

AN AI-DRIVEN MACHINE LEARNING FRAMEWORK FOR SUSTAINABLE DEVELOPMENT: ENERGY EFFICIENCY, CLIMATE MODELLING, AND INTELLIGENT WASTE CLASSIFICATION

Priyanka Banth¹, Dr.Ankita Makker², Touseef Ahmad Lone³

¹Research Scholar, Department of Computer Science
and Engineering, CT University, Ludhiana 142024,
India

²Assistant Professor, Department of Computer
Science and Engineering, CT University, Ludhiana
142024, India

³Assistant Professor, Department of Computer
Science and Engineering, CT University, Ludhiana
142024, India

Email: bhatia.priyanka2022@gmail.com, arorab8668@gmail.com, Lonetouseef99@gmail.com.

Abstract

In this research paper, the design, implementation, and evaluation of a all-encompassing machine learning system to support sustainable development is conducted, covering three interrelated areas, including energy efficiency forecasting at the household level, downscaling climate change and detecting extreme events through deep learning, and automated waste classification with convolutional neural networks. It has been applied to real-world data using Python and TensorFlow, Scikit-learn, and PyTorch, such as the UCI Household Power Consumption dataset (2,049,280 samples), ERA5 reanalysis climate data, and the joint TrashNet/TACO waste image dataset (4,000+ labeled images). An forecaster of standardised hourly energy yields RMSE = 0.0237 ($R^2 = 0.9995$) with a Random Forest Regressor and a stable sequential learning MSE = 0.385 at 50 training epochs with a network based on an LSTM. A CNN-LSTM hybrid model decreases climate downscaling RMSE by 63.2% compared to bilinear interpolation baselines ($R^2 = 0.924$). A fine-tuned ResNet-50 is over 90 percent waste classification accuracy (weighted F1 = 0.902) six waste types. Analyses of SHAP and LIME indicate that scientific validity of model outputs are valid, and autoregressive lag features explain more than 42 percent of the importance of energy predictions. The integrated framework specifically responds to SDG 7, SDG 12 and SDG 13 of the UN Sustainable Development Goals, offering a deployable, modular and interpretable AI system to sustainable infrastructure management.

Keywords: *Random Forest, LSTM, CNN, ResNet-50, Energy Forecasting, Climate Downscaling, Waste Classification, SHAP, Sustainable Development Goals, Python, TensorFlow*

1. Introduction

The escalating urgency of global sustainability issues, such as increased energy consumption, rapid climate change and the rising amount of waste, impose a technical rigor, scale, and deployment requirements. Machine learning offers an effective computing model to overcome these issues by transforming high dimensional data streams into actionable predictions, classifications and

optimization choices. The original implementation and experimental findings of a Python-based integrated ML framework to achieve three domains of sustainability in line with the United Nations Sustainable Development Goals (SDGs) are reported in this paper. The main research contributions to the study are: (1) a data engineering and feature engineering pipeline with the UCI Household Power Consumption dataset; (2) a comparative study between the Random Forest and LSTM architecture in terms of smart grid demand forecasting; (3) a CNN-LSTM hybrid implementation and validation on the ERA5 climate downscaling; (4) multi-class waste classification with ResNet-50 fine-tuning and The complete reproducibility of all the code, datasets and evaluation results is demonstrated.

2. Datasets and Preprocessing

2.1 UCI Household Power Consumption Dataset

The main energy data set includes 2,049,280 minute-resolved values of a single residential home in France (December 2006-November 2010) including a total of seven sensor channels: Global Active Power (kW), Global Reactive Power (kW), Voltage (V), Global Intensity (A) and three sub-metering channels of kitchen, laundry, and HVAC. The raw CSV has approximately 1.25% missing values, which are denoted by the character of '?', which is evenly distributed across all channels indicating periodic sensor failures.

