Of bots and bottles

Are we hallucinating what AI can do for beer?

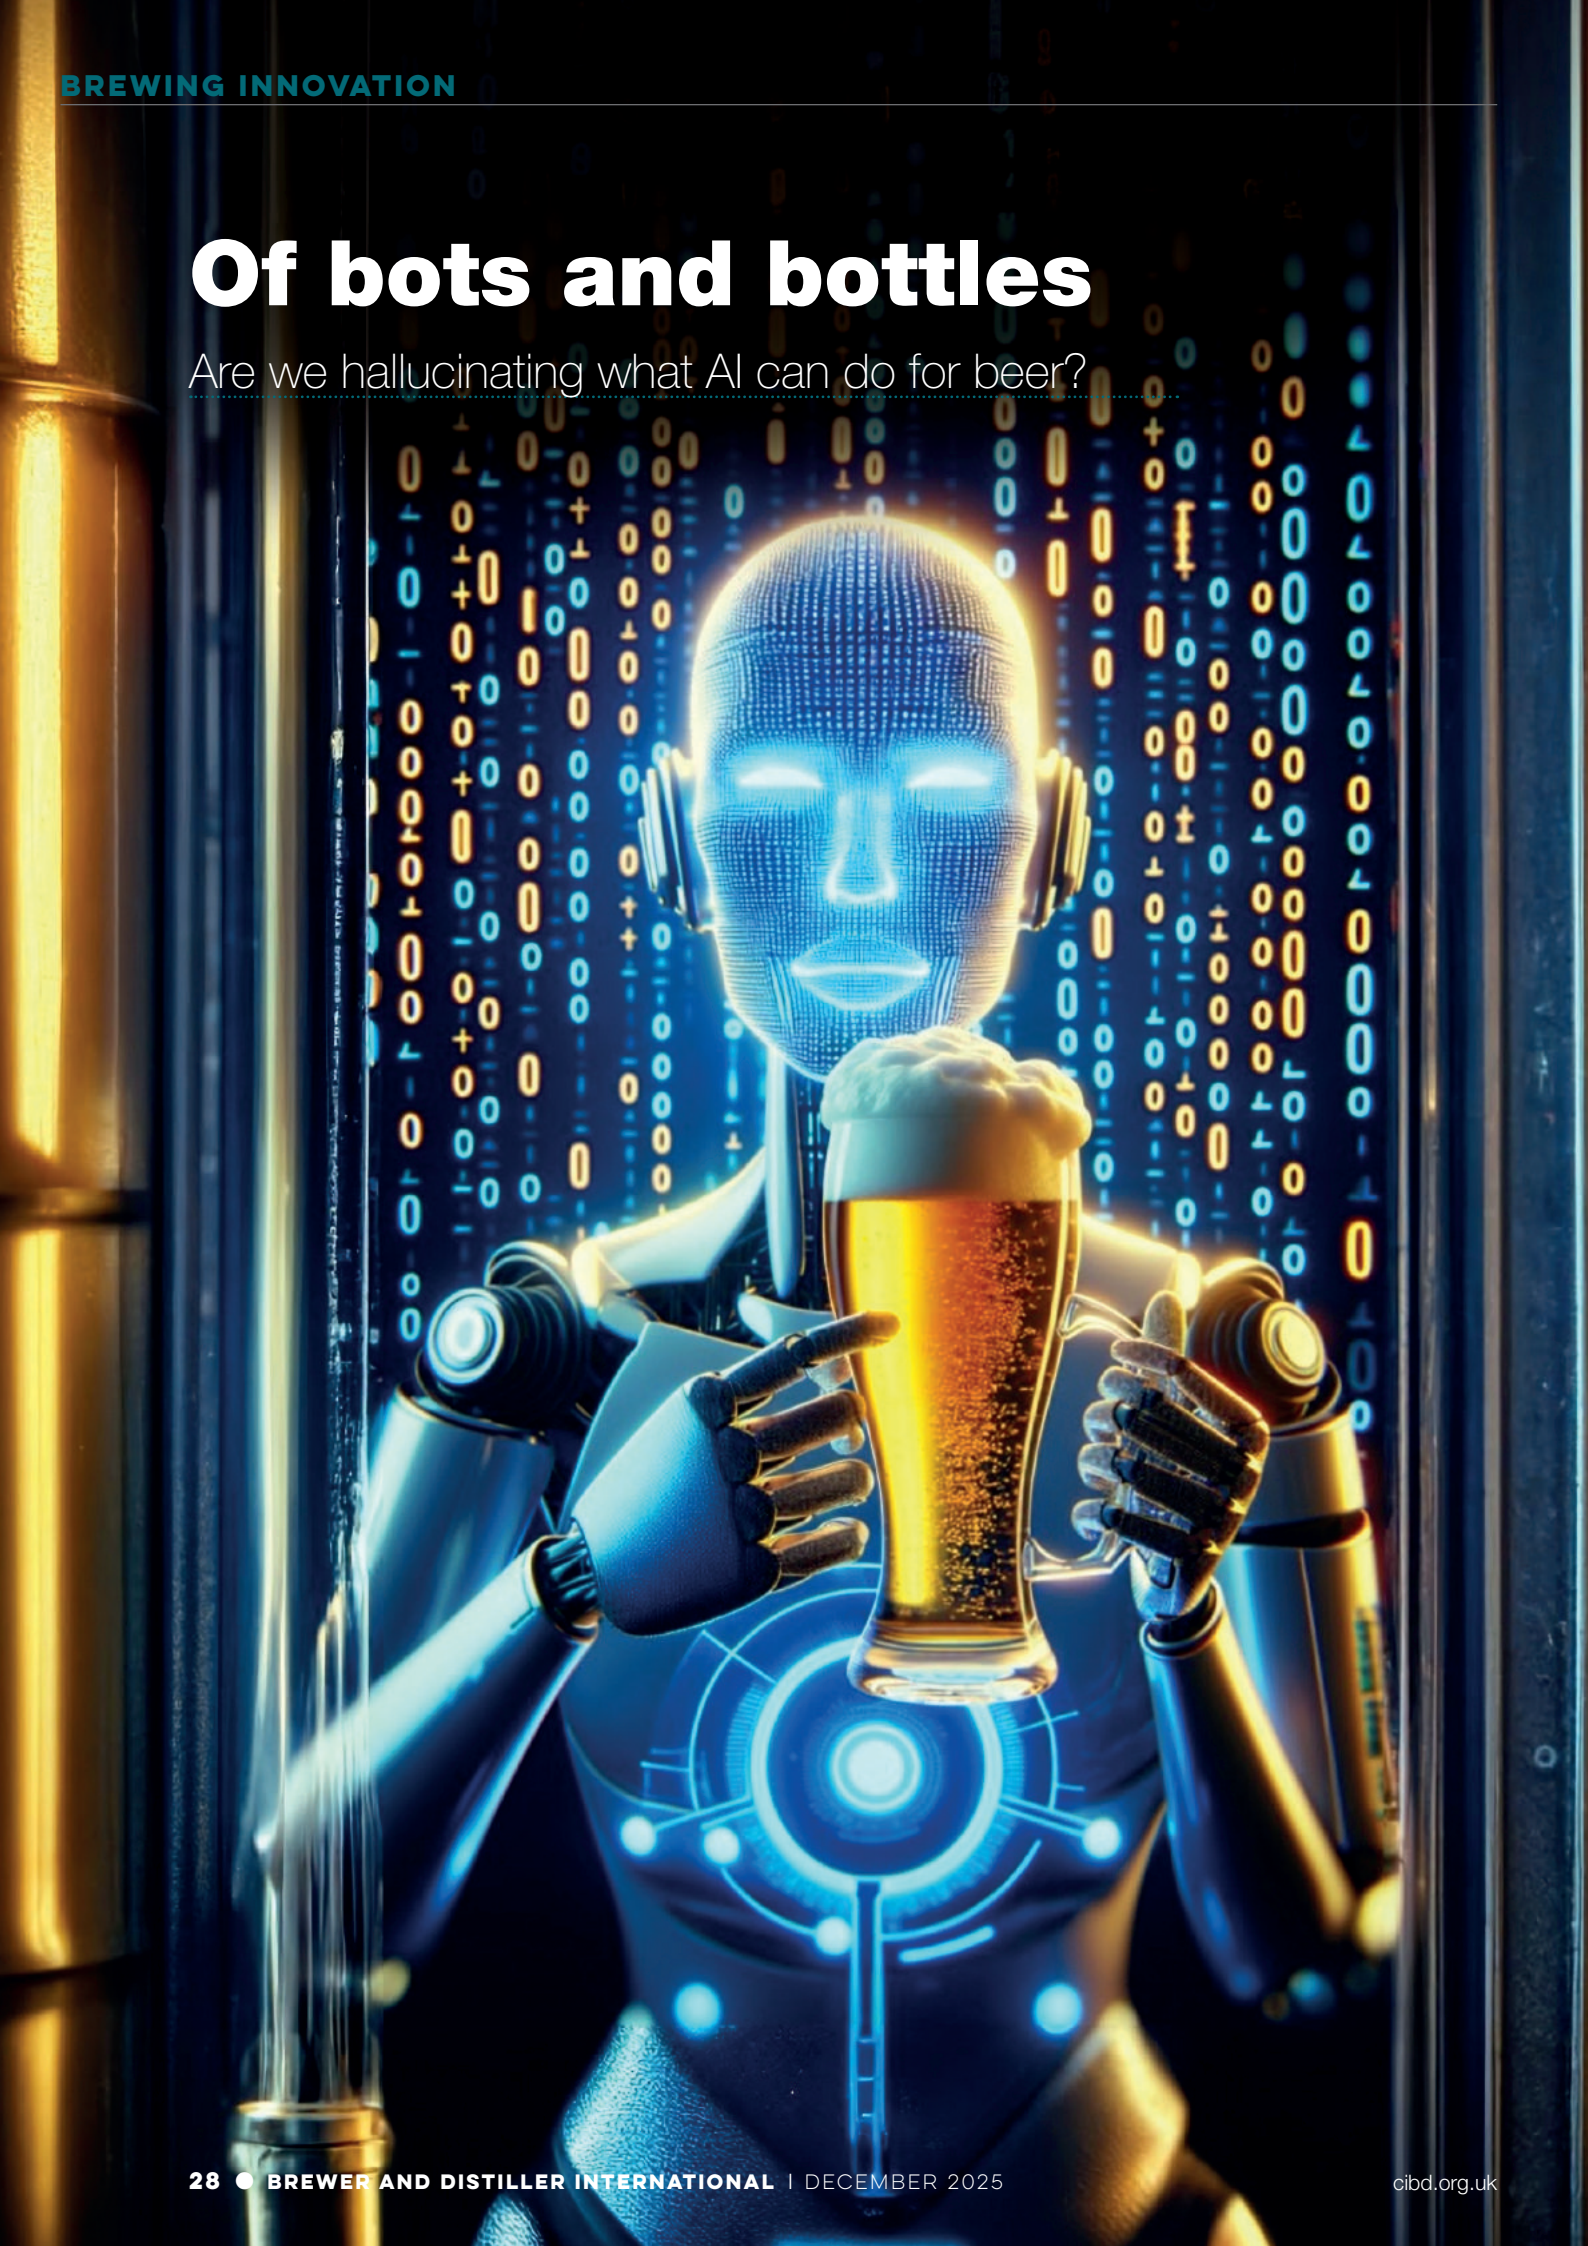

BY ANTHONY GLADMAN

AI is a polarising juggernaut we cannot ignore. For some it's the answer to everything everywhere all at once. For others it's the magic slopbox we love to hate – and often with good reason.

Developing a new beer recipe can involve a lot of testing and tasting…AI might speed up this process (Photo: iStock/werajoe)

Brewing is not immune. AI's promises to the sector have more froth than a fobbing pint: it will design better recipes, watch over cleaner brews, run more efficient operations, pump out smoother marketing, even absorb the slings and arrows of customer service.

