https://doi.org/10.1038/s41590-025-02415-0

Immunological knowledge

Adelaide Gelineau, Brian D. Brown, Edward J. Hall, Lucy L. Wang, Ronald N. Germain & Christophe Benoist

What is the nature of the knowledge sought by immunology research? In this Comment, we discuss what we are collectively aiming to achieve, beyond accumulating packets of information, and how AI-driven revolutions in data handling and access modify this craft.

The 'immune system' is defined on a wide range of scales: the atomic scale that governs antigen binding specificity, through the molecular network scale, to cellular organization and interplay, and finally to organismal integration in tissues. Understanding how the immune system operates is therefore integrative physiology — knowing how these scales interact and synergize to maintain homeostasis and fulfill its different functions.

Epochal changes have affected scientific practice. The power of our experimental tools to analyze, modify and perturb the immune system is astounding compared to what was recently possible, resulting in a dramatic information overload and our inability as individual humans to keep abreast of data and publications. This firehose is accompanied by an unprecedented access to information. The emergence of artificial intelligence (AI) tools, however incompletely reliable they may be today, promises to radically alter how we generate and access knowledge in the future[1,2].

Here we examine the dual explosions of big data and AI, and how they affect the creation of, reporting of and access to knowledge. We place this analysis in the context of what 'scientific knowledge' actually means (Fig. 1). We distinguish field knowledge, which encompasses everything known today by scientific enterprise, stored in books or databases, from experiential knowledge, the individualized slice of knowledge that each of us has stored in our memory, based on readings and encounters. Experiential knowledge is incomplete and fluid, but it supports inferences and experimental design that allow us to function as scientists. Knowledge can also be distinguished along a second axis: factual knowledge is the unit level of information while conceptual knowledge connects and integrates strings of facts into processes, function or causality.

Producing scientific knowledge

Experimental production of data, facts and concepts is at the core of knowledge generation in immunology. It is arguably what makes science worthwhile. The adrenaline rush of discovery or the aesthetic of an experimental design or algorithm drive our motivation as creative scientists. Classical philosophy of science theorized experiments around hypothesis testing[3], but it is debatable whether these philosophical frameworks really capture the scientific process. Francois Jacob distinguished *science de jour*, where experiments play out logically around hypothesis testing, from *science de nuit*, a mix of lucky mistakes and vague hunches that turn out to be correct[4]. The former is the stuff of publications and grant applications; the latter is how science really works. There has always been more to experimental strategy than mere hypothesis testing. Valid designs such as "What is X made of?" (biochemical dissection) or "What happens if?" (perturbation testing) have long sidestepped hypothesis-driven research.

How do these notions apply to knowledge creation in the age of big data? Many big data experiments are hypothesis-free 'landscape' studies generating charts of expressed genomes, cells or tissues. These can have great foundational value, but atlas projects can quickly become mind-numbing. Big data techniques are amenable to elegant experimental designs; however, it may be our limitations that restrict their application to the purely descriptive.

Advances in biology do not occur by sequential addition of interlocking pieces of a flat puzzle, but through a multidimensional and dynamic process that requires novel approaches to data integration, and coexistence of multiple explanatory narratives. There should be rich opportunities for AI here[5] (Fig. 2). Its flexibility should help sidestep the requirement for data harmonization, a heavy task that has long been the bane of human-curated data integration. Foundation models that include different types of data (imaging, transcriptome, immune receptors or combinations thereof), in ways that support free-form interaction, could be instrumental. In immunology, AI strategies are already applied to questions whose complexity cannot be fathomed by human intelligence: cell identification and annotation in single-cell RNA sequencing datasets[6], sequence-to-function models that identify DNA sequence motifs that explain specific gene expression[7], predictions of B cell receptor (BCR) or T cell receptor (TCR) specificity from sequence[8], and prediction of 3D protein structures[9]. A major goal for AI in concept-building is to arrive not merely at relationship cartoons, but at quantitative models that are testable experimentally. For AI to fulfill that function, it needs to be interpretable, such that the variables and connections being considered by the model can be extracted. Therein lies our knowledge.

