# Learning Curves in Orbit: Progress with AI in Space Science

Michael T. Stuart and Sabine Winters

Forthcoming in André Curtis-Trudel, Darrell Rowbottom, and David Barack (eds.). *The Role of Artificial Intelligence in Science: Methodological and Epistemological Studies*. London: Routledge

## Abstract

AI methods are being touted as a powerful new source of scientific progress. Are they? If so, what kind of progress do they facilitate? To find out, we employed qualitative research methods to explore how space scientists conceive of AI. We show that space scientists are mainly concerned with whether AI can help them solve specific problems, and more generally, to extend their abilities in useful ways. Inspired by our qualitative data, we propose a new account according to which (at least one type of) scientific progress is improving scientific abilities. We differentiate this view from others, address some objections, and show how it flexibly integrates insights from existing work.

## 1. Introduction

Views on scientific progress may diverge over any of the following issues:[1]

(i) who or what makes progress,
(ii) what progress consists in the increase of,
(iii) what counts as an "increase,"
(iv) what the "bearer" or "vehicle" of progress is,
(v) what the right scale of analysis is,
(vi) what kinds of progress there are, and
(vii) how progress is related to other issues in philosophy of science.

For example, concerning (i), we might want to know whether it is the individual scientist, the community of scientists, an entire scientific discipline, or the human species as a whole that makes progress. Concerning (ii), we might want to know whether progress consists in increasing truth, knowledge, understanding, problem-solving power, or something else. Concerning (iii), we might want to know whether "increase" is best understood as being related to, e.g., the *degree* of truthlikeness or justification, the *accuracy* of representations, or the *number* and *significance* of problems that scientists are able to solve. Concerning (iv),

[1] For recent discussions about how we should think about the debate on scientific progress, see, e.g., Dellsén (2023, 2025), Dellsén et al. (2022), and Rowbottom (2023).

we might want to know whether progress is primarily evinced in changes to theories, models, concepts, abilities, or something else. Concerning (v), we might have different views about how to delineate "episodes" of scientific progress, i.e., when an episode of progress begins and ends. Concerning (vi), we might want to distinguish between things which *are* progress and things which *lead to* progress. Concerning (vii), we might want to know how progress relates to, e.g., the aims of science, arguments for and against scientific realism, and the debate about values in science, among other things.

Some of the above-mentioned issues are closely connected, others less so. A position on (ii) will strongly constrain positions on (iii), but a position on (i) need not constrain positions on (ii). Accounts of scientific progress are expected to address several of (i)-(vii), but not all contemporary accounts address all of them. This is fine, because accounts that only disagree about, e.g., (ii) might nevertheless agree about the rest.

The current debate on scientific progress revolves around four main accounts. The first is the "semantic" account (Bird 2007). The best developed version of this account argues that we should pay attention theory change when evaluating science for progress. And the kind of change that is important is getting closer to the truth. The main proponents of this view are Popper ([1963] 2002, 1972) and Niiniluoto (1987, 2017, 2014). On Niiniluoto's version, considering possible worlds that are close to the actual world, theory A is closer to the truth compared to theory B if theory A is true in more of those worlds, or closer to the truth in more of those worlds. As Norton et al. (forthcoming) put it, the semantic account sees progress as focusing on the balance between the degree of approximation to the truth and informativeness. They give the following example. One theory claims there are 9 billion people currently living on earth. At the time of writing there are only 8 billion, so this is, strictly speaking, false. A second theory claims that there are less than 100 billion people. This is true. The former theory is nevertheless more truthlike as it more closely approximates the actual world and picks out many possible worlds that are closer to the actual world, while the latter theory is less truthlike because, even though it is literally true, it picks out very many worlds which are quite far from ours. According to Bird, this is the "least demanding" of the four accounts of progress (Bird 2022b).

Another account of scientific progress is the "epistemic" account. On this account, science makes progress when knowledge is increased. Knowledge is a cognitive epistemic state, so instead of theories being the primary vehicle of progress, the epistemic account measures progress wherever we find knowledge (e.g., in the minds of individual scientists, or in groups, or perhaps wherever data is stored – see Birch 2025). The main proponent of this account is Bird (2007, 2022a, 2022b). A central motivating intuition for this account is that arriving at a true claim using unreliable methods should not count as progress. In other words, what the semantic account misses is justification, and some defence against epistemic luck. "Progress" that lacks justification, or gets things right merely as a matter of luck, is not real progress. This account is debatably more demanding than the semantic account, as it adds other necessary conditions associated with knowledge (e.g., justification and an anti-luck condition).[2]

[2] Niiniluoto might disagree that the epistemic account is more demanding for including justification as he has argued that justification is "built in" to the semantic account (2014). Rowbottom has argued that justification is merely of instrumental value towards truth, and so, again, its inclusion in the epistemic account might not make it more demanding. For other ways of developing the semantic account, see Rowbottom (2008, 2010) and Mizrahi (2013).

A third account of scientific progress is the "noetic" account. On this account, science makes progress when understanding is increased. Understanding is a cognitive epistemic state, so progress will be a property of changes to the cognitive-epistemic states of scientists. On the most developed version of this account, science makes progress when scientists grasp models that accurately and comprehensively represent the dependencies in a system. Progress can be increased via a *better* grasp of a comprehensive and accurate dependency model, or equal grasp of a *more comprehensive*, or *more accurate*, dependency model (see Dellsén 2016, 2018, 2020, 2022, 2023; see also Norton et al. forthcoming; Dellsén, Lawler, and Norton 2022; Potochnik 2017). Like the epistemic account, the noetic account is potentially more demanding than the semantic account, and depending on how we conceive of the relationship between knowledge and understanding, this account may or may not be more demanding than the epistemic account (Bird 2022b). The three above accounts tend to be "factive" in the sense that truth is required by each of them as a necessary component for progress to take place. But truth can play various roles. This will be even more pronounced when we turn our attention below to other versions of the accounts, including non-factive ones.

Finally, there are "functional" accounts of scientific progress. The basic idea here is that progress concerns increasing the efficiency or effectiveness of carrying out some particular function. Traditionally, that function has been thought of as problem-solving (Kuhn 1962; Laudan 1978). On this specification of the account, a scientific community makes progress when it gets better at solving problems, and this is measured in terms of changes in the number of significant unsolved problems: Progress takes place when a community solves a significant problem or "downgrades" an unsolved problem from significant to insignificant. Kuhn and Laudan are explicit that progress of this kind is independent of truth and propositional knowledge (see, e.g., Laudan 1978, 22, 126ff).

Until recently, this view has been "taken for granted as indefensible" (Shan 2019, 740).[3] Shan identifies several concerns. First, it has a "skeptical" and "antirealist" flavour, as it consciously starts from a position "internal" to scientific practice, such that what counts as a significant problem is relative to a research tradition (Niiniluoto 2014; Dellsén et al. 2022). This is supposed to lean antirealist, because scientific realism, for some, seems to involve a commitment to science as a generally progressive enterprise (but see Rowbottom 2019b), yet when a paradigm changes, the notion of significance changes, and this makes it difficult or impossible to compare progress across paradigms. If we cannot track progress, we cannot say that science makes progress, and this is a problem. A related problem is that even if we could define a cross-paradigm notion of significance, the number of significant problems solved by a brand-new paradigm will typically be very low compared to the previous paradigm. This makes paradigm shifts seem anti-progressive by definition, which is counterintuitive. Third, progress is a measure of the *number* of significant problems solved, and it has proven difficult to "count" problems (Collingwood 1946; Kleiner 1993; Rescher 1984). Finally, the functional account measures progress in terms of socially recognized solutions to significant problems. This means that erstwhile unrecognized solutions, like Mendel's work on the mechanism of

[3] Though Bird (2008) criticizes the view, which shows that its indefensibility is not entirely taken for granted, at least, not by everyone.

inheritance, do not count as progress until they are recognized by the community, and this is taken to be unintuitive (Shan 2019, 743).

A new version of the functional account has been proposed by Shan (2019, 2022), which attempts to avoid these problems while maintaining the spirit of traditional functional accounts. The initial statement of the account was as follows: "Science progresses if more useful research problems and their corresponding solutions are proposed" (Shan 2019, 744). Shan's motivations for this view are several. First, he agrees with Kuhn and Laudan that problem solving is a focal point of much scientific activity. Like Kuhn, he recognizes that for any given problem, the way it is defined and framed can change over time (e.g., in response to new experiments, data, concepts/models, techniques, instruments, etc.), and so a better definition and framing for a problem should also count as progress. Shan focuses on proposed solutions to problems that suggest general frameworks for solving a class of similar problems (2019, 746). Shan's view doesn't measure progress in terms of socially recognized solutions, or socially recognized usefulness, but simply in terms of whether problem definitions and solutions are repeatable and suggest generally reliable frameworks for solving new problems. In this way, it aims to avoid the alleged relativism of the traditional account. Despite recent work applying functional accounts to seismology (Miyake 2022) and economics (Boumans and Herfeld 2022), the functional account is still "insufficiently assessed" (Shan 2022, 2; see Rowbottom 2023, 27-30 for criticism).

In this paper, we present new evidence concerning how AI is changing one corner of science, and extract from this a new account of scientific progress. Our main claim is that our account is different from, and interacts productively with, the others. The evidence is published here for the first time and was produced using qualitative methods. The motivation for using qualitative methods will be given in the next section. Some final introductory words might still be useful here, however, to motivate our focus on a) the use of AI, and b) space science.

The reason we are focusing on AI in the context of scientific progress is that very strong claims are being made concerning the potential of AI to affect scientific progress, not just by big tech companies, but also by scientists themselves (Sourati and Evans 2023; Alvarez et al. 2024; Messeri and Crockett 2024). Optimists claim that AGI is just around the corner, and it will accelerate innovation by producing new, good ideas for how to address global issues like climate change, hunger, disease, and energy use. Pessimists worry that AI will flood the scientific marketplace with boring useless papers; change incentive structures (e.g., by motivating scientists to work on problems that AI is good for solving rather than problems that really matter), and cause important scientific skills to atrophy (Cheng and Zhang 2025). While there is growing philosophical interest in the effects of AI on science, there is nothing yet written from a philosophy of science perspective on AI and scientific progress, and this paper aims to help fill that gap. The kind of information we require concerns how AI is already affecting scientific progress, and as there is little information publicly available on this, so we decided to collect the data ourselves. One thing that is important to note at the outset is that we choose not to define progress in advance, but to extract an account from what scientists say and do. This creates a risk of erroneously taking scientists to be saying things that are actually not relevant to the philosophical debate. We will discuss the strategies we took to avoid this in sections 2 and 4.

The reason we are focusing on space science is because the achievements there are many and easy to list. Most of us can name the first satellite in orbit, the first dog in space, the first human in space, the first person on the moon, and at least one of the Mars rovers. In addition to technological achievements, the number of scientific discoveries that any given space mission is responsible for is usually in the hundreds.[4] National and international space programs like NASA keep close track of the amount and quality of the science they are responsible for, including which programs have led to more progress and why (Pirtle and Moore 2019). Further, space agencies are themselves interested in the potential of AI to accelerate progress. Joe Pellicciotti, NASA's chief engineer, claims that "some of the biggest changes, at least more recently, have been in the digital area and AI world…We've seen advances already, and that's just going to continue to grow. And as it gets bigger and bigger, and we qualify and validate more of these systems…it'll just accelerate exponentially. So, I think there's huge advancements…there's a very positive future in it" (Almeida 2025). Katherine Van Hooser, NASA's deputy chief engineer, says "some of our analytical solutions would take days to run on a supercomputer…And now those same level of solutions, and probably even better ones, are coming from, you know, somebody's laptop that they're, they're running, you know, multiple times in a day. And so, the tools have gotten a lot better, which is great because it lets us explore more problems and find, you know, multiple solutions or options for programs. That's where I think we've gotten a lot stronger…We've got to figure out how to use AI and figure out how to make more of our digital tools work together more efficiently, or else we're going to get left behind" (Almeida 2025, see also Izzo et al. 2022 and Antonsen et al. 2025). Given the intuitive connection between space science and progress, as well as growing interest in AI methods in that field, this seemed like a good case study for better understanding the effects of AI on progress, and perhaps also the nature of progress itself.

