



# A Bayesian perspective on single-shot laser characterization

J. Esslinger<sup>a,1</sup> , N. Weiße<sup>a,1</sup>, C. Eberle<sup>a</sup>, J. Schröder<sup>a</sup> , S. Howard<sup>a,b</sup>, P. Norreys<sup>b</sup> , S. Karsch<sup>a</sup>, and A. Döpp<sup>a,b,2</sup>

Affiliations are included on p. 9.

Edited by Victor Malka, Weizmann Institute of Science, Rehovot, Israel; received April 29, 2025; accepted September 2, 2025 by Editorial Board Member Hui Cao

We introduce a Bayesian framework for measuring spatiotemporal couplings in ultraintense lasers that reconceptualizes what constitutes a “single-shot” measurement. Moving beyond traditional distinctions between single- and multishot devices, our approach provides rigorous criteria for determining when measurements can truly resolve individual laser shots rather than statistical averages. By contextualizing single measurements, this framework shows that single-shot capability is not an intrinsic device property but emerges from the relationship between measurement precision and predictability. Implementing this approach with a custom measurement device at the ATLAS-3000 petawatt laser, we provide quantitative uncertainty bounds on pulse front tilt and curvature. Notably, we observe that our Bayesian method reduces uncertainty by up to 60% compared to traditional approaches. Through this analysis, we reveal how the interplay between measurement precision and intrinsic system variability defines achievable resolution-insights that have direct implications for applications where precise control of laser–matter interaction is critical.

ultraintense lasers | laser metrology | Bayesian inference | spatiotemporal couplings

The characterization of complex physical systems presents a fundamental challenge in metrology: how to capture high-dimensional information using inherently limited measurement devices. This challenge is exemplified in ultraintense lasers, where petawatt-scale powers and focused intensities exceeding  $10^{23}$  W/cm<sup>2</sup> push the boundaries of measurement capabilities (1, 2). The field of a laser pulse in vacuum is characterized by its intensity and phase along the coordinates perpendicular to its propagation axis and measured along time or frequency. Many diagnostics focus on measuring specific projections or cross-sections, such as a nearfield profile and spectrum for intensity, or a wavefront and chirp for phase. However, such measurements are oblivious to correlations between dimensions. The simplest example of such spatiotemporal couplings (STCs) is the laser pulse front tilt (PFT), which compares the orientation of the pulse to its propagation direction. It is equal to angular dispersion in the spectral domain and hence only apparent when a spectrally resolved wavefront is measured (3). Traditional approaches using sequential scanning techniques (4–13) are strictly only suitable when laser systems exhibit neither drifts nor significant shot-to-shot fluctuations. When these conditions are not met, “few-shot” methods like FALCON (14) and IMPALA (15) trade resolution for acquisition time. For true instantaneous characterization, the field has pursued “single-shot” measurements, though this ideal often proves elusive and comes with significant trade-offs (16–20). Crucially, all current methods lack uncertainty quantification (21)—particularly important as the measurement precision is typically near the observable level of intrinsic fluctuations.

These challenges prompt us to reconsider laser characterization through the lens of Bayesian inference. For an isolated pulse absent any prior knowledge of its properties, capturing all information in a single shot offers clear advantages. However, many real-world scenarios involve contextual knowledge about the laser, be it from system specifications, previous measurements, or characteristic drift patterns. This is why, in our framework, each measurement represents an update to our understanding of the system, building upon prior information rather than starting from scratch.

Consider FALCON or other techniques (3, 22, 23), that measure wavefronts sequentially at different wavelengths: Is it “single-shot” because it takes an image every shot, or “multishot” because it needs multiple measurements for full spectral coverage? After how many filters is a measurement “complete”? Such questions have no clear answer in traditional frameworks. Rather than attempting to capture all information in one shot, we recognize that laser systems typically exhibit dynamics at separated time scales: slow drifts over minutes to hours and rapid fluctuations above the laser’s

## Significance

Ultraintense lasers are a key technology behind multiple transformative technologies such as laser-driven particle acceleration or inertial confinement fusion. To date, characterization of these lasers relies mostly on multishot scanning approaches or isolated single-shot techniques.

We introduce a probabilistic framework that treats individual measurements in the context of experimental setup and predictive models derived from other measurements. This allows us to quantify key laser parameters, their trends, as well as their intrinsic fluctuations while also reducing measurement uncertainty. This approach shows the intricate connection between measurements and data-driven modeling, which paves the way to a new generation of diagnostics for lasers and other metrology applications.

Author contributions: J.E., C.E., J.S., S.H., P.N., S.K., and A.D. designed research; J.E., N.W., and A.D. performed research; J.E. and A.D. analyzed data; and J.E. and A.D. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission. V.M. is a guest editor invited by the Editorial Board.

Copyright © 2025 the Author(s). Published by PNAS. This article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

<sup>1</sup>J.E. and N.W. contributed equally to this work.

<sup>2</sup>To whom correspondence may be addressed. Email: a.doepp@lmu.de.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2510645122/-/DCSupplemental>.

Published October 24, 2025.

repetition rate. While slow evolution can be predicted using time series models (24), high-frequency components manifest as stochastic shot-to-shot variations requiring instantaneous measurement. This distinction provides rigorous criteria for determining when measurements can truly resolve individual laser states rather than statistical averages. Understanding these shot-to-shot variations is crucial, especially for applications that make use of tailored STCs, ranging from particle acceleration (25–27) to attosecond pulse generation (28, 29).

In this paper, we present a rigorous mathematical framework demonstrating how “single-shot” capability emerges from the relationship between measurement precision and intrinsic system variability. Section 1 provides analytical tools for studying measurement system behavior and parameter estimation, while Section 2 showcases a demonstration for spatiotemporal coupling measurements in ultraintense lasers using a simple, custom-made single-shot diagnostic that is a derivative of FALCON, including a full probabilistic treatment.

## 1. Theory

As mentioned earlier, we describe the evolution of any measured parameter as a combination of deterministic and stochastic components. More formally, the state of any parameter  $x$  at times  $t_k$  can be modeled as

$$x_k = f(t_k) + \varepsilon_k \quad [1]$$

as shown in Fig. 1. There  $f(t_k)$  represents predictable changes in the system. In later sections in this paper, we will approximate this function by its instantaneous state plus the first-order linear instantaneous trend.  $\varepsilon_k \sim \mathcal{N}(0, \sigma_{\text{stoch}}^2)$  are independent, normally distributed random variables characterizing shot-to-shot fluctuations. The normal distribution naturally arises here as the aggregate effect of many small independent fluctuations tends to be normally distributed due to the Central Limit Theorem. The variance parameter  $\sigma_{\text{stoch}}^2$  characterizes the intrinsic stochasticity of the system, that is, fluctuations that cannot be predicted ahead.

Each measurement  $y_k$  comes with its own uncertainty,

$$y_k = x_k + e_k = f(t_k) + \varepsilon_k + e_k, \quad [2]$$

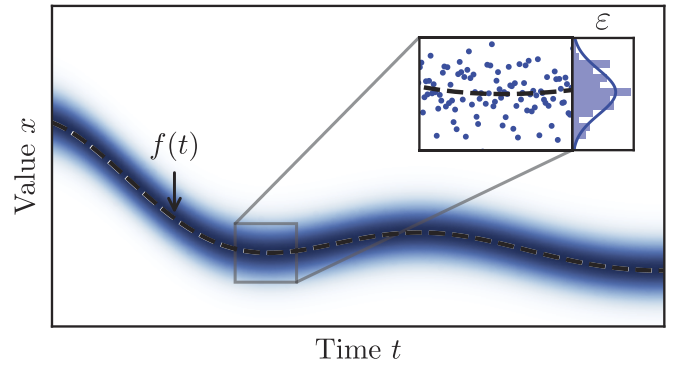
where  $e_k \sim \mathcal{N}(0, \sigma_{\text{meas}}^2)$  is the measurement noise. This measurement uncertainty, combined with the intrinsic stochasticity, determines our ability to resolve changes in the system state. Understanding this interplay requires a Bayesian perspective, where each measurement updates our knowledge of both the deterministic and stochastic components.

**1.1. Bayesian State Updates.** This section will go through basic concepts of Bayesian inference (30) required for our framework. The subsequent sections present our original work.

