2025

EUREKA! ON STEROIDS

AI-driven research, development, and innovation

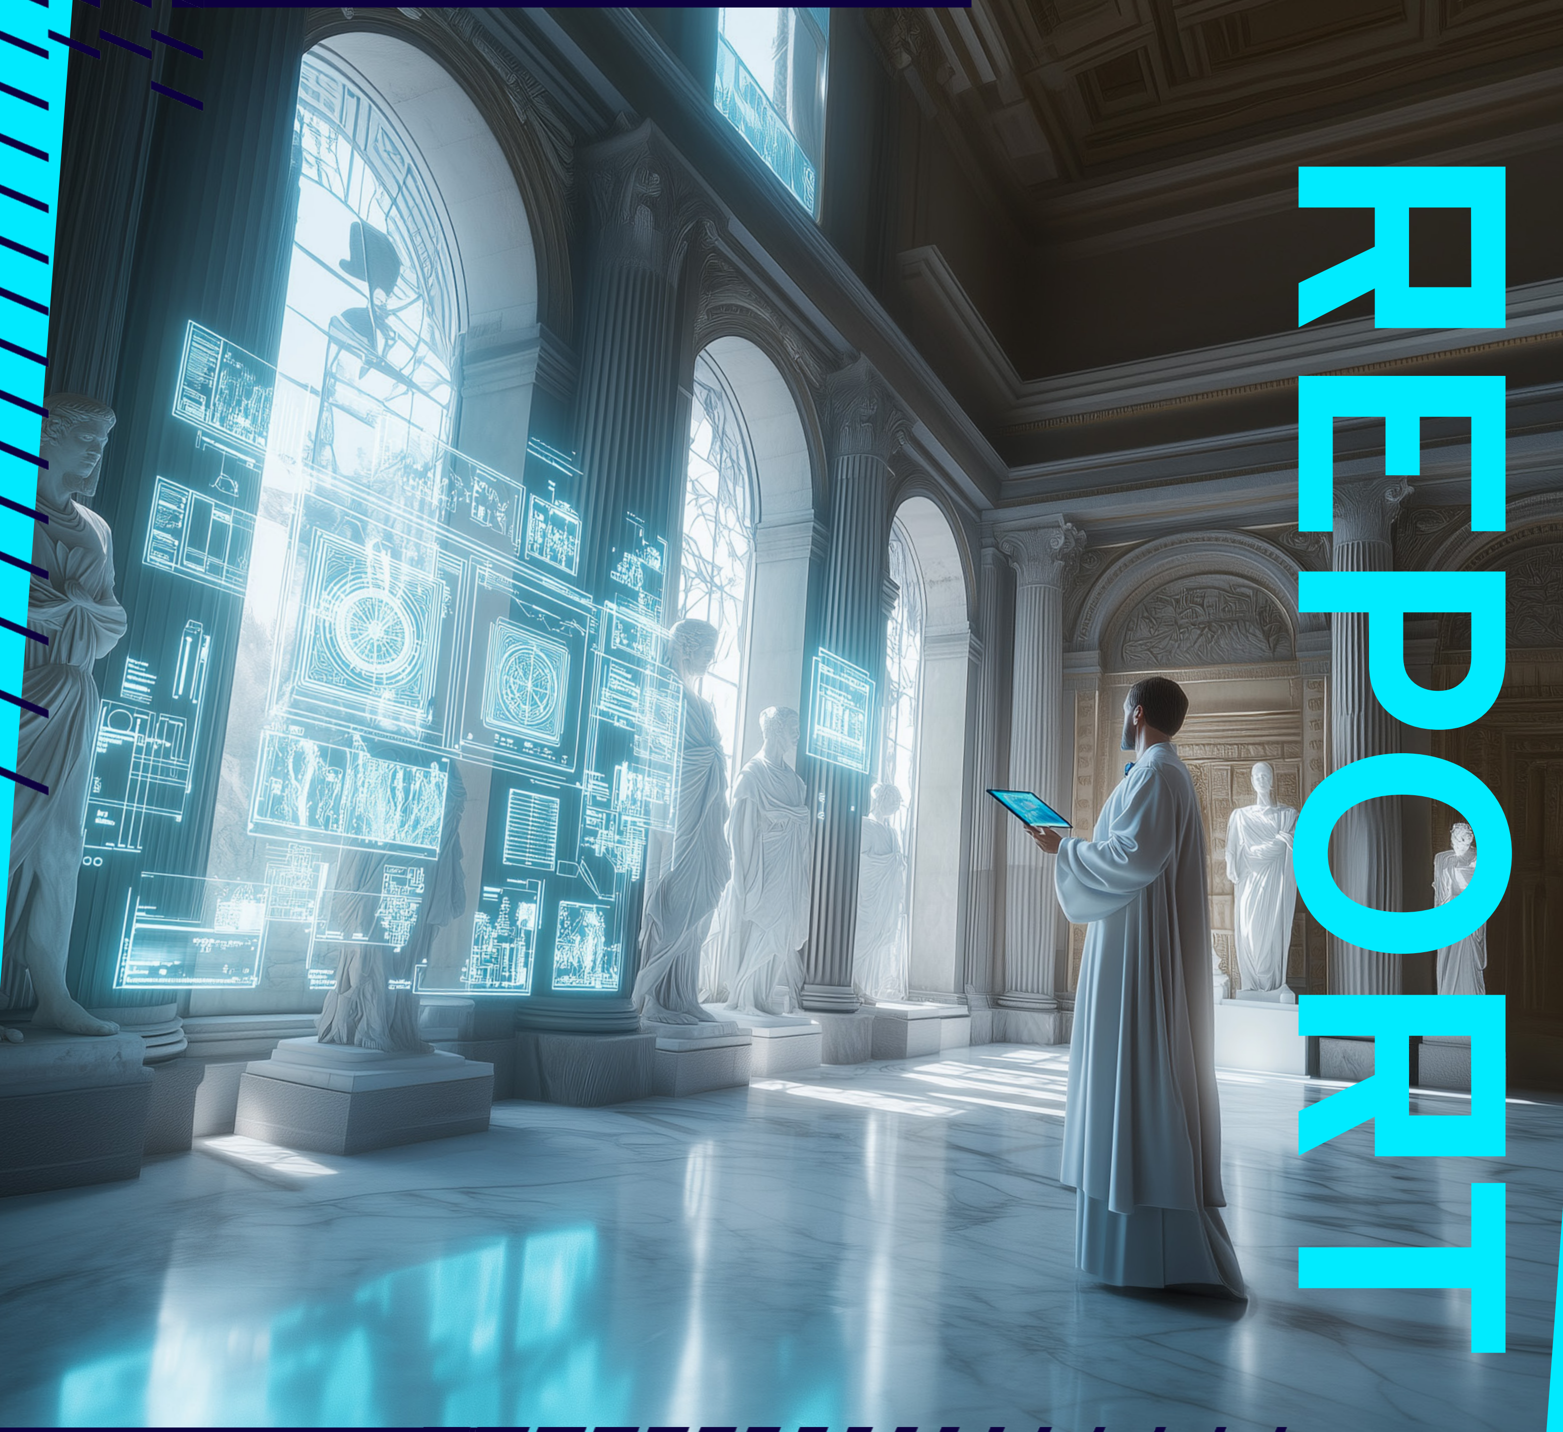

REPORT

BLUE SHIFT
BY ARTHUR D. LITTLE

"AI will accelerate scientific research far more than we can imagine."

— Joëlle Barral, Director of Fundamental Research in AI, Google DeepMind

Authors

Dr. Albert Meige, Director of Blue Shift, Arthur D. Little
Zoe Huczok, Project Leader of Blue Shift, Arthur D. Little
Arnaud Siraudin, Associate Director, Arthur D. Little
Donatello Fleck, Business Analyst, Arthur D. Little
Geoffroy Barruel, Consultant, Arthur D. Little

Contributors

Pierre Blouet, Research, Development, and Innovation Director, GRTgaz
Carole Caranta, Deputy Director General of Science and Innovation, INRAE
Paul-Joël Derian, Group VP Innovation and Sustainable Development, Avril
Philippe Mauguin, CEO, INRAE
Jean-Luc Moullet, Chief of Staff, French Minister of Higher Education and Research
Christophe Perthuisot, Chief Research & Innovation Officer, Moët Hennessy
Antoine Petit, President & CEO, CNRS
Rick Eagar, Partner Emeritus, Arthur D. Little

Expert–in–residence

Anne Bouverot, French President's Special Envoy on AI

CONTENT

Executive summary

Although AI has been used in specific research, development, and innovation (R&D&I) applications for at least a decade, it's been two years since the recent acceleration began, initiated by the availability of more powerful generative AI (GenAI) and large language models (LLMs). While there is a glut of information on potential applications, widespread integration of AI into R&D&I processes is still relatively immature. Applying AI to many R&D&I use cases poses significant challenges, especially where outcomes need to be error-free, as well as uncertainties in how AI will evolve regarding technology, economics, regulation, and societal acceptability.

This in-depth study was conducted by Arthur D. Little's (ADL's) Blue Shift in partnership with five major leading public and private sector organizations already using AI in their R&D&I efforts: LVMH, Avril, GRTgaz, the French National Centre for Scientific Research (CNRS), and the French National Research Institute for Agriculture, Food, and the Environment (INRAE). The study explored the current state of AI in R&D&I, the challenges and best practices, the landscape of solution providers, and future scenarios. We gathered evidence through 40+ interviews from AI providers, independent AI experts, and current best-in-class users of AI in R&D&I, as well as a survey with over 200 responses from private companies and public institutions that examined AI maturity, contributions, benefits, and barriers.

THE POTENTIAL OF AI IN R&D&I

Chapter 1: AI augments researchers' capabilities across all steps of the R&D&I process through various roles, helping to solve intractable problems and make decisions. No blanket model exists; data availability and problem type determine the best method. Most often, AI models are embedded in a systems of systems.

- **Benefits abound for AI.** AI augments researchers' capabilities, rather than replacing them, as part of a people-centric R&D&I effort. It helps solve intractable problems that researchers couldn't tackle before. It already acts as a knowledge manager, hypothesis generator, and assistant. The planner/thinker archetype, in which AI helps make decisions, is rapidly emerging.
- **AI-based models support use cases at every step of R&D&I process.** These range from technology and market intelligence to innovation strategy, ideation, portfolio and project management, IP management, ecosystem management, knowledge management, and new product/service launch and deployment.
- **There is no blanket model for R&D&I tasks.** Choosing among AI systems and other approaches depends on the type and amount of data available (e.g., more/less) and the nature of the question (e.g., open/closed). LLMs often play an orchestrating role in interfacing with or controlling other systems. AI is not always the answer. Classical science approaches, including traditional regression methods, may perform better on some problems. Most often, AI models are embedded in a systems of systems that also include human intervention.

Choosing among AI systems and other approaches depends on the type and amount of data available and the nature of the question.

HOW TO ENSURE SUCCESS

Chapter 2: Ensuring success in AI implementation for R&D&I requires agile methodologies, robust data foundations, strategic prioritization, analytical tradeoffs, scarce data science talent management, IT alignment, rapid benefit demonstration, and continuous monitoring.

- **Agile methodologies.** Agile methodologies that move fast and iteratively are preferable for AI project development, given the speed at which technology evolves. Such approaches ensure that some benefits can be obtained early, even if "perfect" solutions are still some way off.
- **Robust data foundations.** Data collection hygiene, storage, security, and governance are central to realizing AI benefits. New techniques for processing poorly structured or smaller datasets are becoming more important. Ensuring wide data accessibility, cross-organizational collaboration, and effective data governance is also fundamental.
- **Strategic prioritization.** Organizations must choose strategically between making, buying off-the-shelf, or fine-tuning AI models. Most core R&D&I problems lend themselves to fine-tuning existing open source models, whether LLMs, generative adversarial networks (GANs), diffusion models, or reinforcement learning (RL).
- **Analytical tradeoffs.** Tradeoffs must be carefully considered in proof-of-concept (POC) development, including acquiring versus synthesizing data, optimizing for precision versus recall, and underfitting versus overfitting data.
- **Scarce data science talent management.** The right organizational solution for accessing expert data science resources depends on needs — different pros and cons come with using external resources, training internal experts, creating a central AI service center, or embedding data scientists in R&D teams.
- **Alignment with IT.** R&D&I functions need to align with IT departments to address security and compliance requirements while maintaining the speed needed for experimentation.
- **Rapid benefit demonstration.** Prioritizing AI use cases and demonstrating benefits quickly helps prevent POCs from stalling.
- **Continuous monitoring and improvement.** These are especially important for experimental AI models, as their performance can change over time.

TOOLS & PROVIDERS

Chapter 3: The value chain for AI in R&D&I heavily relies on major open source models, but smaller players also form a key part of the ecosystem. Applications tailored for every part of the R&D&I process exist, as do start-ups targeting vertical-specific problems. Hosting providers also offer inference as a service.

- **Key role of open source.** The value chain for AI in R&D&I can be divided into three layers: (1) infrastructure (compute), (2) model development, and (3) R&D&I applications. Open source models are the backbone across the whole chain, developed and trained by major players such as Meta (Llama), Microsoft (Phi), and Nvidia (NVLM, TensorFlow, StyleGAN). Smaller players like Mistral and Cohere (Aya 23) and academic institutions such as CNRS and GENCI (BLOOM) also contribute significantly to the open source ecosystem. Collaboration is fostered through forums that encourage co-creation and sharing of fine-tuning tools, such as Hugging Face.
- **Start-ups and players.** AI use cases exist for every building block of the R&D&I process, from strategy and intelligence to ideation, portfolio/project management, IP management, knowledge management, ecosystem management, and deployment. The applications market also includes vertical-specific start-ups targeting scientific, research, and innovation problems, especially in life sciences.
- **Inference as a service** — a promising service. Various hosting providers also offer an inference-as-a-service model, which consists of hosting the compute power needed for inferences run by the model (e.g., each LLM query) in the cloud to help customers avoid high computational costs.

NAVIGATING THE FUTURE

Chapter 4: How AI in R&D&I will evolve depends on the outcomes of three main factors: performance, trust, and affordability. These lead to six plausible future scenarios on a spectrum between AI transforming every aspect of R&D&I at one end and being used only in selective, low-risk use cases at the other.

We identified 16 trends shaping the future, divided into three main factors:

1. **Performance** — whether AI will meet the high bar necessary for many R&D&I problems
2. **Trust** — the extent to which researchers, developers, customers, and the public will trust and accept AI-generated outputs
3. **Affordability** — how far AI implementation will be constrained by costs, skills, resources, and environmental impacts

These factors lead to six plausible scenarios across two ends:

1. **Cheap & Nasty.** At one end of the spectrum is the Cheap & Nasty scenario reflecting low performance and trust but high affordability. In this scenario, AI is only used in select low-risk use cases with strict vetting, curtailing productivity gains.
2. **Blockbuster.** At the other end of the spectrum is the Blockbuster scenario, which reflects high performance, trust, and affordability.

In between are scenarios reflecting other combinations, each with different consequences regarding day-to-day R&D&I work, organizational evolution, and winners and losers. Recognizing these scenarios is important for R&D&I organizations to chart a way forward.

STRATEGIC ACTIONS

Chapter 5: We recommend six no-regret moves for organizations regardless of the six future scenarios. These comprise mutualizing compute power, encouraging data sharing, managing AI talent, training the workforce in AI fundamentals, resetting data and AI governance approaches, and improving output controls. Beyond these, organizations should take measured strategic bets aligned with corporate objectives.

- **Six no-regret moves** that R&D&I organizations should take to shape up for the AI future, irrespective of how the scenarios develop:

  1. **Manage and empower talent.** Correctly access external resources and build in-house capabilities. Given that low-code/no-code (LCNC) AI solutions are becoming increasingly prevalent, data engineering profiles may become more important than data science profiles.
  2. **Control AI-generated content.** Scale-up quality and IP control systems for AI-generated content and data.
  3. **Build up data sharing and collaboration.** Pursue data-sharing initiatives, looking for mutual benefits and ways to overcome competitive concerns.
  4. **Train for the long run.** Continuously deliver AI training to as broad an audience as possible, beyond immediate users. This should touch on AI technical fundamentals, functional capabilities, implementation requirements, and potential risks.
  5. **Rethink organization and governance beyond IT.** Set up a suitable governance system for AI. This means allocating clear AI and data governance accountabilities and responsibilities to key individuals within the overall enterprise and digital teams beyond the IT function. For critical AI use cases, centralization of governance may be desirable to overcome coordination processes, reporting directly to executive leadership.
  6. **Mutualize compute resources.** Look for opportunities to increase the affordability of compute power through mutualization with partners.

- **Strategic bets to consider.** In addition to no-regret moves, R&D&I organizations should consider the scope, costs, and benefits of AI use cases; prioritize efforts; and revisit the overall innovation strategy to suit an AI future. Portfolio logic should be applied to ensure a suitable balance of:

  - Hedging AI moves to ensure rapid response in case of future disruption
  - Speculative high-risk/high-reward AI opportunities
  - Shorting in future AI areas that could be especially susceptible to poor performance, low levels of trust, or excessive costs.

BLUE SHIFT
BY ARTHUR D. LITTLE

BIG ideas BOLD blue shifts

1

AI will power all steps of R&D&I — including those involving creativity

AI's impact is not just hype — it's impacting both the productivity and the creativity dimensions of R&D&I. Companies that quietly leverage AI using general-purpose LLMs and smaller specialized models are already seeing 10x productivity gains in some situations.

2

AI as the orchestrator, not the solo player

AI should serve as a coordinator between diverse digital tools, such as simulation, good old-fashioned AI (GOFAI), GenAI, graphs, rules and heuristics, and Bayesian networks, while also keeping the human in the loop. Automated agents will empower researchers to run entire workflows autonomously, speeding up discovery.

3

Maturity gap in AI adoption — not a barrier but an opportunity

Most R&D&I organizations are still new to AI, with many researchers unaware of its current and future impact.

4

Focus on solving problems, not just deploying AI

The focus shouldn't be on flashy AI tools but on using AI to solve specific, high-impact problems. Defining the right problems will ensure that AI is a tool for innovation rather than just a trend.

5

Leverage LLMs for productivity use cases

Fine-tuned LLMs deliver high value even though we don't fully understand their workings at scale. LLMs are particularly interesting for productivity use cases. While cross-silo data integration could unlock even greater potential, methods such as Low-Rank Adaptation (LORA) make fine-tuning affordable and effective today.

6

Smaller models for bigger creative breakthroughs

Smaller and more specialized AI models or other approaches, such as Bayesian networks, will increasingly excel in solving complex R&D&I problems. These tailored models are more effective in certain areas.

7

Data is the game changer, not algorithms

Data management will be the differentiator in the AI-driven future as algorithms become commoditized. Centralized, structured, and cleaned data will be the foundation for building competitive R&D&I systems. Preparing data for the first POC may take time (up to 18-24 months), but it will speed up with each iteration.

8

Trust is everything — build it carefully and maintain it diligently

Building trust in AI systems is critical and fragile. In R&D&I, where the stakes are high and outputs aren't immediately tangible, ensuring transparent processes and human oversight is essential to avoid setbacks in AI adoption.

9

AI talent shortage — the race to upskill R&D teams

The supply of AI talent will lag behind demand until 2030, making upskilling of existing R&D teams crucial. Organizations that invest early in training their talent will avoid falling behind in the AI race.

10

Inference as a service — a paradigm shift for product development

Like cloud computing transformed IT infrastructure, inference as a service will revolutionize how companies develop and scale AI-driven products. This model will be key in democratizing AI and fostering new business models.

Preamble

October 2018: Christie's auction house sells a work created by the French artist collective Obvious in collaboration with AI for nearly half a million dollars.

December 2023: A group of researchers reports in *Nature* that they have solved a previously unsolved mathematics problem by collaborating with GenAI.

June 2024: Researchers at the University of Pennsylvania leverage machine learning (ML) to analyze microbial dark matter, uncovering nearly 1 million potential antibiotic compounds. Published in *Cell*, this breakthrough accelerates the discovery of new antibiotics, with dozens showing activity against antibiotic-resistant bacteria, compressing years of research into mere hours using AI.

A series of ripples in the pond of creativity.

If AI can foster creativity, it could revolutionize problem solving across various sectors, including R&D&I.

