GitHub Link: https://github.com/glendonC/mallards

🚨 The Challenge

AI-driven financial systems **fail to recognize cultural financial behaviors**, leading to:

- **Unnecessary rejections & financial exclusion** due to biased models.
- **A reactive approach to bias** instead of proactively preventing it.
- **Lack of explainability** —lenders & policymakers don't know *why* AI makes certain decisions.

💡 The Solution

We built an **AI-powered financial intelligence platform** that:

- **Proactively identifies lending biases before they happen.**
- **Explains financial decisions in real-time using an AI Assistant** powered by **LLMs + TTS.**
- **Simulates policy adjustments** to show the impact of financial rule changes.

Instead of just detecting bias **after** it happens, our system **predicts and prevents it.**

Current Market Gaps

1. **AI Monitoring Tools (Datadog, Splunk, New Relic)**
    a. **What they do:**
        - Monitor **system performance** & detect **generic anomalies.**
        - Identify technical failures, **not financial or cultural biases.**
    b. **What they lack:**
        - No **financial decision intelligence.**
        - No **explainability—just anomaly flags.**
        - No **cultural adaptation or proactive learning.**
2. **AI Fairness Tools (IBM AI Fairness 360, Aequitas)**
    a. **What they do:**
        - Analyze **demographic bias in AI models.**
        - Monitor protected categories (e.g., race, gender).
        - Provide **basic fairness metrics.**

b. What they lack:

- No **real-time decision adjustments.**
- No **predictive simulations to prevent bias before it happens.**
- No **AI-driven explanations—only statistical fairness reports.**

Mallard's Unique Focus

Instead of just monitoring AI bias, we're creating an "AI-Driven Cultural Intelligence Layer" that:

- **Predicts & explains financial decisions.**
- **Prevents AI-driven bias before it occurs.**
- **Provides interactive AI insights to financial institutions.**

Key Features

1. **AI-Powered Financial Simulation (Simulation Sandbox)**
   - **Allows users to modify financial parameters and run AI-driven impact assessments.**
   - **Instantly displays fairness, risk, and approval rate changes.**
   - **Uses before-and-after comparisons to visualize financial impact.**
2. **AI Insights Panel**
   - **Analyzes real-time decision patterns.**
   - **Provides LLM-powered explanations for transaction approvals, denials, and risk factors.**
   - **Users can ask AI questions for further clarification.**
3. **Predictive Modeling & Anomaly Detection**
   - **Uses forecasting models (Prophet, ARIMA) to detect cultural financial patterns.**
   - **Flags potential AI decision conflicts before they become an issue.**
4. **Bias Prevention & Fairness Visualization**
   - **Instead of just detecting bias after the fact, the system proactively suggests policy optimizations before they negatively impact users.**
   - **AI-powered fairness simulations help institutions adjust lending models.**

Real-World Example

Problem: A bank's AI model flags high-value jewelry purchases as suspicious due to fraud prevention algorithms.

Challenge:

- **During Diwali, an Indian cultural festival, luxury jewelry purchases spike.**
- **Instead of simply "allowing more transactions", our system would:**
    1. **Recognize Diwali as an upcoming high-value transaction period.**
    2. **Pre-adjust fraud models BEFORE misclassifications happen.**
    3. **AI Assistant explains these adjustments to financial decision-makers.**
    4. **Simulation Sandbox allows banks to test policy changes in real-time.**
    5. **Post-event AI insights provide fairness & financial inclusion analysis.**

Impact Areas

1. **Religious Financial Practices**
    - **Detect patterns in Islamic banking compliance (i.e no interest lending rules)**
    - **Identify seasonal donation trends in faith-based communities**
    - **Predict financial behaviors during religious events and adjust risk models**
2. **Community-Based Credit Systems**
    - **Model cooperative lending structures like ROSCAs (Rotating Savings and Credit Associations)**
    - **Analyze approval trends for microfinance programs**
    - **Provide AI-generated recommendations on optimizing local lending policies.**
3. **Cultural Views on Debt & Wealth**
    - **Identify cultural spending habits and adjust AI risk thresholds accordingly.**
    - **Simulate the impact of policy changes on culturally specific financial behaviors**
    - **Provide real-time AI-driven insights on regional financial behaviors.**

