Data Sheet

fiddler | Domino

Integrated AI Governance, Observability, and MLOps

Mission-driven AI at scale — secure, governed, and trusted.

Fiddler: AI Observability for LLM and ML

Fiddler offers a unified AI Observability and Security platform for LLM and MLOps, empowering government agencies to operationalize high-stakes AI with trust.

- Monitor performance, safety, accuracy, and privacy.
- Diagnose model issues and root causes to improve ML models and LLM applications.
- Protect from prompt injections and jailbreaking attacks.

Domino: AI Platform and MLOps

Domino's enterprise AI and MLOps platform helps government agencies integrate AI into their missions rapidly, safely, and cost-effectively.

- Build, deploy, and manage AI on a unified platform.
- Manage access to on-demand infrastrastructure, data, tools, models, and projects across environments.
- Improve collaboration and governance.

Fiddler and Domino Overview

Scalable and Collaborative AI Development

- Develop, train, and deploy models at scale across environments with Domino.
- Validate for model transparency, interpretability, and trustworthiness in production with Fiddler.

Continuous Model Governance and Monitoring

- Domino streamlines model lifecycle management with reproducibility and versioning at every stage.
- Fiddler provides model monitoring to keep AI systems performant, accurate, and trustworthy.

Operationalize Trusted AI at Scale

- Embed fairness, explainability, and compliance directly into AI workflows.
- Confidently scale AI, meet regulatory requirements and ethical standards without slowing innovation.

Challenges Around Responsible Deployment of AI

Performance: How satisfactory is the model's response?

Safety: Are the user prompt and model responses safe?

Correctness: How accurate is the response?

Bias: Is the model's response biased?

Security: Is the model leaking private data?

Quality: How is the data quality?

Cost: Where is the best ROI?

Transparency: Why did the model say that?

A/B Test: Is the model changing across versions?

Robustness: How sensitive is the model's response?

What Agencies Need for Trustworthy AI at Scale

Industrialization	• Deploy and monitor tabular, computer vision, and deep learning models across on premises, cloud(s) and edge infrastructures. • Automate the AI lifecycle while achieving resilience and cost avoidance.
Security and Governance	• Secure all work with contractors at the data, infrastructure and mode levels. • Govern the AI model lifecycle to provide traceability, observability, interpretability, and replaceability
Collaboration	• Bring together all relevant parties across the AI lifecycle: Data Science, Model Developers, IT, DevOps. • Provide templates, best practices, and guardrails to ensure predictability.
Innovation	• Flexibility to integrated the latest innovations in AI space — both commercial and open source. • Transform internal and cross-agencies innovation in mission impact.

Use Case: Automatic Target Recognition (ATR) at the US Department of Defense

For years, the U.S. Navy has relied on computer vision machine learning (ML) models to enhance underwater threat detection by unmanned underwater vehicles (UUVs). However, monitoring and improving ML performance post-deployment proved challenging. Models struggled to adapt to evolving underwater conditions and enemy tactics, increasing the risk of mission-critical systems delivering inaccurate or unreliable intelligence.

Project AMMO, leveraging Domino and Fiddler, transformed this process. By accelerating ATR model training, and enhancing decision-making speed and accuracy, the initiative enabled faster and more reliable autonomous warfighting, reducing model retraining time from 12 months to just 2 weeks.

fiddler | AI Observability Platform

Fiddler enables **model developers** in Project AMMO to measure subtle differences in tranches of sonar imagery training data that their ATR models care about. Additionally, it allows them identify and isolate common failure cases with semantic clustering and model explainability.

Fiddler will provide **tactical operators**, "the human in the loop", with visual explanations of model decisions that fortifies trust in their tools and allows AI to enhance their expert decision making, rather than overriding it.

Validate, Evaluate, and Explain

- Observe Model Behavior
- Explain Outcomes
- Visualize Data Patterns and Outliers
- Insights to Improve Model Performance

Monitor and Analyze

- Measure Operational Metrics (ie. Drift)
- Custom Dashboards and Reports
- Receive Alerts
- Root Cause Analysis
- Explain Outcomes

Domino | Enterprise AI Platform

Domino's enterprise AI platform serves as the factory for integrating technologies, enabling **model developers** to leverage the best open source and commercial tools of their choice in a highly governed way, on one platform, to generate and analyze ATR field-generated insights.

Domino's built-in model portability allows these teams to easily convert and export models into external formats to far-edge platforms, and it allows users to securely access and share structured and unstructured data from anywhere.

AI Factory

- Orchestrate Open Source and Commercial Tools in a Unified Platform
- Self-serve, Secure Compute and Data
- Flexible Model Deployment and Portability

Secure IP and Collaboration

- Rapidly Onboard Users and Contractors
- Share Data and Models Centrally
- Automatic Model Lineage (versioning of code, data, models, and environments)

High-Level Computer Vision Model Workflow

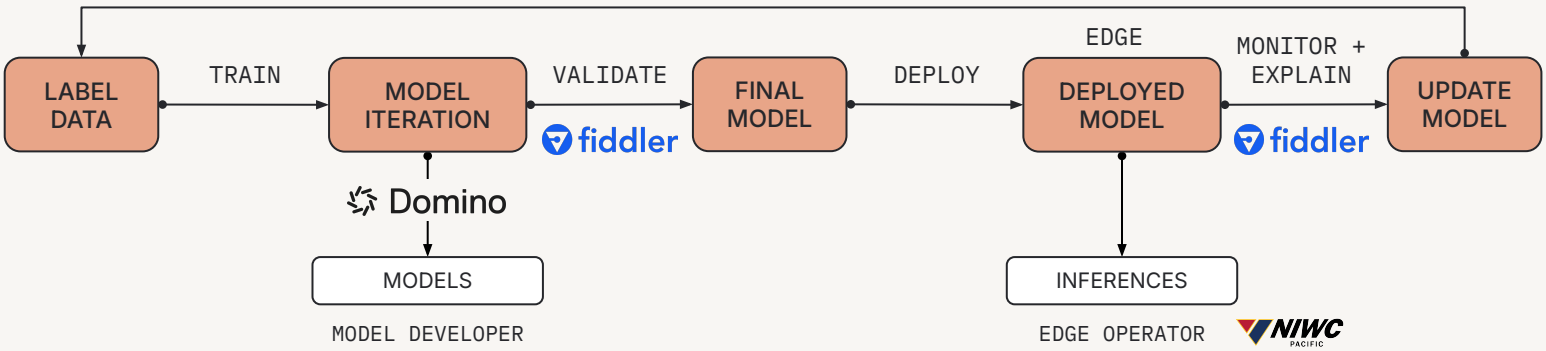

Fiddler and Domino are Trusted by:

DEFENSE INNOVATION UNIT

AMERICA'S NAVY

Fiddler and Domino are available via Carahsoft Contract Vehicles, Tradewinds Marketplace, and Cloud marketplaces.

Visit: fiddler.ai | domino.ai