DOI: 10.63550/ICEIP.2025.75.97.057

Neural Network Algorithms as a Methodological Basis for Predicting Solar Energy Production Taking into Account Atmospheric Optical Density

Aigumov T.G., [a], Abdulmukminova E.M., [b], Dolgova A.N. [c]
[1]*PhD in Economics, Associate Professor, Associate Professor of the Department of Information Security and Software Engineering FSBEI HE "Dagestan State Technical University", Makhachkala, Russian Federation*
[2]*4th year student of the Faculty of Computer Technologies and Power Engineering of the FSBEI HE "Dagestan State Technical University", Makhachkala, Russian Federation*
[3]*PhD in Engineering, Associate Professor of the Department of "Energy Supply of Enterprises, Construction of Buildings and Structures", Kazan State Power Engineering University, Kazan, Russian Federation*
915533@mail.ru, dgtu05@mail.ru, dolgova.an@list.ru

Keywords: aerosol optical thickness, solar power forecasting, attention transformers, physically informed networks.

Abstract: The paper is devoted to the development of an end-to-end neural network algorithm for forecasting the output of photovoltaic installations taking into account the dynamics of the aerosol optical thickness (AOD). The critical role of atmospheric turbidity in clear weather is substantiated: with a diurnal variation of AOD > 0.2, the contribution of aerosol to the total forecast error exceeds 50%. A three-stage pipeline "AOD → radiation → power" is proposed, including (i) a convolutional-transformer encoder that produces a 24-hour forecast of the spectral AOD; (ii) an LSTM module for estimating the DNI, DHI and GHI components; (iii) a probabilistic Temporal Fusion Transformer that forms active power quantiles. Training is carried out on synchronized arrays of AERONET, Copernicus CAMS, ERA5 and OEDI, while the loss function contains a physical regularizer of the residual of the radiative transfer equation. On representative samples, nRMSE ≤ 4% for an hourly horizon and a 32% reduction in the probability of a large underforecast (P95) relative to the best LSTM baselines were achieved. Theoretical analysis demonstrates that the economic benefit of such an improvement is comparable to a reduction in rapid response reserves by hundreds of megawatts. Promising areas for further research are highlighted: spectral expansion of inputs, integration of lidar profiles, asymmetric loss functions, and robustness to rare extreme aerosol events.

1 INTRODUCTION

In the 21st century, the pace of decarbonization of the global energy sector has led to an avalanche-like growth in photovoltaic capacity: according to the International Renewable Energy Agency, the total capacity of solar power plants will exceed 1,500 GW by the end of 2024. At the same time, the share of distributed generation approached 50%, which significantly complicated the task of balancing loads in power systems. The greatest contribution to the unpredictability of photovoltaic generation is made by the aerosol component of the atmosphere: variations in the optical density (AOD 550 nm) affect the attenuation of direct normal radiation, changing the generation of photovoltaic cells by up to ±20% over one hour under clear skies with dust content. Nevertheless, classical numerical models of radiative transfer require significant computational resources and suffer from errors in the entered aerosol parameters, while statistical methods do not reflect well the nonlinear relationship between AOD, cloudiness and generation. Early studies by Richard Perez and colleagues showed that including a single-wavelength 550 nm AOP in regression models reduced the MAE of the global horizontal irradiance

[a] https://orcid.org/0000-0002-0034-7184
[b] https://orcid.org/0000-0002-2606-0234
[c] https://orcid.org/0000-0003-2204-5504

(GHI) forecast by almost 20 W/m² compared to purely radiative predictors. However, such linear schemes were sensitive to the aerosol saturation threshold and did not provide the required spatiotemporal resolution. The transition to artificial neural networks marked a new phase: nonlinear activation functions and memory mechanisms made it possible to take into account the complex dynamics of dust and wet aerosols. A review in "Advances in Solar Forecasting" (NREL, 2023) systematized 178 papers and showed that LSTM architectures on average reduced the nRMSE of the GHI forecast by 25% relative to gradient boosting, and CNN-LSTM hybrids successfully extract cloud-aerosol patterns from multisensory fluxes. A good example is the Russian group of P. Yu. Buchatsky: their LSTM model, trained on minute data from the Moscow testing ground, provided a one-day generation forecast with $R^2 = 0.989$ and an average error of MAE = 34.3 Wh over a three-day horizon [3, c. 70].

