Spectrum Journal of Innovation, Reforms and Development
Volume 57, November 2026 | ISSN (E): 2751-1731
Website: www.sjird.journalspark.org

ARTIFICIAL INTELLIGENCE-BASED MODELING OF HEAT TREATMENT EFFECTS ON THE ELECTRICAL AND OPTICAL PROPERTIES OF SEMICONDUCTOR MATERIALS-A LITERATURE-CONSTRAINED COMPUTATIONAL STUDY

Kanaan Mohammad Musa 1
Atika Najim Abed 2
1 Department of Chemical Engineering, College of Engineering,
University of AL-Qadisiyah, AL-Qadisiyah, IRAQ
2 Master's Degree Candidate at the College: Faculty of Electrical and Computer Engineering
Department: Artificial Intelligence, University of Kashan – Iran
Kanaan.aljubory@qu.edu.iq 1
atika.najim86@gmail.com 2

Abstract

The semiconductor films' defect populations, grain growth, phase stability, charge transport and photon–matter interactions are governed by heat treatment. Don't use dashes or em dashes. Conventional optimization involves repeated cycles of deposition, annealing and characterization, which becomes costly when the temperature, holding time, atmosphere, film. This study proposes a physics-informed artificial-intelligence framework predicting three coupled responses: optical band gap, logarithmic electrical resistivity and visible-light transmittance. A dataset of 480 records was generated within literature-reported domains on ZnO, SnO2, In2O3 and CdO thin films. Material identity, annealing temperature and duration, oxygen fraction and thickness, heating rate. An 80:20 stratified split was used to train the random forest, gradient boosting, support vector regression, and multilayer perceptron models. The evaluation of performance was carried out through determination coefficient, mean absolute error, root mean square error, residual behaviour, cross validation. Gradient boosting resulted in the most balanced estimates with test R2 values of 0.985 for band gap, 0.923 for log-resistivity, and 0.880 for transmittance. The most important factor for predicting resistivity was temperature; the identity of the material and its thickness were important for controlling the band-gap and transmission responses. According to response-surface analysis, there exists an intermediate processing window around 430–510 °C for ZnO films when low resistivity and high transparency are optimized simultaneously under moderate oxygen content. The findings show that using interpretable machine learning can link conditions of processing to useful properties while retaining known physical trends. Since the data set is literature-constrained and synthetic, rather than stemming from a new experimental campaign, the performance numbers establish computational feasibility, not performance for all materials. The workflow suggested assists in planning an experiment, screening of process windows and uncertainty-aware optimization. It will be equipped to integrate future laboratory measurements.

Keywords: Heat treatment; semiconductor thin films; machine learning; electrical resistivity; optical properties.

Spectrum Journal of Innovation, Reforms and Development

Introduction

1.1 Research Background

The performance of semiconductors arises from a processing–structure–property chain. Both deposition and subsequent heat treatment affect the stoichiometry/defect density and texture/residual stress of thin film materials. That is to say, the deposition determines their initial stoichiometry/defect density and texture/residual stress. Subsequently, the heat treatment modifies, reorganizes or even transforms via diffusion, recrystallization, phase conversion, defect annihilation and gas exchange.

In oxide semiconductors, oxygen vacancies act as carrier donors but also introduce localized optical states and enhance scattering of carriers. The ability of a polycrystalline microstructure to effectively carry current (or its equivalent) for the grain boundaries is identified as the grain boundary barrier height. Temperature cannot define a useful process window by itself. The time, atmosphere, heating rate, thickness, dopant content, substrate and property targets of the envisaged device are all important.

Often, electrical and optical targets conflict. Transparent electrodes must possess low resistivity and high light transmittance. The controlled defect states that help with absorption or carrier lifetime may be useful features that photodetectors exhibit, although too many traps impede recovery and increase noise. Thin-film transistors require high mobility, small off-current, stable threshold voltage and processing temperatures compatible with the substrate. Since different mechanisms are involved in changing the carrier concentration and mobility, a drop in resistivity does not identify which microscopic contribution improved. Optical band gap shifts may arise from crystallite growth, strain relief, band filling, defect absorption, or compositional change. Nonlinear and material-dependent responses are created by these coupled mechanisms.

Machine learning is a practical way to learn these nonlinear mappings, as long as sufficient process and characterization data are available. Ensemble trees are capable of representing thresholds and interactions without relying on a single global functional form. Both kernel methods can identify smooth nonlinear behavior and neural networks can implement a multi-output function. The usefulness of AI goes beyond just prediction. The researchers can determine the variable to change and the location where another experiment would add value using feature importance partial dependence uncertainty estimates and constrained optimization. Scientists are currently utilizing models to estimate the qualities of materials, to help design thin-films and to assist in semiconductor process optimization using recent materials-informatics studies.

