

In Praise of Suffering?

Posted – December 22, 2025

A month ago, on The Conversation, Shiri Melumad, a professor in the marketing department at the University of Pennsylvania, published an article reviewing a study she and a colleague conducted that compared the “depth of learning” from information obtained through traditional Google searches with information obtained through the use of AI tools like ChatGPT. The study’s conclusion appears in the article’s title:

[Learning with AI falls short compared to old-fashioned web search](#)

I didn’t know what to expect from this study, although even from reading the title alone it wasn’t hard to guess. I assumed that since search is a relatively active activity, the study would find that learning that occurs through Google searches for information is deeper than through LLM tools. For me, at least, it’s very easy to adopt Melumad’s assertion that using LLM, which is often effortless, creates a poor and superficial sense of “learning.” Indeed, Melumad writes:

However, a new paper I co-authored offers experimental evidence that this ease may come at a cost: When people rely on large language models to summarize information on a topic for them, they tend to [develop shallower knowledge](#) about it compared to learning through a standard Google search.

The study, which, as we recall, finds value in searching for information through Google, was published in October, and it is interesting that I came across it near [the previous article here](#) in which I quoted Justin Reich, who claimed that too many students do not know how to conduct an informed information search through Google. Regardless of Reich’s criticism, I prefer “traditional” search to the use of AI, but it is precisely because of this preference that I felt the need to read [the full study](#). I was afraid that, like [the study by Nataliya Kosmyna](#) at MIT, which [I also wrote about](#), I would discover that, at least in my opinion, the study’s conclusions are not exactly supported by the study itself. Happily, I was wrong.

I found a study that made efforts to examine whether there really is a difference in learning, or the quality of learning, between searching the web in the “traditional” manner and using AI to obtain information. It was built from four different parts. In the first part, in which over a thousand students participated, the students were divided into two groups who were asked to prepare instructions for a colleague on the topic of growing a vegetable garden. To prepare the instructions, one group researched the topic using “regular” Google searches, while the other group used ChatGPT. The results were clear - the participants who conducted searches via Google before preparing their instructions felt that they had learned more about the topic, and an examination of what the participants wrote showed richer and more diverse content. In the second part (with about 2000 participants) the participants were again divided into a Google search group and a ChatGPT group, but there was a certain trick here - in essence, the researchers presented the same results to both groups, with only the means of accessing the information differing. Participants in the ChatGPT group reported that they learned less than the Google group.

The third part sought to isolate the differences even further. 250 participants were divided into one group that searched via Google, and a second group that used Google’s AI Overview. (In this experiment, the task

was not to explain how to grow a vegetable garden, but how to live a healthier life.) Again, the results were similar. The involvement in preparing the instructions for the search group was greater, and the results were also richer. In the fourth part, the results of the 250 participants in the third part were presented to 1,500 readers who were asked to rate the quality of the information that appeared in the various summaries. Again, the summaries of those who used the “regular” Google search were considered of higher quality than those of the AI Overview. (It is worth emphasizing here that I have tried to summarize a fairly complex study faithfully, and I recommend reading the full study rather than relying on my summary.)

From all this, it seems to me that there is coverage for Melumad's statement when she writes:

learning from synthesized LLM responses led to shallower knowledge compared to gathering, interpreting and synthesizing information for oneself via standard web links.

Precisely because I initially tend to agree with the results of Melumad's research, these results caused me a certain discomfort. We are always happy to find references that support our opinions, and a little caution regarding the conclusions we wish to draw cannot hurt. But my discomfort is also related to an age-old educational question – whether one should suffer (or at least make an effort) in order to learn something thoroughly.

Melumad notes that one of the most basic principles of skill development is active involvement in what one wishes to learn. She writes that when learning about a topic is done through a Google search:

we face much more “friction”: We must navigate different web links, read informational sources, and interpret and synthesize them ourselves

However, when we turn to an LLM, the work is done for us and the learning process is passive. The learner is not personally involved in the process, and the result is superficial learning. Melumad uses the term disfluency (or "lack of fluency") and explains how it relates to the research she conducted:

learning through web search involves a particular type of disfluency: the need to engage in trial-and-error navigation among result links, discovering which are suitable for the question at hand, and then interpreting and synthesizing the different pieces of information for oneself.

It was interesting, by the way, to see that in the footnotes, Melumad links to an article from 2013 that I came across, I think, in 2015 and really liked - [The Benefits of Cognitive Disfluency](#) by Adam Alter. This is a very nice article, although not particularly deep. Alter's main argument is that in order to learn something, it's desirable to encounter a certain degree of difficulty. And here is my discomfort. I do not advocate that for the sake of learning, teachers and lecturers should force their students to suffer. Still, it makes sense that when we encounter a certain difficulty, or have to invest effort, the chance of internalizing what we are learning increases.

But this situation confronts me with two dilemmas – is it even right to mobilize technological tools to make things easier for learners, and secondly, if it **is** actually desirable to make things easier, why require tools from a decade ago but deny easier tools, such as LLM tools, of today.

It's clear that the pocket calculator has made things easier for students, and so have other tools – from spell checking their transcripts to quick access to information via the Internet. Perhaps the use of these was a

mistake, and as many today claim, in the end, these easements harm the learners' learning ability? As far as I am concerned, the answer is no. After all, there are also clear educational advantages to these "easements". These "easements" allow for a focus on the workflow (when there is no need to be careful with spelling, for example), and they invite expansion and deepening in many areas. It's doubtful that we would want to give up these and instead require learners to sweat more, even if we understand that there are cases in which sweating can lead to better learning.

But if I don't deny the ICT tools of a decade (or more) ago, why am I skeptical about new tools that also, at least according to their developers, contribute a lot to learning? Maybe I was wrong about the old tools too? It's worth noting that recently there has been quite harsh criticism of almost everything "technological" in education, even going so far as to demand that computers be removed from classrooms altogether. (Of course, the pencil is also a technology, and it's possible that there was once opposition to its integration in the classroom, but I'm not aware of any criticism today against the use of a pencil.) Because of this opposition, it is important to protect technology. But when we do this, it is important to distinguish between what promotes learning and what hinders it. In addition to opponents of ICT, there are technology companies that do not distinguish between what contributes to learning and what harms it, and they seek to convince us that progress cannot be stopped. Therefore, it's important to clarify that the march of technology is not a sequence of continuous improvement. The technologies we adopt reflect the perceptions and values of the society into which they grow, as well as the perceptions and values of those who develop them. Education may, and indeed should, embrace what can promote learning (and of course there is no unanimity among educators about what "learning" really is). Melumad's research, which focuses on the positive value of making an effort, illuminates an aspect of "old" technology that is precisely absent in new technologies that make it easier, and shows us that there is no shame in not always keeping up with "progress."