

Net carbon losses in Central African forests revealed by high-resolution biomass change maps

Corresponding Author: Dr Liang Wan

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)

Considering that I have used several notations and formulas in my review, I have uploaded a document titled "Remarks for the Authors" (remarks_to_the_authors.pdf).

(Remarks on code availability)

I did not have access to the code and data.

Reviewer #2

(Remarks to the Author)

Key results

The key results of this paper are that the majority of the net carbon gains and losses in the Tropical Moist Forest in the Congo basin occur in very small patches, but because there are so many of them the overall impact on the carbon balance is substantial.

Validity

For most parts, the analysis methods and data interpretation are robust and the conclusions are justified by the results.

However, I have one major concern that the authors should address before the paper can be published:

Lines 111-114: One of the main points of novelty of this paper is, as the authors argue in the abstract, the 10 m resolution of the canopy height and 30 m resolution of the biomass maps. Near lines 111-114, the paper should provide the comparison of canopy height change on the 30 m spatial resolution in addition to the aggregated analysis that is currently here (this is currently provided in Supplementary Figures 10 and 11).

Supplementary Figure 11 suggests that there is no (or almost no) correlation between the difference in canopy height from GEDI and the deep learning model at 30 m resolution. Is RH100 a good choice of GEDI product or would you find a better relationship if you had used RH99 or RH95, for example?

Significance

The potential significance of the conclusions of the paper lies in the high importance of small-scale deforestation, forest degradation and regrowth patches in Tropical Moist Forests in Africa for the overall carbon balance. This is not recognised and addressed by most current studies of the carbon balance of Africa.

Line 83: This part of the review of previous work omits other studies that have used higher resolution biomass change maps, e.g. Rodríguez-Veiga et al. at 100 m resolution. Relevant studies of higher resolution biomass change maps should be reviewed and cited here briefly.

Reference:

Rodríguez-Veiga et al. "Loss of tropical moist broadleaf forest has turned Africa's forests from a carbon sink into a source." Scientific Reports 15.1 (2025): 41744.

Data and methodology

My opinion is that the approach of using deep learning to estimate canopy height from Sentinel-1 and 2 trained on GEDI RH100 data is valid. The quality of the presentation of the data and results is excellent.

Analytical approach

The analytical approach and statistical tests applied follow best practice in the discipline.

Suggested improvements

Line 37: The abstract mentions some causes of deforestation in the Congo rainforest but fails to mention mining, which I suggest should be added.

Line 39: In the abstract, it is also not clear how the authors get from canopy height maps to estimates of carbon gains and losses. This needs to be made clear here.

Line 39 also states the spatial resolution of 10 m for the canopy height maps. This suggests to the reader that the paper presents an analysis at that scale, when in fact the main analysis was done at 30 m resolution with the biomass maps. This needs to be stated in the abstract for clarity.

Line 40: The abstract is unclear as to which Sentinel satellite data were used exactly, presumably Sentinel-1 and 2. This needs to be stated here.

Lines 69-74: The review of existing approaches to carbon accounting in Africa omits the RECCAP and RECCAP2 carbon budgets, which use not only models but also Earth observation data. Relevant papers should be included and cited here for completeness.

Line 83: This part of the review of previous work omits other studies that have used higher resolution biomass change maps, e.g. Rodríguez-Veiga et al. at 100 m resolution. Relevant studies of higher resolution biomass change maps should be reviewed and cited here briefly.

Lines 111-114: One of the main points of novelty of this paper is, as the authors argue, the 10 m resolution of the canopy height and biomass maps. However, the analysis was done at 30 m resolution, even though the paper presents evidence that there is no correlation between the canopy height differences from airborne LiDAR and the data produced by the authors at that spatial resolution. This needs to be fully addressed before the paper can be published.

Lines 252-253 contain an incomplete sentence that needs to be corrected.

Line 394: NISAR and BIOMASS are not upcoming missions anymore. Both have been launched already.

Clarity and context

The text is generally clear to read, with a few exceptions that I have described as issues for improvement.

References

The manuscript references the previous literature appropriately, but it omits some publications of higher resolution biomass and canopy height maps, which should be added.

Your expertise

I feel competent to comment on the paper in its entirety.

(Remarks on code availability)

I have not reviewed the code.

Version 1:

Reviewer comments:

Reviewer #1

(Remarks to the Author)

Please see my attached documents with the comments for the authors.

(Remarks on code availability)

The code repository is well organised and includes a computing environment (<https://codeocean.com>) for running the code. All data and code required to reproduce the article's figures are available in the repository. However, I encourage the authors to further improve the study's reproducibility by providing the complete set of processing pipelines used to generate the deep learning models for canopy height and biomass mapping. These training and model development pipelines are not yet available in the current repository. Making these workflows publicly accessible would substantially enhance the transparency, reproducibility, and long-term value of the study.

Reviewer #2

(Remarks to the Author)

What are the noteworthy results?

The key results of this paper are that the majority of the net carbon gains and losses in the Tropical Moist Forest in the Congo basin occur in very small patches, but because there are so many of them the overall impact on the carbon balance is substantial. My concerns regarding the role of the spatial resolution, the interpretation of canopy height change at 30 m, and the clarity of the description of the methods have been addressed in great detail by the authors. As a result, the revised manuscript is of much improved quality and I consider the methods, data, approach, and conclusions drawn from the work as robust and clear now.

My most significant concern I raised in the first round of reviews was that the 10 m resolution of the canopy height and 30 m resolution of the biomass maps were described as the main novelty of the paper, but a comparison of canopy height change on the 30 m spatial resolution in addition to the aggregated analysis at a coarser resolution was missing. The authors have clarified the role of spatial resolution in the revised manuscript and moved the comparison of canopy height change at 30 m resolution from the Supplementary Figures 10 and 11 to the main text, and expanded the validation to include both 10 m and 30 m comparisons against airborne LiDAR (ALS), as well as GEDI-based comparisons at 30 m. This, together with further substantial work, has fully addressed my previous concern. The accuracy gains of going to a coarser resolution are now much more visible in the paper and have been treated more robustly in the analysis.

Will the work be of significance to the field and related fields? How does it compare to the established literature? If the work is not original, please provide relevant references. Does the work support the conclusions and claims, or is additional evidence needed?

The revised manuscript is much improved in how it presents the significance of this study to the understanding of the global terrestrial carbon cycle and the role of African forest ecosystems. My minor comments have been fully addressed. The additional work presented in the revised paper fully supports the conclusions.

Are there any flaws in the data analysis, interpretation and conclusions? Do these prohibit publication or require revision? Is the methodology sound? Does the work meet the expected standards in your field? Is there enough detail provided in the methods for the work to be reproduced?

The quality of the data analysis, interpretation and conclusions now meets my expectations fully and the authors have addressed all my detailed suggested revisions.

My only open issue that I want to raise here is that in the main manuscript and the supplementary materials, all presented maps should have a scale bar and/or lat/lon or geographic coordinates, as is good practice.

(Remarks on code availability)

n/a

Version 2:

Reviewer comments:

Reviewer #1

(Remarks to the Author)

Comments for the Authors

I congratulate the authors of this latest review. I am very satisfied with the authors' response, the revisions they have made, and the overall improvements to the manuscript. The new analyses and additional data substantially strengthen the study. Specifically, the authors addressed my last concerns by acknowledging the high and still-present uncertainty in validating small canopy-height changes (<10-20 m) and the confidence with which these changes can be interpreted as real disturbances contributing to carbon loss and gain. Among the new additions, including new analyses, figures, and tables, I am pleased to see in the main article the changes to the previous statement about the uncertainty from "Independent visual validation using high-resolution PlanetScope imagery confirmed that 84% of these additional small-loss events represent real canopy disturbances (Extended Data Table 1)" to "Independent visual validation using high-resolution PlanetScope imagery confirmed that ~72% of these small height losses represent real canopy disturbances (Supplementary Table 2). The contributions of different height change classes to the affected area and AGC change are summarised in Supplementary Table 3, highlighting the substantial contribution of small height losses".

Overall, the article now addresses the validation of small canopy-height changes (<10-20 m), a key driver of carbon loss and gain across African forests. I have no further concerns about this article.

(Remarks on code availability)

The code has been provided, and I understand that the complete deep-learning training pipelines used to generate the annual canopy height and biomass products are under access restrictions.

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/>

Response to reviewers

EDITOR COMMENTS

As you will see from the reports copied below, the reviewers raise important concerns. We find that these concerns limit the strength of the study, and therefore we ask you to address them with additional work. Without substantial revisions, we will be unlikely to send the paper back to review. In particular, it is essential to verify and strengthen the robustness and reliability of methodology and conclusions. Independent assessment of the small-disturbance accuracy and sensitivity analysis of filtering thresholds are highly needed. In addition, the uncertainty should be quantified, and the novelty should be addressed. Finally, improving the overall narrative will be important for better conveying the study's scientific rigor.

Response: Thank you for the opportunity to revise and resubmit our manuscript. We also appreciate reviewers' constructive comments and suggestions for further helping us to improve our work. By considering all the comments carefully, we revised our manuscript accordingly. In the revised manuscript, we have made four major improvements.

1. We strengthened the independent validation of canopy height gain and loss patches using high-resolution PlanetScope imagery (5 m). We visually interpreted whether independently detected patches exhibited clear canopy decrease or increase signals consistent with the mapped disturbance direction. We conducted a stratified validation across disturbance size classes, including small disturbances below 1 ha. The previous version included only small disturbance patches (<1 ha), comprising 150 loss patches and 100 gain patches. In the revised manuscript, we expanded the validation to additional patch-size classes and increased the number of independently validated small patches to 176 loss patches and 169 gain patches, resulting in a total of 422 validated loss patches and 498 validated gain patches. This fully independent validation does not rely on existing disturbance datasets and therefore provides a more direct assessment of detection reliability.
2. We restructured the quantification of uncertainty in aboveground carbon (AGC) change. We revised the uncertainty framework to provide an improved assessment that accounts for multiple sources of error across the processing chain.
3. We performed a sensitivity analysis of the filtering parameters used in canopy height change processing. We evaluated how key thresholds influence both disturbance detection and the resulting AGC change estimates, and incorporated these results into the manuscript. The analysis showed that variations in filtering parameters did not alter the overall conclusion of widespread net carbon loss, while providing a more robust and realistic assessment of uncertainty.
4. We revised and refined the overall narrative of the manuscript. We clarified the Abstract, Introduction, Results, and Discussion to improve clarity and avoid overstatement, and incorporated additional relevant references that were previously missing.

We believe these revisions further strengthen the robustness of the methodology and the reliability of the conclusions. The line numbers referenced by the reviewers correspond to the originally submitted manuscript, while those cited in our responses refer to the revised manuscript with

tracked changes. In the responses below, our replies are shown in blue, and the corresponding text, figures, and tables from the revised manuscript are shown in red. All modifications in the revised manuscript are consistently highlighted in red.

REVIEWER COMMENTS

Reviewer #1 (Remarks to the Author):

General overview

This study aims to quantify the carbon balance of Central African dense humid forests (~2.1 Mkm²) from 2019 to 2022 by generating annual 10 m resolution canopy height (CH) maps using a U-Net deep learning model that fuses GEDI spaceborne LiDAR data with Sentinel-1 radar and Sentinel-2 optical imagery. Canopy height changes are converted to aboveground carbon (AGC) stock changes at 30 m resolution via an empirical allometric model calibrated on ~400 field plots. The authors report a net AGC loss of -58 ± 9 Tg C yr⁻¹, composed of gross losses of -126 ± 7 Tg C yr⁻¹ partially offset by gains of 68 ± 6 Tg C yr⁻¹, and emphasize that 87% of all disturbance patches are smaller than 1 ha — a finding they claim has been largely missed by previous satellite products (TMF, GFC, RADD). They argue that these results represent the first biome-wide, high-resolution quantification of forest biomass dynamics in Central Africa, with implications for national carbon accounting, conservation evaluation, and climate policy. They found that Central African forests are a net carbon source (-58 ± 9 Tg C yr⁻¹), with small-scale disturbances (<1 ha) playing a dominant and previously undetected role in both carbon losses (37% of deforestation losses) and gains (48% of regrowth gains).

The claim that Central African forests are a net carbon source over 2019–2022 is plausible and broadly consistent with independent bookkeeping models and inventory data. However, the specific magnitude depends on a chain of uncertain transformations — from raw satellite imagery, through a deep learning- based height prediction, through noise filtering, through an allometric biomass model with $R^2 = 0.40$, and through a disturbance classification scheme. Each step introduces errors that may compound in non-trivial ways, yet the uncertainty analysis treats many of these steps as independent and Gaussian, which is unlikely to be strictly true.

Response: Thank you for this comprehensive summary of our study. We agree that our estimates rely on a sequence of processing steps, including canopy height prediction, biomass conversion, and disturbance detection, and that uncertainties propagate across these steps in potentially non-independent and non-linear ways. We also understand that a central concern is whether the magnitude of the estimated carbon balance, as well as the prominent role of small-scale disturbances, could be affected by compounded uncertainties or detection errors. In the revised manuscript, we have therefore strengthened both the independent validation of disturbance detection and the uncertainty assessment framework to directly address this issue as below:

To assess the reliability of mapped CH changes, we conducted an independent validation using high-resolution monthly PlanetScope imagery (5 m) (**Supplementary Text 6**). We adopted a stratified random sampling design following commonly used good practices for land-change

accuracy assessment⁴³, based on patch-size classes (0–1, 1–5, 5–10 and 10–100 ha) and change type (gain and loss). Our initial sampling target included 200 small-disturbance patches (<1 ha) and 150 patches for each larger size class. Samples were randomly selected across the study region, and only locations with sufficiently cloud-free PlanetScope imagery in both 2019 and 2022 were retained for interpretation. We visually interpreted whether sampled patches exhibited canopy structural decrease or increase consistent with the mapped CH change direction. Interpretation was based on comparisons between monthly PlanetScope imagery from 2019 and 2022 together with the mapped CH change layer, including changes in canopy cover, shadowing, canopy openings and vegetation recovery signals. Compared to the previous version, which included only small disturbance patches (<1 ha), we expanded the validation to additional patch-size classes and increased the number of independently validated small patches to 176 loss patches and 169 gain patches, resulting in a total of 422 validated loss patches and 498 validated gain patches. Small disturbance patches, which are particularly difficult to identify because of subtle canopy structural changes, were independently interpreted by both an experienced forest remote sensing expert (Martin Brandt) and a trained analyst following the same interpretation protocol. Interpretation consistency across the same validation samples reached 94%. Larger disturbance patches (1–100 ha) were interpreted by the trained analyst using the same protocol. Finally, we derived user's accuracy for each size class as well as area-weighted accuracy⁴³ to account for the disproportionate contribution of large patches to total mapped area change (**Extended Data Table 1**). In addition, omission errors were assessed using randomly sampled 2-ha reference patches selected from undisturbed forest areas. These reference areas were visually confirmed as stable forest using PlanetScope imagery together with independent disturbance datasets.

This analysis indicated that the user's accuracy (UA) was 84% for loss and 77% for gain across the 0–100 ha range (area-weighted), and 79% and 69%, respectively, for small disturbances (<1 ha) (**Extended Data Table 1**). Omission and commission errors were ~6% and ~5%, respectively, when evaluated relative to the full landscape, corresponding to undetected changes within no-change areas and false detections in stable forests. (Lines 163–168)

Extended Data Table 1: Visual validation of canopy height loss and gain, and estimation of false detection rates using visually confirmed stable forest areas from 5-m Planet imagery.


Patch size	Type	Area (%)	Sample size	TP	FP	UA
0–1 ha	Loss	29.48%	176	139	35	78.98%
1–5 ha	Loss	27.52%	71	57	14	80.28%
5–10 ha	Loss	11.98%	79	67	12	84.81%
10–100 ha	Loss	31.02%	96	89	8	92.70%
Area-weighted user's accuracy (UA) for loss (0–100 ha)						84.29%
0–1 ha	Gain	45.63%	169	117	52	69.23%
1–5 ha	Gain	30.94%	113	92	21	81.42%

5–10 ha	Gain	8.79%	112	93	19	83.04%
10–100 ha	Gain	14.64%	104	91	13	87.50%
Area-weighted user’s accuracy (UA) for gain (0–100 ha)						76.89%

Second, we improved the uncertainty analysis framework by explicitly quantifying three major components: (1) biomass uncertainty arising from allometric modelling and measurement errors (σ_{bio}); (2) classification uncertainty associated with disturbance attribution and detection accuracy (σ_{cls}); and (3) uncertainty related to the choice of a CH change processing scheme (σ_{scheme}), including trend detection, Sentinel-1 filtering, and canopy cover filtering. Independent validation using stratified sampling further quantified the accuracy of CH change detection across different patch sizes. However, this source of uncertainty was not treated as a separate component because it had already been incorporated into the classification uncertainty framework. We quantified σ_{bio} , σ_{cls} , and σ_{scheme} separately using dedicated simulation and error propagation approaches, and combined these components under the assumption of independence to obtain a tractable estimate of total uncertainty. Full methodological details of uncertainty formulation, parameterization, and propagation are provided in **Supplementary Text 10, Supplementary Fig. 30**. (Lines 797–811)

Importantly, we further evaluated the sensitivity of our results to key parameter choices in the canopy height change processing (**Supplementary Text 4.5**). Sensitivity analyses of parameters of our change detection algorithm including trend detection, Sentinel-1 filtering, and canopy cover filtering (**Supplementary Table 8**) show that, while the overall conclusion of a net carbon loss remains unchanged, the magnitude of individual components is affected (**Supplementary Table 9, Table 10**). In particular, gross losses show a higher sensitivity to parameter choices of our change detection algorithm, with our preferred configuration yielding relatively lower net carbon loss compared to alternative thresholds, primarily due to differences in the attribution of “other loss” components. These differences are largely associated with small height losses, which are inherently more difficult to detect robustly and are more sensitive to processing assumptions.

As a result, uncertainty is higher for these small loss components. Based on the updated sensitivity analysis and uncertainty estimation, we now estimate aboveground carbon change in Central African dense forests with both revised and original uncertainty estimates in **Table A1**. These variations do not alter the overall conclusion that Central African forests acted as a net carbon source over 2019–2022.



Supplementary Fig. 30: Framework for uncertainty estimation and propagation of aboveground carbon (AGC) change.

Supplementary Table 8: Sensitivity analyses of key parameter choices of the canopy height.

Canopy height change processing	Description of change processing	key parameter choices	
		Parameter used in main analysis	Sensitivity analysis
trend detection	We applied a Mann–Kendall trend test to annual CH gain data from 2019 to 2022 to identify gradual and consistent canopy height increases, retaining only pixels with S-statistic ($S \geq 6$) to reduce noise and avoid false regrowth signals.	$S \geq 6$	$S \geq 6$; $S \geq 5$
Sentinel-1 filtering	Sentinel-1 filtering was applied to remove potentially false CH loss pixels exhibiting very low temporal variability in the Sentinel-1 VV signal ($ CV \leq 5\%$) between 2019 and 2022.	$ CV \leq 5\%$	$ CV \leq 1\%-6\%$
Canopy cover filtering	We used annual canopy cover change ($ \Delta CC $) data as an additional reference to validate CH change, treating pixels with $ \Delta CC \leq 2\%$ as weak or uncertain disturbances and removing them only when ΔCC and ΔCH showed inconsistent change directions.	$ \Delta CC \leq 2\%$	$ \Delta CC \leq 1\%-3\%$

Table A1: Comparison of revised and original uncertainty estimates for aboveground carbon change in Central African dense forests (for response letter only).

		AGC change (original uncertainty)	AGC change (revised uncertainty)
Loss	Other	-22.4 ± 3.7	-22.4 ± 6.0
	Deforestation	-23.3 ± 1.7	-23.3 ± 6.6
	Degradation	-80.4 ± 6.0	-80.4 ± 11.5
	Total	-126.1 ± 7.3	-126.1 ± 20.6
Gain	Other	18.8 ± 1.6	18.8 ± 4.1
	Degradation	34.4 ± 5.0	34.4 ± 5.6
	Regrowth	14.5 ± 2.4	14.5 ± 2.6
	Total	67.8 ± 5.8	67.8 ± 10.0
Net		-58.3 ± 9.3	-58.3 ± 13.5

Where the argument may break down:

The main novelty claim — that small disturbances (<1 ha) have been systematically missed — rests on the assumption that the additional detections in their maps are real disturbances rather than residual noise surviving the filtering pipeline. While the authors invest considerable effort in multi-step noise reduction (spatiotemporal connectivity, Mann-Kendall tests, Sentinel-1 CV thresholds, canopy cover consistency), the validation of these small signals is indirect and incomplete. The direct validation against airborne LiDAR has $R^2 = 0.53$ at 30 m (for height change, not carbon), and against GEDI, only $R^2 = 0.22$ at 10 km. The visual validation with Planet imagery (80% agreement for losses, 66% for gains, after including TMF-consistent "uncertain" cases) includes a substantial fraction of ambiguous cases that are generously counted as confirmatory. This makes the quantitative partitioning of carbon fluxes into small vs. large disturbances less robust than presented.

Response: To address this concern, we conducted an independent validation using high-resolution (5 m) PlanetScope imagery, with a particular focus on disturbance size. We adopted a stratified sampling design⁴³ based on patch size (0–1, 1–5, 5–10 and 10–100 ha) and change type (gain and loss), and visually interpreted randomly selected samples. We included a total of 422 validated loss patches and 498 validated gain patches. For each sample, we assessed whether the detected CH change was supported by observed changes, and derived user's accuracy for each size class as well as area-weighted accuracy⁴³ to account for the disproportionate contribution of larger patches to total change (**Extended Data Table 1**). This analysis indicated that the user's accuracy (UA) was 84% for loss and 77% for gain across the 0–100 ha range (area-weighted), and 79% and 69%, respectively, for small disturbances (<1 ha) (**Extended Data Table 1**). Omission and commission errors were ~6% and ~5%, respectively, when evaluated relative to the full landscape, corresponding to undetected changes within no-change areas and false detections in stable forests. Our validation is fully independent of existing disturbance products, and only clear disturbance

signals observed in PlanetScope imagery were counted as agreement. This significantly enhanced size-stratified validation allows us to directly assess detection reliability across scales. The results show that the detected CH gain and losses have a good agreement with high-resolution imagery, supporting that these detections are not dominated by residual noise.

We have added these results and revised the corresponding text (Lines 161–168, Lines 623–643). [Please refer to our response to the related comment in the General Overview section above.](#)

1. Validation of Small-Scale Disturbance Detection

The single most consequential weakness of this manuscript concerns the evidence offered in support of its central novelty claim: that 87% of all forest disturbances in Central Africa are smaller than one hectare, and that these sub-hectare patches account for 37% of deforestation carbon losses and 48% of regrowth gains. These are striking figures with direct implications for national carbon accounting and conservation policy. Yet the chain of evidence linking the satellite-derived height-change maps to these quantitative assertions is, at present, insufficiently robust to bear the weight of the conclusions drawn from it.

Every downstream result in the paper — the biome-wide carbon budget, the national-scale comparisons with inventories and bookkeeping models, the assessment of protected areas, and the policy recommendations — depends on the accuracy of these small-disturbance detections. If even a modest fraction prove to be false positives, the headline numbers shift substantially, and the narrative of the paper changes with them.

The validation evidence, though multifaceted, falls short of what is needed. Direct comparison with airborne LiDAR yields an R^2 of 0.53 at 30 m resolution, but this is drawn from a single site in the DRC (Yangambi Biosphere Reserve) where the ALS acquisition dates (2015 and 2023) do not coincide with the study period, requiring the authors to impose a linear interpolation assumption for height gains and to exclude pixels showing loss between 2015 and 2019. The most spatially extensive independent comparison — against GEDI footprint data at 10 km resolution — produces an R^2 of only 0.22, a figure that the authors themselves acknowledge reflects non-independence between the training and evaluation data. Visual interpretation of Planet imagery over 250 randomly selected patches yields 66% clear agreement for loss and 56% for gain; figures that rise to 80% and 66%, respectively, only after the generous inclusion of cases classified as "uncertain" but deemed consistent with TMF disturbance trends. These inflated agreement rates are then presented as the primary evidence for the reliability of sub-hectare detections.

Perhaps most critically, a structural circularity undermines the comparison that underpins Figure 2. The noise-filtering pipeline retains all pixels classified as disturbed by the Tropical Moist Forests dataset, regardless of whether they meet the authors' noise criteria. The final product is therefore guaranteed to be a superset of TMF, and the comparison showing that the new maps detect "more" disturbances is a foregone conclusion. The genuinely novel contribution lies solely in the additional detections — the pixels that TMF missed but the deep learning model found — and these are precisely the subset that has survived the noise filters by the narrowest of margins, having been drawn from a population in which 86% of pixels with height changes smaller than 2 m were themselves rejected as noise by the authors' own criteria.

To address this issue, the authors should conduct a formal accuracy assessment using a stratified random sample of sub-hectare detections, interpreted independently against very-high-resolution imagery by multiple analysts, with a proper confusion matrix and area-adjusted accuracy metrics following established protocols such as those of Olofsson et al. (2014). The false-positive rate should be reported for each patch-size class. A null-hypothesis test — applying the detection pipeline to temporally shuffled data or to known-stable forest areas — would help establish the expected rate of spurious detections. If the false-positive rate proves substantial, all quantitative claims regarding the contribution of small disturbances to the carbon budget must be revised accordingly.

Reference:

Olofsson, P., Foody, G. M., Herold, M., Stehman, S. V., Woodcock, C. E., & Wulder, M. A. (2014). Good practices for estimating area and assessing accuracy of land change. *Remote Sensing of Environment*, 148, 42–57.

Response: Thanks a lot for the comments on the data validation. We have made related revisions based on three major comments.

(1) CH change processing using the TMF dataset

We want to clarify that TMF was not used as a strict inclusion mask. Instead, it was used at specific steps to retain pixels that showed consistent disturbance signals with our CH gains and losses maps when other thresholds might remove them. However, we agree that this introduces a degree of dependence, and therefore comparisons with TMF should not be interpreted as an independent validation. We have clarified the text as:

To avoid excluding real disturbances during the threshold-based filtering steps, the TMF disturbance data⁸ was only used in steps (2)–(4) to re-evaluate the pixels for which our CH change and TMF disturbance trends were consistent. TMF pixels classified as disturbed were not automatically retained; instead, they were kept only when their disturbance signal matched our CH change trends. (Lines 689–693)

(2) Validation of CH change using the ALS data

We acknowledge that this validation was made based on simple assumptions and a single dataset. We have discussed the impact and added relevant results in [Supplementary Text 5](#) as:

The effects of temporal interpolation and sample filtering on the comparison between satellite-derived CH change (2019–2022) and ALS observations (2015–2023) were evaluated at both 10 m and 30 m resolutions (**Supplementary Figs. 12, 13**). We first evaluated the consistency between satellite-derived and ALS-derived CH maps at individual years (**Supplementary Fig. 2**). Our CH map in 2019 showed good agreement with ALS CH observations acquired in 2015 ($R^2 = 0.76$, MAE = 4.35 m, Bias = -1.08 m), with comparable performance between our CH map and ALS observations acquired in the same year 2023, although the correlation was lower ($R^2 = 0.55$, MAE = 4.42 m, Bias = 0.29 m). The shift in bias between the two comparisons likely reflects temporal mismatches between ALS acquisition dates and satellite observations, as forest growth and disturbance may occur between acquisition periods. Applying a linear interpolation to pixels with height gain (condition 1) resulted in only minor changes in agreement metrics, indicating that

assumptions on growth dynamics have a limited influence on the comparison. In contrast, excluding pixels with ALS-detected height loss prior to the 2019–2022 study period (condition 2) led to a substantial improvement in agreement, increasing R^2 from 0.06 to 0.22 at 10 m resolution and to 0.52 at 30 m resolution, while also reducing error metrics at both spatial resolutions. This reflects the effect of temporal mismatch between the ALS acquisition (2015 and 2023) and our shorter study period, whereby disturbances occurring before our observation window introduce inconsistencies in the raw comparison. Removing these pixels reduces discrepancies associated with misaligned disturbance timing rather than altering the underlying distribution of canopy height change. Importantly, consistent patterns across 10 m and 30 m analyses indicate that the observed improvements are not driven by scale-dependent smoothing or selective bias, but by improved temporal alignment between datasets. These results suggest that differences in the comparison between ALS-observed CH change and our dataset are primarily attributable to timing and sampling inconsistencies, rather than systematic overestimation of height change in our satellite dataset.

