PARTNERSHIP ON AI

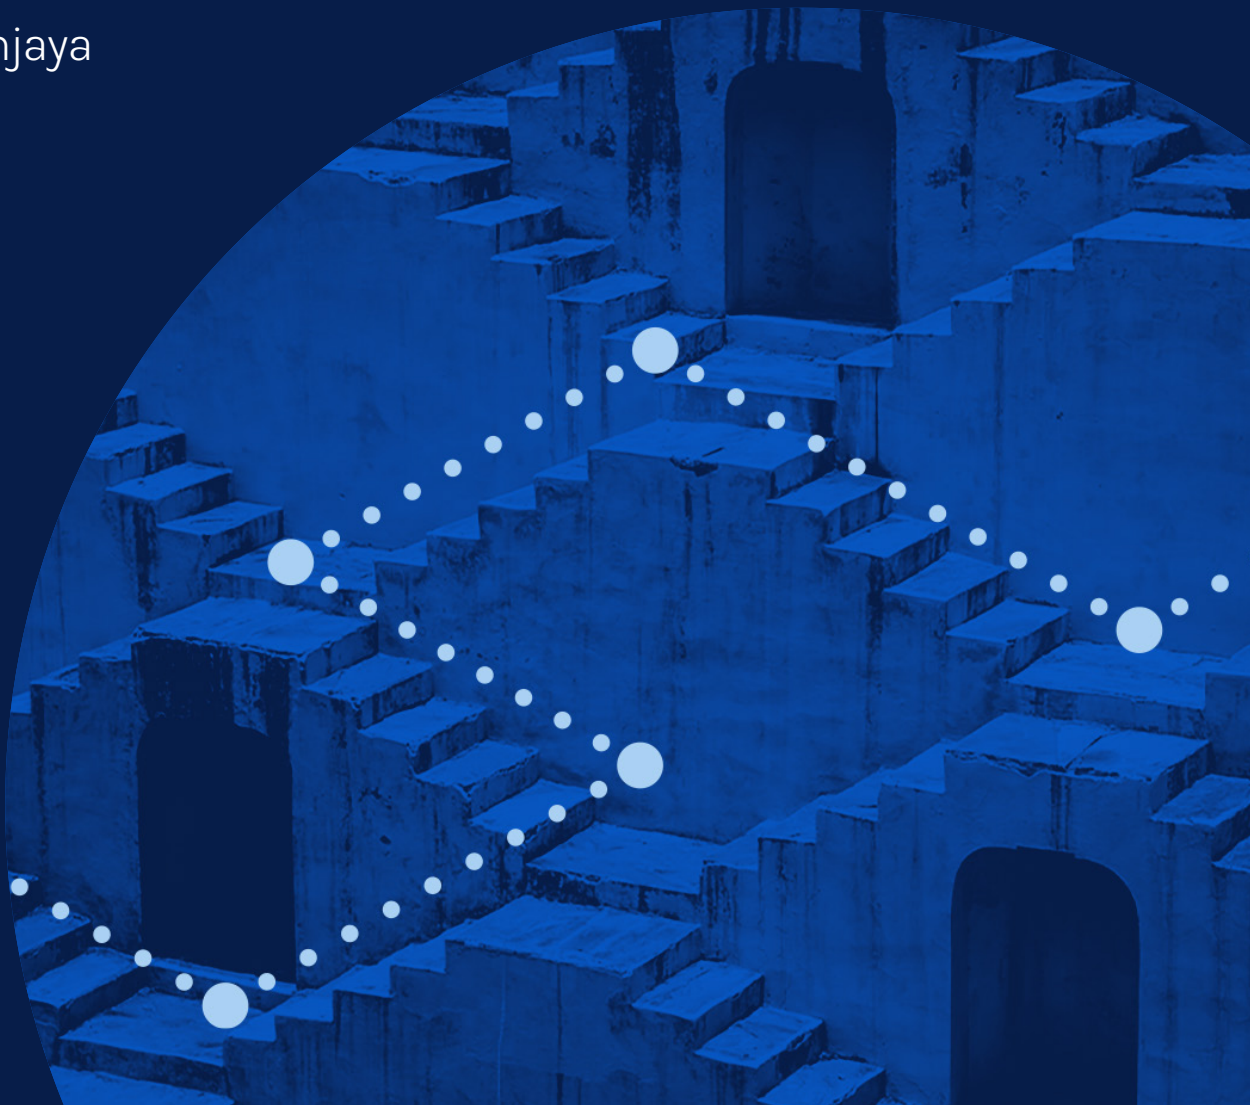
PILOT SUMMARY

ABOUT ML in Practice

The Example of Intuit

intuit

Albert Tanjaya



Executive Summary

Transparency is a widely recognized key principle of AI ethics which enables other AI ethics goals. The Partnership on AI's (PAI) ongoing ABOUT ML initiative focuses on one particularly promising aspect of transparency: the documentation of machine learning (ML) systems.