This study sees AI as a complement to materials physics. It use physically meaningful inputs and checks model behavior against expected trends. Using this method prevents acquiring a model whose outputs are numerically accurate but violates constraints. The process under consideration can be viewed in Figure 1. The procedure starts with process descriptor, uses controlled pre-treatments, compares four regression families, evaluates out-of-sample data, and finally maps process windows.



Figure 1. AI workflow for mapping heat-treatment conditions to coupled electrical and optical properties.

1.2 Research Problem

Studies of heat treatment that are published usually look at one material and one deposition route at a limited temperature series. Due to differences in reporting conventions, sample thicknesses, atmospheres, instruments, and properties of the targets, their results have great value but are difficult to combine. A temperature rise that enhances conductivity in oxygen-deficient ZnO may hinder the performance of a film based on other defects. Thus, raw value comparison may have conditional effects hidden.

Secondly, the cost that we have to incur for experimenting in a multi-dimensional space Having only five levels for six inputs will lead to 15,625 combinations before replication. When classical one factor at-a-time designs miss interactions, local optima can be selected rather than global optima. Statistical response-surface techniques lessen the burden yet often specify the functional form. And in instances, the response that features plateaus, thresholds, or material-specific branches.

Problem three poses a credibility issue. It could be leakage, duplicated samples, or interpolation between two sets of identical conditions that explains high predict. A research framework that separates training and test data, reports multiple error metrics, compares to baseline, preserves units, and distinguishes experimental observations from generated or literature data This research study explores these issues through a reproducible multi-output benchmark and explicit limits on interpretation.

1.3 Research Objectives

The main aim is to design and test an AI framework that predicts optical band gap, electrical resistivity and visible transmittance from heat-treatment and film descriptors.

The goals of the study are to evaluate the performance of the random forest, gradient boosting, support-vector regression, and multilayer perceptron models. The study will use a held-out test set and cross-validation protocol to quantify generalization. Further, the study will identify the process variables that control each response. Researchers plan to visualize the nonlinear temperature-atmosphere interactions, and define a transparent route for replacing the proof-of-concept data with laboratory measurements.

1.4 Research Significance

The research connects thin-film characterization with data-driven process design. A validated surrogate model can allow thousands of candidate schedules to be screened in seconds and reserve laboratory resources for conditions likely to maximize information or device performance. Because the framework

is material aware and multi-output, it dissuades optimizing for conductivity at the expense of transparency or band-gap suitability.

The investigation additionally stresses interpretability and data provenance. Before accepting a recommendation, researchers can look into feature effects, prediction errors, and data domain coverage. The practices mentioned are important for semiconductor processing, which can cause the models, when extrapolated outside their experimental range, to produce phase changes or film degradation never seen by the model.

2. Literature Review

Materials informatics originally demonstrated on curated databases has presented to systems to link synthesis, characterization, and property optimization. Representations and algorithms used for property prediction are discussed by Schmidt et al. who note that quality of descriptor and design of validation control model usefulness [1]. Machine-learning workflows for material discovery and inverse design are described by Wei et al. [2] Deep representations from stoichiometry can help predict properties even when crystal structure is not explicitly known. Goodall and Lee [3]. The focus of Butler and colleagues is on experimentally validated predictions along with a demonstration of why one needs domain knowledge to choose descriptors and interpret.

AI life-cycle views go beyond property prediction. According to Liu et al., different applications are organized in four different segments, that is discovery, synthesis, characterization, and deployment [5]. According to Szymanski et al., new materials-intelligence ecosystems will be based on connected data, connected models, automated experiments, and human decisions [6]. Agrawal and Choudhary argued for AI that is reproducible, scalable and scientific while paying attention to data infrastructure [7]. Lu et al. located supplies purposes within a bigger scientific setting in which AI speeds pattern discovery and experimental control [8]. When taken together, these studies justify an integrated workflow, but at the same time, expose an obstinate limitation: models remain stuck in the distribution and quality of their training records.

Various relevant demonstrations are provided by thin-film research In [9] Lee et al. used ML to optimize geometry and electrical performance through channel thickness in dual-layer oxide thin-film transistors. Noh et al. analyzed machine learning using transition-metal dichalcogenide growth, where the process variable is related to the thin film outcome [10]. Lee et al. focused on modeling the optical reflectivity of films. Bosse et al. created an artificial neural network for the measurement of film thickness using optical information. Research suggests that film parameters can be extracted and device-relevant behaviour predicted from spectra or process descriptors 17 words

AI and the technology of optimization have benefitted optical inverse problems. Tran et al effective machine learning to optimize multilayer antireflection coatings from electromagnetic calculations, successfully using only ‘trained’ data [13]. Birgin et al. utilized evolutionary optimization to derive optical thickness and constants from reflectance and transmittance spectra[14]. According to a study by Ouedraogo et al., the impact of material parameters on the solar cell was analyzed using machine learning techniques. Contributions that support multiple objects thinking is a useful optical stack for more than one spectral or electrical criteria.