At the heart of Bayesian inference lies Bayes’ theorem, which provides a formal method for updating our beliefs about a state variable  $x_k$  given new measurement data  $y_k$ :

$$p(x_k|y_k) = \frac{p(y_k|x_k)p(x_k)}{p(y_k)} \quad [3]$$

Here,  $p(x_k|y_k)$  is the posterior probability (our updated belief),  $p(y_k|x_k)$  is the likelihood (the probability of observing the measurement given the state),  $p(x_k)$  is the prior probability (our initial belief), and  $p(y_k)$  is the evidence (a normalization factor). We make the recursive Markov assumption for the prior, i.e. that



**Fig. 1.** Decomposition of measurements into deterministic and stochastic components. The dashed black line shows the predictable trend  $f(t)$ , while the blue density represents the measurement spread due to intrinsic stochasticity. The inset shows a zoomed region with individual measurements (blue dots) and their statistical distribution (histogram), demonstrating how the stochastic component  $\varepsilon_k \sim \mathcal{N}(0, \sigma_{\text{stoch}}^2)$  manifests in the measurements.

each new state depends only on the immediately previous state, not the entire history of states. The evidence can be expressed as an integral over all possible values of the state

$$p(y_k) = \int p(y_k|x_k)p(x_k)dx_k. \quad [4]$$

In practice, solving this integral can be challenging. But given a Gaussian model for both the stochastic component and measurement uncertainty, we can solve Eq. 3 analytically due to the self-conjugacy of Gaussian distributions. For Gaussian distributions, with prior  $p(x_k) = \mathcal{N}(\mu_{k|k-1}, \sigma_{k|k-1}^2)$  and likelihood  $p(y_k|x_k) = \mathcal{N}(x_k, \sigma_{\text{meas}}^2)$ , the posterior is proportional to:

$$p(x_k|y_k) \propto \exp\left(-\frac{(x_k - \mu_{k|k-1})^2}{2\sigma_{k|k-1}^2} - \frac{(x_k - y_k)^2}{2\sigma_{\text{meas}}^2}\right)$$

Here,  $\mu_{k|k-1}$  and  $\sigma_{k|k-1}^2$  represent our prediction (prior) estimate and uncertainty before the  $k$ -th measurement, while  $y_k$  is the observed measurement, and  $\sigma_{\text{meas}}^2$  represents the measurement uncertainty. Completing the square in the exponent, such that it takes the form of a singular Gaussian, shows that the mean  $\mu_{k|k}$  and variance  $\sigma_{k|k}^2$  of this posterior Gaussian distribution—hence after the measurement—are given by

$$\mu_{k|k} = (1 - \gamma_k) \cdot \mu_{k|k-1} + \gamma_k \cdot y_k \quad [5]$$

$$\sigma_{k|k}^2 = \left(\frac{1}{\sigma_{k|k-1}^2} + \frac{1}{\sigma_{\text{meas}}^2}\right)^{-1} = \gamma_k \sigma_{\text{meas}}^2 \quad [6]$$

Here, we introduced the update weight  $\gamma_k$

$$\gamma_k = \frac{\sigma_{k|k-1}^2}{\sigma_{k|k-1}^2 + \sigma_{\text{meas}}^2}. \quad [7]$$

Note that the prior mean  $\mu_{k|k-1}$  and variance  $\sigma_{k|k-1}^2$  are based on our prediction of what to expect; see Section 1.6. The lower limit on the prior variance is given by the unpredictable intrinsic stochasticity  $\sigma_{\text{stoch}}$  of the system.

**1.2. Asymptotic Behavior and Measurement Regimes.** The update weight defined in Eq. 7 encapsulates several key insights. First, it quantifies how our uncertainty evolves with new measurements, accounting for both our prior knowledge and the system's inherent variability. Second, it reveals the interplay between measurement precision and intrinsic stochasticity in determining our final uncertainty. The behavior of this equation in limit cases reveals its consistency with established principles. As our prior variance  $\sigma_{k|k-1}^2 \rightarrow \infty$ , representing complete initial ignorance, we recover the classical result for the posterior variance:  $\sigma_{k|k}^2 \approx \sigma_{\text{meas}}^2$ . When  $\sigma_{\text{stoch}}^2 \rightarrow 0$ , indicating a deterministic system,  $n$  repeated measurements can reduce uncertainty indefinitely following  $\sigma_{\text{posterior}}^2 = \sigma_{\text{meas}}^2/n$ . This can be deduced by using the posterior of the previous measurement  $\sigma_{k-1|k-1}^2$  as the new prior  $\sigma_{k|k-1}^2$  in Eq. 6, starting with the  $\sigma_{\text{meas}}^2$  and repeating it  $n$  times.

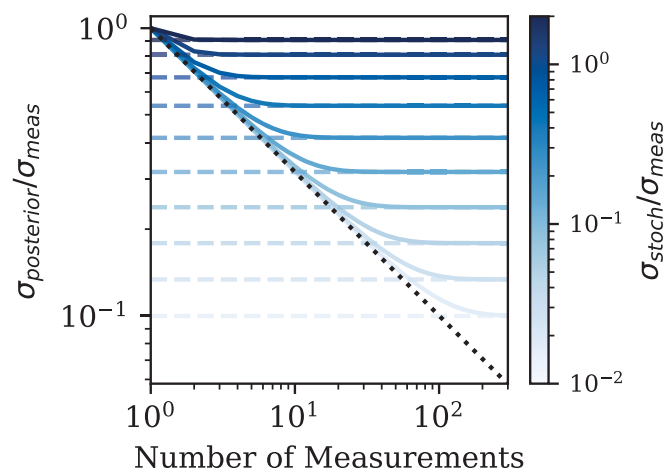
Between these cases, for  $\sigma_{\text{stoch}}^2 > 0$  and a sequence of “single-shot” measurements, the asymptotic behavior of our system reveals a stationary solution for which the posterior variance ceases to reduce further:

$$\begin{aligned} \sigma_{\infty|\infty}^2 &= \frac{\sqrt{1 + 4(\sigma_{\text{meas}}^2/\sigma_{\text{stoch}}^2)} - 1}{2} \cdot \sigma_{\text{stoch}}^2 \\ &= \frac{1 - \gamma}{\gamma} \cdot \sigma_{\text{stoch}}^2 \end{aligned} \quad [8]$$

This is shown in Fig. 2.

This asymptotic limit delineates a critical threshold when  $\sigma_{\text{meas}}^2 = 2\sigma_{\text{stoch}}^2$ , distinguishing two regimes of measurement capabilities:

**Distribution-Centric Regime** ( $\sigma_{\text{meas}}^2 > 2\sigma_{\text{stoch}}^2$ ): In this regime, the prediction relies more on the prior and the posterior variance remains always greater than the inherent system variability. For  $\sigma_{\text{meas}}^2 \gg \sigma_{\text{stoch}}^2$  the limit is approximately  $\sigma_{\infty|\infty}^2 \approx \sigma_{\text{stoch}} \cdot \sigma_{\text{meas}} = \sigma_{\text{stoch}}^2 \cdot (\sigma_{\text{meas}}/\sigma_{\text{stoch}})$ . In this regime of precision, there is no fundamental advantage of a single-shot device except for shorter measurement time, because a



**Fig. 2.** Evolution of the normalized posterior SD ( $\sigma_{\text{posterior}}/\sigma_{\text{meas}}$ ) over iterations for different ratios of process to measurement noise ( $\sigma_{\text{stoch}}/\sigma_{\text{meas}}$ , shown by color). Solid lines represent the Bayesian update process, while dashed lines indicate the corresponding asymptotic limits. The convergence rate and final value depend strongly on the noise ratio, with the dotted line indicating the ideal trend for vanishing intrinsic stochasticity.

multishot device can accumulate equivalent information about the average properties of the laser over time.

**State-Resolving Regime** ( $\sigma_{\text{meas}}^2 < 2\sigma_{\text{stoch}}^2$ ): In this regime, we achieve sufficient precision to meaningfully distinguish single shots. Here, a measurement can provide information about the instantaneous system state that allows our knowledge to surpass its average behavior, i.e.  $\sigma_{\infty|\infty}^2 < \sigma_{\text{stoch}}^2$ . Note that the stationary solution of the update weight in this regime  $\gamma = \sigma_{\infty|\infty}^2/(\sigma_{\infty|\infty}^2 + \sigma_{\text{meas}}^2) > 0.5$ , meaning during Bayesian inference, the measurement weight amounts to more than 50% and the prior contribution falls below 50%; so, a single measurement in this regime takes less advantage of prior information.

