IBIMA Publishing
Communications of International Proceedings
https://ibimapublishing.com/p-articles/47AI/2026/4724326/
Vol. 2026 (12), Article ID 4724326

Hybrid Physics–Machine Learning Retrieval of AOD550 from HYPSO Hyperspectral Imagery*

Dawid JURCZYŃSKI

Silesian University of Technology in Gliwice, Poland

Correspondence should be addressed to: Dawid JURCZYŃSKI, dawid.jurczynski@polsl.pl

*Presented at the 47th IBIMA International Conference, 29-30 June 2026, Madrid, Spain

Abstract

Atmospheric aerosols significantly influence air quality, climate processes, visibility, and environmental hazard assessment, making aerosol optical depth at 550 nm (AOD550) an important parameter in satellite-based monitoring systems. Existing aerosol retrieval approaches are often computationally expensive, dependent on complex radiative-transfer inversion methods, and difficult to apply in rapid-response workflows. At the same time, despite the growing availability of hyperspectral satellite systems such as HYPSO, relatively little research has focused on using HYPSO data for practical AOD550 retrieval in near-real-time environmental and threat-assessment applications. This study addresses this gap by proposing a hybrid physics–machine learning framework for rapid scene-level AOD550 estimation from HYPSO hyperspectral imagery. The methodology consists of two stages. First, a physics-inspired baseline estimate is generated using compact spectral, reflectance-derived, and geometric descriptors extracted from HYPSO observations. Second, a supervised machine-learning model is applied to estimate and correct the systematic residual bias of the baseline retrieval relative to a CAMS-based reference aerosol field. The framework was evaluated using Leave-One-Scene-Out cross-validation across 33 hyperspectral scenes. The results demonstrate a substantial improvement after the bias-correction stage, reducing the out-of-fold mean absolute error (MAE) from 0.234 to 0.024 and the root mean square error (RMSE) from 0.243 to 0.039. The findings indicate that HYPSO observations contain meaningful aerosol-related information and that hybrid physics–machine learning workflows can support rapid and computationally lightweight AOD550 estimation for satellite-data processing and environmental threat-assessment systems.

Keywords: HYPSO, aerosol optical depth, AOD550, hyperspectral imagery, machine learning, remote sensing, atmospheric aerosols, bias correction, satellite data processing, environmental threat assessment.

Introduction

Atmospheric aerosols are a major component of environmental risk because they influence air quality, visibility, radiative transfer, and, indirectly, human health and climate processes. For this reason, aerosol optical depth at 550 nm (AOD550) is widely used as a compact indicator of columnar aerosol loading in satellite-based environmental monitoring, climate analysis, and air pollution assessment[1]. In the broader context of processing satellite data for threat assessments, aerosol retrieval is particularly relevant because elevated aerosol load may accompany or indicate hazardous conditions such as wildfire smoke transport, dust outbreaks, industrial pollution

[1] [1] Holben, B.N.; Eck, T.F.; Slutsker, I.; Tanré, D.; Buis, J.P.; Setzer, A.; Vermote, E.; Reagan, J.A.; Kaufman, Y.J.; Nakajima, T.; et al. AERONET—A Federated Instrument Network and Data Archive for Aerosol Characterization. *Remote Sens. Environ.* 1998, *66*, 1–16. https://doi.org/10.1016/S0034-4257(98)00031-5

Cite this Article as: Dawid JURCZYŃSKI Vol. 2026 (12) “Hybrid Physics–Machine Learning Retrieval of AOD550 from HYPSO Hyperspectral Imagery " Communications of International Proceedings, Vol. 2026 (12), Article ID 4724326, https://doi.org/10.5171/2026.4724326

episodes, and regional haze events. Reliable and timely AOD550 estimates are therefore valuable not only as atmospheric products, but also as analytical inputs for threat-oriented remote sensing workflows[2].

Satellite observations have enabled regional and global aerosol monitoring for decades, leading to the development of retrieval algorithms for sensors such as MODIS, VIIRS, and related systems[3]. However, aerosol retrieval remains a difficult inverse problem because the measured top-of-atmosphere signal depends not only on aerosol concentration, but also on surface reflectance, cloud contamination, viewing geometry, sensor calibration, and assumptions embedded in radiative transfer and aerosol models[4]. Even mature satellite products therefore retain systematic uncertainty and scene-dependent bias. Previous studies have shown that retrieval performance varies with surface type, aerosol regime, and AOD level, with low-AOD conditions often dominated by surface-related uncertainties and higher-AOD scenes becoming increasingly sensitive to aerosol model assumptions and multiple-scattering effects[5].

