

AI-Driven Drug Development

**Striking a balance between innovation and risk
through robust AI governance**

Ajay Tomar, SAS Institute, Delhi, India

Ashwini Reddy, SAS Institute, Hyderabad, India



Where AI is Showing up in Life Sciences

Contextualizing AI in life sciences for competitive advantage

Experimentation

Utilization

Commercialization

Internal

External

Examples

Digital twin
for cellular
response

Molecule
docking
prediction

Automated
radiology
interpreter

Real-time
clinical
decision
support

Automated
regulatory
compliance

Real-time
health
monitoring
and analysis

Genomic
variant
classifier

Protein
folding
simulation

AI pathology
assistant

Readmission
risk predictor


Supply chain
optimization

AI-based drug
discovery

What are Regulators Expecting?



Standardized Risk Management



Remediation Plans in Place



Self Assessments
Embedded in ERMF




Integration of Non-Model Risks



3 Lines of Defense Framework




High-Quality Reporting
Informs Decisions



Model Validation over
Performance Management



Standard Evaluation of
Ethical Decisions




The Challenge

AI models can be complicated and difficult to explain to stakeholders

The Solution: Model Cards in SAS Model Manager

Trustworthy AI Landscape



Trustworthy AI Landscape

ACTIVATION

Oversight

AI Governance, Strategy,
and Enforcement

Operations

SOPs w/ Supporting
Infrastructure

Compliance

Performance and
Risk Management

Culture

Ethically Systemic
Norms & Practices

Data Management

Data
Quality

Variable
Metadata

Data
Preparation

Data Asset
Catalog

Explanation

Natural
Language
Explanation

Explainable ML

Counterfactual
Explanation

Surrogate Model
Interpretation

Causal
Inference

TECHNOLOGY

Detection

Bias
Detection

Fairness
Assessment

Privacy and Security

Privacy Preservation

Model
Security

Autonomy
Preservation

Consent & Control

MODEL CARDS

Includes features that allows users to document the intended use case, performance, limitations, model information/settings

Model
Cards

Decisioning

Lifecycle
Management

Metric
Monitoring

Model
Robustness


Model
Oversight

Model Ops

What are Model Cards?

“Nutrition labels for models” provide critical insight

Model cards generate a summary of a model’s training data, intended use and performance - demonstrating whether models are accurate, transparent and fair.



Fosters Trustworthy AI ● Provides Consumable Insights ● Enables Ethical Decision-Making




Overview Tab

Key Insights and Use Cases

Summary of model performance, reliability compared to selected metrics and key influential variables.

Helps decide whether to put the model in production.



Model Usage Tab

Key Insights and Use Cases

Document intended use of the model, appropriate and out of scope use cases.

Prevent misuse and promote responsible deployment by making users aware of inappropriate applications and the model's limitations.

The screenshot displays the SAS Model Manager interface for a model named "Gradient Boosting (1) (AutoML) (1.1)" under the project "Salary Classification (Version 1)". The "Model Usage" tab is selected, showing a "Not published" status and a "Yes" response for "available:". The "Model Usage" section includes fields for "Intended Use", "Expected Benefit", "Out-of-Scope Use Cases", and "Limitations". The "Data Summary" section shows 11 columns, 32561 rows, and a size of 7 MB, with a status of "Approved" and a completeness of 98%. The "Data Tags" section includes "Public" and "UCIrvine". The "Training Data" section describes the dataset as "ADULT.sashdat" and lists entities: age, individual, race, and 2 others, noting the storage format is CAS and that the data contains outliers and sensitive values.

SAS® Model Manager - Manage Models

< Models

Gradient Boosting (1) (AutoML) (1.1) Project: Salary Classification (Version 1)

Model Card Files Variables Properties Versions

Overview

Model Usage

Data Summary

Model Summary

Model Audit

Tags:

DS2 multi-type

Modeler:

sasadm

Responsible party:

Sophia Rowland

Not published

Yes

available:

Yes

SAS

Model Usage

Intended Use

This model will be used to target our marketing efforts towards individuals who are likely to meet the thresholds for our services.

Expected Benefit

To increase the ROI of our marketing spend.

Out-of-Scope Use Cases

Determining who meets the requirements for our services.

Limitations

This model was trained on public census data, curated by UC Irvine in their machine learning repository.

Data Summary

Date last modified: Apr 1, 2025, 12:46:56 PM

Date last analyzed: Apr 1, 2025, 12:48:05 PM

Columns: 11

Rows: 32561

Size: 7 MB

Status: Approved

Completeness: 98%

Data Tags

Public UCIrvine

Training Data

ADULT.sashdat

This dataset describes information about the following entities: **age**, **individual**, **race**, and 2 others. This dataset contains information from multiple **country** locations. The storage format is **CAS**. The data contains outliers and has values that could be considered **sensitive**.

Data Summary

Key Insights and Use Cases

Analyze training data for insight into the AI's decision-making process.

Includes data source, data privacy considerations, data tags to indicate data types, data description, data relevance in training the model.

Includes analyses of data completeness, approval status and potential outliers.

The screenshot displays the SAS Model Manager interface for a project named "Salary Classification (Version 1)". The left sidebar shows a navigation menu with options: Overview, Model Usage, Data Summary (selected), Model Summary, and Model Audit. Below the menu, there are tags for "DS2 multi-type", "Modeler: sasadm", and "Responsible party: Sophia Rowland".

The main content area is titled "Data Summary" and includes the following information:

- Model Card:** Gradient Boosting (1) (AutoML) (1.1) | Project: Salary Classification (Version 1)
- Overview:** This model was trained on public census data, curated by UC Irvine in their machine learning repository.
- Data Summary Metrics:**
 - Columns: 11
 - Rows: 32561
 - Size: 7 MB
 - Status: Approved
 - Completeness: 98%
- Data Tags:** Public, UCIrvine
- Description:** This is a public data set from UC Irvine's Machine Learning Repository. It contains census data, including various characteristics and a salary indicator.
- Information Privacy:**
 - Private Variables:** Age, Occupation, WorkClass
 - Age contains 143 outliers that might impact its analysis.
 - Age is strongly correlated with the variable "Relationship". (Correlation: 0.66)
 - Age is strongly correlated with the variable "MaritalStatus". (Correlation: 0.71)
 - Sensitive Variables:** Race, Sex
 - Sex does not contain any outliers.
 - Sex is strongly correlated with the variable "Relationship". (Correlation: 0.74)
 - Candidate Variables:** NativeCountry
 - NativeCountry does not contain any outliers.
 - NativeCountry is not correlated with other variables within the data.
- Training Data:** ADULT.sashdat
 - This dataset describes information about the following entities: age, individual, race, and 2 others.
 - This dataset contains information from multiple country locations.
 - The storage format is CAS.
 - The data contains outliers and has values that could be considered sensitive.
- Model Summary:** Tool: Model Studio, Tool version: V2024.09


Model Summary

Key Insights and Use Cases

Understand granular model performance results, fairness assessment results, and generalizability metrics.

Review model interpretability, which indicates the variables most useful for model predictions.

Include governance details like the modeler, responsible party and model version.




Model Audit

Key Insights and Use Cases

Monitor an AI model over time to ensure accuracy, fairness, and avoid model drift.

Understand when the AI model was last evaluated, and which data was used for performance evaluation.



Thank you

Ajay Tomar (ajay.tomar@sas.com)

Ashwini Reddy (Ashwini.reddy@sas.com)

