PARTNERSHIP ON AI

PILOT SUMMARY

ABOUT ML in Practice

The Example of Biologit

biologit

Jiyoo Chang



Introduction

Transparency has been widely recognized as a key principle of AI ethics, one that can also enable other AI ethics goals. The Partnership on AI's (PAI) ongoing ABOUT ML* initiative has focused on one particularly promising approach to operationalizing transparency: the documentation of machine learning (ML) systems.

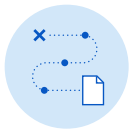
* Annotation and Benchmarking on Understanding and Transparency of Machine-learning Lifecycles

Through ABOUT ML, PAI has been bringing together a diverse range of perspectives to develop, test, and implement ML system documentation practices at scale. This Pilot Summary serves as description of how the ABOUT ML team sought to improve documentation practice at Biologit and the lessons we learned along the way.

Biologit created an extensive end-to-end ML documentation framework prior to launching its first product. After reviewing their documentation artifacts and processes, PAI led Biologit's team members through a series of exercises designed to deepen their understanding of the why, how, and who of ML documentation in order to better achieve their documentation goals. Through this collaboration, we identified several opportunities to make Biologit's already extensive documentation framework even more effective in practice:



Building a shared understanding on the value of documentation



Communicating and coordinating internal documentation processes



Clarifying the documentation audience and their needs

Background

ABOUT ML is a multi-year, multi-stakeholder initiative led by PAI that aims to promote transparent machine learning systems by identifying best practices for machine learning documentation. We have published and iterated on a Reference Document with implementation instructions as well as other tools in our Resource Library. PAI is now putting these guidelines and resources into practice by conducting Pilots with a select group of organizations who are looking to improve their documentation approaches across the ML lifecycle.

Our first pilot was with an Ireland-based startup, Biologit, which is the subject of this summary. Launched two years ago, Biologit is building an AI product that collects drug safety information from medical journal articles. This technology is related to a field called "pharmacovigilance," which is highly regulated in the European

PHARMACOVIGILANCE

Pharmacovigilance is defined by the World Health Organization as "the science and activities relating to the detection, assessment, understanding and prevention of adverse effects or any other medicine/vaccine related problem."

Organization at a quick glance

NAME	Biologit
WEBSITE	www.biologit.com
DOMAIN	Pharmacovigilance
YEAR FOUNDED	2019
ORGANIZATION MATURITY	Early stage startup
GEOGRAPHIC LOCATION	Headquartered in Ireland
INTENDED AUDIENCE OF AI SYSTEM	Drug safety professionals
ALGORITHMIC SYSTEM TYPE	Natural Language Processing

Union where the product is to be deployed. This motivated the team to prepare for downstream regulations and audits.

Beyond such requirements, Biologit's CTO wanted the company to practice transparency from its inception to distinguish it and gain trust from the public. He envisioned the company leading by example in the pharmacovigilance space with its AI transparency efforts. In line with this goal, he was highly motivated to build robust documentation practices throughout the development process and bring the rest of the team onboard. Putting this into practice, Biologit faced challenges from being at the forefront of innovation in a field that is relatively conservative towards technology adoption. They also faced challenges from implementing documentation that currently lacks industry-wide standardization in the AI field.

Challenges

Tools & resources that Biologit was already using

- Assessment List for Trustworthy Artificial Intelligence
- GAMP 5 Guide: Compliant GxP Computerized Systems
- Validating Intelligent Automation Systems in Pharmacovigilance
- ABOUT ML Resources
- IBM FactSheets

Prior to launching their first product, Biologit had developed an extensive set of documentation artifacts. These included public-facing blog posts, AI FactSheets, and compliance statements as well as internal documentation of the AI and software development processes. The Biologit team put in significant effort building these documents from the ground up, using online resources and guides such as PAI's ABOUT ML resources and the IBM FactSheet guide as references.

Even while following these guidelines meticulously, however, the team still had many questions come up: Who is the target audience of the documentation? What do they need to know? What is the value proposition of documentation? How do we effectively create and maintain documentation as a team? Having clear answers to these questions would help the team sharpen their focus when creating key documentation artifacts, ultimately better serving their audiences.

Ultimately, PAI saw an opportunity to help Biologit more clearly identify *why* and *how* we document and *for whom*.

Our Approach

At the time of our engagement, the organization had six team members who worked remotely in two different time zones. The team members brought a wide range of expertise in AI, pharmacovigilance, system validation, and quality assurance. As a startup preparing a product launch and juggling multiple priorities, Biologit did not have an opportunity to engage in in-depth discussions about their documentation practices.

Our hypothesis was that by bringing the team together to talk about and think about the bigger picture of documentation, we could make meaningful changes in the way the team thought about documentation and how they might implement it in their day-to-day work.