Core Innovation

We (or just me) are pioneering the first AI-powered Cultural Intelligence Layer for financial systems, which:

1. **Proactively Predicts & Explains Cultural Financial Trends**
    - **Anticipates financial behavior shifts before they happen.**
    - **Users can ask the AI Assistant why certain financial trends occur.**
2. **Automates Cultural Adaptation with Predictive AI**

- **Risk scoring and fraud detection models adjust dynamically during cultural financial periods**
- **Simulation Sandbox allows financial institutions to preview the impact of policy changes before deployment.**

3. **Measures Community Impact & Drives Fairness in Lending**
    - **Quantifies how AI lending decisions affect different cultural groups.**
    - **Provides adaptive policy recommendations to financial institutions.**
    - **AI-generated insights explain why approval rates shift during cultural events**

Measurable Impact

1. **Reduced False Positives & Smarter Fraud Detection**
    - **Fewer incorrect transaction declines during cultural financial spikes.**
    - **AI adapts fraud detection dynamically based on real-time spending trends.**
2. **Increased Approval Rates with Adaptive Risk Models**
    - **More accurate AI decision-making for legitimate cultural transactions**
    - **Simulation Sandbox allows lenders to test & optimize approval policies in advance**
3. **Better Financial Inclusion with AI-Driven Fairness Adjustments**
    - **Lending policies adapt based on data-driven insights, not demographic assumptions.**
    - **AI quantifies lending fairness & ensures underrepresented groups have access to credit**

The Future Vision

By implementing this solution, financial institutions can move beyond basic demographic bias checks and instead embrace true AI-driven cultural intelligence, creating an ecosystem where:

- **AI-driven decisions proactively respect cultural financial diversity.**
- **Businesses operate with greater confidence in diverse markets.**
- **Communities have equitable access to financial opportunities.**
- **AI doesn't just detect bias—it explains, adapts, and prevents it before it happens**

Component Details

1. Dashboard

a. Active Cultural Periods

i. Core Purpose
   1. This component tracks ongoing and upcoming cultural financial periods and provides a summarized view of their impact on financial behaviors.

ii. Key Functions
   1. **Monitors Active and AI-Predicted Cultural Periods**
      a. Displays currently active financial periods with their estimated impact.
      b. Predicts upcoming periods using AI and provides expected financial behavior changes.
   2. **Expandable Detailed View**
      a. Users can **expand/collapse** to see full details of active & upcoming periods.
      b. **Pagination system** to prevent data overload (10 events at a time).
   3. **Impact Analysis**
      a. Shows **expected transaction volume change (%)**.
      b. Categorizes the impact level (e.g., **Low, Medium, High**).
      c. AI-predicted events are clearly labeled.

iii. Use Cases
   1. Risk Management → Identify upcoming transaction surges.
   2. Fraud Prevention → Adjust fraud detection thresholds proactively.
   3. Business Strategy → Plan financial policies around cultural periods

b. Community Impact Map

i. Core Purpose
   1. Provides a global visualization of cultural financial impact by region, tracking trends in spending, approvals, and bias.

ii. Key Functions
   1. **Interactive Globe for Regional Insights**
      a. Displays financial behaviors **by country/region**.
      b. Allows filtering by:
         i. **Spending Hotspots** (high transaction volumes).
         ii. **Decision Patterns** (loan approval rates).
         iii. **Bias Factors** (fairness disparities in AI decisions).
   2. **Data-Driven Analysis**
      a. **Color-coded heatmaps** to visualize risk levels.
      b. **Threshold-based scaling** for spending/approval metrics.

c. Regional trend detection (rising/stable/declining financial behavior).

