

Artificial intelligence for food innovation

Received: 25 September 2025

Accepted: 1 May 2026

Published online: 3 July 2026

 Check for updates

Bianca Datta¹, Markus J. Buehler^{2,3}, Yvonne Chow⁴, Kristina Gligorić⁵, Dan Jurafsky⁶, David L. Kaplan⁷, Rodrigo Ledesma-Amaro⁸, Giorgia Del Missier⁸, Lisa Neidhardt⁸, Karim Pichara⁹, Benjamin Sanchez-Lengeling¹⁰, Miek Schlangen¹¹, Skyler R. St. Pierre¹², Ilias Tagkopoulos¹³, Anna Thomas⁶, Nik Watson¹⁴ & Ellen Kuhl¹²✉

Global food systems must deliver nutritious, sustainable foods while sharply reducing environmental impact. Yet, food innovation remains slow, empirical and fragmented. Artificial intelligence (AI) offers a transformative path to link molecular composition to functional performance, connect chemical structure to sensory outcomes and accelerate cross-disciplinary innovation across the production pipeline. While it is broadly applicable to food systems, we focus on sustainable proteins—plant-based, fermentation-derived and cultivated—as a high-impact test bed for AI-driven closed-loop design. We review the applications, opportunities and challenges of AI for food as an emerging discipline that integrates ingredient design, formulation development, fermentation and production, texture analysis, sensory science, manufacturing and recipe generation. We identify four priorities: advancing scientific machine learning with embedded domain priors, treating food as a programmable biomaterial, building self-driving laboratories for automated discovery and developing deep reasoning models that integrate nutrition and sustainability. Integrating AI responsibly into the food innovation cycle can accelerate the transition to sustainable food systems and establish a predictive, design-driven science of food for human and planetary health.

Food is among the most powerful determinants of human health, planetary sustainability and cultural identity¹; yet, research and development for food remains slow, fragmented and driven by trial and error rather than rational design. Radical, urgent and creative food solutions are needed to combat the current climate crisis². Notably, across farmland, livestock and production, the food sector generates 35% of all global greenhouse gas emissions³ and contributes heavily to biodiversity loss and public health challenges⁴. To meet climate goals and feed a growing population, we need alternative sources of protein that are delicious, affordable, scalable and nutritionally compelling⁵. Today, however, the scientific bottlenecks in texture, flavour, scalability and cost pose barriers to adoption⁶. Artificial intelligence (AI) has reshaped domains from protein design and drug discovery to materials engineering⁷. AI is now poised to revitalize how we grow and make food—from optimizing

plant-based meats to generating climate-conscious recipes⁸. Enabling protein diversification will have profound implications, and AI can substantially accelerate scientific progress if applied responsibly and towards the correct challenges⁹. But the success of AI is by no means guaranteed. When deployed in a scattershot way, AI can generate low-quality outputs, increase environmental burdens or undermine consumer trust, aspects that are especially sensitive in food¹⁰. The goal is thus not to deploy AI indiscriminately. Instead, we need to partner with AI where it measurably reduces time, cost and uncertainty for problems that matter¹¹. Responsible application—focused on the right challenges—is essential. In this Review, we identify areas across the food system where we can constructively engage with AI: we provide background on the current state of the art, outline new opportunities and discuss the challenges that AI is facing in the quest to accelerate the future of food.

A full list of affiliations appears at the end of the paper. ✉e-mail: ekuhl@stanford.edu




Fig. 1 | Historical milestones in AI for food. The birth of AI in the 1950s–1960s, marked by Alan Turing’s concepts of computing machinery and intelligence, established the foundations of machine intelligence. The rise of machine learning and neural networks in the 1980s–1990s, driven by backpropagation and statistical learning, shifted AI from rule-based systems to data-driven approaches. In parallel, the Human Genome Project and the omics revolution in the 1990s–2000s, enabled by high-throughput sequencing, created a culture of large-scale biological and nutritional data. The expansion of big data and cloud computing in the 2000s further empowered the analysis of massive ingredient, supply chain and consumer preference datasets. Since the beginning of the twenty-first century, advances in computational chemistry and cheminformatics in the 2010s have enabled the prediction of flavour molecules, aroma profiles and chemical interactions. Breakthroughs in deep learning,

exemplified by AlexNet in 2012, catalysed applications in food recognition and image-based dietary tracking. Since 2014, generative models have expanded into food design, flavour and recipe generation, and ingredient discovery, with companies such as NotCo applying AI to plant-based food development. Since 2015, advanced fermentation and alternative protein models have started to include AI-driven strain engineering and process optimization to accelerate precision fermentation and industrial adoption. The current decade is witnessing the integration of wearable and microbiome data for AI-driven personalized nutrition since 2016, the emergence of digital twins for real-time monitoring and optimization of food manufacturing processes and the rise of foundation models and LLMs that leverage unstructured data for formulation and recipe innovation since the 2020s. Together, these advances are fuelling the emergence of AI for food as a distinct scientific discipline.

The intellectual and technological trajectory that has shaped the emergence of AI for food as a discipline spans more than seven decades, weaving together advances in AI, food science and biotechnology (Fig. 1). From the birth of symbolic AI in the 1950s¹², to the rise of machine learning and neural networks in the 1980s–1990s¹³, to the deep learning breakthroughs in the 2010s¹⁴ and the advent of generative models¹⁵ and foundation models¹⁶ in the 2020s, AI has progressively expanded its ability to model, design and predict complex systems. In parallel, developments in genomics¹⁷, computational chemistry¹⁸, fermentation and personalized nutrition¹⁹ have created unprecedented data resources and biological insights directly relevant to food design. Their convergence has given rise to a new field in which AI is now central to ingredient discovery, recipe formulation, manufacturing optimization and consumer personalization¹¹.

Designing food materials poses unique challenges, as their properties emerge from multiscale interactions among proteins, carbohydrates, lipids and processing conditions²⁰. Yet, AI is shifting food innovation from trial and error to data-driven discovery. Modern models link chemical composition to sensory outcomes and enable targeted formulation and rapid iteration. Machine learning frameworks predict consumer preferences from chemical-sensory panels²¹, while graph neural networks map molecular structure to odour quality²². Advances in protein modelling and molecular learning accelerate functional ingredient discovery^{23,24}, and precision fermentation provides scalable routes to novel proteins and textures²⁵. Together, these advances address long-standing challenges—engineering egg-free emulsions, meltable dairy alternatives or fibrous plant-based meats—and have translated into market-ready products from companies such as NotCo, Perfect Day and Climax Foods²⁶. This transition from conceptual promise to industrial deployment underscores AI’s growing role as a driver of culinary innovation and sustainability.

We focus on the post-harvest stages and follow an AI-powered cycle through ingredient design, formulation development, fermentation and production, texture analysis, sensory properties, manufacturing and recipe generation (Fig. 2). In the following sections, we highlight how AI-powered approaches, from generative models to digital twins, drive innovation across these steps.

AI across food innovation

Ingredients

Novel food ingredients must balance sustainability, nutrition, sensory appeal and functionality. Functionality includes water-holding, emulsification, gelation, viscosity and stability, which shape processing and structure. Many ingredients aim to replicate animal-derived

functions, such as foaming in egg replacers; yet, matching performance remains challenging. Structure–function relationships link molecular architecture to activity. In food, macromolecular features of proteins, lipids and fibres encode functionality. This raises the question to what extent AI can bridge this scale gap and translate molecular features into techno-functional predictions for a functionality-first design–build–test–learn cycle.

