

Article

Not peer-reviewed version

Generative AI-Enhanced Digital Twins for Predictive Ecosystem Management and Conservation

[Pablo Vicente-Martínez](#) , [Adrián Chust-Ros](#) , Ismerai David Gutiérrez-Rodríguez , [Emilio Soria-Olivas](#) , [María Ángeles García-Escrivà](#) * , [Edu William-Secin](#)

Posted Date: 29 May 2026

doi: 10.20944/preprints202605.2046.v1

Keywords: ecosystem management; digital twin; generative AI; predictive ecology; environmental decision support; machine learning; explainable AI



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC, OpenAlex.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Generative AI-Enhanced Digital Twins for Predictive Ecosystem Management and Conservation

Pablo Vicente-Martínez ¹, Adrián Chust-Ros ¹, Ismerai David Gutiérrez-Rodríguez ¹, Emilio Soria-Olivas ², María Ángeles García-Escrivà ^{3,*} and Edu William-Secin ⁴

¹ SPV Scala, Gran Canaria, 35100 San Bartolomé de Tirajana, Spain

² Intelligent Data Analysis Laboratory (IDAL), Department of Electronic Engineering, Universitat de València, 46022 Valencia, Spain

³ Fundación Canaria Living Lab, 35017 Las Palmas de Gran Canaria, Spain

⁴ Department of Economics and Business Management, Institute of Tourism and Sustainable Development (TIDES), Universidad de Las Palmas de Gran Canaria, 35017 Las Palmas de Gran Canaria, Spain

* Correspondence: coordinacionit@canariaslivinglab.org

Abstract

The escalating impacts of climate change and anthropogenic pressures on vulnerable ecosystems demand digital tools that make advanced modeling more accessible to conservation practitioners. This study presents a TRL-4 prototype that integrates a configurable Digital Twin core with a generative AI conversational interface for conservation-oriented modeling in Doñana National Park, Spain, a UNESCO World Heritage site facing significant environmental challenges. The main contribution is not the training of specific ecological forecasting models, but the validation of an end-to-end workflow that allows users to configure, execute, inspect, and interpret a predictive system through natural language. The architecture connects a structured YAML configuration, heterogeneous environmental and biological datasets, automated machine-learning training, database-backed traceability, dashboard visualization, and SHAP-based interpretability. Through representative executions, the prototype demonstrates that non-technical users can select target and explanatory variables, configure preprocessing options, launch model training, generate predictions, and review their outputs without directly editing configuration files or running code. Although the predictive metrics obtained in selected runs remain preliminary and should be interpreted as diagnostics rather than evidence of general forecasting skill, the results show that conversational Digital Twins can substantially reduce technical barriers to ecological modeling. By combining generative AI, cloud infrastructure, reproducible machine-learning workflows, and explainable AI, the proposed architecture provides a strong foundation for future conservation decision-support systems that augment expert judgment while preserving human oversight, transparency, and critical interpretation.

Keywords: ecosystem management; digital twin; generative AI; predictive ecology; environmental decision support; machine learning; explainable AI

1. Introduction

1.1. Context and Motivation

Protected ecosystems worldwide face unprecedented pressures from climate change, habitat fragmentation, and intensifying anthropogenic activities [1,2]. The transition from reactive monitoring to proactive, predictive management has become imperative for biodiversity conservation [3,4]. Doñana National Park in southwestern Spain exemplifies this challenge: as a UNESCO World Heritage site and critical wetland ecosystem hosting over 300 bird species, it experiences severe water stress, illegal groundwater extraction, and climate-driven hydrological variability [5,6]. Traditional management approaches, relying on periodic field surveys and retrospective analysis, are insufficient to anticipate ecosystem state transitions or evaluate intervention scenarios before implementation [7,8].

The ecological complexity of such systems (characterized by nonlinear dynamics, threshold effects, and multi-scale interactions) demands sophisticated modeling frameworks capable of integrating diverse data streams (meteorological, hydrological, biological) and generating actionable forecasts [9]. However, the technical knowledge required to construct and operate such models often limits their practical implementation in conservation contexts [10,11].

1.2. State of the Art

1.2.1. Ecosystem Modeling and Digital Twins

Digital Twins (DTs), originating in engineering and manufacturing [12], have recently emerged as promising tools for environmental science [13]. A Digital Twin is a virtual representation of a physical system that continuously integrates real-world data to simulate current states and predict future conditions [14]. In ecology, DTs have been proposed for forest management, urban ecosystems, and marine environments, enabling scenario exploration and adaptive decision-making.

Recent work has demonstrated the feasibility of coupling process-based ecological models with data-driven machine learning (ML) approaches to improve predictive accuracy [15]. Hybrid models combining mechanistic understanding with statistical learning can capture both known ecological processes and emergent patterns from observational data [16]. Time-series forecasting techniques, including Autoregressive Integrated Moving Average (ARIMA), Seasonal ARIMA (SARIMA), and ensemble methods such as Random Forest and Gradient Boosting (e.g., XGBoost), have shown efficacy in predicting ecological variables like species abundance, water levels, and phenological events [17,18].

1.2.2. Machine Learning for Biodiversity Monitoring and Prediction

The application of ML to ecological forecasting has accelerated in recent years [19,20]. Supervised learning algorithms have been successfully applied to predict species distributions [21], population dynamics [22], and habitat suitability under climate scenarios [23]. Tree-based models (Decision Trees, Random Forests) offer interpretability alongside predictive power, which is critical for gaining stakeholder trust in conservation applications [24]. Boosting algorithms like XGBoost have demonstrated superior performance in complex, high-dimensional ecological datasets [25].

Despite these advances, a persistent gap exists between model development in research settings and operational deployment by environmental practitioners [10,26]. Conservation managers often lack the technical expertise or computational resources to configure, execute, and interpret advanced predictive models [27,28].

1.2.3. Generative AI for Scenario Analysis and Decision Support

The recent emergence of Large Language Models (LLMs) and generative AI has introduced new possibilities for human-computer interaction in scientific workflows [29,30]. Generative AI systems can interpret natural language queries, translate them into structured tasks, and synthesize outputs in accessible formats [31,32]. In environmental science, applications have begun to explore generative AI for literature synthesis [33], data annotation [34], and educational tools [35].

However, the integration of generative AI with predictive ecological models, specifically to enable non-expert users to configure simulation scenarios conversationally, remains largely unexplored [36]. This represents a significant opportunity: by bridging the gap between complex modeling infrastructures and intuitive natural language interfaces, generative AI could democratize access to advanced decision-support tools in conservation.

1.3. Research Gap and Proposed Solution

While Digital Twins and ML-based predictive models offer substantial potential for ecosystem management, their practical utility is constrained by accessibility barriers. Configuring simulations typically requires expertise in programming, model parameterization, and data preprocessing; skills not commonly held by field ecologists, park managers, or policy-makers [10,11]. Existing Environmen-

tal Decision Support Systems (EDSS) often present steep learning curves or rigid interfaces that do not accommodate the exploratory, iterative nature of conservation planning [37,38].

This study proposes an architecture that addresses these limitations by integrating a generative AI conversational interface with a configurable Digital Twin core. The system enables users to express modeling intentions in natural language, such as selecting a target variable, changing preprocessing options, choosing predictors, launching a training process, or generating predictions from an existing model. The conversational agent maps these requests to a structured YAML configuration, invokes predefined backend tools, and presents results through a dashboard-oriented workflow. In this paper, the emphasis is on the accessibility, traceability, and interpretability of this configuration process.

1.4. Contributions

This work makes the following key contributions:

1. **Conversational configuration of ecological modeling workflows:** We present a generative AI interface that allows users to configure the core of a conservation-oriented digital twin. It allows to inspect, modify and validate, and persist a structured YAML file to configure the system without directly editing code or complex files.
2. **Configurable digital twin pipeline:** We describe a modular prototype that connects data preprocessing, target and feature selection, model training, prediction generation, traceability metadata, and database-backed visualization in a reproducible workflow.
3. **Explainability for conservation decision support:** We incorporate SHAP-based feature attribution for the trained models, enabling users to complement predictions with interpretable information about the variables that most influence model outputs.
4. **TRL-4 feasibility assessment in a National Park:** We evaluate the system as a laboratory prototype applied to a real conservation context, focusing on functional capability, agent-assisted configuration, and qualitative usefulness rather than claiming operational predictive performance.

1.5. Paper Structure

The remainder of this paper is organized as follows: Section 2 describes the study area, system architecture, configurable modeling pipeline, conversational interface, and prototype infrastructure. Section 3 presents a lightweight evaluation of the agent-assisted configuration workflow, selected model training and prediction capabilities, SHAP-based interpretability, and traceability mechanisms. Section 4 discusses the implications for accessible conservation modeling, the value of explainability, and the limitations of the current TRL-4 prototype. Section 5 summarizes the main contributions and outlines future work required for operational validation.

2. Materials and Methods

2.1. Study Area: Doñana National Park

Doñana National Park is located in southwestern Spain, at the estuary of the Guadalquivir river. Covering approximately 54,252 hectares of marshland, dunes, and Mediterranean scrubland, Doñana is recognized as a UNESCO World Heritage Site and Biosphere Reserve, hosting one of Europe's most important wetland ecosystems [5]. The park serves as a critical stopover for migratory birds along the East Atlantic Flyway, supporting over 300 avian species and multiple threatened populations of the Spanish Imperial Eagle (*Aquila adalberti*) and the Iberian Lynx (*Lynx pardinus*) [39].

Doñana faces multiple conservation challenges: intensive agriculture in surrounding areas drives illegal groundwater extraction, depleting aquifers that sustain the marsh [40]. Climate change has contributed to a decline in precipitation across much of the Iberian Peninsula, particularly during the summer months, thereby exacerbating drought risk and water scarcity in southern Mediterranean regions [41]; and anthropogenic pressures from tourism, infrastructure development, and pollution threaten ecosystem integrity. These stressors make Doñana an ideal testbed for predictive conserva-

tion technologies, as timely interventions informed by robust forecasting can significantly influence management outcomes [6].

