

A Bayesian approach to time-domain Photonic Doppler Velocimetry analysis

J. R. Allison,¹ R. Bordas,¹ J. Read,¹ G. Burdiak,¹ V. Beltrán,¹ N. Joiner,¹ H. Doyle,¹ N. Hawker,¹ J. Skidmore,¹ T. Ao,² A. Porwitzky,² D. Dolan,³ B. Farfan,² C. Johnson,² and A. Hansen²

¹*First Light Fusion Ltd., Unit 10 Oxford Pioneer Park, Yarnton, OX5 1QU, UK*

²*Sandia National Laboratories, Albuquerque, New Mexico 87185, USA*

³*Physics Department and Institute for Shock Physics, College of Arts and Sciences, Washington State University, WA 99163, USA*

(*Electronic mail: james.allison@firstlightfusion.com.)

(Dated: 14 July 2025)

Photonic Doppler Velocimetry (PDV) is an established technique for measuring the velocities of fast-moving surfaces in high-energy-density experiments. In the standard approach to PDV analysis, a short-time Fourier transform (STFT) is used to generate a spectrogram from which the velocity history of the target is inferred. The user chooses the form, duration and separation of the window function. Here we present a Bayesian approach to infer the velocity directly from the PDV oscilloscope trace, without using the spectrogram for analysis. This is clearly a difficult inference problem due to the highly-periodic nature of the data, but we find that with carefully chosen prior distributions for the model parameters we can accurately recover the injected velocity from synthetic data. We validate this method using PDV data collected at the STAR two-stage light gas gun at Sandia National Laboratories, recovering shock-front velocities in quartz that are consistent with those inferred using the STFT-based approach, and are interpolated across regions of low signal-to-noise data. Although this method does not rely on the same user choices as the STFT, we caution that it can be prone to misspecification if the chosen model is not sufficient to capture the velocity behavior. Analysis using posterior predictive checks can be used to establish if a better model is required, although more complex models come with additional computational cost, often taking more than several hours to converge when sampling the Bayesian posterior. We therefore recommend it be viewed as a complementary method to that of the STFT-based approach.

I. INTRODUCTION

The velocity of fast-moving surfaces (such as a shock front or a projectile) is an important diagnostic for experiments in the field of shock physics¹. Two velocimetry methods are widely used, the Velocity Interferometer System for Any Reflector (VISAR^{2,3}) and Photonic Doppler Velocimetry (PDV^{4,5}); both use laser interferometry, but differ in the details of how velocity is inferred from the interference pattern. In the case of VISAR, which acts as a delay-leg interferometer, the variation in light intensity is directly proportional to the velocity of the surface, while for PDV the interference fringes are proportional to displacement and so must be differentiated with respect to time (usually by Fourier analysis). This requires a PDV system to have sufficient recording bandwidth to measure high velocities, although the advent of digital acquisition systems in excess of 10 GHz has somewhat negated this issue. VISAR has the advantage of velocimetry resilience in the presence of rough surfaces (hence, ‘any reflector’) while PDV can exhibit signal drop-out due to rough surfaces or multi-mode degraded fiber delivery. Conversely, PDV can be used to measure multiple velocities simultaneously, which is not possible with VISAR. Ultimately, since PDV is operationally simpler it has been widely adopted by the research community as an adjunct or replacement velocimetry method to VISAR⁶.

A full description of the standard approach to PDV analysis was given in the 2020 review by Dolan⁷, which we briefly summarize here. Near-infrared laser light (typically $\lambda \approx 1550$ nm) is reflected off a moving target and mixed with reference light (either from the same source or from a second laser) to generate a time-varying intensity signal at the receiver, the read-out

of which is captured by a digital oscilloscope. A short-time Fourier transform (STFT) is then applied to the voltage versus time data to create a time-frequency representation of the signal (a spectrogram), from which dominant spectral components can be more easily identified and a velocity history extracted. The inferred velocity history is dependent on the choices made by the user that include the form, duration and separation of the STFT window function, and the method to identify and characterize the time-varying spectrogram signal (typically using either peak finding, curve fitting or centroid analysis). The implicit time-scale is set by the STFT window function, which necessarily leads to temporal broadening of the signal that can obscure behavior at smaller time intervals. This is particularly evident at points of rapid acceleration/deceleration or discontinuities. The precision of the velocity measurement is ultimately limited by the sampling rate, signal noise fraction, and analysis duration⁸, although other sources of error, such as laser speckle interference⁹, are expected to contribute¹⁰.

Bayesian inference is a probabilistic approach to data analysis that has gained traction in many research disciplines because it provides a robust and easily interpretable method for determining, with uncertainty, model parameters from experimental data. Crucially, the use of Bayes’ theorem allows the practitioner to naturally incorporate information from both prior knowledge and experimental evidence, optimizing the precision of the inferred parameter. In the domain of laser Doppler velocimetry (LDV), a closely-related diagnostic technique to PDV, Bayesian methods have been used with success to infer the dynamic behavior of vibrating solid objects^{11,12} and fluid flows^{13–15}. These methods often assume a simple

model for the velocity behavior of the object under study, such as, in the case of acoustic oscillations, a time-varying sinusoid of constant phase and amplitude. This limits the number of model parameters and so greatly reduces the computational cost of numerically sampling the joint probability distribution. By contrast, PDV analysis poses unique challenges to the Bayesian approach because of the extreme physical conditions under study, including continuously time-varying supersonic velocity, the behavior of which is often unknown beforehand. Furthermore, PDV data are often highly-periodic and contain several thousand fringes in a single measurement, rendering a highly multi-modal probability distribution that is challenging to accurately sample using standard numerical approaches.

In this paper, we use recent advancements in numerical sampling techniques for Bayesian inference to explore the potential of a time-domain approach to PDV analysis as a complementary method to the standard STFT-based approach. Specifically, we directly model the time-dependent oscilloscope signal and thereby avoid using the spectrogram, and associated user choices, as an analysis tool. We use a general parameterized time-series model for the velocity history, from which we generate synthetic PDV oscilloscope traces that can be compared with the observed data. Used within a Bayesian framework, we can automatically recover velocity histories from the data with uncertainties that are determined naturally from an input noise model. Such an approach is susceptible to model misspecification and requires that we either capture possible errors in our model or that these are dealt with in data pre-processing. To avoid greatly increasing the number of required parameters in our model, we filter and normalize the data to remove any DC or low-frequency behavior (which provide no information about the velocity) and, subsequently, we construct a noise covariance model that captures the correlated behavior due to the filtering. We verify this approach by recovering an input synthetic velocity history and then validate it against standard PDV analysis of a shock front propagating through a quartz sample from data obtained with the STAR two-stage light gas gun at Sandia National Laboratories¹⁶.

We structure this paper as follows: in Section II we describe in detail our approach to Bayesian inference for PDV analysis, including verification against synthetic data, in Section III we present results from analysis of PDV data from the STAR two-stage light gas gun, and in Section IV we summarize our conclusions and recommendations.

II. BAYESIAN PDV ANALYSIS

A. Bayesian inference

We use Bayesian inference to recover the velocity histories of moving targets from PDV data. This approach is advantageous over other techniques because it allows us to obtain probability distributions for the model parameters and derivative variables, and to include information from prior knowledge and experiments. Furthermore, by forward-modeling the system we are able to reliably capture the effect of systematic errors and uncertainties in the experiment. The disadvantages of this approach are that the result is highly dependent on the model and so any misspecification of the data, either in the signal or noise, will be propagated through as an unknown

error. Knowledge of the limitations of the model are therefore important for interpreting any result obtained using this approach.