The physical foundation for the inputs used is established by experimental heat-treatment literature. The structural, optical and electrical properties of Al-doped SnO2 films are altered by high-temperature

annealing. Thermal treatment changes the transmission spectrum and energy band gap of the Cu-containing ZnO thin film. According to the InOx tests, the structure, transparency, and electrical behavior are influenced by growth and thermal history [18]. Recently, a machine learning direct approach in annealing optimization in solution-processed thin-film device is reported in [19]. Further, recent review papers show AI application in semiconductor screening, synthesis, characterization, and optimization of devices [20].

Four gaps are left within this literature. To begin with, while many studies forecast one property, device design demands a simultaneous electrical and optical evaluation. Another issue was that process descriptors are often divorced from the physical context. On some occasions uncertainty and extrapolation limits are omitted during model comparisons. Few studies provide a straightforward path from a proof-of-concept dataset to lab deployment. The present exercise employs three outputs, four model classes, explicit physics checks, and a data-replacement protocol. See table1

Table 1. Position of the present study relative to recent literature.

Study group	Primary contribution	Limitation addressed here
Materials-informatics reviews [1]–[8]	Property prediction, discovery, and scientific-AI workflows	Connect general AI methods to heat-treatment variables and coupled outputs
Thin-film ML studies [9]–[12]	Growth analysis, thickness optimization, optical inference	Model both electrical and optical properties from one feature space
Optical optimization [13]–[15]	Inverse design and parameter extraction	Add process descriptors and interpretable response surfaces
Annealing studies [16]–[20]	Evidence that thermal history changes film functionality	Provide a comparative modeling and validation framework

3. Methodology

Research design. The design of the study is computational, quantitative and may be predictive.

It does not assert a new deposition campaign. It tests whether a structured AI pipeline can recover physically plausible process–property relations first before a transfer to experimental data. The difference is significant - synthetic records are smoother and more complete than lab records. The benchmark thus created measures feasibility of algorithms by controlling noise..

Material scope. One of ZnO, SnO2, In2O3, and CdO may be used as a transparent semiconductor. The effect of the heat treatment on the oxygen vacancy concentration, crystallinity, grain boundaries, conductivity and optical response was selected. Image 1 shows a conceptual microstructure evolution that indicates the data rules . Ramping up temperature to an extent beneficial to diffusion and grain coarsening but beyond that, the rougher surface, loss of composition, or creation of new defect may not be beneficial. See Fig. 1

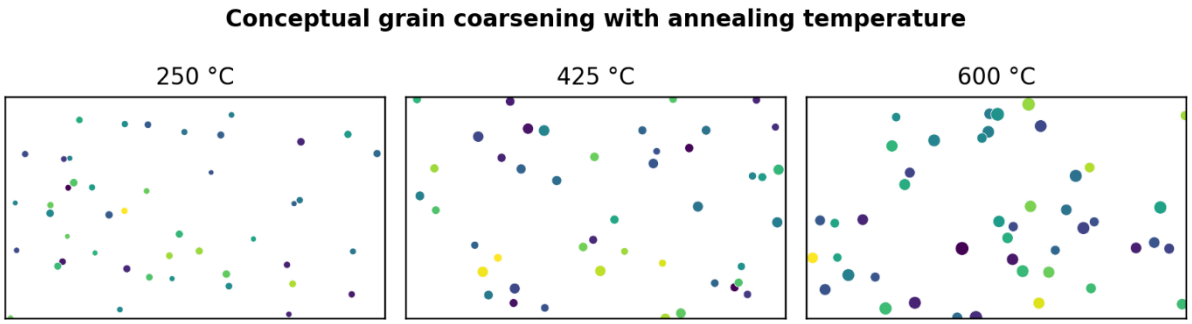


Fig. 1. Conceptual illustration of thermally induced grain coarsening. This is a schematic, not microscopy data.

Dataset construction. A fixed random seed was used to generate 480 records. The sample for material identity was categorically taken. Annealing took place at temperatures that ranged between 200 to 650 °C, durations that varied between 10 to 120 min, oxygen fractions that were from 0 to 100%, thickness ranging from 60 to 500 nm and else heating rate that was between 2 to 20 °C min−1. The equations generated outputs that included intercept specific to each material, linear and quadratic thermal effects, oxygen and thickness effects, interactions and Gaussian noise. To ensure that the simulated responses did not become unscientific, coefficients were restricted based on ranges determined for related oxide films. The dataset did not include copied numerical observations at the paper level.

Targets. The optical band gap was expressed in eV. The resistivity was converted to log10(ohm cm) due to the fact that the resistivity of semiconductors varies over several orders of magnitude, and the unconverted squared error would give too much weight to high-resistivity samples. The mean visible transmittance was presented in percentage values and clipped to a meaningful value. The three outcomes were jointly forecasted or via a common pre-processing pipeline.