Conceptual knowledge is also built through literature reviews that perform data integration, assessing old and new concepts and offering alternatives. But these can suffer from authors' bias and from the 'echo chamber' phenomenon — the unquestioning propagation of statements from one review to the next. Condensing knowledge in review articles is an area ripe for AI, and agents are being developed to generate literature reviews[10,11], albeit to mixed reception[1,12]. Beyond the common concerns about falsehood in AI constructs (hallucinations, inaccuracies, manipulative AI, risk of scientific monoculture[1]), it will be essential to manage error propagation and reinforcement bias that might run amok with large language model (LLM)-based approaches. The diversity of viewpoints in human-authored reviews should help, and hybrid 'AI-assisted humans' are likely to be the preferred solution.

The creation of experiential knowledge is in a different realm altogether. It evolves haphazardly over a lifetime, from basic schooling through our experience as practicing scientists: reading of the literature, seminars and meetings, and chance coffee machine conversations. Much of this will likely remain, but one can imagine harnessing AI tools to funnel information in user-optimized ways.

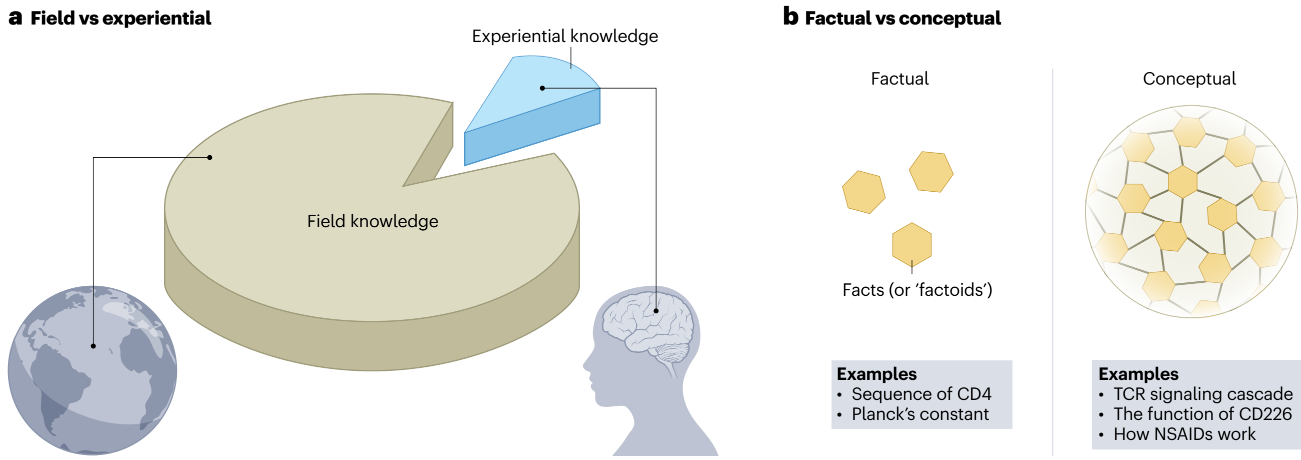

Fig. 1 | Different kinds of knowledge. a, Field versus experiential. We distinguish two broad categories of knowledge. Field knowledge (taupe) encompasses everything known today from the scientific enterprise, stored in books or databases. Field knowledge in immunology has increased explosively over the last decades, dwarfing human capacity, at an increasing rate of expansion. Experiments providing massive numbers of data points whose interpretations are challenging, and they often end up unused. Experiential knowledge (blue) is the individualized slice of this knowledge that each of us has stored in our memory, based on observations, readings and encounters. It accounts for only a small and variable proportion of field knowledge, which is probably desirable: like Borges' character in '*Funes the Memorious*', a scientist remembering every detail perfectly would be paralyzed and unable to think or experiment. Experiential knowledge is essential for a scientist to make inferences, formulate hypotheses, apply intuition and engage in lateral thinking. It is fluid, evolving with daily input that replaces fading memories. It is incomplete, subjective, fallible and thus profoundly human. Advances in scientific understanding lie at the intersection between field and experiential knowledge. While building on the broad corpus of field knowledge, it is the experimenter's experiential knowledge that provides the key intuitive spark to drive advances. **b**, Factual versus conceptual. Factual knowledge is the atomic level of information (for example, the sequence of CD4, or Planck's constant), sometimes referred to as 'factoids'. It is expected to be mostly stable over time. Conceptual knowledge connects and integrates strings of facts and deals with processes, function and causality (for example, the TCR signaling cascade, the function of the CD226 receptor or how non-steroidal anti-inflammatory drugs (NSAIDs) work). At their best, these concepts encompass the emergent properties of connected systems in non-intuitive ways that transcend smaller-scale drawings such as quantum mechanics. Conceptual knowledge is really a model that we project onto reality, and hence is subjective and can at times become disconnected from the reality it is meant to represent, akin to Baudrillard's simulacra. But conceptual knowledge has predictive power: it is the knowledge that can fly planes or cure disease.

