



ICONOGRAPHIC SEMANTIC STACK: FIVE LAYERS FOR ICONS THAT SERVE BOTH HUMAN AND AGENT READERS

Hrozian S., Staff Product Designer, Speak, San Francisco, California, USA

Morozova L., Kharkiv National University of Radio Electronics, Department of MST

Abstract. *This thesis introduces the Iconographic Semantic Stack (ISS) – a five-layer framework for icons in interfaces now read by both humans and AI agents. Drawing on Reddit’s icon governance, Material Symbols, and emerging Model Context Protocol patterns, the framework names visual form, accessible label, behavior token, provenance, and intent metadata as the layers that determine whether an icon survives agent inference.*

Keywords: *iconography, machine experience, design systems, accessibility, model context protocol, semantic design.*

For thirty years an icon was a picture. The designer’s job was to make it readable to a human eye and a human screen reader. Then 2026 happened. An icon is now also a token an AI agent reads before any human sees the screen [1-3]. I noticed this first in agent traces. An agent asked to upvote a thread inside a community had to choose between two visually similar arrow controls. It picked the more prominent one. It was the wrong one. The user paid the cost.

This thesis introduces the Iconographic Semantic Stack (ISS). Five layers that decide whether an icon survives agent inference. The layers are not new individually. The contribution is naming them as one stack and identifying what fails when a layer is missing (Table 1, Figure 1).

Table 1 – Iconographic Semantic Stack (ISS): five layers and the cost of omitting each

Layer	Definition and cost of omission
Visual form	The rendered glyph: stroke, terminal, optical correction. Omission cost: human illegibility.
Accessible label	ARIA / screen-reader name. Omission cost: WCAG failure plus LLM mislabeling.
Behavior token	Semantic identifier tying the icon to an action. Omission cost: agent triggers wrong action.
Provenance	Origin marker – system, custom, or AI-generated. Omission cost: brand drift, attribution failure.
Intent metadata	Natural-language description of when and why the icon is used. Omission cost: agent picks visually similar but semantically wrong icon.

The first four layers are conventional but newly consequential. Visual form is the historical cost of unreadability [4]. Accessible label was once an accessibility requirement [5]. Now it is the string an LLM agent parses as the icon’s name. Behavior token ties the icon to a system action, not just an English word. Provenance marks the icon as system, custom, or AI-generated. This last one matters more than it used to, especially as generative tools begin producing icons on demand and model-reporting practices expand to UI artifacts [6, 7].

Intent metadata is the fifth layer. Most teams have not shipped it. It is a natural-language description of when and why the icon is used. Agents need it most. Design systems carry it least. It is the difference between an agent that triggers the wrong



action and one that asks for clarification. The Model Context Protocol [8] provides a transport for this metadata. It does not specify the shape. Teams will need to invent local conventions before standards emerge.

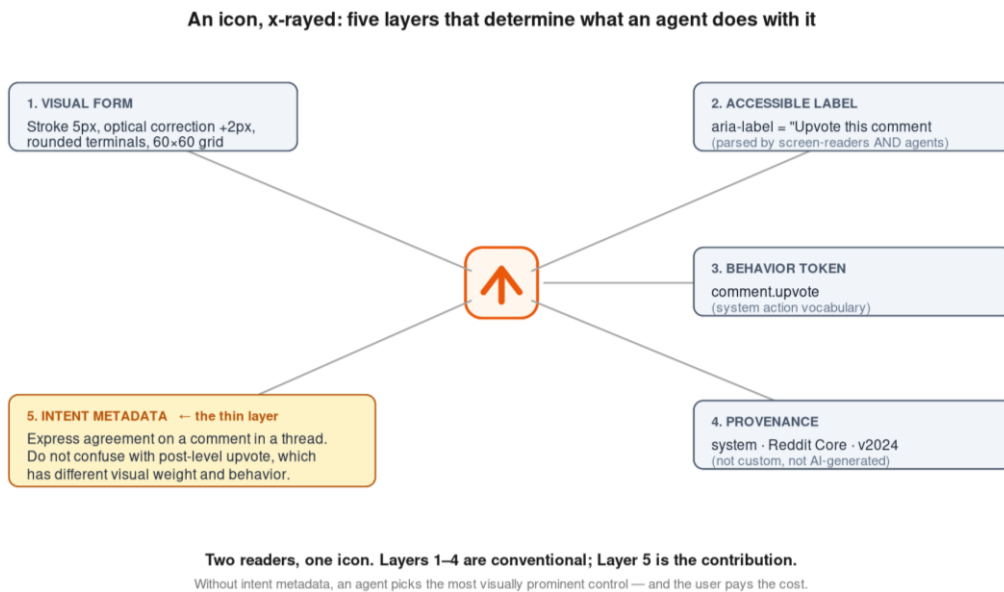


Figure 1 – A Reddit upvote control x-rayed: layers 1-4 are conventional; layer 5 is the contribution

None of this is free. Every layer adds bytes to the bundle and governance overhead someone has to maintain. Intent metadata is the most expensive layer, because someone has to actually write it, and keep writing it as the system evolves. The payoff is asymmetric. An agent that takes the wrong action damages trust faster than one that takes the right one builds it. The honest answer is to adopt incrementally. Visual form and accessible label are already standard. Behavior token and provenance can ship next. Intent metadata is the long-term work.

Every icon now has two readers. Most design systems still ship for one. I think this changes within a year, mostly because the consequences of an agent triggering the wrong control are obvious the first time it happens to someone important. Until then the work is unglamorous. Write the metadata. Name the tokens. Audit the provenance. There is no clever shortcut here that I have found.

References

1. Borovynska, Y., & Vovk, O. (2024). Investigating the vision of AI driven website builder in user interface components. *Jóvenes en la ciencia*, (26). <https://www.jovenesenlaciencia.ugto.mx/index.php/jovenesenlaciencia/article/view/4233/3714>.
2. Kaluhin, N., Vovk, O., & Chebotarova, I. (2024). The impact of artificial intelligence on future of humanity. *Jóvenes en la ciencia*, (26). <https://www.jovenesenlaciencia.ugto.mx/index.php/jovenesenlaciencia/article/view/4235/3716>.
3. Chebotarova, I., & Cherkashyna, H. (2025). Font management with artificial intelligence. *Memorias de SYNTOPIA*. (p. 32-33).
4. Apple Inc. (2024). *Human Interface Guidelines: SF Symbols*. <https://developer.apple.com/design/human-interface-guidelines/sf-symbols>.
5. W3C. (2023). *WAI-ARIA Authoring Practices Guide (Version 1.2)*. World Wide Web Consortium. <https://www.w3.org/WAI/ARIA/apg/>.
6. Google. (2024). *Material Symbols and icons*. Google Design. <https://fonts.google.com/icons>.
7. Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model cards for model reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (pp. 220–229). ACM. <https://doi.org/10.1145/3287560.3287596>.
8. Anthropic. (2024). *Model Context Protocol specification*. <https://modelcontextprotocol.io>.