nature portfolio

Peer Review File

Low-power integrated optical amplification via second-harmonic resonance

Corresponding Author: Professor Amir Safavi-Naeini

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

Version 0:

Reviewer comments:

Referee #1

(Remarks to the Author)
A. Summary of the key results

Although optical parametric amplification (OPA) is not a new topic in itself, the field has long been waiting for a truly groundbreaking representative work. What makes a work "groundbreaking," as rightly emphasized by the authors, lies in two critical aspects: high gain and low pump power—which fundamentally determine its practical applicability.
In industrial practice, researchers heavily rely on erbium-doped fiber amplifiers (EDFAs), which has confined frequency conversion and parametric resonance research largely to the telecommunications band. While significant progress has been made within this band, we have also been forced to accept inherent limitations, such as noise issues. More critically, this EDFA-dependent research paradigm—frankly speaking, including my own work—has inherently stifled the development of richer optical bands for future applications in communications, optical computing, and microwave photonics.

B. Originality and significance: if not novel, please include reference

This work has made significant original progress, and it is the originality of the underlying devices, which is of great significance.

C. Data & methodology: validity of approach, quality of data, quality of presentation

The authors demonstrated their commitment to breaking away from this constraint as early as in their previous OPO study published in Nature Photonics, which already showed remarkable conversion efficiency. From a methodological perspective, this represents a very positive shift in research logic, and the present results are even more encouraging. Achieving 17 dB gain with 2000 mW pump power can be regarded as the first crucial step of progress in the industry. Most notably, the simultaneous realization of three key features—fiber–chip–fiber net gain, a bandwidth exceeding 100 nm, and microwatt-level pump power—all implemented in a CMOS-compatible lithium niobate on insulator (LNOI) platform—constitutes a milestone achievement.

D. Appropriate use of statistics and treatment of uncertainties

The authors' rigorous approach to statistics and their meticulous treatment of uncertainties are highly commendable and significantly strengthen the credibility of their findings. Rather than merely reporting representative data, they consistently provide clear error bars and confidence intervals across multiple measurements, which is essential for a meaningful interpretation of the results. Furthermore, their methodology for propagating uncertainties—particularly in key metrics such as the parametric gain, conversion efficiency, and bandwidth—is thoroughly documented. This transparency not only allows for a robust assessment of the performance claims but also sets a high standard for reproducibility in the field. Such statistical diligence is a cornerstone of reliable scientific reporting and is particularly crucial for a milestone work that is likely to serve as a benchmark for future studies."

E. Conclusions: robustness, validity, reliability

The conclusions presented are compelling and well-supported, effectively deriving their robustness, validity, and reliability directly from the extensive evidence compiled in the Extended Data. The authors have successfully transcended a mere proof-of-concept, demonstrating the reproducibility and statistical significance of their key results through rigorous and multifaceted validation. This comprehensive approach firmly establishes their work as a dependable foundation for future research and applications in the field.

F. Suggested improvements: experiments, data for possible revision

(1) “line 140” Regarding the potential for co-optimizing broadband guided-wave and resonant-pumped amplification, I remain skeptical that this approach fundamentally addresses the bandwidth limitation. A more precise interpretation might be that the exceptional device efficiency (which mitigates the typical efficiency-bandwidth trade-off) primarily creates design headroom for future bandwidth enhancement, rather than representing a direct optimization for bandwidth itself. I would appreciate the authors' perspective on this point.
(2) “Figure 1e” The authors' design is notably sound and pragmatic. The direct use of filtering and resonance is a strategic choice—one that we have also prioritized in our own work to avoid introducing excess noise or interference, which is paramount for quantum applications such as the generation of squeezed states. I am in complete agreement with this approach.
That said, focusing specifically on the OPA performance metrics of bandwidth and efficiency (and setting aside the quantum prospects for a moment), I have a question: would it be feasible to introduce an interferometric mechanism in this architecture as a means to further enhance the operational bandwidth?
(3) “line 207” The difficulty of achieving uniform poling increases almost exponentially with the device's effective length (L), which represents a daunting challenge for future scaling. While the authors have outlined a prospective path, the viable routes are limited and hinge on rigorous device optimization. At this juncture, it remains unclear whether the on-chip OPA approach can unequivocally surpass traditional bulk-optic solutions.
To better contextualize the promise of the integrated approach, I would recommend that the authors provide a projection based on their current results and existing fabrication capabilities. A quantitative forecast regarding the achievable effective nonlinear interaction length while maintaining high performance would be highly informative. Specifically, adding a figure that maps the predicted performance (e.g., gain or conversion efficiency) against device length, while overlaying current and anticipated fabrication milestones, would powerfully illustrate the roadmap and potential superiority of the on-chip scheme.

