

Classification of Worker Productivity and Resource Allocation Optimization with Machine Learning: Garment Industry

A'isyah Nur Aulia Yusuf¹, Zakiyyan Zain Alkaf², Elsa Sari Hayunah Nurdiniyah³,
Tri Wisudawati⁴, Muhammad Ihsan Fawzi^{*5}

^{1,3} Electrical Engineering, Fakultas Teknik, Universitas Jenderal Soedirman, Indonesia

^{2,4} Industrial and Mechanical Engineering, Fakultas Teknik, Universitas Jenderal Soedirman, Indonesia

⁵ Informatics, Fakultas Teknik, Universitas Jenderal Soedirman, Indonesia

Email: ihswfz@incomso.com

Received : Aug 11, 2025; Revised : Aug 31, 2025; Accepted : Sep 2, 2025; Published : Oct 16, 2025

Abstract

This study presents an integrated predictive-prescriptive framework for improving workforce management in the garment industry by combining machine learning classification with linear programming optimization. Using a publicly available dataset of 1,197 production records, productivity levels were categorized into low, medium, and high classes. Data preprocessing included handling missing values, one-hot encoding of categorical variables, and class balancing using SMOTE. Eleven classification algorithms were evaluated, with LightGBM achieving the highest performance (accuracy 78.3%, weighted F1-score 78.3%, Cohen's Kappa 63.4%) after hyperparameter tuning via Bayesian Optimization. The optimized model's predictions were then incorporated into a linear programming model, implemented with PuLP, to maximize the allocation of high-productivity workers across production departments under capacity constraints. The results yielded an allocation plan assigning 117 high-productivity workers, significantly enhancing potential operational efficiency. The novelty of this work lies in integrating an optimized ensemble learning model with mathematical programming for end-to-end productivity classification and resource allocation, a combination rarely explored in labor-intensive manufacturing contexts. This framework offers a scalable decision-support tool for data-driven workforce planning and could be adapted to other manufacturing domains with similar operational structures.

Keywords : Bayesian Optimization, Garment Industry, Linear Programming, Machine Learning, Productivity Classification, Random Forest

This work is an open access article and licensed under a Creative Commons Attribution-Non Commercial 4.0 International License



1. INTRODUCTION

The garment industry remains a vital driver of economic growth in many developing countries, characterized by labor-intensive production processes where worker productivity plays a pivotal role in determining operational efficiency, cost control, and competitiveness in global markets [1]. Conventional productivity assessment methods—such as manual observation and basic statistical analysis—are often limited in their ability to process large-scale, complex datasets, which may lead to inaccurate evaluations and suboptimal resource allocation [2]. In the era of Industry 4.0, data-driven approaches supported by advanced machine learning (ML) algorithms provide a promising pathway to address these challenges.

Recent studies have demonstrated the effectiveness of machine learning models in predicting and classifying worker productivity within garment manufacturing. For instance, Balla et al. [3] applied the Random Forest algorithm to predict garment worker productivity, achieving superior performance in terms of mean absolute error (MAE $\approx$ 0.0787) compared to linear regression and neural network

approaches. Similarly, Saad [4] reported that ensemble techniques such as bagging, when applied to decision-tree-based models, significantly improved accuracy in textile worker performance evaluation, achieving an accuracy rate of approximately 99.2%. Parallel to predictive modeling, optimization methods have been explored to enhance resource allocation in garment production. Shakirullah et al. [5] developed a linear programming (LP) model for the garment industry in Bangladesh, resulting in a 22% increase in profit and a 37% reduction in operational costs. More recently, Kong et al. [6] employed mixed-integer linear programming (MILP) for garment line balancing under Lean manufacturing principles, achieving over 50% reductions in labor costs while meeting production constraints. On the machine learning optimization front, Bayesian Optimization has emerged as an efficient and accurate approach for hyperparameter tuning. Zlobin and Bazylevych [7] demonstrated that Bayesian Optimization improved the performance of best models while reducing computational overhead compared to grid search or random search methods. This aligns with broader literature identifying Bayesian Optimization as a leading hyperparameter tuning method due to its ability to achieve strong results with fewer, more informative evaluations [8].

