Active Fourier spectroscopy: From fixed sampling protocols to adaptive measurement

J. Schröder,[1,2] S. Howard,[1,2,3] C. Eberle,[1,2] J. Esslinger,[1] N. Leopold-Kerschbaumer,[1,4,5] K. V. Kepesidis,[1,4,5] and A. Döpp[1,2,3,4,*]

[1]*Fakultät für Physik, Ludwig-Maximilians-Universität München, Am Coulombwall 1, 85748 Garching, Germany*
[2]*Munich Center for Machine Learning (MCML), München, Germany*
[3]*Department of Physics, Clarendon Laboratory, University of Oxford, Oxford OX1 3PU, United Kingdom*
[4]*Max Planck Institut für Quantenoptik, Hans-Kopfermann-Strasse 1, 85748 Garching, Germany*
[5]*Center for Molecular Fingerprinting (CMF), Budapest, Hungary*

(Received 16 July 2025; accepted 23 October 2026; published 17 November 2026)

Fourier-transform spectroscopy (FTS) is based on the Nyquist-Shannon theorem, which dictates a strict sampling protocol for the measurement of continuous signals. Here, we introduce active Fourier spectroscopy (AFS), a generalization of FTS in which measurements are still taken in the Fourier-conjugate delay domain, but where the most informative acquisition points are selected adaptively and the spectrum is reconstructed by sequential, uncertainty-aware inference rather than by fixed-grid Fourier inversion. Our approach relies on a theoretical insight that the Nyquist condition applies only under complete ignorance about the system under investigation, and shows how to systematically improve upon it when prior knowledge exists. AFS recovers conventional FTS when the prior is uninformative and falls back to it when prior knowledge is not supported by the data. We demonstrate AFS for optical spectroscopy in three settings vital to research and medicine: clinical FTIR (Fourier-transform infrared spectroscopy)-based blood plasma analysis, optical metrology of structured laser beams, and hyperspectral imaging with an RGB camera. In each case, AFS reaches a higher target reconstruction quality from far fewer acquisitions and provides complete uncertainty quantification in real time, propagated rigorously through downstream analysis.

DOI: 10.1103/shq1-j249

Fourier-transform spectroscopy (FTS) is one of the most widely used analytical methods in science, medicine, and industry today. Fourier-transform infrared (FTIR) spectroscopy [1], for instance, is ubiquitous in chemical laboratories [2,3], industrial quality control [4], medical diagnostics [5], and environmental monitoring [6]. At its core, FTS measures the fluence of light, F, as a function of the time delay, τ, between two identical copies of the optical signal. Depending on the delay, different wavelengths add constructively or destructively, so that each measurement is a linear projection of the spectrum, and the combined signal forms an autocorrelation trace described by

$$F(\tau) = \int S(\omega)(1 + \cos(\omega\tau))\, d\omega. \tag{1}$$

Conventionally, the spectrum $S(\omega)$ is then retrieved by an inverse Fourier transform of the fluences sampled on a uniform grid, as prescribed by the Nyquist-Shannon sampling theorem [7]. This prescription is so well established that it is rarely questioned—yet it is worth asking where it comes from, and whether it is the most one can do.

Measuring any continuous-domain quantity fundamentally constitutes a signal reconstruction problem that involves deciding how and where to collect data. The guiding principle for these decisions, at least in the one-dimensional ("time-valued") case, has long been the Nyquist-Shannon sampling theorem, which sets fundamental sampling constraints for sufficient reconstruction of bandlimited functions in signal processing or spectroscopy [7]. However, the emergence of sub-Nyquist sampling methods [8], most prominently compressed sensing [9], suggests that Nyquist-Shannon theory—while broadly applicable—is not the most general possible framework. Specifically, these newer approaches highlight that fewer samples can suffice if prior information about the signal structure is known or assumed [9].

We can formulate this as an adaptive, decision theoretic view on measurement: one in which the observer actively seeks to maximize information about the system under investigation. From this perspective, determining an *optimal* measurement strategy equates precisely to identifying actions that yield the greatest *information gain*: As each new datum arrives, an estimate of the signal is updated, and one can explicitly calculate how much information each possible measurement is expected to provide, assigning a quantifiable information gain to every potential measurement action.

As we discuss rigorously in the Supplemental Material [10], Nyquist sampling emerges in this framework not merely

*Contact author: a.doepp@lmu.de

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Open access publication funded by Max Planck Society.

as sufficient but as uniquely optimal when we have no knowledge about the signal under consideration apart from it being constrained by both finite observation duration T and bandwidth B. This dual constraint defines a finite informational "area"; picture a rectangle bounded horizontally by time and vertically by frequency. Within it, the Heisenberg-Gabor uncertainty principle rigorously defines minimal distinguishable information cells—Gabor's "logons" [46]—with smallest possible area 1/2. These cells represent independent pieces of information about the signal. Intuitively—accurate up to a small refinement discussed in the Supplemental Material—the total number of these distinguishable cells, the signal's degrees of freedom, is then $2BT$.

Crucially, the observer's initial state of maximal ignorance corresponds mathematically to assigning equal prior value to every information cell: Each is equally likely to contain significant signal energy. Under this condition, the rational measurement strategy should distribute samples such that each cell is measured exactly once, without creating information gaps (unmeasured cells) or redundancies (multiple measurements of the same cell). Translating this intuition rigorously into information-theoretic terms, we find that maximizing information gain reduces in this scenario to maximizing the determinant of the Fisher information matrix: a criterion recognized in experimental design theory as "D-optimality". Classical results from this theory show that reconstruction within a trigonometric model is optimized when samples are spaced evenly [47]. For our $2BT$ degrees of freedom, this then requires a spacing of $1/(2B)$: exactly the Nyquist rate. Hence, Nyquist sampling is not just a minimal criterion for signal reconstruction, but is the uniquely optimal information-acquisition strategy under conditions of maximal uncertainty. In turn, when prior knowledge becomes available, meaning the content of some information cells is partially known while others are suspected to contain more novel information, the optimal measurement spacing adjusts away from uniformity.