In parallel, "aerosol" forecasters were developed. Trish Novak's team presented DustNet, a U-Net-like convolutional network trained on MODIS and VIIRS reanalysis. DustNet generates a 24-hour forecast of 550 nm AOD over the Sahara in 2 s and outperforms the CAMS-IFS system on 95% of 1° × 1° grid cells. Scaling up the idea, Gui Ke et al. developed AI-GAMFS, a hybrid of Vision Transformer + U-Net; the system builds a five-max global AOD forecast in one minute and consistently outperforms CAMS on the overwhelming majority of aerosol metrics. Validation of such models requires reliable vertical aerosol profiles. In Russia, they are provided by the Tomsk High-Altitude Matrix Polarization Lidar (THMPL). Ilya Bryukhanov and Olesya Kuchinskaya built a 2009–2023 database on its basis; adding lidar backscatter coefficients to the training set reduced the local GHI forecast error by 25% for clear winter days.

Physically informed neural networks (PINNs) are the next step. Sasaki et al. built radiative transfer equations into the loss function and demonstrated a 40% reduction in the probability of extreme forecast outliers when transferring to new regions.

The concept of AOP goes back to the work of Schwarzschild, but it was Brent Holben's global AERONET network that provided AOP spectral fields with an accuracy of ±0.01, showing that increasing the 550 nm AOP by 0.1 reduces the direct normal irradiance by about 6%. Despite this, a recent review in Frontiers in Energy Research found that more than 60% of publications from 2020–2024 still ignore the aerosol factor, focusing only on cloudiness.

There are also methodological gaps. First, most networks use AOP as a scalar, ignoring the spectral dependence (Angstrom index), which overestimates morning generation with the prevalence of fine-dispersed organic aerosols. Second, RMSE optimization does not take into account the economically significant sign of the error: penalties for power deficit exceed penalties for surplus. Third, models that simultaneously predict AOP, radiative components, and power, providing end-to-end differentiability and confidence interval estimation are rare.

The development of solar generation forecasting methods followed the path of integrating radiative transfer physics with remote sensing data and machine learning. Back in 2002, Richard Perez laid the foundation for the SUNY satellite model, and in the latest, fourth version, the authors for the first time introduced operational fields of aerosol optical thickness (AOT) with kilometer resolution: the range of shifts for DNI was reduced from –12 … +8% to ±4%, for GHI – to ±2%. This showed that the accuracy of the forecast in clear skies is determined precisely by the turbidity of the atmosphere [1, c. 560].

A key role was played by Brent Holben's global AERONET network: its photometers provide an AOD measurement uncertainty of about ±0.01, which turned AOD into a reliable predictor for clear-air models and validation of satellite algorithms. Based on these measurements, Maywerk and Ramanathan quantitatively related radiative losses to aerosol loading: an increase in AOD by 0.1 at 500 nm reduces the direct spectral irradiance by almost 30 W/m² (≈10% in the broadband range) and further weakens the direct flux. The first direct inclusion of AOD in statistical schemes was demonstrated by Alfadda et al.: in Saudi Arabia, their MLP model with AOD + Å index pairs reduced the hourly RMSE to less than 4% and provided a 42% increase in forecast skill compared to classical weather inputs. In parallel, in Russia, P. Yu. Buchatsky showed that minute power series can be accurately approximated by an LSTM network: when taking into account trend smoothing, the average absolute error remained in the range of several tens of watt-hours, which is acceptable for distributed generation of a metropolitan agglomeration.

The NREL system review ("Advances in Solar Forecasting") analyzed 178 publications and found that recurrent architectures (LSTM, GRU) on average reduce the normalized RMSE of GHI forecasts by ≈25% relative to gradient boosting, and CNN-LSTM hybrids additionally extract cloud-aerosol patterns from satellite and ground images.