G. References: appropriate credit to previous work?

The authors' citations are comprehensive and their evaluations are fair and to the point.

H. Clarity and context: lucidity of abstract/summary, appropriateness of abstract, introduction and conclusions

The content of the article is clear and its structural logic is reasonable.

Referee #2

(Remarks to the Author)
This manuscript presents a outstanding and impactful work on integrated photonics. The authors demonstrate a second-harmonic-resonant optical parametric amplifier (OPA) on thin-film lithium niobate (TFLN) that achieves remarkable performance. The core innovation—decoupling the resonant pump enhancement from the broadband signal amplification—is elegant and effectively solves the critical power consumption problem that has long hampered the integration of practical OPAs. The results represent a significant leap forward and will undoubtedly influence future developments in both classical and quantum photonic systems. I think the paper can be accepted after the following concerns are addressed.

1. In phase-sensitive amplification mode, the gain is highly sensitive to pump wavelength and phase. Did the authors observe resonance frequency instability caused by thermal drift or environmental disturbances? Furthermore, such gain and phase fluctuations would directly translate into additional phase noise, potentially degrading the output signal's noise figure beyond the quantum limit. Have the authors considered the noise introduced by resonator instability?

124 mW of FH power and a 95% conversion efficiency indicates that 117.8 mW of SH light is generated in the microring. Such a high SH power propagating in the TFLN waveguide would likely induce quite significant instability. How did the authors deal with this problem?

Does the test setup diagram in Figure 3a of the article include tests for phase-sensitive amplification?

2. In Figure 2b, the second harmonic resonance peak at high power exhibits significant broadening, which the authors attribute to thermal effects. This is reasonable. However, we note that the lineshape of this resonance peak does not appear entirely symmetrical. Can you comment on this?

3. In Fig. 3b, why is the gain near zero when the pump and signal wavelengths are set to 1575 nm and 1580 nm, respectively?

Referee #3

(Remarks to the Author)
The manuscript reports an integrated optical parametric amplifier (OPA) on thin-film lithium niobate (TFLN) using a second harmonic resonant architecture. To the best of my knowledge, this is the first demonstration of an OPA with over 10 dB gain at 0.2 W pump power on this platform.
Moreover, the authors use the device to demonstrate a broadband gain of over 110 nm. In detail, a higher gain of 7-20 dB is demonstrated under one pump condition. Then a moderate gain with a flat near-quantum-limited noise is realized under another.

The manuscript is clearly written and technically competent, and the experimental data is presented with care. However, the work does not meet the conceptual or experimental standard for a high-impact journal. The core concept—resonant enhancement of second-harmonic power to reduce pump requirements—is incremental, and several of the main claims are insufficiently validated experimentally. The manuscript lacks a convincing demonstration of either a fundamental breakthrough or a system-level application.

My comments are given below:

1. The central novelty claimed in the abstract— “an order-of-magnitude increase in effective pump power via resonant enhancement”—is technically interesting but conceptually incremental. Resonant enhancement of nonlinear conversion efficiency is well known in $\chi(2)$ systems. Moreover, it raises a more fundamental concern that broadband performance appears inconsistent with the resonant nature of the design.
This is clearly shown in Fig. 6a, where a period of 2.5 nm in signal wavelength can be observed in the measured gain which is likely related to the free spectral range of the used resonator (This "gain ripple" is inherent to the resonant architecture and is distinct from the much finer Fabry-Perot ripples discussed in the manuscript in section VIII.F, which are on the order of 0.05 nm). When the pump/signal is in resonance the SHG efficiency/gain is comparably high, but when it’s out of resonance the performance degrades up to 5-10 dB.
In comparison, a waveguide based OPA provides non-resonant based flat gain across a wide band of over 100 nm. This limitation becomes more critical for frequency-comb applications. A flat gain is required to uniformly amplify the entire comb spectrum to maintain a good operational performance as in (i) Nikolai K., et al. (Nature 2025).

