to (C.12). This section presents a specific set of design choices based on the current state-of-the-art. In particular, we rely on a policy parametrization through cost parametrization technique, borrowed from the DP literature, in synergy with an RNN architecture. Then we make use of backpropagation through time (BPTT) [48], a gradient computation technique borrowed from the deep learning literature [47]. Our proposed solution can effectively solve the problem formulated in Sec. C.3.6 for $\rho \leq 2$. The remaining configurations manifest instability issues, and even though they may be solvable, they lie outside of the scope of the current paper as further discussed in the following Sec. C.6.2.

Future developments in the DP or deep learning areas, such as new policy approximation approaches, neural architectures, or optimizers, can possibly render the techniques proposed in this section obsolete but will not affect the validity of the problem formulated in Sec. C.3.6.

C.4.1   Policy form

Parametric policy approximation via parametric cost function approximation (CFA) [45] is a method that seeks through the policy space, in our case Π, among those

Figure C.2: Visual representation of our RNN architecture satisfying the relation in (C.15). The elements onto the gray shaded area constitute the mapping $R_{\theta'}$. The CAT cell stands for a concatenation operation.

policies defined as an optimization problem with parametrized objectives. In this work, we are interested in CFA-based policies of the form

$$\boldsymbol{\mu}_{\boldsymbol{\theta}}(\boldsymbol{s}_{m,t}) = \arg\min_{\boldsymbol{a} \in \mathcal{A}(\boldsymbol{s}_{m,t})} \left\{ \kappa(\boldsymbol{s}_{m,t}, \boldsymbol{a}) + J_{\boldsymbol{\theta}}(\boldsymbol{s}_{m,t}, \boldsymbol{a}; \boldsymbol{h}_{m,t}) \right\}, \tag{C.13}$$

where the map κ denotes the cost described in (C.10), and the mapping $J_{\boldsymbol{\theta}} : \mathcal{S} \times \mathcal{A} \times \mathbb{R}^H \to \mathbb{R}$ is a parametric *cost-to-go* approximation involving P parameters contained in the vector $\boldsymbol{\theta}$. Regarding the vector $\boldsymbol{h}_{m,t} \in \mathbb{R}^H$, it represents a *latent state* value at the tth time step of an mth example series. The latent state may encode relevant information from past observations and can be viewed as a policy memory [133, 134, 135].

We aim for a cost-to-go approximation $J_{\boldsymbol{\theta}}$, which penalizes those actions that are distant from the output of a certain RNN. The main reason behind this approach is that an RNN that successfully captures the temporal dynamics of the environment has the potential to pull towards actions that yield a low expected total cost. So, it is constructed as follows

$$J_{\boldsymbol{\theta}}(\boldsymbol{s}_{m,t}, \boldsymbol{a}; \boldsymbol{h}_{m,t}) = \lambda \left\| \boldsymbol{a} - \begin{bmatrix} \boldsymbol{0}_{\varphi+1} \\ \boldsymbol{r}_{m,t} \end{bmatrix} \right\|_2^2, \tag{C.14}$$

where $\lambda \in \mathbb{R}_+$. The vectors $\boldsymbol{r}_{m,t} \in \mathbb{R}^{d-\varphi}$, and $\boldsymbol{h}_{m,t} \in \mathbb{R}^H$ represent the outputs and latent state of an RNN, $R_{\boldsymbol{\theta}'} : \mathcal{S}_t \times \mathbb{R}^H \to \mathbb{R}^{d-\varphi} \times \mathbb{R}^H$, respectively. They are obtained from the following relation

$$R_{\boldsymbol{\theta}'}(\boldsymbol{s}_{m,t}; \boldsymbol{h}_{m,t}) = \begin{bmatrix} \boldsymbol{r}_{m,t} \\ \boldsymbol{h}_{m,t+1} \end{bmatrix}, \tag{C.15}$$

with $\boldsymbol{\theta} = [\lambda, {\boldsymbol{\theta}'}^\top]^\top$ and $\boldsymbol{\theta}' \in \mathbb{R}^{P-1}$, exemplified in Fig. C.2. From now on, we refer to the policy in (C.13) with cost-to-go as in (C.14) as the RNN-based policy. Finally, notice that besides being parametric, the RNN-based policy is admissible and stationary by design.

Figure C.3: Rolled representation of the environment-agent system. The sampler cell samples the random process modeling the dynamics of the environment. The interpolator cell performs the reconstruction, evaluates the cost, and updates the states. The agent cell contains the policy and the RNN. The blue-shaded area encompasses the environment. The cost (in gray) is used during the training phase but not for evaluation.

C.4.2   Policy evaluation

Evaluating the proposed policy involves solving the optimization problem stated in (C.13). Notice that both the cost κ built as in (C.10), and the cost-to-go approximation $J_{\boldsymbol{\theta}}$ described in (C.14), are convex with respect to the actions and hence, the objective in (C.13) is convex too. Moreover, the admissible action set, described in Sec. C.3.3, is convex. Therefore, the optimization problem in (C.13) is convex thus, any locally optimal action is globally optimal [113].

Additionally, the optimization problem in (C.13) has been designed to admit a closed-form solution. Closed-form evaluations can usually be computed faster and more precisely than solutions obtained from numerical methods, and thus, they are more suitable under zero-delay requirements. See Appendix C.10 for the derivation of the closed-form evaluation.

C.4.3   Policy training

As explained in Sec. C.4.1, we have reduced the search space of problem (C.12) by restricting the policy space to a family of policies of the form given in (C.13). Specifically, from searching a function $\boldsymbol{\mu_\theta}$ in the function space Π, we have narrowed the problem down to that of finding a vector $\boldsymbol{\theta}$ in the vector space $\mathbb{R}^P$. In fact, tuning the proposed policy parameters by solving the optimization problem (C.12) is commonly referred to as policy training. Unfortunately, the objective (C.12a) is non-convex with respect to the parameters in $\boldsymbol{\theta}$. As a reasonable solution, we rely on a gradient-based optimizer aiming to converge to a high-performance local minimum.

From a deep learning perspective, the policy evaluation presented in Sec. C.4.2 can be understood as a forward pass of a DCOL on top of an RNN, and hence, it is trained using BPTT via automatic differentiation [136]. This point of view is schematized in Fig. C.3, where traveling the given mth series, by following a policy $\boldsymbol{\mu_\theta}$, allows to construct the cumulative objective in (C.12a) used for training. Ad-

ditionally, and thanks to the closed-form policy evaluation discussed in Sec. C.4.2, computing and propagating the gradient of the tth action $\boldsymbol{a}_t$ with respect to the parameters contained in $\boldsymbol{\theta}$ is done avoiding the need of unrolling numerical optimizers [137] or using specific numerical tools for DCOLs such as CVXPY Layers [14].

C.5   Benchmark and baseline methods

Recall from Sec. C.2.4 that the batch formulation provides the optimal reconstruction with hindsight. The batch solution can be found by solving the optimization problem (C.6), but only once all time-series data are available. Thus, it cannot be used for zero-delay interpolation. Conceptually, online methods achieve a zero-delay response at the expense of incurring higher or equal loss than the batch solution. For this reason, the batch solution is used here as a baseline.

On the other hand, as stated in the Introduction and to the best of our knowledge, there is no related work to our trainable zero-delay smoothing interpolation approach in the literature. One could consider that the closest approach is the interpolation method known as myopic. This is a local method in the sense that it only focuses on the last received data sample while completely ignoring the distribution of future arriving data. For this reason, the myopic method is used here as a benchmark. In this sense, our proposed method must outperform the myopic method to be deemed acceptable.

C.5.1   Myopic benchmark

A policy that chooses the action that minimizes the current or instantaneous cost is commonly referred to as myopic. It can be constructed as

$$\boldsymbol{\mu}(\boldsymbol{s}_t) = \arg \min_{\boldsymbol{a} \in \mathcal{A}(\boldsymbol{s}_t)} \left\{ \kappa(\boldsymbol{s}_t, \boldsymbol{a}) \right\}, \tag{C.16}$$

with cost κ as in (C.10) and admissible action set as described in Sec. C.3.3. Notice that since the myopic policy does not contain trainable parameters, it does not need to be trained. Moreover, the myopic approach is carried out as a parameterless CFA-based policy, hence, becoming a particular case of (C.13). For this reason, it also admits a unique and closed-form evaluation. See Appendix C.10 for more details.

C.6   Experiments

In this section, we experimentally validate the effectiveness of the proposed RNN-based policy, introduced in Sec. C.4. To this end, we first describe the time-series datasets used. Then, we outline the possible policy configurations, i.e., the possible types of splines as well as the RNN architecture. Afterward, we report how the experiments have been carried out. Finally, we present and comment on the experimental results.

(a) (4,2), $\eta = 10$, and synthetic dataset.

(b) (3,1), $\eta = 1$, and R1.

(c) (4,2), $\eta = 10$, and R2.

(d) (4,2), $\eta = 1$, and R3.

(e) (4,2), $\eta = 1$, and R4.

(f) (3,1), $\eta = 10$, and R5.

Figure C.4: Some of the training-validation curves for the considered RNN-based policy configurations (d, φ) and η values for each dataset. The legends (T) and (V) refer to the training and validation partitions, respectively. The loss metric is the average total cost per function section (see Sec. C.3.5). The shaded areas represent one standard deviation.

Table C.1: Performance metrics averaged over the test partitions, for different η values and policy configurations with $\rho = 2$. Recall that the improvement metric (C.17) reports the gain of any RNN(d, φ) configuration over its corresponding Myopic(d, φ) (as a benchmark) and batch(d, φ) (as the baseline) configurations.

Dataset	Configuration	$\eta = 0.1$			$\eta = 1$			$\eta = 10$		
		MSE	MAE	Improvement	MSE	MAE	Improvement	MSE	MAE	Improvement
Synthetic	Myopic(3, 1)	0.49	0.60		0.31	0.46		0.64	0.63	
	Batch(3, 1)	0.41	0.56	$71.5\% \pm 2.8\%$	0.26	0.44	$51.1\% \pm 3.9\%$	0.22	0.39	$78.1\% \pm 2.5\%$
	RNN(3, 1)	0.29	0.47		0.25	0.42		0.28	0.43	
	Myopic(4, 2)	0.50	0.60		0.38	0.50		0.83	0.71	
	Batch(4, 2)	0.41	0.56	$64.8\% \pm 2.6\%$	0.26	0.44	$53.0\% \pm 3.6\%$	0.22	0.39	$81.0\% \pm 2.6\%$
	RNN(4, 2)	0.28	0.44		0.27	0.43		0.30	0.44	
R1	Myopic(3, 1)	0.09	0.13		0.20	0.22		0.44	0.40	
	Batch(3, 1)	0.06	0.11	$59.4\% \pm 20.3\%$	0.07	0.12	$81.3\% \pm 8.2\%$	0.10	0.16	$75.5\% \pm 8.2\%$
	RNN(3, 1)	0.07	0.12		0.10	0.15		0.19	0.24	
	Myopic(4, 2)	0.13	0.16		0.29	0.27		0.51	0.43	
	Batch(4, 2)	0.06	0.11	$59.6\% \pm 17.0\%$	0.07	0.12	$78.6\% \pm 8.6\%$	0.10	0.16	$72.9\% \pm 8.5\%$
	RNN(4, 2)	0.08	0.13		0.12	0.17		0.23	0.27	
R2	Myopic(3, 1)	0.42	0.50		0.40	0.48		0.32	0.42	
	Batch(3, 1)	0.27	0.40	$67.3\% \pm 6.4\%$	0.26	0.39	$52.2\% \pm 7.1\%$	0.23	0.37	$68.8\% \pm 5.7\%$
	RNN(3, 1)	0.23	0.36		0.23	0.37		0.20	0.34	
	Myopic(4, 2)	0.65	0.62		0.58	0.59		0.39	0.47	
	Batch(4, 2)	0.27	0.40	$70.1\% \pm 3.9\%$	0.26	0.39	$52.0\% \pm 5.4\%$	0.23	0.37	$73.5\% \pm 4.0\%$
	RNN(4, 2)	0.21	0.35		0.22	0.36		0.21	0.35	
R3	Myopic(3, 1)	4.60	1.80		2.90	1.40		1.80	1.09	
	Batch(3, 1)	3.10	1.40	$76.3\% \pm 4.4\%$	2.35	1.30	$81.3\% \pm 5.3\%$	1.66	1.06	$65.6\% \pm 11.6\%$
	RNN(3, 1)	2.56	1.30		1.56	1.02		1.40	0.96	
	Myopic(4, 2)	6.30	2.10		3.02	1.43		1.89	1.10	
	Batch(4, 2)	3.10	1.40	$69.7\% \pm 4.4\%$	2.36	1.26	$88.9\% \pm 3.9\%$	1.66	1.06	$80.8\% \pm 9.0\%$
	RNN(4, 2)	1.80	1.00		1.43	0.95		1.52	1.00	
R4	Myopic(3, 1)	4.5e-3	2.8e-2		4.1e-3	2.7e-2		3.6e-3	2.5e-2	
	Batch(3, 1)	2.7e-3	2.3e-2	$13.7\% \pm 57.3\%$	2.6e-3	2.2e-2	$-61.8\% \pm 94.4\%$	2.4e-3	2.0e-2	$16.5\% \pm 54.5\%$
	RNN(3, 1)	7.1e-3	6.8e-2		6.0e-3	5.6e-2		3.2e-3	3.1e-2	
	Myopic(4, 2)	6.9e-3	3.6e-2		5.7e-3	3.2e-2		4.9e-3	3.0e-2	
	Batch(4, 2)	2.7e-3	2.3e-2	$-35.5\% \pm 78.5\%$	2.6e-3	2.2e-2	$65.7\% \pm 25.5\%$	2.4e-2	2.0e-2	$19.5\% \pm 51.5\%$
	RNN(4, 2)	3.7e-3	3.3e-2		4.1e-3	3.1e-2		3.5e-3	2.7e-2	
R5	Myopic(3, 1)	2.7e-4	9.3e-3		6.0e-4	1.3e-2		7.1e-3	3.8e-2	
	Batch(3, 1)	1.8e-4	7.8e-3	$-55.4\% \pm 101.6\%$	1.7e-4	7.7e-3	$26.3\% \pm 59.8\%$	2.6e-4	9.1e-3	$88.2\% \pm 7.1\%$
	RNN(3, 1)	2.3e-4	8.7e-3		3.1e-4	1.0e-2		3.2e-3	3.0e-2	
	Myopic(4, 2)	4.7e-4	1.1e-2		9.7e-4	1.7e-2		1.1e-2	4.8e-2	
	Batch(4, 2)	1.8e-4	7.8e-3	$-13.3\% \pm 68.2\%$	1.7e-4	7.7e-3	$39.8\% \pm 41.9\%$	2.6e-4	9.1e-3	$85.1\% \pm 7.5\%$
	RNN(4, 2)	2.9e-4	9.7e-3		6.1e-4	1.4e-2		5.4e-3	3.6e-2	

C.6.1 Problem data description

For these experiments, we use a synthetic dataset and five real datasets. Each dataset consists of a time series of 28800 signal samples which has been split into 288 series of 100 samples each, except for the first real dataset which contains 57600 samples split into 576 series.

The synthetic dataset is first generated as a uniformly arranged realization of a given autoregressive process AR(2) with white Gaussian noise $\mathcal{N}(0, 0.1)$. Then, the resulting series is compressed via PI [138] with CompDev = 0.1, CompMax = ∞ and CompMin = 0. As a result, the series time stamps are not uniformly distributed anymore.