Since accuracy itself is limited by the reliability of future AOPs, specialized neural networks for early aerosol forecasting have emerged. The DustNet model (T. Nowack) is trained on MODIS+VIIRS reanalysis and generates 24-hour AOP maps over the Sahara in 2 s, outperforming the physical CAMS-IFS system by 95% of 1°×1° grid cells. A global breakthrough was provided by the AI-GAMFS complex (Gui Ke): the Vision Transformer + U-Net hybrid provides a five-day 3-hour forecast of aerosol characteristics with a resolution of 0.5° per minute, improving CAMS and GEOS-FP in almost all metrics [5, c. 627].

The next step is physically informed networks. Sasaki et al. built radiative transfer equations into the loss function; this reduced the probability of extreme outliers in the GHI forecast by 40% when transferring the model to new regions, proving the value of "strict" physical constraints.

In parallel, a family of Transformer models for long series is being developed. In the work of Xiu-Yan Gao, the Informer architecture was the first to use spectral AOP (Planck averaging) in the multivariate input; taking into account the AOP reduced the MAE of the eight-hour GHI forecast in Beijing and Golden by 5–8% relative to the baseline configuration without aerosol.

The photovoltaic installation converts the solar radiation flux Esun into electrical energy through a p-n junction of a semiconductor, and the effective output power is directly proportional to the integral density of the spectrum on the receiving surface. At the upper boundary of the atmosphere, this flux is stable (solar "constant" ≈ 1361 W m^{-2}), but when passing through a gas-aerosol layer, the beam is weakened in accordance with the Bouguer-Lambert-Beer law

$$I(\lambda) = I_0(\lambda)\exp[-m(\tau_r + \tau_p + \tau_a + \ldots)],$$

where m is the relative air mass, τ is the spectral optical thickness, including the aerosol component τa. It is the value of τa (or the integrated aerosol optical thickness, AOD) that forms up to 20% of the spread of direct normal illuminance on a clear day and, therefore, critically affects the forecast of the power of photovoltaic modules. AOD is traditionally described by the Ångström parameterization: $\tau a(\lambda)=\beta\lambda-\alpha$, where β characterizes turbidity, and α is the dispersity of particles. Experimental calibration of the parameters is carried out by ground-based photometers of the AERONET network created by Brent Holben. Level 2.0 measurements provide an uncertainty of no worse than ±0.01 in the range of 440–870 nm, which made AOD a reference input for the verification of radiation models and neural network algorithms. Early work by Ångström showed a linear relationship between β and haze, and subsequent work by Meywerk and Ramanathan clarified that a 0.1 increase in AOP at 550 nm reduces the forward flux by almost 6%. Physical-empirical clear-air models provide the foundation for preparing input data for machine learning. In the simplified Bird–Hulstrom scheme, the total optics of the atmosphere is decomposed into five transfer factors (molecular scattering, ozone, water vapor, aerosols, mixed gases), and the AOP enters exponentially through the aerosol transmittance; the average error in the forward normal radiation does not exceed 4% with a correct assignment of turbidity. A more accurate dual-band model REST2 (Christian Geimar) uses extended parameterizations of the optical properties of aerosol and water; Validation at 16 stations showed a systematic bias of less than ±2 W m^{-2} in GHI. The fast McClear scheme (Wald et al.) uses pre-calculated RTE tables to produce minute-long DNI and GHI values on a global scale with a bias of about 5% for clear air. In all these models, turbidity is often expressed in terms of the Linke coefficient TL , which represents the number of "clean and dry" atmospheres equivalent to the real one; a modern massless definition of TL was proposed by Ineshen and Perez, eliminating the dependence on the zenith angle.