Indeed, one seldom measures entirely unknown samples, and thus, in virtually every use case *some* prior information on the spectrum is known—ranging from physical constraints such as the illumination profile and detector response to structural expectations about absorption peaks, bands, and overall smoothness. Conventionally, this knowledge is ignored at the acquisition stage and partially reintroduced *a posteriori* using heuristic data processing such as apodization, smoothing, baseline correction, or peak fitting [48,49].

Here, we invert this reasoning and instead adopt a sequential Bayesian inference scheme that uses available knowledge *ab initio* during the measurement, adaptively choosing the delays that yield the greatest information gain relative to the prior. While this requires more complex computations than a fast Fourier transformation, highly optimized hardware and algorithms for matrix operations—the same techniques powering modern AI systems—now make these more sophisticated mathematical workflows feasible in real time. This mirrors the historical context in which FTS itself emerged, becoming broadly adopted only in the 1970s once affordable computing power enabled rapid Fourier transforms [50]. In this sense, our proposed method—active Fourier spectroscopy (AFS)—represents a timely evolution of spectroscopic measurement, capitalizing on vastly expanded computational resources to implement more generally information-optimal measurement strategies.

I. ACTIVE FOURIER SPECTROSCOPY

Technically, AFS can be understood as a sequential Bayesian treatment of the FTS measurement. It works by first encoding our initial belief in the spectrum $S(\omega)$ as a probability distribution, characterized by a mean μ over the spectrum on a discrete frequency grid and a covariance Σ that quantifies the remaining uncertainty and the correlations across frequencies. Each measurement taken at a particular delay τ can be viewed as a linear functional of the spectral vector $S(\omega)$. Discretizing the frequency domain, the relationship between the spectral estimate μ_{prior} and the measured quantity F at delay τ can be expressed through a known linear response operator R that emerges naturally from Eq. (1). Given F and R, and assuming Gaussian measurement noise, the updated estimate of the spectrum after incorporating the new observation is

$$\begin{aligned} \mu_{\text{posterior}} &= \mu_{\text{prior}} + \gamma(F - R\mu_{\text{prior}}), \\ \Sigma_{\text{posterior}} &= (I - \gamma R)\Sigma_{\text{prior}}, \end{aligned} \tag{2}$$

where γ is an update weight (see the "Materials and Methods" section in the Supplemental Material [10]) that weighs prior belief against measurement noise. Rather than passively assembling the entire interferogram before inverting it, AFS continually refines the spectrum as each measurement arrives, merging existing information with new data in a principled manner.

This adaptive, incremental updating process can be guided by information theory to choose each measurement such that it maximizes the expected reduction in uncertainty [51,52]. Because the estimate explicitly tracks how uncertain it remains in each region of the spectrum, we can write the expected information gain from a single measurement in closed form as

$$\text{information gain} = \frac{1}{2} \ln \frac{|\Sigma_{\text{prior}}|}{|\Sigma_{\text{posterior}}|}. \tag{3}$$

This means that rather than taking measurements at fixed intervals, AFS can select delays that optimally reduce uncertainty. As discussed in the introduction, under total ignorance this objective reduces to uniform Nyquist sampling, which we consider as our baseline case; with prior information, the optimal delays redistribute toward regions where the prior is least constraining.

When prior knowledge is available, adaptive measurement faces a critical challenge: How to leverage this knowledge when it is correct, but avoid pathologies that arise when the model assumptions are wrong [53–55]. To address this, we evaluate the predictive ability of different candidate models—for example, individualized priors derived from different patient cohorts in a clinical study—by scoring each according to how well it predicted the running set of measurements $\mathcal{D} = \{F(\tau_1), F(\tau_2), \ldots\}$ before they were taken, a quantity known as the model evidence. This naturally penalizes overly complex models that must spread their probability mass over many possible outcomes. Models with higher

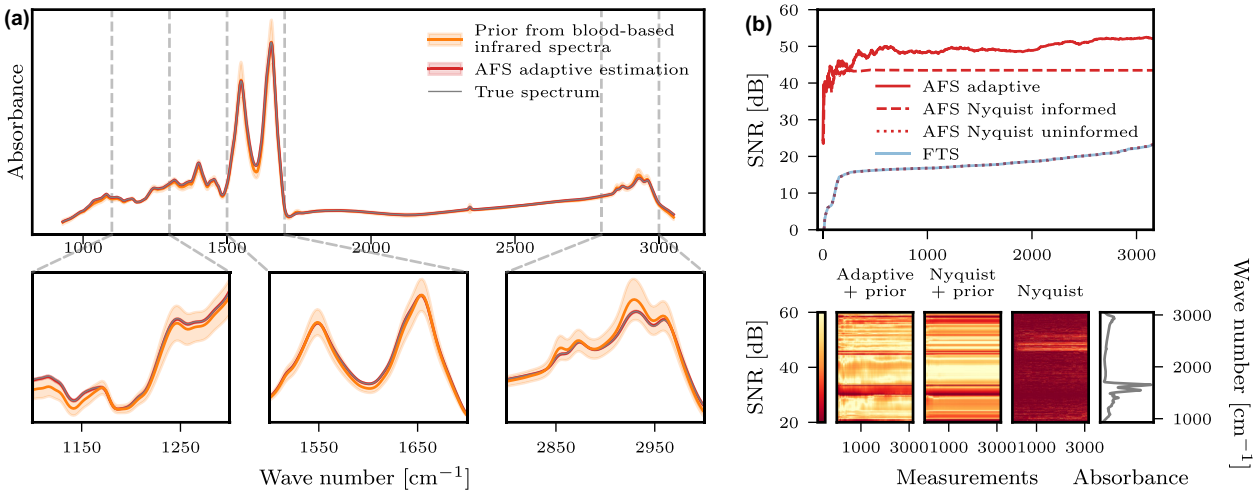


FIG. 1. AFS reconstruction of a biological FTIR absorbance spectrum. (a) The prior distribution (orange band) derived from a set of blood-based infrared spectra exhibits small variance, indicating strong constraints on possible spectra. AFS adaptive estimation (red) closely matches the true spectrum (gray) across the full wave number range. Enlarged panels highlight key spectral regions showing detailed agreement between reconstruction and ground truth. (b) SNR evolution with measurement count shows AFS adaptive sampling (solid red) reaching a high SNR fast and continuously refining the estimate, followed by AFS with prior knowledge and fixed Nyquist sampling (dashed red) lacking this refinement, while uninformed reconstruction (dotted) performs significantly worse. Wave-number-resolved SNR maps reveal sampling dynamics, with adaptive AFS exploring spectral regions most efficiently. The right panel shows the reconstructed absorbance spectrum for reference.

