THE DECISION SCIENCE SERIES

joe worden marketing

The Decision Sequence

Designing Engagement Strategies Based on How Consumers Decide.

Your Buyer Outsourced Their Memory. Now What?



DS 17

Cognitive Offloading

–A Strategic Resource by Joe Worden | More @ www.jworden.com

DECISION ARCHITECTURE RESEARCH LIBRARY

The AI-Era Research · Entry 17

Cognitive Offloading and the Illusion of Understanding

Sparrow, Liu & Wegner · Risko & Gilbert · Fisher, Goddu & Keil · Rozenblit & Keil · Carr · Wolf

Questions this paper answers

- How can we tell whether our team understands the AI-assisted material we publish?
- How do we handle a well-briefed prospect who has not worked through implementation?
- What should our content explain beyond a summary a buyer can get from AI?
- How can we use AI for production while keeping the underlying judgment in-house?

Definition. Cognitive offloading is what happens when we hand a chunk of thinking to something outside our heads, jotting the number instead of holding it, setting the reminder instead of remembering. The illusion of understanding is the familiar, stubborn habit of feeling like we've got something nailed down when, in truth, our grasp is thin. Put the two together and a pattern appears: the very tools that lighten the mental lift can also quietly puff up our confidence. When a machine produces a tidy, plausible explanation on demand, it becomes easy to confuse possession of an answer with real comprehension.

The Glow

Let me start with a confession, since this entry is going to spend some time on other people's overconfidence.

Not long ago I asked an AI assistant to explain a structured product I needed to understand for a client conversation. The explanation came back clean, organized, confident. I read it twice, nodded along, and felt sharp as a tack. Genuinely informed. Pleasantly expert.

Two days later, the client asked me one follow-up question, one level deeper, and I discovered what I'd retained: a feeling. The shape of an understanding with nothing load-bearing inside it. I'd experienced what I now think of as the glow, the warm sensation of comprehension that reading a good explanation produces, which turns out to be a different thing entirely from comprehension.

Here's why that story belongs in a marketing library: a buyer can have that same experience about your category, your services, or your firm. They arrive at meetings glowing. They summarize your

industry back to you fluently. They feel informed in a way their questions don't support. And the research says this isn't a character flaw or a new laziness. It's fifty years of documented human wiring, recently connected to a machine that runs on it.

We Remember the Path, Not the Page

Back in 2011, Betsy Sparrow, Jenny Liu, and Daniel Wegner published "Google Effects on Memory." Their finding was plain and a little unsettling: when people expect they can retrieve information later, they store less of the information itself, yet they hold onto where to locate it. Search starts behaving like an external memory, less a reference than a place we deposit what we might otherwise carry. So we learn the route instead of the destination. Efficient, yes, but it comes with a quiet exchange. What sits in the machine doesn't automatically settle into you.

Offloading Is Human, and Now Supercharged

None of this began with the internet. Evan Risko and Sam Gilbert, writing in their review "Cognitive Offloading," describe how ordinary it has always been to recruit the environment to ease mental strain: fingers for counting, paper for notes, a tilted page for legibility, an alarm for recall. Most of the time it helps you do the immediate task better. Still, the bill arrives elsewhere. Dependence can creep in, and the abilities you no longer rehearse can dull.

For centuries, offloading mostly swallowed memory and arithmetic. The shift now is that AI pushes the same idea into reasoning and explanation, parts of thinking that used to demand more effort from the person using the tool.

We Already Think We Understand More Than We Do

Even without any device in the loop, people routinely overestimate what they understand. Leonid Rozenblit and Frank Keil gave this tendency a name: the illusion of explanatory depth. Ask someone how well they think they understand a zipper, a toilet, a bicycle, and confidence runs high. Then ask for a step-by-step explanation, and that confidence tends to crumble. What they had was a feeling of knowing, not the structure of knowledge.

And the illusion bites hardest right where it matters, in the "how" and "why" that sit under real explanations.

Easy access makes it worse. Matthew Fisher, Mariel Goddu, and Frank Keil show, in "Searching for Explanations," that internet searching can inflate people's sense of their own knowledge, even about questions they never looked up. After a search session, people feel smarter than they are. They slide from "I can find it" to "I understand it." The tool that holds the knowledge starts to feel like the mind that holds it.

That slide is close to what I mean by the glow. The experiments concerned internet search, not my AI-assisted client conversation, but they give me a reason to examine the resemblance.

The Medium Shapes the Mind

Two writers take the story beyond memory and into the contours of attention itself. In The Shallows, Nicholas Carr argues that the internet conditions us for skimming and jumping, while the slower, more linear attention that deep thought depends on gets squeezed out. Maryanne Wolf, in Reader, Come Home, makes a related case: deep reading, the kind that supports analysis, reflection, and insight, is learned and fragile, and digital routines can erode it.

