

Time-Resolved Laser Diagnostics in Hall Thrusters with Time-Delay Embedding Data Fusion

IEPC-2024-596

*Presented at the 38th International Electric Propulsion Conference, Toulouse, France
June 23-28, 2024*

Parker J. Roberts,*

University of Michigan, Ann Arbor, Michigan, 48109, U.S.A.

Vernon H. Chaplin,†

Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California, 91109, U.S.A.

and

Benjamin A. Jorns‡

University of Michigan, Ann Arbor, Michigan, 48109, U.S.A.

Numerical data fusion algorithms are applied to the reconstruction and denoising of laser plasma diagnostic data in a Hall thruster plume. The performance of two algorithms—average linear transfer function estimation, and nonlinear shadow manifold interpolation based on Takens’s theorem—are compared based on their ability to synchronize and denoise time-resolved laser-induced fluorescence measurements of ion velocity fluctuations. Hyperparameters are optimized for both models in order to minimize error based on testing on synthetic, ground truth datasets. Convergence studies are then performed to assess the optimal performance of both algorithms with respect to dataset sample depth and periodicity. These techniques are then applied to real laser-induced fluorescence data taken in a magnetically shielded Hall thruster for two operating voltages, corresponding to low-frequency discharge oscillations in both a highly oscillatory, near-periodic regime as well as an aperiodic regime. It is found that the nonlinear time-delay embedding model is able to converge to a high-fidelity reconstruction with roughly 10 times fewer samples than required for the Fourier-domain method, and may be less prone to artifacts such as overshoots. However, the transfer function method is generally able to achieve a lower total error when large training datasets are used. The nonlinear model is also more computationally expensive than the linear algorithm. Following these comparisons, four advanced applications of these algorithms for laser plasma diagnostics are demonstrated: inference of nonlinear dynamics, reference data fusion, high-frequency time-resolution via background subtraction, and application to pulsed diagnostics. The feasibility and utility of both algorithms for these applications are discussed in the context of sampling requirements, experiment time, and model assumptions.

*Doctoral Candidate, Department of Aerospace Engineering, pjob@umich.edu

†Technologist, Electric Propulsion Group, vernon.h.chaplin@jpl.nasa.gov

‡Associate Professor, Department of Aerospace Engineering, bjorns@umich.edu

Nomenclature

AEPS	= Advanced Electric Propulsion System
AOM	= Acousto-optic modulator
EVDF	= Electron velocity distribution function
FFT	= Fast Fourier Transform
ITS	= Incoherent Thomson scattering
IVDF	= Ion velocity distribution function
LIF	= Laser-induced fluorescence
PSD	= Phase-sensitive detection
SMI	= Shadow manifold interpolation reconstruction algorithm
SNR	= Signal-to-noise ratio
TALIF	= Two-photon absorption laser-induced fluorescence
TF	= Transfer function estimation reconstruction algorithm
$A(t)$	= Training reference time series
$B(t)$	= Reconstruction reference time series
β	= Normalization constant for SMI weighted sum
$C(t)$	= Training LIF time series for a particular wavelength
$D(t)$	= Reconstruction LIF time series
$d(t_1, t_2)$	= Euclidean distance between time delay vectors at times t_1 and t_2
E	= time-delay manifold dimension
$F(t, \lambda)$	= Laser-induced fluorescence signal
$\tilde{F}(\omega, \lambda)$	= Fourier transform of laser-induced fluorescence signal
$\tilde{H}(\omega)$	= Transfer function
k	= Number of nearest neighbors
m_i	= sample index for the i th nearest neighbor
$N(t)$	= Noise added to signal
N_{chunks}	= Number of chunks to subdivide transfer function time series
ω	= Angular frequency
$R(t)$	= Reference signal
$\tilde{R}(\omega)$	= Fourier transform of reference signal
t	= time
$t_{m_i}^A$	= Training time corresponding to the m_i th sample
τ	= Time-delay parameter
\vec{M}_X	= Shadow manifold/time-lag vector for signal $X(t)$
w_j	= weights for weighted average
V_D	= Discharge Voltage
λ	= laser wavelength

I. Introduction

THE implementation of electric propulsion (EP) technologies on space missions relies on accurate prediction of on-orbit performance. For EP missions, which typically trade wet mass for extended burn times due to higher achievable specific impulse, device lifetime and failure prediction are of particular concern.^{1,2} Lifetime qualification efforts are historically based on ground testing in vacuum chambers, but the improvement of physics-based models for these devices enables an increased reliance on simulations.³ These simulations can significantly decrease cost and qualification time if sufficiently reliable. However, achieving fully predictive plasma simulations is complicated by the strong impacts on performance and lifetime from poorly understood, time-dependent plasma behavior across a range of spatiotemporal scales.

For example, Hall thrusters, the most commonly flown type of in-space electric propulsion, exhibit time-dynamic behavior at frequencies spanning orders of magnitude. These phenomena include the relatively low-frequency breathing mode (< 20 kHz),^{4,5} intermediate coherent cathode oscillations ($\sim 50 - 100$ kHz),^{6,7} and high-frequency, microscopic plasma instabilities (> 1 MHz).^{8,9} These spatiotemporal fluctuations can strongly influence the local discharge physics, and in turn impact predictions of thrust, efficiency, and erosion rates.¹⁰⁻¹³ Therefore, models used for design and qualification must rely on experimental characterization of these time-dependent effects.

Non-invasive laser scattering diagnostics are a key experimental tool for characterizing the local plasma physics in Hall thrusters. Laser-induced fluorescence (LIF) velocimetry, a method for directly measuring the ion velocity distribution, has been used extensively to investigate the discharge physics in Hall thrusters, and to calibrate unknown parameters in models.^{14,15} More recently, incoherent Thomson scattering (ITS) has also provided key insights into Hall thruster physics by resolving the electron velocity distribution,^{16,17} and two-photon absorption laser-induced fluorescence (TALIF) has likewise supplied measurements of the ground-state neutral gas density.¹⁸ Despite the utility of these diagnostics, measuring time-resolved behavior is difficult due to their low signal-to-noise ratio (SNR). Because background light and random noise typically dwarf the signal of interest, ensemble averaging techniques over long periods must be used to discern the signal—however, this process generally destroys high-frequency information.¹⁹

Despite these difficulties, several methods exist for achieving time resolution for these low-signal diagnostics. These techniques generally work by mapping the signal of interest to some time-resolved global reference measurement, such as the discharge current. Several methods exist in the literature for generating such a mapping when thruster oscillations are either naturally periodic or forced to be thus.²⁰⁻²² However, passive Hall thruster oscillations are not generally periodic,¹⁹ and there is thus a need for a more general method to generate mappings between a reference signal and optical diagnostic data.

Two methods have emerged in the literature to perform this type of denoising operation on low-signal, time-resolved diagnostic data based on learning a mapping from an aperiodic reference measurement. First, the technique of empirical transfer function (TF) estimation accomplishes this by mapping between frequency-domain representations of the signals. This technique has been successfully applied to time-resolved probe measurements⁴ as well as LIF velocimetry.^{19,23} However, transfer function estimation requires significant experimental time to build large training datasets, and relies on the potentially questionable assumption that the Hall thruster dynamics are linear and time-invariant.

The second method is known as shadow manifold interpolation (SMI) or time-delay embedding reconstruction, which has more recently been demonstrated on Hall thruster measurements from a variety of electrical probes.²⁴⁻²⁶ This technique reconstructs the diagnostic signal from a more general, fully nonlinear mapping, based on an approximation of the higher-dimensional state space of the thruster, and may be more robust to noise than the TF method.^{25,27} However, to the authors' knowledge, this SMI analysis technique has not yet been applied and validated for the types of optical diagnostic data from LIF measurements which are the state of the art for calibration of Hall thruster simulation codes. There is then a need for the rigorous comparison of performance between these two aperiodic signal reconstruction techniques for LIF data, as well as an exploration of the utility of SMI for achieving time resolution for diagnostics to which transfer function method is more difficult to apply, such as the pulsed diagnostics of ITS and TALIF.

In this work, we implement and demonstrate both the TF and SMI algorithms on time-resolved LIF datasets from the Advanced Electric Propulsion System (AEPS) thruster.²⁸ In Section II, we describe the details of time-resolved LIF acquisition, followed by the working principles of both analysis techniques. Following this, Section III presents the results of performance comparison on both synthetic datasets in order to determine optimal hyperparameters, as well as with real LIF data. Section IV then explores four advanced applications of time-resolved reconstruction with both synthetic and real datasets: inferring nonlin-

ear dynamics, performing data fusion of multiple reference time series, allowing for background subtraction to improve the time-resolution, and reconstructing sparse samples from pulsed laser diagnostics. Finally, Section V summarizes our key conclusions from this body of work.

II. Methodology

In this section, we introduce the laser-induced fluorescence (LIF) diagnostic which we use as the testbed for exploring the performance of algorithms for reconstruction and denoising of LIF data based on a reference signal measurement. We then summarize the underlying principles and implementation of two reconstruction methods: a linear frequency domain technique, and a nonlinear time-domain based technique.

A. Overview of Laser-Induced Fluorescence and Phase-Sensitive Detection

Laser-induced fluorescence (LIF) velocimetry is a key diagnostic widely used to measure the ion velocity distribution in electric propulsion plasmas. To accomplish this, a tunable laser is used to excite an electronic transition from a metastable state of the target ion or atom. This excited electron subsequently decays, emitting fluorescence which can be detected by a photomultiplier tube (PMT). As the laser wavelength is varied, the beam selectively excites ions moving at particular velocities along the laser axis, which correspond to the Doppler-shifted transition wavelength for those velocities.²⁹ As a result, the fluorescence intensity as a function of laser wavelength provides a proxy for the ion velocity distribution function (IVDF) projected along the laser injection axis. Non-Doppler effects which broaden the lineshape include Zeeman or Stark splitting and hyperfine splitting, but these effects are typically so small as to have a negligible impact on the mean velocity of the distribution, such that they are typically ignored for the singly-charged xenon LIF scheme used in the present work.³⁰

Due to the intense background light from the plasma, the signal-to-noise ratio (SNR) of this PMT measurement is too low to resolve a meaningful signal in most cases.¹⁹ To mitigate this, a technique known as phase-sensitive detection (PSD) is typically applied. PSD relies on modulating the laser at a known frequency, for example with a mechanical chopper or acousto-optic modulator (AOM). A lock-in amplifier then uses frequency mixing and low-pass filtering to extract the portion of the PMT signal which varies at the modulation frequency. Since the majority of the noise due to background light should occur at frequency components other than this modulation frequency, this procedure significantly improves the SNR. However, the time-constant τ_{PSD} of the low-pass filter used for PSD is essentially an averaging window, and time-resolved information which varies faster than this timescale is destroyed.

To attain time-resolved LIF signal and infer the fluctuations in the IVDF, it is thus necessary to reduce τ_{PSD} to be shorter than the phenomenon of interest. For example, to resolve the 10-100 kHz oscillations which dominate global Hall thruster current, Durot et al. sample the LIF signal at 20-30 MHz, modulate the laser at 1 MHz, and perform PSD digitally in post-processing with $1/\tau_{PSD} = 500$ kHz.¹⁹ However, the adverse effect of shortening the time constant is that the SNR worsens. We must therefore perform further ensemble averaging in a way that preserves time-resolved information.

One way to accomplish this is by recording a “reference” signal, which encodes the time-resolved global state of the Hall thruster, and generating a mapping between the reference and LIF signals in order to further reduce the noise.⁴ The discharge current to the anode, as measured by inductive probes on the transmission lines, is a natural choice for this purpose. Figure 1 shows a diagram of the Hall thruster operating in a ground test with these data streams, demonstrating the noisy LIF signal as a function of time for a particular wavelength (velocity). If the discharge current oscillates in a periodic manner, the reference and noisy LIF signals can straightforwardly be broken up into segments corresponding to several phase bins of the oscillation - averaging these segments can result in “phase resolution” of the LIF signal. This idea is behind several periodic reconstruction techniques, known for example as sample-and-hold averaging, boxcar averaging, and photon counting, which are detailed in Ref. 31. However, for the quasiperiodic and aperiodic oscillations which are more commonly present in Hall thrusters, it is necessary to learn a more sophisticated mapping between the reference and LIF signal. This information can provide a metric for identifying times with “similar behavior” in a way that preserves time-resolved information when ensemble averaging is applied.