This threshold helps us to better understand what truly constitutes “single-shot” measurement devices. Rather than being an intrinsic property of the measurement apparatus, the single-shot capability is now understood as a relationship between measurement precision and intrinsic stochasticity. Instead of relying on arbitrary definitions of measurement completeness, we can now precisely quantify when a device can meaningfully resolve individual system states. While there is certainly a significant practical advantage in acquiring as much diverse information as possible within a single shot, additional information only emerges as the measurement precision crosses this critical threshold relative to the system's inherent variability. Our analysis also confirms the intuitive notion that a state-resolving device must inherently rely more on the data from a single measurement than on prior information. Note that this is a general statement that includes any type of prior information, including sophisticated artificial neural network prediction.

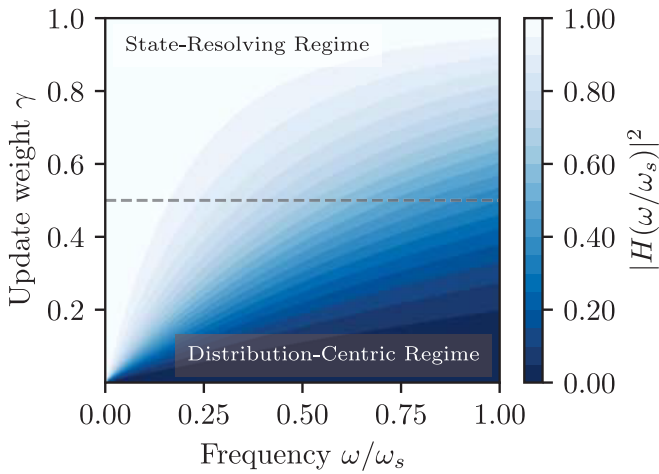
**1.3. Relationship Between Noise and Sampling Frequency.** A critical aspect in laser metrology is understanding the interplay between intrinsic stochasticity  $\sigma_{\text{stoch}}^2$  and our measurement paradigm. What appears as stochastic noise in a laser system often represents deterministic processes occurring at timescales beyond our measurement resolution. This phenomenon can lead to a deterministic, dynamic system masquerading as a stationary, stochastic one if the sampling frequency is insufficient. This is analogous to aliasing in signal processing, where undersampling can lead to misinterpretation of high-frequency components as lower-frequency variations. To better understand this, we consider the relationship between our measurement frequency and the characteristic frequencies of the laser system's dynamics.

In its asymptotic limit, the Bayesian update Eq. 5 behaves like an exponential moving average filter with the update weight  $\gamma$  as time constant. The frequency response of the system is then given as

$$\|H(\omega/\omega_s)\|^2 = \gamma^2 / (1 - 2(1 - \gamma) \cos(\omega/\omega_s) + (1 - \gamma)^2), \quad [9]$$

where  $\omega_s$  is the sampling rate. From this equation, we see how a smaller update weight leads to the attenuation of higher frequencies and thereby obfuscates dynamics; see Fig. 3. At our regime threshold, we find that the Nyquist frequency  $\omega_s/2$  is attenuated to about 67%. We can roughly quantify the ability to resolve frequencies using the cutoff frequency  $\omega_c = \arccos((\gamma^2 + 2\gamma - 2)/(2\gamma - 2))\omega_s$ , for which the spectral power decays to  $\|H(\omega_c/\omega_s)\|^2 = 1/2$ . Shot-to-shot dynamics at frequencies above this cutoff tend to not be resolved and will appear stochastic.

The concept of intrinsic stochasticity, therefore, emerges as a relative construct, intricately tied to our measurement capabilities



**Fig. 3.** Frequency response  $|H(\omega/\omega_s)|^2$  as a function of normalized frequency  $\omega/\omega_s$  and update weight  $\gamma$ . This plot illustrates how the system's frequency response changes with different update weights, highlighting the trade-off between noise reduction and dynamic response in the Bayesian inference process.

rather than an absolute property of the laser system. In practical terms, a 1 Hz laser system may for instance exhibit fluctuations at tens or hundreds of Hz due to mechanical vibrations. As the laser's repetition rate, and with it the minimum sampling rate, lies below these frequencies, these deterministic oscillations manifest as apparent randomness in our measurements. This effect propagates across timescales: As we increase reliance on prior measurements, we inevitably sacrifice frequency resolution, potentially misclassifying even slower drifts as “noise.”

This perspective reframes our understanding of intrinsic stochasticity within the context of laser metrology as a delineation between deterministic and probabilistic modeling. This boundary is not fixed but dynamic, shaped by our measurement capabilities and analysis methods. As we enhance our ability to resolve and model faster laser dynamics, we expand the deterministic realm of our understanding. Conversely, unresolved dynamics—whether due to sampling limitations or genuinely stochastic processes—fall into the domain of intrinsic stochasticity. This perspective transforms the intrinsic stochasticity estimation from a mere quantification of uncertainty into a crucial tool for defining the limits of our predictive capabilities. It determines whether we can make precise statements about individual laser shots or must resort to statistical predictions about consecutive shots.

**1.4. Estimating Intrinsic Stochasticity.** With this refined understanding of noise in laser systems, we can now discuss how to learn to delineate the boundary between resolved dynamics and unresolved fluctuations in the laser's behavior.

Given knowledge of the measurement noise, typically obtained through device calibration (Sections 2–4.3), we can estimate the intrinsic stochasticity by leveraging the discrepancy between what we predict and what we observe. This residual serves as a measure of model misspecification and forms the foundation of our adaptive estimation framework. Note that this step uses a quasi-deterministic prediction, meaning that uncertainty of trend components in the reconstruction is assumed as negligible compared to the instantaneous stochastic noise, which allows for real time analysis. Otherwise, a joint inference would be necessary, which is computationally far more expensive. For simplicity, we again sketch the derivation for a single variable with  $\mu_{\text{meas}}$  and  $\sigma_{\text{meas}}^2$  as the mean and variance of the current

measurement, as well as  $\mu_{\text{posterior}}$  and  $\sigma_{\text{posterior}}^2$  as the updated estimates for the parameter and its variance.

The residual  $\Delta\mu$  is the difference between the predicted state mean and the measurement mean

$$\Delta\mu = \mu_{\text{meas}} - \mu_{\text{pred}} \quad [10]$$

and the variance of the residual combines the variances of the prediction and the measurement

$$\sigma_{\Delta\mu}^2 = \sigma_{\text{pred}}^2 + \sigma_{\text{meas}}^2 \approx \sigma_{k|k}^2 + \sigma_{\text{stoch}}^2 + \sigma_{\text{meas}}^2 \quad [11]$$

The approximation used here is motivated in Section 1.6. From here, the likelihood of the measurement given the current model state is

$$p(\Delta\mu, \sigma_{\Delta\mu}^2) = \frac{1}{\sqrt{2\pi\sigma_{\Delta\mu}^2}} \exp\left(-\frac{(\Delta\mu)^2}{2\sigma_{\Delta\mu}^2}\right) \quad [12]$$

The question is how to change  $\sigma_{\text{stoch}}^2$  to maximize the likelihood. For this, we consider the roots of the log-likelihood

$$\frac{\partial \log p(\Delta\mu, \sigma_{\Delta\mu}^2)}{\partial \sigma_{\text{stoch}}^2} = 0$$

Solving this equation, while dropping the dependency of  $\sigma_{k|k}^2$  on the intrinsic stochasticity itself in Eq. 6 as higher-order effect, we find the rather intuitive result

$$\sigma_{\text{stoch}}^2 \approx (\Delta\mu)^2 - \sigma_{k|k}^2 - \sigma_{\text{meas}}^2. \quad [13]$$

However, this example only considers a single measurement, which is clearly insufficient to estimate the actual value of  $\sigma_{\text{stoch}}^2$ ; even for a perfect prediction of the mean and no measurement noise the residual is only a sample from the distribution  $\Delta\mu \sim \mathcal{N}(0, \sigma_{\text{stoch}}^2)$ . Assuming that the intrinsic stochasticity is constant or slowly varying, we can use a sequence of  $n$  measurements to achieve an estimate with variance  $\sigma_{\text{stoch}}^2/n$ .