(3) Accuracy assessment of CH change

We have added a formal accuracy assessment using a stratified random sample of sub-hectare detections, interpreted independently against very-high-resolution imagery by referring to Olofsson et al. (2014) as. [Please refer to our response to the related comment on small disturbance validation in the General Overview section above.](#)

To assess the reliability of mapped CH changes, we conducted an independent validation using high-resolution monthly PlanetScope imagery (5 m) (**Supplementary Text 6**). We adopted a stratified random sampling design following commonly used good practices for land-change accuracy assessment⁴³, based on patch-size classes (0–1, 1–5, 5–10 and 10–100 ha) and change type (gain and loss). Samples were randomly selected across the study region, and only locations with sufficiently cloud-free PlanetScope imagery in both 2019 and 2022 were retained for interpretation. We visually interpreted whether sampled patches exhibited canopy structural decrease or increase consistent with the mapped CH change direction. Interpretation was based on comparisons between monthly PlanetScope imagery from 2019 and 2022 together with the mapped CH change layer, including changes in canopy cover, shadowing, canopy openings and vegetation recovery signals. We included a total of 422 validated loss patches and 498 validated gain patches. Small disturbance patches, which are particularly difficult to identify because of subtle canopy structural changes, were independently interpreted by both an experienced forest remote sensing expert and a trained analyst following the same interpretation protocol. Interpretation consistency across the same validation samples reached 94%. Larger disturbance patches (1–100 ha) were interpreted by the trained analyst using the same protocol. Finally, we derived user's accuracy for each size class as well as area-weighted accuracy⁴³ to account for the disproportionate contribution of large patches to total mapped area change (**Extended Data Table 1**). In addition, omission errors were assessed using randomly sampled 2-ha reference patches selected from undisturbed forest areas. These reference areas were visually confirmed as stable forest using PlanetScope imagery together with independent disturbance datasets. (Lines 623–643)

This analysis indicated that the user's accuracy (UA) was 84% for loss and 77% for gain across the 0–100 ha range (area-weighted), and 79% and 69%, respectively, for small disturbances (<1 ha) (**Extended Data Table 1**). Omission and commission errors were ~6% and ~5%, respectively, when evaluated relative to the full landscape, corresponding to undetected changes within no-change areas and false detections in stable forests. (Lines 163–168)

2. Uncertainty Estimation

The reported uncertainty of ± 9 Tg C yr⁻¹ on a net flux of -58 Tg C yr⁻¹ — roughly 16% in relative terms — almost certainly understates the true range of plausible values. The error budget is constructed from a sequence of assumptions, several of which are either demonstrably violated or insufficiently justified, and whose cumulative effect is to produce confidence intervals that are narrower than the evidence warrants.

The core issue is the allometric model that connects canopy height to aboveground biomass. With an R^2 of 0.40 and a relative error of 32.6%, this model explains only 60% of the biomass variation, yet it remains the only method used to convert height-change observations into the carbon fluxes which are the main focus of the paper. The uncertainty framework considers this model as a fixed, known relationship rather than a significant source of error propagation. Adding to this, the two annual maps, whose difference indicates the carbon change, share the same deep learning architecture, the same allometric equation, the same wood density inputs, and similar sensor conditions. Errors from these shared components are perfectly correlated across years and do not cancel out when differenced, yet the propagation framework assumes inter-annual independence—an assumption the authors openly state but do not verify.

The calibration uncertainty is set at a flat 20%, a figure borrowed from Chave et al. (2001), a study of biomass variability within a single ten-hectare plot in French Guiana (Lines 60-601 of the article). No explanation is provided for why this figure should be relevant to Central African forests, satellite-derived biomass estimates, or the spatial scales discussed here. It acts as a placeholder where a properly calibrated estimate is needed. The classification uncertainty is more problematic: an XGBoost model trained on the authors' own disturbance class labels is then evaluated against those same labels, producing an overall accuracy of 0.87 that measures internal consistency rather than true correctness. The subsequent decision to average two class-specific error rates — 1% for degradation and 14% for deforestation — into a single 7.5% value, applied uniformly to both classes, lacks statistical grounding and conceals real differences in classification difficulty. Lastly, the propagation from pixel-level to regional uncertainty entirely ignores spatial autocorrelation, despite the fact that errors in canopy height prediction, biomass estimation, and disturbance classification are expected to exhibit strong spatial dependence due to shared atmospheric conditions, sensor geometry, and training data.

In revision, the authors should substitute the circular classification uncertainty assessment with one derived from independent reference data. Address spatial autocorrelation using geostatistical or bootstrap-based error propagation methods. The 20% calibration uncertainty should be justified with Central African data or replaced with a more precise value. Include a sensitivity analysis demonstrating how the net flux estimate changes within a reasonable range of noise-filtering thresholds—such as the Mann-Kendall S-statistic, Sentinel-1 coefficient of variation, and canopy

cover change—to define the uncertainty range. The authors should also consider the possibility that the revised uncertainty may include zero net change in some sub- regions and adjust their language accordingly.

Response: In the revised manuscript, we have restructured entirely the uncertainty estimation framework to reduce inappropriate assumptions and improve uncertainty propagation. We have updated the biomass uncertainty calculation, and have not used a constant RMSE or an XGBoost-derived residual uncertainty. Instead, biomass uncertainty is modeled as a heteroscedastic error proportional to AGC, based on independent validation of the biomass model (relative error). This allows the AGC uncertainty to vary spatially with biomass magnitude, rather than assuming a uniform error across all forests. We further accounted for intertemporal dependence in biomass estimation errors using a mixed error structure when propagating uncertainty in AGC change, instead of assuming independent errors between years.

We have removed the previous approach, using XGBoost-derived accuracy evaluated on the same labels and averaging error rates across disturbance types. Specifically, the classification uncertainty is now quantified separately for each disturbance class by combining (1) detection uncertainty derived from independent validation data (high-resolution monthly PlanetScope images at 5 m), (2) threshold sensitivity in distinguishing deforestation and degradation losses, and (3) Uncertainty associated with regrowth and degradation gain estimation. These components are treated explicitly and are not averaged across classes. Uncertainty is propagated through Monte Carlo simulations and class-specific coefficients of variation, allowing differences in classification difficulty and ambiguity to be preserved. This provides a more realistic representation of classification uncertainty compared to the original framework.

We now decomposed the total uncertainty into three components: (1) biomass uncertainty arising from allometric modelling and measurement errors (σ_{bio}); (2) classification uncertainty associated with disturbance attribution and detection accuracy (σ_{cls}); and (3) uncertainty related to the choice of a CH change processing scheme (σ_{scheme}), including trend detection, Sentinel-1 filtering, and canopy cover filtering. Independent validation using stratified sampling further quantified the accuracy of CH change detection across different patch sizes. However, this source of uncertainty was not treated as a separate component because it had already been incorporated into the classification uncertainty framework. We quantified σ_{bio} , σ_{cls} , and σ_{scheme} separately using dedicated simulation and error propagation approaches, and combined these components under the assumption of independence to obtain a tractable estimate of total uncertainty. Full methodological details of uncertainty formulation, parameterization, and propagation are provided in **Supplementary Text 10, Supplementary Fig. 30**. The resulting total relative uncertainties are 23.2%, 16.3%, and 14.7% for net AGC change, gross AGC loss, and AGC gain in Central African forests, respectively (**Extended Data Table 2**).

We have clarified these points in the revised manuscript (Lines 797–811, Supplementary Text 10, Supplementary Text 4.5). [Please also refer to our response to the comment on the uncertainty estimation in the General Overview section above.](#)

3. Interpretive Overreach

A recurring pattern throughout the manuscript is the tendency for interpretive claims to surpass the supporting evidence. This is not merely a matter of isolated overstatements; it is a systematic framing issue that influences the abstract, the results, the discussion, and the conclusions.

The abstract states that the remote-sensing-derived national carbon budgets “align well with country-level inventories and bookkeeping model estimates.” In practice, the agreement is close for Gabon and Equatorial Guinea — two relatively small countries with stable forest sectors — reasonable for Cameroon, and notably inconsistent for the Democratic Republic of the Congo, which accounts for over 70% of the regional forest area and dominates the overall carbon flux. Describing this varied pattern as “aligning well” is somewhat optimistic. More concerningly, the discrepancy in the DRC is interpreted asymmetrically: where there is agreement, both methods are considered validated; where there is disagreement, the national inventory is presumed to be incorrect. This asymmetrical interpretation characterising the national budget comparison should be addressed.

The characterisation of the DRC as a “small net source” of carbon at -46 Tg C yr^{-1} is misleading in a different way. This figure represents approximately 79% of the total regional net loss. The word “small” appears to be intended in relation to the DRC’s vast total carbon stock, but this framing is not made explicit, and a reader encountering the phrase in the abstract would reasonably conclude that the flux is unimportant. It is not.

The claim that protected areas show “uneven conservation outcomes” based on net biomass loss is an interpretive leap that the data cannot justify without identifying potential causes. Biomass losses within protected areas could be due to natural disturbance events — such as treefalls, droughts, or fires — just as much as they could indicate conservation failure. They might also stem from edge effects spreading inward from nearby deforestation or from prior degradation before the areas were designated. The authors admit elsewhere that they cannot distinguish between natural and human-induced drivers, yet the abstract language suggests a judgment about conservation effectiveness.

The headline statistic that 87% of all disturbances are smaller than one hectare, while mathematically correct, is partly a trivial artefact caused by high-resolution rasterisation. In any landscape with dispersed disturbances, the number of small patches will far exceed that of larger ones. The ecologically and policy-relevant question — whether these small patches contribute a disproportionate share of carbon flux — is partly addressed by the 37% and 48% figures of carbon losses and gains, respectively, but these rely on the very small-disturbance detections whose validation is questionable.

The concluding policy recommendations — on carbon crediting schemes, support for local communities, and prioritising young-forest regeneration over large-scale plantations — go beyond what the data can substantiate. The study maps carbon changes from space; it does not examine the socioeconomic drivers, institutional contexts, or ecological trade-offs that would be necessary to descriptions. Similarly, the finding that undisturbed forests show near-zero or negative growth rates, presented as evidence that the IPCC default growth rates are too high, could alternatively reflect a regression-to-the-mean bias in the deep learning model, which would systematically

underestimate predicted heights in the tallest forests and thus diminish apparent gains. This alternative explanation is not discussed.

In revision, the authors should rewrite the abstract and discussion to clearly distinguish between what is reliably established — such as the existence and general significance of small-scale disturbances, and the approximate magnitude and sign of the regional carbon imbalance — and what remains quantitatively uncertain, including exact proportional attributions, precise national budgets, and the growth-rate comparison with IPCC defaults. The term “align well” should be removed or heavily qualified. The DRC should not be described as a “small” net source. Policy recommendations should either be removed or explicitly reframed as speculative observations requiring further investigation.

Response: Thank you for your suggestions. We agree that some interpretations in the original manuscript were too strong compared to the supporting evidence. We appreciate this comment and have revised the abstract, results, and discussion to clearly separate well-supported findings and more uncertain speculations, and to use cautious language throughout. Specifically, we have (1) removed general statements such as “align well” and instead show differences among countries, especially for the DRC; (2) avoided one-sided interpretation of differences by acknowledging uncertainties in both our estimates and national inventories; (3) removed potentially misleading wording (e.g., “small” for the DRC net source) and clarified how net carbon change magnitudes should be interpreted; (4) revised the discussion about protected areas to avoid implying causality; and (5) reduced emphasis on patch counts and focused on their contribution to carbon change, while noting related uncertainties. We have also removed or clarified policy-related statements and added discussion of key limitations, including possible model-related biases in estimating growth rates. We have provided point-to-point revision and reply to each related comment described below.

4. The Processing Pipeline and Its Sensitivity

The multi-step noise-filtering pipeline that transforms raw canopy height differences into the final change product — passing through spatiotemporal connectivity analysis, a Mann-Kendall trend test, Sentinel-1 coefficient-of-variation filtering, and canopy cover consistency checks — is one of the study's more thoughtful methodological contributions. It reflects a genuine effort to grapple with the formidable challenge of separating real structural change from sensor artefacts in persistently cloudy tropical forests. Yet the pipeline also introduces a serious vulnerability: it relies on a series of thresholds, each drawn from a different study in a different region, and the sensitivity of the headline results to these threshold choices has not been tested.

The Mann-Kendall S-statistic threshold of 6 for height gains, the Sentinel-1 |CV| threshold of 5 for height losses, and the canopy cover change threshold of $|\Delta CC| \leq 2\%$ for pinpointing undisturbed forests were each derived from published research focusing on different forest ecosystems. None has been proven to be suitable specifically for Central African dense humid forests. The $|CV| = 5$ threshold, for example, was chosen despite the lack of any point at that value in the authors' own data (Supplementary Fig. 8), where canopy height loss varies steadily with the coefficient of variation. The canopy cover threshold of 2% may be too strict in closed-canopy forests where selective logging of a single tree could easily cause a cover change below this level, thereby risking the exclusion of genuine low-intensity degradation events.

Another issue is the uneven treatment of existing and new detections. TMF-classified disturbance pixels are kept, regardless of whether they meet the noise criteria, whereas recently detected pixels must pass all four filtering steps. This dual approach means the product's benefit over TMF is somewhat exaggerated — it is guaranteed to detect everything TMF does, with new detections subject to a more rigorous evidentiary standard. The significance of this asymmetry for the comparisons in Figure 2 is not explained.

The spatiotemporal connectivity analysis, while conceptually sound, operates across three time layers ($\Delta CH_{2021-2019}$, $\Delta CH_{2022-2019}$, $\Delta CH_{2023-2019}$) that all share the 2019 baseline map. Correlated noise in the 2019 layer — such as from a single cloud-contaminated Sentinel-2 composite — could propagate into all three layers as a consistent apparent gain, creating a spatiotemporally connected cluster that survives the filter despite lacking any physical basis.

The authors should include a sensitivity analysis in which each key threshold is varied within a reasonable range, and the resulting changes in net carbon flux, gross gains, gross losses, and the distribution of sub-hectare disturbances are documented. This analysis could be displayed as a figure or table in the supplementary materials. If the results are sensitive to the threshold choices, the reported proportional coded uncertainty should be broadened to account for this parameter dependence, or the authors should consider using a multi-threshold approach.

Response: We clarify that TMF was not used as a strict inclusion mask. Instead, it was used at specific processing steps to retain pixels if TMF disturbance signals are consistent with our CH-derived gains and losses even when other thresholds might otherwise exclude them.

To avoid excluding real disturbances during the threshold-based filtering steps, the TMF disturbance data⁸ was only used in steps (2)–(4) to re-evaluate the pixels for which our CH change and TMF disturbance trends were consistent. TMF pixels classified as disturbed were not automatically retained; instead, they were kept only when their disturbance signal matched our CH change trends. Further, we evaluated the sensitivity of our results to key parameter choices in the CH change processing including trend detection, Sentinel-1 filtering, and canopy cover filtering (**Supplementary Table 8**). The resulting variability in AGC gain, loss, net change, and associated area (**Supplementary Table 9, Table 10**) was propagated as part of the structural uncertainty in the uncertainty framework. Details of the sensitivity analysis are provided in **Supplementary Text 4.5**. (Lines 689–698).

Details of the sensitivity analysis are provided in [Supplementary Text 4.5](#) as:

To assess the robustness of our CH change detection, we evaluated the sensitivity of our results to key parameter choices in the canopy height change processing. Sensitivity analyses of parameters of our change detection algorithm including trend detection, Sentinel-1 filtering, and canopy cover filtering as below (**Supplementary Table 8**):

(1) Trend detection: We applied a Mann–Kendall trend test^{22–24} to assess the temporal trend of CH gain ($\Delta CH > 0$) using annual data from 2019 to 2022. Because the available time series is short, this test mainly identifies pixels with gradual and consistent canopy height increases. Only pixels with S-statistic (S) ≥ 6 were retained to ensure that detected increases represent consistent growth signals, as lower values are insufficient to robustly distinguish true regrowth from noise. Sensitivity

analyses using alternative thresholds ($S \geq 5$) resulted in only minor differences in AGC and area estimates.

(2) Sentinel-1 filtering: To reduce noise of the remaining valid CH loss pixels, we collocated the annual-mean Sentinel-1 descending images with the pixels exhibiting valid CH changes. Pixels showing very low variability in the Sentinel-1 VV signal ($|CV| \leq 5\%$) between 2019 and 2022 were excluded, as these are likely associated with false structural canopy changes. This threshold was selected empirically based on model uncertainty, as height loss within this range is generally small (<5 m) and subject to higher relative uncertainty. In addition, we evaluated the sensitivity of our estimated AGC loss and disturbed area to alternative Sentinel-1 variability thresholds ($|CV| \leq 1\% - 6\%$).

(3) Canopy cover filtering: We used annual canopy cover change (ΔCC) data as an additional reference to validate CH change. Forest pixels with valid CH change (ΔCH) after the previous processing but with $|\Delta CC| \leq 2\%$ ^{27,28} were considered weak or uncertain disturbances. In this step, pixels were removed only when ΔCC and ΔCH showed inconsistent change trends, such that disturbances with consistent ΔCH signals were retained even when ΔCC was small. Furthermore, we evaluated the sensitivity of our estimated AGC loss and disturbed area to alternative $|\Delta CC|$ thresholds ($|\Delta CC| \leq 1\% - 3\%$).

Across all tested configurations (**Supplementary Tables 9, 10**), variations in these parameters primarily affected estimates of net AGC change (up to 12.43%) and the corresponding area (up to 10.00%) as well as gross AGC gain (6.65%) and the corresponding area (15.58%). At the disturbance-class level, the largest sensitivity was observed for the “other loss” category, with AGC estimates varying by up to 21.80% and area by up to 31.51%, while other disturbance classes showed comparatively smaller variations. These results indicate that parameter choices mainly influence less constrained disturbance types, whereas overall AGC budgets remain relatively stable. The resulting variability was incorporated into the structural uncertainty of AGC and area change estimates.

5. The Allometric Model and Biomass Estimation

The allometric equation that converts canopy height and wood density into aboveground biomass — $AGB = 6.5981 \times (CH \times WD)^{1.0356}$ — is the sole quantitative link between the height-change maps and the carbon budget. Its performance, with an R^2 of 0.40 and a relative error of 32.6% against approximately 400 field calibration plots, indicates that 60% of biomass variability at the pixel level remains unexplained. For a study that reports national carbon budgets with uncertainties of 7–15%, this is a significant limitation.

Several features of the model deserve closer examination than the manuscript does. The exponent of 1.04 is notably close to one, indicating a nearly linear relationship between the product of canopy height and wood density, and aboveground biomass. Conversely, allometric theory predicts a distinctly non-linear scaling between tree height and biomass when considering stand-level distributions. The nearly linear exponent may therefore be a compensatory artefact: GEDI RH100, which captures the top-of-canopy return, is an imperfect proxy for stand-level biomass in structurally diverse canopies, and the model fit might perform best by compressing the underlying allometric relationship. If this is true, the model may be reasonably effective for estimating average

biomass over large areas but could introduce systematic biases for certain disturbance types or canopy structures — precisely where the study's change estimates are most critical.

The wood density input, derived from a 1 km-resolution map resampled to 30 m via bilinear interpolation, introduces an additional constraint. Because each 30 m pixel within a 1 km neighbourhood effectively shares the same wood density value, all fine-scale variation in predicted biomass is driven solely by canopy height. A genuine change in species composition following disturbance — for example, the replacement of slow-growing hardwoods with fast-growing pioneer species — would alter wood density at the stand level but would not be reflected in the model, potentially biasing the estimated carbon change. Perhaps the AGB map "one" produced by the lead author without WD, and published in 2025 in *Frontiers in Remote Sensing*, is, in my view, much more compelling than the current AGB map, which does not use coarse WD data at 1 km.

The extension of this dense-forest-calibrated model to non-dense ecosystems such as Miombo woodlands is presented casually, as though it were a minor generalisation. In fact, Miombo woodlands have fundamentally different canopy architecture, species composition, fire regimes, and height-biomass relationships. The validation offered — a comparison with ESA CCI AGB, which is itself a satellite-derived product — demonstrates agreement between two models, not accuracy against the physical world. Without field-based validation in non-dense forests, the country-level carbon budgets that incorporate these areas should be interpreted with considerable caution.

It is also worth noting that all validation in the manuscript addresses height change, not carbon change. Given the allometric model's limited explanatory power, the height-change maps may be broadly correct while the carbon budget remains substantially biased. The authors should discuss these limitations more transparently, provide a residual analysis stratified by forest type, disturbance severity, and canopy structure, and clearly acknowledge that validation against ESA CCI AGB constitutes cross-comparison with another model rather than ground truth.

Response: We agree that the canopy height–biomass relationship is a central component of our framework and that its uncertainty should be more clearly and transparently discussed. In the revised manuscript, we first clarify the interpretation of the reported model performance ($R^2 = 0.40$). This value is derived from a limited set of field inventory plots (~400) that are unevenly distributed across Central Africa and are disproportionately located in high-biomass forests. As a result, this statistical metric does not fully represent model performance across the entire range of forest conditions, and must be interpreted with caution. We further clarify that the moderate agreement with field data implies substantial uncertainty at the pixel scale. However, uncertainties in aggregated carbon estimates are reduced through spatial averaging, and the reported uncertainties in national-scale carbon budgets therefore primarily reflect aggregated variability rather than pixel-level errors.

We have expanded the discussion about the uncertainty in biomass mapping in the Method section: Although biomass–height scaling relationships are theoretically expected to be non-linear, we evaluated several alternative formulations, including machine-learning approaches²⁶. While some models achieved higher predictive accuracy of biomass on validation samples, they produced spatial artifacts and unstable temporal patterns when applied to wall-to-wall mapping at 30 m

resolution. To ensure spatial coherence and temporal consistency across annual maps, we therefore adopted this simpler height-based formulation. The WD map was resampled from 1 km to 30 m for AGB estimation. Although this resampling does not introduce fine-scale variation in WD, including WD helps mitigate biomass saturation effects when converting CH to AGB. In practice, most spatial and temporal variability in AGB is driven by CH, while WD acts primarily as a regional modifier. Because the WD dataset did not fully cover all forest pixels in either 2019 or 2022, missing values were filled using the nearest available WD value. These filled pixels occur primarily in disturbed areas and low-density forest regions and account for only 1.5% of all pixels. The AGB model was calibrated using more than 400 field observations of AGB from African dense forests and showed moderate agreement with available field inventory data, with R^2 of 0.40 and relative error of 32.6%. This validation statistic reflects scale mismatch and representativeness of the available field inventory plot validation dataset and should not be interpreted as the variance explained across the entire mapped domain. Field plots are sparse across Central Africa and are disproportionately located in high-biomass forests, which affects the validation statistics. (Lines 539–556)

The AGB model selected in this study was specifically calibrated for African dense forests. To enable country-level carbon accounting, we applied the same CH–AGB relationships to non-dense forest systems. This represents an important assumption, as structural differences in canopy architecture, species composition, and disturbance regimes may affect the transferability of these relationships (**Supplementary Text 1**). However, our previous ALS-based evaluation²⁶ included a wide range of forest structures, including relatively short forests in wet regions that are not classified as dense forests, and showed generally consistent performance across these conditions. Due to the limited availability of field inventory data for non-dense forests, a direct recalibration of the model was not feasible. Instead, we evaluated the robustness of the resulting AGB estimates through comparison with the ESA-CCI AGB product as an independent satellite-based dataset. The two datasets showed broadly consistent spatial patterns (**Supplementary Text 12, Supplementary Figs. 34–36**), supporting the large-scale applicability of the approach while acknowledging potential uncertainties at local scales. (Lines 564–576)

6. Framing and Literature

The introduction creates a narrative arc where current methods for estimating the carbon balance of African forests — including national inventories, bookkeeping models, and dynamic global vegetation models — are shown to have significant limitations. This ultimately motivates the authors’ satellite-based stock-change approach as the solution. The critique of existing methods is mostly fair. However, it lacks any recognition that the stock-change approach presents its own challenges: sensor noise, allometric uncertainty, difficulty detecting understory biomass changes, and reliance on empirical height-to-biomass relationships that may not apply across different forest types. This one-sided framing suggests that spatially explicit remote sensing automatically overcomes the limitations of other methods, when in fact each approach has strengths and weaknesses that are best utilised together.

The intact forest carbon sink is introduced by citing Lewis et al. (2009), a landmark study that reported large and increasing carbon uptake in African tropical forests. However, since that

publication, Hubau et al. (2020, Nature)—a study not cited in the manuscript—has shown that the African intact forest sink peaked in the 2000s and has since been declining, with the possibility of further weakening under continued warming. The paper's own data reinforce this more recent picture: Supplementary Table 4 reports near-zero or even negative mean carbon change in forests classified as undisturbed by TMF. The introduction should engage with this tension rather than citing a fifteen-year-old reference as though it represented the current state of knowledge.

The research question, as framed — "Do our high-resolution maps detect more small-scale disturbances than previously found?" — is self-fulfilling. A 10 m product will, by construction, resolve smaller features than a 30 m product, all else being equal. The scientifically meaningful question is not whether more detections occur, but whether they are real and whether they materially alter the estimated carbon budget. Reframing the question in these terms would strengthen the paper's intellectual contribution and align the introduction with the validation challenges identified throughout this review.

Finally, the extension of the study's methods to non-dense forests, including Miombo woodlands, is introduced in the final sentence of the introductory section as a subordinate clause. This understates what is, in practice, a substantial methodological assumption — the application of a canopy height model and an allometric equation calibrated on dense humid forests to ecosystems with entirely different vertical structure, species pools, and disturbance regimes. This decision deserves a more prominent and transparent treatment.