Their shared warning is simple enough to linger: understanding grows in a mode of attention that offloading and constant scanning can gradually crowd out. If deep work is where understanding is built, then a tool that removes the deep work removes more than inconvenience.

Why This Belongs in Decision Architecture

Decision Architecture reduces the point to a single line, access to information isn't understanding, and understanding isn't judgment. What follows here is the support for that sentence. The same easy access that makes us feel informed can leave us no better at deciding what any of it means.

It also matters that the struggle to learn was never separate from expertise. Often, the struggle produced the expertise. Offloading research helps explain why. Take away the friction and you can keep the output, the explanation, the summary, the confident-sounding answer, while losing the comprehension that effort used to generate.

This is where the thread loops back to trust in automation. Automation bias is overtrust in the machine's judgment. The illusion of understanding is overtrust in your own. They mirror each other, and AI can provoke both in the same moment: it hands you fluent prose that feels authoritative, from a system you're tempted to depend on.

That is why the white paper can plausibly argue that everyone can sound like an expert while fewer people become one. Earlier tools pushed facts outward but still left you to do the thinking. AI offers to do the thinking as well, then returns it in crisp, confident language that can make understanding

feel easier than it is. You read the clean explanation, feel the glow of comprehension, and move on, having manufactured the sensation of knowing without paying the cost that makes it true.

So the advantage shifts, quietly but decisively. When producing explanations is basically free, the scarce skill becomes judging whether an explanation is right. Yet that judgment demands the very understanding the machine makes it tempting to skip. AI rewards the person who brings judgment to it, and it punishes the person who tries to swap it in for judgment. The day everyone began to sound like an expert is also the day genuine expertise became harder to spot, and more worth having.

The Marketing Read

Some prospects will arrive briefed by a machine and fluent in your category's vocabulary. Find out how they are thinking about the actual decision before assuming that fluency is either deep understanding or empty confidence.

Correcting them fails, because nobody thanks you for puncturing their expertise, and the glow feels identical to knowledge from the inside. Playing along fails too, because a buyer who believes they understand has no need for what you sell. My application of Rozenblit and Keil's explanation task is practical: ask for the step-by-step. "Walk me through how you're thinking about implementation." "What happens in month two?" Respectful process questions let the buyer discover the edges of their own understanding privately, without you ever saying they're wrong, and the gap they find is the exact shape of your value. The firms that thrive with glowing buyers aren't the ones with better rebuttals. They're the ones with better questions.

Build your content for the "how," because that's where the illusion lives and where machine-generated fluency runs thin. Summaries of what and why are free now; everyone's got them. The piece that walks through how, in the order a practitioner would do it, with the failure points marked, demonstrates the depth the glow can't fake, and gives the buyer who just discovered their gap somewhere to land.

Then run the same test inside your own shop, because your team is offloading too. The strategy deck nobody can defend one level deeper, the thought leadership the named author hasn't fully understood, the analysis that's really the machine's analysis with the firm's logo, these ship the feeling of knowing, and your buyers can now generate that themselves, free, in seconds. My rule: never publish an argument the named author can't extend one question past the text. Use the machines for production all day long. Keep the thinking reps in-house, because the reps were never busywork. They were where the expertise came from.

What backfires: weaponizing the gap. There's a version of the process question that's really a pop quiz designed to humiliate, and buyers know the difference instantly. The gap is your opening only if the buyer walks through it with their dignity attached.

The hard call: pull up the last piece your firm published and ask the named author one follow-up question, one level deeper than the text. If the answer isn't there, you're not marketing expertise. You're marketing the glow, into a market where the glow is free.

Primary Sources

- Betsy Sparrow, Jenny Liu & Daniel M. Wegner, "Google Effects on Memory: Cognitive Consequences of Having Information at Our Fingertips." Science, vol. 333, 2011, pp. 776–778.
- Evan F. Risko & Sam J. Gilbert, "Cognitive Offloading." Trends in Cognitive Sciences, vol. 20, 2016, pp. 676–688.
- Matthew Fisher, Mariel K. Goddu & Frank C. Keil, "Searching for Explanations: How the Internet Inflates Estimates of Internal Knowledge." Journal of Experimental Psychology: General, vol. 144, 2015, pp. 674–687.
- Leonid Rozenblit & Frank Keil, "The Misunderstood Limits of Folk Science: An Illusion of Explanatory Depth." Cognitive Science, vol. 26, 2002, pp. 521–562.
- Nicholas Carr, The Shallows: What the Internet Is Doing to Our Brains. W. W. Norton, 2010.
- Maryanne Wolf, Reader, Come Home: The Reading Brain in a Digital World. Harper, 2018.