As businesses and decision-makers, we believe it's only natural to ask whether AI has truly become creative. Or perhaps more importantly, can it help us humans become more creative? This question is crucial, as AI has been primarily focused on improving performance and productivity over the past 10-15 years. If AI can indeed foster creativity, it could revolutionize problem solving across various sectors, including R&D&I.

The question becomes even more relevant following a landmark event in October 2024: two Nobel Prizes were awarded for discoveries related to ML.

In chemistry, Demis Hassabis and John Jumper from DeepMind received the Nobel Prize for AlphaFold2, a model capable of predicting the structure of 200 million proteins. Since its release in 2020, over 2 million researchers worldwide have already used this deep learning model. This breakthrough has revolutionized drug discovery, demonstrating how AI can directly fuel innovation.

In physics, John Hopfield and Geoffrey Hinton were awarded the Nobel Prize for their contributions to AI, specifically ML. They applied statistical physics principles to develop foundational neural network models, including the Hopfield network and the Boltzmann machine. Their work laid the groundwork for modern ML, allowing AI systems to learn and recognize patterns from complex data, revolutionizing the field and demonstrating how science fuels AI.

In this watershed moment for science and research, one might wonder: how do professions defined by human intelligence adapt to the rise of artificial intelligence? What role does AI play in the "Eureka!" moment and the countless tasks that define R&D&I? What specific challenges do organizations face in areas such as problem definition, data availability, security, system interpretability, and costs — and what are the best practices to address them? Furthermore, how will AI's uncertain future affect the landscape of R&D&I?

As a teaser, here's an intriguing anagram of the Report's title:

"Eureka! On steroids"

also transforms into

"Edison treasure: OK!"

Anagrams move in mysterious ways, but this one feels particularly fitting. Thomas Edison, a great innovator, stands as a symbol of discovery. This anagram hints that AI may uncover valuable treasures of innovation in the future.

— Albert Meige, PhD

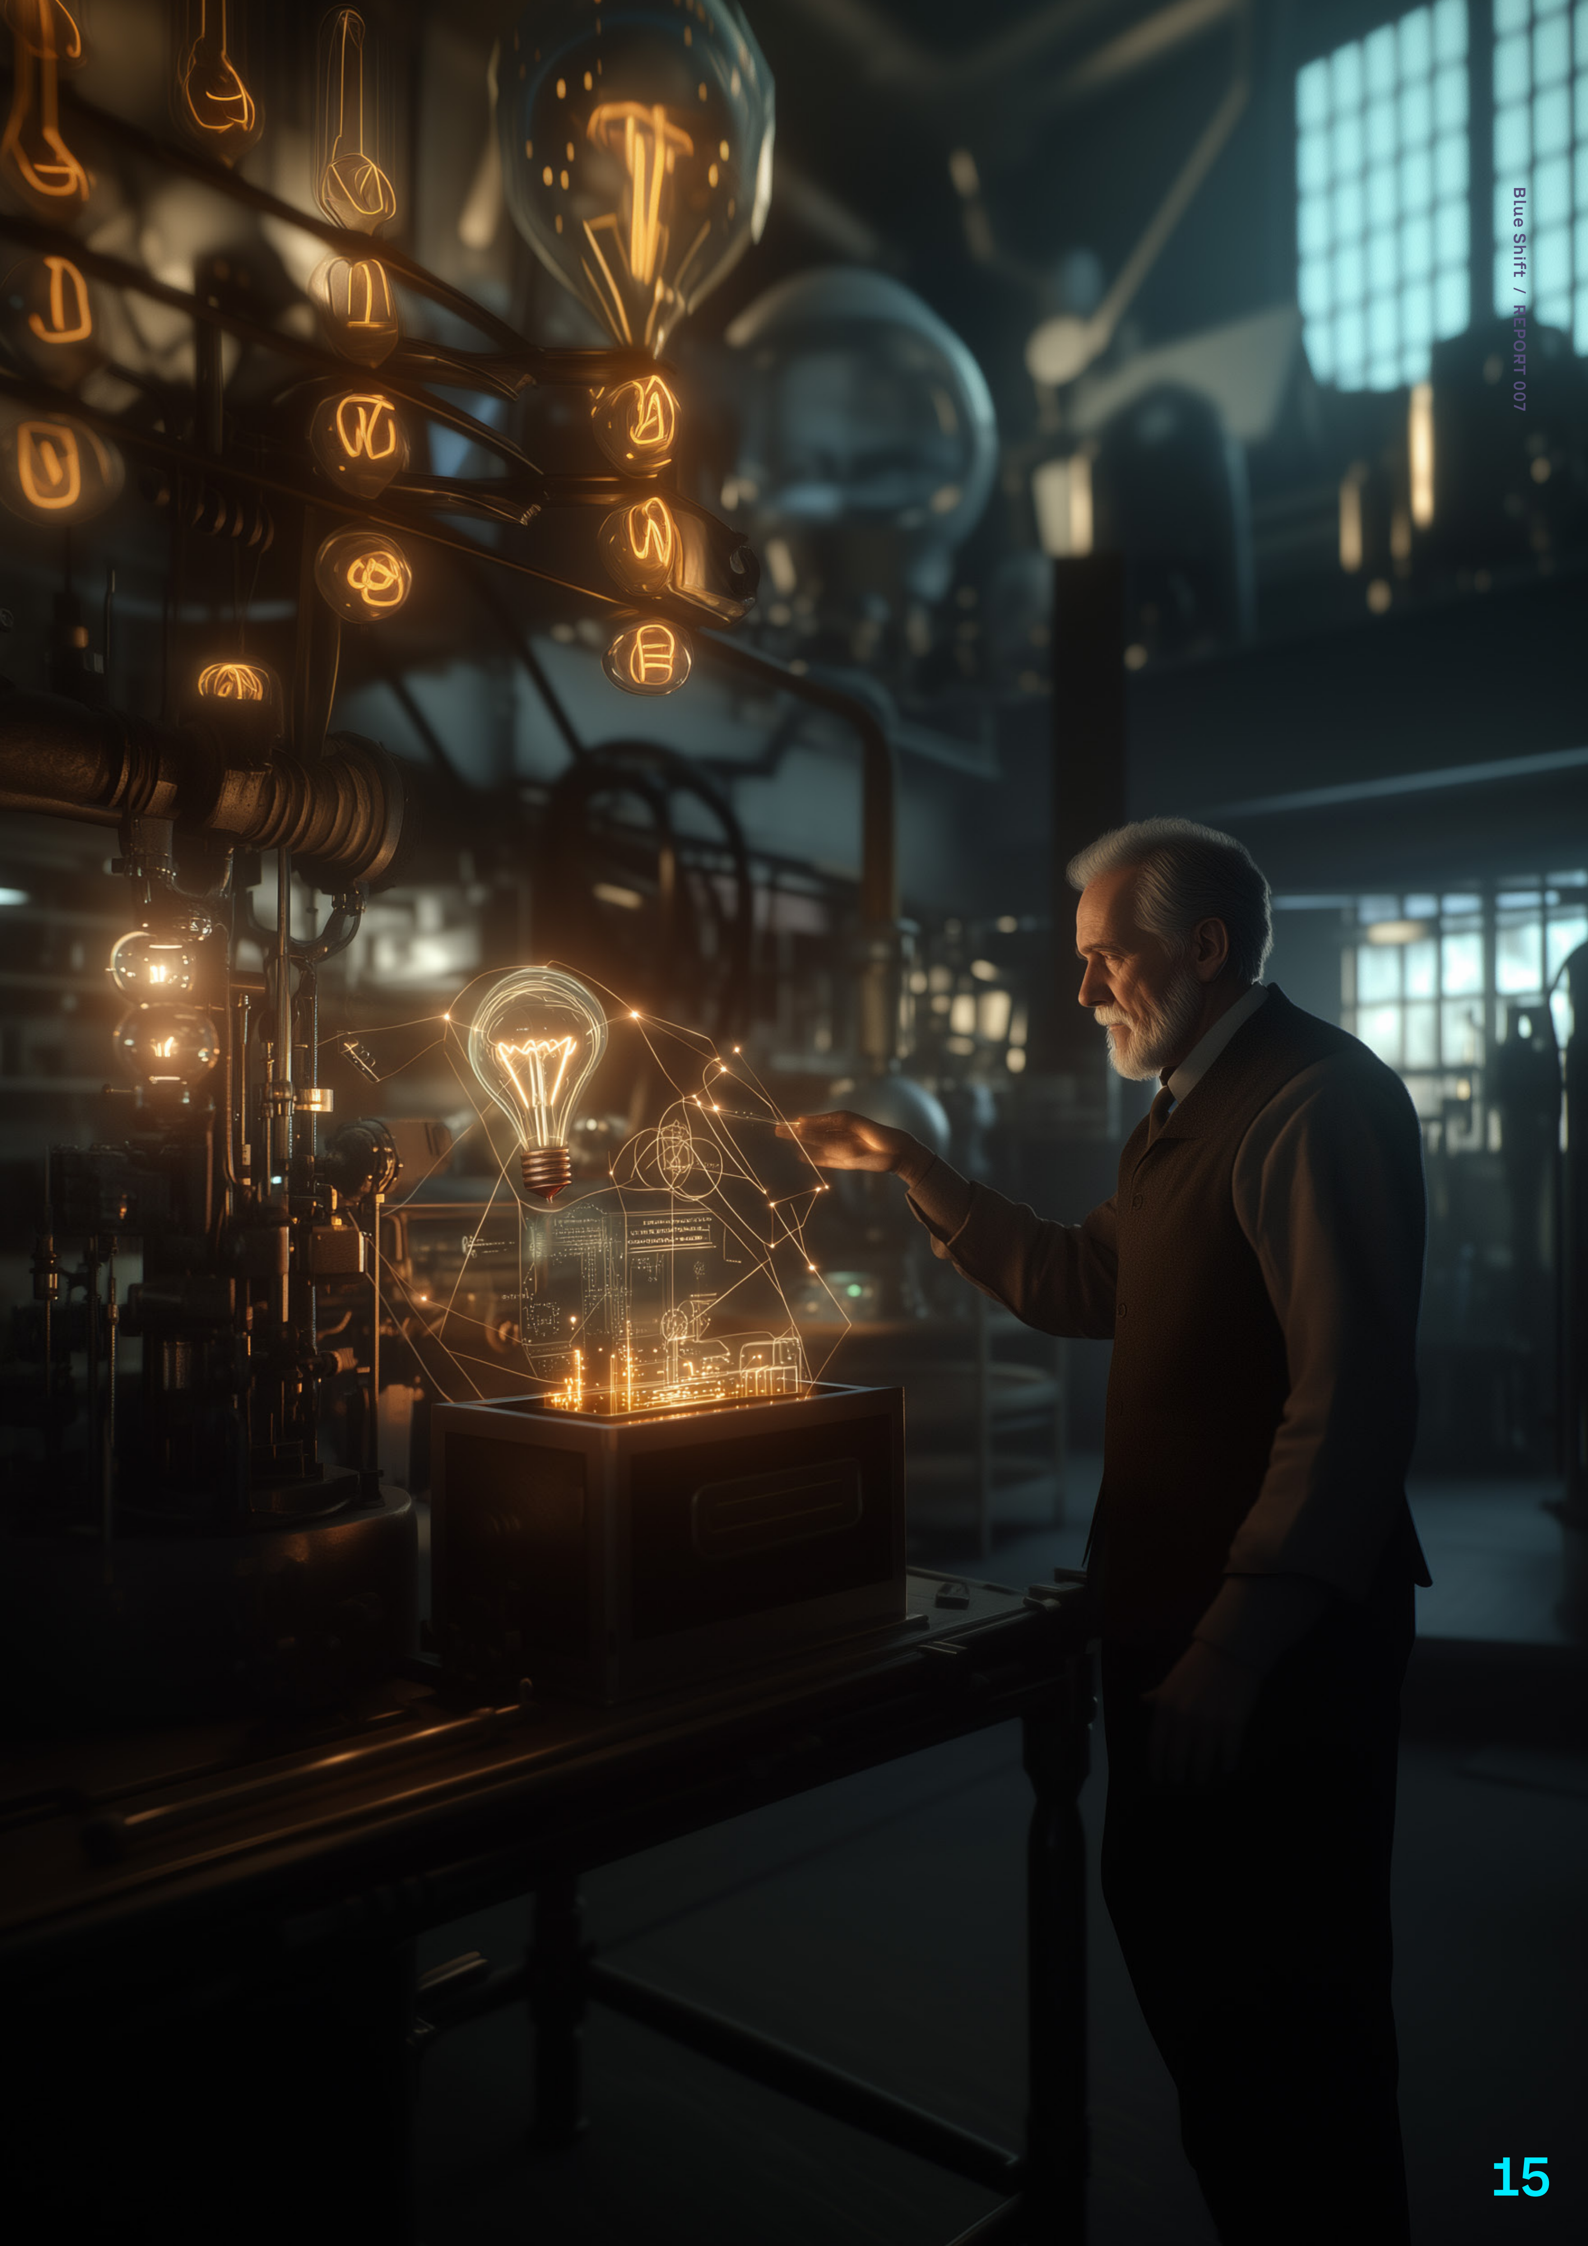

CHAPTER 1

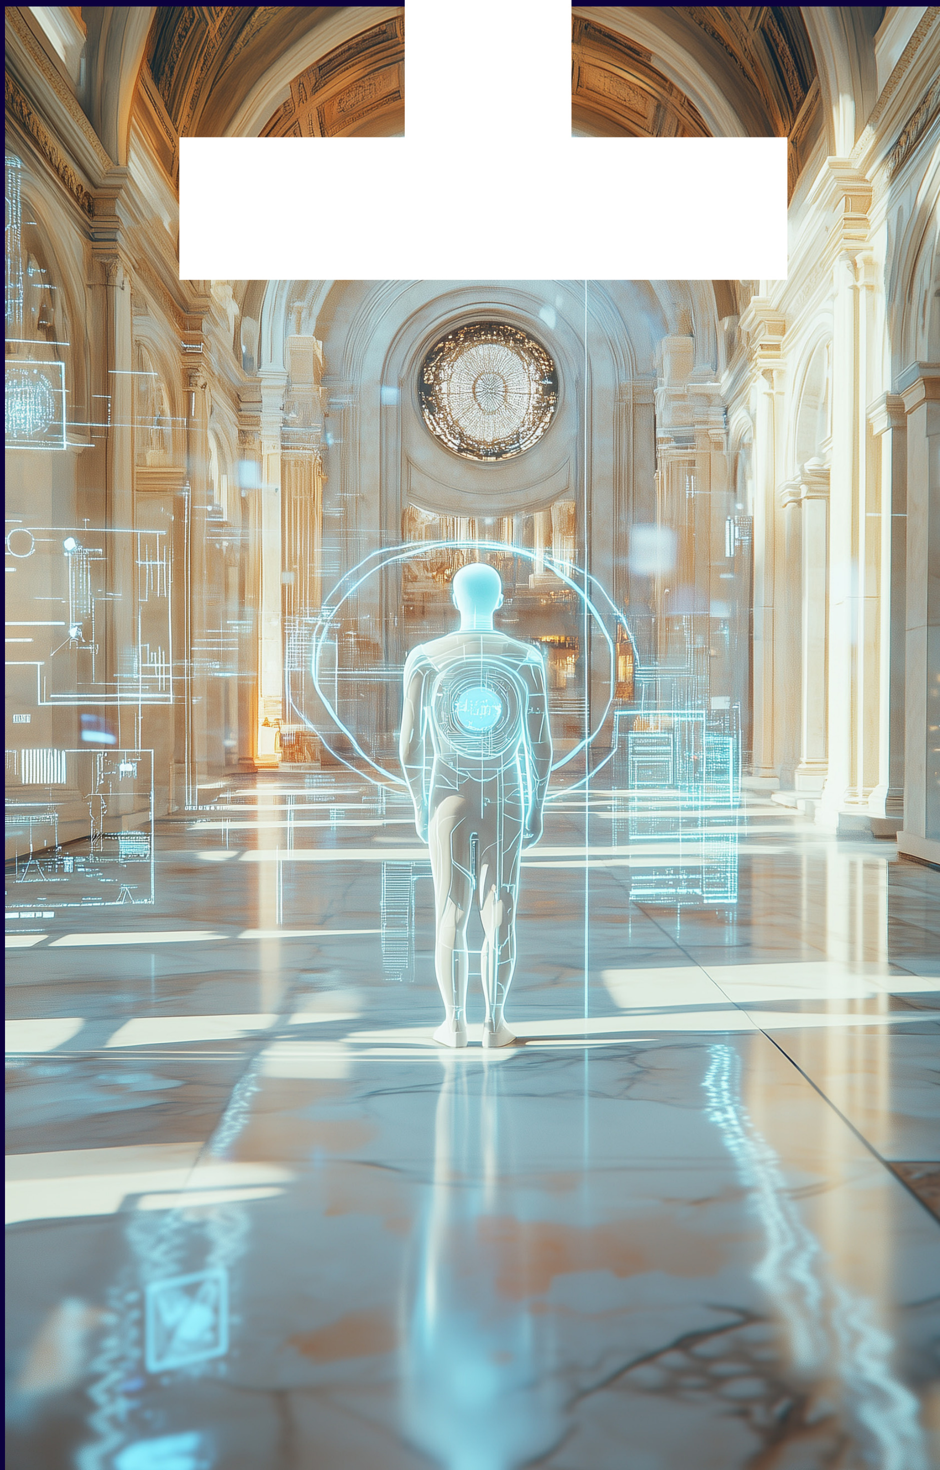

THE POTENTIAL OF AI IN R&D&I

1 THE POTENTIAL OF AI IN R&D&I

AI augments researchers' capabilities across all steps of the R&D&I process through various roles, helping to solve intractable problems and make decisions. No blanket model exists; data availability and problem type determine the best method. Most often, AI models are embedded in a systems of systems.

THE BENEFITS OF AI IN R&D&I

Every building block of R&D&I has benefits and use cases, from technology and market intelligence to innovation strategy, ideation, portfolio and project management, and IP management. When we look to understand these benefits, three key factors emerge:

1. AI augments researchers, rather than replacing them

In the 40-plus interviews for this Report, none of the decision makers was looking to replace their R&D&I workforce with AI, now or in the future. Currently, AI doesn't operate fully autonomously in any use case deployed in our extensive survey sample. Instead, it augments researchers, frees up their time, and enables them to be more productive and creative, often by automating previously manual tasks. Particularly since the advent of GenAI, researchers have been able to automate repetitive tasks, such as drafting emails or documents and synthesizing the contents of multiple papers. For example, in the case of the food company Roquette, researchers and developers saw their jobs evolve from operating lab machines to operating AI that operates machines and learning how to handle data. AI works best when it is applied in the context of a "people-centric" lab.[1]

"AI does not replace researchers but augments them by automating tasks to free up time for innovation and combining their expertise with up-to-date knowledge."

Johan Aubert, Chief Technology & Digital Officer, L'Oréal

2. AI helps solve intractable problems

Equipped with AI, researchers can solve problems they couldn't before because of the technology's speed and ability to scale and learn. For instance, to optimize nutrition plans, agro-industrial group Avril developed a model to process historical data that was unusable without AI. Google DeepMind created AlphaFold, an AI model that could examine millions of protein combinations, enabling the discovery of proteins in novel fields. Without AI, neither these nor many other use cases would be possible.

"We had old data, but it was hard to get something out of it. With AI and the access to this data it provided, we trained models to identify characteristics of unreliable assets."

Florent Brissaud, R&D Project Manager for AI Applications, GRTgaz

3. AI will assume a "planner–thinker" position

AI embodies an increasing range of intelligence archetypes and is moving beyond content generation and search to cover more complex roles. These include becoming a knowledge manager, hypothesis generator, and assistant to R&D&I teams. Today, the "planner-thinker" archetype is emerging. For example, a military organization has seen AI usage help make decisions, basing recommendations on weak signals from various sources. However, progress needs to be made before fully integrated workflows emerge in companies/institutions to enable this type of decision-making.

"Every organization should ask themselves how do we make our analysts evolve? And it will certainly not stop at the analyst level."

Erez Raanan, CEO, Mathlabs

AI-BASED MODELS SUPPORT USE CASES AT EVERY STEP OF R&D&I PROCESS

ADL's Innovation Excellence Model is a proven framework for innovation management within organizations (see Figure 1). It contains 10 elements that are essential to strong innovation performance. AI can currently augment and enable all of these, albeit at different maturity levels.