For this work, we seek to explore the performance of these denoising algorithms by reanalyzing the datasets presented in Refs. 23 and 12. This LIF data consists of time-resolved IVDFs acquired in the plume

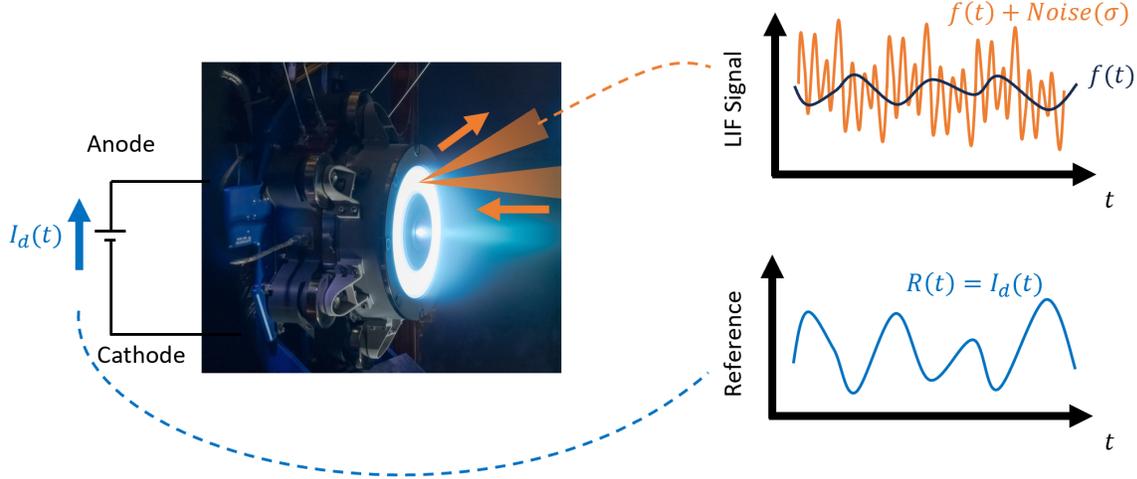


Figure 1. Schematic of the fundamental data streams for time-resolved laser induced fluorescence and similar optical diagnostics for a Hall thruster, consisting of the discharge current as an example global reference, and the noisy time-resolved laser diagnostic output. Our goal is to train a model to predict/denoise the LIF signal, $f(t)$, based on the information about the thruster encoded within the reference signal, $R(t)$.

of the Advanced Electric Propulsion System (AEPS) Hall thruster at the Jet Propulsion Laboratory (JPL), California Institute of Technology. These measurements were performed with the thruster operating on xenon propellant, with a discharge current of 20.83 A and discharge voltages of 300 V and 600 V. A transition in singly-charged xenon ions from the metastable $5p^4 ({}^3P_2) 5d^2 [4]_{7/2}$ state to the $5p^4 ({}^3P_2) 6p^2 [3]_{5/2}^o$ state at 834.953 nm (vacuum) was targeted by a tunable diode laser.³² The corresponding fluorescence wavelength was 542.06 nm, and an optical bandpass filter reduced background light from other wavelengths before the scattered signal reached the PMT. For these datasets, the light was injected axially to measure the axial ion velocity distribution within the thruster acceleration region.

At the 300-V operating condition, this thruster exhibits weaker, aperiodic fluctuations, compared to the 600-V case’s strong, nearly periodic 50-60 kHz breathing mode oscillations.^{23,33} These two datasets therefore provide interesting test cases for validating these algorithms in qualitatively different dynamical situations. For these measurements, the sample rate of the digitized PMT output was 25 MS/s, with each laser wavelength worth of training data extending 30 s in length. This corresponds to the significant storage requirement of at least 3 GB per wavelength, per position. In the remaining subsections, we describe in detail the transfer function analysis algorithm which was previously used to analyze this dataset and map the breathing mode oscillation to thruster performance in Refs. 23 and 12, in addition to the newer shadow manifold technique based on Takens’s embedding theorem.

B. Fourier-domain method (Average Transfer Function Estimation)

The first denoising algorithm we focus on in this work is known as average transfer function (TF) estimation. This technique makes the assumption that the thruster is a linear-time-invariant (LTI) system. The consequence of this assumption is that the Fourier transform of the output measurement time series, in this case the LIF signal at a particular laser wavelength $F(t, \lambda)$, is related to the reference signal $\tilde{R}(\omega)$ by a linear transfer function, $\tilde{H}(\omega, \lambda)$, assumed to be constant in time:

$$\tilde{F}(\omega, \lambda) = \tilde{H}(\omega, \lambda) \tilde{R}(\omega). \quad (1)$$

Note that each quantity in this expression is complex, so that the transfer function relates both amplitude and phase information between the reference and LIF signal at each frequency. In our context, this reference signal is typically the thruster discharge current, measured with a Pearson coil. If the assumption in Eq. 1 holds, it is possible to estimate the transfer function $H(\omega)$ from a sufficiently long time series for F and R .

Detailed explanations of the method for inferring the transfer function and its performance against other methods for denoising of periodic signals are found in the work of Lobbia⁴ and Durot¹⁹, who leveraged TF estimation to synchronize spatiotemporal Langmuir probe maps and achieve time-resolved laser-induced fluorescence during aperiodic oscillations, respectively. However, we provide a brief overview of this technique below.

We assume that the reference signal (discharge current) is measured without noise; however, the LIF signal has a noise component consisting of background light and electrical fluctuations in the detector. This leads to the measured time series $F_{MEAS}(t, \lambda) = F_{LIF}(t, \lambda) + N(t)$. The noise time series $N(t)$ can be approximated as the realization of a random variable, but we note that this may not be entirely accurate because the variation of the background light is also a result of the dynamical system governing the plasma state. We discuss this point further in Sec. IIIC. Regardless, after performing phase-sensitive detection, only the frequency component of the deterministic part of the background laser light which is near the modulation frequency should survive. We then divide the time series, which may be hundreds of millions of values long in practice, into several thousand shorter "chunks", in order to facilitate averaging. For each chunk (indexed here by j), the discrete Fourier transforms of both the reference signal and LIF signal are computed, yielding (with noise) $F_{MEAS,j}(\omega) = F_{LIF,j}(\omega) + N_j(\omega)$, and $R_j(\omega)$. An empirical transfer function can then be computed by the simple division $H_j(\omega) = F_{MEAS,j}(\omega)/R_j(\omega)$. The average transfer function is then computed by simply averaging the $H_j(\omega)$ curves for each chunk. In practice, this is achieved with a weighting scheme w_j , i.e. $H_{est}(\omega) = \sum_{j=1}^{N_{chunks}} w_j H_j(\omega)$, where $\sum w_j = 1$. Provided that $N(\omega)$ has zero mean, this formulation converges to the true transfer function as the number of chunks becomes large, for constant sample rate and chunk length. For this work, we follow Durot¹⁹ in using the Welch weighting scheme,

$$w_j(\omega) = \frac{|R_j(\omega)|^2}{\sum_j |R_j(\omega)|^2}, \quad (2)$$

which enforces finite variance for the transfer function estimator by weighting the sum by frequency components which are present with large amplitudes in the data.³⁴ Our numerical estimate for H_{est} can now be substituted into Eq. 1 for a given reference signal, leading to a denoised reconstruction of the $F_{LIF}(\omega)$; finally, simply taking the inverse discrete Fourier transform results in the denoised time series. This reconstruction algorithm is illustrated in flowchart form in Fig. 2.

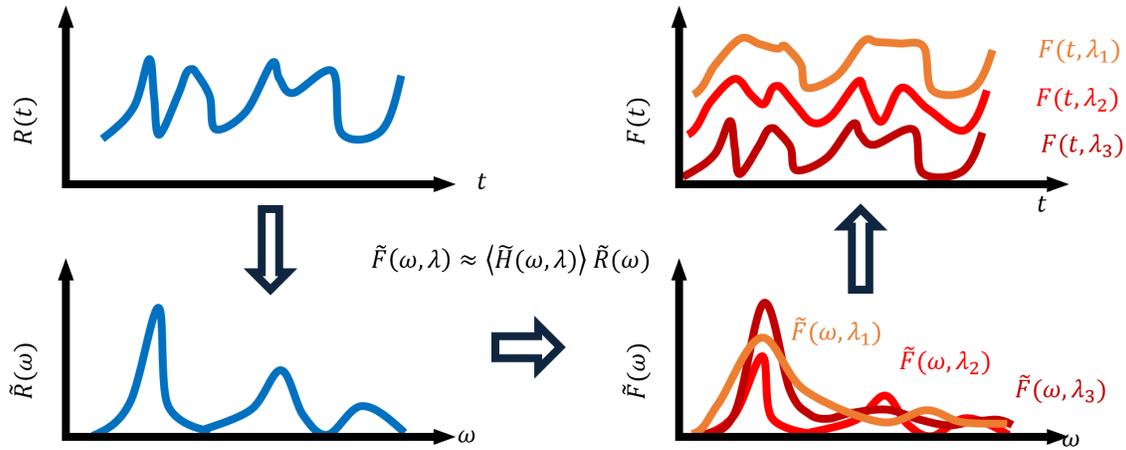


Figure 2. Illustration of the transfer function estimation (TF) algorithm. The reference signal $R(t)$ is transformed to the frequency domain, from which an average linear mapping is learned to the frequency-domain representation of LIF signal at each wavelength. The inverse transforms of these reconstructions to the time domain make up the time-resolved IVDF.

This technique can be applied to measure the transfer functions for several distinct output signals from the same input at different times, for example for different LIF wavelengths spanning the Doppler-broadened

lineshape of the targeted transition. Provided that the same reference signal trace is used for the reconstruction and the system is truly linear and time-invariant, this enables the reconstruction of synchronized instances of all output variables, in turn enabling simultaneous reconstruction of the entire IVDF. This transfer function technique is computationally powerful as it relies on the many decades of development of fast Fourier transform computations. However, the assumption that the system is linear may fail, especially given that Hall thrusters have been known to show nonlinear and/or chaotic dynamics. It is therefore possible that this reconstruction method could miss relevant features of the true LIF signal when these assumptions are violated. Also, the number of chunks to average over is an unspecified hyperparameter for which there is no obvious choice.

C. Shadow Manifold Interpolation

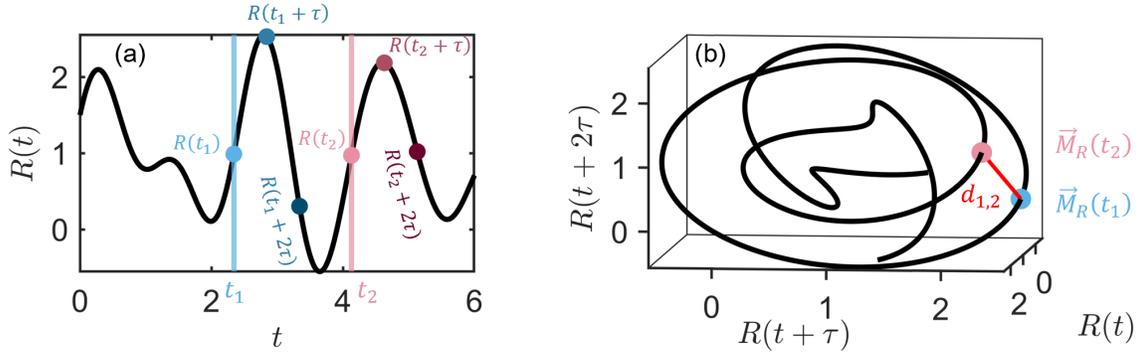


Figure 3. Example of 3D shadow manifold construction for a 1D signal. (a) 1D example signal demonstrating the values of a time lag vector at two times, t_1 (shades of blue) and t_2 (shades of red). (b) Three-dimensional shadow manifold representation of the signal in (a). The Euclidean distance d (red) between the points on the manifold at times t_1 and t_2 is a measure of the similarity of the system’s state at these two times.

