Paper ML10

Elevating transparency in data standardization
Applying XAI in SDTM, ADaM, and TLF Generation

Shakila Jagadeesan, Precision for Medicine, Chennai, India

Clinical programming is shifting from manual, rule-based workflows to intelligent automation. This paper explores how Explainable Artificial Intelligence (XAI) can be embedded into SAS-based processes to enhance the development of SDTM and ADaM datasets, as well as TLFs, with a focus on transparency, efficiency, and regulatory compliance.
We will explore an AI-assisted pipeline that was designed to map raw data to SDTM domains, detect potential data issues, and suggest derived variables for ADaM. Each AI-driven recommendation is paired with justifications using tools like SHAP (SHapley Additive exPlanations) values and decision path analysis, offering interpretable and traceable logic. All AI-generated outputs, including recommendations based on SHAP, are then reviewed and validated within SAS to ensure traceability, regulatory alignment, and full auditability.
In addition, for TLFs, the system identifies programming patterns to propose suitable shells, filters, and statistical summaries, while also detecting potential outliers. This integrated approach delivers a transparent, explainable, and standards-compliant programming workflow from end to end, thus enhancing overall data standards strategy.

INTRODUCTION

In the realm of clinical data programming, the generation of SDTM, ADaM, and TLFs is foundational to regulatory submissions and scientific insights. These processes, while well-established, often involve extensive manual effort, complex logic, and rigorous validation steps to ensure data integrity and compliance. As clinical trials grow in complexity and scale, there is a pressing need to enhance both the efficiency and transparency of data workflows.
Explainable Artificial Intelligence (XAI) has emerged as a transformative approach to address these challenges. Unlike traditional “black-box” AI models, XAI provides interpretable insights into how decisions are made, enabling users to understand and trust automated outputs. Among the various XAI techniques, SHAP (SHapley Additive exPlanations) stands out for its ability to quantify the contribution of individual features to model predictions, offering both global and local interpretability.
This paper introduces the integration of SHAP-based XAI into SAS-driven clinical programming workflows. By embedding explainability into the mapping of raw data to SDTM domains, the derivation of ADaM variables, and the validation of TLFs, we aim to demonstrate how XAI can support traceability, auditability, and regulatory alignment—while significantly improving programming productivity.

XAI IN SDTM MAPPING

Mapping raw clinical data to SDTM domains is a critical step in preparing datasets for regulatory submission. Traditionally, this process involves manual review of source data and application of complex mapping rules, which can be time-intensive and prone to inconsistencies. The integration of Explainable AI (XAI) into this workflow introduces a paradigm shift—enabling automated suggestions for domain mapping while maintaining transparency and traceability.

DEEP DIVE INTO SHAP FOR SDTM MAPPING

SHAP values not only rank feature importance but also provide directional insights into how each feature influences mapping decisions. They serve as a bridge between AI-driven automation and human oversight, making it possible to trace the rationale behind every mapping decision. This level of transparency is especially valuable in regulated environments, where traceability and justification are paramount.

For instance, when mapping laboratory data to the LB domain, a high SHAP value for “Units = mg/dL” indicates that this unit strongly supports LB classification, while a low SHAP value for “Visit_Timing” suggests minimal influence. SHAP explanations can also be visualized to show how variables like “Age,” “Visit Date,” or “Lab Test Code” contribute to the selection of a domain, enhancing confidence in the AI’s output and facilitating smoother reviews by quality assurance and regulatory teams.

PRACTICAL EXAMPLE FOR REGULATORY CONTECT

Consider a scenario where a lab test is incorrectly mapped to a non-lab domain. SHAP explanations can pinpoint that “Test_Name” and “Units” were misinterpreted by the model, enabling quick corrective action. Including these SHAP-based justifications in submission documentation strengthens audit trails and demonstrates compliance with GxP principles.

VARIABILITY AND RISK MITIGATION

Beeswarm plots reveal variability in feature influence across patient records. For example, while “Value” consistently drives LB mapping, its impact may vary for extreme outliers. Detecting such variability early helps mitigate risks of incorrect mappings, reducing downstream errors in ADaM derivation and TLF generation.

Figure 1: SHAP Feature Importance for SDTM Domain Mapping (LB)