During the Pilot, the PAI team and the organization convened remotely in a series of biweekly workshops, which involved:

- In-depth reviews of current documentation artifacts and processes
- Facilitated discussions amongst cross-functional team members
- Human-centered design exercises such as persona- and user-journey-mapping to better understand stakeholder needs and experiences

Asynchronously, participants completed an online diary study each week to share updates and capture insights. Participants also completed an online survey at three points over the course of the Pilot to reflect on the value of documentation and assess their own internal documentation processes and artifacts.

Findings

We worked closely with the participating organization to align the project scope to the major challenges that they were facing with documentation. Here we outline the three key areas that we focused on and how we navigated each challenge.

1 Building a shared understanding of the "why" of documentation

For a practice that takes a lot of time and effort like documentation, implementing without a clear sense of purpose can leave opportunities on the table. For example, a participant said of one document, "We put a lot of time into this document and so far it's been useless." Here we saw an opportunity to zoom out and deepen the understanding of why documentation is important and necessary in order to make documentation more meaningful and valuable to the team members in their day-to-day work.

First, we discussed how documentation connects to the larger goal of bringing

transparency to ML systems. We began with an exercise of enumerating the success metrics and potential risks of documentation. For instance, we discussed how misuse of AI could be a risk of poor product documentation. As one participant remarked, "Poor product docs can lead to dropping the AI features being used, and can also lead to misuse of AI — a big risk!" On the other hand, high-quality documentation can help mitigate potential AI adoption risks and enable potential clients to use the AI product with a high degree of autonomy. We discovered that many of the successes and risks of documentation identified by the team were aligned with Responsible AI principles, such as accountability, trust, and explainability.

"The most valuable takeaway that we gained is the ability and process to look at documentation from multiple different lenses to judge its 'value.'"

We also discussed the specifics of how each documentation artifact brings value to its users. For example, product documentation increases customer trust and ensures compliance to industry standards by effectively communicating requirements and processes. Software development documentation helps the engineering team make onboarding and handoff quicker and training of new models easier. Discussing the benefits of specific documentation artifacts for specific audiences helped some of the team members evaluate and understand their value. As one participant reflected: "The most valuable takeaway that we gained is the ability and process to look at documentation from multiple different lenses to judge its "value.'"

2 Clarifying the documentation audience and their needs

Potential audiences for ML documentation can vary widely. Section 1.1.4.3 of the ABOUT ML Reference Document lists the following potential audiences:

- ML system developers
- ML system procurers
- Users of ML system APIs
- End users
- Internal Compliance teams
- External auditors
- Marketing groups

When the Biologit team was developing a wide range of documentation artifacts, it was not always clear who the intended audience was. One of the team members said, "It isn't very systematic who we are writing this documentation for." We saw a need to clarify the intended audience for the documentation artifacts and explore how documentation could best meet the needs of that audience.

One of the barriers was that the team had limited opportunity to directly interact with users before launching their products. It was difficult to clearly identify whom the documentation was written for and what information the reader might need from documentation.

During one of our workshops, we created user journey maps to deepen empathy for different users' needs and identify pain points throughout the journey of interacting with

documentation. Even without directly engaging with real users, this exercise gave the team an opportunity to think in the roles of specific users who might engage with their documentation artifacts, such as an end user, a new engineer joining the team, or a client quality assurance lead. This kind of journey-mapping exercise can be useful to teams that might not have easy access to their users and, to an extent, help the team build a better understanding of different users.

Our journey-mapping workshop inspired the team to test their documentation with their beta users. A participant shared a shift in understanding after discussions with their beta users:

> "From beta discussions, it's becoming clear [that] a good part of our audience will rely on a different type of documentation than FactSheets. We might need documentation that is more task-specific, matches people's understanding of AI or otherwise educates them, uses the right language, etc. As of now, FactSheets may be too geared towards the data science community in scope, terminology, etc."

Using this feedback, the participant said the team shifted their strategy to center the documentation around users' goals and understanding of AI:

> "In our documentation we discuss AI a lot in terms of what we did and how it works, but not enough in terms of how to approach adoption from the client side."

Lastly, while clarifying the intended audience for documentation, we made sure to impart a flexible mindset about audiences. Building documentation for flexibility is important in a constantly changing landscape. As one participant observed, "If the team grows, so will the audience and the author and reviewer."

