



HAL
open science

Risk Estimate under a Nonstationary Autoregressive Model for Data-Driven Reproduction Number Estimation

Barbara Pascal, Samuel Vaïter

► **To cite this version:**

Barbara Pascal, Samuel Vaïter. Risk Estimate under a Nonstationary Autoregressive Model for Data-Driven Reproduction Number Estimation. 2024. hal-04705913

HAL Id: hal-04705913

<https://hal.science/hal-04705913v1>

Preprint submitted on 23 Sep 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Risk Estimate under a Nonstationary Autoregressive Model for Data-Driven Reproduction Number Estimation

Barbara Pascal *Member, IEEE*, Samuel Vaïter.

Abstract—COVID-19 pandemic has brought to the fore epidemiological models which, though describing a rich variety of behaviors, have previously received little attention in the signal processing literature. During the pandemic, several works successfully leveraged state-of-the-art signal processing strategies to robustly infer epidemiological indicators despite the low quality of COVID-19 data. In the present work, a novel nonstationary autoregressive model is introduced, encompassing, but not reducing to, one of the most popular models for the propagation of viral epidemics. Using a variational framework, penalized likelihood estimators of the parameters of this new model are designed. In practice, the main bottleneck is that the estimation accuracy strongly depends on hyperparameters tuning. Without available ground truth, hyperparameters are selected by minimizing specifically designed data-driven oracles, used as proxy for the estimation error. Focusing on the nonstationary autoregressive Poisson model, the Stein’s Unbiased Risk Estimate formalism is generalized to construct asymptotically unbiased risk estimators based on the derivation of an original autoregressive counterpart of Stein’s lemma. The accuracy of these oracles and of the resulting estimates are assessed through intensive Monte Carlo simulations on synthetic data. Then, elaborating on recent epidemiological models, a novel weekly scaled Poisson model is proposed, enabling to better account for intrinsic variability of the contamination while being robust to reporting errors. Finally, the overall data-driven procedure is particularized to the estimation of COVID-19 reproduction number and exemplified on real COVID-19 infection counts in different countries and at different stages of the pandemic, demonstrating its ability to yield consistent estimates.

Index Terms—Autoregressive models, hyperparameter selection, Stein’s Unbiased Risk Estimate, Poisson noise, Variational estimators, Epidemiology, Reproduction Number, COVID-19.

I. INTRODUCTION

Context. Inverse problems are ubiquitous in signal and image processing [1–4], with a wealth of domains of application as diverse as nonlinear physics [5, 6], astronomy [7], hyperspectral imaging [8], tomography [9], cardiology [10] and epidemiology [11]. A general inverse problem consists in estimating underlying quantities of interest from direct or indirect observations. The intricate measurement processes

necessary to obtain *physical*¹ observations can be the source of several difficulties in estimating the quantities of interest, among which: acquisition performed in a transformed domain requiring backward transformation to access the quantity of interest [9], linear or nonlinear deformation of observations through the measurement process [12, 13], and corruption by stochastic perturbations, either stemming from the physical phenomenon at stake or from the measurement device [2]. Mathematically, without loss of generality, the quantity of interest and observations can be represented by real-valued vectors $\mathbf{X} \in \mathbb{R}^T$ and $\mathbf{Y} \in \mathbb{R}^S$ respectively, with $T, S \in \mathbb{N}^*$, and a generic inverse problem writes

$$\mathbf{Y} \sim \mathcal{B}(\mathbf{A}(\mathbf{X})) \quad (1)$$

where $\mathbf{A} : \mathbb{R}^T \rightarrow \mathbb{R}^S$ is a possibly nonlinear and non-invertible continuous transformation and \mathcal{B} is a stochastic degradation, possibly *data-dependent*, i.e., neither additive nor multiplicative. For example, in low-photon imaging [4], \mathbf{A} is a singular blur operator and the observations are corrupted by Poisson noise, i.e., for $s \in \{1, \dots, S\}$, the random variable Y_s follows a Poisson distribution of intensity $(\mathbf{A}\mathbf{X})_s$.

Most inverse problems of the form (1) are *ill-posed*, e.g., when $S < T$ leading to non-injective \mathbf{A} , hence compromising uniqueness of the quantity of interest given observations under Model (1), or *ill-conditioned*, e.g., when \mathbf{A} is a linear operator with a large conditioning number, inducing numerical instability which, in the presence of noise, translates into prohibitively large estimation variance. Numerous strategies have been designed to address these challenging limitations, which can be divided into two categories: supervised machine learning [14–16] and unsupervised variational and Bayesian techniques [2, 17–19]. The present work focuses on the later category, which is particularly adapted when tackling fundamental research problems [5, 6] or newly emerging phenomena [11], for which annotated data are not available. While the negative log-likelihood associated to Model (1) provides a measure of the fidelity of the observations to the model, a regularization term is designed to lift up the ambiguity in the solution of ill-posed and ill-conditioned problems, e.g., leveraging negative log-prior on the quantity of interest [20, 21]. Balancing the fidelity to the data and the

Manuscript received September 23, 2024. *Corresponding author: Barbara Pascal.*

B. Pascal is with Nantes Université, École Centrale Nantes, CNRS, LS2N, UMR 6004, F-44000 Nantes, France (e-mail: barbara.pascal@cnrs.fr); S. Vaïter is with CNRS, Université Côte d’Azur, LJAD, Nice, France (e-mail: samuel.vaïter@cnrs.fr).

B. Pascal is supported by ANR-23-CE48-0009 “*OptiMoCSI*” and S. Vaïter by ANR-18-CE40-0005 “*GraVa*”. The authors thank Patrice Abry for insightful discussions.

¹*Physical* is to be understood in a very broad sense, encompassing biological, epidemiological as well as human data.

regularity constraints then amounts to compute the maximum a posteriori estimate of the quantity of interest as

$$\widehat{\mathbf{X}}(\mathbf{Y}; \boldsymbol{\lambda}) \in \underset{\mathbf{X} \in \mathbb{R}^T}{\text{Argmin}} \mathcal{D}(\mathbf{Y}, \mathbf{A}\mathbf{X}) + \mathcal{R}(\mathbf{X}; \boldsymbol{\lambda}), \quad (2)$$

where \mathcal{D} is the negative log-likelihood used as a measure of discrepancy between the observations and the model, \mathcal{R} is the negative log-prior penalizing highly irregular solution, and $\boldsymbol{\lambda} \in \Lambda$ is a set of hyperparameters controlling the level of regularity enforced in the estimate. For example, in the study of solid friction targeting the characterization of the stick-slip regime [5, 6], observations consists in a force signal measured across time corrupted by independent identically distributed (i.i.d.) additive Gaussian noise of variance σ^2 . Hence, $T = S$ and $\mathbf{A} = \mathbf{I}_T$, with the data fidelity term in (2) reducing to the negative log-likelihood of a Gaussian random vector of mean \mathbf{X} and scalar covariance matrix $\sigma^2 \mathbf{I}_T$, that is $\mathcal{D}(\mathbf{Y}, \mathbf{A}\mathbf{X}) = \|\mathbf{Y} - \mathbf{X}\|_2^2 / \sigma^2$. Physicists expect the stick-slip regime to induce an almost piecewise linear force signal, a behavior that is favored in the estimate through the penalization $\mathcal{R}(\mathbf{X}; \lambda) = \lambda \|\mathbf{D}_2 \mathbf{X}\|_1$, where $\mathbf{D}_2 : \mathbb{R}^T \rightarrow \mathbb{R}^{T-2}$ is the discrete Laplacian operator.² the ℓ_1 -norm enforces sparsity of the second order derivative of the estimate, and the level of sparsity is controlled by the regularization parameter $\lambda > 0$. As in most ill-posed inverse problems, the tuning of λ is key to obtain an accurate estimate: for small λ , some noise remains, while large λ might cause significant information loss due to over-regularization.

Related works. Ideally, the hyperparameters would be selected by minimizing the *estimation error*, that is by choosing

$$\lambda^\dagger \in \underset{\lambda \in \Lambda}{\text{Argmin}} \left\| \widehat{\mathbf{X}}(\mathbf{Y}; \lambda) - \overline{\mathbf{X}} \right\|_2^2. \quad (3)$$

Though, in practice, the ground truth is not available, and hence the estimation error cannot be evaluated. To tackle this issue, several classes of hyperparameter tuning strategies have been developed during the past decades. For additive Gaussian noise models, leading to quadratic data-fidelity terms, regularized with quadratic penalizations, the L-curve criterion consists in selecting the regularization level by close inspection of the plot of the residual against the regularity of the solution as the regularization parameter is varied [22]. Not only the L-curve method is restricted to Gaussian models under Tikhonov regularization, but also it has been shown to be inaccurate when the targeted ground truth is very smooth [23] and to behave inconsistently as the size of the problem increases [24]. An alternative way is to model the hyperparameters as random variables and to estimate them via hierarchical Bayesian techniques [2, 20]. Such strategies come at the price of an extra complexity, as they require to specify the a priori hyperparameters distribution, and a significant computational cost, as Monte Carlo sampling is often necessary. Recent deep learning methods take advantage of large training databases to learn the hyperparameters of variational estimators [25–27]. While being very accurate as soon as enough annotated data are available, these methods are

not adapted to tackle fundamentally new problems for which the generation of training database would either be too costly or even not possible due to lack of available expert knowledge. A widely used class of unsupervised methods consists in the construction of an *oracle* \mathcal{O} not depending explicitly on the ground truth, which can thus be evaluated in practice, and whose minimization yields an approximation of the optimal hyperparameter in terms of the true error (3) defined as

$$\lambda_{\mathcal{O}} \in \underset{\lambda \in \Lambda}{\text{Argmin}} \mathcal{O}(\mathbf{Y}; \lambda) \quad (4)$$

where $\Lambda \subset \mathbb{R}^L$ denotes the set of admissible hyperparameters. For linear parametric estimates under additive Gaussian noise hypothesis, the Generalized Cross Validation strategy [28] consists in minimizing a data-dependent criterion constructed as the ratio between the residual sum of squares and the trace of a model-dependent linear operator. Though, it has been shown that the Generalized Cross Validation function might not have a unique properly defined minimizer [29], leading the Generalized Cross Validation method to fail catastrophically by producing grossly underestimated regularization parameters. Moreover, it is not flexible enough to handle sparsity-inducing penalizations, which are very popular in the inverse problem literature [30]. Among the methods relying on the design of a tractable oracle, *Stein's Unbiased Risk Estimate* based strategies, elaborating on the seminal work [31] to construct an approximation of the estimation risk, were initially formulated for i.i.d. Gaussian noise model, but have then been extended to more general noise models [32], notably including a data-dependent Poisson contribution [33–35]. In the past decades, Stein-based strategies have demonstrated their ability to provide accurate hyperparameter selection in numerous applications, reaching state-of-the-art performance in inverse problem resolution, both in the variational framework as multispectral image deconvolution [36], denoising of force signal in nonlinear physics [6], and more recently in unsupervised deep learning for image denoising [37].

Contributions and outline. The recent COVID-19 pandemic crisis has triggered massive research efforts on epidemic modeling and surveillance, way beyond the scientific community of epidemiologists [38–45]. Notably, the challenging estimation of the COVID-19 transmissibility in real-time, and with high accuracy despite the low quality of data collected day-by-day by health agencies, has been reformulated as a look-alike inverse problem, enabling to leverage the state-of-the-art variational estimators to get very accurate estimates of the *reproduction number*, a crucial indicator quantifying the intensity of an epidemic [11, 38]. The popular model for viral epidemics proposed in [46] states that the number of new infections at time t , Z_t , follows a Poisson distribution whose time-varying intensity is the product of the *global infectiousness*, defined as a weighted sum of past infection counts $\Phi_t(\mathbf{Z}) = \sum_{s \geq 1} \varphi_s Z_{t-s}$, and of the effective reproduction number at time t , R_t ,

$$Z_t | Z_1, \dots, Z_{t-1} \sim \mathcal{P}(\Phi_t(\mathbf{Z})R_t). \quad (5)$$

Model (5) is very reminiscent of Problem (1), the unknown quantity \mathbf{X} being (R_1, \dots, R_T) , the role of \mathbf{A} being played by

² \mathbf{D}_2 consists in the discrete second order derivative defined such that $(\mathbf{D}_2 \mathbf{X})_t = X_{t+2} - 2X_{t+1} + X_t, \forall t \in \{1, \dots, T-2\}$.

the linear operator:

$$(R_1, \dots, R_T) \mapsto (\Phi_1(\mathbf{Z})R_1, \dots, \Phi_T(\mathbf{Z})R_T) \quad (6)$$

and the degradation \mathcal{B} consisting in data-dependent Poisson noise. Elaborating on this formal resemblance, the variational framework (2) has been fruitfully leveraged to design COVID-19 instantaneous reproduction number R_t estimators [11, 38]. Up to now, the fine-tuning of the regularization parameters of these variational estimators has been done manually, based on expert knowledge, which not only impairs the analysis of huge amount of data, but also might reflect subjective biases of the users. To derive a fully data-driven hyperparameter selection strategy, the main challenge lies in the design of an adapted oracle \mathcal{O} , for example of a Stein estimator. Indeed, meticulous comparison of Models (1) and (5) shows that the autoregressive nature of the epidemiological model (5), which induce a dependency of the linear operator $(R_1, \dots, R_T) \mapsto (\Phi_1(\mathbf{Z})R_1, \dots, \Phi_T(\mathbf{Z})R_T)$ in the observation vector \mathbf{Z} , definitively excludes direct use of generalized Stein Unbiased Risk Estimators, requiring \mathbf{A} to be statistically independent of \mathbf{Z} , if not deterministic [6, 31, 32, 47]. Developing a novel Stein paradigm, adapted to autoregressive models, is a challenging, though crucial, step toward the design of fully data-driven, hence objective, strategies for reproduction number estimation.

Section II first proposes a formal description of a new class of models, namely the generalized *nonstationary autoregressive models*, providing a general framework encompassing Model (5); then the variational framework is leveraged to design estimators of generalized nonstationary autoregressive models unknown parameters under several commonly encountered noise distributions. The proposed original autoregressive Stein paradigm is developed in Section III, a novel Stein's-type lemma is first derived in Section III-B for generalized nonstationary autoregressive models involving Poisson noise, and then used to derive a prediction and an estimation unbiased risk estimators in Section III-C, finally Finite Differences and Monte Carlo strategies are implemented to yield tractable *Autoregressive Poisson Unbiased Risk Estimates*. The accuracy of the derived risk estimates is supported by intensive numerical simulations on synthetic data, presented in Section IV. Then, in Section V, a novel *weekly scaled Poisson* epidemiological model, accounting more precisely for the intrinsic variability of the pathogen propagation while being robust to administrative noise, is introduced. Finally, the Autoregressive Poisson Unbiased Risk Estimate is particularized to this new model to design an original data-driven COVID-19 reproduction number estimator, which is exemplified on real data from different countries worldwide and at different pandemic stages.

Notations. \mathbb{R} denotes the set of real numbers, \mathbb{R}_+ the non-negative real numbers and \mathbb{R}_+^* the positive real numbers. \mathbb{N} denotes the set of nonnegative integers and \mathbb{N}^* the positive integers. Matrices are denoted in roman bold characters, e.g., \mathbf{L} , vectors in upper case sans serif bold characters, e.g., \mathbf{Y} , and scalars in upper case sans serif plain characters, e.g., Y . To avoid unnecessary complications, deterministic and random variables are denoted in the same font; explanations are provided in case the context induces any ambiguity. For

$\mathbf{Y} \in \mathbb{R}^T$ a vector of length T , $\text{diag}(\mathbf{Y}) \in \mathbb{R}^{T \times T}$ denotes the diagonal square matrix of size T , with diagonal consisting in the components of \mathbf{Y} . The entrywise product (resp. division) between two vectors is denoted by \odot (resp. \oslash).