The success of animal-free ingredients depends on functionality and the ability to meet formulation and consumer requirements²⁷. These traits arise from molecular features, environmental conditions such as pH and temperature, and interactions with fats, fibres and solutes. Proteins, lipids and fibres provide core functionality, with proteins being the most versatile. The most promising alternative proteins today are wheat gluten and soy; they can form fibrous textures under extrusion and shear²⁷. New sources include legumes, oilseeds, grains, algae and leaves²⁸. Processing strongly shapes functionality: dry fractionation preserves native structures, while wet fractionation produces isolates that control solubility and gelation²⁹. Lipids and fibres add complementary features, from stabilizing interfaces to modulating viscosity and lubrication. Since plant heterogeneity limits predictability, developers modify ingredients chemically, physically or enzymatically. The field is now shifting towards functionality-enriched fractions for more consistent, cost-effective performance²⁸.

AI and informatics enable functionality-driven ingredient design by linking molecular and mesoscopic features to macroscopic outcomes under processing constraints. Classification models support discovery; for example, a quantitative structure–activity relationship model identified glycyrrhizin as a natural emulsifier³⁰. Regression models predict properties such as pectin viscosity to reduce cost and accelerate formulation³¹. Machine learning can predict techno-functional traits across crops with $R^2 > 0.80$ (ref. 32), while AI models can capture plant protein rheology with $R^2 > 0.99$ (ref. 33). Deep learning increasingly links functionality to sensory design³⁴. A recent study optimized starch-based gels for saltiness and strength using e-tongue and oral-processing data³⁵. Multi-objective optimization can predict solubility, gel strength and emulsification towards scalable formulation strategies³². Together, these advances support closed-loop design–build–test–learn cycles. High-throughput screening, active learning and digital twins operationalize models, while generative approaches target multiple criteria and bridge lab-to-plate gaps.

Ingredient design still struggles with consistent performance across food matrices, and key challenges remain. First, models often fail to generalize across ingredients and conditions. Second, nonlinear interactions in multicomponent systems complicate prediction and




Fig. 2 | The AI-powered production cycle. Seven key stages where AI can accelerate the time to market: ingredient design, formulation development, fermentation and production, texture analysis, sensory properties, manufacturing and recipe generation. The production cycle illustrates the tight integration of food science, biotechnology, materials science, behaviour science and AI to guide the development of new foods, from raw ingredients to commercial products.

require hybrid physics-informed approaches. Third, sensory preferences vary across cultures, which demands diverse datasets. Fourth, models must link molecular features such as sequence, charge and hydrophobicity to functionality in a transparent and scalable way. Addressing these challenges requires open, collaborative tools and an AI-driven ecosystem that predicts techno-functional performance from molecular features.

Formulations

Ingredient functionality sets the foundation, but effective food design depends on how those ingredients interact in complex systems; formulation brings these components together. Meeting the dual mandate of sustainability and sensory fidelity requires moving beyond slow trial-and-error formulation development. AI offers a compelling approach to uncover hidden patterns that connect ingredients, formulations, production, texture and sensory properties, and propose candidate formulations that are predicted to meet nutritional, functional or sensorial targets^{36–38}. Today, AI practitioners can draw on a broad collection of public resources including Recipe1M³⁹, FlavorGraph⁴⁰, FlavorDB⁴¹, USDA FoodData Central and FAO/INFOODS, as well as domain-specific databases such as Phenol-Explorer, LipidBank and the Volatile Compounds in Food database. Inputs include molecular descriptors, nutritional and functional properties, recipe structures and process sequences; outputs are complete candidate formulations aligned with target criteria. AI-powered approaches combine latent-space learning through autoencoders with conditional generation, constraint-aware optimization and expert-in-the-loop refinement (Fig. 3a). But can they truly capture the complex, nonlinear interactions between ingredients, processes and perception—even under sparse and noisy data—and unlock a new generation of food formulations?

Computational gastronomy has long shown that recipes and flavour chemistry form structured networks that we can mine for

predictive patterns. The Flavor Network formalized ingredient compatibility via shared volatiles⁴², while FlavorGraph integrated one million recipes with ~1,500 flavour molecules to generate graph embeddings for pairing prediction and ingredient clustering⁴⁰. These studies have successfully shown that high-dimensional, multimodal food representations can capture the complexity of culinary design. In alternative proteins, industrial disclosures from NotCo patents^{36–38} illustrate hybrid generative optimization pipelines. These systems embed ingredients and recipes in latent spaces using autoencoders and recurrent networks, condition sampling on feature vectors of target product profiles, refine candidates with constraint solvers such as sparse regression and iteratively incorporate expert sensory feedback. This architecture mirrors multimodal generative trends in other fields and enables analogue discovery. It can identify non-obvious ingredient substitutions while maintaining process feasibility. Other companies, including Climax Foods and Eat Just, deploy similar proprietary approaches, while academic research continues to build on Recipe1M³⁹, FlavorGraph⁴⁰ and FlavorDB⁴¹. Together, these efforts highlight a state of the art that blends representation learning with constraint-aware generation to accelerate formulation discovery.

Generative AI expands opportunities across protein diversification. For example, it can design plant-based cheeses that melt and stretch by identifying proteins that act like casein, the key driver of structure and meltability in animal cheese. It can discover ingredient blends that replicate meat textures and pinpoint volatile combinations that evoke familiar flavours⁴³. It can also improve plant-based meats from sensory panel feedback and match expert predictions of panel rankings⁴⁴. Large-scale screening lets scientists focus on a few promising prototypes instead of hundreds of trials, which shortens the development cycle. In the near future, these tools will be able to tailor formulations to individual nutritional needs and taste preferences and enable personalized food design.

But important challenges remain before AI can transform formulation discovery at scale. High-quality sensory datasets are limited and often fail to generalize across different food types. Even with detailed chemical and nutritional data, predicting human perception remains an open problem. Modelling of ingredient performance under varying heat, shear or storage conditions is incomplete. Cost, supply chain, allergen control and clean-label requirements must be built into models as hard constraints. Trust among regulators, industry and consumers will depend on transparent, auditable design processes. The near-term path will be hybrid, where AI proposes and humans refine, towards transparent formulation platforms that integrate constraints and predict performance at scale.

Fermentation

Molecular and formulation design define the building blocks, but realizing performance requires precise control at the cellular and microbial scale. Precision fermentation—where engineered microorganisms are programmed to produce functional proteins such as casein or haem—has become central to protein diversification and sustainable food production. Its success depends on tightly coupled variables, strain stability, metabolic fluxes, nutrient composition and bioreactor dynamics. AI has the potential to fuse these heterogeneous data with mechanistic insight to enable predictive, scalable and cost-effective fermentation and production to address critical bottlenecks such as yield optimization and resource efficiency (Fig. 3c).