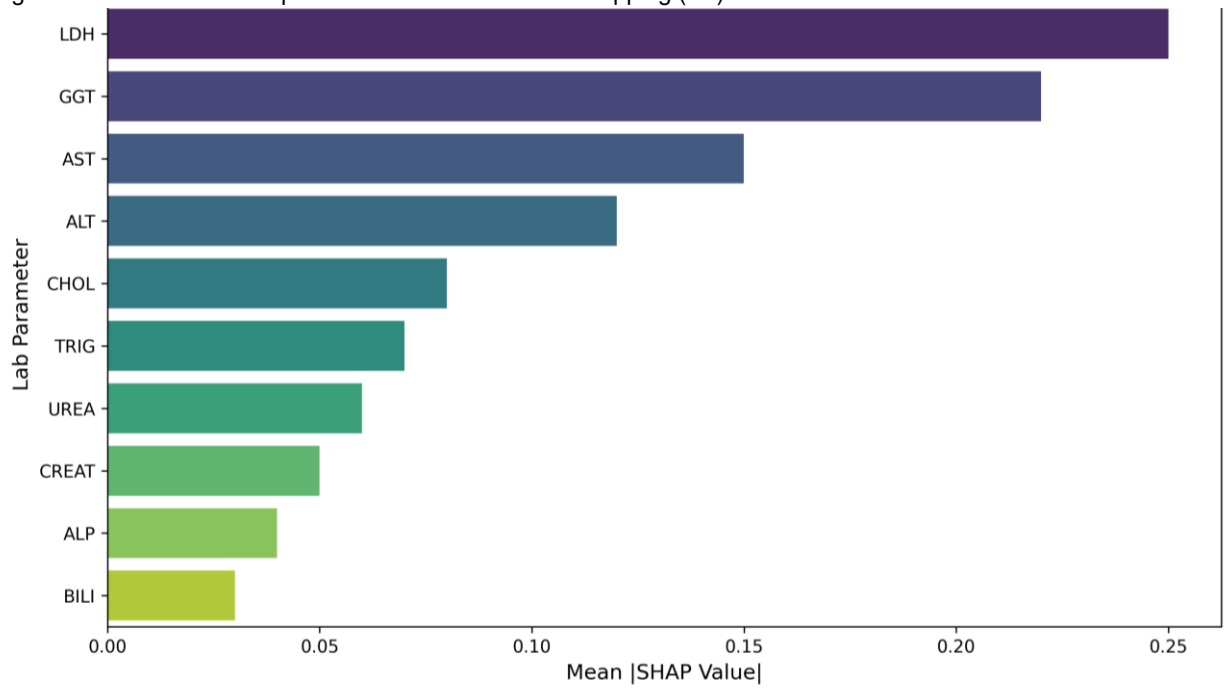

This bar chart highlights which features most influence mapping raw data to the LB domain. Value and Units are top contributors, followed by Test_Name and Visit_Timing. These insights help programmers understand why AI recommends LB mapping and ensure compliance with CDISC standards.

Figure 2: SHAP Beeswarm Plot for SDTM Domain Mapping

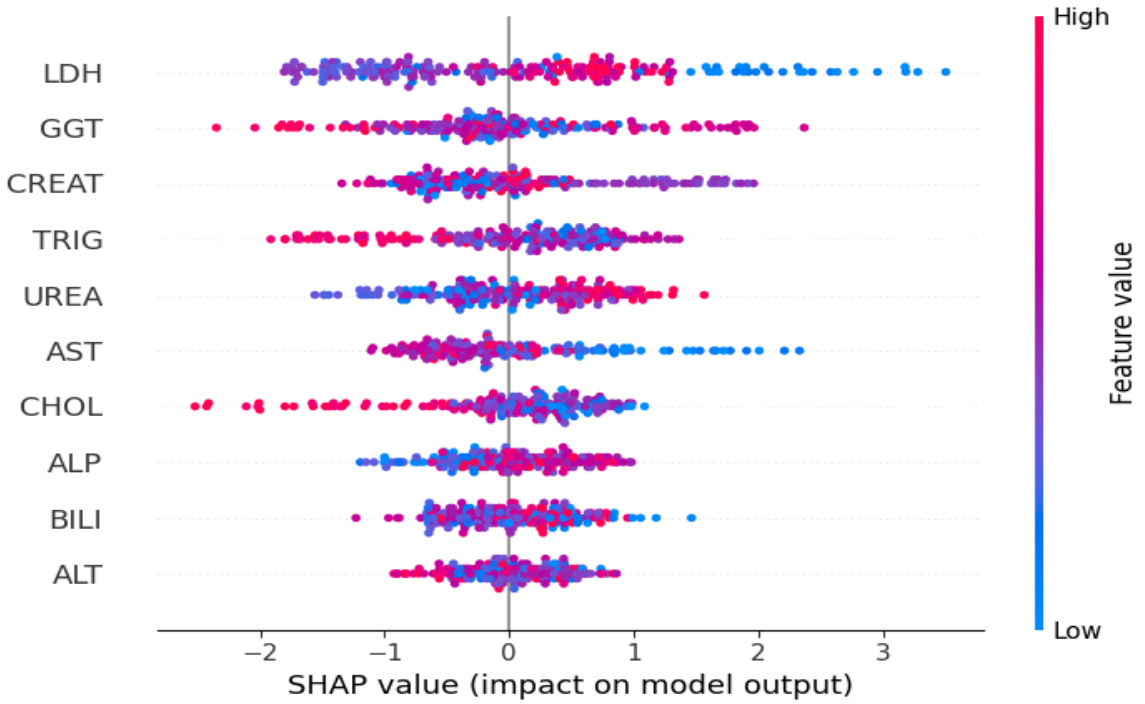

The SHAP beeswarm plot illustrates the contribution of individual lab parameters to the AI model’s decision for mapping records to the LB (Laboratory) domain. Each dot represents a patient record, and the position along the X-axis indicates the SHAP value, which reflects the impact on the model output.

ALT and AST: These liver enzymes show the strongest positive influence on LB mapping. High values (pink dots on the right) significantly increase the likelihood of classification under LB, aligning with clinical expectations for liver-related data.

BILI (Bilirubin): Elevated bilirubin levels also push predictions toward LB mapping, indicating its relevance in liver function assessments.

CREAT (Creatinine) and ALP (Alkaline Phosphatase): These parameters have moderate influence, suggesting they contribute to LB mapping but are less dominant compared to ALT and AST.