PAI's third ABOUT ML pilot was conducted in collaboration with Intuit, a large consumer-facing financial software and services company. The pilot was conducted by PAI's ABOUT ML team in partnership with members of the Intuit team to explore how their responsible AI intake review process can be further improved to streamline artifact and process creations. Initial conversations gave insight that technical teams had the autonomy to use a variety of documentation artifacts like JIRA and GitHub, and build various artifact-creating processes at different instances of a given model's lifecycle. Teams also used ad hoc handoff processes for information requests from both technical and non-technical teams, making knowledge sharing a decentralized process.

Further introspection was mainly done through workshops that covered these topics:

- In-depth reviews of current documentation artifacts and processes.
- Identifying and explaining gaps in documentation artifacts and processes.
- Facilitating discussions across various team members and departments.
- Human-centered design exercises such as persona- and user-journey-mapping to better understand external stakeholder needs and experiences.
- Start to build out the ideal state of documentation in theory based on the time spent together.

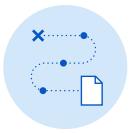
Ultimately, the pilot surfaced three actionable themes:



Standardization of the creation of documentation processes and artifacts across internal teams



Considerations for the variety of stakeholders' needs when interfacing with documentation



Embracing additional responsible AI values throughout the entirety of the model lifecycle

ABOUT ML Overview

Partnership on AI (PAI) is a non-profit partnership of academic, civil society, industry, and media organizations creating solutions so that AI advances positive outcomes for people and society. By convening diverse international stakeholders, PAI seeks to pool collective wisdom to make change.

PAI's ongoing ABOUT ML* initiative has focused on one promising approach to operationalizing transparency: the systematic documentation of machine learning (ML) systems. Transparency involves making a system's properties, purpose, and origin clear and explicit to users, practitioners, and other impacted stakeholders. Transparency is valuable for many purposes. Not only can transparency help these groups understand under what conditions it is appropriate to use the systems, transparency is a foundation for both internal accountability among the developers, deployers, and API users of an ML system and external accountability to customers, impacted non-users, civil society organizations, and policymakers.

* Annotation and Benchmarking on Understanding and Transparency of Machine-Learning Lifecycles

ABOUT ML aims to establish new norms on transparency through the identification of best practices for documenting and characterizing key components and phases throughout the ML system lifecycle from design to deployment, including annotations of datasets, algorithmic models, performance assessments, and maintenance requirements.

Having established a detailed reference document of these best practices, the most recent phase of the ABOUT ML initiative was to drive experimental pilots with members of our Partner community to:

1. Test and improve the usability of the ABOUT ML recommendations found in our reference document with implementation instructions as well as other tools in the resource library in real-world settings.
2. Share lessons learned from a range of use case scenarios to support further implementation by a variety of organizations.

RESOURCES

ABOUT ML Reference Document

ABOUT ML Resource Library

Intuit Pilot Overview

Documentation Overview at Intuit

As Intuit operates as a prominent industry player and a large-scale corporation, it is imperative to maintain comprehensive documentation to ensure its operations are seamless and functioning efficiently. In understanding their processes and values through several meetings before the pilot itself, pilot participants acknowledged that documentation ultimately:

- Enables effective knowledge transfer (e.g., onboarding new members/handing off projects)
- Improves collaboration and communication across teams
- Improves traceability of data and technical model information
- Improves reproducibility of model systems
- Improves ease of integration of RAI practices as teams build out models

Intuit's documentation approaches involve all teams participating in a Responsible AI intake process that makes the flow and discoverability of information and processes easier to reference back to. Part of this process is ensuring that their Responsible AI leadership team can create global rules of documentation for all teams so that information is interoperable across teams that might need it, especially for their legal teams who evaluate and assess for high-risk impacts towards their employees and customers.