evidence are weighted more strongly, while models that perform poorly relative to the uninformed baseline trigger an automatic fallback to uniform sampling, returning to the Nyquist protocol. With several candidate priors, the expected information gain has no closed form; we use a weighted approximation that is exact for a single prior and is analyzed in Sec. S3.4.3 of the Supplemental Material [10]. This mechanism resolves the central tension of using prior information: It delivers substantial gains when the prior is meaningful, while falling back to the conventional FTS baseline when it is not, at the cost of the few measurements needed to accumulate evidence against it.

A distinctive feature of this formulation is that it introduces no free parameters describing the physical system itself. The effective dimension of the problem is fixed by the time-bandwidth product, and the basis in which the spectrum is represented is fixed by the interferometric measurement of Eq. (1); both follow from physical constraints rather than from a modeling decision. No structural model selection is required because this representation is complete for the measurement at hand—a completeness confirmed by the emergence of Nyquist sampling under maximal ignorance, which is a consequence of the formulation rather than an assumption fed into it. What remains is the choice of prior beliefs about the sample—assumptions about spectral smoothness or expected features, encoded through the covariance structure—and even these are adjudicated by evidence, allowing the data to determine appropriate regularization rather than requiring the practitioner to commit *a priori*.

Finally, the fully integrated analytical pipeline—from sequential updates and multimodel weighting to active (optimal) sampling—can be executed in real time on modern hardware. This analytical tractability is a consequence of formulating spectral inference as a linear-Gaussian model, the same model structure implicit in conventional FTS, of which Nyquist sampling is the maximally uninformed limit. AFS therefore does not introduce a new approximation but rather makes the existing one explicit and improvable. (This also situates the method within the broader language of optimal experimental design, although here the construction follows entirely from the physics of the measurement; see Sec. S2 of the Supplemental Material [10].) The Gaussian noise model is appropriate for detector-noise-dominated spectroscopy, the regime of our demonstrated applications (FTIR with thermal detectors, wavefront sensing, and imaging with charge-coupled device arrays). Extension to photon-limited regimes, where Poisson statistics dominate, requires Monte Carlo sampling for posterior updates and information gain estimation, as demonstrated for parametric spectral models in x-ray photoelectron spectroscopy [56] (see Sec. S3.5.2 of the Supplemental Material [10]).

II. MOLECULAR FINGERPRINTING

One of the most prominent applications of FTS is in the medical domain [5,57], with one area of increasing interest being metabolic fingerprinting for disease diagnosis [58]. To demonstrate our approach in this context, we use synthetic data modeled after real infrared measurements of human blood plasma in the framework of the Munich-based Lasers4Life (L4L) clinical study (see the Supplemental Material [10]) [59]. Based on these data, previous studies have demonstrated that molecular fingerprinting through vibrational signatures enables the noninvasive detection of cancer using blood plasma samples [60].

Figure 1 demonstrates AFS reconstruction of a biological FTIR absorbance spectrum. The prior distribution, derived from patient data, exhibits small variance in regions with

strong biochemical constraints, indicating substantial preexisting knowledge about possible spectral features. Building on this informative prior, the posterior estimation achieves even higher precision, evidenced by an uncertainty band too narrow to be visible in the plot. The reconstruction shows exceptional agreement with the true spectrum across the full wave number range (1000–3000 cm^{-1}), with precise reconstruction in key diagnostic regions highlighted in the enlarged panels.

The convergence behavior is analyzed in Fig. 1 through signal-to-noise ratio (SNR) calculations, defined as the ratio between signal power and reconstruction error power. AFS demonstrates superior performance in two configurations: Adaptive sampling achieves the highest SNR and fastest convergence, while informed Nyquist sampling also shows significant improvement over traditional methods. We observe that repeated measurements at specific positions can yield more information than the same amount of equally distributed measurements. Note that the uninformed Nyquist case performs equivalently to standard FTS, confirming our theoretical prediction of FTS as a limiting case of AFS.

These results demonstrate how AFS can enhance FTIR workflows by incorporating prior knowledge and adaptive sampling strategies while retaining traditional FTS as a fallback.

III. OPTICAL VORTICES

One core strength of Bayesian methods is their rigorous treatment of uncertainty. This is especially important for measurement techniques where the desired quantities are not observed directly but reconstructed from indirect or raw data. Such reconstructions risk compounding errors if uncertainties are not fully analyzed or accounted for, potentially misleading subsequent decisions. A complete uncertainty analysis should therefore also include correlations among reconstructed quantities, represented explicitly through a covariance matrix.

Space-time characterization of ultrafast laser pulses is an example for a complex system where direct measurements of all the relevant quantities are not feasible [61]. Of particular interest are optical vortices [62] with spectrally dependent orbital angular momentum [63]. This type of structured beam has applications in various fields, from relativistic light-matter interactions in the high-intensity regime [64] to quantum computing [65,66]. Wavelength-dependent distortions that affect beam propagation invoke techniques to spectrally resolve the phase structure. Specifically, the ultrafast metrology community has been using various incarnations of imaging FTS to spectrally resolve wavefronts [67–70]. But in this configuration there exists no principled propagation of correlated uncertainties from their respective spectral reconstructions to quantities calculated further down the analysis chain such as the spectrally resolved wavefront.