Reference: Hubau, W., Lewis, S. L., Phillips, O. L., Affum-Baffoe, K., Beeckman, H., Cuní-Sanchez, A., ... & Zemagho, L. (2020). Asynchronous carbon sink saturation in African and Amazonian tropical forests. *Nature*, 579(7797), 80-87.

Response: Thank you for this constructive comment. We have revised the manuscript accordingly, and respond to each point below.

(1) On the one-sided framing of the stock-change approach

In the revised manuscript, we now explicitly acknowledge key sources of uncertainty of satellite-based stock-change estimates, including sensor noise, allometric uncertainty, limited sensitivity to understory biomass changes, and the potential lack of transferability of height–biomass relationships across all forest types. We have also clarified that satellite-based stock-change approaches are complementary to, rather than suggested as a replacement for, inventories and model-based methods, and that these approaches provide distinct but complementary information on forest carbon dynamics. (Lines 73–80, 88–100) Nevertheless, our data provide one of the very first estimate of stock changes, at high resolution and now with a rigorous validation and uncertainty analysis

(2) On the outdated representation of the African forest carbon sink

The introduction has been revised to incorporate more recent evidence on the evolution of the African forest carbon sink (Hubau et al., 2020, Rodríguez-Veiga et al., 2025, Sagang et al., 2026). We also added a brief discussion about the uncertainty in estimated net carbon source. (Lines 58–60, 507–511)

(3) On the formulation of the research question

We now emphasize whether the additional small-scale disturbances identified in our high-resolution maps represent real disturbance events compared with existing Landsat/Sentinel-1-based disturbance datasets. (Lines 116–118)

(4) On the treatment of non-dense forests (e.g., Miombo woodlands)

We agree that the extension of the method to non-dense forests represents a substantial methodological assumption. In the revised manuscript, we have made this point more explicit and prominent. In the introduction, we now clearly identify this extension as a key methodological assumption and explicitly acknowledge the associated uncertainties related to differences in forest structure, species composition, and disturbance regimes. (Lines 119–125)

In addition, we have strengthened the Methods section to clarify the implications of applying models calibrated for dense forests to non-dense systems, and to describe the validation strategy used to assess the robustness of the resulting carbon estimates. (Lines 564–576)

7. Priorities for Revision and Recommendations

The core finding — that sub-hectare disturbances have a greater impact on Central African forest carbon dynamics than current monitoring systems reflect — is almost certainly correct in its overall direction, even if the exact measurement remains uncertain. Therefore, the formal accuracy assessment of small- disturbance detection is the most important revision to make. This assessment will likely require new analysis — stratified sampling, independent VHR interpretation, confusion matrices — and is consequently the most time-consuming revision to undertake. It should be initiated first.

The sensitivity analysis of filtering thresholds is the second priority. It is essential for establishing the robustness of the headline results and is methodologically straightforward: the existing pipeline is rerun with varied parameters, and the outputs are compared. If the results prove stable across reasonable threshold ranges, this analysis provides strong support for the approach; if not, the uncertainty estimates must be increased.

Third in priority is the uncertainty framework itself, which requires replacing the circular classification assessment, incorporating spatial autocorrelation, and justifying or replacing the borrowed calibration uncertainty. These are analytical revisions that strengthen the paper's quantitative foundation.

The fourth priority — tempering the interpretive claims throughout the abstract, results, and discussion — is fairly straightforward to implement and should accompany the analytical revisions. The remaining issues — a more transparent discussion of the allometric model's limitations and an updated, balanced introduction — are strongly recommended improvements that would improve the overall quality and credibility of the manuscript, even though they are less likely than the first four to determine the editorial outcome on their own.

In summary, the revision must deliver four things. First, a formal, statistically rigorous accuracy assessment of sub-hectare disturbance detection, grounded in independent reference data, with false-positive rates reported by patch-size class. This is the single analysis most likely to determine whether the paper can be published. Second, a sensitivity analysis demonstrating the dependence — or, ideally, the independence — of the headline results on the noise-filtering thresholds. Third,

a revised uncertainty framework that eliminates the circularity in the classification assessment, accounts for spatial autocorrelation, and replaces the borrowed calibration uncertainty with a properly justified value. Fourth, recalibrate the interpretive language throughout the manuscript to clearly and consistently distinguish between what is robustly established and what remains quantitatively uncertain.

Response: We appreciate this summary of the key revisions required. We have revised the manuscript to address all four points. Specifically, we have (i) conducted an independent accuracy assessment of disturbance detection, (ii) added sensitivity analyses of the filtering thresholds, (iii) revised the uncertainty framework to remove circularity and better account for spatial dependence, and (iv) adjusted the interpretive language throughout the manuscript to clearly distinguish robust findings from more uncertain results. Detailed responses to each point are provided below.

Specific Comments (Line-by-Line)

All comments refer to the line numbers in the submitted manuscript. Comments are organised by section.

ABSTRACT (Lines 36–52)

Line 36–38. The term "clearings" is vague and overlaps with logging and deforestation. If it refers to agricultural conversion, please specify it. If it includes shifting cultivation, that is a distinct process with its own regrowth dynamics. Framing all disturbances as external "pressure" does not recognise that many small-scale clearings are part of long-standing subsistence land-use systems.

Response: We have revised the sentence accordingly to clarify the disturbance types and removed the wording referring to external "pressure" as:

Dense humid forests in Central Africa hold vast biomass carbon stocks but are increasingly affected by logging, smallholder agricultural clearing, mining activities, and fire disturbances. (Lines 35–36)

Line 38–39. "Missed by current satellite observations" conflates two separate issues: (a) disturbances genuinely below the detection limit of existing sensors, and (b) disturbances whose structural signals are too small to reliably distinguish from noise in any satellite system, including 10 m Sentinel data. The paper struggles with this distinction throughout, but does not highlight it here.

Response: We have revised the abstract to explicitly distinguish between these two cases:

Many of these disturbance events occur at small spatial scales and produce small structural changes that are often below the detection limits of current satellite observations or cannot be reliably distinguished from noise. (Lines 36–39)

Lines 41–43. The 87% figure is expected from any high-resolution disturbance map where rasterised disturbance patches are spatially dispersed. Small patches will always vastly outnumber larger patches. The key policy question is whether these patches contribute disproportionately to carbon flux. The 37% and 48% figures partially address this, but are presented without clarification about the noise-filtering thresholds that removed roughly 86% of the smallest height-change pixels.

Response: We have clarified that this percentage refers to detected disturbances from processed change maps and have focused on their contribution to carbon flux:

The processed and noise-filtered change maps reveal widespread small disturbance patches, with 87% of the detected disturbances being smaller than 1 ha. Despite their small size, these events account for 48% of biomass carbon gains in regrowing forests and 37% of carbon losses from deforestation. (Lines 42–45)

Lines 43–45. Uncertainty in Net Flux is likely to be incorrect. The reported uncertainty on the net flux (± 9 Tg C yr⁻¹) happens to match the value you would obtain by combining the gross gain (± 6) and gross loss (± 7) uncertainties in quadrature (i.e., $\sqrt{6^2 + 7^2} \approx \pm 9.2$). Quadrature addition, however, is only valid when the two error sources are independent — when the things that can go wrong in one estimate have no bearing on what goes wrong in the other. The reported uncertainty on the net flux (± 9 Tg C yr⁻¹) equals the value obtained by adding the gross gain (± 6) and gross loss (± 7) uncertainties in quadrature:

$$\sigma_{net}^{indep} = \sqrt{\sigma_{loss}^2 + \sigma_{gain}^2} = \sqrt{7^2 + 6^2} \approx \pm 9.2$$

This formula assumes the two uncertainties are independent. That assumption is violated here: gross gains and gross losses are derived from the same canopy height maps, allometric model, noise-filtering pipeline, and land-cover classification. Systematic errors in these shared components affect both quantities simultaneously and in the same direction, introducing a positive correlation ($\rho > 0$) between their errors. The correct propagation for the uncertainty on a difference (Net = Loss – Gain) is:

$$\sigma_{net} = \sqrt{\sigma_{loss}^2 + \sigma_{gain}^2 - 2\rho\sigma_{loss}\sigma_{gain}}$$

When $\rho > 0$, the corrective term is positive, so σ_{net} exceeds the quadrature sum — the opposite of what the current estimate implies. The reported ± 9 Tg C yr⁻¹ therefore represents a lower bound, not a best estimate. The fact that the reported ± 9 aligns with the independence assumption, therefore, suggests the correlation between gain and loss errors has not been accounted for. The true uncertainty on the net flux should be meaningfully larger than ± 9 Tg C yr⁻¹, and the authors should either derive a propagated uncertainty that accounts for this covariance or explicitly acknowledge that ± 9 is a lower bound. Either propagate the uncertainty accounting for the covariance between gross gain and gross loss errors, or explicitly state that the reported value is a lower bound and discuss the expected direction and magnitude of the bias.

Response: We acknowledge the fact that gain and loss are not independent. Both are derived from shared data sources and modeling components. We have revised our uncertainty analysis to explicitly account for multiple sources of uncertainty, including biomass estimation model, disturbance classification, and change detection algorithm (filtering threshold parameters). These component uncertainties were jointly propagated through a Monte Carlo framework, in which both model parameters and class assignments were perturbed. This approach produces an ensemble of realizations of biomass gain, loss, and net change across disturbance size classes, from which the covariance structure among components emerges naturally. In addition, the Monte Carlo simulations allow us to quantify correlations among different disturbance and regrowth classes, rather than assuming their independence. These empirically derived correlations are then

incorporated into the uncertainty propagation, providing a more realistic estimate of uncertainty at both class and aggregated (gain, loss, and net) levels.

The results about the carbon change have been updated as:

We found a net carbon loss of $-58 \pm 14 \text{ Tg C yr}^{-1}$ composed of a gross loss of $-126 \pm 21 \text{ Tg C yr}^{-1}$ partially offset by a gain of $68 \pm 10 \text{ Tg C yr}^{-1}$, implying a turnover of biomass carbon losses of $0.52 \pm 0.10\%$ per year. (Lines 45–47)

Full methodological details of uncertainty formulation, parameterization, and propagation are provided in **Supplementary Text 10, Supplementary Fig. 30**. The resulting total relative uncertainties are 23.2%, 16.3%, and 14.7% for net AGC change, gross AGC loss, and AGC gain in Central African forests, respectively (**Extended Data Table 2**). (Lines 808–811)

Lines 46–47. Characterising -46 Tg C yr^{-1} in the DRC as a "small net source" is misleading. This is ~79% of the total regional net loss. The word "small" appears intended relative to the DRC's total stock, but this is not stated, and a reader would interpret it as unimportant.

Response: We agree and have removed the term “small” and updated the text as: **The Democratic Republic of the Congo is a net source of carbon ($-46 \pm 11 \text{ Tg C yr}^{-1}$), representing 79% of the regional loss** (Lines 47–48)

Lines 48–49. "Uneven conservation outcomes" constitutes an interpretive leap. Losses in protected areas might stem from natural disturbances, edge effects, or legacy degradation. Without the ability to attribute these drivers — which the authors cannot do — this framing is unsuitable for the abstract.

Response: We have removed the interpretive wording (**highlighting uneven conservation outcomes**).

Lines 50–52. “Align well” overstates the agreement. It is nearly accurate for Gabon and Equatorial Guinea, somewhat reasonable for Cameroon, but clearly inconsistent for the DRC (>70% of regional forest area and carbon flux). This varied agreement does not justify saying “align well.”

Response: We have revised the text to avoid overstating the agreement with national inventories and bookkeeping estimates as: **Our remote-sensing data aggregated into national carbon budgets show broad consistency with country-level inventories and bookkeeping model estimates.** (Lines 51–53)

INTRODUCTION (Lines 55–104)

Lines 55–58. The claim that intact forests are "large carbon sinks" cites Lewis et al. (2009), which is now outdated. Hubau et al. (2020, Nature) — not cited — shows the African intact forest sink has been declining. The paper's own data (Supplementary Table 4: near-zero or negative growth in undisturbed forests) contradicts this framing. The introduction should address this issue.

Response: We agree that the carbon sink in intact African forests has weakened in recent decades. We have revised the text to reflect this updated and add new references from Hubau et al. (2020), Rodríguez-Veiga et al. (2025), and Sagang et al. (2026) as:

Central African forests play a vital role in the global carbon cycle, storing approximately 10% of the world's forest biomass carbon ¹⁻³. Although intact forests have historically acted as a major

carbon sink⁴, recent studies suggest that this sink is weakening, tending toward carbon neutrality or even becoming a net source⁵⁻⁷. (Lines 57–60)

Lines 55–58. The parenthetical "(see definition in Methods)" for degradation is inadequate. If the definition is critical to the paper — and it is — it should be stated in the introduction, at least in abbreviated form.

Response: We have included a brief definition in the Introduction as:

Disturbances from deforestation, degradation (canopy disturbances that reduce biomass without converting forest to non-forest⁸), and fragmentation are increasingly affecting these forests^{9,10}. (Lines 60–62)

Lines 58–62. Grouping sustainable forestry, restoration, plantations, and shifting-cultivation regrowth together obscures fundamentally different processes with distinct carbon trajectories and policy implications. The paper's classification scheme does not distinguish among them, which should be acknowledged.

Response: Our intention was to provide examples of processes that can lead to forest carbon gains rather than to classify them. We have revised the text as:

Carbon gains arise from sustainable forestry, restoration, plantations and forest regrowth following clearing events from shifting cultivation and changes of environmental factors like rainfall, yet these gains remain difficult to detect and quantify using satellite observations. (Lines 62–65)

We did not explicitly attribute the drivers of biomass change to fire, windthrow, plantation expansion, or afforestation⁶³, due to the short time window covered by our maps and the lack of detailed and spatially explicit ancillary data. (Lines 475–477)

Lines 62–64. The sentence "logging facilitated by road expansion and land conversion to agriculture have accelerated the clearing of forests" suggests a temporal trend ("accelerated"), but the study only covers 2019–2022 and does not provide data for multiple decades. The word "accelerated" is not supported by the data in this paper.

Response: We have removed the sentence to avoid suggesting a trend beyond the scope.

Lines 64–65. "Intensified forest fires activity" overstates the cited source (Zhao et al. 2021, ref 12), which focused on fire-driven degradation at forest edges, not a general intensification across the region. Fire is not explicitly detected or attributed in this study, making the reference somewhat misleading.

Response: We have removed this sentence including the cited source.

Lines 69–74. The critique of existing approaches (inventories, bookkeeping models, DGVMs) is fair but biased. No limitations of the authors' own stock-change approach are acknowledged. The reader is led to believe that remote sensing inherently solves these issues. A more balanced framing would recognise that each approach has its own strengths and limitations. Additionally, "DGVM models" is redundant, as "model" is already included in the acronym.

Response: We have revised the relevant description in the Introduction. We also corrected the wording “DGVM models” to “DGVMs”. The sentences have been rewritten as:

NGHGs provide nationally consistent estimates defining the UNFCCC reporting, while bookkeeping models track carbon fluxes associated only with land-use transitions¹⁶. Both approaches are generally not spatially explicit, limiting their verification on the ground. DGVMs simulate large-scale vegetation dynamics and climate–carbon interactions, but they have unconstrained parameters in Africa and oversimplify disturbances, for example neglecting forest degradation. Building on these approaches, recent synthesis efforts such as the RECCAP2¹⁸ analysis by Ernst et al. ¹⁹, although important uncertainties remain, particularly in resolving disturbance processes. (Lines 73–80)

Direct mapping of biomass carbon stock change from satellite observations can complement inventories and models based on growth-rates by providing direct spatially explicit observations of forest carbon changes ^{8,23,24}. However, these approaches are also subject to uncertainties from sensor noise, errors in height-biomass allometric relationships, limited sensitivity to understory changes, and potential biases when the same algorithms are applied across structurally diverse forest types^{25,26}. Despite these limitations, maps of stock changes offer a unique opportunity to resolve the spatial distribution, size, and type of disturbance events^{27,28}, which are often not captured by existing approaches. (Lines 88–95)

Lines 76–80. The critique of Harris et al. is valid but does not address the recent update (Gibbs et al. 2025, ref 22, cited elsewhere in this paper). Ignoring the latest version may make the gap seem larger than it actually is.

Response: We agree that the recent update presented in Gibbs et al. (2025) provides an important refinement of the Harris et al. framework. We have revised this section to acknowledge the updated estimates and improvements as:

Harris *et al.* ²⁰ estimated a net carbon sink of -61 Tg C yr^{-1} in Central African forests over the last two decades by using fixed forest growth rates and prescribed forest area losses from Landsat-derived tree cover loss. Gibbs *et al.* ²¹ further refined this framework by improving model inputs, reducing uncertainty in carbon removal factors, and extending the time series of forest carbon fluxes. (Lines 82–86)

Lines 80–88. "Direct mapping of biomass carbon stock change from satellite observations can overcome limitations" is too strong. "Can complement" or "can partially overcome" would be more accurate. The claim that previous maps show "conflicting results" is not elaborated — the reader does not learn how or why they conflict.

Response: We have clarified the inappropriate statement as:

Direct mapping of biomass carbon stock change from satellite observations can complement inventories and models based on growth-rates by providing direct spatially explicit observations of forest carbon changes ^{8,23,24}. (Lines 88–90)

Existing biomass change maps relying on coarse-resolution satellites (500 m to 25 km) have shown inconsistent estimates of carbon sources and sinks across tropical regions ^{3,29,30}. (Lines 95–97)

Lines 90–95. The resolution switch from "10 m" CH to "30 m" AGC is introduced without explanation. The word "integrate" for the deep learning approach overstates what happens: the model is trained to predict GEDI RH100 from Sentinel imagery, which is prediction/interpolation, not data fusion.

Response: We clarified both points in the revised text. Canopy height is predicted at 10 m resolution and then aggregated to 30 m before converting to AGC, because biomass calibration relies on field plot observations and large tropical tree crowns often exceed ~20 m, making estimates of biomass at 10 m unstable. We also replaced the word "integrate": GEDI RH100 is used as the training target in the deep learning model to predict canopy height from Sentinel-1 and Sentinel-2 imagery, rather than being directly integrated with the imagery.

Lines 95–97. "Detailed validation" and "to ensure reliable estimates" are self-congratulatory language for an introduction. Whether validation is detailed and whether it ensures reliability are for reviewers to judge. The validation results ($R^2 = 0.53$ at ALS, $R^2 = 0.22$ at GEDI) do not "ensure" reliability.

Response: We have revised the sentence to avoid overly strong wording as: **First, we validated the CH and AGC change estimates for the period 2019–2022 using different reference datasets (Supplementary Text 5–Text 7).** (Lines 113–114)

Lines 97–104. The research question "Do our high-resolution maps detect more small-scale disturbances?" is self-fulfilling — a 10 m product will always detect more small features than a 30 m product. The meaningful question is whether these additional detections are real and whether they affect the carbon budget. The extension to Miombo woodlands is introduced casually but represents a significant methodological assumption (the use of a dense-forest-calibrated allometric model across structurally distinct ecosystems).

Response: We have revised this section to clarify the research questions as:

First, we validated the CH and AGC change estimates for the period 2019–2022 using different reference datasets (Supplementary Text 5–Text 7). We then analysed carbon losses and gains across different sizes of disturbances to address the following questions: Do the small-scale disturbances revealed by our high-resolution maps represent real disturbance events that were overlooked by existing Landsat/Sentinel-1-based datasets^{8,23,24}; How do these small disturbances contribute to carbon losses and gains associated with degradation, deforestation and regrowth. Finally, we extended the approach to non-dense forest systems such as Miombo woodlands to obtain national budgets. This extension is based on the key methodological assumption that our height–biomass relationships calibrated for dense forests hold for non-dense systems, which introduces additional uncertainty in Miombos and Woodlands north and south of the Congo basin. To address this, as we lack field data in Miombos and non-dense systems, we assessed the robustness of our biomass estimates against independent satellite-based biomass products (Supplementary Text 12). (Lines 113–125)

RESULTS — Evaluation of Canopy Height Changes (Lines 106–127)

Lines 112–114. Reporting $R^2 = 0.53$ (MAE = 3.19 m) against ALS and $R^2 = 0.22$ (MAE = 4.39 m) against GEDI without explicitly stating that these values indicate moderate-to-weak agreement may create the false impression that they are acceptable. For change detection — where the signal of interest (small height changes of 2–5 m) is similar to the MAE — these errors have significant implications.

Response: We revised the text to clarify the interpretation of the validation results as:

Validation was primarily conducted at 30 m resolution due to scale compatibility with ALS data and consistency with biomass change data used in subsequent analyses, $R^2 = 0.53$ and MAE = 3.19 m (**Extended Data Fig. 2, Supplementary Fig. 11**). Agreement at 10 m was lower ($R^2 = 0.22$, MAE = 3.70 m) than at 30 m due to increased fine-scale heterogeneity and unresolved sub-pixel canopy gaps relative to ALS data, although MAE remained comparable across resolutions. Aggregation to 30 m improves the agreement through spatial smoothing, whereas the 10 m product preserves fine-scale canopy disturbances. Comparisons with changes from GEDI discrete footprint samples (RH100) between the first and last years of our study period showed weak agreement at 30 m but improved when aggregated to 10 km ($R^2 = 0.23$, MAE = 4.35 m) (**Extended Data Fig. 2, Supplementary Fig. 14**). This reflects several factors, including scale mismatch, sparse GEDI footprint sampling with no revisit in time, and geolocation uncertainty⁴², as well as the difficulty of detecting small height changes and uncertainties in our biomass mapping at fine spatial scales. (Lines 138–150)

Lines 115–117. The weaker agreement with GEDI in smaller areas is attributed to "the small number of footprints and geolocation errors." While true, this also means the product has essentially no validated accuracy at the spatial scales where the novel claims are made (<1 ha patches).

Response: We agree that validation of canopy height change at sub-hectare scales using GEDI is inherently limited, due to the sparse footprint sampling and geolocation uncertainties, which make it difficult to assess high-resolution changes directly. Although we also used high-resolution airborne LiDAR (ALS) data for validation, this comparison still relies on several processing assumptions and therefore cannot provide a fully independent assessment of fine-scale changes. To address this limitation, we conducted an independent validation using high-resolution (5 m) PlanetScope imagery with a stratified sampling design⁴¹ across patch sizes and change types

Lines 123–125. Visual validation over "250 random small areas", finding "80% agreement for loss and 66% for gain patches when accounting for trends consistent with TMF disturbances" — the phrase "when accounting for" is doing substantial work. Without the TMF-consistency adjustment, the raw agreement is 66% for losses and 56% for gains, with 14–8% of cases showing changes opposite to the CH signal. These raw numbers should be prominently stated.

Response: We have updated the visual validation in the revised manuscript as:

To assess the reliability of mapped CH changes, we conducted an independent validation using high-resolution monthly PlanetScope imagery (5 m) (**Supplementary Text 6**). We adopted a stratified random sampling design following commonly used good practices for land-change accuracy assessment⁴³, based on patch-size classes (0–1, 1–5, 5–10 and 10–100 ha) and change type (gain and loss). Samples were randomly selected across the study region, and only locations

with sufficiently cloud-free PlanetScope imagery in both 2019 and 2022 were retained for interpretation. We visually interpreted whether sampled patches exhibited canopy structural decrease or increase consistent with the mapped CH change direction. Interpretation was based on comparisons between monthly PlanetScope imagery from 2019 and 2022 together with the mapped CH change layer, including changes in canopy cover, shadowing, canopy openings and vegetation recovery signals. We included a total of 422 validated loss patches and 498 validated gain patches. Small disturbance patches, which are particularly difficult to identify because of subtle canopy structural changes, were independently interpreted by both an experienced forest remote sensing expert and a trained analyst following the same interpretation protocol. Interpretation consistency across the same validation samples reached 94%. Larger disturbance patches (1–100 ha) were interpreted by the trained analyst using the same protocol. Finally, we derived user's accuracy for each size class as well as area-weighted accuracy⁴³ to account for the disproportionate contribution of large patches to total mapped area change (**Extended Data Table 1**). In addition, omission errors were assessed using randomly sampled 2-ha reference patches selected from undisturbed forest areas. These reference areas were visually confirmed as stable forest using PlanetScope imagery together with independent disturbance datasets. (Lines 623–643)

This analysis indicated that the user's accuracy (UA) was 84% for loss and 77% for gain across the 0–100 ha range (area-weighted), and 79% and 69%, respectively, for small disturbances (<1 ha) (**Extended Data Table 1**). Omission and commission errors were ~6% and ~5%, respectively, when evaluated relative to the full landscape, corresponding to undetected changes within no-change areas and false detections in stable forests. (Lines 163–168)

RESULTS — High Resolution Maps Reveal More Small Disturbances (Lines 142–169)

Lines 147–151. The statement "80% of the areas where we detect large height losses (> 20 m) are matching TMF events" is reassuring, but the main concern pertains to small losses rather than large ones. The following sentence states "65% larger area with small height losses (< 10 m) than reported in the TMF" — however, whether these additional detections are genuine is supported only by visual interpretation of Planet imagery (80% confirmed), including an admittedly generous classification of uncertain cases.

Response: We have clarified the interpretation of these results and emphasized that the comparison with TMF mainly shows consistency for large disturbances. We also clarified that the evidence for additional small disturbances relies on visual interpretation and acknowledged that some uncertainty remains for individual cases.

Quantitative analysis confirms that 80% of the areas where we detect large height losses (> 20 m) are matching TMF events at 30 m resolution (**Supplementary Fig. 17**), indicating consistency for large disturbances. The main difference arises for small height losses (<10 m), where our maps detect 65% more affected area than reported in the TMF deforestation and degradation products. Independent visual validation using high-resolution PlanetScope imagery confirmed that 84% of these additional small-loss events represent real canopy disturbances (**Extended Data Table 1**). However, only a subset of these validated small CH loss patches overlapped with TMF disturbances (26%–78%), depending on the overlap threshold applied (**Supplementary Table 1**). (Lines 190–198)

Lines 156–161. The finding that "90% larger area affected by height loss than the Global Forest Change (GFC) map and the radar-based deforestation alerts (RADD)" could indicate genuine improvement or increased false positives. The claim that these products are "designed to detect tree cover loss and thus more prone to seasonal noise" is not well-supported — signal saturation is a sensor property, not a product design feature, and seasonal noise affects Sentinel-based products equally.

Response: We have clarified the comparison with existing disturbance datasets as:

We also found a 90% larger area affected by height loss in our product than in the Global Forest Change (GFC) tree-cover map²³ and the radar-based deforestation alerts (RADD)²⁴ (**Fig. 2a**). Here, independent visual validation using high-resolution PlanetScope imagery showed that GFC and RADD captured less than 53% and 26%, respectively, of validated small CH loss patches (<1 ha) (**Supplementary Table 1**), indicating limited sensitivity of existing products to partial canopy loss and small-scale degradation^{46–48}. (Lines 205–211)

Lines 163–164. "87% of the pixels experiencing a height decrease from 2019 to 2022 are located in small patches ranging from 0.09 to 1 ha." This statistic should be put into context: what proportion of the total area or total carbon loss do these patches account for? The number of patches is less informative than their combined contribution.

