

Efficient ultrafast homodyne detection of quantum light

Corresponding Author: Professor Young-Sik Ra

This file contains all reviewer reports in order by version, followed by all author rebuttals in order by version.

Version 0:

Reviewer comments:

Reviewer #1

(Remarks to the Author)

Summary

The manuscript formulates quadrature extraction in pulsed homodyne detection as the choice of a temporal weight $w(t)$, shows that maximizing the SNR is a generalized Rayleigh quotient solved by $Sw = \gamma Ew$, and recovers the area-integration and peak-value methods as special weights. The optimal weight, obtained from the measured shot-noise and electronic-noise matrices, increases the SNR and improves the observed squeezing/anti-squeezing of an ultrafast squeezed vacuum, in agreement with a loss model. The work appears technically sound, carefully checked, and useful as a practical contribution to pulsed homodyne detection. However, the core method is mathematically an established result of classical optimal detection, and the directly demonstrated improvement in the quantum-optical measurement is modest. I am therefore not convinced that the work reaches the level of novelty and broad importance expected by Communications Physics. In my view, the manuscript would be better suited to a more specialized journal after appropriate reframing. The points below are offered as constructive suggestions to improve the manuscript wherever it is submitted.

Major

1. Relation to classical optimal filtering and the novelty. In the white-noise limit the optimal weight becomes $w \propto r(t)$, i.e. the matched filter, with E providing the noise whitening in the colored-noise case. This is the same generalized Rayleigh quotient / generalized eigenvalue problem that is standard in radar and communication, gravitational-wave template matching, and signal processing (Fisher's discriminant, MVDR beamforming). The manuscript cites only quantum parameter-estimation works. I suggest citing the classical optimal-detection literature and stating clearly that the contribution is not the method itself but its application to the pulsed ultrafast CV regime, together with the unification of the area-integration and peak-value methods. For the classical literature, see for example D. O. North, Proc. IEEE 51, 1016 (1963) [reprint of the 1943 RCA report]; G. L. Turin, IRE Trans. Inf. Theory 6, 311 (1960); and B. Allen et al., Phys. Rev. D 85, 122006 (2012) (FINDCHIRP).

Minor: motivation and framing

2. Motivation (saturation) vs. method. The Introduction motivates the problem by detector saturation limiting the LO power, but the method is post-processing at a fixed, sub-saturation operating point (consistent with the linear scaling in Fig. 6(e)). It would be clearer to state that the method improves the SNR by noise rejection at a fixed operating point, separate from the saturation issue.

3. "Broadening the detector response to reduce the peak height" (Introduction) is ambiguous. Reducing the peak requires slowing the temporal response, i.e. narrowing the bandwidth; widening the bandwidth would sharpen the response and raise the peak. Please clarify that the slowing of the temporal response is meant, and link it to the stated drawback (worse pulse discrimination / larger crosstalk), citing Ref. [31].

Minor: timescale assumptions

4. Pulse duration $\ll$ detector response time (rank-one approximation). Eq. (6) assumes $r(t)$ varies slowly compared with $g(t)$, making R rank-one (Eq. 15). This holds well here (75 fs vs ~ns), and the dominant leading eigenvalue of R in Fig. 5(f) confirms it. It would help to state this assumption explicitly in the main text and cite Fig. 5(f).

5. Detector response $\ll$ repetition period T (bandwidth $\gg$ repetition rate). This is the condition for the inter-pulse crosstalk $\bar{r}_{[j,k]}$ ($j \neq k$) to be negligible (Eq. 8, 15). Fig. 4 shows it holds (~200 MHz vs 80 MHz), but it is not stated as an assumption. Please state it and cite Fig. 4. This also clarifies the Introduction's remark that slowing the detector degrades pulse

discrimination — that is precisely where this condition fails.

Minor: interpretation and presentation

6. Quantify the required efficiency in the Introduction. The SNR maps onto an effective efficiency via $\eta_e = 1 - 1/\text{SNR}$ (Eq. 10), but the required level is not stated. Loss tolerance differs greatly between fields (a 50% loss is acceptable in classical communication but fatal for strong-squeezing CV experiments). I suggest adding the SNR–loss correspondence (clearance $\sim 13/15/18/20$ dB $\leftrightarrow$ loss $\sim 5.0/3.2/1.6/1.0\%$) and the target efficiencies (e.g. the near-unity efficiency for strong squeezing, Ref. [38]). It is worth stating that, in quantum information processing, even a few percent of efficiency can decide whether the nonclassical resource remains usable — which also explains why the method matters most at high η_0 and large r (Fig. 2).

7. Comparable pulse/response timescales (optional). With broadband homodyne detectors now reaching the GHz range, the regime where the pulse duration and the response time become comparable is increasingly realistic. There R becomes multimode, and a trade-off appears between SNR maximization and temporal-mode fidelity (with non-negligible crosstalk). Since Eqs. (11)–(13) are model-independent, the method still applies; a sentence on how to proceed (e.g. keeping several leading modes) would clarify the range of validity. No calculation is needed.

8. dB vs. percentage; measured vs. projected. The “+135% / +44%” in the Conclusion correspond to 3.7 dB and 1.6 dB; showing the dB values would be more consistent. Also, the directly measured squeezing improvement (Fig. 7, $-2.45 \rightarrow -2.74$ dB, ~ 0.3 dB) and the larger Fig. 2 gains (up to 2.46 dB) are at different operating points — Fig. 7 at a measured SNR of $9.68 \rightarrow 13.10$ dB with the actual r and η_0 , the Fig. 2 values as loss-model projections at higher SNR, larger r , and higher η_0 . Please state this difference so the two are not conflated.

9. Title. “Ultrafast homodyne detection” is ambiguous, as it can also denote a CW-LO broadband-detector scheme. Since this work uses a pulsed LO and pulse-by-pulse measurement, making this explicit in the title (e.g. “pulsed” or “pulse-by-pulse”) would help.

Reviewer #2

(Remarks to the Author)

Continuous-variable (CV)-based quantum information protocols harness the rich modal structure of quantum light towards the development of scalable quantum technologies. In particular, the spectral correlations of ultrafast light sources can be used to generate large, entangled states. However, CV quantum states are vulnerable to decoherence, making their implementation in practical quantum information protocols challenging. This sensitivity to loss is one argument against discrete-variable based approach, i.e., using the photon number or the polarization of light. Nevertheless, CV encoded states, e.g., squeezed states and non-gaussian states, with their rich modal structure, are promising candidates for fast quantum communication and metrology.