## 2. Methodology

This study employed a qualitative, interview-based approach. It is part of an on-going multi-year qualitative project. Qualitative methods allow for deep, context-rich insights into participants' experiences, values, and interpretive frameworks—factors that are often crucial in scientific decision-making but difficult to quantify in discussions of technological practice and epistemology.[5]

The core aim is to understand how practitioners interpret, justify, and navigate the practicalities of working with AI systems. These kinds of questions are not easily answered through surveys or quantitative metrics. Instead, semi-structured interviews provide the flexibility to probe underlying assumptions and the social and organizational dynamics that shape practice.

[4] Some major scientific discoveries made by space science institutions like NASA and ESA concern the existence of Earth's radiation belts, the structure of the sun, the planets and their moons, the existence and features of black holes, steadily burning cool flames, drug discoveries relevant for cancer, gum disease, muscular dystrophy, Alzheimer's, Parkinson's, asthma and heart disease. And a great deal of work is done on Earth's climate from space which is relevant for understanding climate change, including predicting and addressing natural disasters.

[5] The need for sociological investigation in the debate on scientific progress is nicely motivated by Rowbottom (2023, 45).

At the time of writing, twenty-five participants had been interviewed, all of whom were involved in space research within one of several different space organizations. Participants held diverse roles within the field of space science, representing various teams across space organizations and institutes. Our participant group includes 8 project managers, 7 permanent research staff, 4 postdoctoral researchers, and 6 graduate students. To preserve anonymity while allowing the reader to distinguish among perspectives, we use coded identifiers: M1, M2, etc., for project managers; R1, R2, etc., for permanent researchers; P1, P2, etc., for postdoctoral researchers; and G1, G2, etc., for graduate students. In terms of gender representation, 7 participants identified as female and 18 as male. While this distribution reflects current imbalances in the field, we acknowledge the importance of gender equity and actively strove toward a balanced representation in the research.

Participants were recruited via email invitations sent by Winters, who had previously spent time as a guest within one of the space organizations in an ethnographic capacity, and who had also developed connections with individuals from other organizations through her ongoing ethnographic fieldwork. This pre-existing relationship helped facilitate openness and trust during the interviews, which were typically conducted physically within the organizations themselves. This allowed for context-rich conversations grounded in the participants' everyday professional environments. The interview process aims to ensure a comprehensive representation of perspectives across different roles and teams within space organizations, and to investigate other topics. The long-term nature of the study helps to capture evolving insights and practices within the field.

This paper reports findings from the first twelve months of the study. Follow-up interviews (with M1, G1, G2, P3, P2), were informed by preliminary analysis, and tailored to explore particular themes in greater depth. Participants were affiliated with three distinct space research institutions. Three participants were employed as a professor at a university, one founded an independent space-related foundation, and one started a space-related start-up. The remaining 20 participants were affiliated with a single, large space organization, representing 10 different internal teams. Of these, 11 participants worked within the same team, while the other 9 were part of separate teams within the organization. At the time of writing, five participants are no longer employed by the organization at which they were interviewed, due to the scheduled completion of their research projects. This distribution reflects both the organizational diversity and the transitional nature of careers in space science research, allowing for a nuanced understanding of differences and commonalities in practices, perspectives, and institutional roles.

Each participant took part in one or several semi-structured interviews lasting between 45 and 90 minutes. Interviews were conducted either in person or via video call, depending on availability and location. All interviews were audio-recorded with participant consent and subsequently transcribed. The quotations have been revised with attention to authenticity and diction, to enhance clarity. The research project was approved by the Science-Geo Ethics Review Board of Utrecht University.

Winters conducted the interviews and did the coding and preliminary analysis. The transcript coding proceeded iteratively using a grounded theory approach: initial descriptive codes were developed based on repeated readings of the transcripts, then grouped into higher-level themes, for example, related to epistemic

positioning in relation to various cognitive processes, research-related and practical decisions, and different institutional roles. Stuart offered ongoing feedback on the emerging themes, guided the development of the coding approach, and supported the structuring of key quotations and interpretive framing throughout the analysis. This collaborative process ensured analytical rigor while retaining the close empirical grounding characteristic of qualitative research.

# 3. AI and Progress in Space Science: Results

In general, we were interested in finding out whether, when, and in what senses AI was having an effect on work in space science. Interviews revealed several areas in which AI was having an impact. These include autonomous on-board machine learning (including rover navigation and satellite control), data analysis (both for science and for policy use), medicine in space, mission design, estimating the shapes of comets and asteroids, materials science, and spacecraft design. Based on our analysis of the interviews, we present participant responses in two sections, relating to a) problem-solving and b) general capacity-building. These are not mutually exclusive categories since uses of AI that solve particular problems typically also build more general capacities on which researchers later draw, and general capacities are built because they are expected to solve problems (among other things). Still, we can roughly distinguish the two categories by the level of detail the participants present concerning the uses of AI: the first section concerns cases where the problems are known and quite specific, whereas the second concerns general, often merely anticipated, issues.

The fact that the quotations fit entirely into sections only concerning problem-solving and capacity-building is surprising, insofar as we do not find direct mention of other markers of progress that we would expect to find if the semantic, epistemic or noetic accounts were correct. We come back to this in section 4.

## 3.1 When AI is progressive: Solving specific problems

Most participants discussed AI in the context of specific problems that AI methods could be used to solve, and they criticize AI when it is introduced in the absence of such problems.

We begin with several extended quotations from an interview with M3, who is transitioning from a high-profile innovation-center into a new role within a space agency, focusing on the application of AI and software in space exploration missions. Specifically, their work involves identifying mission-critical challenges—such as rover navigation on the Moon—and determining where AI or advanced software solutions can meaningfully address those challenges. This fact makes M3 a key source of information for how AI is affecting space science.

> Yesterday I had a call with someone from a consultancy - I won't name them, who was really trying hard to sell me his AI algorithms. But he didn't understand what I actually needed them for, or what our challenges in exploration are. He came with a very technical pitch, talking about knowledge graphs and neural networks and I know what those are, but I had to ask: what problem is this solving?

Often there's a disconnect, because AI is hyped. We get lots of proposals: large language models, neural networks, but they don't really address our needs. I don't want to make a long list of algorithms just so we can say we've used them. What I'm trying to do is identify real problems in our missions, whether it's a rover, a space station, or an astronaut on the Moon without GPS, and then ask: is this a problem that needs software? Or AI? It's not always AI.

So I'm trying to separate the hype from actual use cases, because there are people who want to use AI just so they can say they've used AI, because it sounds cooler. But it's always tricky, working within a hype cycle...

I've found that the less someone understands what AI is, the more likely they are to think of it as a kind of magical solution...Sometimes, statistics and AI are quite close and for some problems, all you need is a smart Excel sheet with macros or a pivot table. That's not AI, but for someone unfamiliar with Excel, it can still feel like magic...

I don't want to demotivate people. It's great that they're excited about the technology, but it can also be dangerous if they think they understand it when they don't. What I try to do is break the problem down and ask: what exactly are you trying to use? Are we talking about machine learning, deep learning, supervised or reinforcement learning? And then, often, they realize they don't actually know...

It really depends on who you talk to, but in management, AI is often seen as something we need to adopt just because everyone else is doing it. It's like, 'Let's put AI into this, we need something with AI in it.' And then when I ask, 'What exactly did you have in mind? Is there a specific algorithm you're thinking of? What would be the advantage of using AI here?' - they usually realize they don't have the answer...I don't want to embarrass anyone, so I just lightly touch on the topic. But once I start asking those kinds of questions, they'll often admit: 'Actually, I'm not the expert. Maybe you should take a look at this...'

This is a situation everyone is dealing with, and honestly, we're probably a bit late in running into this problem. Many companies started looking into AI ten or fifteen years ago. For some, AI is already fully integrated. They use it like a battery or a chip or a laptop, just another great innovation. For them, AI is normal. Then there are others who are either scared or not informed. So we're often caught in the middle, - between enthusiasts who don't understand AI but are eager to use it, and those who are resistant or intimidated by it...

What I prefer is to focus on the problem - not talk too much about the AI, and then propose a solution, whatever it may be called. That approach probably comes from my background in university and software development. I'm part of a generation that was really pushed to think about user needs. There's just so much software out there that was built by developers who wanted

> to build something, but it doesn't meet any user requirements. It's not user-friendly. That has really influenced how I work: I always try to put the user at the center. (M3, 23-08-2024)[6]

M3 is cautious about the tendency to apply AI solutions without first adequately determining whether there is a problem it can help solve. They emphasize that this risk stems from a lack of conceptual clarity about what AI actually is. Overall, AI is justified when needed, and it is needed when it can solve a specific, existing problem. As we will see, this is not the only justification for introducing AI, but it is the one that occurs first and most powerfully to M3.

M3 uses the above metric to judge whether AI should be used, and this helps them navigate AI hype. While AI is acknowledged as potentially productive, M3 warns that excitement can be harmful if it fosters uncritical implementation that wastes resources and introduces new failure points. As a corrective, they advocate for a problem-first approach, in which the choice of technological tools, including AI, is guided by a clear understanding of the use case. In the next subsection, we present examples to concretize these ideas and show that they are to be found in different parts of space science.

### *Examples*

The following are specific examples where AI was thought to be useful in solving specific problems.

M1's team is working on spaceflight-associated neuro-ocular syndrome (SANS). This is a condition that affects astronauts during long-duration space missions — e.g., those aboard the International Space Station (ISS). There are several problems to solve concerning SANS, which M1's team has considered using AI to address.

> A lot of people have been asking: how can we detect it early, and how can we protect against it? Our team took that as a challenge and developed a solution using a mobile phone camera combined with an artificial intelligence system to detect the onset of symptoms.
>
> Essentially, it's possible to use a phone camera to image the back of the eye, the retina. We developed an AI system that analyzes those images and can identify early signs of this condition. To train the AI, we used a large, medically robust dataset from people here on Earth who have this syndrome. So the model is grounded in real clinical data.
>
> This isn't about using AI for the sake of novelty. It's about giving astronauts practical, reliable tools. They don't need AI that talks to them, they need tools that help them understand what's happening to their bodies, so they can act on it. This tool is pragmatic, targeted, and addresses a real need.

[6] Following usual social scientific conventions, we cite quotations by reference to the participant and the date of the interview.

> We've already flown it to space once, and it worked. We're aiming to fly it again in about a year, and the goal is to eventually make it an operational tool, – something astronauts can have in their standard inventory to monitor their health autonomously. (M1, 14-08-2024)

Here, M1 distances (in a nuanced manner) their project from those which employ AI for its own sake. In contrast, what they want to emphasize is the system's utility as a practical diagnostic aid rather than a technological showpiece. The goal is not to impress with AI, but to equip astronauts with tools that enable in-space health management. In this context, it is particularly important that M1's team maintains direct contact with astronauts (who are literally across the hallway) from whom they receive first-hand feedback while they are developing the tool.