Response: We have added relevant results and revised the sentence as:

In our maps, ~87% of patches (~21% of the total loss area) exhibiting CH decrease from 2019 to 2022 fall within the smallest height-loss size class (<1 ha; **Fig. 2b**). (Lines 213–214)

Lines 165–166. "Among 10 million such small height loss patches of 0.09 ha, only 56% are seen in the TMF dataset." The fact that 44% are novel detections is presented as evidence of TMF's limitations, but it could equally indicate a 44% false-positive rate. Without independent ground truth, both interpretations are equally valid.

Response: We have revised this sentence to avoid attributing the differences solely to limitations of the TMF dataset as: Among the ~10 million small height-loss patches of 0.09 ha identified in our maps, 56% coincide with disturbances recorded in the TMF dataset. The remaining differences may reflect either additional small disturbances not captured by TMF³⁴ or uncertainties in our detection of small changes. (Lines 214–217)

RESULTS — Size-Dependence Patterns (Lines 182–223)

Lines 188–191. "The size distribution of deforestation and degradation events is dominated by small patches (mean = 0.9 ha), with very few large-scale clearings of up to 10,000 ha." The mean patch size of 0.9 ha is stated without confidence intervals. Considering the dominance of very small patches (0.09 ha), the mean is probably heavily skewed. The median patch size would provide more informative insights.

Response: We have now included the median patch size in the revised manuscript as:

The size distribution of deforestation and degradation events is dominated by small patches (mean = 0.9 ha; median = 0.18 ha), indicating a strongly right-skewed distribution, with very rare

extremely large-scale clearings (maximum size 10,000 ha) and many small clearings. (Lines 245–248)

Lines 192–193. "The fraction of biomass lost ranges from 83% to 88%." This applies to full-canopy patches without edge removal. Supplementary Fig. 23 shows that the apparent loss fraction is much higher (up to 89%) when edge pixels are excluded, suggesting that edge-pixel mixing significantly dilutes the deforestation signal. This has implications for the accuracy of patch-level deforestation carbon loss estimates.

Response: We agree that edge-pixel mixing can influence estimates of biomass loss and have revised the sentence as:

Relative to the initial biomass density in 2019, the fraction of biomass lost after deforestation increases with patch size. Biomass loss is even higher when excluding edge pixels, reaching up to ~89%, suggesting that patch edge effects dilute the biomass loss signal because edge pixels contain cut and intact trees (**Supplementary Text 8, Supplementary Fig. 27**). (Lines 248–252)

Lines 196–198. "Moderate degradation events of ≈ 10 ha and small deforestation events of ≈ 1 ha make the largest contribution to the total biomass loss." This is a key finding but depends entirely on the accuracy of the disturbance classification at these scales. The classification scheme (CH-based thresholds + XGBoost extension of TMF labels) has not been validated against independent data for these intermediate patch sizes.

Response: We first clarify that the original estimate for small deforestation events is 1.48 ha, which we now report as ~1.5 ha in the revised manuscript to ensure consistency and precision. In our framework, deforestation and degradation losses are defined using physically interpretable canopy height (CH) change detection with a minimum threshold of CH in 2022 (CH_{2022}), while gain after degradation is derived from TMF data and an XGBoost-based extension. We acknowledge that the gain component has additional uncertainty.

To directly address the concern regarding classification accuracy at these spatial scales, we conducted a large independent validation using high-resolution (5 m) PlanetScope imagery (**Supplementary Text 6**). We adopted a stratified sampling design based on patch size (0–1, 1–5, 5–10, and 10–100 ha) and change type (gain and loss), and visually interpreted randomly selected samples. This analysis indicated that the user's accuracy (UA) was 84% for loss and 77% for gain across the 0–100 ha range (area-weighted), and 79% and 69%, respectively, for small disturbances (<1 ha). Omission and commission errors were ~6% and ~5%, respectively, when evaluated relative to the full landscape, corresponding to undetected changes within no-change areas and false detections in stable forests.

Please refer to visual validation results in Extended Data Table 1 and our response to the related comment in the General Overview section above.

In addition, we assessed the robustness of the identified patch sizes through a set of sensitivity analyses (Supplementary Text 4.5; Supplementary Table 2).

(1) Sensitivity to the CH_{2022} threshold. We varied the CH_{2022} threshold from 10 to 20 m while keeping the CH loss minimum threshold fixed ($CH\ loss < 5\ m$). This range is consistent with commonly used definitions of forests (e.g., $\geq 5\ m$ canopy height) and previous studies. Across this range, small deforestation events of $\sim 1.5\ ha$ consistently make the largest contribution to total biomass loss, indicating that this dominant patch-size signal is robust to threshold choice. For degradation events, the dominant patch size is $\sim 10.2\ ha$ for CH_{2022} thresholds between 10 and 15 m, while greater variability emerges at higher thresholds (15–20 m), where the dominant patch size decreases to $\sim 5.4\ ha$.

(2) Sensitivity to the definition of areas with gain after degradation events (TMF only vs. TMF + XGBoost). We compared results derived from (i) TMF only and (ii) TMF combined with our XGBoost classification extension. The inclusion of the XGBoost-based gain leads to more consistent estimates of characteristic patch size, typically within $\sim 5\text{--}10\ ha$ across thresholds, whereas TMF-only results show larger variability. This suggests that the extended number of gain areas brought by the XGBoost classification extension dataset improves the consistency of ‘degradation gains’ characterization by capturing regrowth signals not represented in TMF alone.

We revised the original sentence as: The key finding is that moderate degradation events from 1.5 to $\sim 10\ ha$ and small deforestation events of $\sim 1.5\ ha$ make the largest contribution to the total biomass loss in the study region, although uncertainty remains due to limitations in disturbance detection and classification, particularly for small disturbances (**Supplementary Table 2**). (Lines 255–258)

Lines 217–219. "High AGC growth rates are sustained for two years, followed by declining growth rates thereafter." This temporal pattern depends on TMF legacy data, which has known detection limitations in the early 2000s due to poor Landsat availability in Central Africa (acknowledged in lines 582–584). The growth rates in the early period may therefore be based on a biased subset of disturbances (only those large or clear enough to be detected despite data limitations).

Response: This limitation about the early TMF data has been added to the manuscript:

For regrowing forests, high AGC growth rates are sustained for two years after regrowth is detected, followed by declining growth rates thereafter (**Fig. 3b**). This temporal dynamics of growth rates is however sensitive to TMF legacy data, as early-period TMF estimates may be biased toward easier to detect larger disturbances due to limited Landsat data availability in the early 2000s. (Lines 279–283)

Lines 220–222. "Forests that experienced former degradation events show a large initial loss of AGC density which attenuates over time, yet with no return to a carbon gain regime even after 20 years." This is a significant finding, but the observation period for any individual pixel is only 3 years (2019–2022). The "20- year trajectory" is reconstructed by observing different cohorts of disturbances at various ages — a space- for-time substitution that assumes no secular trends in degradation severity, recovery conditions, or climate over the past two decades.

Response: We have clarified this limitation in the manuscript as:

In contrast, forests that experienced degradation in the past show a large initial loss of AGC density that attenuates over time. Based on this space-for-time substitution approach that assumes no secular trends in degradation severity, recovery conditions, or climate over the past two decades, we found no clear evidence of a return to a carbon gain regime in the first ~20 years after degradation (**Fig. 3d**). (Lines 285–289)

RESULTS — Canopy height and biomass changes and Forest Management (Lines 239–273)

Lines 240–243. The net AGC loss of $-58.3 \pm 9.3 \text{ Tg C yr}^{-1}$ is reported with two significant figures for the uncertainty (9.3), suggesting a level of precision that is not justified given the numerous unquantified sources of error discussed above.

Response: We have improved the uncertainty analysis framework by explicitly quantifying three major components: (1) biomass uncertainty arising from allometric modelling and measurement errors (σ_{bio}); (2) classification uncertainty associated with disturbance attribution and detection accuracy (σ_{cls}); and (3) uncertainty related to the choice of a CH change processing scheme (σ_{scheme}), including trend detection, Sentinel-1 filtering, and canopy cover filtering.

Please refer to our responses to the [related comments in the General Overview section and Comment 2](#) above regarding uncertainty estimation.

Based on this revised framework, the estimated net AGC loss is $-58.3 \pm 13.5 \text{ Tg C yr}^{-1}$, reflecting a more comprehensive quantification of uncertainty. The original sentence has been revised as:

Scaling up to the entire Central African forest biome (**Fig. 4a**), our data show a small net AGC loss of $-58.3 \pm 13.5 \text{ Tg C yr}^{-1}$ (see uncertainty estimates in **Methods, Supplementary Text 10**). (Lines 308–309)

Lines 244–245. "The turnover of carbon with respect to losses from disturbances is thus 0.52%." This derived value incorporates all the uncertainty of both the gross loss and the total stock estimates but is presented without an uncertainty range.

Response: Thank you for this comment. We have now included an uncertainty estimate for the turnover rate by propagating uncertainties in annual carbon loss and total carbon stock using error propagation. Specifically, we assume that uncertainties in annual AGC loss and baseline AGC stock are independent. Although both quantities are derived from related biomass maps, their uncertainties arise from different components of the modelling framework, and this assumption provides an approximation for uncertainty propagation in ratio estimates. Based on the revised uncertainty estimation, the result has been updated as:

With a total stock of 24.3 Pg C in 2019, the turnover rate associated with disturbance-driven losses is $0.52\% \pm 0.10\%$ per year (**Supplementary Text 11**). (Lines 311–312)

Lines 248–250. "The AGC budget of other forests, classified as undisturbed in TMF, is nearly neutral ($-3.6 \pm$ although these forests still exhibit dynamic gross gains and losses." The uncertainty range includes zero, so the direction of change remains uncertain. More importantly, Extended Data Table 1 shows gross losses of $-22.4 \text{ Tg C yr}^{-1}$ and gross gains of $18.8 \text{ Tg C yr}^{-1}$ in these

"undisturbed" forests — significant fluxes that could reflect natural dynamics or classification noise. This ambiguity deserves more discussion.

Response: We agree that, given the associated uncertainty, the net AGC change in forests classified as undisturbed is not statistically distinguishable from zero. We have revised the text to clarify:

The AGC budget of so-called ‘other forests’, classified as apparently undisturbed in TMF but not exempt of AGC gains and losses, is close to neutral and not distinguishable from zero within uncertainty (**Extended Data Table 2**), although these forests exhibit substantial localized gross gains and losses. Their loss and gain fluxes in ‘other forests’ likely reflect a combination of real features (gaps, natural disturbances, human disturbances missed by TMF) and uncertainties in our disturbance classification and CH change data (**Extended Data Table 1**). This result highlights the overlooked dynamics of carbon in forests classified as apparently undisturbed. (Lines 316–323)

Lines 252–254. "A net biomass carbon loss of -3.8 ± 1.2 Tg C yr⁻¹ is found in protected areas, and a loss of -5.4 ± 1.2 Tg C yr⁻¹ in concessions." These are stated without comparison to the baseline loss rate outside protected areas and concessions. Without this comparison, the reader cannot assess whether protected areas are performing better or worse than the unprotected landscape.

Response: To provide a clearer baseline for comparison, we added a new table (**Supplementary Table 1**) summarizing aboveground carbon (AGC) changes across different land-use categories, including protected areas, concessions, non-protected areas, and non-concession areas. The results show that protected areas exhibit substantially lower biomass density carbon loss (-0.09 ± 2.40 Mg C ha⁻¹ yr⁻¹) compared to non-protected areas (-0.31 ± 3.55 Mg C ha⁻¹ yr⁻¹), indicating that protected areas are indeed effective in reducing carbon loss, although they still experience net declines.

We have revised the main text (Lines 325–329) to explicitly include these comparisons as:

Protected forest areas and forestry concessions together represent 34% (73.5 Mha) of the total dense forest area. In protected areas, we assess a net biomass carbon loss of -3.8 ± 1.4 Tg C yr⁻¹ (mean rate = -0.09 ± 2.40 Mg C ha⁻¹ yr⁻¹), substantially lower than that observed in non-protected areas (-0.31 ± 3.55 Mg C ha⁻¹ yr⁻¹). In concessions, we found a net loss of -5.4 ± 1.6 Tg C yr⁻¹ (mean rate = -0.15 ± 2.88 Mg C ha⁻¹ yr⁻¹) (**Fig. 4c–f; Supplementary Table 3**).

Lines 262–263. "The mean loss rate of biomass per unit forest area is 61% larger in concessions than in protected areas, indicating the effectiveness of protection measures." This is a confounded comparison. Concessions and protected areas differ in location, accessibility, forest type, surrounding land use, and disturbance history. Attributing the difference to "effectiveness of protection measures" without controlling for these confounders is not justified.

Response: Thank you for this important comment. We have revised the manuscript to avoid attributing observed differences directly to the effectiveness of protection measures. We have updated the estimated difference based on the revised analysis. The mean biomass loss rate per unit forest area is now ~67% higher in concessions than in protected areas (previously reported as 61%). The sentences are rewritten as:

The mean loss rate of biomass per unit forest area is ~67% larger in concessions than in protected areas (**Supplementary Table 3**). Concessions and protected areas differ in environmental and socio-economic contexts, including accessibility, forest types, surrounding land use, and disturbance history (**Supplementary Text 9**). Therefore, this comparison does not imply a causal attribution to protection status or effectiveness alone. (Lines 338–342)

Lines 268–273. The discussion of why concessions have higher AGC density than protected areas offers two explanations (government allocation of biomass-rich forests to concessions; prior degradation in some protected areas) but does not consider selection bias in protected area placement (some areas may have been protected specifically because of biodiversity values in naturally lower-biomass ecosystems).

Response: We have revised the manuscript to explicitly acknowledge that protected areas may be preferentially located in ecosystems selected for biodiversity richness or other ecosystem services rather than for a high biomass, which can include naturally lower-biomass forests. This provides an alternative explanation for the lower mean AGC density observed in protected areas relative to concessions. In addition, we clarified the fact that these differences in biomass are not intended to imply a causal effect of land-use designation on biomass levels. The original text has been updated as: (Lines 346–355)

Interestingly, protected areas have a smaller mean AGC density ($127 \pm 47 \text{ Mg C ha}^{-1}$) than concessions ($144 \pm 44 \text{ Mg C ha}^{-1}$) (**Extended Data Fig. 7**). However, this difference is modest relative to the uncertainty in AGC estimates and is accompanied by substantial overlap between distributions, indicating that the two land-use types are only weakly separable in terms of obvious influence on mean biomass density. This pattern may partly reflect non-random spatial placement of protected areas, which may include lower-biomass ecosystems due to non-carbon centric choices for deciding a protection status (cultural value, animal and plant diversity). It may also be influenced by differences in management regimes, for example where logging in concessions is relatively low-intensity or in more fertile forests, and allows trees to retain or recover biomass¹².

RESULTS — National Carbon Budgets (Lines 290–336)

Lines 303–306. "Net biomass carbon changes at country scale show close agreement with national inventories in Gabon and Equatorial Guinea, with differences of 1 and 0.2 Tg C yr⁻¹, and good consistency in Cameroon (within 5 Tg C yr⁻¹)." These are small countries with relatively small and stable forest sectors. The agreement in these cases does not guarantee accuracy in the DRC, which dominates the regional budget.

Response: We agree that the agreement observed in Gabon, Equatorial Guinea, and Cameroon does not necessarily hold in larger countries such as the Democratic Republic of the Congo (DRC), which dominates the regional carbon budget. We have updated the sentences as:

Net biomass carbon changes at country scale (**Fig. 5a**) show broadly consistent magnitudes with national inventories¹⁷ harmonized by Grassi *et al.*¹⁶ in Gabon and Equatorial Guinea which have relatively small and stable forest sectors, with differences of 1 and 0.2 Tg C yr⁻¹, and of 5 Tg C yr⁻¹ in Cameroon where there is more deforestation. Differences between our estimates and Grassi *et al.*

are larger in large countries such as DRC that dominate the regional carbon budget. (Lines 388–392)

Lines 306–307. "In the DRC, our estimate indicates a smaller net carbon loss than the inventory." The discrepancy is attributed to "an overestimation of the carbon sink in 'forest remaining forest' in the inventory," but the alternative explanation — that the authors' approach underestimates losses in the DRC — is not considered. This asymmetric interpretation represents a significant framing bias.

Response: We have revised the text to provide a more balanced interpretation of this discrepancy: In the DRC, our estimate indicates a smaller net carbon loss than the inventory (**Supplementary Table 4**). This discrepancy may reflect differences between the harmonized inventory framework¹⁶ and our satellite-based biomass change estimates, particularly in the partitioning of gross carbon losses and gains, as well as uncertainties in our estimates, including possible underestimation of biomass losses in low-canopy non-dense forests. (Lines 392–397)

Lines 325–330. The claim that national inventories overestimate gross carbon gains by >50% in Gabon, Cameroon, and the DRC due to IPCC Tier 1 default growth rates is potentially important. However, the near- zero growth rates found by the authors for undisturbed forests could also reflect regression-to-the-mean bias in the deep learning model, which would systematically attenuate height gains in the tallest forests. This alternative explanation must be addressed.

Response: We agree that the near-zero growth rates estimated for undisturbed forests may partly reflect methodological limitations, including potential regression-to-the-mean effects in the canopy height model and saturation of satellite data in very high biomass forests that may attenuate height gains. We have revised the text as:

To evaluate this assumption, we calculated the average forest carbon growth rate from our maps and found that our growth rates over intact (undisturbed) forests were close to zero or even negative (**Supplementary Table 5**), lower than the refined IPCC default values⁵⁸. This result suggests that the gross carbon gains reported in national inventories may be overestimated by more than 50% in Gabon, Cameroon, and the DRC (**Fig 5b**). At face value, this difference may also reflect methodological limitations in our satellite-based height change estimates, including potential regression-to-the-mean effects in deep-learning models and saturation of satellite data in very high biomass forests that may attenuate height gains in old-growth forests. (Lines 417–425)

Lines 331–335. "Conversely, national inventories tend to underestimate carbon losses." This is plausible, but again, the comparison is between fundamentally different quantities: inventories report "deforested land," while this study reports all canopy-height losses, including natural dynamics and degradation.

Response: We have revised the text to clarify this methodological difference and to frame the comparison more cautiously:

National inventories also report lower gross carbon losses than our estimates (**Fig. 5c**). This difference partly reflects methodological differences between the two approaches. National inventories primarily report emissions associated with deforestation and do not explicitly report

human-induced degradation losses, whereas biomass losses in our study include both deforestation and forest degradation, as well as natural disturbance processes. By explicitly separating carbon gains and losses, our approach provides a complementary perspective on the dynamic of all disturbances and associated biomass carbon changes. (Lines 427–433)

RESULTS — Uncertainty and implications (Lines 359–410)

Lines 359–361. "...can support national inventories by linking carbon changes to specific local processes." This overstates operational readiness. The authors have not demonstrated that their uncertainty budget is fit for direct uptake into IPCC Tier 2/3 inventory frameworks; a more defensible claim would be that the maps could inform or complement inventories as an exploratory dataset.

Response: We have revised this sentence to avoid implying direct operational use in national inventory frameworks as:

Our high-resolution maps of forest carbon change provide new insights into fine-scale carbon dynamics in Central African forests and can help inform or complement national inventories by linking carbon changes to local processes such as small-scale clearing and regrowth. (Lines 455–457)

Lines 363–366. "Our gross carbon losses...are broadly consistent with...L-band vegetation optical depth data for 2010–2019." The comparison spans a decade-long temporal mismatch (2010–2019 vs. 2019–2022) in a region of rapid land-use change, which is not discussed; additionally, the phrase "strongly does not allow" is grammatically incorrect and must be revised.

Response: We revised this sentence to clarify the temporal mismatch and correct the wording.

Our estimates of gross carbon losses (2019–2022) from degradation and deforestation are broadly comparable to previous coarse-resolution estimates based on L-band vegetation optical depth data for 2010–2019⁵⁹ (**Supplementary Table 7**), although the comparison should be interpreted with caution due to differences in the analysed time periods and the coarse spatial resolution (25 km) of the L-band data, which does not allow detection of individual disturbances. (Lines 457–462)

Lines 366–368. "Our patch-scale AGC recovery rates...are consistent with field inventory in African rainforests (ref. 56)." Reference 56 (Csillik et al. 2024) is an Amazon study, not an African rainforest study; attributing it as African field evidence is a misattribution that must be corrected, and Cook-Patton et al. (2020), the other cited comparator, is a global modelled estimate rather than a Central African field measurement.

Response: We have corrected the discussion as:

Our patch-scale AGC recovery rates from regrowing forests are broadly consistent with recent remote-sensing estimates of Amazonian forest growth⁶⁰ and with global modeled estimates of tropical forest recovery, including those reported for African forests⁶¹, but exceed the refined IPCC guidelines 2006 default growth rates⁵⁸. (Lines 462–466)

Lines 368–371. "We also find high relative recovery rates in the first three years of regrowth...surpassing previous 20-year recovery estimates." Comparing 3-year rates against 20-

year mean rates is methodologically unsound: biomass recovery is fastest immediately after disturbance, so exceeding a multi-decade average is trivially expected and does not constitute a novel finding without a like-for-like temporal comparison.

Response: To avoid this methodological inconsistency, we have revised the text to remove the comparison with 20-year recovery estimates and instead emphasize the rapid early-stage recovery captured by our high-resolution observations of young regrowth forests.

We also find high relative recovery rates in the first three years of regrowth after degradation or temporary deforestation followed by regrowth. This highlights the important contribution of young regrowth forests in Central Africa (1–8 years old; Fig. 3), which exhibit faster biomass accumulation during early succession than older secondary forests⁶¹. (Lines 466–469)

Lines 371–373. "Some of these young forests may include agroforestry systems...which are difficult to distinguish from natural regrowth." This caveat is introduced as a brief remark, but its implications could be significant: the main claim that sub-hectare regrowth accounts for 48% of total biomass gains depends on the assumption of genuine forest recovery, and the authors should include at least an order-of-magnitude sensitivity estimate of how agroforestry interference might alter this figure.

Response: To evaluate the potential influence of agroforestry systems on our regrowth estimates, we conducted an additional analysis using the Spatial Database of Planted Trees (SDPT) dataset (V2.1)⁶², which maps tree crops (e.g., oil palm and rubber) and planted forests. We quantified their contribution to total aboveground carbon (AGC) gain within areas classified as regrowth. We find that tree crops and planted forests together account for only ~3% of total AGC gain within regrowth forests during 2019–2022. This indicates that the vast majority of biomass gains attributed to regrowth are not associated with agroforestry systems. Therefore, potential agroforestry interference is expected to have a limited effect on our estimates, and the conclusion that sub-hectare regrowth accounts for 48% of total biomass gains remains robust. We have clarified this point in the revised manuscript.

Some of these areas may include a small proportion of agroforestry systems (e.g., shaded cacao plantations) that are difficult to distinguish from natural regrowth. Based on the Spatial Database of Planted Trees⁶², these systems account for ~3% of our total AGC gain within regrowth forests and do not substantially affect regional biomass recovery estimates. (Lines 469–473)

Lines 375–377. "One study did not explicitly attribute the drivers of biomass change...due to the short time window covered by our maps." The phrase "One study" almost certainly refers to the present study, but the third-person wording suggests an external third party; this should be corrected to "This study" or "We were unable to," as the authors seem to be distancing themselves from their own limitation.

Response: We have revised the sentence accordingly by replacing "One study" with "We".

Lines 378–382. "Our approach...collectively enhances the robustness of change detection and reduces false positives." This claim is made without quantification: the authors do not specify the residual false-positive rate, the baseline they are improving upon, or the method of verification —

notably, their own visual validation shows 14% of loss patches were classified as "no" by expert interpreters, yet this figure is not mentioned here.

Response: We updated the independent validation (please refer to our response to Comment 1 above regarding the validation of small-scale disturbances) and revised the text to avoid overstatement and clarify the validation results.

Our approach integrates annual height composites, strict detection thresholds, and external forest loss datasets to improve the robustness of change detection. Visual validation indicates that most detected loss patches correspond to real canopy disturbances, although ~16% do not represent true loss signals based on expert interpretation (**Extended Data Table 1**). (Lines 480–484)

Lines 382–387. "A key next step will be to distinguish natural disturbances from anthropogenic drivers." This directly conflicts with Lines 404–405, where the authors already attribute most carbon loss to shifting cultivation and wood harvesting; if driver attribution is a future goal here, it cannot simultaneously be a settled conclusion at the end of the same section.

Response: We have removed the original statement about the driver attribution.

Lines 389–392. "Our height mapping primarily reflects top-of-canopy structure and may overlook understory losses." The acknowledgement is correct, but the discussion stops there; considering that selective fuelwood collection — one of the most widespread pressures in the Congo Basin — operates almost entirely beneath the canopy, the authors should estimate the directional bias and likely magnitude of this gap in their net flux figure.

Response: Thank you. We have revised the text as:

Our height mapping primarily reflects top-of-canopy structure and does not probe understory trees, which can be preferentially cut by fuelwood collection, selective logging of small trees, agroforestry practices, low-intensity fires, or die from drought⁶⁵. Because such disturbances often occur without substantially altering top canopy height, their contribution to biomass loss may be underestimated in our analysis. In particular, small-scale and sub-canopy disturbances that do not produce detectable height reductions may not be captured, contributing to uncertainty in biomass loss estimates. More broadly, uncertainties in canopy height change estimation itself, arising from detection limits, temporal sampling, and signal variability, further complicate the interpretation of both loss and gain. Future work should expand field inventories and leverage higher-resolution satellite data and deep learning-based models able to map full vertical waveforms from GEDI⁶⁶. (Lines 492–504)

Lines 393–395. "Upcoming missions like NISAR and BIOMASS will enhance monitoring of forest vertical structure and biomass dynamics." The reference is superficial: given that the allometric bridge ($R^2 = 0.40$, 32.6% relative error) is identified as a key limitation, the explanation should specifically address how BIOMASS's P-band sensitivity to woody biomass would address it, rather than simply claiming potential future improvement. authors should

Response: We have revised the text to clarify how upcoming missions, particularly BIOMASS, could address current limitations in estimating biomass from canopy height alone.

Recent missions such as BIOMASS, whose P-band radar is sensitive to large stems and will provide forest vertical tomography, together with NISAR, help reduce uncertainties associated with height-to-biomass conversion and improve monitoring of biomass dynamics. (Lines 504–507)

Lines 395–396. "Protecting the carbon sink function of African forests requires addressing both large-scale degradation and widespread small-scale disturbances." This conflicts with the paper's main finding: Central African forests were a net carbon source ($-58 \pm 9 \text{ Tg C yr}^{-1}$) during 2019–2022, not a sink; the wording must be changed to match the reported results.

Response: We have removed the sentence to avoid repeated discussion.

Lines 402–404. "This study provides the first biome-wide, high-resolution quantification of forest biomass dynamics." The unqualified "first" claim needs verification against Bourgoin et al. (2024) and Zhao et al. (2024), both cited elsewhere in the manuscript and both providing biome-wide Central African carbon assessments; the authors should specify the precise aspect — resolution, temporal coverage, carbon budget components — in which this claim is genuinely made.

Response: We have removed this summary paragraph to simplify the discussion.