In this manuscript, the authors solve the generalized Rayleigh quotient coefficient to improve the signal-to-noise ratio (SNR) of homodyne detectors (HD) by applying optimized temporal weights to quadrature measurements. The paper includes a detailed theoretical background on homodyne detection, temporal correlations in homodyne measurements and implementation of the optimization protocol. More details on the homodyne detector and the development of the optimization protocol are included in the Supplementary Materials. The theory is supported throughout by the appropriate references, to which the results are compared. The theory is followed by a thorough description of the experimental setup, measurement protocol and results. The results show a clear benefit of the proposed optimization method in extracting quadrature measurements.

Therefore, I think that research presented in this article shows a concrete benefit to the development of CV-based quantum information protocols, and that Efficient ultrafast homodyne detection of quantum light should definitely be accepted for publication in Communications Physics.

Below, I have added a few questions and remarks:

Section III – Page 3: Could you clarify what are sampling errors in the estimation of the shot and electronic noise?

Section III – Page 4:

- The theory shows that the optimization performance is sensitive to increasing frequency; does it lead to a restriction of the bandwidth over which a temporal-correlation-improved homodyne detector can be designed or a restriction on the laser than can be used (i.e., repetition rate)?

- What do the different η_0 values correspond to? A different HD design or modified mode-matching?

Section IV – Page 5: Although the data collection protocol is thorough, can the authors clarify what the definition of an “experimental” pulse is, the relation between consecutive experimental pulses compared to the theoretical definition, and what is the relation between the choice of T and f_c , if any.

Conclusion – Page 7: potentially “detector signals” instead of “detected” signals.

Reviewer #3

(Remarks to the Author)

I co-reviewed this manuscript with one of the reviewers who provided the listed reports. This is part of the Communications Physics initiative to facilitate training in peer review and to provide appropriate recognition for Early Career Researchers who co-review manuscripts.

Reviewer #4

(Remarks to the Author)

This manuscript presents a method to improve ultrafast homodyne detection by optimizing the temporal weight used to extract pulse-resolved quadrature outcomes. The optimization is formulated as a generalized Rayleigh quotient problem based on experimentally estimated autocorrelation matrices of the total detector output and the electronic noise. The authors then demonstrate experimentally that the optimized temporal weight increases the homodyne signal-to-noise ratio (SNR) relative to conventional constant and peak-shaped weights, and that this improvement leads to enhanced observed squeezing and anti-squeezing levels.

Overall, I find the manuscript technically sound, clearly written, and potentially interesting for the community working on ultrafast continuous-variable measurements and pulsed homodyne detection. The central idea is physically sensible: the useful homodyne signal and the electronic noise have different temporal covariance structures, and this difference can be exploited to define a more efficient temporal readout mode. The theoretical formulation is coherent, and the experimental demonstration is relevant.

However, in its present form, I think that several aspects require clarification before the conclusions can be considered fully convincing. My main concerns are listed below.

1. The central statement of the paper is that “the pulsed detector signals in ultrafast homodyne detection exhibit strong temporal correlations that are absent in the electronic noise.” This point is essential for the whole manuscript, but its physical origin is not explained clearly enough.

As I understand it, these correlations should not be interpreted as correlations between successive optical pulses. Rather, they arise from the finite impulse response of the homodyne detector: each ultrafast optical pulse is much shorter than the electronic response time of the detector and is therefore converted into an extended voltage waveform.

If correct, this physical interpretation should be stated explicitly, preferably already in the Introduction, since it is the basic reason why temporal weighting can improve the SNR. The present wording may be misleading because the electronic noise is not necessarily completely free of temporal correlations: finite electronic bandwidth generally produces stationary autocorrelations. What seems to be absent in the electronic noise is instead the pulse-synchronous, localized, approximately rank-one covariance structure associated with the optical homodyne signal.

2. Related to the previous point, the manuscript should more clearly distinguish between intra-pulse temporal correlations and correlations between different pulses. In Fig. 6(d), the authors report a residual correlation of approximately 6% at $d=1$ for the constant-weight strategy, while the peak-shaped and optimized weights show negligible correlations. This is an important experimental observation, but its origin is not sufficiently discussed.

Is this residual correlation due to a small overlap of the detector impulse responses from adjacent pulses, baseline recovery, electronic ringing, finite-bandwidth effects, or the use of a nonoptimal integration window? Since the method is partly motivated by the need to avoid correlations between pulse-resolved quadrature samples, the authors should explicitly discuss the mechanism behind the observed $d=1$ correlation and explain why the optimized weight suppresses it.

3. Although I understand that the main focus of the manuscript is the optimized homodyne detection method, the description of the squeezed-state generation in the experimental section is too minimal. Since the final experimental demonstration relies on measuring squeezed vacuum generated by a SPOPO, the reader needs more information to assess the quality and reproducibility of this part of the experiment. In particular, the authors should provide additional experimental details on:

- the SPOPO cavity design and operating point;
- the output coupler reflectivity and approximate escape efficiency;
- the cavity linewidth or finesse;
- the pump power used for the squeezing measurements and its value relative to threshold;
- the measured homodyne visibility;
- the optical losses between the SPOPO and the homodyne detector;
- the photodiode quantum efficiency already quoted in the text, but included in a full efficiency budget;
- the total detection efficiency used to interpret the measured squeezing levels;
- the estimated contribution of electronic noise to the effective homodyne efficiency.

At present, Fig. 7 reports measured squeezing and anti-squeezing levels, and the authors infer consistency with the expected efficiency improvement. However, without a clear efficiency budget, it is difficult to judge whether the observed squeezing levels are limited mainly by the SNR, optical losses, mode matching, phase noise, or properties of the SPOPO itself.