When promises are this lofty, the reality is usually more nuanced. The question is whether we're seeing AI's genuine potential, or hallucinating capabilities that don't yet exist – or may never come to pass.

Untangling the alphabet spaghetti

One sticky barrier to understanding what AI can do for beer (or anything else) is the muddiness of the language used. "AI" gets trotted out as an umbrella term that boils down to "computer does something clever" but all its izzy-whizzy is not just down to one kind of busy. It encompasses everything from simple rule-based routines to advanced neural networks.

"AI and machine learning are such hypes that they are becoming over-used by scientists as well as marketeers," says Professor Kevin Verstrepen of KU Leuven and VIB -Flanders Institute for Biotechnology.

To cut through the hype, it helps to understand the different terms. There's generative AI like ChatGPT that creates text, images or code from prompts. These are the "creative" Large Language Models (LLMs) that get all the headlines – the ones that have teachers despairing of their pupils' homework, or get lawyers in hot water by feeding them fictitious legal citations.

Machine learning (ML) is much less sexy. These are the nerdy algorithms that pick out patterns from data and use them to make predictions or classifications. Verstrepen and his PhD student Michiel Schreurs use ML in their research at the Leuven Institute for Beer Research to analyse hundreds of commercial beers, creating models that can predict flavour characteristics from chemical composition.

Their work is published in *Nature Communications*. ML is also what most brewing applications rely on. Deep learning is an ever-nerdier subset of this. It's more sophisticated but requires massive datasets. No, bigger than that.

Building on machine learning, computer vision is a flavour of AI that "sees" and interprets images – either still or video. This is sometimes used in brewing for quality inspection, packaging defects, and so on.

Predictive analytics uses historical data to forecast outcomes. Brewers use this like a crystal ball for patterns in customer demand or when to schedule equipment maintenance.

And finally there's process automation, which is not AI but often gets lumped into the conversation. There's no interpretation or prediction here, just rules-based systems that follow predetermined logic: bot sees this, bot does that, for ever – or until you turn it off.

Understanding these distinctions matters, but there's another wrinkle to get your head around: it's all about the data. "The current AI models are already powerful enough for tasks like recipe development and process optimisation," Verstrepen says. "The only limitation is data availability."

For brewers, this means the question isn't whether AI could revolutionise brewing, but whether they can gather enough high-quality data to make it work – and whether they want to.

The silence of the labs

Developing a new beer recipe can involve a lot of testing, tasting and iterating. At its most simple, the process goes something like this: brew a small batch, drink some, chat about it with the other brewers, tweak the recipe and brew again, repeating the cycle until everyone's reasonably happy. I say "reasonably" because when has a brewer ever been fully satisfied with their own beer?

Taste is subjective though, so if your brewery has a bit more money behind it you might also throw in some chromatography and spectrometry to get a more scientific look at your samples – what's in there that might deliver particular flavours or textures? Even when beefed up like this, getting a beer out of a brewer's brain and into a customer's glass can take months, sometimes even years.

In 2017, Carlsberg (or more specifically the boffins in its research laboratory) reckoned they could speed this up. They teamed up with iNano at Aarhus University, DTU Chemical Engineering, Innovation Fund Denmark (a government-backed funding body) and Microsoft to launch a beer fingerprinting project that would use sensors and AI to create "flavour fingerprints" of beer samples – of which the Carlsberg

lab can produce up to 1,000 per day.

Carlsberg thought this might shave up to 30% off the time it takes to develop a new beer. The researchers also planned to use this new understanding of their beer's flavour to select and develop novel yeast strains at a much higher speed and quality. While there was much fanfare around the launch of this three-year project, the news articles quickly became repetitive, then dried up altogether around 2019. 2020 came and went with no reporting of results or even progress made.

The silence is telling. For a body that built its reputation on openly sharing innovations like the pH scale, the lack of follow-up suggests the Beer Fingerprinting Project didn't deliver the time savings it promised. Carlsberg declined to be interviewed for this article.