The second technique we explore in this work is known as time-delay embedding reconstruction, or alternatively as shadow manifold interpolation (SMI). Unlike the transfer function estimation procedure, this technique is a method for finding a fully nonlinear mapping between the reference and output signals. Rather than assuming a linear system, SMI instead relies on the assumption that the measurements are governed by an unknown, potentially high-dimensional dynamical system. However, the cost of relaxing the linearity assumption is significant computational expense and complexity, especially in regards to the large space of hyperparameters which must be tuned to achieve a proper reconstruction.

The SMI reconstruction algorithm can be understood intuitively by recalling our goal of performing ensemble averaging over subsets of times which preserve time-resolved information, rather than averaging over the entire signal. To achieve this, any reconstruction method must contain some definition of “similarity” between the system behavior at different times. For example, the TF method described in the preceding *implicitly* identifies similar states of the system as those in which the phases of the reference signal’s constituent frequency components lead to repeatable patterns of interference. The SMI method instead provides an explicit topological measure for this notion of similarity by looking to Takens’s embedding theorem, which describes the relationships between so-called “shadow manifolds” or time-lagged representations of a measurement.

To illustrate this idea by example, we show an artificial, aperiodic time-series reference signal $R(t)$ in Fig. 3(a). Suppose we would like to find times at which the dynamical system is in a similar state as at time t_1 . Our first step might be to look for points at which the signal has the same value, i.e. to identify times t for which $R(t) \approx R(t_1)$. While such points exist within the time series shown, for several of them the signal obviously has a very different slope and concavity – more generally, the system is not in a similar location in phase space in this case. We can improve this by constructing a two-dimensional “time-delay vector”, which not only considers the current state of the system, $R(t_1)$, but also includes the future state after some time lag τ , i.e. $R(t_1 + \tau)$. More generally, we can construct an E -dimensional time-delay vector,

$\vec{M}_R(t)$, by considering the value of the time series at the next E time-lags of duration τ :

$$\vec{M}_R(t) = [R(t), R(t + \tau), \dots, R(t + (E - 1)\tau)]. \quad (3)$$

For example, in Fig. 3(a), the components of a time-delay vector with dimension $E = 3$ at the time t_1 would be the y-value of the three blue points, while at time t_2 the three red points would make up the time-delay vector components, for a particular choice of τ . The so-called “shadow manifold” is then the path in E -dimensional space which $\vec{M}_R(t)$ traces out as the time t evolves. We show in Fig. 3(b) the three-dimensional shadow manifold corresponding to the time series in Fig. 3(a), with the blue and red points highlighting the value of the time-lag vector at times t_1 and t_2 , respectively.

The utility of constructing such a shadow manifold is due to Takens’s embedding theorem, which is described in full rigor in Refs. 35 and 27. In lay terms, this theorem states that if at two times, t_1 and t_2 , the corresponding points on the shadow manifold are “nearby” each other, according to the Euclidean distance metric,

$$d(t_1, t_2)^2 = \left(\vec{M}_R(t_1) - \vec{M}_R(t_2) \right)^2, \quad (4)$$

then the system’s location in its higher-dimensional phase-space is also nearby at the two corresponding times. This is generally true if a sufficiently large manifold dimension, E , is chosen.²⁷ In other words, the distance metric d in the shadow manifold space provides the measure of “similarity” which we are looking for. For example, in Fig. 3(a), the qualitative similarity between the signal behavior after times t_1 and t_2 (a sharp increase followed by a decrease) is reflected by the closeness of these two points in the shadow manifold space in Fig. 3(b). We can thus denoise the spectrum while preserving time resolution by taking an ensemble average over points in the time series which are nearby each other according to this metric. We note that many texts adopt a convention in terms of past time lags, $(t - m\tau)$, etc. instead of the future times, $(t + m\tau)$ which we use here—however, the theorem is symmetric with respect to time, and we do not lose generality with this choice.

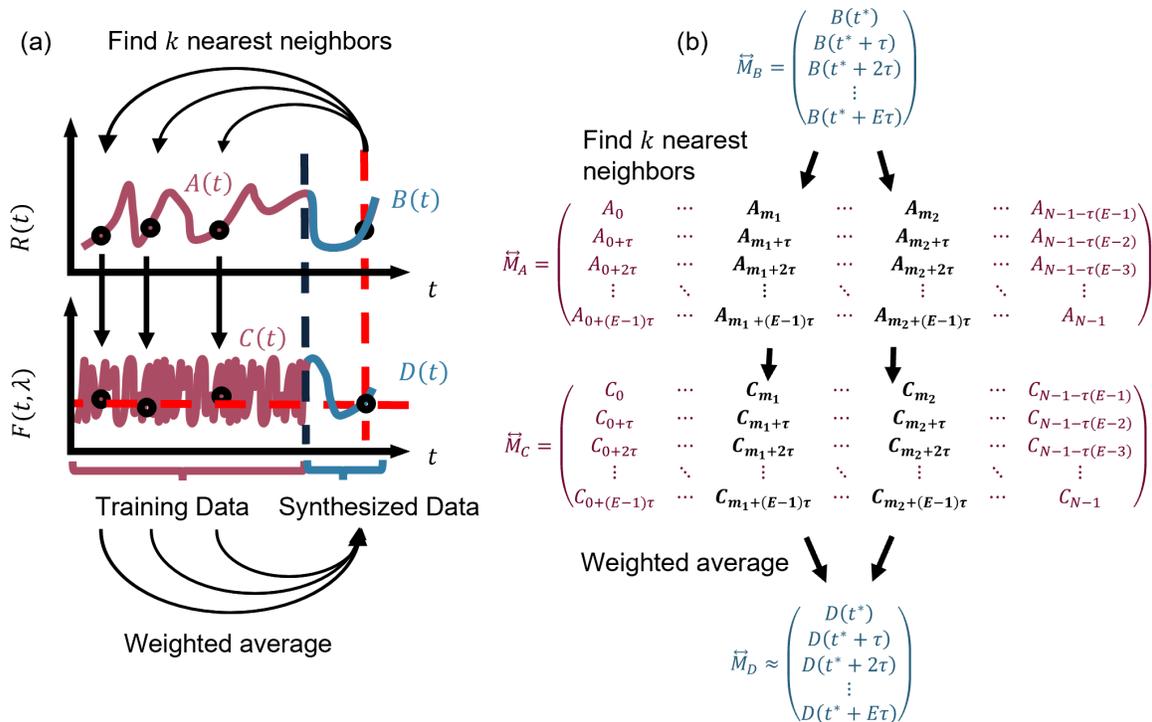


Figure 4. (a) Illustration of the SMI reconstruction algorithm, with the reconstruction time base chosen to fall at the end of the training time series. (b) Example of the discrete representation of the shadow manifold reconstruction of the signal at time t^* .

In practice, we must leverage these properties of shadow manifolds to denoise an LIF signal $F(\lambda, t)$ at several wavelengths, λ , from the reference signal $R(t)$. Just as for the TF analysis, we accomplish this by acquiring simultaneous LIF and reference data over a long “training” time series for each wavelength. We also choose one shorter portion of the reference signal over which to reconstruct the denoised signal at all wavelengths, as shown in Fig. 4(a). For each LIF wavelength, the time lag vectors are evaluated for the reference training data, $\vec{M}_A(t)$, and for the noisy LIF training data, $\vec{M}_C(t, \lambda)$, at each time t_m^A in the training time series, where m is an index ranging from 0 to $N - 1$, the total sample depth. We also compute the time-lag vectors for the reconstruction reference data, $\vec{M}_B(t)$, for each time t^B in the reconstruction time base.

For each time t_j^B in the reconstruction time series, we use Eq. 4 to find $d(t_j^B, t_m^A)$, the distance between the reconstruction reference and the training reference shadow manifolds, for all training times t_m^A . We then select the set of training time indices $m_i = \{m_0, m_1, \dots, m_{k-1}\}$ corresponding to the k nearest neighbors (lowest values of d). Figure 4(a) and (b) show a graphical depiction of this process and the structure of the resulting matrices, respectively. The shadow manifold $\vec{M}_D(t_j^B)$ can now be reconstructed by performing a weighted ensemble average of the LIF shadow manifold values at these k times. We follow Eckhardt et al. in performing this average with weights specified by a decaying exponential kernel,

$$\vec{M}_D(t_j^B) \approx \sum_{i=1}^k w_i \vec{M}_C(t_{m_i}^A), \quad (5)$$

where $w_i = \beta \exp(-d(t_j^B, t_{m_i}^A))$. The normalization constant β is chosen such that $\sum w_j = 1$.^{24–27} Equation 5 defines the reconstructed time-lag vector as the simple average of the time-lag vectors at the k times corresponding to the most similar behavior found within the training dataset, with the exponential weighting guaranteeing that the times with less similar behavior have a smaller effect on the sum. When this averaging is performed for all reconstruction times t^B , each of the E rows of the reconstructed shadow manifold matrix represents a distinct, time-lagged copy of the LIF time series^a – one column of this matrix is shown in the final step of Fig. 4(b). We thus can obtain further noise reduction by averaging each of these copies, resulting in a single denoised reconstruction of the LIF signal $F(t, \lambda)$. This is accomplished by “unwrapping” the reconstructed shadow manifold matrix by shifting each row to undo the time lags. For example, in the final step of Fig. 4(b), unwrapping the neighboring columns would result in E realizations of the value $D(t^*)$. Then, we compute the mean value of $D(t^*)$ as our final reconstruction at that time.

This procedure is in general less restrictive than the transfer function method discussed in the preceding, as there is no assumption of linearity required — only that the measurements arise from a deterministic dynamical system. However, due to the large amount of training data to search through, the k -nearest neighbor search step can be computationally expensive. Also, this function requires three hyperparameters: the time delay τ , the shadow manifold dimension E , and the number of nearest neighbors to use, k . While some methods exist for guessing at the optimal hyperparameters, the best values to use can vary between datasets and have a significant impact on the reconstruction quality.

III. Performance Comparison

In this section, we describe our framework for directly comparing the SMI and TF analysis algorithm performance and the results of these comparisons. We first address the need to choose hyperparameters for each analysis method and our strategy for determining optimal values. We then describe our method for generating artificial datasets to obtain direct performance comparison between the algorithms based on a least-squares error metric. Finally, we demonstrate the performance of the algorithms on real LIF data in the AEPS thruster at two operating conditions with different dynamics, followed by exploring two unique capabilities of the SMI algorithm: reconstruction of nonlinear mappings and fusing multiple reference data streams.

^aTechnically, there is a buffer consisting of the first $\tau \cdot (E - 1)$ points at the beginning and end of this time series matrix which have less than E copies due to boundary effects. In this work, we redefine our reconstruction time axis to remove these points. Thus, the length of the reconstruction is slightly shorter than the original time series.

A. Parameter Optimization for both techniques

In order to implement the analysis algorithms described in the previous section, several hyperparameters related to each algorithm and the data size must be chosen. For both methods, the dataset must be pre-conditioned based on the available computational time for analysis - this includes making a choice of the down-sampling ratio n_{ds} , which represents the factor by which we down-sample the data before processing, as well as the sample depth, N_{samp} . In this work, we generally hold the downsample ratio fixed at $n_{ds} = 50$ to facilitate faster computations, leading to an effective sample rate of 500 kHz. This is still well above the bandwidth limitation of ~ 100 kHz from the PSD time constant of 700 ns which we used. Past studies which analyzed the dataset we focus on in this work used the full, 30-second record at each LIF wavelength ($N_{samp} = 7.5 \cdot 10^8$), however in the following we demonstrate convergence of the analysis for shorter record times. After selecting these hyperparameters related to dataset length and sample rate, the parameters must be specified for each analysis method: For the transfer function analysis, the primary parameter is the chunk length by which to subdivide the dataset for training and averaging, while for shadow manifold reconstruction, the three parameters τ , E , and k (c.f. Section II.C) must be specified.