These limitations are important in applications requiring rapid interpretation of satellite data. Threat assessment workflows benefit from products that are not only physically meaningful, but also computationally efficient and available with minimal latency. This requirement motivates interest in compact retrieval strategies that can support near-real-time analysis without relying on heavy inversion frameworks or delayed auxiliary products. In this context, hyperspectral Earth observation offers a promising opportunity. Compared with conventional multispectral systems, hyperspectral sensors provide denser spectral sampling, which can improve scene characterization and potentially enhance the retrieval of aerosol-sensitive information. Recent and current missions such as PRISMA, EnMAP, PACE, and HYPSO illustrate the growing relevance of hyperspectral observations in remote sensing[6].

Among these platforms, HYPSO is particularly interesting as a small hyperspectral satellite system capable of delivering spectrally rich observations in an operationally flexible format. However, the existing HYPSO-related literature has mainly focused on mission design, hyperspectral image analysis, segmentation, multimodal methods, and onboard intelligence rather than on dedicated AOD550 retrieval for rapid practical use[7]. Thus, despite clear potential, the role of HYPSO in aerosol monitoring and threat-oriented atmospheric assessment remains insufficiently explored.

At the same time, machine learning has increasingly been used in aerosol research to merge products, correct model outputs, and post-process satellite retrievals[8]. These studies show that aerosol retrieval errors often contain structured, learnable components rather than purely random noise. Consequently, machine learning can be used not to replace physical retrieval entirely, but to correct its systematic residual bias. This hybrid strategy is attractive because it combines physical interpretability with flexible data-driven error reduction[9].

This study addresses that gap by proposing a hybrid framework for scene-level HYPSO-based AOD550 retrieval in the context of satellite-data processing for threat assessment. A physics-inspired baseline estimate is first derived from spectral and geometric scene information, and a supervised machine-learning model is then used to

---

[2] [5] Inness, A.; Ades, M.; Agustí-Panareda, A.; et al. The CAMS Reanalysis of Atmospheric Composition. *Atmos. Chem. Phys.* 2019, *19*, 3515–3556. https://doi.org/10.5194/acp-19-3515-2019

[3] [6] Levy, R.C.; Mattoo, S.; Munchak, L.A.; et al. The Collection 6 MODIS Aerosol Products over Land and Ocean. *Atmos. Meas. Tech.* 2013, *6*, 2989–3034. https://doi.org/10.5194/amt-6-2989-2013

[4] [2] Popp, T.; de Leeuw, G.; Martynenko, D.; et al. Aerosol Retrieval Experiments in the ESA Aerosol_cci Project. *Atmos. Meas. Tech. Discuss.* 2013, *6*, 2353–2401. https://doi.org/10.5194/amtd-6-2353-2013

[5] [4] Gumber, A.; Reid, J.S.; Holz, R.E.; et al. Assessment of Severe Aerosol Events from NASA MODIS and VIIRS Aerosol Products for Data Assimilation and Climate Continuity. *Atmos. Meas. Tech.* 2022. https://doi.org/10.5194/amt-2022-290

[6] Alvarez Justo, J.; Langer, D. D.; Berg, S.; et al. Hyperspectral Image Segmentation for Optimal Satellite Operations: In-orbit Deployment of 1D-CNN. *Preprints* 2024. https://doi.org/10.20944/preprints202411.0139.v1

[7] Penne, C.; Orlandic, M.; Garrett, J.; et al. Multimodal Analysis of Hyperspectral Images from HYPSO-1/2 Cube Satellites. *Proceedings of SPIE* 2025. https://doi.org/10.1117/12.3069723

[8] Gupta, P.; Remer, L. A.; Levy, R. C.; Mattoo, S. Validation of MODIS 3 km Land Aerosol Optical Depth from NASA's EOS Terra and Aqua Missions. *Atmospheric Measurement Techniques* 2018, 11, 3145–3159. https://doi.org/10.5194/amt-11-3145-2018

[9] Lipponen, A.; Reinvall, J.; Väisänen, A.; et al. Deep Learning Based Post-Process Correction of the Aerosol Parameters in the High-Resolution Sentinel-3 Level-2 Synergy Product. *Atmospheric Measurement Techniques* 2021. https://doi.org/10.5194/amt-2021-262

correct its residual bias relative to a reference aerosol product. In Leave-One-Scene-Out cross-validation over 33 scenes, the bias-correction stage reduced the out-of-fold MAE from 0.234 to 0.024 and the RMSE from 0.243 to 0.039, indicating a substantial reduction in retrieval error . These results suggest that HYPSO observations contain useful aerosol-related information and that a hybrid physics–machine learning workflow can support rapid AOD550 estimation as a building block for broader satellite-based threat assessment systems.