II. NONSTATIONARY AUTOREGRESSIVE MODELS

A. Observation model

The proposed *nonstationary autoregressive* model encompasses both the standard autoregressive model of finite order and the epidemiological model introduced in [46] and generalizes them in several directions. First, the memory term Ψ_t is not necessarily a linear function, it is only assumed to be *causal*, that is depending only on past observations, and *smooth*. Second, observations are no longer restricted to follow an independent Gaussian distribution with constant variance, but instead any distribution with prescribed mean, possibly depending on additional time-varying parameters. Finally, and this is the major originality, observations are externally driven by a time-varying *reproduction coefficient*.

Definition 1 (Driven autoregressive model). Let $T \in \mathbb{N}^*$ be a time horizon, $\bar{\mathbf{X}} = (\bar{X}_1, \dots, \bar{X}_T) \in \mathbb{R}_+^T$ a time-varying reproduction coefficient, $Y_0 \in \mathbb{R}_+^*$ an initial state, and for each $t \in \{1, \dots, T\}$, $\Psi_t : \mathbb{R}^{t-1} \rightarrow \mathbb{R}$ a smooth function, with by convention $\Psi_1 = Y_0$. Observations $\mathbf{Y} = (Y_1, \dots, Y_T)$ follow a *driven autoregressive* model with reproduction coefficient $\bar{\mathbf{X}}$ and memory functions $\{\Psi_t, t = 1, \dots, T\}$ if and only if

$$\forall t \in \{1, \dots, T\}, \quad Y_t \sim \mathcal{B}_{\alpha_t}(\bar{X}_t \Psi_t(Y_1, \dots, Y_{t-1})), \quad (7)$$

where $\mathcal{B}_\alpha(U)$ denotes a probability distribution of mean $U \in \mathbb{R}$ depending on an additional parameter $\alpha \in \mathbb{R}$.

The memory functions $\{\Psi_t, t = 1, \dots, T\}$ encapsulate how the memory of past observations impacts the process at time t . In this work, they are assumed to be perfectly known. As an external source of nonstationarity, the parameter of the probability distribution α_t is allowed to vary with time.

Example 1 (Driven linear autoregressive model). The class of *linear* driven autoregressive models corresponds to *linear* functions $\{\Psi_t, t = 1, \dots, T\}$ defined as

$$\Psi_t(Y_1, \dots, Y_{t-1}) = \sum_{s=1}^{\min(\tau, t-1)} \psi_s Y_{t-s} \quad (8)$$

where $\tau \in \mathbb{N}^*$ is a finite memory horizon. The sequence $\{\psi_s\}_{s=1}^\tau$, encoding the dynamical characteristics of the system, then fully characterizes the entire family of memory functions. If, moreover, $\{\psi_s\}_{s=1}^\tau$ is normalized, that is if $\sum_{s=1}^\tau \psi_s = 1$, then the global trend in the behavior of Y_t is governed by \bar{X}_t : if $\bar{X}_t > 1$, Y_t is exponentially growing, while if $\bar{X}_t < 1$, Y_t decreases exponentially fast.

Linear driven autoregressive models are widely used in epidemiology, where the sequence $\{\psi_s\}_{s=1}^\tau$ corresponds to the *serial interval distribution*, accounting for the randomness of the time delay between primary and secondary infections, and \bar{X}_t corresponds to the *effective reproduction number* [38, 46]. Figure 1 provides a synthetic example of linear driven autoregressive observations under the Poisson model (11), with ψ

corresponding to a discretized Gamma distribution of mean 6.6 and standard deviation 3.5 truncated at $\tau = 25$, mimicking the serial interval function of COVID-19 [48, 49] used in [11, 38, 46]. The three time periods corresponding to $\bar{X}_t > 1$ are represented as light blue areas (first row), and results in three temporally separated bumps in Y_t (second row).

Remark (Autoregressive model of order τ). The driven autoregressive model of Definition 1 encompasses the *standard* autoregressive model of order $\tau \in \mathbb{N}^*$ defined as

$$Y_t = \sum_{s=1}^{\tau} \psi_s Y_{t-s} + \Xi_t, \quad \Xi_t \sim \mathcal{N}(0, \alpha^2) \quad (9)$$

where $(\Xi_t)_{t \in \mathbb{N}^*}$ is a sequence of i.i.d. Gaussian variables of zero mean and variance α^2 , for some $\alpha > 0$ [50, Section 2.2]. This standard autoregressive model indeed corresponds to the *linear* model described in Example 1 with *constant* reproduction coefficient $\bar{X}_t = 1$, and Gaussian noise with constant variance $\alpha_t^2 = \alpha^2$.

It is worth noting that the present work focuses on autoregressive processes which are *driven* by an external *unknown and time-varying* reproduction coefficient, which constitutes a paradigm drastically different from the thoroughly studied autoregressive processes, preventing from using the standard tool described, e.g., in [50].

Example 2 (Noise models). Numerous probability distributions \mathcal{B} are encountered in the inverse problem literature [4, 34, 51–53]. Three major representative examples adapted to the driven autoregressive model introduced in Definition 1 are: i) the additive Gaussian noise of variance α_t^2

$$Y_t \mid Y_1, \dots, Y_{t-1} \sim \mathcal{N}(\bar{X}_t \Psi_t(\mathbf{Y}), \alpha_t^2) \quad (10)$$

where $\mathcal{N}(U, \alpha^2)$ denotes the Gaussian distribution of mean U and variance α^2 ;

ii) the scaled Poisson distribution with scale parameter $\alpha_t > 0$

$$Y_t \mid Y_1, \dots, Y_{t-1} \sim \alpha_t \mathcal{P}\left(\frac{\bar{X}_t \Psi_t(\mathbf{Y})}{\alpha_t}\right), \quad (11)$$

where $\mathcal{P}(U)$ denotes the Poisson distribution of intensity U , which has mean and variance both equal to U ; [3]

iii) the multiplicative Gamma noise of shape parameter $\alpha_t > 0$

$$Y_t \mid Y_1, \dots, Y_{t-1} \sim \mathcal{G}\left(\alpha_t, \frac{\bar{X}_t \Psi_t(\mathbf{Y})}{\alpha_t}\right), \quad (12)$$

where $\mathcal{G}(\alpha, U)$ refers to the Gamma distribution of shape parameter α and scale parameter U ; [3]

B. Variational estimators

Notations. From now on, for the sake of compactness, the driving term at t will be denoted $\Psi_t(\mathbf{Y}) := \Psi_t(Y_1, \dots, Y_{t-1})$ and the collection of all terms will be referred to as $\Psi(\mathbf{Y}) = (\Psi_1(\mathbf{Y}), \dots, \Psi_T(\mathbf{Y}))$.

Given some observations $\mathbf{Y} = (Y_1, \dots, Y_T)$, the most straightforward way to estimate the time-varying reproduction

³By convention, whatever $\alpha > 0$, if $U \leq 0$, $\mathcal{P}(U)$ and $\mathcal{G}(\alpha, U)$ are Dirac distributions, i.e., $Y_t = 0$ deterministically.

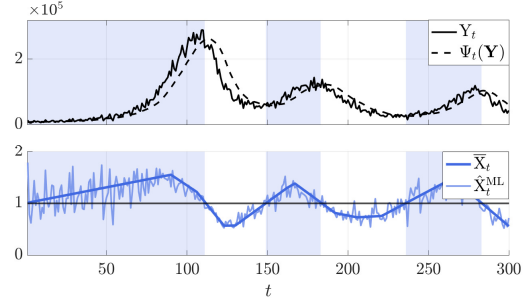


Fig. 1: Driven Poisson autoregressive data. The process Y_t (top plot, solid curve) follows Model (7) with piecewise linear reproduction coefficient X_t (bottom row, deep blue curve); linear memory functions (8) with horizon $\tau = 25$; Poisson degradation (11), with *constant* scale parameter $\alpha_t \equiv 10^3$, initialized at $Y_0 = 10^4$ and run for $T = 300$ steps. Blue areas corresponds to $X_t \geq 1$, driving exponential growth of Y_t .

coefficient $\bar{\mathbf{X}} = (\bar{X}_1, \dots, \bar{X}_T)$ consists in maximizing the likelihood associated with Model (7), yielding the Maximum Likelihood (ML) estimator:

$$\hat{\bar{\mathbf{X}}}^{\text{ML}} = \underset{\mathbf{X} \in \mathbb{R}^T}{\text{argmin}} \mathcal{D}_\alpha(\mathbf{Y}, \mathbf{X} \odot \Psi(\mathbf{Y})) \quad (13)$$

where the discrepancy function $\mathcal{D}_\alpha(\mathbf{Y}, \mathbf{X} \odot \Psi(\mathbf{Y})) = -\ln(\mathbb{P}(\mathbf{Y}|\mathbf{X}; \alpha))$ is the opposite log-likelihood of Model (7).⁴

Example 3 (Discrepancies, Example 2 continued). Under the additive Gaussian noise Model (10) the discrepancy is quadratic

$$-\ln(\mathbb{P}(\mathbf{Y}|\mathbf{X}; \alpha)) = \sum_{t=1}^T \frac{1}{\alpha_t^2} (Y_t - X_t \Psi_t(\mathbf{Y}))^2. \quad (14)$$

Under the scaled Poisson noise Model (11), the discrepancy coincides with the so-called *Kullback-Leibler divergence*

$$-\ln(\mathbb{P}(\mathbf{Y}|\mathbf{X}; \alpha)) = \sum_{t=1}^T d_{\text{KL}}\left(\frac{Y_t}{\alpha_t} \middle| \frac{X_t \Psi_t(\mathbf{Y})}{\alpha_t}\right), \quad (15)$$

where $d_{\text{KL}}(Y|U) =$

$$\begin{cases} Y \ln\left(\frac{Y}{U}\right) + U - Y & \text{if } Y > 0, U > 0 \\ U & \text{if } Y = 0, U \geq 0 \\ \infty & \text{otherwise.} \end{cases} \quad (16)$$

Finally under the multiplicative Gamma noise Model (12), the discrepancy is, up to a term independent of U , the Itakura-Saito divergence [51]

$$-\ln(\mathbb{P}(\mathbf{Y}|\mathbf{X}; \alpha)) = \sum_{t=1}^T d_{\text{IS}}\left(Y_t | \alpha_t, \frac{X_t \Psi_t(\mathbf{Y})}{\alpha_t}\right) \quad (17)$$

where $d_{\text{IS}}(Y|\alpha, U) =$

$$\begin{cases} \frac{Y}{U} - \alpha \ln\left(\frac{Y}{U}\right) + \ln(\Gamma(\alpha)) + \ln(Y) & \text{if } Y > 0, U > 0 \\ \infty & \text{otherwise,} \end{cases} \quad (18)$$

⁴Note that, according to Model (7), the probability distribution of the random vector $\mathbf{Y} = (Y_1, \dots, Y_T)$ depends on the memory functions and on the initialization Y_0 which are assumed known and deterministic, hence are not mentioned in the conditional probability.

where Γ denotes the Euler gamma function.

For all three models, Gaussian (10), Poisson (11) and Gamma (12), provided that $\Psi_t(\mathbf{Y})$ is positive for all $t \in \{1, \dots, T\}$, the maximum likelihood estimator (13) writes

$$\hat{\mathbf{X}}_t^{\text{ML}} = \mathbf{Y}_t / \Psi_t(\mathbf{Y}). \quad (19)$$

An example of Maximum Likelihood estimate on synthetic observations following the Poisson noise model is provided in Figure 1. Due to the presence of noise in the observations (top plot, black solid curve,) the straightforward Maximum Likelihood estimate of the instantaneous reproduction coefficient (bottom plot, light blue curve) suffers from erratic fluctuations compared to ground truth (bottom plot, deep blue curve). Such noisy estimate severely impairs the diagnostic of exponential growth based on $\mathbf{X}_t > 1$; e.g., around $t = 50$, although ground truth is clearly above one, some values $\hat{\mathbf{X}}_t^{\text{ML}} < 1$ are observed.

Obtaining an accurate estimate of the reproduction coefficient thus requires to use additional information. Widely used strategies consists in enforcing a priori constraints, such as, e.g., piecewise linearity [6, 38] and/or sparsity [11, 54]. To that aim, the variational framework consists in augmenting the negative log-likelihood objective of Equation (13) with a regularization term enforcing a priori desirable properties on the estimate leading to parametric estimators of the form

$$\hat{\mathbf{X}}(\mathbf{Y}; \boldsymbol{\lambda}) \in \underset{\mathbf{X} \in \mathbb{R}^T}{\text{Argmin}} \mathcal{D}_\alpha(\mathbf{Y}, \mathbf{X} \odot \Psi(\mathbf{Y})) + \mathcal{R}(\mathbf{X}; \boldsymbol{\lambda}), \quad (20)$$

where $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_L) \in \Lambda$ is a vector of regularization parameters, balancing the overall regularization level as well as the relative importance of the different constraints encoded in the penalization. Commonly used regularization terms are composite [13, 55, 56] and are expressed as

$$\mathcal{R}(\mathbf{X}; \boldsymbol{\lambda}) = \sum_{\ell=1}^L \lambda_\ell \|\mathbf{L}_\ell \mathbf{X}\|_{q_\ell}^{q_\ell} \quad (21)$$

where for each $\ell \in \{1, \dots, L\}$, \mathbf{L}_ℓ is a linear operator, q_ℓ a positive exponent, and $\lambda_\ell \geq 0$ is a regularization parameter balancing the importance of the ℓ th constraint with respect to the other constraints in (21) and to the data-fidelity term of Equation (20). Each term of the functional enforces a specific constraint, hence enabling to take into account several regularity and sparsity properties simultaneously. For example, when choosing the discrete Laplacian $\mathbf{L} = \mathbf{D}_2$, $q_\ell = 1$, favors sparsity of the second order temporal derivative, and hence results in *piecewise linear* estimates, while $q = 2$ yields smooth estimates. The penalized likelihood strategy sketched in Equation (20) is highly flexible and adapts to a large collection of noise models and constraints, hence, by favoring a priori behavior, it has the ability to provide consistent and accurate regularized estimates. The excellent performances of variational estimators are at the price of a cautious fine-tuning of the regularization parameters associated to each term of the penalization. Not only this task is very tough to perform manually but also, and more importantly, in practice ground truth is not available and it is necessary to resort to data-driven oracles to approach optimal hyperparameters.

III. UNBIASED RISK ESTIMATORS

A. General framework

Given an observation model and a parametric estimator, e.g., a variational estimator of the form (20), the ideal hyperparameter selection strategy would consists in minimizing the *estimation risk*, defined as

$$\mathcal{E}(\hat{\mathbf{X}}(\cdot; \boldsymbol{\lambda}), \bar{\mathbf{X}}) := \mathbb{E}_{\mathbf{Y}} \left[\left\| \hat{\mathbf{X}}(\mathbf{Y}; \boldsymbol{\lambda}) - \bar{\mathbf{X}} \right\|_2^2 \right] \quad (22)$$

where $\mathbb{E}_{\mathbf{Y}}$ denotes the expectation over realizations of \mathbf{Y} . For inverse problems of the form (1) with an ill-conditioned or non-injective operator \mathbf{A} , the estimation risk is potentially numerically instable; an alternative is to shift the emphasis on the reconstruction error and to consider the *prediction risk*

$$\mathcal{P}(\hat{\mathbf{X}}(\cdot; \boldsymbol{\lambda}), \bar{\mathbf{X}}) := \mathbb{E}_{\mathbf{Y}} \left[\left\| \hat{\mathbf{X}}(\mathbf{Y}; \boldsymbol{\lambda}) \odot \Psi(\mathbf{Y}) - \bar{\mathbf{X}} \odot \Psi(\mathbf{Y}) \right\|_2^2 \right]. \quad (23)$$

Both the *estimation* and *prediction* risks depends on the ground truth $\bar{\mathbf{X}}$, which in practice is not available. The purpose of this section is thus to devise *oracles* for the quality an estimate $\hat{\mathbf{X}}$, which are independent of the unknown ground truth, and whose minimization provides approximately *optimal* hyperparameters, where *optimal* is to be understood as reaching low estimation or prediction risk.

B. A novel autoregressive Poisson lemma

The remaining of the paper focuses on the Poisson model (11), which is commonly used for modeling the pathogen spread during epidemics, with the aim of applying the developed tools to the estimation of COVID-19 reproduction number from real infection counts in Section V.