2. Moreover, even at resonance as shown in the paper, the gain performance is only comparable to that of existing waveguide-based works. For example, in (i) Yong Meng S., et al. (Optics Letters 2018), a 15.4 dB gain was achieved with a pump power of 30 dBm. Furthermore, Fig. 3c suggests a gain saturation above 500 mW pump power, implying that the resonator will be fundamentally limited to gains below ~25 dB gain.
A WG-based OPA can easily achieve 35 dB gain or higher with a 100 nm or wider bandwidth, as in (i) Mengwen C., et al. (OPTICA 2025) (ii) Yong Meng S., et al. (Optics Letters 2018) and (iii) Luis L., et al. (OPTICA 2022). The authors need to provide proof of realizing a comparable gain of 35 dB or higher in this design.
Also, the statement in Section VIII.I that “saturation effects are out of scope” is a major oversight. Saturation is a critical performance constraint for resonant OPAs and should be addressed experimentally, as likely explains the gain saturation shown in the experiment.
Overall, the manuscript does not convincingly show that this approach provides a performance advantage over existing waveguide-based OPAs.

3. Furthermore, in Abstract, as the second novelty the author claims that “We demonstrate flat near-quantum-limited noise performance over 110 nm”.
Noise figure is a key metric for OPAs, especially in quantum information applications. However, this claim is not substantiated by direct experimental evidence. The presented noise figure is calculated from measured spontaneous parametric fluorescence and gain, rather than being directly measured via the standard and definitive method of measuring the degradation of the signal-to-noise ratio, such as the Y-factor technique or, even better for a phase-sensitive amplifier, a homodyne detection setup.
Claims of quantum-limited operation are exceptionally high-stakes and require the highest standard of proof. Relying on theoretical models and indirect calculations (as detailed in Section VIII.H) is insufficient. Without direct measurement, the claim of "near-quantum-limited" performance remains an unverified assertion, significantly diminishing the impact and reliability of the results.
In general, the work contains critical flaws in experimental validation that fundamentally undermine its central claims and limit its scientific impact.

4. Moreover, due to the lack of a breakthrough in performance compared to other chip-level OPAs, the claimed application scenario in Conclusion is also lacking support, “… enable a fully integrated OPA subsystem suitable for quantum and classical applications”.
Recent high-impact works have integrated OPAs into system-level demonstrations such as (i) Nikolai K., et al. (Nature 2025). (ii) Ping Z., et al. (Nature 2025) and (iii) Mengwen C., et al. (OPTICA 2025). All have demonstrated a system enabling coherent communications. Additionally, the first of these realized a direct amplification of frequency combs (both EO combs and Kerr combs) with a flat net gain over 20 dB across the spectrum. Compared to these works, the actual performance of the demonstrated OPA was not fully characterized, thus lacking proof for the claim.
The paper would benefit from a clearly defined application. Depending on the desired application, the performance of the OPA could be sufficient and for others it is not.

In summary, this paper demonstrates a resonance-based OPA on lithium niobate. The platform is powerful, the combination of SHG and OPA is impressive, but too much is missing in this report to represent a major advancement. How will the OPA demonstrated be valuable for these applications? The authors are encouraged to demonstrate an application in their work to prove the robustness of their device. As a result, the work is better suited for a photonics-focused or less selective journal. Alternatively, the authors could continue to develop this platform to demonstrate an integrated coherent modulator or comb with the OPA in the same chip with more compelling evidence — which would certainly be novel and impactful.

I therefore recommend rejection at this stage.

Referee #1 (Remarks to the Author):

A. Summary of the key results

Although optical parametric amplification (OPA) is not a new topic in itself, the field has long been waiting for a truly groundbreaking representative work. What makes a work "groundbreaking," as rightly emphasized by the authors, lies in two critical aspects: high gain and low pump power—which fundamentally determine its practical applicability.
In industrial practice, researchers heavily rely on erbium-doped fiber amplifiers (EDFAs), which has confined frequency conversion and parametric resonance research largely to the telecommunications band. While significant progress has been made within this band, we have also been forced to accept inherent limitations, such as noise issues. More critically, this EDFA-dependent research paradigm—frankly speaking, including my own work—has inherently stifled the development of richer optical bands for future applications in communications, optical computing, and microwave photonics.

B. Originality and significance: if not novel, please include reference

This work has made significant original progress, and it is the originality of the underlying devices, which is of great significance.

C. Data & methodology: validity of approach, quality of data, quality of presentation

The authors demonstrated their commitment to breaking away from this constraint as early as in their previous OPO study published in Nature Photonics, which already showed remarkable conversion efficiency. From a methodological perspective, this represents a very positive shift in research logic, and the present results are even more encouraging. Achieving 17 dB gain with 2000 mW pump power can be regarded as the first crucial step of progress in the industry. Most notably, the simultaneous realization of three key features—fiber–chip–fiber net gain, a bandwidth exceeding 100 nm, and microwatt-level pump power—all implemented in a CMOS-compatible lithium niobate on insulator (LNOI) platform—constitutes a milestone achievement.