The challenge probably wasn't the AI models themselves. Verstrepen and Schreurs say their ML model can predict with 80% accuracy how a beer will be scored by their in-house sensory panel, or what its appreciation score on the online consumer platform RateBeer would be – when there's enough data.

The bottleneck is gathering that data. "Obtaining a sufficient amount of such data is not trivial, not even for the largest breweries," Verstrepen says. Even with access to 1,000 samples per day, creating the comprehensive datasets needed for truly transformative AI applications takes years, not months.

The yeast whisperer speaks up

Out in the real world, some aspects of brewing resist algorithmic prediction. Matthew Peetz is CTO & Founder of Propagate Lab, a "Yeast Whisperer" and fermentation scientist who created "Yeast Buddy," an AI chatbot for yeast selection.

Despite (or perhaps rather because of) his use of AI, Peetz remains sceptical about claims it will replace brewing expertise anytime soon. Rather than seeing it as a do-everything-bot, he compares it to something more modest: a calculator.

"It's a powerful tool for handling repetitive, data-heavy tasks," he explains. "This includes things like predicting fermentation curves, optimising raw-material usage, or flagging potential off-flavours from sensor data. But, just like a calculator doesn't 'understand' mathematics, AI doesn't understand' brewing."

The limitations are biological as much as technical. Yeast is a living organism with complex genetics and context-dependent behaviour. "Factors such as subtle genetic drift, spontaneous mutations, or stress responses to unique combinations of nutrients and environmental changes can't be fully captured in a dataset," Peetz says. "No matter how much data you feed an AI system, it can only interpolate from what it has seen – it can't foresee a novel adaptation or an unexpected interaction between strains and ingredients."

This shows up in practice. "I have seen AI suggest recipes that include ingredients like turmeric or weird spice additions," Peetz says. The recommendations looked statistically sound but made no practical brewing sense, a reminder that pattern recognition isn't the same as understanding.

AI can throw out suggestions lickety-split but, Peetz says, we must treat these as hypotheses to be validated through controlled pilot brews, sensory panels, and monitoring of key metrics like attenuation and cell health.

"Sensory skills such as aroma evaluation and balance of sweetness to bitterness remain irreplaceable," he says. "So do core brewing sciences: microbiology, fermentation biochemistry, and process engineering. Understanding how to design an experiment, interpret anomalies, and make judgment calls when data conflict with taste or tradition is where human expertise shines."

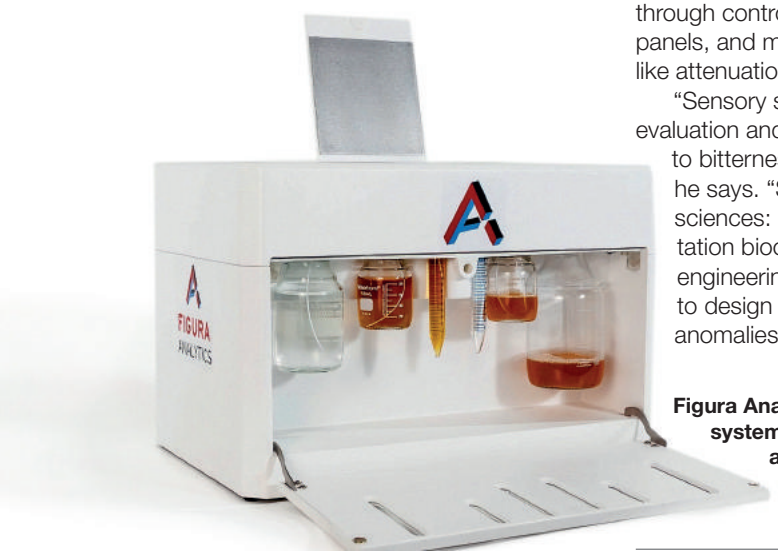

Figura Analytics' multi-sensor system can 'particle fingerprint' a liquid and will apply machine learning as datasets grow (Image: Loughborough University)

Counting on today to bring tomorrow

Figura Analytics is a food technology startup that spun out of Loughborough University in 2021, co-founded by CEO Nick Whitehurst alongside Dr Mark Platt and Rhushabh Maugi, the academic researchers who developed the core technology. While the company positions itself as having applications across life sciences, healthcare and semiconductors, its current focus is firmly on the food and beverage sector, particularly brewing.

