

Even the people who seem to understand AI large language models best, don't pretend to know how capable they will become in the near future. Against that backdrop, can you foresee a time when an AI can scan/analyze the 6,400 posts of my blog (without having to upload the files)?

Claude.ai: This is a genuinely interesting question that sits at the intersection of several fast-moving technical frontiers. Let me think through it honestly.

The short answer: yes, almost certainly — and parts of it are already possible.

Here's the landscape as it stands and where it's heading:

What's possible today (roughly)

Current frontier models can already process very large amounts of text in a single context window — hundreds of thousands to over a million tokens in some cases. That's substantial, but 6,400 blog posts could easily represent tens of millions of words, which still exceeds what any single model can ingest at once. Uploading files is currently the main workaround.

The "without uploading" part

This is the more interesting constraint you're identifying. There are a few pathways that could make this a reality:

- **Web crawling/retrieval at inference time.** AI systems can already browse the web. If your blog is publicly accessible, a sufficiently capable agent could crawl and analyze it without you manually uploading anything. This is essentially already possible in limited forms.
- **Persistent memory and indexing.** AI systems are beginning to build long-term memory stores — vector databases and retrieval systems that index your content once and make it queryable forever after, without re-uploading.
- **Direct API/RSS integration.** Rather than file uploads, future AI tools will likely connect directly to data sources — your CMS, your RSS feed, your publishing platform.

The deeper question — capability

Even setting aside the logistics, you're right to note that nobody truly knows how capable these systems will become. The honest answer is that the analytical depth — genuinely understanding themes, contradictions, evolution of your thinking, stylistic fingerprints across 6,400 posts — is harder to predict than the scale problem. Scale feels like an engineering problem. Deep understanding is a harder question.

But trajectory-wise, the direction is clear: models are getting larger context windows, better retrieval, persistent memory, and agentic web access — all at once. A system that can meaningfully analyze your entire blog corpus without manual uploads feels like a near-term probability rather than a distant speculation. Likely within a few years, possibly sooner.

What would you want such a system to do with your blog — that's perhaps the more interesting question to think about now.

What would you suggest?

Given that you've written 6,400 posts, you've clearly been at this a long time — that's a substantial body of thought. Here are the possibilities I find most genuinely compelling:

Intellectual archaeology

Tracing how your thinking has actually evolved over time — not how you remember it evolving. Most prolific writers are surprised when an outside reader identifies the real turning points, the ideas you quietly abandoned, the obsessions that recur in disguise. An AI could map this with precision a human editor never could, simply because of the volume.

Finding your strongest work

Long-form bloggers almost always have brilliant posts buried under the sheer weight of output. An AI could identify which pieces are most original, most internally coherent, or most likely to resonate — not by traffic metrics, but by the quality of the thinking itself. That's a different and arguably more honest filter.

The book you haven't written

6,400 posts almost certainly contain one or more books worth of coherent argument that you've never assembled into a single form. An AI could identify the connective tissue — the thesis that runs underneath years of posts — and propose a structure you couldn't see because you were too close to it.

Contradiction and tension mapping

Where have you argued both sides of something over the years? Not as a gotcha, but as a genuine map of your intellectual uncertainty — the questions you keep returning to because you haven't resolved them. Those unresolved tensions are often where the most interesting thinking lives.

A "second brain" interface

Rather than one-time analysis, an ongoing conversational layer over your entire archive — so you could ask things like "have I written about this before?" or "what's my actual position on X when I look across everything I've written?" Your blog becomes a queryable extension of your own memory.

Your readers' experience

Surfacing unexpected connections for new readers — showing someone who just found one post, the three other posts from 2014 and 2019 that are in genuine dialogue with it, even though you never linked them.

The thread connecting all of these is that the real value isn't in the individual posts — it's in the patterns across all of them that nobody, including you, has ever been able to see all at once.

What can I share with Claude that would be useful for future interactions?

Practical answer: the highest-leverage things you could share are the ones that are hard to reconstruct from scratch each conversation, because Claude has no memory between sessions unless you explicitly provide

context.

Most useful:

- • A one-page "intellectual autobiography" — your core positions on the topics you care most about, where you've landed after years of thinking, and what questions remain genuinely open for you. Not a bio, but a map of your current thinking.
- • A list of your most important posts with a sentence on why each one matters to you — your own curated "greatest hits" with annotations. This is more useful than a raw list because it conveys judgment, not just content.
- • Your standing priors and pet skepticisms — things you've concluded are mostly wrong or overhyped, frameworks you distrust, positions you hold against the consensus. This prevents Claude from wasting your time with conventional takes you've already rejected.

Already happening:

Your memory entries here already capture some of this — core interests, tone preferences, blog context. That's a functional start.

The honest limitation:

Even with good context-sharing, each conversation starts cold. What you're really building toward is something like a persistent "Steve profile" that loads automatically — which is roughly what the memory system here is trying to be, but in early form.