Optimizing fermentation requires a deep understanding of bio-processes to maximize yield and efficiency while maintaining homogeneity, sterility and scalability⁴⁵. Kinetic bioprocess models remain widely used but are constrained by steady-state assumptions, and metabolic flux models for engineered strains often face computational bottlenecks. Experimental optimization through orthogonal design⁴⁶ and response surface methodology⁴⁷ remains the industry standard, while multiscale feature engineering is still limited⁴⁸. Machine learning has




Fig. 3 | AI for food. a, AI for formulations. An iterative workflow to discover new formulations encodes target profiles such as sensory properties, cost and nutritional profiles; generates formulation sets with deep learning; predicts properties and scores; and integrates laboratory prototyping and sensory panels in continuous feedback loops. **b**, AI for texture. Machine learning integrates multimodal data, from ingredients to process conditions, to optimize plant-based meat analogues, from microstructural appearance to macrostructural texture. Future machine learning tools could help to reverse-engineer ingredients and process conditions to achieve a desired appearance and texture. **c**, AI for fermentation and production. AI and LLMs can accelerate

the fermentation and production cycle during microbial strain selection, data collection, process optimization, real-time monitoring and adaptive control, upscaling of prediction and simulation, and product profiling and customization. **d**, AI for sensory properties. Optimizing the sensory perception of food is traditionally a costly, time-consuming process involving expert tasting panels and instrumental analysis, including mass spectrometry and gas chromatography. By contrast, an AI-powered pipeline has the potential to predict sensory perceptions from only the chemical or molecular make-up of food, eliminating the need for costly human panels and reducing the time to iterate product formulations.

proved valuable for modelling nonlinear fermentation dynamics⁴⁹ and for real-time monitoring of anaerobic digestion⁵⁰. For example, deep learning soft sensors that combine convolutional neural networks with bidirectional long short-term memory architectures have achieved R^2 values up to 0.98 in predicting key anaerobic digestion state variables from real-time pH, electrical conductivity and oxidation–reduction potential measurements⁵⁰. While AI and reinforcement learning have been successfully applied to sustainable protein fermentation⁵¹ and co-culture optimization⁵², the large-scale adoption of AI, and especially generative AI, remains limited.

AI can shift fermentation and production from reactive optimization to predictive management, with customized strains and self-optimizing processes. Machine learning models can optimize strain design, nutrient feeds and bioreactor control, and deep learning approaches can capture nonlinear interactions better than traditional regression⁵³. Automation and AI-enhanced imaging⁵⁴, soft sensors⁵⁵ and spectroscopy⁵¹ can improve monitoring and adaptive control. Algorithms such as support vector machines⁵⁶ and fuzzy inference systems⁵⁷ support robust predictions, while evolutionary optimization⁵⁸ aids process tuning. Hybrid mechanistic–AI models⁵⁹ and explainable AI frameworks⁶⁰ can improve transparency and decision-making. Training large language models (LLMs) on unstructured bioprocess data—including text, images and sensor logs—could support knowledge

retrieval, troubleshooting and experimental design and bridge biology, chemistry, engineering and design through actionable insights, such as predicting protein structures or cell morphology⁶¹.

To realize this potential, the field must overcome the key barriers of data scarcity, data integration and model transferability. Generating large experimental datasets remains labour intensive, but scaled-down bioreactors and small-data algorithms now offer new possibilities. At the same time, robust biosecurity and ethical frameworks will be critical to safely use generative AI. Integration with genome-scale metabolic models and digital twins of bioreactor performance could create predictive platforms for design and scale-up. The collective goal is an integrated fermentation platform that links strain design, bioreactor dynamics and real-time sensing to support predictive scale-up and reliable process optimization.

Texture

Texture is a key quality attribute that shapes consumer acceptance. It emerges from coupled physical, mechanical and sensory properties across scales—from molecular structure to oral processing—which complicates quantification and prediction. We distinguish between physical texture, measured through mechanical and rheological tests, and sensory texture, perceived through attributes such as chewiness or fibrousness. This distinction links measurable material behaviour

to human perception and defines how AI can connect mechanics with sensory response. Current approaches rely on costly sensory panels and labour-intensive tests such as cutting, tension, compression, shear and rheology. The core challenge is to map intuitive descriptors such as crispiness or chewiness onto measurable parameters that guide product design (Fig. 3b). This raises the question of how AI can integrate multimodal datasets to predict and engineer consumer-preferred textures and accelerate the adoption of sustainable proteins.

Every ingredient and product possesses a characteristic texture, which we can engineer through processing techniques such as extrusion, shear cell structuring or 3D printing to mimic animal-based foods. Recent advances show the promise of machine learning for decoding texture–structure relationships. Constitutive neural networks can map mechanical signatures of plant-based meat alternatives onto interpretable and generalizable physics-based models⁶² that achieve R^2 values from 0.92 to 0.99 for commercially available plant-based meats⁶³. Autoencoders can predict sensory texture attributes from rheological data and advance data analysis beyond traditional correlation methods³⁴. Classifiers such as support vector machines and neural networks can identify snack freshness from mechanical and acoustic signals with up to 92% accuracy⁶⁴. Researchers are using machine learning to predict viscosity, gelation and foaming of plant proteins^{32,65}. Meanwhile, texture analysis from images is rapidly growing. Transfer learning with models such as VGGNet, ResNet and DenseNet allows us to predict cooking times for fish with potential extension to plant-based foods⁶⁶, while machine-learning-driven microscopic image analysis can help us to identify droplets in colloidal systems to link microstructure to texture⁶⁷. Beyond academic applications, commercial platforms already harness AI for texture optimization: for example, NotCo^{36–38} and Climax Foods successfully deploy AI to mine ingredient databases, combine plant ingredients and design formulations that replicate animal textures.

AI accelerates prototyping by predicting textural outcomes before physical trials and reduces time- and resource-intensive trial-and-error cycles. In addition, ingredient design and process optimization can balance target textures with environmental impact—for example, to reduce greenhouse gas emissions. Another important application is personalization, such as designing softer foods for older populations to improve swallowing safety. LLMs can also convert consumer feedback such as ‘too chewy’ or ‘not crispy enough’⁴⁴ into quantifiable parameters for product adjustment and bridge the gap between perception and engineering⁶⁸. To address data sparsity, LLMs can generate synthetic training data by pairing realistic texture descriptors with mechanical properties. Ideally, AI can integrate data across multiple length scales, from molecular interactions to macroscopic structure, to capture the hierarchical origins of texture. Inversely, future machine learning tools could help to reverse-engineer ingredients and process conditions to achieve a desired texture. More broadly, generative design tools, similar to those used for inorganic materials⁷, hold the potential to propose ingredient–process combinations that meet desired textural constraints.

Significant obstacles remain before AI-driven texture engineering will truly become routine. The biggest challenge is the lack of large, open datasets that couple textural, sensory and processing data. Industrial data sharing is rare, and data fusion remains technically challenging, as mechanical, rheological, visual and compositional data are inherently multimodal and often unstructured. The ultimate vision would be an AI system for inverse design—with a specific target texture as input and optimal ingredients and processes as output—that transforms texture development into a predictive, model-guided science.

Sensory properties

Ultimately, food is not a material design challenge but a sensory experience. Aroma, flavour and texture are the strongest determinants of consumer acceptance. The central challenge is to map measurable physical

and chemical properties of ingredients to the multidimensional space of human perception. Machine learning provides a framework to learn this mapping and enables both the prediction of sensory profiles and the computational design of novel ingredients with desired characteristics (Fig. 3d). But how far are current AI models from accurately predicting human sensory perception and from generating novel flavour and aroma experiences?