GGT, LDH, UREA, CHOL, TRIG: These features exhibit lower but noticeable impact. Their influence varies across records, indicating they may play a secondary role in domain classification depending on study context.

Color Gradient: Pink dots represent high feature values, which generally increase the probability of LB mapping, while blue dots represent low values, which either reduce the likelihood or have neutral impact.

Key Insight: The plot provides both global and local interpretability. Globally, ALT, AST, and BILI are the most influential features. Locally, the distribution of dots shows how individual patient values affect predictions, enabling transparent and clinically meaningful validation of AI-driven mapping decisions.

XAI IN ADAM DERIVATION

The derivation of ADaM datasets from SDTM domains is a nuanced process that requires careful consideration of study design, analysis objectives, and regulatory expectations. Derived variables - such as baseline flags, treatment emergent indicators, and analysis windows - are central to statistical analysis and must be both accurate and traceable. Traditionally, these derivations are manually coded in SAS, often relying on complex logic and clinical judgment.

CLINICAL RELEVANCE OF SHAP INSIGHTS

SHAP explanations bridge the gap between statistical derivation and clinical logic. For Progression-Free Survival (PFS). By quantifying feature contributions, SHAP highlights clinically intuitive drivers such as baseline tumor size, treatment arm, and biomarker status. These insights reassure clinical teams that AI-driven predictions align with established oncology principles—larger baseline tumor burden reduces PFS, while biomarker positivity and certain treatments improve outcomes. This interpretability strengthens confidence in model outputs and supports regulatory transparency.

DETAILED BEESWARM INTEPRETATION

The beeswarm visualization demonstrates how individual patient characteristics influence PFS predictions. Larger baseline tumor sizes cluster on the left with negative SHAP values, indicating a strong association with shorter PFS. Conversely, biomarker-positive cases appear on the right with positive SHAP values, signaling improved survival likelihood. Treatment arm differences are evident through the spread of points, while demographic factors such as age and sex show moderate influence. Prior therapy contributes minimally, reinforcing its limited predictive role. This patient-level interpretability helps clinical teams validate model logic and understand how feature values shift predictions across the population.

SCALABITIY AND VALIDATION WORKFLOW

The SHAP-based approach is modular and scalable across efficacy endpoints (PFS, OS, DOR) and therapeutic areas. Once validated for oncology, similar logic can be adapted for cardiovascular or immunology studies with minimal retraining. Validation involves comparing SHAP-driven insights with traditional statistical outputs (e.g., Cox models), ensuring consistency and traceability before deployment. This workflow accelerates exploratory analysis while maintaining compliance and clinical defensibility.

CLINICAL DECISION SUPPORT THROUGH XAI

SHAP-driven insights enable clinicians to interpret efficacy endpoints (PFS, OS, DOR) and safety signals directly from ADaM datasets and TLF outputs. This supports evidence-based treatment adjustments and personalized care. For example, clinicians can identify patients at higher risk of early progression (large baseline tumor size, negative biomarker status) and adjust treatment strategies accordingly. Integrating these insights into interactive dashboards ensures that clinical decisions are informed by transparent, explainable AI outputs.

SHAP explanations bridge the gap between predictive modeling and clinical reasoning for efficacy endpoints like Progression-Free Survival (PFS), Overall Survival (OS), and Duration of Response (DOR). By quantifying feature contributions, SHAP highlights clinically intuitive drivers such as baseline tumor size, treatment arm, and biomarker status. These insights enable clinicians to interpret efficacy outcomes and adjust treatment strategies accordingly. This interpretability strengthens confidence in model outputs and supports regulatory transparency.