The first real dataset (R1) consists of a series of household minute-averaged active power consumption (in Kilowatts) [139]. The second real dataset (R2) is a quantized and PI-compressed (and hence not uniformly sampled) time series measuring an oil separation deposit pressure[3] (in Bar). For the third real dataset (R3) [140], a cooling fan with weights on its blades is used to generate vibrations which are recorded by an attached accelerometer. The vibration samples are recorded every 20 milliseconds. We use the accelerometer recorded x-values (which are standardized) for the rotation speeds ranging from 5 to 40 rpm. The fourth real dataset (R4) [141] monitors the skin temperature (in Celsius degrees) of a volunteer subject through a wearable device every 4 minutes. The fifth and last real dataset (R5) [142] consists of a sensor within a sensor network deployed in a lab, collecting the temperature-corrected relative humidity in percentage. The sampling rate is non-uniform and ranges from deciseconds to tens of seconds. Finally, it is worth mentioning that the datasets R4 and R5 contain gaps (several orders of magnitude wider than the average sampling period) of missing data that we have shortened to avoid instability in the reconstruction. In similar cases where the available raw data is of low quality, thorough and task-specific data preprocessing techniques are assumed. This can improve the performance results as described in the ensuing Sec. C.6.4.

C.6.2 Policy configuration

We experimentally observe that the myopic policy described in Sec. C.5.1 is not stable for values of $\rho > 2$. Recall that the value of ρ affects the policy cost, set as in (C.10), and delimits the order and degree of smoothness of the spline signal estimate, as explained in Sec. C.2.3. We also observe instability under the myopic policy for $\rho = 2$ with a spline signal estimate of order $d = 3$ and degree of smoothness $\varphi = 2$. Consequently, our proposed RNN-based policy is unstable for the same ρ values and spline configurations since it implicitly uses the myopic policy as a guided starting point. This can be seen by comparing (C.16) and (C.13) with a near-zero initial value of λ. Although further theoretical instability studies, alternative policy architectures, or low-delay approaches can contribute to solving the instability issue, they lie outside of the scope of this paper. Nonetheless, we have maintained the general problem formulation as a starting point for future works to take over. On the other hand, the interpolation problem with $\rho = 1$ is not interesting since it leads to linear interpolation. Therefore, in the present work, we focus on the smoothing interpolation problem with $\rho = 2$ and with the remaining stable spline

---

[3]Data collected from Lundin's offshore oil and gas platform Edvard-Grieg.

configurations, within the search function space described in Sec. C.2.4, which can lead to optimal reconstructions. Those spline configurations, hereinafter specified by the shorthand notation (d, φ) of the order and degree of smoothness of the spline, correspond to (3,1) and (4,2). Accordingly, the notation Myopic(d, φ) or RNN(d, φ) refers to the type of policy besides the spline configuration.

Regarding the RNN architecture shaping the approximated cost-to-go within the RNN-based policy, introduced in Sec. C.4.1 and illustrated in Fig. C.2, we set a preprocessing step that forwards the time length, i.e. $u_t = x_t - x_{t-1}$, of the tth time section $\mathcal{T}_t$, instead of directly using the time stamps. This preprocessing step makes the architecture invariant to time shifts in the set of time stamps. In our experiments, the recurrent unit consists of two stacked gated recurrent unit (GRU) layers [51], with a latent state (hidden state) of size 16 and an input of size 16. The input and output layers are set as linear layers to match the required dimensionality, i.e., to match the input size after the preprocessing step and to match the order of the spline minus the number of constrained coefficients as output size.

C.6.3   Experimental setup

The datasets are randomly divided into 192 series for training, 64 for validation, and 32 for testing. Except for the R1 dataset, which has been divided in the same proportion but in relation to its data size. The benefits of this train-validation-test partition are two-fold: i) the policy becomes more robust against unknown initial conditions, and ii) we can validate the reconstruction against an optimal batch solution (shorter sequences are computationally tractable using batch optimization). All series within a dataset are standardized for implementation convenience. To avoid data leaking, the mean and standard deviation of their respective training partition are used for the standardization. In other words, we compute the mean and variance of the training partition and assume them to be the moments of the true data distribution. The standardization of series is useful to enforce the RNN unit to focus on the fluctuations of the signal values rather than on their magnitude. Finally, the RNN-based policy has been trained using the adaptive moments (Adam) optimizer [143], with $\beta_1 = 0.9$, $\beta_2 = 0.999$, without weight decay, and a learning rate of 0.001 over mini-batches of 32 time series each (double mini-batch size in the case of R1).

C.6.4   Results and discussion

Some of the training-validation curves are presented in Fig. C.4. As expected, we observe that randomly initialized RNN-based policies (except for the parameter λ, which controls the length of the initial performance gap, as discussed in Sec. C.6.2, and is manually initialized) only outperform the myopic policy after training. We also observe wider (in relative terms) standard deviations in those datasets with more abrupt changes, either from the nature of the data, as in R1, or due to missing data and posterior preprocessing, as in the case of R4 and R5. This phenomenon appears also to be caused by highly non-uniform sampling rates, as in R5. But in this case, the width seems to decrease as the policy yields more accurate estimates. This implies that the RNN-based policy is able to learn how to adapt under non-uniform sampling rates properly.

Figure C.5: Training curves of the parameter $\lambda \in \mathbb{R}_+$ introduced in (C.14). The elements displayed coincide with those shown in Fig. C.4.

Figure C.6: Box plot of the policy execution time per interpolation step over the test partitions and η values $\{0.1, 1, 10\}$. It illustrates (excluding outliers) the minimum, first quartile, median, third quartile, and maximum.

Once the RNN-based policy has been trained, we measure its performance with respect to the myopic policy (as the benchmark) and the batch reconstruction (as the baseline) through an *improvement* metric defined as

$$I = \frac{\ell_M - \ell_R}{\ell_M - \ell_B}, \tag{C.17}$$

where ℓ_M, ℓ_R, ℓ_B denote the loss metric displayed in Fig. C.4 but over the test partition for the myopic, the RNN-based policies and the batch solution, respectively. Acceptable performances yield improvement values in $(0, 1]$, being $I = 1$ the best possible value, whereas nonpositive improvement values indicate a deficient performance. The standard deviation of the improvement metric is then estimated through error propagation, i.e.,

$$\sigma_I = \sqrt{\left(\frac{\partial I}{\partial \ell_M}\right)^2 \sigma_M^2 + \left(\frac{\partial I}{\partial \ell_R}\right)^2 \sigma_R^2 + \left(\frac{\partial I}{\partial \ell_B}\right)^2 \sigma_B^2}, \tag{C.18}$$

with σ_M, σ_R, and σ_B denoting the standard deviation of the respective loss metrics over the test partition. The improvement results are summarized in Table C.1. From Fig. C.4 and Table C.1, it can be observed that the policy configurations with the highest improvement scores over each of the considered dataset test partitions are in agreement with their corresponding validation curves. Table C.1 also shows standard performance descriptors such as the mean squared error (MSE) and mean absolute error (MAE). See Appendix C.11 for their computation. Note that for most of the experiments that we have carried out, the RNN-based policy outperforms, in terms of the MSE and MAE metrics, the myopic policy while it falls behind the batch policy. This observation experimentally justifies the smoothness assumption in our formulation.

Regarding the parameter $\lambda \in \mathbb{R}_+$ introduced in (C.14), it can be understood as the confidence of the RNN-based policy in its ability to foresee incoming data samples. In this way, it also quantifies the importance of the RNN architecture (detailed in Fig. C.2) in the reconstruction task. As an illustration, Fig. C.5 shows the training curves corresponding to the parameter λ for the policy configurations presented in Fig. C.4.

On the other hand, we observe a competitive performance in terms of the execution time of the RNN-based policy evaluation (forward pass) as compared to their myopic counterpart. Our evaluation time results are summarized in Fig. C.6, where the policies are implemented in Python 3.8.8. and the experiment is done in a 2018 laptop with a 2.7 GHz Quad-Core Intel Core i7 processor and 16 GB 2133 MHz LPDDR3 memory. Regarding memory complexity, the myopic policy is parameterless (see Sec. C.5.1), and our configuration of the RNN-based policy (see Sec. C.6.2) contains approximately 3400 trainable parameters, which is arguably a reduced model size for most tasks.

Finally, and for the sake of completeness, Fig. C.7 shows a snapshot of a zero-delay smooth signal reconstruction alongside its two first derivatives using our proposed method.

Figure C.7: Snapshot of a reconstructed time series from the test partition of the synthetic dataset. Both of the presented policies have been trained over the training partition before reconstruction.

C.7 Conclusion

In this paper, we propose a method for zero-delay smoothing spline interpolation. Our method relies on a parametric policy, named the RNN-based policy, specifically engineered for the zero-delay interpolation task. As new data samples arrive, this policy yields piecewise polynomial functions used for smooth signal reconstruction. Our experiments show that the RNN-based policy can learn the dynamics of the target signal and efficiently incorporate them (in terms of improved accuracy and reduced response time) into the reconstruction task.

This work can be seen as a proof of concept with several immediate follow-ups. The flexibility in our policy design allows extending this work to multivariate time series with a moderate increase in complexity. It is also possible to generalize the problem data, e.g., quantization intervals instead of data points, as well as to accommodate additional constraints as long as the convexity of the policy evaluation problem is preserved. Lastly, we notice that our work provides the foundation and can be tailored effectively for reconstructing non-stationary signals by borrowing reinforcement learning techniques.

C.8 Proof of Proposition 1

Recall from Sec. C.2.3 that every spline f_T, as in (C.1), is composed of T function sections and $T - 1$ contact points. We say that two consecutive function sections have a contact of order φ if they have φ equal derivatives at the contact point. Then, guaranteeing a degree of smoothness φ for a given spline f_T is equivalent to ensuring that all its contact points are at least of order φ since every tth function section g_t,

as in (C.2), is already smooth over the interior of its domain $\mathcal{T}_t$. In practice, this can be ensured by imposing the following equality constraints

$$\lim_{x \to x_{t-1}^-} D_x^k\, g_{t-1}(x) = \lim_{x \to x_{t-1}^+} D_x^k\, g_t(x), \tag{C.19}$$

for every $k \in \mathbb{N}^{[0,\varphi]}$ and $t \in \mathbb{N}^{[2,T]}$. From here, notice that the kth derivative of every tth function section g_t can be computed as

$$D_x^k\, g_t(x) = \boldsymbol{a}_t^\top \left[D_x^k\, [\boldsymbol{p}_t(x)]_1, \ldots, D_x^k\, [\boldsymbol{p}_t(x)]_{d+1} \right]^\top. \tag{C.20}$$

Also, notice from the definition in (C.3) that the ith component of the tth basis vector function $\boldsymbol{p}_t$ equals

$$[\boldsymbol{p}_t(x)]_i = (x - x_{t-1})^{i-1}, \tag{C.21}$$

for all $i \in [1, d+1]$. From this point, the kth derivative of each ith component of the basis vector function $\boldsymbol{p}_t$ can be straightforwardly computed as

$$D_x^k\, [\boldsymbol{p}_t(x)]_i = (x - x_{t-1})^{i-1-k} \prod_{j=1}^k (i - j). \tag{C.22}$$

Now observe that

$$\lim_{x \to x_{t-1}^+} D_x^k\, [\boldsymbol{p}_t(x)]_i = \begin{cases} k! & \text{if } i = k+1, \\ 0 & \text{otherwise.} \end{cases} \tag{C.23}$$

Therefore, from the relations in (C.20) and (C.23), the right hand term in (C.19) can be equivalently computed as

$$\lim_{x \to x_{t-1}^+} D_x^k\, g_t(x) = \sum_{i=1}^{d+1} [\boldsymbol{a}_t]_i \lim_{x \to x_{t-1}^+} D_x^k\, [\boldsymbol{p}_t(x)]_i \tag{C.24a}$$

$$= k!\, [\boldsymbol{a}_t]_{k+1}. \tag{C.24b}$$

Separately, we can define a vector $\boldsymbol{e}_t \in \mathbb{R}^{\varphi+1}$ whose components are constructed as

$$[\boldsymbol{e}_t]_{k+1} \triangleq \frac{1}{k!} \lim_{x \to x_t^-} D_x^k\, g_t(x) \tag{C.25a}$$

$$= \frac{1}{k!} \sum_{i=1}^{d+1} [\boldsymbol{a}_t]_i\, u_t^{i-1-k} \prod_{j=1}^k (i - j), \tag{C.25b}$$

for every $k \in \mathbb{N}^{[0,\varphi]}$ and $t \in \mathbb{N}^{[1,T]}$ with $u_t \triangleq x_t - x_{t-1}$, and where the step (C.25b) uses the relations described in (C.20) and (C.22). On the other hand, $\boldsymbol{e}_0$ encodes the initial boundary conditions of the reconstruction and can be set by the user in advance or calculated. Finally, by dividing both sides of the equality constraint in (C.19) by $k!$, using the relations derived in (C.24) and (C.25), and appropriately renaming the indices we obtain the **Proposition 1**.

C.9 Proof of Proposition 2

Recall from Sec. C.2.4 that the solution to the optimization problem stated in (C.6) is a spline function in $\mathcal{W}_\rho$. This fact allows us to reduce the search function space without loss of optimality. In fact, we can incorporate the spline form of the solution into the objective functional as far as we ensure the required minimum degree of smoothness of the solution, for example, via (C.19). From here, we can equivalently compute the regularization term in the objective in (C.6) (second term) as

$$\int_{\bigcup_{t=1}^{T} \mathcal{T}_t} \left(D_x^\rho f_T(x)\right)^2 dx = \sum_{t=1}^{T} \int_{\mathcal{T}_t} \left(D_x^\rho g_t(x)\right)^2 dx. \tag{C.26}$$

Separately, and making use of the definition of function section in (C.2), we obtain the following relation

$$\int_{\mathcal{T}_t} \left(D_x^\rho g_t(x)\right)^2 dx = \int_{\mathcal{T}_t} \left(D_x^\rho \boldsymbol{a}_t^\top \boldsymbol{p}_t(x)\right)^2 dx \tag{C.27a}$$

$$= \boldsymbol{a}_t^\top \boldsymbol{M}_t \boldsymbol{a}_t, \tag{C.27b}$$

with

$$[\boldsymbol{M}_t]_{i,j} = \int_{\mathcal{T}_t} D_x^\rho \left[\boldsymbol{p}_t(x)\right]_i \, D_x^\rho \left[\boldsymbol{p}_t(x)\right]_j \, dx. \tag{C.28}$$

From the relation in (C.22), it is clear that the first ρ rows and columns of the matrix defined in (C.28) are zero valued. Then, we can compute the rest of the elements in the matrix $\boldsymbol{M}_t \in \boldsymbol{S}_+^{d+1}$ as follows

$$[\boldsymbol{M}_t]_{i,j} = \prod_{k=1}^{\rho} (i-k)(j-k) \int_{x_{t-1}}^{x_t} (x - x_{t-1})^{i+j-2(\rho+1)} dx \tag{C.29a}$$

$$= \frac{(x_t - x_{t-1})^{i+j-2\rho-1}}{i+j-2\rho-1} \prod_{k=1}^{\rho} (i-k)(j-k). \tag{C.29b}$$

On the other hand, the sum of squared residuals in the objective in (C.6) (first term) can be equivalently computed as

$$\sum_{t=1}^{T} \left(f_T(x_t) - y_t\right)^2 = \sum_{t=1}^{T} \left(g_t(x_t) - y_t\right)^2, \tag{C.30}$$

from the definition of spline, see relation (C.1).

Summing up, the result stated in **Proposition 2** can be reached starting from the objective in (C.6) then following the relations in (C.26), (C.27) alongside (C.29) and (C.30).