Traditional sensory science has relied on human panels for perceptual evaluation and on analytical instruments such as gas chromatography and mass spectrometry for chemical composition. Panels remain the gold standard but are slow and expensive, while instruments offer precision without perceptual insight. Early machine learning approaches combined biomimetic sensors, such as e-noses, with classifiers⁶⁹, but required heavy feature engineering and were limited to narrow classification tasks⁷⁰. The field has since shifted towards deep representation learning⁷¹, which learns general-purpose sensory maps directly from data. The Principal Odor Map exemplifies this trend: a graph neural network organizes molecules into a vector space where distance encodes perceptual similarity with human-level accuracy in odour labelling. On a prospective validation set of 323 novel odorants, the model’s predictions closely matched a trained panel mean, with per-molecule correlations approaching $R \approx 0.6$ in representative cases²². Researchers now use attention-based architectures to extend this mapping to molecular mixtures⁷². Algorithmic tools for predictive sensory modelling are maturing rapidly, but progress is constrained by the scarcity and cost of high-quality sensory data.

Learned sensory representations open powerful opportunities for food innovation. High-throughput virtual screening of molecular libraries can identify novel flavour compounds for plant-based meats and accelerate product development¹¹. Integrating predictive models with low-cost, miniaturized sensors enables real-time quality monitoring in manufacturing and packaging⁷³—for example, with wireless freshness indicators that reduce food waste⁷⁴. These tools also provide systematic frameworks for biomimicry to predict how modifications of ingredients or processes alter the final sensory profile. Still, a feedback loop remains essential: computational predictions must be continuously validated and refined with real-world experiments to ensure consumer relevance.

Today, the bottleneck is no longer algorithmic capacity but high-quality data acquisition⁷⁵. Deep learning models require large cross-scale datasets, yet sensory panels cannot generate data at scale to capture the vast chemical space of food⁴⁴. The path forward lies in closing the loop between modelling and experiment. Generative models such as transformers and diffusion architectures are algorithmically ready, but their success depends on new data pipelines. Developing robust, low-cost hardware and embedding it into automated platforms will be critical to realize the full potential of AI for sensory science. The ultimate vision is a self-driving sensory lab that integrates chemical synthesis, sensory characterization and machine learning to autonomously generate, predict and validate flavour, aroma and texture perception across diverse food systems.

Manufacturing

Even the best foods will fail if they cannot be produced consistently and affordably. Manufacturing convincing meat analogues is inherently complex because technical goals such as texture, flavour, appearance and shelf life interact nonlinearly with environmental and economic constraints. Relevant datasets span composition, structure, rheology, thermal behaviour and processing steps such as extrusion or fermentation, as well as telemetry, sensory and consumer data. Through digital twins, computer vision and other modes, AI is emerging as a key technology to integrate these multimodal datasets and streamline process optimization. Typical model inputs include formulation and process profiles, while outputs can range from texture, colour and flavour to cost and environmental impact. But can these models deliver reliable,

scalable and cost-effective production paths without compromising sensory quality or sustainability goals?

AI and machine learning span the manufacturing pipeline, from formulation and process design to in-line quality control. High-moisture extrusion produces fibrous structures that mimic muscle, but outcomes depend on moisture, temperature and screw configuration. Bayesian optimization explores this space more efficiently than trial-and-error methods⁷⁶. In a pilot study, it identified optimal parameters in 10 trials versus 15 for response-surface methodology and reduced validation errors from 61% to below 15%⁷⁷. Digital twins support soft sensing and control, but most applications remain at the pilot scale⁷⁸. Data fragmentation persists across proprietary logs and repositories with inconsistent metadata. To address these limitations, FAIR data principles⁷⁹ and LLM-based extraction⁸⁰ have successfully used knowledge graphs to link formulation, process, structure and sensory performance.

AI is opening multiple opportunities for advancing alternative protein manufacturing. Engineers can use machine learning and Bayesian optimization to tune high-moisture extrusion parameters to target fibrousness, yield and energy efficiency, while digital twins evaluate what-if simulations for scale-up and real-time sensing^{76,81,82}. Food scientists can use informatics pipelines to rank plant proteins and design blends at scale using rheology–functionality maps that allow ingredient discovery well beyond commodity sources^{77,79}. LLMs further act as formulation copilots, transform unstructured texts into structured knowledge graphs and suggest experiments to resolve uncertainty⁷⁹. At the same time, computer vision and hyperspectral analytics can advance in-line quality assurance by linking colour and structure to sensory proxies to enable continuous assessment during production⁸². Together, these approaches accelerate prototyping, reduce waste and bridge lab and industrial manufacturing.

Scaling these tools requires high-quality, shareable datasets that link formulations, processes, structures and sensory outcomes with standardized metadata. Proprietary and fragmented data currently hinder model transfer and benchmarking and underscore the need for domain-specific FAIR implementations, shared test beds and community baselines⁷⁸. Ingredient variability remains a crucial bottleneck, and robust hybrid models are needed to achieve reliable process control^{76,81}. Industrial adoption also demands capital investment, workforce training and validation on real production lines since most food digital twins have not been tested at scale⁸². Shared data standards and LLM-enabled knowledge bases can bridge lab and industrial scales and enable consistent, cost-efficient manufacturing at or below price parity.

Recipes

Beyond production, food preparation shapes consumer experience. Individuals and food-service operations must balance enjoyment with health, cost, convenience and allergens while advancing broader goals such as human and planetary health. LLMs offer strong potential because they draw on large recipe datasets and integrate sources such as RecipesIM, USDA FoodData Central, the HESTIA climate database and Food.com reviews. The key challenge is to embed human knowledge so that models generalize across diverse users, dietary needs and cultural traditions.

Recent research evaluates and refines LLMs for recipe generation, revision and preference prediction. Progress relies on fine-tuning, retrieval-augmented generation and integration with optimization tools. Benchmarks on RecipeIM show that off-the-shelf models produce unreliable outputs⁸³. Fine-tuning improves performance⁸⁴ and supports revision tasks⁸⁵. A generative model trained on 2,216 curated burger recipes explored 8.9×10^{43} candidate formulations and identified Pareto-efficient recipes that reduced environmental impact without compromising taste⁴³. In blinded restaurant surveys, some AI-generated burgers outperformed the Big Mac benchmark in flavour, texture and overall liking⁸⁶. Other

work demonstrates retrieval-augmented generation for recipes⁸⁷ and sustainable menu design⁸⁸. Some LLMs now match expert food scientists in preference modelling⁸⁸. Combined with optimization algorithms, these models can generate recipes that meet nutritional and sustainability targets.

The ability to generate recipes under explicit constraints creates opportunities across individual and institutional contexts. For individuals, fine-tuned LLMs could power AI health coaches that teach users how to modify meals to be more sustainable, nutritious or culturally appropriate, while respecting restrictions such as allergies or religious requirements. Chefs and students could practice recipe generation in educational applications, receiving AI-driven feedback to refine their skills⁸⁹. In institutional settings, LLMs could support food-service operations in offices, universities or hospitals, tailoring descriptions, ingredients and preparation methods under strict constraints. This is particularly valuable in programmes such as Food as Medicine, which deliver cost- and time-sensitive meals for patients with chronic conditions after hospital discharge⁹⁰. In all these scenarios, AI could accelerate constrained recipe generation, broaden access to healthy and sustainable foods and scale public health interventions that rely on tailored meal planning.