The cornerstone of the design of the Stein's Unbiased Risk Estimate is the seminal Stein's lemma [31], turning an expectation explicitly involving ground truth, into another one in which the explicit dependency is completely removed. Under the Poisson noise model (11), the standard Stein's lemma cannot be used. Further, due to the dependency of the memory term in past observations stated in Equation (7), none of the Stein's lemmas generalized to Poisson noise [32, 33, 35] apply. The major challenge in designing oracles adapted to the *driven autoregressive Poisson* model is thus to derive a new *autoregressive Poisson* Stein's lemma counterpart.

Proper definition of the estimation and prediction risks and formal derivation of the proposed *autoregressive Poisson* lemma require further hypotheses: Assumption 1 ensures integrability of all quantities involved, while the technical Assumption 2, easily checked in the practical application of Section V-B, is required to handle autoregressive models.

Assumption 1. Let $\Theta : \mathbb{R}^T \rightarrow \mathbb{R}$, the real-valued function defined on \mathbb{N}^T as

$$\mathbf{k} \mapsto \Theta(\alpha_1 k_1, \dots, \alpha_T k_T) \bar{\mathbf{X}}_t \Psi_t(\alpha_1 k_1, \dots, \alpha_{t-1} k_{t-1})$$

is summable with respect to the driven autoregressive Poisson distribution obtained plugging the driven autoregression model (7) into the Poisson distribution (11)

$$\mathbb{P}(k_1, \dots, k_T) = \prod_{s=1}^T \frac{(\bar{X}_s \Psi_s(\mathbf{Y}))^{k_s}}{k_s!} e^{-\bar{X}_s \Psi_s(\mathbf{Y})}, \quad (24)$$

where for all $s \in \{1, \dots, T\}$ $Y_s = \alpha_s k_s$.

Assumption 2. For all $t, s \in \{1, \dots, T\}$,

$$\forall \mathbf{Y} \in \mathbb{R}_+^T, \quad |\partial_{Y_t} \Psi_s(\mathbf{Y}) \times \alpha_t| \ll |\Psi_s(\mathbf{Y})|. \quad (25)$$

It is worth insisting on the fact that, due to the dependency of the memory functions in the past observations, the components of \mathbf{Y} are not independent Poisson random variables. This is visible in Equation (24) where, because $\Psi_s(\mathbf{Y})$ depends on Y_1, \dots, Y_{s-1} , the right-hand side is not a separable product of independent Poisson distribution and cannot be reframed as such. This is precisely this major difference with standard inverse problems of the form (1) which impairs the application of the standard Poisson counterpart of Stein's lemma [33-35] 5

Notations. Let $\Theta : \mathbb{R}^T \rightarrow \mathbb{R}$ and $\alpha \in \mathbb{R}^T$, for $t \in \{1, \dots, T\}$ the function Θ^{-t} is defined as

$$\Theta^{-t}(\mathbf{Y}) = \Theta(Y_1, \dots, Y_t - \alpha_t, \dots, Y_T), \quad \mathbf{Y} \in \mathbb{R}^T. \quad (26)$$

Lemma 1 (Autoregressive Poisson lemma). Let $\mathbf{Y} = (Y_1, \dots, Y_T)$ observations following the driven autoregressive model (7) with ground truth time-varying reproduction coefficient $\bar{\mathbf{X}} = (\bar{X}_1, \dots, \bar{X}_T) \in \mathbb{R}_+^T$ and memory functions Ψ_t satisfying Assumption 2 corrupted by scaled Poisson noise (11) of time-varying scale parameter $\alpha = (\alpha_1, \dots, \alpha_T) \in \mathbb{R}_+^T$. Then, for $\Theta : \mathbb{R}^T \rightarrow \mathbb{R}$ satisfying Assumption 1 $\forall t \in \{1, \dots, T\}$,

$$\mathbb{E}_{\mathbf{Y}} [\Theta(\mathbf{Y}) \bar{X}_t \Psi_t(\mathbf{Y})] \stackrel{\alpha \rightarrow 0}{=} \mathbb{E}_{\mathbf{Y}} [\Theta^{-t}(\mathbf{Y}) Y_t]. \quad (27)$$

Proof. Proof of lemma 1 is detailed in Appendix VII. \square

Thanks to Lemma 1, the ground truth-dependent expectation in the left-hand side of (27) is approached by a fully data-dependent one, which will turn out key in the derivation of unbiased risk estimates.

C. Autoregressive Poisson Unbiased Risk Estimators

Expanding the estimation and prediction risks of Equations (22) and (23) respectively, and applying Lemma 1 to remove the explicit dependency in the ground truth, Theorem 1 yields novel estimation and prediction Autoregressive Poisson Unbiased Risk Estimates.

Theorem 1. Let $\mathbf{Y} = (Y_1, \dots, Y_T)$ be observations satisfying the requirements enunciated in Lemma 1. Let $\hat{\mathbf{X}}(\mathbf{Y}; \lambda)$ be a parametric estimator of $\bar{\mathbf{X}}$, such that $\forall t \in \{1, \dots, T\}$, $\forall \lambda \in \Lambda$, $\mathbf{Y} \mapsto \hat{X}_t(\mathbf{Y}; \lambda)$ satisfies Assumption 1. Define the data-dependent prediction risk estimate

$$\begin{aligned} \text{APURE}^{\mathcal{P}}(\mathbf{Y}; \lambda | \alpha) &= \|\hat{\mathbf{X}}(\mathbf{Y}; \lambda) \odot \Psi(\mathbf{Y})\|_2^2 \\ &- 2 \sum_{t=1}^T \hat{X}_t^{-t}(\mathbf{Y}; \lambda) \Psi_t(\mathbf{Y}) Y_t + \sum_{t=1}^T (Y_t^2 - \alpha_t Y_t) \end{aligned} \quad (28)$$

⁵This remark on data-dependent Poisson noise would be true as well for the additive Gaussian and multiplicative Gamma noises of Example 2

where $\hat{X}_t^{-t}(\mathbf{Y}; \lambda) = \hat{X}_t(Y_1, \dots, Y_t - \alpha_t, \dots, Y_T; \lambda)$. Then, $\text{APURE}^{\mathcal{P}}$ is an asymptotically unbiased estimate of the prediction risk in the small scale parameters limit, that is

$$\mathbb{E}_{\mathbf{Y}} [\text{APURE}^{\mathcal{P}}(\mathbf{Y}; \lambda | \alpha)] \stackrel{\alpha \rightarrow 0}{=} \mathcal{P}(\hat{\mathbf{X}}, \bar{\mathbf{X}}). \quad (29)$$

Further assuming that $\forall t \in \{1, \dots, T\}$, $\Psi_t(\mathbf{Y}) \neq 0$, define the data-dependent estimation risk estimate

$$\begin{aligned} \text{APURE}^{\mathcal{E}}(\mathbf{Y}; \lambda | \alpha) &= \|\hat{\mathbf{X}}(\mathbf{Y}; \lambda)\|_2^2 \\ &- 2 \sum_{t=1}^T \frac{\hat{X}_t^{-t}(\mathbf{Y}; \lambda)}{\Psi_t(\mathbf{Y})} Y_t + \sum_{t=1}^T \left(\frac{Y_t^2}{\Psi_t(\mathbf{Y})^2} - \frac{\alpha_t Y_t}{\Psi_t(\mathbf{Y})^2} \right). \end{aligned} \quad (30)$$

Then, $\text{APURE}^{\mathcal{E}}$ is an asymptotically unbiased estimate of the prediction risk in the small scale parameters limit, that is

$$\mathbb{E}_{\mathbf{Y}} [\text{APURE}^{\mathcal{E}}(\mathbf{Y}; \lambda | \alpha)] \stackrel{\alpha \rightarrow 0}{=} \mathcal{E}(\hat{\mathbf{X}}, \bar{\mathbf{X}}). \quad (31)$$

Proof. Proof of Theorem 1 is developed in Appendix VIII. \square

D. Finite difference Monte Carlo estimators

Although fully data-driven with an explicit formula, both the estimation and prediction risk estimates $\text{APURE}^{\mathcal{E}}$ and $\text{APURE}^{\mathcal{P}}$ turn out to be complicated to evaluate in practice. Indeed, they both involve all T functions $\mathbf{Y} \mapsto \hat{X}_t^{-t}(\mathbf{Y})$. Since, in general, the estimator $\hat{\mathbf{X}}(\mathbf{Y}; \lambda)$ is not separable in t , it is thus necessary to evaluate the estimator T times. For parametric estimators designed using the variational framework of Equation (20), the involved minimization can be very costly. Consequently, as T is growing, the computational burden of the direct evaluation of $\text{APURE}^{\mathcal{E}}$ and $\text{APURE}^{\mathcal{P}}$ from (22) and (23) rapidly becomes prohibitive. To circumvent this difficulty, the Finite Difference and Monte Carlo strategies [36, 57-59] are combined to yield tractable asymptotically unbiased estimation and prediction risk estimates, requiring further assumptions on the parametric estimator $\hat{\mathbf{X}}(\mathbf{Y}; \lambda)$.

Assumption 3. For any hyperparameters $\lambda \in \Lambda$, the function $\mathbf{Y} \mapsto \hat{\mathbf{X}}(\mathbf{Y}; \lambda)$ is continuously differentiable on \mathbb{R}_+^T .

Theorem 2. Let $\mathbf{Y} = (Y_1, \dots, Y_T)$ be observations satisfying the requirements of Lemma 1 and $\hat{\mathbf{X}}(\mathbf{Y}; \lambda)$ be a parametric estimator of $\bar{\mathbf{X}}$ whose components satisfy Assumption 1 as stated in Theorem 1 and satisfying Assumption 3. Let $\zeta \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ a zero-mean Gaussian vector with covariance matrix the identity in dimension T . Define the data-dependent Finite Difference Monte Carlo prediction risk estimate

$$\begin{aligned} \text{APURE}_{\zeta}^{\mathcal{P}}(\mathbf{Y}; \lambda | \alpha) &= \|\hat{\mathbf{X}}(\mathbf{Y}; \lambda) \odot \Psi(\mathbf{Y})\|_2^2 - 2 \sum_{t=1}^T \hat{X}_t(\mathbf{Y}; \lambda) \Psi_t(\mathbf{Y}) Y_t \\ &+ 2 \left\langle \text{diag}(\alpha \odot \Psi(\mathbf{Y})) \partial_{\mathbf{Y}} \hat{\mathbf{X}}[\zeta], \text{diag}(\mathbf{Y}) \zeta \right\rangle + \sum_{t=1}^T (Y_t^2 - \alpha_t Y_t) \end{aligned} \quad (32)$$

where $\partial_{\mathbf{Y}} \hat{\mathbf{X}}[\zeta]$ denotes the differential of $\mathbf{Y} \mapsto \hat{\mathbf{X}}(\mathbf{Y}; \lambda)$ at the current point $(\mathbf{Y}; \lambda)$ applied to the random vector ζ . Then, $\text{APURE}_{\zeta}^{\mathcal{P}}$ is an asymptotically unbiased estimate of the prediction risk in the small scale parameters limit, that is

$$\mathbb{E}_{\mathbf{Y}, \zeta} [\text{APURE}_{\zeta}^{\mathcal{P}}(\mathbf{Y}; \lambda | \alpha)] \stackrel{\alpha \rightarrow 0}{=} \mathcal{P}(\hat{\mathbf{X}}, \bar{\mathbf{X}}). \quad (33)$$

Further assuming that $\forall t \in \{1, \dots, T\}$, $\Psi_t(\mathbf{Y}) \neq 0$, define the data-dependent Finite Difference Monte Carlo estimation risk estimate

$$\begin{aligned} \text{APURE}_\zeta^\mathcal{E}(\mathbf{Y}; \lambda | \alpha) &= \|\widehat{\mathbf{X}}(\mathbf{Y}; \lambda)\|_2^2 - 2 \sum_{t=1}^T \frac{\widehat{X}_t(\mathbf{Y}; \lambda)}{\Psi_t(\mathbf{Y})} Y_t \\ &+ 2 \left\langle \text{diag}(\alpha \cdot / \Psi(\mathbf{Y})) \partial_{\mathbf{Y}} \widehat{\mathbf{X}}[\zeta], \text{diag}(\mathbf{Y}) \zeta \right\rangle \\ &+ \sum_{t=1}^T \left(\frac{Y_t^2}{\Psi_t(\mathbf{Y})^2} - \frac{\alpha_t Y_t}{\Psi_t(\mathbf{Y})^2} \right). \end{aligned} \quad (34)$$

Then, $\text{APURE}_\zeta^\mathcal{E}$ is an asymptotically unbiased estimate of the estimation risk in the small scale parameters limit, that is

$$\mathbb{E}_{\mathbf{Y}, \zeta} \left[\text{APURE}_\zeta^\mathcal{E}(\mathbf{Y}; \lambda | \alpha) \right]_{\alpha \rightarrow 0} = \mathcal{E}(\widehat{\mathbf{X}}, \bar{\mathbf{X}}). \quad (35)$$

Proof. Proof of Theorem 2 is developed in Appendix IX. \square

Using only one realization of the Monte Carlo vector ζ in the evaluation of $\text{APURE}_\zeta^\mathcal{E}$ and $\text{APURE}_\zeta^\mathcal{P}$ might lead to noisy estimates of \mathcal{E} and \mathcal{P} ; using them directly as oracles for hyperparameters selection according to (4) hence might result in suboptimal and unstable hyperparameter choices [60]. To circumvent this issue, it is possible to average over several independent realizations of ζ to stabilize both the risk estimates and the resulting hyperparameter choice [59, 60].

Proposition 1. Let $\mathbf{Y} = (Y_1, \dots, Y_T)$ be observations satisfying the requirements enunciated in Lemma 1 and $\widehat{\mathbf{X}}(\mathbf{Y}; \lambda)$ be a parametric estimator of $\bar{\mathbf{X}}$ satisfying the assumptions listed in Theorem 2. Let $N \in \mathbb{N}^*$ and $(\zeta^{(1)}, \dots, \zeta^{(N)})$ be independent realizations of the Monte Carlo vector $\zeta \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. The robustified risk estimates defined as

$$\begin{aligned} \overline{\text{APURE}_\zeta^\mathcal{E}}^N &= \frac{1}{N} \sum_{n=1}^N \text{APURE}_{\zeta^{(n)}}^\mathcal{E} \\ \overline{\text{APURE}_\zeta^\mathcal{P}}^N &= \frac{1}{N} \sum_{n=1}^N \text{APURE}_{\zeta^{(n)}}^\mathcal{P} \end{aligned} \quad (36)$$

are asymptotically unbiased estimation (resp. prediction) risk estimates.

Proof. By linearity of the expectations $\mathbb{E}_{\mathbf{Y}, \zeta}$ in Equation (36) and of the limit $\alpha \rightarrow 0$ in Equations (33) and (35)

$$\begin{aligned} \mathbb{E}_{\mathbf{Y}, \zeta} \left[\overline{\text{APURE}_\zeta^\mathcal{E}}^N \right]_{\alpha \rightarrow 0} &= \mathcal{E}(\widehat{\mathbf{X}}, \bar{\mathbf{X}}) \\ \mathbb{E}_{\mathbf{Y}, \zeta} \left[\overline{\text{APURE}_\zeta^\mathcal{P}}^N \right]_{\alpha \rightarrow 0} &= \mathcal{P}(\widehat{\mathbf{X}}, \bar{\mathbf{X}}). \end{aligned} \quad (37)$$

\square

IV. APPLICATION TO PIECEWISE LINEAR ESTIMATION

The purpose of this section is twofold: first, to assess the ability of the robustified Finite Difference Monte Carlo risk estimates deriving from Theorem 2 and Proposition 1, to approximate faithfully the true estimation and prediction risks; second, to demonstrate numerically that the hyperparameters selected by minimizing these risk estimates yields accurate estimates of the reproduction coefficients from observations

following the driven autoregressive model with data-dependent Poisson noise (11). To that aim, intensive simulations are run on synthetic data generated according to (7) and (11).