It should be noted that the sample variance of the residuals can sometimes be smaller than the sum of the known variances. This can lead to physically meaningless results. Even with this, a variance matrix needs to be positive semidefinite, which might not be the case all the time with this simple approximation. Therefore, we use this analytic equation as a starting point to fit the intrinsic stochasticity by maximizing the log-likelihood numerically via Stochastic Gradient Descent. For simplicity we do not model correlations between modes and hence assume the intrinsic stochasticity matrix to be diagonal.

Taking into account that our model  $\hat{f}(t_k)$  is imperfect, we can interpret Eq. 13 as

$$\hat{\sigma}_{\text{stoch}}^2 = \text{Var}(x_k - \hat{f}(t_k)) - \sigma_{\text{meas}}^2 \quad [14]$$

to find our imperfect description  $\hat{\sigma}_{\text{stoch}}^2$  of the true stochastic noise  $\sigma_{\text{stoch,true}}^2$ . It can be deduced that

$$\hat{\sigma}_{\text{stoch}}^2 \geq \sigma_{\text{stoch,true}}^2 \quad [15]$$

Hence, the reconstructed  $\hat{\sigma}_{\text{stoch}}^2$  is an upper bound for the true value. Intuitively speaking, the repetition rate of measurements influences whether some deterministic processes might

be absorbed into the stochastic noise but could have been resolved if measured at smaller timescales. However, we obtain a conservative estimate of stochasticity, which translates directly to a stricter bound on measurement requirements, ensuring our conclusions remain valid.

This estimation of intrinsic stochasticity is deeply intertwined with our ability to separate deterministic and stochastic components. The decomposition is only meaningful if we can accurately estimate the deterministic component  $f(t_k)$  - otherwise, errors in estimating  $f(t_k)$  will be absorbed into our estimate of  $\epsilon_k$ , leading to biased estimates of  $\sigma_{\text{stoch}}^2$ . The challenge becomes particularly acute in the state-resolving regime identified in Section 1.2: When our measurement precision is sufficient to resolve individual states, each measurement captures not just  $f(t_k)$  but also (part of) the instantaneous stochastic fluctuation  $\epsilon_k$ . Using such measurements directly for prediction would erroneously treat these random fluctuations as part of the deterministic evolution. Because of this, pure state-space models like the adaptive Kalman filter (31) are not applicable to our problem; they systematically overestimate the system's noise. To address this, in the next section, we focus on estimating the deterministic component of the system evolution, before combining it with our stochasticity estimates to form a complete prediction model.

**1.5. Local Linear Approximation.** For sufficiently small time intervals  $\Delta t$ , we can approximate the deterministic evolution locally by a linear function:

$$f(t_k + \Delta t) \approx x_k + v_k \Delta t, \quad [16]$$

where  $x_k$  represents the true state and  $v_k$  the instantaneous rate of change. This linear approximation is valid when changes in the deterministic component are small compared to stochastic fluctuations, i.e.,  $|v_k \Delta t| \ll \sigma_{\text{stoch}}$ . Under this condition, any errors from nonlinear behavior are negligible compared to the intrinsic stochasticity we aim to characterize. Using Holt's linear exponential smoothing (32) with additive trends as robust framework for tracking evolving systems while filtering out noise, our best estimate  $\hat{x}_k$  of this state evolves as

$$\hat{x}_{k+1} = (1 - \alpha) \cdot [\hat{x}_k + \hat{v}_k \Delta t] + \alpha \cdot \mu_{k|k} \quad [17]$$

$$\hat{v}_{k+1} = (1 - \beta) \cdot \hat{v}_k + \beta \cdot \frac{\hat{x}_{k+1} - \hat{x}_k}{\Delta t}, \quad [18]$$

where  $\mu_{k|k}$  is the posterior mean from our measurements up to time  $t_k$ . The parameters  $\alpha$  and  $\beta$  are learned by maximizing the log-likelihood over a finite window:

$$\begin{aligned} (\hat{\alpha}, \hat{\beta}) &= \operatorname{argmax}_{\alpha, \beta} \sum_{k=1}^N \log p(y_k | \hat{x}_k(\alpha, \beta)) \\ &= \operatorname{argmin}_{\alpha, \beta} \sum_{k=1}^N (y_k - \hat{x}_k(\alpha, \beta))^2, \end{aligned} \quad [19]$$

where the equality follows from our Gaussian assumptions. By using a finite window, we assume local stationarity of the parameters while allowing the model to adapt to changes in system behavior. Note that our model assumes that variances are approximately constant, which is the case once the measurement system approaches the steady state discussed in Section 1.2.

More sophisticated approaches to learn  $f(t)$  could be considered. Higher-order models or neural networks, for instance, could learn complex temporal patterns in the system's evolution.

However, such approaches come with their own drawbacks: The more complex the model, the more additional parameters need to be estimated, potentially making the decomposition between deterministic and stochastic components less stable. Fitting more parameters allows for a seemingly larger fit accuracy, however the precision would go down as the overall uncertainty would rise simultaneously. This would counteract the requirement for the adaptive intrinsic uncertainty estimation, which requires a quasi-deterministic prediction. A related effect of more complex models is the potential confusion between genuine patterns and spurious correlations in the noise. As our experimental results in Section 2.2 demonstrate, the local linear model provides a pragmatic balance that can accurately separate predictable changes from stochastic variations, while its simple, recursive structure makes it compatible with real-time evaluation.

**1.6. Prediction Model.** Combining our local linear approximation from Section 1.5 with the estimated stochastic component from Section 1.4 according to Eq. 1 yields a probabilistic prediction model, whose mean follows the deterministic model, i.e.,  $\mu_{k+1|k} = \hat{x}_{k+1}$ .

The prediction's variance is calculated as the combination of the uncorrelated uncertainties of the noise and the local linear model:

$$\sigma_{k+1|k}^2 = \sigma_{\text{stoch}}^2 + \sigma_{\text{pred}}^2(\Delta t) \quad [20]$$

Here, the prediction uncertainty would generally be

$$\sigma_{\text{pred}}^2 = \sigma_{k|k}^2 + (\Delta t)^2 \sigma_v^2 + 2\Delta t \sigma_{xv}, \quad [21]$$

where  $\sigma_v^2$  and  $\sigma_{xv}$  are model parameter uncertainties. Here, we assume that the model noise is dominated by our uncertainty in the true state, given by our posterior variance  $\sigma_{k|k}^2$ , and we can drop the uncertainty on the trend of the local linear model. This is justified because a) we use a simple model with small parameter uncertainty and b) when the time interval  $\Delta t$  is sufficiently small. This condition is typically satisfied in our application, where measurement rates are high compared to the timescale of systematic drifts. Even when this approximation introduces some error, it manifests as an apparent increase in the estimated intrinsic stochasticity, ensuring our uncertainty estimates remain conservative (cf. Eq. 15). From this follows:

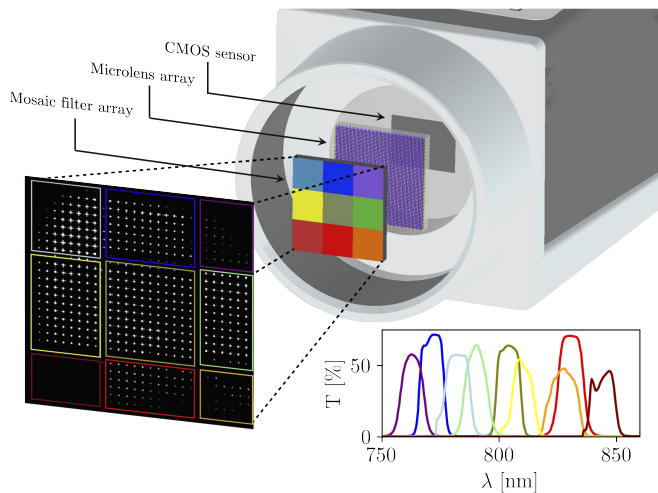
$$\sigma_{k+1|k}^2 \approx \sigma_{\text{stoch}}^2 + \sigma_{k|k}^2 \quad [22]$$

The probabilistic predictions from this combined model serve as our prior for subsequent Bayesian state updates, completing the loop back to Section 1.1.

