Surviving the Flood
by Matthew Neary

It's quite likely that machine-generated texts will soon outnumber human-written ones. With social media we got machine-generated feeds, but the content itself is still man-made. Soon most of the content might be machine-generated, too.

Let's take the example of news. First social media newsfeeds automated the editor, and now AI has the potential to automate journalists, too. When we replaced editors with algorithmic feeds, we gained certain advantages. Compared to newspapers, social media is more dynamic and more personalized. If automation is extended further and the content shared online is increasingly generated by AI reporters, these advantages will increase as well.

However, it's important to remember that in addition to these benefits, social media also introduced new kinds of problems—such as fake news, toxic discourse, and political polarization. At base, there's one big problem with social media: people live increasingly in information bubbles that reinforce their biases. This makes it very hard to reach broad consensus. In recent years, we've seen the effects of this: historical events have been interpreted in opposite ways by different groups, without ever really reaching a resolution.

Considered at a higher level, these trends are not limited to just social media or just text. The media environment of the Internet is, in general, highly individualized. This is often a good thing. However, at the limit, we all end up atomized—unplugged from community and disconnected from the broader world. With AI this trend toward atomization will accelerate. Our information silos will shrink further, asymptotically approaching one-man bubbles. In the future, we might encounter the world primarily as a fantasy land of hallucinated, machine-generated content.

For society to survive the coming flood of machine-generated content, there needs to be a countervailing force: technologies that can bolster collective sense-making and counter the gravitational pull of atomization. Otherwise, social reality will be so completely fractured that there won't be any society left. As Alan Kay once said, "The best way to predict the future is to invent it." Similarly, the best way to avoid a terrible future is to articulate a better alternative.

There's already an overwhelming amount of content online. The near-future of AI-generated content promises to greatly exacerbate this problem. Luckily, the same LLMs that power content generation can help sift through an increasingly noisy information ecosystem. However, this merely transforms a quantity problem into a quality problem. As we get more information second-hand, filtered through LLMs, we must be more careful to discern the truth and context of what's being conveyed. To survive and thrive in the era of AI-generated content, we need tools that leverage AI to help in all of these areas. It's essential that we adopt tools that help us not only to get through more content, but also to validate claims and contextualize ideas.

Solving the general case is a huge undertaking, but a new workspace for documents is a good place to start. I've started working on one. It lives, in its current form, at MagicPaper.ai—here you'll find the beginnings of an online workspace for distilling, summarizing, and correlating texts. The main insight is that the best way to validate texts and clarify context is by tracking the provenance of where ideas come from. Furthermore, when AI is leveraged to process information automatically, provenance can provide the missing link between the cursory view and full-length sources. Let me walk through some examples:

Question Answering: First, load a document and then ask a question. The margins will fill with answers, each attached to a specific part of the document that backs it up. You then have the option to open the answers as a page of their own. Even once they've been lifted into their own document, the answers are all linked to the corresponding parts of the source.

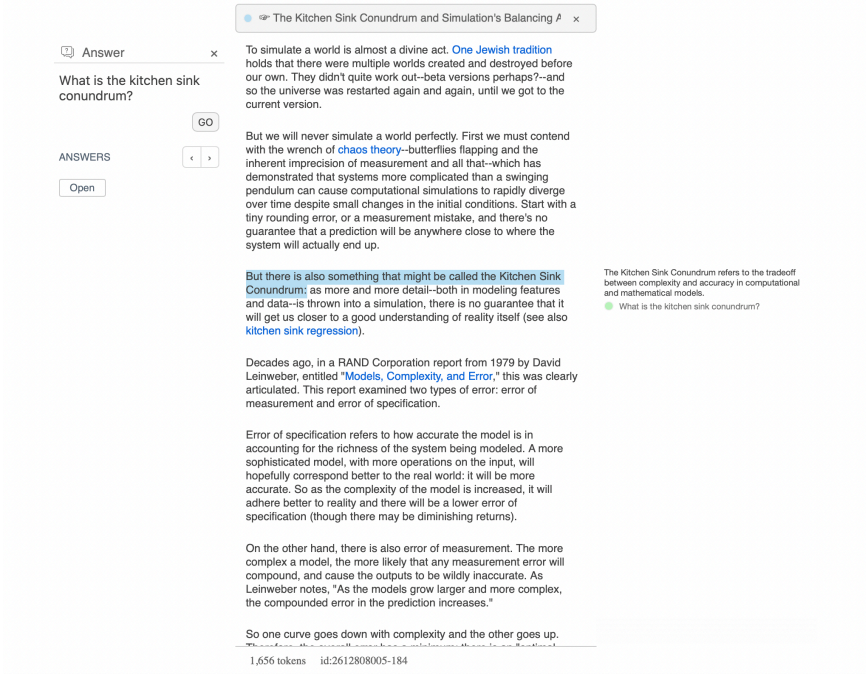

Extractive Summarization: Again, load a document into the tool. Click distill and two things will happen: fragments of the document will be automatically highlighted and you'll be presented with a histogram. The histogram visualizes the distribution of fragment salience. You can use this chart to adjust the filter threshold. As you change the threshold, the highlights in the document will change. You can open the highlights as a page and they will remain linked to the source.