A. Synthetic data

To prepare for the application to epidemiological indicator estimation, developed in Section V the ground truth reproduction coefficient $\bar{\mathbf{X}}$ is designed piecewise linear [11, 38], imitating temporal evolution of the reproduction number of COVID-19 observed from real-world data [61]. All synthetic data in this section are of length $T = 70$ and share the same ground truth, represented by the deep blue curve in Figure 1, bottom plot, alternating expansion and recession phases, as represented by the blue areas indicating exponential growth period characterized by $\bar{X}_t > 1$. The initial state is set according to the observed real COVID-19 infection counts during the imitated period at $Y_0 = 3395$. Mimicking the epidemiological model proposed in [46] particularized to COVID-19 pandemic, the memory functions are chosen *linear*, as described in Example 1, with a constant memory horizon of $\tau = 25$ and a sequence $\{\psi_s\}_{s=1}^\tau$ chosen as the daily discretization of the serial interval distribution of COVID-19 modeled as a Gamma distribution of mean 6.6 days and standard deviation 3.5 days [48, 49]. The scale parameter of the data-dependent Poisson noise is constant through time, that is $\forall t, \alpha_t = \alpha > 0$. Seven values of α logarithmically spaced between $\alpha = 10^2$, corresponding to low noise level, to a very high noise level of $\alpha = 10^5$, are explored. Figure 2 provides in top row examples of synthetic observations with the same underlying ground truth reproduction coefficient, displayed in blue in the second row plots, for three scale parameter values, corresponding, from first to third columns, to low $\alpha = 10^2$, medium $\alpha = 10^3$ and high $\alpha = 10^4$ noise levels.

B. Estimation strategy

The estimation of the reproduction coefficient is performed through a variational procedure consisting in the minimization of the penalized negative log-likelihood functional

$$\widehat{\mathbf{X}}(\mathbf{Y}; \lambda) = \underset{\mathbf{X} \in \mathbb{R}_+^T}{\text{argmin}} \mathcal{D}_\alpha(\mathbf{Y}, \mathbf{X} \odot \Psi(\mathbf{Y})) + \lambda \|\mathbf{D}_2 \mathbf{X}\|_1 \quad (38)$$

where \mathcal{D} is the Kullback-Leibler divergence customized to the driven autoregressive Poisson model as described in Equations (15) and (16), \mathbf{D}_2 is the discrete Laplacian, which, combined with the ℓ_1 norms favors piecewise linear behavior of the estimate, and $\lambda > 0$ is a regularization parameter balancing the data-fidelity and the penalization [11, 38]. As the minimization problem (38) is convex but nonsmooth, it is solved using the Chambolle-Pock proximal algorithm [11, 38, 62] leveraging the closed-form expression of the proximity operator of the Kullback-Leibler divergence of Equation (15) [38].

Proposition 2. Let \mathbf{Y} observations following the driven autoregressive Poisson model described in Section IV-A. Then, the variational estimator $\widehat{\mathbf{X}}(\mathbf{Y}; \lambda)$ defined in Equation (38) satisfies the assumptions enunciated in Theorems 1 and 2. Hence, for any $N \in \mathbb{N}^*$, the data-driven robustified oracles

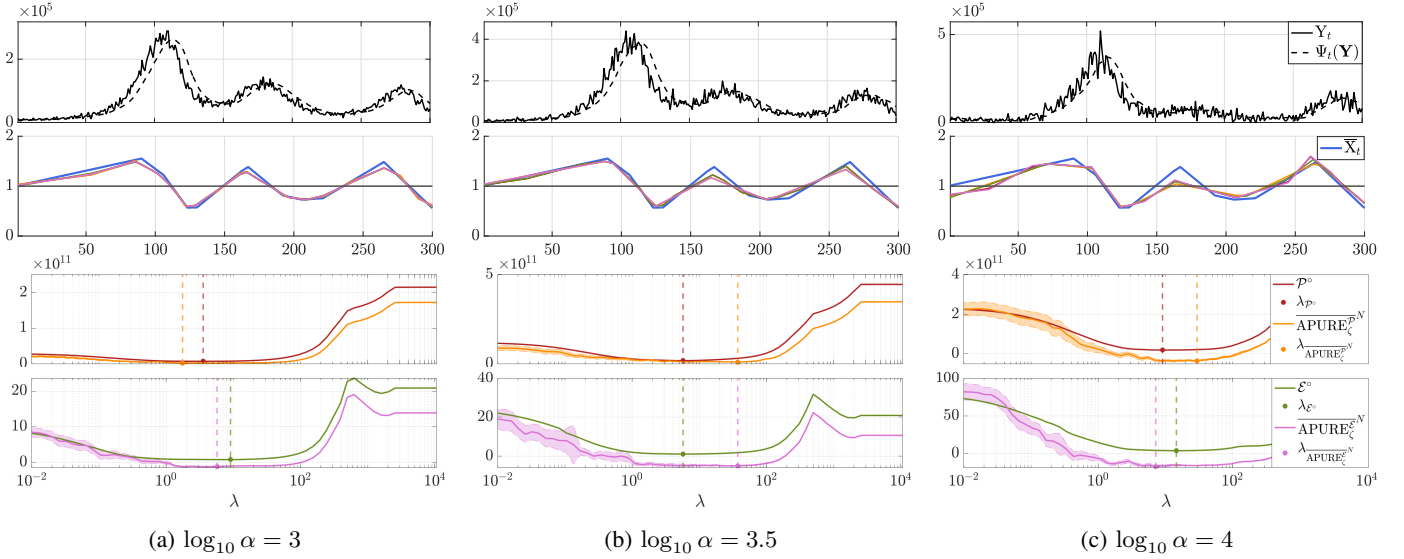


Fig. 2: **Estimation of the reproduction coefficient from synthetic driven autoregressive Poisson observations using the variational estimator (38) coupled with the oracle-based hyperparameter selection strategy (4).** First row: One realization of synthetic observations drawn following the setup described in Section IV-A. Second row: Underlying ground truth reproduction coefficient in blue, and estimates obtained from the variational estimator (38) combined with the oracle-based regularization parameter selection strategy, for the two ground truth-dependent oracles \mathcal{P}° and \mathcal{E}° in red and green respectively, and for the two data-driven oracles $\overline{\text{APURE}}_\zeta^{\mathcal{P}^N}$ and $\overline{\text{APURE}}_\zeta^{\mathcal{E}^N}$ in pink and orange respectively. Third and fourth rows: Robustified data-driven prediction and estimation risk estimates computed on a logarithmic grid of regularization parameters λ , accompanied with their Gaussian confidence regions computed from the $N = 10$ Monte Carlo vector realizations, and exact prediction and estimation errors; vertical dashed lines indicate optimal hyperparameters selected for each oracle.

$\overline{\text{APURE}}_\zeta^{\mathcal{E}^N}$ and $\overline{\text{APURE}}_\zeta^{\mathcal{P}^N}$ are asymptotically unbiased estimators of the estimation and prediction risks respectively.

Given a realization of synthetic observations \mathbf{Y} , the fine-tuning of the regularization parameter λ is performed through the oracle minimization strategy of Equation (4). Four oracles are considered: the exact estimation and prediction errors

$$\begin{aligned} \mathcal{E}^\circ(\mathbf{Y}; \lambda) &= \|\widehat{\mathbf{X}}(\mathbf{Y}; \lambda) - \overline{\mathbf{X}}\|_2^2, \\ \mathcal{P}^\circ(\mathbf{Y}; \lambda) &= \|\widehat{\mathbf{X}}(\mathbf{Y}; \lambda) \odot \Psi(\mathbf{Y}) - \overline{\mathbf{X}} \odot \Psi(\mathbf{Y})\|_2^2 \end{aligned} \quad (39)$$

which, by definition of \mathcal{E} and \mathcal{P} , are ground truth-unbiased estimates of the estimation and prediction risk respectively; and the proposed $\overline{\text{APURE}}_\zeta^{\mathcal{E}^N}$ and $\overline{\text{APURE}}_\zeta^{\mathcal{P}^N}$ introduced in Section III which are fully data-driven. Oracles are minimized through exhaustive search over a logarithmic grid of λ from $10^{-2} \times \text{std}(\mathbf{Y})$ to $10^4 \times \text{std}(\mathbf{Y})$, providing four hyperparameter choices $\lambda_{\mathcal{O}}$, and resulting in four estimates $\widehat{\mathbf{X}}(\mathbf{Y}; \lambda_{\mathcal{O}})$. Numerical simulations aim at assessing the quality of the estimates obtained from the proposed data-driven oracles compared to estimates based on ground truth oracles, which are not usable in real-world problems, such as epidemic monitoring described in Section V.

C. Performance evaluation

For each scale parameter explored, performances are evaluated on a collection of $Q = 10$ independent synthetic observations $\{\mathbf{Y}^{(q)}, q = 1, \dots, Q\}$ generated following the setup of Section IV-A. For $\mathcal{O} \in \{\mathcal{E}^\circ, \mathcal{P}^\circ, \overline{\text{APURE}}_\zeta^{\mathcal{E}^N}, \overline{\text{APURE}}_\zeta^{\mathcal{P}^N}\}$, the

hyperparameters selected through the oracle strategy applied to the q th realization are denoted⁶

$$\lambda_{\mathcal{O}}^{(q)} \in \underset{\lambda \in \mathbb{R}_+}{\text{Argmin}} \mathcal{O}(\mathbf{Y}^{(q)}; \lambda). \quad (40)$$

Several performance criteria are considered. First, the *Minimal Mean Squared Error* over the Q realizations, defined as

$$\text{MMSE} = \frac{1}{Q} \sum_{q=1}^Q \|\widehat{\mathbf{X}}(\mathbf{Y}^{(q)}; \lambda_{\mathcal{O}}^{(q)}) - \overline{\mathbf{X}}\|_2^2, \quad (41)$$

quantifies the overall accuracy of the estimate obtained when selecting hyperparameters so as to minimize the oracle \mathcal{O} . It is accompanied by its 95% Gaussian confidence interval:

$$\text{CI} = \frac{1.96}{\sqrt{Q}} \times \frac{1}{Q} \sum_{q=1}^Q \left(\|\widehat{\mathbf{X}}(\mathbf{Y}^{(q)}; \lambda_{\mathcal{O}}^{(q)}) - \overline{\mathbf{X}}\|_2^2 - \text{MMSE} \right)^2. \quad (42)$$

Let $\langle \widehat{\mathbf{X}} \rangle_{\mathcal{O}}$ the mean estimate obtained using the oracle \mathcal{O} over the Q realizations,⁷ then the Minimal Mean Squared Error

⁶The dependency of $\lambda_{\mathcal{O}}$ in \mathbf{Y} is omitted for readability.

⁷The mean estimate writes $\langle \widehat{\mathbf{X}} \rangle_{\mathcal{O}} = \frac{1}{Q} \sum_{q=1}^Q \widehat{\mathbf{X}}(\mathbf{Y}^{(q)}; \lambda_{\mathcal{O}}^{(q)})$.

Metric	Minimized oracle	$\log_{10} \alpha = 2$	$\log_{10} \alpha = 2.5$	$\log_{10} \alpha = 3$	$\log_{10} \alpha = 3.5$	$\log_{10} \alpha = 4$	$\log_{10} \alpha = 4.5$	$\log_{10} \alpha = 5$
MMSE \pm CI	Estimation error \mathcal{E}°	0.51 ± 0.03	0.59 ± 0.07	1.14 ± 0.36	1.53 ± 0.76	1.69 ± 0.73	1.28 ± 0.36	4.20 ± 3.35
	Prediction error \mathcal{P}°	0.52 ± 0.03	0.61 ± 0.07	1.35 ± 0.62	1.71 ± 0.74	1.85 ± 0.76	1.39 ± 0.38	5.22 ± 4.01
	$\overline{\text{APURE}}_{\zeta}^{\mathcal{P}^{\circ N}}$	0.52 ± 0.03	0.62 ± 0.07	1.37 ± 0.63	1.82 ± 0.74	1.84 ± 0.80	1.46 ± 0.43	20.03 ± 26.27
	$\overline{\text{APURE}}_{\zeta}^{\mathcal{E}^{\circ N}}$	0.52 ± 0.03	0.61 ± 0.07	1.57 ± 0.77	6.39 ± 7.31	6.98 ± 6.25	16.69 ± 19.51	209.61 ± 286.66
Bias – Variance	Estimation error \mathcal{E}°	$0.49 - 0.02$	$0.52 - 0.07$	$0.76 - 0.38$	$0.78 - 0.75$	$0.85 - 0.84$	$0.53 - 0.75$	$1.59 - 2.61$
	Prediction error \mathcal{P}°	$0.49 - 0.03$	$0.51 - 0.09$	$0.64 - 0.71$	$0.72 - 0.99$	$0.76 - 1.08$	$0.48 - 0.91$	$1.11 - 4.11$
	$\overline{\text{APURE}}_{\zeta}^{\mathcal{P}^{\circ N}}$	$0.50 - 0.02$	$0.54 - 0.07$	$0.68 - 0.69$	$0.75 - 1.07$	$0.97 - 0.87$	$0.59 - 0.87$	$2.67 - 17.36$
	$\overline{\text{APURE}}_{\zeta}^{\mathcal{E}^{\circ N}}$	$0.50 - 0.03$	$0.58 - 0.08$	$0.66 - 0.91$	$1.16 - 5.23$	$1.28 - 5.69$	$1.91 - 14.78$	$25.81 - 183.78$

TABLE I: **Performance of the variational estimate (38) combined with oracle-based selection of the regularization parameter.** Ground truth-dependent oracles, \mathcal{E}° and \mathcal{P}° defined in Equation (39), are compared with the proposed fully data-driven robustified $\overline{\text{APURE}}_{\zeta}^{\mathcal{E}^{\circ N}}$ and $\overline{\text{APURE}}_{\zeta}^{\mathcal{P}^{\circ N}}$ derived in Proposed (I), averaged over $N = 10$ realizations of the Monte Carlo vector for seven logarithmically spaced scale parameters α . Second to fifth rows: Minimal Mean Squared Error accompanied with 95% Gaussian confidence intervals. Sixth to ninth rows: Estimation bias and estimation variance. Performances are computed on $Q = 10$ realizations of synthetic observations drawn following the setup described in Section IV-A.

further decomposes into a squared bias term and a variance term, $\text{MMSE} = \text{Bias} + \text{Variance}$, with

$$\text{Bias} = \left\| \langle \widehat{\mathbf{X}} \rangle_{\mathcal{O}} - \overline{\mathbf{X}} \right\|_2^2, \quad (43)$$

$$\text{Variance} = \frac{1}{Q} \sum_{q=1}^Q \left\| \widehat{\mathbf{X}} \left(\mathbf{Y}^{(q)}; \lambda_{\mathcal{O}}^{(q)} \right) - \langle \widehat{\mathbf{X}} \rangle_{\mathcal{O}} \right\|_2^2 \quad (44)$$

reported together with the MMSE for further comparison.

D. Results

Figure 2 compares the four oracle-based estimation strategies on one realization of driven autoregressive Poisson synthetic observations for three values of the scale parameter, corresponding to low, medium and high noise levels. Synthetic observations and memory terms are displayed in the first row; the underlying ground truth reproduction coefficient and its estimates obtained using the four different oracles are plotted in blue on the second row; third (resp. fourth) row compares the prediction (resp. estimation) ground truth and data-dependent oracles, and the associated optimal hyperparameters.

First of all, for the three noise levels, both the prediction and estimation data-dependent oracles, displayed as the orange curves in the third row plots and pink curves in the fourth row plots respectively, approximate very closely the true prediction and estimation errors, displayed as the red curves in the third row plots and green curves in the fourth row plots respectively, throughout the large range of regularization parameter explored. For low and medium noise levels, first and second columns in Figure 2, the four oracles are flat enough in the optimal regularization parameter region so that the small differences between the selected hyperparameters are not visible on the obtained estimates, displayed in the second row, which are all of equal quality. When the noise level gets higher, all oracles are more picked around their minima and the hyperparameters selected by ground truth dependent or data-driven oracles are very close, yielding similarly good estimates. As expected, the 95% Gaussian confidence regions around the robustified data-driven oracles, computed from the

$N = 10$ realizations of the Monte Carlo vector and displayed in deemed colors, thicken as the noise level increases, while remaining of reasonable size, demonstrating the stability of the parameter selection strategy relying on the data-driven oracles.