Materials and Methods

This study was designed as a hybrid framework for rapid scene-level AOD550 estimation from HYPSO hyperspectral satellite data. The workflow consisted of two main stages. First, a physics-inspired baseline AOD550 estimate was derived directly from HYPSO observations using spectral and geometric scene descriptors. Second, a supervised machine-learning model was used to estimate and correct the residual bias of that baseline retrieval relative to an external reference aerosol field. The purpose of this design was not to replace physically motivated retrieval with a fully data-driven predictor, but to preserve a physically interpretable first-order estimate and improve it through learned residual correction.

The analysis was conducted at the scene level rather than at the pixel level. Each HYPSO acquisition was represented by a compact set of aggregated descriptors extracted from calibrated hyperspectral products. These descriptors included reflectance-related spectral information, scene-summary statistics, and geometric variables such as solar zenith angle, sensor zenith angle, and relative azimuth angle. This choice followed the operational objective of the study, namely rapid aerosol estimation from compact hyperspectral observations without performing a full radiative-transfer inversion for each pixel.

To evaluate and correct the HYPSO-based retrieval, a reference AOD550 field was assigned to each scene using CAMS-based aerosol information. The reference data were spatially matched to the HYPSO footprint and temporally aligned with the closest available product time. In this study, the reference field was treated as an external benchmark and supervision target rather than as absolute ground truth, since assimilated aerosol products may also contain uncertainty.

The baseline retrieval was physics-inspired in the sense that it used radiometrically and geometrically meaningful scene descriptors expected to respond to aerosol loading. Conceptually, the approach assumed that aerosol variability modifies the magnitude and spectral shape of the measured signal in ways that can be partially captured by selected hyperspectral features. Instead of solving the full inverse problem explicitly, the method used a compact set of physically motivated summary variables to generate an initial AOD550-like proxy at scene level.

For each scene, a feature vector was then constructed from the baseline estimate, selected spectral descriptors, reflectance-derived summary statistics, and scene geometry variables. These features were used as inputs to the bias-correction model. The role of the machine-learning stage was to learn structured discrepancies between the baseline HYPSO estimate and the reference AOD550 field. In this way, the correction model acted as a post-processing layer rather than as an unconstrained black-box retrieval.

Model performance was assessed using Leave-One-Scene-Out cross-validation, which directly tested generalization to unseen scenes. The dataset comprised 33 scenes, and the evaluation metrics included mean absolute error (MAE), root mean square error (RMSE), and the coefficient of determination (R^2). The baseline retrieval achieved an out-of-fold MAE of 0.234 and RMSE of 0.243, whereas the bias-corrected retrieval reduced these values to 0.024 and 0.039, respectively . These results indicate that the correction stage substantially reduced scene-level absolute error, although the framework should still be regarded as a proof-of-concept scene-level method rather than a validated pixel-wise operational aerosol product.

Results and Discussion

The proposed hybrid framework produced a substantial improvement in scene-level agreement between the HYPSO-based retrieval and the reference AOD550 field. Before bias correction, the baseline retrieval exhibited a pronounced systematic deviation from the reference data. This behaviour is visible in the scatter plot shown in Figure 1, where the predicted values remain consistently displaced from the 1:1 line and the overall relationship is dominated by overestimation.