The complexity and high-dimensionality of modern manufacturing datasets pose significant challenges for traditional machine learning models. Consequently, ensemble methods, particularly gradient boosting algorithms like LightGBM, have gained prominence due to their robustness and ability to capture intricate non-linear relationships without extensive feature engineering [9]. However, the predictive power of these models is highly dependent on the careful tuning of their numerous hyperparameters, a task that is computationally expensive and often intractable with conventional methods like grid search. This is precisely where Bayesian Optimization offers a distinct advantage, providing a principled and sample-efficient strategy to navigate the complex hyperparameter space and identify optimal model configurations, thereby ensuring that the full predictive potential of the underlying model is harnessed for industrial applications [10].

Zhang et al. [11] utilized XGBoost with Gaussian Process-based Bayesian tuning to predict maximum stress in lattice structures produced by additive manufacturing, achieving high accuracy and generalizability. Similarly, Mahani et al. [12] combined Bayesian Optimization with Singular Value Decomposition (SVD) and XGBoost in optical design, reducing computational burden while maintaining predictive precision. More broadly, Bayesian-optimized XGBoost has been shown to efficiently optimize complex, expensive-to-evaluate processes in manufacturing, such as configuring atmospheric plasma spraying and fused deposition modeling, through a batch Bayesian optimization framework [13]. Recent advances also include multi-objective Bayesian optimization frameworks that enable autonomous decision-making in smart manufacturing by guiding sequential experimentation [14], and constrained Bayesian optimization techniques applied to curing process control in manufacturing settings [15].

In workforce planning and production scheduling, mathematical optimization methods remain central. A 2025 survey in *Annals of Operations Research* [16] reviewed workforce modeling methodologies—covering recruitment, training, attrition, scheduling—and highlighted the continued relevance of analytical and simulation-based decision models. Reconfigurable manufacturing systems have leveraged Mixed-Integer Linear Programming (MILP) to simultaneously allocate workers, machines, and job schedules [17]. Workforce size determination using dynamic programming demonstrated cost savings and staffing efficiency in manufacturing [18]. Moreover, workforce scheduling in shift-based workshops has successfully applied MILP to minimize labor cost while respecting regulations and demand constraints [19].

Beyond predictive accuracy, a significant barrier to the adoption of advanced models like gradient boosting in high-stakes industrial settings is their inherent 'black box' nature. For workforce management, where decisions directly impact personnel and operational outcomes, this lack of

transparency can lead to mistrust and hinder effective implementation by managers who need to understand the rationale behind a model's predictions. This challenge has spurred the development of Explainable AI (XAI) techniques, such as SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-agnostic Explanations), which provide crucial insights into which features most influence a model's output [20]. By illuminating the key drivers of productivity predictions—such as a worker's tenure, department, or specific task performance—XAI not only enhances model transparency and fosters stakeholder trust but also provides actionable insights for targeted training and process improvements [21].

Integration of predictive analytics with prescriptive optimization has emerged in related domains. An explainable, Bayesian-optimized model improved component failure detection in predictive maintenance, enhancing decision transparency [22]. In the context of resource allocation beyond manufacturing, a decision-support framework applied Bayesian-optimized XGBoost and Optuna for budget distribution, demonstrating potential applicability to workforce allocation problems [23]. These studies highlight three converging trends:

1. The effectiveness of Bayesian-optimized in modeling complex manufacturing processes and structural predictions.
2. The robustness of optimization methods (LP, MILP, dynamic programming) in workforce/resource scheduling under constraints.
3. The growing interest in integrated predictive–prescriptive frameworks that support real-time decision-making in industrial operations.

Despite these advances, most studies have treated predictive modeling and optimization as separate domains, with limited research integrating the two into a cohesive, end-to-end decision-support framework. Furthermore, only a few works have incorporated both advanced ensemble learning methods, such as LightGBM, with Bayesian Optimization for model enhancement, followed by mathematical programming for actionable resource allocation [24]. Addressing these gaps, this study proposes an integrated framework that classifies garment worker productivity using an optimized model and subsequently applies integer linear programming via PuLP to optimize workforce allocation [25]. Specifically, the research objectives are:

- (1) to develop a high-accuracy worker productivity classification model,
- (2) to enhance its performance through Bayesian Optimization, and
- (3) to implement resource allocation optimization to maximize high-productivity workforce deployment.

This combined predictive–prescriptive approach offers both methodological novelty and practical relevance to the data-driven management of labor-intensive manufacturing operations. The research addresses the problem of inconsistent productivity assessment and inefficient resource allocation, which are prevalent in labor-intensive industries. By leveraging a dataset of garment worker performance, we employ advanced machine learning models to predict productivity and use linear programming to allocate workers optimally across departments. This study contributes to the growing body of literature on data-driven decision-making in industrial settings.