Despite this promise, important challenges remain. Off-the-shelf LLMs perform unevenly across cultures, fail to capture diverse preferences and produce numerically inconsistent or unsafe outputs. They lack molecular-level understanding and have limited knowledge of food safety. Training models from scratch requires substantial resources. These limitations introduce risks of impractical or unsafe recipes and increase environmental costs. Mitigation strategies include integrating external datasets, using open-source pre-trained models and coupling LLMs with optimization tools. Beyond technical concerns, systems must preserve user agency and support human expertise. A more human-centred approach would allow LLMs to assist with drafting and revision, while users adapt recipes to their needs, preferences and cultures.

Conclusions

AI is beginning to transform food science across the entire innovation cycle, from ingredient design, formulation development, and fermentation and production to texture analysis, sensory properties, manufacturing and recipe generation (Fig. 2). Yet, the field's true potential lies not only in optimizing today's processes but also in charting a bold trajectory for the future of food as a programmable, predictive science. Building on recent advances, we envision four major areas that define the frontier of AI for sustainable food production.

Future materials for food innovation

Treating food as a programmable biomaterial creates opportunities to design ingredients from first principles: texture, functionality, stability and sensory perception emerge from multiscale interactions among proteins, carbohydrates, lipids and processing conditions²⁰. This perspective enables the design of edible films and gels with controlled nutrient release and plant-based meat analogues with tailored mechanical and sensory performance. Advances in protein design, including AlphaFold and diffusion-based generative models⁹¹, point towards *de novo* edible proteins with tunable functionality. Early work in cellular agriculture, such as engineered fat and muscle cells with modified composition⁹², demonstrates this potential, while microbial and fungi-based fermentation offer complementary routes. Mycelium fermentation can upcycle agricultural side-streams such as brewers' spent grain or fruit pomace into protein- and fibre-rich ingredients with distinctive textures and clear links to circularity⁹³. A key opportunity lies in integrating molecular design with tunable processing, including fermentation and upcycling, to build predictive platforms that generate new food materials with targeted properties and reduced waste.




Fig. 4 | Next-generation AI systems for scientific discovery in food.

a, Multimodal AI integrates diverse data types such as text, images, graphs and molecular structures. **b**, First-principles AI reasoning models perform iterative thinking and hypothesis generation. **c**, A possible workflow applies these capabilities to edible protein design and predicts structure, mechanics

and sensory outcomes by explicitly modelling disparate concepts and their relationships. **d**, The transition from conventional retrieval-based AI to reasoning-enabled, self-reflective AI enables exploration, discovery and ultimately the generation of novel, functional foods.

Scientific machine learning for food systems

Scientific machine learning⁹⁴ hardwires physical, chemical and biological principles directly into AI architectures and provides a foundation for trustworthy and scalable food innovation. Physics-informed neural networks⁹⁵ enforce conservation laws in data-sparse settings⁹⁶ such as fermentation and extrusion. Constitutive artificial neural networks⁹⁷ embed mechanical laws to discover physically interpretable models of food texture and rheology⁶³. Sparse identification of nonlinear dynamics⁹⁸ discovers compact governing equations from experimental data and can reveal the kinetics of ingredient interactions or microbial growth. The opportunity ahead is to couple physics-based machine learning with data-driven models to achieve interpretable, generalizable predictions under sparse data and enable mechanism-aware food design.

Automation and self-driving labs

Empirical trial-and-error methods based on rheology, sensory panels and compositional analysis have long shaped food science, but they remain slow, costly and hard to scale. Generative AI, robotics and high-throughput experimentation enable self-driving labs that design, test and refine food materials within a closed-loop design–build–test–learn cycle, accelerating systematic exploration across scales. Recent

advances in automated model discovery with custom-designed constitutive neural networks show how AI can autonomously identify the best models, parameters and even experimental protocols directly from data⁹⁷—a method already successfully applied to compare plant-based and animal meats⁶³, deli products⁹⁹ and fungi-based steak⁹³. Agentic AI systems such as Sparks¹⁰⁰, BioDiscoveryAgent¹⁰¹ and the Virtual Lab¹⁰² demonstrate the capacity to generate hypotheses, perform simulations and report results without human intervention. In food, such systems could integrate genome-scale metabolic models, digital twins of bioreactors and in-line quality sensors^{61,103}. A fully automated discovery pipeline that links synthesis, characterization and modelling would substantially accelerate innovation in sustainable, personalized and scalable food production.

Deep reasoning models

AI models with reasoning capabilities will become central to self-driving lab infrastructure (Fig. 4). The emerging discipline of AI for food will benefit from rapid progress in reasoning models and agentic systems that design and execute tasks with external tools¹⁰⁴. Current generative models propose molecules and recipes, but they rely on static datasets and struggle to assess feasibility, manufacturability and emergent properties. Closing this gap requires systems that generate

protein sequences, simulate folding and mechanics and predict sensory outcomes such as flavour and texture within an interpretable, self-improving loop. Deep reasoning models define the next frontier beyond pattern recognition and instead hypothesize, simulate and generate data. They integrate protein sequences, structural mechanics, metabolic fluxes and sensory responses to predict how molecular composition translates into taste, texture and stability²⁰. They can reduce the simulation-to-reality gap by validating predictions against experiments and updating models in real time. Such approaches can guide experiment selection⁶², integrate results into finite element simulations of oral processing¹⁰⁵ and predict sensory perception from physical principles⁶³. By embedding sustainability constraints, cultural diversity and personalized health data, these systems can design foods that meet functional, nutritional and societal goals. The ultimate goal is a closed-loop, reasoning-driven platform that designs, tests and refines food materials with predictive accuracy across scales.

Limitations and risks

This Review uses sustainable proteins as a test bed for AI-enabled food systems and focuses on post-harvest product development, from ingredients to recipes. It does not cover upstream agriculture, supply chains, packaging, safety testing or retail analytics, domains that also adopt AI and merit dedicated reviews. Benchmarking remains limited by heterogeneous datasets, proprietary silos and weak metadata standards, which constrain reproducibility and global representation. AI integration also introduces technical and societal risks, including limited transferability, domain shift and bias from sparse data. In food systems, these risks carry direct safety implications: inconsistent outputs may yield unsafe formulations, opaque optimization may hide allergen or nutritional trade-offs and automated processes raise safety and biosecurity concerns. Responsible use of AI requires transparent benchmarks, physics-informed constraints, human-in-the-loop validation, regulatory alignment and inclusive data design. AI must augment expert oversight to ensure safe, scalable and sustainable innovation.

Outlook

Taken together, these four areas define a path towards a predictive and design-driven science of food. AI can link molecular composition to structure, processing and perception and integrate these levels into closed-loop systems that design, test and refine foods across scales. From programmable biomaterials and physics-informed models to self-driving laboratories and reasoning-based AI, the field now has the foundations to move from empirical optimization to mechanism-aware design. Realizing this vision will require interoperable data standards, rigorous benchmarking and platforms that integrate experiments, simulations and learning. Progress will also depend on human-centred design that embeds cultural context, nutrition and safety into AI systems. If these challenges are addressed, AI will not only accelerate food innovation but also establish a new paradigm in which food design becomes predictive, scalable and aligned with both human and planetary health.