3. **Filter and View Mode Customization**
    a. Toggle between:
        i. **Heatmap Mode** → Shows intensity of financial activity.
        ii. **Bubble Mode** → Displays transaction hotspots.

iii. Use Cases
1. **Financial Institutions** → Compare financial inclusion across regions.
2. **Regulatory Bodies** → Detect bias in lending policies.
3. **Market Strategy** → Identify underserved regions for financial products.

c. Cultural Alignment Score

i. Core Purpose
1. Evaluates how aligned financial decisions are with cultural patterns, detecting deviations from expected approval rates.

ii. Key Functions
1. **Cultural vs. Normal Financial Patterns**
    a. Compares **approval rates in cultural periods** vs. normal periods.
    b. Tracks long-term shifts in **financial behavior trends**.
2. **AI-Powered Forecasting**
    a. Uses **historical transaction data** to predict **future approval trends**.
    b. **Confidence intervals** show forecast reliability.
3. **Anomaly Detection**
    a. Highlights significant **deviations in approval trends**.
    b. Detects patterns in **denials, approvals, and policy shifts**.

iii. Use Cases
1. **Loan Approvals** → Detect unfair cultural rejections.
2. **Fraud Detection** → Identify unexpected approval spikes.
3. **Financial Inclusion Analysis** → Evaluate policy fairness.

d. Cultural Decision Impact

i. Core Purpose
1. Measures how cultural events influence financial decision-making, including approval rates, loan volumes, and decision shifts.

ii. Key Functions
1. **Tracks Approval Rate Changes Over Time**
    a. Monitors **how approvals vary across cultural periods**.
    b. **Region-specific** financial impact insights.
2. **Detects Anomalous Decision Patterns**
    a. **Highlights significant shifts** in approval/rejection rates.

        b. Tracks **p-value and confidence intervals** for statistical accuracy.
    3. **Regional Fairness Metrics**
        a. Measures **disparities in financial decisions across cultural groups**.
        b. Provides **AI-driven recommendations** to adjust policies.
iii. Use Cases
    1. **Regulatory Compliance** → Ensuring fairness in financial AI models.
    2. **Bias Detection** → Identifying unexplained lending disparities.
    3. **Market Expansion** → Evaluating approval trends in underserved areas.

e. Cultural Event Analytics

i. Core Purpose
    1. Analyzes how specific cultural events impact financial transactions historically and in forecasts.
ii. Key Functions
    1. **Hierarchical Event Analysis**
        a. Groups events **by Year → Quarter → Individual Events**.
        b. Sorts events based on **financial impact**.
    2. **AI-Powered Forecasting**
        a. Predicts **future trends** in cultural spending/approval rates.
        b. Provides **upper/lower confidence bounds** for forecasts.
    3. **Historical Comparison & Trend Tracking**
        a. Displays **past vs. predicted financial outcomes**.
        b. Highlights **events with the largest financial deviations**.
iii. Use Cases
    1. **Risk Assessment** → Predict upcoming financial shifts.
    2. **Business Planning** → Allocate resources for high-transaction periods.
    3. **Fraud Prevention** → Detect transaction anomalies before they occur.

f. Cultural Pattern Alerts

i. Core Purpose
    1. Detects real-time alerts for financial anomalies that may indicate unexpected cultural transaction patterns.
ii. Key Functions
    1. **Monitors Financial Patterns Across Focus Areas**
        a. **Spending Anomalies** → Unusual spikes/drops in transactions.
        b. **Approval Deviations** → Changes in loan acceptance rates.
        c. **Bias Indicators** → AI fairness drift detection.
    2. **Adaptive AI Learning**

        a. Alerts evolve over time as AI **learns from past patterns**.
        b. Recommendations for **real-time policy adjustments**.
    3. **Severity Categorization**
        a. Alerts grouped into **Low, Medium, High**.
        b. **Suggested Actions** for each alert based on risk level.
iii. Use Cases
    1. **Fraud Prevention** → Detect early signs of fraudulent activity.
    2. **Financial Decision-Making** → Prevent misclassifications in cultural spending.
    3. **Regulatory Oversight** → Ensure fair AI-driven lending decisions.

