

Closed-Loop Photoreactor Design Enabled by Machine Learning and Digital Twins

Corresponding Author: Professor Timothy Noel

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)

This manuscript by Noël and co-workers presents a workflow combining ray-tracing simulations, Bayesian optimization, and additive manufacturing for the design of batch photoreactors. The authors build a detailed digital twin of a multi-vial photoreactor, calibrated against actinometric measurements, and apply multi-objective Bayesian optimization to identify reactor geometries that maximize photon flux into vials whilst maintaining irradiation homogeneity. The optimized reactor is subsequently validated experimentally and applied to a quantum yield screening study.

The work described in this manuscript is technically competent, clearly written, and addresses an important challenge in photochemical reactor design. However, this reviewer does have several concerns about the formulation and interpretation of the optimization problem and the strength of several conceptual claims. In its current form, the manuscript overstates the role and impact of machine learning-enabled optimization and, more importantly, appears to base key conclusions on claims that may not be defensible given their current methodologies.

My consideration and comments are outlined below.

1. The first relates to the definition of the homogeneity objective. The authors define irradiation homogeneity as the ratio of minimum to maximum photon flux across vials:

$$H = q_{(i, \max)} / q_{(i, \min)}$$

All geometric parameters explored in the optimization enforce strict rotational symmetry; (i) symmetric placement of all vials (identical radial position and depth); (ii) a rotationally symmetric reflector generated by revolved curves; (iii) symmetric light source arrangement; and (iv) predominantly diffuse reflection from the internal aluminum foil.

Under these assumptions, the physical system is invariant under rotation, and all vials are physically equivalent. In an idealized model (such as that simulated), inter-vial photon flux should therefore be identical by construction. The homogeneity across vials must therefore be unity regardless of the reflector geometry selected ($H=1$ always). Any deviation of H from unity is unlikely to reflect genuine physical inhomogeneity and must instead arise from numerical artifacts of the ray-tracing simulation (finite ray counts, meshing error, surface discretization, or stochastic noise). The “homogeneity” objective is dominated by numerical error and does not relate to reactor performance. This explains several observations: (i) Homogeneity barely improves during optimization; (ii) Scatter dominates the homogeneity axis; (iii) Pareto fronts are dense and shallow; (iv) Best “homogeneity” designs have the smoothest, simplest geometries (those expected to minimize numerical variance in the simulations), and (v) Near identical performance between reactor α and γ in real experiments.

In my view, the homogeneity metric would more accurately be described as a measure of numerical robustness for the simulation as opposed to a measure of the uniformity of photon delivery. This significantly weakens the multi-objective optimization framework and casts doubt on the interpretation of the Pareto fronts and trade-offs presented. If the homogeneity metric instead measures the uniformity of delivery to a single vial (i.e. variance across the surface of a single vial) then the manuscript needs to be altered to make this clear. Homogeneity cannot be measured by comparing photon flux between different vials in the presented system.

2. The second concern relates to the reliance on simulated data without robust validation during optimization. The Bayesian optimization is driven almost entirely by simulated photon flux data. Whilst the authors validate the digital twin against a single arbitrary geometry, this validation does not establish reliability over the full design space explored during optimization. Convergence testing is performed only post-hoc on three selected geometries with no a priori justification that these geometries span the numerical difficulty of the search space. There is no systematic assessment of numerical accuracy (i) across the optimization trajectory, (ii) for geometries with sharper curvature or stronger optical focusing, or (iii) with respect to preservation of relative rankings under mesh or ray-count refinement. Numerical accuracy should be demonstrated where the optimizer operates, not after the fact on selected exemplars. Because the optimization itself appears to suppress numerically challenging geometries, post-hoc convergence on the final designs does not validate the simulation as a reliable objective function. If geometry-dependent numerical errors are present, even at modest levels, the optimization trajectory and resulting optima may be dominated by numerical bias rather than physical signal. In a simulation-driven Bayesian optimization framework, this omission is significant. If the homogeneity metric was measured between vials, then its value gives a good measure of numerical error. The levels this would indicate are not suitable for simulation driven optimization.

3. The third major consideration optimization problem collapses to a trivial design choice. Analysis of the optimization trajectory and results indicate that photon flux is largely dominated by vial proximity and insertion depth, with the Bayesian optimizer consistently pushing these parameters to their bounds. This behavior is entirely intuitive and consistent with basic radiometric considerations. In contrast, most reflector parameters show weak correlations with performance and broad, poorly defined distributions. The optimization therefore appears to collapse onto a trivial geometric trade-off (maximizing proximity to the light source while maintaining spacing), with reflector curvature acting only as a minor secondary effect under predominantly diffuse optical conditions. This is entirely expected if the homogeneity metric is not physically informative. The geometric designs were likely not given a sufficient degree of freedom to meaningfully contribute to the limited. Given this collapse, it is difficult to conclude that meaningful geometric optimization has occurred beyond what would be achieved via straightforward heuristic design.

4. The final consideration concerns terminology. The manuscript repeatedly describes the workflow as “closed-loop” and “autonomous.” However, the optimization loop is closed entirely within simulation, with experimental data used only for one-time calibration and post-optimization validation. In the context of autonomous chemistry and reactor optimization, “closed loop” typically denotes iterative experimental feedback informing the optimization in real time. That is not the case here as the work represents offline, simulation-based optimization with experimental confirmation. While this is still valuable, the terminology overstates the degree of autonomy and robustness, particularly given the central reliance on a fixed digital twin whose accuracy is not continuously constrained by experiment.

Overall, the authors demonstrate an interesting attempt at incorporating ML into reactor design using simulated digital twins. However, the optimization problem as formulated is physically under-constrained, and one of the central objectives appears to lack meaningful physical content due to enforced symmetry. Combined with the heavy reliance on simulation and the absence of robust validation during optimization, this impedes the conceptual claims regarding machine-learning-guided reactor design. In my view, the work would be more appropriately as a simulation-assisted reactor design study rather than a demonstration of closed-loop, multi-objective optimization enabled by machine learning. Major revisions addressing the above concerns would be required to to realign the claims with what is actually demonstrated in this manuscript.

Reviewer #2

(Remarks to the Author)

Photoreactor design requires balancing photon capture with irradiation uniformity. This work integrates ray-tracing digital twins with multi-objective Bayesian optimization to autonomously design reactor geometries. The framework is validated against potassium ferrioxalate actinometry and applied to a photocatalytic Suzuki–Miyaura cross-coupling, demonstrating predictive accuracy and enabling quantitative mechanistic insights across diverse conditions.

The manuscript presents a closed-loop workflow integrating ray-tracing digital twins with Bayesian optimization to autonomously design photoreactors. While digital twins and optimization are established in other fields, their specific application to multi-objective photoreactor co-design for uniform irradiation represents a methodological advance, which is also experimentally validated using actinometry and a photocatalytic reaction.