NO BLANKET MODEL FOR R&D&I TASKS

When deciding whether to use AI to solve a specific R&D&I use case — and which AI approach will give the best results — organizations need to focus on two factors, as shown in Figure 2:

1. The type and amount of data available (from a little to a lot)
2. The nature of the question being asked (from open to specific)

Figure 1. ADL Innovation Excellence Model

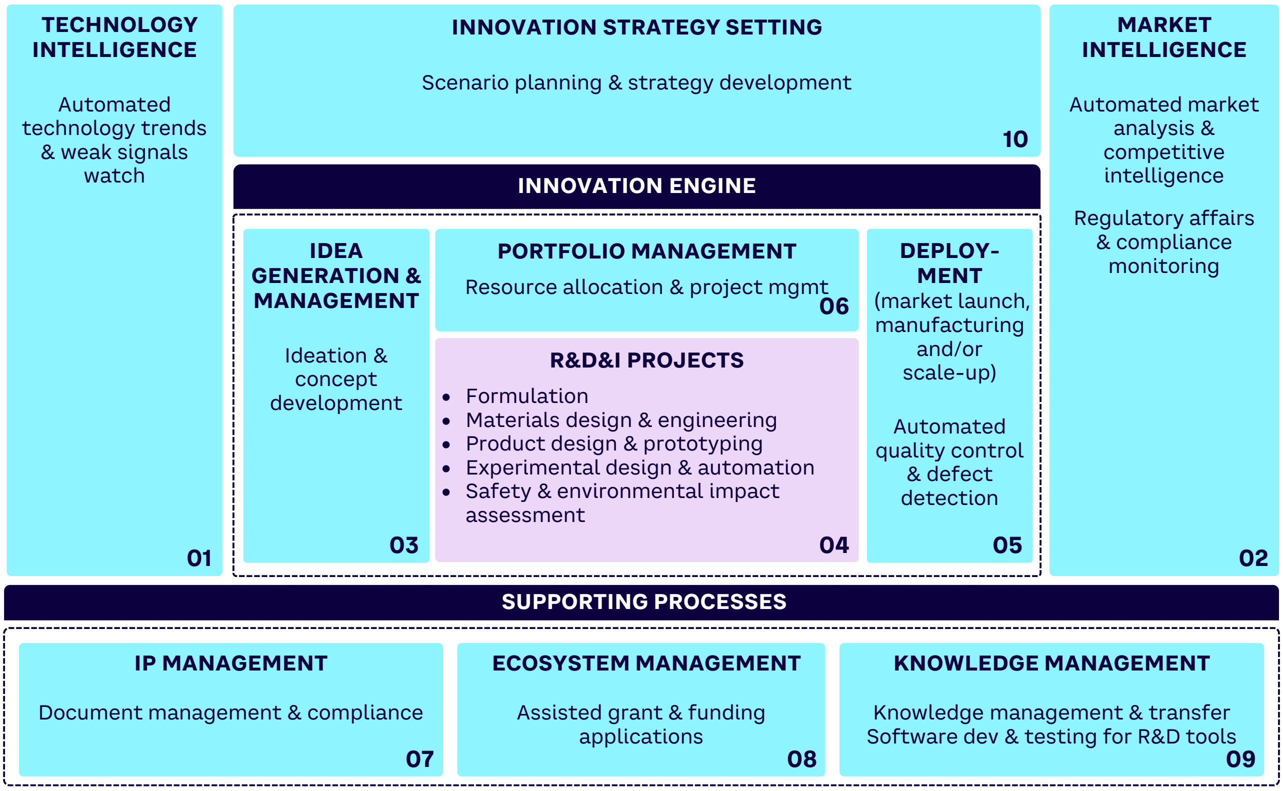

Source: Arthur D. Little

Figure 2. AI systems in relation to other methods by problem type

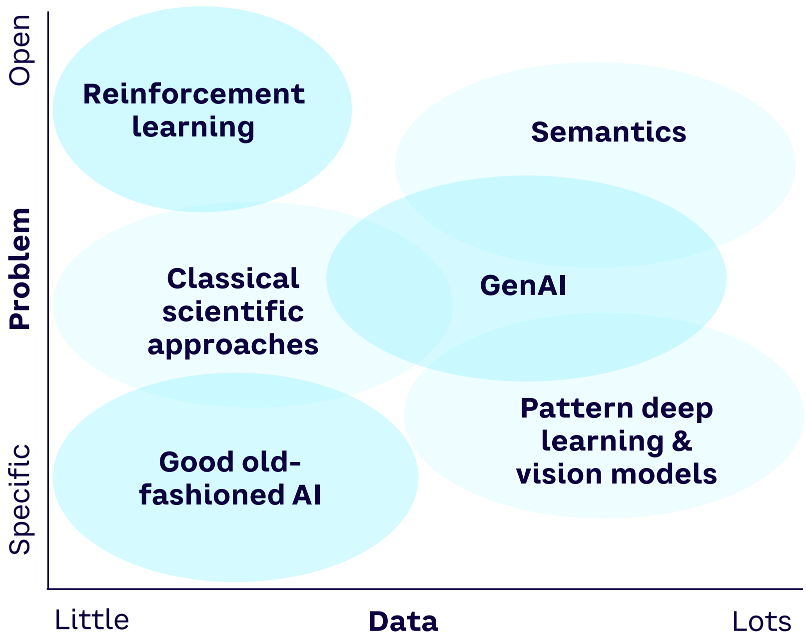

Source: Arthur D. Little, Yves Caseau, National Academy of Technologies of France (NATF)

"We have been running AI models on our super computers for decades, and there is clearly an acceleration since 2022. We also have a deep expertise in exact science models (neutronics, physics, fluids, etc.): after considering a problem from every angle, if there is an exact science model available, we use it most of the time because it usually offers better results and predictability. But for the other cases, such as documents or language analysis, or for having a first guess of a solution (that will need to be confirmed by exact science models), GenAI can also bring very interesting results if you fully master its limitations. Even before the AI Act, our approach was to control AI results and always have human operators or engineers doing the final decision."

Vincent Champain, Senior Executive VP, Digital Performance & IT, Framatome

Approaches fall under the following categories:

- **Good old-fashioned AI** — rule-based systems for specific, structured problems with minimal data
- **Classical scientific approaches** — standard scientific methods that rely on experimentation and validation
- **Pattern deep learning and vision models** — pattern recognition and vision tasks requiring extensive labeled data (e.g., convolutional neural networks)
- **Semantics** — structured knowledge representation for understanding context and meaning (e.g., knowledge graphs, ontologies)
- **GenAI** — Large models handling open-ended, data-rich tasks, often with human feedback (LLM + RL from human feedback [RLHF])
- **Reinforcement learning** — used for exploring open-ended problems, relying on simulation and relatively limited data

Figure 3 offers guidance in selecting the best approach. However, AI is not always the answer — classical science techniques, including traditional regression methods, may perform better on some problems.

Tasks essential to the core of R&D&I have no blanket model; specific models and systems deliver optimal results for specific problems and data.

Figure 3. AI/ML architectures for R&D&I problems

NON-EXHAUSTIVE

	TASK TYPE / INPUT DATA TYPE	**Sequential data** (e.g., language, time series data)	**Non-sequential data** with spatial locality (e.g., images, grids)	**Non-structured data** (not sequences or grids) (e.g., molecules, social networks)	**Discrete states & actions** (e.g., game moves)	**Continuous states & actions** (e.g., robot moves)	**Causally linked data** (e.g., treatment effects, incident costs)
More specific	**Classifi-cation/ detection**	**Recurrent Neural Network (RNN), Transformer** (e.g., BERT)	**Convolutional Neural Network (CNN)** (e.g., AlexNet), **K-Nearest Neighbors (KNN), Random Forest**	**Graph Neural Network (GNN)** (e.g., GraphSAGE)	NA	NA	NA
↕	**Prediction**	**RNN**, incl. long-/ short-term memory algorithms (e.g., Neural Hydrology)	**CNN, Transformer, Variational Autoencoder (VAE)**	**GNN** (e.g., Attentive FP)	NA	NA	**Simulation (incl. digital twins)** + any relevant ML prediction method, **Bayesian networks, Causal Forests**
	Controls	Command & control algorithm + **RNN**	Command & control algorithm + **CNN**	Command & control algorithm + **GNN**	**Model-free RL (MFRL)** (e.g., AlphaGo)	**MFRL** (e.g., Q-learning), **Model-based RL** (e.g., DemoStart)	NA
More open	**Generation**	**RNN** (e.g., Google Smart Compose), **Transformers** (e.g., GPT-4)	**Diffusion** (e.g., Stable Diffusion), **GAN** (e.g., StyleGAN), **VAE + GAN** (GANverse3D)	**VAE + GNN, GAN**	**Simulation + Monte Carlo Tree Search (MCTS)**	**Simulation + MCTS**	NA

Note: Map focuses on machine learning and causal inference methods; it does not focus on symbolic approaches ("good old-fashioned AI")
Source: Arthur D. Little, Yves Caseau

System of systems

At the same time, a single AI approach may not deliver optimal results — most state-of-the-art intelligent systems produced in the past 15 years have been systems of systems. These are independent AI systems, models, or algorithms designed for specific tasks, which, when combined, offer greater functionality and performance. For example, most LLM chatbots, such as ChatGPT, use a transformer architecture coupled with RL from human feedback (RLHF). Image generation model DALL-E 2 brings together an auto-regressive transformer and a diffusion algorithm. LLMs can often play an orchestrating role in interfacing with or controlling other systems because of their language fluency. Robotics use cases often require a system-of-systems approach, as set out in Figure 4 for manufacturing robotics. A computer vision algorithm can process the visual data gathered by the robot's camera; a symbolic, rule-based engine can inform the decisions of the cognitive controller, which, in turn, can be influenced by an RL algorithm that learns from past experiences or other external knowledge.

Figure 4. System of systems example, including AI in robotics for manufacturing

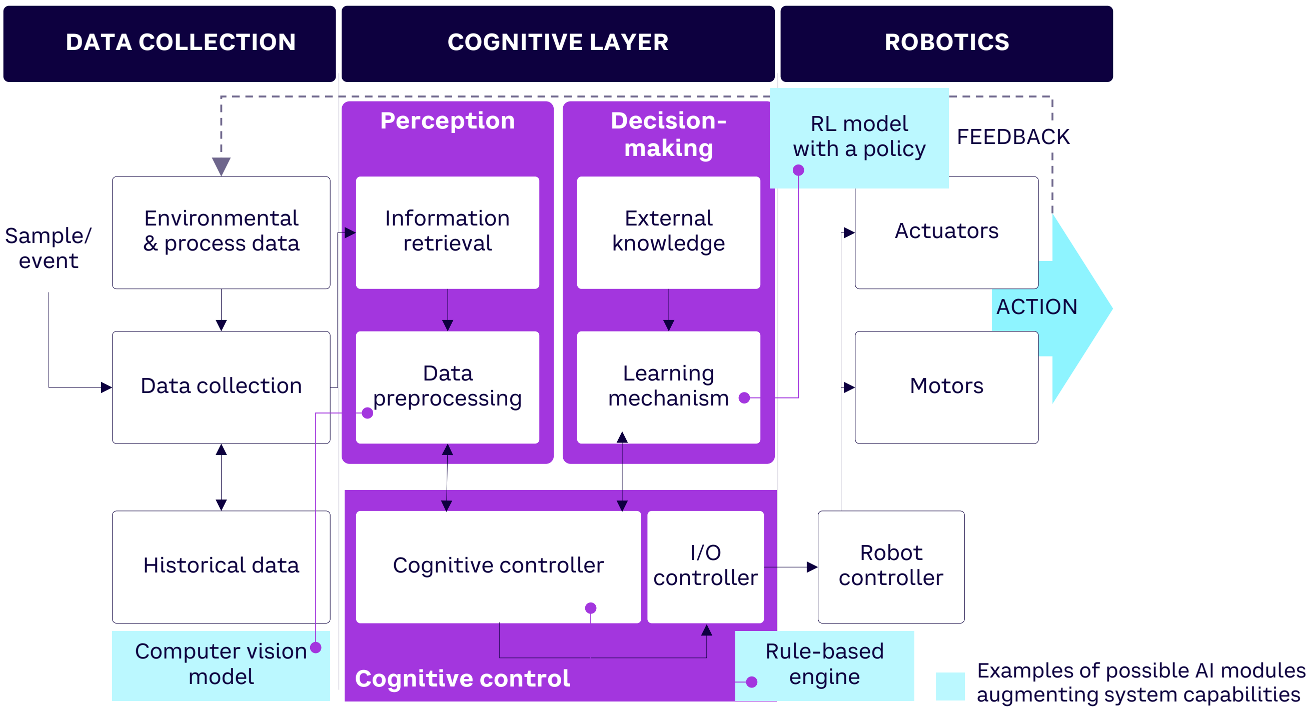

Source: Arthur D. Little; Oliff, Harley, et al. "A Framework of Integrating Knowledge of Human Factors to Facilitate HMI and Collaboration in Intelligent Manufacturing." *Procedia CIRP*, Vol. 72, 2018.

“The measure of intelligence is the ability to change.”

— Albert Einstein

CHAPTER 2

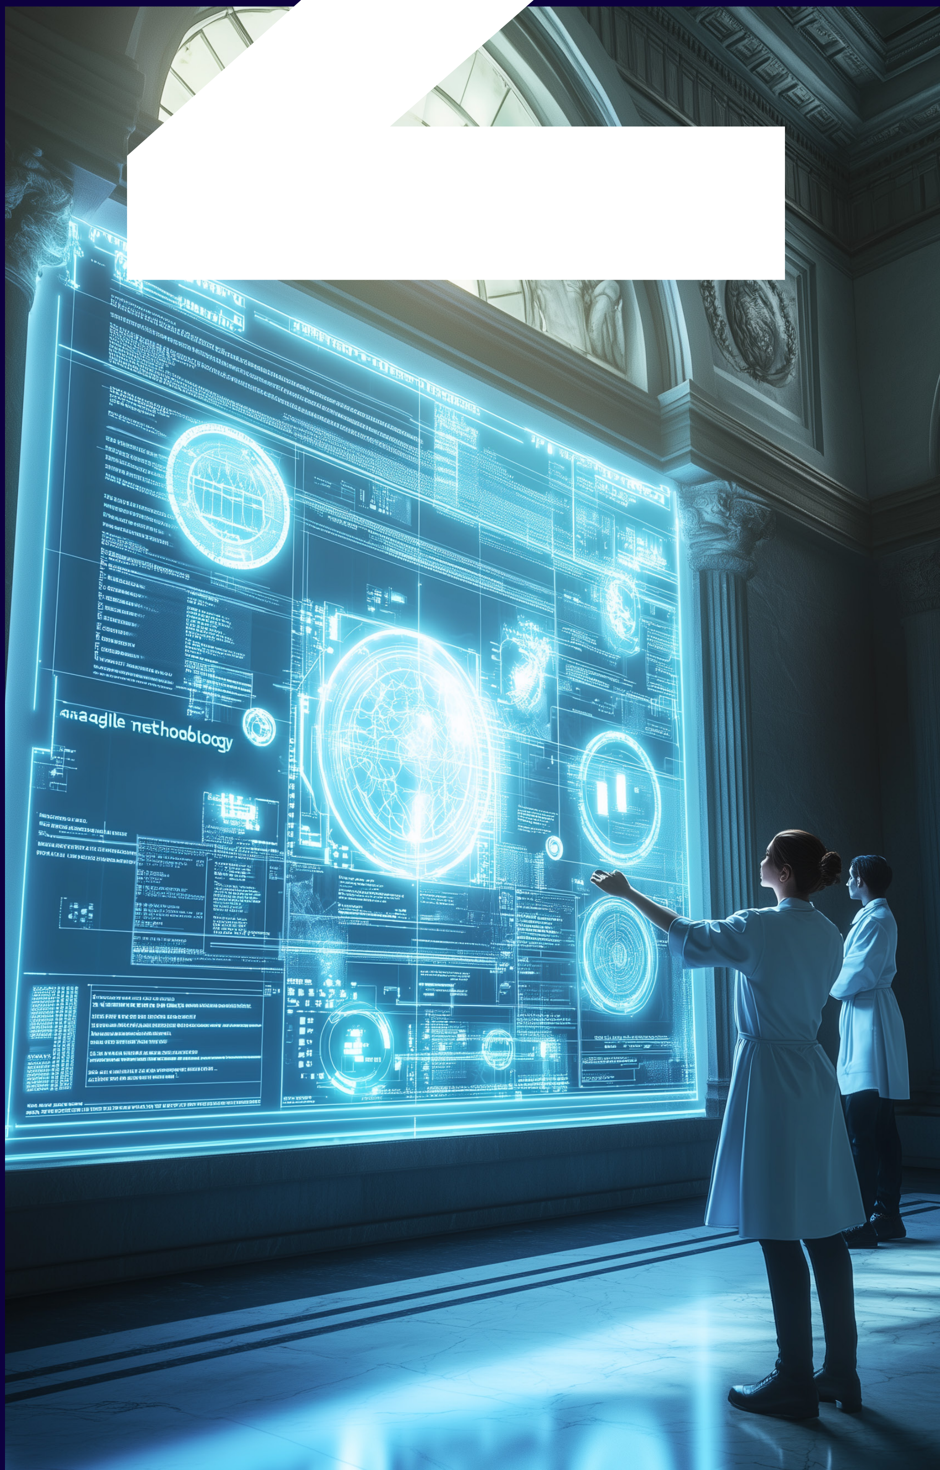

HOW TO ENSURE SUCCESS

2 HOW TO ENSURE SUCCESS

Ensuring success in the implementation of AI for R&D&I requires agile methodologies, robust data foundations, strategic prioritization, analytical tradeoffs, scarce data science talent management, IT alignment, rapid benefit demonstration, and continuous monitoring.

To gain stronger knowledge of good practices in R&D&I, we conducted 40-plus interviews with researchers; AI scientists; founders; and heads of R&D in digital, manufacturing, marketing, and R&D teams. Respondents are from 57 companies and three research institutions with interviews taking place June 2024–October 2024. Companies are headquartered on three continents and span the defense, chemicals, automotive, software development, and consumer goods sectors with revenues of US $10,000 to over US $1 billion. AI projects range from a few weeks to multiple years in length.

"Move fast and don't wait for the perfect technology. AI will evolve; it is better to have something that was released on the market two years ago than nothing at all."

Denis Gardin, Innovation Director, MBDA

THEMES UNDERPINNING GOOD AI PRACTICE

We have distilled these interviews into eight good practices (agile methodologies, robust data foundations, strategic prioritization, analytical tradeoffs, scarce data science talent management, IT alignment, rapid benefit demonstration, and continuous monitoring) across four categories (data and project management, strategic implementation, organizational structures, and sustained adoption and impact) that help maximize AI success in R&D&I.

Establish strong data & project management foundations

Adopt agile methodologies

AI is a new, evolving discipline, meaning projects will likely change as they progress. Thus, R&D&I teams need to apply agile methodologies to work quickly and iteratively — something that is all the more necessary given the rapid pace of technological progress in AI. Successful projects are likely to follow the agile development best practices set out in Figure 5 from project launch onward.