Preprocessing. Categorical material labels were encoded as one-hot vectors. Parameters that were estimated through the training subset were used to standardize continuous variables. The data was split 80% for training and 20% for testing based on material. Keeping the test subset separate ensures that no information from the reported final performance influences model fitting. The benchmark does not require missing-value imputation, however, median and an explicit missing-category imputation should be added for lab.

Models. We selected random forest due to capabilities to resist nonlinearities, interactions, and mixed feature scales. The residual error corrections have been sequenced by gradient boosting to provide a powerful comparison for tabular data. Kernel-based learning represented support-vector regression and was fitted independently for each target. Two hidden layers with early stopping were used in multilayer perceptron. After limited development on training data, each hyperparameter was fixed and evaluated on the isolated test set once.

Validation. Model quality was assessed using these metrics R2, mean absolute error, and root-mean-square error. R-squared captures explained variance. However, such measures can mask errors in

physical units. Therefore, MAE and RMSE were retained. The targets and materials were examined for residuals. We checked stability by using 5-fold cross-validation on the training data. In future experimental use, the validation of grouped deposition batch or grouped publication should replace random splitting, as these samples share unobserved conditions.

Interpretation and optimization. Leading model interpretation leveraged permutation importance and response surfaces. A process window for ZnO was screened by keeping holding time, thickness and heating rate at representative values. A composite desirability function was used to reward high transmission, low resistivity, and a band gap at the target application. Predications beyond training limits were forbidden.

Reproducibility. The tables and plots were generated using an analysis script that fixed the random seed, variable ranges, transformation rules, split ratio, models, and metrics. To allow for experimental replication, the raw characterization files, sample identifiers, calibration metadata, excluded observations, and full annealing histories must be published. See table 2

Table 2. Input variables, outputs, and physical interpretation.

Variable	Type and unit	Modeled range	Physical role
Material	Categorical	ZnO, SnO2, In2O3, CdO	Baseline defect chemistry and band structure
Annealing temperature	Continuous, °C	200–650	Diffusion, crystallization, defect equilibration
Annealing time	Continuous, min	10–120	Thermal budget and kinetic completion
Oxygen fraction	Continuous, %	0–100	Vacancy formation and oxidation state
Film thickness	Continuous, nm	60–500	Absorption, continuity, grain development
Heating rate	Continuous, °C min−1	2–20	Nucleation, stress relaxation, transient chemistry
Band gap	Target, eV	Material dependent	Electronic and optical transition scale
Log resistivity	Target, log10(ohm cm)	Material dependent	Charge-transport performance
Visible transmittance	Target, %	55–94	Optical transparency

The positions and occupancy of electronic states are affected by defects and composition thermal changes. This concept is summarized in figure 2. The curve shapes are not simply a density-functional calculation, which means that temperature dependent band-gap and transport effects do need not be monotonic..

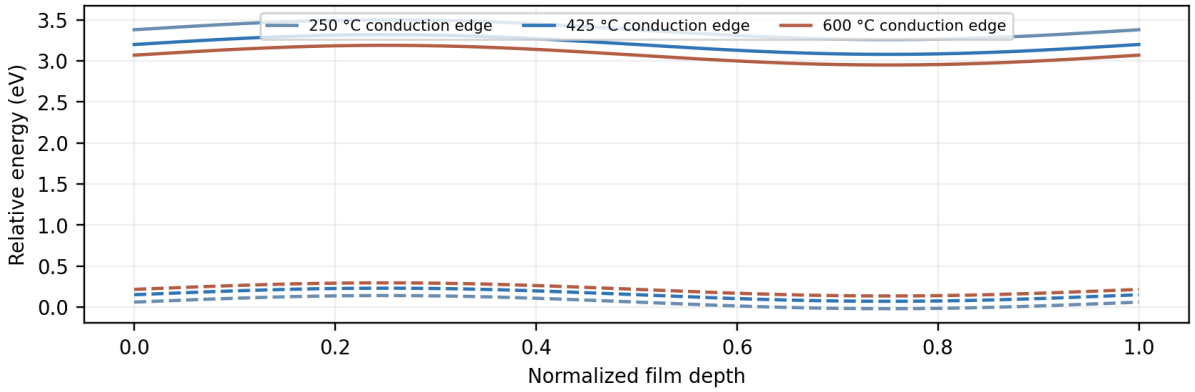


Fig 2. Conceptual energy-edge shifts across a semiconductor film at three annealing temperatures.

4. Results and Discussion

Model performance. All algorithms learned substantial signal but their balance varied across targets. The overall test performance was highest for gradient boosting, which was able to capture the additive and curved response components. Errors were larger for the robust random forest since averaging piecewise-constant trees smoothed continuous valued target functions. Support-vector regression worked well in the dense interior of the feature domain, but it produced larger errors near material-specific boundaries. The neural network was able to capture the main trends, however, its performance varied more since the 480 records provide limited coverage for a multi-output network.