4. The experimental scheme in Fig. 3 and the corresponding description leave some ambiguity regarding the weak seed beam used to lock the relative phase between the squeezed state and the local oscillator. From the text, it appears that a weak coherent seed beam is injected into the SPOPO and that the relative phase is locked by monitoring a tapped output power at a photodiode. The manuscript should clarify more precisely the propagation direction, optical path, and role of this

seed beam. In particular:

- Does the seed propagate in the same direction as the squeezed vacuum?
- Does any residual seed field reach the homodyne detector?
- If so, how is its contribution removed or made negligible?
- Does the seed introduce a coherent displacement of the squeezed state, or is the measured state effectively an undisplaced squeezed vacuum?
- What exact error signal is used for phase stabilization?
- Is the phase lock performed on resonance or off resonance?
- How is degradation of squeezing due to residual seed leakage, phase noise, or cavity detuning avoided?

5. The paragraph describing the timing correction is not sufficiently clear. The authors state that the oscilloscope sampling rate is not an exact integer multiple of the pulse repetition rate, so that the measurement window drifts relative to the pulse train. They then mitigate this effect using a polyphase finite-impulse-response resampling algorithm, interpolating the 125 samples in each period to 1000 samples, compensating the drift, and downsampling back to 125 samples.

Why was the acquisition not directly synchronized or triggered from the laser oscillator, the LO, or the weak seed?

6. Figures 5(g), 5(h), and 5(i) should be improved visually. The caption states that the three leading eigenvectors are shown, but the meaning of the three colors is not immediately clear from the figure itself. The authors should either add a legend directly in the panels or label the curves more clearly. The same applies to the normalization and sign convention of the eigenvectors, since eigenvectors are defined only up to an arbitrary sign.

7. I suggest inserting the following references.

a. A recent and directly relevant reference on time-resolved balanced homodyne detection for high-repetition-rate pulsed sources is Altiglia et al., “A Simple and Robust Balanced Homodyne Detector for High-Repetition-Rate Pulsed Sources”.
<https://arxiv.org/abs/2604.06994>

b. The recent experimental paper “Generation and detection of squeezed states via a synchronously pumped optical parametric oscillator” (Applied Physics B, 2026) would be a particularly relevant addition in the paragraph introducing SPOPO-based pulsed squeezing.

<https://link.springer.com/article/10.1007/s00340-025-08617-6>

I recommend publication of this paper after major revisions.

Version 1:

Reviewer comments:

Reviewer #1

(Remarks to the Author)

I thank the authors for the careful revision. The added statements of the underlying assumptions, the quantitative SNR–efficiency correspondence in the Introduction, the dB values in the Conclusion, and the clarification of the different operating points in Figs. 2 and 7 address my earlier comments well. The new Methods section, added in response to another referee, also makes the efficiency budget easy to follow. I now support publication.

I would like the authors to check one remaining point before publication. It concerns internal consistency of the revised text.

After Eq. (16), the revised manuscript states: “In realistic detectors, however, the rank-one kernel model of Eq. (15) is generally insufficient because imperfections such as timing jitter and detector-response shape variations increase the rank of the detector response matrix R .” In the discussion of Fig. 5(f), however, the manuscript states that the dominance of the leading eigenvalue “is consistent with the ideal detector response model in Eq. (15), indicating that contributions from nonideal detector responses are negligible”, and that the contribution of the 1st eigenvector associated with timing jitter is negligible. These two statements do not agree with each other. For the detector used in this work, R is effectively rank-one, and the optimal weight is then the whitened matched filter for the measured colored electronic noise.

I do not think this affects the main claim of the paper. As I understand it, the contribution is the formulation of pulsed ultrafast homodyne detection that leads to the generalized eigenvalue problem, the unification of the area-integration and peak-value methods within it, the high-frequency truncation procedure, and the experimental verification. None of these depends on R having rank higher than one.

I therefore suggest to revise the sentence after Eq. (16) as follows: state that for the present detector the rank-one condition is well satisfied, so that the optimal weight coincides with the classical whitened matched filter [47–49], and that the generalized Rayleigh quotient formulation remains valid when the rank-one condition fails, for example for detectors with larger timing jitter or in the regime where the pulse duration and the detector response time become comparable. In my view this is a stronger position than the present one: the agreement with the classical result then works as a consistency check of the method, instead of a point the authors need to argue against.

Reviewer #2

(Remarks to the Author)

In response to the initial review, the authors have substantially revised their manuscript. They thoroughly addressed the referees' comments and questions, and updated their manuscript accordingly. They expanded the measurement-protocol description to add key definitions and clarifications. They also added a detailed description of their experimental setup and data-processing methodology. These additions significantly improve the manuscript's clarity; therefore, I recommend it for publication.

Reviewer #4

(Remarks to the Author)

The authors have adequately addressed all the concerns raised in my previous report. The revisions have significantly improved the clarity and completeness of the manuscript, and I am satisfied with the authors' responses. I therefore recommend the manuscript for publication.

Open Access This Peer Review File is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

In cases where reviewers are anonymous, credit should be given to 'Anonymous Referee' and the source.

The images or other third party material in this Peer Review File are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

To view a copy of this license, visit <https://creativecommons.org/licenses/by/4.0/>

Reviewer #1:

The manuscript formulates quadrature extraction in pulsed homodyne detection as the choice of a temporal weight $w(t)$, shows that maximizing the SNR is a generalized Rayleigh quotient solved by $Sw = \gamma Ew$, and recovers the area-integration and peak-value methods as special weights. The optimal weight, obtained from the measured shot-noise and electronic-noise matrices, increases the SNR and improves the observed squeezing/anti-squeezing of an ultrafast squeezed vacuum, in agreement with a loss model. The work appears technically sound, carefully checked, and useful as a practical contribution to pulsed homodyne detection. However, the core method is mathematically an established result of classical optimal detection, and the directly demonstrated improvement in the quantum-optical measurement is modest. I am therefore not convinced that the work reaches the level of novelty and broad importance expected by Communications Physics. In my view, the manuscript would be better suited to a more specialized journal after appropriate reframing. The points below are offered as constructive suggestions to improve the manuscript wherever it is submitted.

Reply: We thank the reviewer for the constructive and thoughtful comments. Below, we address the reviewer's comments point by point.