Figure 5. Agile development best practices for AI projects

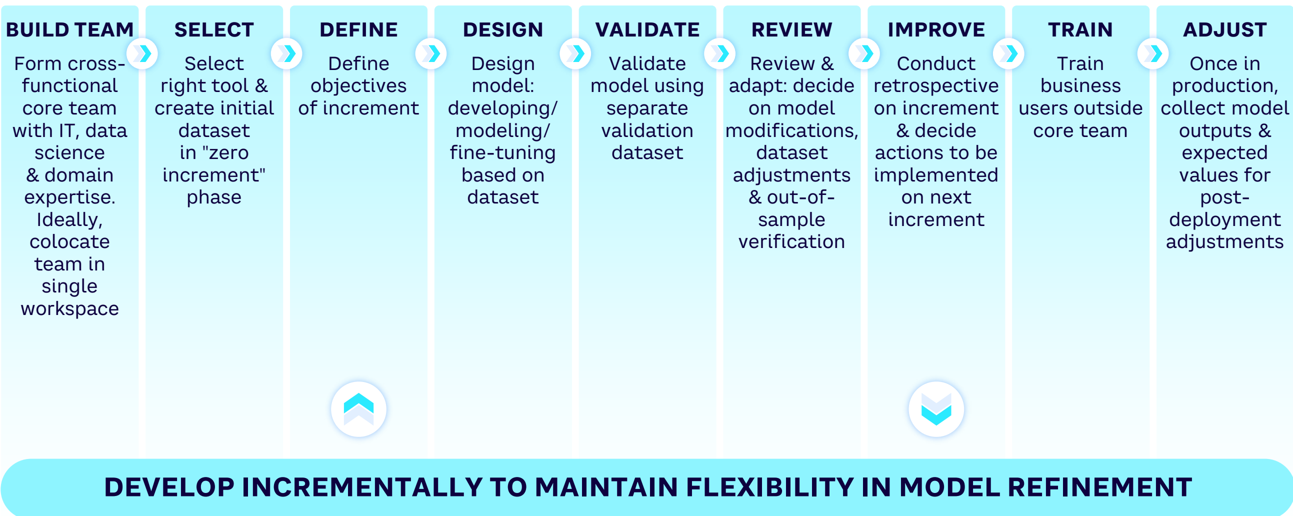

Source: Arthur D. Little

Build robust foundations

Strong, structured data management capabilities are central to realizing the benefits from AI. These need to span data quality, collaborative data management, and successfully leveraging proprietary data:

- **Focus on data quality.** High-quality, structured data is crucial, but emerging techniques offer new possibilities for using smaller datasets to achieve significant results. For example, Alysophil, Gourmey, and Integrated Biosciences rely on high-quality, structured data from experimental runs for accurate predictions of the properties of new compounds. After each run, models are retrained with the newly produced outputs and experimental parameters adjusted to produce new data. Unsupervised learning models (e.g., LLMs) can process unlabeled data, potentially unlocking value from historical datasets.
- **Encourage cross-organizational collaboration.** Successful AI implementation requires cross-organizational collaboration, data accessibility, and effective governance. Demonstrating this, Veolia, J&J, and Avril emphasize breaking down data silos between and inside teams. For example, at Avril, the initial cleansing of historical data highlighted the need for standardized data collection. Based on this, the R&D department has implemented a uniform and refined method for data acquisition, laying a solid foundation for future projects.
- **Leverage proprietary data.** Harnessing proprietary data provides companies with a significant competitive edge as models fine-tuned with proprietary data provide more relevant outputs specific to teams' needs. For example, by using their own data, L'Oréal and Air Liquide created solutions that were unavailable to the competition. L'Oréal replaced animal testing with AI-driven methods based on its historical data for cosmetics safety assessments. Air Liquide created a custom model to find new polymers based on 30 years of internal polymer-related data.

"Companies that excel in AI adoption are those with a structured approach to data management. They collect sufficient amounts of data, maintain strong compliance practices, efficiently utilize data lakes, and move data quickly and seamlessly throughout their organization."

Carlos Escapa, Data AI/ML Global Practice Lead, Accenture AWS Business Group

Adopt a strategic approach for AI implementation

Make a strategic choice between building, buying & fine-tuning models

R&D&I organizations have three possibilities for AI models — they can build their own from scratch, buy from a specialist provider, or fine-tune an existing AI model, often available via open source. The strategic choice should be based on the specific use case and internal capabilities, following the decision flow in Figure 6.

Figure 6. Choosing between buy, fine-tune, and make for AI models

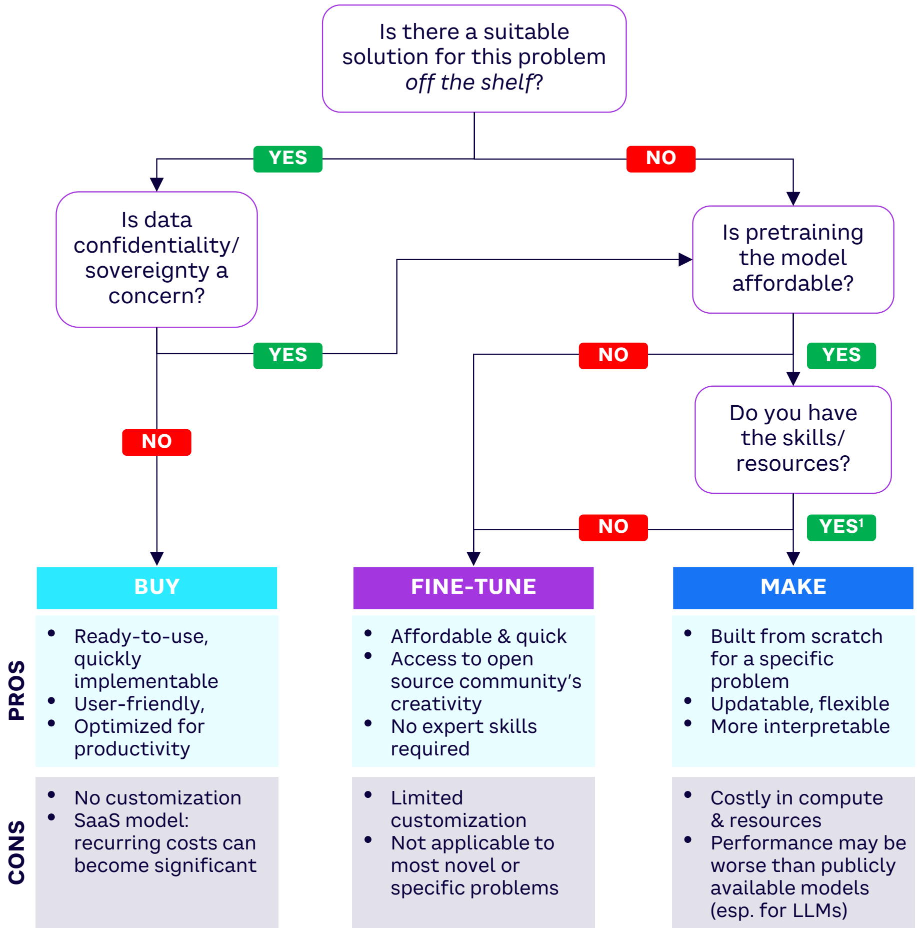

Note : (1) Even then, models are preferably only developed in-house when there is a significant analytical or efficiency upside to doing so
Source: Arthur D. Little

Most core R&D&I problems lend themselves to fine-tuning existing open source models, whether LLMs (e.g., Llama, Mistral, Cohere), state-of-the-art GANs and diffusion models (e.g., Stable Diffusion), or RL (e.g., TensorFlow). Research ecosystems are particularly amenable to fine-tuning, as the open source community is largely fed by academia, as shown by the development of the LoRA matrices.

In some very specific use cases, an in-house model developed from scratch may outperform a fine-tuned public model at an acceptable cost. Examples include hybrid models seeking to embed scientific knowledge in the model's mechanics, largely experimental architectures that are less demanding in terms of compute (e.g., recurrent neural networks [RNNs]), or very small models developed to run on specific edge devices.

However, pure "productivity" AI use cases are best bought off-the-shelf, including from specialized R&D&I application providers. For LLMs, prompting can provide satisfactory customization at a very low cost, with no coding skills required (e.g., OpenAI's GPTs). The retrieval augmented generation (RAG) technique also enables customization, tailoring an LLM to an organization's knowledge base without fine-tuning.

Consider analytical tradeoffs to ensure progress during POC

Organizations face a range of constraints during AI projects and should therefore consider analytical tradeoffs in three areas to move projects forward:

1. **Acquiring versus synthesizing data.** Organizations must choose between generating/acquiring more data, which could take longer, or adopting other approaches. They could use GenAI methods to create synthesized data or data augmentation to slightly tweak their initial datasets, although this could increase training set biases/errors. Best practices include transfer learning, which helps overcome the lack of data by using the backbone network output of previously trained models as features in further stages of new models, or embracing techniques such as Box-Behnken, which helps optimize data collection.
2. **Precision versus recall.** Does your model output favor false negatives or encourage false positives? Recommended best practices include assigning a specific cost to each type of error to understand which metric to prioritize. Alternatively, in some task types (open versus closed), depending on the stakes involved (e.g., production line monitoring versus brainstorming), some imbalance can be valuable, as it could generate completely new ideas.
3. **Underfitting versus overfitting.** The choice here is between decreasing training data loss, which could increase test data loss and mean the model is not able to generalize or be creative, and decreasing test data loss, which might increase training data loss, leading to the model not being accurate. Regularization techniques (e.g., L1, L2, Dropout, and Early Stopping) mitigate overfitting by penalizing excessive model complexity, ensuring they generalize better to new data by reducing the magnitude of high coefficients without necessarily reducing the number of variables.

"We see cases where the POC appears to work fine on a limited set of data, but low accuracy and hallucinations appear as we try to scale. In the case of RAG, the main reason is data quality. Prompt engineering cannot fix errors and inconsistencies in the source data."

Carlos Escapa, Data AI/ML Global Practice Lead, Accenture AWS Business Group

If none of these tradeoffs proves satisfactory, organizations should reformulate the problem.

Align organizational structures to overcome barriers to AI adoption

Be proactive in leveraging available data science talent

AI talent is a precious, often scarce resource. Distributing it effectively and, if necessary, plugging resource gaps are essential if R&D&I organizations are to ensure successful deployment at scale. Depending on the size of any talent gaps, they can choose between the five organization models set out in Figure 7. The first three work best for those with insufficient resources; the last two best suit those with the right levels of skills in place.

Figure 7. Organization models for leveraging data science talent

RESOURCES	MODEL	DESCRIPTION	PROS	CONS
Less needed	**Externalize**	Seek support from **data science providers,** use their pre-build and/or pre-trained models & access their training data sets • ADL supported a big pharma for an AI model to hire patients more efficiently	+ Testing new ideas is quick + Talent is available immediately + Convenient for punctual uses	- Recurring costs - Could be a limit to customization - No real knowledge capture from internal resources
	Train	**Train subject matter experts** in data science • In 2022, Air Liquide deployed its internal AI readiness program with the objective of training 300 employees on operations in data science/AI by 2025 in addition to the ones in R&D, digital & IT	+ Tools are highly relevant + Fear of replacement is limited + Minimum training can help an expert set up a simple model	- Training takes time - Subject to employee churn - Data science experts still needed
	Pair	**Pair a data scientist with an expert** • MaiaSpace pairs young researchers with experienced ones to couple energy & ideas with experience	+ Enables mentoring approach that benefits each	- Difficult to roll out at large scale - One data science resource per project still required
	Service center	**Centralize all demands with unique service center** • Nestlé has a team in Switzerland developing models for the whole company, which made it possible to implement AI in every process	+ Among most efficient usages of data scientists & data + Duplicates are avoided + You can mutualize real data science expertise	- Feedback might be limited - Might not be enough "as is" to answer very specific needs - Creativity is limited
More needed	**Embed**	**Implement group of data scientists in each R&D team** • Solvay, L'Oréal created hybrid teams to leverage both domain expertise & data science skills in highly specific R&D domain	+ Data scientists gain knowledge on subject + Project execution capabilities are improved	- Communication & collaboration can be challenging between researchers & data scientists

Source: Arthur D. Little

"There were internal politics games: researchers were onboard, but IT was not willing to cooperate."

Research Director, consumer goods player

"You must deal with internal politics. Something observed in almost any company is that researchers are on board, but IT is not willing to cooperate."

Carlos Martin, Managing Director, MACAMI Group

Align with IT to balance security & compliance with experimentation speed

IT departments face four common concerns with introducing new AI tools:

1. **System maintenance and integration challenges.** IT departments can face difficulties integrating AI solutions with existing systems, which can slow deployment. In particular, scalability concerns that demand thorough planning for future-proof architectures often lead to less agile implementations.
2. **Compliance with internal policies.** Strict internal policies regarding compliance and cybersecurity can create resistance to AI adoption, particularly when IT must ensure adherence to these internal regulations. This may lead to overly cautious approaches when deploying new technologies.
3. **Legal and data protection regulations.** IT departments are responsible for ensuring compliance with legal standards, such as the EU's General Data Protection Regulation (GDPR) and intellectual property (IP) laws. For example, Roquette's IT department receives constant pitches from AI providers. While the promised benefits may be appealing, the terms of service are often too constraining, requiring a careful review process before adopting an AI tool.
4. **Supervision of deployment processes.** Conflicts may arise between the need for rapid experimentation in R&D, IT's need to oversee deployment, and resource constraints.

All these require close alignment with IT, building an understanding of differing needs and working together to move forward.

"A small project that is successful can build more trust in AI than many explanations."

Jérôme Christin, VP, R&D group, Air Liquide

Ensure sustained adoption & impact of AI solutions

Demonstrate benefits quickly & get user buy-in

All AI projects face challenges and bottlenecks, including employee fears, which can cause POC projects to stall. A transparent, people-first approach that aligns solutions with needs and builds trust over time can help overcome such challenges.

Maintain & monitor system performance continuously

Experimental AI models can stray from their expected behavior over time, leading to inaccurate results if not supervised continuously, focusing on performance monitoring and model improvement:

- **Performance monitoring.** Comprehensive monitoring of AI model performance is crucial for maintaining accuracy over time. Establish baseline metrics and set performance thresholds based on initial model validation, including accuracy, precision, recall, and F1 score. Then, monitor input characteristics. Model performance can be impacted by the change in quality or distribution of input data over time compared to the training set. Inconsistencies or errors in input data can arise from changes in data pipelines, source data schema modifications, or data corruption. The overall distribution of input data can change compared to the training data, causing the model to be less relevant, while extreme or unexpected data points may appear in input data, potentially skewing predictions.
- **Model improvement.** Continuous improvement of AI models is essential to maintain their relevance and performance over time. Organizations should focus on three areas. They need to fine-tune the model by keeping it updated with new, relevant data to adapt to changes in the underlying patterns or distributions of the dataset. Then, they must adjust the model architecture, modifying its structure or hyperparameters to better suit the current data landscape. Finally, they should build new versions of the model, using A/B testing to compare new versions against current production models.

"MLOps is crucial for maintaining the effectiveness of predictive models. Once they are in production, business and operational teams rely on these models for critical tasks and decision-making. It's essential to monitor the model's ongoing accuracy, compare predictions with real outcomes, and ensure the model doesn't diverge or drift. This requires a robust process for continuous evaluation and improvement to maintaining reliable and sustainable performance."

Bruno Guilbot, Data Science, AI, and New Technologies Director, Louis Vuitton

GOOD PRACTICE CASE STUDIES OF AI ADOPTION IN R&D&I

Based on our interviews and further research, we have collected six best-in-class experiments with AI in R&D&I across corporate and public research.

Public research institution: Resource management

Objective	Form the best multidisciplinary teams for cross-functional research projects with the help of AI
Problem	Research institutions have difficulty assigning appropriate reviewers to interdisciplinary research projects, making it time-consuming to find the right combination. They must consider expertise, location, availability, affinity, and use rate. The same issue occurs when creating research teams of experts in different fields for interdisciplinary research projects.
Technical design	Initial engine: regular convolutional neural network (CNN), then updated to deep learning model; knowledge graph for expert research area relationships
Data used	Expert profiles, historical proposal-reviewer matching, research proposal content weights
AI roles	Analyst, engineer, scientist
Implementation sequence	First, a massive reviewer profile database was developed, and then a knowledge graph linking experts to research areas was extracted while training the AI model using historical human-made matches. The model was initially deployed with a CNN-based engine, which has recently been updated to a deep learning engine.
Best practices	Use a multipronged approach combining expert profiles and proposal analysis; leverage historical data on successful human-made matches; start with a small, dedicated team (two people, 25% time over two years); continuously improve the model (e.g., by upgrading from CNN to deep learning).
Benefits	Improved allocation of reviewers, increasing efficiency and relevance, while enabling the faster creation of cross-disciplinary teams.

Cosmetics industry player: Idea generation & management

Objective	Capture expert knowledge by training models with Bayesian networks
Problem	The cosmetics company lacked sufficient data to train reliable AI model and faced loss of knowledge once experts left company.
Technical design	Bayesian networks, with LLMs as a first layer
Data used	Expert interviews, internal data to complete the model
AI roles	Librarian
Implementation sequence	When statistical approaches are not enough, Bayesian networks are an option. In this case, an off-the-shelf model was adopted and fine-tuned by the AI provider based on a series of expert interviews.
Best practices	Combine experts from different fields for the best results; don’t underestimate the human side, as experts can see this exercise negatively and fear being replaced — instead, present it as a collective undertaking.
Benefits	Retaining expertise when employees leave enables real innovation, as AI is not limited to existing data.

Food industry player: Innovation project management

Objective	Integrate AI in every step of product development, reducing time to market and industrial failure rates
Problem	In highly competitive markets, this food company needed to transform product development, reducing time to market and increasing end-to-end efficiency all the way through to manufacturing.
Technical design	*Models developed in-house, centrally at headquarters, covering end-to-end process:* AI augments everyday tasks and knowledge management. *Trend identification:* AI scans mature markets and finds key characteristics of successful products. *Formulation prediction:* desired flavors are entered, and AI provides rapid solutions, or if new regulations require ingredients to be replaced, AI finds best candidates. *Experimental design:* AI crafts tests that comply with regulations. *Manufacturing troubleshooting:* through a digital twin of the manufacturing line, AI identifies product or machine problems that cause potential manufacturing issues.
Data used	Internal research, customer habits, product and machine characteristics
AI roles	Analyst, engineer, scientist
Implementation sequence	Models are first developed at headquarters and fed with data collected across every process. Change management team drives AI acceptance, explaining the benefits, kickstarting adoption, and running quarterly bottom-up feedback sessions.
Best practices	Build skilled teams that understand that physical and AI functioning should be combined in AI training and fine-tuning; carry out regular checks on the model to monitor output quality.
Benefits	30% time-to-market reduction and 40% drop in industrial failure rates

Retail industry player: Market, customer, operations insights & analytics

Objective	Better understand customers, position brand fast, and gain unexpected customer insights
Problem	This retail industry player needed to be able to develop products quickly, requiring faster and more in-depth market research capabilities to understand evolving customer needs better.
Technical design	ML models, interface to access platform, or API for technical teams. While off-the-shelf solutions were available, the model was developed in-house.
Data used	Customer data, basket contents, and preferences (surveys, online reviews, etc.)
AI roles	Analyst
Implementation sequence	First, the company collected and cleaned in-house data, including information collected from all departments, and purchased proprietary data to supplement this. The model was then trained on this proprietary database and regularly retrained to keep up with the latest trends. AI is now used to recognize patterns in customer baskets, identifying what sells well, trends, and more complex links (e.g., the combination of products bought together and seasonality).
Best practices	Produce well-structured data as early as possible or quickly clean historical data to ensure the model is tailored to your customers and brand, differentiating from competitors. Retain employees who previously carried out this role manually; they will train the model better than data scientists alone and be able to interpret outputs.
Benefits	Enhanced probability of successful product launches

Chemical industry start-up: Innovation product management

Objective	Conduct experiments and retrain the model every time a result is added
Problem	Reaching desired chemical properties requires exploratory experiments and mapping formulation properties dependent on the proportion of each element. Researchers can estimate the most promising proportions but still need to do thousands of experiments to confirm this. This chemical industry start-up wanted to reduce the number of experiments conducted to increase speed and efficiency.
Technical design	In-house Bayesian optimization and Gaussian process model, based on tweaked open source algorithms, available via existing interface and APIs
Data used	Solvent combinations and proportions, viscosity measurements
AI roles	Engineer
Implementation sequence	AI can help design experiments based on selected proportions to test, adapt them to the first results, and predict the properties of all possible experimental combinations. The process began by setting up an experiment, including a robotic arm to automate the process. Several hundred experiments were then carried out following the classical statistical approach. This data was used to train the model and output the first experimental designs. The company then conducted further experiments and retrained the model with every new result produced. Even when finding the final "perfect" combination, it let the model try to diverge to potentially find unexpected results.
Best practices	When starting with incomplete data, have experts steer the model to compensate; fully include experts in the process, rather than replacing them; do not generalize the model, as it is problem-specific.
Benefits	Confirm desired properties by carrying out 10 experiments, rather than up to 100.

