

5 future side effects of AI we need to talk about

Key note from Jan Birkemose, Medietrends.dk

September 2nd 2024

Thank you for inviting me to speak today.

Since ChatGPT made its breakthrough in November 2022, it's become clear that artificial intelligence is set to change the world, especially within the media industry.

We have all heard about the obvious potentials and benefits, the clear challenges, and the threats it poses.

But what about the indirect side effects? We haven't really talked enough about those, likely because they're harder to spot since they tend to emerge more gradually

Over the next half hour, I'm going to try to outline some possible contours of these side effects.

As a naturally curious and optimistic person, I've been captivated by the exciting possibilities that generative AI offers since it became available to the public. Some might even label me as an AI evangelist—and while that's not entirely inaccurate, it's not the whole truth either.

The reality is that the more I've worked with AI and witnessed its rapid development, the more concerned I've become. And that's what I want to share with you over the next half hour.

I apologize in advance if I end up dampening the positive atmosphere.

Before we dive into the substance, I want to emphasize one crucial point: we are in the midst of this evolution, and absolutely no one can predict exactly what will happen next.

To illustrate, let me take you back to 1779, when the world's first iron bridge was constructed in Shropshire, England. Cast iron is an incredible material, vastly different from wood in almost every way.

But what did they do when they constructed the bridge with this new innovative material?

They treated the iron as if it were wood, using wedges and dowels just like they would in building a wooden bridge.

The takeaway here is that we often lack the imagination, intelligence, and creativity to fully exploit new things right away.

We tend to stick with what we know, solving familiar problems and meeting familiar needs while being thrilled that we've found a new tool or material that can shave off a few percent of our time or effort.

But eventually, after some time, some sharp minds come along who connect the dots in surprising ways.

This is the kind of innovation that transforms the way we live—change that is irreversible, and with a high impact on our lives and societies.

AI undoubtedly has the power to change our lives, but we're not quite there yet.

The technology is impressive, but it hasn't fundamentally changed us as human beings—at least not yet.

But I believe we're getting close, and the scary part is that we might not realize how our lives have been changed until long after it's already happened.

Despite that disclaimer I will share a few cautious scenarios of what might unfold.

First, a brief introduction about myself:

I'm a trained journalist and started my career as an investigative reporter at *Ekstra Bladet*. I then moved on to the now-defunct *Ugebrevet A4*, which for many years was Denmark's most agenda-setting and widely cited media outlet.

After 15 years there—eight of which I spent as editor-in-chief—I left in 2016 and have since been working full-time as an independent media analyst.

I run a company called Medietrends, which focuses on the future—not in the sense of future studies, but as an inspiring knowledge base grounded in real experiences and ideas that can be used to navigate specific challenges and opportunities.

Since my heart is in journalism and media, an important part of Medietrends involves offering training courses for journalists. Every year, I or my partners provide further education to several hundred journalists.

In recent years, the main focus of this training has, of course, been AI. I was actually the first in Denmark to develop workshops on the responsible and constructive use of AI in journalism. I've conducted internal training sessions for all journalists at dozens of Danish media houses and had hundreds of participants in my external courses.

So, I closely follow AI developments and have a wealth of experience in how people are adapting to these new opportunities.

As I mentioned earlier, I'm not a doomsayer. I'm excited about AI. But like many others, I can't help but worry about what might happen—or what could happen.

Please consider the following scenarios as potential side effects that may arise as AI is fully rolled out. These aren't definitive claims, just some open-ended thoughts.

Sound good?

#1

Summary Fatigue

I like to call this scenario "summary fatigue," as it might be too extreme to label it a "summary plague."

However, there's a potential downside to this remarkable ability AI has given us to summarize long texts and articles.

Summaries are not a new feature; they've been successfully integrated into media outlets like CNN and Business Insider for years. But now, AI has made it easy and accessible for everyone to generate summaries.

Let's first consider the many benefits:

Behind the scenes, journalists can now handle larger volumes of text than ever before, potentially leading to better research and higher-quality journalism.

Of course, there are obvious challenges in outsourcing the prioritization of key points in reports and documents to AI. AI has its biases, and there's always the risk of hallucinations or critical omissions — issues that have already received significant attention.

For users, the advantages of summaries are clear. They fit into our busy lives, allowing us to quickly grasp the essence of a news story. If we want to delve deeper into the story, we can do it on an informed decision since we already have seen the key points.

Summaries have proven successful in many Nordic media outlets, boosting engagement and even being considered valuable enough to be placed behind paywalls.

