

The Duke AI Health Community of Practice: A Model for Increasing AI Proficiency and Critical Understanding

Whitney Welsh, PhD^{1,2}, and Shelley Rusincovitch, MMCI, FAMIA^{1,3}


¹Duke AI Health, ²Duke Social Science Research Institute, ³Duke Clinical and Translational Science Institute



Introduction

The **2024 “State of AI” landscape survey** conducted for Duke University by the Center for Computational Thinking revealed a **U-shaped curve** of AI research at Duke—most AI research being either recent (less than 1 year) or over 10 years old, but not many projects in the middle (see Figure 1). We wondered, **how do people move from one side of the curve to the other** (from novice to expert), and **how can we build up the middle** so that more people have both the skills to use AI responsibly as well as the critical thinking skills to assess emerging trends in AI?

Figure 1: AI research projects at Duke by project length (from State of AI survey)



These results were very pertinent to our early work supporting the Duke community in AI applications and understanding foundational AI knowledge, and we hypothesized that the **Duke AI Health Community of Practice**—that is, a group of learners, practitioners, and experts in the fields of machine learning, artificial intelligence, and health data science who meet regularly and learn from each other—might be a mechanism by which people could **increase their understanding of AI** and **stay current** in a rapidly evolving field.


Using Wenger-Trayner et al.'s (2023) framework for evaluating communities of practice, which describes how social learning impacts practice and creates value, we developed a **logic model** that maps the **current activities** of the Duke AI Health Community of Practice, their **anticipated impacts**, and **how all the components relate to one another**.

Early results from recent evaluations of several component activities suggest that the Duke AI Health Community of Practice is achieving many of its short-term objectives, and may serve as a blueprint for other programs seeking to increase AI understanding and, ultimately, to support responsible AI development and thoughtful adoption.

What is a Logic Model?

A **logic model** is a visualization that maps the hypothesized causal chains between (and within) program elements and anticipated outcomes (see Figure 2).

Figure 2: Basic logic model structure



A logic model consists of four basic components:

- Inputs:** The resources required to do the activities (e.g., funding, staff, equipment, etc.)
- Activities:** Actions the program undertakes to bring about the desired results
- Outputs:** The products, goods, or services produced by the activities
- Outcomes:** The changes or benefits that are expected to result from the activities and outputs; models often differentiate between short- and long-term or direct and indirect impacts

Logic models can be used for **program planning**, by working backward from desired outcomes to develop interventions; to guide **evaluation**, by helping evaluators make strategic choices about which outcomes are most important and how to measure them; and to **communicate** both what a program is trying to accomplish and how to assess its performance (Frechtling 2007; McLaughlin & Jordan 2010; Wyatt Knowlton & Phillips 2009)

Evaluating a Community of Practice

A **community of practice** is a group of active practitioners in a given field who regularly socialize with each other and thereby increase their knowledge and expertise of that field by learning from one another's experiences and engaging in collective problem solving (Wenger, McDermott, & Snyder 2002; Wenger et al. 2023).

Evaluating a community of practice can be difficult because its impact usually manifests in settings outside of the community of practice (i.e., when members are independently playing their skills), the effects are often the result of incremental changes that add up over time, and there is often a time lag between interaction, implementation, and results.

Evaluating a Community of Practice (cont'd)

Wenger-Trayner et al. (2023) offer a framework for evaluating communities of practice in terms of value creation. They identify four basic types of value that a community of practice can generate:


- Immediate value:** The value of engaging with colleagues in and of itself, regardless of whether or not anything further comes from it.
- Potential value:** The value of new insights, ideas, social connections, and so forth, generated through engagement that have the potential to impact practice.
- Applied value:** The value gained from applying new ideas or methods to practice, regardless of results (e.g., experience and skills gained from testing ideas or adapting methods to new contexts).
- Realized value:** The value of the actual results that stem from changes in practice

They also highlight four additional types of value that may arise from the interaction of the community of practice and the context in which it is situated:

- Orienting value:** The value of enriching the community of practice by incorporating the larger social landscape (e.g., what others are doing, recent trends and research, potential audiences)
- Enabling value:** The value gained from improving skills (e.g., fundraising, organizing, networking, etc.) which facilitate community of practice activities
- Strategic value:** The value that derives from engaging with key stakeholders to expand the reach and impact of the community of practice (e.g., fostering a broader understanding of the community of practice's work and relevance, creating and nurturing relationships)
- Transformative value:** The value that comes from transforming people and systems beyond the community of practice, which in turn transforms its context

Learning in a community of practice creates flows across the different types of value (see Figure 3). For instance, convening a community of practice creates immediate value. The conversations that occur during that convening may give rise to a new idea (potential value), which a member may then test (applied value) and discover it produces positive results (realized value). These flows, in turn, describe how social learning can impact practice, and serve as a foundation for evaluation.

Figure 3: How communities of practice create value (from Wenger-Trayner et al. 2023)




Duke AI Health Community of Practice Logic Model

Creating the logic model (Figure 4) was an **iterative process**, involving much discussion, and multiple rounds of revisions to and rearrangement of all the components.