Lines 404–405. "Most of the carbon loss in Central African forests is due to small disturbances, including shifting cultivation and wood harvesting." This mixes the dominance of small patches with the attribution of causes: the unattributed "Other forests" loss category ($-22.4 \pm 3.7 \text{ Tg C yr}^{-1}$) remains unexplained, and claiming shifting cultivation and wood harvesting as the main causes goes beyond what the data support, given the limitations of driver attribution acknowledged three paragraphs earlier.

Response: We have removed this statement.

Lines 405–410. "Reducing carbon losses will require...carbon crediting schemes...facilitating the regeneration of young forests...will be much more effective than creating massive new plantation projects." These policy recommendations are not based on any analysis in the paper: carbon crediting is suggested without considering additionality, permanence, or leakage; and the claim that regeneration is preferable to plantation rests on no comparative evidence presented in this study — all three should be removed, reframed as research questions, or moved to a dedicated policy commentary.

Response: We agree that these policy recommendations extend beyond the scope of the analyses presented in this study. To avoid overinterpretation, we have removed these statements from the manuscript.

METHODS — Annual Canopy Height and Biomass Mapping (Lines 413–434)

Lines 414–416. The deep learning model is described too briefly in the main text. Key missing information includes the number of Sentinel images per tile per year, the cloud masking procedure, the method for generating annual composites, the loss function, and the spatial sampling strategy. Since the study's conclusions rely entirely on the quality of these maps, a self-contained description in the main Methods is essential.

Response: We have revised the **Methods** section to provide additional details on the deep learning model used to generate the canopy height maps. Specifically, we now describe the preprocessing steps for Sentinel-1 and Sentinel-2 data, the number of images used per tile per year, the cloud masking and annual compositing procedures, the spatial sampling strategy for model training, and the loss function used to optimize the U-Net model.

Further implementation details are provided in [Supplementary Text 1](#). The full model description train/validation/test set up is described in ref ²⁶ (<https://doi.org/10.3389/frsen.2025.1724950>).

Lines 420–422. The allometric equation $AGB = 6.5981 \times (CH \times WD)^{1.0356}$ has several unaddressed issues: (a) the near-linear exponent is ecologically unexpected, as theory predicts non-linear height-biomass scaling; (b) the 1 km WD map resampled to 30 m means all pixels within neighbourhoods essentially share the same WD, so fine-scale AGB variation is driven by CH alone; (c) the R^2 of 0.40 indicates that 60% of biomass variance remains unexplained, and residual analysis stratified by forest type is absent.

Response: We thank the reviewer for these helpful comments regarding the canopy height–biomass relationship. We agree that biomass–height scaling relationships are often expected to be non-linear from an ecological perspective. When developing the biomass map (Wan et al., 2025) ¹², we tested several alternative formulations, including machine-learning approaches. Although some of these models achieved higher predictive accuracy on validation samples, they produced spatial artifacts and unstable spatial patterns when applied to wall-to-wall mapping at 30 m resolution. This is mainly due to the heterogeneous quality and limited temporal consistency of many potential predictor variables at this spatial resolution. To ensure spatial stability and temporal consistency across annual maps, we therefore adopted the simpler height-based formulation used here.

Second, regarding wood density (WD), we acknowledge that the available WD dataset has a coarse spatial resolution (~1 km). While resampling to 30 m does not introduce fine-scale variation in WD, incorporating WD helps mitigate biomass saturation effects when converting canopy height to biomass. In practice, most of the spatial and temporal variability in the AGB maps is driven by canopy height changes, while WD mainly acts as a regional modifier.

Third, the reported R^2 (0.40) reflects scale mismatch and representativeness of the available field inventory plot validation dataset and should not be interpreted as the variance explained across the entire mapped domain. Field plots are sparse across Central Africa and are disproportionately located in high-biomass forests, which affects the validation statistics. Because of the limited number of field plot samples, it was not possible to robustly evaluate the height–biomass relationship separately across forest types. Nevertheless, we have collected 400 plots data with all site PI being authors of our previous study, which is the largest dataset ever used for African biomass validation. We therefore assessed the uncertainties of the canopy height model and the biomass map separately. Larger uncertainties were found in montane forests, mainly due to the influence of terrain on GEDI lidar measurements.

We have included a detailed description of the allometric model and uncertainty calculation in the Methods section (Lines 535–576) and [Supplementary Text 1](#).

Lines 420–421. The AGB map "one" produced by the lead author without WD, and published in 2025 in *Frontiers in Remote Sensing* (Ref. 36), is, in my view, much more compelling than the current AGB map, which does not use coarse WD data at 1 km.

Response: We would like to clarify that the biomass map presented in our previous study (*Frontiers in Remote Sensing*, 2025) did incorporate wood density (WD) in the final model. In that study, we systematically evaluated several alternative approaches for estimating aboveground biomass from canopy height, including allometric models with and without WD, as well as machine-learning-based models. Although some of the more complex models achieved higher predictive accuracy on validation samples, they produced spatial artifacts and less stable spatial patterns when applied to wall-to-wall mapping at 30 m resolution. Based on these comparisons, we ultimately selected an allometric model that includes both canopy height and WD, as it provided the most spatially consistent and interpretable biomass maps. The inclusion of WD also helped reduce biomass saturation at higher canopy heights.

In the present study, we apply the same biomass estimation model to derive annual AGB maps from the canopy height time series. Although the available WD dataset has a relatively coarse spatial resolution (~1 km), it primarily acts as a regional modifier, while the spatial and temporal variability in the resulting AGB maps is driven by canopy height changes.

To avoid confusion, we have clarified the description of our biomass estimation model in the *Methods* section. (Lines 535–556)

Lines 425–428. The carbon conversion factor ($AGC = 45.6\%$ of AGB) and the forest definition ($CH \geq 5$ m) are both reasonable choices, but should be justified. The 5 m threshold may include non-forest land covers (shrublands, agroforestry) in a region where shaded cacao and other tree crops are widespread.

Response: We have clarified the justification for both the forest definition ($CH \geq 5$ m) and the carbon conversion factor used to derive AGC from AGB in the manuscript. Specifically, we now note that the $CH \geq 5$ m threshold follows the definition used in the GLAD tree height product and is broadly consistent with FAO forest definitions, while acknowledging that it may include some tree-based agricultural systems (e.g., shaded cacao agroforestry), although their influence on regional biomass estimates is expected to be limited in the dense forest regions analyzed here. We also clarified that the conversion factor of 45.6% lies within the commonly reported range of carbon fractions for tropical woody biomass.

Lines 429–430. Gap-filling of the WD map by "assigning the nearest available WD" is ad hoc. No assessment is provided of the fraction of pixels requiring gap-filling or of the spatial distribution of gaps.

Response: We have now added a description of the fraction of pixels requiring gap filling in the *Methods* section as: **The WD map was resampled from 1 km to 30 m for AGB estimation. Although this resampling does not introduce fine-scale variation in WD, including WD helps mitigate biomass saturation effects when converting CH to AGB. In practice, most spatial and temporal variability in AGB is driven by CH, while WD acts primarily as a regional modifier. Because the WD dataset did not fully cover all forest pixels in either 2019 or 2022, missing values were filled**

using the nearest available WD value. These filled pixels occur primarily in disturbed areas and low-density forest regions and account for only 1.5% of all pixels. (Lines 544–550)

Lines 432–434. Applying the dense-forest AGB model to non-dense forests and claiming "overall good validation performance" based on comparison with ESA CCI AGB (another satellite product) is not field validation. This cross-comparison may simply demonstrate that two satellite models agree.

Response: We agree that the comparison with ESA-CCI AGB should not be interpreted as field validation. In the revised manuscript, we have clarified that the comparison with ESA CCI AGB was therefore intended only as a consistency check with an independent satellite-based biomass product rather than as a formal validation.

The AGB model selected in this study was specifically calibrated for African dense forests. To enable country-level carbon accounting, we applied the same CH–AGB relationships to non-dense forest systems. This represents an important assumption, as structural differences in canopy architecture, species composition, and disturbance regimes may affect the transferability of these relationships (**Supplementary Text 1**). However, our previous ALS-based evaluation²⁶ included a wide range of forest structures, including relatively short forests in wet regions that are not classified as dense forests, and showed generally consistent performance across these conditions. Due to the limited availability of field inventory data for non-dense forests, a direct recalibration of the model was not feasible. Instead, we evaluated the robustness of the resulting AGB estimates through comparison with the ESA-CCI AGB product as an independent satellite-based dataset. The two datasets showed broadly consistent spatial patterns (**Supplementary Text 12, Supplementary Figs. 34–36**), supporting the large-scale applicability of the approach while acknowledging potential uncertainties at local scales. (Lines 564–576)

METHODS — Reference Data for CH Change Validation (Lines 435– 487)

Lines 437–449. The only ALS validation site (Yangambi, DRC) has repeat scans from 2015 and 2023, not 2019 and 2022. The assumption of linear interpolation for gains and the exclusion of pixels with loss between 2015 and 2019 are both problematic. The spatial representativeness of a single site for a 210 Mha study region is a major limitation.

Response: Due to the lack of repeated airborne lidar observations in Central Africa, the Yangambi Biosphere Reserve currently provides the only site in the region with multi-year ALS measurements suitable for evaluating canopy height change. Although the ALS observations (2015–2023) do not perfectly match the satellite observation period (2019–2022), we applied a processing approach to derive comparable CH changes from the ALS data.

We also acknowledge that this comparison is subject to limitations, including temporal mismatches, the need for interpolation and sample filtering, and limited spatial representativeness. To assess the potential effects of temporal interpolation and sample filtering, we conducted sensitivity analyses (**Supplementary Text 5**). (Lines 595–599)

We have added the discussion in [Supplementary Text 5](#):

The effects of temporal interpolation and sample filtering on the comparison between satellite-derived CH change (2019–2022) and ALS observations (2015–2023) were evaluated at both 10 m and 30 m resolutions (**Supplementary Figs. 12, 13**). We first evaluated the consistency between satellite-derived and ALS-derived CH maps at individual years (**Supplementary Fig. 2**). Our CH map in 2019 showed good agreement with ALS CH observations acquired in 2015 ($R^2 = 0.76$, MAE = 4.35 m, Bias = -1.08 m), with comparable performance between our CH map and ALS observations acquired in the same year 2023, although the correlation was lower ($R^2 = 0.55$, MAE = 4.42 m, Bias = 0.29 m). The shift in bias between the two comparisons likely reflects temporal mismatches between ALS acquisition dates and satellite observations, as forest growth and disturbance may occur between acquisition periods.

Applying a linear interpolation to pixels with height gain (condition 1) resulted in only minor changes in agreement metrics, indicating that assumptions on growth dynamics have a limited influence on the comparison. In contrast, excluding pixels with ALS-detected height loss prior to the 2019–2022 study period (condition 2) led to a substantial improvement in agreement, increasing R^2 from 0.06 to 0.22 at 10 m resolution and to 0.52 at 30 m resolution, while also reducing error metrics at both spatial resolutions. This reflects the effect of temporal mismatch between the ALS acquisition (2015 and 2023) and our shorter study period, whereby disturbances occurring before our observation window introduce inconsistencies in the raw comparison. Removing these pixels reduces discrepancies associated with misaligned disturbance timing rather than altering the underlying distribution of canopy height change. Importantly, consistent patterns across 10 m and 30 m analyses indicate that the observed improvements are not driven by scale-dependent smoothing or selective bias, but by improved temporal alignment between datasets. These results suggest that differences in the comparison between ALS-observed CH change and our dataset are primarily attributable to timing and sampling inconsistencies, rather than systematic overestimation of height change in our satellite dataset.

Lines 441–443. Using the maximum ALS value within a 30×30 m area as a proxy for GEDI RH100 is imperfect — GEDI's 25 m footprint with specific waveform processing is not equivalent to a wall-to-wall ALS maximum.

Response: The original ALS CH maps, generated at 2 m resolution, were aggregated to coarser grids by extracting the maximum value within each cell, representing the upper canopy envelope. Although this approach does not fully reproduce the GEDI waveform-derived RH100 metric, it provides a practical approximation of top-of-canopy height in the absence of waveform information. (Lines 584–588)

Lines 450–469. The GEDI comparison is recognised as non-independent. The argument that "changes are smaller signals than the spatial gradients used for training" is weak — model biases (e.g., regression to the mean) would reduce change signals, possibly leading to artificially high large-area agreement. $R^2 = 0.22$ at 10 km is troubling and warrants further discussion.

Response: We have removed the statement that “changes are smaller signals than the spatial gradients used for training.” We have also clarified the role of the GEDI-based comparison. Because the canopy height model is trained on GEDI RH100 observations, this comparison is not

fully independent and is therefore intended primarily as a consistency check of large-scale RH100 change patterns rather than as an independent validation. (Lines 608–610)

We have additionally expanded the discussion of the relatively low R^2 value in the Results section. Comparisons with changes from GEDI discrete footprints samples (RH100) between the first and last years of our study period showed weak agreement at 30 m but improved when aggregated to 10 km ($R^2 = 0.23$, $MAE = 4.35$ m) (**Extended Data Fig. 2, Supplementary Fig. 14**). This reflects several factors, including scale mismatch, sparse GEDI footprint sampling with no revisit in time, and geolocation uncertainty⁴², as well as the difficulty of detecting small height changes and uncertainties in our biomass mapping at fine spatial scales. (Lines 145–150)

Lines 470–487. Comparisons with TMF, GFC, and RADD reflect context rather than validation. These are other satellite products with their own errors. Showing more detections than these products cannot distinguish genuine improvement from increased false positives without independent ground truth.

Response: We have revised the text to clarify that these comparisons are intended to examine the consistency of disturbance patterns as:

These datasets provide independent satellite-based references for disturbance patterns, although they are not used here as formal validation because each product is based on different methods and has its own detection uncertainties. (Lines 612–615)

METHODS — CH Change Processing (Lines 488–530)

Lines 493–503. The 3D connected-component analysis employs three time layers linked to the 2019 baseline, which means that correlated noise in 2019 could generate false spatiotemporal clusters that pass the filter. The temporal consistency check is lenient — requiring only one adjacent year with the same trend.

Response: The purpose of the 3D connectivity analysis is to remove isolated noise pixels by identifying spatiotemporally coherent clusters of canopy height change. We have clarified this filtering logic in the Methods section. (Lines 654–665)

To identify and filter out isolated noise pixels, we constructed a 3D data cube of CH changes voxels across three years and applied a 26-connectivity method⁷⁹ to detect spatiotemporally coherent clusters of change: $\Delta CH_{2021-2019}$, $\Delta CH_{2022-2019}$, and $\Delta CH_{2023-2019}$. Similar connectivity analysis methods have been widely used to identify forest disturbances and fragmentation patterns^{50,80}. We retained only the $\Delta CH_{2022-2019}$ layer, treating a pixel as reliable if it was part of a 3D-connected cluster; isolated voxels were considered noise and removed. We then applied a temporal consistency check: if the sign of change in $\Delta CH_{2022-2019}$ was opposite to both $\Delta CH_{2021-2019}$ and $\Delta CH_{2023-2019}$, the pixel was excluded. Although the three layers share the same 2019 baseline, they are influenced by noise in that baseline. The additional temporal consistency check can reduce such effects by removing pixels whose change direction in 2022 is inconsistent with both neighboring years, thereby filtering out isolated or inconsistent fluctuations.

Lines 504–508. The Mann-Kendall test with only 4–5 time points and a threshold of $S < 6$ would exclude genuine abrupt gain events (e.g., rapid post-disturbance pioneer flush). TMF-classified regrowth is still considered, resulting in a dual standard in which new detections face stricter criteria than TMF pixels.

Response: The Mann–Kendall filter was designed to identify pixels exhibiting consistent and sustained canopy height increase over the 2019–2022 period, thereby reducing false positives from short-term fluctuations. We acknowledge that, with a limited number of time points, this approach may be less sensitive to abrupt or recently initiated regrowth events, particularly when recovery begins late in the observation period. To mitigate this limitation, we incorporated regrowth information from the TMF dataset when it coincides with observed canopy height gain. This step is intended to complement, rather than replace, the Mann–Kendall filter by capturing regrowth signals that may not yet exhibit a strong temporal trend in the height time series. Importantly, both datasets are required to indicate gain, ensuring consistency in the direction of change and avoiding the inclusion of spurious signals. We therefore do not consider this a dual standard, but rather a complementary use of independent information to balance omission and commission errors in detecting regrowth.

Lines 509–513. The $|CV| \leq 5$ threshold on Sentinel-1 VV lacks solid physical justification. Supplementary Fig. 8 illustrates continuous variation with no clear inflexion at $|CV| = 5$. The threshold is arbitrary, and its impact on actual changes (noted: "some large ΔCH values may also be excluded") is not quantified.

Response: We acknowledge that the CV threshold does not represent a strict physical boundary, as the distribution of CV values is continuous.

Pixels showing very low variability in the Sentinel-1 VV signal ($|CV| \leq 5\%$) between 2019 and 2022 were excluded, as these are likely associated with false structural canopy changes. This threshold was selected empirically based on model uncertainty, as height loss within this range is generally small (<5 m) and subject to higher relative uncertainty. The threshold is therefore intended to reduce noise from low-magnitude and uncertain changes rather than to define a sharp physical transition. (Lines 675–680)

To evaluate the sensitivity of this choice, we conducted additional analyses using a range of CV thresholds (1–6%). We found that gross carbon loss estimates remain relatively stable ($\sim 10\%$) for thresholds between 3% and 6% when other parameters were fixed, whereas lower thresholds (1–2%) lead to larger deviations ($>20\%$) by allowing more low-magnitude and potentially spurious height changes to be retained. These results indicate that the selected threshold represents a balance between reducing noise and preserving meaningful canopy change signals. In addition, the CV threshold is not applied in isolation but is combined with temporal consistency checks (Mann–Kendall test) and disturbance information from the TMF dataset. We have clarified this multi-criteria filtering strategy and added the sensitivity analysis in the [Methods, Supplementary Text 4.5, Supplementary Tables 9, 10](#).

Lines 514–519. The $|\Delta CC| \leq 2\%$ threshold for undisturbed forests is borrowed from studies in different regions. In Central African closed-canopy forests (80–100% cover), selective logging could easily produce $<2\%$ cover change at 10 m, meaning this filter could systematically exclude low-intensity degradation.

Response: The $|\Delta CC| \leq 2\%$ threshold follows previous studies^{84,85} and is applied conservatively in this framework. We acknowledge that, in dense closed-canopy forests, some low-intensity disturbances (e.g., selective logging) produce only small changes in canopy cover and could therefore be partially excluded by this threshold. However, this criterion is not applied in isolation. Pixels are removed only when ΔCC and ΔCH show inconsistent change directions, meaning that disturbances associated with coherent canopy height loss are still retained even if canopy cover changes are small. In addition, disturbance information from the TMF dataset is used as supporting evidence to further reduce the risk of excluding real disturbances. Sensitivity analyses show that varying the ΔCC threshold between 1% and 3% results in only limited changes in AGC estimates, indicating that the results are not highly sensitive to this parameter. (Lines 681–688)

Lines 521–524. Maintaining TMF-classified disturbance pixels regardless of noise criteria ensures the final product is a superset of TMF. This structural assurance weakens the comparison in Figure 2 and should be clearly disclosed.

Response: TMF pixels are not automatically retained in the final CH change dataset. TMF disturbance information is used only in steps (2)–(4) to re-evaluate the pixels for which our CH change and TMF disturbance trends were consistent. TMF pixels classified as disturbed were not automatically retained; instead, they were kept only when their disturbance signal matched our CH change trends. Therefore, the final CH change dataset is not a superset of TMF detections. (Lines 689–693)

METHODS — CH Change Classification (Lines 531–565)

Lines 533–538. The height-based (structural) definitions of deforestation and degradation may not align with land-use definitions used in UNFCCC reporting or national inventories, leading to comparability issues in the Fig. 5 national budget comparison.

Response: Disturbance classes defined here are based on canopy height changes rather than land-use definitions used in national inventories. We have clarified this point in the **Methods** section as: **This classification is based on height changes derived from satellite observations rather than land-use definitions used in national inventories.** (Lines 706–707)

Lines 547–555. The deforestation threshold ($\Delta CH < -5$ m AND $CH_{2022} < 10$ m) is borrowed from Bornean forest studies (ref 78). Applicability to Central African forests is not demonstrated. Supplementary Fig. 31 shows many TMF-deforestation pixels with a 2022 CH of 15–45 m, indicating that the $CH < 10$ m criterion would exclude a substantial fraction of what TMF considers deforestation.

Response: We have included more explanations and comparisons in the revised manuscript as: **Following definitions used in ALS-based canopy gap studies⁸⁷, we defined deforestation as pixels with $\Delta CH < -5$ m and a final CH below 10 m in 2022, while remaining loss pixels were classified**

as degraded forests. Although these thresholds were initially informed from other tropical regions, our analysis combining TMF classifications and CH change shows that the separation between degradation and deforestation is more clear using these thresholds (**Supplementary Fig. 37**). Our ΔCH threshold is consistent with previous studies showing that severe canopy loss is associated with mean height reductions of 5 m or more^{88,89}. This CH-based definition does not fully overlap with TMF classifications. TMF deforestation represents persistent forest-cover disruption detected from long-term Landsat observations (>900 days)⁸, which serves as a proxy for deforestation processes, whereas our CH-based framework emphasizes the structural change. Consequently, some TMF-deforestation pixels retained relatively high post-disturbance CH or experienced rapid recovery, making them structurally more consistent with partial canopy disturbance than complete canopy removal in our framework. These pixels were therefore classified as degraded forests in the CH-based analysis (**Supplementary Table 12**). (Lines 722–736)

In our comparison, nearly all TMF degradation pixels were classified as CH degradation, whereas only ~24% of TMF-deforestation pixels satisfied our structural deforestation criterion. Many remaining TMF-deforestation pixels retained relatively high post-disturbance canopy height (e.g., 15–45 m). Although uncertainties in satellite-derived CH estimates may affect low-height or subtle disturbances, they are unlikely to account for the consistently high residual CH observed in these pixels. Such pixels are therefore more consistent with partial canopy disturbance or structurally degraded forests than with complete canopy removal, and were consequently classified as CH degradation in our framework. (**Supplementary Table 12**)

To further evaluate the applicability of this threshold in Central African forests, we conducted sensitivity analyses by varying the CH_{2022} threshold from 10 to 20 m (**Supplementary Table 2**). Across this range, the dominant deforestation patch size contributing to total biomass loss remained highly stable (~1.5 ha), indicating that our main conclusions are robust to threshold choice. Taken together, our CH-based framework provides a physically interpretable characterization of structural canopy loss and distinguishes severe canopy removal from partial canopy disturbance, complementing spectral disturbance products such as TMF.

Supplementary Table 12: Correspondence between TMF disturbance classes and canopy-height-based disturbance classes. Percentages represent the proportion of TMF disturbance pixels overlapping with canopy-height-based (CH-based) disturbance classes derived from canopy height change and post-disturbance canopy height.

TMF disturbance class	Percentage of TMF disturbed pixels overlapped with CH deforestation	Percentage of TMF disturbed pixels overlapped with CH degradation	Outside CH degradation and deforestation
TMF degradation	6.3%	93.7%	0.0%
TMF deforestation	24.4%	75.5%	0.1%

Lines 556–565. The XGBoost classifier for "uncertain" pixels (12.7% of gains) is trained on TMF labels and evaluated against TMF labels. It can only reproduce TMF's logic, not improve upon it. Features include the same CH and Δ CH data being classified, creating circularity. The reported accuracy (0.87) measures TMF consistency, not absolute accuracy.

Response: We acknowledge that the XGBoost classifier is trained using TMF-derived labels and therefore does not provide an independent classification. Its purpose is to assign “TMF uncertain pixels” to regrowth or degradation categories based on canopy height features, thereby extending the TMF disturbance classification to pixels lacking a definitive label. Because the model is trained on TMF labels, its performance primarily reflects consistency with the TMF labeling scheme rather than an absolute classification accuracy. We incorporate this consistency into the class-level uncertainty as an additional source of disturbance detection error in the subsequent analyses. (Lines 743–750)

METHODS — Disturbance patterns (Lines 566–588)

Lines 566–569. "We applied an 8-connected component labelling algorithm to identify contiguous clusters of disturbed 30 m pixels averaged from 10-m CH changes." The patch delineation unit (30 m) is inconsistent with the detection resolution (10 m): averaging 10-m CH changes to 30 m before applying the labelling algorithm can either merge adjacent genuine sub-hectare disturbances into a single larger patch or suppress isolated 10-m signals that do not survive the averaging step. This choice directly conditions the headline claim that 87% of patches are sub-hectare, yet no justification is given for why 30 m — rather than the native 10 m — was chosen as the patch unit, and no sensitivity analysis tests how the statistic changes under alternative choices.

Response: We clarify that patch delineation is based on canopy height (CH) change derived from annual CH maps at 30 m resolution, rather than from the 10 m CH change maps aggregated to 30 m. Annual CH maps were first generated at 10 m resolution and subsequently aggregated to 30 m, from which both CH change and AGC change were consistently derived. This ensures full spatial consistency between disturbance delineation and carbon estimates.

To assess sensitivity to spatial resolution, we compared this approach with patch delineation based on native 10 m resolution CH change maps (**Supplementary Fig. 21**). In addition, we evaluated an alternative approach in which 10 m CH change maps were first aggregated to 30 m prior to patch identification. This approach produced CH change distributions consistent with those from our primary 30 m-based method, but missed some small disturbance patches that are not represented in the 30 m AGC change maps. In contrast, using native 10 m CH change maps results in a substantially larger number of small disturbance patches due to increased sensitivity to fine-scale variability, with the number of small patches differing by orders of magnitude compared to the 30 m-based analysis. (Lines 760–768)

Lines 566–569. "Each connected region was labelled as a patch...patch sizes were calculated based on the number of connected pixels at 30 m resolution (0.09 ha per pixel)." The 8-connected component algorithm requires a prior binary classification of each pixel as disturbed or not disturbed, but the threshold defining this binarisation is not stated here. Given that the noise-filtering pipeline involves multiple tunable thresholds (Mann–Kendall trend test, Sentinel-1 CV,

canopy cover consistency), small changes in the final binarisation step can substantially fragment or merge patches, altering both patch count and size distribution. The sensitivity of all patch statistics to this threshold should be reported.