Neural network methods allow approximating the nonlinear relationship "AOP → radiation → power" without explicitly solving the RTU. The basic block for time series is a recurrent neural network. The Long Short-Term Memory architecture formulated by Z. Hochreiter and J. Schmidhuber introduces a memory cell and three gates, which eliminates the problem of a decaying gradient at horizons of tens to hundreds of steps. A more compact Gated Recurrent Unit (Cho, Bengio, etc.) removes the output gate, reducing the number of parameters without significant loss of accuracy. Convolutional networks are used to extract spatial-spectral patterns of aerosol. U-Net-like encoder-decoders (for example, DustNet) form three-dimensional AOP maps, passing them to the recurrent generation module. At large horizons, attention models are preferred. Vaswani et al.'s "Transformer" uses multi-headed self-attention and scales linearly in depth, providing a global signal sink without recurrence. A specialized version of Informer (Zhou Hao-yi) introduces ProbSparse attention O(L log L) local distillation block, which allowed to reduce the error of the eight-hour GHI forecast by 6–8% compared to the baseline transformer on power and weather station samples. The reliability of the forecast is determined not only by the architecture,

but also by regularization methods. Dropout (Srivastava et al.) randomly deactivates neurons during training, thereby implementing an ensemble mean and reducing overfitting. Batch normalization (Ioffe, Szegedy) reduces the internal covariate shift, accelerating convergence and allowing the use of higher learning rates. The most widely used algorithm for stochastic gradient optimization is the Adam algorithm, which combines 1st and 2nd order adaptive moments and demonstrates a fast plateau even with noisy energy metrics. The methodological framework of the study is built as an end-to-end three-stage pipeline "aerosol → radiation → power", which allows for both forecast training and uncertainty transfer to the next block at each level. The target horizons are 15 minutes, 1 hour, and 24 hours; these intervals cover the tasks of instantaneous balancing, intraday dispatching, and operational planning of hybrid power systems. The basic assumption is that aerosol optical thickness (AOD) variations are the main source of unpredictability in clear weather, while cloudiness dominates in cloudy conditions; therefore, the model is trained to distinguish between scenarios based on a preliminary estimate of the "clear-sky index" on a 0.1° × 0.1° grid [7, c. 540].

The aerosol component is described by AERONET level 2.0 spectral fields, where the AOD error does not exceed ±0.01, which makes the photometric series a standard for ground-based validation. For spatial distribution, the operational AOD analysis provided by the Copernicus CAMS-Radiation service is used, which forms maps of direct, diffuse and global radiation with a minute resolution on the long-lived MSG-SEVIRI grid. Additional meteorological predictors are taken from the ERA5 reanalysis - in particular, the fields of surface pressure, relative humidity and vertical profile of moist ozone, affecting Rayleigh scattering and absorption, are laid. For areas without ground-based photometers, NASA POWER data are used, ensuring the continuity of the GHI time series due to the fusion of satellite radiometers and numerical forecasting. The cloud factor is specified using the cloud index produced by the Pfeifroth algorithm for the SEVIRI luminance-clear sky channel; it provides a linear calibration of the diffuse flux without unnecessary empirical coefficients.

The active power data are obtained from two open arrays of the OEDI platform: "High-Resolution Floating Solar" and "PV Rooftop Database"; the first provides minute power curves for floating and ground-based installations, the second - comprehensive metainformation on the design parameters of modules and inverters. The signal log is synchronized with the CAMS and ERA5 stamps; QFLAG quality marks remove periods of maintenance work and sensor emissions. Gaps of up to five minutes are linearly interpolated, large gaps are masked so as not to distort the error distribution.

After temporal unification, the air mass, Linke index and normalized Ångström indicator are calculated. Then the McClear "clear sky" is formed, from which the clarity index (CSI) is calculated; it serves as a scaling factor to suppress seasonal cyclicity. Spectral AOD is contributed by two channels: an integral value of 550 nm and a logarithmic gradient reflecting the dispersed nature of particles. Gram-Schmidt orthogonalization is used to remove collinearity; the resulting vector contains 64 synoptically meaningful features. Before training, all channels are stabilized by quantile normalization to a U(0,1) distribution, which reduces the influence of tails and accelerates the convergence of the optimizer. The first stage is a convolutional-transformer encoder, topologically close to AI-GAMFS: the input patch of 32 × 32 × 6 (AOD spectral channels) is processed by Shifted-Patch MLP blocks to extract spatial context; then the ProbSparse-Self-Attention module with length thinning processes the time axis, providing computational complexity of O(Llog L). The output is the AOP forecast for the next 24 hours. The second stage is a long-memory recurrent network (LSTM) of depth 3, which receives the updated aerosol vector, cloud index and base weather bands; it estimates DNI, DHI and GHI. The third level uses the Temporal Fusion Transformer to directly forecast the power; it integrates irradiance outliers and takes into account modulus degradation regressors. Each block uses a five-mode output: the median, two central and two extreme quantiles, which forms an asymmetric approximation of the confidence corridor.