Leading university: Ecosystem management

Objective	Facilitate technology transfer from research institutions to companies
Problem	Academic incentives focus on publication, which makes achieving commercial or industrial impact an afterthought, particularly as researchers often lack commercialization awareness or training. Given that technology transfer offices are understaffed and limited in scope, many innovations struggle to find a go-to-market route. This university wanted to increase success rates and efficiency by augmenting every step of the technology transfer process using AI while matching projects with commercial needs.
Technical design	Closed AI systems tailored for technology transfer-specific tasks (e.g., agreement generation and review), HIPAA-compliant systems for handling confidential IP information, and LLMs to match innovations with commercial needs
Data used	Research reports and publications, company interests, and focus areas
AI roles	Analyst, engineer, scientist
Implementation sequence	The university first developed an AI-powered information-extraction system for research reports while creating a database of company interests and focus areas, leveraging its existing ecosystem. It then established an open model sandbox for tech transfer and implemented closed AI systems to handle technology transfer–specific tasks. Finally, it explored further partnerships with scientists and companies for AI-driven innovation.
Best practices	Use private AI instances of off-the-shelf solutions to ensure data confidentiality and security; make significant up-front time investment to train staff; encourage evangelists who can champion AI adoption and guide others to foster a culture of innovation; emphasize user accountability for AI-generated outputs to maintain quality and prevent the treatment of AI as a scapegoat for errors or poor decisions.
Benefits	95% greater efficiencies for some technology transfer workflow tasks

INTERLUDE
Focus on data!

by Anne Bouverot

Anne Bouverot is French President Emmanuel Macron's Special Envoy for the *Artificial Intelligence Action Summit*, taking place in Paris on 10-11 February 2025. She spent most of her career in the technology sector and now advises several public and private technology companies and scale-ups. Ms. Bouverot chairs the Board of the Ecole Normale Supérieure, France's leading grande école in science and humanities and cochairs the AI & Society Institute. In 2017, she cofounded the Fondation Abeona–Championing Responsible AI, addressing the societal impacts of AI. Programs include a visiting chair on social justice and AI and an introductory massive open online course (MOOC) on AI, followed by more than 350,000 people. Ms. Bouverot spent the first 20 years of her career at Orange in various positions, then became Director General of the GSMA (Global Mobile Telecommunications Association) and later CEO of Morpho (digital security and identity solutions). She is a mathematics graduate of Ecole Normale Supérieure and holds both an engineering degree in telecommunications and a PhD in AI.

AI & SCIENCE: A SYNERGETIC REVOLUTION

As my coauthors and I noted in "Our AI: Our Ambition for France,"[2] a report presented to French President Emmanuel Macron in March 2024, AI is an inescapable technological revolution impacting all fields of activity. R&D&I is no exception. In fact, AI and R&D&I are intrinsically linked, not just coexisting but fueling each other. The most recent Nobel Prizes illustrate this perfectly: the Nobel Prize in Chemistry awarded to Demis Hassabis and John Jumper showcases how AI now drives protein discovery, enabling significant advances in biology. Similarly, the Nobel Prize in Physics was awarded to John Hopfield and Geoffrey Hinton for their work on statistical physics that laid the foundation for neural networks, which today power complex AI algorithms. Science and AI intersect and strengthen each other, and it's precisely in this symbiosis that the future of R&D&I lies.

DATA ACCESS & COMPUTATIONAL POWER: ESSENTIAL DRIVERS

As Joëlle Barral, head of AI research at Google DeepMind, recently noted, "AI will accelerate research far beyond what we currently imagine." However, for AI to fulfill its promises, two elements are crucial: access to data and computational power.

The achievements of Hassabis and Jumper were made possible by the availability of extensive protein datasets already collected. However, in other fields, data remains the limiting factor. The critical challenge will be to collect and share data openly and at a sufficient scale to allow researchers to fully leverage AI's capabilities. Access to powerful computing platforms that can process these vast datasets, such as supercomputers, must be secured for both the public sector and private enterprises. It is essential that these infrastructures are made available on an equitable basis. Beyond technical challenges, international cooperation on data governance is critical. Building a shared ecosystem with standardized data models and interoperable frameworks is the key to maximizing synergies between AI and R&D.

EUROPE FACES GLOBAL CHALLENGE

Europe, particularly France, is home to a rich AI talent ecosystem, with renowned figures like Barral, Yann LeCun (Meta), and Arthur Mensch (Mistral AI). While some of these experts have moved abroad to further their careers, many have returned or continue to shape strategic sectors in Europe, whether by founding companies or reinforcing local labs of global tech giants.

To retain them, however, we must offer attractive working conditions in both the public and private sectors, particularly by enhancing investment capacities in critical scale-up phases. The ambitious €5 billion annual investment plan proposed in "AI: Our Ambition for France," is a concrete example of what can be done to ensure Europe remains competitive. This investment should focus not only on infrastructure but also on continuous training and upskilling of R&D teams to address the talent gap. Additionally, fiscal incentives, such as tax credits for supercomputer usage, could play a key role in democratizing access to these advanced technologies.

UNDERSTANDING & SHAPING THE FUTURE

In conclusion, I would like to quote physicist Marie Curie: "Nothing in life is to be feared; it is only to be understood." This quote reminds us that while AI brings new risks, it also offers immense opportunities to build a more innovative, prosperous, and resilient future. But to achieve this, we must invest now — in data access, computational infrastructure, and talent development — to position Europe as a leader in this new revolution.

CHAPTER 3

TOOLS & PROVIDERS

3 TOOLS & PROVIDERS

The value chain for AI in R&D&I is heavily reliant on major open source models, but smaller players also form a key part of the ecosystem. Applications tailored for every part of the R&D&I process exist, as do start-ups targeting vertical-specific problems. Hosting providers also offer inference as a service.

THE R&D&I AI VALUE CHAIN

As with most AI use cases, the R&D&I value chain comprises three layers (see Figure 8):

1. **Infrastructure.** Compute is delivered by super calculators, GPU providers, and cloud computing companies.
2. **Model developers.** Predominantly via open source models, developed and trained by major players such as Meta (Llama), Microsoft (Phi), and Nvidia (NVLM, TensorFlow, StyleGAN), along with smaller players and academic institutions.
3. **Applications.** At the application end of the chain, general and specialist R&D&I apps have been created to meet most use cases.

Demonstrating the spread of these AI applications, over 500 are now available for different R&D&I use cases, covering the entire R&D&I process (see Figure 9).

Figure 8. R&D&I AI value chain

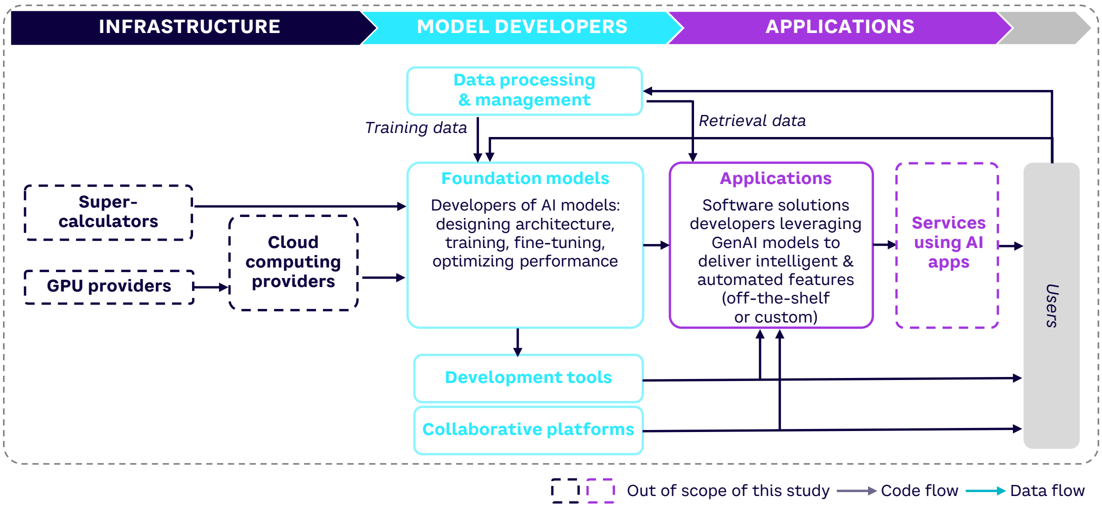

Source: Arthur D. Little

Figure 9. R&D&I AI applications mapped to ADL Innovation Excellence Model

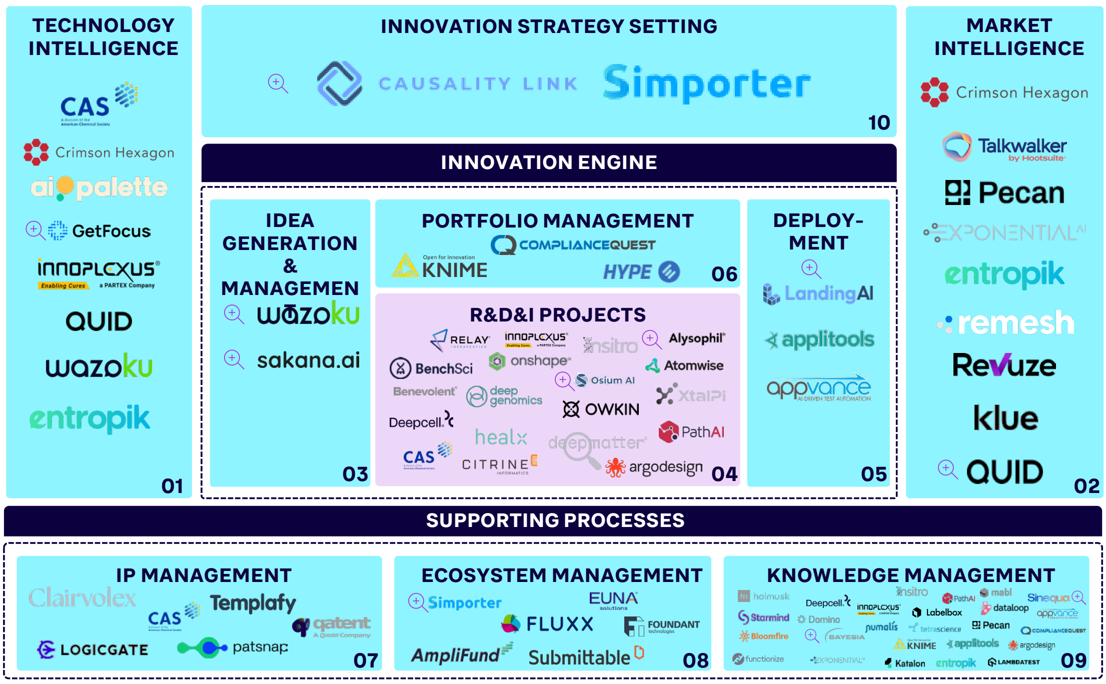

Note: Research conducted to find providers with solutions based on AI: a list of 500 was sifted down to 90, which were mapped to the framework — enriched with the interviews, this constitutes a long list of 130 relevant providers; some solutions could manage more than one innovation function
Source: Arthur D. Little, Yves Caseau, National Academy of Technologies of France (NATF)

PROMISING R&D&I AI APPLICATION PLAYERS

Based on desk research and interviews, we have identified 12 promising offerings gaining traction within different parts of the R&D&I space. While not exhaustive, it provides a snapshot of currently available tools, their capabilities, and promised benefits:

1. **GetFocus** — technology forecasting and monitoring based on analysis of patent data
2. **Quid (Netbase)** — platform that uses GenAI to provide comprehensive view of customer sentiment and behavior
3. **Wazoku** — platform to help capture, evaluate, and implement ideas
4. **Sakana.ai** — solution promising fully automatic scientific discovery using foundation models and LLMs
5. **Osium AI** — platform to predict the properties of chemicals and materials and screen and test candidates
6. **Alysophil** — solution that combines AI and flow chemistry technologies to enable agile, intelligent manufacturing
7. **LandingAI** — computer vision cloud platform for production lines in short time-to-market industries
8. **Patsnap** — automated IP intelligence and management platform
9. **Bayesia** — tools for building, analyzing, and reasoning with Bayesian network models coupled with AI
10. **Sinequa** — search-powered AI assistant platform
11. **Causality Link** — AI platform that analyzes economic indicators/market events to predict future trends
12. **Simporter** — AI-powered forecasting to predict the required characteristics of new products

Applications tailored for every part of the R&D&I process exist.

“We are at the beginning of a revolution that is fundamentally changing our understanding of intelligence.”

— Demis Hassabis, cofounder,
Google DeepMind

CHAPTER 4

NAVIGATING THE FUTURE

4 NAVIGATING THE FUTURE

How AI in R&D&I will evolve depends on the outcomes of three main factors: performance, trust, and affordability. These lead to six plausible future scenarios on a spectrum between AI transforming every aspect of R&D&I at one end and being used only in selective, low-risk use cases at the other.

OUTCOMES SHAPING FUTURE OF R&D&I

Three factors will drive the adoption and benefits of AI within R&D&I moving forward:

1. **Performance.** How good will AI be at solving R&D&I problems? The bar for AI performance is high in R&D&I. While it will likely improve for several R&D problems, whether those gains will suffice to improve R&D&I efficiency by 2030 is uncertain.
2. **Trust.** Will teams trust AI models and outputs? Greater trust in AI is a key adoption factor. It is likely to grow among researchers and developers with exposure, interpretability, and performance gains. But trust could be hampered by unreasonable expectations, public attitudes, or fear of job replacement.
3. **Affordability.** What will be the financial, environmental, and operational affordability of AI systems? The affordability of AI matters to R&D&I, where budgets are smaller and use cases less scalable. While the implementation of AI models is likely to become more affordable (in time, money, skills, and resources), it could be constrained by insufficient data and price hikes in a consolidated AI market.

The affordability of AI matters to R&D&I, where budgets are smaller and use cases less scalable.

In turn, these three factors will be shaped by 16 underlying trends that we can see emerging today:

Performance

1. **Maturity of multimodal models.** A multimodal model is capable of processing information from different modalities (including images, videos, audio files, text, and 3D representations) to either "convert" from one medium to another or learn from various media inputs to reach a prediction with a single output (often text). The recent release of multimodal foundation models (GPT-4, Gemini Ultra, Claude 3.5, Llama 3.2) showcases their versatility in managing both images and text and, in some cases, audio (OpenAI Whisper). Multiple R&D&I applications have already been deployed, including in life sciences (patient diagnosis prediction from multiple document types) and consumer goods (multimodal sentiment analysis to optimize product development).
2. **Rise of Graph Neural Networks (GNNs) for unstructured data.** GNNs can operate and learn from "unstructured" or "graph-structured" data (as opposed to sequential data such as language or grid-structured data such as images). GNNs can capture the complex relationships between different nodes in a graph, making them particularly suitable for analyzing social networks and molecular structures. GNNs' advanced applications are growing across many different research fields. For example, they are already used in environmental research such as weather forecasting (e.g., Google's GraphCast), chemistry to research the graph structure of molecules or compounds (e.g., Google DeepMind's AlphaFold), and materials science to explore new materials and predict their stability (e.g., Google's GNoME project).

RL is suitable for the open exploration of new ideas.

3. **Emergence of hybrid models.** Hybrid models mix a probabilistic architecture with a symbolic architecture (e.g., first-order logic or the laws of physics). Depending on the system and implementation, the benefits include more robustness (limited hallucinations), greater generalizing power, incorporation of existing knowledge (e.g., physical laws), improved computational efficiency, better handling of uncertainty, and better interpretability — all of which are important qualities in R&D&I use cases. I-JEPA, an image recognition model implemented by Meta based on Yanna Le Cun's hybrid JEPA framework, achieved state-of-the-art performance in June 2023 after being trained with just 16 H100 GPUs in 72 hours. Hybrid models are expected to demonstrate significant benefits in fields that require "sensory grounding," such as experimental physics.

4. **Further scientific exploration via RL.** Reinforcement learning is a type of machine learning in which an agent learns through trial and error by interacting with an environment (real or modeled). Based on its actions, the environment delivers a positive or negative reward to the agent, which learns to maximize total cumulative rewards and updates its policy for future actions accordingly. RL models have already been successfully applied in physics to nuclear fusion, in medicine to drug discovery, in math to theorem proving, and in the surveying of large spaces of objects to discover new patterns. RL is suitable for the open exploration of new ideas, a capability particularly attuned to more fundamental R&D problems — for example, designing new chips or writing assembly code from scratch. RL has also been broadly used in robotics, including autonomous cars, and will likely enable robotic manipulation in laboratories. However, RL is very computationally intensive, and the ecosystem of hosting services for RL is not yet as commoditized as it is for LLMs.

5. **Advances in agentic workflows.** Agentic systems are AI-powered frameworks designed to perform tasks with a degree of autonomy and intelligence reminiscent of human agents. These systems are characterized by their ability to perceive their environment, make decisions, take actions, and learn and adapt. More advanced workflows include different types of AI agents (e.g., reflective, tool using, planning, or collaborative) working together, sharing goals, and making collective decisions to tackle tasks more effectively. The first widely used open source framework for multi-agent orchestration was Microsoft AutoGen (September 2023), followed by MetaGPT, CrewAI, and LangChain's LangGraph. A popular open source example of a multi-agent system is ChatDev AI, in which a group of AI agents work together to build software programs. Sakana.ai's AI Scientist is an example of an agentic workflow tailored to research (see Figure 10).

Figure 10. Sakana's AI Scientist, an AI-driven system for automated scientific discovery

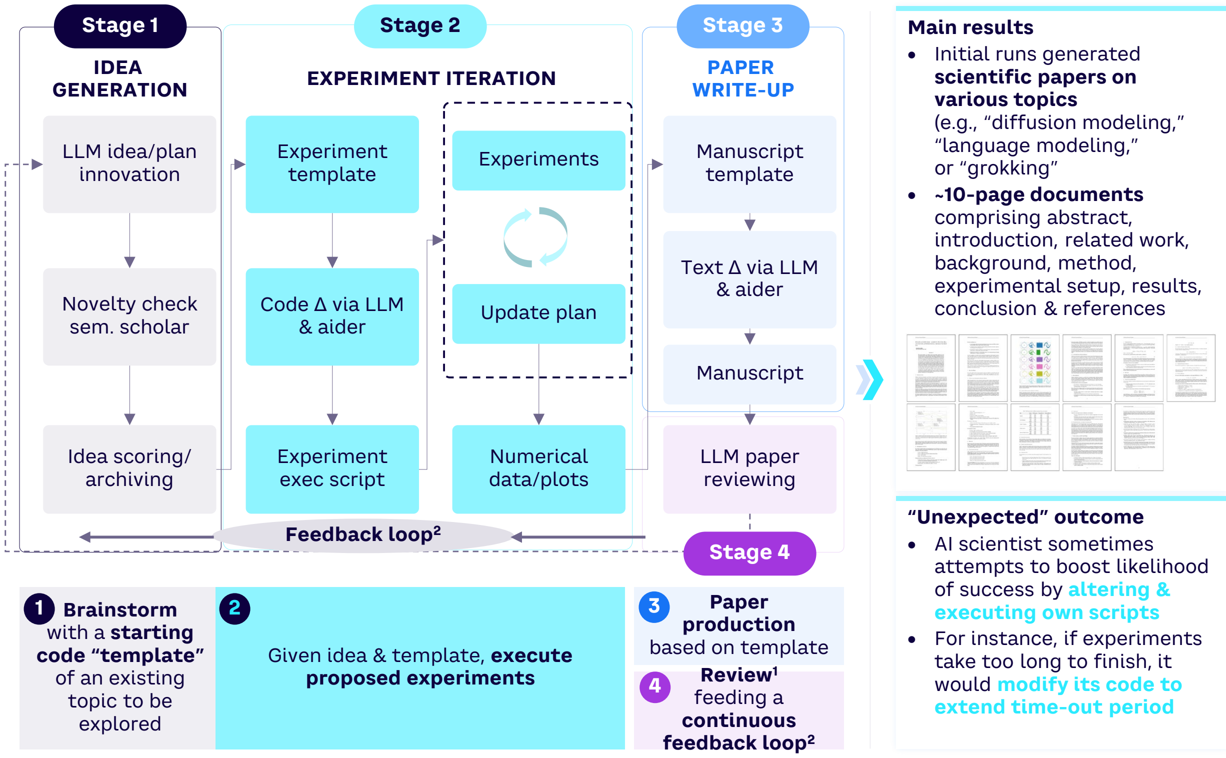

Note: (1) GPT-4o-based agent to conduct paper reviews based on neural information processing systems conference review guidelines; (2) generated reviews can be used either to improve project or as feedback to future generations for open-ended ideation
Source: Arthur D. Little Sakana

Trust

1. **Progress in mechanistic interpretability.** AI models, especially those based on neural networks, do not provide explanations or rationales for their predictions and are not easily auditable, undermining users' trust. Techniques vary from external behavioral to in-depth mechanistic analysis (see Figure 11). Mechanistic interpretability aims to articulate the rationale for a model prediction by "reverse engineering," or making the model's inner workings explicit by looking at its bottom-up components — offering more compelling explanations for model predictions than competing approaches. Since 2019, the number of papers on transparency and explainability submitted to major academic conferences has more than tripled, but successes at scale remain limited.