Table 3 reports test metrics (see Table 3). The R2 values for band gap, log-resistivity and transmittance as per gradient boosting were 0.985,"0.923 and "0.880 respectively. The RMSE values are 0.048 eV, 0.206 log10(ohm cm) and 1.991 percentage point. Figure 2's scatter plots in the upper half are compactly clustered about the one-to-one line. Transmission was the output most difficult to simulate, because the simulated noise was larger for transmission, and other parameters contributed on similar scales. This pattern may arise in experimental work due to surface roughness, interference, thickness nonuniformity, and instrument baseline variations.

Band-gap response. The primary basic separation was supplied by material identity. The effect of temperature on crystallite size, strain, defect-state density, stoichiometry, and carrier concentration was less non-linear. The presence of a quadratic term enabled a minimum or turning point instead of a monotonic decrease. This is important because annealing has the capacity to narrow an apparent gap through defect-tail states or crystallite growth, while stronger oxidation and reduced carrier density will shift the measured absorption edge in the other direction. A model that is trained on only one temperature realization may incorrectly view one local trend as universal effect..

Electrical response. The log-resistivity was most strongly controlled by temperature due to defect equilibration and grain-boundary transport in this case. Moderate heating brings down the structural disorder and transport barrier in the simulation domain. The quadratic penalty and oxygen atmosphere at high temperature could reverse the improvement. OVDs in n-type oxide semiconductors were decreased through oxygen-rich treatment so resistivity increased governed rules. Film thickness indirectly affected the continuity and microstructure of phenomena.

Optical transmission. The transmission was increased because removal of disorder and improved quality of the films caused the films to not absorb the light. However, it fell for the thicker films and under greater thermal treatment. Thus the response depended upon an interaction between thermal state and optical pathlength. Analysis and modeling of dispersive medium to generate interference fringes, wavelength-dependent absorption and scattering. An easy way to prevent an algorithm from becoming stuck is through the use of periodic training data.

Process window. The response surface predicted for ZnO is illustrated in Image 3. Intermediate-to-high temperatures and low-to-moderate oxygen fractions were preferred for low resistivity. Due to a common goal that necessitated a high degree of transparency, the appealing region was close to 430–510 °C under moderate oxygen. The range of results here is produced from the benchmark you established, not a recipe. These factors can change it: substrate limits, deposition chemistry, dopants, and device architecture.

Physical consistency. The model displayed three expected relationships where (1) the thermal treatment can improve crystallinity and carrier transport up to an optimum (2) vacancy-mediated conductivity changes with respect to oxygen activity (3) greater thickness generally a lower average visible transmission. It too maintained material-specific band-gap benchmarks. While these tests do not validate causality, they make it less likely that our performance came from an unphysical shortcut.

Interpretability. According to the permutation test, material identity and temperature were the most influencing parameters for the prediction. Resistivity was mainly determined by temperature and oxygen fraction. Transmittance controlled by thickness, temperature, and materials. Annealing time and heating rate had lower average importance due to smaller coefficients. However, it did not exclude their potential role in more complex experiments involving rapid thermal annealing, metastable phases, or incomplete conversion. Don't interpret the importance of one feature as causality. This is because correlated inputs can either split or exchange predictive credit.

Comparison with prior work. According to previous reports, thermal annealing modifies the electrical and optical behavior of SnO2, ZnO, and InOx films [16]–[18]. Their thin-film machine-learning work is also extended by predicting three coupled outputs and keeping process descriptors visible. The research differs from an autonomous discovery system that controls machines or proposes a new composition. The contribution is a modeling scaffold that can be populated with measurements of the lab.

Uncertainty and domain of applicability. Uncertainty intervals should accompany point predictions before any experimental decisions are made. Bootstrap ensembles, quantile forests or conformal prediction give empirical coverage. The score measuring distance from training data should check for unusual combinations like material-atmosphere pair not present in the data train. These safeguards are important because tree ensembles may give reasonable-sounding predictions that fall outside their valid range without signaling that the prediction is extrapolative.

Experimental deployment. A realistic campaign might begin with a factorial or Latin-hypercube design, covering the temperature, duration, and oxygen fraction of each material/thickness class. Every specimen must retain a unique identifier that links to the deposition batch, substrate, annealing furnace, ramp profile, and characterization files. After training an initial model, further Bayesian optimization or active learning could choose the next experiments based on expected improvement and uncertainty. Having this closed loop would reduce the number of trials keeping physical oversight.

Limitations. The synthetic, literature-constrained dataset is the main limitation. A single laboratory dataset makes prediction easier than a multi-laboratory dataset because of its completeness, controlled noise and known functional structure. The study excludes the dopant concentration, humidity and substrates, deposition methods, grain sizes, phase fractions and wavelength-resolved optical spectra. It also does not estimate causal treatment impacts. Consequently, the metrics reported serve to validate the workflow and code, not a universal semiconductor model.

