



AI in the Loop: A Summary Report

Dr. Jerry Laurienti wrote a thought-provoking article, '[How the Intelligence Community Can Shape AI Trust With New Terminology](#),' for *The Cipher Brief* on 30 September 2026. He is a former Senior Intelligence Service officer, having served as CIA Deputy Associate Director and a member of the CIA Leadership Board, and is currently an NGA account manager at Leidos.

While much of his discussion revolves around the Intelligence Community and AI, the following comments focus on the nomenclature being used. The term "Human-in-the-loop", which is in vogue, subordinates humans to AI; the term implies that "the loop" belongs to AI and that humans will somehow be included. Laurienti prefers the term "AI-in-the-loop" that represents a purposeful attempt to reflect what we want to achieve: using AI when and where we elect to do so.

"AI-in-the-loop" is not a new expression and has been used in various contexts over the past four years, mainly to argue that humans should remain in charge while AI operates as part of a human-directed process. Combining the strengths of data-driven machine learning models with the contextual understanding of human experts was initially referred to as "Human-in-the-loop", but academics started pointing out that this was a misnomer, as the human is in control of the system, while the AI is there to support the human. The over-emphasis on the machine learning component's performance tended to neglect the human expert's critical role; hence, the more appropriate term would be "AI-in-the-loop", or a similar term.

Much of the current discourse treats "human-in-the-loop" as an AI governance approach, yet academia, government, and industry have not clearly defined the mechanisms by which humans are expected to oversee or collaborate with AI-enabled systems. Indeed, inconsistent uses of terms like "human-in-the-loop" and "human-AI teaming" create confusion about what type of oversight and collaboration is intended or even needed.

Increasingly, we are seeing documents stating that personnel remain responsible and accountable for the analysis, decisions, and outcomes derived from insights gleaned using AI. Users are required to understand and be accountable for AI use. Personnel must be able to review, challenge, and reject or replace AI-derived recommendations and findings when appropriate. Standards and policies on AI and AI Assurance are being developed as we seek the responsible adoption of the AI across all domains, while also elevating the human role. The mantra has become: "AI does not replace human judgment".

At the same time, we are seeing rapid procurement and a desire to speed up the adoption of innovative technologies, with increasing unease about just what is being unleashed.

Adopting new terminology now would be more than symbolic; it would help shape how we approach and use AI and ensure AI augmentation at speed and scale is part of the human toolkit, not that humans are just “in-the-loop.”

Thus, it is important to formalise AI-in-the-loop, or a related term, to set a tone of expectation and clarity that agentic autonomy must begin with human intent and be wrapped in human decision. Procurement decisions should make it clear that development and implementation must include a workflow that explicitly highlights human ownership of AI adoption. Even when AI does the heavy lifting across data sets that humans can’t process alone at speed and scale, it must do so on our behalf.

Jerry Laurienti is prepared to be flexible about just what term is used; it doesn’t have to be AI-in-the-loop, and he offers other suggestions:

- *Human-Led AI (HLAI)* — simple, direct, and clear about who leads the process.
- *Human-Guided AI (HGAI)* — emphasises direction and oversight, though it may sound less decisive.
- *Human-AI Teaming (HAIT)* — widely used, though it can imply a more equal partnership than intended, and an awkward acronym.
- *Human-Directed AI (HDAI)* — stronger than “guided,” emphasising human authority.
- *Human-Commanded AI (HCAI)* — strong human primacy, though perhaps too military in tone for broad use.
- *Human-Owned AI Workflow (HOAIW)* — inelegant, but explicit that the workflow belongs to the human.

Gary Waters

13 October 2026