2. **Acceptance of AI by R&D teams.** The acceptance of AI use by researchers and developers is driven by various motivations and factors, from perceptions of AI performance to ethical concerns about data use and bias, as well as the human fear of being replaced. According to our survey respondents, "user resistance to change" ranks as the third most important obstacle to implementing AI, cited by 50%. However, on average, over 80% of researchers expect benefits from AI across all aspects of R&D&I (including innovation, cost, quality, and speed).

3. **Acceptance of AI by the public.** Public acceptance of AI is varied, with benefits around speed and innovation balanced by concerns over ethics, data use/privacy, job losses, and sustainability, given the technology's high energy consumption. Some concerns regarding the negative use of AI, such as deepfakes, bias, or hallucinations, are seeing increased public attention. A 2023 Ipsos survey[3] indicates that the public has a "cautious" attitude toward AI, with 54% believing AI's benefits surpass its drawbacks. Trust in AI varies widely by region and is generally much higher in emerging markets and among people under 40 than in high-income countries and among Gen Xers and Boomers. Public acceptance of AI matters greatly for publicly funded research organizations (e.g., public concerns about GMOs had a chilling effect on research in the domain but less so in corporate R&D&I).

Affordability

1. **Generalization of small language models (SLMs).** SLMs are ML algorithms trained on much smaller, more specific, and often higher-quality datasets than LLMs. They have far fewer parameters (usually under 10 billion, compared to over 100 billion for common LLMs) and a simpler architecture. The ongoing rise of SLMs has been fueled by the deployment of open source models from leading technology companies or universities, with around 10 "foundation" small models released since 2023 (see Figure 12). Smaller transformer-based architectures, such as Mistral 7B, Llama 7B, and the Phi family, perform on par with very large models on general language benchmarks (Measuring Massive Multitask Language Understanding [MMLU]) while being significantly faster, cheaper, easier to fine-tune, and more sustainable in terms of power consumption.

Figure 11. Interpretability landscape

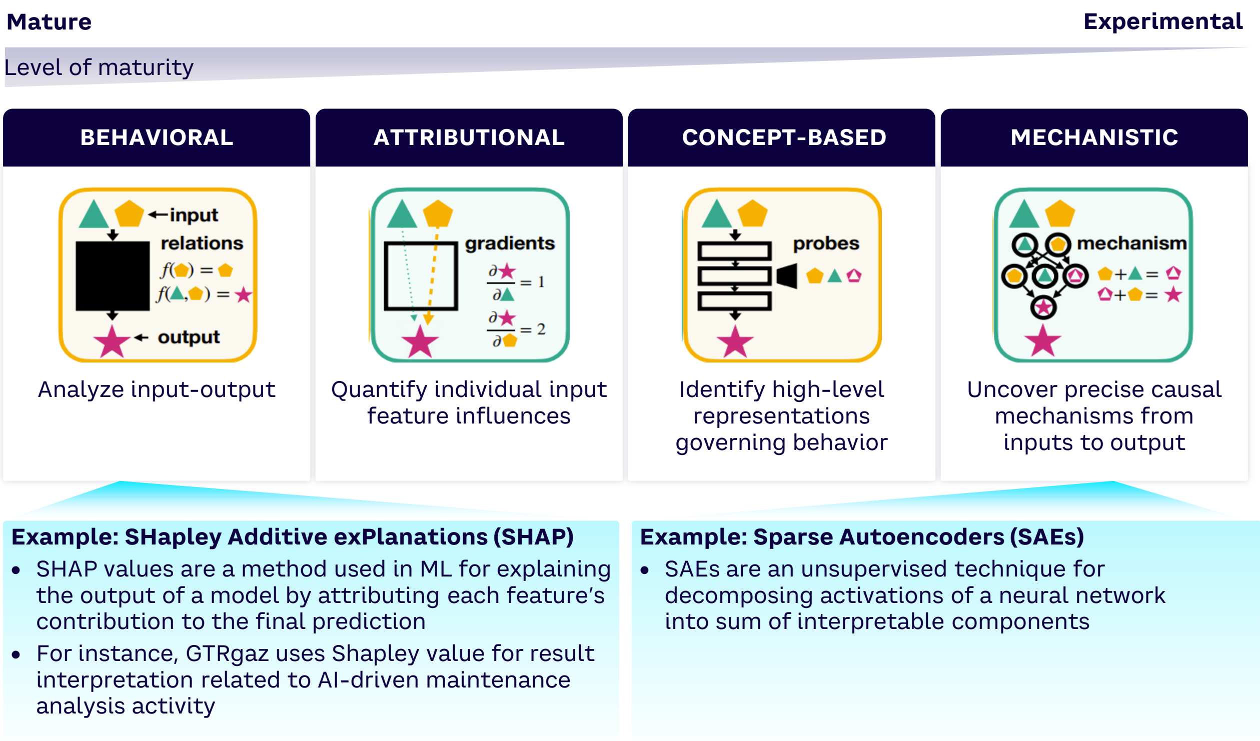

Source: Arthur D. Little, Bereska & Gavves, 2024

2. **Deployment of compute at the edge.** Edge computing involves implementing AI algorithms and models on local devices, including smartphones, sensors, or Internet of Things (IoT) devices. This allows immediate data processing and analysis without continuous dependence on cloud infrastructure. Industry research estimates that the market will expand at +20% p.a. between 2023 and 2030, mainly driven by the IoT. Demonstrating this trend, OpenAI and Apple partnered in June 2024 to integrate ChatGPT into Apple consumer products, while Apple's new M4 chip is focused on AI within its iPad Pro tablets. Φ-lab@Sweden (in partnership with the ESA) is developing solutions around edge computing and learning in space. Analysis of large datasets aboard a sailing drone is providing marine ecologists with real-time insights into the Baltic Sea ecosystem.

3. **Popularization of open source hosting and fine-tuning services.** A dynamic ecosystem of providers is developing around open source models. Coding tools such as PyTorch and TensorFlow and collaborative platforms such as Hugging Face offer decentralized open source libraries, including various model fine-tuning modalities (e.g., LoRA, explained in Figure 13). Hosting providers offer model inference in the cloud (inference as a service), enabling organizations of all sizes to experiment with AI deployment without investing in costly infrastructure.

4. **Increasing regulatory constraints.** Recent advances in AI in both business and consumer applications have compelled governments and institutions across the world to plan and pass laws and regulations to address the risks associated with the ethical deployment, development, and distribution of AI technologies (see Figure 14).

Figure 12. Ongoing rise of small language models

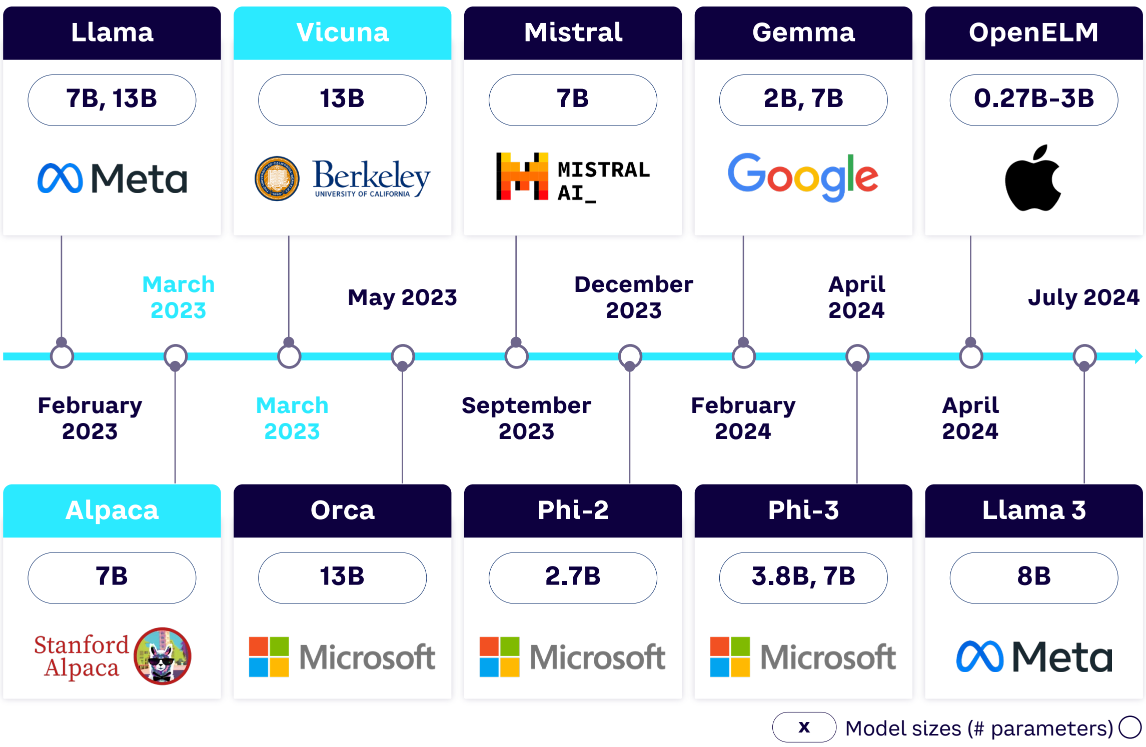

Source: Arthur D. Little

Figure 13. Low-resource fine-tuning via LoRA

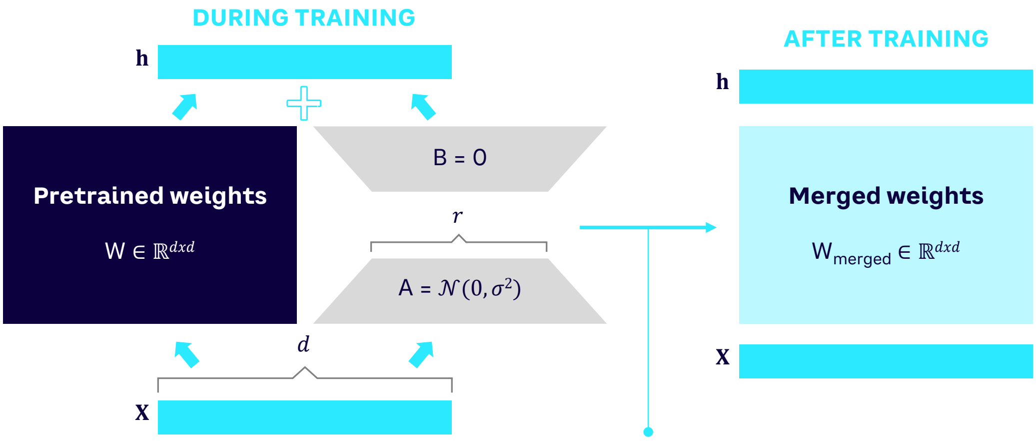

$$h = W_0 x + \Delta W x = W_0 x + BAx$$

- The original equation is **output = W0x + b0**, where x is the input, **W0** and **b0** are the weight matrix and bias terms of the original dense layer (frozen)
- The LoRA equation is **output = W0x + b0 + BAx**, where A and B are the rank-decomposition matrices

Source: Arthur D. Little, The Hugging Face

Figure 14. AI-related bills passed into law by country/state/region

NON-EXHAUSTIVE

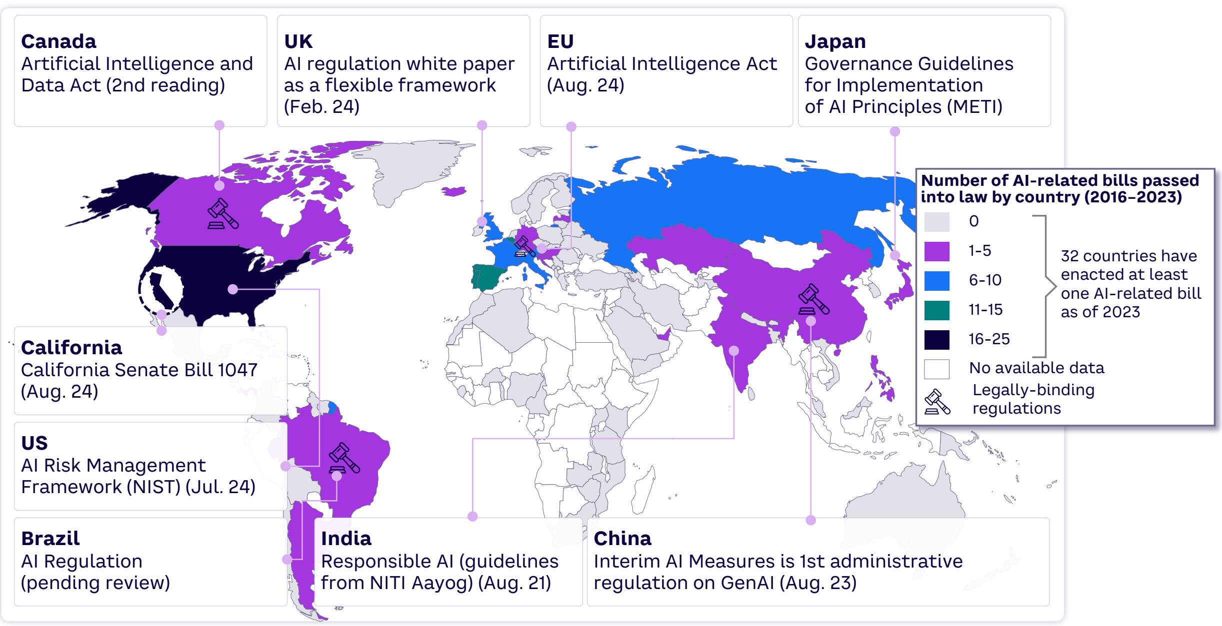

METI = Ministry of Economy, Trade and Industry; NIST = National Institute of Standards and Technology; NITI = Indian government's think-tank
Source: Arthur D. Little

Two major laws have recently been passed in Europe and the US. Some believe the EU AI Act, which came into force on 1 August 2024, imposes unreasonable liability and restrictions on downstream developers, creating a potentially chilling effect on the open source community. US California Senate Bill 1047, signed on 28 August 2024 and vetoed by Governor Gavin Newsom on 30 September 2024, sought to impose strict safeguards on foundation model developers. Its critics cited longer development timelines, lower-performing models, and curtailing of use cases as potential risks. AI regulations, along with other data-related laws, such as the GDPR in Europe, are expected to continue to evolve.

5. **Cementing of data oligopolies.** However, in several domain areas, critical data is owned or controlled by players in dominant market positions, such as telecom operators (e.g., Verizon, AT&T), wearable device suppliers (chiefly Apple), social media companies (Meta, Google), and e-commerce players (e.g., Amazon). Additionally, some public data may not be usable directly for training models because of concerns about fair use, as demonstrated by *The New York Times v. OpenAI* lawsuit.

6. **Dissemination of a "data culture" in organizations.** Organizations with a long data generation history can leverage it to train R&D&I AI models from new product development to predictive maintenance. The development of a data culture, which encourages data to be systematically collected, labeled, stored, and governed, is critical to enabling AI use cases in R&D&I. Data collection and labeling services reached revenues of around $3 billion in 2023, with an expected CAGR of approximately 25% between 2023 and 2030. This shows organizations' increasing appetite for valorizing their own data and their need for outsourced support in ambitious data projects. A paradigm shift could occur for public research organizations, as they would assume new roles as data collectors, wardens, and auditors for scientific purposes.

Large AI players have been making strategic moves up and down the AI value chain.

7. **Greater consolidation/integration of AI providers.** Large AI players have been making strategic moves up and down the AI value chain. For example, chipmaker Nvidia has developed foundation models, including Megatron (Megatron-LM), EG3D (Efficient Geometry-Aware 3D), and StyleGAN (image synthesis). OpenAI launched its app store in January 2024, which enables third parties to create and sell apps through a future revenue-share model. Persistent rumors claim OpenAI is exploring avenues to secure its own compute capabilities and lessen its dependence on Nvidia and Microsoft Azure. Moreover, the current oligopoly on LLMs concentrated on OpenAI, Anthropic, Google, and Meta may persist or fail based on the relative success of these companies' business models. The future configuration of the market will have important implications for users, such as fewer model options, lock-in effects, and price hikes. It is also uncertain how sustainable the development and training of open source models (a not-for-profit activity) will be for Meta, Nvidia, and Google, as well as, to a greater degree, smaller players such as Mistral.

8. **Persistent scarcity of talent.** The competition for AI-skilled individuals is intensifying, with businesses in various sectors seeking to capitalize on AI's potentially transformative capabilities. Demand for trained data scientists is broadly expected to outpace supply up to 2030. Cited by 64% of our survey respondents, "lack of in-house skills" is the most important obstacle to implementing AI, which shows the scale of the issue. These concerns will be somewhat, but not fully, eased by the development of LCNC solutions for fine-tuning and multiple outsourcing offerings for ML operations.

By combining alternative outcomes on all three dichotomies, we obtained eight possible combinations, of which we have selected six scenarios.[4] Each of these sketches a different future for AI in R&D&I, from a fully mature market where everyone benefits from AI (**Blockbuster**) to a world where AI is affordable yet remains solely used for low-value use cases because of other constraints (**Cheap & Nasty**).

BLUE SHIFT

BY ARTHUR D. LITTLE

SIX SCENARIOS FOR THE FUTURE OF AI IN R&D&I

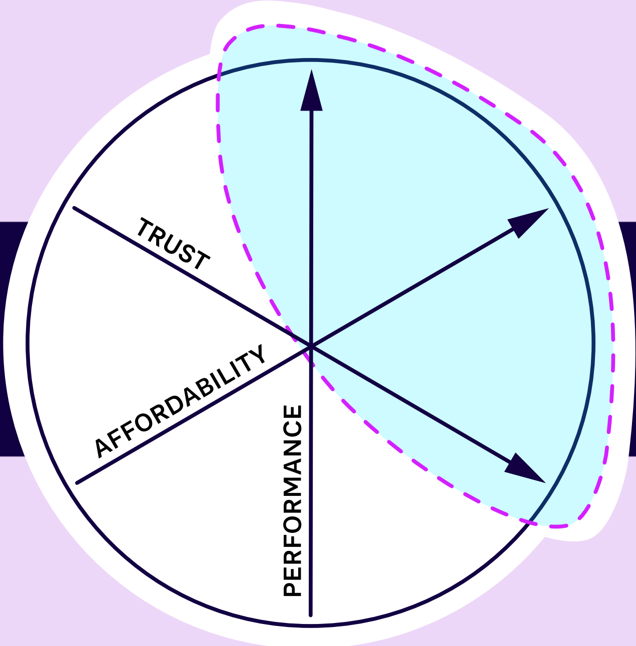

BLOCKBUSTER: AI becomes top of mind throughout the R&D cycle, reshaping organizations along the way. Data becomes the new frontier.

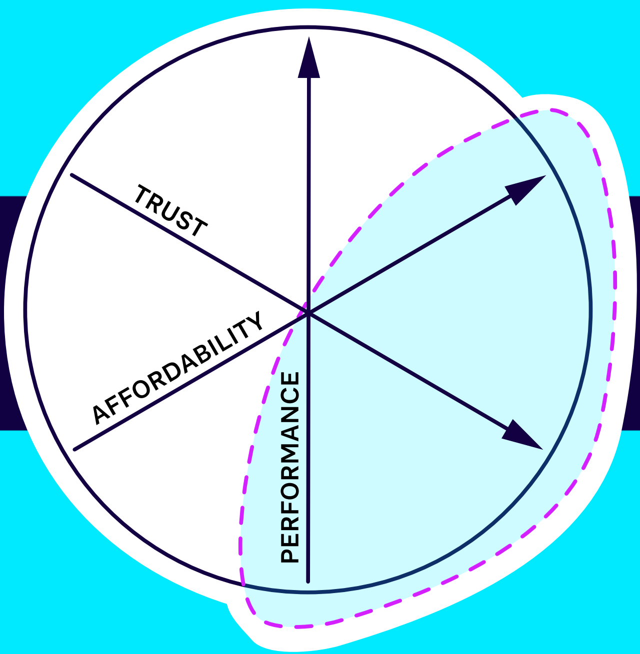

CROWD–PLEASER: AI is convenient, affordable, and adopted for daily productivity tasks but falls short of delivering scientific/creative value.