## 2. Experiment

We now apply the Bayesian framework developed in the last section to the measurement of spatiotemporal couplings in a petawatt laser. We demonstrate single-shot sensitivity to spatiotemporal couplings of the ATLAS-3000 petawatt laser using a new, simple measurement device. Information about the experimental setup at the ATLAS-3000 laser as well as the applied Bayesian modal reconstruction, the calibration and measurement noise estimation can be found in *Materials and Methods*.

**2.1. Single-Shot Measurements Using Mosaic Filter Array.** In ref. 14, we introduced FALCON, an approach to reconstruct spatial temporal couplings based on measurements of the entire



**Fig. 4.** Single-shot FALCON, consisting of a mosaic bandpass filter array in front of a microlens array placed in effective focal length in front of the camera. The *Inset on the Bottom Right* shows the corresponding spectral transmission of the mosaic filter array, measured with a spectrophotometer. Magnified is a raw image with filter positions marked in their respective colors. Note that the weaker signal behind the filters centered at around 765 and 845 nm is due to the spectrum of the ATLAS laser.

wavefront at frequencies selected with narrowband filters. While simple and fast compared to established schemes, the method requires on the order of  $\sim 10$  to 50 shots, to a) scan through different spectral filters and b) to have enough statistics to suppress noise due to pointing jitter within the measurement set.

To alleviate this problem, we introduce a single-shot variation of the technique that promises easy experimental implementation. The key idea is that the Zernike modes cover the full spatial domain of the measurement—as opposed to the pixel basis (which is often applied in the zonal reconstruction approach). This means that one can *in principle* (given a noise-free environment) use a subaperture measurement to deduce the entire wavefront. By filtering each subaperture with different spectral filters, we can then measure spatiotemporal couplings. *In practice* the inevitable noise limits our measurement capabilities to low Zernike order STCs, which depends on the size of the subapertures. However, the fact that we measure a pulse that is confined to a small spatial and temporal volume in the focal

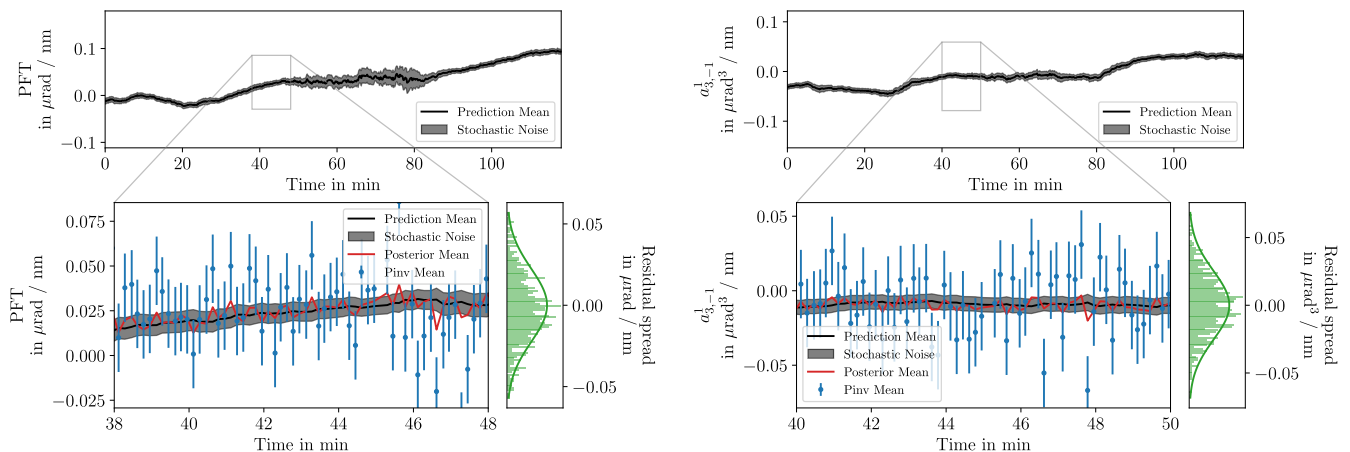
plane (it is focus-able) guarantees that low-order STCs are the dominant ones.

In this subaperture measurement approach, a mosaic filter array divides the wavefront into a series of subapertures  $A_i$  (where  $i$  indexes the subaperture), each with its own narrowband filter characterized by a central wavelength  $\omega_i$ . The array consists of nine bandpass filters by Omega Optical, glued together into a  $3 \times 3$  array (33). The device is sketched in Fig. 4. The nine bandpass filters cover the spectral range from 760 nm to 840 nm; see *Bottom Right* of Fig. 4 for spectrophotometer measurements of their transmission curves. The arrangement allows us to measure the wavefront in different sections of the beam for different colors. Our analysis shows that this configuration is well suited to measure pulse front tilt and curvature, in particular.


To ensure the accuracy of this Single-Shot FALCON device, it was validated against FALCON measurements in test scenarios with a broadband fs-oscillator. It should be noted that, before measuring with the Single-Shot FALCON, it is beneficial to measure the wavefront once to ensure that no high-order spatial distortions are present that could be misinterpreted as lower-order spatio-spectral effects. More specifically, if one uses a FALCON or similar device to measure the full spatio-spectral phase at the start, this can be used as a prior on the Single-Shot FALCON measurements instead of a flat phase. A more detailed explanation about the validity of Single-Shot FALCON reconstructions can be found in *SI Appendix*.

Note that an alternative implementation of a single-shot device could be to use a beamsplitter, e.g., a diffractive optical element (DOE) (33), and to spectrally filter individual copies of the entire beam. However, the introduction of a DOE would increase complexity and cost of the setup. In contrast, the solution presented here is based on the addition of a single optical element, the mosaic filter, in front of an existing wavefront sensor, making it simple to implement and use.

**2.2. Single-Shot Measurements of STCs.** To demonstrate our diagnostic, we monitored the change of the STCs of the ATLAS-3000 Laser during continuous operation for about two hours. During this time, the laser system was not actively changed, but still subject to thermalization and other types of drifts. The preprocessed local phase gradients from the Single-Shot FALCON supporting this paper are available (38).



**Fig. 5.** The evolution of the pulse front tilt (*Left*) and linear-frequency dependent-coma (*Right*) prediction (=prior) values together with the intrinsic stochasticity. The zoomed sections add the posterior mean as well as the pseudoinverse (Pinv) as introduced in *Materials and Methods* in Eqs. 28 and 29. The Pinv solution also shows the measurement noise in the coefficient space as error bars. The histograms show the spread of the residuals as well as the equivalent Gaussian distribution obtained from our estimate Eq. 11.



**Fig. 6.** Estimated intrinsic stochasticity and update weight for PFT (*Top*) and linear coma (*Bottom*) as a function of time. The estimated intrinsic stochasticity  $\hat{\sigma}_{\text{stoch}}^2$  is calculated as described in Section 1.4 and is equivalent to the size of the gray area in Fig. 5. The update weight is calculated as shown in Eq. 7. In both plots it can be seen that the update weight and the estimated intrinsic stochasticity correlate as one would expect and a higher value for the latter guarantees that the new measurement is weighted more.