2. METHOD

The methodological framework adopted in this study combines predictive analytics and prescriptive optimization to address the dual objectives of classifying worker productivity and determining optimal resource allocation in the garment industry. The process begins with data preparation, including preprocessing, feature transformation, and handling of class imbalance, followed by the development and benchmarking of multiple machine learning models. The best-performing model is then fine-tuned through hyperparameter optimization to maximize predictive performance [26].

The output of this classification stage—predicted productivity classes—is subsequently used as input to a linear programming model designed to optimize workforce distribution across production departments.




Figure 1. Proposed Method

2.1. Data Preparation and Model Development

The dataset utilized in this study comprised 1,197 production records obtained from a publicly available garment industry dataset compiled by Imran et al. [27]. Each record included multiple attributes relevant to worker productivity assessment, namely department (finishing or sewing), production quarter, day of the week, work-in-progress (WIP), number of workers assigned, and actual productivity achieved during the recorded period. To enable the classification process, the continuous productivity variable was discretized into three classes: Low (< 0.5), Medium ($0.5\text{--}0.8$ inclusive), and High (> 0.8), based on established industry performance thresholds. This categorization allowed for targeted analysis and more interpretable decision-making in the subsequent optimization stage.

The preprocessing stage was conducted to ensure data integrity and compatibility with machine learning algorithms [28]. Initially, date attributes were converted to a standardized datetime format and excluded from further processing, as temporal identifiers were not considered predictive features. Missing WIP values were replaced with zero to preserve dataset completeness. Nominal categorical features, including quarter, department, and day, were transformed into numerical form via one-hot encoding. This transformation allowed algorithms to interpret categorical distinctions without imposing ordinal relationships.

The preprocessed data was then divided into training (80%) and testing (20%) sets using stratified sampling to maintain class proportions in both sets. A critical challenge in the dataset was class imbalance, where certain productivity classes, particularly High, were underrepresented. To mitigate the bias introduced by this imbalance, the Synthetic Minority Oversampling Technique (SMOTE) [29] was applied. SMOTE generated synthetic samples for minority classes by interpolating between existing minority class instances in feature space, resulting in a balanced dataset of 1,341 records.

A comprehensive model benchmarking was conducted to identify the most effective classification algorithm. Eleven classifiers were evaluated, encompassing linear models (Logistic Regression), non-linear algorithms (Support Vector Machine, K-Nearest Neighbors, Gaussian Naive Bayes, Decision Tree) [30], ensemble methods (Random Forest, Gradient Boosting, AdaBoost) [31], and advanced gradient boosting implementations (XGBoost, CatBoost, LightGBM) [32]. Model performance was evaluated using three complementary metrics: overall accuracy, weighted F1-score, and Cohen's Kappa. The weighted F1-score was prioritized as the primary metric due to its ability to account for both class imbalance and precision–recall trade-offs, while Cohen's Kappa was used to assess classification

agreement beyond chance. The equation of F1-score is shown in equation (1) where precision, $P = TP/(TP + FP)$, and recall, $R = TP/(TP + FN)$ [33], [34].

$$F1 - score = 2 \times \frac{Precision \times Recall}{Precision + Recall} \quad (1)$$

The algorithm achieving the highest weighted F1-score—was selected for further refinement through Bayesian Optimization [35]. This hyperparameter tuning process explored a defined search space for four key parameters: the number of boosting iterations (`n_estimators`), maximum tree depth (`max_depth`), learning rate (`learning_rate`), and subsample ratio (`subsample`). Bayesian Optimization began with five randomly sampled initial points to establish prior knowledge, followed by 15 guided iterations using the Expected Improvement (EI) acquisition function. This approach allowed the search to efficiently exploit promising hyperparameter configurations while continuing to explore less-visited regions of the search space. The optimal configuration identified through this process was then used to retrain the best model on the full training dataset.

2.2. Optimization for Resource Allocation

Following the development of the optimized classification model, its predictions were employed as the input for a prescriptive analytics stage using Linear Programming (LP). The objective of the LP model was to determine the optimal allocation of workers across the two main departments—finishing and sewing—while maximizing the deployment of high-productivity workers, as identified by the classification model.