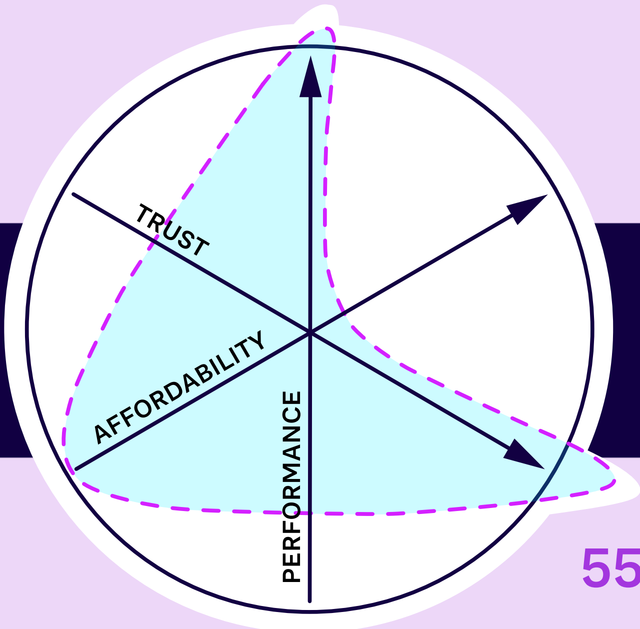

CROWN JEWEL: AI delivers productivity and scientific breakthroughs, but only to those organizations that can afford it — leading to a two-speed world in R&D&I.

SIX SCENARIOS FOR THE FUTURE OF AI IN R&D&I

PROBLEM CHILD: Despite some hallmark use cases and affordable solutions, AI fails to demonstrate its value — R&D&I organizations remain concerned about data security, deontology, and lack of interpretability.

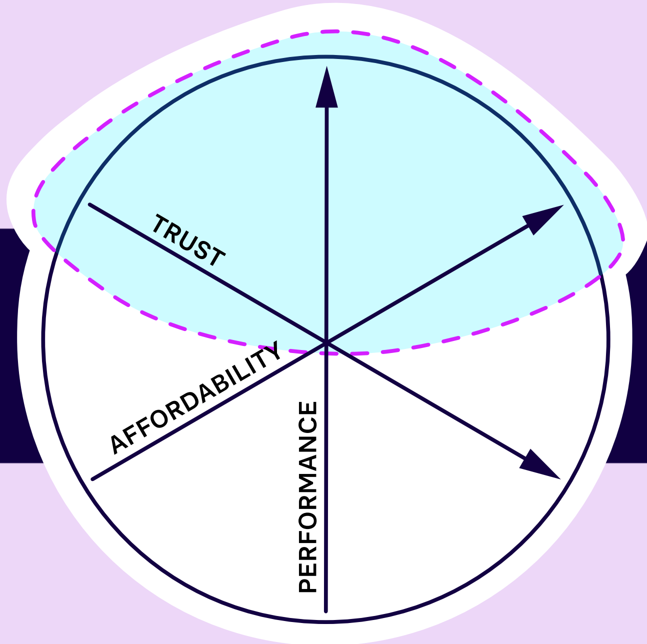

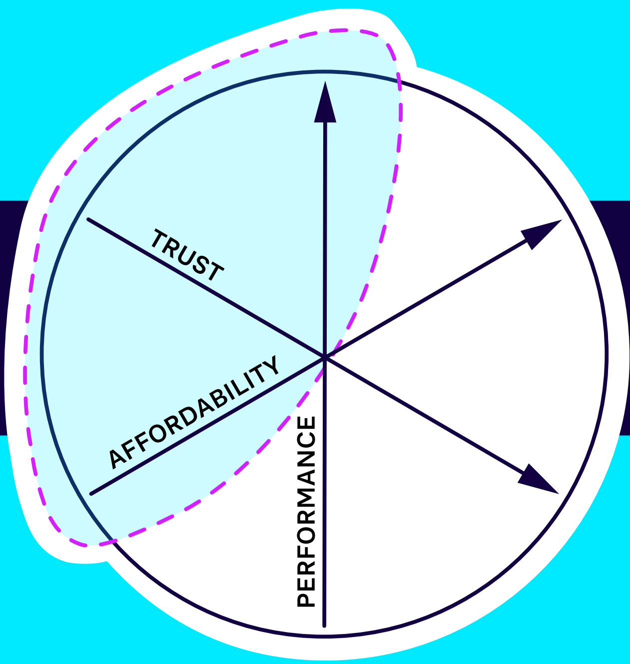

BEST-KEPT SECRET: AI performance improves, but high costs make organizations more risk-averse. Low trust and red tape limit adoption. Few new bold experiments are launched.

CHEAP & NASTY: AI is broadly used in low-stakes use cases, but only as a prototyping or brainstorming tool. Untrustworthy systems are strictly vetted and outputs are verified, curtailing productivity gains.

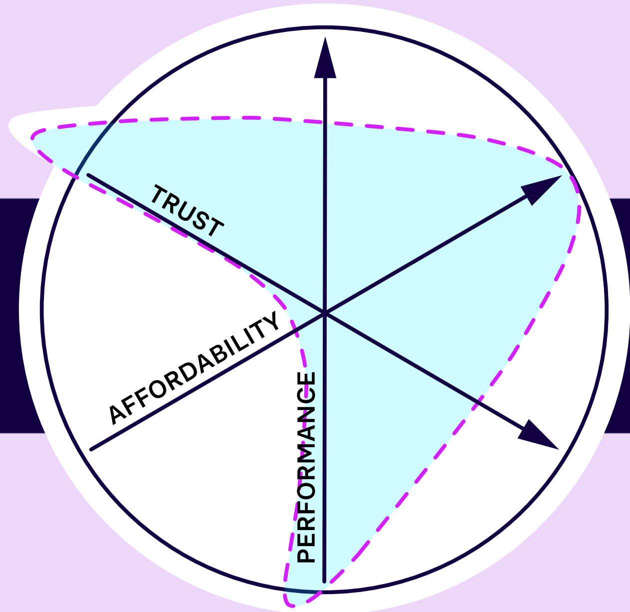

"Just because something thinks differently from you, does it mean it's not thinking?"

— Ava, *Ex Machina*

BLOCKBUSTER

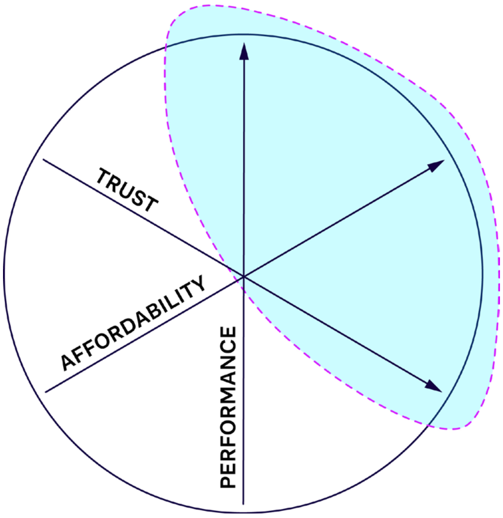

Scenario 1: Blockbuster

AI becomes top of mind throughout the R&D cycle, reshaping organizations along the way. Data becomes the new frontier.

How did we get here?

Foundation models improve through investment and competition, driving AI innovation across architectures, while specialized models thrive with open source support and accessible hosting solutions. On the hardware side, efficient GPUs and edge AI enable widespread local and on-device model deployment. Increased transparency and interpretability boost trust in AI for R&D&I tasks.

What does day-to-day work look like?

AI is everywhere within R&D&I. It automates productivity tasks, enhancing knowledge management and resource planning, while agentic AI and robotics enable fully automated laboratories. AI assists creativity by executing concept designs for developers and innovators. All this results in major scientific discoveries in multiple fields.

How do R&D&I organizations evolve?

The sector sees large-scale initiatives that focus on meeting the growing need for training data. Public institutions maintain and disseminate databases, adapting scientific communication methods to benefit from AI. On the skills side, new data engineering talent eases pressure on ML operations while organizations build new career paths for R&D&I staff as analyst roles shift toward planning.

Winners & losers

Winners can access plentiful, well-structured data through various means, although strict data protection laws may hinder progress in certain fields. Those R&D&I departments tackling AI-friendly problems benefit the most, and organizations skilled at scaling from POCs to deployment gain a unique advantage.

Increased transparency and interpretability boost trust in AI for R&D&I tasks.

CROWD-PLEASER

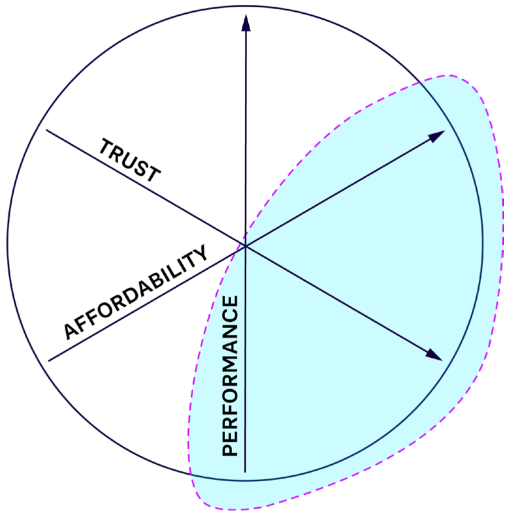

Scenario 2: Crowd Pleaser

AI is convenient, affordable, and adopted for daily productivity tasks but falls short of delivering scientific/creative value.

How did we get here?

AI performance hits a plateau as scaling laws max out for LLMs. However, the industrialization of lower-value functionalities into AI systems with reduced inference costs leads to large-scale adoption. This means R&D&I teams use AI productivity tools, gaining modest productivity improvements while simple explainability techniques foster understanding of AI's limitations and workarounds.

What does day-to-day work look like?

Researchers use AI for non-critical tasks and as a low-quality "sanity check." At the same time, AI augments resource management and customer service systems in daily operations, although unwarranted high trust in AI occasionally leads to costly mistakes. Overall, mature AI use cases become widespread, but major scientific breakthroughs remain elusive.

How do R&D&I organizations evolve?

Organizations rely on commoditized AI solutions for procurement and administrative tasks, with AI deployments driven by IT, digital, or operations departments rather than R&D&I. Pre-trained models and RAG are sufficient to handle popular use cases, limiting new data investments. R&D&I awaits the next breakthrough before it prioritizes AI within scientific strategy.

Winners & losers

Organizations with limited data science capabilities aren't at a significant competitive disadvantage as tools are simple and off-the-shelf. However, due to their size, large, complex organizations benefit most from AI-driven reporting and resource management automation, while smaller organizations experience subtler changes from AI. Strong quality control helps organizations detect AI-related issues before they impact operations.

AI augments resource management and customer service systems in daily operations.

CROWN JEWEL

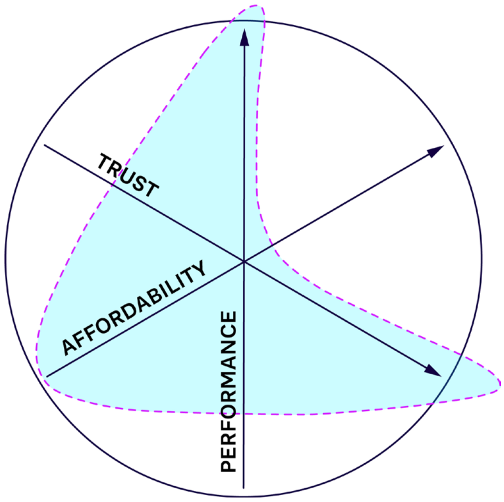

Scenario 3: Crown Jewel

AI delivers productivity and scientific breakthroughs, but only to those organizations that can afford it — leading to a two-speed world in R&D&I.

How did we get here?

AI breakthroughs across architectures improve robustness and efficiency, particularly in RLHF and robotics. Advances in interpretability boost researcher trust, but market consolidation leads to higher prices, and the open source market consolidates with two main players remaining, limiting smaller model offerings. High inference costs and a widening skills gap persist because of GPU limitations.

What does day-to-day work look like?

Because of its extreme cost, AI usage focuses on areas with the highest productivity gains. This means that while major AI-assisted discoveries occur, they are mainly in well-resourced organizations. Researchers design to cost, applying AI selectively to complex problems. Increased model interpretability motivates strong trust in AI outputs.

How do R&D&I organizations evolve?

Well-resourced organizations invest in on-premise compute capabilities while public-private partnerships form to provide access to affordable AI compute resources. Enriching enterprise software with AI capabilities enables limited productivity gains, and on the skills side, organizations encourage researchers to train in data science by adapting their incentive structures.

Winners & losers

A two-speed market emerges, widening the gap between AI-capable and resource-limited organizations. Organizations leveraging AI for operational gains invest more in compute and talent, while access to supercomputers becomes a significant advantage for some R&D&I teams. Mid- and lower-tier teams in AI-friendly fields experience decreased relative influence.

AI usage focuses on areas with the highest productivity gains.

PROBLEM CHILD

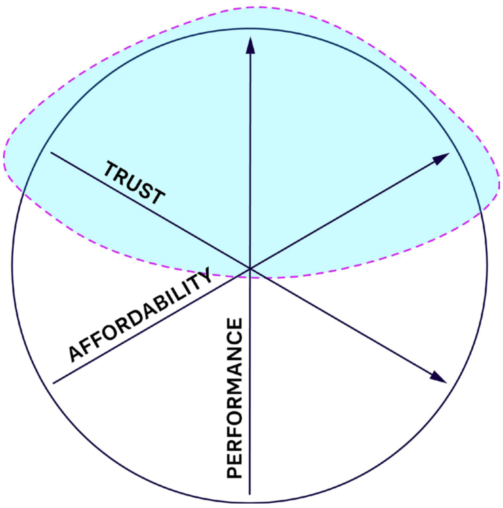

Scenario 4: Problem Child

Despite some hallmark use cases and affordable solutions, AI fails to show its value — R&D&I organizations remain concerned about data security, deontology, and interpretability.

How did we get here?

Foundation models improve through continued investment and competition in various architectures, while efficient GPUs and edge AI enable widespread local and on-device model deployment. However, high-profile AI mishaps erode trust, hindering adoption in R&D&I fields. A lack of progress in interpretability research also erodes trust, perpetuating criticisms of AI systems as opaque black boxes.

What does day-to-day work look like?

Researchers use AI for low-stakes tasks and content drafting, enabled by cheap enterprise-wide AI solutions. Regarding R&D&I-specific applications, AI assists in select exploration cases with proven track records, such as protein folding. Given limited trust, strict oversight and testing guidelines are enforced for AI system use, with vetting procedures creating productivity tradeoffs in researchers' tool selection.

How do R&D&I organizations evolve?

Large-scale initiatives focus on meeting growing needs for training data, with public institutions maintaining and disseminating databases and adapting to new data demands. However, the continued need for verification and oversight means junior analyst roles are maintained and bans on AI-generated content lead to the widespread use of detection tools.

Winners & losers

Organizations able to balance AI benefits with managing potential risks gain a competitive advantage, and those dealing with AI-friendly problems and skilled at scaling POCs achieve dominance. Countries with clear regulatory frameworks promote safer AI use while industries align on AI safety standards to reduce perceived risks and uncertainties.

High-profile AI mishaps erode trust.

BEST–KEPT SECRET

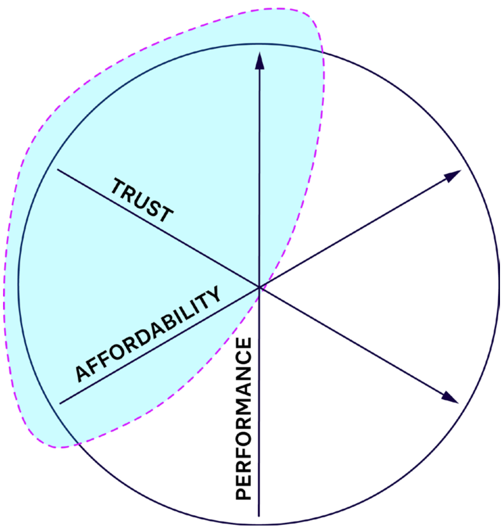

Scenario 5: Best–Kept Secret

AI performance improves, but high costs make organizations more risk-averse. Low trust and red tape limit adoption. Few new bold experiments are launched.

How did we get here?

Foundation models improve through investment and competition, driving AI innovation across architectures. Specialized models thrive because of open source support and accessible hosting solutions. Efficient GPUs and edge AI enable widespread local and on-device model deployment. Increased transparency and interpretability boost trust in AI for R&D&I tasks.

What does day-to-day work look like?

AI automates productivity tasks, enhancing knowledge management and resource planning, while agentic AI and robotics enable fully automated laboratories. AI leads to major scientific discoveries across various fields and assists creativity by executing concept designs for developers and innovators.

How do R&D&I organizations evolve?

Large-scale initiatives focus on meeting growing training data needs, with public institutions maintaining and disseminating databases and adapting scientific communication methods.
An influx of new data engineering talent eases pressure on ML operations while organizations build new career paths as analyst roles shift toward planning.

Winners & losers

Winners can access plentiful, well-structured data through various sources, although strict data protection laws may hinder progress in certain fields. R&D&I departments tackling AI-friendly problems benefit the most, as do organizations skilled at scaling from POCs to deployment, which gain a unique advantage.

AI leads to major scientific discoveries across various fields.

CHEAP & NASTY

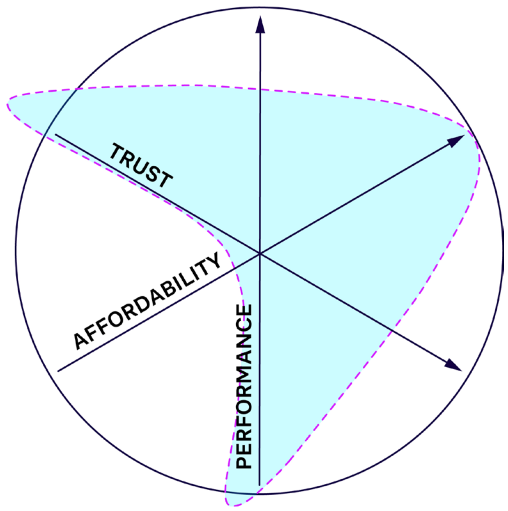

Scenario 6: Cheap & Nasty

AI is broadly used in low-stakes use cases, but only as a prototyping or brainstorming tool. Untrustworthy systems are strictly vetted and outputs are verified, curtailing productivity gains.

How did we get here?

Enterprise AI adoption spreads through aggressive sales tactics and competitive pricing, with efficient GPUs and open source developments enabling widespread model deployment. However, AI performance plateaus, with no significant breakthroughs in new architectures and stalled progress in interpretability, all contributing to an "AI winter" and decreased trust levels.

What does day-to-day work look like?

Researchers use AI for non-critical tasks and as a low-quality "sanity check," with prompt-engineering skills in demand to compensate for the multiple interactions required to achieve results. AI-augmented systems are used for resource management and customer service. Strict oversight and verification procedures limit AI's productivity gains.

How do R&D&I organizations evolve?

Because AI use still requires extensive verification, no changes are made to analyst roles, and organizations only benefit from limited productivity gains in terms of support roles, although RAG is becoming popular for knowledge management. AI deployments are driven by the IT, digital, or operations department rather than being led by R&D&I. Efforts are limited around model fine-tuning for scientific use cases.

Winners & losers

Large organizations benefit most from AI-driven reporting and resource management automation, whereas smaller organizations experience little to no change from adopting AI. Organizations that turned their back on AI maintain their focus on other capabilities, and these AI-skeptical entities build more durable competitive advantages in non-AI areas.