1. Relation to classical optimal filtering and the novelty. In the white-noise limit the optimal weight becomes $w \propto r(t)$, i.e. the matched filter, with E providing the noise whitening in the colored-noise case. This is the same generalized Rayleigh quotient / generalized eigenvalue problem that is standard in radar and communication, gravitational-wave template matching, and signal processing (Fisher's discriminant, MVDR beamforming). The manuscript cites only quantum parameter-estimation works. I suggest citing the classical optimal-detection literature and stating clearly that the contribution is not the method itself but its application to the pulsed ultrafast CV regime, together with the unification of the area-integration and peak-value methods. For the classical literature, see for example D. O. North, Proc. IEEE 51, 1016 (1963) [reprint of the 1943 RCA report]; G. L. Turin, IRE Trans. Inf. Theory 6, 311 (1960); and B. Allen et al., Phys. Rev. D 85, 122006 (2012) (FINDCHIRP).

Reply: Mathematically, the matched-filter problem is a special case of the generalized Rayleigh quotient problem considered in our work, when the detector response matrix ($\mathbf{R}$) has rank one. However, this rank-one condition does not necessarily hold for realistic homodyne detectors whose responses exhibit imperfections such as timing jitter or pulse-shape variations. We have revised the manuscript to clarify this connection in the discussion of Eq. (15) and Eq. (16). We have also cited the three suggested papers.

Importantly, the main contribution of our work is not just the optimization algorithm itself, but the formulation of ultrafast homodyne detection that leads naturally to the generalized Rayleigh-quotient problem. Despite extensive studies of ultrafast homodyne detection (Refs. 30, 31, and 33-38 in the revised manuscript), this formulation has not, to our knowledge, been recognized in the literature. Our work provides a unified framework for understanding and extending ultrafast homodyne detection methods.

First, we present a systematic formulation of the quadrature-measurement process in pulsed ultrafast homodyne detection. Within this framework, conventional methods, including area-integration and peak-value detection, are identified as particular choices of the weighting

function applied to the detector signal within a single measurement window. By expressing the measurement process in this form, we show that optimizing the weighting function for increasing the clearance can be reduced to a generalized eigenvalue problem.

Second, we present a practical procedure for implementing this optimization in an ultrafast homodyne experiment. To mitigate the influence of sampling errors in the experimentally estimated matrices, particularly in the high-frequency regime, we introduce a high-frequency truncation method that enables stable estimation of the optimal weighting function.

Finally, we experimentally validate the optimization method. We demonstrate that the quadrature outcomes obtained using the optimized weighting function retain the essential characteristics of shot-noise-limited quadrature measurements and reproduce these characteristics more faithfully than those obtained using conventional methods: a Gaussian quadrature distribution, negligible pulse-to-pulse correlations, and linear scaling of the shot-noise variance with local-oscillator power. We further observe an enhancement in the measured squeezing level when the optimized weighting function is used.

2. Motivation (saturation) vs. method. The Introduction motivates the problem by detector saturation limiting the LO power, but the method is post-processing at a fixed, sub-saturation operating point (consistent with the linear scaling in Fig. 6(e)). It would be clearer to state that the method improves the SNR by noise rejection at a fixed operating point, separate from the saturation issue.

Reply: As the reviewer correctly noted, our method uses the same LO power and does not alter the detector saturation power. The SNR enhancement arises from using the optimized weighting function instead of using the conventional area-integration or peak-value detection methods. For clarity, we have revised the sentence in the Introduction, “This optimization leads to a substantial increase in the SNR,” to “This optimization leads to a substantial increase in the SNR at the same LO power.” The preceding sentence, “An optimized temporal weight can efficiently extract quadrature outcomes from the detector signals while suppressing the contribution of electronic noise,” clarifies the origin of this SNR enhancement.

3. “Broadening the detector response to reduce the peak height” (Introduction) is ambiguous. Reducing the peak requires slowing the temporal response, i.e. narrowing the bandwidth; widening the bandwidth would sharpen the response and raise the peak. Please clarify that the slowing of the temporal response is meant, and link it to the stated drawback (worse pulse discrimination / larger crosstalk), citing Ref. [31].

Reply: Thank you for the comment. Here, “broadening” refers to broadening of the detector response in the time domain, rather than to broadening of the spectral bandwidth. We have revised the introduction to clarify the point.

4 and 5. Pulse duration $\ll$ detector response time (rank-one approximation). Eq. (6) assumes $r(t)$ varies slowly compared with $g(t)$, making R rank-one (Eq. 15). This holds well here (75 fs vs $\sim$ ns), and the dominant leading eigenvalue of R in Fig. 5(f) confirms it. It would help to state

this assumption explicitly in the main text and cite Fig. 5(f). Detector response $\ll$ repetition period T (bandwidth $\gg$ repetition rate). This is the condition for the inter-pulse crosstalk $\bar{r}_{\{j,k\}}$ ($j \neq k$) to be negligible (Eq. 8, 15). Fig. 4 shows it holds (~ 200 MHz vs 80 MHz), but it is not stated as an assumption. Please state it and cite Fig. 4. This also clarifies the Introduction's remark that slowing the detector degrades pulse discrimination — that is precisely where this condition fails.

Reply: As the reviewer suggested, these time-scale comparisons are important for supporting the assumptions underlying Eqs. (6,8,15). We have revised the manuscript to state explicitly that both assumptions are satisfied under our experimental conditions.

6. Quantify the required efficiency in the Introduction. The SNR maps onto an effective efficiency via $\eta_e = 1 - 1/\text{SNR}$ (Eq. 10), but the required level is not stated. Loss tolerance differs greatly between fields (a 50% loss is acceptable in classical communication but fatal for strong-squeezing CV experiments). I suggest adding the SNR–loss correspondence (clearance $\sim 13/15/18/20$ dB $\leftrightarrow$ loss $\sim 5.0/3.2/1.6/1.0\%$) and the target efficiencies (e.g. the near-unity efficiency for strong squeezing, Ref. [38]). It is worth stating that, in quantum information processing, even a few percent of efficiency can decide whether the nonclassical resource remains usable — which also explains why the method matters most at high η_0 and large r (Fig. 2).

Reply: We have revised the Introduction to include practical applications that require high detection efficiency. These examples include quantum error correction with GKP states [Phys. Rev. A 107, 052414 (2023)], loophole-free Bell tests [Phys. Rev. Lett. 93, 130409 (2004)], and high-quality quantum-enhanced measurements [Phys. Rev. Lett. 117, 110801 (2016)]. Moreover, for quantitative analysis, we have added the detection efficiencies corresponding to the clearance levels discussed in Figs. 2 and 6.