References

- Lappé, F. M. *Diet for a Small Planet* (Ballantine Books, 1971).
- Clark, M. A. et al. Global food system emissions could preclude achieving the 1.5° and 2°C climate change targets. *Science* **370**, 705–708 (2020).
- Xu, X. et al. Global greenhouse gas emissions from animal-based foods are twice those of plant-based foods. *Nat. Food* **2**, 724–732 (2021).
This study shows that animal-based foods produce roughly twice the greenhouse gases of plant-based foods. It underscores the environmental benefits of plant and microbial proteins central to sustainable food design.
- Keesing, F. Diet for a small footprint. *Proc. Natl Acad. Sci. USA* **119**, e2204241119 (2022).
- Friedrich, B. Transforming a 12,000-year-old technology. *Nat. Food* **3**, 807–808 (2022).
- Barabasi, A. L. et al. The unmapped chemical complexity of our diet. *Nat. Food* **1**, 33–37 (2020).
This perspective reveals the vast and largely unexplored chemical diversity in everyday foods. It highlights the need for computational and AI-driven mapping of food molecules to understand nutrition, flavour and health interactions.
- Zeni, C. et al. A generative model for inorganic materials design. *Nature* **639**, 624–632 (2025).
- Al-Sarayreh, M. et al. Inverse design and AI/deep generative networks in food design: a comprehensive review. *Trends Food Sci. Technol.* **138**, 215–228 (2023).
- Datta, A. et al. Computer-aided food engineering. *Nat. Food* **3**, 894–904 (2022).
This review outlines how computation, modelling and automation can transform food engineering. It establishes a framework for integrating AI, digital twins and process optimization into food design and manufacturing.
- King, A. Four ways to power-up AI for drug discovery. *Nature* <https://doi.org/10.1038/d41586-025-00602-5> (2025).
- Kuhl, E. AI for food: accelerating and democratizing discovery and innovation. *npj Sci. Food* **9**, 82 (2025).
- Turing, A. M. Computing machinery and intelligence. *Mind* **59**, 433–460 (1950).
- Rumelhart, D. E. et al. Learning representations by back-propagating errors. *Nature* **323**, 533–536 (1986).
- Krizhevsky, A. et al. ImageNet classification with deep convolutional neural networks. *Commun. ACM* **60**, 84–90 (2017).
- Goodfellow, I. et al. Generative adversarial networks. *Adv. Neural Inf. Process. Syst.* **27**, 2672–2680 (2014).
This work presents the GAN framework for generating realistic synthetic data. It underpins AI applications in food imaging, formulation design and sensory data augmentation.
- Vaswani, A. et al. Attention is all you need. *Adv. Neural Inf. Process. Syst.* **30**, 5998–6008 (2017).
This landmark paper introduces the transformer architecture of today's large language models. It enables multimodal AI systems for recipe generation, ingredient substitution and scientific discovery.
- International Human Genome Sequencing Consortium Initial sequencing and analysis of the human genome. *Nature* **409**, 860–921 (2001).
- Gilmer, J. et al. Neural message passing for quantum chemistry. *Proc. Int. Conf. Mach. Learn.* **34**, 1263–1272 (2017).
This paper introduces message-passing neural networks for predicting molecular properties from graph structures. It lays the foundation for AI models in food chemistry and ingredient design.
- Zeevi, D. et al. Personalized nutrition by prediction of glycemic responses. *Cell* **163**, 1079–1094 (2015).
- Gordon, E. B. et al. Biomaterials in cellular agriculture and plant-based foods for the future. *Nat. Rev. Mater.* **10**, 500–518 (2025).
- Schreurs, M. et al. Predicting and improving complex beer flavor through machine learning. *Nat. Commun.* **15**, 2368 (2024).
- Lee, B.-K. et al. A principal odor map unifies diverse tasks in olfactory perception. *Science* **381**, 999–1006 (2023).
This work creates a continuous map linking molecular structure to human odour perception. It enables AI models to predict and generate new flavour and aroma profiles for food design.
- Jumper, J. et al. Highly accurate protein structure prediction with AlphaFold. *Nature* **596**, 583–589 (2021).
This breakthrough model achieves near-experimental accuracy in protein structure prediction. It inspires the design of functional food proteins and enzymes for alternative protein production.