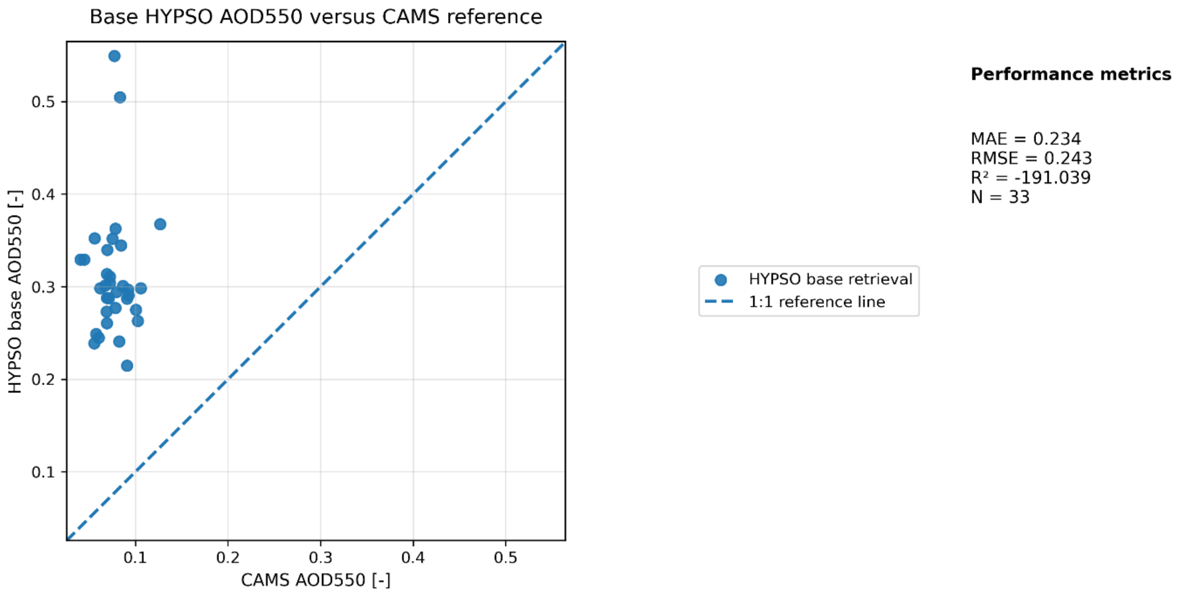

Figure 1 Baseline HYPSO-derived scene-level AOD550 estimates versus reference CAMS AOD550 values. Source: own elaboration.

The corresponding Leave-One-Scene-Out cross-validation metrics confirm this pattern: the baseline retrieval yielded an out-of-fold MAE of 0.234 and an RMSE of 0.243 for 33 analysed scenes, with a strongly negative R^2 . These results indicate that the initial HYPSO-derived proxy, when used alone, did not provide an adequate quantitative approximation of the reference field. Nevertheless, the existence of a stable and directional error pattern suggested that the baseline estimate retained physically meaningful aerosol-related information that could be exploited by the correction stage.

After machine-learning-based bias correction, the scene-level agreement improved markedly. As shown in Figure 2, the corrected estimates cluster much more closely around the 1:1 line, and the dispersion relative to the reference field is substantially reduced.

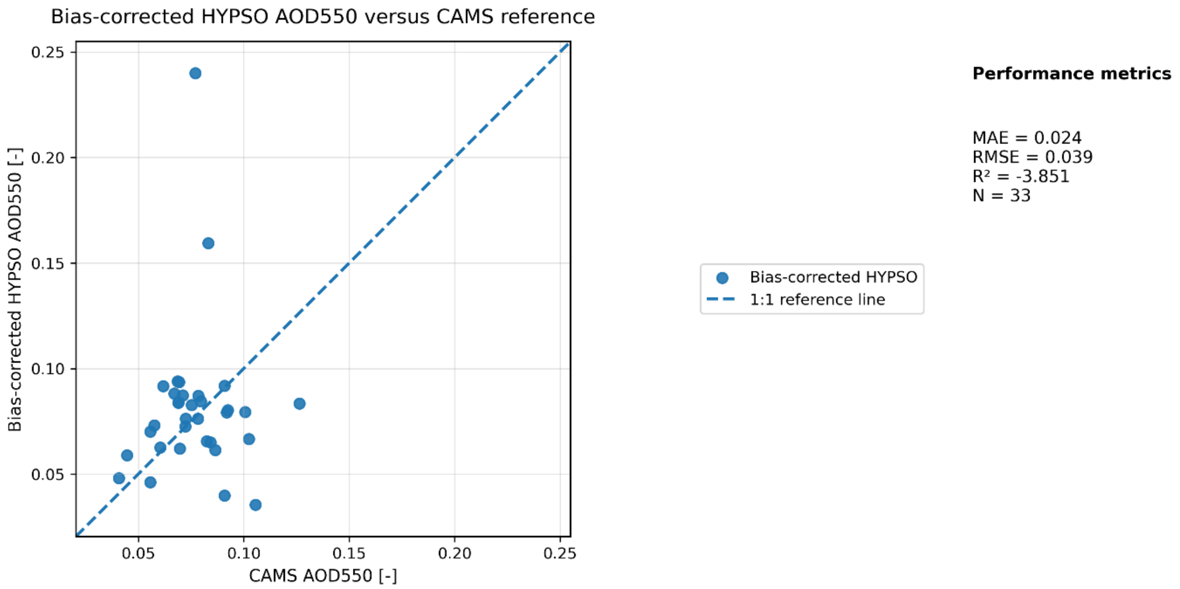

Figure 2 Bias-corrected HYPSO-derived scene-level AOD550 estimates versus reference CAMS AOD550 values. Source: own elaboration.

This visual improvement is supported by the numerical evaluation, which shows that the out-of-fold MAE decreased from 0.234 to 0.024 and the RMSE from 0.243 to 0.039 . In practical terms, these reductions indicate that the correction model was able to capture and remove a large fraction of the scene-dependent systematic bias present in the baseline retrieval. The magnitude of the reduction in RMSE is particularly relevant, because it suggests improvement not only in the average discrepancy but also in the larger deviations that dominated the raw retrieval.

The redistribution of residuals is illustrated in Figure 3.