M3, who we quoted from at length above, is also working on several AI applications in the team. One of these projects concerns autonomous navigation for lunar rovers.

> You don't want to get stuck in the shade, and then your battery dies, because there are fourteen days of night on the Moon. So, if you have solar panels, and you get stuck in the shade, that's very bad, because you're dead. It's critical for rovers on the Moon to navigate quickly into the sunlight. If they get stuck, it's a serious issue.
>
> That's where we need better software. For example, having a camera that can recognize the terrain, the rocks, the elevation, the sun's position and then calculate how to get from point A to point B without getting stuck. In this case, AI is part of the vision-based system for hazard avoidance. It analyzes the camera feed to automatically label terrain features and propose a route.
>
> My role isn't to build the software, but to define the requirements based on mission needs, identify a company that can develop it, and then manage the project to make sure it's delivered and integrated into the rover. (M3, 23-08-2024)

This case of using AI to navigate difficult terrain illustrates M3's method, from selecting a problem, understanding (defining and framing) the problem, identifying key points in the solution-space, and identifying the best tools (and then creating and refining them) as part of a solution to the problem.

Finally, we turn to P4, who is applying formal knowledge from causal inference to problems in space science. Specifically, they are exploring how causal inference tools can be used in applications such as improving exoplanet detection by removing instrumental noise from telescope data. In P4's work, the causal relations probed concern the connections between instrumental and environmental factors and observational noise in telescope data. For example, variations in sensor temperature, mechanical vibrations, or detector drift can create patterns in the data that obscure or mimic true astrophysical signals. The goal is to understand how these factors causally influence measurements so that real planetary or stellar signals can be distinguished from artifacts.

Machine learning, particularly causal inference methods, captures these relations by modelling the dependencies between variables explicitly, rather than simply identifying correlations. Techniques such as causal graphs or structured probabilistic models allow P4 to represent assumptions about underlying

mechanisms, simulate interventions (e.g., “what would the signal look like if parameter A were corrected?”), and estimate the effect of individual factors on observational outcomes. This is different from traditional statistical analysis, which typically focuses on correlations or aggregate patterns without accounting for causal structure.

The approach is more useful than standard statistical methods because it enables targeted noise reduction and enhanced signal detection, improving the reliability of astrophysical inferences. By separating causal influences from spurious correlations, P4’s models generate cleaner data and allow for more confident identification of faint exoplanets. In this sense, AI is valuable because it provides better solutions, not by automating scientific discovery, but by enhancing the quality of the evidence on which scientific conclusions are based.

P4 summarizes some of the other main areas of research, which, in their opinion, AI is impacting.

> In the space sector, a major focus right now is on the management of data, particularly on-board satellite data. We already have small satellites equipped with NVIDIA components for real-time data processing. The idea is that the satellite can autonomously identify which areas are most interesting and prioritize data from those zones. By processing data intelligently on board, we can make better use of the limited bandwidth available for transmitting data to the ground.
>
> Another huge area of investment is making all this Earth observation data accessible and usable for decision-makers. If you're a politician making decisions about climate, agriculture, or pollution, the key question is: which data should you be looking at? So there’s a strong push to make these data streams available and intelligible to both the scientific community and policy stakeholders.
>
> A third major initiative is what we call 'digital green', building massive machine-learning-driven simulation models of the entire Earth. These models integrate satellite data with geographical, social, and environmental dimensions. The goal is to have a system that continuously improves itself and simulates the planet as a whole. This work is also being pursued at the European Commission level, especially through projects like Copernicus, where I’m currently working in Earth observation. (P4, 13-12-2024)

P4 outlines three major areas (other than their own) where machine learning is currently being applied in the space sector. The first is autonomy in satellite data collection. Satellites can only transmit data when they have a line of sight to ground stations that can receive data, and those transmissions have limited bandwidth, so it’s important to make sure the satellites are getting and transmitting useful data. Rather than program data collection protocols in advance, or try to react in real-time to new developments, satellites can use AI to make decisions about where to focus their data collection. Second, P4 describes an institutional shift toward making Earth observation data more accessible to decision-makers. Given the volumes of satellite data generated – often measured in petabytes – the challenge lies in transforming this raw data into actionable insights for policymakers in domains such as climate policy, agriculture, and pollution control. And AI can be useful here, by making patterns of interest cognitively available to those who need them. Third, they highlight the emergence of “digital green” initiatives, which aim to build comprehensive machine-learning simulation models of Earth. These models integrate geospatial, environmental, and social data to simulate

the planet's systems dynamically. This effort aligns with broader goals at the European Commission level, particularly through projects like Copernicus, to support sustainability, governance, and public utility through AI-enhanced Earth observation. As with M1 and M3, P4's work coheres with the idea that the introduction of AI is justified, and its employment is at least potentially progressive, when it helps scientists generate *useful* data, or make raw data more *usable*, or support sustainability, governance, and public utility by addressing specific needs, in the sense of overcoming specific problems (bandwidth limitations, short signalling windows, and massive and complex datasets).

The work of M1, M3, and P4 exemplifies M3's position on the importance of AI for solving particular problems. In the cases mentioned, participants feel that the question should predate the uptake of a new tool. The team has specific goals and directives, and putting these into practice leads to challenges, and overcoming those challenges requires solving specific problems. Insofar as AI can be used to help overcome some of these problems, it can be included as a useful or progressive methodology.

Finally, it is important to note that although both M1 and M3 hold managerial roles, they serve different functions. M1 leads a team that conceptualizes and implements new developments internally within the organization, whereas M3 is more closely affiliated with an AI innovation function and liaises with industry. The function of M3 appears to align more closely with management's broader strategic agenda than with open-ended exploration. And yet each of these participants express a similar view about when AI implementation is justified, hinting at a similar view about what progress is, and how AI can help further it.

But this is not the only way we found participants speaking positively about AI.

## 3.2 When AI is progressive: Improving abilities

In this section, we will see how AI plays an increasingly supportive role in scientific and engineering work, particularly as a tool that can assist in accelerating learning, problem-solving, and creative thinking. AI tools – such as retrieval-augmented generators, computer vision models, graph neural networks, reinforcement learning, and predictive analytics – enhance the above-mentioned processes by enabling faster analysis, generating ideas, and helping scientists mentally probe complex problems. The key benefit lies not in AI offering answers to existing well-defined problems, but in improving the quality (e.g., speed and precision) of the inferential and exploratory processes that researchers already employ. We can think of these processes as related to individual physical or cognitive-epistemic capacities, which might be distributed across individuals and instruments, at the lab or institutional level (Nersessian 2022). At the individual level, AI can serve as a powerful enhancement for those who already have well-developed epistemic capacities (e.g., for analytical or imaginative reasoning). For others, it risks becoming a crutch that may diminish or prevent the development of epistemic capacities if not used reflectively. First, we focus on the positive use of AI to extend epistemic abilities.

R4 is a researcher at a space organization, working on highly interdisciplinary projects that combine numerical computing, physical computing, and physics. Their focus includes integrated computational materials engineering, such as designing and 3D-printing metal components like rocket nozzles based on computer models and adapting complex plasma simulation codes developed by external researchers for use

on various computing platforms. The work involves troubleshooting software and hardware issues, rewriting code and equations, and applying DevOps principles for software portability.

> I had a situation with NVIDIA where they came to me and said, 'I'll start something now, and in five minutes I'll tell you what it is.' After five minutes, they said, 'We've just imagined 140 new COVID-19 vaccines using our computer, without any prior knowledge.' Of course, there's a simulation behind it, but it's not physics-based. It's called Stable Diffusion, which means it just takes particles and tries to arrange them to see if the configuration looks sensible. So it's more imagination than physics, because you're moving objects into shapes that appear reasonable, and then you explore whether they actually work.
>
> So yes, this is really imagination. And that plays a role in the whole AI discussion, that you can imagine configurations without knowing whether they're realistic, and that definitely opens up new design possibilities.
>
> For myself, I would say AI enhances imagination, because I had the experience before using the tools. I'm used to imagining things without external help. But I would agree that for people who haven't developed those imaginative skills, it can dull your capabilities. You see an answer and think it's the golden goose, and then you rely on it. But that's not true. You need to challenge the computer. It's just your partner - you still need your own capability to question what it gives you. (R4, 04-10-2024)

R4 presents a balanced perspective: AI can enhance creativity for those with an existing foundation in imaginative thinking, acting as a partner in ideation. Crucially, AI enhances rather than replaces the intuitive and exploratory processes central to scientific inquiry. This introduces a second rationale for the value of AI: its potential to enhance general cognitive capacities, such as abstract modelling, pattern recognition, and adaptive reasoning, that are thought to be indispensable in scientific and technological work. A second reason AI is considered valuable is its ability to enhance general cognitive and epistemic capacities that are expected to be crucial for future scientific and technological progress. Rather than replacing human cognition, AI can accelerate learning, problem-solving, and creative exploration, especially for individuals already skilled in independent reasoning. Through tools like retrieval-augmented generation, AI supports complex tasks in science and engineering by enhancing, rather than automating, intuitive and imaginative processes.

However, R4 emphasizes that AI should be treated as a partner in thought – not a substitute for thought. Without active cognitive engagement, AI may diminish rather than enhance scientific capacities, particularly for individuals lacking prior training or over relying on a single model, such as ChatGPT. As in section 3.1, we will now turn to some examples which concretize and demonstrate the generality of these ideas.

### *Examples*

The following are examples where AI was valued for improving general scientific abilities.

P3 is currently engaged in two primary research projects that integrate principles from AI, computational neuroscience, and biomimetic modeling. The first project involves the development of a computational model of the human retina, aiming to simulate how different retinal cell layers and circuits process visual information such as motion, depth, and global features. The objective is to construct a biologically faithful representation that could eventually be implemented in hardware, contributing to more efficient and perceptually grounded machine vision systems. The second project focuses on continual learning—also referred to as lifelong or online learning—which entails developing algorithms capable of adapting themselves to new data in real time. This approach mirrors the learning processes of humans and animals and addresses a key challenge in creating adaptive, resilient AI systems.

> Part of the reason I've been working on this continued learning project for the past three years is that I think it's very relevant for space. One of the big challenges is that space is incredibly inhospitable, so we'll likely have to rely on machines to explore, to send back interesting data, and to operate without being explicitly controlled by humans.
>
> The further away these missions go, the harder it becomes to communicate, control, or debug systems, they have to function entirely on their own. Even the Moon is challenging, and Mars more so. But now there are missions planned to Jupiter's icy moons, or to Titan, and we don't really know what those environments are like.
>
> We have some observations, some sketchy data, but we don't know what the probes will encounter. Imagine a probe landing on Titan. We have no idea what that surface will actually be like. You can't pre-train a controller or an AI system on data that doesn't exist.
>
> It needs to learn on the fly. It needs to adapt after deployment and incorporate new knowledge as it encounters the unknown. That's a huge challenge, and that's why I'm working on it. (P3, 04-10-2024)

P3 here reflects on the importance of continued learning systems in the context of space exploration. As missions venture farther from Earth, the feasibility of human oversight diminishes and it becomes desirable for AI systems to be capable of operating autonomously and learning in real time. P3 emphasizes that we cannot rely solely on pre-trained models using Earth-based data; instead, onboard systems must be able to update and adapt dynamically after deployment. This capacity for autonomous, situational learning represents a critical and urgent challenge for space robotics.