But what happens to us—users or journalists—when we get used to not having to exert effort to consume new information?

We already know that our attention spans are in decline. Data from the 50 largest American media outlets show that long before AI made summaries possible, our time spent on these news sites was getting shorter and shorter.

So what happens when we no longer need to read long texts to understand the essence?

Are we at risk of producing more superficial, unnuanced journalism that loses details and promotes polarization?

Will public debate lose nuance, depth, and the appreciation of balanced arguments? And what about us as news consumers?

Will this be like what happened with GPS and calculators, which have diminished our abilities to navigate or do mental math? Research clearly shows that we struggle to learn and retain skills unless we continuously practice them.

So are we actually at risk of losing the ability to read long texts?

Or at least long-form news?

If that's the case, how will we tell complex stories?

How do we convey nuance in bullet points?

I'm confident a solution can be found, but it requires conscious attention to the risk that summaries might potentially erode the written format we've known since Gutenberg invented the printing press.

A positive sign, however, is that people don't always opt for quick fixes.

As you know, there's still a strong trend towards podcasts, which offer depth and long-term engagement.

Perhaps we just need to accept that long reads are evolving into "long listenings."

#2

AI is devaluing content

For a long time, news has been a commodity that's difficult to profit from because there's simply more supply than demand.

But now, we're seeing a shift. The amount of content is exploding, and if people can't find what they're looking for, they can just create it themselves using the same AI tools that professionals use.

From the perspective of experts and professionals, we can agree that the quality of homemade articles, images, audio, and video usually doesn't match the level of content produced by the media organizations you represent.

But does that really matter if more and more people feel like they can do it themselves?

Perception is truth.

Some might even think that media content is generated by a robot in 10 seconds. Everyone knows it's possible and it's happening. This erodes the recognition that journalism is something special.

It's a fair consideration whether one wants to spend three minutes reading an article if they assume it was produced by a journalist in 10 seconds.

If curated media outlets want to maintain their market value in a world of synthetic, quickly produced content, they need a powerful branding narrative.

Users must be practically indoctrinated to understand that content from DR, SVT, NRK, YLE, RUV, etc., is fundamentally different from everything else out there—including what people can create themselves.

At every touchpoint of contact, it must be made clear that there are real people behind the content. Not because humans automatically are a proof of quality, but because they bring something different than algorithms.

This is particularly about judgment, ethics, common sense, and the ability to put news and events into a relevant context.

This isn't a new analysis, but it's more urgent now than ever before.

The New York Times has recognized this challenge and has started to emphasize their human element much more consciously than before. And what better way to show that human-created content is made by humans than by highlighting the people behind the journalism?

They have developed extensive and very detailed biographies linked to bylines. When you click on a byline—at least most of them— you're taken to a page full of information about the journalist. Who is she, what are her ethics, what does she typically cover, and how can you reach her?

It might not solve everything, but it's a start if we want to stand out in a chaotic and very disruptive future.

#3

We are running out of content

This spring, I had a bit of an aha moment while I was attending the Nordic AI Media Summit here in Copenhagen.

During a session with Aftenposten from Sweden, one of their editors shared that personalization had increased their click-through rate by 7.1 percent.

A success story, no doubt. But with a somewhat absurd side effect. The personalization worked so well that Aftenposten almost ran out of content.

“We are running out of content,” said the Aftenposten editor on stage.

Let me emphasize that I’m a supporter of mixing news in a way that ensures as many people as possible are presented with what matters most to them. It makes great sense both journalistically and economically.

But that isolated quote also contains a potential self-created, self-reinforcing loop that media diving into personalization could find themselves trapped in.

Here’s the situation:

A media outlet has a wealth of journalism they want to present to as many people as possible. But if all readers see the same homepage, very few will discover and click on what’s important to them.

Also there is no room for 150 different articles on a single front page. Only a fraction of the content can be featured. So the outlet has to prioritize. And then they might just as well use algorithms to select the articles that make the most sense for each individual reader.

So far, so good. But the idea of “running out of content” is worth a closer look.

The issue arises because the media outlet already has “too much” content. If they only had 10 articles a day, there wouldn’t be a need to personalize anything. Anyone visiting the homepage would easily be able to scan all of them and decide what they wanted to read.

There are, of course, many good reasons why a large media outlet produces more content than even the most loyal core reader can consume in a single day. And that creates the need for personalization.

The starting point of personalization is therefore that the outlet, in principle, has too much content.

Conversely, personalization also means that the outlet commits to actually having a broad range of content, so there’s something for everyone. Otherwise, there’s nothing to personalize.

