Responsible AI Profile Legend

General Background

Description	A brief description of the AI feature or model's primary function and its intended use case.
Is PII used in the training or operation of this model?	Indicates whether Personally Identifiable Information (PII) is used in the training data or during the operation of the model.
Base Model	Specifies the foundational model on which the AI system is built. For example, "Azure OpenAI - GPT-4" or "Anthropic - Claude 2."
Model Type	Describes what kind of AI model is being used. For example, a Generative AI model vs a Predictive AI model.
Model Customization	Refers to the process of tailoring the model's application or functionality to meet specific product or customer needs, without changing the underlying model itself. This may involve configuring model parameters, adding proprietary data, or integrating additional features to optimize its use within the Alteryx ecosystem. Model customization does not imply changes to the core architecture or design of the third-party LLM.

Third-Party LLM Responsibility

To the extent that this product or feature utilizes a third-party LLM, please refer to the respective provider's documentation for information on their data handling practices. This document describes how Alteryx's product interacts with and uses the LLM, but the model's management of data is governed by the third-party provider.

Transparency and Explainability

Model Outputs Explained	Refers to the explanation of how the AI model arrives at its answers or decisions. This involves providing users with insights into the model's decision-making process, including the logic and factors influencing the model's output.

Human Agency and Oversight

Is the Feature Optional?	Indicates whether the customer can opt-in or opt-out of this specific feature.
Human in the Loop?	Explains whether there are mechanisms or tools in place that allow a human to review, change, or block the output of the AI before any final action is taken?

Trust and Accountability

Base model trained with customer data?	Was any customer data used to train the base model?
Training data anonymized?	Is the data being anonymized before it is being used to train the model?
Customer data shared with model vendor?	Is any customer data sent to the vendor that trains the base model?
Data deletion?	Refers to the process of removing specific pieces of data that have been stored within the model or used in its training. This could include removing training examples, user interactions, or other inputs that have been used to fine-tune or operate the model. If the product utilizes third-party LLMs, please refer to the provider's policies.
Data retention?	How long data related to the feature will be retained. This includes metadata about inputs or weights, interaction logs, and other relevant information, often for legal, auditing, or operational purposes. If the product utilizes third-party LLMs, please refer to the provider's policies.
Data processing location?	Refers to the physical or cloud-based location where the AI system processes data.
Data storage details?	Describes how and where data is stored within the system, whether on-premise, in the cloud, or on specific servers. For data stored by the third-party LLM, please refer to the provider's policies.
Data encrypted in transit and at rest?	Confirms whether the data is encrypted both during transmission and while stored.

Reliability and Safety

Logging and Auditing Mechanisms Available?	Does the feature provide logging and auditing tools that help a user understand what output the AI has produced and how the output was produced?
Guardrails?	Are there checks and guardrails in place that check for harmful output including ethical, bias, violence, hate, etc.?
Impact Assessment Conducted?	Confirms whether a systematic evaluation has been carried out to identify and mitigate risks associated with an AI system. This includes assessing ethical, social, legal, and operational impacts to ensure compliance with regulations, prevent potential harms, and implement safeguards that protect users and stakeholders.
Compliant with Applicable Regulations?	Verifies that the AI system adheres to relevant laws, standards, and regulatory requirements, including data protection, privacy laws, and AI-specific regulations, to operate legally and ethically.
Input/Output Consistency?	Evaluates whether the feature consistently produces the same output when provided with the same input.

Fairness and Inclusivity

Data Sources?	Identifies the origin of the data used by the AI system, such as internal databases, third-party providers, or user-generated content. For products utilizing third-party LLMs, this refers only to the data sources under Alteryx's control. Please refer to the third-party LLM provider for their data sources and usage practices.
Bias detection and mitigation in place?	Confirms whether mechanisms and tools are in place to detect, measure, and reduce biases in the AI system. This includes checking for potential bias in data, algorithms, and outputs, as well as applying strategies or safeguards to ensure fairness, equity, and non-discrimination in the system's decisions and actions.

Empower Social Good

Designed for Ethical Use?	Refers to the intentional design and development of the Alteryx assembled AI system, incorporating safeguards to ensure fairness, inclusivity, transparency, and accountability. Ethical use also considers the broader societal impact of the AI system, incorporating safeguards to mitigate risks to both its deployment and operation.