This project reflects P3's broader philosophical and methodological approach to AI in terms of building useful capacities. As they explain:

> The whole argument about the efficiency of how we learn language and so on - okay, that's interesting. But most of my inspiration comes from biology. I usually start approaching a problem by thinking about how I, or someone else, or even an animal, would try to solve it. How would an animal work around it?

> I imagine the process. Sometimes I watch videos of actual animals solving actual problems. I do that often. The octopus, for example, is a fascinating animal, really underrated and understudied in terms of how it solves problems.
>
> Then I look at the models we have — maybe machine learning solutions or other AI systems, and ask: what's the difference? How are we trying to solve this from a machine learning perspective, and how does the animal actually do it?
>
> I try to imagine that difference and then think about how something like the animal's approach could be implemented in a computer or machine. So, in a nutshell, that's kind of the process. (P3, 04-10-2024)

This appeal to biological problem-solving underpins much of P3's work, suggesting a future where machines not only compute but adapt, improvise, and respond—more like animals than algorithms. P3 is interested in how animals approach problems, not because P3 is interested in the particular solutions, but because they are interested in general problem-solving strategies and abilities which might be developed for robotics.

P3 contrasts this with standard machine learning approaches, reflecting critically on the differences between biological and artificial strategies. It is not merely mimicking animal behavior, but imagining how general strategies might be abstracted, translated, and implemented into computational systems. This suggests a design logic that is grounded in analogical reasoning and embodied cognition, rather than either pure formalism or hands-off machine learning, revealing a creative, cross-domain mode of capacity-building within AI development.

Building on this perspective, P3 turns to the challenges of embodiment, emphasizing that intelligence cannot be understood or designed in isolation from the physical systems through which it perceives and acts. They highlight a core issue in AI and robotics: the limitations posed not only by algorithms, but also by the physical interfaces between machines and their environments. As P3 explains:

> Challenges typically arise even in routine, everyday contexts, and many of the things we do without thinking remain extremely difficult, if not impossible, for AI systems to replicate.
>
> The limitations extend beyond computational models to hardware, particularly to sensory systems and effectors. Our biosensors, such as vision and hearing, are only imperfectly approximated by cameras and microphones. Similarly, when it comes to interacting with the environment, nothing in robotics yet matches the dexterity of the human hand or the sophistication of the human eye. This means that, at the interface between the system and its environment, a significant amount of information is already lost or degraded.
>
> Now, imagine a hypothetical system that *did* overcome those limitations, say, an advanced android equipped with highly refined sensors and effectors. The question then becomes: what sort of controller would be needed to operate such a system? Conversely, consider the opposite scenario,

> we might have a highly capable controller, but it is connected to only basic, rudimentary sensors and actuators, especially when compared to biological systems.
>
> In practice, we never have both: we rarely, if ever, possess a highly advanced physical platform *and* a correspondingly sophisticated controller. So, in most real-world cases, we must work with imperfection on at least one side of the system. The key challenge, then, is determining how a controller can operate effectively under such constraints; how it functions when dealing with incomplete, imprecise, or low-fidelity sensory input. (P3, 04-10-2024)

These quotations demonstrate a sustained attention to a set of cognitive and physical abilities that P3 finds valuable. These abilities might be valuable for solving problems, e.g., providing a specific kind of data. But they might also be valuable even in the absence of a specific problem, e.g., for gathering and interpreting data or navigating new landscapes in an exploratory way (i.e., without there being any "problem" in the sense of there being something that has gone wrong, or something which is at odds with the agent's goals, see Barseghyan et al. 2024). Future work could clarify whether those abilities are best attributed to the machine itself, to individual humans (e.g., operators, see Vertesi 2012, 2015, 2023), or to human-robot teams (e.g., a lab or company). But in any case, it is clear that extending these abilities by developing flexible, robust, reliable, and powerful technological devices is an important goal of these scientists, the achievement of which should count as progress, in some sense.

# 4. Discussion

## 4.1 Little mention of truth, knowledge, or understanding

A notable absence from the interview data was any explicit mention of epistemic good-making features like truth, knowledge, or understanding, and only a few references to representational accuracy. Instead, the participants spoke about what AI was allowing them to do, including what problems it was enabling them to solve, and which capacities it had improved.

One explanation for this might be that we did not ask our participants what they thought of scientific progress in general. Instead, we asked scientists to define for themselves what was good or bad about recent developments, and we did not impose any restrictions on the kinds of goodness we were interested in (e.g., technological, epistemic, ethical, aesthetic, etc.). Still, we identified specifically epistemic notions of goodness in the transcripts and observations, and then asked follow-up questions to zoom in on those. This produces better data than simply asking directly about epistemic progress, since answers to that kind of question typically result in general mottos that don't necessarily reflect actual positions and practices.

It was surprising that we were able to abstract a notion of scientific "goodness" that was shared across all participants, despite large differences in their perspectives, research topics, and institutional roles. As we have seen, participants all tended to frame AI as contributing to science through capability-building rather than by producing new truths or instances of knowledge/understanding. We are characterizing "progress" in terms of movement toward that shared notion of goodness.

Still, we might want to distinguish between success and progress, where "success" might refer to the attainment of specific project-bound objectives and "progress" might refer to more cumulative, general advances.[7] While participants frequently described markers of success, such as implementing a model, achieving technical milestones, or completing project tasks, a broader notion of progress could be abstracted from their responses that transcends individual achievements. They emphasized how AI mediates research practices, opens new methodological avenues, and enables forms of inquiry that extend their epistemic reach. And importantly, much of that progress remains epistemic. For example, P4's work on applying causal inference to telescope data shows how modelling and noise reduction (important successes) contribute to the more general goal of improving the ability to detect exoplanets and interpret astrophysical signals, which seems clearly epistemic.

The kind of progress valued by the participants exemplifies Kitcher's idea of "pragmatic progress," i.e., thinking about scientific progress as progress *from* some particular state, instead of progress *to* some antecedently specified state (Kitcher 2022).[8] Science is not always like building a house according to a blueprint, where progress can be tracked toward achieving a clear goal. Instead, scientists often track progress against past performance. In line with this, here is an extended quotation from P1, whose research focuses on using the motion of stars near the center of the galaxy to investigate dark matter. Although the effect of dark matter on stellar orbits is very small, with long-term and highly precise observations, it should be possible to detect its signature. They developed computer models by testing single dark matter profiles, and then attempted to build more flexible models that could account for a wider range of dark matter density shapes. However, this flexibility required more observational data to constrain the many parameters. Their work therefore aimed to see whether it was possible to balance model sophistication and flexibility with realistic limits on observability.

> When we first started working with [AI], we basically began by throwing data at it. At the beginning, we didn't use real observational data. Instead, we generated synthetic, or 'fake,' data. These models can be used to create such artificial data, which you can then feed back into the model to test how well it performs.
>
> This approach is useful for asking questions like: If I observe for ten more years, will I start to see an effect?

[7] Our distinction between success and progress is related to the distinction in the literature between aims and progress. See, e.g., the discussion between Rowbottom 2023 and Dellsén 2025.

[8] Rowbottom (2023, 44) has argued that Kitcher's distinction between progress-from and progress-to is unhelpful in the context of the debate about scientific progress because we can express each of the four accounts using either concept. E.g., we might say that scientists aim to make yesterday's theory more truthlike today, or that they aim to increase their theory toward maximum truthlikeness in the end. While that might be true, we think it is still a sensible distinction between ways that scientists conceive of progress towards particular goals that they have in their daily work (e.g., progress towards mapping the human genome vs. progress in increasing measurement accuracy at the Laser Interferometer Gravitational-Wave Observatory, see Westerblad, de Regt and Stuart forthcoming).

> The trial-and-error process went something like this: at first, we were very optimistic. We thought that with just one orbit of data, and given the current measurement accuracy, we might be able to constrain the model. But once we tested it, we quickly realized there was no way that would work.
>
> So the next step was to ask: What if I improve the instruments? We narrowed the error bars on the synthetic data, only to realize that the required level of precision was so extreme that it simply wouldn't be achievable in the foreseeable future.
>
> That forced us to find a middle ground: for example, assume we could observe for 20 years and improve accuracy somewhat, not to unrealistic levels, but enough to matter. Then we asked, What can we still squeeze out of the model under those conditions?
>
> When using more flexible models, the trade-off is that they become much harder to constrain with observations. Our initial hope was that the model would directly fit the dark matter profile and reveal how it is distributed, whether like this or like that. Depending on the outcome, we could then infer what kind of dark matter it might be, or at least rule out certain possibilities.
>
> We wanted the model itself to reveal how dark matter is distributed, and from that distribution we can infer what type of dark matter it might be, or at least rule out models that are clearly incompatible. Unfortunately, we have not yet been able to achieve this.
>
> In the end, it turned out that the results weren't as promising as we had hoped. (P1, 28-06-2024)

This is an episode in which the use of AI was deemed not to be very progressive. The ultimate aim is to find out about the nature of dark matter, and the subgoal pursued was trying to build AI models with certain properties. The negative evaluation has less to do with AI's inability to tell us the truth about dark matter, and more to do with the team's inability to do anything more useful or interesting than what they could already do. And this was measured in terms of ability: AI does not increase our ability to discriminate between models of dark matter given the empirical data that we have, and that is unfortunate.

Prima facie, measuring progress in terms of the improvement of abilities might tell against the semantic, epistemic and noetic accounts. But there is a lot of wiggle room. Proponents of those accounts might argue that scientists simply learn to talk in terms of problem-solving and capacity-building, perhaps because of incentive structures, funding mandates, or modesty, while "genuine" progress nevertheless consists in increasing the truthlikeness of theories or the knowledge/understanding of scientists. Consistent with this, it also might be that improving abilities promotes progress, but does not constitute progress (for our reply to this worry see section 4.5). A natural argument to combine with the above two is to say that accounts of progress are purely normative, and therefore if scientists do not speak or act in accordance with some account, so much the worse for science (for responses to these kinds of claims, see Mizrahi 2020 and Rowbottom 2023, 21).

Nevertheless, we think it is worth mentioning that our participants speak of abilities rather than truth, knowledge, or understanding when evaluating the scientific value of AI. While we do not claim this is

generally true across science, insofar as it does generalize, recognizing this emphasis might justify shifting the discussion of scientific progress somewhat towards ability.

## 4.2 Little mention of opacity

Another surprising absence in our interview data was any explicit mention of computational opacity. Given the centrality of this concept in the philosophy of AI (Stuart 2025), one might think that this would be high up on the list of conversation topics. However, at least in the case of space science, this does not seem to be a major concern.

Why might this be? Here is NASA's chief engineer Joe Pellicciotti again: "We've taken the design of a product, and we've used AI to make that design more mass efficient, more stiffness efficient, and so forth. In the end, you know, we can take that design, and we can test it and make sure it still meets all of our qualifications, but that design has…it's taken less time to go through the process. It's more efficient, which lowers cost" (Almeida 2025). Pollicciotti gives a general template for AI-use that we saw many times: Scientists use AI to identify something useful, e.g., a way to make a component lighter, or a flightpath more efficient. They then try out the new idea and verify that it works. For example, they can weigh a model of the component after modifying it or run a proposed flightpath through the usual calculations to make sure that it works.