D. Appropriate use of statistics and treatment of uncertainties

The authors' rigorous approach to statistics and their meticulous treatment of uncertainties are highly commendable and significantly strengthen the credibility of their findings. Rather than merely reporting representative data, they consistently provide clear error bars and confidence intervals across multiple measurements, which is essential for a meaningful interpretation of the results. Furthermore, their methodology for propagating uncertainties—particularly in key metrics such as the parametric gain, conversion efficiency, and bandwidth—is thoroughly documented. This transparency not only allows for a robust assessment of the performance claims but also sets a high standard for reproducibility in the field. Such statistical diligence is a cornerstone of reliable scientific reporting and is particularly crucial for a milestone work that is likely to serve as a benchmark for future studies."

E. Conclusions: robustness, validity, reliability

The conclusions presented are compelling and well-supported, effectively deriving their robustness, validity, and reliability directly from the extensive evidence compiled in the Extended Data. The authors have successfully transcended a mere proof-of-concept, demonstrating the reproducibility and statistical significance of their key results through rigorous and multifaceted validation. This comprehensive approach firmly establishes their work as a dependable foundation for future research and applications in the field.

F. Suggested improvements: experiments, data for possible revision

(1) "line 140" Regarding the potential for co-optimizing broadband guided-wave and resonant-pumped amplification, I remain skeptical that this approach fundamentally addresses the bandwidth limitation. A more precise interpretation might be that the exceptional device efficiency (which mitigates the typical efficiency-bandwidth trade-off) primarily creates design headroom for future bandwidth enhancement, rather than representing a direct optimization for bandwidth itself. I would appreciate the authors' perspective on this point.

- The reviewer is correct in that the parameter space for OPAs is multidimensional as shown in Figure 1a. Improved performance in one parameter does indeed provide more design headroom to optimize the entire system. However, line 140 specifically refers to the amplification bandwidth of signal-resonant OPAs, which confine the gain spectrum into narrowband peaks around the signal resonances and therefore are not suitable for amplifying broadband signals (with the exception of carefully matched frequency combs). The bandwidth limitation we are referring to is the bandwidth limitation in efficient amplifiers that resonate the signal/idler. See, for example Y. Zhao et al. Optica 2023.
- We clarified this line in the text
    - Unlike typical signal-resonant and fully-resonant systems, which achieve ultralow-power nonlinear interactions but limit the bandwidth to narrowband resonances,

(2) "Figure 1e" The authors' design is notably sound and pragmatic. The direct use of filtering and resonance is a strategic choice—one that we have also prioritized in our own

work to avoid introducing excess noise or interference, which is paramount for quantum applications such as the generation of squeezed states. I am in complete agreement with this approach.

That said, focusing specifically on the OPA performance metrics of bandwidth and efficiency (and setting aside the quantum prospects for a moment), I have a question: would it be feasible to introduce an interferometric mechanism in this architecture as a means to further enhance the operational bandwidth?

- The reviewer raises an interesting question. We are not sure an interferometric mechanism could significantly increase the amplification bandwidth, but there has been prior work on dispersion engineering thin-film lithium niobate devices to achieve extremely broad amplification bandwidth by selecting a waveguide geometry where the phase mismatch has reduced wavelength-dependence (see Jankowski et al 2021).

(3) “line 207” The difficulty of achieving uniform poling increases almost exponentially with the device's effective length (L), which represents a daunting challenge for future scaling. While the authors have outlined a prospective path, the viable routes are limited and hinge on rigorous device optimization. At this juncture, it remains unclear whether the on-chip OPA approach can unequivocally surpass traditional bulk-optic solutions.

To better contextualize the promise of the integrated approach, I would recommend that the authors provide a projection based on their current results and existing fabrication capabilities. A quantitative forecast regarding the achievable effective nonlinear interaction length while maintaining high performance would be highly informative. Specifically, adding a figure that maps the predicted performance (e.g., gain or conversion efficiency) against device length, while overlaying current and anticipated fabrication milestones, would powerfully illustrate the roadmap and potential superiority of the on-chip scheme.