Previous studies have compared the performance of the SMI and TF techniques on other types of data, such as from electrical plasma probes, with a relatively high signal to noise ratio. This literature has largely chosen hyperparameters such as the time delay, τ , and manifold dimension, E , based on widely accepted existing theoretical metrics for dynamical systems embedding representations.^{36,37} In this work, however, we work with optical diagnostic data which has a far lower signal to noise ratio (SNR), ranging from roughly 0.01 to 0.1. As a result, we found that these best-practice metrics were not effective in all cases for producing optimal reconstructions, especially for aperiodic, noisy data. In this section, we therefore begin by empirically finding optimal parameters for both the TF and SMI methods based on an error study for synthetic, “ground-truth” datasets. In the following subsection, we discuss our method for generating these simulated datasets in detail, followed by the resulting optimizations for both the SMI and TF method hyperparameters.

B. Error study using artificial ground truth data

To assess the performance and convergence of the two algorithms quantitatively, we required a ground truth dataset similar in character to the LIF data we sought to analyze. To accomplish this, we began with the real discharge current traces which were recorded as reference signals for the LIF analysis. We used these references to first generate artificial linear transfer functions to create a known ground truth, then added random, Gaussian noise. The artificial transfer function that we specified consists of the sum of two large Gaussian functions in the frequency domain, with peaks near 5 kHz and 50 kHz, which is approximately representative of the thruster data:

$$\tilde{H}(2\pi\omega) = [\exp(-(\omega/(2\pi) - 5000)^2/10000^2) + 5 \exp(-(\omega/(2\pi) - 50000)^2/10000^2)] \cdot \exp^{i\phi}, \quad (6)$$

with the phase being chosen as $\phi = -2\pi^2\omega/10^6$. This function was reflected over negative frequencies to produce a real-valued output. This form was chosen to recreate several features of the real data^{12,23} which must be captured by the analysis algorithms: a nonzero DC offset ($\tilde{H}(0) \neq 0$), and wide-band frequency content near 0-10 kHz and 50-60 kHz.

Noise was then added to the resulting LIF waveform with a pseudorandom number generator to match a specified SNR, defined as the ratio of the RMS amplitude of the signal relative to that of the noise. For this study, we generated artificial data in this manner with SNR values of both 0.1 and 0.01, whereas we determined the SNR of the real LIF data to be approximately 0.07 based on previous LIF analysis¹². We then simulated the modulation process by mixing the laser signal with the AOM waveform, which was also recorded in the original dataset. We repeated this process for the full-length datasets at both the operating condition with aperiodic fluctuations (300 V) and the strong oscillations (600 V). For the following analysis, we use the root-mean-square error (RMSE) or L2 norm relative to the known ground truth dataset as our primary loss function. That is, we performed parameter optimization by finding the set of hyperparameters which minimized this function.

We assessed the performance of the transfer function and shadow manifold reconstruction algorithms by applying both to the artificial ground truth data, for several sample depths ranging from 10^6 samples (40 ms of real time) to the full acquisition of $7.5 \cdot 10^8$ samples (30 s of real time). This allowed us to assess not only the best-case reconstruction quality using the full time series (30 s of acquisition per wavelength, corresponding to 1 hour of experimental time to acquire a full IVDF), but also to quantify how each algorithm performs

with a smaller dataset which would not be as expensive to acquire. First, we performed phase-sensitive detection on the artificial LIF data to deconvolve the modulation, as described in Refs. 23 and 19. We ran this analysis on a 24-core processor with full parallelization. For a downsampling ratio of 50, processing with the TF analysis took only a few seconds to minutes per wavelength (scaling with the sample depth), while the SMI analysis generally took up to 5-10 minutes per wavelength. However, in practice, the rate-determining factor comes from loading the large datasets from disk into the 256-GB RAM of the processing computer. Improvements to this process are possible, for example by sequentially loading and processing single chunks of data at once, as discussed in Ref. 19.

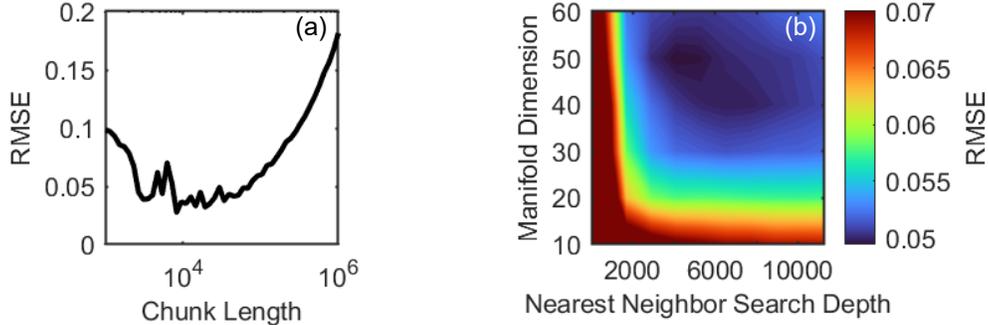


Figure 5. Examples of discrete hyperparameter optimization for both reconstruction methods based on artificially generated ground truth data with a sample depth of $3.2 \cdot 10^8$. (a) TF method chunk length. The root-mean-square error is used to determine optimality. (b) SMI method parameters (number of nearest neighbors, k_{nn} , and manifold dimension, E).

In order to fairly compare the performance of each algorithm, it was necessary to pick hyperparameters which would lead to the lowest error for each sample depth trial within the convergence study. For the SMI analysis, the time-delay index τ was chosen to be one timestep ($2 \mu\text{s}$ after downsampling by 50), due to the fact that using $\tau = 1$ in initial studies with E and k held fixed led to the lowest error across the board. We then performed an approximate discrete optimization of the remaining hyperparameters for each sample depth. This was accomplished by performing the analysis for each ordered pair (E, k) within a discrete grid of possible hyperparameter sets. Figure 5 shows examples of the error as a function of hyperparameter choice for a single sample depth at 300 V, for a sample depth of $3.2 \cdot 10^8$ points. Certain features of this trial were consistent among most of the cases - in particular, the presence of a steep minimum in the error as a function of nearest-neighbor search depth k , as well as chunk length, with a weaker dependence on manifold dimension E .

The existence of an optimal k follows intuitively from the nature of the SMI algorithm: If averaging over too few shadow manifold vectors, the typical $\sim 1/\sqrt{k}$ noise reduction will not be sufficient to improve the SNR. However, if the search is over too many points, the algorithm may be forced to include “false” nearest neighbors for which the dynamical system state is relatively far from the state at the reconstructed time - this will lead to an increase in error. These contributions should be small due to the exponential weighting, however in many cases the false nearest neighbors eventually overwhelm the true neighbors and error begins to increase. A similar statement can be made regarding the optimal chunk length used for transfer function estimation. Too small of a chunk length degrades the quality of the individual FFTs due to poor frequency resolution; however, too large of a chunk length leads to only a small number of chunks, again leading to poor $\sim 1/\sqrt{N_{chunks}}$ noise reduction. The behavior of the error with E at optimum k can be described as a threshold - provided that a sufficiently large manifold dimension is chosen, increasing this parameter further does not significantly impact the result. Increasing E further does cause additional computational expense, which scales roughly linearly with E due to the increased number of matrix rows required. Further, for very large E and τ , the loss of $E(\tau - 1)$ points on either side due to boundary issues with unwrapping the shadow manifold could effectively reduce the sample depth and degrade the reconstruction. This is shown by the slight increase in error with E for $E > 50$.

Figure 6 shows examples of the reconstructed waveforms for the 300-V case shown in Fig. 5 as well as the corresponding sample depth for the 600-V, periodic reference data. The shown data had a pre-PSD SNR of 0.1, and a sample depth of $3.2 \cdot 10^8$ was processed (13 s of acquisition). While the full LIF

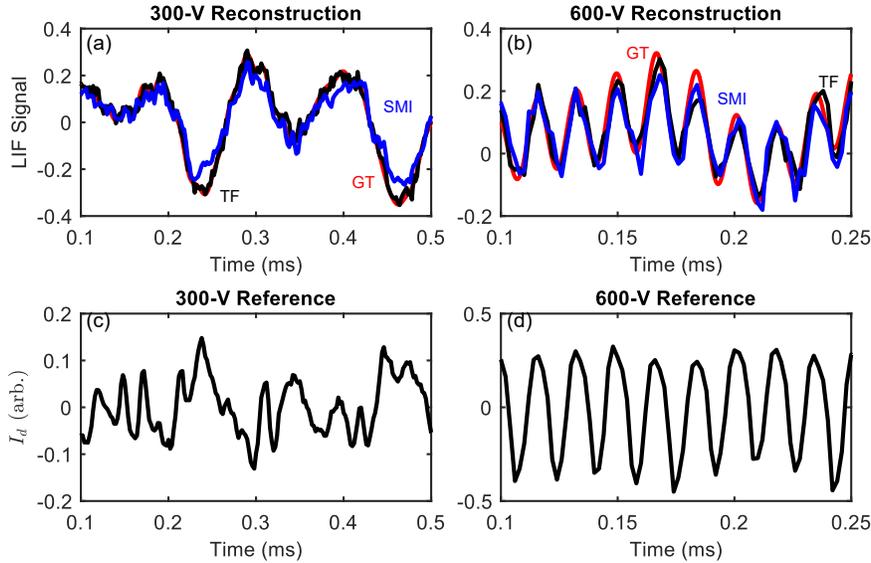


Figure 6. Optimal single-wavelength reconstruction of artificially generated ground truth (GT) data with optimal parameters using the transfer function (TF) and shadow manifold interpolation (SMI) algorithms, for a sample depth of $3.2 \cdot 10^8$. (a) 300 V reconstruction (aperiodic fluctuations). (b) 600 V reconstruction (quasiperiodic fluctuations). (c) Reference for 300-V data (discharge current). (d) Reference for 600-V data (discharge current).

datasets include a sample depth corresponding to 30 seconds, in the interest of computational time we only performed convergence studies up to this point. For both operating conditions, both algorithms are generally successful at reconstructing the waveform, but the transfer function method performs slightly better. For both the aperiodic and quasiperiodic fluctuations, the SMI reconstruction generally captures the shape of the waveform, but in some cases the amplitude is reduced and the reconstruction undershoots the peaks. This may be due to the nature of the averaging process, since extreme values are less common in the range of accessed y -values of the time-series, while the TF method is free from this difficulty since it relies on the frequency domain. This feature generally was not an issue with the noisier data ($\text{SNR} = 0.01$), which is consistent with this interpretation, as the noise increases the codomain of the time series.

Figure 7 displays the results of the full convergence studies with parameter optimization for synthetic data generated from the real discharge oscillations at both of the 300-V 600-V operating conditions, as a function of sample depth. The optimal chunk length increased monotonically as a function of sample depth in all cases, and both datasets with a larger amplitude of added noise required shorter chunks to minimize the error. This can be explained by the trade between resolution in the frequency domain (improved by increasing the chunk length) vs. the need to reduce the noise by averaging over a sufficient number of points (improved by decreasing the chunk length). When more data is supplied, the algorithm can continue to keep the noise level low but further lower error by improving resolution, driving up the chunk length. This hypothesis also explains the need for a shorter chunk length when increasing the noise level for both datasets, since in this case the larger noise might necessitate averaging over more chunks at the expense of frequency resolution.