Figura's pride and joy is its multi-sensor system. It doesn't look like much, just a white box about the size of a beer crate, with a front panel that hinges down to reveal a small row of beakers and tubes within. The device has been designed with practical brewery use in mind.

Right now it sits within laboratory teams, but Figura has conducted trials with the device positioned directly on production lines, and the company's ambition is for line operators to be able to use it without specialist training. The user experience has been simplified considerably. "We've spent an awful lot of time and effort and money making it really easy to run samples," Whitehurst says. "You put your sample in, you press play, you walk away. When it's finished, you get the results straight away."

But the low-tech workflow masks some clever goings-on inside. For instance, it contains a patented nanopore flow cell that uses resistive pulse sensing to analyse particles in liquid samples. "Essentially what happens is you put your beer in and then your software controls the flow of the beer," Whitehurst explains, "and effectively we force the beer into single file." As particles pass through flow cells, two electrodes pass a current across it. "We measure it 25,000 times a second, and when particles go through they block the current. That gives us information about that particle," Whitehurst says.

The system can determine particle size and count, creating what Figura calls a particle fingerprint of the beer. The system can handle particles ranging from 0.1 microns to 150 microns in diameter by varying the apertures in two different sensors, which is a broader range than traditional filtration

A brewery investing in AI-powered quality control or recipe optimisation is, indirectly, adding to the sector's carbon footprint (Photo: iStock/gorodenkoff)

methods can manage effectively.

In practical terms, this means Figura's device can detect both the small contaminants that brewers worry about (bacteria and wild yeast, typically 0.5 to 5 microns) and larger particulates (debris up to 150 microns) in a single pass. The entire process takes less than an hour, compared to five to seven days for traditional microbiological testing.

Figura has been working hard to ensure its device's results can be trusted. Over six months in early 2025 it trialled its system at Greene King, which ran duplicate and triplicate samples side by side, comparing Figura's results with traditional microbiology. Prior to that, Campden BRI and partner breweries conducted more academic validation, inoculating samples with known concentrations of specific contaminants in different beers to verify the correlation between Figura's readings and established methods.

Whitehurst says the system's limits of detection match traditional microbiology, and that it can distinguish between live and dead cells, giving it significant advantages over PCR as a contamination screening tool.

Figura is also developing a version of this system that can be used at the point of dispense, to monitor beer quality in pubs and bars. This addresses a particular concern with low-alcohol and no-alcohol products, which have higher contamination risk. "You can compare what's coming out the tap with brew-fresh in real time," Whitehurst explains. Rather than waiting days for microbiological results to confirm whether lines are clean or product has been put on the wrong line, the portable device can provide immediate alerts. All data sits in the cloud and remains accessible to brewery and venue staff simultaneously.

But here again, it's the data that is the bottleneck – and in Figura's case explains the gap between what might be and what currently is. Figura's website and promotional materials emphasise shape recognition algorithms that can distinguish spherical particles from rod-shaped bacteria in real time, and describe machine learning as an active component of the technology.

However, Whitehurst is more circumspect in person. "It's been built on the AWS platform so that we can start to apply machine learning to the fingerprints and the changes to the fingerprints as it gets populated by more and more data from more and more customers as we start to grow," he explains. In other words, the sophisticated AI features remain on the development roadmap rather than in current deployment.

Shape detection, which Whitehurst describes as "part of our patent," represents a future capability that could allow the system to assign probabilities about whether particles look like specific contaminants. "Given that yeast and bacteria are different to most flavour particulates, as they go in through the sensors there's the potential to apply a probability of whether it looks more like a particular contaminant or not," he explains.

"At a high level, we could build a library of known contaminants based on their size and shape, and then have the potential to see those as they're going through the sensors." But this depends on accumulating sufficient training data and "a little bit of development to be able to recognise many different shapes all at the same time. That's our development pathway."