C.10 Closed-form policy evaluation

Notice that both the proposed policy in (C.13) and the myopic policy in (C.16) can be equivalently evaluated by solving the following quadratic convex problem

$$\boldsymbol{\mu}(\boldsymbol{s}_t) = \arg \min_{\boldsymbol{a} \in \mathcal{A}(\boldsymbol{s}_t)} \left\{\boldsymbol{a}^\top \boldsymbol{A}_t \boldsymbol{a} + \boldsymbol{b}_t^\top \boldsymbol{a}\right\}, \tag{C.31}$$

Table C.2: Terms in (C.31). The notation is shared with the rest of the paper with the incorporation of $\boldsymbol{P}_t \triangleq \boldsymbol{p}_t(x_t)\boldsymbol{p}_t(x_t)^\top$ and $\boldsymbol{v}_t \triangleq [\boldsymbol{0}_{\varphi+1}^\top, \boldsymbol{r}_t{}^\top]^\top$.

	$\boldsymbol{A}_t$	$\boldsymbol{b}_t$
Myopic	$\boldsymbol{P}_t + \eta\boldsymbol{M}_t$	$-2y_t\boldsymbol{p}_t(x_t)$
RNN	$\boldsymbol{P}_t + \eta\boldsymbol{M}_t + \lambda\boldsymbol{I}_{d+1}$	$-2(y_t\boldsymbol{p}_t(x_t) + \lambda\boldsymbol{v}_t)$

where the terms $\boldsymbol{A}_t \in \boldsymbol{S}_+^{d+1}$ and $\boldsymbol{b}_t \in \mathbb{R}^{d+1}$ take different values for the different policy variations as described in the Table C.2. The form in (C.31) is displayed as an intermediate step for the sake of clarity, and the dependencies with example time series (indexed by m) and the policy parameters (contained in $\boldsymbol{\theta}$) have been omitted for the sake of notation. Then, we relocate the equality constraints (presented in (C.4) and satisfied by the actions in the admissible set $\mathcal{A}(\boldsymbol{s}_t)$) in the objective of (C.31), by restating

$$\boldsymbol{a} = \begin{bmatrix} \boldsymbol{e}_{t-1} \\ \boldsymbol{0}_{d-\varphi} \end{bmatrix} + \begin{bmatrix} \boldsymbol{0}_{\varphi+1} \\ \boldsymbol{\alpha} \end{bmatrix}, \tag{C.32}$$

or equivalently, by setting $\boldsymbol{a} = \boldsymbol{B}_1\boldsymbol{e}_{t-1} + \boldsymbol{B}_2\boldsymbol{\alpha}$ where the components of $\boldsymbol{e}_{t-1}$ are described in (C.5), the matrices $\boldsymbol{B}_1 \triangleq [\boldsymbol{I}_{\varphi+1}, \boldsymbol{0}_{(\varphi+1)\times(d-\varphi)}]^\top \in \mathbb{R}^{(d+1)\times(\varphi+1)}$ and $\boldsymbol{B}_2 \triangleq [\boldsymbol{0}_{(d-\varphi)\times(\varphi+1)}, \boldsymbol{I}_{d-\varphi}]^\top \in \mathbb{R}^{(d+1)\times(d-\varphi)}$ are defined for the sake of notation, and where $\boldsymbol{\alpha} \in \mathbb{R}^{d-\varphi}$. After some algebraic steps, both policies can be equivalently evaluated as

$$\boldsymbol{\mu}(\boldsymbol{s}_t) = \begin{bmatrix} \boldsymbol{e}_{t-1} \\ \boldsymbol{\alpha}_t \end{bmatrix}, \tag{C.33}$$

where the vector $\boldsymbol{\alpha}_t \in \mathbb{R}^{d-\varphi}$ is obtained from

$$\boldsymbol{\alpha}_t = \arg\min_{\boldsymbol{a}\in\mathbb{R}^{d-\varphi}} \{\boldsymbol{\alpha}^\top \boldsymbol{B}_2^\top \boldsymbol{A}_t \boldsymbol{B}_2 \boldsymbol{\alpha} \tag{C.34a}$$

$$+ 2\left(\boldsymbol{e}_{t-1}^\top \boldsymbol{B}_1^\top \boldsymbol{A}_t \boldsymbol{B}_2 + \boldsymbol{b}_t^\top \boldsymbol{B}_2\right)\boldsymbol{\alpha}\}, \tag{C.34b}$$

with closed-form solution given by

$$\boldsymbol{\alpha}_t = -\left(\boldsymbol{B}_2^\top \boldsymbol{A}_t \boldsymbol{B}_2\right)^{-1}\left(\boldsymbol{B}_2\boldsymbol{A}_t\boldsymbol{B}_1\boldsymbol{e}_{t-1} + \frac{1}{2}\boldsymbol{B}_2^\top \boldsymbol{b}_t\right), \tag{C.35}$$

being $\boldsymbol{B}_2^\top \boldsymbol{A}_t \boldsymbol{B}_2 \in \boldsymbol{S}_{++}^{d-\varphi}$.

C.11 Boostrap method for estimating the MSE and MAE

When the data distribution, or in our case, the underlying (assumed) smooth process, is unknown, we cannot follow the standard MSE and MAE computation procedure because the original function $\psi \in \mathcal{W}_\rho$ is also unknown. Instead, we only have access to a certain dataset of test samples, e.g., $\mathcal{D} = \{(x_t, \psi(x_t))\}_{t=1}^T$. Following a bootstrap-inspired method [144], we choose a subset $\mathcal{B} \subseteq \mathcal{D}$ for the signal reconstruction and use the complementary set $\overline{\mathcal{B}}$, i.e., $\mathcal{B} \cup \overline{\mathcal{B}} = \mathcal{D}$ and $\mathcal{B} \cap \overline{\mathcal{B}} = \emptyset$

Figure C.8: Illustration of the residuals used to estimate the MSE via bootstrapping. The test partition comes from the synthetic data and has been reconstructed with a Myopic(3,1) policy with $\eta = 1$ and $\rho = 2$. $\mathcal{B}$ contains 90% of the test partition and $\overline{\mathcal{B}}$ the remaining 10%.

to estimate the MSE and MAE performance metrics. Mathematically, this can be expressed as:

$$\text{MSE}\left(f_{\mathcal{B}}\right) = \frac{1}{|\overline{\mathcal{B}}|} \sum_{i \in \overline{\mathcal{B}}} \left(f_{\mathcal{B}}(x_i) - \psi(x_i)\right)^2, \tag{C.36a}$$

$$\text{MAE}\left(f_{\mathcal{B}}\right) = \frac{1}{|\overline{\mathcal{B}}|} \sum_{i \in \overline{\mathcal{B}}} |f_{\mathcal{B}}(x_i) - \psi(x_i)|, \tag{C.36b}$$

where $f_{\mathcal{B}}$ is the signal estimate constructed from the test data subset $\mathcal{B}$. This procedure is illustrated in Fig. C.8.

Notice that it is important to partition the test data because any test data sample used for the signal reconstruction cannot be used to compute the performance metrics. Otherwise, this results in data leakage. On the other hand, due to the lack of data samples, the performance metrics estimated in this way, may not be as accurate as if we had larger test sets, or more specifically, large test sets with higher temporal resolution. Thus, to reduce the variance of the MSE and MAE estimators, we repeat the procedure for several randomly chosen partitions $\mathcal{B}, \overline{\mathcal{B}}$ with replacement (i.e. they may repeat) and average the result. Particularly, we perform 10 repetitions.

Appendix D

Paper D

Consistent Signal Reconstruction from Streaming Multivariate Time Series

Emilio Ruiz-Moreno, Luis Miguel López-Ramos and Baltasar Beferull-Lozano

Abstract — **Digitalizing real-world analog signals typically involves sampling in time and discretizing in amplitude. Subsequent signal reconstructions inevitably incur an error that depends on the amplitude resolution and the temporal density of the acquired samples. From an implementation viewpoint, consistent signal reconstruction methods have proven a profitable error-rate decay as the sampling rate increases. Despite that, these results are obtained under offline settings. Therefore, a research gap exists regarding methods for consistent signal reconstruction from data streams. Solving this problem is of great importance because such methods could run at a lower computational cost than the existing offline ones or be used under real-time requirements without losing the benefits of ensuring consistency. In this paper, we formalize for the first time the concept of consistent signal reconstruction from streaming time series data. Then, we present a signal reconstruction method able to enforce consistency and also exploit the spatiotemporal dependencies of streaming multivariate time series data to further reduce the signal reconstruction error. Our experiments show that our proposed method achieves a favorable error-rate decay with the sampling rate compared to a similar but non-consistent reconstruction.**

D.1   Introduction

Most real-world physical quantities can only be observed by sampling their continuous-time waveforms (analog signals). If the analog signal is bandlimited, i.e., its spectral density has bounded support, we can ensure a reversible discretization in time by sampling at, or above, its Nyquist sampling rate [52, 53]. This is a universal sufficient condition but not necessary for all signals: some non-bandlimited signals can be perfectly reconstructed after uniformly sampled at, or above, its finite rate of innovation. Examples of such signals are streams of Dirac delta functions and *splines* [58]. Nonetheless, we can only observe a modified version of the sampled analog signal in practice because of the physical limitations of acquisition systems. For instance, acquisition systems such as analog-to-digital converters (ADCs) introduce both a sampling and a quantization step. The consequent discretization in amplitude causes an irreversible loss of information, making perfect signal reconstruction no longer possible.

Even though quantization is a deterministic operation, the signal reconstruction error-rate decay dependencies can be analyzed effectively from a stochastic point of view. For example, the reconstruction error caused by a uniform quantizer can be well approximated and modeled as an additive uniform white noise source independent of the input signal when certain assumptions are met [54, 55]. Under these assumptions, the quantization noise spectral density is constant. Thus, if an analog signal ψ is bandlimited, we can reduce the amount of in-band quantization noise by sampling above its Nyquist sampling rate (oversampling). Accordingly, we

(a) Snapshot at time x_t of a multivariate signal being converted from analog to digital. The acquired data samples form a streamed multivariate time series of quantization intervals.

(b) Snapshot at time x_t of the multivariate time series from Fig. D.1a being reconstructed. The consistent signal reconstruction method is represented visually as a microprocessor. The green shaded area indicates the currently reconstructed portion of the signal, while the red represents the future.

Figure D.1: Illustration of the considered signal acquisition-reconstruction process spanning from time 0 to X. Stationary spatial relationships are represented as edges in a graph with black dotted lines.

can reduce the reconstruction error by digitally filtering out the quantization noise above the original signal bandwidth. In this scenario, it is well known that low-pass filtering reduces the mean squared error (MSE) of a reconstructed signal f by a factor proportional to the squared quantization step size Δ^2 and inversely proportional to the oversampling ratio R. Formally,

$$\mathbb{E}\left[\|f - \psi\|_2^2\right] \propto \frac{\Delta^2}{R}, \tag{D.1}$$

where the expectation is taken over the amplitude of the reconstruction error, modeled as a uniform random variable [54]. The result shown in (D.1) is useful for identifying the main factors that influence the signal reconstruction error-rate decay under an oversampling regime. However, it does not reflect the effects of any quantization scheme or reconstruction method other than uniform quantization and linear decoding, i.e., low-pass filtering of the quantized signal at a cut-off frequency.

In contrast, deterministic analyses, based on deterministic bounds over the quantization levels, show that it is possible to reconstruct quantized signals with a squared norm error that is asymptotically proportional to the square of the quantization step size and inversely proportional to the square of the oversampling ratio, that is

$$\|f - \psi\|_2^2 = \mathcal{O}\left(\frac{\Delta^2}{R^2}\right). \tag{D.2}$$

The trend in (D.2) can be achieved by performing signal reconstruction via *consistent signal estimates*, where the term "consistent" here means coherent with all available knowledge about the original analog signal and the acquisition system. This has been shown for bandlimited signals [27, 56], and non-bandlimited signals with finite rate of innovation [28]. The behavior presented in (D.2) suggests that when a uniformly quantized and oversampled signal is reconstructed using consistent estimates, the reconstruction error can be asymptotically reduced at the same rate by either increasing the oversampling ratio or decreasing the quantization step size. This balanced signal reconstruction error-rate decay has important implications. For instance, in practice, it is more convenient to increase the oversampling ratio because the implementation cost of the analog circuitry necessary for halving the quantization step size is much higher than that for doubling the oversampling ratio [57].

Notwithstanding the above, consistent signal reconstructions usually require sophisticated methods that are far more computationally complex than linear decoding, thus hindering their practical implementation. Only a few previous studies have explored how to alleviate the computational complexity of yielding consistent reconstructions while maintaining the asymptotic behavior established in (D.2). For example, in [59], linear decoding computational complexity is achieved by using different quantization step sizes, while the difficulty lies in the design of the quantization scheme. As another example, the authors of [145] design a subtractive dithering method aimed at recursively solving an overdetermined system of linear equations set up from quantized data described by frames[1]. However, the resulting signal estimates are only guaranteed to be consistent with the most recently analyzed frame representation in the recursion. In summary, there seems to be a research gap

---

[1]Frames provide redundant and usually non-unique representations of vectors [146, 147].

regarding low-complexity consistent signal reconstruction methods under uniform quantization schemes.

Online strategies seem a reasonable approach to fill this research gap as they allow for a low runtime complexity with respect to the number of acquired samples. They can also be used to deal with streaming data sequentially, possibly under real-time requirements, in exchange for sacrificing some accuracy in the solution. To the best of our knowledge, there are no consistent signal reconstruction methods designed under online settings or for sequentially reconstructing streaming data samples. Consequently, there are no studies about their error-rate decay dependencies either.

This paper aims to cover the aforementioned gap. To this end, we focus on a class of multivariate, non-bandlimited, temporally and spatially interdependent smooth signals with a finite rate of innovation that can be expressed in terms of splines. This is justified because most signals of interest are smooth, usually elapse within a finite timespan, and rarely occur in complete isolation. Then, we consider an ADC that uniformly samples and quantizes these multivariate signals, resulting in streamed multivariate time series of quantization intervals, as illustrated in Fig. D.1a. From here, our goal is to design a consistent signal reconstruction approach from streaming data that satisfies two key requirements: i) signal smoothness, in agreement with the signals under consideration, and ii) a zero-delay response, i.e., the ability to process streaming data samples as they arrive and before the next sample is available. The considered requirements are of great importance in practice since they arise naturally in a wide range of signal reconstruction problems, ranging from high-speed digital to analog conversion [71] to online trajectory planning [67, 148], among others.

Concretely, this paper presents the first approach to zero-delay consistent signal reconstruction from a streaming multivariate time series of quantization intervals (see Fig. D.1b), as well as an extensive experimental study of its error-rate decay dependencies. Specifically, our approach follows the strategy presented in [3], which describes how to formulate a smoothing spline interpolation problem from streaming data under a sequential decision-making perspective and zero-delay requirements. In [3], a stationary parametric policy, i.e., a time-invariant decision strategy, based on tunable parameters on top of a recurrent neural network (RNN) [130] is devised. Then, the policy is trained over a set of representative streaming data to reduce a given trade-off between the sum of squared residuals and the roughness, i.e., a derivative-based measure of the model complexity [32], of the reconstructed signal. However, that strategy is intended for univariate and noisy data samples and fails to ensure consistency in the reconstruction as well as to exploit the spatial dependencies of multivariate streaming data.

The work in this paper considerably differs from the work in [3]. This is because our zero-delay consistent signal reconstructions are obtained by optimizing a different objective than the one introduced in [3]. Furthermore, the resulting optimization problem contains additional constraints to ensure consistency, such as enforcing the reconstructed signal to pass through quantization intervals, thus requiring to seek solutions in different feasible sets. These facts also substantially change the procedure in which the solution is found. Moreover, the multivariate analysis presented in this paper allows us to reach additional results and conclusions that would not be available after merely extending the approach in [3] to multivariate signals. Concretely, offline consistent signal reconstructions do not benefit from a multivariate

formulation (due to separability), whereas our zero-delay consistent signal reconstruction does improve (in terms of the error-rate decay) when incorporating the spatial relations among the multiple time series of quantization intervals. Therefore, the extension to the multivariate case provides new insights that improve our understanding of online consistent signal reconstructions.