Decision variables were defined to represent the number of workers assigned to each department within each productivity category (Low, Medium, High). The optimization objective function was formulated to maximize the sum of high-productivity worker assignments, thereby aligning with managerial goals of improving overall production efficiency. Two primary constraints governed the model:

1. Departmental Capacity Constraint – The total number of workers allocated to each department could not exceed its historical maximum capacity as observed in the dataset.
2. Total Workforce Constraint – The total number of assigned workers across both departments was required to match the total number of workers available in the dataset, ensuring a feasible and realistic allocation strategy.

The LP problem was implemented using the PuLP optimization library [25] in Python, with the default CBC solver employed to identify optimal solutions. The solver iteratively explored feasible allocation configurations, calculating the corresponding objective values until the global optimum was reached. By integrating the predictive outputs of the optimized model with the prescriptive optimization process, this methodology forms a unified decision-support framework. The approach enables data-driven workforce planning that simultaneously leverages the predictive accuracy of advanced machine learning and the operational efficiency of mathematical optimization.

All stages of data preparation, model training, hyperparameter optimization, and LP-based resource allocation were implemented in Python 3.11 within a Google Colab environment. The primary libraries used included Pandas and NumPy for data handling, scikit-learn for baseline model development, imbalanced-learn for SMOTE, XGBoost/CatBoost/LightGBM for advanced boosting algorithms, bayes_opt for Bayesian Optimization, and PuLP for optimization modeling.

3. RESULT

The comparative performance of eleven classification algorithms is presented in the table. Among all tested models, LightGBM and CatBoost achieved the highest performance with an identical accuracy

of 78.3%. However, LightGBM had a slight edge with an F1-Score of 0.784 and a Cohen's Kappa of 0.634. The Gradient Boosting model followed closely, recording an accuracy of 77.5%. Conversely, SVM and Logistic Regression recorded the lowest accuracies at 26.0% and 46.8%, respectively. This low performance can be attributed to the limited ability of these two models to capture the non-linear relationships present in the dataset.

Table 1. Model Performance Comparison

Model	Accuracy	F1-Score	Kappa
LightGBM	0.783333	0.783529	0.634093
CatBoost	0.783333	0.782122	0.632379
Gradient Boosting	0.775000	0.777669	0.624097
Random Forest	0.766667	0.766462	0.606038
XGBoost	0.766667	0.766171	0.604706
Decision Tree	0.741667	0.741495	0.562250
AdaBoost	0.625000	0.644972	0.401164
KNN	0.512500	0.540565	0.257025
Logistic Regression	0.462500	0.468203	0.139235
Naive Bayes	0.504167	0.382436	0.084381
SVM	0.250000	0.260458	0.036317

The classification report for LightGBM, detailed in Table 2, shows consistently high precision, recall, and F1-scores across the three productivity classes (Low, Medium, High). Notably, the model demonstrated its best performance in the Low productivity class, achieving an F1-score of 0.84.

Table 2. Classification Report for LightGBM

Class	Precision	Recall	F1-Score	Support
Low	0.58	0.56	0.57	27
Medium	0.76	0.80	0.78	101
High	0.86	0.82	0.84	112

Bayesian Optimization was employed to fine-tune the Random Forest model, resulting in optimal hyperparameters: $n_estimators = 158.97778$, $max_depth = 30.674113$. The optimization process, as shown in Table 3, reached a peak F1-score of 78.7%, confirming the model's robustness after tuning.

Table 3. Bayesian Optimization Convergence for LightGBM

Iteration	F1-Score	n estimators	max depth	min samples	subsample
15	0.7874384	158.97778	30.674113	0.0695431	0.8898375
7	0.7866767	259.06126	18.595835	0.1832457	0.7819454
4	0.7858913	258.11066	18.493554	0.0627292	0.5917022
3	0.7792671	200.27875	38.322903	0.0159695	0.9849549

Based on the SHAP interaction plot on the Figure 1, the model's predictions are strongly and directly influenced by the main effects of *smv* and *targeted_productivity*. For both features, higher values (shown as red points) consistently push the prediction higher, while lower values (blue points) push it lower, indicating a clear positive correlation with the outcome. The *team* feature, by itself, shows a much weaker and less distinct impact on the model's predictions. Crucially, the off-diagonal plots reveal no significant interaction effects among these three variables. The influence of one feature does not appear to meaningfully change based on the value of another, as the points are largely scattered

around zero without a clear pattern. This suggests that the model treats *smv*, *team*, and *targeted_productivity* as largely independent predictors rather than features that have a synergistic or combined effect on the final output.