The optimal nearest-neighbor search depth, k , exhibits an increasing trend which can be explained in a similar way. For longer datasets, we would expect a larger number of times at which the system truly passes through similar phase-space states, so that a larger number of nearest neighbors leads to higher accuracy. We also observe that deeper nearest-neighbor searches are optimal for noisier data, which again follows from the need to average over a larger number of points to reduce the noise. We note that these search depths are many orders of magnitude larger than the depths implemented in other recent work, which is likely a result of the fact that we are here analyzing data with a far lower SNR.^{24–26} No salient trends are evident for the scaling of the optimal manifold dimension. This fact is consistent with the relative independence of the error on E in the example shown in Fig. 5. However, for all tested cases, choosing an E -value between 30 and 50 proved effective and computationally efficient for these datasets.

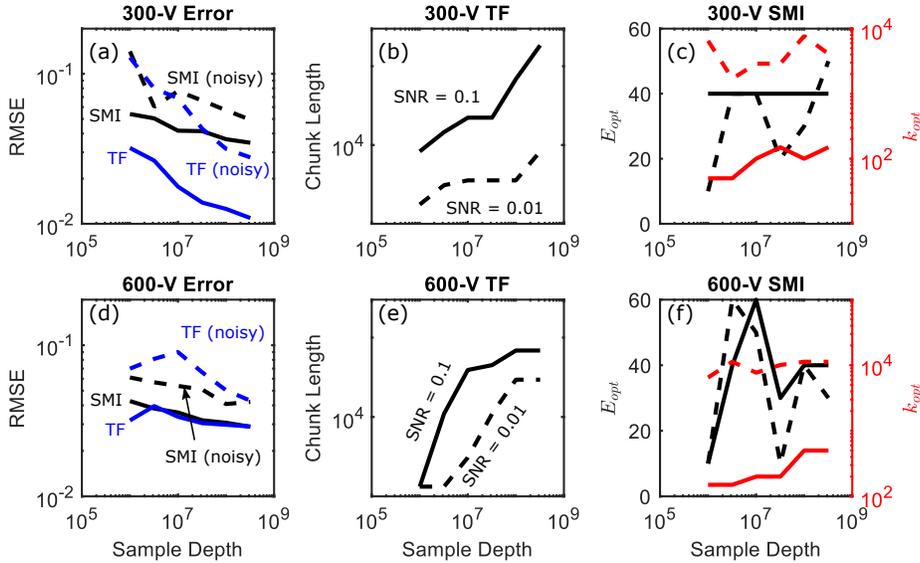


Figure 7. Results of error convergence study with optimal parameters for the artificially generated data. (a) Root mean square error (RMSE) for both reconstruction frameworks on the aperiodic, 300-V data with high noise (SNR = 0.01), and low noise (SNR = 0.1). (b) Optimal chunk length for the TF method applied to the 300-V data. (c) Optimal manifold dimension E (black) and nearest-neighbor search depth k (red) applied to high-noise (dashed lines) and low-noise (solid lines) 300-V data. (d) RMSE for 600-V (quasiperiodic) data. (e) Optimal chunk length for 600 V data. (f) Optimal SMI parameters for 600 V data.

Fig. 7(a) and (d) show the RMSE estimates from the convergence study. Including more points in the sample set generally decreased the error of the reconstruction, with a small number of exceptions. The data with sample depths below 10^7 was likely not sufficient for accurate determination of the signal phase during the PSD process, which could explain the increase in error with sample depth at some of those points. With a smaller amount of noise added, both the TF and SMI methods performed similarly well on the oscillatory, 600-V dataset. When larger noise was added, both methods performed more poorly, but the shadow manifold technique was more resilient to noise for the oscillatory dataset, and converged to the minimum error with less data. For the aperiodic dataset with low noise, the TF method outperformed the other technique by far. When an increased noise level was added, this gap closed, but the transfer function method remained more reliable in this case. One caveat to these results is that the background signal in the real data may not be well-represented by simple white Gaussian noise, since background light from the plasma could occur at specific frequencies and could be coupled to the dynamical system itself.

C. Application to real LIF data

Armed with the optimal hyperparameters determined by the convergence study on artificial data, we then proceeded to perform analysis on all wavelengths of example IVDFs from the LIF experiments. Because the ground truth is not known for the real dataset, we used the optimized hyperparameters from the convergence study with synthetic data at the maximum recorded sample depth, as shown in Table 1. The results of these optimal reconstructions of the real data are shown in Fig. 8. To illustrate the performance on both a limited and full dataset, we present examples with sample depths of $1e7$ and of the full acquisitions, $7.5e8$ samples. The same slicing settings were used as in the ground truth analysis: the time series were downsampled by a factor of 50 and the optimal values for the noisier datasets in Fig. 7 were used. We see that for the periodic dataset (600 V), both methods are effective at reconstructing the IVDF qualitatively, even for a low amount of data. For the low-sample-depth case, there is additional noise present in the TF reconstruction, and the SMI algorithm appears to perform better. However, when trained on the full-length acquisitions, the model reconstructions do not show significant differences. For the 300-V operating condition with aperiodic fluctuations, the SMI reconstruction is significantly better than the TF reconstruction for low sample depth. Again, both reconstructions agree well, with some minor differences, when trained on the full-length time

series. The SMI reconstruction appears to only improve slightly, while the TF method changes significantly with sample depth. This appears to show that the SMI reconstruction performance is better on the real data than suggested by our artificial convergence study.

V_D (V)	SNR	Chunk Length	TF Error	k_{opt}	E_{opt}	SMI Error
300	0.1	184200	0.0109	150	40	0.0347
300	0.01	8300	0.0277	4100	50	0.0494
600	0.1	68700	0.0290	500	40	0.0289
600	0.01	29500	0.0429	11300	30	0.0422

Table 1. Optimal hyperparameters and best error values for the optimization based on the synthetic ground truth LIF data, for the maximum sample depth of $3.2 \cdot 10^8$ at which we performed ground-truth convergence studies.

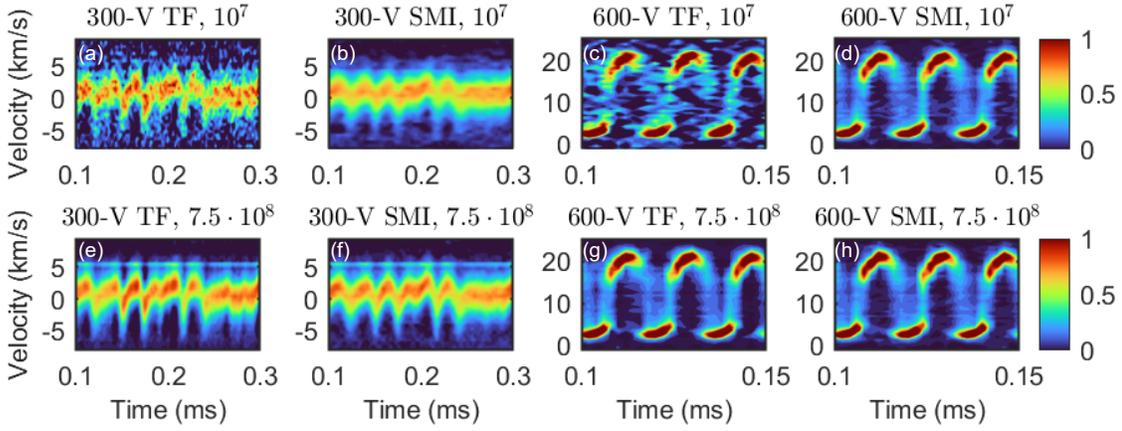


Figure 8. Examples of reconstructions of time-resolved IVDFs from the AEPS thruster on channel centerline in the acceleration region. (a) TF Method, $N_{samp} = 10^7$, (b) SMI Method, $N_{samp} = 10^7$, (c) TF Method, $N_{samp} = 10^7$, (d) SMI Method, $N_{samp} = 10^7$, (e) TF Method, $N_{samp} = 7.5 \cdot 10^8$, (f) SMI Method, $N_{samp} = 7.5 \cdot 10^8$, (g) TF Method, $N_{samp} = 7.5 \cdot 10^8$, (h) SMI Method, $N_{samp} = 7.5 \cdot 10^8$.

To explore the performance on the real data quantitatively, we again performed a convergence study using a wavelength with large LIF signal for each dataset. Since we do not know the actual ground truth for the real data, we computed the RMSE relative to the highest-sample-depth TF reconstruction. We held the hyperparameters constant at the values in Table 1. This shows the approximate convergence rate, but can be misleading if any errors are present from ground truth in that model output. The results of this convergence study are shown in Fig. 9. The trends are similar for the operating conditions with aperiodic oscillations (300 V) and quasiperiodic oscillations (600 V). In both cases, the SMI algorithm achieves a much higher quality reconstruction of the denoised LIF signal for low sample depths, below 10^7 points. However, even for orders of magnitudes shorter acquisitions, the SMI algorithm appears to plateau relatively quickly, while the TF algorithm converges at a faster rate. Clearly there exists a fundamental trade off between these two techniques based on data availability/cost and the needed accuracy of the reconstruction. Referencing Fig. 8, the relatively small improvement offered by the TF method may offer only diminishing returns given its high data rate and depth requirements - in many cases the SMI reconstruction on the smaller dataset may be more than sufficient to gain physical insights about the system with a shorter and less expensive experiment.

Interestingly, this convergence study differs from the ground truth study on artificial noisy datasets, in that the SMI performs better than the TF method on the real data, whereas on the artificial data the TF reconstruction is nearly always better. This difference may stem from our assumption that the noise is Gaussian and random in the artificial dataset - instead, in the real LIF signal, the noise/background light in the PMT may occupy specific frequency ranges which are more effectively filtered out by the PSD process. Alternatively, the SMI method's reliance on the topology of dynamical systems may offer a better

description of the true dynamics, whereas the artificial dataset was created with a transfer function approach that may have been biased toward direct inversion via the TF method. We note also that we obtained these results using the optimal parameters from the simulated dataset parameter study, and further optimization for different dataset generation methods could improve or change the reconstruction quality as a function of sample depth. In the remaining subsections, we discuss unique advantages offered by the SMI technique on a nonlinear dynamical transformation as well as a dataset with multiple fused measurement time series.

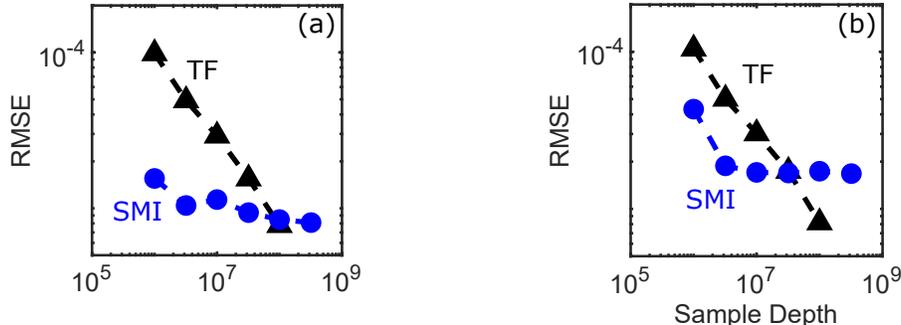


Figure 9. Error convergence for reconstruction of real LIF data from the AEPS thruster, for a single wavelength with high LIF signal. (a) 300 V operating condition with aperiodic oscillations. (b) 600 V operating condition with quasiperiodic operating conditions.

IV. Extended Capabilities of Denoising Algorithms

In this section, we address the capabilities of both denoising algorithms beyond the performance comparison we showed in the preceding. Namely, we investigate the possibility of four unique applications for these denoising algorithms within the context of time-resolved laser diagnostics in an oscillating plasma discharge. First, we characterize the algorithm performance on a synthetic dataset with nonlinear, rather than linear, coupling between the reference and LIF signal. Second, we demonstrate the capability of the SMI technique to perform data stream fusion, i.e. combining multiple reference measurements to enhance the reconstruction quality. Following this, we investigate the ability of these denoising processes to allow direct background subtraction, potentially removing the need for PSD and allowing an increase in diagnostic bandwidth. Finally, we propose a modification of the SMI reconstruction algorithm which performs sparse reconstruction between datasets with different time series, which can be applied to diagnostics with pulsed lasers, such as ITS or TALIF.