In understanding this pilot and the specific model being assessed, this model team leader produced many documentation artifacts like JIRA, tickets, GitHub comments, Google Docs, etc. Documentation exists throughout this model's lifecycle, and there are a variety of ad hoc processes for information transfers between teams like data, legal, and AI governance teams. However,

Organization Profile

NAME	Intuit
WEBSITE	www.intuit.com
GEOGRAPHIC LOCATION	Global (headquartered in US)
COMPANY SIZE	10,000+
YEAR FOUNDED	1983
ORGANIZATION MATURITY	Mature • Established company with global reach and size • $14B+ Revenue (FY 2023)
ORGANIZATION DESCRIPTION	"Intuit helps consumers and small businesses prosper by delivering financial management, compliance, and marketing products and services. [They] also provide specialized tax products to accounting professionals, who are key partners that help [them] serve small business customers." *(From Intuit's September 12, 2023 SEC filing)*
INDUSTRY	Financial technology
AI'S INTEGRATION AND IMPACT AT INTUIT	The strategy is to be an AI-driven expert platform. • Revolutionize speed to benefit. • Connect people to experts • Unlock smart money decisions • Be the center of small business growth • Disrupt the small business mid-market
INDUSTRY REGULATION LEVEL	• Heavily regulated industry • Intuit's products utilize sensitive data containing personally identifiable information

this is not the same throughout other model teams, leading to decentralized documentation efforts within technical teams.

PAI & Intuit's Pilot Goals

Participating staff at Intuit were asked to set goals for their engagement with the pilot. They identified the following as their priorities for the collaboration:

1. Understand their current state of documentation for the ML model lifecycle
2. Learn about other best practices for documentation
3. Identify possible gaps in their documentation practices
4. Create a preliminary strategy to improve ML model documentation
5. Receive guidance on additional resources for ML documentation throughout the model lifecycle

These goals were agreed upon after a round of goal refinements with a shared understanding of the pilot's limitations. In this case, Intuit also sought to understand documentation as a means for compliance purposes for risk management. Unfortunately, this goal was outside the scope of the ABOUT ML pilot program at this stage of the work.

Pilot Assessment Profile

AI SYSTEM	Proprietary Confidence model for ID fields in tax documents, built in-house.
DOCUMENTATION INVOLVED	• Jira tickets • Slack • Google Docs/Sheets • GitHub repositories • Jupyter Notebook
PILOT TEAM COMPOSITION	Responsible AI and legal (1 participant) • Oversees the AI Governance process Model development (1 participant) • Develops in-house models AI Lens/Machine Learning Platform (2 participants) • Works on pillars of observability in the AI space (operations and explainability)
INTENDED AUDIENCE FOR DOCUMENTATION	Internal teams: • Policy teams (responsible AI and legal) • Technical teams (data scientists, model engineers) • Non-technical teams (sales, clients)

Exploring the Gaps

In exploring the gaps, pilot participants identified areas that needed strengthening through Workshop 1 (a deep dive at current documentation practices) and Workshop 2 (identifying and explaining gaps from Workshop 1). While the implementation of the Responsible AI review process can give newer development teams more alignment on what is being documented earlier in the model lifecycle, participants noted that for older model and data teams, it is harder to track model system and dataset information since different teams are at different maturity levels of documentation. As a result, different teams are not on the same level of documentation standards.

Additionally, participants attributed that the misalignment of documentation standards made the transfer of datasets and model systems across teams (technical and non-technical) difficult to track and explain. This brought up the question of how documentation might interface differently within the various teams that exist, from technical teams like

data scientists and model developers to non-technical teams like legal and end users.

This also posed an adjacent obstacle: building organizational-wide buy-in to utilize the Responsible AI review process effectively. From the pilot surveys, one entry noted that the creation of such buy-in should be on some level of consensus from all the involved teams. However, the different documentation practices make consensus-building and facilitating buy-in difficult.

Emergent Themes

1 Opportunities exist to prioritize and standardize documentation processes and artifacts across internal teams

In the pilot, a participant representing a model team expressed a strong commitment to high code documentation standards driven by efficiency, accessibility, and visibility. This team's internal documentation practices were robust despite minor obstacles like tool limitations and slow updates.

Workshops revealed that documentation within Intuit was usually driven by the needs of specific teams with accountability independently resting on team leaders and minimal clarity on standards across teams. Limited instances of knowledge-sharing and feedback exchanges between teams allowed teams autonomy over their processes, leading to a decentralized approach.

To help set a minimal level of standard, teams should be prepared to answer fundamental questions outlined in Section 3 of the ABOUT ML Reference Document, "Preliminary Synthesized Documentation Suggestions." This section can be used as a litmus to help identify what practices teams are already doing, what practices are missing, and what practices that have not been considered before would benefit from existing.

Teams should also embrace sociotechnical considerations to make models and data information more accessible to non-technical audiences within Intuit, as it may aid in identifying social harms and risks for Intuit and its end users. Opportunities include defining accountability (what does it look like? Who enforces it?), identifying teams with best practices to see what works and what could be improved, and defining the measures of success of internal documentation improvements.