Table I provides systematic performances computed on $Q = 10$ realizations of the observations for each of the seven noise levels explored. As expected, the larger α , the larger the Minimal Mean Squared Error. Though, the estimation accuracy when using the two ground truth-dependent oracles, second and third row, and the robustified *prediction* unbiased risk estimate, fourth row, increases very slowly with α . Furthermore, at fixed noise level, using either \mathcal{E}° , \mathcal{P}° or $\overline{\text{APURE}}_{\zeta}^{\mathcal{P}^{\circ N}}$ leads to equivalent Minimal Mean Squared Error. This shows, first, that the estimation accuracy is unaltered when replacing the exact *estimation* error \mathcal{E}° by the exact *prediction* error \mathcal{P}° , and second, that the data-driven oracle $\overline{\text{APURE}}_{\zeta}^{\mathcal{P}^{\circ N}}$ yields estimates of very similar quality. Considering the performances of the robustified unbiased *estimation* risk estimate, reported in the fifth row, they remain similarly good as the three other oracles for low to medium noise levels, but then suddenly drop for $\alpha > 10^3$ while the associated 95% Gaussian confidence intervals are also thickening violently. This shows that, despite the robustification procedure of Proposition 1, the data-driven oracle $\overline{\text{APURE}}_{\zeta}^{\mathcal{E}^{\circ N}}$ is unstable, which is probably due to the division by $\Psi_t(\mathbf{Y})$ which order of magnitude varies significantly with t , between 10^3 and 10^5 in these examples, as can be observed on Figure 2, top row. To gain further insight, the decomposition of the Minimal Mean Squared Error into the squared bias and the variance is reported in sixth to ninth rows of Table I. For low noise levels $\alpha < 10^3$, third and fourth columns, the bias, which is intrinsic to all regularized estimation strategies of the form (2), is responsible of almost all the estimation error. When the noise levels exceeds $\alpha = 10^3$, squared bias and variance contributes equally to the estimation error, as observed in fifth to ninth columns. These results advocates, in a practical context, to use preferably the robustified unbiased *prediction* risk estimate which appears both very accurate in the selection

of the optimal regularization parameter and more robust to medium to high noise levels.

V. APPLICATION TO EPIDEMIOLOGY

A major motivation of the present work lies in the need for data-driven hyperparameter fine-tuning strategies for recently proposed COVID-19 reproduction estimators leveraging the variational framework sketched in Equation (2) [11, 38, 61].

A. Weekly scaled Poisson epidemiological model

The considered epidemiological model, briefly introduction in Section I in Equation (5), was initially proposed in [46] and states that, conditionally to past infection counts Z_1, \dots, Z_{t-1} , the number of new infections at time t , denoted Z_t , follows a Poisson distribution of intensity equal to the product of the effective reproduction number at time t , R_t , and the global infectiousness $\Phi_t(\mathbf{Z}) = \sum_{s \geq 1} \varphi_s Z_{t-s}$ with φ the serial interval distribution.⁸ Although very accurate for a posteriori consolidated data [46], the daily version of this model appeared not suitable for real-time COVID-19 daily infection counts, an example of which is displayed in Figure 3, gray curve. Indeed, COVID-19 data are severely corrupted by administrative noise, taking the form of missing counts during week-ends and holidays, large cumulative counts on Mondays, pseudo-seasonalities, erroneous samples, to name but a few. This noise has two major effects: first, reported COVID-19 daily infections counts contain outlier values, notably during week-ends as illustrated in the gray curve in Figure 3; second, the variance of the reported counts reflects both the intrinsic variance of the propagation process and the additional variance induced by the fluctuating administrative delays and errors. To circumvent the presence of outlier samples, a classical strategy in epidemiology is to consider aggregated data [45, 63, 64], e.g., at the scale of the week as illustrated in Figure 3, black curve, which is far smoother and hence more realistic from an epidemiological point of view. The present work proposes to account for the increased variance observed in real COVID-19 infection counts through a constant in time scale parameter $\alpha > 1$. Up to the authors knowledge, scaled Poisson distributions have not yet been explored in the epidemiology literature, making this an original contribution of the present work. Altogether, the proposed *weekly scaled Poisson* epidemiological model writes

$$Z_t | Z_1, \dots, Z_{t-1} \sim \alpha \mathcal{P} \left(\frac{\Phi_t(\mathbf{Z}) R_t}{\alpha} \right). \quad (45)$$

where time instants t corresponds to weeks and the serial interval distribution used to compute $\Phi_t(\mathbf{Z})$ is *weekly* discretized. The weekly discretized COVID-19 serial interval distribution can be obtained by coarsening the daily discretized distribution provided in [48, 49] into a weekly distribution, e.g., by performing an integration over one-week windows using the left rectangle method with a one-day integration step. Under Model (45), the expected number of infections at week t is $\Phi_t(\mathbf{Z}) R_t$, unchanged compared to the standard Poisson

⁸The *serial interval* is the random delay between primary and secondary infections; its distribution encodes the typical time scales of the propagation.

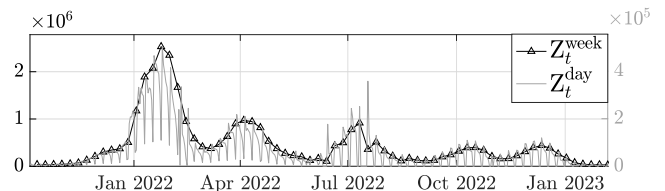


Fig. 3: **COVID-19 daily vs. weekly new infections counts.** French counts from Johns Hopkins University repository, covering 71 weeks from October 4, 2021 to February 6, 2023.

model of Equation (5), but the variance is $\alpha \Phi_t(\mathbf{Z}) R_t$, larger by a factor α than the variance of the standard Poisson model.

B. Data-driven reproduction number estimation strategy

Following [11, 38], R_t is assumed piecewise linear in time, and the COVID-19 *weekly* reproduction number is thus estimated from aggregated infection counts by plugging the weekly discretized propagation model introduced in Equation (45) into the variational estimator of Equation (38) leading to

$$\hat{\mathbf{R}}(\mathbf{Z}; \lambda) = \underset{\mathbf{R} \in \mathbb{R}_+^T}{\operatorname{argmin}} \mathcal{D}_\alpha(\mathbf{Z}, \mathbf{R} \odot \Phi(\mathbf{Z})) + \lambda \|\mathbf{D}_2 \mathbf{R}\|_1. \quad (46)$$

The scale parameter, unknown in practice, is assumed constant in time. Remark that $\Phi_t(\mathbf{Z})$ is linear in \mathbf{Z} , with $\sum_{s \geq 1} \varphi_s = 1$ by normalization of the serial interval distribution, hence if $s > t$, then $\partial_{Z_t} \Phi_s(\mathbf{Z}) = \varphi_{s-t}$ is of order one. Thus, to ensure that α is reflecting the inflated variance observed in COVID-19 data while satisfying Assumption 2, the following data-driven heuristics is proposed $\alpha = 0.1 \times \operatorname{std}(\mathbf{Z})$. The regularization parameter λ controlling the level of regularization in Equation (46) is selected in a data-driven manner through

$$\lambda_{\text{APURE}_\zeta^{\mathcal{P}}^N} \in \underset{\lambda \in \mathbb{R}_+}{\operatorname{Argmin}} \overline{\text{APURE}_\zeta^{\mathcal{P}}^N}(\mathbf{Z}; \lambda) \quad (47)$$

leveraging the robustified unbiased prediction risk estimate, whose accuracy and robustness have been demonstrated on synthetic data in Section IV.

C. Experimental setup

During the entire course of the COVID-19 pandemic, daily infection counts reported by National Health Agencies of 200+ countries worldwide were collected by Johns Hopkins University and made publicly available in an online repository.⁹ To demonstrate the universality of the proposed data-driven reproduction number estimation procedure, four countries, belonging to different continents, and two time periods, corresponding to late and early epidemic stages, are considered. Daily reported counts in France and India, from October 5, 2021 to February 6, 2023, and in Canada and Argentina, from December 22, 2020 to April 25, 2022, are downloaded from JHU repository.⁹ Each sequence of daily counts is then aggregated at the scale of the week, yielding a vector of observed counts \mathbf{Z} of length $T = 70$ weeks. Finally,

⁹<https://coronavirus.jhu.edu/>

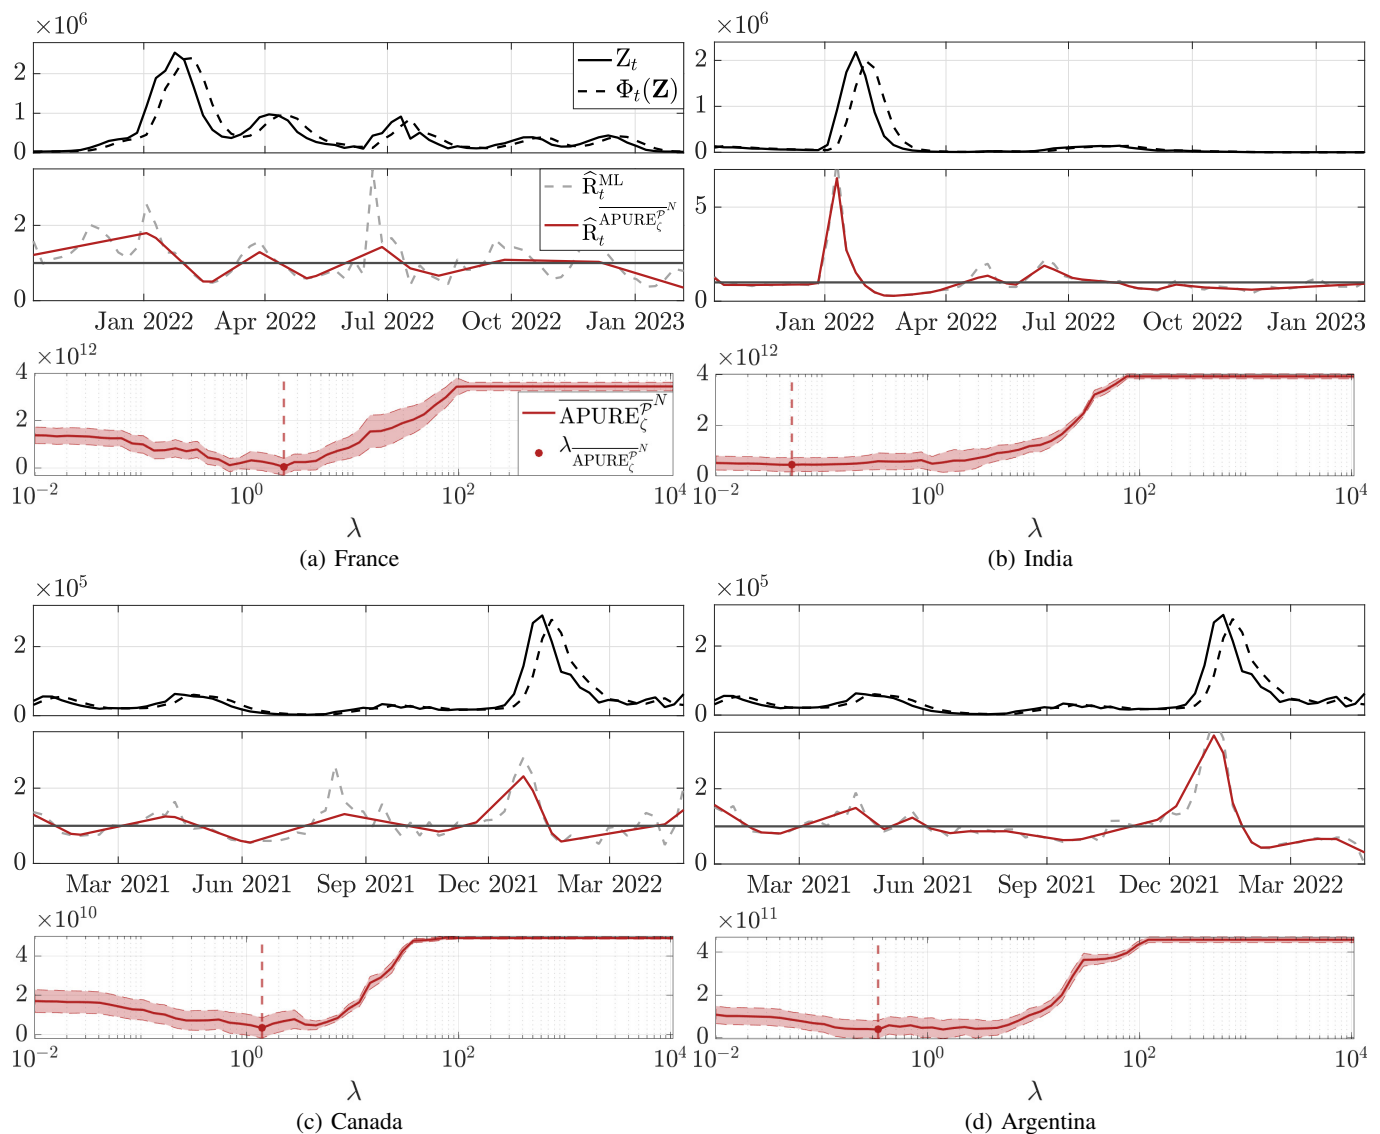


Fig. 4: **Data-driven estimation of the reproduction number of COVID-19 from weekly infection counts.** Top row: weekly aggregated infection counts. Middle row: estimated effective reproduction number with data-driven hyperparameter selection based on the robustified $\overline{\text{APURE}}_{\zeta}^{\mathcal{P}^N}$. Bottom row: robustified autoregressive unbiased prediction risk estimate averaged over $N = 10$ realizations of the Monte Carlo vector with associated 95% Gaussian confidence interval for a logarithmically spaced range of regularization parameters and minimizing regularization parameter obtained through grid search.

the associated global infectiousness $\Phi(\mathbf{Z})$ is computed using $\Phi_t(\mathbf{Z}) = \sum_{s \geq 1} \varphi_s Z_{t-s}$, where φ is the weekly discretized serial interval distribution introduced in Section V-A. Top plots of Figure 4 shows weekly infection counts and associated global infectiousness in France 4a and India 4b, from 2021 to 2023, and Canada 4c and Argentina 4d, from 2020 to 2022.

The minimization in Equation (46) is performed using the Chambolle-Pock algorithm derived in [11, 38], as in Section V-B. Based on Section IV the robustified unbiased prediction risk estimate is computed by averaging over $N = 10$ realizations of the Monte Carlo vector, which yields a robust approximation of the prediction error, and accurate estimates when used as an oracle for hyperparameter selection. The regularization parameter is chosen by solving (47) through an exhaustive search over a logarithmic grid of hyperparameters

λ ranging from $10^{-2} \times \text{std}(\mathbf{Z})$ to $10^4 \times \text{std}(\mathbf{Z})$.

D. Discussion

For each country and pandemic phase configuration, the robustified unbiased prediction risk estimate as a function of the regularization parameter λ , accompanied by the associated 95% Gaussian confidence interval computed from the N realizations of the Monte Carlo vector, is displayed in the bottom plots of Figure 4. The resulting optimal hyperparameter is indicated by the vertical dashed line. For all four configurations, from Figure 4a to 4d, the prediction risk estimate has a clearly identifiable minimizer. The resulting reproduction number estimates are displayed as red curves in the middle plots of Figure 4, together with the naive

Maximum Likelihood Estimator of Equation (13), represented as dashed gray curves. The regularized estimates with data-driven hyperparameter selection, in red, vary more slowly than the Maximum Likelihood estimates, in grey, and are hence more realistic to account for the pandemic spread. This is notably the case for France, Figure 4a, and Canada, Figure 4c, where the Maximum Likelihood estimate present rapid fluctuations leading to significant overestimation of the reproduction number, e.g., in July 2022 in France and August 2021 in Canada. This is all the more important to get accurate estimate of R_t as the sanitary measures have dramatic social and economical impact that decision makers need to balance. In such context, both false alarms and missed events may have highly detrimental consequences. Figure 4b shows that the proposed data-driven reproduction number estimator is able to capture both rapid bursts, corresponding to severe pandemic waves, as experimented by India in January 2022, and smaller waves, e.g., in June 2022, with similar accuracy. Together with the quantitative assessment performed on synthetic data in Section IV, this qualitative assessment of the proposed data-driven reproduction number estimation strategy on real COVID-19 infection counts of different countries in various pandemic stages, demonstrates its ability to be used to monitor closely a viral epidemics, even in a context of degraded reporting suffering from outlier samples and inflated variance.