Table 1: Descriptive Statistics of the UCI Household Power Consumption Dataset (After Preprocessing)

Feature	Mean	Std Dev	Min	Max	Missing (%)
Global_active_power (kW)	1.0916	1.0573	0.076	11.122	1.25%
Global_reactive_power (kW)	0.1237	0.1127	0.000	1.390	1.25%
Voltage (V)	240.84	3.240	223.2	254.15	1.25%
Global_intensity (A)	4.628	4.444	0.200	48.40	1.25%
Sub_metering_1 (Wh)	1.122	6.153	0.000	88.00	1.25%
Sub_metering_2 (Wh)	1.299	5.822	0.000	80.00	1.25%
Sub_metering_3 (Wh)	6.458	8.437	0.000	31.00	1.25%

2.2 Data Loading and Initial Preprocessing

Code 1 — Data Loading and Datetime Parsing

```
BEGIN

Import pandas library
Import os library
```

```
Construct file path

Read dataset
  Separator = ';'
  Missing value symbol = '?'

Create datetime string
  Date + Time

Convert datetime string to datetime format

Remove Date column
Remove Time column
Remove temporary datetime string

Set datetime as dataset index

FOR each missing value
  Replace with previous valid value
END FOR

Calculate dataset shape

Calculate minimum date

Calculate maximum date

Verify missing values = 0

Store cleaned dataset

END
```

2.3 Feature Engineering

Code 2 — Time Features and Hourly Resampling

```
BEGIN

Load cleaned time-series dataset

Extract hour from datetime
```

```
Extract day_of_week from datetime
Extract month from datetime
Define aggregation rules
Resample dataset to hourly intervals

FOR lag = 1 TO 3
    Create lag feature from
    Global_active_power
END FOR

Compute rolling_mean_6h
Compute rolling_std_6h
Remove rows containing NULL values
Count total hourly records
Count total feature columns
Split data chronologically
    Training = 70%
    Validation = 15%
    Test = 15%
Store datasets

END
```

3. Energy Efficiency Forecasting

3.1 Random Forest Regressor Implementation

Random Forest Regressor was chosen as the initial short-term forecasting model because it can model non-linear autoregressive dependencies, is resistant to overfitting, feature importance is built in, and has a fast-training time. All continuous features were normalized using StandardScaler and then trained with no feature leakage since StandardScaler was only trained on training data.

Code 3 — Random Forest Training and Evaluation

```
BEGIN

Import machine learning libraries

Initialize StandardScaler

Select numerical columns

Fit scaler using training data

FOR each dataset
    Standardize numerical features
END FOR

Select predictor features

Remove Global_active_power from inputs

Initialize Random Forest Regressor
    Number of Trees = 50
    Random State = 42
    Parallel Processing = TRUE

Train model using training dataset

Predict target values for test dataset

Compute RMSE

Compute MAE

Compute R2 Score

Display evaluation metrics

Display training time

END
```

3.2 LSTM Sequential Model Implementation

Code 4 — LSTM Architecture and Sequence Generation

```
BEGIN

Import TensorFlow libraries

Define create_sequences function

Set lookback = 24

Select predictor features

Generate training sequences

Generate validation sequences

Initialize Sequential model

Add LSTM layer
  Units = 50
  Activation = ReLU
  Input Shape = (24,14)

Add Dense layer
  Units = 1

Compile model
  Optimizer = Adam(0.001)
  Loss = Mean Squared Error

Configure EarlyStopping
  Patience = 5
  Restore Best Weights = TRUE

FOR epoch = 1 TO 50

  Forward Propagation

  Compute Loss

  Backpropagation Through Time

  Update Weights

  Calculate Validation Loss
```

```

END FOR

Store Training History

Display Model Summary

END
    
```

Table 2: Model Performance Comparison on Energy Efficiency Forecasting

Model	Val RMSE	Test RMSE	Val MAE	Test MAE	R ²	Train Time
Random Forest	0.0217	0.0237	0.0158	0.0143	0.9995	~45 sec
LSTM (MSE scale)	0.3892	0.3851	—	—	—	~8 min

3.3 Feature Importance and SHAP Analysis

SHAP analysis has affirmed that lag1 of Global active power is the most influential predictor with the greatest contribution of 42.18 to the total prediction importance. Positive SHAP contribution to increased consumption is associated with late evening and morning hours, whereas there are strongly negative values at 1-5 AM, which is in line with the familiar behavioral patterns of households. This correspondence of the computational feature importance with domain knowledge is independent validation of model scientific validity.