Accurate uncertainty quantification in spectral reconstructions is critically important because the resulting data serve as inputs for detailed physical simulations that underpin our fundamental understanding of these systems. Without uncertainty estimates, one cannot assure that the accurate parameter regime is modeled, while overestimated uncertainties would inflate the computational costs of already expensive simulations. A rigorous, balanced uncertainty treatment thus ensures that simulations are both reliable and efficient. We address this issue by employing AFS to spectrally resolve measurements from a wavefront sensor, which not only requires fewer samples (as demonstrated in the previous section), but allows for uncertainty quantification and propagation.

Our experimental setup combines scanning interferometry with Shack-Hartmann wavefront sensing. As described in the "Materials and Methods" section of the Supplemental Material [10], microlens images of a vortex and reference beam are recorded on a polarization-sensitive sensor at the information-optimal delays to calculate a spectrally resolved estimate, which was the starting point for phase reconstruction. Thus, this method constitutes a full vector beam measurement up to a spectral phase ambiguity. The AFS reconstruction yields a full spectral covariance matrix for each spatial point, capturing correlations between different wavelength channels. These correlations are crucial for accurate uncertainty propagation through the phase retrieval pipeline. The analytical propagation of the covariance through the centroid calculation and phase reconstruction is computationally efficient, requiring only matrix operations scaling with the number of spectral channels. This approach avoids the computational cost of sampling-based methods while maintaining higher accuracy than simplified diagonal-covariance approximations. The latter would treat spectral channels as independent, leading to overestimated uncertainties in derived quantities. Our method thus provides a complete characterization of structured light beams with rigorous uncertainty quantification at manageable computational cost—see Fig. 2 for a visualization of the results.

IV. HYPERSPECTRAL IMAGING

So far, we have considered use cases in which AFS enhanced (imaging) Fourier transform spectroscopy by incorporating prior knowledge and rigorous uncertainty propagation. Next, we illustrate a new capability emergent from our more general framework: the seamless fusion of multiple, partial information streams. The example we consider is hyperspectral imaging, a technology used for spatiospectral discrimination across diverse domains from medical imaging [71], food safety and agriculture [72,73], materials science [74] to remote sensing [75]. Common implementations require either extensive spatial or spectral scanning or complex geometric arrangements [76] and thus are subject to fundamental trade-offs of spatial and spectral resolution, measurement speed, and computational resources [77].

In our proof-of-concept demonstration, we realize a compact hyperspectral imager by combining coarse spectral information from color filters in an RGB camera with strategically chosen interferometric delay. Our setup uses a liquid crystal retarder as common-path interferometer immediately in front of the sensor. The two delayed components are superimposed on polarization-filtered pixels of an RGB camera to yield the autocorrelation measurements. For each delay step, the camera captures an image where different pixel groups—corresponding to the red, green, and blue filters—encode spectral information through a combination of their inherent filter transmission curves and interference patterns modulated according to Eq. (1). While the color filters alone

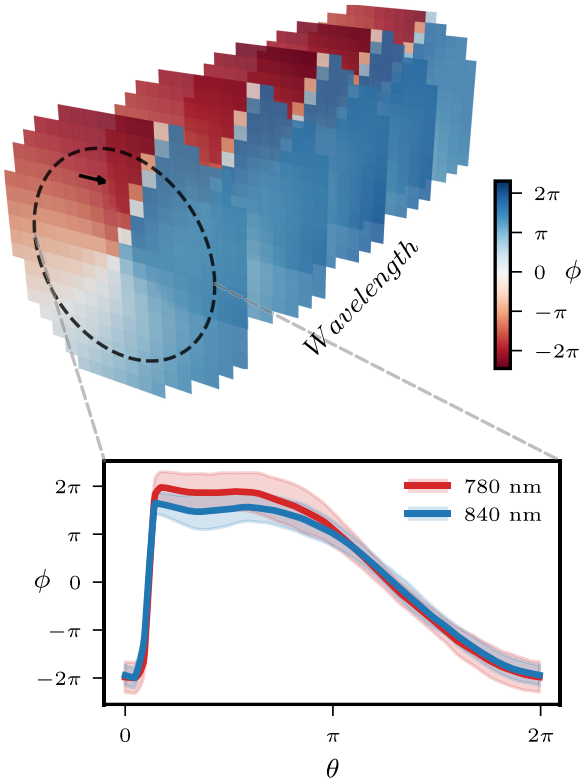


FIG. 2. Wavelength-resolved phase reconstruction of an optical vortex beam. Three-dimensional visualization of the beam's phase structure across wavelengths (750–850 nm), showing the characteristic azimuthal phase variation. The enlarged panel shows the azimuthal phase profiles extracted at 780 nm (design wavelength) and 840 nm with uncertainty bands derived from covariance propagation through the full AFS spectral reconstruction and phase retrieval chain. The observed wavelength dependence reflects the chromatic behavior of the spiral phase plate.

provide only a coarse spectral resolution, the additional time-dependent modulations encode finer spectral details. Building on the basic AFS inference algorithm [Eq. (2)], we then integrate the known spectral responses of the RGB filters with the measured interferometric autocorrelation curves. By iteratively updating the spectral estimate, AFS systematically combines the limited but complementary spectral information from the different color channels and delay positions. The algorithm selects optimal delay steps to maximize information gain, ensuring rapid and precise spectral reconstruction with minimal measurements.

As shown in Fig. 3, the reconstructed spectra rapidly converge to high fidelity within just a few strategically chosen delay measurements. Therefore, this approach effectively transforms the RGB camera into a compact hyperspectral imager, capable of resolving detailed spectral features beyond the native resolution of its color filters while maintaining full spatial resolution. It should be noted that the present measurements included minimal prior information. Contextual priors derived from scene information or metadata can effectively complement this, providing a fertile ground for developing fast and compact hyperspectral imagers.

V. CONCLUSIONS AND OUTLOOK

We have introduced a framework that operationalizes adaptive measurements among different variations of Fourier-based spectroscopy, based on theoretical work that connects classical sampling theory as a special case of Bayesian decision-making when operating with uninformed priors. This conceptual bridge enables substantially accelerated measurements when prior knowledge exists. Since conventional sampling is itself an admissible design, the information-optimal design never has a lower expected information gain, and the evidence-based fallback guards against priors not supported by the data—meaning that our approach can *upgrade* many existing setups to achieve both faster and more reliable, uncertainty quantified results.