Reporting new knowledge

How is new knowledge reported? In biology, the formulaic scientific article has long been the cornerstone. By tracing the sequence of the authors' thoughts and explorations, the article inherits aspects of the age-old tradition of story-telling, hoping to educate but also to entertain the audience with elegant experimental designs, logical progression and eye-popping discoveries. The article distills an experimental trek down to its key elements (one usually does not want all the trains of thought, trials and errors that occurred). The article also enables critical evaluation by reviewers, readers, and lab journal-clubs. The published article is accompanied by uploading its main datasets into databases and repositories.

Unfortunately, articles have become extraordinarily bloated, sometimes over 100 pages. The entertainment value has correspondingly plummeted, with the drudgery of having to scroll through pages of loathsome PDFs to reach the buried supplementary figure 16, for example. Landscape and atlas papers represent the extreme of this evolution, where the data inundation obscures true conceptual learnings, with paltry generation of experiential knowledge. Faster data generation has created expectations of papers that combine a diversity of complementary approaches, elegantly in some cases, but often just to check a requisite number of boxes.

In parallel to overweight articles, Twitter and Bluesky-type publishing has appeared: short announcements that highlight a finding or advertise an article. Does this add to information overload? Or does it offer valuable opportunities such as speed, a broad dissemination of key lessons, and a portent of future parallel publishing? 'Blog publishing' on Substack or the like has more heft, but also lacks pier review bulwarks. Other innovative forms of scientific reporting are under experimentation[13]. 'Living reviews' could support evolving conceptual integration. Micropublications might report a continuing line of investigation in a lab or consortium, as an addition to intermittent articles, providing more flexible nuggets of experiential knowledge and alleviating the dreadful monster paper that takes four years of the first author's life.

So, does the article still work in the age of big data and AI? In theory, contributing to field knowledge does not require publication since uploading new data and concepts into accessible data structures should suffice. Arguably, AI agents could better ingest new results into a corpus of knowledge. But discussing "What is the article of the future?" may be the more useful question. Can we consider a future when integrative AI handles the flow of data into foundation models that hold all field knowledge, while publications revert to a data-light, storytelling purpose of driving experiential knowledge? Such articles would be built around a small and inspiring core of conceptual novelty, the underlying data being pulled up dynamically as the reader progresses.