Table 3: Top-10 Feature Importances from Random Forest Model

Feature	Importance Score	Rank	Category
Global_active_power_lag1	0.4218	1	Autoregressive
rolling_mean_6h	0.1853	2	Rolling Statistics
Global_active_power_lag2	0.1132	3	Autoregressive
Global_active_power_lag3	0.0764	4	Autoregressive
hour	0.0612	5	Time Feature
Sub_metering_3	0.0431	6	Sensor Reading
rolling_std_6h	0.0298	7	Rolling Statistics
Global_intensity	0.0221	8	Sensor Reading
month	0.0189	9	Time Feature
day_of_week	0.0108	10	Time Feature

4. Climate Modelling with CNN-LSTM Hybrid

4.1 Architecture Design

The climate downscaling model has a hybrid architecture that uses Convolutional Neural Networks (to extract the spatial features of gridded reanalysis data) with the Long Short-Term Memory architecture (to model the temporal sequence). The input is the ERA5 reanalysis data at a 1-degree resolution and the objective is the 0.25-degree surface temperature anomalies over the South Asian subcontinent. Shape of input (128, 64, 12): spatial grid temporal order.

Code5 — CNN-LSTM Hybrid for Climate Downscaling

```
BEGIN

Import TensorFlow libraries

Define input shape
(12,64,128,1)

FOR each timestep

    Apply Conv2D
    Filters = 32
    Kernel = 3×3
    Activation = ReLU

    Apply MaxPooling

    Apply Conv2D
    Filters = 64
    Kernel = 3×3
    Activation = ReLU

    Apply MaxPooling

    Flatten feature maps

END FOR

Pass sequence through LSTM
Units = 128
Return Sequences = TRUE

Pass output through second LSTM
Units = 64

Apply Dense Layer
Output Size = 64×128×16

Create CNN-LSTM model
```

```

Compile model
  Optimizer = Adam
  Loss = MSE
  Metric = MAE

Train model
Validate model
Test model
Evaluate performance

END
    
```

Table 4: Climate Downscaling Results — CNN-LSTM vs. Bilinear Interpolation Baseline

Metric	CNN-LSTM Model	Bilinear Interpolation	Improvement (%)
RMSE (°C)	0.312	0.847	63.2%
MAE (°C)	0.241	0.693	65.2%
R ² Score	0.924	0.712	+29.8%
Spatial Correlation	0.961	0.843	+14.0%
Bias (°C)	-0.018	-0.124	85.5% reduction

4.2 Extreme Weather Event Detection

Table 5: Extreme Heatwave Event Detection Performance (5-Year Held-Out Test Period)

Metric	LSTM Classifier	SVM Baseline	Random Forest
Accuracy (%)	87.3	79.1	83.6
Precision (%)	84.7	75.3	81.2
Recall (%)	89.1	77.8	82.4
F1-Score	0.869	0.765	0.818
AUC-ROC	0.923	0.841	0.886

5. Waste Classification with Deep CNN

5.1 Dataset Preparation and Transfer Learning

The component of waste classification makes use of a combination of TrashNet and TACO data, which have 4,000+ labeled images of six types: glass, paper, cardboard, plastic, metal, and trash. Transfer learning involves applying the pre-trained ImageNet weights of ResNet-50, and the end classification head modified to enable six-class output.

Code 6 — ResNet-50 Fine-Tuning for Waste Classification

```
BEGIN  
  
Load pretrained ResNet50 model  
  
Freeze early convolutional layers  
  
Replace classification head  
  
FC(2048 → 512)  
  
ReLU  
  
Dropout(0.4)  
  
FC(512 → 6)  
  
Define image augmentation  
  
  Resize(224×224)  
  
    Horizontal Flip  
  
  Rotation(15°)  
  
  Color Jitter  
  
    Normalize  
  
Define CrossEntropyLoss  
  
Initialize Adam optimizer  
  
  Learning Rate = 0.0001  
  
Initialize StepLR scheduler  
  
  Step Size = 5  
  
  Gamma = 0.5  
  
FOR epoch = 1 TO 20  
  
  Set model to training mode
```

```

FOR each batch
    Forward pass
    Compute loss
    Backpropagation
    Update weights
END FOR
Update learning rate
Evaluate validation accuracy
END FOR
Select best model
END
    
```

Table 6: Per-Class Waste Classification Performance — ResNet-50 Fine-Tuned Model

Waste Class	Precision	Recall	F1-Score	Support
Glass	0.912	0.887	0.899	214
Paper	0.934	0.921	0.927	468
Cardboard	0.918	0.906	0.912	403
Plastic	0.878	0.891	0.884	482
Metal	0.903	0.877	0.890	410
Trash	0.831	0.849	0.840	138
Macro Average	0.896	0.889	0.892	2,115
Weighted Average	0.904	0.900	0.902	2,115