Our demonstrations across three distinct domains of optical spectroscopy—clinical diagnostics, optical vector-field characterization, and hyperspectral imaging—reveal the method's

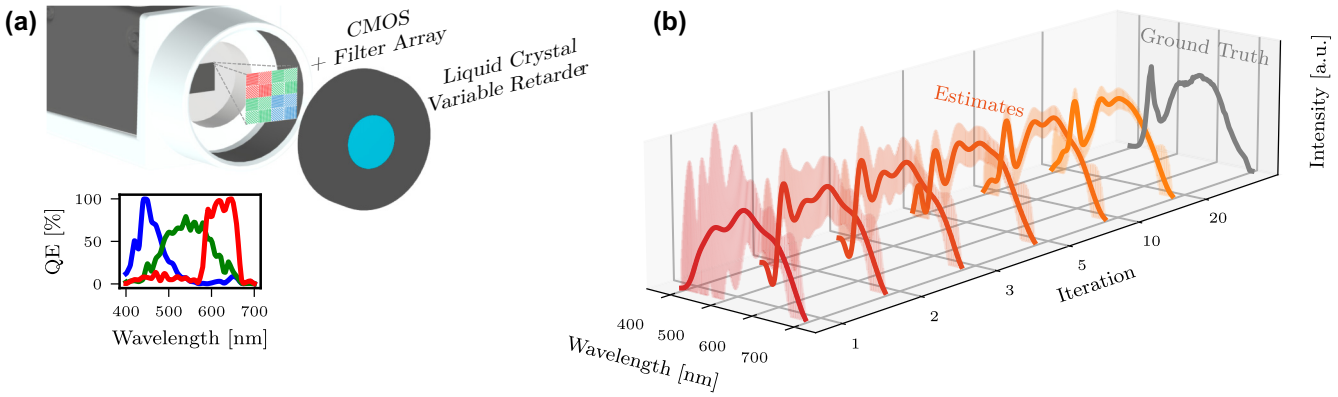


FIG. 3. Spectral resolution enhancement of an RGB camera using a liquid crystal retarder. (a) Device overview, consisting of a liquid crystal retarder for delay control between two polarization components and an RGB camera with a color and polarization filter pixel array, which constitute a compact common-path interferometer with three measurement channels. The inset shows the quantum efficiency of the color filters. (b) Evolution of the AFS estimate with uncertainty bands as a function of iteration, showing progressive refinement that already improves upon the three-channel resolution after only few measurements.

versatility and immediate practical impact. In each case, we observe accelerated measurements and can provide rigorous uncertainty bounds with fully propagated covariances, enabling confident decision-making in scenarios where conventional methods offer only point estimates or simplified error models. This is particularly valuable in data-driven science and medicine, which is increasingly relying on automated processing pipelines and decision systems. Here, our rigorous uncertainty quantification addresses a critical gap between raw measurements and evidence-based decision-making. We also demonstrated the unique potential for new diagnostics that merge information from different data streams, showcased by combining RGB sensors with autocorrelation measurements to realize hyperspectral sensors. This optimally combines the strengths of different measurement modalities into new hybrid measurement schemes.

Our work naturally extends in several directions: The current implementation employs multivariate Gaussian priors, but the mathematical framework supports more sophisticated constructions. Future implementations could incorporate Bayesian neural networks [78] to capture complex spectral relationships conditioned on relevant metadata, e.g., individualized or group-specific priors for precision health care. Similarly, more direct data streams could be integrated, e.g., combining many spectral channels with suitable delay measurements. Furthermore, depending on the application the optimization goal itself may be adapted from pure information gain to more objective-driven information metrics, such as maximum-value entropy search [79]. Beyond specific extensions, our framework exemplifies a fundamental shift in measurement science: from static protocols to dynamic, knowledge-driven strategies. This "agentic" approach recognizes that in many scenarios measurement is not about collecting predetermined data points but about extracting the most informative ones given our prior knowledge. This transparent, bottom-up approach starting from available knowledge complements deep learning approaches. As computational capabilities continue to advance, this approach paves the way for a new generation of scientific instruments that accumulate knowledge in a truly optimal way. In doing so, our work addresses a blind spot in the rapidly growing literature on autonomous experiments, self-driving laboratories, and agentic scientific instruments, such as Refs. [80–84]. Much of this recent work realizes autonomy in a top-down manner: A deep-learning model or large language model (LLM) agent is placed above the experiment and tasked to actively coordinate. Such systems are powerful, but their decisions depend on learned representations and training distributions whose latent biases generally cannot easily be audited or allow for disciplined continual operation. AFS realizes a different type of agency that is bottom-up and purely decision theoretic: The possible actions are fixed by the physics of Fourier spectroscopy, the belief state is an explicit posterior over the spectrum, and every measurement is chosen by a transparent information-gain objective. The resulting autonomy is therefore not achieved via an opaque intelligence acting on the instrument, but an interpretable statistical decision rule embedded in the measurement process itself. This distinction is central for the applications in which FTS is most valuable. In science, medicine, metrology, and diagnostics, faster acquisition alone is not sufficient if the resulting measurement record cannot be trusted, audited, or assigned meaningful uncertainty. AFS provides guarantees of a kind that compressed sensing, deep learning, and LLM-agent control do not generally provide. It therefore contributes an interpretable and reliability-oriented form of agentic measurement: one that actively exploits prior knowledge to improve outcomes, while maintaining the trustworthiness that makes FTS a standard tool in scientific and medical settings.

ACKNOWLEDGMENTS

We thank Prof. Ferenc Krausz for supporting this work. We thank the Federal Republic of Germany and the Free State of Bavaria for funding the CALA infrastructure (15171 E 0002) and its operation. J.S. and J.E. thank the International Max Planck Research School for Advanced Photon Science (IMPRS-APS) for support. 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. C.E. was funded by BMBF through the MACLIP project.