Strengths

- The digital twin is rigorously validated against potassium ferrioxalate actinometry, ensuring reliable simulation of photon fluxes.
- The closed-loop integration of ray-tracing with Bayesian optimization effectively navigates design spaces for reactor geometry.
- The framework successfully predicts experimental reaction yields in a photocatalytic Suzuki–Miyaura cross-coupling, confirming virtual reactor accuracy.

Major comments

1. The claim of high scientific novelty regarding the predictive accuracy of the virtual reactor requires stronger justification.
2. Page 8: "Together, these results demonstrate that the digital twin accurately predicts photonic performance and that the optimized geometries translate directly into improved and more uniform photochemical outcomes." This point comes from a circular argumentation line. The authors state that they operate in the photon limited regime. Thus, performance linearly depends on the photon flux and with the realized circular arrangement of vials in relation to the circularly emitting light source, only deviations from the uniformity would be scientifically surprising. Higher conversions solely stem from a higher photon flux; this logic applies to any arrangement with a more powerful light source as well. The authors must decouple these factors and provide independent evidence for photon limitation.
3. It is unclear to me why exponential functions were used for the reactor design. This should be reasoned.
4. The motivation for the chosen Bayesian optimization algorithm is unclear to me. Why is no conventional optimization algorithm used? I must admit I'm not an expert in developing optimization routines but it appears to me that the found geometries are complicated solutions to a minor issue. The pareto front indicates only 10 % in homogeneity difference. Maybe only the reactor example is less demanding in terms of photon field than required to demonstrate the potential of the approach.
5. Fig 4c: What is the experimental error of the reaction itself? How much of the shown variance comes from the reactor and how much from the radiation field?
6. Equation 4 is presented as accounting for differences of the radiation field, but the link is unclear. The authors must explain precisely how this equation adjusts for differences and ensures that reactivity differences reflect chemical effects?
7. Page10: It is unclear whether the variations of the reaction conditions occur simultaneously or are independent. The authors should clarify this to avoid misinterpretation of the optimization results.
8. Page 10: The experimental procedure lacks sufficient detail, and the impact of spectral properties on the results is unaddressed. The authors should analyze how the different spectral properties of the light-absorbing species influence the observed outcomes.

Minor comments

- The term "distinct reaction condition" is undefined and requires precise specification of the varied parameters.
- Page 10: The statement that the study is a "quantitative mechanistic" one is misleading. The shown analysis provides qualitative insights rather than quantitative kinetics.

Literature completeness

The work by Loh et al. is missing. It covers a wide range of aspects of this manuscript under review, advocating for physics-informed machine learning (which is not addressed here). I suggest that the authors analyze their results within the scope of this work.

Loh, J. Y. Y. et al. (2024), doi:10.1002/advs.202306604 —Essential context for AI-driven multiscale physics in photocatalyst and photoreactor design.

Recommendation

The manuscript by Schuurmans et al. is a solid piece of work, reporting on an interesting coupling of ray tracing simulation with an optimizer to design efficient photoreactors. Nevertheless, considering the presented results and the available literature, this manuscript leaves me with many questions. Due to this, the significance of this work to the community is not very clear to me:

1. I miss a clear motivation for the chosen reactor design, the optimization strategy and the application case. In fact, a multitude of lab reactors has been presented in recent years.
2. Why is an indirect irradiation chosen when a circularly emitting light source is used. I would expect the direct irradiation setup with the used vial arrangement to yield a significantly higher intensity (avoids significant reflection losses) with a similar homogeneity. Small arrangements of the vial positions to the radiation field in Fig 2c would also allow to compensate for the small deviations from an ideal Lambertian emission.
3. How does this advance community work? For which types of reactor designs would this approach generate a leap forward?

Considering my previous comments, I recommend major revisions.

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 Nature Communications 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)

The manuscript describes a closed-loop workflow for photoreactor design in which a ray-tracing digital twin is calibrated experimentally, coupled to multi-objective Bayesian optimization, and then used to identify reactor geometries that balance photon capture and irradiation homogeneity. The optimized reactor is validated experimentally using a photon-limited photocatalytic Suzuki–Miyaura cross-coupling and is subsequently applied to a quantum-yield screening study.

Overall, I find the work well conceived and technically solid. The reactor design and validation are convincing, particularly the close agreement between simulated and actinometrically measured photon fluxes. The combination of calibrated ray tracing, Bayesian optimization, additive manufacturing, and experimental validation provides a useful workflow for photoreactor engineering. The chemistry used for validation is not especially diverse, and the chemical scope is therefore not the main strength of the manuscript. However, the reactor-design methodology is valuable and, in my view, suitable for publication after minor revision.

The main point that should be addressed is the framing of novelty. Ray tracing, chemical actinometry, 3D-printed photoreactors, photon-flux normalization, and Bayesian optimization have all been used previously in photochemical reactor development. The novelty of the present work is not simply “photoreactor design” or “ray-tracing-enabled characterization”, but rather the integrated and experimentally validated closed-loop co-design workflow for a multi-vial screening reactor. I recommend that the authors sharpen this distinction in the Introduction and Discussion and explicitly compare their contribution with prior photoreactor-design and photonic-characterization studies.

The quantum-yield section is useful, but the notation and interpretation should be made clearer. The values reported in Figure 5 and Table S7 appear to be quantum yields in percent, with the best cases around 0.8%, not 80%. This distinction is important, because the absorbed photon fraction is often close to quantitative, whereas the chemical quantum yield remains below 1%. The authors should ensure that the text, figure labels, and SI consistently distinguish between percentage of photons absorbed and quantum yield. The statement that all quantum yields are below 1% and are consistent with the absence of productive radical-chain propagation is reasonable, but it would benefit from a short comparison to established quantum-yield criteria in photoredox catalysis.

I would also encourage the authors to moderate some of the broader wording. The manuscript convincingly demonstrates a strong reactor-design workflow, but it does not yet establish broad chemical generality. The Suzuki–Miyaura reaction is an appropriate photon-limited benchmark, yet the substrate and reaction diversity are limited. The authors could present this more explicitly as a proof-of-concept chemical validation and mechanistic use case, rather than as a broad demonstration of general photochemical applicability.

Specific points for revision

(-) Please clarify throughout the manuscript and SI that the reported QY values such as 0.80 are 0.80%, not 80%. In Figure 5, Table S7, and the discussion, distinguish clearly between absorbed photon flux or percentage photon absorption and quantum yield.