Figure 1. SHAP Result

Following the classification stage, a linear programming model was applied to optimize resource allocation based on predicted productivity categories. The optimization process assigned 117 high-productivity workers across production departments while respecting departmental capacity constraints. As shown in Table 4, the sewing department received the largest allocation (89 workers), followed by finishing sub-departments with 28 workers, respectively. Medium- and low-productivity workers were not assigned in the optimal configuration, reflecting the objective of maximizing production efficiency.

Table 4. Optimal Worker Allocation

Department	High	Medium	Low
Finishing	28.0	0.0	0.0
Sewing	89.0	0.0	0.0

4. DISCUSSIONS

The results from this study show that the combined approach of machine learning classification and linear programming can effectively support decision-making in the garment industry. The Random Forest model, after hyperparameter tuning, achieved an accuracy of 78.7%, a weighted F1-score of 78.6%, and a Cohen's Kappa of 63.9%. Result in this research in harmony with past research in context of health that LightGBM has the best result [36], [37]. These metrics suggest that the model was able to capture meaningful patterns in the data and distinguish well between low, medium, and high productivity workers. The best performance was observed in the Low productivity class ($F1 = 0.57$), which is encouraging because early identification of low productivity can help managers intervene before it impacts output. Performance for the Medium and High classes was slightly lower ($F1 = 0.80$ and 0.82), which indicates some overlap between these categories in the feature space. It is important to note that oversampling using SMOTE was applied after splitting the data into training and testing sets. Splitting the data first, then applying SMOTE only to the training set, results in higher and more realistic accuracy [38].

The hyperparameter tuning using Bayesian Optimization provided a small but meaningful boost in performance [39]. This suggests that careful tuning is worthwhile, as it can make the model more stable and adaptable to different types of input data. However, in the current setup, tuning decisions were based on the test set, which could also lead to overestimating performance. Using a separate validation set or nested cross-validation would help avoid this issue in future implementations. Once the productivity predictions were obtained, they were fed into a linear programming model to allocate workers to departments. The optimization objective was simple: maximize the number of high-productivity workers while staying within department capacity limits. The optimal solution assigned 142 high-productivity workers, with the sewing department receiving the largest share (89 workers), and the two finishing lines receiving 25 and 28 workers, respectively. Medium- and low-productivity workers were not allocated in the optimal plan, which is expected given the single-objective function. However, this also raises a practical point: in a real factory setting, some medium and low performers will still need to be assigned due to skill diversity, training needs, or fairness considerations.

One limitation of the current setup is that the LP model treats the classifier's predictions as absolute. In reality, predictions have uncertainty, and misclassifying even a small number of workers could affect the optimal allocation. A possible improvement would be to use the model's predicted probabilities rather than hard labels, allowing the optimization to work with expected productivity levels instead of assuming perfect classification. This would make the allocation more robust to classification errors. Another area worth exploring is interpretability. While this study focused on predictive accuracy and allocation efficiency, it would be valuable to examine which features most influence the classification. For example, if certain variables—such as work-in-progress levels or specific days of the week—consistently contribute to low productivity, these insights could guide operational changes. The code already contains the option to calculate SHAP values, and using these in practice would help explain the reasoning behind each prediction to managers.

Finally, the results here are promising for a proof-of-concept, but further work is needed before this approach can be deployed in a live production environment. The dataset used is relatively small and limited in scope, so additional data from different factories or time periods would help test generalizability. It would also be useful to extend the LP formulation to handle multiple objectives—such as balancing productivity with worker training or fairness—and to include additional operational constraints that reflect real-world scheduling complexity. Overall, the integration of a tuned LightGBM classifier with an optimization model offers a practical way to link data-driven predictions to actionable workforce planning. With refinements in data handling, model validation, and optimization design, this approach could provide a valuable decision-support tool for improving efficiency in labor-intensive manufacturing environments.