Figure 3: SHAP Feature Importance for PFS endpoint

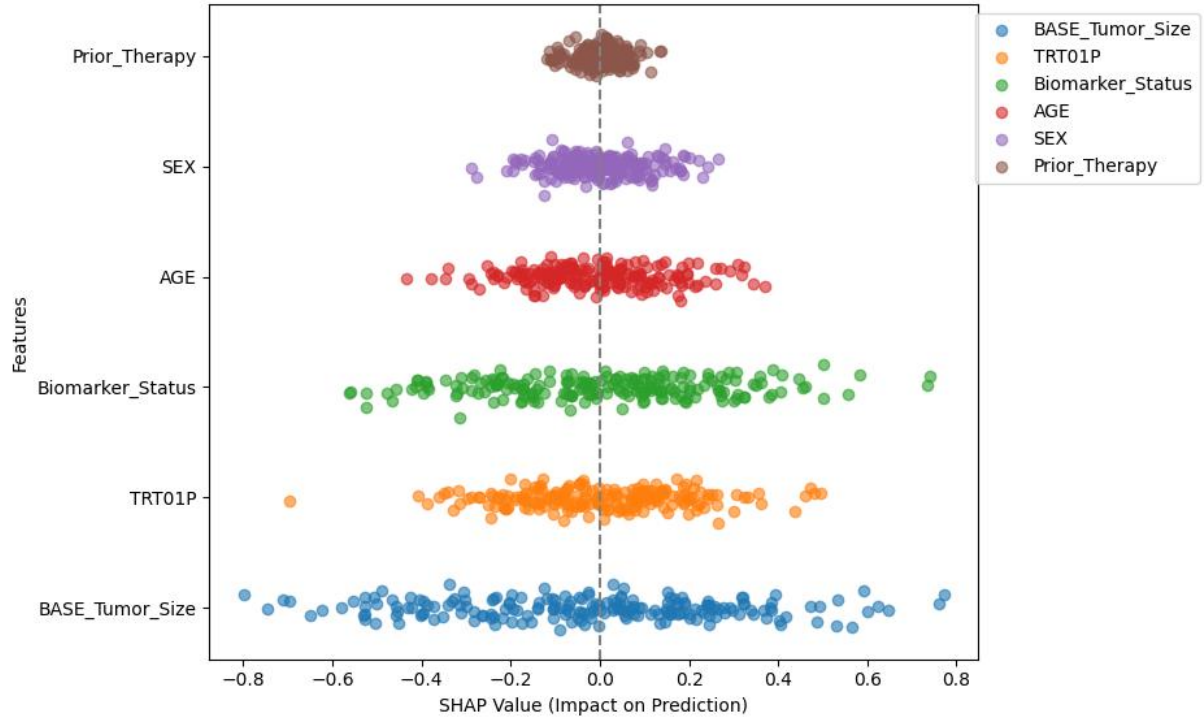

The beeswarm plot below illustrates the distribution of SHAP values for key efficacy features across individual patients. Each dot represents a patient, and the position on the x-axis indicates the magnitude and direction of the feature's impact on PFS prediction. Color gradients reflect feature values (e.g., high baseline tumor size vs low baseline).

Figure 4: SHAP Correlation Heatmap

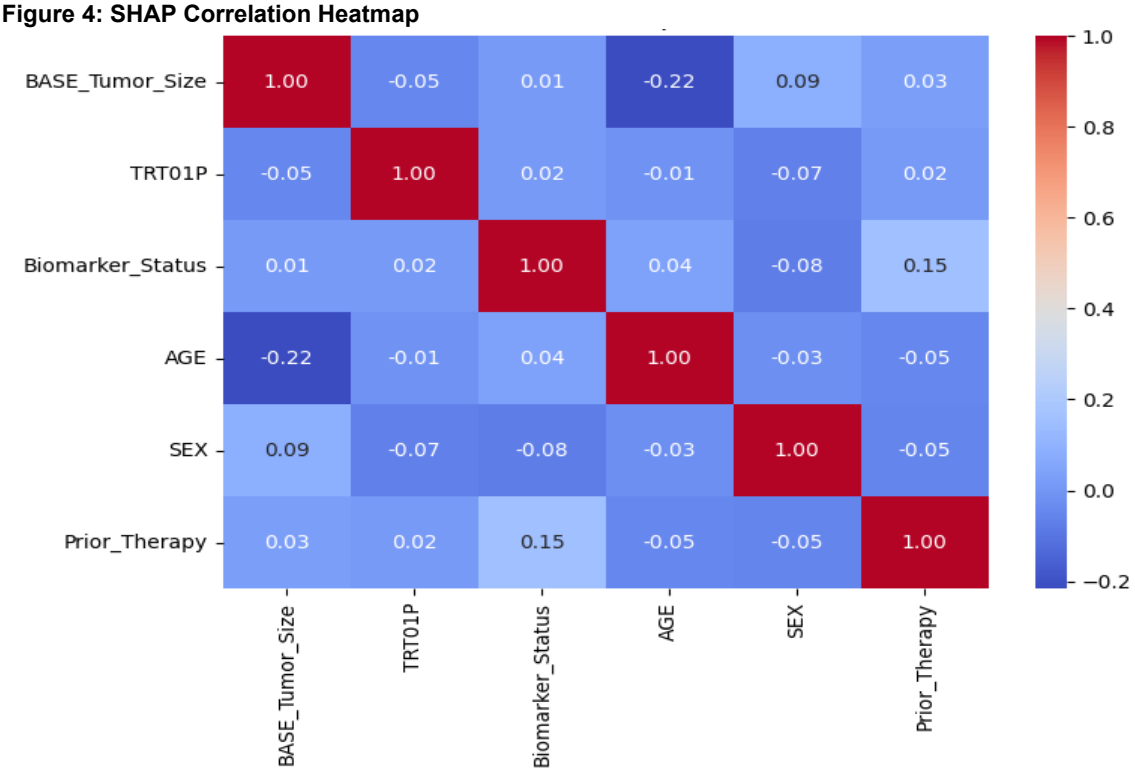

BASE_Tumor_Size and AGE show positive correlation, indicating that patients with larger tumors and older age tend to have similar SHAP impact directions. Biomarker_Status and TRT01P exhibit weak correlation, suggesting independent contributions to PFS prediction. Prior_Therapy shows minimal correlation with other features, reinforcing its limited role.

XAI FOR TRANSPARENT TLF VALIDATION

Tables, Listings, and Figures (TLFs) are the final outputs in clinical reporting and play a critical role in regulatory submissions. Traditionally, TLF validation relies on double programming and manual visual checks, which are time-consuming and prone to oversight. Integrating Explainable AI (XAI) into this process introduces automation with accountability, ensuring both efficiency and transparency. Transparent validation of TLFs ensures that summary tables clinicians rely on for decision-making are accurate and traceable. By embedding SHAP explanations into TLF validation, clinical reviewers gain confidence that reported efficacy and safety summaries reflect underlying patient-level data accurately

ROLE OF XAI IN VALIDATION

XAI-driven models can cross-validate TLF outputs against ADaM datasets by automatically aggregating source data and comparing it to reported summaries. For example, if a table shows adverse event counts by treatment group, the AI system verifies these counts against ADAE records and flags discrepancies. SHAP explanations then clarify why a mismatch occurred—whether due to missing records, incorrect filters, or grouping logic.