In most modeling problems, we are interested in inferring the posterior probability of the parameters θ that describe our model M , given our data D . Using Bayes' theorem, we can express this in terms of the product of the *likelihood function*, $P(D|\theta, M)$, and the *prior probability distribution*, $P(\theta|M)$,

$$P(\theta|D, M) = \frac{P(D|\theta, M)P(\theta|M)}{P(D|M)}, \quad (1)$$

where the denominator, $P(D|M)$, known as the marginal likelihood (or model evidence), is independent of θ and so can be absorbed into the normalization. In general, the posterior is not analytically tractable and so numerical methods, such as Monte Carlo sampling, are used to approximate the probability distribution as a function of the model parameters. From Eq. 1, we can see that choosing the correct likelihood function and priors are critical in accurately representing our belief in the model parameters. In the following sections we present our choices of these for PDV analysis.

B. Likelihood function

The likelihood function, $\mathcal{L}(\theta) = P(D|\theta, M)$, encodes information about the model parameters gained from the experimental data, updating our belief from the prior to the posterior probabilities. In the case of PDV analysis, the dataset is a voltage-time trace acquired with a digital oscilloscope. The ground-truth PDV signal is contaminated by random Gaussian noise due to components in the signal chain, which is dominated by photon noise from the reference laser (see Sections IV C and VII F of Dolan 2020⁷). We therefore use a multivariate Gaussian likelihood function, given by

$$\mathcal{L}(\theta) = \frac{1}{\sqrt{(2\pi)^N |\mathbf{C}|}} \exp \left[-\frac{1}{2} (\mathbf{d} - \boldsymbol{\mu}(\theta))^T \mathbf{C}^{-1} (\mathbf{d} - \boldsymbol{\mu}(\theta)) \right], \quad (2)$$

where \mathbf{C} is the matrix that encodes the covariance due to the noise, \mathbf{d} is the vector (of length N) of sampled voltage data as a function of time, and $\boldsymbol{\mu}$ are the expected values of these data given by the model. The components of \mathbf{C} are $\rho_{i,j}\sigma_i\sigma_j$, where σ_i and σ_j are the standard deviations due to noise at the i 'th and j 'th positions in the data, and $\rho_{i,j}$ are Pearson correlation coefficients. We describe how these are estimated and modeled in Section II C 3.

C. Models

1. Velocity history

We model the velocity history of a single component using a parameterized one-dimensional series based on piecewise linear components. Parameters consist of an initial velocity, v_0 , and n accelerations \dot{v}_i over each subsequent time interval $t_i - t_{i-1}$. By parameterizing the series in terms of acceleration, we do not need to adjust the prior distributions with respect to temporal resolution. The velocity at time $t \in (t_{j-1}, t_j)$ is then




FIG. 1. A synthetic PDV oscilloscope trace (voltage, V versus time, t) generated using the models discussed in Section II C and used for verification of our Bayesian method (see Section II E). The vertical red lines correspond to the start and end times, at 0.12 and 0.16 μs respectively, of the target signal. Although the ground-truth PDV signal is sinusoidal, the presence of noise renders visual identification of the fringes challenging (see inset); hence the need for techniques such as STFT analysis or Bayesian forward-modeling to accurately infer the velocity history from the data.

given by

$$v(t) = v_{j-1} + (v_j - v_{j-1}) \frac{t - t_{j-1}}{t_j - t_{j-1}}, \quad (3)$$

where

$$v_j = v_0 + \sum_{i=1}^j \dot{v}_i (t_i - t_{i-1}). \quad (4)$$

The number of acceleration parameters depends on the desired time resolution but are constrained by the computational complexity of sampling the parameter space for inference. The start and end times of the series could be included as inferred parameters, but we found that model misspecification of real experimental data leads to concentration of fitting in high signal-to-noise regions (see Section III). Therefore, these are fixed to reasonable values that enclose the PDV signal based on visual inspection of the spectrogram.

This model has the advantage of simplicity and interpretability, but forces linear behavior over time scales smaller than Δt_i and is not differentiable at component intersections. This behavior could be avoided by choosing alternative models that are constructed from linear combinations of a finite set of orthogonal basis functions, which include, for example, the Hermite Polynomials or Karhunen-Loève (KL) expansion of a Gaussian Process (GP) covariance function^{17,18}. Likewise, systems that are known to include discontinuities in the velocity history could be explicitly modeled by incorporating a finite number of parameterized step functions that multiply the time series. Such models are beyond the scope of this work but should be considered in future implementations.

2. PDV signal

We generate a synthetic PDV signal by modeling the interference fringe pattern between light from the target and reference lasers. The time-dependent beat frequency of this interference signal is given by

$$f(t) \equiv \frac{d\phi(t)}{dt} = |\nu_T - \nu_R + 2\nu_T\beta^*(t)|, \quad (5)$$

where ϕ is the phase angle, ν_T and ν_R are, respectively, the target and reference laser frequencies, and $\beta^*(t)$ is the ratio of the apparent velocity, $v^*(t)$, to the speed of light in vacuum. Uncertainties in the PDV laser frequencies could be included as a component of the model, but are expected to be at the level of one part in a million⁷ and so we do not include this in our analysis. For a full discussion of the definition of apparent velocity please see section VI B of Dolan (2020)⁷; here, for the simple reflective target surfaces considered in this work, the apparent velocity is related to the lab-frame velocity, $v(t)$, by

$$v^*(t) = n_0 v(t), \quad (6)$$

where n_0 is the refractive index ahead of the target.

For near infrared lasers ($\lambda \approx 1550$ nm) that are typically used in PDV, changes in beat frequency are related to changes in the apparent velocity by

$$\left[\frac{\Delta f}{1.3 \text{ GHz}} \right] \approx G(v^*) \left[\frac{\Delta v^*}{1 \text{ km s}^{-1}} \right], \quad (7)$$

where $G(v^*)$ can be ± 1 , given by

$$G(v^*) = 2H(\nu_T - \nu_R + 2\nu_T\beta^*) - 1. \quad (8)$$

H is the Heaviside step function, which has the property of being zero for arguments less than zero, and unity otherwise. Therefore, for fast-acquisition oscilloscope bandwidths $\gtrsim 10$ GHz, changes in apparent velocities $\gtrsim 8 \text{ km s}^{-1}$ can be measured. If $\nu_T > \nu_R$, then $G(v^*) = 1$ for all positive (that is, approaching) apparent velocities and the diagnostic is known as up-shifted PDV. Alternatively, for down-shifted PDV, the target and reference laser wavelengths are chosen such that $\nu_T < \nu_R$ and so $G(v^*)$ switches sign within the range of positive apparent velocities. This doubles the range of measurable velocities within the available oscilloscope bandwidth, but with ambiguity about the frequency-to-velocity mapping (see figure 4 of Dolan 2020⁷).

Once we have modeled the beat frequency as a function of time, the resulting PDV signal is given by

$$V(t) = \hat{V}(t) \sin\left(\phi_0 + \int_{t_0}^t f(t') dt'\right), \quad (9)$$

TABLE I. Model parameter values used to generate a synthetic PDV dataset for verification of the method.

Model	Parameter	Units	Value
Velocity history	v_0	km s^{-1}	25
	\vdots	km s^{-1} per μs	0
	\dot{v}_1	\vdots	0
	\dot{v}_2		-20
	\dot{v}_3		-50
	\dot{v}_4		-100
	\dot{v}_5		-100
	\dot{v}_6		-100
	\dot{v}_7		-50
	\dot{v}_8		-20
	\dot{v}_9		0
	\dot{v}_{10}		0
	\dot{v}_{11}		20
	\dot{v}_{12}		50
	\dot{v}_{13}		100
	\dot{v}_{14}		50
	\dot{v}_{15}		20
	\dot{v}_{16}		0
	\dot{v}_{17}		-10
	\dot{v}_{18}		-20
	\dot{v}_{19}		-50
Phase offset	ϕ_0	deg	0

where \hat{V} is the amplitude (dependent on the unknown time-varying intensities of the target and reference lasers), and ϕ_0 is the unknown phase at $t = t_0$. We cumulatively evaluate the integral numerically using the composite trapezoid method. We could model the amplitude as a time series, using the models described above, but this would not provide us with further information about the velocity history of the target. Alternatively, we can attempt to capture and factor out the amplitude during a pre-processing step, thereby being left with only the frequency-dependent component of the signal. However, this assumes that only a single coherent component exists in the PDV signal. We discuss this further in Section III. In the case of the unknown phase offset, we simply include ϕ_0 in the model as a nuisance parameter.