(-) The novelty statement should be refined. Prior literature already reports 3D-printed photoreactors (including the ones from the corresponding author team I), actinometric characterization, ray-tracing/radiometric photon-flux determination, and photonic correction of screening reactors. The authors should position the contribution as an integrated, validated digital-twin/Bayesian-optimization workflow for reactor co-design.

(-) The chemical diversity demonstrated in the manuscript is limited. Ideally, the authors should include one or two additional reaction classes, for example a photoredox organocatalytic transformation, to strengthen the claim of general applicability. That said, the current scope is acceptable given that the manuscript is primarily focused on reactor design rather than reaction development. If no additional examples are added, the authors should moderate any claims of broad reaction generality and make clear that the photocatalytic Suzuki–Miyaura coupling serves mainly as a benchmark reaction for validating the reactor-design workflow.

(-) Please provide a clearer discussion of how transferable the optimized geometry is to other wavelengths, vial geometries, solvents, and reaction absorbance profiles. Since the digital twin is calibrated for a specific light source, material reflectivity, and reactor arrangement, the limits of transferability should be stated explicitly.

(-) The authors should clarify whether the optimized reactor is primarily maximizing incident photon flux, absorbed photon flux, or collected optical power, and how these quantities relate under partially absorbing reaction conditions.

(-) In the validation section, it would be helpful to include a concise error estimate for the actinometry-derived photon flux and for the simulation prediction. The agreement between ~25 and ~26 mW per vial is impressive, but the uncertainty should be reported.

(-) The manuscript would benefit from a short practical comparison: how many physical reactor iterations would likely have been required using conventional design-of-experiments or intuition-guided optimization, compared with the digital-twin workflow?

Minor wording corrections:

“facilitates quantitative understanding” should be “facilitate quantitative understanding” if the subject is “machine learning and digital twins”.

“using set simulation time” is awkward; consider “using predefined simulation time, mesh quality, and number of rays”. “programmable” in the SI should be “programmable”.

Check consistency of “aluminum”/“aluminium”.

Check consistency of “quantum yield”, “QY”, and “ Φ/ϕ ” notation across the manuscript and SI.

Subject to these revisions, I would support publication.

Version 1:

Reviewer comments:

Reviewer #1

(Remarks to the Author)

This reviewer thanks the authors for their comprehensive point-by-point response and the considerable effort invested in revising the manuscript. The manuscript has improved substantially since the original submission. In particular, the additional convergence studies, expanded validation of the digital twin, and clarification of the reactor symmetry and Bayesian optimization methodology have strengthened this manuscript. This reviewer also appreciates the revised balanced discussion of the limitations of the optimization procedure and the additional validation presented in the Supporting Information.

The authors have therefore satisfactorily addressed the majority of my technical concerns regarding the implementation of the digital twin and the robustness of the simulation framework. My original concern was not whether the optimization framework was technically correct, but whether it demonstrated that machine learning enabled fundamentally new reactor design principles. In the revised manuscript, the optimization consistently identifies vial position and distance from the light source as the dominant design variables, while reflector geometry contributes comparatively little to reactor performance. Indeed, the authors now explicitly acknowledge this conclusion in their response. While this finding is informative and quantitatively supported, it ultimately reinforces the view that the optimization largely recovers physically intuitive engineering principles rather than revealing unexpected reactor architectures or previously inaccessible design strategies.

Similarly, although the authors provide a convincing explanation for why perfect rotational symmetry is not achieved, this does not substantially alter my broader concern regarding the role of the homogeneity objective. The revised analysis demonstrates that the metric is physically meaningful rather than a numerical artefact, but it remains a relatively weak optimization objective whose influence on the final reactor design appears modest compared with simple geometric proximity to the light source.

This reviewer believes the revised manuscript now represents a technically rigorous study where the methodology is executed well and will undoubtedly be valuable to the reaction engineering community. However the principal advance is, in my opinion, an engineering optimization workflow rather than a conceptual advance in photochemical reactor design or machine-learning-enabled chemistry. The machine learning serves primarily as an optimization tool rather than enabling fundamentally new scientific insight or uncovering non-intuitive reactor design principles.

For these reasons, although I acknowledge that the manuscript has improved significantly since the original submission, my overall assessment remains unchanged; this work lacks the conceptual novelty necessary for publication in Nature Communications and would be more impactful in a specialized journal.

Reviewer #2

(Remarks to the Author)

The authors addressed all comments of the reviewers appropriately. The manuscript can be published in the current form.

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 Nature Communications initiative to facilitate training in peer review and to provide appropriate recognition for Early Career Researchers who co-review manuscripts.

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/>

June 08, 2026

Re: Submission of a revised version of “Closed-Loop Photoreactor Design Enabled by Machine Learning and Digital Twins” to Nature Communications [Manuscript No. NCOMMS-26-013015].

Dear Referees

We would like to thank you all for the careful evaluation of our manuscript and for the constructive feedback provided throughout the review process. We have carefully considered all comments and believe that they have helped us to improve both the clarity and the rigor of the work.

In response to the reviewers' suggestions, we have revised the manuscript extensively. The main changes include a more detailed discussion of the physical basis of the homogeneity metric, additional convergence and validation analyses to further support the robustness of the digital twin and optimization framework, clarification of the scope and terminology used throughout the manuscript, and an improved positioning of the novelty and broader applicability of the workflow.

A detailed point-by-point response to all reviewer comments is provided below. We hope that the revisions adequately address the concerns raised and that the manuscript is now suitable for publication.

Sincerely,

Tim, on behalf of all authors.

Reviewer #1 (Remarks to the Author)

This manuscript by Noël and co-workers presents a workflow combining ray-tracing simulations, Bayesian optimization, and additive manufacturing for the design of batch photoreactors. The authors build a detailed digital twin of a multi-vial photoreactor, calibrated against actinometric measurements, and apply multi-objective Bayesian optimization to identify reactor geometries that maximize photon flux into vials whilst maintaining irradiation homogeneity. The optimized reactor is subsequently validated experimentally and applied to a quantum yield screening study. The work described in this manuscript is technically competent, clearly written, and addresses an important challenge in photochemical reactor design. However, this reviewer does have several concerns about the formulation and interpretation of the optimization problem and the strength of several conceptual claims. In its current form, the manuscript overstates the role and impact of machine learning-enabled optimization and, more importantly, appears to base key conclusions on claims that may not be defensible given their current methodologies.

My consideration and comments are outlined below.

1. The first relates to the definition of the homogeneity objective. The authors define irradiation homogeneity as the ratio of minimum to maximum photon flux across vials:

$$H = q_{(i,max)} / q_{(i,min)}$$

All geometric parameters explored in the optimization enforce strict rotational symmetry; (i) symmetric placement of all vials (identical radial position and depth); (ii) a rotationally symmetric reflector generated by revolved curves; (iii) symmetric light source arrangement; and (iv) predominantly diffuse reflection from the internal aluminum foil.