Accessing knowledge

Accessing knowledge means dealing with the field knowledge accumulated in the literature and databases since the beginning of immunology. Retrieval from memory is innately tied to knowledge, and in fact

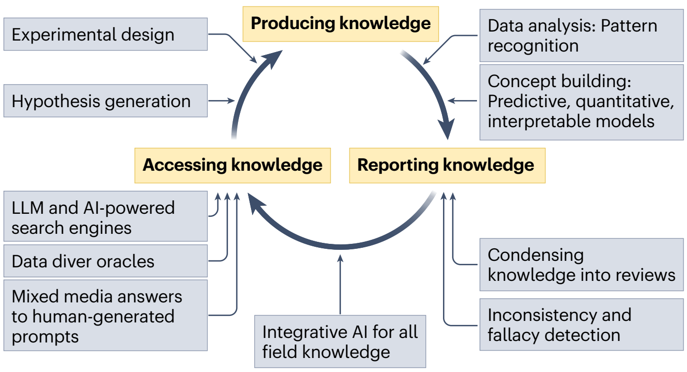

Fig. 2 | AI at different points of the knowledge cycle. Gray boxes represent different interventions of AI agents in the cycle of discovery and access to knowledge. Some of these exist, some are aspirational.

defines it. For field knowledge, paper-based approaches have almost completely given way to electronic solutions. Access through search engines was unwieldy, returning long series of marginally related articles for scientists to wade through. This is an area where the LLM revolution is making high-impact changes, with agents that combine generative AI with search functions, grounding the LLM output in reality while enabling refined queries. There is a blossoming industry of AI engines tailored for the scientific literature, which generate structured reports (for example, Asta or Scite). A diversity of viewpoints may be important to maintain for AI literature-retrieving agents.

A variety of online databases hold vast stores of non-text data, such as sequences and transcriptomes. This is another area of scientific knowledge handling that could benefit from AI breakthroughs. Many such databases only serve links to datasets but do not show the actual data, and even less knowledge. Yet we most often need a piece of information quickly and simply, not a large zipped file where that information might perhaps reside. In addition, we are not even remotely close to exploiting the full potential of these data because the value of what is in each dataset is unknown, and not queryable.

More futuristic are oracles that can dive into disparate data lakes and return integrated answers to human-formulated prompts. The returns should include the direct answer to the query, but ideally also knowledge that the human might not have known was there – in essence, addressing the 'unknown unknown' and returning new insights. This agent should not be a universal retriever, but adapted to the granularity and focus of each users' needs. Ideally, answers would also flag possible deviations or reservations about the conclusions by identifying confounders or by automatically detecting issues learned from prior training – for example, knowing that lipopolysaccharide contamination may color certain results. One can imagine integrative AI able to merge text and data and return mixed-media (text plus image)-based answers that transform the data and wrap a text 'legend' around a data graph.

Finally, and perhaps most importantly, might AI solutions connect field knowledge to experiential knowledge? Beyond *Brave New World* scenarios where we learn in our sleep, might one imagine virtual reality solutions where accessing stored data is so fast that it can partake seamlessly in our thought processes?

Dysknowledge: errors, falsehood, fraud and disinformation

The discussion above operates in an idealized context, where all data are valid and all texts true. This is far from reality, particularly in this age of disinformation, when policy rooted in disinformation has reached frightening levels.

There is a positive place for errors and misinterpretation in science. Some errors are in fact useful approximations: Newtonian physics, although fundamentally wrong, provides a valid framework for biology. Others can drive discoveries. For instance, penicillin was essentially discovered through sloppy microbiology and natural killer T cells when realizing that $CD4^+$ T cells are not all major histocompatability complex (MHC) class II-restricted. Misinterpretations are inherent to scientific debate, with most important concepts only solidifying after iterations of erroneous models. Adaptive immunity, for example, went through beautiful but completely wrong models, starting with Ehrlich, before Burnet got clonal selection right. But many errors – or worse, fabrications – are not so benign, and can have lasting consequences for a field, such as the suppressor factors of the late 1970s[14]. At the extreme, scientific errors have catastrophic humanitarian consequences when they enter societal decision-making, examples being Lysenko genetics in the Soviet Union of the 1930s and HIV denialism in South Africa of the 2000s.