VI. CONCLUSION AND PERSPECTIVES

A novel driven autoregressive model has been proposed, inspired by state-of-the-art viral epidemics models. The estimation of the parameters of the model, namely the time-varying *reproduction coefficient*, has been framed as a nonstandard, highly nonstationary, inverse problem. The proposed rigorous mathematical formulation enabled to leverage the efficient and versatile variational framework to design accurate reproduction coefficient estimators which are robust to high noise levels in the observations. The third major contribution consisted in the design of asymptotically unbiased risk estimates, which are then plugged into an oracle strategy for fully data-driven fine-tuning of the hyperparameters of the variational estimator, removing the major obstacle to its practical use. The resulting data-driven estimation strategy has been assessed through intensive Monte Carlo simulations on synthetic data. Taking advantage of the proposed extended mathematical framework, the standard epidemiological model for viral epidemics has been enriched to account for the low quality of COVID-19 data, leading to a novel *weekly scaled Poisson* model. Finally, the data-driven estimation procedure is shown to yield very consistent estimation of the COVID-19 reproduction number in various countries and pandemic stages despite the low quality of reported data, demonstrating its practical applicability for epidemic monitoring in a crisis context. The data-driven nature of the proposed epidemiological indicator estimation strategy constitutes a major asset for its dissemination beyond signal processing as its use requires no expert knowledge.

Further work will consists in, first, extending the unbiased risk estimates to other noise models, e.g., to additive Gaussian and multiplicative Gamma noises inspiring from

the generalized Stein estimators [32], in order to further enrich the proposed framework and to be able to extend the developed methodology to other models, in epidemiology or beyond. Second, as in epidemiology the memory functions are often parametric, with parameters encapsulating the pathogen transmission characteristics, the proposed framework will be leveraged to perform simultaneously the extraction of the memory functions parameters and the fine-tuning of the estimator hyperparameters.

For the sake of reproducibility and dissemination, a Matlab toolbox has been made publicly available on the GitHub of the corresponding author.¹⁰

REFERENCES

- [1] J. Idier, *Bayesian approach to inverse problems*. John Wiley & Sons, 2013.
- [2] J.-F. Giovannelli and J. Idier, *Regularization and Bayesian Methods for Inverse Problems in Signal and Image Processing*. Wiley Online Library, 2015.
- [3] X. Wei, H. Van Gorp, L. Gonzalez-Carabarin, D. Freedman, Y. C. Eldar, and R. J. van Sloun, "Deep Unfolding With Normalizing Flow Priors for Inverse Problems," *IEEE Trans. Signal Process.*, vol. 70, pp. 2962–2971, 2022.
- [4] S. Melidonis, P. Dobson, Y. Altmann, M. Pereyra, and K. Zygalakis, "Efficient bayesian computation for low-photon imaging problems," *SIAM J. Imaging Sci.*, vol. 16, no. 3, pp. 1195–1234, 2023.
- [5] J. Colas, N. Pustelnik, C. Oliver, J. Géminard, and V. Vidal, "Nonlinear denoising for solid friction dynamics characterization," *Physical Review E*, vol. 100, no. 032803, pp. 88–89, 2019.
- [6] B. Pascal, N. Pustelnik, P. Abry, G. J.-C., and V. Vidal, "Parameter-free and fast nonlinear piecewise filtering: application to experimental physics," vol. 75, no. 11, pp. 655–671.
- [7] L. Denneulin, M. Langlois, É. Thiébaud, and N. Pustelnik, "RHAPSODIE : Reconstruction of High-contrast Polarized SOURces and Deconvolution for Circumstellar Environments," *Astronomy & Astrophysics*, vol. 653, p. A138, 2021.
- [8] R. A. Borsoi, T. Imbiriba, J. C. M. Bermudez, C. Richard, J. Chanussot, L. Drumetz, J.-Y. Tourneret, A. Zare, and C. Jutten, "Spectral variability in hyperspectral data unmixing: A comprehensive review," *IEEE geoscience and remote sensing magazine*, vol. 9, no. 4, pp. 223–270, 2021.
- [9] M. Savanier, E. Chouzenoux, J.-C. Pesquet, and C. Riddell, "Deep unfolding of the dbfb algorithm with application to roi ct imaging with limited angular density," *IEEE Trans. Computational Imaging*, 2023.
- [10] L. R. Bear, Y. S. Dogrusoz, J. Svehlikova, J. Coll-Font, W. Good, E. van Dam, R. Macleod, E. Abell, R. Walton, R. Coronel, M. Haissaguerre, and R. Dubois, "Effects of ECG Signal Processing on the Inverse Problem of Electrocardiography," in *2018 Comput. Cardiol. Conference (CinC)*, vol. 45, 2018, pp. 1–4.
- [11] B. Pascal, P. Abry, N. Pustelnik, S. Roux, R. Gribonval, and P. Flandrin, "Nonsmooth convex optimization to estimate the Covid-19 reproduction number space-time evolution with robustness against low quality data," *IEEE Transactions on Signal Processing*, vol. 70, pp. 2859–2868, 2022.
- [12] M. Cachia, V. Stergiopoulou, L. Calatroni, S. Schaub, and L. Blanc-Féraud, "Fluorescence image deconvolution microscopy via generative adversarial learning (FluoGAN)," *Inverse Problems*, vol. 39, no. 5, p. 054006, 2023.
- [13] A. Repetti, M. Q. Pham, L. Duval, E. Chouzenoux, and J.-C. Pesquet, "Euclid in a Taxicab: Sparse Blind Deconvolution with Smoothed ℓ_1/ℓ_2 Regularization," *IEEE Signal Process. Lett.*, vol. 22, no. 5, pp. 539–543, 2014.
- [14] S. H. Chan, X. Wang, and O. A. Elgendy, "Plug-and-play ADMM for image restoration: Fixed-point convergence and applications," *IEEE Trans. Comput. Imaging*, vol. 3, no. 1, pp. 84–98, 2016.
- [15] H. Vu, G. Cheung, and Y. C. Eldar, "Unrolling of Deep Graph Total Variation for Image Denoising," in *Proc. Int. Conf. Acoust., Speech Signal Process.* IEEE, 2021, pp. 2050–2054.
- [16] R. Laumont, V. D. Bortoli, A. Almansa, J. Delon, A. Durmus, and M. Pereyra, "Bayesian Imaging Using Plug & Play Priors: When Langevin Meets Tweedie," *SIAM Journal on Imaging Sciences*, vol. 15, no. 2, pp. 701–737, 2022.

¹⁰<https://github.com/bpascal-fr/APURE-Estim-Epi>

- [17] E. Chouzenoux, A. Jeziarska, J.-C. Pesquet, and H. Talbot, "A convex approach for image restoration with exact Poisson–Gaussian likelihood," *SIAM J. on Imaging Sci.*, vol. 8, no. 4, pp. 2662–2682, 2015.
- [18] M. Foare, N. Pustelnik, and L. Condat, "Semi-linearized proximal alternating minimization for a discrete Mumford–Shah model," *IEEE Transactions on Image Processing*, vol. 29, pp. 2176–2189, 2019.
- [19] C. Vacar and J.-F. Giovannelli, "Unsupervised joint deconvolution and segmentation method for textured images: a Bayesian approach and an advanced sampling algorithm," *EURASIP Journal on Advances in Signal Processing*, vol. 2019, pp. 1–17, 2019.
- [20] C. P. Robert, *The Bayesian choice: From Decision-Theoretic Foundations to Computational Implementation*. Springer, 2007, vol. 2.
- [21] R. Gribonval, "Should Penalized Least Squares Regression be Interpreted as Maximum A Posteriori Estimation?" *IEEE Trans. Signal Process.*, vol. 59, no. 5, pp. 2405–2410, 2011.
- [22] P. C. Hansen and D. P. O’Leary, "The use of the L-curve in the regularization of discrete ill-posed problems," *J. Sci. Comput.*, vol. 14, no. 6, pp. 1487–1503, 1993.
- [23] M. Hanke, "Limitations of the L-curve method in ill-posed problems," *BIT Numerical Mathematics*, vol. 36, no. 2, pp. 287–301, 1996.
- [24] C. R. Vogel, "Non-convergence of the L-curve regularization parameter selection method," *Inverse problems*, vol. 12, no. 4, p. 535, 1996.
- [25] C. Bertocchi, E. Chouzenoux, M.-C. Corbineau, J.-C. Pesquet, and M. Prato, "Deep unfolding of a proximal interior point method for image restoration," *Inverse Problems*, vol. 36, no. 3, p. 034005, 2020.
- [26] P. Nguyen, E. Soubies, and C. Chaux, "Map-informed unrolled algorithms for hyper-parameter estimation," in *Proc. Int. Conf. Image Process.* IEEE, 2023, pp. 2160–2164.
- [27] M. Gharbi, E. Chouzenoux, and J.-C. Pesquet, "An unrolled half-quadratic approach for sparse signal recovery in spectroscopy," *Sig. Process.*, vol. 218, p. 109369, 2024.
- [28] G. H. Golub, M. Heath, and G. Wahba, "Generalized cross-validation as a method for choosing a good ridge parameter," *Technometrics*, vol. 21, no. 2, pp. 215–223, 1979.
- [29] A. Thompson, J. Kay, and D. Titterton, "A cautionary note about crossvalidatory choice," *J. Stat. Comput. Simul.*, vol. 33, no. 4, pp. 199–216, 1989.
- [30] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D: nonlinear phenomena*, vol. 60, no. 1-4, pp. 259–268, 1992.
- [31] C. M. Stein, "Estimation of the mean of a multivariate normal distribution," *Ann. Stat.*, pp. 1135–1151, 1981.
- [32] Y. C. Eldar, "Generalized SURE for exponential families: Applications to regularization," *IEEE Trans. Signal Process.*, vol. 57, no. 2, pp. 471–481, 2008.
- [33] F. Luisier, T. Blu, and M. Unser, "Image denoising in mixed Poisson–Gaussian noise," *IEEE Trans. Image Process.*, vol. 20, no. 3, pp. 696–708, 2010.
- [34] Y. Le Montagner, E. D. Angelini, and J.-C. Olivo-Marin, "An unbiased risk estimator for image denoising in the presence of mixed Poisson–Gaussian noise," *IEEE Trans. Image Process.*, vol. 23, no. 3, pp. 1255–1268, 2014.
- [35] J. Li, F. Luisier, and T. Blu, "PURE-LET Image Deconvolution," *IEEE Trans. Image Process.*, vol. 27, no. 1, pp. 92–105, 2018.
- [36] R. Ammanouil, A. Ferrari, D. Mary, C. Ferrari, and F. Loi, "A parallel and automatically tuned algorithm for multispectral image deconvolution," *Monthly Notices of the Royal Astronomical Society*, vol. 490, no. 1, pp. 37–49, 2019.
- [37] D. Chen, J. Tachella, and M. E. Davies, "Robust Equivariant Imaging: A Fully Unsupervised Framework for Learning To Image From Noisy and Partial Measurements," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 5647–5656.
- [38] P. Abry, N. Pustelnik, S. Roux, P. Jensen, P. Flandrin, R. Gribonval, C.-G. Lucas, É. Guichard, P. Borgnat, and N. Garnier, "Spatial and temporal regularization to estimate COVID-19 reproduction number $R(t)$: Promoting piecewise smoothness via convex optimization," *Plos one*, vol. 15, no. 8, p. e0237901, 2020.
- [39] IHME COVID-19 Forecasting Team and S. I. Hay, "COVID-19 scenarios for the United States," *medRxiv*, 2020.
- [40] A. Flahault, "COVID-19 cacophony: is there any orchestra conductor?" *The Lancet*, vol. 395, no. 10229, p. 1037, 2020.
- [41] A. Oustaloup, F. Levron, S. Victor, and L. Dugard, "Non-integer (or fractional) power model to represent the complexity of a viral spreading: Application to the COVID-19," *Annu. Rev. Control*, vol. 52, pp. 523–542, 2021.
- [42] R. K. Nash, P. Nouvellet, and A. Cori, "Real-time estimation of the epidemic reproduction number: Scoping review of the applications and challenges," *PLoS Digit. Health*, vol. 1, no. 6, p. e0000052, 2022.
- [43] J. Paireau, A. Andronico, N. Hozé, M. Layan, P. Crepey, A. Roumagnac, M. Lavielle, P.-Y. Boëlle, and S. Cauchemez, "An ensemble model based on early predictors to forecast COVID-19 health care demand in France," *Proceedings of the National Academy of Sciences*, vol. 119, no. 18, p. e2103302119, 2022.
- [44] S. Bhatia, J. Wardle, R. K. Nash, P. Nouvellet, and A. Cori, "Extending EpiEstim to estimate the transmission advantage of pathogen variants in real-time: SARS-CoV-2 as a case-study," *Epidemics*, p. 100692, 2023.
- [45] R. K. Nash, S. Bhatt, A. Cori, and P. Nouvellet, "Estimating the epidemic reproduction number from temporally aggregated incidence data: A statistical modelling approach and software tool," *PLoS Comput. Biol.*, vol. 19, no. 8, p. e1011439, 2023.
- [46] A. Cori, N. M. Ferguson, C. Fraser, and S. Cauchemez, "A new framework and software to estimate time-varying reproduction numbers during epidemics," *American Journal of Epidemiology*, vol. 178, no. 9, pp. 1505–1512, 2013.
- [47] B. Pascal, S. Vaiter, N. Pustelnik, and P. Abry, "Automated data-driven selection of the hyperparameters for Total-Variation based texture segmentation," *J. Math. Imaging Vis.*, pp. 1–30, 2021.
- [48] D. Cereda, M. Tirani, F. Rovida, V. Demicheli, M. Ajelli, P. Poletti, F. Trentini, G. Guzzetta, V. Marziano, A. Barone *et al.*, "The early phase of the COVID-19 outbreak in Lombardy, Italy," *Preprint arXiv:2003.09320*, 2020.
- [49] F. Riccardi, M. Ajelli, X. D. Andrianou, A. Bella, M. Del Manso, M. Fabiani, S. Bellino, S. Boros, A. M. Urdiales, V. Marziano *et al.*, "Epidemiological characteristics of COVID-19 cases and estimates of the reproductive numbers 1 month into the epidemic, Italy, 28 January to 31 March 2020." *Euro Surveillance*, 2020.
- [50] R. H. Shumway, D. S. Stoffer, and D. S. Stoffer, *Time Series Analysis and Its Applications*. Springer, 2000, vol. 3.
- [51] C. Févotte and A. T. Cemgil, "Nonnegative matrix factorizations as probabilistic inference in composite models," in *2009 17th European Signal Processing Conference*. IEEE, 2009, pp. 1913–1917.
- [52] S. Durand, J. Fadili, and M. Nikolova, "Multiplicative Noise Removal Using L1 Fidelity on Frame Coefficients," *J. Math. Imaging Vision*, vol. 36, pp. 201–226, 2010.
- [53] A. Jeziarska, C. Chaux, J.-C. Pesquet, and H. Talbot, "An EM approach for Poisson–Gaussian noise modeling," in *Proc. Eur. Signal Process. Conf.* IEEE, 2011, pp. 2244–2248.
- [54] M. Gharbi, E. Chouzenoux, J.-C. Pesquet, and L. Duval, "Gpu-based implementations of mm algorithms. application to spectroscopy signal restoration," in *2021 29th European Signal Processing Conference (EUSIPCO)*. IEEE, 2021, pp. 2094–2098.
- [55] B. Jin, D. A. Lorenz, and S. Schiffler, "Elastic-net regularization: error estimates and active set methods," *Inverse Problems*, vol. 25, no. 11, p. 115022, 2009.
- [56] A. Repetti and Y. Wiaux, "Variable metric forward-backward algorithm for composite minimization problems," *SIAM Journal on Optimization*, vol. 31, no. 2, pp. 1215–1241, 2021.
- [57] A. Girard, "A fast 'Monte-Carlo cross-validation' procedure for large least squares problems with noisy data," *Numerische Mathematik*, vol. 56, no. 1, pp. 1–23, 1989.
- [58] S. Ramani, T. Blu, and M. Unser, "Monte-carlo SURE: A black-box optimization of regularization parameters for general denoising algorithms," *IEEE Trans. Image Process.*, vol. 17, no. 9, pp. 1540–1554, 2008.
- [59] C.-A. Deledalle, S. Vaiter, J. Fadili, and G. Peyré, "Stein Unbiased Gradient estimator of the Risk (SUGAR) for multiple parameter selection," *SIAM J. Imaging Sci.*, vol. 7, no. 4, pp. 2448–2487, 2014.
- [60] C.-G. Lucas, B. Pascal, N. Pustelnik, and P. Abry, "Hyperparameter selection for Discrete Mumford–Shah," *Signal, Image and Video Processing*, vol. 17, no. 5, pp. 1897–1904, 2023.
- [61] J. Du, B. Pascal, and P. Abry, "Compared performance of Covid19 reproduction number estimators based on realistic synthetic data," Grenoble, France, Aug.28 - Sept. 1 2023.
- [62] A. Chambolle and T. Pock, "A first-order primal-dual algorithm for convex problems with applications to imaging," *J. Math. Imag. Vis.*, vol. 40, no. 1, pp. 120–145, 2011.
- [63] N. M. Ferguson, Z. M. Cucunubá, I. Dorigatti, G. L. Nedjati-Gilani, C. A. Donnelly, M.-G. Basáñez, P. Nouvellet, and J. Lessler, "Countering the Zika epidemic in Latin America," *Science*, vol. 353, no. 6297, pp. 353–354, 2016.
- [64] K. Charniga, Z. M. Cucunubá, M. Mercado, F. Prieto, M. Ospina, P. Nouvellet, and C. A. Donnelly, "Spatial and temporal invasion dynamics of the 2014–2017 Zika and chikungunya epidemics in Colombia," *PLoS Comput. Biol.*, vol. 17, no. 7, p. e1009174, 2021.