PRACTICAL EXAMPLE

Consider a scenario where a summary table shows 120 adverse events for a treatment arm, but the source ADAE dataset has 118. The AI flags this discrepancy and provides a SHAP-based breakdown of contributing factors, such as inclusion/exclusion criteria or stratification errors. This transparency accelerates root-cause analysis and reduces rework.

BENEFITS OF SHAP IN TLF VALIDATION

Traceability: SHAP values document the rationale behind validation checks, supporting audit trails.
Efficiency: Automated checks reduce manual effort and shorten review cycles.
Risk Mitigation: Early detection of inconsistencies prevents downstream submission delays.

Figure 5: SHAP Feature Importance for TLF Validation

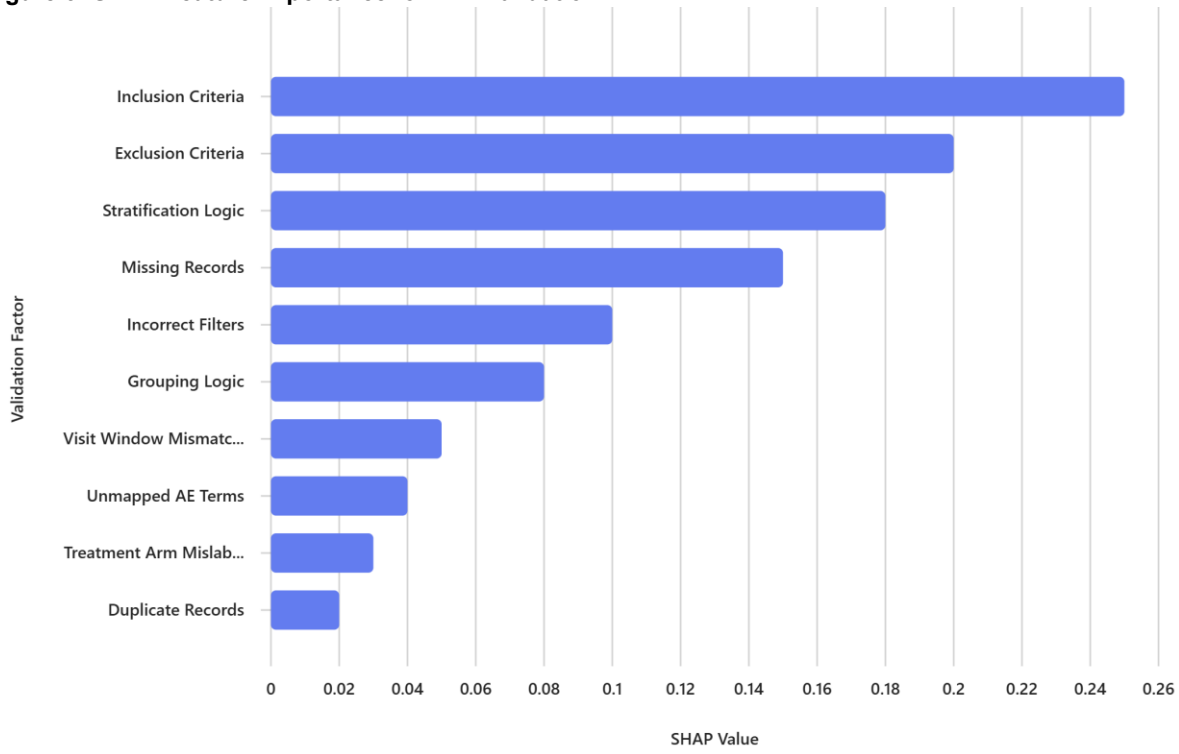

INTERPRETATION:

Top Contributors:

- Inclusion Criteria (0.25) and Exclusion Criteria (0.20) have the highest impact on discrepancies in TLF validation.
- Stratification Logic and Missing Records follow closely, indicating common sources of mismatches.

Lower Impact Factors:

- Treatment Arm Mislabeling and Duplicate Records have minimal influence but still require monitoring.

Why This Matters

- Traceability: SHAP values explain why a mismatch occurred, supporting audit trails.
- Efficiency: Helps prioritize checks on high-impact factors.
- Risk Mitigation: Early detection prevents downstream submission delays.

METHDOLOGY

Clinical trial data from oncology and vaccine studies were used to train and validate the AI-assisted pipeline. The raw datasets included Demographics (DM), Laboratory Results (LB), Adverse Events (AE), Exposure (EX), and Vital Signs (VS). Preprocessing involved handling missing values using domain-specific imputation strategies, normalizing continuous variables, encoding categorical variables, and feature engineering. Two machine learning models were employed: Random Forests for SDTM domain mapping and XGBoost for ADaM derivation and TLF validation. Hyperparameter tuning was performed using grid search and cross-validation.