The main contributions of this paper can be summarized as follows:

- We formalize the concept of consistent signal reconstruction from streaming data under zero-delay response requirements. As a result, our formulation generalizes the concept of consistency beyond offline settings.

- To the best of our knowledge, we devise the first consistent method for zero-delay signal reconstruction from streaming multivariate time series. The method can enforce consistency in a closed-form step allowing for a faster and more convenient implementation than numerical alternatives.

- We show experimentally that the reconstruction error incurred by our proposed method decays at the same rate by decreasing the quantization step size or increasing the oversampling ratio. Moreover, the error-rate decay slope with respect to the oversampling ratio doubles (in logarithmic scale) the one obtained with a similar zero-delay smooth signal reconstruction method that does not enforce consistency.

- Additionally, we observe that the error-rate decay incurred by our method does improve when incorporating the spatial relationships among the multiple time series, whereas offline consistent signal reconstructions do not benefit from a multivariate formulation.

The rest of the paper is organized as follows. Sec. D.2 introduces the main notation used throughout the paper. Sec. D.3 formalizes the notion of consistency and presents the concept of consistent signal reconstruction from streaming data under zero-delay requirements. Then, Sec. D.4 describes our proposed method from a sequential decision-making perspective, and Sec. D.5 analyzes its error-rate decay behavior experimentally. Finally, Sec. D.6 concludes the paper.

D.2   Mathematical notation

This section introduces the mathematical notation most recurrently used throughout the paper.

Vectors and matrices are represented in bold lowercase and bold capital letters, respectively. Given a vector $\boldsymbol{v} = [v_1, \ldots, v_C]^\top$ its cth component is indicated as $[\boldsymbol{v}]_c \triangleq v_c$. Similarly, given a matrix $\boldsymbol{M} \in \mathbb{R}^{R \times C}$, the element in the rth row and cth column is indicated as $[\boldsymbol{M}]_{r,c}$. The notation $[\boldsymbol{v}]_{i:j}$ refers to the sliced vector $[v_i, \ldots, v_j] \in \mathbb{R}^{j-i+1}$. We use Euler's notation for the derivative operator; thus, D_x^k denotes the kth derivative over the variable x. Lastly, we refer to functions f as elements of a function space or explicitly allude to its domain and codomain with the compact notation $f(x)$, depending on the context.

D.3 Consistency

This section describes and models the function space comprising the signals considered in this work and the acquisition system. Accordingly, we identify and define the set of consistent signal estimates and introduce the related concept of consistent signal reconstruction from streaming acquired data.

D.3.1 Function space of multivariate smooth signals

Most real-world physical processes are bounded and smooth due to energy conservation [123]. Moreover, most of them occur within a certain finite time lapse. Therefore, we assume any real-world process-associated measurable quantity to be described by smooth and non-bandlimited signals.

Let us define $\mathcal{W}_\rho$ as the function space of real functions defined over the domain $(0, X] \subseteq \mathbb{R}$ with $\rho - 1$ absolutely continuous derivatives and with the ρth derivative square integrable. Then, any of the aforementioned real-world signals can be accurately modeled by a function $\psi \in \mathcal{W}_\rho$. In practice, ψ can be approximated, or reconstructed, from a set of T_* uniformly sampled measurements $\{\psi(tX/T_*)\}_{t=1}^{T_*}$ as follows

$$f_* \triangleq \arg\min_{f \in \mathcal{W}_\rho} \int_0^X (D_x^\rho f(x))^2 \, dx \tag{D.3a}$$

$$\text{subject to: } f\left(\frac{tX}{T_*}\right) = \psi\left(\frac{tX}{T_*}\right), \ \forall t \in \mathbb{N}^{[1,T_*]}, \tag{D.3b}$$

where the metric used as objective in (D.3a) is usually referred to as roughness and f_* is unique as long as $T_* \geq \rho$ [149, 150]. Actually, any function $\psi \in \mathcal{W}_\rho$ with roughness bounded by an arbitrary constant is known to satisfy

$$\sup_{x \in [X/T_*, X)} \left| D_x^k \psi(x) - D_x^k f_*(x) \right| \leq \text{const} \cdot \left(\frac{X}{T_*}\right)^{\rho-k}, \tag{D.4}$$

for all $k \in \{0, \ldots, \rho - 1\}$ [24]. That is, the reconstruction error, up to the $(\rho - 1)$th derivative, is upper bounded by a term inversely proportional to the number T_* of measurements.

Since the reconstruction error in (D.4) can be made arbitrarily small by increasing the number of samples T_*, we restrict ourselves without loss of generality to those signals in $\mathcal{W}_\rho$ whose approximation, as in (D.3), is exact, i.e., the left-hand term in (D.4) is zero-valued for all allowed kth derivatives. The function space where such signals of interest belong is denoted as $\mathcal{V}_\rho \subseteq \mathcal{W}_\rho$, and turns out to be constituted by splines, i.e., piecewise polynomial functions, of order $2\rho - 1$ with T_* knots (joint points) at $\mathcal{X}_* = \{tX/T_*\}_{t=1}^{T_*}$ [24].

On the other hand, more than one signal often concurrently stems from the same physical process. Given that, one can expect spatial interdependencies among them. For this reason, this work deals with multivariate signals of the form $\boldsymbol{\Psi} = [\psi^{(1)}, \ldots, \psi^{(N)}]^\top$ with $\psi^{(n)} \in \mathcal{V}_{\rho_n}$ for all $n \in \mathbb{N}^{[1,N]}$, namely, $\boldsymbol{\Psi} \in \mathcal{V}_{\rho_1} \times \cdots \times \mathcal{V}_{\rho_N}$. However, and for the sake of simplicity in the notation, we rather focus on those multivariate signals in $\mathcal{V}_\rho \times \cdots \times \mathcal{V}_\rho = \mathcal{V}_\rho^N$, i.e., same order ρ across all N dimensions, since the considered formulation can be easily adapted to the presented general case.

D.3.2 Acquisition system and irreversibility

Thanks to the reconstruction procedure presented in (D.3), a multivariate signal $\boldsymbol{\Psi} \in \mathcal{V}_\rho^N$ can be reversibly sampled despite being non-bandlimited. That is, we can recover the multivariate signal $\boldsymbol{\Psi}$ from finitely many measurements as long as they contain its set of knots, namely, the signal values $\boldsymbol{\Psi}(x)$ sampled at $x \in \mathcal{X}_*$. For example, we can measure the knots by sampling uniformly at a rate $\nu_* = T_*/X$. In fact, any uniform sampling rate $\nu = R\nu_*$, with R being a positive integer oversampling ratio, measures all the knots in $\mathcal{X}_*$ because $\mathcal{X} = \{t\nu^{-1}\}_{t=1}^{RT_*} \supseteq \mathcal{X}_*$.

Ideally, these measurements would be acquired through an infinite amplitude-resolution sampling mechanism, namely, an ideal sampler. However, this is impossible in practice due to the physical limitations of measuring devices [151]. Instead, the usual procedure is to assume an ideal sampler and then model the resulting finite amplitude-resolution measurements as the outcome of a quantization mapping. In this work, we assume an ideal sampler working at a uniform sampling rate ν followed by a uniform midtread quantizer with a quantization step size $\Delta^{(n)}$ for each of the n signals.

The resulting acquired data consists of a streaming multivariate time series of N time series of length $T = RT_*$. We refer to every tth term of the streaming multivariate time series as an *observation*, $\boldsymbol{o}_t$. Every tth observation is constructed from N quantization intervals described by their time stamp $x_t = tX/T \in \mathcal{X}$, common for all intervals of the tth observation, their center $y_t^{(n)} \in \mathbb{R}$, and half quantization step size $\epsilon^{(n)} = \Delta^{(n)}/2 \in \mathbb{R}_+$. Formally, $\boldsymbol{o}_t = [x_t, \boldsymbol{y}_t^\top, \boldsymbol{\epsilon}^\top]^\top$ where $\boldsymbol{y}_t = [y_t^{(1)}, \ldots, y_t^{(N)}]^\top$ and $\boldsymbol{\epsilon} = [\epsilon^{(1)}, \ldots, \epsilon^{(N)}]\top$.

Unfortunately, the acquisition system under consideration is irreversible. That is, the multivariate signals in $\mathcal{V}_\rho^N$ may not be perfectly recovered from a finite set of measurements anymore if acquired as described above. This drawback is two-sourced. First, the sampler works sequentially. Thus, intermediate acquisition stages are irreversible because not all knots, i.e., the multivariate signal values at $\mathcal{X}_*$, have been measured yet. Second, the quantization step is a many-to-one mapping, i.e., non-injective and hence, not invertible. Consequently, incurring a reconstruction error in this context is unavoidable.

D.3.3 Consistent signal estimate

As stated in the previous Sec. D.3.2, it is generally impossible to recover a multivariate signal $\boldsymbol{\Psi} \in \mathcal{V}_\rho^N$ from its stream of observations. Nevertheless, we can select a relatively reduced set of functions where the multivariate signal $\boldsymbol{\Psi}$ lies. Specifically, for every nth time series and for any of its ith quantization intervals we can construct its associated hyperslab $\mathcal{H}_i^{(n)} \subseteq \mathcal{V}_\rho$ as

$$\mathcal{H}_i^{(n)} = \left\{ h \in \mathcal{V}_\rho : \left| h(x_i) - y_i^{(n)} \right| \leq \epsilon^{(n)} \right\}. \tag{D.5}$$

Then, at time step t, we can state that each signal $\psi^{(n)}$ belongs to the intersection of its so-far constructed hyperslabs. That is, $\psi^{(n)} \in \mathcal{H}_{1:t}^{(n)}$ with $\mathcal{H}_{1:t}^{(n)} \triangleq \bigcap_{i=1}^t \mathcal{H}_i^{(n)}$, as illustrated in Fig. D.2a. In this context, any signal estimate described by a function drawn from $\mathcal{H}_{1:t}^{(n)}$ is called a consistent signal estimate at time step t, or simply consistent signal estimate, when clear by context. This is because any reacquisition, i.e., processing the newly estimated signal through the same acquisition system

(a) Intersection of hyperslabs as a Venn diagram. At time step t, the functions within the area enclosed by the dotted black line, $\mathcal{H}_{1:t}^{(n)}$, describe current consistent signal estimates.

(b) Comparison between the ongoing reconstructed signals using a non-consistent (NC) and a consistent (C) method. Fig. D.2a illustrates how both methods produce a sequence of consistent signal estimates, namely, $\{\tilde{f}_t^{(n)}\}_{t=1}^T$ and $\{f_t^{(n)}\}_{t=1}^T$. However, the sequence $\{\tilde{f}_t^{(n)}\}_{t=1}^T$ does not lead to a consistent signal reconstruction, while the sequence $\{f_t^{(n)}\}_{t=1}^T$ does.

Figure D.2: Visual representation of a consistent signal estimate and a consistent signal reconstruction.

until the same time step t, leads to the same set of already received t quantization intervals. Equivalently, any signal estimate is conventionally referred to as consistent if and only if: i) it passes through all the (so-far) received quantization intervals, and ii) it is described by a function in $\mathcal{V}_\rho$ [57].

Recall from Sec. D.3.1 that $\mathcal{V}_\rho$ is constituted by splines of order $2\rho - 1$, with $\rho - 1$ continuous derivatives and T_* knots located at time stamps in $\mathcal{X}_*$. For this reason, we construct every nth signal estimate as a spline composed of T pieces, or function sections, as

$$f^{(n)}(x) = \begin{cases} g_1^{(n)}(x), & \text{if } 0 < x \le x_1 \\ g_2^{(n)}(x), & \text{if } x_1 < x \le x_2 \\ \vdots \\ g_T^{(n)}(x), & \text{if } x_{T-1} < x \le x_T \end{cases} \tag{D.6}$$

where every tth function section $g_t^{(n)} : (x_{t-1}, x_t] \to \mathbb{R}$ is a linear combination of polynomials of the form

$$g_t^{(n)}(x) = \boldsymbol{a}_t^{(n)\top} \boldsymbol{p}_t(x), \tag{D.7}$$

with combination coefficients $\boldsymbol{a}_t^{(n)} \in \mathbb{R}^{2\rho}$ and basis vector function $\boldsymbol{p}_t : (x_{t-1}, x_t] \to \mathbb{R}^{2\rho}$ defined as

$$\boldsymbol{p}_t(x) = \left[1, (x - x_{t-1}), \ldots, (x - x_{t-1})^{2\rho-1}\right]^\top. \tag{D.8}$$

Then, the remaining requirements for $f^{(n)}$ to belong to $\mathcal{V}_\rho$ are that the location of its knots must cover all time stamps in $\mathcal{X}_*$ and that it has to be continuous up to the $(\rho - 1)$th derivative. By design, the knots of $f^{(n)}$ lie at $\mathcal{X}$ hence, covering all time stamps in $\mathcal{X}_*$ since $\mathcal{X}_* \subseteq \mathcal{X}$. Regarding its continuity, for $f^{(n)}$ to be $(\rho - 1)$-smooth, the function sections must satisfy

$$\lim_{x \to x_{t-1}^-} D_x^k g_{t-1}^{(n)}(x) = \lim_{x \to x_{t-1}^+} D_x^k g_t^{(n)}(x), \tag{D.9}$$

for every $k \in [0, \rho - 1]$ and $t \in [2, T]$. This can be ensured by imposing the following equality constraint

$$\left[\boldsymbol{a}_t^{(n)}\right]_{1:\rho} = \boldsymbol{e}_{t-1}^{(n)}, \tag{D.10}$$

for each $t \in \mathbb{N}^{[1,T]}$, where $\boldsymbol{e}_t^{(n)} \in \mathbb{R}^\rho$ is a vector with elements given by

$$[\boldsymbol{e}_t^{(n)}]_i = \frac{1}{(i-1)!} \sum_{j=1}^{2\rho} \left[\boldsymbol{a}_t^{(n)}\right]_j u_t^{j-i} \prod_{k=1}^{i-1} (j - k), \tag{D.11}$$

where $u_t \triangleq x_t - x_{t-1}$, except $u_1 = x_1$, and where $\boldsymbol{e}_0^{(n)}$ determines the initial boundary conditions of the signal estimate ([3, Proposition 1]).

In this way, any signal estimate $f^{(n)}$ constructed as in (D.6), and satisfying (D.9), belongs to $\mathcal{V}_\rho$. Additionally, and as stated before, if at time step t it passes through all the so-far received quantization intervals, i.e., $f^{(n)}$ satisfies $|f^{(n)}(x_i) - y_i^{(n)}| \le \epsilon^{(n)}$ for all $i \in \{1, \ldots, t\}$, or equivalently $f^{(n)} \in \mathcal{H}_{1:t}^{(n)}$, then it is a consistent signal estimate at time step t.

D.3.4 Zero-delay consistent signal reconstruction

Recall that the acquisition system modeled in Sec. D.3.2 generates a stream of observations. In this paper, we pursue its converse reconstruction approach; that is, our goal is to sequentially reconstruct a signal $\boldsymbol{\Psi} \in \mathcal{V}_\rho^N$ as its observations are received. To this end, we devise a zero-delay multivariate signal reconstruction method. In this context, our objective is to yield a sequence of signal estimates in $\mathcal{V}_\rho^N$, where the tth multivariate signal estimate $\boldsymbol{f}_t = [f_t^{(1)}, \ldots, f_t^{(N)}]^\top$ is used to reconstruct the multivariate time series between the last two received observations, i.e., $\boldsymbol{o}_{t-1}$ and $\boldsymbol{o}_t$. Moreover, every reconstruction step has to be performed as soon as the tth observation is received and without knowing the to-be-received observations[2].