A. Application to Nonlinear Dynamics

One of the primary benefits of the SMI algorithm relative to the TF method is that its formulation is not limited to linear relationships between the reference and output signals. Such nonlinear couplings are common in plasma systems due to the coupled nonlinear differential equations that describe the motion of charged particles, so this is a key advantage. For the real data on which we tested these algorithms in the previous section, the high similarity of the SMI and TF reconstructions suggests that at least for those cases, this limitation of TF did not present itself. This is validating for past results based on only linear analysis of this LIF data. However, in order to demonstrate the effect of nonlinearity on the reconstruction, in this section, we generate artificial training data using a different, nonlinear construction.

To generate a dataset with a nonlinear coupling to the real reference (discharge current) data, in general we might define some example dynamical system or nonlinear transfer function to apply in the forward direction. However, to demonstrate an example case, it is sufficient to simply apply the artificial linear transfer function defined in the previous section, followed by a nonlinear transformation of the result. In this case, we simply square the result, i.e. $F_{nonlin}(t) = [F_{lin}(t)]^2$, where $F_{lin}(t)$ is the inverse FFT of $\tilde{H}(\omega, \lambda)\tilde{I}(\omega)$. Not only does this provide a simple nonlinear coupling between the input and output, but it also produces a semipositive-definite output. This is an interesting test case for optical diagnostics, because the measured signal without noise should indeed retain semipositive-definiteness (i.e., the system cannot measure negative

photons).

Figure 10 shows examples of the reconstructions learned for the nonlinear dataset for both methods, for an SNR value of 1, and a sample depth of 10^7 . In the 300-V, aperiodic case, the linear TF method is not able to reconstruct the dynamics in any obvious sense at this sample depth, while the SMI method does successfully recreate the general shape of the ground truth waveform. However, the intensity of the spikes in the reconstruction are only roughly half of the extreme values in the ground truth data. This may be explained by the relative rarity of these more extreme values, which would have limited examples in the training dataset - relatedly, the knn-averaging set may include a number of “false” nearest neighbors which effectively pulls the output closer to more common values. This interpretation is consistent with the fact that for the noisier data, the SMI reconstruction is closer to the height of the ground truth, since in this case the noise increases the codomain of the noisy dataset, allowing it to span a wider range. For the more periodic, 600-V case, both methods are successfully able to match the ground truth with higher fidelity. The nonlinear SMI reconstruction performs slightly better qualitatively in that it more consistently captures the local minimums, but both methods undershoot the highest peaks, likely for similar reasons as the 300-V case. We note that for this example, we used the optimal parameters inferred for the data generated using a linear transfer function model, and further optimization with these dynamics or with a larger dataset could lead to better reconstructions.

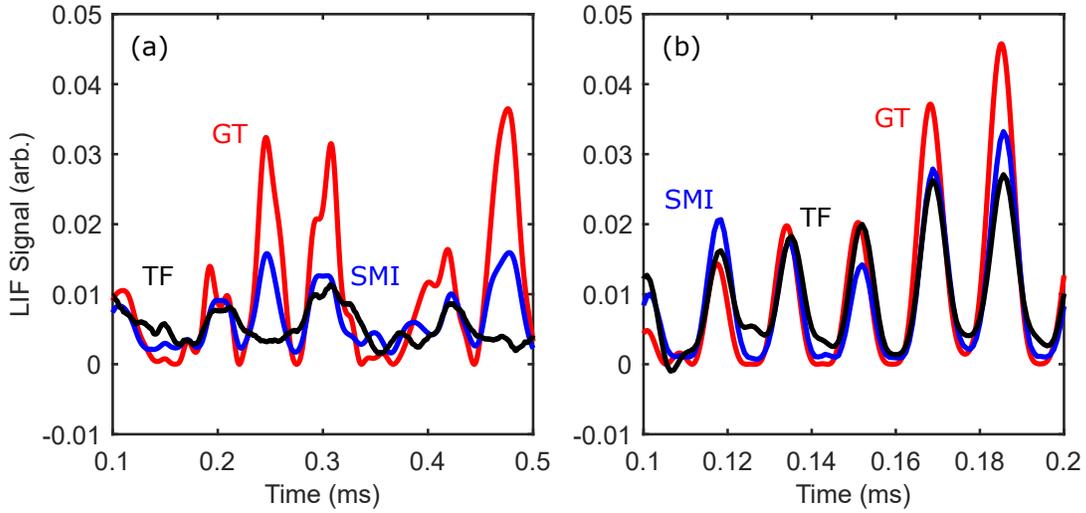


Figure 10. Reconstructions of a synthetic nonlinear mapping at a sample depth of 10^7 and $\text{SNR} = 1$. (a) 300-V (aperiodic) operating condition, (b) 600 V (quasiperiodic) operating condition.

B. Fusing Reference Data Streams

Another advantage offered by the SMI technique, which is demonstrated in Ref. 24, is the flexibility of the shadow manifold technique to the use of multidimensional data streams. Since the reconstruction relies on the Euclidean distance metric, it is symmetric with respect to the order of the components in the time-lag vector (c.f. Eqs. 3 and 4). Therefore, an additional measurements with their own time lag vectors can simply be “stacked” to build a fused shadow manifold. To illustrate this, suppose there are two measurement time series, $R_1(t)$ and $R_2(t)$, measured simultaneously. We can then construct the composite time-lag vector as

$$\vec{M}_{R,\text{Fusion}}(t) = [R_1(t), R_1(t + \tau), \dots, R_1(t + E\tau), R_2(t), R_2(t + \tau), \dots, R_2(t + E\tau)], \quad (7)$$

and carry out the SMI inference as usual. Further modifications of this framework are possible: data streams measured over different sample rates/depths/times could be combined, and the average could include weighting parameters that bias the nearest neighbor search to favor one or the other reference. However, in the present work we restrict to equal weighting and sampling for two reference signals.

This idea has a direct application to the existing AEPS LIF data shown in a previous section. For these datasets, in addition to acquiring thruster discharge current waveforms, the cathode-to-ground voltage

was simultaneously recorded as an alternative reference signal.²³ Figure 11 shows the results of combining the current and voltage waveforms as simultaneous references for the SMI technique according to Eq. 7, with both TF and SMI reconstructions using solely the current or voltage as reference, respectively. It is evident by comparing Figures 11(a) and (b), and (d) and (e), respectively, that slightly different features arise when using the discharge current alone ((a) and (d)) vs. the cathode voltage alone ((b) and (e)), but the waveforms are largely similar. Due to the small sample depth of 10^7 for this study, the SMI analysis ((a) and (b)) tended to over-smooth the data, while the TF method produced a crisper waveform but with larger background noise ((d) and (e)). Figure 11(c), which shows the result of the SMI data fusion trial, does not differ significantly from Fig. 11(a), suggesting that the addition of the second data stream did not change the selected nearest neighbor sets in an important way. However, as is visible in Fig. 11(f), the cathode to ground voltage exhibits some content at higher frequencies than the discharge current. It is therefore possible that if hardware and software limitations on diagnostic bandwidth are removed, as will be discussed in the following section, that this method of data fusion could provide further insight into the higher-frequency behavior of the plume ion dynamics. Further, in future implementations of this technique, the waveforms of both reference signals should be normalized so that the larger signal does not bias the distance calculation.

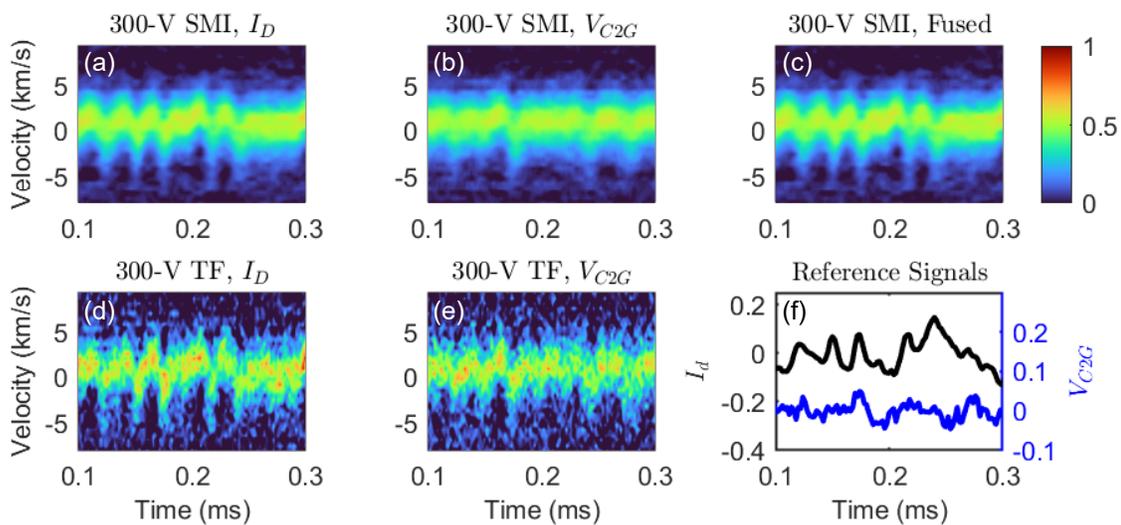


Figure 11. Impact of fusing multiple reference signals to generate the reconstruction, for the 300-V operating condition and a sample depth of 10^7 . (a) SMI method with discharge current reference. (b) SMI method with cathode voltage reference. (c) SMI method with fused discharge current and cathode voltage references. (d) TF method with discharge current reference. (e) TF method with cathode voltage reference. (f) Reference signals: 1 = discharge current, 2 = cathode voltage.

C. Increasing Time Resolution with Background Subtraction

In this section, we explore the possibility of leveraging both the TF and SMI techniques in order to improve limitations on the bandwidth of time-resolved LIF data analyzed using the previously described techniques. We first describe current constraints on time-resolution of past aperiodic, time-resolved optical diagnostics arising from hardware choices as well as the nature of the phase-sensitive detection which is typically used. We then present frameworks using both mapping algorithms to eliminate the need for phase-sensitive detection altogether via a background subtraction scheme.

The heritage acquisition strategy for time-resolved LIF, described in detail earlier in the paper, is designed around the need to isolate the relatively weak LIF signal from the intense background signal, consisting of non-stimulated plasma emission and electrical noise. This is typically accomplished by modulating the incoming laser and processing the measured signal with phase-sensitive detection (PSD), performed either with an analog lock-in amplifier or via digital post-processing. The PSD process involves a low-pass filtering step which effectively averages the signal over some time constant - this low-pass filter sets the limit on time-resolution. Existing time-resolved LIF acquisition strategies significantly improve this upper limit by chopping the laser at a faster rate, for example by using an acousto-optic modulator (AOM), which can

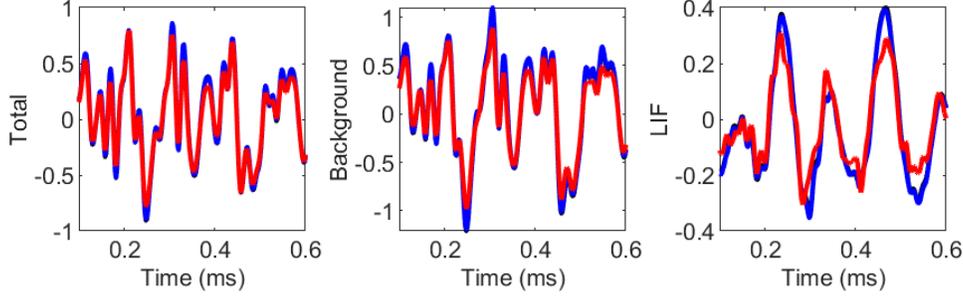


Figure 12. Reconstruction via learned background subtraction with TF and SMI methods for the 300 V operating condition. (a) Reconstruction of total signal. (b) background only. (c) LIF signal only.