- We began by sketching in the **activities and outputs** from the Duke AI Health Community of Practice project portfolio. For this high-level model, the projects themselves serve as the outputs, though in future component evaluations, the specific activities and outputs associated with the project or event will be delineated in greater detail. Which projects to include and how to categorize them was subject to change until near the end of the process.
- We then drew on our previous evaluations of the components to fill in some of the **inputs and outcomes**. To capture any additional **inputs**, we thought about what another group would require to replicate the program. Some were immediately apparent (e.g., funding, expertise, logistical support, strategic partnerships), but others took several iterations to identify.
- To elaborate the **outcomes**, we reviewed the objectives in our project descriptions and considered them in light of Wenger-Trayner et al.'s (2023) **value creation framework**. We then thought more broadly about what we wanted to accomplish with the community of practice. The resulting list contained outcomes with varying degrees of immediacy and abstraction, as well as differing, yet sometimes overlapping, levels of impact.
- To accommodate the range of outcomes, we separated the more immediate, **direct effects** from the longer term, less specific **ripple effects**. We also divided the outcomes into **individual**, **institutional**, and **extra-institutional levels** using different colored boxes to allow for overlap.
- Finally, we wanted to acknowledge that the community of practice is not a static program or one time intervention, but rather continuously evolving and adapting over time as the community learns. To that end, we added a **feedback loop** connecting the outcomes back to the inputs.

The resulting model represents a current snapshot of the Duke AI Health Community of Practice—what we are doing, what we are trying to accomplish, and how it all fits together—that will be an important touchstone for our ongoing evaluation efforts and program planning going forward.

Figure 4: Duke AI Health Community of Practice Logic Model



Evaluation Results

In addition to conducting ongoing evaluations of specific activities*, in March 2025 we invited Community of Practice members to participate in a more general **feedback survey** to **gauge the broader impact** that the AI Health Community of Practice had in the past year, and **identify areas that need improvement**.

Of the 59 respondents who completed surveys, 48 reported interacting with Duke AI Health at least once in the last year. Table 1 shows to what extent (A lot, Some, A little, Not at all) those 48 thought that their interactions with Duke AI Health had contributed to each of the **value-creation measures** from the community of practice evaluation framework. The column headed “Any impact” is the sum of respondents who selected A lot, Some, or A little.

The primary finding is that over **95% of respondents** who interacted with Duke AI Health at least once in the last year reported that those interactions **changed the way they thought about AI or how AI can be used**, at least a little. However, the numbers overall were fairly strong; over 60% reported any impact on each of the measures.

Also of note is that over **85% of respondents** who interacted with Duke AI Health at least once in the last year reported that it **informed their work**. When asked how Duke AI Health had informed their work, the respondents offered a broad spectrum of answers:



These results suggest that the Duke AI Health Community of Practice is succeeding in **increasing participants' understanding of AI**, and furthermore, that they are **applying that knowledge in their work**. As the Dean of the School of Medicine recently communicated, Duke Health and the University are actively investing in AI tools for the larger goal of **supporting our people**. Through our work with the Duke AI Health Community of Practice, we are proud to support that institutional priority and look forward to continuing to grow and evolve in responsible ways in order to continue supporting **what matters most**.

*Impact reports for Duke AI Health Community of Practice activities are available at: <https://aihealth.duke.edu/proj-aih-cop-eval/>

Table 1: To what extent interactions with AI Health contributed to each metric, for respondents who interacted at least once with AI Health in the last year (N=48)

	Any impact	A lot	Some	A little	Not at all	No answer
Changed how you think about AI or how AI could be used	95.8%	39.6%	43.8%	12.5%	4.2%	0.0%
Informed your work	85.4%	14.6%	56.3%	14.6%	10.4%	4.2%
Better access to advice and resources when dealing with challenges	83.3%	16.7%	33.3%	33.3%	16.7%	0.0%
Opportunities for professional growth	81.3%	18.8%	35.4%	27.1%	18.8%	0.0%
Gaining new skills	77.1%	20.8%	33.3%	22.9%	22.9%	0.0%
Gaining confidence	75.0%	16.7%	31.3%	27.1%	20.8%	4.2%
Greater sense of belonging and being part of a team	72.9%	12.5%	31.3%	29.2%	25.0%	2.1%
Networking	68.8%	16.7%	27.1%	25.0%	29.2%	2.1%
Pathways to new professional roles	66.7%	10.4%	37.5%	18.8%	33.3%	0.0%
Greater job satisfaction	62.5%	16.7%	25.0%	20.8%	37.5%	0.0%

References

- Duke State of AI Survey, Center for Computational Thinking. April 2024.
- Frechtling, Joy A. 2007. *Logic Modeling Methods in Program Evaluation*. San Francisco: John Wiley & Sons.
- McLaughlin, John A. and Gretchen B. Jordan. 2010. “Using Logic Models.” Pp. 55-80 in *Handbook of Practical Program Evaluation* (Third Edition), Joseph S. Wholey, Harry P. Hatry, and Kathryn E. Newcomer (eds). San Francisco: Jossey-Bass.
- Wenger, Etienne, Richard McDermott, and William M. Snyder. 2002. *Cultivating Communities of Practice*. Cambridge: Harvard Business Press.
- Wenger-Trayner, Etienne, Beverly Wenger-Trayner, Phil Reid, and Claude Bruderlein. 2023. *Communities of Practice Within and Across Organization: A Guidebook*. Sesimbra, Portugal: Social Learning Lab.
- Wyatt Knowlton, Lisa and Cynthia C. Phillips. 2009. *The Logic Model Guidebook*. Thousand Oaks, CA: Sage.

Special thanks to **Matt Hirsche** and the **CCT** for allowing us to reproduce Figure 1