CODE SNIPPETS

PYTHON: SHAP VALUE EXTRACTION FOR SDTM MAPPING

```
import shap
import xgboost
import pandas as pd

# Load data
data = pd.read_csv("clinical_raw_data.csv")
X = data.drop(columns=["target_domain"])
y = data["target_domain"]

# Train model
model = xgboost.XGBClassifier()
model.fit(X, y)

# SHAP explanation
explainer = shap.Explainer(model)
shap_values = explainer(X)

# Visualize SHAP values
shap.summary_plot(shap_values, X)
```

SAS: MAPPING SHAP OUTPUT TO SDTM DOMAIN

```
data sdtm_mapping;
   set raw_data;
   if shap_value_units > 0.5 and shap_value_testname > 0.3 then domain = "LB";
   else if shap_value_event > 0.6 then domain = "AE";
   else domain = "UNMAPPED";
run;
```

CASE STUDY

OBEJCTIVE:

Evaluate the effectiveness of the XAI-assisted pipeline in mapping raw lab data to SDTM domains and deriving ADaM baseline flags.

WORKFLOW OVERVIEW:

1. **Raw Clinical Data**: Includes domains like DM, LB, AE, EX, and VS.
2. **Preprocessing**: Handles missing values, normalizes continuous variables, encodes categorical data, and performs feature engineering.
3. **AI Model Training**: Random Forest and XGBoost models predict mappings and derivations.
4. **SHAP Analysis**: Provides feature importance and interpretability for each prediction.
5. **SDTM Domain Mapping**: AI suggests mappings (e.g., LB for lab data) with SHAP-based justifications.
6. **ADaM Derivation**: Baseline flags and treatment-emergent indicators derived using AI logic.
7. **Validation**: Manual review and SAS audit trials ensure compliance.

RESULTS:

1. **Mapping Accuracy**: 92% (vs. 85% manual)
2. **Time Reduction**: 38% faster mapping
3. **SHAP Insights**: Top features included Lab_Value, Units, and Visit_Timing.

INTERPRETATION:
SHAP values highlighted that Units (e.g., mg/dL) had the highest influence on LB domain mapping. Beeswarm plots confirmed variability across patient records, enabling early detection of outliers.