To exemplify this, here is an extended quotation from P2, who is working on applying graph-structured data and neural networks to novel problems. In this quote, they describe their current work as developing and testing machine learning systems on simplified experimental setups to better understand how well new ideas perform within a system:

> We use certain metrics to measure how well the system performs. In machine learning, this could be accuracy. Take image recognition: accuracy would mean how often the model correctly identifies the content of an image.
>
> In our case, we train small 'cubes' to arrange themselves into a specific configuration on their own. Another way to think of it would be a swarm of satellites: they should autonomously form a certain constellation as quickly as possible. The performance metric here is how successfully and how quickly they reach that target configuration.
>
> Typically, you start with something simple – you don't go straight to the end goal. Instead, you first test your idea on a very basic setup, where it should definitely work. The metric then acts as a signal: it ranges from 0 to 1, with 1 being very good. If it's close to 1, say 0.8, that suggests the system is working. But if it's stuck at 0, then something is wrong, either in the concept or in the implementation.
>
> These signals guide the experimentation. You then think about what else you can do to better understand what's happening inside the system. For example, one thing we did recently was rerun the experiment under different noise conditions. Another useful approach is plotting various

> outputs, visualizing graphs of the system's behavior, and directly observing what it's doing. In our case, that means watching how it reconfigures itself over time. (P2, 09-09-20)

It doesn't matter that AI is opaque because there are other methods for verifying the outputs of AI models. The way the scientists experiment with models does not aim at discovering much about the computational details of how they work: it is more about getting a clear picture of the model's behavioural profile, so they can use it.

What does this mean for the debate about scientific progress? If the epistemic or noetic accounts are correct, opacity should be a concern for scientists, as opacity stands in the way of progress, by standing in the way of justification (a problem for the epistemic account) and grasp (a problem for the noetic account).

A reply from the epistemic account might be that opacity is not a serious concern precisely because the outputs of AI can be justified by other means. However, on Bird's version of the epistemic account, science makes progress when knowledge is increased, as long as that knowledge is socially accessible. AI models might contain a great deal of knowledge that humans cannot access. For all we know, there might be a great deal of progress hidden in those algorithms (Birch 2025). Scientists should be interested in finding out whether that is the case, given that they value progress. But they are not pursuing algorithmic explainability to access whatever knowledge there might be in the AI models. Rather, explainability, when it is pursued, is pursued in connection with use and performance. In a recent interview, scientists at Google DeepMind were asked whether drug discovery algorithms needed to be explainable. They replied that "the call for explainability" is another way of signalling that the model isn't working well, that is, "we need to make this model better." Explainability is for helping us "understand the pathologies that this model has, the biases that it has." The goal is to "get to this point where, yeah, actually, we can just do end-to-end design purely in silico, and maybe do a final round of verification in the lab at the end" (Fry 2025). Reliably capable algorithms are scientifically desirable despite opacity because of what they enable scientists to do. This is easy to make sense of if progress is somehow related to increasing scientific abilities: the opacity of human abilities (like the ability to generate good new ideas) does not stand in the way of their value, so it should perhaps be no surprise that we see the same in the case of computationally mediated abilities.

Another reply that one might make on behalf of the epistemic account is that AI is only useful in the context of discovery, not the context of justification. Indeed, there are good reasons for thinking that AI is used as a tool of discovery, and does not, on its own, provide full justification for much (Duede 2023). But very little processes do provide complete justification on their own: scientific justification is holistic (Dürr and Dellsén forthcoming; Elgin 2017). Even the best laboratory experiments do not, on their own, justify much: they also require a justified theoretical interpretation of the results and justified inductive inferences to extend the findings to the real world. Likewise for AI use in space science. Here is an example.

M4 researches the fate of planetary systems after their host stars evolve into white dwarfs. They investigate the elemental composition of planetary debris that accretes onto white dwarfs, enabling reconstruction of the original planets' makeup. This work offers insights into planetary chemistry beyond bulk density, something typically inaccessible in studies of intact exoplanets. M4 describes their process as one that relies heavily on imagination, observation, and connection to other fields of research:

> Astronomy, my field, is a particular kind of science, it's observational. You can't run lab experiments in the traditional sense. If something explodes, you can't rewind and repeat the event. It might be a unique occurrence, and you may need to wait another 100,000 years to see it again, which is, of course, beyond the length of a PhD.
>
> Breakthroughs in nearby fields are also exciting. It's satisfying to witness others succeed, and it motivates you in your own work. You're trying to discover something, driven by curiosity and the desire to understand how the universe works, what it's made of, whether we truly understand physics. We think we do, but there are still areas where current laws fall short, and we might need entirely new ones. Understanding how things work and interconnect - that, to me, is incredibly rewarding.
>
> Imagination is about open-mindedness and the ability to connect concepts, items, and ideas. To do that, you need exposure to a variety of thoughts. It's like assembling with Lego bricks: if you only have one type, your possibilities are limited. But when someone brings a different shape from another field, your options expand. This also includes methods, worldviews, and knowledge of boundaries...Sometimes, knowing where not to go is just as important. We borrow from fields like medical imaging or geophysics to develop new techniques.
>
> In research, we're not doing something mechanical or repetitive. The moment you make real progress, you're doing something no one has done before. You can't rely on established patterns; you need imagination to envision the next step. During my PhD, I studied star clusters to understand stellar evolution. But we realized we could use those clusters to study how galaxies formed - a shift that required imagination. We encountered contradictions with existing theories: galaxies seemed younger than the ancient star clusters they contained. Eventually, we realized young stars outshine old ones, masking older structures. At first, we thought other researchers were simply wrong, but then we asked: what if both views are valid? It took cycles of imagining different perspectives to reconcile what we saw.

As this excerpt shows, it's not always possible to cleanly separate discovery and justification, at least in part because of the meandering, holistic nature of progress. Most discoveries are partial justifications, and most justifications are partial discoveries (Buzzoni 2015). Insofar as AI does play a justificatory role in space science, the epistemic account cannot claim that AI is always and only a tool for discovery, in order to explain away the lack of attention to opacity by space scientists.

Turning to the noetic account, the lack of mention of opacity is surprising because opacity is thought to frustrate grasp (Beisbart 2021; Janvid 2018; Lenhard 2018; Stuart and Nersessian 2019), and on most accounts, grasp is necessary for objectual understanding. Specifically on Dellsén's account, understanding is achieved when a sufficiently comprehensive dependency model is grasped, and understanding can be improved (and progress made) when that grasp is improved. Perhaps there is a way to avoid this apparent tension by augmenting the notion of grasp, for example, by weakening it (Khalili 2024), or identifying it with something akin to an ability, skill, or know-how (Belkoniene 2023; Strevens 2024; 2025; Stuart 2025).

But taking this latter option brings the noetic account closer to the functional account. More on this potential unification in section 4.5.

## 4.3 A new proposal: progress as improving abilities

In light of our data, we wish to propose a new account of scientific progress, one that straddles several of the existing accounts. We claim that progress in space science, at least in AI-relevant contexts, concerns improving at least one relevant scientific ability. This is connected to the progress-as-problem-solving account, but importantly different. Solving a single problem can be an instance of what we called "success," above. But such instances might also be thought of as highlights on the way toward improving our abilities. It is this kind of development we want to honour with the label "progress."

Importantly, not all scientific abilities concern problem solving: some relevant abilities might be imaginative or exploratory, e.g., finding new problems, or new concepts, or new methods, even in the absence of the existence of a particular problem.[9] We therefore claim that the central notion is scientific ability, not problem-solving or problem-defining. Improving scientific abilities can be epistemically progressive by providing a new ability to solve a particular problem, but it can also be progressive simply by building or improving capabilities. Here is the proposal:

> The ability account of scientific progress: Science makes progress when scientific abilities are improved.

This needs to be developed and defended. We begin by returning to the beginning of the paper, to briefly sketch out some options for how this account might address points (i)-(vii). Who makes progress? The relevant abilities might be possessed by individuals, groups, or distributed across individuals and instruments. What does progress improve? The quality of the scientific abilities of an agent or group. What counts as an "improvement"? Improvement includes the development of new abilities, or the improvement of existing abilities. Improvement might be measured in terms of the increase in precision, robustness, reliability, accuracy (etc.) of the actions/processes that result from exercising those abilities. What is the "bearer" or "vehicle" of progress? Abilities become the unit of analysis. This is different from information stored in a theory or model, or the cognitive-epistemic state of scientists. What is the right scale of analysis? Perhaps it is permissible to conceptualize abilities as being possessed by very large-scale communities, or even the human species as a whole, developed over millennia. Or perhaps we should focus only on abilities possessed by single individuals, developed over short timespans. We leave this open. What kinds of progress can we track? On the ability account, there can be as many kinds of progress as there are kinds of ability. We might decide which kinds of abilities are relevant to our analysis of scientific progress depending on choices about the scale of our case study, or what kinds of agents are involved, or whether we are focusing on purely theoretical vs. more applied vs. politically sensitive science. E.g., we might distinguish between mathematical

[9] Of course, with some creativity, it is possible to interpret any ability as a problem-solving ability. E.g., the ability to do mental multiplication is a solution to the problem of multiplying numbers. However, this seems like an ad hoc strategy for saving the problem-solving version of the functionalist account that is not sufficiently well motivated. We take problems to be those which are named by scientists (currently or in retrospect) as problems that they are or were trying to solve.

abilities and physical abilities, or between aesthetic, epistemic, ethical, political, and practical abilities. There can be a type of progress for each type of ability. With that said, for the ability account to compete with the semantic, epistemic, and noetic accounts of progress, we might want to restrict our focus to those scientific abilities which are "cognitive" or "cognitive-epistemic." (More on this restriction below). How is progress related to other issues in philosophy of science? We highlight a few directions in the following. But before that, three caveats.

First, we want to address the possibility that science can make progress in different ways (Chang 2004; Goebel 2019; Rowbottom 2019a, 2023). It could be progressive for science to increase the truthlikeliness of theories, or to have scientists possess more knowledge and understanding, or to solve more problems, or to increase scientific ability. Perhaps all of these changes to science are good in their own, distinct, yet fundamental ways. So far, the debate has mostly assumed monism (Dellsén forthcoming), but whether we are pluralists or monists, the ability account can find a place, either as describing one of the fundamental kinds of progress, or as describing the single fundamental kind of progress. We do not take a stance on this issue here.

Second, we should distinguish between instrumental and categorical progress (Rowbottom 2023), or between promoting and constituting progress (Dellsén forthcoming). Following Dellsén, we can say that an event promotes progress to the extent that it raises the chance that future events will constitute progress. For example, a new method might promote progress on the semantic view because it increases the chance of more truthlike theories in the future, even though it does not increase the truthlikeness of any theory. Proponents of factive accounts of progress will allow that improving scientific abilities can promote progress, but would deny that improving abilities could constitute progress. This does not cohere well with the empirical findings presented above, where progress seems to be constituted by improving abilities. Still, it is fair to ask how improving scientific ability could constitute progress in itself, full stop. After all, why are abilities valuable, if not for increasing our stock of truth, knowledge and understanding? We worry that the unintuitive nature of our claim might merely be a felt incongruence with long-ingrained habits in philosophy. Thinking about ability as a fundamental kind of good might require effort, but that effort can be worthwhile. And indeed, there are worries about the value of truth, knowledge, and understanding as well. If these are not themselves grounded in abilities, what are they good for? And as we will argue below, on certain versions of the epistemic and noetic account, knowledge and understanding seem to be composed at least partially of abilities, so improving them would be constitutive of progress, if increasing knowledge and understanding counts as making progress.