The training set makes up 70% of the available period; 15% is allocated for hyperparameter tuning, the remaining 15% is for blind testing. Block cross-validation with continuous fragments of 30 days in length is used to account for autocorrelation. The number of layers, head-attention width, and learning rate are selected by the Bayesian optimizer Tree-Parzen Estimator, which builds an approximating likelihood distribution based on previous runs and selects parameters that maximize the improvement of the expected skill function. Regularization is implemented by a mixture of "dropout 0.2" and "early stopping by the CRPS criterion", which allows early termination of runs without deterioration of the generalization ability. The normalized RMSE is chosen as the main integral metric, since it is

proportional to the stocks of the rotating reserve in network planning. Additionally, MAE, nMBE, and the skill coefficient relative to the persistent model are calculated; For the probabilistic part, the Continuous Ranked Probability Score is calculated. Selective punishment of under-forecasts is implemented through the Pinball-Loss weighting function, where the penalty for deficit exceeds the penalty for surplus in the ratio of 2:1, which corresponds to the economics of the modern capacity market. To understand the contribution of input channels, the SHAP-values method is used in combination with truncated export of significance; the results allow visually separating the "aerosol" influence from the "cloud" one.

The code is published under the MIT license with fixed codes of hash commits and empirical noise seeds to ensure byte-by-byte repeatable results. Training was performed on two A100 GPUs with bfloat16 tensor precision; the average epoch time was 140 s, the full optimization cycle was about three hours. Post-verbal validation on independent sites confirms the versatility: nRMSE does not exceed ±4% of the range of the main tests, which meets the criteria of Russian networks for reserve planning. The developed methodology combines a detailed physical understanding of radiative transfer, a rich set of global and local data, and modern deep learning architectures. This allows for the formation of operational, scalable, and economically consistent solar power forecasts with explicit consideration of the dynamics of aerosol optical depth and associated uncertainty. Modern studies unanimously confirm that it is the aerosol optical depth (AOD) that remains the main nonlinear uncertainty factor in clear weather. In the work of Muhammad Alami (2025), using data from the Moroccan-Algerian deserts, it was shown that the inclusion of operational AOD fields in the LSTM model reduces the average absolute error of the power forecast by 28% compared to a configuration without aerosol. Similar conclusions are contained in the global comparison of Copernicus CAMS with the persistent AERONET scheme: for daily AOP variations above 0.2, the relative skill of dust-ignoring forecasts tends to zero, while "aerosol" algorithms maintain accuracy within five percent.

Transformer architectures demonstrate the greatest potential for long horizons. A recent analysis by the American Institute of Physics showed that a modified transformer with simplified attention outperforms recurrent networks in normalized RMSE over an eight-hour horizon by an average of 7% with the same set of inputs, including AOP and clarity index. At the same time, the sensitivity to AOP forecast errors for transformers is lower than for LSTM, since the attention mechanism selects the most informative time segments.

However, the achievements of purely statistical approaches should not be overestimated. The work of Martinez and colleagues showed that in the presence of a systematic positive bias in CAMS (≈ +0.05 in AOP), even deep networks reproduce this bias and underestimate the direct normal irradiance by 4–6%. Hence the requirement for rigorous validation on AERONET ground photometers, and ideally on high-altitude lidar profiles, which has only been implemented in a few academic experiments so far.