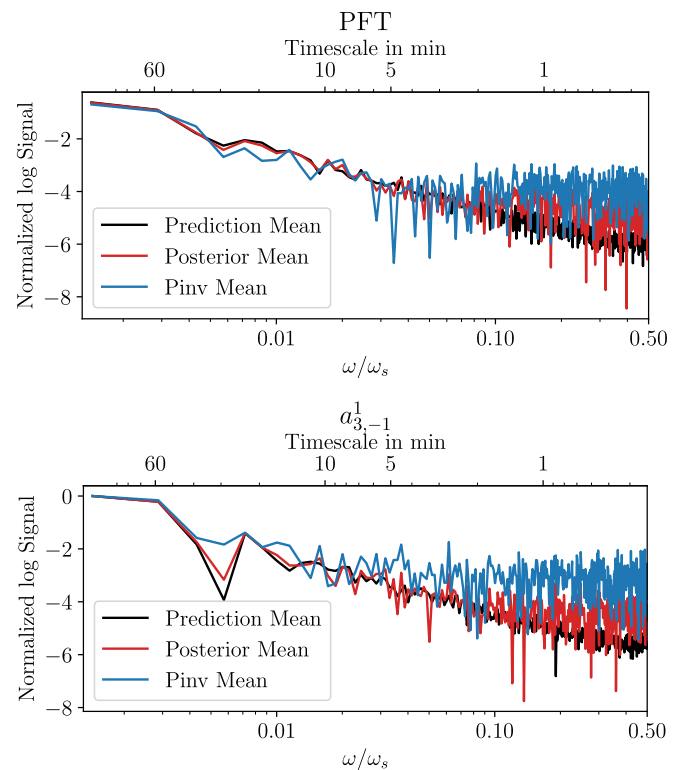
We used a sampling frequency  $\omega_s = 0.1$  Hz and a sliding window size  $N = 20$  to fit both the stochasticity and the local linear model. This covers about 3 min. Parameter optimization yields, on average,  $\bar{\alpha} = 0.1596$  and  $\bar{\beta} = 0.0080$ . These values are easier interpreted as “half-life” periods: The level component ( $\alpha$ ) has a half-life of about 40 s, meaning that the influence of any measurement on the estimated state decays to half its original value after this time. The trend component ( $\beta$ ) has a much longer half-life of approximately 14.4 min, indicating that—on the timescale of the measurements—trend changes slowly. The results for the prediction mean  $\hat{x}$  and the intrinsic stochasticity  $\sigma_{\text{stoch}}$  as uncertainty on that value are exemplary shown for the PFT and linear-frequency dependent-coma (linear coma) in Fig. 5.

It can be seen that there are dynamics on the longer timescales that go beyond the laser’s intrinsic stochasticity. As discussed in Section 1.3, if one would measure with a sampling rate at this timescale the fitted intrinsic stochasticity would be larger, as these dynamics would be incorporated into the stochasticity estimate. To prove the legitimacy of these results, Fig. 5 also shows a ten minute section together with the results of the posterior mean as well as a result if only a least-squares solution  $\mu_k$  to the forward matrix—calculated using a pseudoinverse (Pinv) approach as described in *Materials and Methods* with Eqs. 28 and 29—would have been calculated. It can be seen that the posterior mean follows the least-squares results far more than the predicted mean but its fluctuations stay at the same order as the determined intrinsic stochasticity. This indicates that the predicted mean together with the intrinsic stochasticity accurately represents the intrinsic systems state with its uncertainty. The average improvement of the posterior uncertainty compared to the measurement uncertainty is about 54% and 60%, respectively. The histograms show that the residuals match the distribution (green curve) that one would expect from Eq. 11. Additional plots for other modes can be found in *SI Appendix*.

Additionally, one can look at the change of the update weight compared to the intrinsic stochasticity. The results are shown for PFT and linear coma in Fig. 6. From both, it can be seen that the update weight is directly correlated to the intrinsic stochasticity. The update weights remain between bounds of 0.3 to 0.75 and 0.25 to 0.5 for PFT and linear coma, respectively. This implies that for the PFT with this measurement setup we are at the edge between a state-resolving and a distribution-centric regime. More precisely, according to Eq. 9, state components are attenuated by around 50% in this setting. This implies that there is a limited, but nonzero advantage to having a “single-shot” device with this measurement uncertainty looking. But a few- or multishot device can also be a reasonable choice. However, only this gives us an insight into the intrinsic stochasticity. For the linear coma, we are entirely in the distribution-centric regime, whereas the reconstruction for the Tilt (shown in *SI Appendix*) is entirely in the state-resolving regime. This demonstrates the more nuanced view provided by our framework.

For the PFT one can calculate the intensity loss in focus  $\chi_{\text{intensity}}$  due to the intrinsic stochasticity on the PFT by the following analytic formula, adapted from ref. 34:

$$\chi_{\text{intensity}} = \left( 1 + \left( \frac{\sigma_{\text{stoch}}^{\text{PFT}} \sigma_{\lambda} D}{0.84 \lambda_0} \right) \right)^{-2} \quad [23]$$



**Fig. 7.** Logarithmic spectra for the prediction (=prior) mean  $\mu_{k|k-1}$ , posterior mean  $\mu_{k|k}$  and pseudoinverse mean  $\mu_k$  of the PFT (*Top*) and linear coma (*Bottom*). This plot is obtained by Fourier transforming the corresponding curves in Fig. 5. The frequency at which the curves separate can be seen as the threshold at which long-term trends become less dominating and short-term fluctuations take over. The Pinv mean includes all the measurement noise and intrinsic stochasticity and hence has the strongest signal at short timescales. The posterior mean dampens the influence of the measurement noise but includes the intrinsic stochasticity. The prior mean only includes the long-term trends.

with  $\sigma_\lambda = 30$  nm being the bandwidth of the equivalent Gaussian spectrum,  $D = 27$  cm the laser's diameter and  $\lambda_0 = 800$  nm. As  $\sigma_{\text{stoch}}^{\text{PFT}}$  is between 0.005 and 0.02  $\mu\text{rad}/\text{nm}$ , this corresponds to 0.3 and 5.5% decrease in focus intensity. Hence, if the prediction mean of the PFT would be compensated for by adjusting the compressor grating angle, this would be the decrease compared to an ideal flat-top beam. As these fluctuations can not be compensated for due to their apparent stochastic nature, this provides a lower limit for peak intensity fluctuations of an otherwise perfect laser.

Last, Fig. 7 shows how the spectra for the predicted, posterior and pseudoinverse mean of PFT and linear coma look like. It can be seen that for low sampling rates the three shown curves overlap whereas features at higher rates the pseudoinverse mean gives the strongest signal and the predicted mean the lowest, as expected. The point, where the three curves separate gives, up to which timescales general trends are dominating and from when on statistical fluctuations dominate. For the PFT, this is at approximately 8 min and for the linear coma between 40 and 20 min. Notice also how the posterior's frequency spectrum is below the least square solution's for high frequency modulations, indicative of how the Bayesian update helps us to remove measurement noise while still resolving the intrinsic stochasticity of the system.

### 3. Conclusions and Outlook

In this work, we have introduced a Bayesian perspective on laser characterization. In this approach, the very concept of what constitutes a "single-shot" measurement is challenged and instead, we argue that it is more suitable to distinguish between the capacity to resolve the average state of the system versus resolving stochastic shot-to-shot fluctuations. Along with this, we have developed a rigorous framework that handles Bayesian updates, as well as dynamic fitting of the system's underlying, predictable trends and unpredictable variability.

We have then applied the framework to the problem of monitoring spatiotemporal couplings in a petawatt laser system. By employing a modal reconstruction technique based on Zernike-Legendre polynomials and a mosaic filter array combined with a Shack-Hartmann wavefront sensor, we have demonstrated the retrieval of key STCs from single measurements. Crucially, the Bayesian framework allows us to obtain a posterior uncertainty that is well below the pure measurement uncertainty.

Within this framework, it is also possible to see the measurement problem from the other end. Instead of asking what is the lowest reconstruction uncertainty given the measurement uncertainty that we can achieve, one can ask which measurement precision is necessary to reach a certain reconstruction precision, given the obtained system knowledge. This can be seen as another interpretation of Eq. 6:

$$\sigma_{\text{meas}}^2 = \left( \frac{1}{\sigma_{\text{posterior}}^2} - \frac{1}{\sigma_{\text{prior}}^2} \right)^{-1} \quad [24]$$

In this sense, the intrinsic stochasticity included in the prior can be seen as an inherent upper bound on the achievable posterior uncertainty, with each measurement reducing this toward a desired threshold. In an information theoretical view, the measurement only needs to provide a limited amount of new information, given by the information gain  $IG$ , which is the difference in entropy before and after the measurement:

$$IG = H(\text{prior}) - H(\text{posterior}) \quad [25]$$

$$= \frac{1}{2} \left( \log(\det \Sigma_{\text{prior}}^2) - \log(\det \Sigma_{\text{posterior}}^2) \right) \quad [26]$$

Future work will focus on extending the proposed approach to more complex spatiotemporal field structures. Our framework can both serve as input for predictive control algorithms for beam stabilization and as an estimate for the amount of fluctuations that cannot be corrected for. Crucially, while it might not be possible to correct for all changes a priori, we can use the a posteriori information from a state-resolving diagnostic to remove uncertainty from physical models. For instance, shot-to-shot information on pulse front tilt may serve as crucial input to refine probabilistic models of laser wakefield accelerators (24, 35).