Under these assumptions, the physical system is invariant under rotation, and all vials are physically equivalent. In an idealized model (such as that simulated), inter-vial photon flux should therefore be identical by construction. The homogeneity across vials must therefore be unity regardless of the reflector geometry selected ($H=1$ always). Any deviation of H from unity is unlikely to reflect genuine physical inhomogeneity and must instead arise from numerical artifacts of the ray-tracing simulation (finite ray counts, meshing error, surface discretization, or stochastic noise). The “homogeneity” objective is dominated by numerical error and does not relate to reactor performance. This explains several observations: (i) Homogeneity barely improves during optimization; (ii) Scatter dominates the homogeneity axis; (iii) Pareto fronts are dense and shallow; (iv) Best “homogeneity” designs have the smoothest, simplest geometries (those expected to minimize numerical variance in the simulations), and (v) Near identical performance between reactor α and γ in real experiments

In my view, the homogeneity metric would more accurately be described as a measure of numerical robustness for the simulation as opposed to a measure of the uniformity of photon delivery. This significantly weakens the multi-objective optimization framework and casts doubt on the interpretation of the Pareto fronts and trade-offs presented. If the homogeneity metric instead measures the uniformity of delivery to a single vial (i.e. variance across the surface of a single vial) then the manuscript needs to be altered to make this clear. Homogeneity cannot be measured by comparing photon flux between different vials in the presented system.

Answer: We respectfully disagree with the reviewer's assertion that all vials are physically equivalent by construction and that the homogeneity must therefore be unity irrespective of reactor geometry.


While both the light source and vial arrangement exhibit rotational symmetry, their symmetry orders differ (C6 for the LED source and C16 for the vial arrangement). Consequently, the individual vials do not occupy equivalent optical environments. Vials located closer to one of the six LED chips are expected to receive a different photon flux than vials positioned between two LEDs. As illustrated below, the combination of the C6 light source and C16 vial arrangement gives rise to five distinct vial positions. In addition, the 3D-printed support structure required to mount the light source is located in close proximity to the LEDs and inevitably interacts with emitted light, introducing additional anisotropy and resulting in nine distinct vial positions. A discussion of these symmetry considerations has been added to the Supporting Information.

Although a reasonable degree of homogeneity was expected based on the approximate rotational symmetry of the reactor and the diffuse nature of the reflector, the extent of this homogeneity cannot be determined a priori and must be quantified through simulation or experiment.

To ensure that the observed homogeneity differences do not originate from numerical artefacts, we performed convergence studies with respect to simulation time, number of photons, and mesh elements (further discussed below and added to the Supporting Information). These studies support that the predicted homogeneity values are stable with respect to numerical refinement.

Importantly, the optimization consistently steers away from designs with low homogeneity (<0.90), indicating that the Bayesian optimization framework identifies and exploits physically meaningful trends in the simulated data. We agree that comparing homogeneity values that differ only marginally (e.g., 0.83 versus 0.86) may be of limited practical significance. However, the range observed throughout the optimization campaign spans approximately 0.8–1.0, where differences become meaningful in the context of reactor performance. This interpretation is further supported by the Spearman analysis, which shows a clear correlation between homogeneity and vial distance from the light source. As expected, increasing this distance improves the uniformity of photon delivery at the expense of photon flux, and the optimization framework successfully quantifies this trade-off.





2. The second concern relates to the reliance on simulated data without robust validation during optimization. The Bayesian optimization is driven almost entirely by simulated photon flux data. Whilst the authors validate the digital twin against a single arbitrary geometry, this validation does not establish reliability over the full design space explored during optimization. Convergence testing is performed only post-hoc on three selected geometries with no a priori justification that these geometries span the numerical difficulty of the search space. There is no systematic assessment of numerical accuracy (i) across the optimization trajectory, (ii) for geometries with sharper curvature or stronger optical focusing, or (iii) with respect to preservation of relative rankings under mesh or ray-count refinement. Numerical accuracy should be demonstrated where the optimizer operates, not after the fact on selected exemplars. Because the optimization itself appears to suppress numerically challenging geometries, post-hoc convergence on the final designs does not validate the simulation as a reliable objective function. If geometry-dependent numerical errors are present, even at modest levels, the optimization trajectory and resulting optima may be dominated by numerical bias rather than physical signal. In a simulation-driven Bayesian optimization framework, this omission is significant. If the homogeneity metric was measured between vials, then its value gives a good measure of numerical error. The levels this would indicate are not suitable for simulation driven optimization.

Answer: We agree with the reviewer that, in a simulation-driven optimization framework, the reliability of the objective function must be demonstrated across the design space and not solely for the final optimized geometries.

To address this point, we have expanded the discussion of the simulation settings and convergence analysis in both the manuscript and Supporting Information. The simulation time, number of photons, and meshing parameters were initially selected based on convergence studies performed. Importantly, we deliberately chose simulation settings that exceeded the

minimum requirements identified from these studies, thereby providing an additional safety margin and ensuring convergence throughout the explored design space.

To further support the robustness of the objective function, we have now included convergence studies for additional reactor geometries, including the design exhibiting the lowest homogeneity and therefore the greatest sensitivity to potential numerical artefacts. Across all investigated cases, the objective values showed only minor changes upon further refinement of simulation time, number of rays, and number of mesh elements, indicating that the reported photon fluxes and homogeneity values are stable with respect to simulation settings.

Finally, Bayesian optimization inherently accounts for uncertainty in objective evaluations through its surrogate model. While the residual simulation uncertainty in our case is small (approximately 3% of the objective value), this uncertainty is naturally incorporated into the optimization process and is not expected to affect the overall optimization trajectory or conclusions.

3. The third major consideration optimization problem collapses to a trivial design choice. Analysis of the optimization trajectory and results indicate that photon flux is largely dominated by vial proximity and insertion depth, with the Bayesian optimizer consistently pushing these parameters to their bounds. This behavior is entirely intuitive and consistent with basic radiometric considerations. In contrast, most reflector parameters show weak correlations with performance and broad, poorly defined distributions. The optimization therefore appears to collapse onto a trivial geometric trade-off (maximizing proximity to the light source while maintaining spacing), with reflector curvature acting only as a minor secondary effect under predominantly diffuse optical conditions. This is entirely expected if the homogeneity metric is not physically informative. The geometric designs were likely not given a sufficient degree of freedom to meaningfully contribute to the limited. Given this collapse, it is difficult to conclude that meaningful geometric optimization has occurred beyond what would be achieved via straightforward heuristic design.