3 Communicating and coordinating internal documentation processes

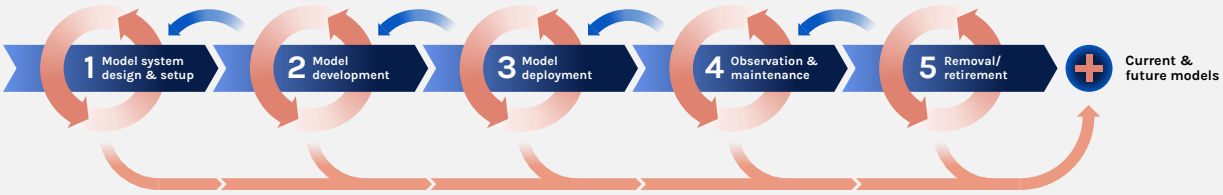
Documentation frequently requires collaboration between cross-functional team members. We explored how the current internal processes and communication patterns can support the creation and maintenance of documentation.

The team members have different roles and disciplinary backgrounds, bringing different knowledge bases, values, and operating modes. One of the participants described the differences between the two fields where their work intersected as follows: "AI is the wild west. Everything pharmacovigilance (PV) people do is informed by risk." These differences in background and role were compounded by the distribution of the team across multiple geographies, reducing opportunities for informal knowledge-sharing. There was no formal coordinated process to update staff from different teams on documentation relevant to

their work. Consequently, staff were often unaware of documentation done outside their team or by the company as a whole.

We began by analyzing current documentation artifacts to understand: 1) when documentation is created during development, and 2) which part of the ML lifecycle is documented. As we aligned documentation artifacts along the development lifecycle, the team noticed that "much of the documentation [was] happening during the release stage," an observation that opened up conversations around creating documentation throughout the ML development process. We also discussed at which points documentation should be created and updated during the AI lifecycle in order to enable implementation of documentation in their day-to-day work.

Key Process Considerations for Documentation
(ABOUT ML Reference Document Section 2.2.2)



In addition, we facilitated the team to openly discuss any ambiguities around documentation, addressing questions from them such as "What is this fuzzy piece of documentation and who is writing it?" Through these conversations, the team was able to achieve differentiation between client specs, product documentation, and user manuals. We enhanced the understanding of which teams develop which key documents.

There was also a need for more coordination for review, update, and maintenance of documentation. As one participant said,

> "It seems we [have not yet] clearly documented in our processes what will trigger the production and refresh of documentation. Most are created on an ad-hoc basis or often long after a release, creating the chance of doc production being missed and getting stale."

The team recognized a need for structured checkpoints and processes, such as a weekly meeting specifically for documentation or a Git-like workflow for releasing documentation along with code.

Conclusion

During the Pilot process, PAI and Biologit team members got together to think about the fundamentals of documentation:

- Why is documentation important?
- Who is the documentation audience and what do they need?
- How is Biologit communicating and coordinating documentation processes?

We then explored each of the questions through the following exercises:

- Building a shared understanding on the why of documentation
  - Connecting to the larger goals and values of the organization
  - Enumerating value to specific users
- Clarifying the intended audience and their needs
  - Mapping out user journeys
  - Gathering user feedback as early as possible
  - Having a flexible mindset about the audience
- Communicating and coordinating internal processes
  - Determining who is documenting what
  - Building in structured checkpoints
  - Covering the ML lifecycle comprehensively

Lastly, we heard from the participants about the importance of embracing ambiguity. As one participant observed:

> "Most often we do not know the end from the beginning and that is okay. What is important is to keep documentation an active process and keep doing it alongside the process even when we are not sure about the audience, stakeholders, and evolving requirements."

In a constantly changing landscape, it is crucial to build a flexible mindset about the shifting audiences and needs, while actively documenting to prepare for the future.

We hope this Pilot Summary will help lower barriers to implementing documentation, build a larger community around ML documentation, and contribute to discussions around responsible AI practices.

In our architecture of the Pilot, we infused the tenets of Inclusive Research and Design to explore the usefulness of ABOUT ML recommendations at organizations attempting to implement responsible AI principles into practice. In accordance with PAI's Theory of Change, we have several programs and workstreams validating research and theory through pilots and we intend to strengthen the institutional approach in the future.

With the completion of this ABOUT ML Pilot, we are continuing to incorporate our lessons learned to solidify and standardize our approach for future Pilot cohorts. Our ongoing learnings will also help identify what practices work best for different organization types and iterate on guidance as additional Pilots are launched.

To learn more about joining our documentation work, please email Jiyoo Chang or Albert Tanjaya.

Acknowledgments

We are grateful to Michael Hind, Meg Mitchell, Kush Varshney, Hanna Wallach, and Jennifer Wortman Vaughan for contributing their expertise and valuable feedback to this work. We would like to thank Kasia Chmielinski for leading the project as a visiting consultant at PAI, Dr. Tina Park for providing insights as the Head of Inclusive Research and Design at PAI, and Dr. Christine Custis for overseeing this project as the Director of Programs and Research at PAI.