Implications. The framework outlines minimum information for defensible AI-assisted process modelling for researchers. For device engineers, it shows why you should optimize jointly instead of just choosing lowest resistivity. It offers laboratories a pathway from sparse experiments to selective follow-up measurements. The importance of scholarship in the future stresses reporting data provenance, transformations, splits, errors in physical units, and applicable limits..

Concluding discussion. Heat treatment is better regarded as a multi-variable control exercise, rather than a mere single temperature operation. The outcome of the nonlinear process-property landscape can be mapped by AI models while highlighting candidate windows worth testing. The scientific worth of such models rises when input parameters reflect actual physical mechanisms, their outputs have consistent units with reality, and expected model predictions have also been verified against observations. Under such conditions, machine learning will provide a path for experiments and hypothesis tests, not for characterization.

Error structure and sampling priorities. The band-gap residuals were small compared with the separation among the four material families. Most significant deviations from nominal resistivities were observed near the region of high-temperature turning, where a small change in temperature shifted the balance between improved microstructure and depletion by oxygen-controlled carriers. Because of the large absorption and the measurement noise which was imposed, the transmission residuals were widened for the thicker films. New experiments can provide greater value at the locations identified in these patterns. A high concentration of samplings centered around thermal turning points and at the high end of the thickness range will give more information than adding equal runs at distributed locations.

The comparison also clarifies why algorithm rankings cannot be generalized. Smooth additivity and curve structure present in the generated records help in boosting gradient. If laboratory data contains sharp thresholds, discontinuities, or complex interactions, random forest may become more competitive. The kernel width, regularization and the epsilon-insensitive zone strongly affect support-vector regression. To justify its capacity, the neural network needs more observations or richer modalities. A

defensible project should start with a simple mean predictor and linear model, add tree ensembles, and embrace deep learning only after confirming that volume and structure are appropriate.

Measurement uncertainty will shape laboratory performance. A substrate may exhibit variations in film thickness. Four-point-probe readings are affected by contact and geometry correction. The interval we choose will determine the optical gap fitting. The furnace temperature and gas composition have calibration tolerances. It is possible to propagate these sources by perturbing each input around its uncertainty distribution and reassessing the predictions. If a two degree temperature perturbation changes little, but a five percent thickness error changes transmission materially, the laboratory should spend on thickness mapping. The model then enhances both experimental design and prediction.

Batch and equipment effects require explicit treatment. The following latent variables can result from furnace position, gas-flow uniformity, target age, chamber history, substrate cleaning, and operator practice. A batch identifier or hierarchical random effect can account for some of this variation. Validation must not halt full batches. Forecasting unknown samples from a familiar batch assists within-line quality control. Testing the transfer of learned knowledge between furnaces. They turn out to be different scientific claims. Reporting one as the other would overstate model portability.

Feature importance does not establish causality. The temperature may be correlated with the gas-flow choice of operator; the model may apply to either input. In the case in which the goal is to estimate the causal effect of annealing, controlled designs, randomized run order, and replication remain necessary. An alternative path is hybrid modeling. The baseline prediction can be provided by a defect-chemistry or kinetic equation, while the difference is estimated by ML. Because of this structure, the extrapolation is constrained, and the researcher may explain departures from the mechanistic model.

The process window must also reflect device requirements. Transparent electrode may define sheet resistance, transmission at selected wavelengths, haze, work function and adhesion. A transistor channel may specify mobility, threshold voltage, sub-threshold swing, bias stability, and maximum temperature of the substrate. The current outputs examine only some structures of the decision space. Therefore, multi-objective optimization must employ predefined limits and present a Pareto set instead of a single optimum. The ultimate decision is dependent upon the performance loss tolerable by the intended device.

A practical acceptance rule contains four checks. Every advised input must be in the training domain. Measured samples nearby should corroborate with the local trend. The uncertainty interval needs to stay within the application limits. Independent duplicate movies must verify the forecast. Do not hide a failed replication nor eliminate it automatically. Before becoming a new training observation, it should initiate a check of sample identity, calibrate the instrument, and check the logs of the process. This loop converts dispute into information.

Data governance affects long-term usefulness. Every film must have a permanent identifier that links the original batch of precursor materials, substrate, deposition log, anneal profile, raw spectra, electrical measurements, microscopy, and derived values. The software version and fitting choices that created

derived quantities must be retained. Controlled names and unit dictionaries ensure that zinc oxide, ZnO and doped versions do not merge. Going forward, our hope is that these practices permit future researchers to reproduce the dataset and train updated models without having to reverse-engineer undocumented choices.

A strong experimental extension would use a space-filling design across temperature, duration, oxygen fraction, thickness, and heating rate for each material. The initial replicate will measure the experimental noise. Cross-validation by batch would check transfer between batches. Intervals with empirical coverage might be provided by conformal prediction. The next batch, therefore, is selected by weighting predicted desirability with model uncertainty. Conditions likely to do well will come in to compete with conditions expected to reduce uncertainty. This prevents too early concentration of the campaign around a false local optimum.