2.2. System Architecture Overview

The proposed architecture integrates four core modules (Figure 1): (1) **Data ingestion module**, responsible for acquiring and preprocessing heterogeneous environmental data; (2) **Predictive engine (digital twin core)**, which trains and deploys machine learning models to forecast ecosystem indicators; (3) **Conversational AI interface**, enabling natural language-driven configuration of simulation scenarios; and (4) **Decision support dashboard**, visualizing predictions and supporting adaptive management workflows. Each module operates as an independently deployable microservice, ensuring modularity, fault isolation, and scalability [42].

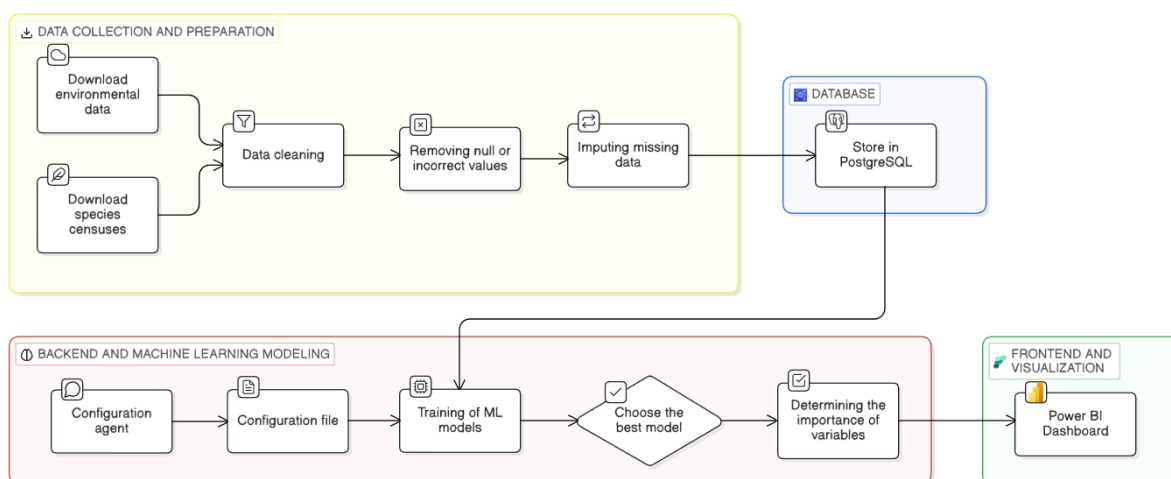


Figure 1. System architecture diagram showing data flow from ingestion through storage and predictive modeling to visualization. The Conversational AI Interface interprets user requests, the digital twin core trains ML models, and Power BI presents interactive dashboards. All components communicate via RESTful APIs, with data persistence in a PostgreSQL database.

2.3. Data Ingestion and Processing

2.3.1. Data Sources

The system integrates heterogeneous environmental, biological, and socio-demographic data sources related to Doñana National Park and its surrounding municipalities. The selection of variables was designed not only to support predictive modeling experiments, but also to increase the visibility of emblematic species and conservation challenges associated with one of Europe's most important protected wetlands.

1. **Meteorological and environmental data:** Historical climate records were obtained from the Spanish State Meteorological Agency (AEMET), including variables such as precipitation, temperature, relative humidity, and wind speed from monitoring stations within and around Doñana [43]. These data were complemented with measurements from the ICTS Doñana monitoring infrastructure [44], including maximum, mean, and minimum values for air pressure, relative humidity, air temperature, accumulated hail, hail duration, accumulated rainfall, rainfall duration, soil temperature, water level, and water temperature.
2. **Waterbird census data:** Long-term bird census data were collected from the ICTS Doñana infrastructure [45] for a set of iconic waterbird species selected for their ecological relevance and their potential to communicate conservation challenges to a broader audience. The selected species were Northern Pintail (*Anas acuta*), Greylag Goose (*Anser anser*), Black Stork (*Ciconia nigra*), Black-tailed Godwit (*Limosa limosa*), Eurasian Wigeon (*Mareca penelope*), Greater Flamingo (*Phoenicopterus roseus*), Eurasian Spoonbill (*Platalea leucorodia*), and Glossy Ibis (*Plegadis falcinellus*).

These species provide a recognizable and policy-relevant biological layer for evaluating how the Digital Twin can support conservation-oriented exploration [5,46].

3. **Iberian lynx population data:** Population indicators for the Iberian lynx (*Lynx pardinus*) in Doñana were extracted from EBD-CSIC reports, including total individuals, reproductive females, and cubs born. Unlike many species in Doñana that show declining or unstable trends, the Iberian lynx is subject to an active reintroduction and recovery programme, which has produced an upward population trend. Including this species therefore allows the system to represent a contrasting conservation trajectory and to test how the Digital Twin handles species with different ecological and management dynamics.
4. **Socio-demographic data:** Population data for municipalities surrounding the park were obtained from the Spanish National Statistics Institute (INE) [47]. The selected municipalities were Almonte, Hinojos, La Puebla del Río, and Aznalcázar. These data were incorporated to complement environmental variables and to explore whether nearby human population size or population density, calculated using municipal surface area, may be associated with changes in the abundance of selected species.

2.3.2. Data Preprocessing Pipeline

Raw data were obtained from heterogeneous sources and formats, including tabular datasets, spreadsheets, and published technical reports. The preprocessing pipeline, implemented in Python 3.11, performs the following operations:

Step 1: Data extraction and standardization. All datasets are converted to a common tabular structure with standardized column names, date formats, temporal references, and measurement units. While most meteorological, environmental, bird census, and socio-demographic data were available in structured tabular formats, the Iberian lynx population indicators had to be manually extracted from EBD-CSIC reports published since 2000. This manual extraction step was necessary because the relevant information was embedded in document-based reports rather than consistently available as a machine-readable time series.

Step 2: Temporal harmonization. The integrated dataset combines variables with different temporal granularities, including daily environmental measurements, periodic bird censuses, annual socio-demographic records, and irregularly reported lynx population indicators. These sources are harmonized into modeling-ready time series according to the temporal frequency selected in the Digital Twin configuration.

Step 3: Missing data imputation. To test the Digital Twin and enable model training across heterogeneous sources, missing values were imputed using linear interpolation. This was particularly important for variables that were not available annually, such as some Iberian lynx population indicators extracted from reports. The resulting gaps made it difficult to apply time-series models such as ARIMA directly without first generating continuous experimental series. Imputed values were therefore used to support prototype testing and should be interpreted as part of the modeling workflow validation, not as replacement for complete observational records.

Step 4: Feature engineering. Additional variables are derived to support predictive experiments. Population density measures for nearby municipalities were calculated to enrich the dataset with complementary socio-demographic information that may help explain variations in the selected conservation variables. In contrast, lag features are generated only for the target variable selected in the user-defined configuration during the configuration session. These lagged values are intended to help non-time-series-specific models, such as Random Forest and XGBoost, incorporate temporal dependencies when predicting the target variable.

Step 5: Data persistence. Cleaned and feature-enriched data are stored in a PostgreSQL relational database hosted on AWS RDS, with tables organized by data type, including climate and environmental variables, biological census records, Iberian lynx population indicators, socio-demographic variables, model metadata, and prediction outputs [48].

2.4. Predictive Engine: Digital Twin Core

2.4.1. Model Library

The Digital Twin's predictive capability resides in a modular Python library encapsulating seven supervised learning algorithms commonly used for time-series forecasting:

1. Moving average (MA): A baseline statistical model computing the arithmetic mean over a sliding temporal window (default: 30 days). Despite its simplicity, MA provides a robust benchmark for evaluating more complex models [49].

2. Linear regression (LR): Ordinary least squares regression modeling the target variable as a linear combination of predictors. LR offers interpretability through coefficient inspection, but assumes linearity and homoscedasticity [50].

3. Decision tree regressor (DTR): A non-parametric model partitioning feature space via recursive binary splits. DTRs capture non-linear relationships and interactions but are prone to overfitting [51].

4. Random forest regressor (RFR): An ensemble of decision trees trained on bootstrap samples with random feature subsets at each split. RFR reduces variance and improves generalization through averaging [52].

5. XGBoost regressor (XGB): An optimized gradient boosting algorithm that sequentially trains trees to correct residuals of prior models. XGBoost employs regularization (L1/L2 penalties) to prevent overfitting and supports parallel computation [53].

6. Autoregressive integrated moving average (ARIMA): A univariate time-series model capturing temporal autocorrelation through autoregressive (AR) and moving average (MA) components, with differencing (I) to achieve stationarity [54]. Model orders (p, d, q) are selected via Akaike Information Criterion (AIC) minimization.

7. Seasonal ARIMA (SARIMA): An extension of ARIMA incorporating seasonal patterns through additional AR, MA, and differencing terms at the seasonal period (e.g., annual cycle: $s=365$ days) [49]. SARIMA is particularly effective for variables exhibiting strong seasonality, such as waterbird abundance.

All models are implemented using `scikit-learn 1.3` (LR, DTR, RFR), `xgboost 2.0` (XGB), and `statsmodels 0.14` (ARIMA, SARIMA) [53,55,56].

2.4.2. Training and Model Selection

Training proceeds as follows:

Data splitting. After the target variable is selected (e.g., flamingo count), the dataset is chronologically divided into training, validation, and test sets to respect temporal ordering and avoid data leakage [57]. The default percentages are 70%, 15% and 15%, although they are configurable by the user through the agent.

Cross-validation. Time-series cross-validation (expanding window) is applied within the training set to tune hyperparameters and assess stability [58]. Five folds are used, with each fold adding one year of data to the training window.

Model training. All seven algorithms are trained on the same preprocessed features. For univariate models (ARIMA, SARIMA, MA), only the target variable's historical values and temporal covariates are used. For multivariate models (LR, DTR, RFR, XGB), all available features (climate, hydrology, lagged variables) serve as predictors.