- The reviewer raises an important question regarding expectations for future capabilities. The paragraph around line 207 is meant to highlight what is currently feasible, and this projected performance is plotted in Figure 1d. In fact, excellent periodic poling across lengths as long as 21 mm have been demonstrated through use of adaptive poling (Chen et al. 2023). Note, however, that our design would only provide strong benefit at such long lengths if second harmonic propagation losses can be reduced further. We have added clarification to the sentence:
    - Attaining the simulated nonlinearity $\eta_0 = 4000 \frac{\%}{\text{Wcm}^2}$ by improved poling, a modest improvement in the finesse (similar to that achieved on another device on the same chip), and extending the nonlinear length to $L=12$ mm (within current capabilities) would reduce power requirements for SH generation by another order of magnitude.

G. References: appropriate credit to previous work?

The authors' citations are comprehensive and their evaluations are fair and to the point.

H. Clarity and context: lucidity of abstract/summary, appropriateness of abstract, introduction and conclusions

The content of the article is clear and its structural logic is reasonable.

Referee #2 (Remarks to the Author):

This manuscript presents a outstanding and impactful work on integrated photonics. The authors demonstrate a second-harmonic-resonant optical parametric amplifier (OPA) on thin-film lithium niobate (TFLN) that achieves remarkable performance. The core innovation—decoupling the resonant pump enhancement from the broadband signal amplification—is elegant and effectively solves the critical power consumption problem that has long hampered the integration of practical OPAs. The results represent a significant leap forward and will undoubtedly influence future developments in both classical and quantum photonic systems. I think the paper can be accepted after the following concerns are addressed.

1. In phase-sensitive amplification mode, the gain is highly sensitive to pump wavelength and phase. Did the authors observe resonance frequency instability caused by thermal drift or environmental disturbances? Furthermore, such gain and phase fluctuations would directly translate into additional phase noise, potentially degrading the output signal's noise figure beyond the quantum limit. Have the authors considered the noise introduced by resonator instability?

- The reviewer raises a good question - relative phase fluctuations between signal and pump are detrimental to any phase-sensitive amplifier. A measure of these relative phase fluctuations are shown by the amplification-deamplification oscillations in Figure 4c - there is both slow second-timescale drift and also fast millisecond-timescale fluctuations. In our measurements, where the signal and pump each pass through several meters of off-chip fiber, we attribute the relative phase fluctuations to temperature drift and path length changes in these fibers. Because the on-chip path lengths are so small (millimeter-scale) and the temperature fluctuations are also small, we do not believe the on-chip environment caused significant fluctuations. This is supported by the fact that we did not observe much drift in the resonance frequency - a phase drift of pi of the round-trip second harmonic path in the resonator would completely detune the device from functioning - but we observed no noticeable resonance drift and so assume that the on-chip fluctuations are small. The below paragraph in the text has been updated to clarify this.

        - Figure 4c shows the phase-sensitive transmission of a degenerate signal. Without any phase locking, the relative phase between the FH pump and degenerate seed drifts in time with temperature fluctuations and vibrations along their fiber paths, resulting in the observed oscillatory amplification and deamplification cycles. We attribute the relative phase fluctuations to off-chip effects, as the stable cavity resonance indicates that the shorter on-chip paths are phase-stable.

124 mW of FH power and a 95% conversion efficiency indicates that 117.8 mW of SH light is generated in the microring. Such a high SH power propagating in the TFLN waveguide would likely induce quite significant instability. How did the authors deal with this problem?

- The reviewer raises a good point. In fact, there is even more than 117.8 mW of power circulating within the resonator due to the resonant enhancement. We found – counter-intuitively – that the higher powers led to thermal-broadening-enhanced stability. Because the photothermal effects shift the resonance towards longer wavelengths, we always tuned from short wavelengths to longer ones for wavelength swept measurements and for tuning onto a particular mode. This way, we set up a negative feedback loop for the laser-resonator frequency detuning and enhance stability. We did not observe photorefractive effects (which would have caused a blue shift in resonance frequency rather than the thermal red shift).
- We have the following comment on this stability in the Methods section
    - We tune the FH pump onto an SH resonance and leave it there without any active feedback. We find the resonances are stable for at least tens of minutes, particularly at higher powers where the resonances are thermally broadened, and therefore wavelength-cavity locking is not necessary for our measurements.

Does the test setup diagram in Figure 3a of the article include tests for phase-sensitive amplification?

- Yes, this is shown by the “switch” where the fiber could be connected to degenerate tap as seed, or the nondegenerate signal laser. We have the following comment on the setup in the figure caption:
    - A nondegenerate signal laser is input on the other side of the chip for phase-insensitive amplification measurements. For phase-sensitive measurements, a tap of the fundamental pump laser is input as degenerate seed instead.