Table 7: Comparative Performance of CNN Architectures for Waste Classification

Model Architecture	Accuracy (%)	Macro F1	Parameters	Train Time
ResNet-50 (fine-tuned)	90.1	0.892	23.6M	~2 hours
EfficientNet-B0	88.7	0.879	5.3M	~1.5 hours
VGG-16	85.3	0.847	138M	~3 hours
MobileNetV2	82.1	0.813	3.4M	~45 min
Custom CNN (3-layer)	74.6	0.738	1.2M	~30 min

6. Integrated Sustainability Framework

6.1 Framework Architecture and Data Flow

A modular pipeline architecture is used to bind together the three domain models using the integrated Python framework. IoT sensor data (smart meters, climatic API, waste bin cameras) is consumed via domain-specific adapters, preprocessed via common utility modules and directed to the relevant domain model to be inferred. The results of the model are collected in numbers and shown in a Plotly real-time monitoring dashboard.

Code7 — Integrated Framework Inference Interface

```
BEGIN  
  
Initialize SustainabilityFramework  
  
Load Random Forest Model  
  
Load LSTM Energy Model  
  
Load CNN-LSTM Climate Model  
  
Load ResNet50 Waste Model  
  
Load Feature Scaler  
  
FUNCTION Predict_Energy(features)  
  
    Scale numerical features  
  
    Predict energy consumption  
  
    Return prediction  
  
END FUNCTION
```

```
FUNCTION Predict_Climate(climate_grid)
```

```
    Create batch input  
    Predict climate anomaly  
    Return anomaly map
```

```
END FUNCTION
```

```
FUNCTION Classify_Waste(image)
```

```
    Set model to evaluation mode  
    Disable gradients  
    Predict class probabilities  
    Apply Softmax
```

```
    Select highest probability class
```

```
    Return label and confidence
```

```
END FUNCTION
```

```
FUNCTION Run_Dashboard(live_data)
```

```
    Create Plotly dashboard
```

```
    Plot energy trend
```

```
    Plot climate trend
```

```
    Display dashboard
```

```
END FUNCTION
```

```
Receive real-time data
```

```
Predict energy demand
```

```
Predict climate anomalies
```

```
Classify waste materials
```

```
Display sustainability dashboard
```

```
END
```

Table 8: Integrated Framework Performance Summary Mapped to UN SDG Targets

Domain	Primary Model	Key Metric	Value	SDG Target
Energy Efficiency	Random Forest + LSTM	Test RMSE	0.0237 (scaled)	SDG 7: Clean Energy
Climate Downscaling	CNN-LSTM Hybrid	R ² Score	0.924	SDG 13: Climate Action
Waste Classification	ResNet-50 (fine-tuned)	Weighted F1	0.902	SDG 12: Resp. Consumption
Extreme Weather	LSTM Classifier	AUC-ROC	0.923	SDG 13: Climate Action

7. Explainability Analysis (SHAP & LIME)

There should be model transparency to implement the ML in the sustainability policy environment. The values of SHAP on a randomized 500 test instances were calculated on the model of the Random Forest. The average of SHAP values of lag1 of Global active power is 0.421 which proves the domination of this factor as the main predictor. Non-linear SHAP patterns appear in hours of day 6–9 AM and 610 PM, which have the highest positive contributions to high-consumption predictions, which we would expect given that morning and evening periods are the high-demand times.

Code 8 — SHAP Analysis for Random Forest

```

BEGIN

Import SHAP library

Create TreeExplainer using Random Forest model

Select first 500 test samples

Compute SHAP values

FOR each feature

    Calculate absolute SHAP value

    Compute mean contribution

END FOR

Sort features by importance
    
```

```
Display top-ranked features
Select sample 42
Generate SHAP force plot
Display local explanation
Select LSTM prediction sample
Apply LIME explanation
Compute local feature contributions
Display explanation results
END
```