The framework itself can be developed further, for instance using other basis functions or using Gaussian processes to model both correlations in time and between variables. The availability of uncertainty estimates allows the development of new, adaptive diagnostic: The principles of Bayesian inference can be extended beyond data analysis to inform the design of experiments themselves as demonstrated for spectroscopy in ref. 36. Bayesian experimental design aims to optimize measurement strategies to maximize information gain, which could be particularly valuable in the context of single-shot diagnostics where each measurement is precious. Additionally, this framework can be adapted to more sophisticated single-shot laser diagnostics like RAVEN (37) that has far more measurement points than the Single-Shot FALCON. It could also include additional data from other diagnostic devices, and hence, a combined reconstruction from multiple devices should be feasible. Similarly, it could also improve few- or multishot diagnostics, where the additional information of each shot is limited. The framework would enable more accurate reconstructions within fewer shots.

In conclusion, our work has introduced a framework for single-shot measurements of STCs, leveraging Bayesian inference and prior knowledge to refine the concept of single-shot measurements. This paradigm shift opens up possibilities for efficient and targeted optimization of high-power laser systems and their applications.

### 4. Materials and Methods

**4.1. Experimental Setup at the ATLAS-3000 Petawatt Laser.** The ATLAS-3000 laser is a Titanium:Sapphire (Ti:Sa) laser system, capable of a maximum pulse energy of 90 J prior to compression. Post compression, the laser pulse has a duration of 27 fs, categorizing its peak power in the Petawatt regime. The beam diameter after the final expander measures 27 cm. For particle acceleration experiments, the laser pulse can be delivered with full energy or, for diagnostic purposes, significantly reduced energy by using a reflective attenuator placed behind the final amplifiers and before the laser beam expander and compressor. For the measurements presented in the following, the laser was operated at 15.0 J before attenuation and compression. Since particle acceleration experiments are highly sensitive to fluctuations in the laser intensity profile, a diagnostic that allows measurement of spatiotemporal couplings on a shot-to-shot basis is desirable.

Measurements are performed on a demagnified image of the laser's near field. For this purpose, we use a telescope consisting of a spherical mirror with  $f_1 = 10$  m focal length, otherwise used for electron acceleration experiments, combined with an  $f_2 = 0.5$  m achromatic doublet lens, followed by a 3:5 reimaging and demagnifying telescope. The laser enters the spherical mirror under a small angle, and the resulting aberrations are precorrected using a deformable mirror. The telescope images a plane of the near field approximately 5 m before the spherical mirror. The intermediate focus of the telescope has a size of approximately 30  $\mu\text{m}$  diameter. To calibrate the phase measurement,

we place a motorized pinhole with 20  $\mu\text{m}$  diameter at the focus position. The beam is fully imaged on the in-vacuum Shack-Hartmann sensor with a beam diameter of approximately 8 mm, corresponding to a demagnification of 1 : 34.

**4.2. Bayesian Modal Reconstruction.** As introduced in ref. 14, the spatio-spectral phase  $\Phi(x, y, \omega)$  of a laser pulse can be expressed in terms of the Zernike-Taylor coefficients  $a_{m,n}$ , which are the equivalent of  $x_k$  in the previous sections:

$$\Phi(x, y, \omega) = \sum_{m,n,i} a_{m,n}^i (\omega - \omega_0)^i \cdot Z_n^m(x, y), \quad [27]$$

where  $a_{m,n}^i = \left( \partial^i a_{m,n} / \partial \omega^i \right)_{\omega=\omega_0}$ ,  $\omega_0$  is the central frequency and  $Z_n^m$  are the Zernike polynomials.

A frequency-resolving Shack-Hartmann wavefront sensor measures the gradient of the spatio-spectral phase, which can be related to coefficients describing spatiotemporal couplings through a transfer matrix  $\mathbf{T}$ :

$$\vec{d}_k = \mathbf{T} \cdot \vec{a}_k + \vec{n}, \quad [28]$$

where  $\vec{d}_k$  is a vector containing the measured centroid positions and  $\vec{a}_k$  is a vector of the retrieved coefficients for spatiotemporal couplings.  $\vec{n}$  describes the noise of our measurement, which is assumed to be Gaussian. The natural approach to solve this equation for  $\vec{a}_k$  is to calculate the inverse of  $\mathbf{T}$ . However, as there are typically more measurement points  $\vec{d}_k$  than coefficients  $\vec{a}_k$  to be reconstructed,  $\mathbf{T}$  is nonsquare and only a so called pseudoinverse  $\mathbf{T}^+$  can be calculated numerically. The resulting reconstructed coefficients

$$\hat{\vec{a}}_k = \mathbf{T}^+ \cdot \vec{d}_k \quad [29]$$

are equivalent to the least-squares solution of Eq. 28. With this simple approach, the subaperture nature of our mosaic filter measurements introduces an additional challenge. In a zonal reconstruction, this would be completely impossible. In the modal approach, the transfer matrix becomes less well-conditioned compared to full-aperture measurements, as we are attempting to reconstruct global modes from a reduced set of local measurements. This poorer conditioning makes the system more sensitive to noise, potentially amplifying measurement uncertainties in the reconstruction. This is where our Bayesian framework proves especially valuable. By incorporating prior knowledge about the wavefront and properly accounting for measurement uncertainties, we can achieve robust reconstruction even with partial aperture measurements. The formalism outlined in the previous sections adapts naturally to the vectorial case of reconstructing multiple Zernike-Taylor coefficients at once through matrix equations, allowing us to properly weight the contribution of each subaperture measurement while maintaining the global coherence of the reconstruction. This approach transforms what might be an ill-conditioned inverse problem in a traditional least-squares framework into a well-posed probabilistic inference task. The equivalent of Eq. 5 becomes

$$\vec{a}_{k|k} = (\mathbf{1} - \mathbf{\Gamma}_k) \vec{a}_{k|k-1} + \mathbf{\Gamma}_k \cdot (\mathbf{T}^+ \vec{d}_k) \quad [30]$$

1. C. Danson, D. Hillier, N. Hopps, D. Neely, Petawatt class lasers worldwide. *High Power Laser Sci. Eng.* **3**, e3 (2015).
2. J. W. Yoon *et al.*, Realization of laser intensity over  $10^{23}$  W/cm<sup>2</sup>. *Optica* **8**, 630–635 (2021).
3. S. W. Jolly, O. Gobert, F. Quéré, Spatio-temporal characterization of ultrashort laser beams: A tutorial. *J. Opt.* **22**, 103501 (2020).
4. P. Bowlan, P. Gabolde, R. Trebino, Directly measuring the spatio-temporal electric field of focusing ultrashort pulses. *Opt. Express* **15**, 10219–10230 (2007).
5. P. Bowlan, P. Gabolde, M. A. Coughlan, R. Trebino, R. J. Levis, Measuring the spatiotemporal electric field of ultrashort pulses with high spatial and spectral resolution. *J. Opt. Soc. Am. B* **25**, A81–A92 (2008).
6. V. Gallet, S. Kahaly, O. Gobert, F. Quéré, Dual spectral-band interferometry for spatio-temporal characterization of high-power femtosecond lasers. *Opt. Lett.* **39**, 4687–4690 (2014).
7. B. Alonso *et al.*, Spatiotemporal amplitude-and-phase reconstruction by fourier-transform of interference spectra of high-complex-beams. *J. Opt. Soc. Am. B* **27**, 933–940 (2010).
8. B. Alonso, M. Miranda, Í. J. Sola, H. Crespo, Spatiotemporal characterization of few-cycle laser pulses. *Opt. Express* **20**, 17880–17893 (2012).

and Eq. 6 becomes

$$\Sigma_{k|k}^2 = (\mathbf{1} - \mathbf{\Gamma}_k) \Sigma_{k|k-1}^2 (\mathbf{1} - \mathbf{\Gamma}_k)^T + \mathbf{\Gamma}_k \Sigma_{\text{meas}}^2 \mathbf{\Gamma}_k^T \quad [31]$$

and the update weight Eq. 7 translates to

$$\mathbf{\Gamma}_k = \Sigma_{k|k-1}^2 \cdot \left( \Sigma_{k|k-1}^2 + \Sigma_{\text{meas}}^2 \right)^{-1}. \quad [32]$$

In this notation, bold capital  $\Sigma$  and  $\mathbf{\Gamma}$  represent the matrix versions of the scalar variance  $\sigma$  and update weight  $\gamma$ , respectively.