2 Intuit has a clear vision for creating model cards and can build on that by considering additional stakeholders' needs for such documents

Theme 1 highlighted that documentation processes were often insular and occurred within silos, hindering explainability and accessibility. Participants agreed that silos created inefficiencies when non-model teams needed to access information about specific models. In the workshops, the team began mapping stakeholders who should have access to documentation artifacts. They identified key internal groups like AI and legal leadership, the AI Governance Committee, end users/customers, and internal/external auditors. Through this participants were able to focus on understanding how each stakeholder group interacted with the ideal model card, their information needs, and the value derived from a robust documentation artifact.

It is essential to consider perspectives from internal stakeholders (e.g., different departments and roles) and external stakeholders (e.g., auditors and customers) throughout the relevant stages of the ML lifecycle. Given that the ML systems Intuit builds rest in the financial industry, Intuit has a heightened need for awareness of how bias and discrimination historically have impacted various communities, necessitating consultation with highly affected but often overlooked groups during development. Collecting and integrating feedback on documentation is crucial for sustainable documentation practices.

3 While risk compliance is a priority for AI documentation, there is space to embrace additional responsible AI values throughout the model lifecycle

Risk compliance and management processes are absolutely essential in Intuit's industry for reasons stated in theme 2. Documentation needed for those processes needs to exist throughout any given model's lifecycle to help uncover instances of both technical and sociotechnical risks. Previous themes addressed the importance of involving the perspectives of internal and external stakeholders throughout the documentation process as well as having a clear minimum level of standard of what and how to document. In order to cultivate a successful documentation practice, teams should be aligned in the values documentation brings. These conversations and levels of involvement begin with setting institutional values, aligning teams' work with those values, and representing them in the transparency processes, including documentation. Documentation for ML lifecycles is not simply about disclosing a list of characteristics about the data sets and mathematical models within an ML system but rather an entire process that an organization needs to incorporate throughout the design, development, and deployment of the ML system being considered.

In the workshops, the team mapped documentation touchpoints where each participant's

team intersected with a model across its lifecycle. The team identified an opportunity to expand documentation practices throughout the model lifecycle. Participants also gathered insights on the overall sentiments, motivations, and barries different teams have to documentation at certain stages of the lifecycle. Here, the team also explored questions on what they already have and what they might need for an ideal state of documentation. Through this, participants built a high level overview of their ideal state of documentation which in this case, was their ideal model card. Sections 3.4.1 ("Suggested Documentation Sections for Datasets") and 3.4.2 ("Suggested Documentation Sections for Models") of the ABOUT ML Reference Document break up the data and model stages into smaller stages more specific than the ones the team explored in the workshops.

Conclusion

In the context of implementing responsible AI documentation, it's essential to view transparency not as a final objective but rather as an ongoing commitment to diligent practice. The Pilot is the beginning of what should be many strategic assessments of what documentation best practices may look like based on the learnings and research provided by the ABOUT ML reference library. It is not meant to be the end of conversations on documentation practices, nor is it meant to provide all the solutions to any given question that arises.

The team invites you to stay up-to-date with the ABOUT ML and documentation work.To learn more about joining our work on documentation, please please email Albert Tanjaya.

In summary, the pilot team was able to address and explore the gaps Intuit posed, such as:

- Who are the various audiences of documentation, and what do they need?
- How are documentation processes communicated and coordinated?
- How can Responsible AI practitioners at Intuit share values and build buy-in?

These questions were explored and discussed through the following workshops:

- In-depth reviews of current documentation artifacts and processes.
- Identifying and explaining gaps in documentation artifacts and processes.
- Facilitating discussions across various team members and departments.
- Human-centered design exercises such as persona- and user-journey-mapping to better understand external stakeholder needs and experiences.
- Start to build out the ideal state of documentation in theory based on the time spent together.

Through a last round of pilot feedback, one participant noted, "I enjoyed the group sessions and hearing about the struggles other teams have with documentation, and to hear their perception of what Intuit needs in order for better documentation to be realized."

Acknowledgments

Albert Tanjaya, Program and Research Associate of the ABOUT ML Program, led the synthesis of the pilot engagement outputs and the development of this summary document. He had support in the pilot execution and document creation from Dr. Christine Custis and Stephanie Bell, who supported some workshop sessions and the creation of the summary document. Hudson Hongo contributed editorial support and Neil Uhl contributed design support to the summary document.