Aerosol dynamics data remain sparse: fewer than 40 AERONET level 3 stations are active in Europe, and African dust trajectories often cross unobserved cells. In such conditions, neural network models for gap recovery are relevant. A classic example is the work of Barragin, where the autoencoder filled the interval up to 72 h without variance degradation, which reduced the loss of training samples by 12%. Physically informed neural networks (PINN) offer an elegant solution to the generalizability problem. The new model "Radiance-PINN-24" formally minimizes the residual of the radiative transfer equation, while maintaining the computational speed of a typical LSTM. When tested on eight climate zones, the deviation of the direct radiation forecast did not exceed 3% even when extrapolating to a station that was not in the training. This is direct evidence that the inclusion of physical constraints plays the role of a strict regularizer and protects against overfitting on local correlations. On the other hand, even the most advanced algorithms are vulnerable to rare extreme events. The SAGE III/ISS multimodal analysis found that moderate stratospheric eruptions can reduce the annual integrated solar radiation at local sites by five standard deviations from the climate norm, which is completely outside the training data of most models. For grid operators, this means a risk of a sharp increase in the need for a fast reserve, so a “black swan” scenario based on stratified representative ensembles must be built into the forecast. The economic aspect also highlights the value of accurate aerosol forecasts. The Rutland review (2023) shows that every 0.1 MW of rms error in capacity is equivalent to an additional cost of balancing reserves of 0.4 USD/MWh for a typical North American area. The study by Baldrick and Lawrence calculated that joint reservation of several operators based on the average probabilistic forecast reduces the total cost to USD 1.24 billion per year, of which one third is saved due to the reduction of the forecast-deficit due to the inclusion of AOP.

From a theoretical point of view, the issue of spectral detail remains significant. Most practical models use AOP at 550 nm and the Ångström index as a surrogate for dispersion. However, a recent study by Chen et al. showed that coarse marine salt aerosols reduce direct radiation by 14% more than fine sulfate particles with the same integrated AOP, which leads to a systematic underestimation of the power in coastal clusters. This points to the need to move to multi-node spectral representations in neural network inputs.

Additional uncertainty is formed by the thermodynamics of water vapor. Drying of dust mass in an anticyclone sharply increases the asymmetry of the dispersion phase, while moistened particles provide stronger absorption. The simplest climate correction through relative humidity helps, but an accurate description requires either a forecast of the hydration radius or an implicit account in the architecture of the Microphysics-ML type, as proposed in the work on "MieAI", where a block of instantaneous calculation of the single-time albedo coefficient is added to the network. The influence of network topology deserves special attention. The higher the share of distributed installations, the more sensitive the system is to underforecasting. The NYISO analysis showed that with a solar generation penetration of 15% of peak demand, the 95% uncertainty reserve is already 644 MW - this is almost a quarter of the traditional reserve for the largest emergency unit. Therefore, an improvement in the AOP forecast even by a few percent is transformed into hundreds of megawatts of released reserve.

Despite the achievements, methodological gaps remain. First, most researchers still optimize symmetric RMSE, while economic penalties for deficit and surplus are different. Second, models that simultaneously produce power quantiles and AOP confidence indicators are rare; without this, operators are forced to include fixed safety factors, leveling out the benefit from accuracy. Third, the vast majority of publications test algorithms at mid-optical stations; for tropical zones with AOP> 0.6, there is little public data, and the adaptation of models based on the principle of transfer learning is still experimental. The next logical step seems to be the creation of a unified multi-level ensemble: a satellite transformer predicts the AOP field, a PINN module specifies radiation components with physical limitations, and a probabilistic power generator closes the loop, optimizing not only the technical but also the economic metric. This solution can simultaneously reduce reserve costs, reduce the carbon footprint of dispatching, and increase the resilience of the power system to rare large aerosol events.

2 CONCLUSION

The multi-level approach "AOF → radiation → power" considered in the work confirmed the key assumption: under clear skies, it is the variations in the aerosol optical thickness that dominate the balance of photovoltaic generation forecast errors. A comparative analysis of recent publications shows that even a simple inclusion of the Copernicus-CAMS AOP daytime forecast reduces the normalized error of DNI and GHI by 15–20% relative to the persistent scheme without taking aerosol into account. The developed theoretical architecture combines a spatio-temporal aerosol encoder, a recurrent radiation block, and a probabilistic power transformer. The rejection of strictly deterministic coefficients in favor of joint training of all three stages ensures end-to-end consistency of the information flow and allows the formation of asymmetric confidence intervals focused on real economic penalties for undergeneration. In practical tests on representative open samples, OEDI has achieved nRMSE ≤ 4% for an hourly horizon, a level comparable to the best published results for transformer networks.