3. PDV noise

We model the noise in the PDV data as a multivariate Gaussian distribution with zero mean and covariance $C_{i,j} = \rho_{i,j} \sigma_i \sigma_j$, where σ_i and σ_j are the standard deviations at the i 'th and j 'th positions in the data, and $\rho_{i,j}$ is the Pearson correlation coefficient. We could include these parameters in the inferred model, but this would greatly increase the complexity of the inference problem. Instead we estimate the noise *a priori* using the following summary statistics based on prior assumptions:

Standard deviations: Under the assumption of homoskedastic (that is, constant variance) noise, we can estimate a single value for the standard deviation across the oscilloscope

trace. We identify a signal-free region in frequency space from a spectrogram of the data, then apply an appropriate noise-only bandpass filter to remove the PDV signal. In the case of the experimental validation data used in Section III, we used beat frequencies between 2 and 5 GHz as representative of signal-free data (see the spectrograms in Fig. 9). We then estimate the standard deviation due to noise using the median absolute deviation (MAD) across the filtered trace. In the presence of any residual signal or artifacts, MAD is a more reliable estimator of the variability due to noise than the standard deviation estimator. Finally, because the noise-only filter acts to reduce the true noise, we apply a correction factor to the standard deviation that is estimated using a Monte Carlo sample of 1000 filtered synthetic Gaussian noise traces.

Correlation coefficients: With regard to the correlation coefficients, we assume that all correlated noise is the result of bandpass filtering due to the oscilloscope bandwidth and subsequent data processing. We identify the frequency band of the oscilloscope trace from the spectrogram and then estimate the correlation coefficients from a Monte Carlo sample of 1000 bandpass filtered synthetic Gaussian-noise traces.

We demonstrate the validity of this model in Appendix A. Time series are shown for Gaussian noise generated with and without a covariance matrix determined using the method above for a 1 GHz low-pass filter. The resulting spectrograms demonstrate that the covariance matrix correctly models the noise correlation due to the filter.

D. Implementation

We implement all data ingestion, processing and modeling in the Python programming language. We use the Bilby Bayesian inference library^{19,20} with the Dynesty Nested Sampling (NS) package^{21,22} to recover an approximation of the posterior probability distribution for the model parameters. NS²³ is an alternative approach to the commonly used Markov Chain Monte Carlo (MCMC) methods that step a chain of samples through the parameter space; instead, a one-dimensional integral is calculated over slices of the prior volume that are constrained by nested iso-likelihood contours. Although originally developed with the goal of estimating the model evidence, NS also produces the sampled posterior distribution that can be used for parameter estimation. Given the number of parameters and inherently multi-modal nature of the PDV likelihood, we use the random walk and multi-modal bounding options for live point replacement sampling (LRPS), with 1500 live points used for robust posterior sampling (for a review of these methods please refer to Buchner 2023²⁴). This produces a chain of sample points that can be used to estimate the joint posterior distribution of the model parameters, and hence recover the PDV velocity history. The run time is dependent on the number of live points and the complexity of the posterior distribution; PDV analysis in the voltage-time space is a difficult inference problem and so the NS algorithm can take many hours to converge with more than 1000 live points.




FIG. 2. Spectrograms (beat frequency, f versus time, t) generated from (a) the synthetic PDV data shown in Fig. 1, (b) the inferred posterior distribution of model PDV signals, (c) the inferred maximum *a posteriori* (MAP; that is, most probable) model and (d) the corresponding residual (MAP model subtracted from the data). These visually demonstrate that the synthetic signal is recovered by the Bayesian forward-modeling method. In the case of the MAP spectrogram, we added random noise drawn from our PDV noise model for visual comparison with the data. The noiseless bands at $f < 2$ and $f > 23$ GHz are the result of bandpass filtering during pre-processing.

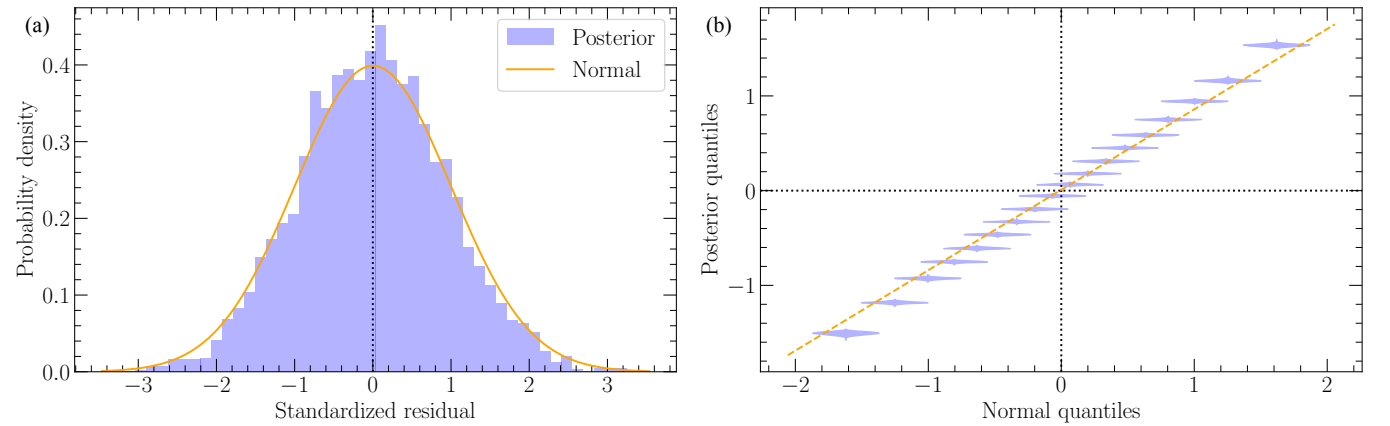


FIG. 3. (a) The posterior distribution of standardized residuals (PDV model subtracted from data) and (b) resulting quantile-quantile (QQ) plot with respect to the normal distribution. The synthetic noise is Gaussian and so residual deviation from a normal distribution is a measure of the error in recovering the synthetic PDV signal; this is quantified in the QQ plot by the degree of deviation from linear.




FIG. 4. The posterior velocity history (velocity, v versus time, t) inferred from the synthetic PDV data shown in Fig. 1. The blue solid line and envelope, respectively, denote the median and 95 % credible interval over the sample distribution of velocity curves. The orange dashed line denotes the injected model from which the synthetic data were generated. The "bow tie" shape is a result of using a piecewise linear model for the time series.

TABLE II. Priors used in this work for each model parameter.

Model	Parameter	Units	Prior distribution	Prior hyperparameters
Velocity history	v_0	km s^{-1}	Gaussian	mean = 25, std = 1
	\dot{v}_i	$\text{km s}^{-1} \text{ per } \mu\text{s}$	Gaussian	mean = 0, std = 100
Phase offset	ϕ_0	deg	Uniform	min = -180, max = 180

E. Verification against synthetic data

1. Synthetic data generation

To verify that our method can accurately recover velocity histories from PDV data, we tested it on a synthetic dataset that is similar to the data used for experimental validation in Section III. We generated a synthetic oscilloscope trace over an interval of $0.1 - 0.2 \mu\text{s}$, with a sample interval of 0.02 ns that is typical of multi-GHz oscilloscopes. We added a PDV signal with start and end times of 0.12 and $0.16 \mu\text{s}$, respectively, amplitude $\hat{V} = 1 \text{ V}$, phase offset $\phi_0 = 0 \text{ deg}$, and laser wavelengths $\lambda_T = 1549.219 \text{ nm}$ and $\lambda_R = 1548.945 \text{ nm}$. These wavelengths correspond to down-shifted PDV with a "bounce" at an apparent velocity of $v^* = 26.5 \text{ km s}^{-1}$.