In that sense, the quote about running out of content makes sense. Aftenposten has squeezed 7.1 percent more value out of the same journalistic investment, and that's impressive. To push beyond that 7 percent, though, they might need more content.

To be fair, the Aftenposten editor didn't say, "... so we need to hurry up and publish more content."

But it's worth noting that the remark was made at a conference where several media outlets talked about how they're developing AI templates that can quickly draft articles in the outlet's own tone of voice for journalists to work on further.

So, increasing the production of articles with the help of AI is a real possibility. Whether it's done to give journalists more time for the important stories, as the typical narrative goes, or to feed the personalization engine should not be a random choice but a conscious and well-thought-out strategic decision.

AI offers media fantastic opportunities. But we should be careful not to fall into the trap of mass-producing journalism just to satisfy a technology that rewards us with isolated KPI boosts.

Media leaders need to keep the big picture in mind, stay true to their outlet's mission, and thoughtfully integrate AI based on that understanding.

To make the picture even more complex, a user-centric focus is also more important than ever, bringing us back to the personalization track. But perhaps the question can still be boiled down to a yes-or-no decision:

Should we have a lot of content for everyone or something particularly high-quality for a slightly smaller audience?

We could also ask: Are you someone who believes "more is more," or do you think "less is more"?

#4

Are we undermining our key difference?

I've just laid out some strong arguments for why media organizations will increasingly leverage AI to boost content output. This isn't just about enhancing user experiences through personalization, but also to maximize the value of the content already stored in the archives.

On the surface, these ideas seem like a win-win. Who wouldn't want to get more mileage out of their content while simultaneously delivering better products for users?

But (as you probably already have guessed) there's a potential downside.

Let me start with the positive aspects.

Every media company has a vast archive of high-quality content that's largely underutilized. Some of it might still hold SEO value and enjoy a long tail of relevance and consuming, but most of this content—no matter how insightful or well-crafted it was at the time it was produced—has been forgotten and therefore has lost its value..

But with AI, we can breathe new life into this content. We can recycle the content by letting AI act as our superbrain, seamlessly refining new insights and extracting value from the archive.

A year ago, I built a chatbot that was fed with all the content from my website. When people asked questions it would generate new articles based on the knowledge already stored in the Medietrends archive.

However, I wasn't satisfied with the quality, so I dumped the project. The problem is that the chatbot operates in real-time, and any errors wouldn't be caught until it was too late.

If I'm going to abandon the long-standing principle of having a human in the loop, I need to be 99.99% sure that it will work without errors and mistakes. Even then, I'm not sure that this is the right path to take.

Meanwhile, much larger players have raced down this road, leading to chatbots from both the *Financial Times* and *Washington Post*.

This summer, *The Washington Post* launched a chatbot called Climate Answers, allowing users to chat directly with content from the last eight years of climate-related articles.

If any articles in the archive can answer the question, the chatbot crafts a well-formulated response and links to the relevant articles.

Since everything happens in real time, there's no human in the loop which is otherwise a central principle in media's AI implementation.

The *Financial Times* did something similar with their *Ask FT* feature, launched in March for a smaller audience. Here too, responses are generated without human oversight.

When media outlets emphasize the importance of a human overseeing AI-generated content, it's not just to avoid misinformation. It's also a crucial signal that there's a fundamental difference between curated, quality journalism and the algorithm-driven content of tech platforms.

Human editors are a key differentiator.

Human beings are actually what makes media stand out on the internet.

Many in the media industry have confidently embraced the narrative that the need for curated, human-driven journalism will grow as the digital world becomes more synthetic.

But does this narrative hold up when leading media outlets themselves start behaving like Google's new AI search engine?

And isn't there a risk that the human-in-the-loop principle will be further diluted now that this promise is not bullet proof anymore?

This isn't said to be provocative, but rather to be entirely objective.

Træk vejret!

The truth is though that chatbots are here to stay, whether the media chooses to engage with this development or not. If the future of search is AI-generated answers instead of links, it might be hard to argue that the media should not make their own chatbots.

And it's arguably better for these searches to pull from quality journalism archives rather than random sources from the vast and unfiltered internet.

But the stakes are high because this touches on core journalistic values and, ultimately, trust and credibility.

Even without compromising the human-in-the-loop principle, media outlets risk eroding the valuable brand identity that sets them apart from the rest of the internet.

I'm a strong advocate for using AI to enhance our workflows so we can create better content. I'm also open to AI producing end-products ready for publication.