A.D. and J.S. developed the mathematical framework. J.S. performed and analyzed experiments with the help of S.H., C.E., and J.E. The L4L dataset was provided by N.L.-K. and K.V.K. A.D. and J.S. wrote the manuscript with input from all co-authors. A.D. supervised the work.

DATA AVAILABILITY

There are no publicly available research data or software supporting this manuscript. Requests for further information or data should be sent to the authors.

[1] P. R. Griffiths and J. A. de Haseth, *Fourier Transform Infrared Spectrometry*, 2nd ed (Wiley–Interscience, Hoboken, NJ, 2007).

[2] A. A. Enders, N. M. North, C. M. Fensore, J. Velez-Alvarez, and H. C. Allen, Functional group identification for FTIR spectra using image-based machine learning models, Anal. Chem. **93**, 9711 (2021).

[3] M. Dendisová, A. Jeništová, A. Parchaňská-Kokaislová, P. Matějka, V. Prokopec, and M. Švecová, The use of infrared spectroscopic techniques to characterize nanomaterials and nanostructures: A review, Anal. Chim. Acta **1031**, 1 (2018).

[4] S. O. Fakayode, G. A. Baker, D. K. Bwambok, N. Bhawawet, B. Elzey, N. Siraj, S. Macchi, D. A. Pollard, R. L. Perez, A. V. Duncan, and I. M. Warner, Molecular (Raman, NIR, and FTIR) spectroscopy and multivariate analysis in consumable products analysis, Appl. Spectrosc. Rev. **55**, 647 (2020).

[5] Z. Movasaghi, S. Rehman, and I. U. Rehman, Fourier transform infrared (FTIR) spectroscopy of biological tissues, Appl. Spectrosc. Rev. **43**, 134 (2008).

[6] C. M. Simonescu, Application of FTIR spectroscopy in environmental studies, in *Advanced Aspects of Spectroscopy*,

edited by M. A. Farrukh (IntechOpen, London, UK, 2012), pp. 49–84.
[7] C. Shannon, Communication in the presence of noise, Proc. IRE **37**, 10 (1949).
[8] M. Mishali and Y. C. Eldar, Sub-Nyquist sampling, IEEE Signal Process. Mag. **28**, 98 (2011).
[9] E. J. Candès, J. K. Romberg, and T. Tao, Stable signal recovery from incomplete and inaccurate measurements, Commun. Pure Appl. Math. **59**, 1207 (2006).
[10] See Supplemental Material at http://link.aps.org/supplemental/10.1103/shq1-j249 for additional information on the methods, related works, and detailed derivations, which includes Refs. [11–45].
[11] C. Beleites, U. Neugebauer, T. Bocklitz, C. Krafft, and J. Popp, Sample size planning for classification models, Anal. Chim. Acta **760**, 25 (2013).
[12] T. Eissa, K. V. Kepesidis, M. Zigman, and M. Huber, Limits and prospects of molecular fingerprinting for phenotyping biological systems revealed through *in silico* modeling, Anal. Chem. **95**, 6523 (2023).
[13] S. Schmuck and J. Svensson, Fourier spectroscopy: A Bayesian way, Int. J. Spectrosc. **2017**, 9265084 (2017).
[14] J. Kiefer, General equivalence theory for optimum designs (approximate theory), Ann. Stat. **2**, 849 (1974).
[15] M. Stone, Application of a measure of information to the design and comparison of regression experiments, Ann. Math. Stat. **30**, 55 (1959).
[16] J. M. Bernardo, Expected information as expected utility, Ann. Stat. **7**, 686 (1979).
[17] K. Chaloner and I. Verdinelli, Bayesian experimental design: A review, Stat. Sci. **10**, 273 (1995).
[18] T. Rainforth, A. Foster, D. R. Ivanova, and F. B. Smith, Modern Bayesian experimental design, Stat. Sci. **39**, 100 (2024).
[19] R. D. McMichael, S. Dushenko, and S. M. Blakley, Sequential Bayesian experiment design for adaptive Ramsey sequence measurements, J. Appl. Phys. **130**, 144401 (2021).
[20] C. E. Granade, C. Ferrie, N. Wiebe, and D. G. Cory, Robust online Hamiltonian learning, New J. Phys. **14**, 103013 (2012).
[21] S. Dushenko, K. Ambal, and R. D. McMichael, Sequential Bayesian experiment design for optically detected magnetic resonance of nitrogen-vacancy centers, Phys. Rev. Appl. **14**, 054036 (2020).
[22] M. Valeri, E. Polino, D. Poderini, I. Gianani, G. Corrielli, A. Crespi, R. Osellame, N. Spagnolo, and F. Sciarrino, Experimental adaptive Bayesian estimation of multiple phases with limited data, npj Quantum Inf. **6**, 92 (2020).