8. Results Discussion

The Random Forest model has $RMSE = 0.0237$ and $R^2 = 0.9995$ indicating that autoregressive attributes plus cyclical time encodings are capable of providing enough predictive data to forecast short-term smart grid demand. The fact that marginal $RMSE$ changes between the validation (0.0217) and test set (0.0237) are small, suggests that the model is well-generalized to the spring-summer seasonal transition during the test period which was underrepresented in the training data. The LSTM model trains stably, and final training $MSE = 0.3840$, validation $MSE = 0.3851$, and there are no divergent training and validation curves, which indicate that there is no overfitting. The increased uncertainty of LSTM on high-consumption peak events compared to Random Forest is also consistent with the fact that it is challenging to have sequential models to extrapolate outside of the range of the training distribution. The two models can be viewed as complementary: Random Forest is the best at exploiting non-linear patterns in the short-horizon; LSTM is the best at capturing medium-horizon sequential dynamics. The 63.2 percent reduction in $RMSE$ by CNN-LSTM climate model compared to bilinear interpolation and near-zero bias (-0.018 C) demonstrates deep learning to be the best approach to downscaling climatic conditions on a regional scale. The operational cost of false negative is much higher than false alarm, which is why the high recall (89.1) of the LSTM heatwave classifier is critical especially in the context of operationally oriented early warning applications. The 7.3-percent accuracy improvement over SVM assures that there are practical benefits in the learning of temporal features using recurrent architectures over their static classifier counterparts. The ResNet-50 waste classifier achieves an accuracy of over 90% (weighted $F1 = 0.902$), which confirms the use of transfer learning as the best approach to apply with small waste image sets. The poorer results in the 'Trash' category ($F1 = 0.840$) is an inherent difficulty: miscellaneous waste does not have a coherent visual signature, which is also found in the literature of various independent studies. Competitive performance of EfficientNet-B0 (88.7) with only 5.3M parameters (as compared to 23.6M in ResNet-50) makes it

the architecture of choice when required to run on an edge.

9. Conclusion and Future Work

In this paper, the entire design, implementation and experimental analysis of an integrated machine learning framework has been described to attain sustainable development. The framework shows that AI-driven solutions are able to generate quantifiably high performance in energy efficiency ($R^2 = 0.9995$), climate modeling (RMSE 0.632), and waste classification ($F1 = 0.902$) -levels of performance that are adequate to deploy AI-based solutions in the real-world in smart grid management, climate early warning systems, and automated recycling centers. Major technical contributions include: reproducible data engineering pipeline on the UCI household dataset; comparing Random Forest with LSTM on energy forecasting; a CNN-LSTM to ERA5 climate downscaling; transfer learning ResNet-50 on waste classification; and explainability using SHAP/LIME to validate scientifically the model. The framework will further be applied in future work to Transformer-based forecasting (PatchTST, Temporal Fusion Transformer), Physics-Informed Neural Networks to model climates, federated learning to classify waste across numerous facilities, and multi-objective reinforcement learning to optimize energy, carbon, and waste measures. Containerized deployment through Docker/Kubernetes with MLOps pipelines to retrieve automated retraining of models is the operational route to production deployment in sustainable infrastructure management.

References

- [1] Kaifeng Bi, et al. 'Accurate medium-range global weather forecasting with 3D neural networks.' *Nature*, 619, 2023.
- [2] Atit Bashyal, Tina Boroukhian. 'Multi-agent deep reinforcement learning based demand response.' *Applied Energy*, 392, 2025.
- [3] Ricardo Vinuesa, et al. 'The role of artificial intelligence in achieving the Sustainable Development Goals.' *Nature Communications*, 2020.
- [4] Bingbing Fang, et al. 'Artificial intelligence for waste management in smart cities: a review.' *Environ. Chemistry Letters*, 21, 2023.
- [5] Gatti A., et al. 'Toward Greener Smart Cities: ML-Based Algorithms for Smart Bin Collection.' *Electronics*, 13(5), 836, 2024.
- [6] Prabadevi Boopathy, et al. 'Deep learning for intelligent demand response and smart grids.' *Computer Science Review*, 2024.
- [7] Biswajit Biswal, et al. 'Review on smart grid load forecasting using ML and DL techniques.' *Energy Reports*, 12, 2024.
- [8] David Weinberg, et al. 'A Review of RL for Controlling Building Energy Systems.' *Sustainable Cities and Society*, 89, 2023.
- [9] Zhong X., et al. 'End-to-End Deep RL Control for HVAC Systems in Office Buildings.' *Designs*, 6(3), 52, 2022.
- [10] He K., Zhang X., Ren S., Sun J. 'Deep Residual Learning for Image Recognition.' *CVPR*, 2016.
- [11] Hochreiter S., Schmidhuber J. 'Long Short-Term Memory.' *Neural Computation*, 9(8), 1735–1780, 1997.
- [12] Lundberg S.M., Lee S.I. 'A Unified Approach to Interpreting Model Predictions.' *NeurIPS*,

2017.