Answer: We agree with the reviewer that the optimization identifies vial position and the distance between the vial and light source as dominant factors governing photon collection. In hindsight, this trend is physically intuitive and consistent with basic radiometric considerations. However, the purpose of the optimization is not merely to identify qualitative trends, but to quantitatively determine the trade-off between photon flux and irradiation homogeneity across a large design space.

While the optimization revealed that vial placement is the dominant factor, it also showed that reflector geometry has a less pronounced influence than initially anticipated. We consider this finding informative rather than trivial. Prior to this study, it was not clear whether the detailed reflector shape would strongly influence reactor performance, especially given the complexity of multiple reflections and the large number of geometric parameters involved. The optimization indicates that, under the diffuse reflection conditions used here, the contribution of reflector geometry is secondary compared to vial positioning.

We would also like to note that the apparent lack of importance of individual reflector parameters should not be interpreted as evidence that reflector geometry is irrelevant. Different combinations of the underlying design parameters can generate similar reflector shapes, making it difficult to infer geometric importance directly from parameter correlations. For this reason, parameter importance does not necessarily translate directly into geometric importance.

More broadly, the optimization provides a quantitative framework for identifying where the balance between photon flux and homogeneity is achieved. Although some of the observed trends may appear intuitive in hindsight, the corresponding design choices are far from uniformly adopted in the photoreactor literature, as evidenced by the large diversity of reported reactor geometries (e.g. refs 1-8). In particular, we observe that placing a diffuse reflector in close proximity to the vials can improve performance while maintaining acceptable homogeneity, a design feature that is not commonly encountered in either academic or commercial photoreactors.

We have revised the manuscript to better emphasize that the primary outcome of the optimization is the quantification of these trade-offs and the identification of design principles, rather than the discovery of a single unexpected geometric feature.

1. T. M. Masson, S. D. A. Zondag, J. H. A. Schuurmans, T. Noël, *React. Chem. Eng.* **2024**, 9, 2218–2225.
2. F. Schiel, C. Peinsipp, S. Kornigg, D. Böse, *ChemPhotoChem* **2021**, 5, 431–437.
3. S. Rossi, M. V. Dozzi, A. Puglisi, M. Pagani, *Chem. Teach. Int.* **2020**, 2, DOI 10.1515/cti-2019-0010.
4. M. R. Penny, S. T. Hilton, *J. Flow Chem.* **2023**, 13, 435–442.
5. T. Aubineau, J. Laurent, L. Olanier, A. Guérinot, *Chemistry–Methods* **2023**, DOI 10.1002/cmtd.202300002.
6. S. Hedquist, J. López, A. M. Moore, A. Asurumunige, D. F. Fernández, M. González-Esguevillas, *Synthesis (Stuttg.)* **2026**, 58, 1155–1160.
7. R. M. Bailey, M. Martínez-Aguirre, B. Schaefer, M. R. Crimmin, P. W. Miller, *React. Chem. Eng.* **2026**, 11, 228–233.
8. C. “Chip” Le, M. K. Wismer, Z.-C. Shi, R. Zhang, D. V. Conway, G. Li, P. Vachal, I. W. Davies, D. W. C. MacMillan, *ACS Cent. Sci.* **2017**, 3, 647–653.

4. The final consideration concerns terminology. The manuscript repeatedly describes the workflow as “closed-loop” and “autonomous.” However, the optimization loop is closed entirely within simulation, with experimental data used only for one-time calibration and post-optimization validation. In the context of autonomous chemistry and reactor optimization, “closed loop” typically denotes iterative experimental feedback informing the optimization in real time. That is not the case here as the work represents offline, simulation-based optimization with experimental confirmation. While this is still valuable, the terminology overstates the degree of autonomy and robustness, particularly given the central reliance on a fixed digital twin whose accuracy is not continuously constrained by experiment.

Answer: We appreciate the reviewer's distinction and agree that, within the autonomous chemistry community, the term “closed-loop” is often associated with experimental feedback-driven optimization, where experimental measurements are iteratively used to guide subsequent experiments.

In the present work, the optimization loop is instead closed through an experimentally validated digital twin. The optimization proceeds autonomously by generating reactor geometries, evaluating their performance through ray-tracing simulations, and feeding the resulting objective values back to the Bayesian optimization algorithm without human intervention. In this sense, the workflow operates as a closed-loop optimization framework, albeit one based on simulation feedback rather than real-time experimental feedback.

To ensure that the simulation feedback is physically meaningful, the underlying model was extensively calibrated using radiometric measurements and subsequently validated experimentally. The close agreement between simulated and measured photon fluxes suggests that the physics-based model can predict photonic performance with sufficient accuracy to guide reactor optimization. Because the model is based on physics-based descriptions of light emission and optical interactions, successful validation for geometries involving multiple reflections supports its applicability across the explored design space.

That said, we agree that our workflow differs from experimentally autonomous platforms in which the digital model is continuously updated using newly acquired experimental data. To avoid ambiguity, we have revised the manuscript to more clearly distinguish simulation-driven closed-loop optimization from experimentally feedback-driven autonomous workflows and to better define the role of the digital twin within the optimization framework.

Overall, the authors demonstrate an interesting attempt at incorporating ML into reactor design using simulated digital twins. However, the optimization problem as formulated is physically under-constrained, and one of the central objectives appears to lack meaningful physical content due to enforced symmetry. Combined with the heavy reliance on simulation and the absence of robust validation during optimization, this impedes the conceptual claims regarding machine-learning-guided reactor design. In my view, the work would be more appropriately as a simulation-assisted reactor design study rather than a demonstration of closed-loop, multi-objective optimization enabled by machine learning. Major revisions addressing the above concerns would be required to realign the claims with what is actually demonstrated in this manuscript.

Answer: We thank the reviewer for their thoughtful assessment of the work and for raising important questions regarding the interpretation of the optimization framework, the role of the homogeneity objective, and the validation of the digital twin.

In response to these concerns, we have substantially revised the manuscript to clarify the homogeneity metric, expand the discussion of simulation convergence and numerical robustness, and more carefully position the workflow within the context of simulation-driven reactor design and optimization. We have also moderated and refined the terminology used throughout the manuscript to clearly distinguish our approach from experimentally feedback-driven autonomous optimization platforms.

While we acknowledge that the optimization relies on an experimentally calibrated and validated digital twin rather than iterative experimental feedback, we believe that the combination of physics-based ray-tracing simulations, multi-objective Bayesian optimization, additive manufacturing, and experimental validation constitutes a meaningful framework for photoreactor co-design. The objective of the work is not to replace experimental optimization, but rather to demonstrate how experimentally validated simulations can be integrated with machine learning to efficiently explore large reactor design spaces and identify improved reactor geometries.