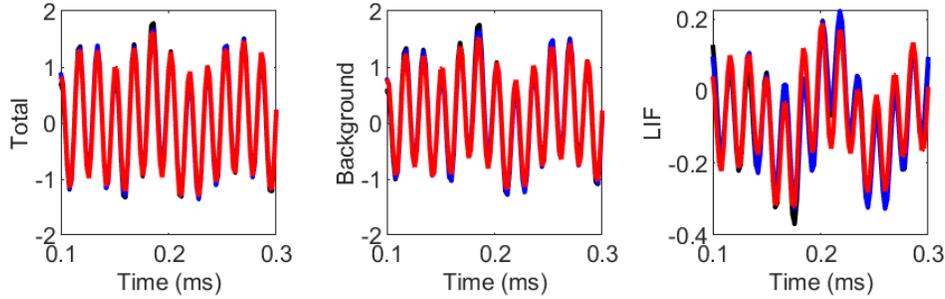


Figure 13. Reconstruction via learned background subtraction with TF and SMI methods for the 600 V operating condition. (a) Reconstruction of total signal. (b) background only. (c) LIF signal only.

achieve MHz-scale modulation, in place of a mechanical chopper, which is limited to several kHz. Even with these faster choppers, the resulting lock-in time constant required to recover sufficient SNR has historically limited the upper limit to frequency resolution of around 100 kHz. That said, higher frequency LIF has been achieved with a photon counting technique,³⁸ but stabilization of the breathing mode by an externally applied potential oscillation was required to provide a reference phase for averaging. However, numerous phenomena in Hall thrusters are known to occur at much higher rates than 100 kHz, including transit time oscillations³⁹, ion acoustic turbulence⁴⁰, and other plasma instabilities⁸. There is therefore a significant motivation to improve the time resolution of time-resolved LIF in a passive Hall thruster with naturally occurring, aperiodic oscillations.

In this section, we propose a method to accomplish this and demonstrate its efficacy on artificially generated data. The premise is as follows: Instead of performing implicit background subtraction by modulating the input signal, we record separate, equal-time acquisitions with the laser on and off, composing an "LIF" and "background" traces respectively. We then apply a denoising mapping based on a reference signal to each acquisition, for example either the TF or SMI method. The resulting mappings enable the reconstruction of both the total (background + LIF) and background-only signals synchronized to a single reference time series. Provided that a sufficiently accurate reconstruction of each signal can be computed and that the data acquisition section possesses sufficient dynamic range, background subtraction can then be performed directly to isolate the LIF signal of interest.

We demonstrated this background subtraction method with a numerical example. To do so, we generated artificial data with the same hypothetical transfer function as used in the preceding section. However, we also added artificial "background light", generated with a unique transfer function with larger average amplitude. This background signal physically represents the variation in the light emitted by the plasma near the fluorescence wavelength due to processes not stimulated by the laser. For example, collisional excitation and fluorescence of the targeted metastable transition (or other transitions with similar fluorescence wavelengths) will occur without the laser present. This background signal would presumably be dynamically tied to global plasma fluctuations via correlated local changes in the electron temperature or density. Without modulating

the laser and performing PSD to remove the background fluctuations that do not occur at the modulation frequency, the additional fluorescence due to laser excitation of the metastable state must be found by subtracting this background signal.

As shown in Figs. 12 and 13, both the TF and SMI methods are able to successfully reconstruct the LIF signal for both the aperiodic and periodic Hall thruster discharge current fluctuations, despite the intense simulated variations in background light, by learning a mapping for the synthetic background light itself. In reality, the background signal in the LIF acquisition likely consists of an additive blend of effectively random noise, which is either loosely or not at all correlated with the dynamical/LTI system of the Hall thruster, and strongly correlated background fluctuations for which a mapping can be efficiently learned and subtracted. Additionally, for sufficiently long datasets, the averaging inherent to both the TF and SMI nearest neighbor algorithms will tend to reduce the effect of any random noise. Regardless, this result demonstrates that in principle, the upper limit on time resolution for LIF can be improved by leveraging asynchronous background subtraction with a simultaneous reconstruction strategy.

D. Sparse Reconstruction for Pulsed Diagnostics

For extremely low-light diagnostics which require a more powerful, Q-switched laser, such as incoherent Thomson scattering (ITS) or two-photon absorption laser-induced fluorescence (TALIF), the sample rate is limited to the repetition rate of the laser pulse. For common laser systems, the pulse rate rarely exceeds tens of Hertz. This limitation makes achieving time resolution difficult when the signal-to-noise is sufficiently low that single-shot acquisition is not possible, as is often the case in low-temperature plasmas. In recent work, Antozzi et al. showed that reconstruction of a pulsed TALIF signal is possible for periodic oscillations by performing phase-binning.²² However, typical Hall thruster oscillations are not periodic in general, as the AEPS examples from the previous sections show. Therefore, aperiodic reconstruction techniques offer a possible path to achieving general time resolution with these informative but low-signal diagnostics. In the following, we describe a framework for adapting the SMI analysis technique to pulsed diagnostics with limited sample rate, and demonstrate a numerical example with synthetic data.

To achieve time resolution with an SMI reconstruction for a pulsed diagnostic, we must measure a reference signal time series corresponding to each laser pulse which captures the plasma dynamics on a much faster scale than the 0.1-s pulse period. One could accomplish this experimentally by triggering a recording of the discharge current onto a fast data acquisition card upon each laser pulse. In our example, we sample the discharge current value at a rate of 1 MS/s, and collect only 1024 samples for each laser shot. Despite the sparsity of our laser pulse rate relative to the faster plasma dynamics of interest, we are still able to construct a shadow manifold and exploit Takens’s theorem to build a reconstruction for this situation. The vector of 1024 samples of the discharge current, or some downsampled subset of it, is effectively the shadow manifold vector for the reference signal measurement. Meanwhile, the count value of each CCD pixel, corresponding to the Thomson spectrum intensity at a particular wavelength, is a very simple, 1-dimensional, shadow manifold vector corresponding to the output. While a shadow manifold vector with dimension 1 would be far too low to practically embed the state space of a realistic dynamical system, this is not a problem for our purpose: referring to Fig. 3, it is the *reference* vector which is used to identify similarity in the state space, and this we are able to record with essentially arbitrary sample rate and depth.

The averaging process then can proceed as usual: for each time within the reference time series chosen for reconstruction, the shadow manifold vector is computed, and the k nearest neighbors are found within the training dataset. Each nearest neighbor corresponds to the timing of a single laser pulse, for which the value of each CCD pixel represents a 1-dimensional, noisy shadow manifold vector for the output signal. The values of each pixel may then be averaged (with exponential weighting) for the set of nearest neighbors, effectively resulting in an ensemble average of the Thomson spectrum during subsets of laser pulses during which the thruster was exhibiting similar behavior.

In Fig. 14, we demonstrate the validation of this technique on a synthetic dataset. We generated simple artificial ground truth reference and Thomson data (for a single wavelength) by composing the sum of several sinusoidal functions at similar frequencies with different weights. We include 18000 laser shots, corresponding to 30 minutes of acquisition time. For the SMI reconstruction of the discharge current, we use a manifold dimension of $E = 10$, sampled with a rate of 1 MS/s, and use a nearest-neighbors search depth of $k = 1000$. With no noise added, the SMI algorithm is able to reconstruct the dataset fairly well despite the sparseness of the laser sampling and the need to map between measurements with different shadow manifold parameters. We note that this method relies on the assumption that the laser pulse rate is uncorrelated with the phase

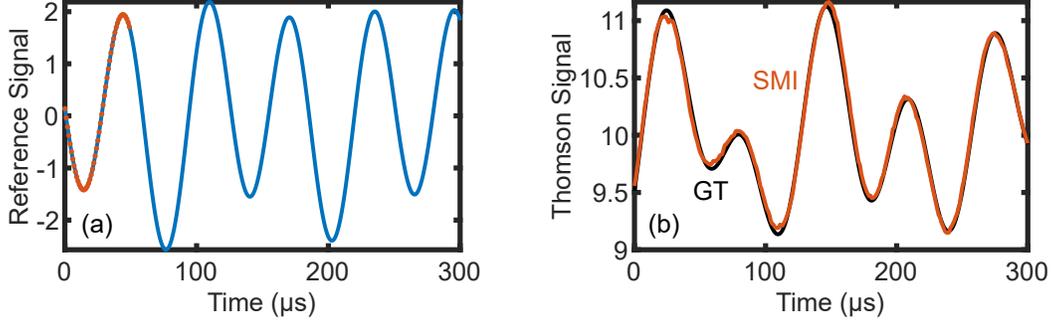


Figure 14. Demonstration of the time-resolved Thomson scattering algorithm on artificially generated data. (a) Reference data for a particular laser shot, with the points used to construct a 50-dimensional shadow manifold vector shown in orange. The simulated laser pulse occurs at the beginning of the vector (first shot at $t = 0$). (b) Ground truth synthetic Thomson signal (black) and sparse SMI reconstruction (orange).

of the discharge current oscillations, in the case that they are sufficiently periodic that such a phase could be well-defined – this is so that the times of the laser pulses must represent a sufficiently fair representation of the phase space topology.

While a feasible proof of concept, to implement this technique on low-signal diagnostics will be subject to difficult-to-overcome constraints due to the high noise level. This can potentially be overcome with long acquisitions, but with diminishing returns on ensemble-averaged noise reduction. However, the SMI technique provides a useful framework for mapping diverse measurements of the complex thruster system, even with orders-of-magnitude variation between the sample rates of the reference signal and scattering signal, and in chaotic or fully aperiodic oscillation regimes.

V. Conclusions

In summary, we have explored in this work the performance and features of two algorithms for reconstructing and denoising time-resolved LIF data in a Hall thruster plume. These methods have direct application to resolving particle dynamics on the time scales of low-frequency plasma fluctuations. These techniques included TF estimation, which works by learning a linear, frequency-domain-based map from a reference signal to a denoised version of the signal; and SMI, a nonlinear, time-domain based method based on approximating the system state space and averaging over times of similar behavior. To test both algorithms, we generated synthetic data as a known ground truth and added pseudorandom noise, then used brute-force parameter optimization to determine the optimal hyperparameter choices for each algorithm. This parameter optimization resulted in the conclusion that two hyperparameters, the TF chunk length and the nearest-neighbor search depth, must increase with the sample depth of the measurement to yield the lowest error. The optimal nearest-neighbor search depth was found to be orders of magnitude larger than that employed in other works, likely due to the reduced SNR characteristic of the LIF data. Based on these error studies, we found that the SMI technique is more robust to noise than the TF technique, but experiences diminishing returns with larger datasets.

Following this parameter optimization, we demonstrated the performance of these algorithms on a real LIF dataset acquired in the plume of the AEPS Hall thruster for two operating conditions: one with aperiodic fluctuations in the discharge current, and one with intense, quasiperiodic oscillations near 50 kHz. Using the optimal parameters from the convergence study on the high-noise simulated data, the SMI algorithm performed even better than expected on the real data for a much smaller amount of sample depth points. While the TF reconstruction was significantly affected by the noise for smaller sample depths and required long, expensive acquisitions to accurately resolve the IVDF, the SMI technique was able to effectively determine the dynamics with a much smaller number of points, albeit with increased computational cost. These findings both validate previous results based on the TF analysis with an independent comparison, and demonstrate that high-quality reconstructions can be obtained with lower requirements on sample rate and depth. This finding lessens the significant experimental time typically needed for time-resolved LIF

acquisition.