From the discussion in Sec. D.3.3, one may expect to obtain a consistent signal reconstruction using a zero-delay signal reconstruction method that generates a sequence of consistent signal estimates. That is, for each nth time series every tth signal estimate $f_t^{(n)}$ belongs to $\mathcal{H}_{1:t}^{(n)}$. However, it is important to notice that a method producing a sequence of consistent signal estimates does not necessarily achieve a consistent reconstruction. This is because guaranteeing a consistent reconstruction requires the sequence of signal estimates to fulfill additional requisites. Specifically, every two consecutive consistent signal estimates, namely $f_{t-1}^{(n)}$ and $f_t^{(n)}$, must ensure smoothness at the takeover time stamp x_{t-1}, as shown in Fig. D.2b. Formally, a consistent signal reconstruction is attainable from a sequence of consistent signal estimates $\{f_t^{(n)}\}_{t=1}^T$ if they satisfy that $c_t^{(n)} \in \mathcal{H}_{1:t}^{(n)}$, where

$$c_t^{(n)}(x) = \begin{cases} f_{t-1}^{(n)}(x) & \text{if } 0 < x \le x_{t-1}, \\ f_t^{(n)}(x) & \text{if } x_{t-1} < x \le X, \end{cases} \tag{D.12}$$

for all $t \in \mathbb{N}^{[2,T]}$.

In our case, the presented zero-delay reconstruction methods can yield a consistent signal reconstruction by sequentially updating the spline-based signal estimate coefficients. Concretely, if at every time step t, we update the remaining $T - t$ spline coefficients, i.e., $\boldsymbol{a}_i^{(n)}$ for $i = t, \ldots, T$ subject to satisfying the corresponding continuity constraints in (D.10), and subject to passing through the last received quantization interval, i.e., the tth spline coefficients obey $|\boldsymbol{a}_t^{(n)\top} \boldsymbol{p}_t(x_t) - y_t^{(n)}| \le \epsilon^{(n)}$, then the resulting signal estimate update $f_t^{(n)}$ is part of a sequence satisfying (D.12). In this way, the overall signal reconstruction, i.e., the reconstructed signal at the last time step T, is completely described by the last consistent signal estimate $f_T^{(n)}$. In practice, updating just the tth spline coefficients at every time step t (temporarily ignoring the $T - t$ remaining coefficients) eventually leads to the same consistent signal reconstruction.

Given this, notice that the core difference between the offline consistent signal reconstruction methods discussed in Sec. D.1 and our approach lies in the fact that we explicitly account for the sequential nature of the acquisition system. This is conceptually more challenging than assuming that all data samples are available in advance since in the sequential case we cannot achieve a consistent signal reconstruction with just a single consistent signal estimate, but with a sequence of them,

---

[2]For the sake of completeness, implementing a zero-delay signal reconstruction method also requires a reduced constant complexity per iteration to guarantee the required execution speed and scalability [124].

subject to (D.12), instead. This significant difference opens the question of whether a method for zero-delay consistent signal reconstruction has the corresponding benefits in the reconstruction error-rate decay behavior as the offline consistent ones have [27, 56, 59, 28]. We anticipate a positive answer to the posed question, but before presenting such an outcome, we introduce our zero-delay consistent signal reconstruction approach in detail.

D.4   Zero-delay spline interpolation

This section introduces a zero-delay spline interpolation approach based on policy training that, differently from [3], can accommodate multivariate quantized data and enforce consistency as defined in Sec. D.3.

D.4.1   A trainable multivariate approach to zero-delay spline interpolation

The problem of reconstructing a multivariate time series of quantization intervals using a sequence of spline-based signal estimates under smoothness and zero-delay requirements can be formalized from a sequential decision-making perspective as follows. At time step t, we encode the condition of the so-far reconstructed signal, and the last received quantization interval in a vector-valued variable referred to as the tth state $\boldsymbol{s}_t \in \mathcal{S}$, being $\mathcal{S}$ the state space. Every tth state is constructed as $\boldsymbol{s}_t = [x_{t-1}, \boldsymbol{o}_t^\top, \boldsymbol{e}_{t-1}^\top]^\top$ with $\boldsymbol{e}_{t-1} = [\boldsymbol{e}_{t-1}^{(1)^\top}, \ldots, \boldsymbol{e}_{t-1}^{(N)^\top}]^\top$ where every nth component $\boldsymbol{e}_{t-1}^{(n)}$ represents the right hand term at (D.10). In fact, every tth state is fully determined once all nth spline coefficients $\boldsymbol{a}_t^{(n)}$ are fixed, and the tth observation $\boldsymbol{o}_t$ is received. This can be explicitly described by a state update mechanism through a deterministic mapping, as $\boldsymbol{s}_{t+1} = F(\boldsymbol{s}_t, \boldsymbol{a}_t, \boldsymbol{o}_{t+1})$, allowing us to identify all visitable states handily. On the other hand, selecting all nth spline coefficients at time step t can be understood as an action. This is because every nth function section $g_t^{(n)}$, as defined in (D.7), is fully determined as soon as $\boldsymbol{a}_t^{(n)}$ is chosen and hence so is the resulting piece of the subsequent signal reconstruction. Explicitly, we construct every tth action by concatenating all nth spline coefficients, i.e., $\boldsymbol{a}_t = [\boldsymbol{a}_t^{(1)^\top}, \ldots, \boldsymbol{a}_t^{(N)^\top}]^\top$, and denote the action space as $\mathcal{A} \subseteq \mathbb{R}^{2N\rho}$. We distinctively indicate the set of actions accessible from a given state $\boldsymbol{s}_t$ as the admissible action set $\mathcal{A}(\boldsymbol{s}_t) \subseteq \mathcal{A}$. In particular, we select the tth action through a stationary parametric policy $\boldsymbol{\mu_\theta} : \mathcal{S} \to \mathcal{A}$. Stationary policies do not change over time, and they are suitable for making decisions in problems governed by stationary spatiotemporal dynamics and with varying time steps. On the other hand, policy parametrization helps to reduce the pool of candidate policies and allows, with some flexibility, to accommodate the problem structure, e.g., spatiotemporal dynamics and constraints, into the policy.

In this context, any reconstructed signal can be assessed and therefore compared by computing a cumulative cost $K : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$ through the traversed state-action pairs by following a given policy $\boldsymbol{\mu_\theta}$. From here, we can tune the policy parameters

Table D.1: Per time-series cost $\kappa(\cdot, \cdot)$ and admissible action set $\mathcal{A}(\cdot)$ for the consistent (■) and smoothing (□) policy variations.

Per time-series cost evaluated at the tth state-action pair
$\boldsymbol{a}_t^{(n)\top} \boldsymbol{M}_t \boldsymbol{a}_t^{(n)}$
$\left(\boldsymbol{p}_t^\top(x_t)\boldsymbol{a}_t^{(n)} - y_t^{(n)}\right)^2 + \eta\,\boldsymbol{a}_t^{(n)\top}\boldsymbol{M}_t\boldsymbol{a}_t^{(n)}$, with $\eta > 0$

Admissible action set for the tth state
$\Big\{\boldsymbol{a} \in \mathcal{A} : \left\|\boldsymbol{p}_t^\top(x_t)\boldsymbol{a}^{(n)} - y_t^{(n)}\right\| \leq \epsilon^{(n)},$ $[\boldsymbol{a}^{(n)}]_{1:\rho} = \boldsymbol{e}_{t-1}^{(n)},\ \forall n \in [1, N]\Big\}$
$\Big\{\boldsymbol{a} \in \mathcal{A} : [\boldsymbol{a}^{(n)}]_{1:\rho} = \boldsymbol{e}_{t-1}^{(n)},\ \forall n \in [1, N]\Big\}$

via policy training by solving the following optimization problem

$$\arg\min_{\boldsymbol{\theta}\in\mathbb{R}^P} \sum_{m=1}^{M}\sum_{t=1}^{T} K\left(\boldsymbol{s}_{m,t}, \boldsymbol{\mu_\theta}(\boldsymbol{s}_{m,t})\right) \tag{D.13a}$$

$$\text{s. to: } \boldsymbol{s}_{m,t} = F\left(\boldsymbol{s}_{m,t-1}, \boldsymbol{\mu_\theta}(\boldsymbol{s}_{m,t-1}), \boldsymbol{o}_{m,t}\right), \forall m, t, \tag{D.13b}$$

$$\boldsymbol{\mu_\theta}(\boldsymbol{s}_{m,t}) \in \mathcal{A}(\boldsymbol{s}_{m,t}), \forall m, t, \tag{D.13c}$$

where the integer M denotes the number of example multivariate time series, indexed by m, and where all the stream of observations $\boldsymbol{o}_{m,t}$ and initial states $\boldsymbol{s}_{m,0}$ are given. Thus, if a sufficiently large collection of representative multivariate time series is available, one expects the trained policy to outperform, on average, any other policy with different parameter values when tested over a previously unseen stream of observations.

D.4.2 Policy architectures and variations

In this work, we are interested in stationary parametric policies of the form

$$\boldsymbol{\mu_\theta}(\boldsymbol{s}_t) = \arg\min_{\boldsymbol{a}\in\mathcal{A}(\boldsymbol{s}_t)} \left\{K(\boldsymbol{s}_t, \boldsymbol{a}) + J_{\boldsymbol{\theta}}(\boldsymbol{s}_t, \boldsymbol{a}; \boldsymbol{h}_t)\right\}, \tag{D.14}$$

where K denotes the same cost used in the objective (D.13a), and the map $J_{\boldsymbol{\theta}} : \mathcal{S} \times \mathcal{A} \times \mathbb{R}^L \to \mathbb{R}$ is a parametric approximation of the cost-to-go [45]. Here, $\boldsymbol{h}_t \in \mathbb{R}^L$ represents a latent state, common for all N time series at the tth time step, and it plays the role of policy memory [133, 134, 135]. Notice that the choice of $J_{\boldsymbol{\theta}}$ determines the architecture of the policy, i.e., the relation between the P parameters within $\boldsymbol{\theta}$. In this work, we adopt three possible policy architectures.

First, we consider a *myopic* policy consisting of a parameterless policy architecture where $J_{\boldsymbol{\theta}}(\boldsymbol{s}_t, \boldsymbol{a}; \boldsymbol{h}_t) = 0$ for all t. The myopic policy is arguably the simplest architecture-wise but ignores the cost-effect of current actions in future decisions. Second, we suggest an *RNN-based* policy corresponding to the following parametric

cost-to-go approximation

$$J_{\boldsymbol{\theta}}(\boldsymbol{s}_t, \boldsymbol{a}; \boldsymbol{h}_t) = \sum_{n=1}^{N} \lambda_n \left\| \boldsymbol{a}^{(n)} - \begin{bmatrix} \boldsymbol{0}_\rho \\ \boldsymbol{r}_t^{(n)} \end{bmatrix} \right\|_2^2, \tag{D.15}$$

where the vectors $\boldsymbol{r}_t^{(n)} \in \mathbb{R}^\rho$, for all $n \in \mathbb{N}^{[1,N]}$, and $\boldsymbol{h}_t$ represent the output value and latent state of an RNN $R_{\boldsymbol{\theta}'} : \mathcal{S} \times \mathbb{R}^L \to \mathbb{R}^{N\rho} \times \mathbb{R}^L$, respectively. They are obtained through the following relation

$$R_{\boldsymbol{\theta}'}(\boldsymbol{s}_t; \boldsymbol{h}_t) = \begin{bmatrix} \boldsymbol{r}_t \\ \boldsymbol{h}_{t+1} \end{bmatrix}, \tag{D.16}$$

where $\boldsymbol{r}_t = [\boldsymbol{r}_t^{(1)^\top}, \ldots, \boldsymbol{r}_t^{(N)^\top}]^\top$, $\boldsymbol{\theta}' \in \mathbb{R}^{P-N}$ and $\boldsymbol{\theta} = [\lambda_1, \ldots, \lambda_N, \boldsymbol{\theta}'^\top]^\top \in \mathbb{R}_+^N \times \mathbb{R}^{P-N}$. The vector $\boldsymbol{0}_\rho$ is used in (D.15) to disregard the components of every $\boldsymbol{a}^{(n)}$ that are already decided by imposing the continuity constraints in (D.10). The RNN-based policy promotes actions close, in the sense of a Euclidean distance, to the output of a certain RNN, in this case, $R_{\boldsymbol{\theta}'}$. Thus, an RNN that successfully captures the spatiotemporal dynamics of the process generating the multivariate time series can lead to admissible actions with a reduced-cost effect in future decisions. Finally, we utilize a *batch* policy, or solution, involving an exact cost-to-go approximation. The name of the policy comes from the fact that you need the whole batch of streaming observations in advance to construct the exact cost-to-go, i.e., solving (D.3). Thus, in practice, it cannot be used under zero-delay requirements. The batch solution is parameterless, and by definition, it yields the signal reconstruction with the lowest possible cost.

From the presented policy architectures, notice that we can build different policy variations depending on the choice of the cost K and admissible set $\mathcal{A}(\cdot)$. In this work, we devise a policy yielding consistent signal reconstructions and present another policy variation designed to reconstruct noisy data smoothly but not consistently. Both presented policy variations, called from now on consistent and smoothing, respectively, are based on the fact that the roughness, as expressed in (D.3a), can be equivalently computed from a trajectory of state-action pairs, as we have shown in our previous work ([3, Proposition 2]). Specifically, the roughness of a spline-based signal estimate $\boldsymbol{f}_T = [f_T^{(1)}, \ldots, f_T^{(N)}]^\top$ over the domain $(0, X]$ can be computed as

$$\sum_{n=1}^{N} \int_0^X \left(D_x^\rho f_T^{(n)}(x) \right)^2 dx = \sum_{t=1}^{T} \sum_{n=1}^{N} \boldsymbol{a}_t^{(n)^\top} \boldsymbol{M}_t \boldsymbol{a}_t^{(n)}, \tag{D.17}$$

where the elements of $\boldsymbol{M}_t \in \boldsymbol{S}_+^{2\rho}$ are obtained through

$$[\boldsymbol{M}_t]_{i,j} = \begin{cases} 0 & \text{if } i \leq \rho \text{ or } j \leq \rho, \\ u_t^{i+j-2\rho-1} \frac{\prod_{k=1}^{\rho}(i-k)(j-k)}{i+j-2\rho-1} & \text{otherwise,} \end{cases} \tag{D.18}$$

with $u_t \triangleq x_t - x_{t-1}$, except $u_1 = x_1$. Similarly, we can construct a roughness-aware penalization from a cost of the form

$$K(\boldsymbol{s}_t, \boldsymbol{a}_t) = \sum_{n=1}^{N} \kappa\left(\boldsymbol{s}_t^{(n)}, \boldsymbol{a}_t^{(n)} \right), \tag{D.19}$$

Figure D.3: Snapshot of an ongoing signal reconstruction using the smoothing policy variation.

Figure D.4: Conceptualization of the hyperplane projection $P_{\widehat{H}_t^{(n)}}$ in the inner product space $\left(\mathbb{R}^\rho, \langle \cdot, \cdot \rangle_{\boldsymbol{D}_t^{(n)}} \right)$.

where $\boldsymbol{s}_t^{(n)}$ is an auxiliar per time-series state constructed from the tth state, $\boldsymbol{s}_t = [x_{t-1}, \boldsymbol{o}_t^\top, \boldsymbol{e}_{t-1}^\top]^\top$, as $\boldsymbol{s}_t^{(n)} = [x_{t-1}, x_t, y_t^{(n)}, \epsilon^{(n)}, \boldsymbol{e}_{t-1}^{(n)}{}^\top]^\top$ and where κ denotes a per time-series cost. Table D.1 summarizes our proposed policy variations.

Briefly, the consistent policy variation ensures smoothness while enforcing that the reconstructed signal passes through the quantization intervals, thus yielding a consistent signal reconstruction. On the other hand, the smoothing policy variation extends the smoothing interpolation approach, introduced in our previous work [3], to multivariate time series. Specifically, it aims to reduce the squared residuals with the center of the intervals while promoting smoothness, but it does not enforce that the ongoingly reconstructed signal passes through the quantization intervals, as illustrated in Fig. D.3, hence neither consistency.