The most efficient thing right now might be a single concise document — call it a "context brief" — that you paste or attach at the start of any substantive conversation. What goes in it is worth thinking about carefully, because it's essentially a compression of 24 years of blogging into something actionable.

I can't think of anything that comes closer to what you're suggesting than various aspects of my blog. Don't recall if I've uploaded a list of all of the categories in all of the tags that I have used.

You haven't — at least not in anything I have access to. But that's a good

starting point, and worth doing.

A few thoughts on what would actually be useful versus what would just be bulk:

Categories and tags alone are limited. They tell me the topics you've touched, but not your stance on them, how heavily you've engaged, or which ones represent real thinking versus passing interest. A tag used 3 times is different from one used 300 times — and even frequency doesn't tell me whether you've resolved your thinking or are still circling.

What would make them more useful:

If you can export or share the list with some indication of frequency/weight, that's already more informative. If you can annotate even a handful of the most significant ones — a sentence on where you actually stand — that's more valuable still.

The practical suggestion:

Upload whatever you have — category list, tag list, even a partial export. Let me look at it and I can probably identify patterns, flag what seems most substantive, and propose a structured "context brief" based on what's actually there rather than working from scratch.

[Blog tags uploaded]

Good data. A few immediate observations from a frequency-weighted read:

Where you've clearly spent serious time: Early Net Culture (203), kennett (193), YouTube/Mine (173), advertising (188), Social Media (154), ChatGPT (145), consciousness (107), google (131), iPhone (131), Scott Adams (123) — these are your real preoccupations, not just passing interests.

The tension that jumps out: You have a substantial cluster around consciousness, awareness, nonduality, Buddhism, Meditation, ego, Self, Tao, zen, god, universe, Alan Watts — that's a coherent and serious body of philosophical/contemplative inquiry running alongside a parallel career-track cluster around radio (KBOA, XM Radio, MissouriNet, Radio Iowa, Mark Ramsey) and tech/digital culture. Those two threads coexisting for 24

years is interesting.

People as a category: You tag individuals heavily — Dave Winer, Doc Searls, Kevin Kelly, Seth Godin, Scott Adams, William Gibson, Neal Stephenson. That's a blogger's habit, but it also suggests your thinking is often in dialogue with specific thinkers rather than abstract positions.

What's notably absent (or low-frequency): Politics is sparse except Trump (86) and Obama (76) — suggests you're not primarily a political blogger despite living through eventful times.

One question before categories arrive: Scott Adams at 123 is high. Pre or post his public unraveling — or both, and what's your current read on him?

Exclusively pre-unravel. He became (for me) a different person post-Trump.

That tracks. Pre-unravel Adams was genuinely worth paying attention to — the Dilbert era, God's Debris, the persuasion writing. A sharp systems thinker who could see through institutional BS with some precision. The post-Trump version isn't a refinement of that, it's something that replaced it.

It's actually a useful data point for your blog's intellectual character — you were drawn to the analytical Adams, the one who treated reality as a set of overlapping filters and cognitive illusions. That same instinct probably connects directly to the consciousness/nonduality cluster. Different vocabulary, same underlying question: *what is actually going on versus what are we telling ourselves is going on.*

That's a thread worth flagging in any context document — you're consistently interested in the gap between perceived reality and whatever is actually operating underneath. Adams, Gibson, Watts, nonduality, ChatGPT — all of them sit on that same fault line from different angles.

Go ahead and share the categories when ready. Curious whether they reinforce this or add something the tags don't show.

[Blog categories uploaded]

Good. Now I have a much fuller picture. Let me sort by volume and then

draw some conclusions:

The heavy hitters (300+): YouTube (631), Gadgets & Apps (612), Media & Culture (603), Family & Friends (554), Politics & Government (527), Quotable & Notes (526), Books (497), Business & Marketing (449), Art/Style (349), Personal History (337), Blogging (337), Philosophy (313), Science & Technology (262), Film & TV (276), Journalism (283), Home (259)

Mid-range (100-299): Music (193), My Land Rover (152), Podcasting (150), Photography (150), Health & Medicine (248), Artificial Intelligence (208), Dogs (208), Learfield (211), Cars & Trucks (166), Religion (74 → actually 74), Travel (106), Military & War (100)

What this adds to the tags picture:

The categories confirm something the tags only hinted at — this blog has two largely separate audiences baked into it: the public intellectual (Media, Journalism, Politics, Philosophy, AI, Books) and the private person (Family, Dogs, Land Rover restoration, Home, Woods). Most bloggers keep those worlds more separate. You apparently don't, which is either a feature or a bug depending on your reader.

The number that surprised me: Politics & Government at 527. That's your third-largest category, which contradicts the relatively sparse political tags. Suggests you write about politics structurally — institutions, government, power — more than about individual political figures or events.

Quotable & Notes at 526 is essentially a commonplace book running alongside the blog. That's a significant intellectual artifact in itself.