Response: Patch delineation is not based on a single threshold defining disturbed versus non-disturbed pixels. Disturbed pixels were defined as valid CH changes following the multi-step filtering procedure described above. Each connected cluster was labeled as a patch, and patch size was calculated based on the number of connected pixels (0.09 ha per pixel at 30 m resolution). The resulting classification defines pixels as gain, loss, or disturbance types, and patch delineation is then applied to these classified pixels. (Lines 756–760)

We evaluated the sensitivity of our results to the CH_{2022} threshold. We conducted a set of sensitivity analyses by varying the threshold from 10 to 20 m (**Supplementary Table 2**). These analyses show that the characteristic patch size of deforestation remains highly stable (~1.5 ha) across thresholds. (Lines 255–258)

We also have evaluated the sensitivity of detected areas of different disturbance types to different filtering configurations and threshold choices within the pipeline. Across all tested configurations, variations in these parameters primarily affected estimates of net AGC change (up to 12.43%) and the corresponding area (up to 10.00%) as well as gross AGC gain (6.65%) and the corresponding area (15.58%). At the disturbance-class level, the largest sensitivity was observed for the “other loss” category, with AGC estimates varying by up to 21.80% and area by up to 31.51%, while other disturbance classes showed comparatively smaller variations. These results indicate that parameter choices mainly influence less constrained disturbance types, whereas overall AGC budgets remain relatively stable. The resulting variability was incorporated into the structural uncertainty of AGC and area change estimates. (Supplementary Text 4.5)

Lines 570–572. "We computed patch sizes for different bins of CH loss and gain...each type exhibiting distinct patch size distributions." The dominance of sub-hectare patches in the count distribution partly results from a mathematical property of rasterised landscape data: geometric probability ensures that small patches outnumber large ones in any spatially dispersed detection output, regardless of ecological reality. The authors do not provide a null expectation — for example, the patch-size distribution expected under random spatial noise at this resolution — to compare with their observed distribution. Without this, the 87% figure cannot be distinguished from a artefact of rasterisation.

Response: We acknowledge that smaller patches are expected to be more numerous in raster-based analyses due to geometric effects, and therefore patch counts alone may overrepresent small disturbances.

To ensure that our results are not an artifact of rasterization or noise, we conducted an independent validation using high-resolution (5 m) PlanetScope imagery ([please refer to our response to Comment 1 above regarding the validation of small-scale disturbances](#)). This analysis shows that detected disturbance patches, including small ones, correspond well to real canopy changes, indicating that they are not random noise.

Importantly, our conclusions are not based solely on patch counts. We additionally quantified the contribution of different patch-size classes to total AGC change. After noise filtering, patches smaller than 1 ha account for 48% of biomass carbon gains in regrowing forests and 37% of carbon losses in deforestation patches. This demonstrates that their importance is not merely a consequence of numerical dominance, but reflects ecosystem-scale carbon dynamics. Furthermore, independent datasets such as TMF show similar patterns, with small disturbance patches (<1 ha) dominating across the study region (87% of patch count and 16% of disturbance area), providing additional support that this pattern is not specific to our methodology (87% of patch count and 21% of loss area).

Lines 573–574. "We used disturbance patch size and types defined by CH change and then calculated AGC change for each class." Patch types (degradation, deforestation, regrowth) are classified using CH change thresholds, and AGC change is then attributed to these same CH-derived classes. This means the classification uncertainty identified elsewhere — notably the 14% error rate for deforestation — propagates directly into the AGC attribution by class. The current section does not acknowledge this linkage, creating a disconnect between the disturbance-typing method and the uncertainty framework that follows.

Response: We have restructured the uncertainty estimation framework. Classification uncertainty is propagated into AGC change estimates by disturbance type and is now explicitly accounted for in the uncertainty framework. We have revised this sentence as:

Finally, we quantified both patch counts and their contributions to total AGC change across different bins of CH loss and gain, as well as across disturbance types defined above. Uncertainty associated with disturbance classification was explicitly propagated into the AGC change estimates for each class. (Lines 768–771)

Lines 574–576. "Road mapping (Extended Data Fig. 5) was derived from 5-meter Planet imagery, and the road fraction was calculated as the ratio of road pixels to total pixels within 30-meter grids." Road detection from optical imagery in dense tropical forests often misses roads due to canopy closure over unpaved tracks — the most ecologically significant road type in the Congo Basin. No validation of the completeness of road mapping is reported, yet road proximity is used as a proxy for anthropogenic attribution. If logging tracks and minor roads are systematically overlooked, the road-proximity analysis will underestimate the human impact of small disturbances, reducing the reliability of any driver inferences made from it.

Response: We have clarified the description as:

Because optical imagery may miss small logging trails or canopy-covered roads, this dataset primarily captures visible roads and should be interpreted as a proxy for accessibility rather than a complete representation of all forest access routes. We note that our map of roads based on 5 m Planet data is the best available to date for the whole region and it is consistent across all the tropics. (Lines 776–779)

Lines 577–580. "To quantify how patch size influences legacy (future) AGC change following forest regrowth and degradation events that occurred prior to our observation period, we used past disturbance information from the TMF dataset collocated with AGC change observed in the current

period." This design considers two independent variables: patch size and time since disturbance. Patches disturbed in 2018 will show a quicker apparent recovery from 2019 to 2022 than patches disturbed in 2002, regardless of their size, because biomass increases linearly and is fastest immediately after disturbance. Without explicitly controlling for time since disturbance, the analysis cannot distinguish patch size effects from temporal recovery effects, and any observed patch size recovery in AGC may be confounded.

Response: In our analysis, disturbance age is explicitly considered in our analysis through the TMF dataset. We examined AGC change as a function of patch size and disturbance age (Fig. 2), which allows comparison of patch-size effects within the same disturbance-age groups. Therefore, the observed relationship between patch size and AGC change is not solely driven by temporal recovery differences, but reflects size-dependent patterns under comparable disturbance ages. (Lines 781–786)

Lines 580–583. "Our results focus on the 20-year period before our observations (2019–2022), as the availability of reliable Landsat data is limited before ~2000." The 20-year window is shown as a deliberate analytical choice, but it is mainly due to data availability limits. Disturbances before 2000 might still be affecting AGC dynamics from 2019 to 2022, especially for species that recover slowly and larger clearings. Cutting off the window at 2000 may consistently underestimate the impact of earlier and larger patches, which could bias the observed link between patch size and recent AGC change in favour of smaller, more recent disturbances.

Response: We agree that disturbances occurring before 2000 may still influence forest biomass dynamics during our observation period. The 20-year window used in this analysis is primarily constrained by the availability of reliable Landsat observations in Central Africa. We have clarified in the Methods section that this choice is mainly driven by data availability rather than an ecological assumption about recovery times. Disturbances occurring earlier than 2000 may still contribute to current biomass dynamics but cannot be consistently detected with available satellite records. To provide additional context, we also report results using the full TMF disturbance record (1990–2022) in [Supplementary Fig. 38](#). (Lines 786–792)

Lines 585–587. "As the annual change collection cannot separate afforestation from secondary forests, we also compared AGC change after excluding afforestation pixels from the Transition Map (Supplementary Fig. 33)." The recognised difficulty in distinguishing afforestation from secondary forest regrowth is a significant limitation for the main conclusion that sub-hectare regrowth patches account for 48% of gross carbon gains. If a meaningful portion of these patches are plantation establishment rather than natural recovery, the ecological interpretation — and the policy suggestion favouring regeneration over plantations — is compromised by the very data used to support it. This sensitivity test should be quantified and discussed in the main text alongside the 48% figure it directly pertains to, rather than being limited to the supplementary materials.

Response: Afforestation may contribute to part of the detected regrowth signal; however, direct comparison of results with and without afforestation pixels shows highly consistent patterns in patch size and AGC change (Fig. 3, Supplementary Fig. 39), indicating that its influence is minimal (Lines 792–796). In addition, independent data from the Spatial Database of Planted Trees indicate

that plantation systems account for only ~3% of total AGC gain within regrowth forests (Lines 471–473). Together, these results suggest that afforestation has a limited impact on our estimates and does not affect the conclusion that small regrowth patches contribute substantially to biomass carbon gains.

METHODS — Uncertainty in AGC Change Estimates (Lines 589–627)

Lines 592–594. The assumptions of independence and Gaussian distribution for inter-annual errors are almost certainly violated. Shared model architecture, overlapping training data, identical allometry, and similar atmospheric conditions lead to correlated errors. Gaussian assumptions are unlikely to hold for satellite-derived height errors, which often display heavy tails.

Response: We agree that the assumptions of independence and Gaussian-distributed inter-annual errors are unlikely to be strictly satisfied in satellite-derived biomass change estimates. As noted by the reviewer, shared model architecture, overlapping training data, and consistent allometric formulations can introduce correlated error structures, while remote sensing retrieval errors often exhibit non-Gaussian behavior.

In the updated manuscript, we have revised the uncertainty framework to better reflect the underlying sources of uncertainty and to reduce implicit dependencies. [Please refer to our responses to the related comments in the General Overview section and Comment 2 above regarding uncertainty estimation.](#)

Specifically, we now decompose total uncertainty into three components: (1) biomass uncertainty arising from allometric modelling and measurement errors (σ_{bio}); (2) classification uncertainty associated with disturbance attribution and detection accuracy (σ_{cls}); and (3) uncertainty related to the choice of a CH change processing scheme (σ_{scheme}), including trend detection thresholds, Sentinel-1 radar-based noise filtering, and canopy cover consistency constraints. Each component was quantified separately using dedicated simulation and propagation approaches tailored to its error characteristics, rather than assuming a common Gaussian error structure. Full methodological details are provided in **Supplementary Text 12 and Supplementary Fig. 21**. While we combine these components assuming independence to provide a tractable estimate of total uncertainty, we acknowledge that some degree of correlation may persist, particularly between σ_{bio} and σ_{scheme} due to shared inputs and processing steps. As such, our estimates should be interpreted as a first-order approximation of uncertainty. Importantly, sensitivity analyses across alternative processing schemes (σ_{scheme}) show that our main conclusions are consistent across different processing parameter choices, suggesting that residual dependencies or deviations from Gaussian assumptions do not materially affect the key results. We have clarified these points in the revised manuscript (Lines 797–811, Supplementary Text 10, Supplementary Text 4.5).

Lines 596–600. Using the allometric RMSE (98.89 Mg ha⁻¹) as a spatially constant prediction uncertainty fails to account for variation in the height-biomass relationship across different forest types, species compositions, and disturbance histories. The XGBoost-predicted residual map partly addresses this issue, but was trained on only about 400 plots.

Response: We agree that using a spatially constant RMSE does not capture variability in the height–biomass relationship across forest types, and that a spatially explicit residual model based on a limited number of field plots would be difficult to justify at regional scale.

In response, we have revised the biomass uncertainty framework. [Please refer to our responses to the related comments in the General Overview section and Comment 2 above regarding uncertainty estimation.](#)

We no longer use a constant RMSE or an XGBoost-derived residual surface. Instead, biomass uncertainty is modeled as a heteroscedastic error proportional to AGC, based on independent validation of the biomass model. This allows uncertainty to vary spatially with biomass magnitude, rather than assuming a uniform error across all forests. We further account for intertemporal dependence in biomass estimation errors using a mixed error structure when propagating uncertainty in AGC change, instead of assuming independent errors between years.

The full description of the revised framework, including uncertainty propagation and aggregation, has been added to the [Supplementary Text 10](#).

Lines 600–601. The 20% calibration uncertainty from Chave et al. (2001) — a single 10 ha plot in French Guiana — is not relevant to Central Africa or to satellite-derived biomass. This is a placeholder, not a calibrated estimate.

Response: We thank the reviewer for this comment and agree that the 20% calibration uncertainty reported by Chave et al. (2001), based on a single plot, is not appropriate for representing biomass uncertainty in this study.

We have removed this term entirely and updated the biomass-related uncertainty framework. Instead of applying an external plot-based calibration uncertainty, we model biomass uncertainty as a heteroscedastic error proportional to AGC ($\sigma_{AGC} = \alpha \times AGC$). This formulation reflects the aggregate prediction uncertainty of the biomass estimation framework. Because this uncertainty term already incorporates the combined effects of allometric relationships and field measurement errors through model validation, we do not introduce an additional plot-level uncertainty term to avoid double counting. We nevertheless tested the inclusion of an additional plot-level uncertainty (8.3%) reported in previous studies (Ploton et al, 2020, Scientific Data) and found that it had only a minor effect on aggregated AGC estimates.

The revised formulation is described in detail in [Supplementary Text 10](#).

Lines 614–623. The classification uncertainty is evaluated in a circular manner (XGBoost is trained and tested on the same labels). Averaging 1% (degradation) and 14% (deforestation) to obtain 7.5% for both classes is statistically unjustifiable and hides actual differences in classification difficulty. These are separate error modes; averaging them is not equivalent to error propagation.

Response: Thank you for this important comment. We agree that the previous approach, using XGBoost-derived accuracy evaluated on the same labels and averaging error rates across disturbance types, was not appropriate, as it introduced circularity and obscured class-specific differences in classification uncertainty.

In the revised manuscript, we have removed this approach entirely and redefined classification uncertainty using a component-based framework. Specifically, classification uncertainty is now

quantified separately for each disturbance class by combining (i) detection uncertainty derived from independent validation data (high-resolution Planet images), (ii) threshold sensitivity in distinguishing deforestation and degradation losses, and (iii) class extension uncertainty for gain-related regrowth and degradation. These components are treated explicitly and are not averaged across classes. Uncertainty is propagated through Monte Carlo simulations and class-specific coefficients of variation, allowing differences in classification difficulty and ambiguity to be preserved. This provides a more realistic representation of classification uncertainty compared to the original framework.

The full formulation and propagation of classification uncertainty are described in [Supplementary Text 10](#).

Lines 624–625. The acknowledgement that $\sigma_{\text{class_cv}}$ (propagated from pixel-level uncertainty) is "considerably smaller" than $\sigma_{\text{class_cc}}$ (classification uncertainty), with the latter dominating, means the total uncertainty is primarily determined by the most poorly quantified component (the circular classification assessment).

Response: As described in our responses above, we have fundamentally revised the classification uncertainty framework by removing the circular XGBoost-based assessment and replacing it with a component-based approach using independent validation, threshold sensitivity, and class-specific uncertainty propagation. As a result, classification uncertainty is no longer dominated by a single poorly constrained component. [We refer the reviewer to our detailed responses above and the revised methods \(Supplementary Text 10\) for the full description of the updated uncertainty framework.](#)

FIGURES

Figure 1, Panel b. Averaging change and non-change pixels within 5 km cells masks the true distribution of disturbances. A cell with 95% undisturbed forest and 5% severe disturbance shows a near-zero mean ΔCH . The colour scale (–2 to +2 m) compresses the signal. Displaying the fraction of changed pixels per cell, or the total carbon change per cell, would be more informative.

Figure 1, Panels c–e. Three "illustrative" regions are presented without outlining the selection criteria. The reader cannot determine whether they are representative or cherry-picked. The TMF overlay combines temporal classes (recent degradation and older degradation showing recovery), which complicates interpretation. The spatial pattern of noise masking is not discussed.

Response: For Fig. 1b, we agree that averaging canopy height change within grid cells could mask disturbance signals. We have therefore revised this panel to show the fraction of pixels exhibiting change within each 5 km cell, which better represents the spatial distribution of disturbances. For Fig. 1c–e, we have clarified the selection criteria in the caption. The three regions are located in areas of high disturbance density and were selected to represent typical disturbance patterns dominated by shifting and permanent agriculture in Central African forests.

We also revised the TMF overlay to display the year of first degradation instead of disturbance categories, which allows clearer differentiation between recent and older disturbance events. Noise pixels were removed through a multi-step filtering procedure designed to exclude no-change

pixels and spurious change signals. As this filtering is based on signal consistency rather than spatial location, it is not expected to introduce systematic spatial bias in the mapped patterns.

Figure 2, Panel a. The log-scale y-axis visually compresses differences. At the critical -2 to -5 m range (where noise filtering removed 67–86% of pixels), the discrepancy between this study and TMF is approximately one order of magnitude. Whether these retained signals are real or residual noise remains unresolved. A false-positive analysis is necessary.

Figure 2, Panel b. The result (more small patches detected) is fundamentally influenced by higher spatial resolution. Comparing total area or total carbon flux per would provide more insight than simply counting patches.

Response: We have added additional analyses to address both concerns.

First, regarding the potential influence of noise in the detected height loss signals (Panel a), we performed a validation analysis based on independent sampling across loss, gain, and stable forest conditions. Rather than focusing on a narrow height range, we evaluated the reliability of detected disturbances at the patch level, particularly for small patches (<1 ha), where noise effects are expected to be strongest. For these small patches, the mean height change is close to 5 m. The validation results show an accuracy of 79% for loss and 69% for gain, indicating that the majority of detected signals correspond to real forest change rather than residual noise.

Second, we agree that higher spatial resolution increases the detection of small disturbance patches (Panel b). To address this, we extended the analysis beyond patch counts by quantifying the total affected area as a function of patch size, which is now shown as an inset in Fig. 2b. While small patches dominate in number, they also contribute substantially to the total disturbed area.

Figure 3, Panels a, c, e. Very large standard deviations (error bars spanning from positive to negative) indicate considerable within-class variability. For degradation at small patch sizes (panel c), the mean AGC change is close to zero or slightly positive, suggesting the class is heterogeneous (with a mixture of recovery and actual losses).

Figure 3, Panels b, d. Both figures display AGC density change plotted against degradation years (b) and regrowth years after deforestation (d), with each trend line annotated with its slope and a p-value reported as $P = 0.00$. With millions of observations, statistical significance is guaranteed, making it meaningless. The scientifically relevant question is how large the trend is, not whether it exists — and $P = 0.00$ does not answer that. A p-value tests only one thing: whether an effect is distinguishable from zero. With the millions of pixels underlying these panels, the answer to that question is always "yes" — even a slope so shallow it has no biological meaning will return $P < 0.0001$. The reported $P = 0.00$ (i.e., $p < 0.001$, rounded) tells the reader nothing about whether the observed trend is ecologically meaningful, only that the dataset is very large. What the figure needs instead is: (1) The slope with its confidence interval (e.g., -0.12 ± 0.03 Mg C ha $^{-1}$ yr $^{-1}$ per year of regrowth or degradation), so readers can judge whether the rate of change is large or negligible relative to total AGC stocks. (2) An effect size (R^2), such as the total AGC change attributable to the trend across the observed range of patch sizes or years, to contextualise the slope in carbon-budget terms. So, in summary: Concretely, the authors should:

1. Report R^2 for each regression line in panels b and d. This is the minimum.

2. Report the slope with a 95% confidence interval for each regression line in panels b and d. For example, "AGC increases at 0.8 ± 0.1 Mg C ha⁻¹ per year of regrowth" — so the practical magnitude is clear.

Figure 3, Panel f. The stacked plot is hard to interpret. Time trends obscure ecological dynamics with TMF detection history (fewer disturbances detected in earlier years due to data limitations).

Response: We have revised the figure and the associated text to improve clarity and interpretability. For panels a, c, and e, we retain the standard deviation to reflect the full variability within each class and clarify its interpretation in the text. The large variability primarily reflects heterogeneity within disturbance classes, particularly for degradation, where small patches often include a mixture of partial canopy loss and recovery processes, resulting in mean changes close to zero but wide dispersion. We have revised the Results section accordingly to avoid overinterpreting mean values: Degradation events show a smaller decrease in biomass density than deforestation events (**Fig. 3c**), with substantial variability across patches, particularly at small sizes, likely reflecting a mixture of partial canopy loss, recovery, and heterogeneous disturbance processes, and resulting in a wide dispersion. (Lines 242–245)

For panels b and d, we agree that p-values are not informative given the very large sample size. We have therefore removed the reported p-values and now present the slope together with its 95% confidence interval, as well as the coefficient of determination (R^2), for each regression. This allows readers to assess the magnitude and ecological relevance of the observed trends rather than their statistical significance alone. For panel f, we acknowledge that temporal trends may be influenced by detection limitations in earlier years. We have added a note in the manuscript to clarify that early-year patterns should be interpreted with caution and are not over-interpreted as:

For regrowing forests, high AGC growth rates are sustained for two years after regrowth is detected, followed by declining growth rates thereafter (**Fig. 3b**). This temporal dynamics of growth rates is however sensitive to TMF legacy data, as early-period TMF estimates may be biased toward easier to detect larger disturbances due to limited Landsat data availability in the early 2000s. Nevertheless, **Fig. 3f** shows that small patches (<1 ha) of regrowth in young secondary forests exhibit substantial AGC gains and play a dominant role in basin-scale carbon gains. In contrast, forests that experienced degradation in the past show a large initial loss of AGC density that attenuates over time. Based on this space-for-time substitution approach that assumes no secular trends in degradation severity, recovery conditions, or climate over the past two decades, we found no clear evidence of a return to a carbon gain regime in the first ~20 years after degradation (**Fig. 3d**). This result is consistent with recent GEDI-based analyses showing saturation in young regrowth and persistent losses in degraded forests¹⁰. (Lines 279–291)

Figure 4, Panels c, e. Individual protected areas and concessions are too small to be visible at this scale. A histogram or cumulative distribution of per-area net AGC change would provide more insight. The finding that 59% of protected areas experienced net losses lacks a benchmark (Lines 254-255) — what is the loss fraction outside protected areas?

Response: We have added an histogram of per-area net AGC change, with the mean value indicated. In addition, we have included a table summarizing the total AGC change, mean AGC density

change, and the fractions of loss and gain in protected and concession areas (**Supplementary Table 3**). We have also revised Panels c and e, which originally displayed total AGC change in each protected and concession area, to instead show mean AGC density change, as this representation is clearer at this scale. The original versions of Panels c and e have been moved to **Supplementary Fig. 31**.

Figure 5, Panel a. Agreement varies between countries but is described as "aligning well." The DRC discrepancy is the most significant and is interpreted asymmetrically.

Figure 5, Panels b, c. The suggestion that inventories overstate gains (due to IPCC defaults) and understate losses (due to omitting degradation) is plausible, but the alternative explanation — that the authors' maps underestimate gains because of model attenuation bias — is not discussed. Uncertainty bars do not account for systematic biases.

Response: We have updated the estimated uncertainty in **Fig. 5** and **Supplementary Table 3**. We also clarified related statement and explanation in the revised manuscript as:

Net biomass carbon changes at country scale (**Fig. 5a**) show broadly consistent magnitudes with national inventories¹⁷ harmonized by Grassi *et al.*¹⁶ in Gabon and Equatorial Guinea which have relatively small and stable forest sectors, with differences of 1 and 0.2 Tg C yr⁻¹, and of 5 Tg C yr⁻¹ in Cameroon where there is more deforestation. Differences between our estimates and Grassi *et al.* are larger in large countries such as DRC that dominate the regional carbon budget. In the DRC, our estimate indicates a smaller net carbon loss than the inventory (**Supplementary Table 4**). This discrepancy may reflect differences between the harmonized inventory framework¹⁶ and our satellite-based biomass change estimates, particularly in the partitioning of gross carbon losses and gains, as well as uncertainties in our estimates, including possible underestimation of biomass losses in low-canopy non-dense forests. (Lines 388–397)

To evaluate this assumption, we calculated the average forest carbon growth rate from our maps and found that our growth rates over intact (undisturbed) forests were close to zero or even negative (**Supplementary Table 5**), lower than the refined IPCC default values⁵⁸. This result suggests that the gross carbon gains reported in national inventories may be overestimated by more than 50% in Gabon, Cameroon, and the DRC (**Fig 5b**). At face value, this difference may also reflect methodological limitations in our satellite-based height change estimates, including potential regression-to-the-mean effects in deep-learning models and saturation of satellite data in very high biomass forests that may attenuate height gains in old-growth forests. National inventories also report lower gross carbon losses than our estimates (**Fig. 5c**). This difference partly reflects methodological differences between the two approaches. National inventories primarily report emissions associated with deforestation and do not explicitly report human-induced degradation losses, whereas biomass losses in our study include both deforestation and forest degradation, as well as natural disturbance processes. By explicitly separating carbon gains and losses, our approach provides a complementary perspective on the dynamic of all disturbances and associated biomass carbon changes. (Lines 417–433)

Extended Data Fig. 1. The framework flowchart highlights various adjustable steps without concurrent optimisation. A sensitivity analysis figure displaying total detected area or total carbon flux relative to key parameters would be useful.

Response: Thank you for this suggestion. Owing to the complexity of the processing steps, it is hard to include the sensitivity analysis within **Extended Data Fig. 1**. Instead, we report the total detected area for each of the four CH change processing steps in this figure, and provide two tables summarizing total AGC and area change under different processing parameters (**Supplementary Tables 9 and 10**). In addition, we provide a framework flowchart for uncertainty estimation (**Supplementary Fig. 30**), which includes necessary information for sensitivity analysis.

Extended Data Fig. 2, Panel a. This is an important figure. Canopy height change: satellite CH vs. ALS at 30 m resolution ($R^2 = 0.53$, $MAE = 3.19$ m, $Bias = +1.75$ m). I have several assumptions about this figure:

Positive bias (+1.75 m): the satellite product systematically overestimates height change relative to ALS. The bias is defined as $\Delta CH_{\text{satellite}} - \Delta CH_{\text{ALS}} = +1.75$ m, meaning the satellite-derived height changes are on average 1.75 m larger in the positive direction than the ALS reference. This has two directional consequences: height gains are overestimated (the satellite records more recovery than ALS measures indicate), and height losses are underestimated (the satellite records shallower losses than ALS measures detect). Both effects compound in the same direction — they jointly cause the net carbon flux to appear less negative than it truly is. The headline figure of -58 ± 9 Tg C yr⁻¹ may therefore be an underestimate of actual net carbon loss. This directional bias is not discussed in the main text and should be, since it affects the central conclusion about whether Central African forests are a small or large net source.

The bias magnitude (1.75 m) is large relative to the signal being detected. The MAE of 3.19 m already implies substantial pixel-level error, but the systematic component (1.75 m) is the more consequential concern: random errors partially average out at regional scale, whereas a systematic positive bias does not. For small disturbances — the paper's central novelty — where CH losses may be only 2–5 m, a +1.75 m bias represents a 35–90% underestimation of the height-loss signal. The contribution of sub-hectare degradation patches to gross carbon losses (reported as 37% of the total) is therefore likely understated.

The temporal interpolation used to align the ALS data introduces a further upward bias risk. The ALS scans from 2015 and 2023 are bridged to the 2019–2022 study window by assuming a linear annual rate of height gain (condition 1, Supplementary Text 2) and by excluding pixels with ALS losses between 2015 and 2019 (condition 2). Condition 2 selectively removes pixels that lost height before the study period, which preferentially eliminates the largest-loss reference points from the comparison. This pre-selection biases the ALS reference distribution toward smaller or positive changes, which would mechanically increase the apparent positive bias of the satellite product relative to a complete ALS reference. The authors do not report how many pixels were excluded under condition 2, nor whether the bias estimate changes when they are retained.

Taken together, a +1.75 m positive bias means the satellite product is more likely to understate carbon losses and overstate carbon gains than to overstate losses, as the previous review comment incorrectly implied. The carbon budget's net figure is therefore a lower bound on actual net loss —

a conclusion with direct implications for the paper's policy-relevant framing that Central African forests are a "small net source."

Response: Thank you for the insightful comments. We note, however, that the interpretation is based on the previous formulation of the comparison, in which the axes were inverted and the bias was defined as $\Delta\text{CH_ALS} - \Delta\text{CH_satellite}$. In the revised analysis, we have corrected this by placing ALS-derived canopy height change on the x-axis and satellite-derived change on the y-axis, such that bias is now consistently defined as: $\text{Bias} = \text{mean}(\Delta\text{CH_satellite} - \Delta\text{CH_ALS})$. Under this definition, the estimated bias is **-1.75 m**, indicating that the satellite product underestimates canopy height change relative to ALS, rather than overestimating it. As a consequence, the directional interpretation proposed by the reviewer is reversed: instead of overstating gains and understating losses, our satellite-derived map tends to estimate more negative canopy height changes overall, implying a tendency toward slightly higher carbon losses and lower carbon gains. However, due to the limitations of the current validation datasets, we do not have sufficient evidence to robustly determine whether the net carbon change is systematically overestimated or underestimated at the regional scale.