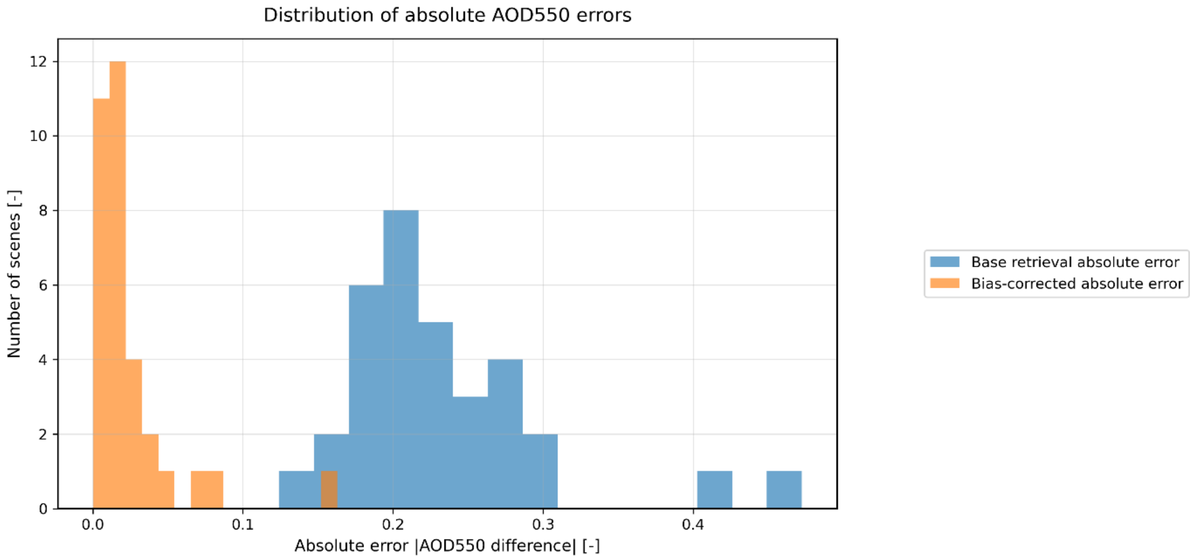

Figure 3 Residual error distributions for the baseline and bias-corrected scene-level AOD550 retrievals. Source: own elaboration.

The baseline residuals are broad and clearly shifted away from zero, which is consistent with a persistent positive bias and substantial error spread. In contrast, the corrected residual distribution is much narrower and centered substantially closer to zero. This change indicates that the correction stage did not merely compress the prediction range, but rather learned a coherent residual structure across the dataset. Such behaviour is desirable in a post-processing framework, because it implies that the machine-learning model acted on a repeatable component of the error rather than on random fluctuations alone.

The per-scene behaviour provides further evidence of model consistency. As shown in Figure 4, the correction stage reduced the error for all analysed scenes.

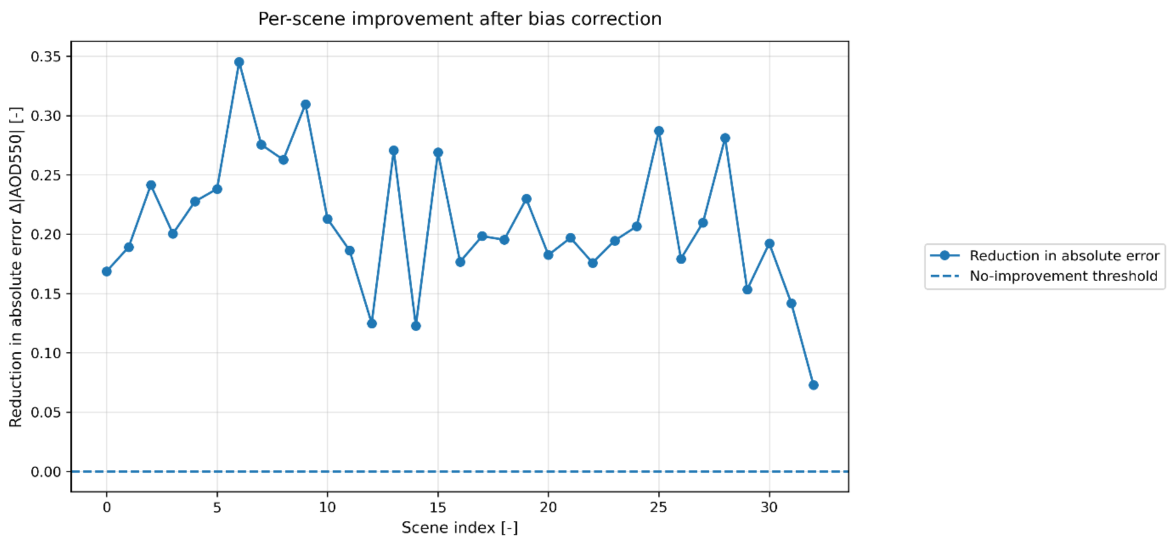

Figure 4 Per-scene reduction in absolute error after machine-learning bias correction of the baseline HYPSO-derived AOD550 retrieval. Source: own elaboration.