Risk Estimate under a Nonstationary Autoregressive Model for Data-Driven Reproduction Number Estimation – Supplementary material

Barbara Pascal *Member, IEEE*, Samuel Vaiter.

VII. AUTOREGRESSIVE POISSON STEIN-LIKE LEMMA

Proof of driven autoregressive Poisson Stein-like lemma [1](#)

Let $\mathbf{Y} \in \mathbb{R}_+^T$ random observations under the driven generalized autoregressive model [\(7\)](#) with ground truth reproduction number $\bar{\mathbf{X}} \in \mathbb{R}_+^T$ and memory functions $\{\Psi_t, t = 1, \dots, T\}$ satisfying Assumption [2](#) following a scaled Poisson distribution [\(11\)](#) of time-varying modulus $\alpha \in (\mathbb{R}_+^*)^T$, and a function $\Theta : \mathbb{R}_+^T \rightarrow \mathbb{R}$ satisfying Assumption [1](#). First, by Assumption [1](#), the expectation in the left-hand side of Equation [\(27\)](#) is well defined. Hence, using the discrete probabilistic density function of the scaled Poisson distribution leads to

$$\begin{aligned} & \mathbb{E}_{\mathbf{Y}} [\Theta(\mathbf{Y}) \bar{\mathbf{X}}_t \Psi_t(\mathbf{Y})] \\ &= \sum_{k_1=0}^{\infty} \dots \sum_{k_T=0}^{\infty} \Theta(\mathbf{Y}) \bar{\mathbf{X}}_t \Psi_t(\mathbf{Y}) \prod_{s=1}^T \frac{(\bar{\mathbf{X}}_s \Psi_s(\mathbf{Y}))^{k_s}}{k_s!} e^{-\bar{\mathbf{X}}_s \Psi_s(\mathbf{Y})}, \end{aligned} \quad (48)$$

where $\forall t, Y_t = \alpha_t k_t, k_t \in \mathbb{N}$. Since $\Psi_s(\mathbf{Y})$ does not depend on Y_s for $s \geq u$, considering only the sums over k_t, \dots, k_T , the above expression factorizes as

$$\begin{aligned} & \sum_{k_t=0}^{\infty} \dots \sum_{k_T=0}^{\infty} \Theta(\mathbf{Y}) \bar{\mathbf{X}}_t \Psi_t(\mathbf{Y}) \prod_{s=1}^T \frac{(\bar{\mathbf{X}}_s \Psi_s(\mathbf{Y}))^{k_s}}{k_s!} e^{-\bar{\mathbf{X}}_s \Psi_s(\mathbf{Y})} \\ &= \prod_{s=1}^{t-1} \frac{(\bar{\mathbf{X}}_s \Psi_s(\mathbf{Y}))^{k_s}}{k_s!} e^{-\bar{\mathbf{X}}_s \Psi_s(\mathbf{Y})} \\ & \quad \times \sum_{k_t=0}^{\infty} G(\mathbf{Y}) \bar{\mathbf{X}}_t \Psi_t(\mathbf{Y}) \frac{(\bar{\mathbf{X}}_t \Psi_t(\mathbf{Y}))^{k_t}}{k_t!} e^{-\bar{\mathbf{X}}_t \Psi_t(\mathbf{Y})} \end{aligned} \quad (49)$$

where $G(\mathbf{Y}) =$

$$\sum_{k_{t+1}=0}^{\infty} \dots \sum_{k_T=0}^{\infty} \Theta(\mathbf{Y}) \prod_{u=t+1}^T \frac{(\bar{\mathbf{X}}_u \Psi_u(\mathbf{Y}))^{k_u}}{k_u!} e^{-\bar{\mathbf{X}}_u \Psi_u(\mathbf{Y})}. \quad (50)$$

$G(\mathbf{Y})$ is obtained by marginalizing over the variables Y_{t+1}, \dots, Y_T , and thus it depends only on Y_1, \dots, Y_t . Then, the summation over k_t appearing in [\(49\)](#) writes

$$\begin{aligned} & \sum_{k_t=0}^{\infty} G(\mathbf{Y}) \bar{\mathbf{X}}_t \Psi_t(\mathbf{Y}) \frac{(\bar{\mathbf{X}}_t \Psi_t(\mathbf{Y}))^{k_t}}{k_t!} e^{-\bar{\mathbf{X}}_t \Psi_t(\mathbf{Y})} \\ &= \sum_{k_t=0}^{\infty} G(\mathbf{Y}) \times (k_t + 1) \times \frac{(\bar{\mathbf{X}}_t \Psi_t(\mathbf{Y}))^{k_t+1}}{(k_t + 1)!} e^{-\bar{\mathbf{X}}_t \Psi_t(\mathbf{Y})} \end{aligned} \quad (51)$$

Manuscript received ; revised. *Corresponding author: Barbara Pascal.* B. Pascal is supported by ANR-23-CE48-0009 ‘‘OptiMoCSP’’.

where, by Definition [1](#), $\Psi_t(\mathbf{Y}) = \Psi_t(\alpha_1 k_1, \dots, \alpha_{t-1} k_{t-1})$ is not depending on k_t , and $G(\mathbf{Y}) = G(\alpha_1 k_1, \dots, \alpha_T k_T)$. Renaming the summation variable in the right-hand side of Equation [\(51\)](#) so that $k_t + 1 \rightarrow k_t$, it follows

$$\begin{aligned} & \sum_{k_t=0}^{\infty} G(\mathbf{Y}) \bar{\mathbf{X}}_t \Psi_t(\mathbf{Y}) \frac{(\bar{\mathbf{X}}_t \Psi_t(\mathbf{Y}))^{k_t}}{k_t!} e^{-\bar{\mathbf{X}}_t \Psi_t(\mathbf{Y})} \\ &= \sum_{k_t=0}^{\infty} G^{-t}(\mathbf{Y}) k_t \frac{(\bar{\mathbf{X}}_t \Psi_t(\mathbf{Y}))^{k_t}}{k_t!} e^{-\bar{\mathbf{X}}_t \Psi_t(\mathbf{Y})} \end{aligned} \quad (52)$$

where $\Psi_t(\mathbf{Y})$ is unchanged as it depends only on the *fixed* variables Y_1, \dots, Y_{t-1} , and, by definition, $G^{-t}(\mathbf{Y}) = G(Y_1, \dots, Y_t - \alpha_t, \dots, Y_t)$.

Remark. The above discrete counterpart of integration by part on variable k_t performed in Equations [\(51\)](#) and [\(52\)](#) can be seen as an application of the standard Poisson Stein’s lemma counterpart [\[1, 2\]](#) on the function G . Though, this reformulation is not enough to design Poisson Unbiased Risk Estimates for the driven autoregressive model [\(7\)](#) as, contrary to the case of standard Poisson Unbiased Risk Estimate, neither the function G nor its translate versions G^{-t} can be reformulated easily as expectations of tractable quantities.

Further computations are thus required. By Assumption [2](#)

$$\forall u \geq t + 1, \quad \Psi_s(\tilde{\mathbf{Y}}) \simeq \Psi_s(\mathbf{Y}) - \partial_{Y_t} \Psi_s(\mathbf{Y}) \alpha_t \simeq \Psi_s(\mathbf{Y}), \quad (53)$$

which, when injected into the expression of $G^{-t}(\mathbf{Y})$, obtained by replacing $\mathbf{Y} = (Y_1, \dots, Y_T)$ by $\tilde{\mathbf{Y}} = (Y_1, \dots, Y_t - \alpha_t, \dots, Y_T)$ in Equation [\(50\)](#), yields $G^{-t}(\mathbf{Y}) \simeq$

$$\sum_{k_{t+1}=0}^{\infty} \dots \sum_{k_T=0}^{\infty} \Theta^{-t}(\mathbf{Y}) \prod_{u=t+1}^T \frac{(\bar{\mathbf{X}}_u \Psi_u(\mathbf{Y}))^{k_u}}{k_u!} e^{-\bar{\mathbf{X}}_u \Psi_u(\mathbf{Y})} \quad (54)$$

where the change of variable $\mathbf{Y} \rightarrow \tilde{\mathbf{Y}}$ only affects the term in Θ but not the Poisson densities. Finally, injecting [\(54\)](#) into [\(48\)](#), yields $\mathbb{E}_{\mathbf{Y}} [\Theta(\mathbf{Y}) \bar{\mathbf{X}}_t \Psi_t(\mathbf{Y})] \simeq$

$$\sum_{k_1=0}^{\infty} \dots \sum_{k_T=0}^{\infty} \Theta^{-t}(\mathbf{Y}) k_t \prod_{s=1}^T \frac{(\bar{\mathbf{X}}_s \Psi_s(\mathbf{Y}))^{k_s}}{k_s!} e^{-\bar{\mathbf{X}}_s \Psi_s(\mathbf{Y})} \quad (55)$$

with $\forall t, Y_t = \alpha k_t$, in which one recognizes $\mathbb{E}_{\mathbf{Y}} [\Theta^{-t}(\mathbf{Y}) Y_t]$ by definition of the expectation on \mathbf{Y} . \square

¹By convention $\Psi_1 = Y_0$, where Y_0 is a deterministic initialization.

VIII. AUTOREGRESSIVE POISSON UNBIASED RISK ESTIMATE

Proof of Theorem 1 Let \mathcal{P} be the prediction risk defined in Equation (23). The derivation of an unbiased estimate of \mathcal{P} relies on the expansion of the prediction error:

$$\begin{aligned} & \left\| \widehat{\mathbf{X}}(\mathbf{Y}; \lambda) \odot \Psi(\mathbf{Y}) - \bar{\mathbf{X}} \odot \Psi(\mathbf{Y}) \right\|_2^2 = \left\| \widehat{\mathbf{X}}(\mathbf{Y}; \lambda) \odot \Psi(\mathbf{Y}) \right\|_2^2 \\ & - 2 \left\langle \widehat{\mathbf{X}}(\mathbf{Y}; \lambda) \odot \Psi(\mathbf{Y}), \bar{\mathbf{X}} \odot \Psi(\mathbf{Y}) \right\rangle + \left\| \bar{\mathbf{X}} \odot \Psi(\mathbf{Y}) \right\|_2^2. \end{aligned} \quad (56)$$

The first term in the right-hand side of (56) only depends on observations and hyperparameters, and hence appears as is in APURE^P in Equation (28). The second and third terms depend on the inaccessible ground truth $\bar{\mathbf{X}}$, and hence need to be reformulated, taking care of not introducing any bias. Considering the second term,

$$\begin{aligned} & \mathbb{E}_{\mathbf{Y}} \left[\left\langle \widehat{\mathbf{X}}(\mathbf{Y}; \lambda) \odot \Psi(\mathbf{Y}), \bar{\mathbf{X}} \odot \Psi(\mathbf{Y}) \right\rangle \right] \\ & = \sum_{t=1}^T \mathbb{E}_{\mathbf{Y}} \left[\widehat{X}_t(\mathbf{Y}; \lambda) \Psi_t(\mathbf{Y}) \bar{X}_t \Psi_t(\mathbf{Y}) \right]. \end{aligned}$$

By hypothesis, the memory functions satisfy Assumption 2. Further, for any $\lambda \in \Lambda$, and any $t \in \{1, \dots, T\}$, the function $\mathbf{Y} \mapsto \widehat{X}_t(\mathbf{Y}; \lambda) \Psi_t(\mathbf{Y})$ satisfies Assumption 1, hence the autoregressive Poisson lemma 1 applies and yields

$$\begin{aligned} & \mathbb{E}_{\mathbf{Y}} \left[\widehat{X}_t(\mathbf{Y}; \lambda) \Psi_t(\mathbf{Y}) \bar{X}_t \Psi_t(\mathbf{Y}) \right] \\ & \stackrel{\alpha \rightarrow 0}{=} \mathbb{E}_{\mathbf{Y}} \left[\left(\widehat{X}_t(\mathbf{Y}; \lambda) \Psi_t(\mathbf{Y}) \right)^{-t} Y_t \right] \end{aligned} \quad (57)$$

where $\left(\widehat{X}_t(\mathbf{Y}; \lambda) \Psi_t(\mathbf{Y}) \right)^{-t} = \widehat{X}_t^{-t}(\mathbf{Y}; \lambda) \Psi_t^{-t}(\mathbf{Y})$. Since Ψ_t does not depend on Y_t , $\Psi_t^{-t}(\mathbf{Y}) = \Psi_t(\mathbf{Y})$, leading to

$$\begin{aligned} & \mathbb{E}_{\mathbf{Y}} \left[\widehat{X}_t(\mathbf{Y}; \lambda) \Psi_t(\mathbf{Y}) \bar{X}_t \Psi_t(\mathbf{Y}) \right] \\ & \stackrel{\alpha \rightarrow 0}{=} \mathbb{E}_{\mathbf{Y}} \left[\widehat{X}_t^{-t}(\mathbf{Y}; \lambda) \Psi_t(\mathbf{Y}) Y_t \right] \end{aligned} \quad (58)$$

in which the reader recognizes the expression of the second term in the definition of APURE^P in Equation (28). As for the third term of Equation (56), it writes

$$\mathbb{E}_{\mathbf{Y}} \left[\left\| \bar{\mathbf{X}} \odot \Psi(\mathbf{Y}) \right\|_2^2 \right] = \sum_{t=1}^T \mathbb{E}_{\mathbf{Y}} \left[|\bar{X}_t \Psi_t(\mathbf{Y})|^2 \right]. \quad (59)$$

By hypothesis, the memory functions satisfy Assumption 2. Moreover for all $t \in \{1, \dots, T\}$, the function $\mathbf{Y} \mapsto \bar{X}_t \Psi_t(\mathbf{Y})$ satisfies Assumption 1, hence the autoregressive Poisson lemma 1 applies and yields

$$\mathbb{E}_{\mathbf{Y}} \left[(\bar{X}_t \Psi_t(\mathbf{Y}))^2 \right] \stackrel{\alpha \rightarrow 0}{=} \mathbb{E}_{\mathbf{Y}} \left[(\bar{X}_t \Psi_t(\mathbf{Y}))^{-t} Y_t \right] \quad (60)$$

where $(\bar{X}_t \Psi_t(\mathbf{Y}))^{-t} = \bar{X}_t \Psi_t^{-t}(\mathbf{Y})$. But, since Ψ_t does not depend on Y_t , one has $\bar{X}_t \Psi_t^{-t}(\mathbf{Y}) = \bar{X}_t \Psi_t(\mathbf{Y})$, and hence

$$\mathbb{E}_{\mathbf{Y}} \left[(\bar{X}_t \Psi_t(\mathbf{Y}))^2 \right] \stackrel{\alpha \rightarrow 0}{=} \mathbb{E}_{\mathbf{Y}} \left[\bar{X}_t \Psi_t(\mathbf{Y}) Y_t \right] = \mathbb{E}_{\mathbf{Y}} \left[Y_t \bar{X}_t \Psi_t(\mathbf{Y}) \right].$$

Then, remarking that the function $\mathbf{Y} \mapsto Y_t$ satisfies Assumption 1, the autoregressive Poisson Stein's lemma 1 applies once again, and it follows that

$$\mathbb{E}_{\mathbf{Y}} \left[(\bar{X}_t \Psi_t(\mathbf{Y}))^2 \right] \stackrel{\alpha \rightarrow 0}{=} \mathbb{E}_{\mathbf{Y}} \left[(Y_t - \alpha_t) Y_t \right]. \quad (61)$$

Equations (58) and (61), combined with the expansion provided in Equation (56) demonstrates the asymptotic unbiasedness of APURE^P stated in Equation (29).