The velocity history was modeled using the piecewise linear model described in Section II C 1, with $v_0 = 25 \text{ km s}^{-1}$ and twenty \dot{v}_i parameters (see Table I) that generate arbitrary, but typical, varying behavior of a propagating shock front in a solid diagnostic block. This was then converted to the apparent velocity using a refractive index of $n_0 = 1.528$, equal to that of the quartz sampled used in the experiment described in Section III. Finally, we added noise with unity standard deviation and applied a bandpass filter between 2 and 23 GHz , to simulate the AC coupling and bandwidth limit of the scope, respectively. The resulting synthetic oscilloscope trace is shown in Fig. 1 and the spectrogram in the top left hand panel of Fig. 2.

2. Model prior selection

It is evident from visual inspection of the synthetic oscilloscope trace that this is a challenging inference problem; we are attempting to infer a model from noisy data that describes the time-varying behavior of a periodic signal with several thousand cycles. The likelihood function is inherently multi-modal because frequency aliases, corresponding to higher velocities, will produce equally good representations of the data. Likewise, the signal periodicity will generate several local likelihood maxima that are likely to prove challenging for numerical samplers. Given these challenges, we adopted Gaussian prior distributions for our velocity model parameters that are weakly informative and therefore able to regularize the inference against finding aliased modes with physically-undesirable velocities. In particular, the prior mean and standard deviation for the initial velocity, v_0 , is chosen based on inspection of the signal in the spectrogram. This is further helped by including the phase offset parameter, which shifts the model to better match the true periodic PDV signal. We summarize the prior distributions for each model parameter in Table II.

3. Verification results

In Fig. 2 we compare spectrograms for the synthetic data with those obtained from the inferred model velocity history. In the top-right panel we show the full posterior spectrogram, which is generated by stacking spectrograms for each set of model parameters sampled from the posterior probability distribution. Visually, this is comparable to the signal seen in the data. We also show the maximum *a posteriori* (MAP) spectrogram, with noise added from our model, which again is visually consistent with the data. The corresponding MAP residual is shown and consistent with noise, again demonstrating that the synthetic signal is recovered by the method.

This is further demonstrated in Fig. 3, where we compare the posterior distribution of residuals in the PDV trace (difference between model and data) with the normal distribution. The synthetic noise is Gaussian and so any deviation from the normal distribution is a metric for errors in recovering the synthetic signal. Quantitative analysis is done using the quantile-quantile (QQ) plot²⁵, which graphically assesses whether the posterior is plausibly consistent with a normal distribution. The linearity of the points suggest that the residuals are indeed drawn from a normal distribution.

As a final assessment of the verification, in Fig. 4 we compare the inferred posterior velocity history with the injected model. They are consistent within the 95 % credible interval of the posterior, demonstrating that the injected model is recovered with high confidence.

III. EXPERIMENTAL VALIDATION

A. Experimental setup

Validation of our Bayesian approach to PDV analysis was achieved by comparing the inferred velocity histories from a real experimental dataset with those obtained using the standard STFT-based method. To this end, we used data from a recent experiment by Skidmore et al.²⁶ with the STAR two-stage light-gas gun (2SLGG) at Sandia National Laboratories¹⁶. Skidmore et al. used down-shifted PDV to measure the velocity of a shock driven into a 3.5 mm diameter quartz sample by a 17 mm diameter tantalum projectile launched at 6.5 km s^{-1} using the 2SLGG. This was carried out as part of a commissioning campaign for a novel hydrodynamic pressure amplifier developed at First Light Fusion (FLF) as a platform for low-cost Equation-of-state (EoS) measurements at terapascal pressures. The geometry of the pressure amplifier (to be presented in a forthcoming paper by Skidmore et al.²⁶) is such that the shock wave is split and recombined to form a Mack stem at the output. The resulting planar shock front is sufficient to uniformly compress the quartz sample, situated at the output,




FIG. 5. Schematic (not to scale) of the experimental set up used to provide a validation dataset for our Bayesian approach to PDV analysis. A 17 mm diameter tantalum projectile was launched at a velocity of 6.5 km s^{-1} using the STAR two-stage light-gas gun at Sandia National Laboratories¹⁶, which drove a pressure-amplified planar shock into a 3.5 mm diameter quartz sample at velocities $\approx 25 \text{ km s}^{-1}$. The velocity history of the shock wave was measured using a bespoke multi-point PDV probe containing an array of 24 fibers, each separated by $125 \mu\text{m}$, and arranged in the configuration shown in the right-hand insert. Four fibers (empty circles) were used to send the laser light to the quartz sample and 15 (numbered circles) to return the reflected light from the shock front. Fiber 16 (crossed out) was broken and so did not return any usable data. The colors denote fiber light of the same wavelength, which are given in Table III.

to pressures exceeding 1 TPa across a 1 mm diameter region (equivalent to velocities $\approx 25 \text{ km s}^{-1}$); similar schemes have been fielded previously on laser platforms²⁷.

To test the amplifier performance, a bespoke multi-point PDV probe of 24 optical fibers was placed in three rows extending $875 \mu\text{m}$ across the quartz sample as shown in Fig. 5. Four fibers were used to send laser light to the sample, and 15 to return the reflected signal (although one return fiber was broken and so did not provide usable data). The incident light used an NKT Koheras BASIK laser source at 40 mW connected to a 4 way light splitter. Each splitter output was connected to an independent attenuator to regulate the incident power so as to obtain a -30 dBm of return light in the PDV probe, measured using a power meter in the return fiber. Multiple target and reference wavelengths were used to minimize the cross-talk between adjacent channels; we summarize these in Table III.

Skidmore et al. carried out PDV analysis using the standard STFT approach, obtaining velocity histories that match the performance requirements of the amplifier. However, they noted that this method prevented velocity recovery within the first 1-2 ns (due to the finite time window). By forward-modeling the PDV trace directly, our method removes the need to choose an STFT window function; effectively extrapolating the velocity behavior, with uncertainty, to the chosen start-time. Skidmore et al. also noted that two probes (15 and 23) had unreliable velocity extraction due to spectral broadening resulting from the numerical aperture of the optical fiber. Our approach would be similarly affected by contamination from multiple velocities, but provides a reliably objective method of inferring the

signal-to-noise weighted velocity in the fiber.

B. Data preparation

The raw oscilloscope traces were recorded with 50 GS s^{-1} Tektronix digitizers with 25 GHz bandwidth, which we ingested into a Python dictionary. We removed DC and low-frequency signal components using a high-pass filter with a 2 GHz cutoff, and then estimated the noise covariance using the technique described in Section II C.

To avoid unnecessarily increasing the complexity of our inference problem, we estimated the time-varying behavior of the amplitude empirically from the data. This was achieved by applying a 3rd-order Savitsky-Golay filter²⁸, with a window size of 10 ns (several hundred PDV fringes for the expected velocity of 25 km s^{-1}), to the absolute values of the PDV data. The data and the noise covariance were then normalised by this estimated amplitude, ideally removing time-varying behavior due to changes in the laser intensities and/or fiber performance. This approach assumes that:

1. the signal-to-noise ratio is sufficiently high such that the noise does not contribute significantly to our estimate of the amplitude
2. temporal variations in amplitude are significantly slower than the fringe period ($\Delta t \gg 0.10 \text{ ns}$)
3. the PDV signal contains a single coherent component and not multiple sinusoids




FIG. 6. Velocity histories (velocity, v versus time, t) inferred from experimental PDV data with the STAR 2SLGG. The solid lines and envelopes, respectively, denote the median and 95 % credible interval from our Bayesian analysis. The dashed lines denote the analysis by Skidmore et al. using the standard STFT-approach, demonstrating similarity between the two methods. The single black error bar (top right) indicates the velocity uncertainty (0.1275 km s^{-1}) and time window (5 ns) of the STFT-based approach.