2. Anomaly

a. Cultural Pattern Violations

i. Core Purpose
    1. This component tracks transaction behaviors that deviate from expected cultural financial patterns. It ensures financial systems adapt to cultural events and prevent misclassifications of transactions (e.g., holiday spending spikes being flagged as fraud).

ii. Key Functions
    1. **Pattern Detection**
        a. Identifies **unexpected spending surges, approval dips, and unusual deviations**.
        b. Associates violations with **cultural events & regional financial patterns**.
    2. **Severity Classification**
        a. Assigns a severity score (`High`, `Medium`, `Low`) based on **impact & confidence level**.
    3. **Time-Based Trend Analysis**
        a. Uses **historical transaction trends** to predict **expected behavior**.
        b. Highlights deviations **in approval rates & transaction amounts**.

iii. Use Cases
    1. **Fraud Prevention – Identify if normal holiday spending is wrongly flagged as fraud.**
    2. **Policy Adjustments – Preemptively modify risk scores during financial events.**
    3. **Regulatory Reporting – Show fairness in financial decision-making.**
    4. **Real-Time Alerts – Detect sudden approval declines and prevent unnecessary rejections.**

b. AI Decision Conflicts

i. Core Purpose

1. This component analyzes conflicts in AI-driven financial decisions, ensuring AI models do not introduce unfair biases against cultural financial behaviors.

ii. Key Functions

1. **Approval vs. Rejection Discrepancy Analysis**
    a. Tracks **approval/rejection ratios across different cultural groups**.
    b. Detects **overly conservative or aggressive lending patterns**.
2. **Regional & Spending Focus**
    a. Users can toggle between:
        i. **Spending Conflicts** (Average transaction amounts by group)
        ii. **Approval Conflicts** (Approval rate trends)
        iii. **Regional Conflicts** (Approval rates by location)
3. **Cultural Context Insights**
    a. Displays the **cultural significance of lending behaviors**.
    b. AI **justifies lending decisions** based on historical & real-time trends.

iii. Use Cases

1. **Bias Detection** – Detect if **AI models unfairly reject transactions from specific groups.**
2. **Fair Lending Compliance** – **Ensure AI follows responsible lending regulations.**
3. **Risk Management** – **Prevent aggressive rejection patterns that harm financial inclusion.**
4. **Cultural Sensitivity Analysis** – **Show how AI lending aligns with regional & cultural financial behaviors.**

c. Unusual Deviations

i. Core Purpose

1. This component detects outlier financial behaviors, highlighting transactions that significantly deviate from historical norms.

ii. Key Functions

1. **Deviation Pattern Analysis**
    a. Categorizes anomalies by:
        i. **Timing** – Unexpected transaction spikes.
        ii. **Amount** – Unusually high or low spending.
        iii. **Frequency** – Rapid transaction bursts.
        iv. **Distribution** – Geographic spending irregularities.
2. **Severity & Trend Classification**
    a. Assigns severity (`High`, `Medium`, `Low`).
    b. Analyzes **historical patterns to detect increasing/decreasing trends**.
3. **AI-Driven Impact Analysis**

        a. Uses **model confidence scores** to **filter false positives**.
        b. Highlights **high-risk deviations that require attention**.
    iii. Use Cases
        1. **Fraud Detection** – Identify **potential fraudulent transaction bursts**.
        2. **Anomaly-Based Alerts** – **Flag sudden transaction spikes for investigation.**
        3. **Adaptive Risk Management** – **Help banks dynamically adjust lending policies.**
        4. **Regulatory Compliance** – **Ensure AI-driven financial models meet fairness standards.**