7. Comparable pulse/response timescales (optional). With broadband homodyne detectors now reaching the GHz range, the regime where the pulse duration and the response time become comparable is increasingly realistic. There R becomes multimode, and a trade-off appears between SNR maximization and temporal-mode fidelity (with non-negligible crosstalk). Since Eqs. (11)–(13) are model-independent, the method still applies; a sentence on how to proceed (e.g. keeping several leading modes) would clarify the range of validity. No calculation is needed.

Reply: We thank the reviewer for this optional suggestion regarding a possible extension of the present work. Our homodyne-detection model assumes that the optical pulse duration is much shorter than the detector response time, or equivalently, that the optical bandwidth is much broader than the detector bandwidth. This assumption remains well satisfied even for GHz-bandwidth detectors used with femtosecond optical pulses, whose bandwidths are typically on the order of tens of terahertz. If the two timescales become comparable, however, the approximation underlying our present theory would no longer be valid, and a different theoretical model would be required. In the Discussion, we have added a sentence emphasizing this underlying assumption and clarifying the scope of the present method.

8. *dB vs. percentage; measured vs. projected. The “+135% / +44%” in the Conclusion correspond to 3.7 dB and 1.6 dB; showing the dB values would be more consistent. Also, the directly measured squeezing improvement (Fig. 7, $-2.45 \rightarrow -2.74$ dB, ~ 0.3 dB) and the larger Fig. 2 gains (up to 2.46 dB) are at different operating points — Fig. 7 at a measured SNR of $9.68 \rightarrow 13.10$ dB with the actual r and η_0 , the Fig. 2 values as loss-model projections at higher SNR, larger r , and higher η_0 . Please state this difference so the two are not conflated.*

Reply: In the revised manuscript, we have added the corresponding improvements in decibels. We have also clarified that the experimentally measured squeezing enhancement in Fig. 7 and the larger improvements predicted in Fig. 2 correspond to different operating conditions. Specifically, Fig. 7 presents the measured improvement obtained at the experimentally realized values of r , η_0 , whereas Fig. 2 shows loss-model projections for larger r and η_0 . We further state that greater improvements are expected as r and η_0 increase.

9. *Title. “Ultrafast homodyne detection” is ambiguous, as it can also denote a CW-LO broadband-detector scheme. Since this work uses a pulsed LO and pulse-by-pulse measurement, making this explicit in the title (e.g. “pulsed” or “pulse-by-pulse”) would help.*

Reply: In the relevant literature, the term “ultrafast” generally refers to optical pulses exhibiting temporal variations on a subpicosecond timescale. Such pulses require not only a broad optical spectrum but also phase locking among the constituent frequency components (i.e., a mode-locked laser). A broadband LO without such mode locking would instead correspond to CW broadband homodyne detection. In our experiment, the LO has both a broad spectrum and phase-locked frequency components and therefore operates in the pulsed regime, enabling pulse-by-pulse quadrature measurements. We therefore believe that the title “Ultrafast homodyne detection” appropriately describes the physical regime considered in this work. To avoid possible ambiguity, we have clarified in the Introduction that the present work concerns homodyne detection using a mode-locked LO.

Reviewer #2:

Continuous-variable (CV)-based quantum information protocols harness the rich modal structure of quantum light towards the development of scalable quantum technologies. In particular, the spectral correlations of ultrafast light sources can be used to generate large, entangled states. However, CV quantum states are vulnerable to decoherence, making their implementation in practical quantum information protocols challenging. This sensitivity to loss is one argument against discrete-variable based approach, i.e., using the photon number or the polarization of light. Nevertheless, CV encoded states, e.g., squeezed states and non-gaussian states, with their rich modal structure, are promising candidates for fast quantum communication and metrology.

In this manuscript, the authors solve the generalized Rayleigh quotient coefficient to improve the signal-to-noise ratio (SNR) of homodyne detectors (HD) by applying optimized temporal weights to quadrature measurements. The paper includes a detailed theoretical background on homodyne detection, temporal correlations in homodyne measurements and implementation of the optimization protocol. More details on the homodyne detector and the development of the optimization protocol are included in the Supplementary Materials. The theory is supported throughout by the appropriate references, to which the results are compared. The theory is followed by a thorough description of the experimental setup, measurement protocol and results. The results show a clear benefit of the proposed optimization method in extracting quadrature measurements.

Therefore, I think that research presented in this article shows a concrete benefit to the development of CV-based quantum information protocols, and that Efficient ultrafast homodyne detection of quantum light should definitely be accepted for publication in Communications Physics.

Below, I have added a few questions and remarks:

Reply: We thank the reviewer for recommending our manuscript for publication in Communications Physics and for providing us constructive comments that have helped us improve the manuscript. We address the reviewer's comments below.

Section III – Page 3: *Could you clarify what are sampling errors in the estimation of the shot and electronic noise?*

Reply: The sampling errors arise from a finite number of measurement data for estimating the matrix elements $S_{lm} = \langle \hat{v}(t_l) \hat{v}(t_m) \rangle$ and $E_{lm} = \langle \hat{v}_e(t_l) \hat{v}_e(t_m) \rangle$. The resulting statistical fluctuations can make the optimization sensitive to sampling errors and lead to an unreliable estimate of the optimal weighting vector. This effect is particularly pronounced at high frequencies, where the detector response and electronic-noise power are weak. Small estimation errors can then introduce spurious high-frequency components into the optimal weight. For example, $\mathbf{w}^T \mathbf{E} \mathbf{w}$ may become nearly zero while $\mathbf{w}^T \mathbf{S} \mathbf{w}$ remains finite, yielding a spuriously diverging SNR.

To mitigate this issue, we restrict the frequency components of $\mathbf{w}$ to frequencies below a cutoff frequency f_c , as described in Eqs. (17)–(20). To clarify the effect of sampling errors, we have added the following sentences to the main text: “This issue is particularly pronounced at

high frequencies, where the detector response and electronic-noise power are weak, making the estimates of $\mathbf{S}$ and $\mathbf{E}$ more susceptible to sampling errors. For example, $\mathbf{w}^T \mathbf{E} \mathbf{w}$ may become nearly zero while $\mathbf{w}^T \mathbf{S} \mathbf{w}$ remains finite, yielding a spuriously diverging SNR.”