These are issues that AI approaches do not deal with easily: LLMs training ingests fantasy as well as fiction and cannot distinguish a reasonable hypothesis from a flawed proposition. Reinforcement bias might propagate falsehoods very efficiently. Sycophantic LLM elaborates dangerously around a user's delusions[15] and would similarly support flaky pseudoscience. Such issues reflect the general debate around AI's dangers. Conversely, AI tools can be trained to recognize plagiarism, image doctoring or LLM-derived text. Retracted studies could be flagged to be ignored, with a gray area for irreproducible studies that are debunked but not formally retracted. More sophisticated natural language processing agents might detect logical fallacies in a scientific text, or compare results to a previous knowledge base, albeit with the clear danger of marking a true discovery as erroneous. Hence, AI systems for research need to be tailored to rigorous epistemic constraints, to distinguish valid paradigm shifts or scientifically productive falsehood from disingenuous content.

This Comment is not a white paper aimed at specific recommendations, but a perspective aimed at provoking thoughts about the changing mission of scientists as knowledge creators and inadequacies of the present frameworks. There is no doubt that an exciting world lies ahead for immunology.

Adelaide Gelineau[1], Brian D. Brown[2], Edward J. Hall[3], Lucy L. Wang[4,5], Ronald N. Germain[6] & Christophe Benoist[1]
[1]Department of Immunology, Harvard Medical School, Boston, MA, USA. [2]Marc and Jennifer Lipschultz Precision Immunology Institute and Department of Immunology and Immunotherapy, Icahn School of Medicine at Mount Sinai, New York, NY, USA. [3]Department of Philosophy, Harvard University, Cambridge, MA, USA. [4]Information School, University of Washington, Seattle, WA, USA. [5]Allen Institute for AI, Seattle, WA, USA. [6]Lymphocyte Biology Section, Laboratory of Immune System Biology, National Institute of Allergy and Infectious Diseases, NIH, Bethesda, MD, USA.
e-mail: cbdm@hms.harvard.edu

Published online: 05 February 2026

References

1. Messeri, L. & Crockett, M. J. *Nature* **627**, 49–58 (2024).
2. Bentley, S. V. *Humanit. Soc. Sci. Commun.* **12**, 409 (2025).
3. Popper, K. R. *The Logic of Scientific Discovery* (Hutchinson, 1959).

4. Jacob, F. *La Statue Interieure* (Odile Jacob, 1987).
5. Wang, H. et al. *Nature* **620**, 47–60 (2023).
6. Lotfollahi, M. et al. *Nat. Biotechnol.* **40**, 121–130 (2022).
7. Novakovsky, G., Dexter, N., Libbrecht, M. W., Wasserman, W. W. & Mostafavi, S. *Nat. Rev. Genet.* **24**, 125–137 (2023).
8. Weber, A., Pélissier, A. & Rodríguez Martínez, M. *Immunoinformatics (Amst.)* **15**, 100040 (2024).
9. Jumper, J. et al. *Nature* **596**, 583–589 (2021).
10. Tóth, B., Berek, L., Gulácsi, L., Péntek, M. & Zrubka, Z. *Syst. Rev.* **13**, 174 (2024).
11. Skarlinski, M. D. et al. Preprint at https://doi.org/10.48550/arXiv.2409.13740 (2024).
12. Pearson, H. *Nature* **635**, 276–278 (2024).
13. Maxwell, J. W. et al. *Mind the Gap: A Landscape Analysis of Open Source Publishing Tools and Platforms* (MIT Press, 2019).
14. Germain, R. N. *Immunology* **123**, 20–27 (2008).
15. Malmqvist, L. Preprint at https://doi.org/10.48550/arXiv.2411.15287 (2024).

Acknowledgements
We thank C. Bargmann, E. Lander, B. Malissen, P. Marrack, D. Mathis, R. Medzhitov, S. Mostafavi and K. Murphy for enlightening discussions.

Competing interests
The authors declare no competing interests.