If these assumptions do not hold to desirable tolerances, then the time-varying amplitude will need to be explicitly included in the model, with a significant increase in the computational complexity of the inference. Spectrograms constructed from the resultant normalised oscilloscope traces are shown in the left panels of Fig. 9.

TABLE III. Summary of the experimental parameters used for each dataset in our analysis of the STAR 2SLGG shot. Fiber identification numbers (IDs) are as shown in Fig. 5. Wavelengths, λ_T and λ_R , are for the target and reference lasers, respectively. Fixed start and end times for the velocity history were chosen from visual inspection of the corresponding spectrograms (Fig. 9).

Fiber ID	λ_T [nm]	λ_R [nm]	Start time [μ s]	End time [μ s]
1	1549.193	1548.923	0.120	0.145
3	1549.193	1548.923	0.120	0.150
5	1554.125	1553.855	0.120	0.160
7	1554.125	1553.855	0.120	0.160
9	1549.193	1548.923	0.120	0.150
10	1549.193	1548.923	0.120	0.150
11	1550.654	1550.383	0.120	0.155
12	1550.654	1550.383	0.120	0.155
13	1554.125	1553.855	0.120	0.155
14	1554.125	1553.855	0.120	0.160
15	1549.218	1548.945	0.120	0.160
16	1549.218	1548.945	0.120	0.160
19	1550.654	1550.383	0.120	0.155
21	1550.654	1550.383	0.120	0.165
23	1549.218	1548.945	0.120	0.160

C. Validation results

We used the method discussed in Section II for analysis of the experimental PDV data, with the parameter priors summarized in Table II and the experimental parameters summarized in Table III. In early tests, we found that the sampling algorithm would infer start and end times for the PDV signal that are significantly different to those evident from the spectrogram. This is due to the finite temporal resolution of the parametrized model; a model that is a good fit to a subset of high signal-to-noise data is preferred over one that is a worse fit to the whole dataset. Therefore, in order to recover a model of the whole signal, we fixed the start and end times to those evident from visual inspection of the corresponding spectrograms.

In Fig. 6 we show the inferred velocity histories for each fiber optic probe using our Bayesian analysis and the STFT approach. Qualitatively, the approaches give similar results and our analysis supports the finding of Skidmore et al. that the amplifier achieves the required design specifications. Minor systematic velocity differences of $\sim 1\%$ for some fibers (for example, Fiber 21) are likely due to details of the extraction. Larger differences in velocity at specific times (for example, Fiber 15 between 0.125 and 0.135 μ s) are the result of spectral broadening or laser speckle. The spectrograms generated from the experimental data and those inferred using the model are shown in Fig. 9, from which we can see the advantage of the Bayesian approach in being able to infer signal from noisy data

(for example, Fiber 3) and interpolate across signal gaps (for example, Fibers 5 and 7). Likewise, the Bayesian approach has allowed us to infer velocity histories at earlier times than was available with the STFT-approach.

IV. CONCLUDING REMARKS

We have presented a methodology for carrying out Bayesian inference of velocity histories directly from PDV oscilloscope data as an adjunct to standard STFT-based analysis. We verified the method using synthetic PDV data generated from a template based on an experiment carried out by Skidmore et al.²⁶ with the STAR 2SLGG at Sandia National Laboratories¹⁶. We were able to recover the injected velocity history, within the uncertainties, and the posterior predictive distribution is consistent with the synthetic data. We then applied it to the real experimental data and compared with velocity histories obtained using the standard STFT-based analysis. We find that in general the two approaches agree, although we caution that the prior probability distributions for model parameters need to be carefully chosen to regularize the inference against, for example, aliasing of higher frequency harmonics. The Bayesian approach is found to be particularly advantageous when choices need to be made about interpolating across regions of missing data, low signal-to-noise or laser speckle.

The inferred velocity history is dependent on the model and priors chosen by the user (here we use a piecewise linear model) and, importantly, the resolution. We therefore caution that this does not capture uncertainty due to model misspecification, for example, of any rapidly time-varying or discontinuous behavior. Likewise, multiple velocities can only be captured by either including them explicitly in the model or, possibly, by using multi-modal sampling techniques. Here, we have assumed a single velocity component so that the PDV amplitude can be captured and normalised in a pre-processing step; multiple velocities would require the amplitudes to be explicitly included in the model. These issues can be identified post-hoc using posterior predictive checks (see e.g. Fig. 3) and overcome by including a sufficiently flexible model with the required parameters to capture the desired behavior. However, increased model complexity has associated computational costs; even with the model used here the numerical sampling can take several hours to converge. We therefore recommend that it be used as a complementary method to the standard STFT-based approach.

ACKNOWLEDGMENTS

The authors thank the anonymous reviewers for their constructive comments and suggestions, which helped improve the quality of this manuscript. The authors would also like to thank the staff at FLF for supporting the work in this paper, which included provision of the HPC for parallel sampling in our analysis. J. R. Allison thanks M. Gorman for useful discussions whilst preparing the paper. Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC (NTESS), a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's

National Nuclear Security Administration (DOE/NNSA) under contract DE-NA0003525.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available from the corresponding author upon reasonable request.

Appendix A: Demonstration of the noise model

The noise model is a critical component of any inference method. In Section II C 3 we described our PDV noise model, which is based on a multivariate Gaussian distribution with mean zero and a covariance matrix estimated from Monte-Carlo simulations of bandpass-filtered white noise. Here we demonstrate the validity of this model; in Fig. 7 we show an example of correlated Gaussian noise, sampled at 50 GHz, that has been generated using this model for a 1 GHz low-pass filter. The correlations between data at high frequency are visually evident when compared with Gaussian noise that has been generated using a model with no covariance. This is further demonstrated by the spectrograms shown in Fig. 8, which show the successful removal of noise frequencies greater than 1 GHz.

Appendix B: Experimental validation data

Experimental validation data were obtained by Skidmore et al. from a shot with the STAR 2SLGG facility at Sandia National Laboratories²⁶. A planar shock front was driven into a quartz sample at about 25 km s^{-1} and PDV data were recorded using fourteen fiber-optic probes in a linear array across a 1 mm region. The oscilloscope traces were bandpass filtered and the signal amplitude estimated using the method described in Section III B. Spectrograms constructed from the resultant normalized oscilloscope traces are shown in the left panels in Fig. 9. The inferred posterior spectrograms are shown in the right panels of the same figure, highlighting the ability of the Bayesian method to recover the signal from noisy data.