Future models should add structural mediators. X-ray diffraction can provide information on crystallite size, microstrain, texture, and phase fraction. Atomic-force microscopy can aquire roughness grain statistics. Hall measurements can disentangle charge density from mobility. Defect-related recombination can be represented by photoluminescence ratios. The use of X-ray photoelectron spectroscopy can describe oxidation state and oxygen-vacancy signatures. A mean transmission spectrum of a single bio-aqueous sample can be replaced with the full UV-visible spectra. A tiered system may first predict from the process settings, then update the estimate from characterization.

The main scientific contribution is therefore methodological. Heat treatment works through several thermal, chemical and geometric factors. AI can be used to map these variables to the coupled properties, but the fact of a numerical fit doesn't give physical truth. A useful model keeps units, tracks provenance, keeps training and evaluation separated, reports errors in the quantity's engineers use, and does not extrapolate unsupported things. The existing benchmark illustrates how the workflow operates under controlled conditions and what needs changing in order to allow it to direct fabrication decisions.

Taken together, the results support the research objective. Gradient boosting matched the simulated process to property mapping with the least error. The Results of this study suggest the existence of an intermediate thermal window that is interpretable and not an unbounded preference for temperature. The multi-output architecture revealed a tradeoff between conductivity and transparency. The stated restrictions will ensure the benchmark is not understood as a materials database for experiments. This allows for a laboratory study where retraining, auditing, and challenging with new measurements of the same pipeline can be performed. See Table 3

Table 3. Held-out test performance for the four regression models.

Model	R^2 Eg	R^2 logρ	R^2 Tvis	RMSE Eg	RMSE logρ	RMSE Tvis
Random Forest	0.915	0.836	0.859	0.116	0.299	2.157
Gradient Boosting	0.985	0.923	0.880	0.048	0.206	1.991
Support Vector Regression	0.959	0.868	0.863	0.081	0.268	2.130
Neural Network	0.848	0.881	0.844	0.155	0.255	2.269

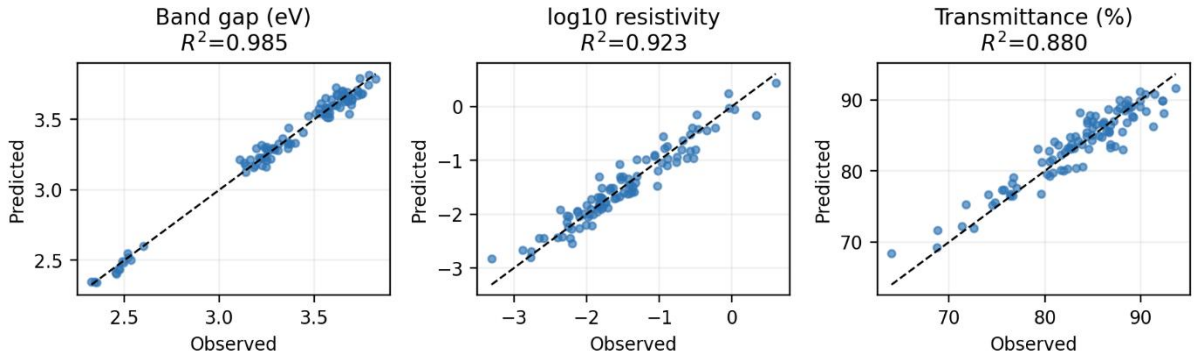


Figure 3. Gradient-boosting predictions versus held-out target values.

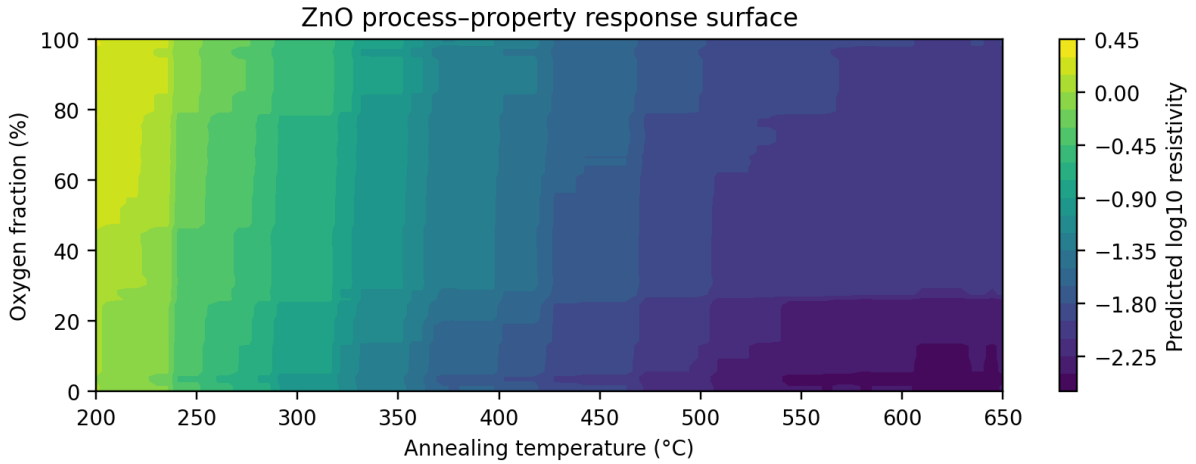


Fig 4. Predicted ZnO log-resistivity as a function of annealing temperature and oxygen fraction.