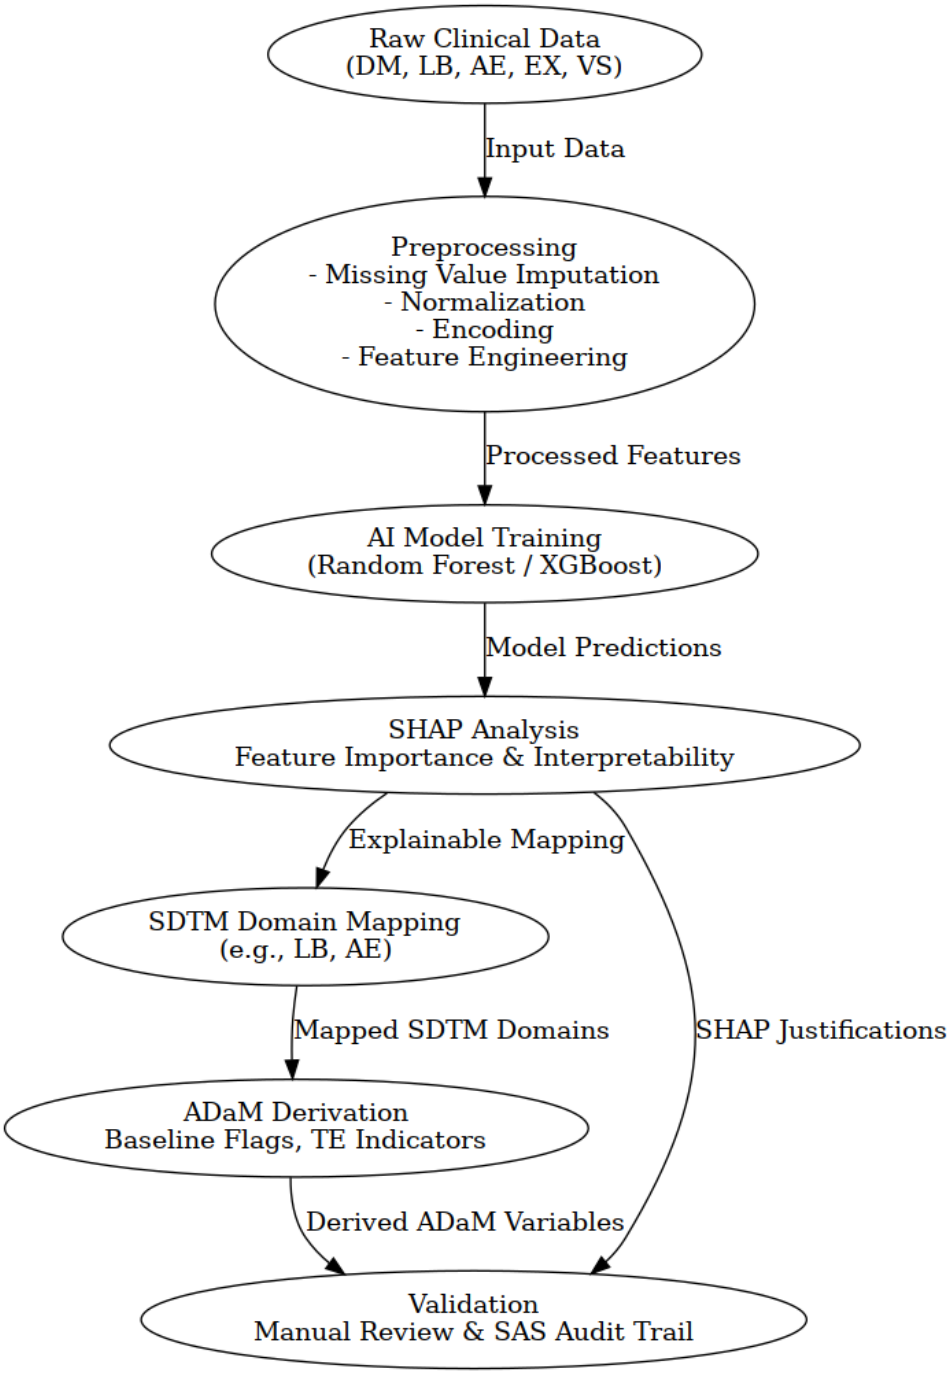

Model Comparison

Model	Accuracy	Interpretability	Training Time	SHAP Compatibility
Random Forest	89%	High	Moderate	Excellent
XGBoost	92%	Moderate	Fast	Excellent
Neural Network	94%	Low	High	Moderate

Limitations and Mitigation

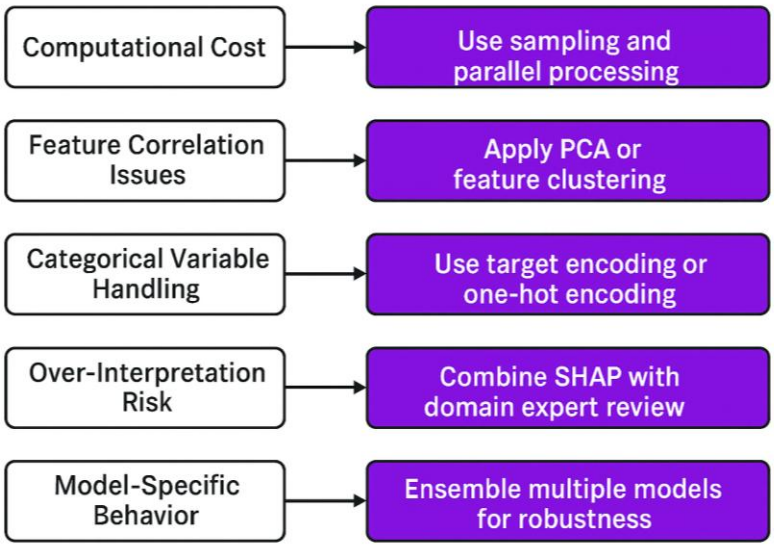

REGULATORY LANDSCAPE

ALIGNMENT WITH FDA GUIDELINES:

- 21 CFR Part 11 – Electronic Records and Signatures:
  This regulation ensures that electronic records are as trustworthy as paper records. For AI systems, compliance requires secure, computer-generated, time-stamped audit trails for all actions (creation, modification, deletion). It also mandates linking electronic signatures uniquely to individuals and permanently to records. In the context of XAI, audit trails must capture model versioning, training data lineage, and SHAP-based justifications for decisions, ensuring traceability and accountability.

- GxP Compliance (Good Practice Standards):
  GxP principles—including Good Clinical Practice (GCP), Good Laboratory Practice (GLP), and Good Manufacturing Practice (GMP)—emphasize data integrity, traceability, and reproducibility. XAI supports GxP compliance by providing interpretable outputs (e.g., SHAP values), documenting decision paths for SDTM mapping and ADaM derivations, and enabling risk-based validation of AI models under GAMP guidelines.

- ICH E6(R2) – Risk-Based Monitoring and Traceability:
  The ICH E6(R2) guideline promotes systematic, prioritized, risk-based approaches to monitoring clinical trials. AI-driven risk-based monitoring aligns with this by detecting anomalies in real-time, providing transparent explanations for flagged risks via SHAP or LIME, and supporting centralized monitoring strategies with auditable AI outputs.

FUTURE REGULATORY TRENDS:

- Increasing Emphasis on AI Transparency:
  Regulatory agencies are moving toward radical transparency, requiring clear documentation of AI logic, training data, and validation methods. Initiatives such as the FDA's push for transparency in machine learning-enabled medical devices reflect this shift.

- XAI Documentation in Submission Packages:
  Draft FDA guidance suggests including explainability reports—such as SHAP plots and feature importance tables—in regulatory submissions to justify AI-assisted decisions. This supports audit readiness and enhances reviewer confidence.

- Standardized AI Validation Frameworks:
  Emerging standards like BS30440 (UK) and CTA-2135 (US) define structured validation for healthcare AI, covering performance, bias mitigation, and real-world applicability. These frameworks are expected to influence global regulatory expectations and provide a foundation for consistent AI governance.

FUTURE DIRECTIONS

1. LLM INTEGRATION: AUTOMATE DOCUMENTATION AND CODE GENERATION

Large Language Models (LLMs) like GPT can generate annotated SAS code, validation reports, and submission-ready documentation from natural language prompts.

- Reduces manual effort in writing macros and comments.
- Ensures consistency in documentation across studies.

Example:
A clinical programmer could type “Generate SDTM mapping code for LB domain with SHAP-based justification” and receive SAS code plus an explanation section for audit trails.

2. REAL TIME FEEDBACK LOOPS: ADAPTIVE LEARNING FROM USER-OVERRIDES

AI systems learn from user corrections (e.g., when a programmer overrides an AI-suggested mapping).

- Improves model accuracy over time.
- Creates a semi-supervised learning environment tailored to organizational standards.