3. Simulation Sandbox

a. Parameter Controls

i. Core Purpose
    1. Parameter Controls allow users to manually adjust AI decision thresholds before applying AI-driven optimizations. These controls help simulate lending, fraud detection, and fairness policies.

ii. Key Functions
    1. Adjustable sliders for approval sensitivity, fraud detection, and cultural impact.
    2. **Allows real-time AI recalibration** based on manual changes.
    3. **Live Feedback:** AI **immediately recalculates lending risks and fairness changes** based on adjustments.

iii. Use Cases
    1. **Financial Analysts:** Fine-tune risk thresholds before real-world deployment.
    2. **Regulatory Teams:** Simulate compliance changes before enforcement.
    3. **AI Researchers:** Test how parameter adjustments affect model fairness.

b. Optimization Panel

i. Core Purpose
    1. The Optimization Panel is responsible for displaying and applying AI-driven parameter optimizations. It provides users with before vs. after analysis and highlights expected improvements based on AI recommendations.

ii. Key Functions
    1. Processes AI-driven parameter optimizations
    2. Compares original and optimized parameters with before vs. after percentage changes.
    3. Explains optimization reasoning with a collapsible AI insights section.

iii. Use Cases

1. **Banking Analysts:** Adjust AI-driven lending policies dynamically.
2. **Regulatory Reviewers:** Understand AI bias adjustments before approval.
3. **Developers:** Test different AI fairness models before production deployment.

c. Parameter Optimizer
   i. Core Purpose
      1. The Parameter Optimizer executes AI-driven lending policy adjustments based on user-defined scenarios. This is where the AI applies financial optimizations to improve fairness, approval rates, and fraud detection.
   ii. Key Functions
      1. Communicates with the backend to generate optimized lending parameters.
      2. Applies AI-generated suggestions to adjust risk thresholds dynamically.
      3. Allows users to review AI-driven changes before applying them.
      4. Displays percentage-based improvements for approval rates, fraud detection, and fairness metrics.
   iii. Use Cases
      1. **Lending Institutions:** Ensure **AI lending policies are fair before deployment.**
      2. **Regulators:** Test how **policy changes impact different demographics.**
      3. **Financial Engineers:** Optimize AI fairness models before production launch.

d. AI Insights Panel (applied to every other page technically)
   i. Core Purpose
      1. Provides AI-driven explanations for changes in financial decision-making.
   ii. Key Functions
      1. AI summarizes why metrics changed (i.e “Approval rates dropped due to stricter fraud detection.”).
      2. Users can ask the AI questions like *"Why did fraud detection trigger more rejections?"*
      3. AI generates text-based recommendations → TTS

e. Results Viewer
   i. Core Purpose
      1. Provides a side-by-side comparison of before vs. after simulation results.
   ii. Key Functions
      1. Displays **Approval Rates, Bias Metrics, Fraud Impact, and Risk Scores** before/after adjustments.

2. Shows **trend indicators** (📈 Increase, 📉 Decrease).
3. AI generates **risk level recommendations** (i.e. *"Consider adjusting fraud thresholds to maintain fairness."*).

iii. Use Cases

1. **Lenders & Banks**: Test impact of new lending rules before rollout.
2. **Regulators**: Assess if financial policies improve fairness & reduce bias.

f. Visualization Panel

i. Core Purpose

1. Generates interactive visualizations to help users understand lending impact at a glance.

ii. Key Functions

1. **Approval Rate Bar Charts** → Shows approval rate changes over time.
2. **Fraud Detection "Heatmap:** → Highlights regions with **increased fraud risk.**

iii. Use Cases

1. **Risk Teams**: See if stricter fraud detection **impacts fair lending.**
2. **Product Teams**: Adjust loan approval strategies for cultural markets.

g. Export Results

i. Self explanatory