2. In Figure 2b, the second harmonic resonance peak at high power exhibits significant broadening, which the authors attribute to thermal effects. This is reasonable. However, we note that the lineshape of this resonance peak does not appear entirely symmetrical. Can you comment on this?

- The lineshape both broadens and becomes asymmetric (shark-fin shaped) due to the photothermal effect - as more optical power is resonated close to resonance, the heating shifts the resonance. If the resonance is approached from the shorter wavelengths, then the heating will shift if further and effectively broaden it. If the resonance is approached from longer wavelengths, then the heating will draw the resonance closer and a bistable point is crossed. We added a reference (Carmon et al. 2004) on photothermal broadening in microresonators.
    - At higher powers, FH depletion deepens and thermally induced resonance shifts\cite{Carmon:04} distort and broaden the SH lineshapes, increasing tolerance to pump detuning and expanding the usable wavelength range.

3. In Fig. 3b, why is the gain near zero when the pump and signal wavelengths are set to 1575 nm and 1580 nm, respectively?

- These regions that lack gain correspond to either minima in the SHG phase matching (horizontal line around pump wavelength of 1575 nm) or OPA phase matching (curved line around signal wavelength of 1580 nm). We added a clarifying sentence to the Methods section's description of this measurement.
    - We measured the OPA gain spectrum of Figure \ref{fig_opa}b by recording signal transmission while the FH pump wavelength was swept across the SHG phase matching bandwidth. The generation of resonant SH pump led to amplification of those signal wavelengths that were phase matched (the horizontal and curved regions that lack gain in Figure \ref{fig_opa}b correspond to phase-mismatched wavelengths for SHG and OPA, respectively).

Referee #3 (Remarks to the Author):

The manuscript reports an integrated optical parametric amplifier (OPA) on thin-film lithium niobate (TFLN) using a second harmonic resonant architecture. To the best of my knowledge, this is the first demonstration of an OPA with over 10 dB gain at 0.2 W pump power on this platform.
Moreover, the authors use the device to demonstrate a broadband gain of over 110 nm. In detail, a higher gain of 7-20 dB is demonstrated under one pump condition. Then a moderate gain with a flat near-quantum-limited noise is realized under another.

The manuscript is clearly written and technically competent, and the experimental data is presented with care. However, the work does not meet the conceptual or experimental standard for a high-impact journal. The core concept—resonant enhancement of second-harmonic power to reduce pump requirements—is incremental, and several of the main claims are insufficiently validated experimentally. The manuscript lacks a convincing demonstration of either a fundamental breakthrough or a system-level application.

My comments are given below:

1. The central novelty claimed in the abstract— “an order-of-magnitude increase in effective pump power via resonant enhancement”—is technically interesting but conceptually incremental. Resonant enhancement of nonlinear conversion efficiency is well known in χ(2) systems. Moreover, it raises a more fundamental concern that broadband performance appears inconsistent with the resonant nature of the design.
This is clearly shown in Fig. 6a, where a period of 2.5 nm in signal wavelength can be observed in the measured gain which is likely related to the free spectral range of the used resonator (This "gain ripple" is inherent to the resonant architecture and is distinct from the much finer Fabry-Perot ripples discussed in the manuscript in section VIII.F, which are on the order of 0.05 nm). When the pump/signal is in resonance the SHG efficiency/gain is comparably high, but when it’s out of resonance the performance degrades up to 5-10 dB. In comparison, a waveguide based OPA provides non-resonant based flat gain across a wide band of over 100 nm. This limitation becomes more critical for frequency-comb applications. A flat gain is required to uniformly amplify the entire comb spectrum to maintain a good operational performance as in (i) Nikolai K., et al. (Nature 2025).

- We respectfully disagree with the reviewer on a few of these points. Regarding the bandwidth concerns, we would like to emphasize that our design is nonresonant, and therefore identical to a broadband traveling wave amplifier in its response to a signal. The pump is resonantly enhanced but *does not limit the signal amplification bandwidth*. We ensure that the signal is nonresonant through the use of four filters, each with up to 25 dB suppression at the signal wavelength. The reviewer noticed a ripple in the measured gain spectrum of Figure 6a, but incorrectly associated it to the SH resonator’s free spectral range. The SH resonator’s free spectral range is also on the order of 0.05 nm, but due to the filters mentioned above, it does not resonate any signal light. Instead, these ripples in the gain spectrum of Figure 6a are in fact due to the fine Fabry-Perot ripples mentioned and a nuance in the measurement. We are comparing signal transmission with and without the high-power pump field. When the high power FH pump is present and generates high power SH pump, it thermally heats the waveguides and by the thermorefractive effect slightly changes the phase of the Fabry-Perot ripples on the signal path. Then, when we calculate gain as the ratio of signal transmission with and without the pump, the slight shift in Fabry-Perot ripples appears as a change in gain. This is evidenced by the fact that the ~ 1 dB ripple is observed even in the absence of gain, so long as the SH pump is present (e.g. FH pump wavelength 1546 nm). This is why we carefully accounted for the gain ripple in the gain measurements shown in Figure 3c.