Finally, the ability account might be extended beyond science. Perhaps art makes progress when artists improve their abilities to express ideas, or stimulate a certain emotional response in an audience, or challenge a tradition. Likewise for society, we might make progress when we improve our abilities to understand one another and treat each other fairly. Having a general account does not preclude our specifying it in ways that interest us. We can distinguish scientific epistemic progress from scientific ethical progress, and focus on whichever we like. This seems to us a happy consequence, and perhaps a point in favour of the ability account, but we restrict our focus to science here and do not defend anything more general for the purposes of this paper.

## 4.4 Similarity to other accounts

Let's consider how this account relates to its two closest relatives: Shan's functional account and those accounts that focus on know-how.

Shan claims that "science progresses if and only if more useful exemplary practices are proposed" (Shan 2022, 50). What does it mean for a practice to be useful? According to Shan, practices are useful iff they propose ways of defining and solving problems, where those ways are new, repeatable, fruitful (lead to new problems), and generalizable. We agree that good practices are often those that enable framing and solving problems, and often, those practices are especially good when they are new, repeatable, fruitful, and generalizable.

However, we deny that progress only ever consists in the proposing of such practices. First, proposing is not always enough: progress requires taking up proposals and acting on them, as well as in achieving our goals by means of them. On our view, the proposal of useful practices often counts as progressive only when acting on such proposals leads to improved scientific abilities, or when those proposals make clear how to apply existing abilities in a new way to achieve scientific goals. Second, our account can explain why the virtues listed by Shan (novelty, repeatability, fruitfulness and generalizability) are virtues of practices, while at the same time allowing us to capture counterexamples where a practice is progressive despite lacking one or more of those virtues. Repeatability, fruitfulness, and generalizability can be seen as virtues of practices because they describe good-making features of abilities. Consider generalizability. Scientific abilities that are not generalizable are less good than abilities that are generalizable. If an ability can be applied to many systems in many contexts, the agent who possesses that ability is, in a sense, more able. Though, importantly, even non-generalizable abilities can be progressive, as long as they enable the achievement of some scientifically relevant goal. So generalizability, while a good-making feature of abilities, is not necessary for progress. In terms of fruitfulness, most abilities desired by scientists will be desired because they are fruitful. As we will see below, "being able" to do something is often cashed out in terms of reliably managing to achieve some goal by means of intentional action.[10] Abilities that cannot be repeated and do not perform some useful function would not count as abilities, according to most accounts of what an ability is. Newness, however, is not a necessary virtue for a proposed practice to be progressive. A proposed practice might be an old one that we forgot about. Proposing it now can still be a good thing, it can still be progressive, insofar as it improves our current set of abilities.

We agree that Shan's definition allows us to capture many genuine instances of progress, and we agree that the case studies he uses to motivate his account refer to genuine examples of progress. But we think the

[10] We might not want to define ability as a success term. For example, imagination is an ability that is "free" by nature, or at least, freer than most other cognitive processes. So an agent might genuinely have the ability to imagine, yet only imagine in ways that lead to useless ideas, as measured by scientific metrics of success. This helps to motivate fruitfulness as a virtue of scientific proposals, because the improvement of abilities should be fruitful in order to say that progress has been made. And indeed, scientists recognize that imagination might be used in non-fruitful ways, and they value it mainly for its fruitful consequences (Stuart 2019, 2022).

virtues of progressive episodes that he cites are better explained on our account, that is, by focusing on improved abilities rather than useful proposed practices.

Nevertheless, we take our account to be close to the functional account, since we take progress to concern what science can do, rather than being something about the quality of the propositional content of science, or the quality of the cognitive relation between epistemic agents and such content. But it differs from Kuhn's and Laudan's original functional accounts because instead of talking about the number or significance of problems solved, we focus on improving abilities, which may or may not be problem-solving abilities. Our account therefore will not face all the objections brought against the traditional problem-solving account. Still, it is worth considering how our account might deal with versions of those traditional objections, and other objections, and we turn to these in the next section.

But there is another family of accounts we should consider, which portray scientific progress at least partially in terms of the accumulation of knowledge-how. Mizrahi (2013) draws on historical case studies to suggest that scientific progress consists in the accumulation of knowledge, but distinguishes four kinds of knowledge, each of which are on a par: empirical knowledge (e.g., of empirical regularities), theoretical knowledge (e.g., of well-confirmed hypotheses), practical knowledge (e.g., as in practical applications), and methodological knowledge (e.g., as in methods). The latter two are not defined in terms of know-how. But the examples Mizrahi gives are. That is, knowing-how to treat certain diseases and disorders, and knowing-how to study anatomy using certain surgical techniques. Mizrahi draws on Baird and Faust (1990, 147) who claim that the knowledge of scientists sometimes "consists in the ability to *do* things with nature, not *say* things about nature" (see also Rowbottom 2023, who also claims that scientific progress can consist in increases in know-how). Below, we will consider the possibility of reducing knowledge-how to knowledge-that. For now, it is important to note that there is no consensus on the relation between know-how and ability. Pavese (2024) argues that we can know-how to swim (move your arms and legs around *like this*) without necessarily being able to swim. As we will see below, some argue that increasing our abilities, e.g., to treat diseases or get useful information from a target system, are better characterized under a non-factive kind of ("pragmatic") understanding. In sum, while no one in the literature explicitly defends the view that scientific progress is increasing know-how, this is a possible view that someone might defend, and it might be very closely related to the view we set out here, depending on how we characterize the relation between know-how and ability. More generally, as we will see below, how such accounts relate to the semantic, epistemic, and noetic views depends on how we structure the relations (and hierarchies) between epistemic goods like truth, knowledge, understanding, and ability.

## 4.5 Response to objections

Sometimes scientists appear to develop abilities that should not be associated with progress because they are not "genuine" abilities. Examples might include manipulating phlogiston, making phrenological/astrological predictions, and creating the alchemist's panacea. If we cannot differentiate genuine from merely apparent abilities, how can we differentiate between genuine and merely apparent progress? There are several ways we might reply. One is to say that scientists can be wrong about what abilities they have. In some cases, scientists merely *thought* they had developed certain abilities. Of course, in most cases, there usually *are* some genuine abilities that were developed, and therefore genuine progress that

was made. Examples include abilities to measure, infer, and manipulate features of some system. Such abilities typically remain even after the removal of the false theory. As Chang has argued (2012, 53-4), phlogiston theorists were able to get quite a lot done, and it was these abilities that scientists built upon (and continue to build on), even after oxygen theory was introduced. For example, it was through Priestly's efforts to "de-phlogisticate" air that chemists gained the ability to make oxygen (Chang 2022, 152). We might therefore say that scientists really were developing abilities, and thus making progress, and notably, those abilities were central in motivating the turn from phlogiston to oxygen. But the scientists were wrong about the nature of their abilities, and thus about the nature of the progress they were making.

The history of science provides some reason to take this proposal seriously. It is said that the 100 most-cited science papers are methods papers, that is, papers that introduce a method. Scientists may stop citing papers about unobservable entities they no longer take to exist, but they don't stop citing papers that introduce techniques that improve our abilities. As Douglas writes,

> Scientific progress can be defined in terms of the increased capacity to predict, control, manipulate, and intervene in various contexts…While paradigm change can create losses in understanding or losses in explanatory unification as clear conceptual structures are swept away, what is not lost is the ability to predict phenomena and/or the ability to control aspects of the world…we are hard pressed to think of a predictive or manipulative capacity that has been lost. (2014)[11]

If there are paradigm shifts, and if some abilities improve across them, this suggests one possible answer to the question of how we can be confident that a given ability is a genuine ability: those that can be used effectively across different times, scientific fields, and background theories, are genuine.

To go further, we should differentiate between the quality of an ability as opposed to the quality of a *description* of an ability. Focusing first on the latter, what we want is a reason to believe that our descriptions of abilities are accurate. A difficulty is that those descriptions can include reference to unobservable or theoretical entities, and in those cases, for the description to be accurate, those entities must exist as described. The pessimistic metainduction makes it difficult to fully justify such claims. Interpreting things this way, we find ourselves re-creating the realism debate. Anti-realists will claim that we ought not commit ourselves to beliefs about having abilities whose descriptions contain reference to unobservable theoretical entities. Realists will attempt to identify the sorts of theoretical posits that can be trusted to survive theory change and which can thus be included in our ability-descriptions. Finally, pragmatists will hold that realism and anti-realism are merely frameworks or stances (as opposed to truth-evaluable claims about the accuracy of ability-representations) which are only better or worse to the extent that they are useful (Boucher and Forbes 2024). The possibility of reframing the realism debate in terms of abilities is not, on its own, a

[11] Douglas goes on to say that increased abilities to destroy humans should not count as progress. We should add some ethical limits or constraints on the notion of progress. For those who agree, we can add to our account that the improvement and deployment of abilities should be done responsibly. One way to put this is that epistemic abilities must be developed in connection with ethical abilities. Thinking of the ethics of science in this way might open fruitful options in the debate about which roles non-epistemic values should play in science (cf. Brown 2020).

problem for our view. It might even be fruitful. The realism debate will continue in some form or another, and having a version of it focused on ability descriptions does not seem like a bad thing.

However, there is another way to investigate whether an ability is genuine. This way does not focus on accurate descriptions of abilities, but on the quality of the abilities themselves, perhaps centred on the connection between abilities and the world. If "ability" is a success term, we succeed in our projects by means of abilities that allow us to manipulate systems (whether material, theoretical, mathematical, fictional, etc.). Vetter (2024) distinguishes between two senses of ability: *simple* abilities delimit the set of all actions that are possible for an agent. In this sense, we're able to do whichever action it is possible for us to do, perhaps in the sense that in at least one possible world, we do it. Then there are *robust* abilities, which delimit actions that are, in some sense, within our power to do, in the actual world. It is toward these latter abilities that epistemologists have recently turned their interest.

What explains this power we have to act in ways that reliably achieve some end? Traditionally-minded philosophers of science might argue that the answer lies somehow with the possession of propositional knowledge or truth. There are at least two ways to make such a claim. On the first, the traditionalist will argue that agents are able to, e.g., measure the value of a variable, because they *know how* to do that, and they know how to do that because they *know that* method *x* is the way to measure the value of that variable. This reduces abilities ontologically to propositional knowledge. Having an ability just is having some know-how, and having know-how is just having propositional knowledge (e.g., about which actions achieve which aims in which circumstances). The second strategy allows that abilities might be ontologically different from epistemic states like propositional knowledge, but contends that an epistemic ability is valuable (when it is) only because the agent has or gains knowledge about the target system to which the ability will be applied. Thus, a radiologist may be able to see cancer tumors in computed tomography (CT) scans, and that is an epistemically valuable ability which is different from any states of propositional knowledge. But the reason the radiologist's ability is valuable is because it is justified by background knowledge or true beliefs about, e.g., what cancer tumours look like in CT scans, how the CT scanning process works, etc. The key move is then to claim that the "real" value of the radiologist's ability is wholly to be located outside the ability itself, e.g., in the propositional background knowledge, or in the new knowledge (or true beliefs) that the ability produces. For example, that this patient's scan results do not display signs of cancer.