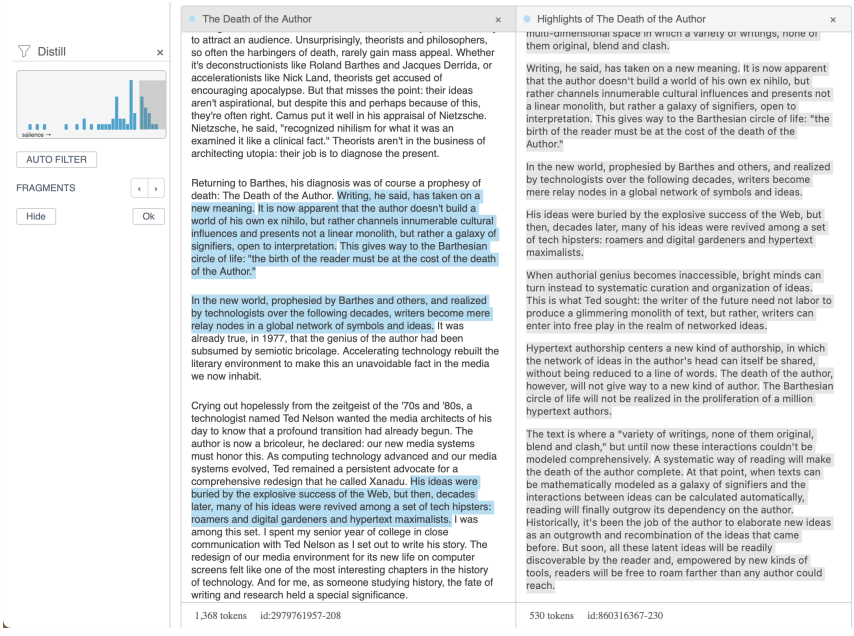

Abstractive Summarization: This is the most complicated workflow, but also the one that I intend to focus on the most going forward. First you distill the key fragments from a document, then you generate a summary based on those fragments. Finally, you use another tool to correlate the summary with the input fragments. Now, every point in the summary will be connected to parts of the source text that support it.

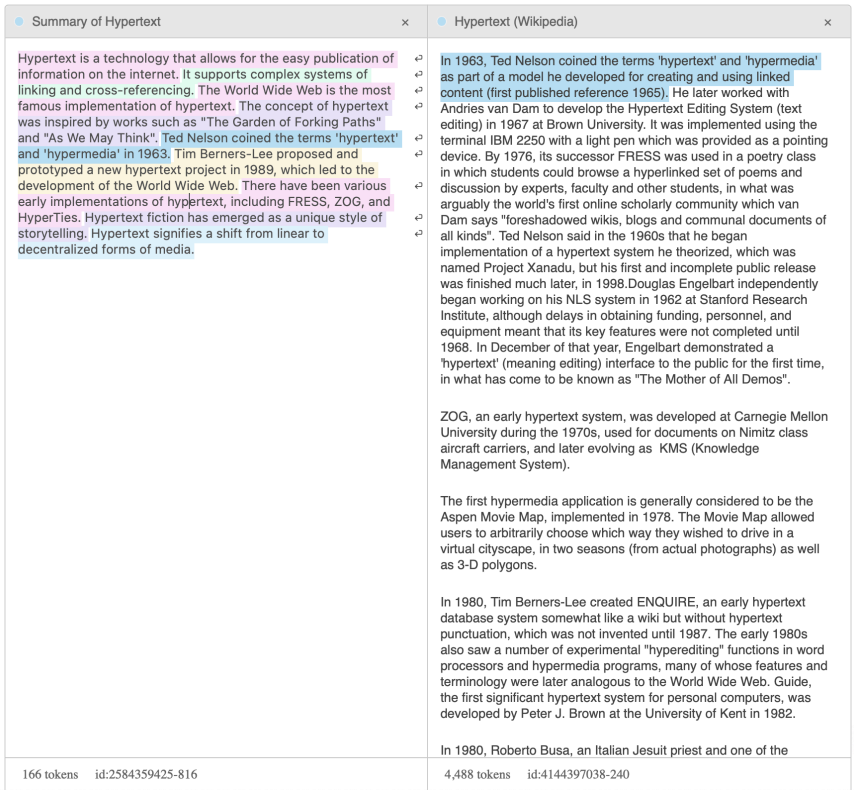

Interestingly, a document workspace like above isn't a new idea. It predates even the personal computer and the Web. Much of the design I've adopted in Magic Paper is inspired by ideas from decades ago. The provenance-tracking and parallel pages, specifically, are similar to designs Ted Nelson shared in *Literary Machines* (1981). These design primitives, I've found, are a perfectly countervailing force for the major challenges presented by AI adoption.

By adopting a tool like Magic Paper, we can upgrade our information bandwidth to get through more content. This will be increasingly necessary as the volume of online content grows. Similarly, we'll need to put greater care into tracing ideas back to their source and context, as more of it comes second-hand through LLMs.

The underlying designs here are nothing new. As we adopt AI in our knowledge work, we need to go back to the future of decades past. Advanced tools for reading and writing have been anticipated since the earliest days of personal computing, but they never really took off. Now, with LLMs, more powerful tools are both a possibility and a necessity.