The company also plans to work more extensively with DNA aptamers, synthetic DNA strands engineered to bind to specific targets. If Figura wants to detect *Lactobacillus brevis*, for instance, it would create an aptamer that binds to that bacteria, making it appear larger as it passes back through the sensor and confirming that specific contaminant's presence.

Gather and keep as much data about ingredients, machines, recipes and beers as possible (Photo: iStock/FG Trade Latin)

Aptamers can be designed as broad-spectrum tools that bind to multiple targets or as highly specific identifiers for single species, but this too is described as something Figura wants to develop "a little bit further down the line" rather than a current operational capability.

What breweries are using now is essentially a sophisticated, rapid particle counter that creates quality fingerprints and flags contamination presence. "I think the benefits for brewers are that we have as good contamination detection as micro, so our limits of detection are as good. We can tell the difference between live and dead, so I think as a direct competitor to PCR we've got significant advantages in that regard," Whitehurst explains. The system doesn't yet offer the specificity of more detailed laboratory methods, "but given that for most brewers, there's zero contamination and they're looking to release products or get those results, we're really good at being able to tell you there's nothing there quickly." As he puts it, if contamination is found and detailed analysis is needed, "all the time you saved on all the clear samples" makes the investment worthwhile.

There's another cost worth considering. AI systems demand significant computing power, and the data centres that support machine learning models consume vast amounts of electricity and water for cooling. While the marginal environmental impact of individual brewing applications may be modest compared to general AI usage, the cumulative effect matters. A brewery investing in AI-powered quality control or recipe optimisation is, indirectly, adding to the sector's carbon footprint. Whether those efficiency gains – a few percentage points shaved off development time, slightly reduced ingredient waste – justify the environmental cost of the computing infrastructure supporting them is a question breweries should weigh alongside the technical and economic considerations.

I, Human

Some researchers are bullish about AI's potential. In response to Peetz's calculator analogy, Verstrepen points to chess computers, which were once thought incapable of creative play but now routinely beat grandmasters. "There is absolutely no reason to think that, when provided with sufficient data, AI would not be able to take over most if not all of the conceptual tasks that modern brewers have, including recipe development and tuning."

That conditional – "when provided with sufficient data" – is doing significant work. In practice, even large breweries find gathering comprehensive, high-quality datasets challenging. The AI revolution in brewing isn't being held back by inadequate algorithms; it's being held back by the unglamorous work of data collection and management.

For brewers evaluating AI tools today, Verstrepen and Schreurs recommend focusing on data infrastructure. Verstrepen explains this means "keeping as much data about their ingredients, machines, recipes and beers as possible, which will allow plugging these into AI applications that are specifically tied to their breweries". The genuinely useful AI applications – like Figura's rapid contamination detection or machine learning models that can predict flavour profiles – do exist. But they work best when they're built on years of careful data collection, not purchased as off-the-shelf solutions to problems breweries haven't properly defined.

There's also the question of cost. For craft and regional breweries, the investment in data infrastructure, sophisticated sensors, and AI systems may exceed any efficiency gains – at least in the near term. The data-driven future Verstrepen describes might be technically achievable but economically viable only for major players with the resources to gather and manage comprehensive datasets over years. Rather than democratising brewing innovation, AI could widen the gap between multinationals and independent brewers.

There's another consideration beyond technical capability. "Consumers often like the idea that a human is in control, and that a product is the result of a combination of scientific knowledge, creativity and craftsmanship," Verstrepen says. Even if AI eventually matches or exceeds human brewers at recipe optimisation, consumers will still have a special place in their heart for products that retain the human touch – that keep at least one foot in the world of real, tangible things. Would you watch AI actors on TV? When you get off the sofa to fetch a beer, would you reach for one brewed by AI?

For brewers, that means the question isn't whether to embrace or reject AI wholesale, but whether investing years in data infrastructure is worth the potential gains, or whether the craft is better served by humans who can make judgment calls no algorithm can replicate from a spreadsheet. The data suggests the latter.