Whether these arguments can be made to work is an open question, and we won't pursue them further here. Still, we think new research in the epistemology of ability points in interesting non-reductionist directions that justify independent attention to the epistemic features of abilities, like control, adaptability, reliability, and success (Mayr and Vetter 2023). Improving our abilities seems to be at least one kind of scientific progress, and it seems possible to define, discuss, and measure such progress independently of increases in propositional knowledge and truth. If this is the case, then, to come back finally to the question raised at the beginning of this section, genuine abilities will be those that enjoy the right kind of connection with the target systems that they are applied to, and the goals which they help us achieve. The epistemology of ability is still at an early stage with respect to specifying the options for how these connections might look. Perhaps it will have to do with embodiment, enactivism, embeddedness, affordances, trial and error, evolution, and much else. And we should not necessarily expect a unified account that works for all abilities, goals, and systems. In the end, then, the response is twofold: 1) It seems possible to tell a story that grounds the

(epistemic) value of abilities without a) reducing those abilities ontologically to propositional epistemic states, or b) reducing that value to the value of some propositional epistemic states, and 2) It seems possible to tell a story that goes in the other direction: that is, to reduce other sources of (epistemic) value to ability. We now turn briefly to this possibility.

There are views of knowledge that define knowledge in terms of ability. For example, Chang's view of "active" knowledge claims that knowledge "is a matter of our ability to engage productively with reality" (2022, 119). It is knowledge "as ability" (2022, 18). If scientific progress consists in the increase of knowledge, and if knowledge just is ability, or if knowledge increases only as abilities are improved, then the epistemic account of progress can be wholly or partially explained by (or grounded in) the ability account. In other words, the value of new knowledge reduces to (or is grounded in) the value of new abilities.

The noetic view claims that science makes progress when scientific understanding is increased. The most detailed version of this view is Dellsén's (Dellsén 2016; 2021; 2022; Norton et al. forthcoming; Dellsén et al. 2022). On Dellsén's development, there are three "hallmarks" of such understanding: "the ability for successful predictions," "the ability to formulate successful explanations of some target phenomenon," and achieving "conceptual integration" (i.e., creating a coherent network of dependency relations between elements of some domain) (Dürr and Dellsén forthcoming). Two of these three hallmarks are abilities. If there is a way to explain the achievement of conceptual integration in terms of abilities, then (the value of) all three hallmarks could be reduced to (the value of) abilities. In addition, as we saw above, insofar as grasp is involved, this might be best characterized in terms of abilities (Strevens 2025). If this is correct, it seems that the noetic account can be reduced to a version of the ability account that defines progress in terms of the improvement of a number of specific abilities.

By looking at the broader literature on understanding, we can easily imagine other versions of the noetic account, and many of these will be friendly to the ability account of progress. For example, it is widely agreed that understanding (whether explanatory, objectual or practical/pragmatic) is, or is characteristically exemplified by, some kind of ability, skill, or cognitive mastery (Le Bihan 2017, Rowbottom 2019a). One kind of understanding, called practical or pragmatic understanding, is argued to simply be an ability (Delarivière and Van Kerkhove 2021; Toon 2015; Currie 2020; Leonelli 2009; Lenhard 2006; 2009; 2019). For Stuart, this kind of understanding is having a praiseworthy skill, and Stuart leaves open the possibility that skills might be developed versions of abilities (2025). If understanding is, centrally involves, or requires having certain abilities, and the value of understanding can be explained wholly or partially by reference to those abilities, then it makes sense to characterize scientific progress fundamentally in terms of the improvement of those abilities. It seems then that the ability account could explain, ground, or support, many different versions of the noetic view of scientific progress.[12]

There are difficult issues lurking in the background here about which we do not have the space to go into detail, e.g., the nature and types of abilities, demarcating or counting abilities, and the possibility of cases where abilities might conflict, etc. We leave these for future discussion. Our main goal has been to suggest a

[12] If truth(likeness) can be reduced to ability, the semantic account could also be explained by (or reduced to) the ability account. We can think of some potential candidate strategies for such a reduction (e.g., in the American pragmatists), but do not pursue the idea further here.

new account of scientific progress, which we claim should be taken seriously as an account of scientific progress. More generally, we hope to increase philosophical interest in scientific abilities.

# 5. Conclusion

This paper presented new qualitative empirical data and proposed a new account of scientific progress for space science. We are happy to categorize it either as a functional or non-factive noetic account. We also suggest that the ability account might (at least partially) explain or ground some of the other accounts.

We stress again that our qualitative evidence and resulting account only bear on the general debate about scientific progress insofar as space science is a kind of science. We think it is. But even if features of space science, e.g., its close connection to technology, prevent broad generalization to science as a whole, the set of practices we *can* generalize to will still form an important set, perhaps including many engineering science projects (including fusion energy projects, particle collider physics, and large telescope astronomy and cosmology), science involving robotics (including automated chemical and pharmaceutical research), and a good deal of medicine and bioengineering. Given that philosophers have already extended at least some of the traditional accounts to these contexts, there is room for meaningful debate at least about those contexts (e.g., the noetic account on medical progress: (Dellsén et al. 2025, Dürr and Dellsén forthcoming). Further, as the evidence given above shows, space science is not merely technological: it is also centrally concerned with testing hypotheses (e.g., about the distribution of dark matter in our galaxy), analyzing data (e.g., satellite data for climate research), and refining theories (e.g., ΛCDM), all of which are core practices of science.

Our account claims that we make progress in science by improving scientific abilities. We claim that this account is a good start for making sense of the fact that scientists find AI to be a progressive tool in space science when it helps them solve specific problems or improve their abilities to, e.g., solve problems, find patterns, explore difficult and unknown environments, imagine in useful ways, etc. If this is correct, much more work now needs to be done to bring in existing insights on the nature of ability (from philosophy of action, mind, and epistemology, see, e.g., Vetter and Schoonen forthcoming) to philosophy of science, including about how abilities are distributed across tools and social groups (Toon 2015; Nersessian 2022). It also suggests that we can further our understanding of epistemic abilities via qualitative research into the practice of scientific progress.

# Acknowledgments

We wish to thank the editors of the collection for their flexibility, encouragement, and helpful comments. We also thank audiences at the Joint meeting of the Asia Pacific Philosophy of Science Association and the Logic, Methodology and Philosophy of Science and Technology chapter at National Yang Ming Chiao Tung University (2025), as well as Zach Pirtle, Sabina Leonelli, and two anonymous referees. Finally, we thank our participants for their trust, and for sharing their experiences and knowledge.

# References

Almeida, Andreas. 2025. *Celebrating NASA's Engineers*. Episode 148. Small Steps, Giant Leaps. March 5. Podcast. https://appel.nasa.gov/podcast/episode-148-celebrating-nasas-engineers/.

Alvarez, Amanda, Aylin Caliskan, M. J. Crockett, Shirley S. Ho, Lisa Messeri, and Jevin West. 2024. 'Science Communication with Generative AI'. *Nature Human Behaviour* 8 (4): 625–27. https://doi.org/10.1038/s41562-024-01846-3.

Antonsen, Erik, Barbara K. Burian, Sylvain V. Costes, et al. 2025. 'Chapter 8 - Artificial Intelligence in Precision Space Health'. In *The Digital Doctor*, edited by Chayakrit Krittanawong. Academic Press. https://doi.org/10.1016/B978-0-443-15728-8.00012-4.

Baird, Davis, and Thomas Faust. 1990. 'Scientific Instruments, Scientific Progress and the Cyclotron1'. The *British Journal for the Philosophy of Science* 41 (2): 147–75. https://doi.org/10.1093/bjps/41.2.147.

Barseghyan, Hakob, Paul Patton, Guillaume Dechauffour, and Carlin Henikoff. 2024. 'What Are Problems?' *Philosophical Studies* 181 (10): 2755–77.

Beisbart, Claus. 2021. 'Opacity Thought through: On the Intransparency of Computer Simulations'. *Synthese* 199 (3): 11643–66. https://doi.org/10.1007/s11229-021-03305-2.

Belkoniene, Miloud. 2023. 'Grasping in Understanding'. *The British Journal for the Philosophy of Science* 74 (3): 603–17. https://doi.org/10.1086/714816.

Birch, Jonathan. 2025. 'Knowing Science, by Alexander Bird'. *Mind* 134 (534): 515–23. https://doi.org/10.1093/mind/fzad048.

Bird, Alexander. 2007. 'What Is Scientific Progress?' *Noûs* 41 (1): 64–89. https://doi.org/10.1111/j.1468-0068.2007.00638.x.

Bird, Alexander. 2022a. *Knowing Science*. Oxford University Press.

Bird, Alexander. 2022b. 'The Epistemic Approach : Scientific Progress as the Accumulation of Knowledge'. In *New Philosophical Perspectives on Scientific Progress*, edited by Yafeng Shan. Routledge.

Boucher, Sandy C., and Curtis Forbes. 2024. 'The Pragmatic Turn in the Scientific Realism Debate'. *Synthese* 203 (4): 1–23. https://doi.org/10.1007/s11229-024-04528-9.

Boumans, Marcel, and Catherine Herfeld. 2022. 'Progress in Economics'. In *New Philosophical Perspectives on Scientific Progress*. Routledge.

Brown, Matthew. 2020. *Science and Moral Imagination*. University of Pittsburgh Press.

Buzzoni, Marco. 2015. 'The Practice Turn in Philosophy of Science: The Discovery/Justification Distinction, and The Social Dimension of Scientific Objectivity'. In *The Practical Turn in Philosophy of Science*, edited by E. Agazzi and G. Heinzmann. Angeli.

Chang, Hasok. 2004. *Inventing Temperature: Measurement and Scientific Progress*. Oxford University Press.

Chang, Hasok. 2012. *Is Water H2O?* Vol. 293. Boston Studies in the Philosophy of Science. Springer Netherlands.

Chang, Hasok. 2022. *Realism for Realistic People: A New Pragmatist Philosophy of Science*. Cambridge University Press.

Cheng, Xusen, and Lulu Zhang. 2025. 'AI-Generated Literature Reviews Threaten Scientific Progress'. *Nature* 641 (8064): 852–852. https://doi.org/10.1038/d41586-025-01603-0.

Collingwood, Robin George. 1946. *The Idea of History*. Oxford University Press. https://archive.org/details/ideaofhistory0000coll/page/n7/mode/2up.

Currie, Adrian. 2020. 'Bottled Understanding: The Role of Lab Work in Ecology'. *The British Journal for the Philosophy of Science* 71 (3): 905–32. https://doi.org/10.1093/bjps/axy047.

Delarivière, Sven, and Bart Van Kerkhove. 2021. 'The Mark of Understanding: In Defense of an Ability Account'. *Axiomathes* 31 (5): 619–48. https://doi.org/10.1007/s10516-020-09529-0.

Dellsén, Finnur. 2016. 'Scientific Progress: Knowledge versus Understanding'. *Studies in History and Philosophy of Science Part A* 56 (April): 72–83. https://doi.org/10.1016/j.shpsa.2016.01.003.