D.4.3 Policy evaluation

By design, evaluating the considered policies involves solving optimization problem (D.14). The myopic and RNN-based policy architectures, in any of their variations given in Table D.1, can be equivalently evaluated by solving the following quadratic

Table D.2: Terms of the quadratic form for the smoothing policy variation. This notation is shared with the rest of the paper with the incorporation of $\boldsymbol{P}_t \triangleq \boldsymbol{p}_t(x_t)\boldsymbol{p}_t(x_t)^\top$ and $\boldsymbol{v}_t^{(n)} \triangleq [\boldsymbol{0}_\rho^\top, \boldsymbol{r}_t^{(n)^\top}]^\top$.

	$\boldsymbol{A}_t^{(n)}$	$\boldsymbol{b}_t^{(n)}$
Myopic	$\boldsymbol{P}_t + \eta \boldsymbol{M}_t$	$-2y_t^{(n)}\boldsymbol{p}_t(x_t)$
RNN	$\boldsymbol{P}_t + \eta \boldsymbol{M}_t + \lambda_n \boldsymbol{I}_{2\rho}$	$-2(y_t^{(n)}\boldsymbol{p}_t(x_t) + \lambda_n \boldsymbol{v}_t^{(n)})$

Table D.3: Terms of the quadratic form for the consistent policy variation. Same notation as Table D.2.

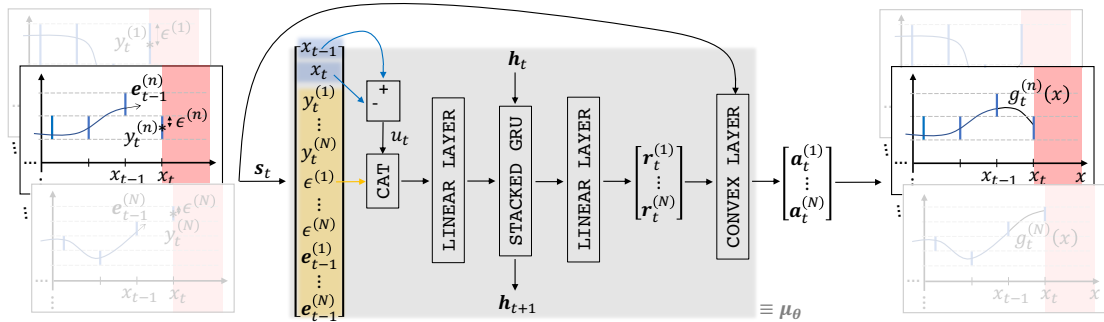|	$\boldsymbol{A}_t^{(n)}$	$\boldsymbol{b}_t^{(n)}$
Myopic	$\boldsymbol{M}_t$	$\boldsymbol{0}_{2\rho}$
RNN	$\boldsymbol{M}_t + \lambda_n \boldsymbol{I}_{2\rho}$	$-2\lambda_n \boldsymbol{v}_t^{(n)}$

Figure D.5: Representation of a tth forward pass of our RNN-based policy. CAT refers to a vector concatenation layer. Evaluating the policy consists of solving a convex optimization problem with respect to $\boldsymbol{a}$, see (D.14) and Sec. D.4.3. Therefore, its last step can be seen as a convex optimization layer [14].

convex problem

$$\boldsymbol{\mu_\theta}(\boldsymbol{s}_t) = \arg\min_{\boldsymbol{a} \in \mathcal{A}(\boldsymbol{s}_t)} \left\{ \sum_{n=1}^{N} \boldsymbol{a}^{(n)\top} \boldsymbol{A}_t^{(n)} \boldsymbol{a}^{(n)} + \boldsymbol{b}_t^{(n)\top} \boldsymbol{a}^{(n)} \right\}, \tag{D.20}$$

where the terms $\boldsymbol{A}_t^{(n)} \in \boldsymbol{S}_+^{2\rho}$ and $\boldsymbol{b}_t^{(n)} \in \mathbb{R}^{2\rho}$ depend on the chosen policy architecture and variation, as displayed in Tables D.2 and D.3. The policy evaluation in (D.20) is presented as an intermediate step for clarity. Then, after some algebraic steps, it can be equivalently reformulated as

$$\boldsymbol{\mu_\theta}(\boldsymbol{s}_t) = \arg\min_{\boldsymbol{a} \in \mathcal{A}(\boldsymbol{s}_t)} \left\{ \sum_{n=1}^{N} \|\boldsymbol{a}^{(n)} - \boldsymbol{\alpha}_t^{(n)}\|_{\boldsymbol{D}_t^{(n)}}^2 \right\} \tag{D.21}$$

where every $\boldsymbol{D}_t^{(n)} \in \boldsymbol{S}_{++}^{\rho}$ and $\boldsymbol{\alpha}_t^{(n)} \in \mathbb{R}^{\rho}$ are computed from the elements of the tth state, through $\boldsymbol{A}_t^{(n)}$ and $\boldsymbol{b}_t^{(n)}$, as

$$\boldsymbol{D}_t^{(n)} = \boldsymbol{B}_2^\top \boldsymbol{A}_t^{(n)} \boldsymbol{B}_2, \tag{D.22}$$

and

$$\boldsymbol{\alpha}_t^{(n)} = -\boldsymbol{D}_t^{(n)^{-1}} \left(\boldsymbol{B}_2 \boldsymbol{A}_t^{(n)} \boldsymbol{B}_1 \boldsymbol{e}_{t-1}^{(n)} + \frac{1}{2} \boldsymbol{B}_2^\top \boldsymbol{b}_t^{(n)} \right), \tag{D.23}$$

for all $n \in \mathbb{N}^{[1,N]}$ and $t \in \mathbb{N}^{[1,T]}$. The auxiliary matrices $\boldsymbol{B}_1 \triangleq [\boldsymbol{I}_\rho, \boldsymbol{0}_{\rho \times \rho}]^\top \in \mathbb{R}^{2\rho \times \rho}$ and $\boldsymbol{B}_2 \triangleq [\boldsymbol{0}_{\rho \times \rho}, \boldsymbol{I}_\rho]^\top \in \mathbb{R}^{2\rho \times \rho}$ are defined for notational brevity. The evaluation form in (D.21) is analytically preferable since it corresponds to a minimum norm problem or a hyperslab projection problem [23, 1], for the smoothing or consistent policy variations, respectively. This is because the aforementioned problems admit the following closed-form solution

$$\boldsymbol{\mu_\theta}(\boldsymbol{s}_t) = \begin{bmatrix} \boldsymbol{a}_t^{(1)^*} \\ \vdots \\ \boldsymbol{a}_t^{(N)^*} \end{bmatrix}, \quad \text{with } \boldsymbol{a}_t^{(n)^*} = \begin{bmatrix} \boldsymbol{e}_{t-1}^{(n)} \\ \zeta_t^{(n)} \left(\boldsymbol{\alpha}_t^{(n)} \right) \end{bmatrix}, \quad \forall n, \tag{D.24}$$

where $\zeta_t^{(n)} : \mathbb{R}^\rho \to \mathbb{R}^\rho$ is the identity operator for the smoothing policy variation and the projection onto the hyperslab

$$\tilde{\mathcal{H}}_t^{(n)} \triangleq \left\{ \boldsymbol{\beta} \in \mathbb{R}^\rho : \left| \boldsymbol{\beta}^\top \boldsymbol{q}_t - w_t^{(n)} \right| \le \epsilon^{(n)} \right\}, \tag{D.25}$$

for the consistent one, where $\boldsymbol{q}_t \triangleq [\boldsymbol{p}_t(x_t)]_{\rho+1:2\rho}$ and $w_t^{(n)} \triangleq y_t^{(n)} - \boldsymbol{e}_{t-1}^{(n)^\top} [\boldsymbol{p}_t(x_t)]_{1:\rho}$ for all $t \in \mathbb{N}^{[1,T]}$ and $n \in \mathbb{N}^{[1,N]}$. Every hyperslab $\tilde{\mathcal{H}}_t^{(n)}$ can be visualized as the set of all points which belong between and onto the hyperplanes

$$\hat{H}_t^{(n)} \triangleq \left\{ \boldsymbol{\beta} \in \mathbb{R}^\rho : \boldsymbol{\beta}^\top \boldsymbol{q}_t = w^{(n)} + \epsilon^{(n)} \right\}, \tag{D.26a}$$

$$\check{H}_t^{(n)} \triangleq \left\{ \boldsymbol{\beta} \in \mathbb{R}^\rho : \boldsymbol{\beta}^\top \boldsymbol{q}_t = w^{(n)} - \epsilon^{(n)} \right\}. \tag{D.26b}$$

Based on this observation, the projection onto the hyperslab $\tilde{\mathcal{H}}_t^{(n)}$ can be computed as

$$\zeta_t^{(n)}(\boldsymbol{\beta}) = \begin{cases} P_{\hat{H}_t^{(n)}}(\boldsymbol{\beta}), & \text{if } \boldsymbol{\beta}^\top \boldsymbol{q}_t > w_t^{(n)} + \epsilon^{(n)}, \\ \boldsymbol{\beta}, & \text{if } \left| \boldsymbol{\beta}^\top \boldsymbol{q}_t - w_t^{(n)} \right| \le \epsilon^{(n)}, \\ P_{\check{H}_t^{(n)}}(\boldsymbol{\beta}), & \text{if } \boldsymbol{\beta}^\top \boldsymbol{q}_t < w_t^{(n)} - \epsilon^{(n)}, \end{cases} \tag{D.27}$$

where $P_{\hat{H}_t^{(n)}} : \mathbb{R}^\rho \to \hat{H}_t^{(n)}$ and $P_{\check{H}_t^{(n)}} : \mathbb{R}^\rho \to \check{H}_t^{(n)}$ denote the projections onto the hyperplanes $\hat{H}_t^{(n)}$ and $\check{H}_t^{(n)}$, respectively. These projections can be computed in closed form as

$$P_{\hat{H}_t^{(n)}}(\boldsymbol{\beta}) = \boldsymbol{\beta} - \left(\boldsymbol{\beta}^\top \boldsymbol{q}_t - w_t^{(n)} - \epsilon^{(n)}\right) \frac{\left(\boldsymbol{D}_t^{(n)}\right)^{-1} \boldsymbol{q}_t}{\boldsymbol{q}_t^\top \left(\boldsymbol{D}_t^{(n)}\right)^{-1} \boldsymbol{q}_t}, \tag{D.28a}$$

$$P_{\check{H}_t^{(n)}}(\boldsymbol{\beta}) = \boldsymbol{\beta} - \left(\boldsymbol{\beta}^\top \boldsymbol{q}_t - w_t^{(n)} + \epsilon^{(n)}\right) \frac{\left(\boldsymbol{D}_t^{(n)}\right)^{-1} \boldsymbol{q}_t}{\boldsymbol{q}_t^\top \left(\boldsymbol{D}_t^{(n)}\right)^{-1} \boldsymbol{q}_t}. \tag{D.28b}$$

See Fig. D.4 for some intuition behind the closed-form projection in (D.28a). The same reasoning applies to the other hyperplane.

D.4.4   Policy training

Unlike the myopic policy architecture and batch solution, which are parameterless, the RNN-based architecture requires to be tuned, in this case, through (D.13), for adequate performance. Since evaluating the RNN-based policy can be seen as performing a forward pass from a deep learning point of view [47], the policy training problem (D.13) can be tackled via backpropagation through time [48]. Here, the closed-form evaluation described in Sec. D.4.3 allows us to compute and propagate the gradient of the tth action $\boldsymbol{a}_t$ with respect to the parameters contained in $\boldsymbol{\theta}$ without the need of unrolling numerical optimizers [137] or other specialized tools [14]. An overall scheme of our chosen RNN-based policy and its components can be found in Fig. D.5.

D.5   Error-rate decay

In this section, we experimentally analyze the signal reconstruction error-rate decay incurred by the consistent variation of the RNN-based policy (see Sec. D.4.2), with respect to the oversampling ratio and quantization step size of the acquisition system. To this end, we first set a specific function space of multivariate smooth signals, as described in Sec. D.3.1. For simplicity, we fix the dimensionality of the multivariate signal to $N = 2$, and specify the roughness by choosing $\rho = 2$. Then, we describe a form to generate multivariate signals from $\mathcal{V}_2^2$. These multivariate signals, along with an acquisition system with given oversampling ratio and quantization stepsize values, shape the time series data of the experiment. Separately, we use the myopic policy architecture as a benchmark, or base method, which does not exploit the spatiotemporal dependencies among time series, and the batch solution as a baseline, or best possible performance indicator. The central goal of this section is to compare the performance of the RNN-based policy architecture that considers the possible spatial relations between the time series (multivariate) against a policy based on the same RNN architecture but ignoring the spatial relations (univariate). This is done for both policy variations: consistent and smoothing.

Figure D.6: Scheme of the acyclic VAR(1) process generating the 2 series of knots for any mth sequence. The thickness of the arrows represents the magnitude of the autoregressive parameters.

D.5.1 Data generation

Unlike real-world acquired signals, which usually include uncertainty due to unidentified external noise sources or lack of knowledge about their underlying generating process, synthetic signals can be generated in a controlled environment. On the one hand, synthetic data is usually a simplified version of real-world data, hence not always accurately representing the real world. On the other hand, the mechanisms used to generate synthetic data are fully known. This is especially convenient in studies involving irreversible transformations such as quantization since there is no other unidentified source of error than the quantization and reconstruction methods at work. In this paper, we use synthetically generated multivariate time series to delimit, as much as possible, the dependencies between the reconstruction error and the reconstruction method used. Moreover, the synthetic data we use is complex enough to accommodate the smooth behavior as well as the spatiotemporal dependencies assumed in Sec. D.3.1.

First, we generate $M = 288$ sequences composed of $N = 2$ series of $T_* = 100$ knots each, where every knot has been computed recurrently from a stable vector autoregressive VAR(1) process $\boldsymbol{z}_{m,t} = \boldsymbol{\Phi}\boldsymbol{z}_{m,t-1} + \mathbf{w}_{m,t}$ with autoregressive parameters $\boldsymbol{\Phi} \in \mathbb{R}^2$, being $[\boldsymbol{\Phi}]_{1,2} = 0$ to avoid self feedback, Gaussian innovation $\mathbf{w}_{m,t} \sim \mathcal{N}\left(\mathbf{0}, \mathrm{diag}([0.1, 0.1]^\top)\right)$ and where $\boldsymbol{z}_{m,0} = \mathbf{0}_2$ for all $m \in \mathbb{N}^{[1,M]}$. In this way, the resulting sequences of knots inherit the spatiotemporal relations of the VAR(1) generator process, as illustrated in Fig. D.6. Second, each series of knots is arranged uniformly, with a unitary period $1/\nu^* = 1$, and is interpolated with cubic natural splines, which are the smoothest interpolators in terms of the $\rho = 2$ roughness [6]. Lastly, the resulting continuous signals are uniformly sampled at a frequency $\nu = R\nu^* = R$ with a natural oversampling ratio R spanning from 1 to 5 and quantized by a uniform midtread quantizer, with a half-step size $\epsilon \in \{0.1, 0.08, 0.06, 0.04, 0.02\}$. As a result, we generate a dataset consisting of 288 multivariate time series with elements of the form $\tilde{y}_{m,t}^{(n)} = \mathrm{round}(z_{m,t}^{(n)}/2\epsilon)2\epsilon$ for every $n \in \{1, 2\}$, $m \in \mathbb{N}^{[1,288]}$ and $t \in \mathbb{N}^{[1,T]}$, where $T = 100R$.