5. CONCLUSION

This study demonstrated the potential of combining machine learning classification with linear programming optimization to support workforce allocation in the garment industry. Using a publicly available dataset, multiple classification algorithms were benchmarked, with LightGBM emerging as the best-performing model. After hyperparameter tuning via Bayesian Optimization, the classifier achieved a weighted F1-score 78.7%. Indicates that the classifier demonstrates strong and balanced performance across all classes. This score is significant because the F1-score itself represents a harmonic mean of Precision (the accuracy of positive predictions) and Recall (the ability to identify all actual positive instances). The optimized predictions were then integrated into a linear programming model to determine the optimal allocation of workers. The resulting plan assigned 117 high-productivity workers across the sewing and finishing departments while respecting capacity constraints, thereby maximizing operational efficiency in line with the model's objective function.

While the results are promising, the study also revealed several considerations for future work. To obtain accurate performance estimates, split the data into training and test sets. Subsequently, the SMOTE technique should be applied exclusively to the training set to handle class imbalance. Addressing these through more rigorous validation techniques will strengthen the reliability of the model. Additionally, incorporating prediction uncertainty into the optimization process, applying explainability techniques such as SHAP, and expanding the optimization objectives to reflect fairness, skill diversity, and training needs would improve both robustness and practical applicability. In summary, the proposed predictive–prescriptive framework offers a scalable and data-driven approach to workforce management in labor-intensive manufacturing. With further refinement and validation on larger, more diverse datasets, it holds significant potential as a decision-support tool for improving productivity and operational efficiency in the garment industry.

CONFLICT OF INTEREST

The authors declare no conflict of interest between the authors or with the research object in this paper.

ACKNOWLEDGEMENT

The authors would like to acknowledge the financial support provided by Universitas Jenderal Soedirman through the Program Penelitian Riset Peningkatan Kompetensi under Grant No. 14.573/UN23.34/PT.01/V/2025.

REFERENCES

- [1] J. S. Mah, “Industrial-Led Economic Development of Cambodia: Implications for Low-Income Developing Countries,” *Southeast Asian Economies*, vol. 39, no. 2, pp. 198–210, 2022, doi: 10.1355/ae39-2e.
- [2] J. VAN BIESEBROECK, “ROBUSTNESS OF PRODUCTIVITY ESTIMATES *,” *J Ind Econ*, vol. 55, no. 3, pp. 529–569, Sep. 2007, doi: 10.1111/j.1467-6451.2007.00322.x.
- [3] I. Balla, S. Rahayu, J. Jaya Purnama, and C. Author, “GARMENT EMPLOYEE PRODUCTIVITY PREDICTION USING RANDOM FOREST,” *Jurnal TECHNO Nusa Mandiri*, vol. 18, no. 1, pp. 49–54, Mar. 2021, [Online]. Available: www.nusamandiri.ac.id
- [4] H. R. Saad, “Use Bagging Algorithm to Improve Prediction Accuracy for Evaluation of Worker Performances at a Production Company,” *Industrial Engineering & Management*, vol. 07, no. 02, 2018, doi: 10.4172/2169-0316.1000257.
- [5] F. M. Shakirullah, M. Uddin Ahammad, and M. Forhad Uddin, “Profit Optimization of an Apparel Industry in Bangladesh by Linear Programming Model,” *American Journal of Applied Mathematics*, vol. 8, no. 4, p. 182, 2020, doi: 10.11648/j.ajam.20200804.13.
- [6] R. W. M. Kong, D. Ning, and T. H. T. Kong, “A Mixed-Integer Linear Programming (MILP) for Garment Line Balancing,” *International Journal of Scientific Research and Modern Technology (IJSRMT)*, Feb. 2025.
- [7] M. ZLOBIN and V. BAZYLEVYCH, “BAYESIAN OPTIMIZATION FOR TUNING HYPERPARAMETERS OF MACHINE LEARNING MODELS: A PERFORMANCE ANALYSIS IN XGBOOST,” *Computer systems and information technologies*, no. 1, pp. 141–146, Mar. 2025, doi: 10.31891/csit-2025-1-16.
- [8] J. Snoek, H. Larochelle, and R. P. Adams, “Practical Bayesian Optimization of Machine Learning Algorithms,” Aug. 2012, [Online]. Available: <http://arxiv.org/abs/1206.2944>
- [9] T. A. Lai, N. Nguyen, and Q. Bui, “Hyper-parameter optimization of gradient boosters for flood susceptibility analysis,” *Transactions in GIS*, vol. 27, no. 1, pp. 224–238, Feb. 2023, doi: 10.1111/tgis.13023.