The economic benefits of such error reduction are confirmed by independent system assessments. A report by Lawrence Berkeley Lab showed that regional reserve coordination alone, based on accurate probabilistic forecasts, can save up to USD 1.24 billion annually with a RES share of 34–65%. An analysis by the California Commission (CEC-500-2023-025) additionally indicates that increasing the accuracy of daily solar generation forecasts by 20% reduces fuel and start-up costs of thermal power plants by almost 3% per system loading cycle.

However, there are still unresolved issues that limit further improvements in reliability:

• spectral detail. Monochrome 550 nm AOP does not differentiate between coarse sea salts and fine sulfates; the difference in attenuation of direct radiation reaches 14% with the same integral thickness. It is necessary to switch to multi-node spectral contributions in the composition of input features;

• rare extreme events. Volcanic eruptions and massive forest fires can reduce the annual DNI resource by several standard deviations, completely going beyond the boundaries of the training samples. For system stability, it is necessary to implement

scenario ensembles of the "black swan" and adaptive PINN correctors;

• observation infrastructure. In Europe, less than 40 AERONET third-level stations are active; gaps can be filled through cheaper photometers and the use of lidar profiles in the data-fusion mode;

• economically weighted metrics. Most works optimize symmetric RMSE, while in capacity markets the penalty coefficient for underforecasting is at least twice the penalty for surplus. A wider application of quantile weighting and CRPS functions is required.

Thus, the theoretical and methodological basis presented in the paper demonstrates that the combination of aerosol load forecasting, physically informed neural networks and probabilistic power inference can provide accuracy sufficient for operational dispatch planning. Further research should be aimed at (i) spectral expansion of inputs, (ii) integration of lidar observations, (iii) development of asymmetric loss functions that directly reflect the structure of market penalties, and (iv) creation of ensembles robust to rare large-scale aerosol anomalies. Fulfillment of these tasks will further reduce the uncertainty of solar generation and reduce the costs of balancing reserves in power systems with a high share of RES.

REFERENCES

Andreev, V.V., Arshinov, M.Yu., 2022. Koncentratsiya troposfernogo ozona nad territoriej Rossii v 2021 g. Optika atmosfery i okeana, 35 (7), 559–571.

Bezrukova, N.A., Chernokulsky, A.V., 2023. Rossiyskie issledovaniya oblakov i osadkov v 2019–2022 gg. Izvestiya RAN. Fizika atmosfery i okeana, 59 (7), 882–914.

Karamov, D.N., 2018. Matematicheskoe modelirovanie solnechnoj radiatsii s ispolzovaniem mnogoletnih meteorologicheskih ryadov. Izvestiya Tomskogo politehnicheskogo universiteta. Inzhiniring georesursov, 329 (1), 69–88.

Mordynsky, A.V., 2024. Metody opredeleniya rasseyannogo solnechnogo izlucheniya s tselyu povysheniya tochnosti prognozirovaniya pochasovoj vyrabotki PV. Problemy energetiki, 48 (2), 73–82.

Sakerin, S.M., Gorbarenko, E.V., Kabanov, D.M., 2008. Osobennosti mnogoletej izmenchivosti aerodzol'noj opticheskoj tolshchiny atmosfery i otsenki vliyaniya razlichnyh faktorov. Optika atmosfery i okeana, 21 (7), 625–631.

Salamov, O.M., Mamedova, L.G., Salmanova, F.A., Yusupov, I.M., 2021. Vybor modeley raschyota summarnoj solnechnoj radiatsii yasnogo neba v naklonnoj ploskosti v vysokih shirotah Rossii. Problemy energetiki, 44 (6), 58–69.

Zavyalov, D.D., Solomakha, T.A., 2021. Parametrizatsiya pogloshcheniya solnechnoj radiatsii snezhno-ledyanym pokrovom v termodinamicheskoj modeli l'da Azovskogo morja. Morskoy gidrofizicheskiy zhurnal, 37 (5), 538–553.