Performance evaluation. Models are evaluated on the validation set using three metrics:

- Mean absolute error (MAE): $MAE = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i|$
- Root mean squared error (RMSE): $RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2}$
- Mean absolute percentage Error (MAPE): $MAPE = \frac{100\%}{n} \sum_{i=1}^n \left| \frac{y_i - \hat{y}_i}{y_i} \right|$

where y_i denotes observed values and \hat{y}_i predicted values.

Automatic model selection. The system ranks models by validation RMSE (primary criterion) and MAE (tiebreaker). The best-performing model is serialized (Python `pickle` or `joblib`) and registered with a unique identifier, timestamp, and metadata (hyperparameters, feature list) in the database.

2.4.3. Prediction Generation and Traceability

Once trained, models generate forecasts via two modes:

Historical backcasting: For validation, models predict over the test set, enabling comparison with ground truth to assess generalization error.

Future forecasting: For operational use, models predict future states (e.g., water levels for the next 30 days). Univariate models (ARIMA/SARIMA) require only the forecast horizon. Multivariate models require future values of predictor variables, which need to be provided by the user (assumed precipitation scenarios imputed using climatological normals or generated by other forecast models).

All predictions are timestamped and stored in the database for further analysis.

2.5. Conversational AI Architecture and Configuration Management

The conversational interface employs Google's **Gemini 2.5 Flash Lite**, a large language model optimized for low-latency, high-throughput applications [59]. Gemini 2.5 Flash Lite was selected for its balance between contextual understanding, response generation quality, and computational cost compared with larger variants such as Flash or Pro. The model is accessed through Google Cloud's Vertex AI API [65], enabling compliance with the European Union's General Data Protection Regulation (GDPR) concerning user-submitted data.

The interface is implemented using Chainlit, a Python framework for developing conversational AI applications [60]. Chainlit provides session management, message history persistence, and user interface components supporting rich media outputs such as plots and tables. The backend coordinates interactions among the user, the Gemini model, and the Digital Twin API.

To ensure effective interaction, a structured prompt engineering strategy was developed to facilitate accurate intent recognition and reliable system behavior. The prompt template consists of three main components: (i) a static system prompt defining the agent's role, capabilities, and operational constraints; (ii) dynamic injection of the YAML configuration template, allowing the model to identify editable parameters; and (iii) tool descriptions enabling function calling capabilities. Gemini 2.5 supports structured function invocation through JSON outputs [61], allowing the conversational agent to execute predefined operations such as updating configurations, training new models using the digital twin core, or generating predictions from existing models. Functions are registered using JSON schemas specifying parameter types, constraints, and valid inputs.

The system maintains a YAML configuration file containing the parameters required for digital twin operation, including data preprocessing options (e.g., imputation methods, scaling procedures, and outlier detection strategies), target variables, predictive feature variables, time-series frequency (daily, monthly, or annual), and the set of machine learning models used during the training phase. Users can modify these configurations conversationally through the interface. When an update request is received, the agent validates the proposed modification by checking parameter existence and value constraints, presents the updated configuration to the user, and requests explicit confirmation before applying persistent changes. This human-in-the-loop design reduces the risk of unintended modifications caused by misinterpretation or ambiguous instructions [62].

2.6. Decision Support Dashboard and User Interaction Workflow

Predictions and model output are visualized in **Microsoft Power BI**, integrated through **Microsoft Fabric** for cloud-hosted reporting [63]. The dashboard connects securely to the AWS RDS PostgreSQL database through a gateway that enables multiple daily data updates. This infrastructure allows users to analyze generated predictions alongside historical observations and supports data-driven decision making when evaluating alternative conservation measures.

The dashboard organizes information into several thematic report views according to data type and modeling outputs. Environmental and climate data visualizations provide temporal analyses of meteorological and environmental variables obtained from AEMET and the ICTS Doñana monitoring infrastructure. These include precipitation, air temperature, relative humidity, air pressure, soil temperature, water temperature, and water level measurements, enabling users to inspect environmental trends and seasonal patterns.

Additional views focus on species population and census data for conservation-relevant species. These include monitored waterbird populations from ICTS Doñana and Iberian lynx indicators extracted from EBD-CSIC reports, such as total individuals, reproductive females, and cubs born. The dashboard enables users to examine long-term historical trends and compare the temporal evolution of different species populations.

The platform also incorporates socio-demographic information from municipalities surrounding Doñana, including Almonte, Hinojos, La Puebla del Río, and Aznalcázar. These visualizations provide contextual information regarding nearby human population dynamics and derived indicators such as population density, which can serve as complementary variables during exploratory modeling tasks.

Model prediction views display outputs generated by the trained machine learning models for selected conservation indicators, particularly species population variables. These visualizations allow users to compare predicted and historical values in an accessible format that facilitates interpretation of model behavior and trends. In the current prototype, these predictions should be interpreted as exploratory outputs of the modeling workflow rather than validated operational forecasts.

The overall interaction workflow combines the conversational interface with the dashboard environment. Users begin by initiating a session through the Chainlit interface, where they receive a welcome message and a summary of the current configuration. Natural-language requests can then be submitted to configure predictive models, for example by specifying target species and selecting environmental or socio-demographic explanatory variables. The conversational agent interprets the request, proposes the corresponding configuration modifications, and requests explicit confirmation before persisting any changes. Once confirmed, the backend triggers model training through the Digital Twin API. After training is completed, the selected model, associated identifier, and diagnostic metrics are returned to the user through the conversational interface. Users can subsequently generate predictions and inspect the results either directly within the Chainlit interface through inline matplotlib visualizations or through the Power BI dashboard, which serves as the primary reporting platform for formal analysis and presentation purposes.

2.7. Implementation Details and Infrastructure

2.7.1. Prototype Deployment Architecture

The prototype was deployed using a cloud-based architecture intended to validate integration among the conversational interface, backend services, relational storage, and dashboard layer. The backend and Chainlit frontend were containerized and deployed on two services (agent and digital twin engine) on AWS Fargate behind an Application Load Balancer. Environmental data, model metadata, configuration records, and prediction outputs were persisted in PostgreSQL using Amazon RDS. The dashboard layer connected to the PostgreSQL database to support inspection of historical observations and generated predictions.

The deployment should be interpreted as a prototype deployment rather than a production-grade operational infrastructure. Features such as autoscaling policies or long-term archival storage are considered deployment extensions and are not treated as evaluated results in this study.

2.7.2. Software Stack

- **Backend:** Python 3.11, FastAPI 0.120.1, Psycogp2 2.9.10, Pandas 2.3.0, NumPy 1.26.4.
- **Machine learning:** scikit-learn 1.7.0, XGBoost 3.0.2, statsmodels 0.14.4, matplotlib 3.10.3, SHAP 0.49.1.
- **Conversational AI:** Chainlit 2.6.2, Google Generative AI SDK 1.26.0.

- **AI generated configuration validation:** Pydantic 2.11.7.
- **Containerization:** Docker 29.0.2.
- **Database:** PostgreSQL 17.5.
- **Deployment:** AWS Copilot CLI for infrastructure-as-code (IaC) provisioning.

2.7.3. Evaluation Protocol

System validation follows a multi-criteria framework aligned with TRL-4 objectives: validating key functions in a controlled environment before operational deployment [66]. The evaluation focuses on four aspects:

Agent-assisted configuration. The conversational agent is evaluated on its ability to interpret natural-language requests, update the YAML configuration, request clarification when needed, and avoid persisting changes without user confirmation.

Pipeline execution. The Digital Twin core is evaluated through a small number of reproducible training and prediction tasks, used to verify that the configured workflow executes end to end and returns model metadata and outputs.

Explainability. SHAP values are generated for models, allowing users to inspect the relative contribution of input variables to individual predictions or aggregate model behavior [64,68].

Traceability. Training runs and predictions are linked to configuration files, model identifiers, timestamps, and data references, supporting auditability and reproducibility.

3. Results

3.1. Overview of Prototype Evaluation

The evaluation was designed to match the actual maturity and evidence available for the system. The results focus on whether the prototype enables non-technical users to configure and execute modeling workflows through a conversational interface. The validation therefore covers four dimensions: agent-assisted configuration, execution of selected training and prediction tasks, SHAP-based interpretation, and traceability of outputs.

This framing is consistent with a TRL-4 prototype, demonstrating successful integration of the different system components and validating the functional feasibility of the proposed architecture within a controlled environment. The platform provides a robust foundation for interactive conservation-oriented digital twin workflows in Doñana, enabling model training, prediction generation, and integrated data exploration through conversational and dashboard-based interfaces. In this context, the reported quantitative metrics should be interpreted as representative diagnostics of the implemented modeling workflows and as evidence of the platform's capability to support exploratory ecological analysis and decision-support tasks.

3.2. Agent-Assisted Configuration Workflow

The main functional result of the prototype is the ability to use natural language to configure a modeling workflow that would otherwise require direct manipulation of a YAML file. The configuration file defines the input dataset, date column, target variable, numerical predictors, train-test split, preprocessing options, lagged features, correlation and VIF filters, scaling choices, prediction horizon, enabled estimators, output directory, and database connection. By injecting this structure into the conversational agent prompt, the system constrains the interaction to parameters that the Digital Twin core can actually use.

Figure 2 illustrates how the agent exposes the active configuration to the user. This view makes the modeling assumptions inspectable, including the selected target variable, numerical predictors, preprocessing options, lagged features, enabled estimators, and output settings.