- [10] A. Candelieri, A. Ponti, and F. Archetti, "Bayesian Optimization in Wasserstein Spaces," 2022, pp. 248–262. doi: 10.1007/978-3-031-24866-5_19.
- [11] Z. Zhang, Y. Zhang, Y. Wen, and Y. Ren, "Data-driven XGBoost model for maximum stress prediction of additive manufactured lattice structures," *Complex & Intelligent Systems*, vol. 9, no. 5, pp. 5881–5892, Oct. 2023, doi: 10.1007/s40747-023-01061-z.
- [12] M. R. Mahani, I. A. Nechepurenko, T. Flisgen, and A. Wicht, "Combining Bayesian Optimization, Singular Value Decomposition, and Machine Learning for Advanced Optical Design," *ACS Photonics*, vol. 12, no. 4, pp. 1812–1821, Apr. 2025, doi: 10.1021/acsphotonics.4c02157.
- [13] X. Guidetti, A. Rupenyan, L. Fassel, M. Nabavi, and J. Lygeros, "Advanced Manufacturing Configuration by Sample-Efficient Batch Bayesian Optimization," *IEEE Robot Autom Lett*, vol. 7, no. 4, pp. 11886–11893, Oct. 2022, doi: 10.1109/LRA.2022.3208370.
- [14] A. S. Asru, H. Khosravi, I. Ahmed, and A. Azeem, "From automation to autonomy in smart manufacturing: a Bayesian optimization framework for modeling multi-objective experimentation and sequential decision making," *The International Journal of Advanced Manufacturing Technology*, vol. 137, no. 9–10, pp. 5027–5057, Apr. 2025, doi: 10.1007/s00170-025-15407-z.
- [15] Y. Li, Q. Zhang, M. Limaye, and G. Li, "Constrained Bayesian Optimization under Bivariate Gaussian Process with Application to Cure Process Optimization," May 2025, [Online]. Available: <http://arxiv.org/abs/2506.00174>
- [16] W. J. Peck, "Workforce planning: a review of methodologies," *Ann Oper Res*, Aug. 2025, doi: 10.1007/s10479-025-06734-1.
- [17] B. Vahedi-Nouri, R. Tavakkoli-Moghaddam, Z. Hanzálek, and A. Dolgui, "Workforce planning and production scheduling in a reconfigurable manufacturing system facing the COVID-19 pandemic," *J Manuf Syst*, vol. 63, pp. 563–574, Apr. 2022, doi: 10.1016/j.jmsy.2022.04.018.
- [18] M. K. So and S. L. Kek, "Workforce size problem in manufacturing with dynamic programming approach," 2020, p. 090005. doi: 10.1063/5.0018444.
- [19] Quantagonia, "How to get started: Workforce Management and Shift Planning Optimization." Accessed: Aug. 02, 2025. [Online]. Available: <https://www.quantagonia.com/post/workforce-management-shift-planning-optimization>
- [20] R. Mitchell, E. Frank, and G. Holmes, "GPUPtreeShap: massively parallel exact calculation of SHAP scores for tree ensembles," *PeerJ Comput Sci*, vol. 8, p. e880, Apr. 2022, doi: 10.7717/peerj-cs.880.
- [21] S. Makubhai, G. R. Pathak, and P. R. Chandre, "Interpreting Healthcare Insights: The Power of Explainable AI and Enhanced Data Analysis," 2024, pp. 413–424. doi: 10.1007/978-981-97-6684-0_33.
- [22] H. Kumar, K. K. Bharti, D. Dhabliya, R. Agarwal, S. Kumar, and S. Tripathi, "Explainable Bayesian-Optimized XGBoost Model for Component Failure Detection in Predictive Maintenance," 2024, pp. 137–155. doi: 10.4018/979-8-3693-1347-3.ch010.
- [23] Saulo11340, "AI-driven budget allocation using machine learning and Bayesian optimization," Mar. 2025. Accessed: Aug. 02, 2025. [Online]. Available: <https://github.com/Saulo11340/AI-Budget-Allocation>
- [24] R. Shi, X. Xu, J. Li, and Y. Li, "Prediction and analysis of train arrival delay based on XGBoost and Bayesian optimization," *Appl Soft Comput*, vol. 109, p. 107538, Sep. 2021, doi: 10.1016/j.asoc.2021.107538.
- [25] N. Bruschi, A. Garofalo, F. Conti, G. Tagliavini, and D. Rossi, "Enabling mixed-precision quantized neural networks in extreme-edge devices," in *Proceedings of the 17th ACM International Conference on Computing Frontiers*, New York, NY, USA: ACM, May 2020, pp. 217–220. doi: 10.1145/3387902.3394038.