- We would also like to point out that the referenced paper (Kuznetsov et al. 2025) exhibited much stronger oscillations in transmission (>10 dB) due to both facet reflections and multimode interactions in the waveguide.
- Regarding the novelty, it is true that resonant enhancements of nonlinear conversion efficiency are well known in χ(2) systems, but, to our knowledge, they almost always have been associated with a corresponding decrease in bandwidth. Our system benefits from a resonant enhancement while retaining the desirable broadband properties of the nonlinear conversion process.

2. Moreover, even at resonance as shown in the paper, the gain performance is only comparable to that of existing waveguide-based works. For example, in (i) Yong Meng S., et al. (Optics Letters 2018), a 15.4 dB gain was achieved with a pump power of 30 dBm. Furthermore, Fig. 3c suggests a gain saturation above 500 mW pump power, implying that the resonator will be fundamentally limited to gains below ~25 dB gain.
A WG-based OPA can easily achieve 35 dB gain or higher with a 100 nm or wider bandwidth, as in (i) Mengwen C., et al. (OPTICA 2025) (ii) Yong Meng S., et al. (Optics Letters 2018) and (iii) Luis L., et al. (OPTICA 2022). The authors need to provide proof of realizing a comparable gain of 35 dB or higher in this design.
Also, the statement in Section VIII.I that "saturation effects are out of scope" is a major oversight. Saturation is a critical performance constraint for resonant OPAs and should be addressed experimentally, as likely explains the gain saturation shown in the experiment. Overall, the manuscript does not convincingly show that this approach provides a performance advantage over existing waveguide-based OPAs.

- The reviewer raises a good point that the gains reported in this work are comparable to previously demonstrated gains. The key demonstration, we argue, is that such gains are achieved at low powers that could be generated on-chip. While in some systems, such as table-top setups and pulsed systems, high powers are readily available to achieve high gains, many chip-scale systems are limited by available power.
- We agree with the reviewer that saturation behavior can be critical for some optical amplifier applications, such as in telecommunications. Saturation is often associated with efficient power conversion from the pump to the signal, which can be desirable. The reviewer claims to infer gain saturation above 500 mW pump power from Figure 3c. Our model does not include gain saturation, and our data agrees with the model, so we are not certain what the reviewer is referring to. The measured data points in Figure 3c agree reasonably well with the model trendline, and if anything we expect the peak gain to increase further as the end-facet Fabry-Perot oscillation threshold is approached as shown in Figure 13a. At some point, however, the oscillation threshold will be surpassed and the oscillation will draw pump power and limit achievable small-signal gain. This is a well-known limitation in optical amplifiers, and in the text we note many established mitigation techniques:
    - These ripples can be mitigated by AR coatings, index matching fluids, and adiabatic or angled end facet couplers.

- We did not pursue gain measurements with on-chip pump powers beyond 450 mW on-chip power out of concern for burning a facet or waveguide (the facet of one device burned at comparable power). We expect saturation behavior depends on the injected seed power, which is application-specific. We agree that a study of the saturation behavior of this SH-resonant OPA design is worthwhile and important for many applications, but maintain that it is out of the scope of this work.

3. Furthermore, in Abstract, as the second novelty the author claims that “We demonstrate flat near-quantum-limited noise performance over 110 nm”.
Noise figure is a key metric for OPAs, especially in quantum information applications. However, this claim is not substantiated by direct experimental evidence. The presented noise figure is calculated from measured spontaneous parametric fluorescence and gain, rather than being directly measured via the standard and definitive method of measuring the degradation of the signal-to-noise ratio, such as the Y-factor technique or, even better for a phase-sensitive amplifier, a homodyne detection setup.
Claims of quantum-limited operation are exceptionally high-stakes and require the highest standard of proof. Relying on theoretical models and indirect calculations (as detailed in Section VIII.H) is insufficient. Without direct measurement, the claim of "near-quantum-limited" performance remains an unverified assertion, significantly diminishing the impact and reliability of the results.
In general, the work contains critical flaws in experimental validation that fundamentally undermine its central claims and limit its scientific impact.

