

{.pulseprint}

# What does useful AI in a blog even look like?

# Notes on Research

Welcome to this {pulseprint}!  
It's great to have you joining us here today.

Today, we'd like to share with you our thoughts on what integrating AI into a blog well really looks like. AI has been used for decades, with LLMs constantly being trained to recognise patterns and act on user requests. However, with the recent launch of Large Language Models with GPT capabilities, like chatGPT and Claude, we feel compelled to examine the impacts these amazing technologies can have on a blog like ours.

Perhaps we'll integrate AI into dogAdvisor really soon. Perhaps it will come in the form of a search field you can ask anything to - or perhaps a chatbot that lives on the side of your content or within it.

Since latest year, we've been developing a custom-built LLM that pulls data from our own articles, makes inferences and conclusions, uses this data in its responses, and understands semantic search with a groundbreaking new semantic index that understands user intents (what we call Active Intents).

This research paper will explore the best ways one could integrate AI into dogAdvisor. FYI, this report was **written by AI** so we could get a sense of how this technology thinks, itself, will be the best way to integrate our custom-built LLM into dogAdvisor.

I hope you enjoy this read!  
All the best,



**Deni D**  
Founder & CEO, dogAdvisor

# What does useful AI in a blog even look like?

1. Useful AI should answer questions directly, rather than act as a traffic manager.

The AI began by distinguishing between two types of AI commonly seen on blogs: one that actually answers the user's question, and one that simply links to other pages and tells them to "go read more." Most current AI integrations fall into the second camp. They mimic the logic of a sitemap or a search bar, acting more like a marketing assistant than an informational partner. In contrast, the AI suggested that truly useful blog-integrated AI should offer direct answers where appropriate - and only recommend further reading when the question is genuinely too complex or personal to answer fully in one reply. That distinction, while seemingly simple, is significant: it separates AI that is optimised for *the reader* from AI that is optimised for *the publisher*.

2. Useful AI must be trained not on the internet - but on the blog's own content, tone, and thinking.

One of the most crucial insights was that most AI implementations fail because they treat the blog like any other corner of the internet. They rely on web-trained language models that might return responses based on what Google says - but not what *this specific blog* believes, prioritises, or stands for. The AI noted that a genuinely useful implementation would involve a model trained exclusively on the blog's own articles. This doesn't mean a surface-level scan of text - it means semantic understanding of argument flow, structure, internal logic, and the nuances of how information is prioritised.

The AI was surprisingly insistent on this point. It noted that without this internal model, any AI assistant would feel like a guest - someone who can talk about the content, but never really *represent* it. Whereas with tight fine-tuning and prompt calibration, the AI becomes a true extension of the blog's voice. It reflects its standards, tone, knowledge depth, and even ethics. On a site like dogAdvisor, this would mean that the AI understands that we speak plainly, never exaggerate danger, avoid jargon, and prioritise calm reassurance in emergencies.

3. Useful AI doesn't interrupt the reading experience - it enhances it subtly and only when needed.

The AI described a vision where AI is not always the main interface. Sometimes, it's just a helpful layer. It might appear in a corner of the screen, or hover beside a paragraph, ready to explain, simplify, or personalise what the reader is seeing. But it doesn't force itself into the experience. It doesn't pop up constantly asking "Need help?" or block the screen with flashy animations. It exists in the margins, not the centre. In the same way a really good footnote in a book doesn't break your flow but gently adds to your understanding, useful AI should behave like a guide - not a narrator.

4. Useful AI knows what not to say.

This was one of the most unexpected - and frankly, insightful - conclusions. The AI noted that truly intelligent assistants know their limits. They don't guess dangerously. They don't speculate on medical advice. They're transparent when they don't have enough data to answer. But more than that, they improve over time. They learn what users ask most often and where misunderstandings occur. They don't repeat the same useless suggestions forever - they adapt. In this sense, a good AI is not just a static tool but an evolving one, constantly shaped by real engagement.

5. Useful AI becomes part of how users trust the blog - not just how they use it.

Finally, the AI drew a subtle but important line: functionality builds trust. It's not just what the AI says but how it *behaves* that teaches users whether the blog deserves to be trusted. An AI that is consistent, modest, and informative over time actually becomes a signal of credibility. And when paired with strong human-written content and transparent authorship, that AI can amplify trust rather than dilute it. The more it reflects the blog's values - accuracy, usefulness, accessibility - the more it strengthens the reader's relationship with the whole platform.

In conclusion, then, the AI's own answer to the question "what does useful AI in a blog look like?" is not just a checklist of features. It is a philosophy of assistance - one grounded in humility, helpfulness, and a deep respect for the reader's time. If we at dogAdvisor choose to integrate our custom AI soon, it will be because we believe it meets those standards. Not because AI is trendy, but because it is finally capable of doing something we care about: making dog owners feel better supported, more informed, and less alone in the everyday (and sometimes scary) decisions they have to make.

So - if we accept the idea that AI should serve the user, speak in the blog's voice, and quietly enhance the reading experience without stealing the spotlight, then what *form* should that AI actually take?

Should it be a chatbot? A search tool? An assistant that floats next to each paragraph? Or something entirely different?

From the AI's own reflections and the design principles we've developed internally, the clearest answer is that **form should follow intent**. The way the AI appears should be shaped by what the user actually needs at that moment - not by what looks trendy or technically impressive.

In some cases, that means a **semantic search bar** may be the most helpful interface. Imagine typing in, "Is it okay if my dog skips one meal?" and being shown not ten articles with the phrase "skip" in them, but one clean, plain-English answer that directly addresses the concern - drawn from the best of dogAdvisor's articles, backed by clear reasoning. This approach doesn't require a chat interface, and doesn't overwhelm the reader. It's just a better search experience, with smarter understanding behind the scenes.

In other cases, especially on pages where the user is already deep in a specific issue - for example, reading about raw feeding or separation anxiety - a **context-aware chatbot** might be more valuable. Not one that sits in the corner and offers vague advice, but one that sees which article you're reading and can answer specific, nuanced follow-ups. For example: "What if I'm gone for 9 hours, not 8?" or "Can I feed raw but still use kibble sometimes?" These kinds of questions deserve

more than hyperlinks - they deserve an actual exchange. And in this setting, a conversational format makes sense. Not to replace the article, but to sit beside it and make it more alive.

There may also be future forms we haven't even tested yet - like **inline expansions**, where you tap on a sentence and instantly get a short AI explanation or context box right in the margin. These are subtle, low-effort ways to deepen understanding without needing to leave the page or start a full conversation. We believe this could be especially powerful for new dog owners, who often hesitate to ask what might feel like "silly" questions. With inline AI, they wouldn't need to.

Ultimately, what emerged from this AI-written report is that **no single format is perfect on its own**. Search, chat, expansion, alerts - each serves a different kind of reader need. So the smartest way forward may not be choosing one format and running with it, but rather creating a **flexible AI layer** that shows up in different ways depending on the moment. That's what we're exploring now at dogAdvisor. A system that doesn't demand attention - it earns it. One that knows when to speak, when to listen, and when to quietly stay out of the way.

If we do this right, users won't say, "Wow, that's impressive AI." They'll say, "That helped." And that's the only metric that matters.

{.pulseprint}

This is the end of the publication