Assuming that $\forall t \in \{1, \dots, T\}$, $\forall \mathbf{Y} \in \mathbb{R}^T$, $\Psi_t(\mathbf{Y}) \neq 0$, the estimation risk \mathcal{E} of Equation (22) is well-defined and the estimation error can be expanded as:

$$\left\| \widehat{\mathbf{X}}(\mathbf{Y}; \lambda) - \bar{\mathbf{X}} \right\|_2^2 = \left\| \widehat{\mathbf{X}}(\mathbf{Y}; \lambda) \right\|_2^2 - 2 \left\langle \widehat{\mathbf{X}}(\mathbf{Y}; \lambda), \bar{\mathbf{X}} \right\rangle + \left\| \bar{\mathbf{X}} \right\|_2^2. \quad (62)$$

The first term of the expansion, which is fully data-dependent, is kept as is in the definition of APURE^E in Equation (22). The second and third terms depend on the ground truth and have to be carefully reformulated, while avoiding to introduce any bias. The second term writes

$$\mathbb{E}_{\mathbf{Y}} \left[\left\langle \widehat{\mathbf{X}}(\mathbf{Y}; \lambda), \bar{\mathbf{X}} \right\rangle \right] = \sum_{t=1}^T \mathbb{E}_{\mathbf{Y}} \left[\widehat{X}_t(\mathbf{Y}; \lambda) \bar{X}_t \right].$$

By hypothesis, the memory functions satisfy Assumption 2. Further, for any $\lambda \in \Lambda$, and any $t \in \{1, \dots, T\}$, the function $\mathbf{Y} \mapsto \widehat{X}_t(\mathbf{Y}; \lambda) / \Psi_t(\mathbf{Y})$ is well-defined and satisfies Assumption 1, hence the autoregressive Poisson lemma 1 applies and yields

$$\begin{aligned} \mathbb{E}_{\mathbf{Y}} \left[\widehat{X}_t(\mathbf{Y}; \lambda) \bar{X}_t \right] & = \mathbb{E}_{\mathbf{Y}} \left[\frac{\widehat{X}_t(\mathbf{Y}; \lambda)}{\Psi_t(\mathbf{Y})} \bar{X}_t \Psi_t(\mathbf{Y}) \right] \\ & \stackrel{\alpha \rightarrow 0}{=} \mathbb{E}_{\mathbf{Y}} \left[\left(\frac{\widehat{X}_t(\mathbf{Y}; \lambda)}{\Psi_t(\mathbf{Y})} \right)^{-t} Y_t \right] \\ & \stackrel{\alpha \rightarrow 0}{=} \mathbb{E}_{\mathbf{Y}} \left[\frac{\widehat{X}_t^{-t}(\mathbf{Y}; \lambda) Y_t}{\Psi_t(\mathbf{Y})} \right] \end{aligned} \quad (63)$$

since $\Psi_t(\mathbf{Y})$ does not depend on Y_t . The third term of Equation (62) writes

$$\mathbb{E}_{\mathbf{Y}} \left[\left\| \bar{\mathbf{X}} \right\|_2^2 \right] = \sum_{t=1}^T \mathbb{E}_{\mathbf{Y}} \left[\bar{X}_t^2 \right] \quad (64)$$

$$= \sum_{t=1}^T \mathbb{E}_{\mathbf{Y}} \left[\frac{\bar{X}_t}{\Psi_t(\mathbf{Y})} \bar{X}_t \Psi_t(\mathbf{Y}) \right]. \quad (65)$$

By hypothesis, the memory functions satisfy Assumption 2. Moreover, for any $\lambda \in \Lambda$, and any $t \in \{1, \dots, T\}$, the function $\mathbf{Y} \mapsto \bar{X}_t / \Psi_t(\mathbf{Y})$ is well-defined and satisfies Assumption 1, hence the autoregressive Poisson lemma 1 applies, using that $\Psi_t^{-t}(\mathbf{Y}) = \Psi_t(\mathbf{Y})$, it leads to

$$\mathbb{E}_{\mathbf{Y}} \left[\frac{\bar{X}_t}{\Psi_t(\mathbf{Y})} \bar{X}_t \Psi_t(\mathbf{Y}) \right] \stackrel{\alpha \rightarrow 0}{=} \mathbb{E}_{\mathbf{Y}} \left[\frac{\bar{X}_t}{\Psi_t(\mathbf{Y})} Y_t \right] \quad (66)$$

$$\stackrel{\alpha \rightarrow 0}{=} \mathbb{E}_{\mathbf{Y}} \left[\frac{Y_t}{\Psi_t^2(\mathbf{Y})} \bar{X}_t \Psi_t(\mathbf{Y}) \right] \quad (67)$$

By hypothesis, the memory functions satisfy Assumption 2. Further remarking that the function $\mathbf{Y} \mapsto \Psi_t(\mathbf{Y})$ satisfies Assumption 1, the autoregressive Poisson lemma 1 applies once again, and using that $\Psi_t^{-t}(\mathbf{Y}) = \Psi_t(\mathbf{Y})$, one gets

$$\begin{aligned} \mathbb{E}_{\mathbf{Y}} \left[\frac{\bar{X}_t}{\Psi_t(\mathbf{Y})} \bar{X}_t \Psi_t(\mathbf{Y}) \right] &\stackrel{\alpha \rightarrow 0}{=} \mathbb{E}_{\mathbf{Y}} \left[\left(\frac{Y_t}{\Psi_t^2(\mathbf{Y})} \right)^{-t} Y_t \right] \\ &\stackrel{\alpha \rightarrow 0}{=} \mathbb{E}_{\mathbf{Y}} \left[\frac{(Y_t - \alpha_t) Y_t}{\Psi_t^2(\mathbf{Y})} \right]. \end{aligned} \quad (68)$$

Equations (63) and (68), combined with the expansion provided in Equation (62) demonstrates the asymptotic unbiasedness of APURE^P stated in Equation (31). \square

IX. FINITE DIFFERENCE MONTE CARLO ESTIMATORS

Proof of Theorem 2 The Finite Difference Monte Carlo strategy applied to the prediction risk estimate APURE^P of Equation (28) consists in rewriting the second term in the definition of APURE^P in Equation (28) in a tractable way. The proof thus focuses on this term and aims at demonstrating that

$$\begin{aligned} \mathbb{E}_{\mathbf{Y}} \left[\sum_{t=1}^T \hat{X}_t^{-t}(\mathbf{Y}; \lambda) \Psi_t(\mathbf{Y}) Y_t \right] &= \mathbb{E}_{\mathbf{Y}, \zeta} \left[\sum_{t=1}^T \hat{X}_t(\mathbf{Y}; \lambda) \Psi_t(\mathbf{Y}) Y_t \right. \\ &\quad \left. - \left\langle \text{diag}(\alpha \odot \Psi(\mathbf{Y})) \partial_{\mathbf{Y}} \hat{\mathbf{X}}[\zeta], \text{diag}(\mathbf{Y}) \zeta \right\rangle \right], \end{aligned} \quad (69)$$

which can be shown by alternatively demonstrating that for any $\lambda \in \Lambda$ and $t \in \{1, \dots, T\}$

$$\begin{aligned} \mathbb{E}_{\mathbf{Y}} \left[\hat{X}_t^{-t}(\mathbf{Y}; \lambda) \Psi_t(\mathbf{Y}) Y_t \right] &= \\ \mathbb{E}_{\mathbf{Y}, \zeta} \left[\hat{X}_t(\mathbf{Y}; \lambda) \Psi_t(\mathbf{Y}) Y_t - \alpha_t \Psi_t(\mathbf{Y}) \left(\partial_{\mathbf{Y}} \hat{\mathbf{X}}[\zeta] \right)_t Y_t \zeta_t \right] \end{aligned} \quad (70)$$

where $\partial_{\mathbf{Y}} \hat{\mathbf{X}}[\zeta] \in \mathbb{R}^T$ is the differential of $\hat{\mathbf{X}}(\mathbf{Y}; \lambda)$ with respect to the variable \mathbf{Y} at $(\mathbf{Y}; \lambda)$ applied to the T -dimensional Monte Carlo vector ζ where for the sake of conciseness the point $(\mathbf{Y}; \lambda)$ at which the differential is applied is omitted.²

To prove (70), first remark that, in the limit of small scale parameter $\alpha_t \rightarrow 0$, Assumption 3 implies that

$$\hat{X}_t^{-t}(\mathbf{Y}; \lambda) \stackrel{\alpha \rightarrow 0}{=} \hat{X}_t(\mathbf{Y}) - \alpha_t \frac{\partial \hat{X}_t}{\partial Y_t}. \quad (71)$$

The first term in the right-hand side of Equation (70) stems directly from the first term in Equation (71). Then, since ζ is a standard Gaussian vector, by definition $\mathbb{E}_{\zeta} [\zeta_s \zeta_t] = \delta_{s,t}$ where $\delta_{s,t}$ is the Kronecker delta.³

$$\frac{\partial \hat{X}_t}{\partial Y_t} = \sum_{s=1}^T \frac{\partial \hat{X}_t}{\partial Y_s} \mathbb{E}_{\zeta} [\zeta_s \zeta_t] = \mathbb{E}_{\zeta} \left[\sum_{s=1}^T \frac{\partial \hat{X}_t}{\partial Y_s} \zeta_s \zeta_t \right], \quad (72)$$

in which one recognizes the t th component of the differential of $\hat{\mathbf{X}}$ with respect to \mathbf{Y} applied to the vector ζ

$$\sum_{s=1}^T \frac{\partial \hat{X}_t}{\partial Y_s} \zeta_s = \left(\partial_{\mathbf{Y}} \hat{\mathbf{X}}[\zeta] \right)_t.$$

²Remind that the differential of a function $f : \mathbb{R}^T \rightarrow \mathbb{R}^T$ at a given point $\mathbf{Z} \in \mathbb{R}^T$ is a linear application $\partial_{\mathbf{Y}} f(\mathbf{Z})[\cdot] : \mathbb{R}^T \rightarrow \mathbb{R}^T$.

³By definition, $\forall s, t \in \mathbb{N}$, $\delta_{s,s} = 1$ and if $s \neq t$, $\delta_{s,t} = 0$.

Then, multiplying Equations (71) and (72) by $\Psi_t(\mathbf{Y}) Y_t$ and combining them, one gets

$$\begin{aligned} \mathbb{E}_{\zeta} \left[\hat{X}_t^{-t}(\mathbf{Y}; \lambda) \Psi_t(\mathbf{Y}) Y_t \right] \\ \stackrel{\alpha \rightarrow 0}{=} \mathbb{E}_{\zeta} \left[\hat{X}_t(\mathbf{Y}) \Psi_t(\mathbf{Y}) Y_t - \alpha_t \left(\partial_{\mathbf{Y}} \hat{\mathbf{X}}[\zeta] \right)_t \zeta_t \Psi_t(\mathbf{Y}) Y_t \right]. \end{aligned} \quad (73)$$

Taking the expectation with respect to $\mathbb{E}_{\mathbf{Y}}$ on both sides demonstrates (70), and finally shows the asymptotic unbiasedness of APURE^P stated in Equation (33).

Assuming that $\forall t \in \{1, \dots, T\}$, $\forall \mathbf{Y} \in \mathbb{R}^T$, $\Psi_t(\mathbf{Y}) \neq 0$, a similar proof holds for the estimation risk estimate. Applying the Finite Difference Monte Carlo strategy applied to the estimation risk estimate APURE^E of Equation (30) amounts to rewrite the second term in the definition of APURE^E in Equation (30) in a tractable way. The proof thus focuses on this term and aims at demonstrating that

$$\begin{aligned} \mathbb{E}_{\mathbf{Y}} \left[\sum_{t=1}^T \frac{\hat{X}_t^{-t}(\mathbf{Y}; \lambda)}{\Psi_t(\mathbf{Y})} Y_t \right] &= \mathbb{E}_{\mathbf{Y}, \zeta} \left[\sum_{t=1}^T \sum_{t=1}^T \frac{\hat{X}_t(\mathbf{Y}; \lambda)}{\Psi_t(\mathbf{Y})} Y_t \right. \\ &\quad \left. - \left\langle \text{diag}(\alpha \cdot / \Psi(\mathbf{Y})) \partial_{\mathbf{Y}} \hat{\mathbf{X}}[\zeta], \text{diag}(\mathbf{Y}) \zeta \right\rangle \right], \end{aligned} \quad (74)$$

which can be shown by alternatively demonstrating that for any $\lambda \in \Lambda$ and $t \in \{1, \dots, T\}$

$$\begin{aligned} \mathbb{E}_{\mathbf{Y}} \left[\frac{\hat{X}_t^{-t}(\mathbf{Y}; \lambda)}{\Psi_t(\mathbf{Y})} Y_t \right] &= \\ \mathbb{E}_{\mathbf{Y}, \zeta} \left[\frac{\hat{X}_t(\mathbf{Y}; \lambda)}{\Psi_t(\mathbf{Y})} Y_t - \frac{\alpha_t}{\Psi_t(\mathbf{Y})} \left(\partial_{\mathbf{Y}} \hat{\mathbf{X}}[\zeta] \right)_t Y_t \zeta_t \right]. \end{aligned} \quad (75)$$

The first term in the right-hand of Equation (75) stems directly from the first term in the Taylor expansion of Equation (71). The second term Equation (71) when injected in (75) and rewritten leveraging the Monte Carlo strategy of Equation (72) yields exactly the second term in Equation (75). Taking the expectation with respect to $\mathbb{E}_{\mathbf{Y}}$ on both sides demonstrates (75), and finally shows the asymptotic unbiasedness of APURE^E stated in Equation (35). \square

REFERENCES

- [1] F. Luisier, T. Blu, and M. Unser, "A new SURE approach to image denoising: Interscale orthonormal wavelet thresholding," *IEEE Trans. Image Process.*, vol. 16, no. 3, pp. 593–606, 2007.
- [2] Y. Le Montagner, E. D. Angelini, and J.-C. Olivo-Marin, "An unbiased risk estimator for image denoising in the presence of mixed Poisson–Gaussian noise," *IEEE Trans. Image Process.*, vol. 23, no. 3, pp. 1255–1268, 2014.