-
- [26] L. Yang and A. Shami, "On hyperparameter optimization of machine learning algorithms: Theory and practice," *Neurocomputing*, vol. 415, pp. 295–316, Nov. 2020, doi: 10.1016/j.neucom.2020.07.061.
- [27] A. Al Imran, S. Rahim, and T. Ahmed, "Mining the productivity data of the garment industry," 2021.
- [28] B. Qolomany *et al.*, "Leveraging Machine Learning and Big Data for Smart Buildings: A Comprehensive Survey," *IEEE Access*, vol. 7, pp. 90316–90356, 2019, doi: 10.1109/ACCESS.2019.2926642.
- [29] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE: Synthetic Minority Over-sampling Technique," *Journal of Artificial Intelligence Research*, vol. 16, pp. 321–357, Jun. 2002, doi: 10.1613/jair.953.
- [30] E. R. Kessler, M. Shah, S. K. Gruschkus, and A. Raju, "Cost and Quality Implications of Opioid-Based Postsurgical Pain Control Using Administrative Claims Data from a Large Health System: Opioid-Related Adverse Events and Their Impact on Clinical and Economic Outcomes," *Pharmacotherapy: The Journal of Human Pharmacology and Drug Therapy*, vol. 33, no. 4, pp. 383–391, Apr. 2013, doi: 10.1002/phar.1223.
- [31] O. FAUST, P. C. A. ANG, S. D. PUTHANKATTIL, and P. K. JOSEPH, "DEPRESSION DIAGNOSIS SUPPORT SYSTEM BASED ON EEG SIGNAL ENTROPIES," *J Mech Med Biol*, vol. 14, no. 03, p. 1450035, Jun. 2014, doi: 10.1142/S0219519414500353.
- [32] V. Borisov, T. Leemann, K. Seßler, J. Haug, M. Pawelczyk, and G. Kasneci, "Deep Neural Networks and Tabular Data: A Survey," *IEEE Trans Neural Netw Learn Syst*, vol. 35, no. 6, pp. 7499–7519, Jun. 2024, doi: 10.1109/TNNLS.2022.3229161.
- [33] C. J. Van Rijsbergen, "Information retrieval: theory and practice," *Proceedings of the Joint IBM/University of Newcastle upon Tyne Seminar on Data Base Systems*, 1979.
- [34] D. J. Hand, P. Christen, and N. Kirielle, "F*: an interpretable transformation of the F-measure," *Mach Learn*, vol. 110, no. 3, pp. 451–456, Mar. 2021, doi: 10.1007/s10994-021-05964-1.
- [35] M. A. Setitra, M. Fan, B. L. Y. Agbley, and Z. E. A. Bensalem, "Optimized MLP-CNN Model to Enhance Detecting DDoS Attacks in SDN Environment," *Network*, vol. 3, no. 4, pp. 538–562, Dec. 2023, doi: 10.3390/network3040024.
- [36] D. Zuo, L. Yang, Y. Jin, H. Qi, Y. Liu, and L. Ren, "Machine learning-based models for the prediction of breast cancer recurrence risk," *BMC Med Inform Decis Mak*, vol. 23, no. 1, p. 276, Nov. 2023, doi: 10.1186/s12911-023-02377-z.
- [37] Thulasi. M and G. Thailambal, "An Ensemble Machine Learning Model for Osteoporosis Risk Prediction from Medical Data," in *2025 International Conference on Machine Learning and Autonomous Systems (ICMLAS)*, IEEE, Mar. 2025, pp. 411–415. doi: 10.1109/ICMLAS64557.2025.10968017.
- [38] S. A. Alex, J. Jesu Vedha Nayahi, and S. Kaddoura, "Deep convolutional neural networks with genetic algorithm-based synthetic minority over-sampling technique for improved imbalanced data classification," *Applied Soft Computing*, vol. 156, p. 111491, May 2024, doi: 10.1016/j.asoc.2024.111491.
- [39] M. Lindauer *et al.*, "SMAC3: A Versatile Bayesian Optimization Package for Hyperparameter Optimization," *Journal of Machine Learning Research*, vol. 23, 2022.
-