[23] M. Caouette-Mansour, A. Solyom, B. Ruffolo, R. D. McMichael, J. Sankey, and L. Childress, Robust spin relaxometry with fast adaptive Bayesian estimation, Phys. Rev. Appl. **17**, 064031 (2022).
[24] L. Sarra and F. Marquardt, Deep Bayesian experimental design for quantum many-body systems, Mach. Learn.: Sci. Technol. **4**, 045022 (2023).
[25] B. de Neeve, A. V. Lebedev, V. Negnevitsky, and J. P. Home, Time-adaptive phase estimation, Phys. Rev. Res. **7**, 023070 (2025).
[26] T. J. Loredo, Bayesian adaptive exploration, in *Bayesian Inference and Maximum Entropy Methods in Science and Engineering: 23rd International Workshop*, edited by G. J. Erickson and Y. Zhai (AIP, Melville, NY, 2004), Vol. 707, pp. 330–346.
[27] P. Paykari and A. H. Jaffe, Sparsely sampling the sky: A Bayesian experimental design approach, Mon. Not. R. Astron. Soc. **433**, 3523 (2013).
[28] N. Houlsby, F. Huszár, Z. Ghahramani, and M. Lengyel, Bayesian active learning for classification and preference learning, arXiv:1112.5745.
[29] D. G. Krige, A statistical approach to some basic mine valuation problems on the Witwatersrand, J. Chem. Metall. Min. Soc. S. Afr. **52**, 119 (1951).
[30] G. Matheron, Principles of geostatistics, Econ. Geol. **58**, 1246 (1963).
[31] H. J. Kushner, A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise, J. Basic Eng. **86**, 97 (1964).
[32] J. Močkus, On Bayesian methods for seeking the extremum, in *Optimization Techniques IFIP Technical Conference Novosibirsk, July 1–7, 1974*, edited by G. I. Marchuk (Springer, Berlin, 1975), pp. 400–404.
[33] J. González, Z. Dai, P. Hennig, and N. Lawrence, Batch Bayesian optimization via local penalization, in *Artificial Intelligence and Statistics* (PMLR, 2016), pp. 648–657.
[34] J. González, Z. Dai, A. Damianou, and N. D. Lawrence, Preferential Bayesian optimization, in *International Conference on Machine Learning* (PMLR, 2017), pp. 1282–1291.
[35] P. Hennig and C. J. Schuler, Entropy search for information-efficient global optimization, J. Mach. Learn. Res. **13**, 1809 (2012).
[36] Y. Ito, Y. Takeichi, H. Hino, and K. Ono, Optimal spectroscopic measurement design: Bayesian framework for rational data acquisition, Mach. Learn.: Sci. Technol. **6**, 025037 (2025).
[37] D. Slepian, On bandwidth, Proc. IEEE **64**, 292 (1976).
[38] I. Daubechies, Time-frequency localization operators: A geometric phase space approach, IEEE Trans. Inf. Theory **34**, 605 (1988).
[39] H. J. Landau and H. O. Pollak, Prolate spheroidal wave functions, Fourier analysis and uncertainty—III: The dimension of the space of essentially time- and band-limited signals, Bell Syst. Tech. J. **41**, 1295 (1962).
[40] C. E. Rasmussen, Gaussian processes in machine learning, in *Advanced Lectures on Machine Learning*, edited by O. Bousquet, U. von Luxburg, and G. Rätsch (Springer, Berlin, Heidelberg, Germany, 2003), pp. 63–71.
[41] L. Wang, Karhunen-Loeve expansions and their applications, Ph.D. thesis, London School of Economics and Political Science, UK, 2008.
[42] E. T. Jaynes, *Probability Theory: The Logic of Science* (Cambridge University Press, Cambridge, UK, 2003).
[43] J. Latz, Bayesian inverse problems are usually well-posed, SIAM Rev. **65**, 831 (2023).
[44] G. Casella and R. L. Berger, *Statistical Inference*, 2nd ed (Duxbury Press, Pacific Grove, CA, 2002).
[45] M. Burger, A. Hauptmann, T. Helin, N. Hyvönen, and J.-P. Puska, Sequentially optimized projections in x-ray imaging, Inverse Prob. **37**, 075006 (2021).
[46] D. Gabor, Theory of communication. Part 1: The analysis of information, J. Inst. Electr. Eng. Part 3 **93**, 429 (1946).
[47] F. Pukelsheim, *Optimal Design of Experiments* (Society for Industrial and Applied Mathematics, Philadelphia, PA, 2006).
[48] M. J. Baker, J. Trevisan, P. Bassan, R. Bhargava, H. J. Butler, K. M. Dorling, P. R. Fielden, S. W. Fogarty, N. J. Fullwood,