24. Butler, K. T. et al. Machine learning for molecular and materials science. *Nature* **559**, 547–555 (2018).
This review outlines how machine learning accelerates discovery in chemistry and materials science. It provides key methodologies directly transferable to AI-based formulation and texture engineering in food.
25. Graham, A. E. & Ledesma-Amaro, R. The microbial food revolution. *Nat. Commun.* **14**, 2231 (2023).
26. Lurie-Luke, E. Alternative protein sources: science-powered startups to fuel food innovation. *Nat. Commun.* **15**, 4425 (2024).
27. Kyriakopoulou, K. et al. in *Sustainable Meat Production and Processing* (ed. Galanakis, C. M.) 103–126 (Academic Press, 2019).
28. Fasolin, L. H. et al. Emergent food proteins—towards sustainability, health and innovation. *Foods* **10**, 600 (2021).
29. Day, L. Proteins from land plants—potential resources for human nutrition and food security. *Trends Food Sci. Technol.* **32**, 25–42 (2013).
30. Liu, M.-Q. et al. Digging natural emulsifiers based on machine learning and exploration their performance for stabilizing dairy products. Preprint at SSRN <https://ssrn.com/abstract=4994781> (2025).
31. Siejak, P. et al. The prediction of pectin viscosity using machine learning based on physical characteristics—case study: aglupectin HS-MR. *Sustainability* **16**, 5877 (2024).
32. Lie-Piang, A. et al. Quantifying techno-functional properties of ingredients from multiple crops using machine learning. *Curr. Res. Food Sci.* **7**, 100601 (2023).
33. Yilmaz, M. T. et al. Explainable AI-driven evaluation of plant protein rheology using tree-based and Gaussian process machine learning models. *Ain Shams Eng. J.* **16**, 103565 (2025).
34. Kraessig, P. M. et al. Sensory-biased autoencoder enables prediction of texture perception from food rheology. *Food Res. Int.* **205**, 116007 (2025).
35. Meng, Y. et al. Advancing salt reduction technologies: AI-assisted structural design of starch-based emulsion gel systems for next-generation low sodium food formulations. *Trends Food Sci. Technol.* **264**, 105234 (2025).
36. NotCo. Latent space method of generating food formulas. US patent 10,915,818 (2021).
37. NotCo. Neural network method of generating food formulas. US patent 10,957,424 (2021).
38. NotCo. Systems and methods to mimic target food items using artificial intelligence. US patent 11,164,478 (2021).
39. Martin, J. et al. Recipe1M+: a dataset for learning cross-modal embeddings for cooking recipes and food images. *IEEE Trans. Pattern Anal. Mach. Intell.* **43**, 187–203 (2021).
40. Park, D. et al. FlavorGraph: a large-scale food-chemical graph for generating food representations and recommending food pairings. *Sci. Rep.* **11**, 931 (2021).
41. Garg, N. et al. FlavorDB: a database of flavor molecules. *Nucl. Acids Res.* **46**, D1210–D1216 (2018).
42. Ahn, Y.-Y. et al. Flavor network and the principles of food pairing. *Sci. Rep.* **1**, 196 (2011).
43. Tac, V., Gardner, C. & Kuhl, E. Generative artificial intelligence creates delicious, sustainable, and nutritious burgers. Preprint at <https://arxiv.org/abs/2602.03092> (2026).
This study uses generative AI to create burgers that outperformed the classic Big Mac in a blinded restaurant survey.
44. van den Bedem, S. D. et al. Open-source benchmarking of plant-based and animal meats. *Foods* **15**, 2112 (2026).
45. Eastham, J. L. & Leman, A. R. Precision fermentation for food proteins: ingredient innovations, bioprocess considerations, and outlook—a mini-review. *Curr. Opin. Food Sci.* **58**, 101194 (2024).
46. Feng, X. et al. Enhanced lipid production by *Chlorella pyrenoidosa* through magnetic field pretreatment of wastewater and treatment of microalgae-wastewater culture solution: magnetic field treatment modes and conditions. *Bioresour. Technol.* **206**, 123102 (2020).
47. Sharma, D. et al. Response surface methodology and artificial neural network modelling for enhancing maturity parameters during vermicomposting of floral waste. *Bioresour. Technol.* **324**, 124672 (2021).
48. Liao, X. et al. Artificial Intelligence: a solution to involution of design–build–test–learn cycle. *Curr. Opin. Biotechnol.* **75**, 102712 (2022).
49. Liu, L. et al. Modeling and optimization of microbial hyaluronic acid production by *Streptococcus zooepidemicus* using radial basis function neural network coupling quantum-behaved particle swarm optimization algorithm. *Biotechnol. Prog.* **25**, 1819–1825 (2009).
50. Jia, R. et al. Exploration of deep learning models for real-time monitoring of state and performance of anaerobic digestion with online sensors. *Bioresour. Technol.* **363**, 127908 (2022).
51. Cheng, Y. et al. Artificial intelligence technologies in bioprocess: opportunities and challenges. *Bioresour. Technol.* **369**, 128451 (2023).
52. Treloar, N. J. et al. Deep reinforcement learning for the control of microbial co-cultures in bioreactors. *PLoS Comput. Biol.* **16**, e1007783 (2020).
53. Singhal, A. et al. Pretreatment of *Leucaena leucocephala* wood by acidified glycerol: optimization, severity index and correlation analysis. *Bioresour. Technol.* **265**, 214–223 (2018).
54. Austerjost, J. et al. A machine vision approach for bioreactor foam sensing. *SLAS Technol.* **26**, 408–414 (2021).
55. Butean, A. et al. A review of artificial intelligence applications for biorefineries and bioprocessing: from data-driven processes to optimization strategies and real-time control. *Processes* **13**, 2544 (2025).
56. Alejo, L. et al. Effluent composition prediction of a two-stage anaerobic digestion process: machine learning and stoichiometry techniques. *Environ. Sci. Pollut. Res. Int.* **25**, 21149–21163 (2018).
57. Akinade, O. O. & Oyedele, L. O. Integrating construction supply chains within a circular economy: an ANFIS-based waste analytics system (A-WAS). *J. Clean. Prod.* **229**, 863–873 (2019).
58. Peng, J. et al. Time-dependent fermentation control strategies for enhancing synthesis of marine bacteriocin 1701 using artificial neural network and genetic algorithm. *Bioresour. Technol.* **138**, 345–352 (2013).
59. Sharma, D. & Singh, K. AI-enhanced bioprocess technologies: machine learning implementations from upstream to downstream operations. *World J. Microbiol. Biotechnol.* **41**, 278 (2025).
60. Holzinger, A. et al. AI for life: trends in artificial intelligence for biotechnology. *New Biotechnol.* **74**, 16–24 (2023).
61. Ghafarollahi, A. & Buehler, M. J. ProtAgents: protein discovery via large language model multi-agent collaborations combining physics and machine learning. *Digit. Discov.* **3**, 1389–1409 (2024).
62. St. Pierre, S. R. et al. Discovering the mechanics of artificial and real meat. *Comput. Methods Appl. Mech. Eng.* **415**, 116236 (2023).
63. St. Pierre, S. R. et al. The mechanical and sensory signature of plant-based and animal meat. *npj Sci. Food* **8**, 94 (2024).
This study quantitatively links mechanical texture measurements with human sensory perception in plant-based and animal meats. It provides critical benchmark data for AI models to predict texture and mouthfeel from physical properties.
64. Sanahuja, M. et al. Classification of puffed snacks freshness based on crispiness-related mechanical and acoustical properties. *J. Food Eng.* **226**, 53–64 (2018).