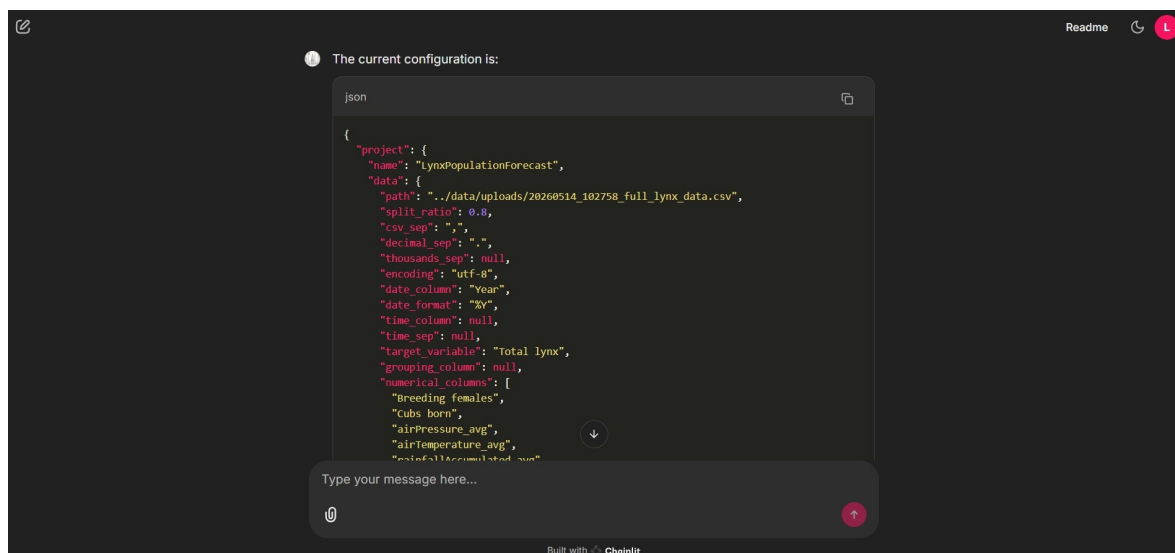


Figure 2. Conversational display of the active Digital Twin configuration, including data, preprocessing, model, and output parameters.

The agent supports the following operations:

- Displaying the current configuration and explaining editable fields;
- Updating one or more configuration parameters, such as the target variable, feature set, time frequency, preprocessing method, or enabled models;
- Checking whether requested values are compatible with the known configuration structure;
- Asking for clarification when the request is ambiguous;
- Saving changes only after explicit user confirmation;
- Launching model training from the active configuration;
- Generating predictions from a trained model when the required model identifier and input information are available;
- Reporting tool errors directly to the user instead of hiding failed execution.

A lightweight evaluation set of natural-language tasks was used to measure the agent's behavior. The task set covers configuration inspection, single-parameter edits, multi-parameter edits, model selection, training requests, prediction requests, and ambiguous instructions. The measured metrics are task completion rate, number of clarification turns, and response time. These metrics evaluate the agent as a useful configuration interface.

Table 1. Evaluation results of the conversational configuration agent. Each task was tested with 20 different prompts.

Task category	Metric	Interpretation
Configuration inspection	Success rate: 90.0%, latency: 2.3 s	Ability to expose the current YAML state to non-technical users.
Single or multiple parameter updates	Correct update rate: 95.0%, clarification turns: 1.5 avg.	Ability to map natural language onto valid configuration fields.
Training requests	Execution success: 90.0%, returned meta-data completeness: 100.0%	Ability to trigger the Digital Twin core and report model identifiers and diagnostics.
Prediction requests	Execution success: 90.0%, required inputs requested: 100.0%	Ability to route prediction requests according to model type and available data.
Error handling	Invalid changes blocked: 100.0%, errors reported: 100.0%	Ability to preserve human control and avoid silent failures.

3.3. Representative Configuration Scenario

Figure 3 illustrates the kind of interaction supported by the system. The first example consists of a full configuration workflow. Its purpose is to show how the agent translates a user request into structured parameters, confirms the relevant assumptions, and triggers backend tools.

This interaction exemplifies the system's core value proposition: enabling non-technical users to rapidly explore "what-if" scenarios without writing code or navigating complex GUIs. The agent's confirmatory prompts (parameter review before execution) implement human-in-the-loop safeguards, reducing the risk of unintended simulations [62].

User: I want to train a model to predict the number of sighted individuals of the species *Anas acuta* using precipitation, temperature, and the observations from the previous 2 years.

Agent: I can configure that workflow. I have updated the target variable, set the selected numerical predictors, configured the lagged features, and used the estimators accumulated rainfall, rainfall duration and soil temperature. Before saving, please confirm the configuration changes.

User: Save the configuration and train the model.

Agent: Training has been successfully completed. The model with the best results is: Model name: RandomForestRegressor ID: 24 Metrics: MAE: 1.05 MAPE: 35.70 RMSE: 1.22 Is there anything else I can help you with?

User: Generate predictions for the next year using the model with ID 24.

Agent: The predictions for next year are:

2025-01-01T00:00:00: 89

To provide greater clarity in analyzing the results, a dashboard powered by these predictions is available and updated several times a day: [Power BI link]

Figure 3. Representative user-agent dialogue for conversational configuration, training, and prediction. The example demonstrates the workflow and human-in-the-loop confirmation pattern, not a validated conservation forecast.

A real execution of this workflow is shown in Figure 4. The interaction demonstrates how the agent reports the selected model and diagnostic metrics, along with the generated prediction and an inline plot in the same conversational environment.

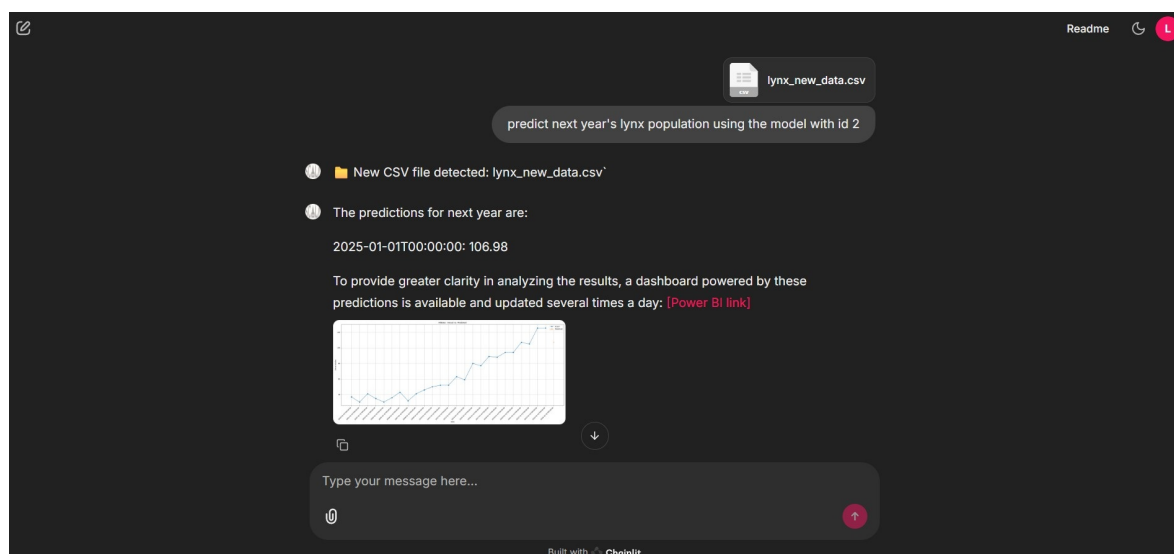


Figure 4. Example of a workflow result in which the agent launched model training, reported the selected model and diagnostic metrics, and generated a one-year-ahead prediction.

Figure 5 shows an inline prediction generated by the agent for the Iberian lynx population indicator. This example was selected because the Iberian lynx represents a high-profile conservation case in Doñana and provides a contrasting trajectory relative to many other monitored species, due to ongoing recovery and reintroduction efforts.

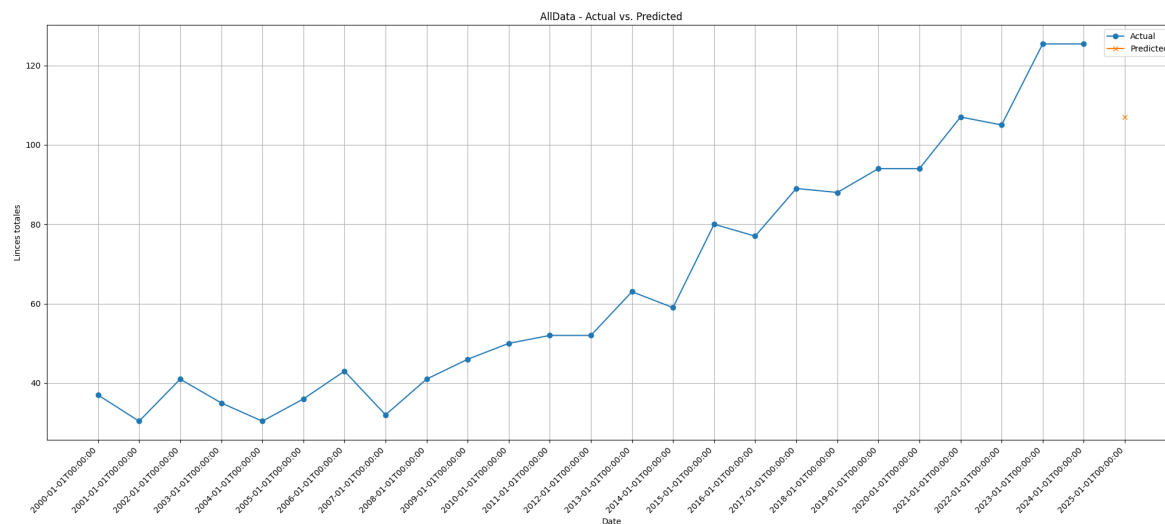


Figure 5. Inline prediction generated by the conversational agent for an Iberian lynx population indicator. The figure illustrates the prototype’s ability to return both numerical and visual outputs directly within the conversational workflow.

3.4. Selected Model Training and Prediction Runs

Table 2 summarizes several representative training runs executed using the Digital Twin core. These examples include both mammal and bird population indicators and illustrate the system’s ability to process the configured variables, train the enabled estimators, select a suitable model, and return diagnostic metrics to the user through the conversational interface. The reported errors remain too high to support operational forecasting, particularly given the limited availability, heterogeneity, and temporal sparsity of some biological datasets. Nevertheless, the executions validate the main technical objective of the prototype: enabling users to configure, launch, and inspect predictive modeling workflows in a reproducible and accessible way. The results therefore support the system’s value as a functional decision-support prototype while identifying predictive accuracy as a priority for future work.