Dellsén, Finnur. 2018. 'Scientific Progress: Four Accounts'. *Philosophy Compass* 13 (11): e12525. https://doi.org/10.1111/phc3.12525.
Dellsén, Finnur. 2020. 'Beyond Explanation: Understanding as Dependency Modelling'. *The British Journal for the Philosophy of Science*, ahead of print. https://doi.org/10.1093/bjps/axy058.
Dellsén, Finnur. 2021. 'Understanding Scientific Progress: The Noetic Account'. *Synthese* 199 (3): 11249–78. https://doi.org/10.1007/s11229-021-03289-z.
Dellsén, Finnur. 2022. 'The Noetic Approach: Scientific Progress as Enabling Understanding'. In *New Philosophical Perspectives on Scientific Progress*. Routledge.
Dellsén, Finnur. 2023. 'Scientific Progress: By-Whom or for-Whom?' *Studies in History and Philosophy of Science* 97 (February): 20–28. https://doi.org/10.1016/j.shpsa.2022.11.006.
Dellsén, Finnur. 2025. 'Scientific Progress: Normative, but Aimless'. *Asian Journal of Philosophy* 4 (1): 37. https://doi.org/10.1007/s44204-025-00264-0.
Dellsén, Finnur. forthcoming. 'Promoting Scientific Progress'. *Philosophical Studies* 182 (8): 2123–43. https://doi.org/10.1007/s11098-025-02335-7.
Dellsén, Finnur, Insa Lawler, and James Norton. 2022. 'Thinking about Progress: From Science to Philosophy'. *Noûs* 56 (4): 814–40. https://doi.org/10.1111/nous.12383.
Duede, Eamon. 2023. 'Deep Learning Opacity in Scientific Discovery'. *Philosophy of Science* 90 (5): 1089–99. https://doi.org/10.1017/psa.2023.8.
Dürr, Patrick, and Finnur Dellsén. forthcoming. 'Scientific Progress and Modern Cosmology'. *European Journal for Philosophy of Science*. https://philsci-archive.pitt.edu/26683/.
Elgin, Catherine. 2017. *True Enough*. MIT Press.
Fry, Hannah. 2025. *A Quest for a Cure: AI Drug Design with Isomorphic Labs*. Good DeepMind the Podcast. June 5. 47:08. https://www.youtube.com/watch?v=XpIMuCeEtSk.
Goebel, Clara. 2019. 'A Hybrid Account of Scientific Progress: Finding Middle Ground Between the Epistemic and the Noetic Accounts'. *KRITERION – Journal of Philosophy* 33 (3): 1–16. https://doi.org/10.1515/krt-2019-330304.
Izzo, Dario, Gabriele Meoni, Pablo Gómez, Dominik Dold, and Alexander Zoechbauer. 2022. 'Selected Trends in Artificial Intelligence for Space Applications'. arXiv:2212.06662. Preprint, arXiv, December 17. https://doi.org/10.48550/arXiv.2212.06662.
Janvid, Mikael. 2018. 'Getting a Grasp of the Grasping Involved in Understanding'. *Acta Analytica* 33 (3): 371–83. https://doi.org/10.1007/s12136-018-0348-5.
Khalili, Mahdi. 2024. 'Against the Opacity, and for a Qualitative Understanding, of Artificially Intelligent Technologies'. *AI and Ethics* 4 (4): 1013–21. https://doi.org/10.1007/s43681-023-00332-2.
Kitcher, Philip. 2022. 'Scientific Progress and the Search for Truth'. In *Current Trends in Philosophy of Science*. Springer, Cham. https://doi.org/10.1007/978-3-031-01315-7_9.
Kleiner, Scott A. 1993. *The Logic of Discovery*. Springer Netherlands. https://doi.org/10.1007/978-94-015-8216-2.
Kuhn, Thomas S. 1962. *The Structure of Scientific Revolutions*. University of Chicago press.
Laudan, Larry. 1978. *Progress and Its Problems: Towards a Theory of Scientific Growth*. University of California Press.
Le Bihan, Soazig. 2017. 'Enlightening Falsehoods: A Modal View of Scientific Understanding'. In *Explaining Understanding: New Perspectives from Epistemology and Philosophy of Science*, edited by Stephen Grimm, Christoph Baumberger, and Sabine Ammon. Routledge.
Lenhard, Johannes. 2006. 'Surprised by a Nanowire: Simulation, Control, and Understanding'. *Philosophy of Science* 73 (5): 605–16. https://doi.org/10.1086/518330.
Lenhard, Johannes. 2009. 'The Great Deluge: Simulation Modeling and Scientific Understanding'. In *Pp. 169-188 in Henk de Regt et Al*, edited by eds).

Lenhard, Johannes. 2018. 'Thought Experiments and Simulation Experiments'. In *The Routledge Companion to Thought Experiments*. Routledge. https://doi.org/10.4324/9781315175027.ch27.
Lenhard, Johannes. 2019. *Calculated Surprises: A Philosophy of Computer Simulation*. Oxford Studies in Philosophy of Science. Oxford University Press.
Leonelli, Sabina. 2009. 'Understanding in Biology: The Impure Nature of Biological Knowledge'. In *Pp. 189-209 in Henk de Regt et al*.
Mayr, Erasmus, and Barbara Vetter. 2023. 'Adaptive Abilities'. *Philosophical Issues* 33 (1): 140–54. https://doi.org/10.1111/phis.12249.
Messeri, Lisa, and M. J. Crockett. 2024. 'Artificial Intelligence and Illusions of Understanding in Scientific Research'. *Nature* 627 (8002): 49–58. https://doi.org/10.1038/s41586-024-07146-0.
Mizrahi, Moti. 2013. 'What Is Scientific Progress? Lessons from Scientific Practice'. *Journal for General Philosophy of Science* 44 (2): 375–90. https://doi.org/10.1007/s10838-013-9229-1.
Miyake, Teru. 2022. 'Progress in Seismology: Turning Data Into Evidence About the Earth's Interior'. In *New Philosophical Perspectives on Scientific Progress*. Routledge.
Mizrahi, Moti. 2020. *Conceptions of Scientific Progress in Scientific Practice: An Empirical Study*. September 18. http://philsci-archive.pitt.edu/18117/.
Nersessian, Nancy J. 2022. *Interdisciplinarity in the Making: Models and Methods in Frontier Science*. MIT Press. https://doi.org/10.7551/mitpress/14667.001.0001.
Niiniluoto, Ilkka. 1987. *Truthlikeness*. Springer Netherlands. https://doi.org/10.1007/978-94-009-3739-0.
Niiniluoto, Ilkka. 2014. 'Scientific Progress as Increasing Verisimilitude'. *Studies in History and Philosophy of Science Part A*, Values and Norms in Modeling The Progress of Science, vol. 46 (June): 73–77. https://doi.org/10.1016/j.shpsa.2014.02.002.
Niiniluoto, Ilkka. 2017. 'Optimistic Realism about Scientific Progress'. *Synthese* 194 (9): 3291–309. https://doi.org/10.1007/s11229-015-0974-z.
Norton, James, Finnur Dellsén, Andrew J. Latham, and Somogy Varga. forthcoming. 'Medical Progress: Science Versus Practice'. *Erkenntnis*, 1–24. https://doi.org/10.1007/s10670-025-00953-5.
Pirtle, Zachery, and Jared Moore. 2019. 'Where Does Innovation Come From?: Project Hindsight, TRACEs, and What Structured Case Studies Can Say About Innovation'. *IEEE Technology and Society Magazine* 38 (3): 56–67. https://doi.org/10.1109/MTS.2019.2930270.
Popper, Karl. 1972. *Objective Knowledge: An Evolutionary Approach*. Oxford University Press.
Popper, Karl. (1963) 2002. *Conjectures and Refutations: The Growth of Scientific Knowledge*. 2nd edition. Routledge.
Potochnik, Angela. 2017. *Idealization and the Aims of Science*. University of Chicago Press. https://doi.org/10.7208/chicago/9780226507194.001.0001.
Rescher, Nicholas. 1984. *The Limits of Science*. University of California Press.
Rowbottom, Darrell P. 2008. 'N-rays and the semantic view of scientific progress'. *Studies in History and Philosophy of Science Part A*, Volume 39, Issue 2, pages 277-278.
Rowbottom, Darrell P. 2010. 'What Scientific Progress Is Not: Against Bird's Epistemic View'. *International Studies in the Philosophy of Science*, 24(3), 241–255
Rowbottom, Darrell P. 2019a. *The Instrument of Science: Scientific Anti-Realism Revitalised*. Routledge. https://doi.org/10.4324/9780429022517.
Rowbottom, Darrell P. 2019b. 'Scientific realism: what it is, the contemporary debate, and new directions'. *Synthese*, 196(2), 451–484.
Rowbottom, Darrell P. 2023. *Scientific Progress*. Cambridge Elements. Cambridge University Press. https://www.cambridge.org/core/elements/scientific-progress/21503F2D64B4381982CFEE6F290E5FC7.
Shan, Yafeng. 2019. 'A New Functional Approach to Scientific Progress'. *Philosophy of Science* 86 (4): 739–58. https://doi.org/10.1086/704980.

Shan, Yafeng. 2022. 'The Functional Approach: Scientific Progress as Increased Usefulness'. In *New Philosophical Perspectives on Scientific Progress*. Routledge.

Sourati, Jamshid, and James A. Evans. 2023. 'Accelerating Science with Human-Aware Artificial Intelligence'. *Nature Human Behaviour* 7 (10): 1682–96. https://doi.org/10.1038/s41562-023-01648-z.

Strevens, Michael. 2024. 'Grasp and Scientific Understanding: A Recognition Account'. *Philosophical Studies* 181 (4): 741–62. https://doi.org/10.1007/s11098-024-02121-x.

Strevens, Michael. 2025. 'Grasp as a Universal Requirement for Understanding'. *Philosophical Studies*, ahead of print, May 17. https://doi.org/10.1007/s11098-025-02342-8.

Stuart, Michael T. 2019. 'Everyday Scientific Imagination: A Qualitative Study of the Uses, Norms, and Pedagogy of Imagination in Science'. *Science & Education* 28 (6–7): 711–30. https://doi.org/10.1007/s11191-019-00067-9.

Stuart, Michael T. 2022. 'Scientists Are Epistemic Consequentialists about Imagination'. *Philosophy of Science* 90 (3): 518–38. https://doi.org/10.1017/psa.2022.31.

Stuart, Michael T. 2025. 'A New Account of Pragmatic Understanding, Applied to the Case of AI-Assisted Science'. *Philosophical Studies*, July 3, 1–25. https://doi.org/10.1007/s11098-025-02336-6.

Stuart, Michael T., and Nancy J. Nersessian. 2019. 'Peeking Inside the Black Box: A New Kind of Scientific Visualization'. *Minds and Machines* 29 (1): 87–107. https://doi.org/10.1007/s11023-018-9484-3.

Toon, Adam. 2015. 'Where Is the Understanding?' *Synthese* 192 (12): 3859–75. https://doi.org/10.1007/s11229-015-0702-8.

Vertesi, Janet. 2012. 'Seeing like a Rover: Visualization, Embodiment, and Interaction on the Mars Exploration Rover Mission'. *Social Studies of Science* 42 (3): 393–414. https://doi.org/10.1177/0306312712444645.

Vertesi, Janet. 2015. *Seeing Like a Rover: How Robots, Teams, and Images Craft Knowledge of Mars*. University of Chicago Press. https://press.uchicago.edu/ucp/books/book/chicago/S/bo18295743.html.

Vertesi, Janet. 2023. *Shaping Science: Organizations, Decisions, and Culture on NASA's Teams*. University of Chicago Press. https://press.uchicago.edu/ucp/books/book/chicago/S/bo49911700.html.

Vetter, Barbara. 2024. 'Abilities and the Epistemology of Ordinary Modality'. *Mind* 133 (532): 1001–27. https://doi.org/10.1093/mind/fzae037.