This result is important because it excludes a scenario in which the average improvement is driven only by a subset of samples while other scenes deteriorate. Instead, the improvement was uniform across the full set of 33 scenes, which supports the interpretation that the baseline HYPSO retrieval contains a stable aerosol-related signal whose systematic bias can be learned at scene level. From a methodological standpoint, this is one of the central findings of the study: the first-stage retrieval is not accurate enough to be used as a final product, but it is sufficiently structured to serve as a meaningful basis for supervised correction.

The feature-importance analysis shown in Figure 5 provides additional insight into the mechanism of improvement.

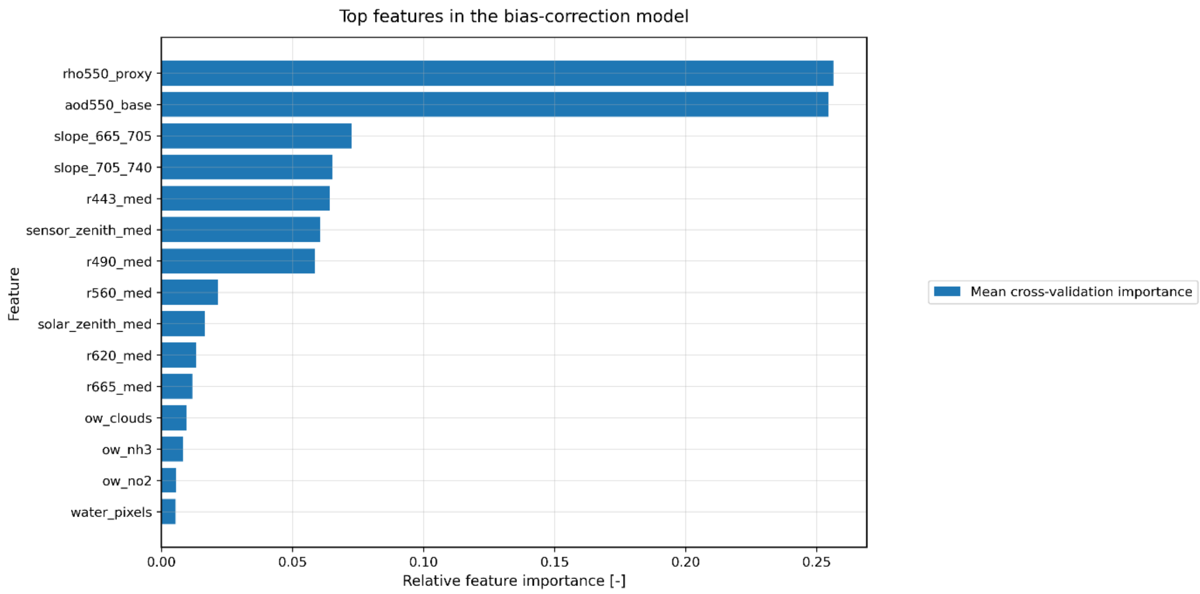

Figure 5 Feature importance ranking for the machine-learning bias-correction model applied to scene-level HYPSO-based AOD550 estimates. Source: own elaboration.

The dominant predictor was the baseline HYPSO-derived estimate itself, followed by selected spectral and scene-summary descriptors. This ranking is informative because it indicates that the corrected solution remained anchored in the physics-inspired first-stage retrieval. The machine-learning model therefore behaved primarily as a residual-correction layer rather than as a disconnected empirical regressor. In other words, the improvement was not driven solely by auxiliary correlations, but by the interaction between the original retrieval and a limited number of additional scene descriptors that helped explain its systematic departures from the reference field.

A compact statistical synthesis is presented in Figure 6, while the direct contrast between the baseline and corrected stages is shown in Figure 7.

Bias correction summary

Cross-validation mode: Leave-One-Scene-Out
Number of scenes: 33

Base MAE: 0.234
Corrected MAE: 0.024

Base RMSE: 0.243
Corrected RMSE: 0.039

Base R^2: -191.039
Corrected R^2: -3.851

Full-fit corrected MAE: 0.011
Full-fit corrected RMSE: 0.020
Full-fit corrected R^2: -0.257

Figure 6 Cross-validation performance metrics for the baseline and bias-corrected scene-level AOD550 retrievals. Source: own elaboration.