- ¹G. E. Duvall and R. A. Graham, *Rev. Mod. Phys.* **49**, 523 (1977).
- ²L. M. Barker, *Experimental Mechanics* **12**, 209 (1972).
- ³W. F. Hemsing, *Rev. Sci. Instrum.* **50**, 73 (1979).
- ⁴O. T. Strand, L. Berzins, D. Goosman, W. Kuhlow, P. Sargis, and T. Whitworth, *Proceedings of the 26th International Congress on High-Speed Photography and Photonics (SPIE)* **5580**, 593 (2004).
- ⁵O. T. Strand, D. R. Goosman, C. Martinez, T. L. Whitworth, and W. W. Kuhlow, *Review of Scientific Instruments* **77**, 083108 (2006).
- ⁶D. H. Dolan, *Meas. Sci. Technol.* **35**, 061001 (2024).
- ⁷D. H. Dolan, *Review of Scientific Instruments* **91**, 051501 (2020).
- ⁸D. H. Dolan, *Rev. Sci. Instrum.* **81**, 053905 (2010).
- ⁹We note that a technique based on polarization signal diversity technique has recently been developed for laser Doppler velocimetry^{29,30} which could be used in PDV to significantly reduce dropouts due to laser speckle noise. However, the relative cost of PDV (including GHz-bandwidth oscilloscopes and multiple channels) and extreme experimental conditions present significant barriers to use.
- ¹⁰W. P. Ambrose, "Precision and accuracy in PDV and VISAR," *Tech. Rep.* (Lawrence Livermore National Lab.(LLNL), Livermore, CA (United States), 2017).
- ¹¹P. M. Goggans and C. R. Smith, in *Detection and Remediation Technologies for Mines and Minelike Targets IV*, Vol. 3710, edited by A. C. Dubey, J. F. Harvey, J. T. Broach, and R. E. Dugan, International Society for Optics and Photonics (SPIE, 1999) pp. 1429 – 1437.
- ¹²W. F. Buell, B. A. Shadwick, and R. W. Farley, in *Laser Radar Technology and Applications V*, Vol. 4035, edited by G. W. Kamerman, U. N. Singh, C. Werner, and V. V. Molebny, International Society for Optics and Photonics (SPIE, 2000) pp. 444 – 455.
- ¹³A. Fischer, T. Pfister, and J. Czarske, *Measurement* **43**, 1556 (2010).
- ¹⁴O. Penttinen and H. Nilsson, *Proceedings from International Symposium of Fluid Flow Measurements* **2015** (2015).
- ¹⁵R. Roncen, F. Méry, E. Piot, and F. Simon, *AIAA Journal* **57**, 1055 (2019).
- ¹⁶L. C. Chhabildas, *AIP Conference Proceedings* **78**, 621 (1982).
- ¹⁷O. P. L. Maitre and O. M. Knio, *Spectral Methods for Uncertainty Quantification* (Springer, 2010).
- ¹⁸C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning* (The MIT Press, 2006).
- ¹⁹G. Ashton et al., *Astrophys. J. Suppl.* **241**, 27 (2019).
- ²⁰Documentation for Bilby can be found at <https://bilby-dev.github.io/bilby/>.
- ²¹J. S. Speagle, *MNRAS* **493**, 3132 (2020).
- ²²Documentation for Dynesty can be found at <https://dynesty.readthedocs.io>.
- ²³J. Skilling, *Bayesian Analysis* **1**, 833 (2006).
- ²⁴J. Buchner, *Statistics Surveys* **17**, 169 (2023).
- ²⁵M. B. Wilk and R. Gnanadesikan, *Biometrika* **55**, 1 (1968).
- ²⁶J. W. Skidmore et al., *J. Appl. Phys.* (submitted).
- ²⁷D. C. Swift and C. R. Ruiz, *AIP Conference Proceedings* **845**, 1297 (2006).
- ²⁸A. Savitsky and M. J. E. Golay, *Analytical Chemistry* **36**, 1627 (1964).
- ²⁹F. Wang, J. Wölck, M. Hess, and C. Rembe, *tm - Technisches Messen* **91**, 44 (2024).
- ³⁰F. Wang, S. Krause, and C. Rembe, *Measurement: Sensors* **22**, 100377 (2022).




FIG. 7. Correlated Gaussian noise (orange; voltage, V versus time, t) generated using a covariance matrix that models the effect of a 1 GHz low-pass filter using the method discussed in Section II C 3. Uncorrelated Gaussian noise (grey) is shown for comparison.

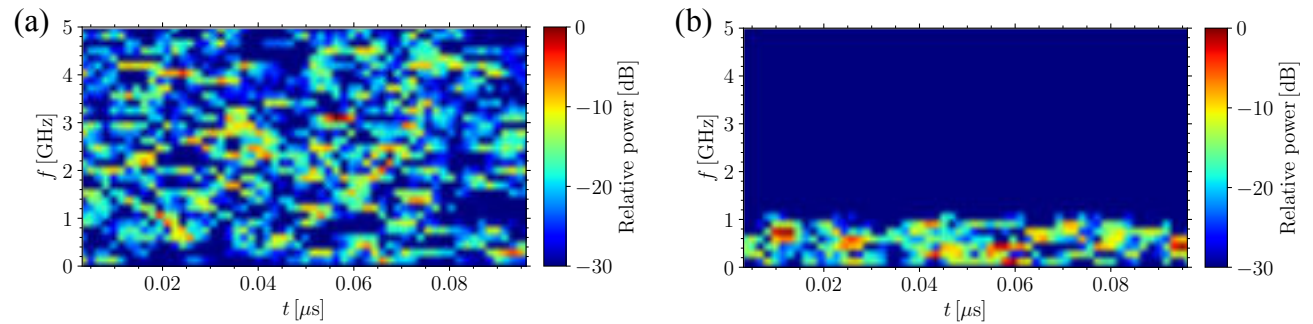


FIG. 8. Spectrograms (beat frequency, f versus time, t) generated from (a) the uncorrelated and (b) correlated noise-only traces shown in Fig. 7, showing the expected removal of frequencies above 1 GHz.

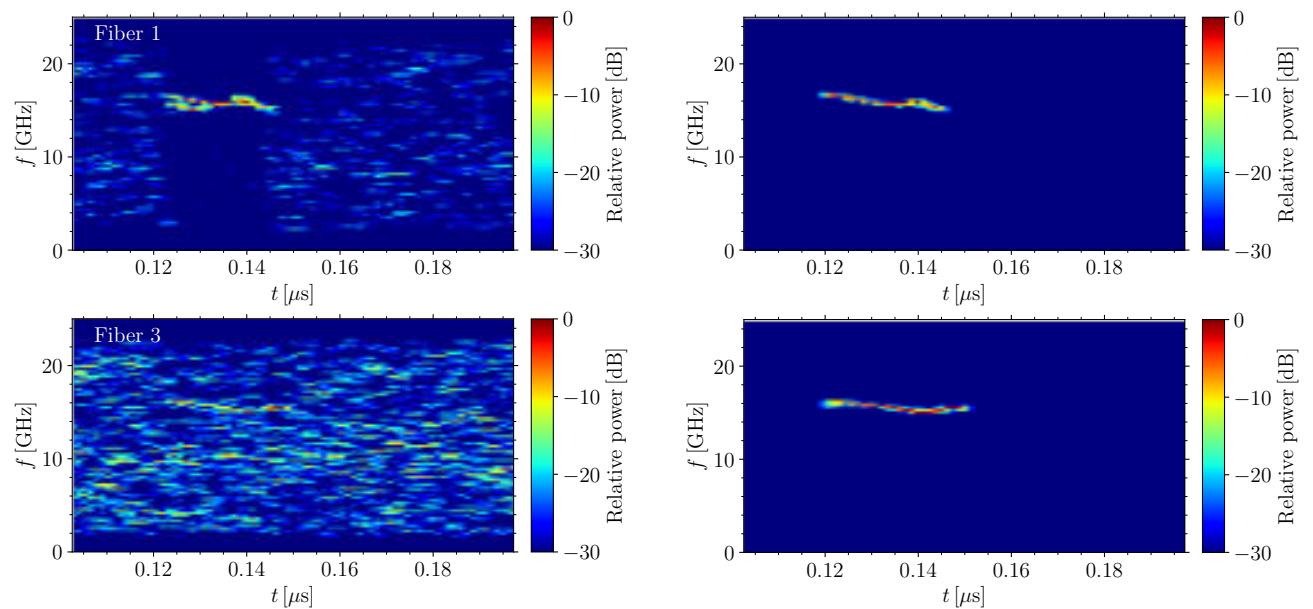


FIG. 9. Spectrograms (beat frequency, f versus time, t) generated from the experimental validation data discussed in Section III (left panels) and the inferred posterior distribution of model PDV signals (right panels).

