Making Structural Labor Visible: A Protocol for Attribution in Modular AI Systems

Muyang Li
muyang.li@mail.mcgill.ca
McGill University

1. Introduction: Who Gets Credit for AI's Reasoning?

When a medical AI system uses a cardiologist's diagnostic flowchart across thousands of patient assessments, who gets credited?

Currently: the prompt engineer.
Not the cardiologist whose logic enables safe triage.
Not the compliance expert whose checklist ensures regulatory alignment.
Not the educator or lawyer whose logic scaffolds personalized tutoring or legal review.

As modular AI systems scale, they increasingly rely on embedded logic authored by subject-matter experts. Yet this structural labor remains attributionally invisible. While prompt engineering has gained recognition, deeper epistemic contributions are routinely abstracted away.

This reflects a core concern in CSCW: collaborative infrastructures often rely on labor that goes uncredited and unrewarded (Star & Strauss, 1999). In response, I propose treating structural logic design, such as triage frameworks, legal classifiers, and policy trees, as a first-class contributor role in AI development.

To support this, I introduce the *Prompt Revenue Mapping Protocol (PRMP):* a lightweight, modular framework for embedding attribution into logic modules reused across systems. PRMP enables traceability, recognition, and optional revenue sharing for the domain experts whose reasoning becomes infrastructure.

2. What Is PRMP?

PRMP allows logic modules, such as workflows, checklists, or taxonomies, to be treated as referenceable, role-tagged, and recursively attributable units. It supports multi-role attribution (e.g., domain expert, adapter, UX formatter), recursive credit flows across forks and downstream reuse, metadata for identity and licensing, and optional dispute resolution.

Example: Logic in Medicine and Law

When Dr. Smith's triage flowchart is embedded in an AI assistant, PRMP:

- Registers the logic module with role-tagged metadata
- Tracks each reuse or fork via a directed graph
- Routes attribution (and optional revenue) back to her
- Ensures traceability across downstream tools

Similarly, when a patent examiner’s logic for novelty classification is used in a legal AI assistant, PRMP ensures their decision tree remains linked to future adaptations, supporting not just reuse but also recognition.

3. Structural Labor as Articulation Work

This proposal builds on CSCW’s analysis of invisible labor. Star and Strauss (1999) describe *articulation work* as the coordination that enables systems to function, yet often goes unrecognized.

In AI workflows, structural logic is articulation work. It ensures reliability, compliance, and contextual fit, but it remains buried inside prompts or pipelines. PRMP makes this labor legible and accountable, bridging design justice with infrastructure studies and epistemic fairness (Birhane, 2021).

4. Why PRMP Matters for Responsible AI and CSCW

PRMP engages key themes in CSCW and responsible AI development:

- Consent and credit: Domain logic is reused without contributor input or visibility.
- Epistemic justice: Prompt engineers are credited; domain experts are not.
- Modular reuse: As logic and datasets become composable, metadata must follow function.
- Infrastructure transparency: Like SBOMs or contributor taxonomies, PRMP embeds recognition into the runtime layer.

5. Anticipating Challenges

PRMP faces practical challenges: integration overhead in current toolchains, incentive misalignment with systems benefiting from uncredited labor, and ambiguity in defining roles and attribution weights. Still, as AI becomes infrastructure, we risk losing expert participation without attribution systems.

6. Goals for the Workshop

I hope to contribute design feedback on PRMP's architecture, ideas for integration with tools like Hugging Face and LangGraph, governance models for licensing disputes, and participatory approaches to embed under-credited contributors. I welcome collaboration with researchers working on dataset documentation, AI transparency, and invisible labor in digital systems.

7. Final Provocation

What if we treated logic scaffolds the way we treat datasets—
as collaborative artifacts, worthy of consent, credit, and licensing?

We’ve built models that learn from human logic.
Now we need systems that honor the humans who created that logic.

References

1. Birhane, A. (2021). *Algorithmic injustice: A relational ethics approach*. Patterns, 2(11), 100347. https://doi.org/10.1016/j.patter.2021.100347

2. Brand, A., Allen, L., Altman, M., Hlava, M., & Scott, J. (2015). *Beyond authorship: A taxonomy of contributor roles*. Information Standards Quarterly, 27(1), 15–19.

3. DataCite Metadata Working Group. (2021). *DataCite Metadata Schema Documentation for the Publication and Citation of Research Data and Other Research Outputs* (Version 4.4). https://schema.datacite.org/

4. Groth, P., & Moreau, L. (2013). *An Overview of the PROV Family of Documents*. W3C Working Group Note. https://www.w3.org/TR/prov-overview/

5. He, H. (2024, January 16). *LangGraph: A Declarative Framework for Building Production-Ready LLM Applications*. LangChain Blog. https://blog.langchain.dev/langgraph/

6. Liu, P., Yuan, W., Fu, J., Jiang, Z., Hayashi, H., Neubig, S., & Neubig, G. (2023). *Pre-train, Prompt, and Predict: A Systematic Survey of Prompting Methods in Natural Language Processing*. ACM Computing Surveys, 55(9), 1–35.

7. National Telecommunications and Information Administration (NTIA). (2021). *The Minimum Elements for a Software Bill of Materials (SBOM)*. U.S. Department of Commerce. https://www.ntia.gov/files/ntia/publications/sbom_minimum_elements_report.pdf

8. Star, S. L., & Strauss, A. (1999). *Layers of Silence, Arenas of Voice: The Ecology of Visible and Invisible Work*. Computer Supported Cooperative Work (CSCW), 8(1–2), 9–30. https://doi.org/10.1023/A:1008651105359