**4.3. Calibration and Estimating Measurement Noise.** A key component to the above formalism is knowledge of the measurement uncertainty  $\Sigma_{\text{meas}}^2$ . The matrix formalism allows us to calculate the variance on the coefficients given the uncertainty of the local phase gradients  $\vec{n}$ :

$$\Sigma_{\text{meas}}^2 = \mathbf{T}^+ \mathbf{n} \mathbf{T}^{+T}, \quad [33]$$

where  $\mathbf{n}$  is the diagonal covariance matrix assigned to  $\vec{n}$ . There are two types of uncertainty that have to be considered. For obvious reasons, the foci spot positions of a Shack-Hartmann can be determined only up to a certain precision due to their finite width on a sensor. We use a combination of a local 2D peak finder for coarse position determination followed by a 2D Gaussian fit for subpixel accuracy. The second reason is that fitting a limited modal model to a zonal detector setup cannot include artifacts at high spatial frequencies. This could only be compensated by fitting high-order Zernike polynomials, which showed to be numerically unstable. Therefore, these high-order fluctuations manifest as an uncorrelated noise term. This adds additional uncertainty to any fit in the pixel base domain.

For our measurements, we used the Thorlabs MLA300-14AR microlens array and the IDS U3-3992SE camera with a pixel size of 2.74  $\mu\text{m}$ . From calibration measurements, we estimate the accuracy in the focal spot position reconstruction to 0.2 pixels. We assume this error for all well-illuminated pixels. While one might want to consider photon noise, its influence is on the focal spot estimate is difficult to estimate, so instead it was decided to remove individual microlens spots with insufficient illumination from the reconstruction set. We estimate the model fit uncertainty as an additional factor of 0.6 pixels.

**Data, Materials, and Software Availability.** Data have been deposited in Zenodo (<https://doi.org/10.5281/zenodo.16311477>) (38). All other data are included in the manuscript and/or *SI Appendix*.

**ACKNOWLEDGMENTS.** This work was supported by the Independent Junior Research Group "Characterization and control of high-intensity laser pulses for particle acceleration," DFG Project No. 453619281.

Author affiliations: <sup>a</sup>Centre for Advanced Laser Applications, Faculty of Physics, Ludwig-Maximilians-Universität München, Garching 85748, Germany; and <sup>b</sup>Department of Physics, Clarendon Laboratory, University of Oxford, Oxford OX1 3PU, United Kingdom

9. B. Alonso *et al.*, Frequency resolved wavefront retrieval and dynamics of diffractive focused ultrashort pulses. *J. Opt. Soc. Am. B* **29**, 1993–2000 (2012).
10. B. Alonso *et al.*, Characterization of sub-two-cycle pulses from a hollow-core fiber compressor in the spatiotemporal and spatio-spectral domains. *Appl. Phys. B* **112**, 105–114 (2013).
11. M. Miranda *et al.*, Spatiotemporal characterization of ultrashort laser pulses using spatially resolved fourier transform spectrometry. *Opt. Lett.* **39**, 5142–5144 (2014).
12. G. Pariente, V. Gallet, A. Borot, O. Gobert, F. Quéré, Space-time characterization of ultra-intense femtosecond laser beams. *Nat. Photon.* **10**, 547–553 (2016).
13. A. Borot, F. Quéré, Spatio-spectral metrology at focus of ultrashort lasers: A phase-retrieval approach. *Opt. Express* **26**, 26444–26461 (2018).
14. N. Weisse *et al.*, Measuring spatio-temporal couplings using modal spatio-spectral wavefront retrieval. *Opt. Express* **31**, 19733 (2023).
15. S. Smartshev *et al.*, Simple few-shot method for spectrally resolving the wavefront of an ultrashort laser pulse. *Opt. Lett.* **49**, 1900–1903 (2024).
16. P. Gabolde, R. Trebino, Single-shot measurement of the full spatiotemporal field of ultrashort pulses with multispectral digital holography. *Opt. Express* **14**, 11460–11467 (2006).

17. C. Dorrer, S. W. Bahk, Spatio-spectral characterization of broadband fields using multispectral imaging. *Opt. Express* **26**, 33387–33399 (2018).
18. Y. G. Kim *et al.*, Single-shot spatiotemporal characterization of a multi-PW laser using a multispectral wavefront sensing method. *Opt. Express* **29**, 19506–19514 (2021).
19. H. Tang *et al.*, Single-shot compressed optical field topography. *Light Sci. Appl.* **11**, 244 (2022).
20. S. Howard, J. Esslinger, R. H. W. Wang, P. Norreys, A. Döpp, Hyperspectral compressive wavefront sensing. *High Power Laser Sci. Eng.* **11**, e32 (2023).
21. B. Alonso, A. Döpp, S. W. Jolly, Space-time characterization of ultrashort laser pulses: A perspective. *APL Photon.* **9**, e3303575 (2024).
22. E. Kueny *et al.*, Wavefront analysis of a white-light supercontinuum. *Opt. Express* **26**, 31299–31306 (2018).
23. A. de Andres Gonzalez, *Enhancement of Few-Cycle Light Fields for Relativistic Nanophotonics* (Umeå University, Umeå, 2023).
24. A. Döpp *et al.*, Data-driven science and machine learning methods in laser-plasma physics. *High Power Laser Sci. Eng.* **11**, e55 (2023).
25. A. Debus *et al.*, Circumventing the dephasing and depletion limits of laser-wakefield acceleration. *Phys. Rev. X* **9**, 031044 (2019).
26. J. P. Palastro *et al.*, Dephasingless laser wakefield acceleration. *Phys. Rev. Lett.* **124**, 134802 (2020).
27. C. Caizergues, S. Smartsev, V. Malka, C. Thaury, Phase-locked laser-wakefield electron acceleration. *Nat. Photon.* **14**, 475–479 (2020).
28. H. Vincenti, F. Quéré, Attosecond lighthouses: How to use spatiotemporally coupled light fields to generate isolated attosecond pulses. *Phys. Rev. Lett.* **108**, 113904 (2012).
29. F. Quéré *et al.*, Applications of ultrafast wavefront rotation in highly nonlinear optics. *J. Phys. B At. Mol. Opt. Phys.* **47**, 124004 (2014).
30. J. K. Blitzstein, J. Hwang, *Introduction to Probability, Texts in Statistical Science* (CRC Press Taylor & Francis group, Boca Raton, 2nd edition, 2019).
31. A. Mohamed, K. Schwarz, Adaptive kalman filtering for INS/GPS. *J. Geod.* **73**, 193–203 (1999).
32. C. C. Holt, Forecasting seasonals and trends by exponentially weighted moving averages. *Int. J. Forecast.* **20**, 5–10 (2004).
33. D. Haffa *et al.*, Temporally Resolved Intensity Contouring (TRIC) for characterization of the absolute spatio-temporal intensity distribution of a relativistic, femtosecond laser pulse. *Sci. Rep.* **9**, 7697 (2019).
34. G. Pretzler, A. Kasper, K. Witte, Angular chirp and tilted light pulses in CPA lasers. *Appl. Phys. B* **70**, 1–9 (2000).
35. F. Irshad *et al.*, Pareto optimization and tuning of a laser wakefield accelerator. *Phys. Rev. Lett.* **133**, 085001 (2024).
36. J. Schroeder *et al.*, Information-optimal measurement: From fixed sampling protocols to adaptive spectroscopy. arXiv [Preprint] (2025). <https://doi.org/10.48550/arXiv.2505.14364> (Accessed 7 October 2025).
37. S. Howard *et al.*, Single-shot spatiotemporal vector field measurements of petawatt laser pulses. *Nat. Photon.* **19**, 898–905 (2025).
38. J. M. Esslinger *et al.*, A Bayesian perspective on single-shot laser characterization - Supplementing Dataset. Zenodo. <https://doi.org/10.5281/zenodo.16311477>. Deposited 25 July 2025.