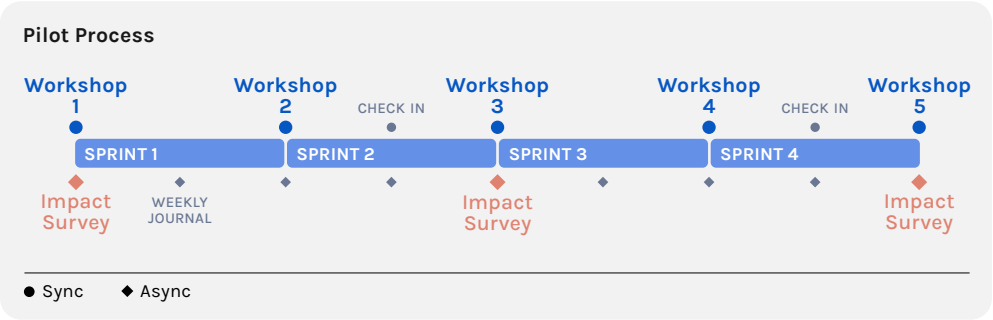
Thanks to Dr. Tina Park, Jiyoo Chang, and Kasia Chmielinski for building the design of the pilot processes that future pilots continued to use. Thanks to Intuit and the pilot participants for their time and resources for the completion of this pilot.

We are grateful to Michael Hind, Meg Mitchell, Kush R. Varshney, Hanna Wallach, and Jennifer Wortman Vaughan for contributing their expertise and valuable feedback as advisory committee members for the ABOUT ML work at Partnership on AI.

Appendix

Pilot Process Overview

The pilot was conducted by PAI's ABOUT ML team in partnership with the Intuit team. We convened five collaborative ● **workshops** in the spring and summer of 2023 to identify opportunities to improve current documentation practices. Workshops were supplemented with one-on-one individual ● **check-ins**. Asynchronously, participants completed pre-work for each workshop, ◆ **impact surveys** at regular intervals over the course of the pilot, and weekly ◆ **journal entries** to document personal insights. These components are described in more detail in the Pilot Content subsection below.



Pilot Team Composition

The pilot team was designed to reflect the stakeholders necessary to design, develop, and deploy responsible ML systems and products. The pilot team consisted of two PAI staff members (facilitators) and four Intuit employees (participants). The participants were drawn from the following teams at Intuit:

Number of Participants	Name of Intuit Team	Descriptor
1	Responsible AI and Legal team	The cross-functional team responsible for AI governance process
1	Model Development Team	The model developing team of the model explored in this pilot
2	AI Lens/ Machine Learning Platform Team	The team building intelligent systems to support data scientists and software engineers

Pilot Content

Workshops

Workshops were structured around pilot goals and designed to build upon previous sessions. The themes and findings are mainly drawn from the discussions and activities of these workshops, which help to identify opportunities to improve current documentation practices, which involved:

- In-depth reviews of current documentation artifacts and processes.
- Identifying and explaining gaps in documentation artifacts and processes.
- Facilitating discussions across various team members and departments.
- Human-centered design exercises such as persona- and user-journey-mapping to better understand external stakeholder needs and experiences.
- Start to build out the ideal state of documentation in theory based on the time spent together.

Impact Surveys

Participants were given longitudinal surveys at three points during the pilot: before the first workshop, after the third workshop, and after the last workshop. The surveys were sent to all pilot participants via Google Forms, and each member answered the surveys individually. The surveys were used as a "pulse check" to assess changes in the following areas throughout the pilot:

- Perceived value of documentation
- Benefits of documentation
- Costs of documentation
- Self-assessment of documentation artifacts and processes
- Challenges currently being faced
- Landscape of current documentation practices
- Effectiveness of current documentation practices

One-on-One Check-ins

PAI's facilitators had individual check-ins with the participants throughout the pilot as well as after its completion. We discussed what was going well, what was going not-so-well, and participants' learnings and insights so far. These optional one-on-one sessions gave participants the opportunity to reflect verbally on the pilot with facilitators.

One check-in occurred during the duration of the pilot process, and two exit one-on-ones occurred post-pilot.

Weekly Journal Entries

Participants were asked to reflect and write down thoughts, "ah ha" moments, or questions

that arose in their own work lives during the week before and after the workshops. These entries were unstructured and optional for pilot participants to help document their own individual experiences with the pilot.

One participant provided a few entries.

Additional Resources

ABOUT ML Reference Document

- Section 2.2 Documentation to Operationalize AI Ethics Goals
- Section 3.4.1 Suggested Documentation Sections for Datasets
- Section 3.4.2 Suggested Documentation Sections for Datasets
- Appendix Compiled List of Documentation Questions

ABOUT ML Resources Library

ABOUT ML AI Ethics values (Google Slides)