But when we do this, we need to be aware that there's a limit to how much we can push it before our core value proposition begins to erode.

Træk vejret!

Smart media leaders were quick to recognize that AI-generated photorealistic images were a no-go. No matter how clearly they're labeled, these images contribute to a loss of trust among users. What you see should be reality, not a constructed version of reality. This is a principle upheld by all quality media.

Yet, mistakes happen, as seen earlier this year when *Sjællandske Medier* in Denmark mistakenly published an AI-generated news image.

What's interesting is that this principle—that users should be able to trust what they see as reality—doesn't seem to apply to audio.

For years, media outlets have offered narrated articles, which makes perfect sense. But with recent advancements in voice cloning and text-to-speech technology, it's now possible to clone a media outlet's most popular voice and use it to narrate all articles.

This results in a much better user experience. But why uphold the principle that what you see must be real if the same doesn't apply to what you hear?

The next level of this will be seen when *Aftonbladet* in Sweden launches AI-generated summaries as a kind of mini-podcast. Here, AI will first read selected news articles from *Aftonbladet*, then edit them into scripts, and finally have a clone of one of their popular podcast hosts to narrate the audio.

It's an exciting prospect, but how does it align with the core value of human-edited media? What you hear is no longer reality. It's a construction, because your favorite anchor never reported or recorded it.

So, where do we draw the line?

The media is at a crossroads: Should we continue to be a bastion of curated, human-driven content, or should we embrace a digital reality where speed and accessibility are valued higher than human insight and judgment?

The answer to this question will shape the future media landscape and may even determine whether media outlets can maintain their relevance in a world where information and access to it are endless.

#5

Infocalypse

If the media is Superman, then declining trust is the kryptonite that could weaken us, or even make us irrelevant.

This lack of trust doesn't necessarily have to target the content produced by media outlets. It could simply be a general atmosphere of zero trust, as the growing amount of artificially created content force us as media users to doubt everything.

This is not a theoretical horror scenario. It is real.

I had my first *infocalypse* experience a few months ago.

Infocalypse is a state where our natural trust in information is undermined by an equally natural skepticism—or rather, mistrust. It's an apathetic condition where we can no longer distinguish truth from falsehood or real from fake.

That will eventually lead us to trust nothing at all.

Better to dismiss one piece of news too many than one too few.

And that was exactly my experience when I was doomscrolling on X and stumbled upon a surprising piece of news about Gaza. The post was accompanied by a short video from BBC's newsroom.

But despite the fact that BBC is one of the most trusted news organizations in the world, I got a gut feeling that the news might be fake and that the video was fabricated using AI.

My skepticism was misplaced, but the damage was done. Because if we don't trust BBC and other editorial media, they de facto lose their value. News that no one trusts is nearly worthless, regardless of whether it's true or not.

If we're to take it to the next level, it's the Fourth Estate that risks losing its power. And with that, a critical pillar of democracy begins to crumble.

This concern is also highlighted in the latest report from the Danish government's expert group on tech giants. Here they address the phenomenon under the term *information credibility*.

Romantics might dream that the infocalypsic chaos trigger a counterreaction where people decide that the safest place to find the truth is to return back home to the good old trusted institutions. Media, museums, government agencies, and other authorized sources.

The latter is true, but it's hard to measure whether this return to credible institutions will happen on a large scale, or if it will remain a niche trend, like those who prefer vinyl records over streaming or dumbphones over smartphones.

The best antidote to infocalypse is, of course, that as many people as possible visit traditional sources. When information is viewed, read, or heard in these domains, it serves as a verification in itself against being fabricated.

We could end the story here with the conclusion that people just need to click on the right media outlets. But we know that the media is already working hard to make this happen.

However, with the launch of new AI platforms like ChatGPT, Microsoft CoPilot, Claude, Google Gemini, and others, there are now even more reasons to bypass traditional media, as you can have the news tailored exactly the way you want through your preferred AI service.

And even when we do turn to the media to find evidence of claims we have seen on social media, we often find a mismatch between fast-paced news cycle and the discussions happening on social media, where conversations might still revolve around events that happened days ago.

In the case of BBC, I couldn't find the specific piece of news content —probably because BBC had already moved on.

This is not a criticism of BBC or other media outlets' search functions. It's simply an observation that media infrastructure is more in tune with the media's own logic than with society's needs.

How the media can manage to both create news and simultaneously be reactive, serving as a useful tool to ensure information credibility in a broader sense than just the individual news story, will be one of the biggest and most complex challenges in the coming times.

Thank you.