- We agree with the reviewer that noise figure is an important measurement that should be performed carefully. The optical measure of noise figure we have demonstrated follows an established technique for characterizing optical amplifiers. See, for example “Theory and Measurement Techniques for the Noise Figure of Optical Amplifiers” (Baney et al 2000) and “Overcoming the quantum limit of optical amplification in monolithic waveguides” (Ye et al. 2021).
- We have modified the wording in Section V to better communicate that this is an optical measurement of noise figure as opposed to one that directly measures the electrical noise spectrum.
    - Comparing the on-chip gain and fluorescence yields an optical measure of the on-chip noise figure (assuming it is limited by signal-spontaneous beat noise and shot noise) shown in Figure \ref{fig_NF}b \cite{OPANFtheory}. The inferred noise figure is quantum limited across most of the 1520 - 1630 nm band.
- We added additional clarification to the Methods section D
    - Although many noise terms can degrade the performance of an optical amplifier, including pump fluctuations and spontaneous-spontaneous beat noise, signal-spontaneous beat noise and shot noise are the dominant and fundamental limitations on performance in many applications\cite{OPANFtheory}.

4. Moreover, due to the lack of a breakthrough in performance compared to other chip-level OPAs, the claimed application scenario in Conclusion is also lacking support, "… enable a fully integrated OPA subsystem suitable for quantum and classical applications".
Recent high-impact works have integrated OPAs into system-level demonstrations such as (i) Nikolai K., et al. (Nature 2025). (ii) Ping Z., et al. (Nature 2025) and (iii) Mengwen C., et al. (OPTICA 2025). All have demonstrated a system enabling coherent communications. Additionally, the first of these realized a direct amplification of frequency combs (both EO combs and Kerr combs) with a flat net gain over 20 dB across the spectrum. Compared to these works, the actual performance of the demonstrated OPA was not fully characterized, thus lacking proof for the claim.
The paper would benefit from a clearly defined application. Depending on the desired application, the performance of the OPA could be sufficient and for others it is not.

We agree with the reviewer that the cited works are impressive and highlight how competitive this field is. In fact, we believe this context underscores the significance of our breakthrough, which directly addresses a key limitation those works do not.

We also respectfully disagree that our work lacks a performance breakthrough.

In our view, **the most important performance metric for integration is power consumption at a given gain**. On this axis our device is substantially ahead of the recent works cited by the reviewer. Our work demonstrates >10x reduction in power consumption compared to all the recent papers noted by the reviewer, while also demonstrating performance better or on par with the cited works on all other metrics. We measure **17 dB gain** for **<200mW power**, while Chen et al show 13.9dB for >2300 mW, Zhou et al show 2.5 dB for 2500 mW, and Kuznetsov et al show 35 dB for 4400 mW. In terms of power efficiency (dB gain per mW pump power) our work *demonstrates* an order of magnitude improvement over these papers. We project even higher performance by outlining a clear path for obtaining another 10x improvement via minor optimizations, a path that is simply not present in the other approaches. A 100x improvement in the most important metric for integration is in our view a significant performance breakthrough.

Also along the lines of improved integration, our design demonstrates a significantly higher level of integration than the cited works. Our device architecture naturally integrates dichroic couplers for on-chip pump-signal multiplexing. The competing works (e.g., Zhao et al., Kuznetsov et al.) explicitly rely on external, off-chip components (like WDMs or filters) to combine and separate the pump and signal. Our approach is thus more scalable and suitable for the "fully integrated" classical and quantum applications we hope to enable.

(Note: The flat net gain the reviewer cites from Kuznetsov et al. 2025 is for the EO-comb with an 8 nm bandwidth. The kerr comb amplification demonstration and full OPA characterization measurements highlighted a non-flat gain spectrum.)

In summary, this paper demonstrates a resonance-based OPA on lithium niobate. The platform is powerful, the combination of SHG and OPA is impressive, but too much is missing in this report to represent a major advancement. How will the OPA demonstrated be valuable for these applications? The authors are encouraged to demonstrate an application in their work to prove the robustness of their device. As a result, the work is better suited for a photonics-focused or less selective journal. Alternatively, the authors could continue to develop this platform to demonstrate an integrated coherent modulator or comb with the OPA in the same chip with more compelling evidence — which would certainly be novel and impactful.

I therefore recommend rejection at this stage.