After assessing the algorithmic performance on LIF data, we demonstrated the application of these algorithms to four advanced optical diagnostic use cases. We first used synthetic data to test the SMI algorithm’s ability to reconstruct nonlinear couplings between the reference and output data, which the TF method is not able to capture. After this, we returned to the AEPS study to explore the ability to seamlessly fuse multiple reference data streams into the reconstruction. We also investigated the ability of both SMI and TF estimation to learn a deterministic signal for the background light, finding that both methods can allow background subtraction and mitigating the need for the bandwidth limitations of PSD. Without the PSD step, LIF time resolution could in this way exceed 1 MHz and reveal higher-frequency thruster physics. However, the problem remains of finding a reference signal with sufficiently strong coupling to these higher-frequency dynamics—for example, local, microscopic turbulence may average out of global metrics such as the discharge current. Finally, we explored the application of the SMI algorithm to time-resolution for pulsed laser diagnostics, such as Thomson scattering or TALIF, in low-temperature/density plasma sources. We demonstrated a method for sparse SMI reconstruction across measurements with different sample rates, which performs well on noiseless data. However, to apply this technique to extremely noisy pulsed diagnostic data, it may be necessary to collect a significant amount of data and reoptimize parameters for this technique. Regardless, these techniques and applications represent a key step for gleaning the maximum information from limited diagnostics and experiments. By accessing novel regimes of temporal behavior non-perturbatively and linking multiple plasma diagnostics through data fusion and synchronization, we grow closer to a more complete description of Hall thruster discharge dynamics which is needed to model and predict behavior and performance with high fidelity.

Acknowledgments

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004), and was supported by a NASA Space Technology Graduate Research Opportunity, Grant 80NSSC20K1229.

References

- ¹ Goebel, D. M., Katz, I., and Mikellides, I. G., *Fundamentals of electric propulsion*, John Wiley & Sons, 2023.
- ² Brophy, J., Polk, J., Randolph, T., and Dankanich, J., “Lifetime Qualification Standards for Electric Thrusters for Deep-Space Missions,” *44th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, American Institute of Aeronautics and Astronautics, Hartford, CT, 2008. <https://doi.org/10.2514/6.2008-5184>.
- ³ Hofer, R., Kamhawi, H., Peterson, P., Polk, J., Welander, B., Inaba, D., Blackner, G., Ferraiuolo, B., Frieman, J., Lobbia, R., Chaplin, V., Zubair, J., and Branch, N., “Development and Qualification of a 12-kW Hall Thruster for Deep-Space Missions,” *73rd International Astronautical Congress (IAC)*, IAF, Paris, France, 2022, pp. IAC-22-C4-5-15.
- ⁴ Lobbia, R. B., “A time-resolved investigation of the Hall thruster breathing mode,” Ph.D. thesis, University of Michigan, 2010.
- ⁵ Laffleur, T., Chabert, P., and Bourdon, A., “The origin of the breathing mode in Hall thrusters and its stabilization,” *Journal of Applied Physics*, Vol. 130, No. 5, 2021.
- ⁶ Jorns, B. A., and Hofer, R. R., “Plasma oscillations in a 6-kW magnetically shielded Hall thruster,” *Physics of Plasmas*, Vol. 21, No. 5, 2014.
- ⁷ Jorns, B. A., Cusson, S. E., Brown, Z., and Dale, E., “Non-classical electron transport in the cathode plume of a Hall effect thruster,” *Physics of Plasmas*, Vol. 27, No. 2, 2020.
- ⁸ Tsikata, S., Héron, A., and Honoré, C., “Hall thruster microturbulence under conditions of modified electron wall emission,” *Physics of Plasmas*, Vol. 24, No. 5, 2017.
- ⁹ Mikellides, I. G., and Lopez Ortega, A., “Growth of the lower hybrid drift instability in the plume of a magnetically shielded Hall thruster,” *Journal of Applied Physics*, Vol. 129, No. 19, 2021.

- ¹⁰ Lopez Ortega, A., Mikellides, I. G., and Katz, I., “Hall2De numerical simulations for the assessment of pole erosion in a magnetically shielded Hall thruster,” *34th International Electric Propulsion Conference*, Electric Rocket Propulsion Society, Hyogo-Kobe, Japan, 2015, pp. IEPC–2015–249.
- ¹¹ Lopez Ortega, A., Mikellides, I. G., Chaplin, V. H., Huang, W., and Frieman, J. D., “Anomalous ion heating and pole erosion in the 12.5-kW Hall effect rocket with magnetic shielding (HERMeS),” *AIAA Propulsion and Energy 2020 Forum*, 2020, p. 3620.
- ¹² Roberts, P. J., Chaplin, V. H., Lopez Ortega, A., and Mikellides, I. G., “Impact of low-frequency oscillations on ion energy in a high-power magnetically shielded Hall thruster,” *Journal of Applied Physics*, Vol. 131, No. 3, 2022.
- ¹³ Roberts, P. J., Jorns, B. A., and Chaplin, V. H., “Experimental characterization of wave-induced azimuthal ion velocities in a hollow cathode plume,” *AIAA SciTech 2022 Forum*, 2022, p. 1561.
- ¹⁴ Chaplin, V. H., Jorns, B. A., Lopez Ortega, A., Mikellides, I. G., Conversano, R. W., Lobbia, R. B., and Hofer, R. R., “Laser-induced fluorescence measurements of acceleration zone scaling in the 12.5 kW HERMeS Hall thruster,” *Journal of Applied Physics*, Vol. 124, No. 18, 2018.
- ¹⁵ Su, L. L., Marks, T. A., and Jorns, B., “Investigation into the efficiency gap between krypton and xenon operation on a magnetically shielded Hall thruster,” *International Electric Propulsion Conference, Massachusetts Institute of Technology, Cambridge, MA, USA*, 2022.
- ¹⁶ Vincent, B., Tsikata, S., and Mazouffre, S., “Incoherent Thomson scattering measurements of electron properties in a conventional and magnetically-shielded Hall thruster,” *Plasma Sources Science and Technology*, Vol. 29, No. 3, 2020, p. 035015.
- ¹⁷ Roberts, P. J., and Jorns, B. A., “Laser Measurement of Anomalous Electron Diffusion in a Crossed-Field Plasma,” *Physical Review Letters*, Vol. 132, No. 13, 2024, p. 135301.
- ¹⁸ Wegner, T., Thompson, S. J., Williams, J., and Yalin, A. P., “Two-photon absorption laser induced fluorescence (TALIF) of neutral xenon in a Hall effect thruster plasma,” *AIAA Propulsion and Energy 2021 Forum*, 2021, p. 3391.
- ¹⁹ Durot, C., Gallimore, A., and Smith, T., “Validation and evaluation of a novel time-resolved laser-induced fluorescence technique,” *Review of Scientific Instruments*, Vol. 85, No. 1, 2014.
- ²⁰ Lucca Fabris, A., Young, C. V., and Cappelli, M. A., “Time-resolved laser-induced fluorescence measurement of ion and neutral dynamics in a Hall thruster during ionization oscillations,” *Journal of Applied Physics*, Vol. 118, No. 23, 2015.
- ²¹ Dale, E. T., and Jorns, B. A., “Experimental characterization of Hall thruster breathing mode dynamics,” *Journal of Applied Physics*, Vol. 130, No. 13, 2021.
- ²² Gottfried, J. A., Antozzi, S., Stienike, J., Thompson, S. J., Williams, J. D., and Yalin, A. P., “Temporally resolved relative krypton neutral density during breathing mode of a hall effect thruster recorded by TALIF,” *Journal of Electric Propulsion*, Vol. 3, No. 1, 2024, p. 9.
- ²³ Chaplin, V., Lobbia, R., Lopez Ortega, A., Mikellides, I., Hofer, R., Polk, J., and Friss, A., “Time-resolved ion velocity measurements in a high-power Hall thruster using laser-induced fluorescence with transfer function averaging,” *Applied Physics Letters*, Vol. 116, No. 23, 2020.
- ²⁴ Eckhardt, D., Koo, J., Martin, R., Holmes, M., and Hara, K., “Spatiotemporal data fusion and manifold reconstruction in Hall thrusters,” *Plasma Sources Science and Technology*, Vol. 28, No. 4, 2019, p. 045005.
- ²⁵ Baird, M., McGee-Sinclair, R., Lemmer, K., and Huang, W., “Time-resolved ion energy measurements using a retarding potential analyzer,” *Review of Scientific Instruments*, Vol. 92, No. 7, 2021.
- ²⁶ Thomas, A., and Lemmer, K., “Time-resolved ion energy measurements using a retarding potential analyzer for electric propulsion applications,” *Review of Scientific Instruments*, Vol. 95, No. 2, 2024.
- ²⁷ George, E., Chan, C. E., Dimand, G., Chakmak, R. M., Falcon, C., Eckhardt, D., and Martin, R., “Decomposing Signals from Dynamical Systems Using Shadow Manifold Interpolation,” *SIAM Journal on Applied Dynamical Systems*, Vol. 20, No. 4, 2021, pp. 2236–2260. <https://doi.org/10.1137/20M1350923>, URL <https://epubs.siam.org/doi/10.1137/20M1350923>.

- ²⁸ Herman, D. A., Tofil, T. A., Santiago, W., Kamhawi, H., Polk, J. E., Snyder, J. S., Hofer, R. R., Picha, F. Q., Jackson, J., and Allen, M., “Overview of the development and mission application of the advanced electric propulsion system (AEPS),” *International Electric Propulsion Conference*, 2018.
- ²⁹ Cedolin, R. J., *Laser-induced fluorescence diagnostics of xenon plasmas*, Stanford University, 1997.
- ³⁰ Huang, W., Gallimore, A., and Smith, T., “Two-axis laser-induced fluorescence of singly-charged xenon inside a 6-kW Hall thruster,” *49th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, 2011, p. 1015.
- ³¹ Young, C., Fabris, A. L., MacDonald-Tenenbaum, N., Hargus, W., and Cappelli, M., “Time-resolved laser-induced fluorescence diagnostics for electric propulsion and their application to breathing mode dynamics,” *Plasma Sources Science and Technology*, Vol. 27, No. 9, 2018, p. 094004.
- ³² Huang, W., Drenkow, B., and Gallimore, A., “Laser-induced fluorescence of singly-charged xenon inside a 6-kW Hall thruster,” *45th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, 2009, p. 5355.
- ³³ Huang, W., Kamhawi, H., and Haag, T., “Plasma oscillation characterization of NASA’s HERMeS Hall thruster via high speed imaging,” *52nd AIAA/SAE/ASEE Joint Propulsion Conference*, 2016, p. 4829.
- ³⁴ Heath, W. P., “Choice of weighting for averaged nonparametric transfer function estimates,” *IEEE transactions on automatic control*, Vol. 52, No. 10, 2007, pp. 1914–1920.
- ³⁵ Takens, F., “Detecting strange attractors in turbulence,” *Dynamical Systems and Turbulence, Warwick 1980: proceedings of a symposium held at the University of Warwick 1979/80*, Springer, 2006, pp. 366–381.
- ³⁶ Cao, L., “Practical method for determining the minimum embedding dimension of a scalar time series,” *Physica D: Nonlinear Phenomena*, Vol. 110, No. 1-2, 1997, pp. 43–50.
- ³⁷ Fraser, A. M., and Swinney, H. L., “Independent coordinates for strange attractors from mutual information,” *Physical Review A*, Vol. 33, No. 2, 1986, pp. 1134–1140. <https://doi.org/10.1103/PhysRevA.33.1134>, URL <https://link.aps.org/doi/10.1103/PhysRevA.33.1134>.
- ³⁸ Vaudolon, J., Balika, L., and Mazouffre, S., “Photon counting technique applied to time-resolved laser-induced fluorescence measurements on a stabilized discharge,” *Review of Scientific Instruments*, Vol. 84, No. 7, 2013.
- ³⁹ Barral, S., Makowski, K., Peradzyński, Z., and Dudeck, M., “Transit-time instability in Hall thrusters,” *Physics of Plasmas*, Vol. 12, No. 7, 2005.
- ⁴⁰ Katz, I., Lopez Ortega, A., Jorns, B., and Mikellides, I. G., “Growth and saturation of ion acoustic waves in Hall thrusters,” *52nd AIAA/SAE/ASEE Joint Propulsion Conference*, 2016, p. 4534.