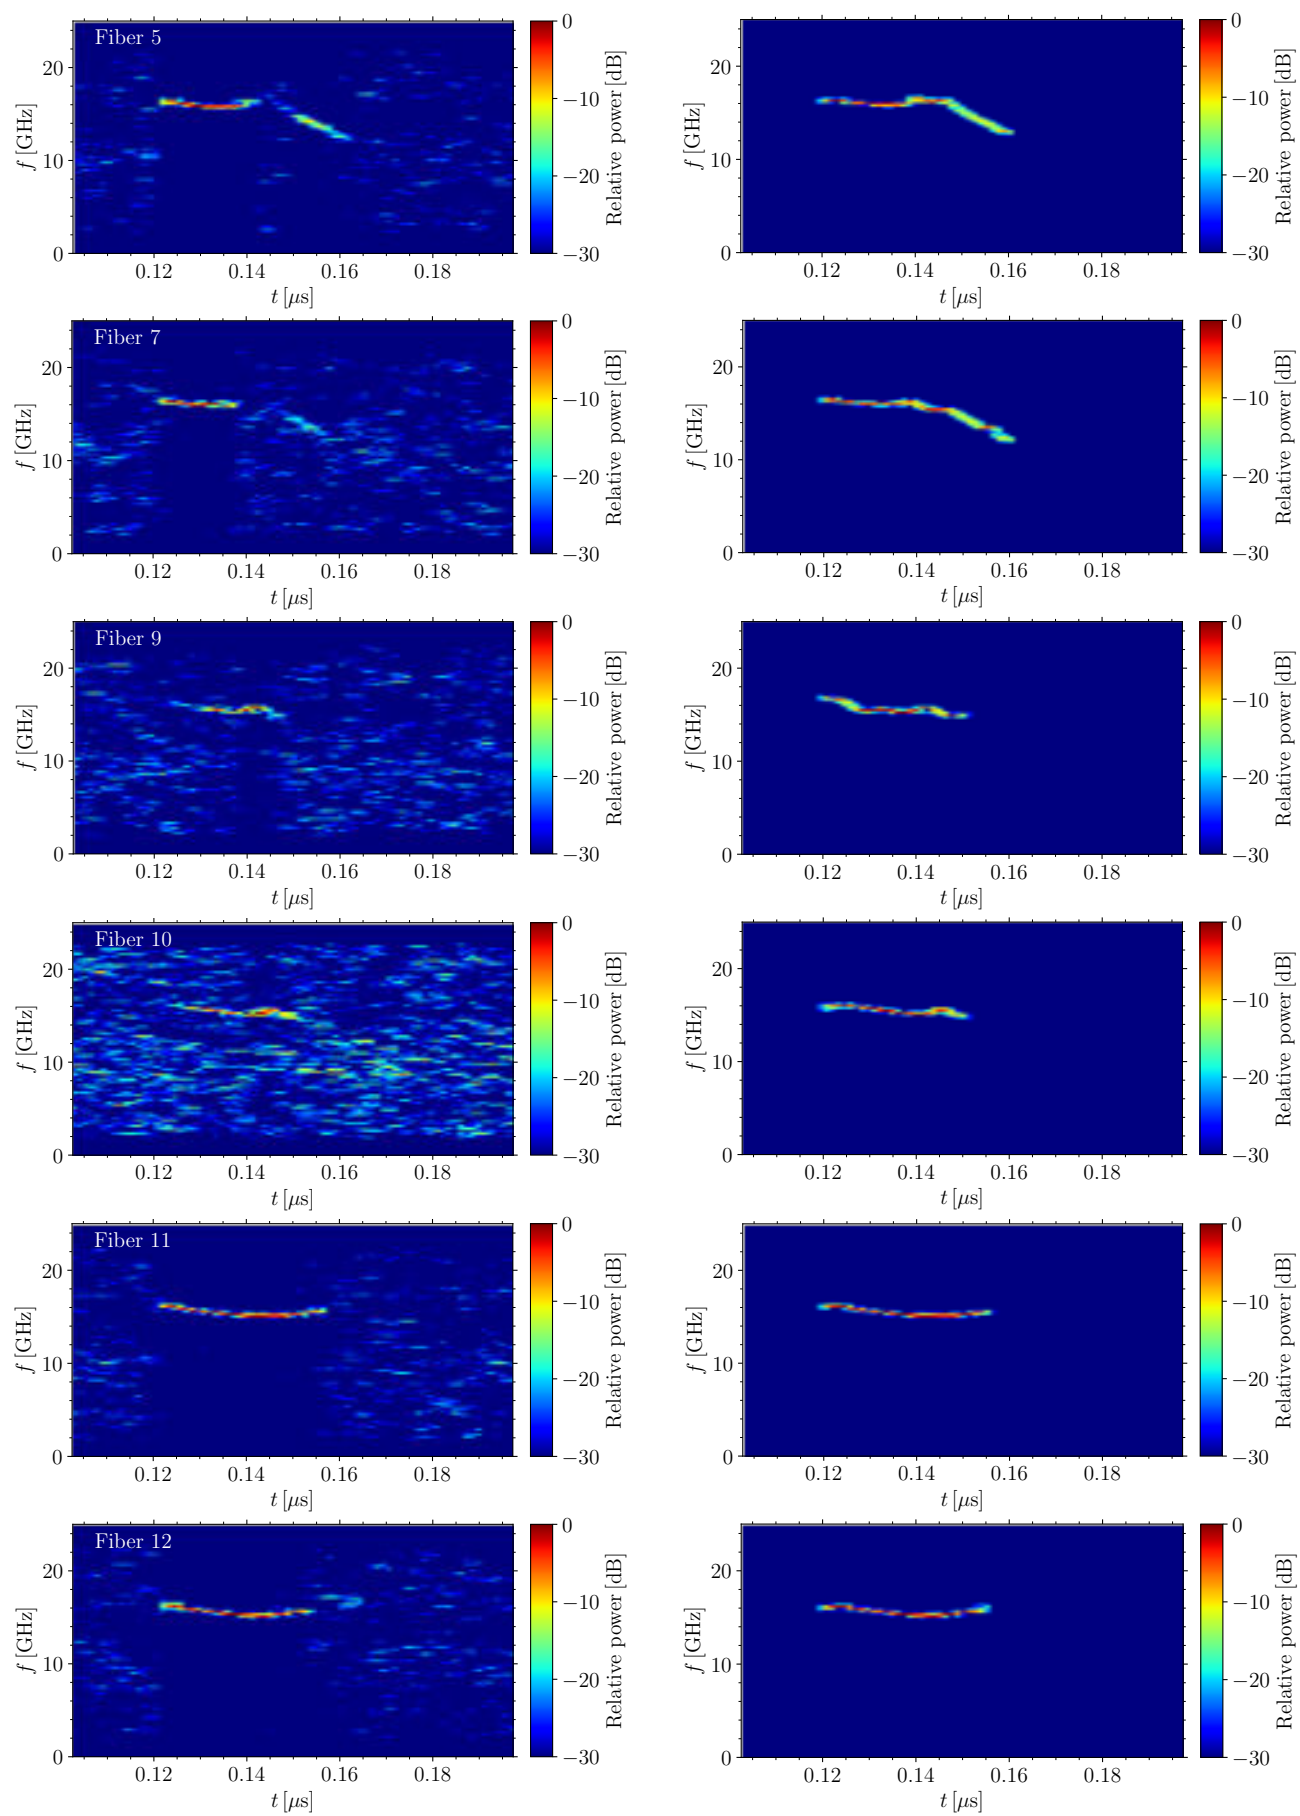


FIG. 9. (Continued) Spectrograms (beat frequency, f versus time, t) generated from the experimental validation data discussed in Section III (left panels) and the inferred posterior distribution of model PDV signals (right panels).