We respectfully disagree with the assertion that the optimization problem is physically under-constrained or that the homogeneity objective lacks physical meaning. As discussed above, the vial positions are not optically equivalent due to the mismatch between the symmetries of the light source and vial arrangement, resulting in differences in photon delivery that can be quantified and optimized. Furthermore, the additional convergence studies and experimental

validations included in the revised manuscript support that the optimization is driven by physical trends rather than numerical artefacts.

Overall, we have revised the manuscript extensively to better align our claims with the scope of what is demonstrated. We hope that the additional analyses, clarifications, and revised framing adequately address the reviewer's concerns and better convey the contribution of the work.

Reviewer #2 (Remarks to the Author)

Photoreactor design requires balancing photon capture with irradiation uniformity. This work integrates ray-tracing digital twins with multi-objective Bayesian optimization to autonomously design reactor geometries. The framework is validated against potassium ferrioxalate actinometry and applied to a photocatalytic Suzuki–Miyaura cross-coupling, demonstrating predictive accuracy and enabling quantitative mechanistic insights across diverse conditions.

The manuscript presents a closed-loop workflow integrating ray-tracing digital twins with Bayesian optimization to autonomously design photoreactors. While digital twins and optimization are established in other fields, their specific application to multi-objective photoreactor co-design for uniform irradiation represents a methodological advance, which is also experimentally validated using actinometry and a photocatalytic reaction.

Strengths

- The digital twin is rigorously validated against potassium ferrioxalate actinometry, ensuring reliable simulation of photon fluxes.
- The closed-loop integration of ray-tracing with Bayesian optimization effectively navigates design spaces for reactor geometry.
- The framework successfully predicts experimental reaction yields in a photocatalytic Suzuki–Miyaura cross-coupling, confirming virtual reactor accuracy.

Major comments

1. The claim of high scientific novelty regarding the predictive accuracy of the virtual reactor requires stronger justification.

Answer: We appreciate the reviewer's comment and agree that the novelty does not lie in any individual component of the workflow. Ray-tracing simulations, Bayesian optimization, additive manufacturing, and experimental photoreactor characterization have all been reported previously, including by our group.

The novelty of the present work lies in the integration of these elements into an experimentally validated workflow for the co-design of photoreactors. To the best of our knowledge, similarly systematic optimization strategies for photoreactor design remain unprecedented, particularly those combining a physics-based digital twin with multi-objective optimization and subsequent experimental validation.

To better reflect this distinction, we have revised the Introduction and Discussion to more clearly position our contribution relative to previous photoreactor-design and photonic-characterization studies.

2. Page 8: "Together, these results demonstrate that the digital twin accurately predicts photonic performance and that the optimized geometries translate directly into improved and more uniform photochemical outcomes." This point comes from a circular argumentation line. The authors state that they operate in the photon limited regime. Thus, performance linearly depends on the photon flux and with the realized circular arrangement of vials in relation to the circularly emitting light source, only deviations from the uniformity would be scientifically surprising.

Higher conversions solely stem from a higher photon flux; this logic applies to any arrangement with a more powerful light source as well. The authors must decouple these factors and provide independent evidence for photon limitation.

Answer: Our intention was not to claim that improved reaction performance constitutes independent proof of the digital twin. Rather, the validation demonstrates that the simulated differences in incident photon flux translate into corresponding differences in photochemical performance under the investigated reaction conditions. Because the reaction was deliberately performed under photon-limited conditions, variations in photon availability can be directly related to variations in conversion, making it possible to assess whether the predicted reactor-to-reactor differences are reflected experimentally.

We agree that increasing the power of the light source would also increase conversion under photon-limited conditions. However, the objective of this work is not simply to maximize conversion through higher light intensity, but rather to improve the efficiency with which photons are delivered to the reaction volume. The optimized reactor geometries achieve higher photon fluxes at the vial positions without increasing the power consumption of the light source, thereby improving photonic efficiency.

Regarding the reviewer's symmetry argument, the vial positions are not optically equivalent despite the rotational symmetry of the setup. The symmetry of the light source (C6) differs from that of the vial arrangement (C16), resulting in distinct optical environments for different vial positions. Furthermore, the light-source support structure introduces additional anisotropy. A discussion of these symmetry considerations has been added to the Supporting Information.

3. It is unclear to me why exponential functions were used for the reactor design. This should be reasoned.

Answer: Exponential curves were chosen because they provide smooth reflector geometries while still allowing a broad range of shapes to be generated from a limited number of parameters. This reduced the risk of producing poorly manufacturable features, while keeping the design space sufficiently flexible for optimization. This rationale has now been added to the manuscript.

4. The motivation for the chosen Bayesian optimization algorithm is unclear to me. Why is no conventional optimization algorithm used? I must admit I'm not an expert in developing optimization routines but it appears to me that the found geometries are complicated solutions to a minor issue. The pareto front indicates only 10 % in homogeneity difference. Maybe only the reactor example is less demanding in terms of photon field than required to demonstrate the potential of the approach.

Answer: Bayesian optimization was chosen because the optimization problem considered here possesses several characteristics that make conventional optimization approaches less suitable.

First, each objective-function evaluation requires a ray-tracing simulation, making exhaustive exploration of the design space impractical. Second, the objective functions are not readily differentiable due to the complex nature of light transport, including diffuse and specular reflections, and refraction, which are mostly surface phenomena. As a result, gradient-based optimization approaches are not straightforward to apply. Bayesian optimization is specifically designed for such expensive black-box optimization problems and can efficiently identify promising regions of the design space with a limited number of evaluations.

We agree that, for the present case study, the optimization of homogeneity is less pronounced than that of photon flux. Nevertheless, the algorithm successfully identifies and quantifies the trade-off between these competing objectives. More importantly, the purpose of this work is not to optimize a single reactor design, but to demonstrate a general workflow that can be applied to more complex photoreactor configurations and alternative optimization objectives. We have clarified this motivation in the revised manuscript.

5. Fig 4c: What is the experimental error of the reaction itself? How much of the shown variance comes from the reactor and how much from the radiation field?

Answer: Based on additional experiments, an experimental error of ~10% was found, this has been added in the manuscript and put in context in terms of comparing homogeneities.

6. Equation 4 is presented as accounting for differences of the radiation field, but the link is unclear. The authors must explain precisely how this equation adjusts for differences and ensures that reactivity differences reflect chemical effects?

Answer: Further explanation of equation 4 has been added.

7. Page 10: It is unclear whether the variations of the reaction conditions occur simultaneously or are independent. The authors should clarify this to avoid misinterpretation of the optimization results.

Answer: We have clarified in the manuscript that the reactions are performed simultaneously, with each vial containing a different reaction condition. Since all vial positions are individually characterized, a single experiment can be used to evaluate up to 16 distinct reaction conditions in parallel.