Section III – Page 4:

- *The theory shows that the optimization performance is sensitive to increasing frequency; does it lead to a restriction of the bandwidth over which a temporal-correlation-improved homodyne detector can be designed or a restriction on the laser than can be used (i.e., repetition rate)?*

Reply: The optimization algorithm itself does not impose a fundamental limitation on the usable bandwidth. Rather, the optimization bandwidth is limited by the detector bandwidth. In our experiment, we set the cutoff frequency to 240 MHz because, above this frequency, the electronic noise becomes comparable to the shot noise, as shown in Fig. 4(a). A detector with a broader bandwidth would therefore allow a higher cutoff frequency for the optimization. The optimization bandwidth is not directly related to the laser repetition rate. The relevant requirement is that the detector bandwidth be sufficiently larger than the laser repetition rate so that successive pulses can be resolved individually.

- *What do the different η_0 values correspond to? A different HD design or modified mode-matching?*

Reply: Here, η_0 denotes the combined efficiency of all factors other than the detection efficiency η_e . It includes the efficiencies of state generation and propagation, the mode matching efficiency between the local oscillator and the quantum state, and the photodiode quantum efficiency. Therefore, different values of η_0 represent different levels of these combined optical efficiencies and do not correspond exclusively to either a different homodyne-detector design or modified mode matching. We have revised the manuscript to clarify the definition of η_0 .

Section IV – Page 5: Although the data collection protocol is thorough, can the authors clarify what the definition of an “experimental” pulse is, the relation between consecutive experimental pulses compared to the theoretical definition, and what is the relation between the choice of T and f_c , if any.

Reply: In the experimental section, the term “pulse” was used to refer to both the femtosecond optical pulse and the corresponding electrical pulse signal produced by the homodyne detector, although they are physically distinct. The consecutive pulses in the experimental part refer to the pulse signals from the detector generated by successive optical pulse; thus, the experimental definition agrees with the theoretical one. To avoid ambiguity, we now distinguish between “optical pulse” or “pulse signal” in the revised manuscript.

T is the measurement window over which a quadrature outcome is obtained. In our experiment, T is set equal to the femtosecond pulse repetition period, 12.5 ns, corresponding to the inverse of the repetition rate, 80 MHz. It therefore defines the time interval assigned to

each repeated quadrature measurement. f_c is the cutoff frequency used to determine the optimal weight. As similarly explained above, there is no direct relation between the cutoff frequency f_c and the measurement window T . The relevant requirement is that the detector bandwidth be sufficiently larger than the laser repetition rate so that successive pulses can be resolved individually.

Conclusion – Page 7: potentially “detector signals” instead of “detected” signals.

Reply: We thank the reviewer for the suggestion. The manuscript uses the phrase “detector signals” rather than “detected signals.”

Reviewer #3:

I co-reviewed this manuscript with one of the reviewers who provided the listed reports. This is part of the Communications Physics initiative to facilitate training in peer review and to provide appropriate recognition for Early Career Researchers who co-review manuscripts.

Reply: We thank the reviewer for their time and effort in reviewing our manuscript.

Reviewer #4:

This manuscript presents a method to improve ultrafast homodyne detection by optimizing the temporal weight used to extract pulse-resolved quadrature outcomes. The optimization is formulated as a generalized Rayleigh quotient problem based on experimentally estimated autocorrelation matrices of the total detector output and the electronic noise. The authors then demonstrate experimentally that the optimized temporal weight increases the homodyne signal-to-noise ratio (SNR) relative to conventional constant and peak-shaped weights, and that this improvement leads to enhanced observed squeezing and anti-squeezing levels.

Overall, I find the manuscript technically sound, clearly written, and potentially interesting for the community working on ultrafast continuous-variable measurements and pulsed homodyne detection. The central idea is physically sensible: the useful homodyne signal and the electronic noise have different temporal covariance structures, and this difference can be exploited to define a more efficient temporal readout mode. The theoretical formulation is coherent, and the experimental demonstration is relevant.

However, in its present form, I think that several aspects require clarification before the conclusions can be considered fully convincing. My main concerns are listed below.

Reply: We thank the reviewer for the constructive comments, which have helped us improve the manuscript. Below, we address the reviewer's comments point by point.

1. The central statement of the paper is that "the pulsed detector signals in ultrafast homodyne detection exhibit strong temporal correlations that are absent in the electronic noise." This point is essential for the whole manuscript, but its physical origin is not explained clearly enough.

As I understand it, these correlations should not be interpreted as correlations between successive optical pulses. Rather, they arise from the finite impulse response of the homodyne detector: each ultrafast optical pulse is much shorter than the electronic response time of the detector and is therefore converted into an extended voltage waveform.

If correct, this physical interpretation should be stated explicitly, preferably already in the Introduction, since it is the basic reason why temporal weighting can improve the SNR. The present wording may be misleading because the electronic noise is not necessarily completely free of temporal correlations: finite electronic bandwidth generally produces stationary autocorrelations. What seems to be absent in the electronic noise is instead the pulse-synchronous, localized, approximately rank-one covariance structure associated with the optical homodyne signal.

Reply: We thank the reviewer for pointing out the need to clarify the physical origin of the temporal correlations. These correlations do not represent correlations between successive optical pulses. Rather, the detector response spreads the quadrature information carried by each ultrafast optical pulse over time, producing pulse-synchronous temporal correlations within the corresponding detector signal.

We also agree that the electronic noise may exhibit temporal correlations owing to the finite electronic bandwidth. The relevant distinction is therefore not the complete absence of temporal correlations in the electronic noise, but the absence of the pulse-synchronous

correlations associated with the detector signal. To clarify these points, we have revised the introduction as follows: “We observe that the detector response spreads the quadrature information of each optical pulse over time, producing strong pulse-synchronous temporal correlations in the detector signals. Such pulse-synchronous correlations are not present in the electronic noise.”

2. Related to the previous point, the manuscript should more clearly distinguish between intra-pulse temporal correlations and correlations between different pulses. In Fig. 6(d), the authors report a residual correlation of approximately 6% at $d=1$ for the constant-weight strategy, while the peak-shaped and optimized weights show negligible correlations. This is an important experimental observation, but its origin is not sufficiently discussed.