Table 2. Example of model training runs executed through the Digital Twin pipeline.

Target variable	Training configuration	Results
Iberian lynxes individuals	Features: target lags (1, 2, 5, and 10 years), cubs born, mean temperature, mean precipitation, and population of nearby municipalities.	MAPE: 26.8%; RMSE: 0.80. Selected model: XGBoost.
<i>Anas acuta</i> individuals	Features: target lags (1, 2, 5, and 10 years), mean precipitation, mean water temperature, mean soil temperature, mean water level, mean water temperature, and mean rainfall duration.	MAPE: 54.15%; RMSE: 1.33. Selected model: XGBoost.
<i>Anser anser</i> individuals	Features: target lags (1, 2, 5, and 10 years), mean precipitation, mean water temperature, mean soil temperature, mean water level, mean water temperature, and mean rainfall duration.	MAPE: 36.76%; RMSE: 0.79. Selected model: Decision-TreeRegressor.

Overall, the selected runs confirm that the Digital Twin pipeline can successfully execute end-to-end training workflows for different conservation variables, including both mammal and bird population indicators. In each case, the system was able to process the configured variables, train the enabled estimators, select a suitable model, and return diagnostic metrics to the user through the conversational interface. The obtained errors remain too high to support operational forecasting, particularly given the limited availability, heterogeneity, and temporal sparsity of some biological datasets. Nevertheless, these experiments validate the main technical objective of the prototype: enabling users to configure, launch, and inspect predictive modeling workflows in a reproducible and accessible way.

3.5. SHAP-Based Interpretability

A distinctive capability of the prototype is the integration of SHAP values for all models. SHAP does not prove causal relationships, but it helps users inspect whether model predictions are influenced by variables that are plausible from a conservation perspective. This is especially relevant in ecological decision-support contexts, where a prediction without interpretability can be difficult for practitioners to trust or challenge.

SHAP summaries can be used to identify which input variables contribute most to model outputs across a dataset, while local explanations can show how variables contributed to an individual prediction. In the Doñana use case, this allows users to examine whether variables such as precipitation, temperature, lagged values of the target variable, or hydrological indicators are driving the model behavior. Such interpretation supports model review, discussion with domain experts, and detection of potentially spurious dependencies.

In one representative run of a model predicting lynx population, the highest mean absolute SHAP values corresponded to *lynx population lag 1*, *rainfall accumulated mean*, and *air temperature mean*. These outputs were used only to inspect model behaviour and were not interpreted as causal ecological relationships.

The appropriate conclusion is therefore that SHAP makes model outputs more inspectable and can help conservation practitioners understand and question predictions. Even though, it should not be presented as evidence that the model has learned true ecological mechanisms.

3.6. Traceability and Dashboard Outputs

The prototype stores outputs with metadata that support later inspection. Training and prediction records can be linked to the active configuration, model identifier, timestamp, selected variables, preprocessing options, and result files. This traceability is important because conversational systems can otherwise make it difficult to reconstruct how a result was produced.

The dashboard layer provides a practical interface for reviewing historical data and generated predictions. In the current prototype, dashboard visualizations should be treated as decision-support artifacts for exploration and communication, not as validated operational recommendations. Their value lies allowing users who may not interact directly with Python scripts to review the modeling workflow, model objects, or database tables.

Figure 6 provides an example of the dashboard layer connected to the prototype database. Although the dashboard is not the main research contribution, it supports traceability and communication by allowing users to inspect historical records, generated predictions, and model output outside the conversational interface, allowing the analysis of trends that may affect the evolution of the variables of interest.

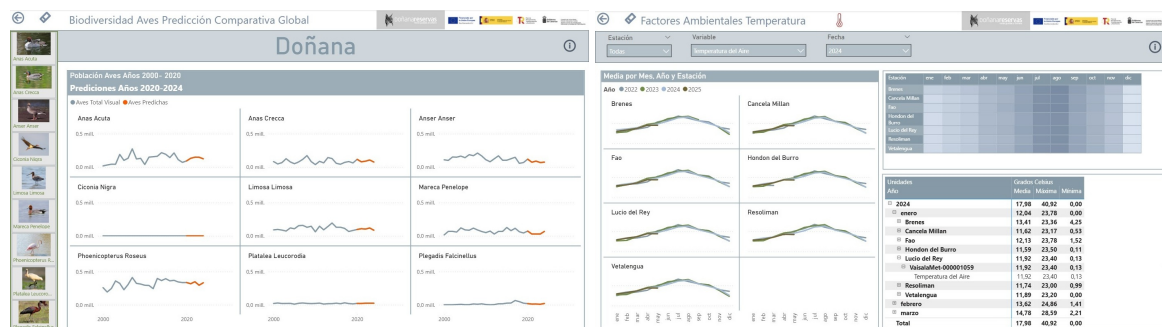


Figure 6. Example Power BI plots connected to the prototype database. The first view shows the population projections generated for the selected waterbird species. The second panel displays historical records of temperature measurements at the different weather stations of the Doñana National Park.

4. Discussion

4.1. Interpretation of Findings

The prototype demonstrates that a generative AI interface can act as a practical configuration layer for a conservation-oriented Digital Twin. Users can interact with the system through natural language while the backend preserves a structured, reproducible configuration. This addresses a real barrier in ecological modeling: many practitioners who could benefit from predictive tools are not comfortable selecting model parameters, running training scripts, or editing configurations using files or complex GUIs.

The evidence demonstrates the capability of the proposed system to support integrated ecological modeling workflows within a coherent and operational TRL-4 prototype. The results highlight the successful integration of conversational configuration, model execution, prediction generation, traceability, and explainability into a unified framework, enabling seamless interaction between users, models, and data sources. This integration establishes a strong foundation for extensible conservation-oriented digital twin systems that can be further developed and deployed in real-world environmental monitoring contexts, including potential future use with real-time sensor data from protected natural reserves. Overall, this section situates the findings within the broader context of conservation technology and outlines their relevance for the advancement of data-driven environmental decision-support systems.

4.2. Conversational Configuration as Access Infrastructure

Traditional Environmental Decision Support Systems often require users to learn fixed interfaces, domain-specific software, or scripting workflows [37,67]. By packaging state-of-the-art predictive models within an intuitive, conversational interface, the proposed system directly addresses usability barriers by allowing users to express modeling intentions in natural language while the agent maps those intentions onto a constrained YAML schema. This is valuable because the configuration file becomes both a machine-readable execution contract and a human-readable representation of modeling assumptions.

The human-in-the-loop design is central to this contribution. The agent can propose changes, but persistence requires confirmation. Tool errors are reported rather than hidden. Ambiguous requests can be routed to clarification. These design choices are especially important in conservation contexts, where model outputs can influence discussions about scarce resources, ecological risk, and management priorities [62].

4.3. The Role of SHAP in Conservation-Oriented Modeling

SHAP-based explanations strengthen the prototype because they shift the interaction from simply asking for predictions to interrogating model behavior. For conservation practitioners, knowing that a model output changed slightly is often less useful than understanding which variables contributed to that change. Feature attributions can support expert review, reveal unexpected dependencies, and help communicate model behavior to stakeholders.

However, SHAP must be used carefully. It explains the behavior of a trained model, not the causal structure of the ecosystem. A variable with high SHAP importance may be a proxy, a correlated measurement, or a consequence of data collection patterns. For this reason, SHAP outputs should be presented as interpretability aids that support discussion and validation, not as ecological conclusions in themselves [64,68].

4.4. Implications for Ecosystem Management

The system's most immediate implication is not automated conservation decision-making, but improved access to exploratory modeling. A park manager, analyst, or project partner can configure variables, run models, inspect predictions, and review explanatory outputs without directly operating the underlying code. This may accelerate hypothesis exploration, make modeling assumptions more visible, and support communication between technical and non-technical stakeholders.

In future operational settings, such a system could contribute to adaptive management by making use of private real-time sensor data, instead of relying on monthly aggregates or public datasets. The training module could also benefit from incorporating better tools to explore hyperparameter spaces, thus leading to better predictions. At the current maturity level, however, the prototype can be used for experimentation, model training, and methodological development.

4.5. Comparison with Existing Approaches

Compared with conventional modeling workflows, the proposed approach emphasizes configurability and accessibility. Rather than replacing established ecological models or specialized Environmental Decision Support Systems, the conversational layer can be understood as an access mechanism that helps users configure and operate modeling components. This distinction is important: the system does not claim to outperform domain-specific hydrological or ecological models. Its contribution lies in connecting configurable modeling, conversational interaction, database persistence, dashboard visualization, and explainability in one prototype.

Compared with broader Digital Twin literature, this work contributes a concrete conservation-oriented implementation pattern. It contributes to filling the gap in the implementation of digital twins focused on ecosystem conservation by demonstrating feasibility in a complex, multi-stressor ecosystem with heterogeneous data and high conservation stakes. The Digital Twin is not presented as a fully autonomous virtual replica of Doñana, but as a configurable data-driven modeling environment that can be updated, queried, and inspected, always keeping a human as a part of the process. This approach mitigates potential risks associated with automated decision-making in an area as sensitive as the protection of endangered species and natural areas of special importance.

4.6. Current Limitations

Several limitations remain. The prototype has been validated in a controlled setting and has not yet been tested as part of routine park-management workflows. Field trials (TRL-5/6) are necessary to assess performance under real circumstances. The user evaluation is also preliminary and should be expanded with a group of conservation practitioners.