D.5.2 Experimental setup

The dataset is divided into 196 samples for training, 64 for validation, and 32 for testing. Then, to avoid data leaking, the whole dataset is standardized along each time series with the mean $\gamma_{\mathrm{tr}}^{(n)}$ and standard deviation $\sigma_{\mathrm{tr}}^{(n)}$ of the training partition. Explicitly, $y_{m,t}^{(n)} = (\tilde{y}_{m,t}^{(n)} - \gamma_{\mathrm{tr}}^{(n)})/\sigma_{\mathrm{tr}}^{(n)}$ and $\epsilon^{(n)} = \epsilon/\sigma_{\mathrm{tr}}^{(n)}$ for all $n \in \{1, 2\}$, $m \in \mathbb{N}^{[1,288]}$ and $t \in \mathbb{N}^{[1,100R]}$, resulting in quantization intervals as described in Sec. D.3.2. The

Figure D.7: Training-validation curves for the consistent variation of the RNN-based policy. The loss metric is the average total cost per function section computed from the per time-series cost presented in Table D.1.

RNN-based policy architecture consists of 2 stacked gated recurrent units (GRU) layers [51], with a latent state of size $L = 48$ and an input of size 32. The input and output layers are set as linear layers matching the required dimensionality, see Fig. D.5 for a visual representation, which accordingly varies depending on whether we reconstruct multivariate time series as a multivariate signal or as separate individual univariate signals. Lastly, the RNN-based policies have been trained using the adaptive moments (Adam) optimizer [143], with an exponential decay rate for the first moment estimates of $\beta_1 = 0.9$ and the second-moment estimates of $\beta_2 = 0.999$, without weight decay, a learning rate of 0.002 and with a gradient norm clip value of 0.1 over mini-batches of 32 samples.

D.5.3 Results and discussion

The training-validation curves obtained for the consistent variation of the RNN-based policy, from data acquired with $R = 1$ and $\epsilon = 0.1$, are presented in Fig. D.7. As expected, the consistent RNN-based policy outperforms the consistent myopic policy for both the multivariate and univariate reconstruction cases. Moreover, the multivariate reconstruction outperforms the univariate one, illustrating the model capabilities to successfully exploit spatial dependencies (and not only the temporal ones) from multivariate time series.

Regarding the reconstruction error-rate decay behavior, we have computed the average MSE over all time series in the test set. This has been done for every considered policy architecture (once they are trained, if applicable) and policy variation. As shown in Fig. D.8, we have presented the results by grouping the consistent and smoothing policy variations versus the quantization half-step size ϵ (with fixed $R = 1$) or versus the oversampling ratio R (with fixed $\epsilon = 0.1$). The slope of the

(a) Consistent policy variation.

(b) Smoothing policy variation.

Figure D.8: Average MSE decay curves over the test set in log-log scale. The fitting curves are log-linear functions adjusted from the multivariate average MSE values. The shaded area represents the confidence in the slope of the fit. The values in Fig. D.8b have been computed for an $\eta = 0.001$.

fitting curve reflects the average error-rate decay in logarithmic scale incurred by the multivariate reconstructions of the RNN-based policies. In the case of the consistent policy variation, the module of the quotient between the slope obtained by varying the quantization step size (0.99 ± 0.09) over the slope obtained by increasing the oversampling ratio (-0.812 ± 0.008) is 1.22 ± 0.11. This quotient ratio is only (at most within) two standard deviations from a unit value, suggesting that the balanced error-rate decay property can be achieved for consistent reconstructions even under zero-delay requirements. As to the smoothing policy variation (with slope -0.435 ± 0.002), we are interested in its relative performance with respect to the consistent one. Here, the quotient ratio between the slopes of the fits is 1.87 ± 0.02, meaning that the error-rate decay for the consistent policy variation is nearly doubled with respect to the smoothing one. Again as expected, a property of offline consistency manifests even under zero-delay requirements.

D.6 Conclusion

This paper designs a method for zero-delay signal reconstruction from streaming multivariate time series of quantization intervals, both for the smoothing and consistent discussed variants. The error-rate decay behavior of the proposed methods has been empirically analyzed showing that the error incurred by the consistent variant decreases nearly twice (in logarithmic scale) as fast as the error incurred by the (non-consistent) smoothing variant, as the oversampling ratio is increased. Moreover, the former error rate is practically the same as the rate the error decreases when the quantization step size is reduced, which is advantageous from an implementation perspective. Finally, we also observe that successfully exploiting the spatiotemporal dependencies of the streaming quantized data, in this case via policy training, helps to reduce the average signal reconstruction error.

Bibliography

[1] E. Ruiz-Moreno and B. Beferull-Lozano, "Tracking of quantized signals based on online kernel regression," in *2021 IEEE 31st International Workshop on Machine Learning for Signal Processing (MLSP)*, pp. 1–6, IEEE, 2021.

[2] E. Ruiz-Moreno and B. Beferull-Lozano, "An online multiple kernel paralleliz-able learning scheme," *IEEE Signal Processing Letters*, 2023.

[3] E. Ruiz-Moreno, L. M. López-Ramos, and B. Beferull-Lozano, "A trainable approach to zero-delay smoothing spline interpolation," *IEEE Transactions on Signal Processing*, vol. 71, pp. 4317–4329, 2023.

[4] E. Ruiz-Moreno, L. M. López-Ramos, and B. Beferull-Lozano, "Zero-delay consistent signal reconstruction from streamed multivariate time series," *arXiv preprint arXiv:2308.12459*, 2023.

[5] S. Mollaebrahim, B. Beferull-Lozano, and E. R. Moreno, "Decentralized sub-space projection for asymmetric sensor networks," in *2020 IEEE 92nd Vehic-ular Technology Conference (VTC2020-Fall)*, pp. 1–5, IEEE, 2020.

[6] S. N. Wood, *Generalized additive models: an introduction with R.* chapman and hall/CRC, 2006.

[7] B. Lim and S. Zohren, "Time-series forecasting with deep learning: a sur-vey," *Philosophical Transactions of the Royal Society A*, vol. 379, no. 2194, p. 20200209, 2021.

[8] Q. T. Ain, M. Ali, A. Riaz, A. Noureen, M. Kamran, B. Hayat, and A. Rehman, "Sentiment analysis using deep learning techniques: a review," *International Journal of Advanced Computer Science and Applications*, vol. 8, no. 6, 2017.

[9] H. Sharma, M. Agrahari, S. K. Singh, M. Firoj, and R. K. Mishra, "Image captioning: a comprehensive survey," in *2020 International Conference on Power Electronics & IoT Applications in Renewable Energy and its Control (PARC)*, pp. 325–328, IEEE, 2020.

[10] R. L. Plackett, "Some theorems in least squares," *Biometrika*, vol. 37, no. 1/2, pp. 149–157, 1950.

[11] M. H. Hayes, *Statistical digital signal processing and modeling.* John Wiley & Sons, 1996.

[12] J. Kivinen, A. J. Smola, and R. C. Williamson, "Online learning with kernels," *IEEE transactions on signal processing*, vol. 52, no. 8, pp. 2165–2176, 2004.

[13] C. De Boor and C. De Boor, *A practical guide to splines*, vol. 27. springer-verlag New York, 1978.

[14] A. Agrawal, B. Amos, S. Barratt, S. Boyd, S. Diamond, and J. Z. Kolter, "Differentiable convex optimization layers," *Advances in neural information processing systems*, vol. 32, 2019.

[15] H. M. Gomes, J. Read, A. Bifet, J. P. Barddal, and J. Gama, "Machine learning for streaming data: state of the art, challenges, and opportunities," *ACM SIGKDD Explorations Newsletter*, vol. 21, no. 2, pp. 6–22, 2019.

[16] I. Roşu, "Fast and slow informed trading," *Journal of Financial Markets*, vol. 43, pp. 1–30, 2019.

[17] M. S. Elbamby, C. Perfecto, M. Bennis, and K. Doppler, "Toward low-latency and ultra-reliable virtual reality," *IEEE Network*, vol. 32, no. 2, pp. 78–84, 2018.

[18] H. Boström, S. F. Andler, M. Brohede, R. Johansson, A. Karlsson, J. Van Laere, L. Niklasson, M. Nilsson, A. Persson, and T. Ziemke, "On the definition of information fusion as a field of research," 2007.

[19] L. Bottou, F. E. Curtis, and J. Nocedal, "Optimization methods for large-scale machine learning," *Siam Review*, vol. 60, no. 2, pp. 223–311, 2018.

[20] A. Mokhtari, S. Shahrampour, A. Jadbabaie, and A. Ribeiro, "Online optimization in dynamic environments: Improved regret rates for strongly convex problems," in *2016 IEEE 55th Conference on Decision and Control (CDC)*, pp. 7195–7201, 2016.

[21] R. Nishihara, P. Moritz, S. Wang, A. Tumanov, W. Paul, J. Schleier-Smith, R. Liaw, M. Niknami, M. I. Jordan, and I. Stoica, "Real-time machine learning: The missing pieces," in *Proceedings of the 16th Workshop on Hot Topics in Operating Systems*, pp. 106–110, 2017.

[22] S. Van Vaerenbergh and I. Santamaría, "Online regression with kernels," *Regularization, Optimization, Kernels, and Support Vector Machines*, pp. 477–501, 2014.

[23] K. Slavakis, P. Bouboulis, and S. Theodoridis, "Online learning in reproducing kernel hilbert spaces," in *Academic Press Library in Signal Processing*, vol. 1, pp. 883–987, Elsevier, 2014.

[24] G. Wahba, *Spline models for observational data*. SIAM, 1990.

[25] J. M. de Carvalho and J. V. Hanson, "Real-time interpolation with cubic splines and polyphase networks," *Canadian Electrical Engineering Journal*, vol. 11, no. 2, pp. 64–72, 1986.

[26] R. Debski, "Real-time interpolation of streaming data," *Computer Science*, vol. 21, no. 4, 2020.

[27] N. T. Thao and M. Vetterli, "Deterministic analysis of oversampled a/d conversion and decoding improvement based on consistent estimates," *IEEE Transactions on signal processing*, vol. 42, no. 3, pp. 519–531, 1994.

[28] I. Jovanovic and B. Beferull-Lozano, "Oversampled a/d conversion and error-rate dependence of nonbandlimited signals with finite rate of innovation," *IEEE Transactions on Signal Processing*, vol. 54, no. 6, pp. 2140–2154, 2006.

[29] A. Gersho, "Principles of quantization," *IEEE Transactions on circuits and systems*, vol. 25, no. 7, pp. 427–436, 1978.

[30] S. Shalev-Shwartz *et al.*, "Online learning and online convex optimization," *Foundations and Trends® in Machine Learning*, vol. 4, no. 2, pp. 107–194, 2012.

[31] Ó. Fontenla-Romero, B. Guijarro-Berdiñas, D. Martinez-Rego, B. Pérez-Sánchez, and D. Peteiro-Barral, "Online machine learning," in *Efficiency and Scalability Methods for Computational Intellect*, pp. 27–54, IGI global, 2013.

[32] P. Craven and G. Wahba, "Smoothing noisy data with spline functions: estimating the correct degree of smoothing by the method of generalized cross-validation," *Numerische mathematik*, vol. 31, no. 4, pp. 377–403, 1978.

[33] M. Unser, "Splines: A perfect fit for signal and image processing," *IEEE Signal processing magazine*, vol. 16, no. 6, pp. 22–38, 1999.

[34] C. Kelley, *Iterative Methods for Optimization*, vol. 18. SIAM, 1999.

[35] S. Ruder, "An overview of gradient descent optimization algorithms," *arXiv preprint arXiv:1609.04747*, 2016.

[36] S. Boyd, L. Xiao, and A. Mutapcic, "Subgradient methods," *lecture notes of EE392o, Stanford University, Autumn Quarter*, vol. 2004, pp. 2004–2005, 2003.

[37] N. Parikh, S. Boyd, *et al.*, "Proximal algorithms," *Foundations and Trends® in Optimization*, vol. 1, no. 3, pp. 127–239, 2014.

[38] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM journal on imaging sciences*, vol. 2, no. 1, pp. 183–202, 2009.

[39] V. Vapnik, "Principles of risk minimization for learning theory," *Advances in neural information processing systems*, vol. 4, 1991.

[40] H. W. Engl, M. Hanke, and A. Neubauer, *Regularization of inverse problems*, vol. 375. Springer Science & Business Media, 1996.

[41] B. Schölkopf, R. Herbrich, and A. J. Smola, "A generalized representer theorem," in *International conference on computational learning theory*, pp. 416–426, Springer, 2001.

[42] J. Mercer, "Xvi. functions of positive and negative type, and their connection the theory of integral equations," *Philosophical transactions of the royal society of London. Series A, containing papers of a mathematical or physical character*, vol. 209, no. 441-458, pp. 415–446, 1909.

[43] J. H. Manton, P.-O. Amblard, *et al.*, "A primer on reproducing kernel hilbert spaces," *Foundations and Trends® in Signal Processing*, vol. 8, no. 1–2, pp. 1–126, 2015.

[44] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.

[45] D. Bertsekas, *Dynamic programming and optimal control: Volume I*, vol. 1. Athena scientific, 2012.

[46] R. Bellman, "Dynamic programming," *Science*, vol. 153, no. 3731, pp. 34–37, 1966.

[47] I. Goodfellow, Y. Bengio, and A. Courville, *Deep learning*. MIT press, 2016.

[48] P. J. Werbos, "Backpropagation through time: what it does and how to do it," *Proceedings of the IEEE*, vol. 78, no. 10, pp. 1550–1560, 1990.

[49] R. Pascanu, T. Mikolov, and Y. Bengio, "On the difficulty of training recurrent neural networks," in *International conference on machine learning*, pp. 1310–1318, Pmlr, 2013.

[50] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[51] K. Cho, B. van Merriënboer, D. Bahdanau, and Y. Bengio, "On the properties of neural machine translation: Encoder–decoder approaches," in *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, pp. 103–111, 2014.

[52] 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, 1977.

[53] H. D. Luke, "The origins of the sampling theorem," *IEEE Communications Magazine*, vol. 37, no. 4, pp. 106–108, 1999.

[54] W. R. Bennett, "Spectra of quantized signals," *The Bell System Technical Journal*, vol. 27, no. 3, pp. 446–472, 1948.

[55] R. M. Gray, "Quantization noise spectra," *IEEE Transactions on information theory*, vol. 36, no. 6, pp. 1220–1244, 1990.

[56] N. T. Thao and M. Vetterli, "Reduction of the mse in r-times oversampled a/d conversion o (1/r) to o (1/r/sup 2/)," *IEEE Transactions on Signal Processing*, vol. 42, no. 1, pp. 200–203, 1994.

[57] T.-T. Nguyen, *Deterministic analysis of oversampled A/D conversion and Sigma Delta modulation, and decoding improvements using consistent estimates*. PhD thesis, Columbia University, 1993.

[58] M. Vetterli, P. Marziliano, and T. Blu, "Sampling signals with finite rate of innovation," *IEEE transactions on Signal Processing*, vol. 50, no. 6, pp. 1417–1428, 2002.

[59] B. Beferull-Lozano and A. Ortega, "Efficient quantization for overcomplete expansions in/spl ropf//sup n," *IEEE Transactions on Information Theory*, vol. 49, no. 1, pp. 129–150, 2003.

[60] B. Schölkopf and A. J. Smola, *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2002.

[61] O. Chapelle, V. Vapnik, O. Bousquet, and S. Mukherjee, "Choosing multiple parameters for support vector machines," *Machine learning*, vol. 46, pp. 131–159, 2002.

[62] M. Gönen and E. Alpaydın, "Multiple kernel learning algorithms," *The Journal of Machine Learning Research*, vol. 12, pp. 2211–2268, 2011.

[63] E. Hazan *et al.*, "Introduction to online convex optimization," *Foundations and Trends® in Optimization*, vol. 2, no. 3-4, pp. 157–325, 2016.