- [13] Burke, M., Lobell, D. B., & Ermon, S. (2016). Using satellite imagery and machine learning to predict poverty. *Science*, 353(6301), 790–794.
- [14] Caruana, R., Lou, Y., Gehrke, J., Koch, P., Sturm, M., & Elhadad, N. (2015). Intelligible models for healthcare: Predicting pneumonia risk and hospital readmission. *Proceedings of the 21st ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1721–1730.
- [15] Chen, M., Mao, S., & Liu, Y. (2014). Big data: A survey. *Mobile Networks and Applications*, 19(2), 171–209.
- [16] Daly, H. E. (1996). *Beyond growth: The economics of sustainable development*. Beacon Press.
- [17] Dong, B., Wang, Z., Rahman, S. M., & Vega, R. (2019). Smart building energy management using artificial intelligence. *Building and Environment*, 158, 268–280.
- [18] Ellen MacArthur Foundation. (2019). *Artificial intelligence and the circular economy*. Ellen MacArthur Foundation.
- [19] Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., & Thrun, S. (2017). Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542(7639), 115–118.
- [20] European Commission. (2021). *Proposal for a regulation laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. European Commission.
- [21] Fang, X., Misra, S., Xue, G., & Yang, D. (2012). Smart grid—The new and improved power grid. *IEEE Communications Surveys & Tutorials*, 14(4), 944–980.
- [22] Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People—An ethical framework for a good AI society. *Minds and Machines*, 28(4), 689–707.
- [23] Gandomi, A., & Haider, M. (2015). Beyond the hype: Big data concepts, methods, and analytics. *International Journal of Information Management*, 35(2), 137–144.
- [24] Geissdoerfer, M., Savaget, P., Bocken, N. M. P., & Hultink, E. J. (2017). The circular economy—A new sustainability paradigm? *Journal of Cleaner Production*, 143, 757–768.
- [25] Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
- [26] Gubbi, J., Buyya, R., Marusic, S., & Palaniswami, M. (2013). Internet of Things (IoT): A vision, architectural elements, and future directions. *Future Generation Computer Systems*, 29(7), 1645–1660.
- [27] Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., & Pedreschi, D. (2018). A survey of methods for explaining black-box models. *ACM Computing Surveys*, 51(5), 1–42.
- [28] IPCC. (2021). *Climate change 2021: The physical science basis*. Cambridge University Press.
- [29] Janssen, M., Charalabidis, Y., & Zuiderwijk, A. (2017). Benefits, adoption barriers and myths of open data and open government. *Information Systems Management*, 29(4), 258–268.
- [30] Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399.
- [31] Kamilaris, A., & Prenafeta-Boldú, F. X. (2018). *Deep learning in agriculture: A survey*. *Computers and Electronics in Agriculture*, 147, 70–90.
- [32] LeCun, Y., Bengio, Y., & Hinton, G. (2015). *Deep learning*. *Nature*, 521(7553), 436–444.
- [33] Mitchell, T. M. (1997). *Machine learning*. McGraw-Hill.
- [34] Rolnick, D., Donti, P. L., Kaack, L. H., Kochanski, K., Lacoste, A., Sankaran, K., Ross, A.

- S., Milojevic-Dupont, N., Jaques, N., Waldman-Brown, A., Luccioni, A., Maharaj, T., Sherwin, E., Mukkavilli, S. K., Kording, K., Gomes, C., Ng, A., Hassabis, D., Platt, J., ... Bengio, Y. (2019). Tackling climate change with machine learning. *ACM Computing Surveys*, 55(2), 1–96.
- [35] Vinuesa, R., Azizpour, H., Leite, I., Balaam, M., Dignum, V., Domisch, S., Felländer, A., Langhans, S. D., Tegmark, M., & Nerini, F. F. (2020). The role of artificial intelligence in achieving the Sustainable Development Goals. *Nature Communications*, 11(1), 233.