8. Page 10: The experimental procedure lacks sufficient detail, and the impact of spectral properties on the results is unaddressed. The authors should analyze how the different spectral properties of the light-absorbing species influence the observed outcomes.

Answer: Additional details regarding the spectral properties of the system have been referenced more clearly in the manuscript. The Supporting Information provides the wavelength-dependent absorbance of the reaction mixtures, the emission spectrum of the light source, and the effective optical path lengths.

These data allow the spectral overlap between the light source and absorbing species to be assessed and provide the basis for interpreting the observed reaction outcomes. Since the optimization itself is based on the incident photon flux delivered to the reaction volume, the reactor-design conclusions are independent of the specific absorbance characteristics of the investigated reaction.

Minor comments

- The term "distinct reaction condition" is undefined and requires precise specification of the varied parameters.

Answer: A reference to Figure 5 has been added.

- Page 10: The statement that the study is a "quantitative mechanistic" one is misleading. The shown analysis provides qualitative insights rather than quantitative kinetics.

Answer: Our use of the term "quantitative" refers to the determination of quantum yields under each reaction condition. By combining reaction outcomes with absorbed photon fluxes, we are able to normalize for differences in photon availability and directly compare the photochemical efficiency of the various conditions. In this sense, the analysis extends beyond a qualitative reaction screen and provides quantitative information on the efficiency with which absorbed photons are converted into product.

Literature completeness

The work by Loh et al. is missing. It covers a wide range of aspects of this manuscript under review, advocating for physics-informed machine learning (which is not addressed here). I suggest that the authors analyze their results within the scope of this work. Loh, J. Y. Y. et al. (2024), doi:10.1002/advs.202306604 —Essential context for AI-driven multiscale physics in photocatalyst and photoreactor design.

Answer: The suggested literature has been cited in the manuscript.

Recommendation

The manuscript by Schuurmans et al. is a solid piece of work, reporting on an interesting coupling of ray tracing simulation with an optimizer to design efficient photoreactors. Nevertheless, considering the presented results and the available literature, this manuscript leaves me with many questions. Due to this, the significance of this work to the community is not very clear to me:

1. I miss a clear motivation for the chosen reactor design, the optimization strategy and the application case. In fact, a multitude of lab reactors has been presented in recent years.

Answer: The primary objective of this work is not to introduce another photoreactor design, but rather to demonstrate a systematic framework for the simulation-based optimization of photoreactors. While a large number of laboratory photoreactors have been reported in recent years, their design is often guided by intuition or practical constraints rather than formal optimization. We aim to show that a physics-based digital twin can be used to efficiently evaluate and optimize a large number of candidate designs prior to fabrication and testing.

Bayesian optimization was selected because it is well suited for black-box optimization with a large design space. The chosen 16-vial photoreactor serves as a representative case study that combines practical relevance with sufficient geometric complexity to demonstrate the workflow.

We have revised the Introduction and Discussion to better communicate these motivations and to place the work in the broader context of photoreactor engineering.

2. Why is an indirect irradiation chosen when a circularly emitting light source is used. I would expect the direct irradiation setup with the used vial arrangement to yield a significantly higher intensity (avoids significant reflection losses) with a similar homogeneity. Small arrangements of

the vial positions to the radiation field in Fig 2c would also allow to compensate for the small deviations from an ideal Lambertian emission.

Answer: Although the LED chips exhibit approximately Lambertian emission profiles, their emission remains directional. Direct irradiation of 16 surrounding vials would therefore require careful matching between the emission pattern and vial positions to efficiently distribute light across all reaction vessels. In practice, this would introduce a similar optimization problem to the one addressed in the present work.

We chose a top-irradiated configuration with the vials suspended within a reflective chamber. This design allows both direct and reflected light to contribute to photon delivery, while providing greater flexibility to redistribute photons throughout the reactor volume. In addition, positioning the light source above the reactor facilitates heat dissipation and simplifies thermal management of the LED assembly.

We agree that alternative reactor configurations, including direct bottom or side irradiation, could be explored and optimized using the same workflow. The objective of the present study is not to demonstrate the superiority of a specific reactor architecture, but rather to establish a general framework for the simulation-based optimization of photoreactors. We have clarified this point in the revised manuscript.

3. How does this advance community work? For which types of reactor designs would this approach generate a leap forward?

Answer: The goal of this work is to provide a general workflow that combines physics-based simulations with machine learning to optimize photoreactor designs for a specific application.

While demonstrated here using a 16-vial screening reactor, the approach is not limited to this particular geometry. We anticipate the greatest benefit for reactor systems in which photon delivery plays a critical role and where design decisions involve complex trade-offs between photon flux, irradiation uniformity, energy efficiency, throughput, and reactor geometry. Examples include parallel and high-throughput photochemical screening platforms, where irradiation homogeneity is essential, as well as larger-scale process reactors, where maximizing photon utilization and minimizing energy consumption become increasingly important.

Traditionally, photoreactors are often designed based on intuition, practical constraints, or iterative experimentation, while systematic optimization remains challenging due to the complexity of light transport and the large design space. We aim to demonstrate that an experimentally validated digital twin can be used to efficiently explore and optimize this design space prior to fabrication.

Considering my previous comments, I recommend major revisions.

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 Nature Communications initiative to facilitate training in peer review and to provide appropriate recognition for Early Career Researchers who co-review manuscripts.

Answer: We thank the reviewer for their time to evaluate the manuscript.

Reviewer #4 (Remarks to the Author):

The manuscript describes a closed-loop workflow for photoreactor design in which a ray-tracing digital twin is calibrated experimentally, coupled to multi-objective Bayesian optimization, and then used to identify reactor geometries that balance photon capture and irradiation homogeneity. The optimized reactor is validated experimentally using a photon-limited photocatalytic Suzuki–Miyaura cross-coupling and is subsequently applied to a quantum-yield screening study.

Overall, I find the work well conceived and technically solid. The reactor design and validation are convincing, particularly the close agreement between simulated and actinometrically measured photon fluxes. The combination of calibrated ray tracing, Bayesian optimization, additive manufacturing, and experimental validation provides a useful workflow for photoreactor engineering. The chemistry used for validation is not especially diverse, and the chemical scope is therefore not the main strength of the manuscript. However, the reactor-design methodology is valuable and, in my view, suitable for publication after minor revision.

The main point that should be addressed is the framing of novelty. Ray tracing, chemical actinometry, 3D-printed photoreactors, photon-flux normalization, and Bayesian optimization have all been used previously in photochemical reactor development. The novelty of the present work is not simply “photoreactor design” or “ray-tracing-enabled characterization”, but rather the integrated and experimentally validated closed-loop co-design workflow for a multi-vial screening reactor. I recommend that the authors sharpen this distinction in the Introduction and Discussion and explicitly compare their contribution with prior photoreactor-design and photonic-characterization studies.