Is this residual correlation due to a small overlap of the detector impulse responses from adjacent pulses, baseline recovery, electronic ringing, finite-bandwidth effects, or the use of a nonoptimal integration window? Since the method is partly motivated by the need to avoid correlations between pulse-resolved quadrature samples, the authors should explicitly discuss the mechanism behind the observed $d=1$ correlation and explain why the optimized weight suppresses it.

Reply: We thank the reviewer for giving us the opportunity to clarify this important advantage of our method. The residual correlation arises from the extended tails of the detector responses from adjacent pulses. In the constant-weight strategy, these tails are integrated uniformly over the measurement window, resulting in a residual pulse-to-pulse correlation. In contrast, the optimal and peak-shaped weighting strategies emphasize the detector signal near the central peak and thereby suppress contributions from neighboring pulses. As a result, the undesired pulse-to-pulse correlation is largely eliminated. We have revised the manuscript to clarify the origin of this residual correlation and its suppression by the weighting strategies.

3. Although I understand that the main focus of the manuscript is the optimized homodyne detection method, the description of the squeezed-state generation in the experimental section is too minimal. Since the final experimental demonstration relies on measuring squeezed vacuum generated by a SPOPO, the reader needs more information to assess the quality and reproducibility of this part of the experiment. In particular, the authors should provide additional experimental details on:

- the SPOPO cavity design and operating point;
- the output coupler reflectivity and approximate escape efficiency;
- the cavity linewidth or finesse;
- the pump power used for the squeezing measurements and its value relative to threshold;
- the measured homodyne visibility;
- the optical losses between the SPOPO and the homodyne detector;
- the photodiode quantum efficiency already quoted in the text, but included in a full efficiency budget;
- the total detection efficiency used to interpret the measured squeezing levels;
- the estimated contribution of electronic noise to the effective homodyne efficiency.

At present, Fig. 7 reports measured squeezing and anti-squeezing levels, and the authors infer consistency with the expected efficiency improvement. However, without a clear

efficiency budget, it is difficult to judge whether the observed squeezing levels are limited mainly by the SNR, optical losses, mode matching, phase noise, or properties of the SPOPO itself.

Reply: We thank the reviewer for pointing out the need for a more complete description of the squeezed-state generation and detection efficiencies. Accordingly, we have expanded the experimental description in the main text and added a new Methods section. The revised manuscript now provides details of the SPOPO, including the cavity configuration, beam radius at the BiBO crystal, reflectances of the input and output couplers, cavity free spectral range, linewidth, finesse, pump power used for the squeezing measurements, and its value relative to the oscillation threshold. We have also included a complete efficiency budget comprising the cavity escape efficiency, propagation efficiency, mode-matching efficiency inferred from the homodyne visibility, photodiode quantum efficiency, the effective homodyne efficiency associated with electronic noise, and phase noise. In addition, we report the total efficiencies calculated from this loss and noise budget and those independently inferred from the measured squeezing and antisqueezing levels.

4. The experimental scheme in Fig. 3 and the corresponding description leave some ambiguity regarding the weak seed beam used to lock the relative phase between the squeezed state and the local oscillator. From the text, it appears that a weak coherent seed beam is injected into the SPOPO and that the relative phase is locked by monitoring a tapped output power at a photodiode. The manuscript should clarify more precisely the propagation direction, optical path, and role of this seed beam. In particular:

- Does the seed propagate in the same direction as the squeezed vacuum?*
- Does any residual seed field reach the homodyne detector?*
- If so, how is its contribution removed or made negligible?*
- Does the seed introduce a coherent displacement of the squeezed state, or is the measured state effectively an undisplaced squeezed vacuum?*
- What exact error signal is used for phase stabilization?*
- Is the phase lock performed on resonance or off resonance?*
- How is degradation of squeezing due to residual seed leakage, phase noise, or cavity detuning avoided?*

Reply: The SPOPO cavity is locked through Pound-Drever-Hall (PDH) locking using an additional locking beam, and a weak coherent seed beam is injected into the SPOPO cavity. The seed, locking beam, and the generated squeezed vacuum are all resonant with the cavity. The output seed beam co-propagates with the squeezed vacuum and is injected into the homodyne detector. The seed serves as a phase reference for the squeezed vacuum. However, this beam does not significantly affect the observed squeezing level because the coherent displacement appears only near DC. As a result, at a sideband frequency of 1.5 MHz, the quantum state is a non-displaced squeezed vacuum. Moreover, we locked the LO phase to measure the p -quadrature, while the mean field is at the x -quadrature. Figure 1 shows the phase-space representation of the squeezed coherent state near DC: green denotes the coherent state, red denotes the phase-squeezed coherent state, and blue denotes the amplitude-squeezed coherent state. The measured mean p -quadrature values are close to zero.




Figure 1: Phase space representation of the coherent state (green), phase-squeezed coherent state (red), and amplitude-squeezed coherent state (green)

Phase locking is experimentally implemented using the Red Pitaya STEMLab 125-14 FPGA and the pyRPL software. The FPGA first records the photodiode signal that monitors the seed intensity and generates the error signal obtained by applying a 1 kHz phase dither. This error signal is used to stabilize the seed-pump phase by minimizing or maximizing the seed intensity. The results are shown in Fig. 2(a). The pump phase is scanned by applying a voltage to the PZT. The gray curve shows the FPGA output voltage, which is amplified before being applied to the PZT. Without the pump, the output intensity remains constant (green curve). When the pump beam is injected, the output intensity varies according to the relative phase. The system is locked to either the amplification point (red curve, phase-squeezed) or the deamplification point (blue curve, amplitude-squeezed). Subsequently, the phase between the generated squeezed coherent light and the LO is locked to the phase (p) quadrature using the homodyne signal. In this process, the FPGA does not generate a separate error signal; instead, it directly adjusts the PZT to maintain the LO phase. The phase between the LO and the amplified or deamplified seed is successfully locked, as shown in Fig. 2(b). Since the p -quadrature is measured, squeezing is observed when the seed is amplified, while anti-squeezing is observed when the seed is deamplified. We have included the method of phase locking in Methods.




Figure 2: (a) Detector output for the measurement of the seed intensity while the seed-pump relative phase is scanned or locked. (b) Near-DC homodyne detection output while the seed-LO relative phase is scanned or locked.