We have added some results and detailed discussion in [Supplementary Text 5](#) as:

To further understand the origin of this bias, we examined the effects of temporal interpolation and sample filtering on the comparison between satellite-derived canopy height (CH) change (2019–2022) and ALS observations (2015–2023) at both 10 m and 30 m resolutions (**Supplementary Figs. 12, 13**). Applying a linear interpolation to pixels with height gain (condition 1) resulted in only minor changes in agreement metrics, indicating that assumptions on growth dynamics have a limited influence on the comparison. In contrast, excluding pixels with ALS-detected height loss prior to the 2019–2022 study period (condition 2) led to a substantial improvement in agreement, with increased R^2 and reduced error metrics at both spatial resolutions. This reflects the effect of temporal mismatch between ALS acquisition (2015–2023) and the study period, whereby disturbances occurring outside the target window introduce inconsistencies in the raw comparison. Removing these pixels reduces discrepancies associated with misaligned disturbance timing rather than altering the underlying distribution of canopy height change. Importantly, consistent patterns across 10 m and 30 m analyses indicate that the observed improvements are not driven by scale-dependent smoothing or selective bias, but by improved temporal alignment between datasets. These results suggest that differences in the raw comparison are primarily attributable to timing and sampling inconsistencies, rather than systematic overestimation of height change in the satellite product.

We also note that the ALS comparison is based on a small spatial extent and may not fully represent the balance of canopy height gain and loss across the broader region. Conversely, validation using GEDI data provides larger spatial coverage but is spatially sparse and subject to known uncertainties on GEDI footprints, including geolocation errors and sampling limitations. Importantly, the GEDI-based validation shows predominantly positive bias, in contrast to the negative bias observed in the ALS comparison. As such, neither dataset alone can fully constrain the regional bias in canopy height change. The contrasting bias patterns observed between ALS and

GEDI further suggests that bias estimates are dataset-dependent, rather than indicative of a consistent directional error in the satellite-derived product.

Therefore, the current validation datasets do not provide sufficient evidence to robustly determine whether the net carbon change is systematically overestimated or underestimated at the regional scale. From a methodological perspective, our framework may tend to produce slightly conservative estimates of forest gain, suggesting that carbon gains could be underestimated, e.g. in mature forests. However, independent validation (**Extended Data Table 1**) indicates that both loss and gain estimates are subject to uncertainty, with relative uncertainties of approximately 16% and 23%, respectively. These uncertainties highlight that both components of the carbon budget remain imperfectly constrained, and further validation with expanded spatial and temporal coverage will be necessary to better quantify potential biases in net carbon change.

Extended Data Fig. 2, Panel b. $R^2 = 0.22$ at 10 km is the key independent validation metric but is buried in Extended Data. This should be in the main text with honest discussion of implications.

Extended Data Fig. 2, Panel c. Only "representative cases" of confirmed losses are shown. False-positive examples (predicted loss, no visible change in Planet) are absent. A balanced presentation requires both. This is important!

Response: Thank you for these helpful suggestions. We note that the reported value in Extended Data Fig. 2b has been updated to $R^2 = 0.23$ in the revised analysis (previously 0.22). We have also added a description and discussion of this independent validation metric in the main Results section to clarify its implications as:

Comparisons with changes from GEDI discrete footprints samples (RH100) between the first and last years of our study period showed weak agreement at 30 m but improved when aggregated to 10 km ($R^2 = 0.23$, $MAE = 4.35$ m) (**Extended Data Fig. 2, Supplementary Fig. 14**). This reflects several factors, including scale mismatch, sparse GEDI footprint sampling with no revisit in time, and geolocation uncertainty⁴², as well as the difficulty of detecting small height changes and uncertainties in our biomass mapping at fine spatial scales. (Lines 145–149)

In addition, to provide a more balanced evaluation, we have reorganized the visual validation examples. The previous Panel c has been moved to a new [Extended Data Fig. 3](#), which now includes both correctly detected (true positive) and falsely detected (false positive) cases for both canopy height loss and gain.

Extended Data Fig. 7. The finding that concessions have higher AGC density (144 vs. 127 Mg C ha⁻¹) than protected areas is explained speculatively. Selection bias in protected area placement (protection of lower- biomass ecosystems for biodiversity) is not considered. A p-value < 0.01 is expected with large samples; practical significance should also be discussed.

Lines 268–272 and Extended Data Fig. 7. "Protected areas have a smaller mean AGC density (127 Mg C ha⁻¹) than concessions (144 Mg C ha⁻¹)...differences between all pairs of groups are statistically significant ($p < 0.01$)."

The 17 Mg C ha⁻¹ difference is statistically significant, but the paper does not ask whether it is large enough to matter. Three things should be discussed:

1. Relative magnitude. 17 Mg C ha⁻¹ represents approximately a 13% difference from a baseline of 127 Mg C ha⁻¹. Given that the allometric model has a relative error of 32.6% (RMSE = 98.89 Mg C ha⁻¹), the observed difference between the two land-use categories is well within the measurement uncertainty of a single estimate. The authors should explicitly compare the 17 Mg C ha⁻¹ difference to their reported uncertainty to determine whether it can be distinguished from noise.
2. Carbon-budget consequence. The key question is not whether concessions have higher AGC stocks, but whether this difference significantly impacts the net carbon loss conclusions. Protected areas and concessions together span 36.7 Mha (17.1% of the total forest). A 17 Mg C ha⁻¹ density difference across this area results in a stock difference of approximately 0.6 Pg C; however, the influence of this on annual loss rates is never quantified.
3. Overlap in distributions. The violin plots in Extended Data Fig. 7 demonstrate that the distributions of AGC density for protected areas and concessions overlap significantly. A mean difference of 17 Mg C ha⁻¹ between two highly overlapping distributions has limited practical significance: a randomly selected concession pixel is hardly distinguishable from a randomly chosen protected-area pixel. Reporting the overlap coefficient or the proportion of concession pixels that exceed the protected-area mean would communicate practical relevance far better than a p-value alone.

Response: We have substantially revised the manuscript with the updated [Extended Data Fig. 7](#) and added a dedicated [Supplementary Text 9](#) to explicitly assess the magnitude, robustness, and implications of the observed difference in AGC density between protected areas and forest concessions.

First, we now explicitly compare the observed difference (17 Mg C ha⁻¹, ~13%) with the uncertainty of the biomass estimates. The relative RMSE of the allometric model is 32.6% (98.89 Mg C ha⁻¹), which substantially exceeds the observed difference. This indicates that the difference in mean AGC density cannot be robustly distinguished at the pixel level.

Second, we evaluated the practical significance of this difference using distribution-based metrics. We find strong overlap between AGC density distributions (overlap coefficient = 0.83), a small standardized effect size (Cohen's d = 0.36), and substantial bidirectional overlap (41% of protected-area pixels exceed the concession mean). These results demonstrate that, despite statistical significance (p < 0.01), the two distributions are only weakly separable in practice.

Third, we explicitly assessed the implications for carbon budgets. A difference of 17 Mg C ha⁻¹ across the total concession area (36.7 Mha) corresponds to a stock difference of ~0.6 Pg C. However, we show that this difference does not translate into measurable differences in annual carbon loss rates. Protected areas and concessions exhibit highly similar carbon dynamics, including comparable total AGC change (−3.8 vs. −5.4 Tg C yr⁻¹), mean AGC density change (−0.09 vs. −0.15 Mg C ha⁻¹ yr⁻¹), and nearly identical loss and gain fractions (6.0% vs. 6.5%, and 4.1% vs. 3.8%).

Finally, we now clarify that the observed difference in AGC density likely reflects non-random spatial placement of protected areas, rather than a causal land-use effect. Taken together, these additional analyses demonstrate that the difference in baseline AGC density is modest, strongly overlapping, and small relative to model uncertainty.

Extended Data Table 1. "Other forests" (TMF-undisturbed) show gross losses of $-22.4 \text{ Tg C yr}^{-1}$ and gross gains of $18.8 \text{ Tg C yr}^{-1}$. These large fluxes in supposedly undisturbed forests could represent real natural dynamics, TMF underestimation of disturbance, or false positives in the CH pipeline. The paper favours the first two interpretations but does not adequately exclude the third.

Response: We have revised the description of the results for "Other forests" as follows:

The AGC budget of so-called 'other forests', classified as apparently undisturbed in TMF but not exempt of AGC gains and losses, is close to neutral and not distinguishable from zero within uncertainty (**Extended Data Table 2**), although these forests exhibit substantial localized gross gains and losses. Their loss and gain fluxes in 'other forests' likely reflect a combination of real features (gaps, natural disturbances, human disturbances missed by TMF) and uncertainties in our disturbance classification and CH change data (**Extended Data Table 1**). This result highlights the overlooked dynamics of carbon in forests classified as apparently undisturbed. (Lines 316–323)

Reviewer #2 (Remarks to the Author):

Key results

The key results of this paper are that the majority of the net carbon gains and losses in the Tropical Moist Forest in the Congo basin occur in very small patches, but because there are so many of them the overall impact on the carbon balance is substantial.

Response: We appreciate the reviewer's constructive evaluation of our manuscript. We agree with the reviewer's overall assessment that small disturbance patches play a disproportionate role in shaping the carbon balance of Central African tropical forests, which is the central result of this study. We also acknowledge the reviewer's key concerns regarding the role of spatial resolution, the interpretation of canopy height change at 30 m, and the clarity of methodological description. In response, we have revised the manuscript extensively to address these points. In particular, we have clarified the role of spatial resolution in both validation and analysis, moved key comparisons from the Supplementary Information to the main text, and improved the presentation of canopy height change. Detailed responses to each comment are provided below.

In addition, following the numerous and excellent comments from the first reviewer, we have strengthened the uncertainty quantification and sensitivity analyses, and provided additional validation to support the robustness of our results.

Validity

For most parts, the analysis methods and data interpretation are robust and the conclusions are justified by the results. However, I have one major concern that the authors should address before the paper can be published:

Lines 111-114: One of the main points of novelty of this paper is, as the authors argue in the abstract, the 10 m resolution of the canopy height and 30 m resolution of the biomass maps. Near lines 111-114, the paper should provide the comparison of canopy height change on the 30 m spatial resolution in addition to the aggregated analysis that is currently here (this is currently provided in Supplementary Figures 10 and 11).

Response: Thank you for the comment. We agree that explicitly demonstrating the effect of spatial resolution is important. We first clarify that canopy height maps were generated at 10 m resolution, while aboveground biomass (AGC) and carbon change were derived at 30 m resolution. This choice reflects both the spatial scale of available field and reference data and the fact that large tropical tree crowns often extend beyond 20 m, making 30 m a more appropriate scale for biomass estimation and validation.

In response, we have revised the manuscript to clarify the role of spatial resolution in both validation and subsequent analyses. We have also moved the comparison of canopy height change at 30 m resolution (previously Supplementary Figures 10 and 11) to the main text, and expanded the validation to include both 10 m and 30 m comparisons against airborne LiDAR (ALS), as well as GEDI-based comparisons at 30 m. We have revised relevant description, results, and Method section as:

To address this knowledge gap, we generated annual 10 m canopy height maps for 2019–2022 using a deep learning model trained using spaceborne lidar height and Sentinel-1 radar and Sentinel-2

optical imagery, aggregated them to 30 m, and derived aboveground biomass changes using an allometric model. (Lines 39–42)

A deep learning model was trained to predict maximum CH measured by the GEDI spaceborne LiDAR^{39,40} from Sentinel-1 radar and Sentinel-2 multispectral imagery, enabling spatially continuous estimation of CH (**Methods, Extended Data Fig. 1**). These 10-m CH estimates were then aggregated to 30 m and converted to aboveground carbon (AGC) using field inventory data and allometric relationships from ref. (26). The 30 m resolution matches the spatial scale of available field plots, knowing that large tropical tree crowns often extend over 20 m⁴¹. (Lines 107–113)

Validation was primarily conducted at 30 m resolution due to scale compatibility with ALS data and consistency with biomass change data used in subsequent analyses, $R^2 = 0.53$ and MAE = 3.19 m (**Extended Data Fig. 2, Supplementary Fig. 11**). Agreement at 10 m was lower ($R^2 = 0.22$, MAE = 3.70 m) than at 30 m due to increased fine-scale heterogeneity and unresolved sub-pixel canopy gaps relative to ALS data, although MAE remained comparable across resolutions. Aggregation to 30 m improves the agreement through spatial smoothing, whereas the 10 m product preserves fine-scale canopy disturbances. Comparisons with changes from GEDI discrete footprints samples (RH100) between the first and last years of our study period showed weak agreement at 30 m but improved when aggregated to 10 km ($R^2 = 0.23$, MAE = 4.35 m) (**Extended Data Fig. 2, Supplementary Fig. 14**). This reflects several factors, including scale mismatch, sparse GEDI footprint sampling with no revisit in time, and geolocation uncertainty⁴², as well as the difficulty of detecting small height changes and uncertainties in our biomass mapping at fine spatial scales. (Lines 138–150)

Supplementary Figure 11 suggests that there is no (or almost no) correlation between the difference in canopy height from GEDI and the deep learning model at 30 m resolution. Is RH100 a good choice of GEDI product or would you find a better relationship if you had used RH99 or RH95, for example?

Response: We used GEDI RH100 as the reference metric because our canopy height model was trained using RH100, ensuring methodological consistency. In our previous work¹², RH100 has also been shown to better represent tall canopies and reduce saturation effects in dense tropical forests, making it suitable for tropical forest regions. Following the reviewer's suggestion, we tested alternative GEDI metrics (RH95 and RH98). Using RH95 or RH98 slightly improves agreement but does not change the overall pattern, indicating that discrepancies are largely independent of the choice of RH metric. The relatively weak correlation at 30 m reflects several factors, including geolocation uncertainty and sparse repeat sampling in GEDI footprints, the difficulty of detecting small-magnitude canopy height changes, and residual model uncertainty at fine spatial scales. As expected, agreement improves at coarser spatial resolution, where these uncertainties are reduced and sampling density increases. In contrast, comparisons with airborne LiDAR data at 30 m show stronger agreement, supporting the robustness of our canopy height change estimates.

We have revised the manuscript and included relevant results as:

Comparisons with changes from GEDI discrete footprints samples (RH100) between the first and last years of our study period showed weak agreement at 30 m but improved when aggregated to

10 km ($R^2 = 0.23$, MAE = 4.35 m) (**Extended Data Fig. 2, Supplementary Fig. 14**). This reflects several factors, including scale mismatch, sparse GEDI footprint sampling with no revisit in time, and geolocation uncertainty⁴², as well as the difficulty of detecting small height changes and uncertainties in our biomass mapping at fine spatial scales. Using the RH95 or RH98 height metrics slightly improves agreement but does not change the overall pattern (**Supplementary Fig. 15**), indicating that discrepancies are largely independent of the choice of height metric. These results highlight that validation at very small spatial scales remains limited^{38,39}, and the maps are primarily interpreted in terms of aggregated spatial patterns rather than individual pixels. (Lines 145–154)

Significance

The potential significance of the conclusions of the paper lies in the high importance of small-scale deforestation, forest degradation and regrowth patches in Tropical Moist Forests in Africa for the overall carbon balance. This is not recognised and addressed by most current studies of the carbon balance of Africa.

Response: We appreciate this comment. We agree that the importance of small-scale disturbance patches for the overall carbon balance is a key implication of this study. In the revised manuscript, we substantially strengthened the independent validation of small disturbances using a greatly expanded collection of high-resolution (5 m) PlanetScope imagery. We adopted a stratified sampling design⁴¹ based on patch size (0–1, 1–5, 5–10 and 10–100 ha) and change type (gain and loss), and visually interpreted randomly selected samples from bi-temporal 2019 and 2022 imagery (Extended Data Table 1). Compared with the previous version, which only included validation of small patches (<1 ha), the revised analysis includes substantially more validated disturbance patches across multiple size classes, providing a more robust assessment of small-scale canopy disturbances. ([Extended Data Table 1](#))

Line 83: This part of the review of previous work omits other studies that have used higher resolution biomass change maps, e.g. Rodríguez-Veiga et al. at 100 m resolution. Relevant studies of higher resolution biomass change maps should be reviewed and cited here briefly.

Reference:

Rodríguez-Veiga et al. "Loss of tropical moist broadleaf forest has turned Africa's forests from a carbon sink into a source." *Scientific Reports* 15.1 (2025): 41744.

Response: Thank you for the suggestion. We have revised the manuscript to include the recommended reference on higher-resolution biomass change maps (e.g. ~100 m), and we incorporated recent studies in the opening paragraph to better reflect the current understanding of carbon dynamics in African forests. The text has been revised to reflect this updated and add new references from Hubau et al. (2020), Rodríguez-Veiga et al. (2025), and Sagang et al. (2026) as:

Although intact forests have historically acted as a major carbon sink⁴, recent studies suggest that this sink is weakening, tending toward carbon neutrality or even becoming a net source⁵⁻⁷. (Lines 58–60)

Existing biomass change maps relying on coarse-resolution satellites (500 m to 25 km) have shown inconsistent estimates of carbon sources and sinks across tropical regions^{3,29,30}. A recent study has explored higher-resolution approaches (for example at ~100 m resolution)⁶. (Lines 95–98)

Data and methodology

My opinion is that the approach of using deep learning to estimate canopy height from Sentinel-1 and 2 trained on GEDI RH100 data is valid. The quality of the presentation of the data and results is excellent.

Analytical approach

The analytical approach and statistical tests applied follow best practice in the discipline.

Response: Thank you. To further strengthen the reliability of the data and methodology, we have improved the quantification of uncertainty in aboveground carbon (AGC) change by explicitly quantifying three major components: (1) biomass uncertainty arising from allometric modelling and measurement errors (σ_{bio}); (2) classification uncertainty associated with disturbance attribution and detection accuracy (σ_{cls}); and (3) uncertainty related to the choice of a CH change processing scheme (σ_{scheme}), including trend detection, Sentinel-1 filtering, and canopy cover filtering.

(Please refer to our responses to the related comments in the General Overview section and to Comment 2 from Reviewer 1 regarding uncertainty estimation.)

In addition, we performed a sensitivity analysis of the filtering parameters used in canopy height change detection, evaluating how key thresholds influence both disturbance detection and the resulting AGC change estimates. These results have been incorporated into the revised manuscript. (Please refer to our response to Comment 4 from Reviewer 1 regarding the processing pipeline and sensitivity analyses)

Suggested improvements

Line 37: The abstract mentions some causes of deforestation in the Congo rainforest but fails to mention mining, which I suggest should be added.

Response: We agree that mining is an important and increasingly significant driver of forest disturbance in the Congo Basin. Accordingly, we have include it in the abstract as:

Dense humid forests in Central Africa hold vast biomass carbon stocks but are increasingly affected by logging, smallholder agricultural clearing, mining activities, and fire disturbances. (Lines 35–36)

Line 39: In the abstract, it is also not clear how the authors get from canopy height maps to estimates of carbon gains and losses. This needs to be made clear here.

Response: We have revised the abstract as:

To address this knowledge gap, we generated annual 10 m canopy height maps for 2019–2022 using a deep learning model trained using spaceborne lidar height and Sentinel-1 radar and Sentinel-2 optical imagery, aggregated them to 30 m, and derived aboveground biomass changes using an allometric model. (Lines 39–42)

Line 39 also states the spatial resolution of 10 m for the canopy height maps. This suggests to the reader that the paper presents an analysis at that scale, when in fact the main analysis was done at 30 m resolution with the biomass maps. This needs to be stated in the abstract for clarity.

Response: We have revised the abstract to explicitly distinguish between the 10 m resolution canopy height maps and the 30 m resolution biomass products used in the main analysis.

To address this knowledge gap, we generated annual 10 m canopy height maps for 2019–2022 using a deep learning model trained using spaceborne lidar height and Sentinel-1 radar and Sentinel-2 optical imagery, aggregated them to 30 m, and derived aboveground biomass changes using an allometric model. (Lines 39–42)

Line 40: The abstract is unclear as to which Sentinel satellite data were used exactly, presumably Sentinel-1 and 2. This needs to be stated here.

Response: We have revised the abstract to explicitly state that both Sentinel-1 radar and Sentinel-2 optical data were used in the model. Please see our reply to the above two comments.

Lines 69–74: The review of existing approaches to carbon accounting in Africa omits the RECCAP and RECCAP2 carbon budgets, which use not only models but also Earth observation data. Relevant papers should be included and cited here for completeness.

Response: We have revised the manuscript to include related RECCAP2 research papers as:

Building on these approaches, recent synthesis efforts such as the RECCAP2¹⁸ analysis by Ernst et al. ¹⁹, although important uncertainties remain, particularly in resolving disturbance processes. (Lines 78–80)

References:

18 Ciais, P. *et al.* Definitions and methods to estimate regional land carbon fluxes for the second phase of the REgional Carbon Cycle Assessment and Processes Project (RECCAP-2). *Geoscientific Model Development Discussions* **2020**, 1–46 (2020).

19 Ernst, Y. *et al.* The African regional greenhouse gases budget (2010–2019). *Global Biogeochemical Cycles* **38**, e2023GB008016 (2024).

Line 83: This part of the review of previous work omits other studies that have used higher resolution biomass change maps, e.g. Rodríguez-Veiga et al. at 100 m resolution. Relevant studies of higher resolution biomass change maps should be reviewed and cited here briefly.

Response: We thank the reviewer for this helpful suggestion. We have revised the manuscript to reflect this update and add new references from Hubau et al. (2020), Rodríguez-Veiga et al. (2025), and Sagang et al. (2026) as:

Although intact forests have historically acted as a major carbon sink⁴, recent studies suggest that this sink is weakening, tending toward carbon neutrality or even becoming a net source ^{5–7}. (Lines 58–60)

Existing biomass change maps relying on coarse-resolution satellites (500 m to 25 km) have shown inconsistent estimates of carbon sources and sinks across tropical regions ^{3,29,30}. A recent study has explored higher-resolution approaches (for example at ~100 m resolution)⁶. (Lines 95–98)

Lines 111-114: One of the main points of novelty of this paper is, as the authors argue, the 10 m resolution of the canopy height and biomass maps. However, the analysis was done at 30 m resolution, even though the paper presents evidence that there is no correlation between the canopy height differences from airborne LiDAR and the data produced by the authors at that spatial resolution. This needs to be fully addressed before the paper can be published.

Response: We agree that the apparent weak agreement at 30 m resolution requires clarification. In the revised manuscript, we explicitly addressed this issue by adding validation results across multiple spatial scales.

Comparisons with changes from GEDI discrete footprint samples (RH100) between the first and last years of our study period showed weak agreement at 30 m but improved when aggregated to 10 km ($R^2 = 0.23$, MAE = 4.35 m) (Extended Data Fig. 2, Supplementary Fig. 14). This reflects several factors, including scale mismatch, sparse GEDI footprint sampling with no revisit in time, and geolocation uncertainty⁴², as well as the difficulty of detecting small height changes and uncertainties in our biomass mapping at fine spatial scales. Using the RH95 or RH98 height metrics slightly improves agreement but does not change the overall pattern (Supplementary Fig. 15), indicating that discrepancies are largely independent of the choice of height metric. These results highlight that validation at very small spatial scales remains limited^{38,39}, and the maps are primarily interpreted in terms of aggregated spatial patterns rather than individual pixels. (Lines 145–154)

To provide independent evaluation at fine spatial scales, we conducted an independent validation using high-resolution (5 m) PlanetScope imagery with a stratified sampling design⁴³ across patch sizes and change types (Extended Data Fig. 3, Supplementary Text 6). This analysis indicated that the user's accuracy (UA) was 84% for loss and 77% for gain across the 0–100 ha range (area-weighted), and 79% and 69%, respectively, for small disturbances (<1 ha) (Extended Data Table 1). Omission and commission errors were ~6% and ~5%, respectively, when evaluated relative to the full landscape, corresponding to undetected changes within no-change areas and false detections in stable forests. (Lines 161–168)

Lines 252-253 contain an incomplete sentence that needs to be corrected.

Response: We have corrected the incomplete sentence as:

Protected forest areas and forestry concessions together represent 34% (73.5 Mha) of the total dense forest area. (Lines 325–326)

Line 394: NISAR and BIOMASS are not upcoming missions anymore. Both have been launched already.

Response: We have revised the text to reflect that BIOMASS and NISAR are no longer upcoming missions.

Recent missions such as BIOMASS, whose P-band radar is sensitive to large stems and will provide forest vertical tomography, together with NISAR, help reduce uncertainties associated with height-to-biomass conversion and improve monitoring of biomass dynamics. (Lines 504–507)

Clarity and context

The text is generally clear to read, with a few exceptions that I have described as issues for improvement.

Response: Thank you. We have addressed the issues mentioned above.

References

The manuscript references the previous literature appropriately, but it omits some publications of higher resolution biomass and canopy height maps, which should be added.

Response: Thank you. We have added relevant references in the revised manuscript. Please see our reply to the above comments.

Dear Editor,

Thank you for the opportunity to revise and resubmit our manuscript. We greatly appreciate the reviewers' constructive comments and suggestions, which have helped us further improve our work. We have carefully considered all comments and revised the manuscript accordingly. The major revisions are summarized below.

1. We included an additional visual assessment using randomly sampled 1-ha and 5-ha image chips extracted from the mapped canopy height (CH) change layer (**Supplementary Table 2**). The image chips were stratified according to their mean estimated CH loss magnitude ($-10 \leq \Delta CH < 0$ m, $-20 \leq \Delta CH < -10$ m and $\Delta CH < -20$ m), and representative examples are shown in **Supplementary Fig. 37**. This assessment was designed to quantify the reliability of CH loss detection (true and false detections) across different CH loss magnitude classes.
2. We added a new table summarizing the contributions of different disturbance patch-size classes and canopy height (CH) loss classes to the affected area and aboveground carbon (AGC) change (**Supplementary Table 3**). We also included an illustrative figure (**Supplementary Fig. 37**) comparing PlanetScope imagery, canopy height change, and AGC change maps across different CH-loss magnitude classes.
3. We updated several figures to avoid using red and green simultaneously within the same plot, thereby improving accessibility for readers with color-vision deficiencies. We also added geographic coordinates (latitude and longitude) to all maps throughout the main manuscript and supplementary materials.
4. We further revised the manuscript to clarify the descriptions related to disturbance patch size and uncertainty assessment.

We believe these revisions further strengthen the robustness of the methodology and the reliability of the conclusions. The line numbers referenced by the reviewers correspond to the originally submitted manuscript, while those cited in our responses refer to the revised manuscript with tracked changes. In the responses below, our replies are shown in blue, and the corresponding text, figures, and tables from the revised manuscript are shown in red. All modifications in the revised manuscript are consistently highlighted in red.