[64] F. Orabona, "A modern introduction to online learning," *arXiv preprint arXiv:1912.13213*, 2019.

[65] H. H. Bauschke, R. Goebel, Y. Lucet, and X. Wang, "The proximal average: basic theory," *SIAM Journal on Optimization*, vol. 19, no. 2, pp. 766–785, 2008.

[66] A. Jadbabaie, A. Rakhlin, S. Shahrampour, and K. Sridharan, "Online optimization: Competing with dynamic comparators," in *Artificial Intelligence and Statistics*, pp. 398–406, 2015.

[67] S. A. Bazaz and B. Tondu, "Minimum time on-line joint trajectory generator based on low order spline method for industrial manipulators," *Robotics and Autonomous Systems*, vol. 29, no. 4, pp. 257–268, 1999.

[68] T. Kröger, *On-Line Trajectory Generation in Robotic Systems: Basic Concepts for Instantaneous Reactions to Unforeseen (Sensor) Events*, vol. 58. Springer, 2010.

[69] J. NILSSON, "Real-time control systems with delay," *PhD thesis, Lund Institute of Technology*, 1998.

[70] A. Gambier, "Real-time control systems: a tutorial," in *2004 5th Asian Control Conference (IEEE Cat. No. 04EX904)*, vol. 2, pp. 1024–1031, IEEE, 2004.

[71] C. Schmidt, C. Kottke, V. Jungnickel, and R. Freund, "High-speed digital-to-analog converter concepts," in *Next-Generation Optical Communication: Components, Sub-Systems, and Systems VI*, vol. 10130, pp. 133–141, SPIE, 2017.

[72] E. Waring, "Vii. problems concerning interpolations," *Philosophical transactions of the royal society of London*, no. 69, pp. 59–67, 1779.

[73] I. J. Schoenberg, "Contributions to the problem of approximation of equidistant data by analytic functions," in *IJ Schoenberg Selected Papers*, pp. 3–57, Springer, 1988.

[74] E. Meijering, "A chronology of interpolation: from ancient astronomy to modern signal and image processing," *Proceedings of the IEEE*, vol. 90, no. 3, pp. 319–342, 2002.

[75] L. Schumaker, *Spline functions: basic theory.* Cambridge University Press, 2007.

[76] M. L. Littman, *Algorithms for sequential decision-making.* Brown University, 1996.

[77] A. Blum and Y. Monsour, "Learning, regret minimization, and equilibria," *Algorithmic Game Theory*, 2007.

[78] A. Nosedal-Sanchez, C. B. Storlie, T. C. Lee, and R. Christensen, "Reproducing kernel hilbert spaces for penalized regression: A tutorial," *The American Statistician*, vol. 66, no. 1, pp. 50–60, 2012.

[79] G. Chiarot and C. Silvestri, "Time series compression survey," *ACM Computing Surveys*, vol. 55, no. 10, pp. 1–32, 2023.

[80] G. Kimeldorf and G. Wahba, "Some results on tchebycheffian spline functions," *Journal of mathematical analysis and applications*, vol. 33, no. 1, pp. 82–95, 1971.

[81] A. Koppel, A. S. Bedi, K. Rajawat, and B. M. Sadler, "Optimally compressed nonparametric online learning: Tradeoffs between memory and consistency," *IEEE Signal Processing Magazine*, vol. 37, no. 3, pp. 61–70, 2020.

[82] Z. Wang, K. Crammer, and S. Vucetic, "Breaking the curse of kernelization: Budgeted stochastic gradient descent for large-scale svm training," *The Journal of Machine Learning Research*, vol. 13, no. 1, pp. 3103–3131, 2012.

[83] A. Singh, N. Ahuja, and P. Moulin, "Online learning with kernels: Overcoming the growing sum problem," in *2012 IEEE International Workshop on Machine Learning for Signal Processing*, pp. 1–6, IEEE, 2012.

[84] S. Van Vaerenbergh, J. Via, and I. Santamaría, "A sliding-window kernel rls algorithm and its application to nonlinear channel identification," in *2006 IEEE International Conference on Acoustics Speech and Signal Processing Proceedings*, vol. 5, pp. V–V, IEEE, 2006.

[85] A. S. Bedi, A. Koppel, K. Rajawat, and B. M. Sadler, "Trading dynamic regret for model complexity in nonstationary nonparametric optimization," in *2020 American Control Conference (ACC)*, pp. 321–326, IEEE, 2020.

[86] A. Mokhtari, S. Shahrampour, A. Jadbabaie, and A. Ribeiro, "Online opti- mization in dynamic environments: Improved regret rates for strongly convex problems," in *2016 IEEE 55th Conference on Decision and Control (CDC)*, pp. 7195–7201, IEEE, 2016.

[87] R. Dixit, A. S. Bedi, R. Tripathi, and K. Rajawat, "Online learning with inexact proximal online gradient descent algorithms," *IEEE Transactions on Signal Processing*, vol. 67, no. 5, pp. 1338–1352, 2019.

[88] Y.-L. Yu, "Better approximation and faster algorithm using the proximal av- erage," in *Advances in neural information processing systems*, pp. 458–466, 2013.

[89] A. Beck, *First-order methods in optimization*, vol. 25. SIAM, 2017.

[90] F. Pérez-Cruz and O. Bousquet, "Kernel methods and their potential use in signal processing," *IEEE signal processing magazine*, vol. 21, no. 3, pp. 57–65, 2004.

[91] H. Takeda, S. Farsiu, and P. Milanfar, "Kernel regression for image processing and reconstruction," *IEEE Transactions on image processing*, vol. 16, no. 2, pp. 349–366, 2007.

[92] A. Pozdnoukhov, *Prior Knowledge in Kernel Methods*. PhD thesis, Verlag nicht ermittelbar, 2006.

[93] M. Stone, "Cross-validatory choice and assessment of statistical predictions," *Journal of the royal statistical society: Series B (Methodological)*, vol. 36, no. 2, pp. 111–133, 1974.

[94] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*, vol. 2. MIT press Cambridge, MA, 2006.

[95] F. Orabona and J. Luo, "Ultra-fast optimization algorithm for sparse multi kernel learning," in *Proceedings of the 28th International Conference on Ma- chine Learning*, 2011.

[96] J. A. Tropp, "Greed is good: Algorithmic results for sparse approximation," *IEEE Transactions on Information theory*, vol. 50, no. 10, pp. 2231–2242, 2004.

[97] Y. Engel, S. Mannor, and R. Meir, "The kernel recursive least-squares algo- rithm," *IEEE Transactions on signal processing*, vol. 52, no. 8, pp. 2275–2285, 2004.

[98] C. Richard, J. C. M. Bermudez, and P. Honeine, "Online prediction of time series data with kernels," *IEEE Transactions on Signal Processing*, vol. 57, no. 3, pp. 1058–1067, 2008.

[99] J. B. Souza Filho and P. S. Diniz, "A fixed-point online kernel principal compo- nent extraction algorithm," *IEEE Transactions on Signal Processing*, vol. 65, no. 23, pp. 6244–6259, 2017.

[100] J. B. Souza Filho, L.-D. Van, T.-P. Jung, P. S. Diniz, *et al.*, "Online component analysis, architectures and applications," *Foundations and Trends® in Signal Processing*, vol. 16, no. 3-4, pp. 224–429, 2022.

[101] A. Rahimi and B. Recht, "Random features for large-scale kernel machines," *Advances in neural information processing systems*, vol. 20, 2007.

[102] R. Money, J. Krishnan, and B. Beferull-Lozano, "Sparse online learning with kernels using random features for estimating nonlinear dynamic graphs," *IEEE Transactions on Signal Processing*, 2023.

[103] O. Dekel, S. Shalev-Shwartz, and Y. Singer, "The forgetron: A kernel-based perceptron on a fixed budget," *Advances in neural information processing systems*, vol. 18, 2005.

[104] D. Sahoo, S. C. Hoi, and B. Li, "Online multiple kernel regression," in *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 293–302, 2014.

[105] Y. Shen, T. Chen, and G. Giannakis, "Online ensemble multi-kernel learning adaptive to non-stationary and adversarial environments," in *International Conference on Artificial Intelligence and Statistics*, pp. 2037–2046, PMLR, 2018.

[106] D. Sahoo, S. C. Hoi, and B. Li, "Large scale online multiple kernel regression with application to time-series prediction," *ACM Transactions on Knowledge Discovery from Data (TKDD)*, vol. 13, no. 1, pp. 1–33, 2019.

[107] U. Marteau-Ferey, D. Ostrovskii, F. Bach, and A. Rudi, "Beyond least-squares: Fast rates for regularized empirical risk minimization through self-concordance," in *Conference on learning theory*, pp. 2294–2340, PMLR, 2019.

[108] S. Shalev-Shwartz and S. Ben-David, *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.

[109] C. A. Micchelli, M. Pontil, and P. Bartlett, "Learning the kernel function via regularization.," *Journal of machine learning research*, vol. 6, no. 7, 2005.

[110] Y. Shen, T. Chen, and G. B. Giannakis, "Random feature-based online multi-kernel learning in environments with unknown dynamics," *The Journal of Machine Learning Research*, vol. 20, no. 1, pp. 773–808, 2019.

[111] J. Gorski, F. Pfeuffer, and K. Klamroth, "Biconvex sets and optimization with biconvex functions: a survey and extensions," *Mathematical methods of operations research*, vol. 66, no. 3, pp. 373–407, 2007.

[112] M. Zinkevich, "Online convex programming and generalized infinitesimal gradient ascent," in *Proceedings of the 20th international conference on machine learning (icml-03)*, pp. 928–936, 2003.

[113] S. P. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.

[114] N. Aronszajn, "Theory of reproducing kernels," *Transactions of the American mathematical society*, vol. 68, no. 3, pp. 337–404, 1950.

[115] S. Shalev-Shwartz and Y. Singer, "Efficient learning of label ranking by soft projections onto polyhedra," *Journal of Machine Learning Research*, 2006.

[116] W. Wang and M. A. Carreira-Perpinán, "Projection onto the probability simplex: An efficient algorithm with a simple proof, and an application," *arXiv preprint arXiv:1309.1541*, 2013.

[117] M. Blondel, A. Fujino, and N. Ueda, "Large-scale multiclass support vector machine training via euclidean projection onto the simplex," in *2014 22nd International Conference on Pattern Recognition*, pp. 1289–1294, IEEE, 2014.

[118] A. D. Mishra and D. Garg, "Selection of best sorting algorithm," *International Journal of intelligent information Processing*, vol. 2, no. 2, pp. 363–368, 2008.

[119] J. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra, "Efficient projections onto the l 1-ball for learning in high dimensions," in *Proceedings of the 25th international conference on Machine learning*, pp. 272–279, 2008.

[120] J. Nocedal and S. J. Wright, *Numerical optimization*. Springer, 1999.

[121] A. Uncini, *Fundamentals of adaptive signal processing*. Springer, 2015.

[122] C. M. Bishop and N. M. Nasrabadi, *Pattern recognition and machine learning*, vol. 4. Springer, 2006.

[123] O. Kosheleva and V. Kreinovich, "Why physical processes are smooth or almost smooth: A possible physical explanation based on intuitive ideas behind energy conservation," *Mathematical Structures and Modeling*, 2021.

[124] H. S. Wilf, *Algorithms and complexity*. AK Peters/CRC Press, 2002.

[125] C. E. Rasmussen, "Gaussian processes in machine learning," in *Summer school on machine learning*, pp. 63–71, Springer, 2004.

[126] Q. Lu, G. V. Karanikolas, and G. B. Giannakis, "Incremental ensemble gaussian processes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.

[127] M. Unser and T. Blu, "Generalized smoothing splines and the optimal discretization of the wiener filter," *IEEE Transactions on Signal Processing*, vol. 53, no. 6, pp. 2146–2159, 2005.

[128] K. Frankish and W. M. Ramsey, *The Cambridge handbook of artificial intelligence*. Cambridge University Press, 2014.

[129] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *nature*, vol. 521, no. 7553, pp. 436–444, 2015.

[130] H. Salehinejad, S. Sankar, J. Barfett, E. Colak, and S. Valaee, "Recent advances in recurrent neural networks," *arXiv preprint arXiv:1801.01078*, 2017.

[131] C. H. Reinsch, "Smoothing by spline functions," *Numerische mathematik*, vol. 10, no. 3, pp. 177–183, 1967.

[132] W. Fan, C.-H. Lee, and J.-H. Chen, "A realtime curvature-smooth interpolation scheme and motion planning for cnc machining of short line segments," *International Journal of Machine Tools and Manufacture*, vol. 96, pp. 27–46, 2015.

[133] L. Peshkin, N. Meuleau, and L. Kaelbling, "Learning policies with external memory," *arXiv preprint cs/0103003*, 2001.

[134] A. M. Schäfer, *Reinforcement learning with recurrent neural networks*. PhD thesis, Osnabrück, Univ., Diss., 2008, 2008.

[135] M. Zhang, Z. McCarthy, C. Finn, S. Levine, and P. Abbeel, "Learning deep neural network policies with continuous memory states," in *2016 IEEE international conference on robotics and automation (ICRA)*, pp. 520–527, IEEE, 2016.

[136] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in pytorch," *NIPS 2017 Autodiff Workshop*, 2017.

[137] V. Monga, Y. Li, and Y. C. Eldar, "Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing," *IEEE Signal Processing Magazine*, vol. 38, no. 2, pp. 18–44, 2021.

[138] OSIsoft, "Pi system," *URL: https://www. osisoft. com/pi-system*, 1980.

[139] A. Hebrail, Georges & Berard, "Individual household electric power consumption." UCI Machine Learning Repository, 2012.

[140] G. S. Sampaio, A. R. de Aguiar Vallim Filho, L. S. da Silva, and L. A. da Silva, "Prediction of motor failure time using an artificial neural network," *Sensors*, vol. 19, p. 4342, Oct. 2019.

[141] J. Huan, J. S. Bernstein, P. Difuntorum, N. V. R. Masna, N. Gravenstein, S. Bhunia, and S. Mandal, "A wearable skin temperature monitoring system for early detection of infections," *IEEE Sensors Journal*, vol. 22, no. 2, pp. 1670–1679, 2022.

[142] P. Bodik, W. Hong, C. Guestrin, S. Madden, M. Paskin, and R. Thibaux, "Intel berkeley research lab data," *URL: http://db. csail. mit. edu/labdata/labdata. html*, 2004.

[143] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.

[144] B. Efron and R. J. Tibshirani, *An introduction to the bootstrap*. CRC press, 1994.

[145] S. Rangan and V. K. Goyal, "Recursive consistent estimation with bounded noise," *IEEE Transactions on Information Theory*, vol. 47, no. 1, pp. 457–464, 2001.

[146] V. K. Goyal, J. Kovačević, and J. A. Kelner, "Quantized frame expansions with erasures," *Applied and Computational Harmonic Analysis*, vol. 10, no. 3, pp. 203–233, 2001.

[147] J. Kovacevic, V. K. Goyal, and M. Vetterli, "Fourier and wavelet signal processing," *Fourier Wavelets. org*, pp. 1–294, 2013.

[148] T. Marcucci, P. Nobel, R. Tedrake, and S. Boyd, "Fast path planning through large collections of safe boxes," *arXiv preprint arXiv:2305.01072*, 2023.

[149] I. J. Schoenberg, "On interpolation by spline functions and its minimal properties," in *On Approximation Theory/Über Approximationstheorie*, pp. 109–129, Springer, 1964.

[150] P. J. Green and B. W. Silverman, *Nonparametric regression and generalized linear models: a roughness penalty approach*. Crc Press, 1993.

[151] B. Widrow and I. Kollár, *Quantization noise: roundoff error in digital computation, signal processing, control, and communications*. Cambridge University Press, 2008.