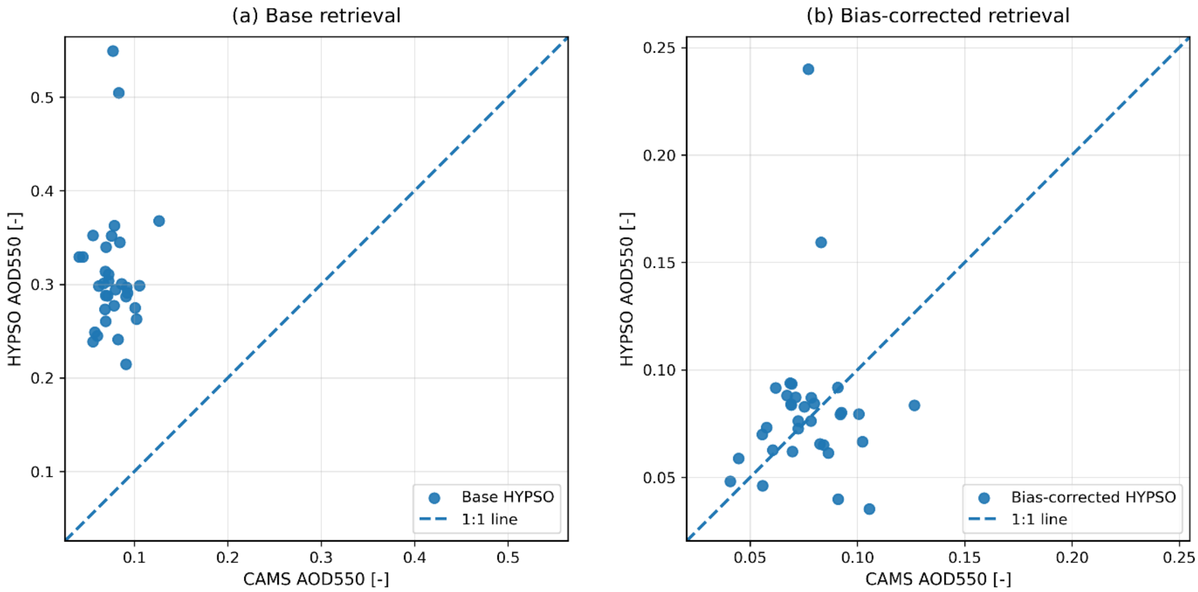

Figure 7 Side-by-side comparison of baseline and bias-corrected HYPSO-derived scene-level AOD550 estimates against reference CAMS AOD550 values. Source: own elaboration.

Both figures are fully consistent with the uploaded cross-validation metrics . Taken together, the figures indicate that the proposed workflow materially altered the relationship between the HYPSO-derived estimate and the reference AOD550 field, rather than merely smoothing the output or truncating outliers.

These findings should, however, be interpreted within the limits of the experimental design. First, the analysis was performed at scene level rather than at pixel level, and the present framework should therefore not be interpreted as a validated operational aerosol retrieval product. Second, the reference field was derived from CAMS-based aerosol information and thus represents an external benchmark rather than absolute ground truth. Third, the sample size remained limited to 33 scenes. These constraints are relevant for interpretation of the negative R^2 values, which remained below zero even after correction . Accordingly, the most robust conclusion is not that the proposed framework fully solves the AOD550 retrieval problem, but that it achieves a strong and systematic reduction of absolute error relative to the selected reference product.

From a broader perspective, the results support the central premise of the study. The baseline HYPSO retrieval alone is physically informative but quantitatively biased. The bias is not random; rather, it is sufficiently structured to be learned through supervised correction using compact scene-level descriptors. This makes the hybrid formulation particularly suitable for rapid-processing applications, where a lightweight and interpretable first estimate can be generated quickly and then refined using a data-driven residual model. In the context of satellite-data processing for threat assessment, this is a useful property, because scene-level aerosol indicators do not necessarily need to achieve the maturity of full atmospheric products in order to support screening, prioritization, or early interpretation workflows.

Overall, the results indicate that HYPSO observations contain usable aerosol-related information at scene level and that a hybrid physics–machine learning approach can substantially improve their consistency with an external reference AOD550 field. The method should therefore be understood as a proof-of-concept framework for rapid scene-level aerosol estimation rather than as a final retrieval product, but the magnitude and consistency of the error reduction suggest that the proposed strategy is a promising basis for further development.

Conclusions

This study presented a hybrid framework for rapid scene-level AOD550 estimation from HYPSO hyperspectral satellite data, combining a physics-inspired baseline retrieval with a machine-learning-based bias-correction stage. The results show that the baseline HYPSO-derived estimate contained a usable aerosol-related signal, but that this signal was affected by a substantial and systematic scene-dependent bias. The subsequent correction stage reduced the out-of-fold MAE from 0.234 to 0.024 and the RMSE from 0.243 to 0.039 in Leave-One-Scene-Out cross-validation over 33 scenes . These reductions indicate that the residual error structure was sufficiently stable to be learned and corrected using compact scene-level descriptors.