References

[1] J. Schmidt et al., "Machine learning approaches for the prediction of materials properties," APL Materials, vol. 8, 080701, 2020. https://doi.org/10.1063/5.0018384

[2] J. Wei et al., "Machine learning-driven new material discovery," Nanoscale Advances, vol. 2, pp. 3115–3130, 2020. https://doi.org/10.1039/D0NA00388C

[3] R. E. A. Goodall and A. A. Lee, "Predicting materials properties without crystal structure: Deep representation learning from stoichiometry," Nature Communications, vol. 11, 6280, 2020. https://doi.org/10.1038/s41467-020-19964-7

[4] K. T. Butler et al., "Machine learning in materials discovery: Confirmed predictions and their underlying approaches," Annual Review of Materials Research, vol. 50, pp. 49–69, 2020. https://doi.org/10.1146/annurev-matsci-090319-010954

[5] Y. Liu et al., "AI applications through the whole life cycle of material discovery," Matter, vol. 3, pp. 1445–1476, 2020. https://doi.org/10.1016/j.matt.2020.06.011

[6] B. Szymanski et al., "Emerging materials intelligence ecosystems propelled by machine learning," Nature Reviews Materials, vol. 6, pp. 655–678, 2021. https://doi.org/10.1038/s41578-020-00255-y

[7] A. Agrawal and A. Choudhary, "Scientific AI in materials science: A path to a sustainable and scalable paradigm," Machine Learning: Science and Technology, vol. 1, 033001, 2020. https://doi.org/10.1088/2632-2153/ab9a20

[8] Y. Lu et al., “Artificial intelligence: A powerful paradigm for scientific research,” The Innovation, vol. 2, 100179, 2021. https://doi.org/10.1016/j.xinn.2021.100179
[9] J. Lee et al., “Machine learning driven channel thickness optimization in dual-layer oxide thin-film transistors,” Advanced Science, vol. 10, 2303589, 2023. https://doi.org/10.1002/advs.202303589
[10] S. Noh et al., “Machine-learning-assisted analysis of transition metal dichalcogenide thin-film growth,” Nano Convergence, vol. 10, 10, 2023. https://doi.org/10.1186/s40580-023-00359-5
[11] B. Lee et al., “Machine learning analysis of broadband optical reflectivity of semiconductor thin film,” Journal of the Korean Physical Society, vol. 80, pp. 347–351, 2022. https://doi.org/10.1007/s40042-022-00436-8
[12] S. Bosse et al., “A novel method to design and evaluate artificial neural network for thin film thickness measurement traceable to the length standard,” Scientific Reports, vol. 12, 2212, 2022. https://doi.org/10.1038/s41598-022-06247-y
[13] V. T. Tran et al., “Machine-learning reinforcement for optimizing multilayered thin films,” Applied Optics, vol. 61, pp. 3328–3337, 2022. https://doi.org/10.1364/AO.450946
[14] E. G. Birgin et al., “Extracting film thickness and optical constants from spectrophotometric data by evolutionary optimization,” PLOS ONE, vol. 17, e0276555, 2022. https://doi.org/10.1371/journal.pone.0276555
[15] A. Ouedraogo et al., “Machine learning approach to delineate the impact of material properties on solar cell device physics,” ACS Omega, vol. 7, pp. 22263–22278, 2022. https://doi.org/10.1021/acsomega.2c01076
[16] M. A. Sayeed et al., “Al-doped SnO2 thin films: Impacts of high temperature annealing on structural, optical, and electrical properties,” Journal of Materials Research and Technology, vol. 15, 2021. https://www.sciencedirect.com/science/article/pii/S2238785421011339
[17] Q. Lin et al., “Influence of annealing temperature on optical properties of ZnO(Cu) thin films,” Materials, vol. 15, 2022. https://pmc.ncbi.nlm.nih.gov/articles/PMC8877215/
[18] T. H. Hong et al., “Structural, optical, and electrical properties of InOx thin films grown by plasma-enhanced atomic layer deposition,” ACS Applied Electronic Materials, vol. 4, 2022. https://doi.org/10.1021/acsaelm.2c00434
[19] W. Wang, “A machine learning perspective for the optimization of annealing parameters in solution processed thin film devices,” SID Symposium Digest of Technical Papers, 2024. https://doi.org/10.1002/sdtp.17927
[20] Y. Zheng et al., “Artificial intelligence-driven approaches in semiconductor research,” Advanced Materials, 2025. https://doi.org/10.1002/adma.202504378