65. Dahl, J. F. et al. Predicting rheological parameters of food biopolymer mixtures using machine learning. *Food Hydrocoll.* **160**, 110786 (2025).
66. Zhang, Y. et al. Construction of an intelligent recognition system for the cooking doneness of deep-fried golden pompano (*Trachinotus ovatus*) based on deep learning. *Food Biosci.* **68**, 106644 (2025).
67. Saalbrink, J. et al. Quantifying microscopic droplets in colloidal systems through machine learning-based image analysis. *Food Hydrocoll.* **166**, 111301 (2025).
68. Dunne, R. A. et al. Texture profile analysis and rheology of plant-based and animal meat. *Food Res. Int.* **205**, 115876 (2025).
69. Munekata, P. E. S. et al. Applications of electronic nose, electronic eye and electronic tongue in quality, safety and shelf life of meat and meat products: a review. *Sensors* **23**, 672 (2023).
70. Karakaya, D. et al. Electronic nose and its applications: a survey. *Int. J. Autom. Comput.* **17**, 179–209 (2019).
71. Bengio, Y. et al. Representation learning: a review and new perspectives. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**, 1798–1828 (2012).
72. Tom, G. et al. From molecules to mixtures: learning representations of olfactory mixture similarity using inductive biases. Preprint at <https://arxiv.org/abs/2501.16271> (2025).
73. Istif, E. et al. Miniaturized wireless sensor enables real-time monitoring of food spoilage. *Nat. Food* **4**, 427–436 (2023).
74. Watson, N. J. et al. Intelligent sensors for sustainable food and drink manufacturing. *Sustain. Food Syst.* **5**, 642786 (2021).
75. Rombach, R. et al. High-resolution image synthesis with latent diffusion models. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition* 10684–10695 (IEEE, 2022).
76. Dinali, M. et al. Fibrous structure in plant-based meat: high-moisture extrusion factors and sensory attributes in production and storage. *Food Rev. Int.* **40**, 2940–2968 (2024).
77. Jiang, Y. et al. Plant-based protein extrusion optimization: comparison between machine learning and conventional experimental design. *Curr. Res. Food Sci.* **11**, 101157 (2025).
78. Abdurrahman, E. E. M. et al. Digital twin applications in the food industry: a review. *Front. Sustain. Food Syst.* **9**, 1538375 (2025).
79. Top, J. et al. Cultivating FAIR principles for agri-food data. *Comput. Electron. Agric.* **196**, 106909 (2022).
80. Bölücü, N. et al. An evaluation of large language models for supplementing a food extrusion dataset. *Foods* **14**, 1355 (2025).
81. Watson, N. & Rafiq, S. *NAPIC Partner Engagement Workshop Report* (Univ. Leeds, 2025); <https://doi.org/10.48785/100/320>
82. *Reducing the Price of Alternative Proteins* (Good Food Institute, 2025); <https://gfi.org/reducing-the-price-of-alternative-proteins>
83. Mohbat, F. & Mohammed, J. Z. Llava-chef: a multi-modal generative model for food recipes. In *Proc. 33rd ACM International Conference on Information and Knowledge Management* (ACM, 2024).
84. Li, P. et al. Cheffusion: multimodal foundation model integrating recipe and food image generation. In *Proc. 33rd ACM International Conference on Information and Knowledge Management* (ACM, 2024).
85. Senath, T. et al. Large language models for ingredient substitution in food recipes using supervised fine-tuning and direct preference optimization. Preprint at <https://arxiv.org/abs/2412.04922> (2024).
86. Tac, V. & Kuhl, E. Generative AI for material design: a mechanics perspective from burgers to matter. *Comp. Meth. Appl. Mech. Eng.* **61**, 119171 (2026).
87. Liu, G. et al. Retrieval augmented recipe generation. In *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)* (IEEE, 2025).
88. Thomas, A. T. et al. What can large language models do for sustainable food? *Proc. 42nd Int. Conf. Mach. Learn* **267**, 59377–59433 (PMLR, 2025).
89. Yang, D. et al. Social skill training with large language models. Preprint at <https://arxiv.org/abs/2404.04204> (2024).
90. Rosas, L. G. et al. The effectiveness of Recipe4Health: a quasi-experimental evaluation. *Am. J. Prevent. Med.* **68**, 377–390 (2025).
91. Abramson, J. et al. Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature* **630**, 493–500 (2024).
92. Yuen, J. S. K. et al. Aggregating in vitro-grown adipocytes to produce macroscale cell-cultured fat tissue with tunable lipid compositions for food applications. *eLife* **12**, e82120 (2023).
93. Vervenne, T. et al. Probing mycelium mechanics and taste: the moist and fibrous signature of fungi steak. *Acta Biomater.* **202**, 341–351 (2025).
94. Rackauckas, C. et al. Universal differential equations for scientific machine learning. Preprint at <https://arxiv.org/abs/2001.04385> (2020).
95. Raissi, M. et al. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **378**, 686–707 (2019).
96. Karniadakis, G. E. et al. Physics-informed machine learning. *Nat. Rev. Phys.* **3**, 422–440 (2021).
97. Linka, K. & Kuhl, E. A new family of constitutive artificial neural networks towards automated model discovery. *Comput. Methods Appl. Mech. Eng.* **403**, 115731 (2023).
98. Brunton, S. L. et al. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proc. Natl Acad. Sci. USA* **113**, 3932–3937 (2016).
99. St. Pierre, S. R. et al. Biaxial testing and sensory texture evaluation of plant-based and animal deli meat. *Curr. Res. Food Sci.* **10**, 101080 (2025).
100. Ghafarollahi, A. & Buehler, M. J. Sparks: multi-agent artificial intelligence model discovers protein design principles. Preprint at <https://arxiv.org/abs/2504.19017> (2025).
101. Roohani, Y. et al. BioDiscoveryAgent: an AI agent for designing genetic perturbation experiments. Preprint at <https://arxiv.org/abs/2405.17631> (2024).
102. Swanson, K. et al. The Virtual Lab of AI agents designs new SARS-CoV-2 nanobodies. *Nature* **646**, 716–723 (2025).
103. Buehler, M. J. Accelerating scientific discovery with generative knowledge extraction, graph-based representation, and multimodal intelligent graph reasoning. *Mach. Learn. Sci. Technol.* **5**, 035083 (2024).
104. Boiko, D. A. et al. Autonomous chemical research with large language models. *Nature* **624**, 570–578 (2023).
This paper demonstrates how large language models can autonomously plan, execute and interpret chemical experiments. It marks a major step towards self-driving discovery systems, directly inspiring similar AI-driven pipelines for food formulation and ingredient design.
105. Peirlinck, M. et al. On automated model discovery and a universal material subroutine for hyperelastic materials. *Comput. Methods Appl. Mech. Eng.* **418**, 116534 (2024).

Acknowledgements

We thank H. Fattaey for stimulating discussions on the future of food, R. Yow for helping design Fig. 3c and G. Atanasovski for illustrating Figs. 3 and 4.

Author contributions

B.D. and E.K. designed the layout. B.D., M.J.B., Y.C., K.G., D.J., D.L.K., R.L.-A., G.D.M., L.N., K.P., B.S.-L., M.S., S.R.S.P., I.T., A.T., N.W. and E.K. wrote and edited the paper.

Funding

D.L.K. discloses support from the USDA (grant number FA9550-23-1-0606). R.L.-A. discloses support from BBSRC (grant number BB/Y008510/1), ERC (grant number DEUSBIO-949080) and Bezos Earth Fund (grant number BCSP/IC/001). M.S. discloses support from the Novo Nordisk Foundation (grant number NNF23OC0085919). S.R.S.P. discloses support from NSF (Graduate Research Fellowship). N.W. discloses support from the UK National Alternative Protein Innovation Centre NAPIC and Innovate UK Grant (grant number BB/Z516119/1). E.K. discloses support from Food@Stanford, Stanford SDSS Accelerator, NSF (grant number CMMI Award 2320933) and ERC (grant number 101141626).

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to Ellen Kuhl.

Peer review information *Nature Food* thanks Hongwei Zhang, Abdo Hassoun and Jiakai Lu for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

© Springer Nature Limited 2026

¹Good Food Institute, Washington, DC, USA. ²Department of Civil & Environmental Engineering, Massachusetts Institute of Technology, Cambridge, MA, USA. ³Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA, USA. ⁴Singapore Institute of Food & Biotechnology Innovation, Agency for Science, Technology and Research, Singapore, Singapore. ⁵Department of Computer Science, Johns Hopkins University, Baltimore, MD, USA. ⁶Department of Computer Science, Stanford University, Stanford, CA, USA. ⁷Department of Biomedical Engineering, Tufts University, Medford, MA, USA. ⁸Department of Bioengineering, Bezos Centre for Sustainable Protein, Microbial Food Hub and Centre for Engineering Biology, Imperial College London, London, UK. ⁹NotCo, San Francisco, CA, USA. ¹⁰Department of Chemical Engineering and Applied Chemistry and Vector Institute for Artificial Intelligence, University of Toronto, Toronto, Ontario, Canada. ¹¹Department of Green Technology, University of Southern Denmark, Odense, Denmark. ¹²Department of Mechanical Engineering, Stanford University, Stanford, CA, USA. ¹³USDA/NIFA AI Institute for Next-Generation Food Systems, University of California, Davis, Davis, CA, USA. ¹⁴School of Food Science and Nutrition, University of Leeds and National Alternative Protein Innovation Centre NAPIC, Leeds, UK. ✉e-mail: ekuhl@stanford.edu