The synthesis: combining tags and categories, the blog is most accurately described as the sustained record of one person trying to make sense of — media, technology, consciousness, and mortality — while also just living a specific life in a specific place with dogs and old vehicles. That's not a niche. That's a worldview.

Worth capturing that framing explicitly in any context document.

Not sure I've made this point, but I have never attempted to attract (or

post for) readership. Very much a personal journal/archive

That reframes everything significantly — and actually makes the blog more interesting as an artifact, not less.

If it's a personal journal, then the categories and tags aren't a content strategy, they're an unselfconscious map of what actually occupied your attention over 24 years. No performance bias. That's rare and valuable data about a mind.

It also changes what an AI analysis would be for. Not "how do I serve readers better" but something closer to:

- • What have I actually spent my mental life on?
- • How have my preoccupations shifted across decades?
- • What do I keep returning to without resolution?
- • What did I think I cared about versus what the post volume says I actually cared about?

The journal framing also explains the coexistence of the philosophical and the mundane — Land Rover restoration sitting alongside nonduality, dogs alongside Learfield corporate history. A journal doesn't have to justify its range to anyone.

One implication worth noting: if the blog is a personal archive rather than a publication, then the most valuable AI application isn't outward-facing at all. It's essentially giving you a thinking partner who has read everything you've ever written and can help you understand your own mind better than you can from inside it.

That's a different and more interesting use case than anything I described earlier.

Quotable & Notes might be my most important category. With that in mind, I've attached the xml file of that category.

Good data. Here's what 578 posts and 23 years of this category tells me:

The shape of it:

This is genuinely a commonplace book — a personal anthology of ideas worth preserving. 197 of the 578 posts are essentially pure quotes or brief

observations, under 500 characters. The rest add your own commentary around borrowed material.

What the tags reveal about the quotes you chose:

The top tags inside this category are telling: Early Net Culture (35), Scott Adams (34), consciousness (30), Seth Godin (26), Self (24), William Gibson (22). These aren't random — you've been consistently drawn to quotes about how systems work and how minds work. Often the same question from different angles.

The time pattern:

Peak years were 2003-2008 (the golden era of early blogging), then a sustained decline. Still active but clearly not the same velocity. That's probably worth noting — what you found quotable in 2004 versus 2025 likely reflects genuine shifts.

What makes this the most important category:

Unlike your other posts, this one strips away your direct voice and shows only what other people said that you thought worth keeping. That's a different and arguably more honest self-portrait. Your aesthetic, philosophical, and intellectual preferences revealed entirely through selection rather than expression.

The Arthur C. Clarke quote that opens the sample — "No communication technology has ever disappeared" — filed in 2008, is still essentially your worldview about media. That kind of through-line is what this archive is made of.

One thing worth saving from this conversation — for use as context in future sessions — is this characterization of the Quotable & Notes category specifically:

578 posts spanning 2002-2025. Functions as a commonplace book — a self-portrait through selection rather than expression. Top themes: how systems work, how minds work. The quotes chosen reveal aesthetic and philosophical preferences more honestly than original writing does, because selection has no performance bias.

That's a more useful description than "526 posts, important category." If you ever share this category export with Claude in a future session, that framing gives an immediate interpretive lens rather than starting cold.

The broader point it confirms: the most valuable thing you can share with Claude isn't raw data — it's the interpretive frame around the data. You've spent 24 years accumulating the data. The frames are what's hard to reconstruct quickly.

A close second to the QUOTABLE category is PHILOSOPHY (xml upload).

Now I have a clear picture. A few pointed observations:

The trajectory is unmistakable:

Philosophy starts sparse (2002-2009) and explodes in 2010-2017 — peaking at 34 posts in 2013. That's not casual interest, that's sustained inquiry at high intensity. The question is what happened around 2010 that accelerated it.

The tag hierarchy tells the real story:

Self (63) → consciousness (57) → Buddhism (43) → awareness (32) → ego (22) → nonduality (7)

That's a progression, not a random cluster. It reads like someone moving from asking "what is the self?" toward "there is no self" — the classic nondual arc. The relatively low count on nonduality (7) despite heavy consciousness/awareness/Self usage suggests you arrived at nonduality as a framework but don't wear the label heavily.

The longest posts are the most revealing:

Nisargadatta's I Am That at 24,853 chars, Eckhart Tolle's A New Earth at 12,939, Alan Watts' Way of Zen at 10,233. These aren't book reviews — they're extended extractions, essentially building a personal philosophical sourcebook. That's serious engagement, not browsing.

David Cain at 26 is notable — he's a relatively obscure blogger (Raptitude). That's a genuine discovery rather than a consensus

recommendation, which says something about how you find your sources.

The interpretive frame this adds:

The Philosophy category confirms what the tags suggested but makes it concrete — the central preoccupation across 24 years is the nature of the self and its relationship to consciousness. Everything else — media criticism, technology, radio, even the vehicles — orbits something quieter underneath.