5. The paragraph describing the timing correction is not sufficiently clear. The authors state that the oscilloscope sampling rate is not an exact integer multiple of the pulse repetition rate, so that the measurement window drifts relative to the pulse train. They then mitigate this effect using a polyphase finite-impulse-response resampling algorithm, interpolating the 125 samples in each period to 1000 samples, compensating the drift, and downsampling back to 125 samples.

Why was the acquisition not directly synchronized or triggered from the laser oscillator, the LO, or the weak seed?

Reply: We simultaneously record the voltage traces from the homodyne detector and the laser synchronization pulses. We have revised the manuscript to clarify the use of the sync pulses. Since the clocks of the oscilloscope and the light source are not synchronized, the relative timing drift between the measurement windows and the sync pulses cannot be avoided. Instead, we use the upsampled sync pulses to realign the homodyne detector signal within each measurement window on a finer temporal grid. Further details are provided in Methods.

6. Figures 5(g), 5(h), and 5(i) should be improved visually. The caption states that the three leading eigenvectors are shown, but the meaning of the three colors is not immediately clear from the figure itself. The authors should either add a legend directly in the panels or label the curves more clearly. The same applies to the normalization and sign convention of the eigenvectors, since eigenvectors are defined only up to an arbitrary sign.

Reply: We have added the labels indicating the index of the eigenvectors on the right side of each plot. Additionally, we have redefined the eigenvectors such that their norm is unity and the sum of all components is positive.

7. I suggest inserting the following references.

a. A recent and directly relevant reference on time-resolved balanced homodyne detection for high-repetition-rate pulsed sources is Altilia et al., “A Simple and Robust Balanced Homodyne Detector for High-Repetition-Rate Pulsed Sources”. <https://arxiv.org/abs/2604.06994>

b. The recent experimental paper “Generation and detection of squeezed states via a synchronously pumped optical parametric oscillator” (Applied Physics B, 2026) would be a particularly relevant addition in the paragraph introducing SPOPO-based pulsed squeezing. <https://link.springer.com/article/10.1007/s00340-025-08617-6>

Reply: Thank you for suggesting relevant references. We have included these papers in the revised manuscript.

I recommend publication of this paper after major revisions.

Reviewer #1:

I thank the authors for the careful revision. The added statements of the underlying assumptions, the quantitative SNR–efficiency correspondence in the Introduction, the dB values in the Conclusion, and the clarification of the different operating points in Figs. 2 and 7 address my earlier comments well. The new Methods section, added in response to another referee, also makes the efficiency budget easy to follow. I now support publication.

Reply: We thank the reviewer for the positive evaluation of our work. Below, we address the reviewer's last remaining comment.

I would like the authors to check one remaining point before publication. It concerns internal consistency of the revised text.

After Eq. (16), the revised manuscript states: “In realistic detectors, however, the rank-one kernel model of Eq. (15) is generally insufficient because imperfections such as timing jitter and detector-response shape variations increase the rank of the detector response matrix R .” In the discussion of Fig. 5(f), however, the manuscript states that the dominance of the leading eigenvalue “is consistent with the ideal detector response model in Eq. (15), indicating that contributions from nonideal detector responses are negligible”, and that the contribution of the 1st eigenvector associated with timing jitter is negligible. These two statements do not agree with each other. For the detector used in this work, R is effectively rank-one, and the optimal weight is then the whitened matched filter for the measured colored electronic noise.

I do not think this affects the main claim of the paper. As I understand it, the contribution is the formulation of pulsed ultrafast homodyne detection that leads to the generalized eigenvalue problem, the unification of the area-integration and peak-value methods within it, the high-frequency truncation procedure, and the experimental verification. None of these depends on R having rank higher than one.

I therefore suggest to revise the sentence after Eq. (16) as follows: state that for the present detector the rank-one condition is well satisfied, so that the optimal weight coincides with the classical whitened matched filter [47–49], and that the generalized Rayleigh quotient formulation remains valid when the rank-one condition fails, for example for detectors with larger timing jitter or in the regime where the pulse duration and the detector response time become comparable. In my view this is a stronger position than the present one: the agreement with the classical result then works as a consistency check of the method, instead of a point the authors need to argue against.

Reply: We thank the reviewer for this additional comment and for pointing out the potential inconsistency in our revised manuscript. We agree that the distinction between the general formulation and the specific behavior of our detector should be stated more clearly.

Accordingly, we have revised the sentence after Eq. (16) to emphasize that, in general, the rank-one kernel model of Eq. (15) may become insufficient when nonideal detector responses, such as timing jitter or variations in the detector-response shape, increase the rank of the detector response matrix (R). We then clarify in the discussion of Fig. 5(c,f) that, for the detector used in our experiment, the response matrix is experimentally found to be close to rank one. In this case, the optimal weighting vector approaches the result of the whitened

matched filter method.

Specifically, the revised sentence after Eq. (16) reads: “In general, however, the rank-one kernel model of Eq. (15) may be insufficient because imperfections such as timing jitter and detector-response shape variations can increase the rank of the detector response matrix R .”

We have also revised the discussion of Fig. 5(c,f) as follows: “In particular, the dominance of the leading eigenvalue in Fig. 5(f) is consistent with the ideal detector response model in Eq. (15), indicating that contributions from nonideal detector responses are negligible and that matched filtering therefore provides a good approximation to the optimal weighting vector [47–49].”

These revisions clarify that our generalized Rayleigh-quotient formulation includes the matched-filter result as the rank-one limit, while remaining applicable when the rank-one condition is not satisfied.

Reviewer #2:

In response to the initial review, the authors have substantially revised their manuscript. They thoroughly addressed the referees' comments and questions, and updated their manuscript accordingly. They expanded the measurement-protocol description to add key definitions and clarifications. They also added a detailed description of their experimental setup and data-processing methodology. These additions significantly improve the manuscript's clarity; therefore, I recommend it for publication.

Reply: We thank the reviewer again for recommending our manuscript for publication in Communications Physics. Constructive comments from the reviewer have been very helpful to improve the manuscript.

Reviewer #4:

The authors have adequately addressed all the concerns raised in my previous report. The revisions have significantly improved the clarity and completeness of the manuscript, and I am satisfied with the authors' responses. I therefore recommend the manuscript for publication.

Reply: We thank the reviewer for recommending our paper for publication in Communications Physics. The reviewer's constructive comments have been very helpful to improve our manuscript.