The digital twin engine is versatile and data-driven, but that also means it may capture correlations without representing real ecological mechanisms. It also forecasts single variables independently, ignoring cross-variable dependencies (e.g., how water level affects flamingo counts). Multivariate time-series models (e.g., Vector Autoregression, Neural ODEs) could capture these interactions [71,72]. Moreover, the models capture correlations but do not explicitly encode ecological processes (e.g., nutrient cycling, predator-prey dynamics). Hybrid models combining mechanistic equations with ML could improve interpretability and extrapolation [15,70].

4.7. Future Research Directions

Future work should address the following directions:

- **Expanded predictive validation:** Predictive experiments should be reproduced on selected conservation variables using richer and more complex datasets. This would allow the system to be evaluated against specific conservation problems and would help determine where the current modeling pipeline is useful, where errors remain too high, and which variables or data sources are needed to improve performance.
- **Integration of remote sensing data:** Incorporating satellite-derived data, such as Sentinel-2 vegetation indices or Sentinel-1 soil moisture indicators, could improve spatial resolution and predictive capacity [75,76]. Recent advances in geospatial AI could also support automated feature extraction from imagery [77].
- **End-user validation and participatory design:** The system should be tested with park managers, conservation practitioners, and other end users to validate usability, trust, and practical value. Participatory modeling workshops could help refine the interface, identify real decision-making needs, and complement quantitative predictions with qualitative expertise from local stakeholders.
- **Domain-aware conversational agent:** The agent should be extended with domain-specific knowledge to improve its ability to guide users, interpret modeling outputs, and identify potential limitations in the results. Retrieval-Augmented Generation (RAG) could be used to incorporate technical documents, project documentation, ecological references, and operational guidelines into the agent's responses.
- **Controlled field deployment and long-term monitoring:** Moving beyond TRL-4 will require deployment in a controlled operational setting within Doñana National Park. Longitudinal studies would be needed to assess how the system integrates into daily workflows, how users rely on its outputs over time, and whether it has measurable effects on decision-making processes or conservation outcomes [74].

5. Conclusions

This paper presents a TRL-4 prototype that integrates a configurable Digital Twin core with a generative AI conversational interface for conservation-oriented modeling in Doñana National Park. The main contribution of this work is a practical architecture that makes ecological modeling workflows easier to configure, execute, inspect, and explain. Rather than requiring users to directly edit configuration files or operate machine-learning scripts, the system allows them to interact with a structured modeling pipeline through natural language.

The prototype demonstrates several advances for conservation technology. First, it provides a modular, cloud-native framework based on AWS Fargate, Amazon RDS, Docker, and dashboard-based visualization, supporting reproducibility, scalability, and maintainability. Second, it connects conversational interaction with a structured YAML configuration, enabling users to select target variables, define explanatory variables, configure preprocessing options, launch model training, generate predictions, and review outputs through database-backed dashboards. Third, it incorporates traceability mechanisms that link predictions to input data, model identifiers, and configuration metadata. Fourth, it adds an interpretability layer through SHAP values, allowing users to examine which variables influence model behaviour and supporting a more transparent dialogue between data scientists and conservation practitioners.

These contributions are relevant because they address one of the persistent barriers to the adoption of predictive modeling in conservation: accessibility. The system does not replace ecological expertise, but it reduces the technical friction involved in configuring and exploring models. In doing so, it can help non-technical stakeholders engage more directly with data-driven analyses, compare modeling assumptions, inspect results, and participate in more informed discussions about conservation scenarios.

The system's design philosophy, positioning AI as a decision-support tool rather than an autonomous decision-maker, reflects a commitment to preserving human judgment and professional

expertise while augmenting analytical capacity [62]. Human-in-the-loop safeguards (parameter confirmation prompts, transparency through SHAP feature importance, comprehensive audit trails) operationalize principles of responsible AI deployment in high-stakes environmental contexts [68,69].

As we confront accelerating biodiversity loss and climate disruption, the urgency of equipping conservation practitioners with advanced analytical tools has never been greater. This study offers evidence that conversational Digital Twins can become valuable access infrastructure for conservation modeling. The Doñana prototype is a meaningful step toward making predictive and explainable modeling tools more transparent, configurable, and usable by the people who need them. By combining generative AI, cloud infrastructure, traceable machine-learning workflows, dashboard visualization, and SHAP-based interpretability, the system provides a strong foundation for future decision-support tools that augment expert judgment rather than replace it. Looking ahead, the main challenge lies in ensuring that predictions remain relevant to specific problems and are integrated into existing conservation efforts. In this endeavor, every innovation must be guided by a singular imperative: serving the flourishing of life on Earth.

Author Contributions: Conceptualization, E.S.-O. and E.W.-S.; methodology, A.C.-R., M.Á.G.-E. and P.V.-M.; software, I.D.G.-R. and A.C.-R.; validation, M.Á.G.-E. and P.V.-M.; writing-original draft preparation, A.C.-R. and E.S.-O. All authors have read and agreed to the published version of the manuscript.

Funding: This work has been carried out within the framework of the Spain Living Lab project (Grant Reference 1/1/2024-0412093852—SLLC16-01), funded by the Canarian Agency for Research, Innovation and the Information Society (ACIISI), Department of Universities, Science, Innovation and Culture of the Government of the Canary Islands, under the RETECH Programme, contributing to milestones 251, 252 and 253 of Component 16 of the Recovery, Transformation and Resilience Plan (PRTR), and co-funded by the European Union—Next Generation EU.

Data Availability Statement: The data and code supporting the findings of this study are available from the corresponding author (coordinacionit@canariaslivinglab.org) upon reasonable request.

Conflicts of Interest: Authors Pablo Vicente-Martínez, Adrián Chust-Ros and Ismerai David Gutiérrez-Rodríguez were employed by the company SPV Scala. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Appendix A. System Prompt and Configuration

Appendix A.1. Configuration Agent System Prompt

The following appendix presents the operational instructions used to guide the functional behavior of the dialogue system. These instructions define the agent's persona, authorized operations, and procedural constraints, thereby ensuring coherent and systematic performance across the integrated framework. The original prompt was written in Spanish and has been translated to facilitate comprehension.

Listing A1: Agent System Prompt

```
You are a specialist assistant in managing configuration files for simulation models.
```

```
You can help users:
```

- Check or display the initial or current configuration (with or without changes)
- Update a configuration parameter
- Update multiple configuration parameters
- View the configuration file structure
- Save changes
- Train models. The function trains various models and returns the model with the best results/metrics
- Make predictions using a trained model

- Provide SHAP values or feature-importance explanations for trained models, helping users inspect which variables most influence model outputs.

Instructions:

- Use the tools when necessary.
- Display the current configuration settings to the user when requested.
- Report any changes made and ask the user if they want to save the changes before exiting.
- Save the changes only if the user requests it.
- Report any errors received by the tools.
- When the user asks you to train a model, use the 'train_model_tool' function with the configuration information. Respond with the information it returns (model name, ID, metrics, etc.). If an error occurs, inform the user of the error.
- If the user wants to change the training CSV, ask them to attach the new CSV. You will need to modify the configuration path (project.data.path) to the new path where the attached CSV is saved.
- Make predictions using a trained model. Ask the user for the ID of the model they want to use for predictions. If the model is ARIMA or SARIMA, you only need the user to provide the time units (steps); if it is a different model, you need the user to provide the CSV file. Your response should consist solely of the predictions. Do not show the user the path where the image is saved.
- When providing SHAP values, remind the user that they explain the behavior of the trained model, not causal relationships in the ecosystem, and that results should be interpreted with domain expertise.

Important rules regarding tools:

- You can chain multiple tools together without asking the user for confirmation at each intermediate step. Execute all necessary steps at once and present the final result.
- Be proactive; don't force the user to confirm every detail unless you have serious doubts.

Configuration structure:

```
{config_structure}
```

Use the available tools when necessary and always explain the changes you make.

Appendix A.2. Default Configuration File

The following file template illustrates the default values of the parameters to configure the simulations. This specification establishes the number of agents simulated, the opening hours of the points of interest and other relevant parameters

Listing A2: Digital twin core configuration file

```
project:
  name: Example
  data:
    path: "path/to/file.csv"
    split_ratio: 0.8

  # --- CSV Format Settings ---
  csv_sep: ","
  decimal_sep: "."
  encoding: "utf-8"

  # --- Column Definitions ---
```

```
date_column: "date_column_name"
date_format: "%Y-%m-%d"
target_variable: "target_column_name"
numerical_columns:
  [
    "column_name_1",
    "column_name_2",
    "column_name_3",
  ]

preprocessing:
  aggregation_method: "last"
  missing_value_method: linear_interpolation
  apply_outlier: true

# Feature engineering
lagged_features: [1, 2]
apply_correlation: true
apply_vif: true

# Scaling options
x_scaling: "standard"
y_scaling: "standard"
frequency_options: "D"
n_periods: 30

results:
  output_dir: "path/to/results/"
  filename: "filename.csv"

models:
  estimators:
    default:
      - MovingAverageModel
      - LinearRegression
      - DecisionTreeRegressor
      - RandomForestRegressor
      - XGBRegressor
      - StatsmodelsARIMAModel
      - StatsmodelsSARIMAModel

database:
  postgres_database_name
```

References

1. IPBES. *Global Assessment Report on Biodiversity and Ecosystem Services*; IPBES Secretariat: Bonn, Germany, 2019. DOI: 10.5281/zenodo.3831673.
2. Steffen, W.; Richardson, K.; Rockström, J.; Schellnhuber, H.J.; Dube, O.P.; Dutreuil, S.; Lenton, T.M.; Lubchenco, J. The emergence and evolution of Earth System Science. *Nature Reviews Earth & Environment* **2020**, *1*, 54–63. DOI: 10.1038/s43017-019-0005-6.
3. Jetz, W.; McGeoch, M.A.; Guralnick, R.; Ferrier, S.; Beck, J.; Costello, M.J.; Fernandez, M.; Geller, G.N.; Keil, P.; Merow, C.; et al. Essential biodiversity variables for mapping and monitoring species populations. *Nature Ecology & Evolution* **2019**, *3*, 539–551. DOI: 10.1038/s41559-019-0826-1.