The quantum-yield section is useful, but the notation and interpretation should be made clearer. The values reported in Figure 5 and Table S7 appear to be quantum yields in percent, with the best cases around 0.8%, not 80%. This distinction is important, because the absorbed photon fraction is often close to quantitative, whereas the chemical quantum yield remains below 1%. The authors should ensure that the text, figure labels, and SI consistently distinguish between percentage of photons absorbed and quantum yield. The statement that all quantum yields are below 1% and are consistent with the absence of productive radical-chain propagation is reasonable, but it would benefit from a short comparison to established quantum-yield criteria in photoredox catalysis.

I would also encourage the authors to moderate some of the broader wording. The manuscript convincingly demonstrates a strong reactor-design workflow, but it does not yet establish broad chemical generality. The Suzuki–Miyaura reaction is an appropriate photon-limited benchmark, yet the substrate and reaction diversity are limited. The authors could present this more explicitly as a proof-of-concept chemical validation and mechanistic use case, rather than as a broad demonstration of general photochemical applicability.

Specific points for revision

(-) Please clarify throughout the manuscript and SI that the reported QY values such as 0.80 are 0.80%, not 80%. In Figure 5, Table S7, and the discussion, distinguish clearly between absorbed photon flux or percentage photon absorption and quantum yield.

Answer: The manuscript and SI have been adjusted to clarify that it is a percentage, used to depict the quantum yield.

(-) The novelty statement should be refined. Prior literature already reports 3D-printed photoreactors (including the ones from the corresponding author team !), actinometric characterization, ray-tracing/radiometric photon-flux determination, and photonic correction of screening reactors. The authors should position the contribution as an integrated, validated digital-twin/Bayesian-optimization workflow for reactor co-design.

Answer: We adjusted the introduction and outlook accordingly.

(-) The chemical diversity demonstrated in the manuscript is limited. Ideally, the authors should include one or two additional reaction classes, for example a photoredox organocatalytic transformation, to strengthen the claim of general applicability. That said, the current scope is acceptable given that the manuscript is primarily focused on reactor design rather than reaction development. If no additional examples are added, the authors should moderate any claims of broad reaction generality and make clear that the photocatalytic Suzuki–Miyaura coupling serves mainly as a benchmark reaction for validating the reactor-design workflow.

Answer: The chosen reaction is indeed a benchmark, primarily to show real-life application of the photoreactors. Since we optimize based on incident photon flux, any chemistry that is photon-limited should support the simulation outcome. We changed the manuscript to convey this message more clearly.

(-) Please provide a clearer discussion of how transferable the optimized geometry is to other wavelengths, vial geometries, solvents, and reaction absorbance profiles. Since the digital twin is calibrated for a specific light source, material reflectivity, and reactor arrangement, the limits of transferability should be stated explicitly.

Answer: The simulation framework is physics-based and therefore remains applicable provided that the relevant optical properties are adequately described. Consequently, the optimized geometry is expected to remain broadly transferable to systems with similar light-source emission characteristics and material reflection/refraction properties. If these properties change significantly, for example through the use of a different light-source architecture, reflector material, or vial geometry, the optimization should be repeated using the corresponding optical parameters.

It is important to note that the optimization itself is based on the incident photon flux delivered to the reaction volume. As such, the optimized geometry is not directly coupled to the absorbance characteristics of a specific reaction mixture. However, reaction absorbance profiles, solvent properties, and optical path lengths will influence the fraction of incident photons that is ultimately absorbed and utilized by a given photochemical system.

(-) The authors should clarify whether the optimized reactor is primarily maximizing incident photon flux, absorbed photon flux, or collected optical power, and how these quantities relate under partially absorbing reaction conditions.

Answer: It has been clarified in the manuscript that the optimization is on the incident photon flux, that directly correlates to the collected optical power (corrected for the energy of the light).

Given the extensive characterization done on the photoreactor, the absorbed photon flux can be determined even under partially absorbing conditions (Equation 3).⁹ This is achieved by combining the simulation results with chemical actinometry.

9. S. D. A. Zondag, J. H. A. Schuurmans, A. Chaudhuri, R. P. L. Visser, C. Soares, N. Padoin, K. P. L. Kuijpers, M. Dorbec, J. van der Schaaf, T. Noël, Nat. Chem. Eng. 2024, 1, 462–471.

(-) In the validation section, it would be helpful to include a concise error estimate for the actinometry-derived photon flux and for the simulation prediction. The agreement between ~25 and ~26 mW per vial is impressive, but the uncertainty should be reported.

Answer: For the experimental measurements, potassium ferrioxalate actinometry is known to provide low uncertainty when performed according to established procedures.¹⁰ In addition, photon fluxes were determined from kinetic measurements, where a separate experiment was performed for each sampling time point, further improving the robustness of the analysis. The standard deviation between vial positions has now been added to the Supplementary Information (Section 2.3) and is reported alongside the measured photon fluxes.

For the simulations, we have added a dedicated convergence analysis (Section 2.4) evaluating the effects of simulation time, number of rays, and number of mesh elements. These studies provide an estimate of the numerical uncertainty associated with the predicted photon fluxes and demonstrate that the simulation results are stable with respect to further refinement.

10. B. Wriedt, D. Ziegenbalg, J. Flow Chem. 2020, 10, 295–306.

(-) The manuscript would benefit from a short practical comparison: how many physical reactor iterations would likely have been required using conventional design-of-experiments or intuition-guided optimization, compared with the digital-twin workflow?

Answer: While such a comparison would certainly be interesting, we believe it is difficult to make in a rigorous and general manner, as the outcome would strongly depend on the specific reactor design problem, the chosen optimization strategy, and the implementation details of the optimization procedure.

The primary objective of this work is not to compare different optimization methodologies, but rather to demonstrate how an experimentally validated digital twin can be integrated into an optimization workflow for photoreactor design. In principle, the digital twin could be coupled to a variety of optimization strategies, including design-of-experiments approaches, evolutionary algorithms, or Bayesian optimization. We selected Bayesian optimization because it is particularly well suited for black-box optimization problems with a large design space.

Minor wording corrections:

“facilitates quantitative understanding” should be “facilitate quantitative understanding” if the subject is “machine learning and digital twins”.

“using set simulation time” is awkward; consider “using predefined simulation time, mesh quality, and number of rays”.

“programmable” in the SI should be “programmable”.

Check consistency of “aluminum”/“aluminium”.

Check consistency of “quantum yield”, “QY”, and “ Φ/ϕ ” notation across the manuscript and SI.

Answer: These corrections (or similar) have been applied in the manuscript and supporting information.

Subject to these revisions, I would support publication.