K. A. Heys, *et al.*, Using Fourier transform IR spectroscopy to analyze biological materials, Nat. Protoc. **9**, 1771 (2014).
[49] C. L. M. Morais, K. M. G. Lima, M. Singh, and F. L. Martin, Tutorial: Multivariate classification for vibrational spectroscopy in biological samples, Nat. Protoc. **15**, 2143 (2020).
[50] E. D. Becker and T. C. Farrar, Fourier transform spectroscopy: New methods dramatically improve the sensitivity of infrared and nuclear magnetic resonance spectroscopy, Science **178**, 361 (1972).
[51] D. V. Lindley, On a measure of the information provided by an experiment, Ann. Math. Stat. **27**, 986 (1956).
[52] V. V. Fedorov, *Theory of Optimal Experiments* (Academic Press, New York, 1972).
[53] S. J. Sloman, A. Bharti, J. Martinelli, and S. Kaski, Bayesian active learning in the presence of nuisance parameters, in *Proceedings of the Fortieth Conference on Uncertainty in Artificial Intelligence* (PMLR 2024), Vol. 244, pp. 3245–3263.
[54] S. J. Sloman, D. M. Oppenheimer, S. B. Broomell, and C. R. Shalizi, Characterizing the robustness of Bayesian adaptive experimental designs to active learning bias, arXiv:2205.13698.
[55] J. Go and T. Isaac, Robust expected information gain for optimal Bayesian experimental design using ambiguity sets, in *Proceedings of the Thirty-Eighth Conference on Uncertainty in Artificial Intelligence* (PMLR, 2022), Vol. 180, pp. 728–737.
[56] T. Nabika, K. Nagata, M. Mizumaki, S. Katakami, and M. Okada, Bayesian active learning with model selection for spectral experiments, Sci. Rep. **14**, 3680 (2024).
[57] A. C. S. Talari, M. A. G. Martinez, Z. Movasaghi, S. Rehman, and I. U. Rehman, Advances in Fourier transform infrared (FTIR) spectroscopy of biological tissues, Appl. Spectrosc. Rev. **52**, 456 (2017).
[58] D. I. Ellis and R. Goodacre, Metabolic fingerprinting in disease diagnosis: Biomedical applications of infrared and Raman spectroscopy, Analyst **131**, 875 (2006).
[59] M. Huber, K. V. Kepesidis, L. Voronina, F. Fleischmann, E. Fill, J. Hermann, I. Koch, K. Milger-Kneidinger, T. Kolben, G. B. Schulz, F. Jokisch, J. Behr, N. Harbeck, M. Reiser, C. Stief, F. Krausz, and M. Zigman, Infrared molecular fingerprinting of blood-based liquid biopsies for the detection of cancer, eLife **10**, e68758 (2021).
[60] K. V. Kepesidis, M.-G. Stoleriu, N. Feiler, L. Gigou, F. Fleischmann, J. Aschauer, S. Eiselen, I. Koch, N. Reinmuth, A. Tufman, J. Behr, and M. Žigman, Assessing lung cancer progression and survival with infrared spectroscopy of blood serum, BMC Med. **23**, 101 (2025).
[61] B. Alonso, A. Döpp, and S. W. Jolly, Space–time characterization of ultrashort laser pulses: A perspective, APL Photonics **9**, 070901 (2024).
[62] N. R. Heckenberg, R. McDuff, C. P. Smith, and A. G. White, Generation of optical phase singularities by computer-generated holograms, Opt. Lett. **17**, 221 (1992).
[63] A. M. Yao and M. J. Padgett, Orbital angular momentum: Origins, behavior and applications, Adv. Opt. Photon. **3**, 161 (2011).
[64] J. Vieira, R. M. G. M. Trines, E. P. Alves, R. A. Fonseca, J. T. Mendonça, R. Bingham, P. Norreys, and L. O. Silva, Amplification and generation of ultra-intense twisted laser pulses via stimulated Raman scattering, Nat. Commun. **7**, 10371 (2016).
[65] R. Fickler, R. Lapkiewicz, S. Ramelow, and A. Zeilinger, Quantum entanglement of complex photon polarization patterns in vector beams, Phys. Rev. A **89**, 060301 (2014).
[66] V. Parigi, V. D'Ambrosio, C. Arnold, L. Marrucci, F. Sciarrino, and J. Laurat, Storage and retrieval of vector beams of light in a multiple-degree-of-freedom quantum memory, Nat. Commun. **6**, 7706 (2015).
[67] G. Pariente, V. Gallet, A. Borot, O. Gobert, and F. Quéré, Space–time characterization of ultra-intense femtosecond laser beams, Nat. Photon. **10**, 547 (2016).
[68] M. Piccardo, M. de Oliveira, V. R. Policht, M. Russo, B. Ardini, M. Corti, G. Valentini, J. Vieira, C. Manzoni, G. Cerullo, and A. Ambrosio, Broadband control of topological–spectral correlations in space–time beams, Nat. Photon. **17**, 822 (2023).
[69] A. Jeandet, Spatio-temporal characterization of femtosecond laser pulses using self-referenced Fourier transform spectroscopy, Ph.D. thesis, Université Paris-Saclay, 2020.
[70] N. Weisse, J. Esslinger, S. Howard, F. M. Foerster, F. Haberstroh, L. Doyle, P. Norreys, J. Schreiber, S. Karsch, and A. Döpp, Measuring spatio-temporal couplings using modal spatio-spectral wavefront retrieval, Opt. Express **31**, 19733 (2023).
[71] G. Lu and B. Fei, Medical hyperspectral imaging: A review, J. Biomed. Opt. **19**, 010901 (2014).
[72] Y.-Z. Feng and D.-W. Sun, Application of hyperspectral imaging in food safety inspection and control: A review, Crit. Rev. Food Sci. Nutr. **52**, 1039 (2012).
[73] L. M. Dale, A. Thewis, C. Boudry, I. Rotar, P. Dardenne, V. Baeten, and J. A. F. Pierna, Hyperspectral imaging applications in agriculture and agro-food product quality and safety control: A review, Appl. Spectrosc. Rev. **48**, 142 (2013).
[74] X. Dong, M. Jakobi, S. Wang, M. H. Köhler, X. Zhang, and A. W. Koch, A review of hyperspectral imaging for nanoscale materials research, Appl. Spectrosc. Rev. **54**, 285 (2019).
[75] A. F. H. Goetz, Three decades of hyperspectral remote sensing of the Earth: A personal view, Remote Sens. Environ. **113**, S5 (2009).
[76] N. A. Hagen and M. W. Kudenov, Review of snapshot spectral imaging technologies, Opt. Eng. **52**, 090901 (2013).
[77] G. R. Arce, D. J. Brady, L. Carin, H. Arguello, and D. S. Kittle, Compressive coded aperture spectral imaging: An introduction, IEEE Signal Process. Mag. **31**, 105 (2014).
[78] L. V. Jospin, H. Laga, F. Boussaid, W. Buntine, and M. Bennamoun, Hands-on Bayesian neural networks—A tutorial for deep learning users, IEEE Comput. Intell. Mag. **17**, 29 (2022).
[79] Z. Wang and S. Jegelka, Max-value entropy search for efficient Bayesian optimization, in *International Conference on Machine Learning* (PMLR, 2017), pp. 3627–3635.
[80] A. V. Tobias and A. Wahab, Autonomous 'self-driving' laboratories: A review of technology and policy implications, R. Soc. Open Sci. **12**, 250646 (2025).
[81] Z. Wang, M. Han, W. Cheng, Y.-K. Huang, P. N. Ji, D. Wu, M. Safari, F. Holtorf, K. Alqubaisi, D. Zhuo, Y. Chen, T. Wang, N. M. Linke, D. Englund, and T. Chen, Agentic AI for scalable

and robust optical systems control, IEEE Access **14**, 146757 (2026).

[82] S. Yang, F. Chen, R. Zhao, J. Wu, Y. Wang, H. Luo, N. Han, Q. Chen, Y. Hu, W. Li, M. Li, H. Chen, and Y. Yang, End-to-end autonomous scientific discovery on a real optical platform, arXiv:2604.27092.

[83] S. Z. Uddin, S. Vaidya, S. Choudhary, Z. Chen, R. K. Salib, L. Huang, D. R. Englund, and M. Soljačić, AI-driven robotics for optics, Sci. Adv. **12**, eaee1381 (2026).

[84] R. Sha, B. Wang, J. Yang, X. Ma, C. Wu, L. Yan, C. Zhou, J. Liu, G. Wang, S. Yan, and L. Zhu, LLM-based multi-agent copilot for quantum sensor, arXiv:2508.05421.