4. González-García, A.; Palomo, I.; González, J.A.; García-Nieto, A.P.; Montes, C.; Martín-López, B. Quantifying spatial supply-demand mismatches in ecosystem services provides insights for land-use planning. *Land Use Policy* **2023**, *124*, 106397. DOI: 10.1016/j.landusepol.2022.106397.
5. Rendón, M.A.; Green, A.J.; Aguilera, E.; Almaraz, P. Status, distribution and long-term changes in the waterbird community wintering in Doñana, south-west Spain. *Biological Conservation* **2008**, *141*, 1371–1388. DOI: 10.1016/j.biocon.2008.03.006.
6. Green, A.J.; Alcorlo, P.; Peeters, E.T.; Morris, E.P.; Espinar, J.L.; Bravo-Utrera, M.A.; Bustamante, J.; Díaz-Delgado, R.; Koelmans, A.A.; Mateo, R.; et al. Creating a safe operating space for wetlands in a changing climate. *Frontiers in Ecology and the Environment* **2017**, *15*, 99–107. DOI: 10.1002/fee.1459.
7. Cabin, R.J. *Intelligent Tinkering: Bridging the Gap Between Science and Practice*; Island Press: Washington, DC, USA, 2018; ISBN 978-1610918688.
8. Pennekamp, F.; Iles, A.C.; Garland, J.; Brennan, G.; Brose, U.; Gaedke, U.; Jacob, U.; Kratina, P.; Matthews, B.; Momo, F.; et al. The intrinsic predictability of ecological time series and its potential to guide forecasting. *Ecological Monographs* **2019**, *89*, e01359. DOI: 10.1002/ecm.1359.
9. Mouquet, N.; Lagadeuc, Y.; Devictor, V.; Doyen, L.; Duputié, A.; Eveillard, D.; Faure, D.; Garnier, E.; Gimenez, O.; Huneman, P.; et al. Predictive ecology in a changing world. *Journal of Applied Ecology* **2015**, *52*, 1293–1310. DOI: 10.1111/1365-2664.12482.
10. Addison, P.F.E.; Rumpff, L.; Bau, S.S.; Carey, J.M.; Chee, Y.E.; Jarrad, F.C.; McBride, M.F.; Burgman, M.A. Practical solutions for making models indispensable in conservation decision-making. *Diversity and Distributions* **2013**, *19*, 490–502. DOI: 10.1111/ddi.12054.
11. Rose, D.C.; Sutherland, W.J.; Amano, T.; González-Varo, J.P.; Robertson, R.J.; Simmons, B.I.; Wauchope, H.S.; Kovacs, E.; Durán, A.P.; Vadrot, A.B.; et al. The major barriers to evidence-informed conservation policy and possible solutions. *Conservation Letters* **2020**, *13*, e12564. DOI: 10.1111/conl.12564.
12. Grieves, M.; Vickers, J. Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems. In *Transdisciplinary Perspectives on Complex Systems*; Kahlen, F.J., Flumerfelt, S., Alves, A., Eds.; Springer: Cham, Switzerland, 2016; pp. 85–113. DOI: 10.1007/978-3-319-38756-7_4.
13. Batty, M. Digital twins. *Environment and Planning B: Urban Analytics and City Science* **2018**, *45*, 817–820. DOI: 10.1177/2399808318796416.
14. VanDerHorn, E.; Mahadevan, S. Digital Twin: Generalization, characterization and implementation. *Decision Support Systems* **2021**, *145*, 113524. DOI: 10.1016/j.dss.2021.113524.
15. Reichstein, M.; Camps-Valls, G.; Stevens, B.; Jung, M.; Denzler, J.; Carvalhais, N.; Prabhat. Deep learning and process understanding for data-driven Earth system science. *Nature* **2019**, *566*, 195–204. DOI: 10.1038/s41586-019-0912-1.
16. Willard, J.; Jia, X.; Xu, S.; Steinbach, M.; Kumar, V. Integrating scientific knowledge with machine learning for engineering and environmental systems. *ACM Computing Surveys* **2022**, *55*, 1–37. DOI: 10.1145/3514228.
17. Moriasi, D.N.; Gitau, M.W.; Pai, N.; Daggupati, P. Hydrologic and water quality models: Performance measures and evaluation criteria. *Transactions of the ASABE* **2015**, *58*, 1763–1785. DOI: 10.13031/trans.58.10715.
18. Harris, D.J.; Taylor, S.D.; White, E.P. Forecasting biodiversity in breeding birds using best practices. *PeerJ* **2018**, *6*, e4278. DOI: 10.7717/peerj.4278.
19. Christin, S.; Hervet, É.; Lecomte, N. Applications for deep learning in ecology. *Methods in Ecology and Evolution* **2019**, *10*, 1632–1644. DOI: 10.1111/2041-210X.13256.
20. Thakur, M.P.; van der Putten, W.H.; Wilschut, R.A.; Veen, G.F.; Kardol, P.; van Ruijven, J.; Allan, E.; Roscher, C.; van Kleunen, M.; Bezemer, T.M. Plant-soil feedbacks and temporal dynamics of plant diversity-productivity relationships. *Trends in Ecology & Evolution* **2021**, *36*, 651–661. DOI: 10.1016/j.tree.2021.03.011.
21. Valavi, R.; Guillera-Arroita, G.; Lahoz-Monfort, J.J.; Elith, J. Predictive performance of presence-only species distribution models: A benchmark study with reproducible code. *Ecological Monographs* **2022**, *92*, e01486. DOI: 10.1002/ecm.1486.
22. Ward, E.J.; Holmes, E.E.; Thorson, J.T.; Collen, B. Complexity is costly: A meta-analysis of parametric and non-parametric methods for short-term population forecasting. *Oikos* **2014**, *123*, 652–661. DOI: 10.1111/j.1600-0706.2014.00916.x.
23. Bellard, C.; Bertelsmeier, C.; Leadley, P.; Thuiller, W.; Courchamp, F. Impacts of climate change on the future of biodiversity. *Ecology Letters* **2012**, *15*, 365–377. DOI: 10.1111/j.1461-0248.2011.01736.x.
24. Cutler, D.R.; Edwards Jr, T.C.; Beard, K.H.; Cutler, A.; Hess, K.T.; Gibson, J.; Lawler, J.J. Random forests for classification in ecology. *Ecology* **2007**, *88*, 2783–2792. DOI: 10.1890/07-0539.1.

25. Mi, C.; Huettmann, F.; Guo, Y.; Han, X.; Wen, L. Why choose Random Forest to predict rare species distribution with few samples in large undersampled areas? Three Asian crane species models provide supporting evidence. *PeerJ* **2017**, *5*, e2849. DOI: 10.7717/peerj.2849.
26. Pullin, A.S.; Knight, T.M. Doing more good than harm - Building an evidence-base for conservation and environmental management. *Biological Conservation* **2009**, *142*, 931–934. DOI: 10.1016/j.biocon.2009.01.010.
27. Cook, C.N.; Possingham, H.P.; Fuller, R.A. Contribution of systematic reviews to management decisions. *Conservation Biology* **2013**, *27*, 902–915. DOI: 10.1111/cobi.12114.
28. Cook, C.N.; Mascia, M.B.; Schwartz, M.W.; Possingham, H.P.; Fuller, R.A. Achieving conservation science that bridges the knowledge-action boundary. *Conservation Biology* **2013**, *27*, 669–678. DOI: 10.1111/cobi.12050.
29. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, Ł.; Polosukhin, I. Attention is all you need. In *Advances in Neural Information Processing Systems 30*; Curran Associates, Inc.: Red Hook, NY, USA, 2017; pp. 5998–6008.
30. Brown, T.B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. Language models are few-shot learners. In *Advances in Neural Information Processing Systems 33*; Curran Associates, Inc.: Red Hook, NY, USA, 2020; pp. 1877–1901.
31. Devlin, J.; Chang, M.W.; Lee, K.; Toutanova, K. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1*; Association for Computational Linguistics: Minneapolis, MN, USA, 2019; pp. 4171–4186. DOI: 10.18653/v1/N19-1423.
32. Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; Liu, P.J. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research* **2020**, *21*, 1–67.
33. Toderas, M. Artificial Intelligence for Sustainability: A Systematic Review and Critical Analysis of AI Applications, Challenges, and Future Directions. *Sustainability* **2025**, *17*, 8049. DOI: 10.3390/su17178049.
34. Liu, P.; Yuan, W.; Fu, J.; Jiang, Z.; Hayashi, H.; Neubig, G. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys* **2023**, *55*, 1–35. DOI: 10.1145/3560815.
35. Kasneci, E.; Seßler, K.; Küchemann, S.; Bannert, M.; Dementieva, D.; Fischer, F.; Gasser, U.; Groh, G.; Günnemann, S.; Hüllermeier, E.; et al. ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences* **2023**, *103*, 102274. DOI: 10.1016/j.lindif.2023.102274.
36. Tondeur, J.; Petko, D.; Christensen, R.; Drossel, K.; Starkey, L.; Knezek, G.; Schmidt-Crawford, D.A. Quality criteria for conceptual technology integration models in education: Bridging research and practice. *Educational Technology Research and Development* **2023**, *71*, 1915–1937. DOI: 10.1007/s11423-023-10243-w.
37. McIntosh, B.S.; Ascough, J.C.; Twery, M.; Chew, J.; Elmahdi, A.; Haase, D.; Harou, J.J.; Hepting, D.; Cuddy, S.; Jakeman, A.J.; et al. Environmental decision support systems (EDSS) development - Challenges and best practices. *Environmental Modeling & Software* **2011**, *26*, 1389–1402. DOI: 10.1016/j.envsoft.2011.09.009.
38. Matthies, M.; Giupponi, C.; Ostendorf, B. Environmental decision support systems: Current issues, methods and tools. *Environmental Modeling & Software* **2007**, *22*, 123–127. DOI: 10.1016/j.envsoft.2005.09.005.
39. Rodríguez-Caro, R.C.; Graciá, E.; Anadón, J.D.; Giménez, A. Maintained effects of fire on individual growth and survival rates in a spur-thighed tortoise population. *European Journal of Wildlife Research* **2019**, *65*, 1–9. DOI: 10.1007/s10344-019-1329-0.
40. Camacho, C.; Negro, J.J.; Elmberg, J.; Fox, A.D.; Nagy, S.; Pain, D.J.; Green, A.J. Groundwater extraction poses extreme threat to Doñana World Heritage Site. *Nature Ecology & Evolution* **2022**, *6*, 654–655. DOI: 10.1038/s41559-022-01763-6.
41. Liu, Y.; Tang, Q.; Liu, X.; Wang, G.; Zhang, X.; Leng, G. Recent decrease in summer precipitation over the Iberian Peninsula closely linked to the weakening of local moisture recycling. *Hydrology and Earth System Sciences* **2022**, *26*, 1925–1936. DOI: 10.5194/hess-26-1925-2022.
42. Newman, S. *Building Microservices: Designing Fine-Grained Systems*; O'Reilly Media: Sebastopol, CA, USA, 2015; ISBN 978-1491950357.
43. Agencia Estatal de Meteorología (AEMET). Climate Data Portal. Available online: https://www.aemet.es/en/datos_abiertos (accessed on 15 December 2025).
44. ICTS-RBD. *Hydromet: Hydrometeorological Data*; Singular Scientific and Technical Infrastructure – Doñana Biological Reserve: Seville, Spain, 2026. Available online: <https://datos-automaticos.icts-donana.es/en/> (last accessed on 14 May 2026).