The findings support the main premise of the work, namely that a physically interpretable first-order aerosol estimate derived from HYPSO data can be meaningfully improved through supervised post-processing. In this sense, the proposed method should not be interpreted as a replacement for established aerosol retrieval products, but rather as a proof-of-concept framework showing that HYPSO observations can contribute to rapid scene-level aerosol assessment. The dominant contribution of the baseline retrieval and selected spectral descriptors in the feature-importance analysis further suggests that the improvement was not purely empirical, but remained anchored in physically meaningful scene information.

At the same time, the present results should be interpreted conservatively. The study was limited to scene-level retrieval rather than pixel-level inversion, the sample comprised 33 scenes, and the supervision target was an external CAMS-based aerosol reference field rather than direct ground truth. In addition, the negative R^2 values indicate that the framework should not yet be regarded as a fully predictive statistical model of scene-level AOD550 . The most robust conclusion is therefore not that the aerosol retrieval problem has been solved, but that the proposed hybrid approach substantially reduced systematic error relative to the selected reference product.

Within the broader context of processing satellite data for threat assessments, these results are nevertheless relevant. Rapid retrieval of aerosol-related indicators may support the early interpretation of environmentally hazardous conditions, including haze, smoke transport, and dust-related events. From this perspective, the proposed framework offers a practical basis for further development of lightweight hyperspectral workflows intended for rapid-response analysis.

Future work should focus on extending the dataset, improving the physical formulation of the baseline retrieval, testing alternative bias-correction models, and evaluating the method against independent ground-based observations such as AERONET. Further development should also examine whether the approach can be generalized beyond scene-level estimation toward more spatially resolved retrieval products. Such steps will be necessary to determine the operational value of HYPSO-based aerosol estimation in broader monitoring and threat-assessment systems.

References

[1] Holben, B.N.; Eck, T.F.; Slutsker, I.; Tanré, D.; Buis, J.P.; Setzer, A.; Vermote, E.; Reagan, J.A.; Kaufman, Y.J.; Nakajima, T.; et al. AERONET—A Federated Instrument Network and Data Archive for Aerosol Characterization. Remote Sens. Environ. 1998, 66, 1–16. https://doi.org/10.1016/S0034-4257(98)00031-5

[5] Inness, A.; Ades, M.; Agustí-Panareda, A.; et al. The CAMS Reanalysis of Atmospheric Composition. Atmos. Chem. Phys. 2019, 19, 3515–3556. https://doi.org/10.5194/acp-19-3515-2019

[6] Levy, R.C.; Mattoo, S.; Munchak, L.A.; et al. The Collection 6 MODIS Aerosol Products over Land and Ocean. Atmos. Meas. Tech. 2013, 6, 2989–3034. https://doi.org/10.5194/amt-6-2989-2013

[2] Popp, T.; de Leeuw, G.; Martynenko, D.; et al. Aerosol Retrieval Experiments in the ESA Aerosol_cci Project. Atmos. Meas. Tech. Discuss. 2013, 6, 2353–2401. https://doi.org/10.5194/amtd-6-2353-2013

[4] Gumber, A.; Reid, J.S.; Holz, R.E.; et al. Assessment of Severe Aerosol Events from NASA MODIS and VIIRS Aerosol Products for Data Assimilation and Climate Continuity. Atmos. Meas. Tech. 2022. https://doi.org/10.5194/amt-2022-290

Alvarez Justo, J.; Langer, D. D.; Berg, S.; et al. Hyperspectral Image Segmentation for Optimal Satellite Operations: In-orbit Deployment of 1D-CNN. Preprints 2024. https://doi.org/10.20944/preprints202411.0139.v1

Penne, C.; Orlandic, M.; Garrett, J.; et al. Multimodal Analysis of Hyperspectral Images from HYPSO-1/2 Cube Satellites. Proceedings of SPIE 2025. https://doi.org/10.1117/12.3069723

Gupta, P.; Remer, L. A.; Levy, R. C.; Mattoo, S. Validation of MODIS 3 km Land Aerosol Optical Depth from NASA's EOS Terra and Aqua Missions. Atmospheric Measurement Techniques 2018, 11, 3145–3159. https://doi.org/10.5194/amt-11-3145-2018

Lipponen, A.; Reinvall, J.; Väisänen, A.; et al. Deep Learning Based Post-Process Correction of the Aerosol Parameters in the High-Resolution Sentinel-3 Level-2 Synergy Product. Atmospheric Measurement Techniques 2021. https://doi.org/10.5194/amt-2021-262