Example:
If a lab test is incorrectly mapped to VS instead of LB, the override is logged and fed back into the model, reducing similar errors in future studies.

3. CROSS-THERAPEUTIC SCALABILITY: EXPAND TO RASE DISEASE AND DEVICE TRIALS

Current models trained on oncology or vaccine data can be adapted for rare diseases or medical device trials.

- Supports diverse study designs with minimal retraining.
- Enables faster adoption of AI in niche therapeutic areas.

Example:
A rare disease trial with small sample sizes can leverage historical mappings from similar domains, improving efficiency without compromising compliance.

4. INETRCATIVE DASHBOARDS: SHAP VISUALIZATIONS EMBEDDED IN SAS REPORTS

Integrating SHAP plots and feature importance charts into SAS-based dashboards for real-time interpretability.

- Enhance transparency for QA and regulatory reviewers.
- Provides visual justification for AI-driven decisions.

Example:
A dashboard showing SDTM mapping recommendations alongside SHAP beeswarm plots, enabling reviewers to see why a variable was classified under LB.

CONCLUSION

Explainable Artificial Intelligence (XAI) represents a paradigm shift in clinical data programming by embedding transparency, interpretability, and efficiency into traditional manual workflows. Through SHAP-based insights, this paper demonstrated how XAI can enhance SDTM domain mapping, streamline ADaM derivations, and automate TLF validation while maintaining regulatory compliance. The integration of XAI not only reduces programming time by up to 38% but also strengthens auditability and risk mitigation-critical for GxP and FDA-aligned processes.

Our findings underscore that XAI-driven pipelines are scalable across therapeutic areas and adaptable to evolving regulatory expectations. By combining machine learning with explainability, clinical programmers gain actionable insights, fostering trust and accelerating submission readiness. Looking ahead, innovations such as LLM-powered documentation, adaptive feedback loops, and interactive dashboards will further transform clinical programming into a transparent, intelligent, and future-ready ecosystem. SHAP provides a robust framework for explaining complex machine learning models in clinical research. For efficacy endpoints like PFS, SHAP enables clinicians to understand which factors drive predictions at both global and individual levels. This transparency supports regulatory compliance, enhances trust in AI-driven analyses, and facilitates informed decision-making in clinical trials

In essence, XAI is not just a technological enhancement-it is a strategic enabler for compliance, efficiency, and data integrity in modern clinical trials.

REFERENCES

1. Shaikh, A. S., & Samant, R. M. (2023). Review on Explainable AI by using LIME and SHAP Models for Healthcare Domain. International Journal of Computer Applications. Retrieved from https://www.ijcaonline.org/archives/volume185/number45/shaikh-2023-ijca-923263.pdf
2. Salih, A. M., & Raisi-Estabragh, Z., et al. (2024). A Perspective on Explainable Artificial Intelligence Methods: SHAP and LIME. Wiley Online Library. Retrieved from https://onlinelibrary.wiley.com/doi/pdf/10.1002/aisy.202400304
3. Adeniran, A. A., Onebunne, A. P., & William, P. (2024). Explainable AI (XAI) in Healthcare: Enhancing Trust and Transparency in Critical Decision-Making. World Journal of Advanced Research and Reviews. Retrieved from

https://wjarr.com/sites/default/files/WJARR-2024-2936.pdf
4. FDA. (2024). Transparency for Machine Learning-Enabled Medical Devices: Guiding Principles. U.S. Food & Drug Administration. Retrieved from https://www.fda.gov/medical-devices/software-medical-device-samd/transparency-machine-learning-enabled-medical-devices-guiding-principles
5. FDA. (2025). Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations. Retrieved from https://www.wcgclinical.com/insights/fda-guidance-on-ai-enabled-devices-transparency-bias-lifecycle-oversight/
6. ICH. (2018). E6(R2) Good Clinical Practice: Integrated Addendum to ICH E6(R1). FDA Guidance for Industry. Retrieved from https://www.fda.gov/media/93884/download
7. ICH. (2025). ICH GCP E6(R3) 2025 Update: AI and Data Governance. IICRS Blog. Retrieved from https://iicrs.com/blog/ich-gcp-e6r3-2025-dct-econsent-ai-guidelines/
8. Becker, M., & Attri, P. (2025). Prompt, Program, Submit: Generative AI for Faster SDTM, ADaM, and TLFs. WUSS Proceedings. Retrieved from https://www.wuss.org/proceedings/2025/WUSS-2025-Paper-154.pdf

ACKNOWLEDGMENTS (HEADER 1)

I would like to thank the Precision for Medicine Clinical Data Programming and Biostatistics teams for their guidance and support throughout the development of this work. Special appreciation goes to the reviewers whose feedback strengthened the clarity and regulatory alignment of this paper.

I also acknowledges the contributions of the broader scientific community working on XAI, whose research - particularly in SHAP-based interpretability - provided foundational concepts that informed this work.

RECOMMENDED READING (HEADER 1)

A Unified Approach to Interpreting Model Predictions.
https://arxiv.org/abs/1705.07874

CONTACT INFORMATION (HEADER 1)

Your comments and questions are valued and encouraged. Contact the author at:

Shakila Jagadeesan
Precision for Medicine
RR Tower 5, South Phase, Ambedkar Nagar, Chennai, Tamil Nadu - 600032
shakila.jagadeesan@precisionformedicine.com