45. ICTS-DOÑANA. *Online Database: Aerial Waterbird Census in the Guadalquivir River Marshes*; Doñana Biological Station (EBD-CSIC): Seville, Spain, 2026. Available online: <https://censos-aereos.icts-donana.es/> (last accessed on 14 May 2026).
46. Santoro, S.; Green, A.J.; Figuerola, J. Immigration enhances fast growth of a newly established source population. *Ecology* **2013**, *94*, 1058–1067. DOI: 10.1890/12-0696.1.
47. Instituto Nacional de Estadística (INE). *Base de Datos de Indicadores Urbanos*; INE: Madrid, Spain, 2026. Available online: <https://www.ine.es/dyngs/DAB/index.htm?cid=1100> (last accessed on 15 May 2026).
48. PostgreSQL Global Development Group. PostgreSQL 15 Documentation. Available online: <https://www.postgresql.org/docs/15/> (last accessed on 15 December 2025).
49. Hyndman, R.J.; Athanasopoulos, G. *Forecasting: Principles and Practice*, 3rd ed.; OTexts: Melbourne, Australia, 2018. Available online: <https://otexts.com/fpp3/> (accessed on 15 December 2025).
50. James, G.; Witten, D.; Hastie, T.; Tibshirani, R. *An Introduction to Statistical Learning with Applications in R*; Springer: New York, NY, USA, 2013; ISBN 978-1461471370. DOI: 10.1007/978-1-4614-7138-7.
51. Breiman, L.; Friedman, J.H.; Olshen, R.A.; Stone, C.J. *Classification and Regression Trees*; Wadsworth International Group: Belmont, CA, USA, 1984.
52. Breiman, L. Random forests. *Machine Learning* **2001**, *45*, 5–32.
53. Chen, T.; Guestrin, C. XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*; San Francisco, CA, USA, 13–17 August 2016; pp. 785–794.
54. Box, G.E.P.; Jenkins, G.M. *Time Series Analysis: Forecasting and Control*; Holden-Day: San Francisco, CA, USA, 1970; ISBN 978-0816211043.
55. Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; et al. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* **2011**, *12*, 2825–2830.
56. Seabold, S.; Perktold, J. Statsmodels: Econometric and statistical modeling with Python. In *Proceedings of the 9th Python in Science Conference (SciPy 2010)*; Austin, TX, USA, 28 June–3 July 2010; pp. 92–96. DOI: 10.25080/Majora-92bf1922-011.
57. Bergmeir, C.; Benítez, J.M. On the use of cross-validation for time series predictor evaluation. *Information Sciences* **2012**, *191*, 192–213. DOI: 10.1016/j.ins.2011.12.028.
58. Tashman, L.J. Out-of-sample tests of forecasting accuracy: An analysis and review. *International Journal of Forecasting* **2000**, *16*, 437–450. DOI: 10.1016/S0169-2070(00)00065-0.
59. Gemini Team; Anil, R.; Borgeaud, S.; Alayrac, J.-B.; Yu, J.; Soricut, R.; Schalkwyk, J.; Dai, A.M.; Hauth, A.; Millican, K.; et al. Gemini: A Family of Highly Capable Multimodal Models. *arXiv* **2024**, arXiv:2312.11805. DOI: 10.48550/arXiv.2312.11805.
60. Chainlit Development Team. Chainlit: Build Production-Ready Conversational AI. Available online: <https://docs.chainlit.io/> (last accessed on 20 December 2025).
61. Schick, T.; Dwivedi-Yu, J.; Dessì, R.; Raileanu, R.; Lomeli, M.; Zettlemoyer, L.; Cancedda, N.; Scialom, T. Toolformer: Language models can teach themselves to use tools. *arXiv preprint arXiv:2302.04761* **2023**.
62. Amershi, S.; Weld, D.; Vorvoreanu, M.; Fournay, A.; Nushi, B.; Collisson, P.; Suh, J.; Iqbal, S.; Bennett, P.N.; Inkpen, K.; et al. Guidelines for human-AI interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*; Glasgow, UK, 4–9 May 2019; pp. 1–13. DOI: 10.1145/3290605.3300233.
63. Microsoft Corporation. Power BI Documentation. Available online: <https://docs.microsoft.com/en-us/power-bi/> (last accessed on 20 December 2025).
64. Lundberg, S.M.; Lee, S.I. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems 30*; Curran Associates, Inc.: Red Hook, NY, USA, 2017; pp. 4765–4774.
65. Google Cloud. Vertex AI Documentation. Available online: <https://docs.cloud.google.com/vertex-ai/docs> (last accessed on 15 May 2026).
66. Mankins, J.C. Technology readiness levels: A white paper. Advanced Concepts Office, Office of Space Access and Technology, NASA, April 1995.
67. Laniak, G.F.; Olchin, G.; Goodall, J.; Voinov, A.; Hill, M.; Glynn, P.; Whelan, G.; Geller, G.; Quinn, N.; Blind, M.; et al. Integrated environmental modeling: A vision and roadmap for the future. *Environmental Modeling & Software* **2013**, *39*, 3–23. DOI: 10.1016/j.envsoft.2012.09.006.
68. Arrieta, A.B.; Díaz-Rodríguez, N.; Del Ser, J.; Bennetot, A.; Tabik, S.; Barbado, A.; García, S.; Gil-López, S.; Molina, D.; Benjamins, R.; et al. Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion* **2020**, *58*, 82–115. DOI: 10.1016/j.inffus.2019.12.012.

69. Saltelli, A.; Bammer, G.; Bruno, I.; Charters, E.; Di Fiore, M.; Didier, E.; Nelson Espeland, W.; Kay, J.; Lo Piano, S.; Mayo, D.; et al. Five ways to ensure that models serve society: A manifesto. *Nature* **2020**, *582*, 482–484. DOI: 10.1038/d41586-020-01812-9.
70. Karpatne, A.; Atluri, G.; Faghmous, J.H.; Steinbach, M.; Banerjee, A.; Ganguly, A.; Shekhar, S.; Samatova, N.; Kumar, V. Theory-guided data science: A new paradigm for scientific discovery from data. *IEEE Transactions on Knowledge and Data Engineering* **2017**, *29*, 2318–2331. DOI: 10.1109/TKDE.2017.2720168.
71. Lütkepohl, H. *New Introduction to Multiple Time Series Analysis*; Springer: Berlin, Germany, 2005; ISBN 978-3540401728. DOI: 10.1007/978-3-540-27752-1.
72. Chen, R.T.Q.; Rubanova, Y.; Bettencourt, J.; Duvenaud, D. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems 31*; Curran Associates, Inc.: Red Hook, NY, USA, 2018; pp. 6571–6583.
73. Gelman, A.; Carlin, J.B.; Stern, H.S.; Dunson, D.B.; Vehtari, A.; Rubin, D.B. *Bayesian Data Analysis*, 3rd ed.; CRC Press: Boca Raton, FL, USA, 2013; ISBN 978-1439840955.
74. Stem, C.; Margoluis, R.; Salafsky, N.; Brown, M. Monitoring and evaluation in conservation: A review of trends and approaches. *Conservation Biology* **2005**, *19*, 295–309. DOI: 10.1111/j.1523-1739.2005.00594.x.
75. Turner, W.; Rondinini, C.; Pettorelli, N.; Mora, B.; Leidner, A.K.; Szantoi, Z.; Buchanan, G.; Dech, S.; Dwyer, J.; Herold, M.; et al. Free and open-access satellite data are key to biodiversity conservation. *Biological Conservation* **2015**, *182*, 173–176. DOI: 10.1016/j.biocon.2014.11.048.
76. Pettorelli, N.; Schulte to Bühne, H.; Tulloch, A.; Dubois, G.; Macinnis-Ng, C.; Queirós, A.M.; Keith, D.A.; Wegmann, M.; Schrod, F.; Stellmes, M.; et al. Satellite remote sensing of ecosystem functions: Opportunities, challenges and way forward. *Remote Sensing in Ecology and Conservation* **2018**, *4*, 71–93. DOI: 10.1002/rse2.59.
77. Zhu, X.X.; Tuia, D.; Mou, L.; Xia, G.S.; Zhang, L.; Xu, F.; Fraundorfer, F. Deep learning in remote sensing: A comprehensive review and list of resources. *IEEE Geoscience and Remote Sensing Magazine* **2017**, *5*, 8–36. DOI: 10.1109/MGRS.2017.2762307.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.