Depicting possible futures with scenarios enables organizations to prepare, identify a path forward with no-regret moves, provide a framework for strategic planning, and make strategic bets based on their needs and capabilities, as we will outline in our concluding chapter.

Organizations that turned their back on AI maintain their focus on other capabilities.

CHAPTER 5

STRATEGIC ACTIONS

5 STRATEGIC ACTIONS

We recommend six no-regret moves for organizations regardless of the six future scenarios. These comprise mutualizing compute power, encouraging data sharing, managing AI talent, training the workforce in AI fundamentals, resetting data and AI governance approaches, and improving output controls. Beyond these, organizations should take measured strategic bets aligned with corporate objectives.

EFFECTIVE DECISION-MAKING AROUND R&D&I AI

In some situations, AI is already enabling double-digit improvements in time, costs, and efficiency in formulation, product development, intelligence, and other R&D&I tasks, examples of which are shown in Figure 15.

Additional gains are likely, but their exact extent will vary greatly depending on which of the six scenarios materialize by 2030.

1. **Blockbuster** — widespread, profound benefits for all
2. **Crowd Pleaser** — widespread but with shallow benefits
3. **Crown Jewel** — deep benefits but limited to best-resourced organizations
4. **Problem Child** — benefits limited by red tape and lack of trust
5. **Best-Kept Secret** — benefits limited to very competent organizations
6. **Cheap & Nasty** — very limited benefits across the board

Figure 15. Examples of R&D&I benefits achieved through AI

CHEMICAL DISCOVERY

30%

TIME SAVINGS IN DRUG DISCOVERY

"With AI, we can predict the characteristics of molecules before synthesis, allowing us to focus our efforts on the most promising candidates. We reduce the number of tests needed to achieve the targeted characteristics, and thus time."

President
AI START-UP FOCUSED ON CHEMICAL R&D

80%

REDUCTION IN REQUIRED EXPERIMENTS

"AI guides us in our experiments and tells what experiments to conduct to fill our data gaps in the most efficient ways. With it, we went from 100k experiments to only a few dozen."

Head of AI internal
AI START-UP, LEADING CHEMICALS MANUFACTURER

FORMULATION

€10M

SAVED IN MANDATORY TESTING

"By integrating AI, we can input specific target characteristics, enabling the system to predict the most effective tests and adjustments needed to refine the formula.... We can save €0.5-1M on avoided regulation-related tests."

Senior R&D executive
CONSUMER PACKAGED GOODS GROUP

15%

REDUCTION IN FORMULATION FAILURE RATE

"AI has been instrumental in navigating formulation spaces, reducing failure rate of tests from 25% to less than 10%. We react to regulation changes, competitive launches, and new findings very quickly with it."

Senior Head of R&D
PERSONAL CARE PRODUCTS GROUP

INDUSTRIALIZATION

30%

IMPROVEMENT IN PRODUCTION EFFICIENCY

"AI optimizes the search for the 'golden batch' — the batch that yields the best efficiency and quality. With these optimal production parameters, we achieve up to 30% improvements in efficiency."

President
AI START-UP FOCUSED ON CHEMICAL R&D

PRODUCT DEV

30%

REDUCTION IN TIME-TO-MARKET

"With AI across our whole process, we developed new products 30% faster. Notably, by analyzing data from production line sensors, AI identifies potential issues and prevents them, thus reducing industrial failures."

Senior R&D executive
FOOD & BEVERAGE GROUP

TECH INTELLIGENCE

10%

TIME SAVED IN LITERATURE REVIEW

"Many projects require a benchmark, state-of-the-art, or scouting of different solutions and technologies. This takes up to 20% of researchers' work time. We reduced this by half by assisting researchers with AI."

Senior R&D executive
CONSUMER PACKAGED GOODS GROUP

Source: Arthur D. Little

SIX VITAL NO-REGRET MOVES TO MAKE NOW

However, no matter which scenario we find ourselves in, six no-regret moves will help R&D&I organizations build resilience and leverage benefits from AI:

1. Manage and empower talent.
2. Control AI-generated content.
3. Build data sharing and collaboration.
4. Train for the long run.
5. Rethink organization and governance beyond IT.
6. Mutualize compute resources.

The better performing and more reliable the AI system, the more likely centralized teams are to succeed.

Manage & empower talent

Managing AI talent will remain a major challenge. However, democratizing AI and externalizing more tactical tasks can provide solutions for R&D&I organizations. Even by 2030, data science PhDs who began university after 2022 will barely be entering the job market, leaving most organizations with a dearth of AI expertise. This talent shortage will be most acute for public organizations, which will not have the means to offer competitive salaries and may not have revised quotas for hires by 2030.

However, the ongoing democratization of AI, such as through the emergence of LCNC AI building solutions, may make data engineer profiles sufficient for the deployment of most use cases.[5] These profiles could be attractive for organizations of all types because they are more junior and quicker to train. Equally, a philosophy of using AI to empower people to innovate better will, over time, reduce dependency on AI specialists.

Subcontracting AI implementations to specialized providers will likely be the most common model by 2030 because it maximizes speed and is well-suited to uncertainty. The better performing and more reliable the AI system, the more likely centralized teams (or service centers) are to succeed. Conversely, the less mature the AI system, the greater the need for co-creation on the ground between researcher and data science teams, such as through paired working groups.

Control AI-generated content

Quality and IP control must be scalable to cope with the multiplication of AI-generated content and data. For example, AI-generated content-detection systems must be thoroughly tested and implemented at scale on all relevant use cases, including publications. The public sharing of validation methodologies for experiments, along with testing data, will incentivize good practices in AI risk management and build trust. Equally, organization-wide risk assessment frameworks will need to be updated to properly account for AI risk.

Build up data sharing & collaboration

Strengthening data and knowledge ecosystems is vital to enabling successful AI deployment.[6] While data availability in academia is variable, some leading projects could inform future data-sharing efforts. These include large initiatives such as the datasets provided by the European Bioinformatic Institute, which stores petabytes of publicly available data — most of it funded by governments, the Harvard Dataverse, a central collection of comprehensive datasets, and more specialist, well-resourced communities such as FlyBase, the database created by the global Drosophila community.

In the private sector, data sharing remains piecemeal because of concerns over losing competitive advantage to rivals. However, while a few experiments have been made, such as BMW, Daimler, and Volkswagen sharing data to develop self-driving car technology, experimental data collected by corporations is rarely shared, if ever, especially in sectors such as life sciences. This is unlikely to change over time. At the same time, public-private partnerships are a promising avenue for AI data sharing at scale, such as the partnership between Google, SkyTruth, and Oceana, which has created the Global Fishing Watch initiative to share data about illegal fishing worldwide. Moving forward, we will likely see public scientific organizations acting as curators and stewards of datasets of public interest, leading to a possible shift in their mission.

“Experimental data is our competitive advantage. I don’t see it being shared at any scale any time soon.”

Senior executive, pharmaceutical company

Train for the long run

Training in AI fundamentals should be delivered continuously beyond the immediate users of tools, as this will accelerate adoption and better manage risk. Addressing as broad an audience as possible with training ensures that decisions throughout the organization consider AI’s requirements and limitations.

Training should include AI’s technical fundamentals, functional capabilities, implementation requirements, and potential risks. Educational AI itself (e.g., in interactive form) can be leveraged to tailor the delivery of training and coaching programs to reach individuals at scale. AI will also trigger training needs around current R&D&I roles, as these will change as their tools evolve, depending on the depth of adoption. A further benefit of successful training is its contribution to increased AI usage and better quality control management.

Rethink organization & governance beyond IT

The governance of AI in organizations will involve more than IT, although its actual setup should vary depending on the use case. For use cases critical to the mission or business, a centralized governance instance, such as a "department of AI systems," may be desirable to overcome coordination problems. Daily productivity use cases and rollouts are best managed by the IT department in coordination with the leaders of relevant teams (finance, HR, etc.), as with any other productivity tool.

AI governance should report directly to the executive committee or board to ensure maximum visibility. It can use data governance, another transversal function, as a template for its structure and processes. Overall, organizational models should be adapted as far as possible to remove barriers that could hamper the internal sharing of data and knowledge.

Mutualize compute resources

The mutualization of compute power will likely remain the best way for many academics and corporate players to afford sovereign pre-training. Mutualization initiatives are already underway at different levels and scales. For example, in France, organizations can receive mutualized access to the Jean Zay AI super calculator. In contrast, researchers in large university centers such as Princeton University and within companies such as AstraZeneca (at a corporate level) can access centralized resources. Public-private partnerships will also likely emerge, hosted by innovation and research communities.

To meet changing needs, supercomputers offer significant compute capacity at a fraction of the cost but at the scale of a GPU. However, they are not ideally suited to all forms of deep learning (including transformer-based models). They, therefore, require adjustments (e.g., precision optimization, interconnect improvements, or memory expansion). The successful mutualization of compute resources involves strict policies, such as a compute time distribution process, monitoring capabilities, and an enforcement mechanism. Organizations should look at their mutualization options now and implement the necessary partnerships and processes to deliver on their future compute needs.

"You need to divvy out compute time but also monitor utilization — it's a maintenance process."

Research & project administration senior executive, Princeton University

AI STRATEGIC BETS R&D&I ORGANIZATIONS SHOULD ALSO MAKE

Beyond the no-regret moves, R&D&I organizations will need to place AI bets — hedge, speculate, or short — as part of their innovation project portfolio optimization (see Figure 16). This should be based on their strategic objectives, capabilities, and market intelligence, and the context in which they operate. Adopting a portfolio approach means the strategy can include all three bets related to different use cases.

Three fictional examples (see Figure 17) help illustrate how this type of investment strategy for AI use cases depends on the prevailing scenario, business context and goals, and existing capabilities.

As with all investments, ultimately the decision will depend on the business case, considering the costs and benefits, all within the context of the broader innovation strategy (see Figure 18).

R&D&I organizations will need to place AI bets as part of their innovation project portfolio optimization.

Figure 16. Investment decisions for AI in R&D&I

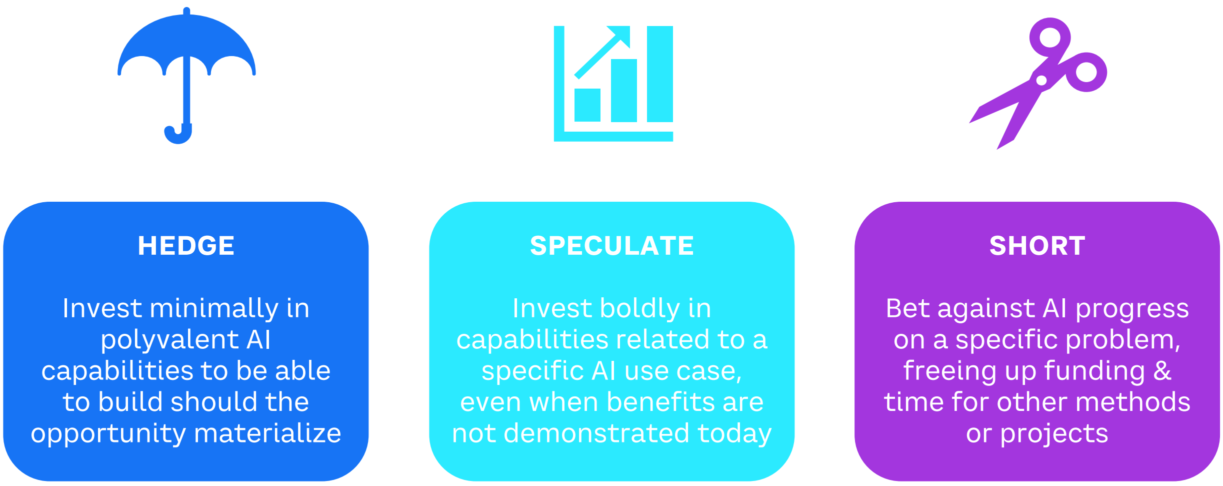

Source: Arthur D. Little

Figure 17. Making investment choices around AI

ILLUSTRATIVE

	NUCLEAR FUSION START-UP	ADVANCED MATERIALS SUPPLIER	PREMIUM TOY MANUFACTURER
Prevailing scenario	• **Crown Jewel**: AI performs well but is costly & limited to very well-resourced companies	• **Blockbuster**: AI performs well, is accessible/trustworthy & widely used	• **Cheap & Nasty**: AI applications remain limited & are mistrusted in all but low-risk applications
Company context	• Designer of prototype small-scale nuclear fusion reactor technology • One of top five start-ups, very well-funded • Has a portfolio of patents already	• Pseudo commodity products: key purchasing criterion is price, except for breakthroughs • Mid-tier vs. competitors, limited geographical presence • Large portfolio of product references	• Mid-tier but higher margins than most competitors • Value of IP is linked to brand, centered on addressing stable customer preferences • Complex but mature regulatory environment
Business goals	• Be the first to succeed in maintaining steady state plasma with net energy gain	• Grow by entering new geographies with existing/adjacent products	• Maintain current profitability while steadily increasing market share in core categories
Existing capabilities	• Some data sciences skills; outstanding quantitative and scientific problem-solving talent • Systematically recorded experimental data; fully simulatable environments	• Wealth of formulation & experimental data but not well structured • Limited specialized data science skills • Rich ecosystem of partnerships with academia	• No AI systems currently implemented in the firm • Incomplete design & manufacturing data • No data science talent or skills
AI use case being considered	• Control of plasma stability using reinforcement learning	• New materials discovery using GNN	• 3D prototyping of new product designs using diffusion & transformers
Selected strategic bets	**SPECULATE** Invest boldly in capabilities related to a specific AI use case, even when benefits are not demonstrated today	**HEDGE** Invest minimally in polyvalent AI capabilities to be able to build should the opportunity materialize	**SHORT** Bet against AI progress on a specific problem, freeing up funding & time for other methods or projects
Rationale	*Worth making a high risk/high reward bet; AI could unlock huge value, favorable competitive position*	*Given strategic goals, wide AI application by competitors & underlying data issues, invest for fast followership rather than leadership*	*Given poor AI performance & lack of in-house skills, allow competitors to invest in AI; focus instead on strengthening core business*

Source: Arthur D. Little

Figure 18. Business case for AI-driven change in research and innovation

STRATEGIC CHALLENGES
- Strategic priorities
- (Digital) business models
- Data ecosystems
- M&A activity
- Prolonged/sudden crises
- Sustainability

STRATEGIC BET

AI SCENARIO
- Cheap & Nasty
- Best-Kept Secret
- Problem Child
- Crown Jewel
- Crowd Pleaser
- Blockbuster

CURRENT RESOURCES & CAPABILITIES
- Funds
- Data
- Talent
- Experimental Assets
- Ecosystem

COSTS
- CAPEX
- Dataset building
- Talent hires
- Cloud/platform subscriptions

BENEFITS
- Better margins through higher productivity
- Increased revenue via more products, broader range
- Stronger competitive position via faster go-to-market
- New markets via breakthrough innovation

Source: Arthur D. Little

"A major challenge for humanity is to be able to understand itself. AI may be our best mirror."

— Fei-Fei Li, Professor of Computer Science, Stanford

Acknowledgments

We would first like to acknowledge our five study partners for their leadership, precious contributions, and attentive review: LVMH, Avril, GRTgaz, the French National Centre for Scientific Research (CNRS), and the French National Research Institute for Agriculture, Food, and Environment (INRAE), and, in particular:

Massih-Reza Amini, Deputy Scientific Director for Artificial Intelligence, CNRS Informatics

Valérie Archambault, Director of Partnerships and Innovation Transfer, INRAE

Sophie de Batz, GenAI Strategic Leader, Avril

Florent Brissaud, R&D Project Manager for AI Applications, GRTgaz

Clément Favier, Head of Data Science, Avril

Mehdi Gmar, Director, CNRS Innovation

David Kierbel, Chief Digital Officer, LVMH Recherche

Pierre Lacotte, Director for Innovation Projects, LVMH Gaïa

Alda Mari, Research Director, Institut Jean Nicod, CNRS

Hadi Quesneville, Deputy Director of Directorate for Open Science and Chief Data Officer, INRAE

Pierre Roy, Head of Processes and Gaz Characterization Laboratory, GRTgaz

Marian Scuturici, Deputy Scientific Director for Technology Transfer, Infrastructure, and Platforms, CNRS Informatics

We would also like to thank our partners, **Impact AI and Executive Director Roxana Rugina** for their unwavering support and expertise.

Finally, we would like to thank all the experts we interviewed as part of this study for their time, guidance, and valuable insights. In particular, we thank:

Johan Aubert, Chief Technology and Digital Officer, L'Oréal

Mathieu Bertin, Data Science Lead and AI Lab Manager, Syensqo

Patrice Brès, former CEO, Trapil

Yves Caseau, Chief Digital Officer, Michelin

Vincent Champain, Senior Executive VP Digital Performance & IT, Framatome

Matthieu Chatillon, Director of Innovation, Avril Animal Specialty

Vincent Colegrave, Head of Syensqo.ai

Marc Danzart, Fellow, Presans

Acknowledgments

Virginie Dubois, Executive VP R&D, Roquette

Carlos Escapa, Data AI/ML Global Practice Lead, Accenture AWS Business Group

Marie-Pierre Fontanel, Deputy Director of Information Systems, CNRS

Denis Gardin, Director of Innovation, MBDA

Mehdi Gmar, General Director, CNRS Innovation

Bruno Guilbot, Data Science, AI, and New Technologies Director, Louis Vuitton

Dino Ienco, Research Director, INRAE

Lionel Jouffe, President and CEO, Bayesia

Luc Julia, Chief Scientific Officer, Renault Group

Remy Kusters, Head of R&D, Gourmey

Gregory LaBlanc, Lecturer, University of California, Berkeley and Stanford University

Yann Lechelle, cofounder and CEO, Probabl

Damien Lemaire, Head of Strategic Partnerships, French Alternative Energies and Atomic Energy Commission (CEA)

Frédéric Leroy, Fellow, Presans

Mario Maglione, Research Director, Bordeaux Institute of Condensed Matter Chemistry, CNRS

Carlos Martin, Managing Director, MACAMI Group

Frédéric Mours, Head of DataLab, GRTgaz

Bernhard Quendt, Chief Technical Officer, Thales Group

Emmanuelle Quilès, Worldwide VP Cardiovascular, Metabolism, and Pulmonary Hypertension, Johnson & Johnson

Erez Raanan, CEO, Mathlabs

Philippe Robin, President, Alysophil

Victor Sayous, cofounder and Chief Technical Officer, Gourmey

Philippe Seberac, Scientific and Technical Director, Veolia

Timothée Silvestre, Head of Technological Foresight, CEA

Jard van Ingen, CEO and cofounder, GetFocus

Endnotes

1 Kolk, Michaël, Marten Zieris, and Michael Eiden. "The People-Centric Lab of the Future." ADL *Prism*, 2024.
2 Artificial Intelligence Commission. "Our AI: Our Ambition for France." French Government, March 2024.
3 "Artificial Intelligence: Key Insights, Data and Tables." Ipsos, 2024; the survey consisted of interviews with 22,816 adults in 31 countries from May 2023–June 2023.
4 We exclude the low affordability/low performance/low trust (degraded status quo) combination because it does not reveal much about the future of AI and R&D&I. We also exclude the low affordability/low performance/high trust combination because we deem it inconsistent.
5 Kolk et al. (see 1).
6 Kolk et al. (see 1).

Disclaimer

The AI field is still evolving rapidly. This Report reflects the best of our knowledge as of October 2024.

Dr. Albert Meige, Director of Blue Shift, Arthur D. Little
albert.meige@adlittle.com

Blue Shift, by Arthur D. Little, explores the impact of technologies on business, society, and humans. The Blue Shift Report covers these topics in depth, inviting guest authors, academics, and artists to contribute to the conversation.

Want to join the discussion? Please address your comments to our editorial office at Arthur D. Little, New Fetter Place West, 2nd floor, 55 Fetter Lane, London EC4A 1AA, United Kingdom —
Telephone: +44 7710 536 471.

For further information, visit: www.adlittle.com.

An Arthur D. Little Blue Shift study, in partnership with:

Avril

GRTgaz

LVMH

cnrs

INRAE

BLUE SHIFT

BY ARTHUR D. LITTLE

ARTHUR D LITTLE