This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.
 PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0267409

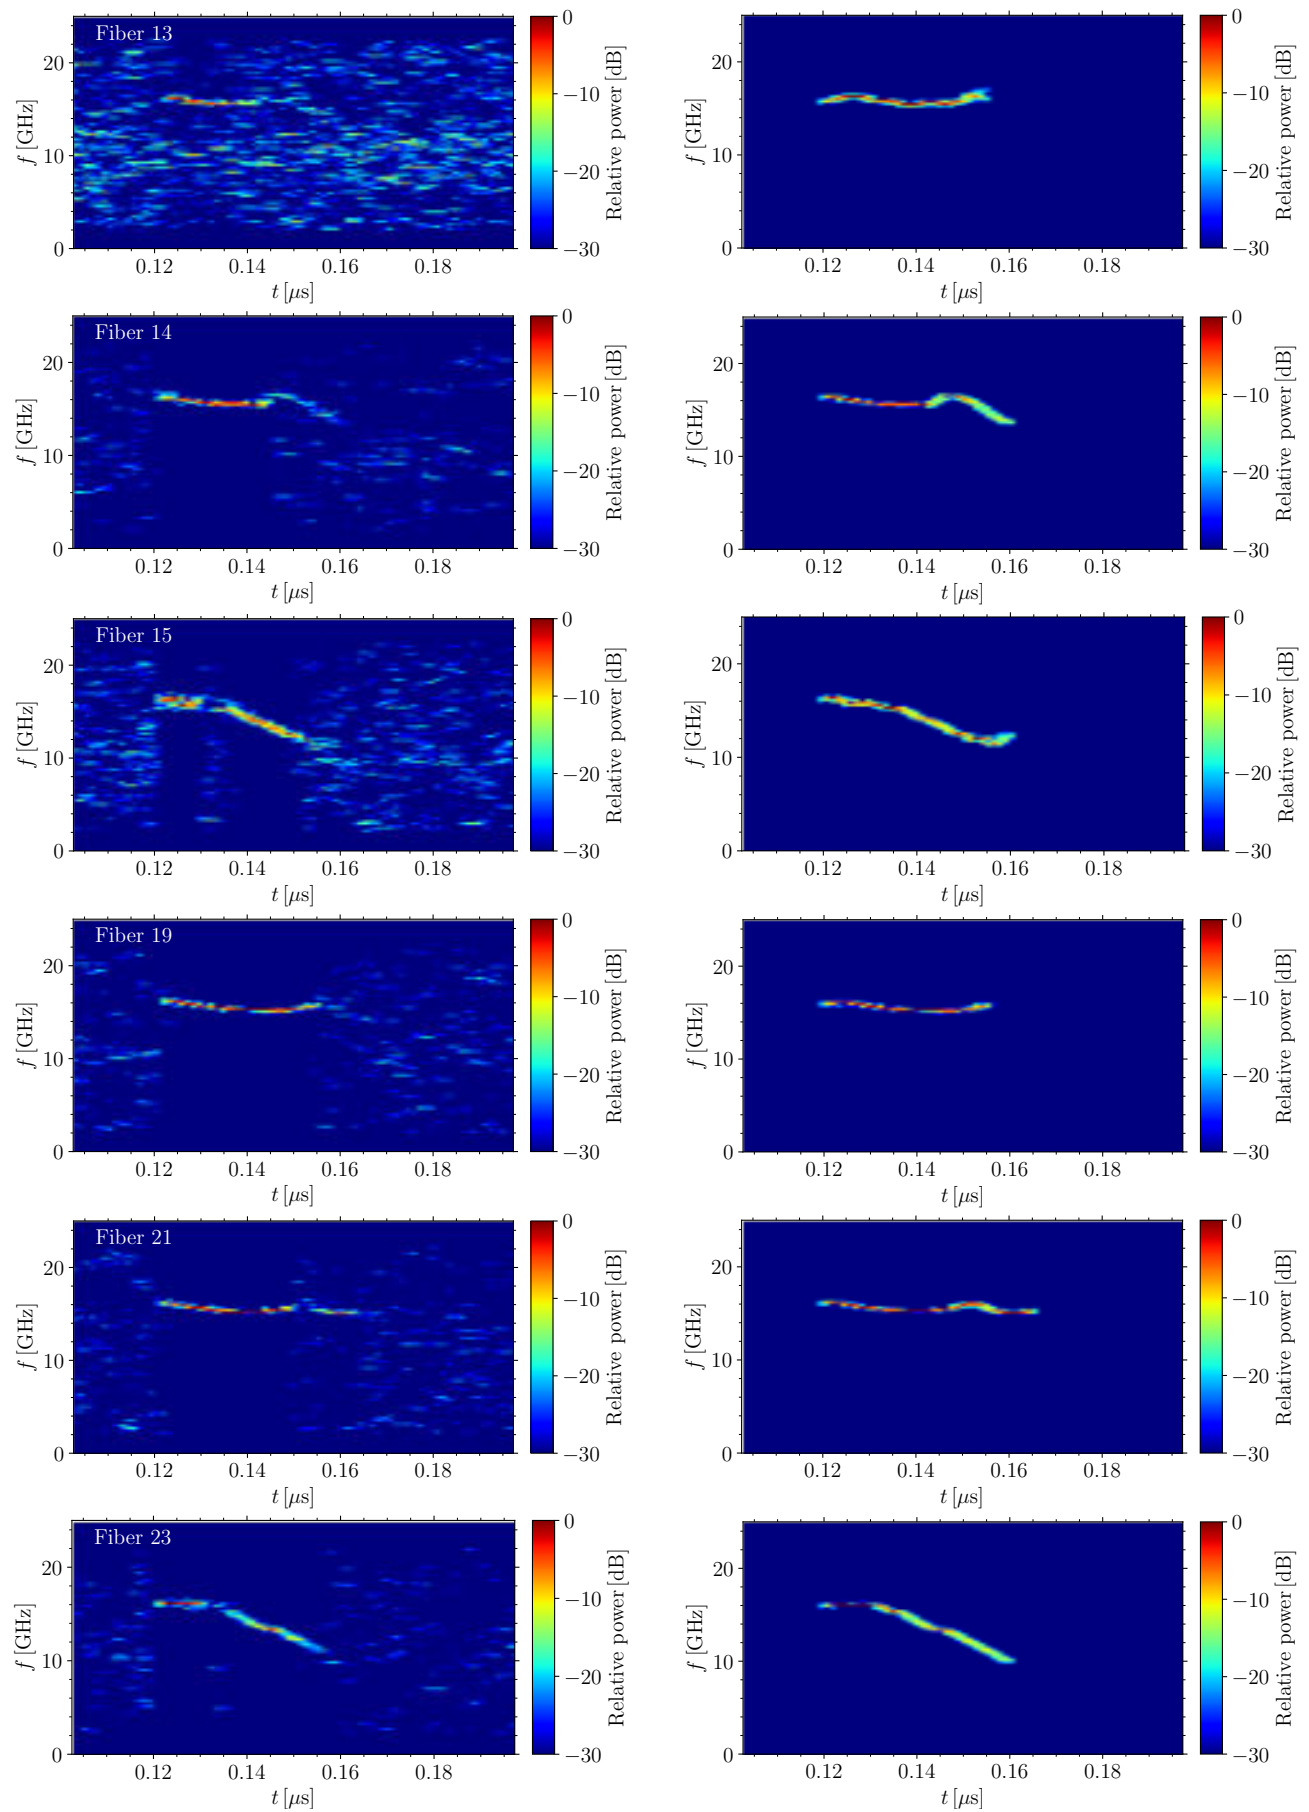


FIG. 9. (Continued) Spectrograms (beat frequency, f versus time, t) generated from the experimental validation data discussed in Section III (left panels) and the inferred posterior distribution of model PDV signals (right panels).