Response to reviewers

Reviewer #1

1. Comments for the Authors and Editor

I am very satisfied with the authors' response, the revisions they have made, and the overall improvements to the manuscript. The new analyses and additional data substantially strengthen the study. However, I remain concerned about one aspect of Section L185–200 (“High-resolution maps of height change reveal more small disturbances”), namely the validation of small canopy-height changes and their contribution to the reported carbon dynamics.

My concern is not related to the frequency distribution of disturbance patch sizes presented in Extended Data Table 1. Rather, it concerns the validation of small canopy-height changes (<10-

20 m) and the confidence with which these changes can be interpreted as real disturbances contributing to carbon loss and gain. This distinction is important because the manuscript appears to conflate two different concepts: (1) the size of disturbance patches and (2) the magnitude of canopy-height change. While the former is well addressed, the latter remains insufficiently demonstrated.

The inclusion of Extended Data Table 1 and Supplementary Table 1 is a valuable addition. Extended Data Table 1 provides convincing visual validation of canopy height loss and gain detections, as well as stable forest areas, using high-resolution PlanetScope imagery. However, it does not indicate the canopy-height change thresholds used for the visual interpretation. Consequently, it remains unclear how small canopy- height changes (<10–20 m), particularly those occurring in apparently stable forests, were independently validated.

This issue is important because the manuscript attributes a substantial fraction of the detected changes to these smaller height classes. For example, Lines 190–191 state:

“Quantitative analysis confirms that 80% of the areas where we detect large height losses (>20 m) are matching TMF events at 30 m resolution, indicating consistency for large disturbances.”

I find this result convincing. Likewise, the manuscript appropriately recognises that:

“The main difference arises for small height losses (<10 m), where our map detects 65% more affected area than reported in the TMF deforestation and degradation products.”

However, I remain unconvinced that the subsequent statement—“Independent visual validation using high-resolution PlanetScope imagery confirmed that 84% of these additional small-loss events represent real canopy disturbances.”—fully resolves this issue. It remains unclear whether this validation refers primarily to small disturbance patches or specifically to small canopy-height changes. Since the manuscript's conclusions regarding carbon dynamics depend on accurately detecting small changes in canopy height, I believe this distinction should be clarified more explicitly. To address this remaining concern, I believe one additional validation analysis would substantially strengthen the manuscript.

Reply: We agree that it is important to distinguish between disturbance patch size and canopy-height (CH) loss magnitude. As pointed out by the reviewer, **Extended Data Table 1** validates the detection of CH loss and gain patches across different patch-size classes, but it does not evaluate whether detections associated with different CH loss magnitudes correspond to real canopy disturbances. To address this concern, we have added a new validation analysis in **Supplementary Table 2**. Rather than stratifying samples by patch size, this analysis stratifies randomly sampled image chips according to their mean estimated CH loss magnitude ($-10 \leq \Delta CH < 0$ m, $-20 \leq \Delta CH < -10$ m, and $\Delta CH < -20$ m). We evaluated both 1-ha and 5-ha image chips using independent 5-m PlanetScope imagery acquired in 2019 and 2022.

We focused this additional analysis on CH loss because the concern was specifically related to the interpretation of small height losses and their contribution to carbon dynamics. An equivalent analysis was not performed for CH gain because ~97% of the mapped CH-gain area falls within the smallest gain class (<10 m), providing little scope for a meaningful stratification by CH gain magnitude. In addition, the comparison with existing forest disturbance products presented in the manuscript focuses exclusively on CH loss, making the additional CH loss magnitude validation the most relevant analysis for addressing the reviewer's concern.

We adopted an image-chip-based validation rather than validating individual CH loss pixels because CH changes cannot be reliably interpreted at the individual-pixel from optical imagery alone. Visual interpretation requires sufficient spatial context to distinguish true canopy disturbances from image noise, seasonal variation, or isolated tree crowns. Therefore, the

sampled image chip, rather than the individual CH loss pixel, provides the appropriate validation unit for assessing whether a canopy disturbance is genuinely present. The new analysis shows that the reliability of CH loss detection increases with increasing CH loss magnitude. User's accuracy ranges from 71.6% to 85.5% for 1-ha image chips and from 80.4% to 87.0% for 5-ha image chips across the three CH loss classes (**Supplementary Table 2**). Importantly, even for the smallest CH loss class ($-10 \leq \Delta CH < 0$ m), the majority of detections correspond to visually identifiable canopy disturbances, supporting our conclusion that a substantial proportion of these small CH loss detections represent real forest disturbances.

To avoid overinterpretation, we also revised the manuscript text to clarify that this validation assesses the reliability of CH loss detection (true versus false detections) across different CH loss magnitude classes, rather than the quantitative accuracy of the estimated CH loss values themselves, which cannot be independently evaluated from optical imagery alone.

We have updated the relevant description and results as:

To further evaluate the reliability of CH loss detections across different loss magnitudes, we conducted an additional visual assessment using randomly sampled 1-ha and 5-ha image chips extracted from the mapped CH change layer (**Supplementary Table 2**). These image-chip sizes refer to the area of the sampled image subsets rather than the size of individual CH loss patches, because CH changes cannot be reliably interpreted at the individual-pixel level using optical imagery alone. Image chips were stratified according to their mean estimated CH loss magnitude ($-10 \leq \Delta CH < 0$ m, $-20 \leq \Delta CH < -10$ m and $\Delta CH < -20$ m) (**Supplementary Fig. 37**). Monthly PlanetScope imagery acquired in 2019 and 2022 was visually interpreted to determine whether canopy loss was present within each sampled image chip. This assessment was designed to quantify the reliability of CH-loss detection (true and false detections) across different loss-magnitude classes, rather than to validate the quantitative accuracy of the estimated CH loss values, which cannot be independently assessed by visual interpretation alone. (Lines 648–660)

The main difference arises for small CH losses (<10 m), where our CH loss estimates identify 65% more affected area than is captured by the TMF deforestation and degradation products within the same disturbed regions. Independent visual validation using high-resolution PlanetScope imagery confirmed that $\sim 72\%$ of these small height losses represent real canopy disturbances (**Supplementary Table 2**). The contributions of different height change classes to the affected area and AGC change are summarized in **Supplementary Table 3**, highlighting the substantial contribution of small height losses. (Lines 194–201)

Supplementary Table 2: Visual validation of canopy-height (CH) loss detections stratified by mean CH loss magnitude using randomly sampled 1-ha and 5-ha image chips interpreted from 5-m Planet imagery. The 1-ha and 5-ha categories refer to the size of randomly sampled image chips extracted from the CH change map rather than the size of individual CH loss patches. Image chips were stratified according to the mean estimated CH loss within each chip. Visual interpretation using 5-m Planet imagery was used only to determine whether canopy loss was present (true positive) or absent (false positive); it does not provide an independent assessment of the quantitative accuracy of CH loss magnitude estimates.

Random image-chip size	CH loss magnitude (m)	Sample number	TP	FP	UA
1 ha	$-10 \leq \Delta CH < 0$	95	68	27	71.58%
	$-20 \leq \Delta CH < -10$	98	80	18	81.63%
	$\Delta CH < -20$	117	100	17	85.47%
5 ha	$-10 \leq \Delta CH < 0$	92	74	18	80.43%
	$-20 \leq \Delta CH < -10$	105	88	17	83.81%
	$\Delta CH < -20$	108	94	14	87.04%

2. Primary recommendation

Add a validation table stratified by canopy-height change classes.

A table following the same format as Extended Data Table 1, but organised by canopy-height change classes rather than patch-size classes, would directly address this issue. Specifically, I suggest replacing "Patch size" with "Canopy Height Change" and reporting results for three height classes (0–10 m, 11–20 m, and >20 m). For each class, please report the sample size, true positives (TP), false positives (FP), and area-weighted user's accuracy (UA) based on independent PlanetScope imagery.

This analysis differs from Supplementary Table 1, which evaluates the proportion of validated canopy-height loss patches that overlap existing forest disturbance products across different overlap thresholds. Instead, the proposed table would directly assess the accuracy of detecting canopy-height changes across different height classes, providing stronger evidence that these reported sub-canopy height changes represent real disturbances in otherwise stable forests.

Such an analysis would considerably strengthen confidence in one of the manuscript's principal conclusions—that small changes in canopy height make an important contribution to carbon dynamics across African forests.

Reply: Thanks a lot for this recommendation. As suggested, we have added an additional validation analysis stratified by canopy height (CH) loss magnitude and present the results in **Supplementary Table 2**. As explained in our response to **above Comment 1**, the validation is based on randomly sampled 1-ha and 5-ha image chips stratified by mean CH loss magnitude, which provides sufficient spatial context for reliable interpretation using PlanetScope imagery. We believe this additional analysis directly addresses the concern by demonstrating the reliability of CH loss detections across different CH loss magnitude classes, particularly for smaller CH losses.


3. Optional suggestions

If the authors consider it feasible, two additional analyses could further strengthen this section, although I do not regard them as essential.

1. An illustrative figure comparing PlanetScope imagery, canopy-height change maps, and biomass change for a representative stable forest area would help readers visually understand how canopy-height changes correspond to biomass changes across the three height classes (0–10, 11–20, and >20 m).

Reply: Thank you for the suggestion. We have included an illustrative figure (**Supplementary Fig. 37**) comparing PlanetScope imagery, canopy height change maps, and aboveground carbon (AGC) change across different canopy height loss classes. The

examples were derived from the additional visual assessment of canopy-height loss magnitudes.



Supplementary Fig. 37. Examples of canopy height (CH) loss and corresponding aboveground carbon (AGC) change across different CH loss magnitude classes. Representative examples of CH loss detections are shown for three CH loss magnitude classes ($-10 \leq \Delta CH < 0$ m, $-20 \leq \Delta CH < -10$ m and $\Delta CH < -20$ m), with two examples for each class (a–f). For each example, the top two rows show 5-m PlanetScope imagery acquired in 2019 and 2022, respectively, followed by the corresponding 30-m CH change map and AGC change map (2019–2022). The latitude and longitude of the center of each image chip are provided below each example. These examples illustrate how visually identifiable canopy disturbances correspond to increasing CH loss and associated AGC loss across the three CH loss magnitude classes.

2. A summary table quantifying the contribution of each canopy-height change class to affected area and aboveground carbon (AGC) across stable forests would provide useful context for interpreting the relative importance of these different height-change classes in the overall carbon budget. These additions would improve the presentation of the results, but are secondary to the validation analysis described above.

Reply: We have included a summary table (**Supplementary Table 3**) quantifying the contribution of each canopy-height change class to affected area and aboveground carbon (AGC).

Supplementary Table 3. Contributions of disturbance patch-size classes and canopy-height (CH) change classes to affected area and aboveground carbon (AGC) change associated with mapped forest loss and gain. Area is expressed in Mha and AGC change in Tg C yr⁻¹. Percentages for forest loss and forest gain were calculated separately relative to the total affected area and total AGC change within the corresponding loss or gain category. Values are rounded to one decimal place

Type	Patch size (ha)	Area (Mha)	Area (%)	AGC change (Tg C yr ⁻¹)	AGC change (%)
Loss	0–1 ha	4.5	23.2	–14.6	11.6
	1–5 ha	4.2	21.7	–21.2	16.8
	5–10 ha	1.8	9.4	–12.0	9.5
	10–100 ha	4.8	24.5	–37.8	30.0
	>100 ha	4.1	21.1	–40.5	32.1
Gain	0–1 ha	5.2	43.7	27.6	40.7
	1–5 ha	3.5	29.7	21.4	31.6
	5–10 ha	1.0	8.4	6.2	9.2
	10–100 ha	1.7	14.0	9.8	14.5
	>100 ha	0.5	4.2	2.8	4.1
	CH change (ΔCH , m)	Area (Mha)	Area (%)	AGC change (Tg C yr ⁻¹)	AGC change (%)
Loss	$-10 \leq \Delta CH < 0$	17.5	90.3	–80.9	64.2
	$-20 \leq \Delta CH < -10$	1.5	7.9	–31.7	25.2
	$\Delta CH < -20$	0.4	1.8	–13.4	10.6
Gain	$0 < \Delta CH \leq 10$	11.4	97.0	60.7	89.5
	$\Delta CH > 10$	0.3	3.0	7.1	10.4

Minor comments

4. Supplementary Table 1

The title of Supplementary Table 1 is difficult to follow. It is not immediately clear what is meant by "Percentages indicate the proportion of canopy-height (CH) loss patches," nor is it obvious which canopy-height classes are being considered. I recommend revising the title to describe more explicitly what is being quantified.

Similarly, the phrase "overlap threshold" used in Lines 196–198 should be defined more clearly. The statement:

"However, only a subset of these validated small CH loss patches overlapped with TMF disturbances (26%–78%), depending on the overlap threshold applied." would benefit from a brief explanation of how the overlap thresholds are defined and why different thresholds are used. This clarification would make the validation procedure much easier for readers to understand.

Reply: Thank you for this helpful suggestion. We have revised the title and caption of [Supplementary Table 1](#) to explicitly describe what is being quantified. The revised caption now states that, for each overlap threshold, the reported percentages represent the proportion of independently validated canopy height (CH) loss patches for which the disturbed area mapped by RADD, GFC, or TMF covers at least a specified percentage (10–50%) of the total CH loss patch area.

In addition, we revised the corresponding text in the main manuscript to define the overlap threshold when first introduced. The text now explains that the reported range (26–78%) depends on the minimum proportion of each validated CH-loss patch required to overlap a mapped TMF disturbance (10–50% of the CH loss patch area; [Supplementary Table 1](#)). These revisions clarify both the meaning of the reported percentages and the rationale for using multiple overlap thresholds.

The relevant description has been updated in the revised manuscript as:

Further comparison showed that only 26–78% of the independently validated small CH loss patches overlapped with TMF disturbances, depending on the minimum patch-overlap criterion applied (10–50% of the CH loss patch area; **Supplementary Table 1**). (Lines 168–170)

Supplementary Table 1: Percentage of independently validated canopy height (CH) loss patches detected by existing forest disturbance products under different minimum overlap thresholds. For each overlap threshold, percentages indicate the proportion of independently validated CH loss patches for which the disturbed area mapped by RADD, GFC or TMF covered at least the specified percentage (10–50%) of the total CH-loss patch area. (Supplementary materials)

5. Clarification of the minimum size of degradation and deforestation events

The manuscript would benefit from a clearer and more consistent definition of the minimum size of degradation and deforestation events. At present, the revised manuscript and the authors' response appear to use two different thresholds.

In the response to the reviewers, the authors state:

"We first clarify that the original estimate for small deforestation events is 1.48 ha, which we now report as ~1.5 ha in the revised manuscript to ensure consistency and precision."

However, Lines 214–217 state:

"Among the ~10 million small height-loss patches of 0.09 ha identified in our maps, 56% coincide with disturbances recorded in the TMF dataset. The remaining differences may reflect either additional small disturbances not captured by TMF or uncertainties in our detection of small changes."

Later, Lines 225–258 state:

"The key finding is that moderate degradation events from 1.5 to ~10 ha and small deforestation events of ~1.5 ha make the largest contribution to the total biomass loss in the study region..."

These statements appear to refer to different minimum disturbance sizes. It is therefore unclear whether the smallest detectable degradation or deforestation event is a single 30×30 m pixel (0.09 ha) or a disturbance patch of approximately 1.5 ha. I encourage the authors to clarify this distinction explicitly. If 1.5 ha is the minimum disturbance size used in the analyses and discussion, then the text referring to the 0.09-ha height-loss patches should explain how these individual pixels relate to the aggregated disturbance patches.

This clarification is important because it affects the interpretation of the reported results. According to Lines 214–217, approximately 10 million height-loss patches of 0.09 ha were identified, of which only 56% overlap disturbances recorded in the TMF dataset. Consequently, approximately 44% (roughly 4.4 million pixels, corresponding to about 396,000 ha) do not overlap existing TMF disturbance records.

The manuscript suggests that these unmatched pixels may represent either previously undetected disturbances or uncertainties in detecting small canopy-height changes. Because these two interpretations have very different implications for estimates of forest degradation and carbon loss, I recommend that the manuscript more clearly explain how to interpret these unmatched pixels. In particular, readers should understand whether these pixels are considered evidence of real degradation events, potential false positives, or candidate disturbances whose status remains uncertain.

Reply: To clarify this distinction, we have revised the manuscript to explicitly state that the ~10 million 0.09 ha CH-loss patches refer to single-pixel (30×30 m) CH loss patches identified in our CH change maps. In contrast, the values of approximately 1.5 ha (deforestation) and 1.5–10 ha (degradation) refer to the patch-size classes contributing the largest share of total biomass loss, rather than the minimum detectable disturbance size. We have revised the corresponding text and [Supplementary Table 3](#) to make this distinction explicit.

We have also clarified the interpretation of the CH loss patches that do not overlap TMF disturbances. These unmatched patches should not be interpreted exclusively as either true disturbances or false positives. Instead, they likely represent a combination of genuine small canopy disturbances that are not mapped by the TMF products and uncertainties associated with detecting small CH changes. The additional independent PlanetScope validation stratified by CH loss magnitude ([Supplementary Table 2](#)) supports this interpretation by showing that the majority of small CH-loss detections correspond to visually identifiable canopy disturbances.

We have revised the title of [Supplementary Table 2](#) and the relevant statement as:

Among the ~10 million 0.09 ha (single-pixel) CH loss patches identified in our maps, 56% coincide with disturbances recorded in the TMF dataset. The remaining differences may reflect either additional small disturbances not captured by TMF³⁴ or uncertainties in our detection of small changes. (Lines 214–217)

For each size category, we calculated the total biomass loss by multiplying the area of each event by its corresponding biomass loss density. The key finding is that degradation patches of approximately 1.5–10 ha and deforestation patches of approximately 1.5 ha contribute the largest share of total biomass loss in the study region, although uncertainty remains due to limitations in disturbance detection and classification ([Supplementary Table 4](#)). (Lines 257–262)

6. Clarification of the uncertainty associated with the disturbance-driven turnover rate

The manuscript would also benefit from a clearer presentation of the uncertainty associated with the estimated disturbance-driven turnover rate.

Lines 311–312 state:

"With a total stock of 24.3 Pg C in 2019, the turnover rate associated with disturbance-driven losses is $0.52\% \pm 0.10\%$ per year (Supplementary Text 11)."

Supplementary Text 11 (page 26) provides a thorough uncertainty analysis, which is a valuable addition to the manuscript. However, I believe the main text could better reflect these results.

First, since Supplementary Text 11 reports a relative uncertainty of 8.65% for the 2019 aboveground carbon (AGC) stock estimate, it would improve consistency and transparency to report this uncertainty in the main text. For example:

"With a total stock of $24.3 \text{ Pg C} \pm 8.65\%$ in 2019, the turnover rate associated with disturbance-driven losses is $0.52\% \pm 0.10\%$ per year (Supplementary Text 11)."

Second, it is not entirely clear how the $\pm 0.10\% \text{ yr}^{-1}$ uncertainty in the turnover rate was derived. Supplementary Text 11 (page 26) reports an estimated relative uncertainty of approximately

18.5% for the turnover rate. Although these values may be consistent, the relationship between the reported relative uncertainty and the absolute uncertainty presented in the main text is not explicitly explained.

I therefore recommend that the authors briefly describe how the $\pm 0.10\%$ uncertainty was calculated (or provide an explicit reference to the corresponding equation or calculation in Supplementary Text 11). This would make the uncertainty analysis more transparent and easier for readers to follow.

Reply: We have revised the corresponding paragraph to explicitly report the relative uncertainty of the 2019 aboveground carbon (AGC) stock estimate (8.65%) together with the reported AGC stock. We have also clarified the origin of the reported turnover-rate uncertainty ($\pm 0.10\% \text{ yr}^{-1}$) by stating that it was derived using first-order error propagation of the uncertainties associated with annual AGC loss and the 2019 AGC stock, as described in [Supplementary Text 11](#). We further specify that this propagation results in an overall relative uncertainty of 18.49% for the turnover-rate estimate, corresponding to an absolute uncertainty of approximately $\pm 0.10\% \text{ yr}^{-1}$ for the estimated turnover rate of $0.52\% \text{ yr}^{-1}$. In addition, we added an explicit reference to the uncertainty propagation described in [Supplementary Text 11](#) to make the calculation more transparent and easier for readers to follow.

We revised the original sentence as:

With a total AGC stock of 24.3 Pg C (8.65% relative uncertainty) in 2019, the turnover rate associated with disturbance-driven losses was $0.52 \pm 0.10\% \text{ yr}^{-1}$. The reported uncertainty was obtained by first-order propagation of the uncertainties in AGC loss (16.34% relative uncertainty) and the 2019 AGC stock, resulting in an overall relative uncertainty of 18.49% for the turnover-rate estimate (**Supplementary Text 11**). (Lines 314–318)

Finally, the results showed the total relative uncertainties of AGC_{loss} and AGC_{2019} are and 8.65%, respectively, resulting in an estimated relative uncertainty in the turnover rate of 18.49%, corresponding to an absolute uncertainty of approximately $\pm 0.10\% \text{ yr}^{-1}$ for the estimated turnover rate ($0.52\% \text{ yr}^{-1}$). (Supplementary Text 11)

7. Remarks on code availability:

The code repository is well organised and includes a computing environment (<https://codeocean.com>) for running the code. All data and code required to reproduce the article's figures are available in the repository. However, I encourage the authors to further improve the study's reproducibility by providing the complete set of processing pipelines used to generate the deep learning models for canopy height and biomass mapping. These training and model development pipelines are not yet available in the current repository. Making these workflows publicly accessible would substantially enhance the transparency, reproducibility, and long-term value of the study.

Reply: Thank you. The current repository already contains all data and code required to reproduce main statistics and figures presented in this study. In response to the reviewer's recommendation, we have further expanded the repository by including the complete processing pipeline used to derive the canopy height (CH) change products from the annual canopy-height maps, enabling readers to fully reproduce the CH change datasets used throughout this study.

We also appreciate the suggestion to provide the complete deep-learning training pipelines used to generate the annual canopy height and biomass products. These models were developed as part of previous collaborative work (<https://doi.org/10.3389/frsen.2025.1724950>) and involve

methodological components that are subject to collaborative agreements and are also being used in ongoing, unpublished research. Consequently, the complete model-training pipelines cannot be made publicly available at this stage. Nevertheless, the methodologies underlying the annual canopy height have been fully described in the corresponding publications and are also summarized in the Methods section of the present manuscript. Importantly, the canopy height and biomass maps used as inputs to this study are publicly available, and the repository now provides all data, code, and processing workflows required to reproduce the analyses, figures, and conclusions presented in this study. We therefore believe that the materials currently provided ensure the full reproducibility of this study while respecting the constraints associated with the underlying model development.

Reviewer #2 (Remarks to the Author):

1. What are the noteworthy results?

The key results of this paper are that the majority of the net carbon gains and losses in the Tropical Moist Forest in the Congo basin occur in very small patches, but because there are so many of them the overall impact on the carbon balance is substantial. My concerns regarding the role of the spatial resolution, the interpretation of canopy height change at 30 m, and the clarity of the description of the methods have been addressed in great detail by the authors. As a result, the revised manuscript is of much improved quality and I consider the methods, data, approach, and conclusions drawn from the work as robust and clear now.

Reply: We sincerely thank the reviewer for the careful re-evaluation of our manuscript and for the positive assessment of the revised version. We greatly appreciate the reviewer's constructive comments throughout the review process, which have substantially strengthened this work.

2. My most significant concern I raised in the first round of reviews was that the 10 m resolution of the canopy height and 30 m resolution of the biomass maps were described as the main novelty of the paper, but a comparison of canopy height change on the 30 m spatial resolution in addition to the aggregated analysis at a coarser resolution was missing.

The authors have clarified the role of spatial resolution in the revised manuscript and moved the comparison of canopy height change at 30 m resolution from the Supplementary Figures 10 and 11 to the main text, and expanded the validation to include both 10 m and 30 m comparisons against airborne LiDAR (ALS), as well as GEDI-based comparisons at 30 m. This, together with further substantial work, has fully addressed my previous concern. The accuracy gains of going to a coarser resolution are now much more visible in the paper and have been treated more robustly in the analysis.

Reply: Thank you.

3. Will the work be of significance to the field and related fields? How does it compare to the established literature? If the work is not original, please provide relevant references. Does the work support the conclusions and claims, or is additional evidence needed?

The revised manuscript is much improved in how it presents the significance of this study to the understanding of the global terrestrial carbon cycle and the role of African forest ecosystems. My minor comments have been fully addressed. The additional work presented in the revised paper fully supports the conclusions.

Reply: Thank you.

4. **Are there any flaws in the data analysis, interpretation and conclusions? Do these prohibit publication or require revision? Is the methodology sound? Does the work meet the expected standards in your field? Is there enough detail provided in the methods for the work to be reproduced?**

The quality of the data analysis, interpretation and conclusions now meets my expectations fully and the authors have addressed all my detailed suggested revisions.

Reply: Thank you for the positive evaluation of the methodology, data analysis, interpretation, and reproducibility of our work.

5. **My only open issue that I want to raise here is that in the main manuscript and the supplementary materials, all presented maps should have a scale bar and/or lat/lon or geographic coordinates, as is good practice.**

Reply: We have added geographic coordinates (latitude and longitude) to all maps in the main manuscript and supplementary materials, including [Supplementary Figs. 2, 25, and 29](#), and [Extended Data Fig. 5](#). For [Extended Data Fig. 3](#) and [Supplementary Fig. 37](#), which presents the PlanetScope validation samples, we added the latitude and longitude of the center of each image chip.

Response to reviewers

Reviewer #1 (Remarks to the Author):

1. Comments for the Authors

I congratulate the authors of this latest review. I am very satisfied with the authors' response, the revisions they have made, and the overall improvements to the manuscript. The new analyses and additional data substantially strengthen the study. Specifically, the authors addressed my last concerns by acknowledging the high and still-present uncertainty in validating small canopy-height changes (<10-20 m) and the confidence with which these changes can be interpreted as real disturbances contributing to carbon loss and gain. Among the new additions, including new analyses, figures, and tables, I am pleased to see in the main article the changes to the previous statement about the uncertainty from "Independent visual validation using high-resolution PlanetScope imagery confirmed that 84% of these additional small-loss events represent real canopy disturbances (Extended Data Table 1)" to "Independent visual validation using high-resolution PlanetScope imagery confirmed that ~72% of these small height losses represent real canopy disturbances (Supplementary Table 2). The contributions of different height change classes to the affected area and AGC change are summarised in Supplementary Table 3, highlighting the substantial contribution of small height losses".

Overall, the article now addresses the validation of small canopy-height changes (<10-20 m), a key driver of carbon loss and gain across African forests. I have no further concerns about this article.

Reply: We sincerely thank the reviewer for the positive assessment and are pleased that our revisions have addressed the remaining concerns.

2. Reviewer #1 (Remarks on code availability):

The code has been provided, and I understand that the complete deep-learning training pipelines used to generate the annual canopy height and biomass products are under access restrictions.

Reply: We thank the reviewer for the comment and understanding regarding the code availability. The code for canopy height change processing and figure generation has been made publicly available, while the deep-learning training pipelines used to generate the canopy height and biomass products were developed in our previous work and remain under access restrictions.